

Implementing Cisco MPLS Traffic Engineering and Other Features

Version 2.0

Student Guide

Text Part Number: 97-2043-01



***Students, this letter describes important
course evaluation access information!***

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco Systems is committed to bringing you the highest-quality training in the industry. Cisco learning products are designed to advance your professional goals and give you the expertise you need to build and maintain strategic networks.

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. We would appreciate a few minutes of your time to complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short post-course evaluation. If there is no Internet access in the classroom, please complete the evaluation within the next 48 hours or as soon as you can access the web

On behalf of Cisco, thank you for choosing Cisco Learning Partners for your Internet technology training.

Sincerely,

Cisco Systems Learning

Table of Contents

Volume 1

<u>Course Introduction</u>	<u>1</u>
Overview	1
Outline	2
Course Objectives	3
Cisco Certifications	4
Learner Skills and Knowledge	5
Learner Responsibilities	6
General Administration	7
Course Flow Diagram	8
Icons and Symbols	9
Learner Introductions	10
<u>MPLS VPN Review</u>	<u>1-1</u>
Overview	1-1
Objectives	1-1
Outline	1-1
<u>Basic MPLS Concepts</u>	<u>1-3</u>
Overview	1-3
Relevance	1-3
Objectives	1-3
Learner Skills and Knowledge	1-3
Outline	1-4
Basic MPLS Concepts	1-5
MPLS Architecture	1-6
MPLS Labels	1-8
Label Switch Routers	1-10
Lesson Summary	1-13
References	1-13
Lesson Review	1-14
Lesson Answer Key	1-15
<u>Label Distribution in Frame-Mode MPLS</u>	<u>1-17</u>
Overview	1-17
Relevance	1-17
Objectives	1-17
Learner Skills and Knowledge	1-17
Outline	1-18
Label Allocation in a Frame-Mode MPLS Network	1-19
Label Distribution and Advertisement	1-21
Populating the LFIB	1-24
Packet Propagation Across an MPLS Network	1-25
Penultimate Hop Popping	1-26
Lesson Summary	1-28
References	1-28
Lesson Review	1-29
Lesson Answer Key	1-30
<u>Configuring MPLS</u>	<u>1-31</u>
Overview	1-31
Relevance	1-31
Objectives	1-31
Learner Skills and Knowledge	1-31
Outline	1-31
Configuring IP CEF	1-32

Configuring MPLS on a Frame-Mode Interface	1-34
Configuring the MTU Size in Label Switching	1-37
Lesson Summary	1-39
References	1-39
Lesson Review	1-40
Lesson Answer Key	1-41
MPLS VPN Routing Model and Packet Forwarding	1-43
Overview	1-43
Relevance	1-43
Objectives	1-43
Learner Skills and Knowledge	1-43
Outline	1-44
Propagation of Routing Information Across the P-Network	1-45
Route Distinguishers	1-47
Route Targets	1-50
End-to-End Routing Information Flow	1-53
MPLS VPNs and Packet Forwarding	1-56
Lesson Summary	1-58
References	1-58
Lesson Review	1-59
Lesson Answer Key	1-60
Configuring Simple MPLS VPNs	1-61
Overview	1-61
Relevance	1-61
Objectives	1-61
Learner Skills and Knowledge	1-61
Outline	1-62
Configuring VRF Tables	1-63
Configuring an MP-BGP Session Between PE Routers	1-65
Configuring Routing Protocols Between PE and CE Routers	1-66
Lesson Summary	1-75
References	1-75
Lesson Review	1-76
Lesson Answer Key	1-77
Module Assessment	1-79
Overview	1-79
Quiz: MPLS VPN Review	1-80
Objectives	1-80
Instructions	1-80
Quiz	1-80
Scoring	1-81
Module Assessment Answer Key	1-82
Carrier Supporting Carrier	2-1
Overview	2-1
Objectives	2-1
Outline	2-1
Carrier's Carrier Overview	2-3
Overview	2-3
Relevance	2-3
Objectives	2-3
Learner Skills and Knowledge	2-3
Outline	2-4
Carrier Supporting Carrier	2-5

Benefits of CSC	2-7
How CSC Works	2-9
Lesson Summary	2-10
References	2-10
Lesson Review	2-11
Lesson Answer Key	2-12
Implementing the CSC Models	2-13
Overview	2-13
Relevance	2-13
Objectives	2-13
Learner Skills and Knowledge	2-13
Outline	2-14
CSC Models	2-15
Customer Carrier: Native IP	2-16
Customer Carrier: MPLS and LDP	2-25
Customer Carrier: MPLS and EBGp	2-30
Customer Carrier: MPLS VPNs	2-35
Lesson Summary	2-38
References	2-38
Next Steps	2-38
Lesson Review	2-39
Lesson Answer Key	2-41
MPLS Traffic Engineering Technology	3-1
Overview	3-1
Objectives	3-1
Outline	3-1
Traffic Engineering Concepts	3-3
Overview	3-3
Relevance	3-3
Objectives	3-3
Learner Skills and Knowledge	3-3
Outline	3-4
What Is Traffic Engineering?	3-5
Business Drivers for Traffic Engineering	3-7
Congestion Avoidance and Traffic Engineering	3-10
Traffic Engineering with a Layer 2 Overlay Model	3-11
Traffic Engineering with a Layer 3 Model	3-15
Traffic Engineering with the MPLS-TE Model	3-17
Lesson Summary	3-19
References	3-19
Lesson Review	3-20
Lesson Answer Key	3-21
MPLS Traffic Engineering Components	3-23
Overview	3-23
Relevance	3-23
Objectives	3-23
Learner Skills and Knowledge	3-23
Outline	3-24
Traffic Tunnels: Concepts	3-25
Traffic Tunnels: Characteristics	3-27
Traffic Tunnels: Attributes	3-28
Network Links and Link Attributes	3-30
Constraint-Based Path Computation	3-31
Role of RSVP in Path Setup Procedures	3-35

Forwarding Table Modifications	3-36
Lesson Summary	3-38
References	3-38
Lesson Review	3-39
Lesson Answer Key	3-40
Constraint-Based Path Computation	3-41
Overview	3-41
Relevance	3-41
Objectives	3-41
Learner Skills and Knowledge	3-41
Outline	3-42
Constraint-Based Path Computation Attributes	3-43
MPLS-TE Link Resource Attributes	3-44
MPLS-TE Tunnel Attributes	3-48
Implementing TE Policies with Affinity Bits	3-53
Propagating MPLS-TE Link Attributes with Link-State Routing	3-57
Constraint-Based Path Computation	3-61
Lesson Summary	3-68
References	3-68
Lesson Review	3-69
Lesson Answer Key	3-70
Path Setup and Maintenance	3-71
Overview	3-71
Relevance	3-71
Objectives	3-71
Learner Skills and Knowledge	3-71
Outline	3-72
Path Setup	3-73
RSVP Usage in Path Setup	3-75
Hop-by-Hop Path Setup with RSVP	3-76
Tunnel and Link Admission Control	3-81
Path Rerouting	3-83
Path Reoptimization	3-84
Path Rerouting: Link Failure	3-86
Lesson Summary	3-88
References	3-88
Lesson Review	3-89
Lesson Answer Key	3-90
Assigning Traffic to Traffic Tunnels	3-91
Overview	3-91
Relevance	3-91
Objectives	3-91
Learner Skills and Knowledge	3-91
Outline	3-92
Traffic Flow Modifications	3-93
Static Routing Pointing to a Tunnel	3-94
Solving the BGP Next-Hop Problem	3-95
IP Forwarding Database Modification with Autoroute	3-96
Autoroute Path Selection Rules	3-97
Autoroute Examples	3-98
Forwarding Adjacency	3-103
Lesson Summary	3-107
References	3-107
Lesson Review	3-108

Lesson Answer Key	3-109
Module Assessment	3-111
Overview	3-111
Quiz: MPLS Traffic Engineering Technology	3-112
Objectives	3-112
Instructions	3-112
Quiz	3-112
Scoring	3-113
Module Assessment Answer Key	3-114
<i>Configuring MPLS Traffic Engineering</i>	<i>4-1</i>
Overview	4-1
Objectives	4-1
Outline	4-1
Configuring MPLS Traffic Engineering on Cisco IOS Platforms	4-3
Overview	4-3
Relevance	4-3
Objectives	4-3
Learner Skills and Knowledge	4-4
Outline	4-4
MPLS-TE Configuration Road Map	4-5
Enabling Device-Level MPLS-TE Support	4-6
Enabling MPLS-TE on an Interface	4-8
Enabling MPLS-TE Support in IS-IS	4-13
Enabling MPLS-TE Support in OSPF	4-16
Creating and Configuring a Traffic Tunnel	4-18
Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency	4-25
Implementing MPLS-TE	4-27
Lesson Summary	4-32
References	4-32
Lesson Review	4-33
Lesson Answer Key	4-34
Advanced MPLS-TE Path Selection	4-35
Overview	4-35
Relevance	4-35
Objectives	4-35
Learner Skills and Knowledge	4-35
Outline	4-36
MPLS-TE Case Study	4-37
Creating Dynamic Path Options for a Traffic Tunnel	4-40
Creating Explicit Path Options for a Traffic Tunnel	4-45
Assigning Traffic to Traffic Tunnels with Static Routes	4-48
Configuring Periodic and Event-Triggered Optimization	4-49
Configuring Path Selection Restrictions	4-53
Example: Configuring Affinity Bits on Tunnels	4-55
Modifying Administrative Weight	4-59
Autoroute Forwarding Engine	4-61
Lesson Summary	4-64
References	4-64
Lesson Review	4-65
Lesson Answer Key	4-66

Advanced MPLS-TE Link Protection	4-67
Overview	4-67
Relevance	4-67
Objectives	4-67
Learner Skills and Knowledge	4-67
Outline	4-68
Improving Convergence Time	4-69
Configuring Parallel Tunnels	4-70
Drawbacks of Parallel Tunnels	4-72
Fast Reroute	4-74
Link Protection with FRR	4-76
Node Protection with FRR	4-78
Fast Reroute Configuration	4-80
Lesson Summary	4-83
References	4-83
Lesson Review	4-84
Lesson Answer Key	4-85
Advanced MPLS-TE Bandwidth Control	4-87
Overview	4-87
Relevance	4-87
Objectives	4-87
Learner Skills and Knowledge	4-87
Outline	4-88
Bandwidth Reservation with Autobandwidth	4-89
Autobandwidth Configuration Commands	4-92
DiffServ-Aware Service	4-94
DiffServ-Aware Service Configuration Commands	4-99
Lesson Summary	4-101
References	4-101
Lesson Review	4-102
Lesson Answer Key	4-103
MPLS-TE Interarea Support	4-105
Overview	4-105
Relevance	4-105
Objectives	4-105
Learner Skills and Knowledge	4-105
Outline	4-106
MPLS-TE Without Interarea Support	4-107
Interarea MPLS-TE	4-108
Interarea MPLS-TE Detailed Operation	4-111
Interarea MPLS-TE Restoration	4-112
Lesson Summary	4-114
References	4-114
Lesson Review	4-115
Lesson Answer Key	4-116
Monitoring and Troubleshooting MPLS-TE	4-117
Overview	4-117
Relevance	4-117
Objectives	4-117
Learner Skills and Knowledge	4-117
Outline	4-118
Monitoring MPLS-TE Tunnels	4-119
Example: Usage of the show mpls traffic-eng tunnel brief Command	4-121
Example: Usage of the show mpls traffic-eng autoroute Command	4-122

Monitoring MPLS-TE	4-123
Example: Usage of the show mpls traffic-eng link-management summary Command	4-123
Example: Usage of the show mpls traffic-eng link-management interfaces Command	4-124
Example: Usage of the show mpls traffic-eng link-management bandwidth-allocation Command	4-126
Example: Usage of the show mpls traffic-eng link-management advertisements Command	4-127
Example: Usage of the show mpls traffic-eng link-management admission-control Command	4-128
Example: Usage of the show mpls traffic-eng link-management igp-neighbors Command	4-130
Example: Usage of the show mpls traffic-eng topology Command	4-130
Example: Usage of the show mpls traffic-eng topology path Command	4-132
Monitoring MPLS-TE Resource Usage on Intermediate Devices	4-133
Example: Usage of the show ip ospf database opaque-area Command	4-133
Example: Usage of the show ip ospf Command	4-134
Example: Usage of the show isis database verbose Command	4-135
Example: Usage of the show isis mpls traffic-eng adjacency-log Command	4-136
Example: Usage of the show isis mpls traffic-eng advertisements Command	4-137
Example: Usage of the show isis mpls traffic-eng tunnel Command	4-138
Troubleshooting Device and Link Setup	4-139
Lesson Summary	4-140
References	4-140
Next Steps	4-140
Lesson Review	4-141
Lesson Answer Key	4-142

Table of Contents

Volume 2

<u>MPLS Quality of Service</u>	<u>5-1</u>
Overview	5-1
Objectives	5-1
Outline	5-1
<u>QoS Models</u>	<u>5-3</u>
Overview	5-3
Relevance	5-3
Objectives	5-3
Learner Skills and Knowledge	5-3
Outline	5-4
QoS Models	5-5
The QoS Pendulum	5-6
Integrated QoS Model	5-7
Differentiated QoS Model	5-8
Lesson Summary	5-10
References	5-10
Lesson Review	5-11
Lesson Answer Key	5-12
<u>MPLS Support for DiffServ</u>	<u>5-13</u>
Overview	5-13
Relevance	5-13
Objectives	5-13
Learner Skills and Knowledge	5-13
Outline	5-14
DiffServ Architecture	5-15
DiffServ Model Features	5-16
DiffServ PHBs and Recommended Codepoints	5-18
DiffServ Scalability by Means of Aggregation	5-21
MPLS Scalability by Means of Aggregation	5-22
Marking MPLS Frames	5-23
Lesson Summary	5-25
References	5-25
Lesson Review	5-26
Lesson Answer Key	5-27
<u>Configuring MPLS QoS</u>	<u>5-29</u>
Overview	5-29
Relevance	5-29
Objectives	5-29
Learner Skills and Knowledge	5-29
Outline	5-30
How to Use the QoS Toolkit	5-31
Configuring QoS in Cisco IOS Modular QoS CLI	5-32
Configuring QoS in Cisco IOS MQC Abstractions	5-33
Configuring QoS in Cisco IOS MQC Syntax	5-35
MPLS QoS Configuration Case Study	5-37
Lesson Summary	5-55
References	5-55
Lesson Review	5-56
Lesson Answer Key	5-57

QoS in MPLS Applications	5-59
Overview	5-59
Relevance	5-59
Objectives	5-59
Learner Skills and Knowledge	5-59
Outline	5-60
MPLS-TE with a Best-Effort Network	5-61
MPLS-TE with a DiffServ Network	5-62
MPLS DS-TE with a DiffServ Network	5-64
QoS-Enabled MPLS VPNs	5-66
QoS Implementation	5-69
Lesson Summary	5-74
References	5-74
Next Steps	5-74
Lesson Review	5-75
Lesson Answer Key	5-76
Any Transport over MPLS	6-1
Overview	6-1
Objectives	6-1
Outline	6-1
Introduction to Any Transport over MPLS	6-3
Overview	6-3
Relevance	6-3
Objectives	6-3
Learner Skills and Knowledge	6-3
Outline	6-4
AToM Overview	6-5
Transport Types	6-8
How AToM Works	6-9
AToM Control Word	6-14
Lesson Summary	6-16
References	6-16
Lesson Review	6-17
Lesson Answer Key	6-18
Configuring AToM on Cisco IOS Platforms	6-19
Overview	6-19
Relevance	6-19
Objectives	6-19
Learner Skills and Knowledge	6-20
Outline	6-20
MTU Issues	6-21
AToM Packet Forwarding with Summarization in the Core	6-22
AToM Configuration	6-23
EoMPLS Configuration	6-28
PPP over MPLS Configuration	6-31
HDLC over MPLS Configuration	6-33
FRoMPLS Port Configuration	6-35
AAL5 over MPLS Configuration	6-39
ATM over MPLS Configuration	6-41
Lesson Summary	6-44
References	6-44
Lesson Review	6-45
Lesson Answer Key	6-46

Monitoring AToM on Cisco IOS Platforms	6-47
Overview	6-47
Relevance	6-47
Objectives	6-47
Learner Skills and Knowledge	6-47
Outline	6-48
Monitoring AToM	6-49
show mpls l2transport vc	6-50
show mpls l2transport vc detail	6-53
show mpls l2transport summary	6-55
debug mpls l2transport vlan control	6-57
Lesson Summary	6-59
References	6-59
Next Steps	6-59
Lesson Review	6-60
Lesson Answer Key	6-61
MPLS IPv6 Support	7-1
Overview	7-1
Objectives	7-1
Outline	7-1
Review of IPv6	7-3
Overview	7-3
Relevance	7-3
Objectives	7-3
Learner Skills and Knowledge	7-3
Outline	7-4
The Benefits of Integrating IPv6	7-5
IPv6 Technology Scope	7-6
IPv6 Address Representation	7-10
Hierarchical Addressing and Aggregation	7-12
Lesson Summary	7-15
References	7-15
Lesson Review	7-16
Lesson Answer Key	7-17
Implementing IPv6 over MPLS	7-19
Overview	7-19
Relevance	7-19
Objectives	7-19
Learner Skills and Knowledge	7-19
Outline	7-20
Benefits of Deploying IPv6 over MPLS Backbones	7-21
IPv6 Using Tunnels on the Customer Edge Routers	7-22
Deploying IPv6 Using Tunnels on the Customer Edge Routers	7-23
IPv6 over a Circuit Transport over MPLS	7-24
Deploying IPv6 over a Circuit Transport over MPLS	7-25
IPv6 on the Provider Edge Routers (Cisco 6PE)	7-26
Deploying Cisco 6PE	7-29
Lesson Summary	7-34
References	7-34
Lesson Review	7-35
Lesson Answer Key	7-36

Monitoring IPv6 over MPLS	7-37
Overview	7-37
Relevance	7-37
Objectives	7-37
Learner Skills and Knowledge	7-38
Outline	7-38
Monitoring IPv6 Support	7-39
show bgp ipv6 Command	7-40
show bgp ipv6 neighbors Command	7-41
show mpls forwarding-table Command	7-42
show bgp ipv6 labels Command	7-43
show ipv6 cef Command	7-44
show ipv6 route Command	7-45
Lesson Summary	7-46
References	7-46
Next Steps	7-46
Lesson Review	7-47
Lesson Answer Key	7-49
<u>Course Glossary</u>	<u>1</u>

Course Introduction

Overview

Service providers today are faced with many challenges as they attempt to deal with current customer demand, including a continuing need for value-added services. Conventional IP packet forwarding has several limitations. Traditional Layer 2 services and traffic engineering (TE) lack the flexibility to meet the demands of the modern network. More and more, service providers realize that something else is needed. Not only must they be concerned with protecting their existing infrastructure, but they must find ways to generate new services that are not currently supportable with existing technologies.

Multiprotocol Label Switching (MPLS) is a high-performance method for forwarding packets through a network. It provides a foundation for a series of applications that integrate the performance and traffic management capabilities of the data link layer (Layer 2) with the scalability and flexibility of network layer (Layer 3) routing. This integration allows service providers the ability to support “value-added” features that are critical for their networks.

Implementing Cisco MPLS Traffic Engineering and Other Features (MPLST) v2.0 is an instructor-led training (ILT) course that extends the knowledge and skills that the learner has obtained from the MPLS Cisco CCIP™ certification track. The course duration is 5 days in length, using a standard ILT format. This course is recommended training for individuals who are seeking advanced MPLS-based knowledge and skills. The course will enable learners to gather information ranging from the technology basics to more updated features and functions such as MPLS Traffic Engineering (MPLS-TE), Fast Reroute (FRR), and Any Transport over MPLS (AToM). The focus of the course is on the technology issues of MPLS from a service provider perspective and how to configure some of the MPLS features and functions in an existing routed environment.

Outline

The Course Introduction includes these topics:

- Course Objectives
- Cisco Certifications
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Flow Diagram
- Icons and Symbols
- Learner Introductions

Course Objectives

This topic lists the course objectives.

Course Objectives

Cisco.com

- Review the MPLS peer-to-peer architecture and explain label allocation, routing update distribution, and packet forwarding
- Configure, monitor, and troubleshoot a typical MPLS network using a Carrier Supporting Carrier implementation
- Identify the MPLS Traffic Engineering architecture and explain how MPLS implements traffic engineering, establishes the constraint-based path, and assigning traffic to traffic trunks
- Configure, monitor, and troubleshoot MPLS Traffic Engineering

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-3

Course Objectives (Cont.)

Cisco.com

- Implement a defined SLA using MPLS QoS services
- Configure, monitor, and troubleshoot Layer 2 services using the Any Transport over MPLS service
- Configure, monitor, and troubleshoot IPv6 support in an MPLS environment

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-4

Cisco Certifications

This topic lists the certification requirements of this course.




Cisco provides three levels of general career certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused Cisco Qualified Specialist (CQS) certifications for designated areas such as cable communications, voice, and security.

There are many paths to Cisco certification, but only one requirement—passing one or more exams demonstrating knowledge and skill. For details, go to <http://www.cisco.com/go/certifications>.

Learner Skills and Knowledge

This topic lists the course prerequisites.



To benefit fully from this course, you must have these prerequisite skills and knowledge:

- Cisco CCNA[®] certification or equivalent level of working knowledge and experience
- *Building Scalable Cisco Internetworks* (BSCI)
- *Implementing Cisco Quality of Service* (QOS)
- *Configuring BGP on Cisco Routers* (BGP)
- *Implementing Cisco MPLS* (MPLS)

Note Practical experience with deploying and operating networks based on Cisco network devices and Cisco IOS software is strongly recommended.

Learner Responsibilities

This topic discusses the responsibilities of the learners.

Learner Responsibilities

Cisco.com



- **Complete prerequisites**
- **Introduce yourself**
- **Ask questions**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—7

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This topic lists the administrative issues for the course.

General Administration

Cisco.com

Class-Related	Facilities-Related
• Sign-in sheet • Course materials • Length and times • Attire	• Break and lunch room locations • Site emergency procedures • Rest rooms • Telephones/faxes

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0-8

The instructor will discuss the administrative issues noted here so you know exactly what to expect from the class.

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Flow Diagram

This topic covers the suggested flow of the course materials.

Course Flow Diagram

Cisco.com

	Day 1	Day 2	Day 3	Day 4	Day 5
A M	Course Introduction	Lab	Configuring MPLS Traffic Engineering	MPLS Quality of Service	MPLS IPv6 Support
	MPLS VPN Review			Lab	Case Study
		Lunch			
P M	Carrier Supporting Carrier	MPLS Traffic Engineering Technology	Lab	Any Transport over MPLS	Wrap-up
	Lab			Case Study	


© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—9

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

This topic shows the Cisco icons and symbols used in this course.



Learner Introductions

This is the point in the course where you introduce yourself.

Learner Introductions

Cisco.com

- **Your name**
- **Your company**
- **Skills and knowledge**
- **Brief history**
- **Objective**



© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0-11

Prepare to share the following information:

- Your name
- Your company
- If you have most or all of the prerequisite skills
- A profile of your experience
- What you would like to learn from this course

MPLS VPN Review

Overview

This module reviews the features of Multiprotocol Label Switching (MPLS) and MPLS Virtual Private Networks (VPNs) that were covered in the *Implementing Cisco MPLS* (MPLS) course.

Objectives

Upon completing this module, you will be able to describe basic MPLS concepts and components. This includes being able to do the following:

- Identify and describe the elements of an MPLS implementation
- Describe how packets traverse an MPLS network
- Identify the commands that are necessary to configure MPLS in a frame-mode network
- Identify the components that are necessary to implement VPNs
- Identify the commands that are necessary to implement a simple MPLS VPN

Outline

The module contains these lessons:

- Basic MPLS Concepts
- Label Distribution in Frame-Mode MPLS
- Configuring MPLS
- MPLS VPN Routing Model and Packet Forwarding
- Configuring Simple MPLS VPNs

Basic MPLS Concepts

Overview

This lesson reviews the basic concepts and architecture of MPLS. The lesson provides information about some of the MPLS components and labels. It lays the foundation for subsequent lessons that cover the key areas, such as MPLS Traffic Engineering (MPLS-TE) and VPNs.

Relevance

It is important for the learner to have a good understanding of the underlying concepts of MPLS before proceeding any further in the course.

Objectives

This lesson reviews basic MPLS concepts, including the concepts of the MPLS architecture. Upon completing this lesson, you will be able to do the following:

- Identify the elements of the MPLS conceptual model
- Identify the elements of the MPLS architecture
- Identify the function of the MPLS label
- Identify the different types of label switch routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Implementing Cisco MPLS* (MPLS) course


Outline

This lesson includes these topics:

- Overview
- Basic MPLS Concepts
- MPLS Architecture
- MPLS Labels
- Label Switch Routers
- Lesson Summary
- Lesson Review

Basic MPLS Concepts

This topic provides a review of the underlying concepts of MPLS.



MPLS is a new switching mechanism that uses labels (numbers) to forward packets.

Labels usually correspond to Layer 3 destination addresses (equal to destination-based routing). Labels can also correspond to other parameters, such as quality of service (QoS), source address, or a Layer 2 circuit.


MPLS was designed to support forwarding of other protocols as well. Label switching is performed regardless of the Layer 3 protocol.

The figure illustrates a situation where the intermediary router does not have to perform a time-consuming routing lookup. Instead, this router simply swaps a label with another label (25 is replaced by 23) and forwards the packet based on the received label (23).

In larger networks, the result of MPLS labeling is that only the edge routers perform a routing lookup. All the core routers forward packets based on the labels.

MPLS Architecture

This topic reviews the main components of the MPLS architecture.



- **Router functionality is divided into two major parts: control plane and data plane.**

MPLS consists of two major components:

- **Control plane:** Takes care of the routing information exchange and the label exchange between adjacent devices.
- **Data plane:** Takes care of the forwarding based on either destination addresses or labels. The data plane is also known as the forwarding plane.

A large number of different routing protocols, such as Open Shortest Path First (OSPF), Interior Gateway Routing Protocol (IGRP), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP), can be used in the control plane.

The control plane also requires protocols such as the label exchange protocols, Tag Distribution Protocol (TDP), MPLS Label Distribution Protocol (LDP), and BGP (used by MPLS VPNs), to exchange labels. Resource Reservation Protocol (RSVP) is used by MPLS TE to accomplish this exchange.

The data plane, however, is a simple label-based forwarding engine that is independent of the type of routing protocol or label exchange protocol. The label forwarding information base (LFIB) table is used to forward packets based on labels. The LFIB table is populated by the label exchange protocols (TDP or LDP or both) that are used.

MPLS implements destination-based forwarding that uses labels to make forwarding decisions.

A Layer 3 routing protocol is still needed to propagate Layer 3 routing information. A label exchange mechanism is simply an add-on to propagate labels that are used for Layer 3 destinations.


The figure illustrates the two components of the control plane:

- OSPF, which receives and forwards IP network 10.0.0.0/8.
- LDP, which receives label 17 to be used for packets with a destination address 10.x.x.x. A local label 24 is generated and sent to upstream neighbors so these neighbors can label packets with the appropriate label. LDP inserts an entry into the data plane LFIB table, where label 24 is mapped to label 17.

The data plane then forwards all packets with label 24 through the appropriate interfaces and replaces label 24 with label 17.

MPLS Labels

This topic reviews the structure of an MPLS label and its format.



MPLS is designed for use with virtually any media and Layer 2 encapsulation. Most Layer 2 encapsulations are frame-based, and MPLS simply inserts a 32-bit label between the Layer 2 and Layer 3 headers (“frame-mode” MPLS).


A 32-bit MPLS label contains four fields:

Label Fields

Field	Description
20-bit label	The actual label.
3-bit experimental field	Used to define a class of service (CoS) (IP precedence).
Bottom-of-stack bit	MPLS allows multiple labels to be inserted; this bit determines if this label is the last label in the packet. If this bit is set (1), it indicates that this is the last label.
8-bit time-to-live (TTL) field	Has the same purpose as the TTL field in the IP header.

MPLS Labels: Stack

Cisco.com



- **Protocol ID in a Layer 2 header specifies that the payload starts with a label (labels) and is followed by an IP header.**
- **Bottom-of-stack bit indicates whether the next header is another label or a Layer 3 header.**
- **Receiving router uses the top label only.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-6

A label does not contain any information about the Layer 3 protocol that is being carried in a packet. A new protocol ID (PID) is used for every MPLS-enabled Layer 3 protocol. Usually only one label is assigned to a packet. Here are some scenarios in which more than one label is used:


- **MPLS VPNs:** Multiprotocol Border Gateway Protocol (MP-BGP) is used to propagate a second label that identifies the VPN in addition to the one that is propagated by TDP or LDP to identify the LSP.
- **MPLS TE:** MPLS TE uses RSVP to establish label switched path (LSP) tunnels. RSVP propagates labels that are used to identify the tunnel LSP in addition to the one that is propagated by LDP or TDP to identify the underlying LSP.
- **MPLS VPNs combined with MPLS TE:** Three or more labels are used to identify the VPN, tunnel LSP and the underlying LSP.

The following Ethertype values are used to identify Layer 3 protocols with most Layer 2 encapsulations:

- **Unlabeled IP unicast:** PID = 0x0800 identifies that the frame payload is a classic unicast IP packet.
- **Labeled IP unicast:** PID = 0x8847 identifies that the frame payload is a unicast IP packet with at least one label preceding the IP header. The bottom-of-stack bit indicates when the IP header actually starts.
- **Labeled IP multicast:** PID = 0x8848 identifies that the frame payload is a multicast IP packet with at least one label preceding the IP header. The bottom-of-stack bit indicates when the IP header actually starts.

Label Switch Routers

This topic reviews label switch routers (LSRs).



In preparation for a detailed description of MPLS, here is some of the terminology that is used in this course:

- **Label switch router (LSR):** A device that forwards packets primarily based on labels
- **Edge LSR:** A device that primarily labels packets or removes labels

LSRs and edge LSRs are usually capable of doing both label switching and IP routing. Their names are based on their positions in an MPLS domain. Routers that have all interfaces enabled for MPLS are called LSRs because they mostly forward labeled packets. Routers that have some interfaces that are not enabled for MPLS are usually at the edge of an MPLS domain (autonomous system, or AS). These routers also forward packets based on IP destination addresses and label them if the outgoing interface is enabled for MPLS.

LSRs of all types must perform these functions:


- Exchange Layer 3 routing information (ATM LSRs must also exchange Layer 3 routing information) (control plane)
- Exchange labels (control plane)
- Forward packets or cells (data plane):
 - Frame-mode MPLS forwards packets based on the 32-bit label.
 - Cell-mode MPLS forwards packets based on labels that are encoded into the virtual path identifier (VPI) and virtual channel identifier (VCI) fields in the ATM header.

The data plane performs the following functions:

- Exchanges routing information regardless of the type of LSR
- Exchanges labels according to the type of MPLS (frame-mode or cell-mode)

Architecture of LSRs

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-1-8

The primary function of an LSR is to forward labeled packets. Therefore, every LSR needs a Layer 3 routing protocol (for example, OSPF, EIGRP, or IS-IS) and a label distribution protocol (for example, LDP or TDP).

LDP populates the LFIB table in the data plane that is used to forward labeled packets.

Note LSRs may not be able to forward unlabeled packets either because they are ATM LSRs or because they do not have all the routing information.

Edge LSRs also forward IP packets based on their IP destination addresses and optionally label them if a label exists.

The following combinations are possible:

- A received IP packet is forwarded based on the IP destination address and sent as an IP packet.
- A received IP packet is forwarded based on the IP destination address and sent as a labeled packet.
- A received labeled packet is forwarded based on the label; the label is changed and the packet is sent.

The following scenarios are possible if the network is not configured properly:

- A received labeled packet is dropped if the label is not found in the LFIB table, even if the IP destination exists in the IP forwarding table (also called the Forwarding Information Base, or FIB).
- A received IP packet is dropped if the destination is not found in the IP forwarding table (FIB table), even if there is an MPLS LSP toward the destination.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **MPLS forwards packets based on labels.**
- **Packets are forwarded using labels from the LFIB table rather than the IP routing table.**
- **Labels are inserted between the Layer 2 and Layer 3 headers in frame-mode networks and use the VPI and VCI field in cell-mode networks.**
- **All LSRs perform three functions:**
 - **Exchange routing information**
 - **Exchange labels**
 - **Forward packets or cells (depending on type)**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—1-9

References

For additional information, refer to this resource:

- RFC 3031, *Multiprotocol Label Switching Architecture*, at the following URL:
http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Technologies:MPLS&s=Overview#Standards_and_Specifications

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The label distribution protocol (either TDP or LDP) is the responsibility of which plane?
- A) data plane
 - B) forwarding plane
 - C) system plane
 - D) control plane
- Q2) Which of the following is NOT true?
- A) MPLS uses labels to forward packets.
 - B) MPLS works only in IP networks.
 - C) MPLS labels can correspond to Layer 3 destination address, QoS, source address, or Layer 2 circuit.
 - D) MPLS does not require a routing table lookup on core routers.
- Q3) The MPLS label field consists of how many bits?
- A) 64 bits
 - B) 32 bits
 - C) 16 bits
 - D) 8 bits
- Q4) Which two of the following statements are true? (Choose two.)
- A) A PE Router (or edge LSR) is a device that primarily inserts labels on packets or removes labels.
 - B) A P Router (or LSR) is a device that primarily labels packets or removes labels.
 - C) A P Router (or LSR) is a device that forwards packets primarily based on labels.
 - D) A PE Router (or edge LSR) is a device that forwards packets primarily based on labels.

Lesson Answer Key

Q1) D

Relates to: MPLS Architecture

Q2) B

Relates to: Basic MPLS Concepts

Q3) B

Relates to: MPLS Labels

Q4) A, C

Relates to: Label Switch Routers

Label Distribution in Frame-Mode MPLS

Overview

This lesson reviews MPLS label allocation in a frame-mode environment, and how distribution of labels works in this environment. Also reviewed are penultimate hop popping (PHP) and how the MPLS data structures are built.

Relevance

It is important for the learner to have a good understanding of the basic fundamentals of how information gets distributed and placed into the appropriate tables for both labeled and unlabeled packet usage.

Objectives

This lesson reviews basic MPLS concepts, including the concepts of the MPLS architecture. Upon completing this lesson, you will be able to do the following:

- Identify how label allocation is performed in a frame-mode MPLS network
- Identify how labels are distributed in a frame-mode network
- Identify how the LFIB table is populated
- Identify packet propagation across an MPLS network
- Identify the function of penultimate hop popping

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Implementing Cisco MPLS* (MPLS) course


Outline

This lesson includes these topics:

- Overview
- Label Allocation in a Frame-Mode MPLS Network
- Label Distribution and Advertisement
- Populating the LFIB
- Packet Propagation Across an MPLS Network
- Penultimate Hop Popping
- Lesson Summary
- Lesson Review

Label Allocation in a Frame-Mode MPLS Network

This topic provides a review of how labels are allocated in a frame-mode MPLS network.



Unicast IP routing and MPLS functionality can be divided into the following steps:

- **Exchange of routing information:** Routing information is exchanged using standard or vendor-specific IP routing protocols (OSPF, IS-IS, EIGRP, and so on).
- **Generation of local labels:** One locally unique label is assigned to each IP destination that is found in the main routing table and is then stored in the label information base (LIB) table.
- **Propagation of local labels to adjacent routers:** These labels can be used as next-hop labels (stored in the FIB and LFIB tables to enable label switching).

The following data structures contain label information:

- The LIB, in the control plane, is the database that is used by LDP to store the labels assigned to IP prefixes. These labels can be either locally significant or learned from a downstream neighbor.
- The LFIB, in the data plane, is the database that is used to forward labeled packets. Local labels, which were previously advertised to upstream neighbors, are mapped to next-hop labels, which were previously received from downstream neighbors.
- The FIB, in the data plane, is the database that is used to forward unlabeled IP packets. A forwarded packet is labeled if a next-hop label is available for a specific destination IP network. Otherwise, a forwarded packet is not labeled.

The figure illustrates how all routers learn about network X via an IGP such as OSPF, IS-IS, or EIGRP. The FIB table on router A contains the entry for network X, which is mapped to the IP next-hop address B. At this time a next-hop label is not available, which means that all packets are forwarded in a traditional fashion (as unlabeled packets).

The figure shows how router B generates a locally significant and locally unique label 25, which is assigned to IP network X. Router B generates this label regardless of other routers (asynchronous allocation of labels).

Note	Labels 0 – 15 are reserved.


When a label is assigned to an IP prefix, it is stored in two tables:

- The LIB table is used to maintain the mapping between the IP prefix (network X), the local label (25), and the next-hop label (not available yet).
- The LFIB table is modified to contain the local label, which is mapped to the pop action (label removal). The pop action is used until the next-hop label is received from the downstream neighbor.

Note	The pop action results in labels being removed rather than swapped. Therefore, if a labeled packet is received, the label is removed and the remainder of the packet (which might be labeled if the incoming packet contained a label stack) is forwarded to the appropriate IP next hop. A similar action is the untagged action, which is equivalent to the pop action with one exception: the resulting packet <i>must be unlabeled</i> ; otherwise, it will be discarded.

Label Distribution and Advertisement

This topic provides a review of how labels are distributed and advertised in a frame-mode MPLS network.




The figure illustrates the next step after a local label has been generated. Router B propagates this label, 25, to all adjacent neighbors, where it can be used as a next-hop label.

Note Because router B cannot predict which routers might use it as the downstream neighbor, it sends its local mappings to all LDP neighbors.

Receiving Label Advertisement

Cisco.com



- **Every LSR stores the received label in its LIB.**
- **Edge LSRs that receive the label from their next-hop LSR also store the label information in the FIB.**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—1-5

Upon receiving an LDP update, router A can fill in the missing piece in its LIB, LFIB, and FIB tables:

- Label 25 is stored in the LIB table as the label for network X that has been received from LSR B.
- Label 25 is attached to the IP forwarding entry in the FIB table to enable MPLS edge functionality (incoming IP packets are forwarded as labeled packets).
- The local label in the LFIB table is mapped to outgoing label 25 instead of the pop action (incoming labeled packets can be forwarded as labeled packets).

Receiving Label Advertisement (Cont.)

Cisco.com



- Every LSR stores received information in its LIB.
- LSRs that receive their label from their next-hop LSR will also populate the IP forwarding table (FIB).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.0—1-6


Because all routers in an MPLS domain asynchronously do the same as routers A and B, an LSP tunnel is generated, which spans from router A to router D.

The figure illustrates how an LDP update from router C, which is advertising label 47 for network X, is sent to all adjacent routers, including router B.

Router B can now map the entry for network X in its FIB, and the local label 25 in its LFIB, to the next-hop label 47 that has been received from the downstream neighbor router C.

Populating the LFIB

This topic reviews how the LFIB table is populated in an MPLS network.




After router C advertises label 47 to adjacent routers, the LSP tunnel for network X has two hops:

- On router A, network X is mapped to the next-hop label 25 (router B).
- On router B, label 25 is mapped to the next-hop label 47 (router C).
- Router C still has no next-hop label. Label 47 is therefore still mapped to the pop action.

Note In the figure label distribution is from right to left and packet forwarding is from left to right.

Packet Propagation Across an MPLS Network

This topic reviews how IP packets traverse an MPLS network.




The figure illustrates how IP packets are propagated across an MPLS domain. The steps are as follows:

1. Router A labels a packet that is destined for network X by using the next-hop label 25 (it does this by performing Cisco Express Forwarding [CEF] switching by using the FIB table).
2. Router B swaps label 25 with label 47 and forwards the packet to router C (it does this by performing label switching by using the LFIB table).
3. Router C removes the label and forwards the packet to router D (it does this by performing label switching by using the LFIB table).

Penultimate Hop Popping

This topic reviews penultimate hop popping (PHP).




The figure illustrates how labels are propagated and used in a typical frame-mode MPLS network. The check marks show which tables are used on individual routers. The egress router in this example must do a lookup in the LFIB table to determine whether the label must be removed and if a further lookup in the FIB table is required.

PHP removes the requirement for a double lookup to be performed on egress LSRs.

Penultimate Hop Popping: After

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-10

The figure illustrates how a predefined label “pop,” which corresponds to the pop action in the LFIB, is propagated on the first hop or the last hop, depending on the perspective. The term “pop” means to remove the top label in the MPLS label stack instead of swapping it with the next-hop label. The last router before the egress router therefore removes the top label.

PHP slightly optimizes MPLS performance by eliminating one LFIB lookup.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **MPLS forwards packets based on labels.**
- **Packets are forwarded by using labels from the LFIB table rather than the IP routing table.**
- **LSRs store received labels in the LIB table.**
- **Outgoing labels are stored in the LFIB table.**
- **PHP optimizes MPLS performance.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—1-11

References

For additional information, refer to these resources:

- RFC 3031, *Multiprotocol Label Switching Architecture*, and RFC 3036, *LDP Specification*, at the following URL: http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Technologies:MPLS&s=Overview#Standards_and_Specifications

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which of the following best describes PHP?
- A) PHP works only for TDP and not for LDP.
 - B) PHP works only for LDP and not for TDP.
 - C) PHP optimizes MPLS performance.
 - D) PHP is configurable and by default is disabled.
- Q2) Which three of the following are contained in the LFIB? (Choose three.)
- A) local generated label
 - B) outgoing label
 - C) incoming label
 - D) next-hop address
- Q3) When an IP packet is to be label-switched as it traverses an MPLS network, which of the following tables will be used to perform the label switching?
- A) LIB
 - B) FIB
 - C) FLIB
 - D) LFIB
- Q4) Which two of the following tables contain label information? (Choose two.)
- A) LIB
 - B) main IP routing label
 - C) FLIB
 - D) LFIB
- Q5) Which of the following generates a label update?
- A) UDP
 - B) OSPF
 - C) EIGRP
 - D) LDP

Lesson Answer Key

Q1) C

Relates to: Penultimate Hop Popping

Q2) A,, B, D

Relates to: Populating the LFIB

Q3) D

Relates to: Packet Propagation Across an MPLS Network

Q4) A, D

Relates to: Label Allocation in a Frame-Mode MPLS Network

Q5) D

Relates to: Label Distribution and Advertisement

Configuring MPLS

Overview

This lesson reviews the steps that are required to successfully implement MPLS.

Relevance

It is important for the learner to have a good understanding of the basic configuration tasks that are necessary to enable MPLS.

Objectives

This lesson reviews the basic MPLS configuration commands. Upon completing this lesson, you will be able to do the following:

- Enable IP CEF
- Enable MPLS on a frame-mode interface
- Configure the MTU size in label switching

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Implementing Cisco MPLS* (MPLS) course

Outline

This lesson includes these topics:

- Overview
- Configuring IP CEF
- Configuring MPLS on a Frame-Mode Interface
- Configuring the MTU Size in Label Switching
- Lesson Summary
- Lesson Review

Configuring IP CEF

This topic reviews the steps that are necessary to enable CEF for MPLS operation.

Configuring IP CEF

Cisco.com

```
Router(config)#  
ip cef [distributed]
```

- Starts CEF switching and creates the FIB table.
- The distributed keyword configures distributed CEF (running on VIP or line cards).
- All CEF-capable interfaces run CEF switching.

```
Router(config-if)#  
no ip route-cache cef
```

- Disables CEF switching on an interface.
- Usually not needed.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—1-3

ip cef

To enable CEF on the route processor card, use the **ip cef** global command in global configuration mode. To disable CEF, use the **no** form of this command:

- **ip cef [distributed]**
- **no ip cef [distributed]**

Syntax Description

distributed (Optional): Enables distributed CEF operation. Distributes the CEF information to the line cards. The line cards perform express forwarding.

ip route-cache cef

To enable CEF operation on an interface after the CEF operation has been disabled, use the **ip route-cache cef** command in interface configuration mode. To disable CEF operation on an interface, use the **no** form of this command:

- **ip route-cache cef**
- **no ip route-cache cef**

Note	Check your router IOS version to determine application we it may be necessary to disable caching.

Syntax Description

This command has no arguments or keywords.

Defaults

When standard CEF or distributed CEF operations are enabled globally, all interfaces that support CEF are enabled by default.

Configuring MPLS on a Frame-Mode Interface

This topic reviews the steps that are necessary to enable MPLS on a Cisco IOS router.

Configuring MPLS on a Frame-Mode Interface

Cisco.com

```
Router(config-if)#  
mpls ip
```

- Enables label switching on a frame-mode interface
- Starts LDP on the interface

```
Router(config-if)#  
mpls label protocol [tdp | ldp | both]
```

- Starts selected label distribution protocol on the specified interface

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—1-4

mpls ip

To enable label switching of IP version 4 (IPv4) packets on an interface, use the **mpls ip** command in interface configuration mode. To disable IP label switching on this interface, use the **no** form of this command:

- **mpls ip**
- **no mpls ip**

Syntax Description

This command has no arguments or keywords.

Defaults

Label switching of IPv4 packets is disabled on this interface.

mpls label protocol [tdp | ldp | both]

To select the label distribution protocol to be used on an interface, use the **mpls label protocol** command in interface configuration mode. To revert to the default label distribution protocol, use the **no** form of this command:

- **mpls label protocol [protocol]**
- **no mpls label protocol [protocol]**

Parameter	Description
tdp	Enables TDP on an interface
ldp	Enables LDP on an interface
both	Enables TDP and LDP on an interface

Defaults

TDP is the default protocol.


Note	For backward compatibility, the MPLS syntax will be entered in the tag-switching syntax in the configuration by the IOS software.

Configuring MPLS on a Frame-Mode Interface:

Cisco.com

Enable MPLS on all core interfaces in your network.

```
ip cef
interface serial 2/1
 mpls ip
interface fastethernet 0/0
 mpls ip
```



Use access lists to prevent customers from running TDP with your routers.

```
ip cef
interface serial 3/1
 ip access-group NoTDP in
interface serial 2/1
 mpls ip

ip access-list NoTdp deny tcp any any eq 711
ip access-list NoTdp permit ip any any
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-1-5

The figure shows the configuration steps that are needed to enable MPLS on an edge LSR. The configuration includes an access list that denies any attempt to establish a TDP session from an interface that is not enabled for MPLS.

You must globally enable CEF switching, which automatically enables CEF on all interfaces that support it. (CEF is not supported on logical interfaces, such as loopback interfaces.)

In this example, nonbackbone interfaces have an input access list that denies TCP sessions on the well-known port number 711 (TDP). This action provides security by blocking TDP advertisements. This action could also be done for LDP.

Configuring the MTU Size in Label Switching

This topic reviews how to change the maximum transmission unit (MTU) size in label switching.

Configuring the MTU Size in Label-Switching

Cisco.com

```
Router(config-if)#  
mpls mtu bytes
```

- Label switching increases the maximum MTU requirements on an interface, because of the additional label header.
- The interface MTU is automatically increased on WAN interfaces; the IP MTU is automatically decreased on LAN interfaces.
- The label-switching MTU can be increased on LAN interfaces (resulting in jumbo frames) to prevent IP fragmentation.
- The jumbo frames are not supported by all LAN switches.

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—1-6

mpls mtu

To set the per-interface MTU for labeled packets, use the **mpls mtu** interface configuration command:

- **mpls mtu** *bytes*
- **no mpls mtu**

Parameter	Description
<i>bytes</i>	MTU in bytes

Defaults


The minimum MTU is 128 bytes. The maximum depends on the type of interface medium.

Configuring the MTU Size in Label-Switching: Example

Cisco.com

Jumbo frames have to be enabled on the switch.

```
set port 1/3 jumbo enable
set port 1/4 jumbo enable
```



MPLS MTU is increased to 1512 to support 1500-byte IP packets and MPLS stack up to three levels deep.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-1.7

One way of preventing labeled packets from exceeding the maximum size (and being fragmented as a result) is to increase the MTU size of labeled packets for all segments in the LSP tunnel. The problem of packets that exceed the maximum size will typically occur on LAN switches, where it is more likely that a device does not support oversized packets (also called jumbo frames or, sometimes, giants or baby giants). Some devices support jumbo frames, and some need to be configured to support them.

The MPLS MTU size is increased automatically on WAN interfaces and needs to be increased manually on LAN interfaces.

The MPLS MTU size has to be increased on all LSRs that are attached to a LAN segment. Additionally, the LAN switches that are used to implement switched LAN segments need to be configured to support jumbo frames. No additional configuration is necessary for shared LAN segments that are implemented with hubs.

A different approach is needed if a LAN switch does not support jumbo frames. The problem may be even worse for networks that do not allow Internet Control Message Protocol (ICMP) MTU discovery messages to be forwarded to sources of packets and if the Don't Fragment (DF) bit is strictly used. This situation can be encountered where firewalls are used.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- CEF is configured globally.
- Use `mpls ip` to enable MPLS (interface level).
- Label switching increases the maximum MTU size on an interface.

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—1-8

References

For additional information, refer to these resources:

- Configuring MPLS and VPN
http://www.cisco.com/en/US/products/hw/modules/ps2797/products_module_installation_guide_chapter09186a0080080cf1.html
- CEF
http://www.cisco.com/en/US/tech/tk827/tk831/tk102/tech_protocol_home.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which command do you use to enable MPLS in Cisco IOS software?
- A) **Router# mpls ip**
 - B) **Router> mpls ip**
 - C) **Router(config)# mpls ip**
 - D) **Router(config-if)# mpls ip**
- Q2) The MPLS MTU is increased to _____ to support 1500-byte IP packets and MPLS stacks up to three levels deep.
- Q3) Which command do you use to enable CEF on a Cisco router?
- A) **Router# ip cef**
 - B) **Router> ip cef**
 - C) **Router(config)# cef**
 - D) **Router(config)# ip cef**

Lesson Answer Key

Q1) D

Relates to: Configuring MPLS on a Frame-Mode Interface

Q2) 1512

Relates to: Configuring the MTU Size in Label Switching

Q3) D

Relates to: Configuring IP CEF

MPLS VPN Routing Model and Packet Forwarding

Overview

This lesson reviews MPLS VPN architecture and operation. It details various customer edge-provider edge (CE-PE) routing options and Border Gateway Protocol (BGP) extensions (route targets and extended community attributes) that allow Internal Border Gateway Protocol (IBGP) to transport customer routes over a provider network. The MPLS VPN forwarding model is also covered together with its integration with core routing protocols.

Relevance

It is important to understand how the MPLS VPN architecture is structured, what the components of that architecture are, and how the components are used.

Objectives

This lesson describes the major architectural components of an MPLS VPN, identifying the functions of route information propagation, route distinguishers (RDs), route targets (RTs), and virtual routing tables. Upon completing this lesson, you will be able to do the following:

- Identify how routing information is propagated across the P-network
- Identify the need for route distinguishers and how they are implemented
- Identify the need for route targets and how they are implemented
- Identify the end-to-end flow of routing updates in an MPLS VPN
- Describe MPLS VPN packet forwarding

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Implementing Cisco MPLS (MPLS)* course


Outline

This lesson includes these topics:

- Overview
- Propagation of Routing Information Across the P-Network
- Route Distinguishers
- Route Targets
- End-to-End Routing Information Flow
- MPLS VPNs and Packet Forwarding
- Lesson Summary
- Lesson Review

Propagation of Routing Information Across the P-Network

This topic reviews the different methods of propagating routing information across the provider network (P-network), and identifies the drawbacks of each method.




The best solution to the customer route propagation issue is to run a single routing protocol between provider edge routers (PE routers) that will exchange all customer routes without the involvement of the provider routers (P routers). This solution is scalable:

- The number of routing protocols running between PE routers does not increase with an increasing number of customers.
- The P routers do not carry customer routes.

The next design decision to be made is the choice of the routing protocol that runs between PE routers. Given that the total number of customer routes is expected to be very large, the only well-known protocol with the required scalability is BGP. And, in fact, BGP is used in MPLS VPN architecture to transport customer routes directly between PE routers.

Propagation of Routing Information Across the P-Network (Cont.)

Cisco.com



Q: How will information about the overlapping subnets of two customers be propagated via a single routing protocol?

A: Extend the customer addresses to make them unique.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-4

MPLS VPN architecture differs in an important way from traditional peer-to-peer VPN solutions: the support of overlapping customer address spaces.

With the deployment of a single routing protocol (BGP) that exchanges all customer routes between PE routers, an important issue arises: how can BGP propagate several identical prefixes, belonging to different customers, between PE routers?

The only solution to this dilemma is the expansion of customer IP prefixes with a unique prefix that makes them unique even if they had previously overlapped. A 64-bit prefix called the route distinguisher (RD) is used in MPLS VPNs to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported between PE routers.

Route Distinguishers

This topic reviews what an RD is and how it is used.

Route Distinguishers

Cisco.com

- The 64-bit route distinguisher (RD) is prepended to an IPv4 address to make it globally unique.
- The resulting address is a VPNv4 address.
- VPNv4 addresses are exchanged between PE routers via BGP.
 - BGP that supports address families other than IPv4 addresses is called Multiprotocol BGP (MP-BGP).

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—1-5


The RD is used only to transform non-unique 32-bit customer IP version 4 (IPv4) addresses into unique 96-bit VPNv4 addresses (also called VPN IPv4 addresses).

VPNv4 addresses are exchanged only between PE routers; they are never used between customer edge (CE) routers. BGP between PE routers must therefore support exchange of traditional IPv4 prefixes as well as exchange of VPNv4 prefixes. A BGP session between PE routers is consequently called a Multiprotocol BGP (MP-BGP) session.

Note	Initial MPLS VPN implementation in Cisco IOS software supports only MPLS VPN services within a single autonomous system (AS). In such a scenario, the BGP session between PE routers is always an IBGP session.

Route Distinguishers (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0-1-6

Customer route propagation across an MPLS VPN network is done using the following process:

- Step 1** The CE router sends an IPv4 routing update to the PE router.
- Step 2** The PE router prepends a 64-bit RD to the IPv4 routing update, which results in a globally unique 96-bit VPNv4 prefix.
- Step 3** The VPNv4 prefix is propagated via a Multiprotocol Internal Border Gateway Protocol (MP-IBGP) session to other PE routers.

Route Distinguishers (Cont.)

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.0—1-7

- Step 4** The receiving PE routers strip the RD from the VPNv4 prefix, which results in an IPv4 prefix.
- Step 5** The IPv4 prefix is forwarded to other CE routers within an IPv4 routing update.

Route Targets

This topic reviews why there is a need for route targets (RTs) and how they are implemented.



The connectivity requirements of the Voice over IP (VoIP) service are illustrated in the figure. Three VPNs are needed to implement the desired connectivity: two customer VPNs and a shared VoIP VPN. Central customer sites participate in the customer VPN as well as in the VoIP VPN.

The RD (again, a single entity that is prepended to an IPv4 route) cannot indicate that a site participates in more than one VPN. A method is needed in which a set of VPN identifiers can be attached to a route to indicate its membership in several VPNs.

RTs were introduced into the MPLS VPN architecture to support this requirement.

Route Targets: What Are They?

Cisco.com

- **RTs are additional attributes attached to VPNv4 BGP routes to indicate VPN membership.**
- **Extended BGP communities are used to encode these attributes.**
 - **Extended communities carry the meaning of the attribute together with its value.**
- **Any number of RTs can be attached to a single route.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.0—1-9

RTs are attributes that indicate VPN membership. When VPNv4 BGP routing updates are sent to peers, the extended BGP communities of each update is used to carry the appropriate RT. Therefore, the RT in the extended community identifies which VPN that the update belongs to.

As with standard BGP communities, a set of extended communities can be attached to a single BGP route, which satisfies the requirements of complex VPN topologies.

Extended BGP communities are 64-bit values. The semantics of the extended BGP community are encoded in the high-order 16 bits of the value, which makes them useful for a number of different applications, such as MPLS VPN RTs.

Route Targets: How Do They Work?

Cisco.com

- **Export RTs**
 - Identifying VPN membership
 - Appended to the customer route when it is converted into a VPNv4 route
- **Import RTs**
 - Associated with each virtual routing table
 - Select routes to be inserted into the virtual routing table

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-10


MPLS VPN RTs are attached to a customer route at the moment that it is converted from an IPv4 route to a VPNv4 route by the PE router. The RTs that are attached to the route are called export RTs and are configured separately for each virtual routing table in a PE router. Export RTs identify a set of VPNs in which sites that are associated with the virtual routing table belong.

When the VPNv4 routes are propagated to other PE routers, those routers need to select the routes to import into their virtual routing tables. This selection is based on import RTs. Each virtual routing table in a PE router can have a number of configured import RTs that identify the set of VPNs from which the virtual routing table is accepting routes.

In overlapping VPN topologies, RTs are used to identify VPN membership. Advanced VPN topologies (for example, central services VPNs) use RTs in more complex scenarios.

End-to-End Routing Information Flow

This topic reviews the end-to-end flow of routing updates in an MPLS VPN implementation.




The next few figures provide an overview of end-to-end routing information flow in an MPLS VPN network. The PE routers receive IPv4 routing updates from the CE routers and install them in the appropriate virtual routing and forwarding (VRF) table.

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers.

Initial MPLS VPN implementation in Cisco IOS software (Cisco IOS Releases 12.0 T and 12.1) supports MPLS VPN services only within the scope of a single AS. The MP-BGP sessions between the PE routers are therefore IBGP sessions and are subject to the IBGP split-horizon rules. Thus, either a full mesh of MP-IBGP sessions is required between PE routers, or route reflectors need to be used to reduce the full-mesh IBGP requirement.

End-to-End Routing Update Flow (Cont.)

Cisco.com



- **PE routers receive IPv4 routing updates from CE routers and install them in the appropriate VRF table.**
- **PE routers export VPN routes from VRF tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-12


An MP-BGP update exchange between PE routers contains the following:

- VPNv4 address
- Extended BGP communities (RTs required; Site of Origin, or SOO, optional)
- Label used for VPN packet forwarding (an explanation of how the label is used is provided later in this lesson).
- Mandatory BGP attributes (for example, the AS path)

Optionally, the MP-BGP update can contain any other BGP attribute; for example, local preference, multi-exit discriminator (MED), or standard BGP community.

End-to-End Routing Update Flow (Cont.)

Cisco.com



- **Receiving PE router imports incoming VPNv4 routes into the appropriate VRF table based on route targets attached to the routes.**
- **Routes installed in VRF table are propagated to CE routers.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-13

The PE routers that receive MP-BGP updates import the incoming VPNv4 routes into their VRF tables based on RTs that are attached to the incoming routes and based on import RTs that are configured in the VRF tables. The VPNv4 routes that are installed in VRF tables are converted to IPv4 routes and then propagated to the CE routers.

MPLS VPNs and Packet Forwarding

This topic reviews the effects of MPLS VPNs on packet forwarding.

MPLS VPNs and Packet Forwarding

Cisco.com

- **The VPN label is understood only by the egress PE router.**
- **An end-to-end LSP tunnel is required between the ingress and egress PE routers.**
- **BGP next hops must not be announced as BGP routes.**
- **LDP labels are not assigned to BGP routes.**
- **BGP next hops announced in IGP must not be summarized in the core network.**
 - **Summarization breaks the LSP tunnel.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-14


For successful propagation of MPLS VPN packets across an MPLS backbone, there must be an unbroken LSP tunnel between PE routers. This need exists because the second label in the stack is recognized only by the egress PE router that has originated it and will not be understood by any other router should it ever become exposed.

Two scenarios could cause the LSP tunnel between PE routers to break:

- If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label and the label stack will not be built correctly.
- If the P routers perform summarization of the address range within which the IP address of the egress PE router lies, the LSP tunnel will be disrupted at the summarization point, as illustrated in the next figure.

MPLS VPNs and Packet Forwarding: Summarization in the Core

Cisco.com



In the figure, the P router summarizes the loopback address of the egress PE router. The LSP tunnel is broken at a summarization point, so the summarizing router needs to perform a full IP lookup. In a frame-based MPLS network, the P router requests PHP for the summary route, and the upstream P router (or a PE router) removes the LDP label, exposing the VPN label to the P router. Because the VPN label is not assigned by the P router but by the egress PE router instead, the label will not be understood by the P router and the VPN packet will be dropped or misrouted.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **BGP is used to exchange customer routes between PE routers.**
- **Route distinguishers transform non-unique 32-bit addresses into 96-bit unique address.**
- **Route targets are used to identify VPN membership when overlapping topologies exist.**
- **Routes are transported using the following:**
 - **IGP (internal core routes)**
 - **BGP IPv4 (core Internet routes)**
 - **BGP VPNv4 (PE-to-PE VPN routes)**
- **LDP labels are not assigned to BGP routes.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-16

References

For additional information, refer to these resources:

- **Configuring MPLS and VPN**
http://www.cisco.com/en/US/products/hw/modules/ps2797/products_module_installation_guide_chapter09186a0080080cf1.html
- **Overview: VPN Service Model**
http://www.cisco.com/en/US/products/sw/netmgts/ps2327/products_programming_usage_guide_chapter09186a00800eab87.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) Which protocol is used to transport customer routes directly between PE routers?

- A) RIP
- B) VPN
- C) BGP
- D) OSPF

Q2) What is the function of the RD in an MPLS VPN?

Q3) What is the function of the RT in an MPLS VPN?

Q4) Which BGP function do MPLS VPNs use to transport RTs?

Q5) What scenarios would cause the LSP tunnel between PE routers to break?

Lesson Answer Key

Q1) C

Relates to: Propagation of Routing Information Across the P-Network

Q2) The RD is used to transform the non-unique IP addresses of the customer into unique VPNv4 addresses.

Relates to: Route Distinguishers

Q3) The RT attaches a set of VPN identifiers to a route that indicate its membership in several VPNs. This capability allows one site to be a member of more than one VPN.

Relates to: Route Targets

Q4) extended communities

Relates to: End-to-End Routing Information Flow

Q5) If the IP address of the PE router is announced as a BGP route, it will have no corresponding LDP label and the label stack will not be built correctly.

If the P routers perform summarization of the address range within which the IP address of the egress PE router lies, the LSP tunnel will be disrupted at the summarization point.

Relates to: MPLS VPNs and Packet Forwarding

Configuring Simple MPLS VPNs

Overview

This lesson reviews the steps that are required to successfully implement a simple MPLS VPN.

Relevance

It is important to know how to configure and successfully implement an MPLS VPN. It is essential to understand the command syntax for the configurations that you want to deploy in your network. This lesson will provide a review of the information that will enable you to succeed at such tasks.

Objectives

This lesson describes the command syntax that is required to configure a simple MPLS VPN. Upon completing this lesson, you will be able to do the following:

- Configure a VRF table
- Configure an MP-BGP session between PE routers
- Configure the routing protocols between PE and CE routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the *Implementing Cisco MPLS (MPLS)* course


Outline

This lesson includes these topics:

- Overview
- Configuring VRF Tables
- Configuring an MP-BGP Session Between PE Routers
- Configuring Routing Protocols Between PE and CE Routers
- Lesson Summary
- Lesson Review

Configuring VRF Tables

This topic provides a review of the tasks that are required to configure VRF tables.



Configuring a VRF table and starting deployment of an MPLS VPN service for a customer consists of four mandatory steps:

- Create a new VRF table.
- Assign a unique RD to the VRF table.

Note You must assign a unique RD to every VRF table that is created in a PE router. The same RD *might* be used in multiple PE routers, based on customer connectivity requirements. Refer to the previous lesson for more details on RD assignment for different VPN topologies.

- Specify import and export RTs for the VRF table.

Note Import and export RTs should be equal to the RD for simple VPN service.

- Assign interfaces to VRF tables.


To illustrate the use of MPLS VPN configuration commands, the figure shows a configuration of the PE router in a sample network with two VPN customers. Customer A (with four sites) is using BGP and RIP as the provider edge-customer edge (PE-CE) routing protocols, and customer B (with two sites) is using only RIP. Both customers are using private IP address space (subnets of network 10.0.0.0).

The configuration steps that you can perform on the PE router so far are as follows:

- Step 1** Configure VRF tables for customer A and customer B.
- Step 2** Assign RDs and RTs to the VRF tables. Only one RD per customer is used on all PE routers in the MPLS VPN backbone because these customers require only simple VPN connectivity. To simplify the configuration and troubleshooting process, the RTs are made equal to the RDs.
- Step 3** Assign PE-CE interfaces to individual VRF tables.

Configuring an MP-BGP Session Between PE Routers

This topic identifies the command syntax that is required to configure BGP address families.



MPLS VPN architecture defines two types of BGP neighbors:

- Global BGP neighbors (other PE routers), with which the PE router can exchange multiple types of routes. These neighbors are defined in the global BGP definition and only have to be *activated* for individual address families.
- Per-VRF-table BGP neighbors (the CE routers), which are configured and activated within the **ipv4 vrf vrf-name** address family.

Configure BGP connectivity between two PE routers in four steps:

- Step 1** Configure the remote PE router as a global BGP neighbor in BGP router configuration mode.
- Step 2** Define the parameters that affect all BGP route exchange (for example, source address for the TCP session) on the global BGP neighbor.
- Step 3** Select the VPNv4 address family and activate the BGP neighbor for VPNv4 route exchange.
- Step 4** Configure additional VPNv4-specific BGP parameters (filters, next-hop processing, route maps) within the VPNv4 address family.

Note IPv4-specific BGP parameters are still configured under the BGP router configuration mode—there is no special IPv4 address family.

Configuring Routing Protocols Between PE and CE Routers

This topic identifies the requirements for configuring PE-CE routing protocols.

Configuring Routing Protocols Between PE and CE Routers

Cisco.com

- PE-CE routing protocols are configured for individual VRF tables.
- Per-VRF-table routing protocols can be configured in two ways:
 - There is only one BGP, EIGRP, or RIP process per router.
 - Per-VRF-table parameters are specified in routing contexts, which are selected with the address-family command.
 - A separate OSPF process has to be started for each VRF table.
- Overall number of routing processes per router is limited to 32. Only 28 are available for VRF assignment.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-5

After you configure the VRF tables and establish MP-IBGP connectivity between PE routers, you have to configure routing protocols between the PE router and the attached CE routers. The PE-CE routing protocols need to be configured for individual VRF tables. Sites in the same VPN but in different VRF tables cannot share the same PE-CE routing protocol.

Note The per-VRF-table configuration of the PE-CE routing protocols is another good reason for grouping as many sites into a VRF table as possible.


The per-VRF-table routing protocols can be configured in two ways:

- As individual address families belonging to the same routing process (similar to what you have already seen for BGP).
- As separate routing processes. This option is used for more complex routing protocols that need to maintain a separate topology database for each VRF table (for example, OSPF).

Note Current Cisco IOS software implementation limits the overall number of routing protocols in a router to 32. Two routing methods are predefined (static and connected), and two routing protocols are needed for proper MPLS VPN backbone operation (BGP and backbone Interior Gateway Protocol [IGP]). The number of PE-CE routing processes is therefore limited to 28.

Configuring Routing Protocols Between PE and CE Routers: BGP

Cisco.com



The figure here shows BGP activated on the CE router. The PE router is defined as a BGP neighbor. Similarly, the CE router is defined as a BGP neighbor and activated under **address-family ipv4 vrf Customer_A**.

Select the VRF table routing context with the **address-family ipv4 vrf vrf-name** command in the RIP and BGP routing processes. All per-VRF-table routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

When you configure BGP as the PE-CE routing protocol, you must start with the per-VRF-table BGP configuration. Use the **address-family ipv4 vrf vrf-name** command in router configuration mode. Enter the address family configuration mode, and then define and activate the BGP neighbors. You also need to configure redistribution from all other per-VRF-table routing protocols into BGP.

Note You always have to configure a BGP address family for each VRF table and configure route redistribution into BGP for each VRF table even if you do not use BGP as the PE-CE routing protocol.


Several BGP options have different default values when you configure the per-VRF-table BGP routing context:

- BGP synchronization is disabled (default = enabled).
- Autosummation (automatic generation of classful networks out of subnets that are redistributed into BGP) is disabled (default = enabled), because the MPLS VPN backbone has to propagate customer subnets unchanged to facilitate transparent end-to-end routing between customer sites.
- Redistribution of internal BGP routes into IGP is enabled (default = disabled).

Note	Common parameters that are defined in router configuration mode are inherited by all address families that are defined for this routing process and can be overridden for each individual address family.

Configuring Routing Protocols Between PE and CE Routers: RIP

Cisco.com



Configuring RIP as the PE-CE routing protocol is even easier than configuring BGP. Start the configuration of individual routing context with the **address-family ipv4 vrf vrf-name** command in router configuration mode. You can enter all standard RIP parameters in the per-VRF-table routing context. Global RIP parameters that are entered in the scope of RIP router configuration are inherited by each routing context and can be overwritten if needed in each routing context.

Note Only RIP version 2 (RIPv2) is supported as the PE-CE routing protocol. It is good configuration practice to configure RIPv2 as a global RIP parameter using the **version 2** command in router configuration mode.

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. It is not copied back into the IGP metric. The IGP metric has to be specified in the **redistribute** command or by using the **default-metric** command in router configuration mode.


The MPLS VPN extension to the **redistribute** command (**metric transparent** option) allows the MED attribute to be inserted as the IGP metric of a route that has been redistributed from BGP back into RIP. This extension gives transparent end-to-end (from the customer perspective) RIP routing:

- By default, the RIP hop count is inserted into the BGP attribute MED when the RIP route is redistributed into BGP by the ingress PE router.
- You can configure the value of the MED attribute (the original RIP hop count) to be copied into the RIP hop count when the BGP route is redistributed back into RIP. This action causes the whole MPLS VPN backbone to appear as a single hop to the CE routers.

Note You should *not* change the MED value within BGP if you use the **redistribute metric transparent** command.

Configuring Routing Protocols Between PE and CE Routers: EIGRP

Cisco.com



```
router eigrp 1
 address-family ipv4 vrf Customer_ABC
  autonomous-system 101
  network 172.16.0.0 255.255.0.0
  redistribute BGP 12703 metric 10000 100 255 1 1500
!
router bgp 12703
 address-family ipv4 vrf Customer_ABC
  redistribute eigrp 101
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-8

MPLS VPN support for EIGRP between PE and CE routers provides EIGRP with the capability to redistribute routes through a BGP VPN cloud. This feature is configured only on PE routers and requires no upgrade or configuration changes to customer equipment. This feature also introduces EIGRP support for MPLS and BGP extended community attributes.

Note Because of current limitations with route redistribution, backdoor links are not supported. If a backdoor link is implemented, it may become active and override selection of the VPN links.

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within standard BGP implementation, the MED attribute is used only as a route selection criterion. It is not copied back into the IGP metric. The metric must be configured for routes from external EIGRP autonomous systems and non-EIGRP networks before these routes can be redistributed into an EIGRP CE router. The metric can be configured in the redistribute statement using the **redistribute (IP)** command or configured with the **default-metric (EIGRP)** command.

Configuring Routing Protocols Between PE and CE Routers: OSPF

Cisco.com

Follow these steps to configure OSPF as the PE-CE routing protocol:

- Configure per-VRF-table copy of OSPF
- Configure redistribution of MP-BGP into OSPF
- Configure redistribution of OSPF into MP-BGP

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1.9

To configure OSPF as a PE-CE routing protocol, you need to start a separate OSPF process for each VRF table in which you want to run OSPF. The per-VRF-table OSPF process is configured in the same way as a standard OSPF process. You can use all the OSPF features that are available in Cisco IOS software.

You need to redistribute OSPF routes into BGP and redistribute BGP routes into OSPF if necessary. Alternatively, you can originate a default route into a per-VRF-table OSPF process by using the **default-information originate always** command in router configuration mode.

MP-BGP propagates more than just OSPF cost across the MPLS VPN backbone. The propagation of additional OSPF attributes into MP-BGP is automatic and requires no extra configuration.

Configuring Routing Protocols Between PE and CE Routers: OSPF (Cont.)

Cisco.com

```
router(config)#
```

```
router ospf process-id vrf name  
... Standard OSPF parameters ...
```

- This command starts the per-VRF-table OSPF routing process.
- The total number of routing processes per router is limited to 32.

```
router(config-router)#
```

```
redistribute bgp as-number subnets
```

- This command redistributes MP-BGP routes into OSPF. The subnets keyword is mandatory for proper operation.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-10

OSPF is the only PE-CE routing protocol that is not fully VPN-aware. A separate OSPF process is run for every VRF table.

router ospf

To configure an OSPF routing process within a VRF table, use the **router ospf** command in global configuration mode. To terminate an OSPF routing process, use the **no** form of this command:

router ospf *process-id* **vrf** *vrf-name*

Syntax Description

Parameter	Description
<i>process-id</i>	Internally used identification parameter for an OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process.
<i>vrf-name</i>	The name of the VRF table where the OSPF process will reside.

Defaults

No OSPF routing process is defined.

Configuring Routing Protocols Between PE and CE Routers: OSPF (Cont.)

Cisco.com

```
router(config)#
```

```
router bgp as-number  
  address-family ipv4 vrf vrf-name  
    redistribute ospf process-id [match [internal]  
[external-1] [external-2]]
```

- **OSPF-to-BGP route redistribution is configured with the redistribute command under the proper address family.**
- **Without the OSPF match keyword specified, only internal OSPF routes are redistributed into OSPF.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—1-11

Use the standard BGP redistribution commands.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- A unique RD must be assigned to every VRF table created in a PE router.
- The MPLS VPN architecture uses the BGP routing protocol in two ways:
 - VPNv4 routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.
 - BGP can be used as the PE-CE routing protocol to exchange VPN routes between the PE routers and the customer edge (CE) routers.
- PE-CE routing protocols need to be configured for individual VRF tables.
- Per-VRF-table routing protocols can be configured in two ways:
 - As individual address families belonging to the same routing process
 - As separate routing processes

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0–1-12

References

For additional information, refer to these resources:

- OSPF Support for Multi-VRF on CE Routers
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1838/products_feature_guide_09186a00800ad917.html
- Overview: VPN Service Model
http://www.cisco.com/en/US/partner/products/sw/netmgts/ps2327/products_programming_usage_guide_chapter09186a00800eab87.html
- MPLS VPN Support for EIGRP Between Provider Edge (PE) and Customer Edge (CE)
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1839/products_feature_guide_09186a0080154db3.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which command do you use to create a VRF table named “VPN_MPLS”?
- A) **ip vrf VPN_MPLS**
 - B) **ip rt vrf VPN_MPLS**
 - C) **ip rd vrf VPN_MPLS**
 - D) **ip vrf forwarding VPN_MPLS**
- Q2) Which of the following commands is used to configure the correct routing context for BGP?
- A) **address-family vpnv4**
 - B) **address-family ipv4**
 - C) **address-family ipv4 vrf vrf-name**
 - D) all of the above
 - E) none of the above
- Q3) When OSPF is the IGP in an MPLS implementation, which mechanism preserves the continuity of OSPF routing?
- A) OSPF superbackbone
 - B) BGP superbackbone
 - C) the ability to use traditional OSPF-to-BGP redistribution
 - D) none of the above

Lesson Answer Key

Q1) A

Relates to: Configuring VRF Tables

Q2) D

Relates to: Configuring an MP-BGP Session Between PE Routers

Q3) A

Relates to: Configuring Routing Protocols Between PE and CE Routers

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers are found in the Module Assessment Answer Key.

Quiz: MPLS VPN Review

Complete the quiz to assess what you have learned in this module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Identify and describe the elements of an MPLS implementation
- Describe how packets traverse an MPLS network
- Identify the commands that are necessary to configure MPLS in a frame-mode network
- Identify the components that are necessary to implement VPNs
- Identify the commands that are necessary to implement a simple MPLS VPN implementation

Instructions

Complete these steps:

- Step 1** Answer all questions in this quiz by selecting the best answer(s) to each question.
- Step 2** Verify your results against the answer key at the end of this quiz.
- Step 3** Review the related lesson for each question that you answered incorrectly.

Quiz

Answer these questions:

- Q1) Which of the following best describes the term “a simple label-based forwarding engine?”
- A) control plane
 - B) ground plane
 - C) data plane
 - D) routing plane
- Q2) What does the term “pop” mean when you are describing penultimate hop popping?
- A) swap the top label with a new label, which is contained in the LIB
 - B) swap the top label with a new label, which is contained in the LFIB
 - C) remove the top label instead of swapping it with the next-hop label
 - D) remove the bottom label instead of swapping it with the next-hop label
- Q3) The MPLS MTU is increased to _____ to support 1500-byte IP packets and MPLS stacks up to three levels deep.
- Q4) Why do MPLS VPNs implement route targets?
- A) to identify different customer VPNs
 - B) to allow a site to participate in more than one VPN
 - C) to convert a customer address to an MP-BGP address
 - D) to convert a non-unique IP address into a unique VPNv4 address

- Q5) What happens to the interface of an existing IP address when you associate the interface with a VRF table?
- A) It remains unchanged.
 - B) It is removed from the interface.
 - C) It is changed to the loopback 0 address.
 - D) It is moved to under the VRF table configuration.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) C

Relates to: Basic MPLS Concepts

Q2) C

Relates to: Label Distribution in Frame-Mode MPLS

Q3) 1512

Relates to: Configuring MPLS

Q4) B

Relates to: MPLS VPN Routing Model and Packet Forwarding

Q5) B

Relates to: Configuring Simple MPLS VPNs

Carrier Supporting Carrier

Overview

It is becoming common for larger service providers to act as the backbone for smaller service providers. Multiprotocol Label Switching (MPLS) supports this environment by using the features and functions of Carrier Supporting Carrier (CSC). MPLS managed services offer a full set of Virtual Private Network (VPN)-aware services that a service provider can offer its customers.

Objectives

Upon completing this module, you will be able to use the Cisco IOS commands that are required to successfully configure, monitor, and troubleshoot service provider support in an MPLS CSC solution. This includes being able to do the following:

- Identify the characteristics of a CSC implementation
- Identify the tasks and commands that are required to implement CSC in supported solutions

Outline

The module contains these lessons:

- Carrier's Carrier Overview
- Implementing the CSC Models

Carrier's Carrier Overview

Overview

It is becoming common for a larger service provider to act as the backbone for a smaller provider. This environment has classically been addressed with overlay VPNs. However, this solution has some serious scalability issues. Earlier attempts to address this environment with MPLS created several issues related to crossing service boundaries. The Carrier Supporting Carrier (CSC) feature was introduced to address these issues.

Relevance

In order to successfully implement the CSC feature of MPLS, it is necessary to understand why a carrier would need to provide backbone support for another carrier and how this support is implemented in an MPLS environment.

Objectives

This lesson describes the need for Carrier's Carrier support in an MPLS environment and how this support operates. Upon completing this lesson, you will be able to do the following:

- Describe the need for CSC
- Describe the benefits of CSC
- At an overview level, describe how CSC operates

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS VPN technology and configuration knowledge
- Advanced Border Gateway Protocol (BGP) knowledge


Outline

This lesson includes these topics:

- Overview
- Carrier Supporting Carrier
- Benefits of CSC
- How CSC Works
- Lesson Summary
- Lesson Review

Carrier Supporting Carrier


This topic describes the need for CSC support.



A carrier network carries traffic between customer sites. The figure illustrates the traditional scenario in which the carrier backbone interconnects two end-user sites.

Carrier Supporting Carrier (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-4

The figure illustrates how the carrier network no longer interconnects end-user sites. Instead, it now interconnects two different point of presence (POP) sites of another carrier, "A." The end-user sites are connected to carrier A. The backbone carrier has become a carrier's carrier, a scenario sometimes referred to as "recursive service providers."

Benefits of CSC

This topic describes the benefits of using CSC support.

Benefits of CSC

Cisco.com

Carrier Supporting Carrier (CSC) has many benefits:

- **Multiple customer carriers can be connected to a single CsC backbone.**
- **Both VPN and Internet services can be provided.**
- **Customer carriers do not have to operate their own long-distance network.**
- **Different addressing schemes can be used by different carriers.**
- **Any link type supported by MPLS can be used.**


© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—2-5

CSC, sometimes called Carrier's Carrier, has many benefits:

- A single CSC backbone can be used to connect many POP sites.
- An MPLS VPN can be used to separate traffic from the POP sites of different carriers.
- A single CSC backbone can provide services to both VPN service providers and to the Internet service provider (ISP).
- The customer service providers do not have to operate their own long-distance network. They purchase that service from the CSC backbone carrier.
- Different customer carriers can use different addressing schemes. The customer carriers will be in separate MPLS VPNs inside the CSC backbone.
- Any link type that is supported by MPLS can be used inside the CSC backbone and as access links between the CSC backbone and customer carriers.

Benefits of CSC: No End-User Routes in CSC Backbone

Cisco.com




In this scenario, the end-user routes are never handled by the CSC backbone. Carrier A is one autonomous system (AS) that is divided over several POPs. The POPs are interconnected by the CSC backbone but only the internal routes, which can possibly be used as BGP next-hop addresses.

Carrier A is using Internal Border Gateway Protocol (IBGP) inside its own AS. End-user routes, like 10.1.1.0/24 in the figure, are carried inside these IBGP sessions. The CSC backbone never sees them.

The CSC backbone must, however, carry the routes that are used by Carrier A as next-hop addresses. In the network in the figure, the address 172.16.1.1 is the next hop. This is the loopback address of A1-2.

How CSC Works

This topic provides an overview of how the CSC feature operates.



The CSC architecture relies on the presence of an MPLS VPN. The CSC backbone is providing an MPLS VPN service to which the customer carriers are connected as VPN sites. MPLS is used between the CSC backbone provider edge (PE) routers and the VPN sites of the customer carriers.

Virtual routing and forwarding (VRF) tables are enabled on the CSC PE routers. The label exchange between PE-X and PE-Y establishes a label switched path (LSP) from CE-A1 via the CSC backbone to CE-A2. Another LSP is also established in the other direction.

Carrier A can now tunnel packets between POP 1 and POP 2 using the LSPs. The CSC backbone does not have to know about end-user sites and their IP addresses.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **The CSC feature is useful when a carrier wants to interconnect several remote POPs without operating its own WAN.**
- **A customer carrier purchases the interconnection service from another carrier.**
- **This system implies a hierarchy of carriers with the CSC backbone on the top.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—2-8

References

For additional information, refer to this resource:

- Access the Cisco website at <http://www.cisco.com> for additional information about the MPLS VPN CSC feature.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) What are the benefits of CSC?

Q2) Do the CSC backbone PE routers need to know the end-user routes? Why or why not?

Q3) How are packets propagated across the CSC backbone?

Q4) How are hierarchical VPNs supported?

Lesson Answer Key

Q1) The benefits of CSC are as follows:

- Multiple customer carriers can be connected to a single CSC backbone.
- Both VPN and Internet services can be provided.
- Customer carriers do not have to operate their own long-distance network.
- Different addressing schemes can be used by different carriers.
- Any link type that is supported by MPLS can be used.

Relates to: Benefits of CSC

Q2) No. The CSC backbone has to carry routing information only about the next-hop address of the customer carrier. The end-user routes are never seen by the CSC routers.

Relates to: Benefits of CSC

Q3) Packets are propagated from one customer POP to another POP along an LSP from CE to CE routers.

Relates to: How CSC Works

Q4) Hierarchical VPNs are supported using a layered label stack.

Relates to: How CSC Works

Implementing the CSC Models

Overview

CSC is implemented using the standard MPLS features. However, the layering of multiple MPLS networks adds a level of complexity. This lesson provides a systematic approach to implementing the different CSC models.

Relevance

In order to successfully implement the CSC feature of MPLS, it is necessary to understand the CSC models and how their support is implemented in an MPLS environment.

Objectives

This lesson describes the different CSC models and how they are implemented. Upon completing this lesson, you will be able to do the following:

- Identify the CSC models
- Configure CSC to support customer carriers that are running native IP
- Configure CSC to support customer carriers that are running MPLS and LDP label distribution
- Configure CSC to support customer carriers that are running MPLS and MP-BGP label distribution
- Configure CSC to support customer carriers that are running MPLS VPNs

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS VPN technology and configuration knowledge
- Advanced BGP knowledge

Outline

This lesson includes these topics:

- Overview
- CSC Models
- Customer Carrier: Native IP
- Customer Carrier: MPLS and LDP
- Customer Carrier: MPLS and EBGp
- Customer Carrier: MPLS VPNs
- Lesson Summary
- Lesson Review

CSC Models

This topic identifies the CSC models.

CSC Models

Cisco.com

- **CSC is implemented using MPLS VPNs.**
- **A customer carrier can be supported in several models:**
 - **Native IP**
 - **MPLS**
 - **CE-PE using BGP**
 - **CE-PE using LDP**
 - **MPLS VPN**


© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—2-3

To support CSC, the CSC backbone must support MPLS VPNs. The customer carrier may run in several modes:

- Native IP (classic VPN support)
- MPLS (label exchange between the customer edge [CE] router of the customer carrier and the CSC provider edge [PE] router):
 - Customer edge-provider edge (CE-PE) label exchange using Border Gateway Protocol (BGP)
 - CE-PE label exchange using Label Distribution Protocol (LDP)
- MPLS VPN (MPLS VPN label exchange)

Customer Carrier: Native IP

This topic describes the two methods of implementing CSC when a customer carrier is running a native IP core with static routing and LDP routing.



The figure shows a customer carrier that is running a classic IP network with static routing between itself and the CSC carrier. The customer carrier is running a native Interior Gateway Protocol (IGP) to propagate its internal routes and BGP to propagate the routes of its customers.

The core routing information that the customer carrier site is advertising to the CSC backbone should include the site internal routes only. No end-user routes should be advertised to the CSC. The customer carrier routing information is propagated to the CSC PE router using any supported IGP CE-PE routing protocol.

Note BGP is not used for routing between the customer carrier and the CSC carrier because it would result in the propagation of the end-customer routing into the VRF table.

The routes are propagated into the VRF table. From the VRF table the routes are assigned a route distinguisher (RD) and are propagated using Multiprotocol Border Gateway Protocol (MP-BGP). At the egress PE router, the MP-BGP routes are imported into the VRF table, and then propagated to the POP 2 of the customer carrier.


This process allows an IBGP session to be established between the routers in the two POPs. To reduce the number of sessions, the most common implementation would be to enable the route reflector function on A1-1 and A2-1.

However, there are still issues with the ability of the CSC PE routers to route customer packets. Because the packets that they receive are in a native mode (that is, they contain the addressing of the end customer), the CSC PE routers must know how to forward customer packets. This goal is accomplished using static routing.

This model has scalability issues that are addressed in the LDP routing model discussed later in this lesson.

Customer Carrier: Native IP – Static Routing (Cont.)


Cisco.com



In this illustration of the static routing method, the end customer forwards an IP packet. This packet is forwarded over the customer carrier network in the native IP mode. When the packet enters the backbone, two labels are imposed; one label identifies the associated VPN, and the other identifies the path to the egress PE router.

Customer Carrier: Native IP – Static Routing (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—2-6

With the static routing method, the configuration on the CSC PE router is very similar to that of a classic VPN PE router. Note the addition of the static routes that point to the prefix of each customer site. This situation results in scaling problems. For this reason, the static routing method is not recommended.

Note For clarity, the MPLS and BGP configurations of the backbone have been omitted.

Customer Carrier: Native IP – Static Routing (Cont.)

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-7

The configuration of the customer carrier Customer Carrier's CE router is very similar to that of a classic VPNCE router. Note that in addition to the internal IBGP neighbor relationships, the routers in one POP also have an IBGP relationship with the routers in the other POP. This relationship can be created with multiple neighbor statements or by the implementation of a route reflector in each POP.

Customer Carrier: Native IP – LDP Routing

Cisco.com



- Customer carrier uses:
 - IGP for internal routes
 - EBGP for propagation between it and the end customer
 - IBGP to the end-customer routes across its network
- Internal routes of customer carrier are propagated to the VRF table using an IGP.
- CSC backbone uses MP-BGP to propagate the routes across the backbone.
- IBGP sessions are established between the customer carrier routers.
- An end-to-end LSP is used to pass the end customer routing.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-8


In this figure, which demonstrates the LDP routing method of implementing CSC in the native IP core, the customer carrier is still running a classic IP network with BGP to propagate its customer routes. The core routing information remains the same.

However, the CSC PE routers no longer need to know how to forward customer packets. An end-to-end LSP has been established between the CE routers. This setup allows them to forward based upon the BGP next hop. The core then packet-switches the packet.

In this situation, the end customer sends an IP packet. This packet is forwarded over the customer carrier network in native IP mode. When the packet leaves the customer carrier router (A1-1), a label is imposed that identifies path over the CSC backbone to the egress CE router (A2-1).

Customer Carrier: Native IP – LDP Routing (Cont.)


Cisco.com



In this illustration of the LDP routing method, the configuration on the backbone PE router is very similar to that of a classic VPN. Note that static routes pointing to the prefix of each customer site have been replaced with LDP sessions.

Customer Carrier: Native IP – LDP Routing (Cont.)

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-10

In the native IP with LDP routing model, the configuration on the CSC PE router is very similar to that of a classic VPN. Note that there are static routes pointing to the prefix of each customer site in the native IP with static routes model have been replaced with LDP sessions.

Customer Carrier: MPLS and LDP

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—2-12

The configuration of the customer carrier CE router is very similar to that of a classic VPN. Note that in addition to the internal IBGP neighbor relationships, the routers in one POP also have an IBGP relationship with the routers in the other POP. This relationship can be created with multiple neighbor statements or by implementation of a route reflector in each POP.

Also note that an LDP session has been established with the CSC PE router.

Customer Carrier: MPLS and LDP

This topic describes the implementation of CSC when a customer carrier is running MPLS and LDP is used to distribute labels between the backbone and the customer carrier.



In this figure, the customer carrier is running a basic MPLS network. An IGP is used to propagate the internal routes, and LDP is used to distribute labels and build LSPs within the customer carrier network. BGP is used by the customer carrier to propagate the end-customer routes.


To propagate the customer carrier routes across the backbone, an IGP is enabled between the customer carrier CE router and the backbone. These routes are propagated into a VRF table, and MP-BGP is used to propagate the routes across the backbone.

LDP is enabled within the VRF table that is supporting the customer carrier. This situation allows the customer carrier and the backbone to distribute labels. The result is the building of LSPs between the two POPs.

The resulting end-to-end LSPs, one in each direction, allow the IBGP session to be established between the routers (A1-2 and A2-2) that are supporting the end customer. The CE routers of the customer carrier (A1-1 and A2-1) no longer need to be a BGP next hop.

Customer Carrier: MPLS and LDP (Cont.)

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-14

In this illustration of the MPLS and LDP customer carrier implementation, the end customer forwards an IP packet. This packet is forwarded over the customer carrier network using MPLS labels. When the packet enters the backbone, two labels are imposed; one label identifies the associated VPN, and the other identifies the path to the egress PE router. At the egress PE router, a label that is provided by the customer carrier is imposed.

Customer Carrier: MPLS and LDP (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—2-15

This configuration differs from the previous configurations seen in the native IP methods in that a loopback address from within the VRF table has been added. This address will be used for the LDP session within the VRF table.

In addition, MPLS and LDP have been enabled on the interface that is supporting the customer carrier.

Customer Carrier: MPLS and LDP (Cont.)


Cisco.com



On the CE routers of the customer carrier (A1-1 and A2-1), MPLS and LDP are enabled on the interface that is connected to the backbone. Note that BGP is no longer needed on these routers because they are performing label swapping only. Their function is now similar to that of the classic PE router.

Customer Carrier: MPLS and EBGP

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-17

The IBGP session is now run between the end routers (A1-2 and A2-2) in the customer carrier network.

Customer Carrier: MPLS and EBG

This topic describes the implementation of CSC when a customer carrier is running MPLS and External Border Gateway Protocol (EBGP) is used to distribute labels between the backbone and the customer carrier.



In this figure, the customer carrier is running a basic MPLS network. An IGP is used to propagate the internal routes, and LDP is used to distribute labels and build LSPs within the customer carrier network. BGP is used by the customer carrier to propagate the end-customer routes.

To propagate the customer carrier routes across the backbone, an IGP is enabled between the CE router of the customer carrier and the backbone. These routes are propagated into a VRF table and MP-BGP is used to propagate the routes across the backbone.

EBGP is enabled within the VRF table that is supporting the customer carrier. This situation allows the customer carrier and the backbone to distribute labels. The result is the building of LSPs between the two POPs.

The resulting end-to-end LSPs allow the IBGP session to be established between the routers (A1-2 and A2-2) that are supporting the end customer. The CE routers of the customer carrier (A1-1 and A2-1) no longer need to be a BGP next hop.

Customer Carrier: MPLS and EBGp (Cont.)

Cisco.com

```
router(config-router)#
```

```
neighbor ip-address send-label
```

- **This command enables a router to use BGP to distribute MPLS labels along with the IPv4 routes to a peer router. You must issue this command on both the local router and the neighboring router.**


© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—2-19

In this illustration of the MPLS and EBGp customer carrier implementation, the end customer forwards an IP packet. This packet is forwarded over the customer carrier network using MPLS labels. When the packet enters the backbone, two labels are imposed; one identifies the associated VPN, and the other identifies the path to the egress PE router. At the egress PE router, a label that is provided by the customer carrier is imposed.

Customer Carrier: MPLS and EBGp (Cont.)

Cisco.com



neighbor send-label

To enable a BGP router to send MPLS labels with BGP routes to a neighboring BGP router, use the **neighbor send-label** command in router configuration mode. To disable the BGP router from sending MPLS labels with BGP routes, use the **no** form of this command.

- **neighbor {ip-address} send-label**
- **no neighbor {ip-address} send-label**

Syntax Description

Parameter	Description
<i>ip-address</i>	IP address of the neighboring router

Defaults

By default, BGP routers distribute only BGP routes.

Command Modes

Router configuration


Usage Guidelines

This command enables a router to use BGP to distribute MPLS labels along with the IP version 4 (IPv4) routes to a peer router. You must issue this command on both the local router and the neighboring router. This command has the following restrictions:

- If a BGP session is running when you issue the **neighbor send-label** command, the command does not take effect until the BGP session is restarted.
- You can use this command only with IPv4 addresses.

Customer Carrier: MPLS and EBG (Cont.)

Cisco.com




This configuration differs from the configuration in the MPLS and LDP method in that although MPLS has been enabled on the backbone, LDP has been disabled on the interface supporting the customer carrier.

In addition, a BGP address family has been added to support the neighbor relationship with the customer carrier. Label propagation has been enabled in this address family.

Customer Carrier: MPLS and EBGP (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—2-21

On the CE routers of the customer carrier (A1-1 and A2-1), MPLS and MP-BGP are enabled on the interface that is connected to the backbone. Note that BGP is now needed on these routers because it is used for label propagation. However, the routers still forward packets using label swapping, so their function is still similar to that of the classic PE router.

Note The configuration of A1-2 and A2-2 is the same as in the MPLS and LDP example.

Customer Carrier: MPLS VPNs


This topic describes the implementation of CSC when a customer carrier is running MPLS VPNs.



This CSC implementation builds on the MPLS and LDP model. The resulting end-to-end LSP enables the customer carrier to establish a peer relationship between its PE routers (A1-2 and A2-2) that are supporting the end customer. It then enables the customer carrier to use MPLS VPNs to support its end customers over an end-to-end VPN.

Customer Carrier: MPLS VPNs (Cont.)

Cisco.com



In this illustration of the MPLS VPNs customer carrier implementation, the end customer forwards an IP packet. When the packet enters the customer carrier network, two labels are imposed; one identifies the associated VPN, and the other identifies the path to the egress PE router (A2-2).


When the packet enters the backbone, one label is imposed that identifies the backbone VPN and the LSP label is swapped.

At the backbone egress PE router, a label that is provided by the customer carrier is imposed over the customer carrier VPN label.

At the customer carrier penultimate hop popping (PHP) router (A2-1), the VPN label is popped and the packet is forwarded via the native label forwarding information base (LFIB).

Customer Carrier: MPLS VPNs (Cont.)

Cisco.com



Note that the configuration here is a standard VPN configuration. The exception is that the **neighbor 37.0.0.2 as-override** command is needed because both end customer sites are using the same AS number.

Note The configuration of A1-1, A2-1, PE-X, and PE-Y is the same as in the MPLS and LDP example.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- The CSC feature is useful when a carrier wants to interconnect several remote POPs without operating its own WAN.
- This feature implies a hierarchy of carriers with the CSC backbone on the top.
- The customer carrier advertises its local routes on each POP to the CSC backbone.
- The CSC backbone operates as an MPLS VPN service and provides connectivity between the POPs.
- End-user routes are:
 - Never advertised to the CSC backbone
 - Are carried by IBGP sessions between the POP sites
- The CSC has to carry only the next-hop routes. The actual end-user routes are never touched by the CSC.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—2-25

References

For additional information, refer to this resource:

- Access the Cisco website at <http://www.cisco.com> for additional information about the MPLS VPN CSC feature.
- MPLS VPN Carrier Supporting Carrier
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1612/products_feature_guide09186a00800e9577.html#xtocid20

Next Steps

For the associated lab exercises, refer to the following sections of the course Lab Guide:

- Lab Exercise 2-1: Implementing CSC in a Native IP Environment
- Lab Exercise 2-2: Implementing CSC Using LDP Label Distribution
- Lab Exercise 2-3: Implementing CSC Using MP-BGP Label Distribution
- Lab Exercise 2-4: Implementing CSC Using MPLS VPNs

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Describe how routing updates are passed end-to-end in the customer carrier: native IP model.

- Q2) Describe how routing updates are passed end-to-end in the customer carrier: MPLS and LDP model.

Q3) Describe how routing updates are passed end-to-end in the MPLS and EBGp model.

Q4) Describe how routing updates are passed end-to-end in the customer carrier: MPLS VPNs model.

Lesson Answer Key

- Q1) The routes are propagated into the VRF table. From the VRF table the routes are assigned an RD and are propagated using MP-BGP. At the egress PE router, the MP-BGP routes are imported into the VRF table and then propagated to the customer carrier POP.
- Relates to:** Customer Carrier: Native IP
- Q2) To propagate the customer carrier routes across the backbone, an IGP is enabled between the CE router of the customer carrier and the backbone. The routes are propagated into a VRF table, and MP-BGP is used to further propagate them across the backbone. LDP is enabled within the VRF table that is supporting the customer carrier. This setup allows the customer carrier and the backbone to distribute labels. The result is the building of LSPs between the two POPs. The resulting end-to-end LSPs, one in each direction, allow the IBGP session to be established between the routers that are supporting the end customer.
- Relates to:** Customer Carrier: MPLS and LDP
- Q3) To propagate the customer carrier routes across the backbone, an IGP is enabled between the CE router of the customer carrier and the backbone. The routes are propagated into a VRF table, and MP-BGP is used to further propagate them across the backbone. MP-BGP is enabled within the VRF table that is supporting the customer carrier. This setup allows the customer carrier and the backbone to distribute labels. The result is the building of LSPs between the two POPs. The resulting end-to-end LSPs allow the IBGP session to be established between the routers that are supporting the end customer.
- Relates to:** Customer Carrier: MPLS and EBGp
- Q4) This CSC implementation builds on the MPLS and LDP model. The resulting end-to-end LSP enables the customer carrier to establish a peer relationship between its PE routers that are supporting the end customer. The customer carrier then uses MPLS VPNs to support its end customers over an end-to-end VPN.
- Relates to:** Customer Carrier: MPLS VPNs

MPLS Traffic Engineering Technology

Overview

This module on MPLS Traffic Engineering (MPLS-TE) technology discusses the requirement for traffic engineering (TE) in modern networks that must attain optimal resource utilization. The traffic-engineered tunnels provide a means of mapping traffic streams onto available networking resources in a way that prevents the overuse of subsets of networking resources while other subsets are underused. All the concepts and mechanics that support TE are presented, including tunnel path discovery with link-state protocols and tunnel path signaling with Resource Reservation Protocol (RSVP). Some of the advanced features of TE, such as autobandwidth and guaranteed bandwidth, are introduced as well.

Objectives

Upon completing this module, you will be able to describe the tasks and commands that are necessary to implement MPLS-TE. This includes being able to do the following:

- Identify the basic TE concepts
- Identify the basic MPLS-TE concepts
- Describe the function of constraint-based path computation
- Describe the operations of LSP path setup and maintenance
- Identify how to assign traffic to traffic trunks

Outline

The module contains these lessons:

- Traffic Engineering Concepts
- MPLS Traffic Engineering Components
- Constraint-Based Path Computation
- Path Setup and Maintenance
- Assigning Traffic to Traffic Tunnels

Traffic Engineering Concepts

Overview

This lesson describes the concepts that allow service providers to map traffic through specific routes to optimize network resources, especially the bandwidth. Traffic engineering (TE) enables backbone networks to be engineered to deliver the total subscribed capacity to service provider customers more efficiently.

Relevance

This lesson is mandatory for learners who are planning to improve the usage of their network resources with MPLS Traffic Engineering (MPLS-TE).

Objectives

This lesson describes the basic TE concepts. Upon completing this lesson, you will be able to do the following:

- Identify the concepts behind TE
- Identify the major business drivers for implementing TE
- Identify congestion avoidance and how TE can reduce some congestion avoidance issues
- Identify how TE is implemented using a Layer 2 overlay model
- Identify how TE is implemented using a Layer 3 model
- Identify how TE is implemented using the MPLS-TE model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Carrier Supporting Carrier” module of this course

Outline

This lesson includes these topics:

- Overview
- What Is Traffic Engineering?
- Business Drivers for Traffic Engineering
- Congestion Avoidance and Traffic Engineering
- Traffic Engineering with a Layer 2 Overlay Model
- Traffic Engineering with a Layer 3 Model
- Traffic Engineering with the MPLS-TE Model
- Lesson Summary
- Lesson Review

What Is Traffic Engineering?

This topic identifies the underlying concepts of TE.

What Is Traffic Engineering?

Cisco.com

- **Term in common use in telephone voice network world**
- **Measures, models, and controls traffic to achieve various goals**
- **Provides an integrated approach to engineering traffic at Layer 3 in the Open Systems Interconnection reference model**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-3

The term “traffic engineering” (TE) is widely used in the telephone voice world. TE means that the traffic is measured and analyzed. A statistical model is then applied to the traffic pattern to make a prognosis and estimations.

If the anticipated traffic pattern does not match well with the network resources, the network administrator will remodel the traffic pattern. Such decisions can be made to achieve a more optimum use of the resources or to reduce costs by selecting a cheaper transit carrier.

In the data communications world, traffic engineering provides an integrated approach to engineering traffic at Layer 3 in the Open Systems Interconnection (OSI) model. The integrated approach means that routers are configured to divert traffic from destination-based forwarding and move the traffic load from congested parts of the network to noncongested parts.

Traditionally, this diversion has been done using overlay networks where routers use carefully engineered ATM PVCs or Frame Relay PVCs to distribute the traffic load on Layer 2.

What Is Traffic Engineering? Traffic Engineering Motivations

Cisco.com

- **Reduce the overall cost of operations by more efficient use of bandwidth resources**
- **Prevent a situation where some parts of a service provider network are overutilized (congested), while other parts remain underutilized**
- **Implement traffic protection against failures**
- **Enhance SLA in combination with QoS**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-4

Cost reduction is the main motivation for TE.

A cost savings, which results from a more efficient use of resources, will help to reduce the overall cost of operations.

Additionally, more efficient use of bandwidth resources means that a provider can avoid a situation where some parts of its network are congested, while other parts are underutilized.

Because TE can be used to control traffic flows, it can also be used to provide protection against link or node failures by providing backup tunnels.

Finally, when combined with QoS functionality, TE can provide enhanced Service Level Agreements (SLAs).

Business Drivers for Traffic Engineering

This topic identifies some business drivers that may lead to the implementation of TE.

Business Drivers for Traffic Engineering

Cisco.com

- **Routers always forward traffic along the least-cost route as discovered by intradomain routing protocol (IGP).**
- **Network bandwidth may not be efficiently utilized:**
 - **The least-cost route may not be the only possible route.**
 - **The least-cost route may not have enough resources to carry all the traffic.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-5

In a Layer 3 routing network, packets are forwarded hop-by-hop. In each hop, the destination address of the packet is used to make a routing table lookup. The routing tables are created by an Interior Gateway Protocol (IGP), which finds the least-cost route, according to its metric, to each destination in the network.

In many networks, this method works well. But in some networks the destination-based forwarding results in the overutilization of some links, while others are underutilized. This imbalance will be the case when there are several possible routes to reach a certain destination. The IGP selects one of them as the best and uses only that route. In the extreme case, the best path may have to carry so large a volume of traffic that packets are dropped while the next best path is almost idle.

One solution to the problem would be to adjust the link bandwidths to more appropriate values. The network administrator could reduce the bandwidth on the underutilized link and increase the bandwidth on the overutilized one. However, making this adjustment is not always possible.

The alternate path is a backup path. In the case of a primary link failure, the backup must be able to forward at least the major part of the traffic volume that is normally forwarded by the primary path. Therefore, it may not be possible to reduce the bandwidth on the backup path. And without a cost savings, the budget may not allow an increase to the primary link bandwidth.

In order to provide better network performance within budget, network administrators move a portion of the traffic volume from the overutilized link to the underutilized link. During normal operations, this move results in fewer packet drops and quicker throughput. In the case of a failure to any of the links, all traffic is forwarded over the remaining link, which then, of course, becomes overutilized.

Moving portions of the traffic volume cannot be achieved by traditional hop-by-hop routing using an IGP for path determination.

Business Drivers for Traffic Engineering (Cont.)

Cisco.com

- **Lack of resources results in congestion in two ways:**
 - **When network resources themselves are insufficient to accommodate offered load**
 - **When traffic streams are inefficiently mapped onto available resources**
- **Some resources are overutilized while others remain underutilized.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-6

Network congestion, caused by too much traffic and not enough network resources, cannot be solved by moving portions of the traffic between different links. Moving the traffic will help only in the case where some resources are overutilized and others are underutilized. The traffic streams in normal Layer 3 routing are inefficiently mapped onto the available resources.

Good mapping of the traffic streams onto the resources creates a better use of the invested money.

Cost savings that result in a more efficient use of bandwidth resources help to reduce the overall cost of operations. These reductions, in turn, help service providers gain an advantage over their competitors. This advantage becomes more and more important as the service provider market gets more and more competitive.

A more efficient use of bandwidth resources means that a provider could avoid a situation where some parts of its network are congested, while other parts are underutilized.

Congestion Avoidance and Traffic Engineering

This topic discusses congestion avoidance and how TE can reduce some of those issues.

Congestion Avoidance and Traffic Engineering

Cisco.com

- **Network congestion can be addressed by either:**
 - **Expansion of capacity or classical congestion control techniques (queuing, rate limiting, etc.)**
 - **Traffic Engineering (TE), if the problems result from inefficient resource allocation**
- **The focus of TE is not on congestion created as a result of a short-term burst, but on congestion problems that are prolonged.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-7


TE does not solve temporary network congestion that is caused by traffic bursts. This type of problem is better handled by an expansion of capacity or by classical techniques such as various queuing algorithms, rate limiting, and intelligent packet dropping.

TE is used when the problems result from inefficient mapping of traffic streams onto the network resources.

In such networks, one part of the network suffers from congestion during long periods of time, possibly continuously, while other parts of the network have spare capacity.

Traffic Engineering with a Layer 2 Overlay Model

This topic discusses the use of a Layer 2 overlay model with TE.



In the overlay model, the routers (Layer 3 devices) are not aware of the physical structure and the bandwidth that is available on the links. The IGP views the PVCs (permanent virtual circuits) or SVCs (switched virtual circuits) as point-to-point links and makes its forwarding decisions accordingly.

Instead, all engineering is done at Layer 2. PVCs are carefully engineered across the network, normally using an off-line management system. SVCs are automatically established using signaling, and their way across the Layer 2 network is controlled by an integrated path determination such as the Private Network-to-Network Interface (PNNI) protocol.

If the Layer 2 network provides a full mesh between all routers, the Layer 3 IGP sees all the other routers as directly connected, and, most likely, uses the direct logical link whenever it forwards a packet to another router. The full mesh gives Layer 2 full control of the traffic load distribution. Manual engineering of PVCs and the configuration of PNNI parameters are the tools that allow a very exact control of how the traffic uses the available bandwidth.

Traffic Engineering with a Layer 2 Overlay Model (Cont.)

Cisco.com

Overlay Model Characteristics

- **Permanent virtual circuits (PVCs) carry traffic across Layer 2.**
- **Switched virtual circuits (SVCs) are established via signaling:**
 - **Example: ATM SVCs:**
 - **The router signals the request to establish a switched virtual circuit to the ATM switch using the User-Network Interface (UNI) protocol.**
 - **The ATM switch opens this SVC using the Private Network-to-Network Interface (PNNI) protocol.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-3-9

In the overlay model, PVCs or SVCs carry the traffic across the network.


In the case of a Frame Relay network, a PVC setup is most often made using a management tool. This tool helps the network administrator calculate the optimum path across the Layer 2 network with respect to available bandwidth and other constraints that may be applied on individual links.

ATM uses either the same type of tools as Frame Relay for PVC establishment, or it may use the SVC approach, where routers use a signaling protocol to dynamically establish a switched virtual circuit.

When SVCs are used, the router merely asks the other router for an SVC with certain attributes by using the signaling protocol that has been specified by the ATM Forum. The Layer 2 network then opens this SVC internally using the PNNI protocol. In the headend ATM switch, PNNI uses link-state information to precalculate a designated transit list (DTL), which describes the suggested total path across the ATM network. This suggested path is then validated across the ATM network by each hop switch, which then provides the SVC.

Traffic Engineering with a Layer 2 Overlay Model: Example

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-10

Traffic engineering in Layer 2, using the overlay model, allows for detailed decisions regarding which link should be used to carry different traffic patterns.

In this example, traffic from R2 to R3 uses a PVC, which takes the shortest path using the upper transit switch.

However, traffic from R1 to R3 also uses a PVC, which does not take the shortest path. TE on Layer 2 has been applied to let the PVC use links that would otherwise have been underutilized. This approach avoids overutilization of the upper path.

Traffic Engineering with a Layer 2 Overlay Model (Cont.)

Cisco.com

Drawbacks of the Overlay Solution

- **Extra network devices**
- **More complex network management:**
 - **Two-level network without integrated network management**
 - **Additional training, technical support, field engineering**
- **IGP routing scalability issue for meshes**
- **Additional bandwidth overhead (“cell tax”)**
- **No differential service (class of service)**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—3-11

Using the overlay model has several drawbacks:

- The routers are not physically connected to other routers. The Layer 2 network introduces the need for an additional device, the ATM or Frame Relay switch.
- Two networks must be managed. The Layer 2 network requires its own management tools, which, among several other tasks, support TE as well. At the same time, the router network (Layer 3), with its IGP and tuning parameters, must be managed. Both these management tasks require trained staff for technical support and in the field.
- The Layer 3 network must be highly meshed in order to take advantage of the benefits that are provided by the Layer 2 network. The highly meshed network may cause scalability problems for the IGP because of the large number of neighbors.
- Overlay networks always require an extra layer of encapsulation. A Frame Relay header must be added to the IP packets, or, when ATM is used, the IP packet must be segmented into cells, each of which must have its own header. The extra layer of encapsulation causes bandwidth overhead.
- The Layer 2 devices do not have any Layer 3 knowledge. After the router has transmitted the IP packet across the physical link to the first switch, all IP knowledge is lost. When congestion does occur in the Layer 2 network, the switches have no ability to selectively discard IP packets or to requeue them. Thus, no IP differentiated services can be used within the Layer 2 switch network.

Traffic Engineering with a Layer 3 Model

This topic discusses the use of a Layer 3 model with TE.



If the same network topology is created using routers (Layer 3 devices), TE must be performed differently. In the example here, if no traffic engineering is applied to the network, traffic from both R8 and R1 toward R5 will use the least-cost path (the upper path). This flow may result in the overutilization of the path R2, R3, R4, R5, while the path R2, R6, R7, R4, R5 will be underutilized.

Traffic Engineering with a Layer 3 Model (Cont.)

Cisco.com

The current forwarding paradigm, centered around “destination-based,” is clearly inadequate:

- **Path computation based just on IGP metric is not enough.**
- **Support for “explicit” routing (source routing) is not available.**
- **Supported workarounds: static routes, policy routing.**
- **Provide controlled backup and recovery.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-13

The destination-based forwarding paradigm that is currently used in Layer 3 networks cannot handle the problem of overutilization of one path while an alternate path is underutilized.

The IGP uses its metric to compute a single best way to reach each destination. Alternate routes with a higher metric are not used at all.


IP source routing could be used to override the IGP-created routing table in each of the intermediate routers. However, in a service provider network, source routing is most often prohibited. The source routing would also require the host to create the IP packets to request source routing. The conclusion is that source routing is not an available tool for TE.

Static routing, which overrides the IGP, can be used to direct traffic to take a different path than that of other traffic. However, static routing does not discriminate between different traffic flows based on the source. Static routing also restricts how redundancy in the network can be used.

Policy-based routing (PBR) is able to discriminate packet flows based on the source, but it suffers from low scalability and the same static routing restrictions as to how redundancy can be used.

Traffic Engineering with the MPLS-TE Model

This topic discusses using TE with the MPLS-TE model.



MPLS-TE means that the routers use the MPLS label-switching and tag-switching model. Labels are assigned and distributed between routers using the Label Distribution Protocol (LDP) or the Tag Distribution Protocol (TDP). Packets are assigned labels by the ingress router, and the packet is then forwarded across the network using label switching, and not on IP header information. At the egress router, the label is removed and the packet is again forwarded as an IP packet.

When full label information is exchanged, any router can reach any other router within the MPLS domain using label switching. In other words, a label switched path (LSP) exists between all routers.

The existing LSPs or newly created ones between the routers are used by MPLS applications such as MPLS Virtual Private Networks (VPNs) and MPLS-TE. A stack of two labels is imposed on the IP packet by the ingress router. The topmost label value is used to let the packet traverse the desired LSP to the router at the other end. The additional label is then used by that router to indicate further actions.

With MPLS VPNs, Multiprotocol Border Gateway Protocol (MP-BGP) is used to distribute the second label in the stack, which is used for telling the egress PE router how to forward the incoming VPN packet.

Traffic Engineering with the MPLS-TE Model (Cont.)

Cisco.com

- **The MPLS-TE LSPs are created by RSVP.**
- **The actual path can be specified:**
 - **Explicitly defined by the system administrator**
 - **Dynamically defined using the underlying IGP protocol**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-15

For MPLS-TE, manual assignment and configuration of the labels can be used to create LSPs to tunnel the packets across the network on the desired path. However, to increase scalability, the Resource Reservation Protocol (RSVP) is used to automate the procedure.

The packets that are forwarded, according to MPLS-TE, have a stack of two labels, imposed by the ingress router. The topmost label identifies a specific LSP to use to reach another router at the other end of the tunnel. The second label indicates what the router at the far end of the LSP should do with the packet.

By selecting the appropriate LSP, a network administrator can direct traffic via explicitly indicated routers. The explicit path across identified routers provides benefits that are similar to the overlay model without introducing a Layer 2 network. This approach also eliminates the risk of running into IGP scalability problems because of the many neighbors that exist in a full mesh of routers.

MPLS-TE provides equivalent mechanisms to those described previously in this lesson in connection with the Layer 2 overlay network. For circuit-style forwarding, instead of using ATM or Frame Relay virtual circuits, the MPLS-TE tunnel is used. For signaling, RSVP is used with various extensions to set up the MPLS-TE tunnels.

For constraint-based routing (CBR), either Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF) with extensions is used to carry resource information like available bandwidth on the link. Both link-state protocols use new attributes to describe the nature of each link with respect to the constraints. A link that does not have the required resource is not included in the LSP, which constitutes the MPLS-TE tunnel.

To actually direct the traffic onto the MPLS-TE tunnels, network administrators need extensions to IS-IS and OSPF. Directing the traffic onto tunnels results in the addition of entries in the Forwarding Information Base (FIB). The IP packets are directed into the MPLS-TE tunnel by imposing the correct label stack.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- Traffic engineering measures, models, and controls traffic to achieve various goals.
- Network challenges: bandwidth may not be efficiently utilized, the least-cost route may not be the only possible route, and it may not have enough resources to carry all traffic.
- With the TE Layer 2 overlay model, routers are not aware of the physical structure and bandwidth available on links.
- With the TE Layer 3 model, the destination-based forwarding paradigm cannot handle the problem of overutilization of one path while an alternate path is underutilized.
- TE with the MPLS-TE model means that the routers use the MPLS label-switching paradigm.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-16

References

For additional information, refer to these resources:

- *RSVP-TE: Extensions to RSVP for LSP Tunnels*, RFC-3209, IETF Network Working Group, December 2001
- *MPLS Traffic Engineering*, RFC-2702, IETF Network Working Group, September 1999

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) _____ is the main motivation for implementing traffic engineering.
- Q2) Which two situations can result in network congestion? (Choose two.)
- A) The least-cost network path is being used to route traffic.
 - B) Traffic streams are efficiently mapped onto available resources.
 - C) Network resources are insufficient to accommodate offered load.
 - D) Traffic streams are inefficiently mapped onto available resources.
- Q3) Short traffic bursts are best handled by which of the following mechanisms? (Choose three.)
- A) rate limiting
 - B) traffic engineering
 - C) queuing algorithms
 - D) intelligent packet dropping
- Q4) When you are using TE with a Layer 2 overlay model, which two of the following transport traffic across a network? (Choose two.)
- A) DVCs
 - B) PVCs
 - C) RVCs
 - D) SVCs
- Q5) When you are using a traffic-engineered Layer 3 model, which two of the following are limitations?
- A) using static routes
 - B) using policy routing
 - C) unavailability of explicit routing
 - D) path computation that is based just on the IGP metric
- Q6) When you are using MPLS-TE, a minimum of _____ labels is assigned to a packet at the ingress router.

Lesson Answer Key

Q1) Cost reduction

Relates to: What Is Traffic Engineering?

Q2) C, D

Relates to: Business Drivers for Traffic Engineering

Q3) A, C, D

Relates to: Congestion Avoidance and Traffic Engineering

Q4) B, D

Relates to: Traffic Engineering with a Layer 2 Overlay Model

Q5) C, D

Relates to: Traffic Engineering with a Layer 3 Model

Q6) two

Relates to: Traffic Engineering with the MPLS-TE Model

MPLS Traffic Engineering Components

Overview

This lesson explains the components of MPLS-TE, such as traffic tunnels (along with associated characteristics and attributes), tunnel path discovery based on link-state protocols, and tunnel setup signaling with RSVP.

Relevance

This lesson is mandatory for learners who are planning to improve the usage of their network resources with MPLS-TE.

Objectives

This lesson describes the basic components of MPLS-TE. Upon completing this lesson, you will be able to do the following:

- Identify, at a conceptual level, how a traffic trunk functions
- Identify traffic trunk characteristics
- Identify traffic trunk attributes
- Identify the relation between network links and link attributes
- Identify the function of constraint-based path computation
- Identify the role of RSVP in path setup procedures
- Identify how using TE modifies the forwarding table mechanisms

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Carrier Supporting Carrier” module of this course

Outline

This lesson includes these topics:

- Overview
- Traffic Tunnels: Concepts
- Traffic Tunnels: Characteristics
- Traffic Tunnels: Attributes
- Network Links and Link Attributes
- Constraint-Based Path Computation
- Role of RSVP in Path Setup Procedures
- Forwarding Table Modifications
- Lesson Summary
- Lesson Review

Traffic Tunnels: Concepts

This topic describes the concept of traffic tunnels.

Traffic Tunnels: Concepts

Cisco.com

The concept of traffic tunnels (MPLS-TE tunnels) was introduced to overcome the limitations of hop-by-hop IP routing:

- **A tunnel is an aggregation of traffic flows that are placed inside a common MPLS label switched path.**
- **Flows are then forwarded along a common path within a service provider network.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-3


The aim of TE is to control the paths along which data flows, rather than relying simply on “normal” destination-based routing. To fulfill this aim, the concept of a “traffic tunnel” has been introduced.

A traffic tunnel is simply a collection of data flows that share some common attribute:

- Most simply, this attribute might be the sharing of the same entry point to the network and the same exit point. In practice, this point would be an Internet service provider (ISP) network, where there is a definable data flow from the points of presence (POPs), where the customers attach to the ISP network. There are also the Internet exchange points (IXPs), where data typically leaves the ISP network to traverse the Internet.
- In a more complex situation, this attribute could be augmented by defining separate tunnels for different classes of service. For example, in an ISP model, leased-line corporate customers could be given a preferential throughput over dial-in home users. This preference might be greater guaranteed bandwidth or lower latency and higher precedence. Even though the traffic enters and leaves the ISP network at the same points, different characteristics could be assigned to these types of users by defining separate traffic tunnels for their data.

Traffic Tunnels: Concepts (Cont.)

Cisco.com



- **Unidirectional single class of service model encapsulates all of the traffic between an ingress and an egress router.**
- **Different classes of service model assigns traffic into separate tunnels with different characteristics.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-4

Defining traffic tunnels requires an understanding of the traffic flows in the network. By understanding the ingress and corresponding egress points, a picture of the traffic flows in the network can be produced.

In the example here, there are two traffic tunnels (TT1 and TT2) that are defined for data from PE1 to PE3. These tunnels are unidirectional; they identify the traffic flows *from* PE1.

Note In practice, there are probably similar tunnels operating in the opposite direction *to* PE.

There may also be tunnels that are defined from all the other routers to each other. Defining tunnels from every router in the network to every other one might sound like an administrative nightmare. However, this is not usually the case for the following reasons:

- The routers that are identified are on the edge of the network. The traffic tunnels link these routers across the core of the network.
- In most networks it is relatively easy to identify the traffic flows, and they rarely form a complete “any-to-any” mesh.
- For example, in ISP networks, the traffic tunnels generally form a number of “star” formations, with their centers at the IXPs and the other points at the POPs. Traffic in an ISP network generally flows from the customers that are connected at the POPs to the rest of the Internet (reached via the IXPs). A starlike formation can also exist in many networks centering on the data center. This tendency is true for both ISP networks (providing web-hosting services) and enterprises.

Traffic Tunnels: Characteristics

This topic describes the characteristics of traffic tunnels.

Traffic Tunnels – Characteristics

Cisco.com

- **Traffic tunnels are routable objects (similar to ATM VCs).**
- **A traffic tunnel is distinct from the MPLS LSP through which it traverses:**
 - **In operational contexts, a traffic tunnel can be moved from one path onto another**
- **A traffic tunnel is assigned attributes influencing its characteristics.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-5


Once the data flows, and therefore the traffic tunnels are defined, the technology that they use to send the data across the network is MPLS. Data that enters a traffic tunnel is assigned an MPLS LSP, which defines the route that is taken through the network. However, traffic tunnels are distinct from the MPLS LSPs that they use in two key ways:

- There is not necessarily a one-to-one mapping of traffic tunnels onto MPLS LSPs. For administrative reasons, two tunnels may be defined between two points and may happen to pick the same path through the network. Therefore, they both have the same MPLS label.
- Also, traffic tunnels are not necessarily bound to a particular path through the network. As resources change in the core, or perhaps as links fail, the traffic tunnel may reroute, picking up a new MPLS LSP as it does.

The configuration of the traffic tunnels includes defining the characteristics and attributes that it requires. In fact, defining the characteristics and attributes of traffic tunnels is probably the most important aspect of TE. Without a specification of the requirements of the data in this traffic tunnel, the data might as well be left to route “normally” based only on destination information over the least-cost path.

Traffic Tunnels: Attributes

This topic describes the attributes of traffic tunnels.



A traffic tunnel is a set of data flows sharing some common feature, attribute, or requirement. If there is no characteristic in the data flow to make it common with some other flow, there is nothing to define that data as part of a flow or group of flows.

Therefore, the traffic tunnel, by its very definition, must include attributes that define the commonality between the data flows making up the tunnel. The attributes that characterize a traffic tunnel include the following:

- **Ingress and egress points:** These points are, fundamentally, the routers at the ends of the tunnel. They are the most basic level of commonality between data flows given that the flows in a tunnel all start in the same place and end in the same place.
- **Complex characteristics of the data flows:** Examples are bandwidth, and latency and precedence requirements.
- **Class of data:** This attribute encompasses what data is part of this tunnel and what is not. This includes such characteristics as traffic flow, class of service, and application class.

The network administrator defines the attributes of a traffic tunnel when the tunnel itself is defined. However, some of these attributes are in part influenced by the underlying network and protocols.

Traffic Tunnels – Attributes (Cont.)

Cisco.com

The administrator enters the relevant information (attributes) at the headend of the traffic tunnel:

- **Traffic parameter**—resources required for tunnel (e.g., required bandwidth)
- **Generic path selection and management**—path can be administratively specified or computed by the IGP
- **Resource class affinity**—include or exclude certain links for certain traffic tunnels
- **Adaptability**—should the traffic tunnel be reoptimized?
- **Priority and pre-emption**—importance of a traffic tunnel and possibility for a pre-emption of another tunnel
- **Resilience**—desired behavior under fault conditions

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-7

The general tunnel characteristics must be configured by the network administrator to create the tunnel. This includes some or all of the following:

- **Traffic parameters:** Traffic parameters are the resources that are required by the tunnel, such as the minimum required bandwidth.
- **Generic path selection and management:** This category refers to the path selection criteria. The actual path that is chosen through the network could be statically configured by the administrator or could be assigned dynamically by the network, based on information from the IGP (IS-IS or OSPF).
- **Resource class affinity:** This category refers to restricting the choice of paths by allowing the dynamic path to choose only certain links in the network rather than being allowed to use any link.

Note This restriction can also be accomplished by using the IP address exclusion feature.

- **Adaptability:** Adaptability is the ability of the path to reroute on failure or to optimize on recovery or discovery of the “better” path.
- **Priority and pre-emption:** Traffic tunnels can be assigned a priority (0 to 7) that signifies their “importance.” When you are setting up a new tunnel or rerouting, a higher-priority tunnel can tear down (pre-empt) a lower-priority tunnel; in addition, a new tunnel of lower priority may fail to set up because some tunnels of a higher priority already occupy the required bandwidth of the lower-priority tunnel.
- **Resilience:** Resilience refers to how a traffic tunnel responds in the event of a failure in the network. Does it attempt to reroute around failures or not?

Network Links and Link Attributes

This topic discusses the relationship between network links and link attributes.

Network Links and Link Attributes

Cisco.com

Resource attributes (link availability) are configured locally on the router interfaces:

- **Maximum bandwidth**
 - The amount of bandwidth available
- **Link affinity string**
 - To allow the operator to administratively include or exclude links in path calculations
- **Constraint-based specific metric**
 - Traffic engineering default metric

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-8

In order for the tunnel to dynamically discover its path through the network, the headend router must be provided with information on which to base this calculation. Specifically, it needs to be provided with the following:

- **Maximum bandwidth:** The maximum bandwidth is the amount of bandwidth that is available on each link in the network. Because there are priority levels for traffic tunnels, the availability information must be sent for each priority level for each link. Including priority levels means that the path decision mechanism is given the opportunity to choose a link with some bandwidth already allocated to a lower-priority tunnel, forcing that lower-priority tunnel to be “bounced” off the link.
- **Link resource class:** For administrative reasons, the network administrator may decide that some tunnels are not permitted to use certain links. To accomplish this goal, for each link, a link resource class must be defined and advertised. The definition of the tunnel may include a reference to particular “affinity bits.” The tunnel affinity bits are matched against the link resource class to determine if a link may or may not be used as part of the LSP.
- **Constraint-based specific metric:** Each link has a cost or metric for calculating routes in the normal operation of the IGP. It may be that, when calculating the LSP for traffic tunnels, the link should use a different metric. Thus, a constraint-based specific metric may be specified.

Constraint-Based Path Computation

This topic introduces constraint-based path computation.

Constraint-Based Path Computation

Cisco.com

- **Constraint-based routing is demand-driven.**
- **Resource-reservation-aware routing paradigm:**
 - **Based on criteria including, but not limited to, network topology**
 - **Calculated at the edge of a network:**
 - **Modified Dijkstra's algorithm at tunnel headend (CSPF [constrained SPF] or PCALC [Path Calculation]).**
 - **Output is a sequence of IP interface addresses (next-hop routers) between tunnel endpoints.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-9

In traditional networks, the IGP calculates paths through the network based on the network topology alone. Routing is destination-based, and all traffic to a given destination from a given source uses the same path through the network. That path is based simply on what the IGP regards as the “least cost” between the two points (source and destination).

Constraint-based routing (CBR) is the term that is used most often for this approach. In some situations it is also referred to as a constrained shortest path first (CSPF) calculation or a path calculation (PCALC).

CBR behaves in the following ways:

- Augments the use of link cost by also considering other factors such as bandwidth availability or link latency when choosing the path to a destination.
- Tends to be carried out at the edge of the network, discovering a path across the core to some destination elsewhere at the other edge of the network. Typically, this discovery uses the CSPF calculation (a version of shortest path first [SPF] that is used by IS-IS and OSPF, but considering other factors besides cost, such as bandwidth availability).
- Produces a sequence of IP addresses that correspond to the routers that are used as the path to the destination; these addresses are the next-hop addresses for each stage of the path.

A consequence of CBR is that, from one source to one destination, many different paths can be used through the network, depending on the requirements of those data flows.

Constraint-Based Path Computation (Cont.)

Cisco.com

- **Constraint-based routing takes into account:**
 - Policy constraints associated with the tunnel and physical links
 - Physical resource availability
 - Network topology state
- **Two types of tunnels can be established across those links with matching attributes:**
 - **Dynamic**—using the least-cost path computed by OSPF/IS-IS
 - **Explicit**—definition of a path by using Cisco IOS configuration commands

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-10

When choosing paths through the network, the CBR system takes into account the following factors:

- The topology of the network, including information about the state of the links (the same information that is used by normal hop-by-hop routing)
- The resources that are available in the network, such as the bandwidth not already allocated on each link and at each of the eight priority levels (priority 0 to 7)
- The requirements that are placed on the constraint-based calculation that is defining the policy or the characteristics of this traffic tunnel


Of course, CBR is a dynamic process, which responds to a request to create a path and calculates (or recalculates) the path based on the status of the network at that time. The network administrator can also explicitly define the traffic tunnel.

By using commands like **exclude-address** or **next-hop loose** in the explicit path configuration, the network administrator can mix static and dynamic computation.

Constraint-Based Path Computation (Cont.)

Cisco.com

What is the best path from R1 to R6 with bandwidth of 30 Mbps?



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-11

An example network is shown in the figure. Each link specifies a link cost for metric calculation and a bandwidth available for reservation; for example, a metric of 10 and an available bandwidth of 100 Mbps is shown for the link between R1 and R2. Other than these criteria, no links are subject to any policy restriction that would disallow their use for creating traffic tunnels.

The requirement is to create a tunnel from R1 to R6 with a bandwidth of 30 Mbps.


Based simply on the link costs, the least-cost path from R1 to R6 is R1-R4-R6 with a cost of 30. However, the link from R4 to R6 has only 20 Mbps of bandwidth available for reservation and therefore cannot fulfill the requirements of the tunnel.

Similarly, the link R5-R6 has only 20 Mbps available as well, so no paths can be allocated via R5.

Constraint-Based Path Computation (Cont.)

Cisco.com

Computed path for a dynamic constraint-based tunnel over the least-cost path.



Administratively defined explicit path Tunnel is still possible over any eligible path.

Path has cost of 45, not the lowest cost.

Explicit and Dynamic Traffic Engineering Tunnels

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-12

The diagram now shows only those links that can satisfy the requirement for 30 Mbps of available bandwidth.

Over this topology, two tunnel paths are shown:

- The network administrator has statically defined the blue colored path (R1-R4-R3-R6). Had the administrator attempted to define a path that did not have the required free bandwidth, the tunnel establishment would have failed. This tunnel does indeed fulfill the minimum bandwidth requirement. However, adding the link costs yields a total of 45, which is not the lowest cost possible.
- The red (upper) path shows the result of a dynamic constraint-based path calculation. The calculation has ignored any links that do not satisfy the bandwidth requirement (those from the last diagram [not shown here], such as the connections to R5) and then executes a CSPF calculation on what remains. This calculation has yielded the path R1-R2-R3-R6 with a path cost of 40.

Role of RSVP in Path Setup Procedures

This topic explains how RSVP is used in the setting up an LSP path.

Role of RSVP in Path Setup Procedures

Cisco.com

- **Once the path has been determined, a signaling protocol is needed:**
 - **To establish and maintain label switched paths (LSPs) for traffic tunnels**
 - **For creating and maintaining resource reservation states across a network (bandwidth allocation)**
- **The Resource Reservation Protocol (RSVP) was adopted by the MPLS workgroup of the IETF.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-13

The result of the constraint-based calculation is a list of routers that form the path to the destination. The path is a list of IP addresses that identify each next hop along the path.

However, this list of routers is known only to the router at the headend of the tunnel that is attempting to build the tunnel. Somehow, this now explicit path must be communicated to the intermediate routers. It is not up to the intermediate routers to make their own CSPF calculations: they merely abide by the path that is provided to them by the headend router.

Therefore, some signaling protocol is required to confirm the path, to check and apply the bandwidth reservations, and finally to apply the MPLS labels to form the MPLS LSP through the routers. RSVP is used to confirm and reserve the path and apply the labels that identify the tunnel. LDP or TDP is used to apply the labels for the underlying MPLS network..

Note RSVP is needed for both explicit and dynamic path setup.

Forwarding Table Modifications

This topic discusses changes that occur in the forwarding table when MPLS-TE is implemented.

Forwarding Table Modifications

Cisco.com

- **IP routing is separate from LSP routing and does not see internal details of the LSP.**
- **The traffic has to be mapped to the tunnel:**
 - **Static routing**—the static route in the IP routing table points to an LSP tunnel interface.
 - **Policy routing**—the next-hop interface is an LSP tunnel
 - **Forwarding adjacency**—the tunnel is announced as a point-to-point link to all other routers within an area
 - **Autoroute**—SPF enhancement:
 - **The headend sees the tunnel as a directly connected interface (for modified SPF only).**
 - **The default cost of a tunnel is equal to the shortest IGP metric regardless of the used path.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-14

The tunnel normally does not appear in the IP routing table. The IP routing process does not see the tunnel, so the tunnel is normally not included in any SPF calculations. The IP traffic can be mapped onto a tunnel in four different ways:

- Using static routes that point to the tunnel interfaces.
- Using policy-based routing (PBR) and setting the next hop for the destination to the tunnel interface.
- Using forwarding adjacency, which allows the tunnel to be announced via OSPF or IS-IS like any other unidirectional link (UDL). In order to be used for data forwarding such a tunnel has to be set up bidirectionally.
- Using the autoroute feature, which is an SPF enhancement that includes the tunnel interface in the route calculation as well. The result of the autoroute feature is that the tunnel is seen at the headend (and only there) as a directly connected interface. The metric (cost) of the tunnel is set to the normal IGP metric from the tunnel headend to the tunnel endpoint (over the least-cost path, regardless if the tunnel is actually using the least-cost path or not).

Note With the autoroute feature, the traffic-engineered tunnel appears in the IP routing table as well, but this appearance is restricted to the tunnel headend only.

The first two options are not very flexible or scalable. The traffic for each destination that needs to use the tunnel must be manually mapped to the tunnel.

For example, when you are using static routes, the tunnel is used only for the explicit static routes. Any other traffic that is not covered by the explicit static routes, including traffic for the tailend router (even though the tunnel terminates on it), will not be able to use the tunnel; instead, it will follow the normal IGP path.

Note	The autoroute and forwarding adjacency features are explained in detail later in this module.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- Traffic tunnels are configured with a set of resource requirements, such as bandwidth and priority.
- CSPF augments the link cost by considering other factors such as bandwidth availability or link latency when you are choosing the path.
- Resource Reservation Protocol (RSVP) with TE extensions is used as a mechanism for establishing and maintaining label switched paths (LSPs).
- TE tunnels do not appear in the IP routing table.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-15

References

For additional information, refer to this resource:

- RFC 2746, *RSVP Operation Over IP Tunnels*

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) A traffic tunnel is a set of _____ that share some common feature, attribute, or requirement.
- Q2) Link resource attributes are distributed to the _____ of traffic tunnels.
- Q3) The constraint-based path computation uses which algorithm?
- A) DUAL algorithm
 - B) modified Dijkstra's algorithm
 - C) modified Bell-Howell algorithm
 - D) none of the above
- Q4) When a traffic tunnel comes up, what will appear in the IP routing table related to the tunnel?
- Q5) What is the role of RSVP in an MPLS-TE implementation?
- A) It identifies the best path for the tunnel.
 - B) It reserves the bandwidth required by the tunnel.
 - C) It performs the CBP calculations for the tunnel setup.
 - D) It assigns the label for the MPLS LSP.

Lesson Answer Key

Q1) data flows

Relates to: Traffic Tunnels: Concepts; Traffic Tunnels: Characteristics; Traffic Tunnels: Attributes

Q2) headend

Relates to: Network Links and Link Attributes

Q3) B

Relates to: Constraint-Based Path Computation

Q4) nothing, unless the tunnel is advertised to IP using static routes, policy-based routing (PBR), forwarding adjacency, or autoroute .

Relates to: Forwarding Table Modifications

Q5) C

Relates to: Role of RSVP in Path Setup Procedures

Constraint-Based Path Computation

Overview

This lesson describes the details of link attribute propagation with an IGP protocol and constraint-based path computation.

Relevance

This lesson is mandatory for learners who are planning to improve usage of their network resources with MPLS-TE.

Objectives

This lesson describes constraint-based path computation. Upon completing this lesson, you will be able to do the following:

- Identify the constraint-based path computation resource attributes
- Explain how MPLS-TE link attributes affect constraint-based path computation
- Explain how MPLS-TE tunnel attributes affect constraint-based path computation
- Identify how TE policies are implemented using affinity bits
- Describe how MPLS-TE link attributes are propagated with link-state routing protocols
- Describe the characteristics of constraint-based path computation as it relates to LSPs

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Carrier Supporting Carrier” module of this course

Outline

This lesson includes these topics:

- Overview
- Constraint-Based Path Computation Attributes
- MPLS-TE Link Resource Attributes
- MPLS-TE Tunnel Attributes
- Implementing TE Policies with Affinity Bits
- Propagating MPLS-TE Link Attributes with Link-State Routing
- Constraint-Based Path Computation
- Lesson Summary
- Lesson Review

Constraint-Based Path Computation Attributes

This topic identifies the constraint-based path computation resource attributes.

Constraint-Based Path Computation

Cisco.com

Constraint-based path computation provides several resource attributes to control LSP path determination.

- **Link resource attributes that provide information on the resources of each link.**
- **Traffic tunnel attributes characterize the traffic tunnel.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-3

Constraint-based path computation, which takes place at the headend of the traffic-engineered tunnel, must be provided with several resource attributes before the LSP path is actually determined. These attributes include the following:

- Link resource attributes that provide information on the resources of each link
- Traffic tunnel attributes that characterize the traffic tunnel

MPLS-TE Link Resource Attributes

This topic describes the MPLS-TE link resource attributes.

MPLS-TE Link Resource Attributes

Cisco.com

- **Maximum bandwidth**
- **Maximum reservable bandwidth**
- **Link resource class**
- **Constraint-based specific link metric**


© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-4

There are four link resource attributes:

- Maximum bandwidth
- Maximum reservable bandwidth
- Link resource class
- Constraint-based specific link metric

MPLS-TE Link Resource Attributes: Maximum Allocation Multiplier

Cisco.com



- **Maximum bandwidth:** the maximum bandwidth that can be used on this link in this direction (physical link)
- **Maximum reservable bandwidth:** The maximum amount of bandwidth that can be reserved in this direction on this link

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-5


Among the link resource attributes, the most important is the maximum allocation multiplier. This attribute deals with the amount of bandwidth that is available on a specified link. “Available” means not already allocated (as opposed to presently in use); the attribute is thus a measure of allocation not utilization. Furthermore, because there are priority levels for traffic tunnels, this availability information needs to be configured for each priority level on the link. Normally, the bandwidth at the upper priority level is always higher than at lower levels (0-7 levels). Because of oversubscriptions, the total amount of bandwidth can exceed the actual bandwidth of the link. There are three components to the link resource attribute:

- **Maximum bandwidth:** This component provides information on the maximum bandwidth that can be used on the link, per direction, given that the traffic tunnels are unidirectional. This parameter is usually set to the configured bandwidth of the link.
- **Maximum reservable bandwidth:** This component provides information on the maximum bandwidth that can be reserved on the link per direction. By default, it is set to 75 percent of the maximum bandwidth.
- **Unreserved bandwidth:** Unreserved bandwidth provides information on the remaining bandwidth that has not yet been reserved.

Note A higher priority can pre-empt a lower priority, but a lower priority can not pre-empt a higher priority.

MPLS-TE Link Resource Attributes: Link Resource Class

Cisco.com



- Link is characterized by a 32-bit resource class attribute.
- Associated with a traffic tunnel in order to include or exclude certain links into or from the path of the traffic tunnel.


© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-6

For each link, another link resource attribute, the link resource class, is provided as well. The link is characterized by a 32-bit link resource class attribute string, which is matched with the traffic tunnel resource class affinity attribute and allows inclusion or exclusion of the link into or from the path of the tunnel.

MPLS-TE Link Resource Attributes: Constraint-Based Specific Link Metric

Cisco.com



This metric is administratively assigned to present a differently weighted topology to traffic engineering SPF calculations:

- Administrative weight (TE metric)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-7

Each link has a cost or metric for calculating routes in the normal operation of the IGP. It may be that, when calculating paths for traffic tunnels, the link should use a different metric than the IGP metric. Hence, a constraint-based specific link metric, the administrative weight, may be administratively assigned as well.

MPLS-TE Tunnel Attributes

This topic describes MPLS-TE tunnel attributes.

MPLS-TE Tunnel Attributes

Cisco.com

- **Traffic parameter**
- **Generic path selection and management**
- **Tunnel resource class affinity**
- **Adaptability**
- **Priority**
- **Pre-emption**
- **Resilience**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-8

Seven tunnel attributes are available to influence path selection:

- Traffic parameter
- Generic path selection and management
- Tunnel resource class affinity
- Adaptability
- Priority
- Pre-emption
- Resilience

MPLS-TE Tunnel Attributes (Cont.)

Cisco.com

- **Traffic parameter:**
 - Indicates the resource requirements (for example, bandwidth) of the traffic tunnel
- **Generic path selection and management:**
 - Specifies how the path for the tunnel is computed:
 - **Dynamic LSP — Constraint-based computed paths based on a combination of bandwidth and policies**
 - **Explicit LSP — administratively specified off line (typically using CLI)**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—3-9

A couple of the MPLS-TE tunnel attributes affect the path setup and maintenance of the traffic tunnel:

- The traffic parameter (bandwidth) attribute specifies (among other traffic characteristics) the amount of bandwidth that is required by the traffic tunnel. The traffic characteristics may include peak rates, average rates, permissible burst size, and so on. From a TE perspective, the traffic parameters are significant because they indicate the resource requirements of the traffic tunnel. These characteristics are useful for resource allocation.
- The path selection and management attribute (path selection policy) specifies the way in which the headend routers should select explicit paths for traffic tunnels. The path can be configured manually or computed dynamically using the constraint-based path computation; both methods take the resource information and policies into account.

MPLS-TE Tunnel Attributes (Cont.)

Cisco.com



Tunnel Resource Class Affinity:

- The properties that the tunnel requires from internal links:
 - 32-bit resource class affinity bit string + 32-bit resource class mask
- Link is included in the constraint-based LSP path when the tunnel resource affinity string or mask matches the link resource class attribute.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-10

The tunnel resource class affinity attribute allows the network administrator to apply path selection policies by administratively including or excluding network links. Each link may be assigned a resource class attribute. Resource class affinity specifies whether to explicitly include or exclude links with resource classes in the path selection process. The resource class affinity is a 32-bit string that is accompanied by a 32-bit resource class mask. The mask indicates which bits in the resource class need to be inspected. The link is included in the constraint-based LSP when the resource class affinity string or mask matches the link resource class attribute.

MPLS-TE Tunnel Attributes (Cont.)

Cisco.com

- **Adaptability:**
 - If reoptimization is enabled, then a traffic tunnel can be rerouted through different paths by the underlying protocols:
 - Primarily due to changes in resource availability
- **Priority:**
 - Relative importance of traffic tunnels
 - Determines the order in which path selection is done for traffic tunnels at connection establishment and under fault scenarios:
 - Setup priority: Priority for taking a resource
- **Pre-emption:**
 - Determines whether another traffic tunnel can pre-empt a specific traffic tunnel:
 - Hold priority: Priority for holding a resource

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-11

The adaptability attribute indicates whether the traffic tunnel should be reoptimized and consequently rerouted to another path primarily because of the changes in resource availability.

The priority and pre-emption tunnel attributes are closely associated and play an important role in competitive situations where traffic tunnels compete for link resources. Two types of priorities are assigned to each traffic tunnel:

- Setup priority (priority) defines the relative importance of traffic tunnels and determines the order in which path selection is done for traffic tunnels at connection establishment and during rerouting because of faulty conditions. Priorities are also important at implementation; they can permit pre-emption in that they can be used to impose a partial order on a set of traffic tunnels according to which pre-emptive policies can be actualized.
- Holding priority (pre-emption) defines the pre-emptive rights of competing tunnels and specifies the priority for holding a resource. This attribute determines whether a traffic tunnel can pre-empt another traffic tunnel from a given path, and whether another traffic tunnel can pre-empt a specific traffic tunnel. Pre-emption can be used to ensure that high-priority traffic tunnels can always be routed through relatively favorable paths within a Differentiated Services (DiffServ) environment. Pre-emption can also be used to implement various prioritized restoration policies following fault events.

MPLS-TE Tunnel Attributes (Cont.)

Cisco.com

Resilience:

- **Determines the behavior of a traffic tunnel under fault conditions:**
 - **Do not reroute the traffic tunnel**
 - **Reroute through a feasible path with enough resources**
 - **Reroute through any available path regardless of resource constraints**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-12

Two additional tunnel attributes define the behavior of the tunnel in faulty conditions or in cases when the tunnel becomes noncompliant with tunnel attributes (for example, the required bandwidth):

The resilience attribute determines the behavior of the tunnel under faulty conditions; it can specify the following behavior:

- Not to reroute the traffic tunnel at all
- To reroute the tunnel through a path that can provide the required resources
- To reroute the tunnel through any available path irrespective of available link resources

Implementing TE Policies with Affinity Bits

This topic describes how to implement TE by using the tunnel affinity bits.

Implementing TE Policies with Affinity Bits

Cisco.com

- Link is characterized by the link resource class
 - Default value of bits is 0
- Tunnel is characterized by:
 - Tunnel resource class affinity
 - Default value of bits is 0
 - Tunnel resource class affinity mask
 - (0=do not care, 1=care)
 - Default value of the tunnel mask is 0x0000FFFF

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0–3-13

The policies during LSP path computation can be implemented using the resource class affinity bits of the traffic tunnel and the resource class bits of the links over which the tunnel should pass (following the computed LSP path).

Each traffic tunnel is characterized by a 32-bit resource class affinity string, which is accompanied by a respective resource class mask. The zero bits in the mask exclude the respective link resource class bits from being checked.

Each link is characterized by its resource class 32-bit string, which is set to 0 by default. The matching of the tunnel resource class affinity string with the resource class string of the link is performed during the LSP path computation.


Note You can also exclude links or nodes by using the IP address exclusion feature when you are configuring tunnels using the **explicit-path** command.

Implementing TE Policies with Affinity Bits (Cont.)

Cisco.com

Setting a link bit in the lower half drives all tunnels off the link, except those specially configured.

Tunnel Affinity: bits = 0000, mask = 0011



Using Affinity Bits to Avoid Specific Links

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-14

The example here shows a sample network with tunnel resource class affinity bits and link resource bits. For simplicity, only the four affinity and resource bits (of the 32-bit string) are shown. The tunnel should be established between routers A (headend) and B (tailend).

With the tunnel resource class affinity bits and the link resource class bits at their default values of 0, the constraint-based path computation would have two possible paths: A-D-E-B or A-D-C-E-B.

Because it is desirable to move all dynamically computed paths away from the link D-E, the link resource class bits were set to a value 0010 and the tunnel mask was set to 0011.


In the example, the tunnel mask requires only the lower two bits to match. The 00 of the traffic affinity does not match the 10 of the link D-E resource class and results in the exclusion of this link as a possible path for the tunnel. The only remaining alternative path is D-C-E, on which the default values of the resource class string (all zeros) match the tunnel affinity bits.

Implementing TE Policies with Affinity Bits (Cont.)

Cisco.com

A specific tunnel can then be configured to allow all links by clearing the bit in its affinity attribute mask.

Tunnel Affinity: bits = 0000, mask = 0001



Using the Affinity Bit Mask to Allow all Links

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-15

In the sample network here, only the lower bit has been set in the tunnel mask. The tunnel affinity bits remain unchanged as well as the resource class bits on the D-E link.

The matching between the tunnel resource class affinity bits and the link resource class bits is done on the lowest bit only (because of the mask setting of 0001). The 0 of the tunnel affinity bit (the lowest bit) matches with the 0 of the link resource class bit (the lowest bit) and therefore the link D-E remains in the possible path computation (along with the D-C-E link).


Which path will actually be used depends on other tunnel and link attributes, including the required and available bandwidth.

Implementing TE Policies with Affinity Bits (Cont.)

Cisco.com

A specific tunnel can be restricted to only some links by turning on the bit in its affinity attribute bits.

Tunnel Affinity: bits = 0010, mask = 0011



Using Affinity Bits to Dedicate Links to Specific Purposes

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-16

This last example deals with setting the tunnel resource class affinity bits and the link resource class bits to force the tunnel to follow a specific path. Links A-D-E-B are all configured with the resource class value 0010.


The tunnel resource class affinity bits are set to a value of 0010 and the mask to 0011. Only the lower two bits will be compared in the constraint-based path computation.

The 10 of the tunnel resource class affinity matches the 10 of the link resource class on all links that are configured with that value.

The 10 does not match the 00 that is set on the path D-C-E, and thus only one possible LSP path remains (A-D-E-B).

Propagating MPLS-TE Link Attributes with Link-State Routing

This topic describes how link resource attributes are propagated using a link-state protocol, OSPF or IS-IS.



The link resource attributes must be propagated throughout the network to be available at the headend of the traffic tunnel when LSP path computation takes place.

Because the propagation (flooding) of the attributes can be achieved only by link-state routing protocols (IGPs), OSPF and IS-IS were extended to support the MPLS-TE features.

OSPF uses new link-state advertisements (opaque LSAs), and IS-IS uses new type, length, value (TLV) attributes in its link-state packets.

Another important factor in LSP computation is the available bandwidth on the link over which the traffic tunnel will pass. This bandwidth is configured per priority level (8 levels, 0 being the highest, 7 the lowest) and communicated in respective IGP link-state updates, again per priority.

When a certain amount of the bandwidth is reserved at a certain priority level, this amount is subtracted from the available bandwidth at that level and at all levels below. The bandwidth at upper levels remains unchanged.

In the figure, the maximum bandwidth is set to the bandwidth of the link, which is 100 (on a Fast Ethernet link). The system allows the administrator to set the available bandwidth (AB) to a higher value than the interface bandwidth. When the administrator is making a reservation, any bandwidth above the interface bandwidth will be rejected. The available bandwidth is advertised in the link-state packets of router D. The value is 100 at all priority levels before any tunnel is set up.

In the next part of the figure, a tunnel at priority level 3 that requires 30 units of bandwidth is set up across the link L. The available bandwidth at all priority levels above level 3 (0, 1, and 2) remains unchanged at 100. On all other levels, 30 is subtracted from 100, which results in an available bandwidth of 70 at priority level 3 and below (4-7).

Lastly, another tunnel is set up at priority level 5 that requires 30 units of bandwidth across the link L. The available bandwidth at all priority levels above level 5 remains unchanged (100 on levels 0 to 2, and 70 on levels 3 and 4). On all other levels, 30 is subtracted from 70, which results in an available bandwidth of 40 at priority level 5 and below (6-7).

Note	All bandwidth reservation is done in the control plane and does not affect the actual traffic rates in the data plane.

Propagating MPLS-TE Link Attributes with Link-State Routing (Cont.)

Cisco.com

IGP resource flooding takes place in the following situations:

- **Link-state changes**
- **Resource class of a link changes:**
 - **Manual reconfiguration**
 - **Amount of available bandwidth crosses one of the preconfigured thresholds**
- **Periodic (timer-based):**
 - **A node checks attributes; if they are different, it floods its update status**
- **On LSP setup failure**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-18

The flooding of resource attributes by the IGP takes place in connection with certain conditions and events:


- When the link changes its state (up, down).
- When the resource class of the link changes because of a manual reconfiguration or because a preconfigured threshold is crossed by the available bandwidth.
- Periodically (based on a timer), a node checks resource attributes, and if the resource attributes were changed, the update is flooded.
- When the LSP setup fails.

Propagating MPLS-TE Link Attributes with Link-State Routing (Cont.)

Cisco.com

Significant Change and Preconfigured Thresholds

- For stability reasons, rapid changes should not cause rapid generation of updates:



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-19

For stability purposes, significant rapid changes in available link resources should not trigger the updates immediately.

There is a drawback, however, in not propagating the changes immediately. In some cases the headend sees the link as available for the LSP and includes the link in its path computation even though the link may be down or does not have the required resource available. When the LSP is actually being established, a node with a link that lacks the required resources cannot establish the path and floods an immediate update to the network.

The thresholds for resources are set both for an up direction (resources exceeding the threshold) and a down direction (resources dropping below the threshold). When the threshold is crossed (in either direction), the node generates an update that carries the new resource information.

The figure shows the threshold values for the up direction (100 percent, 92 percent, 85 percent, 70 percent, and 50 percent) and two updates being sent out. In this example one update is immediately sent when the 50 percent margin is crossed. The second is sent when the 70 percent margin is sent.

Constraint-Based Path Computation

This topic describes constraint-based path computation.

Constraint-Based Path Computation

Cisco.com

When establishing a tunnel, the edge routers have knowledge of both network topology and link resources within its area:

- **Two methods for establishing traffic tunnels:**
 - **Static and dynamic path setup**
- **In both cases the result is an explicit route expressed as a sequence of interface IP addresses (for numbered links) or TE router IDs (for unnumbered links) in the path from tunnel endpoints.**
- **RSVP is used to establish and maintain constraint-based label switched paths for traffic tunnels along an explicit path.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-20

The headend of the traffic tunnel has visibility of both the network topology and network resources. This information is flooded throughout the network via a link-state IGP.

The LSP path for the traffic tunnel can be statically defined or computed dynamically. The computation takes the available resources and other tunnel and link attributes into account (thus, it represents constraint-based path computation). The result of the constraint-based path computation is a series of IP addresses that represent the hops on the LSP path between the headend and tailend of the traffic tunnel.

For LSP signaling and the final establishment of the path, RSVP is used.

Constraint-Based Path Computation (Cont.)

Cisco.com

Dynamic constraint-based path computation is triggered by the headend of the tunnel:

- **For a new tunnel**
- **For an existing tunnel whose current LSP has failed**
- **For an existing tunnel when you are doing reoptimization**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-21

Constraint-based path computation is always performed at the traffic tunnel headend. The computation is triggered for the following situations:

- A new tunnel
- An existing tunnel whose LSP setup has failed
- The reoptimization of an existing traffic tunnel

The LSP path computation is restricted by several factors (constraint-based). The LSP path can be computed only if the following conditions are met:

- The endpoints of the tunnel are in the same OSPF or IS-IS area (due to link-state flooding of resources).
- The links that are explicitly excluded via the link resource class bit string, or that cannot provide the required bandwidth, are pruned from the computation.

Constraint-Based Path Computation (Cont.)

Cisco.com

Constraint-Based Path Selection

- **Path selection:**
 - CBR uses its own metric (administrative weight, or TE cost; by default equal to the IGP cost)—used only during constraint-based computation
 - In case of a tie, select the path with:
 - The highest minimum bandwidth
 - The smallest hop count
 - If everything else fails, then pick a path at random
- **LSP path setup**—an explicit path is used by RSVP to reserve resources and establish LSP path
- **Final result:** unidirectional MPLS-TE tunnel, seen only at the headend router

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-22

Constraint-based path computation selects the path that the traffic tunnel will take based on the administrative weight (TE cost) of each individual link. This administrative weight is by default equal to the IGP link metric. The value is used only during the constraint-based path computation.

If there are more candidates for the LSP path (several paths with the same metric), then the selection criteria is as follows (in sequential order):

- The highest minimum bandwidth on the path takes precedence.
- The smallest hop count takes precedence.

If, after applying both of these criteria, more than one path still exists, the path is randomly chosen.

When the LSP path is computed, RSVP is used to actually reserve the bandwidth, to allocate labels for the path, and finally to establish the path.

The result of a constraint-based path computation is a unidirectional MPLS-TE tunnel (traffic tunnel) that is seen only at the tunnel endpoints (headend and tailend).

Constraint-Based Path Computation (Cont.)

Cisco.com

- **MPLS-TE tunnel is not a link for link-state adjacency:**
 - Establishment of a tunnel does not trigger any LSA announcements or a new SPF calculation (unless the forwarding adjacency feature is enabled).
 - Cisco IOS software uses tunnel interface for MPLS-TE tunnel creation and visualization, but behavior of MPLS-TE tunnels is fairly different from other tunnel protocols (for example, GRE).
- **Only traffic entering at headend router will use tunnel**
- **IP cost: If autoroute is used, MPLS-TE tunnel in the IP routing table has a cost of the shortest IGP path to the tunnel destination (regardless of the LSP path)**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-23

From the perspective of IGP routing, the traffic tunnel is not seen as an interface at all and is not included in any IGP route calculation (apart from other IP tunnels like generic routing encapsulation [GRE] tunnels). The traffic-engineered tunnel, when established, does not trigger any link-state update or any SPF calculation.

This behavior can be changed by using the **mpls traffic-eng forwarding-adjacency** command and by defining two tunnels in a bidirectional way.

Cisco IOS software uses the tunnel mainly for visualization. The rest of the actions, associated with the tunnel, are done by MPLS forwarding and other MPLS-TE-related mechanisms.

The IP traffic that will actually use the traffic-engineered tunnel is forwarded to the tunnel only by the headend of the tunnel. In the rest of the network, the tunnel is not seen at all (no link-state flooding).

With the autoroute feature, the traffic tunnel has the following characteristics:

- Appears in the routing table
- Has an associated IP metric (cost equal to the best IGP metric to the tunnel endpoint)
- Is also used to forward the traffic for destinations behind the tunnel endpoint

Even with the autoroute feature, the tunnel itself is not used in link-state updates and the rest of the networks do not have any knowledge of it.


Constraint-Based Path Computation (Cont.)

Cisco.com

Request by tunnel:

From R1 to R6; Priority 3, BW = 30 Mbps
Resource Affinity: bits = 0010, mask = 0011

Link R4-R3
is excluded.



Path Selection Considering Policy Constraints

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-24

This example of constraint-based path computation and LSP path selection requires that the traffic tunnel be established between R1 (headend) and R6 (tailend). The traffic tunnel requirements are as follows:

- The required bandwidth at priority level 3 is 30 Mbps.
- The resource class affinity bits are set to 0010, and the tunnel mask is 0011. The checking is done only on the lower two bits.

The link R4-R3 should be excluded from the LSP path; therefore, its resource class bit string is set to 0011. When the traffic tunnel resource class affinity bits are compared to the link R4-R3 resource class bits, there is no match, and the link is effectively excluded from the LSP path computation.


Constraint-Based Path Computation (Cont.)

Cisco.com

Request by tunnel:

From R1 to R6; Priority 3, BW = 30 Mbps
Resource Affinity: bits = 0010, mask = 0011

Not enough
bandwidth



Path Selection Considering Available Resources

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-25


The next parameter that is checked during constraint-based path computation is the TE cost (administrative weight) of each link through which the tunnel will possibly pass. The lowest cost is calculated across the path R1-R4-R6; the overall cost is 30. All other possible paths have a higher overall cost.

When resources are taken into account, constraint-based path computation finds that on the lowest-cost path there is not enough bandwidth to satisfy the traffic tunnel requirements (30 Mbps required, 20 Mbps available). As a result, the link R4-R6 is effectively excluded from the LSP path computation.

Constraint-Based Path Computation (Cont.)

Cisco.com

The headend router has two possible paths with a total cost of 40: R1 – R2 – R3 – R6 and R1 – R5 – R6, both offering at least 50 Mbps (minimum bandwidth). Because of the smaller hop count, R1 – R5 – R6 is selected.



The resulting LSP paths (after exclusion of the links that do not satisfy the traffic tunnel requirements) in the example are: R1-R2-R3-R6 and R1-R5-R6. Both paths have a total cost of 40, and the tie has to be resolved using the tiebreak rules.

First, the highest minimum bandwidth on the path is compared. After the comparison, both paths are still candidates because both can provide at least 50 Mbps of the bandwidth (the minimum bandwidth).

The next rule, the minimum number of hops on the LSP path, is applied. Because the lower path (R1-R5-R6) has a lower hop count, this path is finally selected and the constraint-based computation is concluded.

The next step toward final establishment of the LSP path for the traffic-engineered tunnel is the signaling of the path via RSVP.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **The most important link resource attribute is the maximum allocation multiplier.**
- **Traffic tunnel attributes affect how the path is set up and maintained.**
- **Affinity bits determine what policies are used during the LSP path computation.**
- **MPLS-TE link attributes are propagated with a link-state routing protocol.**
- **Dynamic constraint-based path computation is triggered by the headend of the tunnel.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-27

References

For additional information, refer to this resource:

- Search the Cisco website at <http://www.cisco.com> for more information on the topics that were covered in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which of the following is NOT a component of the maximum allocation multiplier attribute?
- A) maximum bandwidth
 - B) unreserved bandwidth
 - C) minimum available bandwidth
 - D) maximum reservable bandwidth
- Q2) The MPLS-TE tunnel attribute _____ allows the network administrator to apply path selection policies.
- Q3) Which two of the following is true regarding the tunnel resource class affinity mask? (Choose two.)
- A) If bit is 0 = do care
 - B) If bit is 1 = do care
 - C) If bit is 0 = do not care
 - D) If bit is 1 = do not care
- Q4) Which two of the following can be used to propagate MPLS-TE link attributes? (Choose two.)
- A) BGP
 - B) OSPF
 - C) EIGRP
 - D) IS-IS
- Q5) In the case of a tie after CBR has been run, which two of the following are used to break the tie? (Choose two.)
- A) larger hop count
 - B) smallest hop count
 - C) highest minimum bandwidth
 - D) highest maximum bandwidth

Lesson Answer Key

- Q1) C
Relates to: MPLS-TE Link Resource Attributes
- Q2) resource class affinity (or tunnel resource class affinity)
Relates to: MPLS-TE Tunnel Attributes
- Q3) B, C
Relates to: Implementing TE Policies with Affinity Bits
- Q4) B, D
Relates to: Propagating MPLS-TE Link Attributes with Link-State Routing
- Q5) B, C
Relates to: Constraint-Based Path Computation

Path Setup and Maintenance

Overview

This lesson describes the details of MPLS-TE tunnels, including path setup procedures and path maintenance.

Relevance

This lesson is mandatory for learners who are planning to improve usage of their network resources with MPLS-TE.

Objectives

This lesson describes LSP path setup and maintenance. Upon completing this lesson, you will be able to do the following:

- Identify MPLS-TE path setup procedures
- Identify how RSVP is used in path setup
- Identify the process that RSVP uses to assist in MPLS-TE path setup
- Identify the functions of tunnel and link admission control
- Identify why LSP path rerouting is necessary
- Identify path reoptimization
- Identify how a path is rerouted as a result of link failure

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Carrier Supporting Carrier” module of this course

Outline

This lesson includes these topics:

- Overview
- Path Setup
- RSVP Usage in Path Setup
- Hop-by-Hop Path Setup with RSVP
- Tunnel and Link Admission Control
- Path Rerouting
- Path Reoptimization
- Path Rerouting: Link Failure
- Lesson Summary
- Lesson Review

Path Setup

This topic describes how the LSP is set up.

Path Setup

Cisco.com

- **LSP path setup is initiated at the headend of a tunnel.**
- **The route (list of next-hop routers) is either:**
 - **Statically defined**
 - **Computed by CBR**
- **The route is used by RSVP to:**
 - **Assign labels**
 - **Reserve bandwidth on each link**
- **Tunnel attributes that affect path setup:**
 - **Bandwidth**
 - **Priority**
 - **Affinity attributes**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-3

LSP setup is always initiated at the traffic tunnel headend. The explicit route for the traffic tunnel is composed of a list of next-hop routers toward the tunnel endpoint (or tailend). The LSP tunnels can be statically defined or computed with CBR and thus routed away from network failures, congestion, and bottlenecks.

The explicit route is used by RSVP with TE extensions to assign labels and to reserve the bandwidth on each link.


Labels are assigned using the downstream-on-demand allocation mode.

Path setup is affected by the following tunnel attributes:

- Bandwidth
- Priority
- Affinity attributes

Path Setup (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-3-4

The figure shows a conceptual block diagram of the components of CBR and path computation. In the upper left corner is a TE control module, where the control algorithms run. The module looks at the tunnels that have been configured for CBR.

The TE control module periodically checks the CBR topology database (shown in the middle of the block diagram) to calculate the best current path from the current device to the tunnel destination.

After the path is calculated, the module passes the path off to the RSVP module to signal the circuit setup across the network. If the signaling succeeds, the signaling message eventually returns to the device, and RSVP announces back to the TE control module that the tunnel has been established.

Consequently, the TE control module tells the IGP routing module that the tunnel is available for use.

The IGP routing module includes the tunnel information in its routing table calculation and uses it to affect what routes are put into the routing table.

RSVP Usage in Path Setup

This topic describes the role of RSVP in setting up the LSP tunnels.

RSVP Usage in Path Setup

Cisco.com

RSVP makes resource reservations for both unicast and multicast applications:

- **RSVP provides support for dynamic membership changes and automatic adaptation to routing changes.**
- **RSVP sends periodic refresh messages to maintain the state along the reserved path.**
- **RSVP sessions are used between routers, not hosts.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-5

RSVP plays a significant role in path setup for LSP tunnels and supports both unicast and multicast applications.

RSVP dynamically adapts to changes either in membership (for example, multicast groups) or in the routing tables.

Additionally, RSVP transports traffic parameters and maintains the control and policy over the path. The maintenance is done by periodic refresh messages that are sent along the path to maintain the state.


In the normal usage of RSVP, the sessions are run between hosts. In TE, the RSVP sessions are run between the routers on the tunnel endpoints. The following RSVP message types are used in path setup:

- Path
- Resv
- PathTear
- ResvErr
- PathErr

When the RSVP Resv message flows back toward the sender, the intermediate nodes reserve the bandwidth and allocate the label for the tunnel. The labels are placed into the Label object of the Resv message.

Hop-by-Hop Path Setup with RSVP

This topic describes what happens at each router as an LSP tunnel is set up using RSVP.



In the example here, the LSP tunnel path setup is started by the RSVP Path message, which is initiated by the tunnel headend (Router R1). (Some of the important contents of the message are explained and monitored in the next example.)

The RSVP Path message contains several objects, including the session identification (R3-lo0, 0, R1-lo0 in the example), which uniquely identifies the tunnel. The traffic requirements for the tunnel are carried in the session attribute. The label request that is present in the message is handled by the tailend router, which allocates the respective label for the LSP path.

The explicit route object (ERO) is populated by the list of next hops, which are either manually specified or calculated by CBR (where R2-1 is used to represent the interface labeled “1” at the R2 router in the figure). The previous hop (PHOP) is set to the outgoing interface address of the router. The record route object (RRO) is populated with the same address as well.

Note The sender template is used in assigning unique LSP path identifiers (R1-lo0 = loopback interface 0, which identifies the tunnel headend; 00 = the LSP ID). The same tunnel can take two possible LSP paths (one primary and another secondary). In such a case the headend must take care of assigning unique IDs to these paths.

As the next-hop router (R2) receives the RSVP Path message, the router checks the ERO and looks into the L bit (loose) regarding the next-hop information. If this bit is set and the next hop is not on a directly connected network, the node performs a CBR calculation (PCALC) using its TE database and specifies this loose next hop as the destination.

In this way the ERO is augmented by the new results and forms a hop-by-hop path up to the next loose node specification.

Then the intermediate routers along the path (indicated in the ERO) perform the traffic tunnel admission control by inspecting the contents of the session attribute object. If the node cannot meet the requirements, it generates a PathErr message. If the requirements are met, the node is saved in the RRO.

Router R2 places the contents of the ERO into its path state block and removes itself from the ERO (R2 removes the R2-1 entry from the ERO). R2 adjusts the PHOP to the address of its own interface (the “2” interface at R2, R2-2) and adds the address (R2-2) to the RRO. The Path message is then forwarded to the next hop in the ERO.

Note	Several other functions are performed at each hop as well, including traffic admission control.

Hop-by-Hop Path Setup with RSVP (Cont.)

Cisco.com



Path State:
Session(R3-Io0, 0, R1-Io0)
PHOP(R2-2)
Label_Request(IP)
ERO ()
Session_Attribute (...)
Sender_Template(R1-Io0, 00)
Record_Route (R1-2, R2-2, R3-1)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-3-7

When the RSVP Path message arrives at the tailend router (the endpoint of the tunnel), the label request message triggers the path label allocation. The label is placed in the corresponding label object of the RSVP Resv message that is generated. The RSVP message is sent back to the headend following the reverse path recorded in the ERO and is stored at each hop in its path state block.


When the RSVP Path message arrives at the tailend router (R3), the path state block is created and the ERO becomes empty (after removing the address of the router itself from the list), which indicates that it has reached the tailend of the tunnel. The RRO at this moment contains the entire path from the headend router.

The RSVP Resv message must be generated.

The label request object in the RSVP Path message requires the tailend router to allocate a label for the specified LSP tunnel (session).

Hop-by-Hop Path Setup with RSVP (Cont.)

Cisco.com



Because R3 is the tailend router, it does not allocate a specific label for the LSP tunnel. The implicit null label is used instead (the value “POP” in the label object).

The PHOP in the RSVP Resv message is populated by the interface address of the tailend router, and this address is copied to the RRO as well.

Note The RRO is reinitiated in the RSVP Resv message.

The Resv message is forwarded to the next-hop address in the path state block of the tailend router. The next-hop information in the path state block was established when the Path message was traveling in the opposite direction (headend to tailend).

The RSVP Resv message travels back to the headend router. On each hop (in addition to the admission control itself) label handling is performed. As you can see from the RSVP Resv message that is shown in the figure, the following actions were performed at the intermediate hop (R2):

- The interface address of R2 was put into the PHOP field and added to the beginning of the RRO list.
- The incoming label (5) was allocated for the specified LSP path.

Note The label switch table is not shown but contains the information for label switching. In this particular case the label “5” is replaced with an implicit null label (“POP”).

The Resv message is forwarded toward the next hop that is listed in the path state block of the router. The next-hop information in the path state block was established when the Path message was traveling in the opposite direction (headend to tailend).

Hop-by-Hop Path Setup with RSVP (Cont.)

Cisco.com



Resv state:
Session(R3-Io0, 0, R1-Io0)
PHOP(R2-1)
Sender_Template(R1-Io0, 00)
Label=5
Record_Route(R1-2, R2-1, R3-1)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-3-9

When the RSVP Resv message arrives at the headend router (R1), the LSP path setup is concluded. The label (5) that is allocated by the next-hop router toward the tunnel endpoint (PHOP = R2-1) is stored, and the explicit route that was taken by the tunnel is now present in the RRO. The LSP tunnel is established.

Tunnel and Link Admission Control

This topic describes how tunnel and link admission control are performed at each hop during tunnel setup.

Tunnel and Link Admission Control

Cisco.com

Invoked by RSVP Path message:

- **Determines if resources are available**
- **If bandwidth is not available:**
 - **Link-level call admission control (LCAC) says no to RSVP**
 - **PathErr message is sent**
- **If bandwidth is available, this bandwidth is put aside in a waiting pool (waiting for the Resv message):**
 - **Triggers IGP information distribution when resource thresholds are crossed**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-10

One of the essential steps that is performed at each hop of the route to the LSP tunnel endpoint (the tunnel) is admission control, which is invoked by the RSVP Path message traveling from the headend to the tailend router.

Each hop on the way determines if the available resources that are specified in the session attribute object are available. Two situations can appear:

- There is not enough bandwidth on a specified link through which the traffic tunnel should be established. The link-level call admission control (LCAC) module informs RSVP about the lack of resources, and RSVP generates an RSVP PathErr message with the code “Requested bandwidth unavailable.” Additionally, the flooding of the node resource information (by the respective link-state IGP) can be triggered as well.
- If the requested bandwidth is available, the bandwidth is reserved and is put into a waiting pool for the Resv message to confirm the reservation. Additionally, if the resource threshold is exceeded, the respective IGP triggers the flooding of the resource information.

Tunnel and Link Admission Control (Cont.)

Cisco.com

Pre-emption

- The process of LSP path setup may require the pre-emption of resources.
- LCAC notifies RSVP of the pre-emption.
- RSVP sends PathErr or ResvErr or both for the pre-empted tunnel.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-11

During admission control, the priorities are checked as well.

If the requested bandwidth is available, but is in use by lower-priority sessions, then the lower-priority sessions (beginning with the lowest priority) may be pre-empted to free the necessary bandwidth. There are eight levels of priority, 0 being the highest, and 7 being the lowest.

When pre-emption is supported, each pre-empted reservation triggers a ResvErr or a PathErr message or both with the code “Policy control failure.”

Path Rerouting

This topic describes why LSP path rerouting is necessary.

Path Rerouting

Cisco.com

Path rerouting may result from:

- Reoptimization due to a change in network resources
- Link failures that affect the LSP path

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-12

Because the tunnel is not linked to the LSP that is carrying it, the actual path can dynamically change without affecting the tunnel.

Path rerouting may result from either of the following two circumstances:

- Reoptimization due to a change in network resources
- Link failures that affect the LSP path

Path Reoptimization

This topic describes path reoptimization.

Path Reoptimization

Cisco.com

- **Problem:** Some resources become available, which results in a nonoptimal path of traffic tunnels
- **Solution: Reoptimization:**
 - A periodic timer checks for the most optimal path
 - If a better LSP seems to be available:
 - The device attempts to signal the better LSP
 - If successful, replaces the old and inferior LSP with the new and better LSP

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-13

The LSP path must be rerouted when there are physical (topology) failures or when certain changes in resource usage require it. As resources in another part of the network become available, the traffic tunnels may have to be reoptimized.


The reoptimization is done on a periodic basis. At certain intervals, a check for the most optimal paths for LSP tunnels is performed and if the current path is not the most optimal, tunnel rerouting is initiated.

The device (headend router) first attempts to signal a better LSP, and only after the new LSP path setup has been established successfully will the traffic be rerouted from the former tunnel to the new one.

Path Reoptimization (Cont.)

Cisco.com

Nondisruptive Rerouting — Reoptimization



- Current Path (ERO = R1 -> R2 -> R6 -> R7 -> R4 -> R9).
- New Path (ERO = R1 -> R2 -> R3 -> R4 -> R9)—shared with current path and reserved for both paths.
- ← Until R9 gets new Path message, current Resv is refreshed—PathTear can then be sent to remove old path (and release resources).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-14

The figure shows how the nondisruptive rerouting of the traffic tunnel is performed. Initially, the ERO lists the LSP path R1-R2-R6-R7-R4-R9, with R1 as the headend and R9 as the tailend of the tunnel.

The changes in available bandwidth on the link R2-R3 dictate that the LSP path be reoptimized. The new path R1-R2-R3-R4-R9 is signaled, and parts of the path overlap with the existing path. Still, the current LSP path is used.

Note On links that are common to the current and new LSP, resources that are used by the current LSP tunnel should not be released before traffic is moved to the new LSP tunnel, and reservations should not be counted twice (doing so might cause the admission control to reject the new LSP tunnel).

After the new LSP path is successfully established, the traffic is rerouted to the new path and the reserved resources of the previous path are released. The release is done by the tailend router, which initiates an RSVP PathTear message.

The labels that are allocated during the RSVP path setup are shown as well. The tailend router assigns the implicit null (POP) label.

Path Rerouting: Link Failure

This topic explains path rerouting due to a link failure.

Path Rerouting: Link Failure

Cisco.com

The Goal

- **Repair at the headend of the tunnel in the event of failure of an existing LSP:**
 - IGP or RSVP alarms the headend.
- **New path for LSP is computed, and eventually a new LSP is signaled.**
- **Tunnel interface goes down if there is no LSP available for 10 sec.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-15

When a link, that makes up the certain traffic tunnel, fails, the headend of the tunnel detects the failure in one of two ways:

- The IGP (OSPF or IS-IS) sends a new link-state packet with information about changes that have happened.
- RSVP alarms the failure by sending an RSVP PathTear message to the headend.

Link failure detection, without any preconfigured or precomputed path at the headend, results in a new path calculation (using a modified SPF algorithm) and consequently in a new LSP path setup.

Note	The tunnel interface that is used for the specified traffic tunnel (LSP path) goes down if the specified LSP path is not available for 10 sec. In the meantime, the traffic that is intended for the tunnel continues using a broken LSP path, which results in black hole routing.

Path Rerouting: Link Failure (Cont.)

Cisco.com

Link failure – What happens:

(Example: One link along a dynamic tunnel LSP path goes down.)

- RSVP PathTear causes the headend to flag LSP as dead
- RSVP session is cleared
- PCALC triggered:
 - No alternative path:
 - Headend sets the tunnel down
 - Alternative path found:
 - New LSP directly signaled
 - Adjacency table updated for the tunnel interface
 - CEF table updated for all entries that resolve on this tunnel adjacency

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-16

When the router along the dynamic LSP path detects a link failure, it sends an RSVP PathTear message to the headend.

This message signals to the headend that the tunnel is down. The headend clears the RSVP session, and a new PCALC is triggered using a modified SPF algorithm.

There are two possible outcomes of the PCALC calculation:

- No new path is found. The headend sets the tunnel interface down.
- An alternative path is found. The new LSP path setup is triggered by RSVP signaling, and adjacency tables for the tunnel interface are updated. Also, the CEF table is updated for all the entries that are related to this tunnel adjacency.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **LSP path setup is initiated at the headend of a tunnel.**
- **LSP tunnel path setup is started by the RSVP Path message.**
- **Admission control is invoked by the RSVP Path message.**
- **After LSP path establishment, the path is constantly monitored.**
- **LSP paths are checked at intervals for the most optimal paths. If the current path is not the most optimal tunnel, rerouting is initiated.**
- **Path reoptimization solves the problem of configuring two tunnels and reserving resources twice.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-17

References

For additional information, refer to this resource:

- Search the Cisco.com website at <http://www.cisco.com> for more information on the topics that were covered in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which three of the following affect how a path is set up? (Choose three.)
- A) priority
 - B) bandwidth
 - C) affinity attributes
 - D) MPLS label stack
- Q2) Which of the following statements is correct as related to the RSVP sessions in an MPLS-TE application?
- A) The sessions are run between the routers at the tunnel endpoints.
 - B) The sessions are run between hosts.
 - C) Two set of sessions are run, one between the hosts and their tunnel end routers another between the tunnel end routers.
- Q3) Admission control is invoked by the _____ message.
- Q4) In case of a network failure, traffic tunnels rerouted by which of the following?
- A) the headend router
 - B) the upstream router that is nearest to the point of failure
 - C) the downstream router that is nearest to the point of failure
 - D) the tunnel end router
- Q5) During path reoptimization, which router first attempts to identify a better LSP?
- A) the headend router
 - B) any router that has identified that it has new resources available
 - C) the tunnel end router
- Q6) Which of the following solves the problem of setting up two tunnels and having the resources counted twice?
- A) path reuse
 - B) path monitoring
 - C) path rerouting
 - D) path reoptimization

Lesson Answer Key

- Q1) A, B, C
Relates to: Path Setup
- Q2) A
Relates to: RSVP Usage in Path Setup
- Q3) Path
Relates to: Tunnel and Link Admission Control
- Q4) A
Relates to: Path Rerouting: Link Failure
- Q5) headend
Relates to: Path Reoptimization
- Q6) D
Relates to: Path Reoptimization

Assigning Traffic to Traffic Tunnels

Overview

This lesson describes several well-known drawbacks of traditional IP routing: full-mesh requirements in the IP-over-ATM overlay model and lack of traffic engineering (TE).

Relevance

This lesson is mandatory for learners who are planning to improve usage of their network resources with MPLS-TE.

Objectives

This lesson describes how to assign traffic to traffic tunnels. Upon completing this lesson, you will be able to do the following:

- Identify the traffic flow modifications that are available with static routes and policy-based routing
- Identify how tunnel access is implemented using static routing pointing to a tunnel
- Identify how MPLS-TE solves the BGP next-hop problem
- Configure autoroute to modify the IP forwarding database
- Identify the autoroute path selection rules
- Identify what happens in certain autoroute situations
- Identify the function of forwarding adjacency as it relates to MPLS-TE

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “Carrier Supporting Carrier” module of this course

Outline

This lesson includes these topics:

- Overview
- Traffic Flow Modifications
- Static Routing Pointing to a Tunnel
- Solving the BGP Next-Hop Problem
- IP Forwarding Database Modification with Autoroute
- Autoroute Path Selection Rules
- Autoroute Examples
- Forwarding Adjacency
- Lesson Summary
- Lesson Review

Traffic Flow Modifications

This topic describes the traffic flow modifications that are available with static routes and PBR.

Traffic Flow Modifications

Cisco.com

- CBR used to find the path for an LSP tunnel.
- IP routing does not see internal details.
- Tunnels can be used for routing only if they are explicitly specified:
 - Static route in the IP routing table points to a selected LSP tunnel interface.
 - Advertise the tunnel to IP using autoroute.
 - Policy routing—the next-hop interface is an LSP tunnel.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-3


The LSP is computed by CBR, which takes the resource requirements into consideration as well. When the LSP path is established for the tunnel, the traffic can flow across it. From the IP perspective, an LSP path is a simple tunnel.

These engineered tunnels can be used for IP routing only if the tunnels are explicitly specified for routing in one of two ways:

- Via static routes that point to the tunnel
- Via policy routing that sets a next-hop interface to the tunnel

Static Routing Pointing to a Tunnel

This topic describes how to implement static routing that points to a tunnel.



The example topology here shows two engineered tunnels: T1 (between R1 and R4) and T2 (between R1 and R5). The loopback addresses on each router are in the form i.i.i.i where “i” is the router number (for example, the R5 loopback address is 5.5.5.5). The metric on each of the interfaces is set to 1.

R1 has two physical interfaces: I1 and I2, and two neighboring routers (next hops) with addresses of A1 and A2, respectively.

The routing table lists all eight loopback routes and associated information. Only the statically configured destinations (R4 and R5) list tunnels as their outgoing interfaces. For all other destinations the normal IGP routing is used and results in physical interfaces (along with next hops) as the outgoing interfaces towards these destinations. The metric to the destination is the normal IGP metric.

Note Even for the destination that is behind each of the tunnel endpoints (R8), the normal IGP routing is performed if there is no static route to the traffic-engineered tunnel.

The SPF algorithm calculates paths to destinations in its usual way with the exception of the paths for the tunnels where a constraint-based computation is performed.

Solving the BGP Next-Hop Problem

This topic describes how MPLS-TE solves the Border Gateway Protocol (BGP) next-hop problem.

Solving the BGP Next-Hop Problem

Cisco.com

Problem: Assume that BGP next hop is not set to *self* on R2 and R3; the BGP local preference cannot be used to prefer R2 over R3 for exit.

```
graph LR
    R1((R1)) -- T1 --> R2((R2))
    R1 -- T2 --> R3((R3))
    R2 --- R4((R4))
    R3 --- R4
    subgraph Cloud
        R4
    end
```

- **Solution 1:** Tunnel T1 with relative IP MPLS-TE metric -6 and static route for next hop on T1.
- **Solution 2:** Tunnel T1 with absolute IP MPLS-TE metric 29 and static route for next hop on T1.
- **All traffic from R1 to the Ethernet will go via R2.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-5

There are several interesting problems that are associated with MPLS-TE routing. In the example in the figure, there are two internal routers (R2 and R3) serving as exits to external BGP destinations on R4. These two routers are also the endpoints of two MPLS-TE tunnels. The BGP **next-hop-self** is not configured on R2 and R3 (R4 remains the BGP next hop), which prevents the use of local preference to prefer R2 to R3 as an exit to external destinations. The best IGP metric is used instead, and R3 becomes the preferred exit.

There are two possible solutions to the problem. Both solutions use the static routes to MPLS-TE tunnels and some metric modifications (relative or absolute metrics):

- The metric for tunnel T1 is set to a relative value decreased by 6 ($35-6=29$), and the static route for the Ethernet is configured to the tunnel T1
- The metric for tunnel T1 is set to an absolute value of 29, and the static route for the Ethernet is configured to the tunnel T1

In both cases all the traffic for the Ethernet will flow via R2 because its metric is less than 30.

IP Forwarding Database Modification with Autoroute

This topic describes how the IP forwarding database is modified when the autoroute feature is implemented.

IP Forwarding Database Modification with Autoroute

Cisco.com

- **Autoroute feature enables the headend to see the LSP as a directly connected interface:**
 - Only for the SPF route determination, not for the constraint-based path computation.
 - All traffic directed to prefixes topologically behind the tunnel endpoint (tailend) is forwarded onto the tunnel.
- **Autoroute affects the headend only; other routers on the LSP path do not see the tunnel.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—3-6

To overcome the problems that result from static routing configurations onto MPLS-TE tunnels, the autoroute feature of Cisco IOS software was introduced.

The autoroute feature enables the headend router to see the MPLS-TE tunnel as a directly connected interface and to use it in its modified SPF computations.

The MPLS-TE tunnel is used only for normal IGP route calculation (at the headend only) and is not included in any constraint-based path computation.

The autoroute feature enables all the prefixes that are topologically behind the MPLS-TE tunnel endpoint (tailend) to be reachable via the tunnel itself (unlike with static routing, where only statically configured destinations are reachable via the tunnel).

The autoroute feature affects the headend router only and has no effect on intermediate routers. These routers still use normal IGP routing for all the destinations.

Autoroute Path Selection Rules

This topic describes the autoroute path selection rules.

Autoroute Path Selection Rules

Cisco.com

- **The cost of the TE tunnel is equal to the shortest IGP metric to the tunnel endpoint; the metric is tunable.**
- **If the tunnel metric is:**
 - **Equal or lower than the native IGP metric, the tunnel replaces existing next hops; otherwise, the tunnel is not considered for routing**
 - **Equal to other TE tunnels, the tunnel is added to the existing next hops (parallel paths)**
- **Tunnels can be load-balanced (CEF mechanism)—tunnel bandwidth factor considered.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-7


Because the autoroute feature includes the MPLS-TE tunnel into the modified SPF path calculation, the metric of the tunnel plays a significant role. The cost of the tunnel is equal to the best IGP metric to the tunnel endpoint regardless of the LSP path. The tunnel metric is tunable using either relative or absolute metrics.

During installation of the best paths to the destination, the tunnel metric is compared to other existing tunnel metrics and to all the native IGP path metrics. The lower metric is better and if the MPLS-TE tunnel has an equal or lower metric than the native IGP metric, it is installed as a next hop to the respective destinations.

If there are tunnels with equal metrics, they are installed in the routing table and provide for load balancing. The load balancing is done proportionally to the configured bandwidth of the tunnel.

Autoroute Examples

This topic describes some examples of using the autoroute feature.




The example topology shows two engineered tunnels: T1 (between R1 and R4) and T2 (between R1 and R5). The loopback addresses on each router are in the form i.i.i.i where “i” is the router number (for example, the R5 loopback address is 5.5.5.5). The metric on each of the interfaces is set to 10.

R1 has two physical interfaces, I1 and I2, and two neighboring routers (next hops) with addresses of A1 and A2, respectively.

Autoroute – Default Metric (Cont.)


Cisco.com



This example shows the resulting shortest-path tree from router R1. In this situation the tunnels are seen for routing purposes only by the headend (A). Intermediate routers do not see the tunnel nor do they take it into consideration for route calculations.

Autoroute – Default Metric (Cont.)

Cisco.com



The routing table shows all destinations at the endpoint of the tunnel and behind it (R8) as reachable via the tunnel itself. The metric to the destination is the normal IGP metric.


The LSP path for T1 follows the path R1-R2-R3-R4 and the tunnel cost is the best IGP metric (30) to the tunnel endpoint. The metric to R8 is 40 (T1 plus one hop).

The LSP path for T2 follows the path R1-R6-R7-R4-R5. Although the LSP path passes through the R7-R4 link, the overall metric of the tunnel is 40 (the sum of metrics on the best IGP path R1-R2-R3-R4-R5).

In the routing table all the networks that are topologically behind the tunnel endpoint are reachable via the tunnel itself. Because, by default, the MPLS-TE tunnel metric is equal to the native IGP metric, the tunnel is installed as a next hop to the respective destinations. This is the effect of the autoroute feature.

Autoroute – Absolute Metric

Cisco.com




Depending on the ability of the IGP to support load sharing, the native IP path may also show up in the routing table as a second path option. In this example, there appear to be two paths to R4, while there is only one physical path. In the case of R5 there appear to be three paths, two of which do not follow the desired tunnel path.

The tunnel metrics can be tuned, and either relative or absolute metrics can be used to resolve this issue.

Autoroute – Relative and Absolute Metric

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-12

In this example, the relative metric is used to control path selection. T1 is given a value of the native IPG value -2 (28), which gives it preference to the native IP path. When the tunnel is considered in the IGP calculation, the native IGP metric (30) is greater than the tunnel metric (28) for all the destinations that are topologically behind the tunnel endpoint. As a result, all the destination networks are reachable via the TE tunnel.

T2 could have been given the native IPG value -4 (36), giving it preference to the native IP path and the T1-R4 path. However, in this example, it was chosen to give it an absolute value of 35.

Forwarding Adjacency

This topic describes the function of forwarding adjacency as it relates to MPLS-TE.

Forwarding Adjacency

Cisco.com

- **Mechanism for:**
 - **Better intra- and inter-POP load balancing**
 - **Tunnel sizing independent of inner topology**
- **Allows the announcement of established tunnel via link-state (LSP) announcements**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—3-13

The MPLS-TE forwarding adjacency feature allows a network administrator to handle a traffic-engineered LSP tunnel as a link in an IGP network based on the SPF algorithm.

A forwarding adjacency can be created between routers regardless of their location in the network.

Forwarding adjacency is a mechanism to allow the announcement of established tunnels via IGP to all nodes within an area.


By using forwarding adjacency, you can achieve the following goals:

- Better load balancing when you are creating POP-to-POP tunnels
- Use of tunnels from any upstream node independent of the inner topology of the network
- Use of tunnels independent of topology changes within the tunneled network area

Forwarding Adjacency (Cont.)

Cisco.com

Traffic flow without Forwarding Adjacency



- Tunnels created and announced to IP with autoroute with equal cost to load-balance.
- All the POP-to-POP traffic exits via the routers on the IGP shortest path:
 - No load balancing
 - All traffic flows on tunnel: A → B → D → F

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-14

Before looking into the real benefits of forwarding agency, you need to clearly see the limitations of autoroute in certain network topologies.

In this example, established tunnels exist from B to D, from B to E, from C to E, and from C to D; the preferred tunnels are B to D and C to E.

- The path metric from B to D for IS-IS is 30 (assuming default metric).
- The path metric from C to E is 20.

But traffic is entering at router A. Router A has no knowledge about the existence of tunnels between B and D and C and E. All it has is its IGP information, indicating that the better path to F leads via router B and D.

The results are as follows:

- There will be no load balancing.
- All traffic will flow via B and D.


Any change in the core topology will affect the path metric and thus any load balancing for POP-to-POP traffic.

Note You can theoretically prevent this problem by creating tunnels from any router to any other router, but this design does not scale in medium and big networks.

Forwarding Adjacency (Cont.)

Cisco.com

Traffic flow without Forwarding Adjacency



- All the POP-to-POP traffic exits via the routers on the IGP shortest path.
- Change in the core topology does affect the load balancing in the POP:
 - Normal state: All traffic flows $A \rightarrow B \rightarrow D \rightarrow F$
 - Link failure: All traffic flows $A \rightarrow C \rightarrow E \rightarrow F$

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-15

What happens in case a link on the path between B and D gets broken?


Even though rerouting will possibly take place, the IGP metric for the complete path A – B – intermediates – D – F will change.

In the example here, the change in the metric will result in a possibly unplanned switchover of the traffic from A to F from the upper to the lower path.

This switchover may result in a possible congestion on the path from C to E, whereas the protected path from B to D gets idled out.

Forwarding Adjacency (Cont.)

Cisco.com



POP to POP traffic is better load balanced:

- In the POP: The two core routers are used
- In the core: At least, two tunnels are used
- As long as the IGP metric for a path with the FA (e.g. 25) is shorter than the FA-free path (e.g., 30)

Inner Topology does not affect Tunnel Sizing:

- Change in the core topology does not affect the load balancing in the POP

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-16

By using forwarding adjacency, you can create POP-to-POP tunnels where traffic paths and load balancing can be designed independent of the inner (core) topology of the network and independent of link failures.

To advertise a TE tunnel as a link in an IGP network, use the **tunnel mpls traffic-eng forwarding-adjacency** command in interface configuration mode.

tunnel mpls traffic-eng forwarding-adjacency {holdtime value}

Syntax Description

Parameter	Description
holdtime <i>value</i>	(Optional) Time in milliseconds (ms) that a TE tunnel waits after going down before informing the network. The range is 0 to 4,294,967,295 ms. The default value is 0.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- TE tunnels can be used for IP routing only if the tunnels are explicitly specified for routing.
- SPF calculates paths to destinations except for paths to tunnels where constraint-based computation is performed.
- Static routes to MPLS-TE tunnels (and some metric modifications) can be used to solve the BGP next-hop issue.
- The **autoroute** feature enables headend routers to see MPLS-TE tunnels as a directly connected interface.
- The cost of the TE tunnel is equal to the shortest IGP metric to tunnel endpoint (metric is tunable).
- Use **autoroute** for destination network directly connected to the tunnel endpoints.
- The **MPLS-TE forwarding adjacency** feature allows a network administrator to handle a traffic engineering, label switched path (LSP) tunnel as a link in an Interior Gateway Protocol (IGP) network.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—3-17

References

For additional information, refer to this resource:

- Search the Cisco.com website at <http://www.cisco.com> for more information on the topics that were covered in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The LSP is calculated using?
- A) CBR
 - B) DUAL algorithm
 - C) SPF algorithm
 - D) No calculation is used for the establishment of an LSP.
- Q2) Statically configured destinations list their tunnels in the routing table as what?
- A) incoming interfaces
 - B) loopback interfaces
 - C) outgoing interfaces
 - D) Statically configured tunnels are not listed in the routing table.
- Q3) Which two terms are used when the IGP metric is modified? (Choose two.)
- A) absolute
 - B) negative
 - C) positive
 - D) relative
- Q4) The feature that enables head-end routers to see the MPLS-TE tunnel as a directly connected interface is called _____.
- Q5) The cost of the TE tunnel is equal to the shortest _____ to the tunnel endpoint.
- Q6) The IGP metric is 50; the tunnel metric has been set to a relative +2. Which path will be used for routing? (Each path contains six routers.)
- A) IGP path
 - B) LSP path
 - C) Both paths will be used, and load balancing will be implemented.
 - D) You cannot tell from the information that has been given.
- Q7) The mechanism that enables the announcement of established tunnels via IGP to all nodes within an area is called _____.

Lesson Answer Key

- Q1) A
Relates to: Traffic Flow Modifications
- Q2) C
Relates to: Static Routing Pointing to a Tunnel
- Q3) A, D
Relates to: Solving the BGP Next-Hop Problem
- Q4) autoroute
Relates to: IP Forwarding Database Modification with Autoroute
- Q5) IGP metric
Relates to: Autoroute Path Selection Rules
- Q6) A
Relates to: Autoroute Examples
- Q7) forwarding adjacency
Relates to: Forwarding Adjacency

Module Assessment

Overview

Use this assessment to test what you learned in this module. The correct answers are found in the Module Assessment Answer Key.

Quiz: MPLS Traffic Engineering Technology

Complete the quiz to assess what you have learned in this module.

Objectives

This activity tests your knowledge on how to meet these objectives:

- Identify the basic TE concepts
- Identify the basic MPLS-TE concepts
- Describe the function of constraint-based path computation
- Describe the operations of LSP path setup and maintenance
- Identify how to assign traffic to traffic trunks

Instructions

Complete these steps:

- Step 1** Answer all questions in this quiz by selecting the best answer(s) to each question.
- Step 2** Verify your results against the answer key at the end of this quiz.
- Step 3** Review the related lesson for each question that you answered incorrectly.

Quiz

Answer these questions:

- Q1) Which two situations can result in network congestion? (Choose two.)
- A) The least-cost network path is being used to route traffic.
 - B) Traffic streams are efficiently mapped onto available resources.
 - C) Network resources are insufficient to accommodate offered load.
 - D) Traffic streams are inefficiently mapped onto available resources.
- Q2) When you are using MPLS-TE, a minimum of _____ labels is assigned to a packet at the ingress router.
- Q3) The constraint-based path computation uses which algorithm?
- A) DUAL algorithm
 - B) modified Dijkstra's algorithm
 - C) modified Bell-Howell algorithm
 - D) none of the above
- Q4) Which two of the following is true regarding the tunnel resource class affinity mask? (Choose two.)
- A) If bit is 0 = do care
 - B) If bit is 1 = do care
 - C) If bit is 0 = do not care
 - D) If bit is 1 = do not care

- Q5) Which two of the following can be used to propagate MPLS-TE link attributes?
(Choose two.)
- A) BGP
 - B) OSPF
 - C) EIGRP
 - D) IS-IS
- Q6) The feature that enables headend routers to see the MPLS-TE tunnel as a directly connected interface is called _____.
- Q7) The cost of the TE tunnel is equal to the shortest _____ to the tunnel endpoint.

Scoring

You have successfully completed the quiz for this lesson when you earn a score of 80 percent or better.

Module Assessment Answer Key

Q1) C, D

Relates to: Traffic Engineering Concepts

Q2) two

Relates to: MPLS Traffic Engineering Concepts

Q3) B

Relates to: MPLS Traffic Engineering Components

Q4) B, C

Relates to: Constraint-Based Path Computation

Q5) B, D

Relates to: Constraint-Based Path Computation

Q6) autoroute

Relates to: Assigning Traffic to Traffic Tunnels

Q7) IGP metric

Relates to: Assigning Traffic to Traffic Tunnels

Configuring MPLS Traffic Engineering

Overview

This module discusses the requirement for implementing MPLS Traffic Engineering (MPLS-TE) in modern networks. It includes configuration of routers to enable basic traffic tunnels, assignment of traffic to a tunnel, control of path selection, and performance of tunnel protection and tunnel maintenance.

Objectives

Upon completing this module, you will be able to describe the tasks and commands that are necessary to implement MPLS-TE. This includes being able to do the following:

- Describe how to configure MPLS-TE on Cisco IOS platforms
- Describe advanced MPLS-TE path selection
- Describe advanced MPLS-TE link protection
- Describe advanced MPLS-TE bandwidth control
- Describe MPLS-TE multiarea support
- Monitor and troubleshoot MPLS-TE implementations

Outline

The module contains these lessons:

- Configuring MPLS Traffic Engineering on Cisco IOS Platforms
- Advanced MPLS-TE Path Selection
- Advanced MPLS-TE Link Protection
- Advanced MPLS-TE Bandwidth Control
- MPLS-TE Interarea Support
- Monitoring and Troubleshooting MPLS-TE

Configuring MPLS Traffic Engineering on Cisco IOS Platforms

Overview

This lesson describes MPLS-TE commands for the implementation of Multiprotocol Label Switching (MPLS) traffic tunnels. The configuration commands that are needed to support MPLS-TE are explained, and sample setups are presented.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes how to configure MPLS-TE on Cisco IOS platforms. Upon completing this lesson, you will be able to do the following:

- Identify the tasks that are required to implement MPLS-TE
- Enable device-level MPLS-TE support
- Enable MPLS-TE on an interface
- Enable MPLS-TE support in an IS-IS environment
- Enable MPLS-TE support in an OSPF environment
- Create and configure a traffic tunnel
- Enable traffic tunnels with autoroute or forwarding adjacency

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course

Outline

This lesson includes these topics:

- Overview
- MPLS-TE Configuration Road Map
- Enabling Device-Level MPLS-TE Support
- Enabling MPLS-TE on an Interface
- Enabling MPLS-TE Support in IS-IS
- Enabling MPLS-TE Support in OSPF
- Creating and Configuring a Traffic Tunnel
- Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency
- Implementing MPLS-TE
- Lesson Summary
- Lesson Review

MPLS-TE Configuration Road Map

This topic describes the tasks that are required to implement MPLS-TE on a Cisco IOS device.

MPLS-TE Configuration Road Map

Cisco.com

- **Globally (device)**
 - **ip cef (distributed)**
 - **mpls traffic-eng tunnels**
- **Interface**
 - **mpls traffic-eng tunnels**
- **MPLS support in the IGP**
 - **OSPF**
 - **IS-IS**
- **Create and configure traffic tunnels**
 - **Tunnel configuration commands**
- **Traffic tunnels with autoroute and forwarding adjacency**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-3

The figure here describes the minimum tasks that are required to implement MPLS-TE on a Cisco IOS router.

This lesson will cover each of these tasks in detail.

The command **ip cef** has to be enabled globally on the device.

MPLS-TE has to be supported in the Interior Gateway Protocol (IGP). Currently, only Open Shortest Path First (OSPF) and Intermediate System-to-Intermediate System (IS-IS) have MPLS-TE support available.

Enabling Device-Level MPLS-TE Support

This topic describes the command syntax that is required to enable device-level MPLS-TE support on a Cisco IOS device.

Enabling Device-Level MPLS-TE Support

Cisco.com

```
router(config)#
```

```
ip cef
```

- Enables IP CEF switching

```
router(config)#
```

```
mpls traffic-eng tunnels
```

- Enables the MPLS Traffic Engineering tunnel feature on a device

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-4

ip cef

To enable Cisco Express Forwarding (CEF) on the route processor card, use the **ip cef** command in global configuration mode. To disable CEF, use the **no** form of this command.

ip cef [distributed]

Syntax Description

Parameter	Description
distributed	(Optional) Enables distributed CEF (dCEF) operation. Distributes CEF information to line cards. Line cards perform express forwarding.

Usage Guidelines

CEF is an advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with dynamic, topologically dispersed traffic patterns, such as those associated with web-based applications and interactive sessions.

mpls traffic-eng tunnels (configuration)

To enable MPLS-TE tunnel signaling on a device, use the **mpls traffic-eng tunnels** command in global configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng tunnels

Usage Guidelines

Enables the MPLS-TE feature on a device. To use the feature, MPLS-TE must also be enabled on the desired interfaces.

Enabling MPLS-TE on an Interface

This topic describes the command syntax that is required to enable MPLS-TE on an interface on a Cisco IOS device.

Enabling MPLS-TE on an Interface

Cisco.com

```
router(config-if)#  
mpls ip
```

- Enable the MPLS on an interface

```
router(config-if)#  
mpls traffic-eng tunnels
```

- Enable the MPLS traffic engineering tunnel feature on an interface

```
router(config-if)#  
ip rsvp bandwidth [total-flow-kbps] [single-flow-kbps]
```

- Enable RSVP for IP on an interface and specify amount of bandwidth to be reserved

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-5

mpls ip

To enable MPLS forwarding of IP version 4 (IPv4) packets along normally routed paths for a particular interface, use the **mpls ip** command in interface configuration mode. To disable this feature, use the **no** form of this command.

Usage Guidelines

MPLS forwarding of IPv4 packets along normally routed paths is sometimes called “dynamic label switching.” Dynamic label switching has been enabled for the platform when this command is issued on an interface, and label distribution begins with the periodic transmission of neighbor discovery Hello messages. When the outgoing label for a destination that is routed out the interface is known, packets for the destination are labeled with that outgoing label and forwarded out the interface.

mpls traffic-eng tunnels (interface)

To enable MPLS-TE tunnel signaling on an interface, assuming that it is enabled for the device, use the **mpls traffic-eng tunnels** command in interface configuration mode.

Usage Guidelines

Enables the MPLS-TE feature on the interface. To use the feature, MPLS-TE must also be enabled on the device globally. An enabled interface has its resource information flooded into the appropriate IGP link-state database, and accepts traffic engineering (TE) tunnel signaling requests.

ip rsvp bandwidth

To enable Resource Reservation Protocol (RSVP) for IP on an interface, use the **ip rsvp bandwidth** interface configuration command. To disable RSVP, use the **no** form of this command.

ip rsvp bandwidth [*interface-kbps* [*single-flow-kbps*]]

Syntax Description

Parameter	Description
<i>interface-kbps</i>	(Optional) Maximum amount of bandwidth, in kbps, that may be allocated by RSVP flows. The range is from 1 to 10,000,000.
<i>single-flow-kbps</i>	(Optional) Maximum amount of bandwidth, in kbps, that may be allocated to a single flow. The range is from 1 to 10,000,000.

Usage Guidelines

RSVP is disabled by default to allow backward compatibility with systems that do not implement RSVP. Weighted random early detection (WRED) or fair queuing must be enabled first.

RSVP cannot be configured with Versatile Interface Processor (VIP) dCEF.

Enabling MPLS-TE on an Interface (Cont.)

Cisco.com

```
router(config-if)#
```

```
mpls traffic-eng attribute-flags attributes
```

- Assigns attributes that will be compared to the affinity bits of a tunnel during selection of a path

```
router(config-if)#
```

```
mpls traffic-eng flooding thresholds { down | up }  
percent [ percent ...]
```

- Sets the reserved bandwidth thresholds of a link on an interface

```
router(config-if)#
```

```
mpls traffic-eng link-management timers periodic-  
flooding interval
```

- Sets the length of the interval for periodic flooding

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-6

mpls traffic-eng attribute-flags

To set the user-specified attribute flags for the interface, use the **mpls traffic-eng attribute-flags** command in interface configuration mode. The interface is flooded globally so that it can be used as a tunnel headend path selection criterion. To disable this feature, use the **no** form of this command.

mpls traffic-eng attribute-flags *0x0-0xFFFFFFFF*

Syntax Description

Parameter	Description
<i>0x0-0xFFFFFFFF</i>	Represents 32 bits. This mask is compared with the affinity bits of a tunnel during dynamic path selection.

Usage Guidelines

The purpose of this command is to assign attributes to a link in order to cause tunnels with matching attributes (as represented by their affinity bits) to prefer this link to others that do not match.

mpls traffic-eng flooding thresholds

To set the reserved bandwidth thresholds of a link, use the **mpls traffic-eng flooding thresholds** command in interface configuration mode. If a bandwidth threshold is crossed, the bandwidth information for the link is immediately flooded throughout the network. To return to the default settings, use the **no** form of this command.

mpls traffic-eng flooding thresholds {down | up} percent [percent...]

Syntax Description

Parameter	Description
down	Sets the thresholds for decreased resource availability. The range is from 0 to 99 percent.
up	Sets the thresholds for increased resource availability. The range is from 1 to 100 percent.
<i>percent [percent...]</i>	Specifies the bandwidth threshold level.

Defaults

The default for **down** is

100, 99, 98, 97, 96, 95, 90, 85, 80, 75, 60, 45, 30, 15.

The default for **up** is

15, 30, 45, 60, 75, 80, 85, 90, 95, 97, 98, 99, 100.

Usage Guidelines

When a threshold is crossed, MPLS-TE link management advertises updated link information. Similarly, if no thresholds are crossed, changes may be flooded periodically unless periodic flooding has been disabled.

mpls traffic-eng link-management timers periodic-flooding

To set the length of the interval that is used for periodic flooding, use the **mpls traffic-eng link timers periodic-flooding** command in global configuration mode. To disable this feature, use the **no** form of this command.

- **mpls traffic-eng link-management timers periodic-flooding** *interval*
- **no mpls traffic-eng link-management timers periodic-flooding**

Syntax Description

Parameter	Description
<i>interval</i>	Length of interval that is used for periodic flooding (in seconds). The range is from 0 to 3600. If you set this value to 0, you turn off periodic flooding. If you set this value anywhere in the range from 1 to 29, it is treated as 3 (default is 3 minutes).

Usage Guidelines

Use this command to set the length of the interval that is used for periodic flooding to advertise link-state information changes that do not trigger immediate action (for example, a change to the amount of bandwidth that is allocated but that does not cross a threshold).

Enabling MPLS-TE Support in IS-IS

This topic describes the command syntax that is required to enable MPLS-TE support in an IS-IS implementation.

Enabling MPLS-TE Support in IS-IS

Cisco.com

```
router(config-router)#  
mpls traffic-eng {level-1 | level-2}
```

- Turns on MPLS traffic engineering for IS-IS Level 1 or Level 2

```
router(config-router)#  
mpls traffic-eng router-id {interface}
```

- Specifies the interface to be associated with the traffic engineering router identifier (also endpoint of a tunnel)

```
router(config-router)#  
metric-style wide
```

- Configures a router to generate and accept only new-style TLV objects

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-7

mpls traffic-eng

To turn on flooding of MPLS-TE link information into the indicated IS-IS level, use the **mpls traffic-eng** command in router configuration mode.

mpls traffic-eng {*level-1* | *level-2*}

Syntax Description

Parameter	Description
<i>level-1</i>	Floods MPLS-TE link information into IS-IS Level 1
<i>level-2</i>	Floods MPLS-TE link information into IS-IS Level 2

Usage Guidelines

This command appears as part of the routing protocol tree and causes link resource information (for instance, the bandwidth available) for appropriately configured links to be flooded in the IS-IS link-state database.

mpls traffic-eng router-id

To specify the TE router identifier for the node that is to be the IP address that is associated with the given interface, use the **mpls traffic-eng router-id** command in router configuration mode. To disable this feature, use the **no** form of this command.

mpls traffic-eng router-id *interface*

Syntax Description

Parameter	Description
<i>interface</i>	The MPLS-TE router identifier is taken from the IP address of the supplied interface. This MPLS-TE router identifier should be configured as the tunnel destination for tunnels that originate at other routers and terminate at this router. This interface should be a stable interface that will not go up and down, such as a loopback interface.

Usage Guidelines

This router identifier acts as a stable IP address for the TE configuration. This stable IP address is flooded to all nodes. For all TE tunnels that originate at other nodes and end at this node, the tunnel destination must be set to the TE router identifier of the destination node, because that identifier is the address that the TE topology database at the tunnel head uses for its path calculation.

metric-style wide

To configure a router running Intermediate System-to-Intermediate System (IS-IS) so that it generates and accepts only new-style type, length, and value objects (TLVs), use the **metric-style wide** command in router configuration mode. To disable this function, use the **no** form of this command.

- **metric-style wide** [**transition**] [**level-1** | **level-2** | **level-1-2**]
- **no metric-style wide** [**transition**] [**level-1** | **level-2** | **level-1-2**]

Syntax Description

Parameter	Description
transition	(Optional) Instructs the router to accept both old- and new-style TLV objects
level-1	Enables this command on routing Level 1
level-2	Enables this command on routing Level 2
level-1-2	Enables this command on routing Levels 1 and 2

Usage Guidelines

If you enter the **metric-style wide** command, a router generates and accepts only new-style TLV objects. The router uses less memory and other resources rather than generating both old-style and new-style TLV objects.

This style is appropriate for enabling MPLS-TE across an entire network.

Note	This discussion of metric styles and transition strategies is oriented toward TE deployment. Other commands and models may be appropriate if the new-style TLV objects are desired for other reasons. For example, a network may require wider metrics but may not use TE.

Enabling MPLS-TE Support in OSPF

This topic describes the command syntax that is required to enable MPLS-TE support in an OSPF implementation.

Enabling MPLS-TE Support in OSPF

Cisco.com

```
router(config-router)#  
mpls traffic-eng area area-num
```

- Turns on OSPF traffic engineering for the respective area

```
router(config-router)#  
mpls traffic-eng router-id interface
```

- Specifies the interface to be associated with the traffic engineering router identifier (also endpoint of a tunnel)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-8

mpls traffic-eng area

To turn on MPLS-TE for the indicated OSPF area on which MPLS-TE is enabled, use the **mpls traffic-eng area** command in router configuration mode.

mpls traffic-eng area *1-n*

Syntax Description

Parameter	Description
<i>1-n</i>	Indicates the OSPF area on which MPLS-TE is enabled

Usage Guidelines

This command affects the operation of MPLS-TE only if MPLS-TE is enabled for that routing protocol instance. Currently, only a single area may be enabled for TE.

mpls traffic-eng router-id

To specify the TE router identifier for the node that is to be the IP address that is associated with the given interface, use the **mpls traffic-eng router-id** command in router configuration mode.

mpls traffic-eng router-id *interface*

Syntax Description

Parameter	Description
<i>interface</i>	The MPLS-TE router identifier is taken from the IP address of the supplied interface. This MPLS-TE router identifier should be configured as the tunnel destination for tunnels that originate at other routers and terminate at this router. This interface should be a stable interface that will not go up and down, such as a loopback interface.

Usage Guidelines

This router identifier acts as a stable IP address for the TE configuration. This stable IP address is flooded to all nodes. For all TE tunnels that originate at other nodes and end at this node, the tunnel destination must be set to the TE router identifier of the destination node, because that identifier is the address that the TE topology database at the tunnel head uses for its path calculation.

Creating and Configuring a Traffic Tunnel

This topic describes the command syntax that is required to create and configure a traffic tunnel on a Cisco IOS device.

Creating and Configuring a Traffic Tunnel

Cisco.com

```
router(config)#  
interface tunnel number
```

- Configures a tunnel interface

```
router(config-if)#  
tunnel destination ip-address
```

- Specifies the destination for the tunnel

```
router(config-if)#  
tunnel mode mpls traffic-eng
```

- Sets the tunnel encapsulation mode to MPLS Traffic Engineering

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-9

interface tunnel

The **interface tunnel** command is used to declare a tag-switched path (TSP) tunnel interface. The tunnel interface *number* is in the range of 0 to 2147483647.

tunnel destination

To specify the destination for a tunnel interface, use the **tunnel destination** interface configuration command.

tunnel destination {*hostname* | *ip-address*}

Syntax Description

Parameter	Description
<i>hostname</i>	Name of the host destination
<i>ip-address</i>	IP address of the host destination expressed in decimal in four-part, dotted notation

Usage Guidelines

Two tunnels cannot use the same encapsulation mode with exactly the same source and destination address. The solution is to create multiple loopback interfaces and source tunnels off each loopback interface. Refer to Cisco IOS AppleTalk and the Novell Internetwork Packet Exchange (IPX) Configuration Guide for more information on AppleTalk Cayman tunneling.

tunnel mode mpls traffic-eng

To set the mode of a tunnel to MPLS for TE, use the **tunnel mode mpls traffic-eng** command in interface configuration mode.

Usage Guidelines

This command specifies that the tunnel interface is for an MPLS-TE tunnel and enables the various tunnel MPLS configuration options.

Creating and Configuring a Traffic Tunnel (Cont.)

Cisco.com

Configuring a Traffic Tunnel

```
router(config-if)#
```

```
tunnel mpls traffic-eng bandwidth bandwidth
```

- Configures the requested **bandwidth** for the MPLS Traffic Engineering tunnel

```
router(config-if)#
```

```
tunnel mpls traffic-eng priority setup-priority [ hold-priority ]
```

- Configures the setup and reservation priority for an MPLS Traffic Engineering tunnel

```
router(config-if)#
```

```
tunnel mpls traffic-eng affinity properties [ mask mask value ]
```

- Configures an affinity (the properties the tunnel requires in its links) for an MPLS Traffic Engineering tunnel

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-10

tunnel mpls traffic-eng bandwidth

To configure the bandwidth that is required for an MPLS-TE tunnel, use the **tunnel mpls traffic-eng bandwidth** command in configuration mode. To disable this feature, use the **no** form of this command.

- **tunnel mpls traffic-eng bandwidth** [*sub-pool* | **global**] *bandwidth*
- **no tunnel mpls traffic-eng bandwidth** [*sub-pool* | **global**] *bandwidth*

Syntax Description

Parameter	Description
<i>bandwidth</i>	The bandwidth required for an MPLS-TE tunnel. Bandwidth is specified in kbps.

tunnel mpls traffic-eng priority

To configure the setup and reservation priority for an MPLS-TE tunnel, use the **tunnel mpls traffic-eng priority** command in interface configuration mode. To remove the specified setup and reservation priority, use the **no** form of this command.

- **tunnel mpls traffic-eng priority** *setup-priority* [*hold-priority*]
- **no tunnel mpls traffic-eng priority** *setup-priority* [*hold-priority*]

Syntax Description

Parameter	Description
<i>setup-priority</i>	The priority that is used when signaling a label switched path (LSP) for this tunnel to determine what existing tunnels are eligible for pre-emption. The range is from 0 to 7, where a lower numeric value indicates a higher priority.
<i>hold-priority</i>	(Optional) The priority that is associated with an LSP for this tunnel once it is established to figure out if it should be pre-empted by other LSPs that are being signaled. The range is from 0 to 7, where a lower numeric value indicates a higher priority.

Usage Guidelines

When an LSP is being signaled and an interface does not currently have enough bandwidth available for that LSP, the call admission software preempts lower-priority LSPs so that the new LSP can be admitted. (LSPs are preempted if that allows the new LSP to be admitted.)

In the described determination, the new LSP's priority is its setup priority, and the existing LSP's priority is its hold priority. The two priorities make it possible to signal an LSP with a low setup priority (so that the LSP does not preempt other LSPs on setup) but a high hold priority (so that the LSP is not preempted after it is established).

Setup priority and hold priority are typically configured to be equal, and setup priority cannot be better (numerically smaller) than the hold priority.

Defaults

- setup priority = 7
- hold priority = 7

tunnel mpls traffic-eng affinity

To configure tunnel affinity (the properties that the tunnel requires in its links), use the **tunnel mpls traffic-eng affinity** command in interface configuration mode. To disable this feature, use the **no** form of this command.

tunnel mpls traffic-eng affinity *properties* [**mask** *mask*]

Syntax Description

Parameter	Description
<i>properties</i>	Attribute values that are required for links that carry this tunnel (values of bits are either 0 or 1).
mask <i>mask</i>	Attribute values that should be checked. If a bit in the mask is 0, that bit or the attribute value of the link is irrelevant. If a bit in the mask is 1, the link attribute value and the required affinity of the tunnel for that bit must match.

Defaults

- *properties*—0X00000000
- *mask*—0X0000FFFF

Creating and Configuring a Traffic Tunnel (Cont.)

Cisco.com

Tunnel Configuration Commands

```
router(config-if)#
```

```
tunnel mpls traffic-eng path-option number {dynamic |  
explicit {name path-name | path-number}} [lockdown]
```

- Configures the tunnel to use a named IP explicit path or a path dynamically calculated from the traffic engineering topology database

```
router(config)#
```

```
ip explicit-path { name word | identifier number }  
[ {enable | disable } ]
```

- Enters the subcommand mode for IP explicit paths and creates or modifies the specified path

```
router(config-if)# *** Physical ****
```

```
mpls traffic-eng administrative-weight weight
```

- Overrides the Interior Gateway Protocol cost of the link

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-11

tunnel mpls traffic-eng path-option

To configure a path option for an MPLS-TE tunnel, use the **tunnel mpls traffic-eng path-option** command in interface configuration mode. To disable the specified path option, use the **no** form of this command.

- **tunnel mpls traffic-eng path-option** *number* {**dynamic** | **explicit** {*name path-name* | *path-number*}} [**lockdown**]
- **no tunnel mpls traffic-eng path-option** *number* {**dynamic** | **explicit** {*name path-name* | *path-number*}} [**lockdown**]

Syntax Description

Parameter	Description
<i>number</i>	When multiple path options are configured, lower numbered options are preferred.
dynamic	Path of the LSP is dynamically calculated
name <i>path-name</i>	Path name of the IP explicit path that the tunnel uses with this option.
<i>path-number</i>	Path number of the IP explicit path that the tunnel uses with this option.
lockdown	(Optional) The LSP cannot be reoptimized.

Usage Guidelines

Multiple path setup options may be configured for a single tunnel. For example, you can configure several explicit paths and a dynamic option for one tunnel. Path setup prefers options with lower numbers to options with higher numbers, so option 1 is the most preferred option.

ip explicit-path

To enter the subcommand mode for IP explicit paths to create or modify the named path, use the **ip explicit-path** command in global configuration mode. An IP explicit path is a list of IP addresses, each representing a node or link in the explicit path.

ip explicit-path {*name word* | *identifier number*} [{**enable** | **disable**}]

Syntax Description

Parameter	Description
name <i>word</i>	Specifies explicit path by name.
identifier <i>number</i>	Specifies explicit path by number. You can specify a number from 1 to 65535.
enable	Sets the state of the path to be enabled.
disable	Prevents the path from being used for routing while it is being configured.

mpls traffic-eng administrative-weight

To override the IGP administrative weight (cost) of the link, use the **mpls traffic-eng administrative-weight** command in interface configuration mode (default matches IGP cost).

mpls traffic-eng administrative-weight *weight*

Syntax Description

Parameter	Description
<i>weight</i>	Cost of the link

Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency

This topic identifies the command syntax that is required to deploy traffic tunnels with autoroute or forwarding adjacency on Cisco IOS devices.

Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency

Cisco.com

```
router(config-if)#  
tunnel mpls traffic-eng autoroute announce
```

- Causes the IGP to use the tunnel in its enhanced SPF

```
router(config-if)#  
tunnel mpls traffic-eng autoroute metric { absolute |  
relative } value
```

- Enables you to specify the MPLS traffic engineering tunnel metric that the IGP enhanced SPF calculation uses

```
router(config-if)#  
tunnel mpls traffic-eng forwarding-adjacency [holdtime  
value ]
```

- Causes the IGP to announce the tunnel as a point-to-point link to other routers within an area

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-12

tunnel mpls traffic-eng autoroute announce

To instruct the IGP to use the tunnel in its SPF calculation (if the tunnel is up), use the **tunnel mpls traffic-eng autoroute announce** command in interface configuration mode. To disable this feature, use the **no** form of this command.

Defaults

The tunnel is not used by the IGP in its SPF calculation.

Usage Guidelines

Currently, the only way to cause traffic to be forwarded onto a tunnel is by enabling this feature or by configuring forwarding explicitly with an interface static route.

tunnel mpls traffic-eng autoroute metric

To specify the MPLS-TE tunnel metric to be used by IGP autoroute, use the **tunnel mpls traffic-eng autoroute metric** command in interface configuration mode. To disable this feature, use the **no** form of this command.

tunnel mpls traffic-eng autoroute metric {absolute | relative} value

Syntax Description

Parameter	Description
metric	The MPLS-TE tunnel metric
absolute	The MPLS-TE tunnel metric mode absolute: a positive metric value can be supplied.
relative	The MPLS-TE tunnel metric mode relative: a positive, negative, or zero value can be supplied.

Defaults

The default is metric relative 0.

tunnel mpls traffic-eng forwarding-adjacency

To advertise a TE tunnel as a link in an IGP network, use the **tunnel mpls traffic-eng forwarding-adjacency** command in interface configuration mode.

tunnel mpls traffic-eng forwarding-adjacency {holdtime value}

Syntax Description


Parameter	Description
holdtime value	(Optional) Time in ms that a TE tunnel waits after going down before informing the network. The range is 0 to 4,294,967,295 ms.

Defaults


The default value is 0.

Implementing MPLS-TE

This topic provides several examples of MPLS-TE implementations.




This figure shows a typical implementation of MPLS-TE on an ingress provider edge router (PE router).



This figure shows a typical implementation of MPLS-TE on a provider router (P router).

Implementing MPLS-TE: The Foundation (Cont.)

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-15

This figure shows a typical implementation of MPLS-TE on an egress PE router.

Implementing MPLS-TE: Building a Dynamic Tunnel

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-16

This figure shows a typical implementation of a dynamic MPLS-TE tunnel. The actual tunnel path could be through either P1 or P2 as selected by the IGP. The tunnel is announced to the IGP using `autoroute`.

Implementing MPLS-TE: Explicit Path

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-17

This figure shows a typical implementation of an explicit path MPLS-TE tunnel. The tunnel is announced to the IGP using `autoroute`.

Implementing MPLS-TE: VPNs – All Traffic

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-18

This figure shows a typical implementation of an MPLS-TE tunnel that is being used to transport all Virtual Private Network (VPN) traffic between point of presence (POP) sites. The tunnel is announced to the IGP using autoroute.

VPN traffic is forwarded to the Border Gateway Protocol (BGP) next hop; therefore, traffic from CE-A1 that is leaving PE-X will be forwarded to PE-Y. Because the destination of the tunnel is the BGP next hop (PE-Y), the IGP will see the tunnel as the best path. As a result, the VPN traffic is forwarded over the tunnel.


However, for VPN forwarding to be successful, there must be an end-to-end LSP. To create this LSP, MPLS must be enabled on the tunnel interface. Enabling MPLS accomplishes several things. First, a Label Distribution Protocol (LDP) relationship is established between the two PE routers. Normally, it would be necessary to enable directed LDP hellos to establish an LDP relationship between two routers that are not directly connected. However, because the relationship is being established over a tunnel, Cisco IOS software automatically handles this relationship.

Once the LDP session is established, labels are propagated and an LSP is created over the tunnel. The LSP allows a label forwarding information base (LFIB) to be created, linking all of the prefixes in the virtual routing and forwarding (VRF) table to the tunnel.

This configuration will forward all traffic from all VPN customers that are destined for PE-Y over the tunnel. To limit tunnel access to a single VRF table or to specific traffic with a VRF table, static routing must be used. This subject is covered in the next example.

Implementing MPLS-TE: VPNs – Selected Traffic

Cisco.com



In this example two tunnels have been created. Tunnel 1 is through P1 and will be used as the default path for all VPN traffic. Tunnel 2 is through P2 and will be used only for traffic for the 10.1.10.0 prefix in VRF A. You need for the IGP to know about Tunnel 2 so that it is announced using autoroute. However, so that Tunnel 2 is not available to other traffic, the metric is set to 4.

With Tunnel 2 known to the IGP but not available as a preferred router to PE-Y, traffic does not normally pass over it. However, by use of a static router for prefix 10.1.10.0, this behavior can be overridden. Traffic that is destined for prefix 10.1.10.0 can be forwarded using the static route to Tunnel 2. All other traffic will be forwarded over Tunnel 1.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- For MPLS-TE support ip cef has to be enabled globally on the device.
- CEF optimizes network performance and scalability for networks with dynamic, topologically dispersed traffic patterns.
- MPLS forwarding of IPv4 packets along normally routed paths is sometimes called dynamic label switching.
- The mpls traffic-eng isis-level command causes link resource information (for instance, bandwidth available) for appropriately configured links to be flooded in the IS-IS link-state database.
- Support of MPLS-TE in OSPF can presently be enabled only in a single area.
- When you are using the tunnel destination command, remember that two tunnels cannot use the same encapsulation mode with exactly the same source and destination address.
- Use the tunnel mpls traffic-eng autoroute announce command to instruct the IGP to use the tunnel in its SPF or next hop calculation (if the tunnel is up).

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-20

References

For additional information, refer to this resource:

- Search the Cisco.com website at <http://www.cisco.com> for more information on the topics that were discussed in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) To enable MPLS-TE tunnel signaling on a device, you must use the _____ command.
- Q2) Which command enables dynamic label switching to be enabled for a Cisco IOS platform?
- A) **ip cef**
 - B) **mpls ip**
 - C) **mpls-te enable**
 - D) **mpls traffic-eng tunnels**
- Q3) The command to allow a router to generate and accept only new-style TLV objects is the _____ command. (Base command only is acceptable.)
- Q4) Which of the following enables MPLS-TE in an OSPF implementation?
- A) **mpls-te enable**
 - B) **mpls traffic-eng**
 - C) **metric-style wide**
 - D) **mpls traffic-eng area**
- Q5) The command that is used to declare a TSP tunnel interface is the _____ command.
- Q6) The command that is used to instruct the IGP to use the tunnel in its SPF or next-hop calculation is the _____ command.

Lesson Answer Key

- Q1) **mpls traffic-eng tunnels**
Relates to: Enabling Device-Level MPLS-TE Support
- Q2) B
Relates to: Enabling MPLS-TE on an Interface
- Q3) **metric-style wide**
Relates to: Enabling MPLS-TE Support in IS-IS
- Q4) D
Relates to: Enabling MPLS-TE Support in OSPF
- Q5) **interface tunnel**
Relates to: Creating and Configuring a Traffic Tunnel
- Q6) **tunnel mpls traffic-eng autoroute announce**
Relates to: Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency

Advanced MPLS-TE Path Selection

Overview

This lesson describes advanced MPLS-TE commands that are used in path selection. The configuration commands are accompanied by usage guidelines and sample setups.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes advanced MPLS-TE path selection. Upon completing this lesson, you will be able to do the following:

- Describe a typical MPLS-TE implementation
- Identify the command syntax to create dynamic path options for a traffic tunnel
- Identify the command syntax to create explicit path options for a traffic tunnel
- Identify the command syntax to assign traffic to traffic tunnels with static routes
- Identify the command syntax to configure periodic and event-triggered optimization
- Identify the command syntax to configure path selection restrictions
- Identify the command syntax to modify the administrative weight of a traffic tunnel
- Identify the command syntax to enable the autoroute forwarding engine

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course

Outline

This lesson includes these topics:

- Overview
- MPLS-TE Case Study
- Creating Dynamic Path Options for a Traffic Tunnel
- Creating Explicit Path Options for a Traffic Tunnel
- Assigning Traffic to Traffic Tunnels with Static Routes
- Configuring Periodic and Event-Triggered Optimization
- Configuring Path Selection Restrictions
- Modifying Administrative Weight
- Autoroute Forwarding Engine
- Lesson Summary
- Lesson Review

MPLS-TE Case Study

This topic will discuss a typical MPLS-TE implementation.

MPLS-TE Case Study: Company Network Description

Cisco.com

The “company” is a service provider that has the following components:

- **Core network:**
 - **Core routers in locations are heavily meshed**
 - **Two permanent connections to the Internet**
- **Local or remote access networks:**
 - **Connected to core in a resilient way, allowing a remote access via point of presence**
- **Access networks:**
 - **BGP customers**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-3

The example in the figure shows a classic Internet service provider (ISP) architecture based on three levels of hierarchy. The design should bring together some of the aspects of traffic engineering and routing design that are discussed in this module, namely, the following:


- Routing protocol choice and interaction between different routing protocols
- Support for MPLS-TE tunnels

MPLS-TE Case Study: IGP Design

Cisco.com

IS-IS is used for routing inside and between the core and POPs:

- A simple flat design using one IS-IS Level 2 backbone
- Loopbacks are advertised with 32-bit network mask
- Use a new flavor of IS-IS TLVs (wide metric)



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-4.4

The network is based on a three-level hierarchy, consisting of the following:

- **Core network:** Highly meshed central sites with high bandwidth requirements between them
- **POP sites:** A distribution layer of regional sites, which are connected back to the core over redundant links and which provide access for remote sites
- **BGP peers:** Upstream Internet providers and customer networks that are connected to the distribution sites via leased lines

The core and regional networks are a complex mesh of routers, and require efficient, scalable routing with fast convergence. A link-state protocol is ideally suited to this situation. Therefore, Integrated IS-IS is the choice.

The proposed structure of the IS-IS protocol is as follows:


- Simply enabling Integrated IS-IS on Cisco routers sets their operation as Level 2 routers.
- All IP subnets would be visible individually to all routers on the IS-IS network from their Level 2 advertisements.
- “Wide metric” is used, which allows for greater granularity of path selection.

MPLS-TE Case Study: BGP Design

Cisco.com

All POP sites are fully meshed with IBGP:

- Multiprotocol Label Switching (MPLS) in the core, but not inside the POP.
- Route reflectors are used at POP sites, to ease configuration of future POP routers.
- Route reflectors are connected in a full mesh.



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-5

The most recent method of core transport uses MPLS as a transport mechanism in the core network. Packets are switched through the core of the network at Layer 2, bypassing the traditional Layer 3 routing process.

The issue with the edge-only BGP peer design is that the routers in between the edges need to have routes to process the packets. If these packets pass through these routers with MPLS tags, they no longer need IP routes to these destinations.

In all cases, BGP relies on another protocol to resolve its next-hop address. An IGP is required to do this, so IS-IS must contain routes to the edge routers and to their attached (public) subnets.

To reduce the number of Internal Border Gateway Protocol (IBGP) sessions that are required between POPs, route reflectors are used.

These tools allow the full-mesh requirement of traditional IBGP operation to be relaxed.

Creating Dynamic Path Options for a Traffic Tunnel

This topic describes the command syntax to create dynamic path options for a traffic tunnel.

[illegible]

This sample configuration shows how to implement traffic engineering (TE) on top of an existing MPLS network.

The sample shows the implementation of two TE tunnels of 250 kbps; the tunnels are either being automatically set up by the ingress label switch routers, or LSRs (POP A, POP B, POP C routers) or manually configured with the explicit paths.


“TE” is a generic term that refers to the use of different technologies to optimize the utilization of the capacity and topology of a given backbone.

Creating Dynamic Path Options for a Traffic Tunnel (Cont.)

Cisco.com

MPLS-TE Platform

Provide an underlying platform for MPLS-TE by configuring the extended IS-IS and RSVP for bandwidth assurance— $\frac{3}{4}$ of the bandwidth is reservable for traffic tunnels.



MPLS-TE uses an extension to existing protocols such as RSVP, IS-IS, and OSPF to calculate and establish unidirectional tunnels that are set according to the network constraint. Traffic flows are mapped on the different tunnels depending on their destination.

With Cisco, MPLS-TE is built on the following Cisco IOS mechanisms:

- A link-state IGP (such as IS-IS) with extensions for the global flooding of resource information, and extensions for the automatic routing of traffic onto LSP tunnels as appropriate.
- An MPLS-TE path calculation module (constraint-based routing, or CBR), which determines the paths to use for LSP tunnels.
- LSP tunnels, which are signaled through RSVP. LSP tunnels that are represented as IOS tunnel interfaces are unidirectional.
- An MPLS-TE link management module that does link admission and bookkeeping of the resource information to be flooded.
- Label switching forwarding, which provides routers with a Layer 2-like ability to direct traffic across multiple hops as directed by the resource-based routing algorithm.

The meaning of some relevant commands is as follows:


- Global configuration commands:
 - **ip cef:** Enables standard CEF operation
 - **mpls traffic-eng tunnels:** Enables the MPLS-TE tunnel feature on a device
- IS-IS configuration mode:
 - **mpls traffic-eng level 2:** Turns on MPLS-TE for IS-IS Level 2
 - **mpls traffic-eng router-id loopback0:** Specifies the TE router identifier for the node to the IP address that is associated with interface loopback0.
 - **metric-style wide:** Configures a router to generate and accept only new-style TLV objects
- Interface configuration mode:
 - **mpls traffic-eng tunnels:** Enables the MPLS-TE tunnel feature on an interface.
 - **ip rsvp bandwidth** [*total-reservable-flow*] [*single-flow*]: Enables RSVP for IP on an interface and specifies the amount of bandwidth to be reserved

Creating Dynamic Path Options for a Traffic Tunnel (Cont.)

Cisco.com

Sample path setup:

- Based on the default parameters and the assumption that the links are free; the best path for both would be the same as computed by IS-IS.
- Therefore, the label switched paths for both tunnels use the middle Core 1 and Core 6 routers.



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-8

The example in the figure shows how the CBR algorithm proposes a path between tunnel endpoints that satisfies the initial requests at the headend of the tunnel.

Based on the assumption that all TE links are free, the traffic from POP A to POP C and from POP B to POP C is directed along the same least-cost path (Core 1 – Core 6) because it is used by IS-IS for native IP routing.

The reason is very simple; CBR is a routing process that takes into account the following two considerations:

- The best route is the least-cost route with enough resources. CBR uses its own metric (administrative weight, or TE cost), which is by default equal to the IGP:
 - In case of a tie, the path with the highest minimum available bandwidth is selected.
 - If a tie continues to exist, then the path with the smallest hop count is selected.
 - Finally, if there is still a tie, then a random path is selected.
- The result of CBR is an explicit route, which is used by RSVP to reserve resources and establish an LSP path.

Creating Dynamic Path Options for a Traffic Tunnel (Cont.)

Cisco.com

#POP-A and POP-B configuration are identical

```
interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination POP-C
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 250
 tunnel mpls traffic-eng path-option 1 dynamic
```

Endpoint of tunnel (pref. loopback)

Tunnel establish and hold priority

Bandwidth needed for tunnel

#POP-C configuration

```
interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 250
 tunnel mpls traffic-eng path-option 1 dynamic
```

Path computation option

interface Tunnel2

```
 ip unnumbered Loopback0
 tunnel source Loopback0
 tunnel destination POP-B
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 250
 tunnel mpls traffic-eng path-option 1 dynamic
```

Traffic tunnels are unidirectional, so a similar configuration should be applied in the opposite direction.

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0-4-9

To set up a dynamic TE tunnel (assuming that the IGP platform has been prepared), follow these command steps in the tunnel interface configuration mode:

- **ip unnumbered loopback0:** Gives the tunnel interface an IP address. An MPLS-TE tunnel interface should be unnumbered because it represents a unidirectional link.
- **tunnel source:** Specifies the source for a tunnel. The source of the tunnel must be the destination of the tunnel in the opposite direction, usually the local loopback address.
- **tunnel destination:** Specifies the destination for a tunnel. The destination of the tunnel must be the source of the tunnel in the opposite direction, usually a loopback address.
- **tunnel mode mpls traffic-eng:** Sets the tunnel encapsulation mode to MPLS-TE.
- **tunnel mpls traffic-eng priority:** Configures the setup and reservation priority for an MPLS-TE tunnel.
- **tunnel mpls traffic-eng bandwidth:** Configures the bandwidth for the MPLS-TE tunnel.
- **tunnel mpls traffic-eng path-option number dynamic:** The LSP path is dynamically calculated.

Creating Explicit Path Options for a Traffic Tunnel

This topic describes the command syntax to create explicit path options for a traffic tunnel.



The example in the figure shows how to avoid the first step in the CBR algorithm by manually setting the explicit path between tunnel endpoints. This path might be derived from the least-cost path.

The best route might not be the least-cost route with enough resources. The best route might be any sequence of next-hop routers that are configured at the headend of the tunnel.

Such a route, as proposed by the network administrator, is then checked against the extended link-state database that is carrying information on currently available resources.

If this check is successful, CBR honors the route and RSVP is initiated to reserve some bandwidth and establish an LSP path. Otherwise, the tunnel stays down.

Creating Explicit Path Options for a Traffic Tunnel (Cont.)

Cisco.com

Explicit Path Configuration

```
router(config)#
```

```
ip explicit-path { name word | identifier number }  
[ {enable | disable } ]
```

- Enter the subcommand mode for IP explicit paths and create or modify the specified path

```
router(config)#
```

```
next-address strict ip-address
```

- Configures the next hop in a strictly routed, explicit MPLS-TE path

```
router(config-if)#
```

```
tunnel mpls traffic-eng path-option number {dynamic |  
explicit {name path-name | path-number}} [lockdown]
```

- Configures the tunnel so that it uses a named IP explicit path or a path dynamically calculated from the Traffic Engineering topology database

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-11

next-address [loose | strict] A.B.C.D

The next-address has two keyword: loose and static (default). This section focuses on the static keyword used to specify an explicit path, where the path can include only the addresses that are specified.

An explicit path can be created by specifying each address in sequence by using the **next-address** command with the strict optional keyword. However, since static is the default mode for path creation, specifying each address in sequence by using the **next-address** command with **NO** optional keyword will accomplish the same thing.

To remove an address from an explicit path, use the **no next-address** command in IP explicit path configuration mode.

Parameter	Description
loose	Specifies that the previous address (if any) in the explicit path need not be directly connected to A.B.C.D, and that the router is free to determine the path from the previous address (if any) to A.B.C.D.
strict	Specifies that the previous address (if any) in the explicit path must be directly connected to A.B.C.D.
<i>A.B.C.D</i>	Next IP address in the explicit path.

Defaults

For **next-address strict** the address has to be one of the following:

- The interface address on a connected network for numbered links
- The TE router identifier of the next-hop router for unnumbered links

For explicit paths to be used for TE tunnels within an IGP area, you may specify a combination of both loose and strict hops via **next-address** commands.

Creating Explicit Path Options for a Traffic Tunnel (Cont.)

Cisco.com

#POP-A configuration

```
interface Tunnel1
  ip unnumbered Loopback0
  tunnel destination POP-C
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng priority 1 1
  tunnel mpls traffic-eng bandwidth 250
  tunnel mpls traffic-eng path-option 1 explicit name Core2-3
!
ip explicit-path name Core2-3 enable
  next-address {strict} Core-2
  next-address {strict} Core-3
  next-address {strict} POP-C
```


Explicit path represents the explicit route object, encoded as incoming IP addresses of the routers along the path.

To set up a static TE tunnel (assuming that the IGP platform has been prepared), use these additional steps:

- Tunnel interface configuration mode:
 - **tunnel mpls traffic-eng path-option number explicit:** This command enables you to configure the tunnel to use a named IP explicit path from the TE topology database.
- Global configuration mode:
 - **ip explicit-path:** An IP explicit path is a list of IP addresses, each representing a node or link (incoming IP address) in the explicit path.

Assigning Traffic to Traffic Tunnels with Static Routes

This topic describes the command syntax that is needed to assign traffic to tunnels with static routes.




When the LSP path is established for the tunnel, the MPLS traffic can flow across it. However, the tunnel is not yet usable because it is not a generic routing encapsulation (GRE) tunnel and does not show up in the routing table.

For this reason the engineered tunnels can be used for IP routing only if the tunnel is explicitly specified for routing in one of the following ways:

- Via static routes that point to the tunnel
- Via policy routing that sets a next-hop interface to the tunnel

Configuring Periodic and Event-Triggered Optimization

This topic describes the command syntax to configure periodic and event-triggered optimization.



The LSP path is constantly monitored to maintain the network traffic tunnel in a desired state. When the path is broken and the tunnel has been set up dynamically, the headend router tries to find an alternative solution. This process is referred to as rerouting.

Reoptimization occurs when a device examines tunnels with established LSPs to see if better LSPs are available. If a better LSP seems to be available, the device attempts to signal the better LSP and, if successful, replaces the old and inferior LSP with the new and better LSP.

This reoptimization might be triggered by a “link up” event or might occur at configurable intervals (the default is 1 hour). Instability and oscillations can result if the reoptimization interval is set too small. However, the network will not react to unexpected shifts in traffic if the interval is too great. One hour is a reasonable compromise. With reoptimization, traffic is routed so that it sees the lightest possible loads on the links that it traverses.

Unfortunately, reoptimization does not bring any improvements for a tunnel that has been established statically. In this case, the path is explicitly determined, which ties the headend router to strictly follow the explicit path.

Configuring Periodic and Event-Triggered Optimization (Cont.)

Cisco.com

Configuring Optimization

```
router(config)#
```

```
mpls traffic-eng reoptimize events {link-up}
```

- Triggers automatic reoptimization whenever an interface becomes operational

```
router(config)#
```

```
mpls traffic-eng reoptimize timers frequency seconds
```

- Controls the frequency with which tunnels with established label switched paths (LSPs) are checked for better LSPs

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-15

mpls traffic-eng reoptimize events

To trigger automatic reoptimization whenever an interface becomes operational, use the **mpls traffic-eng reoptimize events** router configuration command. Use the **no** form of this command to disable this feature.

mpls traffic-eng reoptimize events {link-up}

Syntax Description

Parameter	Description
link-up	Triggers automatic reoptimization whenever an interface becomes operational

Defaults

Event-based reoptimization is disabled.

Command Modes

Global configuration

mpls traffic-eng reoptimize timers frequency

To control the frequency with which tunnels with established LSPs are checked for better LSPs, use the **mpls traffic-eng reoptimize timers frequency** command in global configuration mode.

mpls traffic-eng reoptimize timers frequency *seconds*

Syntax Description

Parameter	Description
<i>seconds</i>	Sets the frequency of reoptimization, in seconds. A value of 0 disables reoptimization.

Defaults

3600 sec (1 hour), with a range of 0 to 604,800 sec (1 week).

Command Modes

Global configuration


Usage Guidelines

This command enables TE tunnels to periodically examine tunnels with established LSPs to see if better LSPs are available. If a better LSP seems to be available, the device attempts to signal the better LSP and, if successful, replaces the old and inferior LSP with the new and better LSP.

Configuring Periodic and Event-Triggered Optimization (Cont.)

Cisco.com

Configuring Event-Triggered Optimization



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-16

The example in the figure shows how traffic can be engineered across a path in the network and how a backup route for that traffic-engineered path can be established.

The primary path is explicit with a manually specified RSVP explicit route object (ERO). If this path suddenly cannot be followed, the MPLS-TE engine uses the next path option, which in this case is a dynamic route.


The drawback of this solution is the time that is needed to establish a backup TE route for the lost LSP path and the time that is needed to revert to the primary path once it becomes available again.

Though the search for an alternate path is immediately triggered by a “link down” event, there is still downtime while the alternate path is being built.

The convergence back to the primary tunnel may be sped up if you configure the triggering of a constraint-based LSP computation on any LSP change (link up event) or by decreasing the time of periodic LSP checks.

Configuring Path Selection Restrictions

This topic describes the command syntax to configure path selection restrictions.



In many cases, some links will need to be excluded from the constraint-based shortest path first (SPF) computation. This exclusion can be implemented using the resource class affinity bits of the traffic tunnel and resource class bits of the links over which the tunnel should pass (following the computed LSP path).


A 32-bit resource class affinity string that is accompanied by a corresponding resource class mask characterizes each traffic tunnel. The “0” bits in the mask exclude the respective link resource class bits from being checked.

Each link is characterized by its resource class 32-bit string, which is set to 0 by default. The matching of the tunnel resource class affinity string with the resource class string of the link is performed during the LSP path computation.

Configuring Path Selection Restrictions (Cont.)

Cisco.com

Configuring Affinity Bits on Links



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-18

The figure shows a sample network with the tunnel resource class affinity bits and link resource bits. The main goal is to force the CBR algorithm to use only links that are explicitly dedicated to certain tunnels for its path computation.

Because it is desirable to move all blue tunnels away from POP A interfaces and red tunnels from POP B interfaces, different link resource class bits are set: 0x00000001 for red interfaces and 0x00000002 for blue interfaces. Those link resource class attribute bits then become a part of the LSP advertisements, which allows all participants to include this information when computing paths for TE tunnels.

Configuring Path Selection Restrictions (Cont.)

Cisco.com

Configuring Affinity Bits on Tunnels

```
#POP-C configuration
interface Tunnel1
  description Red tunnel
  ip unnumbered Loopback0
  tunnel destination POP-A
  tunnel mpls traffic-eng priority 1 1
  tunnel mpls traffic-eng affinity 0x00000001 mask 0x00000001
  .....

interface Tunnel2
  description Blue tunnel
  ip unnumbered Loopback0
  tunnel destination POP-B
  tunnel mpls traffic-eng priority 2 2
  tunnel mpls traffic-eng affinity 0x00000002 mask 0x00000002
  .....
```

Link attributes for tunnel
(mask 1 means it must match)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-19

In the example for Tunnel 1, the tunnel requires a match only on the last bit, whereas Tunnel 2 checks the setting of the second bit from the right-hand side.

With tunnel resource class affinity bits and link resource class bits set, the constraint-based path computation considers only the paths where the match is found.

Example: Configuring Affinity Bits on Tunnels

For Tunnel 1 (red tunnel), the affinity bits are set at 0x00000001, and the mask is set at 0x00000001; the attributes of the red links (POP A interfaces) match (attributes 0x00000001). The black links (core interfaces) also match (attributes 0x00000011). The blue links (POP B interfaces) that are marked with the attribute 0x00000002 are excluded from constraint-based SPF computation because they do not match.

Configuring Path Selection Restrictions (Cont.)

Cisco.com

Configuring Link or Node Exclusion

```
router(config)#
```

```
ip explicit-path { name word | identifier number } [
{enable | disable} ]
```

- Enter the subcommand mode for IP explicit paths and create or modify the specified path.

```
router(config)#
```

```
exclude-address ip-address
```

- Configures explicit path with address exclusion.
- PCALC uses CBR to calculate the path, avoiding specified link (if the configured IP address refers to a link IP address) or node (if the configured IP address specifies a router ID).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-20

Link or node exclusion is accessible using the **ip explicit-path** command, which enables you to create an IP explicit path and enter a configuration submode for specifying the path. Link or node exclusion uses the **exclude-address** submode command to specify addresses that are excluded from the path.

If the exclude address for an MPLS-TE LSP identifies a flooded link, the constraint-based SPF routing algorithm does not consider that link when computing paths for the LSP. If the exclude address specifies a flooded MPLS-TE router ID, the constraint-based SPF routing algorithm does not allow paths for the LSP to traverse the node that is identified by the router ID.

exclude-address

An IP explicit path is a list of IP addresses, each representing a node or link in the explicit path. To exclude an address from an IP explicit path, use the **exclude-address** command after entering explicit path configuration mode via the **ip explicit-path** command.

exclude-address *A.B.C.D*

Syntax Description

Parameter	Description
<i>A.B.C.D</i>	Excludes an address from subsequent partial path segments. You can enter the IP address of a link or the router ID of a node.

Defaults

Addresses are not excluded from an IP explicit path unless they are explicitly excluded by the **exclude-address** command.

If you enter the **exclude-address** command and specify the IP address of a link, the constraint-based SPF routine does not consider that link when it sets up MPLS-TE paths. If the excluded address is a flooded MPLS-TE router ID, the constraint-based SPF routine does not consider that entire node. The person performing the configuration must know the router IDs of the routers because it will not be apparent whether the specified number is for a link or for a node.

Note	MPLS-TE will accept an IP explicit path that is composed of either all exclude addresses configured by the exclude-address command or all include addresses configured by the next-address command, but not a combination of both.

Configuring Path Selection Restrictions (Cont.)

Cisco.com

Excluding Links or Nodes for Traffic Tunnels

```
#POP-C configuration
interface Tunnel1
  description Red tunnel
  ip unnumbered Loopback0
  tunnel destination POP-A
  tunnel mpls traffic-eng priority 1 1
  tunnel mpls traffic-eng path-option 1 explicit name ex1
  .....
ip explicit-path name ex1 enable
  exclude-address POP-B
interface Tunnel2
  description Blue tunnel
  ip unnumbered Loopback0
  tunnel destination POP-B
  tunnel mpls traffic-eng priority 2 2
  tunnel mpls traffic-eng path-option 1 explicit name ex2
  .....
ip explicit-path name ex2 enable
  exclude-address POP-A
```

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—4-21

In a previous example, affinity bits were used to restrict the possible paths that could be used for tunnel creation. This example shows the usage of the IP exclude address feature to accomplish the same function. By excluding the desired nodes, tunnels will not use any links that lead to those nodes.

Note	POP A and POP B in our example have to be the router IDs of the specified nodes. Typically, those addresses conform to the loopback address of the node.

Modifying Administrative Weight

This topic describes the command syntax to modify the administrative weight of a traffic tunnel.



The constrained-based path computation selects the path that the dynamic traffic tunnel will take, based on the administrative weight (TE cost) of each individual link.


This administrative weight is, by default, equal to the IGP link metric (cost). Increase the TE cost on the link if it is desirable to exclude a certain link from any path computation while keeping the link available in the event that the link represents the only available path.

Modifying Administrative Weight (Cont.)

Cisco.com

Configuring Administrative Weight

In the absence of a static path due to a failed POP C-to-Core 3 link, from the POP C perspective, the preferred path for a dynamic tunnel should go through Core 6, Core 3, and Core 2.



```
#Core-6 configuration on link towards Core-7
int ser 0/0
bandwidth 4000
mpls traffic-eng administrative-weight 55
ip rsvp bandwidth 2000 1000
```

The assignment of higher cost impacts only the CSPF calculations.


© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-23

In the example in the figure, the TE cost of the link between Core 1 and Core 7 is increased to 55, which makes links providing alternative paths more economical and more attractive for backup purposes.

Autoroute Forwarding Engine

This topic describes the command syntax to enable the autoroute forwarding engine.



The autoroute feature enables the headend routers to see the MPLS-TE tunnel as a directly connected interface and use it in their modified SPF computations. The MPLS-TE tunnel is used only for normal IGP route calculation (at the headend only) and is not included in any constraint-based path computation.

With the autoroute feature, the traffic tunnel does the following:

- Appears in the routing table
- Has an associated IP metric (cost equal to the best IGP metric to the tunnel endpoint) and is also used for forwarding the traffic to destinations behind the tunnel endpoint

The autoroute feature enables all the prefixes that are topologically behind the MPLS-TE tunnel endpoint (tailend) to be reachable via the tunnel itself (unlike static routing where only statically configured destinations are reachable via the tunnel).

Even with the autoroute feature, the tunnel itself is not used in link-state updates and the rest of the network does not have any knowledge of it.

Autoroute Forwarding Engine (Cont.)

Cisco.com

Changing Autoroute Metrics

```
#POP-C configuration
interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng autoroute metric absolute 1
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 250
 tunnel mpls traffic-eng path-option 1 explicit name Core3-2

interface Tunnel2
 ip unnumbered Loopback0
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng autoroute metric absolute 2
 tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng bandwidth 125
 tunnel mpls traffic-eng path-option 1 dynamic
!
ip explicit-path name Core3-2 enable
 next-address Core-3
 next-address Core-2
 next-address POP-A
```

Announce the presence of the tunnel to the routing protocol and select the primary path.

The metric determines the secondary path.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-25

Because the autoroute feature includes the MPLS-TE tunnel in the modified SPF path calculation, the metric of the tunnel plays a significant role. The cost of the tunnel is equal to the best IGP metric to the tunnel endpoint, regardless of the LSP path.

The tunnel metric is tunable using either relative or absolute metrics, which is the case in the example. When the routing process is selecting the best paths to the destination, the tunnel metric is compared to other existing tunnel metrics and to all the native IGP path metrics. The lower metric is better, and if the MPLS-TE tunnel has a lower metric, it is installed as a next hop to the respective destinations.

If there are tunnels with equal metrics, they are installed in the routing table and provide for load balancing. The load balancing is done proportionally to the configured bandwidth of the tunnel.

Autoroute Forwarding Engine (Cont.)

Cisco.com

Configuring Forwarding Adjacency

#POP-C configuration

```
interface Tunnel12
 ip unnumbered Loopback0
 isis metric 5 level-2
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng forwarding-adjacency
 ...
```

Announce the presence of the tunnel to the routing protocol locally and within the area.

#POP-A configuration

```
interface Tunnel21
 ip unnumbered Loopback0
 isis metric 5 level-2
 tunnel destination POP-C
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng forwarding-adjacency
 ...
```

For the point-to-point link between POP A and POP C to be announced into IS-IS, a tunnel must exist from POP A to POP C with forwarding adjacency enabled.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-26

To advertise a TE tunnel as a link in an IGP network, use the **tunnel mpls traffic-eng forwarding-adjacency** command in interface configuration mode. To disable the functionality, use the **no** form of this command.

- **tunnel mpls traffic-eng forwarding-adjacency {holdtime value}**
- **no tunnel mpls traffic-eng forwarding-adjacency**

Syntax Description

Parameter	Description
holdtime <i>value</i>	(Optional) Time in milliseconds (ms) that a TE tunnel waits after going down before informing the network. The range is 0 to 4,294,967,295 ms. The default value is 0.

Note You must configure a forwarding adjacency on two LSP tunnels bidirectionally, from A to B and B to A. Otherwise, the forwarding adjacency is advertised but not used in the IGP network.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- Traffic tunnels are unidirectional.
- A Network Administrator may set an Explicit Path that will be checked against the extended link-state database that contains information on available resources.
- Explicitly specified tunnels do not show up in routing tables and are used via static routes or via policy routing that sets a next-hop interface to the tunnel.
- If tunnel is explicitly established, automatic reoptimization does not occur.
- Use affinity bits to exclude links from the constraint-based SPF computation.
- Increase the TE cost on the link if it is desirable to exclude a certain link from any path computation.
- The autoroute feature enables the headend routers to see the MPLS-TE tunnel as a directly connected interface and use it in modified SPF computations.

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-27

References

For additional information, refer to this resource:

- Search the Cisco website at <http://www.cisco.com> for more information on the topics that were discussed in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Answer Key.

- Q1) The command that enables the MPLS Traffic Engineering tunnel feature on a device is the _____ command.
- Q2) The command that is used to specify the next IP address in an explicit path is the _____ command.
- Q3) Engineered tunnels can be used for IP routing only if the tunnel is explicitly specified for routing via _____ and _____.
- Q4) To trigger automatic reoptimization whenever an interface becomes operational, use the _____ router configuration command.
- Q5) Links can be excluded from the constraint-based SPF computation by using the _____ and _____ over which the tunnel should pass.
- Q6) The constraint-based path computation selects the path that the dynamic traffic tunnel will take, based on the administrative weight (TE cost), which is, by default, equal to the _____.
- Q7) With the autoroute feature enabled, the traffic tunnel can do which two of the following things? (Choose two.)
- A) appear in the routing table
 - B) have an associated IP metric
 - C) be used in link-state updates
 - D) allow the rest of the network to be aware of the traffic tunnel

Lesson Answer Key

- Q1) **mpls traffic-eng tunnels**
Relates to: Creating Dynamic Path Options for a Traffic Tunnel
- Q2) **next-address** (also **next-address strict** or **next-address loose**)
Relates to: Creating Explicit Path Options for a Traffic Tunnel
- Q3) via policy routing that sets a next-hop interface to the tunnel, and via static routes that point to the tunnel
Relates to: Assigning Traffic to Traffic Tunnels with Static Routes
- Q4) **mpls traffic-eng reoptimize events**
Relates to: Configuring Periodic and Event-Triggered Optimization
- Q5) resource class affinity bits of the traffic tunnel and resource class bits of the links
Relates to: Configuring Path Selection Restrictions
- Q6) IGP link metric (cost)
Relates to: Modifying Administrative Weight
- Q7) A, B
Relates to: Autoroute Forwarding Engine

Advanced MPLS-TE Link Protection

Overview

This lesson describes the advanced MPLS-TE commands for link protection. The configuration commands are accompanied by usage guidelines and sample setups.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes advanced MPLS-TE link protection. Upon completing this lesson, you will be able to do the following:

- Identify issues that are related to improving convergence time in an MPLS-TE network
- Configure a typical parallel tunnel solution
- Identify the drawbacks of using parallel tunnels
- Identify the characteristics of the Fast Reroute feature
- Identify the characteristics of link protection with the Fast Reroute capability
- Identify the characteristics of node protection with the Fast Reroute capability
- Configure fast rerouting

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course


Outline

This lesson includes these topics:

- Overview
- Improving Convergence Time
- Configuring Parallel Tunnels
- Drawbacks of Parallel Tunnels
- Fast Reroute
- Link Protection with FRR
- Node Protection with FRR
- Fast Reroute Configuration
- Lesson Summary
- Lesson Review

Improving Convergence Time

This topic describes issues that are related to improving the convergence time in an MPLS-TE network.



At first, the need for fast convergence is not obvious. Historically, applications were designed to recover from network outages. However, with today's increased usage of voice and video applications, network down time must be kept at a minimum. Even the same amount of time it takes a network to converge after a link failure can now be unacceptable.

In a simple case, a flapping link can result in headend routers being constantly involved in constraint-based computations. Because the time that elapses between link failure detection and the establishment of a new LSP path can cause delays for critical traffic, there is a need for alternative pre-established paths (backups).

Two tunnels are used between the same endpoints at the same time. The main requirement in this case is that preconfigured tunnels between the same endpoints must use diverse paths. In this scenario, as soon as the primary tunnel fails, the traffic is transitioned to the backup tunnel. The traffic is returned to the primary tunnel if conditions provide for the re-establishment of traffic.

Having two pre-established paths is the simplest form of MPLS-TE path protection. Several preparation steps must be taken in order for there to be effective switching between the tunnels. These steps include routing to the proper tunnel.

Configuring Parallel Tunnels

This topic describes the command syntax that is used to configure parallel tunnels.

Configuring Parallel Tunnels

Cisco.com

```
#POP-C configuration

interface Tunnel1
 ip unnumbered Loopback0
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 250
 tunnel mpls traffic-eng path-option 1 explicit name Core3-2

interface Tunnel2
 ip unnumbered Loopback0
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng bandwidth 125
 tunnel mpls traffic-eng path-option 1 dynamic
!
 ip explicit-path name Core3-2 enable
 next-address Core-3
 next-address Core-2
 next-address POP-A

#A pair of floating static routes for primary/backup selection

ip route POP-A1 mask tunnel 1 10
ip route POP-A1 mask tunnel 2 11
```

A lower priority is used, and due to the double counting of reservations, one half of the initial bandwidth is requested by the backup tunnel.

Option: Assign the traffic tunnel and link affinity bits to ensure the usage of the alternative to the primary path.

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-4

The example shows two configured tunnels: Tunnel 1 (following the LSP path Core 3 — Core 2 — POP A) and Tunnel 2 (using a dynamic path).

In the presence of two tunnels, static routing is deployed with two floating static routes pointing to the tunnels.

As soon as the primary tunnel (Tunnel 1) fails, the static route is gone and the traffic is transitioned to the secondary tunnel. The traffic is returned back to the primary tunnel if the conditions provide for the re-establishment of traffic.

Other options are spreading the load proportionally to the requested bandwidth (CEF mechanism) by load balancing, or having one group of static routes pointing to Tunnel 1 and another to Tunnel 2.

Configuring Parallel Tunnels (Cont.)

Cisco.com

#POP-C configuration

```
interface Tunnel12
 ip unnumbered Loopback0
 isis metric 5 level-2
 tunnel destination POP-A
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng forwarding-adjacency
 ...
```

Announce the presence of the tunnel to the routing protocol locally and within the area.

#POP-A configuration

```
interface Tunnel21
 ip unnumbered Loopback0
 isis metric 5 level-2
 tunnel destination POP-C
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng forwarding-adjacency
 ...
```

For the point-to-point link between POP A and POP C to be announced into IS-IS, a tunnel must exist from POP A to POP C with forwarding adjacency enabled.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-5

Any of the available methods can be used to advertise the tunnel. In this example, forwarding adjacency is used to advertise the TE tunnel as a link in an IGP network by means of the **tunnel mpls traffic-eng forwarding-adjacency** command in interface configuration mode.

- **tunnel mpls traffic-eng forwarding-adjacency {holdtime value}**
- **no tunnel mpls traffic-eng forwarding-adjacency**


Syntax Description

Parameter	Description
holdtime <i>value</i>	(Optional) Time in milliseconds (ms) that a TE tunnel waits after going down before informing the network. The range is 0 to 4,294,967,295 ms. The default value is 0.

You must configure a forwarding adjacency on two LSP tunnels bidirectionally, from A to B and from B to A. Otherwise, the forwarding adjacency is advertised but not used in the IGP network.

Drawbacks of Parallel Tunnels

This topic identifies the drawbacks of using parallel tunnels.



Because the time that elapses between link failure detection and the establishment of a new LSP path can cause delays for critical traffic, there is a need for alternative pre-established paths (backups). Therefore, two tunnels are used between the same endpoints at the same time.

Note Preconfigured tunnels between the same endpoints must use diverse paths.

As soon as the primary tunnel fails, traffic is transitioned to the backup tunnel. The traffic is returned to the primary tunnel if conditions provide for the re-establishment of traffic.

Note Having two pre-established paths is the simplest form of MPLS-TE path protection. Another option is to use the precomputed path only and establish the LSP path on demand. In the latter case, there is no overhead in resource reservations.

The figure shows two preconfigured tunnels: Tunnel 1 (LSP 1) is a primary tunnel, and Tunnel 2 (LSP 2) is a backup tunnel. Their physical paths are diverse.


The switchover to the backup tunnel is done at the headend as soon as the primary tunnel failure is detected (via RSVP or via IGP).

There is an obvious benefit to having a preconfigured backup tunnel. However, the solution presents some drawbacks as well:

- The backup tunnel requires all the mechanisms that the primary one requires. The labels must be allocated and bandwidth must be reserved for the backup tunnel as well.
- From the RSVP perspective, the resource reservations (bandwidth) are counted twice.

Fast Reroute

This topic describes the characteristics the of the Fast Reroute (FRR) capability.



In this case study, a company has decided to use only dynamic tunnels. A new high-speed link has been introduced between Core 1 and Core 6. This link influences CBR and native path selection and speeds up transport across the network. The result, however, is that this new high-speed link is now heavily used by traffic tunnels and may cause a serious disruption if it fails.

Fast Reroute

Cisco.com

Fast Reroute allows for temporary routing around a failed link or a failed node while the headend is rerouting the LSP:

- **Controlled by the routers with preconfigured backup tunnels around the protected link or node (link or node protection).**
- **The headend is notified of the failure through the IGP and through RSVP.**
- **The headend then attempts to establish a new LSP that bypasses the failure (LSP rerouting).**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—4-8

FRR is a mechanism for protecting MPLS-TE LSPs from link and node failures by locally repairing the LSPs at the point of failure, allowing data to continue to flow on them while their headend routers attempt to establish new end-to-end LSPs to replace them. FRR locally repairs the protected LSPs by rerouting them over backup tunnels that bypass failed links or nodes.

The headend is notified of the failure through the IGP and through RSVP. The headend then attempts to establish a new LSP that bypasses the failure (LSP rerouting).

Link Protection with FRR

This topic describes link protection using the FRR capability.



Backup tunnels that bypass only a single link of the LSP path provide link protection. They protect LSPs if a link along their path fails by rerouting the LSP traffic to the next hop (bypassing the failed link). These tunnels are referred to as next-hop backup tunnels because they terminate at the next hop of the LSP beyond the point of failure.

This process gives the headend of the tunnel time to re-establish the tunnel along a new, optimal route.

Link Protection with FRR (Cont.)

Cisco.com

“Link Down” Event

- The router, realizing the link is down:
 - Issues an IGP advertisement
 - Issues an RSVP message with session attribute flag 0x01=ON (do not break the tunnel; you may continue to forward packets during the reoptimization)
- In the event of a failure, an LSP is intercepted and locally rerouted using a backup tunnel.
 - Original LSP nested within protection LSP
 - Minimum disruption of LSP flow (under 50 ms - time to detect and switch)
- The headend is notified by RSVP PathErr and by IGP
 - Special flag in RSVP PathErr (reservation in place) indicates that the path states must not be destroyed, so the LSP flow is not interrupted.
 - The headend of the tunnel smoothly re-establishes the tunnel along a new route.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-10

Paths for LSPs are calculated at the LSP headend. Under failure conditions, the headend determines a new route for the LSP. Recovery at the headend provides for the optimal use of resources. However, because of messaging delays, the headend cannot recover as fast as making a repair at the point of failure.

To avoid packet flow disruptions while the headend is performing a new path calculation, the FRR option of MPLS-TE is available to provide protection from link or node failures (failure of a link or an entire router).

The function is performed by routers that are directly connected to the failed link, because they reroute the original LSP to a preconfigured tunnel and therefore bypass the failed path.


Note	In terms of forwarding, it can be said that the original LSP is nested within the protection LSP.

The reaction to a failure, with such a preconfigured tunnel, is almost instant. The local rerouting takes less than 50 ms, and a delay is caused only by the time that it takes to detect the failed link and to switch the traffic to the link protection LSP.

When the headend of the tunnel is notified of the path failure through the IGP or through RSVP; the headend then attempts to establish a new LSP.

Node Protection with FRR

This topic describes node protection using the FRR capability.



FRR provides node protection for LSPs.

Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop backup tunnels because they terminate at the node following the next-hop node of the LSP path, thereby bypassing the next-hop node. They protect LSPs if a node along their path fails by enabling the node upstream of the failure to reroute the LSPs and their traffic around the failed node to the next-next hop.

FRR supports the use of RSVP hellos to accelerate the detection of node failures. Next-next-hop backup tunnels also provide protection from link failures, because they bypass the failed link as well as the node.

Node Protection with FRR (Cont.)

Cisco.com

- **Router node fails; the router detects this failure by an “interface down” notification.**
 - It switches LSPs going out that interface onto their respective backup tunnels (if any).
- **RSVP hellos can also be used to trigger Fast Reroute.**
 - Messages are periodically sent to the neighboring router.
 - If no response is received, hellos declare that the neighbor is down.
 - Causes any LSPs going out that interface to be switched to their respective backup tunnels.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-12

When a router link or a neighboring node fails, the router often detects this failure by an “interface down” notification. On a Gigabit Switch Router (GSR) Packet over SONET (POS) interface, this notification is very fast. When a router notices that an interface has gone down, it switches LSPs going out that interface onto their respective backup tunnels (if any).

RSVP hellos can also be used to trigger FRR. If RSVP hellos are configured on an interface, messages are periodically sent to the neighboring router. If no response is received, the hellos declare that the neighbor is down. This action causes any LSPs going out that interface to be switched to their respective backup tunnels.

Fast Reroute Configuration

This topic describes the command syntax of the FRR configuration commands. The information in this topic is useful for when you want to add FRR protection to a network in which MPLS-TE LSPs are configured.

Fast Reroute Configuration

Cisco.com

- **Enabling Fast Reroute on LSPs**
- **Creating a Backup Tunnel to the Next Hop or to the Next-Next Hop**
- **Assigning Backup Tunnels to a Protected Interface**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-13

This section assumes that you want to add FRR protection to a network in which MPLS-TE LSPs are configured.

Before performing the configuration tasks in this topic, it is assumed that you have done the following tasks:

- Enabled MPLS-TE on all relevant routers and interfaces
- Configured MPLS-TE tunnels

Here are the tasks that are required to use FRR to protect LSPs in your network from link or node failures:

- Enabling FRR on LSPs
- Creating a backup tunnel to the next hop or to the next-next hop
- Assigning backup tunnels to a protected interface

Fast Reroute Configuration: Enabling Fast Reroute on LSPs

Cisco.com

```
router(config-if)#
```

```
mpls traffic-eng backup tunnel num
```

- Configures the interface to use a backup tunnel in the event of a detected failure on the interface

```
router(config-if)#
```

```
tunnel mpls traffic-eng fast-reroute
```

- Enables a traffic engineering tunnel to use a backup tunnel in the event of a link failure if a backup tunnel exists

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-14

mpls traffic-eng backup-path tunnel

To configure the interface to use a backup tunnel in the event of a detected failure on the interface, use the **mpls traffic-eng backup tunnel** interface command.

mpls traffic-eng backup-path tunnel *interface*

Syntax Description


Parameter	Description
<i>interface</i>	IP address that is associated with the given interface

tunnel mpls traffic-eng fast-reroute

To enable an MPLS-TE tunnel to use a backup tunnel in the event of a link failure if a backup tunnel exists, use the **tunnel mpls traffic-eng fast-reroute** interface configuration command.

Fast Reroute Configuration: Link Protection Example

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-15

The example in the figure lists both sets of configuration commands that are needed when you are provisioning a backup for a link over a tunnel:

- Configuration of Core 6 to provide a backup tunnel around the protected link
- POP C configuration of a tunnel and FRR assignment

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **Preconfigured tunnels between the same endpoints must use diverse paths.**
- **Traffic is returned back to the primary tunnel if the conditions provide for the re-establishment of traffic.**
- **To advertise a TE tunnel as a link in an IGP network, the tunnel mpls traffic-eng forwarding-adjacency command is used.**
- **Fast Reroute provides link protection to LSPs by establishing a backup LSP tunnel for a troubled link.**
- **tunnel mpls traffic-eng fast-reroute is the Cisco IOS command that enables an MPLS-TE tunnel to use a backup tunnel in the event of a link failure if backup tunnel exists.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-16

References

For additional information, refer to this resource:

- **MPLS Traffic Engineering (TE): Link and Node Protection, with RSVP Hellos Support**
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/products_feature_guide_09186a008012db55.html#1032351

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The requirement for two tunnels between the same endpoints is that they must use _____ paths.
- Q2) If conditions are corrected and provide for the re-establishment of traffic, the traffic is _____ to the primary tunnel.
- Q3) If you do not configure forwarding adjacency on two LSP tunnels bidirectionally, from A to B and from B to A, the _____ is advertised but not used in the IGP network.
- Q4) _____ provides link protection to LSPs by establishing a backup LSP tunnel for the troubled link.
- Q5) To enable an MPLS-TE tunnel to use a backup tunnel in the event of a link failure if a backup tunnel exists, use the _____ command.

Lesson Answer Key

- Q1) diverse
Relates to: Improving Convergence Time
- Q2) returned
Relates to: Drawbacks of Parallel Tunnels
- Q3) forwarding adjacency
Relates to: Configuring Parallel Tunnels
- Q4) Fast Reroute
Relates to: Link Protection with FRR
- Q5) **tunnel mpls traffic-eng fast-reroute**
Relates to: Fast Reroute Configuration

Advanced MPLS-TE Bandwidth Control

Overview

This lesson describes the advanced MPLS-TE commands for bandwidth control. The configuration commands are accompanied by usage guidelines and sample setups.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes advanced MPLS-TE bandwidth control. Upon completing this lesson, you will be able to do the following:

- Identify the characteristics of bandwidth reservation with autobandwidth
- Configure the autobandwidth feature
- Identify the characteristics of the DiffServ-aware service
- Configure the DiffServ-aware service

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course


Outline

This lesson includes these topics:

- Overview
- Bandwidth Reservation with Autobandwidth
- Autobandwidth Configuration Commands
- DiffServ-Aware Service
- DiffServ-Aware Service Configuration Commands
- Lesson Summary
- Lesson Review

Bandwidth Reservation with Autobandwidth

This topic describes the characteristics of bandwidth reservation using autobandwidth.



After initial TE is complete, network administrators may need an effective way to continually adjust tunnel routes and bandwidth reservations without doing any redesigning.

Cisco IOS software has a new feature called “autobandwidth” that measures utilization averages and dynamically adjusts tunnel bandwidth reservations to meet actual application resource requirements.

This powerful feature creates "self-tuning tunnels," which free network administrators from many of the daily hands-on management tasks that are necessary with other TE techniques.

Bandwidth Reservation with Autobandwidth (Cont.)

Cisco.com

Traffic engineering automatic bandwidth feature adjusts the bandwidth allocation for TE tunnels based on their measured traffic load:

- **This feature periodically changes tunnel bandwidth reservation based on traffic out tunnel.**
- **The average output rate is sampled for each tunnel.**
- **The allocated bandwidth is periodically adjusted to be the largest sample for the tunnel since the last adjustment.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-4

TE automatic bandwidth adjustment provides the means to automatically adjust the bandwidth allocation for TE tunnels based on their measured traffic load.

TE autobandwidth samples the average output rate for each tunnel that is marked for automatic bandwidth adjustment. For each marked tunnel, it periodically (for example, once per day) adjusts the allocated bandwidth of the tunnel to be the largest sample for the tunnel since the last adjustment.


The frequency with which tunnel bandwidth is adjusted and the allowable range of adjustments are configurable on a per-tunnel basis. In addition, the sampling interval and the interval over which to average the tunnel traffic to obtain the average output rate are user-configurable on a per-tunnel basis.

The benefit of the autobandwidth feature is that it makes it easy to configure and monitor the bandwidth for MPLS-TE tunnels. If automatic bandwidth is configured for a tunnel, TE automatically adjusts the bandwidth of the tunnel.

The automatic bandwidth adjustment feature treats each tunnel for which it has been enabled independently. In other words, it adjusts the bandwidth for each such tunnel according to the adjustment frequency that is configured for the tunnel and the sampled output rate for the tunnel since the last adjustment, without regard for any adjustments previously made or pending for other tunnels.

Bandwidth Reservation with Autobandwidth (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-5

The diagram shows the load on the tunnel and intervals of measurement. The input and output rates on the tunnel interfaces are averaged over a predefined interval (load interval). In the example, the interval is the previous 5 minutes.

The automatic bandwidth adjustments are done periodically, for example, once per day. For each tunnel for which automatic bandwidth adjustment is enabled, the platform maintains information about sampled output rates and the time remaining until the next bandwidth adjustment.

When the adjustments are done, the currently allocated bandwidth (shown as horizontal solid lines in the diagram) is reset to the maximum of one of the following:

- The largest average rate that has been sampled since the last bandwidth adjustment
- The configured maximum value

If the new bandwidth is not available, the previously allocated bandwidth is maintained.

Autobandwidth Configuration Commands

This topic describes the commands and syntax that are necessary to configure the autobandwidth feature.

Autobandwidth Configuration Commands

Cisco.com

```
router(config)#  
mpls traffic-eng auto-bw timers [frequency sec]
```

- Enables automatic bandwidth adjustment and specifies the frequency of the sampling interval, in seconds

```
router(config-if)#  
tunnel mpls traffic-eng auto-bw {frequency seconds}  
{max-bw kbs} {min-bw kbs}
```

- Specifies the frequency of adjustments, minimum and maximum automatic bandwidth allocations, in kilobits per second, that can be applied to the tunnel

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-6

For every MPLS-TE tunnel that is configured for Cisco MPLS Autobandwidth, the average output rate is sampled based on various configurable parameters. The tunnel bandwidth is then readjusted automatically based on the largest average output rate that was noticed during a certain interval or a configured maximum bandwidth value.

The Cisco MPLS Autobandwidth allocator monitors the X minutes (default = 5 min) average counter, keeping track of the largest average over some configurable interval Y (default = 24 hours), and then readjusting a tunnel bandwidth based upon the largest average for that interval.

The Cisco MPLS Autobandwidth feature is implemented with the following commands:

- **mpls traffic-eng auto-bw timers [frequency seconds]**
This is a global command to define the interval during which to sample the X average for each tunnel. By default, this value is set to 300 sec (5 min).
- **clear mpls traffic-eng auto-bw timers**
This command is used to clear the timers that have been defined by the previous command.
- **tunnel mpls traffic-eng auto-bw {frequency seconds} {max-bw kbs} {min-bw kbs}**
By default, the frequency is 24 hours.

The last command controls the Y interval between bandwidth readjustments and is tunnel-specific. Setting the **max-bw** limits the maximum bandwidth that a tunnel can adjust to. Similarly, setting the **min-bw** provides the smallest amount that the bandwidth can adjust to. When both **max-bw** and **min-bw** are specified, the tunnel bandwidth will remain between these values.

Autobandwidth Configuration Commands (Cont.)

Cisco.com

#POP-A configuration

```
mpls traffic-eng auto-bw timers frequency 300

interface tunnel1
ip unnumbered loopback 0
tunnel destination POP-C
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng bandwidth 2500
tunnel mpls traffic-eng priority 1 1
tunnel mpls traffic-eng path-option 1 dynamic
tunnel mpls traffic-eng auto bw frequency 3600 max-bw 3000 min-bw 1000
```

Sample (every 5 min) the output rate for tunnels that have been configured for automatic bandwidth adjustment.

Initial tunnel bandwidth, which will be adjusted by the autobandwidth mechanism.

Specifies the minimum and maximum automatic bandwidth allocations, in kilobits per second, that can be applied to the tunnel and adjusted via RSVP (every hour).

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-7

The example in the figure shows the setting of MPLS traffic-engineered tunnels that can actually tune their own bandwidth requirements to increase or decrease their RSVP reservations, as warranted by changing network conditions.

When readjusting bandwidth constraint on a tunnel, a new Resource Reservation Protocol-Traffic Engineering (RSVP-TE) path request is generated, and if the new bandwidth is not available, the last good LSP will continue to be used.

The network experiences no traffic interruptions.

DiffServ-Aware Service

This topic describes the characteristics of the differentiated-services-aware (DiffServ-aware) service.

DiffServ-Aware Service

Cisco.com

Further enhancement to satisfy the stricter requirements:

- **Separate pools to allow for different limits for the amount of traffic admitted on any given link:**
 - **One subpool, for tunnels that carry traffic requiring strict bandwidth guarantees or delay guarantees**
 - **The global pool, for best-effort or DiffServ traffic**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-8

Providing Strict QoS Guarantees Using DiffServ-Aware TE Subpool Tunnels

A tunnel using the subpool bandwidth can satisfy the stricter QoS requirements if you do all of the following:

1. Select a queue—or in DiffServ terminology, select a per-hop behavior (PHB)—to be used exclusively by the strict guarantee traffic..”

If delay or jitter guarantees are sought, the DiffServ Expedited Forwarding PHB (EF PHB) is used. On the Cisco 12000 series routers, that is the “low latency queue.” You must configure the bandwidth of the queue to be at least equal to the bandwidth of the subpool.

If only bandwidth guarantees are sought, the DiffServ Assured Forwarding PHB (AF PHB) is used. On the Cisco 12000 series routers, you use one of the existing Modified Deficit Round Robin (MDRR) queues.

2. Ensure that the guaranteed traffic that is sent through the subpool tunnel is placed in the GB queue at the outbound interface of every tunnel hop, and that no other traffic is placed in this queue.

You do this by marking the traffic that enters the tunnel with a unique value in the mpls exp bits field, and steering only traffic with that marking into the GB queue.

3. Ensure that this GB queue is never oversubscribed; that is, see that no more traffic is sent into the subpool tunnel than the GB queue can handle.

You do this by rate-limiting the guaranteed traffic before it enters the subpool tunnel. The aggregate rate of all traffic entering the subpool tunnel should be less than or equal to the bandwidth capacity of the subpool tunnel. Excess traffic can be dropped (in the case of delay or jitter guarantees) or can be marked differently for preferential discard (in the case of bandwidth guarantees).

4. Ensure that the amount of traffic that is entering the GB queue is limited to an appropriate percentage of the total bandwidth of the corresponding outbound link. The exact percentage to use depends on several factors that can contribute to accumulated delay in your network: your quality of service (QoS) performance objective, the total number of tunnel hops, the amount of link fan-in along the tunnel path, the burstiness of the input traffic, and so on.

You do this by setting the subpool bandwidth of each outbound link to the appropriate percentage of the total link bandwidth (that is, by adjusting the *z* parameter of the **ip rsvp bandwidth** command).

Providing Differentiated Service Using DiffServ-Aware TE Global Pool Tunnels

You can configure a tunnel using global pool bandwidth to carry best-effort as well as several other classes of traffic. Traffic from each class can receive DiffServ service if you do all of the following:

1. Select a separate queue (a distinct DiffServ PHB) for each traffic class. For example, if there are three classes (gold, silver, and bronze), there must be three queues (DiffServ AF2, AF3, and AF4).
2. Mark each class of traffic using a unique value in the MPLS experimental bits field (for example, gold = 4, silver = 5, bronze = 6).
3. Ensure that packets marked as “gold” are placed in the gold queue, “silver” in the silver queue, and so on. The tunnel bandwidth is set based on the expected aggregate traffic across all classes of service.

To control the amount of DiffServ tunnel traffic that you intend to support on a given link, adjust the size of the global pool on that link.

DiffServ-Aware Service (Cont.)

Cisco.com

- **An extension of MPLS-TE (basically, the signaling feature).**
- **Allows CBR of tunnels to use more restrictive bandwidth constraints.**
 - **Dual-bandwidth pool traffic engineering**
 - **Global pool**
 - **Subpool**
- **DiffServ ensures that bandwidth for DiffServ-aware TE tunnels is set aside on each link in the network.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-9

DiffServ-Aware Traffic Engineering (DS-TE) extends the current MPLS-TE capabilities to introduce the awareness of a particular class of traffic, which is the guaranteed bandwidth traffic. DS-TE enables the service provider to perform a separate admission control and route computation of the guaranteed bandwidth traffic. DS-TE is another signaling feature of IGP and RSVP.

With only a single bandwidth pool on the link in traditional MPLS-TE, when the bandwidth is reserved for the tunnel, the traffic within the tunnel is considered as a single class. For example, when voice and data are intermixed within the same tunnel, the QoS mechanisms cannot ensure better service for the voice. Usually, class-based weighted fair queuing (CBWFQ) can be performed for the tunnel.

The idea behind DS-TE is to *guarantee* the bandwidth for DS-TE tunnels across the network. For critical applications (for example, voice), a separate DS-TE tunnel is created. Thus, two bandwidth pools are used, one for traditional MPLS-TE tunnels and one for DS-TE tunnels. The DiffServ QoS mechanisms (low latency queuing, or LLQ) ensure that bandwidth is dedicated for DS-TE tunnels. In the initial phase, the DS-TE supports a single class of bandwidth. It is expected that subsequent phases of DS-TE will provide new capabilities, such as the support of multiple classes of bandwidth and the dynamic reprogramming of queuing or scheduling mechanisms.

DiffServ-Aware Service (Cont.)

Cisco.com

MPLS-TE has the following extensions for DS-TE:

- **Two types of bandwidth limits per interface**
- **IGP advertises both types of bandwidth**
- **Tunnel configured with appropriate bandwidth type**
- **Appropriate bandwidth type considered in path calculations**
- **Tunnel signaled (via RSVP) with the appropriate bandwidth type**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-10

DS-TE tunnels are similar to regular TE tunnels. To support DS-TE, the following modifications to regular MPLS-TE mechanisms have been made:


- There are two types of bandwidth per each link in the network (two bandwidth pools: the global pool and the subpool).
- These bandwidths are both announced in the link-state updates that carry resource information.
- The traffic tunnel parameters include the bandwidth type that the tunnel will use.
- The constraint-based path calculation (PCALC) is done with respect to the type of the bandwidth that the tunnel requires. In RSVP messages, it is always indicated whether the LSP to be set up is a regular MPLS-TE tunnel or a DS-TE tunnel. Intermediate nodes perform admission control and bandwidth allocation (“locking” for the DS-TE) for the appropriate bandwidth pool.

DiffServ-Aware Service (Cont.)

Cisco.com

DS-TE Dual-Bandwidth Pools

- Global pool tracks the true available bandwidth (takes into account the bandwidth used by both types of tunnels).
- Subpool keeps track only of the constraint for the DS-TE.



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-11

On each link in the network, two bandwidth pools are established:

- Global (main) pool, which keeps track of the true available bandwidth. The pool takes into account the bandwidth that is used by both of the tunnels.
- Subpool (DS-TE), which tracks only the bandwidth for the DS-TE tunnels.

The bandwidths that are specified for both pools are independent of the actual physical bandwidth of the link (providing for oversubscription). The same situation applies to traditional MPLS-TE with one bandwidth pool.

The only constraint for the two pools is that the bandwidth of the subpool (dedicated to DS-TE tunnels) must not exceed the bandwidth of the global pool.

DiffServ-Aware Service Configuration Commands

This topic describes the commands and syntax that are necessary to configure the guaranteed bandwidth service.

DiffServ-Aware Service Configuration

Cisco.com

```
router(config-if)#  
ip rsvp bandwidth interface-kbps single-flow-kbps [sub-pool kbps]
```

- The sum of bandwidth used by all tunnels on this interface cannot exceed *interface-kbps*, and the sum of bandwidth used by all subpool tunnels cannot exceed *sub-pool kbps*.

```
router(config-if)#  
tunnel mpls traffic-eng bandwidth {sub-pool | [global]} bandwidth
```

- Configures bandwidth of the tunnel and assign it either to the subpool or the global pool.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-12

ip rsvp bandwidth

To enable RSVP for IP on an interface, use the **ip rsvp bandwidth** interface configuration command.

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

Syntax Description

Parameter	Description
<i>interface-kbps</i>	Amount of bandwidth (in kbps) on interface to be reserved. The range is 1 to 10000000.
<i>single-flow-kbps</i>	Amount of bandwidth (in kbps) that is allocated to a single flow (ignored in DS-TE). The range is 1 to 10000000.
sub-pool <i>kbps</i>	Amount of bandwidth (in kbps) on interface to be reserved to a portion of the total. The range is from 1 to the value of <i>interface-kbps</i> .

tunnel mpls traffic-eng bandwidth

To configure bandwidth required for a Multiprotocol Label Switching (MPLS) traffic engineering tunnel, use the **tunnel mpls traffic-eng bandwidth** command in interface configuration mode. To disable this bandwidth configuration, use the **no** form of this command.

tunnel mpls traffic-eng bandwidth {sub-pool | [global]} *bandwidth*

Syntax Description

Parameter	Description
sub-pool	(Optional) Indicates a subpool tunnel.
global	(Optional) Indicates a global pool tunnel. Entering this keyword is not necessary, for all tunnels are "global pool" in the absence of the keyword sub-pool .
<i>bandwidth</i>	The bandwidth, in kilobits per second, that is set aside for the MPLS-TE tunnel. The range is between 1 and 4294967295.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **Autobandwidth** measures utilization averages and dynamically adjusts tunnel bandwidth reservations to meet actual application resource requirements.
- **The Autobandwidth feature is implemented with the following commands:**
 - `mpls traffic-eng auto-bw timers [frequency seconds]`
 - `tunnel mpls traffic-eng auto-bw`
- **The global pool is used for best-effort or DiffServ traffic.**
- **The `ip rsvp bandwidth` command is used to enable RSVP for IP on an interface.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-13

References

For additional information, refer to this resource:

- Search the Cisco website at <http://www.cisco.com> for more information on the topics that were discussed in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The Cisco IOS feature that measures utilization averages and dynamically adjusts tunnel bandwidth reservations is _____.
- Q2) The Cisco IOS command that defines the interval during which to sample the X average for each tunnel is the _____ command.
- Q3) The _____ is used for tunnels that carry traffic that requires strict bandwidth guarantees or delay guarantees.
- Q4) The Cisco IOS command that enables RSVP for IP on an interface is the _____ command.

Lesson Answer Key

Q1) autobandwidth

Relates to: Bandwidth Reservation with Autobandwidth

Q2) **mpls traffic-eng auto-bw timers** [*frequency seconds*]

Relates to: Autobandwidth Configuration Commands

Q3) subpool

Relates to: DiffServ-Aware Service

Q4) **ip rsvp bandwidth**

Relates to: DiffServ-Aware Service Configuration Commands

MPLS-TE Interarea Support

Overview

This lesson describes the advanced MPLS-TE topics and commands for interarea support. The configuration commands are accompanied by usage guidelines and sample setups.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes MPLS-TE interarea support. Upon completing this lesson, you will be able to do the following:

- Describe the characteristics of MPLS-TE without interarea support
- Describe how MPLS-TE provides interarea support
- Describe MPLS-TE operations in an interarea environment
- Describe interarea MPLS-TE restoration

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course


Outline

This lesson includes these topics:

- Overview
- MPLS-TE Without Interarea Support
- Interarea MPLS-TE
- Interarea MPLS-TE Detailed Operation
- Interarea MPLS-TE Restoration
- Lesson Summary
- Lesson Review

MPLS-TE Without Interarea Support

This topic describes the characteristics of MPLS-TE without interarea support.



Without the capability of interarea TE support, there are only two possibilities for designing a traffic-engineered network:


- Restrict the whole network to a single area
- Make a tunnel per area and use autoroute and the IGP to find the path from the source to the destination

Both designs have drawbacks. In the first case a big single-area network increases the amount of routing protocol overhead; in the second case any bandwidth-related reservation is very difficult to engineer and implement.

Note IS-IS does not make use of an Area Border Router (ABR) and an area 0.0.0.0 like OSPF. For IS-IS any Level 1-2 router has the same role as an ABR, and the network of all Level 2 routers forms the backbone area.

Interarea MPLS-TE

This topic describes how MPLS-TE provides interarea support.



Without the capability of interarea TE support, there are only two possibilities to design a TE network.

The MPLS-TE interarea tunnels feature allows you to establish MPLS-TE tunnels that span IGP areas and levels, removing the restriction that had required the tunnel headend and tailend routers to both be in the same area.


The IGP can be either IS-IS or OSPF.

To configure an interarea tunnel, you specify on the headend router a loosely routed explicit path for the tunnel LSP that identifies each ABR that the LSP should traverse using the **next-address loose** command.

The headend router and the ABRs, along the specified explicit path, expand the loose hops, each computing the path segment to the next ABR or tunnel destination.

Interarea MPLS-TE (Cont.)

Cisco.com



```
interface Tunnel0
  ip unnumbered Loopback0
  no ip directed-broadcast
  tunnel destination R12-loopback-address
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng priority 3 3
  tunnel mpls traffic-eng bandwidth 10000
  tunnel mpls traffic-eng path-option 1 explicit name path1

ip explicit-path name path1 enable
next-address loose R3
next-address loose R11
next-address loose R12
```

The configured addresses no longer need to be directly connected.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-5

Configuring Explicit Paths

To configure explicit paths, use the following commands, beginning in global configuration mode:

- Step 1** Router(config)# **ip explicit-path identifier number enable**
Enters the subcommand mode for IP explicit paths and creates or modifies the specified path.
- Step 2** Router(cfg-ip-expl-path)# **next-address loose A.B.C.D**
Specifies the next loose IP address in the explicit path. Each ABR that the path must traverse should be specified in a **next-address loose** command.

Syntax

To specify the next IP address in an explicit path, use the **next-address** command in IP explicit path configuration mode.

next-address [**loose** | **strict**] *A.B.C.D*

To remove an address from an explicit path, use the **no index** configuration command in IP explicit path configuration mode.

Parameter	Description
loose	Specifies that the previous address (if any) in the explicit path need not be directly connected to A.B.C.D, and that the router is free to determine the path from the previous address (if any) to A.B.C.D.
strict	Specifies that the previous address (if any) in the explicit path must be directly connected to A.B.C.D.
<i>A.B.C.D</i>	Next IP address in the explicit path.


Usage Guidelines

To configure an explicit path where the path may include only the addresses that are specified, indicate each address in sequence by using the **next-address** command without the **loose** optional keyword.

To configure an interarea TE tunnel, you must configure the tunnel path options as loose explicit paths, for which each ABR that is traversed by the tunnel LSP is specified as a loose hop by entering the **loose** option of the **next-address** command.

Interarea MPLS-TE Detailed Operation

This topic describes MPLS-TE operations in an interarea environment.



For explicit paths to be used for TE tunnels within an IGP area, you may specify a combination of both loose and strict hops via **next-address** commands.

Note Because any router in an interarea environment has the necessary topology and TE information only for its own area, for OSPF you need to specify each ABR as a loose next hop in the explicit path configuration. For IS-IS you need to specify all Level 1-2 routers on the path as loose next hops in the explicit path configuration.

Interarea MPLS-TE Restoration

This topic describes interarea MPLS-TE restoration.

Interarea MPLS-TE Restoration

Cisco.com

- **Once the link or node failure has been detected anywhere along the LSP, the TE LSP is rerouted at the headend level.**
- **Once the headend has received the Resv Tear or Path Error:**
 - **The headend will mark the corresponding path option as “not valid” for a “holddown” time.**
 - **Will go to the next path option.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-7

In an interarea TE environment, the headend has *no* information about topology details of other areas. Therefore, the headend, whenever the path is broken, cannot rely on a successful recalculation of the path.

The headend will also not get any link-state information from the IGP when a link outside its area has gone down.

Resv Tear and Path Error are the only messages that the headend gets signaled about when there is a link or node failure along the path.

Upon receipt of such messages, the headend will mark the corresponding path option as invalid for a certain amount of time and will try to establish a new path using the next available path option.

Interarea MPLS-TE Restoration (Cont.)

Cisco.com


```
Interface tunnel 1
...
tunnel mpls traffic-eng path-option 1
explicit name path1
tunnel mpls traffic-eng path-option 2
explicit name path2
```

```
ip explicit-path name path1 enable
next-address loose R3
next-address loose R11
next-address loose R12
```

```
ip explicit-path name path2 enable
next-address loose R4
next-address loose R10
next-address loose R12
```

Suppose link R3-R6 fails:

- R3 will send a Resv Tear or Path Error to R0:
 - With the current code the headend would keep on trying the same path option because R0 does not have any insight about the link failure in area 0.0.0.0.
- R0 will mark the path option 1 as invalid for a certain amount of time (holddown).
- R0 will try the path option 2.



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-8

When there are choices for which ABRs to use, explicit paths are recommended, each of which identifies a different sequence of ABRs in order to increase the likelihood that a path to the destination will be found. Thus, a network administrator will typically define more than one path option when creating interarea tunnels.

Interarea MPLS-TE Restoration (Cont.)

Cisco.com

A possible final configuration could be:


```
interface tunnel 1
tunnel mpls traffic-eng path-option 1
explicit name path1
tunnel mpls traffic-eng path-option 2
explicit name path2
tunnel mpls traffic-eng path-option 3
explicit name path3
tunnel mpls traffic-eng path-option 4
explicit name path4
```

```
ip explicit-path name path1 enable
next-address loose R3
next-address loose R11
next-address loose R12
```

```
ip explicit-path name path2 enable
next-address loose R3
next-address loose R10
next-address loose R12
```

```
ip explicit-path name path3 enable
next-address loose R4
next-address loose R10
next-address loose R12
```

```
ip explicit-path name path4 enable
next-address loose R4
next-address loose R11
next-address loose R12
```



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-9

In this example, in which there are redundant ABRs at the border of each area, the network administrator has optimally configured four different explicit paths using a combination of ABRs.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **Without interarea TE support there are only two possibilities to design a TE network:**
 - **Restrict network to one area**
 - **Make a tunnel per area**
- **Interarea TE allows the tunnel headend and tailend routes to be in different areas.**
- **Use the next-address loose or strict command to specify the next IP address in an explicit path.**
- **When there are choices for which ABRs are to be used, multiple explicit paths are recommended.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-10

References

For additional information, refer to this resource:

- Search the Cisco website at <http://www.cisco.com> for more information on the topics that were discussed in this lesson.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) What are two design drawbacks when interarea MPLS-TE is not supported? (Choose two.)
- A) RSVP is not available.
 - B) The number of nodes in an LSP is limited.
 - C) Bandwidth-related reservation is difficult.
 - D) Reliance on a single network increases routing protocol overhead.
- Q2) MPLS-TE interarea support allows tunnels to span _____ areas and levels.
- Q3) The command that specifies that the previous address in the explicit path need not be directly connected and that the router is free to determine the path from the previous address is the _____ command.
- Q4) When there are choices for which ABRs to use, _____ are recommended.

Lesson Answer Key

- Q1) C, D
Relates to: MPLS-TE Without Interarea Support
- Q2) IGP
Relates to: Interarea MPLS-TE
- Q3) **next-address loose**
Relates to: Interarea MPLS-TE
- Q4) explicit paths
Relates to: Interarea MPLS-TE Restoration

Monitoring and Troubleshooting MPLS-TE

Overview

This lesson describes the monitoring and troubleshooting commands of Cisco IOS software that support MPLS-TE. The troubleshooting guidelines are explained by means of sample command outputs.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS-TE on Cisco IOS platforms.

Objectives

This lesson describes the monitoring and troubleshooting of MPLS-TE implementations. Upon completing this lesson, you will be able to do the following:

- Identify the command syntax that is used to monitor MPLS-TE tunnels
- Identify the command syntax that is used to monitor MPLS-TE
- Identify the command syntax that is used to monitor MPLS-TE resource usage on intermediate devices
- Identify the command syntax that is used to troubleshoot device and link setup

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” module of this course

Outline

This lesson includes these topics:

- Overview
- Monitoring MPLS-TE Tunnels
- Monitoring MPLS-TE
- Monitoring MPLS-TE Resource Usage on Intermediate Devices
- Troubleshooting Device and Link Setup
- Lesson Summary
- Lesson Review

Monitoring MPLS-TE Tunnels

This topic describes the command syntax that is used to monitor MPLS-TE tunnels.

Monitoring MPLS-TE Tunnels

Cisco.com

```
router#  
show mpls traffic-eng tunnels tunnel_interface [brief]
```

- Shows summary information about a tunnel

```
router#  
show mpls traffic-eng tunnels [parameters]
```

- Shows detailed information about a tunnel

```
router#  
show mpls traffic-eng autoroute
```

- Shows tunnels that are announced to IGP, including interface, destination, and bandwidth

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-3

show mpls traffic-eng tunnels

To show information about tunnels, use the **show mpls traffic-eng tunnels** command.

show mpls traffic-eng tunnels *tunnel-interface* [brief]

show mpls traffic-eng tunnels

[*destination address*]

[*source-id* {*number* | *ip-address* | *ip-address number*}]

[*role* {**all** | **head** | **middle** | **tail** | **remote**}]

[**up** | **down**]

[*name string*]

[**suboptimal constraints** {**none** | **current** | **max**}]

[**[interface in** *physical-interface*] [**interface out** *physical-interface*] | [**interface** *physical-interface*] [**brief**]

Parameter	Description
<i>tunnel-interface</i>	Displays information for the specified tunneling interface.
brief	(Optional) Displays the information in brief format.
destination <i>address</i>	(Optional) Restricts the display to tunnels that are destined to the specified IP address.
source-id	(Optional) Restricts the display to tunnels with a matching source IP address or tunnel number.

Parameter	Description
<i>number</i>	(Optional) Tunnel number.
<i>ip-address</i>	(Optional) Source IP address.
<i>ip-address number</i>	(Optional) Source IP address and tunnel number.
role	(Optional) Restricts the display to tunnels with the indicated role (all , head , middle , tail , or remote).
all	(Optional) Displays all tunnels.
head	(Optional) Displays tunnels with their heads at this router.
middle	(Optional) Displays tunnels with their midpoints at this router.
tail	(Optional) Displays tunnels with their tails at this router.
remote	(Optional) Displays tunnels with their heads at another router; this display is a combination of the middle and tail keyword values.
up	(Optional) Displays tunnels if the tunnel interface is up. Tunnel midpoints and tails are typically up or not present.
down	(Optional) Displays tunnels that are down.
name <i>string</i>	(Optional) Displays tunnels with the specified name. The tunnel name is derived from the interface description, if specified; otherwise, it is the interface name. The tunnel name is included in the signaling message so that it is available at all hops.
suboptimal constraints none	(Optional) Displays tunnels whose path metric is greater than the shortest unconstrained path. Selected tunnels have a longer path than the shortest path of the IGP.
suboptimal constraints current	(Optional) Displays tunnels whose path metric is greater than the current shortest path, constrained by the configured tunnel options. Selected tunnels would have a shorter path if they were reoptimized immediately.
suboptimal constraints max	(Optional) Displays tunnels whose path metric is greater than the current shortest path, constrained by the configured tunnel options, and considering only the capacity of the network. Selected tunnels would have a shorter path if no other tunnels were consuming network resources.
interface in <i>physical-interface</i>	(Optional) Displays tunnels that use the specified input interface.
interface out <i>physical-interface</i>	(Optional) Displays tunnels that use the specified output interface.
interface <i>physical-interface</i>	(Optional) Displays tunnels that use the specified interface as an input or output interface.
brief	(Optional) Specifies one line per tunnel.

Example: Usage of the show mpls traffic-eng tunnel brief Command

The following example shows output from the **show mpls traffic-eng tunnel brief** command:

```
Router#show mpls traffic-eng tunnel brief
```

```
Router#show mpls traffic-eng tunnel brief
```

Signaling Summary:

LSP Tunnels Process: running

RSVP Process: running

Forwarding: enabled

Periodic reoptimization: every 180 seconds, next in 108 seconds

TUNNEL NAME STATE	DESTINATION	STATUS
tagsw-r4_t1	10.0.0.11	admin-down down
tagsw-r4_t10011	10.0.0.11	up up
...		
al7500-sw12_t20004	10.0.0.4	signaled up

Displayed 16 (of 16) heads, 0 (of 0) midpoints, 1 (of 1) tails

show mpls traffic-eng autoroute

To show tunnels that are announced to IGP, including interface, destination, and bandwidth, use the **show mpls traffic-eng autoroute** command in privileged EXEC mode.

show mpls traffic-eng autoroute

Usage Guidelines

The SPF, or next-hop, calculation of the IGP has been modified to incorporate TE tunnels. This command shows which tunnels are currently being used by the IGP in its SPF, or next-hop, calculation (tunnels that are up and have autoroute configured).

Example: Usage of the show mpls traffic-eng autoroute Command

The following example shows output from the **show mpls traffic-eng autoroute** command.

Note that the list of tunnels is organized by destination. All tunnels to a destination will carry a share of the traffic to that destination.

```
Router#show mpls traffic-eng autoroute
MPLS TE autorouting enabled
  destination 0002.0002.0002.00 has 2 tunnels
    Tunnel1021 (traffic share 10000, nexthop 2.2.2.2, absolute metric
11)
    Tunnel1022 (traffic share 3333, nexthop 2.2.2.2, relative metric -
3)
  destination 0003.0003.0003.00 has 2 tunnels
    Tunnel1032 (traffic share 10000, nexthop 3.3.3.3)
    Tunnel1031 (traffic share 10000, nexthop 3.3.3.3, relative metric -
1)
```

Monitoring MPLS-TE

This topic describes the command syntax that is used to monitor MPLS-TE.

Monitoring MPLS-TE

Cisco.com

```
router#  
show mpls traffic-eng link-management summary  
[interface-name]
```

- Shows a summary of link management information

```
router#  
show mpls traffic-eng link-management interfaces  
[interface-name]
```

- Shows interface resource and configuration information

```
router#  
show mpls traffic-eng link-management bandwidth-  
allocation [interface-name]
```

- Shows current local link bandwidth allocation

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-4

show mpls traffic-eng link-management summary

To show a summary of link management information, use the **show mpls traffic-eng link-management summary** command in EXEC mode.

show mpls traffic-eng link-management summary [*interface name*]

Example: Usage of the show mpls traffic-eng link-management summary Command

The following example shows output from the **show mpls traffic-eng link-management summary** command:

```
Router#show mpls traffic-eng link-management summary atm0/0.1
```

System Information::

```
Links Count:          3  
Flooding System:      enabled
```

IGP Area ID:: isis level-1

```
Flooding Protocol:    ISIS  
Flooding Status:      data flooded  
Periodic Flooding:    enabled (every 180 seconds)  
Flooded Links:        1
```

```

IGP System ID:          0001.0000.0001.00
MPLS TE Router ID:     10.106.0.6
IGP Neighbors:         3
Link ID:: AT0/0.1 (10.32.0.6)
Link Status:
  Physical Bandwidth:   155520000 bits/sec
  MPLS TE Bandwidth:    5000000 bits/sec (reserved:0% in, 80% out)
  MPLS TE Link State:   MPLS TE on, RSVP on, admin-up, flooded
  Inbound Admission:    allow-all
  Outbound Admission:   allow-if-room
  Admin. Weight:        10 (IGP)
  IGP Neighbor Count:   1

```

show mpls traffic-eng link-management interfaces

To show per-interface resource and configuration information, use the **show mpls traffic-eng link-management interfaces** command in EXEC mode.

show mpls traffic-eng link-management interfaces [*interface*]

Syntax Description

Parameter	Description
<i>interface</i>	(Optional) Specifies the name of a single interface for which information is to be displayed

Example: Usage of the show mpls traffic-eng link-management interfaces Command

The following example shows output from the **show mpls traffic-eng link-management interfaces** command:

```

Router#show mpls traffic-eng link-management interfaces
  System Information::
Links Count:          3
Link ID:: Et1/1/1 (10.1.0.6)
Link Status:
  Physical Bandwidth:  10000000 bits/sec
  MPLS TE Bandwidth:   5000000 bits/sec (reserved:0% in, 0% out)
  MPLS TE Link State:  MPLS TE on, RSVP on
  Inbound Admission:   reject-huge
  Outbound Admission:  allow-if-room
  Admin. Weight:       10 (IGP)
  IGP Neighbor Count:  2
  IGP Neighbor:        ID 0000.0000.0000.02, IP 0.0.0.0 (Up)

```

```

IGP Neighbor:          ID 0001.0000.0001.02, IP 0.0.0.0 (Down)
Flooding Status for each configured area [1]:
IGP Area[1] isis level-1: not flooded
                        (Reason:Interface has been administratively
disabled)
Link ID:: AT0/0.1 (10.32.0.6)
Link Status:
Physical Bandwidth:    155520000 bits/sec
MPLS TE Bandwidth:     5000000 bits/sec (reserved:0% in, 80% out)
MPLS TE Link State:    MPLS TE on, RSVP on, admin-up, flooded
Inbound Admission:     allow-all
Outbound Admission:    allow-if-room
Admin. Weight:         10 (IGP)
IGP Neighbor Count:    1
IGP Neighbor:          ID 0001.0000.0002.00, IP 10.32.0.10 (Up)
Flooding Status for each configured area [1]:
IGP Area[1] isis level-1: flooded

```

show mpls traffic-eng link-management bandwidth-allocation

To show current local link information, use the **show mpls traffic-eng link-management bandwidth-allocation** command in EXEC mode.

show mpls traffic-eng link-management bandwidth-allocation [*interface name*]

Syntax Description

Parameter	Description
<i>interface name</i>	(Optional) Shows only those tunnels that have been admitted on the specified interface

Example: Usage of the show mpls traffic-eng link-management bandwidth-allocation Command

The following example shows output from the **show mpls traffic-eng link-management bandwidth-allocation** command:

```
Router#show mpls traffic-eng link-management bandwidth-allocation
atm0/0.1
```

System Information::

Links Count: 3

Bandwidth Hold Time: max. 15 seconds

Link ID:: AT0/0.1 (10.32.0.6)

Link Status:

Physical Bandwidth: 155520000 bits/sec

MPLS TE Bandwidth: 5000000 bits/sec (reserved:0% in, 80% out)

BW Descriptors: 1

MPLS TE Link State: MPLS TE on, RSVP on, admin-up, flooded

Inbound Admission: allow-all

Outbound Admission: allow-if-room

Admin. Weight: 10 (IGP)

IGP Neighbor Count: 1

Up Thresholds: 15 30 45 60 75 80 85 90 95 96 97 98 99 100
(default)

Down Thresholds: 100 99 98 97 96 95 90 85 80 75 60 45 30 15
(default)

Outbound Bandwidth Information (bits/second):

KEEP	PRIORITY	BW HELD	BW TOTAL	HELD	BW LOCKED	BW TOTAL	LOCKED
0		0	0		0	0	
	1	0	0	4000000		4000000	
	2	0	0	0		4000000	
	3	0	0	0		4000000	
	4	0	0	0		4000000	
	5	0	0	0		4000000	
	6	0	0	0		4000000	
	7	0	0	0		4000000	

Monitoring MPLS-TE (Cont.)

Cisco.com

```
router#
```

```
show mpls traffic-eng link-management advertisements
```

- Shows local link information that MPLS Traffic Engineering link management is currently flooding into the global traffic engineering topology

```
router#
```

```
show mpls traffic-eng link-management admission-control  
[interface-name]
```

- Shows which tunnels were admitted locally and their parameters (such as priority, bandwidth, incoming and outgoing interface, and state)

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-5

show mpls traffic-eng link-management advertisements

To show local link information that is currently being flooded by MPLS-TE link management into the global TE topology, use the **show mpls traffic-eng link-management advertisements** command in EXEC mode.

Example: Usage of the show mpls traffic-eng link-management advertisements Command

The following example shows output from the **show mpls traffic-eng link-management advertisements** command:

```
Router#show mpls traffic-eng link-management advertisements
```

```
Flooding Status:      ready
Configured Areas:    1
IGP Area[1] ID:: isis level-1
System Information::
  Flooding Protocol:  ISIS
Header Information::
  IGP System ID:      0001.0000.0001.00
  MPLS TE Router ID:  10.106.0.6
  Flooded Links:      1
Link ID:: 0
  Link IP Address:    10.32.0.6
  IGP Neighbor:       ID 0001.0000.0002.00, IP 10.32.0.10
```

```

Admin. Weight:      10
Physical BW:        155520000 bits/sec
Reservable BW:      5000000 bits/sec
Output Bandwidth::
  BW Unreserved[0]:  5000000 bits/sec
  BW Unreserved[1]:  1000000 bits/sec
  BW Unreserved[2]:  1000000 bits/sec
  BW Unreserved[3]:  1000000 bits/sec
  BW Unreserved[4]:  1000000 bits/sec
  BW Unreserved[6]:  1000000 bits/sec
  BW Unreserved[7]:  1000000 bits/sec
Affinity Bits       0x00000000

```

show mpls traffic-eng link-management admission-control

To show which tunnels have been admitted locally, and their parameters (such as priority, bandwidth, incoming and outgoing interface, and state), use the **show mpls traffic-eng link-management admission-control** command in EXEC mode.

show mpls traffic-eng link-management admission-control [*interface name*]

Syntax Description

Parameter	Description
<i>interface name</i>	(Optional) Shows only those tunnels that are admitted on the specified interface

Example: Usage of the show mpls traffic-eng link-management admission-control Command

The following example shows output from the **show mpls traffic-eng link-management admission-control** command

```

Router#show mpls traffic-eng link-management admission-control

System Information::
  Tunnels Count:      1
  Tunnels Selected:   1
TUNNEL ID      UP IF      DOWN IF      PRIORITY      STATE
BANDWIDTH
3.3.25.3 1_1    -            PO1/0/0      1/1           Resv Admitted  10000  R

```

Monitoring MPLS-TE (Cont.)

Cisco.com

```
router#
```

```
show mpls traffic-eng link-management igp-neighbors [{  
  igp-id {isis isis-address | ospf ospf-id} | ip A.B.C.D }]
```

- Shows TE IGP neighbors

```
router#
```

```
show mpls traffic-eng topology [ { A.B.C.D | igp-id {  
  isis nsapaddr | ospf A.B.C.D } ] [ brief ]
```

- Shows the MPLS Traffic Engineering global topology currently known at this node

```
router(config)#
```

```
show mpls traffic-eng topology path tunnel-intf  
[parameters]
```

- Shows the properties of the best available path to a specified destination that satisfies certain constraints

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-6

show mpls traffic-eng link-management igp-neighbors

To show IGP neighbors, use the **show mpls traffic-eng link-management igp-neighbors** command in privileged EXEC mode.

show mpls traffic-eng link-management igp-neighbors [{**igp-id** {*isis isis-address* | **ospf** *ospf-id*} | **ip** *A.B.C.D*}]

Syntax Description

Parameter	Description
igp-id	Shows the IGP neighbors that are using a specified IGP identification
isis <i>isis-address</i>	Specifies an IS-IS neighbor to display when you are displaying neighbors by IGP ID
ospf <i>ospf-id</i>	Specifies an OSPF neighbor to display when you are displaying neighbors by IGP ID
ip <i>A.B.C.D</i>	Shows the IGP neighbors that are using a specified IGP IP address

Example: Usage of the show mpls traffic-eng link-management igp-neighbors Command

The following example shows output from the **show mpls traffic-eng link-management igp-neighbors** command:

```
Router#show mpls traffic-eng link-management igp-neighbors
Link ID:: Et0/2
    Neighbor ID: 0000.0024.0004.02 (area: isis level-1, IP:
0.0.0.0)
Link ID:: PO1/0/0
    Neighbor ID: 0000.0026.0001.00 (area: isis level-1, IP:
170.1.1.2)
```

show mpls traffic-eng topology

To show the MPLS-TE global topology as currently known at this node, use the **show mpls traffic-eng topology** command in privileged EXEC mode.

show mpls traffic-eng topology [{*A.B.C.D* | **igp-id** {*isis nsapaddr* | **ospf** *A.B.C.D*}] [**brief**]

Syntax Description

Parameter	Description
<i>A.B.C.D</i>	Specifies the node by the IP address (router identifier to interface address)
igp-id	Specifies the node by IGP router identifier
isis <i>nsapaddr</i>	Specifies the node by router identification (<i>nsapaddr</i>) if you are using IS-IS
ospf <i>A.B.C.D</i>	Specifies the node by router identifier if you are using OSPF
brief	(Optional) Provides a less detailed version of the topology

Example: Usage of the show mpls traffic-eng topology Command

The following example shows output from the **show mpls traffic-eng topology** command:

```
Router#show mpls traffic-eng topology

My_System_id: 0000.0025.0003.00

IGP Id: 0000.0024.0004.00, MPLS TE Id:24.4.4.4 Router Node
link[0 ]:Intf Address: 150.1.1.4
    Nbr IGP Id: 0000.0024.0004.02,
    admin_weight:10, affinity_bits:0x0
    max_link_bw:10000 max_link_reservable: 10000
```

reservable	allocated	reservable	allocated	
---	-----	-----	-----	-----
bw[0] : 0		10000	bw[1] : 0	10000
bw[2] : 0		10000	bw[3] : 0	10000
bw[4] : 0		10000	bw[5] : 0	10000
bw[6] : 0		10000	bw[7] : 0	10000

show mpls traffic-eng topology path

To show the properties of the best available path to a specified destination that satisfies certain constraints, use the **show mpls traffic-eng topology path EXEC** command. You specify the constraints in this command.

show mpls traffic-eng topology path { *tunnel-interface* [**destination address**] | **destination address** } [**bandwidth value**] [**priority value** [*value*]] [**affinity value** [**mask mask**]]

Syntax Description

Parameter	Description
<i>tunnel-interface</i>	Name of an MPLS-TE interface (for example, Tunnel 1) from which default constraints should be copied.
destination address	(Optional) IP address specifying the path destination.
bandwidth value	(Optional) Bandwidth constraint: the amount of available bandwidth that a suitable path requires. This constraint overrides the bandwidth constraint that is obtained from the specified tunnel interface. You can specify any positive number.
priority value [<i>value</i>]	(Optional) Priority constraints: the setup and hold priorities that are used to acquire bandwidth along the path. If specified, these constraints override the priority constraints that are obtained from the tunnel interface. Valid values are from 0 to 7.
affinity value	(Optional) Affinity constraints: the link attributes for which the path has an affinity. If specified, these constraints override the affinity constraints that are obtained from the tunnel interface.
mask mask	(Optional) Affinity constraints: the mask that is associated with the affinity specification.

Example: Usage of the show mpls traffic-eng topology path Command

The following is sample output from the **show mpls traffic-eng topology path** command:

```
Router1#show mpls traffic-eng topology path Tunnel1 bandwidth 1000
```

Query Parameters:

Destination:10.112.0.12

Bandwidth:1000

Priorities:1 (setup), 1 (hold)

Affinity:0x0 (value), 0xFFFF (mask)

Query Results:

Min Bandwidth Along Path:2000 (kbps)

Max Bandwidth Along Path:5000 (kbps)

Hop 0:10.1.0.6 :affinity 00000000, bandwidth 2000 (kbps)

Hop 1:10.1.0.10 :affinity 00000000, bandwidth 5000 (kbps)

Hop 2:10.43.0.10 :affinity 00000000, bandwidth 2000 (kbps)

Monitoring MPLS-TE Resource Usage on Intermediate Devices

This topic describes the command syntax that is used to monitor MPLS-TE resource usage on intermediate devices.

Monitoring MPLS-TE Resource Usage on Intermediate Devices

Cisco.com

```
router#  
show ip ospf database opaque-area
```

- Displays lists of traffic engineering opaque link-state advertisements (LSAs)

```
router#  
show ip ospf [ process-id [ area-id ] ] mpls traffic-eng  
[ link ] | [ fragment ]
```

- Displays information about the links available on the local router for traffic engineering

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—4-7

show ip ospf database opaque-area

To display lists of TE opaque link-state advertisements (LSAs), also known as Type-10 opaque link area link states, use the **show ip ospf database opaque-area** EXEC command.

Example: Usage of the show ip ospf database opaque-area Command

The following is sample output from the **show ip ospf database opaque-area** command:

```
Router#show ip ospf database opaque-area  
OSPF Router with ID (25.3.3.3) (Process ID 1)  
Type-10 Opaque Link Area Link States (Area 0)  
LS age: 12  
Options: (No TOS-capability, DC)  
LS Type: Opaque Area Link  
Link State ID: 1.0.0.0  
Opaque Type: 1  
Opaque ID: 0  
Advertising Router: 24.8.8.8  
LS Seq Number: 80000004
```



```
Checksum: 0xD423
Length: 132
Fragment number : 0
MPLS TE router ID: 24.8.8.8
Link connected to Point-to-Point network
Link ID : 26.2.2.2
Interface Address : 198.1.1.1
```

show ip ospf

To display general information about OSPF routing processes, use the **show ip ospf EXEC** command.

show ip ospf [*process-id*]

Syntax Description

Parameter	Description
<i>process-id</i>	(Optional) Process ID. If this argument is included, only information for the specified routing process is displayed.

Example: Usage of the show ip ospf Command

The following is sample output from the **show ip ospf** command when it is entered without a specific OSPF process ID:

```
Router#show ip ospf
Routing Process "ospf 201" with ID 192.42.110.200
Supports only single TOS(TOS0) route
It is an area border and autonomous system boundary router
Redistributing External Routes from,
    igrp 200 with metric mapped to 2, includes subnets in
redistribution
    rip with metric mapped to 2
    igrp 2 with metric mapped to 100
    igrp 32 with metric mapped to 1
Number of areas in this router is 3
Area 192.42.110.0
    Number of interfaces in this area is 1
    Area has simple password authentication
    SPF algorithm executed 6 times
```

Monitoring MPLS-TE Resource Usage on Intermediate Devices (Cont.)

Cisco.com

```
router#
```

```
show isis database verbose
```

- Displays more information (including sub-TLV information) about the database

```
router(config)#
```

```
show isis mpls traffic-eng adjacency-log
```

- Displays a log of 20 entries of MPLS Traffic Engineering IS-IS adjacency changes

```
router(config)#
```

```
show isis mpls traffic-eng advertisements
```

- Displays the last flooded record from MPLS Traffic Engineering

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-8

show isis database verbose

To display more information about the database, use the **show isis database verbose** command in EXEC mode.

Example: Usage of the show isis database verbose Command

The following is sample output from the **show isis database verbose** command:

```
Router#show isis database verbose
```

```
IS-IS Level-1 Link State Database
```

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime

```
ATT/P/OL
```

ntp-5.00-00 0/0/0	* 0x000000E6	0xC9BB	1042

```
Area Address:49.0001
```

```
NLPID: 0xCC
```

```
Hostname:ntp-5
```

```
Router ID: 5.5.5.5
```

```
IP Address: 172.21.39.5
```

```
Metric:10 IP 172.21.39.0/24
```

ntp-5.00-01 0/0/0	* 0x000000E7	0xAB36	1065

```
Metric:10 IS-Extended ntp-5.01
```

```
Affinity:0x00000000
```

```
Interface IP Address:172.21.39.5
```

```

Physical BW:10000000 bits/sec
Reservable BW:1166000 bits/sec
BW Unreserved[0]: 1166000 bits/sec, BW Unreserved[1]: 1166000
bits/sec
BW Unreserved[2]: 1166000 bits/sec, BW Unreserved[3]: 1166000
bits/sec
BW Unreserved[4]: 1166000 bits/sec, BW Unreserved[5]: 1166000
bits/sec
BW Unreserved[6]: 1166000 bits/sec, BW Unreserved[7]: 1153000
bits/sec
Metric:0          ES dtp-5

```

show isis mpls traffic-eng adjacency-log

To display a log of 20 entries of MPLS-TE IS-IS adjacency changes, use the **show isis mpls traffic-eng adjacency-log** command in EXEC mode.

Example: Usage of the show isis mpls traffic-eng adjacency-log Command

The following is sample output from the **show isis mpls traffic-eng adjacency-log** command:

```

Router#show isis mpls traffic-eng adjacency-log
IS-IS RRR log

```

When	Neighbor ID	IP Address	Interface	Status	Level
04:52:52	0000.0024.0004.02	0.0.0.0	Et0/2	Up	level-1
04:52:50	0000.0026.0001.00	170.1.1.2	PO1/0/0	Up	level-1
04:52:37	0000.0024.0004.02	0.0.0.0	Et0/2	Up	level-1

show isis mpls traffic-eng advertisements

To display the last flooded record from MPLS-TE, use the **show isis mpls traffic-eng advertisements** command in EXEC mode.

Example: Usage of the show isis mpls traffic-eng advertisements Command

The following is sample output from the **show isis mpls traffic-eng advertisements** command:

```
Router#show isis mpls traffic-eng advertisements
System ID:dtp-5.00
Router ID:5.5.5.5
Link Count:1
Link[1]
  Neighbor System ID:dtp-5.01 (broadcast link)
  Interface IP address:172.21.39.5
  Neighbor IP Address:0.0.0.0
  Admin. Weight:10
  Physical BW:10000000 bits/sec
  Reservable BW:1166000 bits/sec
  BW unreserved[0]:1166000 bits/sec, BW unreserved[1]:1166000
    bits/sec
  BW unreserved[2]:1166000 bits/sec, BW unreserved[3]:1166000
    bits/sec
  BW unreserved[4]:1166000 bits/sec, BW unreserved[5]:1166000
    bits/sec
  BW unreserved[6]:1166000 bits/sec, BW unreserved[7]:1153000
    bits/sec
  Affinity Bits:0x00000000
```

Monitoring MPLS-TE Resource Usage on Intermediate Devices (Cont.)

Cisco.com

```
router#
```

```
show isis mpls traffic-eng tunnel
```

- Displays information about tunnels considered in the IS-IS next-hop calculation

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—4-9

show isis mpls traffic-eng tunnel

To display information about tunnels that are considered in the IS-IS next-hop calculation, use the **show isis mpls traffic-eng tunnel** command in EXEC mode.

Example: Usage of the show isis mpls traffic-eng tunnel Command

The following is sample output from the **show isis mpls traffic-eng tunnel** command:

```
Router#show isis mpls traffic-eng tunnel
```

Station Id Mode	Tunnel Name	Bandwidth	Nexthop	Metric
kangpa-router1.00 Relative	Tunnel1022	3333	2.2.2.2	-3
Absolute	Tunnel1021	10000	2.2.2.2	11
tomklong-route.00 Relative	Tunnel1031	10000	3.3.3.3	-1
	Tunnel1032	10000	3.3.3.3	

Troubleshooting Device and Link Setup

This topic describes the command syntax that is used to troubleshoot device and link setup.

Troubleshooting Device and Link Setup

Cisco.com

```
router#  
debug mpls traffic-eng  
  areas          Area configuration change events  
  autoroute      Automatic routing over tunnels  
  link-management Link management  
  load-balancing Unequal-cost load balancing over tunnels  
  path           Path calculation events  
  topology       Show topology events  
  tunnels        MPLS Traffic Engineering tunnels
```

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-10

The **debug mpls traffic-eng** command has various options similar to the **show** commands.

The **debug mpls traffic-eng tunnels** command is a useful command to show tunnel events and errors.

The **path** option shows the path calculation, and the **autoroute** option shows the advertisement of the tunnel as a potential route through the network.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **The following commands are used to monitor MPLS-TE tunnels:**
 - show mpls traffic-eng tunnels
 - show mpls traffic-eng autoroute
- **The show mpls traffic-eng link-management and show mpls traffic-eng topology commands are used to monitor MPLS-TE.**
- **Use the show ip ospf and show isis commands to monitor MPLS-TE resource usage.**
- **The debug mpls traffic-eng command is the main that is used for troubleshooting device and link setup.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—4-11

References

For additional information, refer to this resource:

- Search the Cisco website at <http://www.cisco.com> for more information on topics that were discussed in this lesson.

Next Steps

For the associated lab exercises, refer to the following sections of the course Lab Guide:

- Lab Exercise 4-1: Implementing MPLS Traffic Engineering
- Lab Exercise 4-2: Implementing Link Protection for MPLS Traffic Engineering

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The command that is used to show tunnels that are announced to the IGP is the _____ command.
- Q2) The command that is used to show per-interface resource and configuration information is the _____ command.
- Q3) The command that is used in IS-IS implementation to display the last flooded record from MPLS-TE is the _____ command.
- Q4) The command that displays path calculation events is the **debug mpls traffic-eng** _____ command.

Lesson Answer Key

- Q1) **show mpls traffic-eng autoroute**
Relates to: Monitoring MPLS-TE Tunnels
- Q2) **show mpls traffic-eng link-management interfaces**
Relates to: Monitoring MPLS-TE
- Q3) **show isis mpls traffic-eng advertisements**
Relates to: Monitoring MPLS-TE Resource Usage on Intermediate Devices
- Q4) **path**
Relates to: Troubleshooting Device and Link Setup

MPLS Quality of Service

Overview

This module covers the concepts and features of Multiprotocol Label Switching (MPLS) quality of service (QoS) and presents techniques that network administrators can apply to help the service provider meet the service level agreements (SLAs) of customers. The module also covers how QoS is configured in various situations.

Objectives

Upon completing this module, you will be able to describe the tasks and commands that are necessary to implement MPLS Traffic Engineering (MPLS-TE). This includes being able to do the following:

- Describe the QoS models
- Describe how MPLS supports DiffServ
- Describe how to configure MPLS to support QoS
- Describe how QoS works with MPLS applications

Outline

The module contains these lessons:

- QoS Models
- MPLS Support for DiffServ
- Configuring MPLS QoS
- QoS in MPLS Applications

QoS Models

Overview

This lesson describes the concepts, features, and techniques of quality of service (QoS) that are used within the service provider network. QoS tunneling modes are also be discussed, as well as how Differentiated Services Code Point (DSCP) services are handled by MPLS.

Relevance

This lesson is mandatory for learners who are planning to improve their usage of network resources by implementing QoS in their MPLS networks.

Objectives

This lesson describes the QoS models. Upon completing this lesson, you will be able to do the following:

- Identify the MPLS QoS models
- Identify the characteristics of the QoS pendulum
- Identify the characteristics of the integrated QoS model
- Identify the characteristics of the differentiated QoS model

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Prior knowledge and experience in implementing and configuring QoS in Cisco IOS networks
- Successful completion of the “MPLS Traffic Engineering Technology” and “Configuring MPLS Traffic Engineering” modules of this course

Outline

This lesson includes these topics:

- Overview
- QoS Models
- The QoS Pendulum
- Integrated QoS Model
- Differentiated QoS Model
- Lesson Summary
- Lesson Review

QoS Models

This topic describes the MPLS QoS models.

QoS Models

Cisco.com

- **Best-effort model** – no QoS is applied to packets (default behavior)
- **Integrated Services model** – applications signal to the network that they require special QoS
- **Differentiated Services model** – the network recognizes classes that require special QoS

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-3


During the course of its development thus far, the Internet has experienced three QoS-related models:

- **Best-effort model:** This model is very much a part of the original design of the Internet as a medium for best-effort, no-guarantee delivery of packets. The best-effort approach is still predominant in the Internet today.
- **Integrated Services (IntServ) model:** Introduced to supplement best-effort delivery by setting aside some bandwidth for applications that require bandwidth and delay guarantees. The IntServ model expects applications to signal their QoS requirements to the network. Resource Reservation Protocol (RSVP) is used to signal the requirements.
- **Differentiated Services (DiffServ) model.** Added to provide more scalability in providing QoS to IP packets. The main difference between the DiffServ model and the IntServ model is that, with the DiffServ model, the network recognizes packets (no signaling is needed) and provides the appropriate services to them.

The IP networks of today can use all three models at the same time.

The QoS Pendulum

This topic describes the “QoS pendulum.”



The figure here shows the “QoS pendulum.” Originally, there was no QoS. As stated previously, the Internet was designed for best-effort, no-guarantee delivery of packets.

The first efforts to implement QoS swung the pendulum too far. It was implemented with IntServ and RSVP. While these implementations addressed the QoS issues, they came at a high cost in network bandwidth usage.

The pendulum has finally come to rest in the middle, because QoS is implemented with DiffServ.

Because per-flow QoS is difficult to achieve end-to-end in a network without adding significant complexity, cost, and scalability issues, it naturally leads one to think about classifying flows into aggregates (classes), and providing appropriate QoS for the aggregates.

For example, all TCP flows can be grouped as a single class, and bandwidth can be allocated for the class rather than for individual flows. In addition to the need to classify traffic with aggregated flows, signaling and state maintenance requirements on each network node should be minimized. The Internet Engineering Task Force (IETF) has realized this fact, and defined a mechanism to use the type of service (ToS) field in the IP version 4 (IPv4) header to prioritize packets; any network node along the path of the packet knows the relative importance (priority level) of the packet, and can apply preferential forwarding to packets with higher priority levels.

The ultimate goal is to optimize bandwidth for end-to-end SLAs by using a combination of IntServ, DiffServ, and Traffic Engineering (TE) tools.

Integrated QoS Model

This topic describes the integrated QoS model.

Integrated QoS Model

Cisco.com

- **Application requests a specific kind of QoS service, through explicit signaling.**
- **Resource Reservation Protocol (RSVP) is used by applications to signal their QoS requirements to the router.**
- **Complex to use.**
- **Difficult to support with a large number of RSVP connections, due to:**
 - **the amount of state information required for every flow**
 - **the amount of control traffic**
- **Fine grain, providing strict QoS.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—5-5

The characteristics of the integrated QoS model are as follows:

- It signals QoS requests per individual flow. The network can then provide guarantees to these individual flows. The problem with this situation is that it does not scale to large networks because of the large numbers of concurrent RSVP flows.
- It informs network devices of flow parameters (IP addresses and port numbers). Some applications use dynamic port numbers, which can be difficult for network devices to recognize. Network-based application recognition (NBAR) is a mechanism that has been introduced to supplement RSVP for applications that use dynamic port numbers but do not use RSVP.
- Continuous signaling takes place because of the stateless operation of RSVP.
- RSVP is not scalable to large networks where per-flow guarantees would have to be made to thousands of flows.
- It supports admission control that allows a network to reject (or downgrade) new RSVP sessions if one of the interfaces in the path has reached the limit (all reservable bandwidth is booked).

The IntServ model is described in RFC 1633, *Integrated Services in the Internet Architecture: An Overview* (<http://www.ietf.org/rfc/rfc1633.txt>).

Differentiated QoS Model

This topic describes the differentiated QoS model.

Differentiated QoS Model

Cisco.com

- **QoS is provided by differential treatment to each packet or class of packets.**
- **No explicit signaling from the application.**
- **This model is appropriate for aggregate flows.**
- **Coarse grain, not strict QoS (no guarantees).**


© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-6

The characteristics of the differentiated QoS model are as follows:

- The DiffServ model describes services that are associated with traffic classes. Traffic classes are identified by the value of the DSCP (the DSCP replaces IP precedence in the ToS field of the IP header). Here are more details:
 - While the DSCP replaces IP precedence, it maintains interoperability with non-DiffServ-compliant devices (those that still use IP precedence). Because of this backward compatibility, DiffServ can be gradually deployed in large networks.
 - The DiffServ field (DS field) is the former 8-bit ToS field. The main difference is that the DSCP supports more classes (64) than IP precedence (8).
- The main goals of the DiffServ model are to provide scalability and a similar level of QoS to the IntServ model, without having to do it on a per-flow basis. The network simply identifies a class (not application) and applies the appropriate per-hop behavior, or PHB (QoS mechanism).
- A traffic aggregate is a collection of all flows that require the same service. A service is implemented using different QoS mechanisms (a QoS mechanism implements a PHB).
- The idea is for the network to recognize a class without having to receive any request from applications. This capability allows the QoS mechanisms to be applied to other applications that do not have the RSVP functionality, which is the case for 99 percent of applications that use IP.

The DiffServ model and associated standards are described in the following IETF standardization documents (RFCs):

- RFC 2475, *An Architecture for Differentiated Services* (<http://www.ietf.org/rfc/rfc2475.txt>)
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* (<http://www.ietf.org/rfc/rfc2474.txt>)



This figure shows how with the QoS DiffServ model, the traffic can be divided into different classes that have different requirements.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **There are three QoS models:**
 - Best-effort
 - Integrated Services
 - Differentiated Services
- **DiffServ is at the middle of the QoS pendulum; it has simplicity and the ability to scale.**
- **IntServ uses RSVP so that applications can signal their QoS requirements to the router.**
- **DiffServ provides differential treatment for each packet or class of packets.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-8

References

For additional information, refer to these resources:

- RFC 2475, *An Architecture for Differentiated Services* (<http://www.ietf.org/rfc/rfc2475.txt>)
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* (<http://www.ietf.org/rfc/rfc2474.txt>)
- RFC 1633, *Integrated Services in the Internet Architecture: An Overview*, which provides a description of the IntServ model (<http://www.ietf.org/rfc/rfc1633.txt>)

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) The IntServ QoS model expects applications to _____ their bandwidth and guarantee requirements to the network.
- Q2) The _____ QoS model is simpler and has the ability to scale when providing QoS services to a network.
- Q3) Which two of the following are true of the Integrated Services model? (Choose two.)
- A) It is scalable.
 - B) It is not scalable.
 - C) It does not require RSVP.
 - D) It is difficult to support when there are a large number of RSVP connections.
- Q4) Which two of the following are true of the Differentiated Services model? (Choose two.)
- A) QoS is provided by differential treatment to each packet.
 - B) DiffServ requires explicit signaling.
 - C) DiffServ requires no explicit signaling.
 - D) DiffServ is not appropriate for aggregate flows.

Lesson Answer Key

Q1) signal

Relates to: QoS Models

Q2) differentiated, or DiffServ

Relates to: The QoS Pendulum

Q3) B, D

Relates to: Integrated QoS Model

Q4) A, C

Relates to: Differentiated QoS Model

MPLS Support for DiffServ

Overview

This lesson describes the DiffServ model, including how MPLS supports DiffServ.

Relevance

This lesson is mandatory for learners who need to deploy QoS in their networks.

Objectives

This lesson describes how MPLS supports DiffServ. Upon completing this lesson, you will be able to do the following:

- Identify the characteristics of the DiffServ architecture
- Identify the features of the DiffServ model
- Identify the DiffServ PHBs and recommended codepoints
- Identify the characteristic of DiffServ scalability by means of aggregation
- Identify the characteristic of MPLS scalability by means of aggregation
- Identify how MPLS marks frames

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” and “Configuring MPLS Traffic Engineering” modules of this course


Outline

This lesson includes these topics:

- Overview
- DiffServ Architecture
- DiffServ Model Features
- DiffServ PHBs and Recommended Codepoints
- DiffServ Scalability by Means of Aggregation
- MPLS Scalability by Means of Aggregation
- Marking MPLS Frames
- Lesson Summary
- Lesson Review

DiffServ Architecture

This topic presents an overview of the DiffServ architecture.



DiffServ is a multiple service model that can satisfy differing QoS requirements. However, unlike the IntServ model, an application using DiffServ does not explicitly signal the router before sending data.

For DiffServ, the network tries to deliver a particular kind of service based on the QoS that is specified by each packet. This specification can occur in different ways, for example, by using the IP precedence bit settings in IP packets or by using source and destination addresses. The network uses the QoS specification to classify, mark, shape, and police traffic, and to perform intelligent queuing.

The DiffServ model is used for several mission-critical applications and for providing end-to-end QoS. Typically, this service model is appropriate for aggregate flows because it performs a relatively coarse level of traffic classification.

DiffServ Model Features

This topic identifies features of the DiffServ model.

DiffServ Model Features

Cisco.com

Classification

- **Marking**
- **Policing and Shaping**
- **Congestion Avoidance**
- **Congestion Management**

© 2004, Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-4

Classification

Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, by using the three precedence bits in the ToS field of the IP packet header—two of the values are reserved for other purposes—you can categorize packets into a limited set of up to six traffic classes. After you classify packets, you can use other QoS features to assign the appropriate traffic handling policies, including congestion management, bandwidth allocation, and delay bounds for each traffic class.

Packets can also be classified by external sources, that is, by a customer or by a downstream network provider. You can either allow the network to accept the classification or override it and reclassify the packet according to a policy that you specify.

Packets can be classified based on policies that are specified by the network operator. Policies can be set that include classification based on physical port, source or destination IP or MAC address, application port, IP protocol type, and other criteria that you can specify by using access lists or extended access lists.

You can use Cisco IOS QoS policy-based routing (PBR) and the classification features of Cisco IOS QoS committed access rate (CAR) to classify packets. You can use Border Gateway Protocol (BGP) policy propagation to propagate destination-based packet classification policy throughout a large network via BGP routing updates.

Marking

The IPv4 ToS octet has been redefined from the 3-bit IP precedence field to a 6-bit DSCP field. Packets can be marked with an arbitrary DSCP value or standard values, corresponding to the appropriate Assured Forwarding (AF), Expedited Forwarding (EF), or user-defined class.

Cisco IOS software also supports class-selector codepoints, which provide a way of marking the six DSCP bits. These codepoints are of the form “xyz000,” where x, y, and z can represent a 1 or 0. The codepoint for best-effort traffic will be set to “000000.” The Cisco implementation of codepoints provides added value by allowing you to mark packets with an arbitrary DSCP and map them to a locally significant (non-AF, non-EF, or default) PHB. This implementation allows for the construction of new services.

Policing and Shaping

Cisco IOS QoS includes traffic policing capabilities that are implemented through the rate-limiting aspects of CAR and traffic-shaping capabilities provided by the Generic Traffic Shaping (GTS) and Frame Relay Traffic Shaping (FRTS) protocols.

Congestion Avoidance

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network and internetwork bottlenecks before it becomes a problem. These techniques are designed to provide preferential treatment for premium (priority) class traffic under congestion situations while concurrently maximizing network throughput and capacity utilization and minimizing packet loss and delay. Weighted random early detection (WRED) and its counterpart for the Versatile Interface Processor (VIP), VIP-distributed WRED (DWRED), are the Cisco IOS QoS congestion avoidance features.

Router behavior allows output buffers to fill during periods of congestion, using tail drop to resolve the problem when WRED is not configured. During tail drop, a potentially large number of packets from numerous connections are discarded because of lack of buffer capacity. This behavior can result in waves of congestion followed by periods during which the transmission link is not fully used. WRED obviates this situation proactively by providing congestion avoidance. Instead of waiting for buffers to fill before dropping packets, the router monitors the buffer depth and performs early discards on selected packets that are sent over selected connections.


WRED is the Cisco implementation of the random early detection (RED) class of congestion avoidance algorithms. When RED is used and the source detects the dropped packet, the source slows its transmission. RED is primarily designed to work with TCP in IP internetwork environments.

Congestion Management

Congestion management features operate to control congestion once it occurs. One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic, and then determine some method of prioritizing it onto an output link. Each queuing algorithm was designed to solve a specific network traffic problem and has a particular effect on network performance. The Cisco IOS software congestion management, or queuing, features include FIFO, priority queuing (PQ), custom queuing (CQ), flow-based weighted fair queuing (WFQ), distributed WFQ (DFWFQ), class-based WFQ (CBWFQ), IP Real-Time Transport Protocol (RTP) Priority, Frame Relay IP RTP Priority, and low latency queuing (LLQ).

DiffServ PHBs and Recommended Codepoints

This topic explains DiffServ PHBs and recommended codepoints.



RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, adopts this set of PHBs and values by creating the Class Selector PHB Group. The low-order three bits of the DSCP can identify the class selector or the low-order five bits of the DS field can all be 0.

The DiffServ model uses the DS field in the IP header to mark packets according to their classification into Behavior Aggregates (BAs). The DS field occupies the same eight bits of the IP header that were previously used for the ToS field.

Each DSCP value identifies a BA. Each BA is assigned a PHB. Each PHB is implemented using the appropriate QoS mechanism or a set of QoS mechanisms.

The low-order bit of the DSCP identifies whether the DSCP value identifies a standard action (PHB) or a user-defined action.

The default value of the DSCP is 0. The associated PHB is FIFO service with a tail drop. FIFO queuing is discussed in the “IP QoS – Queuing Mechanisms” module of the *Implementing Cisco Quality of Service (QoS)* course.

The default DSCP value seamlessly maps to the default IP precedence value, which is also 0 according to RFC 1812.

The DSCP selects the PHB throughout the network.

The following PHBs are defined by IETF standards:

- **Default PHB:** Used for best-effort service
 - Normal Mail
- **Class Selector PHB:** Used for backward compatibility with non-DiffServ-compliant devices (devices compliant with RFC 1812 and, optionally, devices compliant with RFC 791)
 - $FP(\text{Precedence}(x+1))^3 FP(\text{Precedence}(x))$
 - Compare to $FP(\text{Express Mail})^3 FP(\text{Priority Mail})$
- **Expedited Forwarding PHB:** Used for low-delay service, low jitter, and assured bandwidth
 - Compare to Express Mail, with Overnight Delivery
- **Assured Forwarding PHB:** Used for guaranteed bandwidth service
 - IETF has defined four AF classes
Compare to registered mail – Very safe and assured

Assured Forwarding

The rough equivalent of the IntServ Controlled Load Service is the Assured Forwarding PHB (AF PHB). It defines a method by which BAs can be given different forwarding assurances. For example, traffic can be divided into gold, silver, and bronze classes, with gold being allocated 50 percent of the available link bandwidth, silver 30 percent, and bronze 20 percent.

The AF_{xy} PHB defines four AF_x classes; namely, AF1, AF2, AF3, and AF4. Each class is assigned a certain amount of buffer space and interface bandwidth, dependent on the SLA with the service provider or the policy. Within each AF_x class, it is possible to specify three drop precedence values.

Thus, the variable “y” in AF_{xy} denotes the drop precedence within an AF_x class. This concept of drop precedence is useful, for example, to penalize flows within a BA that exceed the assigned bandwidth. Packets within these flows can be re-marked by a policer to have a higher drop precedence.

Assured Forwarding Class	Drop Probability	DSCP Value
AF class 1	Low	001 01 0
	Medium	001 10 0
	High	001 11 0
AF class 2	Low	010 01 0
	Medium	010 10 0
	High	010 11 0
AF class 3	Low	011 01 0
	Medium	011 10 0
	High	011 11 0
AF class 4	Low	100 01 0
	Medium	100 10 0
	High	100 11 0

Expedited Forwarding

The EF PHB fulfills the following functions:


- Ensures a minimum departure rate to provide the lowest possible delay to delay-sensitive applications
- Guarantees bandwidth to prevent starvation of the application if there are multiple applications using EF PHB
- Polices bandwidth to prevent starvation of other applications or classes that are not using this PHB

Packets requiring EF should be marked with the DSCP binary value “101110” (46 or 0x2E)

Non-DiffServ-compliant devices will regard the EF DSCP value as IP precedence 5 (101), which is the highest user-definable IP precedence and the one that is typically used for delay-sensitive traffic such as Voice over IP (VoIP).

DiffServ Scalability by Means of Aggregation

This topic describes DiffServ scalability by means of aggregation.



- DiffServ scalability comes from the following:
 - Aggregation of traffic on the edge
 - Processing of aggregates only in the core


The main goals of the DiffServ model are to provide scalability and a similar level of QoS to the IntServ model, without having to do it on a per-flow basis. The network simply identifies a class (not application) and applies the appropriate PHB (QoS mechanism).

DiffServ offers application-level QoS and traffic management in an architecture that incorporates mechanisms to control bandwidth, delay, jitter, and packet loss. The Cisco DiffServ complements the Cisco IntServ offering by providing a more scalable architecture for an end-to-end QoS solution. This scalability is achieved by the mechanisms controlling QoS at an aggregate level.

Application traffic can be categorized into multiple classes (aggregates), with QoS parameters defined for each class. A typical arrangement would be to categorize traffic into premium, gold, silver, bronze, and best-effort classes.

MPLS Scalability by Means of Aggregation

This topic describes the characteristic of MPLS scalability by means of aggregation.



MPLS scalability comes from the following:

- Aggregation of traffic on the edge
- Processing of aggregates only in the core

MPLS does not define a new QoS architecture. Most of the work on MPLS QoS has focused on supporting current IP QoS architectures. There are two QoS architectures defined for IP: IntServ and DiffServ.

IntServ defines per-flow QoS and uses RSVP as the signaling mechanism that is used by applications to request QoS from the network. MPLS can support per-flow QoS with the extensions that have been made to RSVP to propagate bindings between flows and labels. The LABEL_REQUEST and LABEL objects that are added to RSVP enable downstream label allocation using the Path and Resv messages. These extensions are commonly used for implementing resource reservation for flow aggregates in MPLS-TE. They are not used for per-flow QoS because of the limited scalability that such an approach would have in a service provider backbone.

On the other hand, DiffServ defines a QoS architecture based on flow aggregates that requires traffic to be conditioned and marked at the network edges and internal nodes to provide different QoS treatment to packets based on their markings. MPLS packets need to carry the packet marking in their headers because label switch routers (LSRs) do not examine the IP header during forwarding. A three-bit field in the MPLS shim header is used for this purpose. The DiffServ functionality of an LSR is almost identical to that provided by an IP router with respect to the QoS treatment that is given to packets (or PHB in DiffServ terms).

Marking MPLS Frames

This topic describes how MPLS marks frames.

Marking MPLS Frames

Cisco.com

How is DiffServ information conveyed to LSRs?

- **E-LSP**
 - “Queue” inferred from EXP field
 - “Drop priority” inferred from EXP field
- **L-LSP**
 - “Queue” inferred exclusively from label
 - “Drop priority” may be inferred from EXP field

© 2004, Cisco Systems, Inc. All rights reserved.MPLST v2.0—5-8

When packets enter the network, they are marked based on classification policies at the network boundary nodes. The boundary nodes also apply traffic-conditioning functions to control the amount of traffic that enters the network. Traffic conditioning includes shaping (smoothing the rate at which packets are sent into the network) and policing (dropping packets that are in excess of a subscribed-to rate, or recoloring the ones that exceed the rate so that the probability of dropping them increases when there is congestion in the core). Each node within the network then applies different queuing and dropping policies on every packet based on the marking that the packet carries.

What Is an E-LSP?

An E-LSP is a label switched path (LSP) on which nodes infer the QoS treatment for the MPLS packet exclusively from the experimental (EXP) bits in the MPLS header. Because the QoS treatment is inferred from the EXP field (both class and drop precedence), several classes of traffic can be multiplexed onto a single LSP (use the same label). A single LSP can support up to eight classes of traffic because the EXP field is a three-bit field. The maximum number of classes would be less after reserving some values for control plane traffic or if some of the classes have a drop precedence that is associated with them. E-LSPs are not an option for ATM LSRs. On those devices, MPLS packets use the original ATM cell encapsulation, and no EXP field exists in the cell header.

What Is an L-LSP?

An L-LSP is an LSP on which nodes infer the QoS treatment for MPLS packets from the packet label and the EXP bits (or the cell loss priority [CLP] bit for cell-mode MPLS). In particular, the label is used to encode the class that a packet belongs to, and the EXP field (or the CLP bit for cell-mode MPLS) is used to encode the drop precedence of the packet. A separate LSP can be established for each combination of forwarding equivalence class (FEC) and class. For example, three separate LSPs can be established to a single destination if there are packets belonging to three different classes that reach that destination. The class that is associated with an L-LSP needs to be signaled explicitly during label establishment so that each LSR can subsequently infer the packet class from the label. A new RSVP object (DiffServ) and a new LDP type, length, value (TLV) object (DiffServ) are defined for this purpose.

LSPs supporting DiffServ may be established with bandwidth reservation. That is, bandwidth requirements for an LSP can be signaled at LSP establishment time. Bandwidth reservation can be used to perform admission control on the DiffServ resources that have been provisioned. Though admission control can be performed on an LSP basis, the QoS design within the MPLS network is DiffServ-based, taking advantage of the scalability benefits that are implicit in that QoS architecture.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- DiffServ provides end-to-end QoS
- DiffServ provides a more scalable architecture and is achieved by the mechanisms controlling QoS at an aggregate level
- MPLS scalability comes from aggregation of traffic on the edge and processing of aggregates only in the core
- Two methods are used to mark MPLS traffic for QoS handling: E-LSP and L-LSP
 - E-LSP – One LSP, multiframe; EXP bits indicate queuing
 - L-LSP – One LSP per flow; EXP bits indicate drop priority

© 2004, Cisco Systems, Inc. All rights reserved.

MPLS v2.0—5-9

References

For additional information, refer to these resources:

- RFC 2598, *An Expedited Forwarding PHB* (<http://www.ietf.org/rfc/rfc2598.txt>)
- RFC 2597, *Assured Forwarding PHB Group* (<http://www.ietf.org/rfc/rfc2597.txt>)
- RFC 2475, *An Architecture for Differentiated Services* (<http://www.ietf.org/rfc/rfc2475.txt>)
- RFC, 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* (<http://www.ietf.org/rfc/rfc2474.txt>)

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Applications using the DiffServ QoS model _____ explicitly signal the router before sending data.
- Q2) Which of the following is not a feature of DiffServ?
- A) policing and shaping
 - B) congestion avoidance
 - C) congestion elimination
 - D) congestion management
- Q3) The IETF has defined how many AF classes?
- A) 2
 - B) 4
 - C) 16
 - D) 32
- Q4) When using DiffServ, the network identifies a _____ and not an _____, and applies the appropriate QoS mechanism.
- Q5) An E-LSP has the queue inferred from _____ fields and the drop priority inferred from the _____ fields.
- Q6) An L-LSP has the queue inferred from _____ field and the drop priority inferred from the _____ field.

Lesson Answer Key

- Q1) does not
Relates to: DiffServ Architecture
- Q2) C
Relates to: DiffServ Model Features
- Q3) B
Relates to: DiffServ PHBs and Recommended Codepoints
- Q4) class, application
Relates to: DiffServ Scalability by Means of Aggregation
- Q5) label and EXP, label and EXP
Relates to: Marking MPLS Frames
- Q6) label, EXP
Relates to: Marking MPLS Frames

Configuring MPLS QoS

Overview

This lesson describes the configuration tool, Modular QoS CLI (MQC), which is used when applying the QoS parameter on a router. The lesson also discusses how to configure MPLS QoS, including configuration syntax and descriptions.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS QoS on Cisco IOS platforms.

Objectives

This lesson describes how to configure MPLS to support QoS. Upon completing this lesson, you will be able to do the following:

- Identify the features that are available in the QoS toolkit
- Identify the features of MQC
- Identify IOS QoS MQC abstractions
- Configure QoS in IOS MQC syntax

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” and “Configuring MPLS Traffic Engineering” modules of this course


Outline

This lesson includes these topics:

- Overview
- How to Use the QoS Toolkit
- Configuring QoS in Cisco IOS Modular QoS CLI
- Configuring QoS in Cisco IOS MQC Abstractions
- Configuring QoS in Cisco IOS MQC Syntax
- MPLS QoS Configuration Case Study
- Lesson Summary
- Lesson Review

How to Use the QoS Toolkit

This topic describes the QoS toolkit.



Cisco IOS software provides a variety of QoS tools to provide the service levels that are presented in the figure. These tools are typically used within a single network element.

Typically, these tools are turned on at an interface to provide the right QoS characteristics for a specific network application.

The Cisco IOS QoS tools provide three major functions:

- Congestion management (queuing and scheduling)
- Congestion avoidance
- Traffic shaping and policy making

In addition, the Cisco IOS tools provide link efficiency mechanisms that integrate with the other three functions to provide additional improved QoS service.

Configuring QoS in Cisco IOS Modular QoS CLI

This topic describes the features of MQC.

Configuring QoS in Cisco IOS Modular QoS CLI

Cisco.com

- **Template-based command syntax for QoS**
- **Separates classification engine from the policy**
- **Uniform CLI for QoS features**
- **Cisco platform-independent**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—5-4

The Modular Quality of Service Command-Line Interface (MQC) was introduced to allow any supported classification to be used with any QoS mechanism.

The separation of classification from the QoS mechanism allows new IOS versions to introduce new QoS mechanisms and reuse all available classification options. In addition, old QoS mechanisms can benefit from new classification options.

Another important benefit of the MQC is the reusability of configuration. MQC allows the same QoS policy to be applied to multiple interfaces. CAR, by contrast, requires entire configurations to be copied and pasted between interfaces, and modifying configurations is tiresome.

The MQC, therefore, is a consolidation of all the QoS mechanisms that have so far been available only as standalone mechanisms.

Configuring QoS in Cisco IOS MQC Abstractions

This topic describes IOS QoS MQC abstractions.

Configuring QoS in Cisco IOS MQC Abstractions

Cisco.com

- **Class maps**
 - Define traffic classification criteria (for example, ACL, DSCP and IP precedence, MPLS EXP, etc.)
- **Policy maps**
 - Define QoS policy to apply to classes (marking, policing, shaping, queuing, dropping, etc.)
- **Service policy**
 - Apply QoS policy to interface for input or output traffic

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—5-5

Implementing QoS by using the MQC consists of the three steps that are presented in the following table.

Step	Action
1.	Configuring classification by using the class-map command
2.	Configuring traffic policy by associating the traffic class with one or more QoS features by using the policy-map command
3.	Attaching the traffic policy to inbound or outbound traffic on interfaces, subinterfaces, or virtual circuits by using the service-policy command

Class maps are used to create classification templates that are later used in policy maps where QoS mechanisms are bound to classes.

Routers can be configured with a large number of class maps (currently limited to 256). Each traffic policy, however, may support a limited number of classes (for example, CBWFQ and class-based LLQ are limited to 64 classes).

A class map is created using the **class-map** global configuration command. Class maps are identified by case-sensitive names. Each class map contains one or more conditions that determine if the packet belongs to the class.

There are two ways of processing conditions when there is more than one condition in a class map:

- **Match all:** All conditions have to be met to bind a packet to the class
- **Match any:** At least one condition has to be met to bind the packet to the class

The default match strategy of class maps is “match all.”

Class maps can classify packets by using the following classification tools:

- Access lists for any protocol can be used within the class-map configuration mode. The MQC can be used for other protocols, not only IP.
- IP packets can be classified directly by specifying IP precedence values.
- IP packets can also be classified directly by specifying IP DSCP values. DiffServ-enabled networks can have up to 64 classes if DSCP is used to mark packets.
- A QoS group parameter can be used to classify packets in situations where up to 100 classes are needed or there is a need to use the QoS group parameter as an intermediary marker (for example, MPLS-to-QoS-group translation on input and QoS-group-to-class translation on output).
- Packets can also be matched based on the value in the EXP bits of the MPLS header of labeled packets.

Classification can also be performed by identifying a Layer 3 or Layer 4 protocol. Where dynamic protocols are identified, advanced classification is also available by using the NBAR tool and inspecting higher-layer information.

There are many other classification options:

- Another class map can be used to implement template-based configurations.
- Packets can be matched based on the underlying Frame Relay discard eligible (DE) bit.
- Packets can be matched based on the information that is contained in the three class of service (CoS) bits (when you are using IEEE 802.1Q encapsulation) or priority bits (when you are using Inter-Switch Link [ISL] encapsulation).
- Packets can be classified according to the input interface.
- Packets can be matched based on their source or destination MAC addresses.
- RTP packets can be matched based on a range of User Datagram Protocol (UDP) port numbers.

MQC can also be used to implement a QoS mechanism for all traffic in which case classification will put all packets into one class.

Configuring QoS in Cisco IOS MQC Syntax

This topic describes the configuration of QoS in Cisco IOS MQC syntax.

Configuring QoS in Cisco IOS MQC Syntax

Cisco.com

```
class-map [match-any | match-all] class-name
```

Enters configuration submode for class definition

```
policy-map policy-name
```

Enters configuration submode for policy definition (marking, policing, shaping, queuing, etc.)

```
service-policy {input | output} policy-name
```

Applies QoS policy for input or output traffic in interface configuration submode

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-6

class-map

To create a class map to be used for matching packets to a specified class, use the **class-map** global configuration command. To remove an existing class map from the router, use the **no** form of this command.

- **class-map** *class-map-name*
- **no class-map** *class-map-name*

Syntax Description

Parameter	Description
<i>name</i>	Name of the class for the class map. The class name is used for both the class map and to configure policy for the class in the policy map.

policy-map

To create or modify a policy map that can be attached to one or more interfaces to specify a service policy, use the **policy-map** global configuration command. To delete a policy map, use the **no** form of this command.

- **policy-map** *policy-map-name*
- **no policy-map** *policy-map-name*

Syntax Description


Parameter	Description
<i>name</i>	Name of the policy map

service-policy

To apply QoS policy for input or output traffic, use the **service-policy** command in interface configuration submode.

MPLS QoS Configuration Case Study

This topic discusses the network diagram that will be used during the remainder of this lesson.



The figure above shows the complete network that will be broken down into various subnetworks for the following configuration study.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

Case Study Specifications

- Customer uses DiffServ codepoints for three classes of traffic, including VoIP.
- Service provider offers three classes of service:
 - Premium: Minimum bandwidth, low latency, no loss
 - Business: Minimum bandwidth, low loss
 - Best-effort: No guarantees


© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-8

The specifications for our case study are presented here.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-5-9

The figure shows the requirements for the Frame Relay part of the case study network.

The customer edge router (CE router) will perform the following on outbound packets:


- FRTS
- LLQ
- WRED
- FRF.12 (also known as FRF.11 Annex C)
- compressed Real-Time Transport Protocol (cRTP)

And the provider edge router (PE router) will perform the following on inbound packets:

- Policing
- Marking

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



CE-to-PE QoS for Frame Relay Access

- Traffic classified by IP precedence or DSCP (IP QoS)
- Limit bursting above CIR
- Traffic classified by IP LLQ for minimum bandwidth guarantees
- Fragmentation and cRTP on slow links

CE Outbound
FRTS
LLQ
WRED
FRF.12
cRTP

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—5-10

On the outbound CE router, the configuration will have the following characteristics:

- Traffic will be classified by the DSCP.
- Traffic bursting above the committed information rate (CIR) will be limited.
- LLQ will be used to guarantee minimum bandwidth.
- Packet fragmentation and cRTP will be used on slower links.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



CE Outbound

FRTS

LLQ

WRED

FRF.12

cRTP

CE-to-PE QoS for Frame Relay Access

```
class-map match-any PREMIUM
  match ip rtp 16384 16383
  match ip dscp ef
class-map match-all BUSINESS
  match ip dscp af21 af22 af23

policy-map OUT-POLICY
  class PREMIUM
    priority 128
    set ip dscp ef
  class BUSINESS
    bandwidth 256
    random-detect dscp-based
  class class-default
    fair-queue
    random-detect dscp-based
!
```

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—5-11

The figure shows the configuration of the outbound CE router.

```
!
class-map match-all PREMIUM
  match ip rtp 16384 16383
class-map match-all BUSINESS
  match ip dscp af21 af22 af23
!
policy-map OUT-POLICY
  class PREMIUM
    priority 128
    set ip dscp ef
  class BUSINESS
    bandwidth 256
    random-detect dscp-based
  class class-default
    fair-queue
    random-detect dscp-based
!
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



CE Outbound
FRTS
LLQ
WRED
FRF.12
cRTP

CE-to-PE QoS for Frame Relay Access

```
!  
interface Serial0/0.1 point-to-point  
ip address 10.10.1.2 255.255.255.0  
frame-relay interface-dlci 16  
class FR-class  
!  
map-class frame-relay FR-class  
  
frame-relay cir 512000  
frame-relay bc 512  
frame-relay mincir 512000  
frame-relay fair-queue  
service-policy output OUT-POLICY  
frame-relay fragment 512
```

The figure and the text below show a continuation of the configuration of the outbound CE router. Note the frame-relay cir and mincir statements.

```
!  
interface Serial0/0.1 point-to-point  
ip address 10.10.1.2 255.255.255.0  
frame-relay interface-dlci 16  
class FR-class  
!  
map-class frame-relay FR-class  
frame-relay cir 512000  
frame-relay bc 512  
frame-relay mincir 512000  
frame-relay fair-queue  
service-policy output OUT-POLICY  
frame-relay fragment 512
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com


CE-to-PE QoS for Frame Relay Access

- Mark and police traffic according to contract.
- Define IP precedence or DSCP mapping to EXP if needed (queue, drop precedence).
- P routers will service traffic based on EXP marking.

PE Inbound

Policing

Marking



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-13

On the inbound PE router, the configuration will have the following characteristics:


- Traffic will be marked and policed according to the service contract.
- If necessary, DSCP class (IP precedence) mapping will be defined for queue and drop precedence.
- P routers will service traffic based on the EXP bits marking.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

CE-to-PE QoS for Frame Relay Access

```
!
class-map match-all PREMIUM-IP
match ip dscp ef
class-map match-all BUSINESS-IP
match ip dscp af31 af32 af33
!
policy-map IN-POLICY
class PREMIUM-IP
  police 1280000 32000 32000
  conform-action set-mpls-exp-transmit 5
  exceed-action drop
class BUSINESS-IP
  police 22000000 550000 550000
  conform-action set-mpls-exp-transmit 4
  exceed-action set-mpls-exp-transmit 3
class class-default
  set mpls experimental 0
```



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-14

The figure shows the configuration of the inbound PE router.


```
!
class-map match-all PREMIUM-IP
  match ip dscp ef
class-map match-all BUSINESS-IP
  match ip dscp af31 af32 af33
!
policy-map IN-POLICY
  class PREMIUM-IP
    police 1280000 32000 32000
    conform-action set-mpls-exp-transmit 5
    exceed-action drop
  class BUSINESS-IP
    police 22000000 550000 550000
    conform-action set-mpls-exp-transmit 4
    exceed-action set-mpls-exp-transmit 3
  class class-default
    set mpls experimental 0
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

CE-to-PE QoS for Frame Relay Access

```
!  
interface Serial0/0.1 point-to-point  
ip address 10.32.14.2 255.255.255.0  
frame-relay interface-dlci 16  
  class FR-class  
!  
map-class frame-relay FR-class  
  frame-relay cir 512000  
  frame-relay bc 512  
  frame-relay mincir 512000  
  frame-relay fair-queue  
  service-policy input IN-POLICY  
  frame-relay fragment 512  
!
```



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-15

The figure and the text below show a continuation of the configuration of the inbound PE router. Note the service-policy statement.


```
!  
interface Serial0/0.1 point-to-point  
ip address 10.32.14.2 255.255.255.0  
frame-relay interface-dlci 16  
  class FR-class  
!  
map-class frame-relay FR-class  
  frame-relay cir 512000  
  frame-relay bc 512  
  frame-relay mincir 512000  
  frame-relay fair-queue  
  service-policy input IN-POLICY  
  frame-relay fragment 512  
!
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

PE-to-P QoS for Frame-Mode MPLS

7200 Series
7500 Series
10000 ESR
12000 GSR




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-16

Provider routers (P routers) will service traffic based on the EXP bits marking.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



PE-to-P QoS for Frame-Mode MPLS

PE Outbound

- Traffic classified by EXP bits (MPLS QoS)
- LLQ for queuing MPLS packets
- WRED based on the EXP bits to implement dropping precedence
- IP precedence copied to MPLS the EXP bits if no mapping defined in input policy

MPLS

LLQ

WRED

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-17


The configuration of an outbound PE router to a P router will have the following characteristics:

- The traffic will be classified by the EXP bits (MPLS QoS).
- LLQ will be used for queuing MPLS packets.
- WRED will be based on the EXP bits to implement dropping precedence.
- IP precedence will be copied to the MPLS EXP bits if no mapping is defined in the input policy.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

PE-to-P QoS for Frame-Mode MPLS



MPLS

LLQ

WRED

© 2004 Cisco Systems, Inc. All rights reserved.

PE Outbound

```
!
class-map match-all PREMIUM
  match mpls experimental 5
!
class-map match-all BUSINESS
  match mpls experimental 4
!
policy-map OUT-POLICY
  class PREMIUM
    priority 16384
  class BUSINESS
    bandwidth 65536
    random-detect
  class class-default
    random-detect
!
interface POS1/0
  ip address 10.150.1.1 255.255.255.0
  service-policy output OUT-POLICY
  mpls ip
```

MPLST v2.0-5-18


The figure and the text below show the configuration of the outbound PE router to the P router.

```
!
class-map match-all PREMIUM
  match mpls experimental 5
!
class-map match-all BUSINESS
  match mpls experimental 4
!
policy-map OUT-POLICY
  class PREMIUM
    priority 16384
  class BUSINESS
    bandwidth 65536
    random-detect
  class class-default
    random-detect
!
interface POS1/0
  ip address 10.150.1.1 255.255.255.0
  service-policy output OUT-POLICY
  tag-switching ip
```


MPLS QoS Configuration Case Study (Cont.)

Cisco.com

P-to-P QoS for Frame-Mode MPLS



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-19

The configuration of a P router to another P router will have the following characteristics:

- LLQ will be used for queuing MPLS packets and will incorporate Modified Deficit Round Robin (MDRR).
- WRED will be used to implement dropping precedence.


MPLS QoS Configuration Case Study (Cont.)

Cisco.com

P-to-P QoS for Frame-Mode MPLS

P Outbound

- Traffic classified by the EXP bits (MPLS QoS)
- LLQ (MDRR) for queuing MPLS packets
- WRED based on the EXP bits
- Same CLI on GSR for IP and MPLS packets (MQC in the future)
- No need for inbound policy



P Router

LLQ (MDRR)

WRED

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-20

The configuration of a P router to another P router will have the following characteristics:


- The traffic will be classified by the EXP bits (MPLS QoS).
- LLQ will be used for queuing MPLS packets and will incorporate MDRR.
- WRED will be based on the EXP bits.
- The same command-line interface (CLI) will be used on the Gigabit Switch Router (GSR) for IP and MPLS packets (MQC in the future).
- There will be no need for an inbound policy.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

```
!
P Outbound
class-map match-all PREMIUM
match mpls experimental 5
!
class-map match-all BUSINESS
match mpls experimental 4
!
policy-map OUT-POLICY
class PREMIUM
priority 16384
class BUSINESS
bandwidth 65536
random-detect
class class-default
random-detect
!
interface POS2/0
ip address 10.150.1.1 255.255.255.0
service-policy output OUT-POLICY
mpls ip
```

P-to-P QoS for Frame-Mode MPLS



P Router

LLQ or MDRR

WRED

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-21


The figure and the text below show the configuration of an outbound P router.

```
interface POS2/0
ip add 10.64.12.1 255.255.255.252
tag-switching ip
tx-cos OUT-POLICY
!
cos-queue-group OUT-POLICY
precedence 3 queue 1
precedence 4 queue 1
precedence 5 queue low-latency
precedence 0 random-detect-label 0precedence 3 random-
detect-label 1
precedence 4 random-detect-label 2
random-detect-label 0 300 500 1
random-detect-label 1 100 300 1
random-detect-label 2 300 500 1
queue 0 50queue 1 50
queue low-latency strict-priority!
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

P-to-PE QoS for Frame-Mode MPLS



© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—5-22

This figure identifies the features that should be implemented in the outbound P router's link to PE router.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com

P-to-PE QoS for Frame-Mode MPLS



P Outbound

- Traffic classified by EXP bits (MPLS QoS)
- LLQ for queuing MPLS packets
- WRED based on the EXP bits to implement dropping precedence
- Upstream PE should advertise explicit NULL if PHP exposes IP packet

MPLS

LLQ

WRED

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—5-23

The configuration of an outbound P router to a PE router will have the following characteristics:

- Traffic will be classified by using the EXP bits (MPLS QoS).
- LLQ will be used for queuing MPLS packets.
- WRED, based on the EXP bits, will be used to implement dropping precedence.
- The upstream PE should advertise explicit NULL if penultimate hop popping (PHP) exposes the IP packet.

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



MPLS
LLQ
WRED

P-to-PE QoS for Frame-Mode MPLS


```
!
class-map match-all PREMIUM
match mpls experimental 5
!
class-map match-all BUSINESS
match mpls experimental 4
!
policy-map OUT-POLICY
class PREMIUM
priority 16384
class BUSINESS
bandwidth 65536
random-detect
class class-default
random-detect
!
interface POS1/0
ip address 10.150.1.1 255.255.255.0
service-policy output OUT-POLICY
tag-switching ip
```

The figure and the text below show the configuration of the outbound P router to the PE router.

```
!
class-map match-all PREMIUM
match mpls experimental 5
!
class-map match-all BUSINESS
match mpls experimental 3 4
!
policy-map OUT-POLICY
class PREMIUM
priority 16384
class BUSINESS
bandwidth 65536
random-detect
class class-default
random-detect
!
interface POS1/0
ip address 10.150.1.1 255.255.255.0
service-policy output OUT-POLICY
tag-switching ip
```

MPLS QoS Configuration Case Study (Cont.)

Cisco.com



QoS Summary

- **Marking, policing, and shaping at the edges of the network**
- **IP precedence or DSCP mapping to EXP if needed**
- **Queuing and dropping of MPLS packets in the core based on MPLS EXP**
- **One class for low latency traffic and a small number of additional classes to implement SLAs**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-25

Here are several things to remember when you are implementing QoS in MPLS Virtual Private Networks (VPNs):

- Marking, policing, and shaping should be performed on the edge of the network.
- If it is needed, map IP precedence or DSCP to the EXP bits.
- Any queuing or dropping of MPLS packets in the core is done based on the MPLS EXP bits.
- Recommend the use of one class for low latency traffic and a smaller number of additional classes to implement SLAs

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- Cisco IOS QoS provides for the following major functions:
 - Congestion management (queuing and scheduling)
 - Congestion avoidance
 - Traffic shaping and policy making
- MQC (Modular QoS CLI) is a consolidation of all QoS mechanisms that have previously been available only as standalone
- MQC is implemented by the following:
 - class-map command
 - policy-map command
 - service-policy command
- The service-policy command is used in interface configuration submode

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—5-26

References

For additional information, refer to this resource:

- Search the Cisco.com website at <http://www.cisco.com> for more information about implementing MPLS QoS in Cisco IOS software.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which of the following is NOT a feature of the Cisco IOS QoS tools?
- A) congestion avoidance
 - B) congestion ,marking
 - C) congestion management
 - D) traffic shaping and policy making
- Q2) MQC was introduced to allow which of the following?
- A) any supported classification to be used with any QoS mechanism
 - B) any new classification to be used with any QoS mechanism
 - C) any supported QoS classification be used with any queuing mechanism
 - D) any new queuing classification to be used with any QoS mechanism
- Q3) Which of the following is NOT a step when you are implementing QoS using MQC?
- A) class maps
 - B) link maps
 - C) policy maps
 - D) service policy
- Q4) The command that is used in interface configuration submode to apply a QoS policy for input or output traffic is the _____ command.

Lesson Answer Key

Q1) B

Relates to: How to Use the QoS Toolkit

Q2) A

Relates to: Configuring QoS in Cisco IOS Modular QoS CLI

Q3) B

Relates to: Configuring QoS in Cisco IOS MQC Abstractions

Q4) **service-policy**

Relates to: Configuring QoS in Cisco IOS MQC Syntax

QoS in MPLS Applications

Overview

This lesson focuses on how QoS is integrated in MPLS implementations.

Relevance

This lesson is mandatory for learners who are planning to implement MPLS QoS on Cisco IOS platforms.

Objectives

This lesson describes how QoS works with MPLS applications. Upon completing this lesson, you will be able to do the following:

- Identify how MPLS-TE is implemented in a best-effort network
- Identify how MPLS-TE is implemented in a DiffServ network
- Identify how MPLS DS-TE is implemented in a DiffServ network
- Identify how QoS is enabled in an MPLS VPN environment
- Identify some QoS implementation tasks

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Successful completion of the “MPLS Traffic Engineering Technology” and “Configuring MPLS Traffic Engineering” modules of this course


Outline

This lesson includes these topics:

- Overview
- MPLS-TE with a Best-Effort Network
- MPLS-TE with a DiffServ Network
- MPLS DS-TE with a DiffServ Network
- QoS-Enabled MPLS VPNs
- QoS Implementation
- Lesson Summary
- Lesson Review

MPLS-TE with a Best-Effort Network

This topic describes how MPLS-TE is implemented in a best-effort network.




Best-effort is a single service model in which an application sends data whenever it must, in any quantity, and without requesting permission or first informing the network. For best-effort service, the network delivers data if it can, without any assurance of reliability, delay bounds, or throughput.

The Cisco IOS QoS feature that implements best-effort service is first-in, first-out (FIFO) queuing.

Best-effort service is suitable for a wide range of networked applications such as general file transfers or e-mail.

MPLS-TE with a DiffServ Network

This topic describes how MPLS-TE is implemented in a DiffServ network.



In this model, packets entering a DiffServ-enabled network are grouped into a small number of classes. Each class has a color or mark that is associated with it (use of the DSCP bits). This marking makes packet classification extremely scalable and ensures appropriate bandwidth and delay guarantees in the network core. Each node within the core network is applied to different queuing and dropping policies on every packet, based on the marking that the packet carries (PHB).

In the MPLS + DiffServ architecture, packets marked with the DSCP will enter the MPLS network and the PHB is enforced by every LSR along the path. Because LSRs do not have any knowledge of the IP header, the PHB needs to be achieved by looking at different information.

Two general approaches are used to mark MPLS traffic for QoS handling within an MPLS network.

In the first method, the DiffServ coloring information is mapped in the EXP field of the MPLS shim header. This field allows up to eight different QoS markings versus 64 for the DSCP. The packet scheduling (PHB) at each hop (in the MPLS clouds) is done based on the EXP field. LSPs that use this approach are called E-LSPs, where QoS information is inferred from the EXP bits.


Alternatively, the label that is associated with each MPLS packet carries the portion of the DiffServ marking that specifies how a packet should be queued. The dropping precedence portion of the DiffServ marking is carried in the EXP bits (if an MPLS shim header is being used) or on fields that are available for this purpose on underlying technologies (CLP bit in ATM, DE bit for Frame Relay).

The ingress LSR examines the DSCP in the IP header (the CLP bit for ATM and the DE bit for Frame Relay) and selects an LSP that has been provisioned for that QoS level. At the egress point, as the label is removed, the packet is sent to the next IP hop with its original DSCP. LSPs that use this approach are called L-LSPs, where QoS information is inferred in part from the MPLS label.

Caution TE does not differentiate among traffic types. 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, to provide the necessarily stricter QoS guarantees.

MPLS DS-TE with a DiffServ Network

This topic describes how MPLS DS-TE is implemented in a DiffServ network.



MPLS-TE allows constraint-based routing (CBR) of IP traffic. One of the constraints that is satisfied by CBR is the availability of required bandwidth over a selected path. DiffServ-aware TE extends MPLS-TE to enable you to perform CBR of “guaranteed” traffic, which satisfies a more restrictive bandwidth constraint than that satisfied by CBR for regular traffic. The more restrictive bandwidth is termed a “subpool,” while the regular TE tunnel bandwidth is called the “global pool.” (The subpool is a portion of the global pool.)

This ability to satisfy a more restrictive bandwidth constraint translates into an ability to achieve higher QoS performance (in terms of delay, jitter, or loss) for the guaranteed traffic.

For example, DiffServ-Aware Traffic Engineering (DS-TE) can be used to ensure that traffic is routed over the network so that, on every link, there is never more than 40 percent (or any assigned percentage) of the link capacity of guaranteed traffic (for example, voice), while there can be up to 100 percent of the link capacity of regular traffic.

DS-TE is an enhancement to MPLS-TE that introduces the concept of classes (class types to be more precise) to TE. Each participating link advertises the amount of available bandwidth of each class type on that link. When the CBR process is executed for a new tunnel, a bandwidth constraint of a particular class type can be defined as one of the criteria to be used for the path selection.

The admission control process that is carried using RSVP at each hop is performed against the available bandwidth of the specific class type. As an example, you can define a tunnel between point A and B of the network with 10 Mbps of available bandwidth of a particular class type.

The current IOS implementation supports two class types. These two bandwidth amounts are not treated as two separate amounts. Instead, each link defines a bandwidth pool and bandwidth subpool. The bandwidth subpool is part of the bandwidth pool. As an example that does not over-subscribe the line, an OC-3 or STM-1 link can have a bandwidth pool with 155 Mbps and a bandwidth subpool with 55 Mbps.

Available bandwidth, in both the global pool and the subpool, is advertised by the IGP link-state advertisement (LSA) or TLV objects, which ensure 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 SLA requirement, can choose to overbook lower-priority classes or even underbook higher-priority traffic to meet tight QoS requirements.


Assuming QoS mechanisms are also used on every link to queue guaranteed traffic separate from regular traffic, it then becomes possible to enforce separate “overbooking” ratios for guaranteed and regular traffic. (In fact, for the guaranteed traffic it becomes possible to enforce no overbooking at all—or even an “underbooking”—so that very high QoS can be achieved end-to-end for that traffic, even while for the regular traffic a significant overbooking continues to be enforced.)

Also, through the ability to enforce a maximum percentage of guaranteed traffic on any link, the network administrator can directly control the end-to-end QoS performance parameters without having to rely on overengineering or on expected shortest path routing behavior. This ability is essential for transport of applications that have very high QoS requirements (such as real-time voice, virtual IP leased lines, and bandwidth trading), where overengineering cannot be assumed to be present everywhere in the network.

DS-TE involves extending the Open Shortest Path First (OSPF) routing protocol so that the available subpool bandwidth at each pre-emption level is advertised in addition to the available global pool bandwidth at each pre-emption level. And DS-TE modifies CBR to take this more complex advertised information into account during path computation.

QoS-Enabled MPLS VPNs

This topic describes how QoS is enabled in an MPLS VPN environment.




The two models listed in the figure are models that are used in the context of MPLS VPNs that refer to how QoS guarantees are offered.

In the point-to-cloud model, each site receives a single QoS guarantee for traffic that is sent to and received from all other VPN sites. Two parameters are defined for this purpose, Ingress Committed Rate (ICR) and Egress Committed Rate (ECR). In a VPN with three sites, each site will have a single QoS guarantee for all incoming traffic (regardless of the traffic source) and all outgoing traffic (regardless of the traffic destination).

In the point-to-point model, VPN sites have specific QoS guarantees to other VPN sites. In a VPN with three sites, each site will have specific QoS guarantees to the other two sites that belong to the VPN. This approach can be used to offer hard QoS guarantees that are equivalent to the guarantees available with ATM or Frame Relay.

QoS-Enabled MPLS VPNs (Cont.)

Cisco.com



Point-to-Cloud Connection

- Per-VPN QoS policies at the edge.
- Same MPLS QoS policies for packets of all VPNs in the core.
- QoS can be implemented with point-to-network guarantees.
- QoS can also be implemented with point-to-point.

In a point-to-cloud topology, each site receives a single QoS guarantee for traffic that is sent to and received from all other VPN sites.

Service providers that offer QoS services will want to provide an ICR guarantee and an ECR guarantee, possibly for each service class that is offered.

ICR refers to the traffic rate coming into the service provider network, given a particular treatment (premium, gold, and so on) throughout the service provider network.

ECR refers to the traffic rate that is given a particular treatment from the service provider to the customer site.

As long as traffic does not exceed ICR and ECR limits, the network provides bandwidth and delay guarantees. For example, as long as HTTP traffic does not exceed 1 Mbps (into the network and out of the network to the customer site), bandwidth and low delay are guaranteed.

This is the point-to-cloud model because, for QoS purposes, the service provider does not need to keep track of traffic destinations, as long as the destinations are within the ICR and ECR bounds.


With Cisco IOS MPLS, the QoS guarantees of a service provider can be “transparent” to customers. In other words, a service provider can provide these guarantees in a nonintrusive way. Customer sites can deploy a consistent, end-to-end DiffServ implementation without having to adapt to the QoS implementation of the service provider. The service provider can prioritize customer priority traffic without remarking the DSCP field of the IP packet. A separate marking is used to provide QoS within the MPLS network, and it is discarded when the traffic leaves the MPLS domain. The QoS marking that is delivered to the destination network corresponds to the marking that was received when the traffic entered the MPLS network.

QoS-Enabled MPLS VPNs (Cont.)

Cisco.com

Point-to-Point Connection

- Point-to-point (site-to-site) guarantees.
- DS-TE is required to offer hard point-to-point guarantees.
- Point-to-network and point-to-point model are not mutually exclusive.



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-5-8

In the point-to-point model, VPN sites have specific QoS guarantees with other VPN sites.

For the more stringent applications, where a customer desires a point-to-point guarantee, a virtual data pipe needs to be constructed to deliver the highly critical traffic. For example, an enterprise may want two hub sites or data centers that are connected with high SLA guarantees.

DS-TE supports its QoS model by automatically choosing a routing path that satisfies the bandwidth constraint for each service class that is defined (such as premium, gold, silver, or bronze). DS-TE also relieves the service provider from having to compute the appropriate path for each customer, and each service class, per customer. This is referred to as the point-to-point model.

The two MPLS VPN QoS models (point-to-point and point-to-cloud) are not mutually exclusive. A service provider can offer VPN services to a customer where some sites have a hard (point-to-point) guarantee to other sites while yet other sites have just a soft (point-to-cloud) guarantee.

QoS Implementation

This topic describes some of the QoS implementation tasks.

QoS Implementation

Cisco.com

Create and use four or five main traffic classes:

- **Mission-Critical Applications**
 - **Business-critical client-server applications**
- **Transactional and Interactive Applications**
 - **All other client-server or interactive (e.g., messaging) applications**
- **Bulk Data Applications**
 - **FTP, e-mail, backup operations, content distribution**
- **Best-Effort Applications**
 - **Default class**
- **Optional: Scavenger Applications**
 - **Peer-to-peer applications, gaming traffic**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-9

Traffic classes are defined by the network administrator to meet the business requirements of an organization or SLAs. They represent a way to categorize network traffic.

Here are some recommended guidelines for setting up traffic classes:

- Assign as few as possible applications to the mission-critical traffic class.
- Seek executive endorsement of mission-critical application assignments prior to rolling out QoS policies for data.
- Most applications fall into the best-effort class. Make sure that adequate bandwidth is provisioned for this default class.


Some Examples:

Application Class	Example Applications	Application Traffic Properties	Packet, or Message, Sizes
Interactive	Telnet, Citrix, Oracle Thin Clients, AOL Instant Messenger, Yahoo Instant Messenger, PlaceWare (Conference), NetMeeting Whiteboard	Highly interactive applications with tight user feedback requirements.	Average Message Size < 100 B Max message size < 1 KB
Transactional	SAP, PeopleSoft – Vantive, Oracle – Financials + Internet Procurement + Business to Business + Supply Chain Management + Application Server, Oracle 8i Database, Ariba Buyer, I2, Siebel, E.piphany, Broadvision, IBM Business to Business, Microsoft SQL, Lotus Notes, Microsoft Outlook, BEA Systems, E-mail Download (SMTP), DLSw+	Transactional applications typically use a client-server protocol model. User-initiated, client-based queries followed by server response. Query response may consist of many messages between client and server. Query response may consist of many TCP and FTP sessions running simultaneously (e.g., HTTP-based applications).	Depends on application Could be anywhere from 1 KB to 50 MB
Bulk Data	Database syncs, network-based backups, video content distribution, large FTP file transfers	Long file transfers. Always invokes TCP congestion management.	Average message size 64 KB or greater
Best-Effort	All noncritical traffic, HTTP web browsing, and other miscellaneous traffic		

QoS Implementation (Cont.)

Cisco.com

Minimal Baseline Class >> Queue Mapping Five Classes of Service



In the example here, “application class” refers to the type of traffic. The DSCP column shows the DSCP code for the traffic type, and the queues are the user-defined traffic classes.

The example shows that the network administrator has defined five traffic classes and that the real-time traffic class consists of the following traffic types:

- Voice
- Call signaling

The video traffic class consists of the following traffic types:

- Videoconferencing
- Streaming video

The business traffic class consists of the following traffic types:

- Routing
- Locally defined
- Transactional data
- Network management

The bulk traffic class consists of the following traffic type:

- Bulk data


And the best-effort traffic class consists of the following traffic types:

- Scavenger
- Best-effort

QoS Implementation (Cont.)

Cisco.com

Minimal Baseline Class >> Queue Mapping Four Classes of Service



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-11

This example shows that the network administrator has defined four traffic classes and that the bulk traffic class from the previous example has been eliminated.

The real-time traffic class consists of the following traffic types:

- Voice
- Call signaling

The video traffic class consists of the following traffic types:

- Videoconferencing
- Streaming video

The business traffic class consists of the following traffic types:

- Routing
- Locally defined
- Transactional data
- Network management


And the best-effort traffic class consists of the following traffic types:

- Scavenger
- Best-effort
- Bulk data

QoS Implementation (Cont.)

Cisco.com

Minimal Baseline Class >> Queue Mapping Three Classes of Service



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—5-12

This example shows that the network administrator has defined three traffic classes and that the video traffic class from the previous example has been eliminated.

The real-time traffic class consists of the following traffic types:

- Voice
- Call signaling
- Videoconferencing

The business traffic class consists of the following traffic types:

- Routing
- Locally defined
- Transactional data
- Network management
- Streaming video
- Bulk data

And the best-effort traffic class consists of the following traffic types:

- Scavenger
- Best-effort

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

This lesson presented these key points:

- **Best-effort service does not provide any assurance of reliability, delay bounds, or throughput.**
- **In an MPLS-TE network with DiffServ, packets are scheduled at every hop according to EXP marking regardless of LSP.**
- **In an MPLS DS-TE network with DiffServ, LSRs advertise multiple available bandwidths via IGP.**
- **In a QoS-enabled MPLS VPN point-to-cloud implementation, each site receives a single QoS guarantee for traffic that is sent and received from all other VPN sites.**
- **When implementing QoS, the network administrator should recommend no more than four or five main traffic classes.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—5-13

References

For additional information, refer to these resources:

- Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.1
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/qos_c/index.htm
- Cisco IOS Quality of Service Solutions Command Reference, Release 12.1
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/qos_r/index.htm
- Configuring Multiprotocol Label Switching
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/switch_c/xcprt4/xcdtagc.htm
- MPLS Class of Service Enhancements
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t5/mc1214t.htm>

Next Steps

For the associated lab exercises, refer to the following sections of the course Lab Guide:

- Lab Exercise 5-1: Implementing MPLS QoS
- Lab Exercise 5-2: Implementing MPLS QoS and DS-TE

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) In a best-effort network, LSRs advertise which of the following?
- A) two bandwidth pools via IGP
 - B) single available bandwidth via BGP
 - C) single available bandwidth via IGP
 - D) single available bandwidth via LDP
- Q2) In an MPLS-TE network with DiffServ, packets are scheduled at every hop according to the _____ marking regardless of the LSP.
- Q3) In an MPLS DS-TE network with DiffServ support enabled, what IGP functions are used to advertise the available bandwidth (both the global pool and the subpool)?
- Q4) In a point-to-network VPN implementation, QoS is implemented _____ but not at the _____.
- Q5) In MPLS QoS implementations, _____ define traffic classes (which categorize network traffic).

Lesson Answer Key

- Q1) C
Relates to: MPLS-TE with a Best-Effort Network
- Q2) EXP
Relates to: MPLS-TE with a DiffServ Network
- Q3) LSA, TLVs (order is not important)
Relates to: MPLS DS-TE with a DiffServ Network
- Q4) at the edge, core
Relates to: QoS-Enabled MPLS VPNs
- Q5) network administrators
Relates to: QoS Implementation

Any Transport over MPLS

Overview

Multiprotocol Label Switching (MPLS) offers a set of advanced features to provide enhanced services. This module looks at one of the most recent supported features: Any Transport over MPLS (AToM). AToM provides transport of Layer 2 frames and cells over an MPLS network.

Objectives

Upon completing this module, you will be able to use a diagram of a typical simple MPLS service provider solution to identify the appropriate Cisco IOS MPLS advanced feature and describe how that feature is implemented. This includes being able to do the following:

- Identify the reasons behind implementing AToM support, the types of Layer 2 protocols that are supported, and how AToM provides transport for these protocols
- Configure AToM on Cisco IOS platforms to implement supported solutions
- Identify the Cisco IOS commands and syntax that is used to verify AToM operations

Outline

The module contains these lessons:

- Introduction to Any Transport over MPLS
- Configuring AToM on Cisco IOS Platforms
- Monitoring AToM on Cisco IOS Platforms

Introduction to Any Transport over MPLS

Overview

There is an ever-increasing demand for the transport of Layer 2 and Layer 3 over a common backbone. This lesson introduces the MPLS application Any Transport over MPLS (AToM). AToM allows an MPLS network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, High-Level Data Link Control (HDLC), Frame Relay, and ATM.

Relevance

To successfully implement AToM, it is necessary to understand which Layer 2 protocols are supported and how AToM provides transport for these protocols.

Objectives

This lesson describes the reasons behind implementing AToM support, the types of Layer 2 protocols that are supported, and how AToM provides transport for these protocols. Upon completing this lesson, you will be able to do the following:

- Describe the benefits of AToM
- At a high level, describe how AToM transports different Layer 2 frame types and cells across an MPLS network
- Describe how AToM establishes end-to-end connectivity for a Layer 2 frame
- Describe the AToM control word

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS VPN technology and configuration knowledge
- Advanced BGP knowledge

Outline

This lesson includes these topics:

- Overview
- AToM Overview
- Transport Types
- How AToM Works
- AToM Control Word
- Lesson Summary
- Lesson Review

AToM Overview

This topic describes the benefits of using AToM and the reasons why a service provider would implement AToM support.

AToM Overview

Cisco.com

- **Many service providers offer Layer 2 transport services.**
- **Overlay VPNs were built using ATM or Frame Relay PVCs.**
- **Building both a Layer 2 and a Layer 3 network is costly.**
- **Service providers want to unify VPN services and Internet services in a single network: an MPLS-based network.**
- **However, some customers still want Layer 2 connections (Ethernet VLANs or PVCs).**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-3

Many service providers currently offer Layer 2 transport services to their customers. These services are offered over a circuit-based infrastructure to build Layer 2 Virtual Private Networks (VPNs).

Initially, VPNs were built using leased lines. Later, service providers offered Layer 2 VPNs that were based on point-to-point data link layer connectivity, using ATM or Frame Relay virtual circuits. Customers built their own Layer 3 networks to accommodate IP traffic. As a result, separate networks exist for Layer 2 and Layer 3 traffic. But maintaining separate networks for Layer 2 VPNs and Internet traffic is difficult and costly. So service providers want a single IP-based network to provide both Layer 2 and Layer 3 services.

The MPLS VPN was introduced to meet the requirement for a unified network for Layer 3 VPN services. But some customers still want Layer 2 connections. These connections can be in the form of Ethernet VLAN extensions across a metropolitan area or ATM services. The AToM was introduced to facilitate Layer 2 connectivity across an MPLS backbone.

AToM benefits service providers that offer Layer 2 connectivity to customers with traditional offerings such as ATM, Frame Relay, and serial or PPP services. Additionally, it serves providers who are specializing in Ethernet connectivity in metropolitan areas. Services for Layer 2 VPNs also appeal to the enterprise customers of service providers—customers who may already run many of these networks and want just point-to-point connectivity.

AToM Overview (Cont.)

Cisco.com

Benefits

- **Incorporates Layer 2 and Layer 3 services over a common infrastructure**
- **Scalability: MPLS does not keep state information about virtual circuits inside core**
- **Maintains support for existing services while migrating to MPLS**
- **Customer sites independent of service provider backbone**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-4

One way that service providers can create a better network is to incorporate Layer 2 and Layer 3 services over a common infrastructure, such as an MPLS backbone. MPLS-enabled VPNs (Layer 2 or Layer 3) also offer performance, scalability, and new value-added services using other MPLS enhancements such as traffic engineering (TE), Fast Reroute (FRR), and quality of service (QoS).


Many service providers have currently implemented connection-oriented ATM networks in the core. ATM provides QoS, bandwidth, and the ability to perform TE. However, an ATM network without MPLS does not scale well for VPNs because it relies on virtual circuit state information in the core. For every virtual circuit that is created as part of a Layer 2 VPN, the core switches keep the complete state for the virtual circuit.

As service providers migrate to MPLS-based networks, they need to maintain support for existing services, including the ability to transport Layer 2 traffic. AToM is the architectural framework for accomplishing that, regardless of whether the Layer 2 service is ATM, Frame Relay, Ethernet, PPP, or HDLC. AToM encapsulates packets at the provider edge router (PE router), transports them over the backbone to a PE router on the other side of the cloud, unencapsulates them, and sends them to their destination.

The upgrade from a real Layer 2 ATM or Frame-Relay-based network to an MPLS-based network that provides the ATM or Frame Relay services by using AToM is transparent to customers. Unlike the Layer 3 IP-based VPNs using MPLS, the service provider does not participate in the Layer 3 routing of the customer. The service provider provides Layer 2 connectivity only.

AtoM Overview (Cont.)

Cisco.com



- **AToM is not designed specifically for switch replacement.**
- **It is used for interconnecting and transporting a traffic POP across an MPLS core network.**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-5

This example illustrates how the MPLS network is used to extend a Frame Relay point of presence (POP).

Frame Relay over MPLS, as a flavor of AToM, is not designed specifically for Frame Relay switch replacement. It is used for interconnecting and transporting a Frame Relay POP across an MPLS core network.

Transport Types

This topic provides a high-level description of how AToM transports different Layer 2 frame types and cells across an MPLS network.

Transport Types

Cisco.com

AToM enables the following types of Layer 2 frames and cells to be directed across an MPLS backbone:

- **Ethernet, Ethernet VLAN**
- **ATM adaptation layer 5 (AAL5)**
- **ATM cell relay**
- **Frame Relay**
- **Point-to-Point Protocol (PPP)**
- **High-Level Data Link Control (HDLC)**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—6-6

AToM enables the following types of Layer 2 frames and cells to be directed across an MPLS backbone:

- Ethernet, Ethernet VLAN
- ATM adaptation layer 5 (AAL5)
- ATM cell relay
- Frame Relay
- PPP
- HDLC

The first phase of AToM development in Cisco IOS software supports Ethernet and AAL5 like-to-like connectivity. This support requires that the same transport type be at each end of the network. In the future, AToM will be enhanced to support more transport types and also to provide interworking functions that can connect disparate transport types at each end, such as Frame Relay at one end connecting to an Ethernet VLAN at the other.

How AToM Works

This topic describes how AToM establishes end-to-end connectivity for a Layer 2 frame.

How AToM Works

Cisco.com

AToM frames are carried across an MPLS backbone in the following manner:

- Ingress and egress interfaces are non-MPLS interfaces.
- Ingress PE encapsulates frame into MPLS; egress PE decapsulates.
- Label stack of two labels is used (similar to an MPLS VPN).
 - Topmost label (“tunnel label”) used for LSP PE to PE.
 - Second label (“VC label”) identifies outgoing interface in the egress PE.
- LDP has been extended to carry virtual circuit forwarding equivalence class (VC FEC).
- A directed (multihop) LDP session is used from PE to PE.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-7

AToM is used to forward Layer 2 frames. Frames are received on an ingress interface by the ingress PE router. At this point, the frame is a raw Layer 2 frame. The ingress PE router encapsulates it into MPLS and tunnels it across the backbone to the egress PE router. The egress PE router decapsulates the packet and reproduces the raw Layer 2 frame on the egress interface.


The AToM frames are carried across the MPLS backbone using a label stack of two labels. The topmost label is used to propagate the packet from the ingress PE router to the correct egress PE router. The second label is used by the egress PE router to forward the packet out on the correct interface. This process is somewhat similar to an MPLS/VPN, where the egress PE router uses a VPN label.

In the AToM framework, the topmost label is called the “tunnel label.” This name indicates that its use is to tunnel the packet across the MPLS backbone to the egress PE router. The second label is called the “VC label.” The name indicates that its use is to map the packet to an outgoing virtual circuit or link.

A protocol is required between the PE routers so that they can exchange the VC labels. In AToM, the Label Distribution Protocol (LDP) is used for this purpose. A directed multihop LDP session is established between the PE routers. The egress PE router sends an LDP message in which it indicates the label value to use for a virtual circuit forwarding equivalence class (VC FEC). That label value is then used by the ingress PE router as the second label in the label stack that is imposed to the frames of the indicated VC FEC.

How AToM Works (Cont.)

Cisco.com



- The IGP and the LDP between directly connected LSRs establish one LSP in each direction.
- A directed LDP session between PE routers is established.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-6-8

In this example, the big cloud in the middle is the MPLS backbone that provides the AToM services.


An Interior Gateway Protocol (IGP) is used inside the MPLS backbone. This setup means that all routers that are inside the backbone have routing information about how to send IP packets to each other. LDP is also used between directly connected neighbors. Local labels are assigned to each IGP-derived route. The label values are then propagated to the neighbor across the LDP session.

The IGP together with the LDP sessions between directly connected neighbors establishes label switched paths (LSPs) from any router inside the backbone to any other router inside the backbone. In the figure, two LSPs are illustrated, one LSP in each direction between the ingress and egress PE router. The tunnel label is used to propagate the packets along the LSP to the correct egress PE router.

The figure also shows the directed multihop LDP session between ingress and egress PE router that is used to exchange the VC label. Any ingress-egress PE router pair will need such an LDP session. Only one is illustrated in the figure.

How AToM Works (Cont.)

Cisco.com



- LDP between directly connected LSRs generates an LSP.
- The egress PE allocates VC label 17.
- The directed LDP session between PE routers propagates the VC label.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-9


The IGP and LDP in combination are used to create an LSP from ingress PE router to egress PE router. The figure shows the label assignment and advertisement. The egress PE router advertises the pop label for its own loopback address. The backbone router that is closest to it advertises label value 23. The next router advertises value 22, and the last label advertisement shown is value 21. An LSP from ingress router to egress router is now established.

The egress PE router now allocates a local label to be the VC label for the specific circuit in this example. It selects the label value 17 for this. The VC label is advertised to the ingress PE router using the directed LDP session between them.

The ingress PE router now forms a label stack. The topmost label, the tunnel label, has the value 21 and is used to guide the packets to the egress PE router. The second label, the VC label, has the value 17 and is used by the egress PE router to propagate the packets out on the correct interface.

How AToM Works (Cont.)

Cisco.com



- The ingress PE receives a frame on DLCI 101.
- The frame is encapsulated and forwarded along the LSP.
- The egress PE reconstructs the frame.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-10


The ingress PE receives a Frame Relay frame on data-link connection identifier (DLCI) 101 on the incoming interface. The DLCI is mapped to the AToM tunnel across the backbone. The Frame Relay frame is therefore encapsulated into MPLS using the label stack with label 21 as the topmost label and label 17 as the second label.

The packet is then forwarded along the LSP. The topmost label is used for label swapping in the next hop. The top label is changed to the value 22. In the next hop, label swapping results in label value 23 being the top label. In the router just before the egress router, the incoming label value 23 indicates pop. That label therefore performs penultimate hop popping (PHP). The topmost label is removed, and the packet is propagated to the egress PE router with the label value 17, the VC label, which is now the only label left.

When the PE router receives the packet with label value 17, that label value instructs the PE router to decapsulate the packet and send it out on the associated Frame Relay DLCI. In this case the DLCI value is 202. The Frame Relay frame is now reconstructed and transmitted.

VC Label Withdrawal Procedure

Cisco.com



- If a PE router detects a condition that affects normal service, it must withdraw the corresponding VC label by means of LDP signaling.
- A PE router may provide circuit status signaling:
 - Frame Relay must use LMI procedures; ATM should use ILMI procedures.

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—6-11

If a PE router detects a condition that affects normal service, it *must* withdraw the corresponding VC label. This withdrawal is done by means of LDP.

A PE router may provide circuit status signaling on the interface where the customer connects. A PE router that provides Frame Relay services *must* use Local Management Interface (LMI) procedures with the customer equipment. A PE router that provides ATM services *should* use Integrated Local Management Interface (ILMI) procedures.

AToM Control Word

This topic describes the format of the AToM control word.



The figure illustrates AToM encapsulation. The topmost label is the tunnel label. This label is followed by the VC label. Following the VC label is the optional control word. Next in the packet is the Layer 2 protocol data unit (PDU).

The ingress PE router encapsulates the Layer 2 frames into MPLS. It does that by removing any preamble, any frame checksums, and also some header information. Essential header information is carried across the MPLS backbone in a control word, a shim header.

The control word is also used for sequence numbering to guarantee sequenced delivery, if it is required.

The control word is optional. Some Layer 2 protocols make use of it, while others do not. Both endpoints (ingress PE and egress PE) must, of course, agree to either use or not use the control word. It is transmitted after the label stack but before the Layer 2 PDU. It can be used to carry important Layer 2 header information and to guarantee sequenced delivery (if required).

The control word and the Layer 2 PDU are MPLS-encapsulated with the label stack of two labels.

The control word is 32 bits. It is divided into four fields. The first field is four reserved bits, which must always be set to zero. The next field is a four-bit flag field. The flags have different uses depending on the Layer 2 protocol that is being forwarded. The third field is an eight-bit length field, which is used only if the Layer 2 PDU is shorter than the minimum MPLS packet and padding is required. If no padding is required, the length field is not used. The fourth field is a 16-bit sequence number. Sequence numbering is used only on Layer 2 protocols, which

guarantees ordered delivery. A special value of 0 in the sequence field indicates that there is no guaranteed sequenced delivery.

When AToM is used for Frame Relay over MPLS (FRoMPLS), the Frame Relay header is removed and the forward explicit congestion notification (FECN), backward explicit congestion notification (BECN), discard eligible (DE), and command/response (C/R) bits are carried in the control word flag field.

When AToM is used for ATM over MPLS, the first flag in the control word flag field is used to indicate whether it is ATM adaptation layer 5 (AAL5) frames or raw ATM cells that are being transported by AToM. The other three flags are used for explicit forward congestion indication (EFCI), cell loss priority (CLP), and C/R. If one of these flags is set in any of the ATM cells that are being transported in the MPLS packet, then the corresponding flag is set in the control word.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **AToM provides scalable Layer 2 and Layer 3 services over a common infrastructure.**
- **AToM can provide transport for:**
 - Ethernet
 - PPP
 - HDLC
 - Frame Relay
 - ATM
- **AToM uses the IGP and LDP in combination to create an LSP from ingress PE router to egress PE router.**
- **This LSP creates a virtual circuit.**
- **The ingress PE router will encapsulate the Layer 2 frames into MPLS and forward them over the virtual circuit.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-13

References

For additional information, refer to this resource:

- Access the Cisco website at <http://www.cisco.com> for additional information about the unified VPN suite solution for service providers.

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) What are the benefits of AToM?

Q2) How can Layer 2 frames be propagated by AToM across an IP MPLS-based backbone?

Q3) Which Layer 2 transport types are currently supported by AToM?

Lesson Answer Key

Q1) AToM allows the service provider to incorporate Layer 2 and Layer 3 services over a common infrastructure.

Relates to: AToM Overview

Q2) AToM uses the IGP and LDP in combination to create an LSP from ingress PE router to egress PE router. This LSP is assigned a VC label. The Layer 2 frame is encapsulated in an MPLS frame and forwarded over the virtual circuit.

Relates to: How AToM Works

Q3)	Ethernet	PPP
	HDLC	Frame Relay
	ATM (AAL5)	ATM (cell relay)

Relates to: Transport Types

Configuring AToM on Cisco IOS Platforms

Overview

While the configuration of AToM support is not complicated, network administrators must take care to ensure that the proper configuration tasks, parameters, and syntax are used. In addition, AToM provides commands to verify the operation of AToM.

Relevance

To successfully implement an AToM network, network administrators need to be able to identify the command syntax that is required. This syntax includes not only the syntax of the actual commands but also includes the proper configuration level and order.

Objectives

This lesson lists the tasks, parameters, syntax, and descriptions for configuring AToM on Cisco IOS platforms, and also covers how to verify the operation of AToM. Upon completing this lesson, you will be able to do the following:

- Identify the effects of AToM on the MTU
- Identify PE loopback summarization problems
- Enable AToM
- Enable a typical EoMPLS configuration
- Enable a typical PPP configuration
- Enable a typical HDLC configuration
- Enable a typical FRoMPLS configuration
- Enable a typical AAL5 over MPLS configuration
- Enable a typical ATM over MPLS configuration

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS VPN technology and configuration knowledge
- Advanced BGP knowledge

Outline

This lesson includes these topics:

- Overview
- MTU Issues
- AToM Packet Forwarding with Summarization in the Core
- AToM Configuration
- EoMPLS Configuration
- PPP over MPLS Configuration
- HDLC over MPLS Configuration
- FRoMPLS Port Configuration
- AAL5 over MPLS Configuration
- ATM over MPLS Configuration
- Lesson Summary
- Lesson Review

MTU Issues

This topic describes the effects of AToM on the maximum transmission unit (MTU).

MTU Issues

Cisco.com

- **AToM transport of Frame Relay, Ethernet, and AAL5 does not allow packets to be fragmented and reassembled.**
- **Ensure that the MTU of all intermediate links between endpoints is sufficient to carry the largest Layer 2 frame received.**
- **The ingress and egress PE routers must have the same MTU value.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-3

Unlike IP, most Layer 2 protocols (for example, Frame Relay, Ethernet, and AAL5) do not allow fragmentation of frames. This fact has two implications:


- All intermediate links between the ingress PE router and the egress PE router must be able to carry the largest Layer 2 frame that has been received, including the imposed label stack and the 4-byte control word (if it is used).
- The ingress PE interface and the egress PE interface must have the same MTU value.

Failure to comply with the first rule means that the larger frames, where the label stack and the control word contribute to creating a larger size than can be carried, will be dropped by the backbone.

Failure to comply with the second rule means that frames that are received on one PE router will be dropped by the other if the frame size is too big for the egress PE interface.

AToM Packet Forwarding with Summarization in the Core

This topic describes the problems that you can encounter when the PE loopback addresses are summarized.



A requirement for successful propagation of AToM packets across an MPLS backbone is an unbroken LSP between PE routers. The second label in the stack is recognized only by the egress PE router that has originated it and would not be understood by any other router, should it ever become exposed.

There are two scenarios that can cause the LSP between PE routers to break:

- If the IP address of the PE router is announced as a Border Gateway Protocol (BGP) route, it has no corresponding LDP label and the label stack cannot be built correctly.
- If the provider routers (P routers) perform summarization of the address range within which the IP address of the egress PE router lies, the LSP will be disrupted at the summarization point, as illustrated in the figure.

In the figure, the P router summarizes the loopback address of the egress PE router. LSP is broken at a summarization point because the summarizing router needs to perform a full IP lookup. In a frame-based MPLS network, the P router would request PHP for the summary route and the upstream P router (or a PE router) would remove the LDP label, exposing the VPN label to the P router. However, because the VC label was not assigned by the P router, but by the egress PE router, the label will not be understood by the P router and the AToM packet will be dropped or misrouted.

AToM Configuration

This topic identifies the command syntax that is required to enable AToM support.

AToM Configuration

Cisco.com

AToM is configured using the following steps:

- **The PE routers must have a /32 address assigned to their loopbacks.**
- **MPLS must be enabled in the core.**
- **Make sure MTU is large enough in the core.**
- **Enable Layer 2 frame transport in both endpoint PE routers.**
- **Make sure MTU is same on both endpoint interfaces.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-5

AToM is configured in the PE routers using the steps in the following table.

Step	Action
1.	The PE routers must have a /32 address assigned to their loopbacks. This address will be the transport address for the directed LDP session.
2.	MPLS must be enabled and correctly configured in the network core.
3.	Make sure that the MTU is large enough on all the core links. Any possible path between the ingress and egress PE routers must be able to carry the largest packet size that is anticipated.
4.	Configure Layer 2 frame transport in both endpoint PE routers.
5.	Make sure that the MTU is the same on both endpoint interfaces.

AToM Configuration: Configuring the LDP ID

Cisco.com

```
router(config)#
```

```
mpls ldp router-id interface [force]
```

Specifies a preferred interface for determining the Label Distribution Protocol (LDP) router ID.

Parameters:

- *interface* Causes the IP address of the specified interface to be used as the LDP router ID, provided that the interface is operational.
- *force* Alters the behavior of the `mpls ldp router-id` command to force the use of the named interface as the LDP router ID .

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-6

mpls ldp router-id

To specify a preferred interface for determining the LDP router ID, use the **mpls ldp router-id** command in global configuration mode. To remove the preferred interface for determining the LDP router ID, use the **no** form of this command.

- **mpls ldp router-id *interface* [force]**
- **no mpls ldp router-id**

Syntax Description

Parameter	Description
<i>interface</i>	Causes the IP address of the specified interface to be used as the LDP router ID, provided that the interface is operational.
<i>force</i>	(Optional) Alters the behavior of the mpls ldp router-id command to force the use of the named interface as the LDP router ID. For more information, see the mpls ldp router-id command "Usage Guidelines" section.

Defaults

The **mpls ldp router-id** command is disabled.

Usage Guidelines

If the **mpls ldp router-id** command is not executed, the LDP router ID is determined as follows:

1. The IP addresses of all operational interfaces are examined.
2. If these IP addresses include loopback interface addresses, the largest such loopback address is selected as the LDP router ID.
3. Otherwise, the largest IP address that pertains to an operational interface is selected as the LDP router ID.

The normal (default) method for determining the LDP router ID may result in a router ID that is not usable in certain situations. For example, an IP address that is selected as the LDP router ID might be an address that the routing protocol is not able to advertise to a neighboring router. The **mpls ldp router-id** command provides a means for specifying an interface whose IP address is to be used as the LDP router ID. The specified interface must be operational for its IP address to be used as the LDP router ID.

When it is executed without the **force** keyword, the **mpls ldp router-id** command modifies the method for determining the LDP router ID by causing selection of the IP address of the specified *interface* argument (provided that the interface is operational) the next time that it is necessary to select an LDP router ID. The effect of the command is thus delayed until this need arises, which is typically the next time that the interface whose address is the current LDP router ID is shut down or the address itself is not configured.

When it is executed with the **force** keyword, the effect of the **mpls ldp router-id** command depends on the current state of the specified interface:

- If the interface is up (operational) when the **mpls ldp router-id force** command is issued and if its IP address is not currently the LDP router ID, the LDP router ID is forcibly changed to the IP address of the interface. This forced change in the LDP router ID tears down any existing LDP sessions, releases label bindings that have been learned via the LDP sessions, and interrupts MPLS forwarding activity that is associated with the bindings.
- If the interface is down (not operational) when the **mpls ldp router-id force** command is issued, when the interface transitions to up, the LDP router ID is forcibly changed to the IP address of the interface. This forced change in the LDP router ID tears down any existing LDP sessions, releases label bindings that have been learned via the LDP sessions, and interrupts MPLS forwarding activity that is associated with the bindings.

AToM Configuration: Configuring L2 Transport Route

Cisco.com

```
router(config-if)#
```

```
mpls l2transport route destination vc-id
```

Enables Layer 2 routing over an emulated VC across the MPLS cloud

Parameters:

- ***destination*** Specifies the IP address of the remote PE router loopback.
- ***vc-id*** Assigns a VC ID to an interface. The VC ID must be unique to each virtual circuit.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-7

mpls l2transport route

To enable routing of Layer 2 VLAN packets over a specified virtual circuit, use the **mpls l2transport route** command in interface configuration mode. To disable routing over the specified virtual circuit, use the **no** form of this command.

- **mpls l2transport route** *destination vc-id*
- **no mpls l2transport route** *destination vc-id*

Syntax Description

Parameter	Description
<i>destination</i>	Specifies the IP address of the router to which the virtual circuit is destined.
<i>vc-id</i>	Assigns a virtual circuit ID (VC ID) to a router. The VC ID must be unique to each virtual circuit. Specify the same VC ID on both ends of the virtual circuit.

Defaults

No default behavior or values.

Usage Guidelines


An MPLS Layer 2 VLAN virtual circuit runs across an MPLS cloud to connect VLAN interfaces on two PE routers.

Use this command on the VLAN interface of each PE router to route Layer 2 VLAN packets across the MPLS cloud to the interface of the other PE router. Specify the IP address of the other PE router for the *destination* parameter. Do not specify the IP address of the router from which you are issuing the command.

You can choose any number for the VC ID. However, the VC ID must be unique to the virtual circuit. Therefore, in large networks, it may be necessary to track the VC ID assignments to ensure that a VC ID does not get assigned twice.

EoMPLS Configuration

This topic identifies the commands to enable a typical Ethernet over MPLS (EoMPLS) implementation.



The figure illustrates how the MPLS backbone provides a Layer 2 Ethernet port-to-port connection between the two customer devices CE-A1 and CE-A2.

CE-A1 and CE-A2 are exchanging Ethernet frames. PE-Site-X takes whatever Ethernet frame that is produced by CE-A1 on the link to PE-Site-X, encapsulates it into MPLS, and forwards it across the backbone to PE-Site-Y.


PE-Site-Y decapsulates the packet and reproduces the Ethernet frame on its link toward CE-A2.

The AToM feature does not include any MAC layer address learning and filtering. That means that the PE-Site-X and PE-Site-Y routers do not filter any frames that are based on those addresses.

In addition, the AToM feature does not use Spanning Tree Protocol (STP). Bridge protocol data units (BPDUs) are propagated transparently and not processed. LAN loop detection must be performed by other functions or avoided by design.

EoMPLS Configuration (Cont.)

Cisco.com



- The two customer devices, CE-A1 and CE-A2, exchange Ethernet frames.
- PE-Site-X takes whatever Ethernet frames that arrive on the VLAN subinterface and forwards them as an MPLS packet across the MPLS backbone to PE-Site-Y.
- The EoMPLS service does not do MAC address learning and filtering and does not take part in the Spanning Tree Protocol. These functions can be performed by the Cisco catalyst 6000 series switches.
- PE-Site-Y removes MPLS headers and propagates the Ethernet frame out on the VLAN subinterface on its side.

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.0—6-9

This figure shows how the two customer devices CE-A1 and CE-A2 are connected to Ethernet LAN switches (Cisco Catalyst 6000 series). These LAN switches are then connected to the MPLS edge routers PE-Site-X and PE-Site-Y.


The interface encapsulation between the Cisco Catalyst 6000 series LAN switches and the PE routers supports VLANs. Different subinterfaces in the PE router are used to connect to different VLANs. The PE-Site-X subinterface to the VLAN where the CE-A1 device is connected is used for AToM forwarding. Whatever Ethernet frames that arrive on that VLAN subinterface are encapsulated into MPLS and forwarded across the backbone to PE-Site-Y.

PE-Site-Y decapsulates the packet and reproduces the Ethernet frame on the outgoing subinterface.

Although MAC address learning and STP are not features of AToM, combining AToM with a LAN switch with those features allows the service provider to make use of them.

EoMPLS Configuration (Cont.)

Cisco.com



In this example configuration, PE-Site-X is configured to use EoMPLS to propagate VLAN 200 across the backbone to PE-Site-Y.

PE-Site-X has a loopback interface with IP address 2.2.2.2. PE-Site-Y is using address 4.4.4.4.

The subinterface GigabitEthernet0/0.2 is used to connect to VLAN 200 on the 802.1Q encapsulated Ethernet interface. This subinterface is configured for EoMPLS. The Ethernet frames are MPLS-encapsulated and forwarded to the PE router with IP address 4.4.4.4, which is PE-Site-Y.


The VC identifier, value 123, is used to associate the connection with the other end. It is required that both endpoints use the same VC identifier value.

Note There is no requirement that the VLAN identifier be the same on both ends. In this example the dot1Q VLAN ID 200 is used in PE-Site-X while it is rewritten to VLAN ID 201 when frames are transmitted by PE-Site-Y on the dot1Q trunk.

The most important detail is the VC identifier. The value 123 is used in both PE-Site-X and in PE-Site-Y. That is the information that is used by Cisco IOS software to connect the two VLANs.

PPP over MPLS Configuration

This topic identifies the commands to enable a typical PPP implementation.




PPP over MPLS enables two CE devices to use a PPP link-level connection across the MPLS backbone.

The PPP link-level frames that are received by PE-Site-X from CE-A1 are encapsulated into MPLS and propagated across the backbone to PE-Site-Y, which decapsulates them and reproduces them on the interface toward CE-A2.

PPP over MPLS Configuration (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-12

In this example configuration, PE-Site-X is configured to use AToM to propagate PPP frames across the backbone to PE-Site-Y.

PE-Site-X has a loopback interface with IP address 2.2.2.2. PE-Site-Y is using address 4.4.4.4.


The subinterface Serial1/0 is used to connect to CE-A1 using PPP encapsulation. This subinterface is configured to transport PPP frames. The PPP frames are MPLS-encapsulated and forwarded to the PE router with IP address 4.4.4.4, which is PE-Site-Y.

The VC identifier, value 123, is used to associate the connection with the other end. It is required that both endpoints use the same VC identifier value.

The most important detail is the VC identifier. The value 123 is used in both PE-Site-X and in PE-Site-Y.

HDLC over MPLS Configuration


This topic identifies the commands to enable a typical Cisco HDLC implementation.



HDLC over MPLS performs the same operations as PPP over MPLS. However, it does it with Cisco HDLC frames.

HDLC over MPLS Configuration (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-14

In this example configuration, PE-Site-X is configured to use AToM to propagate Cisco HDLC frames across the backbone to PE-Site-Y.

PE-Site-X has a loopback interface with IP address 2.2.2.2. PE-Site-Y is using address 4.4.4.4.


The subinterface Serial1/0 is used to connect to CE-A1 using Cisco HDLC encapsulation. This subinterface is configured to transport HDLC frames. The HDLC frames are MPLS-encapsulated and forwarded to the PE router with IP address 4.4.4.4, which is PE-Site-Y.

The VC identifier, value 123, is used to associate the connection with the other end. It is required that both endpoints use the same VC identifier value.

The most important detail is the VC identifier. The value 123 is used in both PE-Site-X and in PE-Site-Y.

FRoMPLS Port Configuration

This topic identifies the commands to enable a typical FRoMPLS implementation.




FRoMPLS works in either PVC-to-PVC mode or port-to-port mode.

Frame Relay frames that are received by PE-Site-X from CE-A1 are encapsulated into MPLS and forwarded across the backbone to PE-Site-Y, where they are decapsulated and reproduced on the interface to CE-A2.

FRoMPLS: Port-to-Port Configuration (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-16


In this example, the PE-Site-X and PE-Site-Y use FRoMPLS to connect two Frame Relay ports.

In this implementation the Frame Relay LMI session takes place between the two end devices (CE-A1 and CE-A2). The entire Frame Relay frame is transported over the MPLS network. Therefore, the encapsulation on the PE interface is *not* set to Frame Relay as might be expected. The encapsulation is set for HDLC.

The qualifier **sequence** to the **mpls l2transport** command is used to enable the sequencing feature of AToM.

FRoMPLS Port Configuration: DLCI-to-DLCI

Cisco.com



In this example, PE-Site-X and PE-Site-Y use FRoMPLS to connect two Frame Relay DLCIs. Frame Relay is enabled on the physical port. In this implementation the Frame Relay LMI sessions take place between the end devices (CE-A1 and CE-A2) and their associated PE router. Therefore, the encapsulation on the PE interface is set to Frame Relay to provide Frame Relay support. The **frame-relay intf-type** can be either DTE or DCE.

This implementation differs from those that have been presented previously in this lesson in that the **mpls l2transport** command is not entered under the interface sublevel. The **connect** (Frame Relay) command is used to define connections between Frame Relay permanent virtual circuits (PVCs).

Another difference is the number of connections that must be set up for each interface. Because the connection is DLCI-based rather than port-based, a connection must be established for each DLCI that is to be transported. While this requirement adds to the complexity of the configuration, it allows you to establish an AToM transport terminal for each DLCI on a different PE router.

connect (Frame Relay)

To define connections between Frame Relay PVCs, use the **connect** command in global configuration mode. To remove connections, use the **no** form of this command.

- **connect** *connection-name interface dlci {interface dlci | l2transport}*
- **no connect** *connection-name interface dlci {interface dlci | l2transport}*

Syntax Description

Parameter	Description
<i>interface</i>	Interface on which a PVC connection will be defined
<i>dlci</i>	DLCI number of the PVC that will be connected
l2transport	Specifies that the PVC will not be a locally switched PVC, but will be tunneled over the backbone network

Defaults

No default behavior or values.

Command Modes

Global configuration

Usage Guidelines

When Frame Relay switching is enabled, the **connect** command creates switched PVCs in Frame Relay networks.


Example

The following example shows how to enable Frame Relay switching and define a connection called “frompls1” with a DLCI of 100 on serial interface 5/0.

```
PE1_router(config)# connect frompls1 Serial5/0 100 l2transport
```

AAL5 over MPLS Configuration

This topic identifies the commands to enable a typical AAL5 over MPLS implementation.



CE-A1 and CE-A2 are exchanging ATM cells. But in this scenario, all cells that are received on the interface between provider edge (PE) and customer edge (CE) are forwarded, regardless of the virtual path identifier (VPI) or virtual channel identifier (VCI) values in the cell headers. There is also no concept of ATM adaptation layers. It is a raw port-to-port connection.

An ingress PE router receives an ATM AAL5 service data unit (SDU) and removes the header.

The PE router copies the control word elements from the header to the corresponding fields in the control word of the SDU. The control word contains the following information:


- **EFCI bit:** The EFCI bit is used by ATM switches to indicate congestion that is being experienced by forwarded data cells.
- **CLP bit:** The CLP bit indicates whether a cell should be dropped if it encounters extreme congestion as it moves through the ATM network.

The PE router adds a VC label and an LSP tunnel label to the packet for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish ATM AAL5 traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

All cells that are received by PE-Site-X on the ATM interface from CE-A1 are encapsulated into MPLS and propagated across the backbone to PE-Site-Y. At that point, PE-Site-Y decapsulates the packet and reproduces the cells on the outgoing interface to CE-A2.

AAL5 over MPLS Configuration (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-19

In this example, PE-Site-X (the leftmost PE) and PE-Site-Y (the rightmost PE) are using AAL5 over MPLS.

The **pvc** command and the **encapsulation** command work together. How you use the commands for AToM is slightly different than for all other applications.

The following list highlights the differences:


- **pvc command:** For most applications, you create a PVC by using the **pvc vpi/vci** command. For AToM, you must add the **l2transport** keyword to the **pvc** command. The **l2transport** keyword enables the PVC to transport Layer 2 packets.
- **encapsulation command:** The **encapsulation** command for AToM has only two keyword values: **aal5** or **aal0**. You cannot specify an encapsulation type. In contrast, the **encapsulation aal5** command that you use for most other applications requires you to specify the encapsulation type, such as **aal5snap**.
- **pvc command and encapsulation command:** The AToM **encapsulation** command works only with the **pvc** command. You cannot create switched virtual circuits or virtual circuit bundles to transport Layer 2 packets. You can use only PVCs to transport Layer 2 packets.

When you use the **aal5** keyword, incoming cells (except Operation, Administration, and Maintenance [OAM] cells) on the PVC are treated as AAL5-encapsulated packets. The router reassembles the packet from the incoming cells. The router does not check the contents of the packet, so it does not need to know the encapsulation type (such as **aal5snap**, **aal5mux**, and so on). After imposing the MPLS label stack, the router sends the reassembled packet over the MPLS core network.

When you use the **aal0** keyword, the router strips the Header Error Control (HEC) byte from the cell header and adds the MPLS label stack. The router sends the cell over the MPLS core network.

ATM over MPLS Configuration

This topic identifies the commands to enable a typical ATM over MPLS implementation.



ATM Cell Relay over MPLS transports single ATM cells over the MPLS backbone. The AToM circuit is configuring on PVCs. ATM Cell Relay over MPLS encapsulates ATM cells in MPLS packets and forwards them across the MPLS network. Each MPLS packet contains one ATM cell. In other words, each ATM cell is transported as a single packet. The following steps outline the process of encapsulating the ATM cell.

Ingress PE Router

The ingress PE router receives an ATM cell and removes the header. The following items are not removed from the ATM cell:

- **The control word:** The control word contains the following information:
 - **EFCI bit:** The EFCI bit is used by ATM switches to indicate congestion that is being experienced by forwarded data cells.
 - **CLP bit:** The CLP bit indicates whether a cell should be dropped if it encounters extreme congestion as it moves through the ATM network.
- **The VPI and VCI:** The VPI and VCI identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI and VCI fields to identify the next virtual channel link (VCL) that a cell needs to transit on its way to its final destination.

The PE router adds a VC label and an LSP tunnel label to the packet for normal MPLS routing through the MPLS backbone. The P routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish ATM cell relay traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.


Egress PE Router

At the other edge of the MPLS backbone, the egress PE router receives the packet and removes the LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router has removed that label. The PE router also removes the control word and VC label from the packet.

The PE router adds an ATM header and sends the packet out the appropriate customer-facing interface.

ATM over MPLS Configuration (Cont.)

Cisco.com



In this example PE-Site-X (leftmost PE) and PE-Site-Y (rightmost PE) are using ATM Cell Relay over MPLS.

The **pvc** command and the **encapsulation** command work together. How you use the commands for AToM is slightly different than for all other applications.

The following list highlights the differences:

- **pvc command:** For most applications, you create a PVC by using the **pvc vpi/vci** command. For AToM, you must add the **l2transport** keyword to the **pvc** command. The **l2transport** keyword enables the PVC to transport Layer 2 packets.
- **encapsulation command:** The **encapsulation** command for AToM has only two keyword values: **aal5** or **aal0**. You cannot specify an encapsulation type. In contrast, the **encapsulation aal5** command that you use for most other applications requires you to specify the encapsulation type, such as **aal5snap**.
- **pvc command and encapsulation command:** The AToM **encapsulation** command works only with the **pvc** command. You cannot create switched virtual circuits or virtual circuit bundles to transport Layer 2 packets. You can use only PVCs to transport Layer 2 packets.

When you use the **aal5** keyword, incoming cells (except OAM cells) on the PVC are treated as AAL5-encapsulated packets. The router reassembles the packet from the incoming cells. The router does not check the contents of the packet, so it does not need to know the encapsulation type (such as **aal5snap**, **aal5mux**, and so on). After imposing the MPLS label stack, the router sends the reassembled packet over the MPLS core network.

When you use the **aal0** keyword, the router strips the HEC byte from the cell header and adds the MPLS label stack. The router sends the cell over the MPLS core network.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Most Layer 2 protocols do not allow fragmentation of frames.**
 - Both the ingress and egress interface must have the same MTU value.
 - All links between PE routers must be able to carry the largest frame and the imposed labels without fragmentation.
- **There must be an unbroken label switched path between ingress PE and egress PE router. Summarization of PE router loopbacks is one of the common causes of breaking the LSP.**
- **AToM is configured in the PE routers using the following steps:**
 - Make sure that the PE routers have a /32 address assigned to their loopbacks.
 - Enable MPLS and configure it correctly in the network core.
 - Make sure that the MTU is large enough on all the core links.
 - Configure Layer 2 frame transport in both endpoint PE routers.
 - Make sure that the MTU is the same on both endpoint interfaces.

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-22

References

For additional information, refer to these resources:

- Access the Cisco website at <http://www.cisco.com> for additional information about the unified VPN suite solution for service providers.
- MPLS AToM: Configuring at http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/products_feature_guide_chapter09186a0080103d7a.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) What MTU issues must be considered when you are implementing AToM?

Q2) Why can you not summarize the loopback addresses of the PE routers by using the IGP?

Q3) What are the required steps to configure AToM?

Q4) Which command is used to enable AToM support on an interface?

Q5) Which command is used to enable AToM transport of the following Layer 2 protocols on an interface?

Ethernet VLANs	
PPP	
HDLC	
Frame Relay	
AAL5	

Lesson Answer Key

- Q1) Both ingress and egress interface must have the same MTU value. All links between PE routers must be able to carry the largest frame and the imposed labels without fragmentation.

Relates to: MTU Issues

- Q2) There must be an unbroken LSP between ingress PE and egress PE router. Summarization of PE router loopbacks is one of the common causes of breaking the LSP.

Relates to: AToM Packet Forwarding with Summarization in the Core

- Q3) Make sure that the PE routers have a /32 address assigned to their loopbacks.
Enable MPLS and configure it correctly in the network core.
Make sure that the MTU is large enough on all the core links.
Configure Layer 2 frame transport in both endpoint PE routers.
Make sure that the MTU is the same on both endpoint interfaces.

Relates to: AToM Configuration

- Q4) **mpls l2transport route destination vc-id**

Relates to: AToM Configuration

Q5) Ethernet VLANs	<u>encapsulation dot1Q</u>
PPP	<u>encapsulation ppp</u>
HDLC	<u>encapsulation hdlc</u>
Frame Relay	<u>encapsulation frame-relay</u>
AAL5	<u>encapsulation aal5</u>

Relates to: EoMPLS Configuration
PPP over MPLS Configuration
HDLC over MPLS Configuration
FRoMPLS Port Configuration
AAL5 over MPLS Configuration

Monitoring AToM on Cisco IOS Platforms

Overview

While configuring AToM support is not complicated, network administrators must take care to ensure proper operation. This lesson describes the commands to verify the operation of AToM.

Relevance

To successfully support an AToM network, network administrators need to be able to identify the commands that are available to monitor and troubleshoot AToM operations.

Objectives

The previous lesson listed the configuration tasks, parameters, syntax, and descriptions for configuring AToM on Cisco IOS platforms. This lesson covers how to monitor the operation of AToM. Upon completing this lesson, you will be able to do the following:

- Identify the commands that are available to monitor AToM operation
- Identify the function, syntax, and output of the **show mpls l2transport vc** command
- Identify the function, syntax, and output of the **show mpls l2transport vc detail** command
- Identify the function, syntax, and output of the **show mpls l2transport summary** command
- Identify the function, syntax, and output of the **debug mpls l2transport vlan control** command

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Core MPLS knowledge
- MPLS and VPN technology and configuration knowledge
- Advanced BGP knowledge

Outline

This lesson includes these topics:

- Overview
- Monitoring AToM
- show mpls 12transport vc
- show mpls 12transport vc detail
- show mpls 12transport summary
- debug mpls 12transport vlan control
- Lesson Summary
- Lesson Review

Monitoring AToM

This topic identifies the commands that are available to monitor AToM operations.

Monitoring AToM

Cisco.com

- **Virtual-circui-based:**
 - show mpls l2transport vc
 - show mpls l2transport vc detail
 - show mpls l2transport summary
 - debug mpls l2transport vlan control
- **AAL5-based:**
 - **Existing AAL5 commands**
 - http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1612/products_feature_guide09186a00800e9570.html

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-3

The existing MPLS **show** and **debug** commands can be used to monitor the underlying MPLS core network. In addition, a new set of AToM-specific commands have been introduced.

show mpls l2transport vc

This topic describes the function, syntax, and output of the **show mpls l2transport vc** command.

show mpls l2transport vc

Cisco.com

```
router#  
show mpls l2transport vc {summary} | {vc-id} | {{vc-id-min} {vc-id-max}} [detail]
```

- Displays the state of VCs on a router

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-4

show mpls l2transport vc

To display the state of virtual circuits on a router, use the **show mpls l2transport vc EXEC** command.

show mpls l2transport vc [**vcid** *vc-id*] | [*vc-id-min* *vc-id-max*] [**interface** *name* [*local-circuit-id*]] [**destination** *ip-address* | *name*] [**detail**]


Syntax Description

Parameter	Description
vcid	(Optional) The VC ID that is assigned to the router.
<i>vc-id</i>	Displays information about the virtual circuit that you specify.
<i>vc-id-min</i> and <i>vc-id-max</i>	Displays information about a range of VC IDs that you specify. The range is from 0 to 429,467,295.
interface	(Optional) The interface or subinterface of the router that has been enabled to transport Layer 2 packets. This keyword lets you display information about the virtual circuits that have been assigned VC IDs on that interface or subinterface.
<i>name</i>	(Optional) The name of the interface or subinterface.

Parameter	Description
<i>local-circuit-id</i>	(Optional) The number that is assigned to the local circuit. This argument value applies only to the following transport types: • For Frame Relay, enter the DLCI of the PVC. • For AAL5 and cell relay, enter the VPI and VCI of the PVC. • For Ethernet VLANs, enter the VLAN number.
destination	(Optional) Information about the virtual circuits that have been assigned VC IDs for the remote router that you specify.
<i>ip-address</i>	(Optional) The IP address of the remote router.
<i>name</i>	(Optional) The name that is assigned to the remote router.
detail	(Optional) Displays detailed information about the virtual circuits on a PE router.

show mpls l2transport vc (Cont.)

Cisco.com



```

PE-site-X# 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
  
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-5

show mpls l2transport vc Field Descriptions

Syntax Description

Field	Description
Transport VC ID	The virtual circuit identifier that is assigned to one of the interfaces on the router.
Client Intf	The ingress or egress interface through which the Layer 2 VLAN packet travels. For EoMPLS, VLAN interfaces are used.
VC State	The status of the virtual circuit. The status can be one of the following: UP—The virtual circuit is in a state where it can carry traffic between the two virtual circuit endpoints. A virtual circuit is up when both imposition and disposition interfaces are programmed. The disposition interface is programmed if the virtual circuit has been configured and the client interface is up. The imposition interface is programmed if the disposition interface is programmed and there is a remote VC label and an IGP label. The IGP label can be implicit null in a back-to-back configuration. (An IGP label means there is an LSP to the peer.) DOWN—The virtual circuit is not ready to carry traffic between the two virtual circuit endpoints.
Local VC Label	The VC label that a router signals to its peer router, which is used by the peer router during imposition. The local VC label is a disposition label. The local VC label determines the egress interface of an arriving packet from the MPLS backbone.
Remote VC Label	The disposition VC label of the remote peer router.
Tunnel Label	An IGP label that is used to route the packet over the MPLS backbone to the destination router with the egress interface.

show mpls l2transport vc detail

This topic describes the syntax and output of the **show mpls l2transport vc detail** command.

show mpls l2transport summary

Cisco.com

MPLS Backbone

```
graph LR; CE-A1 --- PE-Site-X; PE-Site-X --- PE-Site-Y; PE-Site-Y --- CE-A2
```

```
PE-Site-X# show mpls l2transport vc 2 detail
VC ID: 2, Local Group ID: 8, Remote Group ID: 8 (VC is up)
Client Intf: V12 is up, Destination: 172.21.0.1, Peer LDP Ident: 172.20.0.1:0
Local VC Label: 21, Remove VC Label: 21, Tunnel Label: 22
Outgoing Interface: Gi3/2, Next Hop: 4.0.0.1
Local MTU: 1500, Remote MTU: 1500
Imposition: LC Programmed,
Current Imposition/Disposition Slot: 3/32
Packet Totals(in/out): 803713123/802954183
Byte Totals(in/out): 2067870672/942882144
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—6-7


show mpls l2transport vc detail Field Descriptions

Field	Description
VC ID	The virtual circuit identifier that is assigned to one of the interfaces on the router.
Local Group ID	The ID that is used to group virtual circuits locally. EoMPLS groups virtual circuits by the hardware port, which is unique for each port on a router.
Remote Group ID	The ID that is used by the peer to group several virtual circuits.
VC is up or VC is down	The status of the virtual circuit. The status can be one of the following: UP—The virtual circuit is in a state where it can carry traffic between the two virtual circuit endpoints. A virtual circuit is up when both imposition and disposition interfaces are programmed. The disposition interface is programmed if the virtual circuit has been configured and the client interface is up. The imposition interface is programmed if the disposition interface is programmed and there is a remote VC label and an IGP label. The IGP label can be implicit null in a back-to-back configuration. (An IGP label means there is an LSP to the peer.) DOWN—The virtual circuit is not ready to carry traffic between the two virtual circuit endpoints.
Client Intf	The ingress or egress interface through which the Layer 2 VLAN packet travels. For EoMPLS, VLAN interfaces are used.

Field	Description
Destination	The destination that you specify for this virtual circuit. You specify the destination IP address as part of the mpls l2transport route command.
Peer LDP ID	The LDP IP address of the targeted peer.
Local VC Label	The VC label that a router signals to its peer router, which is used by the peer router during imposition. The local VC label is a disposition label. The local VC label determines the egress interface of an arriving packet from the MPLS backbone.
Remote VC Label	The disposition VC label of the remote peer router.
Tunnel Label	An IGP label that is used to route the packet over the MPLS backbone to the destination router with the egress interface.
Outgoing Interface	The egress interface of the virtual circuit.
Next Hop	The IP address of the next hop.
Local MTU	The maximum transmission unit that you specify for the client interface.
Remote MTU	The maximum transmission unit that you specify for the client interface of the remote router.
Imposition	The status of the line card, which can be one of the following: ■ LC programmed ■ LC not programmed
Current Imposition/Last Disposition Slot	The current imposition is the outgoing interface that is used for imposition. The last disposition slot is the interface where packets for this virtual circuit arrive.
Packet Totals (in/out)	The total number of packets that have been forwarded in each direction.
Byte Totals (in/out)	The total number of bytes that have been forwarded in each direction.

show mpls l2transport summary

This topic describes the syntax and output of the **show mpls l2transport summary** command.



To display summary information about virtual circuits that have been enabled to route AToM Layer 2 packets on a router, use the **show mpls l2transport summary** command in privileged EXEC mode.

The following sample output shows summary information about virtual circuits that have been enabled to transport Layer 2 packets:

```
Router# show mpls l2transport summary
```

```
Destination address: 172.16.24.12 Total number of VCs: 60
0 unknown, 58 up, 0 down, 2 admin down
5 active vc on MPLS interface PO4/0
```

show mpls l2transport summary Field Descriptions

Field	Description
Destination address	The IP address of the remote router to which the virtual circuit has been established.
Total number of VCs	The number of virtual circuits that have been established.
unknown	The number of virtual circuits that are in an unknown state.
up	The number of virtual circuits that are operational.
down	The number of virtual circuits that are not operational.
admin down	The number of virtual circuits that have been disabled.

debug mpls l2transport vlan control

This topic identifies the command syntax of the **debug mpls l2transport vlan control** command.

debug mpls l2transport vlan control

Cisco.com

```
router#  
debug mpls l2transport vlan control
```

- To enable debug messages about the control of traffic transported between Layer 2 VLAN and MPLS, use the **debug mpls l2transport vlan control EXEC** command.
- To disable the debug messages about transport control, use the **no** form of this command.

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—6-8

debug mpls l2transport vlan control

To enable debug messages about the control of traffic that is transported between Layer 2 VLAN and MPLS, use the **debug mpls l2transport vlan control EXEC** command. To disable the debug messages about transport control, use the **no** form of this command.

[no] debug mpls l2transport vlan control

Example

The following example shows how the **debug mpls l2transport vlan control** command enables debug messages about the control of traffic between Layer 2 VLANs and MPLS:

```
Router# debug mpls l2transport vlan control  
Ethernet VLAN transport over MPLS, Control interactions debugging is on  
Router# config t  
Enter configuration commands, one per line. End with CNTL/Z.  
Router(config)# int vlan2  
Router(config-if)# shut  
Router(config-if)#  
*May 25 12:18:33: ATOM_TRANS: atom_if_state_change from Vlan2  
*May 25 12:18:33: ATOM_TRANS: Withdrawing all EoMPLS vcs for lgroupid 8
```

```

*May 25 12:18:33: ATOM_TRANS: sending extended withdraw_bind for vcid
0
to 12.0.0.1, local groupid 8
*May 25 12:18:33: ATOM_TRANS: holding down local label 21
*May 25 12:18:33: ATOM_TRANS: clearing imp. rewrite for vcid = 2
Router(config-if)#
*May 25 12:18:35: ATOM_TRANS: freeing held label 21
Router(config-if)#
*May 25 12:18:35: %LINK-5-CHANGED: Interface Vlan2, changed state to
administratively down
*May 25 12:18:36: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Vlan2,
changed state to down
Router(config-if)# no shut
Router(config-if)#
*May 25 12:19:57: ATOM_TRANS: atom_if_state_change from Vlan2
*May 25 12:19:57: ATOM_TRANS: Withdrawing all EoMPLS vcs for lgroupid
8
*May 25 12:19:57: ATOM_TRANS: sending extended withdraw_bind for vcid
0
to 12.0.0.1, local groupid 8
Router(config-if)#
*May 25 12:19:59: %LINK-3-UPDOWN: Interface Vlan2, changed state to
up
Router(config-if)#
*May 25 12:19:59: ATOM_TRANS: sending bind for vcid 2 to 12.0.0.1,
local
groupid 8
*May 25 12:19:59: ATOM_TRANS: sending request_bind for vcid 2 to
12.0.0.1, local groupid 8
*May 25 12:19:59: ATOM_TRANS: Stale tfib event discarded (12.0.0.1)
*May 25 12:19:59: ATOM_TRANS: Stale tfib event discarded (12.0.0.1)
*May 25 12:19:59: ATOM_TRANS: Stale tfib event discarded (12.0.0.1)
*May 25 12:19:59: ATOM_TRANS: processing tfib event for 12.0.0.1
*May 25 12:19:59: ATOM_TRANS: received BIND from 12.0.0.1:0 remote
group id: 8 vc_id: 2 label: 21
*May 25 12:19:59: ATOM_TRANS: Done setting imp.rewrite for vcid = 2
parent dest 12.0.0.1
*May 25 12:20:00: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Vlan2,
changed state to up

```

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Virtual-circuit-based:**
 - show mpls l2transport vc
 - show mpls l2transport vc detail
 - show mpls l2transport summary
 - debug mpls l2transport vlan control
- **AAL5-based:**
 - Check www.cisco.com for details

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—6-9

References

For additional information, refer to this resource:

- Access the Cisco website at <http://www.cisco.com> for additional information about the unified VPN suite solution for service providers.

Next Steps

For the associated lab exercises, refer to the following sections of the course Lab Guide:

- Lab Exercise 6-1: Implementing a Port-to-Port FRoMPLS Configuration (Case Study)
- Lab Exercise 6-2: Implementing a DLCI-to-DLCI FRoMPLS Configuration (Case Study)

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) Which commands do you use to monitor an AToM session?

Q2) Which command do you use to display the status of all of the active virtual circuits on a router?

- A) **show mpls l2transport vc**
- B) **show mpls l2transport vc summary**
- C) **show mpls l2transport vc detail**
- D) **debug mpls l2transport vlan control**

Q3) Which command do you use to display summary information about virtual circuits that have been enabled to route AToM Layer 2 packets on a router?

- A) **show mpls l2transport vc**
- B) **show mpls l2transport vc summary**
- C) **show mpls l2transport vc detail**
- D) **debug mpls l2transport vlan control**

Q4) Which command do you use to display comprehensive information about the virtual circuits that have been assigned VC IDs?

- A) **show mpls l2transport vc**
- B) **show mpls l2transport vc summary**
- C) **show mpls l2transport vc detail**
- D) **debug mpls l2transport vlan control**

Q5) Which command do you use to display messages about the control of traffic that is transported between Layer 2 VLAN and MPLS?

- A) **show mpls l2transport vc**
- B) **show mpls l2transport vc summary**
- C) **show mpls l2transport vc detail**
- D) **debug mpls l2transport vlan control**

Lesson Answer Key

- Q1) **show mpls l2transport vc, show mpls l2transport vc summary, show mpls l2transport vc detail, and debug mpls l2transport vlan control**
Relates to: Monitoring AToM
- Q2) A
Relates to: show mpls l2transport vc
- Q3) B
Relates to: show mpls l2transport vc summary
- Q4) C
Relates to: show mpls l2transport vc detail
- Q5) D
Relates to: debug mpls l2transport vlan control

MPLS IPv6 Support

Overview

Is the need for IP version 6 (IPv6) support compelling? Yes, the soaring requirement for new IP addresses and enhanced security have made IPv6 a requirement in many implementations. Is IPv6 being accepted worldwide? Again, the answer is yes. In addition, 50 percent of the nodes in the IPv6 production network, the “6bone,” are based in the United States. Implementations involve dual stack networks, tunneling networks, and native IPv6 networks. To meet this market, Multiprotocol Label Switching (MPLS) has implemented support for IPv6 in several modes.

Objectives

Upon completing this module, you will be able to use a typical simple MPLS service provider network to implement support for IPv6. This includes being able to do the following:

- Identify the characteristics of the IPv6 protocol that are related to an MPLS implementation
- Configure IPv6 support in an MPLS solution
- Describe the function of the Cisco IOS commands that are used to verify IPv6 operation in an MPLS solution

Outline

The module contains these lessons:

- Review of IPv6
- Implementing IPv6 over MPLS
- Monitoring IPv6 over MPLS

Review of IPv6

Overview

IPv6 is becoming more and more accepted. To meet this need, MPLS has implemented support of IPv6 in several modes.

Relevance

To successfully implement MPLS support for IPv6, you must first understand the basic characteristics of IPv6.

Objectives

This lesson provides the learner with a brief review of IPv6 operation. Upon completing this lesson, you will be able to do the following:

- Identify the benefits of integrating IPv6 into a network
- Compare IPv4 and IPv6 solutions
- Identify the need for a larger IP address space
- Identify the characteristics of an IPv6 address
- Identify how an IPv6 address is represented
- Identify the hierarchical addressing and aggregation scheme that is used by IPv6

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- IP routing technology

Outline

This lesson includes these topics:

- Overview
- The Benefits of Integrating IPv6
- IPv6 Technology Scope
- IPv6 Address Representation
- Hierarchical Addressing and Aggregation
- Lesson Summary
- Lesson Review

The Benefits of Integrating IPv6

This topic describes the need for integrating IPv6 into a network.

A Need for IPv6?

Cisco.com

- **IPv4 32-bit address = 4 billion hosts**
 - **Approximately 40 percent of the IPv4 address space is still unused.**
- **But IP is everywhere**
 - **Data, voice, audio, and video integration is a reality.**
 - **Regional registries apply a strict allocation control.**
- **So, the one compelling reason to move to IPv6: more IP addresses!**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-3

What is the need for IPv6? Most presentations on this topic list reasons that range from address space to security, but what is the real reason?

Currently about 40 percent of IP version 4 (IPv4) addresses are unused. Does that mean that we are safe for 15 more years? Actually not, because applications like voice, data, and Voice over IP (VoIP) are a reality.

In the short term, solutions like classless interdomain routing (CIDR) and Network Address Translation (NAT) have been developed to address this issue. However, these solutions have presented their own issues. For example, NAT broke the end-to-end IPv4 model.

In addition, IP is leaving the confined sphere of the PC. Now IP is everywhere. There are theaters with IP appliances, IP faxes, and IP industrial sensors.

Currently, IPv4 is used on mobile phones. The average size of a mobile phone service provider is about 20 million users. To support this implementation, service providers require two levels of NAT. A phone call from one mobile IP phone to another mobile IP phone requires both levels, which is an unmanageable situation. Service providers cannot make progress without IPv6 in this area.

There is one truly compelling reason to move to IPv6. Users simply need more IP addresses than IPv4 can provide.

IPv6 Technology Scope

This topic compares IPv4 and IPv6 solutions.

IPv6 Technology Scope		
Cisco.com		
IP Service	IPv4 Solution	IPv6 Solution
Addressing range	32-bit, Network Address Translation	128-bit, Multiple Scopes
Autoconfiguration	DHCP	Serverless, Reconfiguration, DHCP
Security	IPSec	IPSec Mandated, Works End-to-End
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality of Service	Differentiated Services, Integrated Services	Differentiated Services, Integrated Services
IP Multicast	IGMP/PIM/Multicast BGP	MLD/PIM/Multicast BGP, Scope Identifier

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—7.4

Address Space

What does IPv6 bring to the IP solution? First of all, more IP addresses.

Autoconfiguration

While autoconfiguration is available with IPv4, it is clearly easier with IPv6. The address autoconfiguration feature is built into the IPv6 protocol to facilitate intranet-wide address management, which enables a large number of IP hosts to easily discover the network and get a new and globally unique IPv6 address that is associated with their location. The autoconfiguration feature enables “plug-and-play” Internet deployment of new consumer devices, such as cell phones, wireless devices, home appliances, and so on. As a result, network devices can connect to the network without manual configuration and without any servers, such as DHCP servers.

A router on the local link will send network-type information, such as the prefix of the local link and the default route, in its router advertisements. The router provides this information to all the nodes on the local link. As a result, a host can autoconfigure itself by appending its 48-bit link-layer address (MAC address) in an extended universal identifier 64-bit (EUI-64) format to the 64 bits of the local link prefix that is being advertised by the router.

Security

There is IP Security (IPSec) in IPv4. But with IPv6 the IPSec support is embedded in the initial specification of the protocol so you do not need to support another option.

Mobility

Clearly, you have mobility with IPv4, but with IPv6 it is embedded in the protocol headers. With the built-in mobility of IPv6, any IPv6 node can use mobility as needed. Mobility is becoming an important and critical feature in networks. Mobile IP is an Internet Engineering Task Force (IETF) standard that allows mobile devices to move around without breaking their existing connections. In IPv4, the mobility function must be added as a new feature. Mobility support in IPv6 is discussed in the latest version of Internet Draft draft-ietf-mobileip-ipv6-17.txt.

IPv6 packets that are addressed to the home address of a mobile node are transparently routed to the care-of address through the caching of the binding of the home address with the care-of address. This binding allows any packets that are destined for the mobile node to be directed to it at the care-of address. Mobile IPv6 defines four new IPv6 destination options: binding update option, binding acknowledgment option, binding request option, and home address option.

The routing headers in IPv6 make Mobile IPv6 much more efficient for end devices than Mobile IPv4. The use of the routing header for Mobile IP, rather than IP encapsulation, enables Mobile IP to avoid triangle routing, making it much more efficient in IPv6 than in IPv4.

Quality of Service

QoS in IPv6 is handled in the same way that it is currently handled in IPv4. Support for class of service (CoS) is available through the Traffic Class field compliant with the IETF Differentiated Services (DiffServ) model.

However, the IPv6 header has a new field named “Flow Label,” which can contain a label that identifies a specific flow, such as a video stream or videoconference. The source node generates this flow label. Having a flow label enables QoS devices in the path to take appropriate actions that are based on this label. However, the flow label itself is not a feature of QoS.

Multicast Scope Address

One of the salient features of IPv6 is that it does not use broadcasts at all. Functions that are supported by IPv4 broadcasts, such as router discovery and router solicitation requests, are handled by IPv6 multicast. Multicast allows IP packets such as a video stream to be sent to multiple destinations at the same time, which saves network bandwidth. Multicast improves the efficiency of a network by limiting the broadcast requests to a smaller number of interested nodes only. IPv6 also uses specific multicast group addresses for its various functions. Thus, IPv6 multicast prevents the problems that are caused by broadcast storms in IPv4 networks.

IPv4 networks use administratively scoped IP multicast addresses as described in RFC 2365, *Administratively Scoped IP Multicast*, to allow packets to be addressed to a specific range of multicast addresses (for example, 239.0.0.0 to 239.255.255.255). By specifying a multicast scope, the packets are prevented from crossing the configured administrative boundaries. IPv4 uses one broadcast address for a particular scoped zone or IP multicast boundary, and the broadcasts are received by all hosts in this scoped zone. IPv6 uses a 4-bit scope ID to specify address ranges that are reserved for multicast addresses within each scoped zone. Thus, only those hosts in a specified scope address range that are configured to listen to a specific multicast address receive the multicast. However, a host can be a member of several workgroups and can listen to several multicast addresses at the same time.

IPv6 provides a larger range of multicast addresses compared to IPv4. So, allocation of addresses for multicast groups will not be limited for the foreseeable future.

IPv4 and IPv6 Header Comparison

Cisco.com

IPv4 Header

Version	IHL	Type of Service	Total Length			
Identification			Flags	Fragment Offset		
Time to Live	Protocol		Header Checksum			
Source Address						
Destination Address						
Options				Padding		

IPv6 Header

Version	Traffic Class	Flow Label		
Payload Length			Next Header	Hop Limit
Source Address				
Destination Address				

- Legend**
- Field name kept from IPv4 to IPv6
 - Fields not kept in IPv6
 - Name and position changed in IPv6
 - New field in IPv6

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-5

When you quickly compare IPv4 and IPv6 headers, you first realize that the IPv4 header is smaller than the IPv6 header. The larger size of the IPv6 header results from the fact that the addresses are very long. However, there are actually fewer fields in the IPv6 header and they are optimized for compression with a 64-bit alignment.

The IPv6 header length is fixed at 40 bytes to provide for easier processing and the elimination of the need for a header length field.

Besides the header length field, IPv6 has eliminated a number of the IPv4 fields.


Because of differing link maximum transmission units (MTUs), IPv4 supports a complex fragmentation process for routers. With IPv6, the router does not need to do fragmentation. Link MTUs are calculated from the end-to-end path. The MUT for the entire path is set to the lowest MUT encountered over the path. After MTU path discovery, the end host takes care of the fragmentation.

IPv4 uses a header checksum to ensure the integrity of the IP header. IPv6 has eliminated this third level of checking. Instead it depends on the error checking that is provided by other network layers. For example, with a TCP communication, Layer 2 and Layer 4 checksums exist. So there is no need for a third checksum at Layer 3.

IPv4 supports several options. But, in reality, most packets do not carry options, which results in less CPU processing. IPv6 has kept the concept of options, but they exist as specific headers that are daisy-chained to the first header. Thus, all options are now transferred in daisy-chained headers.

IPv6 Address Representation

This topic identifies how an IPv6 address is represented.



IPv6 uses 16-bit hexadecimal number fields separated by colons (:) to represent the 128-bit addressing format, which makes address representation less cumbersome and error-prone than it was with IPv4. The hexadecimal numbers are not case-sensitive. Here is an example of a valid IPv6 address: 2031:0000:130F:0000:0000:09C0:876A:130B.

Additionally, in order to shorten the IPv6 address and make the address easier to represent, IPv6 uses the following conventions:

- Leading zeros in the address field are optional and can be compressed. For example, the following hexadecimal numbers can be represented as shown in a compressed format:
 - **Example 1:** 2031:0000:130F:0000:0000:09C0:876A:130B = 2031:0:130F:0:0:9C0:876A:130B (compressed form)
 - **Example 2:** 0000 = 0 (compressed form)
- A pair of colons (::) represents successive fields of 0. However, the pair of colons is allowed only once in a valid IPv6 address.
 - **Example 1:** 2031:0:130F:0:0:9C0:876A:130B = 2031:0:130F::9C0:876A:130B (compressed form)
 - **Example 2:** FF01:0:0:0:0:0:1 = FF01::1 (compressed form)

An address parser could easily identify the number of missing zeros in an IPv6 address by pulling the two parts of the address apart and filling with 0s until the 128-bit address is complete. However, if two ::s are placed in the same address, then there is no way to identify the size of each block of zeros. The use of the :: makes many IPv6 addresses very small.

The Internet Assigned Numbers Authority (IANA) allocates 2001::/16 to registries from the full address space. From IANA, each registry gets a /23 prefix within the 2001::/16 space, as follows:

- 2001:0200::/23 and 2001:0C00::/23 allocated to Asia Pacific Network Information Center (APNIC) for use in Asia.
- 2001:0400::/23 allocated to American Registry for Internet Numbers (ARIN) for use in the Americas.
- 2001:0600::/23 and 2001:0800::/23 allocated to Réseaux IP Européens - Network Coordination Center (RIPE NCC) for use in Europe and the Middle East.

The registries then allocate an initial /32 prefix to the IPv6 ISPs, and the ISPs allocate a /48 prefix (out of the /32) to each customer or site. The /48 prefix of the site can be further allocated to each LAN using a /64 prefix. This allocates a maximum of 64 bits for ID host addressing on each LAN. Each site could subnet the site into a maximum of 65,535 LANs.

A site should make an address plan prior to beginning allocation of its /48 space. In order to receive a /32 prefix address block from a registry, an ISP must have an exterior routing protocol that is peering with at least three other ISPs and either have at least 40 customers or demonstrate a clear intent to provide an IPv6 service within 12 months.

For the latest information about allocation of IPv6 address space to the registries by IANA, refer to the IANA assignments URL at <http://www.iana.org/assignments/ipv6-tla-assignments>.

IPv6 Address Prefix

The IPv6 prefix is the leftmost bits of the address that have a fixed value. This prefix represents the network identifier. The IPv6 prefix is represented using the IPv6-prefix/prefix-length format just like an IPv4 address that is represented in the CIDR notation. The IPv6-prefix variable must conform to RFC 2373.

The /prefix-length variable is a decimal value that indicates the number of high-order contiguous bits of the address that form the prefix, which is the network portion of the address. For example, 1080:6809:8086:6502::/64 is an acceptable IPv6 prefix. If the address ends in a double colon, the trailing double colon can be omitted. So, the same address can be written as 1080:6809:8086:6502/64. In either case, the prefix length is written as a decimal number 64 and represents the leftmost bits of the IPv6 address.

Hierarchical Addressing and Aggregation

This topic describes the hierarchical addressing and aggregation scheme that is used by IPv6.

Hierarchical Addressing

Cisco.com

- **Unicast** - An address for a single interface
 - Global (2000::/3 - E000::/3)
 - Site-local (FEC0::/10)
 - Link-local (FE80::/10)
- **Anycast** - An address for a set of interfaces that typically belong to different nodes
- **Multicast** - An address for a set of interfaces (in a given scope) that typically belong to different nodes

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—7-7

There is a major difference in the IP address requirements between an IPv4 node and an IPv6 node. An IPv4 node typically uses one IP address, while an IPv6 node requires more than one IP address.

The three major types of IPv6 addresses are as follows:

- **Unicast:** An address for a single interface. A packet that is sent to a unicast address is delivered to the interface that is identified by that address.
- **Anycast:** An address for a set of interfaces that typically belong to different nodes. A packet that is sent to an anycast address is delivered to the closest interface (as defined by the routing protocols in use) that is identified by that address.
- **Multicast:** An address for a set of interfaces (in a given scope) that typically belong to different nodes. A packet that is sent to a multicast address is delivered to all interfaces that are identified by that (in a given scope).

IPv6 Address Assignment

An IPv6 address is assigned to a single interface, not a node. However, a single interface could be assigned multiple IPv6 addresses. Hence, it is easy to identify a node by any of its unicast addresses. The following are notable exceptions to this general rules:

- Multiple interfaces can have a single unicast address assigned to them when they are used for load sharing over multiple physical interfaces. The same is true when multiple physical interfaces are treated as a single interface at the Internet layer.

- Routers using unnumbered interfaces on point-to-point links are not assigned IPv6 addresses, because the interfaces do not function as a source or destination for IP datagrams.

What Is an IPv6 Global Unicast Address?

The IPv6 global unicast address is the equivalent of the IPv4 global unicast address. A global unicast address is an IPv6 address from the global unicast prefix. The structure of global unicast addresses enables aggregation of routing prefixes that limits the number of routing table entries in the global routing table. Global unicast addresses used on links are aggregated upward through organizations and eventually to the Internet service providers (ISPs).

Global unicast addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Except for addresses that start with binary 000, all global unicast addresses have a 64-bit interface ID. The current global unicast address allocation uses the range of addresses that start with the binary value 001 (2000::/3).

2000::/3 is the global unicast address range and uses one-eighth of the total IPv6 address space. It is the largest block of assigned block addresses.

What Is the Structure of a Global Unicast Address?

A fixed prefix of 2000::/3 (001) indicates a global IPv6 address. Addresses with a prefix of 2000::/3 (001) through E000::/3 (111), excluding the FF00::/8 (1111 1111) multicast addresses, are required to have 64-bit interface identifiers in the EUI-64 format. The IANA allocates the IPv6 address space in the ranges of 2001::/16 to the registries.

The global unicast address typically consists of a 48-bit global routing prefix and a 16-bit subnet ID. In the IPv6 aggregatable global unicast address format document (RFC 2374), the global routing prefix included two other hierarchically structured fields called Top-Level Aggregator and Next-Level Aggregator. Because these fields were policy-based, the IETF decided to remove them from the RFCs. However, some existing IPv6 networks deployed in the early days might still be using networks that are based on the older architecture.

A 16-bit subnet field called the Subnet ID can be used by individual organizations to create their own local addressing hierarchy and to identify subnets. This field allows an organization to use up to 65,535 individual subnets.

What Is an IPv6 Site-Local Unicast Address?

Site-local unicast addresses are similar to the private addresses such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16 used in IPv4 networks. Private addresses can be used to restrict communication to a specific domain, or to assign addresses in a site that is not connected to the global Internet, without requiring a globally unique prefix. IPv6 routers must not advertise routes or forward packets that have site-local source or destination addresses outside the site boundary. If the site requires global connectivity in the future, a global unicast prefix must be assigned to that site. The site-local addressing plan that has been initially defined for site-local addressing can be directly applied using the global unicast prefix.

The site-local unicast address is an IPv6 unicast address that uses the prefix range FEC0::/10 (1111 1110 11) and concatenates the subnet identifier (the 16-bit Subnet ID field) with the interface ID in the EUI-64 format. The site-local unicast address range uses 1/1024 of the total address space.

What Is an IPv6 Link-Local Unicast Address?

A link-local unicast address is an IPv6 unicast address that is automatically configured on an IPv6 node interface by using the link-local prefix FE80::/10 (1111 1110 11) and the interface ID in the EUI-64 format. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. They are typically used to connect devices on the same local link network without the need for global addresses. Hence, link-local addresses are useful only in the context of the local link network.

Nodes on a local link can use link-local addresses to communicate with each other without the need for a router. IPv6 nodes do not need site-local or globally unique addresses to communicate. IPv6 routers must not forward to other links packets that have link-local source or destination addresses. FE80::/10 is the link-local unicast address range and uses 1/1024 of the IPv6 address space.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Increased address space.**
- **Three address types:**
 - **Unicast**
 - **Multicast**
 - **Anycast**
- **Addresses are ISP-dependent.**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—7-8

References

For additional information, refer to these resources:

- Access the Cisco website at <http://www.cisco.com> for additional information about the unified VPN suite solution for service providers.
- ABCs of Cisco IOS Software: The ABCs of IP Version 6 at http://www.cisco.com/en/US/partner/products/sw/iosswrel/ios_abcs_ios_the_abcs_ip_version_6_listing.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

Q1) What is the main reason for implementing IPv6 into a network?

- A) to increase network security
- B) to implement end-to-end tunnels
- C) to provide support for mobile video applications
- D) to increase the number of addresses that are available

Q2) How do you assign a host address when using the autoconfiguration feature of IPv6?

Q3) Which of the following is a valid IPv6 address?

- A) 192.168..10.10.5.6
- B) 0FF01:0:0:0:0:1
- C) 2031:0:130F:9C0:876A::130B
- D) 2031:0::130F:0:0:9C0::876A:130B

Q4) Which three of the following are valid IPv6 address types? (Choose three.)

- A) unicast
- B) anycast
- C) autocast
- D) multicast
- E) broadcast
- F) localcast

Lesson Answer Key

Q1) D

Relates to: The Benefits of Integrating IPv6

Q2) The autoconfiguration feature enables “plug-and-play” Internet deployment of new consumer devices. A router on the local link sends network-type information, such as the prefix of the local link and the default route, in its router advertisements. The router provides this information to all the nodes on the local link. As a result, a host can autoconfigure itself by appending its 48-bit link-layer address (MAC address) in an EUI-64 format to the 64 bits of the local link prefix that is being advertised by the router.

Relates to: IPv6 Technology Scope

Q3) C

Relates to: IPv6 Address Representation

Q4) A, B, D

Relates to: Hierarchical Addressing and Aggregation

Implementing IPv6 over MPLS

Overview

MPLS is deployed by many service providers in their IPv4 networks. Service providers want to introduce IPv6 services to their customers, but changes to their existing IPv4 infrastructure can be expensive and the cost benefit for a small amount of IPv6 traffic does not make economic sense. Several integration scenarios have therefore been developed to leverage an existing IPv4 MPLS infrastructure and to add IPv6 services without requiring any changes to the network backbone.

Relevance

To successfully implement MPLS support for IPv6, you must understand the different implementation models and the commands that are used to implement these models.

Objectives

After the brief review of IPv6 operations in the previous lesson, this lesson provides the learner with knowledge needed to implement IPv6 support in an MPLS-based network.. Upon completing this lesson, you will be able to do the following:

- Identify the benefits of deploying IPv6 over MPLS backbones
- Identify the characteristics of an IPv6 model that uses tunnels on the customer edge routers
- Deploy IPv6 using tunnels on the customer edge routers
- Identify the characteristics of a model that deploys IPv6 over a circuit transport over MPLS
- Deploy IPv6 over a circuit transport over MPLS
- Identify the characteristics of an IPv6 model (Cisco 6PE) that uses the provider edge routers
- Deploy IPv6 on the provider edge routers

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- IP routing technology

Outline

This lesson includes these topics:

- Overview
- Benefits of Deploying IPv6 over MPLS Backbones
- IPv6 Using Tunnels on the Customer Edge Routers
- Deploying IPv6 Using Tunnels on the Customer Edge Routers
- IPv6 over a Circuit Transport over MPLS
- Deploying IPv6 over a Circuit Transport over MPLS
- IPv6 on the Provider Edge Routers (Cisco 6PE)
- Deploying Cisco 6PE
- Lesson Summary
- Lesson Review

Benefits of Deploying IPv6 over MPLS Backbones

This topic describes the need for integrating IPv6 into a network.

Benefits of Deploying IPv6 Over MPLS Backbones

Cisco.com

- **Enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 core network.**
- **Requires only a few backbone infrastructure upgrades.**
- **No reconfiguration of core routers.**
- **Allows IPv6 networks to be combined into IPv4 VPNs or extranets over an infrastructure supporting IPv4 VPNs and MPLS-TE.**

© 2004 Cisco Systems, Inc. All rights reserved.


MPLST v2.0—7-3

The deployment of IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires only a few backbone infrastructure upgrades and no reconfiguration of core routers, because forwarding is based on labels rather than the IP header itself. This method provides a very cost-effective strategy for the deployment of IPv6.

Additionally, the inherent Virtual Private Network (VPN) and MPLS Traffic Engineering (MPLS-TE) services that are available within an MPLS environment allow IPv6 networks to be combined into IPv4 VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS-TE. IPv6 VPNs are not supported.

IPv6 Using Tunnels on the Customer Edge Routers

This topic describes the characteristics of an IPv6 model that uses tunnels on the customer edge (CE) routers.



Using tunnels on the CE routers is the simplest way of deploying IPv6 over MPLS networks with no impact on the operation or infrastructure of MPLS. Also, no configuration changes to the core or provider edge (PE) routers are needed. Communication between the remote IPv6 domains uses standard tunneling mechanisms and requires the CE routers to be configured to run dual IPv4 and IPv6 protocol stacks.

Deploying IPv6 Using Tunnels on the Customer Edge Routers

This topic describes the process that is used to implement IPv6 using tunnels on the CE routers.

Deploying IPv6 Using Tunnels on the Customer Edge Routers

Cisco.com

- **Enable a dual IPv4 and IPv6 stack on the CE routers**
- **Enable standard IPv4 MPLS connectivity between the CE and PE routers**
- **Create IPv4 tunnels between CE routers**
- **Refer to Cisco documentation for further details**

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-5

The first step is to enable a dual IPv4 and IPv6 stack on the CE routers. The IPv6 portion of the stack will support the customer IPv6 network. The IPv4 portion of the stack will support the CE-to-PE connection.

Next, enable a standard IPv4 MPLS session between the CE and PE routers.


Finally, create IPv4 tunnels between CE routers to carry the IPv6 traffic.

Refer to Cisco Systems documentation for further details:

- http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d65f5.html
- http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d6604.html

IPv6 over a Circuit Transport over MPLS

This topic describes the characteristics of a model that deploys IPv6 over a circuit transport over MPLS.



Using any circuit transport for deploying IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS, and requires no configuration changes to the core or PE routers.

The remote IPv6 domains run native IPv6 protocols over a dedicated link, where the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled using the Any Transport over MPLS (AToM) or Ethernet over MPLS (EoMPLS) feature with the routers connected through an ATM OC-3 or Ethernet interface, respectively.

Deploying IPv6 over a Circuit Transport over MPLS

This topic describes the process that is used to implement IPv6 over a circuit transport over MPLS.

Deploying IPv6 Over a Circuit Transport Over MPLS

Cisco.com

- **Enable an IPv6 stack on the CE router**
- **Establish Layer 2 connectivity between the CE and PE router**
- **Create an AToM circuit (tunnel) between the PE routers**
- **Refer to AToM module for further details on AToM**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—7-7


The first task is to enable an IPv6 stack on the CE router. This stack will support the customer IPv6 network. Connectivity between the CE and PE router is provided by establishing Layer 2 connectivity between the PE routers using AToM.

Connectivity to the PE router and across the MPLS network is provided by establishing a Layer 2 session. Connectivity between the PE routers is provided by using AToM circuits (tunnels) between the PE routers.

Refer to the module “Any Transport over MPLS” for further details on AToM.

IPv6 on the Provider Edge Routers (Cisco 6PE)

This topic describes the characteristics of an IPv6 model that uses the PE routers.



The Cisco implementation of IPv6 on the PE routers over MPLS is referred to as Cisco 6PE and enables IPv6 sites to communicate with each other over an MPLS IPv4 core network using MPLS label switched paths (LSPs). MPLS decoupling of the control plane and data plane provide an interesting alternative to the integration and coexistence of IPv4, IPv6, and ATM over a single infrastructure.


The Cisco 6PE feature is particularly applicable to service providers that already run an MPLS network or plan to do it. One of the Cisco 6PE advantages is that there is no need to upgrade the hardware, software, or configuration of the core network. Thus, it eliminates the impact on operation and revenue that is generated by existing IPv4 traffic. MPLS has been chosen by many service providers as a vehicle to deliver services to customers. As a multiservice infrastructure technology, MPLS is able to provide Layer 3 VPN, quality of service (QoS), traffic engineering (TE), Fast Reroute (FRR), and integration of ATM and IP switching. MPLS is thus a natural choice to ease IPv6 introduction in existing production networks.

In the figure, all Cisco 6PE and core routers—provider routers (P routers) in this case—within the MPLS domain share a common IPv4 Interior Gateway Protocol (IGP) such as Open Shortest Path First (OSPF) or Integrated Intermediate System-to-Intermediate System (IS-IS).

The Cisco 6PE routers are configured as dual stack routers that are able to route both IPv4 and IPv6 traffic. Each Cisco 6PE router is configured to run Label Distribution Protocol (LDP) or Resource Reservation Protocol (RSVP) (if TE is configured) to bind the IPv4 labels.

Cisco 6PE MP-BGP Sessions

Cisco.com



The Cisco 6PE feature relies on Multiprotocol Border Gateway Protocol (MP-BGP) extensions in the IPv4 network configuration on the PE router to exchange IPv6 reachability information in addition to an MPLS label for each IPv6 address prefix to be advertised.

Edge routers are configured to be dual stack running both IPv4 and IPv6, and use the IPv4-mapped IPv6 address for IPv6 prefix reachability exchange.

Cisco.com




MPLST v2.0—7-10

Within the MPLS network, IPv6 traffic is forwarded using label switching, which makes the IPv6 traffic transparent to the core of the MPLS network. No IPv6 over IPv4 tunnels or Layer 2 encapsulation methods are required.

Deploying Cisco 6PE

This topic describes the process that is used to implement the Cisco 6PE model.



Each Cisco 6PE router—6PE1 and 6PE2 in the figure—is assumed to be running IPv4 routing and Cisco Express Forwarding (CEF).


In this example, the two customer sites have been assigned an IPv6 address of 2001:0db8:ffff::/48 and 2001:0db8:dddd::/48. The CE routers are connected to the Cisco 6PE routers over S0/0.

The MPLS core is using IS-IS with an address range of 192.168.99.0/16. The Cisco 6PE routers are connected to the P routers over E0/0.

The P routers in the core of the network is assumed to be running MPLS, a label distribution protocol, an IPv4 IGP, and CEF or distributed Cisco Express Forwarding (dCEF), and does not require any new configuration to enable the Cisco 6PE feature.

IP Addressing

Cisco.com



The Cisco 6PE router is prepared by enabling IPv4 and IPv6 CEF, enabling IPv6 routing, and specifying the interface from which locally generated packets take their source IPv6 address.


Router 6PE1 exchanges IPv6 routing information with router 6PE2 using Internal Border Gateway Protocol (IBGP) established over an IPv4 connection so that all the **neighbor** commands use the IPv4 address of the 6PE2 router. All the Border Gateway Protocol (BGP) peers are within autonomous system (AS) 65000, so synchronization with IGP is turned off for IPv4. In IPv6 address family configuration mode, synchronization is disabled by default.

IPv6 and Cisco Express Forwarding for IPv6 (CEFv6) are enabled, the 6PE2 neighbor is activated, and aggregate label binding and advertisement is enabled for IPv6 prefixes using the **neighbor send-label** command. Connected and static IPv6 routes are redistributed using BGP. If IPv6 packets are generated in the local router, the IPv6 address for MPLS processing will be the address of loopback interface 0.

In the example here, serial interface 0/0 connects to the customer, and the IPv6 prefix that is delegated to the customer is 2001:0DB8:ffff::/48, which is determined from the service provider IPv6 prefix. A static route is configured to route IPv6 packets between the Cisco 6PE route and the CE router.

Cisco 6PE Routing

Cisco.com




The 6PE router (6PE1) is configured for both IPv4 and IPv6 traffic. Ethernet interface 0/0 is configured with an IPv4 address and is connected to a router in the core of the network. Integrated IS-IS and TDP configurations on this router are similar to a standard MPLS implementation.

A default static route is installed using the IPv6 address of serial interface 0/0 on the CE1 router.

CE-to-6PE Static Routing

Cisco.com




The customer edge router (CE1) is configured to forward its IPv6 traffic to the 6PE1 router.

In the example, serial interface 0/0 of CE1 is connected to the service provider and is assigned an IPv6 address. IPv6 is enabled and a default static route is installed using the IPv6 address of serial interface 0/0 on the 6PE1 router.

CE-to-6PE BGP Routing

Cisco.com




© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-15

CE1 is defined as a BGP neighbor and then activated under the IPv6 address family.

CE-to-6PE BGP Routing (Cont.)

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-16

6PE1 is defined as a BGP neighbor and then activated under the IPv6 address family.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

IPv6 can be transported over an MPLS core using:

- **IPv4 tunnels established between customer edge routers**
- **Layer 2 tunnels created between PE routers using the AToM feature**
- **Cisco 6PE support**
 - **Customer runs IPv6 stack**
 - **Cisco 6PE runs dual IPv4 and IPv6 stack**
 - **IPv6 labels are distributed over IBGP session**

© 2004 Cisco Systems, Inc. All rights reserved. MPLST v2.0—7-17

References

For additional information, refer to these resources:

- **ABCs of Cisco IOS Software: The ABCs of IP Version 6**
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ios_abcs_ios_the_abcs_ip_version_6_listing.html
- **IPv6: Providing IPv6 Services over an IPv4 Backbone Using Tunnels**
http://www.cisco.com/en/US/partner/tech/tk872/tk373/technologies_design_guide09186a00800a3e91.shtml
- **6bone Connection Using 6to4 Tunnels for IPv6**
http://www.cisco.com/en/US/partner/tech/tk872/tk373/technologies_configuration_example09186a00801f3b4f.shtml
- **Implementing Tunneling for IPv6**
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d6604.html

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) Which of the following is true concerning the implementation of IPv6 support in an MPLS network?
- A) End-to-end tunnels must be configured.
 - B) A dual IPV4 and IPv6 stack must be enabled on all routers.
 - C) All routers must be running Cisco IOS version 12.3 or later.
 - D) The edge router may need to be upgraded, but no changes are required on the core routers.
- Q2) Which of the following is true concerning the implementation of IPv6 using tunnels on the customer edge routers?
- A) Tunneling support must be enabled on the core routers.
 - B) A dual IPV4 and IPv6 stack must be enabled on the CE routers.
 - C) The CE routers must be running Cisco IOS version 12.3 or later.
 - D) End-to-end AToM circuit tunnels must be created between CE routers.
- Q3) Which of the following is true concerning the implementation of IPv6 over a circuit transport over MPLS?
- A) AToM support must be enabled on the core routers.
 - B) A dual IPV4 and IPv6 stack must be enabled on the PE routers.
 - C) The CE routers must be running Cisco IOS version 12.3 or later.
 - D) End-to-end AToM Layer 2 circuits must be created between CE routers.
- Q4) Which of the following is true concerning the implementation of Cisco 6PE?
- A) AToM support must be enabled on the core routers.
 - B) A dual IPV4 and IPv6 stack must be enabled on the PE routers.
 - C) The P routers must be running Cisco IOS version 12.3 or later.
 - D) End-to-end Layer 3 circuits must be created between PE routers.

Lesson Answer Key

Q1) D

Relates to: Benefits of Deploying IPv6 over MPLS Backbones

Q2) B

Relates to: Deploying IPv6 Using Tunnels on the Customer Edge Routers

Q3) D

Relates to: Deploying IPv6 over a Circuit Transport over MPLS

Q4) B

Relates to: Deploying Cisco 6PE

Monitoring IPv6 over MPLS

Overview

The IPv6 models for using tunnels on CE routers and using IPv6 over a circuit transport over MPLS are implemented using existing IPv6 and AToM technologies. Therefore, these models use existing **show** and **debug** commands.

The Cisco 6PE model provides additional support over existing MPLS technology. Many of the existing MPLS **show** and **debug** commands are available to monitor Cisco 6PE operations [for information related to these operations, refer to the *Implementing Cisco MPLS* (MPLS) course]. This lesson will also discuss an additional set of 6PE-specific commands that the Cisco 6PE model provides.

Relevance

With the introduction of IPv6 support in an MPLS network, it is important to be able to monitor network operations to ensure successful deployment.

Objectives

This lesson provides the learner with a brief review of IPv6 operations. Upon completing this lesson, you will be able to do the following:

- Identify the commands that are available for monitoring IPv6 support
- Identify the syntax of the **show bgp ipv6** command and the important fields in its output
- Identify the syntax of the **show bgp ipv6 neighbors** command and the important fields in its output
- Identify the syntax of the **show mpls forwarding-table** command and the important fields in its output
- Identify the syntax of the **show bgp ipv6 labels** command and the important fields in its output
- Identify the syntax of the **show ipv6 cef** command and the important fields in its output
- Identify the syntax of the **show ipv6 route** command and the important fields in its output

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- IP routing technology

Outline

This lesson includes these topics:

- Overview
- Monitoring IPv6 Support
- show bgp ipv6 Command
- show bgp ipv6 neighbors Command
- show mpls forwarding-table Command
- show bgp ipv6 labels Command
- show ipv6 cef Command
- show ipv6 route Command
- Lesson Summary
- Lesson Review

Monitoring IPv6 Support

This topic identifies the commands that are available to monitor IPv6 support.

Monitoring IPv6 Support

Cisco.com

- **IPv6 Using Tunnels on the Customer Edge Routers**
 - Used existing IPv4 and IPv6 show and debug commands
- **IPv6 over a Circuit Transport over MPLS**
 - Used existing AToM show and debug commands
- **IPv6 on the Provider Edge Routers (6PE)**
 - show bgp ipv6
 - show bgp ipv6 neighbors
 - show bgp ipv6 label
 - show bgp ipv6 labels
 - show ipv6 cef
 - show ipv6 route

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—7-3

Support for IPv6 was added over existing MPLS technology.

The IPv6 model of using tunnels on the CE routers can be monitored using existing IPv4 and IPv6 commands that are related to tunneling IPv6 over an IPv4 network.

The model of using IPv6 over a circuit transport over MPLS is implemented using existing IPv6 and AToM technologies. Therefore, this model uses existing **show** and **debug** commands.

The Cisco 6PE model provides additional support over existing MPLS technology. Many of the existing MPLS **show** and **debug** commands are available to monitor Cisco 6PE operations. An additional set of 6PE-specific commands has been provided.

When Cisco 6PE is running, you can monitor the following components:

- **MP-BGP:** **show bgp ipv6** and **show bgp ipv6 neighbors**
- **MPLS:** **show mpls forwarding-table** and **show mpls ipv6 label**
- **CEFv6:** **show ipv6 cef**
- **IPv6 routing table:** **show ipv6 route**

show bgp ipv6 Command

This topic identifies the syntax of the **show bgp ipv6** command and the important fields in its output.

show bgp ipv6 Command

Cisco.com

```
Router> show bgp ipv6 2001:0DB8:DDDD::/48

BGP routing table entry for 2001:0DB8:DDDD::/48, version 15
Paths: (1 available, best #1, table Global-IPv6-Table)
  Not advertised to any peer
    Local
      ::FFFF:192.168.99.70 (metric 20) from 192.168.99.70 (192.168.99.70)
        Origin IGP, localpref 100, valid, internal, best
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-4

In this example, output information about an IPv6 route is displayed using the **show bgp ipv6** user EXEC command with an IPv6 prefix.

show bgp ipv6 neighbors Command

This topic identifies the syntax of the **show bgp ipv6 neighbors** command and the important fields in its output.

show bgp ipv6 neighbors Command

Cisco.com

```
Router> show bgp ipv6 neighbors 192.168.99.70
```

```
BGP neighbor is 192.168.99.70, remote AS 65000, internal link
BGP version 4, remote router ID 192.168.99.70
BGP state = Established, up for 00:05:17
Last read 00:00:09, hold time is 0, keepalive interval is 60 seconds
Neighbor capabilities:
  Route refresh: advertised and received
  Address family IPv6 Unicast: advertised and received
  ipv6 MPLS Label capability: advertised and received
Received 54 messages, 0 notifications, 0 in queue
Sent 55 messages, 1 notifications, 0 in queue
Default minimum time between advertisement runs is 5 seconds
```

```
For address family: IPv6 Unicast
BGP table version 21880, neighbor version 21880
Index 1, Offset 0, Mask 0x2
Route refresh request: received 0, sent 0
77 accepted prefixes consume 4928 bytes
Prefix advertised 4303, suppressed 0, withdrawn 1328
Number of NLRI in the update sent: max 1, min 0
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLS v2.0—7-5

In this example, output information about a BGP peer, including the “IPv6 label capability” is displayed using the **show bgp ipv6 neighbors** user EXEC command with an IP address.

show mpls forwarding-table Command

This topic identifies the syntax of the **show mpls forwarding-table** command and the important fields in its output.

show mpls forwarding-table Command

Cisco.com

```
Router> show mpls forwarding-table
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
16	Aggregate	IPv6	0		
17	Aggregate	IPv6	0		
18	Aggregate	IPv6	0		
19	Pop tag	192.168.99.64/30	0	Se0/0	point2point
20	Pop tag	192.168.99.70/32	0	Se0/0	point2point
21	Pop tag	192.168.99.200/32	0	Se0/0	point2point
22	Aggregate	IPv6	5424		
23	Aggregate	IPv6	3576		
24	Aggregate	IPv6	2600		

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-6

In this example, output information that links the MPLS label with prefixes is displayed using the **show mpls forwarding-table** user EXEC command. If the Cisco 6PE feature is configured, the labels are aggregated because there are several prefixes for one local label, and the prefix column contains "IPv6" instead of a target prefix.

show bgp ipv6 labels Command

This topic identifies the syntax of the **show bgp ipv6 labels** command and the important fields in its output.

show bgp ipv6 labels Command

Cisco.com

```
Router> show bgp ipv6 labels
```

Network	Next Hop	In tag/Out tag
2001:0DB8:DDDD::/64	::FFFF:192.168.99.70	notag/20

© 2004 Cisco Systems, Inc. All rights reserved.MPLS v2.0—7-7

In this example, output information about the top-of-the-stack label with label-switching information is displayed using the **show bgp ipv6** user EXEC command with the **labels** keyword.

show ipv6 cef Command

This topic identifies the syntax of the **show ipv6 cef** command and the important fields in its output.

show ipv6 cef Command

Cisco.com

```
Router> show ipv6 cef 2001:0DB8:DDDD::/64

2001:0DB8:DDDD::/64
  nexthop ::FFFF:192.168.99.70
  fast tag rewrite with Se0/0, point2point, tags imposed {19 20}
```

© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0—7-8

In this example, output information about labels from the CEF table is displayed using the **show ipv6 cef** user EXEC command with an IPv6 prefix.

show ipv6 route Command

This topic identifies the syntax of the **show ipv6 route** command and the important fields in its output.

show ipv6 route Command

Cisco.com

```
Router> show ipv6 route

IPv6 Routing Table - 10 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
B 2001:0DB8:DDDD::/64 [200/0]
   via ::FFFF:192.168.99.70, IPv6-mpls
B 2001:0DB8:DDDD::/64 [200/0]
   via ::FFFF:192.168.99.70, IPv6-mpls
L 2001:0DB8:FFFF::1/128 [0/0]
   via ::, Ethernet0/0
C 2001:0DB8:FFFF::/64 [0/0]
   via ::, Ethernet0/0
S 2001:0DB8:FFFF::/48 [1/0]
   via 2001:B00:FFFF::2, Ethernet0/0
```

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—7-9

In this example, output information from the IPv6 routing table is displayed using the **show ipv6 route** user EXEC command. The output shows the IPv6 MPLS virtual interface as the output interface of IPv6 routes that are forwarded across the MPLS cloud.

The 6PE2 router has advertised the IPv6 prefix of 2001:0DB8:DDDD::/48 configured for the CE2 router. The next-hop address is the IPv4-compatible IPv6 address ::FFFF:192.168.99.70, where 192.168.99.70 is the IPv4 address of the 6PE2 router.

Lesson Summary

This topic summarizes the key points discussed in this lesson.

Summary

Cisco.com

- **Use IPv4 and IPv6 tunnel-related show and debug commands to monitor IPv6-to-IPv4 tunnels.**
- **Use AToM show and debug commands to monitor AToM operations.**
- **Use MPLS show and debug commands to monitor MPLS operations.**
- **Use IPv6 CEF show and debug commands to monitor CEF operations.**
- **Use IPv6 route show and debug commands to monitor routing updates.**

© 2004 Cisco Systems, Inc. All rights reserved.MPLST v2.0—7-10

References

For additional information, refer to these resources:

- 6bone Connection Using 6to4 Tunnels for IPv6
http://www.cisco.com/en/US/partner/tech/tk872/tk373/technologies_configuration_example09186a00801f3b4f.shtml
- Implementing Tunneling for IPv6
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d6604.html
- Implementing IPv6 over MPLS
http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d6636.html#54831

Next Steps

For the associated lab exercise, refer to the following section of the course Lab Guide:

- Lab Exercise 7-1: Implementing Cisco 6PE Support (Case Study)

Lesson Review

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Lesson Answer Key.

- Q1) When you are deploying the IPv6 model for using tunnels on the customer edge routers, which command set do you use to monitor operations?
- A) Cisco 6PE commands
 - B) AToM Layer 2 commands
 - C) IPv4-to-IPv6 tunnel commands
 - D) IPv6-to-IPv4 tunnel commands
- Q2) When you are deploying the model for using IPv6 over a circuit transport over MPLS, which command set do you use to monitor operations?
- A) Cisco 6PE commands
 - B) AToM Layer 2 commands
 - C) IPv4-to-IPv6 tunnel commands
 - D) IPv6-to-IPv4 tunnel commands
- Q3) Which command do you use to display the BGP information that is related to a specific prefix?
- A) **show bgp ipv6**
 - B) **show ipv6 cef**
 - C) **show ipv6 route**
 - D) **show bgp ipv6 labels**
 - E) **show bgp ipv6 neighbors**
 - F) **show mpls forwarding-table**
- Q4) Which command do you use to display the status of a BGP session between two peer routers?
- A) **show bgp ipv6**
 - B) **show ipv6 cef**
 - C) **show ipv6 route**
 - D) **show bgp ipv6 labels**
 - E) **show bgp ipv6 neighbors**
 - F) **show mpls forwarding-table**
- Q5) Which command do you use to display the information that links an MPLS label with the prefixes that are known by a router?
- A) **show bgp ipv6**
 - B) **show ipv6 cef**
 - C) **show ipv6 route**
 - D) **show bgp ipv6 labels**
 - E) **show bgp ipv6 neighbors**
 - F) **show mpls forwarding-table**

- Q6) Which command do you use to display the information about the top-of-the-stack label with label-switching information?
- A) **show bgp ipv6**
 - B) **show ipv6 cef**
 - C) **show ipv6 route**
 - D) **show bgp ipv6 labels**
 - E) **show bgp ipv6 neighbors**
 - F) **show mpls forwarding-table**
- Q7) Which command do you use to display the information about the interface that will be used to forward packets to a prefix and about the label that will be used?
- A) **show bgp ipv6**
 - B) **show ipv6 cef**
 - C) **show ipv6 route**
 - D) **show bgp ipv6 labels**
 - E) **show bgp ipv6 neighbors**
 - F) **show mpls forwarding-table**

Lesson Answer Key

- Q1) D
Relates to: Monitoring IPv6 Support
- Q2) B
Relates to: Monitoring IPv6 Support
- Q3) A
Relates to: show bgp ipv6 Command
- Q4) E
Relates to: show bgp ipv6 neighbors Command
- Q5) F
Relates to: show mpls forwarding-table Command
- Q6) D
Relates to: show bgp ipv6 labels Command
- Q7) B
Relates to: show ipv6 cef Command

Course Glossary

The Course Glossary for *Implementing Cisco MPLS Traffic Engineering and Other Features* (MPLST) v2.0 highlights and defines key terms and acronyms used throughout this course. Many of these terms are also described in the Cisco Internetworking Terms and Acronyms resource, available via <http://www.cisco.com>.

Acronym or Term	Definition
ACL	access control list. Filter list used for services such as security, QoS, and routing. A list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router).
AF	There are four AF classes, AF1x through AF4x. Within each class, there are three drop probabilities. Depending on the policy of a given network, packets can be selected for a PHB based on required throughput, delay, jitter, or loss, or according to priority of access to network services.
ATM	Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3.
ATM-LSR	A label switch router with a number of LSC-ATM interfaces. The router forwards the cells among these interfaces using labels carried in the VPI/VCI field.
ATM edge LSR	A router that is connected to the ATM-LSR cloud through LSC-ATM interfaces. The ATM edge LSR adds labels to unlabeled packets and strips labels from labeled packets..
AToM	Any Transport over MPLS (AToM) allows service providers that offer Layer 2 connectivity to expand service offerings by connecting Ethernet, ATM, Frame Relay, and serial or PPP networks through an MPLS backbone. AToM is a scalable architecture based on label switching that allows multiplexing of connections. It is also a standards-based open architecture and can be extended to other transport types.
BA	Behavior Aggregate. A collection of packets with the same codepoint crossing a link in a particular direction. The terms "aggregate" and "behavior aggregate" are used interchangeably.
BGP	Border Gateway Protocol. Interdomain routing protocol that replaces exterior gateway protocol (EGP). BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.
CAR	committed access rate. A traffic policing and marking mechanism. The CAR and DCAR (distributed CAR) services limit the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.
CBWFQ	class-based weighted fair queuing. CBWFQ extends the standard WFQ functionality to provide support for user-defined traffic classes. Allows the user to define traffic classes based on customer-defined match criteria such as ACLs, input interfaces, protocol, and QoS label. When traffic classes have been defined, they can be assigned a bandwidth, queue limit, or drop policy such as weighted random early detection (WRED).
CE router	customer edge router. A router that is part of a customer network and that interfaces to a PE router.
CEF	Cisco Express Forwarding. An advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns. VRF tables use CEF technology; therefore, MPLS VPNs must be CEF-enabled.
CIR	committed information rate. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. CIR, measured in bits per second, is one of the key negotiated tariff metrics.
Cisco IOS software	Cisco operating system software that runs on routers and switches. Cisco system software that provides common functionality, scalability, and security for all products under the CiscoFusion architecture. Cisco IOS software allows centralized, integrated, and automated installation and management of internetworks while ensuring support for a wide variety of protocols, media, services, and platforms.

Acronym or Term	Definition
classification	Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, by using the three IP precedence bits in the ToS field of the IP packet header—two of the values are reserved for other purposes—packets can be categorized into a limited set of up to six traffic classes. After packet classification, other QoS features can be used to assign the appropriate traffic-handling policies, including congestion management, bandwidth allocation, and delay bounds for each traffic class.
CLI	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.
CLP	cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions to free up resources to handle insured traffic.
congestion	Traffic in excess of network capacity.
congestion avoidance	Mechanism by which an ATM network controls the traffic entering the network to minimize delays. To use resources most efficiently, lower-priority traffic is discarded at the edge of the network if conditions indicate that it cannot be delivered.
CoS	class of service. CoS refers to the methods that provide differentiated service, in which the network delivers a particular kind of service based on the class of service specified for each packet. CoS provides specific categories of service such as gold, silver, and best-effort service classes. CoS is a set of concrete device features in which a single network router treats traffic in different classes differently. CoS techniques provide a means of specifying policies to control network resource allocation in support of customer and applications requirements. The implementation of CoS techniques delivers measurable QoS.
CR-LDP	Constraint Route-Label Distribution Protocol
DiffServ	Differentiated Services. A paradigm for providing QoS on the Internet by employing a small, well-defined set of building blocks from which a variety of services can be built.
DLCI	data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the Local Management Interface (LMI) extended specification, DLCIs are globally significant (DLCIs specify individual end devices).
DSCP	The Differentiated Services Code Point (DSCP) octet. In the IP header, DSCP classifies the packet service level. The DSCP maps to a particular observable forwarding behavior called a per-hop behavior (PHB). The DSCP replaces the ToS octet in the IP version 4 (IPv4) header, and the Class octet in the IP version 6 (IPv6) header. Currently, only the first six bits are used, allowing up to 64 different classifications for service levels. The DSCP is unstructured, but it does reserve some values to maintain limited backward compatibility with the IP precedence bits in the ToS octet.
EBGP	External Border Gateway Protocol. EBGP (see "BGP") communicate among different network domains or autonomous systems. The primary function of an EBGP is to exchange network reachability information between autonomous systems, including information about the list of autonomous system routes. The autonomous systems use EBGP border edge routers to distribute the routes, which include label-switching information. Each border edge router rewrites the next-hop and MPLS labels.
EF	Expedited Forwarding. The EF PHB can be used to build a low-loss, low-latency, low-jitter, assured-bandwidth, end-to-end service through DiffServ domains. Such a service appears to the endpoints like a point-to-point connection or a "virtual leased line." This service has also been described as premium service. Codepoint 101110 is recommended for the EF PHB.
EXP	MPLS EXP bits. MPLS is a switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on pre-established IP routing information. The EXP bit is part of the label header. These bits can be used to provide optional QoS marking.

Acronym or Term	Definition
FEC	forwarding equivalence class. An FEC is a group of IP packets that are forwarded in the same manner, over the same path, and with the same forwarding treatment. An FEC might correspond to a destination IP subnet, but it also might correspond to any traffic class that the edge LSR considers significant. For example, all traffic with a certain IP precedence value might constitute an FEC.
FIB	Forwarding Information Base. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information that is contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information that is based on the information in the IP routing table. Because there is a one-to-one correlation between FIB entries and routing table entries, the FIB contains all known routes and eliminates the need for route cache maintenance that is associated with switching paths such as fast switching and optimum switching.
Frame Relay	Industry-standard, switched, data link layer protocol that handles multiple virtual circuits using High-Level Data Link Control (HDLC) encapsulation between connected devices. Frame Relay is more efficient than X.25, the protocol for which it generally is considered a replacement.
FRF.11	Frame Relay Forum implementation agreement for Voice over Frame Relay, or VoFR (v1.0 May 1997). This specification defines multiplexed data, voice, fax, dual tone multifrequency (DTMF) digit relay, and channel-associated signaling (CAS), or robbed-bit signaling, frame formats but does not include call setup, routing, or administration facilities. See www.frforum.com .
FRF.12	Frame Relay Forum implementation agreement for Frame Relay Fragmentation. The FRF.12 Implementation Agreement (also known as FRF.11 Annex C) was developed to allow long data frames to be fragmented into smaller pieces and interleaved with real-time frames. In this way, real-time voice and data frames that are not in real time can be carried together on lower-speed links without causing excessive delay to the real-time traffic. See www.frforum.com .
FRTS	Frame Relay Traffic Shaping. Queuing method that uses queues on a Frame Relay network to limit surges that can cause congestion. Data is buffered and sent into the network in regulated amounts to ensure that the traffic can fit within the promised traffic envelope for the particular connection.
GRE	generic routing encapsulation. Tunneling protocol that was developed by Cisco Systems that can encapsulate a wide variety of protocol packet types inside IP tunnels and thus create a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment.
GTS	Generic Traffic Shaping (GTS). Provides a mechanism to control the traffic flow on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate (also known as the token bucket approach), while queuing bursts of the specified traffic. Thus, traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data-rate mismatches.
IETF	Internet Engineering Task Force. Task force that consists of more than 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of the Internet Society (ISOC).
IntServ	Integrated Services
IP	Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, ToS specification, fragmentation and reassembly, and security. Defined in RFC 791.
IP precedence	A 3-bit value in the TOS byte that is used for assigning precedence to IP packets.

Acronym or Term	Definition
IPSec	IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPSec provides these security services at the IP layer. IPSec uses Internet Key Exchange (IKE) to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPSec. IPSec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.
ISP	Internet service provider. Provider of internet access and services through a single BGP autonomous system (AS).
L2F	Layer 2 Forwarding Protocol. Protocol that supports the creation of secure, virtual, private dialup networks over the Internet.
L2TP	Layer 2 Tunneling Protocol. Protocol that is used for implementing virtual private dial-up networks (VPDNs) and VPNs by tunneling PPP with multivendor interoperability and acceptance. This protocol was proposed as an alternative to IPSec, but is often used with IPSec for authentication. This protocol merges the Microsoft PPTP and Cisco L2F technologies.
label	A label is a header that is used by an LSR to forward packets. The header format depends upon network characteristics. In router networks, the label is a separate, 32-bit header. In ATM networks, the label is placed into the virtual channel identifier/virtual path identifier (VCI/VPI) cell header. In the core, LSRs read only the label, not the packet header. One key to the scalability of MPLS is that labels have only local significance between two devices that are communicating.
label imposition	The act of putting the first label on a packet.
LAN	local-area network. High-speed, low-error data network that covers a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the Open Systems Interconnection (OSI) model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
LC-ATM	label-controlled ATM. An interface on a router or switch that uses label distribution procedures to negotiate label virtual circuits.
LDP	Label Distribution Protocol. Provides communication between edge and core devices. It assigns labels in edge and core devices to establish LSPs in conjunction with routing protocols such as Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Enhanced Interior Gateway Routing Protocol (EIGRP), or BGP.
LFIB	label forwarding information base. A data structure and way of managing forwarding in which destinations and incoming labels are associated with outgoing interfaces and labels.
LIB	label information base. A database that is used by an LSR to store labels that are learned from other LSRs, as well as labels that are assigned by the local LSR.
LSP	label switched path. A sequence of hops (R0...Rn) in which a packet travels from R0 to Rn through label-switching mechanisms. A label switched path can be established dynamically, based on normal routing mechanisms, or it can be established through configuration.
LSP tunnel	label switched path tunnel. A configured connection between two routers in which MPLS is used to carry the packet.
LSR	label switch router. The core device that switches labeled packets according to precomputed switching tables. It can also be a switch or a router.

Acronym or Term	Definition
MIB	Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as Simple Network Management Protocol (SNMP) or Common Management Information Protocol (CMIP). The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a graphical user interface (GUI) network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.
MP-BGP	Multiprotocol Border Gateway Protocol. An extension to the BGP protocol that allows VPN information to remain unique within the MPLS VPN backbone. MP-BGP also allows BGP speakers to identify routing updates that do not carry standard IP version 4 (IPv4) prefix information. Multiprotocol (MP-BGP) and VPN IPv4 routing information provide these extensions. IBGP (Internal BGP) refers to BGP running within a single autonomous system (AS). EBGP (External BGP) refers to BGP running between autonomous systems.
MPLS	Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on pre-established IP routing information.
MPLS-TE	MPLS Traffic Engineering. The techniques and processes that are used to cause routed traffic to travel through the network on a path other than the one that would have chosen if standard routing methods had been used.
MPLS VPN	Multiprotocol Label Switching Virtual Private Network. The MPLS VPN solution is a set of PE routers that are connected via a common “backbone” network to supply private IP interconnectivity between two or more customer sites for a given customer. Each VPN has a set of provisioning templates and policies and can span multiple provider administrative domains (PADs). Customer edge routing communities (CERCs) in a VPN break down complex topologies into manageable subgroups.
MQC	Modular QoS CLI. The Modular QoS CLI is a CLI structure that allows users to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the traffic policy determine how to treat the classified traffic.
NBAR	network-based application recognition (NBAR). A new classification engine that can recognize a wide variety of application-level protocols, including HTTP via URL/Multipurpose Internet Mail Extensions (MIME) type and protocols that use dynamic port assignments. When the traffic is classified by NBAR, appropriate QoS policies can be applied to the traffic classes using existing Cisco IOS QoS features.
packet	Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms datagram, frame, message, and segment also are used to describe logical information groupings at various layers of the Open Systems Interconnection (OSI) reference model and in various technology circles.
payload	Portion of a cell, frame, or packet that contains upper-layer information (data).
PBR	policy-based routing. Routing scheme that forwards packets to specific interfaces based on user-configured policies. Such policies might specify that traffic sent from a particular network should be forwarded out one interface, while all other traffic should be forwarded out another interface.
PE router	provider edge router. A router that is part of a service provider network and that is connected to a CE router. The PE router function is a combination of an MPLS edge LSR function with some additional functions to support VPNs.
PHB	per-hop behavior. RFC 2475 defines PHB as the externally observable forwarding behavior applied at a DiffServ-compliant node to a DiffServ BA. With the ability of the system to mark packets according to DSCP setting, collections of packets with the same DSCP setting, and sent in a particular direction, can be grouped into a BA. Packets from multiple sources or applications can belong to the same BA. The PHB is the packet scheduling, queuing, policing, or shaping behavior of a node on any given packet belonging to a BA, as configured by an SLA or a policy map.

Acronym or Term	Definition
PHP	penultimate hop popping
PIM	Protocol Independent Multicast. Multicast routing architecture that allows the addition of IP multicast routing on existing IP networks. PIM is independent of unicast routing protocols and can be operated in two modes: dense and sparse.
POP	point of presence. Service points of presence (PoPs) groom traffic from the customer's network, perform edge packet switching when IP services are enabled, and perform backbone switching when PoPs are interconnected. Service PoPs are also hubs of high-value Internet services such as web content, Domain Name Service (DNS) servers to Internet service providers, Virtual Private Network (VPN) services, and other applications deployed on a range of flexible, scalable, high-performance IP routers.
PPP	Point-to-Point Protocol. Successor to Serial Line Internet Protocol (SLIP) that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, Internetwork Packet Exchange (IPX), and AppleTalk Remote Access (ARA). PPP also has built-in security mechanisms, such as Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP). PPP relies on two protocols: link control protocol (LCP) and Network Control Protocol (NCP).
PPTP	Point-to-Point Tunneling Protocol. RFC 2637 describes the PPTP protocol.
PVC	permanent virtual circuit (or connection). Virtual circuit that is permanently established. PVCs save bandwidth that is associated with circuit establishment and teardown in situations where certain virtual circuits must exist all the time. In ATM terminology, called a "permanent virtual connection."
QoS	quality of service (QoS). The goal of QoS is to provide better and more predictable network service by providing dedicated bandwidth, controlled jitter and latency, and improved loss characteristics. QoS achieves these goals by providing tools for managing network congestion, shaping network traffic, using expensive wide-area links more efficiently, and setting traffic policies across the network.
RD	route distinguisher. A route distinguisher (RD) creates routing and forwarding tables and specifies the default route distinguisher for a VPN. The RD is an 8-byte value that is added to the beginning of the IP version 4 (IPv4) prefixes of the customer to change them into globally unique VPN IPv4 prefixes. An RD is either ASN-relative, in which case it is composed of an autonomous system (AS) number and an arbitrary number, or it is IP-address-relative, in which case it is composed of an IP address and an arbitrary number. Each VPN VRF table has an RD. Prefixes should use the same RD if they are associated with the same set of RTs. The community of interest association is based on the RT extended community attributes that are distributed with Network Layer Reachability Information (NLRI). The RD value must be a globally unique value to avoid conflict with other prefixes.
RED	random early detection. This class of algorithms is designed to avoid congestion in internetworks before it becomes a problem. RED works by monitoring traffic load at points in the network and randomly discarding packets if the congestion begins to increase. The result of the drop is that the source detects the dropped traffic and slows its transmission. RED is designed to work primarily with TCP in IP internetwork environments.
RSVP	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams that they want to receive. RSVP depends on IP version 6 (IPv6). Also known as Resource Reservation Setup Protocol.
RT	route target. A 64-bit value by which Cisco IOS software discriminates routes for route updates in VRF tables. A route target (RT) defines the destination of the route by specifying a target VPN community. RTs identify a group of sites that exchange routing information. BGP uses the RT value when determining the eligibility of routes for installation to a local routing table.

Acronym or Term	Definition
RTP	Real-Time Transport Protocol. Commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications that transmit real-time data, such as audio, video, or simulation data, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, time stamping, and delivery monitoring to real-time applications.
service class	Collection of service types that are required for a specific service that is being offered. Each service class includes the attributes and values that define the type or QoS that is associated with a given class. For example, data connectivity is a service class that you might define that includes the service type "data-bandwidth."
SLA	service level agreement. Service level agreements (SLAs) are negotiated contracts between VPN providers and their subscribers. An SLA defines the criteria for the specific services that the subscriber expects the provider to deliver. The SLA is the only binding mechanism that the subscriber has to ensure that the VPN provider delivers the services as agreed.
SOO	Site of Origin (SOO) is a concept in MPLS VPN architecture that prevents routing loops in a site that is multihomed to the MPLS VPN backbone and in a site that uses AS-OVERRIDE. SOO is a BGP extended community attribute that is used to identify an IP address that originated from a site to prevent that IP address from being advertised back to the site. This attribute uniquely identifies the site from which the PE router learned the route. The SOO is tagged at a PE router in peering with BGP neighbors using an inbound route map. It works in conjunction with the BGP customer edge-provider edge (CE-PE) routing protocol. The SOO must be unique for each customer site for each VPN. Therefore, the same value of SOO must be used on all PE routers that are connected to the same CE router or to the same customer site.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack.
TDP	Tag Distribution Protocol. The protocol that is used to distribute label bindings to LSRs (Cisco proprietary version of LDP).
ToS	type of service. A byte in the IP version 4 (IPv4) header.
traffic engineering tunnel	A label-switched path tunnel that is used for engineering traffic. It is set up through means other than normal Layer 3 routing and is used to direct traffic over a path different from the one that Layer 3 routing would cause it to take.
tunnel	Secure communication path between two peers, such as two routers.
tunneling	Architecture that provides the services that are necessary to implement any standard point-to-point data encapsulation scheme.
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
VC merge	virtual circuit merge. VC merge allows the switch to transmit cells coming from different virtual channel identifiers (VCIs) over the same outgoing VCI toward the same destination. In other words, it allows multipoint-to-point connections. VC merge is accomplished by queuing complete ATM adaptation layer 5 (AAL5) frames in input buffers until the end of the frame has been received. The cells from the same AAL5 frame are all transmitted before sending cells from any other frame. This queuing requires sufficient buffering capability inside the switch, but no more buffering than is required in IP networks.
VIP	Versatile Interface Processor. Interface card that is used in Cisco 7000 and Cisco 7500 series routers. The VIP provides multilayer switching and runs Cisco IOS software. The most recent version of the VIP is VIP2.

Acronym or Term	Definition
virtual circuit	virtual circuit. Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either a PVC or a switched virtual circuit (SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a "virtual channel." Sometimes abbreviated "VC."
VLAN	virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
VoIP	Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the digital signal processor (DSP) segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with International Telecommunication Union Telecommunication Standardization Sector (ITU-T) specification H.323.
VPI/VCI	virtual path identifier/virtual circuit identifier. An 8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. VCI - virtual channel identifier. A 16-bit field in the header of an ATM cell. The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination.
VPN	Virtual Private Network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses "tunneling" to encrypt all information at the IP level.
VPNv4	Used as a keyword in commands to indicate VPN IPv4 prefixes. These prefixes are customer VPN addresses, each of which has been made unique by the addition of an 8-byte RD.
VRF table	virtual routing and forwarding table. A VRF table consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF table includes the routing information that defines a customer VPN site that is attached to a PE router.
WAN	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices that are provided by common carriers. Frame Relay, Switched Multimegabit Data Service (SMDS), and X.25 are examples of WANs.
WFQ	weighted fair queuing. Congestion management algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that the capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission.

Lab Guide

Overview

Use the tasks here to complete the lab exercises for this course. The solutions for the lab exercises are found in the Lab Exercise Answer Key.

Outline

This Lab Guide includes these exercises:


- Lab Exercise 2-1: Implementing CSC in a Native IP Environment
- Lab Exercise 2-2: Implementing CSC Using LDP Label Distribution
- Lab Exercise 2-3: Implementing CSC Using MP-BGP Label Distribution
- Lab Exercise 2-4: Implementing CSC Using MPLS VPNs
- Lab Exercise 4-1: Implementing MPLS Traffic Engineering
- Lab Exercise 4-2: Implementing Link Protection for MPLS Traffic Engineering
- Lab Exercise 5-1: Implementing MPLS QoS
- Lab Exercise 5-2: Implementing MPLS QoS and DS-TE
- Lab Exercise 6-1: Implementing a Port-to-Port FRoMPLS Configuration (Case Study)
- Lab Exercise 6-2: Implementing a DLCI-to-DLCI FRoMPLS Configuration (Case Study)
- Lab Exercise 7-1: Implementing Cisco 6PE Support (Case Study)

Lab Exercise 2-1: Implementing CSC in a Native IP Environment

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

It is common for lower-tier Internet service providers (ISPs) to interconnect their POPs through a higher-tier ISP. In the past this task has been accomplished by using an EBGP session to pass routing information. With the advent of higher-level services, however, this method has started to suffer from severe limitations. The CSC feature was introduced to overcome these limitations.



In this exercise, carrier Yellow needs to provide connections for its customers, Green and Orange. To do this, Yellow has contracted with the carrier Blue to provide a connection between its POP sites.

For this lab exercise, your workgroup will be divided into two teams.

Team A will be responsible for the following:

- Installing the initial configurations in routers R1, R5, R7, and R3
- Enabling CSC support on routers R1, R5, R7, and R3


Team B will be responsible for the following:

- Installing the initial configurations in routers R2, R6, R8, and R4
- Enabling CSC support on routers R2, R6, R8, and R4

After completing this exercise, you will be able to meet these objectives:


- Configure the initial configuration on all routers
- Configure the PE routers of the backbone carrier (R7 and R8)
- Configure the CE routers of the customer carrier (R1, R2, R5, and R6)

Visual Objective



For this exercise, there are four distinctive IP address ranges. The 192.168.0.0 range is used to provide addressing for the core MPLS network of carrier Blue. The 150.0.0.0 range is used to provide addressing for the links between carrier Blue and the customer carrier Yellow. The 10.0.0.0 is used to provide addressing for the core of carrier Yellow. The 172.11.0.0 and 172.12.0.0 are used to provide network addressing for customers Green and Orange.

Note The addresses of customers Green and Orange are simulated using loopback addresses.



Customer carrier Yellow is running a native network. Open Shortest Path First (OSPF) is being used as the IGP, and BGP is being used to carry customer routes. As this lab starts, OSPF and BGP have been enabled, but the IBGP sessions have not been established between the POP sites of carrier Yellow.

Note To simplify the lab, the EBGP sessions with the sites of customers Green and Orange have been simulated with network statements in the BGP processes.

Carrier Blue is running an MPLS VPN network, with Intermediate System-to-Intermediate System (IS-IS) used as the IGP, and MP-BGP used to carry VPN labels. As this lab starts, IS-IS has been enabled, and all MPLS and MP-BGP sessions have been activated.

During this lab you will configure CSC support on the routers of carrier Blue and IGP and BGP connectivity on the routers of customer carrier Yellow.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Command	Description
address-family ipv4 [multicast unicast vrf vrf-name] no address-family ipv4 [multicast unicast vrf vrf-name]	To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the address-family ipv4 command in router configuration mode. To disable address family configuration mode, use the no form of this command.
ip vrf forwarding vrf-name no ip vrf forwarding vrf-name	To associate a VRF instance with an interface or subinterface, use the ip vrf forwarding command in interface configuration mode. To disassociate a VRF instance, use the no form of this command.
ip vrf vrf-name no ip vrf vrf-name	To configure a VRF routing table, use the ip vrf command in global configuration mode or router configuration mode. To remove a VRF routing table, use the no form of this command.
neighbor {ip-address peer-group-name ipv6- address} activate no neighbor {ip-address peer-group-name ipv6- address} activate	To enable the exchange of information with a BGP neighbor, use the neighbor activate command in address family configuration mode or router configuration mode. To disable the exchange of an address with a BGP neighbor, use the no form of this command.
rd route-distinguisher	To create routing and forwarding tables for a VRF instance, use the rd command in VRF configuration mode.
redistribute protocol [process-id] {level-1 level-1-2 level-2} [as- number] [metric metric- value] [metric-type type-value] [match {internal external 1 external 2}] no redistribute protocol [process-id] {level-1 level-1-2 level-2} [as- number] [metric metric- value] [metric-type type-value] [match {internal external 1 external 2}]	To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. To disable redistribution, use the no form of this command.
router bgp as-number no router bgp as-number	To configure the BGP routing process, use the router bgp command in global configuration mode. To remove a routing process, use the no form of this command.
router ospf process-id no router ospf process-id	To configure an OSPF routing process, use the router ospf command in global configuration mode. To terminate an OSPF routing process, use the no form of this command.
route-target {import export both} route- target-ext-community no route-target {import export both} route- target-ext-community	To create a route target (RT) extended community for a VRF instance, use the route-target command in VRF configuration mode. To disable the configuration of an RT community option, use the no form of this command.

Task 1: Implement the Initial Configurations on the Routers

The backbone of carrier Blue is configured as a standard MPLS network with IBGP peering between the PE routers. The cores of carrier Yellow and customer Orange core are configured as a standard native IP network. To simplify the lab, the EBGP sessions with the sites of customer Green and Orange have been simulated with loopback addresses. To establish this environment, configure all routers as indicated here.

Exercise Procedure

Complete the following commands:

Step 1 On each router erase the start configuration.

Step 2 Reload each router without saving the running configuration.

Team A

Step 1 On R1 install the following configuration:

```
hostname R1

ip cef

interface Loopback0
 ip address 10.1.11.49 255.255.255.255

interface loopback11
 ip address 172.11.1.16 255.255.255.0

interface loopback12
 ip address 172.12.1.16 255.255.255.0

interface Serial0/0
 bandwidth 128000
 no ip address
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R2
 ip address 10.1.11.129 255.255.255.240
 frame-relay interface-dlci 101

router ospf 1
 log-adjacency-changes
 router-id 10.1.11.49
 network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
 no synchronization
 bgp log-neighbor-changes
 network 172.11.1.0 mask 255.255.255.0
 network 172.12.1.0 mask 255.255.255.0
 neighbor 10.1.12.49 remote-as 65010
 neighbor 10.1.12.49 update-source Loopback0
 no auto-summary

line con 0
```

```
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login
```

Step 2 On R5 install the following configuration:

```
hostname R5

ip cef

interface Loopback0
 ip address 10.1.12.49 255.255.255.255

interface Serial0/0
 bandwidth 128000
 no ip address
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R1
 ip address 10.1.11.130 255.255.255.240
 frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
 description Link to R7
 ip address 150.1.11.17 255.255.255.240
 frame-relay interface-dlci 102

router ospf 1
 log-adjacency-changes
 router-id 10.1.12.49
 network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.11.49 remote-as 65010
 neighbor 10.1.11.49 update-source Loopback0
 neighbor 10.1.11.49 route-reflector-client
 no auto-summary

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
no login
```

Step 3 On R7 install the following configuration:

```
hostname R7

ip cef

interface Loopback0
 ip address 192.168.1.17 255.255.255.255
 ip router isis
```

```

interface Serial0/0
  bandwidth 128000
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.102 point-to-point
  desc link to R5
  ip address 150.1.11.18 255.255.255.240
  frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
  desc link to R3
  ip address 192.168.1.49 255.255.255.240
  ip router isis
  mpls label protocol tdp
  tag-switching ip

  frame-relay interface-dlci 111

router isis
  net 39.0440.0000.0000.0500
  metric-style wide

router bgp 65001
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.1.33 remote-as 65001
  neighbor 192.168.1.33 update-source Loopback0
  no auto-summary

  address-family vpnv4
  neighbor 192.168.1.33 activate
  neighbor 192.168.1.33 next-hop-self
  neighbor 192.168.1.33 send-community both
  no auto-summary
  exit-address-family
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Step 4 On R3 install the following configuration:

```

hostname R3

ip cef

interface Loopback0
  ip address 192.168.1.81 255.255.255.255
  ip router isis

interface Serial0/0

```

```

no ip address
bandwidth 128000
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
no shut

interface Serial0/0.111 point-to-point
desc link to R7
ip address 192.168.1.50 255.255.255.240
ip router isis
mpls ip
mpls label protocol tdp
frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
desc link to R4
ip address 192.168.1.113 255.255.255.240
ip router isis
mpls ip
mpls label protocol tdp
frame-relay interface-dlci 112
exit

router isis
net 39.0440.0000.0000.0700
metric-style wide

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

Team B

Step 1 On R2 install the following configuration:

```

hostname R2

ip cef

interface Loopback0
ip address 10.2.11.49 255.255.255.255

interface loopback11
ip address 172.11.2.16 255.255.255.0

interface loopback12
ip address 172.12.2.16 255.255.255.0

interface Serial0/0
bandwidth 128000
no ip address
encapsulation frame-relay
frame-relay lmi-type cisco
no shut

interface Serial0/0.101 point-to-point

```



```

description Link to R6
ip address 10.2.11.129 255.255.255.240
frame-relay interface-dlci 101

router ospf 1
log-adjacency-changes
router-id 10.2.11.49
network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
no synchronization
bgp log-neighbor-changes
network 172.11.2.0 mask 255.255.255.0
network 172.12.2.0 mask 255.255.255.0
neighbor 10.2.12.49 remote-as 65010
neighbor 10.2.12.49 update-source Loopback0
no auto-summary

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

Step 2 On R6 install the following configuration:

```

hostname R6

ip cef

interface Loopback0
ip address 10.2.12.49 255.255.255.255

interface Serial0/0
bandwidth 128000
no ip address
encapsulation frame-relay
frame-relay lmi-type cisco
no shut

interface Serial0/0.101 point-to-point
description Link to R2
ip address 10.2.11.130 255.255.255.240
frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
description Link to R8
ip address 150.1.12.17 255.255.255.240
frame-relay interface-dlci 102

router ospf 1
log-adjacency-changes
router-id 10.2.12.49
network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
no synchronization
bgp log-neighbor-changes
neighbor 10.2.11.49 remote-as 65010
neighbor 10.2.11.49 update-source Loopback0

```

```

neighbor 10.2.11.49 route-reflector-client
no auto-summary
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

Step 3 On R8 install the following configuration:

```

hostname R8

ip cef

interface Loopback0
ip address 192.168.1.33 255.255.255.255
ip router isis

interface Serial0/0
bandwidth 128000
no ip address
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
no shut

interface Serial0/0.102 point-to-point
desc link to R6
ip address 150.1.12.18 255.255.255.240
frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
desc link to R4
ip address 192.168.1.65 255.255.255.240
ip router isis
tag-switching ip
mpls label protocol tdp
frame-relay interface-dlci 111

router isis
net 39.0440.0000.0000.0600
metric-style wide

router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 192.168.1.17 remote-as 65001
neighbor 192.168.1.17 update-source Loopback0
no auto-summary

address-family vpnv4
neighbor 192.168.1.17 activate
neighbor 192.168.1.17 next-hop-self
neighbor 192.168.1.17 send-community both
no auto-summary
exit-address-family
exit

```

```

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  login

```

Step 4 On R4 install the following configuration:

```

hostname R4

ip cef

interface Loopback0
  ip address 192.168.1.97 255.255.255.255
  ip router isis

interface Serial0/0
  bandwidth 128000
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut
!
interface Serial0/0.111 point-to-point
  desc link to R8
  ip address 192.168.1.66 255.255.255.240
  ip router isis
  tag-switching ip
  mpls label protocol tdp
  frame-relay interface-dlci 111
!
interface Serial0/0.112 point-to-point
  desc link to R3
  ip address 192.168.1.114 255.255.255.240
  ip router isis
  tag-switching ip
  mpls label protocol tdp
  frame-relay interface-dlci 112

router isis
  net 39.0440.0000.0000.0800
  metric-style wide
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Exercise Verification

You have completed this exercise when you attain these results:

Team A

- On R1, use the **sh ip rou** command to verify the routing table.

```
R1#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

    172.11.0.0/24 is subnetted, 1 subnets
C      172.11.1.0 is directly connected, Loopback11
    172.12.0.0/24 is subnetted, 1 subnets
C      172.12.1.0 is directly connected, Loopback12
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C      10.1.11.49/32 is directly connected, Loopback0
O      10.1.12.49/32 [110/65] via 10.1.11.130, 00:05:01, Serial0/0.101
C      10.1.11.128/28 is directly connected, Serial0/0.101
```

- On R5, use the **sh ip rou** command to verify the routing table.

```
R5#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

    172.11.0.0/24 is subnetted, 1 subnets
B      172.11.1.0 [200/0] via 10.1.11.49, 00:10:13
    172.12.0.0/24 is subnetted, 1 subnets
B      172.12.1.0 [200/0] via 10.1.11.49, 00:10:13
    10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O      10.1.11.49/32 [110/2] via 10.1.11.129, 00:09:41, Serial1/0.101
C      10.1.12.49/32 is directly connected, Loopback0
C      10.1.11.128/28 is directly connected, Serial0/0.101
    150.1.0.0/28 is subnetted, 1 subnets
C      150.1.11.16 is directly connected, Serial0/0.102
```

- On R5, use the **sh ip bgp** command to verify the routing table.

```
R5#sh ip bgp
BGP table version is 1, local router ID is 10.1.12.49
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
* 172.11.1.0/24      10.1.11.49             0      100      0 i
* 172.12.1.0/24      10.1.11.49             0      100      0 i
```

- On R7, use the **sh ip rou** command to verify the routing table.

```
R7#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

Gateway of last resort is not set

192.168.1.0/24 is variably subnetted, 5 subnets, 2 masks
i L1   192.168.1.97/32 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.112/28 [115/20] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.64/28 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.33/32 [115/40] via 192.168.1.50, Serial0/0.111
C      192.168.1.48/28 is directly connected, Serial0/0.111
C      192.168.1.17/32 is directly connected, Loopback0
150.1.0.0/28 is subnetted, 1 subnets
C      150.1.11.16 is directly connected, Serial0/0.102
```

- On R7, use the **sh ip bgp sum** command to verify the status of the BGP session.

```
R7#sh ip bgp sum
BGP router identifier 192.168.1.17, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
192.168.1.33  4 65001    15     15       1    0  0 00:11:57
0
```

Team B

- On R2, use the **sh ip rou** command to verify the routing table.

```
R2#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      172.11.0.0/24 is subnetted, 1 subnets
C      172.11.2.0 is directly connected, Loopback11
      172.12.0.0/24 is subnetted, 1 subnets
C      172.12.2.0 is directly connected, Loopback12
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C      10.2.11.49/32 is directly connected, Loopback0
O      10.2.12.49/32 [110/2] via 10.2.11.130, 00:05:54, Serial1/0.101
C      10.2.11.128/28 is directly connected, Serial1/0
```

- On R6, use the **sh ip rou** command to verify the routing table.

```
R6#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      172.11.0.0/24 is subnetted, 1 subnets
B      172.11.2.0 [200/0] via 10.2.11.49, 00:10:52
      172.12.0.0/24 is subnetted, 1 subnets
B      172.12.2.0 [200/0] via 10.2.11.49, 00:10:52
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
O      10.2.11.49/32 [110/2] via 10.2.11.129, 00:01:31, Serial0/0/101
C      10.2.12.49/32 is directly connected, Loopback0
C      10.2.11.128/28 is directly connected, Serial0/0.101
      150.1.0.0/28 is subnetted, 1 subnets
C      150.1.12.16 is directly connected, Serial0/0.102
```

- On R8, use the **sh ip rou** command to verify the routing table.

```
R8#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
192.168.1.0/24 is variably subnetted, 5 subnets, 2 masks
i L1    192.168.1.97/32 [115/20] via 192.168.1.66, Serial0/0.111
i L1    192.168.1.112/28 [115/20] via 192.168.1.66, Serial0/0.111
C       192.168.1.64/28 is directly connected, Serial1/0.111
C       192.168.1.33/32 is directly connected, Loopback0
i L1    192.168.1.48/28 [115/30] via 192.168.1.66, Serial0/0.111
i L1    192.168.1.17/32 [115/40] via 192.168.1.66, Serial0/0.111
150.1.0.0/28 is subnetted, 1 subnets
C       150.1.12.16 is directly connected, Serial0/0.102
```

- On R8, use the **sh ip bgp sum** command to verify the status of the BGP session.

```
R8#sh ip bgp sum
```

```
BGP router identifier 192.168.1.33, local AS number 65001
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
State/PfxRcd									
192.168.1.17	4	65001	17	17	1	0	0	00:13:57	0

Task 2: Configure the PE Router of the Backbone Carrier

The backbone of carrier Blue has already been configured in the same manner as the core network in a typical MPLS VPN network. An IGP has been implemented, in this case IS-IS, MPLS has been enabled on the core interfaces, and MP-BGP sessions have been established between PE routers. The following steps explain how to configure the backbone carrier PE router that links to the edge router of the customer carrier.

Exercise Procedure

Complete the following commands:

- Step 1** On your assigned PE router (R7 and R8, respectively), create a carrier Yellow VRF table that will support the customer carrier. An RD and RT of 1:10 should be used.
- Step 2** On your assigned PE router, assign the VRF table to the associated interface.
- Step 3** On your assigned PE router, create the appropriate OSPF instance to support the VRF table of the customer carrier. An instance of 1 should be used.
- Step 4** On your assigned PE router, create the appropriate BGP address families to support the VRF table of the customer carrier.
- Step 5** On your assigned PE router, enable the redistribution between OSPF and BGP that is necessary to support the VRF table of the customer carrier.

Note	For the purposes of this lab, the metric values are unimportant. However, in an actual implementation, the metric values must be set to meet network needs.

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

Team A

On R7, use the **sh ip rou** command to verify the routing table.

Note	Ensure that all of the core routers of carrier Blue are in the routing table. Verify that none of the other routes are in the IGP routing table of carrier Blue. Notice also that the addresses in the 150.1.0.0 range are no longer in the core IGP routing table because they are not in the VRF routing tables.

```
R7#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      192.168.1.0/24 is variably subnetted, 5 subnets, 2 masks
i L1   192.168.1.97/32 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.112/28 [115/20] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.64/28 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.33/32 [115/40] via 192.168.1.50, Serial0/0.111
C      192.168.1.48/28 is directly connected, Serial0/0.111
C      192.168.1.17/32 is directly connected, Loopback0
```

■ On R7, use the **sh ip rou vrf Yellow** command to verify the routing table.

Note	Verify that the VRF routing table contains the routes for local (connected) and remote (BGP) routes for interfaces that support the customer carrier.

```
R7#sh ip rou vrf Yellow
```

Routing Table: Yellow

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      150.1.0.0/28 is subnetted, 2 subnets
B      150.1.12.16 [200/0] via 192.168.1.33, 00:07:45
C      150.1.11.16 is directly connected, Serial1/0.102
```


Team B

- On R8, use the **sh ip rou** command to verify the routing table.

Note Ensure that all of the core routers of carrier Blue are in the routing table. Verify that none of the other routes are in the IGP routing table of carrier Blue. Notice also that the addresses in the 150.1.0.0 range are no longer in the core IGP routing table because they are not in the VRF routing tables.

```
R8#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      192.168.1.0/24 is variably subnetted, 5 subnets, 2 masks
L1      192.168.1.97/32 [115/20] via 192.168.1.66, Serial0/0.111
i L1    192.168.1.112/28 [115/20] via 192.168.1.66, Serial1/0.111
C       192.168.1.64/28 is directly connected, Serial1/0
C       192.168.1.33/32 is directly connected, Loopback0
i L1    192.168.1.48/28 [115/30] via 192.168.1.66, Serial1/0.111
i L1    192.168.1.17/32 [115/40] via 192.168.1.66, Serial1/0.111
```

- On R8, use the **sh ip rou vrf Yellow** command to verify the routing table.

Note Verify that the VRF routing table contains the routes for local (connected) and remote (BGP) routes for interfaces that support the customer carrier.

```
R8#sh ip rou vrf Yellow
```

Routing Table: Yellow

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```
      150.1.0.0/28 is subnetted, 2 subnets
C       150.1.12.16 is directly connected, Serial0/0
B       150.1.11.16 [200/0] via 192.168.1.17, 00:26:45
```

Task 3: Configure the CE Router of the Customer Carrier

The following steps explain how to configure the routers (R5 and R6) in the customer carrier that links to the edge router of the backbone carrier. These steps involve the establishment of an IBGP peer session between the CE routers that will carry the customer routes.

Exercise Procedure

Complete the following commands:

- Step 1** On your assigned router (R5 and R6, respectively), enable routing between carriers Yellow and Blue by placing the appropriate network statement into the OSPF process.
- Step 2** On your assigned router (R5 and R6, respectively), create an IBGP peering session between the loopback addresses of the routers of carrier Yellow.

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

Team A

- On R1, use the **sh ip bgp** command to verify the status of the end-to-end connections.

Note Verify that the local and remote routes (BGP) for customer Green are in the BGP table. Verify that none of the other routes are in the routing table.

```
R1#sh ip bgp
BGP table version is 5, local router ID is 172.12.1.16
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 172.11.1.0/24	0.0.0.0	0		32768	i
*>i172.11.2.0/24	10.2.11.49	0	100	0	i
*> 172.12.1.0/24	0.0.0.0	0		32768	i
*>i172.12.2.0/24	10.2.11.49	0	100	0	i

- On R1, perform a traceroute to the loopback address of R2. Note that while the packet is in the network of carrier Yellow, it is a native IP packet. In the network of carrier Blue, it is an MPLS VPN packet (two labels until the PHP takes place).

```
R1#trace 10.2.11.49
```

```
Type escape sequence to abort.
Tracing the route to 10.2.11.49
```

```
 1 10.1.11.130 60 msec 60 msec 68 msec
 2 150.1.11.18 60 msec 72 msec 68 msec
 3 192.168.1.50 [MPLS: Labels 19/23 Exp 0] 120 msec 200 msec 160
msec
 4 192.168.1.114 [MPLS: Labels 19/23 Exp 0] 120 msec 112 msec 60
msec
 5 150.1.12.18 [MPLS: Label 23 Exp 0] 72 msec 68 msec 68 msec
 6 150.1.12.17 68 msec 52 msec 48 msec
```

7 10.2.11.129 48 msec 52 msec *

Team B

- On R2, use the **sh ip bgp** command to verify the status of the end-to-end connections.

Note Verify that the local and remote routes (BGP) for customer Green are in the BGP table.
 Verify that none of the other routes are in the routing table.

```
R2#sh ip bgp
BGP table version is 3, local router ID is 172.12.2.16
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* i172.11.1.0/24	10.1.11.49	0	100	0	i
*> 172.11.2.0/24	0.0.0.0	0		32768	i
* i172.12.1.0/24	10.1.11.49	0	100	0	i
*> 172.12.2.0/24	0.0.0.0	0		32768	i

- On R2, perform a traceroute to the loopback address of R1. Note that while the packet is in the network of carrier Yellow, it is a native IP packet. In the network of carrier Blue, it is an MPLS VPN packet (two labels until the PHP takes place).

```
R2#trace 10.1.11.49
```

```
Type escape sequence to abort.
Tracing the route to 10.1.11.49
```


```
 1 10.2.11.130 68 msec 60 msec 48 msec
 2 150.1.12.18 52 msec 92 msec 100 msec
 3 192.168.1.66 [MPLS: Labels 18/22 Exp 0] 88 msec 108 msec 180
msec
 4 192.168.1.113 [MPLS: Labels 16/22 Exp 0] 188 msec 120 msec 148
msec
 5 150.1.11.18 [MPLS: Label 22 Exp 0] 60 msec 72 msec 52 msec
 6 150.1.11.17 68 msec 112 msec 48 msec
 7 10.1.11.129 52 msec 72 msec *
```

Lab Exercise 2-2: Implementing CSC Using LDP Label Distribution

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, carrier Yellow has decided to convert its network to an MPLS core. This conversion will allow LDP to be used to distribute established connectivity.



For this lab exercise, your workgroup will be divided into two teams.


Team A will be responsible for routers R1, R5, R7, and R3.

Team B will be responsible for routers R2, R6, R8, and R4.

After completing this exercise, you will be able to meet these objectives:

- Configure the PE routers of the backbone carrier (R7 and R8)
- Configure the CE routers of the customer carrier (R1, R5, R2, and R6)
- Disable unneeded BGP processes


Visual Objective



For this exercise, there are four distinctive IP address ranges. The 192.168.0.0 range is used to provide addressing for the core MPLS network of carrier Blue. The 150.0.0.0 range is used to provide addressing for the links between carrier Blue and the customer carrier Yellow. The 10.0.0.0 is used to provide addressing for the core of carrier Yellow. The 172.11.0.0 and 172.12.0.0 ranges are used provide network addressing for customers Green and Orange.

Note The addresses of customers Green and Orange are simulated using loopback addresses.

Lab 2-2 Logical Connections



Customer carrier Yellow wants to run an MPLS network. OSPF will be used as the IGP, and BGP will be used to carry the customer routes. As this lab starts, OSPF and BGP have been enabled. IBGP sessions have been established within the network of carrier Yellow.

Note To simplify the lab, the EBGp sessions with the sites of customers Green and Orange have been simulated with network statements in the BGP processes.

Carrier Blue is running an MPLS VPN network. IS-IS is being used as the IGP, and MP-BGP is being used to carry VPN labels. As this lab starts, IS-IS has been enabled, and all MPLS and MP-BGP sessions have been activated.

During this lab you will configure LDP support on the routers of carrier Blue and IGP, LDP, and BGP connectivity on the routers of customer carrier Yellow.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Command	Description
<code>ip vrf vrf-name</code> <code>no ip vrf vrf-name</code>	To configure a VRF routing table, use the ip vrf command in global configuration mode or router configuration mode. To remove a VRF routing table, use the no form of this command.
<code>mpls ip</code> <code>no mpls ip</code>	To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the mpls ip command in interface configuration mode. To disable MPLS forwarding of IPv4 packets, use the no form of the command.
<code>mpls label protocol {ldp tdp both}</code> <code>no mpls label protocol</code>	To specify the label distribution protocol to be used on a given interface, use the mpls label protocol command in interface configuration mode. To remove the label distribution protocol, use the no form of this command.
<code>neighbor {ip-address / peer-group-name ipv6- address} activate</code> <code>no neighbor {ip-address / peer-group-name ipv6- address} activate</code>	To enable the exchange of information with a BGP neighbor, use the neighbor activate command in address family configuration mode or router configuration mode. To disable the exchange of an address with a BGP neighbor, use the no form of this command.
<code>router bgp as-number</code> <code>no router bgp as-number</code>	To configure the BGP routing process, use the router bgp command in global configuration mode. To remove a routing process, use the no form of this command.

Task 1: Configure the PE Routers of the Backbone Carrier (R7 and R8)

The backbone of Carrier blue was configured to support a native IP network for the last lab. To prepare it to support an MPLS network, you will need to make a few changes to the routers that support the network of carrier Yellow. In this task you will configure the backbone routers. In the next task you will configure the CE routers. You will be able to use the VRF table that you created in the last lab, but you will need to enable the interfaces that support the network of carrier Yellow.

Exercise Procedure

Complete the following commands:

- Step 1** On your assigned router (R7 and R8, respectively), enable MPLS and LDP in the interface that is connected to the carrier Yellow.

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

On R7 and R8 use the **sh mpls int** command to verify the status of the MPLS interfaces.

Example:

```
R7#sh mpls int
Interface          IP          Tunnel    Operational
Serial0/0.111      Yes (tdp)   No        Yes
```

- On R7 and R8 use the **sh ip vrf int** command to verify the status of the VRF interfaces.

Example:

```
R7#sh ip vrf int
Interface          IP-Address    VRF
Protocol
Serial0/0.102      150.1.11.18   Yellow
Loopback100        10.1.13.49    Yellow
```

up
up

Task 2: Configure the CE Routers of the Customer Carriers (R1, R5, R2, and R6)

To prepare it to support an MPLS network, you will need to make a few changes to the routers that support the network of carrier Yellow. In this section you will complete the LDP configuration.

Exercise Procedure

Complete the following commands:

- Step 1** On your assigned routers (R1 or R2, respectively), configure loopback 0 as the tag-switching router ID.

Note Depending on the Cisco IOS version and hardware configuration, it may be necessary to reload the router to activate this command.

- Step 2** On your assigned routers (R1 or R2, respectively), enable MPLS and TDP on the interfaces that are connected to the remote router (R5 or R6, respectively) .

- Step 3** On your assigned routers (R5 or R6, respectively), enable MPLS and TDP on the interfaces that are connected to the remote router (R1 or R2, respectively).

- Step 4** On your assigned routers (R5 or R6, respectively), enable MPLS and LDP on the interfaces that are connected to the core router (R7 or R8, respectively).

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

- On all customer carrier routers use the **sh mpls int** command to verify the status of the MPLS interfaces.

Example

On R1 and R2 use the **sh mpls int** command to verify the status of the MPLS interfaces.

```
R1#sh mpls int
Interface          IP      Tunnel  Operational
Serial0/0.101      Yes     No       Yes
```

```
R2#sh mpls int
Interface          IP      Tunnel  Operational
Serial0/0.101      Yes     No       Yes
```

Task 3: Change the BGP Sessions of the Customer Carrier

Because the customer carrier is now running MPLS in its core, BGP is needed only on the ingress routers of the customer carrier.

Exercise Procedure

Complete the following commands:

- Step 1** On your assigned routers (R1 or R2, respectively), remove the BGP neighbor peering session with the local router (R5 or R6, respectively).
- Step 2** On your assigned routers (R1 or R2, respectively), enable the BGP neighbor peering session with the remote router (R2 or R1, respectively).
- Step 3** On you assigned routers (R5 or R6, respectively), remove the BGP neighbor peering session with the local router (R1 or R2, respectively).

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

- On your assigned routers (R1 or R2, respectively), use the **sh ip bgp** command to verify the status of the existing end-to-end connections.

Example:

```
R1# 1#sh ip bgp sum
BGP router identifier 172.12.1.16, local AS number 65010
BGP table version is 7, main routing table version 7
2 network entries and 2 paths using 266 bytes of memory
1 BGP path attribute entries using 60 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP activity 4/38 prefixes, 4/2 paths, scan interval 60 secs
```


Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
State/PfxRcd									
10.2.11.49	4	65010	56	56	5	0	0	00:51:36	2

Lab Exercise 2-3: Implementing CSC Using MP-BGP Label Distribution

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, carrier Blue has decided to convert the label exchange from LDP to BGP on the interface that is supporting carrier Yellow.




For this lab exercise, your workgroup will be divided into two teams.

- Team A will be responsible for routers R5 and R7.
- Team B will be responsible for routers R6 and R8.

After completing this exercise, you will be able to meet this objective:


- Configure MP-BGP label distribution

Visual Objective



For this exercise, there are four distinctive IP address ranges. The 192.168.0.0 range is used to provide addressing for the core MPLS network of carrier Blue. The 150.0.0.0 range is used to provide addressing for the links between carrier Blue and customer carrier Yellow. The 10.0.0.0 is used to provide addressing for the core of carrier Yellow. The 172.11.0.0 and 172.12.0.0 ranges are used to provide network addressing for customers Green and Orange.

Note The addresses of customers Green and Orange are simulated using loopback addresses.



The basic implementation in this lab is very similar to the last lab; the implementation differs only in the protocol that is being used to exchange labels between the carriers Yellow and Blue.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Command	Description
<pre>address-family ipv4 [multicast unicast vrf vrf-name] no address-family ipv4 [multicast unicast vrf vrf-name]</pre>	To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the address-family ipv4 command in router configuration mode. To disable address family configuration mode, use the no form of this command.
<pre>mpls ip no mpls ip</pre>	To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the mpls ip command in interface configuration mode. To disable MPLS forwarding of IPv4 packets, use the no form of the command.
<pre>mpls label protocol {ldp tdp both} no mpls label protocol</pre>	To specify the label distribution protocol to be used on a given interface, use the mpls label protocol command in interface configuration mode. To remove the label distribution protocol, use the no form of this command.

Command	Description
<pre>neighbor {ip-address peer-group-name} remote- as as-number no neighbor {ip-address peer-group-name} remote-as as-number</pre>	To add an entry to the BGP or MP-BGP neighbor table, use the neighbor remote-as command in router configuration mode. To remove an entry from the table, use the no form of this command.
<pre>neighbor {ip-address / peer-group-name ipv6- address} activate no neighbor {ip-address / peer-group-name ipv6- address} activate</pre>	To enable the exchange of information with a BGP neighbor, use the neighbor activate command in address family configuration mode or router configuration mode. To disable the exchange of an address with a BGP neighbor, use the no form of this command.
<pre>neighbor ip-address as- override no neighbor ip-address as-override</pre>	To configure a PE router to override the AS number of a site with the AS number of a provider, use the neighbor as-override command in router configuration mode. To remove VPN IPv4 prefixes from a specified router, use the no form of this command.
<pre>neighbor ip-address send- label no neighbor ip-address send-label</pre>	To enable a BGP router to send MPLS labels with BGP routes to a neighboring BGP router, use the neighbor send-label command in address family configuration mode or router configuration mode. To disable this feature, use the no form of this command.
<pre>redistribute protocol [process-id] {level-1 level-1-2 level-2} [as- number] [metric metric- value] [metric-type type-value] [match {internal external 1 external 2}] no redistribute protocol [process-id] {level-1 level-1-2 level-2} [as- number] [metric metric- value] [metric-type type-value] [match {internal external 1 external 2}]</pre>	To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. To disable redistribution, use the no form of this command.
<pre>router bgp as-number no router bgp as-number</pre>	To configure the BGP routing process, use the router bgp command in global configuration mode. To remove a routing process, use the no form of this command.
<pre>router ospf process-id no router ospf process-id</pre>	To configure an OSPF routing process, use the router ospf command in global configuration mode. To terminate an OSPF routing process, use the no form of this command.

Task 1: Configure MP-BGP Label Distribution

The label distribution protocol on R5, R6, R7, and R8 must be converted from LDP to MP-BGP. The first step is to remove the MPLS, LDP, and OSPF redistribution that is currently used to distribute routes and labels. Then an EBGP session can be established to distribute routes and labels.

Exercise Procedure

Complete the following commands on your assigned routers:

R7 and R8

- Step 1** Disable LDP and MPLS on the interface that is connected to the carrier Yellow.
- Step 2** Remove the OSPF process that is supporting the carrier Yellow.
- Step 3** In the address family that is supporting carrier Yellow, create an EBGP peer session with the serial interface of the remote router (R5 or R6). Do not forget to activate the session.
- Step 4** In the address family that is supporting carrier Yellow, enable the BGP AS-override feature.
- Step 5** In the address family that is supporting carrier Yellow, enable the BGP **send-label** option.

R5 and R6

- Step 6** Disable LDP and MPLS on the interface that is connected to the carrier Blue.
- Step 7** To get the IGP routes into BGP, use the redistribute connected command in OSPF. Use a metric of 4 and the **match internal** option.
- Step 8** Create an EBGP peer session with the serial interface of the remote router (R7 or R8).
- Step 9** Enable the BGP **send-label** option.
- Step 10** To get the BGP routes into the IGP, redistribute BGP. Use a metric of 4 and do not forget subnets.

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

Team A

- On R1, use the **sh ip bgp** command to verify the status of the existing end-to-end connections.

Example

```
R1#sh ip bgp
BGP table version is 5, local router ID is 172.12.1.16
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```


Network	Next Hop	Metric	LocPrf	Weight	Path
*> 172.11.1.0/24	0.0.0.0	0		32768	i
*>i172.11.2.0/24	10.2.11.49	0	100	0	i
*> 172.12.1.0/24	0.0.0.0	0		32768	i
*>i172.12.2.0/24	10.2.11.49	0	100	0	i

Lab Exercise 2-4: Implementing CSC Using MPLS VPNs

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this exercise, the carrier Blue has decided to convert its network to provide MPLS VPN support.




For this lab exercise, your workgroup will be divided into two teams.

- Team A will be responsible for router R1.
- Team B will be responsible for router R2.

After completing this exercise, you will be able to meet this objective:

- Configure MPLS VPN support


Visual Objective



For this exercise, there are four distinctive IP address ranges. The 192.168.0.0 range is used to provide addressing for the core MPLS network of carrier Blue. The 150.0.0.0 range is used to provide addressing for the links between carrier Blue and customer carrier Yellow. The 10.0.0.0 is used to provide addressing for the core of carrier Yellow. The 172.11.0.0 and 172.12.0.0 range are used provide network addressing for customers Green and Orange.

Note The addresses of customers Green and Orange are simulated using loopback addresses.

Lab 2-4 Logical Connections



The basic implementation in this lab is very similar to the last lab; the implementation differs only in that now the two customers will be in separate VPNs.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands used in this exercise are described in the table here.

Command	Description
<pre>address-family ipv4 [multicast unicast vrf vrf-name] no address-family ipv4 [multicast unicast vrf vrf-name]</pre>	To enter address family configuration mode for configuring routing sessions, such as BGP, that use standard IPv4 address prefixes, use the address-family ipv4 command in router configuration mode. To disable address family configuration mode, use the no form of this command.
<pre>ip vrf forwarding vrf- name no ip vrf forwarding vrf- name</pre>	To associate a VRF table with an interface or subinterface, use the ip vrf forwarding command in interface configuration mode. To disassociate a VRF table, use the no form of this command.
<pre>ip vrf vrf-name no ip vrf vrf-name</pre>	To configure a VRF routing table, use the ip vrf command in global configuration mode or router configuration mode. To remove a VRF routing table, use the no form of this command.
<pre>mpls ip no mpls ip</pre>	To enable MPLS forwarding of IPv4 packets along normally routed paths for a particular interface, use the mpls ip command in interface configuration mode. To disable MPLS forwarding of IPv4 packets, use the no form of the command.

Command	Description
mpls label protocol {ldp tdp both} no mpls label protocol {ldp tdp both}	To specify the label distribution protocol to be used on a given interface, use the mpls label protocol command in interface configuration mode. To remove the label distribution protocol, use the no form of this command.
neighbor ip-address [priority number] [poll-interval seconds] [cost number] [database-filter all] no neighbor ip-address [priority number] [poll-interval seconds] [cost number] [database-filter all]	To configure OSPF routers that are interconnecting to nonbroadcast networks, use the neighbor command in router configuration mode. To remove a configuration, use the no form of this command.
rd route-distinguisher	To create routing and forwarding tables for a VRF instance, use the rd command in VRF configuration mode.
redistribute protocol [process-id] {level-1 level-1-2 level-2} [as-number] [metric metric-value] [metric-type type-value] [match {internal external 1 external 2}] no redistribute protocol [process-id] {level-1 level-1-2 level-2} [as-number] [metric metric-value] [metric-type type-value] [match {internal external 1 external 2}]	To redistribute routes from one routing domain into another routing domain, use the redistribute command in router configuration mode. To disable redistribution, use the no form of this command.
router bgp as-number no router bgp as-number	To configure the BGP routing process, use the router bgp command in global configuration mode. To remove a routing process, use the no form of this command.
router ospf process-id no router ospf process-id	To configure an OSPF routing process, use the router ospf command in global configuration mode. To terminate an OSPF routing process, use the no form of this command.
route-target {import export both} route-target-ext-community no route-target {import export both} route-target-ext-community	To create an RT extended community for a VRF table, use the route-target command in VRF configuration mode. To disable the configuration of an RT community option, use the no form of this command.

Task 1: Configure MPLS VPN Support

The underlying infrastructure that is needed to support the VPNs of carrier Yellow has already been established in the previous labs. In this lab, the actual VPNs are enabled.

Exercise Procedure

Complete the following commands on your assigned routers (R1 or R2):

- Step 1** Create a VRF table to support the customer Green using the following information:
- VRF name: Green
 - RD: 100:1
 - Import and export router targets: 100:1
- Step 2** Place the loopback address (loopback11) into the VRF table of customer Green.
- Step 3** Create a VRF table to support the customer Orange using the following information:
- VRF name: Orange
 - RD: 100:2
 - Import and export router targets: 100:2
- Step 4** Place the loopback address (loopback12) into the VRF table of customer Orange.
- Step 5** In the BGP process (BGP 65010), remove the network statement for the network of customer Green.
- Step 6** In the BGP process (BGP 65010), remove the network statement for the network of customer Orange.
- Step 7** In the BGP process (BGP 65010), create a VPN IPv4 (VPNv4) peering with the remote router.
- Step 8** In the BGP process (BGP 65010), create an address family to support the VRF table of customer Green and add a network statement for the network of customer Green.
- Step 9** In the BGP process (BGP 65010), create an address family to support the VRF table of customer Orange. and add a network statement for the network of customer Qrange.

Note In a standard implementation it would be necessary to place the associated physical or logical interface into the VRF table and enable a routing protocol between the CE and PE routes. However, because the end customer is simulated in this lab, the CE-PE routing protocol is not needed.

Exercise Verification

The exercise verification should be done when both teams have completed their configuration tasks. You have completed this exercise when you attain these results:

Team A

- On R1, use the **sh ip rou vrf Green** command to verify the routing table.

```
R1#sh ip rou vrf Green
```

```
Routing Table: Green
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
```

```
area
```

```
* - candidate default, U - per-user static route, o - ODR
```

```
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
172.11.0.0/24 is subnetted, 2 subnets
```

```
B 172.11.2.0 [200/0] via 10.2.11.49, 00:09:29
```

```
C 172.11.1.0 is directly connected, Loopback11
```

- On R1, use the **sh ip rou vrf Orange** command to verify the routing table.

```
R1#sh ip rou vrf Orange
```

```
Routing Table: Orange
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
```

```
area
```

```
* - candidate default, U - per-user static route, o - ODR
```

```
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
172.12.0.0/24 is subnetted, 2 subnets
```

```
C 172.12.1.0 is directly connected, Loopback12
```

```
B 172.12.2.0 [200/0] via 10.2.11.49, 00:10:47
```

Team B

- On R2, use the **sh ip rou vrf Green** command to verify the routing table.

```
R2#sh ip rou vrf Green
```

```
Routing Table: Green
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
```

```
area
```

```
* - candidate default, U - per-user static route, o - ODR
```

```
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
172.11.0.0/24 is subnetted, 2 subnets
```

```
C 172.11.2.0 is directly connected, Loopback11
```

```
B          172.11.1.0 [200/0] via 10.1.11.49, 01:54:18
```

- On R2, use the **sh ip rou vrf Orange** command to verify the routing table.

```
R2#sh ip rou vrf Orange
```

```
Routing Table: Orange
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
```

```
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter  
area
```

```
* - candidate default, U - per-user static route, o - ODR
```

```
P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
172.12.0.0/24 is subnetted, 2 subnets
```

```
B          172.12.1.0 [200/0] via 10.1.11.49, 01:55:15
```

```
C          172.12.2.0 is directly connected, Loopback12
```

Next Step


- Module 3: MPLS Traffic Engineering Technology

Lab Exercise 4-1: Implementing MPLS Traffic Engineering

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In this lab carrier Yellow has a network that consists of five routers: R5, R6, R7, R8, and R9. R7, R8, and R9 have been configured to provide an MPLS core. This core supports two PE routers (R5 and R6). The PE routers are located at separate POP sites: POP 1 and POP 2, respectively. The PE routers have been configured to provide MPLS VPN support to the two customers of carrier Yellow: customers Green and Orange.




This service provider has decided to implement MPLS-TE.

In this exercise, you will establish MPLS traffic tunnels to support the needs of the service provider. After completing this exercise, you will be able to meet these objectives:

- Log existing traffic flow
- Configure MPLS-TE support on your routers
- Configure MPLS-TE tunnels


Visual Objective



For accounting reasons, the service provider has decided to route all customer traffic through R9. The local link between the PE routers must be retained for backup purposes. There are several methods to implement this scheme: static routing, metric adjustment, Border Gateway Protocol (BGP) attributes, and so on. However, all of these methods suffer from some type of scaling issue. The service provider has decided to implement MPLS-TE to overcome these issues.

Lab 4-1: IP Addressing

Cisco.com



© 2004 Cisco Systems, Inc. All rights reserved.

MPLST v2.0-17

For this exercise, there are three distinctive IP address ranges. The 192.168.0.0 range is used to provide addressing for the core MPLS network of carrier Yellow. The 150.0.0.0 range is used to provide addressing for the links between carrier Yellow and customers Green and Orange. The 10.1.0.0 and 10.2.0.0 are used provide network addressing for the networks of customers Green and Orange.

Note Because both workgroups share R9, it is necessary for you to substitute your workgroup number in the appropriate places in the IP addressing (for example, in 192.168.wg.81/32).

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<code>interface Tunnel0</code>	Creates a tunnel interface
<code>ip rsvp bandwidth max-reservable single-request</code>	Enables RSVP on an interface
<code>metric-style wide</code>	Specifies the IS-IS metric type
<code>mpls traffic-eng level</code>	Configures IS-IS TE support
<code>mpls traffic-eng logging</code>	Configures MPLS-TE logging

Command	Description
<code>mpls traffic-eng reoptimize timers frequency <i>seconds</i></code>	Defines MPLS-TE reoptimization frequency
<code>mpls traffic-eng router-id <i>interface</i></code>	Defines the IP address that is used as the tunnel destination for this router
<code>mpls traffic-eng tunnels</code>	Enables MPLS-TE globally or on an interface
<code>router isis</code>	Starts IS-IS configuration
<code>show mpls traffic-eng autoroute</code>	Displays autoroute information
<code>show mpls traffic-eng link-management admission-control</code>	Displays tunnels that have been admitted through this MPLS-TE node
<code>show mpls traffic-eng link-management advertisements</code>	Displays advertisement information that is sent from this node into the IS-IS routing process
<code>show mpls traffic-eng link-management igp-neighbors</code>	Displays IS-IS neighbors that support MPLS-TE
<code>show mpls traffic-eng link-management summary</code>	Displays overall MPLS-TE node characteristics
<code>show mpls traffic-eng tunnels</code>	Displays detailed tunnel information
<code>show mpls traffic-eng tunnels brief</code>	Displays the overall status of MPLS-TE tunnels
<code>tag-switching ip</code>	Starts MPLS support and Tag Distribution Protocol (TDP) on an interface
<code>tunnel destination <i>ip-address</i></code>	Specifies the tunnel tailend router
<code>tunnel mode mpls traffic-eng</code>	Specifies tunnel encapsulation mode
<code>tunnel mpls traffic-eng autoroute announce</code>	Enables autoroute on an MPLS-TE tunnel
<code>tunnel mpls traffic-eng bandwidth <i>bandwidth</i></code>	Specifies the bandwidth that is reserved by the MPLS-TE tunnel
<code>tunnel mpls traffic-eng path-option 1 dynamic</code>	Specifies MPLS-TE path options
<code>tunnel mpls traffic-eng priority <i>setup hold</i></code>	Specifies MPLS-TE tunnel setup and hold priority

Task 1: Implement the Initial Configurations on the Routers

As this lab starts, carrier Yellow has established an IGP. Because a link-state protocol is required to support TE, the service provider has decided to implement IS-IS. Also, all MPLS and Multiprotocol Border Gateway Protocol (MP-BGP) sessions need to support the customer MPLS VPNs that have been established. To establish this environment, configure all routers as indicated in the following procedure.

Exercise Procedure

Complete the following commands:

Step 1 On each router erase the start configuration.

Step 2 Reload each router without saving the running configuration.

Note In the following configurations it will be necessary to replace the “wg” with the number of your workgroup.

Team A

Step 1 On R1 install the following configuration:

```
hostname R1

interface Loopback0
 ip address 10.1.1.49 255.255.255.255

interface Loopback10
 ip address 10.1.1.17 255.255.255.240

interface Loopback20
 ip address 10.1.1.65 255.255.255.240

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R5
 ip address 150.wg.1.17 255.255.255.240
 frame-relay interface-dlci 101
 no shut

router ospf 1
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.1.1.49 0.0.0.0 area 0
 network 10.1.1.16 0.0.0.15 area 1
 network 10.1.wg1.64 0.0.0.15 area 1
 redistribute bgp 65011 subnets
 exit

router bgp 65011
 nei 150.wg.1.18 remote-as 65001
```

```
no auto-summary
redistribute ospf 1
exit
```

```
line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login
```

Step 2 On R2 install the following configuration:

```
hostname R2

interface Loopback0
ip address 10.2.1.49 255.255.255.255

interface Loopback10
ip address 10.2.1.17 255.255.255.240

interface Loopback20
ip address 10.2.1.65 255.255.255.240

interface Serial0/0
no ip address
bandwidth 128000
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
no shut

interface Serial0/0.102 point-to-point
desc link to R5
ip address 150.wg.1.33 255.255.255.240
frame-relay interface-dlci 102

router ospf 2
network 150.wg.0.0 0.0.255.255 area 0
network 10.2.1.49 0.0.0.0 area 0
network 10.2.1.16 0.0.0.15 area 1
redistribute bgp 65012 subnets
exit

router bgp 65012
nei 150.wg.1.34 remote-as 65001
no auto-summary
redistribute ospf 2

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login
```

Step 3 On R5 install the following configuration:

```
hostname R5

hostname R5

ip cef

ip vrf Green
  rd wg:10
  route-target both wg:10
exit

ip vrf Orange
  rd wg:20
  route-target both wg:20
exit

interface Loopback0
  ip address 192.168.wg.17 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  ip vrf forwarding Green
  ip address 150.wg.1.18 255.255.255.240
  frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
  ip vrf forwarding Orange
  ip address 150.wg.1.34 255.255.255.240
  frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
  ip address 192.168.wg.49 255.255.255.240
  ip router isis
  mpls label protocol ldp
  mpls ip
  frame-relay interface-dlci 111

router isis
  net 39.0440.0000.0000.wg500
  metric-style wide
exit

router bgp 65001
  no synchronization
  neighbor 192.168.wg.33 remote-as 65001
  neighbor 192.168.wg.33 update-source Loopback0
  no auto-summary
  address-family vpnv4
    neighbor 192.168.wg.33 activate
    neighbor 192.168.wg.33 next-hop-self
    neighbor 192.168.wg.33 send-community both
```

```

        no auto-summary
    exit

    address-family ipv4 vrf Orange
        nei 150.wg.1.33 remote-as 65012
        nei 150.wg.1.33 activate
        nei 150.wg.1.33 as-override
        no auto-summary
        no synchronization
    exit

    address-family ipv4 vrf Green
        nei 150.wg.1.17 remote-as 65011
        nei 150.wg.1.17 activate
        nei 150.wg.1.17 as-override
        no auto-summary
    exit

exit

line con 0
    exec-timeout 999 0
    logging synchronous
line aux 0
line vty 0 4
    no login

```

Step 4 On R7 install the following configuration:

```

hostname R7

ip cef

interface Loopback0
    ip address 192.168.wg.81 255.255.255.255
    ip router isis

interface Serial0/0
    no ip address
    bandwidth 128000
    encapsulation frame-relay
    no fair-queue
    frame-relay lmi-type cisco
    no shut

interface Serial0/0.111 point-to-point
    ip address 192.168.wg.50 255.255.255.240
    ip router isis
    mpls label protocol ldp
    mpls ip
    frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
    ip router isis
    ip address 192.168.wg.113 255.255.255.240
    mpls label protocol ldp
    mpls ip
    frame-relay interface-dlci 112

int s0/0.2wg1 point-to-point

```

```

ip address 192.168.wg.129 255.255.255.248
ip router isis
mpls ip
mpls label protocol ldp
frame-relay interface-dlci 2wg1
no shut

router isis
net 39.0440.0000.0000.wg700
metric-style wide

exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Team B

Step 1 On R3 install the following configuration:

```

hostname R3

interface Loopback0
  ip address 10.1.2.49 255.255.255.255

interface Loopback10
  ip address 10.1.2.17 255.255.255.240

interface Loopback20
  ip address 10.1.2.65 255.255.255.240

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  no desc link to R8
  ip address 150.wg.2.17 255.255.255.240
  frame-relay interface-dlci 101
  no shut

router ospf 1
  network 150.wg.0.0 0.0.255.255 area 0
  network 10.1.2.49 0.0.0.0 area 0
  network 10.1.2.16 0.0.0.15 area 1
  redistribute bgp 65011 subnets
exit

router bgp 65011
  nei 150.wg.2.18 remote-as 65001
  no auto-summary
  redistribute ospf 1
exit

```

```

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Step 2 On R4 install the following configuration:

```

hostname R4

interface Loopback0
  ip address 10.2.2.49 255.255.255.255

interface Loopback10
  ip address 10.2.2.17 255.255.255.240

interface Loopback20
  ip address 10.2.2.65 255.255.255.240

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.102 point-to-point
  desc link to R5
  ip address 150.wg.2.33 255.255.255.240
  frame-relay interface-dlci 102

router ospf 2
  network 150.wg.0.0 0.0.255.255 area 0
  network 10.2.2.49 0.0.0.0 area 0
  network 10.2.2.16 0.0.0.15 area 1
  redistribute bgp 65012 subnets
exit

router bgp 65012
  nei 150.wg.2.34 remote-as 65001
  no auto-summary
  redistribute ospf 2
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```


Step 3 On R6 install the following configuration:

```
hostname R6

ip cef

ip vrf Green
  rd wg:10
  route-target both wg:10
exit

ip vrf Orange
  rd wg:20
  route-target both wg:20
exit

interface Loopback0
  ip address 192.168.wg.33 255.255.255.255
  ip router isis
  !

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  ip vrf forwarding Green
  ip address 150.wg.2.18 255.255.255.240
  frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
  ip vrf forwarding Orange
  ip address 150.wg.2.34 255.255.255.240
  frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
  ip address 192.168.wg.65 255.255.255.240
  ip router isis
  mpls label protocol ldp
  mpls ip
  frame-relay interface-dlci 111

interface FastEthernet0/1
  no ip address
  shutdown
  duplex auto
  speed auto
  no cdp enable

router isis
  net 39.0440.0000.0000.wg600
  metric-style wide
exit

router bgp 65001
```

```

no synchronization
neighbor 192.168.wg.17 remote-as 65001
neighbor 192.168.wg.17 update-source Loopback0
no auto-summary

address-family vpnv4
neighbor 192.168.wg.17 activate
neighbor 192.168.wg.17 next-hop-self
neighbor 192.168.wg.17 send-community both
no auto-summary
exit-address-family

address-family ipv4 vrf Orange
neighbor 150.wg.2.33 remote-as 65012
neighbor 150.wg.2.33 activate
neighbor 150.wg.2.33 as-override
no auto-summary
no synchronization
exit-address-family

address-family ipv4 vrf Green
neighbor 150.wg.2.17 remote-as 65011
neighbor 150.wg.2.17 activate
neighbor 150.wg.2.17 as-override
no auto-summary
no synchronization
exit-address-family
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Step 4 On R8 install the following configuration:

```

hostname R8

ip cef

interface Loopback0
  ip address 192.168.wg.97 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
no shut

interface Serial0/0.111 point-to-point
  ip address 192.168.wg.66 255.255.255.240
  ip router isis
  mpls label protocol ldp
  mpls ip
  frame-relay interface-dlci 111

```

```

interface Serial0/0.112 point-to-point
 ip address 192.168.wg.114 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 112

int s0/0.2wg2 point-to-point
 ip address 192.168.wg.145 255.255.255.248
 ip router isis
 mpls ip
 mpls label protocol ldp
 frame-relay interface-dlci 2wg2
 no shut

router isis
 net 39.0440.0000.0000.wg800
 metric-style wide

exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login

```

Step 5 On R9 install the following configuration:

Note	R9 is shared by both workgroups. Only one workgroup needs to add its configuration.

```

hostname R9

enable password slpm

ip cef
 mpls traffic-eng tunnels
 mpls traffic-eng logging lsp setups
 mpls traffic-eng reoptimize timers frequency 30

interface Loopback0
 ip address 192.168.10.17 255.255.255.255
 ip router isis

interface Serial0/0
 no ip address
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 ip rsvp bandwidth 128 128
 no shut

interface Serial0/0.211 point-to-point
 desc link to R7 in wg 1
 ip address 192.168.1.130 255.255.255.248
 ip router isis
 mpls label protocol ldp

```

```

tag-switching ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 211

interface Serial0/0.212 point-to-point
desc link to R8 in wg 1
ip address 192.168.1.146 255.255.255.248
ip router isis
mpls label protocol ldp
tag-switching ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 212

interface Serial0/0.221 point-to-point
desc link to R7 in wg 2
ip address 192.168.2.130 255.255.255.248
ip router isis
mpls label protocol ldp
tag-switching ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 221

interface Serial0/0.222 point-to-point
desc link to R8 in wg 2
ip address 192.168.2.146 255.255.255.248
ip router isis
mpls label protocol ldp
tag-switching ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 222

router isis
net 49.0001.0000.0000.0009.00
metric-style wide
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-2
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

Exercise Verification

You have completed this exercise when you attain these results:

Team A

- On R1, use the **sh ip rou** command to verify the routing table.

```
R1#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter
area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B       10.1.2.16/28 [20/0] via 150.1.1.18, 00:00:07
C       10.1.1.16/28 is directly connected, Loopback10
B       10.1.2.49/32 [20/0] via 150.1.1.18, 00:00:07
C       10.1.1.49/32 is directly connected, Loopback0
150.1.0.0/28 is subnetted, 2 subnets
B       150.1.2.16 [20/0] via 150.1.1.18, 00:00:07
C       150.1.1.16 is directly connected, Serial0/0.101
```

- On R3, use the **sh ip rou** command to verify the routing table.

```
R3#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C       10.1.2.16/28 is directly connected, Loopback10
B       10.1.1.16/28 [20/0] via 150.1.2.18, 00:01:07
C       10.1.2.49/32 is directly connected, Loopback0
B       10.1.1.49/32 [20/0] via 150.1.2.18, 00:01:07
150.1.0.0/28 is subnetted, 2 subnets
C       150.1.2.16 is directly connected, Serial0/0.101
B       150.1.1.16 [20/0] via 150.1.2.18, 00:01:07
```

Team B

- On R2, use the **sh ip rou** command to verify the routing table.

```
R2#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
C    10.2.1.16/28 is directly connected, Loopback10
B    10.2.2.16/28 [20/0] via 150.1.1.34, 01:53:03
C    10.2.1.49/32 is directly connected, Loopback0
B    10.2.2.49/32 [20/0] via 150.1.1.34, 01:53:03
150.1.0.0/28 is subnetted, 2 subnets
B    150.1.2.32 [20/0] via 150.1.1.34, 01:53:03
C    150.1.1.32 is directly connected, Serial0/0.102
```

- On R4, use the **sh ip rou** command to verify the routing table.

```
R4#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

Gateway of last resort is not set

```

10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B    10.2.1.16/28 [20/0] via 150.1.2.34, 01:53:26
C    10.2.2.16/28 is directly connected, Loopback10
B    10.2.1.49/32 [20/0] via 150.1.2.34, 01:53:26
C    10.2.2.49/32 is directly connected, Loopback0
150.1.0.0/28 is subnetted, 2 subnets
C    150.1.2.32 is directly connected, Serial0/0.102
B    150.1.1.32 [20/0] via 150.1.2.34, 01:53:26
```

Task 2: Log Existing Traffic Flow

In this task, you will log the existing traffic flow pattern to establish the current path through the network.

Exercise Procedure

Complete these steps:

Step 1 Perform a trace between your two assigned customer edge routers (CE routers).

Exercise Verification

You have completed this exercise when you attain these results:

Team A

Trace path for R1 to R3

Trace path for R3 to R1

Team B

Trace path for R2 to R4

Trace path for R4 to R2

Task 3: Configure MPLS Traffic Engineering Support on Your Routers

In this task, you will enable global features that will allow TE tunnels to be built.

Exercise Procedure

Complete these steps:

Step 1 Configure global MPLS-TE support.

Step 2 Configure MPLS-TE support and reservable bandwidth on all links between the backbone routers. Allow up to 128 kbps of reservable bandwidth on any single subinterface and up to 128 kbps of reservable bandwidth on physical interfaces.

Note To make MPLS-TE work over a subinterface, you must configure RSVP on the main interface with the **ip rsvp bandwidth** command even though the main interface is not used for MPLS-TE.

Step 3 Because the network is all Level 1, configure only Level 1 IS-IS support for MPLS-TE. Use “loopback0” as your router ID.

Step 4 Configure MPLS-TE logging and a reoptimization frequency of 30 sec.

Exercise Verification

You have completed this exercise when you attain these results:

- Use **show ip rsvp interface** command to display the RVSP information that has been configured. Verify the proper interface and bandwidths have been enabled.

```
R5#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      0
Se0/0.111    0           128K     128K      0
```

```
R6#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      0
Se0/0.111    0           128K     128K      0
```

```
R7#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      0
Se0/0.112    0           128K     128K      0
Se0/0.2wg1   0           128K     128K      0
Se0/0.111    0           128K     128K      0
```

```
R8#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      0
Se0/0.111    0           128K     128K      0
Se0/0.112    0           128K     128K      0
Se0/0.2wg2   0           128K     128K      0
```


Task 4: Configure MPLS Traffic Engineering Tunnels

In this task, you will create the traffic tunnels between PE routers.

Exercise Procedure

Complete these steps:

- Step 1** On each PE router (R5 and R6), configure an MPLS-TE tunnel toward the other provider edge (PE) in your workgroup (R5 → R7 → R9 → R8 → R6). Use the following tunnel parameters:

Parameter	Value
IP address	Loopback 0
Tunnel bandwidth	100 kbps
Path options	Explicit path selection
Tunnel priority	Setup=0, hold=0
Autoroute	Enabled

- Step 2** Enable MPLS forwarding and LDP on the MPLS-TE tunnels to establish end-to-end LSPs to be used for MPLS VPN traffic.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **show mpls traffic-engineering tunnels brief** command to display tunnels that are going through a particular router and verify the operational state of the tunnels that are originating in the router. You should get a printout similar to the one here:

Team A

```
R5#sh mpls traffic-eng tunnels brief
Signaling Summary:
  LSP Tunnels Process:      running
  RSVP Process:             running
  Forwarding:               enabled
  Periodic reoptimization:  every 30 seconds, next in 11 seconds
  Periodic auto-bw collection: disabled
TUNNEL NAME                DESTINATION    UP IF      DOWN IF
STATE/PROT
R5_t1                      192.168.1.33   -          Se0/0.11
up/up
R6_t1                      192.168.1.17   Se0/0.11   -
up/up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tails
```

Team B

```
R6#sh mpls traffic-eng tunnels brief
```

```
Signaling Summary:
```

```
  LSP Tunnels Process:      running
  RSVP Process:             running
  Forwarding:                enabled
  Periodic reoptimization:   every 30 seconds, next in 2 seconds
  Periodic auto-bw collection: disabled
```

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R6_t1	192.168.1.17	-	Se0/0.11	up/up
R5_t1	192.168.1.33	Se0/0.11	-	up/up

Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tails

- Use **show ip route** to verify that the tunnel is in the global routing table.

```
R5#sh ip rou
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
192.168.10.0/32 is subnetted, 1 subnets
i L2   192.168.10.17 [115/30] via 192.168.1.50, Serial0/0.111
192.168.1.0/24 is variably subnetted, 9 subnets, 3 masks
i L1   192.168.1.97/32 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.112/28 [115/20] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.64/28 [115/30] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.81/32 [115/20] via 192.168.1.50, Serial0/0.111
S      192.168.1.33/32 is directly connected, Tunnel1
C      192.168.1.48/28 is directly connected, Serial0/0.111
C      192.168.1.17/32 is directly connected, Loopback0
i L1   192.168.1.128/29 [115/20] via 192.168.1.50, Serial0/0.111
i L1   192.168.1.144/29 [115/30] via 192.168.1.50, Serial0/0.111
```

```
R6#sh ip rou
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
192.168.10.0/32 is subnetted, 1 subnets
i L2   192.168.10.17 [115/30] via 192.168.1.66, Serial0/0.111
192.168.1.0/24 is variably subnetted, 9 subnets, 3 masks
i L1   192.168.1.97/32 [115/20] via 192.168.1.66, Serial0/0.111
i L1   192.168.1.112/28 [115/20] via 192.168.1.66, Serial0/0.111
C      192.168.1.64/28 is directly connected, Serial0/0.111
i L1   192.168.1.81/32 [115/30] via 192.168.1.66, Serial0/0.111
C      192.168.1.33/32 is directly connected, Loopback0
```

```
i L1 192.168.1.48/28 [115/30] via 192.168.1.66, Serial0/0.111
S 192.168.1.17/32 is directly connected, Tunnel1
i L1 192.168.1.128/29 [115/30] via 192.168.1.66, Serial0/0.111
i L1 192.168.1.144/29 [115/20] via 192.168.1.66, Serial0/0.111
```

- Use the **show ip cef** command to verify that the IP traffic toward the tunnel destination gets forwarded through the tunnel interface.

```
R5#sh ip cef 192.168.wg.33
192.168.wg.33/32, version 36, epoch 0, attached
0 packets, 0 bytes
tag information set, unshareable
local tag: 20
fast tag rewrite with Tu1, point2point, tags imposed: {22}
via Tunnel1, 6 dependencies
valid adjacency
tag rewrite with Tu1, point2point, tags imposed: {22}
```

```
R6#sh ip cef 192.168.wg.17
192.168.wg.17/32, version 33, epoch 0, attached
0 packets, 0 bytes
tag information set, unshareable
local tag: 23
fast tag rewrite with Tu1, point2point, tags imposed: {24}
via Tunnel1, 6 dependencies
valid adjacency
tag rewrite with Tu1, point2point, tags imposed: {24}
```

- Use the **show ip cef vrf name detail** command to verify that the MPLS VPN traffic gets forwarded through the tunnel interface.

```
R5#sh ip cef vrf Green 10.1.2.49
10.1.2.49/32, version 8, epoch 0
0 packets, 0 bytes
tag information set
local tag: VPN-route-head
fast tag rewrite with Tu1, point2point, tags imposed: {22 25}
via 192.168.1.33, 0 dependencies, recursive
next hop 192.168.1.33, Tunnel1 via 192.168.1.33/32
valid adjacency
tag rewrite with Tu1, point2point, tags imposed: {22 25}
```

```
R6#sh ip cef vrf Green 10.1.1.49
10.1.1.49/32, version 8, epoch 0
0 packets, 0 bytes
tag information set
local tag: VPN-route-head
fast tag rewrite with Tu1, point2point, tags imposed: {24 25}
via 192.168.1.17, 0 dependencies, recursive
next hop 192.168.1.17, Tunnel1 via 192.168.1.17/32
valid adjacency
tag rewrite with Tu1, point2point, tags imposed: {24 25}
```

- Repeat the traces that were done in Task 2. Verify that the trace now bypasses the direct P-router-to-P-router link and uses the R9 links.

Lab Exercise 4-2: Implementing Link Protection for MPLS Traffic Engineering

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective


If the link to R9 fails, normal IGP and MPLS processing will declare the tunnel that was created in the last lab to be down and reroute the traffic over the direct R7-to-R8 link. When the failed link recovers, the tunnel will be re-established and become active. However, the time-frame for these transfers can be unacceptable. To address this issue, carrier Yellow has decided to enable link protection with FRR.

After completing this exercise, you will be able to meet these objectives:

- Configure MPLS-TE backup tunnels

Note This feature is currently supported only on high-end routers. Therefore, this exercise will be done as a case study rather than a hands-on lab.

Visual Objective



Two unidirectional backup tunnels need to be created between R7 and R8. Then the logical interface needs to be enabled to use these tunnels.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<code>interface Tunnel0</code>	Creates a tunnel interface
<code>tag-switching ip</code>	Starts MPLS support and TDP on an interface
<code>tunnel destination ip-address</code>	Specifies the tunnel tailend router
<code>tunnel mode mpls traffic-eng</code>	Specifies tunnel encapsulation mode
<code>tunnel mpls traffic-eng bandwidth bandwidth</code>	Specifies the bandwidth that is reserved by the MPLS-TE tunnel
<code>tunnel mpls traffic-eng path-option 1 dynamic</code>	Specifies MPLS-TE path options
<code>tunnel mpls traffic-eng priority setup hold</code>	Specifies MPLS-TE tunnel setup and hold priority

Task 1: Configure the Backup Tunnels

In this task, you will enable the features that allow backup TE tunnels to be used.

Exercise Procedure

Complete these steps:

Step 1 On your assigned P router (R7 or R8), configure the backup tunnel to the other P router using the information in this table.

Parameter	Value
IP address	Loopback 0
Tunnel bandwidth	100 kbps
Path options	Explicit path selection
Tunnel priority	Setup=7, hold=7

Enter your commands here:

Path statements

Tunnel interface

Step 2 Assign the backup tunnel to the interface to R9.

Step 3 On your assigned PE router (R5 or R6), enable FRR in the primary tunnel interface.

Exercise Verification

You have completed this exercise when you attain these results:

- See lab answer.

Next Step

- Module 5: MPLS Quality of Service

Lab Exercise 5-1: Implementing MPLS QoS

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective


The service provider has decided to implement MPLS QoS in the core in order to meet its SLAs.

In this exercise, you will implement MPLS QoS in the core. After completing this exercise, you will be able to meet these objectives:

- Implement an inbound PE QoS policy
- Implement an outbound PE QoS policy
- Implement an outbound QoS policy from a local P router toward a remote P router
- Implement an outbound QoS policy on a P router toward a PE router
- Implement an outbound QoS policy on the PE router toward a CE router

Team A will configure the QoS parameters on the path from R5 to R6. Team B will configure the QoS parameters on the path from R6 to R5.

Visual Objective



Service Definition

Three classes of service should be provided: premium, business, and best-effort. Premium traffic should have a minimum bandwidth guarantee with low latency, low jitter, and no packet loss. Premium traffic that is received from CE devices should be limited to a certain amount. Any excess traffic should be dropped. Business traffic should have a minimum bandwidth guarantee with low packet loss for nonexcess traffic. Business traffic that is received from CE devices should be limited to a certain amount. Any traffic that exceeds the limits should be dropped first, rather than the traffic that meets the limits, in case of congestion. Best-effort traffic should not provide any guarantees, and CE devices should be limited by link capacity in terms of the amount of best-effort traffic that they can transmit.

In this case, you are not defining explicit thresholds for the service guarantees. Delay, jitter, and packet loss should be measured for the different classes during a period of time to be able to define actual thresholds.

Class	Guarantees	Is Traffic Limited?
Premium	Minimum bandwidth, low delay, low jitter, no packet loss	Yes. Excess traffic is dropped.
Business	Minimum bandwidth, low packet loss	Yes. Excess traffic is dropped first during congestion.
Best-effort	None	No.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
class-map <i>class-map-name</i> [match-all match-any] no class-map <i>class-map-name</i> [match-all match-any]	To create a class map to be used for matching packets to a specified class, use the class-map command in global configuration mode. To remove an existing class map from the router, use the no form of this command.
frame-relay class <i>name</i> no frame-relay class <i>name</i>	To associate a map class with an interface or subinterface, use the frame-relay class interface configuration command. To remove the association between the interface or subinterface and the named map class, use the no form of this command.
frame-relay traffic-shaping no frame-relay traffic-shaping	To enable both traffic shaping and per-virtual circuit queuing for all permanent virtual circuits (PVCs) and switched virtual circuits (SVCs) on a Frame Relay interface, use the frame-relay traffic-shaping interface configuration command. To disable traffic shaping and per-virtual circuit queuing, use the no form of this command.
map-class frame-relay <i>map-class-name</i>	To specify a map class to define QoS values for an SVC, use the map-class frame-relay global configuration command.
police <i>bps burst-normal burst-max conform-action action exceed-action action</i> [violate-action action] no police <i>bps burst-normal burst-max conform-action action exceed-action action</i> [violate-action action]	To configure the traffic policing feature, use the police policy map configuration command. To remove the traffic policing feature from the configuration, use the no form of this command.
policy-map <i>policy-map-name</i> no policy-map <i>policy-map-name</i>	To create or modify a policy map that can be attached to one or more interfaces to specify a service policy, use the policy-map command in global configuration command. To delete a policy map, use the no form of this command.
service-policy input <i>policy-map-name</i> no service-policy input <i>policy-map-name</i>	Use the service-policy input command to attach a policy map to an interface. Use the no form of this command to remove a policy map from an interface.

Task 1: Implement an Inbound PE QoS Policy

CE devices are expected to transmit premium traffic with an EF marking and business traffic with an AF2y marking. Any other DSCP is treated as best-effort. Premium traffic should be carried in the backbone with the experimental bits set to 5. In-profile business traffic should be carried in the backbone with the experimental bits set to 2. Out-of-profile business traffic should be carried with the experimental bits set to 1. In case of congestion, out-of-profile packets should be dropped before any in-profile packets. Best-effort traffic should be carried in the backbone with the experimental bits set to 0.

Class	Experimental Bits
Premium	5
Business (in profile)	2
Business (out of profile)	1
Best-effort	0

Implement an inbound QoS policy on your assigned PE to enforce the service contract. Make sure that you classify traffic based on the DSCP values that are expected from the CE. The policy should define the mapping between the DSCP values in the IP packets and the experimental bits in the MPLS packets.

In this particular case, you will focus on the QoS policies that are needed for traffic that is flowing from R1 toward R3 and from R3 toward R1. A complete QoS solution should consider QoS policies for the traffic that could flow from both CEs.

Exercise Procedure

Complete these steps:

- Step 1** On the PE router that is supporting your point of presence (POP), create a class map for the class “PREMIUM-IP.” CE devices are expected to transmit premium traffic with an EF marking.
- Step 2** On the PE router that is supporting your POP, create a class map for the class “BUSINESS-IP.” CE devices are expected to transmit business traffic with an AF21, AF22, or AF23 marking.
- Step 3** On the PE router that is supporting your POP, create a policy map that is named “IN-POLICY” and assign the proper conform action and exceed action to each class.
- PREMIUM-IP
 - Police: cir = 1280000, bc = 32000, and be = 32000
 - Conform action: set mpls exp in the top label to 5 and transmit
 - Exceed action: drop
 - BUSINESS-IP
 - Police: cir = 22000000, bc = 550000, and be = 550000
 - Conform action: set mpls exp in the top label to 2 and transmit
 - Exceed action: set mpls exp in the top label to 1 and transmit

- Class-default
 - **set mpls experimental imposition 0**

- Step 4** On the PE router that is supporting your POP , assign the service policy to the proper inbound interface.
- Step 5** On the PE router that is supporting the other POP , assign the service policy to the proper inbound interface .

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **sh class-map** command on your assigned router (R5 or R6) to verify the class map.

```
PEpop#sh class-map
Class Map match-any class-default (id 0)
  Match any

Class Map match-all BUSINESS-IP (id 2)
  Match ip dscp af21 af22 af23

Class Map match-all PREMIUM-IP (id 1)
  Match ip dscp ef
```

- Use the **sh policy-map interface** command on your assigned router (R5 or R6) to verify the policy map and interface assignments.

```
PEpop#show policy-map interface Serial0/0.101

Service-policy input: IN-POLICY

Class-map: PREMIUM-IP (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: ip dscp ef
  police:
    cir 1280000 bps, bc 32000 bytes
    conformed 0 packets, 0 bytes; actions:
      set-mpls-exp-topmost-transmit 5
    exceeded 0 packets, 0 bytes; actions:
      drop
    conformed 0 bps, exceed 0 bps

Class-map: BUSINESS-IP (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: ip dscp af21 af22 af23
  police:
    cir 22000000 bps, bc 550000 bytes
    conformed 0 packets, 0 bytes; actions:
      set-mpls-exp-topmost-transmit 2
    exceeded 0 packets, 0 bytes; actions:
      set-mpls-exp-topmost-transmit 1
    conformed 0 bps, exceed 0 bps

Class-map: class-default (match-any)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: any
  QoS Set
```

```
mpls experimental imposition 0
Packets marked 0
```

- On your CE router, use the extended **ping** command to send a packet with varied ToS values (for example, AF21 = 0x48, AF22 = 0x50, AF23 = 0x58, and EF = 0xb8). On one of the attached PE routers, use the **sh policy-map interface** command to verify that the policy map is assigning the packet to the proper class.

Task 2: Implement an Outbound PE QoS Policy

Implement an outbound QoS policy on the PE router toward your P routers to guarantee at least 5 percent of the link bandwidth to premium traffic and at least 50 percent to business traffic. The policy should drop out-of-profile business packets before dropping any in-profile business packets.

In case of congestion, best-effort traffic should experience random packet drops. This policy should be acting on MPLS packets, because IP packets have already gone through label imposition.

Exercise Procedure

Complete these steps:

- Step 1** On the PE router that is supporting your POP, create a class map for the class “PREMIUM-MPLS.” This class should match the packets that are marked by the PREMIUM-IP class (for example, MPLS exp = 5).
- Step 2** On the PE router that is supporting your POP, create a class map for the class “BUSINESS-MPLS” (for example, MPLS exp = 1 and 2).
- Step 3** On the PE router that is supporting your POP, create a policy map named “OUT-POLICY” and assign the proper bandwidth and queuing to each class.
- PREMIUM-MPLS
 - Priority queuing with a bandwidth reservation of 5 percent
 - BUSINESS-MPLS
 - Bandwidth reservation of 50 percent
 - Random detect discard
 - For packets marked with an IP precedence of 1, WRED randomly drops 1 out of 100 packets when a threshold of 100 packets is reached and begins dropping all packets when a threshold of 300 packets is reached.
 - For packets marked with an IP precedence of 2, WRED randomly drops 1 out of 10 packets when a threshold of 300 packets is reached and begins dropping all packets when a threshold of 500 packets is reached.
 - Class-default
 - Fair queuing
 - Random detect discard
- Step 4** On the PE router that is supporting your POP, assign the service policy to the proper outbound interface.
- Step 5** On the PE router that is supporting the other POP, assign the service policy to the proper outbound interface.

Note Because this lab exercise uses Frame Relay subinterfaces, you will need to implement this QoS functionality using the frame-relay traffic shaping, map-class, and frame-relay class commands.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **sh class-map** command on your assigned router (R5 or R6) to verify the class map.

```
R6#sh class-map
Class Map match-any class-default (id 0)
  Match any

Class Map match-all BUSINESS-MPLS (id 4)
  Match mpls experimental topmost 1 2

Class Map match-all BUSINESS-IP (id 2)
  Match ip dscp af21 af22 af23

Class Map match-all PREMIUM-IP (id 1)
  Match ip dscp ef

Class Map match-all PREMIUM-MPLS (id 3)
  Match mpls experimental topmost 5
```

- Use the **sh policy-map interface** command to verify the policy map and interface assignments.

```
R6#sh policy-map int s0/0.111
Serial0/0.111: DLCI 111 -
```

Service-policy output: OUT-POLICY

```
Class-map: PREMIUM-MPLS (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: mpls experimental topmost 5
  Queueing
    Strict Priority
    Output Queue: Conversation 24
    Bandwidth 5 (%)
    Bandwidth 6 (kbps) Burst 150 (Bytes)
    (pkts matched/bytes matched) 0/0
    (total drops/bytes drops) 0/0
```

```
Class-map: BUSINESS-MPLS (match-all)
  0 packets, 0 bytes
  5 minute offered rate 0 bps, drop rate 0 bps
  Match: mpls experimental topmost 1 2
  Queueing
    Output Queue: Conversation 25
    Bandwidth 50 (%)
    Bandwidth 64 (kbps)
    (pkts matched/bytes matched) 0/0
    (depth/total drops/no-buffer drops) 0/0/0
    exponential weight: 9
    mean queue depth: 0
```

class	Transmitted pkts/bytes	Random drop pkts/bytes	Tail drop pkts/bytes	Minimum thresh	Maximum thresh	Mark prob
0	0/0	0/0	0/0	20	40	1/10
1	0/0	0/0	0/0	100	300	1/100
2	0/0	0/0	0/0	300	500	1/10
3	0/0	0/0	0/0	26	40	1/10

4	0/0	0/0	0/0	28	40	1/10
5	0/0	0/0	0/0	30	40	1/10
6	0/0	0/0	0/0	32	40	1/10
7	0/0	0/0	0/0	34	40	1/10
rsvp	0/0	0/0	0/0	36	40	1/10

```

Class-map: class-default (match-any)
  2184 packets, 653385 bytes
  5 minute offered rate 1000 bps, drop rate 0 bps
  Match: any
  Queueing
    Flow Based Fair Queueing
    Maximum Number of Hashed Queues 16
    (total queued/total drops/no-buffer drops) 0/0/0
    exponential weight: 9

```

class	Transmitted pkts/bytes	Random drop pkts/bytes	Tail drop pkts/bytes	Minimum thresh	Maximum thresh	Mark prob
0	466/518936	0/0	0/0	20	40	1/10
1	0/0	0/0	0/0	22	40	1/10
2	0/0	0/0	0/0	24	40	1/10
3	0/0	0/0	0/0	26	40	1/10
4	0/0	0/0	0/0	28	40	1/10
5	0/0	0/0	0/0	30	40	1/10
6	1718/135888	0/0	0/0	32	40	1/10
7	0/0	0/0	0/0	34	40	1/10
rsvp	0/0	0/0	0/0	36	40	1/10

- On your CE router, use the extended **ping** command to send a packet with varied ToS values (for example, AF21 = 0x48, AF22 = 0x50, AF23 = 0x58, and EF = 0xb8). On one of the attached PE routers, use the command **sh policy-map interface** to verify that the policy map is assigning the packet to the proper class.

Task 3: Implement an Outbound QoS Policy on the Local P Router Toward a Remote P Router

The P router does not need an inbound policy because the packets have already been marked at the PE router and now are being process-switched. The outbound policy is identical to that of the PE router.

Exercise Procedure

Complete these steps:

- Step 1** Replicate the PE (R5 or R6) outbound QoS policy by copying and pasting it onto the P router (R7 or R8) that is supporting your PE (R5 or R6) router.
- Step 2** Assign the service policy to the proper outbound interface (for example, s0/0.112 and s0/0.2wg1).

Exercise Verification

You have completed this exercise when you attain these results:

- Repeat the verification steps in Task 2.

Task 4: Implement an Outbound QoS Policy on the P Router Toward the PE Router

The P router does not need an inbound policy because the packets have already been marked at the PE router and now are being process-switched. The outbound policy is identical to that of the other P router.

Exercise Procedure

Complete these steps:

- Step 1** Replicate the PE (R5 or R6) outbound QoS policy by copying and pasting it onto the P router that is supporting your PE router.

Note If the other team has already created the QoS policy, it will not be necessary to create the QoS policy yourself.

- Step 2** Assign the service policy to the proper outbound interface (for example, s0/0.111).

Exercise Verification

You have completed this exercise when you attain these results:

- Repeat the verification steps in Task 2.

Task 5: Implement an Outbound QoS Policy on the PE Router Toward the CE Router

At this point the packets have been returned to their native state. The policy implemented in this task should be acting on IP packets, because the MPLS has already been removed. Implement an outbound QoS policy on your PE router toward the CE to guarantee at least 50 percent of the link bandwidth to premium traffic and at least 25 percent to business traffic. The policy should drop out-of-profile business packets before dropping any in-profile business packets. In case of congestion, best-effort traffic should experience random packet drops.

Team A will enable R6. Team B will enable R5.

Exercise Procedure

Complete these steps:

- Step 1** On the PE router that is supporting your POP, create a class map for the class "PREMIUM-IP." CE devices are expected to transmit premium traffic with an EF marking.
- Step 2** On the PE router that is supporting your POP, create a class map for the class "BUSINESS-IP." CE devices are expected to transmit business traffic with an AF21, AF22, or AF23 marking.

- Step 3** On the PE router that is supporting your POP, create a policy map that is named “IP-OUT” and assign the proper conform action and exceed action to each class.
- PREMIUM-IP
 - Priority percent 50
 - BUSINESS-IP
 - Bandwidth percent 25
 - Random-detect dscp-based
 - Random-detect dscp 18 30 40 1
 - Random-detect dscp 20 20 30 1
 - Random-detect dscp 22 10 20 1
 - Class-default
 - Fair-queue
 - Random-detect dscp-based
- Step 4** On the PE router that is supporting your POP, assign the service policy to the proper outbound interface.
- Step 5** On the PE router that is supporting the other POP, assign the service policy to the proper inbound interface.

Exercise Verification

You have completed this exercise when you attain these results:

- Repeat the verification steps in Task 2.

Lab Exercise 5-2: Implementing MPLS QoS and DS-TE

Complete this lab exercise to practice what you learned in the related lesson.

Note	This lab exercise requires hardware that supports MPLS-TE and DS-TE.

Exercise Objective

The customer has activated a new application that runs between servers on its sites.


Note	To simulate this application, you will need to add loopback addresses. Customer Green will use the addresses 10.1.1.64 and 10.1.2.64. Customer Orange will use the addresses 10.2.1.64 and 10.2.2.64.

The service provider has discovered that it cannot meet the service agreement by routing the premium traffic through P1. To correct this issue, the service provider has decided to implement DS-TE and route this traffic over the direct Pwg-to-Pwg router link. The QoS policies remain the same; only the traffic routing must be changed.

In this exercise, you will implement MPLS QoS in the core. After completing this exercise, you will be able to meet these objectives:

- Update existing tunnel parameters
- Implement DS-TE on the interfaces
- Create a DS-TE tunnel
- Assign traffic to tunnels

Visual Objective



Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<pre>ip route vrf vrf-name prefix mask [next-hop- address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]</pre>	To establish static routes for a VPN virtual routing and forwarding (VRF) instance, use the ip route vrf command in global configuration mode. To disable static routes, use the no form of this command.
<pre>no ip route vrf vrf-name prefix mask [next-hop- address] [interface {interface-number}] [global] [distance] [permanent] [tag tag]</pre>	
<pre>ip rsvp bandwidth interface-kbps single- flow-kbps [sub-pool kbps]</pre>	Enables RSVP for IP on an interface and specifies the amount of interface bandwidth (in kbps) that is allocated for RSVP flows (for example, TE tunnels).

Task 1: Update Existing Tunnel Parameters

The first issue that needs to be addressed is related to the existing tunnel. Its bandwidth reservation does not match the actual service agreement.

The QoS policy that was implemented in the previous lab exercise allocated 5 percent of the link bandwidth to PREMIUM-MPLS traffic and 50 percent of the link bandwidth to BUSINESS-MPLS traffic. Remaining traffic was to be best-effort. A tunnel, however, should not reserve bandwidth for best-effort traffic. It should reserve 5 percent (7 kbps) for the PREMIUM traffic and 50 percent (64 kbps) for the BUSINESS traffic for a total of 71 kbps.

Exercise Procedure

Complete these steps:

- Step 1** On the PE router (R5 for Team A or R6 for Team B) that is supporting your POP, change the existing tunnel bandwidth requirements to a global bandwidth of 64 kbps.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **sh ip rsvp interface** command to verify the RSVP configuration.

```
R6#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      10K
Se0/0.111    64K         128K     128K      10K
```

Note It may be necessary to manually shut down and then enable the tunnel interface to force the reservation to update before the reoptimization timer expires.

Task 2: Implement DS-TE on the Interfaces

Exercise Procedure

- Step 1** DS-TE enables a subpool of bandwidth. Update the reservable bandwidth on all links between the backbone routers (R5 and R7 for Team A, R6 and R8 for Team B). Allow up to 128 kbps of reservable bandwidth and 10 kbps of reservable bandwidth in the subpool on any single subinterface. Allow up to 128 kbps of reservable bandwidth and 10 kbps of reservable bandwidth in the subpool on physical interfaces.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **sh ip rsvp interface** command to verify (on either R5 or R6) the RSVP configuration.

```
PEpop#sh ip rsvp int
interface    allocated  i/f max  flow max  sub max
Se0/0        0           128K     128K      10K
Se0/0.111    64K         128K     128K      10K
```

Task 3: Create a DS-TE Tunnel

Exercise Procedure

- Step 1** On the PE router (R5 for Team A or R6 for Team B) that is supporting your POP, create an explicit path that uses your local P-router-to-P-router link.
- Step 2** On the PE router (R5 for Team A or R6 for Team B) that is supporting your POP, create a premium tunnel with the following features:
- A setup and pre-emption priority of 1.
 - Configuration of a bandwidth reservation of 7 kbps from the subpool.
 - Autoroute announcement with a relative metric of 4. You cannot announce the new tunnel with a default metric because it will be load-balanced with the previous tunnel.

Exercise Verification

You have completed this exercise when you attain these results:

- Use the **sh mpls traffic-eng tunnels br** command to verify that tunnels are active.

```
R5#sh mpls traffic-eng tunnels br
```

Signaling Summary:

```
LSP Tunnels Process:      running
RSVP Process:             running
Forwarding:               enabled
Periodic reoptimization:  every 30 seconds, next in 29 seconds
Periodic auto-bw collection: disabled
```

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R5_t1	192.168.wg.33	-	Se0/0.11	up/up
R5_t1	192.168.wg.33	-	Se0/0.11	up/up
R6_t1	192.168.wg.17	Se0/0.11	-	up/up
R6_t1	192.168.wg.17	Se0/0.11	-	up/up

Displayed 2 (of 2) heads, 0 (of 0) midpoints, 2 (of 2) tails

- Use the **sh mpls traffic-eng tunnels name** command to verify that the tunnel is using the proper path.

```
R5#sh mpls traffic-eng tunnels name R5_t1
```

```
Name: R5_t1 (Tunnel2) Destination: 192.168.wg.33
```

Status:

```
Admin: up      Oper: up      Path: valid      Signaling: connected
```

```
path option 1, type explicit R5_R6 (Basis for Setup, path weight 30)
```

Config Parameters:

```
Bandwidth: 7 kbps (Global) Priority: 1 1 Affinity: 0x0/0xFFFF
Metric Type: TE (default)
AutoRoute: disabled LockDown: disabled Loadshare: 7 bw-based
auto-bw: disabled
```

```
InLabel : -
```

```
OutLabel : Serial0/0.111, 37
```

RSVP Signaling Info:

```
Src 192.168.wg.17, Dst 192.168.wg.33, Tun_Id 2, Tun_Instance 2
```

RSVP Path Info:

```
My Address: 192.168.wg.17
```

```
Explicit Route: 192.168.wg.50 192.168.wg.114 192.168.wg.65 192.168.wg.33
```

```

Record Route:  NONE
Tspec: ave rate=7 kbits, burst=1000 bytes, peak rate=7 kbits
RSVP Resv Info:
Record Route:  NONE
Tspec: ave rate=7 kbits, burst=1000 bytes, peak rate=7 kbits
History:
Tunnel:
Time since created: 16 minutes, 55 seconds
Time since path change: 7 minutes, 43 seconds
Current LSP:
Uptime: 7 minutes, 43 seconds
Prior LSP:
ID: path option 1 [1]

Removal Trigger: path option updated

```

Task 4: Assign the Traffic to the Tunnels

The PREMIUM and BUSINESS traffic are following the tunnel that was created and announced in the previous lab exercise. In this task you will redirect the PREMIUM traffic over the DiffServ tunnel.

To simulate the prefix that contains the PREMIUM traffic, loopback interfaces have been created.

Router	Loopback Address
R1	10.1.1.65 255.255.255.240
R2	10.2.1.65 255.255.255.240
R3	10.1.2.65 255.255.255.240
R4	10.2.2.65 255.255.255.240

Exercise Procedure

- Step 1** On the PE routers that support your VRF tables, create an IP VRF route that points the PREMIUM traffic for your customer (Customer Green for Team A and Customer Orange for Team B) to the new tunnel.

Exercise Verification

You have completed this exercise when you attain these results:

- On each CE router, use the **trace** command to verify that tunnels are active.

Example:

```
R1#trace 10.1.2.65
```

Type escape sequence to abort.

Tracing the route to 10.wg.wg2.65

```

 1 150.wg.wg1.18 16 msec 16 msec 12 msec
 2 192.168.wg.50 [MPLS: Labels 16/44 Exp 0] 300 msec 369 msec 300 msec
 3 192.168.wg.114 [MPLS: Labels 19/44 Exp 0] 224 msec 188 msec 201 msec
 4 150.wg.42.18 [AS 650wg2] [MPLS: Label 44 Exp 0] 116 msec 117 msec 116 msec
 5 150.wg.42.17 [AS 650wg2] 232 msec 76 msec *

```

Next Step:

- Module 6: Any Transport over MPLS

Lab Exercise 6-1: Implementing a Port-to-Port FRoMPLS Configuration (Case Study)

The AToM feature is currently limited to selected routers. Because of this limitation, this lab will be performed using a case study. It will use the basic topology that you have been using in the prior lab exercises, but here you will assume that the CE-to-PE routing has been disabled.


Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

There are two methods of transporting Frame Relay over an MPLS core. In this exercise, you will implement the port-to-port Frame Relay transport method for Customer A. After completing this exercise, you will be able to meet these objectives:

- Enable the PE routers to transport AToM packets
- Enable the PE routers to transport Frame Relay with port-to-port connections

Visual Objective



The following are the relevant portions of the initial router configurations:

```
Router: R5
ip cef
!
interface Loopback0
 ip address 192.168.wg.17 255.255.255.255

interface Serial0/0
```



```

no ip address
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco

interface Serial0/0.111 point-to-point
description ***** DLCI to the P router *****
ip address 192.168.wg.49 255.255.255.240
frame-relay interface-dlci 111
mpls label protocol ldp
tag-switching ip
!
interface Serial0/1
description **** link to the Frame Relay Switch ****

!
router isis
net 49.0001.0000.0000.0111.00
is-type level-2-only
metric-style wide
!
router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 192.168.100.129 remote-as 65001
neighbor 192.168.100.129 update-source Loopback0
no auto-summary

```

Router: R6

```

ip cef
!
interface Loopback0
ip address 192.168.wg.33 255.255.255.255

interface Serial0/0
no ip address
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
!

interface Serial0/0.111 point-to-point
description ***** DLCI to the P router *****
ip address 192.168.wg.49 255.255.255.240
frame-relay interface-dlci 111
mpls label protocol ldp
tag-switching ip
!
interface Serial0/1
description **** link to the Frame Relay Switch ****
!
router isis
net 49.0001.0000.0000.0111.00
is-type level-2-only
metric-style wide
!
router bgp 65001
no synchronization
bgp log-neighbor-changes

```

```
neighbor 192.168.100.129 remote-as 65001
neighbor 192.168.100.129 update-source Loopback0
no auto-summary
```

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<code>mpls label protocol {ldp tdp}</code> <code>no mpls label protocol</code>	To specify the default LDP for a platform, use the mpls label protocol command in global configuration mode. To restore the image default, use the no form of this command.
<code>mpls ldp router-id interface [force]</code> <code>no mpls ldp router-id</code>	To specify a preferred interface for determining the LDP router ID, use the mpls ldp router-id command in global configuration mode. To remove the preferred interface for determining the LDP router ID, use the no form of this command.
<code>mpls l2transport route destination vc-id</code>	Creates the virtual circuit to transport Layer 2 packets.

Task 1: Enable the PE Routers to Transport AToM Packets

Before configuring AToM, perform the following configuration:

- Provide LSPs between PE routers by enabling dynamic MPLS labeling (with the **mpls ip** command) on all paths between the imposition and disposition PE routers.
- Enable and configure Cisco Express Forwarding (CEF) or distributed CEF (dCEF) before configuring ATM AAL5 over MPLS, ATM Cell Relay over MPLS, or EoMPLS. Enable and configure CEF before configuring FRoMPLS, HDLC over MPLS, or PPP over MPLS.

There are a couple steps that are required to prepare the PE to transport AToM packets. These steps are as follows:

- Enabling the MPLS label protocol at the global level
- Setting the LDP router ID

For this exercise, you will use loopback0 as the LDP ID.

Exercise Procedure

Which commands are used to perform the following?:

- Enabling LDP as the MPLS label protocol: _____
- Setting the LDP router ID to loopback0: _____

Exercise Verification

You have completed this exercise when you attain these results:

- Your configuration should match the configuration that is provided in the Lab Exercise Answer Key.

Task 2: Enable PE Routers to Transport Frame Relay with Port-to-Port Connections

When you are transporting Frame Relay in the port-to-port mode, the AToM port operates in an HDLC encapsulation mode and the Frame Relay signaling is end-to-end. This mode is commonly used to connect the customer premises equipment (CPE) (Frame Relay DTE) to a Frame Relay switch (Frame Relay DCE).

Exercise Procedure

Which commands are necessary to enable bidirectional port-to-port transport between the two CE routers of customer A using a VC ID of 50?

R5

R6

Exercise Verification

You have completed this exercise when you attain these results:

- Your configuration should match the configuration that is provided in the Lab Exercise Answer Key.

Lab Exercise 6-2: Implementing a DLCI-to-DLCI FRoMPLS Configuration (Case Study)


Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

In the previous exercise you implemented a port-to-port transport. In this exercise, you will implement the DLCI-to-DLCI Frame Relay transport method. After completing this exercise, you will be able to meet these objectives:

- Enable PE routers to transport Frame Relay with DLCI-to-DLCI connections

Visual Objective



The following are the relevant portions of the initial router configurations:

```
Router: R5
ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip
!
interface Loopback0
 ip address 192.168.wg.17 255.255.255.255

interface Serial0/0
 no ip address
 encapsulation frame-relay
 no fair-queue
```

```

    frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
    description ***** DLCI to the P router *****
    ip address 192.168.wg.49 255.255.255.240
    frame-relay interface-dlci 111
    mpls label protocol ldp
    tag-switching ip
!
interface Serial1/0
    description **** link to the CE router for customer A *****

!
interface Serial1/1
    description **** link to the Customer B's CE router *****

router isis
    net 49.0001.0000.0000.0111.00
    is-type level-2-only
    metric-style wide
!
router bgp 65001
    no synchronization
    bgp log-neighbor-changes
    neighbor 192.168.100.129 remote-as 65001
    neighbor 192.168.100.129 update-source Loopback0
    no auto-summary

```

Router: R6

```

ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip
!
interface Loopback0
    ip address 192.168.wg.33 255.255.255.255

interface Serial0/0
    no ip address
    encapsulation frame-relay
    no fair-queue
    frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
    description ***** DLCI to the P router *****
    ip address 192.168.wg.49 255.255.255.240
    frame-relay interface-dlci 111
    mpls label protocol ldp
    tag-switching ip
!
interface Serial1/0
    description **** link to the CE router for customer A *****
!
interface Serial1/1
    description **** link to the Customer B's CE router *****
!
router isis
    net 49.0001.0000.0000.0111.00
    is-type level-2-only

```

```

metric-style wide
!
router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 192.168.100.129 remote-as 65001
neighbor 192.168.100.129 update-source Loopback0
no auto-summary
!

```

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<code>connect connection-name</code> <code>interface dlc</code> <code>l2transport</code>	Defines connections between Frame Relay PVCs. Using the l2transport keyword specifies that the PVC will not be a locally switched PVC but will be tunneled over the backbone network. The argument <i>connection-name</i> is a text string that you provide. The argument <i>interface</i> is the interface on which a PVC connection will be defined. The argument <i>dlci</i> is the DLCI number of the PVC that will be connected.
<code>mpls label protocol {ldp tdp}</code> <code>no mpls label protocol</code>	To specify the default LDP for a platform, use the mpls label protocol command in global configuration mode. To restore the image default, use the no form of this command.
<code>mpls ldp router-id interface [force]</code> <code>no mpls ldp router-id</code>	To specify a preferred interface for determining the LDP router ID, use the mpls ldp router-id command in global configuration mode. To remove the preferred interface for determining the LDP router ID, use the no form of this command.
<code>mpls l2transport route destination vc-id</code>	Creates the virtual circuit to transport the Layer 2 packets.

Task 1: Enable PE Routers to Transport Frame Relay with DLCI-to-DLCI Connections

In the previous lab exercise, you prepared the routers to transport AToM packets. In this task you will add Frame Relay DLCI-to-DLCI support for customer B.

Exercise Procedure

Complete the commands necessary to enable bidirectional DLCI-to-DLCI transport between the two CE routers of customer B by using a connection name of “fr1,” a DLCI of 210, and a VC ID of 303.

- R5

[illegible]

[illegible]

You have completed this exercise when you attain these results:

- ## Next Step:

- 90 Implementing Cisco MPLS Traffic Engineering and Other Features (MPLST) v2.0 Copyright © 2004, Cisco Systems, Inc.
The PDF files and any printed representation for this material are the property of Cisco Systems, Inc.,
for the sole use by Cisco employees for personal study. The files or printed representations may not be
used in commercial training, and may not be distributed for purposes other than individual self-study.
<https://it-ebooks.net>


Lab Exercise 7-1: Implementing Cisco 6PE Support (Case Study)

Complete this lab exercise to practice what you learned in the related lesson.

Exercise Objective

IPv6 is becoming more and more accepted. To meet this need, MPLS has implemented support of IPv6 in several modes. This lab exercise focuses on the Cisco 6PE model.

Note Because the hardware that was issued in the lab setup for this course does not support the Cisco 6PE feature, this lab exercise will be done as a case study.




In this exercise, the carrier Yellow needs to provide IPv6 connections for its customers, Green and Orange. To do this, carrier Yellow is going to implement Cisco 6PE functionality.


After completing this exercise, you will be able to meet these objectives:

- Review the initial configuration on all routers
- Enable IPv6 support on the PE routers (R3 and R4)
- Enable IPv6 BGP connectivity between PE routers
- Enable IPv6 EBGp routing between routers

Visual Objective



For this exercise, the core IP address ranges remain the same as in the previous lab exercises. The 192.168.0.0 range is used to provide addressing for the core MPLS network of the carrier.



Note The addresses of customer Orange are simulated using loopback addresses.

Required Resources

These are the resources and equipment required to complete this exercise:

- Cisco IOS documentation

Command List

The commands that are used in this exercise are described in the table here.

Command	Description
<code>ipv6 unicast-routing</code> <code>no ipv6 unicast-routing</code>	To enable the forwarding of IPv6 unicast datagrams, use the ipv6 unicast-routing command in global configuration mode. To disable the forwarding of IPv6 unicast datagrams, use the no form of this command.
<code>mpls ipv6 source-interface type number</code> <code>no mpls ipv6 source-interface type number</code>	To specify an IPv6 address of an interface to be used as the source address for locally generated IPv6 packets to be sent over a Multiprotocol Label Switching (MPLS) network, use the mpls ipv6 source-interface command in global configuration mode. To disable this feature, use the no form of this command.
<code>ipv6 address [prefix-name ipv6-prefix/prefix-length]</code> <code>no ipv6 address [prefix-name ipv6-prefix/prefix-length]</code>	To configure an IPv6 address based on an IPv6 general prefix and enable IPv6 processing on an interface, use the ipv6 address command in interface configuration mode. To remove the address from the interface, use the no form of this command.
<code>bgp default ipv4-unicast</code> <code>no bgp default ipv4-unicast</code>	To enable the IP version 4 (IPv4) unicast address family on all neighbors, use the bgp default ipv4-unicast command in address family or router configuration mode. To disable the IPv4 unicast address family on all neighbors, use the no form of this command.
<code>address-family ipv6 [unicast multicast]</code> <code>no address-family ipv6 [unicast multicast]</code>	To enter address family configuration mode for configuring routing sessions such as Border Gateway Protocol (BGP) that use standard IPv6 address prefixes, use the address-family ipv6 command in router configuration mode. To disable address family configuration mode, use the no form of this command.
<code>ipv6 route ipv6-prefix/prefix-length {ipv6-address interface-type interface-number [ipv6-address]} [administrative-distance] [administrative-multicast-distance unicast multicast]</code> <code>no ipv6 route ipv6-prefix/prefix-length {ipv6-address interface-type interface-number [ipv6-address]} [administrative-distance] [administrative-multicast-distance unicast multicast]</code>	To establish static IPv6 routes, use the ipv6 route command in global configuration mode. To remove a previously configured static route, use the no form of this command.

Task 1: Implement the Initial Configurations on the Routers

As this lab exercise begins, the IGP and MPLS have been configured in the core. The customer IPv6 routing has been configured.

Exercise Procedure

Step 1 Review the initial router configurations.

R1

```
hostname R5

ip cef
ipv6 unicast-routing
ipv6 cef

interface Loopback0
  no ip address
  ipv6 address 2001:401:111:8::1/64

interface Loopback100
  no ip address
  ipv6 address 2001:401:111:2::1/64

interface Serial0/0
  bandwidth 128000
  no ip address
  encapsulation frame-relay
  frame-relay lmi-type cisco

interface Serial0/0.101 point-to-point
  description Link to R3
  ipv6 address 2001:401:111:1::2/64
  frame-relay interface-dlci 101
exit

router bgp 65111
  no synchronization
  bgp router-id 1.1.1.1
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  no auto-summary
  !
  address-family ipv4 multicast
  no auto-summary
  no synchronization
  exit-address-family
  !
  address-family ipv6
  no synchronization
  redistribute connected
  exit-address-family
  !
  address-family ipv4
  no auto-summary
  no synchronization
  exit-address-family
exit
```

```
line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login
```

R2

```
hostname R2

ip cef
ipv6 unicast-routing
ipv6 cef

interface Loopback0
  no ip address
  ipv6 address 2001:401:112:8::1/64

interface Loopback100
  no ip address
  ipv6 address 2001:401:112:2::1/64

interface Serial0/0
  bandwidth 128000
  no ip address
  encapsulation frame-relay
  frame-relay lmi-type cisco

interface Serial0/0.101 point-to-point
  description Link to R4
  ipv6 address 2001:401:112:1::2/64
  frame-relay interface-dlci 101
exit

router bgp 65112
  no synchronization
  bgp router-id 3.3.3.3
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  no auto-summary
  !
  address-family ipv4 multicast
  no auto-summary
  no synchronization
  exit-address-family
  !
  address-family ipv6
  no synchronization
  redistribute connected
  exit-address-family
  !
  address-family ipv4
  no auto-summary
  no synchronization
  exit-address-family
exit
line con 0
  exec-timeout 999 0
```

```
logging synchronous
line aux 0
line vty 0 4
no login
```

```
exit
exit
```

R3

```
hostname R3

ip cef

tag-switching tdp router-id Loopback0

interface Loopback0
 ip address 192.168.1.17 255.255.255.255
 ip router isis

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 ip rsvp bandwidth 128 128
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R1
 frame-relay interface-dlci 101

interface Serial0/0.111 point-to-point
 description Link to R5
 ip address 192.168.1.49 255.255.255.240
 ip router isis
 mpls label protocol tdp
 mpls ip
 frame-relay interface-dlci 111

router isis
 net 39.0440.0000.0000.1500
 metric-style wide
exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
no login
```

R4

```
hostname R4

ip cef

tag-switching tdp router-id Loopback0
```

```

interface Loopback0
  ip address 192.168.1.33 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  description Link to R2
  frame-relay interface-dlci 101

interface Serial0/0.111 point-to-point
  description Link to R6
  ip address 192.168.1.65 255.255.255.240
  ip router isis
  mpls label protocol tdp
  mpls ip
  frame-relay interface-dlci 111

router isis
  net 39.0440.0000.0000.1600
  metric-style wide

exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

R5

```

hostname R5

ip cef

tag-switching tdp router-id Loopback0

interface Loopback0
  ip address 192.168.1.81 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.111 point-to-point
  description Link to R3
  ip address 192.168.1.50 255.255.255.240

```



```

ip router isis
tag ip
frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
description Link to R6
ip router isis
ip address 192.168.1.113 255.255.255.240
tag ip
frame-relay interface-dlci 112

router isis
net 39.0440.0000.0000.0700
metric-style wide
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

R6

```

hostname R6

ip cef

tag-switching tdp router-id Loopback0

interface Loopback0
ip address 192.168.1.97 255.255.255.255
ip router isis

interface Serial0/0
no ip address
bandwidth 128000
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
no shut

interface Serial0/0.111 point-to-point
description Link to R4
ip address 192.168.1.66 255.255.255.240
ip router isis
tag ip
frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
description Link to R5
ip address 192.168.1.114 255.255.255.240
ip router isis
tag ip
frame-relay interface-dlci 112

router isis
net 39.0440.0000.0000.0800
metric-style wide

```

```

exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

Exercise Verification

You have completed this exercise when you attain these results. Here are typical outputs from the **show** commands:

- On R1, the **sh ipv6 rou** command is used to verify the routing table.

```

R1#sh ipv6 rou
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
C   2001:401:111:1::/64 [0/0]
    via ::, Serial0/0.101
L   2001:401:111:1::2/128 [0/0]
    via ::, Serial0/0.101
C   2001:401:111:2::/64 [0/0]
    via ::, Loopback100
L   2001:401:111:2::1/128 [0/0]
    via ::, Loopback100
C   2001:401:111:8::/64 [0/0]
    via ::, Loopback0
L   2001:401:111:8::1/128 [0/0]
    via ::, Loopback0
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0

```

- On R3, the **sh ip rou** command is used to verify the routing table.

```

R3#sh ip rou
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route

```

Gateway of last resort is not set

```

      192.168.10.0/32 is subnetted, 1 subnets
i L2   192.168.10.17 [115/30] via 192.168.wg.50, Serial0/0.111
      192.168.wg.0/24 is variably subnetted, 9 subnets, 3 masks
i L1   192.168.wg.97/32 [115/30] via 192.168.wg.50, Serial0/0.111
i L1   192.168.wg.112/28 [115/20] via 192.168.wg.50, Serial0/0.111
i L1   192.168.wg.64/28 [115/30] via 192.168.wg.50, Serial0/0.111
i L1   192.168.wg.81/32 [115/20] via 192.168.wg.50, Serial0/0.111
i L1   192.168.wg.33/32 [115/40] via 192.168.wg.50, Serial0/0.111
C      192.168.wg.48/28 is directly connected, Serial0/0.111
C      192.168.wg.17/32 is directly connected, Loopback0

```

```
i L1 192.168.wg.128/29 [115/20] via 192.168.wg.50, Serial0/0.111
i L1 192.168.wg.144/29 [115/30] via 192.168.wg.50, Serial0/0.111
```

- On R5, the **sh ip rou** command is used to verify the routing table.

```
R5#sh ip rou
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       * - candidate default, U - per-user static route, o - ODR
       P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
192.168.10.0/32 is subnetted, 1 subnets
i L1 192.168.1.97/32 [115/20] via 192.168.1.114, Serial0/0.112
C    192.168.1.112/28 is directly connected, Serial0/0.112
i L1 192.168.1.64/28 [115/20] via 192.168.1.114, Serial0/0.112
C    192.168.1.81/32 is directly connected, Loopback0
i L1 192.168.1.33/32 [115/30] via 192.168.1.114, Serial0/0.112
C    192.168.1.48/28 is directly connected, Serial0/0.111
i L1 192.168.1.17/32 [115/20] via 192.168.1.49, Serial0/0.111
```

- On R2, the **sh ipv6 rou** command is used to verify the routing table.

```
R2#sh ipv6 rou
```

```
IPv6 Routing Table - 8 entries
```

```
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
C    2001:401:112:1::/64 [0/0]
    via ::, Serial0/0.101
L    2001:401:112:1::2/128 [0/0]
    via ::, Serial0/0.101
C    2001:401:112:2::/64 [0/0]
    via ::, Loopback100
L    2001:401:112:2::1/128 [0/0]
    via ::, Loopback100
C    2001:401:112:8::/64 [0/0]
    via ::, Loopback0
L    2001:401:112:8::1/128 [0/0]
    via ::, Loopback0
L    FE80::/10 [0/0]
    via ::, Null0
L    FF00::/8 [0/0]
    via ::, Null0
```

- On R4, the **sh ip rou** command is used to verify the routing table.

```
R4#sh ip rou
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
192.168.10.0/32 is subnetted, 1 subnets
i L2 192.168.10.17 [115/30] via 192.168.1.66, Serial0/0.111
192.168.1.0/24 is variably subnetted, 9 subnets, 3 masks
i L1 192.168.1.97/32 [115/20] via 192.168.1.66, Serial0/0.111
i L1 192.168.1.112/28 [115/20] via 192.168.1.66, Serial0/0.111
C 192.168.1.64/28 is directly connected, Serial0/0.111
i L2 192.168.1.81/32 [115/40] via 192.168.1.66, Serial0/0.111
C 192.168.1.33/32 is directly connected, Loopback0
i L2 192.168.1.48/28 [115/40] via 192.168.1.66, Serial0/0.111
i L2 192.168.1.17/32 [115/50] via 192.168.1.66, Serial0/0.111
i L2 192.168.1.128/29 [115/30] via 192.168.1.66, Serial0/0.111
i L1 192.168.1.144/29 [115/20] via 192.168.1.66, Serial0/0.111
```

- On R6, the **sh ip rou** command is used to verify the routing table.

```
R6#sh ip rou
```

```
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
        * - candidate default, U - per-user static route, o - ODR
        P - periodic downloaded static route
```

```
Gateway of last resort is not set
```

```
192.168.1.0/24 is variably subnetted, 9 subnets, 3 masks
C 192.168.1.97/32 is directly connected, Loopback0
C 192.168.1.112/28 is directly connected, Serial0/0.112
C 192.168.1.64/28 is directly connected, Serial0/0.111
i L1 192.168.1.81/32 [115/20] via 192.168.1.113, Serial0/0.112
i L1 192.168.1.33/32 [115/20] via 192.168.1.65, Serial0/0.111
i L1 192.168.1.48/28 [115/20] via 192.168.1.113, Serial0/0.112
i L1 192.168.1.17/32 [115/30] via 192.168.1.113, Serial0/0.112
i L1 192.168.1.128/29 [115/20] via 192.168.1.113, Serial0/0.112
```

Task 2: Enable IPv6 on the PE Routers

The first step to implementing Cisco 6PE is to enable IPv6 on the PE routers.

Exercise Procedure

Complete the following commands.

On the PE router R3, complete the following steps:

- Step 1** Enable IPv6 unicast routing.
- Step 2** Enable IPv6 CEF.
- Step 3** Specify the IPv6 address of interface loopback0 as the source address for locally generated IPv6 packets to be sent over the MPLS network.
- Step 4** Assign an IPv6 address to loopback0 (2001:401:1:1::1 /64).
- Step 5** Enable an IPv6 address on interface s0/0.101(2001:401:111:1::1 /64).

On the PE router R4, complete the following steps:

- Step 6** Enable IPv6 unicast routing.
- Step 7** Enable IPv6 CEF.
- Step 8** Specify the IPv6 address of interface loopback0 as the source address for locally generated IPv6 packets to be sent over the MPLS network.
- Step 9** Assign an IPv6 address to loopback0 (2001:401:1:2::1 /64).

Step 10 Enable an IPv6 address on interface s0/0.101(2001:401:112:1::1 /64).

Exercise Verification

The exercise verification should be done by comparing your configuration additions to those in the Lab Exercise Answer Key.

You have completed this exercise when you attain these results. Here are typical outputs from the **show** commands:

On R3, use the **sh ipv6 rou** command to verify the routing table.

```
R3#sh
*Mar 25 03:29:23.419: %SYS-5-CONFIG_I: Configured from console by console
R5#sh ipv6 rou
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
C   2001:401:1:1::/64 [0/0]
    via ::, Loopback0
L   2001:401:1:1::1/128 [0/0]
    via ::, Loopback0
C   2001:401:111:1::/64 [0/0]
    via ::, Serial0/0.101
L   2001:401:111:1::1/128 [0/0]
    via ::, Serial0/0.101
L   FE80::/10 [0/0]
    via ::, Null0
L   FF00::/8 [0/0]
    via ::, Null0
```

■ On R4, use the **sh ipv6 rou** command to verify the routing table.

```
R4#sh ipv6 rou
IPv6 Routing Table - 8 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
```

```

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
C 2001:401:1:2::/64 [0/0]
  via ::, Loopback0
L 2001:401:1:2::1/128 [0/0]
  via ::, Loopback0
C 2001:401:112:1::/64 [0/0]
  via ::, Serial0/0.101
L 2001:401:112:1::1/128 [0/0]
  via ::, Serial0/0.101
C 2001:401:122:1::/64 [0/0]
  via ::, Serial0/0.102
L FE80::/10 [0/0]
  via ::, Null0
L FF00::/8 [0/0]
  via ::, Null0

```

Task 3: Establish IBGP Connectivity Between PE routers

The following steps enable you to establish an Internal Border Gateway Protocol (IBGP) peer session between the PE routers that will carry the IPv6 routes. In this type of implementation, a router reflector is typically used. However, for the sake of simplicity, you will establish direct peering.

Exercise Procedure

Complete the following commands:

On the PE routers (R7 or R8), complete the following steps:

Step 1 Enable the BGP 65001 process.

Note	If the BGP session will be carrying only IPv6 routes, you may disable IPv4 unicast routing updates to reduce overhead traffic. For this lab exercise, you will disable the IPv4 updates.

Step 2 Disable IPv4 updates.

Step 3 Configure neighbor peering with the remote PE router (R4 or R3).

Step 4 Activate neighbor peering under the IPv6 address family.

Step 5 Configure the BGP process to send a label to R9 under the IPv6 address family.

R3

R4

Exercise Verification

The exercise verification should be done by comparing your configuration additions to those in the Lab Exercise Answer Key.

You have completed this exercise when you attain these results. Here are typical outputs from the **show** commands:

- On R3, use the **sh bgp ipv6 sum** command to verify the status of the peer connection.

```
R3#sh bgp ipv6 sum
BGP router identifier 192.168.1.17, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ  OutQ Up/Down  State/PfxRcd
192.168.1.33  4  65001      15       15         1    0    0 00:11:44      0
```

- On R4, use the **sh bgp ipv6 sum** command to verify the status of the end-to-end connections.

```
R4#sh bgp ipv6 sum
BGP router identifier 192.168.1.33, local AS number 65001
BGP table version is 1, main routing table version 1

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ  OutQ Up/Down  State/PfxRcd
192.168.1.17  4  65001       4        4         1    0    0 00:00:29      0
```


Task 4: Enable IPv6 Routing Between PE Routers

The following steps enable you to establish IPv6 routing between the PE routers.

Exercise Procedure

Complete the following commands:

On R7, complete the following steps:

Step 1 Enable the address range of loopback0 in the IPv6 address family.

Step 2 Enable the address range of s0/0.101 in the IPv6 address family.

On R8, complete the following steps:

Step 3 Enable the address range of loopback0 in the IPv6 address family.

Step 4 Enable the address range of s0/0.101 in the IPv6 address family.

Exercise Verification

The exercise verification should be done by comparing your configuration additions to those in the Lab Exercise Answer Key.

You have completed this exercise when you attain these results. Here are typical outputs from the **show** commands:

- On R4, use the **sh ipv6 rou** command to verify that the routes from R3 (that is, the BGP routes) are present.

```
R4#sh ipv6 rou
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2
B    2001:401:1:1::/64 [200/0]
     via ::FFFF:192.168.1.17, IPv6-mpls
C    2001:401:1:2::/64 [0/0]
     via ::, Loopback0
L    2001:401:1:2::1/128 [0/0]
     via ::, Loopback0
B    2001:401:111:1::/64 [200/0]
     via ::FFFF:192.168.1.17, IPv6-mpls
C    2001:401:112:1::/64 [0/0]
     via ::, Serial0/0.101
L    2001:401:112:1::1/128 [0/0]
     via ::, Serial0/0.101
L    FE80::/10 [0/0]
     via ::, Null0
L    FF00::/8 [0/0]
     via ::, Null0
```

- On R3, use the **sh bgp ipv6 labels** command to verify the BGP label that is assigned to the routes.

```
R3#sh bgp ipv6 labels
Network          Next Hop      In label/Out label
2001:401:1:1::/64
                  ::          27/nolabel
2001:401:111:1::/64
                  ::          28/nolabel
2001:401:112:1::/64
```

::FFFF:192.168.1.33

nolabel/28

- On R3, use the **sh ipv6 cef** command to verify the forwarding for the routes.

```
R3#sh ipv6 cef
2001:401:1:1::1/128
    Receive
2001:401:1:1::/64
    attached to Loopback0
2001:401:1:2::/64
    nexthop ::FFFF:192.168.1.33
    fast tag rewrite with Se0/0.111, point2point, tags imposed: {21
27}
2001:401:111:1::1/128
    Receive
2001:401:111:1::/64
    attached to Serial0/0.101
2001:401:112:1::/64
    nexthop ::FFFF:192.168.1.33
    fast tag rewrite with Se0/0.111, point2point, tags imposed: {21
28}
FE80::/10
    Receive
FF00::/8
    Receive
```

- On R3, use the **sh ipv6 rou** command to verify that the routes from R4 (that is, the BGP routes) are present.

```
R3#sh ipv6 rou
IPv6 Routing Table - 11 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2
C   2001:401:1:1::/64 [0/0]
    via ::, Loopback0
L   2001:401:1:1::1/128 [0/0]
    via ::, Loopback0
B   2001:401:1:2::/64 [200/0]
    via ::FFFF:192.168.1.33, IPv6-mpls
C   2001:401:111:1::/64 [0/0]
```

```

        via ::, Serial0/0.101
L   2001:401:111:1::1/128 [0/0]
        via ::, Serial0/0.101
B   2001:401:112:1::/64 [200/0]
        via ::FFFF:192.168.1.33, IPv6-mpls
L   FE80::/10 [0/0]
        via ::, Null0
L   FF00::/8 [0/0]
        via ::, Null0 0

```

- On R4, use the **sh bgp ipv6 labels** command to verify the BGP label that is assigned to the routes.

```

R4#sh bgp ipv6 labels
Network                Next Hop          In label/Out label
2001:401:1:1::/64
                        ::FFFF:192.168.1.17
                        nolabel/27
2001:401:1:2::/64
                        ::
                        27/nolabel
2001:401:111:1::/64
                        ::FFFF:192.168.1.17
                        nolabel/28
2001:401:112:1::/64
                        ::
                        28/nolabel
                        ::
                        29/nolabel

```

- On R4, use the **sh ipv6 cef** command to verify the forwarding for the routes.

```

R4#sh ipv6 cef
2001:401:1:1::/64
    nexthop ::FFFF:192.168.1.17
    fast tag rewrite with Se0/0.111, point2point, tags imposed: {21
27}
2001:401:1:2::1/128
    Receive
2001:401:1:2::/64
    attached to Loopback0
2001:401:111:1::/64
    nexthop ::FFFF:192.168.1.17
    fast tag rewrite with Se0/0.111, point2point, tags imposed: {21
28}
2001:401:112:1::1/128
    Receive

```

```

2001:401:112:1::/64
    attached to Serial0/0.101
2001:401:121:1::/64
    nexthop ::FFFF:192.168.1.17
    fast tag rewrite with Se0/0.111, point2point, tags imposed: {21
29}
FE80::/10
    Receive
FF00::/8
    Receive

```

Task 5: Enable IPv6 Routing Between CE and PE Routers

The following steps enable you to establish IPv6 routing between the CE and PE routers. For this lab exercise, you will use BGP. It should be noted that static routing, Routing Information Protocol (RIP), or IPv6 OSPF could also be used.

Exercise Procedure

Complete the following commands:

In the existing BGP process on R5, complete the following steps:

- Step 1** Configure neighbor peering with the R7 (2001:401:111:1::1).
- Step 2** Activate neighbor peering under the IPv6 address family.

In the exiting BGP process on R7, complete the following steps:

Step 3 Configure neighbor peering with the R5 (2001:401:111:1::2).

Step 4 Activate neighbor peering under the IPv6 address family.

In the exiting BGP process on R6, complete the following steps:

Step 5 Configure neighbor peering with the R8 (2001:401:112:1::1).

Step 6 Activate neighbor peering under the IPv6 address family.

In the exiting BGP process on R8, complete the following steps:

Step 7 Configure neighbor peering with the R6 (2001:401:112:1::2).

Step 8 Activate neighbor peering under the IPv6 address family.

Exercise Verification

The exercise verification should be done by comparing your configuration additions to those in the Lab Exercise Answer Key.

If the hardware was available, here are typical outputs from the **show** commands:

- On R1, use the **sh ipv6 rou** command to verify that the routes from R2 (that is, the BGP routes) are present.

```
R1#sh ipv6 rou
IPv6 Routing Table - 15 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -
        OSPF ext 2
        .
        .
        .
B    2001:401:112:1::/64 [20/0]
```

```

        via FE80::7FF:FE53:835, Serial0/0.101
B    2001:401:112:2::/64 [20/0]
        via FE80::7FF:FE53:835, Serial0/0.101
B    2001:401:112:8::/64 [20/0]
        via FE80::7FF:FE53:835, Serial0/0.101
.
.
.

```

- On R2, use the **sh ipv6 rou** command to verify that the routes from R1 (that is, the BGP routes) are present.

```

R2#sh ipv6 rou
IPv6 Routing Table - 15 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
        U - Per-user Static route
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2
.
.
.
B    2001:401:111:1::/64 [20/0]
        via FE80::7FF:FE53:836, Serial0/0.101
B    2001:401:111:2::/64 [20/0]
        via FE80::7FF:FE53:836, Serial0/0.101
B    2001:401:111:8::/64 [20/0]
        via FE80::7FF:FE53:836, Serial0/0.101
.
.
.

```


Lab Exercise Answer Key

Lab Exercise 2-1: Implementing CSC in a Native IP Environment

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Implement the Initial Configurations on the Routers

See configurations in the lab section.

Task 2: Configure the PE Router of the Backbone Carrier

■ Team A (R7)

```
ip vrf Yellow
rd 1:10
route-target export 1:10
route-target import 1:10

interface Serial0/0.102 point-to-point
ip vrf forwarding Yellow
ip address 150.1.11.18 255.255.255.240
frame-relay interface-dlci 102

router ospf 1 vrf Yellow
log-adjacency-changes
redistribute bgp 65001 metric-type 1 subnets
network 150.1.11.18 0.0.0.0 area 0

router bgp 65001

address-family ipv4 vrf Yellow
redistribute ospf 1
no auto-summary
no synchronization
exit-address-family
```

■ Team B (R8)

```
ip vrf Yellow
rd 1:10
route-target export 1:10
route-target import 1:10

interface Serial0/0.102 point-to-point
ip vrf forwarding Yellow
ip address 150.1.12.18 255.255.255.240
frame-relay interface-dlci 102

router ospf 1 vrf Yellow
log-adjacency-changes
redistribute bgp 65001 metric-type 1 subnets
network 150.1.12.18 0.0.0.0 area 0

router bgp 65001

address-family ipv4 vrf Yellow
```

```
redistribute ospf 1
no auto-summary
no synchronization
exit-address-family
```

Task 3: Configure the CE Router of the Customer Carrier

■ R5 (Team A)

```
router ospf 1
net 150.1.11.17 0.0.0.0 area 0

router bgp 65010
neighbor 10.2.12.49 remote-as 65010
neighbor 10.2.12.49 update-source Loopback0
```

■ R6 (Team B)

```
router ospf 1
net 150.1.12.17 0.0.0.0 area 0

router bgp 65010
neighbor 10.1.12.49 remote-as 65010
neighbor 10.1.12.49 update-source Loopback0
```

Lab Exercise 2-2: Implementing CSC Using LDP Label Distribution

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configure the PE Routers of the Backbone Carrier (R7 and R8)

■ Team A (R7)

```
interface Serial0/0.102 point-to-point
mpls ip
mpls label protocol ldp

interface Loopback100
ip vrf forwarding Yellow
ip address 10.1.13.49 255.255.255.255
```

■ Team B (R8)

```
interface Serial0/0.102 point-to-point
mpls ip
mpls label protocol ldp

interface Loopback100
ip vrf forwarding Yellow
ip address 10.2.13.49 255.255.255.255
```

Task 2: Configure the CE Routers of the Customer Carriers (R1, R5, R2, and R6)

■ R5 (Team A)

```
mpls ldp router-id loopback 0
interface Serial0/0.101 point-to-point
mpls label protocol tdp
mpls ip

interface Serial0/0.102 point-to-point
mpls label protocol tdp
mpls ip
```

■ R1 (Team A)

```
tag-switching tdp router-id lo0

interface Serial0/0.101 point-to-point
mpls ip
```

■ R6 (Team B)

```
interface Serial0/0.101 point-to-point
mpls label protocol tdp
tag-switching ip

interface Serial0/0.102 point-to-point
mpls label protocol ldp
tag-switching ip
```

■ R2 (Team B)

```
tag-switching tdp router-id lo0

interface Serial0/0.101 point-to-point
tag-switching ip
```

Task 3: Change the BGP Sessions of the Customer Carrier

- R1 (Team A)

```
router bgp 65010
neighbor 10.2.11.49 remote-as 65010
neighbor 10.2.11.49 update-source loo 0
```

Note	Remove the neighbor statement with the following command.

```
no neighbor 10.1.12.49 remote-as 65010
```

- R5 (Team A)

```
no router bgp 65010
```

- R2 (Team B)

```
router bgp 65010
neighbor 10.1.11.49 remote-as 65010
neighbor 10.1.11.49 update-source loo 0
```

Note	Remove the neighbor statement with the following command.

```
no neighbor 10.2.12.49 remote-as 65010
```

- R6 (Team B)

```
no router bgp 65010
```

Lab Exercise 2-3: Implementing CSC Using MP-BGP Label Distribution

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configure MP-BGP Label Distribution

■ Team A (R7)

```
interface Serial0/0.102 point-to-point
```

Note	Remove the MPLS and LDP statements with the following commands.

```
no mpls ip
no mpls label protocol ldp

router bgp 65001

address-family ipv4 vrf Yellow
neighbor 150.1.11.17 remote-as 65010
neighbor 150.1.11.17 activate
neighbor 150.1.11.17 as-override
neighbor 150.1.11.17 send-label
```

■ Team A (R5)

```
interface Serial0/0.102 point-to-point
```

Note	Remove the MPLS and LDP statements with the following commands.

```
no mpls ip
no mpls label protocol ldp
router bgp 65010
neighbor 150.1.11.18 remote-as 65001
address-family ipv4
redistribute connected
redistribute ospf 1 metric 4 match internal
neighbor 150.1.11.18 activate
neighbor 150.1.11.18 send-label

router ospf 1
redistribute bgp 65010 metric 4 subnets
```

■ Team B (R8)

```
interface Serial0/0.102 point-to-point
```

Note Remove the MPLS and LDP statements with the following commands.

```
no mpls ip
no mpls label protocol ldp

router bgp 65001
 address-family ipv4 vrf Yellow
  neighbor 150.1.12.17 remote-as 65010
  neighbor 150.1.12.17 activate
  neighbor 150.1.12.17 as-override
  neighbor 150.1.12.17 send-label
```

■ Team B (R2)

```
interface Serial0/0.102 point-to-point
```

Note Remove the MPLS and LDP statements with the following commands.

```
no mpls ip
no mpls label protocol ldp

router bgp 65010
  neighbor 150.1.12.18 remote-as 65001
  address-family ipv4
    redistribute connected
    redistribute ospf 1 metric 4 match internal
  neighbor 150.1.12.18 activate
  neighbor 150.1.12.18 send-label

router ospf 1
  redistribute bgp 65010 metric 4 subnets
```

Lab Exercise 2-4: Implementing CSC Using MPLS VPNs

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configure MPLS VPN Support

■ Team A (R1)

```
ip vrf Green
 rd 100:1
  route-target export 100:1
  route-target import 100:1

ip vrf Orange
 rd 100:2
  route-target export 100:2
  route-target import 100:2

interface Loopback11
 ip vrf forwarding Green
 ip address 172.11.1.16 255.255.255.0
```

```
interface Loopback12
 ip vrf forwarding Orange
 ip address 172.12.1.16 255.255.255.0

router bgp 65010
```

Note	Remove the network statements with the following commands.

```
no network 172.11.1.0 mask 255.255.255.0
no network 172.12.1.0 mask 255.255.255.0
```

```
address-family vpnv4
 neighbor 10.2.11.49 activate

address-family ipv4 vrf Orange
 network 172.12.1.0 mask 255.255.255.0

address-family ipv4 vrf Green
 network 172.11.1.0 mask 255.255.255.0
```

■ Team B (R2)

```
ip vrf Green
 rd 100:1
 route-target export 100:1
 route-target import 100:1

ip vrf Orange
 rd 100:2
 route-target export 100:2
 route-target import 100:2

interface Loopback11
 ip vrf forwarding Green
 ip address 172.11.2.16 255.255.255.0

interface Loopback12
 ip vrf forwarding Orange
 ip address 172.12.2.16 255.255.255.0

router bgp 65010
```

Note	Remove the network statements with the following commands.

```
no network 172.11.2.0 mask 255.255.255.0
no network 172.12.2.0 mask 255.255.255.0
```

```
address-family vpnv4
 neighbor 10.1.11.49 activate

address-family ipv4 vrf Orange
 network 172.12.2.0 mask 255.255.255.0

address-family ipv4 vrf Green
```

Lab Exercise 4-1: Implementing MPLS Traffic Engineering

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod). There are no answers for tasks 1 and 2.

Task 3: Configure MPLS Traffic Engineering Support on Your Routers

R5 and R6:

```
mpls traffic-eng tunnels
mpls traffic-eng logging lsp setups
mpls traffic-eng reoptimize timers frequency 30

router isis
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-1

int s0/0
 ip rsvp bandwidth 128 128

int s0/0.111
 ip rsvp bandwidth 128 128
mpls traffic-eng tunnels
```

Team A

R7:

```
mpls traffic-eng tunnels
mpls traffic-eng logging lsp setups
mpls traffic-eng reoptimize timers frequency 30

router isis
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-1

int s0/0
 ip rsvp bandwidth 128 128

int s0/0.111
mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128

int s0/0.112
mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128

int s0/0.2wg1
mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128
```


Team B

R8:

```
mpls traffic-eng tunnels
mpls traffic-eng logging lsp setups
mpls traffic-eng reoptimize timers frequency 30

router isis
  mpls traffic-eng router-id Loopback0
  mpls traffic-eng level-1

int s0/0
  ip rsvp bandwidth 128 128

int s0/0.111
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128

int s0/0.112
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128

int s0/0.2wg2
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128
```

Task 4: Configure MPLS Traffic Engineering Tunnels

Team A

R5:

```
ip explicit-path name via_R9 enable
  next-address 192.168.wg.81
  next-address 192.168.10.17
  next-address 192.168.wg.97
  next-address 192.168.wg.33

int tunnel1
  ip unnumbered Loopback0
  tunnel destination 192.168.wg.33
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng priority 0 0
  tunnel mpls traffic-eng bandwidth 100
  tunnel mpls traffic-eng path-option 1 explicit name via_R9
  tunnel mpls traffic-eng autoroute announce
  mpls ip
  exit
```

Team B

R6:

```
ip explicit-path name via_R9 enable
  next-address 192.168.wg.97
  next-address 192.168.10.17
  next-address 192.168.wg.81
  next-address 192.168.wg.17
```

```

int tunnel1
 ip unnumbered Loopback0
 tunnel destination 192.168.wg.17
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 0 0
 tunnel mpls traffic-eng bandwidth 100
 tunnel mpls traffic-eng path-option 1 explicit name via_R9
 mpls ip
 exit

ip route 192.168.wg.17 255.255.255.255 tunnel 1

```

Lab Exercise 4-2: Implementing Link Protection for MPLS Traffic Engineering

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Configure the Backup Tunnels

Team A

R7:

```

ip explicit-path name Backup_Tunnel enable
 next-address 192.168.wg.97

int tunnel10
 ip unnumbered Loopback0
 tunnel destination 192.168.wg.97
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 100
 tunnel mpls traffic-eng path-option 1 explicit name Backup_Tunnel

int s0/0.2wg1
 mpls traffic-eng backup tunnel 10

```

R5:

```

interface Tunnel1
 tunnel mode mpls traffic-eng fast-reroute

```

Team B

R8:

```

ip explicit-path name Backup_Tunnel enable
 next-address 192.168.wg.81

int tunnel10
 ip unnumbered Loopback0
 tunnel destination 192.168.wg.81
 tunnel mode mpls traffic-eng

```

```

tunnel mpls traffic-eng priority 7 7
tunnel mpls traffic-eng bandwidth 100
tunnel mpls traffic-eng path-option 1 explicit name Backup_Tunnel

int s0/0.2wg1
  mpls traffic-eng backup tunnel 10

```

R6:

```

interface Tunnel1
  tunnel mode mpls traffic-eng fast-reroute

```

Lab Exercise 5-1: Implementing MPLS QoS

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task 1: Implement an Inbound PE QoS Policy

Configuration on R5 (Team A) and R6 (Team B):

```

class-map match-all PREMIUM-IP
  match ip dscp ef
class-map match-all BUSINESS-IP
  match ip dscp af21 af22 af23

policy-map IN-POLICY
  class PREMIUM-IP
    police cir 1280000 bc 32000 be 32000
    conform-action set-mpls-exp-topmost-transmit 5
    exceed-action drop
  class BUSINESS-IP
    police cir 2200000 bc 550000 be 550000
    conform-action set-mpls-exp-topmost-transmit 2
    exceed-action set-mpls-exp-topmost-transmit 1
  class class-default
    set mpls experimental imposition 0

int s0/0.101
  service-policy input IN-POLICY
int s0/0.102
  service-policy input IN-POLICY
int s0/0.111
  service-policy input IN-POLICY

```

Task 2: Implement an Outbound PE QoS Policy

Configuration on R5 (Team A) and R6 (Team B):

```

class-map match-all BUSINESS-MPLS
  match mpls experimental topmost 1 2
class-map match-all PREMIUM-MPLS
  match mpls experimental topmost 5

policy-map OUT-POLICY
  class PREMIUM-MPLS
    priority percent 5
  class BUSINESS-MPLS
    bandwidth percent 50

```

```

random-detect
random-detect precedence 1    100    300    100
random-detect precedence 2    300    500    10
class class-default
  fair-queue
  random-detect

map-class frame-relay FRAME-OUT
  frame-relay cir 128000
  frame-relay mincir 128000
  service-policy output OUT-POLICY

int s0/0
  frame-relay traffic-shaping

int s0/0.111
  frame-relay class FRAME-OUT

```

Task 3: Implement an Outbound QoS Policy on the Local P Router Toward a Remote P Router

Configuration on R7 (Team A) and R8 (Team B):

```

class-map match-all PREMIUM-IP
  match ip dscp ef
  class-map match-all BUSINESS-IP
  match ip dscp af21 af22 af23

policy-map IN-POLICY
  class PREMIUM-IP
    police cir 1280000 bc 32000 be 32000
    conform-action set-mpls-exp-topmost-transmit 5
    exceed-action drop
  class BUSINESS-IP
    police cir 22000000 bc 550000 be 550000
    conform-action set-mpls-exp-topmost-transmit 2
    exceed-action set-mpls-exp-topmost-transmit 1
  class class-default
    set mpls experimental imposition 0

class-map match-all BUSINESS-MPLS
  match mpls experimental topmost 1 2
class-map match-all PREMIUM-MPLS
  match mpls experimental topmost 5

policy-map OUT-POLICY
  class PREMIUM-MPLS
    priority percent 5
  class BUSINESS-MPLS
    bandwidth percent 50
  random-detect
  random-detect precedence 1    100    300    100
  random-detect precedence 2    300    500    10
class class-default
  fair-queue
  random-detect

map-class frame-relay FRAME-OUT
  frame-relay cir 128000
  frame-relay mincir 128000

```

```

service-policy output OUT-POLICY

policy-map IN-POLICY
class PREMIUM-IP
  police cir 1280000 bc 32000 be 32000
  conform-action set-mpls-exp-topmost-transmit 5
  exceed-action drop
  class BUSINESS-IP
  police cir 22000000 bc 550000 be 550000
  conform-action set-mpls-exp-topmost-transmit 2
  exceed-action set-mpls-exp-topmost-transmit 1
  class class-default
  set mpls experimental imposition 0
policy-map OUT-POLICY
class PREMIUM-MPLS
  priority percent 5
  class BUSINESS-MPLS
  bandwidth percent 50
  random-detect
  random-detect precedence 1    100    300    100
  random-detect precedence 2    300    500    10
class class-default
  fair-queue
  random-detect
map-class frame-relay FRAME-OUT
frame-relay cir 128000
frame-relay mincir 128000
service-policy output OUT-POLICY
int s0/0
  frame-relay traffic-shaping
int s0/0.112
  frame-relay class FRAME-OUT
int s0/0.212
  frame-relay class FRAME-OUT

```

On R7 use this config:

```

int s0/0.212 wg1
frame-relay class FRAME-OUT

```

On R8 use this config:

```

int s0/0.212 wg2
frame-relay class FRAME-OUT

```

Task 4: Implement an Outbound QoS Policy on the P Router Toward the PE Router

Configuration on R7 (Team B) and R8 (Team B):

```

policy-map IN-POLICY
class PREMIUM-IP
  police cir 1280000 bc 32000 be 32000
  conform-action set-mpls-exp-topmost-transmit 5
  exceed-action drop
  class BUSINESS-IP
  police cir 22000000 bc 550000 be 550000
  conform-action set-mpls-exp-topmost-transmit 2
  exceed-action set-mpls-exp-topmost-transmit 1
  class class-default
  set mpls experimental imposition 0
policy-map OUT-POLICY
class PREMIUM-MPLS

```

```

    priority percent 5
    class BUSINESS-MPLS
    bandwidth percent 50
    random-detect
    random-detect precedence 1    100    300    100
    random-detect precedence 2    300    500    10
class class-default
    fair-queue
    random-detect
map-class frame-relay FRAME-OUT
frame-relay cir 128000
frame-relay mincir 128000
service-policy output OUT-POLICY
int s0/0
    frame-relay traffic-shaping
int s0/0.111
    frame-relay class FRAME-OUT

```

Task 5: Implement an Outbound QoS Policy on the PE Router Toward the CE Router

Configuration on R5 (Team A) and R6 (Team B):

```

class-map match-all BUSINESS-IP
    match ip dscp af21 af22 af23
class-map match-all PREMIUM-IP
    match ip dscp ef
    match ip dscp af21 af22 af23
class-map match-all PREMIUM-IP
    match ip dscp ef
policy-map IP-OUT
class PREMIUM-IP
    priority percent 50
class BUSINESS-IP
    bandwidth percent 25
    random-detect dscp-based
    random-detect dscp 18 30 40 1
    random-detect dscp 20 20 30 1
    random-detect dscp 22 10 20 1
class class-default
    fair-queue
    random-detect dscp-based
map-class frame-relay PE-OUT
frame-relay cir 128000
frame-relay mincir 128000
service-policy output IP-OUT
int s0/0
    frame-relay traffic-shaping
int s0/0.101
    frame-relay class PE-OUT
int s0/0.102
    frame-relay class PE-OUT

```

Lab Exercise 5-2: Implementing MPLS QoS and DS-TE

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod).

Task1: Update Existing Tunnel Parameters

Configuration on R5 (Team A) and R6 (Team B):

```
int tunnel1
 tunnel mpls traffic-eng bandwidth 64
 no tunnel mpls traffic-eng autoroute announce
 exit
```

Task 2: Implement DS-TE on the Interfaces

Configuration on R5 (Team A) and R6 (Team B):

```
int s0/0
 ip rsvp bandwidth 128 sub-pool 10
int s0/0.111
 ip rsvp bandwidth 128 sub-pool 10
```

Configuration on R7 (Team A) and R8 (Team B):

```
int s0/0
 ip rsvp bandwidth 128 sub-pool 10
int s0/0.111
 ip rsvp bandwidth 128 sub-pool 10
int s0/0.112
 ip rsvp bandwidth 128 sub-pool 10
int s0/0.2pop
 ip rsvp bandwidth 128 sub-pool 10
```

Task 3: Create a DS-TE tunnel

Configuration on R5 (Team A):

```
ip explicit-path name R5_R6 enable
 next-address 192.168.wg.81
 next-address 192.168.wg.97
 next-address 192.168.wg.33
int tunnel20
 ip unnumberedLoopback0
 tunnel destination 192.168.wg.33
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 7
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng autoroute metric relative 4
 tunnel mpls traffic-eng path-option 1 explicit name R5_R6
 mpls ip
 exit
```

Configuration on R6 (Team B):

```
ip explicit-path name R6_R5 enable
 next-address 192.168.wg.97
 next-address 192.168.wg.81
 next-address 192.168.wg.17
```

```

int tunnel20
 ip unnumbered Loopback0
 tunnel destination 192.168.wg.17
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 7
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng autoroute metric relative 4
 tunnel mpls traffic-eng path-option 1 explicit name R6_R5
 mpls ip
 #exit

```

Task 4: Assign the Traffic to the Tunnels

Team A—VRF Green

Configuration on R5:

```
ip route 10.1.2.64 255.255.255.255 Tunnel20
```

Configuration on R6:

```
ip route 10.1.1.64 255.255.255.255 Tunnel20
```

Team B—VRF Orange

Configuration on R5:

```
ip route 10.1.2.64 255.255.255.255 Tunnel20
```

Configuration on R6:

```
ip route 10.1.1.64 255.255.255.255 Tunnel20
```

Lab Exercise 6-1: Implementing a Port-to-Port FRoMPLS Configuration (Case Study)

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```

Router: R5
ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip
!
interface Loopback0
 ip address 192.168.wg.17 255.255.255.255

```



```

interface Serial0/0
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
  description ***** DLCI to the P router *****
  ip address 192.168.wg.49 255.255.255.240
  frame-relay interface-dlci 111
  mpls label protocol ldp
  tag-switching ip
!
interface Serial0/1
  description **** link to the Frame Relay Switch *****
  encapsulation hdlc
  mpls l2transport route 192.168.33 123 sequence
!
router isis
  net 49.0001.0000.0000.0111.00
  is-type level-2-only
  metric-style wide
!
router bgp 65001
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.100.129 remote-as 65001
  neighbor 192.168.100.129 update-source Loopback0
  no auto-summary

Router: R6
ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip
!
interface Loopback0
  ip address 192.168.wg.32 255.255.255.255

interface Serial0/0

```

```

no ip address
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
description ***** DLCI to the P router *****
ip address 192.168.wg.49 255.255.255.240
frame-relay interface-dlci 111
mpls label protocol ldp
tag-switching ip
!
interface Serial0/1
description **** link to the CE router for customer A ****
ip address 150.wg.pop.18 255.255.255.240
encapsulation hdlc
mpls l2transport route 192.168.wg.17 123 sequence
!
Router Isis
Net 49.0001.0000.0000.0111.00
is-type level-2-only
metric-style wide
!
router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 192.168.100.129 remote-as 65001
neighbor 192.168.100.129 update-source Loopback0
no auto-summary

```

Lab Exercise 6-2: Implementing a DLCI-to-DLCI FRoMPLS Configuration (Case Study)

When you complete this lab exercise, your router configuration will be similar to the following, with differences that are specific to your pod.

```
Router: R5
ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip

frame-relay switching

!
interface Loopback0
 ip address 192.168.wg.17 255.255.255.255

interface Serial0/0
 no ip address
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
 description ***** DLCI to the P router *****
 ip address 192.168.wg.49 255.255.255.240
 frame-relay interface-dlci 111
 mpls label protocol ldp
 tag-switching ip
!
interface Serial1/0
 description **** link to the CE router for customer A ****
 encapsulation frame-relay IETF
 load-interval 30
 clockrate 128000
 frame-relay lmi-type cisco
 frame-relay intf-type dce
!
interface Serial1/1
 description **** link to the Customer B's CE router ****
 encapsulation frame-relay IETF
 load-interval 30
 clockrate 128000
 frame-relay lmi-type cisco
 frame-relay intf-type dce

router isis
 net 49.0001.0000.0000.0111.00
 is-type level-2-only
 metric-style wide
!
```

```

router bgp 65001
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.100.129 remote-as 65001
  neighbor 192.168.100.129 update-source Loopback0
  no auto-summary

connect fr101 Serial1/0 101 l2transport
  xconnect 192.168.wg.33 encapsulation mpls
connect fr102 Serial1/1 102 l2transport
  xconnect 192.168.wg.33 encapsulation mpls

Router: R6
ip cef
mpls label protocol ldp
mpls ldp router-id Loopback0 force
mpls ip

frame-relay switching

!
interface Loopback0
  ip address 192.168.wg.33 255.255.255.255

interface Serial0/0
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
  description ***** DLCI to the P router *****
  ip address 192.168.wg.49 255.255.255.240
  frame-relay interface-dlci 111
  mpls label protocol ldp
  tag-switching ip
!
interface Serial1/0
  description **** link to the CE router for customer A ****
  encapsulation frame-relay IETF
  load-interval 30
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce

!
interface Serial1/1
  description **** link to the Customer B's CE router ****
  encapsulation frame-relay IETF
  load-interval 30
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce

```

```

!
router isis
 net 49.0001.0000.0000.0111.00
 is-type level-2-only
 metric-style wide
!
router bgp 65001
 no synchronization
 bgp log-neighbor-changes
 neighbor 192.168.100.129 remote-as 65001
 neighbor 192.168.100.129 update-source Loopback0
 no auto-summary
!

connect fr101 Serial1/0 101 l2transport
 xconnect 192.168.wg.17 encapsulation mpls
connect fr102 Serial1/1 102 l2transport
 xconnect 192.168.wg.17 encapsulation mpls

```

Lab Exercise 7-1: Implementing Cisco 6PE Support (Case Study)

When you complete this lab exercise, your router will have the following incremental configuration (with differences that are specific to your pod). There are no answers for task 1.

Task 1: Enable IPv6 on the PE Routers

Team A

R7

```

ipv6 unicast-routing
ipv6 cef
mpls ipv6 source-interface Loopback0

interface Loopback0
 ip address 192.168.1.17 255.255.255.255
 ip router isis
 ipv6 address 2001:401:1:1::1/64

interface Serial0/0.101 point-to-point
 ipv6 address 2001:401:111:1::1/64

```

Team B

R8

```

ipv6 unicast-routing
ipv6 cef
mpls ipv6 source-interface Loopback0

interface Loopback0
 ip address 192.168.1.33 255.255.255.255
 ip router isis
 ipv6 address 2001:401:1:2::1/64

interface Serial0/0.101 point-to-point

```

ipv6 address 2001:401:112:1::1/64

Task 2: Establish IBGP Connectivity Between PE Routers

Team A

R7

```
router bgp 65001
  no synchronization
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  neighbor 192.168.1.34 remote-as 65001
  neighbor 192.168.1.34 update-source Loopback0
  no auto-summary

address-family ipv6
  neighbor 192.168.10.17 activate
  neighbor 192.168.10.17 send-label
exit-address-family
```

Team B

R8

```
router bgp 65001
  no synchronization
  no bgp default ipv4-unicast
  bgp log-neighbor-changes
  neighbor 192.168.1.17 remote-as 65001
  neighbor 192.168.1.17 update-source Loopback0
  no auto-summary

address-family ipv6
  neighbor 192.168.1.17 activate
  neighbor 192.168.1.17 send-label
exit-address-family
```

Task 3: Enable IPv6 Routing Between PE Routers

Team A

R7

```
router bgp 65001
  address-family ipv6
    network 2001:401:1:1::/64
    network 2001:401:111:1::/64
```

Team B

R6

```
router bgp 65001
  address-family ipv6
    network 2001:401:1:2::/64
    network 2001:401:112:1::/64
```

Task 4: Enable IPv6 Routing Between CE and PE Routers

Team A

R5

```
router bgp 65111
 neighbor 2001:401:111:1::1 remote-as 65001

 address-family ipv6
 neighbor 2001:401:111:1::1 activate
```

R7

```
router bgp 65001
 neighbor 2001:401:111:1::2 remote-as 65111

 address-family ipv6
 neighbor 2001:401:111:1::2 activate
```

Team B

R6

```
router bgp 65112
 neighbor 2001:401:112:1::1 remote-as 65001

 address-family ipv6
 neighbor 2001:401:112:1::1 activate
```

R8

```
router bgp 65001
 neighbor 2001:401:112:1::2 remote-as 65112

 address-family ipv6
 neighbor 2001:401:112:1::2 activate
```

Course Administration Guide

Overview

Service providers are faced with many challenges as they attempt to deal with customer demand, including a continuing need for value-added services. Conventional IP packet forwarding has several limitations. Traditional Layer 2 services and traffic engineering (TE) lack the flexibility to meet the demands of the modern network. More and more service providers realize that something else is needed. Not only must they be concerned with protecting their existing infrastructure, but they must find ways to generate new services that are not currently supportable with existing technologies.

Multiprotocol Label Switching (MPLS) is a high-performance method for forwarding packets through a network. It provides a foundation for a series of applications that integrate the performance and traffic management capabilities of the data link layer (Layer 2) with the scalability and flexibility of network layer (Layer 3) routing. This integration allows service providers the ability to support “value-added” features that are critical for their networks.

Implementing Cisco MPLS Traffic Engineering and Other Features (MPLST) v2.0 is an instructor-led training (ILT) course that extends the knowledge and skills that the learner has obtained from the MPLS Cisco CCIP™ certification track. The course duration is 5 days in length, using a standard ILT format. The course reflects the latest developments in network design and technologies, and uses real-world scenarios to help reinforce key objectives.

The MPLST course is recommended training for individuals who are seeking advanced MPLS-based knowledge and skills. The course will enable learners to gather information ranging from the technology basics to more updated features and functions such as MPLS Traffic Engineering (MPLS-TE), Fast Reroute (FRR), and Any Transport over MPLS (AToM). The focus of the course is on the technology issues of MPLS from a service provider perspective and how to configure some of the MPLS features and functions in an existing routed environment.

Course Version

This course is the first version of the MPLST course. It carries a label of “version 2.0” to keep it in synchrony with the current version of the *Implementing Cisco MPLS* (MPLS) v2.0 course.

Objectives

Upon completing this course, the learner will be able to meet these overall objectives:

- Identify the MPLS peer-to-peer architecture and explain label allocation, routing update distribution, and packet forwarding as they relate to this architecture
- Use the Cisco IOS commands that are required to successfully configure, monitor, and troubleshoot service provider support using the MPLS Carrier Supporting Carrier service
- Identify the MPLS Traffic Engineering architecture and explain how MPLS implements traffic engineering, establishes the constraint-based path, and assigns traffic to traffic trunks
- Use the Cisco IOS commands that are required to successfully configure, monitor, and troubleshoot MPLS Traffic Engineering
- Use the Cisco IOS commands that are required to successfully implement a defined SLA using MPLS QoS services
- Use the Cisco IOS commands that are required to successfully configure, monitor, and troubleshoot Layer 2 services using the Any Transport over MPLS service
- Use the Cisco IOS commands that are required to successfully configure, monitor, and troubleshoot IPv6 support in an MPLS environment

Target Audience

The target audience for the MPLST course consists of network administrators and technicians responsible for implementing and troubleshooting basic IP multicast-enabled networks within a single domain, as well as Cisco CCNP[®], CCIP, CCDP[®], and CCIE[®] routing and switching candidates.

The course is also targeted at pre- and post-sales technical engineers who are responsible for designing, implementing, and troubleshooting MPLS networks or solutions that are based on MPLS technology.

Learner Skills and Knowledge

The prerequisite for learners who wish to participate in the MPLST course is the CCIP certification or an equivalent level of working knowledge and experience.

Practical experience with deploying and operating networks that are based on Cisco network devices and Cisco IOS software is strongly recommended.

Course Instruction Details

This section provides the information that you need to prepare the course materials and set up the classroom environment.

Instructor Certification Requirements

To teach this course, instructors must have attended the following training or completed the following requirements:

- Attended an MPLST Train-the-Trainer session

or

- Attended an MPLST ILT class

Note Submit questions concerning instructor certification to icad@external.cisco.com.

Classroom Reference Materials

These items should be available for the learner during presentation of the course:

- Access to the Cisco website

Class Environment

This information describes class size and classroom setup:

- The class is designed for 4 to 16 students.
- The class requires a standard lecture classroom with whiteboard and PowerPoint projection system.
- The classroom should contain 8 Windows-based PCs with Internet capability for access to remote lab equipment

Course Flow

This is the *suggested* course schedule. You may make adjustments based on the skills, knowledge, and preferences of the learners in attendance. The presentation of all topics *for noncertification offerings* is optional, but you are encouraged to use them because they are designed to reinforce the lesson concepts and ensure that learners apply some of the concepts.

Day 1: MPLS VPN Review and Carrier Supporting Carrier

	8:30–9:30	Course Introduction
MPLS VPN Review	9:30–12:00	Basic MPLS Concepts
		Label Distribution in Frame-Mode MPLS
		Configuring MPLS
		MPLS VPN Routing Model and Packet Forwarding
		Configuring Simple MPLS VPNs
	12:00–1:00	Lunch
Carrier Supporting Carrier	1:00–3:00	Carrier's Carrier Overview
		Implementing the CSC Models
Lab Exercise	3:00–5:00	Lab Exercise 2-1: Implementing CSC in a Native IP Environment
	5:00	Day ends

Day 2: Carrier Supporting Carrier and MPLS Traffic Engineering Technology

Lab Exercise	8:30–10:00	Lab Exercise 2-2: Implementing CSC Using LDP Label Distribution
Lab Exercise	10:00–11:00	Lab Exercise 2-3: Implementing CSC Using MP-BGP Label Distribution
Lab Exercise	11:00–12:00	Lab Exercise 2-4: Implementing CSC Using MPLS VPNs
	12:00–1:00	Lunch
MPLS Traffic Engineering Technology	1:00–5:00	Traffic Engineering Concepts
		MPLS Traffic Engineering Components
		Constraint-Based Path Computation
		Path Setup and Maintenance
		Assigning Traffic to Traffic Tunnels
	5:00	Day ends

Day 3: Configuring MPLS Traffic Engineering

	8:30–9:00	Review
Configuring MPLS Traffic Engineering	9:00–12:00	Configuring MPLS Traffic Engineering on Cisco IOS Platforms Advanced MPLS-TE Path Selection Advanced MPLS-TE Link Protection Advanced MPLS-TE Bandwidth Control MPLS-TE Interarea Support Monitoring and Troubleshooting MPLS-TE
	12:00–1:00	Lunch
Lab Exercise	1:00–4:00	Lab Exercise 4-1: Implementing MPLS Traffic Engineering
Lab Exercise	4:00–5:00	Lab Exercise 4-2: Implementing Link Protection for MPLS Traffic Engineering
	5:00	Day ends

Day 4: MPLS Quality of Service and Any Transport over MPLS

	8:30–9:00	Review
MPLS Quality of Service	9:00–10:30	QoS Models MPLS Support for DiffServ Configuring MPLS QoS QoS in MPLS Applications
Lab Exercise	10:30–12:00	Lab Exercise 5-1: Implementing MPLS QoS
	12:00–1:00	Lunch
Lab Exercise	1:00–2:00	Lab Exercise 5-2 Implementing MPLS QoS and DS-TE
Any Transport over MPLS	2:00–3:00	Introduction to Any Transport over MPLS Configuring AToM on Cisco IOS Platforms Monitoring AToM on Cisco IOS Platforms
Lab Exercise	3:00–4:00	Lab Exercise 6-1: Implementing a Port-to-Port FRoMPLS Configuration (Case Study)
Lab Exercise	4:00–5:00	Lab Exercise 6-2: Implementing a DLCI-to-DLCI FRoMPLS Configuration (Case Study)
	5:00	Day ends

Day 5: MPLS IPv6 Support

	8:30–9:00	Review
MPLS IPv6 Support	9:00–10:30	Review of IPv6 Implementing IPv6 over MPLS Monitoring IPv6 over MPLS
Lab Exercise	10:30–12:00	Lab Exercise 7-1: Implementing Cisco 6PE Support (Case Study)
	12:00–1:00	Lunch
	1:00–1:45	Wrap-up

High-Level Course Overview

This section provides a general idea of how the course is organized. The course contains these components:

- Course Introduction
- MPLS VPN Review
- Carrier Supporting Carrier
- MPLS Traffic Engineering Technology
- Configuring MPLS Traffic Engineering
- MPLS Quality of Service
- Any Transport over MPLS
- MPLS IPv6 Support
- Lab Guide
- Course Glossary

Detailed Course Overview

This in-depth outline of the course structure describes the objectives and content of each module.

Course Introduction

The Course Introduction provides learners with the course objectives, prerequisite learner skills and knowledge, and general administrative information. The Course Introduction presents the course roadmap and the icons used in the course illustrations and figures. This section also designates time for the learners to introduce themselves and describe their backgrounds, giving the instructor valuable information about the learners' knowledge and experience levels.

- Overview
- Course Objectives
- Cisco Certifications
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Flow Diagram
- Icons and Symbols
- Learner Introductions

Module 1: MPLS VPN Review

This module reviews the features of MPLS and MPLS Virtual Private Networks (VPNs) that were covered in the *Implementing Cisco MPLS (MPLS)* course.

Lesson 1-1: Basic MPLS Concepts

This lesson reviews the basic concepts and architecture of MPLS. The lesson provides information about some of the MPLS components and labels. It lays the foundation for subsequent lessons that cover the key areas, such as MPLS Traffic Engineering (MPLS-TE) and Virtual Private Networks (VPNs). The lesson includes these topics:

- Basic MPLS Concepts
- MPLS Architecture
- MPLS Labels
- Label Switch Routers
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the elements of the MPLS conceptual model
- Identify the elements of the MPLS architecture
- Identify the function of the MPLS label
- Identify the different types of label switch routers

Lesson 1-2: Label Distribution in Frame-Mode MPLS

This lesson reviews MPLS label allocation in a frame-mode environment, and how distribution of labels works in this environment. Also reviewed are penultimate hop popping (PHP) and how the MPLS data structures are built. The lesson includes these topics:

- Label Allocation in a Frame-Mode MPLS Environment
- Label Distribution and Advertisement
- Populating the LFIB
- Packet Propagation Across an MPLS Network
- Penultimate Hop Popping
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify how label allocation is performed in a frame-mode MPLS network
- Identify how labels are distributed in a frame-mode network
- Identify how the LFIB table is populated
- Identify packet propagation across an MPLS network
- Identify the function of penultimate hop popping

Lesson 1-3: Configuring MPLS

This lesson reviews the steps that are required to successfully implement MPLS. The lesson includes these topics:

- Configure IP CEF
- Configuring MPLS on a Frame-Mode Interface
- Configuring the MTU Size in Label Switching
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Enable IP CEF
- Enable MPLS on a frame-mode interface
- Configure the MTU size in label switching

Lesson 1-4: MPLS VPN Routing Model and Packet Forwarding

This lesson reviews the MPLS VPN architecture and operation. It details various customer edge-provider edge (CE-PE) routing options and Border Gateway Protocol (BGP) extensions (route targets and extended community attributes) that allow Interior Border Gateway Protocol (IBGP) to transport customer routes over a provider network. The MPLS VPN forwarding model is also covered together with its integration with core routing protocols. The lesson includes these topics:

- Propagation of Routing Information Across the P-Network
- Route Distinguishers
- Route Targets
- End-to-End Routing Information Flow
- MPLS VPNs and Packet Forwarding
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify how routing information is propagated across the P-network
- Identify the need for route distinguishers and how they are implemented
- Identify the need for route targets and how they are implemented
- Identify the end-to-end flow of routing updates in an MPLS VPN
- Describe MPLS VPN packet forwarding

Lesson 1-5: Configuring Simple MPLS VPNs

This lesson reviews the steps that are required to successfully implement a simple MPLS VPN. The lesson includes these topics:

- Configuring VRF Tables
- Configuring an MP-BGP Session Between PE Routers
- Configuring Routing Protocols Between PE and CE Routers
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Configure a VRF table
- Configure an MP-BGP session between PE routers
- Configure the routing protocols between PE and CE routers

Module Assessment

Module 2: Carrier Supporting Carrier

It is becoming common for larger service providers to act as the backbone for smaller service providers. This environment has classically been addressed with overlay VPNs. However, this solution has some serious scalability issues. Earlier attempts to address this environment with MPLS created several issues related to crossing service boundaries. The Carrier Supporting Carrier (CSC) feature was introduced to address these issues.

Lesson 2-1: Carrier's Carrier Overview

This lesson discusses the need for CSC and how MPLS provides this support. The lesson includes these topics:

- Carrier Supporting Carrier
- Benefits of CSC
- How CSC Works
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Describe the need for CSC
- Describe the benefits of CSC
- At an overview level, describe how CSC operates

Lesson 2-2: Implementing the CSC Models

This lesson identifies the different models that are used to implement CSC support and details the configuration steps that are necessary to implement each model. The lesson includes these topics:

- CSC Models
- Customer Carrier: Native IP
- Customer Carrier: MPLS and LDP
- Customer Carrier: MPLS and EBGp
- Customer Carrier: MPLS VPNs
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the CSC models
- Configure CSC to support customer carriers that are running native IP
- Configure CSC to support customer carriers that are running MPLS and LDP label distribution
- Configure CSC to support customer carriers that are running MPLS and MP-BGP label distribution
- Configure CSC to support customer carriers that are running MPLS VPNs

Module 3: MPLS Traffic Engineering Technology

This module on MPLS-TE discusses the requirement for TE in modern networks that must attain optimal resource utilization. The traffic-engineered tunnels provide a means of mapping traffic streams onto available networking resources in a way that prevents the overuse of subsets of networking resources while other subsets are underused. All the concepts and mechanics that support TE are presented, including tunnel path discovery with link-state protocols and tunnel path signaling with Resource Reservation Protocol (RSVP). Some of the advanced features of TE, such as autobandwidth and guaranteed bandwidth, are introduced as well.

Lesson 3-1: Traffic Engineering Concepts

This lesson describes the concepts that allow service providers to map traffic through specific routes to optimize network resources, especially the bandwidth. TE enables backbone networks to be engineered to deliver the total subscribed capacity to service provider customers more efficiently. The lesson includes these topics:

- What Is Traffic Engineering?
- Business Drivers for Traffic Engineering
- Congestion Avoidance and Traffic Engineering
- Traffic Engineering with a Layer 2 Overlay Model
- Traffic Engineering with a Layer 3 Model
- Traffic Engineering with the MPLS-TE Model
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the concepts behind TE
- Identify the major business drivers for implementing TE
- Identify congestion avoidance and how TE can reduce some congestion avoidance issues
- Identify how TE is implemented using a Layer 2 overlay model
- Identify how TE is implemented using a Layer 3 model
- Identify how TE is implemented using the MPLS-TE model

Lesson 3-2: MPLS Traffic Engineering Components

This lesson explains the components of MPLS-TE, such as traffic tunnels (along with associated characteristics and attributes), tunnel path discovery based on link-state protocols, and tunnel setup signaling with RSVP. The lesson includes these topics:

- Traffic Tunnels: Concepts
- Traffic Tunnels: Characteristics
- Traffic Tunnels: Attributes
- Network Links and Link Attributes
- Constraint-Based Path Computation
- Role of RSVP in Path Setup Procedures
- Forwarding Table Modifications
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify, at a conceptual level, how a traffic trunk functions
- Identify traffic trunk characteristics
- Identify traffic trunk attributes
- Identify the relation between network links and link attributes
- Identify the function of constraint-based path computation
- Identify the role of RSVP in path setup procedures
- Identify how using TE modifies the forwarding table mechanisms

Lesson 3-3: Constraint-Based Path Computation

This lesson describes the details of link attribute propagation with an Interior Gateway Protocol (IGP) protocol and constraint-based path computation. The lesson includes these topics:

- Constraint-Based Path Computation Attributes
- MPLS-TE Link Resource Attributes
- MPLS-TE Tunnel Attributes
- Implementing TE Policies with Affinity Bits
- Propagating MPLS-TE Link Attributes with Link-State Routing
- Constraint-Based Path Computation
- DiffServ-Aware TE Subpool
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the constraint-based path computation resource attributes
- Explain how MPLS-TE link attributes affect constraint-based path computation
- Explain how MPLS-TE tunnel attributes affect constraint-based path computation
- Identify how TE policies are implemented using affinity bits
- Identify how MPLS-TE link attributes are propagated with link-state routing protocols
- Describe the characteristics of constraint-based path computation as it relates to LSPs
- Describe the operation of the DiffServ-aware TE subpool as it relates to MPLS-TE

Lesson 3-4: Path Setup and Maintenance

This lesson describes the details of MPLS-TE tunnels, including path setup procedures and path maintenance. The lesson includes these topics:

- Path Setup
- RSVP Usage in Path Setup
- Hop-by-Hop Path Setup with RSVP
- Tunnel and Link Admission Control
- Path Rerouting
- Path Reoptimization
- Path Rerouting: Link Failure
- Path Adjustment with Autobandwidth
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify MPLS-TE path setup procedures
- Identify how RSVP is used in path setup
- Identify the process that RSVP uses to assist in MPLS-TE path setup
- Identify the functions of tunnel and link admission control
- Identify why LSP path rerouting is necessary
- Identify path reoptimization
- Identify how a path is rerouted as a result of link failure
- Identify how traffic tunnels are reoptimized and how the bandwidth requirement is adjusted

Lesson 3-5: Assigning Traffic to Traffic Tunnels

This lesson describes several well-known drawbacks of traditional IP routing: full-mesh requirements in the IP-over-ATM overlay model and lack of TE. The lesson includes these topics:

- Traffic Flow Modifications
- Static Routing Pointing to a Tunnel
- Solving the BGP Next-Hop Problem
- IP Forwarding Database Modification with Autoroute
- Autoroute Path Selection Rules
- Autoroute Examples
- Forwarding Adjacency
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the traffic flow modifications that are available with static routes and policy-based routing
- Identify how tunnel access is implemented using static routing pointing to a tunnel
- Identify how MPLS-TE solves the BGP next-hop problem
- Configure autoroute to modify the IP forwarding database
- Identify the autoroute path selection rules
- Identify what happens in certain autoroute situations
- Identify the function of forwarding adjacency as it relates to MPLS-TE

Module Assessment

Module 4: Configuring MPLS Traffic Engineering

This module discusses the requirement for implementing MPLS-TE in modern networks. It includes configuration of routers to enable basic traffic tunnels, assignment of traffic to a tunnel, control of path selection, and performance of tunnel protection and tunnel maintenance.

Lesson 4-1: Configuring MPLS Traffic Engineering on Cisco IOS Platforms

This lesson describes MPLS-TE commands for the implementation of MPLS traffic tunnels. The configuration commands that are needed to support MPLS-TE are explained, and sample setups presented. The lesson includes these topics:

- MPLS TE Configuration Road Map
- Enabling Device-Level MPLS-TE Support
- Enabling MPLS-TE on an Interface
- Enabling MPLS-TE Support in IS-IS
- Enabling MPLS-TE Support in OSPF
- Creating and Configuring a Traffic Tunnel
- Deploying Traffic Tunnels with Autoroute and Forwarding Adjacency
- Implementing MPLS-TE
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the tasks that are required to implement MPLS-TE
- Enable device-level MPLS-TE support
- Enable MPLS-TE on an interface
- Enable MPLS-TE support in an IS-IS environment
- Enable MPLS-TE support in an OSPF environment
- Create and configure a traffic tunnel
- Enable traffic tunnels with autoroute or forwarding adjacency

Lesson 4-2: Advanced MPLS-TE Path Selection

This lesson describes advanced MPLS-TE commands that are used in path selection. The configuration commands are accompanied by usage guidelines and sample setups. The lesson includes these topics:

- MPLS-TE Case Study
- Creating Dynamic Path Options for a Traffic Tunnel
- Creating Explicit Path Options for a Traffic Tunnel
- Assigning Traffic to Traffic Tunnels with Static Routes
- Configuring Periodic and Event-Triggered Optimization
- Configuring Path Selection Restrictions
- Modifying Administrative Weight
- Autoroute Forwarding Engine
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to use a case study to:

- Describe a typical MPLS-TE implementation
- Identify the command syntax to create dynamic path options for a traffic tunnel
- Identify the command syntax to create explicit path options for a traffic tunnel
- Identify the command syntax to assign traffic to traffic tunnels with static routes
- Identify the command syntax to configure periodic and event-triggered optimization
- Identify the command syntax to configure path selection restrictions
- Identify the command syntax to modify the administrative weight of a traffic tunnel
- Identify the command syntax to enable the autoroute forwarding engine

Lesson 4-3: Advanced MPLS-TE Link Protection

This lesson describes the advanced MPLS-TE commands for link protection. The configuration commands are accompanied by usage guidelines and sample setups. The lesson includes these topics:

- Improving Convergence Time
- Configuring Parallel Tunnels
- Drawbacks of Parallel Tunnels
- Fast Reroute
- Link Protection with FRR
- Node Protection with FRR
- Fast Reroute Configuration
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify issues that are related to improving convergence time in an MPLS-TE network
- Configure a typical parallel tunnel solution
- Identify the drawbacks of using parallel tunnels
- Identify the characteristics of the Fast Reroute feature
- Identify the characteristics of link protection with the Fast Reroute capability
- Identify the characteristics of node protection with the Fast Reroute capability
- Configure fast rerouting

Lesson 4-4: Advanced MPLS-TE Bandwidth Control

This lesson describes advanced MPLS-TE commands for bandwidth control. The configuration commands are accompanied by usage guidelines and sample setups. The lesson includes these topics:

- Bandwidth Reservation with Autobandwidth
- Autobandwidth Configuration Commands
- DiffServ-Aware Service
- DiffServ-Aware Service Configuration Commands
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the characteristics of bandwidth reservation with autobandwidth
- Configure the autobandwidth feature
- Identify the characteristics of the DiffServ-aware service
- Configure the DiffServ-aware service

Lesson 4-5: MPLS-TE Interarea Support

This lesson describes the advanced MPLS-TE topics and commands for interarea support. The configuration commands are accompanied by usage guidelines and sample setups. The lesson includes these topics:

- MPLS-TE Without Interarea Support
- Interarea MPLS-TE
- Interarea MPLS-TE Detailed Operation
- Interarea MPLS-TE Restoration
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Describe the characteristics of MPLS-TE without interarea support
- Describe how MPLS-TE provides interarea support
- Describe interarea MPLS-TE restoration

Lesson 4-6: Monitoring and Troubleshooting MPLS-TE

This lesson describes the monitoring and troubleshooting commands of Cisco IOS software that support MPLS-TE. The troubleshooting guidelines are explained by means of sample command outputs. The lesson includes these topics:

- Monitoring MPLS-TE Tunnels
- Monitoring MPLS-TE
- Monitoring MPLS-TE Resource Usage on Intermediate Devices
- Troubleshooting Device and Link Setup
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the command syntax that is used to monitor MPLS-TE tunnels
- Identify the command syntax that is used to monitor MPLS-TE
- Identify the command syntax that is used to monitor MPLS-TE resource usage on intermediate devices
- Identify the command syntax that is used to troubleshoot device and link setup

Module 5: MPLS Quality of Service

This module covers the concepts and features of MPLS quality of service (QoS) and presents techniques that network administrators can apply to help the service provider meet the service level agreements (SLAs) of customers. The module also covers how QoS is configured in various situations.

Lesson 5-1: QoS Models

This lesson describes the concepts, features, and techniques of QoS that are used within the service provider network. QoS tunneling modes are also discussed, as well as how Differentiated Services Code Point (DSCP) services are handled by MPLS. The lesson includes these topics:

- QoS Models
- The QoS Pendulum
- Integrated QoS Model
- Differentiated QoS Model
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the MPLS QoS models
- Identify the characteristics of the QoS pendulum
- Identify the characteristics of the integrated QoS model
- Identify the characteristics of the differentiated QoS model

Lesson 5-2: MPLS Support for DiffServ

This lesson describes the Differentiated Services (DiffServ) model, including how MPLS supports DiffServ. The lesson includes these topics:

- DiffServ Architecture
- DiffServ Model Features
- DiffServ PHBs and Recommended Codepoints
- DiffServ Scalability by Means of Aggregation
- MPLS Scalability by Means of Aggregation
- Marking MPLS Frames
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the characteristics of the DiffServ architecture
- Identify the features of the DiffServ model
- Identify the DiffServ PHBs and recommended codepoints
- Identify the characteristic of DiffServ scalability by means of aggregation
- Identify the characteristic of MPLS scalability by means of aggregation
- Identify how MPLS marks frames

Lesson 5-3: Configuring MPLS QoS

This lesson describes the configuration tool, Modular QoS CLI (MQC), which is used when applying the QoS parameters on a router. The lesson also discusses how to configure MPLS QoS, including configuration syntax and descriptions. The lesson includes these topics:

- How to Use the QoS Toolkit
- Configuring QoS in Cisco IOS Modular QoS CLI
- Configuring QoS in Cisco IOS MQC Abstractions
- Configuring QoS in Cisco IOS MQC Syntax
- MPLS QoS Configuration Case Study
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the features that are available in the QoS toolkit
- Identify the features of MQC
- Identify IOS QoS MQC abstractions
- Configure QoS in IOS MQC syntax

Lesson 5-4: QoS in MPLS Applications

This lesson focuses on how QoS is integrated in MPLS implementations. The lesson includes these topics:

- MPLS-TE with a Best-Effort Network
- MPLS-TE with a DiffServ Network
- MPLS DS-TE with a DiffServ Network
- QoS-Enabled MPLS VPNs
- QoS Implementation
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify how MPLS-TE is implemented in a best-effort network
- Identify how MPLS-TE is implemented in a DiffServ network
- Identify how MPLS DS-TE is implemented in a DiffServ network
- Identify how QoS is enabled in an MPLS VPN environment
- Identify some QoS implementation tasks

Module: Any Transport over MPLS

MPLS offers a set of advanced features to provide enhanced services. This module looks at one of the most recent supported features: Any Transport over MPLS (AToM). AToM provides transport of Layer 2 frames and cells over an MPLS network.

Lesson 6-1: Introduction to Any Transport over MPLS

There is an ever-increasing demand for the transport of Layer 2 and Layer 3 over a common backbone. This lesson introduces the MPLS application AToM. AToM allows an MPLS network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, High-Level Data Link Control (HDLC), Frame Relay, and ATM. The lesson includes these topics:

- AToM Overview
- Transport Types
- How AToM Works
- AToM Control Word
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Describe the benefits of AToM
- At a high level, describe how AToM transports different Layer 2 frame types and cells across an MPLS network
- Describe how AToM establishes end-to-end connectivity for a Layer 2 frame
- Describe the AToM control word

Lesson 6-2: Configuring AToM on Cisco IOS Platforms

While the configuration of AToM support is not complicated, network administrators must take care to ensure that the proper configuration tasks, parameters, and syntax are used. In addition, AToM provides commands to verify the operation of AToM. The lesson includes these topics:

- MTU Issues
- AToM Packet Forwarding with Summarization in the Core
- AToM Configuration
- EoMPLS Configuration
- PPP over MPLS Configuration
- HDLC over MPLS Configuration
- FRoMPLS Port Configuration
- AAL5 over MPLS Configuration
- ATM over MPLS Configuration
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the effects of AToM on the MTU
- Identify PE loopback summarization problems
- Enable AToM
- Enable a typical EoMPLS configuration
- Enable a typical PPP configuration
- Enable a typical HDLC configuration
- Enable a typical FRoMPLS configuration
- Enable a typical AAL5 over MPLS configuration
- Enable a typical ATM over MPLS configuration

Lesson 6-3: Monitoring AToM on Cisco IOS Platforms

While configuring AToM support is not complicated, network administrators must take care to ensure proper operation. This lesson describes the commands to verify the operation of AToM. The lesson includes these topics:

- Monitoring AToM
- `show mpls 12transport vc`
- `show mpls 12transport vc detail`
- `show mpls 12transport summary`
- `debug mpls 12transport vlan control`
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the commands that are available to monitor AToM operation.
- Identify the function, syntax, and output of the **`show mpls 12transport vc`** command.
- Identify the function, syntax, and output of the **`show mpls 12transport vc summary`** command.
- Identify the function, syntax, and output of the **`show mpls 12transport vc detail`** command.
- Identify the function, syntax, and output of the **`debug mpls 12transport vlan control`** command.

Module 7: MPLS IPv6 support

IP version 6 (IPv6) is becoming more and more accepted. To meet this market, MPLS has implemented support of IPv6 in several modes. This module discusses the different modes for implementing IPv6 in an MPLS environment and the necessary configuration effort.

Lesson 7-1: Review of IPv6

This lesson provides the learner with a brief review of IPv6 operation. The lesson includes these topics:

- The Benefits of Integrating IPv6
- IPv6 Technology Scope
- IPv6 Address Representation
- Hierarchical Addressing and Aggregation
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the benefits of integrating IPv6 into a network
- Compare IPv4 and IPv6 solutions
- Identify the need for a larger IP address space
- Identify the characteristics of an IPv6 address
- Identify how an IPv6 address is represented
- Identify the hierarchical addressing and aggregation scheme that is used by IPv6

Lesson7-2: Implementing IPv6 over MPLS

MPLS is deployed by many service providers in their IP version 4 (IPv4) networks. Service providers want to introduce IPv6 services to their customers, but changes to their existing IPv4 infrastructure can be expensive and the cost benefit for a small amount of IPv6 traffic does not make economic sense. Several integration scenarios have therefore been developed to leverage an existing IPv4 MPLS infrastructure and to add IPv6 services without requiring any changes to the network backbone. This lesson discusses their implementation. The lesson includes these topics:

- Benefits of Deploying IPv6 over MPLS Backbones
- IPv6 Using Tunnels on the Customer Edge Routers
- Deploying IPv6 Using Tunnels on the Customer Edge Routers
- IPv6 over a Circuit Transport over MPLS
- Deploying IPv6 over a Circuit Transport over MPLS
- IPv6 on the Provider Edge Routers (Cisco 6PE)
- Deploying Cisco 6PE
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the benefits of deploying IPv6 over MPLS backbones
- Identify the characteristics of an IPv6 model that uses tunnels on the customer edge routers
- Deploy IPv6 using tunnels on the customer edge routers
- Identify the characteristics of a model that deploys IPv6 over a circuit transport over MPLS
- Deploy IPv6 over a circuit transport over MPLS
- Identify the characteristics of an IPv6 model (Cisco 6PE) that uses the provider edge routers
- Deploy IPv6 on the provider edge routers

Lesson 7-3: Monitoring IPv6 over MPLS

The IPv6 models for using tunnels on CE routers and using IPv6 over a circuit transport over MPLS are implemented using existing IPv6 and AToM technologies. Therefore, these models use existing **show** and **debug** commands.

The Cisco 6PE model provides additional support over existing MPLS technology. Many of the existing MPLS **show** and **debug** commands are available to monitor Cisco 6PE operations (for information related to there operations, refer to the *Implementing Cisco MPLS* (MPLS) course). An additional set of 6PE-specific commands has also been provided. These commands will be discussed in this lesson.

The lesson includes these topics:

- Monitoring IPV6 Support
- show bgp ipv6 Command
- show bgp ipv6 neighbors Command
- show mpls forwarding-table Command
- show bgp ipv6 labels Command
- show ipv6 cef Command
- show ipv6 route Command
- Lesson Summary
- Lesson Review

After completing this lesson, the learner will be able to:

- Identify the commands that are available for monitoring IPv6 support
- Identify the syntax of the **show bgp ipv6** command and the important fields in its output
- Identify the syntax of the **show bgp ipv6 neighbors** command and the important fields in its output
- Identify the syntax of the **show mpls forwarding-table** command and the important fields in its output
- Identify the syntax of the **show bgp ipv6 labels** command and the important fields in its output
- Identify the syntax of the **show ipv6 cef** command and the important fields in its output
- Identify the syntax of the **show ipv6 route** command and the important fields in its output

Lab Equipment Configuration

The lab equipment for this course is divided into two 8-student labs that share a common Frame Relay and terminal server infrastructure. The two labs are logical mirror views of each other. The labs may be run independently to support two separate 8-student classes or in parallel to support a single 16-student class. The equipment configuration information is necessary for initial setup of the lab configuration.

Equipment List

This table lists the recommended equipment to support the lab exercises.

Lab Equipment List

Number of Units	Part Number	Product Description
2	CISCO3725 * Note	Cisco 3700 router (terminal server routers and Frame relay switch)
5	NM-8A/S	8-port sync interface (WAN interface for Frame Relay switch)
2	NM-32AS	32-port aysnc Interface for terminal server
5	CAB-OCTAL-ASYNC=	Terminal server cables

Note Cisco 3600s may be substituted for Cisco 3700s as the common equipment.

Lab 1 Equipment List

Number of Units	Part Number	Product Description
8	CISCO2610XM	Cisco 2600 router (CE routers)
9	CISCO3725	Cisco 3700 router (PE and P routers)
9	MEM3725-32U64CF	16- to 32-MB DRAM Flash upgrade for 3725
17	WIC-1T	WAN interface for CE, PE, and P routers
17	CAB-V35MT	WAN cable for CE, PE, and P routers to Frame Relay switch
17	CAB-V35FC	WAN cables for Frame Relay switch to router

Lab 2 Equipment List


Number of Units	Part Number	Product Description
8	CISCO2610XM	Cisco 2600 router (CE routers)
9	CISCO3725	Cisco 3700 router (PE and P routers)
9	MEM3725-32U64CF	16- to 32-MB DRAM Flash upgrade for 3725
17	WIC-1T	WAN interface for CE, PE, and P routers
17	CAB-V35MT	WAN cable for CE, PE, and P routers to Frame Relay switch
17	CAB-V35FC	WAN cables for Frame Relay switch to routers

Lab Cisco IOS Image

Router	Image
CE	c2600-js-mz.122-10a.bin or later
PE	c3725-js-mz.122-15.t.bin or later
P	c3725-js-mz.122-15.t.bin or later
Frame Relay switch	c3660-js-mz.122-8.T5.bin or c3725-js-mz.122-15.t.bin or later

Lab Environment

This section contains the information about the lab setup, details of the physical and logical connectivity in the lab, and information on the addressing scheme and IGP routing. The class will be divided into workgroups. Each workgroup will have its own CE, PE, and P routers.



The names of all routers in your workgroup follow the naming convention that is presented in the figure.

Equipment Setup (Cabling)

These instructions describe how to cable the lab configuration.

Step 1 Rack-mount all routers.

Step 2 Connect each router to a router that is acting as a terminal server as indicated below.

Router Console Connections

Router ID	Type	Terminal Server Card	Terminal Server Port
Lab 1			
P1	3725	S1	0
WG1 R1	2610	S1	1
WG1 R2	2610	S1	2
WG1 R3	2610	S1	3
WG1 R4	2610	S1	4

Router ID	Type	Terminal Server Card	Terminal Server Port
WG2 R1	2610	S1	5
WG2 R1	2610	S1	6
WG2 R3	2610	S1	7
WG2 R4	2610	S1	8
WG1 R5	3725	S1	17
WG1 R6	3725	S1	18
WG2 R5	3725	S1	19
WG2 R6	3725	S1	20
WG1 R7	3725	S1	25
WG1 R8	3725	S1	26
WG2 R7	3725	S1	27
WG2 R8	3725	S1	28
R9	3275	S2	4

Lab 2

WG1 R1	2610	S1	9
WG1 R2	2610	S1	10
WG1 R3	2610	S1	11
WG1 R4	2610	S1	12
WG2 R1	2610	S1	13
WG2 R2	2610	S1	14
WG2 R3	2610	S1	15
WG2 R4	2610	S1	16
WG1 R5	3725	S1	21
WG1 R6	3725	S1	22
WG2 R5	3725	S1	23
WG2 R6	3725	S1	24
WG1 R7	3725	S1	29
WG1 R8	3725	S1	30
WG2 R7	3725	S1	31
WG2 R8	3725	S2	0
P1	3275	S2	3

Common Equipment

FRSW	3632	S2	1

Step 3 Connect the first serial interface (S0/0) of each router to the router acting as the Frame Relay switch.

Router WAN Interface Connections

Router ID	Type	Frame Relay Card	Frame Relay Port

Lab 1

P1	3275	S1	0
WG1 R1	2610	S2	0
WG1 R2	2610	S2	1
WG1 R3	2610	S2	2
WG1 R4	2610	S2	3
WG2 R1	2610	S2	4
WG2 R1	2610	S2	5
WG2 R3	2610	S2	6
WG2 R4	2610	S2	7
WG1 R5	3725	S4	0
WG1 R6	3725	S4	1
WG2 R5	3725	S4	2
WG2 R6	3725	S4	3
WG1 R7	3725	S5	0
WG1 R8	3725	S5	1
WG2 R7	3725	S5	2
WG2 R8	3725	S5	3

Lab 2

P1	3275	S1	1
WG1 R1	2610	S2	0
WG1 R21	2610	S3	1
WG1 R3	2610	S3	2
WG1 R4	2610	S3	3
WG2 R1	2610	S3	4
WG2 R2	2610	S3	5
WG2 R3	2610	S3	6
WG2 R4	2610	S3	7
WG1 R5	3725	S4	4
WG1 R6	3725	S4	5
WG2 R5	3725	S4	6

Router ID	Type	Frame Relay Card	Frame Relay Port
WG2 R6	3725	S4	7
WG1 R7	3725	S5	4
WG1 R8	3725	S5	5
WG2 R7	3725	S5	6
WG2 R8	3725	S5	7

Equipment Setup (Terminal Server Configuration)

The following baseline configuration (click the link) is located in the Lab Configuration Files:


[Terminal Server Baseline Configuration File](#)

Equipment Setup (Baseline Configuration, Lab Exercise 2)

The following baseline configuration (click the link) is located in the Lab Configuration Files:

[Lab Exercise 2 Baseline Configuration](#)

Frame Relay Switch Configuration



Note	Ensure that all frame-relay route statements have been cleared for the ports that are supporting the lab that you are using: Lab 1 – ports S0/1, s2/0 - 7, 4/0 - 3, and 5/0 – 3; Lab 2 – ports S1/1, s3/0 - 7, 4/4 -7, and 5/4 – 7.

Frame Relay Switch Baseline Configuration File

Click the links here for the following files:

- Lab 1
- Lab 2

Logical connectivity is provided by preconfigured permanent virtual circuits (PVCs).

The data-link connection identifier (DLCI) values refer to the learner lab guide.

Router Baseline Configuration Files

Click the links here for the following files:

Note	These configurations are provided to the learner in the Lab Guide. It is left to the discretion of the instructor as to whether they should be loaded by the learner at the start of the lab or by the instructor before the learner starts the lab.


- R1 Baseline
- R2 Baseline
- R3 Baseline
- R4 Baseline
- R5 Baseline
- R6 Baseline
- R7 Baseline
- R8 Baseline

Equipment Setup (Baseline Configuration, Lab Exercise 4)

The following baseline configuration (click the link) is located in the Lab Configuration Files:

Lab Exercise 4 Baseline Configuration

Frame Relay Switch Configuration



Frame Relay Switch Baseline Configuration File

Note You must clear the frame-relay routes from Exercise 2 by loading the “clear” configuration before installing the routes for Exercise 4 using the “set” configuration.

- Lab 1 (click the links):
 - Clear Exercise 2 Frame Relay Routes
 - Set Exercise 4 Frame Relay Routes
- Lab 2 (click the links):
 - Clear Exercise 2 Frame Relay Routes
 - Set Exercise 4 Frame Relay Routes

Logical connectivity is provided by preconfigured PVCs.

The DLCI values refer to the learner lab guide.

Router Baseline Configuration Files

Click the links here for the following files:

Note	These configurations are provided to the learner in the Lab Guide. It is left to the discretion of the instructor as to whether they should be loaded by the learner at the start of the lab or by the instructor before the learner starts the lab.

- R1 Baseline
- R2 Baseline
- R3 Baseline
- R4 Baseline
- R5 Baseline
- R6 Baseline
- R7 Baseline
- R8 Baseline
- R9 Baseline

Equipment Setup (Baseline Configuration, Lab Exercise 5)

The following baseline configuration (click the link) is located in the Lab Configuration Files:

Lab Exercise 5 Baseline Configuration

Frame Relay Switch Configuration

Note	This lab uses the same Frame Relay configuration as Exercise 4.

Router Baseline Configuration Files

Click the links here for the following files:

Note	These configurations are provided by the instructor for the purpose of recovery. They represent the configurations on the routers at the end of Exercise 4-1 and are therefore the starting configuration for this exercise.

- R1 Baseline
- R2 Baseline
- R3 Baseline
- R4 Baseline
- R5 Baseline
- R6 Baseline
- R7 Baseline
- R8 Baseline
- R9 Baseline

Teardown Procedures

These instructions describe how to clean up after completion of the class:

- Step 1** Erase the startup configurations on all CE, PE, and P routers.
- Step 2** Reload all CE, PE, and P routers.

Lab Configuration Files

Terminal Server Baseline

Note	You will need to set your own password and default gateway.

Lab Exercise 2 Baseline

Frame Relay Switch Baseline

Lab 1 Configuration

```
en
conf t

hostname FrameSwitch

enable password slpm

frame-relay switching

ip cef

interface Serial1/0
  description Interface supporting R9 in lab 1
  bandwidth 128
  no ip address
  encapsulation frame-relay
  no fair-queue
  clockrate 128000
  shut

interface Serial2/0
  description Link to R1 router in lab 1 wg1
  bandwidth 128
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 101 interface Serial2/1 101
  no shut

interface Serial2/1
  description Link to R2 router in lab 1 wg1
  bandwidth 128
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 102 interface Serial4/0 102
  frame-relay route 101 interface Serial2/0 101
  no shut

interface Serial4/0
  description Link to R5 router in lab 1 wg1
  no ip address
  encapsulation frame-relay
```

```

clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial2/1 102
frame-relay route 111 interface Serial5/0 111
no shut

interface Serial5/0
description Link to R7 router in lab 1 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/0 111
frame-relay route 112 interface Serial5/1 112
no shut

interface Serial5/1
description Link to R8 router in lab 1 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/1 111
frame-relay route 112 interface Serial5/0 112
no shut

interface Serial4/1
description Link to R6 router in lab 1 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial2/3 102
frame-relay route 111 interface Serial5/1 111
no shut

interface Serial2/3
description Link to all R4 router in lab 1 wg1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/1 102
frame-relay route 101 interface Serial2/2 101
no shut

interface Serial2/2
description Link to all R3 router in lab 1 wg1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce

```

```

frame-relay route 101 interface Serial2/3 101
no shut

interface Serial2/4
description Link to all R1 router in lab 1 wg2
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/5 101
no shut

interface Serial2/5
description Link to R2 router in lab 1 wg2
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/2 102
frame-relay route 101 interface Serial2/4 101
no shut

interface Serial4/2
description Link to R5 router in lab 1 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial2/5 102
frame-relay route 111 interface Serial5/2 111
no shut

interface Serial5/2
description Link to R7 router in lab 1 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/2 111
frame-relay route 112 interface Serial5/3 112
no shut

interface Serial5/3
description Link to R8 router in lab 1 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/3 111
frame-relay route 112 interface Serial5/2 112
no shut

```

```

interface Serial4/3
description Link to R6 router in lab 1 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial2/7 102
frame-relay route 111 interface Serial5/3 111
no shut

interface Serial2/7
description Link to R4 router in lab 1 wg2
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/3 102
frame-relay route 101 interface Serial2/6 101
no shut

interface Serial2/6
description Link to R3 router in lab 1 wg2
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/7 101
no shut

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
login

exit

copy run start

```

Lab 2 Configuration

```

en
conf t

no service password-encryption

hostname FrameSwitch

enable password slpm

frame-relay switching

ip cef
interface Serial1/1
description Interface supporting R9 in lab 2
bandwidth 128

```



```

no ip address
encapsulation frame-relay
no fair-queue
clockrate 128000
no shut

interface Serial3/0
description Link to R1 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/1 101
no shut

interface Serial3/1
description Link to R2 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/0 101
no shut

interface Serial3/2
description Link to R3 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/3 101
no shut

interface Serial3/3
description Link to R4 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/5 102
frame-relay route 101 interface Serial3/2 101
no shut

interface Serial3/4
description Link to R1 router in lab 2 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/5 101
no shut

interface Serial3/5
description Link to R2 router in lab 2 wg2
no ip address
encapsulation frame-relay

```

```

clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/6 102
frame-relay route 101 interface Serial3/4 101
no shut

interface Serial3/6
description Link to R3 router in lab 2 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/7 101
no shut

interface Serial3/7
description Link to R4 router in lab 2 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/7 102
frame-relay route 101 interface Serial3/6 101
no shut

interface Serial4/4
description Link to R5 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial3/1 102
frame-relay route 111 interface Serial5/4 111
no shut

interface Serial4/5
description Link to R6 router in lab 2 wg1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial3/3 102
frame-relay route 111 interface Serial5/5 111
no shut

interface Serial4/6
description Link to R5 router in lab 2 wg2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial3/5 102
frame-relay route 111 interface Serial5/6 111
no shut

```

```

interface Serial4/7
  description Link to all R6 router in lab 2 wg2
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 102 interface Serial3/7 102
  frame-relay route 111 interface Serial5/7 111
  no shut

interface Serial5/4
  description Link to R7 router in lab 2 wg1
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 111 interface Serial4/4 111
  frame-relay route 112 interface Serial5/5 112
  no shut

interface Serial5/5
  description Link to R8 router in lab 2 wg1
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 111 interface Serial4/5 111
  frame-relay route 112 interface Serial5/4 112
  no shut

interface Serial5/6
  description Link to R6 router in lab 2 wg2
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 111 interface Serial4/6 111
  frame-relay route 112 interface Serial5/7 112
  no shut

interface Serial5/7
  description Link to R8 in lab 2 wg2
  no ip address
  encapsulation frame-relay
  clockrate 128000
  frame-relay lmi-type cisco
  frame-relay intf-type dce
  frame-relay route 111 interface Serial4/7 111
  frame-relay route 112 interface Serial5/6 112
  no shut

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  login

```

```
exit

copy run start
```

Router Baselines

R1 Baseline

```
hostname R1

ip cef

interface Loopback0
 ip address 10.1.11.49 255.255.255.255

interface loopback11
 ip address 172.11.1.16 255.255.255.0

interface loopback12
 ip address 172.12.1.16 255.255.255.0

interface Serial0/0
 bandwidth 128000
 no ip address
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R2
 ip address 10.1.11.129 255.255.255.240
 frame-relay interface-dlci 101

router ospf 1
 log-adjacency-changes
 router-id 10.1.11.49
 network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
 no synchronization
 bgp log-neighbor-changes
 network 172.11.1.0 mask 255.255.255.0
 network 172.12.1.0 mask 255.255.255.0
 neighbor 10.1.12.49 remote-as 65010
 neighbor 10.1.12.49 update-source Loopback0
 no auto-summary

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
```

R2 Baseline

```
hostname R2

ip cef

interface Loopback0
 ip address 10.1.12.49 255.255.255.255

interface Serial0/0
 bandwidth 128000
 no ip address
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R2
 ip address 10.1.11.130 255.255.255.240
 frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
 description Link to R5
 ip address 150.1.11.17 255.255.255.240
 frame-relay interface-dlci 102

router ospf 1
 log-adjacency-changes
 router-id 10.1.12.49
 network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
 no synchronization
 bgp log-neighbor-changes
 neighbor 10.1.11.49 remote-as 65010
 neighbor 10.1.11.49 update-source Loopback0
 neighbor 10.1.11.49 route-reflector-client
 no auto-summary

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
```

R3 Baseline

```
hostname R3

ip cef

interface Loopback0
 ip address 10.2.11.49 255.255.255.255

interface loopback11
 ip address 172.11.2.16 255.255.255.0

interface loopback12
 ip address 172.12.2.16 255.255.255.0

interface Serial0/0
```

```

bandwidth 128000
no ip address
encapsulation frame-relay
frame-relay lmi-type cisco
no shut

interface Serial0/0.101 point-to-point
description Link to R4
ip address 10.2.11.129 255.255.255.240
frame-relay interface-dlci 101

router ospf 1
log-adjacency-changes
router-id 10.2.11.49
network 10.0.0.0 0.255.255.255 area 0

router bgp 65010
no synchronization
bgp log-neighbor-changes
network 172.11.2.0 mask 255.255.255.0
network 172.12.2.0 mask 255.255.255.0
neighbor 10.2.12.49 remote-as 65010
neighbor 10.2.12.49 update-source Loopback0
no auto-summary

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

R4 Baseline

```

hostname R4

ip cef

interface Loopback0
ip address 10.2.12.49 255.255.255.255

interface Serial0/0
bandwidth 128000
no ip address
encapsulation frame-relay
frame-relay lmi-type cisco
no shut

interface Serial0/0.101 point-to-point
description Link to R3
ip address 10.2.11.130 255.255.255.240
frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
description Link to R6
ip address 150.1.12.17 255.255.255.240
frame-relay interface-dlci 102

router ospf 1
log-adjacency-changes
router-id 10.2.12.49
network 10.0.0.0 0.255.255.255 area 0

```

```

router bgp 65010
  no synchronization
  bgp log-neighbor-changes
  neighbor 10.2.11.49 remote-as 65010
  neighbor 10.2.11.49 update-source Loopback0
  neighbor 10.2.11.49 route-reflector-client
no auto-summary
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

```

R5 Baseline

```

hostname R5

ip cef

interface Loopback0
  ip address 192.168.1.17 255.255.255.255
  ip router isis

interface Serial0/0
  bandwidth 128000
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.102 point-to-point
  desc link to R2
  ip address 150.1.11.18 255.255.255.240
  frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
  desc link to R7
  ip address 192.168.1.49 255.255.255.240
  ip router isis
  mpls label protocol ldp
  tag-switching ip
  frame-relay interface-dlci 111

router isis
  net 39.0440.0000.0000.0500
  metric-style wide

router bgp 65001
  no synchronization
  bgp log-neighbor-changes
  neighbor 192.168.1.33 remote-as 65001
  neighbor 192.168.1.33 update-source Loopback0
  no auto-summary

  address-family vpnv4
  neighbor 192.168.1.33 activate
  neighbor 192.168.1.33 next-hop-self

```

```

neighbor 192.168.1.33 send-community both
no auto-summary
exit-address-family
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

R6 Baseline

```

hostname R6

ip cef

interface Loopback0
ip address 192.168.1.33 255.255.255.255
ip router isis

interface Serial0/0
bandwidth 128000
no ip address
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco

interface Serial0/0.102 point-to-point
desc link to R4
ip address 150.1.12.18 255.255.255.240
frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
desc link to R8
ip address 192.168.1.65 255.255.255.240
ip router isis
mpls label protocol ldp
tag-switching ip
frame-relay interface-dlci 111

router isis
net 39.0440.0000.0000.0600
metric-style wide

router bgp 65001
no synchronization
bgp log-neighbor-changes
neighbor 192.168.1.17 remote-as 65001
neighbor 192.168.1.17 update-source Loopback0
no auto-summary

address-family vpnv4
neighbor 192.168.1.17 activate
neighbor 192.168.1.17 next-hop-self
neighbor 192.168.1.17 send-community both
no auto-summary
exit-address-family
exit

line con 0

```



```
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
login
```

R7 Baseline

```
hostname R7

ip cef

interface Loopback0
 ip address 192.168.1.81 255.255.255.255

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.111 point-to-point
 ip address 192.168.1.50 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
 ip address 192.168.1.113 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 112
exit

router isis
 net 39.0440.0000.0000.0700
 metric-style wide

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
```

R8 Baseline

```
hostname R8

ip cef

interface Loopback0
 ip address 192.168.1.97 255.255.255.255
 ip router isis

interface Serial0/0
 bandwidth 128000
 no ip address
 encapsulation frame-relay
```

```

no fair-queue
frame-relay lmi-type cisco
!
interface Serial0/0.111 point-to-point
ip address 192.168.1.66 255.255.255.240
ip router isis
mpls label protocol ldp
tag-switching ip
frame-relay interface-dlci 111
!
interface Serial0/0.112 point-to-point
ip address 192.168.1.114 255.255.255.240
ip router isis
mpls label protocol ldp
tag-switching ip
frame-relay interface-dlci 112
router isis
net 39.0440.0000.0000.0800
metric-style wide
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

Lab Exercise 4 Baseline Configurations

Frame Relay Switch Baseline Configuration File

Lab 1

Clear Exercise 2 Frame Relay Routes.

```

en
conf t

interface Serial2/0
description Link to R1 router in lab 1 wg1
no frame-relay route 101 interface Serial2/1 101
shut

interface Serial2/1
description Link to R2 router in lab 1 wg1
no frame-relay route 102 interface Serial4/0 102
no frame-relay route 101 interface Serial2/0 101
shut

interface Serial4/0
description Link to R5 router in lab 1 wg1
no frame-relay route 102 interface Serial2/1 102
no frame-relay route 111 interface Serial5/0 111
shut

interface Serial5/0
description Link to R7 router in lab 1 wg1
no frame-relay route 111 interface Serial4/0 111
no frame-relay route 112 interface Serial5/1 112
shut

```

```

interface Serial5/1
  description Link to R8 router in lab 1 wg1
  no frame-relay route 111 interface Serial4/1 111
  no frame-relay route 112 interface Serial5/0 112
  shut

interface Serial4/1
  description Link to R6 router in lab 1 wg1
  no frame-relay route 102 interface Serial2/3 102
  no frame-relay route 111 interface Serial5/1 111
  shut

interface Serial2/3
  description Link to all R4 router in lab 1 wg1
  no frame-relay route 102 interface Serial4/1 102
  no frame-relay route 101 interface Serial2/2 101
  shut

interface Serial2/2
  description Link to all R3 router in lab 1 wg1
  no frame-relay route 101 interface Serial2/3 101
  shut

interface Serial2/4
  description Link to all R1 router in lab 1 wg2
  no frame-relay route 101 interface Serial2/5 101
  shut

interface Serial2/5
  no frame-relay route 102 interface Serial4/2 102
  no frame-relay route 101 interface Serial2/4 101
  shut

interface Serial4/2
  no frame-relay route 102 interface Serial2/5 102
  no frame-relay route 111 interface Serial5/2 111
  shut

interface Serial5/2
  description Link to R7 router in lab 1 wg2
  no frame-relay route 111 interface Serial4/2 111
  no frame-relay route 112 interface Serial5/3 112
  shut

interface Serial5/3
  description Link to R8 router in lab 1 wg2
  no frame-relay route 111 interface Serial4/3 111
  no frame-relay route 112 interface Serial5/2 112
  shut

interface Serial4/3
  description Link to R6 router in lab 1 wg2
  no frame-relay route 102 interface Serial2/7 102
  no frame-relay route 111 interface Serial5/3 111
  shut

interface Serial2/7
  description Link to R4 router in lab 1 wg2

```

```

no frame-relay route 102 interface Serial4/3 102
no frame-relay route 101 interface Serial2/6 101
shut

interface Serial2/6
description Link to R3 router in lab 1 wg2
bandwidth 128
no frame-relay route 101 interface Serial2/7 101
shut

end

copy run start

```

Set Exercise 4 Frame Relay Routes

```

en
conf t

no service password-encryption

hostname FrameSwitch

enable password slpm

frame-relay switching

ip cef

interface Serial1/0
description Interface supporting R9 in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
no fair-queue
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 211 interface Serial5/0 211
frame-relay route 212 interface Serial5/1 212
frame-relay route 221 interface Serial5/2 221
frame-relay route 222 interface Serial5/3 222
no shut

interface Serial2/0
description Link to wg 1 R1 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/0 101
no shut

interface Serial2/1
description Link to wg 1 R2 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco

```

```

frame-relay intf-type dce
frame-relay route 102 interface Serial4/0 102
no shut

interface Serial2/2
description Link to all wg 1 R3 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/1 101
no shut

interface Serial2/3
description Link to all wg 1 R4 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/1 102
no shut

interface Serial2/4
description Link to all wg 2 R1 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/2 101
no shut

interface Serial2/5
description Link to wg 2 R2 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/2 102
no shut

interface Serial2/6
description Link to wg 2 R3 router in lab 1
bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/3 101
no shut

interface Serial2/7
description Link to wg 2 R4 router in lab 1

```

```

bandwidth 128
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/3 102
no shut

interface Serial4/0
description Link to wg 1 R5 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/0 101
frame-relay route 102 interface Serial2/1 102
frame-relay route 111 interface Serial5/0 111
no shut

interface Serial4/1
description Link to wg 1 R6 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/2 101
frame-relay route 102 interface Serial2/3 102
frame-relay route 111 interface Serial5/1 111
no shut

interface Serial4/2
description Link to wg 2 R5 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/4 101
frame-relay route 102 interface Serial2/5 102
frame-relay route 111 interface Serial5/2 111
no shut

interface Serial4/3
description Link to wg 2 R6 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial2/6 101
frame-relay route 102 interface Serial2/7 102
frame-relay route 111 interface Serial5/3 111
no shut

interface Serial5/0

```

```

description Link to wg 1 R7 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/0 111
frame-relay route 112 interface Serial5/1 112
frame-relay route 211 interface Serial1/0 211
no shut

interface Serial5/1
description Link to wg 1 R8 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/1 111
frame-relay route 112 interface Serial5/0 112
frame-relay route 212 interface Serial1/0 212
no shut

interface Serial5/2
description Link to wg 2 R7 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/2 111
frame-relay route 112 interface Serial5/3 112
frame-relay route 221 interface Serial1/0 221
no shut

interface Serial5/3
description Link to wg 2 R8 router in lab 1
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/3 111
frame-relay route 112 interface Serial5/2 112
frame-relay route 222 interface Serial1/0 222
no shut

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
login

exit

copy run start

```

Lab 2

Clear Exercise 2 Frame Relay Routes

```
en
conf t

interface Serial3/0
description Link to R1 router in lab 2 wg1
no frame-relay route 101 interface Serial3/1 101
shut

interface Serial3/1
description Link to R2 router in lab 2 wg1
no frame-relay route 101 interface Serial3/0 101
shut

interface Serial3/2
description Link to R3 router in lab 2 wg1
no frame-relay route 101 interface Serial3/3 101
shut

interface Serial3/3
description Link to R4 router in lab 2 wg1
no frame-relay route 102 interface Serial4/5 102
no frame-relay route 101 interface Serial3/2 101
shut

interface Serial3/4
description Link to R1 router in lab 2 wg2
no frame-relay route 101 interface Serial3/5 101
shut

interface Serial3/5
description Link to R2 router in lab 2 wg2
no frame-relay route 102 interface Serial4/6 102
no frame-relay route 101 interface Serial3/4 101
shut

interface Serial3/6
description Link to R3 router in lab 2 wg2
no frame-relay route 101 interface Serial3/7 101
shut

interface Serial3/7
description Link to R4 router in lab 2 wg2
no frame-relay route 102 interface Serial4/7 102
no frame-relay route 101 interface Serial3/6 101
shut

interface Serial4/4
description Link to R5 router in lab 2 wg1
no frame-relay route 102 interface Serial3/1 102
no frame-relay route 111 interface Serial5/4 111
shut

interface Serial4/5
description Link to R6 router in lab 2 wg1
no frame-relay route 102 interface Serial3/3 102
no frame-relay route 111 interface Serial5/5 111
```



```

shut

interface Serial4/6
description Link to R5 router in lab 2 wg2
no frame-relay route 102 interface Serial3/5 102
no frame-relay route 111 interface Serial5/6 111
shut

interface Serial4/7
description Link to all R6 router in lab 2 wg2
no frame-relay route 102 interface Serial3/7 102
no frame-relay route 111 interface Serial5/7 111
shut

interface Serial5/4
description Link to R7 router in lab 2 wg1
no frame-relay route 111 interface Serial4/4 111
no frame-relay route 112 interface Serial5/5 112
shut

interface Serial5/5
description Link to R8 router in lab 2 wg1
no frame-relay route 111 interface Serial4/5 111
no frame-relay route 112 interface Serial5/4 112
shut

interface Serial5/6
description Link to R6 router in lab 2 wg2
no frame-relay route 111 interface Serial4/6 111
no frame-relay route 112 interface Serial5/7 112
shut

interface Serial5/7
description Link to R8 in lab 2 wg2
no frame-relay route 111 interface Serial4/7 111
no frame-relay route 112 interface Serial5/6 112
shut

end
copy run start

```

Set Exercise 4 Frame Relay Routes

```

en
conf t

no service password-encryption

hostname FrameSwitch

enable password slpm

frame-relay switching

ip cef
interface Serial1/1
description Interface supporting R9 in lab 2
bandwidth 128
no ip address
encapsulation frame-relay
no fair-queue
clockrate 128000

```

```

frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 211 interface Serial5/4 211
frame-relay route 212 interface Serial5/5 212
frame-relay route 221 interface Serial5/6 221
frame-relay route 222 interface Serial5/7 222
no shut

interface Serial3/0
description Link to wg 1 R1 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/4 101
no shut

interface Serial3/1
description Link to wg 1 R2 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/4 102
no shut

interface Serial3/2
description Link to wg 1 R3 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/5 101
no shut

interface Serial3/3
description Link to wg 1 R4 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/5 102
no shut

interface Serial3/4
description Link to wg 2 R1 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/6 101
no shut

interface Serial3/5
description Link to wg 2 R2 router in lab 2
no ip address

```

```

encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/6 102
no shut

interface Serial3/6
description Link to wg R3 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial4/7 101
no shut

interface Serial3/7
description Link to wg 2 R4 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 102 interface Serial4/7 102
no shut

interface Serial4/4
description Link to wg 1 R5 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/0 101
frame-relay route 102 interface Serial3/1 102
frame-relay route 111 interface Serial5/4 111
no shut

interface Serial4/5
description Link to wg 1 R6 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/2 101
frame-relay route 102 interface Serial3/3 102
frame-relay route 111 interface Serial5/5 111
no shut

interface Serial4/6
description Link to wg 2 R5 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/4 101
frame-relay route 102 interface Serial3/5 102
frame-relay route 111 interface Serial5/6 111

```

```

no shut

interface Serial4/7
description Link to all wg 2 R6 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 101 interface Serial3/6 101
frame-relay route 102 interface Serial3/7 102
no shut

interface Serial5/4
description Link to wg 1 R7 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/4 111
frame-relay route 112 interface Serial5/5 112
frame-relay route 211 interface Serial1/1 211
no shut

interface Serial5/5
description Link to wg 1 R8 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/5 111
frame-relay route 112 interface Serial5/4 112
frame-relay route 212 interface Serial1/1 212
no shut

interface Serial5/6
description Link to wg 2 R7 router in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/6 111
frame-relay route 112 interface Serial5/7 112
frame-relay route 221 interface Serial1/1 221
no shut

interface Serial5/7
description Link to wg 2 R8 in lab 2
no ip address
encapsulation frame-relay
clockrate 128000
frame-relay lmi-type cisco
frame-relay intf-type dce
frame-relay route 111 interface Serial4/7 111
frame-relay route 112 interface Serial5/6 112
frame-relay route 222 interface Serial1/1 222
no shut

```

```

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  login

exit

copy run start

```

Router Baseline Configuration Files

Note	The variable <i>wg</i> must be replaced with the appropriate workgroup number.

R1 Baseline

```

en

conf t

hostname R1

interface Loopback0
  ip address 10.1.1.49 255.255.255.255

interface Loopback10
  ip address 10.1.1.17 255.255.255.240

interface Loopback20
  ip address 10.1.1.65 255.255.255.240

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  description Link to R5
  ip address 150.wg.1.17 255.255.255.240
  frame-relay interface-dlci 101
  no shut

router ospf 1
  network 150.wg.0.0 0.0.255.255 area 0
  network 10.1.1.49 0.0.0.0 area 0
  network 10.1.1.16 0.0.0.15 area 1
  network 10.1.wg1.64 0.0.0.15 area 1
  redistribute bgp 65011 subnets
exit

router bgp 65011
  nei 150.wg.1.18 remote-as 65001
  no auto-summary
  redistribute ospf 1
exit

```

```

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

exit
exit
copy run start

```

R2 Baseline

```

ena
conf t

hostname R2

interface Loopback0
  ip address 10.2.1.49 255.255.255.255

interface Loopback10
  ip address 10.2.1.17 255.255.255.240

interface Loopback20
  ip address 10.2.1.65 255.255.255.240

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.102 point-to-point
  desc link to R5
  ip address 150.wg.1.33 255.255.255.240
  frame-relay interface-dlci 102

router ospf 2
  network 150.wg.0.0 0.0.255.255 area 0
  network 10.2.1.49 0.0.0.0 area 0
  network 10.2.1.16 0.0.0.15 area 1
  redistribute bgp 65012 subnets
exit

router bgp 65012
  nei 150.wg.1.34 remote-as 65001
  no auto-summary
  redistribute ospf 2

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login
exit
exit
copy run start

```

R3 Baseline

```
en

conf t

hostname R3

interface Loopback0
 ip address 10.1.2.49 255.255.255.255

interface Loopback10
 ip address 10.1.2.17 255.255.255.240

interface Loopback20
 ip address 10.1.2.65 255.255.255.240

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 no desc link to R8
 ip address 150.wg.2.17 255.255.255.240
 frame-relay interface-dlci 101
 no shut

router ospf 1
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.1.2.49 0.0.0.0 area 0
 network 10.1.2.16 0.0.0.15 area 1
 redistribute bgp 65011 subnets
 exit

router bgp 65011
 nei 150.wg.2.18 remote-as 65001
 no auto-summary
 redistribute ospf 1
 exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
!
exit
exit
copy run start
```

R4 Baseline

```
ena
conf t

hostname R4

interface Loopback0
 ip address 10.2.2.49 255.255.255.255

interface Loopback10
 ip address 10.2.2.17 255.255.255.240

interface Loopback20
 ip address 10.2.2.65 255.255.255.240

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.102 point-to-point
 desc link to R5
 ip address 150.wg.2.33 255.255.255.240
 frame-relay interface-dlci 102

router ospf 2
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.2.2.49 0.0.0.0 area 0
 network 10.2.2.16 0.0.0.15 area 1
 redistribute bgp 65012 subnets
 exit

router bgp 65012
 nei 150.wg.2.34 remote-as 65001
 no auto-summary
 redistribute ospf 2
 exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
exit
exit
copy run start
```

R5 Baseline

```
en

conf t

hostname R5

ip cef
```



```

ip vrf Green
  rd wg:10
  route-target both wg:10
exit

ip vrf Orange
  rd wg:20
  route-target both wg:20
exit

interface Loopback0
  ip address 192.168.wg.17 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  ip vrf forwarding Green
  ip address 150.wg.1.18 255.255.255.240
  frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
  ip vrf forwarding Orange
  ip address 150.wg.1.34 255.255.255.240
  frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
  ip address 192.168.wg.49 255.255.255.240
  ip router isis
  mpls label protocol ldp
  mpls ip
  frame-relay interface-dlci 111

router isis
  net 39.0440.0000.0000.wg500
  metric-style wide
exit

router bgp 65001
  no synchronization
  neighbor 192.168.wg.33 remote-as 65001
  neighbor 192.168.wg.33 update-source Loopback0
  no auto-summary
  address-family vpnv4
    neighbor 192.168.wg.33 activate
    neighbor 192.168.wg.33 next-hop-self
    neighbor 192.168.wg.33 send-community both
    no auto-summary
  exit

  address-family ipv4 vrf Orange
    nei 150.wg.1.33 remote-as 65012
    nei 150.wg.1.33 activate
    nei 150.wg.1.33 as-override
    no auto-summary

```

```

        no synchronization
    exit

    address-family ipv4 vrf Green
        nei 150.wg.1.17 remote-as 65011
        nei 150.wg.1.17 activate
        nei 150.wg.1.17 as-override
        no auto-summary
    exit

exit

line con 0
    exec-timeout 999 0
    logging synchronous
line aux 0
line vty 0 4
    no login

exit
exit

copy run start

```

R6 Baseline

```

ena
conf t

hostname R6

ip cef

ip vrf Green
    rd wg:10
    route-target both wg:10
exit

ip vrf Orange
    rd wg:20
    route-target both wg:20
exit

interface Loopback0
    ip address 192.168.wg.33 255.255.255.255
    ip router isis
!

interface Serial0/0
    no ip address
    bandwidth 128000
    encapsulation frame-relay
    no fair-queue
    frame-relay lmi-type cisco
    no shut

interface Serial0/0.101 point-to-point
    ip vrf forwarding Green
    ip address 150.wg.2.18 255.255.255.240
    frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point

```

```

ip vrf forwarding Orange
ip address 150.wg.2.34 255.255.255.240
frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
ip address 192.168.wg.65 255.255.255.240
ip router isis
mpls label protocol ldp
mpls ip
frame-relay interface-dlci 111

interface FastEthernet0/1
no ip address
shutdown
duplex auto
speed auto
no cdp enable

router isis
net 39.0440.0000.0000.wg600
metric-style wide
exit

router bgp 65001
no synchronization
neighbor 192.168.wg.17 remote-as 65001
neighbor 192.168.wg.17 update-source Loopback0
no auto-summary

address-family vpnv4
neighbor 192.168.wg.17 activate
neighbor 192.168.wg.17 next-hop-self
neighbor 192.168.wg.17 send-community both
no auto-summary
exit-address-family

address-family ipv4 vrf Orange
neighbor 150.wg.2.33 remote-as 65012
neighbor 150.wg.2.33 activate
neighbor 150.wg.2.33 as-override
no auto-summary
no synchronization
exit-address-family

address-family ipv4 vrf Green
neighbor 150.wg.2.17 remote-as 65011
neighbor 150.wg.2.17 activate
neighbor 150.wg.2.17 as-override
no auto-summary
no synchronization
exit-address-family
exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

```

```
exit
exit
copy run start
```

R7 Baseline

```
en

conf t

hostname R7


ip cef

interface Loopback0
 ip address 192.168.wg.81 255.255.255.255
 ip router isis

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.111 point-to-point
 ip address 192.168.wg.50 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
 ip router isis
 ip address 192.168.wg.113 255.255.255.240
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 112

int s0/0.2wg1 point-to-point
 ip address 192.168.wg.129 255.255.255.248
 ip router isis
 mpls ip
 mpls label protocol ldp
 frame-relay interface-dlci 2wg1
 no shut

router isis
 net 39.0440.0000.0000.wg700
 metric-style wide

exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
```

```
no login

exit
exit

copy run start
```

R8 Baseline

```
ena
conf t

hostname R8


ip cef

interface Loopback0
 ip address 192.168.wg.97 255.255.255.255
 ip router isis

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.111 point-to-point
 ip address 192.168.wg.66 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
 ip address 192.168.wg.114 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 frame-relay interface-dlci 112

int s0/0.2wg2 point-to-point
 ip address 192.168.wg.145 255.255.255.248
 ip router isis
 mpls ip
 mpls label protocol ldp
 frame-relay interface-dlci 2wg2
 no shut

router isis
 net 39.0440.0000.0000.wg800
 metric-style wide

exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
```

```
no login
exit
exit
copy run start
```

R9 Baseline

```
en

conf t

hostname R9

enable password slpm

ip cef
mpls traffic-eng tunnels
mpls traffic-eng logging lsp setups
mpls traffic-eng reoptimize timers frequency 30

class-map match-all BUSINESS-MPLS
  match mpls experimental topmost 1 2
class-map match-all PREMIUM-MPLS
  match mpls experimental topmost 5

policy-map OUT-POLICY
  class PREMIUM-MPLS
    priority percent 5
  class BUSINESS-MPLS
    bandwidth percent 50
  random-detect
    random-detect precedence 1 100 300 100
    random-detect precedence 2 300 500 10
  class class-default
    fair-queue
    random-detect

map-class frame-relay FRAME-OUT
  frame-relay cir 128000
  frame-relay mincir 128000
  service-policy output OUT-POLICY

interface Loopback0
  ip address 192.168.10.17 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  ip rsvp bandwidth 128 128
  no shut

interface Serial0/0.211 point-to-point
  desc link to R7 in wg 1
  ip address 192.168.1.130 255.255.255.248
  ip router isis
  mpls label protocol ldp
  tag-switching ip
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128
```

```

    frame-relay class FRAME-OUT
    frame-relay interface-dlci 211

interface Serial0/0.212 point-to-point
    desc link to R8 in wg 1
    ip address 192.168.1.146 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    ip rsvp bandwidth 128 128
    frame-relay class FRAME-OUT
    frame-relay interface-dlci 212

interface Serial0/0.221 point-to-point
    desc link to R7 in wg 2
    ip address 192.168.2.130 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    frame-relay class FRAME-OUT
    ip rsvp bandwidth 128 128
    frame-relay interface-dlci 221

interface Serial0/0.222 point-to-point
    desc link to R8 in wg 2
    ip address 192.168.2.146 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    ip rsvp bandwidth 128 128
    frame-relay class FRAME-OUT
    frame-relay interface-dlci 222

router isis
    net 39.0440.0000.0000.0009.00
    metric-style wide
    mpls traffic-eng router-id Loopback0
    mpls traffic-eng level-1
exit

line con 0
    exec-timeout 999 0
    logging synchronous
line aux 0
line vty 0 4
    no login

exit
exit

copy run start

```

Lab Exercise 5 Baseline Configurations

Router Baseline Configuration Files

R1 Baseline

```
en

conf t

hostname R1

interface Loopback0
 ip address 10.1.1.49 255.255.255.255

interface Loopback10
 ip address 10.1.1.17 255.255.255.240

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.101 point-to-point
 description Link to R5
 ip address 150.wg.1.17 255.255.255.240
 frame-relay interface-dlci 101
 no shut

router ospf 1
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.1.1.49 0.0.0.0 area 0
 network 10.1.1.16 0.0.0.15 area 1
 network 10.1.wg1.64 0.0.0.15 area 1
 redistribute bgp 65011 subnets
 exit

router bgp 65011
 nei 150.wg.1.18 remote-as 65001
 no auto-summary
 redistribute ospf 1
 exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login

exit
exit
copy run start
```


R2 Baseline

```
ena
conf t

hostname R2

interface Loopback0
 ip address 10.2.1.49 255.255.255.255

interface Loopback10
 ip address 10.2.1.17 255.255.255.240

interface Serial0/0
 no ip address
 bandwidth 128000
 encapsulation frame-relay
 no fair-queue
 frame-relay lmi-type cisco
 no shut

interface Serial0/0.102 point-to-point
 desc link to R5
 ip address 150.wg.1.33 255.255.255.240
 frame-relay interface-dlci 102

router ospf 2
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.2.1.49 0.0.0.0 area 0
 network 10.2.1.16 0.0.0.15 area 1
 redistribute bgp 65012 subnets
 exit

router bgp 65012
 nei 150.wg.1.34 remote-as 65001
 no auto-summary
 redistribute ospf 2

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
exit
exit
copy run start
```

R3 Baseline

```
en

conf t

hostname R3

interface Loopback0
 ip address 10.1.2.49 255.255.255.255

interface Loopback10
 ip address 10.1.2.17 255.255.255.240
```

```

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.101 point-to-point
  no desc link to R8
  ip address 150.wg.2.17 255.255.255.240
  frame-relay interface-dlci 101
  no shut

router ospf 1
  network 150.wg.0.0 0.0.255.255 area 0
  network 10.1.2.49 0.0.0.0 area 0
  network 10.1.2.16 0.0.0.15 area 1
  redistribute bgp 65011 subnets
  exit

router bgp 65011
  nei 150.wg.2.18 remote-as 65001
  no auto-summary
  redistribute ospf 1
  exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login
!
exit
exit
copy run start

```

R4 Baseline

```

ena
conf t

hostname R4

interface Loopback0
  ip address 10.2.2.49 255.255.255.255

interface Loopback10
  ip address 10.2.2.17 255.255.255.240

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.102 point-to-point
  desc link to R5

```

```

ip address 150.wg.2.33 255.255.255.240
frame-relay interface-dlci 102

router ospf 2
 network 150.wg.0.0 0.0.255.255 area 0
 network 10.2.2.49 0.0.0.0 area 0
 network 10.2.2.16 0.0.0.15 area 1
 redistribute bgp 65012 subnets
exit

router bgp 65012
 nei 150.wg.2.34 remote-as 65001
 no auto-summary
 redistribute ospf 2
exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
exit
exit
copy run start

```

R5 Baseline

```

en

conf t

hostname R5

ip cef
mpls traffic-eng tunnels

ip vrf Green
 rd wg:10
 route-target both wg:10
exit

ip vrf Orange
 rd wg:20
 route-target both wg:20
exit

interface Loopback0
 ip address 192.168.wg.17 255.255.255.255
 ip router isis

interface Tunnel1
 ip unnumbered Loopback0
 tag-switching ip
 tunnel destination 192.168.1.33
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 0 0
 tunnel mpls traffic-eng bandwidth 100
 tunnel mpls traffic-eng path-option 1 explicit name via_R9

interface Serial0/0
 no ip address

```

```

bandwidth 128000
encapsulation frame-relay
no fair-queue
ip rsvp bandwidth 128 128
frame-relay lmi-type cisco
no shut

interface Serial0/0.101 point-to-point
ip vrf forwarding Green
ip address 150.wg.1.18 255.255.255.240
frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
ip vrf forwarding Orange
ip address 150.wg.1.34 255.255.255.240
frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
ip address 192.168.wg.49 255.255.255.240
ip router isis
mpls label protocol ldp
mpls ip
ip rsvp bandwidth 128 128
mpls traffic-eng tunnels
frame-relay interface-dlci 111

router isis
net 39.0440.0000.0000.wg500
metric-style wide
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-2
exit

router bgp 65001
no synchronization
neighbor 192.168.wg.33 remote-as 65001
neighbor 192.168.wg.33 update-source Loopback0
no auto-summary
address-family vpnv4
neighbor 192.168.wg.33 activate
neighbor 192.168.wg.33 next-hop-self
neighbor 192.168.wg.33 send-community both
no auto-summary
exit

address-family ipv4 vrf Orange
nei 150.wg.1.33 remote-as 65012
nei 150.wg.1.33 activate
nei 150.wg.1.33 as-override
no auto-summary
no synchronization
exit

address-family ipv4 vrf Green
nei 150.wg.1.17 remote-as 65011
nei 150.wg.1.17 activate
nei 150.wg.1.17 as-override
no auto-summary
exit

exit

```

```

ip route 192.168.1.33 255.255.255.255 Tunnel1

ip explicit-path name via_R9 enable
  next-address 192.168.1.81
  next-address 192.168.10.17
  next-address 192.168.1.97
  next-address 192.168.1.33
exit
line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

exit
exit

copy run start

```

R6 Baseline

```

ena
conf t

hostname R6

ip cef
mpls traffic-eng tunnels

ip vrf Green
  rd wg:10
  route-target both wg:10
exit

ip vrf Orange
  rd wg:20
  route-target both wg:20
exit

interface Loopback0
  ip address 192.168.wg.33 255.255.255.255
  ip router isis
!

interface Tunnel1
  ip unnumbered Loopback0
  tag-switching ip
  tunnel destination 192.168.1.17
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng priority 0 0
  tunnel mpls traffic-eng bandwidth 100
  tunnel mpls traffic-eng path-option 1 explicit name via_R9

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  ip rsvp bandwidth 128 128
  no shut

```

```

interface Serial0/0.101 point-to-point
 ip vrf forwarding Green
 ip address 150.wg.2.18 255.255.255.240
 frame-relay interface-dlci 101

interface Serial0/0.102 point-to-point
 ip vrf forwarding Orange
 ip address 150.wg.2.34 255.255.255.240
 frame-relay interface-dlci 102

interface Serial0/0.111 point-to-point
 ip address 192.168.wg.65 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128
 frame-relay interface-dlci 111

router isis
 net 39.0440.0000.0000.wg600
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2

exit

router bgp 65001
 no synchronization
 neighbor 192.168.wg.17 remote-as 65001
 neighbor 192.168.wg.17 update-source Loopback0
 no auto-summary

address-family vpnv4
 neighbor 192.168.wg.17 activate
 neighbor 192.168.wg.17 next-hop-self
 neighbor 192.168.wg.17 send-community both
 no auto-summary
 exit-address-family

address-family ipv4 vrf Orange
 neighbor 150.wg.2.33 remote-as 65012
 neighbor 150.wg.2.33 activate
 neighbor 150.wg.2.33 as-override
 no auto-summary
 no synchronization
 exit-address-family

address-family ipv4 vrf Green
 neighbor 150.wg.2.17 remote-as 65011
 neighbor 150.wg.2.17 activate
 neighbor 150.wg.2.17 as-override
 no auto-summary
 no synchronization
 exit-address-family
exit

ip route 192.168.1.17 255.255.255.255 Tunnel1

```

```

ip explicit-path name via_R9 enable
  next-address 192.168.1.97
  next-address 192.168.10.17
  next-address 192.168.1.81
  next-address 192.168.1.17
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login
exit
exit
copy run start

```

R7 Baseline

```

en

conf t

hostname R7

ip cef
mpls traffic-eng tunnels

interface Loopback0
  ip address 192.168.wg.81 255.255.255.255
  ip router isis

interface Serial0/0
  no ip address
  bandwidth 128000
  encapsulation frame-relay
  no fair-queue
  frame-relay lmi-type cisco
  no shut

interface Serial0/0.111 point-to-point
  ip address 192.168.wg.50 255.255.255.240
  ip router isis
  mpls label protocol ldp
  mpls ip
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128
  frame-relay interface-dlci 111

interface Serial0/0.112 point-to-point
  ip router isis
  ip address 192.168.wg.113 255.255.255.240
  mpls label protocol ldp
  mpls ip
  mpls traffic-eng tunnels
  ip rsvp bandwidth 128 128
  frame-relay interface-dlci 112

```

```

int s0/0.2wg1 point-to-point
ip address 192.168.wg.129 255.255.255.248
ip router isis
mpls ip
mpls label protocol ldp
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 2wg1
no shut

router isis
net 39.0440.0000.0000.wg700
metric-style wide
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-2

exit

line con 0
exec-timeout 999 0
logging synchronous
line aux 0
line vty 0 4
no login

exit
exit

copy run start

```

R8 Baseline

```

ena
conf t

hostname R8

ip cef
mpls traffic-eng tunnels

interface Loopback0
ip address 192.168.wg.97 255.255.255.255
ip router isis

interface Serial0/0
no ip address
bandwidth 128000
encapsulation frame-relay
no fair-queue
frame-relay lmi-type cisco
no shut

interface Serial0/0.111 point-to-point
ip address 192.168.wg.66 255.255.255.240
ip router isis
mpls label protocol ldp
mpls ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
frame-relay interface-dlci 111

```



```

interface Serial0/0.112 point-to-point
 ip address 192.168.wg.114 255.255.255.240
 ip router isis
 mpls label protocol ldp
 mpls ip
 mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128
 frame-relay interface-dlci 112

int s0/0.2wg2 point-to-point
 ip address 192.168.wg.145 255.255.255.248
 ip router isis
 mpls ip
 mpls label protocol ldp
 mpls traffic-eng tunnels
 ip rsvp bandwidth 128 128
 frame-relay interface-dlci 2wg2
 no shut

router isis
 net 39.0440.0000.0000.wg800
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2

exit

line con 0
 exec-timeout 999 0
 logging synchronous
line aux 0
line vty 0 4
 no login
exit
exit
copy run start

```

R9 Baseline

```

en

conf t

hostname R9

enable password slpm

ip cef
 mpls traffic-eng tunnels
 mpls traffic-eng logging lsp setups
 mpls traffic-eng reoptimize timers frequency 30

class-map match-all BUSINESS-MPLS
 match mpls experimental topmost 1 2
class-map match-all PREMIUM-MPLS
 match mpls experimental topmost 5

policy-map OUT-POLICY
 class PREMIUM-MPLS
  priority percent 5
 class BUSINESS-MPLS

```

```

    bandwidth percent 50
    random-detect
    random-detect precedence 1    100    300    100
    random-detect precedence 2    300    500    10
class class-default
    fair-queue
    random-detect

map-class frame-relay FRAME-OUT
    frame-relay cir 128000
    frame-relay mincir 128000
    service-policy output OUT-POLICY

interface Loopback0
    ip address 192.168.10.17 255.255.255.255
    ip router isis

interface Serial0/0
    no ip address
    encapsulation frame-relay
    no fair-queue
    frame-relay lmi-type cisco
    ip rsvp bandwidth 128 128
    no shut

interface Serial0/0.211 point-to-point
    desc link to R7 in wg 1
    ip address 192.168.1.130 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    ip rsvp bandwidth 128 128
    frame-relay class FRAME-OUT
    frame-relay interface-dlci 211

interface Serial0/0.212 point-to-point
    desc link to R8 in wg 1
    ip address 192.168.1.146 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    ip rsvp bandwidth 128 128
    frame-relay class FRAME-OUT
    frame-relay interface-dlci 212

interface Serial0/0.221 point-to-point
    desc link to R7 in wg 2
    ip address 192.168.2.130 255.255.255.248
    ip router isis
    mpls label protocol ldp
    tag-switching ip
    mpls traffic-eng tunnels
    frame-relay class FRAME-OUT
    ip rsvp bandwidth 128 128
    frame-relay interface-dlci 221

interface Serial0/0.222 point-to-point
    desc link to R8 in wg 2
    ip address 192.168.2.146 255.255.255.248

```

```
ip router isis
mpls label protocol ldp
tag-switching ip
mpls traffic-eng tunnels
ip rsvp bandwidth 128 128
  frame-relay class FRAME-OUT
frame-relay interface-dlci 222

router isis
net 49.0001.0000.0000.0009.00
metric-style wide
mpls traffic-eng router-id Loopback0
mpls traffic-eng level-2
exit

line con 0
  exec-timeout 999 0
  logging synchronous
line aux 0
line vty 0 4
  no login

exit
exit

copy run start
```

Post-Course Evaluations

Cisco uses a post-course evaluation system, Metrics That Matter (MTM), for its instructor-led courses. The instructor must ensure that each student is aware of a confidential evaluation process and that all students submit an evaluation for each course. There are two options for students to complete the evaluation.

For Classes with Internet Access

A URL will be made available specific to each CLP. Obtain the URL from your respective MTM system administrator before the last day of class.

1. Upon completion of the course, instruct the students to enter the URL into their browser.
2. Make sure that the students input their e-mail address (used only for a follow-up evaluation).

Note	Sixty days following a learning event, students will receive a brief, follow-up evaluation, and, again, responses will be kept confidential. E-mail addresses will not be used for marketing purposes. (If students do not have e-mail addresses, they may type in a “dummy” address.)

3. Instruct the students to select the appropriate course from the drop-down list.
4. Instruct the students to complete the course evaluation and click **Submit** one time only.
5. Advise the students to wait for “Thank you” to appear on screen before leaving.

For Classes Without Internet Access

A paper-based version of the post-course evaluation is available. Your MTM system administrator can provide you with copies.

1. Distribute paper-based evaluations at the beginning of the last day of class.
2. Instruct the students to complete the survey only after completing the course.
3. Collect the evaluations and submit them to your MTM system administrator.

To View Evaluation Results

To view your post-course evaluation results:

1. Go to www.metricsthatmatter.com/client. (Reminder: *All* data is confidential; you will see only your own data.)
2. Log in using your ID and the password sent to you from MTM or provided by your company MTM system administrator to ensure confidentiality.
3. Choose Menu Option – Learner Evaluation Reports:
 - Evaluation Retrieval Tool
 - Class Evaluation Summary Report
4. Search for and select the appropriate class.

Copyright © 2004, Cisco Systems, Inc. All rights reserved.

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



Copyright © 2004 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered Network* mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, iQ logo, the iQ Net Readiness Scorecard, LightStream, Linksys, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0402R)

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED "AS IS." CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.