SPCORE

Implementing Cisco Service Provider Next-Generation Core Network Services

Student Guide Volume 1

Version 1.2

Part Number: 97-3426-01

https://t.me/learningnets

Americas Headquarters	Asia Pacific Headquarters	Europe Headquarters
Cisco Systems, Inc. San Jose, CA	Cisco Systems (USA) Pte. Ltd. Singapore	Cisco Systems International BV Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the United States and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned 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. (1110R)

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, NONINFRINGEMENT 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.

https://t.me/learningnets

Welcome Students

Students, this letter describes important course evaluation access information!

Welcome to Cisco Systems Learning. Through the Cisco Learning Partner Program, Cisco 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 that 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. Please 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

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

https://t.me/learningnets

Table of Contents

Course Introduction 1
- Overview 1
- Course Goal and Objectives 3
- Course Flow 4
- Your Training Curriculum 5
- Additional References 8

Module 1: Multiprotocol Label Switching 9

Lesson 1: Introducing Cisco IP NGN Architecture 11
- Traditional ISP versus Traditional Telco 12
- Modern Service Provider 14
- Cisco IP NGN Architecture 15
- Basic Transport Technologies Overview 17
- SONET/SDH 18
- DWDM and ROADM 19
- IP over DWDM 20
- 10/40/100 Gigabit Ethernet Standards 21
- Transformation to IP 23
- Traditional IP Routing 25
- Summary 26

Lesson 2: Introducing MPLS 27
- MPLS Introduction 28
- MPLS Features 29
- MPLS Benefits 30
- MPLS Terminology 31
- MPLS Architecture: Control Plane 33
- MPLS Architecture: Data Plane 34
- Forwarding Structures 35
- MPLS Architecture Example 37
- MPLS Labels 38
- MPLS Static Feature 40
- MPLS Packet Flow Basic Example 41
- Summary 44

Lesson 3: Introducing MPLS Applications 45
- MPLS Applications 46
- MPLS Label Stack 47
- MPLS Unicast IP Routing 49
- MPLS Multicast IP Routing 50
- MPLS VPNs 51
- Layer 3 MPLS VPNs 53
- Layer 2 MPLS VPNs 56
- MPLS Traffic Engineering 59
- MPLS QoS 61
- Interaction Between MPLS Applications 62
- Summary 63

Lesson 4: Label Distribution Protocol 65
- Label Distribution Protocol 66
- LDP Neighbor Adjacency Establishment 68
- LDP Link Hello Message 69
- LDP Neighbor Discovery 70

https://t.me/learningnets

LDP Session Negotiation 71
LDP Discovery of Nonadjacent Neighbors 72
LDP Session Protection 73
LDP Graceful Restart and Nonstop Routing 74
Summary 76
Lesson 5: Introducing MPLS Forwarding 77
MPLS Forwarding Structures 78
Label-Switched Path 79
LSP Building 80
Label Allocation and Distribution 82
Label Allocation 84
Label Advertisement 85
Packet Propagation Across an MPLS Domain 88
Packet Propagation Without PHP 89
MPLS Steady State Condition 90
MPLS Label Control Methods 91
Impact of IP Aggregation on LSPs 92
Summary 94
Lesson 6: Operating MPLS Forwarding 95
Loop Detection using the MPLS TTL field 97
Normal TTL Operation 98
Disabling TTL Propagation 99
Steady State Operation 101
Link Failure MPLS Convergence Process 102
Routing Protocol Convergence 103
MPLS Convergence 104
Link Recovery MPLS Convergence Process 106
IP Switching Mechanisms 108
Standard IP Switching Example 109
Cisco Express Forwarding Switching Example 110
Cisco Express Forwarding in IOS XE Software 112
Cisco Express Forwarding in IOS XR 113
Monitoring IPv4 Cisco Express Forwarding 114
Summary 115
Lesson 7: Implementing MPLS in the Service Provider Core 117
MPLS Configuration on Cisco IOS XR versus Cisco IOS or IOS XE 119
MPLS Configuration Tasks 120
Basic MPLS Configuration 122
MTU Requirements 123
MPLS MTU Configuration 125
IP TTL Propagation 126
Disabling IP TTL Propagation 128
LDP Session Protection Configuration 130
LDP Graceful Restart and NSR Configuration 131
LDP IGP Synchronization Configuration 133
LDP Autoconfiguration 134
Label Advertisement Control 135
Label Acceptance Control 138
Monitor MPLS 140
Display MPLS LDP Parameters 141
Display MPLS Configured Interfaces 143
Display Status of LDP Discovery Process 144
LDP Monitoring Commands 146
Display the Status of the LDP Session 147

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Display the Contents of the LIB Table 149
Monitor Label Switching 151
Display the Contents of the LFIB Table 152
Display the Contents of the FIB Table....154
Debugging MPLS and LDP 155
Classic Ping and Traceroute 156
MPLS Ping and Traceroute Overview 158
MPLS Ping....159
MPLS Traceroute....162
Troubleshooting MPLS 163
LDP Session Startup Issues....164
Label Distribution Issues 166
Intermittent MPLS Failures After Interface Failure....167
Packet Propagation Issues 168
Summary....169
Lesson 8: Module Summary 171
Lesson 9: Module Self-Check 173
Module Self-Check Answers....179
Answer Key 179
Module 2: MPLS Traffic Engineering 181
Lesson 1: Introducing MPLS Traffic Engineering Components 183
Traffic Engineering Concepts....185
Traffic Engineering Motivations....186
Business Drivers for Traffic Engineering....187
Congestion Avoidance and Traffic Engineering 189
Traffic Engineering with a Layer 2 Overlay Model....190
Traffic Engineering with a Layer 2 Overlay Model: Example 191
Drawbacks of the Layer 2 Overlay Solution....192
Layer 3 Routing Model Without Traffic Engineering....193
Traffic Engineering with a Layer 3 Routing Model....194
Traffic Engineering with the MPLS TE Model....195
MPLS TE Traffic Tunnels 197
Traffic Tunnels: Attributes....200
Link Resource Attributes 202
Constraint-Based Path Computation....203
Example of Constraint-Based Path Computation 205
MPLS TE Process 207
Role of RSVP in Path Setup Procedures 209
Path Setup and Admission Control with RSVP 210
Forwarding Traffic to a Tunnel....212
Autoroute 213
Autoroute Example 214
Summary....216
Lesson 2: MPLS Traffic Engineering Operations 217
Attributes Used by Constraint-Based Path Computation....219
MPLS TE Link-Resource Attributes 220
MPLS TE Link Resource Attributes: Maximum Bandwidth and Maximum Reservable Bandwidth...221
MPLS TE Link-Resource Attributes: Link-Resource Class 223
MPLS TE Link-Resource Attributes: Constraint-Based Specific Link Metric (Administrative Weight) 224
MPLS TE Tunnel Attributes 225
MPLS TE Tunnel Attributes: Traffic Parameter and Path Selection and Management....226
MPLS TE Tunnel Attributes: Tunnel Resource Class Affinity 227

https://t.me/learningnets

MPLS TE Tunnel Attributes: Adaptability, Priority, Pre-emption 228
MPLS TE Tunnel Attributes: Resilience....229
Implementing TE Policies with Affinity Bits 230
Using Affinity Bits in TE Policies....231
Propagating MPLS TE Link Attributes with Link-State Routing Protocol....234
Constraint-Based Path Computation....237
Constraint-Based Path Computation: Path Selection 239
Example of Constraint-Based Path Computation 241
Path Setup 244
RSVP Usage in Path Setup....246
Hop-by-Hop Path Setup with RSVP....247
Tunnel and Link Admission Control 251
Path Rerouting....253
Path Reoptimization 254
Path Rerouting: Link Failure 256
MPLS TE Soft Pre-emption 258
MPLS TE Path-Option Attributes 259
Assigning Traffic to Traffic Tunnels 260
Using Static Routing to Assign Traffic to Traffic Tunnel 261
Autoroute....262
Autoroute: Path Selection Rules 263
MPLS TE IPv6 Autoroute....264
Autoroute: Default Metric....265
Autoroute: Relative and Absolute Metric....268
Forwarding Adjacency....270
Forwarding Adjacency Traffic Flows....271
Summary....274
Lesson 3: Implementing MPLS TE....277
MPLS TE Configuration Tasks....279
MPLS TE Configuration....280
RSVP Configuration....281
OSPF Configuration....283
IS-IS Configuration....285
MPLS TE Tunnels Configuration 287
Static Route and Autoroute Configurations....289
Verifying RSVP Operation 290
Verify MPLS TE Tunnels 291
Verify a Routing Table with Autoroute Enabled....295
MPLS TE Case Study....296
MPLS TE Case Study: Dynamic MPLS TE Tunnel....299
MPLS TE Case Study: Explicit MPLS TE Tunnel....304
MPLS TE Case Study: Periodic Tunnel Optimization....306
MPLS TE Case Study: Configuring Path-Selection Restrictions 308
MPLS TE Case Study: Modifying the Administrative Weight 312
MPLS TE Case Study: Autoroute and Forwarding Adjacency 314
Summary....318
Lesson 4: Protecting MPLS TE Traffic 319
Improving MPLS TE Convergence Time 321
Configuring Backup MPLS TE Tunnels....322
Drawbacks of Backup MPLS TE Tunnels....324
Fast Reroute Case Study 325
Fast Reroute Case Study: Link Protection....327
Fast Reroute Case Study: Node Protection....329
Fast Reroute Case Study: Fast Reroute Link Protection Configurations 331

https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Backup 334
MPLS Traffic Engineering Autotunnel Backup Configuration 335
MPLS Traffic Engineering Autotunnel Backup Verification 336
MPLS Traffic Engineering Autotunnel Mesh 337
MPLS Traffic Engineering Autotunnel Mesh Configuration 338
MPLS TE End-to-End Path Protection 340
MPLS TE End-to-End Path Protection Configuration 342
BFD over MPLS TE LSPs 343
MPLS TE Bandwidth Control 345
Automatic Bandwidth Configuration 348
DiffServ-Aware MPLS TE Tunnels 350
DiffServ-Aware MPLS Dual-Bandwidth Pools 353
DiffServ-Aware Service Configuration 354
Summary 358
Lesson 5: Module Summary 361
Lesson 6: Module Self-Check 363
Module Self-Check Answers 369
Answer Key 369
Glossary 371

https://t.me/learningnets

https://t.me/learningnets

Course Introduction

Overview

Implementing Cisco Service Provider Next-Generation Core Network Services (SPCORE) v1.2 is an instructor-led course presented by Cisco Learning Partners to their end-user customers. This five-day course provides network engineers and technicians with the knowledge and skills necessary to implement and support a service provider network.

The course is designed to provide service provider network professionals with the information that they need to use technologies in a service provider core network. The goal is to provide network professionals with the knowledge, skills, and techniques that are required to plan, implement, and monitor a service provider core network.

The course also features classroom activities, including remote labs, to teach practical skills on deploying Cisco IOS, IOS XE, and IOS XR features to operate and support a service provider network.

Learner Skills and Knowledge

This subtopic lists the skills and knowledge that learners must possess to benefit fully from the course. The subtopic also includes recommended Cisco learning offerings that learners should first complete to benefit fully from this course.

https://t.me/learningnets

Learner Skills and Knowledge

- Students considered for this training will have attended the following courses or obtained equivalent level training:
  - *Building Cisco Service Provider Next-Generation Networks Part 1* (SPNGN1) v1.2
  - *Building Cisco Service Provider Next-Generation Networks Part 2* (SPNGN2) v1.2
  - *Deploying Cisco Service Provider Network Routing* (SPROUTE) v1.2
  - *Deploying Cisco Service Provider Advanced Network Routing* (SPADVROUTE) v1.2

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Course Goal and Objectives

This topic describes the course goal and objectives.

Upon completing this course, you will be able to:

- Describe the features of MPLS, and how MPLS labels are assigned and distributed
- Discuss the requirement for traffic engineering in modern networks that must attain optimal resource utilization
- Describe the concept of QoS and explain the need to implement QoS
- Classify and mark network traffic to implement an administrative policy requiring QoS
- Compare the different Cisco QoS queuing mechanisms that are used to manage network congestion
- Explain the concept of traffic policing and shaping, including token bucket, dual token bucket, and dual-rate policing

https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.

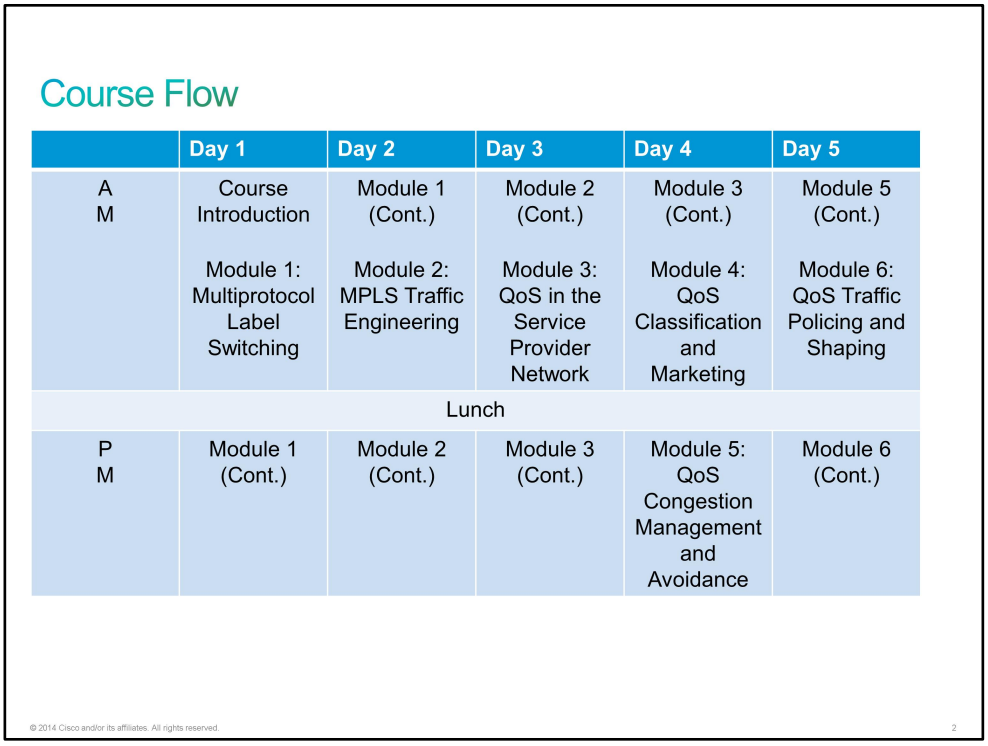

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

https://t.me/learningnets

Your Training Curriculum

This topic presents the training curriculum for this course.

You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, CCNP® Security and CCNP® Voice, or CCSP™). It provides a gathering place for Cisco certified professionals to share questions, suggestions, and information about Cisco Career Certification programs and other certification-related topics. For more information, visit http://www.cisco.com/go/certifications.

https://t.me/learningnets

Training Curriculum

This subtopic presents the training curriculum for this course.

For more information on certifications, visit http://www.cisco.com/go/certifications.

Note	Cisco provides three levels of general certifications for IT professionals with several different tracks to meet individual needs. Cisco also provides focused certifications for designated areas such as cable communications 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 www.cisco.com/go/certifications.

https://t.me/learningnets

Learner Introductions

- Your name
- Your company
- Job responsibilities
- Skills and knowledge
- Brief history
- Objective

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Additional References

This topic presents the Cisco icons and symbols that are used in this course, as well as information on where to find additional technical references.

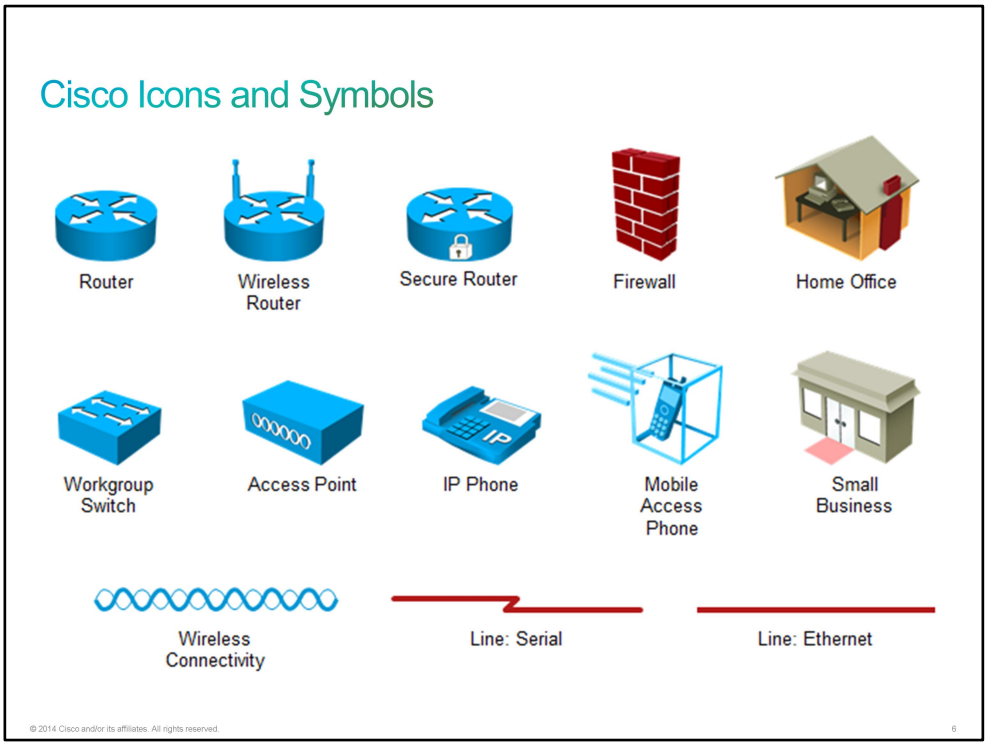

Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the Cisco Internetworking Terms and Acronyms glossary of terms at http://docwiki.cisco.com/wiki/Category:Internetworking_Terms_and_Acronyms_(ITA).

https://t.me/learningnets

Module 1: Multiprotocol Label Switching

Introduction

This module explains the features of MPLS compared with those of traditional hop-by-hop IP routing. MPLS concepts and terminology are explained in this module, along with MPLS label format and LSR architecture and operations. This module describes the assignment and distribution of labels in an MPLS network, including neighbor discovery and session establishment procedures. Label distribution, control, and retention modes will also be explained. This module also covers the functions and benefits of PHP and provides a review of switching implementations, focusing on Cisco Express Forwarding.

The module also covers the details of implementing MPLS on Cisco IOS, IOS XE, and IOS XR platforms, giving detailed configuration, monitoring, and debugging guidelines for a typical service provider network. In addition, this module includes the advanced topics of controlling TTL propagation and label distribution.

Upon completing this module, you will be able to:

- Discuss the basic concepts and architecture of Cisco IP NGN Architecture
- Discuss the basic concepts and architecture of MPLS
- Describe basic MPLS applications
- Describe LDP process and operation in a service provider network
- Describe how label allocation and distribution function in an MPLS network
- Describe how to operate MPLS forwarding
- Configure MPLS on Cisco IOS, IOS XE, and Cisco IOS XR platforms

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Introducing Cisco IP NGN Architecture

Overview

The Cisco IP Next-Generation Network is the next-generation service provider architecture for providing voice, video, mobile, and cloud or managed services to users. This lesson introduces Cisco IP NGN Architecture.

Upon completing this lesson, you will be able to:

- Describe a traditional telco and a traditional ISP
- Describe a modern service provider
- Show the Cisco IP NGN architecture
- Describe the basic service provider transport technologies in the core portion of the Cisco IP NGN model
- Describe SONET/SDH
- Describe DWDM and ROADM
- Describe IPoDWDM
- Describe the 10/40/100 Gigabit Ethernet standards
- Describe the traditional service provider transformation to IP
- Describe traditional IP routing where packet forwarding discussions are based on the IP address

https://t.me/learningnets

Traditional ISP versus Traditional Telco

When talking about the types of communications service providers, discussing the evolution, characteristics, and pros and cons of traditional telco and traditional ISPis also necessary.

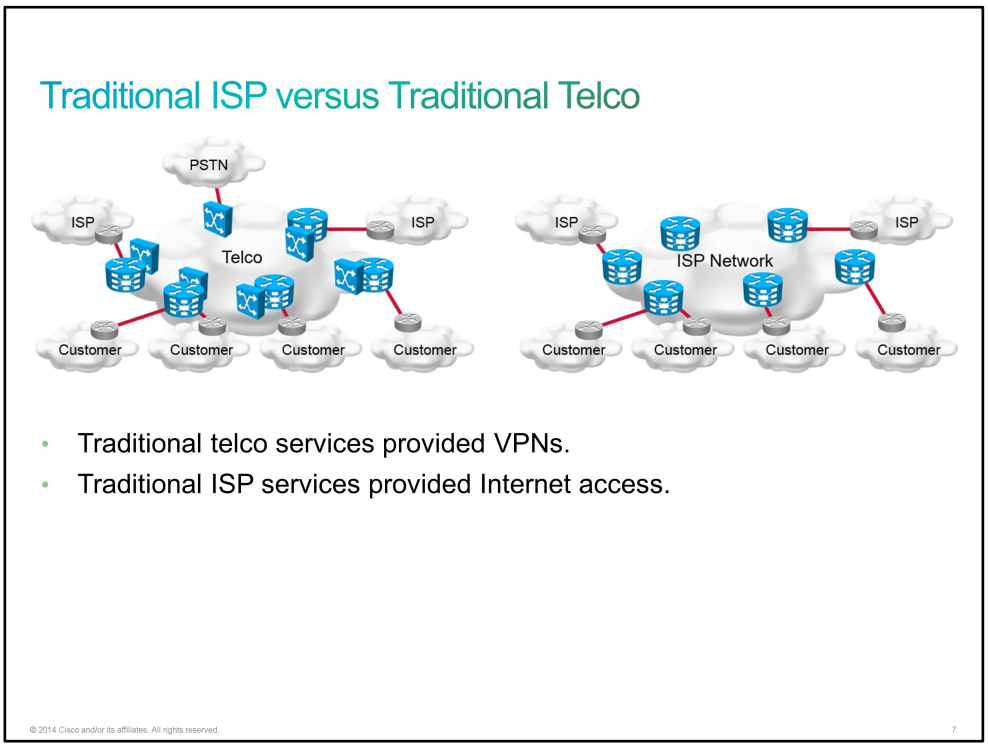

The figure illustrates two types of networks:

- On the left is a traditional telco network that comprises many different devices to offer various services. ATM switches were used to provide VPNs to customers. TDM switches were used to provide circuits or telephony to customers. SONET/SDH was used to carry ATM and TDM across an optical network. Routers were used to provide Internet access to customers.
- On the right is a traditional ISP whose initial focus was only to provide Internet access to their customers. No other services were offered because there was limited capability to offer anything comparable to what telcos could offer through their extensive range of equipment and technologies.

There were advantages with traditional ISPs:

- Employed one technology–IP (cost-effective)
- Had competitive pricing

There were limitations to traditional ISPs:

- Provided only Internet access
- Had no VPNs

https://t.me/learningnets

There were advantages with traditional telcos:

- Employed multiple technologies (IP, ATM, Frame Relay, SONET/SDH)
- Used good QoS mechanisms
- Offered Internet access, VPNs, and telephony

There were limitations with traditional telcos:

- Employed multiple technologies (costly)
- Provided costly services

https://t.me/learningnets

Modern Service Provider

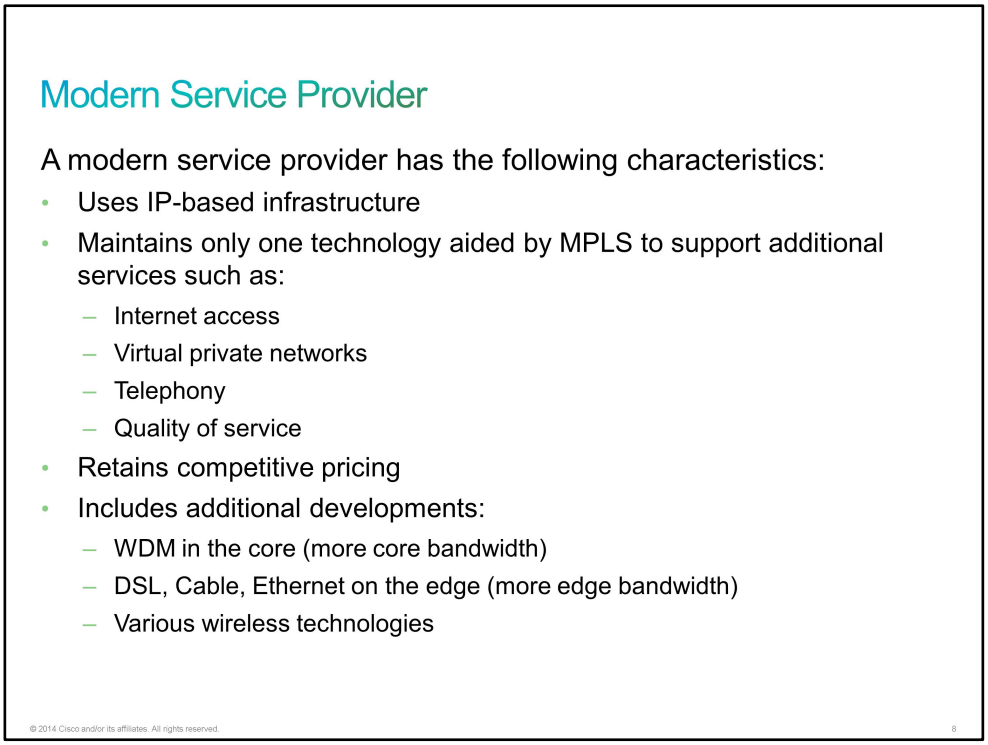

A modern service provider network, whether it evolved from a traditional telco or from a greenfield ISP, can accommodate the same customer requirements as traditional telcos did, without having to use different types of networks and devices. Routers are used to provide Internet access, VPNs, telephony services, and TV services. DWDM is an exception that is often used in addition to routers to increase the amount of throughput that is available via a single strand of optical fiber.

https://t.me/learningnets

Cisco IP NGN Architecture

In earlier days, service providers were specialized for different types of services, such as telephony, data transport, and Internet service. The popularity of the Internet, through telecommunications convergence, has evolved into the usage of the Internet for all types of services. Development of interactive mobile applications, increasing video and broadcasting traffic, and the adoption of IPv6 have pushed service providers to adopt new architecture to support new services on the reliable IP infrastructure with a good level of performance and quality.

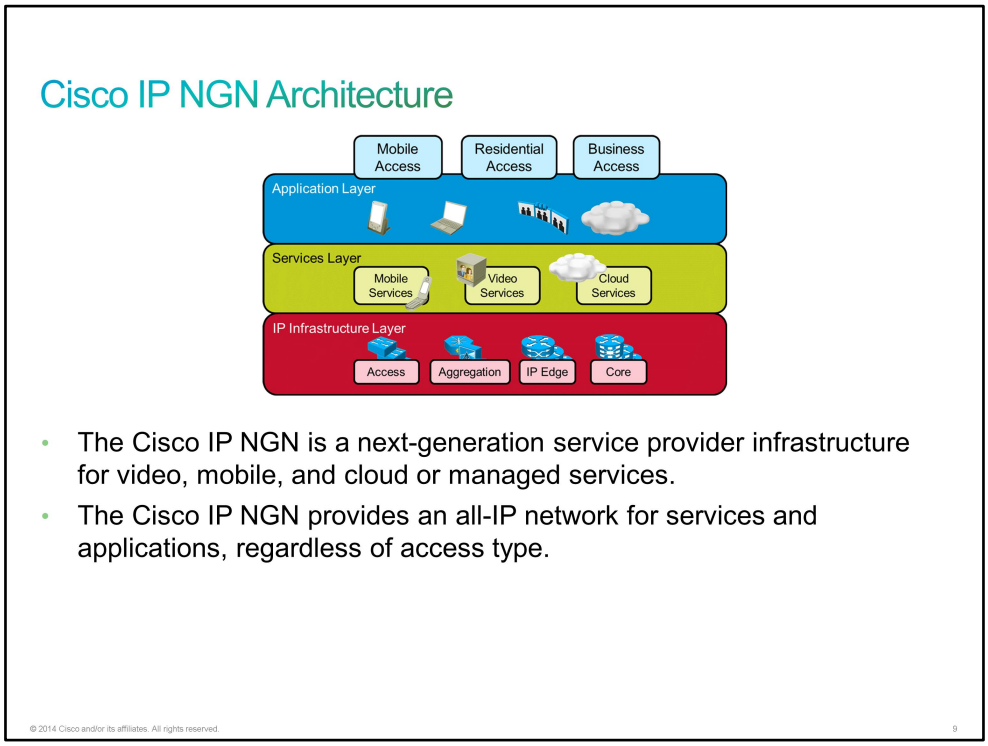

Cisco IP Next-Generation Network is the next-generation service provider architecture for providing voice, video, mobile, and cloud or managed services to users. The general idea of Cisco IP NGN is to provide all-IP transport for all services and applications, regardless of access type. IP infrastructure, service, and application layers are separated in NGN networks, thus enabling addition of new services and applications without any changes in the transport network.

https://t.me/learningnets

Cisco IP NGN Infrastructure Layer

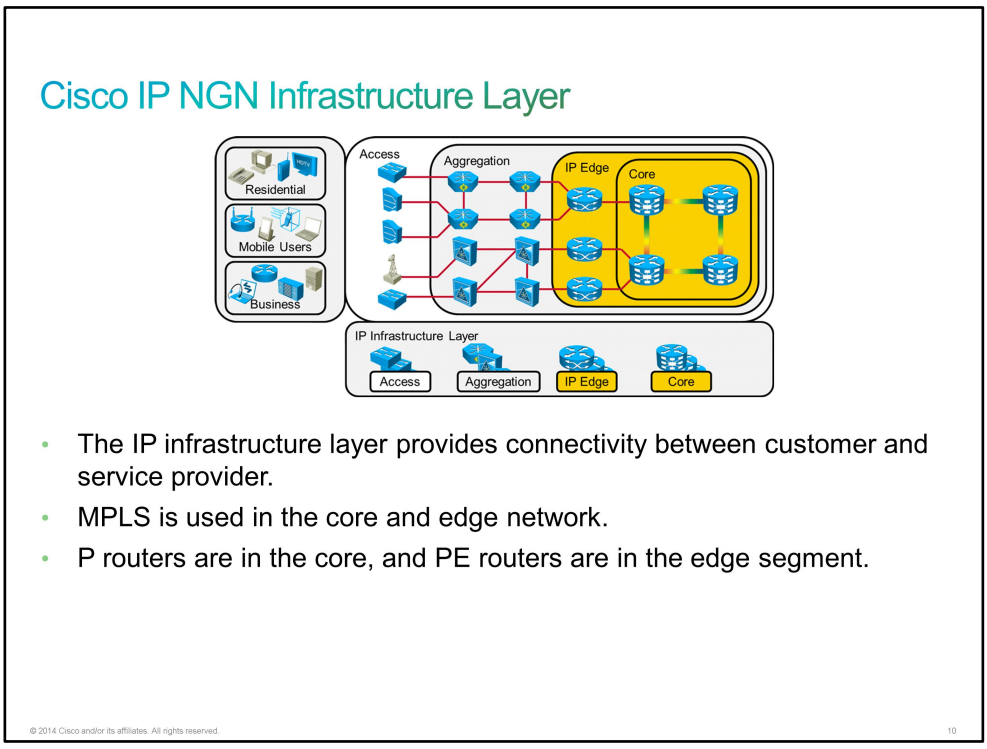

The IP infrastructure layer is responsible for providing reliable infrastructure for running upper layer services. It includes the following:

- Core network
- IP edge network
- Aggregation network
- Access network

It provides the reliable, high-speed, and scalable foundation of the network. End users are connected to a service provider through a CPE device, using any possible technology. Access and aggregation network devices are responsible for enabling connectivity between customer equipment and service provider edge equipment. A core network is used for fast-switching packets between edge devices.

MPLS is a technology that is primarily used in the service provider core and the IP edge portion of the IP infrastructure layer. In service provider networks, the result of using MPLS is that only the routers on the edge of the MPLS domain perform routing lookup; all other routers forward packets based on labels. What really makes MPLS useful in service provider (and large enterprise) networks is that it enhances BGP routing and provides different services and applications, such as Layer 2 and 3 VPNs, QoS, and traffic engineering.

https://t.me/learningnets

Basic Transport Technologies Overview

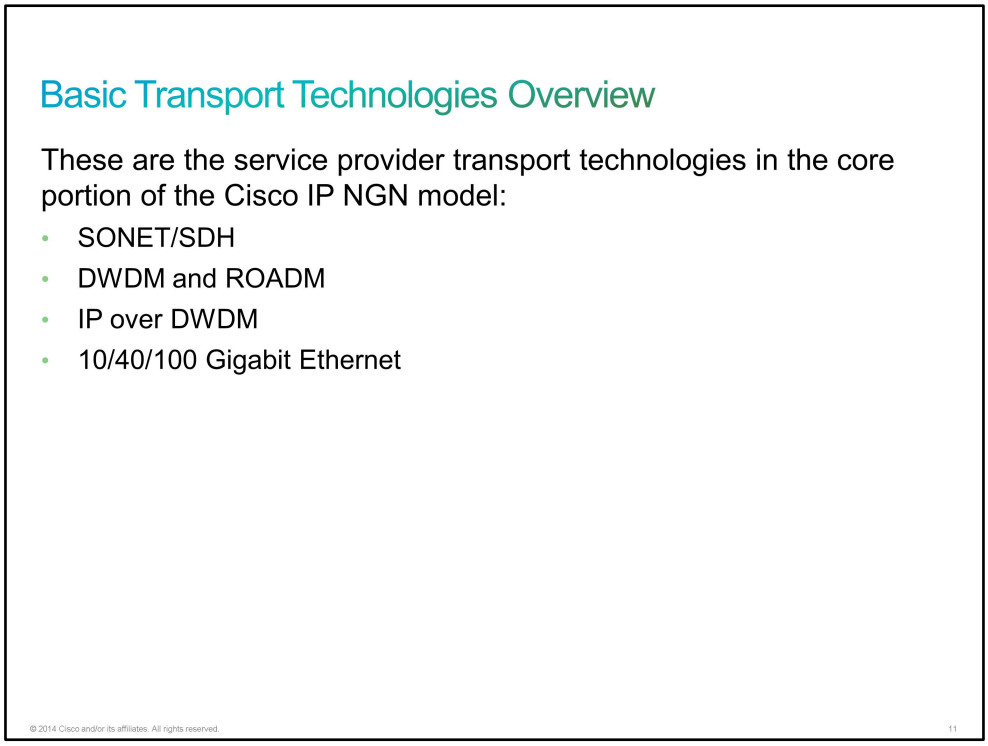

SONET is a set of standards that defines optical signals and a synchronous frame structure for multiplexed digital traffic. A similar standard, SDH, is used in Europe by the International Telecommunication Union Telecommunication Standardization Sector. As SONET and SDH are technically comparable standards, the term SONET is often used to refer to either. Both SONET and SDH are based on a structure that has a basic frame format and speed. Both SONET and SDH have a hierarchy of signaling speeds. Multiple lower-level signals can be multiplexed to form higher-level signals.

In many cases, networks, which have traditionally been based on SONET or SDH technology, are experiencing fiber exhaust as a result of high bandwidth demand. DWDM increases the carrying capacity of the physical medium (fiber); one fiber is transformed into multiple virtual fibers, thus increasing bandwidth over the existing fiber optic backbone. Besides bandwidth, the advantages of DWDM are transparency, scalability, and dynamic provisioning. ROADM allows for switching traffic remotely, dropping or adding wavelengths without having to remove all the active traffic from DWDM system.

The traditional optical network model was based on several layers, including SONET/SDH in addition to IP and DWDM. Implementing IP over DWDM means cutting these intermediate layers and thus reducing operating expenses and capital expenditures. This simplifies the network layout while offering the same features such as grooming, protection, and restoration.

Gigabit Ethernet is a proven technology for easy migration from and integration into traditional Ethernet. As a high-speed Ethernet, 10/40/100 Gigabit Ethernet is being driven by need to interconnect Ethernet LANs operating at lower speeds and can be used for aggregating slower access links in the backbone networks and for WAN access. The standards for 10 Gigabit Ethernet are defined by IEEE 802.3ae and the standards for 40/100 Gigabit Ethernet by IEEE 802.ba.

https://t.me/learningnets

SONET/SDH

SONET/SDH was initially designed to carry 64K PCM voice channels that are commonly used in telecommunication. The basic underlying technology that is used in the SONET/SDH system is synchronous TDM.

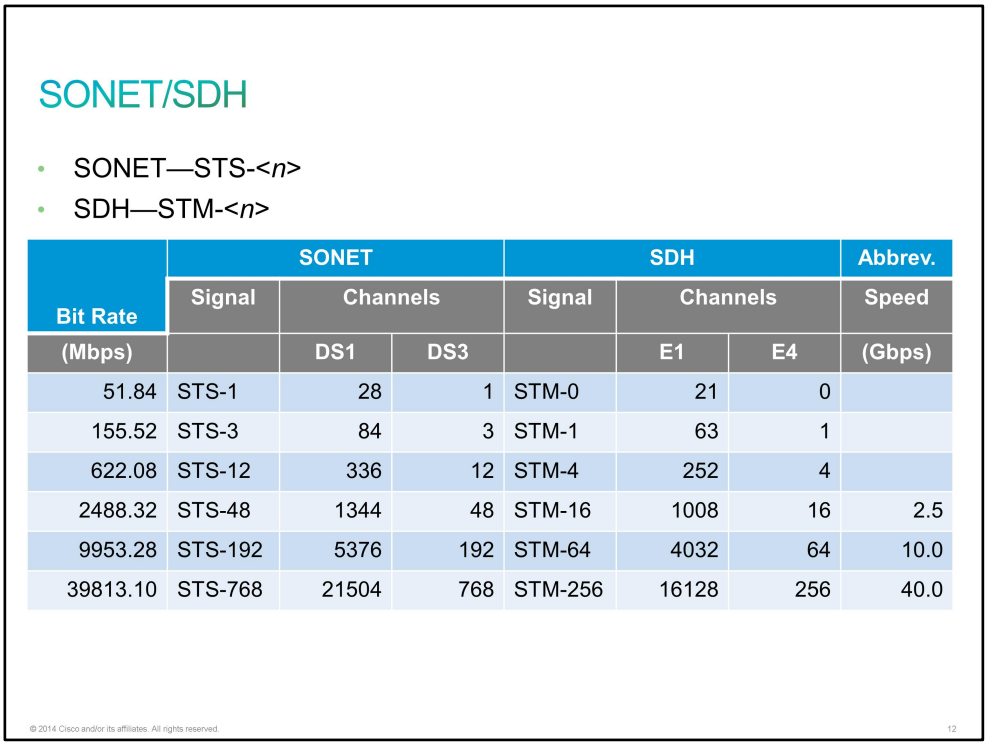

Bit Rate	SONET			SDH			Abbrev.
	Signal	Channels		Signal	Channels		Speed
(Mbps)		DS1	DS3		E1	E4	(Gbps)
51.84	STS-1	28	1	STM-0	21	0	
155.52	STS-3	84	3	STM-1	63	1	
622.08	STS-12	336	12	STM-4	252	4	
2488.32	STS-48	1344	48	STM-16	1008	16	2.5
9953.28	STS-192	5376	192	STM-64	4032	64	10.0
39813.10	STS-768	21504	768	STM-256	16128	256	40.0

The major difference between SONET and SDH is the terminology that is used to describe them. For example, a SONET OC-3 signal is called an SDH STM-1 signal by the ITU-T.

The SONET/SDH standard specifies standards for communication over fiber optics as well as electrical carriers for lower-speed signaling rates (up to 155 Mbps). The standard describes the frame format that should be used to carry the different types of payload signals as well as the control signaling that is needed to keep a SONET/SDH connection operational.

The SONET standard is mainly used in the United States, while the SDH standard is mainly European. In the United States, the SDH standard is used for international connections.

https://t.me/learningnets

DWDM and ROADM

DWDM and ROADM systems are the fiber-optic technologies used in the core portion of the Cisco IP Next-Generation Network model.

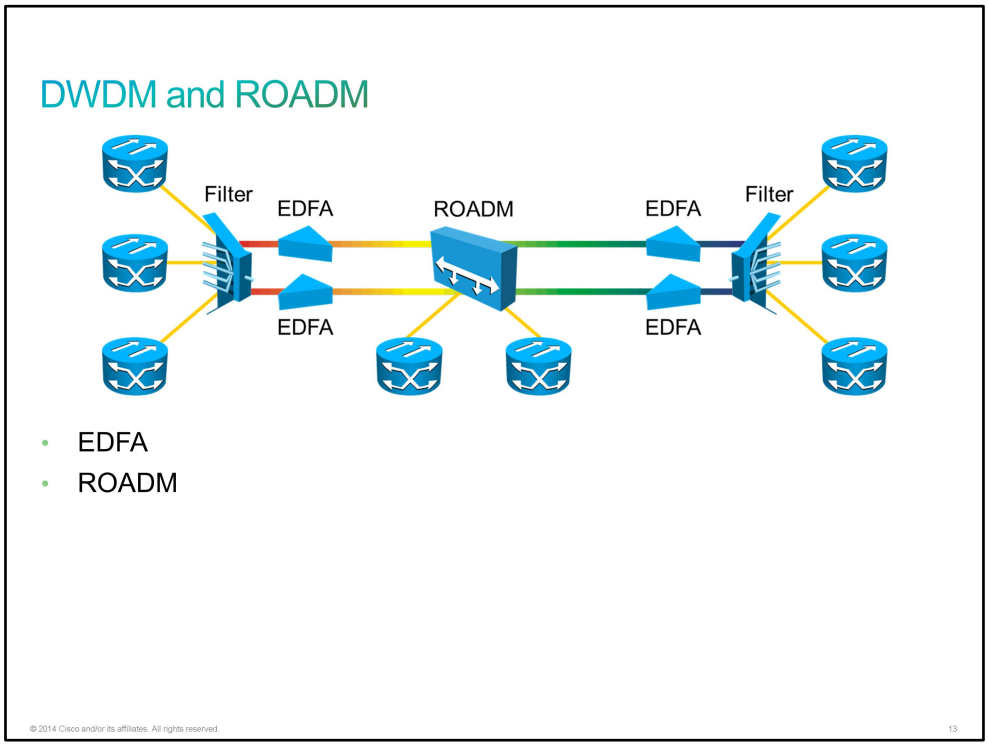

WDM is a technology that multiplexes a number of optical carrier signals into a single optical fiber by using different wavelengths of laser light. Dense WDM, or DWDM, refers to optical signals that are multiplexed within the 1550 nm band so as to leverage the capabilities and cost of EDFAs.

EDFA is a device that amplifies an optical fiber signal. Fiber amplifiers are developed to support DWDM and to expand to the other wavelength bands supported by fiber optics. EDFAs are effective for wavelengths between approximately 1525–1565 nm, or 1570–1610 nm. EDFA fiber optic amplifiers function by inserting a trace of impurity (trivalent erbium ion) to the fiber core material as a dopant. This insertion alters the optical properties of the core and permits signal amplification.

Intermediate optical amplification sites in DWDM systems may allow for the dropping and adding of certain wavelength channels. In earlier systems, adding or dropping wavelengths required manually inserting or replacing wavelength-selective cards. This operation was costly and in some systems required that all active traffic be removed from the DWDM system, because inserting or removing the wavelength-specific cards interrupts the multi-wavelength optical signal.

A ROADM is a form of OADMthat adds the ability to switch traffic remotely, from a WDM system at the wavelength layer. This capability allows individual or multiple wavelengths that are carrying data channels to be added or dropped from a transport fiber without the need to convert the signals on all of the WDM channels to electronic signals, and back again to optical signals.

With a ROADM, network operators can reconfigure the multiplexer remotely by sending soft commands. The architecture of the ROADM is such that dropping or adding wavelengths does not interrupt the “pass-through” channels. Numerous technological approaches are utilized for various commercial ROADMs, the trade-off being between cost, optical power, and flexibility.

https://t.me/learningnets

IP over DWDM

Service providers continue to look for the best economics for increasing network capacity to accommodate the continued growth in IP traffic that is driven by data, voice, and primarily video traffic. The reasons for integrating IP and DWDM are simply to deliver a significant reduction in capital expenditures and improve the operational efficiency of the network.

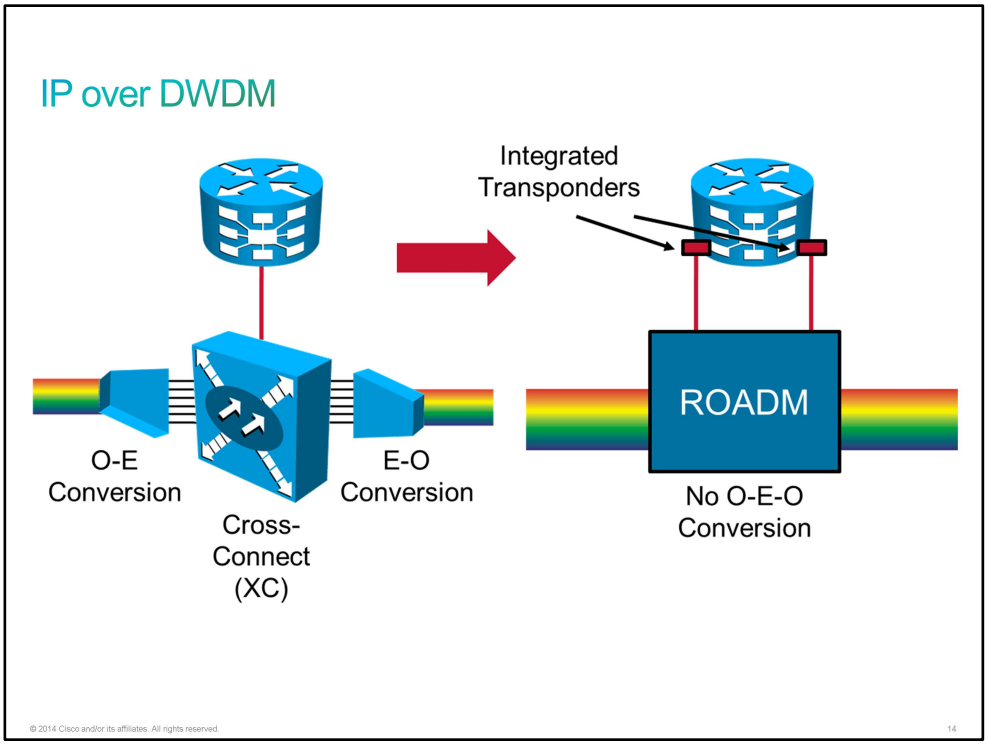

IPoDWDM is a technology that integrates DWDM on routers. Routers must support the ITU-T G.709 standard so that they can monitor the optical path. Element integration refers to the capability to take multiple, separate elements that operate in the network and collapse them into a single device without losing any of the desired functions for continued operation.

The integration of core routers with the optical transport platform eliminates the need for OEO modules (transponders) in the transport platform.

https://t.me/learningnets

10/40/100 Gigabit Ethernet Standards

IEEE 802.3ae is the IEEE standard of the 802.3 family of data link layer standards for Ethernet LAN and WAN applications, whose objective is to support 10 Gigabit Ethernet. 10 Gigabit Ethernet refers to various technologies for transmitting Ethernet frames at a rate of 10 Gbps.

IEEE 802.3ba is the IEEE standard, whose objective is to support speeds faster than 10 Gbps. The standard supports 40 Gbps and 100 Gbps transfer rates. The decision to include both speeds comes from the demand to support the 40 Gbps rate for local server applications and the 100 Gbps rate for Internet backbones.

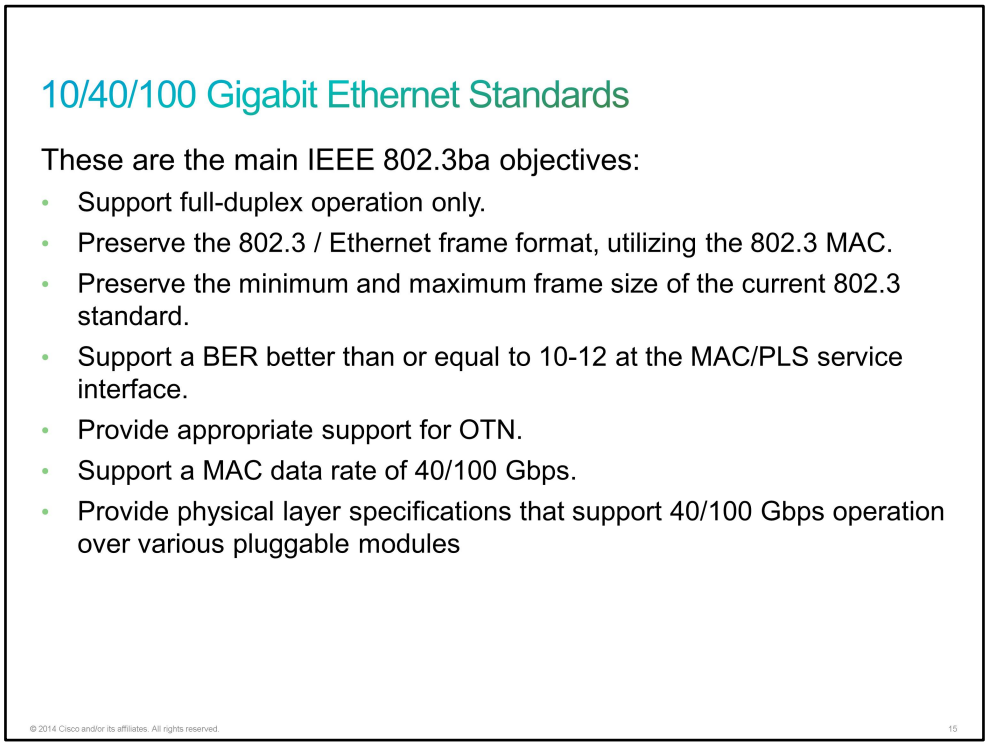

The 40/100 Gigabit Ethernet standards include a number of different Ethernet PHY specifications, so a networking device may support different pluggable modules.

The objective of IEEE 802.3ba is to provide physical layer specifications that support 40 Gbps operation over the following:

- At least 10 km on SMF
- At least 100 m on OM3 MMF
- At least 7 m over a copper cable assembly
- At least 1 m over a backplane

https://t.me/learningnets

The objective of IEEE 802.3ba is to provide physical layer specifications that support 100 Gbps operation over the following:

- At least 40 km on SMF
- At least 10 km on SMF
- At least 100 m on OM3 MMF
- At least 7 m over a copper cable assembly

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Transformation to IP

Traditional service provider architecture was based on numerous parallel services with a simple, stable SDH core, where each service was independent.

Modern service providers usually transform to IP protocol architecture, where everything runs on top of IP protocol. In this scenario, Ethernet replaces ATM or SDH; IP, in combination with MPLS, is used in the core network.

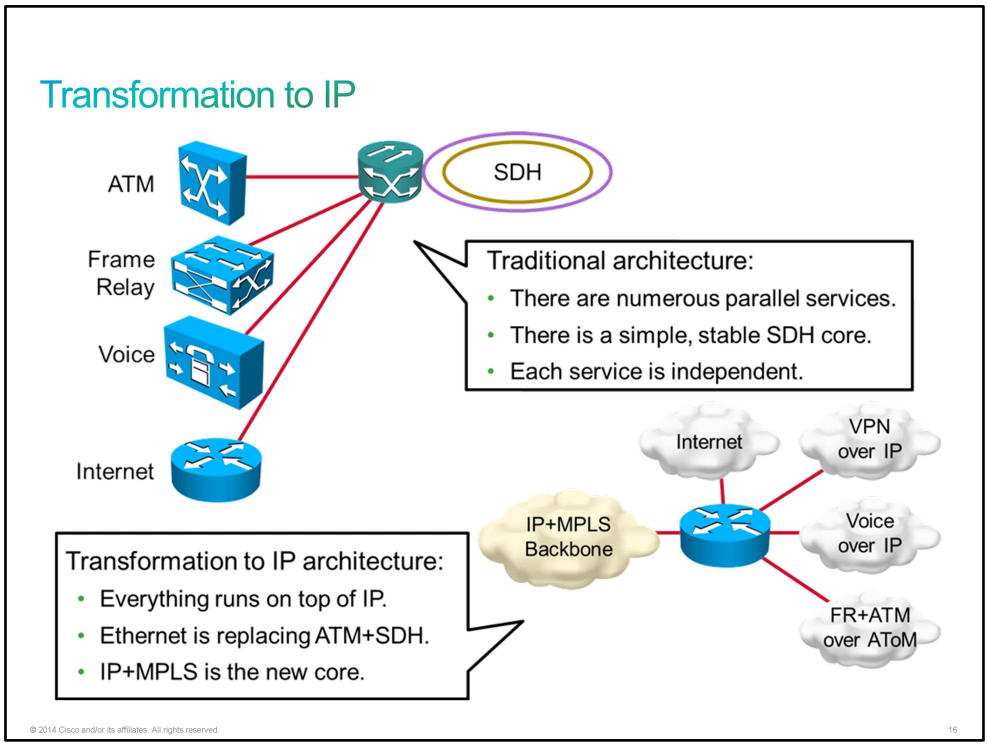

Among the factors that drive service providers to transition to IP could be changed usage patterns such as the following:

- Traditional traffic is growing slowly.
- IP traffic is exploding.
- "Everything-over-IP" makes business sense.
- Customers accept "just-good-enough" solutions.

https://t.me/learningnets

Among the factors that drive service providers to transition to IP could also be technology changes such as the following:

- Packet switching is cheaper than circuit switching.
- MPLS provides reasonable virtual circuit capabilities.
- MPLS recovery times are comparable to SDH.
- Ethernet is cheaper than ATM or SDH.
- Rigid QoS is not needed on very high-speed links.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Traditional IP Routing

In traditional IP routing, packet forwarding discussions are based on the IP address.

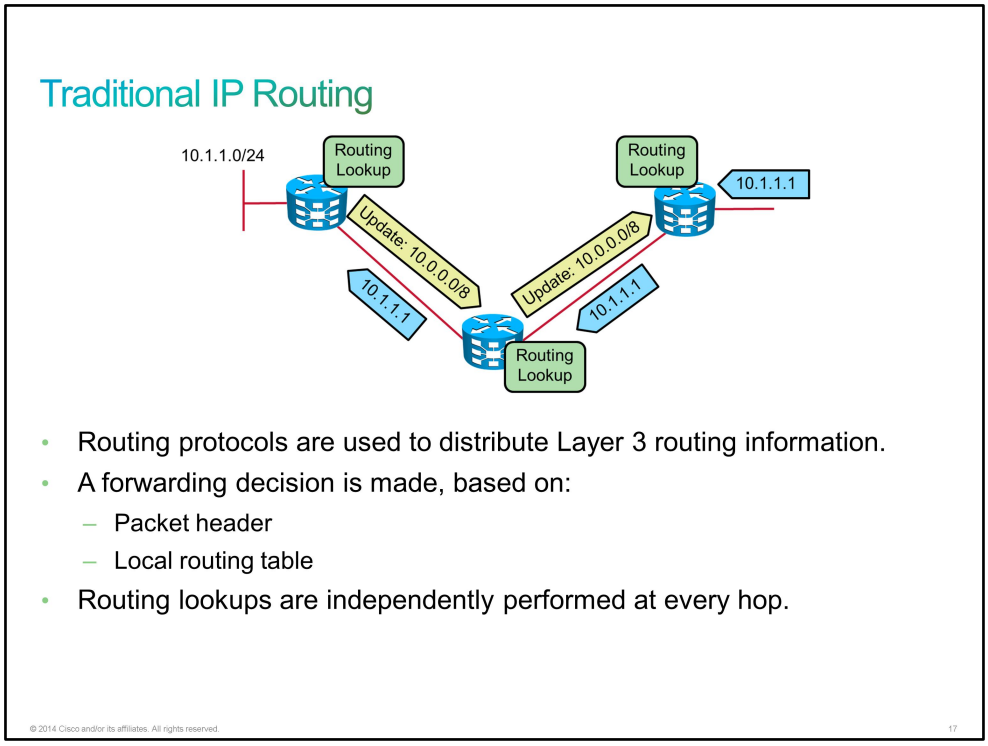

Before basic MPLS functionality is explained, these three foundations of traditional IP routing need to be highlighted:

- Routing protocols are used on all devices to distribute routing information.
- Each router analyzes the Layer 3 header of each packet, compares it to the local routing table, and makes a decision about where to forward the packet. Regardless of the routing protocol, routers forward packets contingent on a destination address-based routing lookup.
- The routing lookup is performed independently on every router in the network.

Note	The exception to this rule is PBR, where routers will bypass the destination-based routing lookup.

- The routing lookup is performed independently on every router in the network.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The traditional telco provides VPN services, telephony, and Internet access, while traditional ISP provides only Internet access.
- Modern service providers provide Internet access, VPN, telephony and QoS using IP and MPLS.
- Cisco IP NGN is the next-generation service provider architecture for providing voice, video, mobile, and cloud or managed services to users.
- The SONET standard is mainly used in the United States, while the SDH standard is mainly European.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- WDM is a technology that multiplexes a number of optical carrier signals into a single optical fiber by using different wavelengths of laser light.
- IPoDWDM is a technology that integrates DWDM on routers.
- IEEE 802.3ba is part of the 802.3 family of data-link layer standards for Ethernet LAN and WAN applications, whose objective is to support transmitting of Ethernet frames at a rate of 40Gbps and 100Gbps.
- Traditional service provider architecture was based on numerous parallel services with a simple, stable SDH core.
- Traditional IP routing forwards packets based on the destination IP address.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 2: Introducing MPLS

Overview

MPLS is a switching mechanism that is often found in service provider environments. MPLS leverages traditional IP routing and supports several services that are required in next-generation IP networks.

Upon completing this lesson, you will be able to:

- Describe MPLS technology
- Describe MPLS forwarding based on the MPLS label
- Describe the benefits of MPLS
- Describe the LSR, edge LSR, and LSP terminologies
- Describe the MPLS control plane
- Describe the MPLS data plane
- Describe the FIB and LFIB
- Describe an example of the protocols used in the MPLS control plane and the LFIB in the data plane
- Describe the MPLS label
- Describe the MPLS static feature
- Explain how an IP packet is forwarded using MPLS label switching

https://t.me/learningnets

MPLS Introduction

MPLS provides an intermediate encapsulation between an OSI Layer 3 IP header and an arbitrary OSI Layer 2 header. This function enables the forwarding of packets through LSPs that can be created using various methods and protocols, depending on the required results. Additionally, the payload can be any Layer 3 or even Layer 2 protocol.

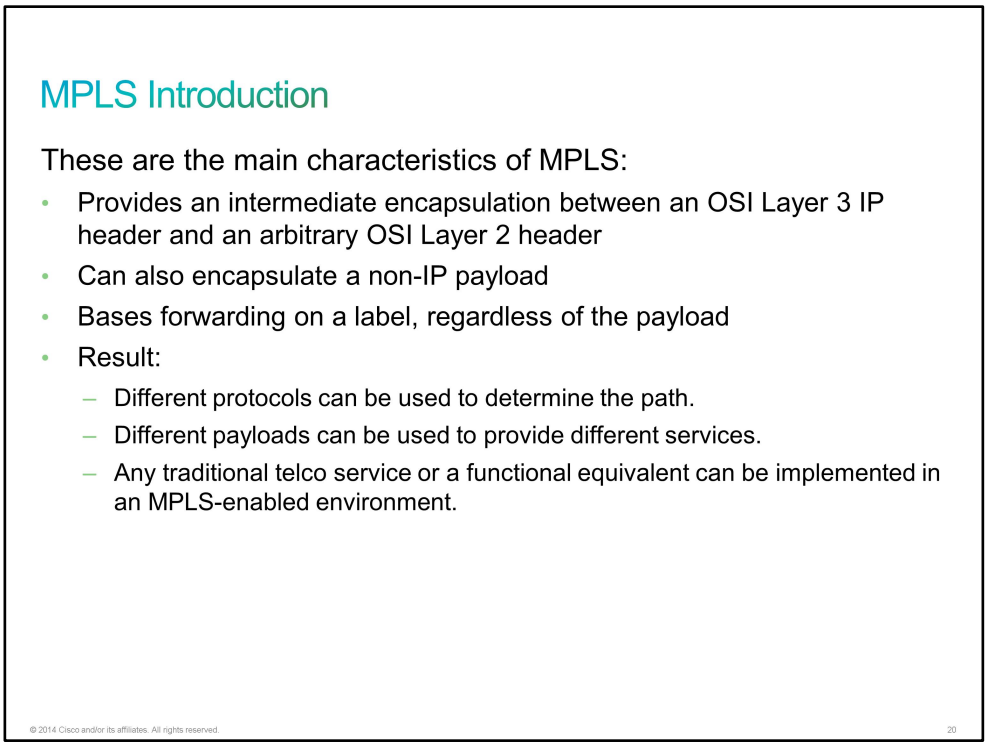

MPLS enables service providers to provide services that are the same or functionally equivalent to service that was available with traditional telcos:

- Internet access can be provided.
- ATM or Frame Relay VPNs can be replaced by Layer 3 MPLS VPNs, or if they are required, they can be retained using Layer 2 MPLS VPNs.
- SONET/SDH can be implemented using DWDM, or the same type of QoS characteristics can be implemented using Layer 2 MPLS VPNs in combination with QoS implementation.

https://t.me/learningnets

MPLS Features

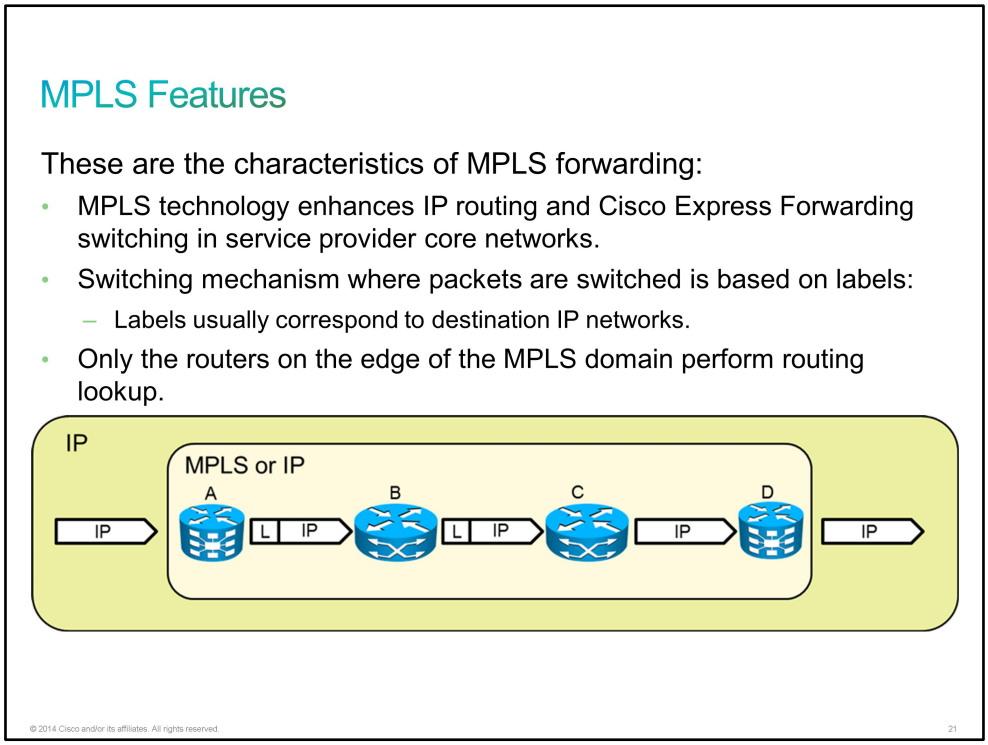

MPLS is a technology that is primarily used in service provider core networks. MPLS improves classic IP routing using Cisco Express Forwarding by introducing an additional header into a packet. This additional header is called the MPLS label. MPLS switches packets based on label lookup instead of IP address lookup. Labels usually correspond to destination IP networks; each destination has a corresponding label on each MPLS-enabled router.

In service provider networks, the result of using MPLS is that only the routers on the edge of MPLS domain perform routing lookup; all other routers forward packets that are based on labels.

MPLS is a packet-forwarding technology that uses appended labels to make forwarding decisions for packets.

- Within the MPLS network, the Layer 3 header analysis is done just once (when the packet enters the MPLS domain). Labels are appended to the packet, and then the packet is forwarded into the MPLS domain.
- Simple label inspection that is integrated with Cisco Express Forwarding switching drives subsequent packet forwarding.

https://t.me/learningnets

MPLS Benefits

In modern routers, MPLS label switching is not much faster than IP routing, but MPLS is not used just because of its switching performance. MPLS decreases the forwarding overhead on the core routers. MPLS also supports the forwarding of non-IP protocols, because MPLS technologies are applicable to any network layer protocol.

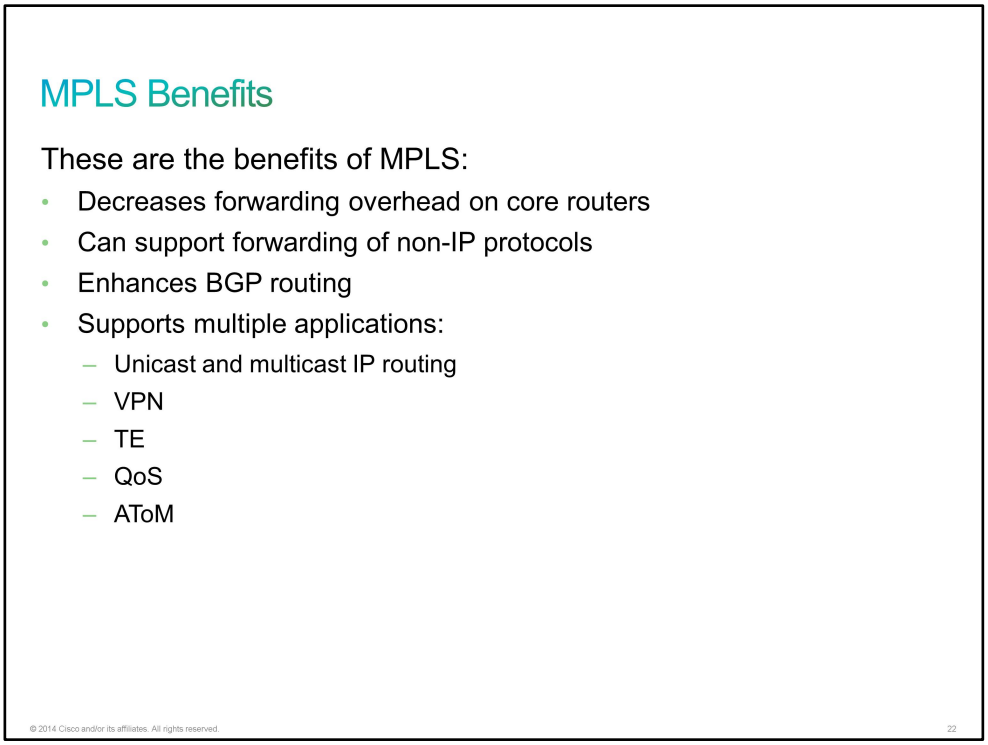

Today, ISPs struggle to scale existing backbone infrastructures for the future and deliver differentiated network services to save costs and generate new revenue streams. ISPs also want to be able to charge premium rates that many customers will pay for special capabilities or levels of service.

ISPs can gain competitive advantages from the MPLS features. MPLS lets ISPs:

- Offer differentiated network services, such as QoS, and to subsequently develop and offer a price model for services.
- Offer MPLS applications to the customers, as MPLS VPNs, MPLS TE and AToM.
- Scale existing network infrastructures to meet future growth requirements

MPLS is very useful in service provider (and large enterprise) networks because it enhances BGP routing and provides different services and applications, such as Layer 2 and 3 VPNs, QoS, and TE.

https://t.me/learningnets

MPLS Terminology

In an MPLS domain, there are three different types of routers:

- **Ingress LSR:** The ingress LSR is at the edge of an MPLS network and is the first to insert an MPLS header and label on a packet.
- **Egress LSR:** The egress LSR is at the edge of the network. It is the last point before leaving the network and thus removes all the MPLS labels and header.
- **Intermediate LSR:** Those are the LSRs which exist within the network and are responsible for pushing, popping, and swapping labels based on the routing with the MPLS network.

Both the ingress and egress LSRs are considered PE routers or edge LSRs. The intermediate LSRs are considered provider or P routers.

Note	Ingress LSR is also called an imposing LSR as it pushes a label onto a packet that was not labeled yet. Egress LSR is also called an disposing LSR as it removes all labels from the labeled packet before getting the packet out of the MPLS domain.

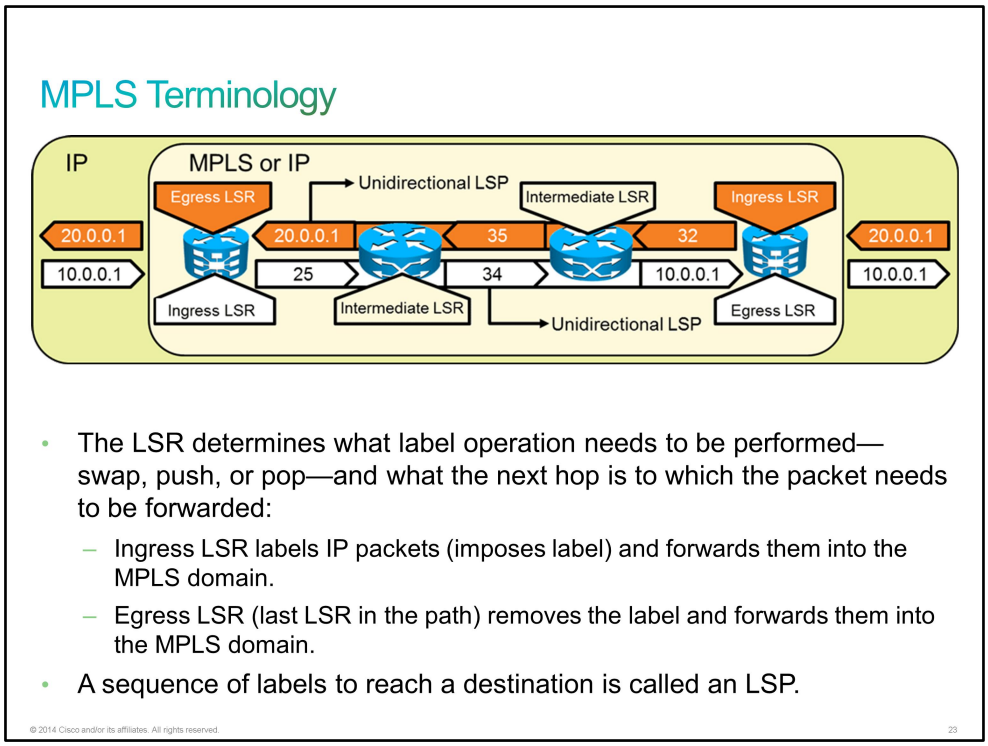

LSRs and edge LSRs are usually capable of doing both label switching and IP routing. Their names are based on their positions in an MPLS domain. Routers that have all interfaces enabled for MPLS are called LSRs because they mostly forward labeled packets (except for the penultimate LSR). Routers that have some interfaces that are not enabled for MPLS are usually at the edge of an MPLS domain. Ingress edge LSR forwards packets that are based on IP destination addresses, and labels them if the outgoing interface is enabled for MPLS. Egress LSR forwards IP packets that are based on routing lookup outside the MPLS domain.

https://t.me/learningnets

A sequence of labels that is used to reach a destination is called a LSP. LSPs are unidirectional; that means that the return traffic uses a different LSP. A penultimate LSR router in an LSP path removes a label and forwards the IP packet to the egress edge LSR router, which routes the IP packet, based on routing lookup. The removing of a label on the penultimate LSR is called PHP.

For example, an edge LSR receives a packet for destination 10.0.0.1, imposes label 25, and forwards the frame to the LSR in the MPLS backbone. The first LSR swaps label 25 for label 34, and forwards the frame. The second (penultimate) LSR removes the label and forwards the IP packet to the edge LSR. The edge LSR forwards the packet, based on IP destination address 10.0.0.1.

Note	PHP is implemented because of the increased performance on the egress edge LSR. Without PHP, the edge LSR would receive a labeled packet, where two lookups would be needed. The first one would be based on labels, and the result would be to remove the label. The second lookup would be to route the IP packet based on destination IP address and routing table.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS Architecture: Control Plane

The control plane builds a routing table (RIB) that is based on the routing protocol. Various routing protocols, such as OSPF, IGRP, EIGRP, IS-IS, RIP, and BGP can be used in the control plane for managing Layer 3 routing.

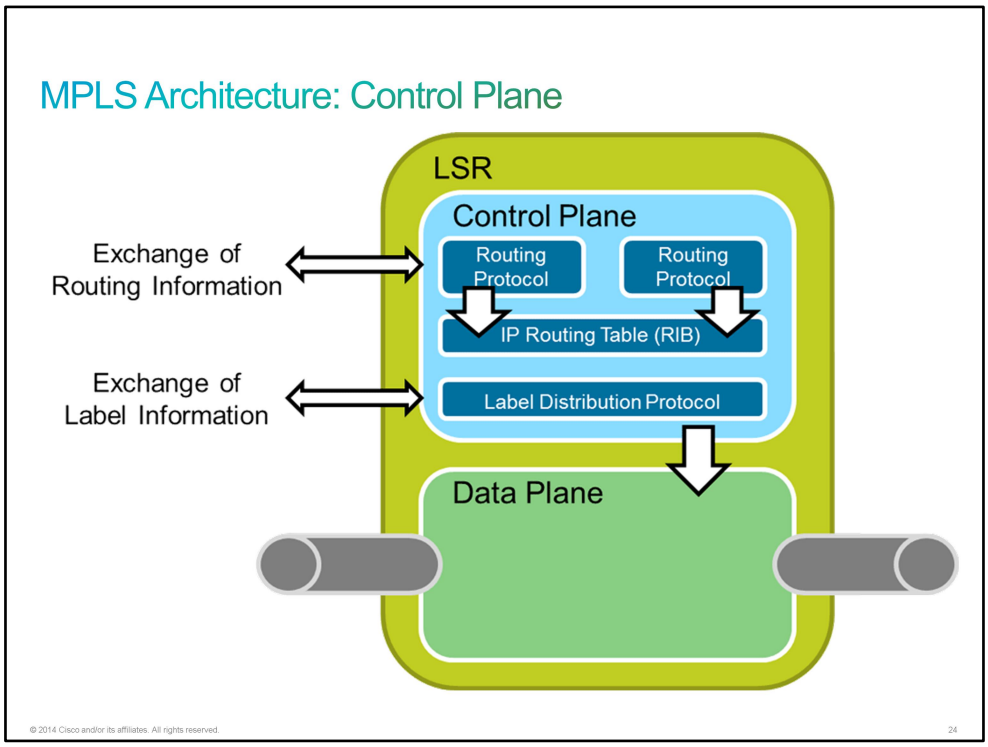

The control plane uses a label exchange protocol to create and maintain labels internally, and to exchange these labels with other MPLS-enabled devices. The label exchange protocol binds labels to networks that are learned via a routing protocol. Label exchange protocols include MPLS LDP, the older Cisco Tag Distribution Protocol, and BGP (used by MPLS VPN). RSVP is used by MPLS TE to accomplish label exchange.

The control plane also builds two forwarding tables, a FIB from the information in the RIB, and a label LFIB table, based on the label exchange protocol and the RIB. The LFIB table includes label values and associations with the outgoing interface for every network prefix.

https://t.me/learningnets

MPLS Architecture: Data Plane

The data plane takes care of forwarding, based on either destination addresses or labels; the data plane is also known as the forwarding plane.

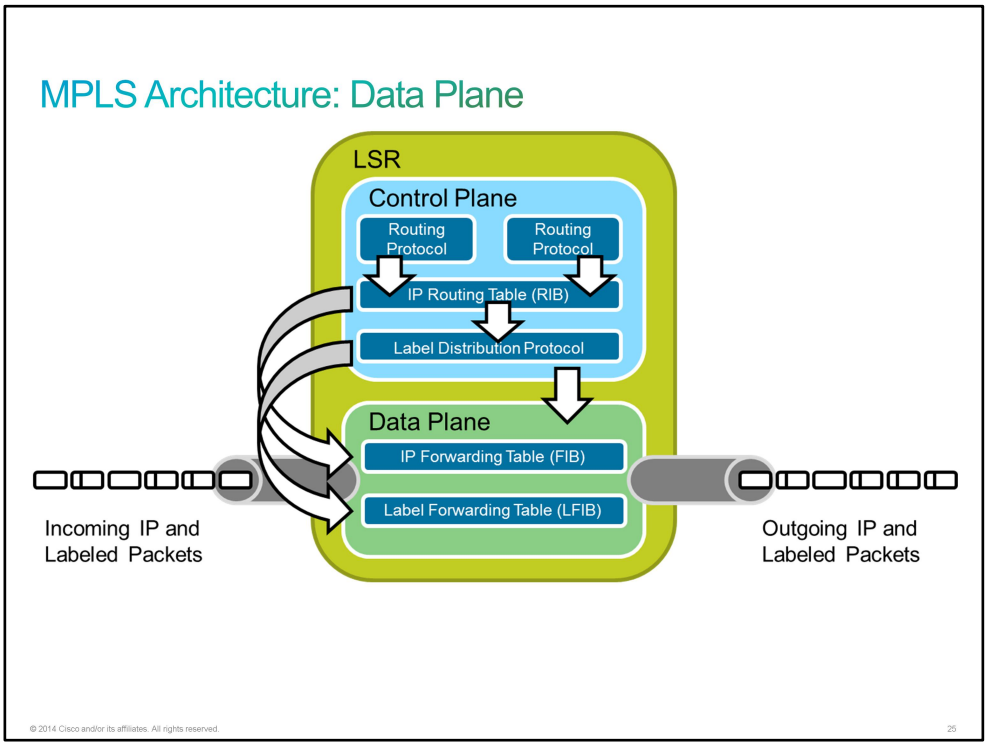

The data plane is a simple forwarding engine that is independent of the type of routing protocol or label exchange protocol being used. The data plane forwards packets to the appropriate interface, based on the information in the LFIB or the FIB tables.

https://t.me/learningnets

Forwarding Structures

The data plane on a router is responsible for forwarding packets, based on decisions done by routing protocols (which run in the router control plane). The data plane on an MPLS-enabled router consists of two forwarding structures: FIB and LFIB.

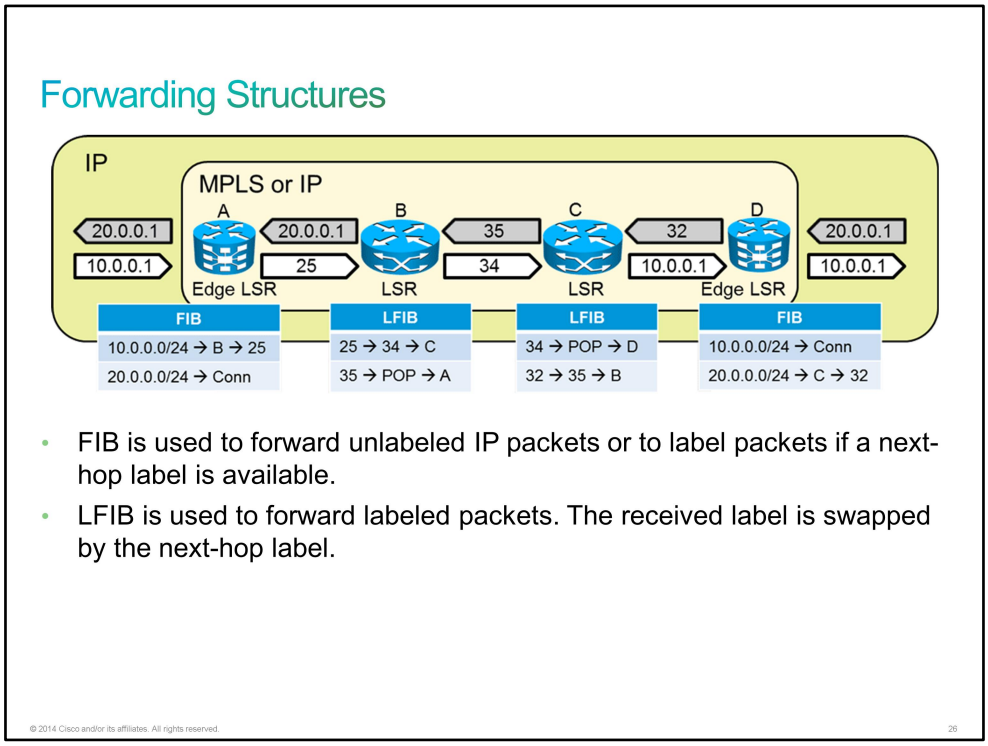

Both FIB and LFIB hold the necessary label information, outgoing interface, and next-hop information:

- **FIB:** When a router is enabled for Cisco Express Forwarding, the FIB is used to forward IP packets, based on decisions made by routing protocols. The FIB is populated from a routing table and includes destination networks, next hops, outgoing interfaces, and pointers to Layer 2 addresses. The FIB on an MPLS-enabled router also contains an outgoing label, if an outgoing interface is enabled for MPLS. FIB lookup is done when an IP packet is received. Based on the result, the router can send out an IP packet or a label can be imposed.
- **LFIB:** The LFIB is used when a labeled packet is received. The LFIB contains the incoming and outgoing label, outgoing interface, and next-hop router information. When an LFIB lookup is done, the result can be to swap a label and send a labeled packet or to remove a label and send an IP packet.

https://t.me/learningnets

These combinations of forwarding packets are possible:

- A received IP packet (FIB) is forwarded, based on the IP destination address, and is sent as an IP packet.
- A received IP packet (FIB) is forwarded, based on the IP destination address, and is sent as a labeled packet.
- A received labeled packet (LFIB) is forwarded, based on the label; the label is changed (swapped) and the labeled packet is sent.
- A received labeled packet (LFIB) is forwarded, based on the label; the label is removed and the IP packet is sent.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS Architecture Example

This example shows the protocols used in the MPLS control plane and the LFIB in the data plane.

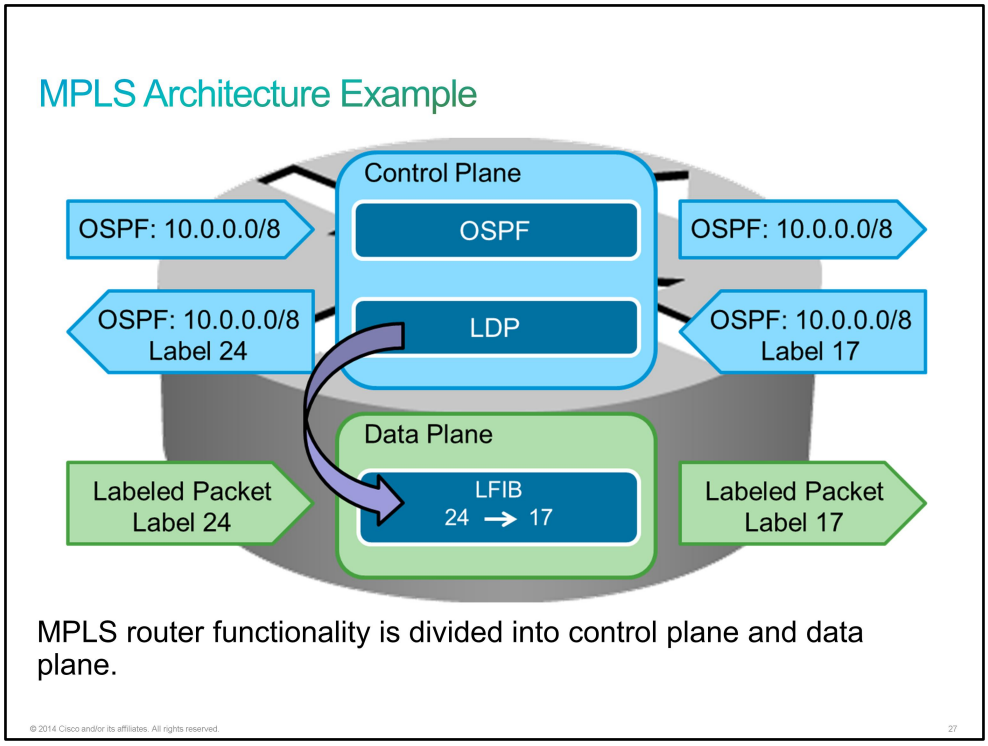

In the example LSR architecture, the control plane uses these protocols:

- A routing protocol (OSPF), which receives and forwards information about IP network 10.0.0.0/8
- A label exchange protocol (LDP), which receives label 24 to be used for packets with destination address 10.0.0.0/8

(A local label 24 is generated and is sent to upstream neighbors so that these neighbors can label packets with the appropriate label.)

The data plane uses an LFIB to forward packets based on labels:

- The LFIB receives an entry from LDP, where label 24 is mapped to label 17. When the data plane receives a packet labeled with a 24, it replaces label 24 with label 17 and forwards the packet through the appropriate interfaces.

Note	In the example, both packet flow and routing and label updates are from left to right.

https://t.me/learningnets

MPLS Labels

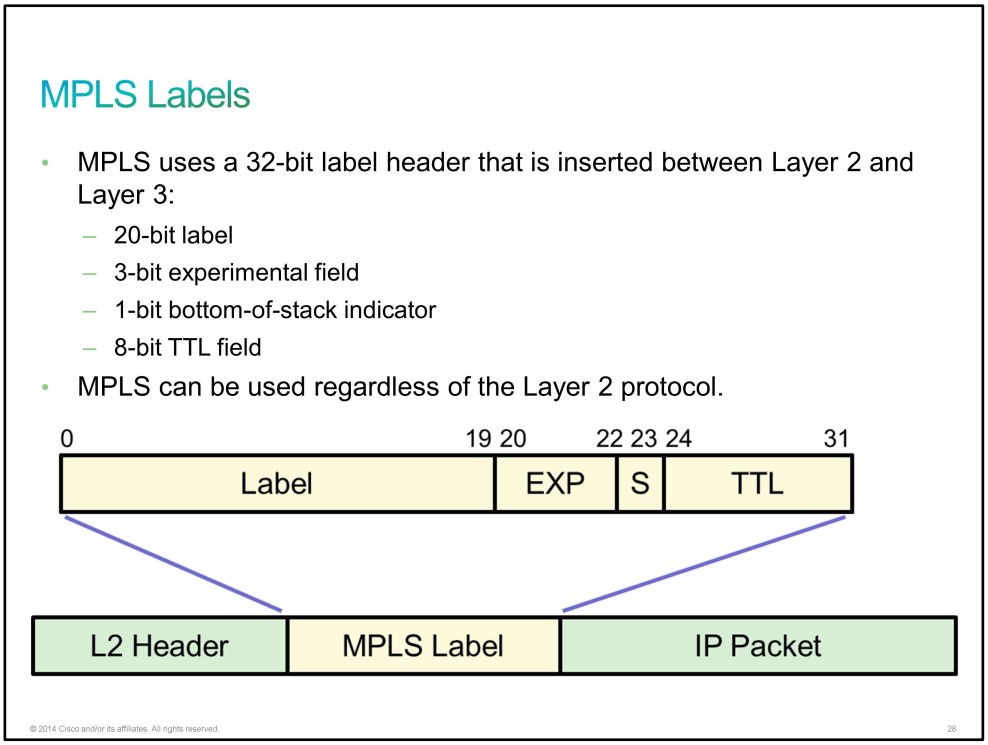

The figure presents an MPLS label that is used for MPLS switching. This label is inserted between the Layer 2 and Layer 3 header and can be used regardless of which Layer 2 protocol is used.

The label is 32 bits long and consists of the following fields:

Field	Description
20-bit label	The actual label used for switching. Values 0 to 15 are reserved.
3-bit EXP field	Undefined in the RFC. Used by Cisco to define a CoS (IP precedence).
Bottom-of-stack bit	MPLS allows multiple labels to be inserted. The bottom-of-stack 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 TTL field	Has the same purpose as the TTL field in the IP header

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

The edge LSR then does these tasks:

- The router performs routing lookup to determine the outgoing interface.
- The router inserts a label between the Layer 2 frame header and the Layer 3 packet header, if the outgoing interface is enabled for MPLS and if a next-hop label for the destination exists. This inserted label is also called the shim header.
- The router then changes the Layer 2 PID or EtherType value in the Layer 2 frame header to indicate that this is a labeled packet. For example, EtherType 0x8847 means a MPLS unicast packet.
- The router sends the labeled packet.

Note	Other routers in the MPLS core simply forward packets based on the received label.

MPLS is designed for use on virtually any media and Layer 2 encapsulation.

https://t.me/learningnets

MPLS Static Feature

The label-switching database provides a platform for applications such as BGPand LDP to allocate local MPLS labels. The allocated labels are stored in the LSD until the owner modifies or deletes the label. These local labels can be allocated from a dynamic pool, and in this case the end user cannot choose a desired label value.

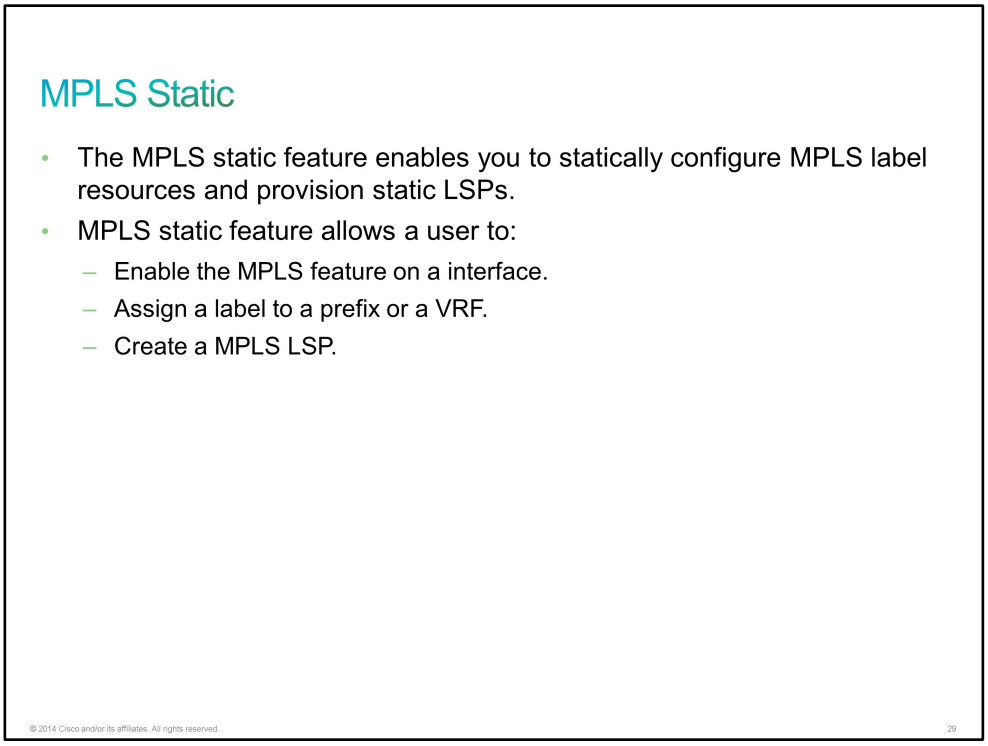

During the development of the LDP-CSC functionality, the MPLS static allocation feature has been proposed as a solution through which an end user could allocate an incoming label statically.

Statically defined labels by an end user will be unique for a prefix for a VRF context, which includes the VRF name and an IPv4 or IPv6 prefix. The incoming labels that a user can configure can be limited to a static range defined by the user.

The MPLS static feature also introduces the ability to create simple label-switched-paths on a router. With this functionality, there is no need for the label signaling application for label distribution. In this case, the user must select a local label and the next-hop information required for packet forwarding.

https://t.me/learningnets

MPLS Packet Flow Basic Example

This example shows how an IP packet is forwarded using MPLS label switching.

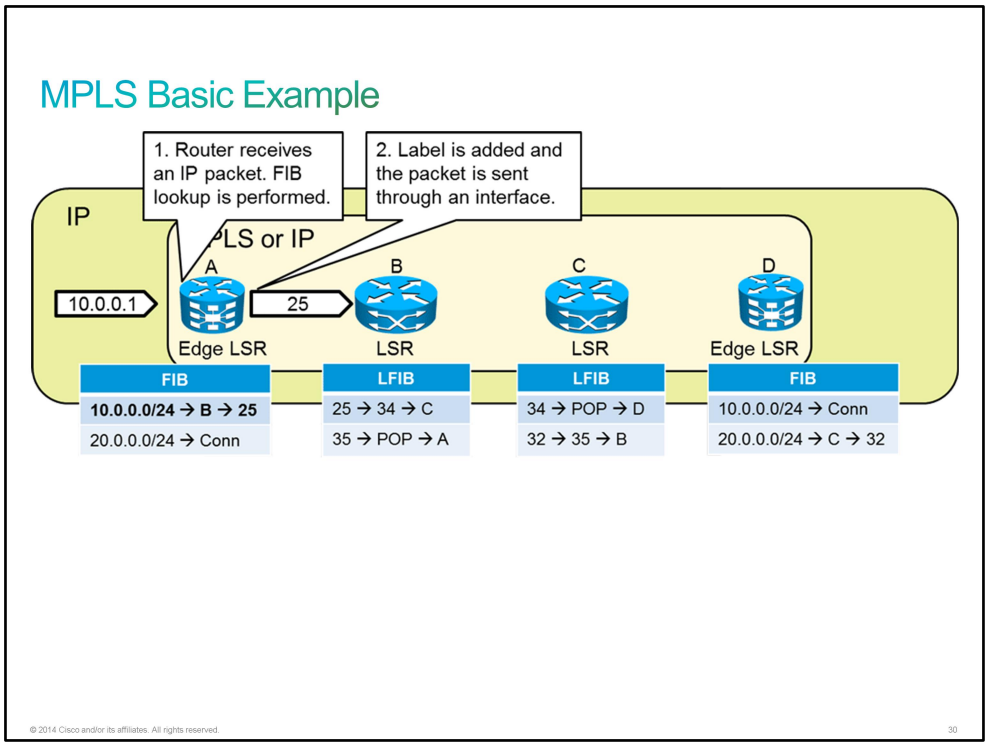

The figure shows an example of the way in which a packet traverses an MPLS-enabled network. Router A receives an IP packet destined for 10.0.0.1.

1. Router A performs a FIB lookup. The FIB for that destination states that the packet should be labeled using label 25 and sent to router B.
2. Router A adds a label 25 and the packet is sent out the interface that connects to router B.
3. Router B receives an IP packet that is labeled with label 25. Router B performs an LFIB lookup, which states that label 25 should be swapped with label 34.
4. The label is swapped and the packet is sent to router C.

https://t.me/learningnets

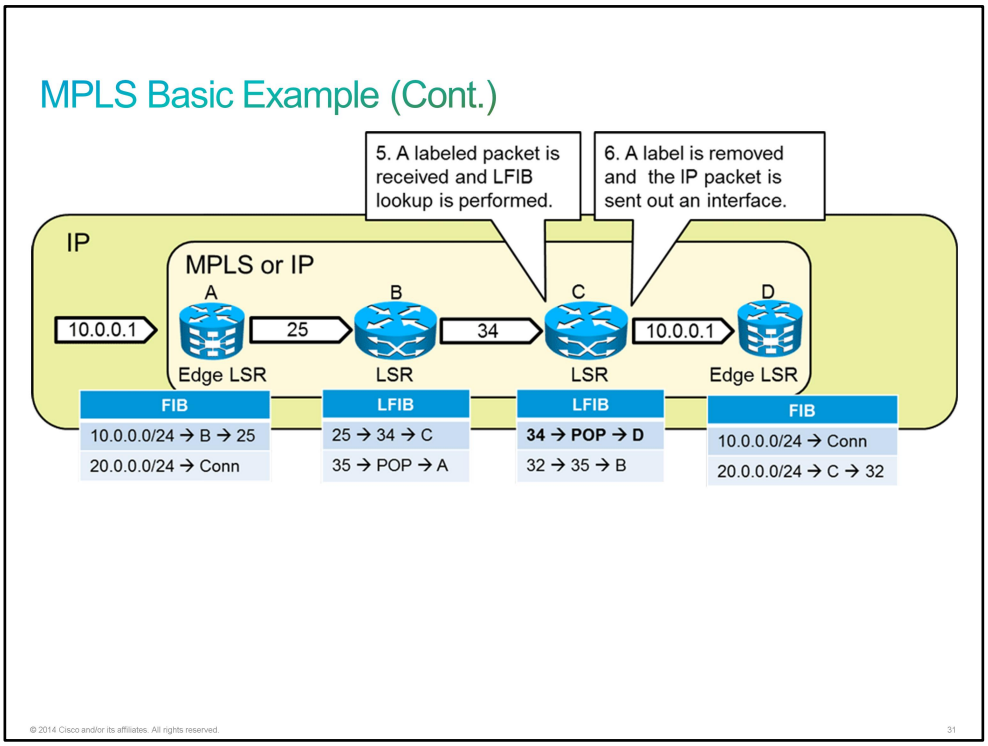

5. Router C receives an IP packet that is labeled with label 34. Router C performs an LFIB lookup, which states that label 34 should be removed (PHP), and the unlabeled IP packet should be sent out the interface that connects to router D. “POP” is often used as a label value that indicates that a label should be removed.
6. The label is removed and the unlabeled IP packet is sent out the interface that connects to router D.

A router will actually display a value of IMP-NULL instead of POP. An implicit null label means that the label should be removed. An IMP-NULL label uses the value 3 from a reserved range of labels.

https://t.me/learningnets

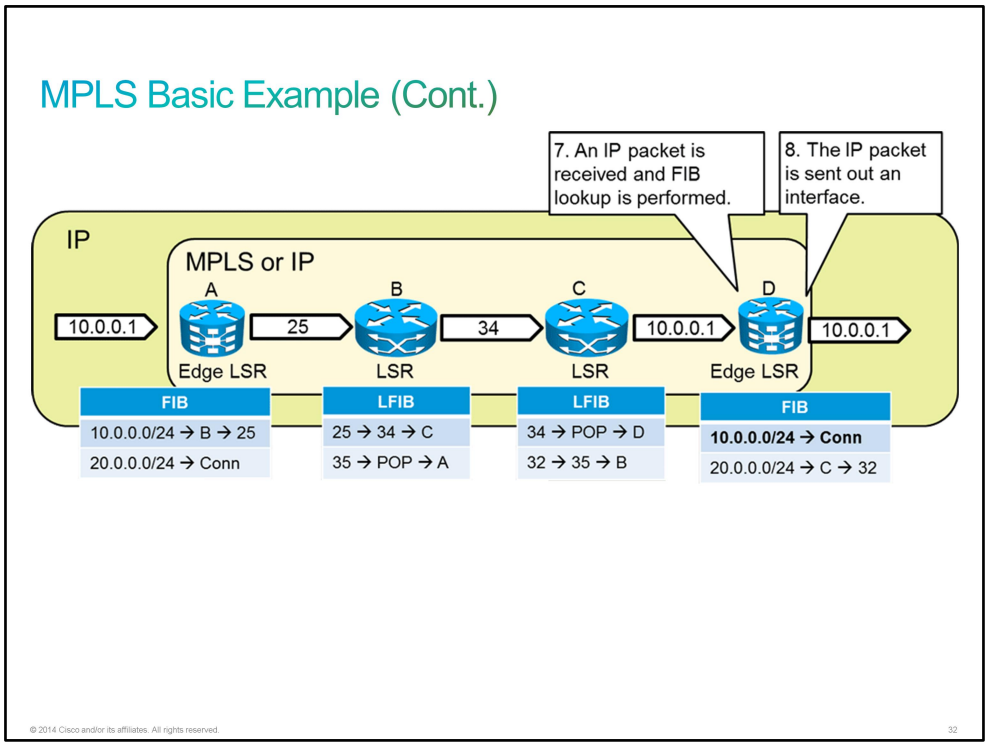

7. Finally, router D receives an IP packet. Router D performs a FIB lookup, which states that the destination network is directly connected.
8. The IP packet is sent out the directly connected interface.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- MPLS enables the forwarding of packets through label-switched paths that can be created using various methods and protocols, depending on the required results.
- MPLS switches packets based on label lookup instead of IP address lookup. Labels usually correspond to destination IP networks.
- MPLS is very useful in service provider (and large enterprise) networks because it enhances BGP routing and provides different services and applications.
- In an MPLS domain, there are two types of routers: LSRs and edge LSRs.
- The control plane builds a routing table that is based on the routing protocol.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- The data plane takes care of forwarding, based on either destination addresses or labels.
- The data plane on an MPLS-enabled router consists of two forwarding structures: FIB and LFIB. FIB is used to forward unlabeled IP packets or to label packets if a next-hop label is available. LFIB is used to forward labeled packets. The received label is swapped by the next-hop label.
- The control plane on an MPLS-enabled router usually uses link-state routing protocol to exchange IP prefixes and LDP to exchange MPLS labels.
- An MPLS label is a 4-byte identifier that is used for making forwarding decisions.
- The MPLS static feature allows the user to allocate incoming MPLS labels statically.
- If an MPLS-enabled router receives a labeled packet, the router performs LFIB lookup.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 3: Introducing MPLS Applications

Overview

It is important to have a clear understanding of the role of MPLS and the makeup of the devices and components. This lesson describes, at a high level, some of the various types of applications with which you can use MPLS.

This understanding will help you have a clear picture of how to differentiate between the roles of certain devices and to understand how information is transferred across an MPLS domain.

Upon completing this lesson, you will be able to:

- Describe the MPLS label stack
- List the MPLS applications in a service provider environment
- Describe MPLS support for unicast IP routing
- Describe MPLS support for multicast IP routing
- Describe MPLS support for VPNs
- Describe MPLS support for Layer 3 VPNs
- Describe MPLS support for Layer 2 VPNs
- Describe MPLS support for traffic engineering
- Describe MPLS support for QoS
- Describe the interaction between MPLS applications

https://t.me/learningnets

MPLS Applications

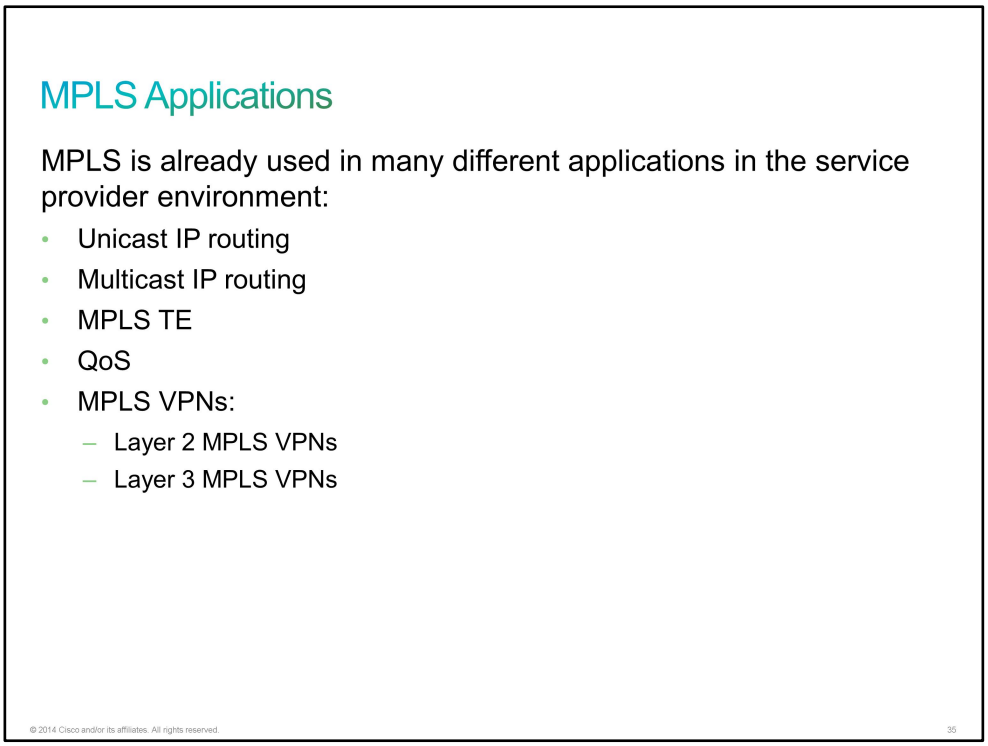

MPLS is a technology that is used for the delivery of IP services. MPLS can be used in different applications, as outlined here:

- Unicast IP routing is the most common application for MPLS.
- Multicast IP routing is treated separately because of different forwarding requirements.
- MPLS TE is an add-on to MPLS that provides better and more intelligent link usages.
- Differentiated QoS can also be provided with MPLS.
- MPLS VPNs are implemented, using labels to allow overlapping address space between VPNs.
- AToM is a solution for transporting Layer 2 packets over an IP or MPLS backbone.

MPLS support for a label stack allows the implementation of enhanced applications, such as VPNs, TE, and enhanced QoS.

https://t.me/learningnets

MPLS Label Stack

MPLS routers encapsulate the packets with special MPLS headers containing one or more labels. Label stacking is the encapsulation of an MPLS packet inside another MPLS packet. The result of stacking is the ability to tunnel one MPLS LSP inside another LSP.

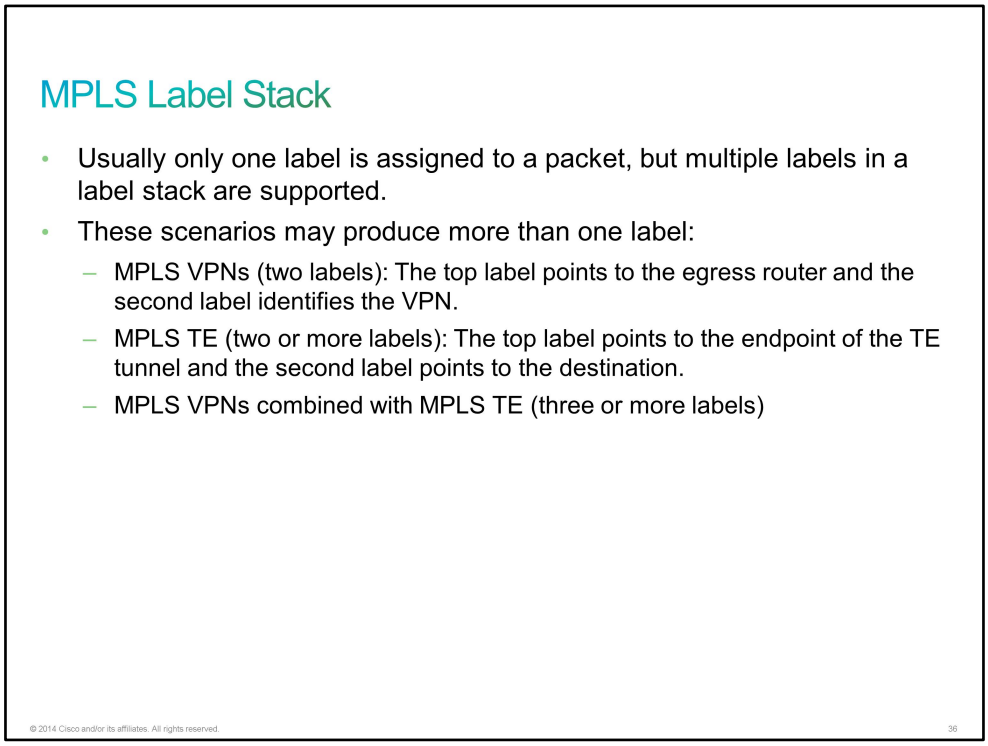

Simple MPLS uses just one label in each packet. However, MPLS does allow multiple labels in a label stack to be inserted in a packet.

These applications may add labels to packets:

- **MPLS VPNs:** With MPLS VPNs, MP-BGP is used to propagate a second label that is used in addition to the one propagated by LDP.
- **Cisco MPLS Traffic Engineering:** MPLS TE uses RSVP to establish LSP tunnels. RSVP also propagates labels that are used in addition to the one propagated by LDP.

A combination of these mechanisms and other advanced features might result in three or more labels being inserted into one packet.

https://t.me/learningnets

MPLS Label Stack (Cont.)

Example:

- The outer label is used for switching the packet in the MPLS network (points to the TE destination).
- Inner labels are used to separate packets at egress points (point to an egress router and identify a VPN).

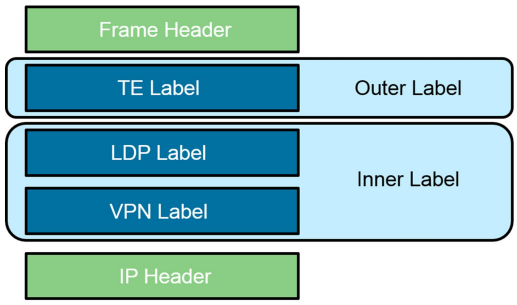

© 2014 Cisco and/or its affiliates. All rights reserved.

The figure shows an example of a label stack where both MPLS TE and MPLS VPN are enabled.

The outer label is used to switch the MPLS packet across the network. In this case, the outer layer is a TE label pointing to the endpoint of a TE tunnel.

The inner labels are ignored by the intermediary routers. In this case, the inner labels are used to point to the egress router and to identify the VPN for the packet.

https://t.me/learningnets

MPLS Unicast IP Routing

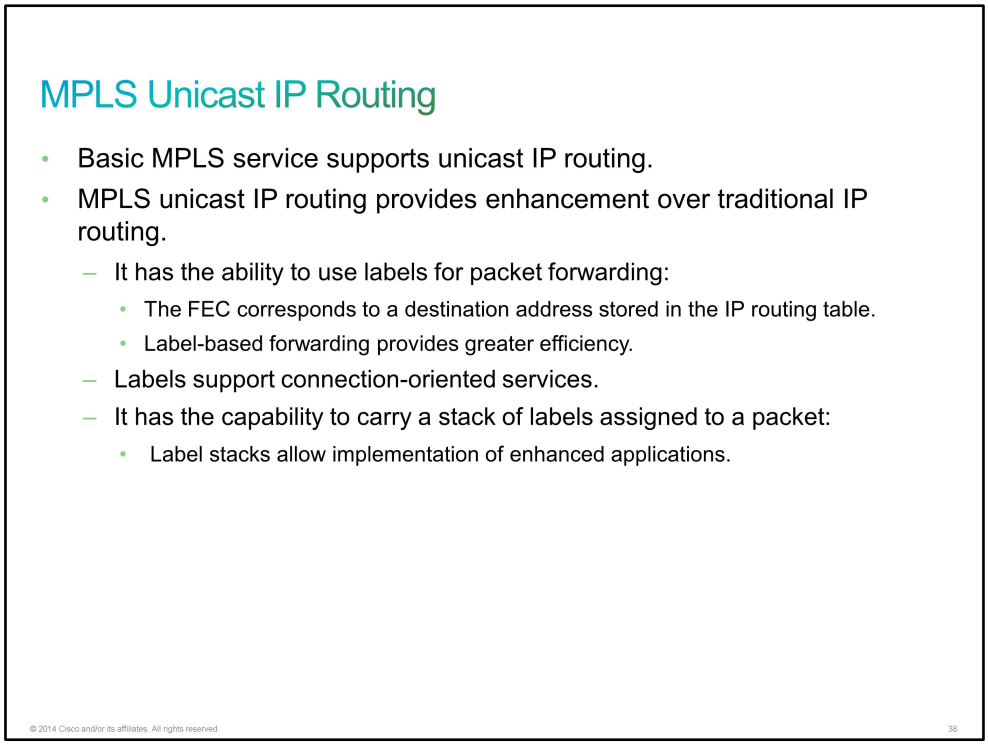

Using labels for packet forwarding increases efficiency in network core devices because the label swapping operation is less CPU-intensive than a routing lookup. MPLS can also provide connection-oriented services to IP traffic due to FEC-based forwarding.

Note	The MPLS unicast IP traffic FEC corresponds to a destination network stored in the IP routing table.

https://t.me/learningnets

MPLS Multicast IP Routing

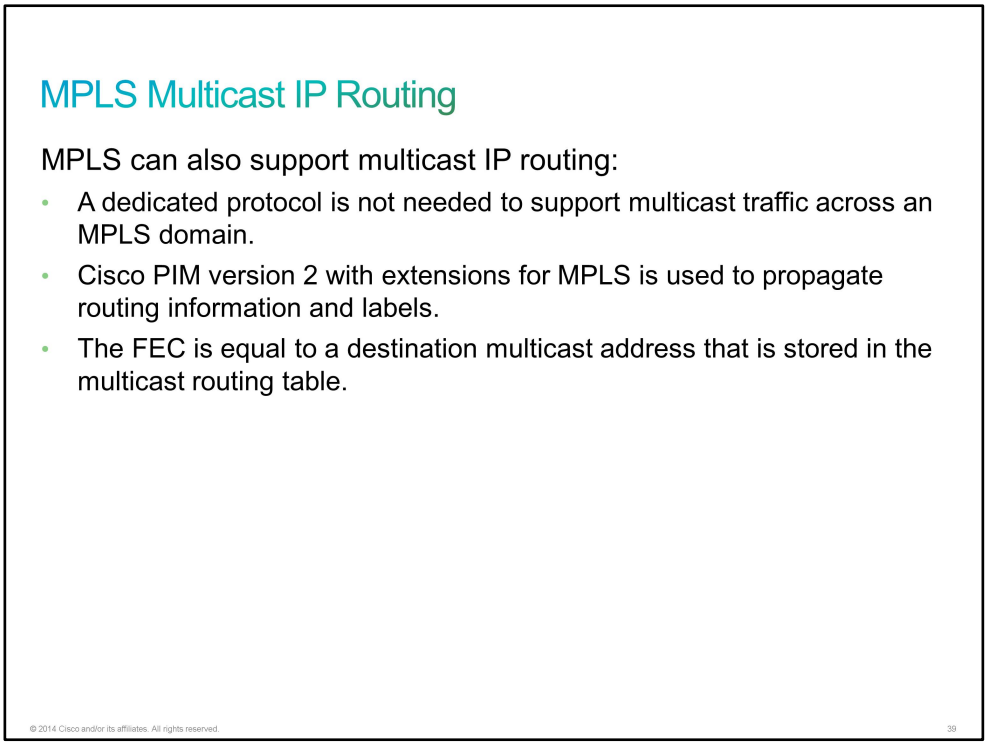

Multicast IP routing can also use MPLS. Cisco PIM version 2 with extensions for MPLS is used to propagate routing information and labels.

The FEC is equal to a destination multicast address.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS VPNs

MPLS enables highly scalable VPN services to be supported.

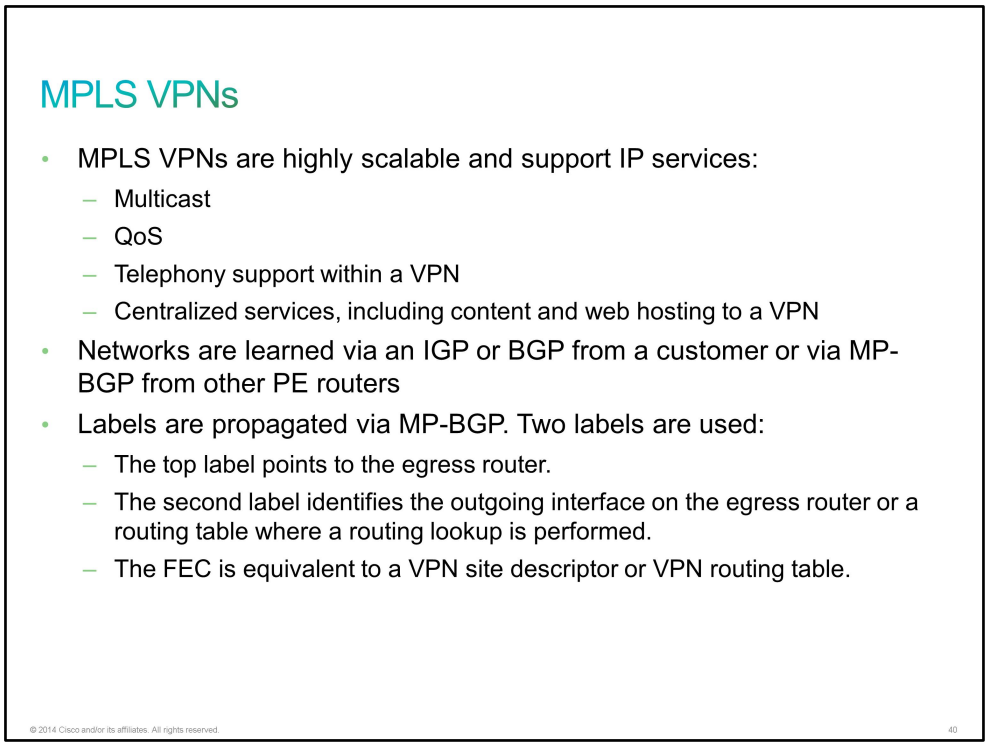

For each MPLS VPN user, the network appears to function as a private IP backbone, over which the user can reach other sites within the VPN organization, but not the sites of any other VPN organization. MPLS VPNs are a common application for service providers. Building VPNs in Layer 3 allows the delivery of targeted services to a group of users represented by a VPN.

MPLS VPNs are seen as private intranets, and they support IP services such as those listed here:

- Multicast
- QoS
- Telephony support within a VPN
- Centralized services including content and web hosting to a VPN

Customer networks are learned via an IGP (OSPF, EBGP, EIGRP, or RIPv2), BGP or via static routing from a customer, or via MP-BGP from other PE routers.

MPLS VPNs use two labels:

- The top label points to the egress router.
- The second label identifies the outgoing interface on the egress router or a routing table where a routing lookup is performed.

https://t.me/learningnets

LDP is needed in the top label to link edge LSRs with a single LSP tunnel. MP-BGP is used in the second label to propagate VPN routing information and labels across the MPLS domain. The MPLS VPN FEC is equivalent to a VPN site descriptor or a VPN routing table.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Layer 3 MPLS VPNs

Layer 3 MPLS VPN services allow businesses to outsource their current network core using a private IP-based service offering from a service provider. MPLS VPNs require that the enterprise peer with the service provider at the IP Layer 3 level.

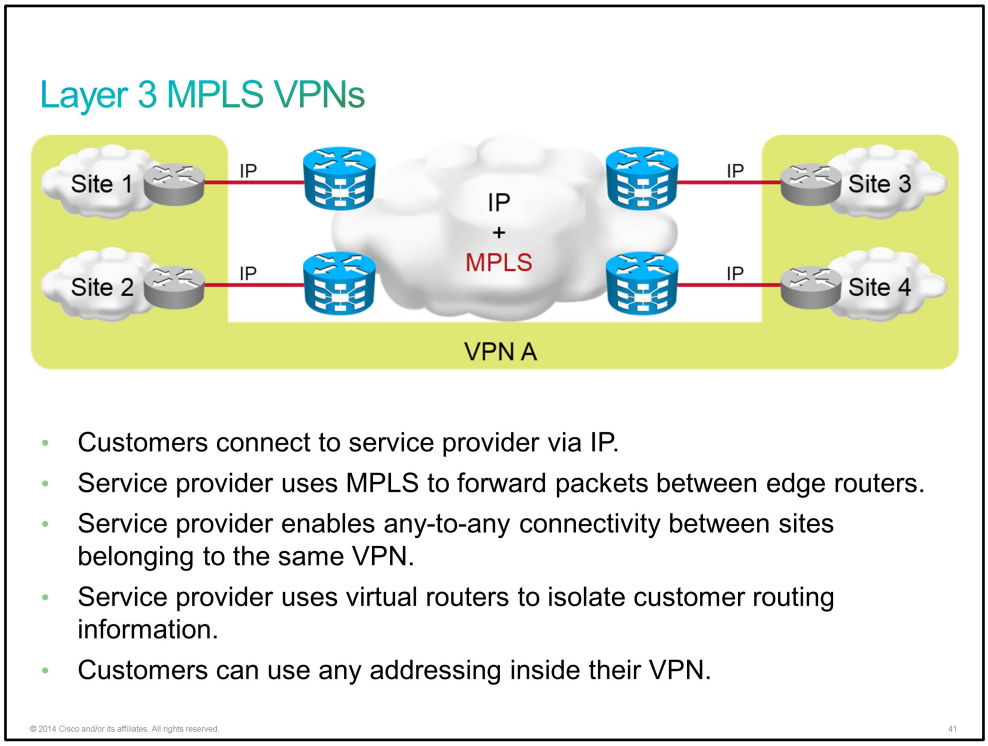

The main characteristic of Layer 3 MPLS VPNs is that customers connect to a service provider via IP. They need to establish IP routing (static or dynamic) in order to exchange routing information between customer sites belonging to the same VPN. As different customers may use the same private IP address ranges, the service provider cannot perform normal IP forwarding. MPLS must be used instead, to ensure isolation in the data plane between packets belonging to different customers, yet potentially having the same IP addresses. Virtual routers (VRF instances) are used on service provider routers to isolate customer routing information. MPLS seamlessly provides any-to-any connectivity between sites belonging to the same VPN.

The most basic VPN is a so-called simple VPN or an intranet. This type of VPN is a collection of sites that are given full connectivity within the VPN while isolating the VPN from any other component in the network (other VPNs, Internet).

https://t.me/learningnets

Layer 3 MPLS VPNs (Cont.)

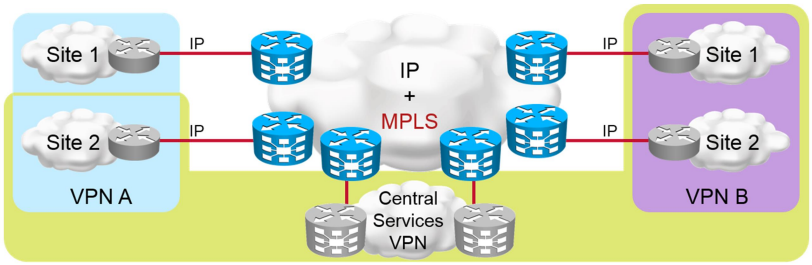

- Service provider enables any-to-any connectivity between VPN sites.
- All or selected sites have access to the central VPN.
- Customers can use any addressing inside their VPNs.
- Customers must use nonoverlapping addresses to access the central VPN.

© 2014 Cisco and/or its affiliates. All rights reserved.

The figure illustrates overlapping Layer 3 MPLS VPNs where multiple customer VPNs are provided access to a central service VPN. Both VPN A (only site 2) and VPN B (both sites) in the example are able to communicate with the central services VPN. VPN A and VPN B are still isolated from each other.

The only requirement in the case of overlapping VPNs is that the VPNs use unique addressing, at least when accessing the resources available in other VPNs.

https://t.me/learningnets

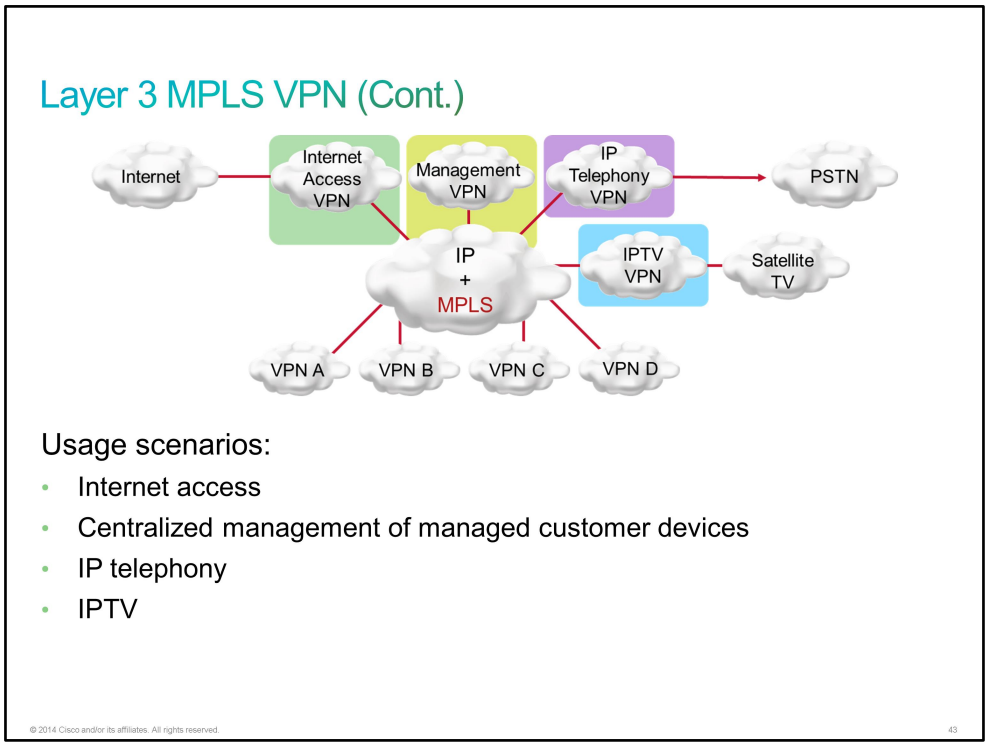

The figure illustrates a few examples of central services VPNs:

- Internet access can be provided through a dedicated Layer 3 MPLS VPN. Service providers can also offer wholesale services to other service providers, where they can choose their upstream Internet service provider.
- A management VPN is often used to manage the network infrastructure and services. This management VPN can also be used to manage customer routers inside Layer 3 MPLS VPNs in case the customer devices are owned and managed by the service provider.
- IP telephony can be isolated and provided to customers through a dedicated VPN.
- Even IPTV can now be isolated and provided through a dedicated VPN.

https://t.me/learningnets

Layer 2 MPLS VPNs

Layer 2 MPLS VPNs enable service providers to offer point-to-point or multipoint Layer 2 connections between distant customer sites. MPLS Layer 2 VPNs consolidate Layer 2 traffic such as Ethernet, Frame Relay, ATM, HDLC, and PPP over an IP or MPLS network. This support, especially for traditional Frame Relay and ATM networks, enables service providers to protect their investments by continuing to provide existing data and voice services while introducing new services and architectures. Service providers also benefit from cost savings in converging services and new service revenues from new, innovative IP services over the newly converged IP or MPLS network.

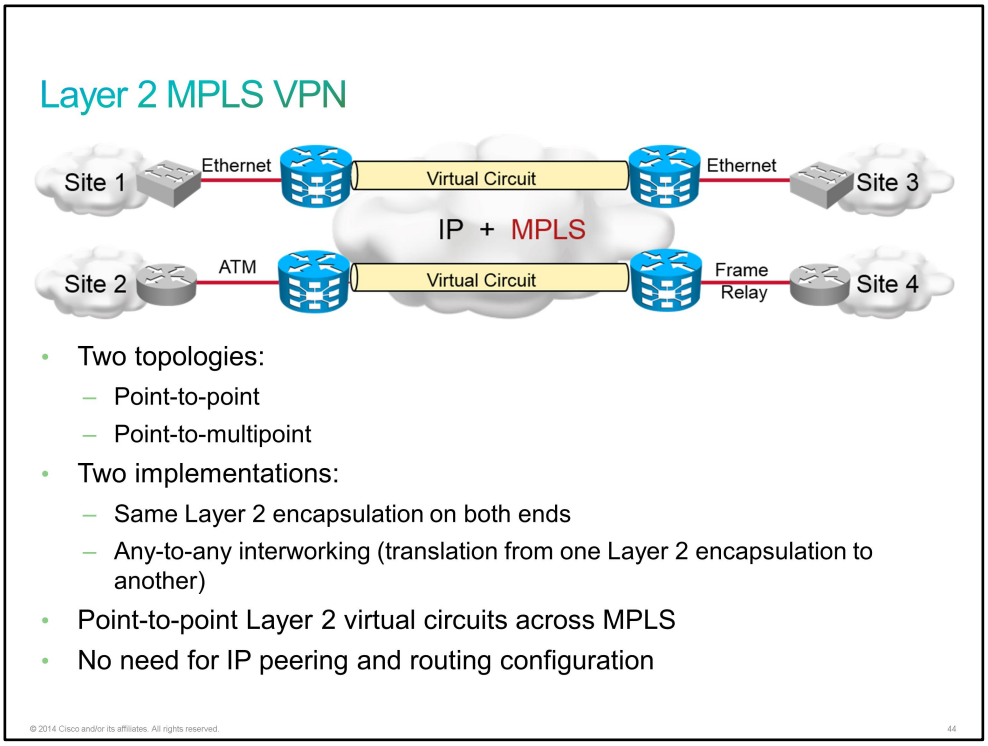

The top figure illustrates a point-to-point Ethernet connection between a pair of customer LAN switches across a virtual circuit that is implemented using MPLS. The other example illustrates interworking where one customer site uses ATM and the other site uses Frame Relay. The MPLS network translates between the two technologies similarly to what most ATM switches were able to do.

The main advantage of Layer 2 MPLS VPNs is that they do not require any IP signaling between the customer and the provider.

EoMPLS can be implemented in two ways:

- Point-to-point Ethernet over MPLS, where all Ethernet traffic is exchanged over a single virtual circuit (LSP)
- VPLS, where multiple sites can be interconnected over a full mesh of virtual circuits

https://t.me/learningnets

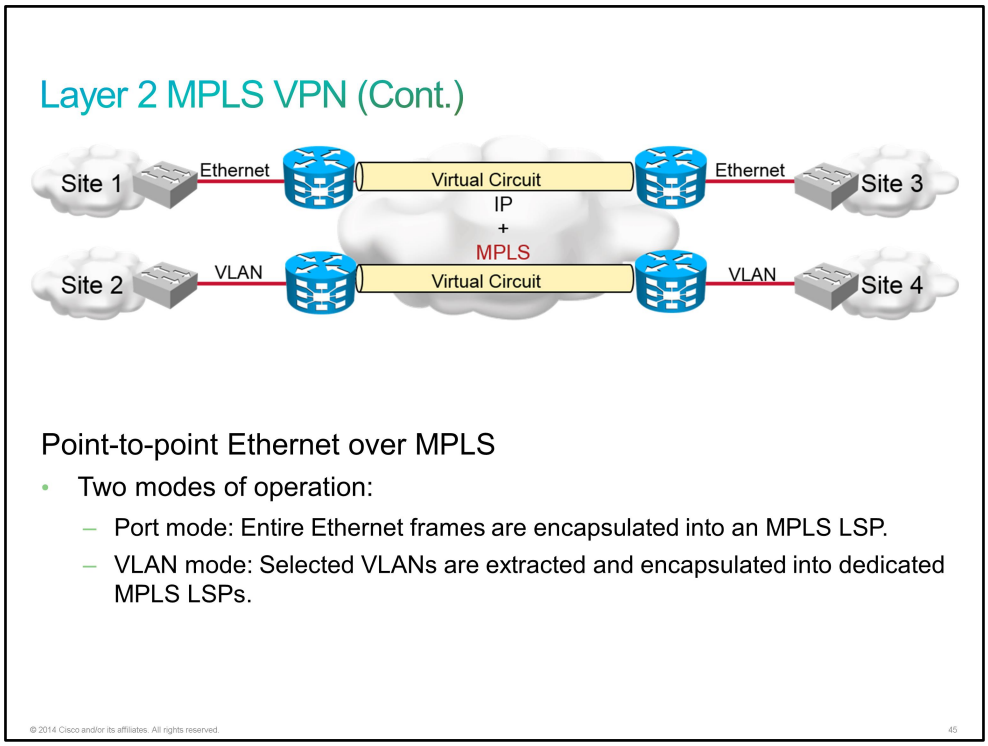

The most common application of Layer 2 MPLS VPNs is to provide point-to-point Ethernet connectivity between customer sites.

EoMPLS can be implemented in two ways:

- **Port mode**: Entire Ethernet frames are encapsulated into an MPLS LSP. This option allows one physical interface to be routed to a single distant remote site, but it can use IEEE 802.1Q VLANs end-to-end.
- **VLAN mode**: Selected VLANs are extracted and encapsulated into dedicated MPLS LSPs. This option allows a central customer site to use a single physical link with multiple VLANs that are then routed to multiple individual remote sites in different locations.

https://t.me/learningnets

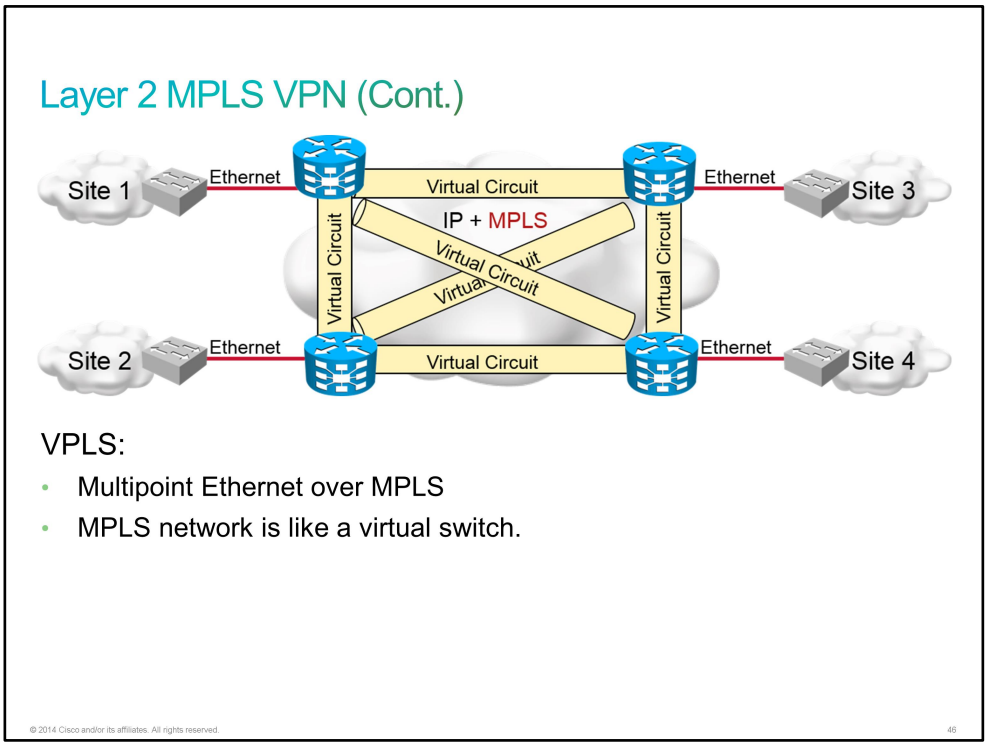

VPLS enables service providers with MPLS networks to offer geographically dispersed EMS, or Ethernet Private LAN Service, as defined by the MEF.

The figure illustrates VPLS implementation between four customer LAN switches in different locations. A full mesh of LSPs ensures optimal forwarding for learned MAC addresses between any pair of sites in the same virtual switch.

Note	Refer to the *Implementing Cisco Service Provider Next-Generation Edge Network Services* course for detailed coverage on Layer 3 and Layer 2 MPLS VPN implementations.

https://t.me/learningnets

MPLS Traffic Engineering

Another application of MPLS is TE. MPLS TE enables an MPLS backbone to replicate and expand upon the TE capabilities of Layer 2 ATM and Frame Relay networks. MPLS TE supports constraint-based routing, in which the path for a traffic flow is the shortest path that meets the resource requirements (constraints) of the traffic flow. Factors such as bandwidth requirements, media requirements, and the priority of one traffic flow versus another can be taken into account. TE capabilities enable the administrator of a network to accomplish these goals:

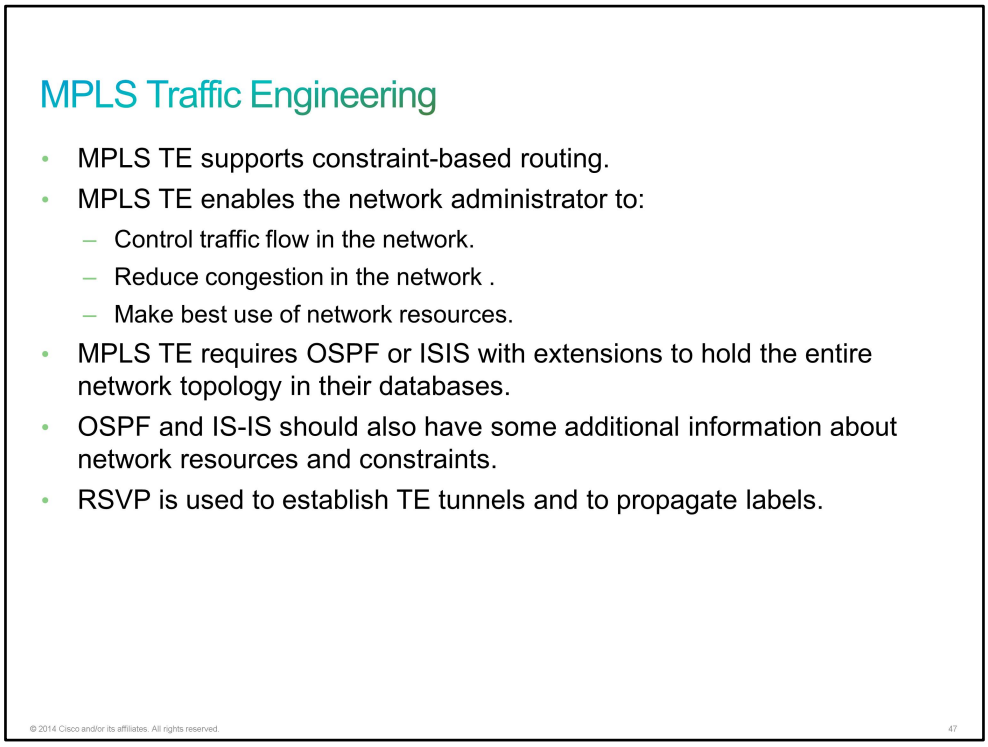

MPLS TE has these special requirements:

- Every LSR must see the entire topology of the network (only OSPF and IS-IS hold the entire topology).
- Every LSR needs additional information about links in the network. This information includes available resources and constraints. OSPF and IS-IS have extensions to propagate this additional information.
- RSVP is used to establish TE tunnels and to propagate the labels.

Every edge LSR must be able to create an LSP tunnel on demand. RSVP is used to create an LSP tunnel and to propagate labels for TE tunnels.

https://t.me/learningnets

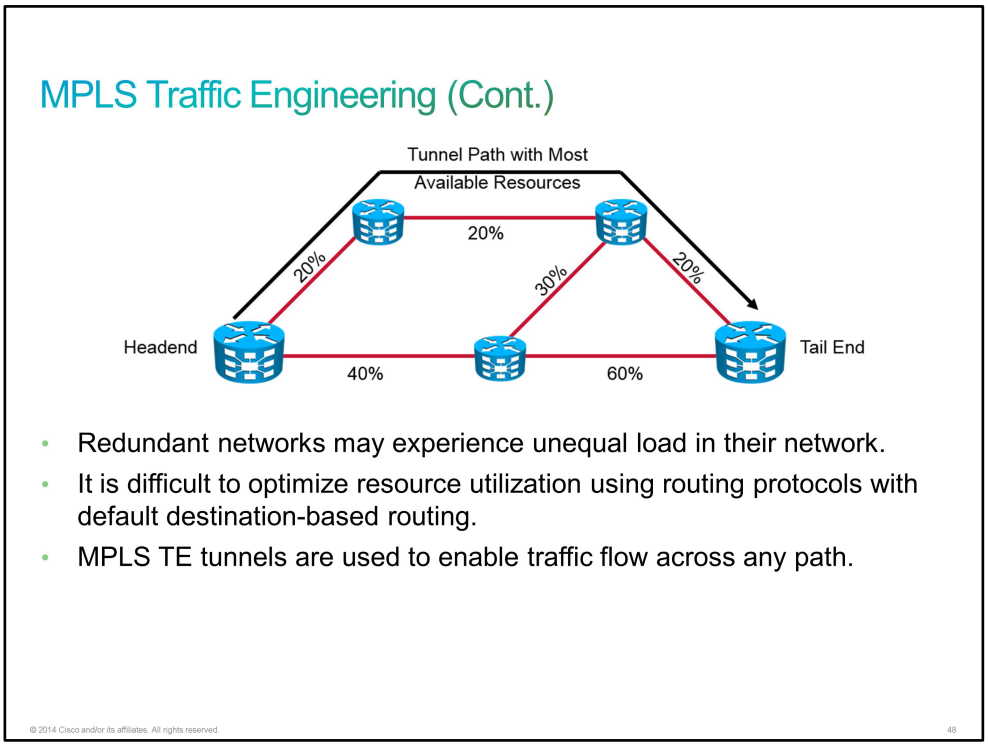

The primary reason for using MPLS TE, as the name suggests, is to engineer traffic paths. Redundant networks may experience unequal loads in their networks, due to the calculated best paths that are typically determined based on IGP metrics. Optimizing resource utilization (link utilization) using routing protocols, with default destination-based routing is difficult. The figure illustrates a scenario where the least-cost path would flow through the most congested link, thus making it even more congested, and possibly resulting in drops and increased delays. MPLS TE can be used to divert some traffic to “less optimal” paths; this capability will result in better utilization of resources (more network throughput) and lower delays.

https://t.me/learningnets

MPLS QoS

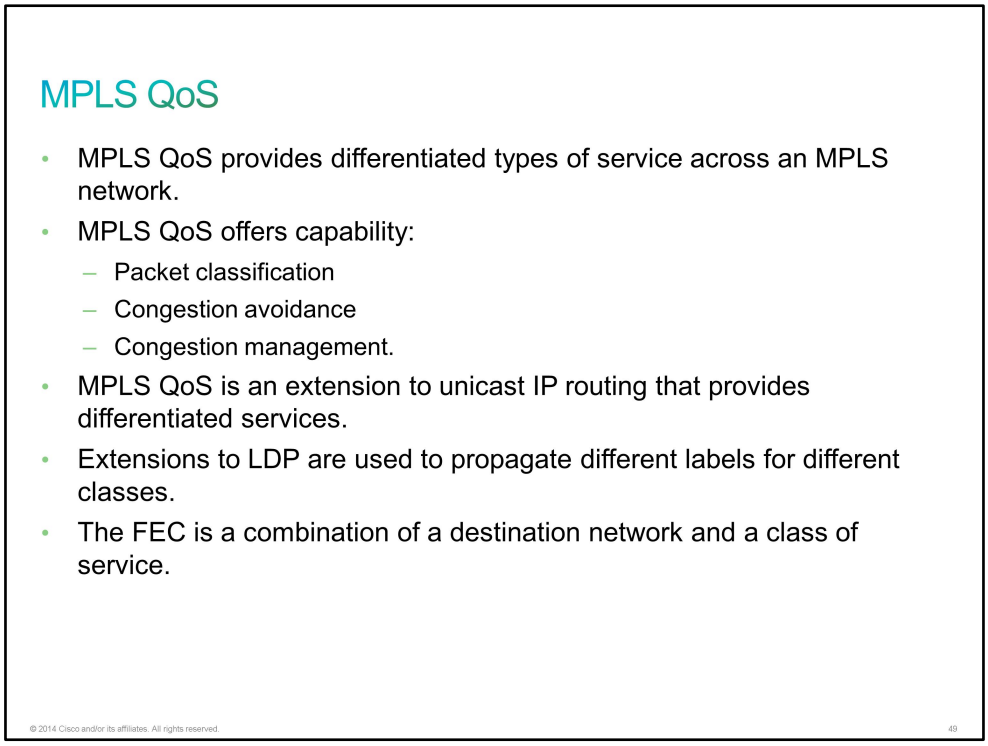

MPLS QoS enables network administrators to provide differentiated types of service across an MPLS network. MPLS QoS offers packet classification, congestion avoidance, and congestion management.

Note	MPLS QoS functions map nearly one-for-one to IP QoS functions on all interface types.

Differentiated QoS is achieved by using MPLS experimental bits or by creating separate LSP tunnels for different classes. Extensions to LDP are used to create multiple LSP tunnels for the same destination (one for each class).

The FEC for MPLS QoS is equal to a combination of a destination network and a CoS.

https://t.me/learningnets

Interaction Between MPLS Applications

The figure shows the overall architecture when multiple MPLS applications are used.

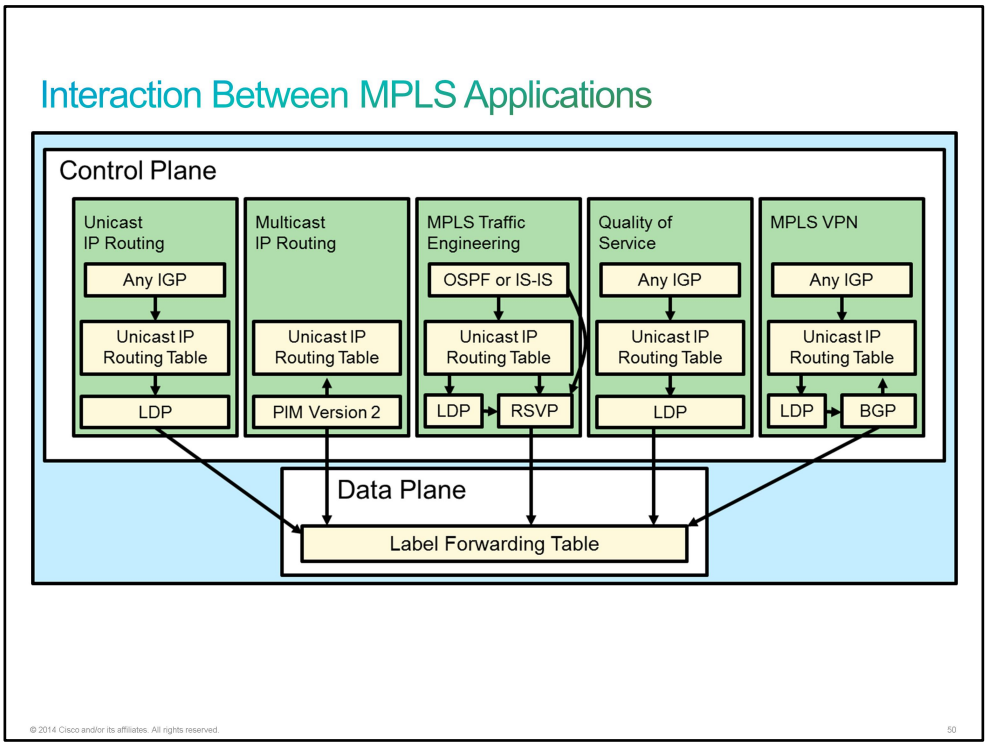

Regardless of the application, the functionality is always split into the control plane and the data (forwarding) plane, as discussed here:

- The applications may use different routing protocols and different label exchange protocols in the control plane.
- The applications all use a common label-switching data (forwarding) plane.
- Edge LSR Layer 3 data planes may differ to support label imposition and disposition.
- Typically, a label is assigned to an FEC.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- MPLS allows multiple labels in a label stack to be inserted in a IP packet.
- MPLS is used in many applications: IP routing, MPLS VPNs, MPLS TE, and QoS.
- One of the significant enhancements of unicast MPLS routing over IP routing is the capability to carry a stack of labels assigned to a packet.
- Multicast IP routing can also use MPLS. Cisco PIM version 2 with extensions for MPLS is used to propagate routing information and labels.
- MPLS supports highly scalable VPN services.
- The main characteristic of Layer 3 MPLS VPNs is that a customer transparently connects his networks through a service provider network via IP.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Layer 2 MPLS VPNs enable service providers to offer point-to-point or multipoint Layer 2 connections between distant customer sites.
- MPLS TE supports constraint-based routing, in which the path for a traffic flow is the shortest path that meets the resource requirements of the traffic flow.
- MPLS QoS provides differentiated types of service across an MPLS network.
- MPLS application may use different protocols at control and data plane.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: Label Distribution Protocol

Overview

This lesson takes a detailed look at the LDP neighbor discovery process via hello messages and the type of information that is exchanged. It also describes the events that occur during the negotiation phase of LDP session establishment, as well as the nonadjacent neighbor discovery process, providing a further understanding of the MPLS technology.

Upon completing this lesson, you will be able to:

- Describe LDP, the protocol used to exchange MPLS labels
- Describe how LDP neighbor adjacency is established
- Describe the LDP link hello message
- Describe the LDP neighbor discovery
- Describe the LDP session negotiation
- Describe the use of the LDP targeted hello message to form LDP neighbor adjacency between nondirectly connected LSRs
- Describe LDP Session Protection using a backup targeted hello
- Describe LDP graceful restart and NSR

https://t.me/learningnets

Label Distribution Protocol

Standard or vendor-specific routing protocols are used to advertise IP routing information. MPLS adds a new piece of information that must be exchanged between adjacent routers.

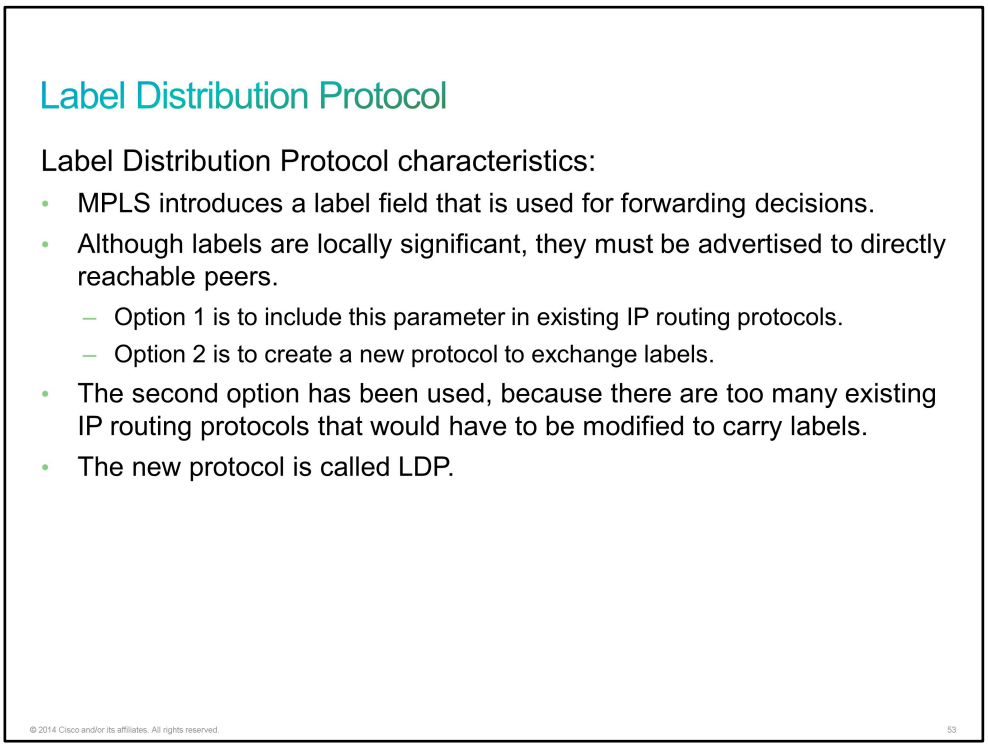

One application of MPLS is unicast IP routing. A label is assigned to destination IP networks and is later used to label packets sent toward those destinations.

Note	In MPLS terminology, an FEC in MPLS unicast IP routing equals an IP destination network.

Standard or vendor-specific routing protocols are used to advertise IP routing information. MPLS adds a new piece of information that must be exchanged between adjacent routers.

Here are the two possible approaches to propagating additional information (labels) between adjacent routers:

- Extend the functionality of existing routing protocols
- Create a new protocol dedicated to exchanging labels

The first approach requires much more time and effort because of the large number of different routing protocols: OSPF, IS-IS, EIGRP, IGRP, RIP, and so on. The first approach also causes interoperability problems between routers that support this new functionality and those that do not. Therefore, IETF selected the second approach and defined LDP in RFC 3036.

https://t.me/learningnets

LDP periodically sends hello messages (every 5 seconds). If the LSR is adjacent or one hop from its neighbor, the LSR sends out LDP link hello messages to all the routers on the subnet as UDP packets with a multicast destination address of 224.0.0.2 or FF02::2 (“all routers on a subnet”) and a destination port number of 646. A neighboring LSR enabled for LDP will respond by opening a TCP session with the same destination port number 646, and the two routers begin to establish an LDP session through unicast TCP.

https://t.me/learningnets

LDP Neighbor Adjacency Establishment

LDP is a standard protocol used to exchange labels between adjacent routers.

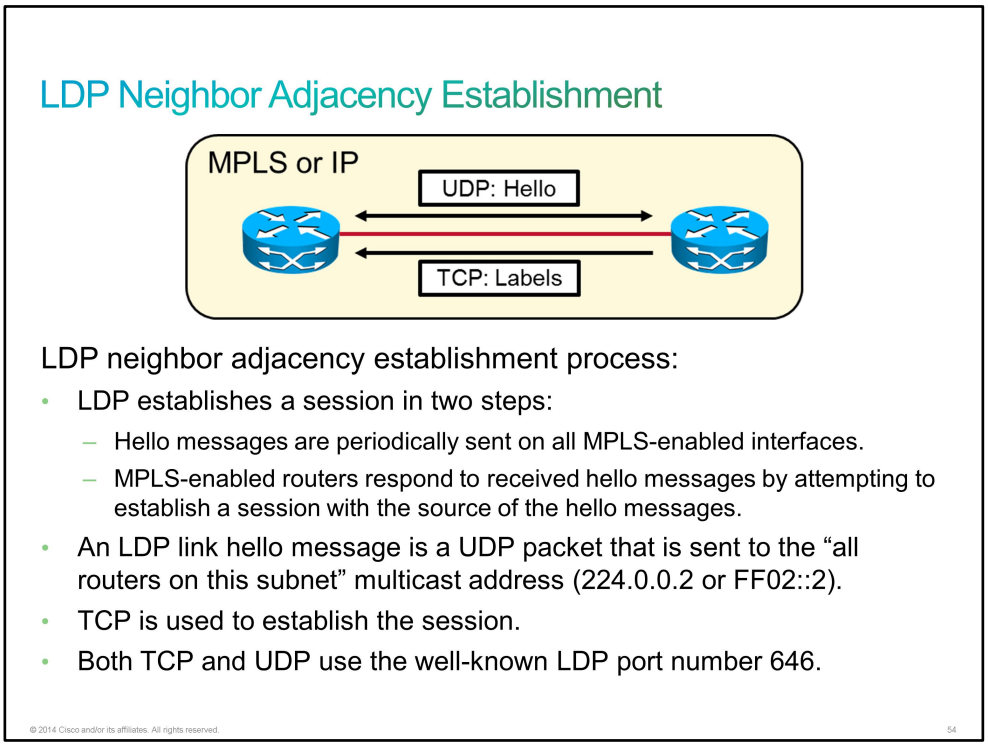

Before labels can be exchanged, MPLS-enabled routers must first establish adjacencies. This process is done in two steps:

- **LDP discovery:** MPLS routers first discover neighbors using hello messages that are sent to all the routers on the subnet as UDP packets with a multicast destination address of 224.0.0.2 (FF02::2 for IPv6) and destination port number of 646.
- **LDP adjacency:** A neighboring MPLS router, that received hello packets, will respond by opening a TCP session with the same destination port number 646, and the two routers begin to establish an LDP session through unicast TCP.

LDP periodically sends hello messages (every 5 seconds). If the LSR is adjacent or one hop from its neighbor, the LSR sends out LDP link hello messages to all the routers on the subnet as UDP packets with a multicast destination address of 224.0.0.2 or FF02::2 ("all routers on a subnet") and a destination port number of 646.

A neighboring LSR enabled for LDP will respond by opening a TCP session with the same destination port number 646, and the two routers begin to establish an LDP session through unicast TCP.

https://t.me/learningnets

LDP Link Hello Message

The figure shows the content of the LDP link hello message.

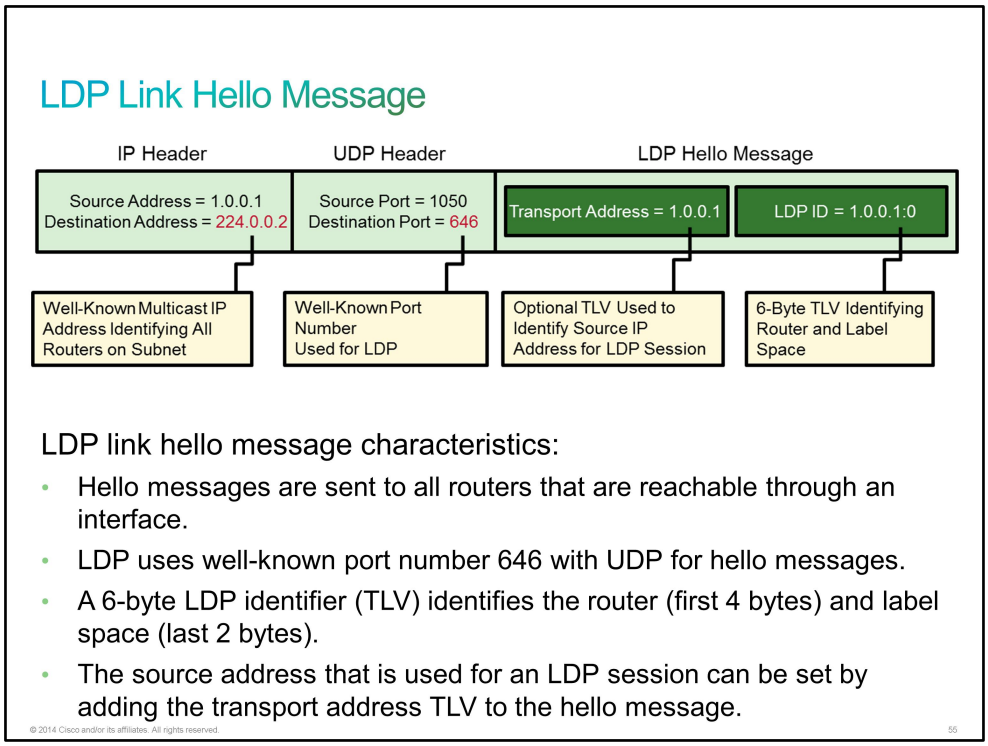

These are the contents of an LDP link hello message:

- Destination IP address (224.0.0.2 for IPv4 or FF02::2 for IPv6) that reaches all routers on the subnetwork
- Destination port that equals the LDP well-known port number 646
- The actual hello message that may optionally contain a transport address TLV to instruct the peer to open the TCP session to the transport address instead of the source address found in the IP header. The LDP ID is used to uniquely identify the neighbor and the label space.

Note	Label space defines the way MPLS assigns labels to destinations. Label space can be either per-platform or per-interface.

On Cisco routers, for all interface types, except the label-controlled ATM interfaces (running cell-mode MPLS-over-ATM), per-platform label space will be used where all the interfaces of the router share the same set of labels. For per-platform label space, the last two bytes of the LDP ID are always both 0. Multiple LDP sessions can be established between a pair of LSRs if they use multiple label spaces. For example, label-controlled cell-mode ATM interfaces use VPIs/VCIs for labels. Depending on the configuration, 0, 1, or more interface-specific label spaces can be used.

https://t.me/learningnets

LDP Neighbor Discovery

In the figure, three out of four routers periodically send out LDP hello messages (the fourth router is not MPLS-enabled).

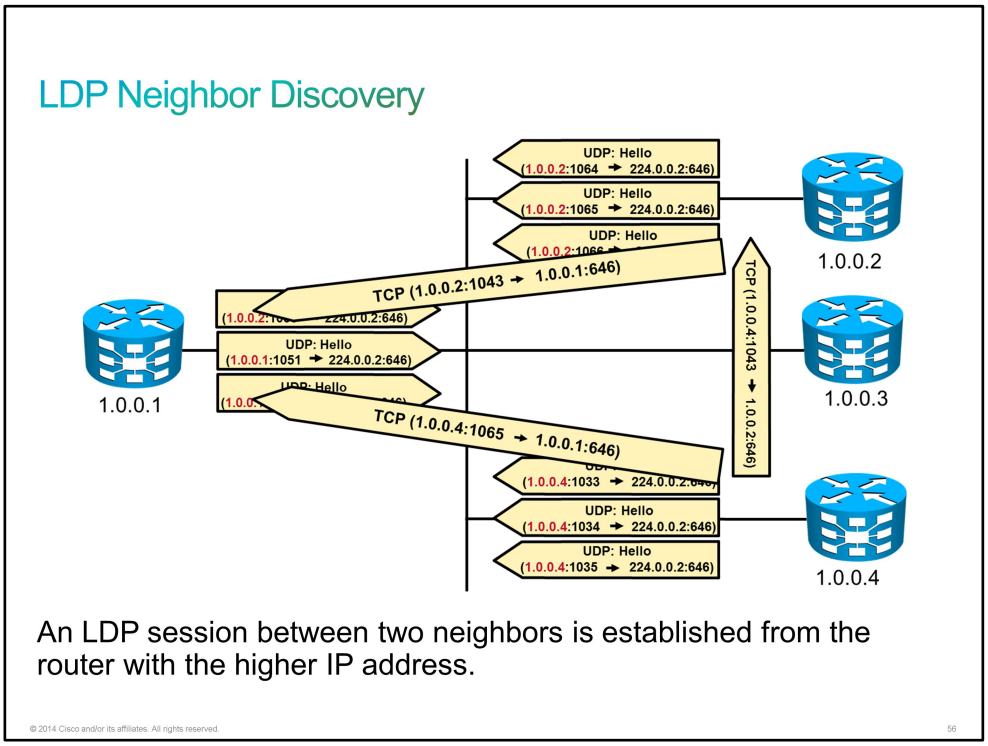

Routers that have the higher LDP router ID must initiate the TCP session. For instance, the router with the LDP router ID 1.0.0.2 initiates a TCP session to the router with LDP router ID 1.0.0.1.

If the LDP router ID is not manually configured, the highest IP address of all loopback interfaces on a router is used as the LDP router ID. If no loopback interfaces are configured on the router, the highest IP address of a configured interface that was operational at LDP startup is used as the LDP router ID.

On Cisco IOS XR Software, if the LDP router ID is not configured, the router can also default to the use of the global router ID as the LDP router ID. After the TCP session is established, routers will keep sending LDP hello messages to potentially discover new peers or to identify failures.

https://t.me/learningnets

LDP Session Negotiation

The figure below explains the LDP session negotiation.

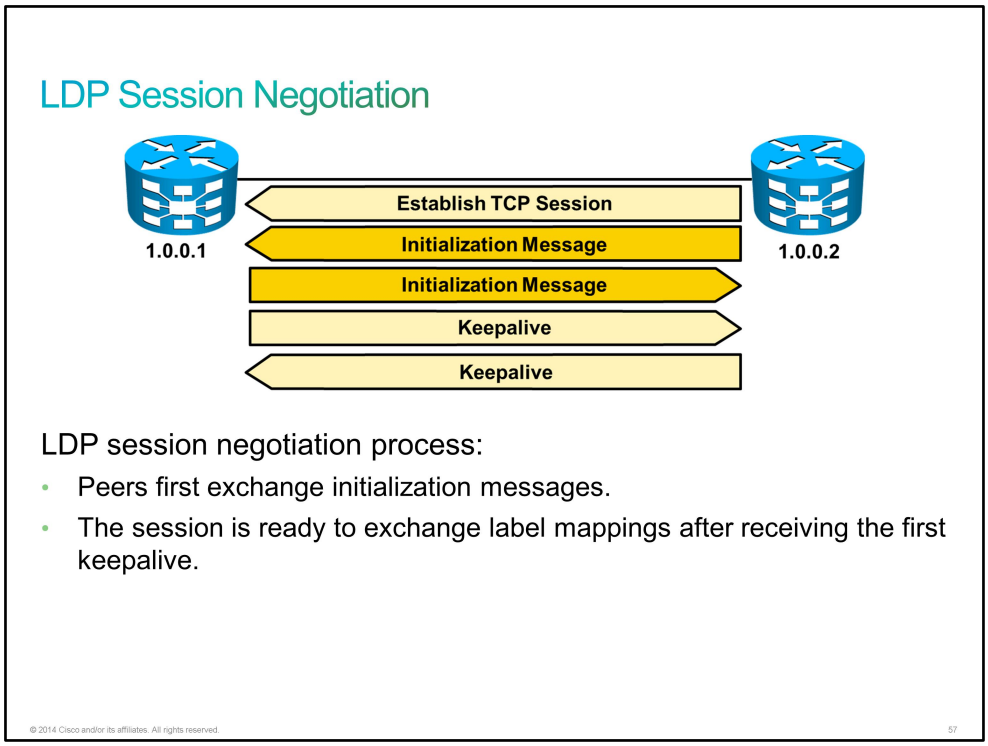

LDP session negotiation is a three-step process:

1. Establish the TCP session.
2. Exchange initialization messages that contain information such as, the label distribution method, the session keepalive time, and the fault-tolerant TLV. The LDP neighbor responds with an initialization message if the parameters are acceptable. If parameters are not acceptable, the LDP neighbor sends an error notification message.
3. Exchange initial keepalive messages.

Note	LDP keepalives are sent every 60 seconds.

After these steps, the two peers start exchanging labels for the networks that they have in their main routing tables.

https://t.me/learningnets

LDP Discovery of Nonadjacent Neighbors

If the LSR is more than one hop from its neighbor, it is not directly connected or adjacent to its neighbor. The LSR can be configured to send a directed hello message as a unicast UDP packet that is specifically addressed to the nonadjacent neighbor LSR. The directed hello message is called an LDP targeted hello.

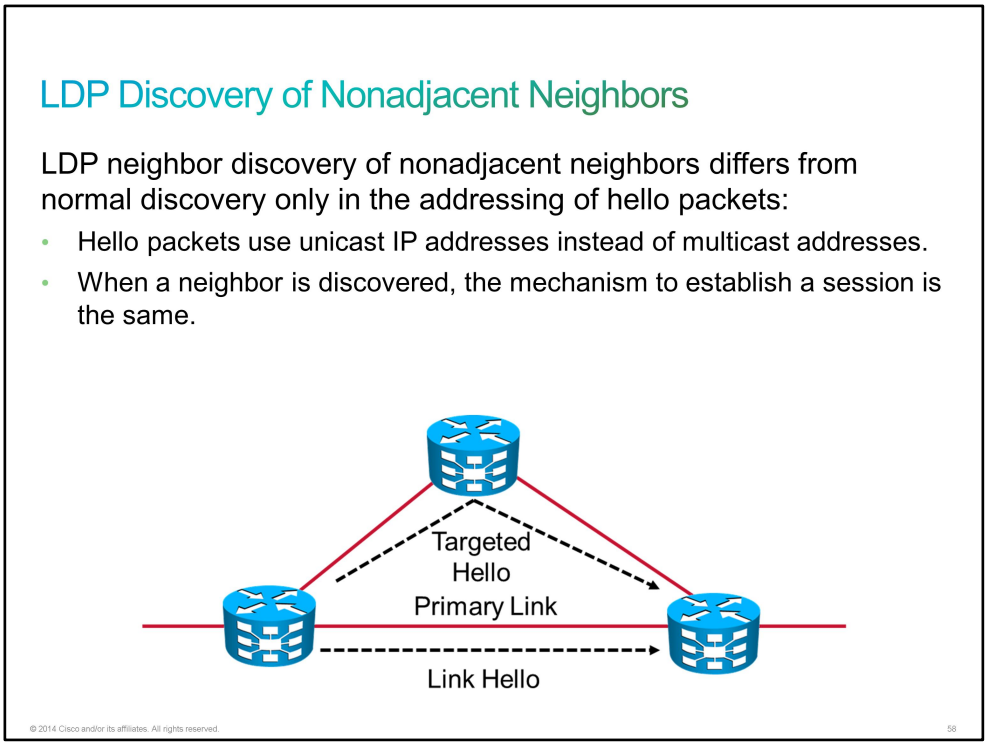

The rest of the session negotiation is the same as for adjacent routers. The LSR that is not directly connected will respond to the hello message by opening a unicast TCP session with the same destination port number 646, and the two routers begin to establish an LDP session.

For example, when you use the MPLS traffic engineering tunnel interface, a label distribution session is established between the tunnel headend and the tunnel tail end routers. To establish this not-directly-connected MPLS LDP session, the transmission of targeted LDP hello messages is used.

https://t.me/learningnets

LDP Session Protection

An example of using targeted LDP hello messages is between directly connected MPLS label switch routers, when MPLS label forwarding convergence time is an issue.

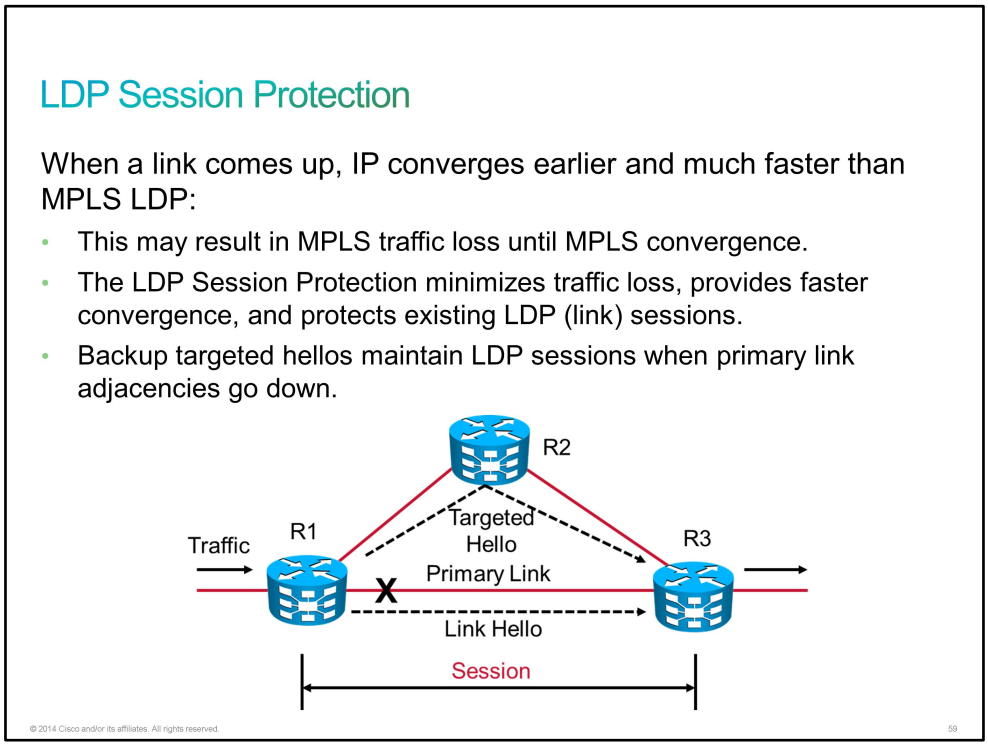

For example, when a link comes up, IP converges earlier and much faster than MPLS LDP and may result in MPLS traffic loss until MPLS convergence. If a link flaps, the LDP session will also flap due to loss of link discovery. LDP Session Protection minimizes traffic loss, provides faster convergence, and protects existing LDP (link) sessions by using a “parallel” source of targeted discovery hello. An LDP session is kept alive and neighbor label bindings are maintained when links are down. Upon re-establishment of primary link adjacencies, MPLS convergence is expedited because LDP does not need to relearn the neighbor label bindings.

LDP Session Protection lets you configure LDP to automatically protect sessions with all or a given set of peers (as specified by the peer-acl). When it is configured, LDP initiates backup targeted hellos automatically for neighbors for which primary link adjacencies already exist. These backup targeted hellos maintain LDP sessions when primary link adjacencies go down.

The Session Protection as shown in the figure illustrates LDP Session Protection between the R1 and R3 LDP neighbors. The primary link adjacency between R1 and R3 is a directly connected link, and the backup targeted adjacency is maintained between R1 and R3 through R2. If the direct link fails, the direct LDP link adjacency is destroyed, but the LDP session is kept functional using targeted hello adjacency (through R2). When the direct link comes back up, there is no change in the LDP session state and LDP can converge quickly and begin forwarding MPLS traffic.

MPLS LDP Session Protection protects a LDP session between directly connected neighbors or an LDP session established for a TE tunnel.

https://t.me/learningnets

LDP Graceful Restart and Nonstop Routing

LDP graceful restart provides a control plane mechanism to ensure high availability and allows detection and recovery from failure conditions while preserving NSF services. Graceful restart is a way to recover from signaling and control plane failures without impacting forwarding.

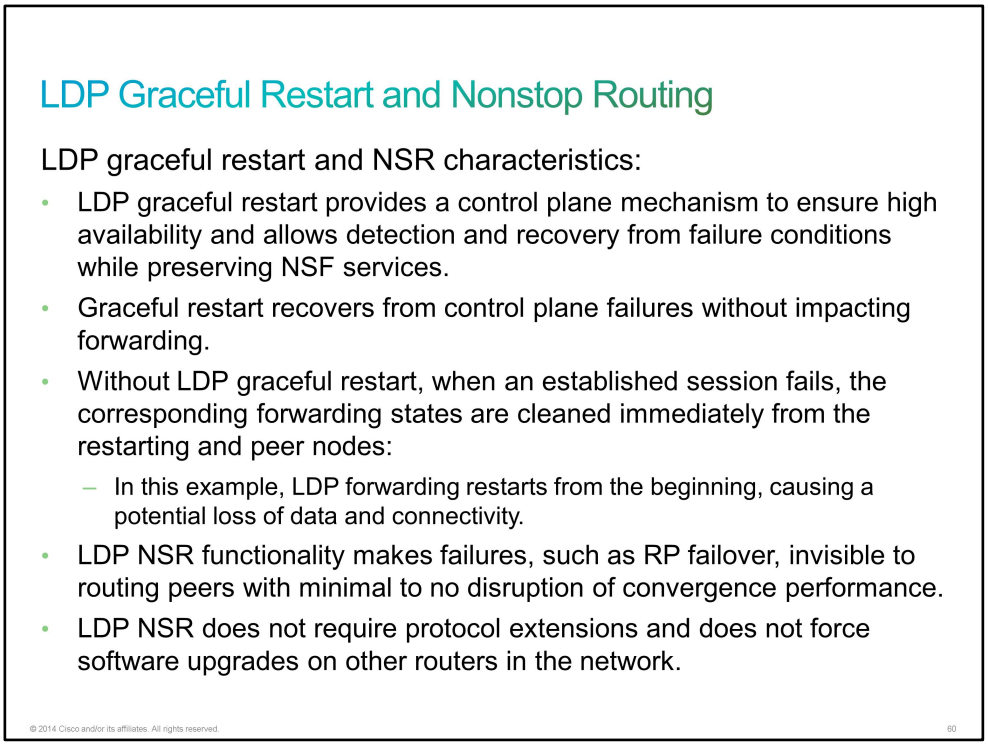

LDP graceful restart provides a control plane mechanism to ensure high availability and allows detection and recovery from failure conditions while preserving NSF services. Graceful restart is a way to recover from signaling and control plane failures without impacting forwarding.

Without LDP graceful restart, when an established LDP session fails, the corresponding forwarding states are cleaned immediately from the restarting and peer nodes. The LDP forwarding restarts from the beginning, causing a potential loss of data and connectivity.

The LDP graceful restart capability is negotiated between two peers during session initialization time, in the FT session TLV. In this TLV, each peer advertises the following information to its peers:

- An LSR indicates that it is capable of supporting LDP graceful restart by including the FT session TLV in the LDP initialization message and setting the L (Learn from Network) flag to 1.
- **Reconnect time:** Advertises the maximum time that the other peer will wait for this LSR to reconnect after control channel failure.
- **Recovery time:** Advertises the maximum time that the other peer will retain its MPLS forwarding state that it preserved across the restart. The recovery time should be long enough to allow the neighboring LSRs to resynchronize their MPLS forwarding state in a graceful manner. This time is used only during session reestablishment after earlier session failure.

Once the graceful restart session parameters are conveyed and the session is functioning, graceful restart procedures are activated.

https://t.me/learningnets

LDP NSR functionality makes failures, such as RP or DRP failover, invisible to routing peers with minimal to no disruption of convergence performance.

Unlike graceful restart functionality, LDP NSR does not require protocol extensions and does not force software upgrades on other routers in the network, nor does LDP NSR require peer routers to support NSR.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- MPLS uses LDP to exchange labels.
- UDP multicast is used to discover adjacent LDP neighbors, while TCP is used to establish a session.
- The LDP link hello message contains the destination IP address, destination port, and the actual hello message.
- The LDP session between two neighbors is established from the router with the higher IP address.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- LDP session negotiation is a three-step process.
- An MPLS-enabled router can be configured to send a directed LDP hello message as a unicast UDP packet that is specifically addressed to the nonadjacent router.
- LDP Session Protection lets you configure LDP to automatically protect sessions with all or a given set of peers.
- LDP graceful restart provides a control plane mechanism to ensure high availability and allows detection and recovery from failure conditions.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 5: Introducing MPLS Forwarding

Overview

This lesson discusses how label allocation and distribution function in an MPLS network, the use of PHP, and how the MPLS data structures are built. These topics are essential for understanding the fundamentals of the way that information gets distributed and placed into the appropriate tables, for both labeled and unlabeled packet usage.

Upon completing this lesson, you will be able to:

- Describe how the forwarding structures used by MPLS are populated
- Explain the LSP
- Explain the building of the LSP
- Explain the MPLS label allocation and distribution process
- Explain the MPLS label allocation process.
- Describe label advertisement
- Describe how IP packets are propagated across an MPLS domain
- Describe how IP packets would be propagated across an MPLS domain without PHP
- Define the steady state condition when all the labels are exchanged by LDP and the LIBs, LFIBs, and FIBs are completely populated
- Explain label advertisement control and label acceptance control
- Explain the how IP aggregation in the core can break an LSP into two segments

https://t.me/learningnets

MPLS Forwarding Structures

Forwarding structures that are used by MPLS need to be populated with labels. LDP, which runs in the router control plane, is responsible for label allocation, distribution, and storage.

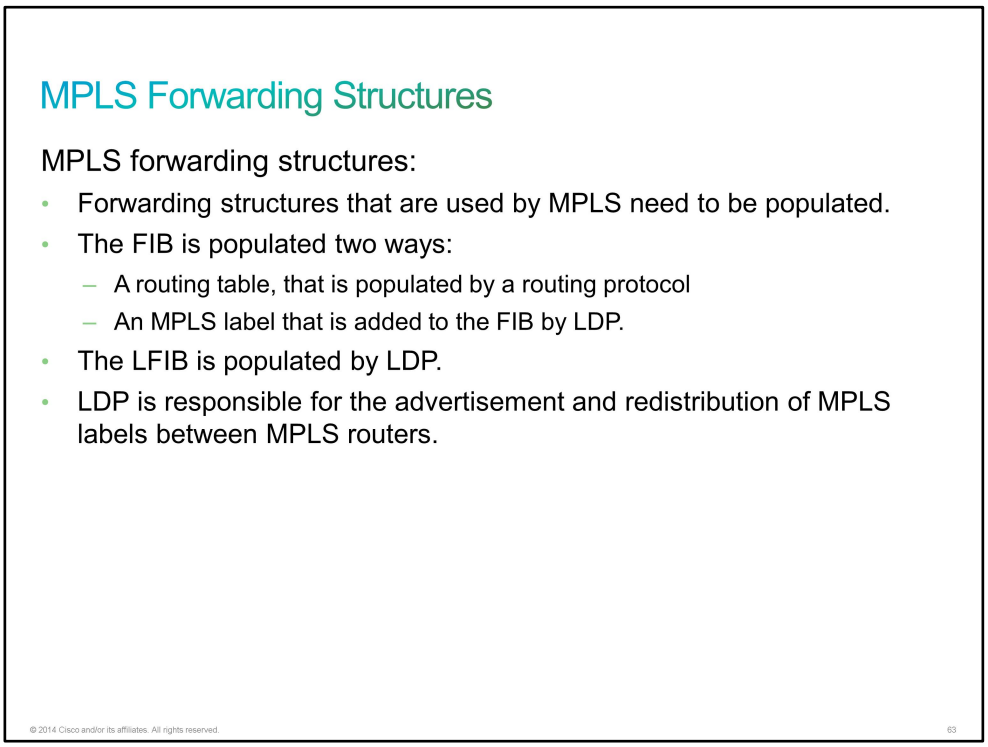

The FIB table, which consists of destination networks, next hops, outgoing interfaces, and pointers to Layer 2 addresses, is populated by using information from the routing table and from the ARP cache. The routing table is, in turn, populated by a routing protocol. Additionally, the MPLS label is added to destination networks if an outgoing interface is enabled for MPLS and a label has been received from the next-hop router. LDP is responsible for adding a label to the FIB table entries.

The LFIB table contains incoming (locally assigned) and outgoing (received from next hop) labels. LDP is responsible for exchanging labels and storing them into the LFIB.

https://t.me/learningnets

Label-Switched Path

An LSP is a sequence of LSRs that forwards labeled packets for a particular FEC. Each intermediate LSR swaps the top label in a packet traversing the LSP.

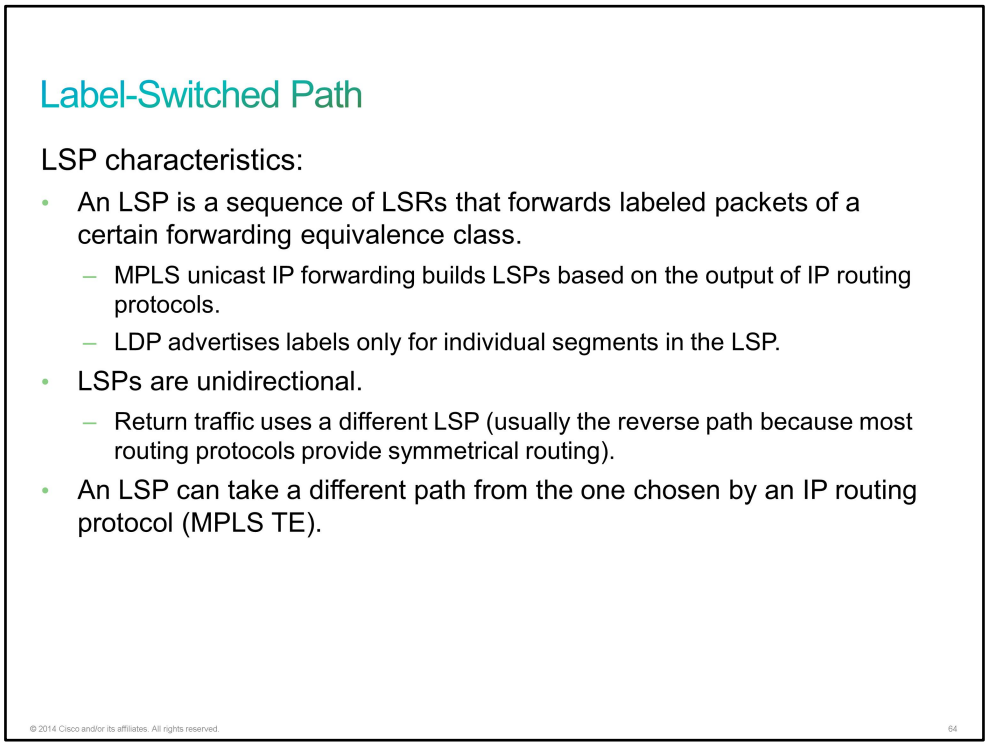

In MPLS unicast IP forwarding, the FECs are determined by destination networks that are found in the main routing table. Therefore, an LSP is created for each entry that is found in the main routing table. BGP entries are the only exceptions for that rule.

An FEC is a group of packets that are forwarded along the same path and get the same forwarding treatment

In an ISP environment where a routing table may contain more than 100,000 routes, to minimize the number of labels that are needed in such networks, an exception was made for BGP-derived routing information. All BGP entries in the main routing table use the same label that is used to reach the BGP next-hop router (the PE router). Only the PE routers are required to run BGP. All the core (P) routers run an IGP to learn about the BGP next-hop addresses. The core routers run LDP to learn about labels for reaching the BGP next-hop addresses. This results in one single label being used for all networks that are learned from the BGP neighbor.

An IGP is used to populate the routing tables in all routers in an MPLS domain. LDP is used to propagate labels for these networks and build LSPs.

LSPs are unidirectional. Each LSP is created over the shortest path, selected by the IGP, toward the destination network. Packets in the opposite direction use a different LSP. The return LSP is usually over the same LSRs, except that the packets forming the LSP are in the opposite order.

Cisco MPLS TE can be used to change the default IGP shortest path selection.

https://t.me/learningnets

LSP Building

The figure illustrates how an IGP, such as OSPF, IS-IS, or EIGRP, propagates routing information to all routers in an MPLS domain. Each router determines its own shortest path.

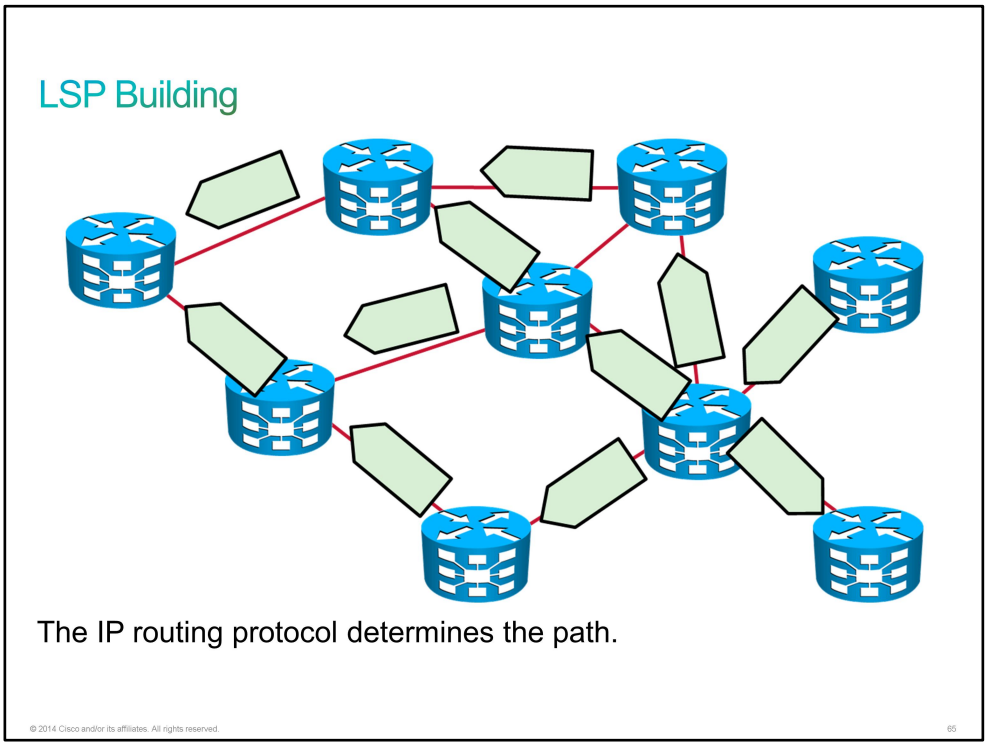

LDP, which propagates labels for those networks, adds labels to the FIB and LFIB tables.

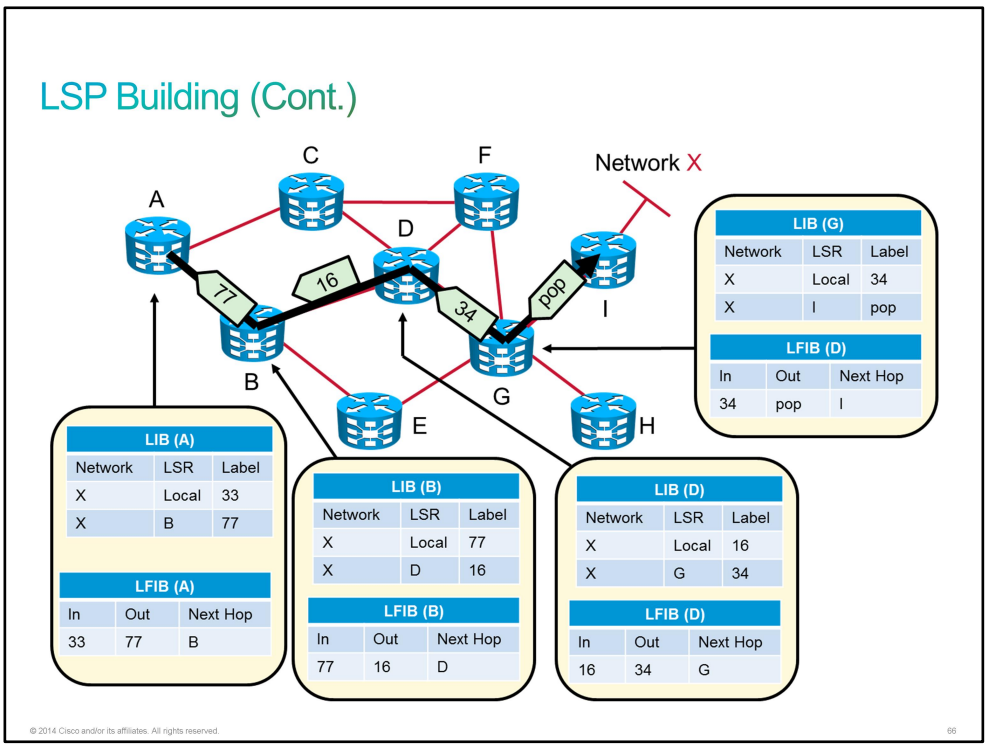

https://t.me/learningnets

The figure shows the contents of LFIB and LIB tables. MPLS uses a liberal label retention mode, which means that each LSR keeps all labels received from LDP peers, even if they are not the downstream peers (the next-hop) for reaching network X. With a liberal retention mode, an LSR can almost immediately start forwarding labeled packets after IGP convergence, but the numbers of labels maintained for a particular destination will be larger and thus will consume more memory.

LIB and LFIB tables are shown on the routers for the LSP A-B-D-G-I. On each router for this path, only local labels, and labels received from adjacent routers forming this path, are shown in LIB tables.

Note	Notice that router G receives a pop label from final destination router I. The pop action results in the removal of the label rather than swapping labels. This allows the regular IP packet to be forwarded out on the router I interface that is directly connected to network X when the packet leaves the MPLS domain.

https://t.me/learningnets

Label Allocation and Distribution

Unicast IP routing and MPLS functionality can be divided into several steps.

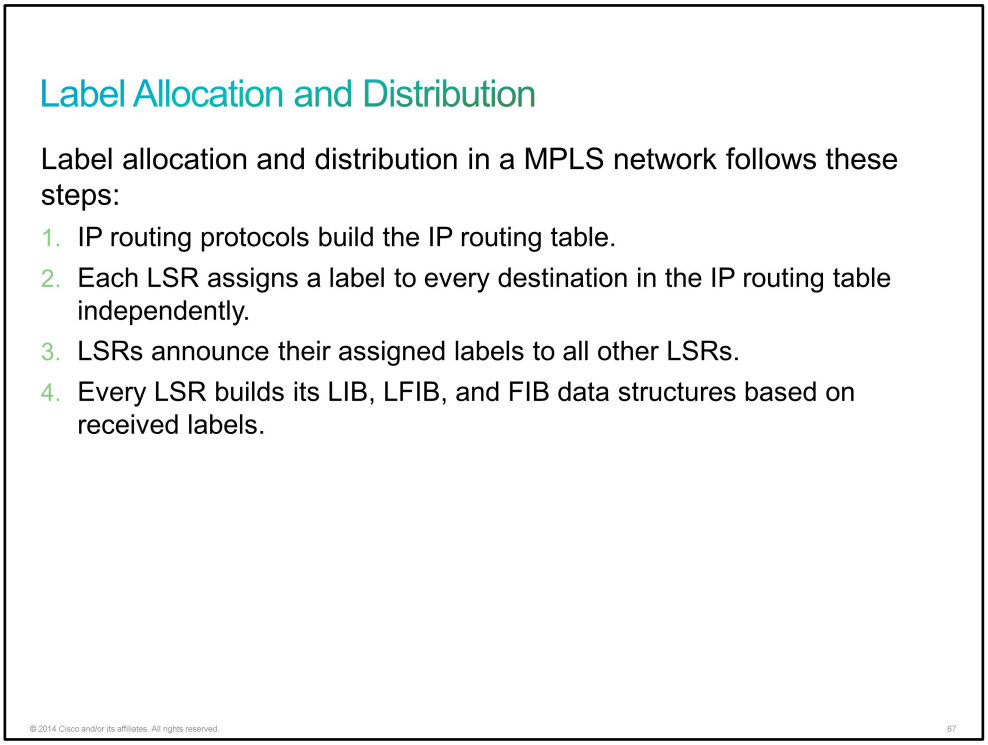

Unicast IP routing and MPLS functionality is divided into these steps:

- Routing information is exchanged using standard or vendor-specific IP routing protocols (OSPF, IS-IS, EIGRP, and so on).
- Local labels are generated. (One locally unique label is assigned to each IP destination found in the main routing table and stored in the LIB table.)
- Local labels are propagated to adjacent routers, where these labels might be used as next-hop labels (stored in the FIB and LFIB tables to enable label switching).

https://t.me/learningnets

These data structures contain label information:

- LSRs store labels and related information inside a data structure called a LIB. The FIB and LFIB contain labels only for the currently used best LSP segment, while the LIB contains all labels known to the LSR, whether the label is currently used for forwarding or not. The LIB in the control plane is the database that is used by LDP; an IP prefix is assigned a locally significant label, which is mapped to a next-hop label that has been learned from a downstream neighbor.
- The LFIB, in the data plane, is the database used to forward labeled packets that are received by the router. Local labels, previously advertised to upstream neighbors, are mapped to next-hop labels, previously received from downstream neighbors.
- The FIB, in the data plane, is the database used to forward unlabeled IP packets that are received by the router. 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.

https://t.me/learningnets

Label Allocation

First, each MPLS-enabled router must locally allocate a label for each network that is known to a router. Labels are locally significant (a label for the same network has a different value on different routers), and allocation of labels is asynchronous (routers assign labels, independent of each other).

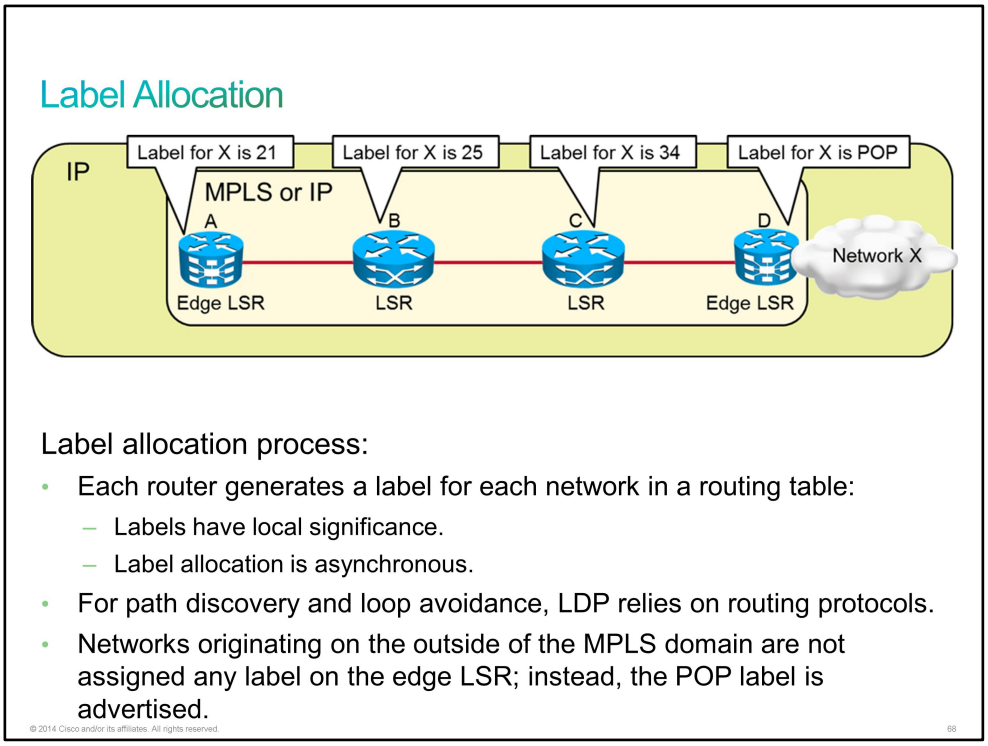

The LDP is not responsible for finding a shortest, loop-free path to destinations. Instead, the LDP relies on routing protocols to find the best path to destinations. If, however, a loop does occur, a TTL field in the MPLS label prevents the packet from looping indefinitely.

On the edge LSR, networks originating on the outside of the MPLS domain are not assigned a label. Instead, the POP (implicit null) label is advertised, which instructs the penultimate router to remove the label.

In the example, all routers except router D assign a label for network X. Router D assigns an implicit null label for that network.

https://t.me/learningnets

Label Advertisement

After a label has been assigned locally, each router must advertise a label to neighbors.

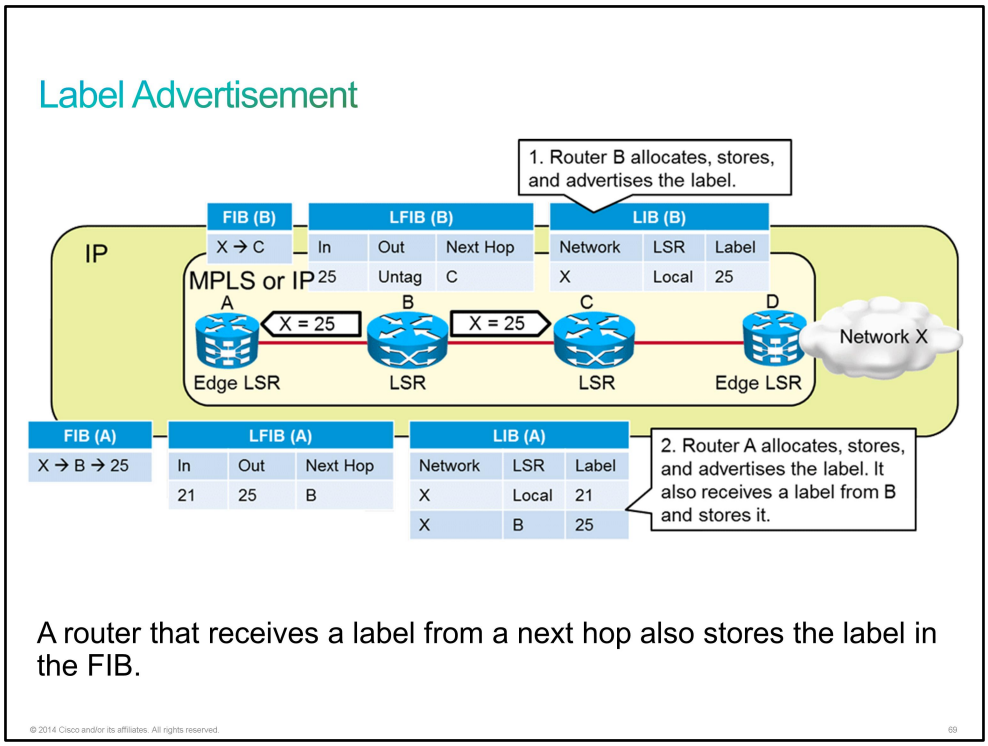

The figure shows how a label is assigned and advertised to neighbors on router B.

1. Router B allocates label 25 for network X. The allocated label is first stored in the LIB, which stores local labels and labels received from neighbors as well. The label is also stored in a LFIB table as an incoming label. The outgoing label has not yet been set, because router B has not received a label for network X from the next-hop router (router C) yet. The allocated label is also advertised to the neighbors (routers A and C), regardless of whether a neighbor actually is a next hop for a destination or not.
2. Router A allocates his own label for network X (21 in the example). This label is again stored in the LIB and in the LFIB as an incoming label. Router A also receives a label 25 from router B and stores the label in the LIB. Because label 25 has been received from a next hop for destination X, router A also stores label 25 in the LFIB as an outgoing label. Router A also sets the label 25 for destination X in the FIB table, because the label has been received from the next hop.

If a packet for network X was received by router A (not shown in the figure), an FIB lookup would be done. The packet would be labeled using label 25 and sent to router B. Router B would perform an LFIB lookup, which would state that the label should be removed, because the outgoing label had not yet been received from the next-hop router (router C).

https://t.me/learningnets

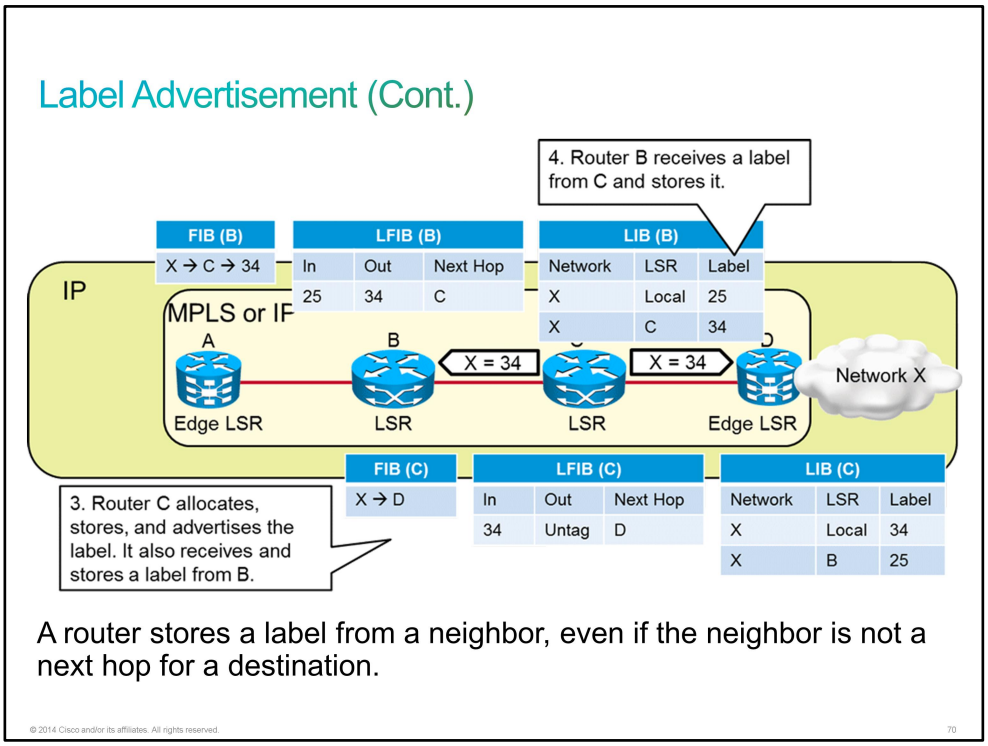

3. Router C allocates label 34 for network X. The allocated label is first stored in the LIB. The label is also stored in the LFIB table as an incoming label. The outgoing label has not yet been set, because router C has not received a label for network X from the next-hop router (router D) yet. The allocated label is also advertised to the neighbors (routers B and D), regardless of whether a neighbor actually is a next hop for a destination or not. Router C also receives a label 25 from router B and stores it in its LIB, even though router B is not a next hop for destination X.
4. Router B receives a label 34 from router C and stores the label in the LIB. Because label 34 has been received from a next hop for destination X, router B also stores label 34 in an LFIB as an outgoing label. Router B also sets the label 34 for destination X in the FIB table, because the label has been received from the next hop.

https://t.me/learningnets

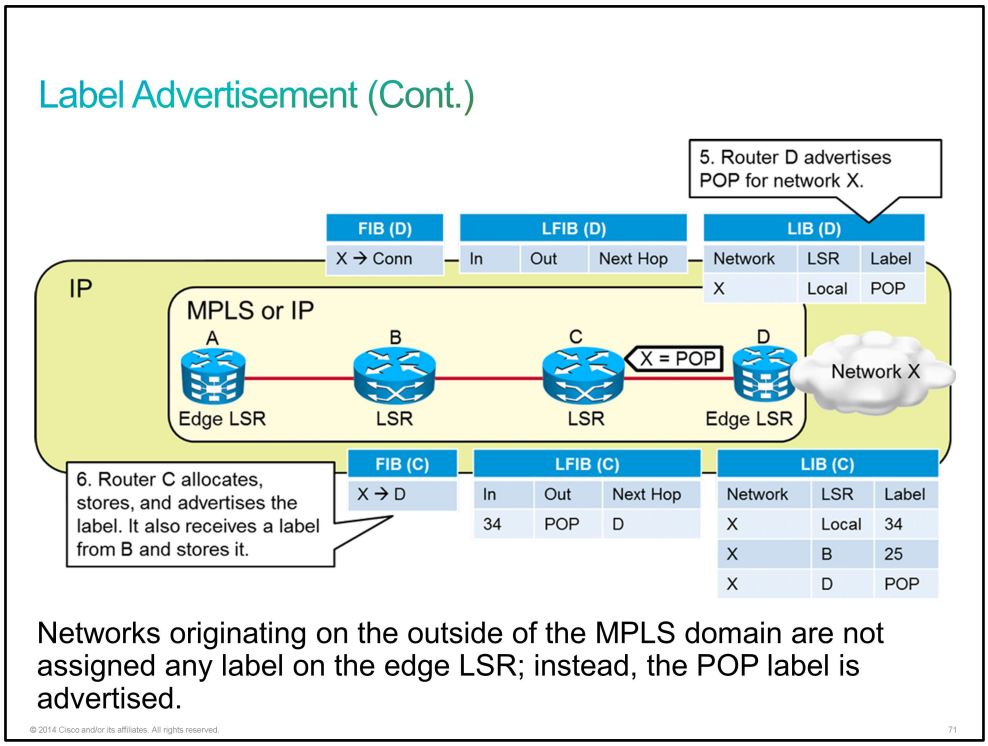

5. Router D allocates the implicit null label for network X. The allocated label is first stored in the LIB. The implicit null label is also advertised to router C. The implicit null label indicates to the upstream router that it should perform a label removal (pop) operation.
6. Router C receives the implicit null label from router D and stores the label in the LIB. Because the label has been received from a next hop for destination X, router C also stores the label in LFIB as an outgoing label. Because the implicit null label indicates that the label should be removed, router C does set the label in the FIB table. The LSP for a network X is now established.

https://t.me/learningnets

Packet Propagation Across an MPLS Domain

The figure illustrates how IP packets are propagated across an MPLS domain.

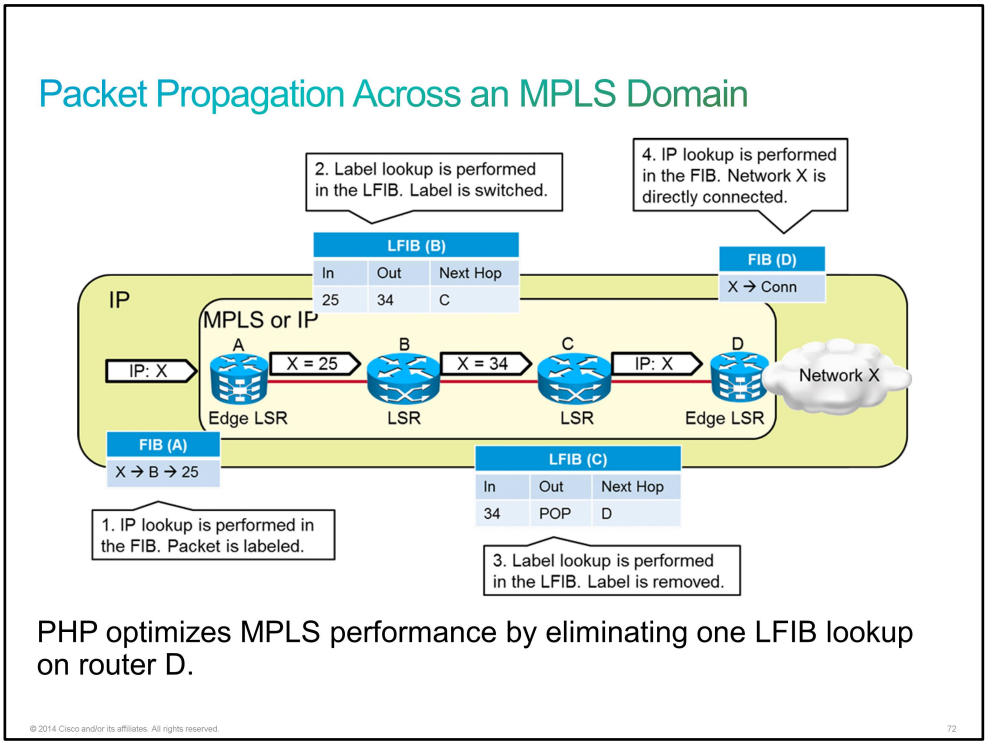

The steps are as follows:

1. Router A labels a packet destined for network X by using the next-hop label 25 (Cisco Express Forwarding, switching by using the FIB table).
2. Router B swaps label 25 with label 34 and forwards the packet to router C (label switching by using the LFIB table).
3. Router C removes (pops) the label and forwards the packet to router D (label switching by using the LFIB table).
4. Router D performs IP lookup in the FIB table. Network X is directly connected.

The figure assumes that the implicit null label, which corresponds to the pop action in the LFIB, has been propagated from the egress router (router D) to router C. 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. The process is called PHP, which is enabled by default on all MPLS-enabled routers.

PHP optimizes MPLS performance by eliminating one LFIB lookup at the egress router, as only FIB lookup is needed on router D, the egress router.

https://t.me/learningnets

Packet Propagation Without PHP

The figure illustrates how IP packets would be propagated across an MPLS domain without the PHP process.

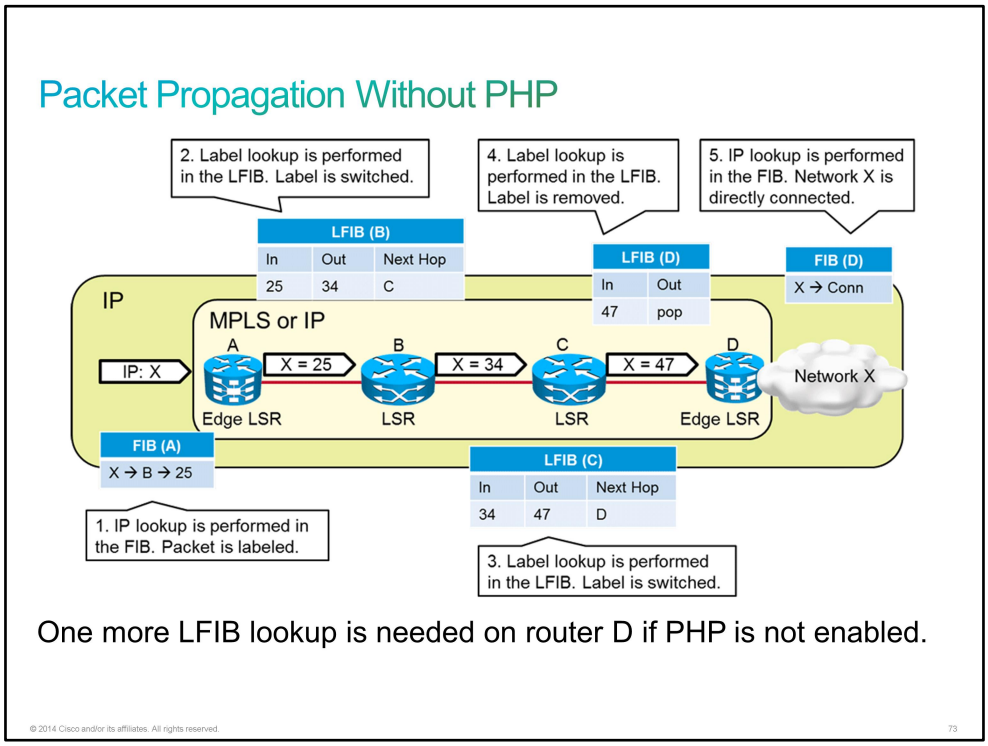

The steps are as follows:

1. Router A labels a packet destined for network X by using the next-hop label 25 (Cisco Express Forwarding switching by using the FIB table).
2. Router B swaps label 25 with label 34 and forwards the packet to router C (label switching by using the LFIB table).
3. As PHP is not enabled, router D does not advertise the implicit null (pop) label to the router C. Therefore, router C swaps label 34 with label 47 and forwards the packet to router D (label switching by using the LFIB table).
4. Router D removes the label (label switching by using the LFIB table).
5. Router D performs IP lookup in the FIB table. Network X is directly connected.

As you can see, one more LFIB lookup is needed on router D if PHP is not enabled.

https://t.me/learningnets

MPLS Steady State Condition

MPLS is fully functional when the routing protocol and LDP have populated all the tables.

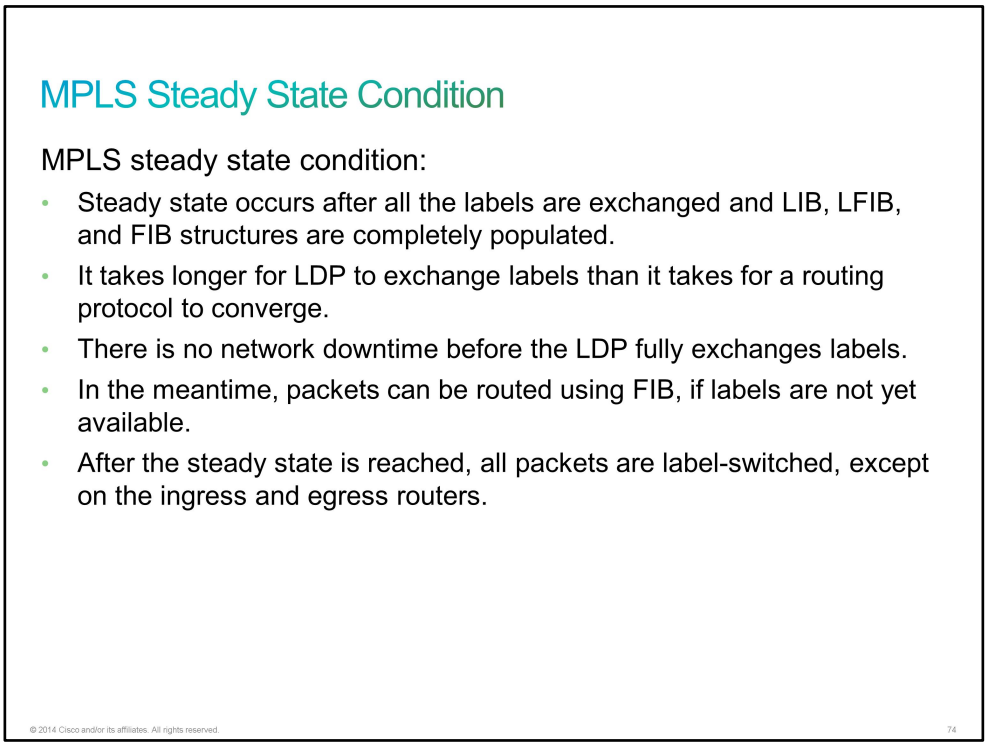

The tables are as follows:

- Main IP routing table
- LIB table
- FIB table
- LFIB table

Such a state is called the steady state. After the steady state, all packets are label-switched, except on the ingress and egress routers (edge LSR).

Although it takes longer for LDP to exchange labels (compared with a routing protocol), a router can use the FIB table in the meantime. Therefore, there is no routing downtime while LDP exchanges labels between adjacent LSRs.

https://t.me/learningnets

MPLS Label Control Methods

By default, LDP advertises labels for all the prefixes to all its neighbors.

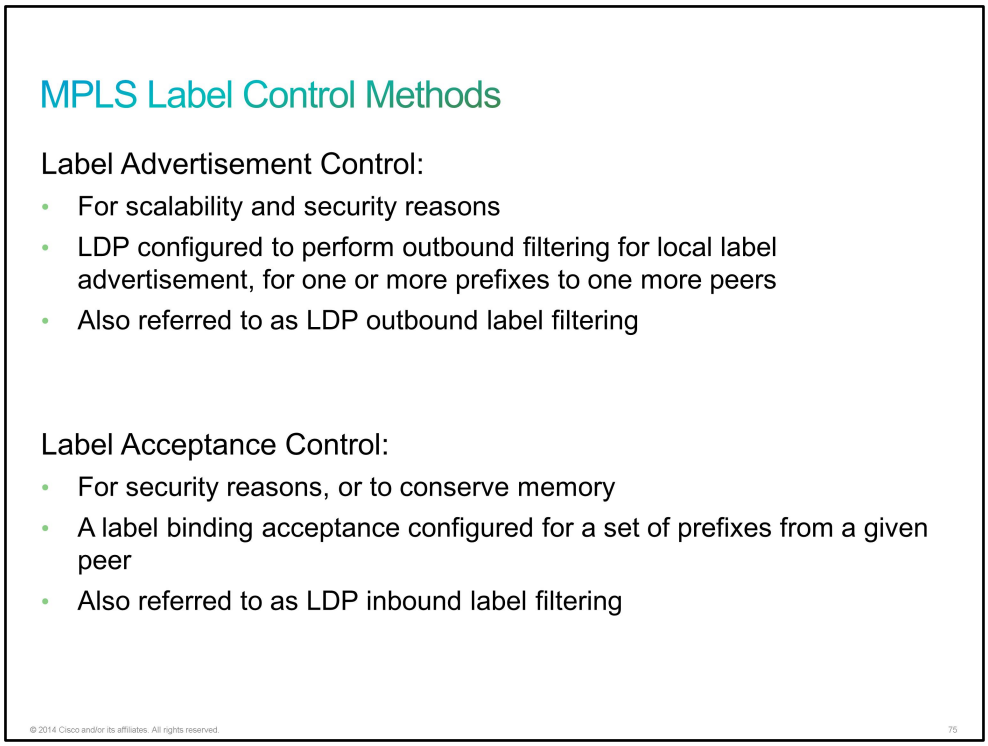

When LDP advertisement of labels for all the prefixes to all its neighbors is not desirable (for scalability and security reasons), you can configure LDP to perform outbound filtering for local label advertisement for one or more prefixes to one or more LDP peers. This feature is known as LDP outbound label filtering, or local label advertisement control.

By default, LDP accepts labels (as remote bindings) for all prefixes from all LDP peers. LDP operates in liberal label retention mode, which instructs LDP to keep remote label bindings from all LDP peers for a given prefix, even if the LDP peer is not the next-hop router. For security reasons, or to conserve memory, you can override this behavior by configuring label binding acceptance for a set of prefixes from a given LDP peer.

The ability to filter remote bindings for a defined set of prefixes is also referred to as LDP inbound label filtering.

https://t.me/learningnets

Impact of IP Aggregation on LSPs

The figure illustrates a potential problem in an MPLS domain.

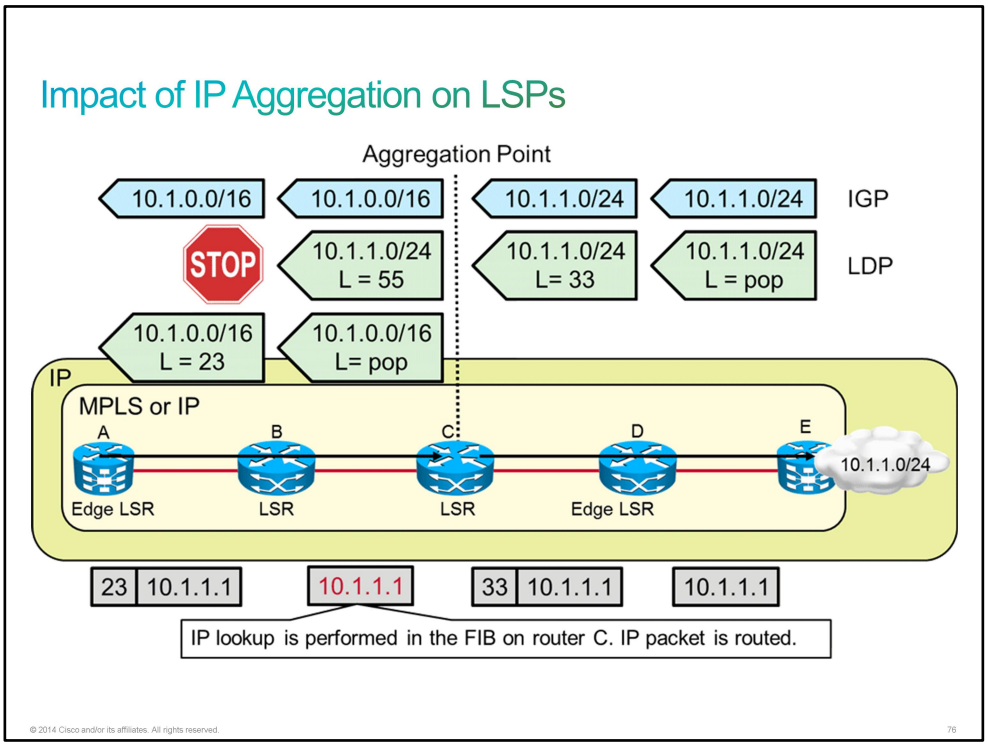

An IGP propagates the routing information for network 10.1.1.0/24 from router E to other routers in the network. Router C uses a route summarization mechanism to stop the proliferation of all subnetworks of network 10.1.0.0/16. Only the summary network 10.1.0.0/16 is sent to routers B and A.

LDP propagates labels concurrently with the IGP. The LSR that is the endpoint of an LSP always propagates the “pop” label.

Router C has both networks in the routing table:

- 10.1.1.0/24 (the original network)
- 10.1.0.0/16 (the summary route)

Router C, therefore, sends a label, 55 in the example, for network 10.1.1.0/24 to router B. Router C also sends an implicit null (pop) label for the new summary network 10.1.0.0/16 that originates on router C. Router B, however, can use the implicit null (pop) label only for the summary network 10.1.0.0/16 because it has no routing information about the more specific network 10.1.1.0/24; this information was suppressed on router C.

The route summarization results in two LSPs for destination network 10.1.1.0/24. The first LSP ends on router C, where a routing lookup is required to assign the packet to the second LSP.

https://t.me/learningnets

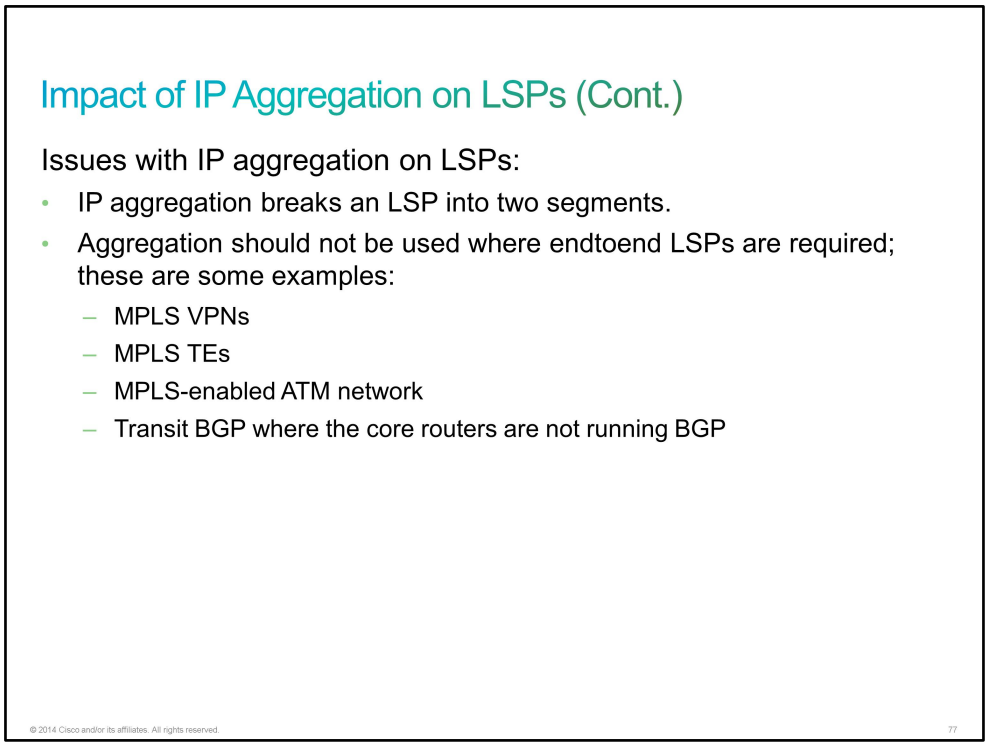

Aggregation should also not be used where an end-to-end LSP is required. Here are some typical examples of networks that require end-to-end LSPs:

- An MPLS VPN backbone
- A network that uses MPLS TE
- An MPLS-enabled ATM network
- A transit BGP AS, where the core routers are not running BGP

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

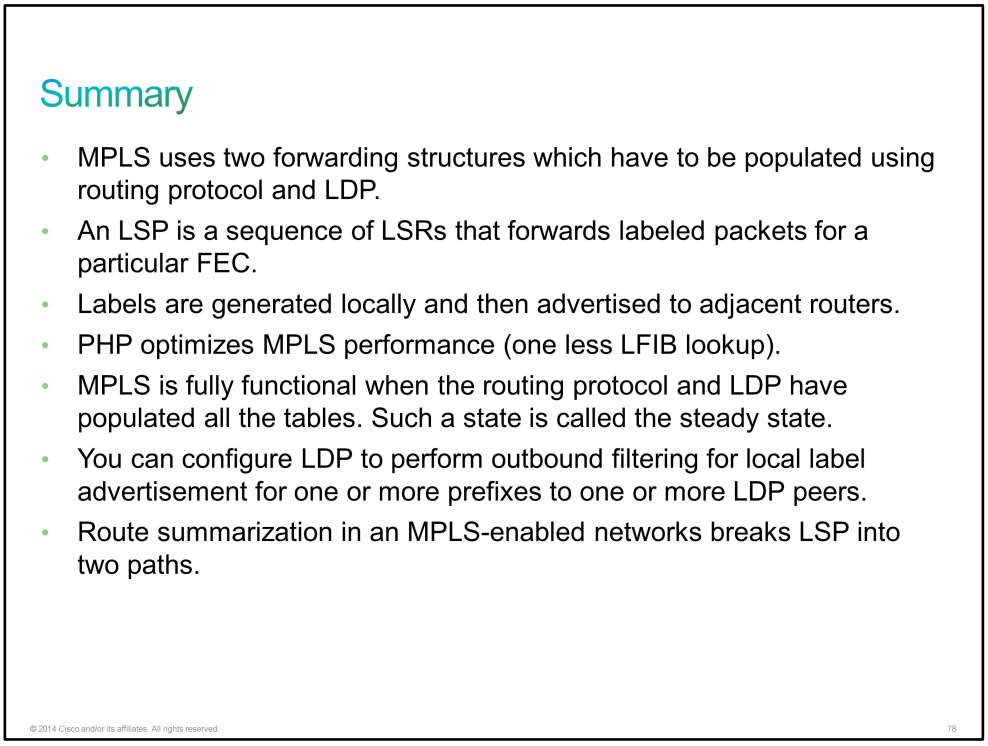

https://t.me/learningnets

Lesson 6: Operating MPLS Forwarding

Overview

This lesson presents LDP convergence issues, describes how routing protocols and MPLS convergence interact, and concludes with an explanation of the switching mechanisms on various Cisco platforms. The lesson also describes loop detection using the MPLS TTL field.

Upon completing this lesson, you will be able to:

- Describe loop detection using the MPLS TTL field
- Describe normal TTL operation
- Describe the disabling of TTL propagation to hide the core routers in the MPLS domain
- Describe a steady state operation in the MPLS domain
- Describe how a link failure is managed in an MPLS domain
- Describe routing protocol convergence in the MPLS network
- Describe the overall convergence in an MPLS network
- Describe how a link recovery is managed in an MPLS domain
- Describe the three IP switching mechanisms (process switching, fast switching, and Cisco Express Forwarding)
- Describe standard IP switching on Cisco IOS platforms
- Describe a Cisco Express Forwarding switching example
- Describe Cisco Express Forwarding on Cisco IOS XE platforms
- Describe Cisco Express Forwarding on Cisco IOS XR platforms
- Describe the **show** commands used to monitor Cisco Express Forwarding operations

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Loop Detection using the MPLS TTL field

Loop detection in an MPLS-enabled network relies on more than one mechanism.

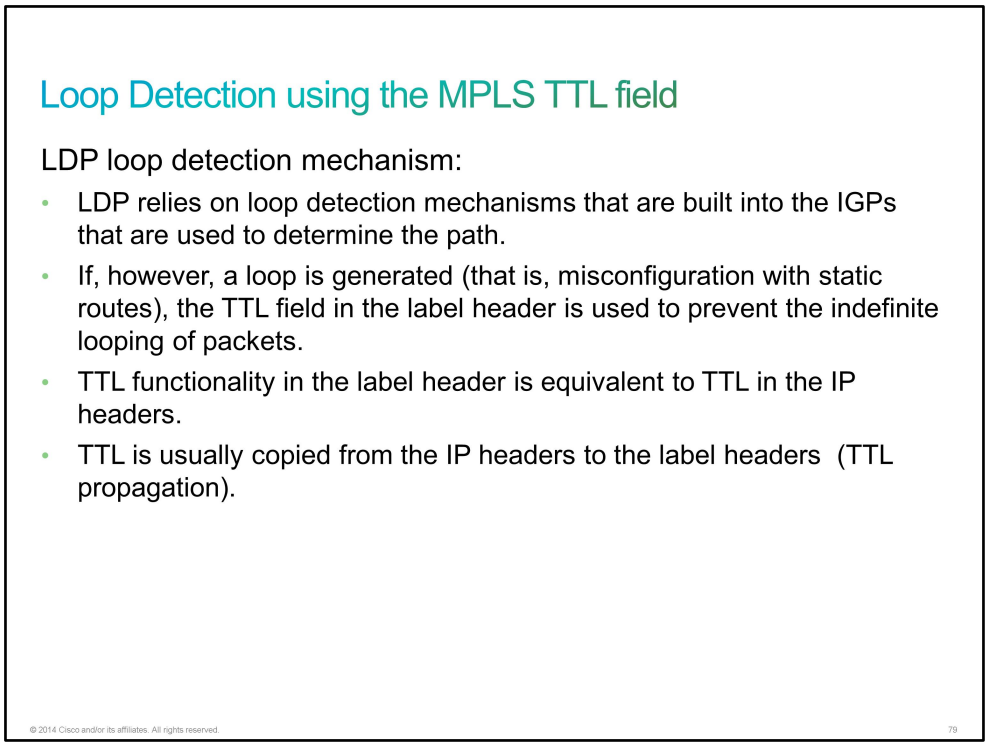

Most routing loops are prevented by the IGP that is used in the network. MPLS for unicast IP forwarding simply uses the shortest paths as determined by the IGP. These paths are typically loop-free.

If, however, a routing loop does occur (for example, because of misconfigured static routes), MPLS labels also contain a TTL field that prevents packets from looping indefinitely.

The TTL functionality in MPLS is equivalent to that of traditional IP forwarding. Furthermore, when an IP packet is labeled, the TTL value from the IP header is copied into the TTL field in the label. This process is called TTL propagation.

Note	The LDP loop detection mechanism is intended for use in networks of non-TTL decrementing devices such as ATM switches.

https://t.me/learningnets

Normal TTL Operation

The figure illustrates how the TTL value 5 in the IP header is decreased and copied into the TTL field of the label when a packet enters an MPLS domain.

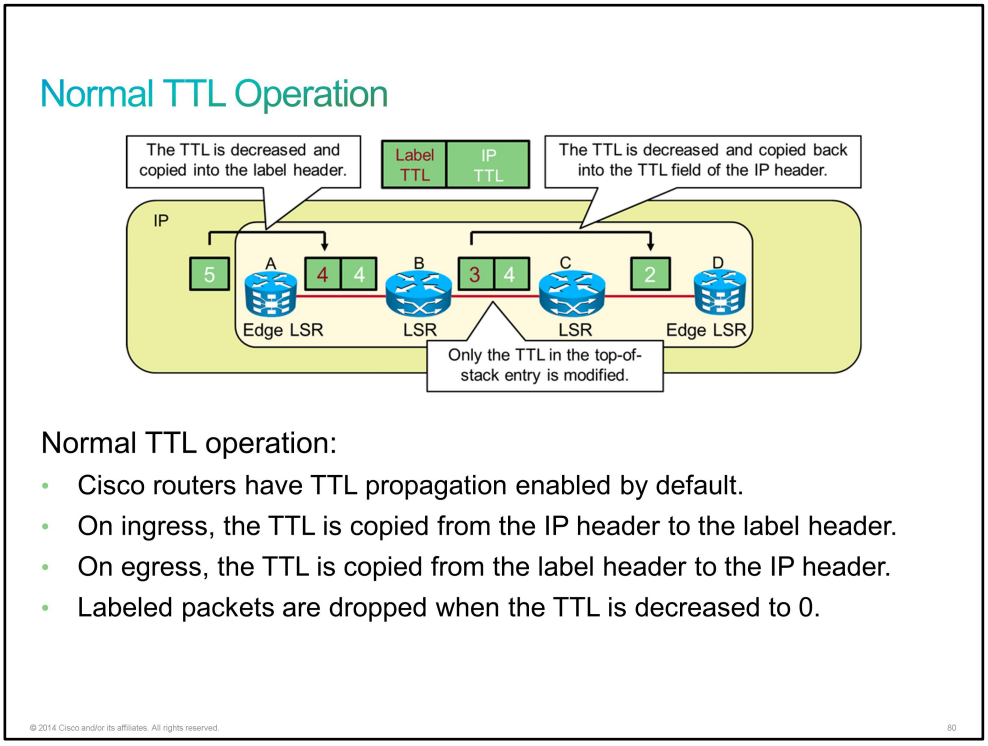

All other LSRs decrease the TTL field only in the label. The original TTL field is not changed until the last label is removed when the label TTL is copied back into the IP TTL.

TTL propagation provides a transparent extension of IP TTL functionality into an MPLS-enabled network.

The packet looping between these two routers is eventually dropped because the value of its TTL field reaches 0.

https://t.me/learningnets

Disabling TTL Propagation

TTL propagation can be disabled to hide the core routers from the end users.

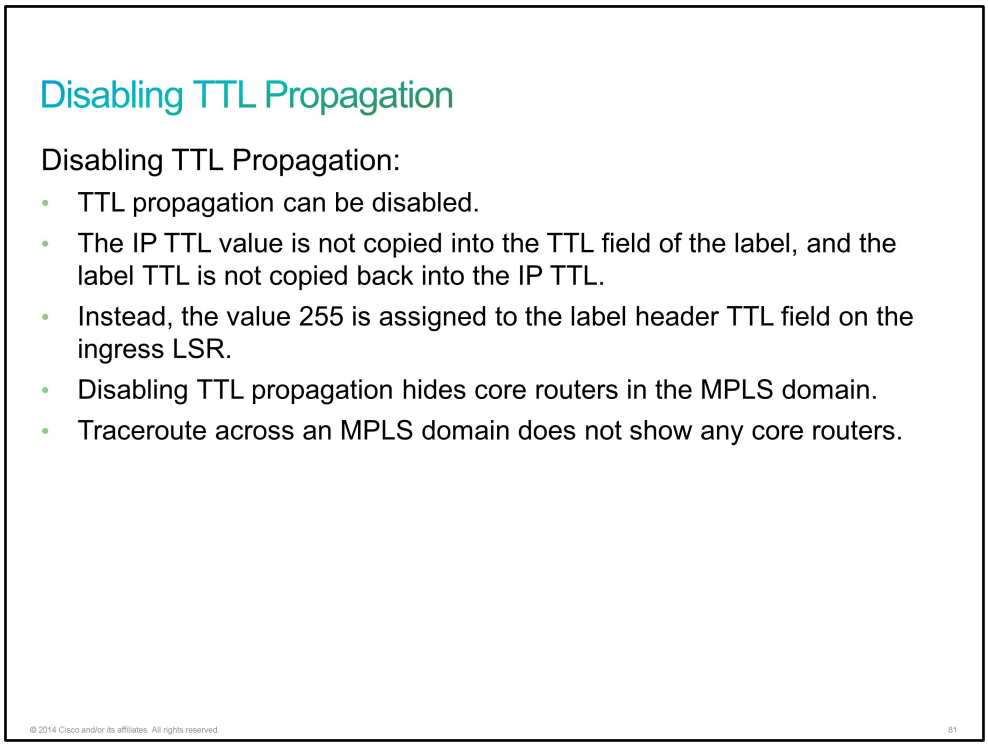

Disabling TTL propagation causes routers to set the value 255 into the TTL field of the label when an IP packet is labeled. The network is still protected against indefinite loops, but it is unlikely that the core routers will ever have to send an ICMP reply to user-originated traceroute packets.

With TTL propagation disabled, the MPLS TTL is calculated independent of the IP TTL, and the IP TTL remains constant for the length of the LSP. Because the MPLS TTL is unlikely to drop from 255 to 0, none of the LSP router hops will trigger an ICMP TTL exceeded message, and consequently these router hops will not be recorded in the traceroute output.

https://t.me/learningnets

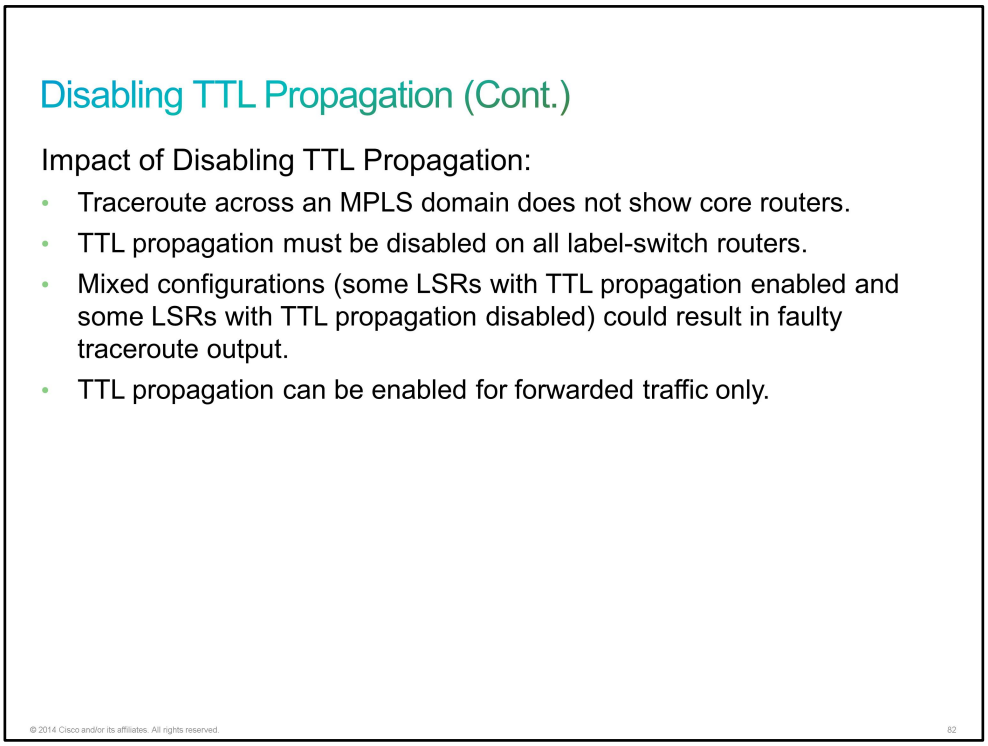

Cisco routers have TTL propagation enabled by default.

If TTL propagation is disabled, it must be disabled on all routers in an MPLS domain to prevent unexpected behavior.

TTL can be optionally disabled for forwarded traffic only, allowing administrators to use traceroute from the routers to troubleshoot problems in the network.

https://t.me/learningnets

Steady State Operation

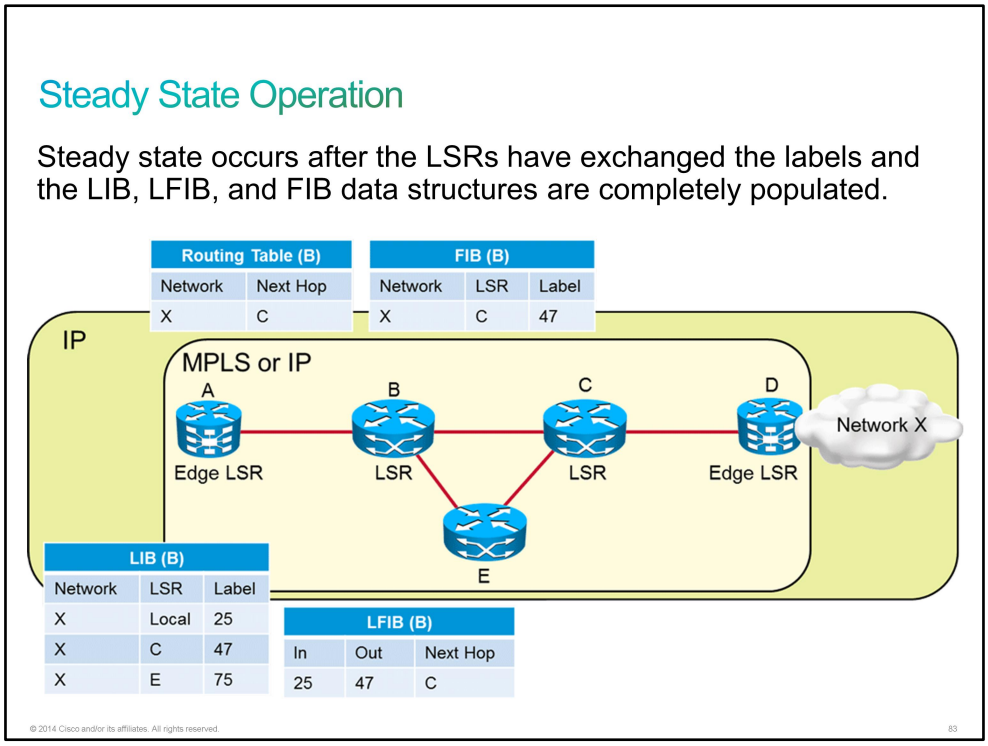

MPLS is fully functional when the IGP and LDP have populated all the tables:

- Main IP routing (RIB) table
- LIB table
- FIB table
- LFIB table

Although it takes longer for LDP to exchange labels (compared with an IGP), a network can use the FIB table in the meantime; therefore, there is no routing downtime while LDP exchanges labels between adjacent LSRs.

https://t.me/learningnets

Link Failure MPLS Convergence Process

The figure illustrates how a link failure is managed in an MPLS domain.

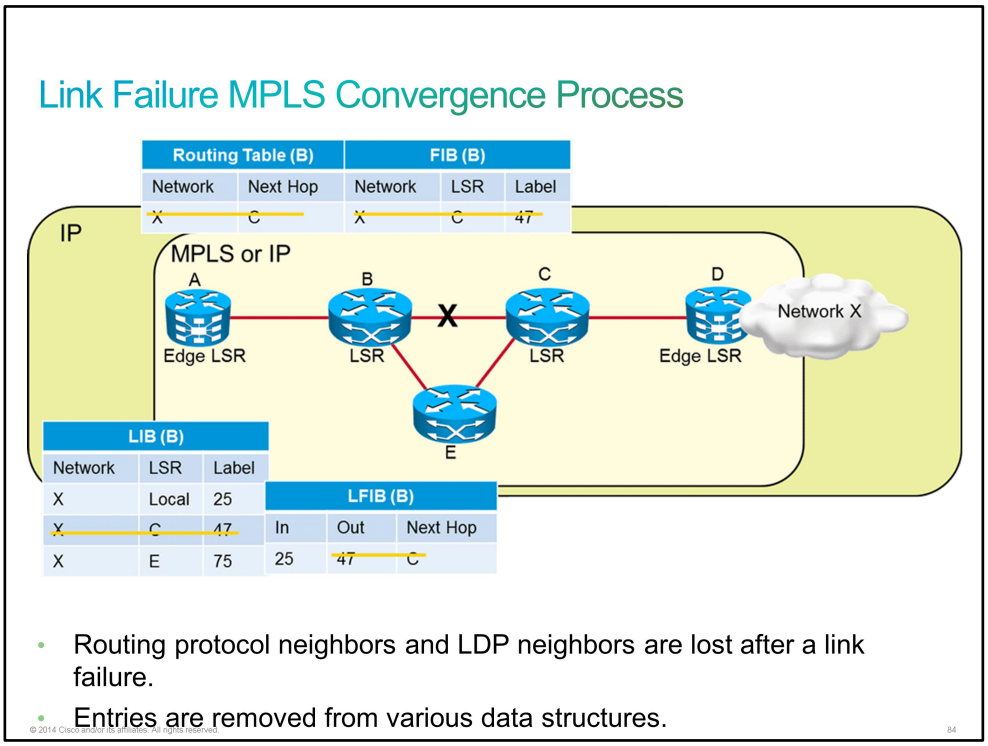

Link failure is managed the following way in an MPLS domain:

- The overall convergence fully depends on the convergence of the IGP used in the MPLS domain.
- The link between router B and router C goes down.
- Entries regarding router C are removed from the LIB, LFIB, FIB, and RIB (routing table).
- When router B determines that router E should be used to reach network X, the label learned from router E can be used to label-switch the packets.

LDP stores all labels in the LIB table, even if the labels are not used, because the IGP has decided to use another path.

This label storage is shown in the figure, where two next-hop labels were available in the LIB table on router B. The following is the label status of router B just before MPLS label convergence:

- Label 47 was learned from router C and is currently unavailable; therefore, because of the failure, label 47 must be removed from the LIB table.
- Label 75 was learned from router E, and can now be used at the moment that the IGP decides that router E is the next hop for network X.

https://t.me/learningnets

Routing Protocol Convergence

The figure illustrates how two entries are removed, one from the LIB table and one from the LFIB table, when the link between routers B and C fails.

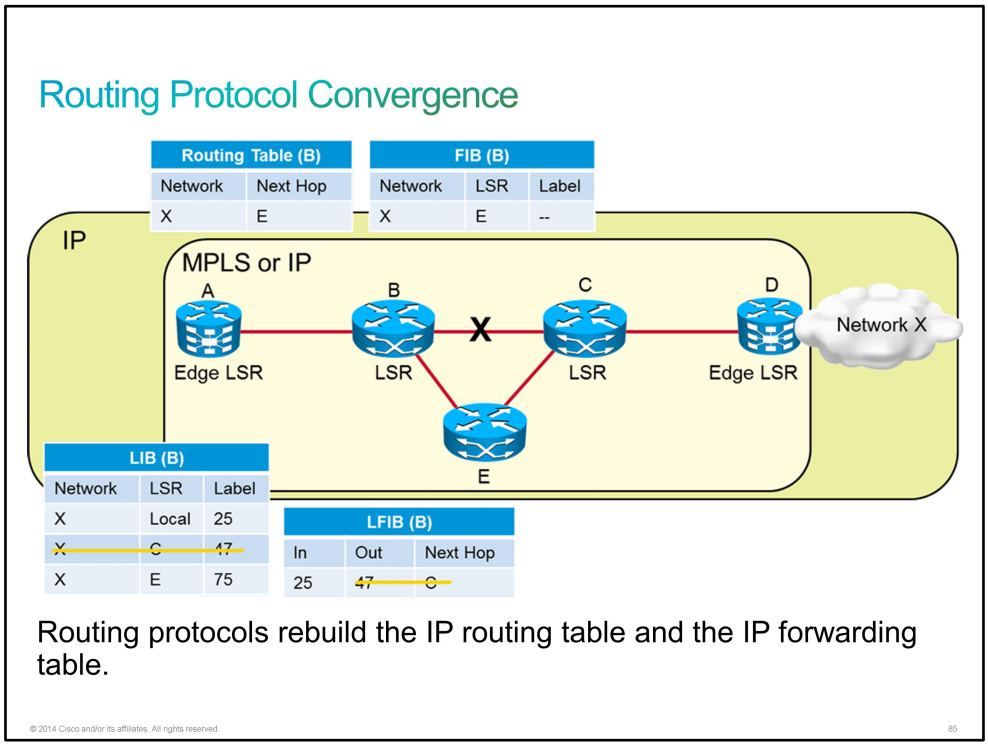

Routing protocol convergence can be described as follows:

- When the IGP determined that the next hop was no longer reachable, Router B removed the entry from the FIB table.
- Router B removed the entry from the LIB table and the LFIB table because LDP has determined that router C is no longer reachable.

https://t.me/learningnets

MPLS Convergence

After the IGP determines that there is another path available, a new entry is created in the FIB table.

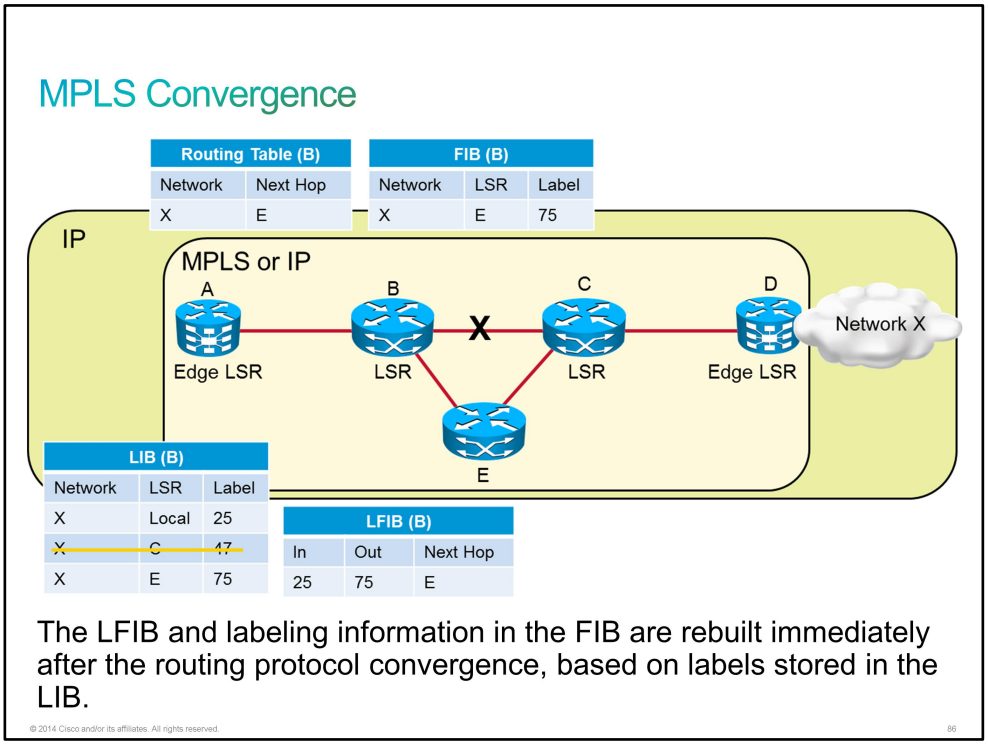

The new entry in the FIB table points toward router E, and there is already a label available for network X via router E in the LIB table. This information is then used in the FIB table and in the LFIB table to reroute the LSP tunnel via router E.

https://t.me/learningnets

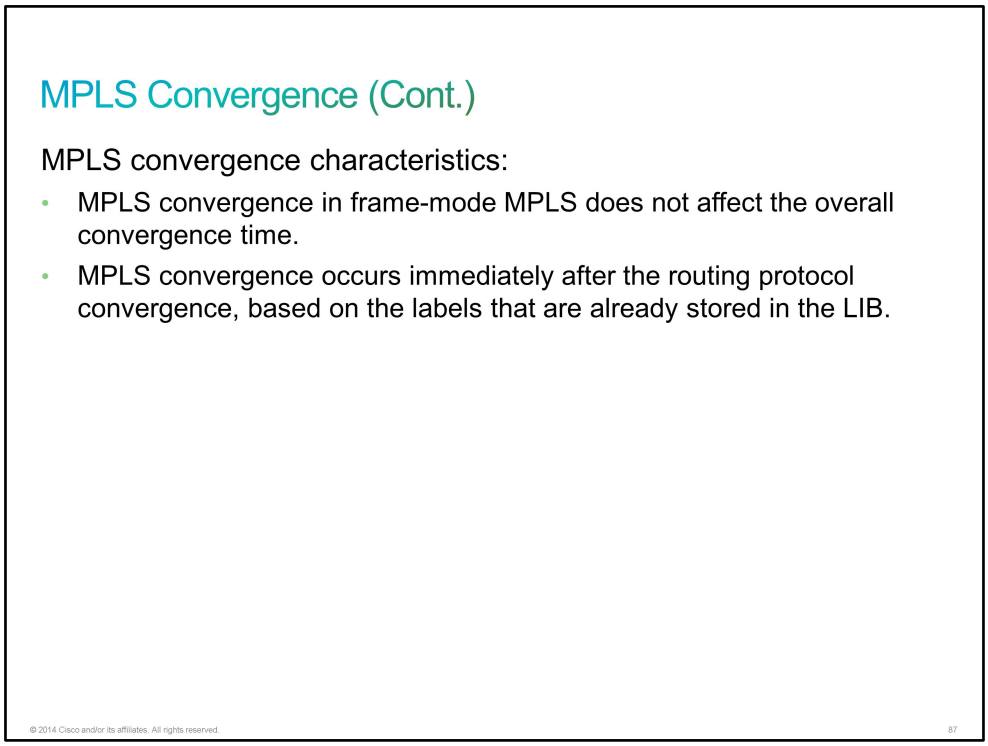

The overall convergence in an MPLS network is not affected by LDP convergence when there is a link failure. Frame-mode MPLS uses liberal label retention mode, which enables routers to store all received labels, even if the labels are not being used. These labels can be used, after the network convergence, to enable immediate establishment of an alternative LSP tunnel

MPLS uses a 32-bit label field that is inserted between Layer 2 and Layer 3 headers (frame-mode MPLS). In frame-mode MPLS, routers that are running MPLS exchange labeled IP packets as well as unlabeled IP packets (PHP) with one another, in an MPLS domain.

https://t.me/learningnets

Link Recovery MPLS Convergence Process

The figure illustrates how a link recovery is managed in an MPLS domain.

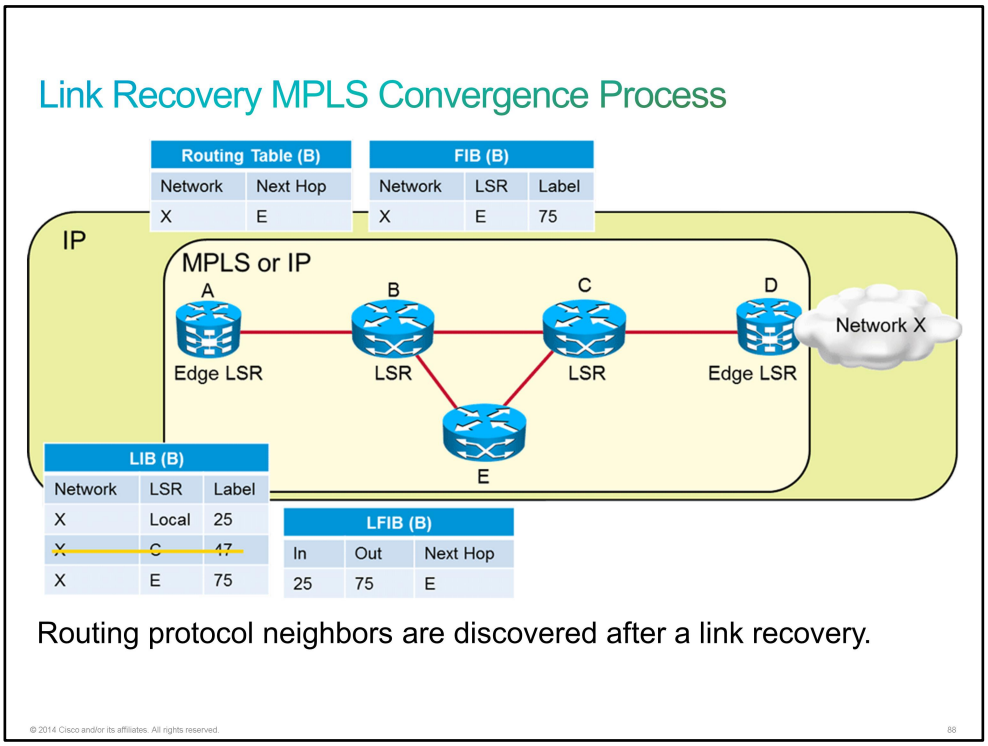

The figure illustrates the state of the B tables of the router at the time the link between routers B and C becomes available again, but before the network reconverges.

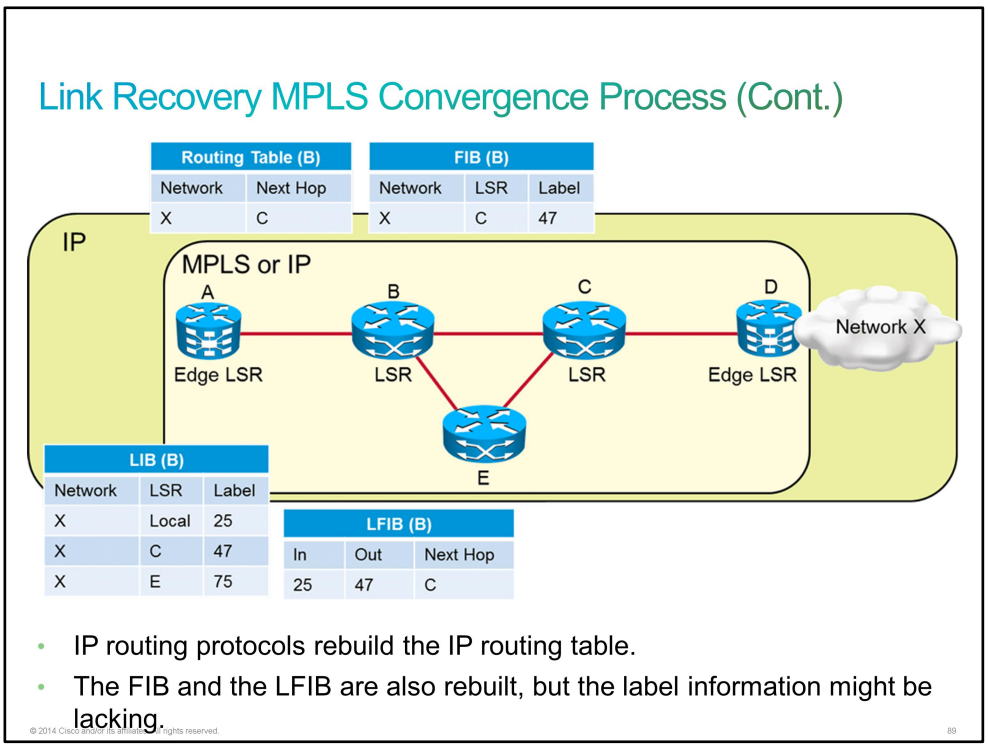

https://t.me/learningnets

The IGP determines that the link between routers B and C is available again, and changes the next-hop address for network X to point to router C. However, when router B also tries to set the next-hop label for network X, it has to wait for the LDP session between routers B and C to be re-established.

A pop action is used in the LFIB table on router B while the LDP establishes the session between routers B and C. This process adds to the overall convergence time in an MPLS domain. The downtime for network X is not influenced by LDP convergence, because normal IP forwarding is used until the new next-hop label is available.

As shown in the figure, after the LDP session between routers B and C is reestablished, router B will update its tables with an outgoing label of 47, with router C as the next-hop for reaching network X.

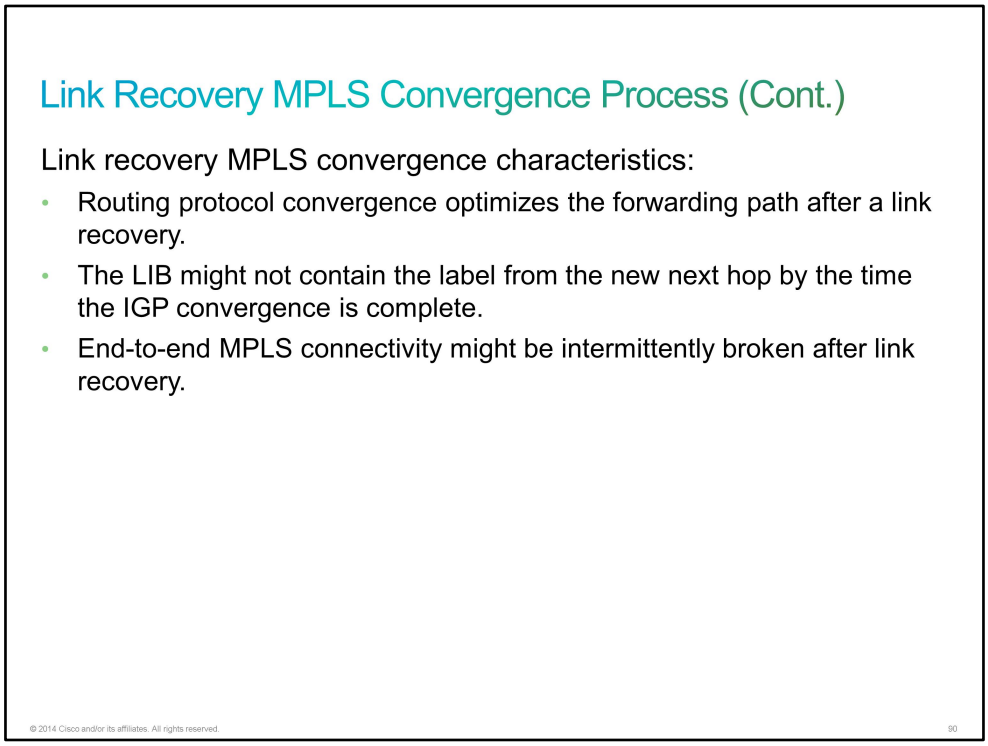

Link recovery requires that an LDP session be established (or reestablished), which adds to the convergence time of LDP. Networks may be temporarily unreachable because of the convergence limitations of routing protocols. Cisco MPLS TE can be used to prevent long downtime when a link fails or when it is recovering.

https://t.me/learningnets

IP Switching Mechanisms

The first and the oldest switching mechanism that is available in Cisco routers is process switching. Because process switching must find a destination in the routing table (possibly a recursive lookup) and construct a new Layer 2 frame header for every packet, it is very slow and is normally not used.

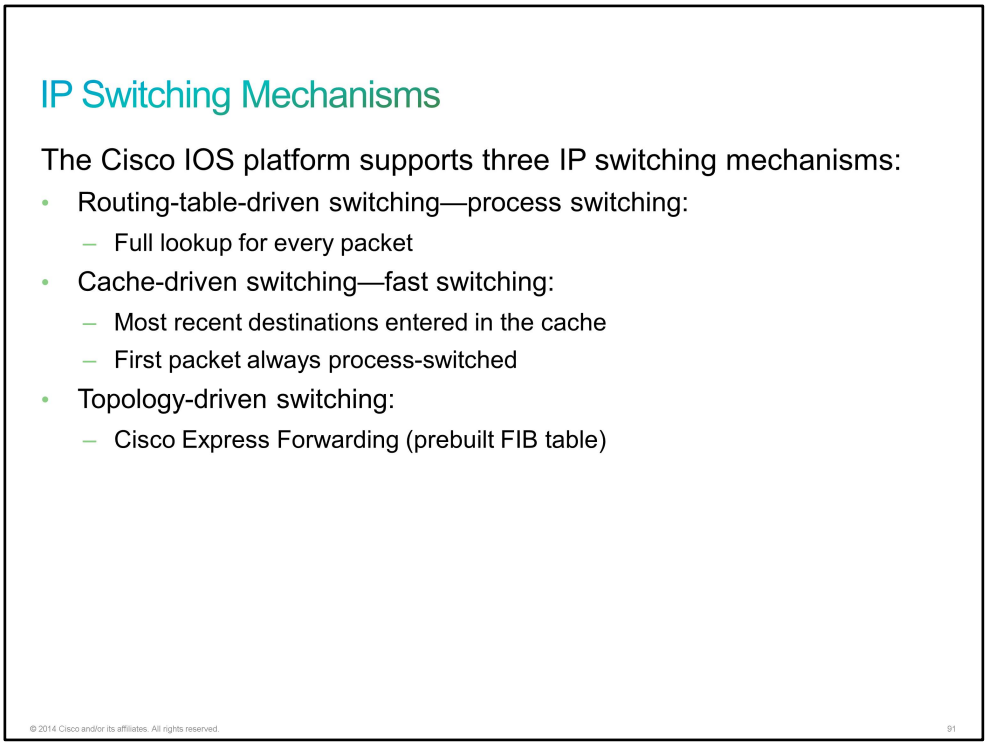

To overcome the slow performance of process switching, Cisco IOS platforms support several switching mechanisms that use a cache to store the most recently used destinations. The cache uses a faster searching mechanism, and it stores the entire Layer 2 frame header to improve the encapsulation performance. The first packet whose destination is not found in the fast-switching cache is process-switched, and an entry is created in the cache. The subsequent packets are switched in the interrupt code using the cache to improve performance.

The latest and preferred Cisco IOS platform-switching mechanism is Cisco Express Forwarding, that incorporates the best of the previous switching mechanisms. Cisco Express Forwarding supports per-packet load balancing (previously supported only by process switching), per-source or per-destination load balancing, fast destination lookup, and many other features not supported by other switching mechanisms.

The Cisco Express Forwarding cache, or FIB table, is essentially a replacement for the standard routing table.

https://t.me/learningnets

Standard IP Switching Example

There is a specific sequence of events that occurs when process switching and fast switching are used for destinations that are learned through BGP.

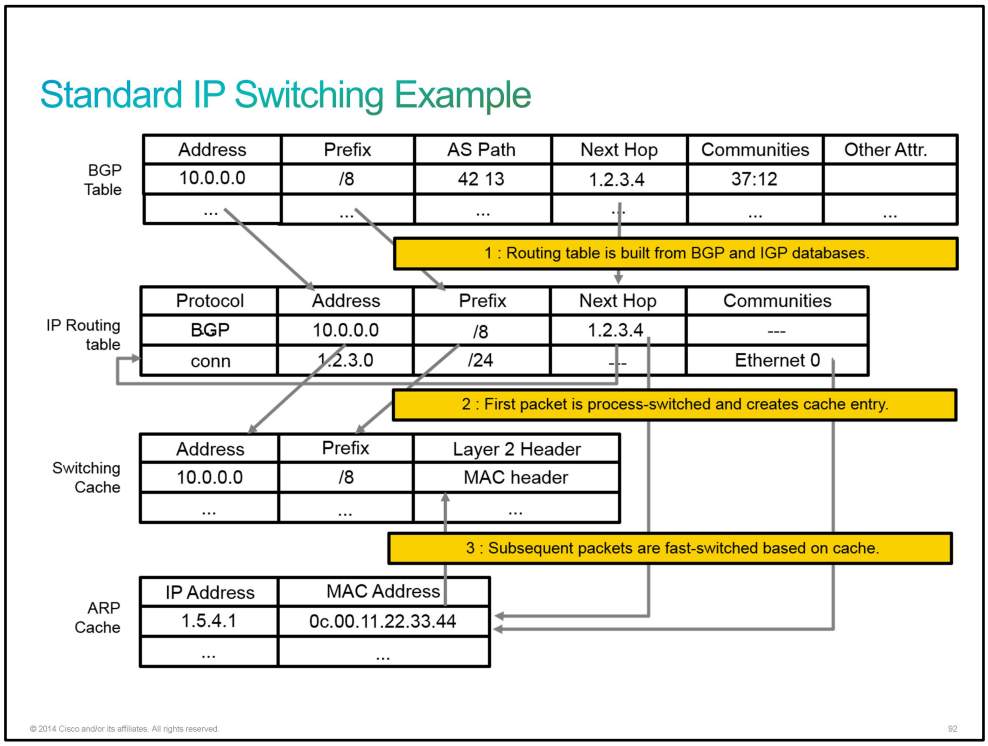

The figure illustrates this process. Here is a description of the sequence of events:

- When a BGP update is received and processed, an entry is created in the routing table.
- When the first packet arrives for this destination, the router tries to find the destination in the fast-switching cache. Because the destination is not in the fast-switching cache, process switching has to switch the packet when the process is run. The process performs a recursive lookup to find the outgoing interface. The process switching may possibly trigger an ARP request or may find the Layer 2 address in the ARP cache. Finally, it creates an entry in the fast-switching cache.
- All subsequent packets for the same destination are fast-switched, as follows:
  - The switching occurs in the interrupt code (the packet is processed immediately).
  - Fast destination lookup is performed (no recursion).
  - The encapsulation uses a pregenerated Layer 2 header that contains the destination and Layer 2 source (MAC) address. (No ARP request or ARP cache lookup is necessary.)

https://t.me/learningnets

Cisco Express Forwarding Switching Example

Cisco Express Forwarding uses a different architecture from process switching or any other cache-based switching mechanism.

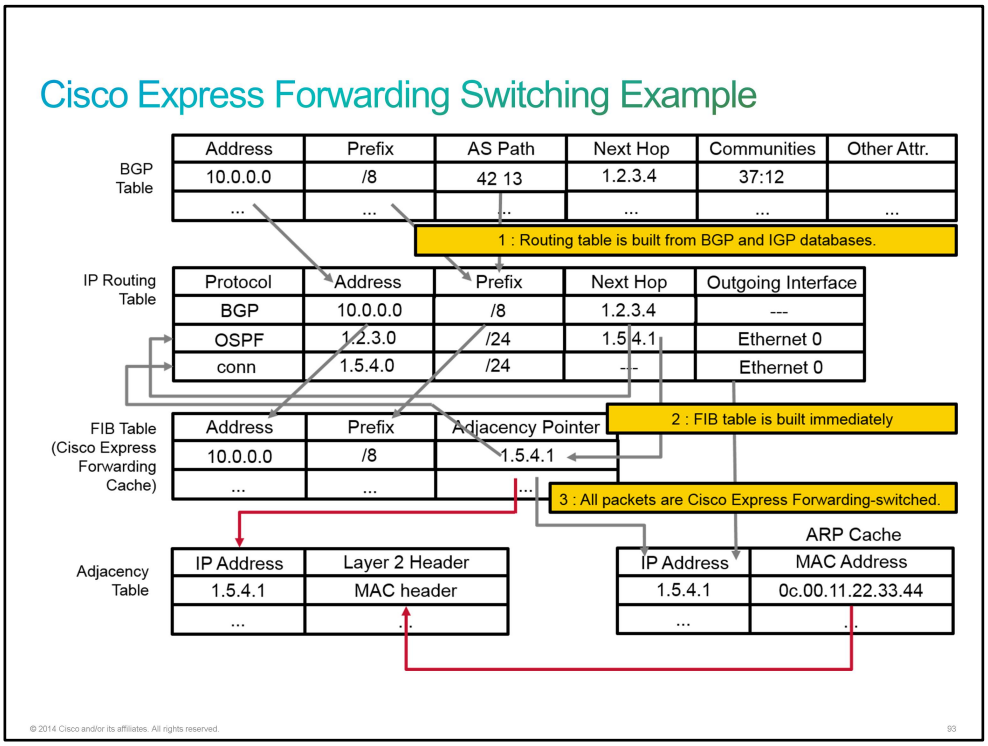

Cisco Express Forwarding uses a complete IP switching table, the FIB table, that holds the same information as the IP routing table. The generation of entries in the FIB table is not packet-triggered but change-triggered. When something changes in the IP routing table, the change is also reflected in the FIB table.

Because the FIB table contains the complete IP switching table, the router can make definitive decisions based on the information in it. Whenever a router receives a packet that should be switched with Cisco Express Forwarding, but the destination is not in the FIB, the packet is dropped.

The FIB table is also different from other fast-switching caches in that, it does not contain information about the outgoing interface and the corresponding Layer 2 header. That information is stored in a separate table, the adjacency table. The adjacency table is similar to a copy of the ARP cache, but instead of holding only the destination MAC address, it holds the Layer 2 header.

Cisco Express Forwarding offers the following benefits:

- **Improved performance**: Cisco Express Forwarding is less CPU-intensive than fast-switching route caching. More CPU processing power can be dedicated to Layer 3 services, such as QoSand encryption.
- **Scalability:** Cisco Express Forwarding offers full switching capacity at each MSC on the Cisco Carrier Routing System routers.

https://t.me/learningnets

- **Resilience:** Cisco Express Forwarding offers an unprecedented level of switching consistency and stability in large dynamic networks. In dynamic networks, fast-switched cache entries are frequently invalidated due to routing changes. These changes can cause traffic to be process-switched, using the routing table, rather than fast-switched using the route cache. Because the FIB lookup table contains all known routes that exist in the routing table, it eliminates route cache maintenance and the fast-switch or process-switch forwarding scenario. Cisco Express Forwarding can switch traffic more efficiently than typical demand caching schemes.

https://t.me/learningnets

Cisco Express Forwarding in IOS XE Software

Cisco Express Forwarding is an advanced Layer 3 IP switching technology. It optimizes network performance and scalability for all kinds of networks: those that carry small amounts of traffic and those that carry large amounts of traffic in complex patterns, such as the Internet, and networks characterized by intensive web-based applications or interactive sessions.

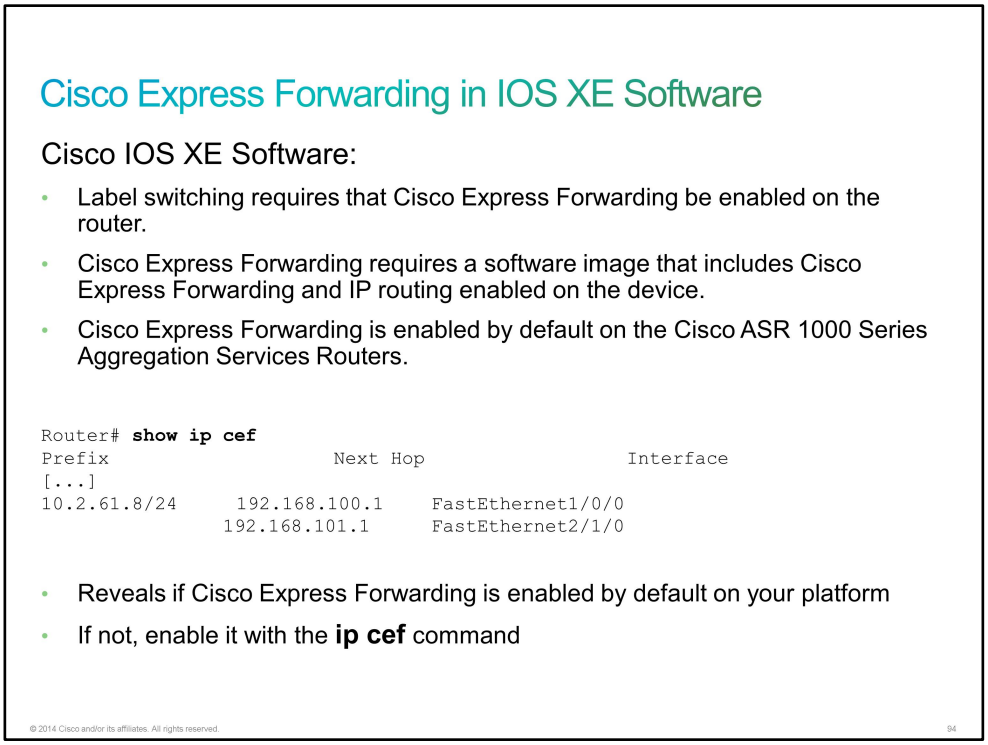

Cisco Express Forwarding requires a software image that includes Cisco Express Forwarding and IP routing enabled on the device.

Cisco Express Forwarding is enabled by default on the Cisco ASR 1000 Series Aggregation Services Routers.

To find out if Cisco Express Forwarding is enabled by default on your platform, enter the **show ip cef** command. If Cisco Express Forwarding is enabled, you receive output that looks like the following:

```
Router# show ip cef
Prefix              Next Hop            Interface
[...]
10.2.61.8/24        192.168.100.1       FastEthernet1/0/0
                    192.168.101.1       FastEthernet2/1/0
[...]
```

If Cisco Express Forwarding is not enabled on your platform, the output for the **show ip cef** command looks like the following:

```
Router# show ip cef
%CEF not running
```

If Cisco Express Forwarding is not enabled on your platform, use the **ip cef** command to enable Cisco Express Forwarding or the **ip cef distributed** command to enable distributed Cisco Express Forwarding.

https://t.me/learningnets

Cisco Express Forwarding in IOS XR

The prerequisite required to implement MPLS forwarding is the installed composite mini-image and the MPLS package, or a full composite image.

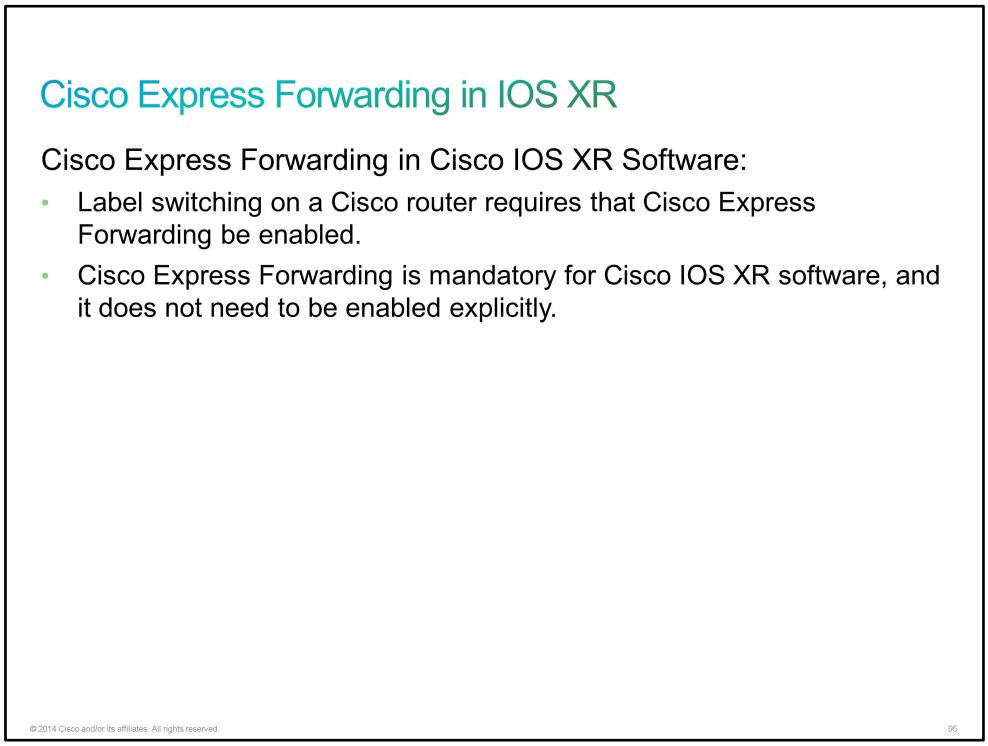

Label switching on a Cisco router requires that Cisco Express Forwarding be enabled. Cisco Express Forwarding is mandatory for Cisco IOS XR Software, and it does not need to be enabled explicitly.

In Cisco IOS XR Software, Cisco Express Forwarding always operates in Cisco Express Forwarding mode with two distinct components: a FIB database and an adjacency table, and a protocol-independent AIB.

Cisco Express Forwarding is a primary IP packet-forwarding database for Cisco IOS XR Software. Cisco Express Forwarding is responsible for these functions:

- Software switching path
- Maintaining forwarding table and adjacency tables (which are maintained by the AIB) for software and hardware forwarding engines

These Cisco Express Forwarding tables are maintained in Cisco IOS XR Software:

- IPv4 Cisco Express Forwarding database
- IPv6 Cisco Express Forwarding database
- MPLS LFD
- MFD

https://t.me/learningnets

Monitoring IPv4 Cisco Express Forwarding

show commands can be used to monitor Cisco Express Forwarding operations.

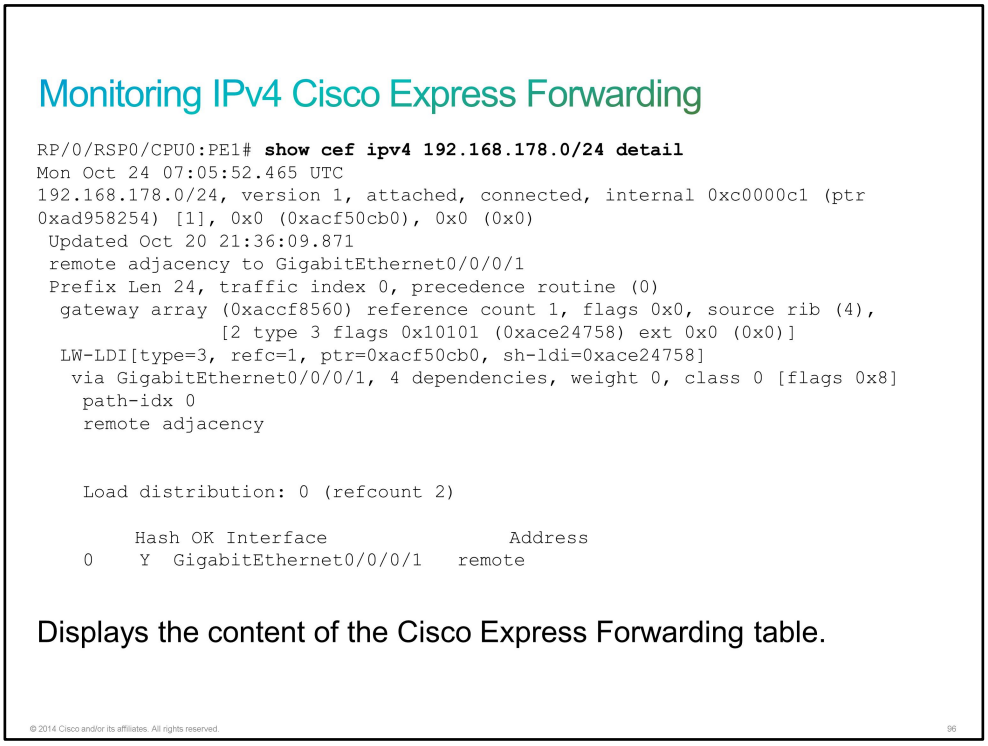

To display the IPv4 Cisco Express Forwarding table, use the **show cef ipv4** command in EXEC mode on Cisco IOS XR Software.

```
RP/0/RSP0/CPU0:PE7# show cef ipv4
Mon Oct 24 07:08:01.177 UTC
Prefix              Next Hop            Interface
0.0.0.0/0           drop                default handler
0.0.0.0/32          broadcast
10.7.1.1/32         receive             Loopback0
10.8.1.0/24         attached            GigabitEthernet0/0/0/1
192.168.178.0/24    attached            GigabitEthernet0/0/0/1
192.168.178.0/32    broadcast           GigabitEthernet0/0/0/1
192.168.178.70/32   receive             GigabitEthernet0/0/0/1
192.168.178.255/32  broadcast           GigabitEthernet0/0/0/1
224.0.0.0/4         point2point
224.0.0.0/24        receive
255.255.255.255/32  broadcast
```

To display unresolved entries in the FIB table or to display a summary of the FIB, use this form of the **show cef ipv4** EXEC command: **show cef ipv4 [unresolved | summary]**.

To display specific entries in the FIB table based on IP address information, use this form of the **show cef ipv4** command in EXEC mode: **show cef ipv4** *[network [mask* [**longer-prefix**]]] [**detail**].

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- The TTL functionality in MPLS is equivalent to that of traditional IP forwarding.
- TTL propagation can be disabled to hide the core routers from the end users.
- Although it takes longer for LDP to exchange labels (compared with an IGP), a network can use the FIB table in the meantime.
- Link recovery requires that an LDP session be established (or reestablished), which adds to the convergence time of LDP.
- The Cisco IOS platform supports three IP switching mechanisms: process switching, fast switching, and Cisco Express Forwarding.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- In standard IP switching, the first packet that arrives is process-switched and all subsequent packets are fast-switched.
- In Cisco Express Forwarding switching, the FIB table is built in advance before a packet for a destination is received.
- On IOS and IOS XE, Cisco Express Forwarding is required for MPLS.
- Cisco Express Forwarding is enabled on Cisco IOS XR and cannot be disabled.
- To monitor Cisco Express Forwarding, you can use various **show** commands.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 7: Implementing MPLS in the Service Provider Core

Overview

This lesson describes how to configure MPLS on various Cisco platforms. The essential configuration tasks and commands, including their correct syntax, are discussed. Also addressed are advanced configurations such as the label-switching MTU, IP TTL propagation, LDP session protection, LDP graceful restart, LDP IGP synchronization, LDP autoconfiguration, and conditional label distribution. The lesson also describes the procedures for monitoring MPLS by using syntax and parameter descriptions, interfaces, neighbor nodes, LIB, and LFIB tables. Also outlined are the usage guidelines for the commands.

The lesson concludes with a look at some of the common issues that arise in MPLS networks. For each issue discussed, there is a recommended troubleshooting procedure to resolve the issue.

Upon completing this lesson, you will be able to:

- Describe the MPLS configuration difference in Cisco IOS XR versus Cisco IOS or IOS XE
- Describe the mandatory and optional MPLS configuration tasks
- Describe a basic MPLS configuration
- Describe the MTU requirements on a label-switching router interface
- Explain the configuration that is used to increase the MPLS MTU size on a label-switching router interface
- Explain IP TTL propagation
- Explain the configuration that is used to disable IP TTL propagation
- Explain LDP Session Protection configuration
- Explain LDP graceful restart and NSR configuration

https://t.me/learningnets

- Explain LDP IGP synchronization configuration
- Explain how to enable LDP autoconfiguration
- Explain label advertisement control configuration and verification
- Explain label acceptance control configuration
- Describe the show commands that are used to monitor MPLS operations
- Display MPLS LDP parameters
- Describe how to display MPLS configured interfaces
- Describe how to display the status of the LDP discovery process
- Describe the **show** commands that are used to monitor LDP operations
- Describe how to display the status of the LDP session
- Describe how to display the content of the LIB table
- Describe the **show** commands used to monitor label-switching
- Describe how to display the contents of the LFIB table
- Describe how to display the contents of the FIB table
- Describe the MPLS and LDP debug commands
- Describe the classic ping and traceroute operations
- Describe the MPLS ping and traceroute operations
- Describe the MPLS ping
- Describe the MPLS ping and traceroute operations
- Describe how to troubleshoot common MPLS issues
- Describe how to troubleshoot LDP session startup issues
- Describe how to troubleshoot label distribution issues
- Describe how to troubleshoot intermittent MPLS failures after interface failure
- Describe how to troubleshoot packet propagation issues

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS Configuration on Cisco IOS XR versus Cisco IOS or IOS XE

The basic configuration of MPLS is simple, as you only need to enable MPLS on the interface.

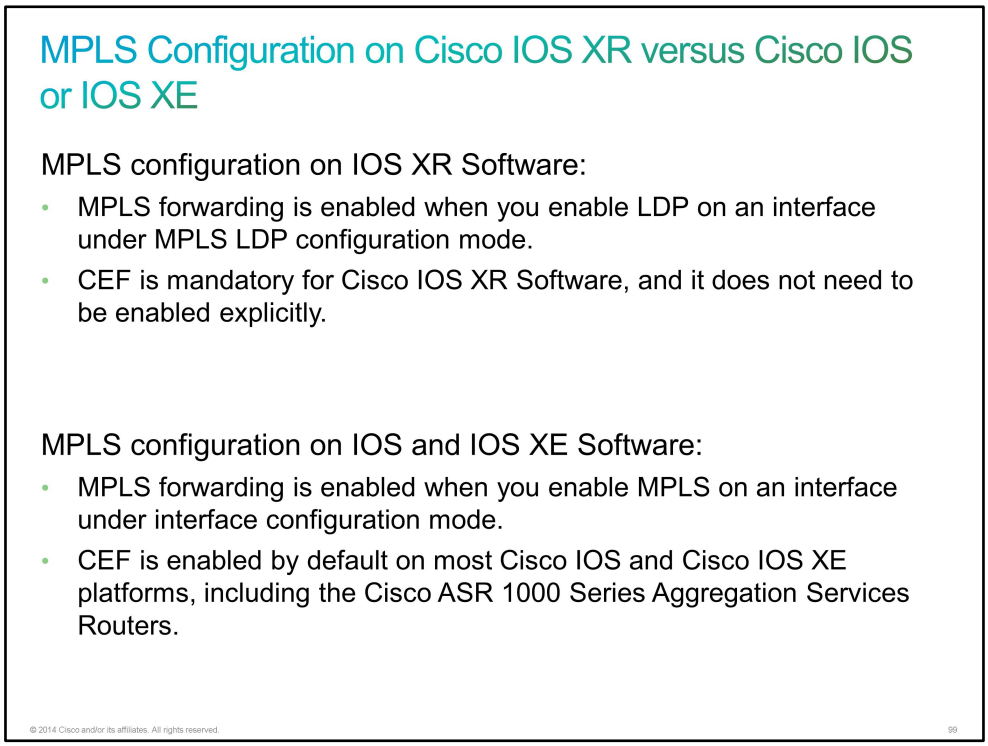

On Cisco IOS XR platforms, MPLS is enabled by enabling LDP on each interface under the MPLS LDP configuration mode. On Cisco IOS and IOS XE platforms, MPLS is enabled on each interface under the interface configuration mode.

https://t.me/learningnets

MPLS Configuration Tasks

To enable MPLS on a router that is running Cisco IOS XR Software, enable LDP on an interface under MPLS LDP configuration mode. Optionally, the maximum size of labeled packets may be changed.

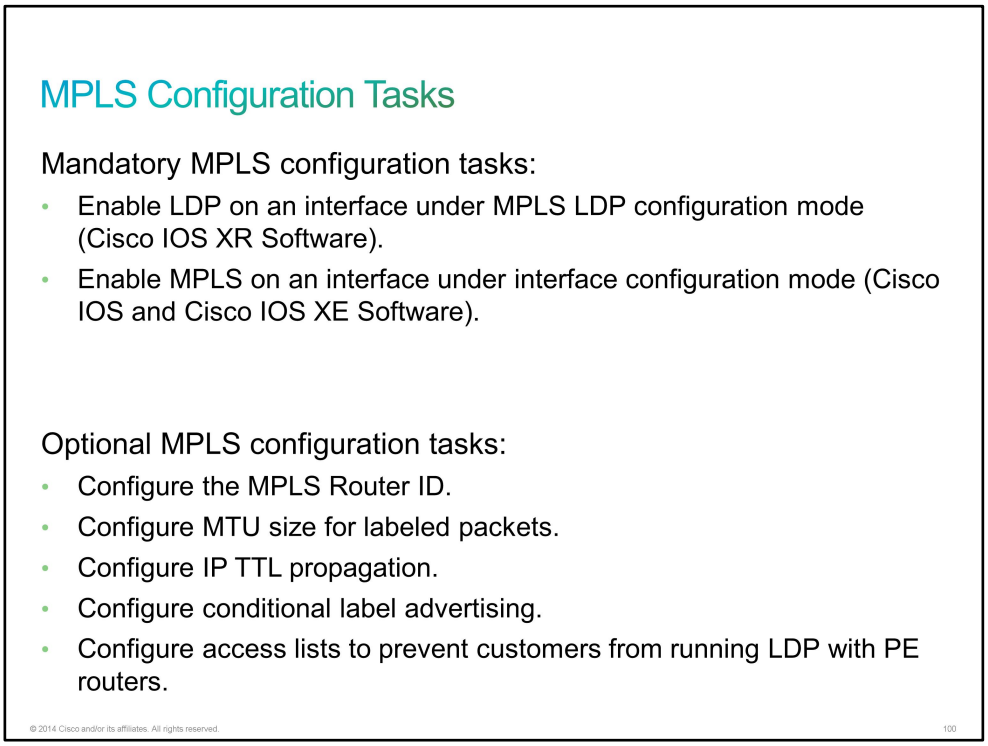

A stable LDP router ID is required at either end of the link to ensure that the link discovery (and session setup) is successful. If you do not manually assign the LDP router ID on the Cisco IOS XR routers, the Cisco IOS XR routers will default to use the global router ID as the LDP router ID. Global router ID configuration is only available on Cisco IOS XR (not available on Cisco IOS and IOS XE Software).

You can override the global **router-id** command in Cisco IOS XR by further configuring a router-id command within a given protocol. However, configuring different router IDs per protocol makes router management more complicated.

By default, the TTL field is copied from the IP header and placed in the MPLS label TTL field when a packet enters an MPLS network. To prevent core routers from responding with (ICMP) TTL exceeded messages, disable TTL propagation. If TTL propagation is disabled, the value in the TTL field of the MPLS label is set to 255.

Note	Ensure that TTL propagation is either enabled in all routers or disabled in all routers. If TTL is enabled in some routers and disabled in others, the result may be that a packet that is leaving the MPLS domain will have a larger TTL value than when it entered.

By default, a router will generate and propagate labels for all networks that it has in the routing table. If label switching is required for only a limited number of networks (for example, only for router loopback addresses), configure conditional label advertising.

https://t.me/learningnets

To prevent customers from running LDP with PE routers, configure access lists that will block the well-known TCP port of an LDP.

https://t.me/learningnets

Basic MPLS Configuration

To enable MPLS on the Cisco IOS XR router, first enter MPLS LDP configuration mode using the **mpls ldp** command. Then specify the interfaces that should be enabled for MPLS by using the **interface** command.

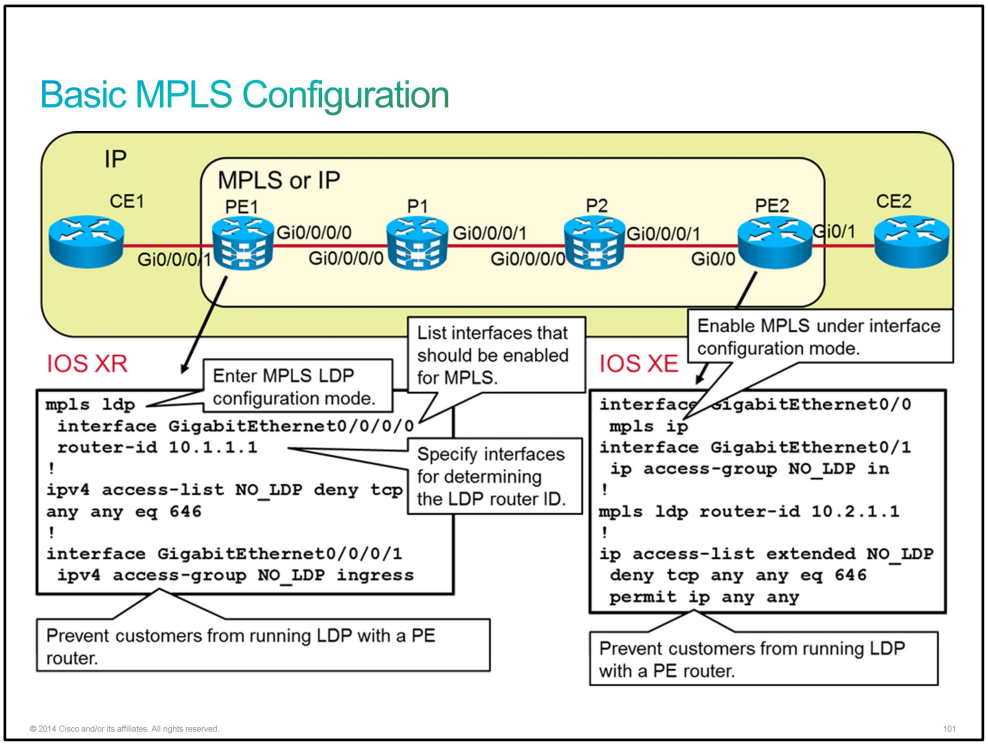

In the example, the MPLS for router PE1 is enabled on the GigabitEthernet0/0/0/0 interface. The configuration includes an ACL that denies any attempt to establish an LDP session from an interface that is not enabled for MPLS. In the example shown in the figure, router PE1 has the NO_LDP access list applied to interface GigabitEthernet0/0/0/1, which is not enabled for MPLS.

Note	Enable MPLS on all core interfaces in your network. On routers P1 and P2, both interfaces GigabitEthernet0/0/0/0 and GigabitEthernet0/0/0/1 should be enabled for MPLS.

To enable MPLS on Cisco IOS and IOS XE routers, first enter the interface configuration mode for a desired interface. Then enable MPLS, using the **mpls ip** command. In the example, MPLS is enabled for router PE2 on the GigabitEthernet0/0 interface. The configuration includes an ACL that denies any attempt to establish an LDP session from an interface that is not enabled for MPLS. In the example, router PE2 has the NO_LDP access list applied to the interface GigabitEthernet0/1, that is not enabled for MPLS.

A stable router ID is required at either end of the link to ensure that the link discovery (and session setup) is successful. In the example, routers PE1 and PE2 have the LDP router ID set to the IP address of interface loopback 0.

https://t.me/learningnets

MTU Requirements

Label switching increases the maximum MTU requirements on an interface because of the additional label header.

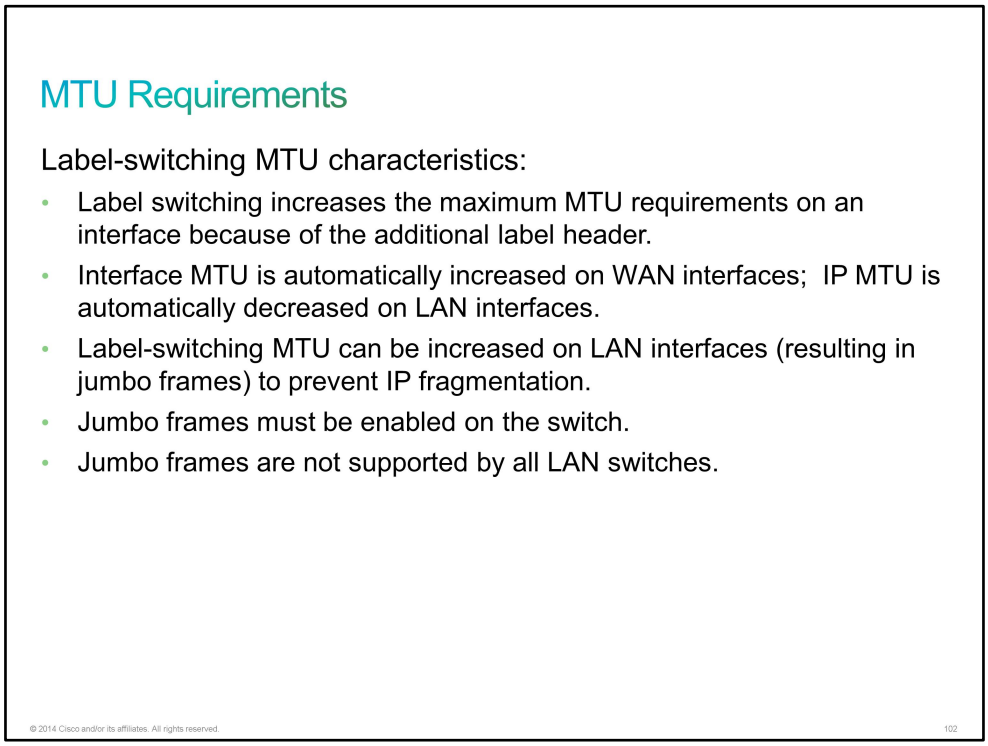

There are three different MTU values:

- The interface MTU is used to determine the maximum size of any packets that can be sent on the interface. In Cisco IOS XR Software, this is the Layer 2 MTU. In IOS and IOS XE Software, this is the Layer 3 payload size.
- The IP MTU is used to determine whether a nonlabeled IP packet that is forwarded through an interface has to be fragmented (the IP MTU has no impact on labeled IP packets).
- The MPLS MTU determines the maximum size of a labeled IP packet (MPLS shim header + IP payload size). If the overall length of the labeled packet (including the shim header) is greater than the MPLS MTU, the packet is fragmented. The default MPLS MTU is the MTU that is configured for the interface.

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 while the IP MTU is automatically decreased on LAN interfaces.

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. This problem will typically occur on LAN switches, where it is more likely that a device does not support oversized packets (also called jumbo frames, giants, or baby giants). Some devices support jumbo frames, and some need to be configured to support them.

The interface MTU size (and therefore also the MPLS MTU) is increased automatically on WAN interfaces but the MPLS MTU needs to be increased manually on LAN interfaces.

https://t.me/learningnets

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.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS MTU Configuration

The figure shows a label-switching MTU configuration on LAN interfaces for routers P1 and PE2. MPLS MTU is increased to 1512 on the Ethernet interfaces of router P1 to support 1500-B IP packets and MPLS stacks up to three levels deep (3 times 4-B label).

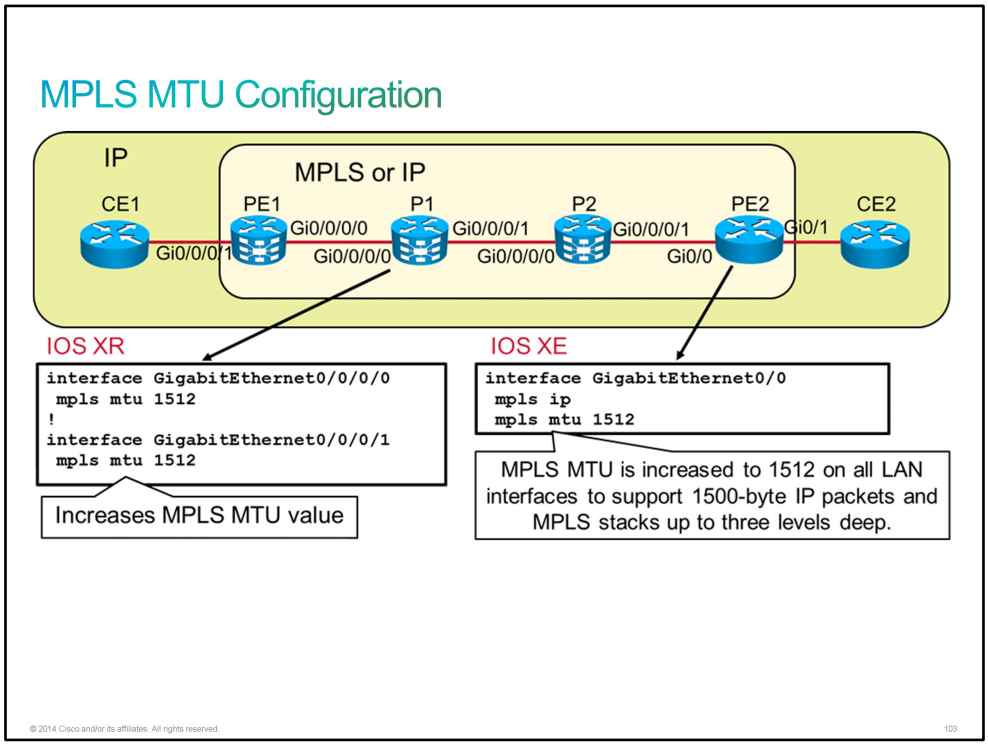

To configure the maximum packet size or MTU size on an MPLS interface (for Cisco IOS XR Software and Cisco IOS and IOS XE Software), use the **mpls mtu** command in interface configuration mode. To disable this feature, use the **no** form of this command.

https://t.me/learningnets

IP TTL Propagation

By default, IP TTL is copied into the MPLS label at label imposition, and the MPLS label TTL is copied (back) into the IP TTL at label removal.

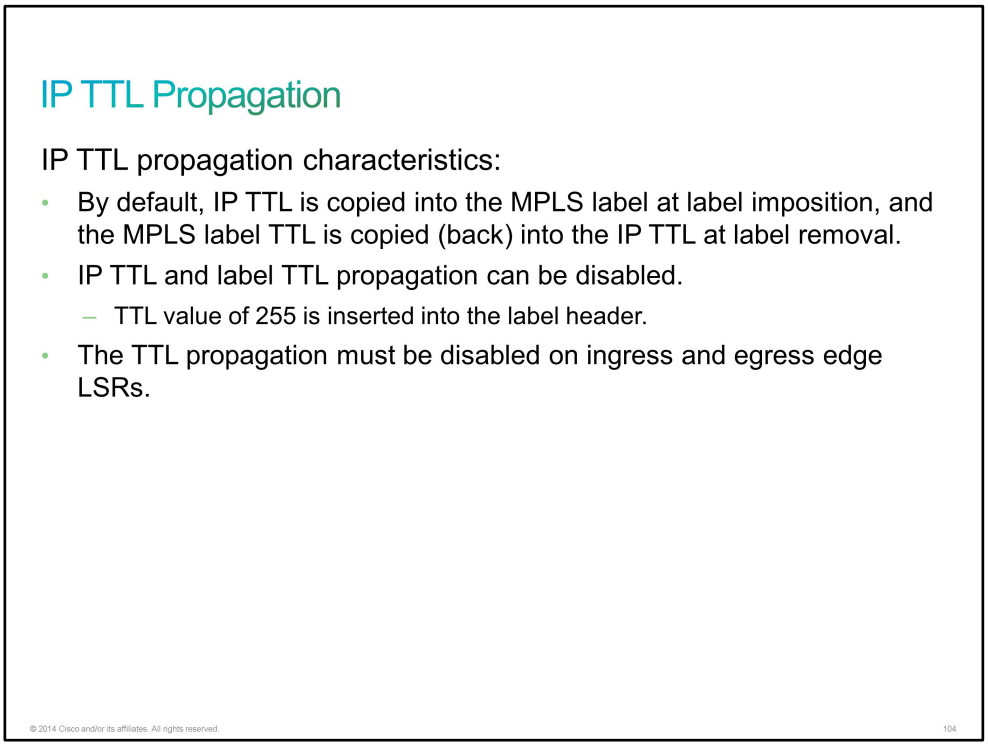

IP TTL and label TTL propagation can be disabled if it is desired to hide the core routers from the traceroute output; a TTL value of 255 is inserted in the label header. The TTL propagation must be disabled, at least on ingress and egress edge LSRs, but it is advisable that all routers have TTL propagation enabled, or all disabled.

https://t.me/learningnets

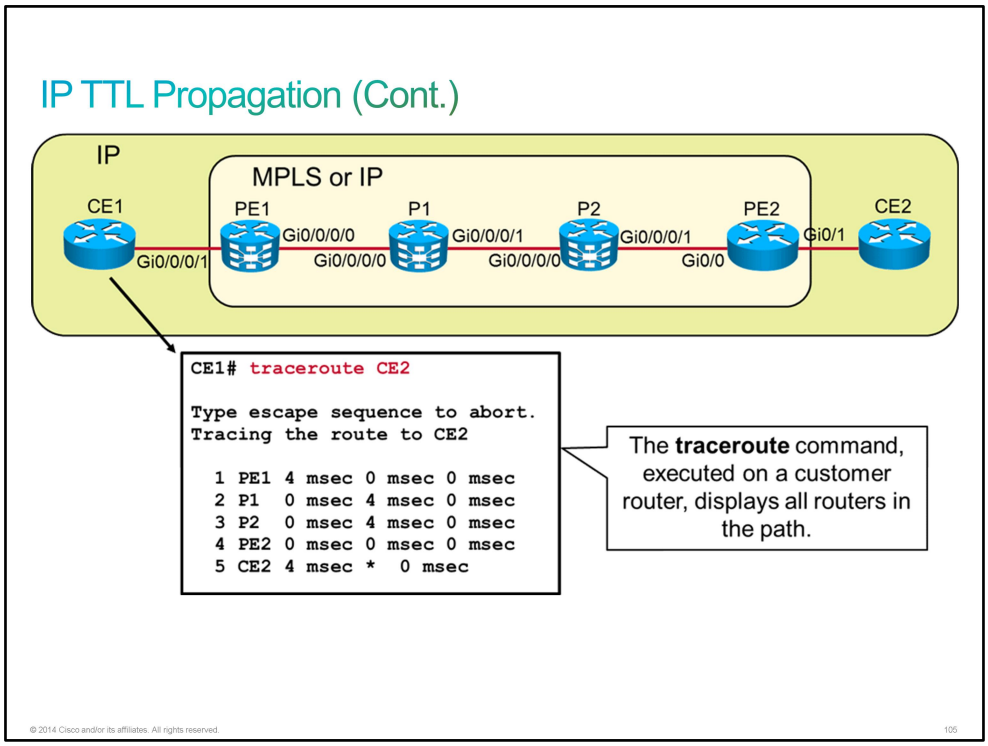

The figure illustrates typical traceroute behavior in an MPLS network. Because the label header of a labeled packet carries the TTL value from the original IP packet, the routers in the path can drop packets when the TTL is exceeded. Traceroute will therefore show all the routers in the path. This is the default behavior.

In the example, router CE1 is executing a traceroute command that results in this behavior. The steps for this process are as follows:

1. The first packet is an IP packet with TTL = 1. Router PE1 decreases the TTL and drops the packet because it reaches 0. An ICMP TTL exceeded message is sent to the source.
2. The second packet sent is an IP packet with TTL = 2. Router PE1 decreases the TTL, labels the packet (the TTL from the IP header is copied into the MPLS label TTL field), and forwards the packet to router P1.
3. Router P1 decreases the MPLS TTL value, drops the packet, and sends an ICMP TTL exceeded message to the source.
4. Processing for the third packet is similar, with TTL = 3. Router P2 sends an ICMP TTL exceeded message to the source.
5. The fourth packet (TTL = 4) experiences processing that is similar to the previous packets, except that router PE2 drops the packet based on the TTL in the IP header. Router P2, because of PHP, had previously removed the label, and the TTL was copied back to the IP header.

The fifth packet (TTL = 5) reaches the final destination, where the TTL of the IP packet is examined.

https://t.me/learningnets

Disabling IP TTL Propagation

If TTL propagation is disabled, the TTL value is not copied into the label header. Instead, the label TTL field is set to 255. The probable result is that the TTL field in the label header will not decrease to 0 for any router inside the MPLS domain (unless there is a forwarding loop inside the MPLS network).

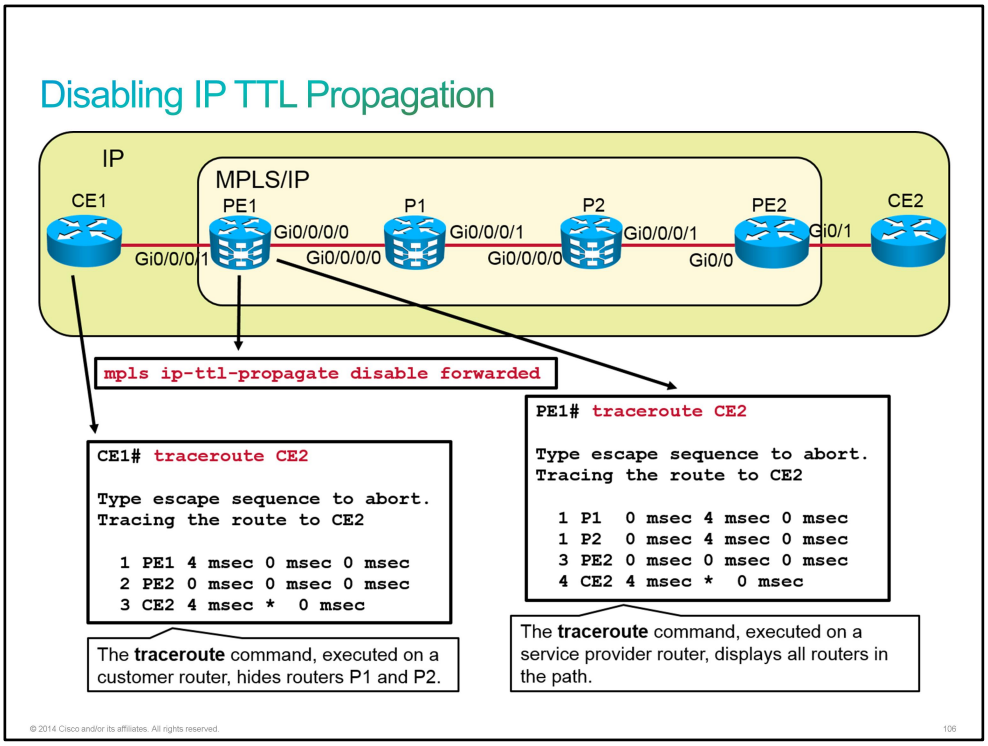

If the **traceroute** command is used, ICMP replies are received only from those routers that see the real TTL that is stored in the IP header.

Typically, a service provider likes to hide the backbone network from outside users, but allows inside traceroute to work for easier troubleshooting of the network.

This goal can be achieved by disabling TTL propagation for forwarded packets only, as described here:

- If a packet originates in the router, the real TTL value is copied into the label TTL.
- If the packet is received through an interface, the TTL field in a label is assigned a value of 255.

The result is that someone using traceroute on a provider router will see all of the backbone routers. Customers will see only edge routers.

Use the **mpls ip-ttl-propagate** (Cisco IOS XR Software) or **mpls ip propagate-ttl** (Cisco IOS XE Software) global configuration command to control generation of the TTL field in the label when the label is first added to the IP packet. By default, this command is enabled, which means that the TTL field is copied from the IP header and inserted into the MPLS label. This aspect allows a **traceroute** command to show all of the hops in the network.

To use a fixed TTL value (255) for the first label of the IP packet, use the **no** form of the **mpls ip propagate-ttl** command on Cisco IOS XE Software. To use a fixed TTL value (255) for the first label of the IP packet, use the **mpls ip-ttl-propagate disable** command on Cisco IOS XR Software. This action hides the structure of the MPLS network from a **traceroute** command.

https://t.me/learningnets

Specify the types of packets for which IP TTL propagation is disabled by using the **forwarded** and **local** arguments. Specifying the **forwarded** parameter allows the structure of the MPLS network to be hidden from customers, but not from the provider. Selective IP TTL propagation hides the provider network from the customer but still allows troubleshooting.

- Cisco IOS or IOS-XE:

```
PE2(config)# no mpls ip propagate-ttl ?
   forwarded  Propagate IP TTL for forwarded traffic
   local      Propagate IP TTL for locally originated traffic
   <cr>
```

- Cisco IOS-XR configuration:

```
RP/0/RSP0/CPU0:PE1(config)# mpls ip-ttl-propagate disable
   forwarded  Disable IP TTL propagation for only forwarded MPLS   packets
   local      Disable IP TTL propagation for only locally generated MPLS packets
   <cr>
```

mpls ip propagate-ttl (mpls ip-ttl-propagate) Syntax Description

Parameter	Description
forwarded	(Optional) Hides the structure of the MPLS network from a **traceroute** command only for forwarded packets; prevents the **traceroute** command from showing the hops for forwarded packets
local	(Optional) Hides the structure of the MPLS network from a **traceroute** command only for local packets; prevents the **traceroute** command from showing the hops only for local packets

https://t.me/learningnets

LDP Session Protection Configuration

LDP Session Protection lets you configure LDP to automatically protect sessions with all or a given set of peers (as specified by the peer ACL). When it is configured, LDP initiates backup targeted hellos automatically for neighbors for which primary link adjacencies already exist. These backup targeted hellos maintain LDP sessions when primary link adjacencies go down.

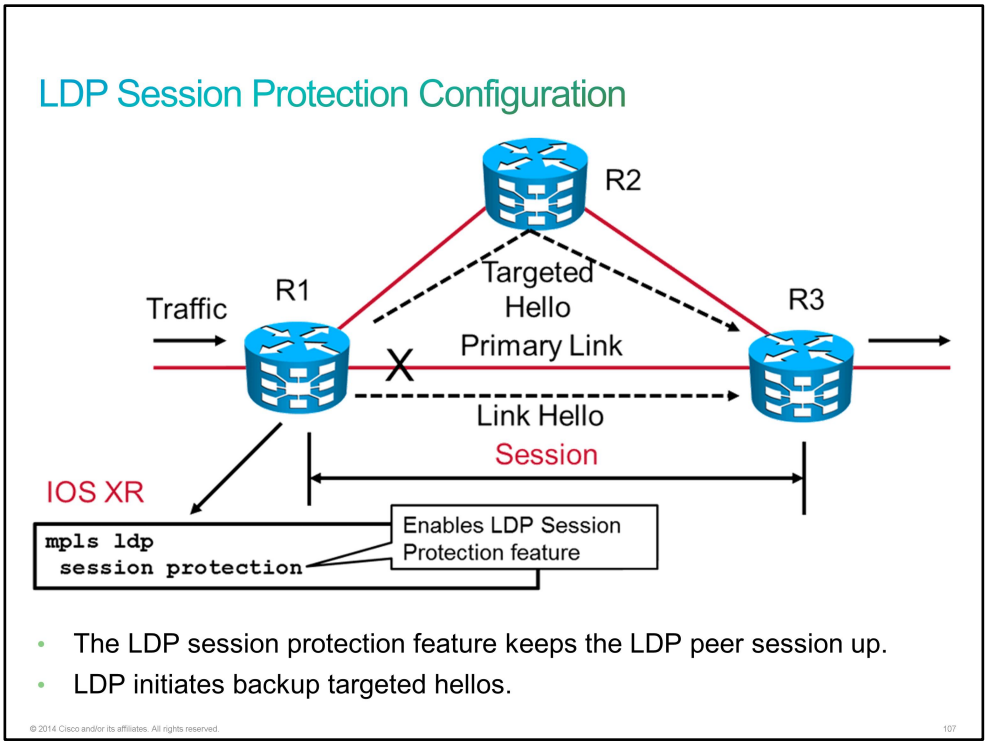

To enable the LDP Session Protection feature for keeping the LDP peer session up by means of targeted discovery following the loss of link discovery with a peer, use the **session protection** command in MPLS LDP configuration mode in Cisco IOS XR Software. To return to the default behavior, use the **no** form of this command:

```
session protection [duration seconds | infinite] [for peer-acl]
no session protection
```

By default, Session Protection is disabled. When it is enabled without peer ACL and duration, Session Protection is provided for all LDP peers, and continues for 24 hours after a link-discovery loss. This LDP Session Protection feature allows you to enable the automatic setup of targeted hello adjacencies with all or a set of peers, and specify the duration for which the session needs to be maintained using targeted hellos, after loss of link discovery. LDP supports only IPv4 standard access lists.

On Cisco IOS and IOS XE Software, the similar command to enable LDP session protection is the **mpls ldp session protection** global configuration command.

https://t.me/learningnets

LDP Graceful Restart and NSR Configuration

LDP graceful restart provides a control plane mechanism to ensure high availability, and allows detection and recovery from failure conditions while preserving NSF services. Graceful restart is a way to recover from signaling and control plane failures without impacting forwarding.

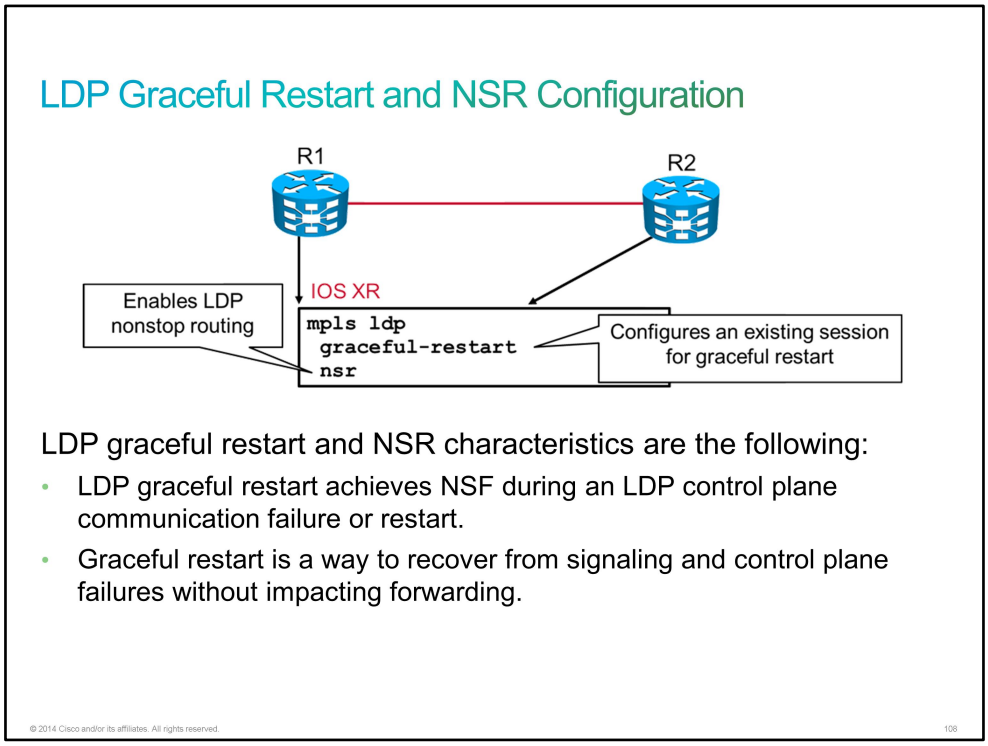

Use the LDP graceful restart capability to achieve nonstop forwarding during an LDP control plane communication failure or restart. To configure graceful restart between two peers, enable LDP graceful restart on both LSRs.

To configure graceful restart, use the **graceful-restart** command in MPLS LDP configuration mode. To return to the default behavior, use the **no** form of this command. To configure graceful restart between two peers, enable LDP graceful restart on both label switching routers.

```
graceful-restart [reconnect-timeout seconds | forwarding-state-holdtime seconds]
```

```
no graceful-restart [reconnect-timeout | forwarding-state-holdtime]
```

graceful-restart Syntax Description

Parameter	Description
forwarding-state-holdtime *seconds*	(Optional) Time that the local forwarding state is preserved (without being reclaimed) after the local LDP control plane restarts. The range is 60 to 600 seconds. The default is 180 seconds.
reconnect-timeout *seconds*	(Optional) Time that the local LDP sends to its graceful "restartable" peer, indicating how long its neighbor should wait for reconnection if there is an LDP session failure. The range is 60 to 300 seconds. The default is 120 seconds.

https://t.me/learningnets

LDP NSR functionality makes failures, such as RP or DRP failover, invisible to routing peers with minimal to no disruption of convergence performance. To enable LDP NSR on Cisco IOS XR Software use the **nsr** command in **mpls ldp** configuration mode.

When you enable MPLS LDP graceful restart on a router that peers with an MPLS LDP SSO- or NSF-enabled router, the SSO- or NSF-enabled router can maintain its forwarding state when the LDP session between them is interrupted. While the SSO- or NSF-enabled router recovers, the peer router forwards packets using stale information. This action enables the SSO- or NSF-enabled router to become operational more quickly.

When an LDP graceful restart session is established and there is control plane failure, the peer LSR starts graceful restart procedures, initially keeps the forwarding state information pertaining to the restarting peer, and marks this state as stale. If the restarting peer does not reconnect within the reconnect timeout, the stale forwarding state is removed. If the restarting peer reconnects within the reconnect time period, it is provided recovery time to resynchronize with its peer. After this time, any unsynchronized state is removed.

The value of the forwarding state hold time keeps the forwarding plane state associated with the LDP control plane in case of a control plane restart or failure. If the control plane fails, the forwarding plane retains the LDP forwarding state for twice the forwarding state hold time. The value of the forwarding state hold time is also used to start the local LDP forwarding state hold timer after the LDP control plane restarts. When the LDP graceful restart sessions are renegotiated with its peers, the restarting LSR sends the remaining value of this timer as the recovery time to its peers. Upon local LDP restart with graceful restart enabled, LDP does not replay forwarding updates to MPLS forwarding until the forwarding state hold timer expires.

To display the status of the LDP graceful restart, use the **show mpls ldp graceful-restart** command in EXEC mode. You can also check to see if the router is configured for graceful restart with the **show mpls ldp neighbor brief** command in EXEC mode.

```
RP/0/RP0/CPU0:router# show mpls ldp neighbor brief
v
Peer               GR Up Time         Discovery Address
------------------ -- --------------- --------- -------
3.3.3.3:0          Y  00:01:04                3       8
2.2.2.2:0          N  00:01:02                2       5

RP/0/RP0/CPU0:router# show mpls ldp graceful-restart

Forwarding State Hold timer : Not Running
GR Neighbors                : 1

 Neighbor ID      Up  Connect Count  Liveness Timer   Recovery Timer
 ---------------  --  -------------  ---------------  --------------
 3.3.3.3          Y        1                -                -
```

On Cisco IOS and IOS XE Software, a similar command that will enable LDP graceful restart is the **mpls ldp graceful-restart** global configuration command.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

LDP IGP Synchronization Configuration

Lack of synchronization between LDP and IGP can cause MPLS traffic loss. Upon link up, for example, IGP can advertise and use a link before LDP convergence has occurred; or, a link may continue to be used in the IGP after an LDP session goes down.

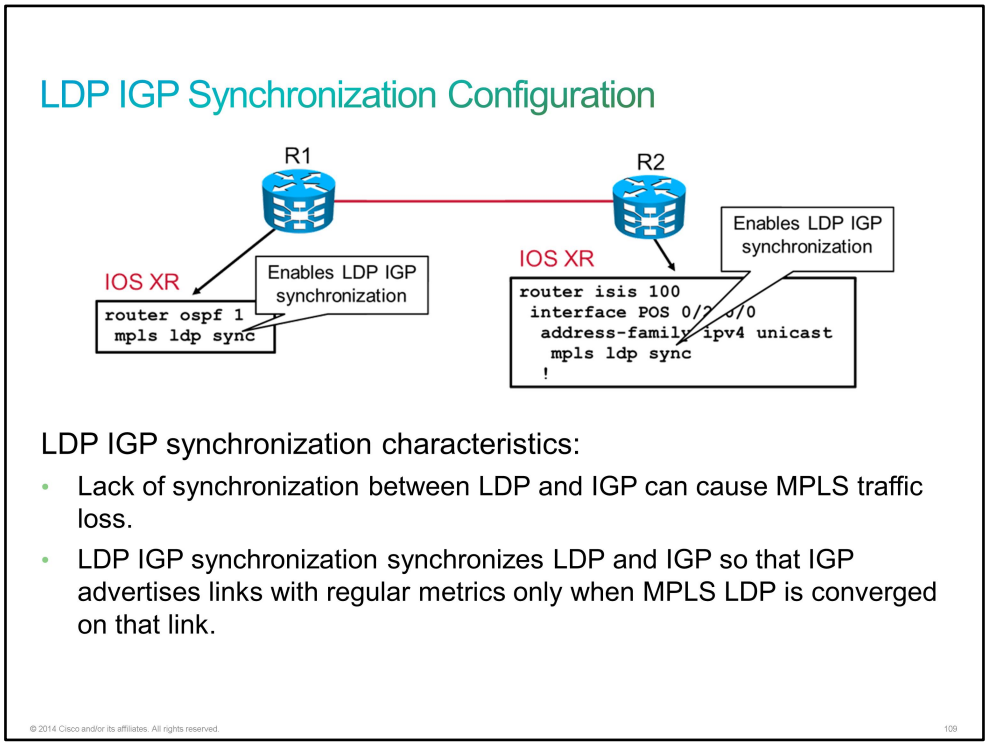

LDP IGP synchronization synchronizes LDP and IGP so that IGP advertises links with regular metrics only when MPLS LDP is converged on that link. LDP considers a link converged when at least one LDP session is operating on the link for which LDP has sent its applicable label bindings and has received at least one label binding from the peer. LDP communicates this information to IGP upon link up or session down events and IGP acts accordingly, depending on the synchronization state.

Normally, when LDP IGP synchronization is configured, LDP notifies IGP as soon as LDP is converged. When the delay timer is configured, this notification is delayed. Under certain circumstances, it might be required to delay declaration of resynchronization to a configurable interval. LDP provides a configuration option to delay synchronization for up to 60 seconds.

The LDP IGP synchronization feature is only supported for OSPF or IS-IS. To enable label LDP IGP synchronization on Cisco IOS, IOS XE, and IOS XR Software, use the **mpls ldp sync** command in the appropriate mode (OSPF or IS-IS configuration mode). To disable LDP IGP synchronization, use the **no** form of this command.

https://t.me/learningnets

LDP Autoconfiguration

To enable LDP on many interfaces, IGP autoconfiguration allows you to automatically configure LDP on all interfaces that are associated with a specified OSPF or IS-IS interface. However, there must be one IGP (OSPF or IS-IS) that is set up to enable LDP autoconfiguration.

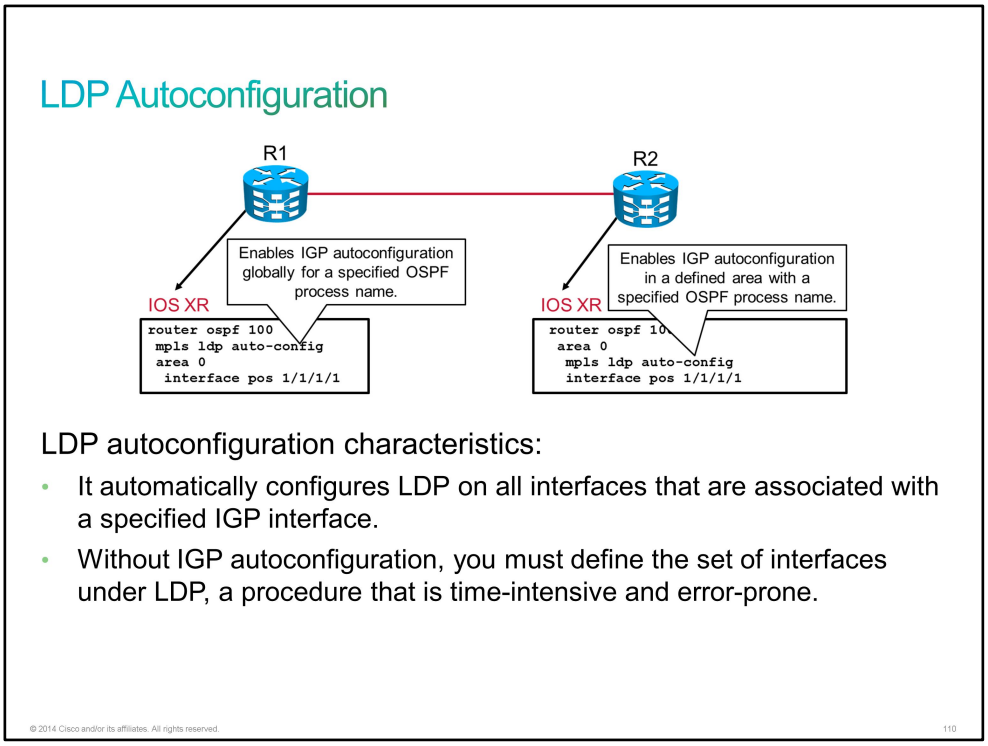

Without IGP autoconfiguration, you must define the set of interfaces under LDP, a procedure that is time-intensive and error-prone.

Similarly, on Cisco IOS and IOS XE Software, the **mpls ldp autoconfig** router OSPF or IS-IS configuration command is used to enable LDP autoconfiguration globally for the OSPF routing process.

To enable LDP autoconfiguration for a defined area in OSPF routing process use the **mpls ldp autoconfig area** *area_id* router OPSF configuration command.

https://t.me/learningnets

Label Advertisement Control

By default, LDP advertises labels for all the prefixes to all its neighbors. When this is not desirable (for scalability and security reasons), you can configure LDP to perform outbound filtering for local label advertisement, for one or more prefixes, to one or more peers. This feature is known as LDP outbound label filtering, or local label advertisement control.

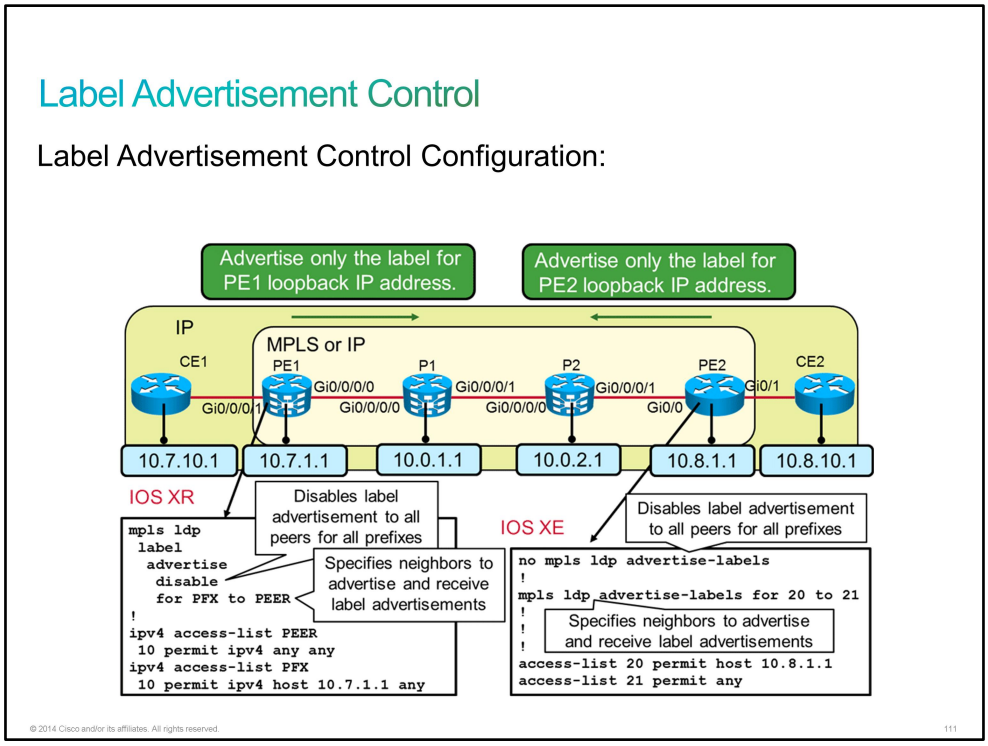

The example describes a situation where conditional label advertising can be used. The existing network still performs normal IP routing, but the MPLS LSP tunnel between the loopback interfaces of the LSR routers is needed to enable MPLS VPN functionality. Using one contiguous block of IP addresses for loopbacks on the PE routers can simplify the configuration of conditional advertising.

In the figure, the PE1 router (running Cisco IOS XR Software) should advertise only the label of the loopback prefix for PE1 (10.7.1.1/32), and not the loopback prefix of CE1 (10.7.10.1/32). In the same manner, the PE2 router (running Cisco IOS XE Software) should advertise only the label for the loopback prefix of PE2 (10.8.1.1/32), and not the loopback prefix of CE1 (10.8.10.1/32).

To control the advertisement of local labels on Cisco IOS XR Software, use the **label advertise** command in MPLS LDP configuration mode. To return to the default behavior, use the **no** form of this command:

```
label advertise {disable | for prefix-acl [to peer-acl] | interface interface}
no label advertise {disable | for prefix-acl [to peer-acl] | interface interface}
```

Example:

```
mpls ldp
label advertise disable
label advertise for PFX to PEER
```

The syntax description is listed here:

https://t.me/learningnets

Parameter	Description
for *prefix-access-list*	(Optional) This parameter specifies which destinations should have their labels advertised.
to *peer-access-list*	(Optional) This parameter specifies which LSR neighbors should receive label advertisements. An LSR is identified by its router ID, that consists of the first 4 bytes of its 6-byte LDP identifier.
interface *interface*	(Optional) This parameter specifies an interface for label allocation and advertisement of its interface IP address.

This command is used to control which labels are advertised to which LDP neighbors. On Cisco IOS and IOS XE Software, use the **mpls ldp advertise-labels** command in global configuration mode. To prevent the distribution of locally assigned labels, use the **no** form of this command, as shown here:

- **mpls ldp advertise-labels** [**for** *prefix-access-list* [**to** *peer-access-list*]]
- **no mpls ldp advertise-labels** [**for** *prefix-access-list* [**to** *peer-access-list*]]

The configuration in the figure for router PE1 disables label advertisement to all peers for all prefixes, except for prefix 10.7.1.1/32. The configuration in the figure for router PE2 disables label advertisement to all peers for all prefixes, except for prefix 10.8.1.1/32.

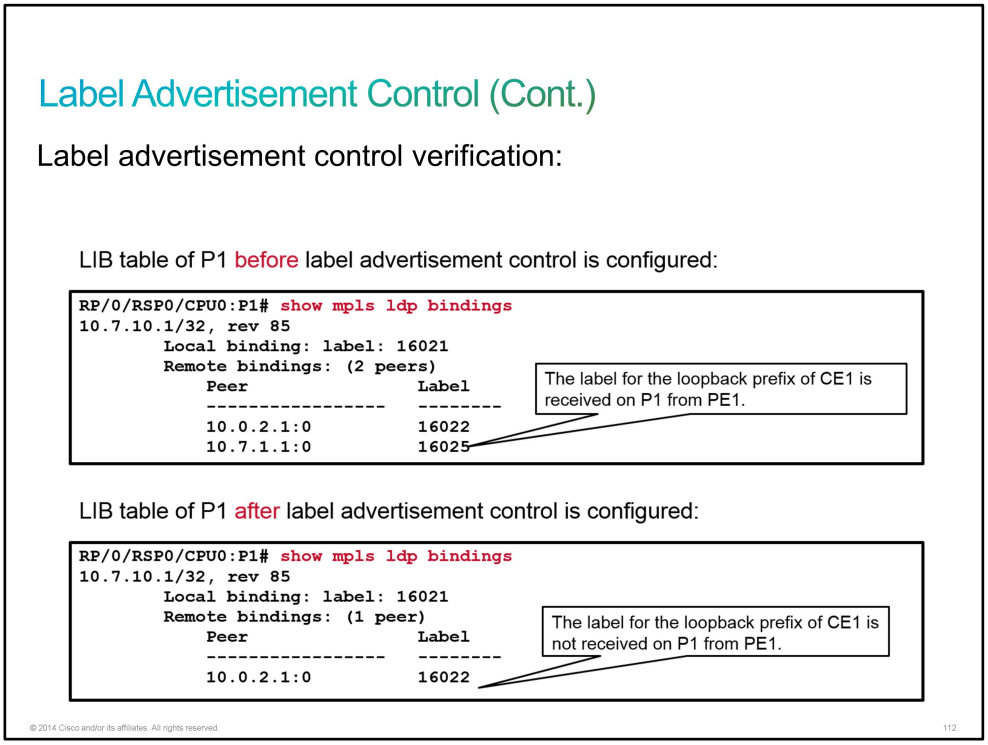

To verify the content of the LIB table, use the **show mpls ldp bindings** command. The output displays local labels for each destination network, as well as the labels that are received from all LDP neighbors.

https://t.me/learningnets

This output was taken from P1 before label advertisement control was configured. In the example, the local label for network 10.7.10.1/32 is 16021, the label received from 10.0.2.1 neighbor (P2) is 16022, and the label received from 10.7.1.1 neighbor (PE1) is 16025. Note that P1 received the label for network 10.7.10.1/32 from neighbor PE1, as highlighted in this output:

```
RP/0/RSP0/CPU0:P1# show mpls ldp bindings
10.7.1.1/32, rev 61
      Local binding: label: 16013
      Remote bindings: (3 peers)
          Peer                Label
          -----------------   --------
          10.0.2.1:0          16013
          10.7.1.1:0          IMP-NULL
10.7.10.1/32, rev 85
      Local binding: label: 16021
      Remote bindings: (3 peers)
          Peer                Label
          -----------------   --------
          10.0.2.1:0          16022
          10.7.1.1:0          16025
```

This output was taken from P1 after label advertisement control was configured. In the example, the local label for network 10.7.10.1/32 is 16021, and the label received from 10.0.2.1 neighbor (P2) is 16022. Note that P1 did not receive the label for network 10.7.10.1/32 from neighbor PE1.

```
RP/0/RSP0/CPU0:P1# show mpls ldp bindings
10.7.1.1/32, rev 61
      Local binding: label: 16013
      Remote bindings: (3 peers)
          Peer                Label
          -----------------   --------
          10.0.2.1:0          16013
          10.7.1.1:0          IMP-NULL
10.7.10.1/32, rev 85
      Local binding: label: 16021
      Remote bindings: (2 peers)
          Peer                Label
          -----------------   --------
          10.0.2.1:0          16022
```


https://t.me/learningnets

Label Acceptance Control

By default, LDP accepts labels (as remote bindings) for all prefixes from all peers. LDP operates in liberal label retention mode, that instructs LDP to keep remote bindings from all peers for a given prefix. For security reasons, or to conserve memory, you can override this behavior by configuring label binding acceptance for a set of prefixes from a given peer.

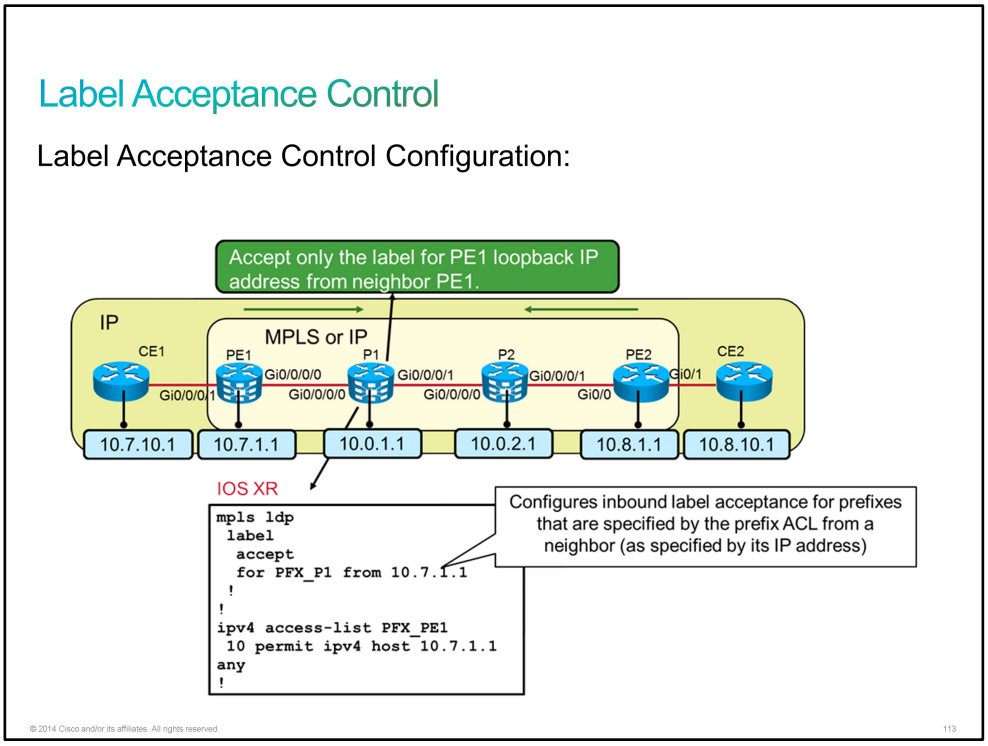

The ability to filter remote bindings for a defined set of prefixes is also referred to as LDP inbound label filtering.

To control the receipt of labels (remote bindings) on Cisco IOS XR Software for a set of prefixes from a peer, use the **label accept** command in MPLS LDP configuration mode. To return to the default behavior, use the **no** form of this command.

```
label accept for prefix-acl from A.B.C.D
```

```
no label accept for prefix-acl from A.B.C.D
```

Example:

```
mpls ldp
 label
  accept
  for PFX_P1 from 10.7.1.1
 !
!
ipv4 access-list PFX_PE1
 10 permit ipv4 host 10.7.1.1 any
!
```


https://t.me/learningnets

label accept for PFX_PE1 from 10.7.1.1 Syntax Description

Parameter	Description
for *prefix-acl*	Accepts and retains remote bindings for prefixes that are permitted by the prefix access list *prefix-acl*.
from *A.B.C.D*	Displays peer IP address

The configuration in the figure for router P1 accepts only the label for the PE1 loopback IP address from neighbor PE1.

https://t.me/learningnets

Monitor MPLS

Several **show** commands are used to monitor the MPLS operation.

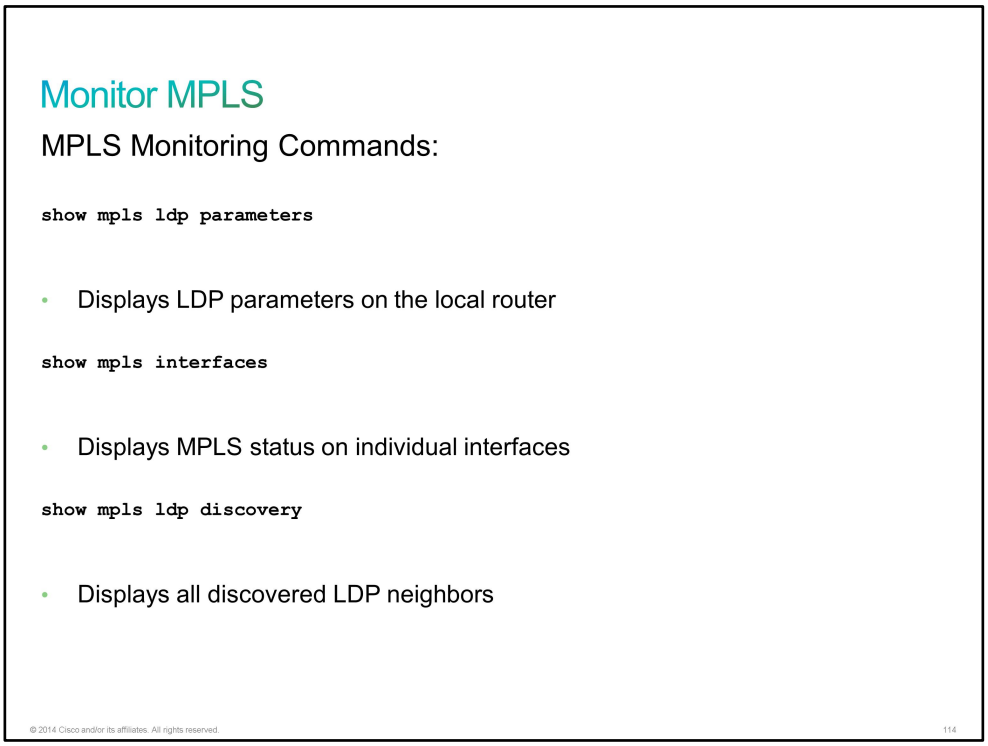

To display the available LDP parameters, use the **show mpls ldp parameters** command in privileged EXEC mode.

To display information about one or more interfaces that have the MPLS feature enabled, use the **show mpls interfaces** [*interfac*e] [**detail]** command in EXEC mode.

To display the status of the LDP discovery process (Hello protocol), use these commands in privileged EXEC mode:

- **show mpls ldp discovery** [**vrf** *vpn-name*]
- **show mpls ldp discover**y [*all*]

The **show mpls ldp discovery** command displays all MPLS-enabled interfaces and the neighbors that are present on the interfaces.

https://t.me/learningnets

Display MPLS LDP Parameters

When monitoring MPLS, one of the important aspects to verify are the LDP parameters.

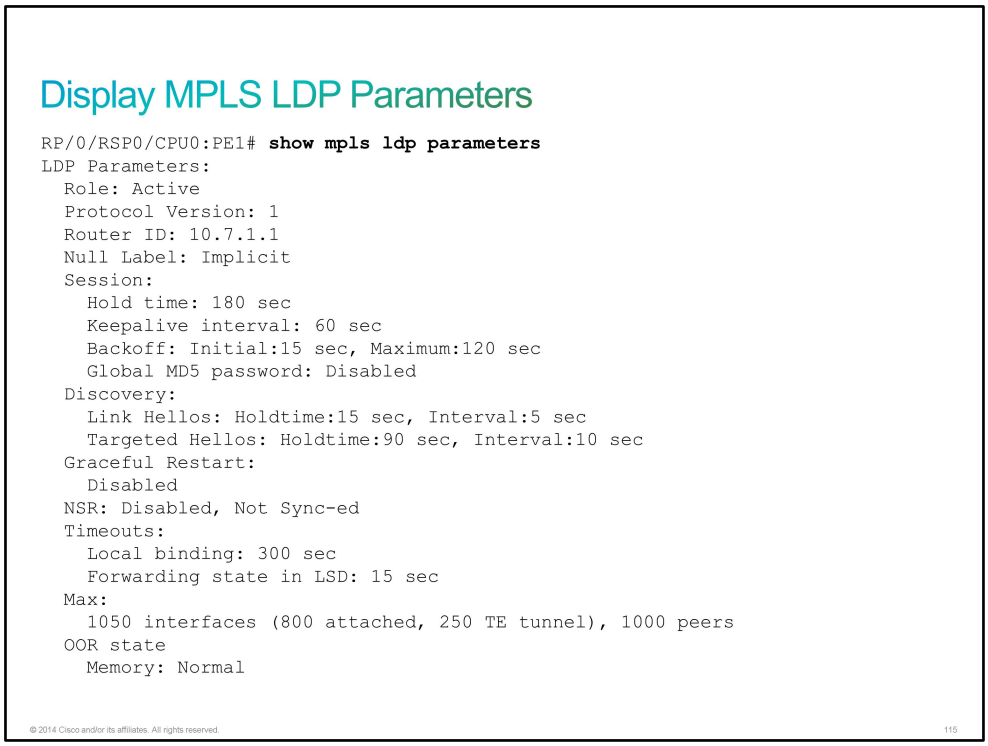

To display the current LDP parameters, use the **show mpls ldp parameters** command in EXEC mode.

https://t.me/learningnets

show mpls ldp parameters Field Description

Field	Description
Protocol version	Identifies the version of LDP that is running on the platform
Router ID	Currently used router ID
Null label	Identifies the LDP use of implicit or explicit null labels for prefixes where it has to use a null label
Session hold time	The time that an LDP session is to be maintained with an LDP peer, without receiving LDP traffic or an LDP keepalive message from the peer
Session keepalive interval	The interval of time between consecutive transmissions of LDP keepalive messages to an LDP peer
Session backoff	Initial maximum backoff time for sessions
Discovery link hellos	The time to remember that a neighbor platform wants an LDP session without receiving an LDP hello message from the neighbor (hold time), and the time interval between the transmission of consecutive LDP hello messages to neighbors (interval)
Discovery targeted hellos	Indicates the time: • To remember that a neighbor platform wants an LDP session when the neighbor platform is not directly connected to the router or the neighbor platform has not sent an LDP hello message. This intervening interval is known as *hold time*. • Interval between the transmission of consecutive hello messages to a neighbor that is not directly connected to the router, and if targeted hellos are being accepted, displaying peer-acl (if any)
Graceful restart	Graceful restart status
NSR	Nonstop routing status
Timeouts	Various timeouts (of interest) that LDP is using. One timeout is *binding no route*, which indicates how long LDP will wait for an invalid route before deleting it. It also shows Restart recovery time for LSD and LDP.
OOR state	Out of resource memory state: Normal, Major, or Critical.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Display MPLS Configured Interfaces

When monitoring MPLS, you may want to find out which interfaces are configured for MPLS.

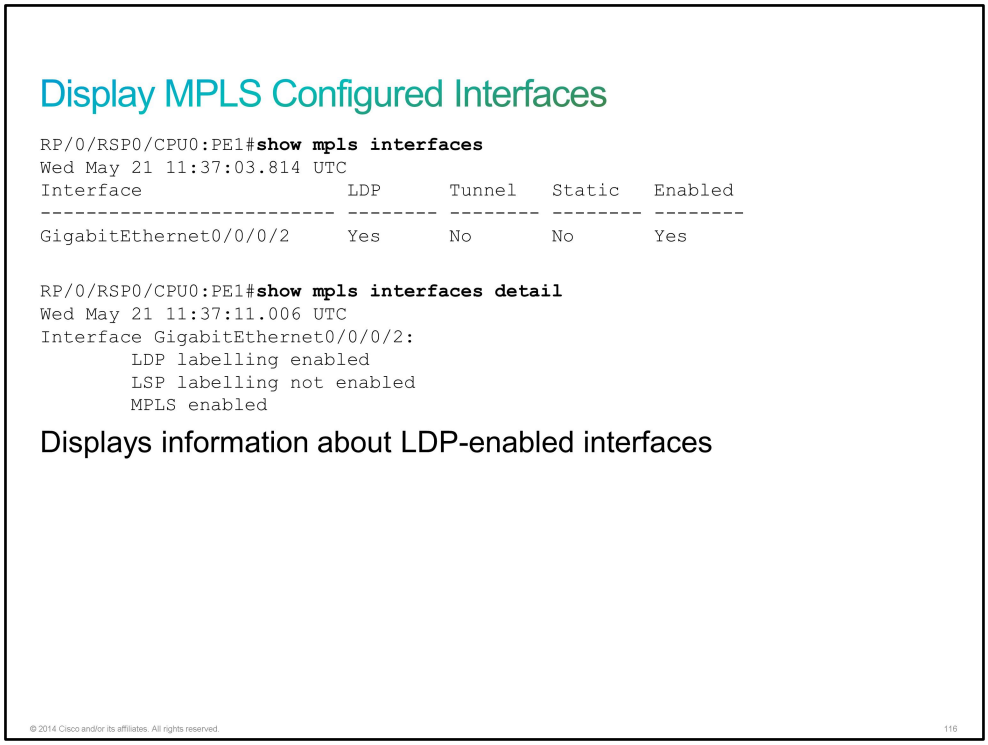

To display information about LDP-enabled interfaces, use the **show mpls ldp interfaces** command in EXEC mode. To display additional information, use the **show mpls ldp interfaces detail** command.

https://t.me/learningnets

Display Status of LDP Discovery Process

When monitoring MPLS, you may want to display the status of the LDP discovery process.

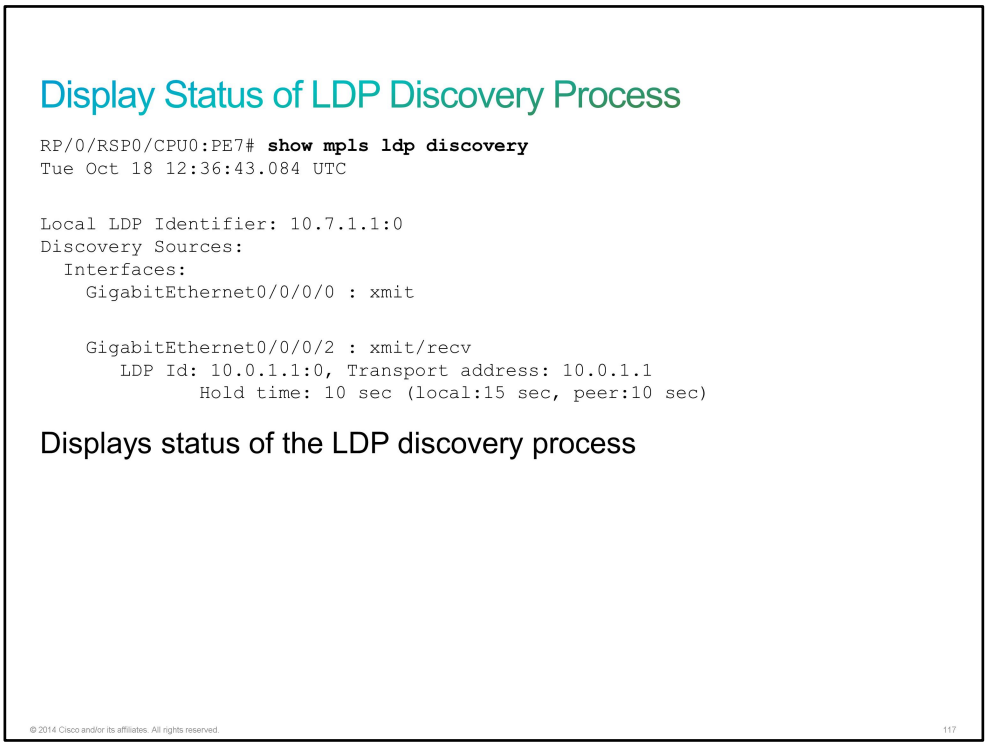

To display the status of the LDP discovery process, use the **show mpls ldp discovery** command in EXEC mode. The **show mpls ldp discovery command** shows both, link discovery and targeted discovery. When no interface filter is specified, this command generates a list of interfaces that are running the LDP discovery process. This command also displays neighbor discovery information for the default routing domain.

https://t.me/learningnets

show mpls ldp discovery Field Description

Field	Description
Local LDP identifier	LDP identifier for the local router. An LDP identifier is a 6-byte construct displayed in the form of an IP address:number. By convention, the first 4 bytes of the LDP identifier constitute the router ID; integers, starting with 0, constitute the final two bytes of the IP address:number construct.
Interfaces	The interfaces that are engaging in LDP discovery activity, as described here: • The xmit field indicates that the interface is transmitting LDP discovery hello packets. • The recv field indicates that the interface is receiving LDP discovery hello packets. The LDP identifiers indicate LDP (or TDP) neighbors discovered on the interface.
Transport address	Address that is associated with this LDP peer (advertised in hello messages)
LDP ID	LDP identifier of the LDP peer
Hold time	State of the forwarding hold timer and its current value

https://t.me/learningnets

LDP Monitoring Commands

Several **show** commands are used to monitor the LDP operation.

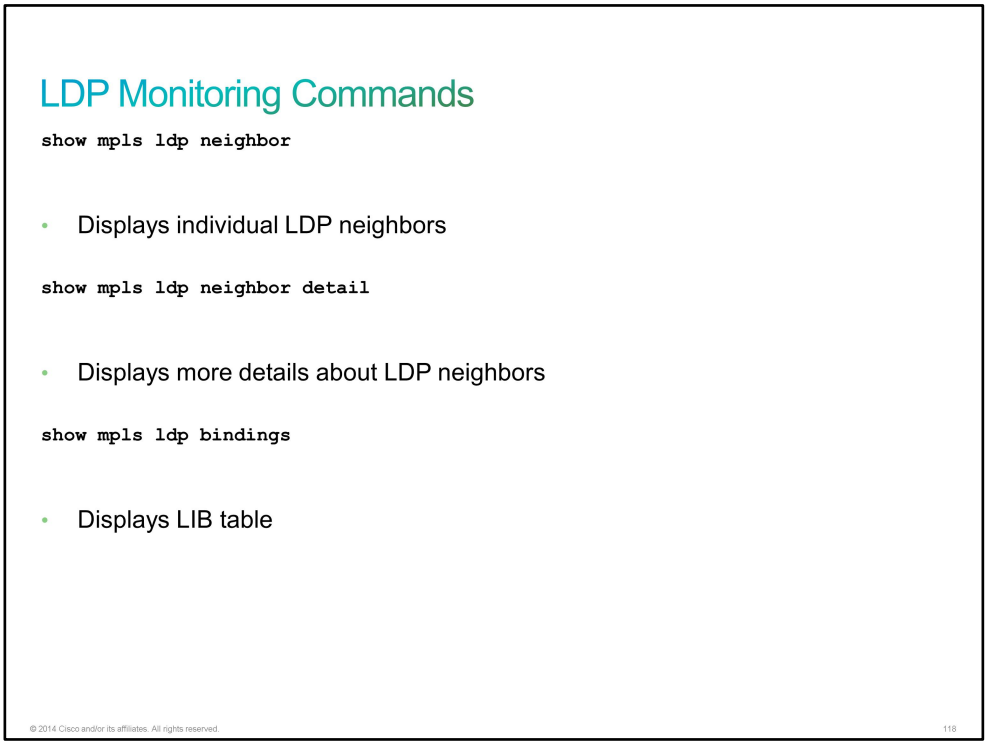

To display the status of LDP sessions, use the **show mpls ldp neighbor** command. To display the contents of the LIB, use the **show mpls ldp bindings** command.

https://t.me/learningnets

Display the Status of the LDP Session

When monitoring LDP, you may want to display the status of the LDP session.

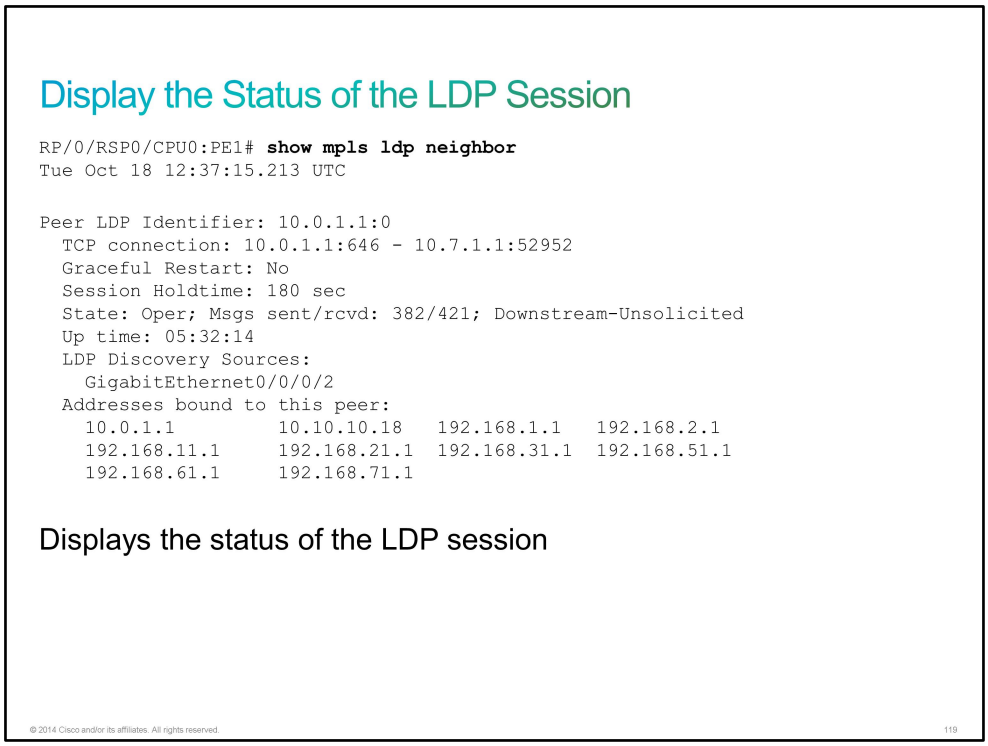

To display the status of LDP sessions, use the **show mpls ldp neighbor** command:

show mpls ldp neighbor [*A.B.C.D* | *type interface-path-id* | **gr** | **non-gr** | **sp** | | **standby** | **brief**] [**detail**]

The status of the LDP session is indicated by “State: Oper” (operational).

The **show mpls ldp neighbor** command provides information about all LDP neighbors in the entire routing domain; conversely, the **show** output is filtered to display:

- LDP neighbors with specific IP addresses
- LDP neighbors on a specific interface
- LDP neighbors that can be gracefully restarted
- LDP neighbors that cannot be gracefully restarted
- LDP neighbors enabled with session protection

https://t.me/learningnets

show mpls ldp neighbor Field Description

Field	Description
Peer LDP identifier	This field is the LDP identifier of the neighbor (peer) for this session
Graceful restart	This field is the graceful restart status (Y or N)
TCP connection	This field displays the TCP connection used to support the LDP session, shown in the format that follows: - neighbor IP address.peer port - local IP address.local port
State	This field displays the state of the LDP session. Generally, this is “Oper” (operational), but “transient” is another possible state.
Msgs sent/rcvd	This field displays number of LDP messages sent to and received from the session peer. The count includes the transmission and receipt of periodic keepalive messages, which are required for maintenance of the LDP session.
Up time	This field displays the length of time that the LDP session has existed.
LDP discovery sources	This field displays the source (or sources) of LDP discovery activity that led to the establishment of this LDP session.
Addresses bound to this peer	This field displays the known interface addresses of the LDP session peer. These are addresses that might appear as next-hop addresses in the local routing table. They are used to maintain the LFIB.

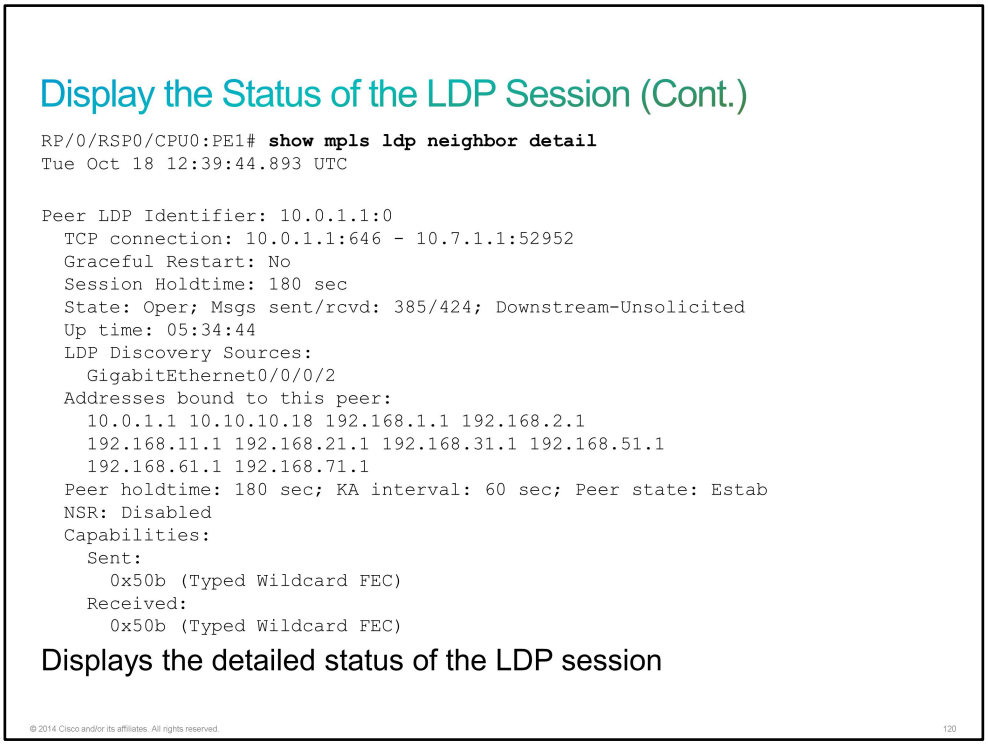

To display the detailed status of LDP sessions, use the **show mpls ldp neighbor detail** command.

https://t.me/learningnets

Display the Contents of the LIB Table

When verifying the operation of LDP, you may want to display the contents of the LIB table.

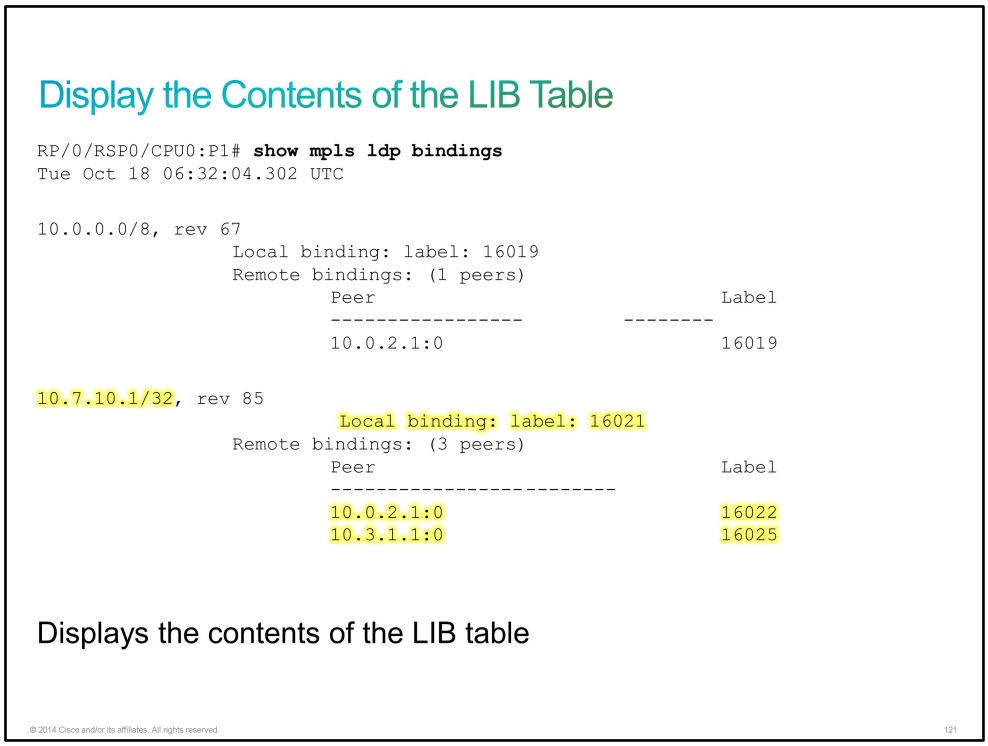

To verify the contents of the LIB table, use the **show mpls ldp bindings** command. The output displays local labels for each destination network, as well as the labels that have been received from all LDP neighbors.

To display the contents of the LIB, use the **show mpls ldp bindings** command in EXEC command:

show mpls ldp bindings [*prefix* {*mask* | *length*}] [**advertisement-acls**] [**detail**] [**local**] [**local-label** *label* [**to** *label*]] [**neighbor** *address*] [**remote-label** *label* [**to** *label*]] [**summary**]

You can choose to view the entire database or a subset of entries according to the following criteria:

- Prefix
- Input or output label values or ranges
- Neighbor advertising the label

In the example, the local label for network 10.7.10.1/32 is 16021, the label for that network received from the 10.0.2.1 neighbor is 16022, and the label received from the 10.3.1.1 neighbor is 16025.

Note	The **show mpls ldp bindings summary** command displays summarized information from the LIB and is used when you are testing scalability or when it is deployed in a large-scale network.

https://t.me/learningnets

show mpls ldp bindings Field Descriptions

Field	Description
a.b.c.d/n	This field is the IP prefix and mask for a particular destination (network/mask).
Rev	This field is the revision number (rev) that is used internally to manage label distribution for this destination.
Local binding	A locally assigned label for a prefix
Remote bindings	Outgoing labels for this destination that are learned from other LSRs. Each item in this list identifies the LSR from which the outgoing label was learned and reflects the label that is associated with that LSR. Each LSR in the transmission path is identified by its LDP identifier.
(Rewrite)	Binding has been written into MPLS forwarding and is in use.
(No route)	Route is not valid. LDP times it out before the local binding is deleted.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Monitor Label Switching

Several **show** commands are used to monitor label switching.

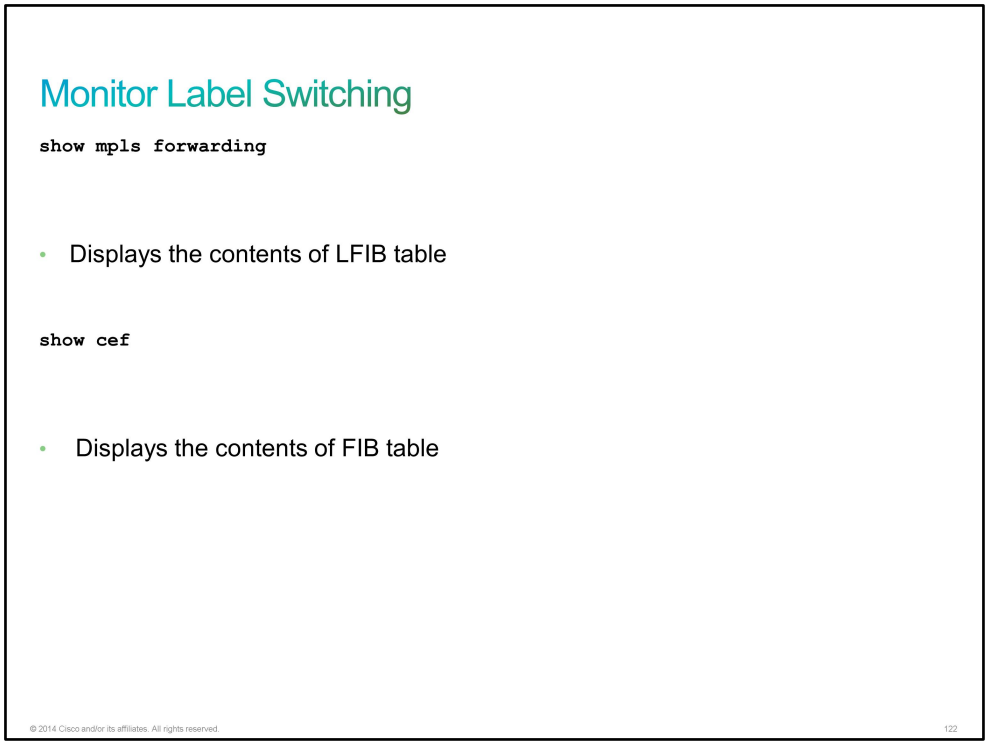

To display the contents of the MPLS LFIB, use the **show mpls forwarding** command in EXEC mode.

To display the contents of the FIB Cisco Express Forwarding table, use the **show cef** command in EXEC mode.

https://t.me/learningnets

Display the Contents of the LFIB Table

When monitoring label switching, you may want to display the content of the LFIB table.

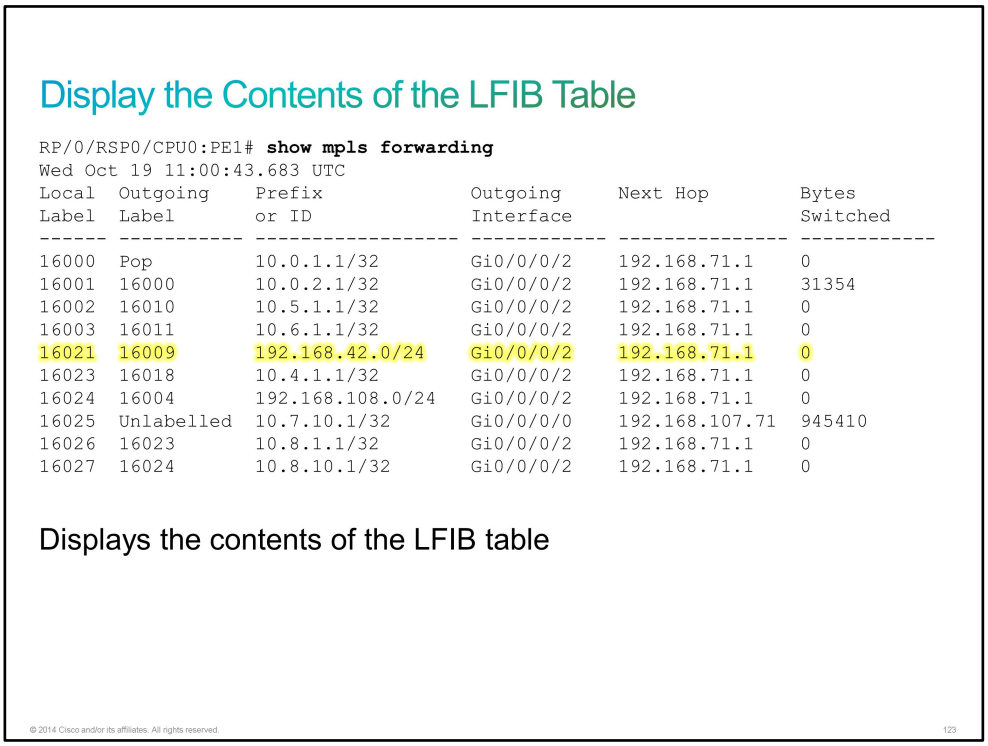

To display the contents of the MPLS LFIB, use the **show mpls forwarding** command in EXEC mode:

show mpls forwarding [**detail** | {**label** *label number*} | **interface** *interface-path-id* | **labels** *value* | **location** | **prefix** [*network/mask* | *length*] | **private** | **summary** | **tunnels** *tunnel-id*]

The output displays the incoming and outgoing label for each destination, together with the outgoing interface and the next hop. In the example, the incoming (Local) label for the 192.168.42.0/24 network is 16021 (allocated by this router), and the outgoing label is 16009 (as advertised by the next hop). In the example, you can also see that the network 10.0.1.1/32 has a POP label set as the outgoing label. This means that the router has learned from a neighbor that a label should be removed from a labeled packet. There is also network 10.7.10.1/32 that does not have an outgoing label set (unlabeled); this means that the label has not yet been received from a neighbor, or that the network is outside the MPLS domain and the router is an edge LSR.

https://t.me/learningnets

show mpls forwarding Field Descriptions

Field	Description
Local label	This field is the IP prefix and mask for a particular destination (network/mask).
Outgoing label	This field is the label that was assigned by the next hop or downstream peer. Some of the entries that display in this column are as follows: • Unlabeled : No label for the destination from the next hop, or label switching is not enabled on the outgoing interface. • Pop Label: Next hop advertised an implicit-null label for the destination.
Prefix or tunnel ID	This field is the address or tunnel where packets with this label are going.
Outgoing interface	This field is the interface through which packets with this label are sent.
Next hop	This field is the IP address of the neighbor that assigned the outgoing label.
Bytes switched	This field is the number of bytes switched with this incoming label.

https://t.me/learningnets

Display the Contents of the FIB Table

When monitoring label switching, you may want to display the contents of the FIB table.

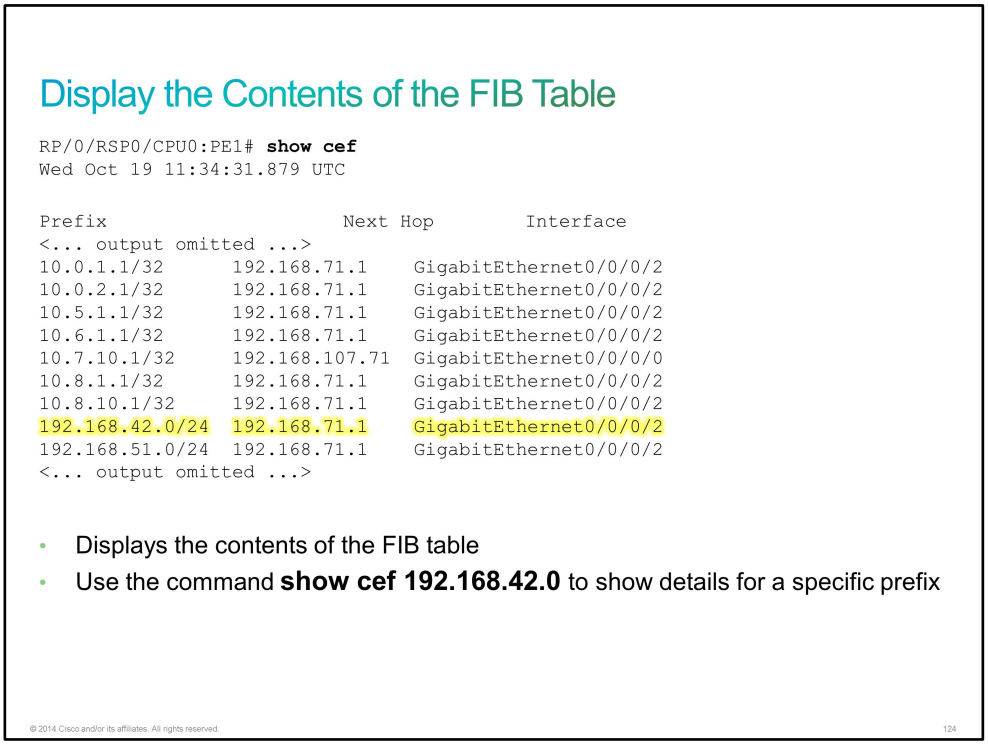

To display information about packets forwarded by FIB Cisco Express Forwarding, use the **show cef** command in EXEC mode:

show cef [*prefix* [*mask*]] [**hardware** {**egress** | **ingress**} | **detail**] [**location** {*node-id* | **all**}]

In the figure, the next hop for network 192.168.42.0/24 is 192.168.71.1 and the outgoing interface is GigabitEthernet0/0/0/2.

To verify the FIB Cisco Express Forwarding table, use the **show cef** command, followed by a desired prefix:

```
RP/0/RSP0/CPU0:PE7# show cef 192.168.42.0
Wed Oct 19 12:08:03.213 UTC
192.168.42.0/24, version 0, internal 0x4004001 (ptr 0xad958e70) [1], 0x0 (0xacf50a94),
0x450 (0xadffc4b0)
 Updated Oct 19 05:51:07.981
 remote adjacency to GigabitEthernet0/0/0/2
 Prefix Len 24, traffic index 0, precedence routine (0)
   via 192.168.71.1, GigabitEthernet0/0/0/2, 4 dependencies, weight 0, class 0 [flags
0x0]
   path-idx 0
   next hop 192.168.71.1
   remote adjacency
     local label 16021      labels imposed {16009}
```


https://t.me/learningnets

Debugging MPLS and LDP

A large number of debug commands are associated with MPLS.

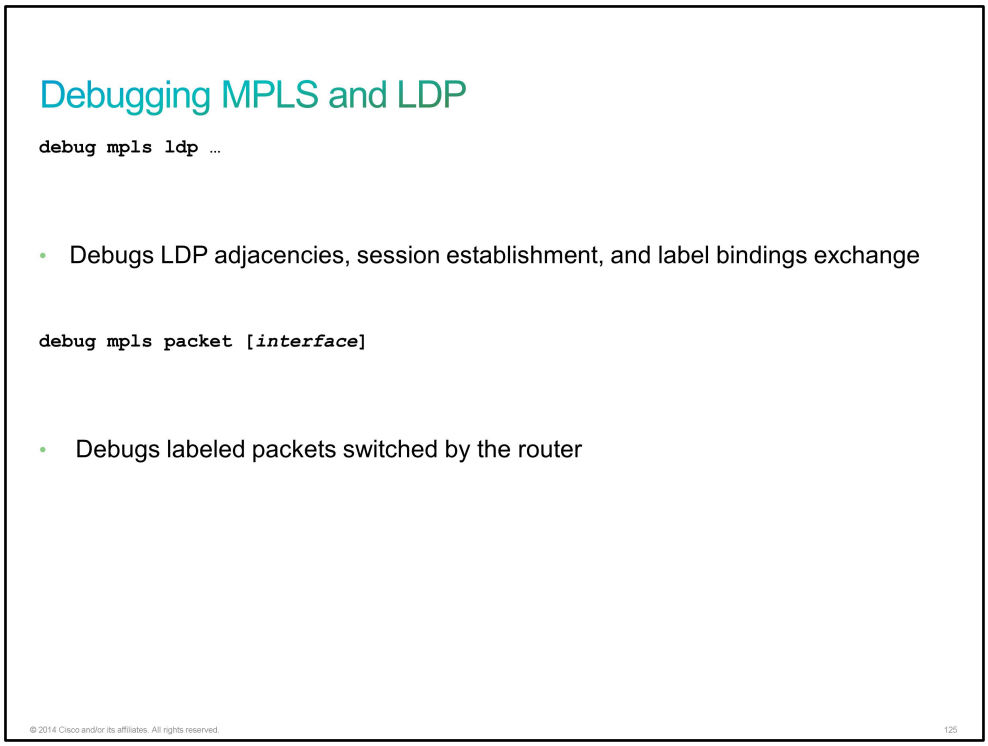

The **debug mpls ldp** commands debug various aspects of the LDP protocol, from label distribution to exchange of the application layer data between adjacent LDP-speaking routers.

Note	Use **debug** commands with caution. Enabling debugging can disrupt the operation of the router under high load conditions. Before you start a **debug** command, always consider the output that the command may generate and the amount of time this may take. You should also look at your CPU load before debugging by using the **show processes cpu** command. Verify that you have ample CPU capacity available before beginning the debugging process.

The **debug mpls packet** command displays all labeled packets that are switched by the router (through the specified interface).

Note	Use the **debug mpls packet** command with care, because it generates output for every packet that is processed. Furthermore, enabling the **debug mpls packet** command causes fast and distributed label switching to be disabled for the selected interfaces. To avoid adversely affecting other system activity, use this command only when traffic on the network is at a minimum.

https://t.me/learningnets

Classic Ping and Traceroute

Standard ping and traceroute tools can be used in MPLS environments to test reachability in different scenarios.

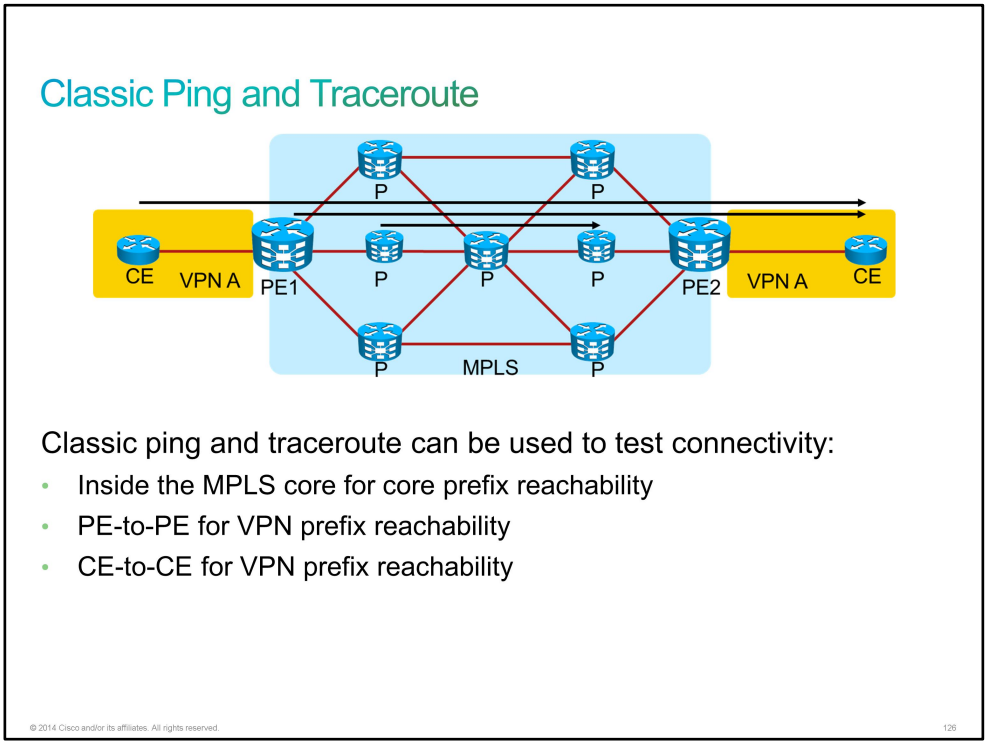

Standard ping and traceroute tools can be used in MPLS environments to test reachability in three different scenarios:

- Test the reachability of prefixes that are reachable through the global routing table via IP forwarding or label switching. The tools can be used on PE and P routers.
- Test the reachability of Layer 3 MPLS VPN prefixes that are reachable through a VRF routing table via label switching. The tools can be used on PE routers configured with the required VRF.
- Customers can use the tools to test Layer 3 MPLS VPN connectivity end-to-end.

https://t.me/learningnets

Classic Ping and Traceroute (Cont.)

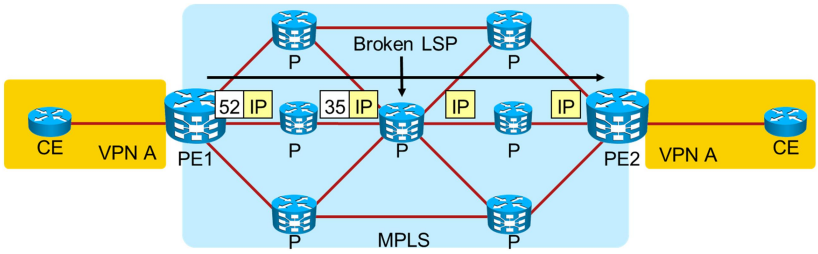

Classic ping and traceroute characteristics:

- Broken LSPs may not always be revealed.
- Broken LSPs revert back to IP forwarding.
- Ping and traceroute succeed.

© 2014 Cisco and/or its affiliates. All rights reserved.

The figure illustrates a redundant network where the standard traceroute tool was used to determine the path from one PE to another PE. Cisco routers in the path will encode some MPLS information into ICMP replies to be displayed by the router that is initiating route tracing.

Classic Ping and Traceroute (Cont.)

```
RP/0/RSP0/CPU0:PE1# traceroute 172.16.1.14

Type escape sequence to abort.
Tracing the route to 172.16.1.14

  1 192.168.1.46 [MPLS: Label 34 Exp 0] 8 msec
    192.168.1.42 [MPLS: Label 38 Exp 0] 12 msec
    192.168.1.46 [MPLS: Label 34 Exp 0] 24 msec
  2 192.168.1.14 [MPLS: Label 37 Exp 0] 48 msec
    192.168.1.18 [MPLS: Label 33 Exp 0] 8 msec
    192.168.1.14 [MPLS: Label 37 Exp 0] 8 msec
  3 192.168.1.38 52 msec
    192.168.1.34 8 msec *
```

Classic ping and traceroute characteristics:

- Cisco IOS Software does encode MPLS information in ICMP replies.
- Even multiple paths can sometimes be detected.

© 2014 Cisco and/or its affiliates. All rights reserved.

The sample output shows how labels are displayed and how even multiple paths can be detected, although not always reliably (equal paths on subsequent hops will typically not be displayed by classic traceroute).

https://t.me/learningnets

MPLS Ping and Traceroute Overview

Special MPLS ping and MPLS traceroute were designed for monitoring and troubleshooting MPLS LSPs. These features provide a means to check connectivity and isolate a failure point, thus providing the MPLS OAM solution.

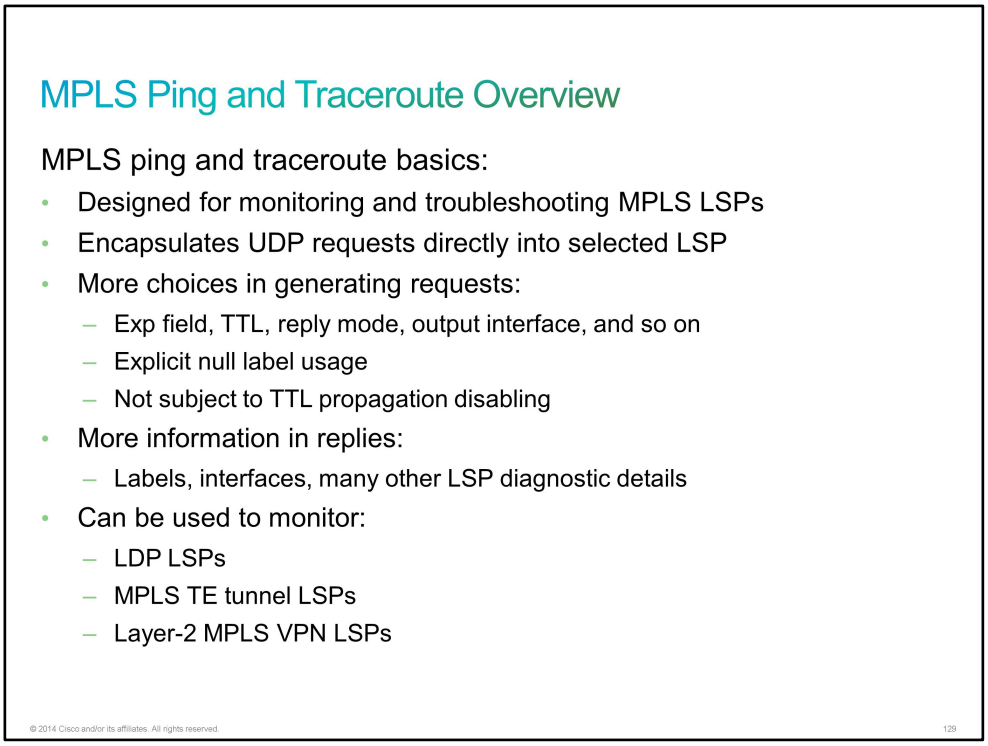

Normal ICMP ping and traceroute are used to help diagnose the root cause when a forwarding failure occurs. However, they may not detect LSP failures because an ICMP packet can be forwarded via IP to the destination when an LSP breakage occurs, whereas MPLS LSP ping and traceroute can be used to identify LSP breakages.

An FEC must be selected to choose the associated LSP. FECs can be any of these:

- IP prefix in the FIB table with label next-hop information
- Layer 2 MPLS VPN virtual circuit (pseudowire)
- MPLS TE tunnel LSP.

MPLS ping and traceroute will use UDP packets with loopback destination addresses to encode requests and label them with the selected FEC label.

Enable MPLS OAM by using the **mpls oam** command on all routers in the MPLS network.

https://t.me/learningnets

MPLS Ping

MPLS LSP ping uses MPLS echo request and reply packets to validate an LSP.

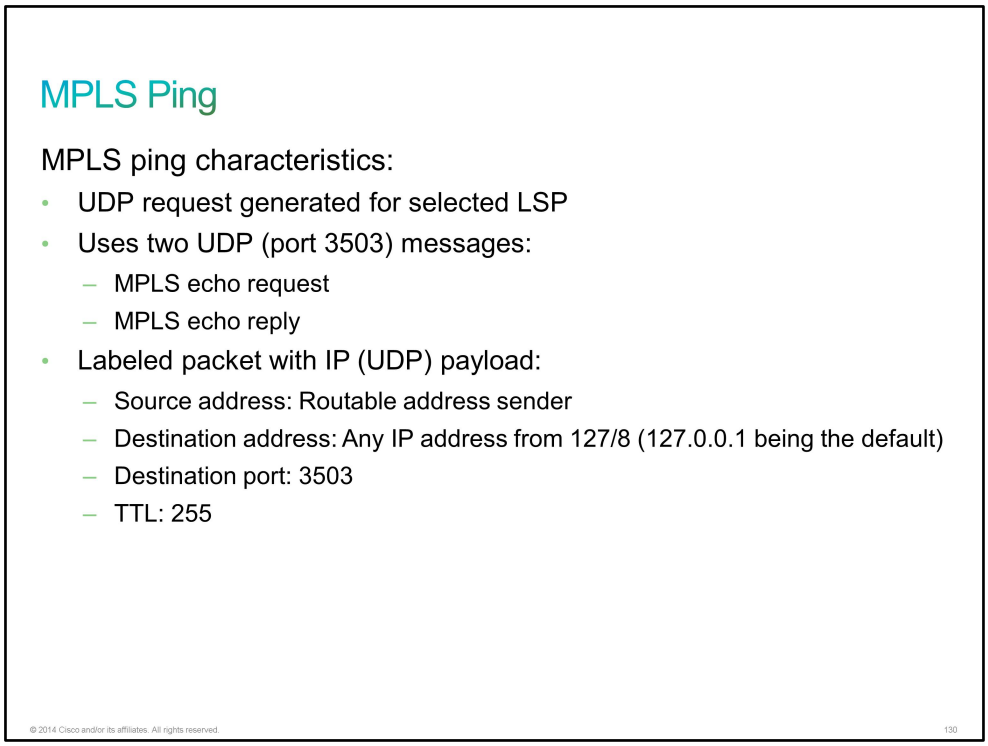

MPLS ping uses UDP on port 3503 to encode two types of messages:

- **MPLS echo request:** The MPLS echo request message includes the information about the tested FEC (prefix) which is encoded as one of the TLVs. Additional TLVs can be used to request additional information in replies. The downstream mapping TLV can be used to request information about additional information such as downstream router and interface, MTU, and multipath information from the router where the request is processed.
- **MPLS echo reply:** The MPLS echo reply uses the same packet format as the request, except that it may include additional TLVs to encode the information. The basic reply is, however, encoded in the reply code field.

The MPLS echo request will use the outgoing interface IP address as the source and a loopback IP address, which is configurable, as the destination (127.0.0.1). The TTL in MPLS ping is set to 255. Using the 127/8 address in the IP header destination address field will cause the packet to not be forwarded by any routers using the IP header, if the LSP is broken somewhere inside the MPLS domain.

The initiating router can also request a reply mode, that can be one of the following:

- The default reply mode uses IPv4 and MPLS to return the reply.
- Router alert mode that forces every router in the return path to perform process switching of the return packet, which in turn forces the use of the IP forwarding table (avoids any confusion if the return LSP is broken). This functionality is achieved by adding label 1 onto the label stack for the reply packet.
- It is also possible to request that no reply be sent.

https://t.me/learningnets

MPLS Ping (Cont.)

MPLS Ping Example 1:

```
RP/0/RSP0/CPU0:PE1# ping mpls ipv4 172.16.1.14 255.255.255.255
Sending 5, 100-byte MPLS Echos to 172.16.1.14/32,
         timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/17/20 ms
```

IPv4 FEC from the global IPv4 routing table.

© 2014 Cisco and/or its affiliates. All rights reserved.

The sample MPLS ping illustrates the command syntax. The main difference from the standard ping is the need to exactly specify the FEC (a prefix in the FIB table) from which the router will learn the next-hop label, output interface, and Layer 2 forwarding information.

A successful reply is also represented by exclamation marks. A number of other results are possible, depending on the return code that can map to any of the characters described in the legend portion of the MPLS ping output.

ping mpls {ipv4 *addr/mask*} **[destination {***start address} {end address} {address increment}***] | [dsmap]** | **[exp** *exp bits in MPLS header***]** | **[force-explicit-null]** | **[interval** *send interval between requests in msec***]** | **[output interface** *echo request output interface***] [pad** *pad TLV pattern***] | [repeat** *repeat count*] | **[reply dscp** *differentiated services codepoint value***] | [reply mode [ipv4** | **router-alert** | **no-reply**] | **[reply pad-tlv]]** | **[revision** *echo packet tlv versioning***] | [{size** *packet size***} | [source** *source specified as an IP address***] | {sweep** {*min value*} {*max value*} {*increment*}**}] | [timeout** *timeout in seconds***] |** [**ttl** *time to live***]** | **[verbose]**

https://t.me/learningnets

MPLS Ping (Cont.)

MPLS Ping Example 2:

```
RP/0/RSP0/CPU0:PE1# ping mpls ipv4 172.16.1.14 255.255.255.255 ttl 1 dsmap
repeat 1
Sending 1, 100-byte MPLS Echos to 172.16.1.14/32,
     timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       <output omitted>
Type escape sequence to abort.
L
Echo Reply received from 192.168.1.2
  DSMAP 0, DS Router Addr 127.0.0.1, DS Intf Addr 0
   Depth Limit 0, MRU 1500 [Labels: 33 Exp: 0]
   Multipath Addresses:
Success rate is 0 percent (0/1)
```

- The dsmap option can be used to retrieve the details for a given hop.
- MPLS traceroute can be used instead to display detailed information for all hops.

© 2014 Cisco and/or its affiliates. All rights reserved.

The example illustrates how you can request the dsmap information and select from a specific hop. The reply contains the downstream information, including MTU.

Below is the full output from the MPLS ping example:

```
RP/0/RSP0/CPU0:PE1# ping mpls ipv4 172.16.1.14 255.255.255.255 ttl 1 dsmap repeat 1
Sending 1, 100-byte MPLS Echos to 172.16.1.14/32,
     timeout is 2 seconds, send interval is 0 msec:
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
       'L' - labeled output interface, 'B' - unlabeled output interface,
       'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
       'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
       'P' - no rx intf label prot, 'p' - premature termination of LSP,
       'R' - transit router, 'I' - unknown upstream index,
       'X' - unknown return code, 'x' - return code 0
Type escape sequence to abort.
L
Echo Reply received from 192.168.1.2
  DSMAP 0, DS Router Addr 127.0.0.1, DS Intf Addr 0
   Depth Limit 0, MRU 1500 [Labels: 33 Exp: 0]
   Multipath Addresses:
Success rate is 0 percent (0/1)
```


https://t.me/learningnets

MPLS Traceroute

MPLS LSP traceroute uses MPLS echo request and reply packets to validate an LSP.

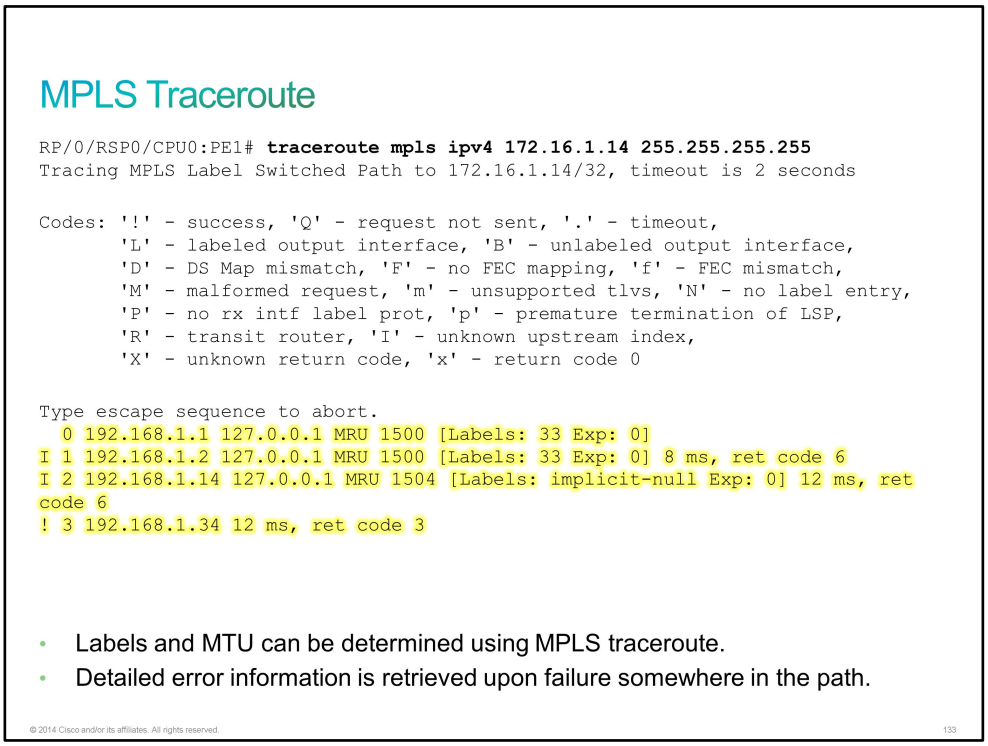

The sample MPLS traceroute shows that the downstream mapping information is reported from each hop where the TTL expired. The output shows the entire path with labels and MRU, as well as the round-trip time and IP addresses on routers in the path. MRU includes the maximum side of the IP packet including the label stack that could be forwarded out of the particular interface.

To learn the routes that packets follow when traveling to their destinations, use the **traceroute mpls** command in EXEC mode.

traceroute mpls {{ipv4 *addr/mask***} | {traffic-eng tunnel** *tunnel intf num***}} [destination** *{start address} {end address} {address increment}***] | [exp** *exp bits in MPLS header***] | [flags fec] | [force-explicit-null] | [output interface** *echo request output interface*] | [**reply dscp** *DSCP bits in reply IP header*] | [**reply mode** *[ipv4 | router-alert | no-reply*]] **[revision** *echo packet tlv versioning***] | [source** *source specified as an IP address***] | [timeout** *timeout in seconds***] | [ttl** *time to live***] | [verbose]**

https://t.me/learningnets

Troubleshooting MPLS

When faced with a troubleshooting ticket involving MPLS, you need to first clear the basics to ensure that everything is working from an LDP perspective.

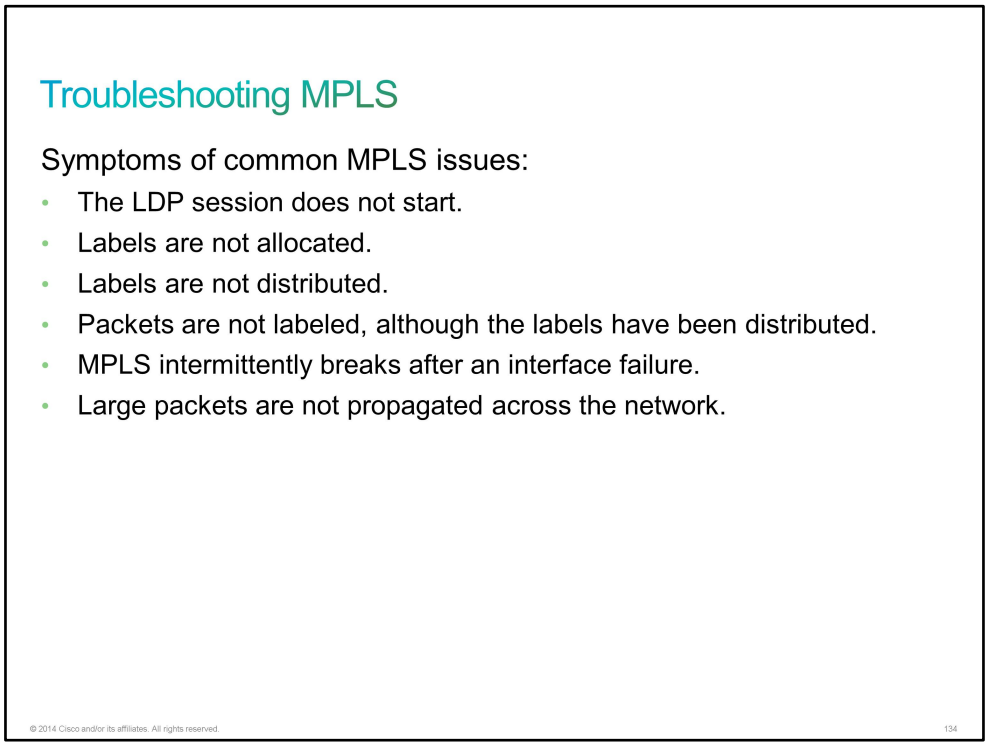

Here are the common issues that can be encountered while you are troubleshooting an MPLS network:

- The LDP session does not start.
- The LDP session starts, but the labels are not allocated or distributed.
- Labels are allocated and distributed, but the forwarded packets are not labeled.
- MPLS stops working intermittently after an interface failure, even on interfaces totally unrelated to the failed interface.
- Large IP packets are not propagated across the MPLS backbone, even though the packets were successfully propagated across the pure IP backbone.

https://t.me/learningnets

LDP Session Startup Issues

Typical examples of LDP issues are when neighbors are not discovered or when the LDP session is not established.

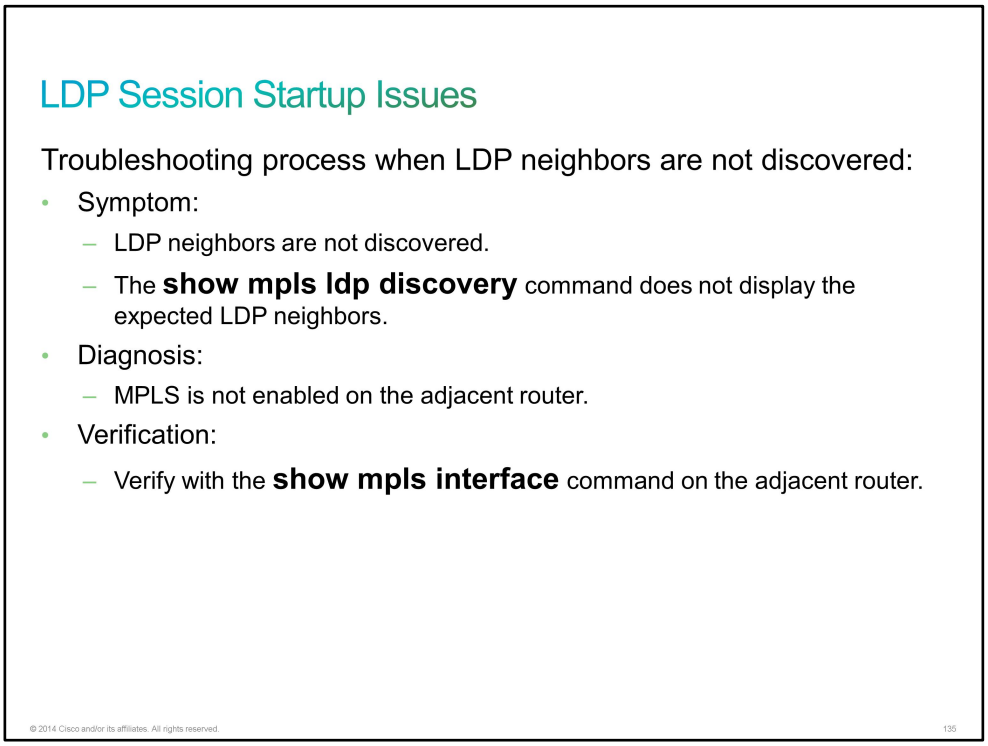

Symptom: If MPLS is enabled on an interface, but no neighbors are discovered, it is likely that MPLS is not enabled on the neighbor.

The router is sending discovery messages, but the neighbor is not replying because it does not have LDP enabled.

Solution: Enable MPLS on the neighboring router.

https://t.me/learningnets

LDP Session Startup Issues (Cont.)

Troubleshooting process when LDP session is not established:

- Symptom:
  - LDP neighbors are discovered; the LDP session is not established.
  - The **show mpls ldp neighbor** command does not display a neighbor in the operational state.
- Diagnosis:
  - The connectivity between loopback interfaces is broken; the LDP session is usually established between loopback interfaces of adjacent LSRs.
- Verification:
  - Verify connectivity with the extended **ping** command.

© 2014 Cisco and/or its affiliates. All rights reserved.

Symptom: LDP neighbors are exchanging hello packets, but the LDP session is never established.

Solution: Check the reachability of the loopback interfaces, because they are typically used to establish the LDP session. Make sure that the loopback addresses are exchanged via the IGP that is used in the network.

https://t.me/learningnets

Label Distribution Issues

A typical example of an LDP issue is a situation when labels are allocated, but not distributed.

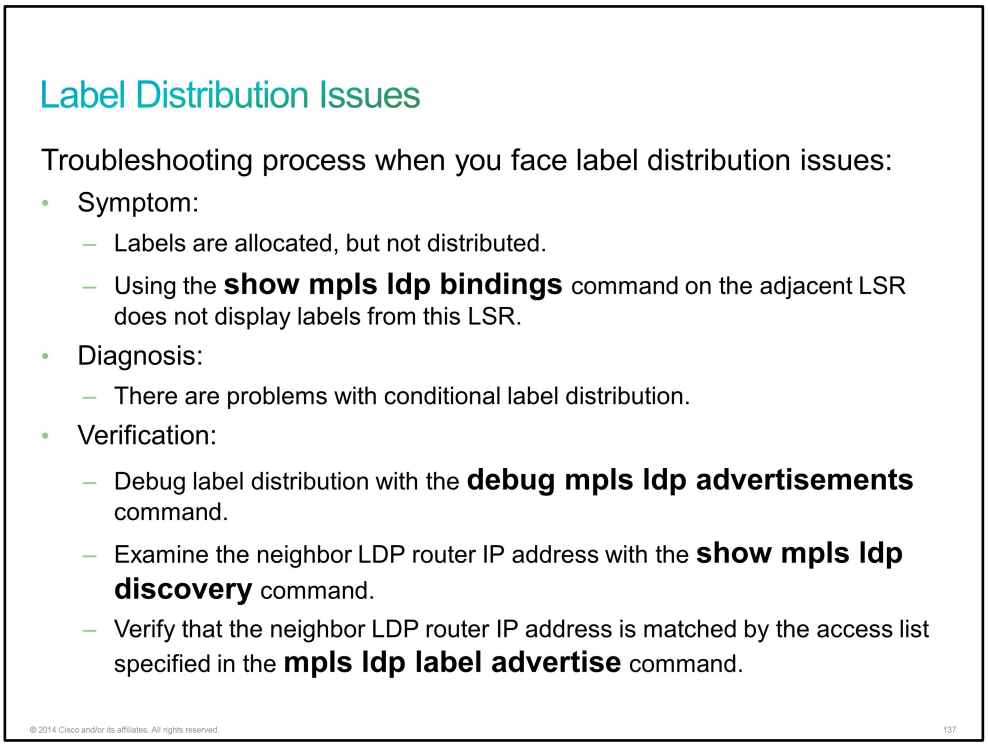

Symptom: Labels are generated for local routes on one LSR but are not received on neighboring LSRs.

Solution: Check whether conditional label advertising is enabled and verify that both access lists that are used with the command.

https://t.me/learningnets

Intermittent MPLS Failures After Interface Failure

A typical example of an MPLS issue is when the overall MPLS connectivity in a router intermittently breaks after an interface failure.

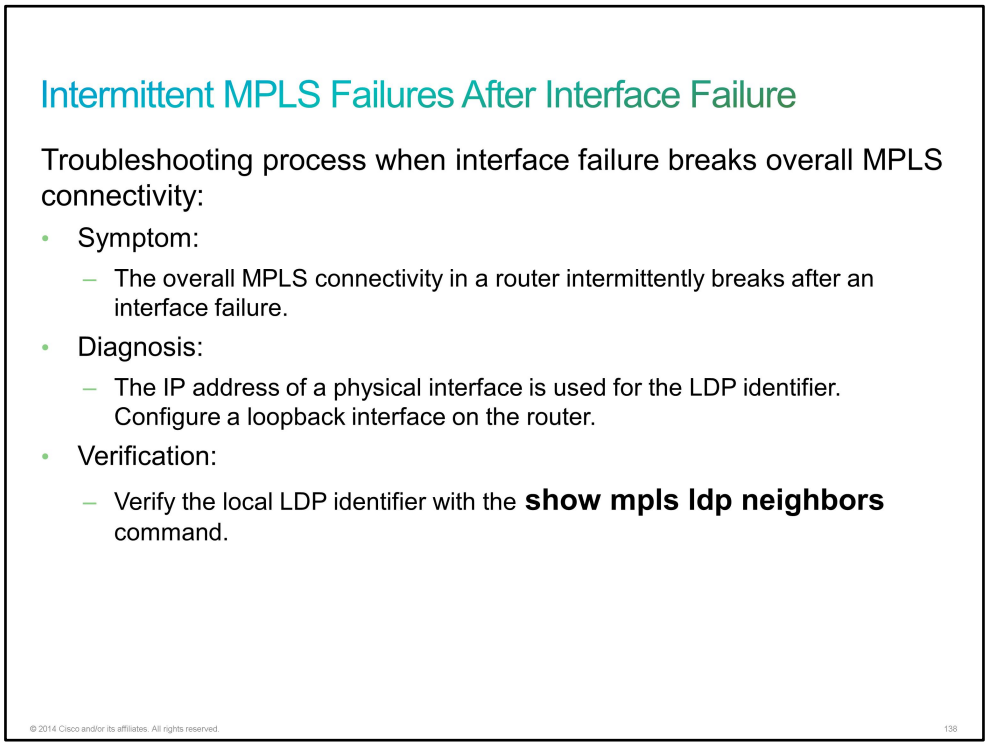

Symptom: MPLS connectivity is established, labels are exchanged, and packets are labeled and forwarded. However, an interface failure can sporadically stop an MPLS operation on unrelated interfaces in the same router.

Details: LDP sessions are established between IP addresses that correspond to the LDP router ID. If the LDP router ID has not been manually configured, the LDP router ID is assigned using the algorithm that is also used to assign an OSPF or a BGP router ID.

This algorithm selects the highest IP address of an active interface if there are no loopback interfaces configured on the router. If that interface fails, the LDP router ID is lost and the TCP session that is carrying the LDP data is torn down, resulting in loss of all neighbor-assigned label information.

The symptom can be easily verified with the **show mpls ldp neighbors** command, which displays the local and remote LDP router ID. Verify that both of these IP addresses are associated with a loopback interface.

Solution: Manually configure the LDP router ID, referencing a loopback interface that is reachable by the IGP.

https://t.me/learningnets

Packet Propagation Issues

A typical example of an MPLS issue is when large packets are not propagated across the network.

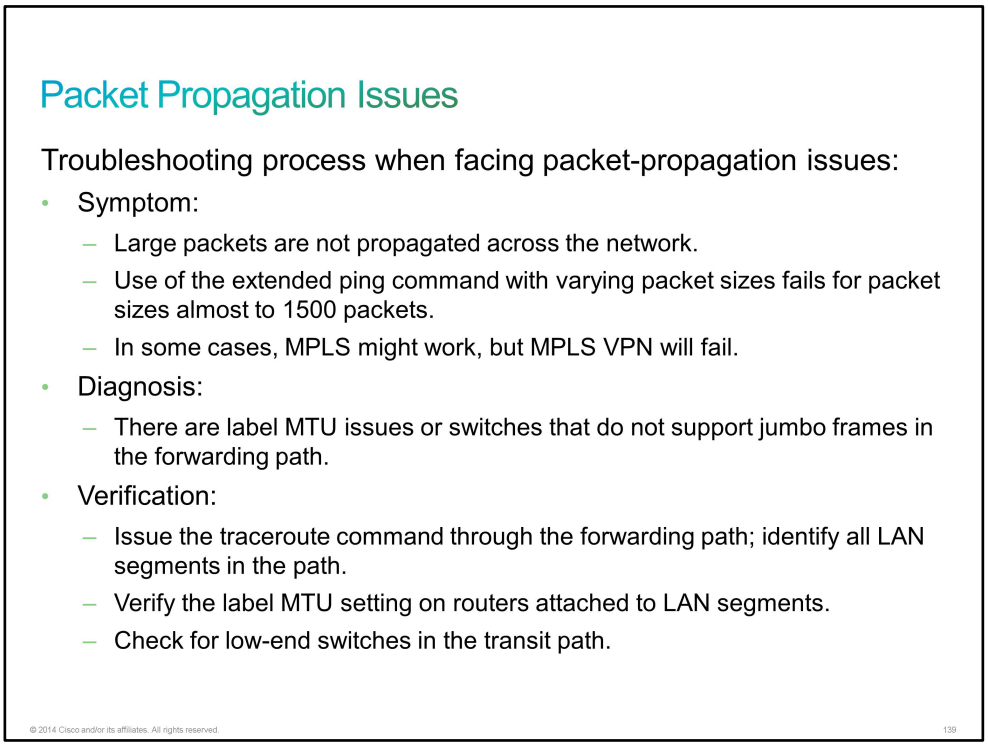

Symptom: Packets are labeled and sent, but they are not received on the neighboring router. A LAN switch between the adjacent MPLS-enabled routers may drop the packets if it does not support jumbo frames. In some cases, MPLS might work, but MPLS VPN will fail.

Solution: Change the MPLS MTU size, taking into account the maximum number of labels that may appear in a packet.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- On Cisco IOS XR platforms, MPLS is enabled by enabling LDP on each interface. On Cisco IOS and IOS XE platforms, MPLS is enabled on each interface.
- It is recommended to manually set the router ID.
- To enable MPLS on the Cisco IOS XR router, first enter MPLS LDP configuration mode and then list the interfaces.
- Label switching increases the maximum MTU requirements on an interface, because of the additional label header.
- To configure the maximum packet size or MTU size on an MPLS interface, use the **mpls mtu** command in interface configuration mode.
- By default, IP TTL is copied into the MPLS label at label imposition, and the MPLS label TTL is copied into the IP TTL at label removal.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- If TTL propagation is disabled, the TTL value is not copied into the label header. Instead, the label TTL field is set to 255.
- When LDP session protection is configured, LDP initiates backup targeted hellos automatically for neighbors for which primary link adjacencies already exist.
- Graceful restart is a way to recover from signaling and control plane failures without impacting forwarding.
- LDP IGP synchronization synchronizes LDP and IGP so that IGP advertises links with regular metrics only when MPLS LDP is converged on that link.
- To enable LDP on many interfaces, IGP autoconfiguration allows you to automatically configure LDP on all interfaces that are associated with a specified OSPF or IS-IS interface.
- LDP outbound label filtering performs outbound filtering for local label advertisement, for one or more prefixes, to one or more peers.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Summary (Cont.)

- The ability to filter remote bindings for a defined set of prefixes is also referred to as LDP inbound label filtering.
- You can use various **show** commands to monitor MPLS.
- When debugging MPLS and LDP in production environments, use the **debug** commands with extreme caution.
- Standard ping and traceroute tools can be used in MPLS environments to test reachability.
- Special MPLS ping and MPLS traceroute were designed for monitoring and troubleshooting MPLS LSPs.
- If an LDP session does not come up, verify if MPLS is enabled on the neighboring router.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 8: Module Summary

This topic summarizes the key points that were discussed in this module.

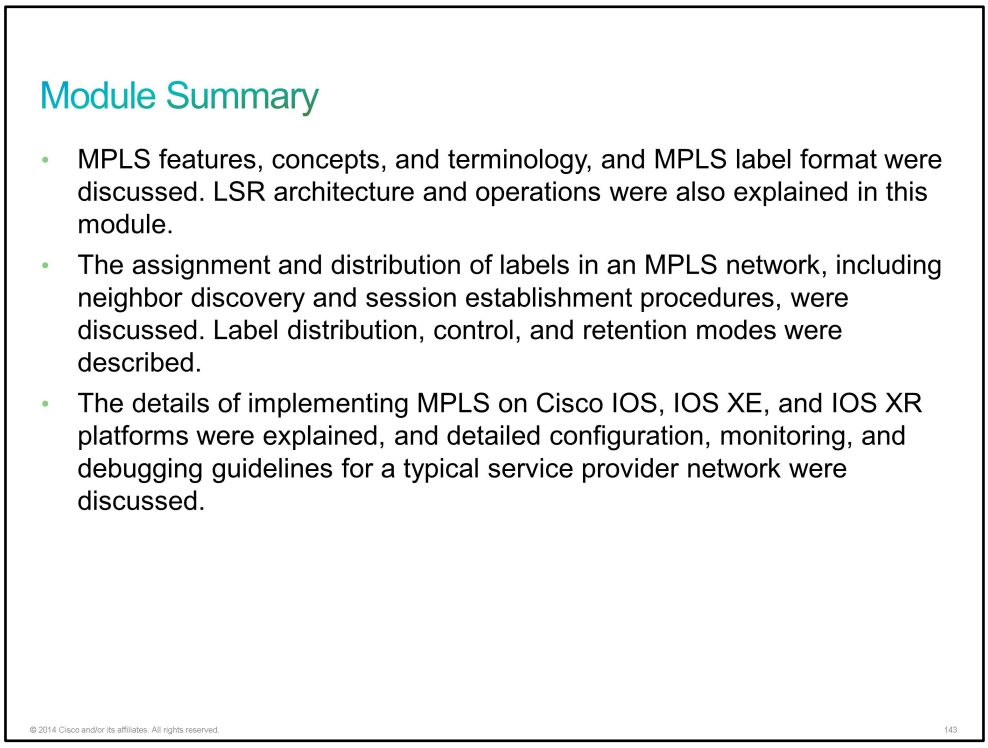

This module explained the features of MPLS compared with those of traditional hop-by-hop IP routing. MPLS concepts and terminology, along with MPLS label format and LSR architecture and operations, were explained in this module. The module also described the assignment and distribution of labels in an MPLS network, including neighbor discovery and session establishment procedures. Label distribution, control, and retention modes were also covered.

The module also explained the process for implementing MPLS on Cisco IOS, IOS XE, and IOS XR platforms, giving detailed configuration, monitoring, and debugging guidelines for a typical service provider network.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 9: Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

1. What are three foundations of traditional IP routing? (Choose three.) (Source: Introducing MPLS)

   A. Routing protocols are used on all devices to distribute routing information.
   B. Regardless of protocol, routers always forward packets based on only the IP destination address (except for using PBR).
   C. Routing lookups are performed on every router.
   D. Routing is performed by assigning a label to an IP destination.
   E. Routing lookup is performed only at the first router on the path.

2. Which three statements about MPLS are true? (Choose three.) (Source: Introducing MPLS)

   A. MPLS uses labels to forward packets.
   B. MPLS works only in IP networks.
   C. MPLS labels can correspond to a Layer 3 destination address, QoS, source address, or Layer 2 circuit.
   D. MPLS does not require a routing table lookup on core routers.
   E. MPLS performs routing lookup on every router.

3. How many bits does the MPLS label field consist of? (Source: Introducing MPLS)

   A. 64 bits
   B. 32 bits
   C. 16 bits
   D. 8 bits

https://t.me/learningnets

4. Which two statements about LSRs are true? (Choose two.) (Source: Introducing MPLS)

   A. An edge LSR is a device that inserts labels on packets or removes labels and forwards packets based on labels.
   B. An LSR is a device that primarily labels packets or removes labels.
   C. An LSR is a device that forwards packets based on labels.
   D. An end LSR is a device that primarily inserts labels on packets or removes labels.

5. Which two statements about RSVP are true? (Choose two.) (Source: Introducing MPLS)

   A. RSVP is used to create an LSP tunnel.
   B. RSVP propagates labels for TE tunnels.
   C. RSVP assigns labels for TE tunnels.
   D. RSVP is not used to create an LSP tunnel.

6. In MPLS VPN networks, which statement is true? (Source: Introducing MPLS)

   A. Labels are propagated via LDP or TDP.
   B. Next-hop addresses, instead of labels, are used in an MPLS VPN network.
   C. Labels are propagated via MP-BGP.
   D. Two labels are used; the top label identifies the VPN, and the bottom label identifies the egress router.

7. Which two statements about interactions between MPLS applications are true? (Choose two.) (Source: Introducing MPLS)

   A. The forwarding plane is the same for all applications.
   B. Differences exist in the forwarding plane depending on the MPLS application.
   C. The control plane is the same for all applications.
   D. Differences exist in the control plane depending on the MPLS application.

8. What does per-platform label space require? (Source: Label Distribution Protocol)

   A. It requires only one LDP session.
   B. It requires one session per interface.
   C. It requires multiple sessions for parallel links.
   D. “Per-platform” is not a proper term in MPLS terminology.

9. Which two well-known port numbers does LDP use? (Choose two.) (Source: Label Distribution Protocol)

   A. TCP 646
   B. UDP 646
   C. TCP 464
   D. UDP 464

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

10. Which three pieces of information are contained in the LFIB? (Choose three.) (Source: Label Distribution Protocol)

    A. locally generated label
    B. outgoing label
    C. local address
    D. next-hop address
    E. destination IP network

11. When an IP packet is to be label-switched as it traverses an MPLS network, which table is used to perform the label switching? (Source: Label Distribution Protocol)

    A. LIB
    B. FIB
    C. FLIB
    D. LFIB

12. Which statement is correct? (Source: Label Distribution Protocol)

    A. An IP forwarding table resides on the data plane; LDP runs on the control plane; and an IP routing table resides on the data plane.
    B. An IP forwarding table resides on the data plane; LDP runs on the control plane; and an IP routing table resides on the control plane.
    C. An IP forwarding table resides on the control plane; LDP runs on the control plane; and an IP routing table resides on the data plane.
    D. An IP forwarding table resides on the control plane; LDP runs on the control plane; and an IP routing table resides on the control plane.

13. Which two tables contain label information? (Choose two.) (Source: Label Distribution Protocol)

    A. LIB
    B. main IP routing table
    C. FLIB
    D. LFIB

14. Which two statements about LSPs are correct? (Choose two.) (Source: Label Distribution Protocol)

    A. LSPs are bidirectional.
    B. LSPs are unidirectional.
    C. LDP advertises labels for the entire LSP.
    D. LDP advertises labels only for individual segments in the LSP.

15. Which statement about TTL propagation being disabled is correct? (Source: Label Distribution Protocol)

    A. The label TTL is copied back into the IP TTL.
    B. The IP TTL is copied back into the TTL of the label.
    C. The IP TTL is not copied back into the TTL of the label.
    D. TTL label propagation can not be disabled.

https://t.me/learningnets

16. Upon a link failure, which three tables are updated to reflect the failed link? (Choose three.) (Source: Label Distribution Protocol)

    A. LIB
    B. LFIB
    C. FIB
    D. FLIB
    E. BIL

17. Which statement correctly describes how a link failure is handled in an MPLS network? (Source: Label Distribution Protocol)

    A. Overall convergence depends on LDP.
    B. Overall convergence depends on the IGP that is used.
    C. Upon a link failure, only LDP convergence is affected.
    D. Upon a link failure, only the IGP convergence is affected.

18. Upon a link recovery, which three tables are updated to reflect the failed link? (Choose three.) (Source: Label Distribution Protocol)

    A. LFIB
    B. FLIB
    C. FIB
    D. LIB
    E. BIL

19. Which statement correctly describes convergence in MPLS network after a link failure has occurred and been restored? (Source: Label Distribution Protocol)

    A. MPLS convergence occurs after IGP convergence.
    B. MPLS convergence occurs before IGP convergence peer-to-peer.
    C. If a failure occurs with the IGP, MPLS convergence is not affected.
    D. If a failure occurs with the IGP, MPLS will not be able to converge after the IGP failure has been corrected unless the MPLS process is bounced.

20. What is another name for topology-driven switching? (Source: Label Distribution Protocol)

    A. Cisco Express Forwarding
    B. fast switching
    C. cache switching
    D. process switching

21. If IP TTL propagation is not allowed, what is the value that is placed in the MPLS header? (Source: Implementing MPLS in the Service Provider Core)

    A. 0
    B. 1
    C. 254
    D. 255

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

22. What needs to be configured to specify which neighbors would selectively receive label advertisements? (Source: Implementing MPLS in the Service Provider Core)

    A. Controlled label distribution needs to be configured.
    B. Conditional label distribution needs to be configured.
    C. Unsolicited label distribution needs to be configured.
    D. No configuration is necessary; all neighbors will receive all labels.

23. Which command is used to display the contents of the LIB table? (Source: Implementing MPLS in the Service Provider Core)

    A. **show mpls ldp labels**
    B. **show mpls ldp bindings**
    C. **show mpls ldp neighbors**
    D. **show mpls forwarding-table**

24. Which is the correct command to enable MPLS in Cisco IOS Software? (Source: Implementing MPLS in the Service Provider Core)

    A. Router(config)# **ip mpls**
    B. Router(config-if)# **ip mpls**
    C. Router(config)# **mpls ip**
    D. Router(config-if)# **mpls ip**

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. A, B, C
2. A, C, D
3. B
4. A, C
5. A, B
6. C
7. A, D
8. A
9. A, B
10. A, B, D
11. D
12. B
13. A, D
14. B, D
15. C
16. A, B, C
17. B
18. A, C, D
19. A
20. A
21. D
22. B
23. B

https://t.me/learningnets

24. D

https://t.me/learningnets

Module 2: MPLS Traffic Engineering

Introduction

This module on MPLSTE technology discusses the requirement for TE in modern service provider 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 RSVP. Some of the advanced features of TE, such as autobandwidth and guaranteed bandwidth, are introduced as well.

This module discusses the requirement for implementing MPLS TE, and includes the 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. Configurations are shown for various Cisco platforms.

Upon completing this module, you will be able to:

- Describe the concepts that allow service providers to map traffic through specific routes to optimize network resources, especially bandwidth
- Describe the details of link attribute propagation with an IGP and constraint-based path computation
- Describe MPLS TE commands for the implementation of MPLS traffic tunnels
- Describe the MPLS TE commands for link and node protection

https://t.me/learningnets

https://t.me/learningnets

Lesson 1: Introducing 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.

Upon completing this lesson, you will be able to:

- Describe basic TE concepts
- Describe TE motivations
- Describe business drivers for TE
- Describe the problem that TE solves: network congestion over a long time period
- Describe TE with a Layer 2 overlay model
- Describe an example of TE with a Layer 2 overlay model
- Describe the drawbacks of TE with a Layer 2 overlay model
- Describe a Layer 3 routing model without TE
- Describe TE with a Layer 3 routing model
- Describe TE with the MPLS TE model
- Describe the basic concept of MPLS TE traffic tunnels
- Describe MPLS TE traffic tunnels attributes
- Describe the link-resource attributes

https://t.me/learningnets

- Describe constraint-based path computation
- Describe an example of constraint-based path computation
- Describe the MPLS TE process
- Describe the role of RSVP in path setup procedures
- Describe the path setup procedures using RSVP
- Describe the three methods to forward traffic to a tunnel
- Describe the Autoroute feature
- Describe an example of Autoroute feature

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Traffic Engineering Concepts

Traffic engineering is widely used, especially in voice telephony. With TE, traffic is measured and analysed. A statistical model is then applied to the traffic pattern to make a prognosis and estimates of traffic matrices.

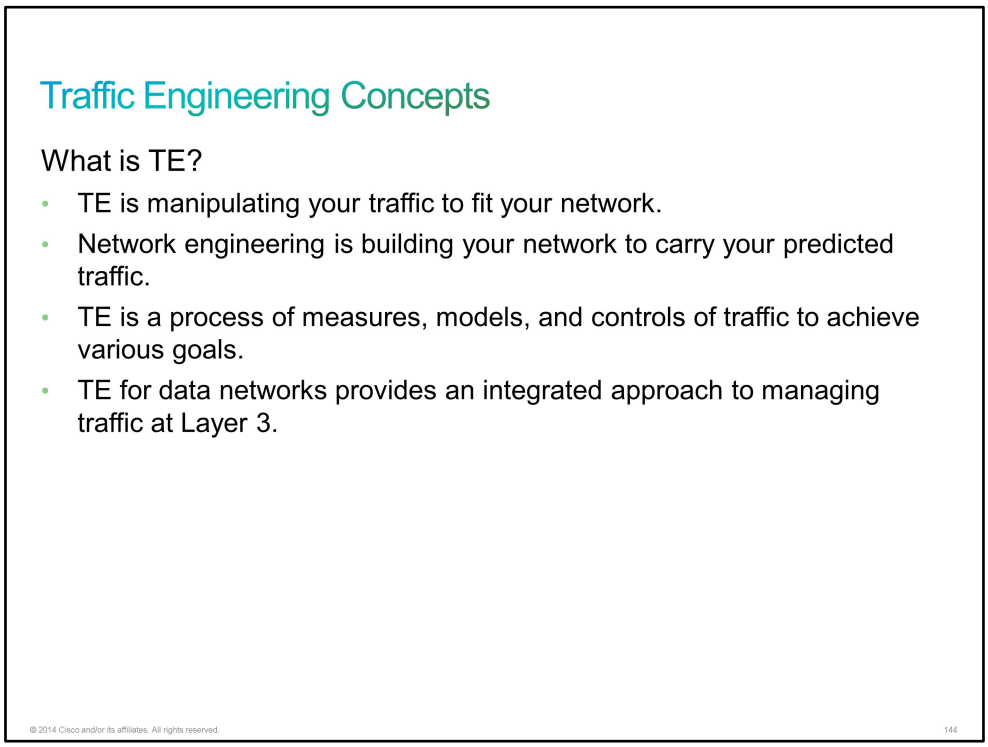

TE can be contrasted with network engineering in the following way:

- TE is manipulating your traffic to fit your network.
- Network engineering is building your network to carry your predicted traffic.

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

In the data communications world, TE provides an integrated approach to engineering traffic at Layer 3 in the 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 the uncongested 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.

https://t.me/learningnets

Traffic Engineering Motivations

The main motivation for using TE is cost reduction.

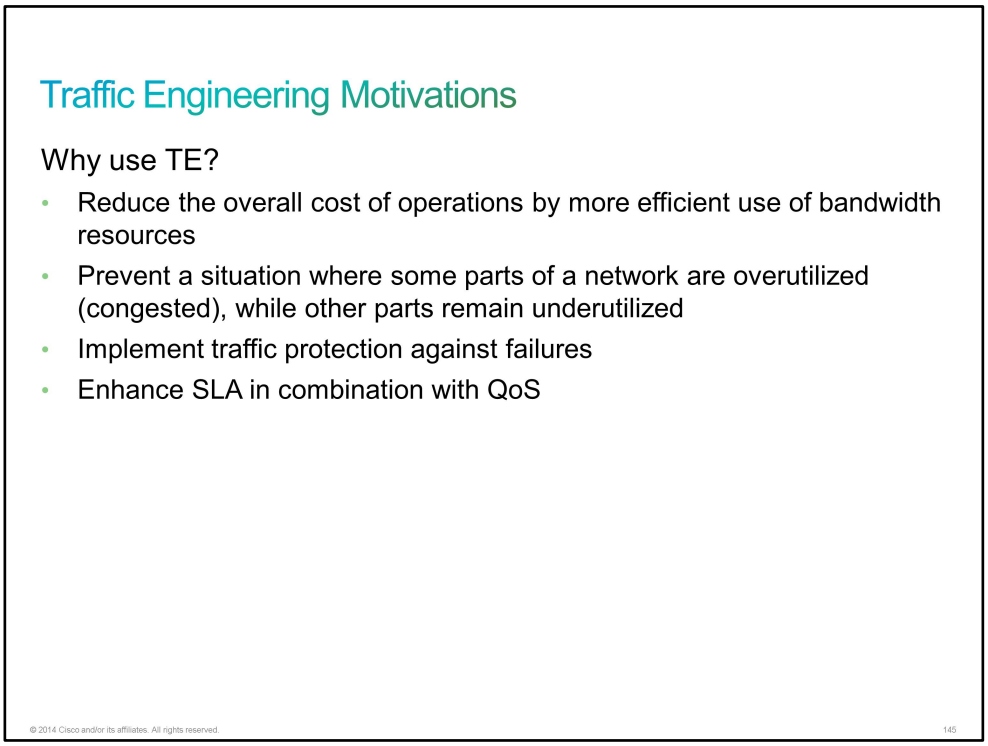

WAN connections are an expensive item in the service provider budget. The cost saving that results from a more efficient use of resources, will help to reduce the overall cost of operations. Additionally, more efficient use of bandwidth resources means that a service provider can avoid a situation where some parts of a 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 SLAs.

https://t.me/learningnets

Business Drivers for Traffic Engineering

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

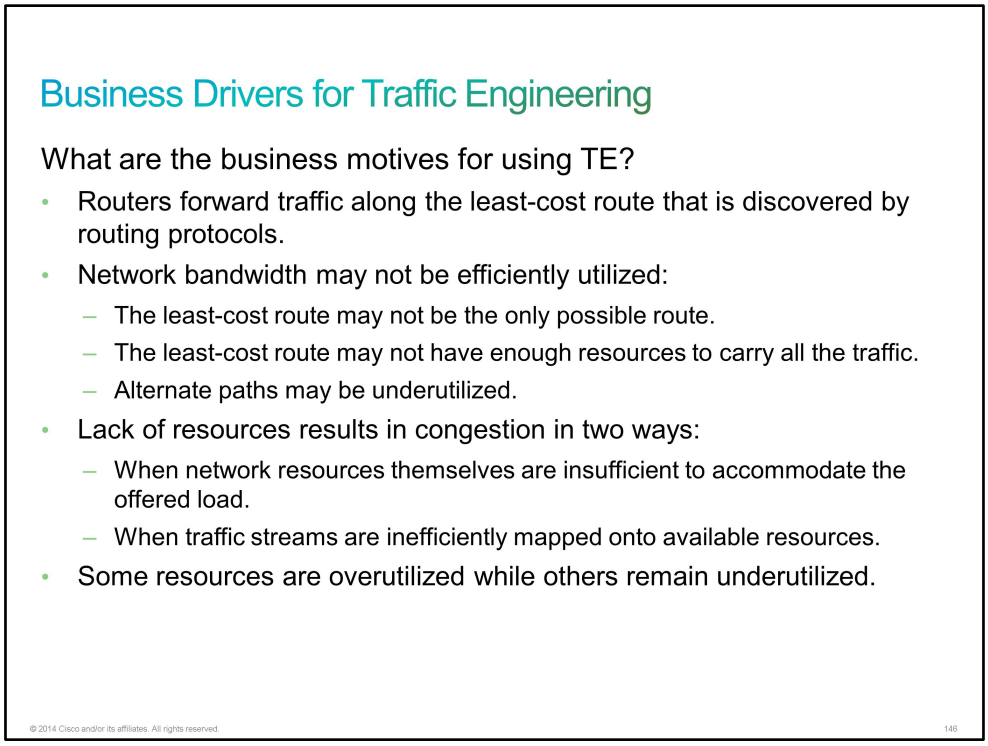

In a Layer 3 routing network, packets are forwarded hop-by-hop. In each hop, the destination address of the packet is used to make a routing table lookup. The routing tables are created by an IGP, which finds the least-cost route, according to its metric, to each destination in the network.

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 happen 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 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.

To provide better network performance within the 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. If there is a failure to any of the links, all traffic is forwarded over the remaining link, which then, of course, becomes overutilized.

https://t.me/learningnets

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

Network congestion, caused by too much traffic and too few 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 constitutes better use of the money invested in the network.

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

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Congestion Avoidance and Traffic Engineering

One of the problems that TE solves is persistent network congestion over a long period of time.

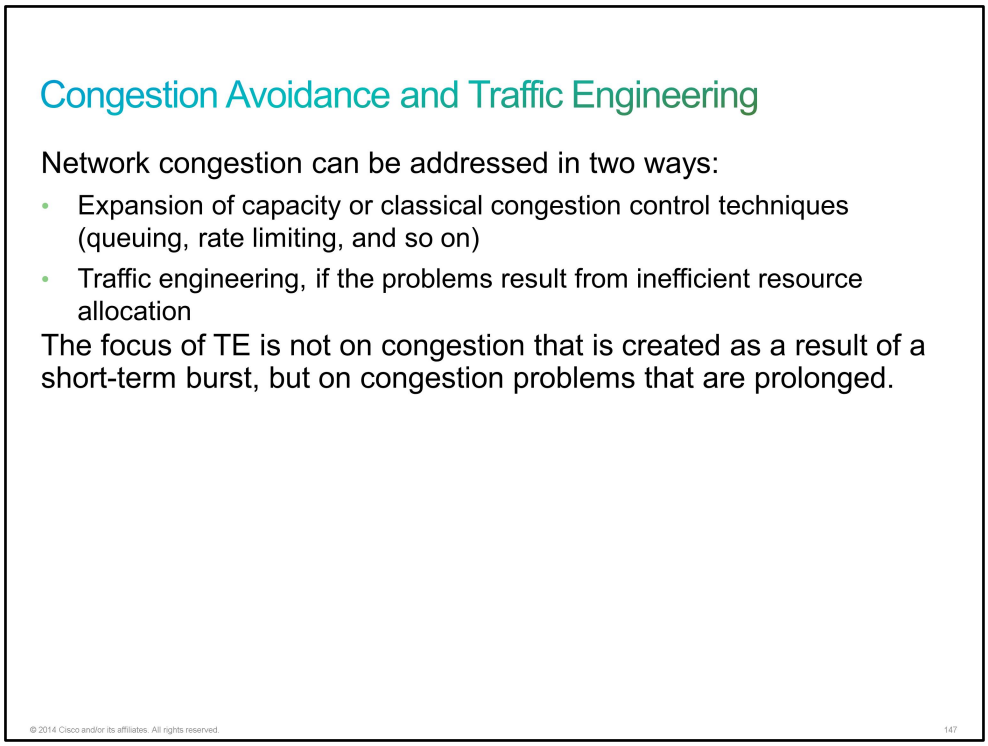

TE does not solve temporary network congestion that is caused by traffic bursts. This type of problem is better managed by an expansion of capacity or by classic techniques such as various queuing algorithms, rate limiting, and intelligent packet dropping. TE does not solve problems when the network resources themselves are insufficient to accommodate the required load.

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 prolonged congestion, possibly continuously, while other parts of the network have spare capacity.

https://t.me/learningnets

Traffic Engineering with a Layer 2 Overlay Model

When using TE in the Layer 2 overlay model, the routers (Layer 3 devices) are overlaid on the Layer 2 topology.

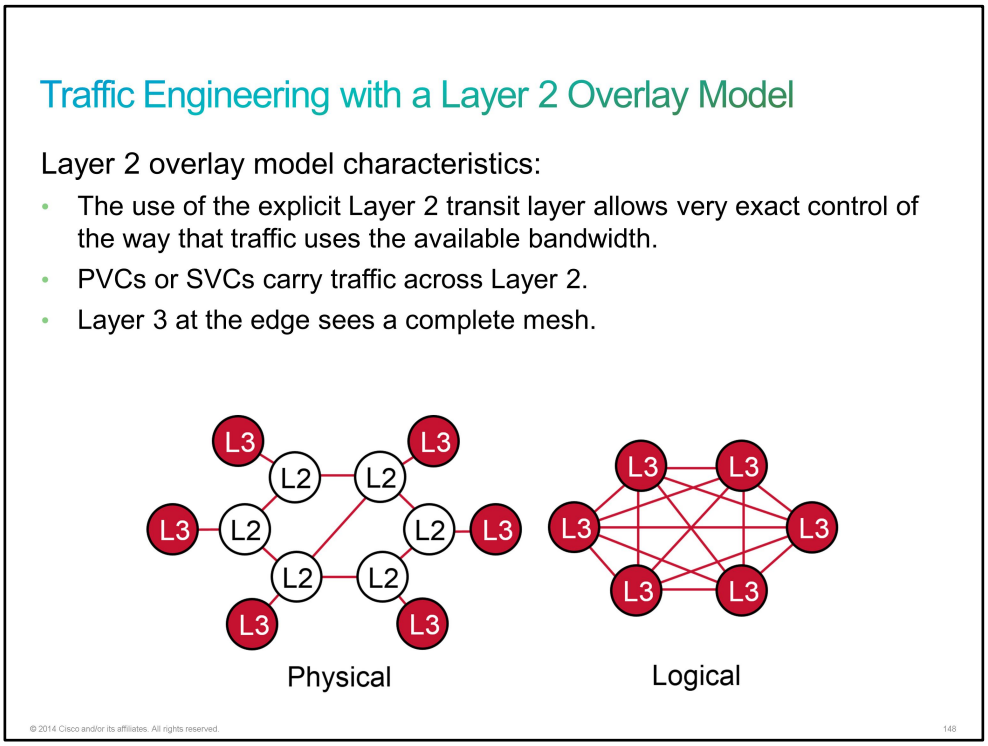

With an overlay model, the routers are not aware of the physical structure and the bandwidth that is available on the links. The IGP views the Layer 2 PVCs or SVCs as point-to-point links and makes its forwarding decisions accordingly.

All traffic engineering is done at Layer 2. PVCs are carefully engineered across the network, normally using an offline management system. SVCs are automatically established by using signaling, and their way across the Layer 2 network is controlled by integrated path determination, such as the PNNI protocol.

In the Layer 2 overlay model, PVCs or SVCs carry the traffic across the network. With a Frame Relay network, PVC setup is usually 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 may use the same type of tools as Frame Relay for PVC establishment, or may use the SVC approach, where routers use a signaling protocol to dynamically establish an SVC.

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 is likely to use 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 very exact control of the way that traffic uses the available bandwidth.

https://t.me/learningnets

Traffic Engineering with a Layer 2 Overlay Model: Example

TE in Layer 2, using the overlay model, allows detailed decisions about which link should be used to carry various traffic patterns.

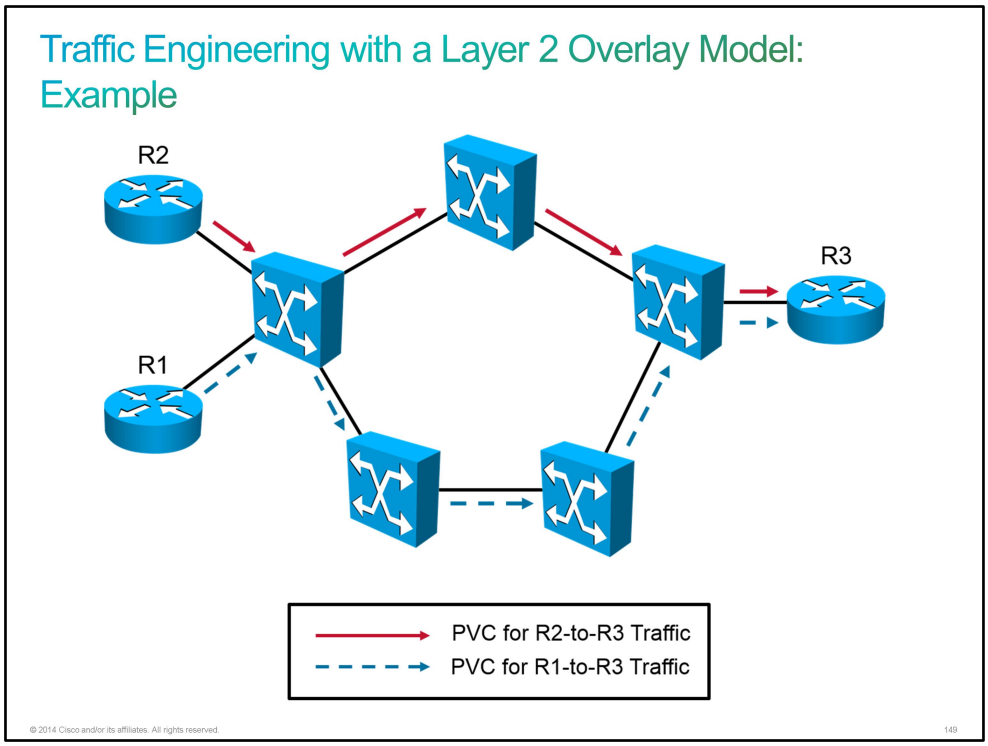

In this example, traffic from R2 to R3 uses the top PVC (solid arrows), which takes the shortest path using the upper transit switch.

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

https://t.me/learningnets

Drawbacks of the Layer 2 Overlay Solution

The Layer 2 overlay model has certain drawbacks.

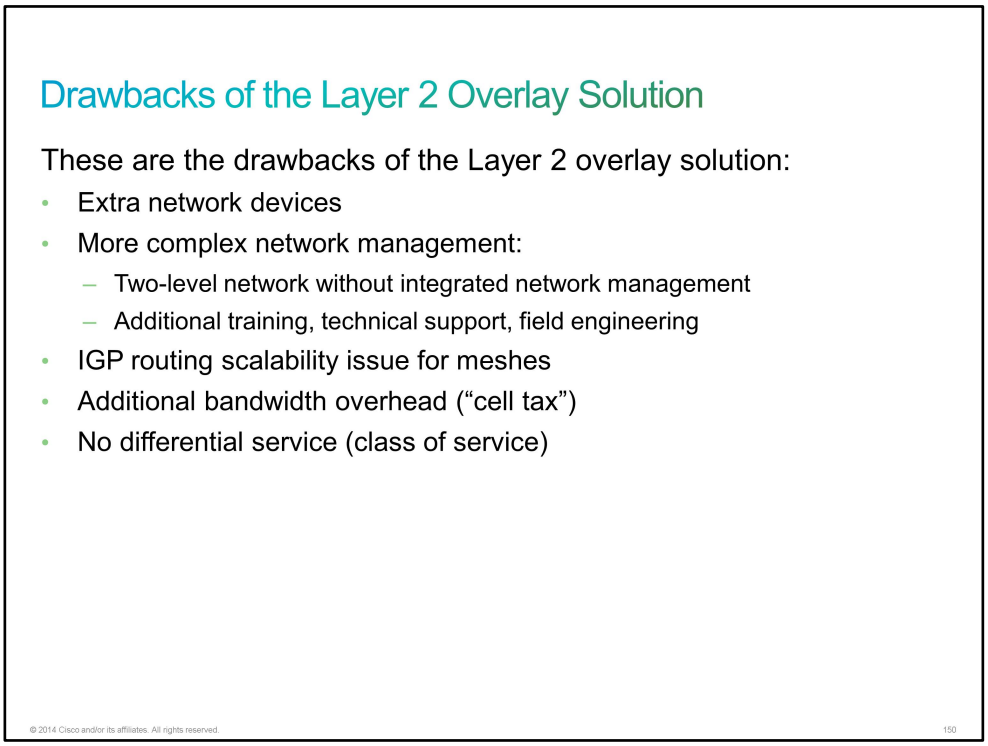

Using the Layer 2 overlay model has several drawbacks:

- 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 that manage several other tasks, and support TE as well. At the same time, the router network (Layer 3), with its IGP and tuning parameters, must be managed. Both of these management tasks require trained staff for technical support and in the field.
- The Layer 3 network must be highly meshed 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 due to the large number of IGP neighbors.
- 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 the IP information is unknown to the Layer 2 devices (ATM or Frame Relay switches). 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.

https://t.me/learningnets

Layer 3 Routing Model Without Traffic Engineering

If the same network topology is created using routers (Layer 3 devices), TE must be performed differently.

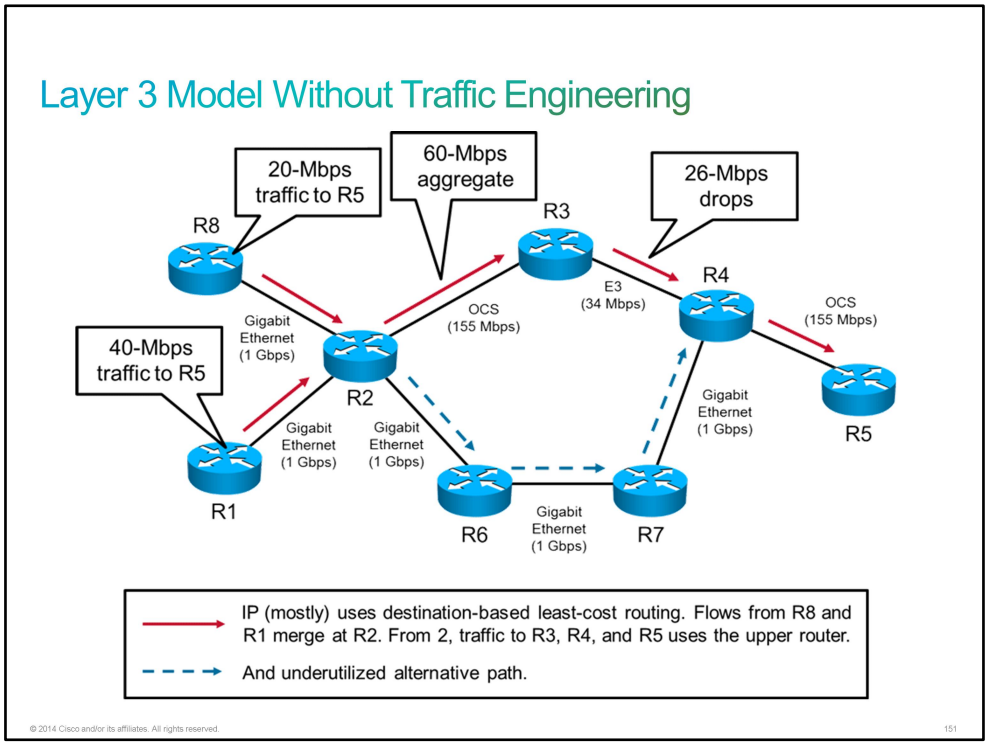

In the example here, if no traffic engineering is applied to the network, traffic from both R8 and R1 toward R5 will use the least-cost path (the upper path, which has one less hop). This flow may result in the overutilization of the path R2, R3, R4, R5, while the lower path R2, R6, R7, R4, R5 (with the one extra hop) will be underutilized.

Note	The assumption is that the routing protocol used in the example is IS-IS, which uses the default metric of 10 for each link, thus making the routing protocol metric to be the hop count.

https://t.me/learningnets

Traffic Engineering with a Layer 3 Routing Model

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

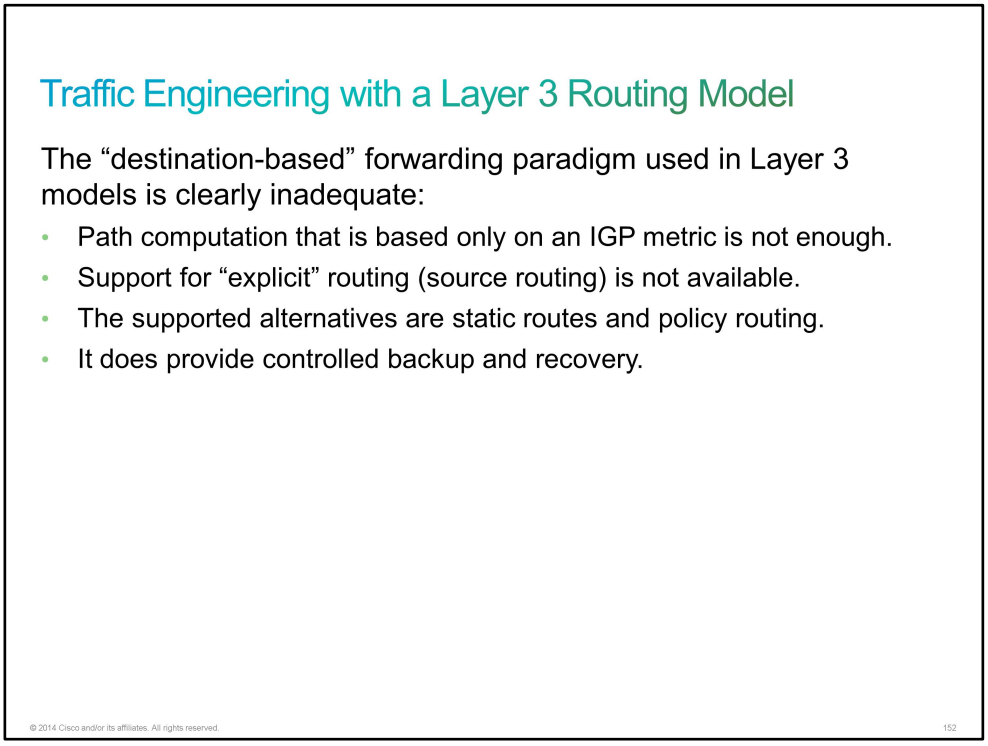

The IGP uses its metric to compute a single best way to reach each destination. There are problems with Layer 3 TE:

- 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 some traffic to take a different path than that of other traffic. However, static routing does not discriminate among various traffic flows based on the source. Static routing also restricts how redundancy in the network can be used, and it is not a scalable solution.
- PBR is able to discriminate among packet flows, based on the source, but it suffers from low scalability and the same static routing restrictions on using redundancy.

https://t.me/learningnets

Traffic Engineering with the MPLS TE Model

MPLS TE uses tunnels to control the data flow path. An MPLS TE tunnel is simply a collection of data flows that share some common attribute. The aim is to control the paths along which data flows, rather than relying simply on destination-based routing. The attribute could be that all traffic is sharing the same entry point to the network and the same exit point.

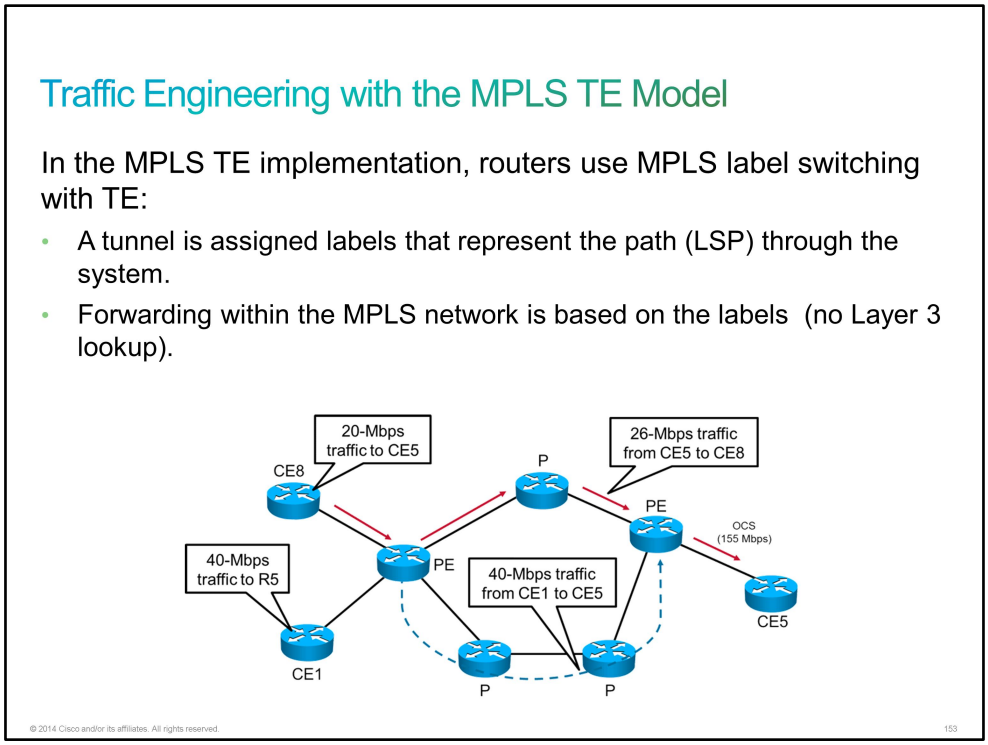

A TE tunnel maps onto an MPLS LSP. After the data flows and the TE tunnels are defined, MPLS technology is used to forward traffic across the network. Data is assigned an MPLS TE, which defines the route for traffic to take through the network. The packets that are forwarded under MPLS TE have a stack of two labels that are imposed by the ingress router. The topmost label identifies a specific LSP or TE tunnel 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 tunnel should do with the packet.

https://t.me/learningnets

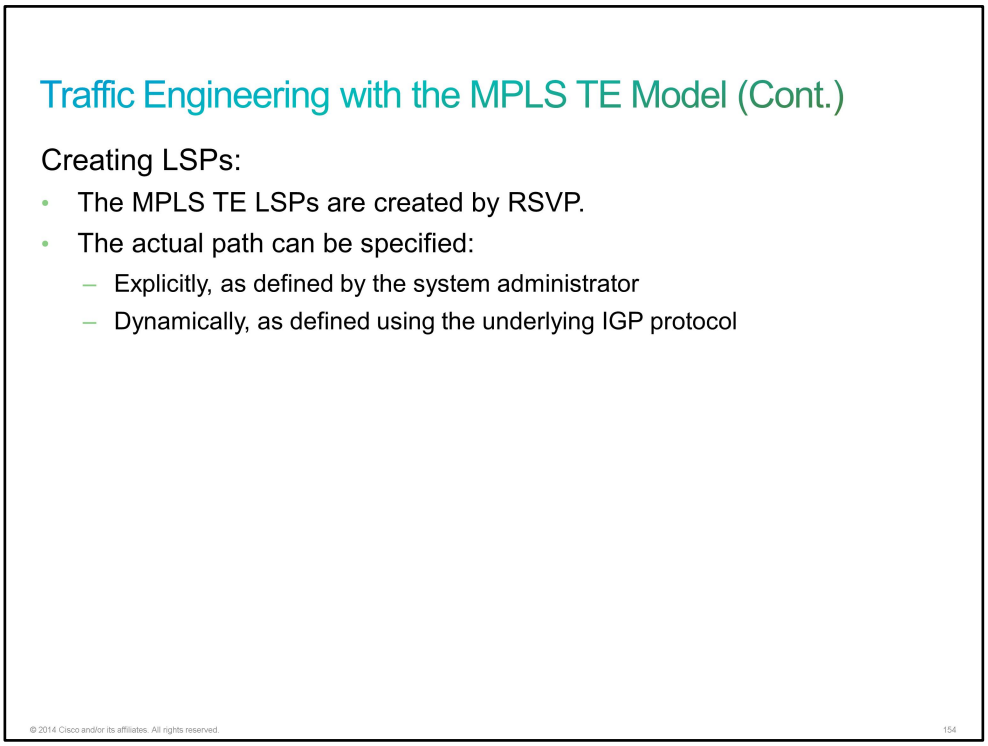

For MPLS TE, manual assignment and configuration of the labels can be used to create LSPs to tunnel the packets across the network on the desired path. However, to increase scalability, the RSVP is used to automate the procedure.

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

MPLS TE provides mechanisms equivalent to those described previously in this lesson, along 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 CBR that is used in MPLS TE, either IS-IS or OSPF with extensions is used to carry resource information, such as the 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 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 into tunnels results in the addition of entries in the FIB. The IP packets are directed into the MPLS TE tunnel by imposing the correct label stack.

https://t.me/learningnets

MPLS TE Traffic Tunnels

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

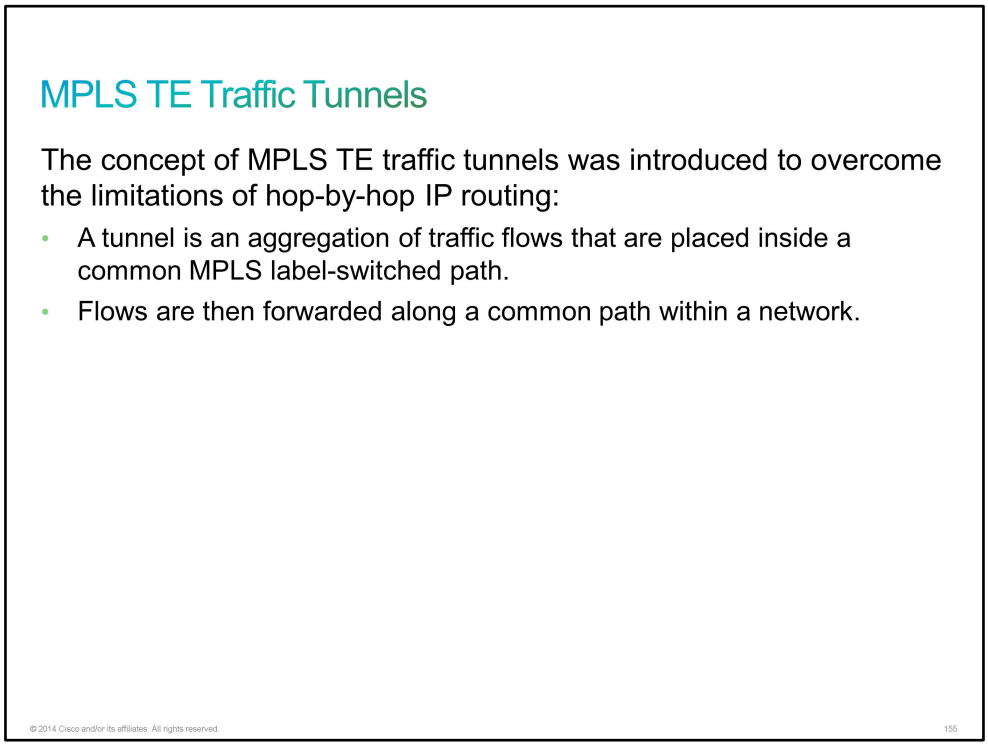

A traffic tunnel is simply a collection of data flows that share some common attribute:

- Most simply, this attribute might be the sharing of the same entry point to the network and the same exit point. In practice, in an ISP network, there is usually a definable data flow from the POPs, where the customers attach to the ISP network. There are also the 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-up home users. This preferential treatment 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.

https://t.me/learningnets

MPLS TE Traffic Tunnels (Cont.)

Class of service model types:

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

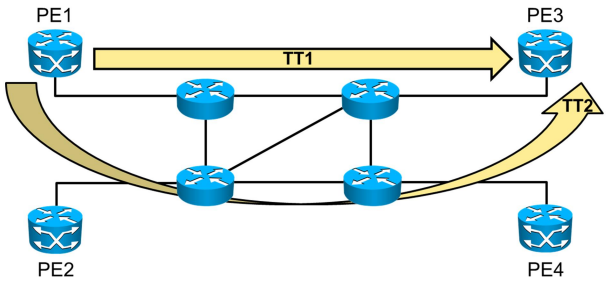

© 2014 Cisco and/or its affiliates. All rights reserved.

Defining traffic trunks (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, 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* PE1 *from* PE3.

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 router might sound like a very difficult administrative task. However, this is not usually the case for the following reasons:

- The routers that are identified as the tunnel headends are usually on the edge of the network. The traffic tunnels link these routers across the core of the network.
- In most networks, identifying the traffic flows is relatively easy, 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 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 star-like formation can also exist in many networks centering on the data center. This tendency is found in both ISP networks (providing web-hosting services) and enterprise networks.

After the data flows, and therefore the traffic tunnels are defined, the technology that they use to forward the data across the network is MPLS. Data that enters a traffic tunnel is assigned an MPLS LSP. The LSP defines the route that is taken through the network.

https://t.me/learningnets

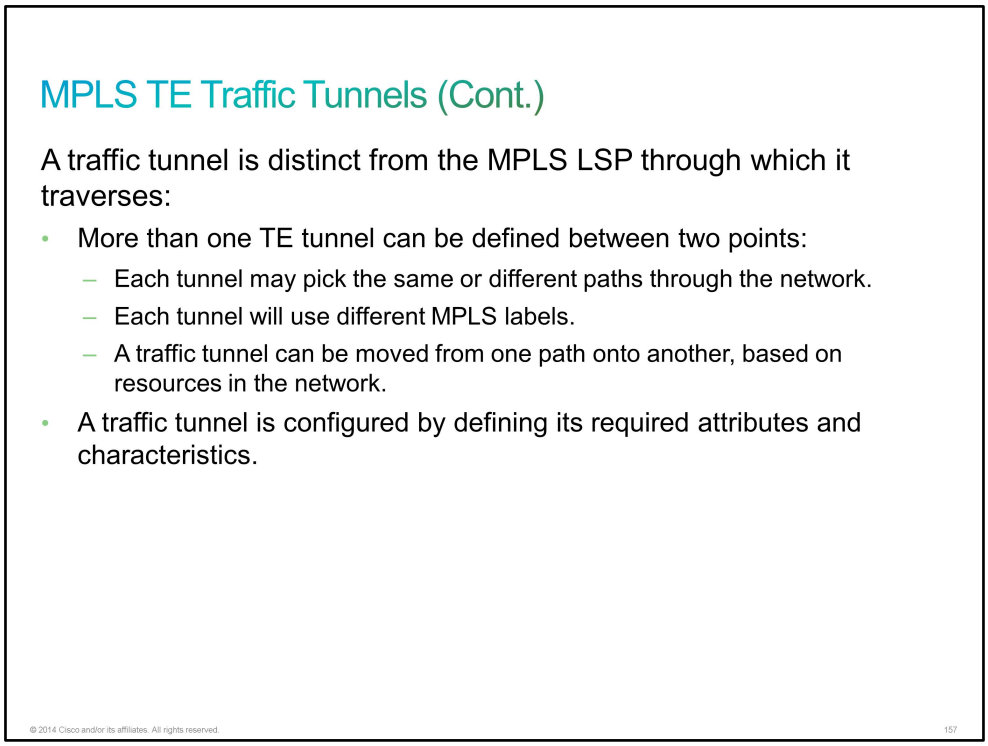

In two important ways, traffic tunnels are distinct from the MPLS LSPs that they use:

- There is a one-to-one mapping of traffic tunnels onto MPLS LSPs. Two tunnels may be defined between two points and may happen to pick the same path through the network. However, they will use different MPLS labels.
- 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.

Configuring 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 a traffic tunnel, the data might as well be left to route as it did previously, based only on destination information over the least-cost path.

https://t.me/learningnets

Traffic Tunnels: Attributes

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 that is in common with some other flow, there is nothing to define that data as part of a flow or group of flows.

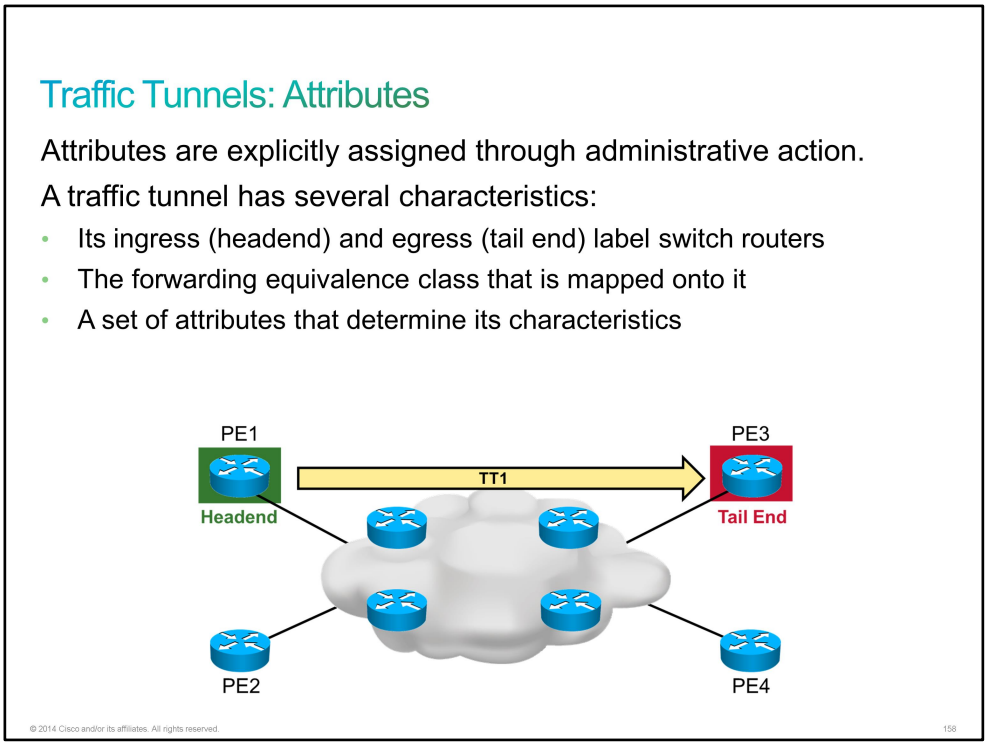

The traffic tunnel must include attributes that define the commonality between the data flows making up the tunnel. The attributes that characterize a traffic tunnel include the following:

- **Ingress and egress points:** These points are, fundamentally, the routers at the ends of the tunnel. They are the most basic level of commonality of 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, latency, and precedence requirements.
- **Class of data:** This attribute defines what data is part of this tunnel and what is not. This definition 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.

Note	MPLS TE setup is a control plane function.

https://t.me/learningnets

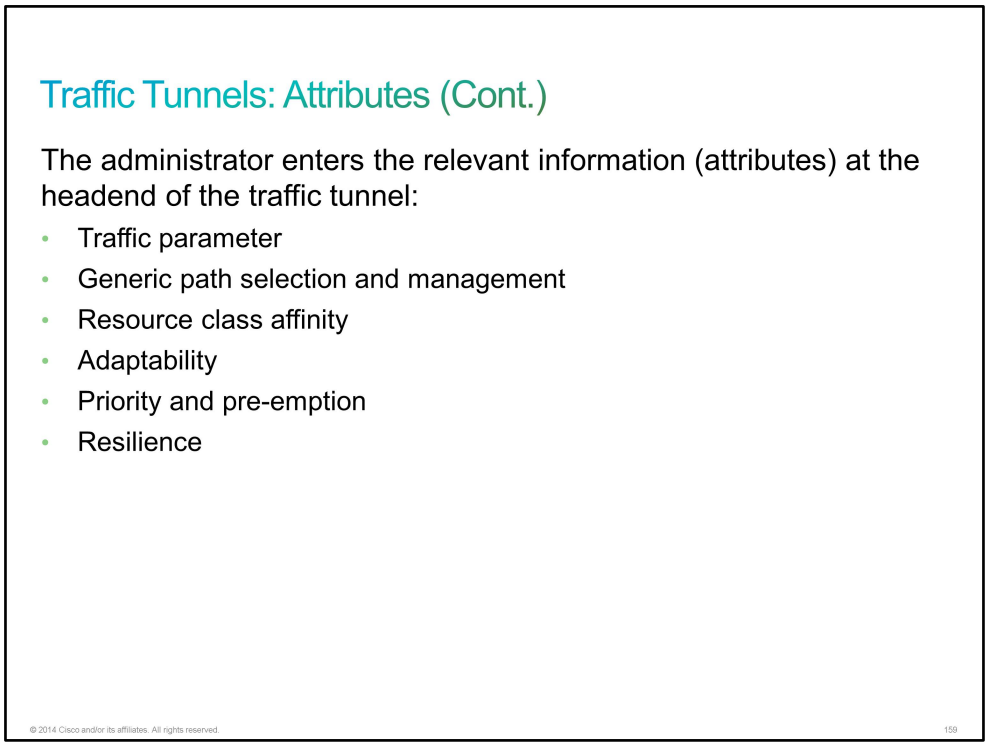

The general tunnel characteristics must be configured by the network administrator to create the tunnel. This configuration includes some or all of these attributes:

- **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, which is 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. 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 a 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 to a failure in the network. Does it attempt to reroute around failures, or not?

https://t.me/learningnets

Link Resource Attributes

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.

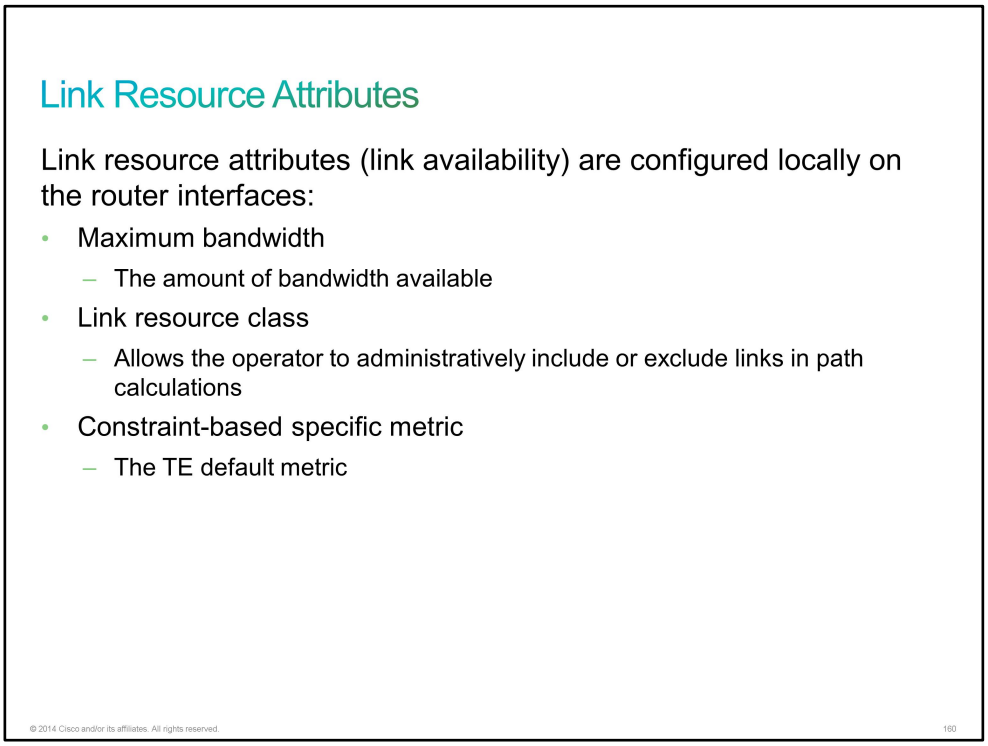

Headend router needs to be provided with this information:

- **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 whether a link may 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.

https://t.me/learningnets

Constraint-Based Path Computation

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).

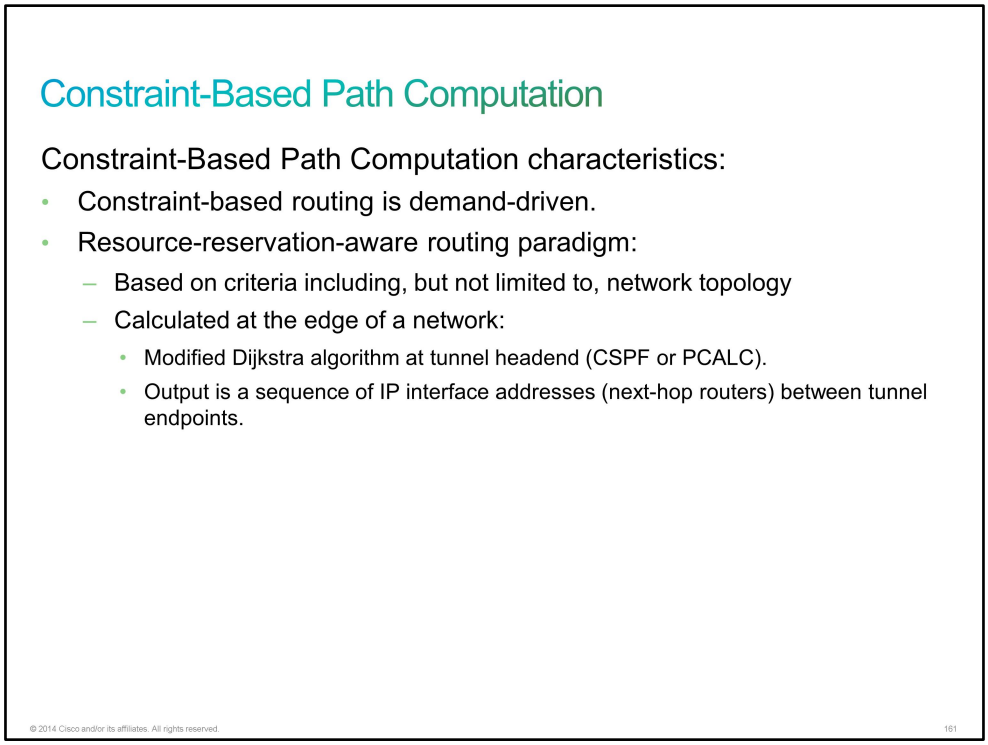

In traditional networks, the IGP calculates paths through the network, based on the network topology alone. Routing is destination-based, and all traffic to a given destination from a given source uses the same path through the network. That path is based simply on what the IGP regards as the “least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

MPLS TE employs CBR in which the path for a traffic flow is the shortest path that meets the resource requirements (constraints) of the traffic flow.

CSPF or PCALC is an extension of SPF algorithms. The path that is computed by using CSPF or PCALC is the shortest path fulfilling a set of constraints.

CBR behaves in these ways:

- It augments the use of link cost by also considering other factors, such as bandwidth availability or link attributes, when choosing the path to a destination.
- It 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 SPF that is used by IS-IS and OSPF, that considers other factors in addition to cost, such as bandwidth availability).
- It 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.

https://t.me/learningnets

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.

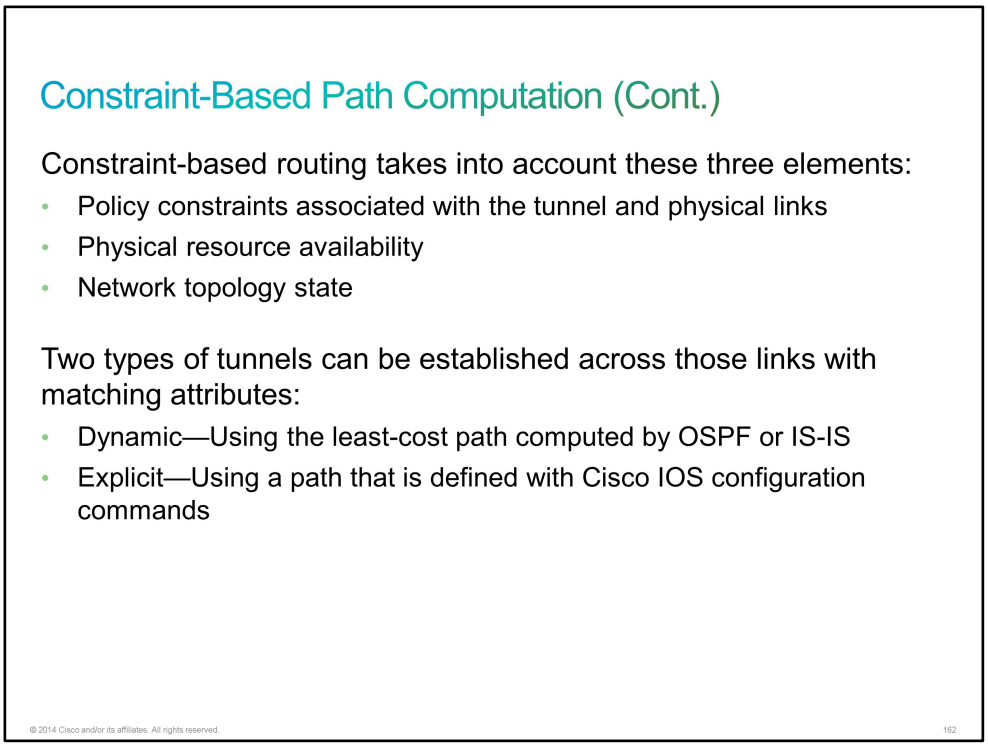

When choosing paths through the network, the CBR system takes into account these factors:

- 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 explicitly define the traffic tunnel and can also mix static and dynamic computation.

https://t.me/learningnets

Example of Constraint-Based Path Computation

An sample 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.

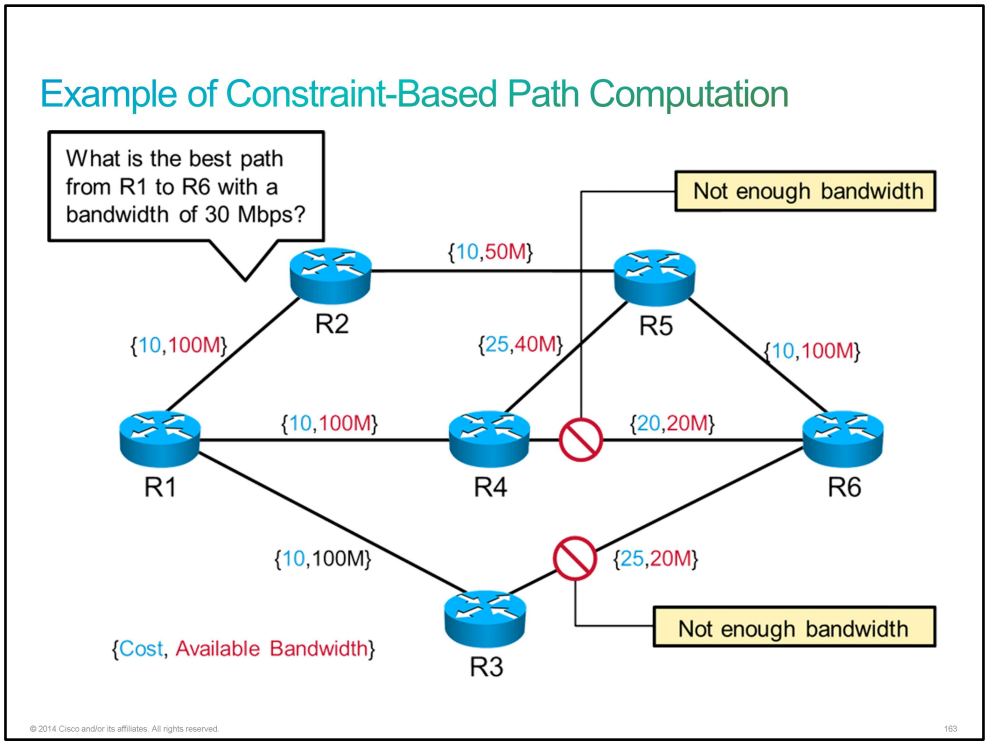

The requirement is to create a tunnel from R1 to R6 with a bandwidth of 30 Mbps.

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 R3-R6 has only 20 Mbps available as well, so no paths can be allocated via R3.

https://t.me/learningnets

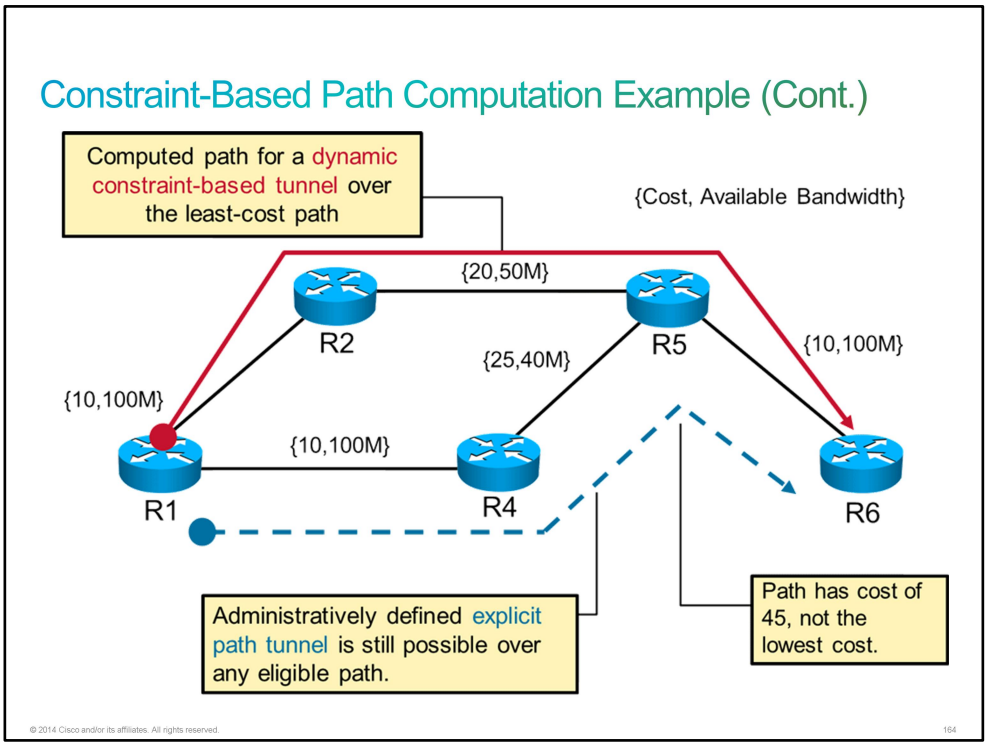

The diagram now shows only those links that can satisfy the requirement for 30 Mbps of available bandwidth.

Over this topology, two tunnel paths are shown:

- The top (solid arrow) path shows the result of a dynamic constraint-based path calculation. The calculation ignores any links that do not satisfy the bandwidth requirement (those shown in the previous figure but not shown here, such as the connections to R3) and then executes a CSPF calculation on what remains. This calculation has yielded the path R1-R2-R5-R6 with a path cost of 40.
- The network administrator has statically defined the bottom (dashed arrow) path (R1-R4-R5-R6). Had the administrator attempted to define a path that did not have the required free bandwidth, 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.

https://t.me/learningnets

MPLS TE Process

There are six TE processes to understand.

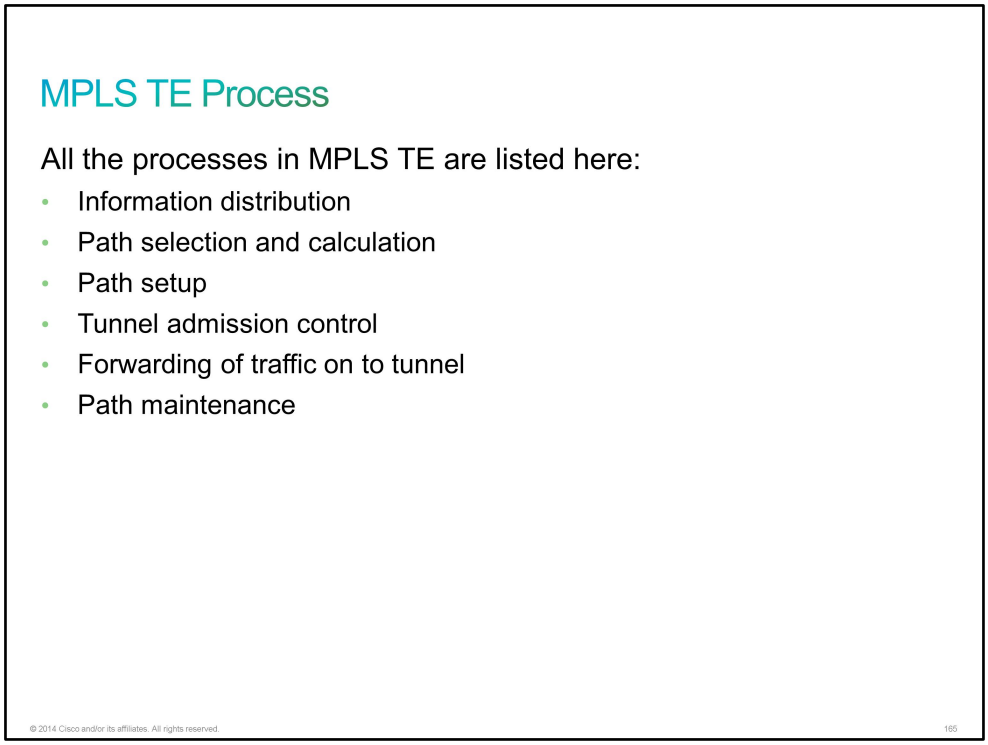

The list of TE processes is as follows:

- **Information distribution:** Because the resource attributes are configured locally for each link, they must be distributed to the headend routers of traffic tunnels. These resource attributes are flooded throughout the network using extensions to link-state intradomain routing protocols, either IS-IS or OSPF. The flooding takes places under these conditions:
  - Link-state changes occur.
  - The resource class of a link changes (this could happen when a network administrator reconfigures the resource class of a link).
  - The amount of available bandwidth crosses one of the preconfigured thresholds.

  The frequency of flooding is bounded by the OSPF and IS-IS timers. There are up thresholds and down thresholds. The up thresholds are used when a new trunk is admitted. The down thresholds are used when an existing trunk goes away.
- **Path selection:** Path selection for a traffic tunnel takes place at the headend routers of the traffic tunnels. Using extended IS-IS or OSPF, the edge routers have knowledge of both network topology and link resources. For each traffic tunnel, the tail-end router starts from the destination of the traffic tunnel and attempts to find the shortest path toward the headend router (using the CSPF algorithm). The CSPF calculation does not consider the links that are explicitly excluded by the resource class affinities of the traffic tunnel or the links that have insufficient bandwidth. The output of the path selection process is an explicit route consisting of a sequence of label switching routers. This path is used as the input to the path setup procedure.

https://t.me/learningnets

- **Path setup:** Path setup is initiated by the headend routers. RSVP is the protocol that establishes the forwarding state along the path that is computed in the path selection process. The headend router sends an RSVP PATH message for each traffic tunnel it originates.
- **Tunnel admission control:** Tunnel admission control manages the situation when a router along a computed path has insufficient bandwidth to honor the resource that is requested in the RSVP PATH message.
- **Forwarding of traffic to a tunnel:** Traffic can be forwarded to a tunnel by several means, including the following:
  - Static routing
  - Policy routing from the global routing table
  - Autoroute
- **Path maintenance:** Path maintenance refers to two operations: path reoptimization and restoration.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Role of RSVP in Path Setup Procedures

Signaling protocol is needed to check the bandwidth available across routers that create paths.

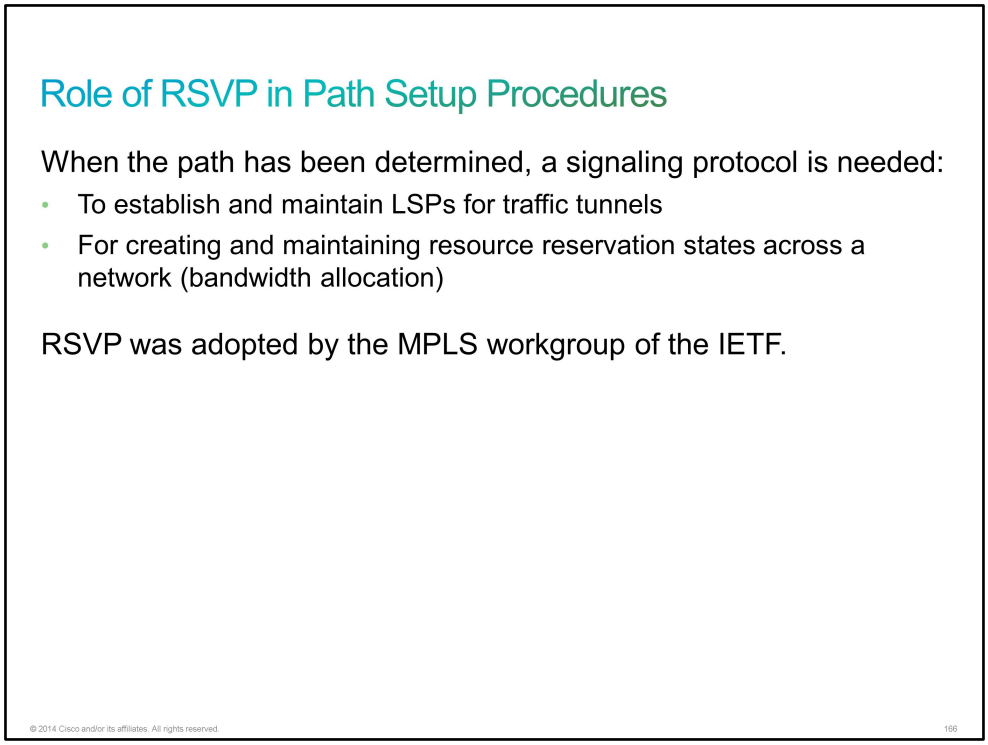

The result of the constraint-based calculation is a list of routers that form the path to the destination. The path is a list of IP addresses that identify each next hop along the path.

However, this list of routers is known only to the router at the headend of the tunnel, the one that is attempting to build the tunnel. Somehow, this now-explicit path must be communicated to the intermediate routers. The intermediate routers do not 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. The MPLS working group of the IETF has adopted RSVP to confirm and reserve the path and apply the labels that identify the tunnel. LDP is used to distribute the labels for the underlying MPLS network.

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

https://t.me/learningnets

Path Setup and Admission Control with RSVP

During path setup procedures, RSVP is signaled through the network.

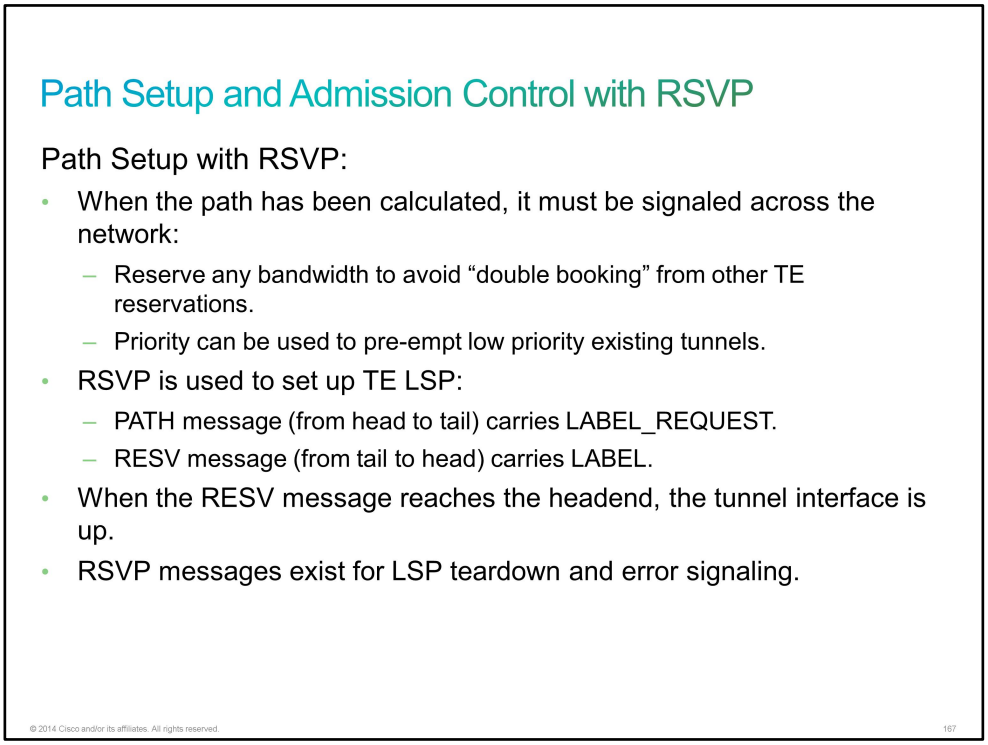

To signal the calculated path across the network, an RSVP PATH message is sent to the tail-end router by the headend router for each traffic tunnel the headend originates.

Note	This process occurs in the MPLS control plane.

The RSVP PATH message carries the explicit route (the output of the path selection process) computed for this traffic tunnel, consisting of a sequence of label switching routers. The RSVP PATH message always follows this explicit route. Each intermediate router along the path performs trunk admission control after receiving the RSVP PATH message. When the router at the end of the path (tail-end router) receives the RSVP PATH message, it sends an RSVP RESV message in the reverse direction toward the headend of the traffic tunnel. As the RSVP RESV message flows toward the headend router, each intermediate node reserves bandwidth and allocates labels for the traffic tunnel. When the RSVP RESV message reaches the headend router, the LSP for the traffic tunnel is established.

RSVP messages also provide support for LSP teardown and error signaling.

https://t.me/learningnets

Path Setup and Admission Control with RSVP (Cont.)

RSVP and Tunnel Admission Control:

- On receipt of PATH message:
  - Router checks whether there is bandwidth available to honor the reservation.
  - If bandwidth is available, then RSVP is accepted.
- On receipt of a RESV message:
  - Router actually reserves the bandwidth for the TE LSP.
  - If pre-emption is required, lower priority LSPs are torn down.
- OSPF or IS-IS updates are triggered.

© 2014 Cisco and/or its affiliates. All rights reserved.

Trunk admission control is used to confirm that each device along the computed path has sufficient provisioned bandwidth to support the resource requested in the RSVP PATH message. When a router receives an RSVP PATH message, it checks whether there is enough bandwidth to honor the reservation at the setup priority of the traffic tunnel. Priority levels 0 to 7 are supported. If there is enough provisioned bandwidth, the reservation is accepted, otherwise the path setup fails. When the router receives the RSVP RESV message, it reserves bandwidth for the LSP. If pre-emption is required, the router must tear down existing tunnels with a lower priority. As part of trunk admission control, the router must do local accounting to keep track of resource utilization and trigger IS-IS or OSPF updates when the available resource crosses the configured thresholds.

https://t.me/learningnets

Forwarding Traffic to a Tunnel

The traffic tunnel normally does not appear in the IP routing table. The IP routing process does not see the traffic tunnel, so the traffic tunnel is generally not included in any SPF calculations.

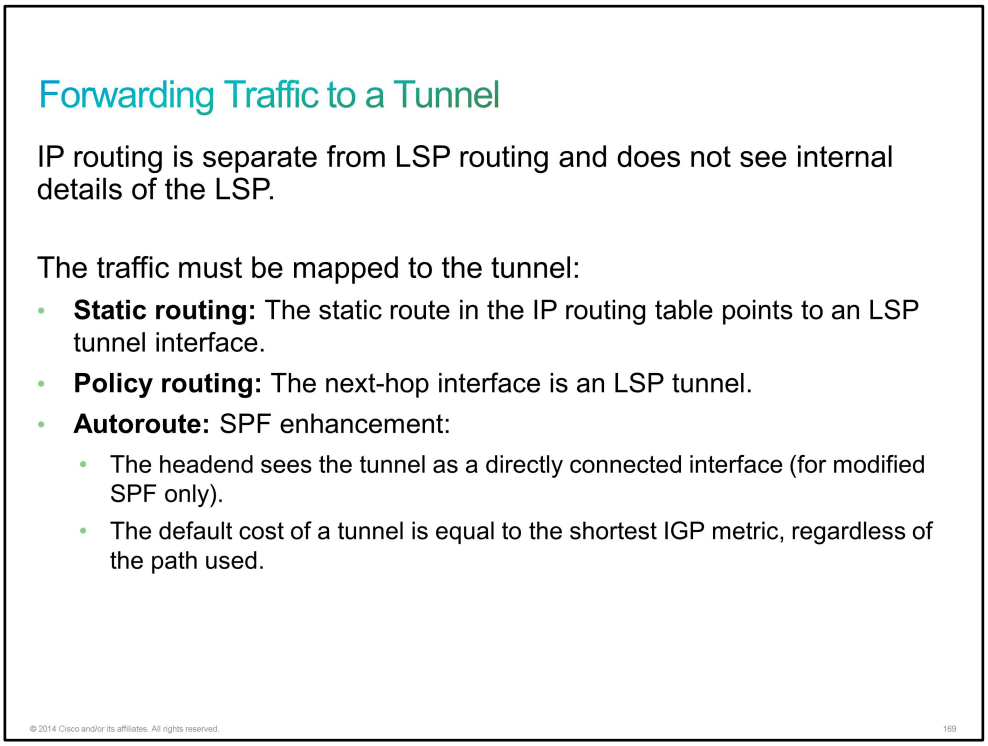

The IP traffic can be mapped onto a traffic tunnel in four ways:

- Use static routes that point to the tunnel interfaces.
- Using PBR, set the next hop for the destination to the tunnel interface.
- Use the autoroute feature, 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 of whether the tunnel is actually using the least-cost path). 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.
- Using forwarding adjacency, which allows the tunnel to be announced via OSPF or IS-IS as a point-to-point link to other routers. To be used for data forwarding, such a traffic tunnel has to be set up bidirectionally.

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 tail-end router (even though the tunnel terminates on that router), will not be able to use the tunnel; instead, it will follow the normal IGP path.

https://t.me/learningnets

Autoroute

To overcome the problems that result from static routing configurations onto MPLS TE tunnels, the autoroute feature was introduced.

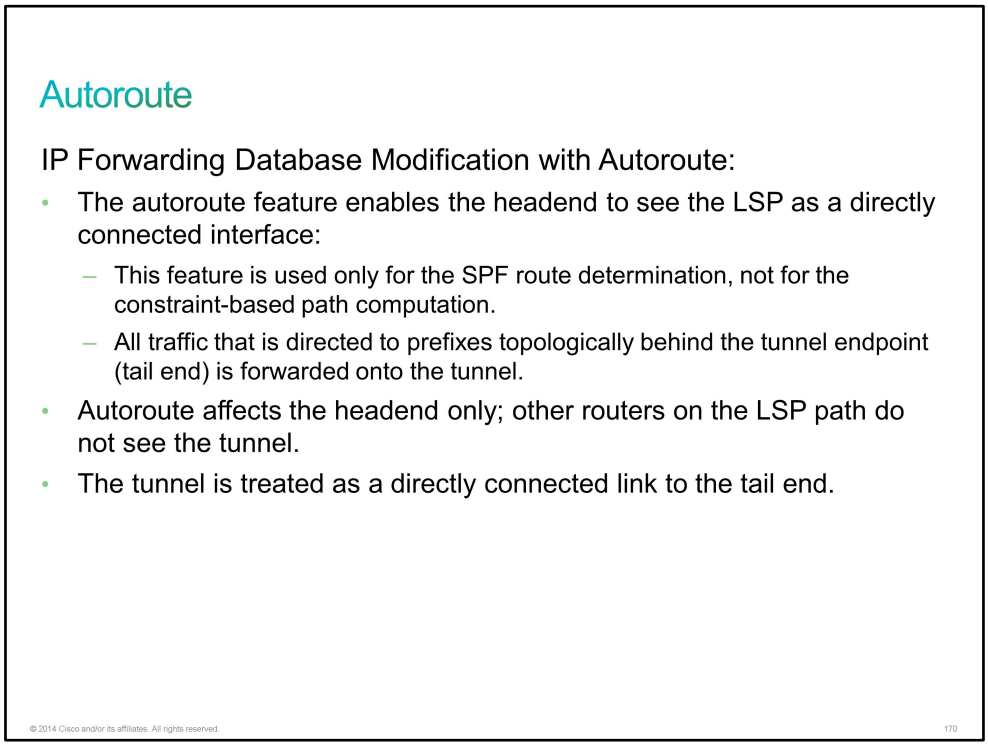

The autoroute feature enables the headend router to see the MPLS TE tunnel as a directly connected interface. The headend uses the MPLS TE tunnel in its modified SPF computations.

Note	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.

When the traffic tunnel is built, there is a directly connected link from headend to tail end.

The autoroute feature enables all the prefixes that are topologically behind the MPLS TE tunnel endpoint (tail end) to be reachable via the tunnel itself. This contrasts 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.

https://t.me/learningnets

Autoroute Example

Lets look at the example of Autoroute feature.

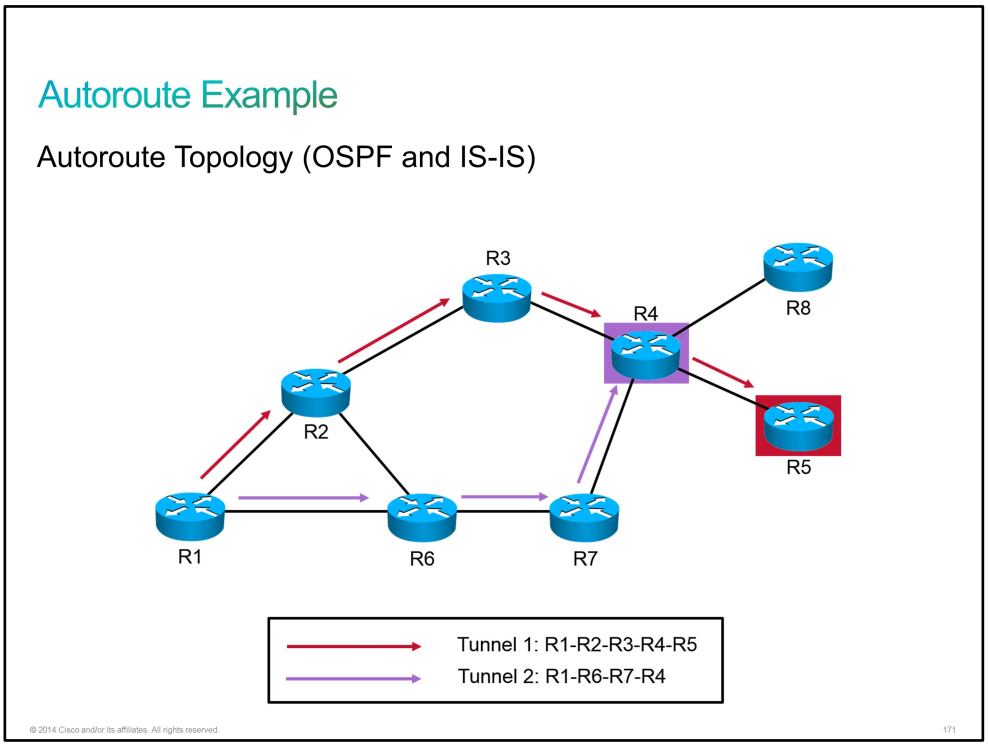

The figure shows an example with two TE tunnels from R1. When the tunnels are up, R4 and R5 appear as directly connected neighbors to R1.

Note	The tunnels are seen for routing purposes only by R1, the headend router. Intermediate routers do not see the tunnel, nor do they take it into consideration for route calculations.

https://t.me/learningnets

Autoroute Example (Cont.)

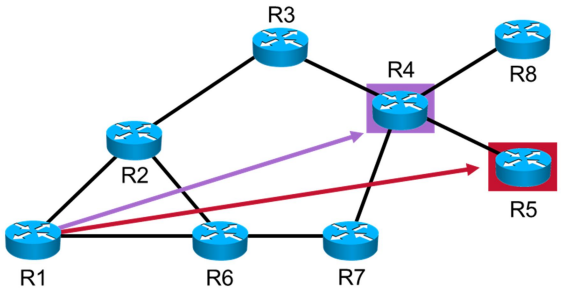

Autoroute Topology (OSPF and IS-IS) from the perspective of R1:

- Next hop to R5 is Tunnel 1.
- Next hop to R4 and R8 is Tunnel 2.

All nodes behind the tunnel are routed via the tunnel.

From the perspective of R1:
Next hop to R5 is Tunnel 1.
Next hop to R4 and R8 is Tunnel2.
All nodes behind the tunnel are routed via the tunnel.

© 2014 Cisco and/or its affiliates. All rights reserved.
172

From the R1 perspective, next hop to router R5 is interface Tunnel 1, and next hop to router R4 and R8 is Tunnel 2. All nodes behind the tunnel are routed via the tunnel.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Traffic engineering is manipulating your traffic to fit your network.
- In TE with a Layer 2 overlay model, PVCs are carefully engineered across the network, normally using an offline management system.
- If TE is not used in a Layer 3 model, some links may be overutilized and others may be underutilized.
- The destination-based forwarding that is currently used in Layer 3 networks cannot resolve the problem of overutilization of one path while an alternate path is underutilized.
- The aim of MPLS TE is to control the path of traffic flow using MPLS labels and LSP.
- A traffic tunnel is a collection of data flows that share some common attribute.
- A traffic tunnel must include attributes that define the commonality between the data flows making up the tunnel.
- A headend router must be provided with network link attributes in order to calculate a path through a network.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Constraint-based path computation finds a path for a traffic flow as the shortest path that meets the resource requirements (constraints) of the traffic flow.
- Tunnel admission control manages the situation when a router along a computed path has insufficient bandwidth to honor the resource that is requested in the RSVP PATH message.
- RSVP is used to confirm and reserve the path and apply the labels that identify the tunnel.
- The RSVP PATH message carries the explicit route (the output of the path selection process) computed for this traffic tunnel, consisting of a sequence of label switching routers.
- The IP routing process does not see the traffic tunnel, so the traffic tunnel is generally not included in any SPF calculations.
- The autoroute feature enables the headend router to see the MPLS TE tunnel as a directly connected interface.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 2: MPLS Traffic Engineering Operations

Overview

This lesson describes the details of link attribute propagation with an IGP and constraint-based path computation. The lesson also describes the details of MPLS TE tunnels, including path setup procedures and path maintenance. The lesson concludes with an explanation of the methods of assigning traffic into MPLS TE tunnels.

Upon completing this lesson, you will be able to:

- Describe the attributes needed for performing constraint-based path computation
- List the four link resource attributes
- Describe the maximum bandwidth and the maximum reservable bandwidth link resource attributes
- Describe the link resource class attribute
- Describe the constraint-based specific link metric attribute (administrative weight)
- List the tunnel attributes
- Describe the traffic parameter and path selection and management tunnel attributes
- Describe the resource class affinity tunnel attributes
- Describe the adaptability, priority, pre-emption tunnel attributes
- Describe the resilience tunnel attributes
- Explain the implementation of TE policies using affinity bits
- Explain the usage of affinity bits in TE policies
- Explain the propagating of MPLS TE link attributes using a link-state routing protocol
- Explain constraint-based path computation

https://t.me/learningnets

- Explain path selection within constraint-based path computation
- Explain example of using constraint-based path computation
- Explain the path setup process
- Explain the RSVP functions in the path setup process
- Explain the hop-by-hop path setup with RSVP
- Explain the tunnel and link admission control process
- Explain the path rerouting process
- Explain path reoptimization
- Explain path rerouting with link failure
- Explain the MPLS TE soft pre-emption
- Explain the MPLS TE path-option attributes
- List the three methods to forward traffic to a tunnel
- Explain using static routing to assign traffic to traffic tunnel
- Explain using autoroute to assign traffic to traffic tunnel
- Explain path selection rules with autoroute
- Explain the MPLS TE IPv6 autoroute
- Explain the need to adjust the tunnel default metric using either an absolute or relative value
- Explain adjusting the tunnel metrics using a relative and absolute value
- Describe the forwarding adjacency feature
- Describe traffic flow without forwarding adjacency and for POP-to-POP traffic

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Attributes Used by Constraint-Based Path Computation

There are two main types of attributes needed for performing constraint-based path computation.

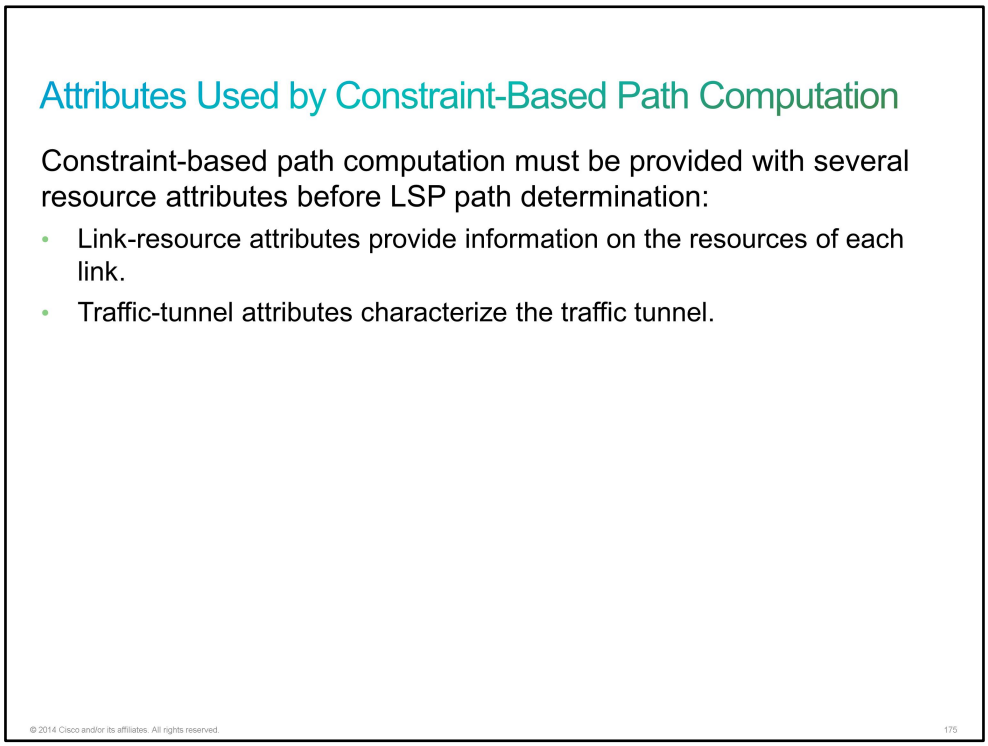

Constraint-based path computation, which takes place at the headend of the traffic-engineered tunnel, must be provided with several resource attributes before the LSP is actually determined. These attributes include the following:

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

https://t.me/learningnets

MPLS TE Link-Resource Attributes

There are four MPLS TE link-resource attributes.

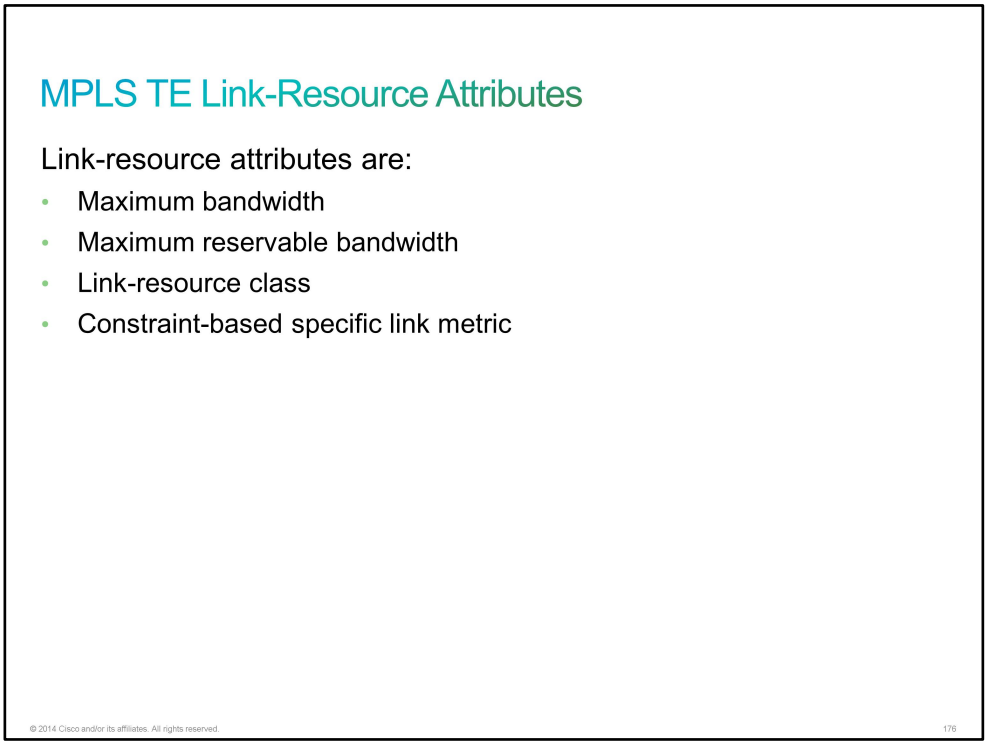

Constraint-based path computation, which takes place at the headend of the traffic-engineered tunnel, must be provided with link-resource attributes. The link-resource attributes provide information on the resources of each link.

Each of these attributes will be discussed in detail.

https://t.me/learningnets

MPLS TE Link Resource Attributes: Maximum Bandwidth and Maximum Reservable Bandwidth

Among the link resource attributes, the most important is the maximum allocation multiplier. This attribute manages the amount of bandwidth that is available on a specified link. “Available” means not yet allocated (as opposed to not 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. The bandwidth at the upper priority level is typically higher than at lower levels (0-7 levels). Because of oversubscriptions, the total amount of bandwidth can exceed the actual bandwidth of the link.

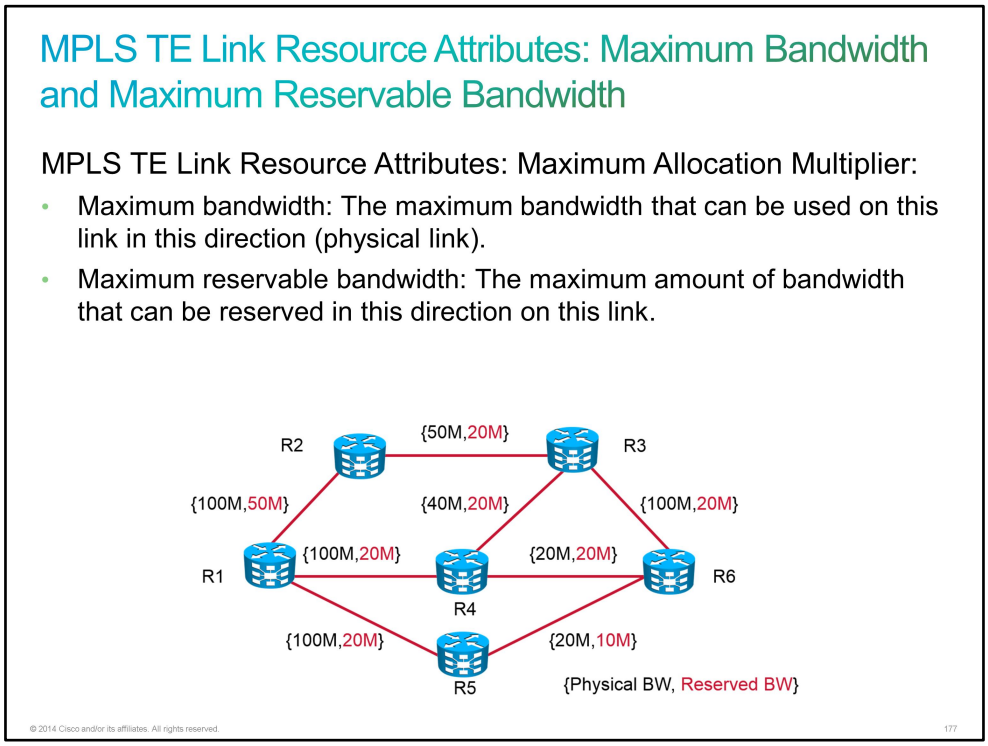

There are three components to the link-resource attribute:

- **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 cannot pre-empt a higher priority.

https://t.me/learningnets

The figure shows topology with few routers and the values of actual physical link bandwidth and reserved bandwidth value for each link.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS TE Link-Resource Attributes: Link-Resource Class

One of the MPLS TE attributes is the link resource class attribute.

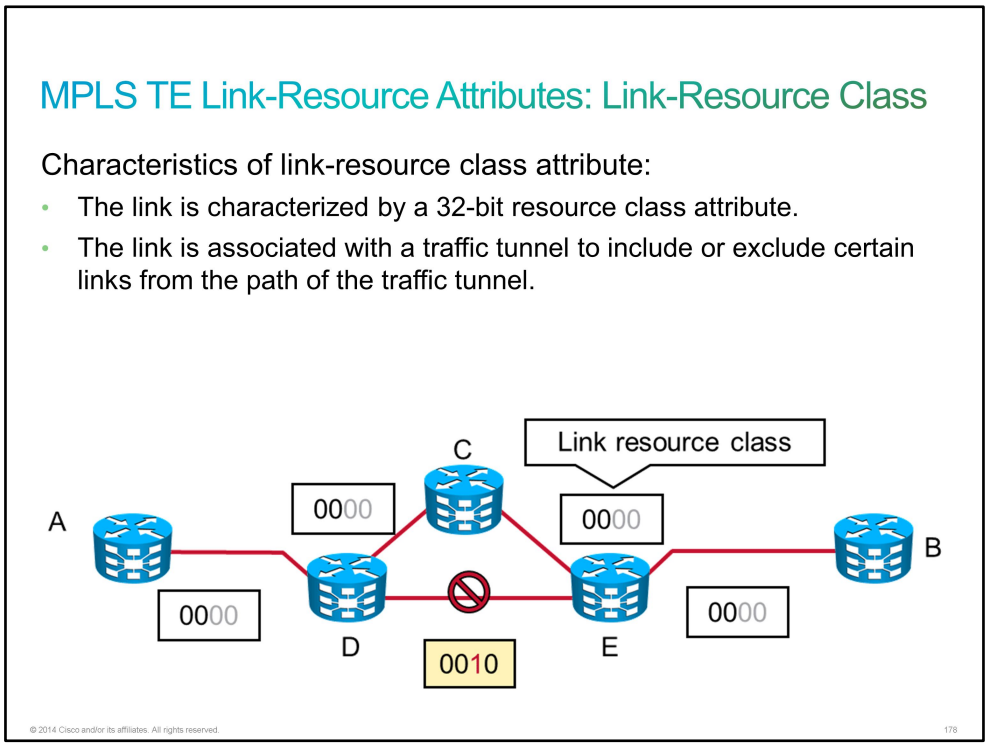

For each link, another link-resource attribute, the link-resource class, is provided. The link is characterized by a 32-bit link resource class attribute string that is matched with the traffic tunnel resource class affinity attribute, and allows inclusion or exclusion of the link into the path of the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

https://t.me/learningnets

MPLS TE Link-Resource Attributes: Constraint-Based Specific Link Metric (Administrative Weight)

Administrative weight is the constraint-based specific attribute that is used to assign cost or metric to links. It is different from the IGP cost or metric.

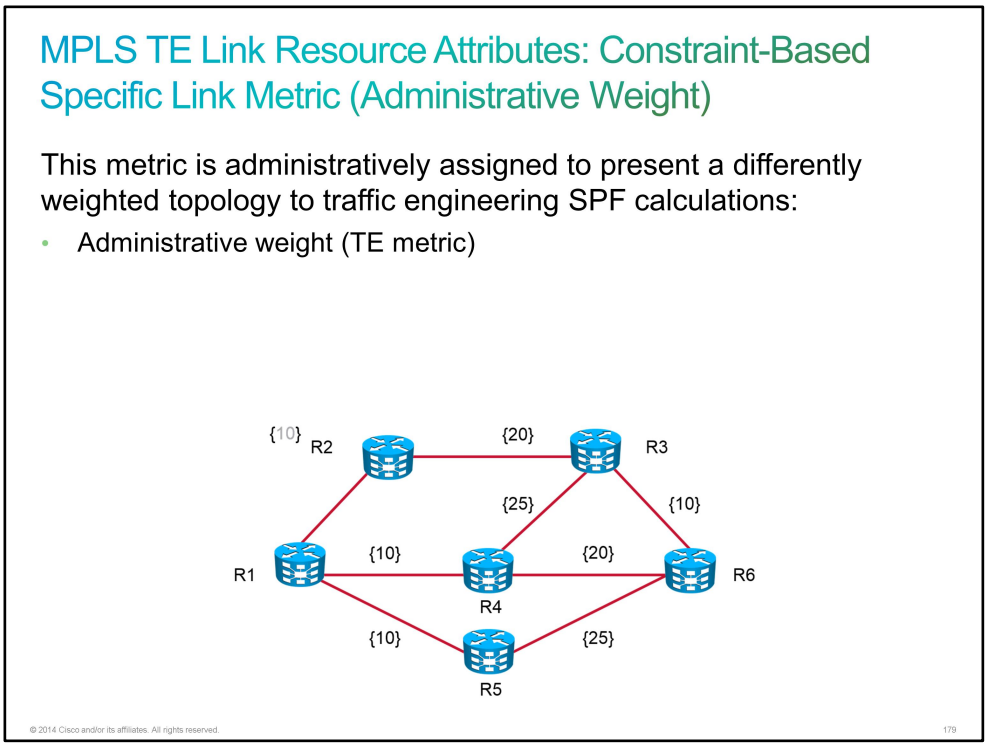

Each link has a cost or metric for calculating routes in the normal operation of the IGP. It may be that, when calculating paths for traffic tunnels, the link should use a different metric than the IGP metric. Hence, a constraint-based specific link metric, the administrative weight, may be administratively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

https://t.me/learningnets

MPLS TE Tunnel Attributes

There are seven tunnel attributes available to influence path selection.

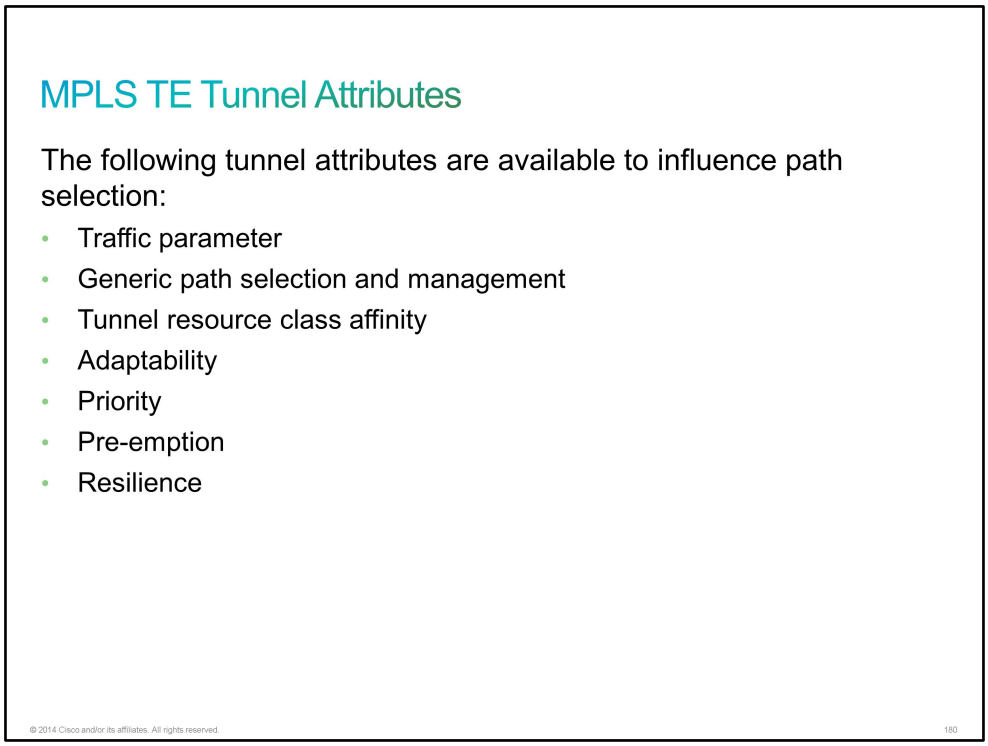

The headend of the tunnel must be provided with MPLS TE tunnel attributes before the LSP is actually determined. The traffic tunnel attributes referred to here characterize the traffic tunnel.

Each of these attributes will be discussed in detail.

https://t.me/learningnets

MPLS TE Tunnel Attributes: Traffic Parameter and Path Selection and Management

There are MPLS TE tunnel attributes that influence path selection of the MPLS TE such as bandwidth and path selection policy.

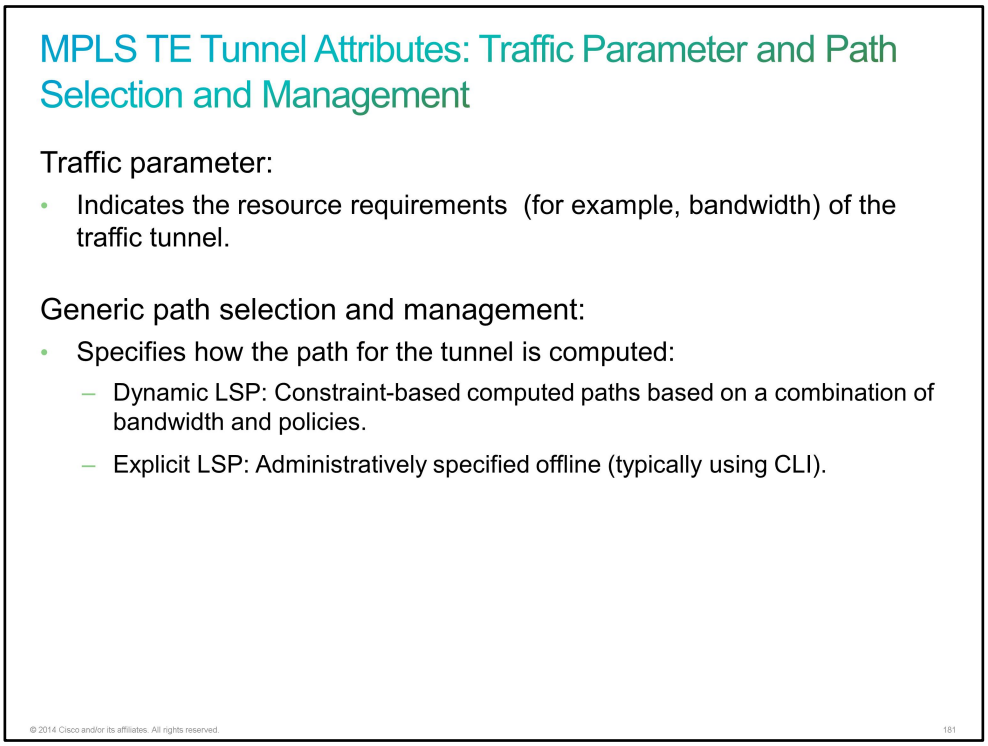

Two of the MPLS TE tunnel attributes affect the path setup and maintenance of the traffic tunnel:

- 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. A path is not considered for an MPLS TE tunnel if it does not have the bandwidth that is required.
- 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 by using the constraint-based path computation; both methods take the resource information and policies into account.

https://t.me/learningnets

MPLS TE Tunnel Attributes: Tunnel Resource Class Affinity

The tunnel resource class affinity attribute allows the network administrator to apply path-selection policies by administratively including or excluding network links.

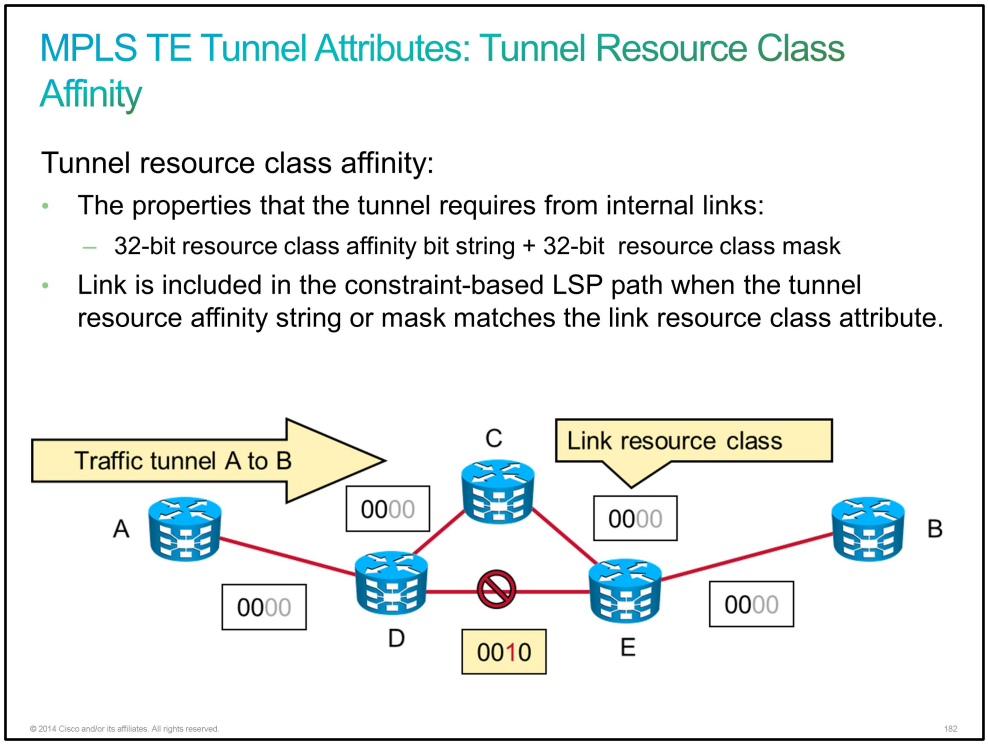

Each link may be assigned a resource class attribute. Resource class affinity specifies whether to explicitly include or exclude links with resource classes in the path selection process. The resource class affinity is a 32-bit string that is accompanied by a 32-bit resource class mask. The mask indicates 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.

https://t.me/learningnets

MPLS TE Tunnel Attributes: Adaptability, Priority, Pre-emption

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.

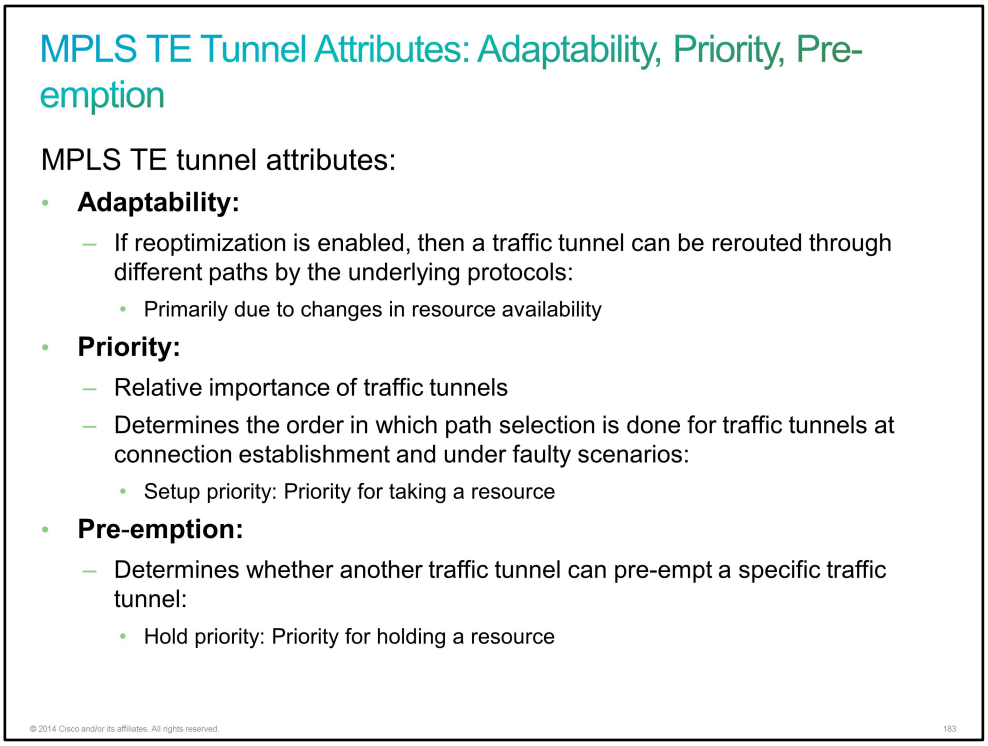

Two types of priorities are assigned to each traffic tunnel:

- 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 because 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 DiffServ environment. Pre-emption can also be used to implement various prioritized restoration policies following fault events.

https://t.me/learningnets

MPLS TE Tunnel Attributes: Resilience

You can define the behavior of the tunnel in faulty conditions or if the tunnel becomes noncompliant with tunnel attributes (for example the required bandwidth).

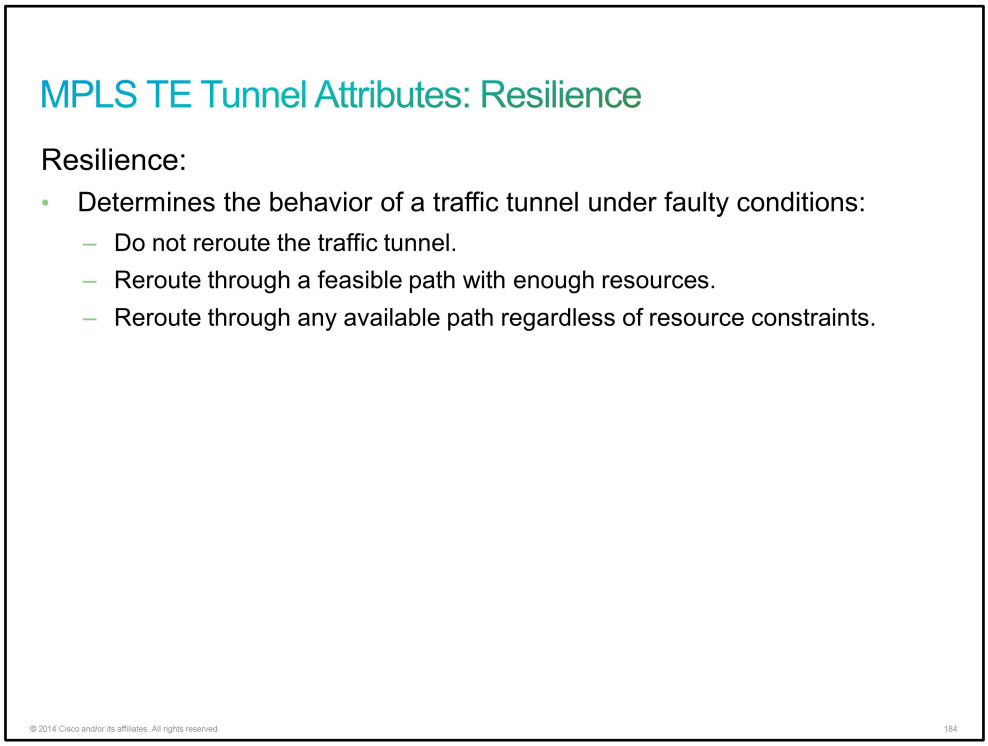

The resilience attribute determines the behavior of the tunnel under faulty conditions; it can specify the following behavior:

- 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 though any available path, irrespective of available link resources

https://t.me/learningnets

Implementing TE Policies with Affinity Bits

The policies during LSP 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).

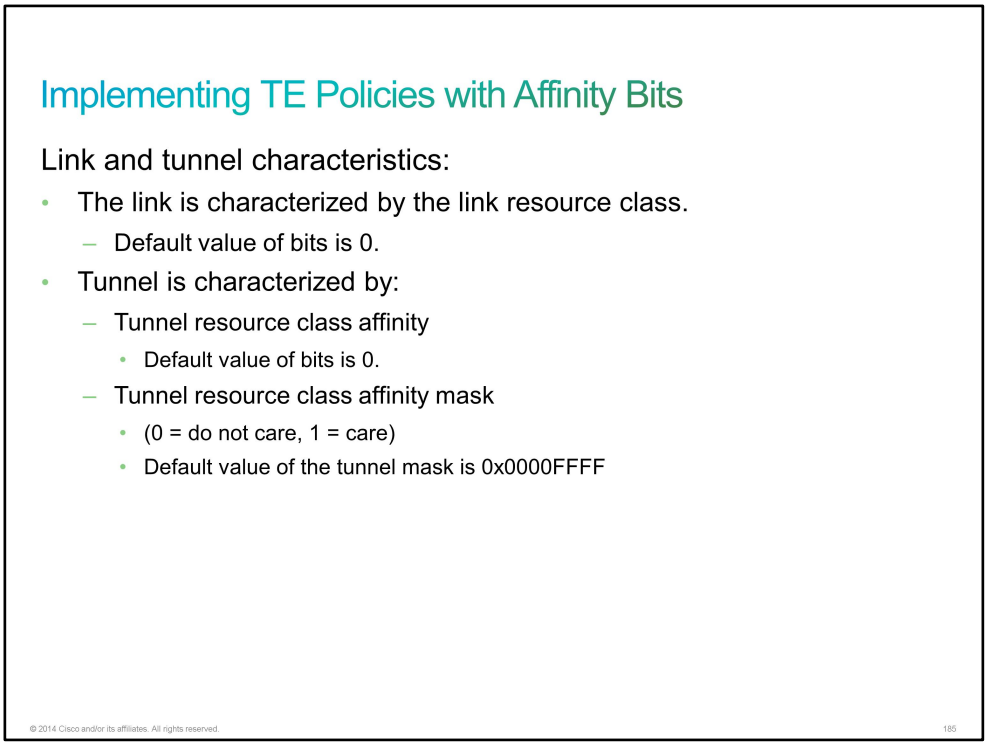

Each traffic tunnel is characterized by a 32-bit resource class affinity string, which is accompanied by a respective resource class mask. The zero bits in the mask exclude the respective link resource class bits from being checked.

Each link is characterized by its resource class 32-bit string, that 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 computation.

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

https://t.me/learningnets

Using Affinity Bits in TE Policies

This example 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 (tail end).

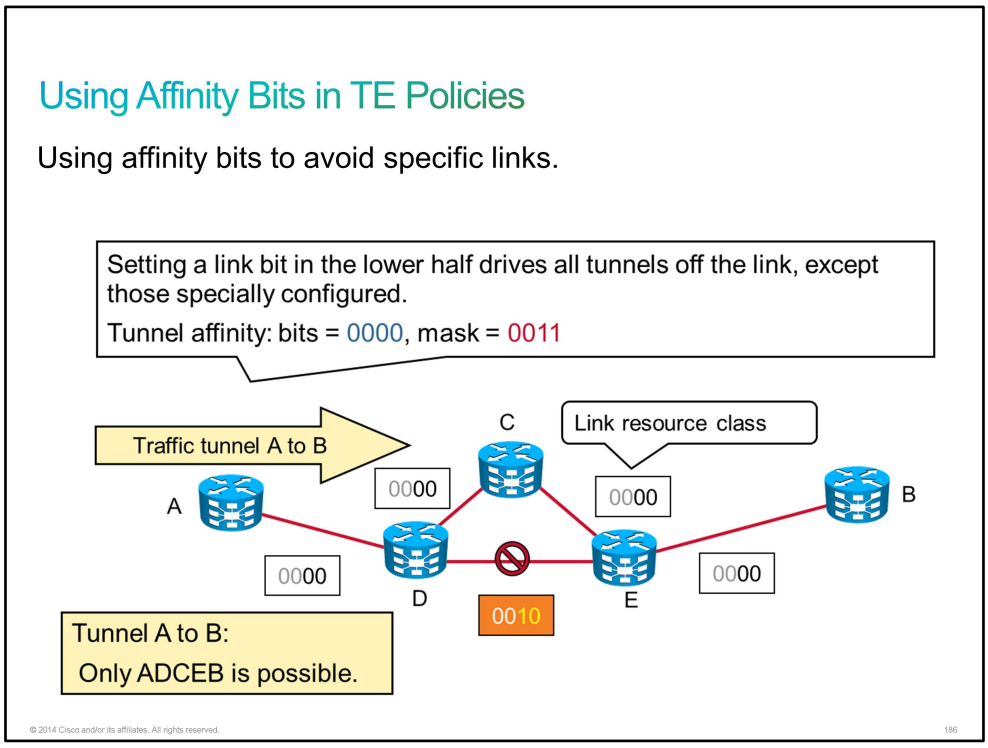

With the tunnel resource class affinity bits and the link resource class bits at their default values of 0, the constraint-based path computation would have two possible paths: A-D-E-B or A-D-C-E-B.

Because moving all dynamically computed paths away from the link D-E is desirable, 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.

https://t.me/learningnets

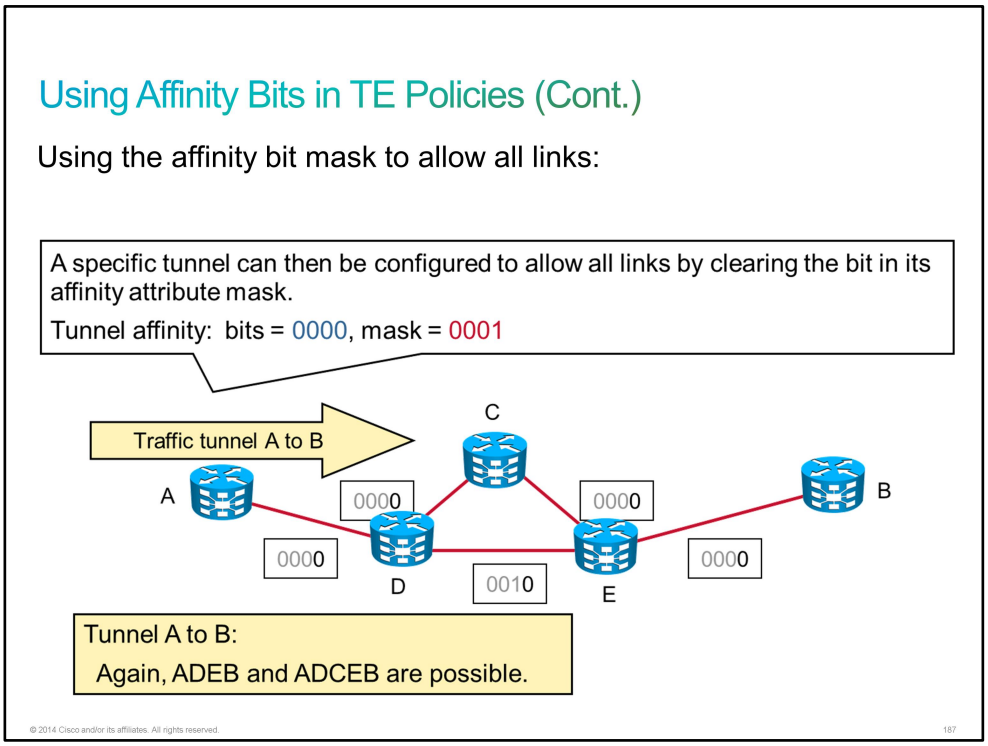

In this sample network, only the lower bit has been set in the tunnel mask. The tunnel affinity bits remain unchanged, as do the resource class bits on the D-E link.

The matching between the tunnel resource class affinity bits and the link resource class bits is done on the lowest bit only (because the mask setting is 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).

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

https://t.me/learningnets

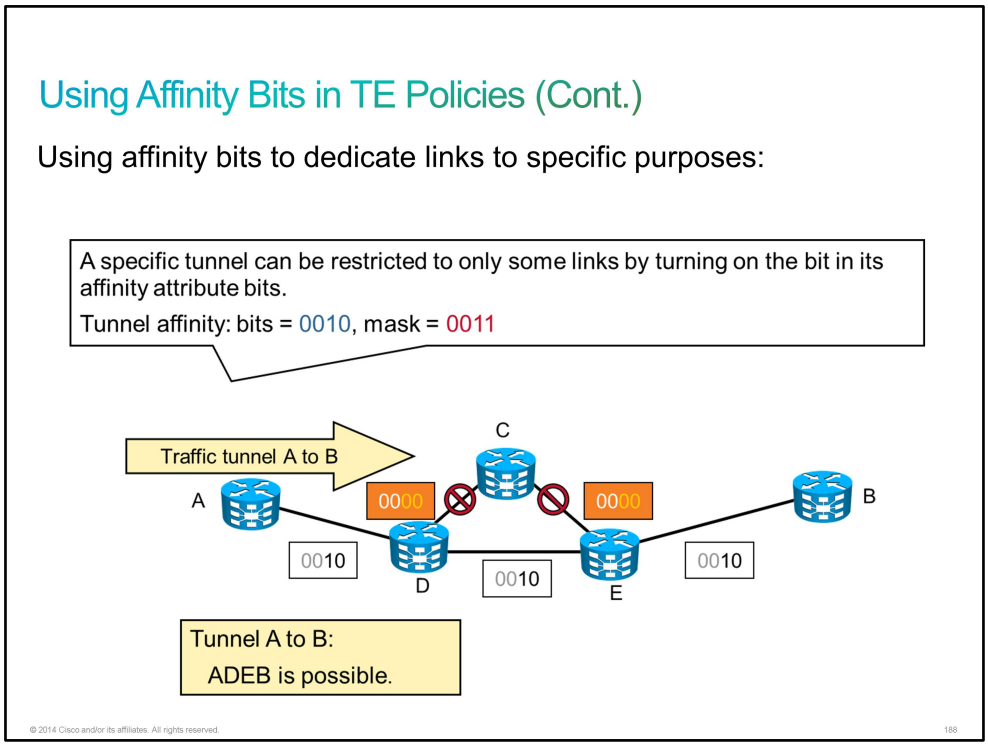

This example shows the setting of the tunnel resource class affinity bits and the link resource class bits to force the tunnel to follow a specific path. Links A-D-E-B are all configured with the resource class value 0010.

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 remains (A-D-E-B).

https://t.me/learningnets

Propagating MPLS TE Link Attributes with Link-State Routing Protocol

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

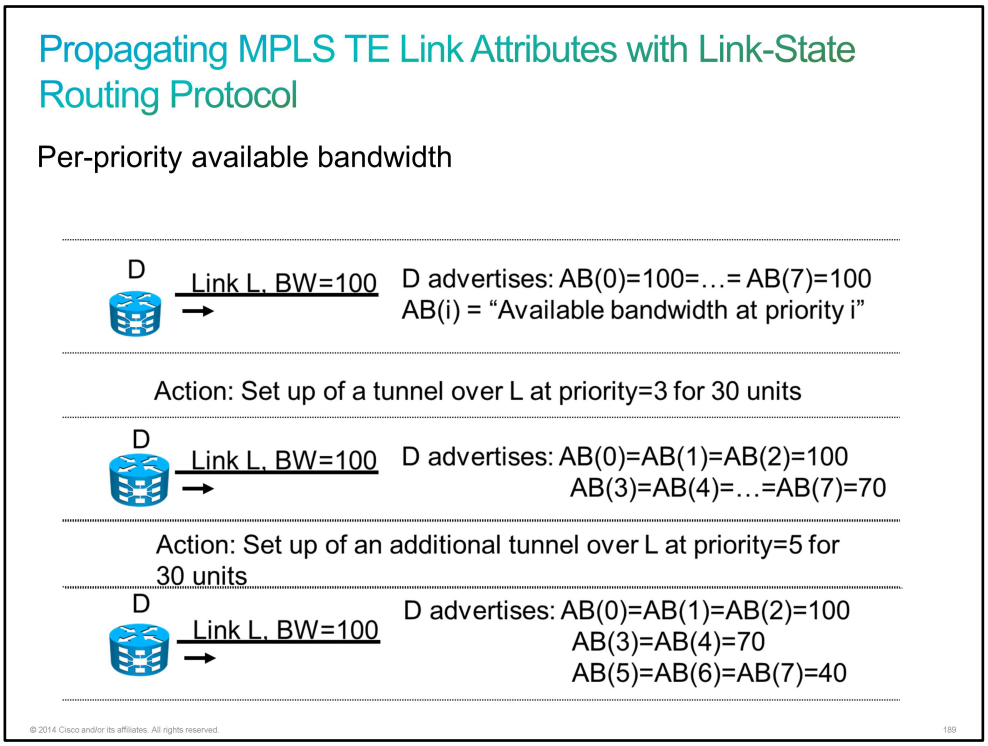

Because the propagation (flooding) of the attributes can be achieved only by IGPs, OSPF and IS-IS were extended to support the MPLS TE features.

OSPF uses new link-state advertisements (opaque LSAs), and IS-IS uses new 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 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).

https://t.me/learningnets

Finally, 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.

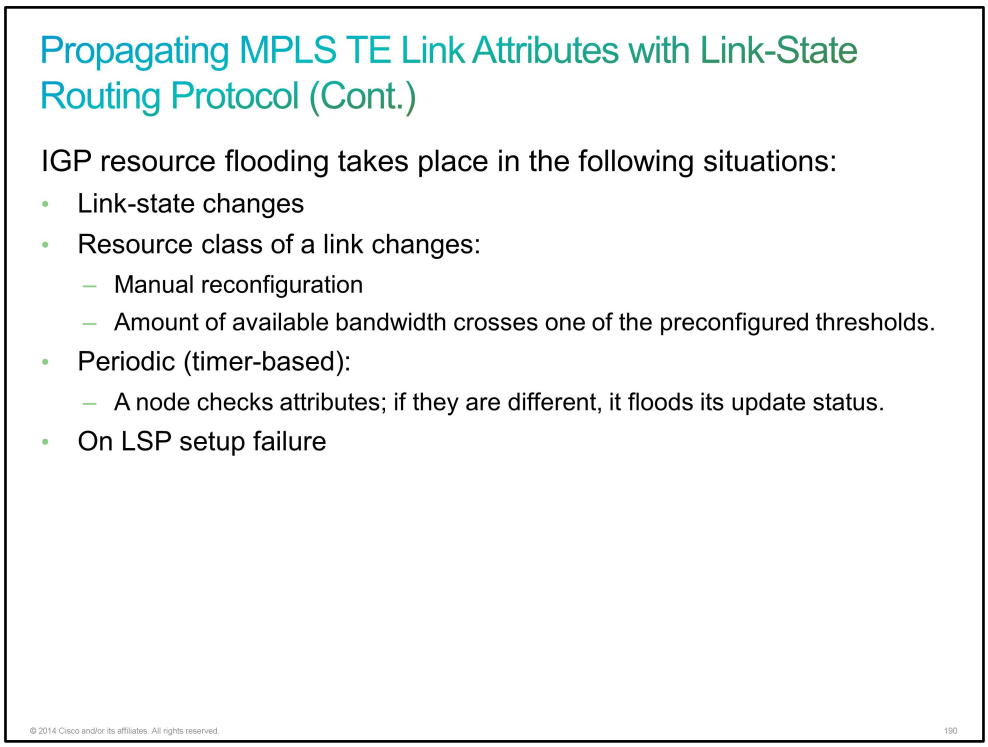

The flooding of resource attributes by the IGP takes place along with certain conditions and events:

- 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.
- When a node periodically checks resource attributes, and if the resource attributes were changed, the update is flooded.
- When the LSP setup fails.

https://t.me/learningnets

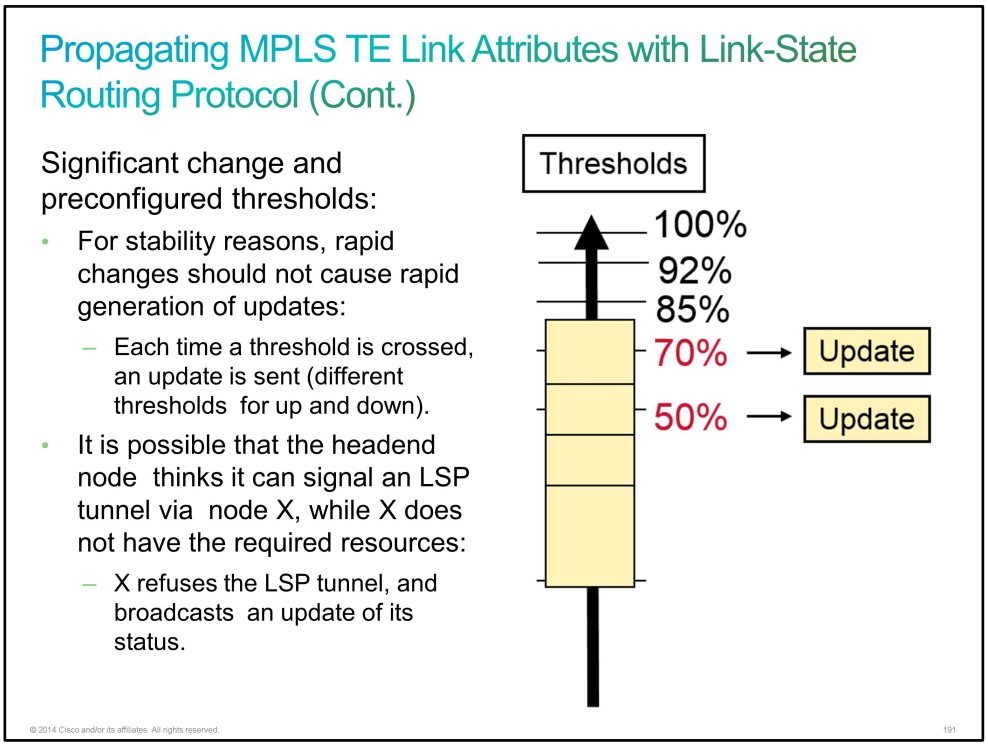

For stability purposes, significant rapid changes in available link resources should not trigger the updates immediately.

There is a drawback, however, in not propagating the changes immediately. Sometimes 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 crossed.

https://t.me/learningnets

Constraint-Based Path Computation

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

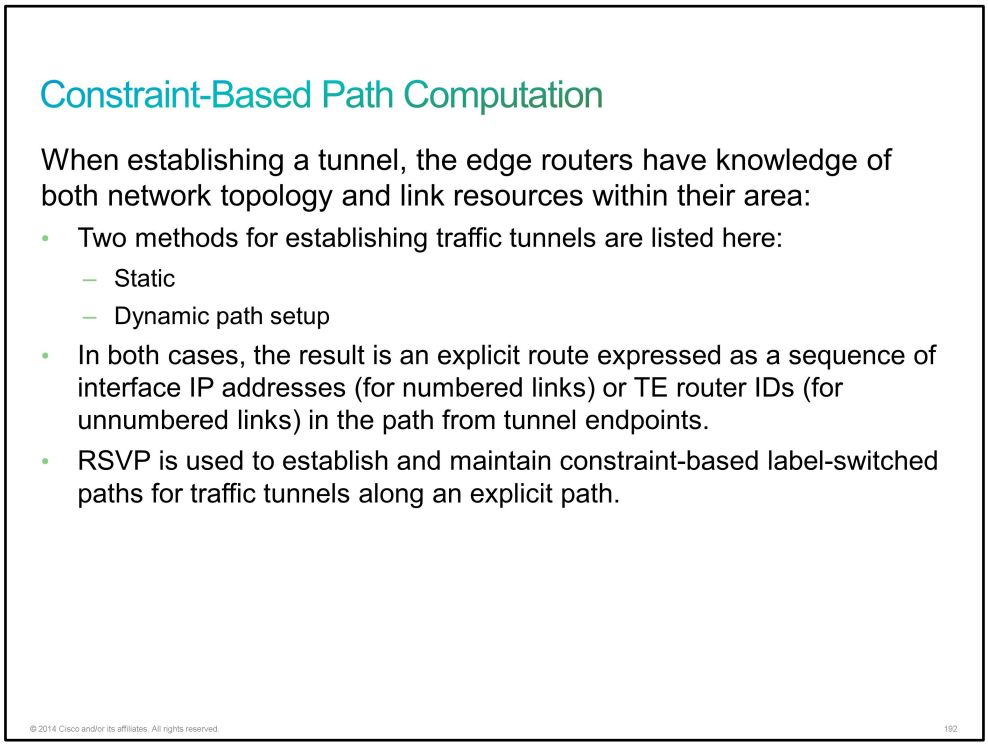

The LSP for the traffic tunnel can be statically defined or computed dynamically. The computation takes the available resources and other tunnel and link attributes into account (thus, it represents constraint-based path computation). The result of the constraint-based path computation is a series of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

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

https://t.me/learningnets

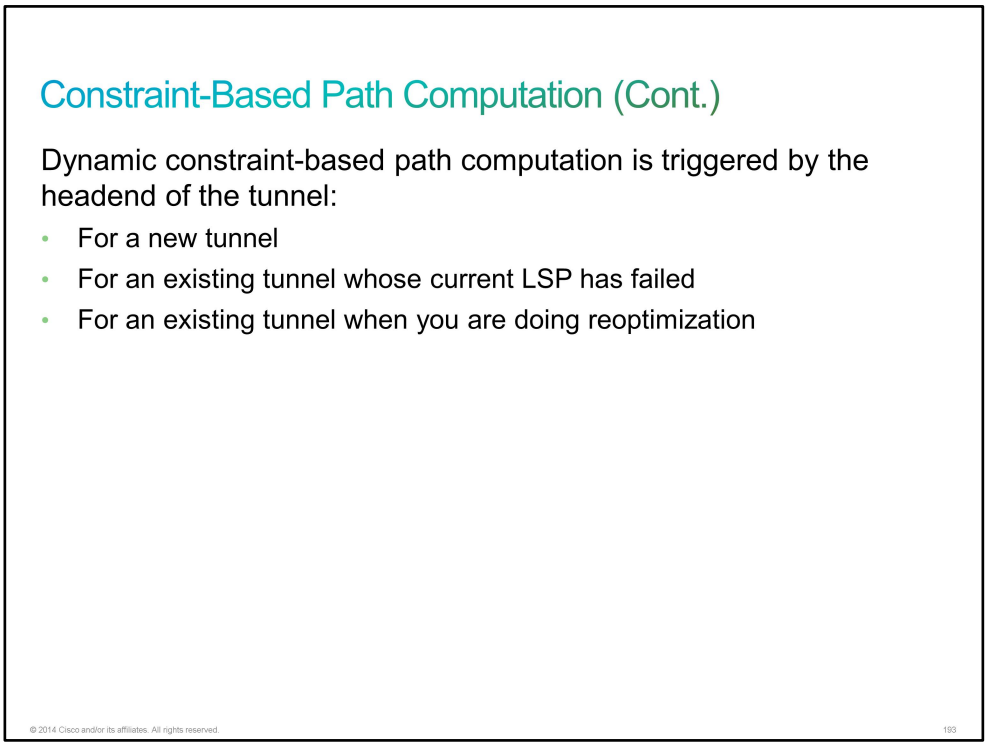

The LSP computation is restricted by several factors (constraint-based). The LSP can be computed only if these conditions are met:

- 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.

https://t.me/learningnets

Constraint-Based Path Computation: Path Selection

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.

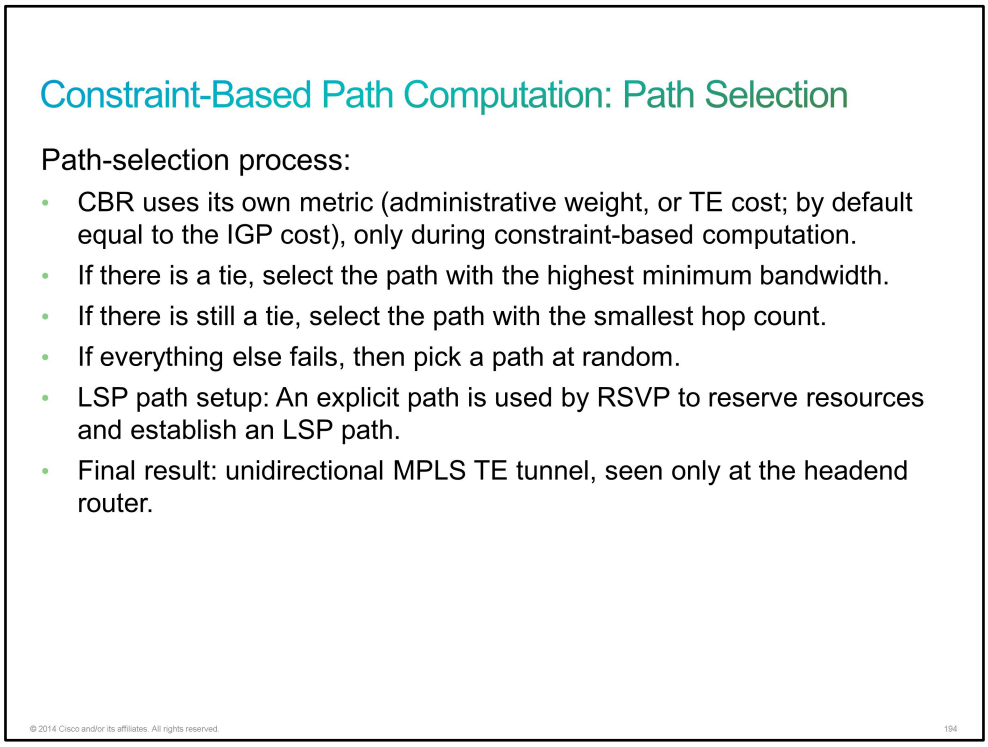

If there are more candidates for the LSP (several paths with the same metric), then the selection criteria is as follows (in sequential order):

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

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

When the LSP 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 tail end).

https://t.me/learningnets

Constraint-Based Path Computation: Path Selection (Cont.)

- 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).
  - The tunnel interface is used for MPLS TE tunnel creation and visualization, but the behavior of MPLS TE tunnels is different from other tunnel protocols (for example, GRE).
- Only traffic entering at the headend router will use the tunnel.
- IP cost: If autoroute is used, an 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).

© 2014 Cisco and/or its affiliates. All rights reserved.

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 calculations (except for other IP tunnels such as 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 defining two tunnels in a bidirectional way.

Cisco IOS Software and Cisco IOS XR Software use the tunnel mainly for visualization. The rest of the actions that are 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 other networks do not have any knowledge of it.

https://t.me/learningnets

Example of Constraint-Based Path Computation

This example of constraint-based path computation and LSP selection requires that the traffic tunnel be established between R1 (headend) and R6 (tail end).

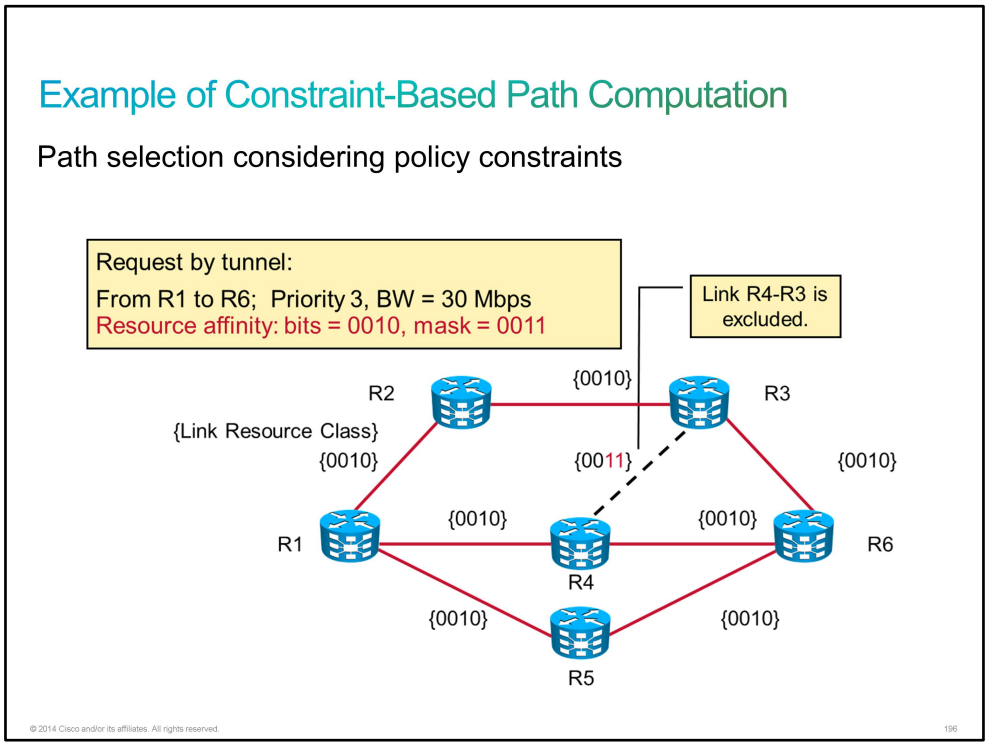

In this example, the traffic tunnel requirements are as follows:

- 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; 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 computation.

https://t.me/learningnets

Example of Constraint-Based Path Computation (Cont.)

Path selection considering available resources

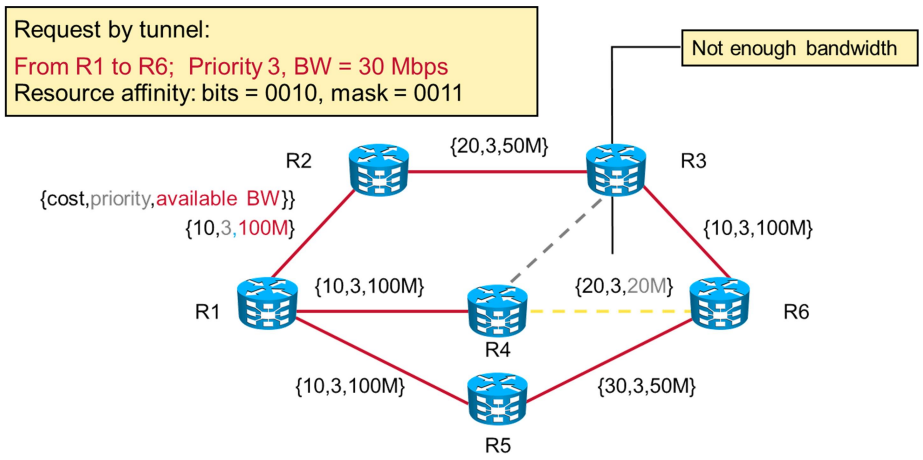

© 2014 Cisco and/or its affiliates. All rights reserved.

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 computation.

Example of Constraint-Based Path Computation (Cont.)

Selecting the best path

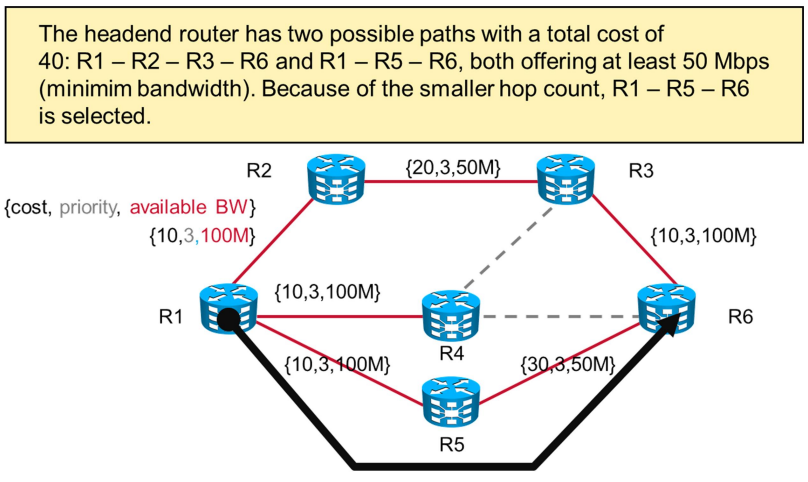

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

The resulting LSPs (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 tie-break 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, 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 for the traffic-engineered tunnel is the signaling of the path via RSVP.

https://t.me/learningnets

Path Setup

LSP setup is always initiated at the traffic tunnel headend.

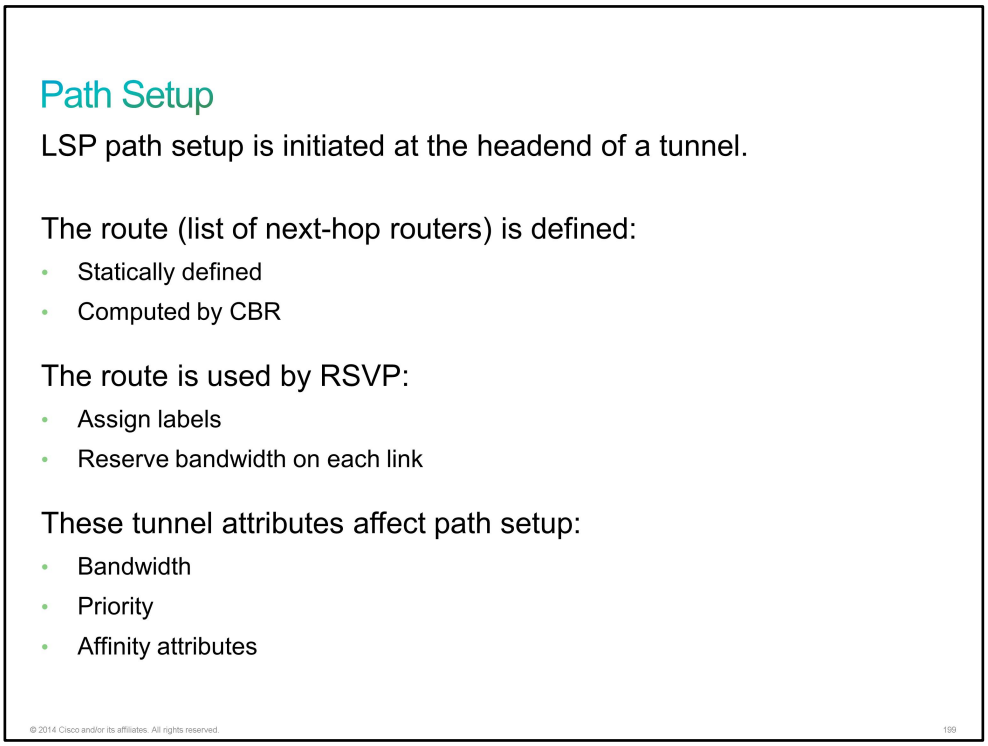

The explicit route for the traffic tunnel is composed of a list of next-hop routers toward the tunnel endpoint (or tail end). The LSP tunnels can be statically defined or computed with CBR and thus routed away from network failures, congestion, and bottlenecks.

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

https://t.me/learningnets

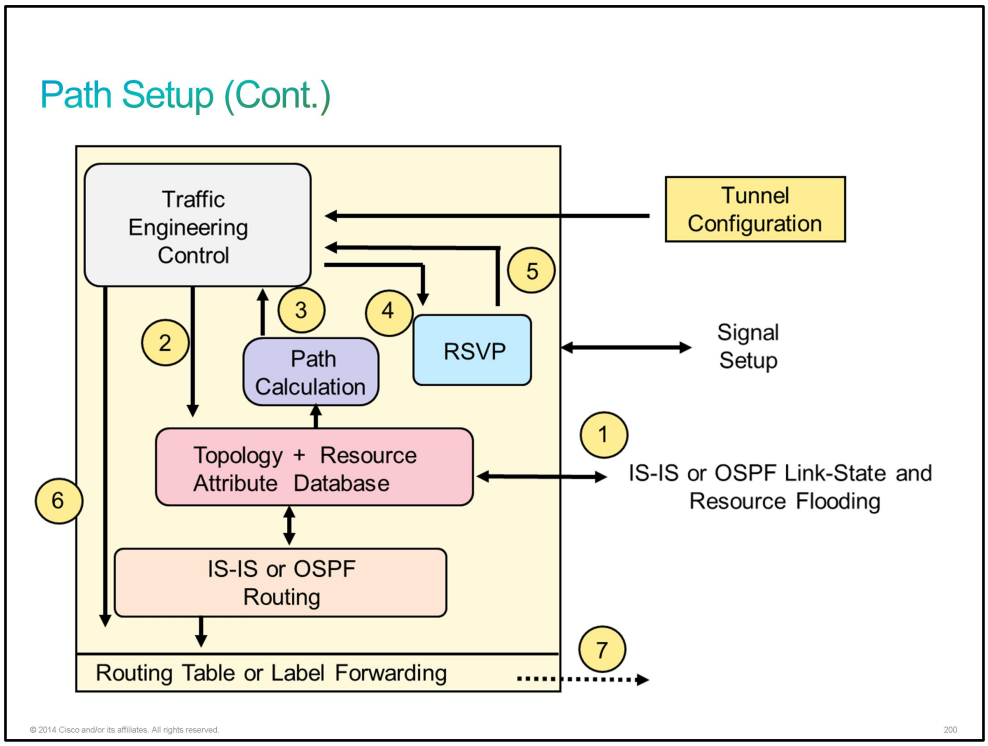

The figure shows a conceptual block diagram of the components of CBR and path computation. In the upper left corner is a TE control module, where the control algorithms run. The module looks at the tunnels that have been configured for CBR.

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 transfers the path 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.

https://t.me/learningnets

RSVP Usage in Path Setup

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

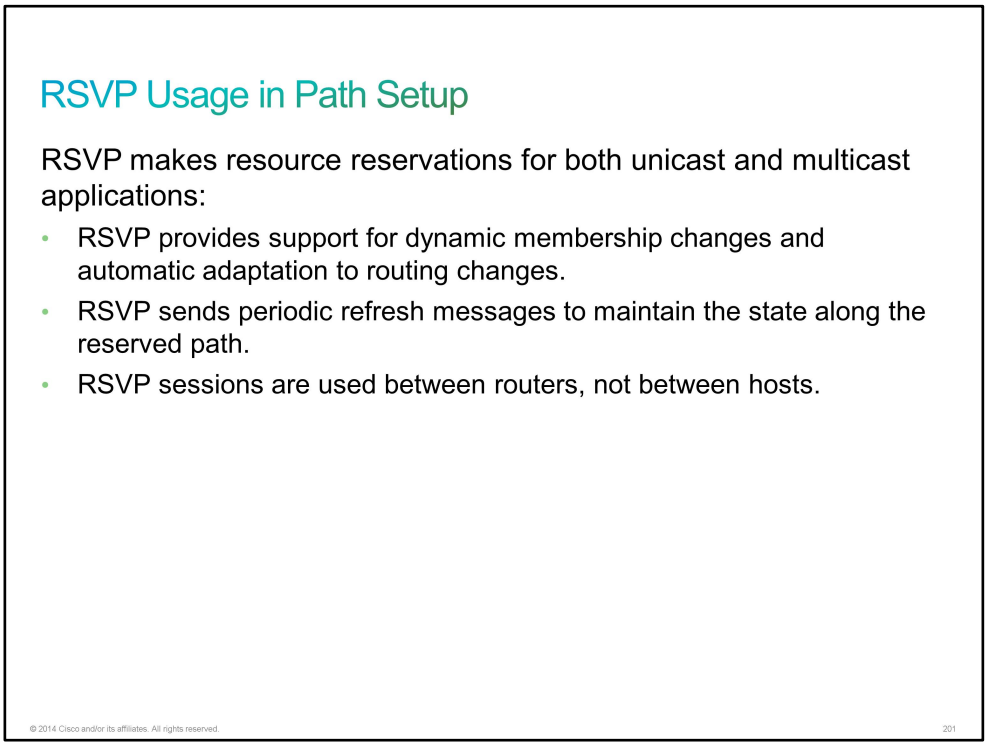

RSVP dynamically adapts to changes either in membership (for example, multicast groups) or in the routing tables.

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
- ResvConf
- ResvTear

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.

https://t.me/learningnets

Hop-by-Hop Path Setup with RSVP

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

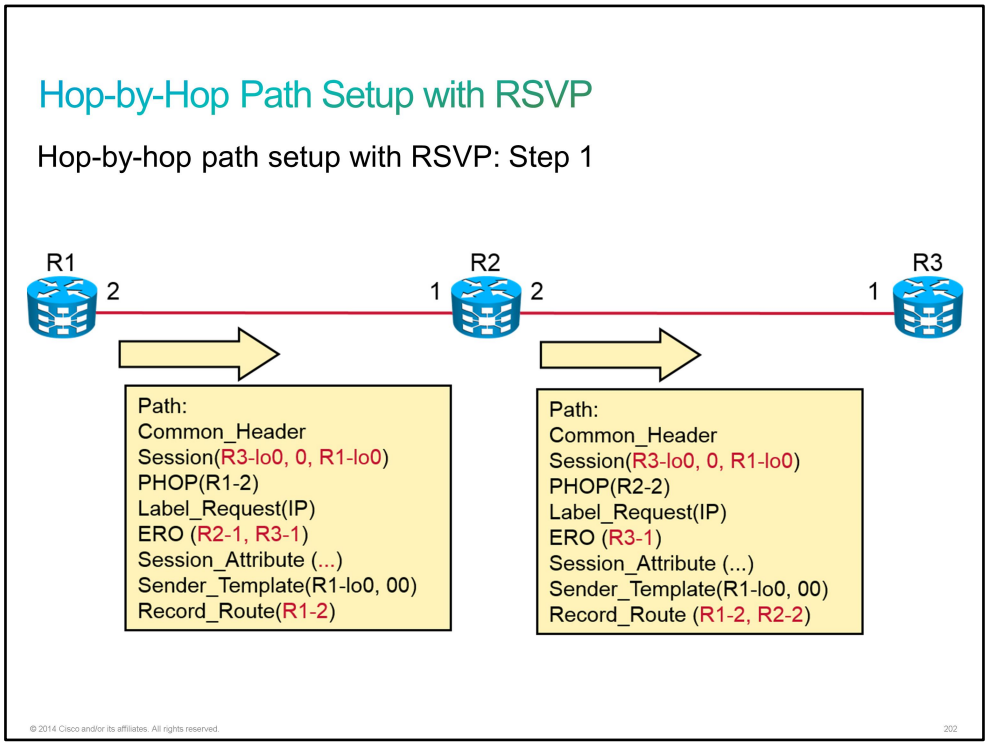

The RSVP Path message contains several objects, including the session identification (R3-lo0, 0, R1-lo0 in the example), which uniquely identifies the tunnel. The traffic requirements for the tunnel are carried in the session attribute. The label request that is present in the message is handled by the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

The ERO is populated by the list of next hops, that 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 PHOP is set to the outgoing interface address of the router. The RRO is populated with the same address as well.

Note	The sender template is used in assigning unique LSP identifiers (R1-lo0 = loopback interface 0, which identifies the tunnel headend; 00 = the LSP ID). The same tunnel can take two possible LSPs (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 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 (path calculation, or 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.

https://t.me/learningnets

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.

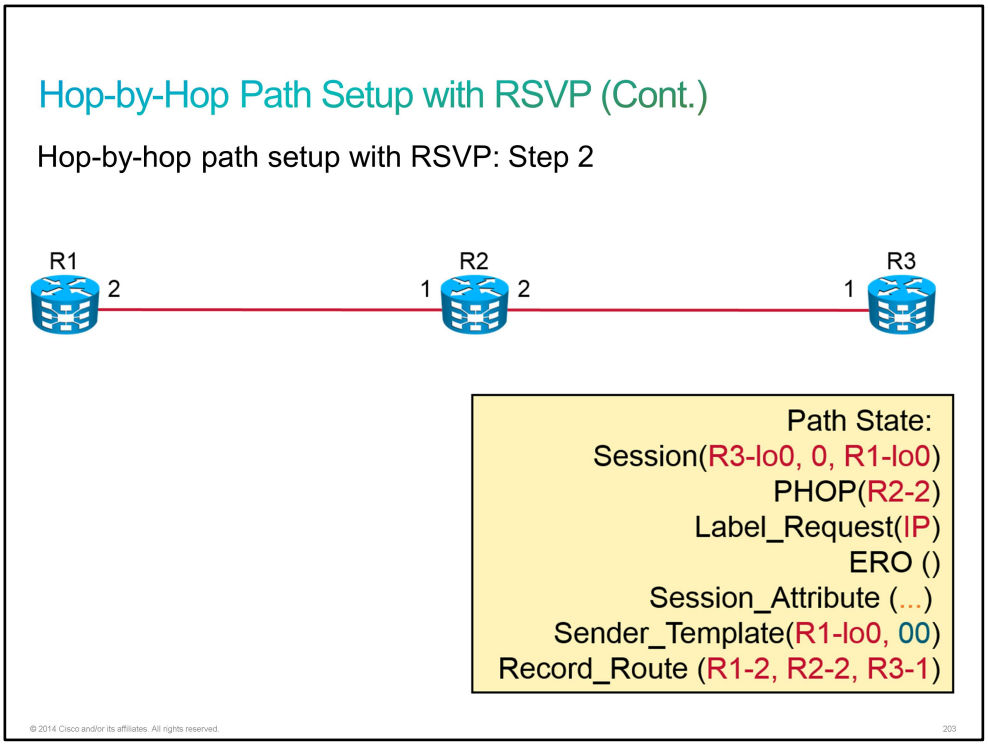

When the RSVP Path message arrives at the tail-end router (the endpoint of the tunnel), the label request message triggers the path label allocation. The label is placed in the corresponding label object of the RSVP Resv message that is generated. The RSVP message is sent back to the headend following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

When the RSVP Path message arrives at the tail-end router (R3), the path state block is created and the ERO becomes empty (after removing the address of the router itself from the list), indicating that it has reached the tail end 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 tail-end router to allocate a label for the specified LSP tunnel (session).

https://t.me/learningnets

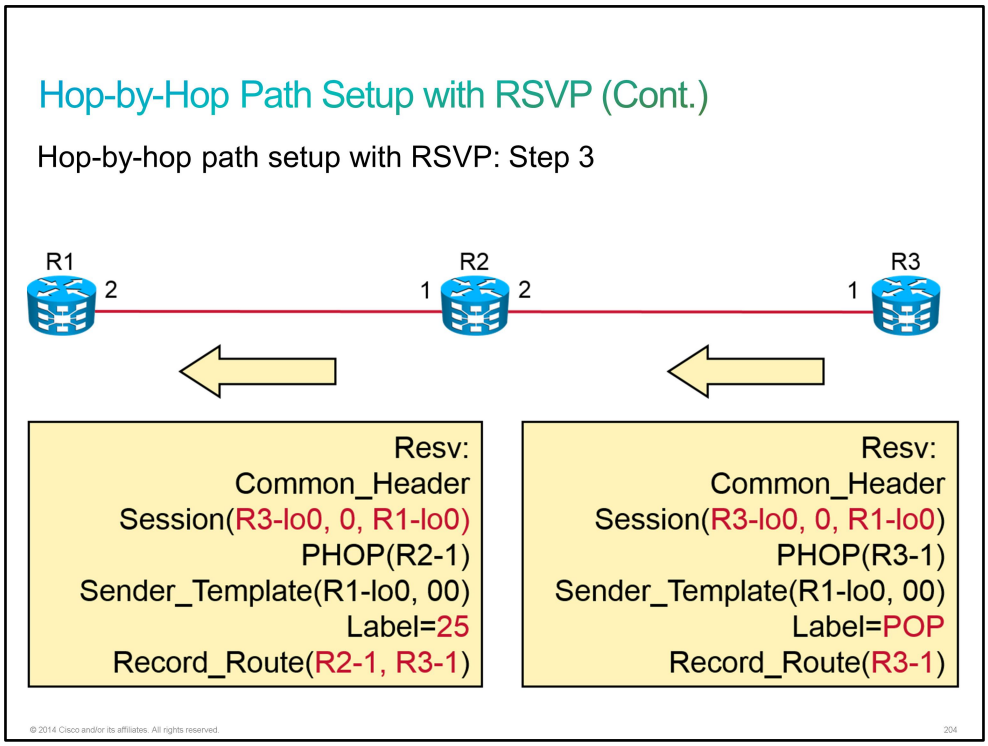

Because R3 is the tail-end router, it does not allocate a specific label for the LSP tunnel. The implicit null label is used instead (the value "POP" in the label object).

The PHOP in the RSVP Resv message is populated by the interface address of the tail-end 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 tail-end router. The next-hop information in the path state block was established when the Path message was traveling in the opposite direction (headend to tail end).

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.

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 tail end).

https://t.me/learningnets

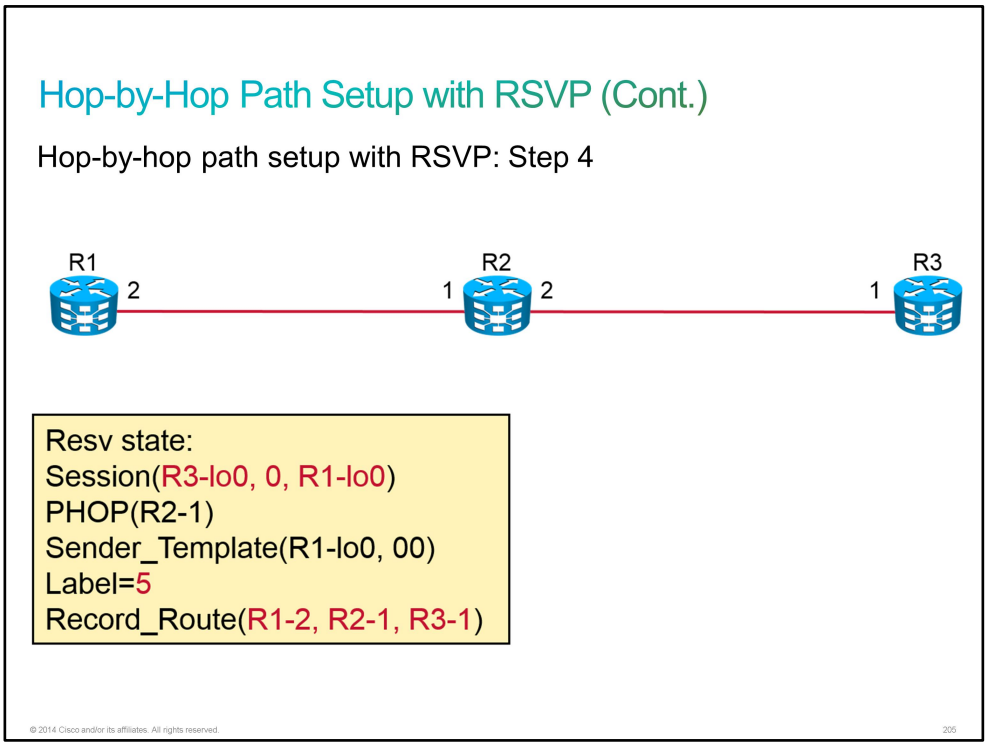

When the RSVP Resv message arrives at the headend router (R1), the LSP setup is concluded. The label (5) that is allocated by the next-hop router toward the tunnel endpoint (PHOP = R2-1) is stored, and the explicit route that was taken by the tunnel is now present in the RRO. The LSP tunnel is established.

https://t.me/learningnets

Tunnel and Link Admission Control

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

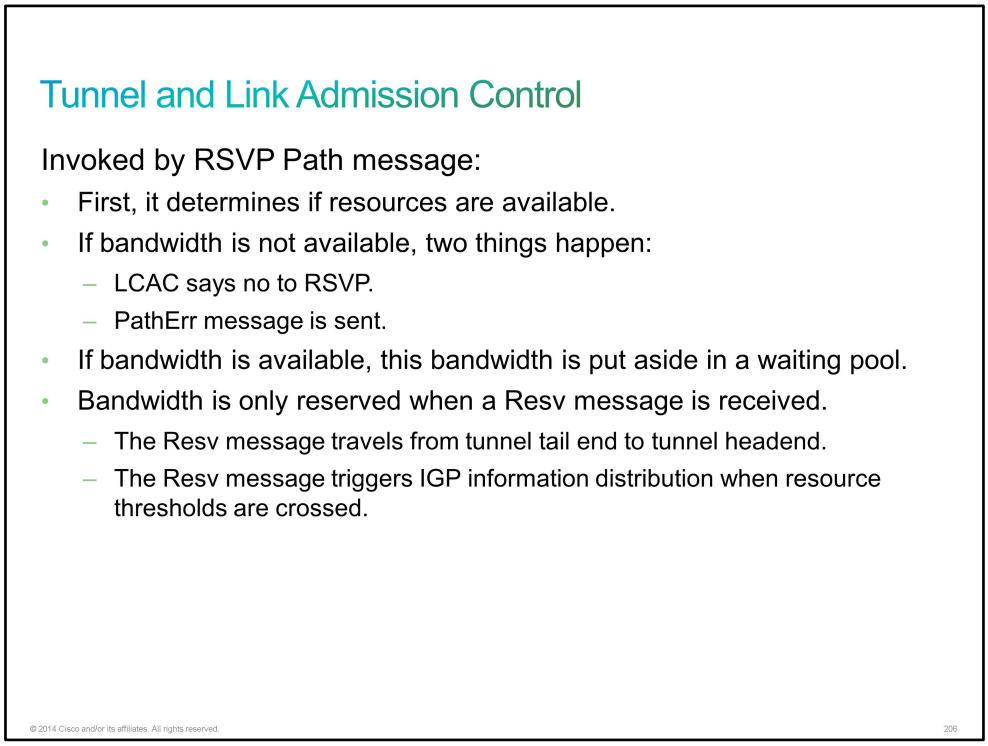

Each hop on the way determines if the available resources that are specified in the session attribute object are available.

If there is not enough bandwidth on a specified link through which the traffic tunnel should be established, the 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.

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.

https://t.me/learningnets

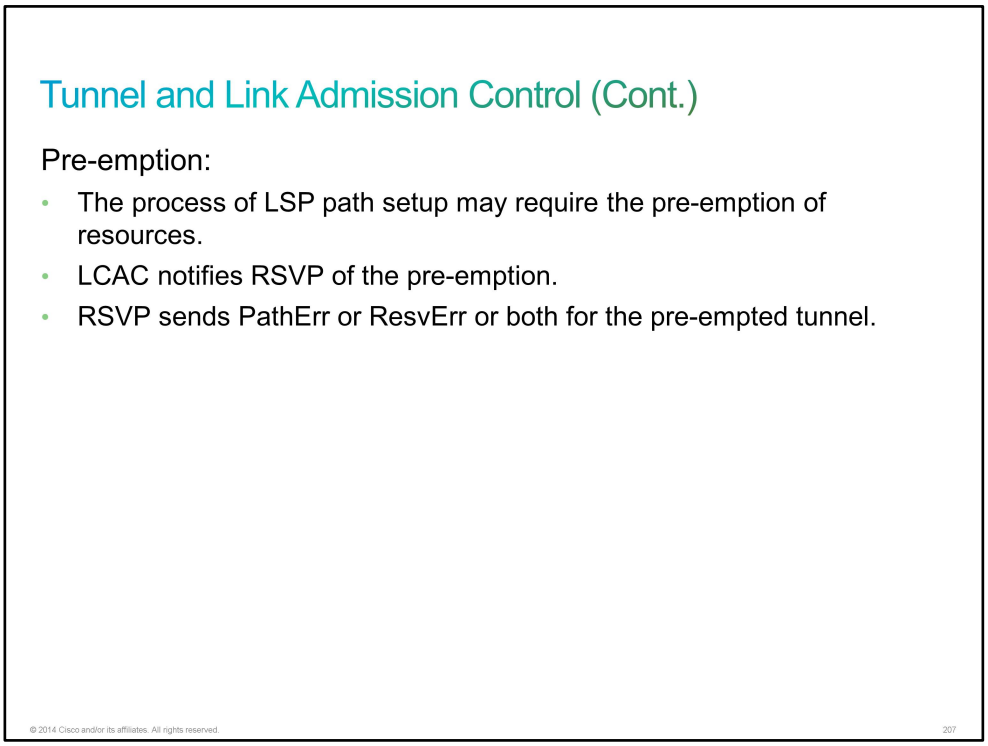

During admission control, the priorities are checked as well.

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.”

https://t.me/learningnets

Path Rerouting

Because the tunnel is not linked to the LSP that is carrying it, the actual path can dynamically change without affecting the tunnel.

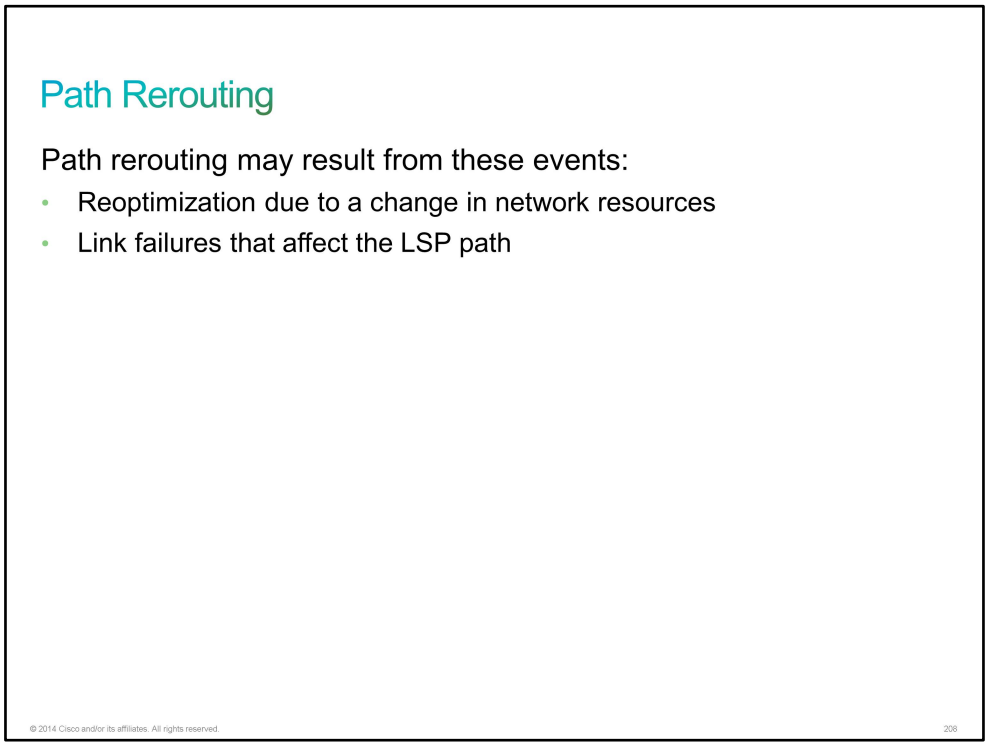

Path rerouting may result from link failures that affect the LSP or from the reoptimization due to a change in network resources.

https://t.me/learningnets

Path Reoptimization

The LSP 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.

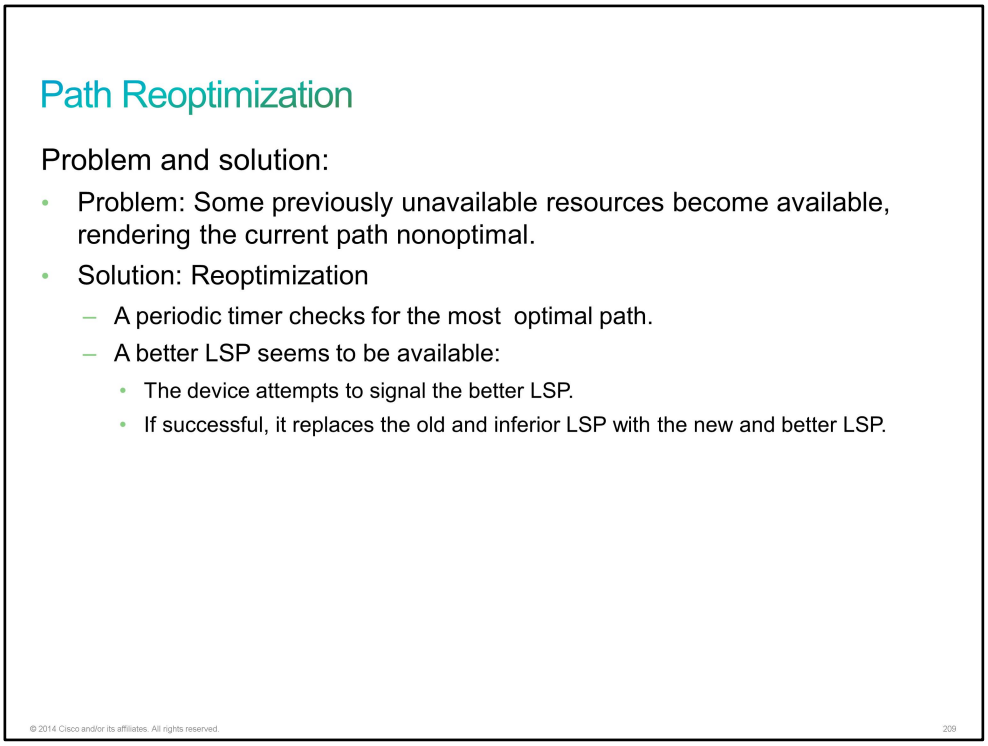

The reoptimization is done on a periodic basis. At certain intervals, a check for the most optimal paths for LSP tunnels is performed and if the current path is not the most optimal, tunnel rerouting is initiated.

The device (headend router) first attempts to signal a better LSP, and only after the new LSP setup has been established successfully, will the traffic be rerouted from the former tunnel to the new one.

Note	Default reoptimization of a TE tunnel is 1 hour.

https://t.me/learningnets

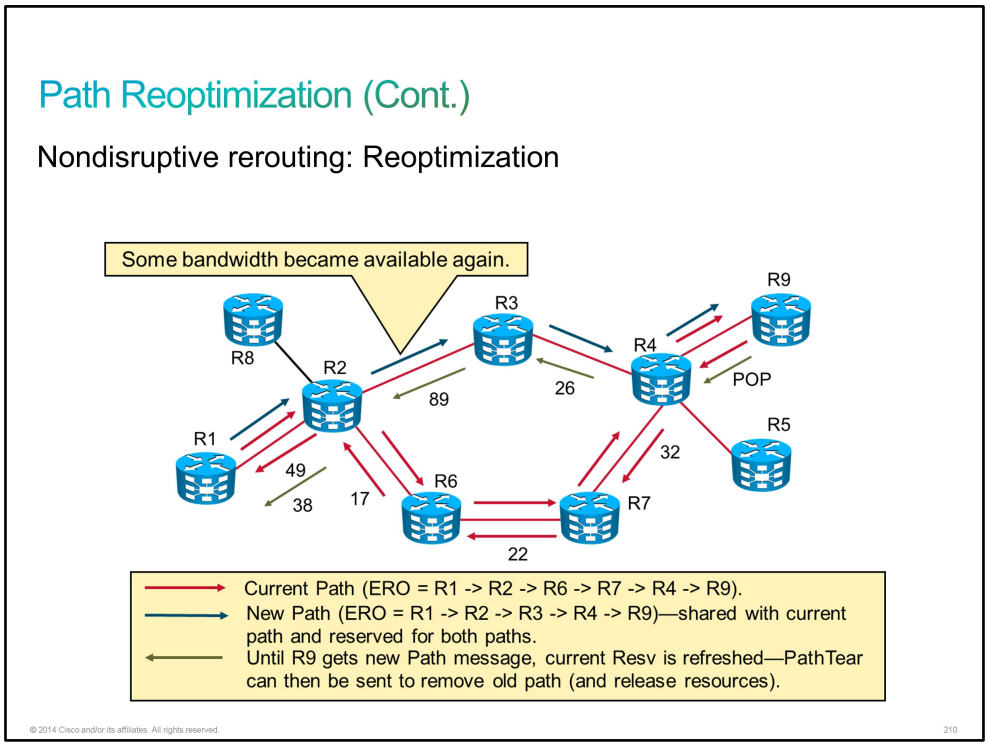

The figure shows how the nondisruptive rerouting of the traffic tunnel is performed. Initially, the ERO lists the LSP R1-R2-R6-R7-R4-R9, with R1 as the headend and R9 as the tail end of the tunnel.

The changes in available bandwidth on the link R2-R3 dictate that the LSP 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 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 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 tail-end router, which initiates an RSVP PathTear message.

The labels that are allocated during the RSVP path setup are shown as well. The tail-end router assigns the implicit null (POP) label.

https://t.me/learningnets

Path Rerouting: Link Failure

When a link failure event occurs, a new path calculation is performed and a new LSP is created.

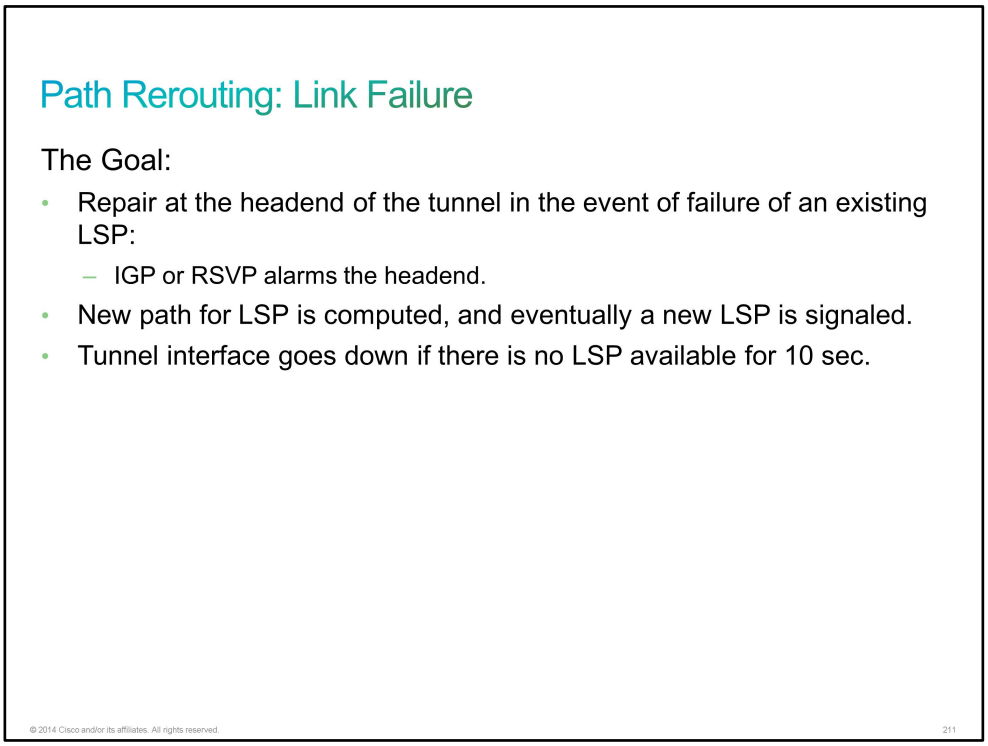

When a link that makes up a certain traffic tunnel fails, the headend of the tunnel detects the failure in one of two ways:

- 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 setup.

Note	The tunnel interface that is used for the specified traffic tunnel (LSP) goes down if the specified LSP is not available for 10 seconds. In the meantime, the traffic that is intended for the tunnel continues using a broken LSP, resulting in black hole routing.

https://t.me/learningnets

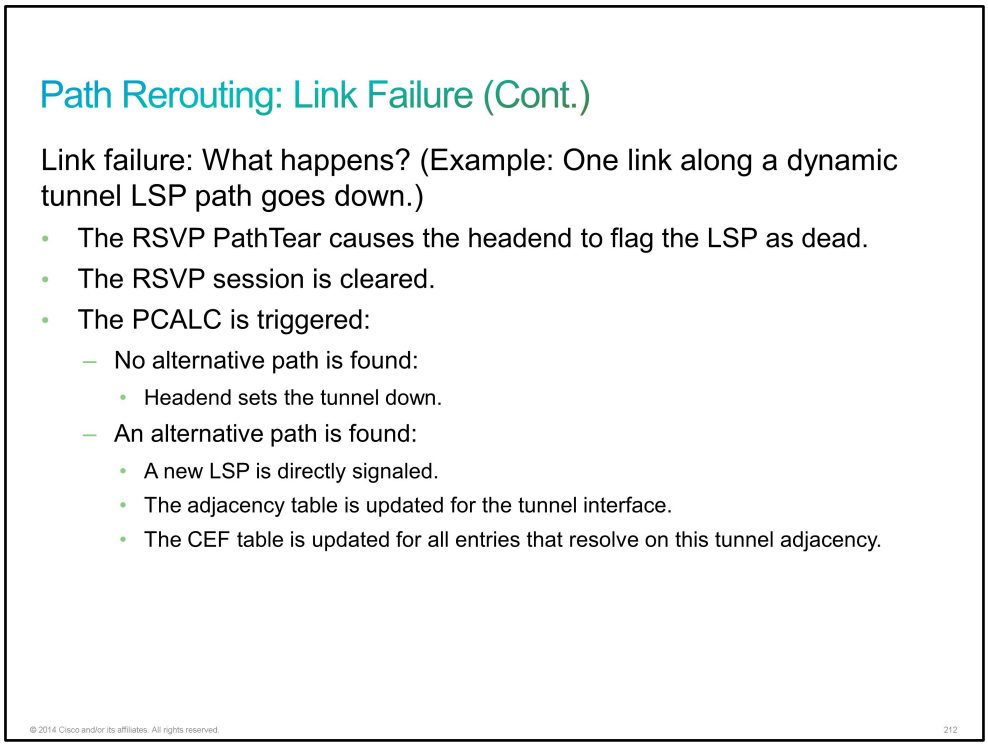

When the router along the dynamic LSP detects a link failure, it sends an RSVP PathTear message to the headend.

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 setup is triggered by RSVP signaling, and adjacency tables for the tunnel interface are updated. Also, the Cisco Express Forwarding table is updated for all the entries that are related to this tunnel adjacency.

https://t.me/learningnets

MPLS TE Soft Pre-emption

The process of freeing resources in the established LSP paths is called MPLS-TE pre-emption.

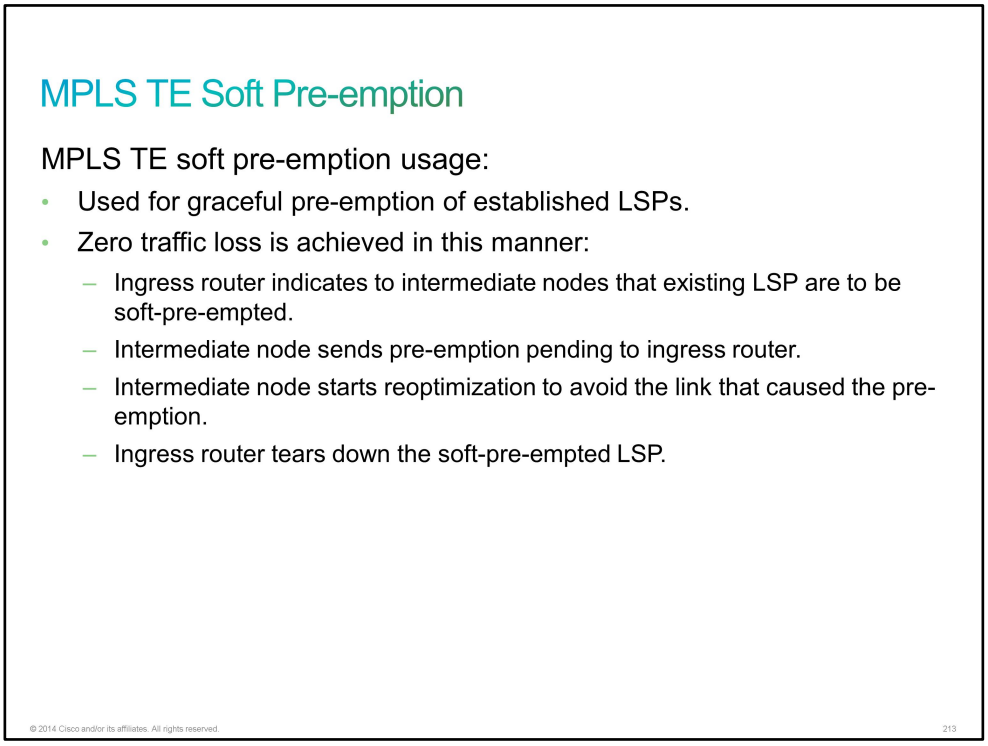

In this process, freed resources are assigned to a new LSP, and it causes a traffic disruption for the LSP that is being pre-empted. In order to minimize a traffic disruption, the soft pre-emption feature is introduced to the RSVP-TE protocol to minimize traffic disrupted on pre-empted LSP.

The soft pre-emption feature tries to pre-empt established LSP in a graceful manner to avoid or minimize traffic loss. In some cases, the link can be oversubscribed for a longer period of time.

No traffic is lost in the network with soft pre-emption, as explained here:

- When sending signaling for a new LSP, ingress router indicates to intermediate nodes that existing LSP is about to be softly pre-empted if resources are needed and are about to be reassigned.
- When a intermediate node needs to soft-pre-empt already established LSP, it will send a special error path message (pre-emption pending) to the ingress router. This intermediate node does not tear down LSP and keeps its state.
- When the error path message (pre-emption pending) is received to ingress router from the intermediate node, it starts reoptimization immediately in order to avoid the link that caused pre-emption.
- Ingress router tears down soft-pre-empted LSP when the reoptimization is completed.

https://t.me/learningnets

MPLS TE Path-Option Attributes

LSPs can be computed with certain path options using attributes specified in the **attribute-set** command. This path options are configurable under template configuration. The attribute-set template is configured globally in the MPLS TE mode. The attribute-set can be applied on a per-LSP basis.

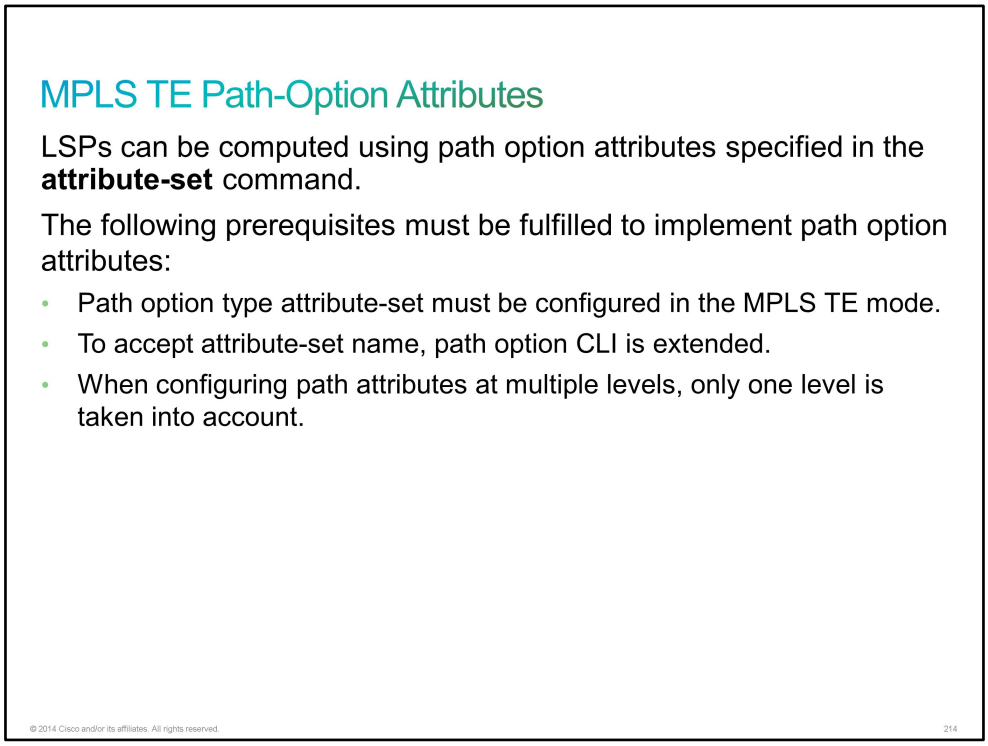

Following prerequisites are necessary to implement path option attributes:

- Path option type attribute-set must be configured in the MPLS TE mode
- To accept the attribute-set name, path option CLI is extended.

When configuring path attributes at multiple levels, such as the LSP level or tunnel level, only one level is taken into account.
Attributes that are not specified within an attribute-set take their values from configuration at the tunnel level.

https://t.me/learningnets

Assigning Traffic to Traffic Tunnels

The LSP is computed by CBR, which takes the resource requirements into consideration as well. When the LSP is established for the tunnel, the traffic can flow across it. From the IP perspective, an LSP is a simple tunnel.

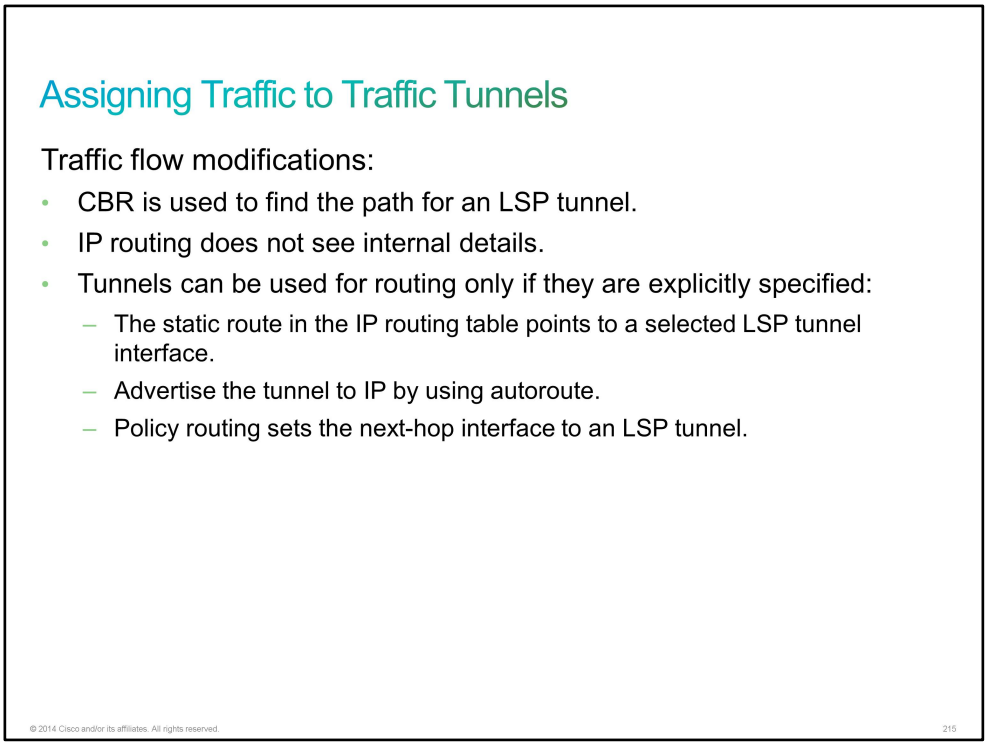

These engineered tunnels can be used for IP routing only if the tunnels are explicitly specified for routing in one of these ways:

- Via static routes that point to the tunnel
- Via the autoroute feature, which advertises the MPLS TE tunnel to IP protocol
- Via policy routing that sets a next-hop interface to the tunnel

https://t.me/learningnets

Using Static Routing to Assign Traffic to Traffic Tunnel

The sample topology below is shown to illustrate the assigning of traffic to a tunnel.

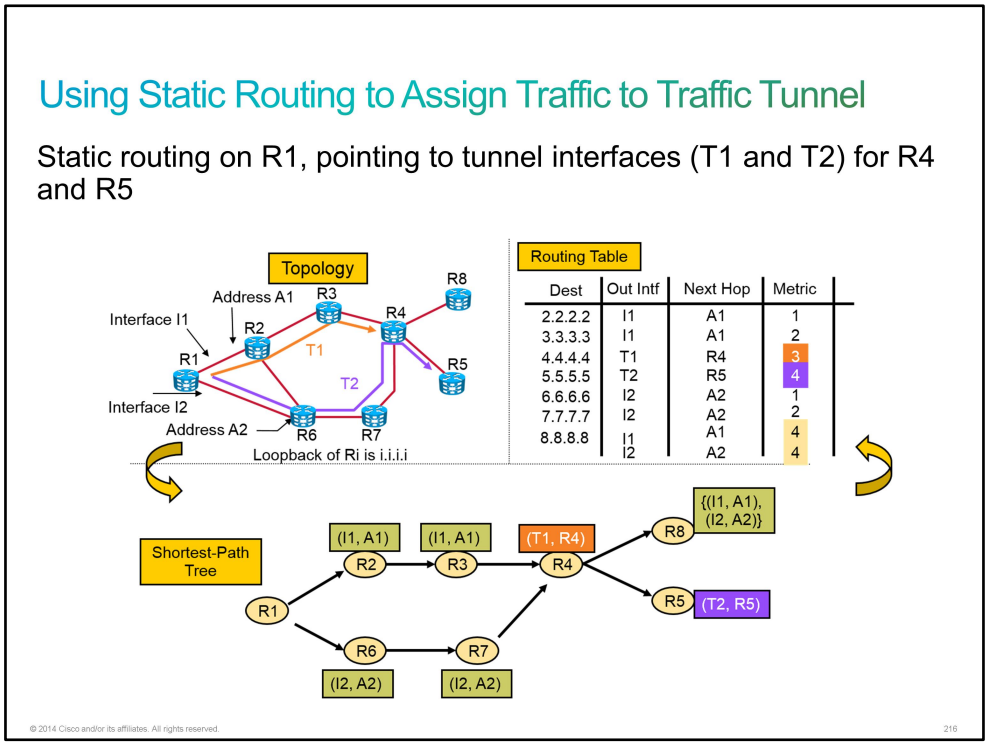

The sample topology here shows two engineered tunnels: T1 (between R1 and R4) and T2 (between R1 and R5). The loopback addresses on each router are in the form i.i.i.i where i is the router number (for example, the R5 loopback address is 5.5.5.5). The metric on each of the interfaces is set to 1.

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; however, a constraint-based computation is performed for the paths and for the tunnels.

https://t.me/learningnets

Autoroute

To overcome the problems that result from static routing configurations onto MPLS TE tunnels, the autoroute feature was introduced.

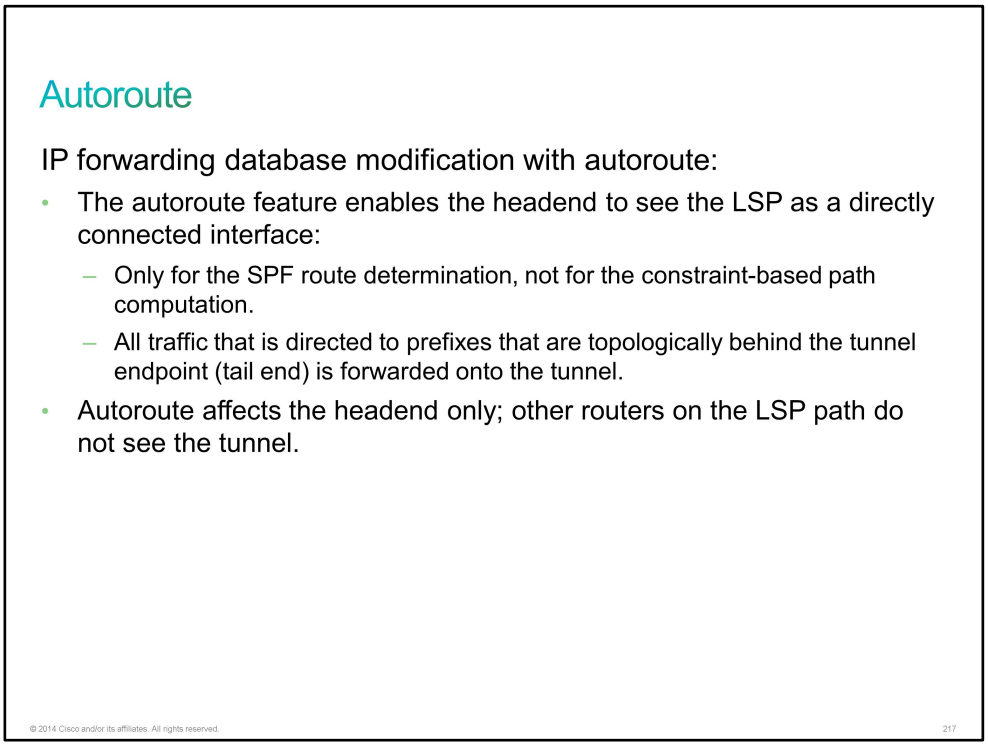

The autoroute feature enables the headend router to see the MPLS TE tunnel as a directly connected interface and to use it in its modified SPF computations.

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 (tail end) 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.

https://t.me/learningnets

Autoroute: Path Selection Rules

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. The tunnel metric is tunable using either relative or absolute metrics.

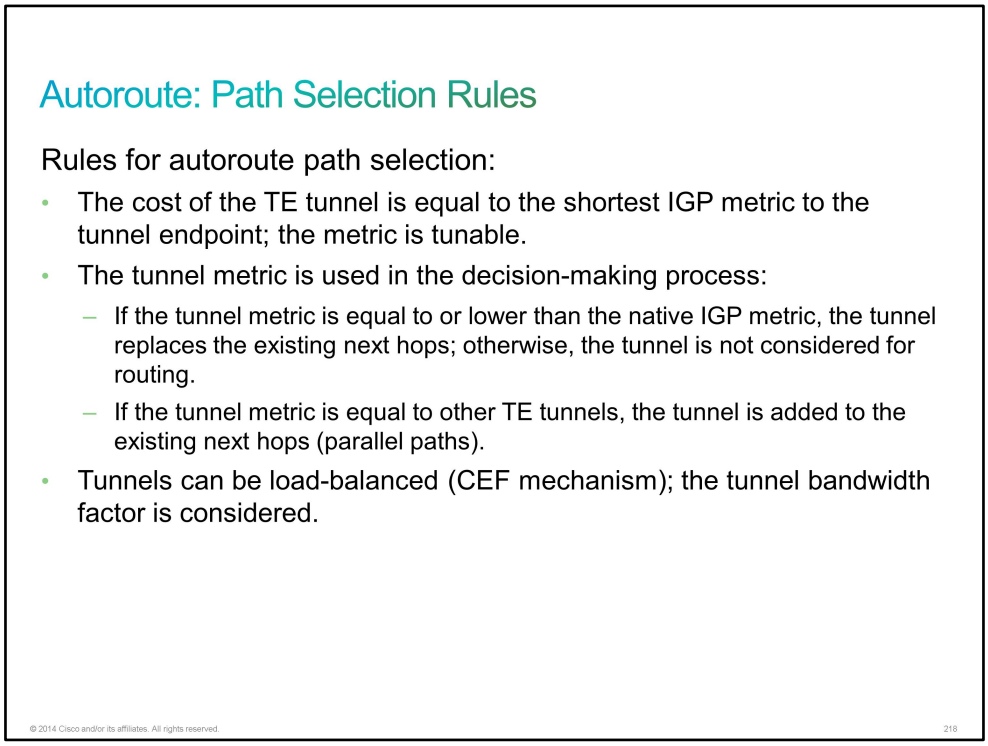

During installation of the best paths to the destination, the tunnel metric is compared to other existing tunnel metrics and to all the native IGP path metrics. The lower metric is better and if the MPLS TE tunnel has an equal or lower metric than the native IGP metric, it is installed as a next hop 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.

https://t.me/learningnets

MPLS TE IPv6 Autoroute

The MPLS TE IPv6 autoroute feature is used for IPv4 MPLS TE tunnels using IPv6 routing. First of all, the tunnel must be IPv6-enabled so it can carry IPv6 traffic. If the tunnel is advertised to carry IPv6 traffic, the routing protocol considers MPLS TEv4 tunnel for IPv6 routing path calculation. In order to advertise the tunnel, IPv6 forwarding adjacency or IPv6 autoroute announce must be configured on the tunnel.

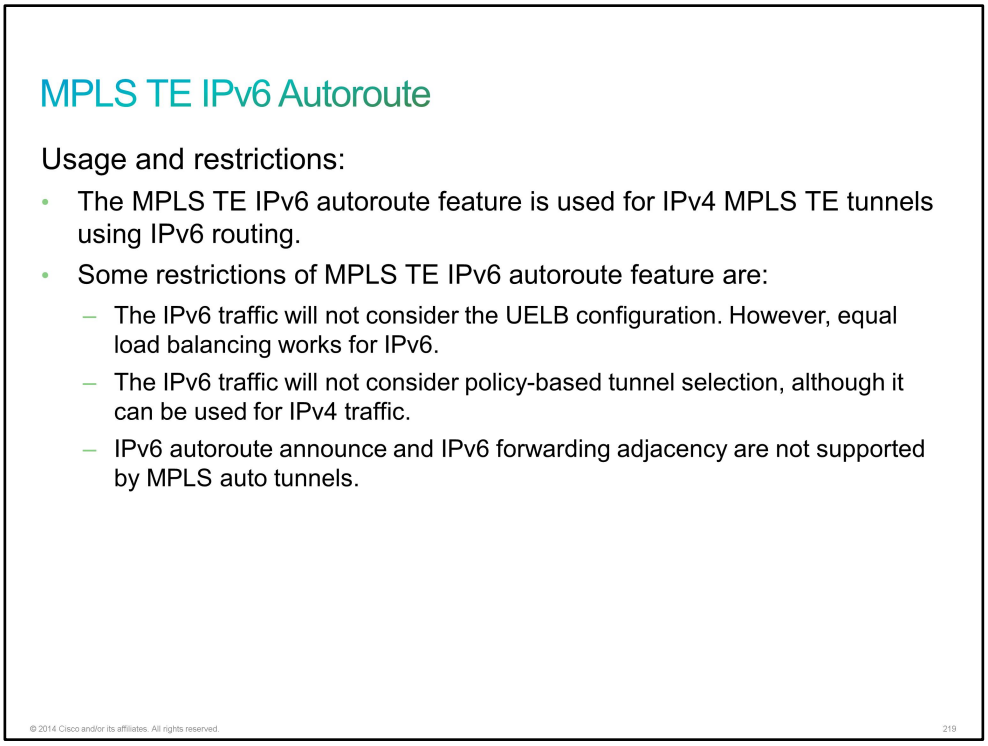

Some restrictions of the MPLS TE IPv6 autoroute feature are:

- The IPv6 traffic will not consider the UELB configuration. However, equal load balancing works for IPv6.
- The IPv6 traffic will not consider policy-based tunnel selection, although it can be used for IPv4 traffic.
- IPv6 autoroute announce and IPv6 forwarding adjacency are not supported by MPLS auto tunnels.

https://t.me/learningnets

Autoroute: Default Metric

The sample topology below illustrates the autoroute default metric with two engineered tunnels.

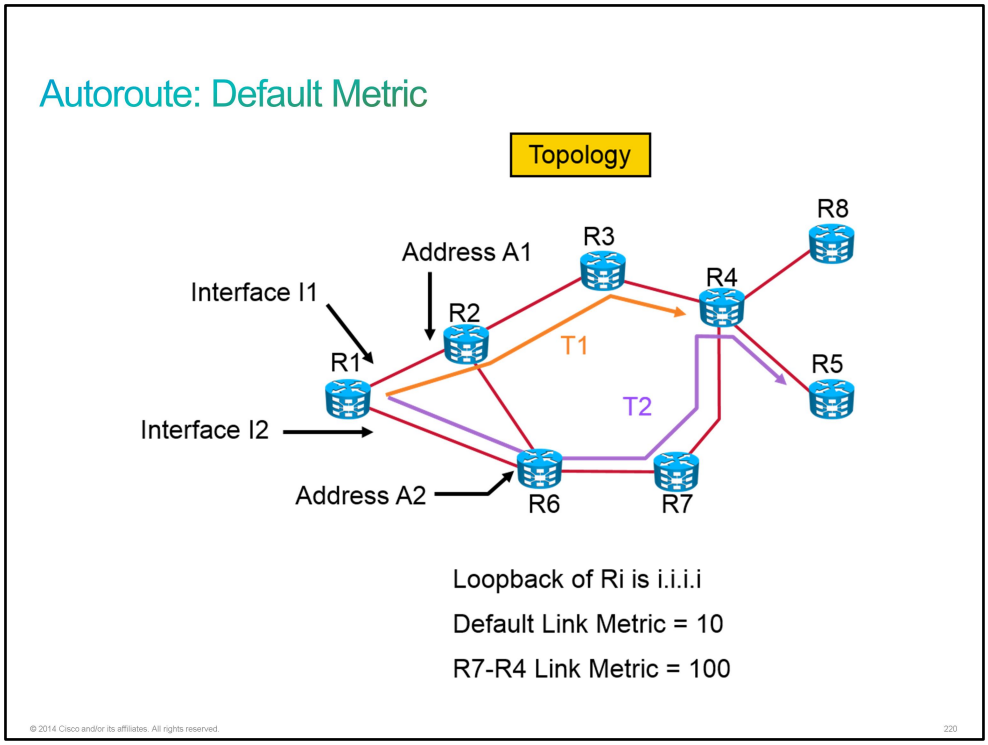

The sample topology shows two engineered tunnels: T1 (between R1 and R4) and T2 (between R1 and R5). The loopback addresses on each router are in the form i.i.i.i where i is the router number (for example, the R5 loopback address is 5.5.5.5). The metric on each of the interfaces is set to 10.

R1 has two physical interfaces, I1 and I2, and two neighboring routers (next hops) with addresses of A1 and A2, respectively.

https://t.me/learningnets

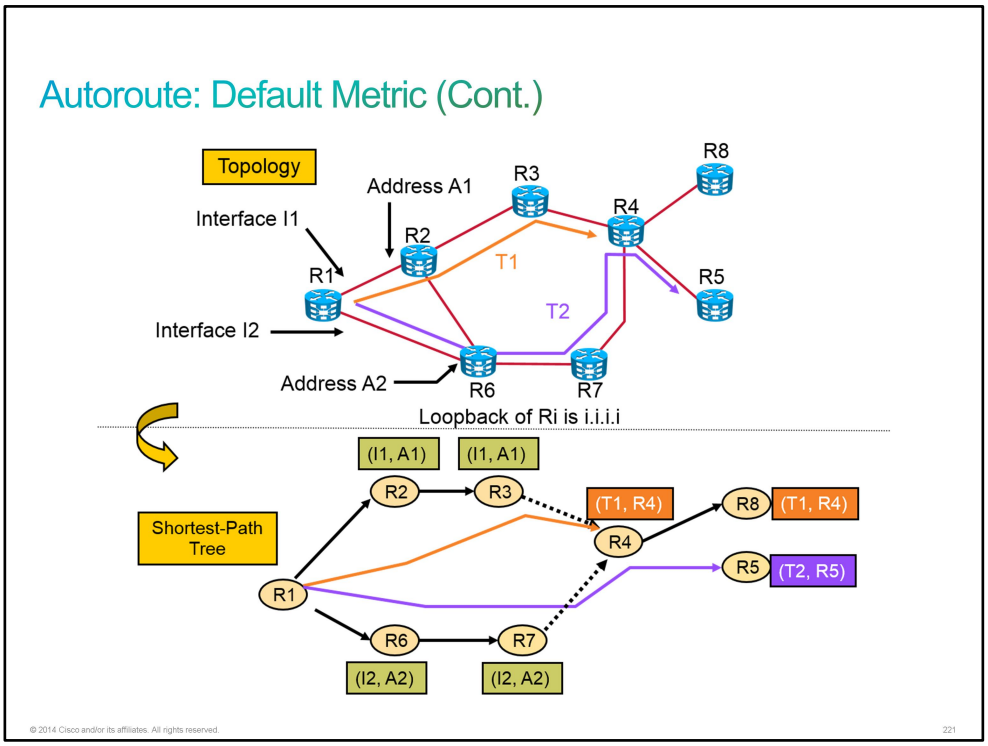

This example shows the resulting shortest-path tree from router R1. In this situation, the tunnels are seen for routing purposes only by the headend. Intermediate routers do not see the tunnel, nor do they take it into consideration for route calculations.

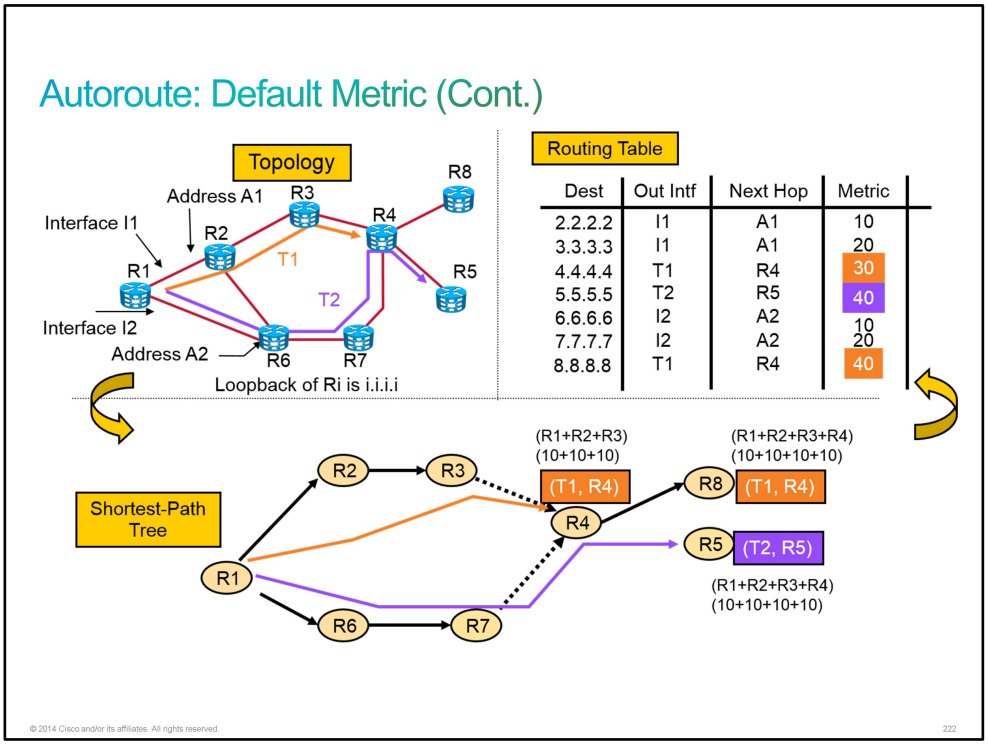

The routing table shows all destinations at the endpoint of the tunnel and behind it (R8) as reachable via the tunnel itself. The metric to the destination is the normal IGP metric.

The LSP 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).

https://t.me/learningnets

The LSP for T2 follows the path R1-R6-R7-R4-R5. Although the LSP 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.

https://t.me/learningnets

Autoroute: Relative and Absolute Metric

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.

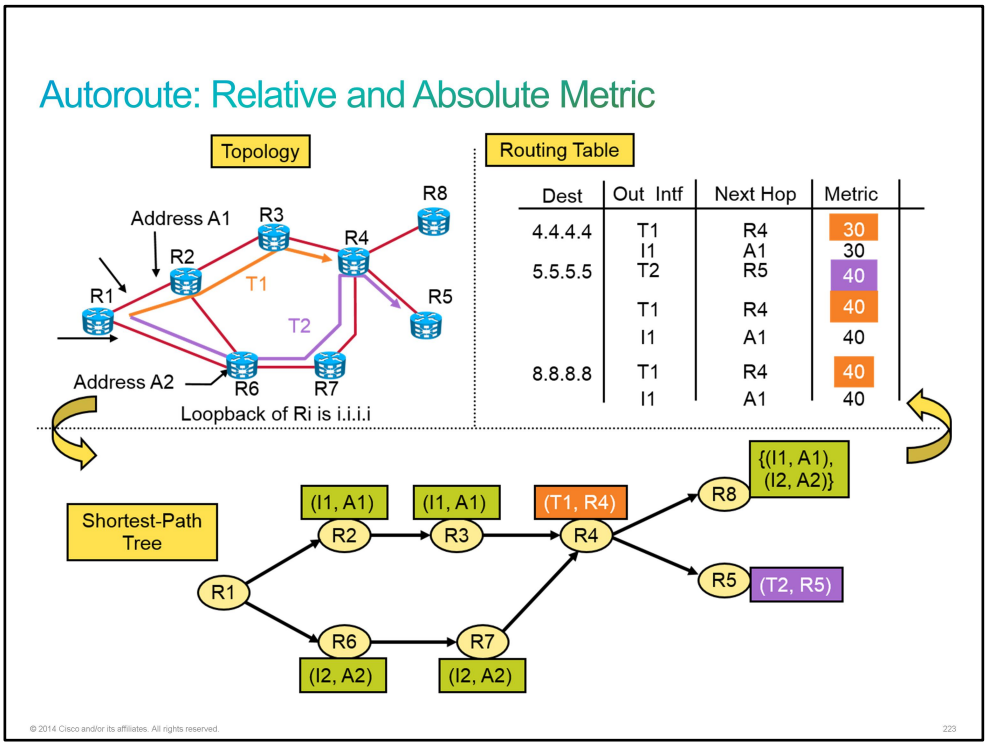

In this example, there appear to be two paths to R4, while there is only one physical path. For R5, there appear to be three paths, two of which do not follow the desired tunnel path.

The tunnel metrics can be tuned, and either relative or absolute metrics can be used to resolve this issue.

https://t.me/learningnets

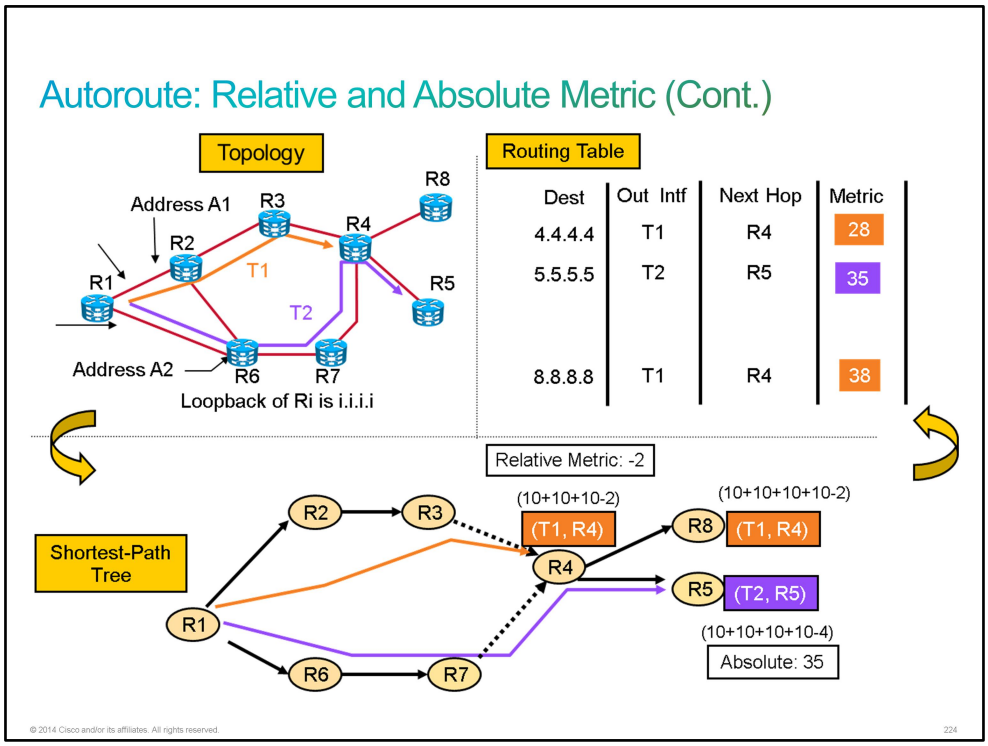

In this example, the relative metric is used to control path selection. T1 is given a value of the native IGP value –2 (28), making it preferred to the native IP path. When the tunnel is considered in the IGP calculation, the native IGP metric (30) is greater than the tunnel metric (28) for all the destinations 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 IGP value –4 (36), giving it preference to the native IP path and the T1-R4 path. However, in this example, it was given an absolute value of 35.

https://t.me/learningnets

Forwarding Adjacency

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.

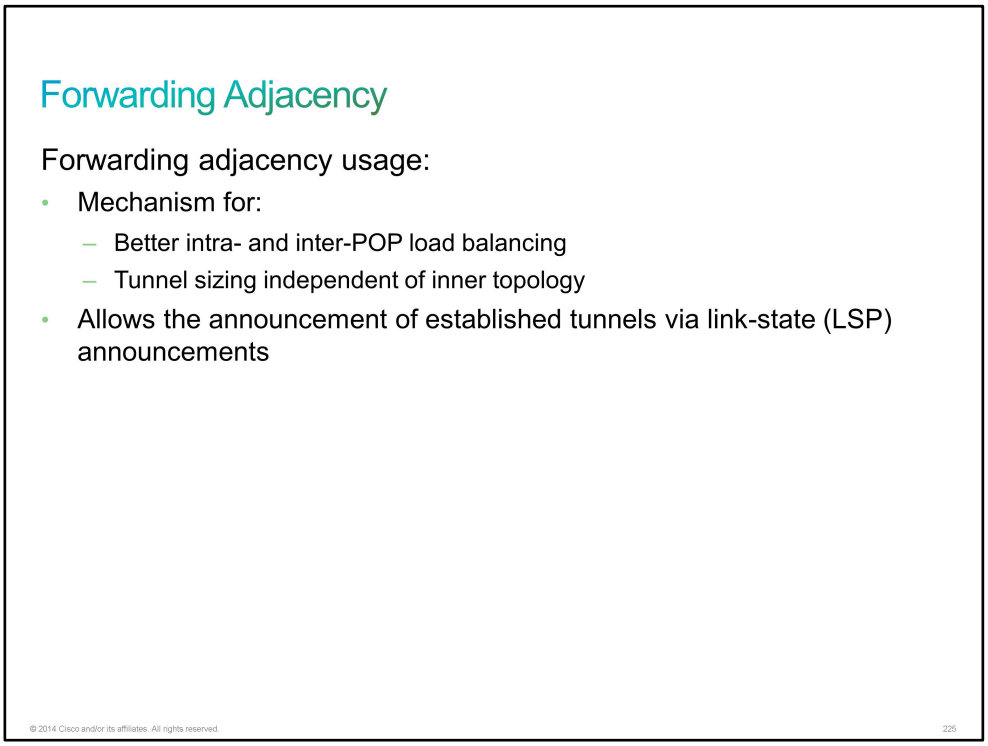

A forwarding adjacency can be created between routers regardless of their location in the network.

A 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

https://t.me/learningnets

Forwarding Adjacency Traffic Flows

Before considering the real benefits of forwarding agency, clearly seeing the limitations of autoroute in certain network topologies is important.

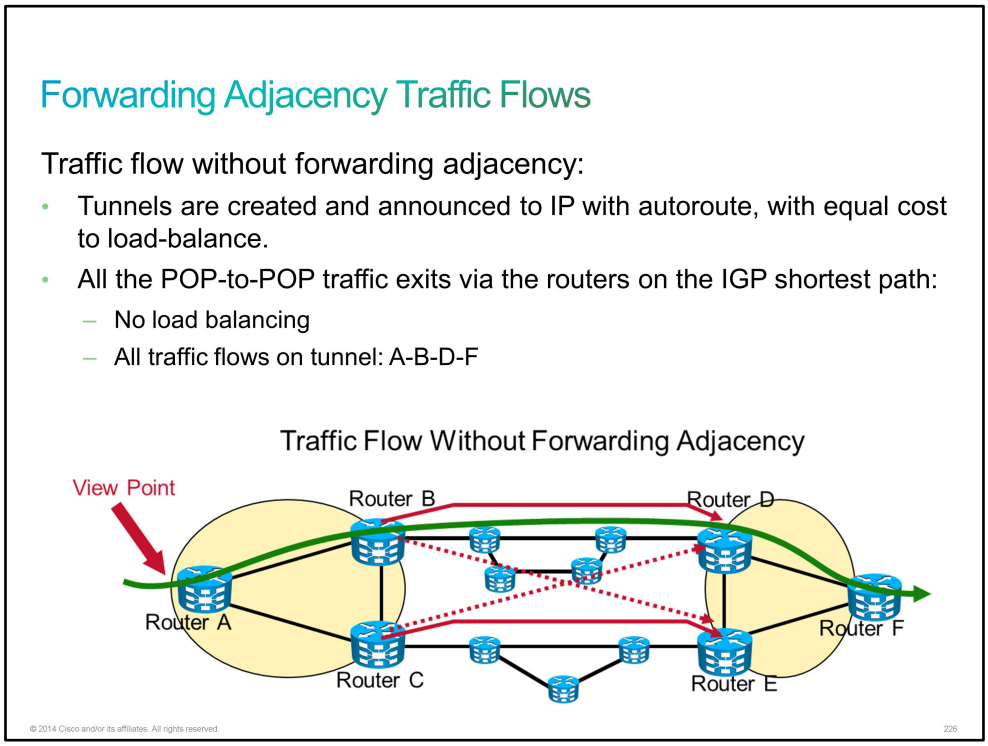

In this example, established tunnels exist from B to D, from B to E, from C to E, and from C to D; the preferred tunnels are B to D and C to E.

- 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. It only has its IGP information, indicating that the better path to F leads via routers 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 will affect any load balancing for POP-to-POP traffic.

Note	You can theoretically prevent this problem by creating tunnels from a router to any other router, but this design does not scale in service provider networks.

https://t.me/learningnets

Forwarding Adjacency Traffic Flows (Cont.)

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-B-D-F
  - Link failure: All traffic flows A-C-E-F

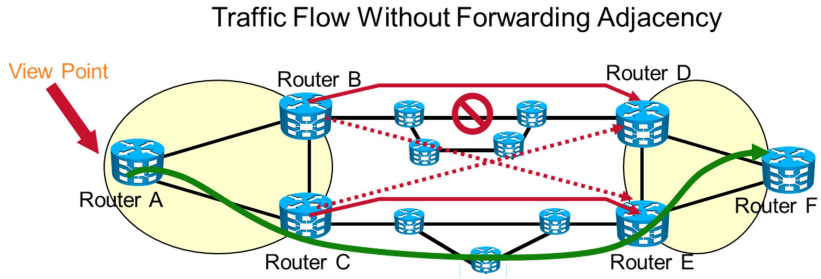

© 2014 Cisco and/or its affiliates. All rights reserved.

What happens if 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 this example, the change in the metric will result in a possibly unplanned switchover of the traffic from A to F from the upper path 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 is idled out.

Forwarding Adjacency Traffic Flows (Cont.)

POP-to-POP traffic is better load-balanced:

- In the POP, if the two core routers are used
- In the core, if at least two tunnels are used
- As long as the IGP metric for a path with the forwarding adjacency (for example, 25) is shorter than the forwarding adjacency-free path (for example, 30).

Inner topology does not affect tunnel sizing:

- Change in the core topology does not affect load balancing in the POP.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

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.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Dynamic constraint-based path computation is triggered by the headend of the tunnel.
- There are four link-resource attributes.
- The most important link resource attribute is the maximum allocation multiplier. Traffic tunnel attributes affect how the path is set up and maintained.
- The link resource class attribute string allows inclusion or exclusion of the link into the path of the tunnel.
- Each link has a cost or metric for calculating routes in the normal operation of the IGP.
- Seven tunnel attributes are available to influence path selection.
- The tunnel path can be configured manually or computed dynamically by using the constraint-based path computation.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- The tunnel resource class affinity attribute allows the network administrator to apply path selection policies by administratively including or excluding network links.
- The adaptability attribute indicates whether the traffic tunnel should be reoptimized, and consequently rerouted to another path.
- The resilience attribute determines the behavior of the tunnel under faulty conditions.
- The policies during LSP 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.
- Important factor in LSP computation is the available bandwidth on the link over which the traffic tunnel will pass.
- The result of the constraint-based path computation is a series of IP addresses that represent the hops on the LSP between the headend and tail end of the traffic tunnel.
- LSP setup is always initiated at the traffic tunnel headend.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Summary (Cont.)

- Each hop on the way determines if the available resources that are specified in the session attribute object are available.
- Path rerouting may result from reoptimization due to a change in network resources or link failures that affect the LSP.
- The MPLS TE path pre-emption enhancement is an extension to the RSVP-TE protocol to minimize and eliminate the traffic disruption over the pre-empted LSP.
- MPLS TE path option attributes are configurable through a template configuration. This template named attribute-set, is configured globally in the MPLS TE mode.
- TE tunnels can be used for IP routing only if the tunnels are explicitly specified for routing.
- Even for the destination that is behind each of the tunnel endpoints, the normal IGP routing is performed if there is no static route to the traffic-engineered tunnel.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- 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.
- When using the autoroute feature, all the networks in the routing table that are topologically behind the tunnel endpoint are reachable via the tunnel itself.
- The tunnel metrics can be tuned, and either relative or absolute metrics can be used to resolve this issue.
- 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.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Implementing MPLS TE

Overview

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 are presented. This lesson describes advanced MPLS TE commands that are used in path selection in a typical service provider environment. The configuration commands are accompanied by usage guidelines and sample setups.

Upon completing this lesson, you will be able to:

- List the MPLS TE configuration tasks
- Explain the commands to enable MPLS TE
- Explain the commands to enable RSVP
- Explain the commands to enable OSPF to support MPLS TE
- Explain the commands to enable IS-IS to support MPLS TE
- Explain the commands to enable MPLS TE tunnels
- Explain the commands to enable static routing and autoroute
- Describe how to verify the RSVP operation
- Describe the how to verify MPLS TE tunnels
- Describe how to verify a routing table with the autoroute feature enabled
- Explain a MPLS TE case study
- Explain creating a dynamic MPLS TE tunnel using a case study
- Explain creating an explicit MPLS TE tunnel using a case study
- Explain the enabling of periodic tunnel optimization using a case study

https://t.me/learningnets

- Explain enabling path selection restrictions using a case study
- Explain modifying the administrative weight using a case study
- Explain enabling autoroute and forwarding adjacency using a case study

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS TE Configuration Tasks

To configure MPLS TE tunnels, MPLS should already be enabled in the core network.

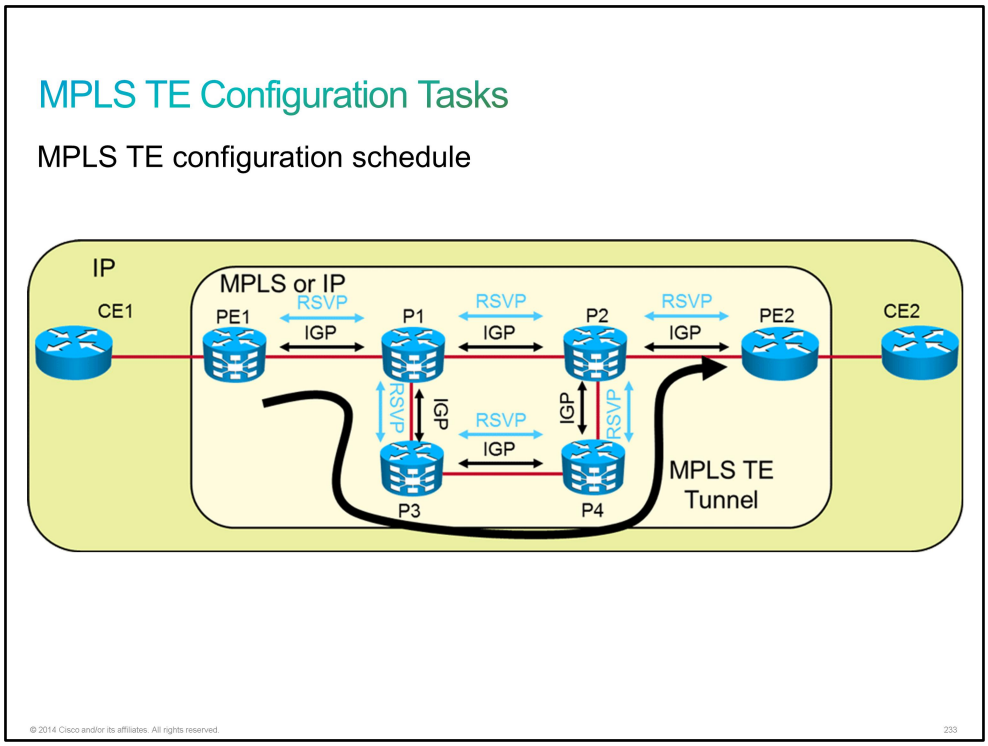

These steps are required to configure MPLS TE tunnels:

- Enable MPLS TE in the core.
- Configure RSVP in the core.
- Enable MPLS TE support in the core IGP with OSPF or IS-IS.
- Configure MPLS TE tunnels.
- Configure routing onto MPLS TE tunnels.

Configuration for each of these required steps for building MPLS TE tunnels between routers will be shown.

https://t.me/learningnets

MPLS TE Configuration

When configuring MPLS TE tunnels, the first step is to enable MPLS TE functionality.

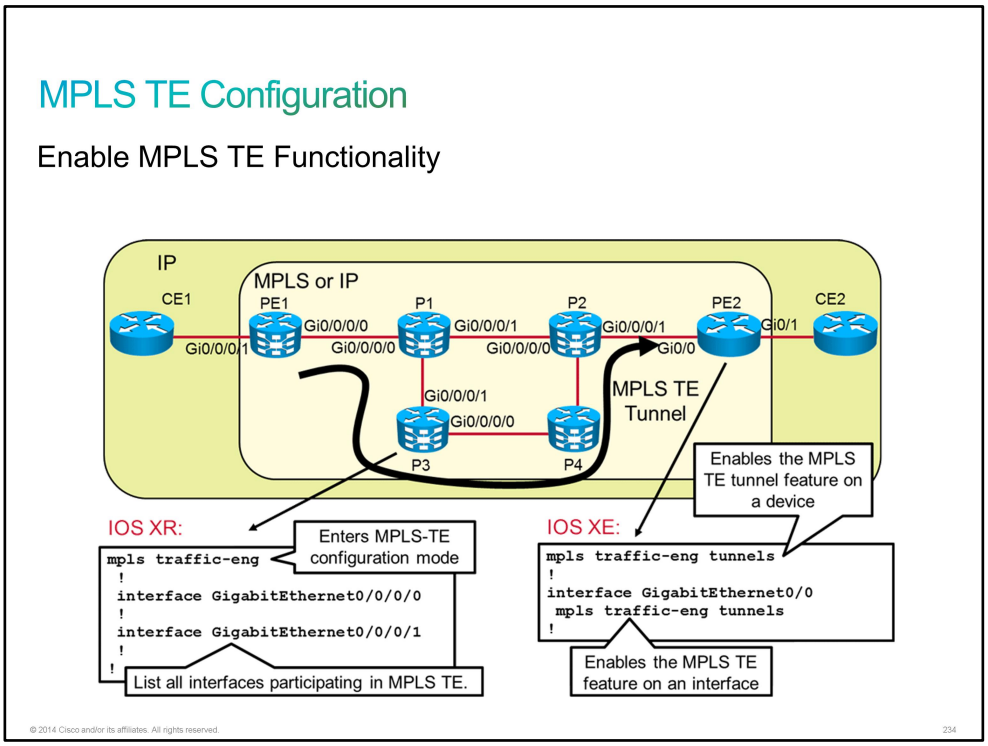

To enter MPLS TE configuration mode on Cisco IOS XR Software, use the **mpls traffic-eng** command in global configuration mode. List all interfaces participating in MPLS TE in mpls traffic engineering configuration mode.

To enable MPLS TE tunnel signaling on a Cisco IOS XE device, use the **mpls traffic-eng tunnels** command in global configuration mode. To disable this feature, use the **no** form of this command.

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. An enabled interface has its resource information flooded into the appropriate IGP link-state database, and accepts TE tunnel signaling requests.

Note	MPLS TE functionality should be enabled on all routers on the path from headend to tail end of the MPLS TE tunnel.

https://t.me/learningnets

RSVP Configuration

When configuring MPLS TE tunnels, you need to configure RSVP support in the core.

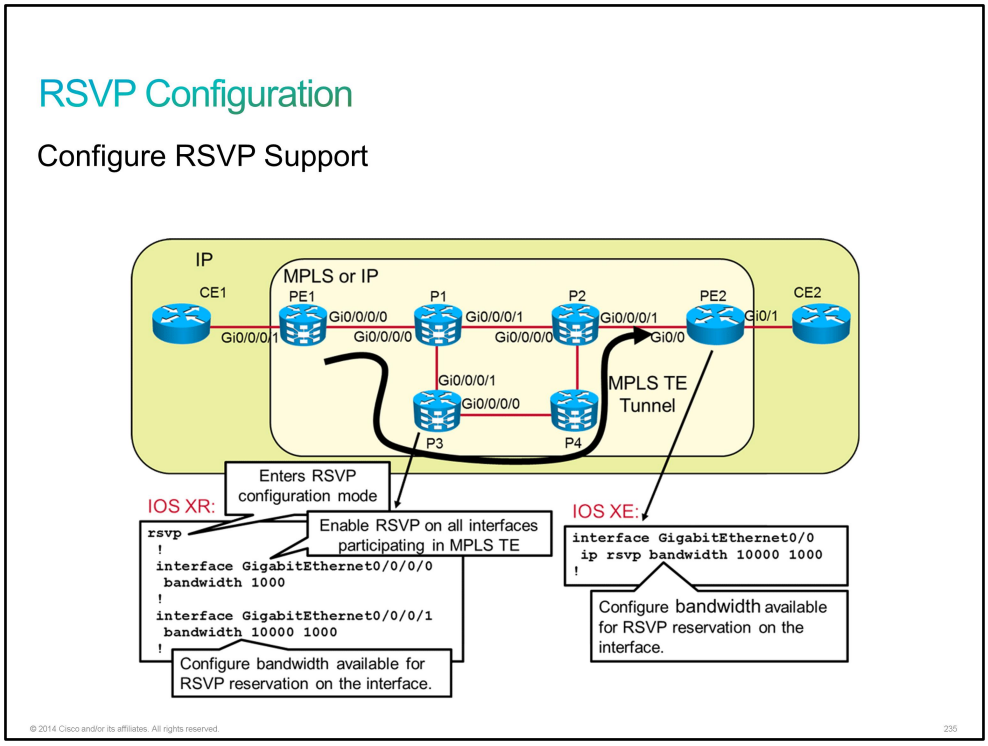

Enable RSVP on all interfaces that are participating in MPLS TE and configure the bandwidth that is available for RSVP reservation.

To enter RSVP configuration mode, use the **rsvp** global configuration command on Cisco IOS XR Software. To enter interface configuration mode for the RSVP protocol, use the **interface** command; use the **bandwidth** command to set the reservable bandwidth, the maximum RSVP bandwidth that is available for a flow and the subpool bandwidth on this interface.

```
bandwidth total-bandwidth max-flow sub-pool sub-pool-bw
```

Syntax Description

Parameter	Description
total-bandwidth (interface-kbps)	(Optional) Maximum amount of bandwidth, in kbps, that may be allocated by RSVP flows
max-flow (single-flow-kbps)	(Optional) Maximum amount of bandwidth, in kbps, that may be allocated to a single flow

To enable RSVP for IP on an interface, use the **ip rsvp bandwidth** interface configuration command on Cisco IOS XE Software. To disable RSVP, use the **no** form of this command.

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


https://t.me/learningnets

Note	RSVP support should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.

https://t.me/learningnets

OSPF Configuration

When configuring MPLS TE tunnels, you need to enable MPLS TE support in routing protocol (OSPF in this case).

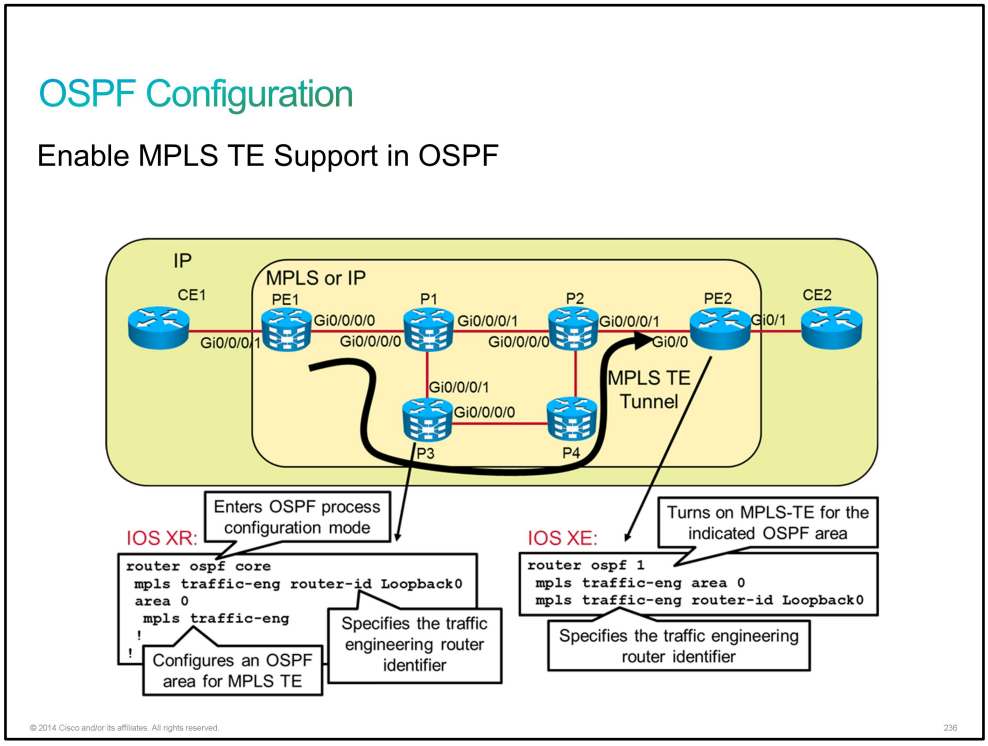

One of the required steps to configure MPLS TE tunnels is enabling MPLS TE support in the IGP routing protocol (OSPF or IS-IS).

To enable MPLS TE support for OSPF routing protocol on Cisco IOS XR Software, enter the OSPF process configuration mode and use the **mpls traffic-eng router-id** *interface* command to specify that the TE router identifier for the node is the IP address associated with a given OSPF interface.

```
mpls traffic-eng router-id {router-id | interface-type interface-instance}
```

A router identifier must be present in IGP configuration. 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.

MPLS TE must be enabled under area configuration. To configure an OSPF area for MPLS TE, use the **mpls traffic-eng** command in the appropriate configuration mode.

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 on Cisco IOS XE Software.

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
```


https://t.me/learningnets

Syntax Description

Parameter	Description
interface	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.

Note	MPLS TE support for IGP should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

IS-IS Configuration

When configuring MPLS TE tunnels, you need to enable MPLS TE support in routing protocol (IS-IS in this case).

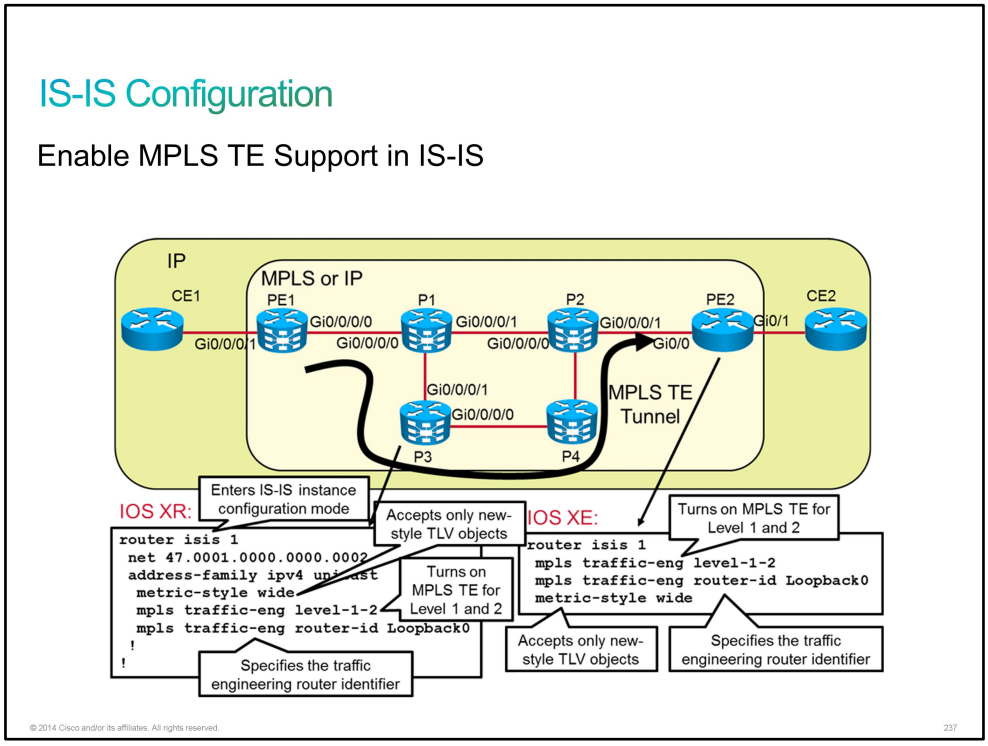

To enable MPLS TE support for IS-IS routing protocol on Cisco IOS XR Software, enter the address family configuration mode for configuring IS-IS routing and use the **mpls traffic-eng router-id** *interface* command to specify that the traffic engineering router identifier for the node is the IP address associated with a given OSPF interface.

```
mpls traffic-eng router-id {router-id | interface-type interface-instance}
```

To configure MPLS TE at IS-IS Level 1 and Level 2 on a router that is running IS-IS, use the **mpls traffic-eng level** command in address family configuration mode:

```
mpls traffic-eng level isis-level
```

To configure the IS-IS software to generate and accept only new-style TLVobjects, use the **metric-style wide** command in address family configuration mode.

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 on Cisco IOS XE Software. 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 {level-1 | level-2}
```

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 on Cisco IOS XE Software. To disable this feature, use the **no** form of this command.

```
mpls traffic-eng router-id interface
```


https://t.me/learningnets

To configure a router running IS-IS so that it generates and accepts only new-style TLVs objects, use the **metric-style wide** command in router configuration mode on Cisco IOS XE Software. To disable this function, use the **no** form of this command.

```
metric-style wide [transition] [level-1 | level-2 | level-1-2]
```

Note	MPLS TE support for IGP should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS TE Tunnels Configuration

After you enabled the MPLS TE functionality and configured RSVP and MPLS TE support for the routing protocol (either OSPF or IS-IS), you can start configuring the MPLS TE tunnel itself and the parameter of the tunnel.

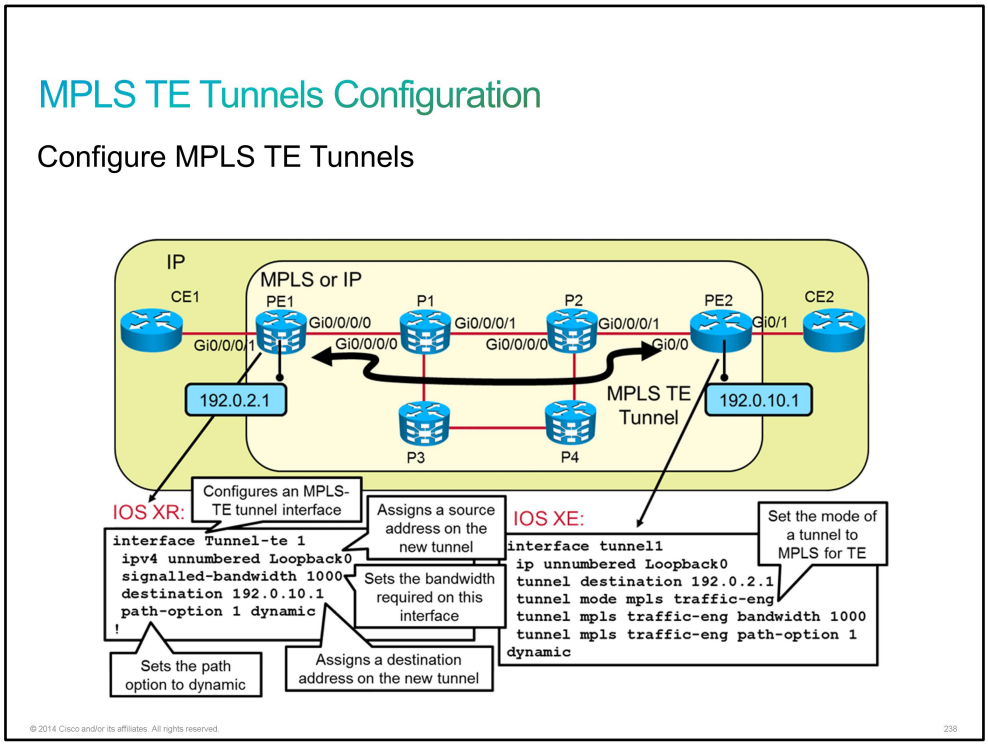

This figure shows a typical implementation of a dynamic MPLS TE tunnel. Two tunnels are created: one from router PE1 to PE2 and one from PE2 to PE1. The actual path could be PE1-P1-P2-PE2 or PE1-P1-P3-P4-P2-PE2, as selected by IGP protocol, because the path option of the tunnel is set to dynamic. Alternatively, you can select the explicit path option, where you can manually specify the desired path of the MPLS TE tunnel.

Note	The MPLS TE tunnel is unidirectional.

Use the **interface tunnel-te** *tunnel-id* command to configure an MPLS TE tunnel interface on Cisco IOS XR Software. You can set several MPLS TE tunnel parameters in interface tunnel-te configuration mode:

- Use the **ipv4 unnumbered** *interface* command in interface tunnel-te configuration mode to assign a source address, so that forwarding can be performed on the new tunnel. Loopback is commonly used as the interface type.
- Use the **destination** *ip-address* command to assign a destination address on the new tunnel. The destination address is the MPLS TE router ID of the remote node.
- Use the **signalled-bandwidth** *bandwidth* command to set the bandwidth that is required on this tunnel-te interface.
- Use the **path-option** *priority* **dynamic** command to set the path option to dynamic.

https://t.me/learningnets

Use the **interface tunnel** command to declare a TSP tunnel interface on Cisco IOS XE Software. You can set several MPLS TE tunnel parameters in interface tunnel configuration mode:

- Use the **ip unnumbered** *interface* command in interface tunnel configuration mode to assign a source address, so that forwarding can be performed on the new tunnel. Loopback is commonly used as the interface type.
- Use the **tunnel destination** *ip-address* command to specify the destination for a tunnel interface.
- Use the **tunnel mode mpls traffic-eng** command to set the mode of a tunnel to MPLS for TE.
- Use the **tunnel mpls traffic-eng bandwidth** command to configure the bandwidth that is required for an MPLS TE tunnel. Bandwidth is specified in kbps.
- Use the **tunnel mpls traffic-eng path-option** command to configure a path option for an MPLS TE tunnel.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Static Route and Autoroute Configurations

After your tunnel is up and running, you need to route traffic into an MPLS TE.

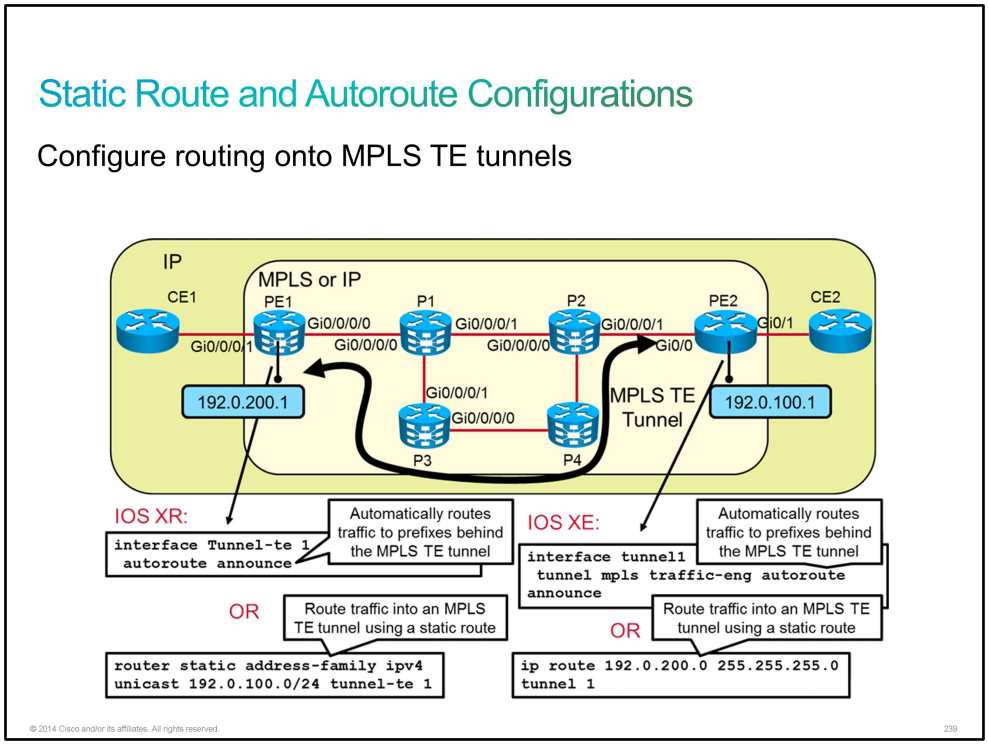

The autoroute feature evaluates MPLS TE tunnel interface with an IGP, and automatically routes traffic to prefixes behind the MPLS TE tunnel, based on IGP metrics.

Use the **autoroute announce** command in interface configuration mode to specify that the IGP should use the tunnel (if the tunnel is up) in its enhanced SPF calculation on Cisco IOS XR Software.

Another option to route traffic into an MPLS TE tunnel is by using a static route. You should route traffic to the prefixes behind MPLS tunnel l to the interface tunnel, in the example tunnel-te1. This configuration is used for static routes when the **autoroute announce** command is not used.

To instruct the IGP to use the tunnel in its SPF calculation (if the tunnel is up) on Cisco IOS XE Software, use the **tunnel mpls traffic-eng autoroute announce** command in interface configuration mode.

https://t.me/learningnets

Verifying RSVP Operation

The first step when monitoring MPLS TE operation is verifying RSVP sessions.

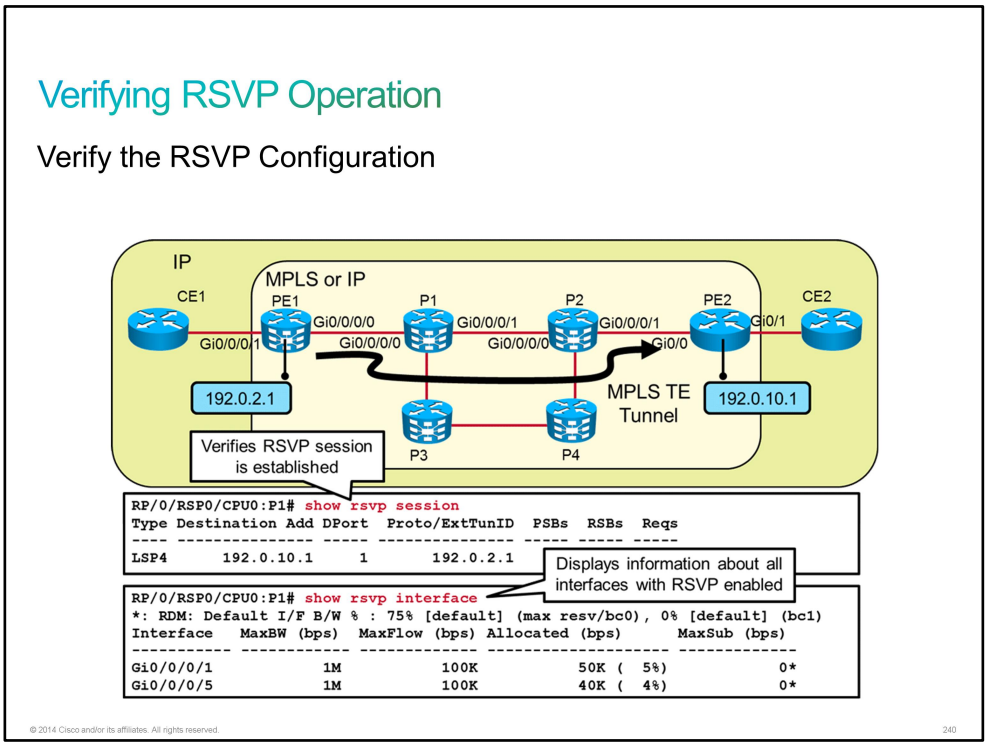

The **show rsvp session** command verifies that all routers on the path of the LSP are configured with at least one PSB and one RSB per session. In the example, the output represents an LSP from the ingress (head) router 192.0.2.1 to the egress (tail) router 192.0.10.1.

To display information about all interfaces with RSVP enabled, use the **show rsvp interface** command in EXEC mode. You can also see the information about the allocated bandwidth to MPLS TE tunnels on RSVP-enabled interfaces.

https://t.me/learningnets

Verify MPLS TE Tunnels

To display information about MPLS TE tunnels, use the **show mpls traffic-eng tunnels** command in EXEC mode.

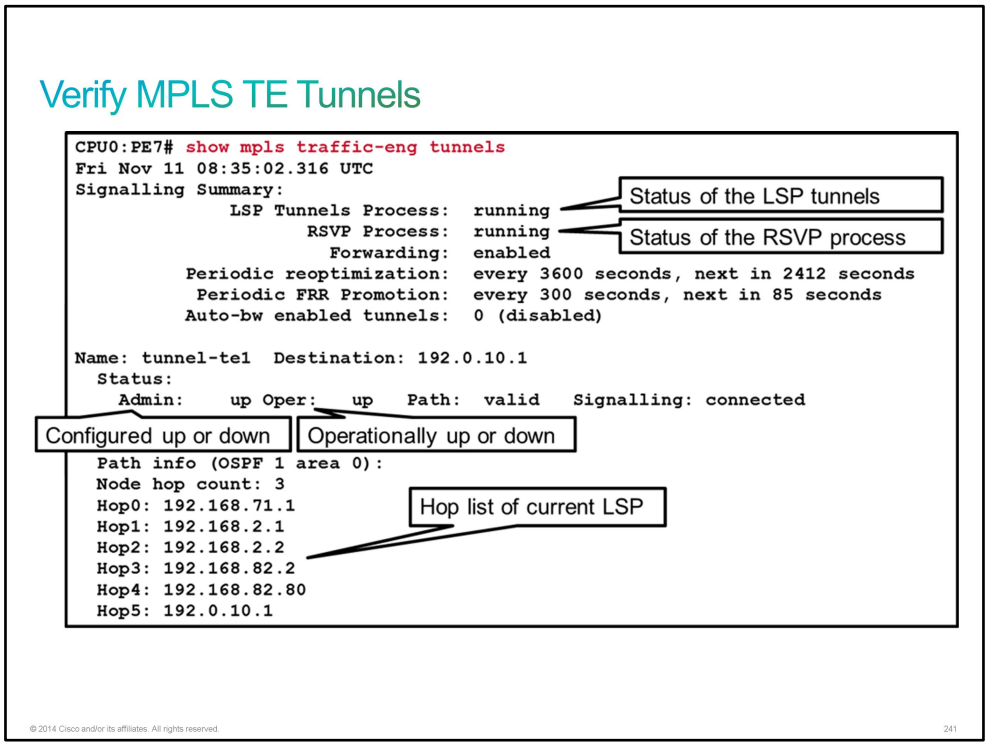

Some output in the figure is omitted. Here is the full output of the command:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1# show mpls traffic-eng tunnels
Fri Nov 11 08:41:16.386 UTC
Signalling Summary:
            LSP Tunnels Process:  running
                   RSVP Process:  running
                     Forwarding:  enabled
         Periodic reoptimization:  every 3600 seconds, next in 2038 seconds
          Periodic FRR Promotion:  every 300 seconds, next in 11 seconds
         Auto-bw enabled tunnels:  0 (disabled)

Name: tunnel-te1  Destination: 192.0.10.1
Status:
  Admin:    up Oper:   up   Path:  valid   Signalling: connected

  path option 1,  type dynamic  (Basis for Setup, path weight 4)
  G-PID: 0x0800 (derived from egress interface properties)
  Bandwidth Requested: 50 kbps  CT0
  Creation Time: Thu Oct 20 16:44:38 2011 (3w0d ago)
Config Parameters:
  Bandwidth:       50 kbps (CT0) Priority:  7  7 Affinity: 0x0/0xffff
  Metric Type: TE (default)
  Hop-limit: disabled
  AutoRoute:  enabled  LockDown: disabled   Policy class: not set
  Forwarding-Adjacency: disabled
  Loadshare:          0 equal loadshares
  Auto-bw: disabled
  Fast Reroute: Disabled, Protection Desired: None
  Path Protection: Not Enabled
History:
  Tunnel has been up for: 20:45:45 (since Thu Nov 10 11:55:31 UTC 2011)
  Current LSP:
    Uptime: 20:01:23 (since Thu Nov 10 12:39:53 UTC 2011)
  Reopt. LSP:
    Last Failure:
      LSP not signalled, identical to the [CURRENT] LSP
      Date/Time: Thu Nov 10 14:15:14 UTC 2011 [18:26:02 ago]
  Prior LSP:
    ID: path option 1 [2]
    Removal Trigger: reoptimization completed

Path info (OSPF 1 area 0):
Node hop count: 3
Hop0: 192.168.71.1
Hop1: 192.168.2.1
Hop2: 192.168.2.2
Hop3: 192.168.82.2
Hop4: 192.168.82.80
Hop5: 10.8.1.1

LSP Tunnel 10.8.1.1 1 [523] is signalled, connection is up
Tunnel Name: PE8_t1 Tunnel Role: Tail
InLabel: GigabitEthernet0/0/0/2, implicit-null
Signalling Info:
  Src 10.8.1.1 Dst 10.7.1.1, Tun ID 1, Tun Inst 523, Ext ID 10.8.1.1
  Router-IDs: upstream   10.0.1.1
              local      10.7.1.1
  Bandwidth: 40 kbps (CT0) Priority:  1  1 DSTE-class: no match
```

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

```
  Path Info:
    Incoming Address: 192.168.71.70
    Incoming:
    Explicit Route:
      Strict, 192.168.71.70
      Strict, 10.7.1.1
    Record Route: Disabled
    Tspec: avg rate=40 kbits, burst=1000 bytes, peak rate=40 kbits
    Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
  Resv Info: None
    Record Route: Disabled
    Fspec: avg rate=40 kbits, burst=1000 bytes, peak rate=40 kbits
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 1 (of 1) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```

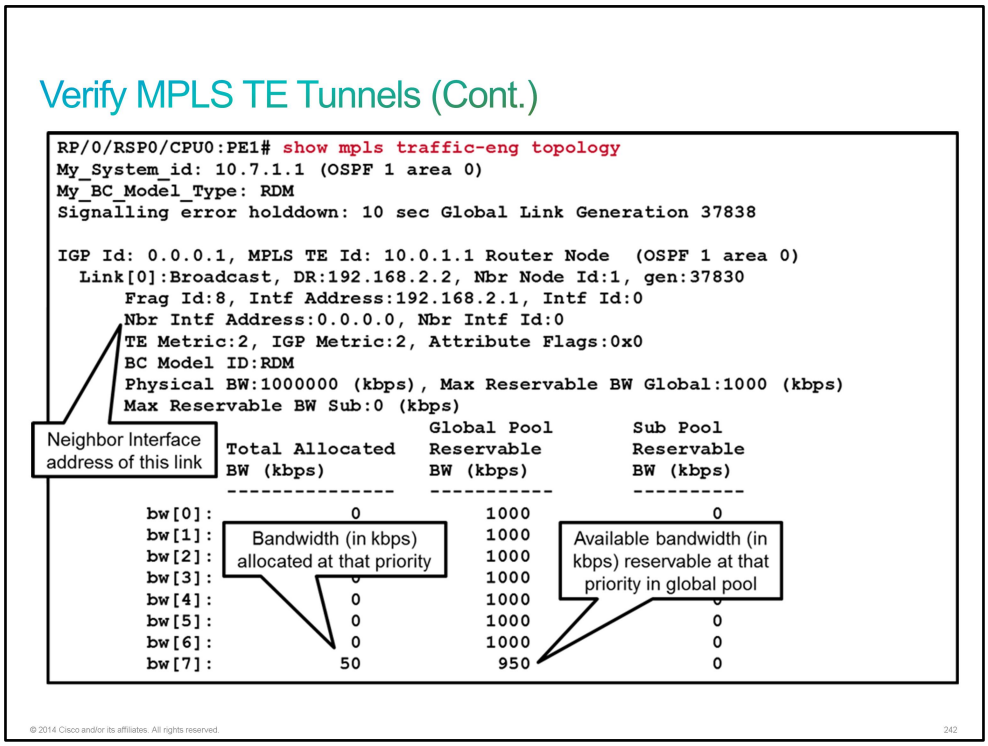

To display the MPLS TE network topology currently known at this node, use the **show mpls traffic-eng topology** command in EXEC mode. Some output in the figure is omitted. Here is the full output of the command:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE7# show mpls traffic-eng topology
Fri Nov 11 09:07:05.630 UTC
My_System_id: 10.7.1.1 (OSPF 1 area 0)
My_BC_Model_Type: RDM

Signalling error holddown: 10 sec Global Link Generation 37838

IGP Id: 0.0.0.1, MPLS TE Id: 10.0.1.1 Router Node  (OSPF 1 area 0)

Link[0]:Broadcast, DR:192.168.2.2, Nbr Node Id:1, gen:37830
    Frag Id:8, Intf Address:192.168.2.1, Intf Id:0
    Nbr Intf Address:0.0.0.0, Nbr Intf Id:0
    TE Metric:2, IGP Metric:2, Attribute Flags:0x0
    Attribute Names:
    Switching Capability:None, Encoding:unassigned
    BC Model ID:RDM
    Physical BW:1000000 (kbps), Max Reservable BW Global:1000 (kbps)
    Max Reservable BW Sub:0 (kbps)
                              Global Pool       Sub Pool
             Total Allocated  Reservable        Reservable
             BW (kbps)        BW (kbps)         BW (kbps)
             ---------------  -----------       ----------
      bw[0]:            0          1000               0
      bw[1]:            0          1000               0
      bw[2]:            0          1000               0
      bw[3]:            0          1000               0
      bw[4]:            0          1000               0
      bw[5]:            0          1000               0
      bw[6]:            0          1000               0
      bw[7]:           50           950               0

Link[1]:Broadcast, DR:192.168.71.70, Nbr Node Id:6, gen:37831
    Frag Id:12, Intf Address:192.168.71.1, Intf Id:0
    Nbr Intf Address:0.0.0.0, Nbr Intf Id:0
    TE Metric:1, IGP Metric:1, Attribute Flags:0x0
    Attribute Names:
    Switching Capability:None, Encoding:unassigned
    BC Model ID:RDM
    Physical BW:1000000 (kbps), Max Reservable BW Global:1000 (kbps)
    Max Reservable BW Sub:0 (kbps)
                              Global Pool       Sub Pool
             Total Allocated  Reservable        Reservable
             BW (kbps)        BW (kbps)         BW (kbps)
             ---------------  -----------       ----------
      bw[0]:            0          1000               0
      bw[1]:           40           960               0
      bw[2]:            0           960               0
      bw[3]:            0           960               0
      bw[4]:            0           960               0
      bw[5]:            0           960               0
      bw[6]:            0           960               0
      bw[7]:            0           960               0
```


https://t.me/learningnets

Verify a Routing Table with Autoroute Enabled

One of the MPLS TE verification steps is to make sure that traffic that is destined for behind the tail end of the tunnel is entering the newly created MPLS TE tunnel.

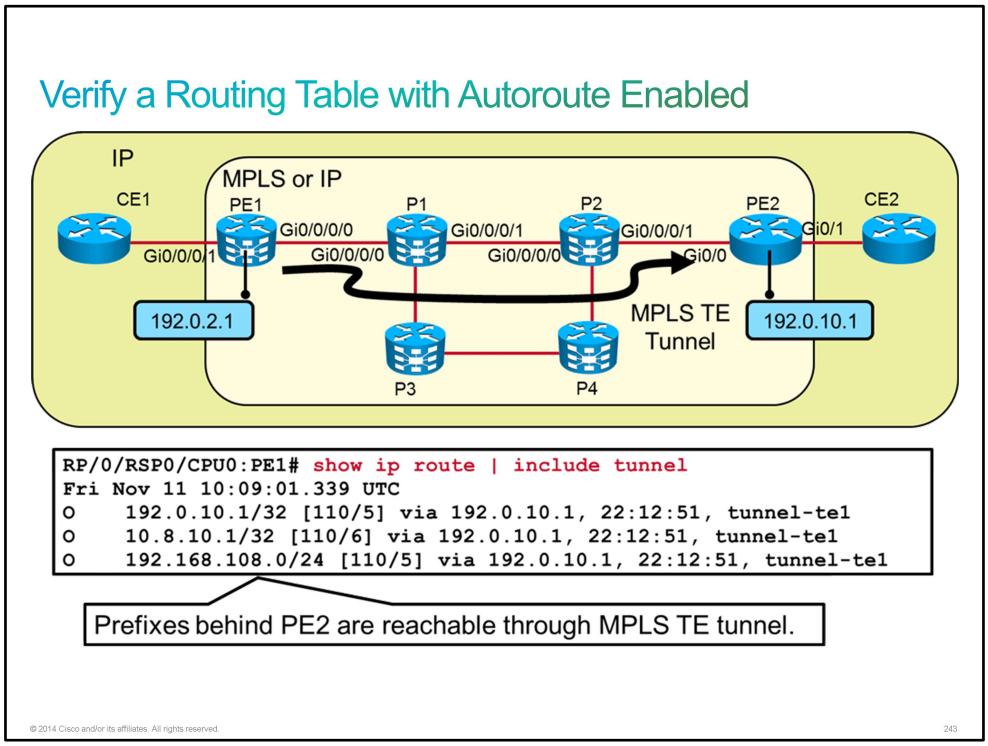

Verify that the prefixes that are behind the tail-end router of the MPLS TE tunnel are reachable through the MPLS TE tunnel interface, as shown in the figure.

https://t.me/learningnets

MPLS TE Case Study

This MPLS TE case study will be used to show the sample configurations when configuring dynamic or explicit MPLS TE tunnels and path-selection restrictions.

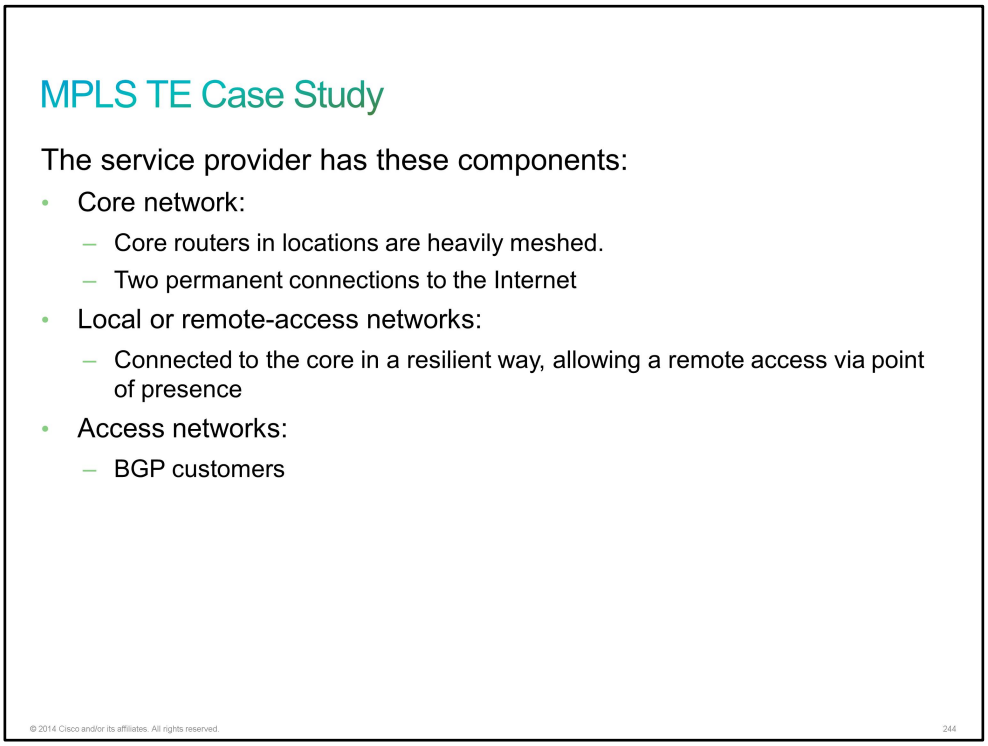

The example in the figure shows a classic ISP architecture based on three levels of hierarchy. The design should bring together some of the aspects of TE and routing design that are discussed in this module:

- Routing protocol choice and interaction between different routing protocols
- Support for MPLS TE tunnels

https://t.me/learningnets

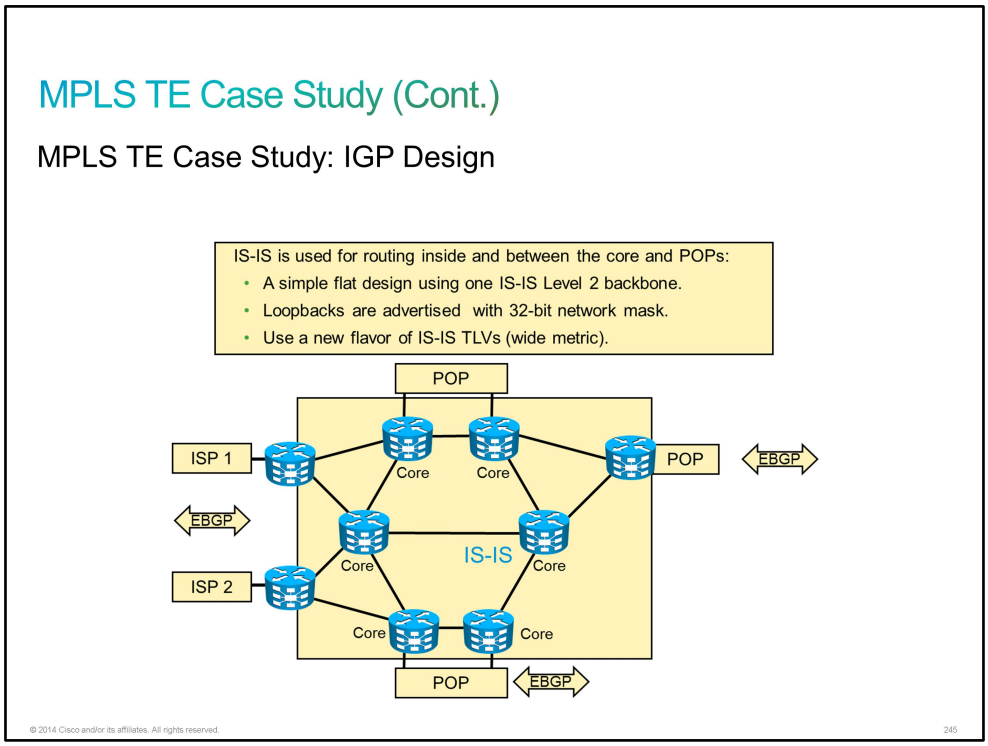

The network is based on a three-level hierarchy:

- **Core network:** Highly meshed central sites with high bandwidth requirements between them
- POP **sites:** A distribution layer of regional sites, that are connected back to the core over redundant links, and 5provide 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.
- A “wide metric” is used, which allows greater granularity of path selection.

https://t.me/learningnets

MPLS TE Case Study (Cont.)

MPLS TE Case Study: BGP Design

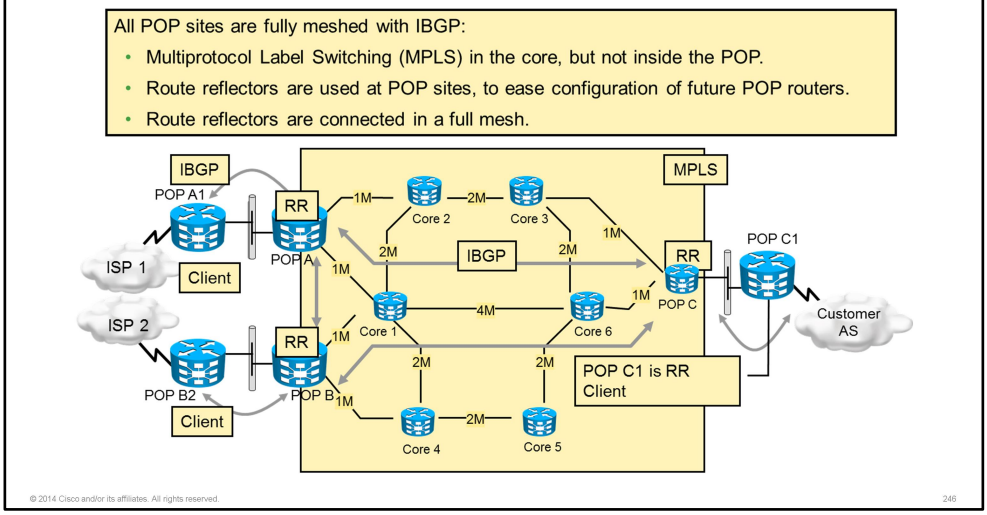

The method of core transport uses MPLS as a transport mechanism in the core network. Packets are switched through the core of the network at Layer 2, bypassing the traditional Layer 3 routing process.

The issue with the edge-only BGP peer design is that the routers 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 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.

Note	The BGP configuration will not be shown in this case study.

https://t.me/learningnets

MPLS TE Case Study: Dynamic MPLS TE Tunnel

The configuration for the dynamic MPLS TE tunnel based on the case study will be shown.

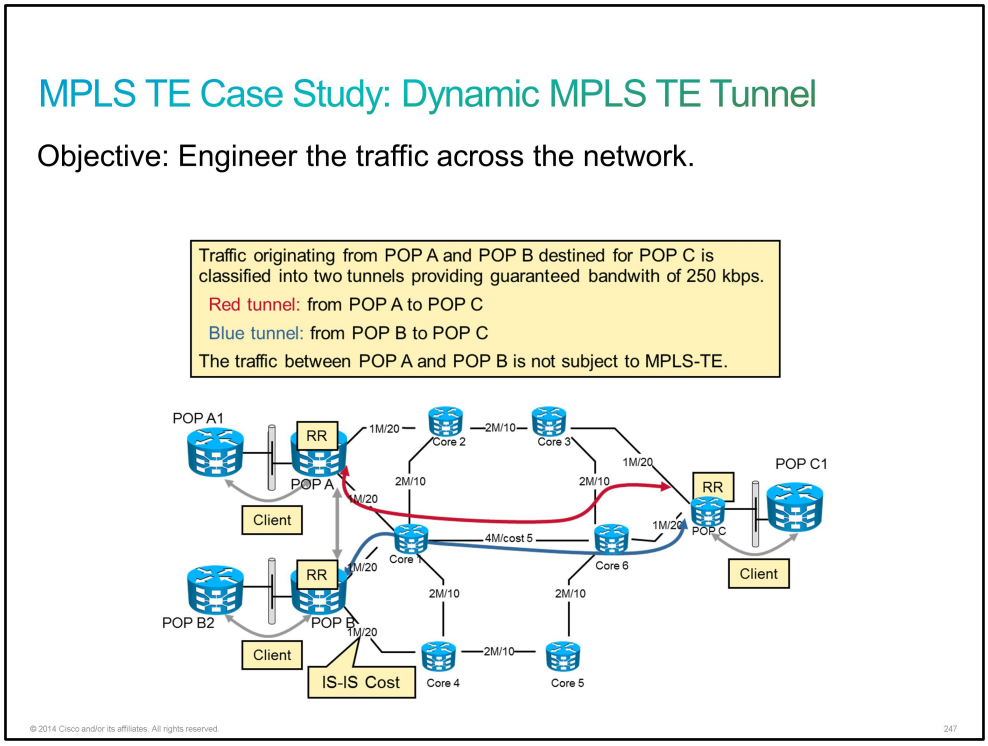

This sample configuration demonstrates how to implement TE with an existing MPLS network.

The figure shows the implementation of two TE tunnels of 250 kbps. Either the tunnels are being automatically set up by the ingress LSRs, (POP A, POP B, and POP C routers), or they are 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.

https://t.me/learningnets

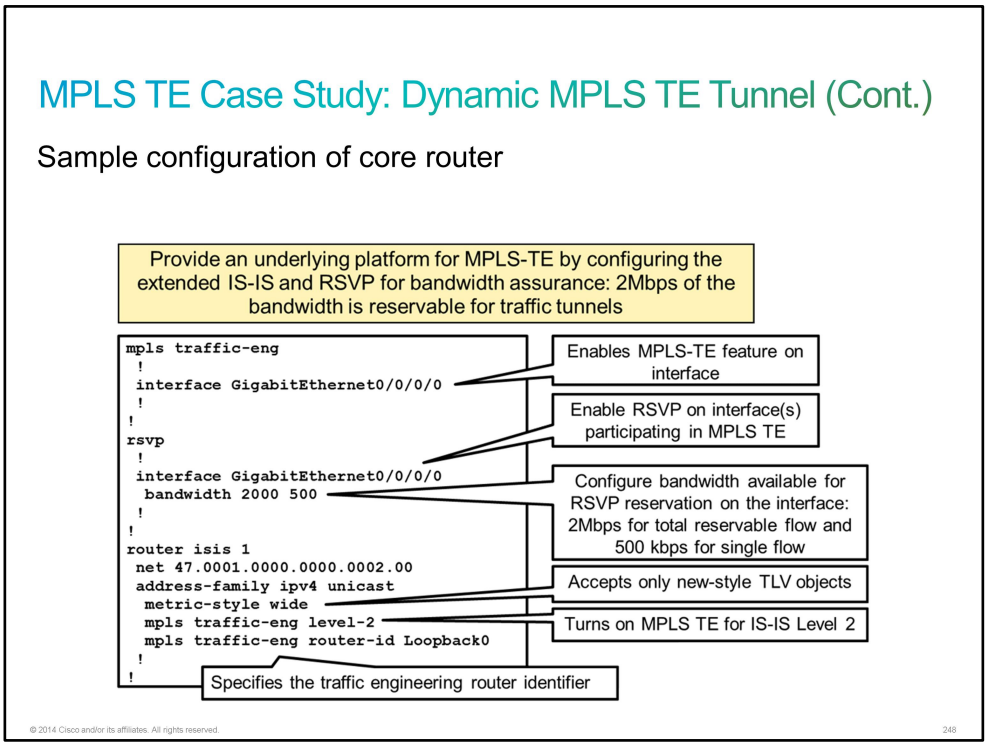

MPLS TE uses an extension to existing protocols such as RSVP, IS-IS, and OSPF to calculate and establish unidirectional tunnels that are set according to the network constraint. Traffic flows are mapped on the different tunnels depending on their destination.

With Cisco, MPLS TE is built on these 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 (CBR), which determines the paths to use for LSP tunnels.
- Unidirectional LSP tunnels, that are signaled through RSVP.
- An MPLS TE link management module that manages link admission and the 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.

```
#Sample Cisco IOS Configuration of Core Router
ip cef
mpls ip
mpls traffic-eng tunnels
interface GigabitEthernet 0/0
mpls traffic-eng tunnels
  ip rsvp bandwidth 1500 500
  ip router isis
router isis
  passive-interface Loopback0
  net 49.0001.0000.0001.0001.00
  is-type level-2-only
  metric-style wide
  mpls traffic-eng router-id Loopback0
  mpls traffic-eng level-2
```


https://t.me/learningnets

MPLS TE Case Study: Dynamic MPLS TE Tunnel (Cont.)

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.

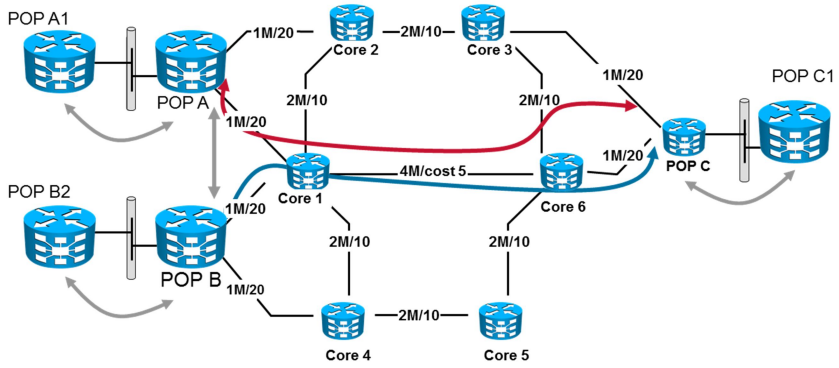

© 2014 Cisco and/or its affiliates. All rights reserved.

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 these two considerations:

- The best route is the least-cost route with enough resources. CBR uses its own metric (administrative weight, or TE cost), that is by default equal to the IGP.

  If there is 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, a path is selected randomly.
- The result of CBR is an explicit route, that is used by RSVP to reserve resources and establish an LSP path.

https://t.me/learningnets

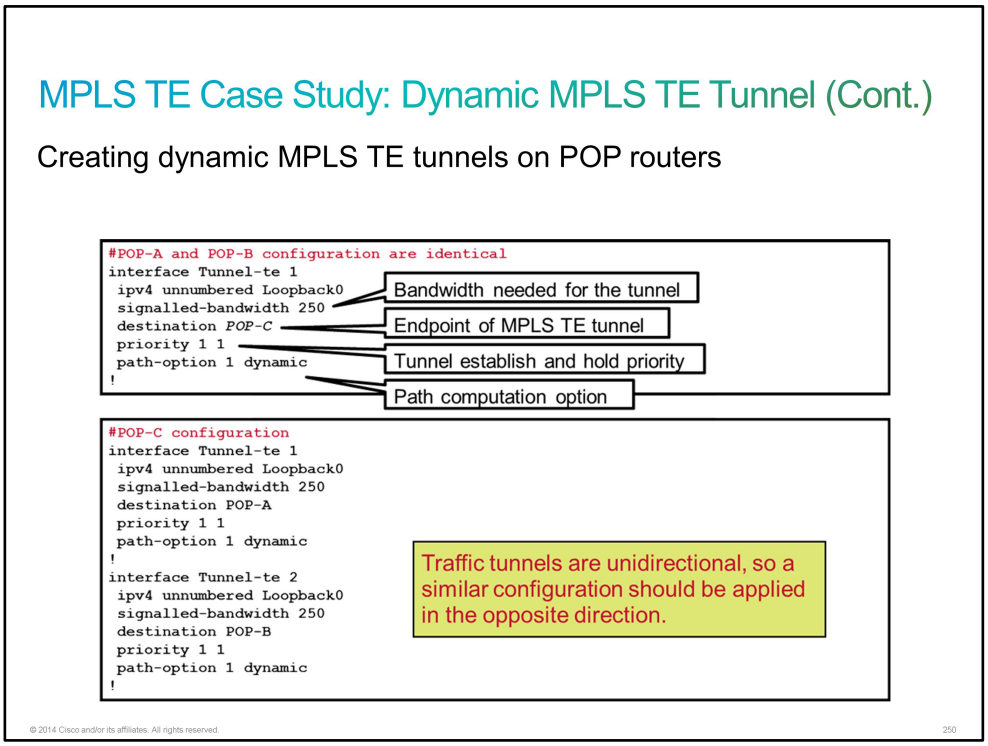

To set up a dynamic TE tunnel (assuming that the IGP platform has been prepared), follow these command steps in the tunnel interface configuration mode:

- **ip unnumbered loopback0**: Gives the tunnel interface an IP address. An MPLS TE tunnel interface should be unnumbered because it represents a unidirectional link.
- **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.
- **priority**:Configures the setup and reservation priority for an MPLS TE tunnel
- **signaled-bandwidth**: Configures the bandwidth for the MPLS TE tunnel
- **path-option** *number* **dynamic**: The LSP path is dynamically calculated.

https://t.me/learningnets

```
#POP-A and POP-B Cisco IOS Configurations are the same
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

#POP-C Cisco IOS 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
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
```


https://t.me/learningnets

MPLS TE Case Study: Explicit MPLS TE Tunnel

The configuration for an explicit MPLS TE tunnel based on the case study will be shown.

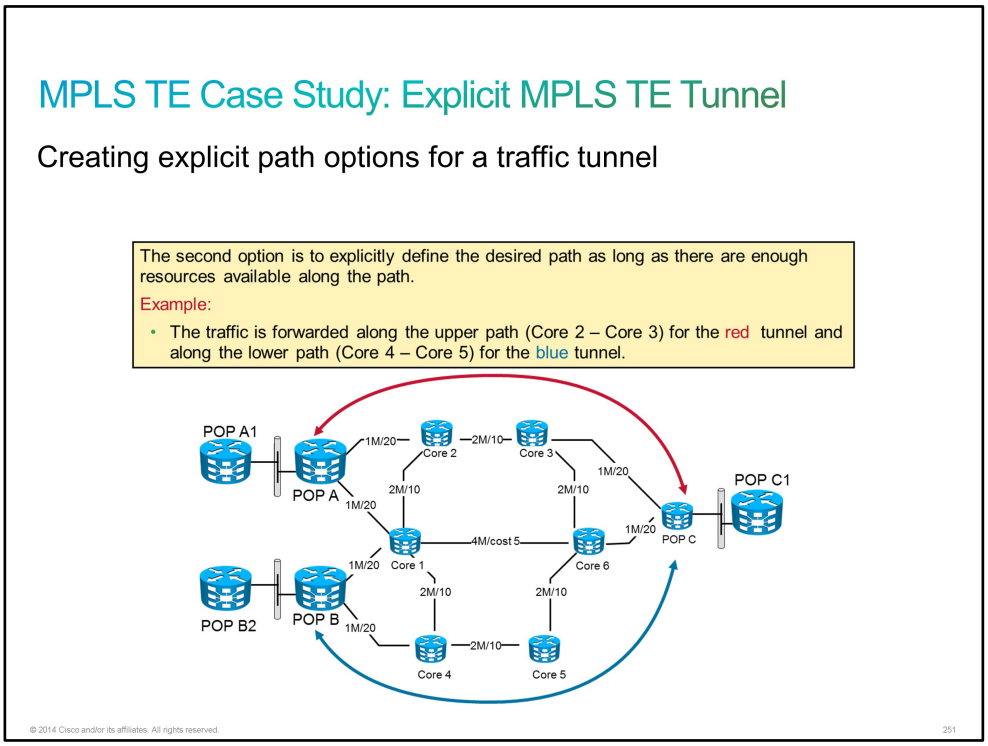

The example in the figure shows how to avoid the first step in the CBR algorithm by manually setting the explicit path between tunnel endpoints. This path might be derived from the least-cost path.

The best route might not be the least-cost route, given 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.

https://t.me/learningnets

MPLS TE Case Study: Explicit MPLS TE Tunnel (Cont.)

Creating explicit MPLS TE tunnel on the POP router:

```
#POP-A configuration
interface Tunnel-te 1
 ipv4 unnumbered Loopback0
 signalled-bandwidth 250
 destination POP-C
 priority 1 1
 path-option 1 explicit name Core2-3
!
explicit-path name Core2-3
 index 1 next-address ipv4 unicast Core-2
 index 2 next-address ipv4 unicast Core-3
 index 3 next-address ipv4 unicast POP-C
```

Explicit path represents the explicit route object, encoded as incoming IP addresses of the routers along the path.

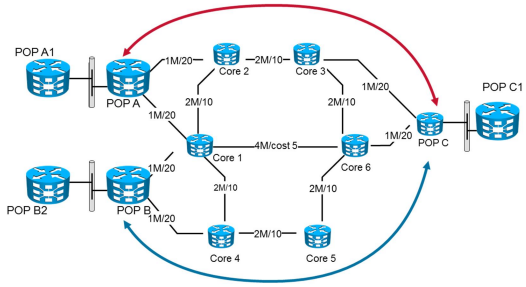

© 2014 Cisco and/or its affiliates. All rights reserved.

To set up a static TE tunnel (assuming that the IGP platform has been prepared), use these additional steps:

- In tunnel interface configuration mode, the following may be used:
  - **path-option** *number* **explicit**: This command enables you to configure the tunnel to use a named IP explicit path from the TE topology database.
- In global configuration mode, the following may be used:
  - **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.
  - To include a path entry at a specific index, use the **index next-address** command in explicit path-configuration mode. To return to the default behavior, use the **no** form of this command: **index** *index-id* {**next-address ipv4 unicast** *A.B.C.D.*}

```
#POP-A Cisco IOS 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
```


https://t.me/learningnets

MPLS TE Case Study: Periodic Tunnel 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 had been set up dynamically, the headend router tries to find an alternative solution. This process is referred to as rerouting.

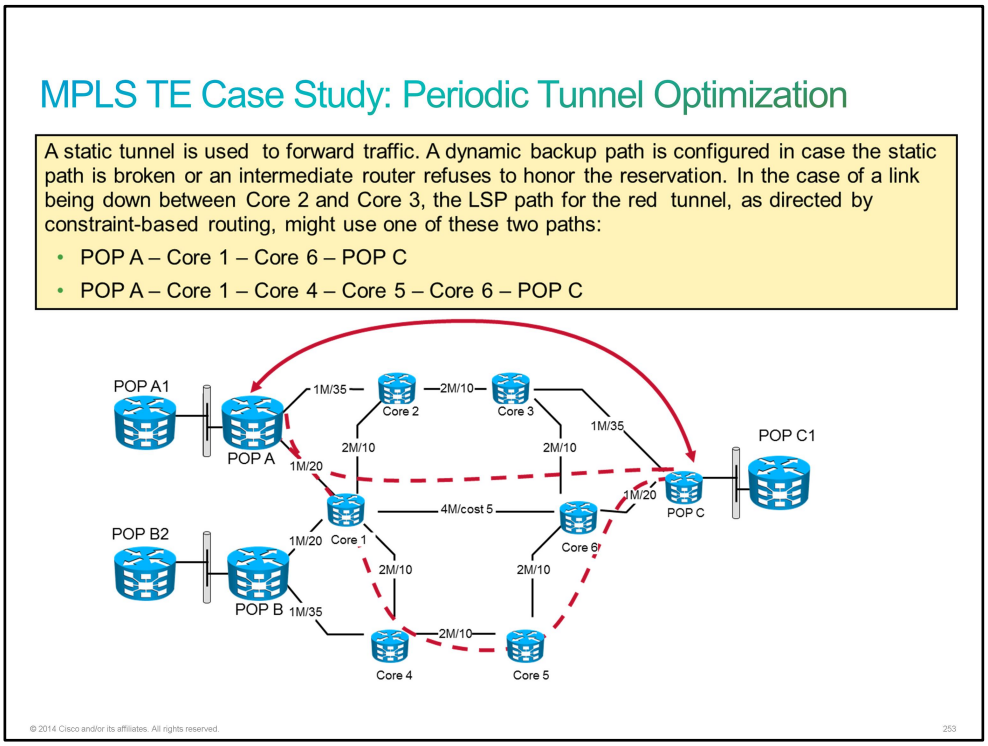

Reoptimization occurs when a device examines tunnels with established LSPs to see if better LSPs are available. If a better LSP seems to be available, the device attempts to signal the better LSP and, if successful, replaces the old and inferior LSP with the new and better LSP.

This reoptimization might be triggered manually, or it might occur at configurable intervals (the default is one 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 instance, the path is explicitly determined, compelling the headend router to strictly follow the explicit path.

https://t.me/learningnets

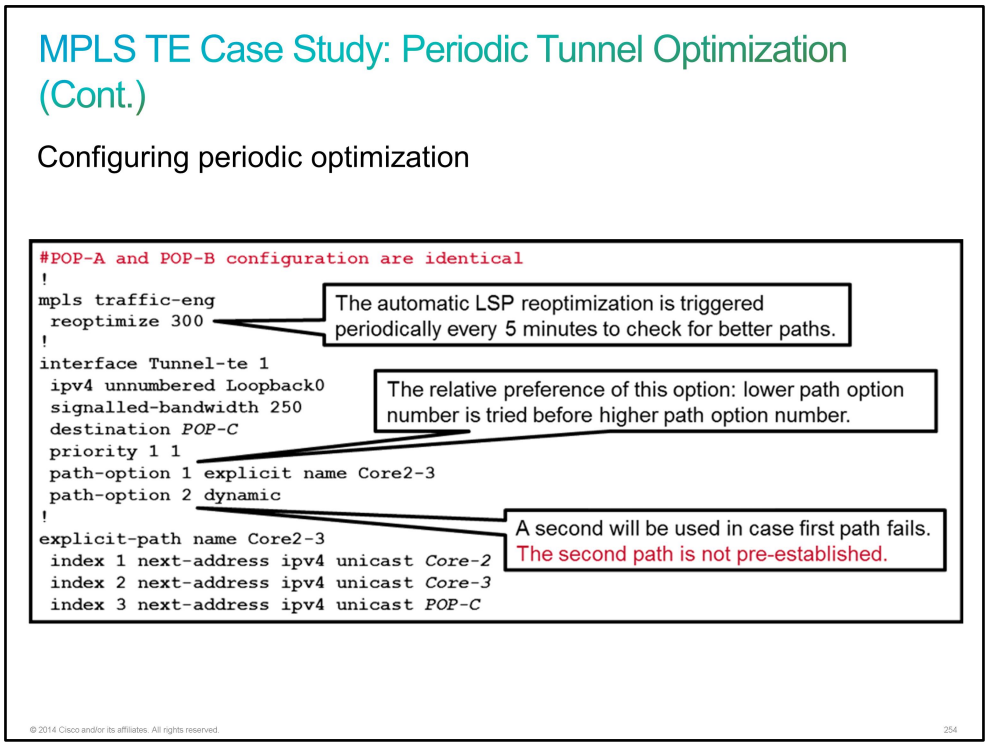

The example in the figure shows how traffic can be engineered across a path in the network and how a backup route for that traffic-engineered path can be established.

The primary path is manually specified, so it is explicit. If this path suddenly cannot be followed, the MPLS TE engine uses the next path option, which in this example is a dynamic route.

The drawback to 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 periodically triggered, there is still downtime while the alternate path is being built.

```
#POP-A nad POP-B Cisco IOS Configuration
mpls traffic-eng reoptimize timers frequency 300
interface Tunnel1
  ip unnumbered Loopback0
  tunnel source 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
  tunnel mpls traffic-eng path-option 2 dynamic
!
ip explicit-path name Core2-3 enable
  next-address strict Core-2
  next-address strict Core-3
  next-address strict POP-C
```


https://t.me/learningnets

MPLS TE Case Study: Configuring Path-Selection Restrictions

In many cases, some links will need to be excluded from the constraint-based SPF computation. This exclusion can be implemented by 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).

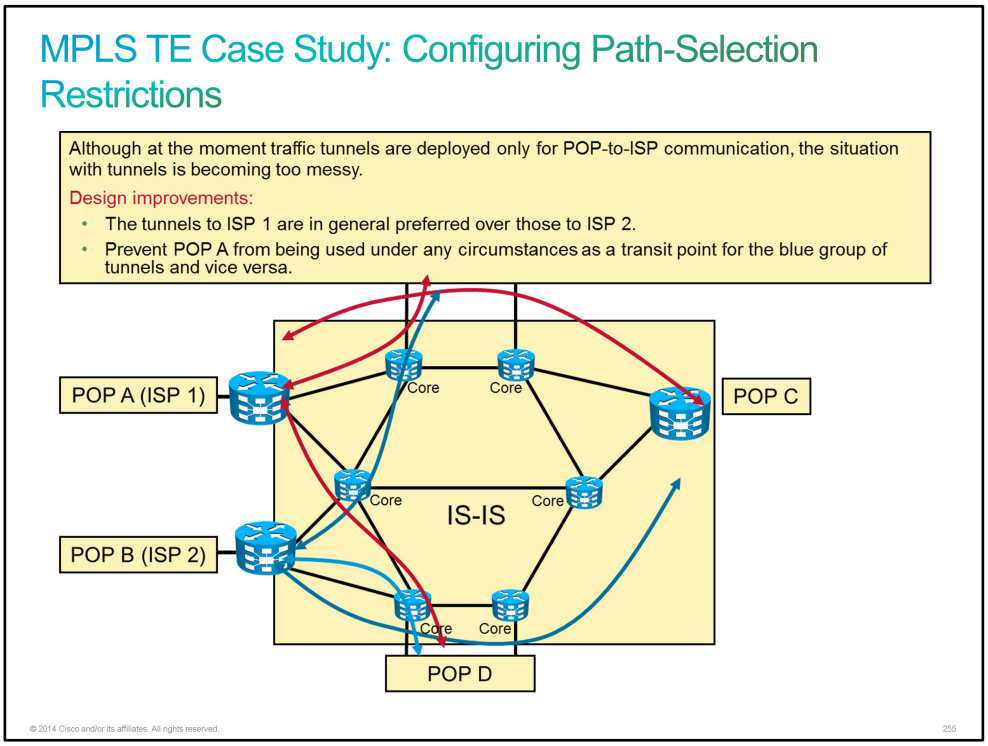

A 32-bit resource class affinity string that is accompanied by a corresponding resource class mask characterizes each traffic tunnel. The “0” bits in the mask exclude the respective link resource class bits from being checked.

Each link is characterized by its resource class 32-bit string, that 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.

https://t.me/learningnets

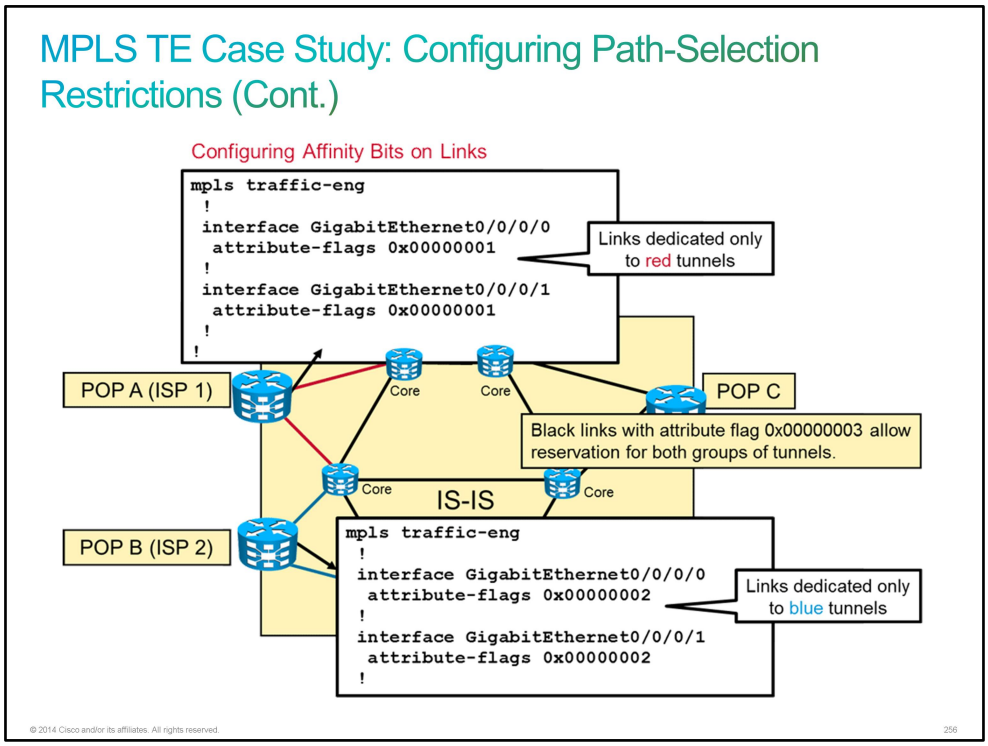

The figure shows a sample network with the tunnel resource class affinity bits and link resource bits. The main goal is to force the CBR algorithm to use only links that are explicitly dedicated to certain tunnels for its path computation.

Because moving all blue tunnels away from POP A interfaces and red tunnels from POP B interfaces is desirable, 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 they compute paths for TE tunnels.

```
#POP-A Cisco IOS Configuration
interface GigabitEthernet 0/0
 mpls traffic-eng attribute-flags 0x00000001
interface GigabitEthernet 0/1
 mpls traffic-eng attribute-flags 0x00000001

#POP-B Cisco IOS Configuration
interface GigabitEthernet 0/0
 mpls traffic-eng attribute-flags 0x00000002
interface GigabitEthernet 0/1
 mpls traffic-eng attribute-flags 0x00000002
```


https://t.me/learningnets

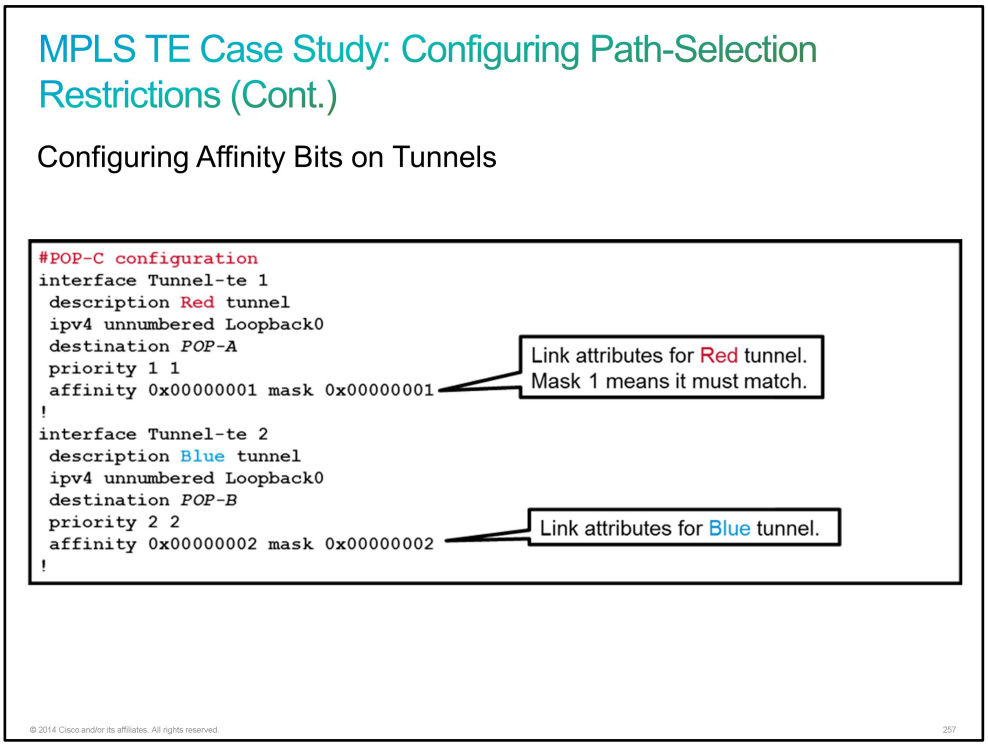

In the example of the first tunnel (Red tunnel), the tunnel requires a match only on the last bit, whereas the second tunnel (Blue tunnel) checks the setting of the next-to-last bit. With tunnel resource class affinity bits and link resource class bits set, the constraint-based path computation considers only the paths where the match is found.

For 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.

```
#POP-C Cisco IOS 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
  ......
```


https://t.me/learningnets

MPLS TE Case Study: Configuring Path-Selection Restrictions (Cont.)

Excluding links or nodes for traffic tunnels

```
#POP-C configuration
interface Tunnel-te 1
 description Red tunnel
 ipv4 unnumbered Loopback0
 destination POP-A
 priority 1 1
 path-option 1 explicit name ex1
!
explicit-path name ex1
 index 1 exclude-address ipv4 unicast POP-B
!
interface Tunnel-te 2
 description Blue tunnel
 ipv4 unnumbered Loopback0
 destination POP-B
 priority 2 2
 path-option 1 explicit name ex2
!
explicit-path name ex2
 index 1 exclude-address ipv4 unicast POP-A
!
```

Address is excluded from the path.

© 2014 Cisco and/or its affiliates. All rights reserved.

Link or node exclusion is accessible using the **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 excluded 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 excluded 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.

Addresses are not excluded from an IP explicit path unless they are explicitly excluded by the **exclude-address** command.

Note	MPLS TE will accept an IP explicit path that is composed of either all excluded addresses configured by the **exclude-address** command, or all included addresses configured by the **next-address** command, but not a combination of both.

In a previous example, affinity bits were used to restrict the possible paths that could be used for tunnel creation. This example shows the use 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 the example are the router IDs of the specified nodes. Typically, those addresses conform to the loopback address of the node.

https://t.me/learningnets

MPLS TE Case Study: Modifying the Administrative Weight

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.

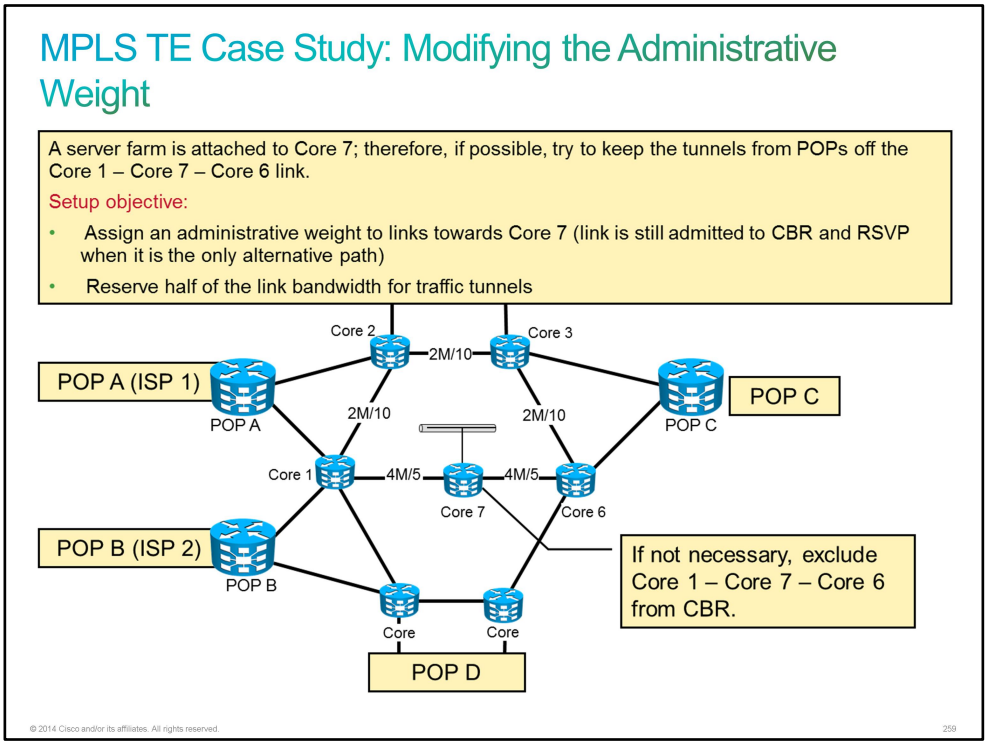

This administrative weight is, by default, equal to the IGP link metric (cost). Increase the TE cost on the link if excluding a certain link from any path computation is desirable, while keeping the link available if that link represents the only available path.

https://t.me/learningnets

MPLS TE Case Study: Modifying the Administrative Weight (Cont.)

Configuring administrative weight

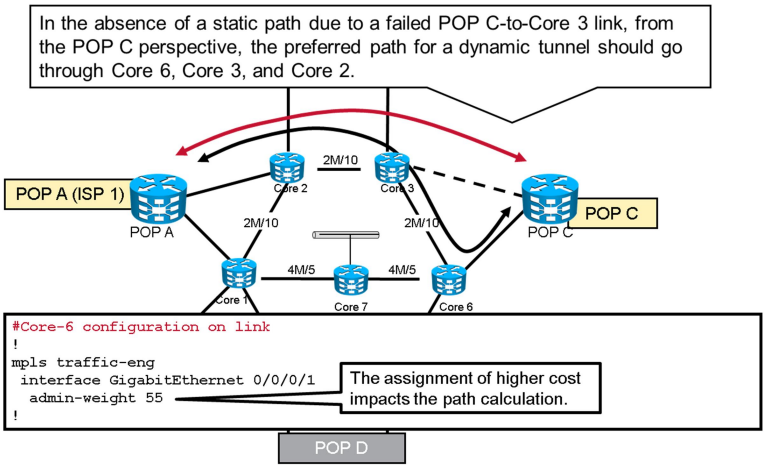

© 2014 Cisco and/or its affiliates. All rights reserved.

In the example in the figure, the TE cost of the link between Core 1 and Core 7 is increased to 55, that makes links providing alternative paths more economical and more attractive for backup purposes.

```
int GigabitEthernet0/0
mpls traffic-eng administrative-weight 55
```


https://t.me/learningnets

MPLS TE Case Study: Autoroute and Forwarding Adjacency

The autoroute feature enables the headend routers to see the MPLS TE tunnel as a directly connected interface and to 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.

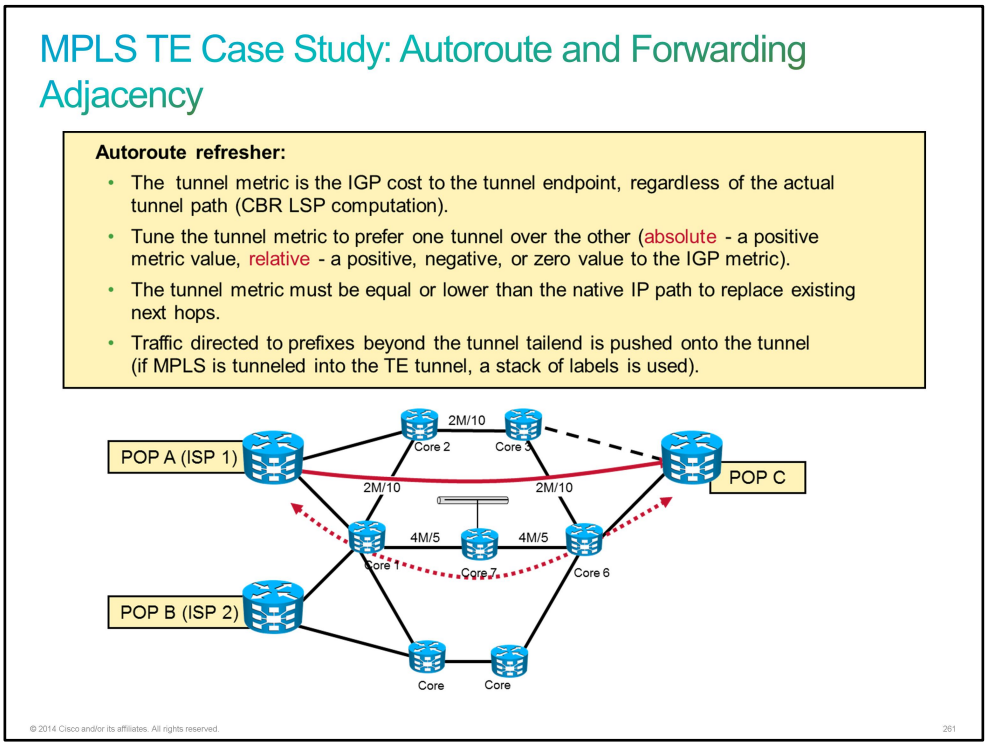

The autoroute feature enables the headend routers to see the MPLS TE tunnel as a directly connected interface and to use it in their modified SPF computations. The MPLS TE tunnel is used only for normal IGP route calculation (at the headend only) and is not included in any constraint-based path computation.

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 (tail end) 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 remainder of the network does not have any knowledge of it.

https://t.me/learningnets

MPLS TE Case Study: Autoroute and Forwarding Adjacency (Cont.)

Changing autoroute metrics

```
#POP-C configuration
interface Tunnel-te 1
 ipv4 unnumbered Loopback0
 signalled-bandwidth 250
 destination POP-A
 priority 1 1
 path-option 1 explicit name Core3-2
 autoroute announce
 autoroute metric absolute 1
!
explicit-path name Core3-2
 next-address Core-3
 next-address Core-2
 next-address POP-A
!
interface Tunnel-te 2
 ipv4 unnumbered Loopback0
 signalled-bandwidth 125
 destination POP-A
 priority 1 1
 path-option 1 dynamic
 autoroute announce
 autoroute metric absolute 2
!
```

Announce the presence of the tunnel to the routing protocol and select the primary path.

The metric determines the secondary path.

© 2014 Cisco and/or its affiliates. All rights reserved.

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, as 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 they provide load balancing. The load balancing is done proportionally to the configured bandwidth of the tunnel.

https://t.me/learningnets

```
#POP-C Cisco IOS 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
```

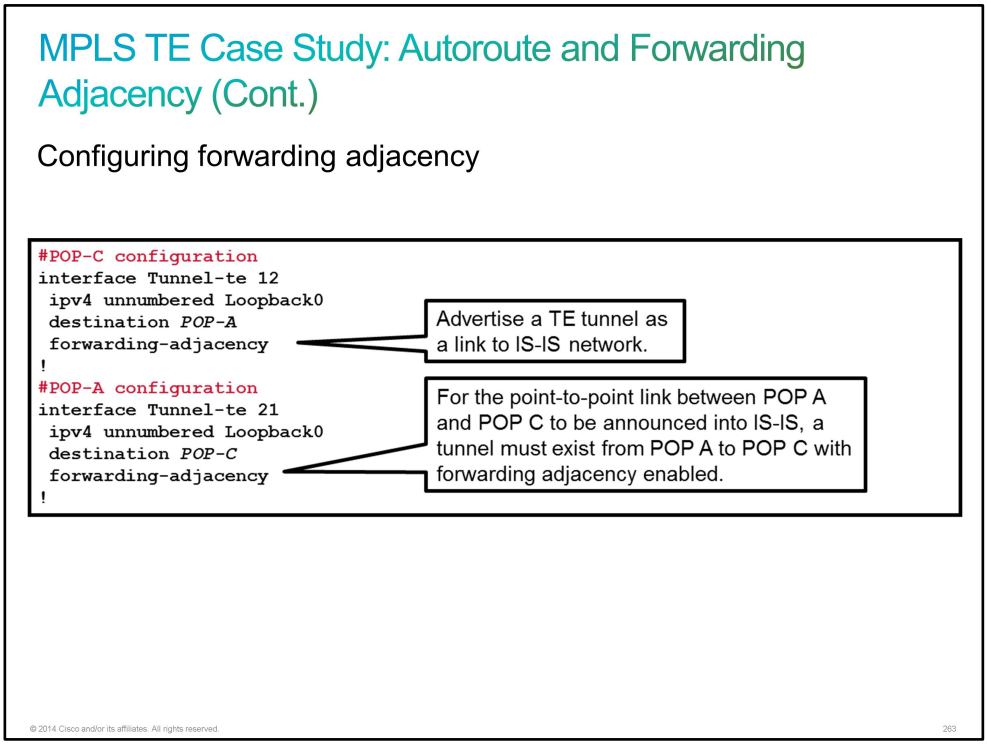

To advertise a TE tunnel as a link in an IGP network, use the forwarding adjacency feature. You must configure a forwarding adjacency on two LSP tunnels bidirectionally, from A to B and from B to A. Otherwise, the forwarding adjacency is advertised but not used in the IGP network.

https://t.me/learningnets

```
#POP-C Cisco IOS Configuration
interface Tunnel12
  ip unnumbered Loopback0
  tunnel destination POP-A
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng forwarding-adjacency
 ...
#POP-C Cisco IOS Configuration
interface Tunnel21
  ip unnumbered Loopback0
  tunnel destination POP-C
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng forwarding-adjacency
  ...
```


https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Several steps are required to configure MPLS TE tunnels.
- MPLS TE functionality should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.
- RSVP support should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.
- MPLS TE support for OSPF should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.
- MPLS TE support for IS-IS should be enabled on all routers on the path from the headend to the tail end of the MPLS TE tunnel.
- The MPLS TE tunnel is unidirectional.
- The autoroute feature automatically routes traffic to prefixes behind the MPLS TE tunnel.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- To display information about MPLS TE tunnels, use the **show mpls traffic-eng tunnels** command.
- Verify that the prefixes that are behind the tail-end router of the MPLS TE tunnel are reachable through the MPLS TE tunnel interface.
- The best route might be any sequence of next-hop routers that are configured at the headend of the tunnel.
- Reoptimization occurs when a device examines tunnels with established LSPs to see if better LSPs are available.
- In many cases, some links will need to be excluded from the constraint-based SPF computation.
- This administrative weight is, by default, equal to the IGP link metric (cost).
- The autoroute feature enables the headend routers to see the MPLS TE tunnel as a directly connected interface and to use it in their modified SPF computations.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Lesson 4: Protecting MPLS TE Traffic

Overview

This lesson describes the advanced MPLS TE commands for link protection and the advanced MPLS TE commands for bandwidth control. The configuration commands are accompanied by usage guidelines and examples.

Upon completing this lesson, you will be able to:

- Describe how to use a backup TE tunnel to improve convergence time
- Explain backup MPLS TE tunnel configurations
- Discuss the drawbacks of using a backup MPLS TE tunnels
- Describe the use of Fast Reroute in a case study
- Explain the use of link protection with Fast Reroute using a case study
- Explain the use of node protection with Fast Reroute using a case study
- Explain the Fast Reroute link protection configurations using a case study
- Describe MPLS TE autotunnel backup
- Explain how to configure the MPLS TE autotunnel backup feature
- Explain how to verify the MPLS TE autotunnel backup feature
- Describe MPLS TE autotunnel mesh
- Explain how to configure the MPLS TE autotunnel mesh feature
- Describe MPLS TE end-to-end path protection
- Explain how to configure MPLS TE path protection feature
- Describe BFD over MPLS TE LSPs

https://t.me/learningnets

- Explain the automatic bandwidth adjustment feature and configuration
- Describes how to configure automatic bandwidth feature
- Describe the basic DiffServ-Aware MPLS TE tunnels concept
- Describe the DiffServ-Aware MPLS TE dual-bandwidth pools
- Describe the basic DiffServ-Aware MPLS TE tunnels configuration

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Improving MPLS TE Convergence Time

At first, the need for fast convergence is not obvious. Historically, applications were designed to recover from network outages. However, with the current increased usage of voice and video applications, network downtime must be kept at a minimum. The amount of time that it previously took for a network to converge after a link failure is now unacceptable.

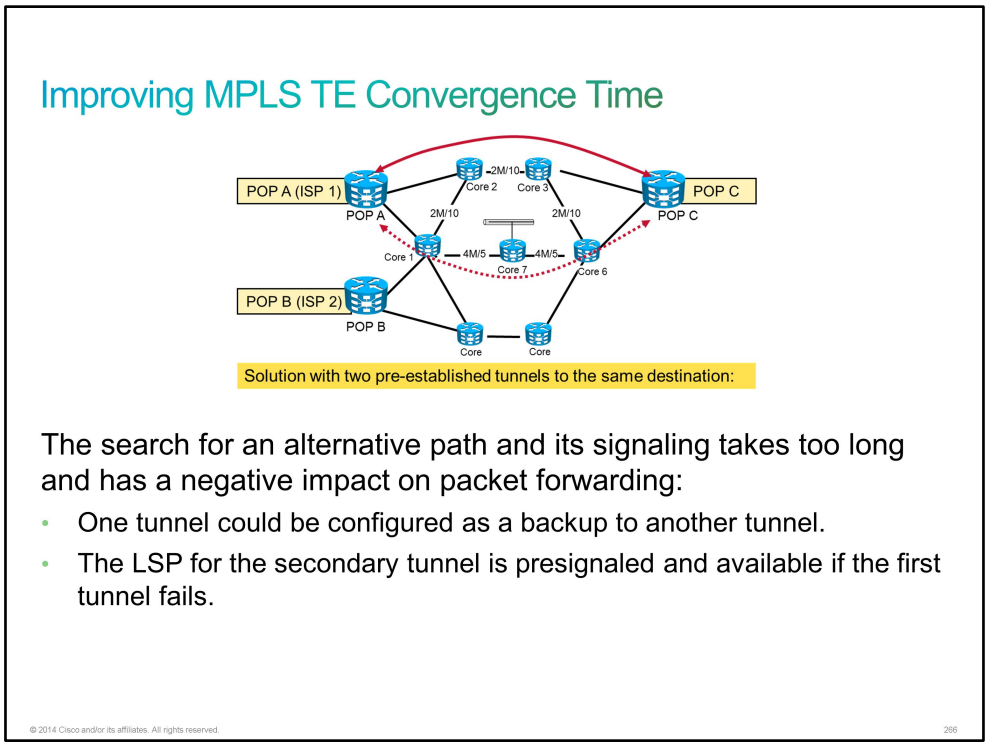

For example, a flapping link can result in headend routers being constantly involved in constraint-based computations. Because the time that elapses between link failure detection and the establishment of a new LSP can cause delays for critical traffic, there is a need for preestablished alternative paths (backups).

Here, two tunnels are used between the same endpoints at the same time. The main requirement in this scenario is that preconfigured tunnels between the same endpoints must use diverse paths. 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 preestablished paths is the simplest form of MPLS TE path protection. Several steps must be taken in preparation for effective switching between the tunnels. These steps include routing to the proper tunnel.

https://t.me/learningnets

Configuring Backup MPLS TE Tunnels

The example shows two configured tunnels: Tunnel-te 1 (following the LSP path Core 3-Core 2-POP A) and Tunnel-te 2 (using a dynamic path).

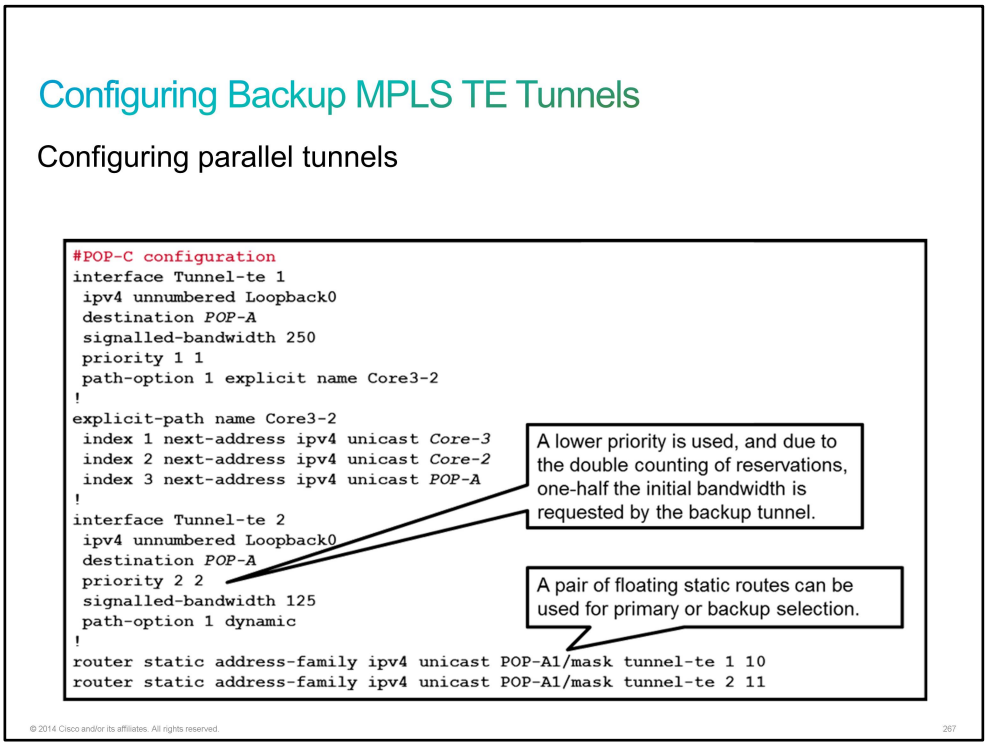

In the presence of two tunnels, static routing is deployed with two floating static routes pointing to the tunnels.

As soon as the primary tunnel (Tunnel-te 1) fails, the static route is gone and the traffic is transitioned to the secondary tunnel. The traffic is returned to the primary tunnel if the conditions support the re-establishment of traffic.

Other options include spreading the load proportionally to the requested bandwidth using the Cisco Express Forwarding mechanism, load balancing, or by having one group of static routes pointing to Tunnel-te 1 and another to Tunnel-te 2.

https://t.me/learningnets

```
#POP-C Cisco IOS 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
ip route POP-A1 mask tunnel 1 10
ip route POP-A1 mask tunnel 2 11
```

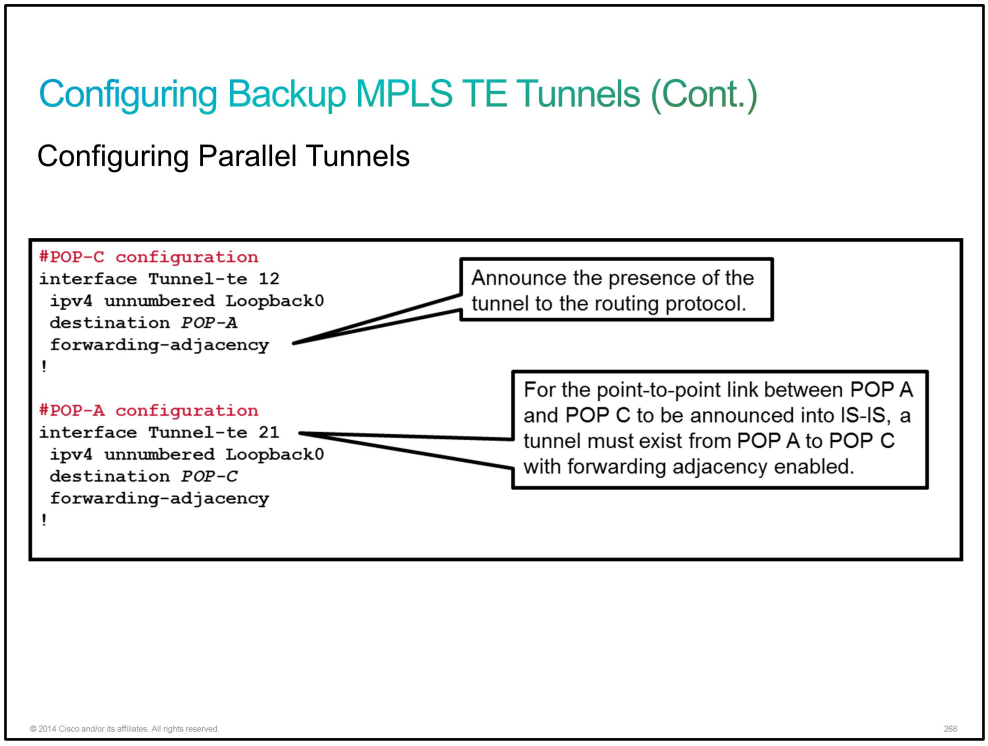

Any of the available methods can be used to advertise the tunnel. In this example, forwarding adjacency is used to advertise the TE tunnel as a link in an IGP network.

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.

https://t.me/learningnets

Drawbacks of Backup MPLS TE 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 preestablished paths (backups). Therefore, two tunnels are used between the same endpoints at the same time.

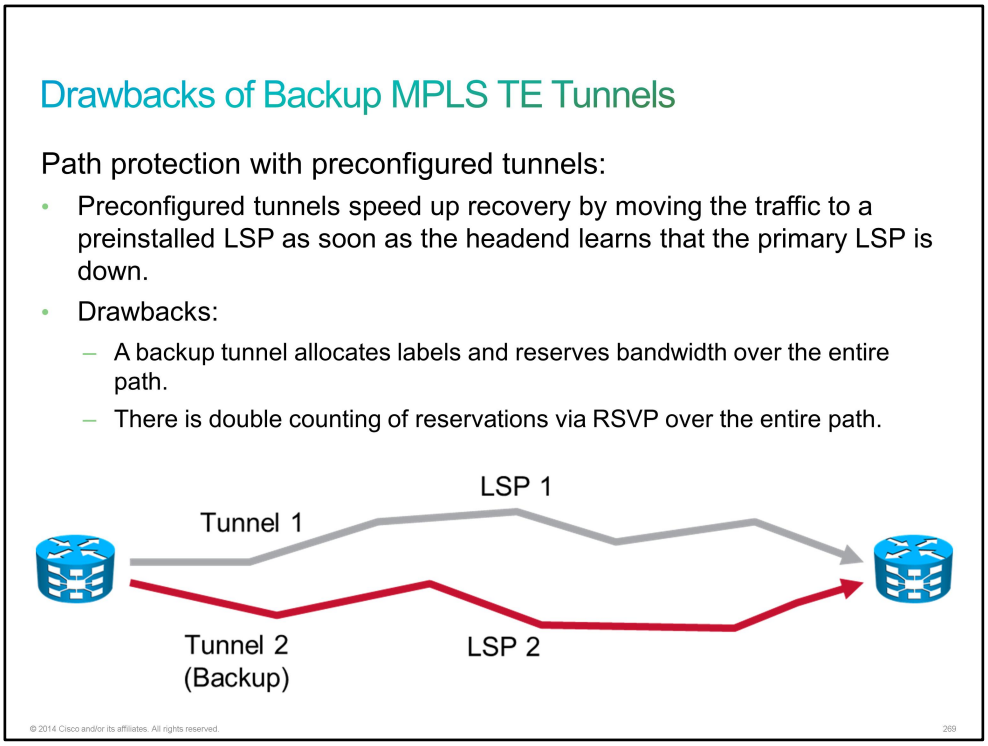

Preconfigured tunnels between the same endpoints must use diverse paths.

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 preestablished 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.

https://t.me/learningnets

Fast Reroute Case Study

In this case study, a company has decided to use only dynamic tunnels.

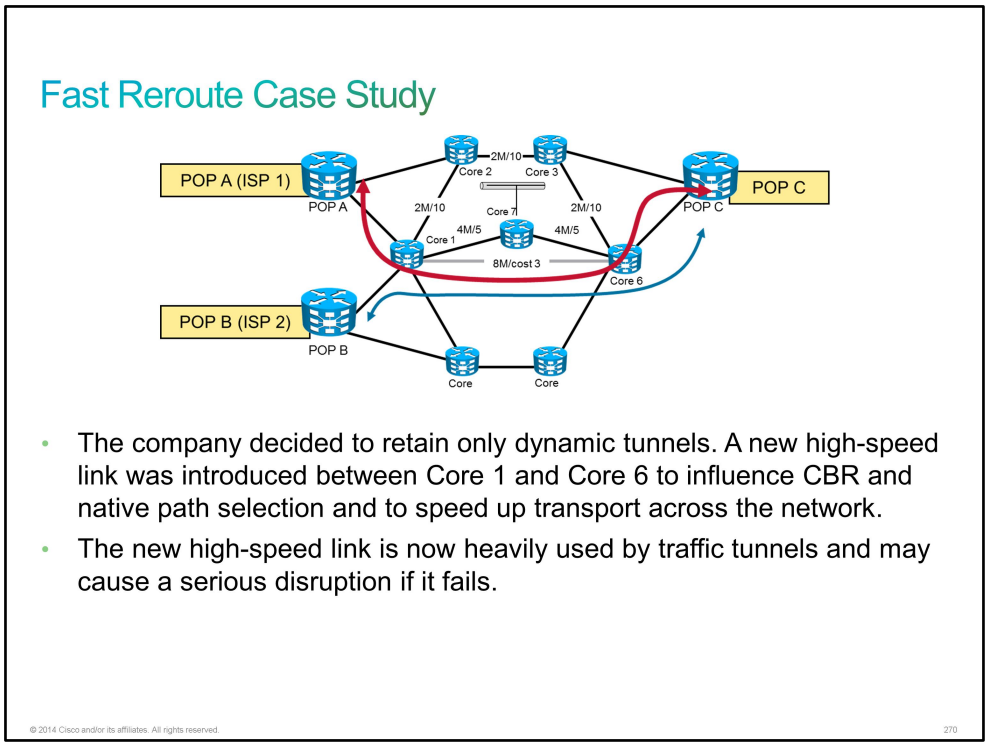

A new high-speed link has been introduced between Core 1 and Core 6. This link influences CBR and native path selection, and speeds up transport across the network. The result, however, is that this new high-speed link is now heavily used by traffic tunnels and may cause a serious disruption if it fails.

https://t.me/learningnets

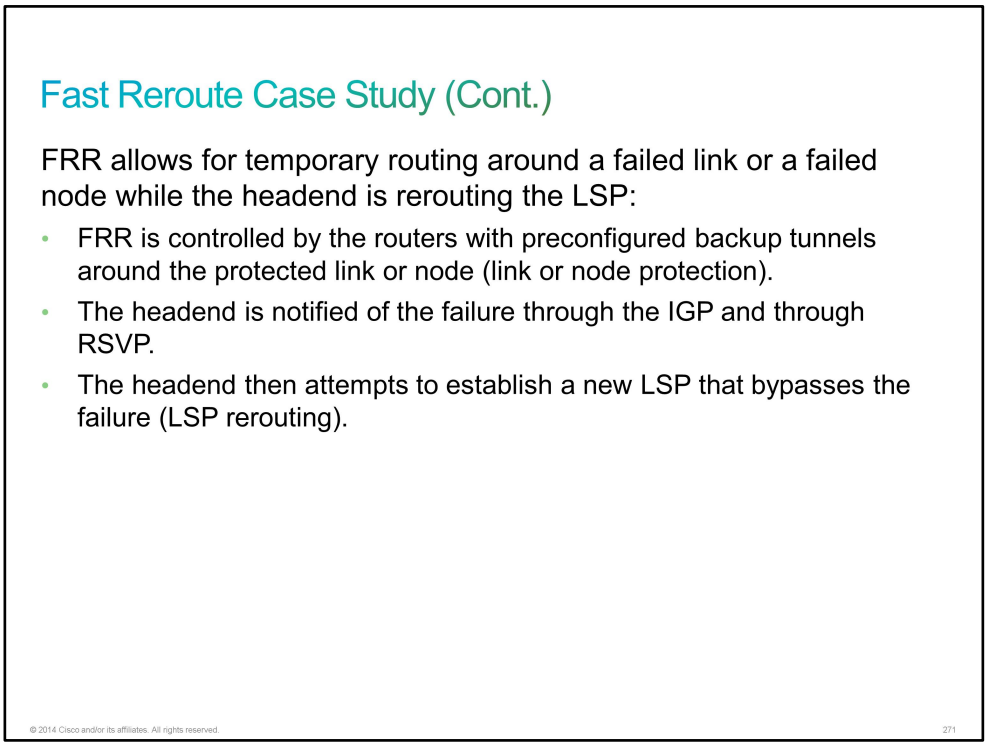

FRR is a mechanism for protecting a MPLS TE LSP from link and node failures by locally repairing the LSP at the point of failure. The FRR mechanism allows data to continue to flow while the headend router attempts to establish a new end-to-end LSP that bypasses the failure. FRR locally repairs any protected 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).

https://t.me/learningnets

Fast Reroute Case Study: Link Protection

Backup tunnels that bypass only a single link of the LSP path provide link protection.

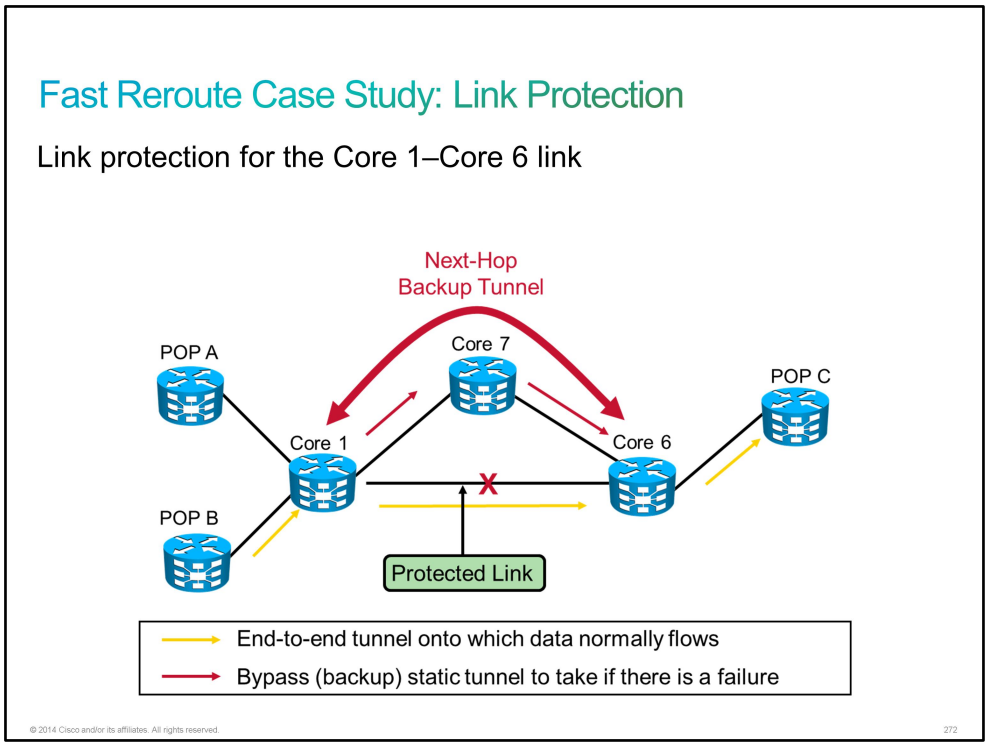

Backup tunnels bypass LSPs if a link along their path fails by rerouting the LSP traffic to the next hop (bypassing the failed link). These tunnels are referred to as next-hop backup tunnels because they terminate at the next hop of the LSP beyond the point of failure.

This process gives the headend of the tunnel time to reestablish the tunnel along a new, optimal route.

The node which redirects the traffic onto the preset backup path is called the PLR, and the node where a backup LSP merges with the primary LSP is called MP. The PLR in the example is the Core-1 router, and MP is the Core-6 router.

https://t.me/learningnets

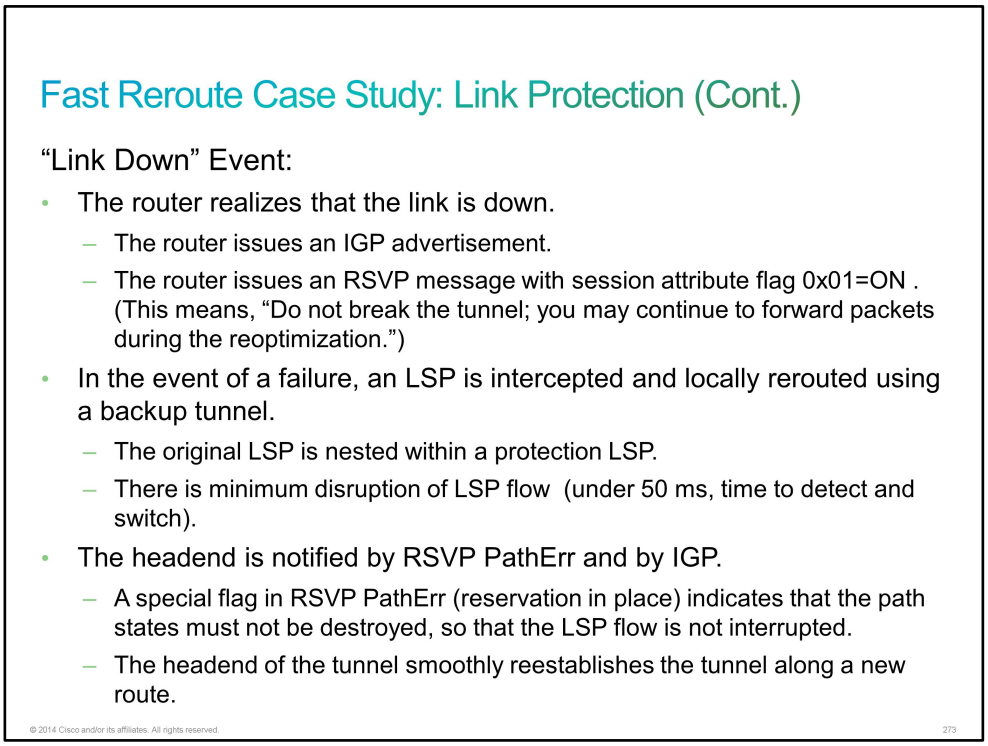

Paths for LSPs are calculated at the LSP headend. Under failure conditions, the headend determines a new route for the LSP. Recovery at the headend provides for the optimal use of resources. However, because of messaging delays, the headend cannot recover as fast as making a repair at the point of failure.

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 only caused 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, it attempts to establish a new LSP.

Normally, NHOP protects a complete link. At a given time, any number of TE LSPs with a certain amount of bandwidth might be crossing the protected link. But, the backup tunnel does not reserve bandwidth, therefore, if the protected tunnels use backup tunnel, there may be chances of some traffic getting dropped.

https://t.me/learningnets

Fast Reroute Case Study: Node Protection

FRR also provides node protection for LSPs.

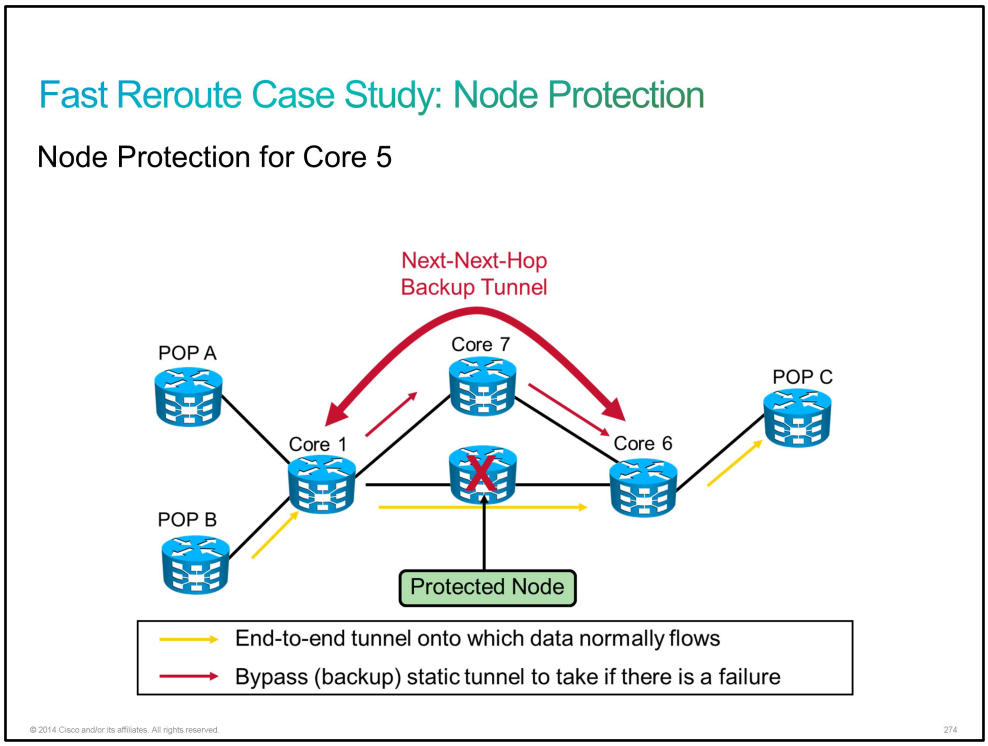

Node protection not only protects the link, but the whole router. Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop backup tunnels because they terminate at the node following the next-hop node of the LSP path, thereby bypassing the next-hop node. They protect LSPs, if 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. In node protection, next-next hop is the MP router.

https://t.me/learningnets

Fast Reroute Case Study: Node Protection (Cont.)

Node protection with FRR:

- 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 FRR:
  - Messages are periodically sent to the neighboring router.
  - If no response is received, hellos declare that the neighbor is down.
  - This causes any LSPs going out that interface to be switched to their respective backup tunnels.

© 2014 Cisco and/or its affiliates. All rights reserved.

When a router link or a neighboring node fails, the router often detects this failure by an “interface down” notification. On a GSRPOS 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.

https://t.me/learningnets

Fast Reroute Case Study: Fast Reroute Link Protection Configurations

This section assumes that you want to add FRR protection to a network in which MPLS TE LSPs are configured.

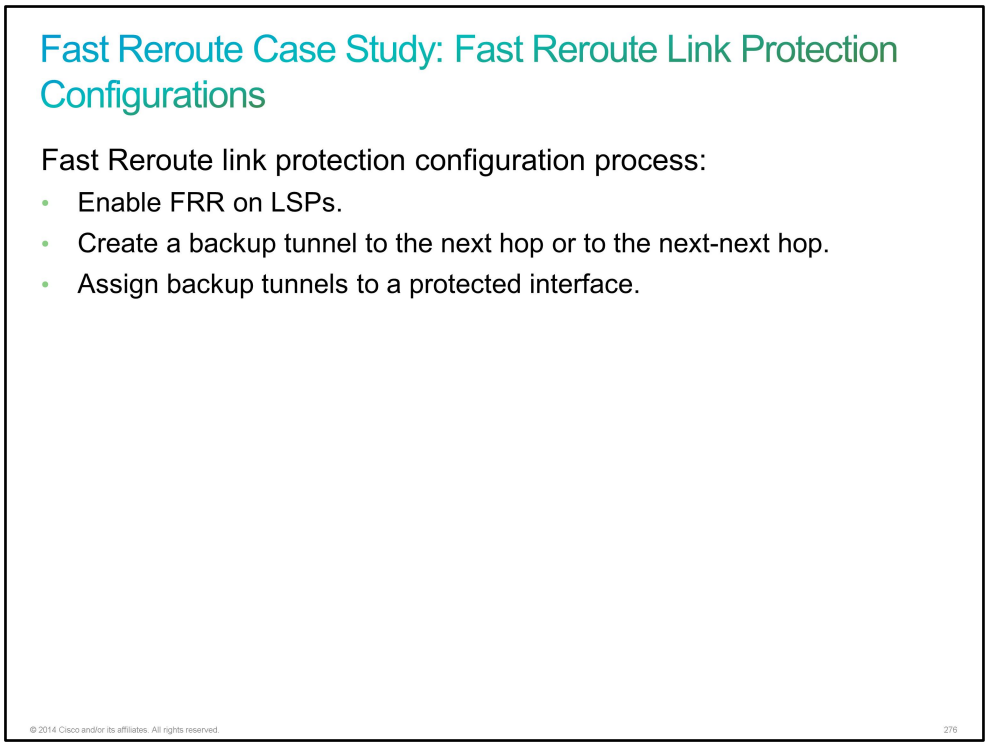

Before performing the configuration tasks in this topic, you must have done these tasks:

- 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:

- Enable FRR on LSPs.
- Create a backup tunnel to the next hop or to the next-next hop.
- Assign backup tunnels to a protected interface.

https://t.me/learningnets

Fast Reroute Case Study: Fast Reroute Link Protection Configurations (Cont.)

FRR configuration: Link protection

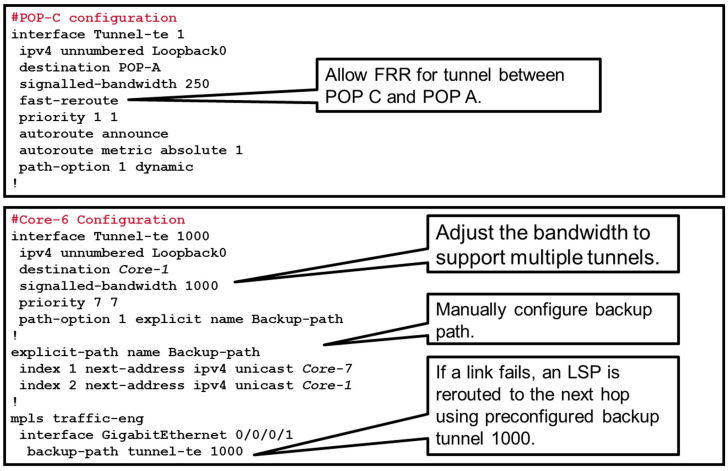

© 2014 Cisco and/or its affiliates. All rights reserved.

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

On Cisco IOS XR Software, use the **fast-reroute** interface configuration command to enable an MPLS TE tunnel to use a backup tunnel in the event of a link failure (if a backup tunnel exists).

On Cisco IOS XR Software, to configure the interface to use a backup tunnel in the event of a detected failure, use the **backup-path tunnel-te** command in the appropriate mode.

https://t.me/learningnets

```
# POP-C Cisco IOS Software 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 fast-reroute
  tunnel mpls traffic-eng priority 1 1
  tunnel mpls traffic-eng bandwidth 250
  tunnel mpls traffic-eng path-option 1 dynamic

#Core-6 Cisco IOS Software Configuration
interface Tunnel 1000
 ip unnumbered Loopback0
 tunnel destination Core-1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 NO tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng path-option 1 explicit name Backup-path
!
ip explicit-path name Backup-path enable
 next-address strict Core-7
 next-address strict Core-1
interface ser0/0
 mpls traffic-eng backup tunnel 1000
```


https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Backup

Configuring static backup tunnels for all core links in a large service provider network requires many statically configured tunnels with explicit paths. This is extremely difficult to implement and maintain. The autotunnel backup feature can be used to simplify and automate the deployment of many backup tunnels for link and node protection, using only a few configuration commands.

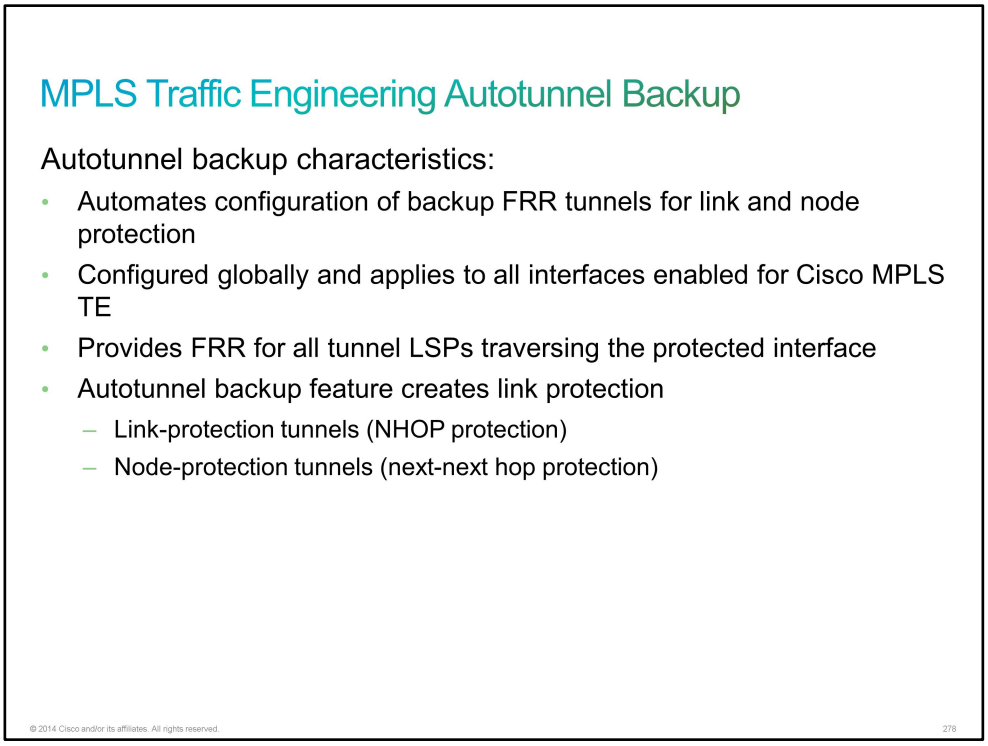

The autotunnel backup feature creates two types of tunnels:

- Link-protection tunnels around all core links (next-hop protection)
- Node-protection tunnels around all adjacent core and edge routers (next-next-hop protection)

https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Backup Configuration

This figure shows how to globally enable the autotunnel backup feature on a core or edge router to protect all MPLS-enabled core links. Defining the tunnel number range simplifies troubleshooting by making it easier for a network operator to distinguish different types of TE tunnels.

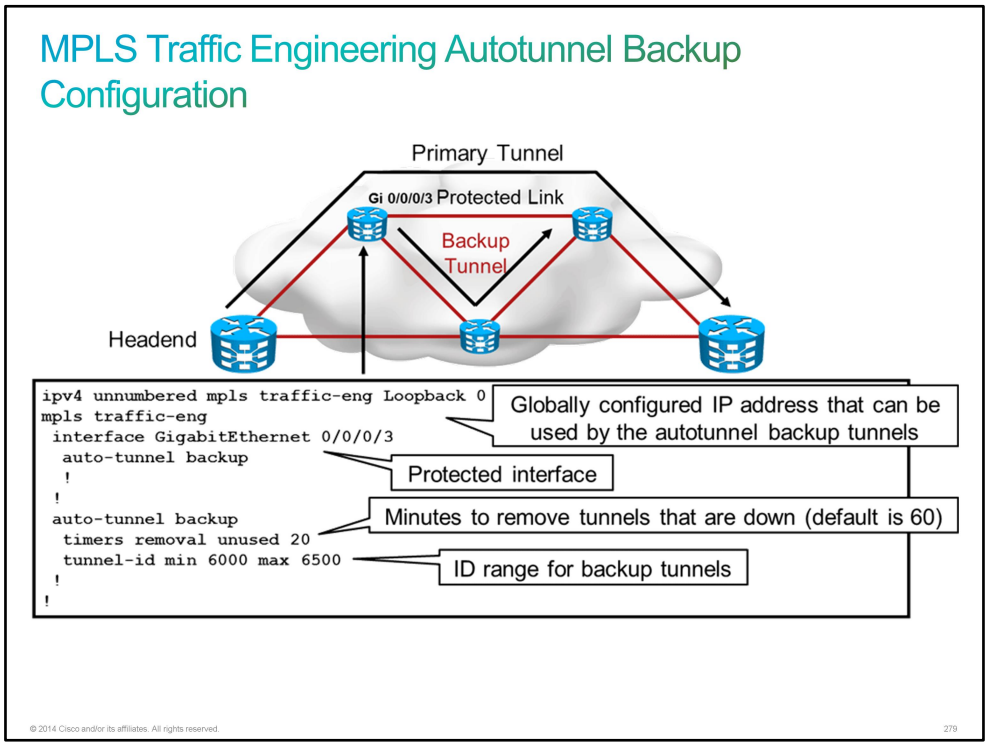

You can tune the autotunnel backup feature by changing two additional parameters:

- **Timers removal unused:** This value specifies a time, in minutes, after which an autotunnel mesh that is down gets deleted if its destination was not in the TE topology. The default value for this timer is 60. The timer gets started when these conditions are met:
  - The tunnel destination node is removed from the topology.
  - The tunnel is in the down state.
- **ID range for the backup tunnels:** You can adjust this setting to the requirements of your deployment

On Cisco IOS XR Software, you also need to configure the globally configured IPv4 address that can be used by the autotunnel backup tunnels. The autotunnel backup tunnels will not come up until a global IPv4 address is configured.

To enable autotunnel backup feature on interface, use the **auto-tunnel backup** Cisco IOS XR Software command in MPLS TE interface configuration mode.

To enable autotunnel feature on Cisco IOS Software, use the **mpls traffic-eng auto-tunnel backup** global configuration command.

https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Backup Verification

There are several ways to verify that backup tunnels have been established.

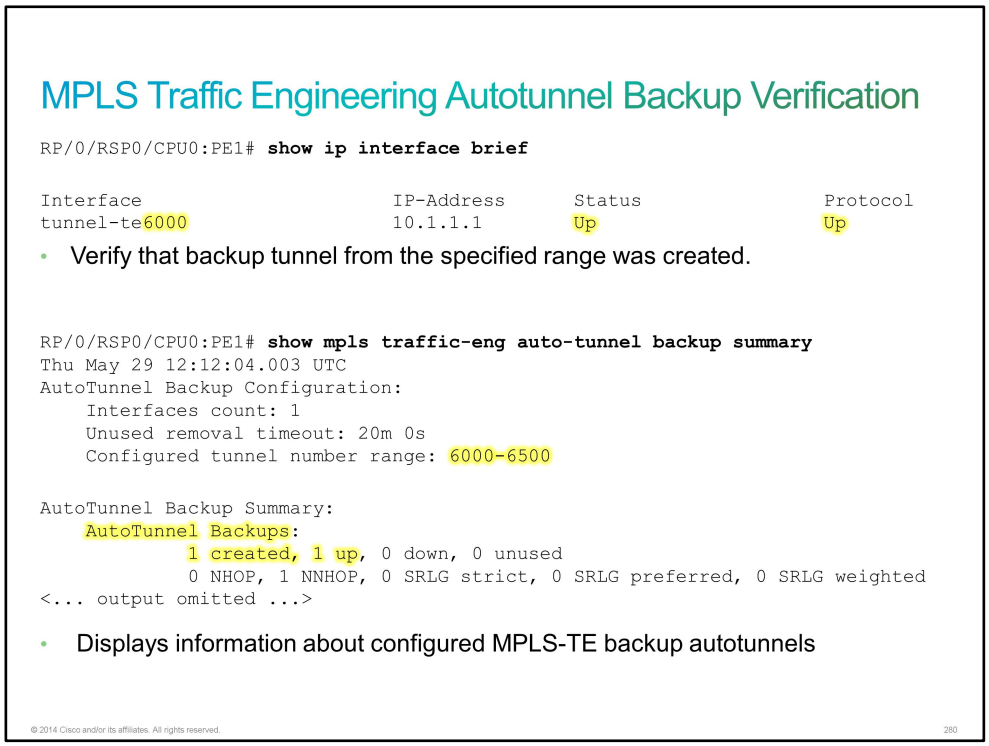

When monitoring the autotunnel backup feature, verify that MPLS TE interfaces from the specified ID range for the backup tunnels were created. In the example, interface tunnel-te6000 was created automatically.

Use the **show mpls traffic-eng auto-tunnel backup summary** command on Cisco IOS XR Software to verify information about configured MPLS-TE backup autotunnels.

Use the **show mpls traffic-eng auto-tunnel backup** command on Cisco IOS Software to verify information about configured MPLS-TE backup autotunnels. The sample output is shown below.

```
PE2# show mpls traffic-eng auto-tunnel backup

State: Enabled
  Auto backup tunnels: 2 (up: 2, down: 0)
  Tunnel ID Range: 65436 - 65535
  Create Nhop Only: No
  Check for deletion of unused tunnels every: 3600 Sec
  SRLG: Not configured

  Config:
    unnumbered-interface: Loopback0
    Affinity/Mask: 0x0/0xFFFF
```


https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Mesh

The Cisco multiprotocol label switching traffic engineering autotunnel mesh (automesh) feature allows you to set up a full mesh of TE point-to-point tunnels automatically, with a minimal set of MPLS TE configurations.

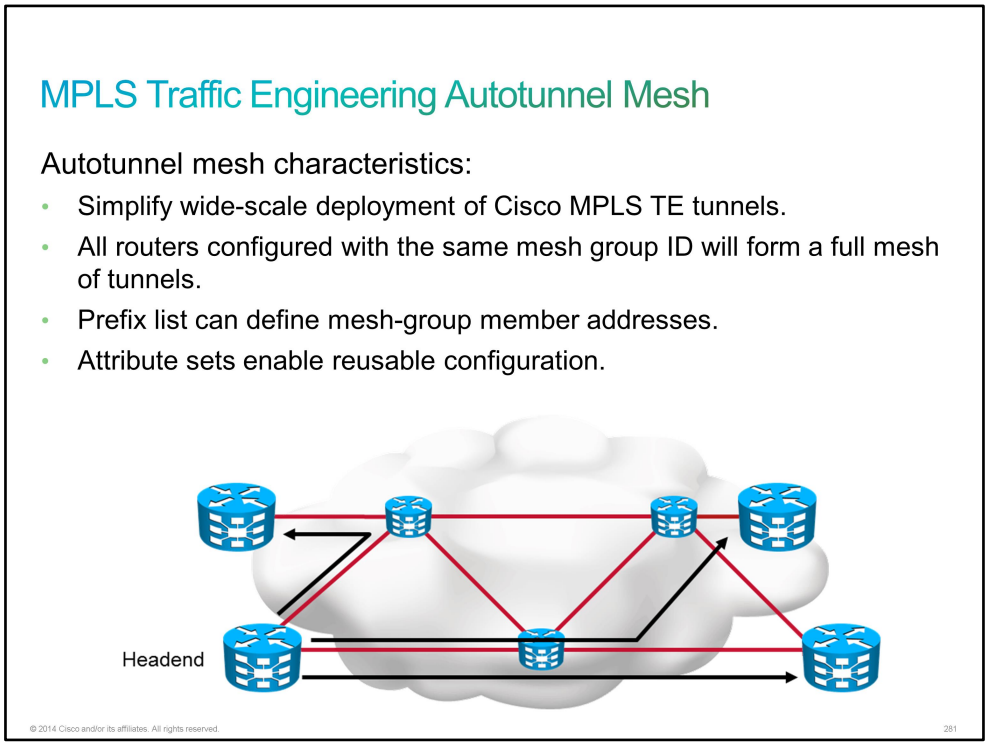

Automesh tunnels provide the following benefits:

- **Minimizes the initial configuration of the network**: You can configure tunnel properties templates and mesh groups or destination lists on each TE LSR, that creates a full mesh of TE tunnels between those LSRs.
- **Minimizes future configurations due to network growth**: It eliminates the need to reconfigure each existing TE LSR in order to establish a full mesh of TE tunnels whenever a new TE LSR is added in the network.

Autotunnel mesh is a set of connections between edge LSRs in service provider network.

You can configure one or more mesh groups. Each mesh group requires a destination-list (IPv4 prefix-list) that lists destinations, which are used as destinations for creating tunnels for that mesh group.

You can configure MPLS TE automesh type attribute sets (templates) and associate them to mesh groups. LSR creates tunnels using the tunnel properties that are defined in the attribute set.

https://t.me/learningnets

MPLS Traffic Engineering Autotunnel Mesh Configuration

An attribute-set template defines settings that can be applied to automesh tunnels. Parameters that you will typically set on the tunnel headend include required bandwidth, autoroute, fast-reroute, AutoBandwidth, and event logging. Event logging can specify state, insufficient bandwidth, reoptimize, or reroute events.

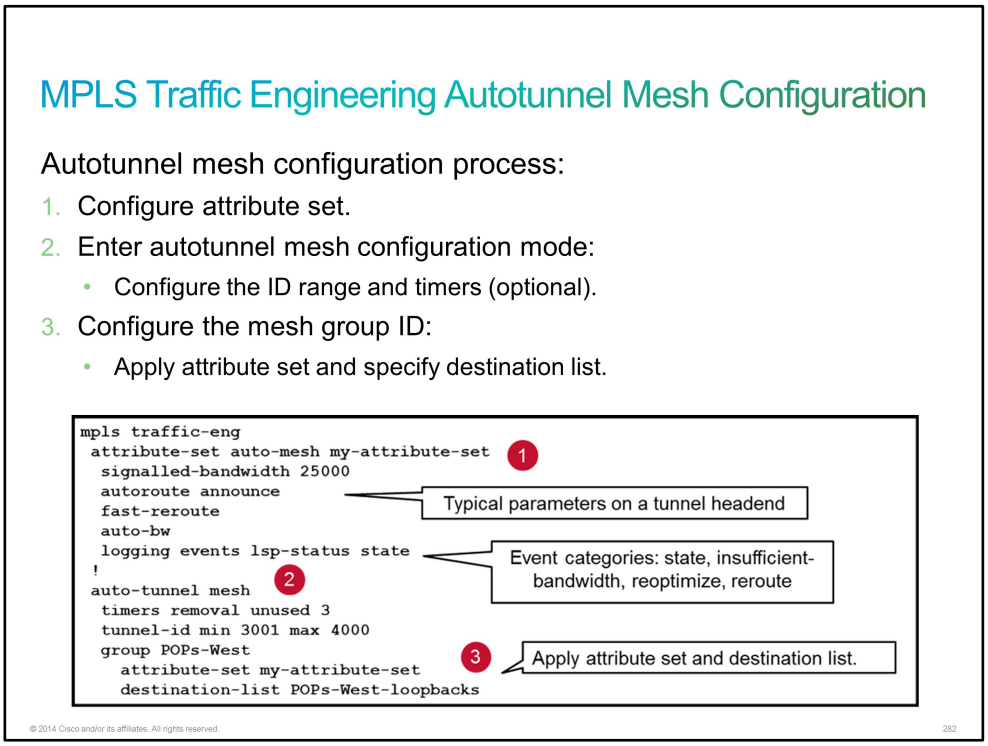

Next, you will enter the autotunnel mesh configuration mode (submode of the MPLS traffic engineering mode), and optionally tune the removal timer and configure the tunnel ID range. Within the autotunnel mesh configuration mode, you will define a group and configure two elements for it:

- **Attribute set:** This command applies the previously defined template to the given group (POPs-West in this example).
- **Destination list:** This list defines all routers that should belong to the mesh group. If a given destination-list is not defined as a prefix-list, this mesh group creates tunnels to all nodes that are available in the TE topology. In the example, destination list is referring to the prefix-list that is named POPs-West-Loopbacks.

Automesh tunnels can be automatically created using the prefix-list. Each TE-enabled router in the network learns about the TE router IDs through an existing IGP extension.

A prefix-list can be configured on each TE router to match a desired set of router IDs. For example, if a prefix-list is configured to match addresses of 100.0.0.0 with wildcard 0.255.255.255, then all 100.x.x.x router IDs are included in the automesh group.

When a new TE router is added in the network and its router ID is also in the block of addresses described by the prefix-list, for example, 100.x.x.x, then it is added in the automesh group on each existing TE router without having to explicitly modify the prefix-list or perform any additional configuration.

https://t.me/learningnets

Automesh does not create tunnels to its own (local) TE router IDs.

https://t.me/learningnets

MPLS TE End-to-End Path Protection

The MPLS TE path protection feature provides an end-to-end failure recovery mechanism (that is, a full path protection) for MPLS-TE tunnels. A secondary LSP is established, in advance, to provide failure protection for the protected LSP that is carrying the tunnel TE traffic.

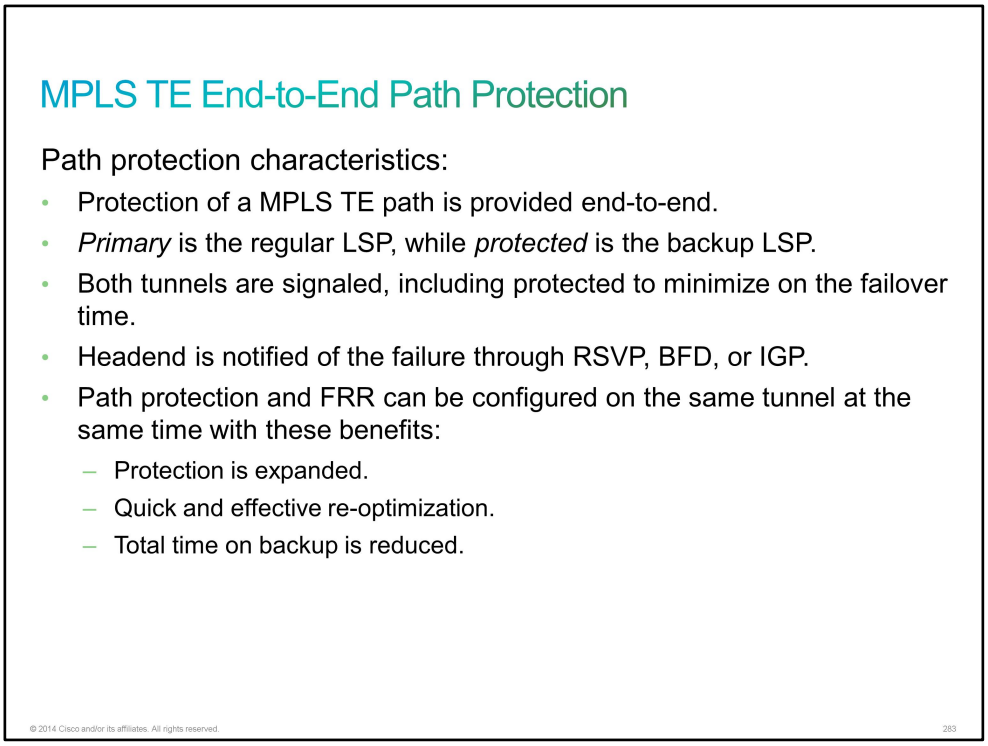

When there is a failure on the protected LSP, the source router immediately enables the secondary LSP to temporarily carry the tunnel traffic. If there is a failure on the secondary LSP, the tunnel no longer has path protection until the failure along the secondary path is cleared. Path protection can be used within a single area (OSPF or IS-IS), eBGP, and static routes.

The failure detection mechanisms triggers a switchover to a secondary tunnel by:

- Path error or resv-tear from RSVP signaling
- Notification from the BFD protocol that a neighbor is lost
- Notification from the IGP that the adjacency is down
- Local teardown of the protected tunnel LSP due to pre-emption in order to signal higher priority LSPs, a POS alarm, OIR, and so on

An alternate recovery mechanism is FRR, which protects MPLS-TE LSPs only from link and node failures, by locally repairing the LSPs at the point of failure. Coexistence of FRR and path protection is supported; this means FRR and path-protection can be configured on the same tunnel at the same time.

Although not as fast as link or node protection, presignaling a secondary LSP is faster than configuring a secondary primary path option, or allowing the tunnel source router to dynamically recalculate a path. The actual recovery time is topology-dependent, and affected by delay factors such as propagation delay or switch fabric latency.

https://t.me/learningnets

Path protection and FRR can be configured on the same tunnel at the same time. The coexistence of path protection and FRR on the same tunnel provides these benefits:

- **Protection is expanded:** having an FRR protected tunnel that is also path-protected ensures that failures of nonprotected links on the primary path are handled more efficiently by a quick switch over to the presignaled standby LSP.
- **Quick and effective re-optimization:** having a precomputed standby LSP allows the system to minimize reoptimization LSP path calculation and signaling, by simply switching over to the presignaled standby LSP. Effectively, path protection switch over replaces the post-FRR LSP down event reoptimization.
- **Total time on backup is reduced:** handling FRR failure using a path protection switchover reduces total time on backup because the traffic is diverted from the backup to the standby, as soon as the headend receives the FRR LSP down notification, without having to wait for a reoptimization LSP.

https://t.me/learningnets

MPLS TE End-to-End Path Protection Configuration

This figure shows how to configure path protection on the MPLS TE tunnel. Path protection for tunnel 10 is configured, where one explicit path is protected by another explicit path.

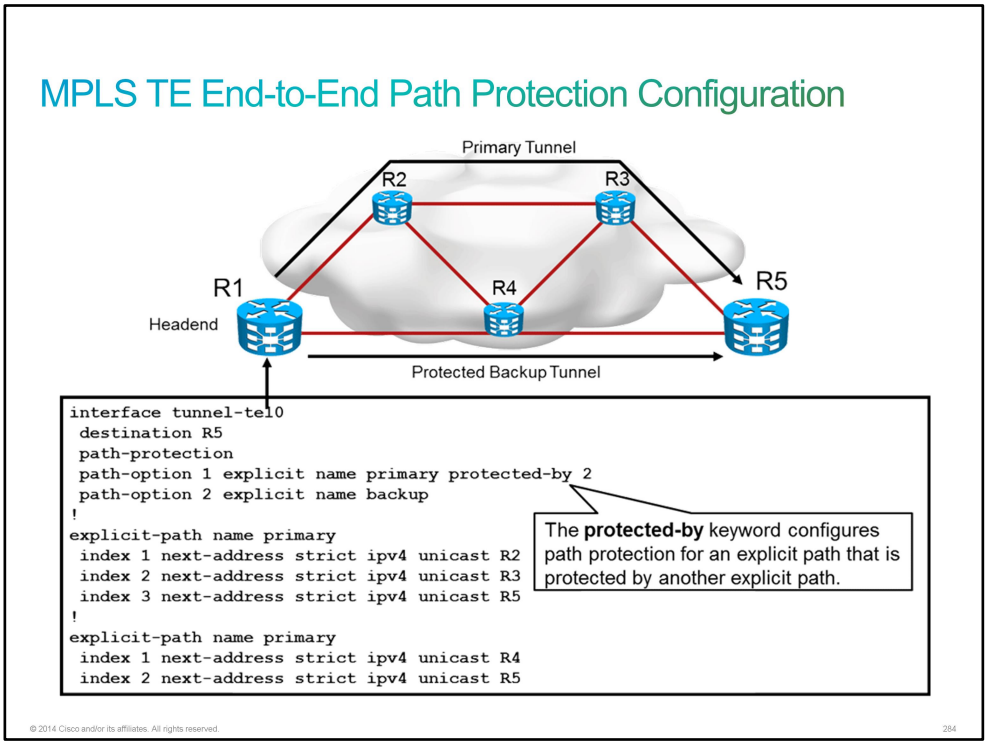

Use the **path-protection** Cisco IOS XR Software command in MPLS TE tunnel configuration mode to enable path protection on the tunnel-te interface.

Use the **protected-by** keyword in **path-option** command to configure path protection for an explicit path that is protected by another explicit path

https://t.me/learningnets

BFD over MPLS TE LSPs

The BFD over MPLS TE LSPs feature in Cisco IOS XR Software detects MPLS LSP data plane failures. Since the control plane processing required for BFD control packets is relatively smaller than the processing required for LSP ping messages, BFD can be deployed for faster detection of data plane failure for a large number of LSPs.

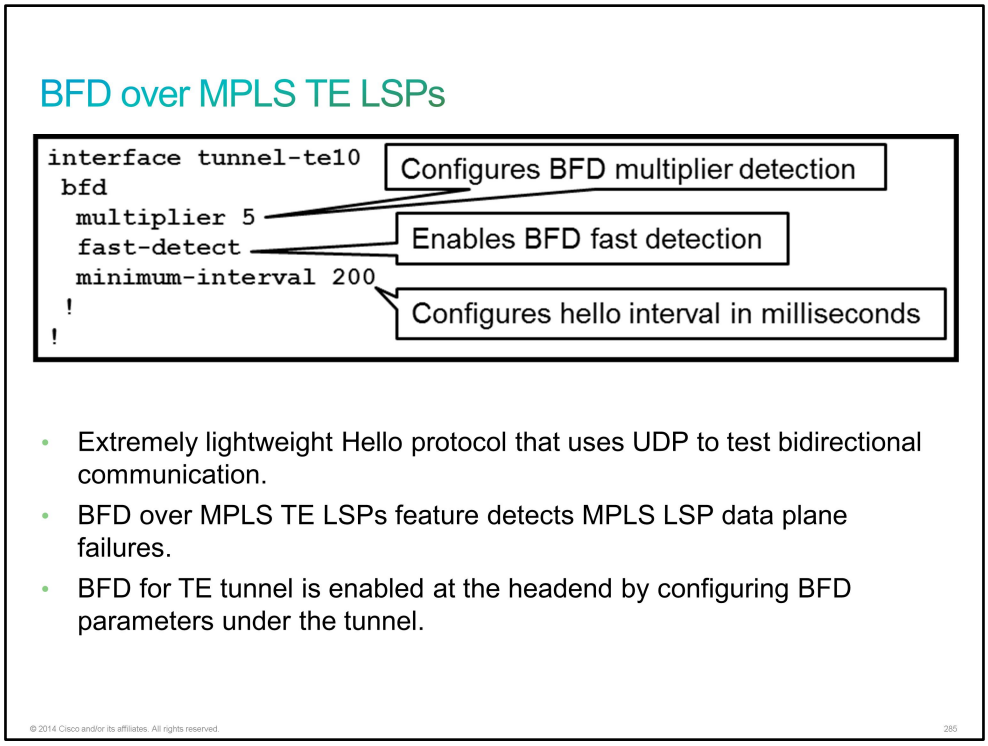

BFD provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers, including the interfaces, data links, and forwarding planes. BFD is a detection protocol that you enable at the interface and routing protocol levels. Cisco supports the BFD asynchronous mode, that depends on the sending of BFD control packets between two systems to activate and maintain BFD neighbor sessions between routers. Therefore, in order for a BFD session to be created, you must configure BFD on both BFD peers. Once BFD has been enabled on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers will begin to send BFD control packets to each other at the negotiated interval.

The BFD over MPLS TE LSPs is used for networks that have deployed MPLS as the multiservice transport and that use BFD as fast failure detection mechanism to enhance network reliability and up time by using BFD as fast failure detection traffic black holing.

BFD for TE tunnel is enabled at the headend by configuring BFD parameters under the tunnel. When BFD is enabled on the already up tunnel, TE waits for the bring-up timeout before bringing down the tunnel. BFD is disabled on TE tunnels by default.

Use the **bfd fast-detect** command in interface MPLS TE tunnel configuration mode to enable BFD fast detection.

You might also change the BFD hello interval or BFD multiplier as follows:

https://t.me/learningnets

- Use the **bfd minimum-interval** command in interface MPLS TE tunnel configuration mode to configure hello interval in milliseconds. The hello interval range is 100 to 30000 milliseconds. The default hello interval is 100 milliseconds
- Use the **bfd multiplier** command in interface MPLS TE tunnel configuration mode to configure BFD multiplier detection. BFD multiplier range is 3 to 10. Default BFD multiplier is 3.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS TE Bandwidth Control

After initial TE is complete, network administrators may need an effective way to continually adjust tunnel routes and bandwidth reservations without doing any redesigning.

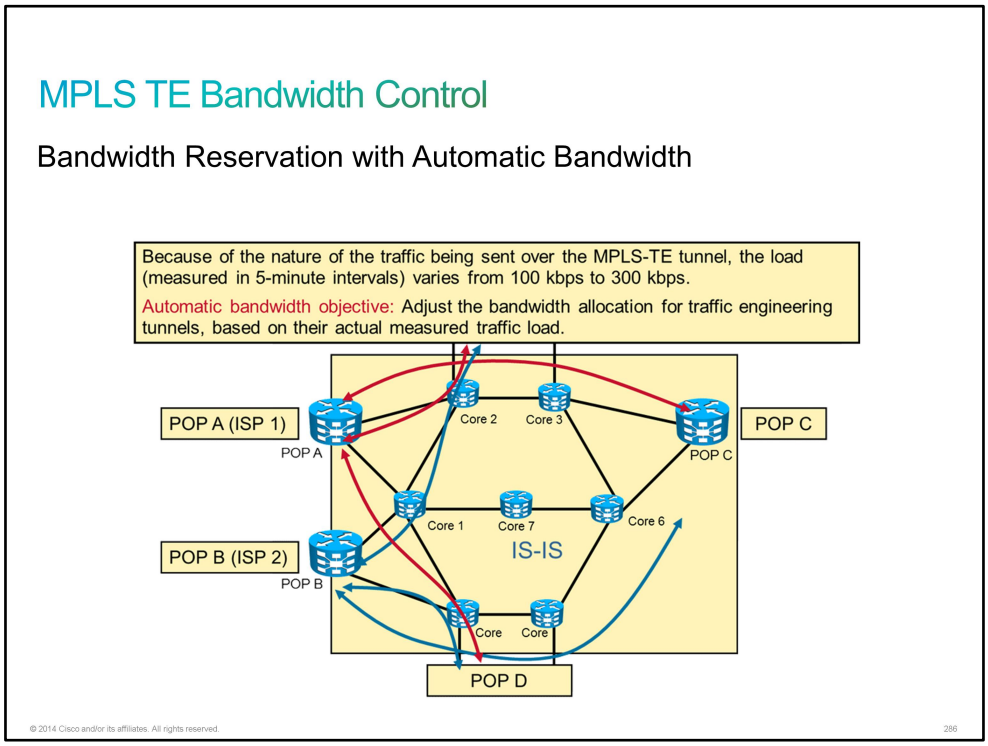

Both Cisco IOS Software and Cisco IOS XR Software have an automatic bandwidth adjustment feature that measures utilization averages and dynamically adjusts tunnel bandwidth reservations to meet actual application resource requirements.

This powerful feature creates “self-tuning tunnels” that relieve network administrators of many of the daily hands-on management tasks that are necessary with other TE techniques.

The MPLS TE automatic bandwidth feature measures the traffic in a tunnel and periodically adjusts the signaled bandwidth for the tunnel. MPLS TE automatic bandwidth is configured on individual LSPs at every headend. MPLS TE monitors the traffic rate on a tunnel interface. Periodically, MPLS TE resizes the bandwidth on the tunnel interface to align it closely with the traffic in the tunnel. The MPLS TE automatic bandwidth feature can perform these functions:

- Monitor periodic polling of the tunnel output rate.
- Resize the tunnel bandwidth by adjusting the highest rate observed during a given period.

https://t.me/learningnets

MPLS TE Bandwidth Control (Cont.)

The 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 the traffic output of the 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.

© 2014 Cisco and/or its affiliates. All rights reserved.

The TE automatic bandwidth adjustment feature provides the means to automatically adjust the bandwidth allocation for TE tunnels based on their measured traffic load.

The TE automatic bandwidth adjustment feature 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, both the sampling interval and the interval used to average the tunnel traffic to determine the average output rate are user-configurable on a per-tunnel basis.

The benefit of the automatic bandwidth 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 enabled tunnel 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.

https://t.me/learningnets

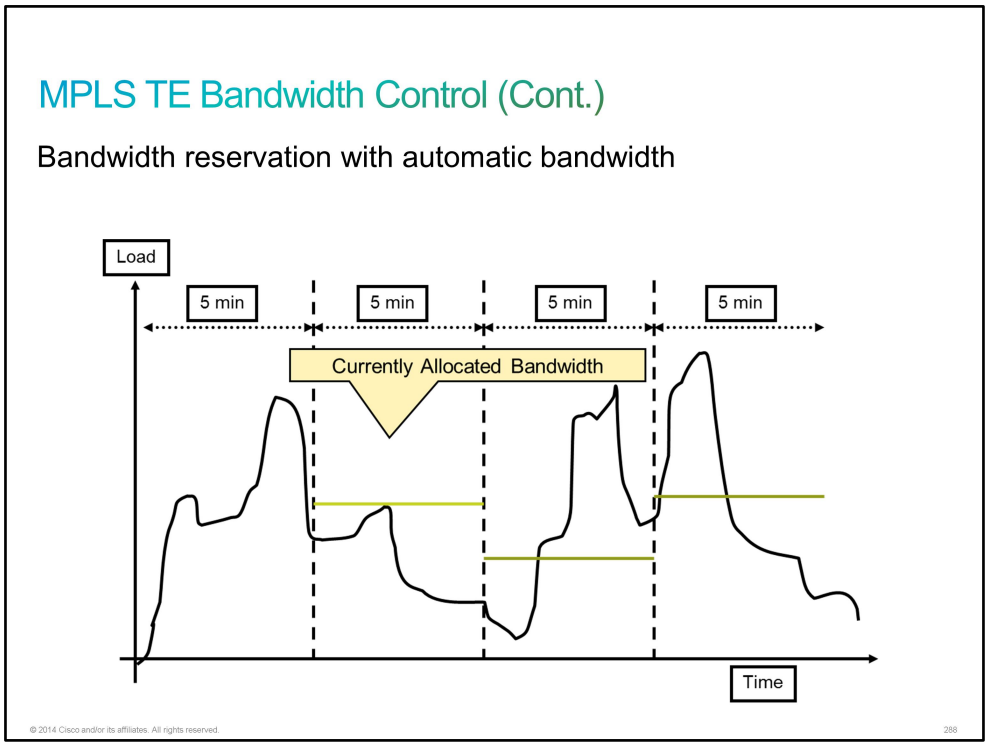

The diagram shows the load on the tunnel and intervals of measurement. The input and output rates on the tunnel interfaces are averaged over a predefined interval (load interval). In the example, the interval is the previous 5 minutes.

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.

https://t.me/learningnets

Automatic Bandwidth Configuration

The example in the figure shows the setting of MPLS TE tunnels that can actually tune their own bandwidth requirements to increase or decrease their RSVP reservations, as warranted by changing network conditions.

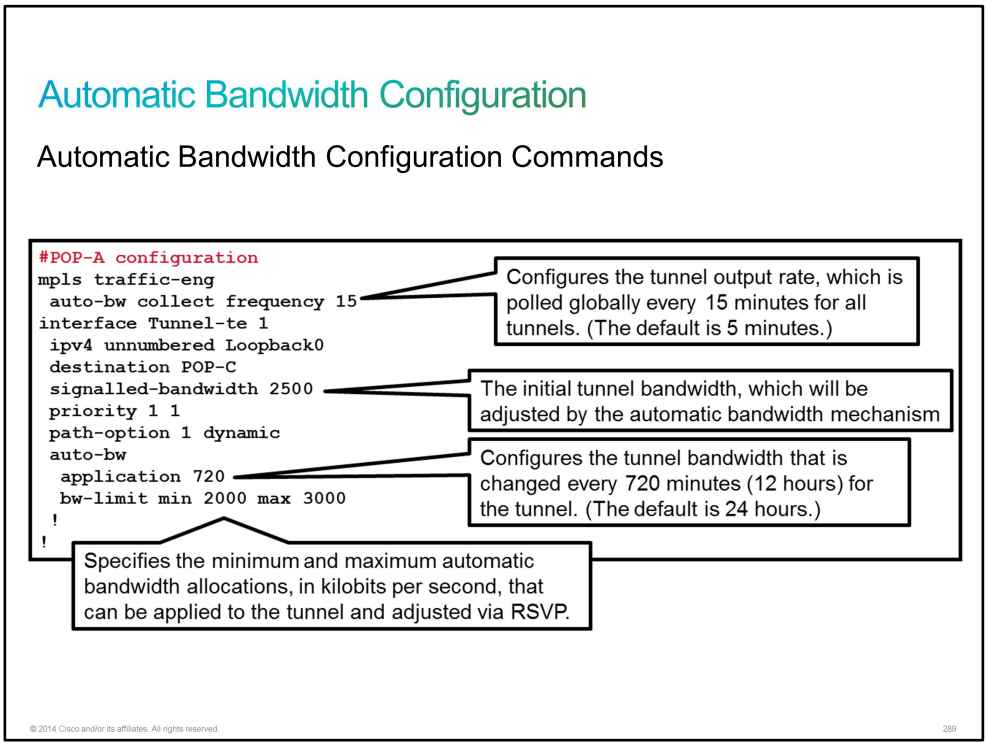

The example in the figure shows the setting of MPLS traffic-engineered tunnels that can actually tune their own bandwidth requirements to increase or decrease their RSVP reservations, as warranted by changing network conditions.

When readjusting bandwidth constraint on a tunnel, a new 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.

For every MPLS TE tunnel that is configured for automatic bandwidth adjustment, 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.

Automatic bandwidth allocation monitors the X minutes (default = 5 minutes) 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 automatic bandwidth feature is implemented with the following commands on Cisco IOS XR Software:

auto-bw collect frequency *minutes*

- Configures the automatic bandwidth collection frequency, and controls the manner in which the bandwidth for a tunnel collects output rate information but does not adjust the tunnel bandwidth. By default, this value is set to 5 minutes.

application *minutes*

https://t.me/learningnets

- Configures the application frequency in minutes for the applicable tunnel. By default, the frequency is 24 hours.

bw-limit {**min** *bandwidth* } {**max** *bandwidth*}

- Configures the minimum and maximum automatic bandwidth in kilobits per second, set on a tunnel.

The automatic bandwidth feature is implemented with these commands on Cisco IOS Software:

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 seconds (5 minutes).

clear mpls traffic-eng auto-bw timers

- This command is used to clear the timers that were 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** value limits the maximum bandwidth that a tunnel can adjust to. Similarly, setting the **min-bw** value provides the smallest bandwidth that the tunnel can adjust to. When both **max-bw** and **min-bw** values are specified, the tunnel bandwidth remains between these values.

```
#POP-A Cisco IOS 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
```


https://t.me/learningnets

DiffServ-Aware MPLS TE Tunnels

MPLS DS-TE is an extension of the regular MPLS TE feature. Regular traffic engineering does not provide bandwidth guarantees to different traffic classes. A single bandwidth constraint is used in regular TE that is shared by all traffic. To support various CoS, users can configure multiple bandwidth constraints. These bandwidth constraints can be treated differently, based on the requirement for the traffic class using that constraint.

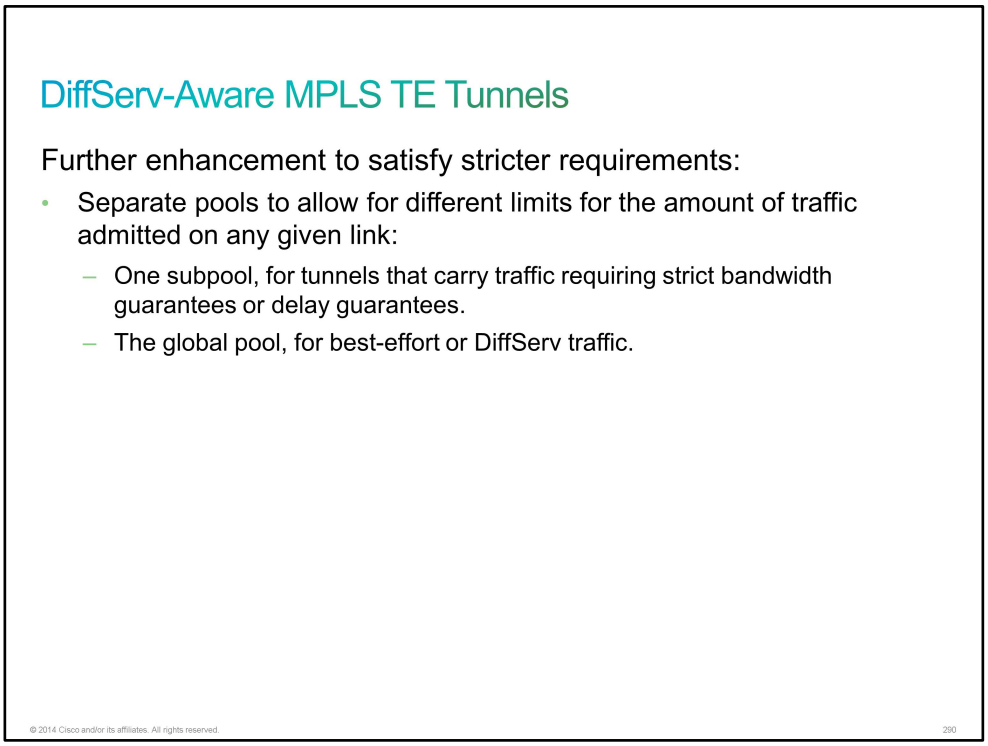

MPLS DS-TE enables you to configure multiple bandwidth constraints on an MPLS-enabled interface. Available bandwidths from all configured bandwidth constraints are advertised using IGP. The TE tunnel is configured with bandwidth value and class-type requirements. Path calculation and admission control take the bandwidth and class-type into consideration. RSVP is used to signal the TE tunnel with bandwidth and class-type requirements.

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.

https://t.me/learningnets

DiffServ-Aware MPLS TE Tunnels (Cont.)

DiffServ-Aware service characteristics:

- 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.

© 2014 Cisco and/or its affiliates. All rights reserved.

DS-TE extends the current MPLS TE capabilities to introduce the awareness of a particular class of traffic, that 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, QoS mechanisms cannot ensure better service for the voice. Usually, 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. 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.

https://t.me/learningnets

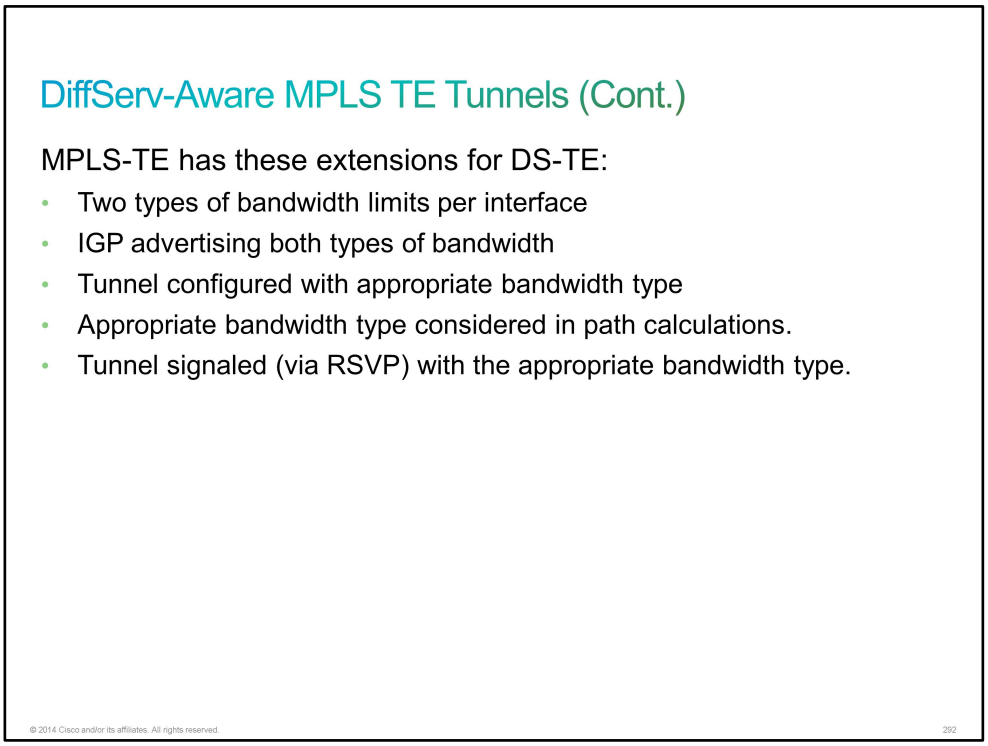

DS-TE tunnels are similar to regular TE tunnels. To support DS-TE, the following modifications to regular MPLS TE mechanisms have been made:

- There are two types of bandwidth per each link in the network (two bandwidth pools: the global pool and the subpool).
- Both of these bandwidths are 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 PCALC is done with respect to the type of the bandwidth that the tunnel requires. In RSVP messages, whether the LSP to be set up is a regular MPLS TE tunnel or a DS-TE tunnel is always indicated. Intermediate nodes perform admission control and bandwidth allocation ("locking" for the DS-TE) for the appropriate bandwidth pool.

https://t.me/learningnets

DiffServ-Aware MPLS Dual-Bandwidth Pools

There are two types of bandwidth per each link in the network (two bandwidth pools: the global pool and the subpool). Both of these bandwidths are announced in the link-state updates that carry resource information.

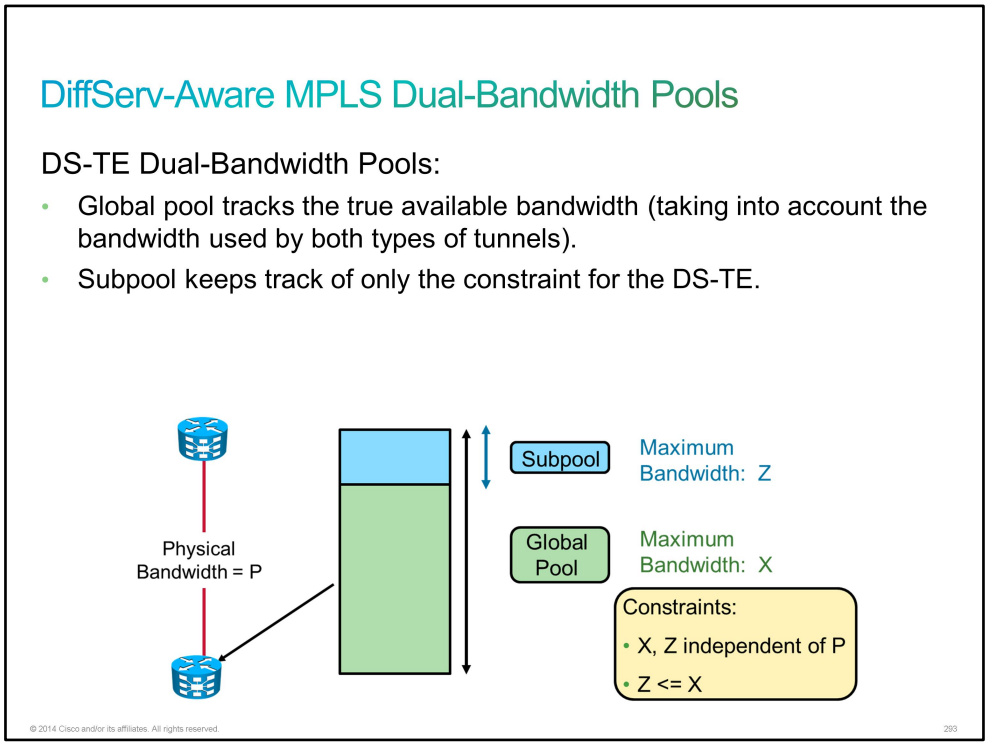

On each link in the network, two bandwidth pools are established:

- The global (main) pool keeps track of the true available bandwidth. The pool takes into account the bandwidth that is used by both of the tunnels.
- The subpool (DS-TE) 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.

https://t.me/learningnets

DiffServ-Aware Service Configuration

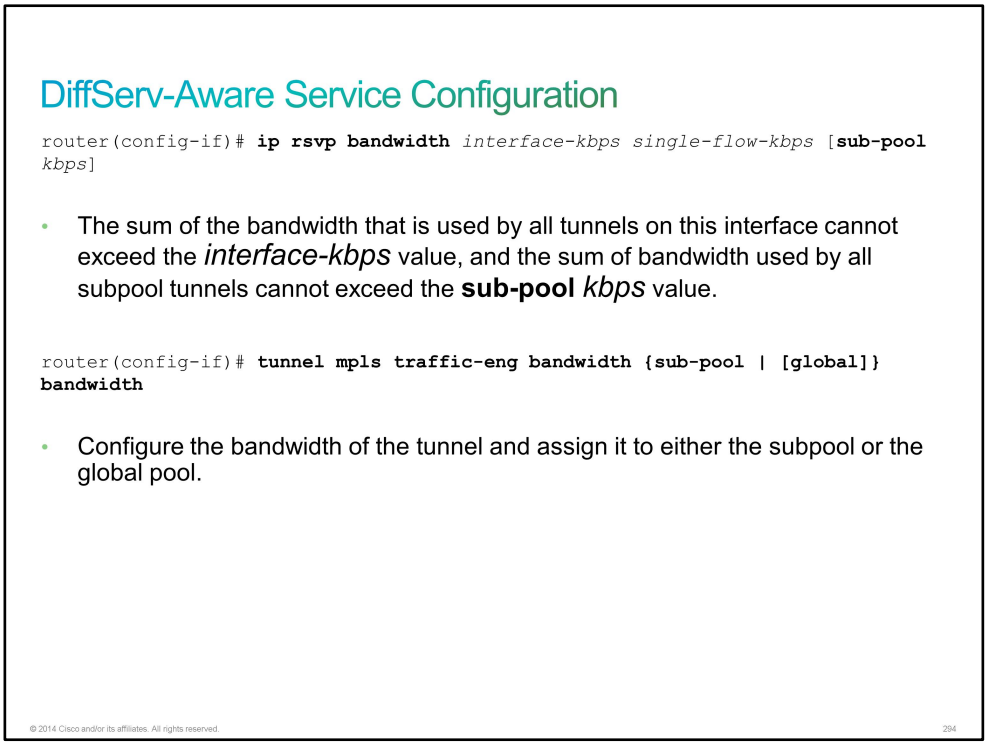

To configure RSVP bandwidth on an interface using prestandard DS-TE mode, use the **bandwidth** command in RSVP interface configuration mode. To reset the RSVP bandwidth on that interface to its default value, use the **no** form of this command.

bandwidth [*total reservable bandwidth*] [**bc0** *bandwidth*] [**global-pool** *bandwidth*] [**sub-pool** *reservable-bw*]

Syntax Description

Parameter	Description
total reservable bandwidth	(Optional) Total reservable bandwidth (in kilobits per second) that RSVP accepts for reservations on this interface. The range is 0 to 4294967295 kbps.
bc0 *bandwidth*	Configures the total reservable bandwidth in the bc0 pool (in kilobits, megabits, or gigabits per second). The default is kilobits per second. The range is 0 to 4294967295.
global-pool *bandwidth*	(Optional) Configures the total reservable bandwidth in the global pool. The range is 0 to 4294967295 kbps.
sub-pool *kbps*	Amount of bandwidth (in kilobits per second) on the interface that is to be reserved to a portion of the total. The range is from 1 to the value of *interface-kbps*.

https://t.me/learningnets

To configure the bandwidth that is required for an MPLS TE tunnel, use the **signalled-bandwidth** command in interface configuration mode. To return to the default behavior, use the **no** form of this command.

```
signalled-bandwidth {bandwidth [class-type ct] | sub-pool bandwidth}
no signalled-bandwidth {bandwidth [class-type ct] | sub-pool bandwidth}
tunnel mpls traffic-eng bandwidth {sub-pool | [global]} bandwidth
```

Syntax Description

Parameter	Description
bandwidth	Bandwidth that is required for an MPLS TE tunnel. Bandwidth is specified in kilobits per second. By default, bandwidth is reserved in the global pool. The range is 0 to 4294967295.
class-type *ct*	(Optional) Configures the class type of the tunnel bandwidth request. The range is 0 to 1. Class-type 0 is strictly equivalent to global-pool. Class-type 1 is strictly equivalent to the subpool.
sub-pool *bandwidth*	Reserves the bandwidth in the subpool instead of the global pool. The range is 1 to 4294967295. A subpool bandwidth value of 0 is not allowed.

DiffServ-Aware Service Configuration on Cisco IOS Software

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
interface-kbps	Amount of bandwidth (in kilobits per second) on the interface to be reserved. The range is 1 to 10000000.
single-flow-kbps	Amount of bandwidth (in kilobits per second) that is allocated to a single flow (ignored in DS-TE). The range is 1 to 10000000.
sub-pool *kbps*	Amount of bandwidth (in kilobits per second) on the interface to be reserved to a portion of the total. The range is from 1 to the value of *interface-kbps*.

To configure the bandwidth that is required for an 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
```


https://t.me/learningnets

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**.
bandwidth	The bandwidth, in kilobits per second, that is set aside for the MPLS TE tunnel. The range is between 1 and 4294967295.

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 PHB, to be used exclusively by the strict guaranteed traffic.

If delay or jitter guarantees are sought, the DiffServ EF PHB is used. 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 AF PHB is used.

2. Ensure that the guaranteed traffic that is sent through the subpool tunnel is placed in the 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 queue.

3. Ensure that this queue is never oversubscribed; that is, see that no more traffic is sent into the subpool tunnel than the queue can manage.

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 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 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.

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).

https://t.me/learningnets

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.

https://t.me/learningnets

Summary

This topic summarizes the key points that were discussed in this lesson.

Summary

- Having two preestablished paths is the simplest form of MPLS TE path protection.
- Any of the available methods can be used to advertise the tunnel.
- Solution with backup tunnels presents some drawbacks (allocation of labels, resource reservation).
- FRR is a mechanism for protecting a MPLS TE LSP from link and node failures by locally repairing the LSP at the point of failure.
- Backup tunnels that bypass only a single link of the LSP path provide link protection.
- Backup tunnels that bypass next-hop nodes along LSP paths are called next-next-hop backup tunnels.
- On Cisco IOS XR Software, use the **fast-reroute** interface configuration command to enable an MPLS TE tunnel.

© 2014 Cisco and/or its affiliates. All rights reserved.

Summary (Cont.)

- Autotunnel backup feature can be used to simplify and automate the deployment of many backup tunnels for link and node protection.
- Autotunnel mesh feature allows you to set up a full mesh of TE point-to-point tunnels automatically.
- MPLS TE path protection feature provides an end-to-end failure recovery mechanism for MPLS-TE tunnels.
- BFD over MPLS TE LSPs feature in Cisco IOS XR Software detects MPLS label-switched path LSP data plane failures.
- The MPLS TE automatic bandwidth feature measures the traffic in a tunnel and periodically adjusts the signaled bandwidth for the tunnel.
- MPLS DS-TE is an extension of the regular MPLS TE feature.

© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

https://t.me/learningnets

https://t.me/learningnets

Lesson 5: Module Summary

This topic summarizes the key points that were discussed in this module.

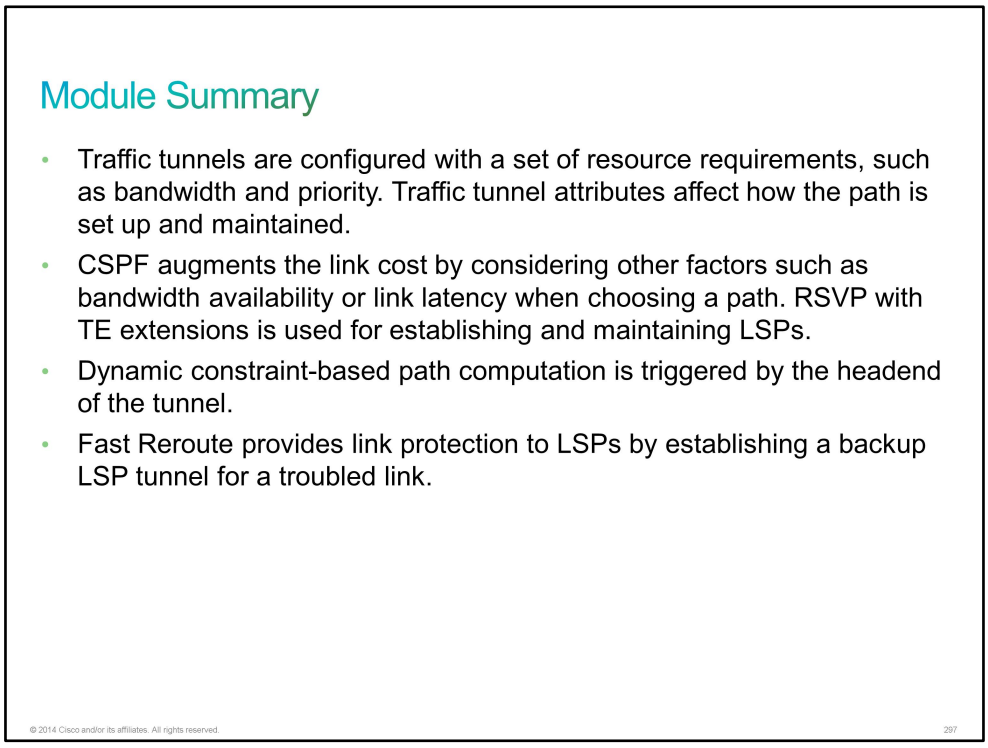

This module discussed the requirements for TE in modern service provider networks that must attain optimal resource utilization. The traffic-engineered tunnels provide a means of mapping traffic streams onto available networking resources in a way that prevents the overuse of subsets of networking resources while other subsets are underused.

All the concepts and mechanics that support TE were presented, including tunnel path discovery with link-state protocols and tunnel path signaling with RSVP. Some of the advanced features of TE, such as automatic bandwidth allocation and guaranteed bandwidth, are introduced as well. LSP setup is always initiated at the headend of a tunnel. TE tunnels can be used for IP routing only if the tunnels are explicitly specified for routing.

https://t.me/learningnets

This module explained the configuration of routers to enable basic traffic tunnels, the assignment of traffic to a tunnel, the control of path selection, and the performance of tunnel protection and tunnel maintenance. Configurations were shown for various Cisco platforms.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 6: Module Self-Check

Use the questions here to review what you learned in this module. The correct answers and solutions are found in the Module Self-Check Answer Key.

1. Which two situations can result in network congestion? (Choose two.) (Source: Introducing MPLS Traffic Engineering Components)

   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 the offered load.
   D. Traffic streams are inefficiently mapped onto available resources.

2. When you are using TE with a Layer 2 overlay model, which two options transport traffic across a network? (Choose two.) (Source: Introducing MPLS Traffic Engineering Components)

   A. DVCs
   B. PVCs
   C. RVCs
   D. SVCs

3. When you are using a traffic-engineered Layer 3 model, which two of the following are limitations? (Choose two.) (Source: Introducing MPLS Traffic Engineering Components)

   A. using static routes
   B. using policy routing
   C. unavailability of explicit routing
   D. path computation that is based on just the IGP metric

https://t.me/learningnets

4. A set of data flows that share some common feature, attribute, or requirement is called _____. (Source: Introducing MPLS Traffic Engineering Components)

   A. a static route
   B. a policy route
   C. a TE tunnel
   D. TE LSP

5. Which two options can be used to advertise a traffic tunnel so that it will appear in the IP routing table? (Choose two.) (Source: Introducing MPLS Traffic Engineering Components)

   A. the autoroute feature
   B. a TE mapping statement
   C. static routes
   D. the **mpls traffic-eng tunnels** command

6. What is the role of RSVP in an MPLS TE implementation? (Source: Introducing MPLS Traffic Engineering Components)

   A. It identifies the best path for the tunnel.
   B. It reserves the bandwidth that is required by the tunnel.
   C. It performs the CBR calculations for the tunnel setup.
   D. It assigns the label for the MPLS LSP.

7. Which three options affect how a path is set up? (Choose three.) (Source: Introducing MPLS Traffic Engineering Components)

   A. priority
   B. bandwidth
   C. affinity attributes
   D. MPLS label stack
   E. MPLS EXP bit

8. Which statement is correct concerning the RSVP sessions in an MPLS-TE application? (Source: Introducing MPLS Traffic Engineering Components)

   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, and another between the tunnel end routers.

9. Which RSVP message.invokes admission control? (Source: Introducing MPLS Traffic Engineering Components)

   A. PATH
   B. RESV
   C. RESV-ERR
   D. PATH-ERR

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

10. If there is a network failure, traffic tunnels are rerouted by _____. (Source: Introducing MPLS Traffic Engineering Components)

    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

11. During path reoptimization, which router first attempts to identify a better LSP? (Source: Introducing MPLS Traffic Engineering Components)

    A. the headend router
    B. any router that has identified that it has new resources available
    C. the tunnel end router

12. Which option solves the problem of setting up two tunnels and having the resources counted twice? (Source: Introducing MPLS Traffic Engineering Components)

    A. path reuse
    B. path monitoring
    C. path rerouting
    D. path reoptimization

13. Which method is used to calculate the LSP? (Source: Introducing MPLS Traffic Engineering Components)

    A. CBR
    B. DUAL algorithm
    C. SPF algorithm
    D. No calculation used.

14. Do statically configured destinations list their tunnels in the routing table? (Source: Introducing MPLS Traffic Engineering Components)

    A. Yes, they are listed as incoming interfaces.
    B. Yes, they are listed as loopback interfaces.
    C. Yes, they are listed as outgoing interfaces.
    D. No, statically configured tunnels are not listed in the routing table.

15. Which two terms are used when the IGP metric is modified? (Choose two.) (Source: Introducing MPLS Traffic Engineering Components)

    A. absolute
    B. negative
    C. positive
    D. relative

https://t.me/learningnets

16. Which feature enables headend routers to see the MPLS-TE tunnel as a directly connected interface? (Source: Introducing MPLS Traffic Engineering Components)

    A. autoroute
    B. forwarding adjacency
    C. Autobandwidth
    D. Fast Reroute

17. The IGP metric is 50; the tunnel metric has been set to a relative +2. Each path contains six routers. Which path will be used for routing? (Source: Introducing MPLS Traffic Engineering Components)

    A. The IGP path will be used.
    B. The LSP path will be used.
    C. Both paths will be used, and load balancing will be implemented.
    D. It is not possible to answer from the information that has been provided.

18. Which mechanism enables the announcement of established tunnels via IGP to all nodes within an area? (Source: Introducing MPLS Traffic Engineering Components)

    A. forwarding neighborship
    B. autoroute
    C. forwarding adjacency
    D. autoroute announce

19. Which option is *not* a component of the maximum allocation multiplier attribute? (Source: Implementing MPLS TE)

    A. maximum bandwidth
    B. unreserved bandwidth
    C. minimum available bandwidth
    D. maximum reservable bandwidth

20. Which two options about the tunnel resource class affinity mask are true? (Choose two.) (Source: Implementing MPLS TE)

    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.

21. Which two protocols can be used to propagate MPLS-TE link attributes? (Choose two.) (Source: Implementing MPLS TE)

    A. BGP
    B. OSPF
    C. EIGRP
    D. IS-IS

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

22. In the case of a tie after CBR has been run, which two values are used to break the tie? (Choose two.) (Source: Implementing MPLS TE)

    A. larger hop count
    B. smallest hop count
    C. highest minimum bandwidth
    D. highest maximum bandwidth

23. Which command enables MPLS-TE in an OSPF implementation on a Cisco IOS Software device? (Source: Implementing MPLS TE)

    A. **mpls-te enable**
    B. **mpls traffic-eng**
    C. **metric-style wide**
    D. **mpls traffic-eng area**

24. Which Cisco IOS XR Software command that is used to instruct the IGP to use the tunnel in its shortest-path first or next-hop calculation? (Source: Implementing MPLS TE)

    A. **autoroute announce**
    B. **announce autoroute**
    C. **auto-route**
    D. **auto-bandwidth**

25. The constraint-based path computation selects the path that the dynamic traffic tunnel will take, based on the administrative weight (TE cost). To which metric is by default TE cost equal? (Source: Implementing MPLS TE)

    A. BGP metric
    B. IGP link metric (cost)
    C. relative metric
    D. absolute metric

26. With the autoroute feature enabled, the traffic tunnel can do which two things? (Choose two.) (Source: Implementing MPLS TE)

    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

27. Which Cisco IOS feature measures utilization averages and dynamically adjusts tunnel bandwidth reservations?

    A. autotunnel mesh
    B. forwarding adjacency
    C. autoroute
    D. automatic bandwidth

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. C, D
2. B, D
3. C, D
4. C
5. A, C
6. C
7. A, B, C
8. A
9. A
10. A
11. A
12. D
13. A
14. C
15. A, D
16. A
17. A
18. C
19. C
20. B, C
21. B, D
22. B, C
23. D

https://t.me/learningnets

24. A, B
25. A, B
26. A, B
27. D

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Glossary

AB

available bandwidth

ACL

access control list. 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 PHB

Assured Forwarding PHB

AIB

adjacency information base

ARP

Address Resolution Protocol. Internet protocol that is used to map an IP address to a MAC address. Defined in RFC 826.

AS

autonomous system. A collection of networks under a common administration sharing a common routing strategy. Autonomous systems are subdivided by areas. An autonomous system must be assigned a unique 16-bit number by the IANA.

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.

AToM

Any Transport over MPLS

BFD

Bidirectional Forwarding Detection

BGP

Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

CBR

constraint-based routing

CBWFQ

class-based weighted fair queuing. Extends the standard WFQ functionality to provide support for user-defined traffic classes.

https://t.me/learningnets

CE

customer edge. Identifies the network devices, connected to a provider network, that are under the administrative control of the customer.

CoS

class of service. An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages. In SNA subarea routing, CoS definitions are used by subarea nodes to determine the optimal route to establish a given session. A CoS definition comprises a virtual route number and a transmission priority field. Also called ToS.

CPE

customer premises equipment. Terminating equipment, such as terminals, telephones, and modems, supplied by the telephone company, installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment residing on the customer site.

CSPF

constraint-based SPF

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.

DRP

distributed route processor

DS-TE

DiffServ-Aware Traffic Engineering

DWDM

dense wavelength division multiplexing. Optical transmission of multiple signals over closely spaced wavelengths in the 1550 nm region. (Wavelength spacings are usually 100 GHz or 200 GHz, which corresponds to 0.8 nm or 1.6 nm.)

EBGP

Exterior Border Gateway Protocol.

EF PHB

Expedited Forwarding PHB

EIGRP

Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency, and combines the advantages of link-state protocols with those of distance vector protocols.

EMS

Ethernet Multipoint Service

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

EoMPLS

Ethernet over MPLS.

ERO

explicit route object

Ethernet

Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mb/s. Ethernet is similar to the IEEE 802.3 series of standards. The most commonly used LAN technology, because its protocol is easy to understand, implement, manage, and maintain. It allows low-cost network implementations, provides extensive topological flexibility for network installation, and guarantees successful interconnection and operation of standards-compliant products, regardless of manufacturer.

EXP

experimental

FEC

forwarding equivalence class

FIB

forwarding information base

Frame Relay

Industry-standard, packet-switched data link layer protocol that handles multiple virtual circuits between connected devices.

FRR

Fast Reroute

FT

fault-tolerant

Gigabit Ethernet

Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.

Greenfield ISP

a new ISP with no major devices

GSR

gigabit switch router

HDLC

High-Level Data Link Control. Bit-oriented synchronous data link layer protocol developed by ISO. Derived from SDLC, HDLC specifies a data encapsulation method on synchronous serial links using frame characters and checksums.

https://t.me/learningnets

IBGP

Interior Border Gateway Protocol.

ICMP

Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792.

IGP

interior gateway protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.

IGRP

Interior Gateway Routing Protocol. Developed by Cisco to address the issues associated with routing in large, heterogeneous networks.

IMP-NULL

implicit null

IP

Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security. Defined in RFC 791.

IP address

32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address.

IP precedence

A 3-bit value in the TOS byte that is used for assigning precedence to IP packets.

IPoDWDM

IP over DWDM

IS-IS

Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, whereby ISs (routers) exchange routing information based on a single metric to determine network topology.

ISP

Internet service provider. Company that provides Internet access to other companies and individuals.

https://t.me/learningnets

ITU-T

International Telecommunication Union Telecommunication Standardization Sector. International body that develops worldwide standards for telecommunications technologies. The ITU-T carries out the functions of the former CCITT.

IXP

Internet exchange point

L bit

loose bit

LDP

Label Distribution Protocol. A standard protocol between MPLS-enabled routers to negotiate the labels (addresses) used to forward packets. This protocol is not supported in Cisco IOS Release 12.0. The Cisco proprietary version of this protocol is the Tag Distribution Protocol (TDP).

LDP ID

LDP identifier

LDP-CSC

LDP Carrier-Supporting-Carrier

LFD

label forwarding database

LFIB

label forwarding information base

LIB

label information base

LLQ

low latency queuing. Provides a strict priority queue mechanism to CBWFQ.

LSA

link-state advertisement. Broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables. Sometimes called an LSP.

LSD

Label Switching Database

LSP

link-state packet

LSP

label-switched path

https://t.me/learningnets

LSR

label switch router

MAC address

Standardized data link layer address that is required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are 6 bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address.

MEF

Metropolitan Ethernet Forum

MFD

Multicast forwarding table

MMF

multimode fiber. Optical fiber supporting the propagation of multiple frequencies of light.

MP

Merge Point

MP-BGP

Multiprotocol Border Gateway Protocol.

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.

MRU

maximum receive unit

MSC

modular services card

MTU

maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.

NGN

Next-Generation Network

NHOP

next-hop

NSF

Nonstop Forwarding

https://t.me/learningnets

NSR

NonStop Routing

OADM

optical add-drop multiplexer

OAM

Operations, Administration, and Maintenance. An acronym that is used to describe the tools, processes, and standards used in the operations, administration, and maintenance of systems. Often used as a general description of these activities or to describe these activities for a specific system or protocol.

OEO

optical-electrical-optical

OIR

online insertion and removal

OSI

Open Systems Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.

OSPF

Open Shortest Path First. Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing. OSPF was derived from an early version of the IS-IS protocol.

P router

provider router

PBR

policy-based routing.

PCALC

path calculation

PCM

pulse code modulation.

PE

provider edge. Identifies the network devices, under the administrative control of the provider, that connect to CE devices.

PHB

per-hop behavior. A description of the externally observable forwarding treatment that is applied at a DiffServ-compliant node to a behavior aggregate.

PHOP

previous hop

https://t.me/learningnets

PHP

penultimate hop popping

PHY

physical layer

PID

protocol ID.

PLR

Point of Local Repair

PNNI

Private Network-to-Network Interface

POP

point of presence. In OSS, a physical location where an interexchange carrier installed equipment to interconnect with a local exchange carrier (LEC).

POS

packet over SONET

PPP

Point-to-Point Protocol. Successor to 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, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: LCP and NCP.

PSB

Path State Block

PVC

permanent virtual circuit (or connection). Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.

QoS

quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

RESV message

reservation message

RIB

routing information base

RIP

Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.

https://t.me/learningnets

RIPv2

Routing Information Protocol version 2

ROADM

Reconfigurable Optical Add/Drop Multiplexer

RP

route processor

RRO

Record route object

RSB

Reservation State Bloc

RSVP

Resource Reservation Protocol. A network-control protocol that allows endpoints to request specific QoS for their data flows.

SDH

Synchronous Digital Hierarchy

SMF

single-mode fiber. Fiber-optic cabling with a narrow core that allows light to enter only at a single angle. Such cabling has higher bandwidth than multimode fiber, but requires a light source with a narrow spectral width (for example, a laser). Also called monomode fiber.

SONET

Synchronous Optical Network. A standard format for transporting a wide range of digital telecommunications services over optical fiber. SONET is characterized by standard line rates, optical interfaces, and signal formats.

SPF

shortest path first. Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms. Sometimes called Dijkstra's algorithm.

SSO

single sign-on.

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.

TDM

time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots. Bandwidth is allocated to each channel regardless of whether the station has data to transmit.

https://t.me/learningnets

TE

traffic engineering

telco

Telephone company.

TEv4

traffic engineering version 4

TLV

type, length, value.

TSP

tag-switched path

TTL

Time to Live. A mechanism that limits the lifespan or lifetime of data in a computer or network.

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.

VPIs/VCI

virtual path identifier/virtual circuit identifier

VPLS

Virtual Private LAN Service.

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.

VRF

A VPN routing/forwarding instance. A VRF 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 includes the routing information that defines a customer VPN site that is attached to a PE router.

WDM

wavelength division multiplexing. Multiple optical wavelengths can share the same transmission fiber. The spectrum occupied by each channel must be adequately separated from the others.

https://t.me/learningnets