SPEDGE

Implementing Cisco Service Provider Next-Generation Edge Network Services

Student Guide

Version 1.2

Part Number: 97-3424-01

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 U.S. 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, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

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 2
- Course Flow 3
- Your Training Curriculum 4
- Additional References 7

Module 1: VPN Technologies 9

Lesson 1: Introducing VPNs 11
- VPN Advantages 13
- VPN Terminology 14
- VPN in the Cisco IP NGN Infrastructure Layer 16
- VPN Types 17
- Overlay Layer 3 VPNs 18
- Overlay Layer 3 VPN: GRE 19
- Overlay Layer 3 VPN: DMVPN 20
- Overlay Layer 3 VPN: L2TPv3 21
- Overlay Layer 3 VPN: IPsec VPN 22
- Overlay Layer 3 VPN: SSL VPN 23
- Peer-to-Peer VPNs 24
- Peer-to-Peer VPN: ACLs (Shared Router) 25
- Peer-to-Peer VPN: Split Routing (Dedicated Router) 26
- Peer-to-Peer VPN: GET VPN 27
- MPLS VPNs 28
- Summary 29

Lesson 2: Introducing MPLS VPNs 31
- MPLS VPN Components 33
- MPLS VPN Architecture 34
- PE Router Architecture 35
- VRF Table Overview 36
- Route Propagation via Backbone 37
- Route Distinguishers 39
- Route Targets 40
- VPN Label 41
- RD, RT, and VPN Label in Layer 3 MPLS VPN 42
- MPLS VPN Routing Requirements 43
- MPLS VPN Routing Model 44
- PE Routing Tables 46
- End-to-End Routing 47
- Packet Forwarding Without a VPN Label 49
- Packet Forwarding with a VPN Label 50
- Penultimate Hop Popping 51
- VPN Label Propagation 52
- VPN Label Requirements 54
- Broken LSP Path 55
- Summary 56

Lesson 3: Module Summary 57

Lesson 4: Module Self-Check 59

Module Self-Check Answers 63
- Answer Key 63

https://t.me/learningnets

Module 2: MPLS Layer 3 VPNs 65
Lesson 1: Implementing MPLS Layer 3 VPN Backbones 67
Layer 3 MPLS VPNs 69
PE Router Routing Contexts 70
VPN-Aware Routing Protocols 71
VRF Table 72
Route Propagation in EBGP and MP-IBGP 74
Route Propagation in IGP and MP-IBGP 76
VRF Configuration Overview (Cisco IOS/XE/XR) 78
Configuring RD 80
Configuring Export and Import RTs 81
VPN Identifier 83
VRF Configuration in Cisco IOS/XE/XR 85
Cisco IOS/XE VRF Common Policies 86
Conditional Route Target Export and Import 87
Migration from Old to New Style VRF CLI 88
Multi-VRF Customer Edge 90
VRF-Lite: VRF and Interface Configuration 91
VRF-Lite: PE-CE Connections Using BGP 92
MP-BGP Configurations 93
MP-BGP Configuration Tasks 94
BGP Address Families 95
BGP Neighbors 97
MP-BGP Community Propagation 99
MP-BGP Configuration in Cisco IOS/XE/XR 100
Disabling IPv4 Route Exchange 101
VPN Label Allocation in Cisco IOS/XE 103
VPN Label Allocation in Cisco IOS XR 104
Summary 105
Lesson 2: Deploying Basic Routing in MPLS VPNs 107
PE-CE Routing 108
Configuring BGP VRF Routing Context 109
Configuring Static PE-CE Routing 110
Configuring PE-CE RIPv2 Routing 112
Configuring PE-CE EIGRP Routing 115
Reinjecting Routes into Core by Multihomed Sites 118
Implementing EIGRP SOO for Loop Prevention 119
Route Filtering 121
Monitoring Layer 3 MPLS VPNs 123
Summary 127
Lesson 3: Deploying OSPF and BGP in MPLS VPNs 129
OSPF as PE-CE Routing Protocol 131
OSPF-BGP Redistribution Issue 132
Caveats of OSPF in MPLS VPN Routing 133
OSPF Superbackbone 134
OSPF Superbackbone Route Propagation 135
OSPF Superbackbone Implementation 136
OSPFv3 as PE-CE Routing Protocol 138
PE-CE OSPF Configuration Procedure 139
Configuring PE-CE OSPFv2 in Cisco IOS XR 140
Configuring PE-CE OSPFv2 in Cisco IOS/XE 141
OSPFv2 PE-CE Routing Example 142
OSPFv3 PE-CE Routing Example 143

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Risk of Routing Loops with BGP and OSPF ... 144
OSPF Down Bit ... 145
Packet Forwarding Without the OSPF Down Bit ... 146
Packet Forwarding with the OSPF Down Bit ... 147
OSPF Sham Link Overview ... 148
Configuring OSPF Sham Link ... 150
BGP as the PE-CE Routing Protocol ... 151
Limiting the Number of Routes in a VRF ... 153
Reusing the Same AS Number for VPN Sites ... 157
Solution: AS-Override ... 158
Example of the Solution: AS-Override ... 159
BGP Exchange with Multihomed Sites ... 160
Solution: allowas-in and Site of Origin ... 161
BGP Filtering Example (Cisco IOS XR) ... 162
Example of BGP Filtering (Cisco IOS/XE) ... 163
Troubleshooting MPLS VPNs ... 164
Summary ... 171
Lesson 4: Module Summary ... 173
References ... 173
Lesson 5: Module Self-Check ... 175
Module Self-Check Answers ... 177
Answer Key ... 177
Module 3: Special Connectivity in MPLS Layer 3 VPNs ... 179
Lesson 1: Implementing Special Connectivity in MPLS VPNs ... 181
Cisco IP NGN Infrastructure Layer ... 183
Overlapping VPNs ... 184
Overlapping VPNs Use Cases ... 185
Routing in Overlapping VPNs ... 186
Packet Flow in Overlapping VPNs ... 187
Overlapping VPNs Configuration Procedure ... 189
VRF Configuration in Overlapping VPNs ... 190
BGP Configuration in Overlapping VPNs ... 192
Central Services VPNs ... 195
Routing in Central Services VPNs ... 196
Packet Flow in Central Services VPNs ... 197
Central Services VPN Configuration Procedure ... 198
Central Services VPN Configuration Example ... 199
Selective Import and Export (Cisco IOS XR) ... 200
Selective Import (Cisco IOS XR) ... 201
Selective Export (Cisco IOS XR) ... 203
Managed CE Router Service ... 205
Managed CE Router Service Requirements ... 206
Managed CE Router Service Configuration ... 207
Summary ... 208
Lesson 2: Implementing Internet Access in MPLS VPNs ... 209
Internet Connectivity Scenarios ... 211
Classical Internet Access from VPN ... 212
Multisite Internet Access ... 213
Wholesale Internet Access ... 214
Service Provider Shared Backbone ... 215
Internet Access Design Models ... 216
Internet Access via Global Routing: Classic Topology ... 217
Internet Access via Global Routing: Separate Interfaces ... 218

https://t.me/learningnets

Internet Access via Global Routing: Static Routes ... 219
Internet Access via Global Routing with BGP ... 221
Global Internet Access with VRF-Aware Tunnels (Cisco IOS XR) ... 222
Global Internet Access with VRF-Aware Tunnels (Cisco IOS/XE) ... 223
Mutisite Internet Access via Global Routing ... 224
Evaluation of Internet Access via Global Routing ... 225
Internet Access Through Separate VPN ... 226
Features of Internet Access via Separate VPN ... 227
Internet Gateway ... 228
Internet Gateway Configuration ... 229
VPN Central Site Configuration ... 230
Separate VPN Internet Access with VRF-Aware Tunnels (Cisco IOS XR) ... 231
Separate VPN Internet Access with VRF-Aware Tunnels (Cisco IOS/XE) ... 232
Redundant Internet Access via Separate VPN ... 233
Multisite Internet Access via Separate VPN ... 234
Wholesale Internet Access via Separate VPN ... 235
Evaluation of Internet Access via Separate VPN ... 236
Summary ... 237
Lesson 3: Module Summary ... 239
Lesson 4: Module Self-Check ... 241
Module Self-Check Answers ... 243
Answer Key ... 243
Module 4: MPLS IPv6 and Interdomain Solutions ... 245
Lesson 1: Deploying IPv6 in an MPLS Environment ... 247
IPv4/IPv6 Interoperability Solutions ... 249
IPv6 in MPLS Environments ... 251
Static IPv6-in-IPv4 Tunnels on CE Routers ... 252
Using Layer 2 MPLS VPNs to Support IPv6 ... 254
6PE Architecture ... 256
6PE Control Plane ... 258
6PE Data Plane ... 259
Configuring 6PE ... 261
Verifying 6PE ... 262
6VPE Architecture ... 263
6VPE Topology ... 264
6VPE Forwarding ... 265
6VPE Configuration Procedure ... 266
Configuring 6VPE ... 267
Verifying 6VPE on PE ... 268
Verifying 6VPE on CE ... 269
Summary ... 270
Lesson 2: Introducing MPLS Interdomain Solutions ... 271
MPLS Interdomain Solutions ... 273
Inter-AS Designs ... 274
Inter-AS: Back-to-Back VRF ... 275
Back-to-Back VRF Characteristics ... 276
Inter-AS: Single-Hop MP-EBGP ... 277
Single-Hop MP-EBGP Characteristics ... 278
Inter-AS: Multihop MP-EBGP ... 279
Multihop MP-EBGP Characteristics ... 280
CSC Overview ... 281
Benefits of CSC ... 282
CSC Operation ... 283

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

CSC Design Options ... 284
CSC Using IGP and LDP ... 285
Data Flow in CSC (IGP/LDP) ... 286
CSC (IGP/LDP) Configuration Overview ... 287
Configure CSC (LDP): MP-IBGP on CSC-PE ... 288
Configure CSC (IGP/LDP): MPLS LDP in PE VRF ... 289
Configure CSC (IGP/LDP): MP-IBGP on PE ... 290
Verify LDP Session in PE VRF ... 291
Verify VPN Label for CSC-PE Loopback: PE ... 292
Verify Label for CSC-PE Loopback: CSC-CE ... 293
Verify End Customer VPN Label: CSC-PE ... 294
CSC Using MP-EBGP ... 295
CSC Data Flow (MP-EBGP) ... 296
Configure CSC (MP-EBGP): CSC-CE ... 297
Configure CSC (MP-EBGP): PE ... 298
Verify CSC (MP-EBGP): MP-EBGP Capabilities ... 299
Verify CSC (MP-EBGP): Egress PE ... 300
Verify CSC (MP-EBGP): Ingress PE ... 301
Verify CSC (MP-EBGP): CSC-CE ... 302
Verify CSC (MP-EBGP): CSC-PE ... 303
Verify CSC (MP-EBGP): Customer Traffic ... 304
Summary ... 306
Lesson 3: Module Summary ... 307
Lesson 4: Module Self-Check ... 309
Module Self-Check Answers ... 311
Answer Key ... 311
Module 5: Layer 2 VPNs ... 313
Lesson 1: Introducing Layer 2 VPNs ... 315
Layer 2 VPNs ... 316
Layer 2 VPN Benefits ... 317
Layer 2 VPN Models ... 318
Layer 2 VPN Architecture ... 319
Layer 2 VPN Control Plane ... 321
Layer 2 VPN Data Plane: L2TPv3 ... 322
Layer 2 VPN Data Plane: MPLS ... 323
Layer 2 VPN Service Types ... 324
AToM Example: EoMPLS ... 325
AToM Interworking ... 327
AToM Interworking Modes ... 328
MEF Service Types ... 329
Summary ... 331
Lesson 2: Implementing AToM ... 333
AToM Operation ... 335
AToM VC Creation Process ... 336
AToM Forwarding ... 337
AToM Signaling ... 339
Virtual Circuit Connectivity Verification ... 340
AToM MTU Considerations ... 342
AToM Control Word ... 343
Inter-AS AToM ... 344
EoMPLS Packet Format ... 345
EoMPLS Redundancy ... 346
One-Way EoMPLS Redundancy ... 347

https://t.me/learningnets

Two-Way EoMPLS Redundancy 348
Cisco EVC Infrastructure Overview 350
AToM Configuration Procedure 358
EoMPLS Interconnecting VLAN 11 359
EoMPLS Interconnecting Untagged Frames 360
EoMPLS Rewriting VLAN Tags 361
EoMPLS Basic Verification 363
EoMPLS Detailed Verification In Cisco IOS XR 364
EoMPLS Detailed Verification In Cisco IOS and IOS XE 365
Monitoring Attachment Circuits 366
Summary 367
Lesson 3: Implementing VPLS 369
Virtual Private LAN Service 371
VPLS Characteristics 373
VPLS Traffic Forwarding 374
VPLS Topology: PE View 375
Virtual Forwarding Instance 376
Routed VPLS 377
VPLS Configuration Procedure 378
VPLS Interconnecting VLAN 111 379
VPLS With VLAN Rewrite 380
Configuring Routed VPLS 381
Verifying VPLS in Cisco IOS XR 382
Monitoring Bridge Domain in Cisco IOS XE 385
Monitoring VFI in Cisco IOS XE 386
Monitoring Attachment Circuits 387
Summary 388
Lesson 4: Enhancing Layer 2 VPN Scalability 389
VPLS Architectures 391
VPLS vs. H-VPLS 392
H-VPLS Edge: 802.1ad and QinQ 393
H-VPLS Edge: MPLS 394
IEEE Carrier Ethernet Standards 396
VPLS and H-VPLS Evaluation 397
VPLS Neighbor Discovery and Signaling 398
H-VPLS and Autodiscovery Configuration 400
H-VPLS with QinQ: Switches 402
H-VPLS with QinQ: S-VLAN Rewrite 403
H-VPLS QinQ: S-VLAN and C-VLAN Rewrite 404
H-VPLS with 802.1ad: Switches 405
H-VPLS with 802.1ad: S-VLAN Rewrite 406
H-VPLS 802.1ad: S-VLAN and C-VLAN Rewrite 407
H-VPLS with BGP Autodiscovery 408
Verify H-VPLS Bridge Domain 410
Verify Attachment Circuits in H-VPLS 411
Verify Autodiscovery: BGP Capabilities 412
Verify Autodiscovery: Layer 2 VPN VPLS Database 413
Verify Autodiscovery: Layer 2 VPN VPLS Attributes 414
Summary 415
Lesson 5: Module Summary 417
References 418
Lesson 6: Module Self-Check 419
Module Self-Check Answers 421
Answer Key 421

https://t.me/learningnets

Glossary..... 423

https://t.me/learningnets

https://t.me/learningnets

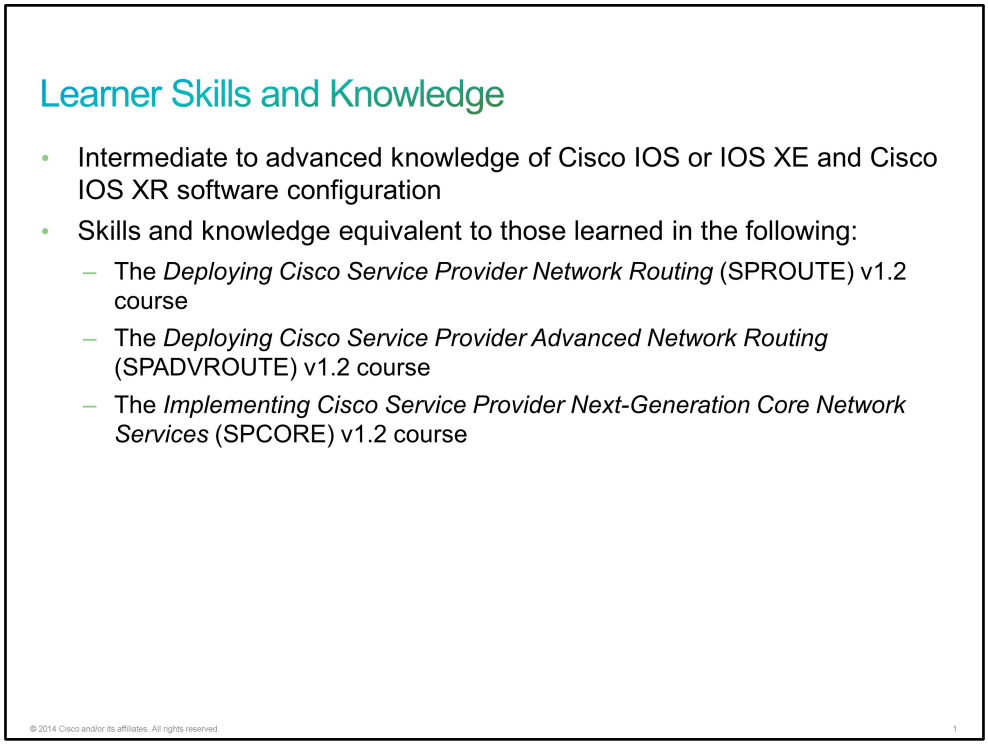
Course Introduction

Overview

This subtopic provides an overview of the course.

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

Course Goal and Objectives

This topic describes the course goal and objectives.

Upon completing this course, you will be able to:

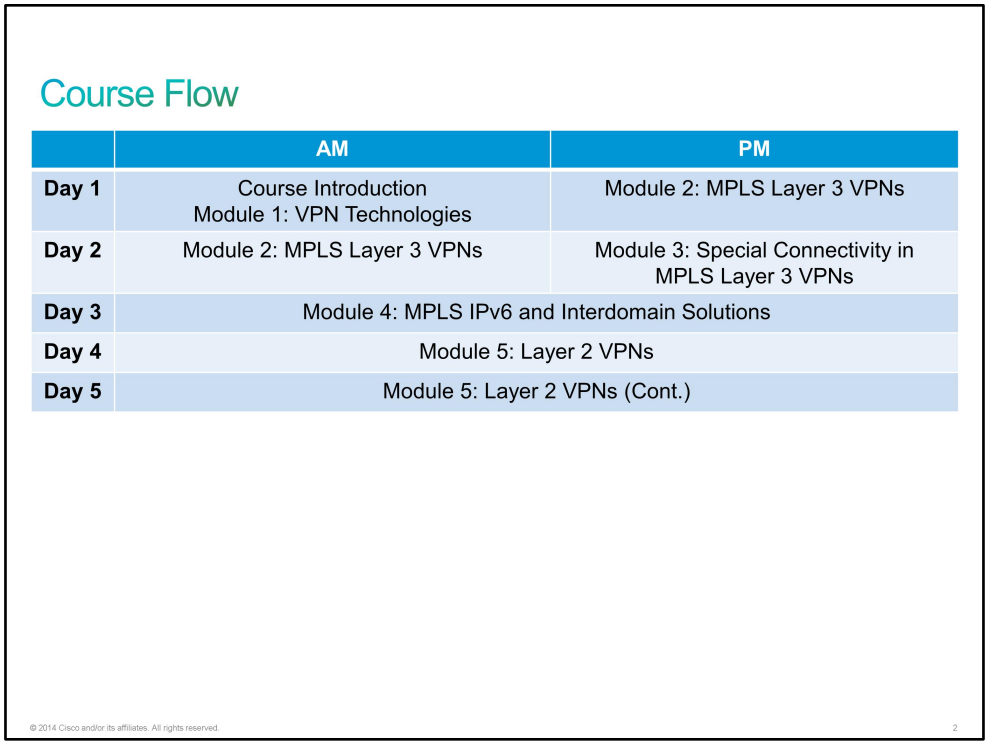
- Describe the VPN technologies that are used in the service provider environment and the MPLS VPN peer-to-peer architecture
- Describe the implementation steps that are needed to provide MPLS Layer 3 VPN service in the service provider network
- Describe how the MPLS Layer 3 VPN model can be used to implement managed services and Internet access
- Describe MPLS solutions for IPv6 and interdomain communication.
- Describe Layer 2 VPNs and Ethernet services

 © 2014 Cisco Systems, Inc.

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.

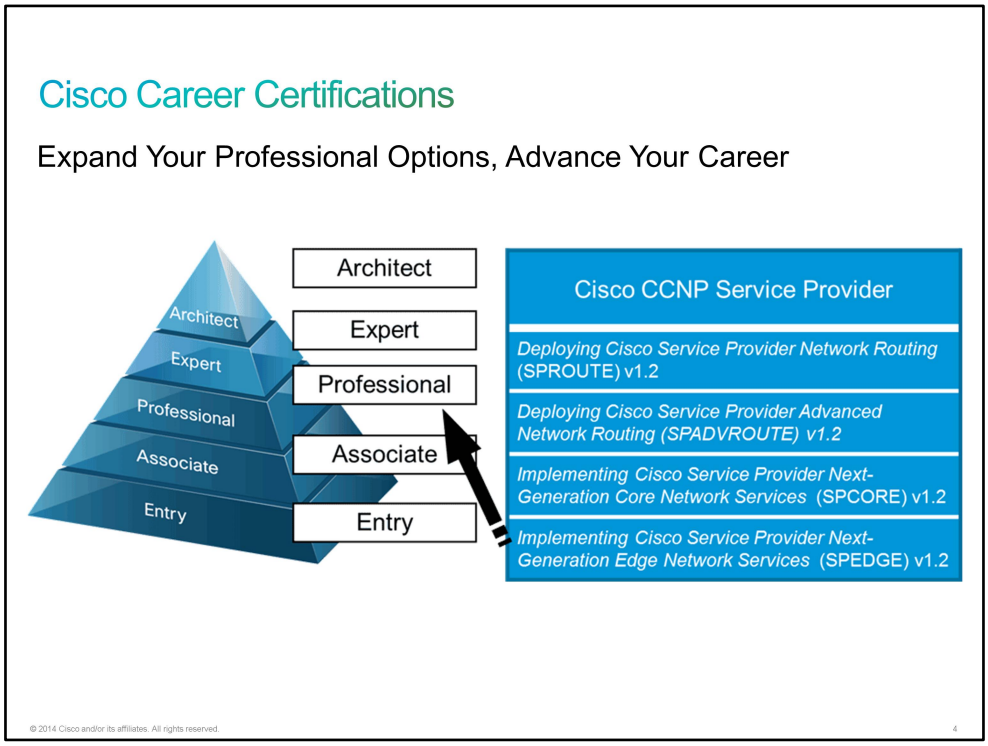


You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, CCNP® Security and CCNP® Voice, or CCSP™). It provides a gathering place for Cisco certified professionals 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.

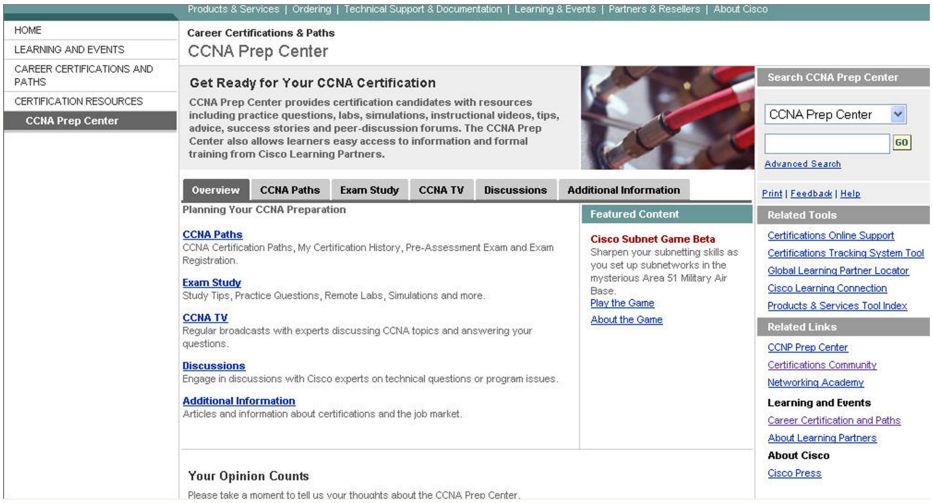


For more information on certifications, visit http://www.cisco.com/go/certifications.

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

CCNA Prep Center



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

Additional information is available at http://www.cisco.com/go/prepcenter.

Learner Introductions

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

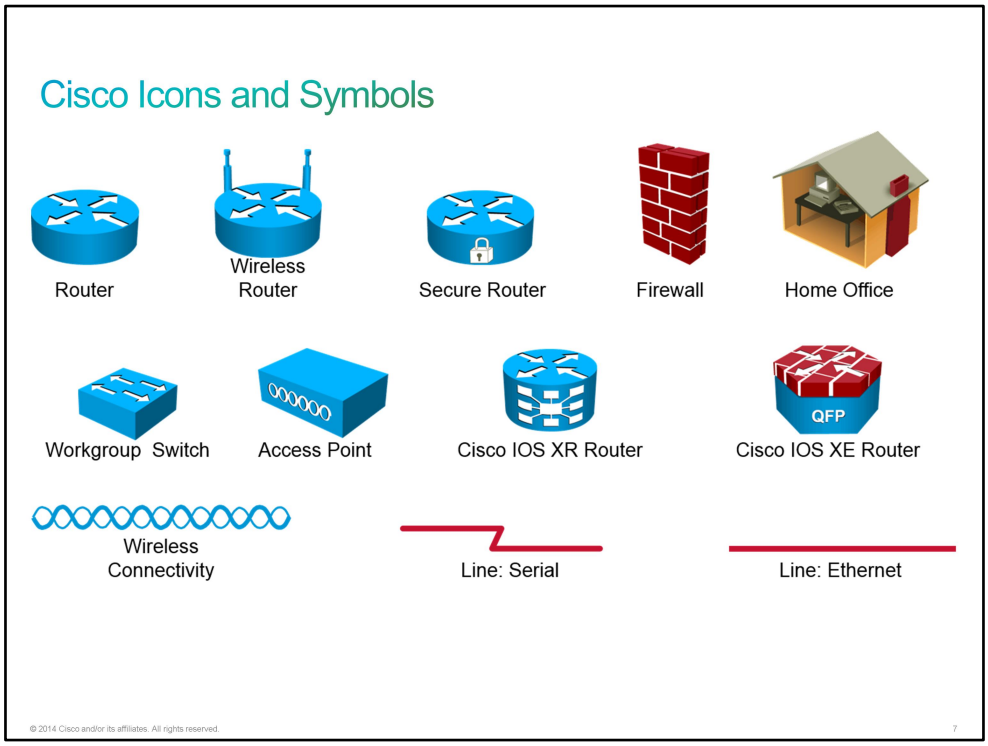


© 2014 Cisco and/or its affiliates. All rights reserved.

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/Internetworking_Terms_and_Acronyms_(ITA)_Guide.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 1: VPN Technologies

Introduction

This module introduces VPNs and two major VPN design options: the overlay VPN and the peer-to-peer VPN. The module also introduces VPN terminology and topologies, and describes MPLS VPN architecture and operations. This module details various CE-PE routing options and BGP extensions (route targets and extended community attributes) that allow IBGP to transport customer routes over a provider network. The MPLS VPN forwarding model is also covered, along with how it integrates with core routing protocols.

Upon completing this module, you will be able to:

- Explain the concept of VPNs and the VPN terminology
- Explain the MPLS VPN architecture, route information propagation, RDs, RTs, and virtual routing tables

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Introducing VPNs

Overview

This lesson introduces the main characteristics of VPNs. This lesson explains the concept of VPNs, the count advantages of VPNs, and the VPN terminology that is also used by the MPLS VPN architecture. The lesson also explains the differences between the overlay and peer-to-peer VPN models, how they are implemented, and the benefits and drawbacks of each implementation.

Understanding the background of VPNs is important because you should be able to determine when an organization might need a VPN and be able to explain how MPLS VPNs can help save time and money. Understanding the different types of VPNs will enable you to recognize where they would be best used in their associated networks.

Upon completing this lesson, you will be able to:

- Describe the advantages of VPNs
- Describe VPN terminology
- Describe the Cisco IP NGN infrastructure layer
- Describe two VPN implementation models (overlay and peer-to-peer)
- List the different VPN types
- Describe the overlay Layer 3 VPNs
- Describe the GRE VPNs
- Describe the DMVPNs
- Describe the L2TPv3 VPNs
- Describe the IPsec VPNs
- Describe the SSL VPNs
- Describe the peer-to-peer VPNs
- Describe the peer-to-peer VPNs using ACLs (shared router)

https://t.me/learningnets

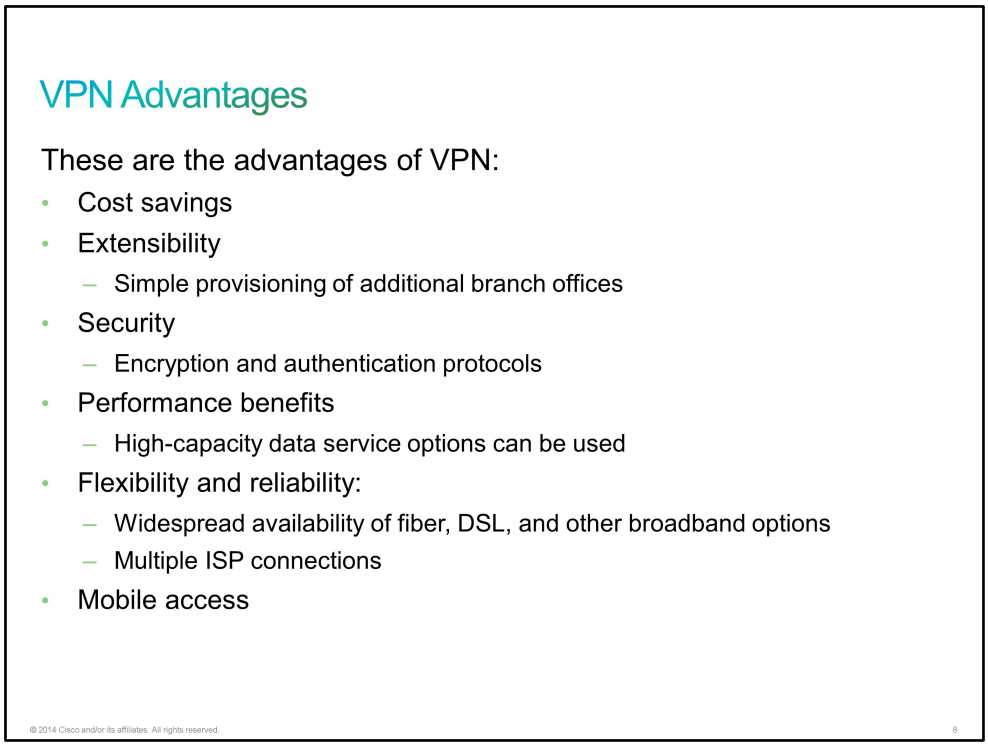
- Describe peer-to-peer VPNs using split routing (dedicated router)
- Describe GET VPNs
- Describe MPLS VPNs

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

VPN Advantages

This topic describes the advantages of VPN.



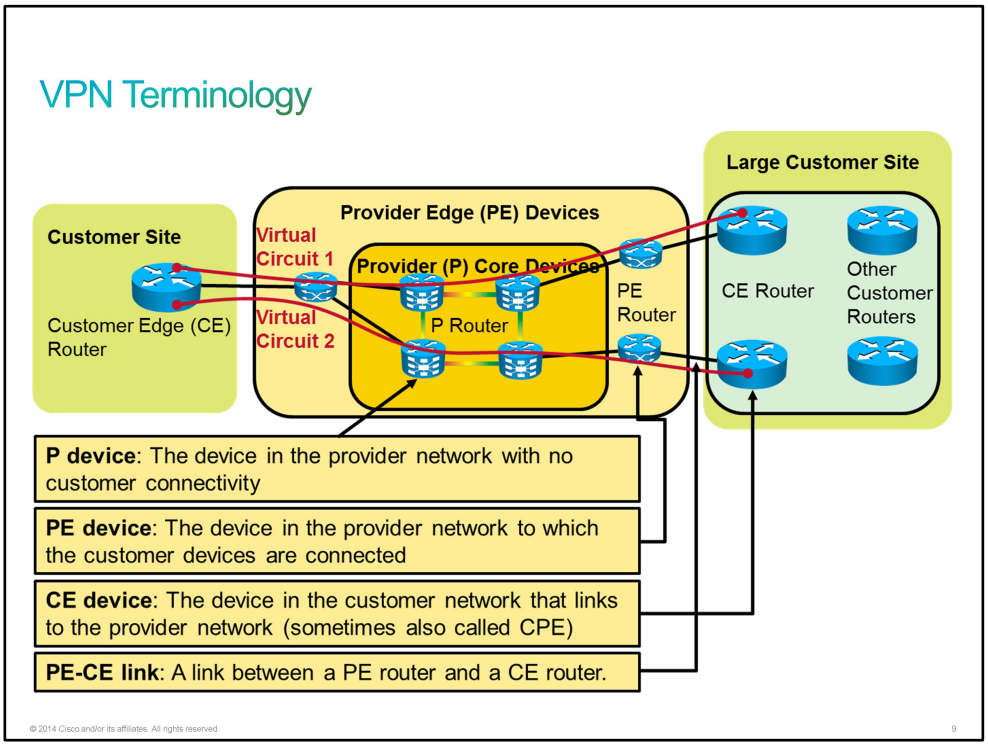
VPNs give an organization the advantage of creating secure channels of communication while at the same time reducing costs, improving security, increasing performance, and providing greater access to remote users:

- **Cost savings:** Dedicated circuits (leased lines) are quite expensive, so replacing leased lines with a much less expensive dedicated connection to the service provider can significantly decrease costs.
- **Scalability:** A company with two branch offices can deploy just one dedicated line to connect the two locations. If a third branch office needs to come online, two additional lines will be required to directly connect that location to the other two. However, by adding more branch offices to the network, the number of leased lines increases dramatically (four branch offices require six lines for full connectivity; five offices require ten lines, and so on). VPNs avoid this problem by simply adopting one link to ISP per branch office.
- **Improved security:** The use of encryption protocols and authentication helps secure the data that is traveling over the VPN channel.
- **Better performance:** VPNs also provide greater bandwidth, thus allowing better performance.
- **Flexibility and reliability:** The widespread availability of fiber, DSL, and other broadband options gives enterprises multiple ways to securely interconnect over a VPN. VPNs can also improve reliability by using data services from several independent ISPs at the same time (having redundant solutions).
- **Greater access to mobile users:** Many workers can work from home or spend a significant amount of time on business trips. By adopting VPN solutions, they are able to connect from anywhere to company servers to access email and data, thus increasing productivity and responsiveness.

https://t.me/learningnets

VPN Terminology

There are many conceptual models and terminologies that describe various VPN technologies and implementations.



The terminology is generic enough to cover nearly any VPN technology or implementation and is thus extremely versatile.

The major parts of an overall VPN solution are always those listed here:

- **P-network:** The common infrastructure that the service provider uses to offer VPN services to customers.
- **C-network:** The part of the overall customer network that is still exclusively under customer control. It consists of the routers at the various customer sites. The routers connecting the sites of individual customers to the service provider network are called CE routers.
- **Customer sites:** These are contiguous parts of the C-network.

A typical C-network implemented with any VPN technology would contain islands of connectivity under customer control (customer sites) connected together via the service provider infrastructure (P-network).

Here is a description of the devices that enable the overall VPN solution. The devices are named based on their position in the network:

- **P device:** Service provider devices that provide only data transport across the service provider backbone, and have no customers that are attached to them. In an MPLS implementation, these devices would be LSRs.
- **PE device:** Service provider devices to which customer devices are attached are called PE devices.

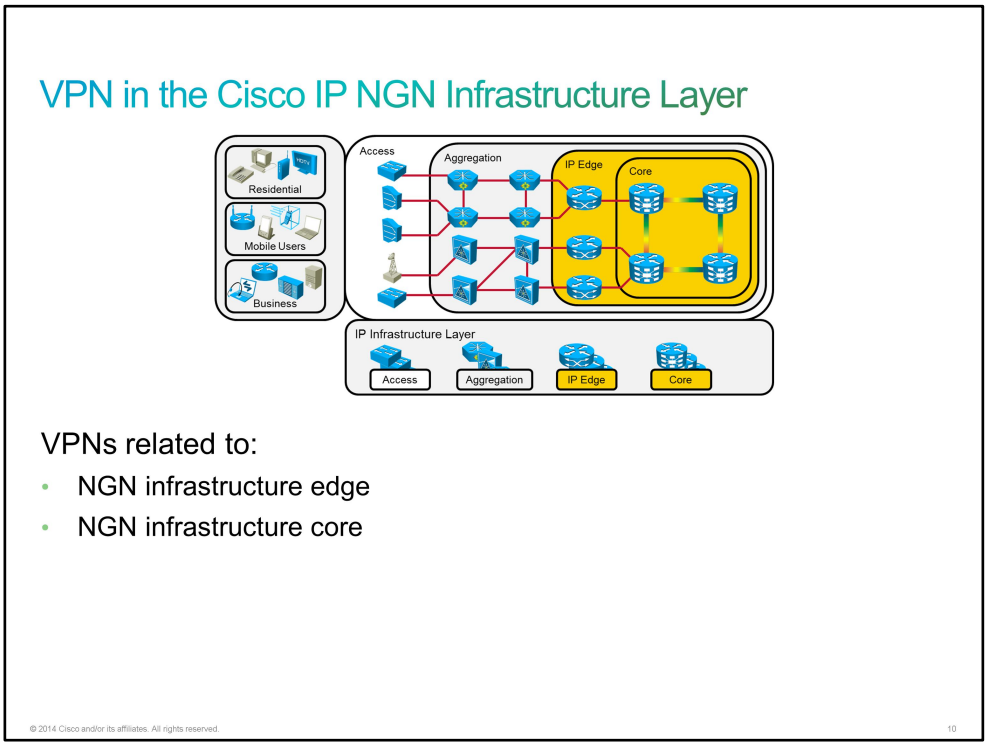
https://t.me/learningnets

- **CE device:** The customer router that connects the customer site to the service provider network is called a CE router, or CE device. This device can also be called CPE.
- **PE-CE link:** A link between a PE router and a CE router.

https://t.me/learningnets

VPN in the Cisco IP NGN Infrastructure Layer

All VPN implementation models relay on the IP edge and core parts of the IP infrastructure layer of the Cisco IP Next-Generation Network.

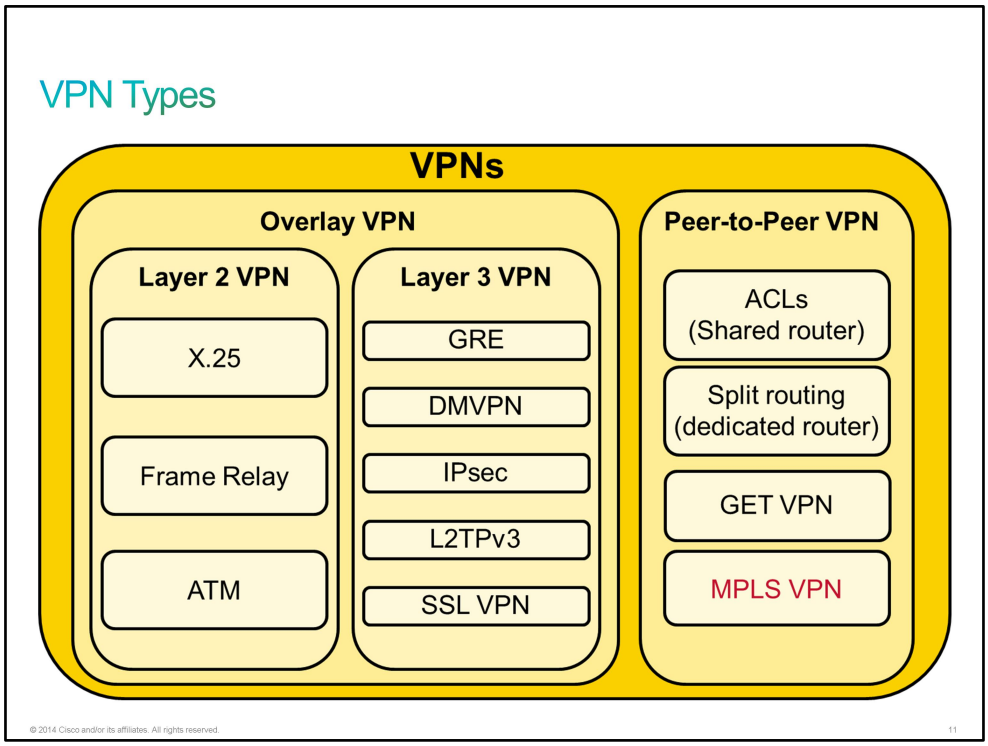


The access layer represents the customer edge devices, most often CE routers. The edge and core provide VPN services to the customer sites. The aggregation layer may play an important role in large environments with many customer connections terminated on the PE devices.

https://t.me/learningnets

VPN Types

Traditional VPN implementations were all based on the overlay model, in which the service provider provided VCs between customer sites as a replacement for dedicated point-to-point links. The overlay model had a number of drawbacks that are identified in this lesson. To overcome these drawbacks, a new model called peer-to-peer VPN was introduced. In the peer-to-peer VPN model, the service provider actively participates in customer routing.



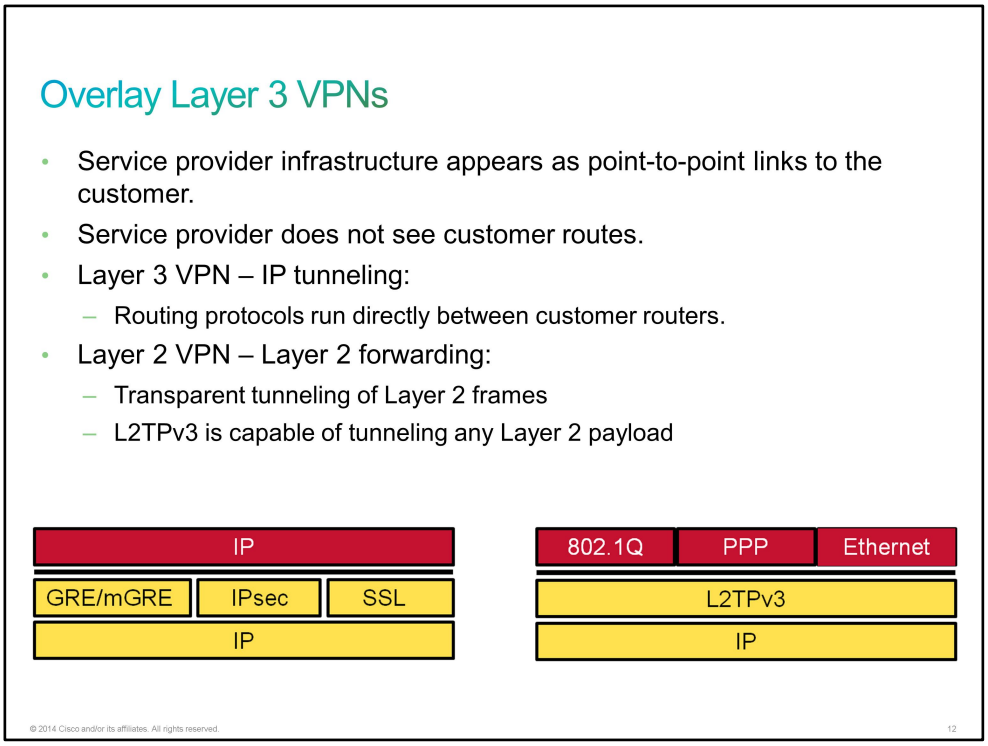
VPNs allow you to use the shared infrastructure of a service provider to implement your private networks. There are basically these two implementation models:

- Overlay VPNs:
  - Layer 2, including technologies such as X.25, Frame Relay, and ATM
  - Layer 3, including GRE, DMVPN, IPsec, SSL VPN, and L2TP
- Peer-to-peer VPNs, implemented with routers and the respective filters, separate routers per customer via GET VPN or with MPLS VPN technology tha is covered in greater detail in later lessons.

https://t.me/learningnets

Overlay Layer 3 VPNs

With the success of IP and associated technologies, some service providers started to implement pure IP backbones to offer VPN services based on IP. In other cases, customers wanted to take advantage of the low cost and universal availability of the Internet to build low-cost private networks over it.



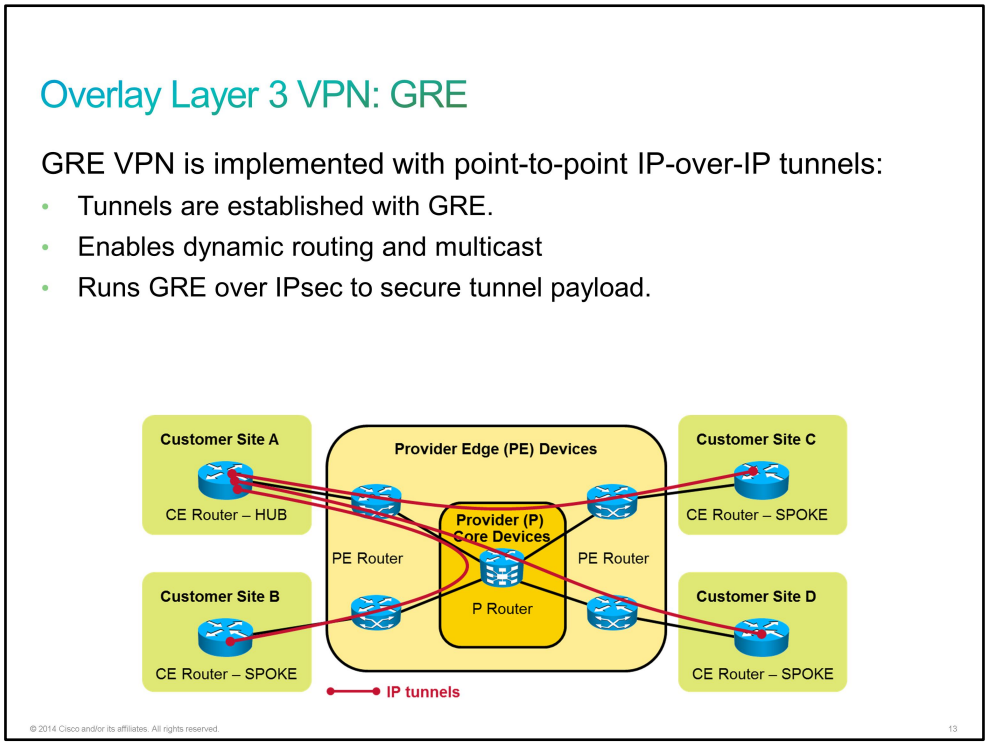
Whatever the business reasons behind it, Layer 3 VPN implementations over the IP backbone always involve tunneling—encapsulation of protocol units at a certain layer of the OSI reference model into protocol units at the same or a higher layer of the OSI model.

Two well-known tunneling technologies are IPsec and GRE. GRE is fast and simple to implement and supports multiple routed protocols, but it provides no security and is thus unsuitable for deployment over the Internet. An alternative tunneling technology is IPsec, which provides network layer authentication and optional encryption to make data transfer over the Internet secure. IPsec supports only the IP routed protocol. SSL is the latest method to make authentication and encryption of data transferred over the Internet secure. It is a remote access solution that replaces IPsec clients and is firewall-friendly (uses SSL as the transport).

L2TPv3 is capable of tunneling any Layer 2 payload over L2TP.

https://t.me/learningnets

Overlay Layer 3 VPN: GRE



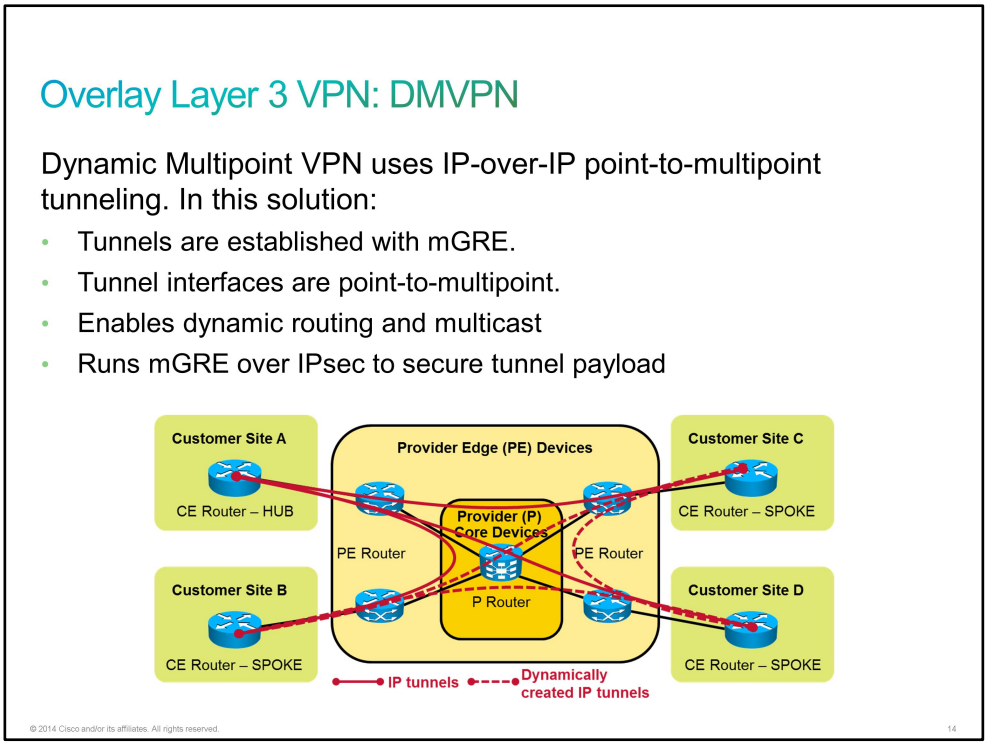
In this case, Layer 3 connectivity is provided between hub and spoke sites by using GRE point-to-point tunnels. The customer needs to connect three sites to Site A (central site, hub site) and orders connectivity between Site A (hub) and Site B (spoke), between Site A and Site C (spoke), and between Site A and Site D (spoke). The service provider implements IP connectivity over its network. Note that spoke-to-spoke traffic has to go through the hub site.

The GRE is a multiprotocol-capable transport protocol (IPv4, IPv6, MPLS, and so on) and enables dynamic routing and multicast over the tunnels.

To secure the tunnel payload, you have to run GRE over IPsec.

https://t.me/learningnets

Overlay Layer 3 VPN: DMVPN



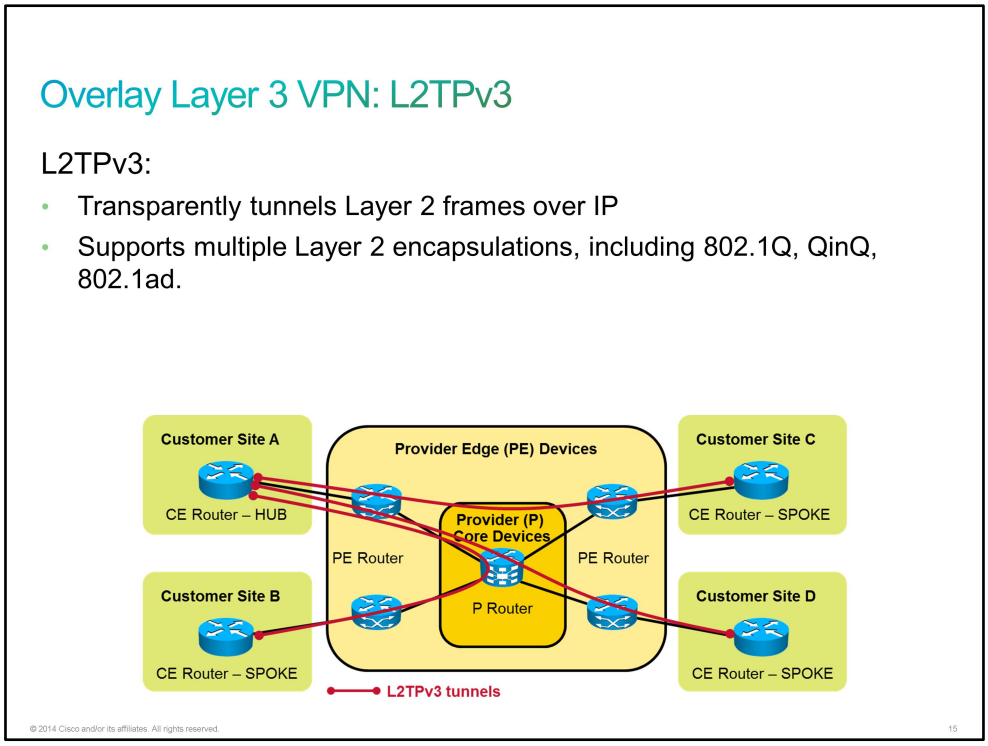
In this DMVPN scenario, point-to-mGRE tunnels are used. The customer needs to connect three sites to Site A (central site, hub site) and orders connectivity between Site A (hub) and Site B (spoke), between Site A and Site C (spoke), and between Site A and Site D (spoke). The service provider implements IP connectivity over its network. Note that in this DMVPN scenario, spoke-to-spoke traffic can flow directly by dynamically establishing GRE tunnels between spokes. To secure the tunnel payload, you have to run mGRE over IPsec.

The GRE is a multiprotocol-capable transport protocol (IPv4, IPv6, MPLS, and so on) and enables dynamic routing and multicast over the tunnels.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Overlay Layer 3 VPN: L2TPv3

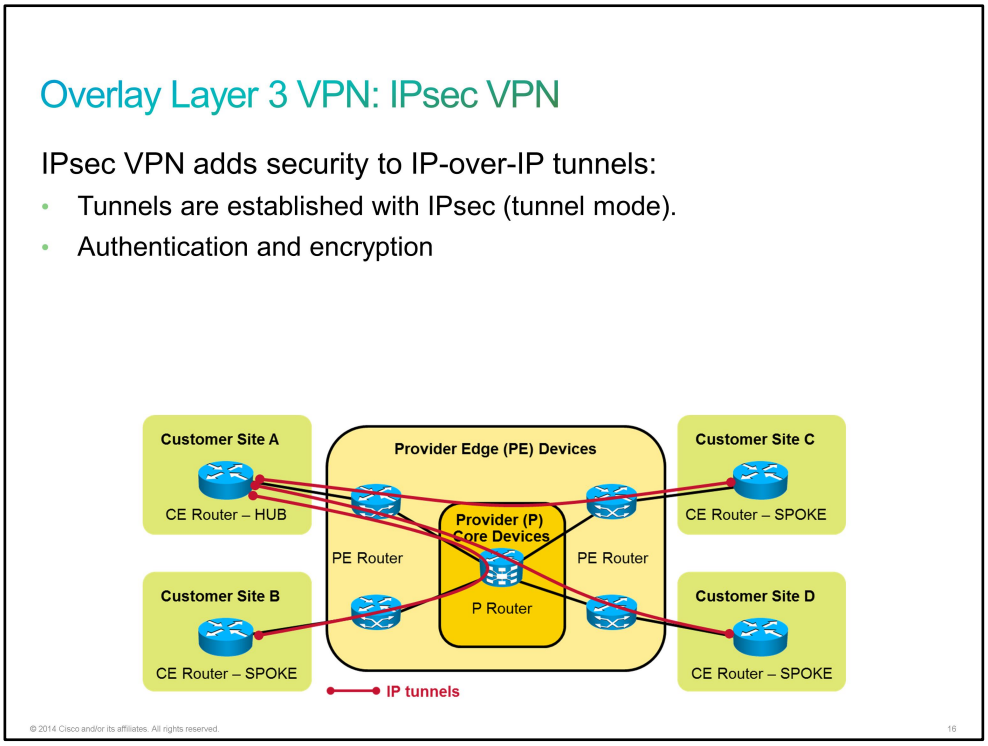


L2TPv3 is capable of tunneling any Layer 2 payload over L2TP. Specifically, L2TPv3 defines the L2TP protocol for tunneling Layer 2 payloads over an IP core network using Layer 2 VPNs. The benefits of this feature include the following:

- L2TPv3 simplifies deployment of VPNs.
- L2TPv3 does not require MPLS.
- L2TPv3 supports Layer 2 tunneling over IP for any payload.

https://t.me/learningnets

Overlay Layer 3 VPN: IPsec VPN

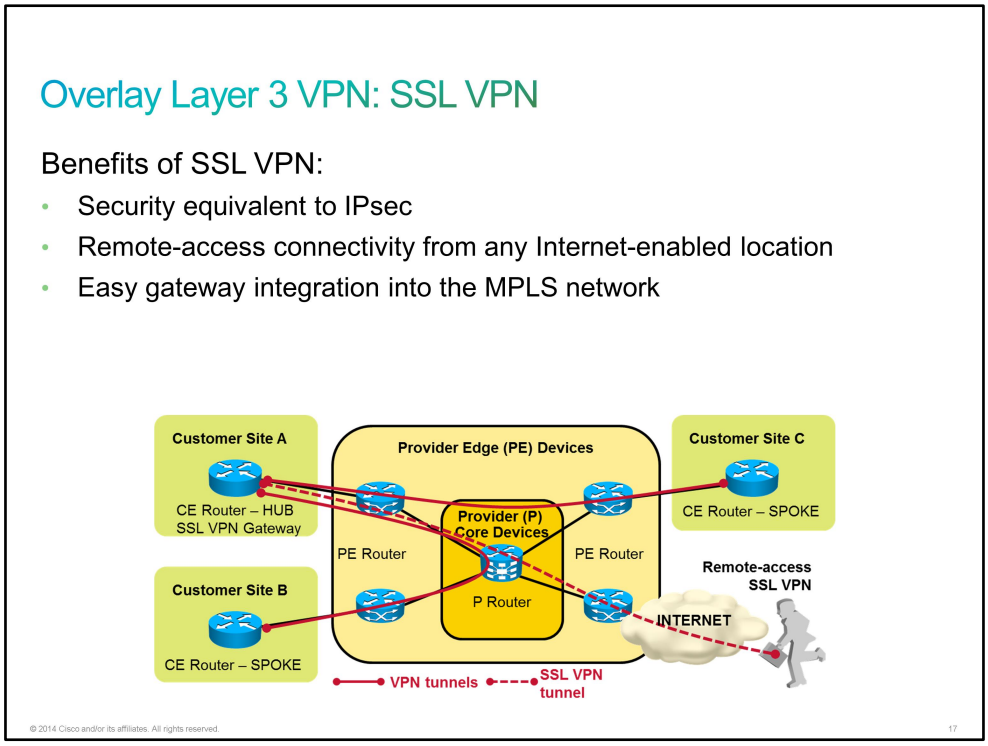


IPsec provides network layer authentication and optional encryption to make data transfer over the Internet secure. This is achieved by creating IP-over-IP tunnels and securing the payload. The limitation of the IPsec tunnels is that they do not offer multicast functionality and instead provide static routing only.

The usage of IPsec, that is, securing the payload, is usually used in securing GRE and mGRE tunnels.

https://t.me/learningnets

Overlay Layer 3 VPN: SSL VPN



SSL is the method to achieve secure authentication and encryption of the data transferred over the Internet. It is a remote access solution that replaces IPsec clients and is firewall-friendly (uses SSL as the transport). It runs in three operational models:

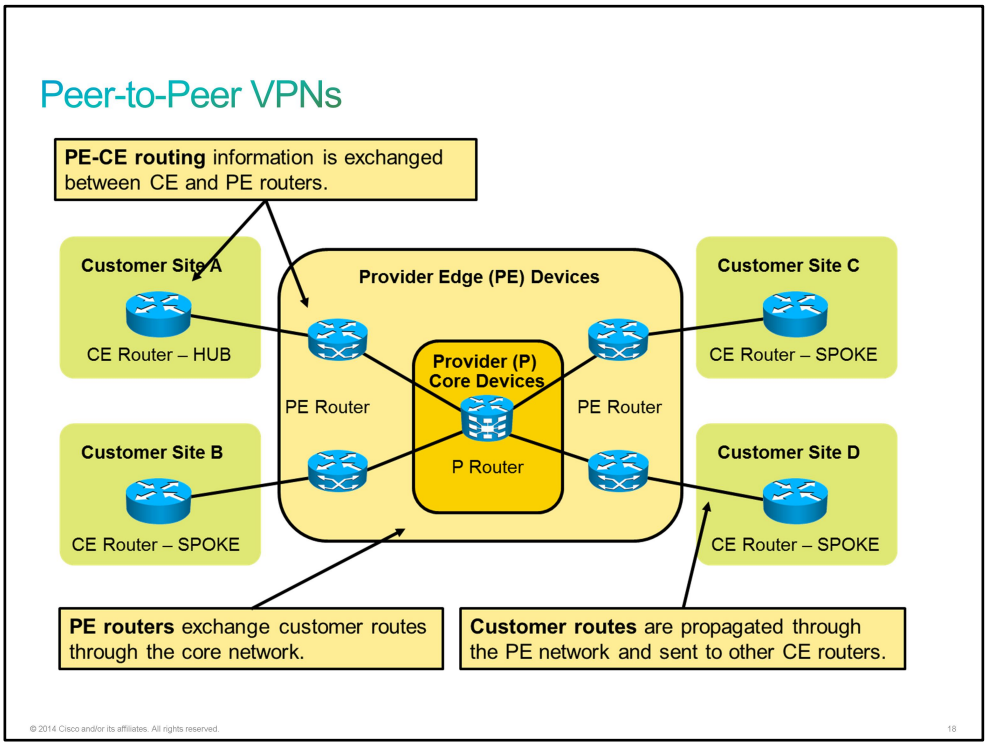
- Clientless, providing access to web servers behind the firewall
- Thin client, providing port forwarding via a Java applet
- Full tunnel with SSL VPN client

Integrating an SSL VPN gateway into an MPLS VPN network is possible.

https://t.me/learningnets

Peer-to-Peer VPNs

The point-to-point overlay VPN model has a number of drawbacks, most significantly, the need for customers to establish point-to-point links or VCs between sites.



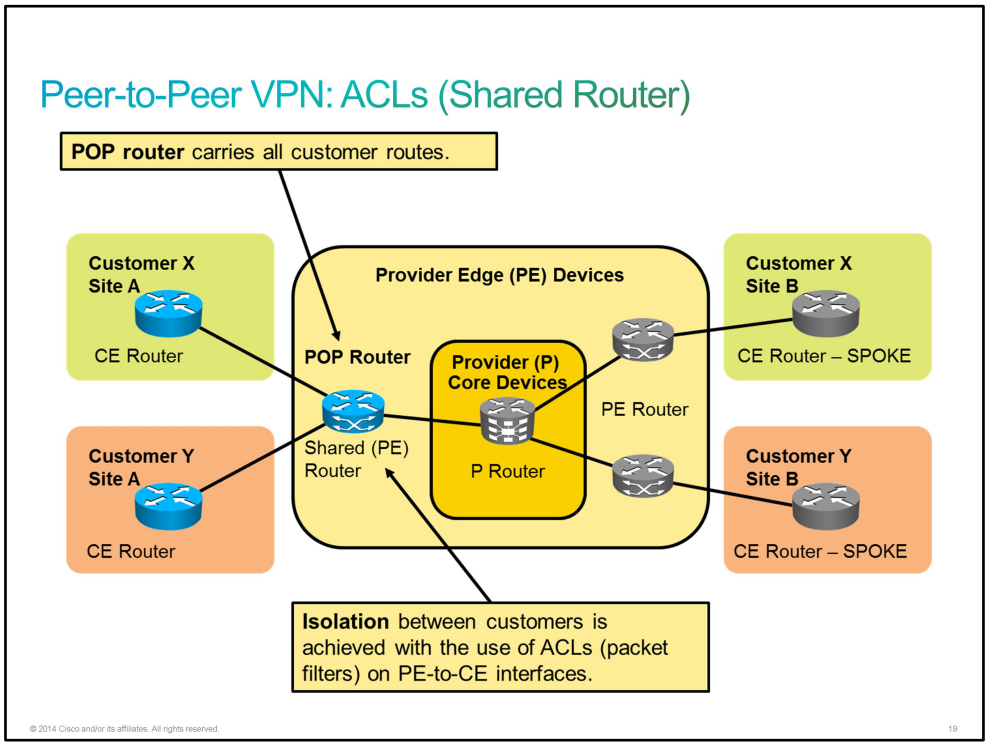
The formula to calculate how many point-to-point links or VCs are needed is ([n]*[n-1])/2, where n is the number of sites to be connected. For example, if a customer wants to have a full mesh between 10 sites, it would need 10*9/2=45 point-to-point links. This would certainly be a scalability issue.

To overcome the scalability issue and provide the customer with optimum data transport across the service provider backbone, the peer-to-peer VPN concept was introduced. Here, the service provider actively participates in customer routing, accepting customer routes, transporting those customer routes across the service provider backbone, and finally propagating them to other customer sites.

https://t.me/learningnets

Peer-to-Peer VPN: ACLs (Shared Router)

The first peer-to-peer VPN solutions appeared with the widespread deployment of IP in service provider networks. Architectures similar to that of the Internet were used to build these VPN solutions.



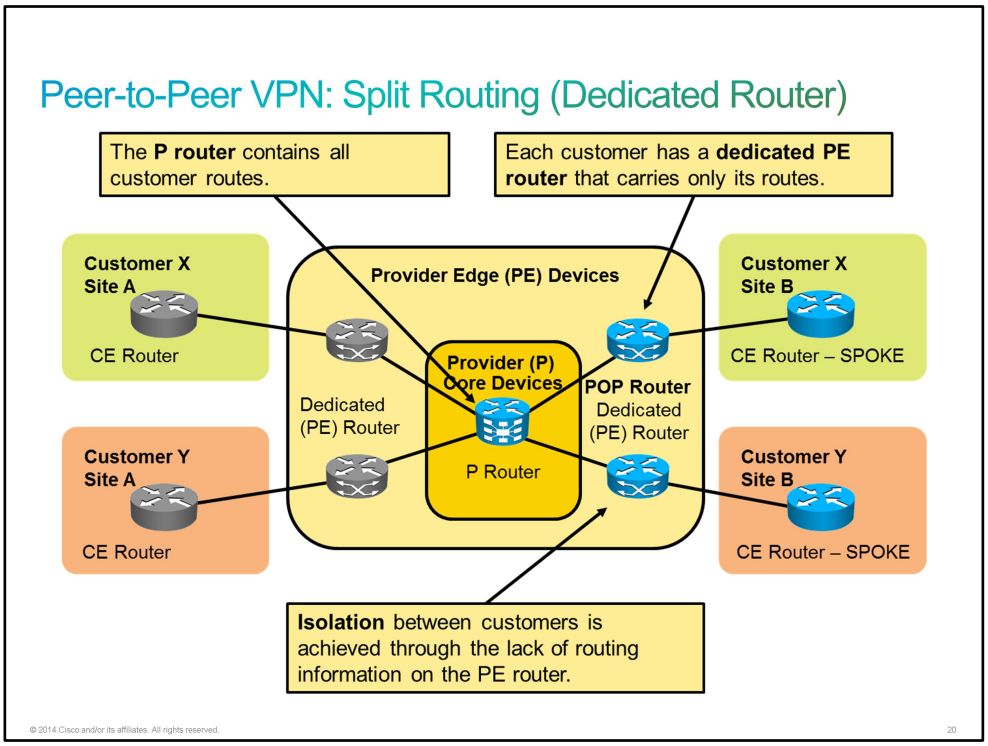
Special provisions were taken into account to transform the architecture, that was targeted toward the public backbones (Internet), into a solution in which customers would be totally isolated and be able to exchange corporate data securely.

The more common peer-to-peer VPN implementation allowed a PE router to be shared between two or more customers. ACLs—that is, packet filters were used on the shared PE routers to isolate the customers. In this implementation, it was common for the service provider to allocate a portion of its address space to each customer and manage the ACLs on the PE routers to ensure full reachability between the sites of a single customer, as well as isolation between separate customers.

https://t.me/learningnets

Peer-to-Peer VPN: Split Routing (Dedicated Router)

Maintaining ACLs is a tedious and error-prone task. Therefore, some service providers have implemented more innovative solutions based on controlled route distribution.



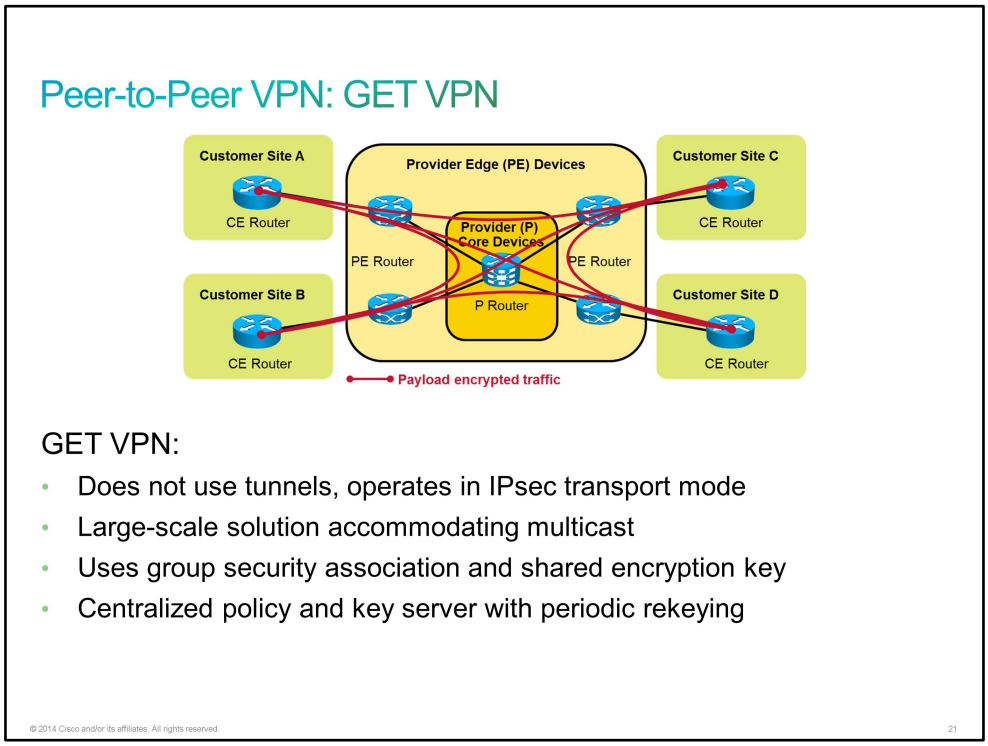
In this approach, the customer has a dedicated PE router. The core service P routers contain all customer routes, and the dedicated PE routers contain only the routes of a single customer. This approach requires a dedicated PE router per customer per POP. Customer isolation is achieved solely through lack of routing information on the PE router.

In the figure, the PE router for customer X, using route filtering between the P router and the PE routers, learns only the routes belonging to customer X, and the PE router for customer Y learns only the outes belonging to customer Y. BGP with BGP communities is usually used inside the provider backbone, because it offers the most versatile route-filtering tools.

https://t.me/learningnets

Peer-to-Peer VPN: GET VPN

GET VPN is a tunnel-less VPN technology that provides end-to-end security for network traffic in a native mode and maintains the fully meshed topology.



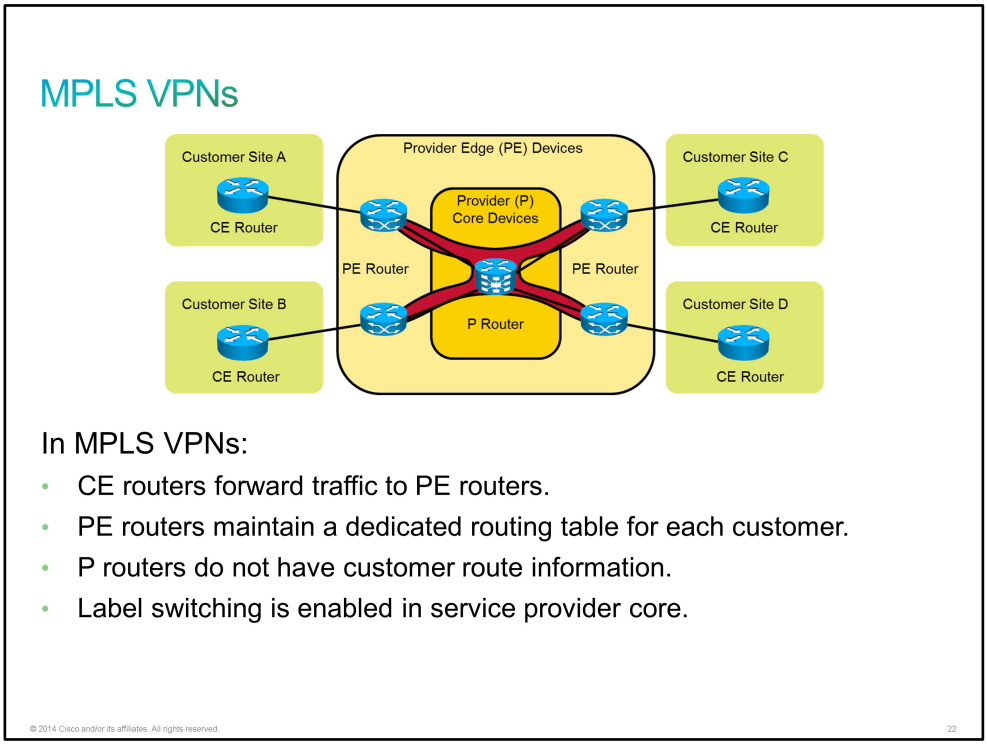
GET VPN preserves the original source and destination IP addresses information in the header of the encrypted packet for optimal routing (like transport mode IPsec). Hence, it is largely suited for an enterprise running over a private MPLS and IP-based core network. It is also better suited to encrypt multicast traffic. GET VPN uses GDOI as the keying protocol and IPsec for encryption. Some of the advantages of GET VPN are as follows:

- Provides high scalability to any meshed topology and eliminates the need for complex peer-to-peer security associations
- For MPLS networks, maintains network intelligence (such as full-mesh connectivity, natural routing path, and QoS). Grants easy membership control with centralized key servers.
- Helps ensure low latency and jitter by enabling full-time, direct communications between sites without requiring transport through a central hub.
- Allows replication of the packets after encryption. This allows the multicast traffic to be replicated at the core, thereby reducing the load and bandwidth requirement on the CPE.
- IP address preservation enables encrypted packets to carry the original source and destination IP addresses in the outer IP header rather than replacing them with tunnel endpoint addresses. This technique is known as IPsec tunnel mode with address preservation. Some of the IP header parameters are also preserved. Many network features such as routing, basic firewall, QoS, and traffic management work based on the information contained in the IP header. Since the IP header is persevered, all the network features will work as before. This eliminates many issues associated with deploying point-to-point encryption in a core network.

https://t.me/learningnets

MPLS VPNs

In the MPLS VPN model, the best features of the overlay and point-to-point models are implemented.



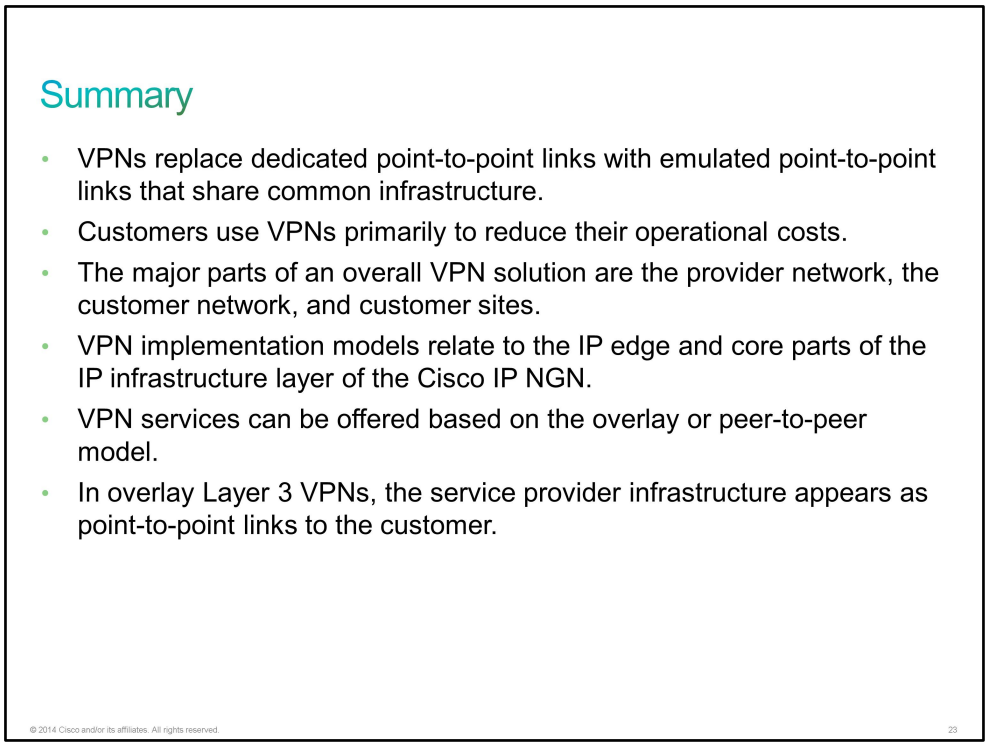
An MPLS-enabled core and edge network provides a very fast and efficient data-switching environment based on MPLS labels. PE routers exchange routing information with customer CE routers and use separate isolated routing tables for each customer. Special routing protocol contexts are used for route exchange between PE and CE routers.

Routes are then exchanged between PE devices using the MP-BGP routing algorithm. For scalability reasons, service provider core routers do not have any customer routing information. PE routers label packets with MPLS labels and P routers use these labels for fast label-switching packets.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Introducing MPLS VPNs

Overview

This lesson explains the MPLS VPN architecture. Understanding how the MPLS VPN architecture is structured, what the components of that architecture are, and how the components are used is important-. This knowledge will help later, when you begin to look at design issues and configuration parameters. The lesson offers address and routing perspectives from the customer and service provider side, and it discusses how routing tables appear on PE routers. This lesson also explains how forwarding across an MPLS VPN backbone occurs, identifies how labels get propagated, and explains the effects of summarization in the core.

Upon completing this lesson, you will be able to:

- Describe the basic MPLS VPN architecture components
- Describe the PE router architecture
- Describe the basic concepts of the VRF table
- Describe the propagation of customer routing information across the service provider network
- Describe the route distinguisher
- Describe the route target
- Describe the VPN label
- Describe the RD and RT functions in Layer 3 MPLS VPN
- Describe the end-to-end routing update flow for Layer 3 MPLS VPNs
- Describe an example of end-to-end routing update flow
- Describe the MPLS forwarding without a VPN label
- Describe the MPLS forwarding with a VPN label
- Describe the penultimate hop popping

https://t.me/learningnets

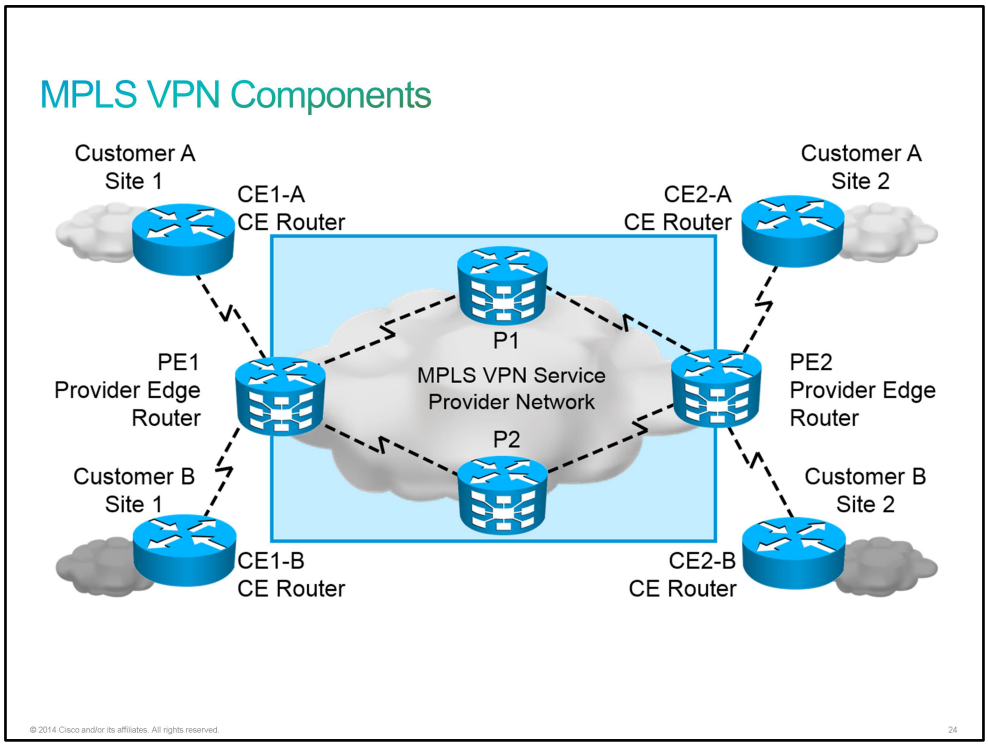
- Describe VPN label propagation between the PE routers
- Describe the problem with a broken LSP between the PE routers

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS VPN Components

The MPLS VPN domain, like the traditional VPN, consists of the customer network and the provider network. The MPLS VPN model is very similar to the dedicated PE router model in a peer-to-peer VPN implementation. However, instead of deploying a dedicated PE router per customer, customer traffic is isolated on the same PE router that provides connectivity into the service provider network for multiple customers.



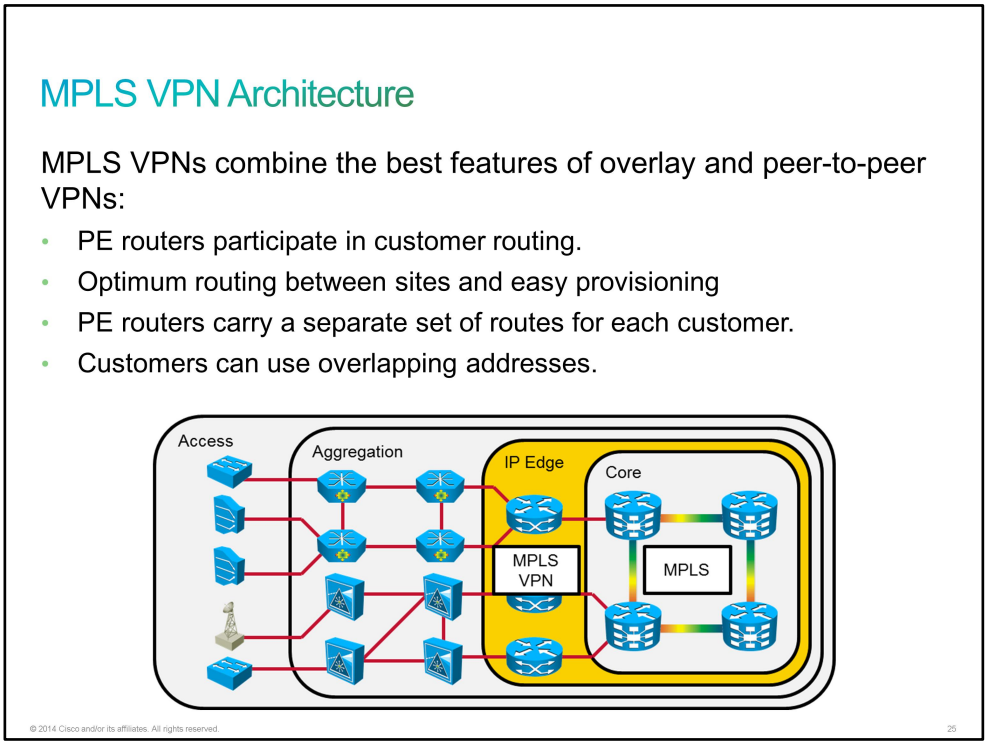
The main components of the MPLS VPN architecture are as follows:

- **C-network:** The C-network is usually a customer-controlled domain consisting of devices or routers spanning multiple sites that belong to the customer.
- **CE routers:** Routers in the customer network that interface with the service provider network are the CE routers.
- **P-network:** The P-networkis the provider-controlled domain consisting of PE and provider core routers that connect sites belonging to the customer on a shared infrastructure. The provider network controls the traffic routing between sites belonging to a customer, along with customer traffic isolation.
- **PE routers:** Routers in the provider network that interface or connect to the customer edge routers in the customer network are the PE routers. In an MPLS VPN implementation, the PE router is also known as the edge LSR.
- **P routers:** Routers in the core of the provider network that interface with either other provider core routers or provider edge routers are the P routers.

https://t.me/learningnets

MPLS VPN Architecture

This topic describes the basic MPLS VPN architecture components.

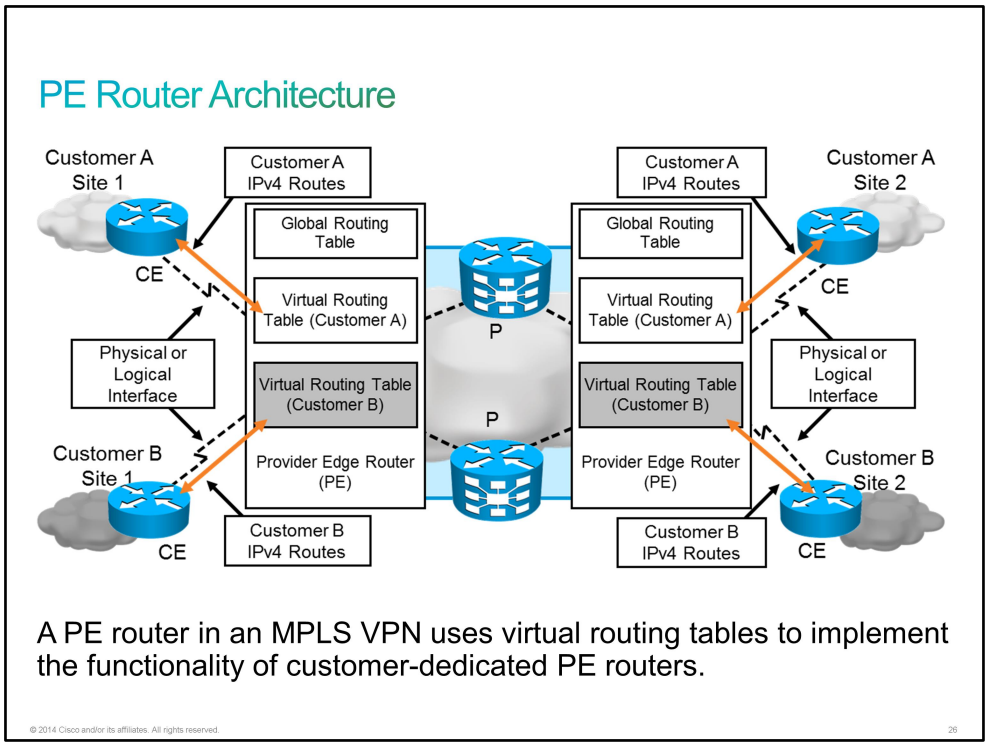


In the MPLS VPN architecture, the edge routers carry customer routing information, providing optimal routing for traffic that belongs to the customer for intersite traffic. The MPLS-based VPN model also accommodates customers who use overlapping address spaces, unlike the traditional peer-to-peer model, in which optimal routing of customer traffic required the provider to assign IP addresses to each of its customers (or the customer to implement NAT) to avoid overlapping address spaces. MPLS VPN is an implementation of the peer-to-peer model; the MPLS VPN backbone and customer sites exchange Layer 3 customer routing information, and data is forwarded between customer sites using the MPLS-enabled service provider IP backbone.

https://t.me/learningnets

PE Router Architecture

An MPLS VPN implementation is very similar to a dedicated router, peer-to-peer model implementation.



From the perspective of a CE router, only routing updates, as well as data, are forwarded to the PE router. The CE router does not need any specific configuration to enable it to be a part of an MPLS VPN domain. The only requirement on the CE router is a routing protocol (or a static default route) that enables the router to exchange IPv4 routing information with the connected PE router.

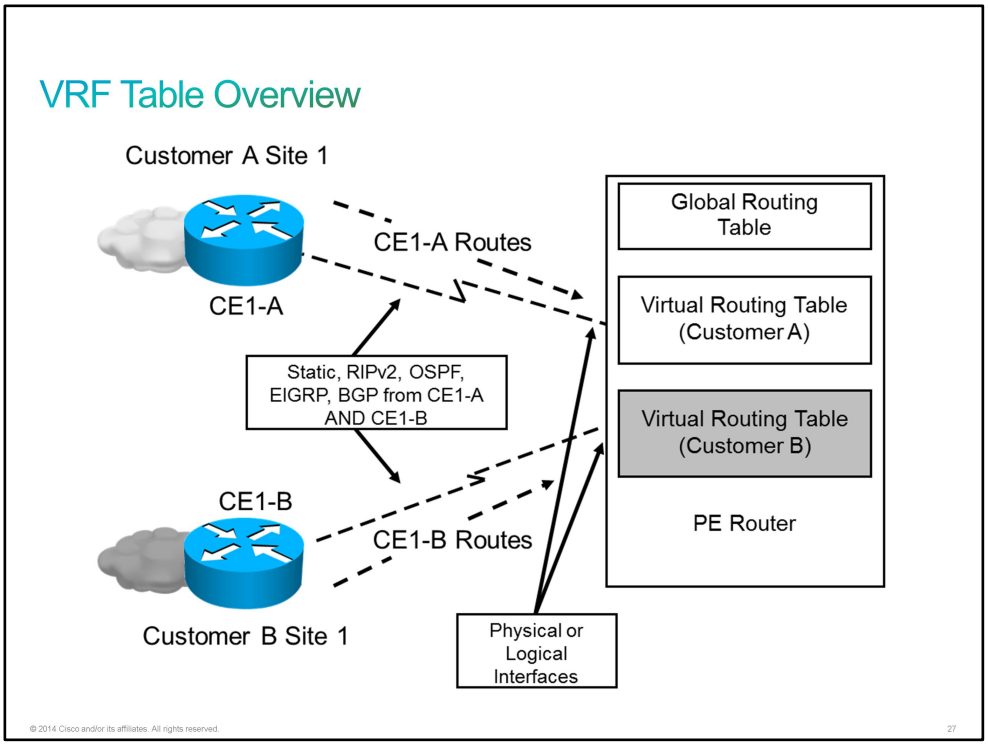
In the MPLS VPN implementation, the PE router performs multiple functions. The PE router must first be capable of isolating customer traffic if more than one customer is connected to the PE router. Each customer, therefore, is assigned an independent routing table (virtual routing table or VRF table) similar to a dedicated PE router in the initial peer-to-peer discussion. Routing across the service provider backbone is performed using a routing process in the global routing table. P routers provide label switching between PE routers and are unaware of VPN routes. CE routers in the customer network are not aware of the P routers, and thus the internal topology of the service provider network is transparent to the customer. The figure shows the functionality of the PE router.

The P routers are responsible only for label-switching of packets. They do not carry VPN routes and do not participate in MPLS VPN routing. The PE routers exchange IPv4 routes with connected CE routers using individual routing protocol contexts. To enable scaling the network to a large number of customer VPNs, MP-BGP is configured between PE routers to carry customer routes.

https://t.me/learningnets

VRF Table Overview

The VRF contains an IP routing table that is analogous to the global IP routing table, a Cisco Express Forwarding table, a list of interfaces that are part of the VRF, and a set of rules defining routing protocol exchange with attached CE routers (routing protocol contexts). In addition, the VRF also contains VPN identifiers and VPN membership information. The interface that is part of the VRF must support Cisco Express Forwarding switching. The number of interfaces that can be bound to a VRF is limited only by the number of interfaces on the router, and a single interface (logical or physical) can be associated with only one VRF.



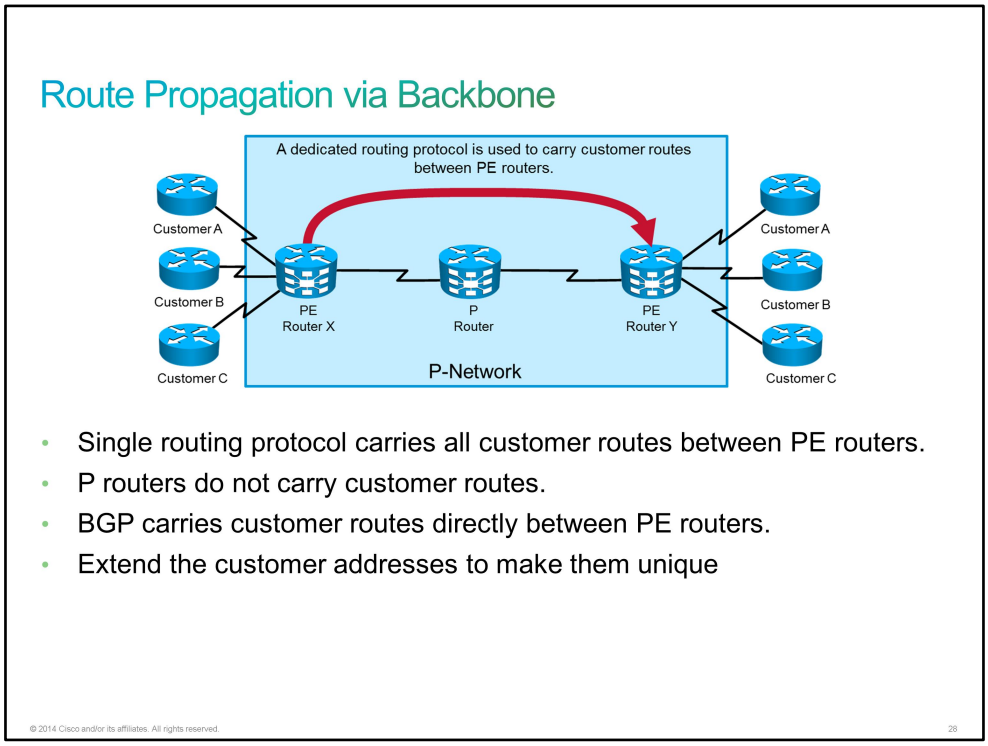
The figure shows the function of a VRF on a PE router to implement customer routing isolation. Cisco IOS Software supports various routing protocols and individual routing processes (OSPF, EIGRP, and so on) per router. However, for some routing protocols, such as RIP and BGP, Cisco IOS Software supports only a single instance of the routing protocol. Therefore, to implement per-VRF routing using protocols that are completely isolated from other VRFs, which might use the same PE-CErouting protocols, the concept of the routing context was developed.

Service providers want to control the size of the VRF table to prevent DoS attacks from customers who advertise too many routes into their VRFs. The maximum supported number of IP prefixes stored in a VRF depends on the platform, the size of the global routing table and other VRFs. You can administratively configure limits for the number of routes in the VRF, and the number of routes learned from a peer. There are also platform-specific features. The Deterministic Large VRF Support feature, for example, available on the Cisco ASR 9000 can be used to allow a number of prefixes per VRF, up to global FIB size.

https://t.me/learningnets

Route Propagation via Backbone

Although virtual routing tables provide isolation between customers, the data from these routing tables still needs to be exchanged between PE routers to enable data transfer between the sites attached to different PE routers.



A routing protocol is needed that will transport all customer routes across the P-network while maintaining the independence of individual customer address spaces.

The best solution to the customer route propagation issue is to run a single routing protocol between PE routers that will exchange all customer routes without the involvement of the P routers. This solution is scalable. Some of the benefits of this approach are as follows:

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

The next design decision to be made is the choice of the routing protocol running between PE routers. Given that the total number of customer routes is expected to be very large, the only well-known protocol with the required scalability is BGP4. In fact, BGP4 is used in the MPLS VPN architecture to transport customer routes directly between PE routers.

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

With the deployment of a single routing protocol (that is, BGP4) exchanging all customer routes between PE routers, an important issue arises: how can BGP4 propagate several identical prefixes, belonging to different customers, between PE routers?

https://t.me/learningnets

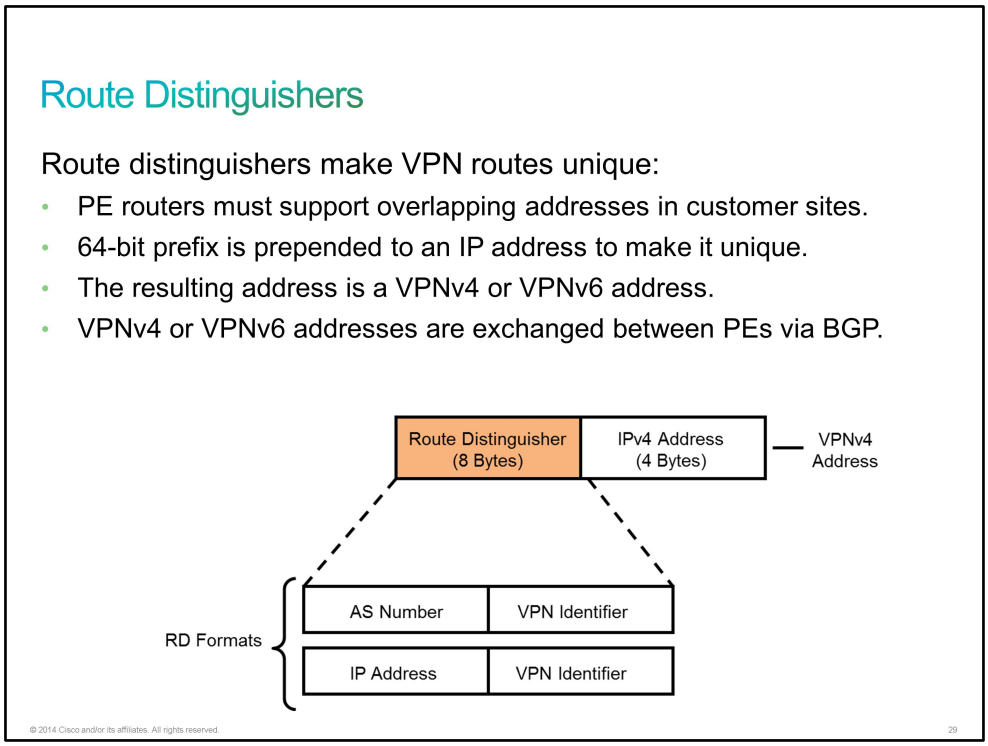
The only solution to this dilemma is the expansion of customer IP prefixes with a unique prefix that makes them unique even if they had previously overlapped. A 64-bit prefix called the RDis used in MPLS VPNs to convert non-unique 32-bit customer addresses into 96-bit unique addresses that can be transported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

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Route Distinguishers

In the MPLS VPNs, the PE routers use VRFs to isolate different customers connected to the same PE.



The PE router must support overlapping address spaces in connected customer networks. The PE routers will learn routes from attached customer networks and propagate this information over the shared provider backbone. Similarly, they will receive customer-specific routes from other PE routers, install them in the customer VRF and forward them to the CE router.

The purpose of a RD is to make the customer prefixes unique within the entire MPLS domain. This 64-bit unique identifier is prepended to the 32-bit customer route available in the customer VRF to make a unique 96-bit network address that can be transported between the PE routers in the MPLS domain. The RD is configured per VRF on a PE router. The resulting address is known as a VPNv4 address (32-bit customer prefix plus 64-bit unique identifier or RD) if it carries IPv4 addresses, or a VPNv6 address if it includes an IPv6 customer prefix.

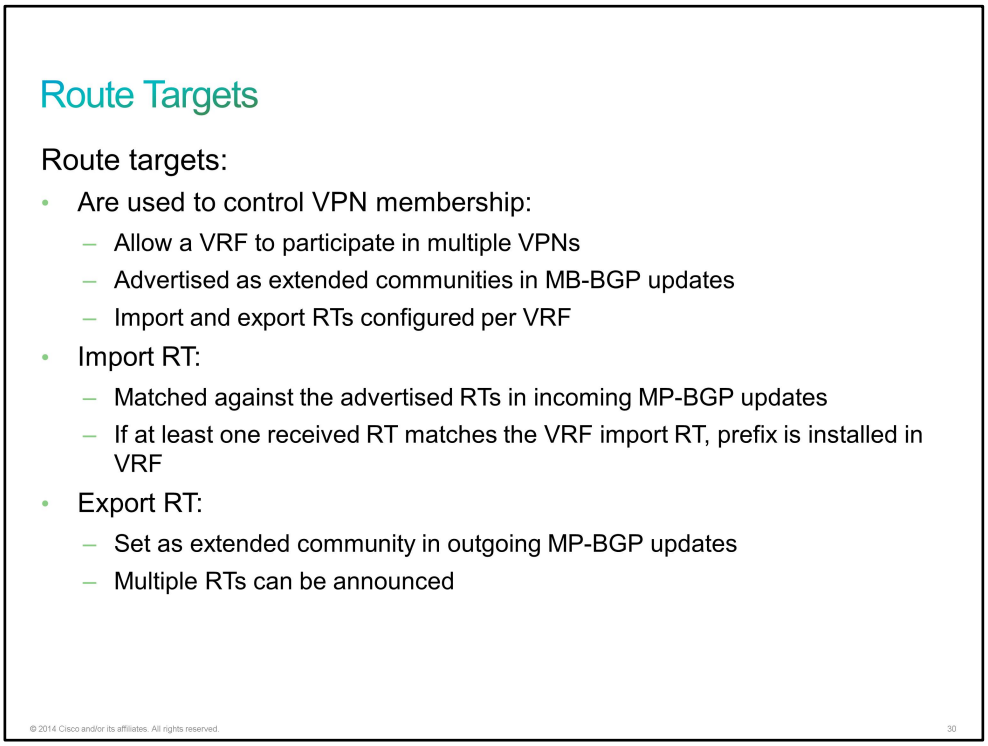
VPNv4 addresses are exchanged only between PE routers; they are never sent to or received from CE routers. BGP used between the PE routers has to support the exchange of traditional IPv4 prefixes and the exchange of VPNv4 prefixes. BGP has been extended to support VPNv4 prefixes. The BGP sessions established between the PE routers, used to carry VPNv4 prefixes, are referred to as MP-BGP sessions.

The format of an RD is shown in the figure. An RD can be of two formats. If the provider does not have a BGP AS number, the IP address format can be used, and if the provider does have an AS number, the AS number format can be used.

https://t.me/learningnets

Route Targets

Consider a scenario where some sites have to participate in more than one VPN. In such a scenario, the RDs cannot identify participation in more than one VPN.



RTs were introduced into the MPLS VPN architecture to support identifying a site that participates in more than one VPN. RTs are additional identifiers used in the MPLS VPN that identify the VPN membership of the routes learned from that particular site. RTs are implemented by the use of extended BGP communities in which the higher-order 16 bits of the BGP extended community (64 total bits) are encoded with a value corresponding to the VPN membership of the specific site. When a VPN route learned from a CE router is injected into VPNv4 BGP, a list of VPN RT extended community attributes is associated with it.

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

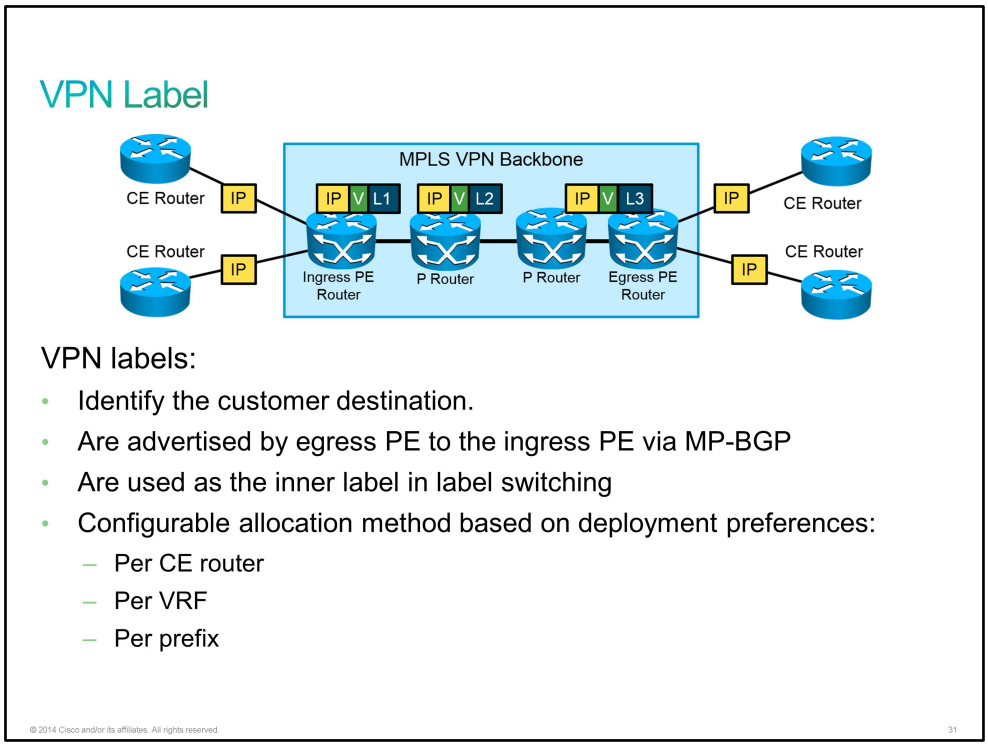
When the VPNv4 routes are propagated to other PE routers, those routers need to select the routes to import into their virtual routing tables. This selection is based on import RTs. Each virtual routing table in a PE router can have a number of configured import RTs that identify the set of VPNs from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

When implementing complex VPN topologies, such as extranet VPN, Internet access VPNs, and network management VPN, using MPLS VPN technology, the RT plays a pivotal role. A single prefix can be associated to more than one export route target when propagated across the MPLS VPN network. The RT can, as a result, be associated to sites that might be a member of more than one VPN.

https://t.me/learningnets

VPN Label

The VPN label (4 bytes) is generated by the egress PE to identify the customer destination.



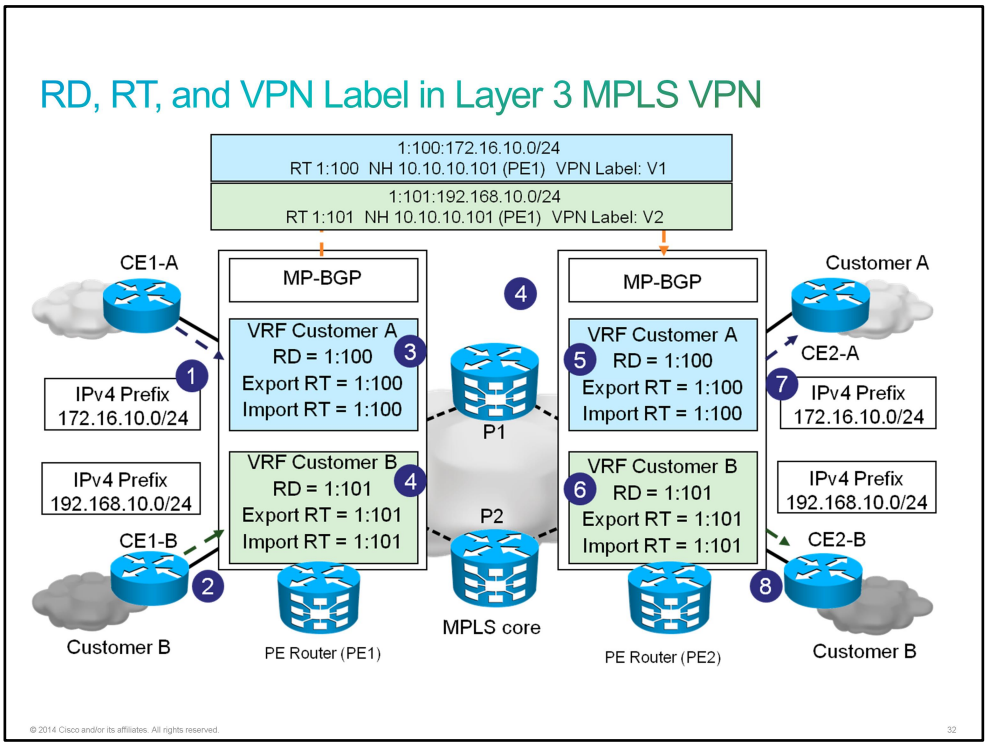
The egress PE uses MP-BGP to advertise the VPNv4 prefix along with the route target and VPN label to the ingress PE. The VPN label is needed as the inner label for the MPLS label switching.

You can choose the label mode based on your deployment preferences. In the per-CE mode, a unique VPN label will be allocated for each CE router. In the per-VRF mode, the same label is generated for all prefixes in the VRF. In the per-VRF mode, each prefix will be assigned a unique VPN label. You can configure the label mode based on arbitrary match criteria such as prefix value and community. The configuration is performed using using a route-policy. This feature is often referred to as the Flexible L3VPN Label Allocation Mode feature.

https://t.me/learningnets

RD, RT, and VPN Label in Layer 3 MPLS VPN

The figure shows the processes that occur during route propagation in an MPLS VPN.

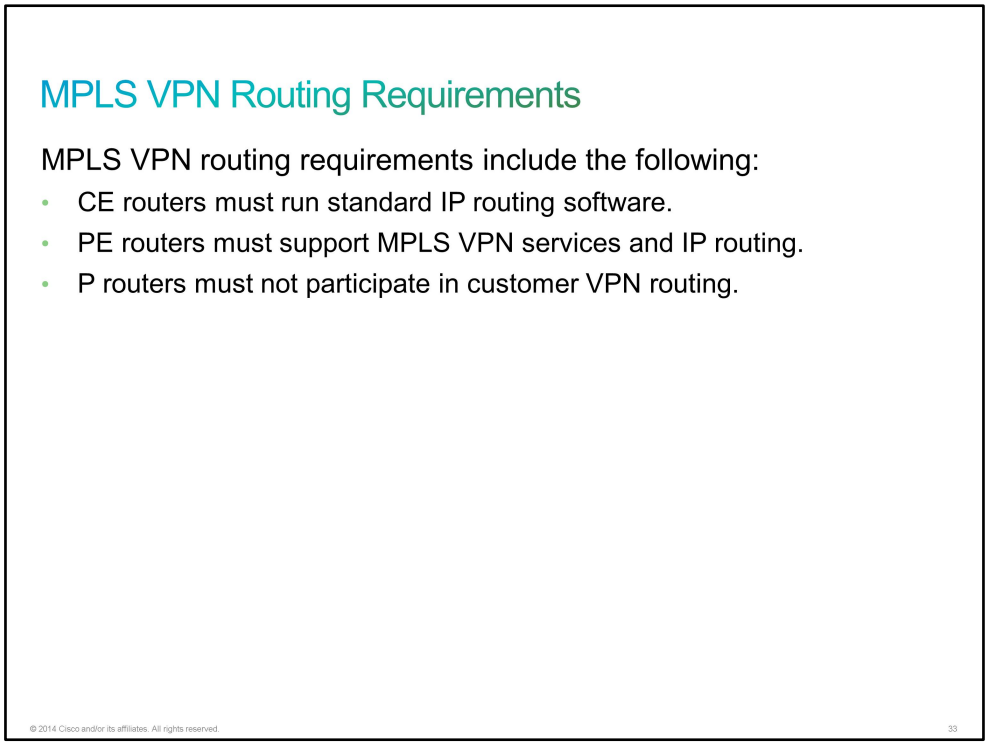


The following processes occur during route propagation in an MPLS VPN, as shown in the figure:

1. The prefix 172.16.10.0/24 is received from CE1-A. This subnet is part of VRF Customer A on PE1. The subnets in this VRF are associated with an RD value of 1:100 and an export RT value of 1:100.
2. The prefix 192.168.10.0/24 is received from CE1-B. This subnet is part of VRF Customer B on PE1. The subnets in this VRF are associated with an RD value of 1:101 and an export RT value of 1:101.
3. Routes learned from CE1-A are redistributed into the MP-BGP process on PE1. The prefix 172.16.10.0/24 is prepended with the RD value of 1:100 and advertised with the RT of 1:100 to PE2.
4. Routes learned from CE1-B are redistributed into the MP-BGP process on PE1. The prefix 192.168.10.0/24 is prepended with the RD value of 1:101 and advertised with the route target of 1:101 to PE2. The VPN labels are depicted by the entries V1 and V2 in the figure.
5. The MP-BGP update is received by PE2, and the routes are stored in the appropriate VRF tables based on the advertised route target. The VRF of Customer A is configured with the import RT of 1:100. The prefix 1:100:192.168.10.0/24, received with RT 1:100 is therefore imported into the Customer A VRF.
6. The VRF of Customer B is configured with the import RT of 1:101. The prefix 1:101:192.168.10.0/24, received with route target 1:101 is therefore imported into the Customer B VRF.
7. The received MP-BGP routes are redistributed into the VRF PE-CE routing process, and the Customer A routes are propagated to CE2-A.
8. Likewise, the Customer B routes are propagated to CE2-B.

https://t.me/learningnets

MPLS VPN Routing Requirements



The designers of MPLS VPN technology were faced with these routing requirements:

- CE routers should not be MPLS VPN-aware; CE routers should run standard IP routing software.
- PE routers must support MPLS VPN services and traditional Internet services.

To make the MPLS VPN solution scalable, P routers must not participate in customer VPN routing. P routers use only label switching.

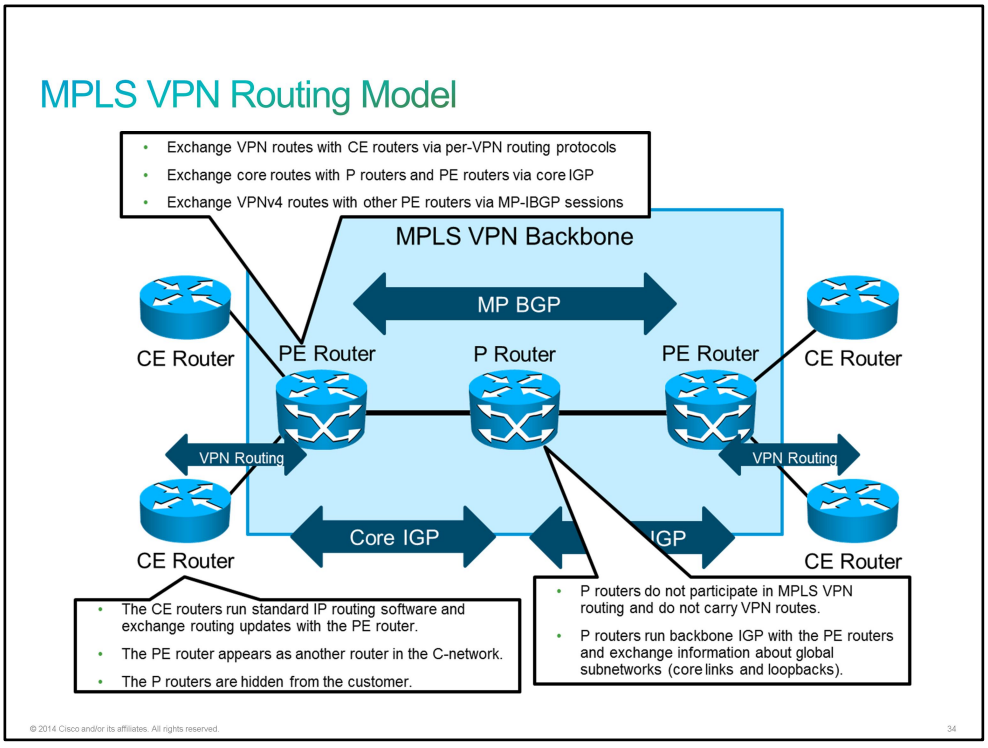
https://t.me/learningnets

MPLS VPN Routing Model

The MPLS VPN backbone should look like a standard corporate backbone to the CE routers. The CE routers run standard IP routing software and exchange routing updates with the PE routers, which appear to them as normal routers in the C-network. An MPLS VPN implementation is very similar to a dedicated router peer-to-peer model implementation. From the perspective of a CE router, only IPv4 updates and data are forwarded to the PE router. The CE router does not need any specific configuration to enable it to be a part of an MPLS VPN domain. The only requirement on the CE router is a routing protocol (or a static default route) that enables the router to exchange IPv4 routing information with the connected PE router that appears as a normal router in the C-network.

From the customer perspective, the MPLS VPN backbone looks like an intracompany BGP backbone with PE routers performing route redistribution between individual sites and the core backbone. The standard design rules that are used for enterprise BGP backbones can be applied to the design of the C-network. The P routers are hidden from customer view; the internal topology of the BGP backbone is transparent to the customer.

From the P router perspective, the MPLS VPN backbone looks even simpler—the P routers do not participate in MPLS VPN routing and do not carry VPN routes. The P routers run only a backbone IGP with other P routers and with PE routers, and exchange information about core subnetworks. BGP deployment on P routers is not needed for proper MPLS VPN operation; BGP deployment might be needed, however, to support traditional Internet connectivity that has not yet been migrated to MPLS.



The PE routers are the only routers in the MPLS VPN architecture that see all routing aspects of the MPLS VPN. PE routers can perform these exchanges:

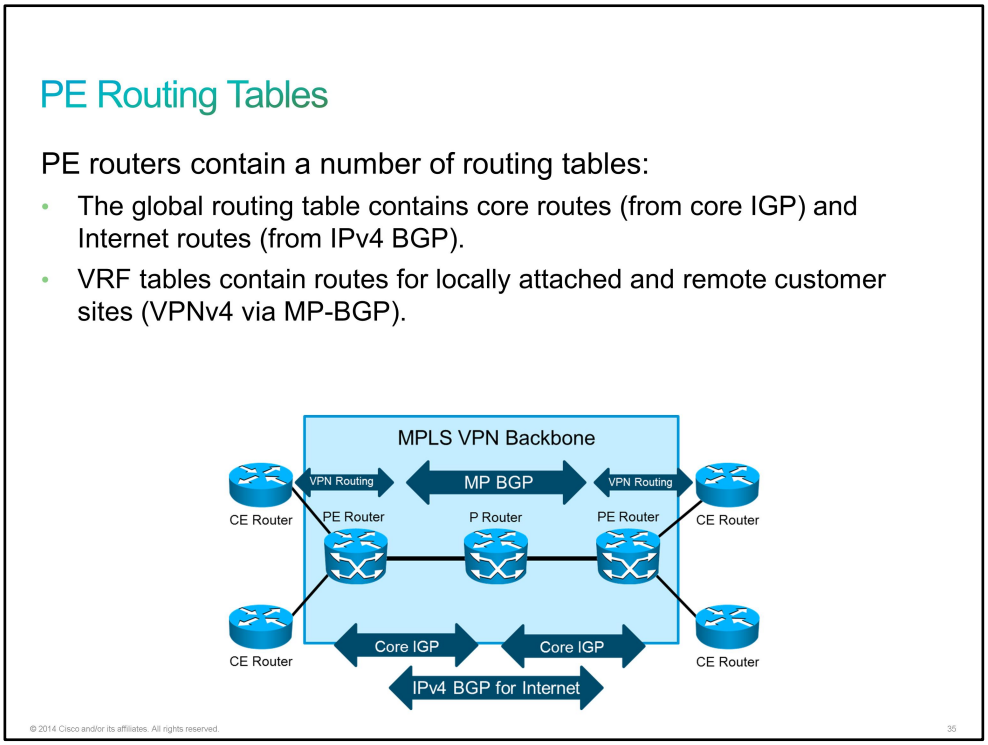
- PE routers exchange IPv4 VPN routes with CE routers via various routing protocols running in the virtual routing tables.
- PE routers exchange VPNv4 routes via MP-BGP sessions with other PE routers.

https://t.me/learningnets

- PE routers exchange core routes with P routers and other PE routers via core IGP.

https://t.me/learningnets

PE Routing Tables



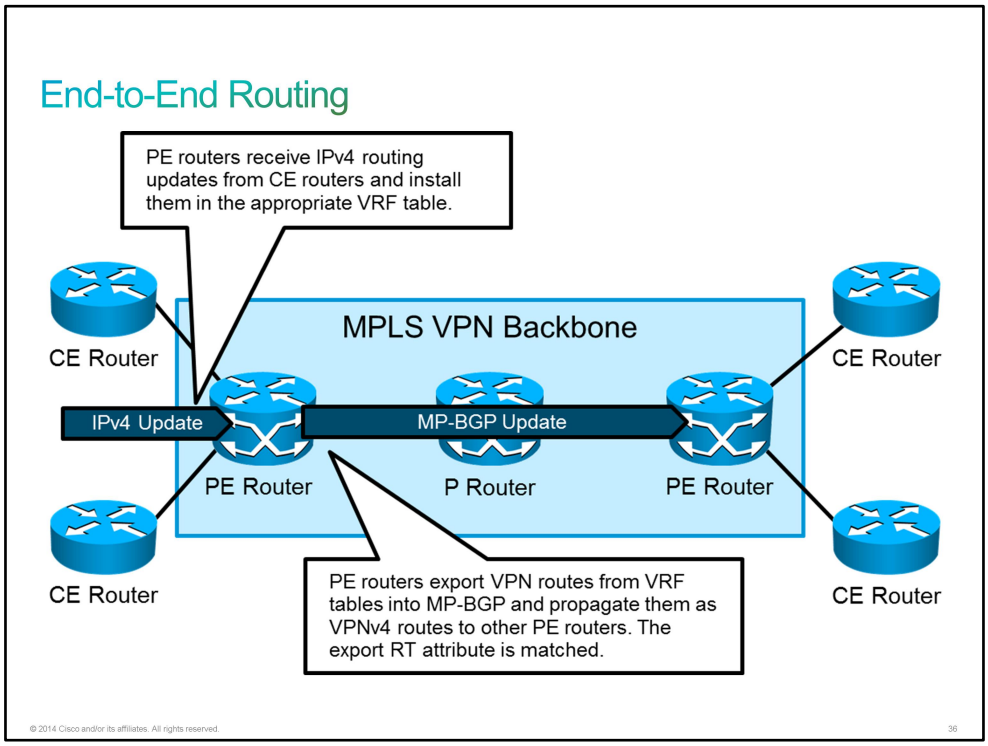
The PE routers fulfill various routing requirements imposed on them by using a number of IP routing tables. Here are some examples:

- The global IP routing table (the IP routing table that is always present in a router even if it is not supporting an MPLS VPN) contains all core routes (inserted by the core IGP) and Internet routes (inserted from the global IPv4 BGP table).
- The VRF tables contain sets of routes for sites with identical routing requirements. The VRFs are filled with intra-VPN IGP information that is exchanged with the CE routers and with VPNv4 routes that are received through MP-BGP sessions from the other PE routers.

https://t.me/learningnets

End-to-End Routing

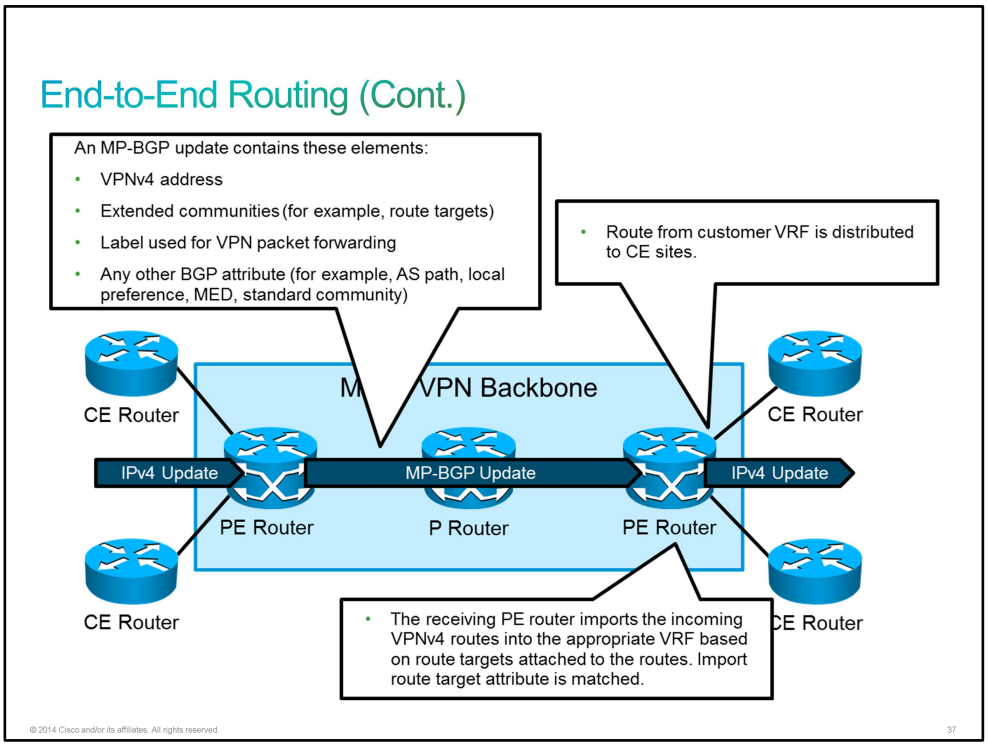
The figure shows how PE routers receive IPv4 routing updates from the CE routers and install the updates in the appropriate VRF table.



Cisco IOS, IOS XE, and IOS XR Software support RIPv2 (multiple contexts), EIGRP (multiple contexts), OSPFv2 (multiple processes), IS-IS (multiple contexts), and BGP4 (multiple contexts) as routing protocols that can be used per-VRF to exchange customer routing information between CE routers and PE routers. The VRF interfaces on PE routers can be either logical or physical, but each interface can be assigned to only one VRF.

The customer routes from VRF tables are exported as VPNv4 routes into MP-BGP and propagated to other PE routers. The MP-BGP sessions between the PE routers are IBGP sessions and are subject to the IBGP split-horizon rules. Either a full mesh of MP-IBGP sessions is required between PE routers, or route reflectors need to be used.

https://t.me/learningnets



An MP-BGP update exchange between PE routers contains these elements:

- VPNv4 address
- Extended BGP communities. At least one route target is required for MPLS VPNs.
- Label used for VPN packet forwarding
- Mandatory BGP attributes (for example, AS path). Optionally, the MP-BGP update can contain any other BGP attribute, such as local preference, MED, or standard BGP community.

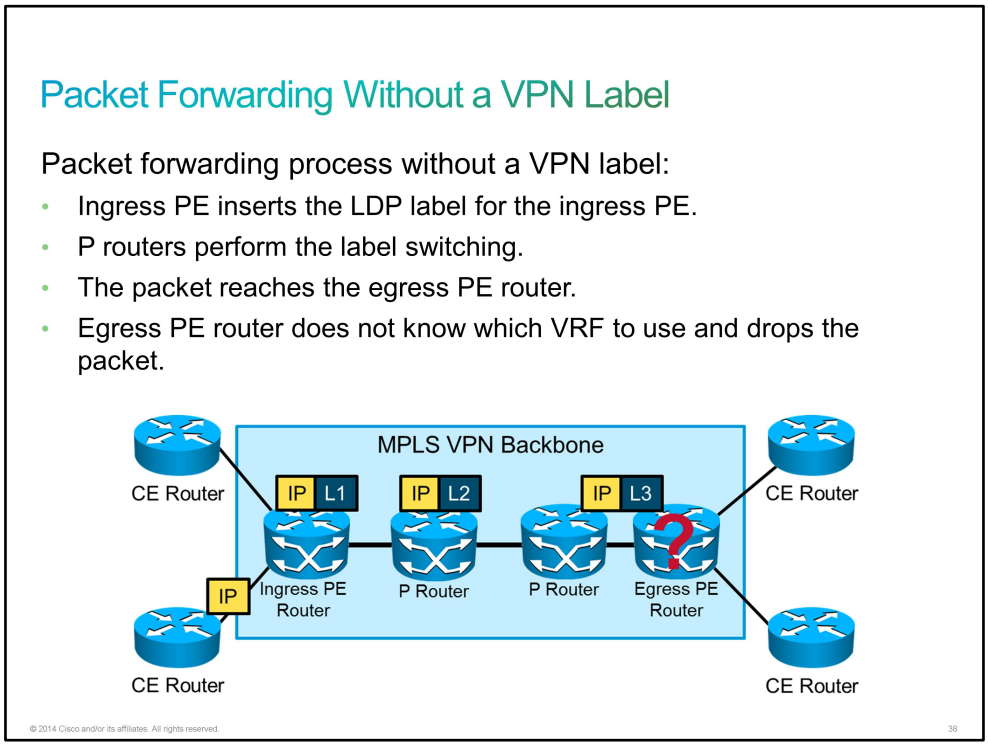
The PE routers receiving MP-BGP updates import the incoming VPNv4 routes into their VRFs based on RTs attached to the incoming routes and based on import RTs configured in the VRFs. The VPNv4 routes installed in the VRFs are converted to IPv4 routes and then propagated to the CE routers.

The RTs that are attached to a route and the import RTs that are configured in the VRF direct the propagation of the routes to the CE router. Incoming VPNv4 routes are imported into VRFs on the receiving PE router, only if at least one RT attached to the route matches at least one import RT that is configured in the VRF.

https://t.me/learningnets

Packet Forwarding Without a VPN Label

A simple MPLS-oriented approach to MPLS VPN packet forwarding across the MPLS VPN backbone would be to label the customer packet with the label assigned by LDP for the egress PE router.

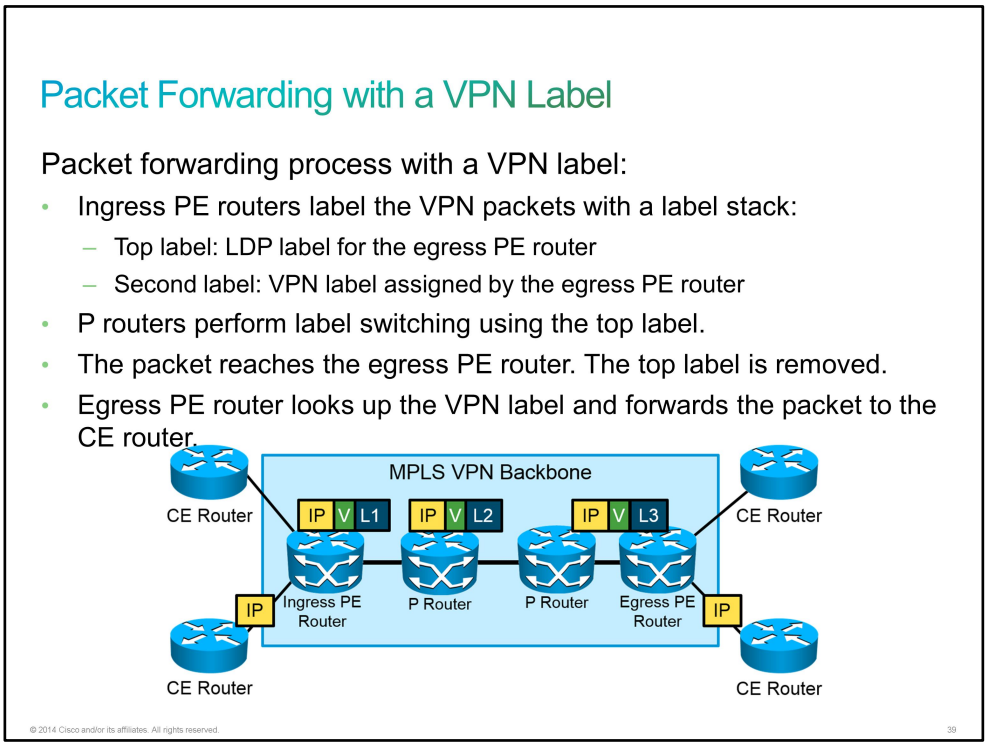


The core routers, consequently, would never see the customer IP packet; instead, the core routers would see just a labeled packet targeted toward the egress PE router. The core routers would perform simple label-switching operations, eventually delivering the customer packet to the egress PE router. Unfortunately, the customer IP packet would contain no VPN or VRF information that could be used to perform VRF lookup on the egress PE router. The egress PE router would not know which VRF to use for packet lookup and would need to drop the packet.

https://t.me/learningnets

Packet Forwarding with a VPN Label

An MPLS label stack can be used to tell the egress PE router what to do with the VPN packet.



When using the label stack, the ingress PE router labels the incoming IP packet with two labels. The top label in the stack is the LDP label for the egress PE router; this label guarantees that the packet will traverse the MPLS VPN backbone and arrive at the egress PE router. The second label in the stack is assigned by the egress PE router, and tells how to forward the incoming VPN packet. The second label could point directly toward an outgoing interface, in which case the egress PE router would perform label lookup only on the VPN packet. The second label could also point to a VRF, in which case the egress PE router would first perform a label lookup to find the target VRF, and then perform an IP lookup within the VRF.

Both methods of implementing second labels are used in Cisco IOS, IOS XE, and IOS XR Software. The second label in the stack points toward an outgoing interface whenever the CE router is the next hop of the VPN route. The second label in the stack points to the VRF table for aggregate VPN routes, VPN routes pointing to a null interface, and routes for directly connected VPN interfaces.

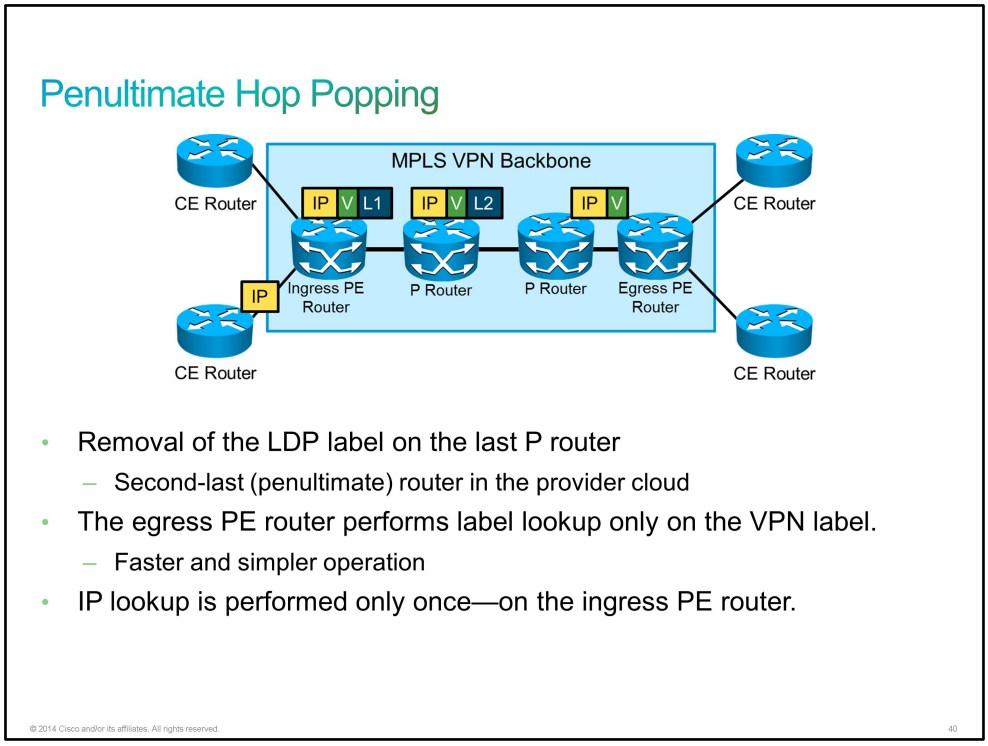
The two-level MPLS label stack satisfies these MPLS VPN forwarding requirements:

- The P routers perform label switching on the LDP-assigned label toward the egress PE router.
- The egress PE router performs label switching on the second label (which it has previously assigned) and either forwards the IP packet toward the CE router or performs another IP lookup in the VRF that is pointed to by the second label in the stack.

https://t.me/learningnets

Penultimate Hop Popping

PHP is the removal of the top label in the stack on the hop prior to the egress router.

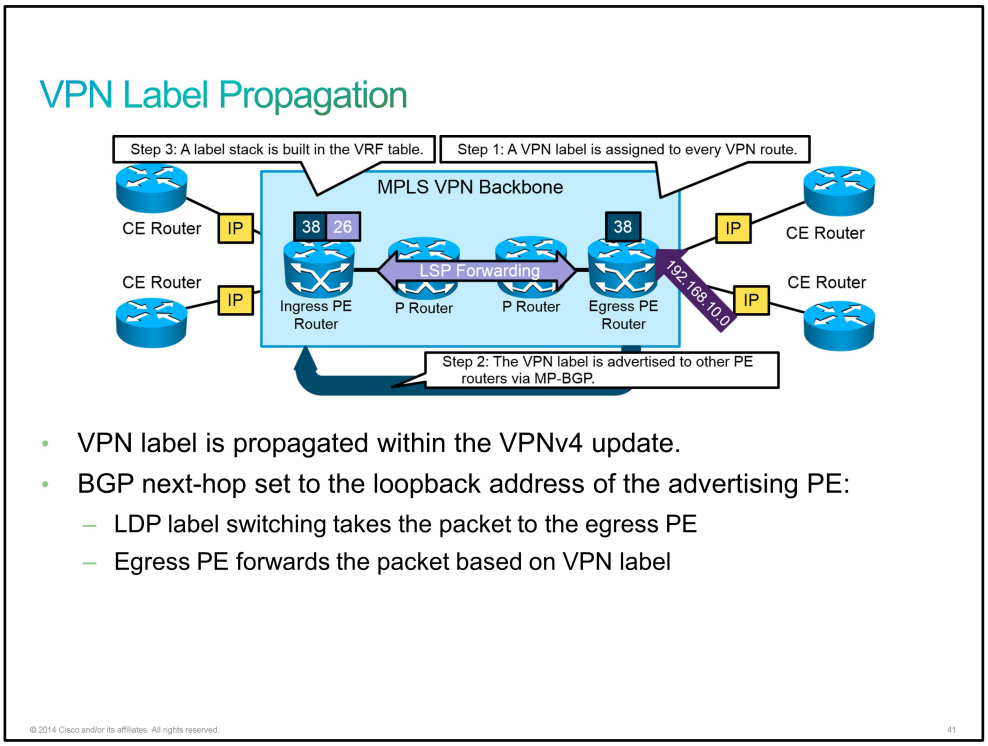


PHP can be performed in frame-based MPLS networks. In these networks, the last P router in the LSPtunnel pops the LDP label (as previously requested by the egress PE router through LDP), and the PE router receives a labeled packet that contains only the VPN label. In most cases, a single label lookup performed on that packet in the egress PE router is enough to forward the packet toward the CE router. The full IP lookup through the FIB is performed only once, in the ingress PE router, even without PHP.

https://t.me/learningnets

VPN Label Propagation

The MPLS label stack needs to contain the VPN label for proper MPLS VPN operation.



The VPN label is allocated by the egress PE router. This label needs to be propagated from the egress PE router to the ingress PE routers to enable proper packet forwarding. MP-BGP is used as the propagation mechanism. The MP-BGP updates carry a label assigned by the egress PE router together with the 96-bit VPNv4 prefix.

These steps describe the label propagation between PE routers:

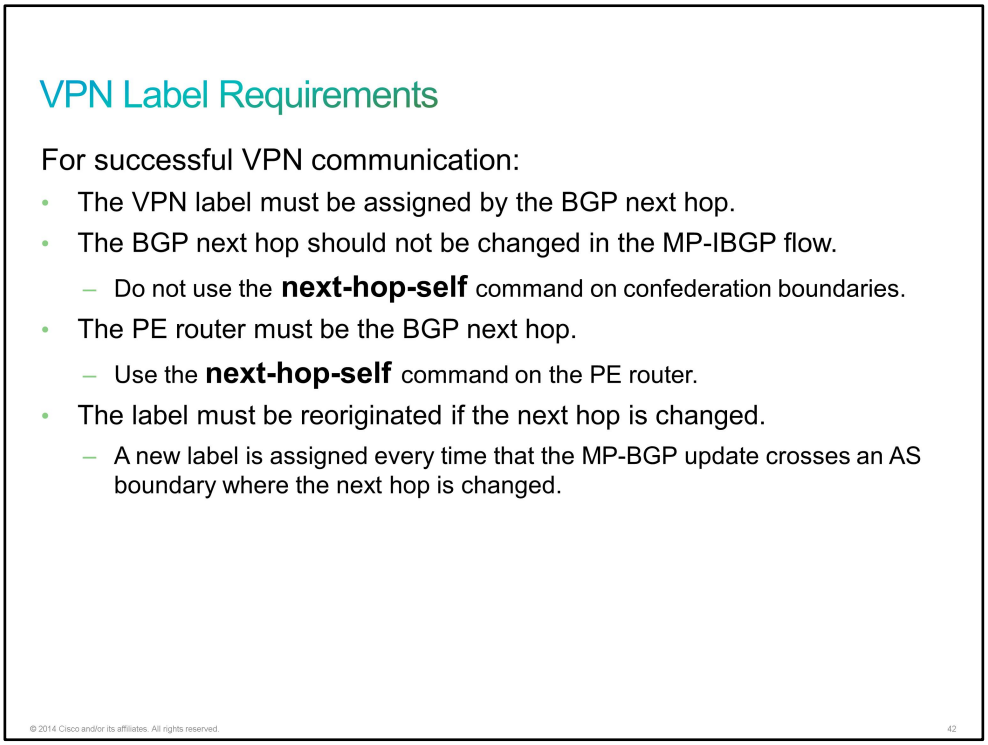
- The egress PE router assigns a label to every VPN route received from the attached CE routers and to every summary route summarized inside the PE router. This label is then used as the second label in the MPLS label stack by the ingress PE routers when labeling VPN packets. In the figure, the VPN label 38 for destination 192.168.10.0 is assigned by the egress PE router. The VPN labels that are assigned locally by the PE router can be inspected with the Cisco IOS, IOS XE, and IOS XR **show mpls forwarding** vrf *vrf-name* command.
- The VPN labels that are assigned by the egress PE routers are advertised to all other PE routers (participating in the VPN) together with the VPNv4 prefix in MP-BGP updates. The labels can be inspected with the Cisco IOS and IOS XE **show ip bgp vpnv4 all labels** command or the Cisco IOS XR **show bgp vpnv4 unicast labels** command on the ingress PE router. The routes that have an input label but no output label are the routes received from the CE routers (and the input label was assigned by the local PE router). The routes with an output label but no input label are the routes received from the other PE routers (and the output label was assigned by the remote PE router).

https://t.me/learningnets

- The ingress PE router has two labels associated with a remote VPN route: a label for the next hop assigned by the next-hop P router via LDP—and taken from the local LIB—and also the label assigned by the remote PE router and propagated via the MP-BGP update. Both labels are combined in a label stack and installed in the VRF table. The label stack in the VRF table can be inspected using the Cisco IOS and IOS XE **show ip cef vrf** *vrf-name detail* command or the Cisco IOS XR **show cef vrf** *vrf-name detail* command. The tag-imposed values in the output display the MPLS label stack. The first label in the MPLS label stack is the LDP label forwarded toward the egress PE router, and the second label is the VPN label advertised by the egress PE router.

https://t.me/learningnets

VPN Label Requirements



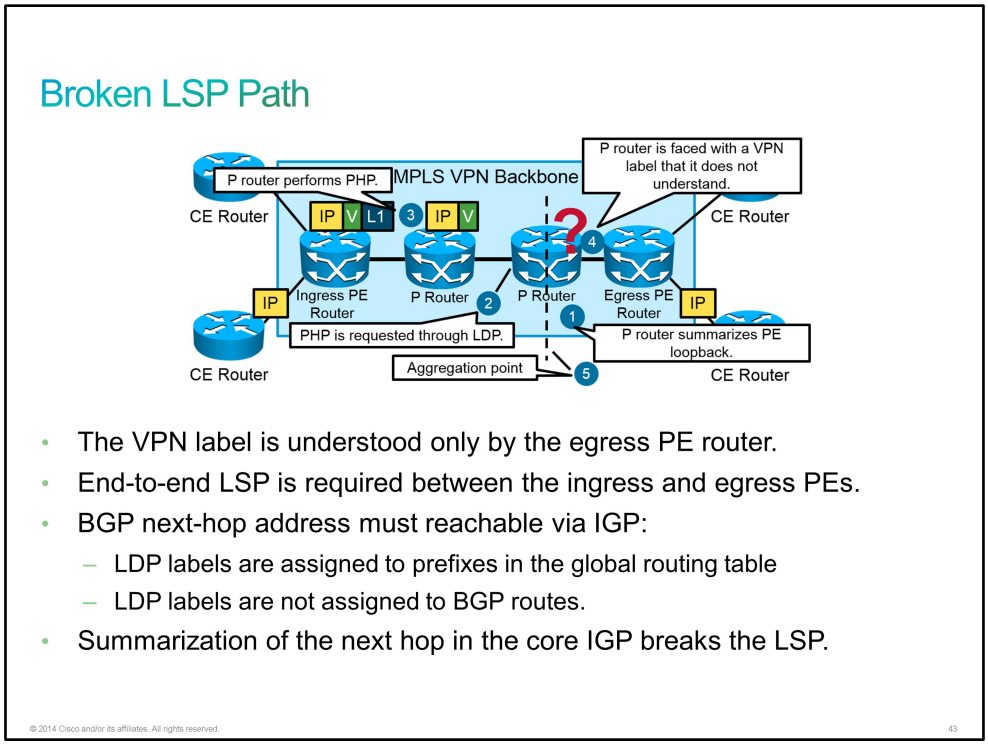
MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same PE router that assigned the second label in the label stack. Here are three scenarios that can cause the BGP next hop to be different from the IP address of the PE router assigning the VPN label:

- If the customer route is received from the CE router via an EBGP session, the next hop of the VPNv4 route is still the IP address of the CE router (the BGP next hop of an outgoing IBGP update is always identical to the BGP next hop of the incoming EBGP update). You must configure the **next-hop-self** command on the MP-BGP sessions between PE routers to make sure that the BGP next hop of the VPNv4 route is always the IP address of the PE router, regardless of the routing protocol used between the PE router and the CE router.
- The BGP next hop should not change inside an AS.
- The BGP next hop is always changed on an EBGP session. If the MPLS VPN network spans multiple public autonomous systems, special provisions must be made in the AS boundary routers to reoriginate the VPN label at the same time that the BGP next hop is changed.

https://t.me/learningnets

Broken LSP Path

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



Listed here are two scenarios that could cause the LSP tunnel between PE routers to break:

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

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

https://t.me/learningnets

Summary

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

Summary

- In MPLS VPN, customer traffic is isolated on the same PE router that provides connectivity into the service provider network for multiple customers.
- A PE router in an MPLS VPN uses virtual routing tables to implement the functionality of customer-dedicated PE routers.
- VRF also contains VPN identifiers and VPN membership information.
- BGP is used to exchange customer routes directly between PE routers.
- The 64-bit route distinguisher is prepended to an IPv4 address to make it globally unique.
- Route targets are additional attributes attached to VPNv4 BGP routes to indicate VPN membership.
- The MP-BGP update is received by the PE and the route is stored in the appropriate VRF table for the customer, based on the VPN label.

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

Summary (Cont.)

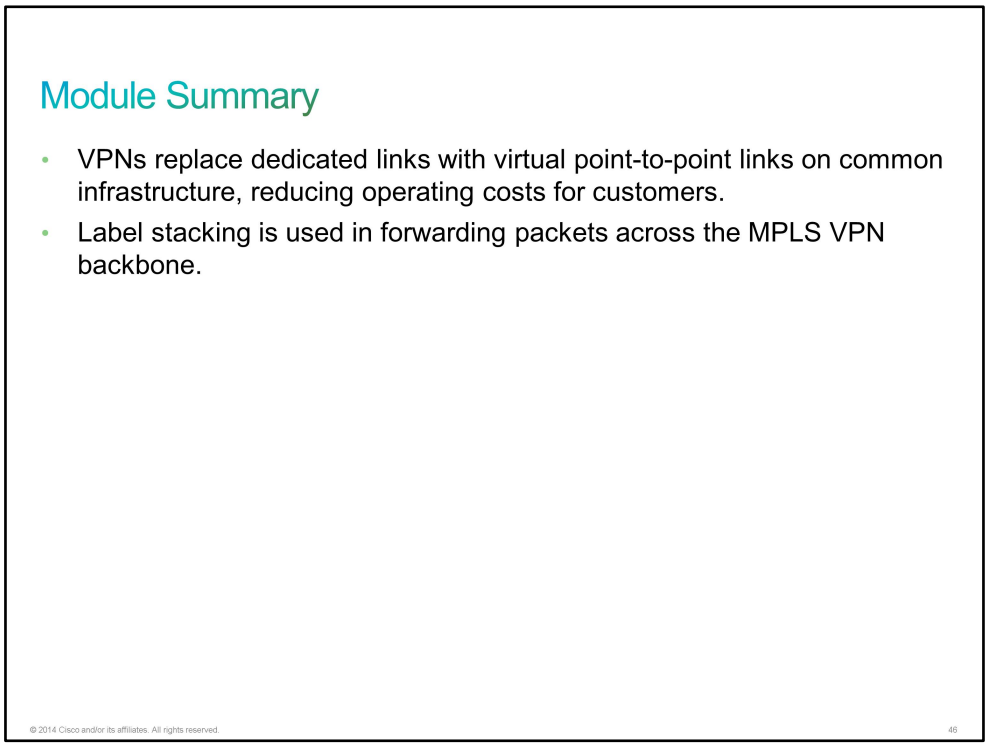
- CE routers should not be MPLS VPN-aware and should run standard IP routing software.
- PE routers export VPN routes from VRF tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.
- PHP is the removal of the top label in the stack on the hop prior to the egress router.
- MPLS VPN packet forwarding works correctly only if the router specified as the BGP next hop in the incoming BGP update is the same PE router that assigned the second label in the label stack.
- For successful propagation of MPLS VPN packets across an MPLS backbone, there must be an unbroken LSP tunnel between PE routers.

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

https://t.me/learningnets

Lesson 3: Module Summary

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



The two major VPN design options—overlay VPN and peer-to-peer VPN—have many benefits and drawbacks. The VPN topology categories and architectural components help determine the method for forwarding packets in a MPLS VPN environment.

https://t.me/learningnets

https://t.me/learningnets

Lesson 4: 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 three options are advantages of VPNs? (Choose three.) (Source: Introducing VPNs)

   A. cost savings
   B. scalability
   C. improved security
   D. complexity
   E. simplified routing

2. Which two options are Layer 3 overlay VPN technologies? (Choose two.) (Source: Introducing VPNs)

   A. Frame Relay
   B. DMVPN
   C. IPsec
   D. ATM

3. Which protocol is used to transport customer routes directly between PE routers? (Source: Introducing MPLS VPNs)

   A. RIP
   B. VPN
   C. BGP
   D. OSPF

https://t.me/learningnets

4. In which two ways do MPLS VPNs support overlapping customer address spaces? (Choose two.) (Source Introducing MPLS VPNs)

   A. by implementing unique RDs for each customer
   B. by implementing unique RTs for each customer
   C. by implementing different LSPs for each customer
   D. by implementing virtual routing spaces for each customer

5. Why do MPLS VPNs implement route targets? (Source: Introducing MPLS VPNs)

   A. to direct the CE-to-PE routing updates into the desired VRF
   B. to allow a site to participate in more than one VPN
   C. to convert a customer address to an MP-BGP address
   D. to convert a non-unique IP address into a unique VPNv4 address

6. Which type of routers exchange VPNv4 routes? (Source: Introducing MPLS VPNs)

   A. P
   B. CE
   C. PE
   D. core

7. Which protocol is used to provide Internet access? (Source: Introducing MPLS VPNs)

   A. IS-IS
   B. EIGRP
   C. BGP4
   D. BGP VPNv4

8. Which requirement do MPLS VPNs impose on CE routers? (Source: Introducing MPLS VPNs)

   A. The CE routers must support BGP.
   B. The CE routers must run a link-state protocol.
   C. The CE routers can run any standard IP routing protocol.
   D. The IGP of the CE routers must be upgraded to a VPN-aware IGP.

9. Why would full IPv4 routing be enabled on a PE router? (Source: Introducing MPLS VPNs)

   A. to support the MPLS VPN route update
   B. to support the MPLS VPN route target exports
   C. to provide Internet access
   D. to support the transport of MP-BGP extended communities

10. Which two types of routes would an MPLS VPN install into the VRF? (Choose two.) (Source: Introducing MPLS VPNs)

    A. routes received via an ISIS update
    B. routes received via a VPNv4 update
    C. routes received via the core IGP update
    D. routes received in a customer update

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

11. Which protocol is used to transport VPN labels between PE routers? (Source: Introducing MPLS VPNs)

    A. LDP
    B. RSVP
    C. MP-BGP
    D. the core IGP

12. How can P routers forward VPN packets if they do not have VPN routes? (Source: Introducing MPLS VPNs)

    A. They forward based on the LSP label.
    B. They forward based on the VPN label.
    C. They forward based on the MP-BGP next hop.
    D. They forward based on a routing table lookup of the IP address.

13. Which router assigns the VPN label? (Source: Introducing MPLS VPNs)

    A. P
    B. egress CE
    C. egress PE
    D. ingress CE
    E. ingress PE

14. What is used to identify the label that will be used to transport the VPN packet to the egress router? (Source: Introducing MPLS VPNs)

    A. the IGP least-cost path
    B. the EBGP next-hop address
    C. the MP-IBGP next-hop address
    D. the VPN label entry in the LFIB

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. A, B, C
2. B, C
3. C
4. A, D
5. B
6. C
7. C
8. C
9. C
10. B, D
11. C
12. A
13. C
14. C

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 2: MPLS Layer 3 VPNs

Introduction

This module covers MPLS VPN implementation on Cisco IOS Software platforms. The module describes the concepts of VRF tables, the interaction between customer-to-provider routing protocols, and MP-BGP in the service provider backbone. It also describes advanced MPLS VPN-specific routing protocol features. The module continues with a description of the MPLS VPN monitoring and debugging commands that are available on Cisco IOS platforms and describes troubleshooting, including failure scenarios, symptoms, and remedial action.

Objectives

Upon completing this module, you will be able to:

- Describe the VRF table and other MPLS VPN attributes that are associated with a VRF instance
- Configure small-scale routing protocols (static, RIP, and EIGRP) between CE and PE routers
- Configure OSPF and BGP as the routing protocol between CE and PE routers and explain how to troubleshoot MPLS VPN operations

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Implementing MPLS Layer 3 VPN Backbones

Overview

This lesson introduces the VRF table, the major data structure that is associated with MPLS VPN implementation on Cisco IOS, IOS XE, and IOS XR platforms. The lesson describes the MPLS VPN attributes that are associated with a VRF instance and explains the need for routing protocol contexts and the interaction of routing protocol contexts, VRFs, and MP-BGP.

Having a clear understanding of how information is exchanged using VRFs and routing protocol contexts will make it easier to configure VRFs in your network.

This lesson also explains how to configure VRF tables, listing the configuration tasks, syntax, and definitions of commands that are used to create VRFs. It provides an example of a VPN configuration.

Knowing how to configure and apply a VRF table onto a routing interface is important. Understanding the command syntax for the configurations that you want to deploy in your network is essential. This lesson provides you with the information that will enable you to succeed at such tasks.

Upon completing this lesson, you will be able to:

- Describe Layer 3 MPLS VPNs
- Describe the use of routing contexts on the PE router
- Describe the VPN-aware routing protocols
- Explain the VRF table
- Explain the propagation of the customer BGP routing information
- Explain the propagation of the customer non-BGP routing information
- Explain the VRF configuration
- Explain the RD configuration

https://t.me/learningnets

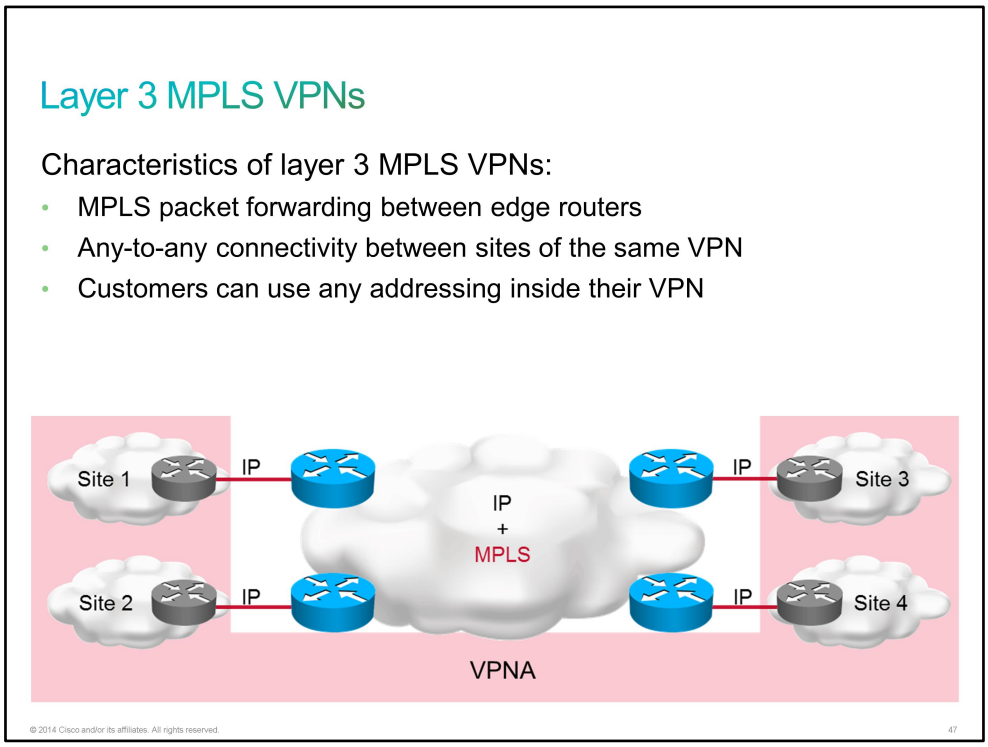
- Explain the import and export RT configuration
- Explain the MPLS VPN ID configuration
- Demonstrate a MPLS VPN configuration example
- Provide a Cisco IOS VRF common policies configuration example
- Describe the conditional route target export and import
- Describe Cisco IOS migration from old to new style VRF configuration
- Describe the VRF-Lite features and configurations
- Describe the VRF-Lite VRF and interface configuration
- Describe the VRF-Lite PE-CE connections using BGP
- Explain MP-BGP configurations
- Explain MP-BGP configuration tasks
- Explain BGP address families
- Explain BGP neighbors
- Explain MP-BGP community propagation
- Describe an example of the MP-BGP configuration
- Describe how to disable IPv4 route exchange
- Describe the VPN label allocation modes on Cisco IOS/XE
- Describe the VPN label allocation modes on Cisco IOS XR

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Layer 3 MPLS VPNs

The main characteristic of Layer 3 MPLS VPNs is that customers connect to the service provider via IP. They need to establish IP routing (static or dynamic) to exchange routing information between customer sites that belong to the same VPN.

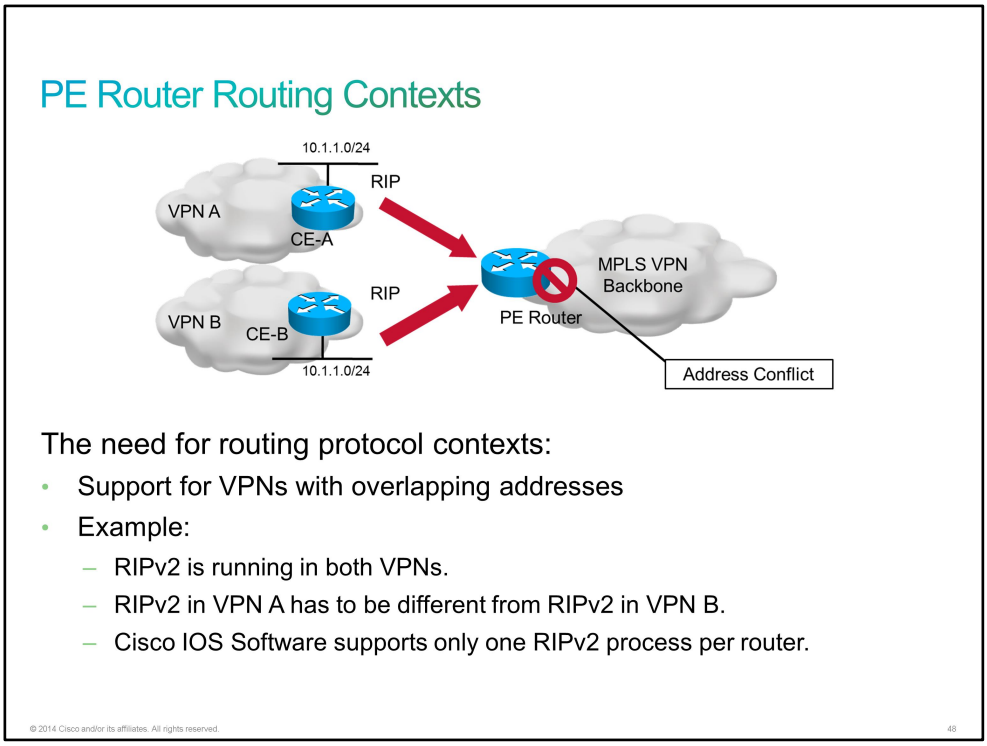


Because different customers might use the same private IP address ranges, the service provider cannot perform normal IP forwarding. Instead, service providers must use MPLS to ensure isolation in the data plane between packets belonging to different customers but potentially having the same IP addresses. Virtual routers (VRF instances) are used on service provider routers to isolate customer routing information. MPLS seamlessly provides any-to-any connectivity between sites that belong to the same VPN.

https://t.me/learningnets

PE Router Routing Contexts

Traditional Cisco IOS Software can support a number of different routing protocols. In some cases, even several completely isolated copies of the same routing protocol are supported. For example, several OSPF processes can be used.

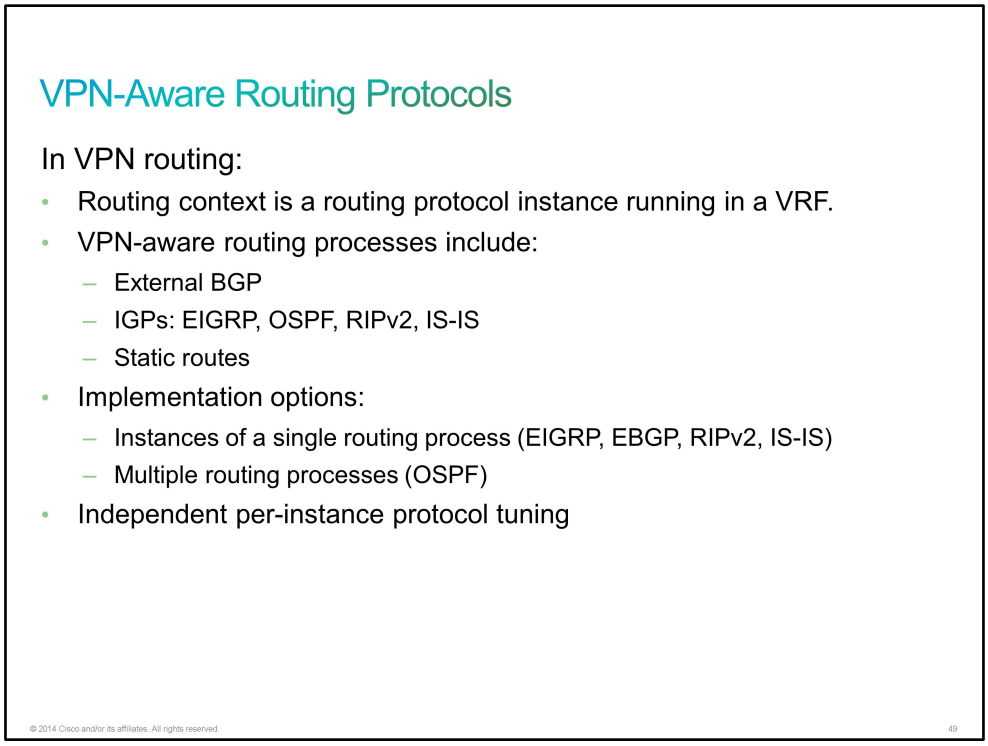


Forr several important routing protocols, such as RIPv2, EIGRP, IS-IS, and BGP, Cisco IOS Software supports only a single copy of the protocol running in the router. These protocols cannot be used directly between PE and CE routers in VPN environments because each VPN (or, more precisely, each VRF) needs a separate, isolated copy of the routing protocol to prevent undesired route leakage between VPNs. Furthermore, VPNs can use overlapping IP address spaces (for example, each VPN could use subnetworks of network 10.1.1.0/24) that would also lead to routing confusion if all VPNs shared the same copy of the routing protocol.

https://t.me/learningnets

VPN-Aware Routing Protocols

Routing contexts were introduced in Cisco IOS Software to support the need for separate isolated copies of VPN routing protocols. Routing contexts can be implemented as separate routing processes (in OSPF), similar to a traditional Cisco IOS Software implementation, or as separate isolated instances of the same routing protocol.

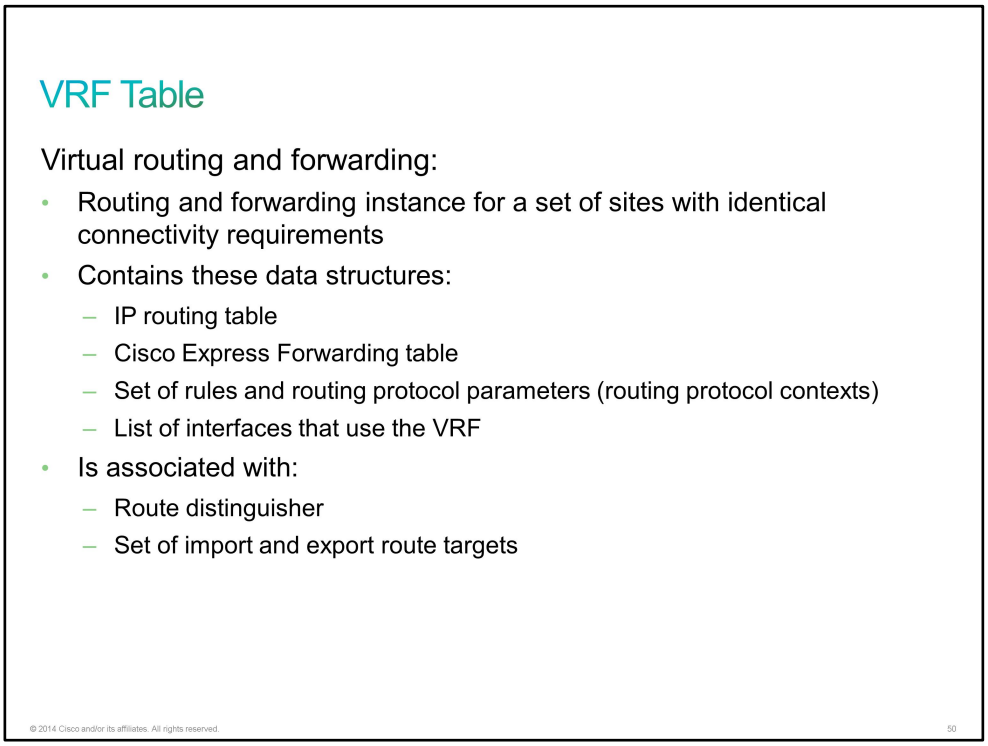


If routing contexts are implemented as instances of the same routing protocol, each instance contains its own independent routing protocol parameters. Examples would include networks over which the routing protocol is run, timers, authentication parameters, passive interfaces, and neighbors. This independence allows the network designer to have maximum flexibility in implementing routing protocols between PE and CE routers.

https://t.me/learningnets

VRF Table

The major data structure that is associated with MPLS VPN implementation on Cisco platforms is the VRF table.



The VRF data structure encompasses these constructs:

- An IP routing table that is identical in function to the global IP routing table in Cisco IOS Software
- A Cisco Express Forwarding table that is identical in function to the global Cisco Express Forwarding table (FIB)
- Specifications for routing protocols running inside the VRF instance

A VRF is a routing and forwarding instance that you can use for a single VPN site or for many sites connected to the same PE router, if and only if these sites share exactly the same connectivity requirements.

Other MPLS VPN attributes that are associated with a VRF table are as follows:

- The RD, that is prepended (for example, RD + IP address) to all routes that are exported from the VRF into the global VPNv4 (also called VPN IPv4) BGP table
- A set of export RTs, that are attached to any route that is exported from the VRF
- A set of import RTs, that are used to select VPNv4 routes that are to be imported into the VRF

https://t.me/learningnets

VRF Table (Cont.)

VRF table characteristics:

- Contains routes for a particular set of sites
- Analogous to the standard Cisco IOS Software routing table:
  - Supports the same set of mechanisms
- Interfaces assigned to VRFs:
  - Many interfaces per VRF
  - Physical interfaces, subinterfaces, or logical interfaces
  - Each interface assignable to only one VRF

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

The routes that are received from VRF routing protocol instances or from dedicated VRF routing processes are inserted into the IP routing table that is contained within the VRF. This IP routing table supports exactly the same set of mechanisms as the standard Cisco IOS Software routing table. These mechanisms include filter mechanisms (distribute lists or prefix lists) and interprotocol route-selection mechanisms (administrative distances).

The per-VRF FIB table is built from the per-VRF routing table. This table is used to forward all the packets that are received through the interfaces that are associated with the VRF. Any interface can be associated with a VRF, whether it is a physical interface, subinterface, or logical interface, as long as it supports Cisco Express Forwarding switching.

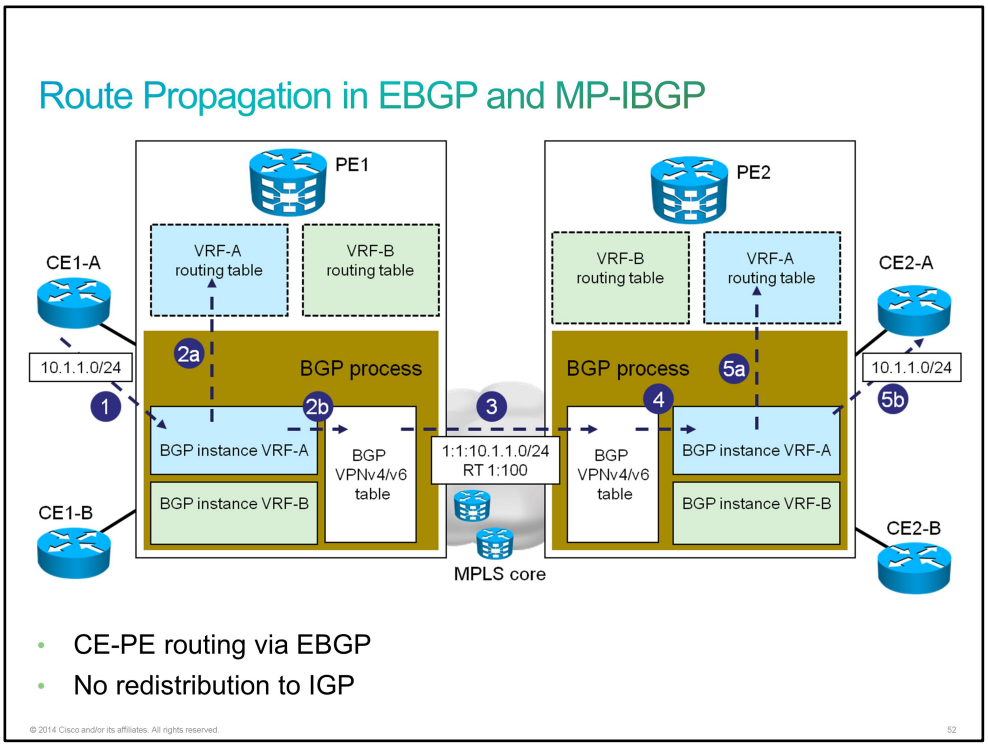
There is no limit to the number of interfaces that can be associated with one VRF (other than the number of interfaces that are supported by the router). However, each interface can be assigned to only one VRF because the router needs to uniquely identify the forwarding table to be used for packets that are received over an interface.

https://t.me/learningnets

Route Propagation in EBGP and MP-IBGP

This sample network contains two VPN customers. This example focuses on only a single PE router, that contains two VRFs—one for each customer. Each customer is connected to the PE router, that is running BGP. CE-BGP-A is the CE router for customer A and is associated with VRF-A (VPN-A). CE-BGP-B is the CE router for customer B and is associated with VRF-B (VPN-B). Both CE routers are using BGP for exchanging routes with the PE router.

This figure illustrates the interactions between VRF instances of routing processes, VRF routing tables, and the global VPNv4 BGP routing process.



There are five steps:

- **Step 1:** The BGP-speaking CE routers announce their networks via EBGP sessions to the PE router. In this example, you see a routing update for network 10.1.1.0/24.
- **Step 2a:** The PE router associates each BGP neighbor relationship with individual VRFs. The routes that are received from each VRF routing protocol instance are inserted into the IP routing table that is contained within that VRF. A per-VRF forwarding table, a FIB, is built from the per-VRF routing table and is used to forward all the packets that are received through the interfaces associated with the VRF.
- **Step 2b:** The BGP routes that are received from BGP-speaking CE routers are copied into the MP-BGP table for further propagation to other PE routers. This process is the export process.
- **Step 3:** The IP prefixes are prepended with the RD, and the set of RTs (extended BGP communities) that are configured as export RTs for the VRF is attached to the resulting VPNv4 route. In this example, the RD is 1:1 and the RT is 1:100.
- **Step 4:** The MP-BGP process in the recipient PE router receives the routes from the remote PEs. The routes are filtered, based on the RT attributes attached to them, and are attached to the proper per-VRF BGP table. In this example, VRF-A has defined import RT 1:100. The routing update is inserted into the VRF-A routing table.

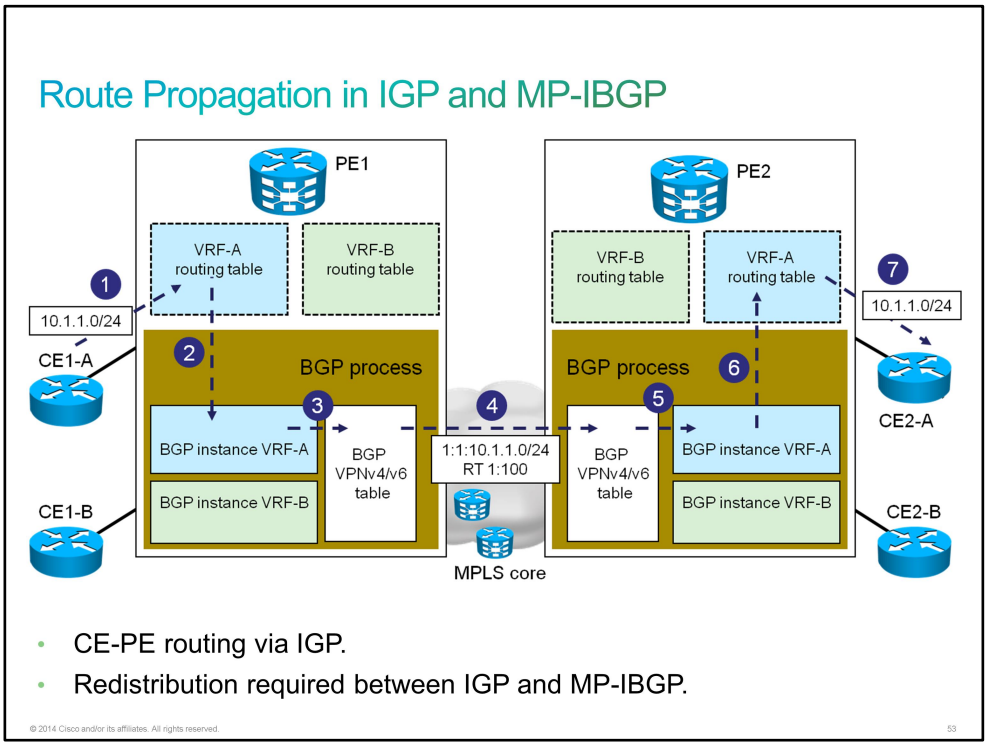
https://t.me/learningnets

- **Step 5a:** Upon additional checks, such as the next-hop reachability and comparison with IGP routes, the routes are inserted into the proper per-VRF routing table. The RD that was prepended by the ingress PE router is removed before the route is inserted into the per-VRF routing table.
- **Step 5b:** The MP-IBGP VPNv4 routes that are selected by the import RTs of a VRF are automatically propagated as 32-bit IPv4 routes to the BGP-speaking CE neighbors.

https://t.me/learningnets

Route Propagation in IGP and MP-IBGP

This example describes the route propagation process in an MPLS VPN implementation. The CE routers speak with the PE routers via an IGP. The VPN routes are transported between the PE routers via MP-BGP.



The routing propagation process occurs in these steps:

1. An IGP update is sent by the CE router and received by the VRF routing context of the IGP running on the PE router. In case of OSPF, a separate VRF-specific OSPF process would be responsible for the adjacency and route exchange with the CE router.
2. The per-VRF IGP routes are redistributed into the per-VRF instance of the BGP process to allow them to be propagated through the backbone MP-BGP process to other PE routers.

Note	If there is an overlap between an IGP update and an inbound EBGP update, the standard route-selection mechanism (administrative distance) is used in the per-VRF IP routing table, and the EBGP route takes precedence over the RIP route. EBGP precedence results from the fact that the administrative distance of EBGP routes (20) is better than the administrative distance of RIP routes (120).

3. The routes are internally marked as MP-BGP routes ready to be advertised via MP-BGP as VPNv4 or VPNv6 routes. A VPNv4 or VPNv6 route contains a RD and has one or more route targets associated with it.
4. MP-BGP is used to advertise the VPN routes (prefixed with the RD) as 96-bit VPNv4 or 160-bit VPNv6 routes to other PE routers.

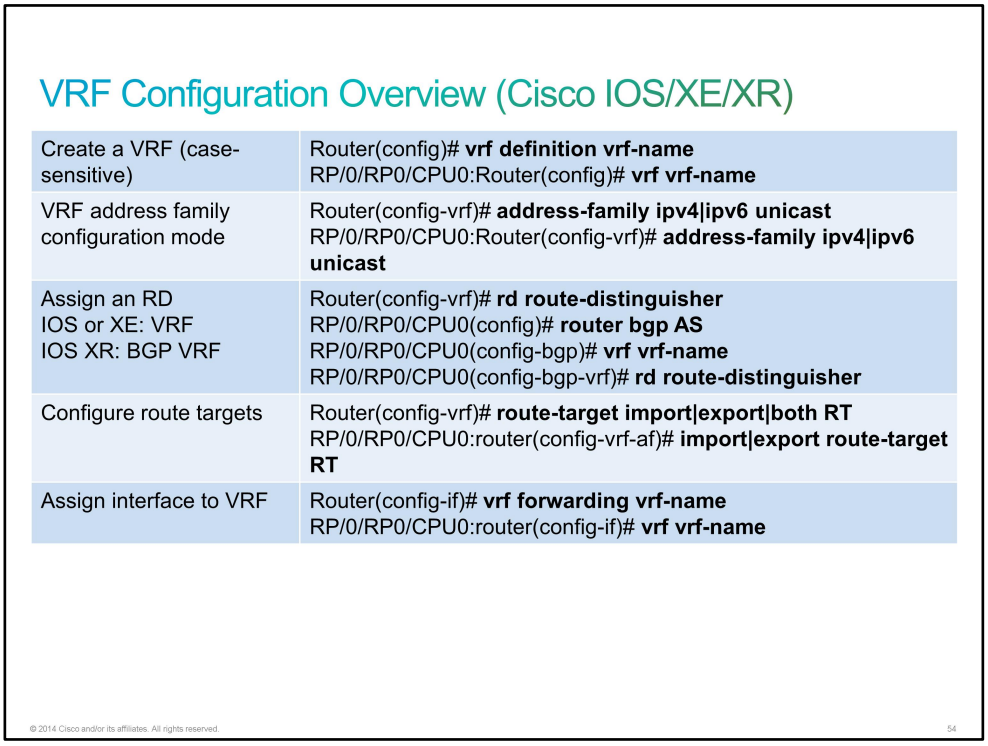
https://t.me/learningnets

5. The routes are filtered, based on the RT attributes attached to them, and are attached to the proper per-VRF BGP table.
6. Upon additional checks, such as the next hop-reachability and comparison with IGP routes, the routes are inserted into the proper per-VRF routing table. The RD that was prepended by the ingress PE router is removed before the route is inserted into the per-VRF routing table.
7. Although the routes have been inserted in the per-VRF IP routing table, they are not propagated to IGP-speaking CE routers automatically. To propagate these MP-IBGP routes to the IGP-speaking CE routers, you must enable redistribution between the per-VRF BGP instance and the per-VRF IGP instance. Then, the IGP process announces these routes to IGP-speaking CE routers, thus achieving transparent end-to-end connectivity between the CE routers.

https://t.me/learningnets

VRF Configuration Overview (Cisco IOS/XE/XR)

The VRF configuration differs between Cisco IOS/XE and Cisco IOS XR platforms.



Configuring a VRF table and starting deployment of an MPLS VPN service consists of these steps:

1. Create a new VRF table. VRF names are case-sensitive. VRF names have only local significance.
2. Assign a RD . In Cisco IOS and IOS XE you assign the route distinguisher to a VRF. In IOS XR you assign the RD in the BGP VRF configuration mode. A VRF is not operational unless you configure an RD. You can use the ASN:nn or A.B.C.D:nn format for RD. Each VRF in a PE router must have a unique RD.
3. You can enter the VRF IPv4 or IPv6 unicast address family configuration using the **address-family ipv4|ipv6 unicast** command. Address families are used within VRF configuration mode to control import and export policies.
4. Specify import and export RTs for the VRF.
5. Assign interfaces to the VRF. When an interface is configured with a particular VRF, its IP address is removed from the interface and from the global routing table. This action is based on the assumption that the address is not valid across multiple routing tables and that the address should be reconfigured after the interface is associated to a VRF.

Note	You must assign a unique RD to every VRF created in a PE router. The same RD might be used in multiple PE routers, based on customer connectivity requirements. The same RD should be used on all PE routers for simple VPN service.

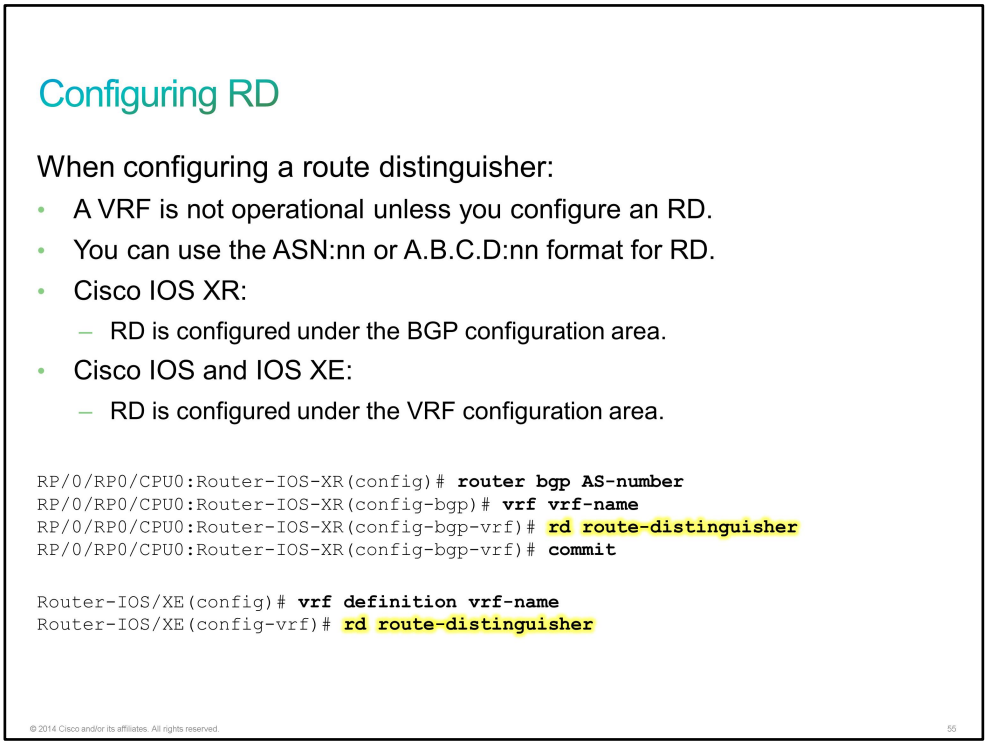
https://t.me/learningnets

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

https://t.me/learningnets

Configuring RD

You need to configure a route distinguisher for a VRF to become operational.



To create routing and forwarding tables for a VRF on the Cisco IOS and IOS XE platforms, use the **rd** command in VRF configuration submode: **rd** *route-distinguisher*. The *route-distinguisher* parameter adds an 8-byte value to an IPv4 prefix to create a VPNv4 prefix.

The RD can be specified in one of these two formats:

- 16-bit AS number followed by a 32-bit decimal number (ASN:nn)
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn)

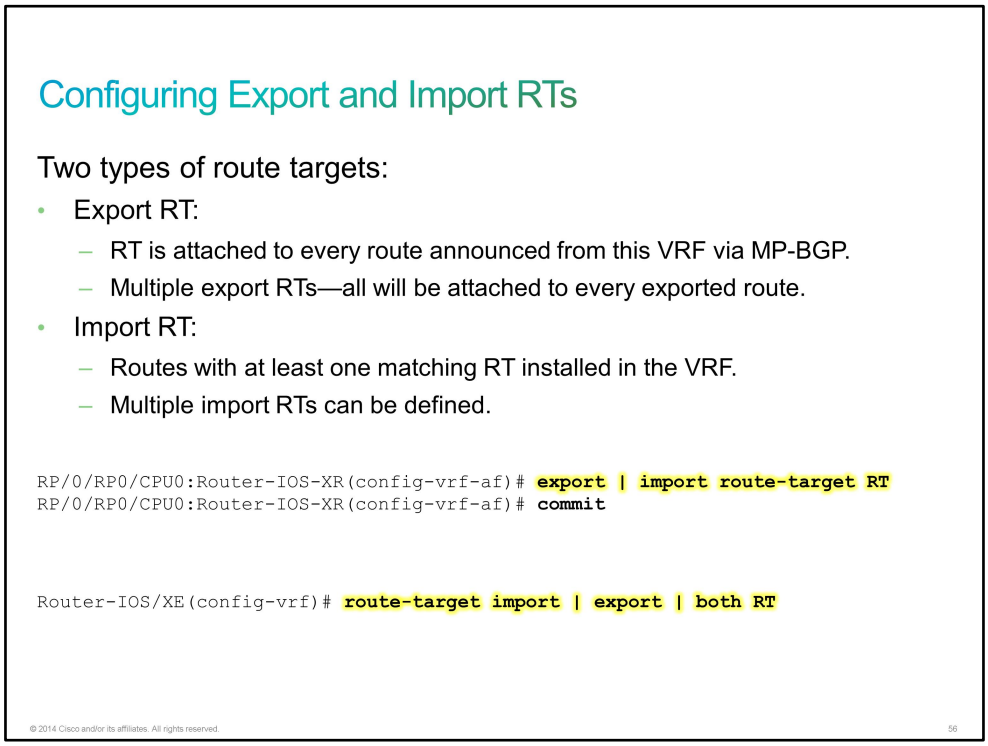
There is no default for this command. An RD must be configured for a VRF table to be functional.

To create routing and forwarding tables for a VRF on a Cisco IOS XR operating system, use the **rd** command in the BGP configuration area in VRF configuration submode.

https://t.me/learningnets

Configuring Export and Import RTs

To create an RT extended community for a VRF on the Cisco IOS and IOS XE platforms, use the **route-target** command in VRF submode.



The **route-target** command supports these options:

- **route-target {import | export | both}** *route-target-ext-community*
- **no route-target {import | export | both}** *route-target-ext-community*

This table describes the parameters for the **route-target** command.

Syntax Description

Parameter	Description
import	Imports routing information from the target VPN extended community
export	Exports routing information to the target VPN extended community
both	Sets the value to be used by both the import and export process to the value that is indicated in the *route-target-ext-community* field
route-target-ext-community	Adds the route target extended community attributes to the VRF list of import, export, or both (import and export) route target extended communities

https://t.me/learningnets

Similar to RDs, RTs can be specified in one of these two formats:

- 16-bit AS number followed by a 32-bit decimal number (ASN:nn).
- 32-bit IP address followed by a 16-bit decimal number (A.B.C.D:nn).

A VRF has no RT extended community attributes associated with it until they are specified by the **route-target** command.

Whenever an RT is both an import RT and an export RT for a VRF, you can use the **route-target both** command to simplify the configuration.

The Cisco IOS XR **export route-target** command associates the local VPN with an RT. When the route is advertised to other PE routers, the export RT is sent along with the route as an extended community.

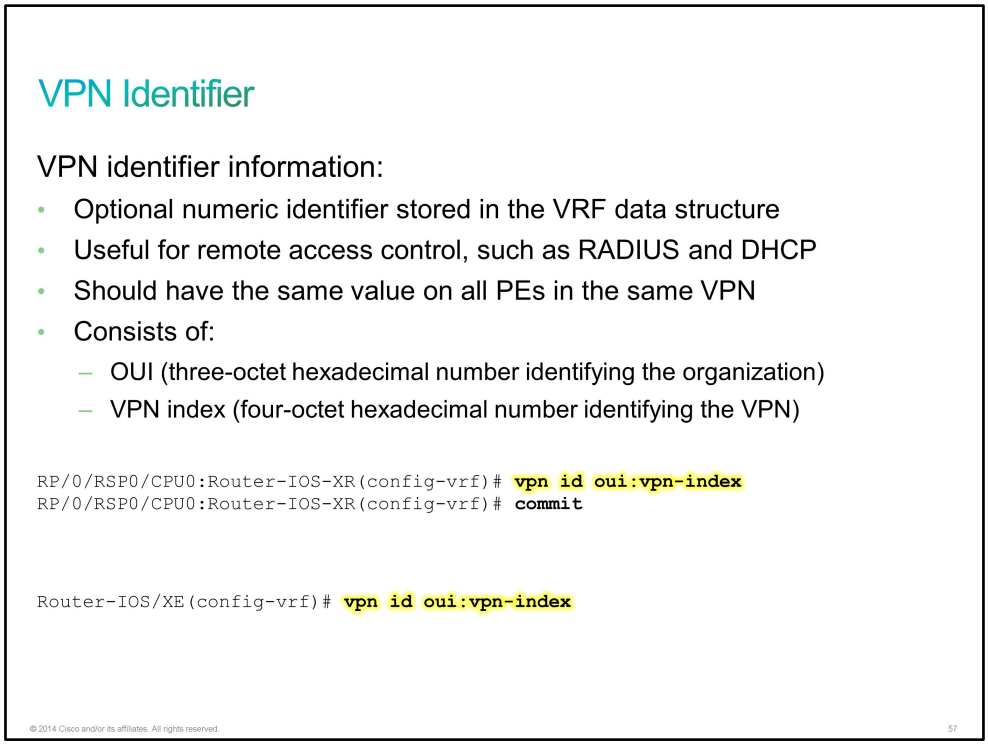
The **import route-target** command allows exported VPN routes to be imported into the VPN, if one of the RTs of the exported route matches one of the local VPN import RTs.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

VPN Identifier

The MPLS VPN ID is an optional feature that allows you to identify VPNs by a VPN identification number. A VPN ID is useful for remote access applications, such as RADIUS and DHCP that can use the MPLS VPN ID to identify a VPN. RADIUS can use the VPN ID to assign dial-in users to the proper VPN, based on the authentication information of each user.



Note	You can use a VRF name (a unique ASCII string) to reference a specific VPN that is configured in the router. Alternatively, you can use a VPN ID to identify a particular VPN that is configured in the router. The VPN name is not affected by the VPN ID configuration.

The VPN ID is stored in the corresponding VRF structure for the VPN. To ensure that the VPN has a consistent VPN ID, assign the same VPN ID to all the routers in the service provider network that service that VPN. The MPLS VPN ID feature is not used to control the distribution of routing information or to associate IP addresses with MPLS VPN ID numbers in routing updates.

You can configure all the PE routers that belong to the same VPN with the same VPN ID. Make sure that the VPN ID is unique to the service provider network.

Each VPN ID that is defined by RFC 2685 consists of these elements:

- An OUI, a three-octet hexadecimal number that is assigned by the IEEE
- A VPN index, a four-octet hexadecimal number that identifies the VPN within the company.

To assign a VPN ID to a VRF, use the **vpn id** command in VRF configuration submode.

https://t.me/learningnets

This table describes the parameters for the **vpn id** command.

Syntax Description

Parameter	Description
oui	Identifies the OUI, which is restricted to three octets and is followed by a colon
vpn-index	Identifies the VPN within the company and is restricted to four octets

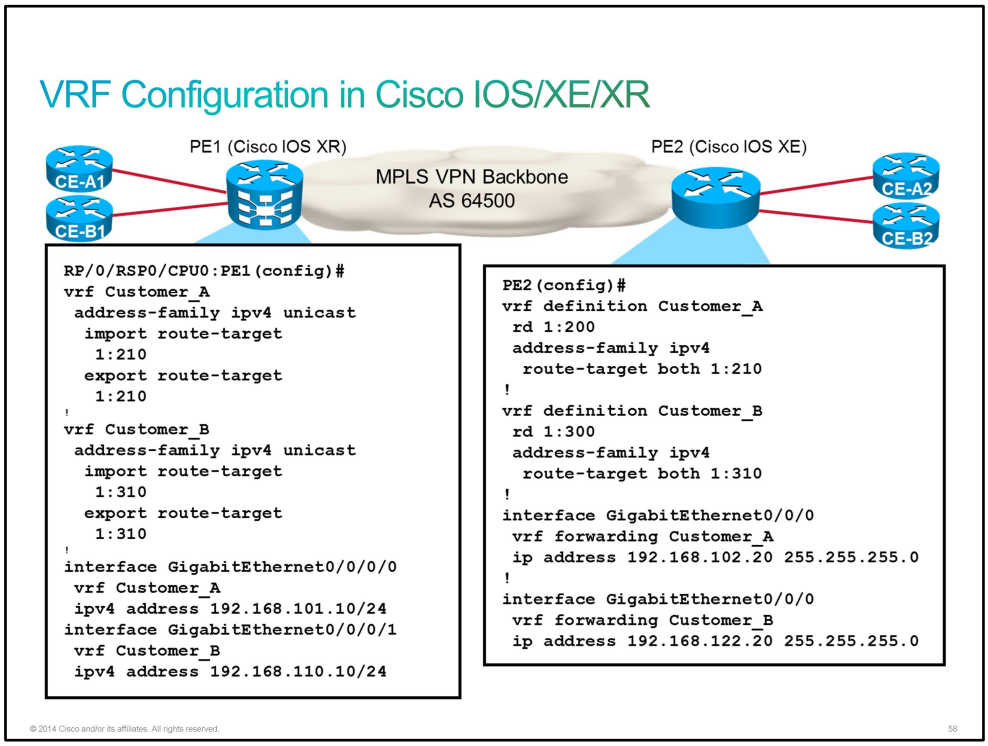
Each VRF configured in a PE router can have a VPN ID configured. Configure all the PE routers that belong to the same VPN with the same VPN ID. The VPN ID should be unique within the service provider network.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

VRF Configuration in Cisco IOS/XE/XR

To illustrate the use of MPLS VPN configuration commands, you can look at a configuration of the PE routers in a sample network.



The figure shows configuration of the PE routers in a sample network with two VPN customers. The PE1 router is running Cisco IOS or IOS XE Software. The PE2 router is running Cisco IOS XR Software.

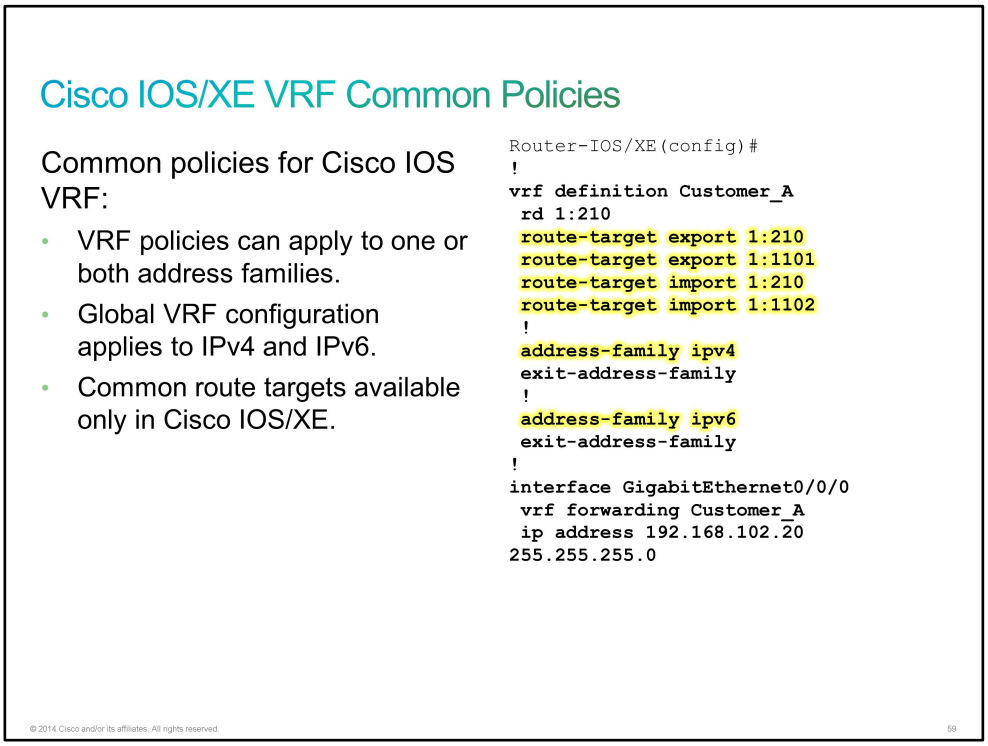
The configuration steps that you perform on the PE router are as follows:

1. Configure VRFs for customer A and customer B.
2. Assign RDs and RTs to the VRFs. In Cisco IOS and IOS XE you can use the **both** keyword to declare a route target as both import and export.
3. Only one RD per customer is required on all PE routers in this MPLS VPN backbone, because these customers require only simple VPN connectivity.
4. On Cisco IOS XR, to configure the RD for the VRF, use the **rd** command in the BGP VRF configuration submode (not shown in the slide).
5. Assign PE-CE interfaces to individual VRFs.

https://t.me/learningnets

Cisco IOS/XE VRF Common Policies

The VRF configuration on Cisco IOS and IOS XE routers supports allows you to choose between common policies, noncommon policies, or a mix of them.

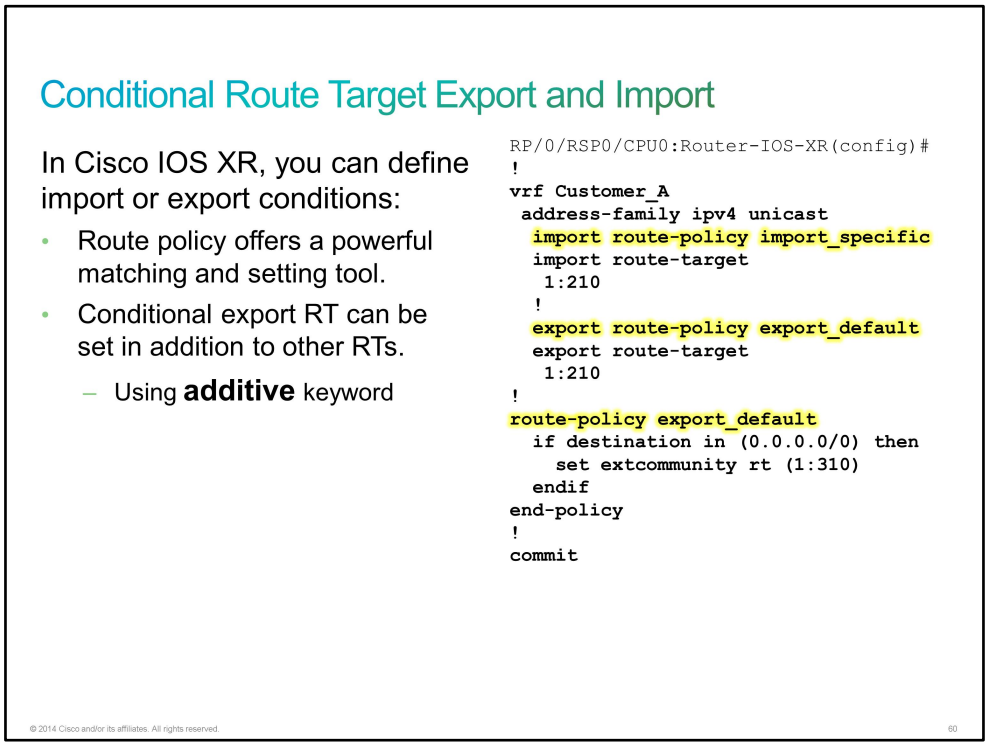


If you configure the route targets in the global VRF configuration, outside of the address family submode, the policies will apply to both IPv4 and IPv6. If you want the route targets to be applied only to a specific address family, configure them in the appropriate submode.

https://t.me/learningnets

Conditional Route Target Export and Import

Cisco IOS XR enables you to define more specific conditions when exporting and importing routes from or into a VRF.



You can use the **export route-policy** command to define the conditions that allow specified routes to be tagged with specified route targets. The **import route-policy** command defines the conditions that allow specified routes to be imported into the VRF instance if the routes are tagged with specified route targets.

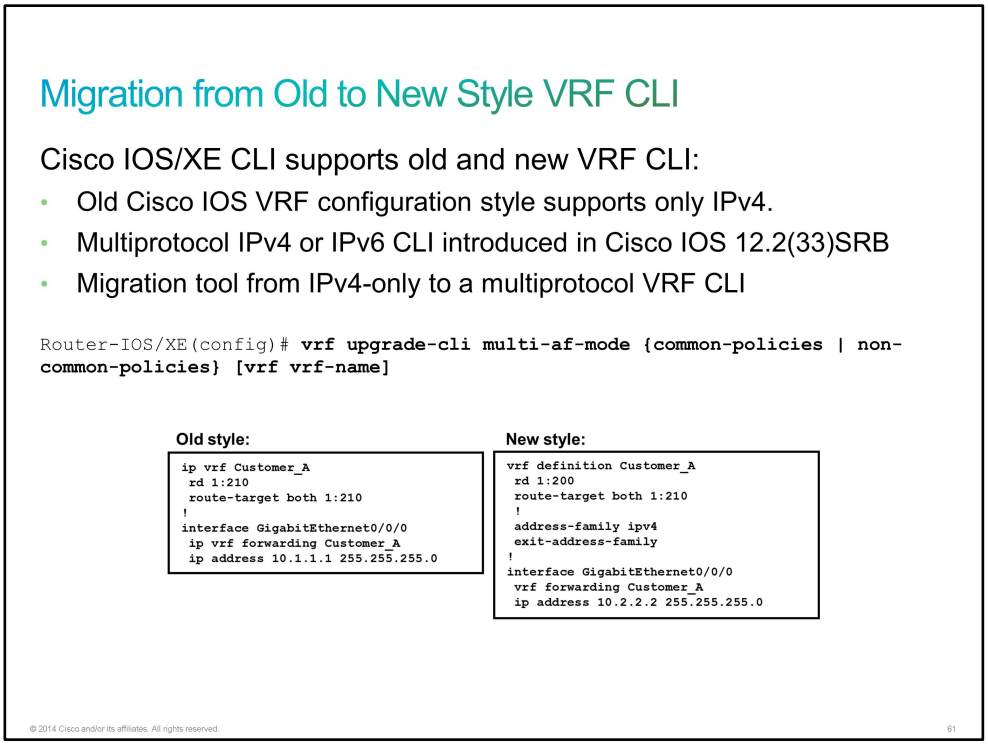
Cisco IOS XR does not use route maps but a more powerful tool for matching routes and setting their parameters. This programmatic approach is called RPL. While IOS route maps operate as a series of statements which are executed sequentially, route-policies not only operate sequentially but provide the ability to invoke other route-policies much like a 'C'-program is able to call separately defined functions. This enables you to create hierarchical policies. In addition, route policies are 'compiled' into a run-time executable portion of code.

In this example, the **route policy import_specific**, not shown here, is used to match routes to be imported into the Customer_A VRF. The matched updates will be imported in addition to the updates that carry RT 1:210. The route policy export_default is used to set the RT 1:310 for the default route when exporting it from the Customer_A VRF. When defining the **export route-policy** behavior, you can use the *additive* keyword to add a route target to other export route targets.

https://t.me/learningnets

Migration from Old to New Style VRF CLI

The multiprotocol VRF configuration allows you to define multiple address families under the same VRF. A given VRF, identified by its name and a set of policies, can apply to both an IPv4 VPN and an IPv6 VPN at the same time. This VRF can be activated on a given interface, even though the routing and forwarding tables are different for the IPv4 and IPv6 protocols.



Cisco IOS releases earlier than 12.2(33)SRB supported only the single-protocol IPv4-only VRF. You created a single-protocol VRF by entering the **ip vrf** command. To activate the single-protocol VRF on an interface, you entered the **ip vrf forwarding** (interface configuration) command. After the introduction of the VRF CLI for IPv4 and IPv6 VPNs feature in Cisco IOS Release 12.2(33)SRB, you can create a multiprotocol VRF by entering the **vrf definition** command. To activate the multiprotocol VRF on an interface, you enter the **vrf forwarding** command.

Cisco IOS offers a migration tool that upgrades a VRF instance or all VRFs configured on the router to support multiple address families under the same VRF. The **vrf upgrade-cli multi-af-mode** {**common-policies** | **non-common-policies**} [**vrf vrf-name**] command is issued in the configuration mode and uses these parameters:

- The **multi-af-mode** keyword specifies an upgrade of a single-protocol VRF or all VRFs to a multiprotocol VRF that supports configuration of multiple address families.
- The **common-policies** keyword specifies copying the route-target policies to the common part of the VRF configuration so that the policies apply to all address families configured in the multiple address family VRF.
- The **non-common-policies** keyword specifies copying the route-target policies to the IPv4 address family part of the VRF configuration so that the policies apply only to IPv4.
- The **vrf** keyword specifies a VRF for the upgrade to a multiple address family VRF configuration.

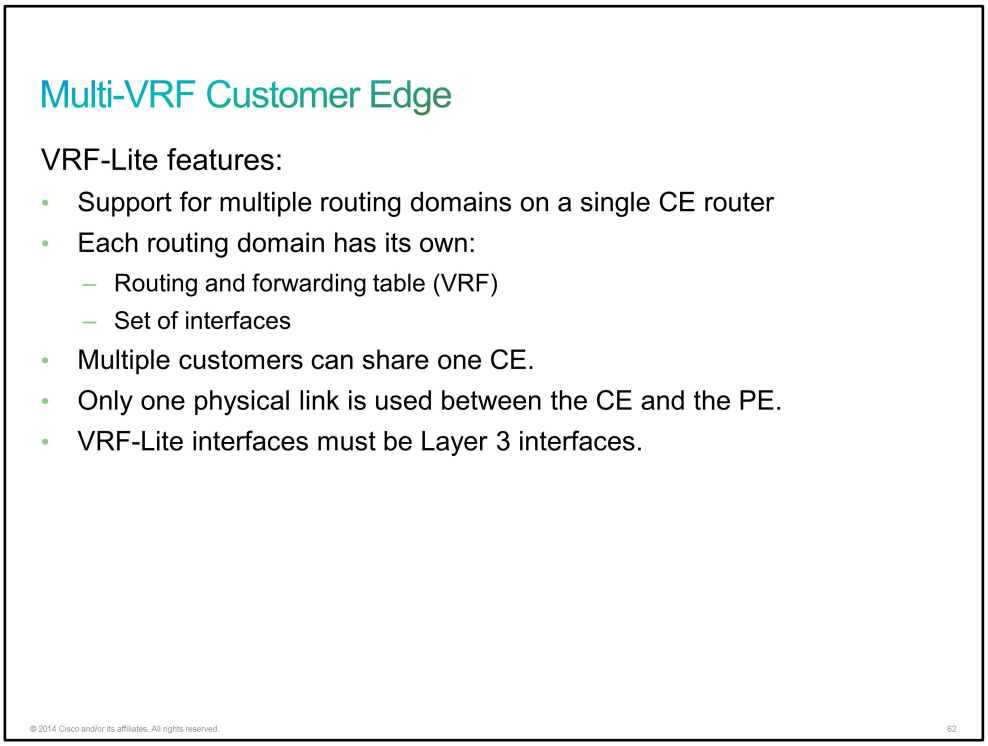
https://t.me/learningnets

- The **vrf-name** argument is the name of the single-protocol VRF to upgrade to a multiple address family VRF configuration.

https://t.me/learningnets

Multi-VRF Customer Edge

VRF-Liteprovides the ability to configure and maintain more than one VRF instance within the same CE router.



VRF-Lite uses input interfaces to distinguish routes for different VPNs and forms VRF tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports, or logical, such as VLAN SVIs. However, a Layer 3 interface cannot belong to more than one VRF at any time. The VRF-Lite feature thus allows an operator to support two or more routing domains on a CE router, with each routing domain having its own set of interfaces and its own set of routing and forwarding tables.

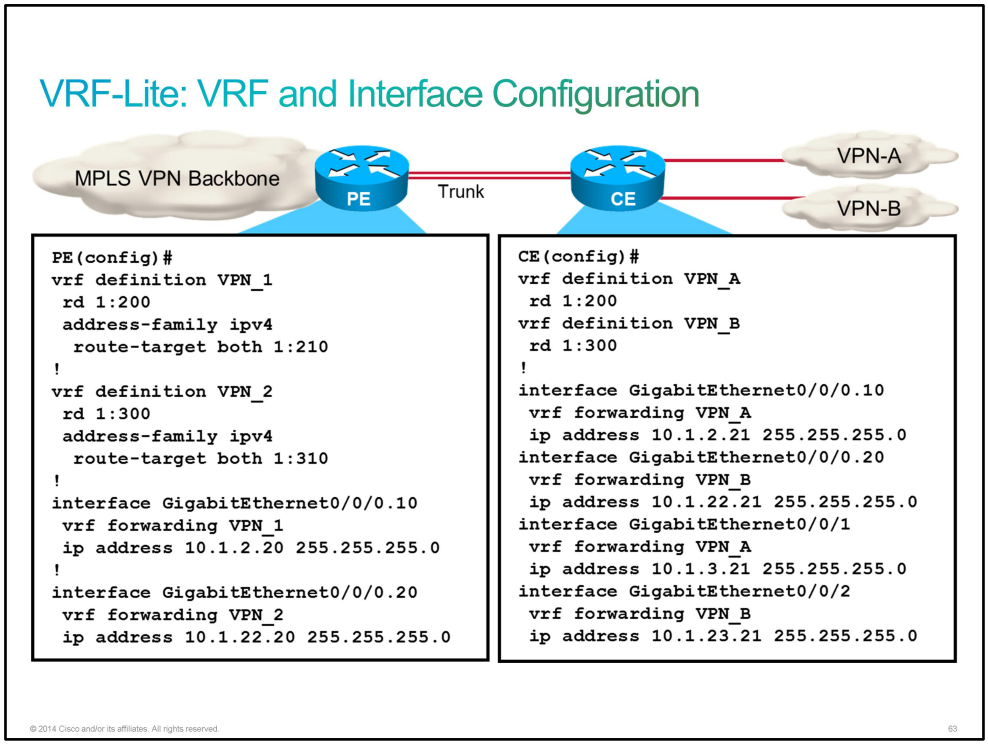
VRF-Lite includes these devices:

- **Customer edge devices:** CE devices provide customer access to the service provider network over a data link to one or more PE routers. The CE device advertises the local routes of the site to the PE router and learns the remote VPN routes from it. A Cisco Catalyst 4500 Series Switch can be a CE.
- **Provider edge routers:** PE routers exchange routing information with CE devices by using static routing or a routing protocol such as BGP,RIPv1, or RIPv2. The PE router is only required to maintain VPN routes for the VPNs to which it is directly attached, eliminating the need for the PE router to maintain all of the service provider VPN routes. Each PE router maintains a VRF for each of its directly connected sites. Multiple interfaces on a PE router can be associated with a single VRF if all of these sites participate in the same VPN. Each VPN is mapped to a specified VRF. After learning local VPN routes from CE devices, a PE router exchanges VPN routing information with other PE routers by using IBGP.

https://t.me/learningnets

VRF-Lite: VRF and Interface Configuration

With VRF-Lite, multiple customers can share one CE, and only one physical link is used between the CE and the PE. The shared CE maintains separate VRF tables for each customer and switches or routes packets for each customer, based on its own routing table. VRF-Lite extends limited PE functionality to a CE device, giving it the ability to maintain separate VRF tables to extend the privacy and security of a VPN to the branch office.



To understand the VRF-Lite configuration, you can look at a configuration of the CE and PE routers in a sample network. First, VRFs must be configured on both the PE and the CE routers. Additionally, you must specify the Layer 3 interface to be associated with the VRF and associate the VRF with

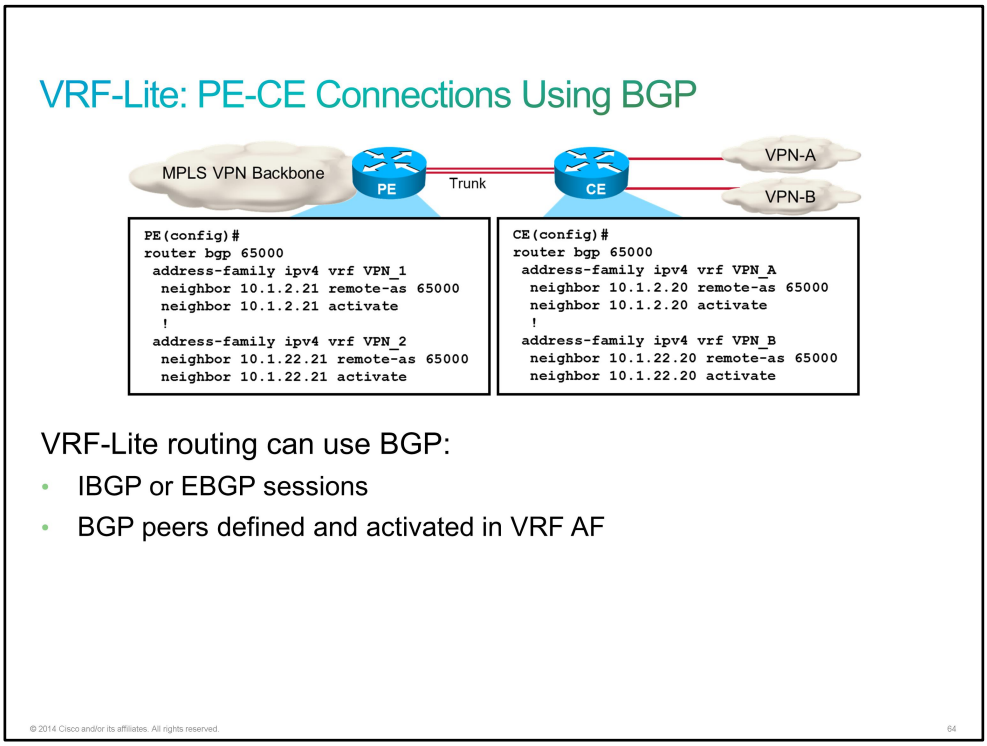
You can implement VRF-Lites on all Cisco routers: Cisco IOS, IOS XE, and IOS XR.

Cisco IOS XR Software Release 4.1.1 introduced the Large VRF support. This Cisco ASR 9000 Series Aggregation Services Routers VRF capability is enhanced to support an increased (up to global FIBsize) number of prefixes per VRF ("big" VRF) for up to 16 user-configured VRF IDs.

https://t.me/learningnets

VRF-Lite: PE-CE Connections Using BGP

BGP is commonly used as the PE-CE routing protocol, as it allows both the provider and customer to tightly control the route exchange.



Most routing protocols can be used between the CE and the PE: BGP, OSPF, EIGRP, RIP, and static routing. However, EBGP is recommended for the following reasons:

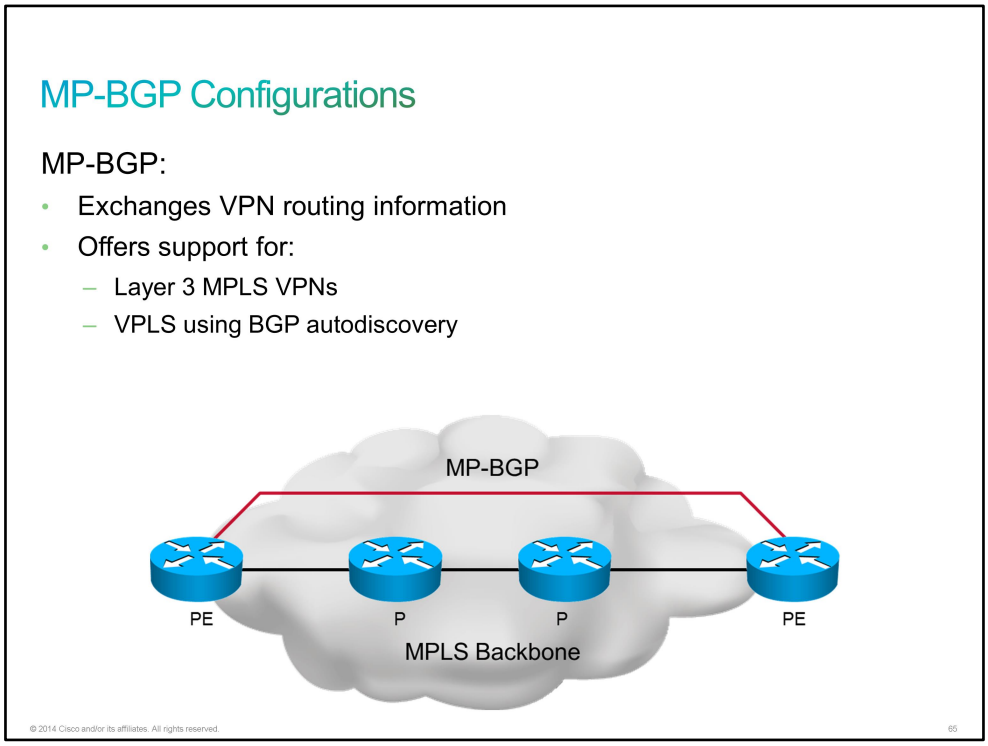
- BGP does not require more than one algorithm to communicate with a multitude of CEs.
- BGP is designed to pass routing information between systems that are run by different administrations.
- BGP makes it easy to pass attributes of the routes to the CE.

Furthermore, when BGP is used as the routing protocol, it can also be used to manage the MPLS label exchange between the PE and CE devices. By contrast, if OSPF, EIGRP, RIP, or static routing is used, LDP must be used to signal labels.

https://t.me/learningnets

MP-BGP Configurations

VPNs based on MPLS require an additional VPN label to distinguish between potentially overlapping prefixes belonging to different VPNs. MP-BGP is BGP version 4 with additional attributes to support the exchange of VPN prefixes.



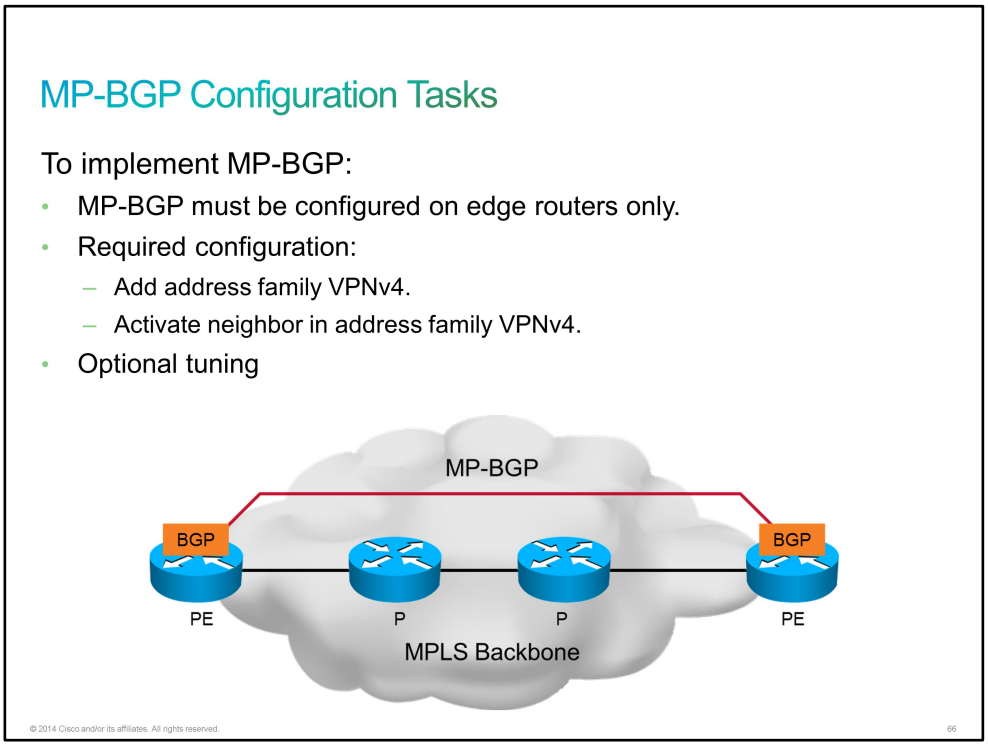
VPLS can also be implemented using the BGP autodiscovery feature to simplify the management of VPLS.

The figure shows an end-to-end MP-IBGP session. This figure is a simplified representation of the BGP capability to propagate VPN routing information between edge LSRs. In real environments with many more PE routers, a route reflector would be used between the edge routers, although that addition would not alter the operation significantly.

https://t.me/learningnets

MP-BGP Configuration Tasks

The MP-BGP extension is used in the MPLS world to relay VPN information between two edge routers. The RD is a 64-bit value that is used to mark prefixes and to separate different customers.



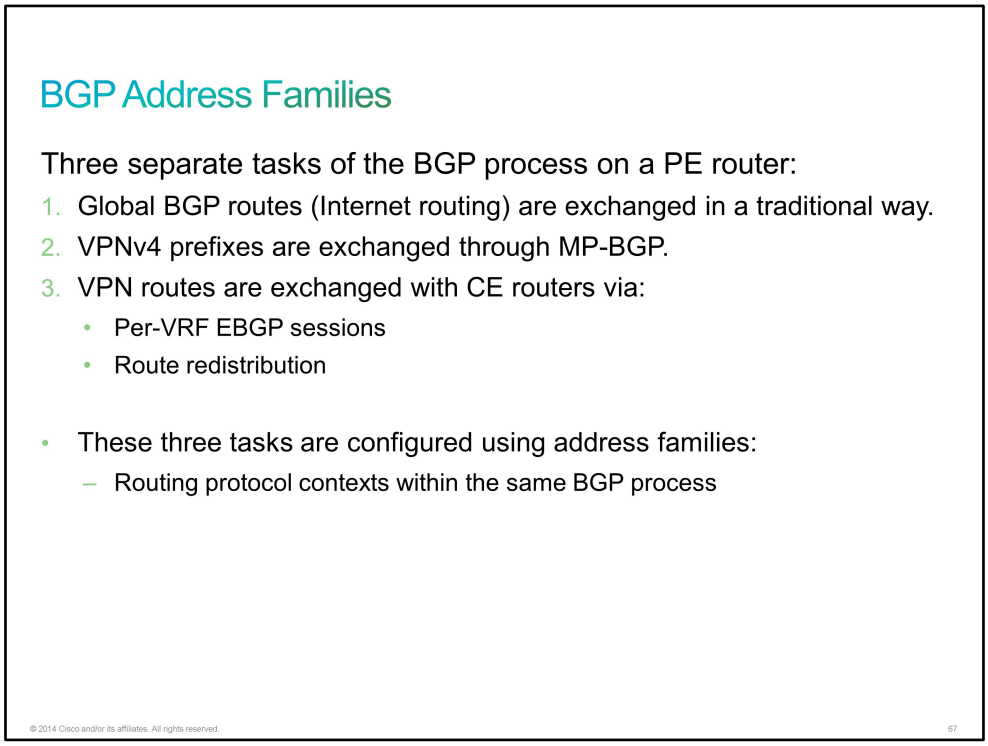
VPN support in BGP is enabled by configuring a VPNv4 address family. This allows MP-BGP neighbor sessions to be established independently from existing IPv4 BGP sessions. These VPNv4 adjacencies are used to relay VPN prefixes together with 64-bit extended communities, where the RT value is stored. The total length of the VPNv4 address is thus 96 bits.

Configuring the VPNv4 address family and activating neighbors in it is the minimum required configuration. Optionally, fine-tuning can be performed by adjusting the BGP timers.

https://t.me/learningnets

BGP Address Families

Independently from the MPLS VPN architecture, the PE router can use BGP IPv4 route updates to receive and propagate Internet routes in situations where the PE routers are also used to provide Internet connectivity to customers.



The MPLS VPN architecture uses the BGP routing protocol in these two ways:

- VPNv4 routes are propagated across an MPLS VPN backbone using MP-BGP between the PE routers.
- BGP can be used as the PE-CE routing protocol to exchange VPN routes between the PE routers and the CE routers.

All three route-exchange mechanisms take place in one BGP process (because only one BGP process can be configured per router). The routing protocol contexts (called address families from the router configuration perspective) are used to configure all three independent route-exchange mechanisms.

To configure the BGP routing process, use the **router bgp** command in global configuration mode. To remove a routing process, use the **no** form of this command.

Command options are listed here:

- **router bgp** *as-number*
- **no router bgp** *as-number*

Use the **address-family** command in router configuration mode to select the routing context that you would like to configure:

https://t.me/learningnets

- Internet routing (global IP routing table) is the default address family that you configure when you start configuring the BGP routing process.
- To configure MP-BGP sessions between the PE routers, use the **address-family vpnv4** command.
- To configure BGP between the PE routers and the CE routers within individual VRF tables, use the **address-family ipv4** *vrf vrf-name* command.

To enter address-family submode for configuring routing protocols, such as BGP, RIP, and static routing, use the **address-family** command in global configuration mode. To disable address-family submode for configuring routing protocols, use the **no** form of this command.

- VPNv4 unicast: **address-family vpnv4 [unicast]**

  This command configures sessions that carry customer VPNv4 prefixes, each of which has been made globally unique by adding an 8-byte RD.
- IPv4 unicast: **address-family ipv4 [unicast]**

  This command configures sessions that carry standard IPv4 address prefixes.
- IPv4 unicast: **address-family ipv4 [unicast] vrf** *vrf-name*

  This command specifies the name of a VPN VRF to associate with submode commands.

Similar to Cisco IOS and IOS XE Software, on Cisco IOS XR Software use the **router bgp** command in global configuration mode.

The VPNv4 address family is configured in the BGP section using the **address-family vpnv4 unicast** command. Afterwards, it will be applied in the neighbor configuration block.

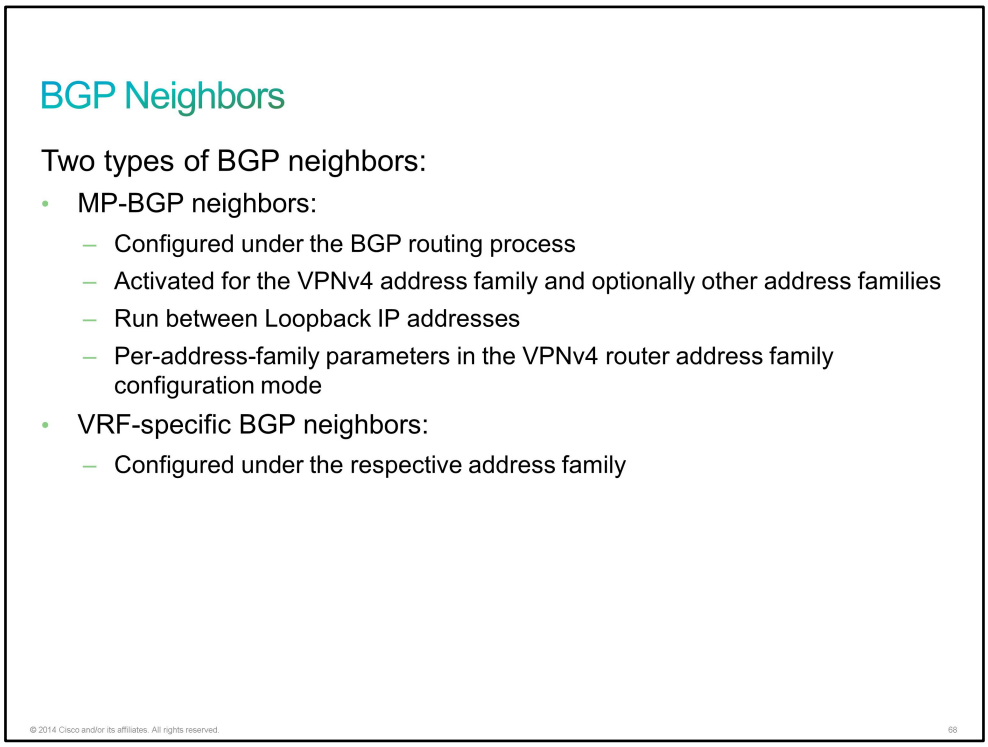
 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Neighbors

MPLS VPN architecture defines the following two types of BGP neighbors:

- Global BGP neighbors (other PE routers) with which the PE router can exchange multiple types of routes. (These neighbors are defined in the global BGP definition and need to be activated only for individual address families.)
- Per-VRF BGP neighbors. (These neighbors are the CE routers.)



The initial commands that are needed to configure an MP-IBGP session between PE routers are as follows:

- The **neighbor** *ip-address* **remote-as** *as-number* command configures the neighboring PE router.
- The **neighbor** *ip-address* **update-source** i*nterface-type interface-number* command configures the source address that is used for the TCP session carrying BGP updates and the IP address that is used as the BGP next hop for VPNv4 routes.
- The **address-family vpnv4** command allows you to enter VPNv4 configuration mode, where additional VPNv4-specific parameters must be configured on the BGP neighbor.

https://t.me/learningnets

After you define the remote PE router as a global BGP neighbor, you must activate it for VPNv4 route exchange. To enable the exchange of information with a BGP neighboring router, use the **neighbor activate** command in router configuration mode. The exchange of addresses with neighbors is enabled by default for the IPv4 address family. For all other address families, address exchange is disabled by default. You can explicitly activate the default command by using the appropriate address family submode.

To enable next-hop processing of BGP updates on the router, use the **neighbor next-hop-self** command in router configuration mode. This command is useful in unmeshed networks (such as Frame Relay or X.25) where BGP neighbors might not have direct access to all other neighbors on the same IP subnet. If you specify a BGP peer group by using the peer-group-name argument, all the members of the peer group inherit the characteristic that is configured with this command. Specifying the command with an IP address overrides the value that is inherited from the peer group.

To configure an MP-IBGP neighbor on devices running Cisco IOS XR Software, enter BGP configuration mode using the command **router bgp** *as-number*.

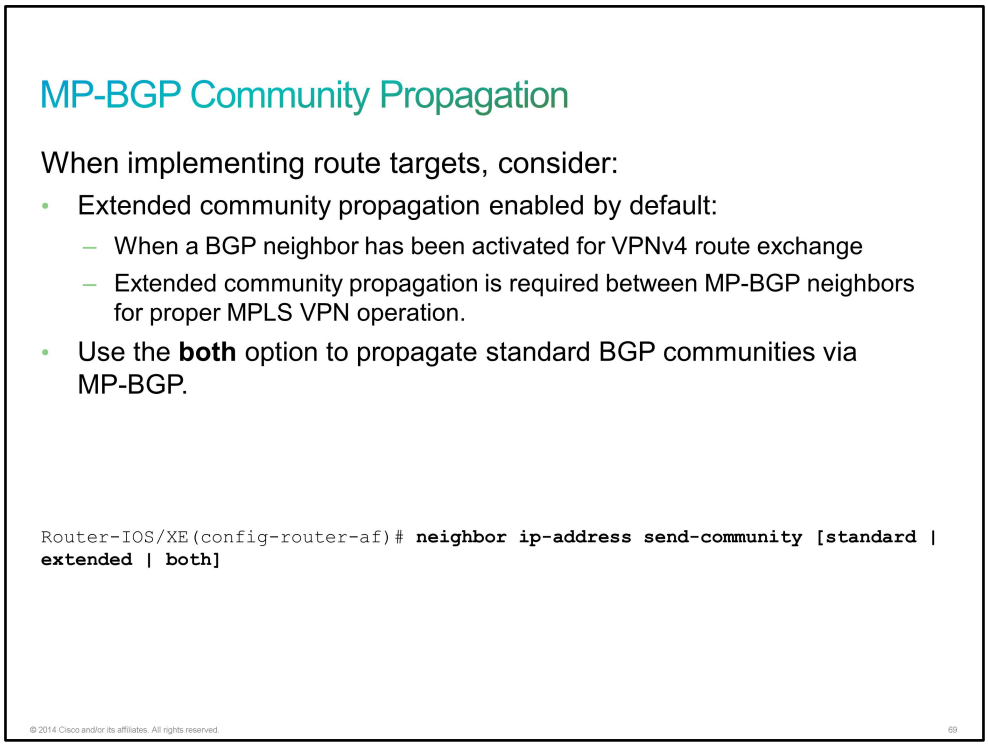
To add a new BGP neighbor, use the command **neighbor** *ip-address* **remote-as** *as-number*.

In each neighbor configuration area, enable the neighbor for the specific address family.

https://t.me/learningnets

MP-BGP Community Propagation

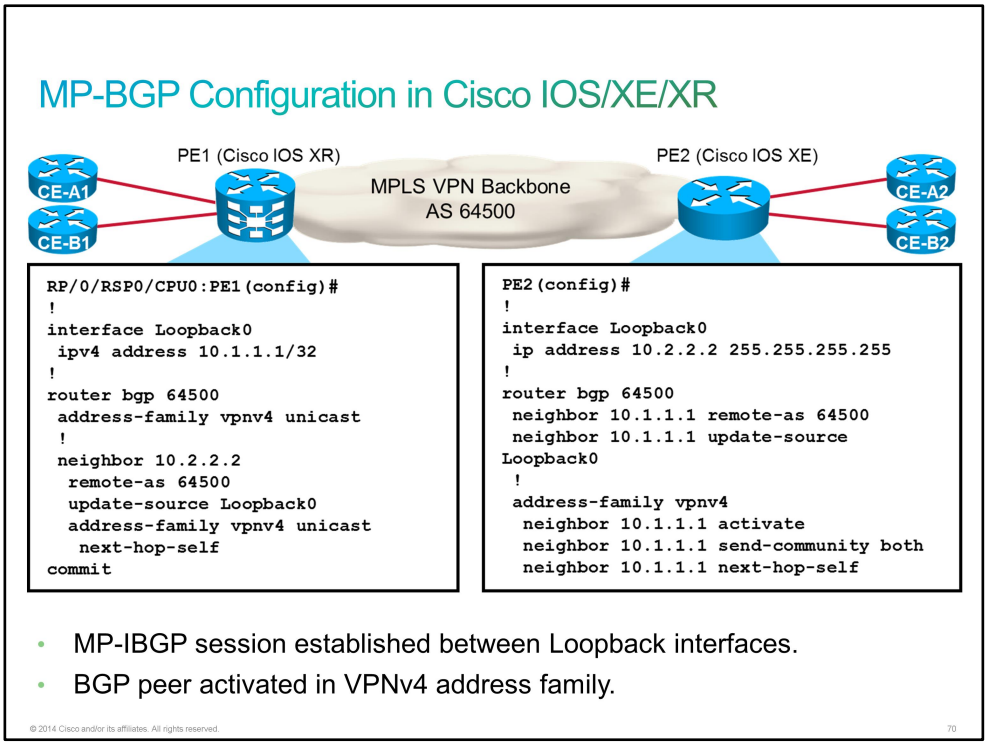
MPLS VPN architecture introduced the extended community BGP attribute. BGP still supports the standard community attribute, which has not been superseded by extended communities. The default community propagation behavior for standard BGP communities has not changed. Community propagation still must be configured manually. Extended BGP communities are propagated by default because their propagation is mandatory for successful MPLS VPN operation.



The **neighbor send-community** command was extended to support standard and extended communities. Use this command to configure propagation of standard and extended communities if your BGP design relies on the use of standard communities. An example of such use would be propagation of QoS information across the network.

https://t.me/learningnets

MP-BGP Configuration in Cisco IOS/XE/XR



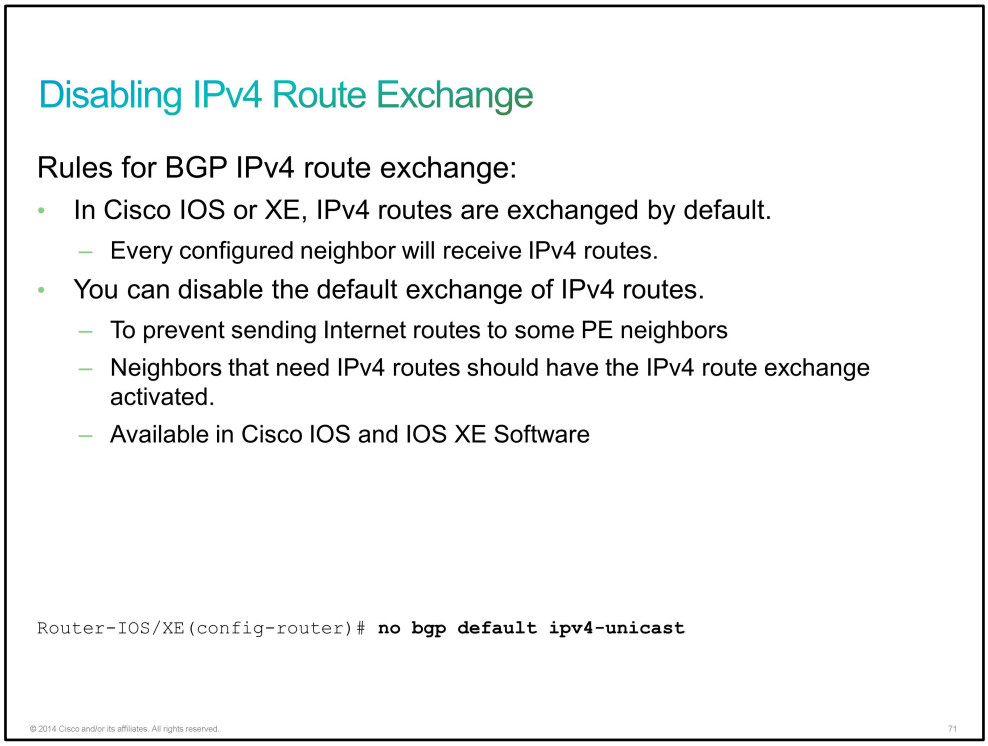
The right box in the figure shows Cisco IOS and IOS XE Software configuration. A neighbor must be configured in the BGP section and then activated in the address family block. The extended community command is added automatically.

The left box shows Cisco IOS XR Software configuration. In this case, the VPNv4 address family is configured in the BGP section and then applied in the neighbor configuration block.

https://t.me/learningnets

Disabling IPv4 Route Exchange

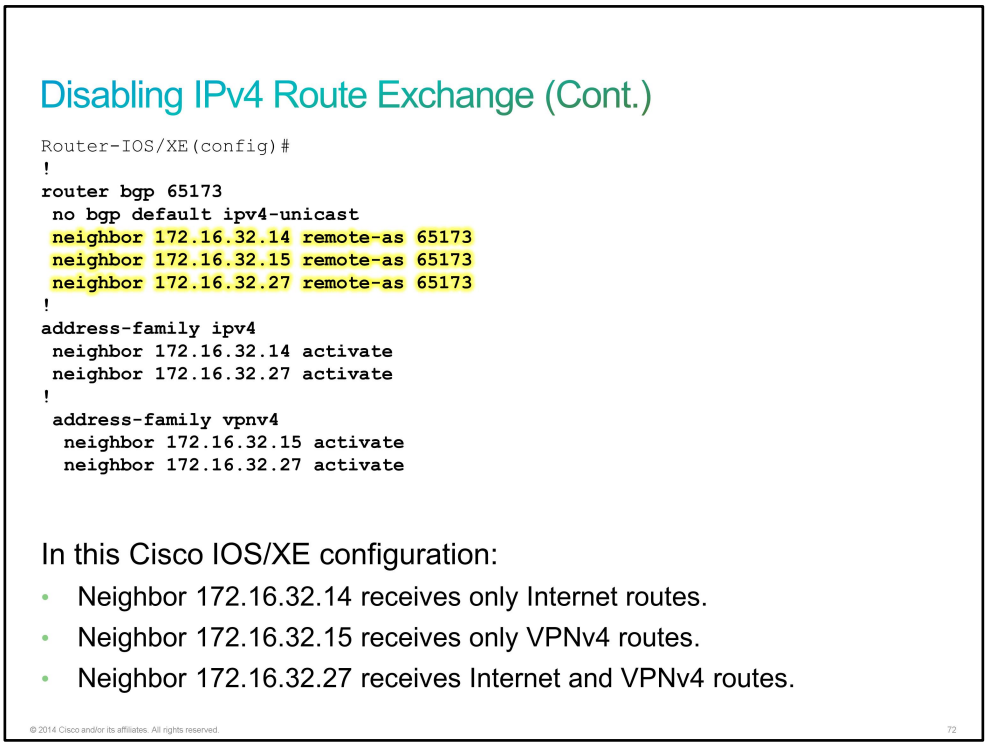
The BGP configuration that has been discussed so far is appropriate for situations where the PE routers provide Internet and VPN connectivity. If the PE routers provide only VPN connectivity, they do not need Internet routing, and the IPv4 route exchange should be disabled.



Here are the two ways of disabling IPv4 route exchange:

- To disable IPv4 route exchange for only a few neighbors, your best option is to disable the IPv4 route exchange on a neighbor-by-neighbor basis by using the **no neighbor activate** command.
- To disable IPv4 route exchange for most (or all) of the neighbors, you can use the **no bgp default ipv4-unicast** command. After you enter this command, you must manually activate IPv4 route exchange for each configured global BGP neighbor.

https://t.me/learningnets



In this example, only a subset of BGP neighbors needs to receive IPv4 routes.

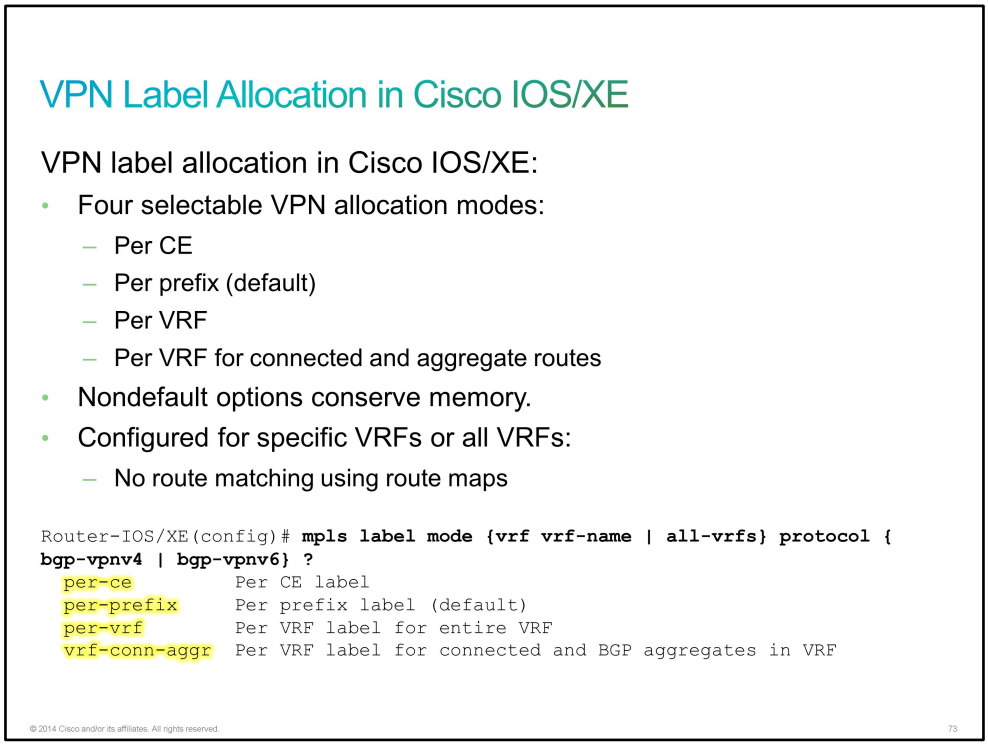
In the figure, the default propagation of IPv4 routes is therefore disabled. IPv4 route exchange—and VPNv4 route exchange—is manually activated on a neighbor-by-neighbor basis:

- Neighbor 172.16.32.14 receives only Internet routes that are based on the IPv4 activation.
- Neighbor 172.16.32.15 receives only VPNv4 routes that are based on the VPNv4 activation.
- Neighbor 172.16.32.27 receives Internet and VPNv4 routes that are based on both IPv4 and VPNv4 activations.

https://t.me/learningnets

VPN Label Allocation in Cisco IOS/XE

The PE devices store both local and remote routes and include a VPN label entry for each VPN route. The Cisco IOS or XE allows you to select one of the VPN allocation modes: per-ce, per-prefix, per-vrf, or per-vrf for connected and aggregate routes.



The default mode, per-prefix, consumes the most memory. When there are many VRF domains and routes, the amount of memory that the per-prefix labels consume can become an issue. The purpose of using per CE label allocation or the other modes is to conserve label space.

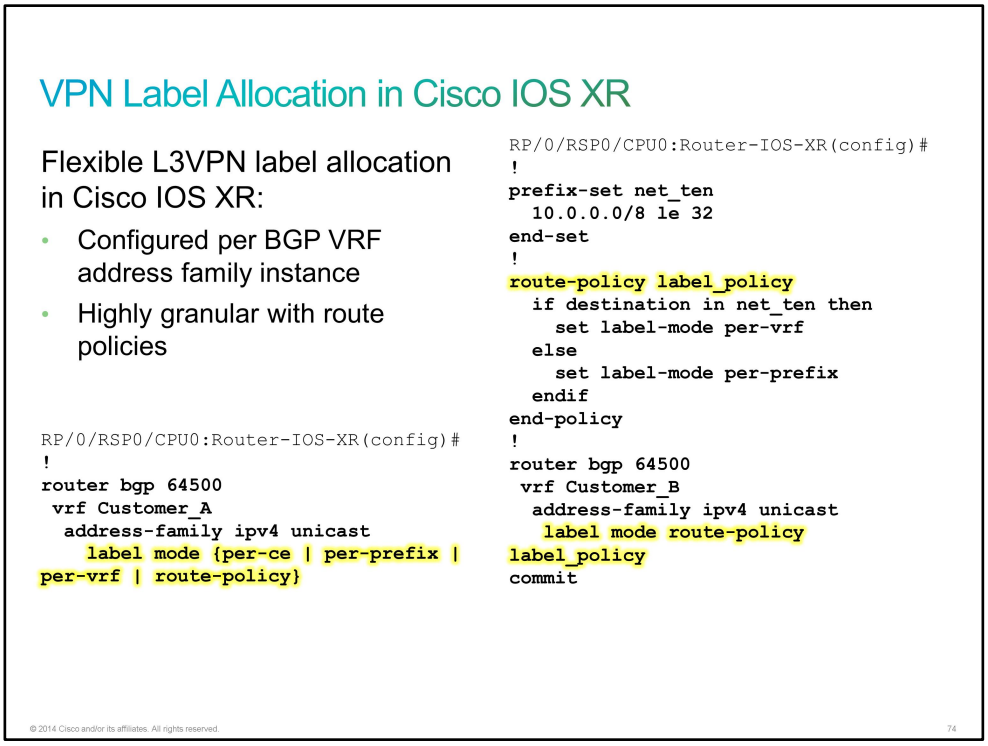
The MPLS VPN allocation mode feature allows the same label to be used for all the routes advertised from a unique CE peer device or VRF.

In the per-ce mode, the PE device allocates one label for every immediate next hop (in most cases, the next hop is a CE router). The label is directly mapped to the next hop so there is no VRF route lookup performed during data forwarding. However, the number of labels allocated is one for each CE rather than one for each prefix. As BGP is aware of all the next hops, it assigns a label for each next hop (not for each PE-CE interface).

https://t.me/learningnets

VPN Label Allocation in Cisco IOS XR

Cisco IOS XR provides the Flexible L3VPN Label Allocation Mode feature. It offers you the ability to set label allocation mode using a route policy, where different allocation modes can be used for different sets of prefixes. Thus, the label mode can be chosen based on arbitrary match criteria such as prefix value and community.



Use the **label mode** command to set the MPLS or VPN label mode based on prefix value. The label-mode attach point enables you to choose the label mode based on any arbitrary criteria.

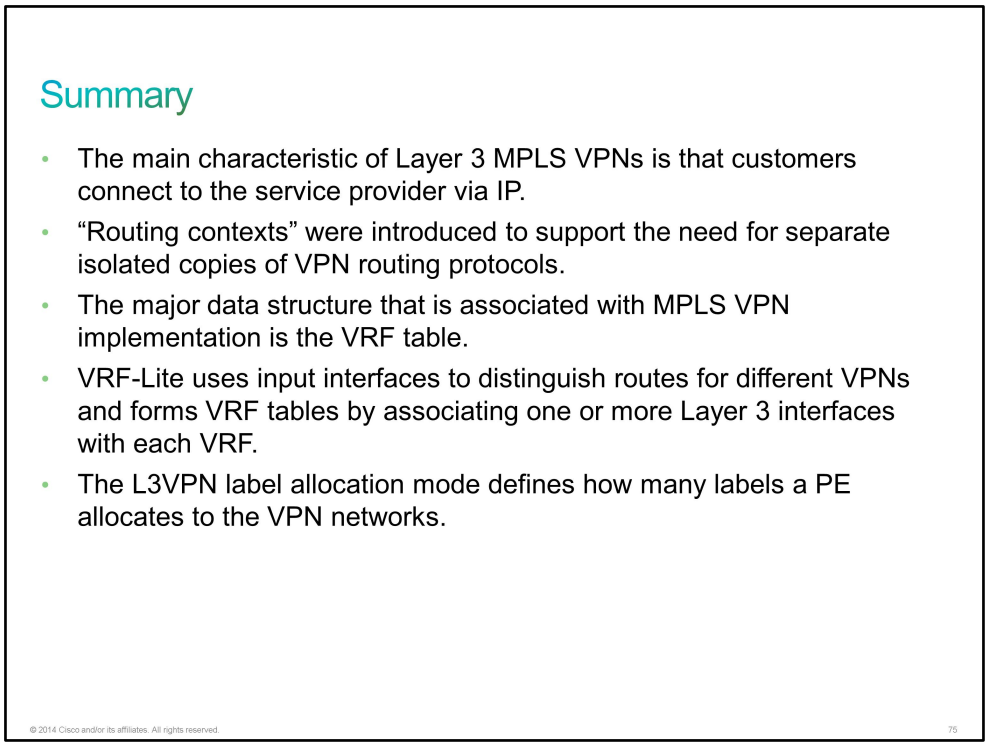
The Cisco IOS XR configuration example uses a prefix set (net_ten) to match any subnets of the 10.0.0.0/8 network. Cisco IOS XR prefix-sets are more flexible equivalents of Cisco IOS prefix-lists.

The route policy (label_policy) defines two L3VPN allocation modes: per-vrf for subnets of the 10.0.0.0/8 network, and per-prefix for all other prefixes. In this scenario, the same L3VPN label will be allocated for all subnets 10.0.0.0/8 network and the connected routes. Individual labels will be generated for all other nonconnected VPN networks.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Deploying Basic Routing in MPLS VPNs

Overview

This lesson explains PE-CE routing protocol configuration steps and the various routing protocols that you can run between PE and CE routers. These protocols include RIP, EIGRP, and static routes.

You must understand not only what you can configure between PE and CE routers when you are setting up MPLS VPNs, but also how to accomplish the configuration successfully. Determining the steps you should take when trying to solve a problem with your MPLS VPN network is also very important.

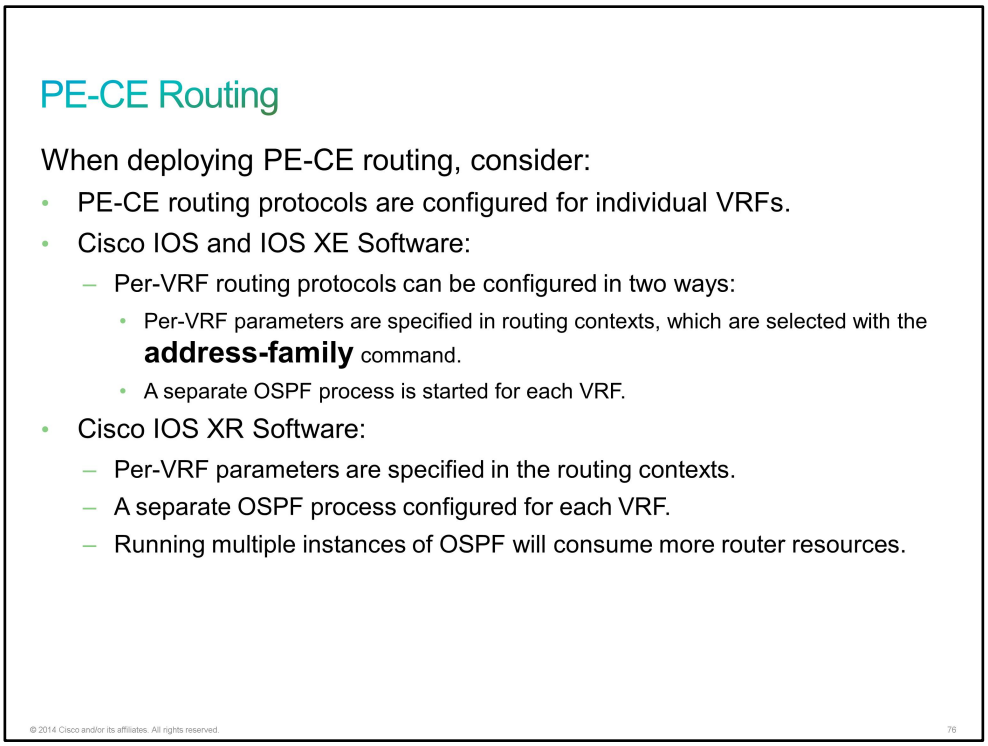
Upon completing this lesson, you will be able to:

- Identify the requirements for configuring PE-CE routing protocols
- Explain VRF routing context configuration within BGP
- Explain static PE-CE routing configuration
- Explain RIPv2 PE-CE routing configuration
- Explain EIGRP PE-CE routing configuration
- Describe reinjecting routes into backbone by multihomed sites
- Describe the implementation of EIGRP SOO for loop prevention
- Describe the inbound and outbound route filtering in VRFs
- Describe how to monitor Layer 3 MPLS operations

https://t.me/learningnets

PE-CE Routing

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



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

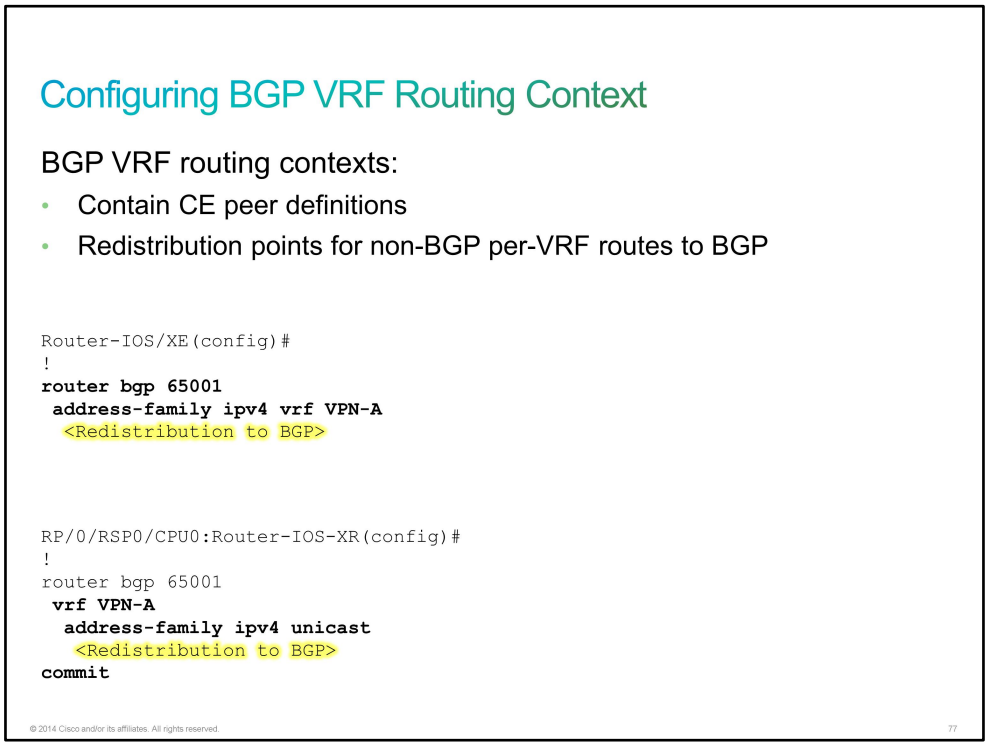
The per-VRF routing protocols can be configured in these ways:

- Cisco IOS and IOS XE Software:
  - Per-VRF routing protocols can be configured in two ways:
    - Per-VRF parameters are specified in routing contexts, which are selected with the **address-family** command.
    - A separate OSPF process has to be started for each VRF.
  - Before Cisco IOS Release 12.3(4)T, the overall number of routing processes per router was limited to 32, of which only 28 were available for VRF assignment.
- Cisco IOS XR Software:
  - Per-VRF parameters are specified in the routing contexts.
  - A separate OSPF process configured for each VRF.
  - Running multiple instances of OSPF will consume more router resources.

https://t.me/learningnets

Configuring BGP VRF Routing Context

The BGP VRF routing context contains the EBGP peering definition if BGP is used as the PE-to-CE routing protocol. It defines IGP-to-BGP route redistribution when an IGP is enabled for PE-to-CE routing.



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

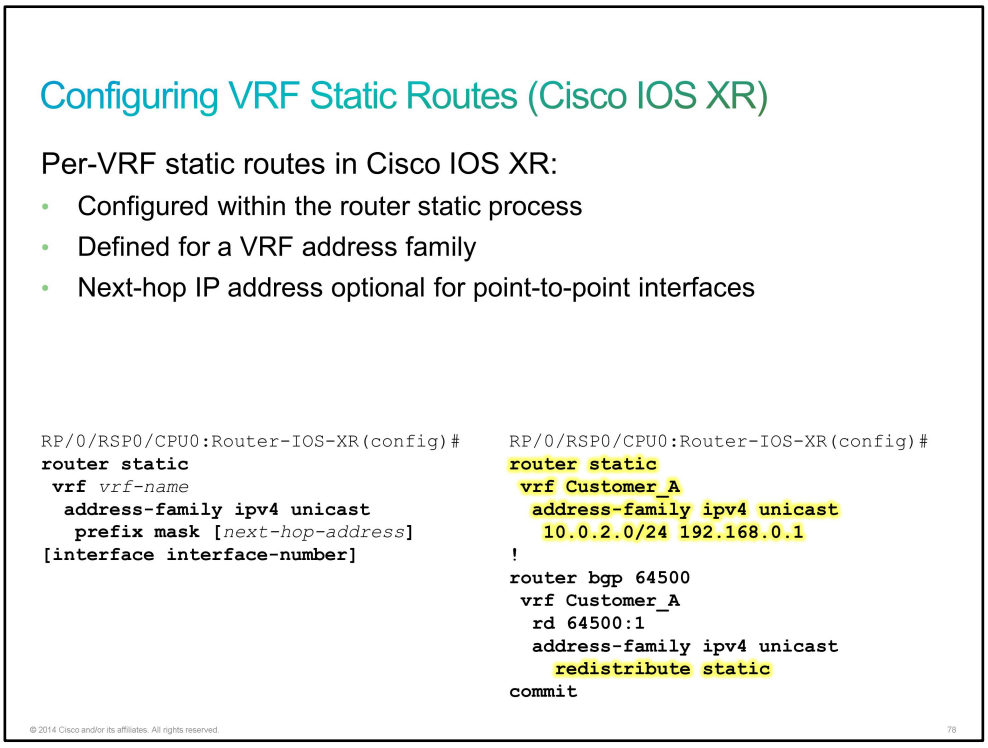
On Cisco IOS XR devices, first define a VRF with the **vrf** *vrf-name* command in the BGP routing processes. Then select the routing context with the **address-family ipv4 unicast** command. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

https://t.me/learningnets

Configuring Static PE-CE Routing

Configuring VRF Static Routes (Cisco IOS XR)

To establish static routes for a VPN VRF instance on Cisco IOS XR devices, you must first define the static router. To enter static router configuration mode, use the **router static** command in global configuration mode. Use the **vrf** command to configure a VRF instance. Enter address family configuration mode with the **address-family ipv4 unicast** command to configure a static route.



Configuring VRF Static Routes (Cisco IOS and IOS XE)

To establish static routes for a VPN VRF instance on Cisco IOS and IOS XE devices, use the **ip route vrf** command in global configuration mode. To disable static routes, use the **no** form of this command.

https://t.me/learningnets

Configuring VRF Static Routes (Cisco IOS/IOS XE)

Per-VRF static routing configuration in Cisco IOS/XE:

- Installs the route in the VRF table
- Next-hop IP address optional for point-to-point interfaces

```
Router-IOS/XE(config)# ip route vrf vrf-name prefix mask [next-hop-address]
[interface interface-number]
```

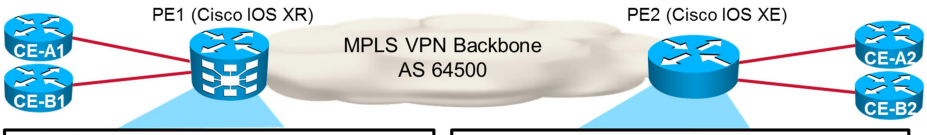
```
Router-IOS/XE(config)#
ip route vrf Customer_A 10.0.0.0 255.0.0.0 10.250.0.2
!
router bgp 65173
 address-family ipv4 vrf Customer_A
  redistribute static
```

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

Example: Configuring VRF Static Routes

The examples in the figure (Cisco IOS and IOS XE and IOS XR CLIs) show how to configure PE-CE routing using static routes.

Example: Configuring VRF Static Routes




```
RP/0/RSP0/CPU0:PE1(config)#
router static
 vrf Customer_A
  address-family ipv4 unicast
   10.0.1.0/24 192.168.1.2
!
router bgp 64500
 vrf Customer_A
  rd 64500:1
  address-family ipv4 unicast
   redistribute static
```

```
PE2(config)#
ip route vrf Customer_A 10.0.2.0
255.255.255.0 192.168.2.2
!
router bgp 64500
 address-family ipv4 vrf Customer_A
  redistribute static
  no auto-summary
```

```
CE-A1(config)#
ip route 0.0.0.0 0.0.0.0 192.168.1.1
```

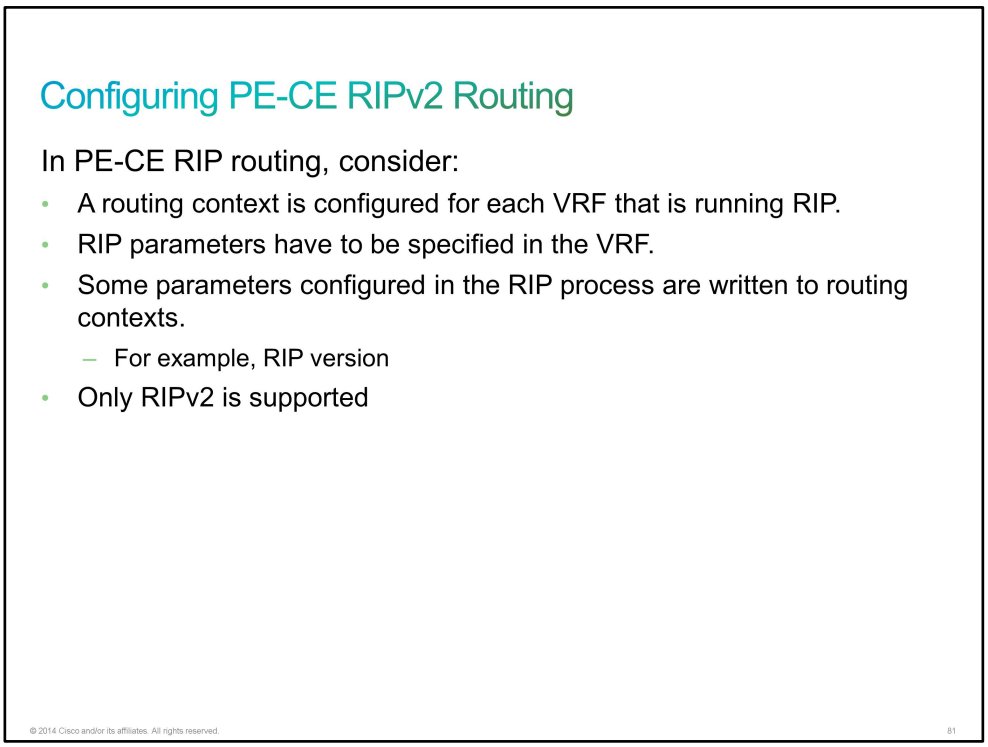
```
CE-A2(config)#
ip route 0.0.0.0 0.0.0.0 192.168.2.1
```

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

https://t.me/learningnets

Configuring PE-CE RIPv2 Routing

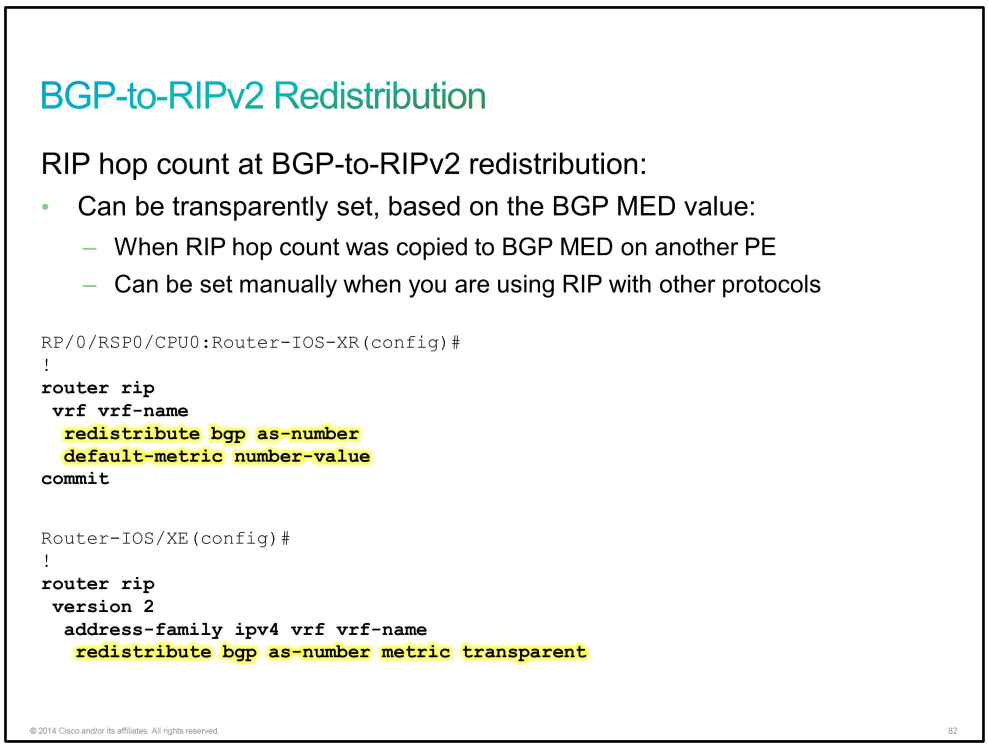
To configure RIP as the PE-CE routing protocol on Cisco IOS and IOS XE devices, start the configuration of the individual routing context with the **address-family ipv4 vrf** *vrf-name* command in router configuration mode. You can enter all standard RIP parameters in the per-VRF routing context. Global RIP parameters that are entered in the scope of RIP router configuration are inherited by each routing context and can be overwritten if needed in each routing context.



Note	Only RIPv2, and not RIPv1, is supported as the PE-CE routing protocol. Configuring the RIP version as a global RIP parameter using the `version 2` command in router configuration mode is a good practice.

https://t.me/learningnets

BGP-to-RIP Redistribution



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

On Cisco IOS and IOS XE devices, the MPLS VPN extension to the **redistribute** command (the **metric transparent** option) allows the MED attribute to be inserted as the IGP metric of a route redistributed from BGP back into RIP. This extension gives transparent end-to-end (from the customer perspective) RIP routing:

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

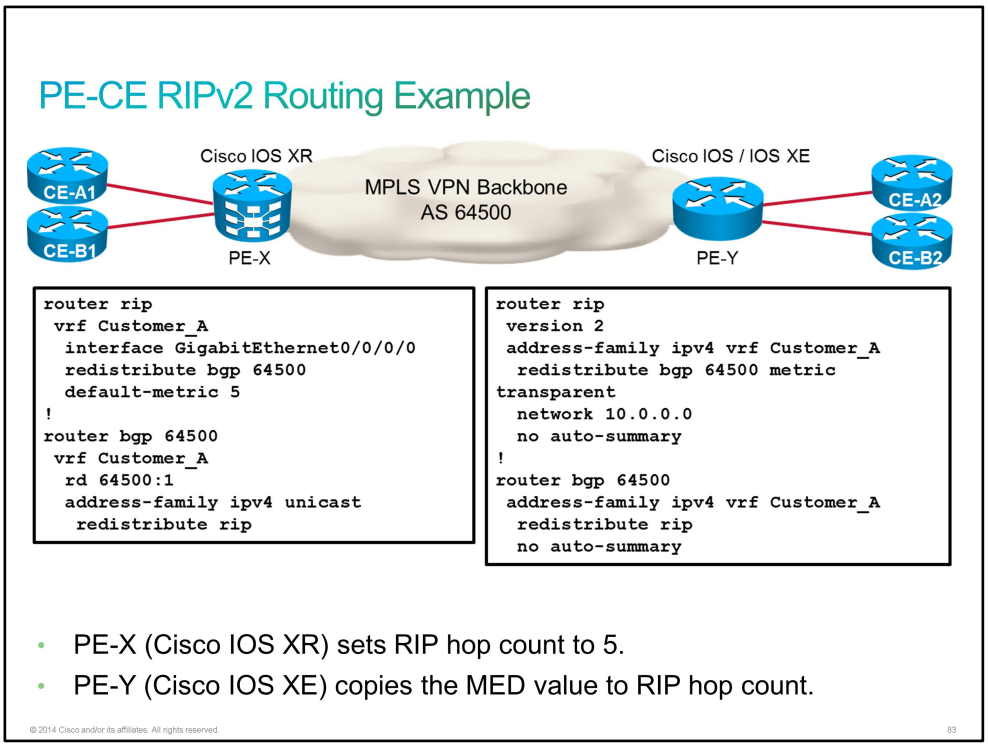
Note	You should not change the MED value within BGP if you use the **redistribute metric transparent** command.

On Cisco IOS XR devices, use the **default-metric** command in VRF configuration mode for RIP routing. The default metric value *(number-value)* is a number from 1 to 15.

https://t.me/learningnets

PE-CE RIPv2 Routing Example

The examples in the figure (Cisco IOS, IOS XE, and IOS XR CLIs) show how to configure PE-CE routing using RIP as a routing protocol.



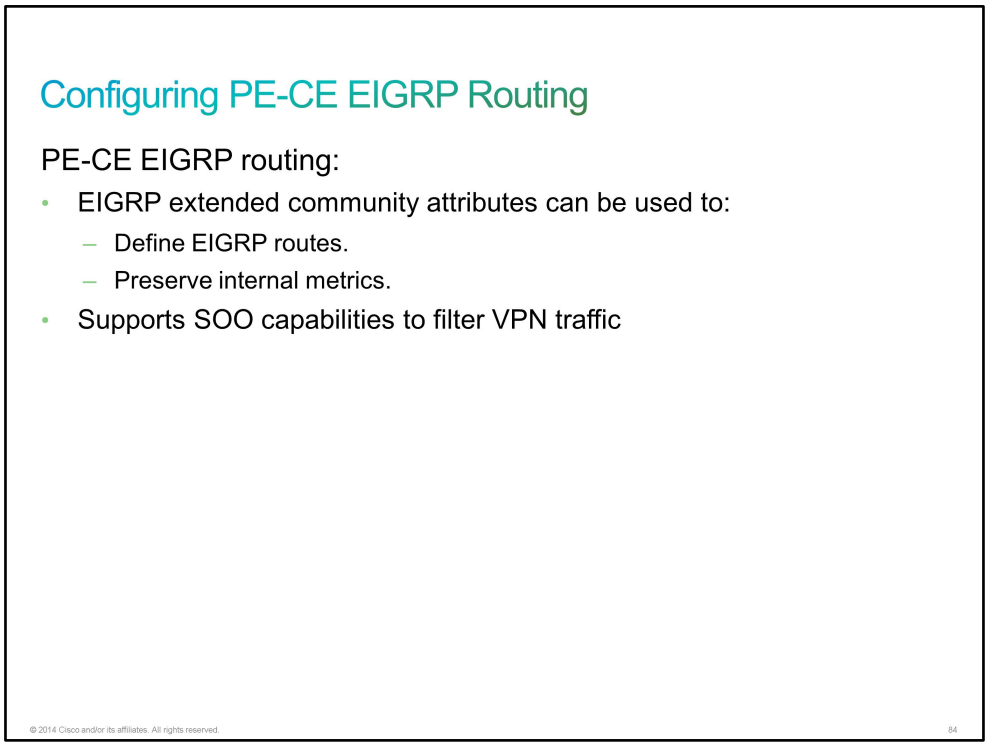
RIPv2 is enabled on both sides as the PE-to-CE routing protocol. Both PE routers redistribute RIPv2 routes into BGP and BGP routes (with appropriate route targets) into RIPv2.

PE-X (Cisco IOS XR) advertises the VPN routes into the RIP domain with the hop metric of 5. PE-Y (Cisco IOS XE) forwards the VPN routes into the RIP domain with the MED copied to the hop count. This reflects the hop-based distance copied into the MED by the originating PE-X. The PE routers copythe RIP hop count into BGP MED when redistributing RIP routes into BGP.

https://t.me/learningnets

Configuring PE-CE EIGRP Routing

The MPLS VPN Support for EIGRP between PE and CE feature provides the capability to transparently connect EIGRP customer networks through an MPLS-enabled BGP core network. EIGRP routes can then be redistributed through the VPN across the BGP network as IBGP routes. The configuration of this feature does not require any customer equipment upgrades or configuration changes. This feature is configured only on PE routers within the service provider network.



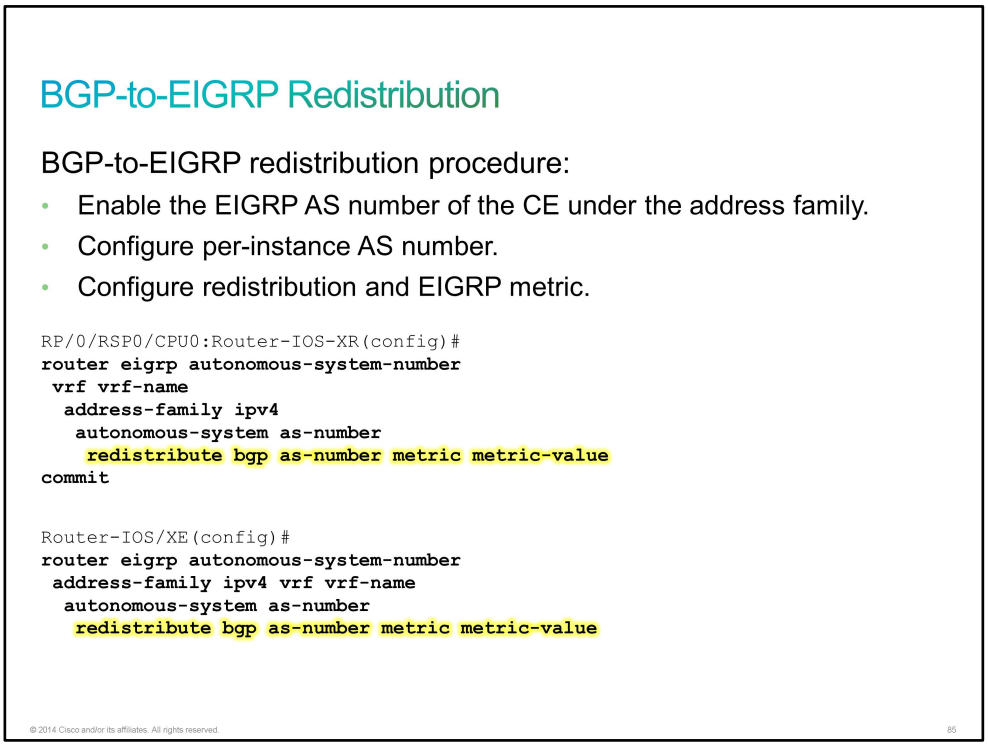
Customer networks and remote sites are connected to each other through the MPLS VPN. The configuration of this feature allows several EIGRP sites to connect seamlessly and appear as a single network. This integration is transparent to the customer sites. When this feature is enabled, EIGRP routes are converted to IBGP routes and transported through the BGP core network. EIGRP extended community attributes are used to define EIGRP routes and preserve internal metrics. These attributes are carried across the core network by MP-BGP.

This feature also introduces EIGRP support for MPLS and BGP extended community attributes such as SOO.

BGP-to-EIGRP Redistribution

The IGP metric is always copied into the MED attribute of the BGP route when an IGP route is redistributed into BGP. Within a standard BGP implementation, the MED attribute is used only as a route-selection criterion. The MED attribute is not copied back into the IGP metric. The metric must be configured for routes from external EIGRP autonomous systems and non-EIGRP networks before these routes can be redistributed into an EIGRP CE router.

https://t.me/learningnets



The EIGRP metric can be configured in the redistribute statement using the **redistribute** (IP) command or configured with the **default-metric** (EIGRP) command.

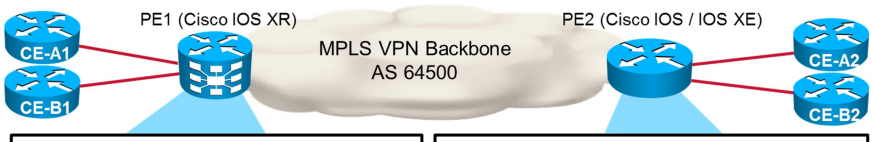
Note	In an MPLS VPN environment, the original EIGRP metrics must be carried inside MP-BGP updates. This configuration is achieved by using BGP extended community attributes to carry and preserve EIGRP metrics when crossing the MP-IBGP domain. These communities define the intrinsic characteristics that are associated with EIGRP, such as the AS number or EIGRP cost metric (bandwidth, delay, load, reliability, and MTU, for example).

PE-CE EIGRP Routing Example

The examples in the figure (Cisco IOS, IOS XE, and IOS XR CLIs) show how to configure PE-CE routing using EIGRP as the routing protocol.

https://t.me/learningnets

PE-CE EIGRP Routing Example




```
RP/0/RSP0/CPU0:PE1(config)#
!
router eigrp 1
 vrf Customer_A
  address-family ipv4
   default-metric 10000 100 255 1 1500
   autonomous-system 1
   redistribute bgp 64500
   interface GigabitEthernet0/0/0/0
!
router bgp 64500
 vrf Customer_A
  rd 64500:1
  address-family ipv4 unicast
   redistribute eigrp 1
```

```
PE2(config)#
!
router eigrp 1
 address-family ipv4 vrf Customer_A
autonomous-system 1
  network 10.0.0.0 255.255.255.0
  redistribute bgp 64500 metric
10000 100 255 1 1500
  no auto-summary
!
router bgp 64500
 address-family ipv4 vrf Customer_A
  redistribute eigrp 1
  no auto-summary
```

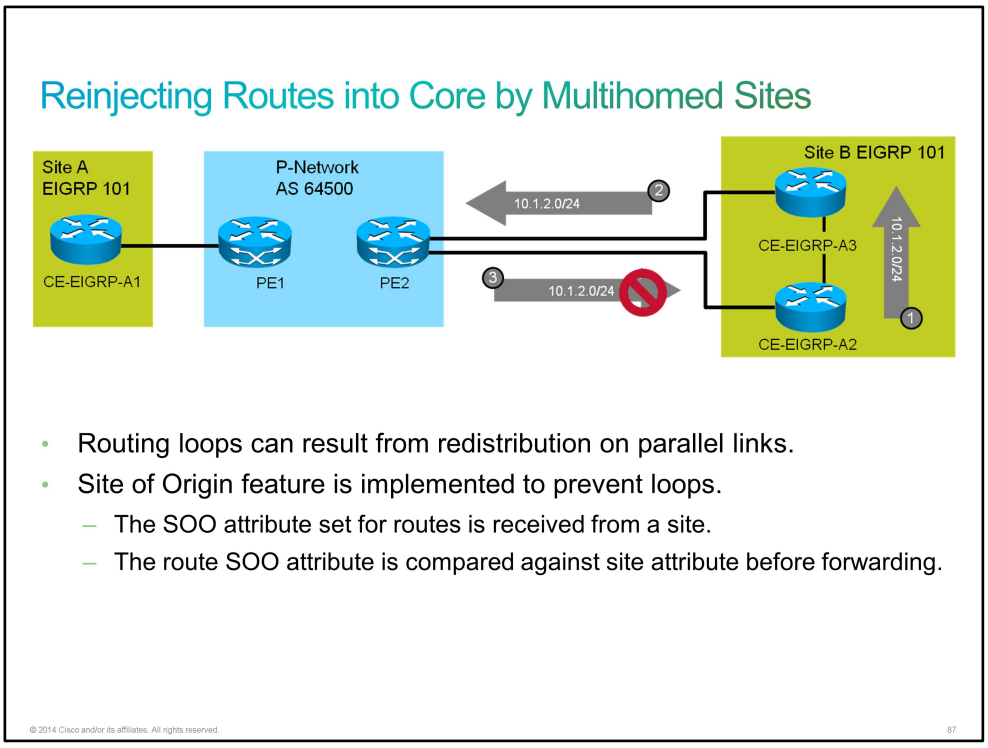
- EIGRP metric must defined for redistribution into EIGRP.

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

https://t.me/learningnets

Reinjecting Routes into Core by Multihomed Sites

The figure shows an MPLS VPN network for a customer that has two sites, Site A and Site B. Site B is multihomed. The figure shows that the EIGRP route 10.1.2.0/24 received by the multihomed site (Site B) is redistributed into the backbone at PE-Site-Y.



Routing loops and suboptimal routing generally occur because of mutual redistribution taking place between EIGRP PE-CE and MP-BGP in an MPLS VPN environment. Routing loops can occur in the following scenarios:

- A route that is received by a multihomed site from the backbone through one link can be forwarded back to the backbone through the other link.
- A route that originated in a multihomed site and that was sent to the backbone through one link can come back through the other link.

Routing loops and suboptimal routing can be avoided by using the following:

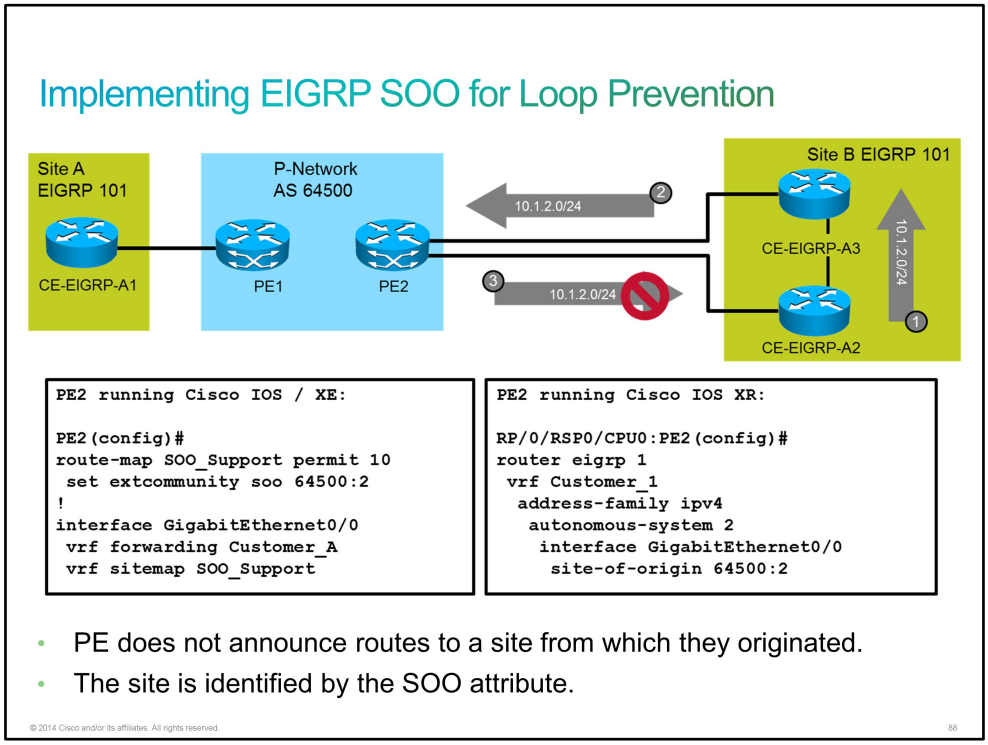
- The BGP Cost Community feature that can be used to force BGP to compare locally originated EIGRP routes and MP-IBGP routes, based on the EIGRP metric
- The EIGRP SOO feature on PE and CE routers that can be used to prevent routing loops

Note	The SOO attribute is needed only for customer networks with multihomed sites. Loops can never occur in customer networks that have only stub sites.

https://t.me/learningnets

Implementing EIGRP SOO for Loop Prevention

The configuration of the SOO extended community allows routers that support this feature to identify the site from which each route originated.



When this feature is enabled, the EIGRP routing process on the PE or CE router checks each received route for the SOO extended community and filters based on the following conditions:

- A received route from BGP or a CE router contains an SOO value that matches the SOO value on the receiving interface:
  - If a route is received with an associated SOO value that matches the SOO value that is configured on the receiving interface, the route is filtered out because it was learned from another PE router or from a backdoor link. This behavior is designed to prevent routing loops.
- A received route from a CE router is configured with an SOO value that does not match:
  - If a route is received with an associated SOO value that does not match the SOO value that is configured on the receiving interface, the route is accepted into the EIGRP topology table so that it can be redistributed into BGP.
  - If the route is already installed in the EIGRP topology table but is associated with a different SOO value, the SOO value from the topology table is used when the route is redistributed into BGP.
- A received route from a CE router does not contain an SOO value:
  - If a route is received without an SOO value, the route is accepted into the EIGRP topology table. The SOO value from the interface that is used to reach the next-hop CE router is appended to the route before it is redistributed into BGP.

https://t.me/learningnets

When BGP and EIGRP peers that support the SOO extended community receive these routes, they also receive the associated SOO values and pass them to other BGP and EIGRP peers that support the SOO extended community. This filtering is designed to prevent transient routes from being relearned from the originating site, which prevents transient routing loops from occurring.

In conjunction with BGP Cost Community, EIGRP, BGP, and the RIB ensure that paths over the MPLS VPN core are preferred over backdoor links.

In the example in the figure, a route map, SOO_Support, with a specific SOO value of 64500:2 is defined. The newly defined route map is then applied to VRF Customer-EIGRP-A1, which connects the EIGRP neighbor (the CE-EIGRP-A2 router) to the PE-Site-Y router.

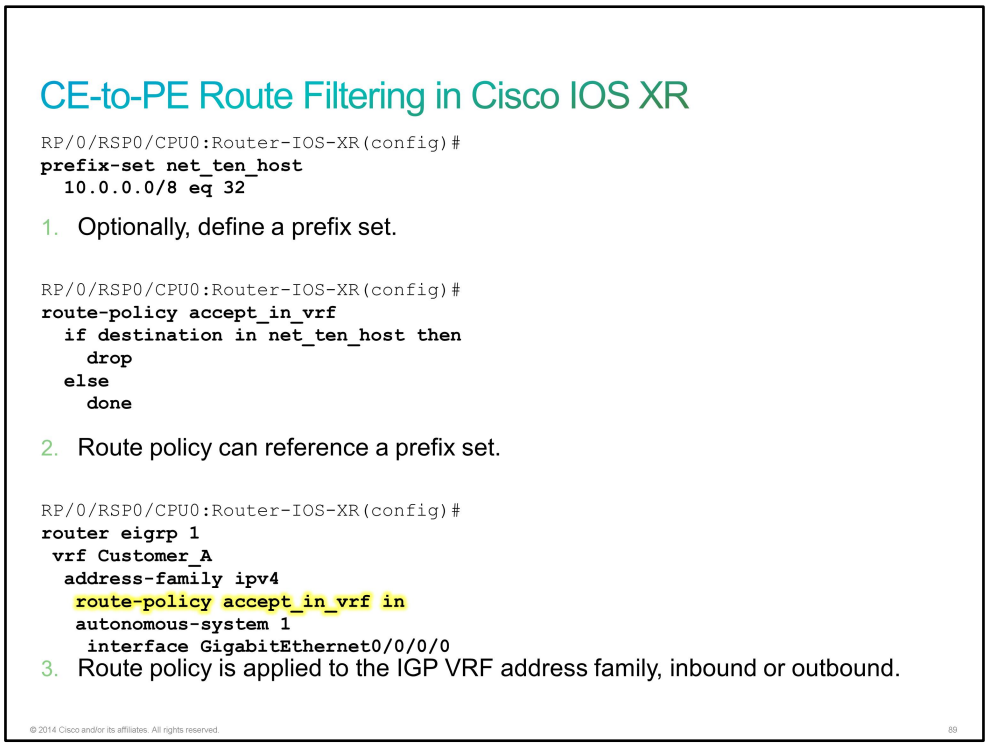
 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Route Filtering

CE-to-PE Route Filtering in Cisco IOS XR

You can filter VPN routes at ingress from the CE to the VRF table or at egress from the VRF to the CE. In Cisco IOS XR, the VRF instances of the IGP processes (except OSPF) allow you to configure inbound or outbound route filtering using a route policy. This example illustrates how to drop /32 subnets of network 10.0.0.0/8 and permit all other EIGRP updates into the VRF Customer_A.



CE-to-PE Route Filtering in Cisco IOS/XE

Cisco IOS and IOS XE offer similar filtering capabilities. You can filter at ingress from the CE to the VRF table or at egress from the VRF to the CE. You can filter using either a route map or a prefix list applied in the inbound or outbound direction. The VRF instances of the IGP processes as well as as the VRF-specific OSPF process allow such filtering configuration. This example illustrates how to drop any subnets of network 10.0.0.0/8 and permit all other EIGRP updates into the VRF Customer_A.

https://t.me/learningnets

CE-to-PE Route Filtering in Cisco IOS/XE

```
Router-IOS/XE(config)#
ip prefix-list net_ten_host_prefix_list seq 5 permit 10.0.0.0/8 ge 32
```

1. Define a prefix list.

```
Router-IOS/XE(config)#
route-map accept_in_vrf deny 10
 match ip address prefix-list net_ten_host_prefix_list
!
route-map accept_in_vrf permit 20
```

2. Optionally, configure a route map.

```
Router-IOS/XE(config)#
router eigrp 1
 address-family ipv4 vrf Customer_A autonomous-system 1
  redistribute bgp 64500 metric 10000 100 255 1 1500
  network 192.168.0.0 0.0.255.255
  distribute-list route-map accept_in_vrf in
```

3. Apply the route map or prefix list, inbound or outbound.

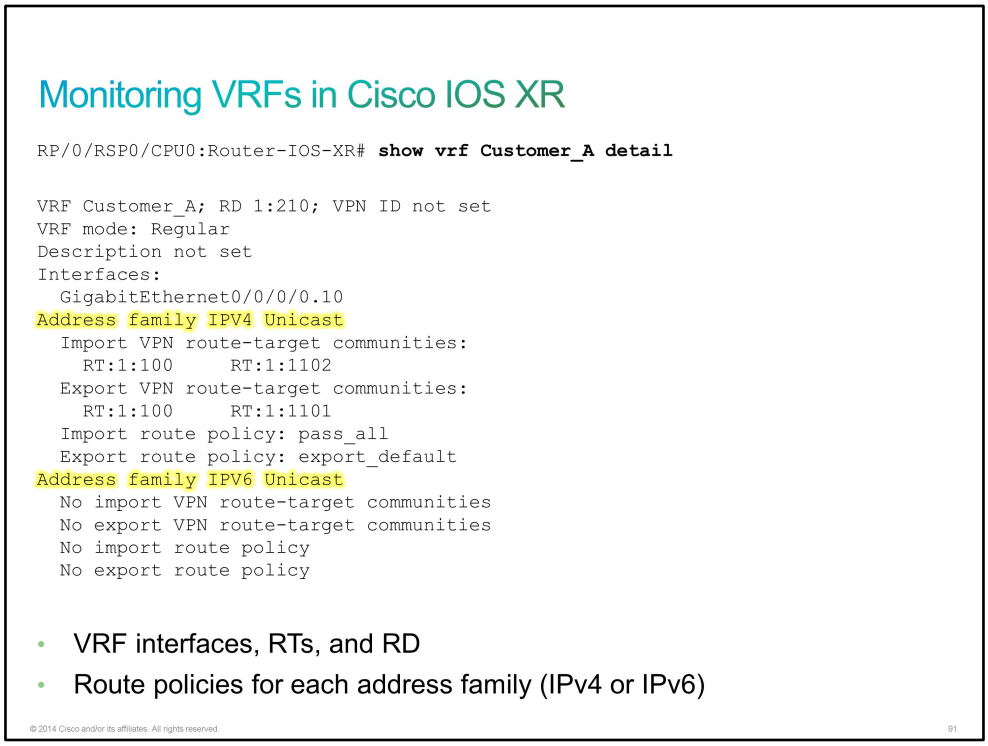
© 2014 Cisco and/or its affiliates. All rights reserved.

https://t.me/learningnets

Monitoring Layer 3 MPLS VPNs

Monitoring VRFs in Cisco IOS XR

To display the set of defined VRFs, use the **show vrf all** (IOS XR) command in EXEC mode. You can view a specific VRF and examine its details, as shown above. The output provides details on the interfaces, route distinguisher, route targets, and import and export route policies.



You can also display the interfaces that are associated with a specific VRF using the **show ipv4 vrf all interface brief** command.

Monitoring VRFs in Cisco IOS/XE

In Cisco IOS/XE you can view the details of a VRF using the **show vrf detail vrf-name** command. The output includes, in addition to the VRF interfaces, the route distinguisher, route targets, import and export route policies, and information about the CLI format (new or old).

https://t.me/learningnets

Monitoring VRFs in Cisco IOS/XE

```
Router-IOS/XE# show vrf detail Customer_A
VRF Customer_A (VRF Id = 2); default RD 1:210; default VPNID <not set>
  New CLI format, supports multiple address-families
  Flags: 0x180C
  Interfaces:
    Gi0/0/0.20
Address family ipv4 (Table ID = 2 (0x2)):
  Flags: 0x0
  Export VPN route-target communities
    RT:1:210                RT:1:1101
  Import VPN route-target communities
    RT:1:210                RT:1:1102               RT:1:310
  No import route-map
  No global export route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-vrf (Label 33)
Address family ipv6 not active
```

- Information about the old or new VRF CLI format
- VRF interfaces, RTs, RD, and policies

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

Monitoring VRF Routing

You can examine the VRF routing table by using the **show route vrf** (IOS XR) or **show ip route vrf** (IOS and IOS XE) commands, respectively.

Monitoring VRF Routing

```
RP/0/RSP0/CPU0:Router-IOS-XR# show route vrf Customer_A
<... output omitted ...>

B    10.2.10.1/32 [200/2] via 10.2.1.1 (nexthop in vrf default), 2d07h
B    10.99.2.1/32 [200/2] via 10.2.1.1 (nexthop in vrf default), 2d07h
L    192.168.101.10/32 is directly connected, 05:43:13,
GigabitEthernet0/0/0/0.10
<... output omitted ...>



Router-IOS/XE# show ip route vrf Customer_A
<... output omitted ...>
Gateway of last resort is 10.1.1.1 to network 0.0.0.0

B*    0.0.0.0/0 [200/1] via 10.1.1.1, 05:41:42
      10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
B        10.1.10.1/32 [200/2] via 10.1.1.1, 05:41:58
<... output omitted ...>
```

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

https://t.me/learningnets

Monitoring MP-BGP Sessions

The **show bgp neighbors** (Cisco IOS XR) and the **show bgp {all | ipv4 | ipv6 | l2vpn | nsap | rtfilter | vpnv4 | vpnv6 | vrf} neighbors** (Cisco IOS and IOS XE) commands are used to monitor BGP sessions with other PE routers and the address families that are negotiated with these neighbors.

Monitoring MP-BGP Sessions

```
RP/0/RSP0/CPU0:Router-IOS-XR# show bgp neighbors

BGP neighbor is 10.0.2.1
 Remote AS 64500, local AS 64500, internal link
 Remote router ID 10.0.2.1
  BGP state = Established, up for 2d07h
<... output omitted ...>

Router-IOS/XE# show bgp all neighbors
For address family: IPv4 Unicast

For address family: VPNv4 Unicast
BGP neighbor is 10.0.2.1,  remote AS 64500, internal link
  BGP version 4, remote router ID 10.0.2.1
  BGP state = Established, up for 2d08h
  Last read 00:00:34, last write 00:00:09, hold time is 180, keepalive interval
is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
<... output omitted ...>
```

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

Monitoring MP-BGP VPNv4 Table

Monitoring MP-BGP VPNv4 Table

- Detailed VPNv4 VRF monitoring shows VPN labels

```
RP/0/RSP0/CPU0:Router-IOS-XR# show bgp vpnv4 unicast vrf A 10.2.1.1/32
BGP routing table entry for 10.2.1.1/32, Route Distinguisher: 1:210
<... output omitted ...>
      Received Label 33
      Origin incomplete, metric 2, localpref 100, valid, internal, best, group-
best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 335
      Extended community: OSPF domain-id:0x5:0x0000000a0200 OSPF route-
type:0:2:0x0 OSPF router-id:192.168.102.20 RT:1:210 RT:1:1101
      Originator: 10.2.1.1, Cluster list: 10.0.2.1
      Source VRF: Customer_1, Source Route Distinguisher: 1:210


Router-IOS/XE# show bgp vpnv4 unicast vrf Customer_1 10.1.10.1/32
BGP routing table entry for 1:210:10.1.10.1/32, version 293
<... output omitted ...>
    10.1.1.1 (metric 4) from 10.0.2.1 (10.0.2.1)
      Origin incomplete, metric 2, localpref 100, valid, internal, best
      Extended Community: RT:1:100 RT:1:210 RT:1:1101 OSPF RT:0.0.0.0:1:0
        OSPF ROUTER ID:10.1.1.1:0
      Originator: 10.1.1.1, Cluster list: 10.0.2.1
      mpls labels in/out nolabel/16014
```

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

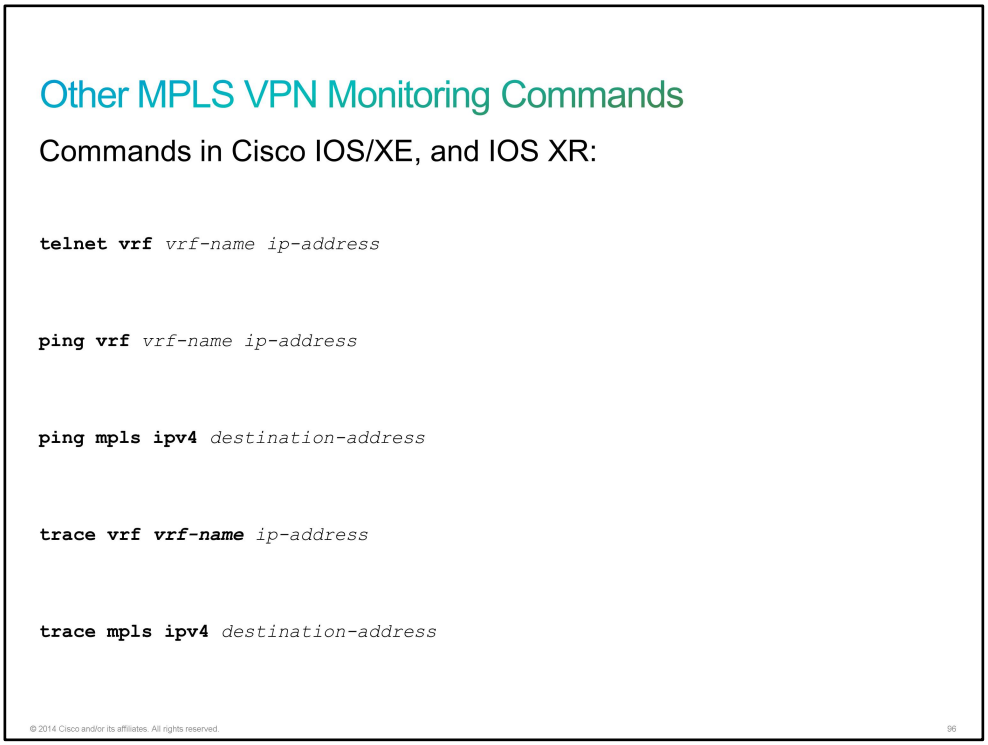
https://t.me/learningnets

The **show ip bgp vpnv4** (Cisco IOS and IOS XE) or **show bgp vpnv4 unicast** (Cisco IOS XR) command displays IPv4 BGP information and VPNv4 BGP information. To display VPNv4 BGP information on devices that are running Cisco IOS or IOS XE, use one of these keywords:

- **all** to display the whole contents of the VPNv4 BGP table
- **vrf** *vrf-name* to display VPNv4 information that is associated with the specified VRF
- **rd** *route-distinguisher* to display VPNv4 information that is associated with the specified RD

If you view the details of BGP VPNv4 routes associated with a VRF, you can examine the VPN labels, as shown here.

Other MPLS VPN Monitoring Commands



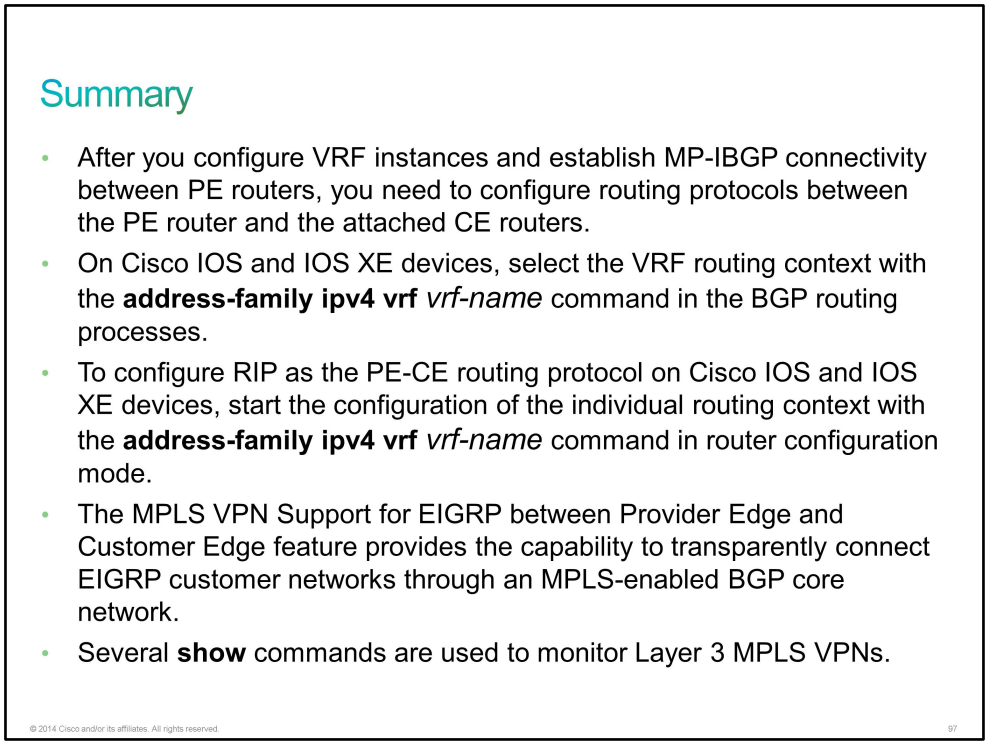
Additional Cisco IOS Software monitoring commands are as follows:

- The **telnet** command can be used to connect to a CE router from a PE router using the **/vrf** option.
- The **ping vrf** command can be used to ping a destination host reachable through a VRF.
- The **ping mpls** command can be used to check MPLS LSP connectivity.
- The **trace vrf** command can be used to trace a path toward a destination reachable through a VRF.
- The **trace mpls** command can be used to discover the MPLS LSP routes that packets actually take when traveling to their destinations.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Deploying OSPF and BGP in MPLS VPNs

Overview

This lesson explains the PE-CE advanced routing protocol configuration steps and the various routing protocols that you can run between PE and CE routers.

The OSPF and BGP routing protocols are discussed. This lesson describes the steps that are needed for MPLS VPN troubleshooting.

Upon completing this lesson, you will be able to:

- Describe OSPF as the routing protocol between PE and CE routers
- Describe the OSPF-BGP redistribution issue in MPLS VPN routing
- Describe the OSPF caveats
- Describe the OSPF superbackbone
- Describe the OSPF superbackbone route propagation
- Describe the OSPF superbackbone implementation
- Describe OSPFv3 as a CE-PE routing protocol
- Describe OSPF as a PE-CE routing configuration procedure
- Describe OSPFv2 as a PE-CE routing configuration in Cisco IOS/XE
- Describe OSPFv2 as PE-CE routing configuration in Cisco IOS
- Describe an OSPFv2 PE-CE configuration example
- Describe an OSPFv3 PE-CE configuration example
- Describe the risk of routing loops between MP-BGP and OSPF
- Describe the OSPF down bit

https://t.me/learningnets

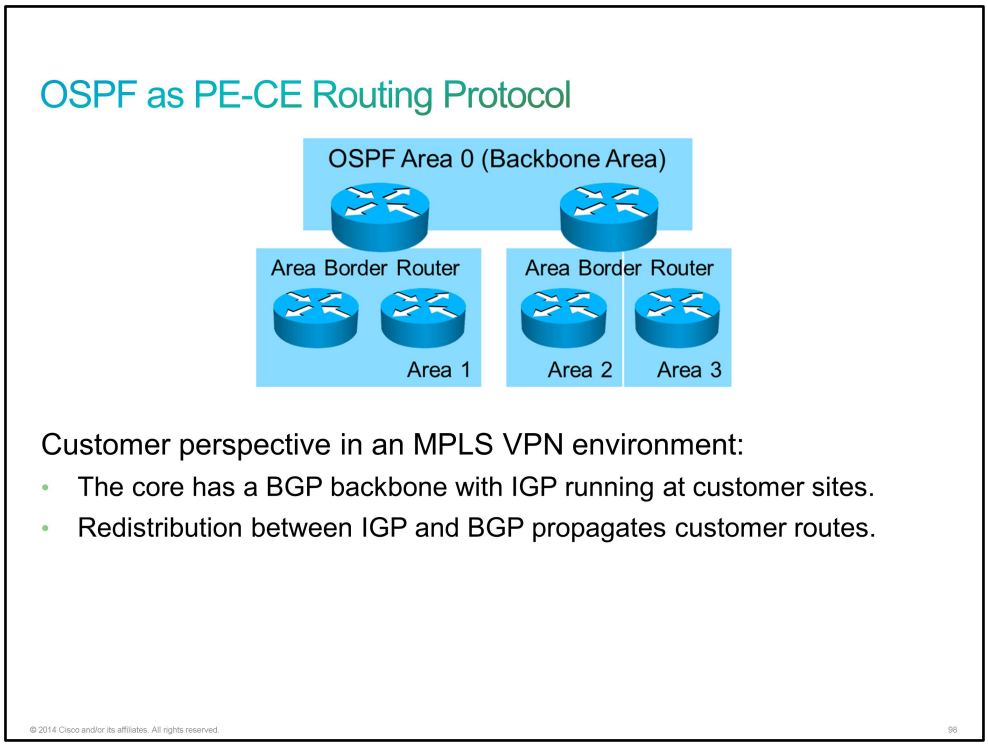
- Describe packet forwarding without the OSPF down bit
- Describe packet forwarding with the OSPF down bit
- Describe the OSPF sham link
- Explain the OSPF sham-link configuration
- Explain BGP PE-CE routing configuration
- Explain how to limit the number of routes in a VRF
- Describe the problem with the reusing same AS number in VPN sites
- Describe the AS-override feature
- Describe an example with the AS-override feature
- Describe the BGP exchange problem with multihomed sites
- Explain the AS-override and allowAS-In features
- Describe BGP filtering in Cisco IOS XR
- Describe BGP filtering in Cisco IOS/XE
- Describe the Layer 3 MPLS VPN troubleshooting steps

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF as PE-CE Routing Protocol

The OSPF routing protocol was designed to support hierarchical networks with a central backbone.



A network that is running OSPF is divided into areas. All areas must be directly connected to the backbone area (Area 0). The whole OSPF network (backbone area and other connected areas) is called the OSPF domain.

The OSPF areas in the customer network can correspond to individual sites, but these other options are often encountered:

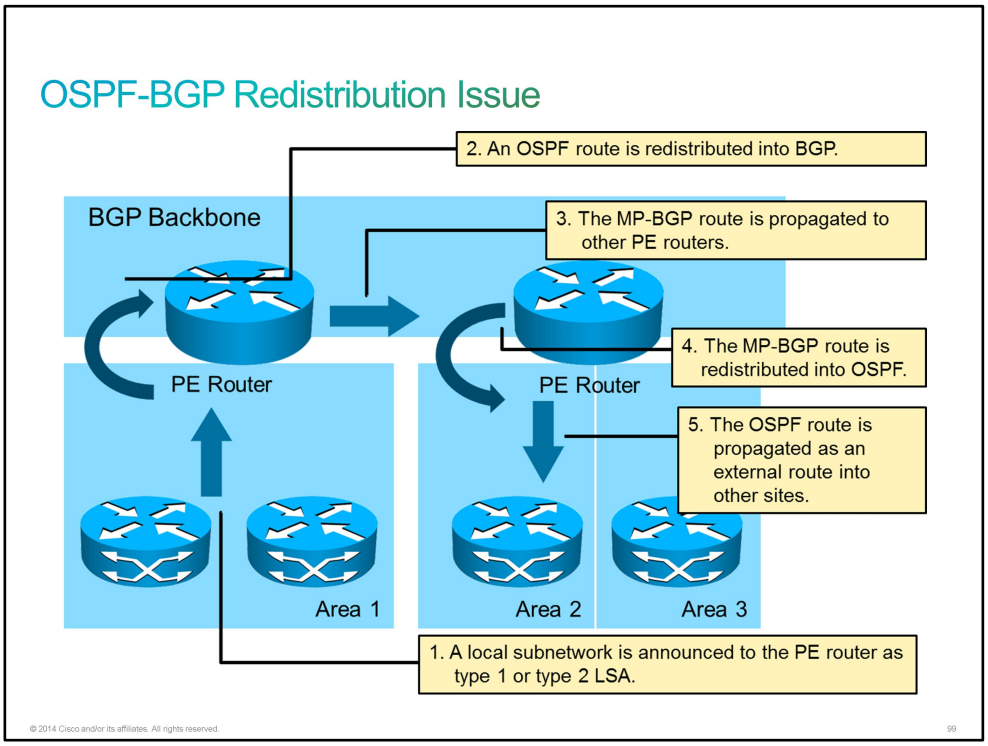
- A single area could span multiple sites (for example, the customer decides to use one area per region, but the region contains multiple sites).
- The backbone area could be extended into individual sites.

The MPLS VPN routing model introduces a BGP backbone into the customer network. Isolated copies of the IGP run at every site, and MP-BGP is used to propagate routes between sites. Redistribution between the customer IGP (running between PE routers and CE routers) and the backbone MP-BGP is performed at every PE router.

https://t.me/learningnets

OSPF-BGP Redistribution Issue

The IGP-BGP redistribution that is introduced by the MPLS VPN routing model does not fit well into customer networks running OSPF.



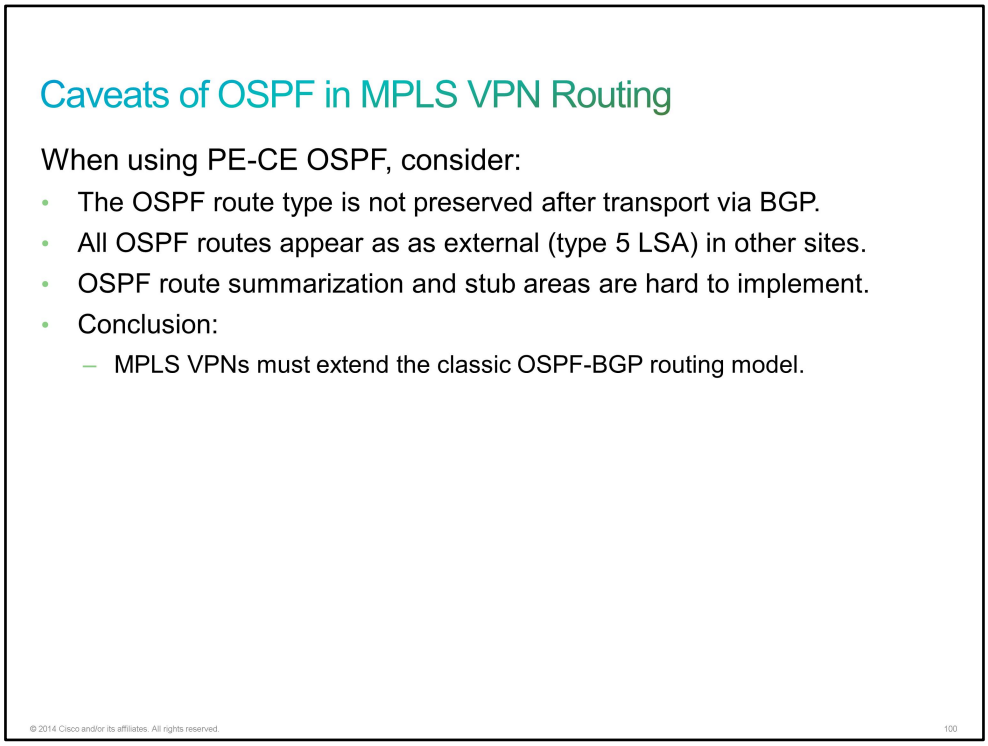
When an OSPF customer is migrated to an MPLS VPN implementation, any route that is redistributed into OSPF from another routing protocol will now be redistributed as an external OSPF route. The OSPF routes received by one PE router are propagated across the MPLS backbone and redistributed back into OSPF at another site as external OSPF routes.

The MPLS VPN routing model introduces a BGP backbone into the customer network. Isolated copies of the IGP run at every site, and MP-BGP is used to propagate routes between sites. Redistribution between the customer IGP (running between PE routers and CE routers) and the backbone MP-BGP is performed at every PE router.

https://t.me/learningnets

Caveats of OSPF in MPLS VPN Routing

With the traditional OSPF-BGP redistribution, the OSPF route type (internal or external route) is not preserved when the OSPF route is redistributed into BGP. When that route is redistributed back into OSPF, it is always redistributed as an external OSPF route.



This list identifies some of the caveats that are associated with external OSPF routes:

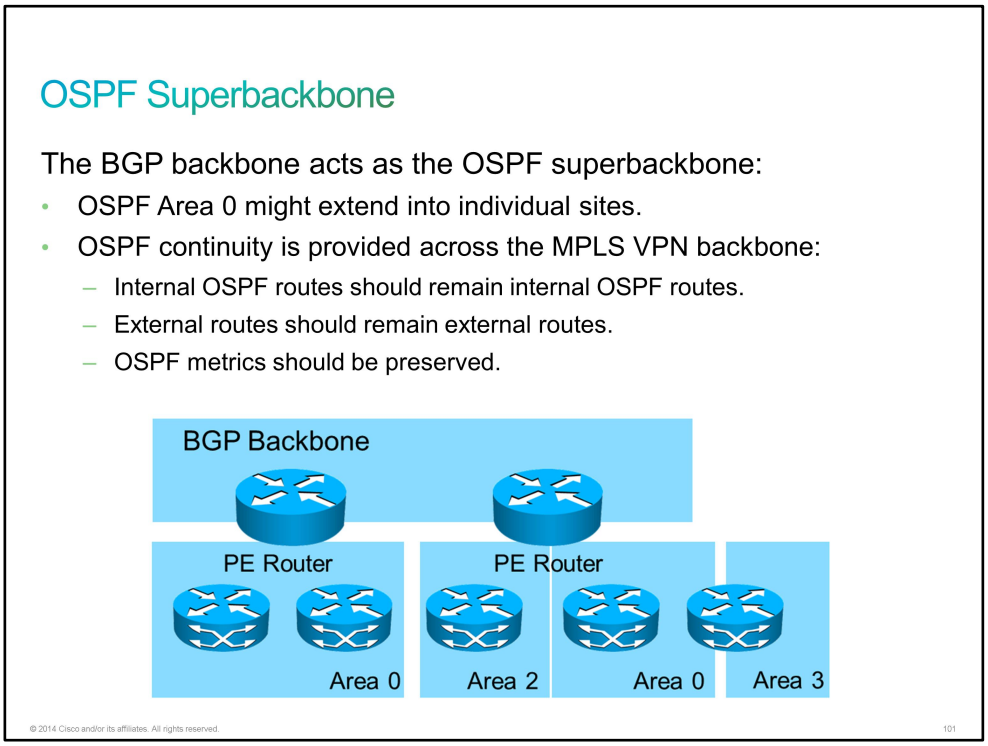
- External routes cannot be summarized.
- External routes are flooded across all OSPF areas.
- External routes could use a different metric type that is not comparable to OSPF cost.
- External routes are not inserted in stub areas or NSSAs.
- Internal routes are always preferred over external routes, regardless of their cost.

Because of all these caveats, migrating an OSPF customer towards an MPLS VPN service might have a severe impact on the routing of that customer. The MPLS VPN architecture must therefore extend the classic OSPF-BGP routing model to support transparent customer migration.

https://t.me/learningnets

OSPF Superbackbone

The MPLS VPN architecture extends the OSPF architecture by introducing another backbone above OSPF Area 0, that is the superbackbone.



The OSPF superbackbone is implemented with MP-BGP between the PE routers but is otherwise transparent to the OSPF routers. The architecture even allows disjointed OSPF backbone areas (Area 0) at MPLS VPN customer sites.

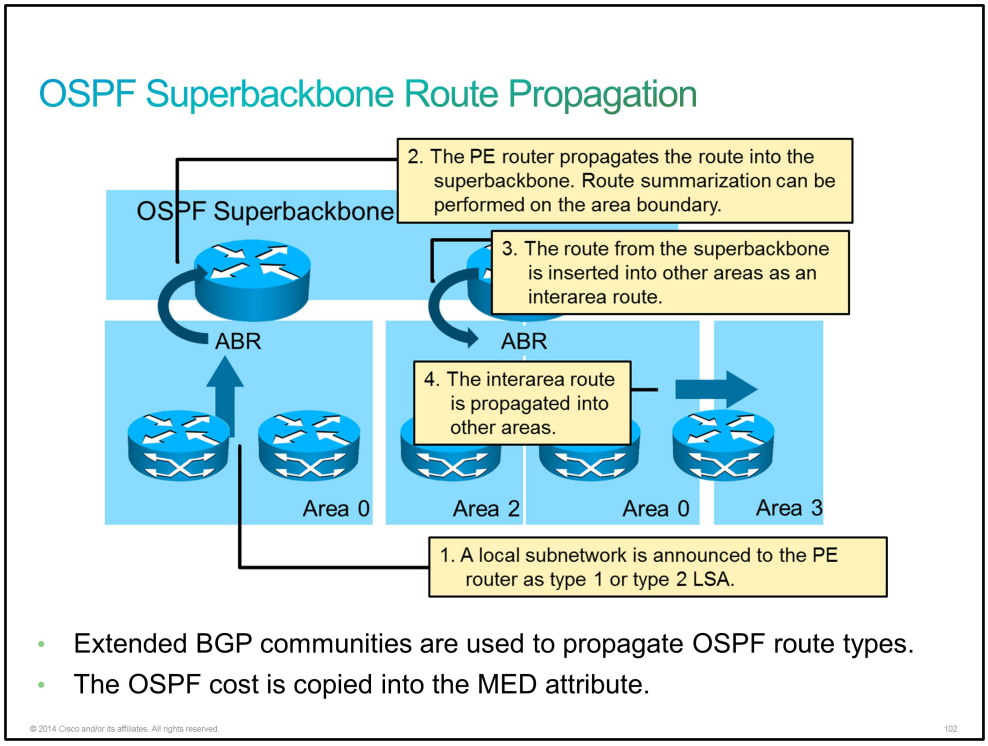
These goals must be met by the OSPF superbackbone:

- The superbackbone will not use standard OSPF-BGP redistribution.
- OSPF continuity must be provided between OSPF sites:
  - Internal OSPF routes must remain internal OSPF routes.
  - External OSPF routes must remain external OSPF routes.
  - Non-OSPF routes that are redistributed into OSPF must appear as external OSPF routes in OSPF.
  - OSPF metrics and metric types (external 1 or external 2) must be preserved.
- The OSPF superbackbone will be transparent to the CE routers that run standard OSPF software.

https://t.me/learningnets

OSPF Superbackbone Route Propagation

The MPLS VPN superbackbone appears as another layer of hierarchy in the OSPF architecture.



The PE routers that connect regular OSPF areas to the superbackbone therefore appear as OSPF ABRs in the OSPF areas to which they are attached. ABRs also appear as ASBRs in nonstub areas.

With the OSPF superbackbone architecture, the continuity of OSPF routing is preserved:

- The OSPF intra-area route, described in the OSPF router LSAor network LSA, is inserted into the OSPF superbackbone by redistributing the OSPF route into MP-BGP. Route summarization can be performed on the redistribution boundary by the PE router.
- The MP-BGP route is propagated to other PE routers and inserted as an OSPF route into other OSPF areas.

The OSPF superbackbone is implemented with the help of several BGP attributes:

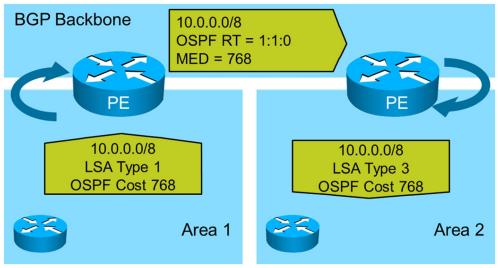
- A new BGP extended community is defined to carry the OSPF route type and OSPF area across the BGP backbone.
- As in the standard OSPF-BGP redistribution, the OSPF cost is carried in the MED attribute.

https://t.me/learningnets

OSPF Superbackbone Implementation

The first figure illustrates the propagation of internal OSPF routes across the MPLS VPN superbackbone.

OSPF Superbackbone Implementation

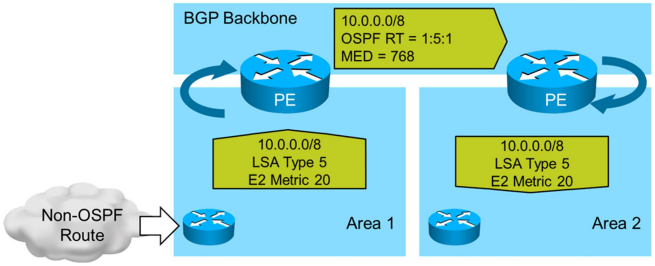


Internal OSPF routes

- The OSPF route type is copied into the extended BGP community on redistribution into BGP.
- The egress PE router performs interarea transformation.

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

OSPF Superbackbone Implementation (Cont.)



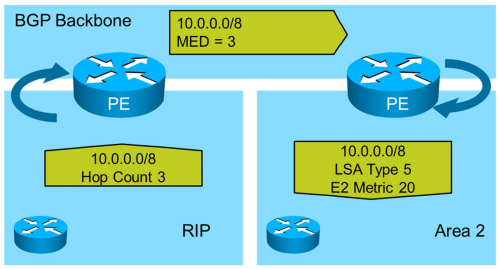
External OSPF routes

- Routes are propagated in the same way as internal OSPF routes across the superbackbone.
- The external metric and route type are preserved.

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

https://t.me/learningnets

OSPF Superbackbone Implementation (Cont.)



Routes from other routing protocols

- Routes from the MP-BGP backbone that did not originate in OSPF are still subject to standard redistribution behavior when inserted into OSPF.

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

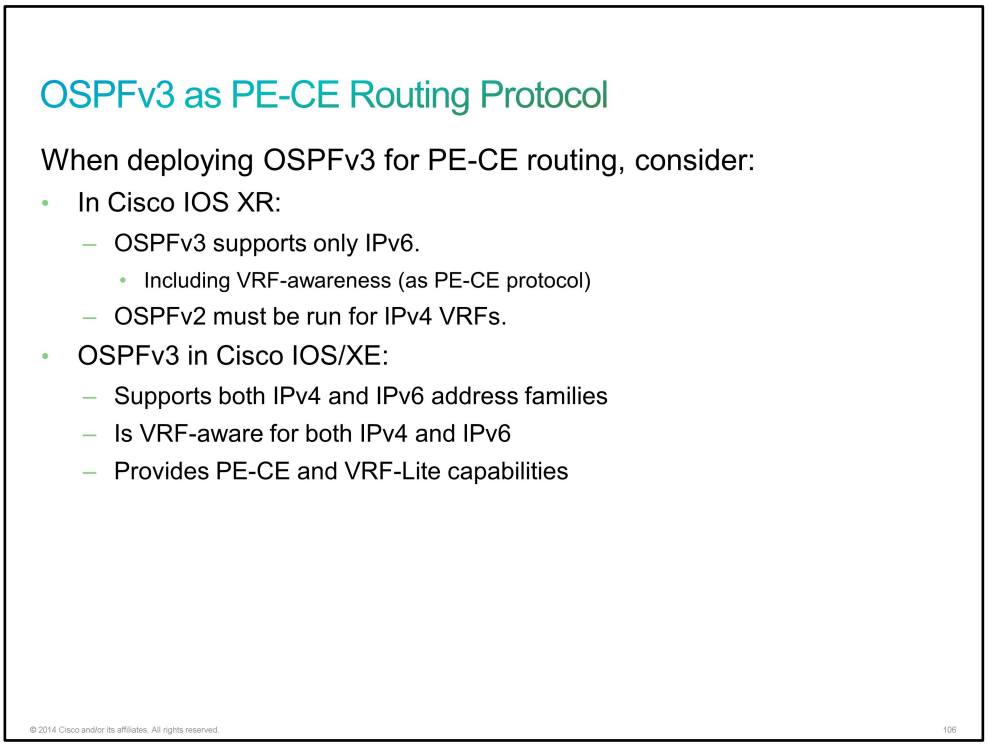
In the second figure, the external OSPF routes are redistributed into MP-BGP in the same way as the internal OSPF routes.

In the third figure, the MPLS VPN superbackbone retains the traditional OSPF-BGP route redistribution behavior for routes that did not originate in OSPF at other sites.

https://t.me/learningnets

OSPFv3 as PE-CE Routing Protocol

OSPFv3 is the successor of OSPFv2.



In Cisco IOS XR, OSPFv3 supports only the IPv6 address family. It is VRF-aware, which means that you can run it as a IPv6 PE-CE routing protocol in Layer 3 MPLS VPNs.

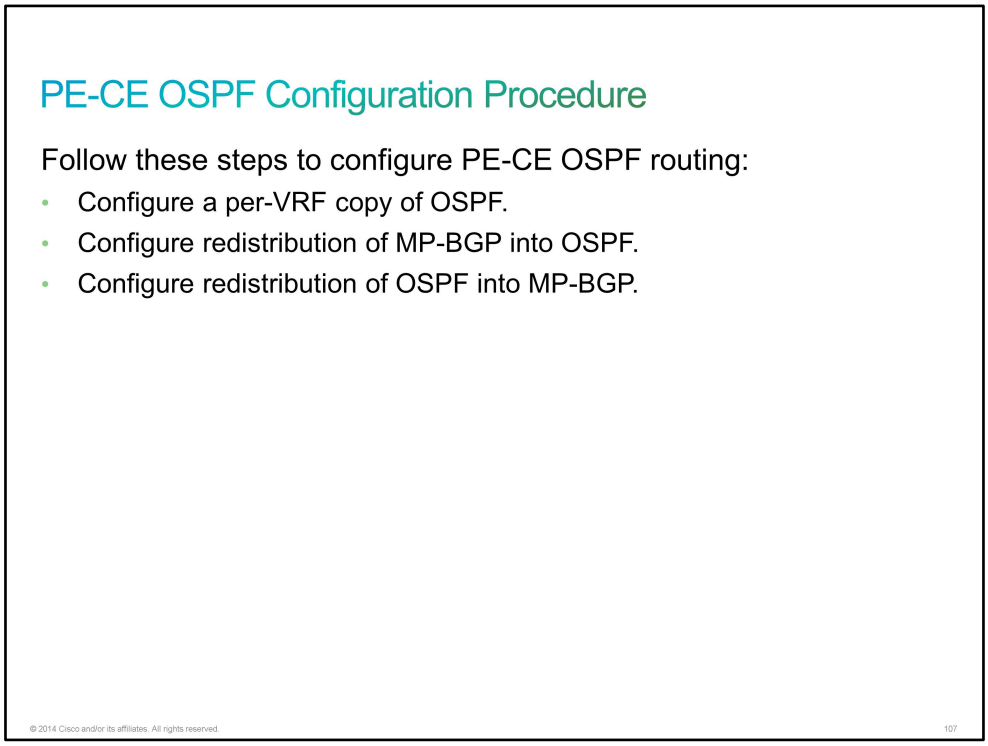
In Cisco IOS/XE, OSPFv3 supports both IPv4 and IPv6 address families, and is VRF-aware.

In Cisco IOS/XE, OSPFv3 operates in nondefault VRF instances for both IPv6 and IPv4 address families and transports the routes across a BGP or a MPLS backbone. On the PE device, customer routes are installed together by OSPFv3 and BGP in a common VRF or address family and each protocol is configured to redistribute the routes of the other. BGP combines the prefixes redistributed into it with a route-distinguisher value defined for the VRF and advertises them to other MPLS-BGP speakers in the same AS using the VPNv4 or VPNv6 address family as appropriate. The OSPFv3 route selection algorithm prefers intra-area routes across the backdoor link over inter-area routes through the MPLS backbone.

https://t.me/learningnets

PE-CE OSPF Configuration Procedure

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



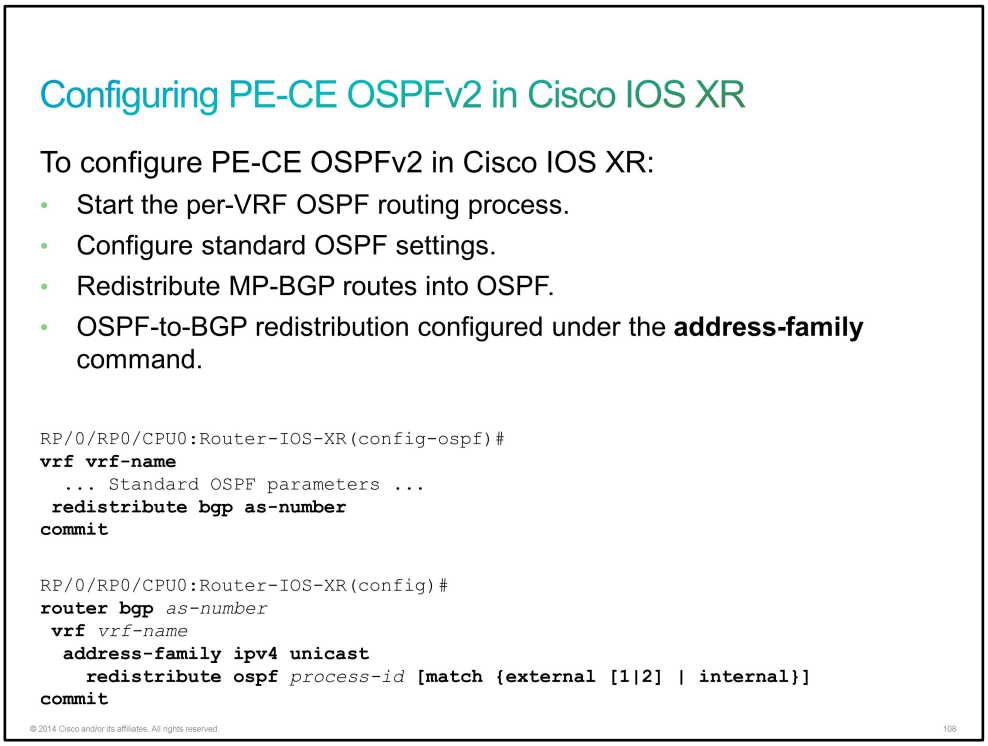
You need to redistribute OSPF routes into BGP and redistribute BGP routes into OSPF if necessary. Alternatively, you can originate a default route into a per-VRF OSPF process by using the **default-information originate always** command in router configuration mode.

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

https://t.me/learningnets

Configuring PE-CE OSPFv2 in Cisco IOS XR

In an MPLS VPN deployment on Cisco IOS/XE, each VPN VRF needs a separate OSPFv2 process when configured to run OSPFv2.

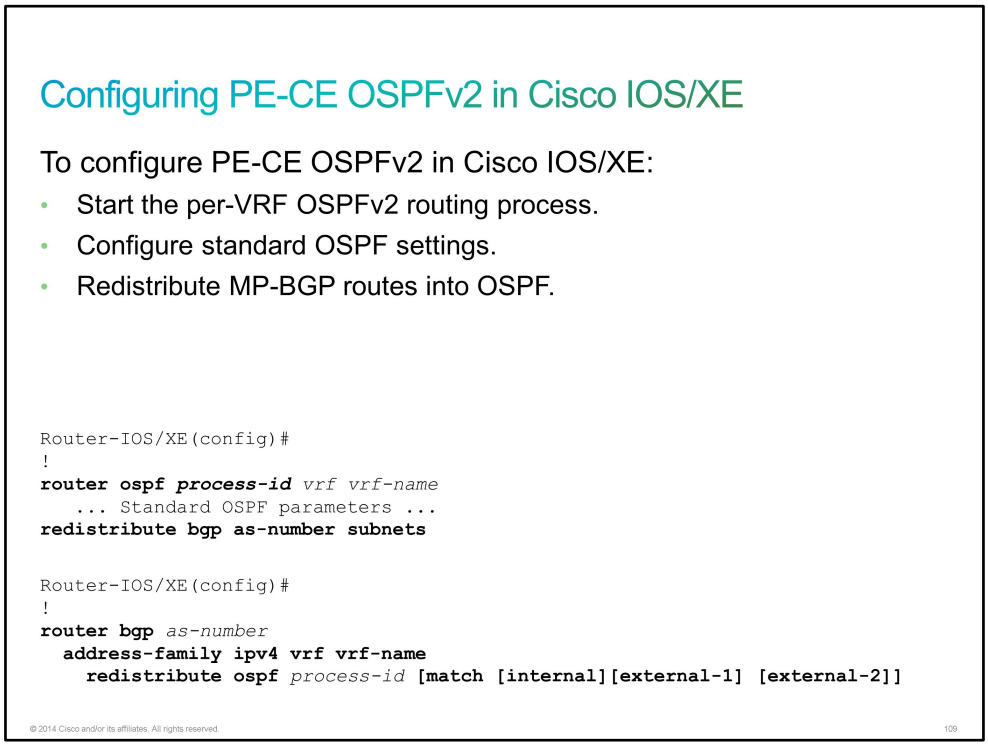


To configure an OSPF routing process within a VRF on Cisco IOS XR devices, you must first enter VRF configuration mode for OSPF routing. OSPF-BGP route redistribution is configured with the **redistribute** command under the proper **address-family** command.

https://t.me/learningnets

Configuring PE-CE OSPFv2 in Cisco IOS/XE

In an MPLS VPN deployment on Cisco IOS/XE, each VPN VRF needs a separate OSPFv2 process when configured to run OSPFv2.



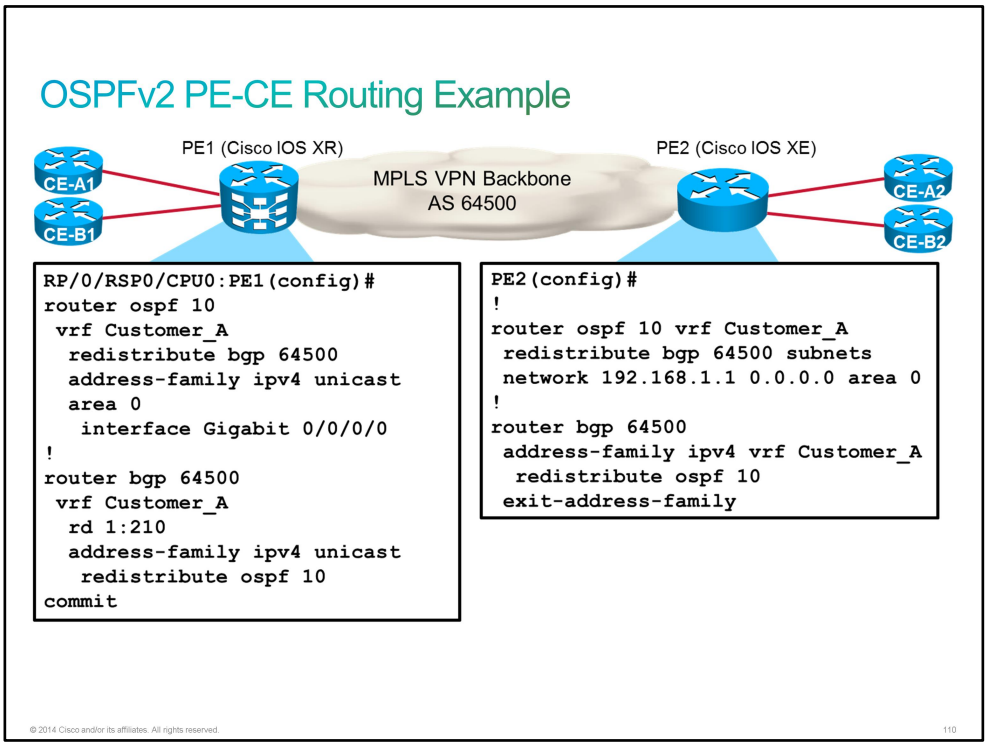
OSPFv2 support for unlimited software VRFs per PE router addresses the scalability issue for OSPF VPNs by eliminating the OSPF VPN limit of 32 processes.

To configure an OSPFv2 routing process within a VRF on Cisco IOS and IOS XE devices, use the **router ospf** command in global configuration mode. OSPF-BGP route redistribution is configured with the **redistribute** command under the proper **address-family** command. Without the OSPF **match** keyword specified, only internal OSPF routes are redistributed into BGP.

https://t.me/learningnets

OSPFv2 PE-CE Routing Example

This example illustrates basic OSPFv2 PE-CE configuration examples in Cisco IOS/XE and IOS XR.



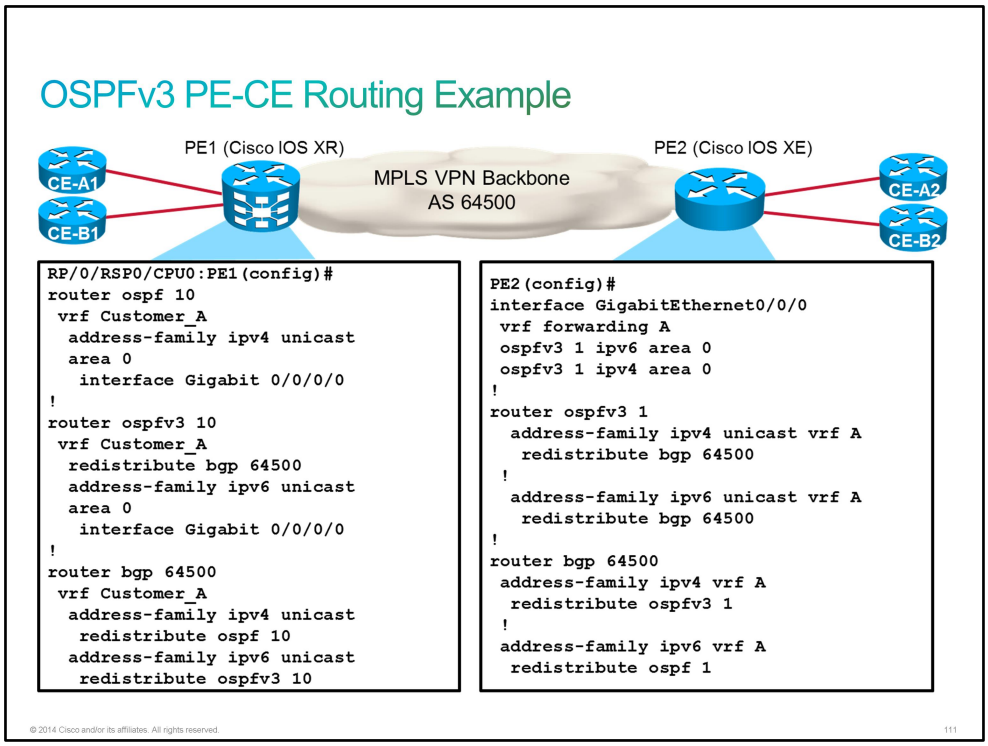
In the Cisco IOS XR OSPFv2 configuration, you can identify the interfaces on which to run OSPFv2 using the **interface** command. In this scenario, you specify the interface GigabitEthernet0/0/0/0 in the **address-family ipv4 unicast** within the VRF instance of the OSPF process.

In the Cisco IOS XE OSPFv2 configuration, you can use the **network** command to identify the interfaces on which the routing process will run.

https://t.me/learningnets

OSPFv3 PE-CE Routing Example

This topic illustrates basic OSPFv3 PE-CE configuration examples in Cisco IOS/XE and IOS XR.



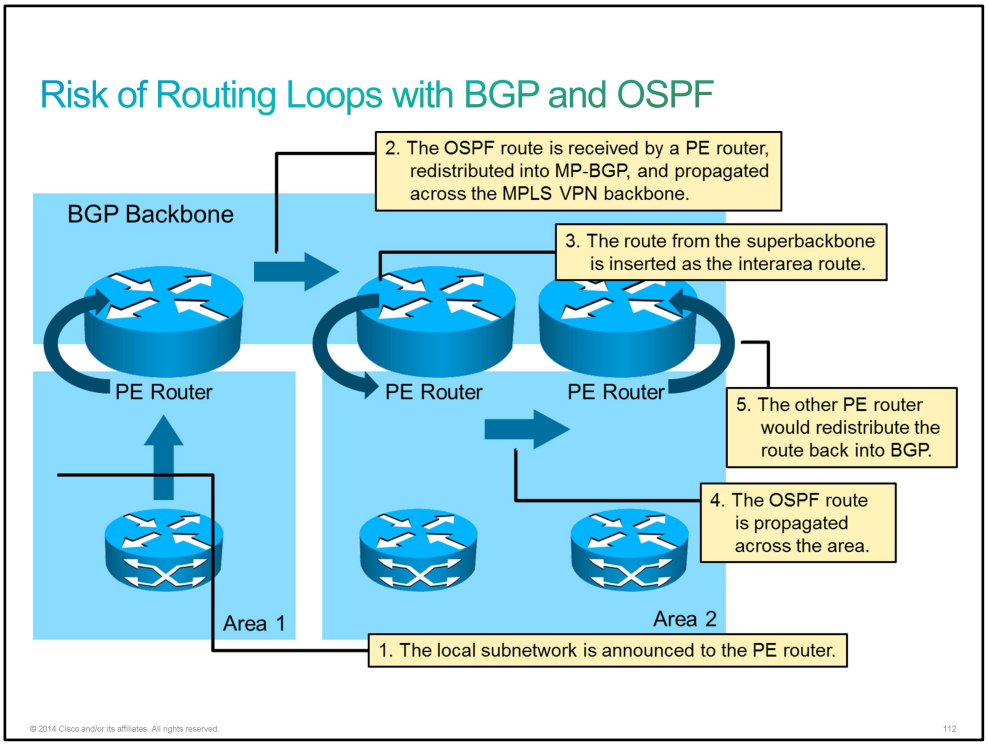
When configuring OSPFv3 in Cisco IOS XR, you can identify the interfaces on which to run OSPFv3 using the **interface** command. In this scenario, you specify the interface GigabitEthernet0/0/0/0 in the **address-family ipv6 unicast** within the VRF instance of the OSPFv3 process. If you want to configure dual-stack OSPF routing, you need to start a separate OSPFv2 process, possibly with the same process ID, to provide IPv4 routing.

To configure OSPFv3 in Cisco IOS and Cisco IOS XE, you enable the required address families (IPv4 or IPv6) for the given VRF and then activate the specific OSPFv3 for the desired address family and area ID in the interface configuration mode.

https://t.me/learningnets

Risk of Routing Loops with BGP and OSPF

OSPF developers have taken many precautions to avoid routing loops between OSPF areas. For example, intra-area routes are always preferred over interarea routes. However, these rules do not work when the superbackbone is introduced.

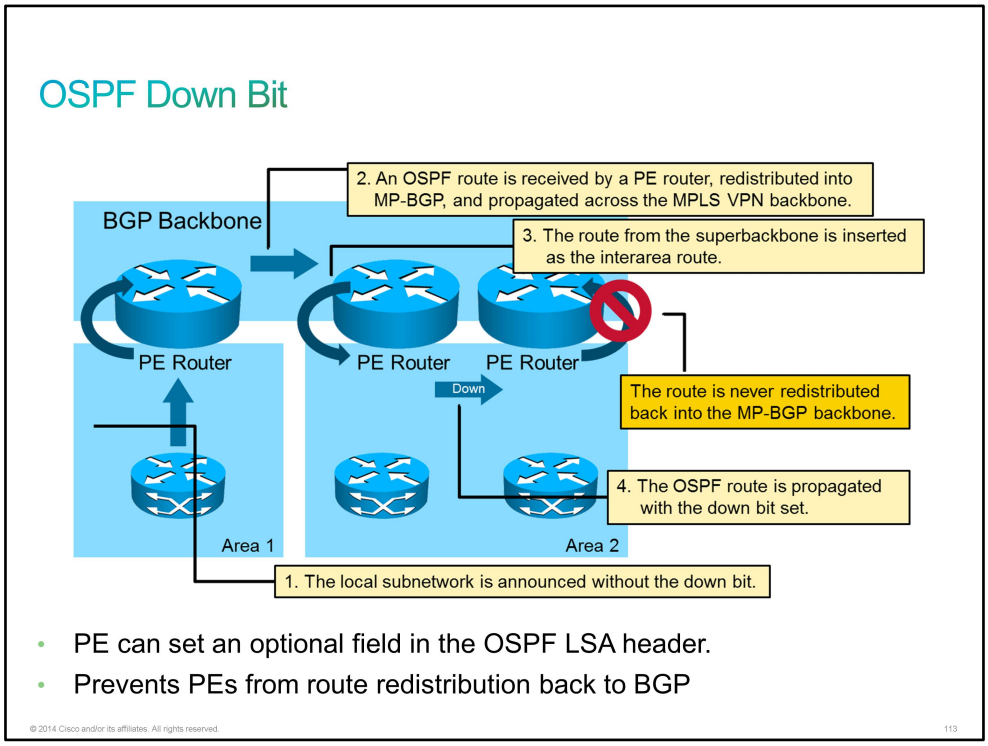


For example, consider the network in the figure, where the receiving OSPF area has two PE routers attached to it.

https://t.me/learningnets

OSPF Down Bit

The down bit was introduced as a mechanism to prevent route redistribution loops between the OSPF and MP-BGP that is running between PE routers. The OSPF down bit is part of the Options field in the OSPF LSA header.

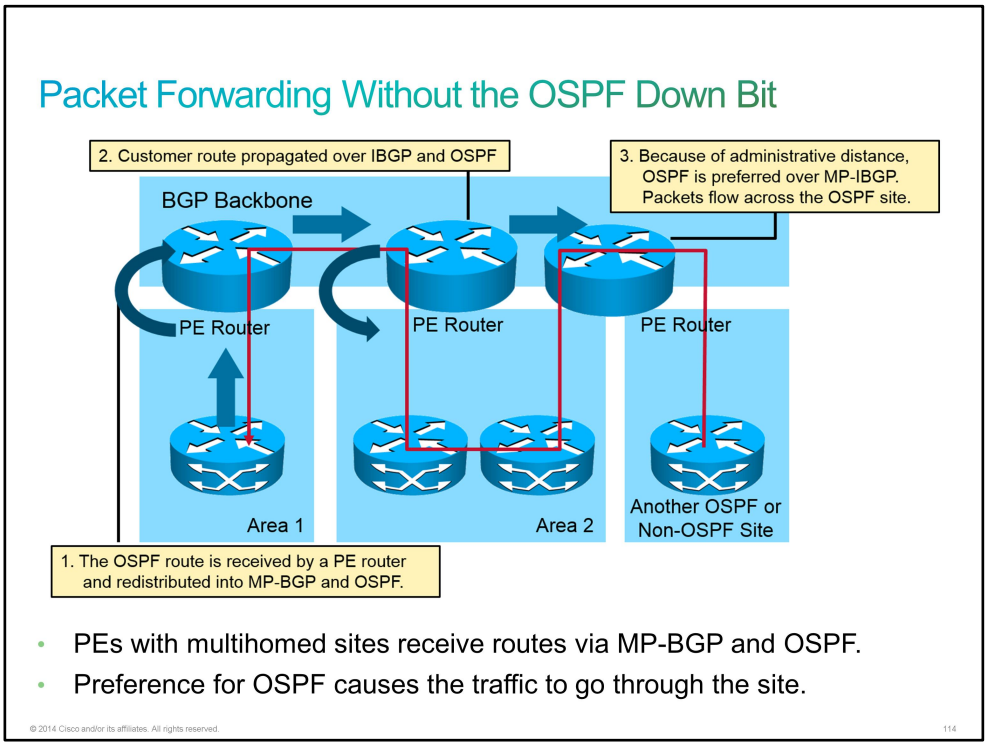


The down bit is used between the PE routers to indicate which routes were inserted into the OSPF topology database from the MPLS VPN superbackbone. Those routes, thus, will not be redistributed back into the MPLS VPN superbackbone. The PE router that redistributes the MP-BGP route as an OSPF route into the OSPF topology database sets the down bit. The other PE routers use the down bit to prevent this route from being redistributed back into MP-BGP.

https://t.me/learningnets

Packet Forwarding Without the OSPF Down Bit

The OSPF superbackbone implementation with MP-BGP has other implications beyond the potential for routing loops between OSPF and BGP.

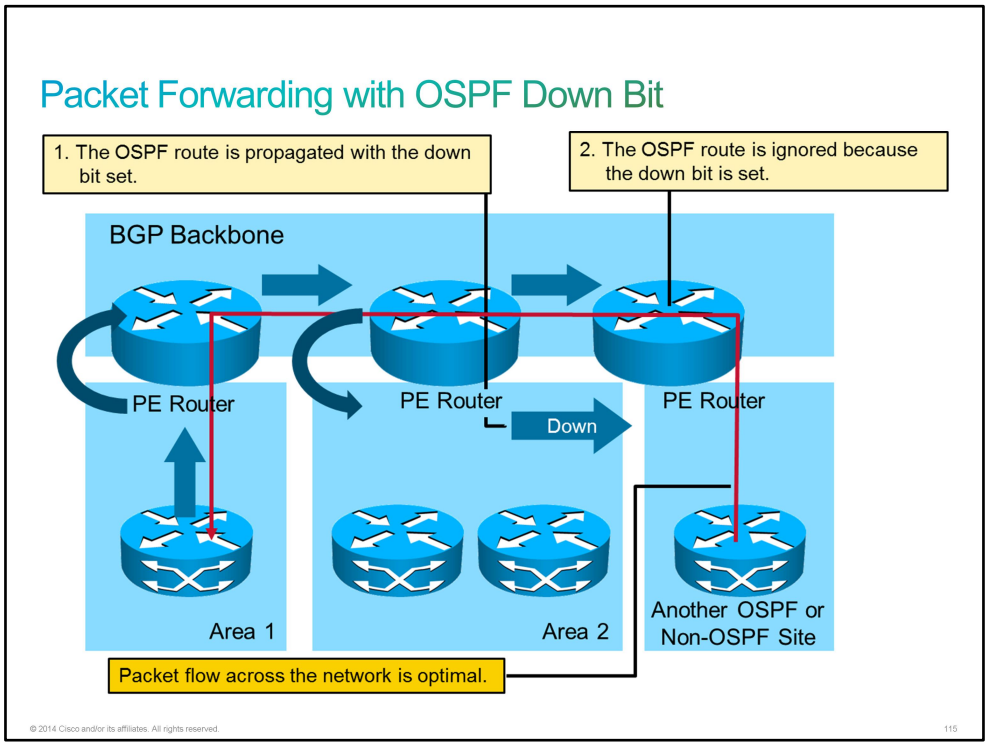


This example shows a typical flow for routing updates. Customer routes are redistributed into BGP on the left PE. The middle PE redistributes them into the customer site via OSPF. The PE on the right receives the customer networks via OSPF and internal MP-BGP. Due to better administrative distance, the OSPF routes are preferred. The traffic would flow via the OSPF site, which is undesirable in most situations.

https://t.me/learningnets

Packet Forwarding with the OSPF Down Bit

To prevent customer sites from acting as transit parts of the MPLS VPN network, the OSPF route-selection rules in PE routers need to be changed.



The PE routers must ignore all OSPF routes with the down bit set, because these routes originated in the MP-BGP backbone, and the MP-BGP route should be used as the optimum route toward the destination.

This rule is implemented with the routing bit in the OSPF LSA. For routes with the down bit set, the routing bit is cleared and these routes never enter the IP routing table, even if they are selected as the best routes by the SPF algorithm.

Note	The routing bit is the Cisco extension to OSPF and is used only internally in the router. The routing bit is never propagated between routers in LSA updates.

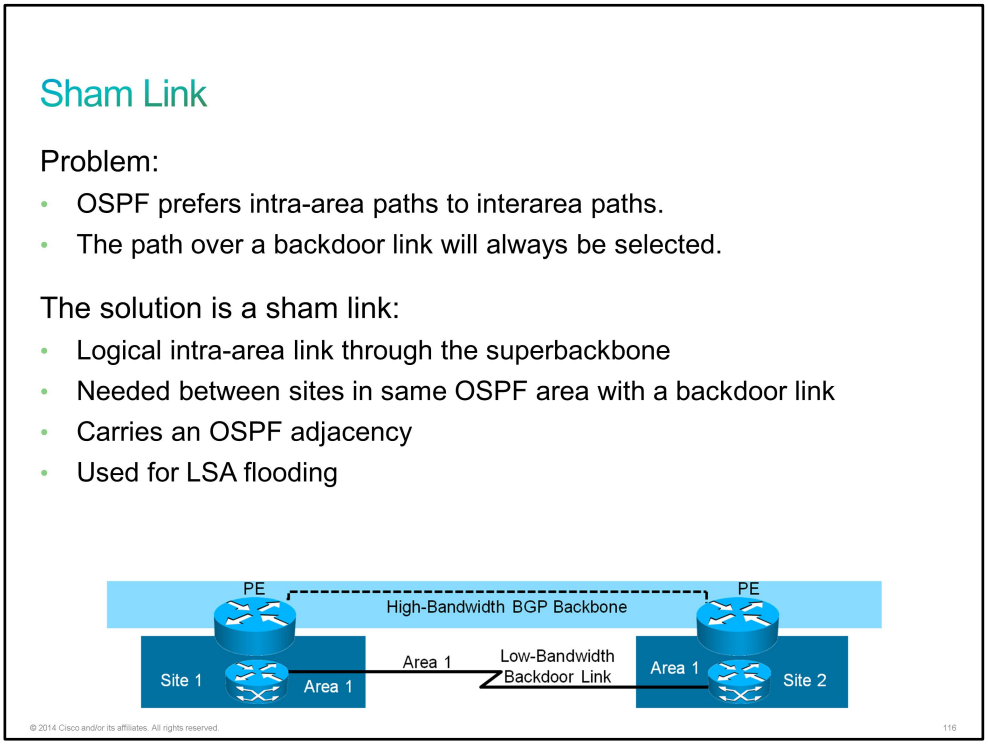
With the new OSPF route-selection rules in place, packet forwarding in the network that is shown in the figure follows the desired path.

https://t.me/learningnets

OSPF Sham Link Overview

Sham Link

Although OSPF PE-CE connections assume that the only path between two client sites is across the MPLS VPN backbone, backdoor paths between VPN sites may exist.



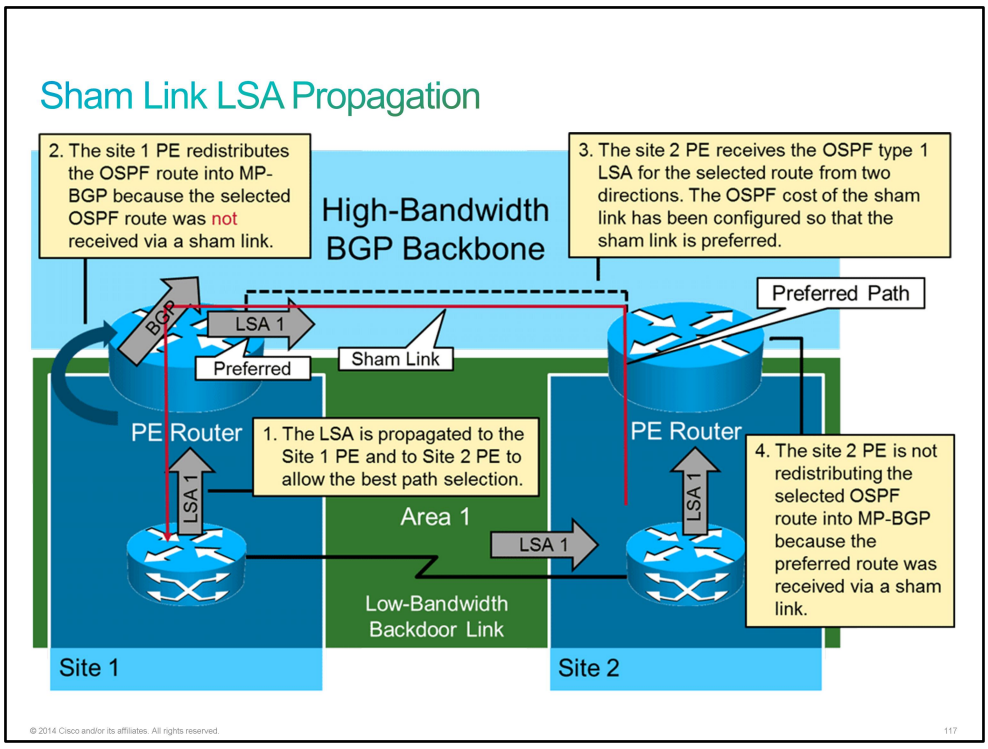
The figure shows backdoor paths between VPN sites. If these sites belong to the same OSPF area, the path over a backdoor link will always be selected, because OSPF prefers intra-area paths to interarea paths. (PE routers advertise OSPF routes learned over the VPN backbone as interarea paths.) For this reason, OSPF backdoor links between VPN sites must be taken into account so that routing is performed based on policy.

Because each site runs OSPF within the same Area 1 configuration, all routing between the sites follows the intra-area path across the backdoor links rather than over the MPLS VPN backbone.

A sham link is required between any two VPN sites that belong to the same OSPF area *and* share an OSPF backdoor link. If no backdoor link exists between the sites, no sham link is required.

https://t.me/learningnets

Sham Link LSA Propagation



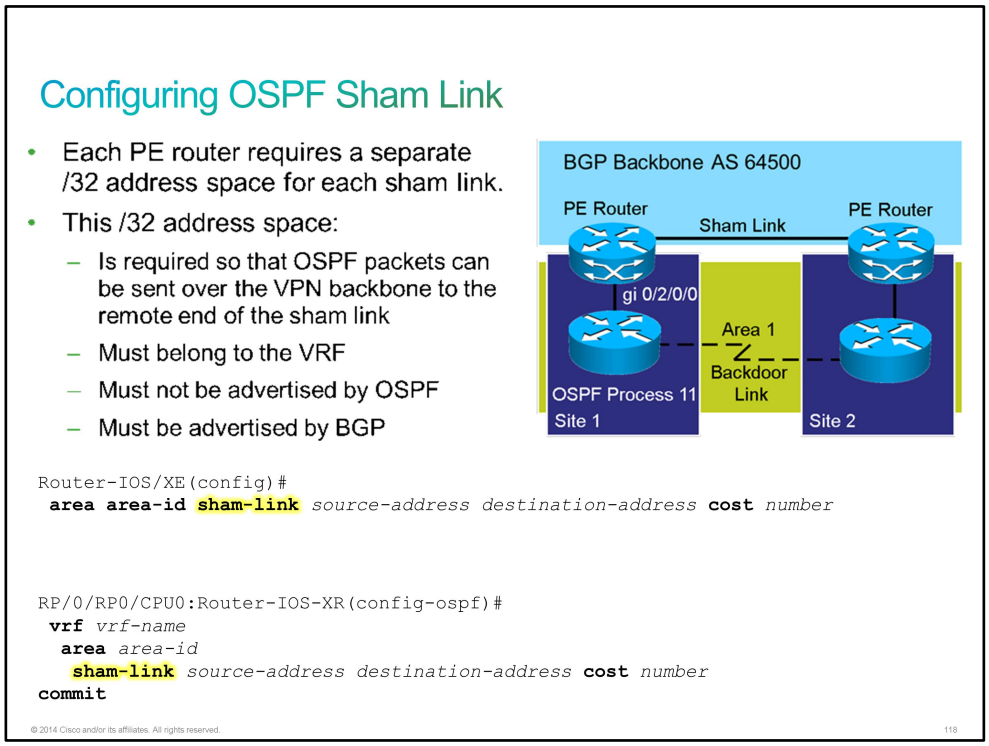
Because the sham link is seen as an intra-area link between PE routers, an OSPF adjacency is created, and a database exchange (for the particular OSPF process) occurs across the link.

After the desired intra-area connectivity is created, the PE routers prefer routes that are learned via the high-bandwidth backbone, because the OSPF cost of the sham link has been configured so that it is preferred over the backdoor link. The implementation results in optimal packet flow.

https://t.me/learningnets

Configuring OSPF Sham Link

When you are configuring a sham link, a separate /32 address space is required in each PE router.



These criteria apply to this /32 address space:

- Required so that OSPF packets can be sent over the VPN backbone to the remote end of the sham link
- Must belong to the VRF
- Must not be advertised by OSPF
- Must be advertised by BGP

To configure a sham-link interface on a PE router that is running Cisco IOS and IOS XE Software in an MPLS VPN backbone, use the **area sham-link cost** command in global configuration mode.

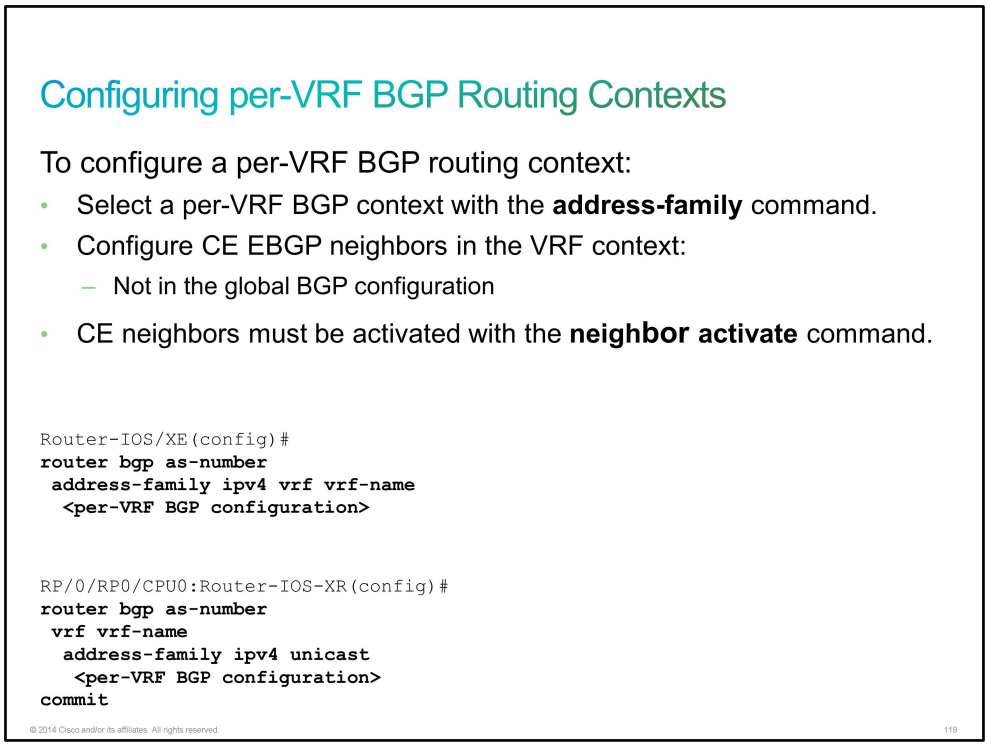
To configure a sham-link interface using Cisco IOS XR Software, you must first enter the VRF configuration mode for OSPF routing and configure an area for the OSPF process using the **area** command. Then you must assign interfaces to the sham link using the **sham-link** command.

https://t.me/learningnets

BGP as the PE-CE Routing Protocol

Configuring per-VRF BGP Routing Contexts

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

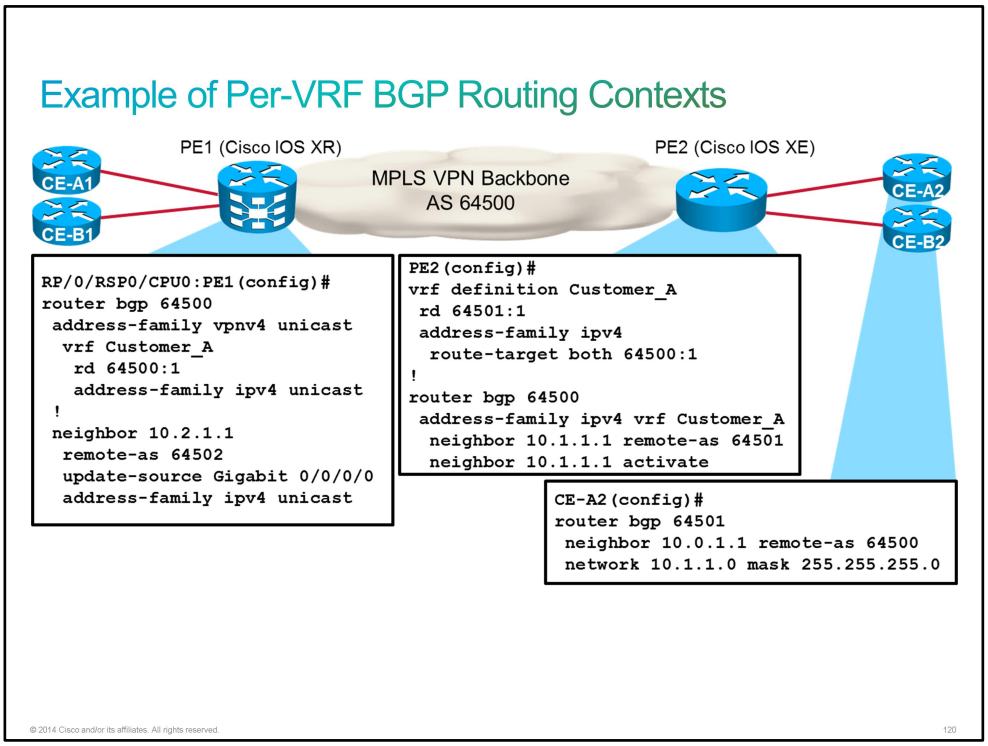


On Cisco IOS XR devices, first define the VRF with the **vrf** *vrf-name* command in the BGP routing processes. Then select the routing context with the **address-family ipv4 unicast** command. All per-VRF routing protocol parameters (network numbers, passive interfaces, neighbors, filters, and so on) are configured under this address family.

Example of Per-VRF BGP Routing Contexts

The figure shows the BGP configurations on the PE routers and a CE router.

https://t.me/learningnets

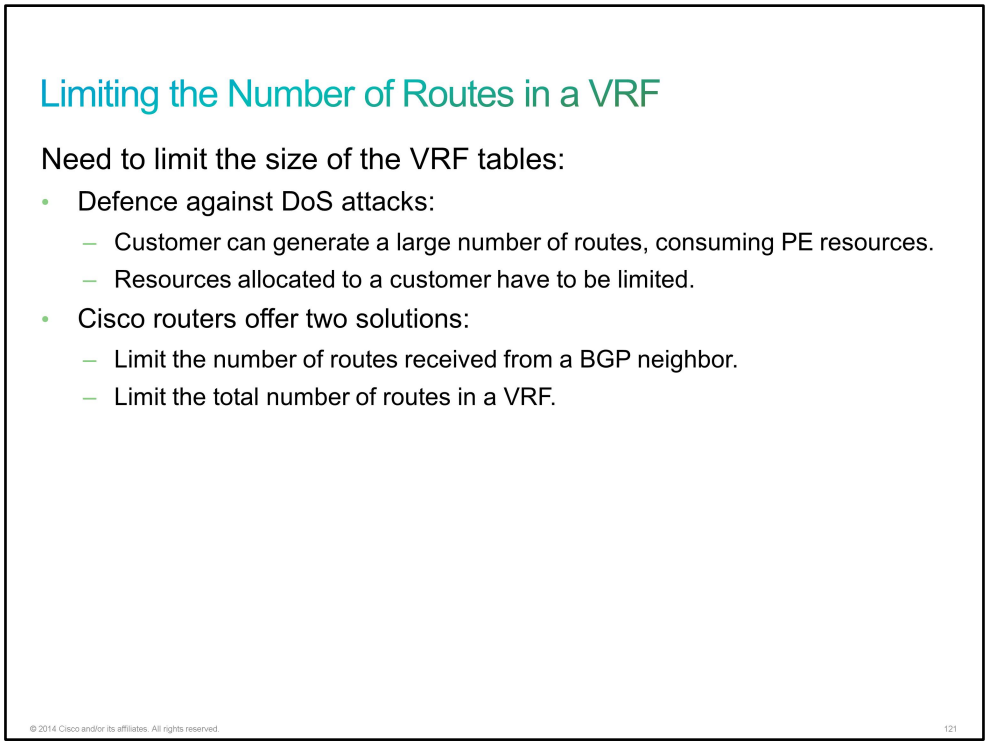


PE1 runs Cisco IOS XR. PE2 runs Cisco IOS XE. Because all CE routers in the example are running Cisco IOS Software, the configuration of the remaining CE routers is omitted.

https://t.me/learningnets

Limiting the Number of Routes in a VRF

MPLS VPN architecture achieves a tight coupling between the customer and the service provider network, resulting in a number of advantages. The tight coupling could also result in a few disadvantages, because the service provider network is exposed to design and configuration errors in customer networks, and a number of new DoS attacks are based on routing protocol behavior.

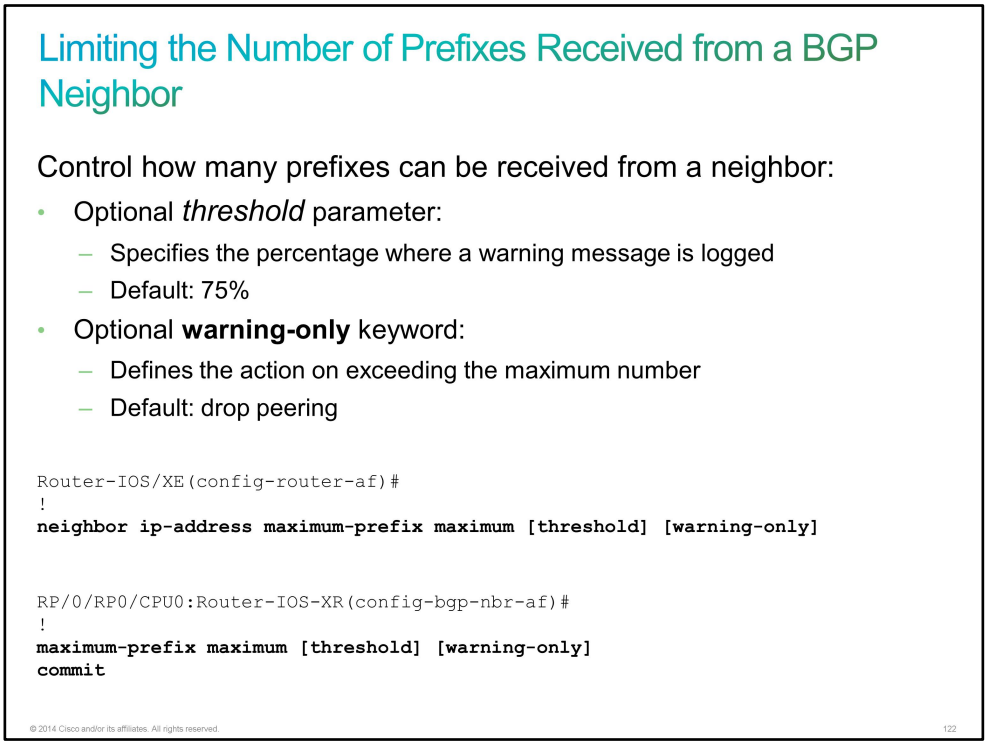


To limit the effect of configuration errors and malicious users, Cisco routers offer two features that limit the number of routes and the resource consumption that are available to a VPN user at a PE router:

- The BGP Maximum-Prefix feature limits the number of routes that an individual BGP peer can send.
- The VRF route limit restricts the total number of routes in a VRF regardless of whether those routes are received from CE routers or from other PE routers via MP-IBGP.

https://t.me/learningnets

Limiting the Number of Prefixes Received from a BGP Neighbor



To control how many prefixes can be received from a neighbor, use the **maximum-prefix** command for the peer for the appropriate address family.

- Cisco IOS and IOS XE Software:

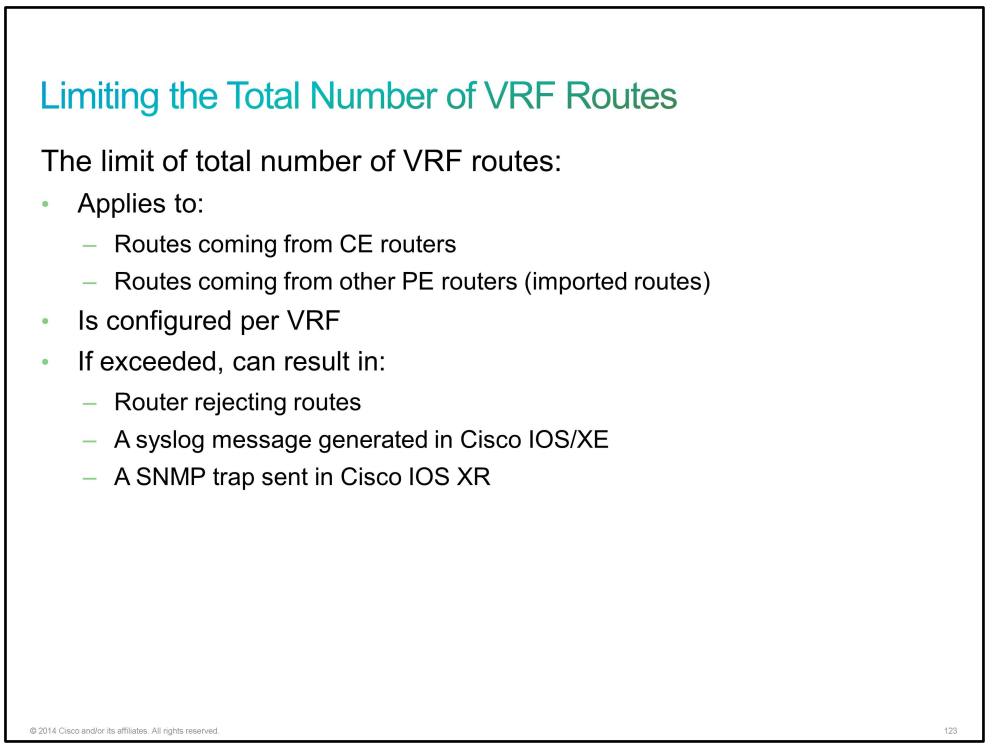
neighbor *{ip-address | peer-group-name}* **maximum-prefix maximum** *[threshold]* [**restart** *restart-interval]* **[warning-only]**

- Cisco IOS and IOS XE Software:

maximum-prefix maximum *[threshold]* **[warning-only]**

https://t.me/learningnets

Limiting the Total Number of VRF Routes



The VRF route limit, unlike the BGP maximum-prefix limit, limits the overall number of routes in a VRF regardless of their origin. As with the BGP Maximum-Prefix feature, the network operator might be warned by a syslog message or SNMP trap when the number of routes exceeds a certain threshold. Additionally, you can configure Cisco IOS and IOS XE and Cisco IOS XR Software to ignore new VRF routes when the total number of routes exceeds the maximum configured limit.

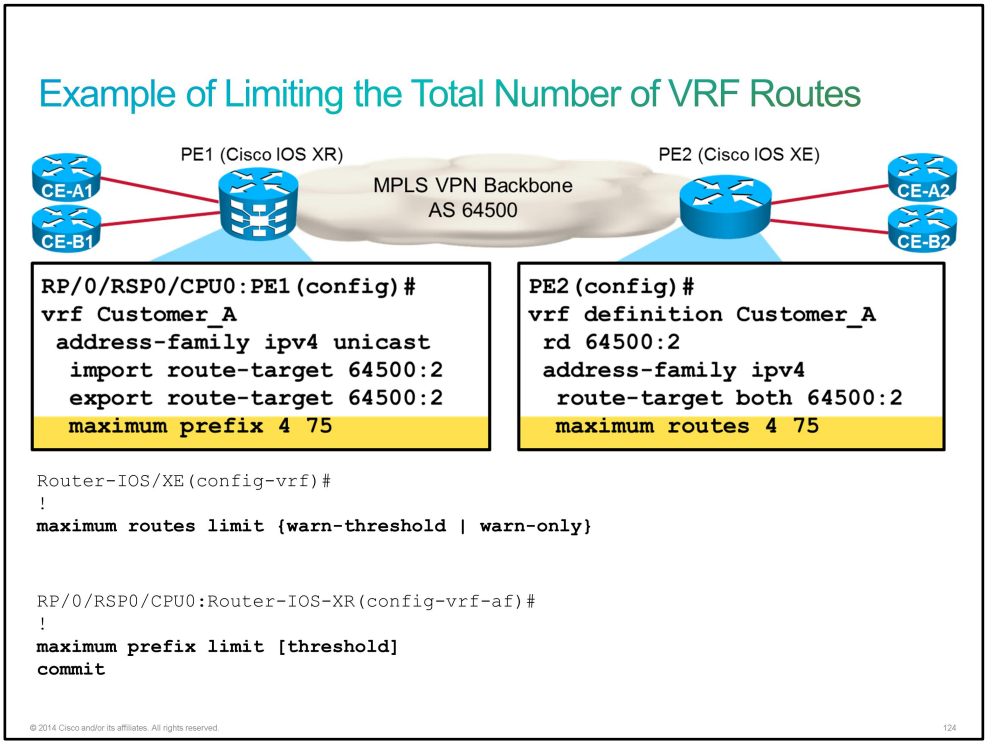
The route limit is configured for each individual VRF, providing maximum design and configuration flexibility.

Note	The per-VRF limit could be used to implement add-on MPLS VPN services. A user wanting a higher level of service might be willing to pay to be able to insert more VPN routes into the network.

https://t.me/learningnets

Example of Limiting the Total Number of VRF Routes

The network designer can decide to limit the number of routes in a VRF.



In the figure, the network designer has decided to limit the number of routes in a VRF to four, with the warning threshold being set at 75 percent (or three routes).

To limit the maximum number of routes in a VRF instance to prevent a PE router from importing too many routes, use the **maximum routes** command in VRF configuration mode (IOS and IOS XE Software) or VRF address family configuration mode (IOS XR Software).

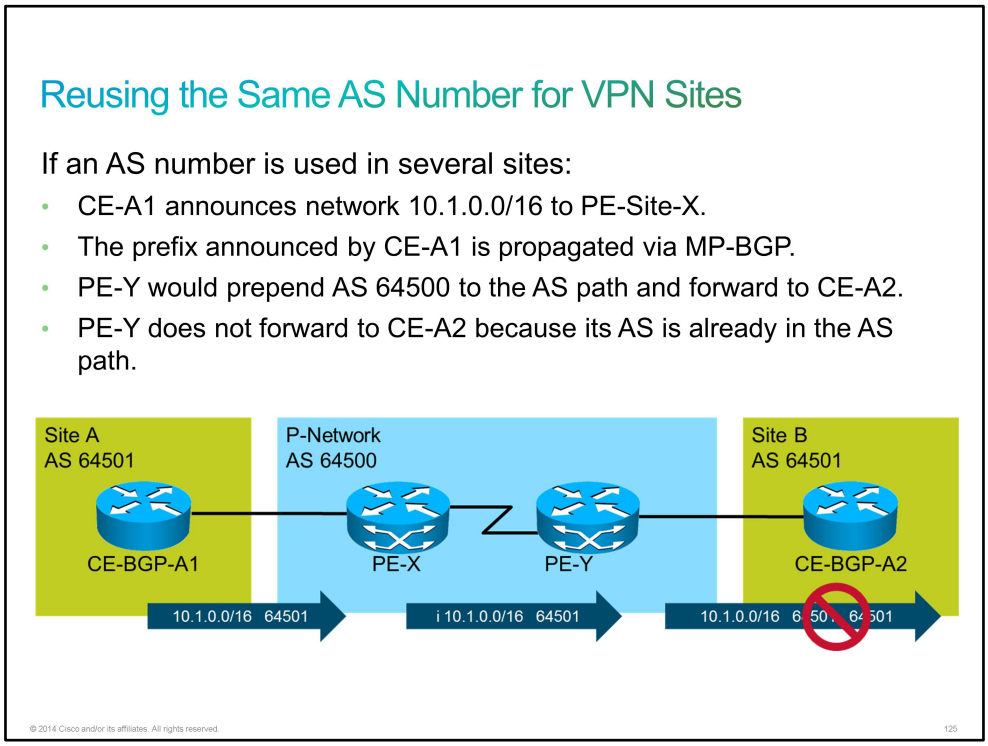
When the first two routes are received and inserted into the VRF, the router accepts them. When the third route is received, a warning message is generated, and the message is repeated with the insertion of the fourth route.

When the PE router receives the fifth route, the maximum route limit is exceeded, and the route is ignored. The network operator is notified through another syslog message.

https://t.me/learningnets

Reusing the Same AS Number for VPN Sites

With BGP used as the PE-CE routing protocol, the customer sites may or may not use different AS numbers.



There are two ways that an MPLS VPN customer can deploy BGP as the routing protocol between PE and CE routers:

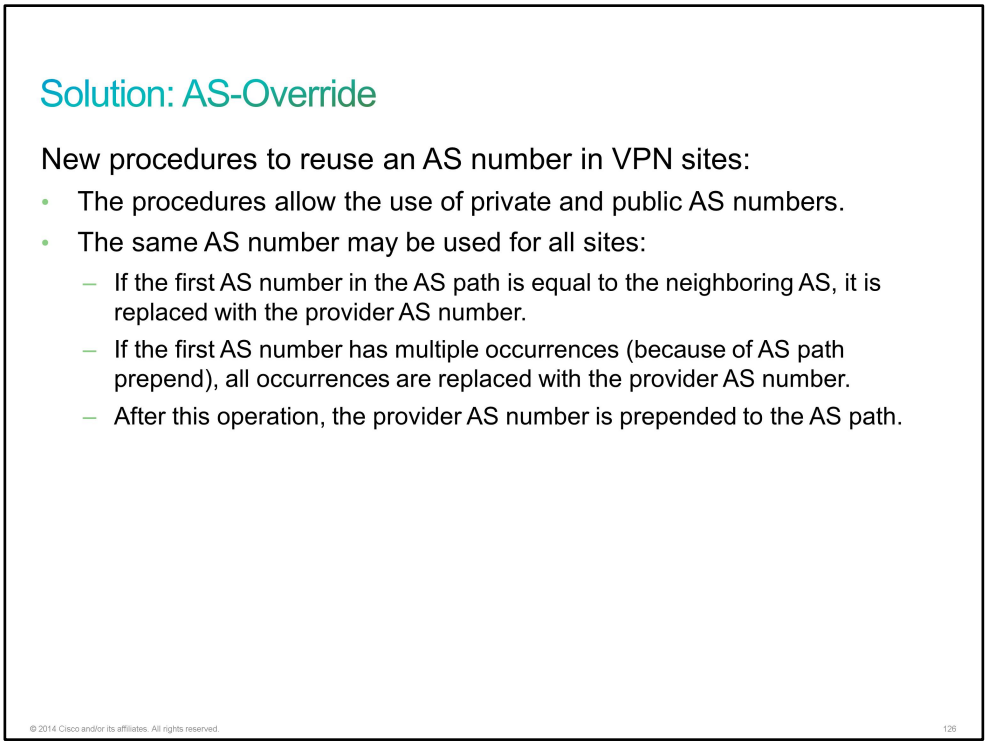
- If the customer has previously used any other routing protocol in the traditional overlay VPN network, there are no limitations on the numbering of the customer autonomous systems. Every site can be a separate AS.
- If the customer has used BGP as the routing protocol before, there is a good chance that all the sites (or a subset of the sites) are using the same AS number.

BGP loop-prevention rules disallow discontiguous autonomous systems. Two customer sites with identical AS numbers cannot be linked by another AS. If such a setup happens (as in the example in the figure), the routing updates from one site will not be forwarded to other VPN sites. There would be no connectivity between the sites.

https://t.me/learningnets

Solution: AS-Override

When you are migrating customers from traditional overlay VPNs to MPLS VPNs, it is not uncommon to encounter a customer topology that requires a customer AS number to be used at more than one site. This requirement can cause problems with the loop-prevention rules of BGP. However, the AS-path update procedure in BGP has been modified to address this issue.



The new AS-path update procedure supports the use of one AS number at many sites (even between several overlapping VPNs) and does not rely on a distinction between private and public AS numbers.

The modified AS-path update procedure is called *AS* override:

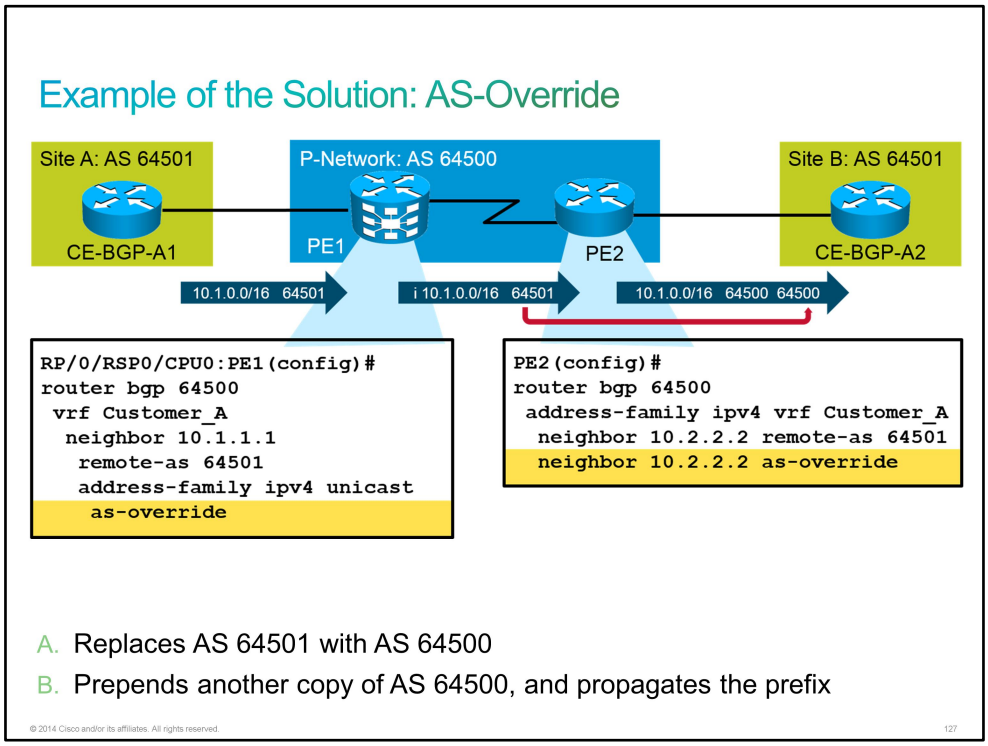
- The procedure is used only if the first AS number in the AS path is equal to the AS number of the receiving BGP router.
- In this case, all leading occurrences of the AS number of the receiving BGP router are replaced with the AS number of the sending BGP router. Occurrences that are farther down the AS path of the AS number of the receiving router are not replaced because they indicate a real routing information loop.
- An extra copy of the sending router AS number is prepended to the AS path. The standard AS number prepending procedure occurs on every EBGP update.

To configure a PE router to override a site AS number with a provider AS number, use the **as-override** command in the appropriate address family.

https://t.me/learningnets

Example of the Solution: AS-Override

In this figure, customer sites A and B use BGP to communicate with the MPLS VPN backbone. Both sites use AS 64501. Site B would drop the update that was sent by site A without the AS-override mechanism.

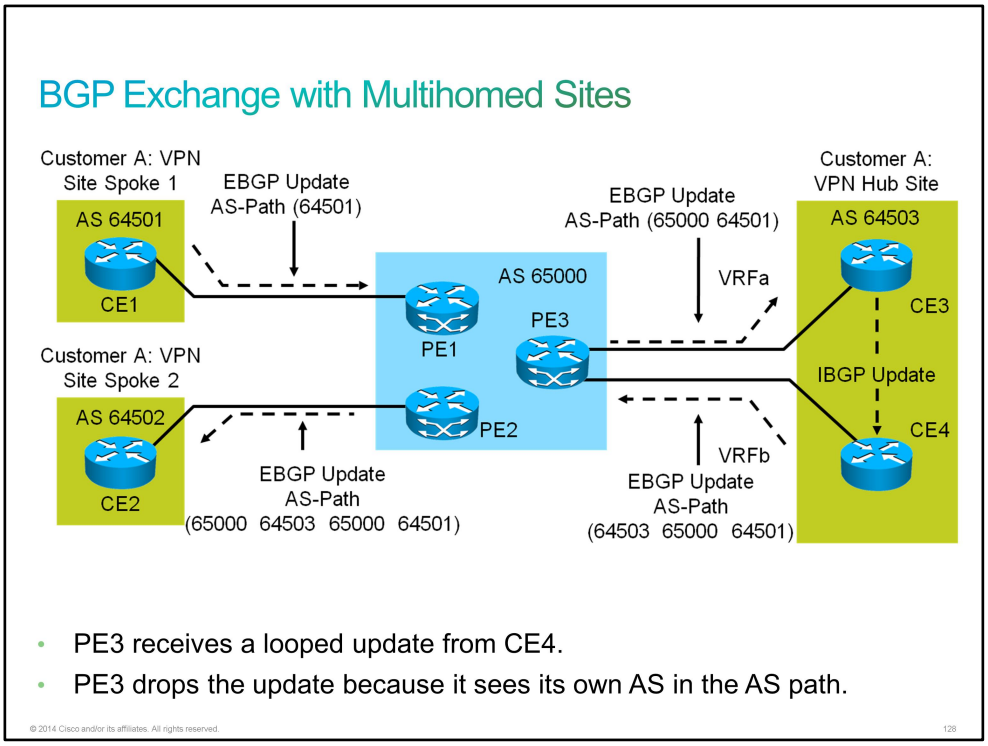


The AS-override mechanism, configured on the PE2 router, replaces the customer AS number (64501) with the provider AS number (64500) before sending the update to the customer site. An extra copy of the provider AS number is prepended to the AS path during the standard EBGP update process.

https://t.me/learningnets

BGP Exchange with Multihomed Sites

Consider a hub-and-spoke scenario that requires you to permit the routes that are coming from the VRF hub site to re-enter the AS of the service provider. To do so requires that the spoke-to-spoke communication happen through the VRF hub site.



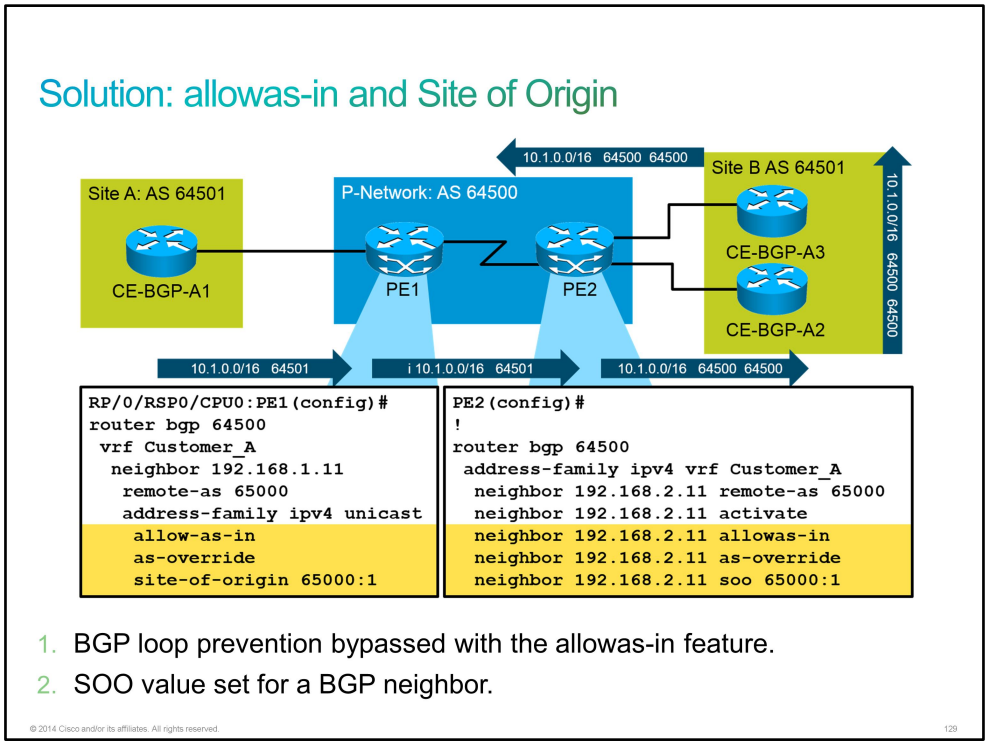
The hub site connects to the provider with two links that belong to two different VRFs on PE3. One link is used to send updates to the hub site, and one link is used to receive updates from the hub site. For BGP, this setup implies that a route traverses the service provider AS from a VRF spoke site to the VRF hub site and traverses it again on the way to another VRF spoke site. The PE3 router that connects to the VRF hub site sees its own AS number in the AS path, so the BGP route is rejected.

To disable the AS path loop check, you can configure the command **neighbor allowas-in number** on the PE3 router that connects to the VRF hub site. The **allowas-in** command permits multiple occurrences of the same AS number (in this case, the AS number of the service provider) as the AS number of the BGP speaker in the AS path without BGP denying the route. You can configure a number from 1 to 10 to specify the number of times that the AS number is allowed in the AS path.

https://t.me/learningnets

Solution: allowas-in and Site of Origin

Most aspects of BGP loop prevention are bypassed when you use either the **as-override** or the **allowas-in** command. Routing information loops can still be detected by manually counting occurrences of an AS number in the AS path in an end-to-end BGP routing scenario and then ensuring that the number field in the **allowas-in** command is set low enough to prevent loops.



The ability to continue to detect loops can present a particular problem when BGP is mixed with other PE-CE routing protocols. The SOO extended BGP community can be used as an additional loop-prevention mechanism in these situations.

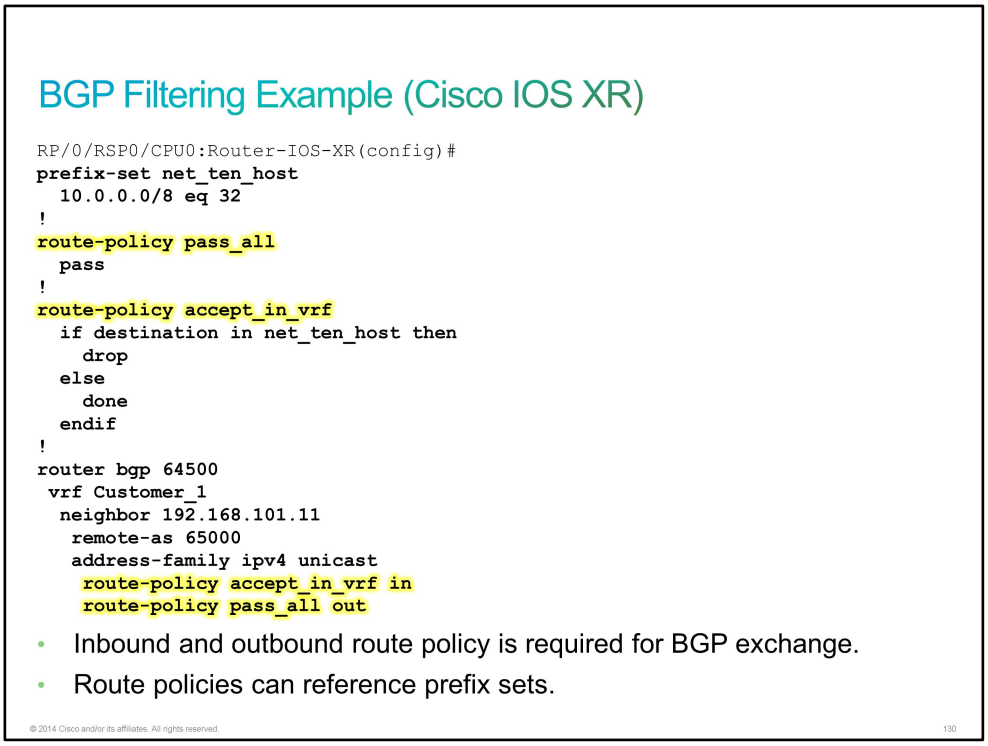
The SOO uniquely identifies the site that originates a route. It is a BGP extended community that prevents routing loops or suboptimal routing, specifically when a back door is present between VPN sites. The SOO provides loop prevention in networks with dual-homed or multihomed sites (sites that are connected to two or more PE routers). You can use it when an IGP is the PE-CE routing protocol. You can also use it when BGP is used between PE and CE, when the AS path loop prevention cannot be trusted anymore. This situation happens when BGP uses **as-override** or **allowas-in**. If the SOO is configured for a CE router and a VPNv4 route is learned with the same SOO, the route must not be put in the VRF routing table on the PE and advertised to the CE.

Use this command to set the SOO value for a BGP neighbor. The SOO value is set under address family IPv4 VRF configuration mode either directly for a neighbor or for a BGP peer group.

https://t.me/learningnets

BGP Filtering Example (Cisco IOS XR)

When you deploy BGP as the PE-CE routing protocol, you may want to filter network updates exchanged with the customer routers. The filtering tool in Cisco IOS XR is based on route policies. You need an inbound and an outbound route policy applied to a VRF EBGP neighbor to exchange any updates.

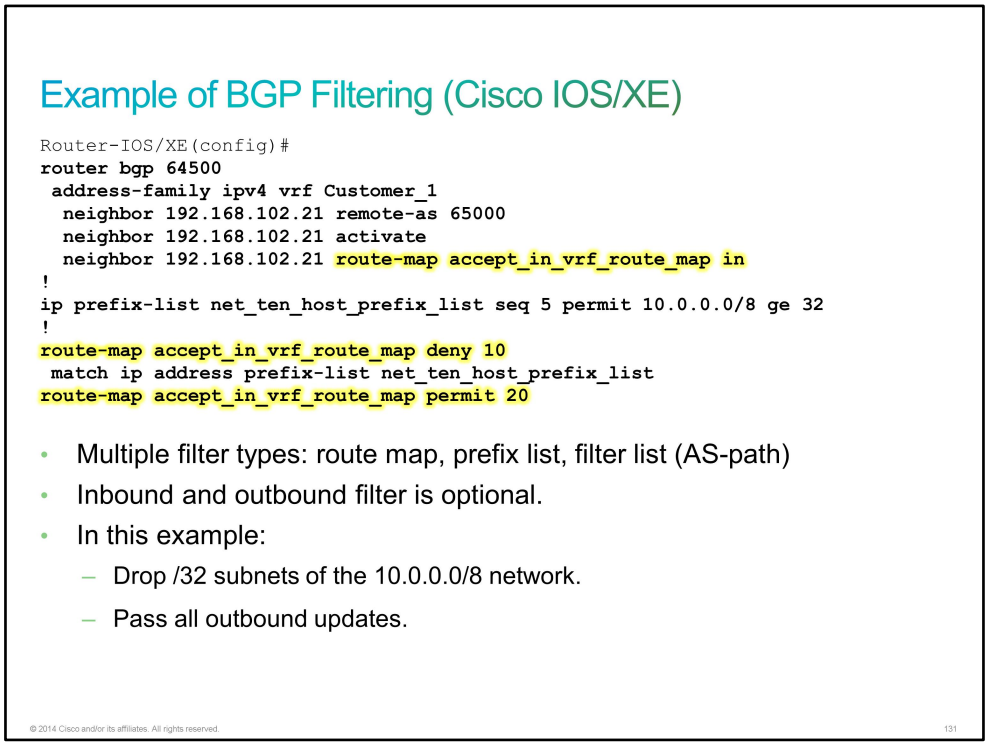


This example shows how to apply a route policy (accept_in_vrf) that references a prefix-set (net_ten_host). In this scenario, all /32 subnets of the 10.0.0.0/8 will be dropped and all others will be accepted into the VRF. The outbound policy allows all network updates to be forwarded to the external BGP peer.

https://t.me/learningnets

Example of BGP Filtering (Cisco IOS/XE)

You may filter incoming or outgoing updates on a PE router running Cisco IOS or Cisco IOS XE software. There are multiple tools that you can use, such as route maps, prefix lists, and filter lists. All of them are optional.



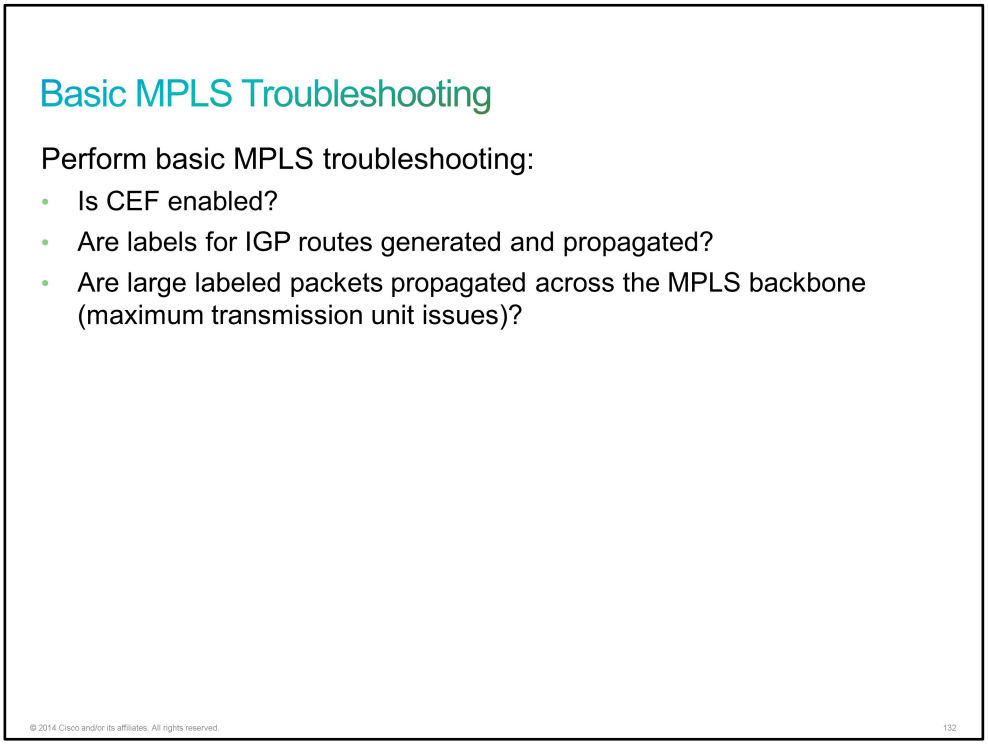
This example shows how to apply a route map (accept_in_vrf_route-map) that references a prefix list (net_ten_host_prefix-list) in coming from the EBGP neighbor. In this scenario, all /32 subnets of the 10.0.0.0/8 will be dropped and all others will be accepted into the VRF. There is no outbound policy and thus all network updates will be forwarded to the external BGP peer.

https://t.me/learningnets

Troubleshooting MPLS VPNs

Basic MPLS Troubleshooting

In basic MPLS troubleshooting, you ensure that LDP and MP-BGP exchange labels appropriately and that traffic reaches its destination along the LSP.



Before you start in-depth MPLS VPN troubleshooting, you should ask the following standard MPLS troubleshooting questions:

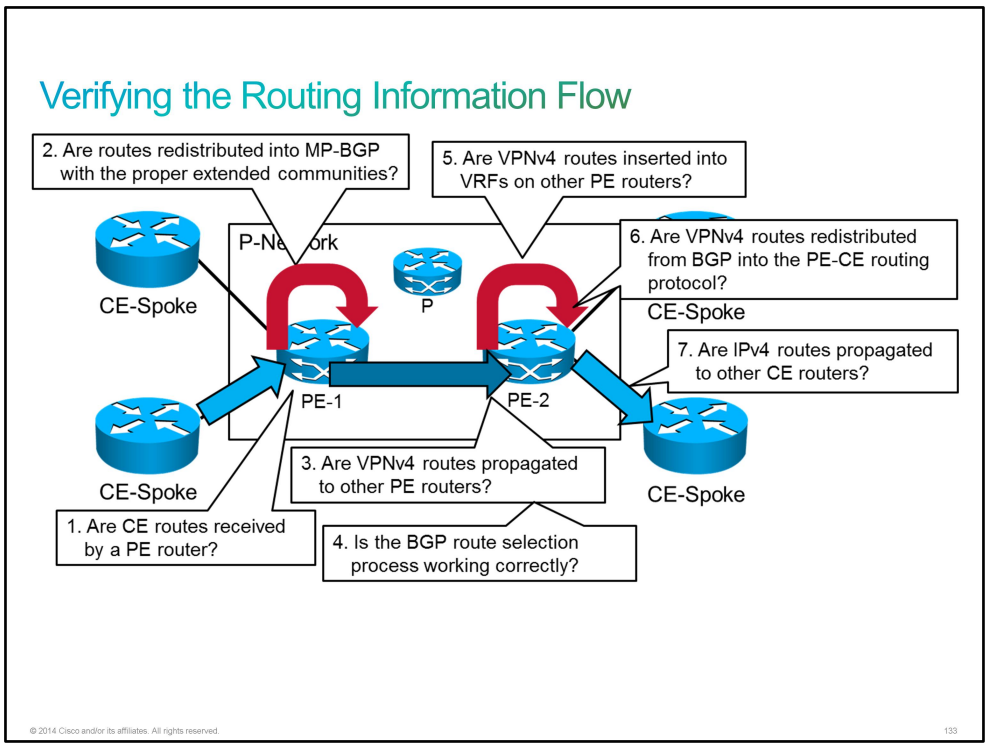
- Is Cisco Express Forwarding enabled on all routers in the transit path between the PE routers?
- Are labels for BGP next hops generated and propagated?
- Are there any MTU issues in the transit path (for example, LAN switches not supporting a jumbo Ethernet frame)?

MPLS VPN troubleshooting consists of these two major steps:

- Verifying the routing information flow, using the checks that are outlined in the figure
- Verifying the data flow, or packet forwarding

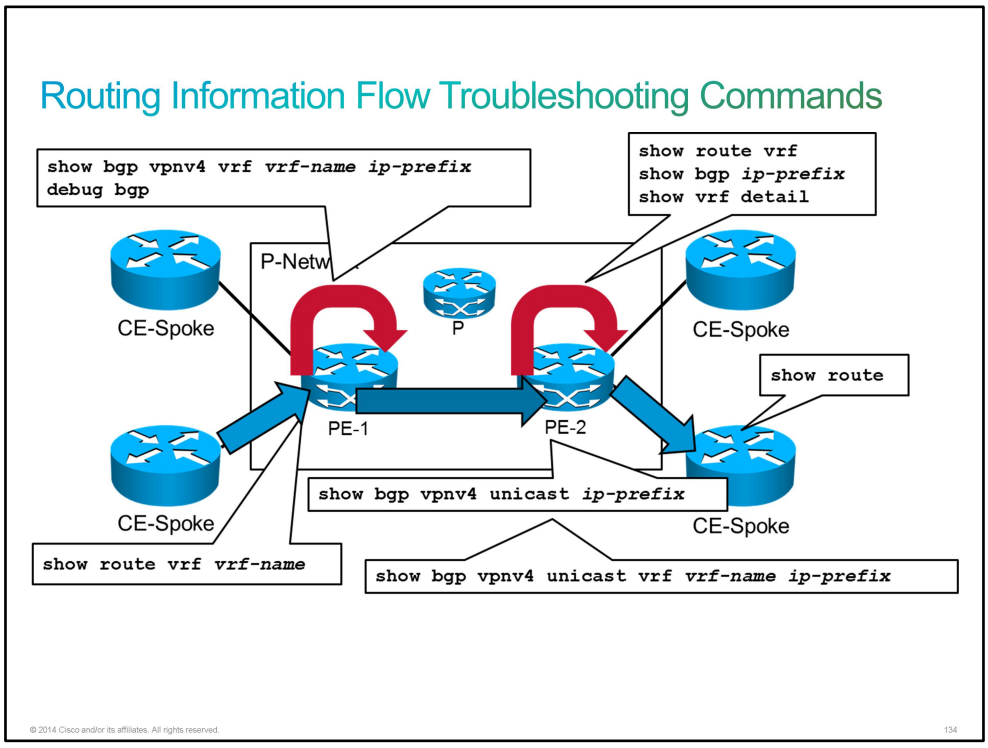
https://t.me/learningnets

Verifying the Routing Information Flow



Verification of the routing information flow should be done systematically, starting at the ingress CE router and moving to the egress CE router.

Routing Information Flow Troubleshooting Commands



https://t.me/learningnets

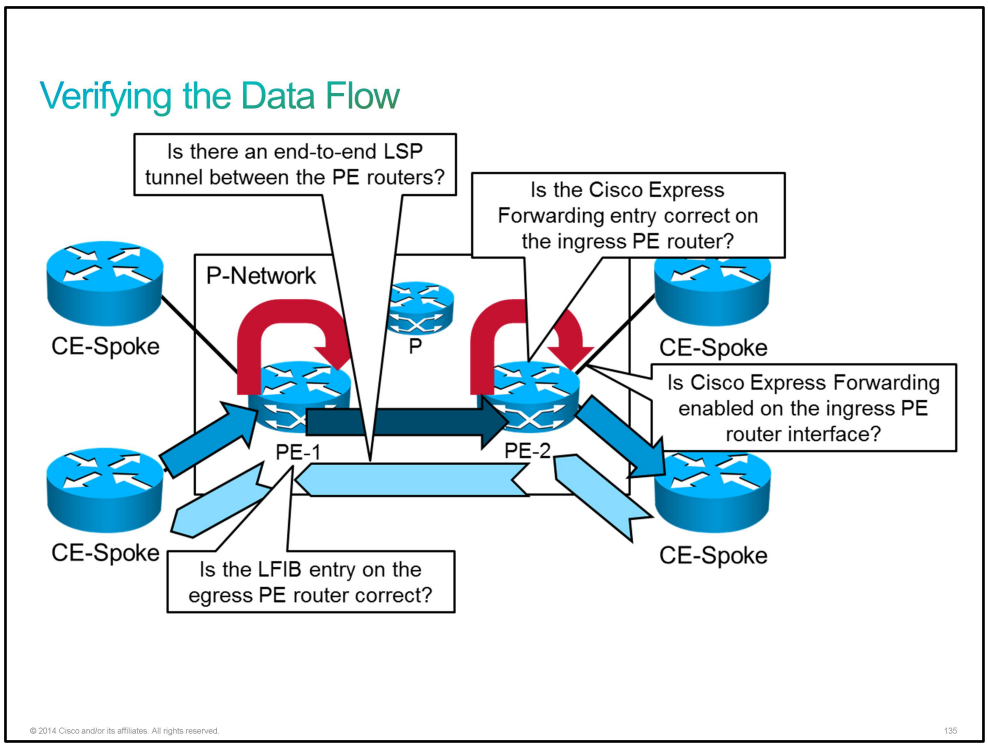
Troubleshooting routing information flow requires the verification of end-to-end routing information propagation between CE routers.

1. The first step is to check the routing information exchange from the CE routers to the PE routers. Use the **show ip route vrf** *vrf-name* (Cisco IOS and IOS XE) or **show route vrf** *vrf-name* (Cisco IOS XR) command to verify that the PE router receives customer routes from the CE router. Use traditional routing protocol troubleshooting if needed.
2. The CE routes that are received by the PE router need to be redistributed into MP-BGP; otherwise, the routes will not be propagated to other PE routers. Proper redistribution of CE routes into a per-VRF instance of BGP can be verified with the **show ip bgp vpnv4 vrf** *vrf-name* (IOS and IOS XE) or **show bgp vpnv4 vrf** *vrf-name* (IOS XR) command. The RD prepended to the IPv4 prefix and the RTs attached to the CE route can be verified with the **show ip bgp vpnv4 vrf** *vrf-name ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 vrf** *vrf-name ip-prefix* (IOS XR) command.
3. The CE routes that are redistributed into MP-BGP need to be propagated to other PE routers. Verify proper route propagation with the **show ip bgp vpnv4 all** *ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 unicast** *ip-prefix* (IOS XR) command on the remote PE router.
4. The VPNv4 routes that are received by the PE router need to be inserted into the proper VRF. This insertion can be verified with the **show ip route vrf** (IOS and IOS XE) or **show route vrf** (IOS XR) command. The validity of the import RTs can be verified with the **show ip bgp vpnv4 all** *ip-prefix* (IOS and IOS XE) or **show bgp vpnv4 unicast** *ip-prefix* (IOS XR) command that displays the RTs attached to a VPNv4 route. You can also verify the validity of the import RTs with the **show ip vrf detail** (IOS and IOS XE) or **show vrf detail** (IOS XR) command that lists the import RTs for a VRF. At least one RT attached to the VPNv4 route needs to match at least one RT in the VRF.
5. Next, the BGP routes that are received via MP-BGP and inserted into the VRF need to be redistributed into the PE-CE routing protocol.
6. Finally, the routes that are redistributed into the PE-CE routing protocol need to be propagated to CE routers.

 © 2014 Cisco Systems, Inc.

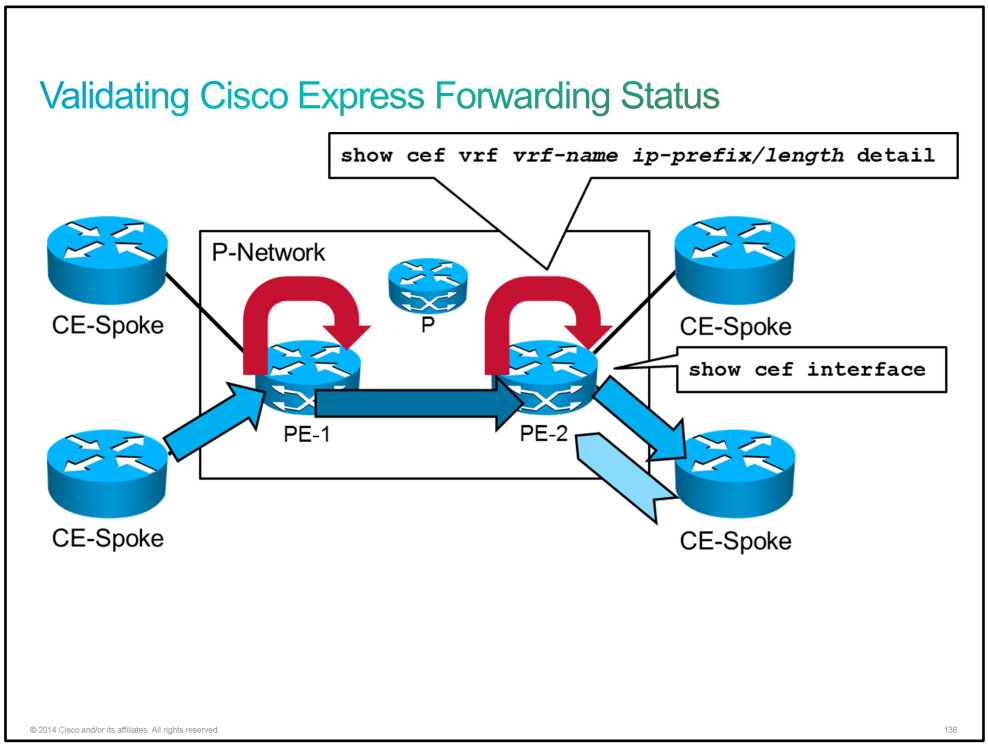
https://t.me/learningnets

Verifying the Data Flow



After you have verified proper route exchange, start MPLS VPN data flow troubleshooting using the checks that are listed in the next figures.

Validating Cisco Express Forwarding Status

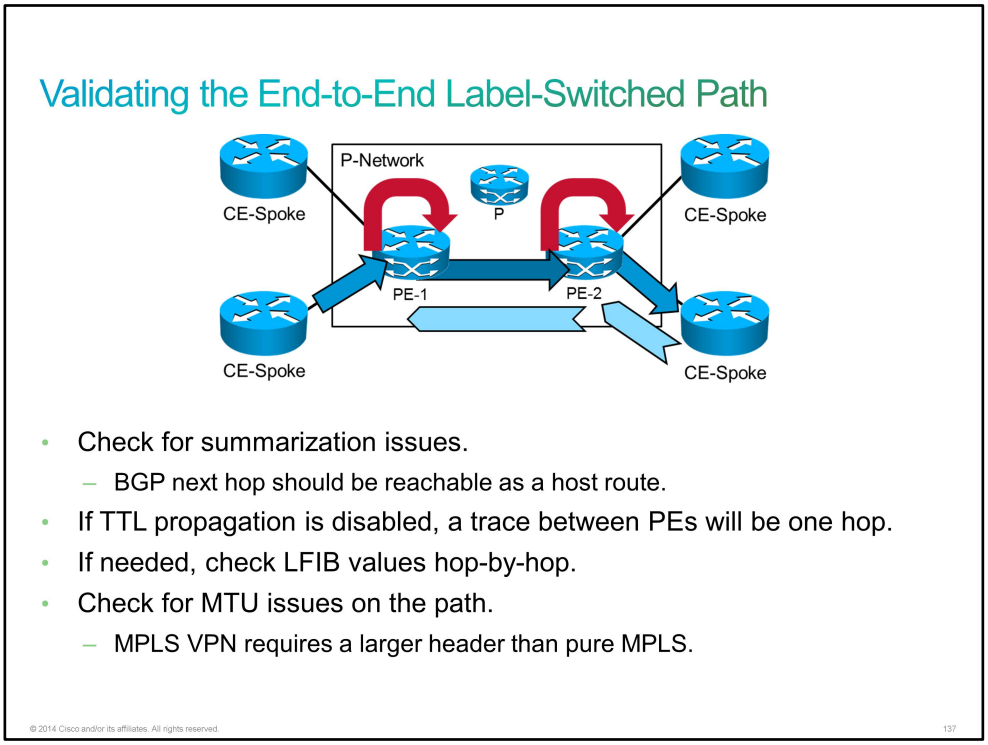


https://t.me/learningnets

One of the most common configuration mistakes related to data flow is the failure to enable Cisco Express Forwarding on the ingress PE router interface. The presence of Cisco Express Forwarding can be verified with the **show cef interface** command on Cisco IOS, IOS XE, and IOS XR devices.

If Cisco Express Forwarding switching is enabled on the ingress interface, you can verify the validity of the Cisco Express Forwarding entry and the associated label stack with the **show ip cef vrf** *vrf-name ip-prefix* **detail** (IOS and IOS XE) or **show cef vrf** *vrf-name ip-prefix* **detail** (IOS XR) command. The top label in the stack should correspond to the BGP next-hop label as displayed by the **show mpls forwarding-table** (IOS and IOS XE) or **show mpls forwarding** (IOS XR) command on the ingress router. The second label in the stack should correspond to the label allocated by the egress router. You can verify this label by using the **show mpls forwarding-table** (IOS and IOS XE) or **show mpls forwarding vrf** (IOS XR) command on the egress router.

Validating the End-to-End Label-Switched Path

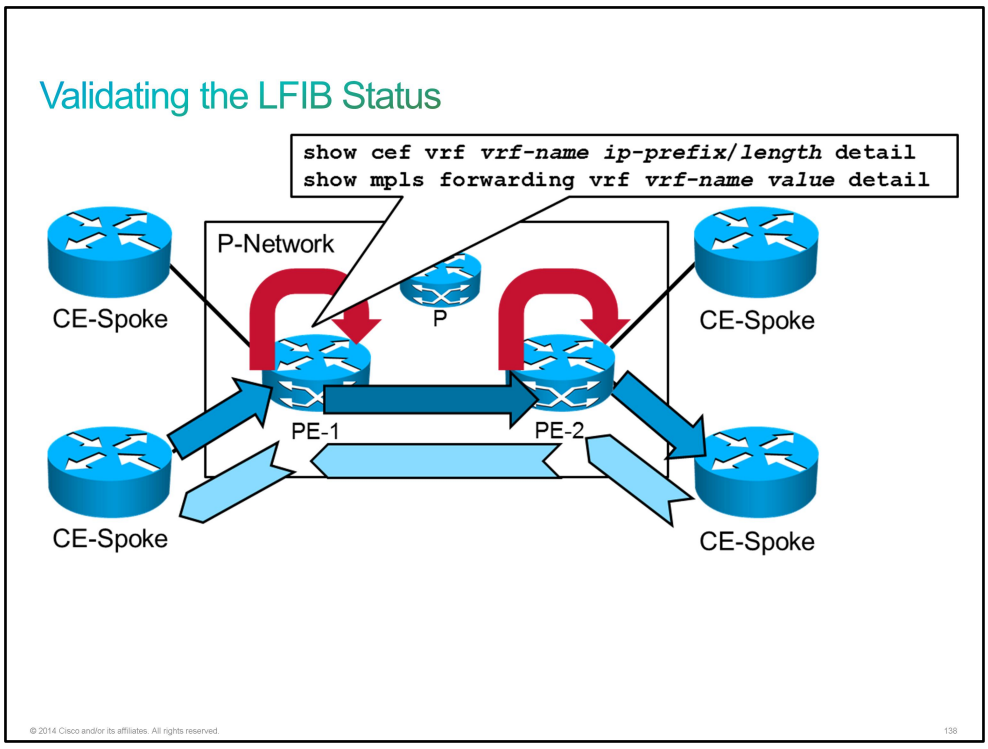


If Cisco Express Forwarding is enabled on the ingress interface and the Cisco Express Forwarding entry contains the proper labels, the data flow problem might lie inside the MPLS core. Two common mistakes include summarization of BGP next hops inside the core IGP and MTU issues.

The quickest way to diagnose summarization problems is to disable IP TTL propagation into the MPLS label header using the **no mpls ip ttl-propagate** (IOS and IOS XE) or **mpls ip ttl-propagate disable** (IOS XR) configuration command on the P router and PE routers. The **traceroute** command from the ingress PE router toward the BGP next hop should display no intermediate hops when TTL propagation is disabled. If intermediate hops are displayed, the LSP tunnel between PE routers is broken at those hops, and the VPN traffic cannot flow.

https://t.me/learningnets

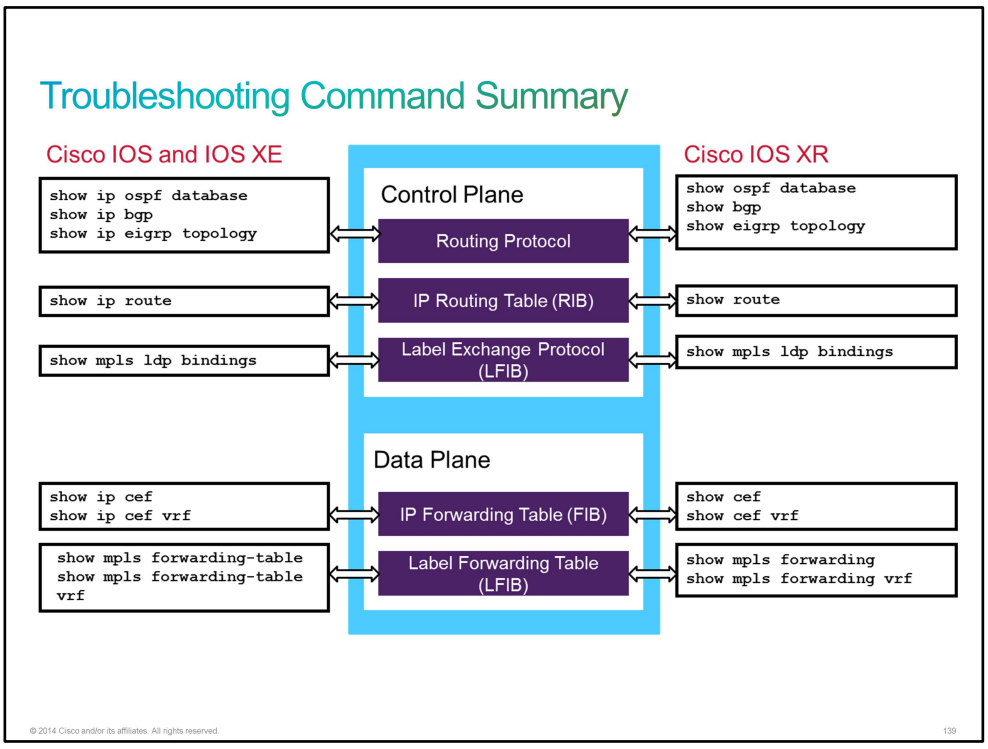
Validating the LFIB Status



As a last troubleshooting measure (usually not needed), you can verify the contents of the LFIB on the egress PE router and compare them with the second label in the label stack on the ingress PE router. A mismatch indicates an internal Cisco IOS Software error that you will need to report to the Cisco Technical Assistance Center.

https://t.me/learningnets

Troubleshooting Command Summary

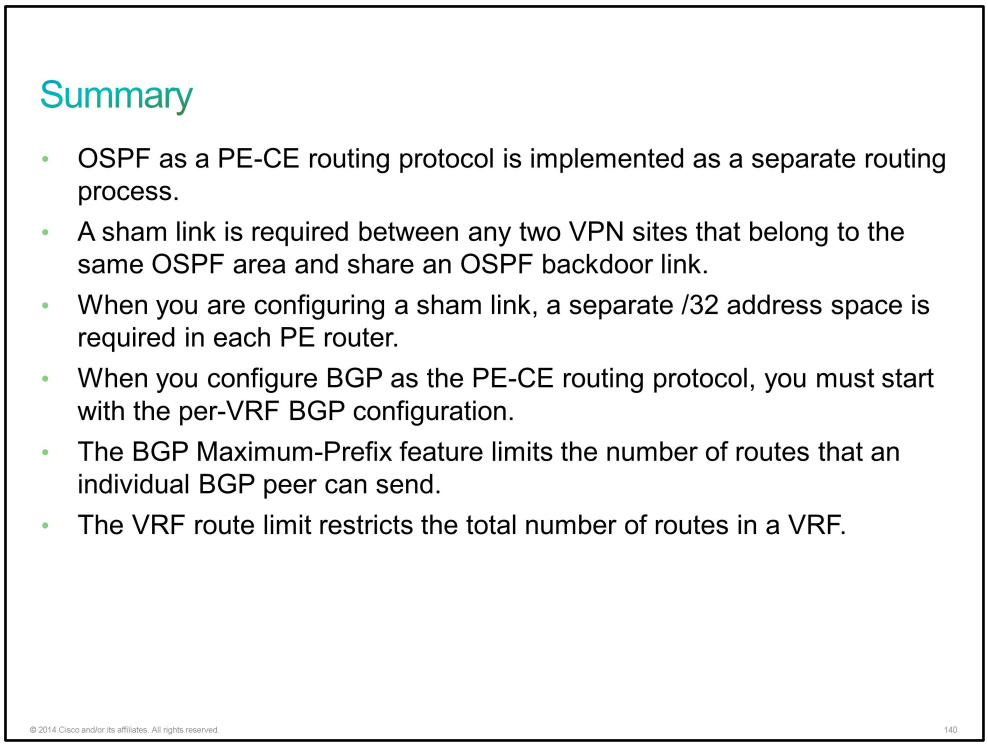


The figure shows the relevant commands for Cisco IOS, IOS XE, and Cisco IOS XR devices to troubleshoot the control plane and the data plane of an MPLS router.

https://t.me/learningnets

Summary

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

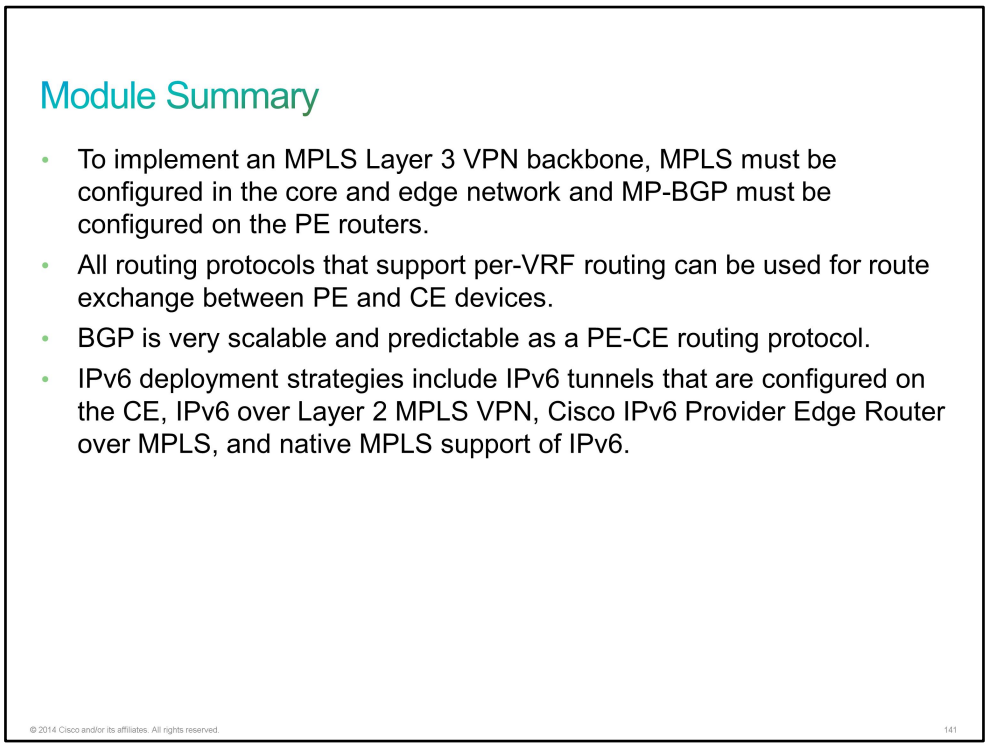


https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: Module Summary



An MPLS VPN implementation involves VRF tables, the interaction between CE-to-P routing protocols, and MP-BGP in the service provider backbone.

References

For additional information, refer to these resources:

- Cisco Systems, Inc. “Configuring MPLS Layer 3 VPNs.” .http://www.cisco.com/en/US/docs/ios/mpls/configuration/guide/mp_cfg_layer3_vpn.html
- Cisco Systems, Inc. “Implementing MPLS Layer 3 VPNs on Cisco IOS XR Software.” http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/mpls/configuration/guide/gc37v3.html.

https://t.me/learningnets

- Cisco Systems, Inc. The “Identifying IPv6 Service Provider Deployment” module in the *IPv6 Fundamentals, Design, and Deployment* (IP6FD) course.
- Cisco Systems, Inc. The “MPLS VPN Implementation” module in the *Implementing Cisco MPLS* (MPLS) v2.3 course.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 5: 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. In an MPLS VPN implementation, what is a VRF? (Source: Implementing MPLS Layer 3 VPN Backbones)

   A. the routing and forwarding instance for all sites belonging to a single customer
   B. the routing and forwarding instance for all sites belonging to a single customer location
   C. the routing and forwarding instance for all sites using a common routing protocol
   D. the routing and forwarding instance for a set of sites with identical connectivity requirements

2. Which two protocols are VPN-aware? (Choose two.) (Source: Implementing MPLS Layer 3 VPN Backbones)

   A. RIPv2
   B. IS-IS
   C. ODR
   D. EIGRP
   E. RIPv1

3. Which command would you use to create a VRF named "VPNA" on Cisco IOS XR devices? (Source: Implementing MPLS Layer 3 VPN Backbones)

   A. **ip vrf VPNA**
   B. **vrf definition VPNA**
   C. **vrf VPNA**
   D. **ip vrf forwarding VPNA**
   E. **vrf forwarding VPNA**

https://t.me/learningnets

4. Which command would you use to associate interface Gigabit Ethernet 0/0/0/0 with a VRF named "VPNA" on Cisco IOS XR devices? (Source: Implementing MPLS Layer 3 VPN Backbones)

   A. **vrf forwarding VPNA**
   B. **ip vrf VPNA interface GigabitEthernet0/0/0/0**
   C. **ip vrf forwarding VPNA**
   D. **vrf VPNA**
   E. **vrf VPNA interface GigabitEthernet0/0/0/0**

5. What happens to the existing IP address of an interface when you associate the interface with a VRF in Cisco IOS XE? (Source: Implementing MPLS Layer 3 VPN Backbones)

   A. It remains unchanged.
   B. It is removed from the interface.
   C. It is changed to the Loopback 0 address.
   D. It is moved under the VRF configuration.

6. Which command configures the redistribution of static VRF routes between PE routers that are running Cisco IOS XR Software? (Source: Connecting Customers Using Simple Routing Protocols)

   A. RP/0/RP0/CPU0:router(config)# **redistribute static**
   B. RP/0/RP0/CPU0:router(config-router)# **redistribute static**
   C. RP/0/RP0/CPU0:router(config-bgp)# **redistribute static**
   D. RP/0/RP0/CPU0:router(config-bgp-vrf)# **redistribute static**
   E. RP/0/RP0/CPU0:router(config-bgp-vrf-af)# **redistribute static**

7. How are routing contexts implemented in OSPFv2? (Source: Connecting Customers Using BGP or OSPF)

   A. by redistributing into MP-BGP
   B. as separate routing processes
   C. by assigning them to interfaces
   D. as separate isolated instances of the same routing protocol

8. In the SOO extended community attribute in MP-BGP, what does "SOO" stand for ? (Source: Connecting Customers Using BGP or OSPF)

   A. Site of Origin
   B. Side of Origin
   C. State of Origin
   D. Source of Origin

9. Which command do you use to display the contents of the LFIB table on a PE router running Cisco IOS XR Software? (Source: Connecting Customers Using BGP or OSPF)

   A. **show mpls forwarding-table**
   B. **show mpls forwarding-table vrf**
   C. **show mpls forwarding**
   D. **show mpls forwarding vrf**

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. D
2. A, D
3. C
4. D
5. B
6. E
7. B
8. A
9. C

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 3: Special Connectivity in MPLS Layer 3 VPNs

Introduction

This module discusses some advanced configuration features that can help increase the stability of the MPLS VPN backbone. The first lesson discusses various MPLS VPN features that a service provider might use to help meet service requirements, and looks at various types of VPN solutions and topologies.

Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. The second lesson provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Upon completing this module, you will be able to:

- Describe the requirements, usage, and solutions associated with overlapping, central services, and managed CE router service VPNs
- Describe common customer Internet connectivity scenarios and identify design models for combining Internet access with MPLS Layer 3 VPN services

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Implementing Special Connectivity in MPLS VPNs

Overview

This lesson discusses some advanced configuration features that can help increase the stability of the MPLS VPN backbone. The lesson also discusses various MPLS VPN features that a service provider might use to help meet service requirements, and looks at various types of VPN solutions and topologies.

Upon completing this lesson, you will be able to:

- Describe the Cisco IP NGN infrastructure layer
- Describe an overview of the overlapping VPNs
- Describe the use cases of overlapping VPNs
- Describe routing in overlapping VPNs
- Describe the packet flow in overlapping VPNs
- Describe the overlapping VPNs configuration procedure
- Describe the VRF configuration in overlapping VPNs
- Explain the BGP configuration in overlapping VPNs
- Describe central services VPNs
- Describe routing in central services VPNs
- Describe the packet flow in central services VPNs
- Describe the central services VPN configuration procedure
- Describe a central services VPN configuration example

https://t.me/learningnets

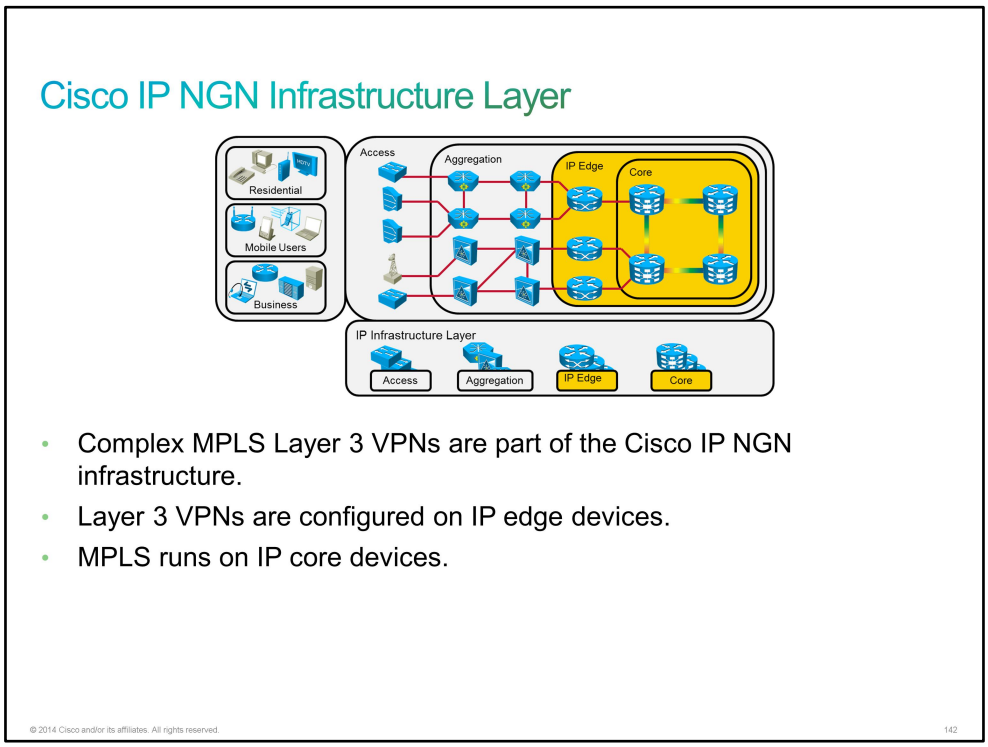
- Describe selective import and export
- Describe the selective import feature
- Describe the selective export feature
- Describe the managed CE routers service
- Describe the requirements when implementing the managed CE routers service
- Describe the managed CE routers service configuration

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Cisco IP NGN Infrastructure Layer

The Cisco IP Next-Generation Network infrastructure layer is typically classified as the area containing the core and edge.



Complex MPLS Layer 3 VPN solutions belong the Cisco IP NGN infrastructure layer. As in standard MPLS VPNs, the VRFs are configured on the edge devices. The logic of the complex VPNs is implemented with a careful selection of the route targets.

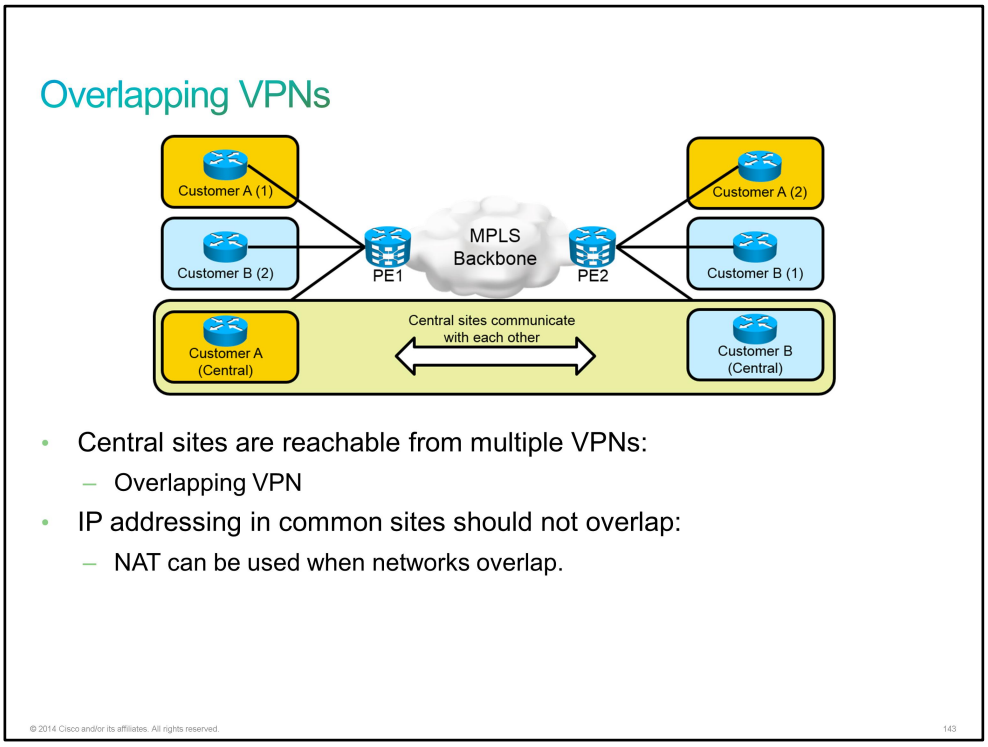
Complex MPLS Layer 3 VPNs include the following traits:

- Selective import or export of route targets or other attributes (situational).
- Layer 3 VPNs are usually configured on IP edge devices.
- MPLS runs on IP core devices.

https://t.me/learningnets

Overlapping VPNs

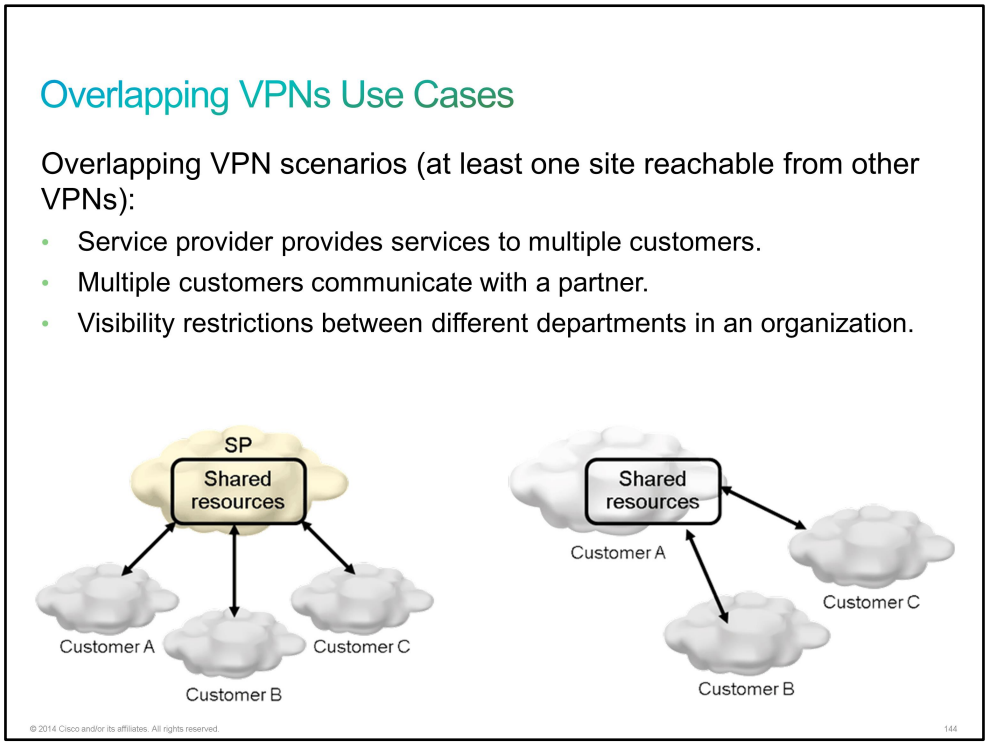
When two VPN customers want to share information, they might decide to interconnect their central sites. To achieve this interconnection, two simple VPNs are created, each containing a customer central site and its remote sites. Then a third VPN, which partially overlaps with the customer VPNs but connects only their central sites, is created. The central sites can talk to each other. Each central site can also talk to the remote sites in its simple VPN, but not to the remote sites belonging to the simple VPN of the other customer. The addresses used in the central sites, however, must be unique in both VPNs.



Another option is to use dual NAT with a registered address to be imported and exported between the two central sites.

https://t.me/learningnets

Overlapping VPNs Use Cases



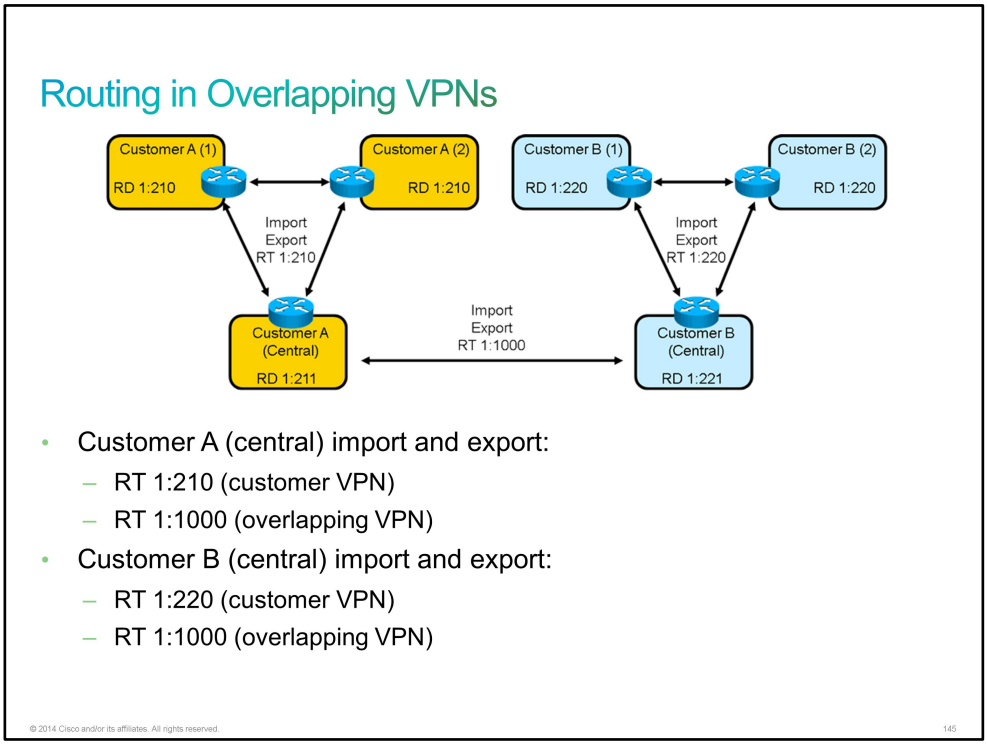
Typical uses for overlapping VPNs include the following:

- A service provider may provide services to its customers.
- Companies that use MPLS VPNs to implement both intranet and extranet services might use overlapping VPNs. In this scenario, each company participating in the extranet VPN would probably deploy a security mechanism on its CE routers to prevent other companies that are participating in the VPN from gaining access to other sites in its VPN.
- A security-conscious company might decide to limit visibility between different departments in the organization. Overlapping VPNs might be used as a solution.

https://t.me/learningnets

Routing in Overlapping VPNs

The figure shows how to implement overlapping VPNs between two customer central sites.

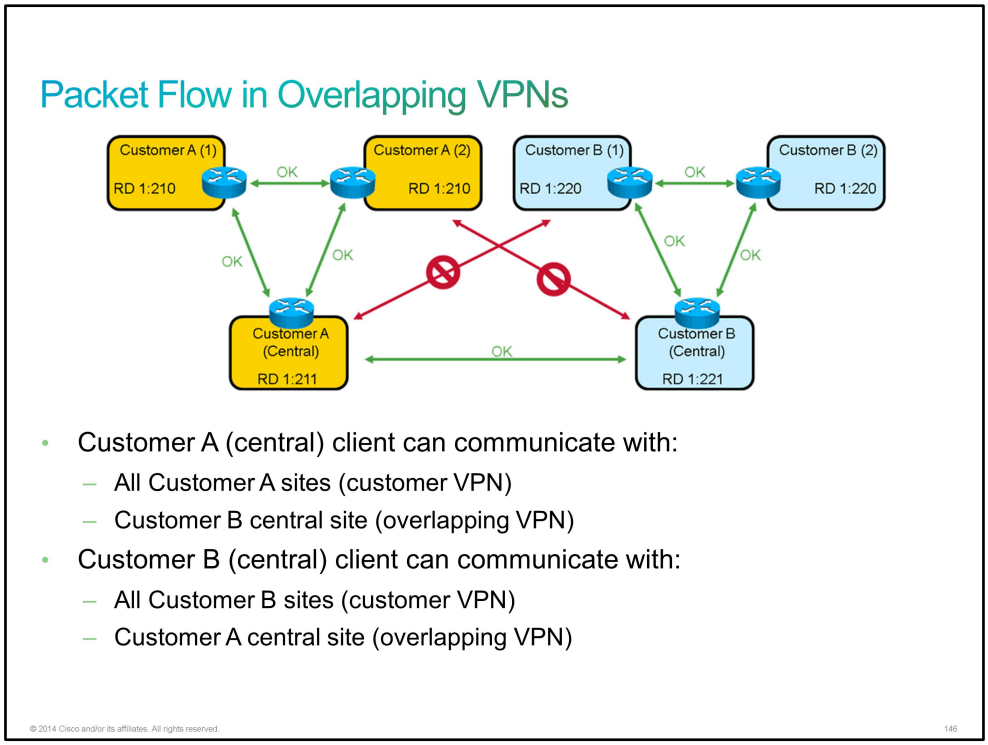


Only the configuration on the customer central-site PE router has to be changed:

- New virtual VRF is created for the customer central site.
- A new RD is configured for customer central site.
- For the Customer A central site, the following occurs:
  - Import and export routes with RT 1:210 (customer routes)
  - Import and export routes with RT 1:1000 (overlapping routes)
- For the Customer B central site, the following occurs:
  - Import and export routes with RT 1:220 (customer routes)
  - Import and export routes with RT 1:1000 (overlapping routes)

https://t.me/learningnets

Packet Flow in Overlapping VPNs



Because sites belonging to different VPNs do not share routing information, they cannot talk to each other. The list shows overlapping VPN data flow:

- The simple VPN for Customer A contains routes that originate from the following:
  - A-Central site
  - A remote sites
- The simple VPN for Customer B contains routes that originate from the following:
  - B-Central site
  - B remote sites
- The overlapping VPN contains routes that originate from the following:
  - A-Central site
  - B-Central site
- All Customer A sites can communicate with each other.
- All Customer B sites can communicate with each other.
- The A-Central site and the B-Central site can communicate with each other.
- The customer A remote site cannot communicate with the customer B remote sites.
- The customer A central site cannot communicate with the customer B remote sites.

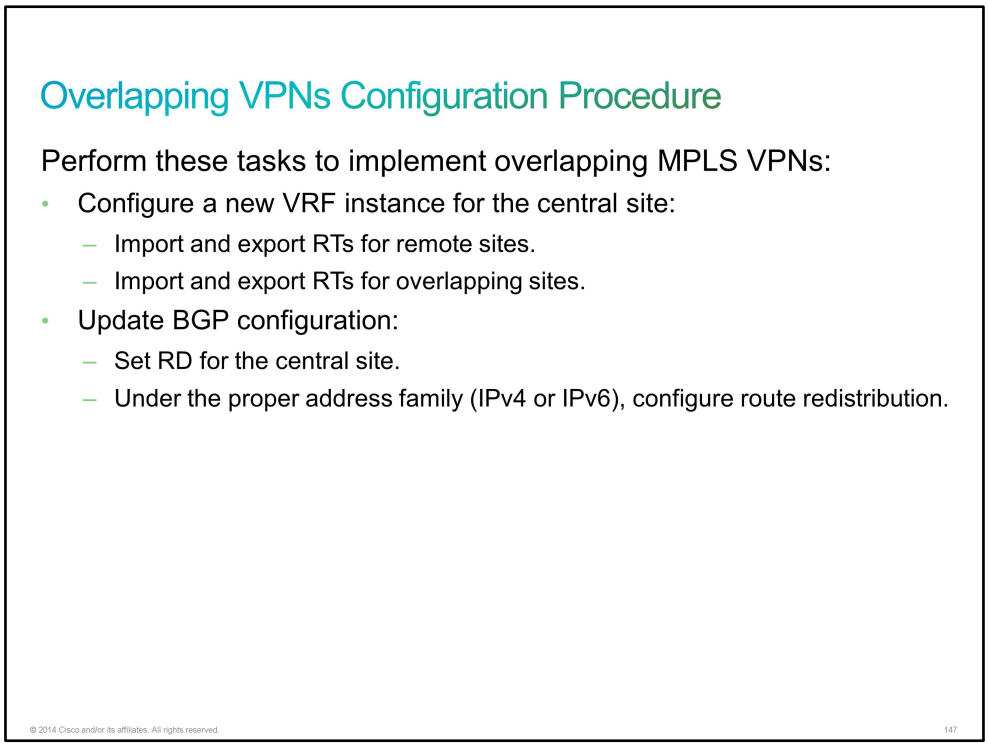
https://t.me/learningnets

Note	If a site participating in more than one VPN is propagating a default route to other sites, it can attract traffic from those sites and start acting as a transit site between VPNs, enabling sites that were not supposed to communicate to establish two-way communication.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Overlapping VPNs Configuration Procedure

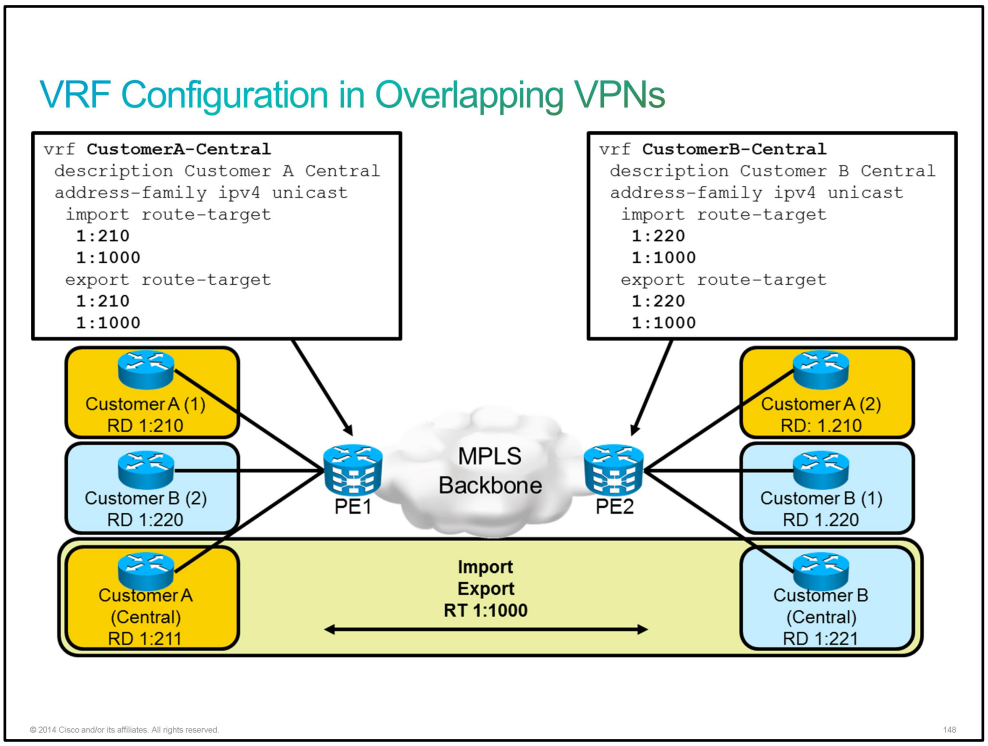


To implement overlapping MPLS VPNs, you will configure the central site VRFs with appropriate RTs to enable route exchange with the central site VRFs in other VPNs and with branch VRFs of their own VPN. The RDs in the central-site VRFs should be set to unique values.

https://t.me/learningnets

VRF Configuration in Overlapping VPNs

The figure shows a Cisco IOS XR sample configuration for configuring VRF instances.



Each PE router contains a CustomerA and a CustomerB VRF. PE1 has the CustomerA-Central VRF. PE2 has the CustomerB-Central VRF.

The relevant PE1 configuration is as follows:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1(config)#
vrf CustomerA-Central
 description Customer A Cent
 address-family ipv4 unicast
  import route-target
   1:210
   1:1000
  export route-target
   1:210
   1:1000
!
vrf CustomerA
 description Customer A
 address-family ipv4 unicast
  import route-target
   1:210
  export route-target
   1:210
!
vrf CustomerB
 description Customer B
 address-family ipv4 unicast
  import route-target
   1:220
  export route-target
   1:220
```

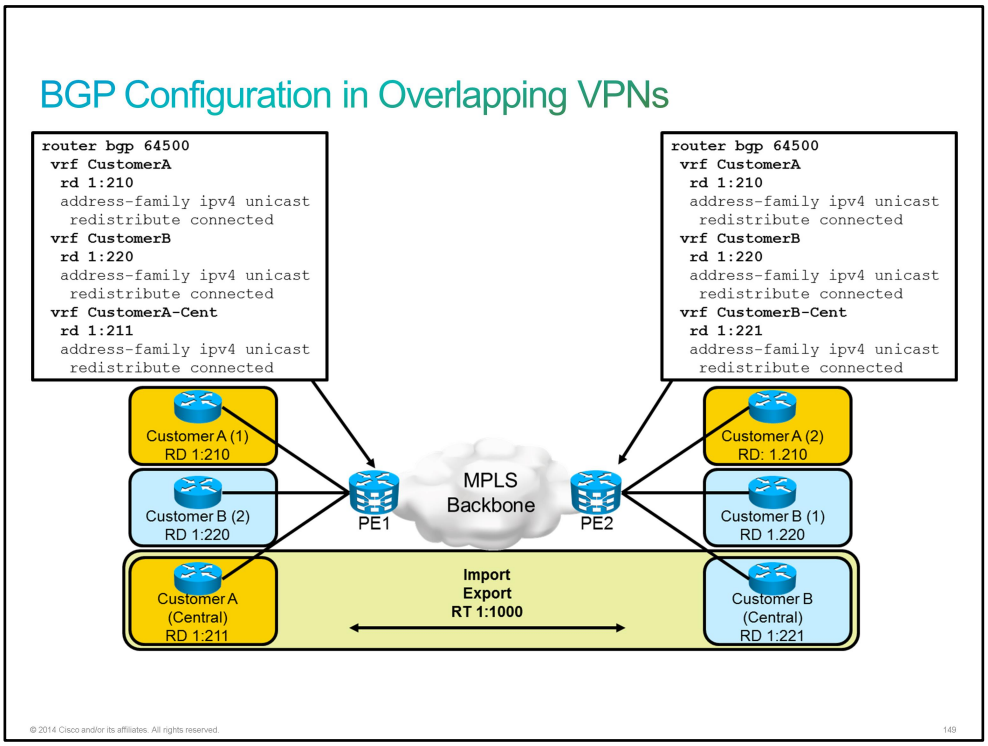
The relevant PE2 configuration is shown here:

```
RP/0/RSP0/CPU0:PE2(config)#
vrf CustomerB-Central
 description Customer B Cent
 address-family ipv4 unicast
  import route-target
   1:220
   1:1000
  export route-target
   1:220
   1:1000
!
vrf CustomerA
 description Customer A
 address-family ipv4 unicast
  import route-target
   1:210
  export route-target
   1:210
!
vrf CustomerB
 description Customer B
 address-family ipv4 unicast
  import route-target
   1:220
  export route-target
   1:220
```


https://t.me/learningnets

BGP Configuration in Overlapping VPNs

The figure shows a Cisco IOS XR sample configuration for configuring the BGP process.



The relevant PE1 configuration is as follows:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1(config)#
!
interface Loopback 0
ipv4 address 10.0.1.1/32
!
router bgp 64500
 address-family vpnv4 unicast
 !
 neighbor 10.0.1.2
  remote-as 64500
  update-source Loopback0
  address-family vpnv4 unicast
   next-hop-self
!
 vrf CustomerA
  rd 1:210
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerB
  rd 1:220
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerA-Central
  rd 1:211
  address-family ipv4 unicast
   redistribute connected
```

The relevant configuration of PE2 is shown here:

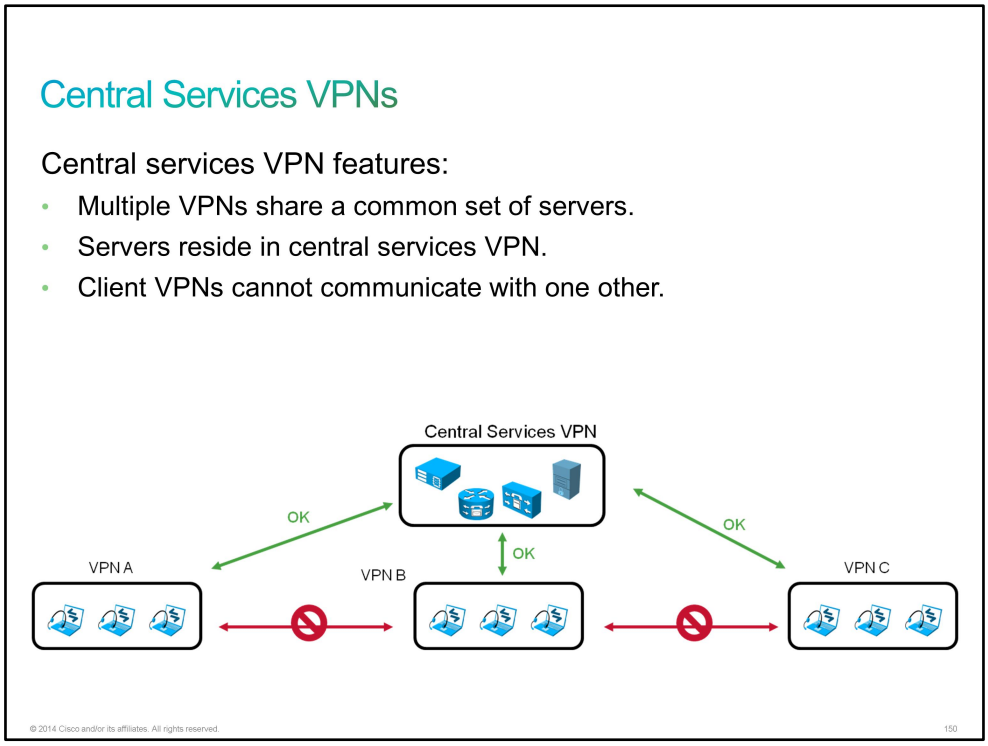
https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE2(config)#
!
interface Loopback 0
ipv4 address 10.0.1.2/32
!
router bgp 64500
 address-family vpnv4 unicast
 !
 neighbor 10.0.1.1
  remote-as 64500
  update-source Loopback0
  address-family vpnv4 unicast
   next-hop-self
!
 vrf CustomerA
  rd 1:210
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerB
  rd 1:220
  address-family ipv4 unicast
   redistribute connected
!
 vrf CustomerB-Central
  rd 1:221
  address-family ipv4 unicast
   redistribute connected
```

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Central Services VPNs



A central services VPN is a topology with these characteristics:

- Some sites (“server sites”) can communicate with all other sites.
- All the other sites (“client sites”) can communicate only with the server sites.

This topology can be used in these situations:

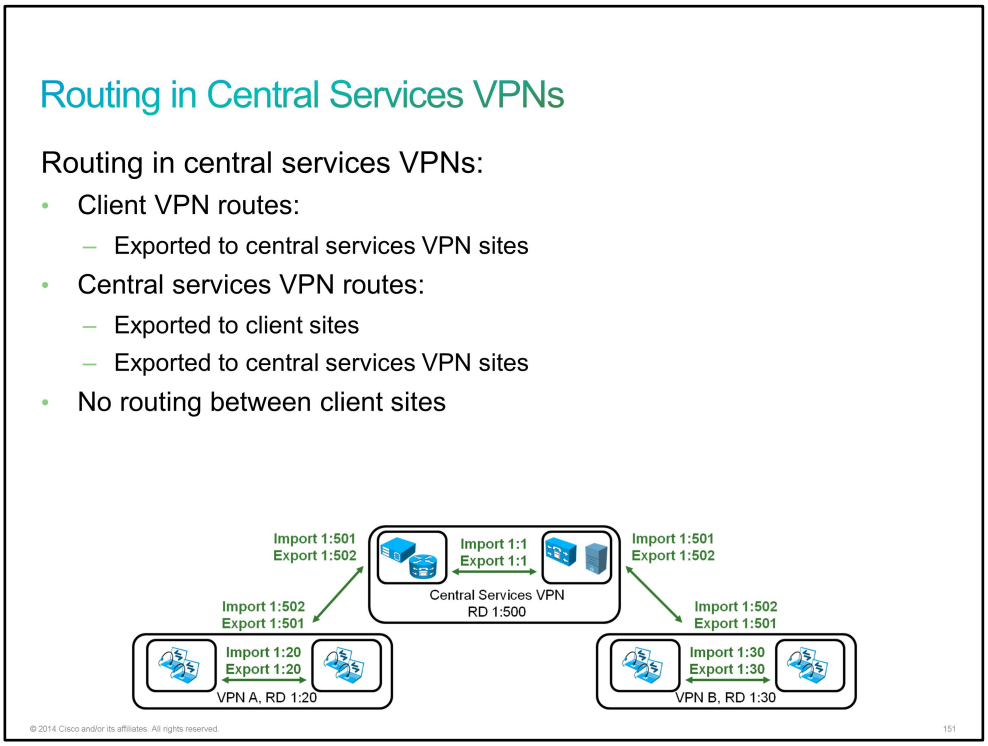
- The service provider offers services to all customers by allowing them access to a common VPN.
- Two (or more) companies want to exchange information by sharing a common set of servers.

A security-conscious company may also separate its departments and allow them access only to common servers.

https://t.me/learningnets

Routing in Central Services VPNs

The figure shows a central services VPN (RD 1:500) with a set of common services. On client-site VPNs, routes are exported with RT 1:501 and imported with RT 1:502.



The figure illustrates the MPLS VPN routing model that is used to implement a central services VPN:

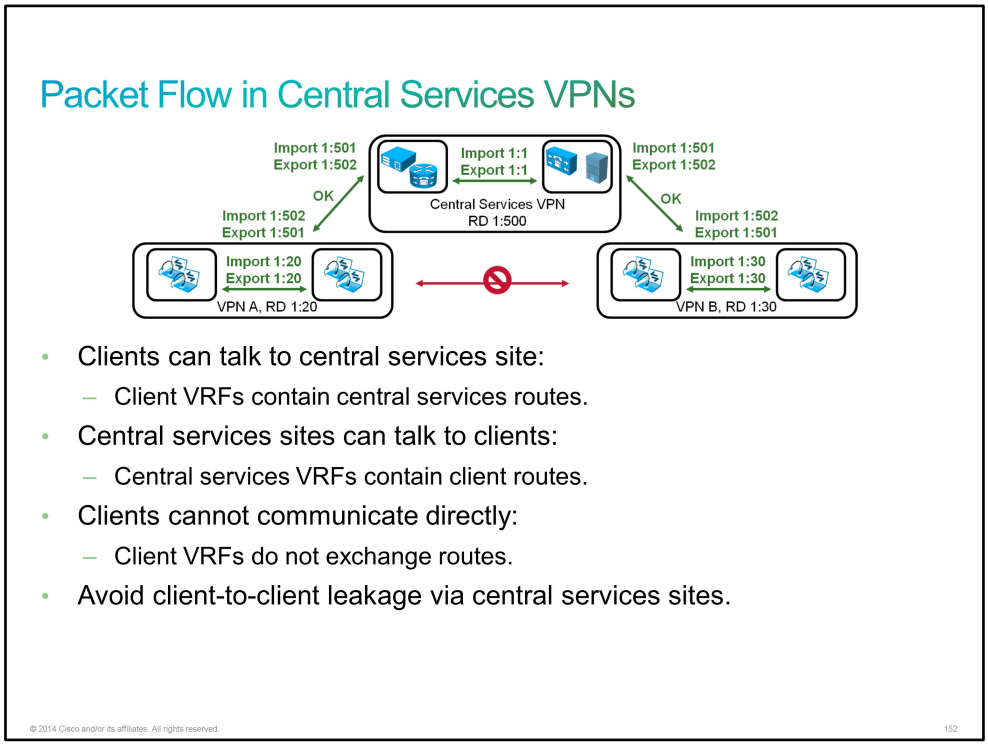
- Client-site VPN routes are exported to the central services VPN.
- Central services VPN routes are exported to client-site VPNs.
- Central services VPN routes are also exchanged with other central services VPN sites.
- There should be no route exchange and connectivity between client sites.

On the central services VPN, routes are exchanged in mirrored fashion: exported with RT 1:502 and imported with RT 1:501.

https://t.me/learningnets

Packet Flow in Central Services VPNs

In the central services VPN topology, the client VRF contains only routes from the client site and routes from the central services sites. This setup precludes the client sites from communicating with other client sites.

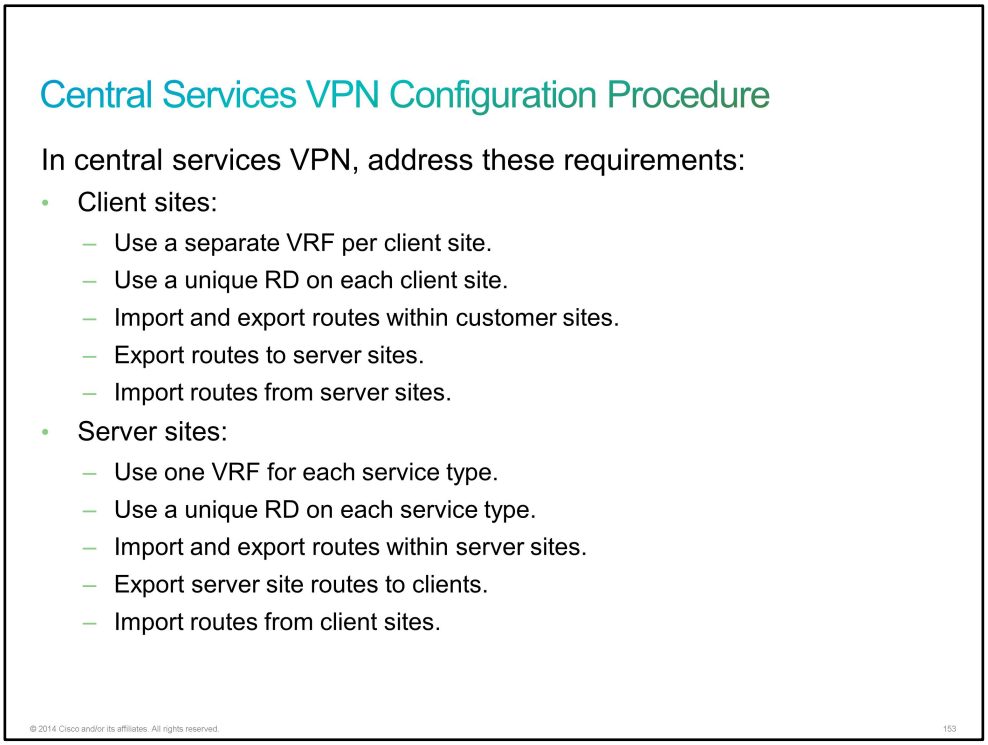


A central services VRF in this topology contains routes from the site or sites attached to the VRF and also routes from all other client and central services sites. Hosts in central services sites can therefore communicate with hosts in all other sites.

Note	If the central site is propagating a default route to other sites, it can result in client sites seeing each other through the CE router in the central site.

https://t.me/learningnets

Central Services VPN Configuration Procedure

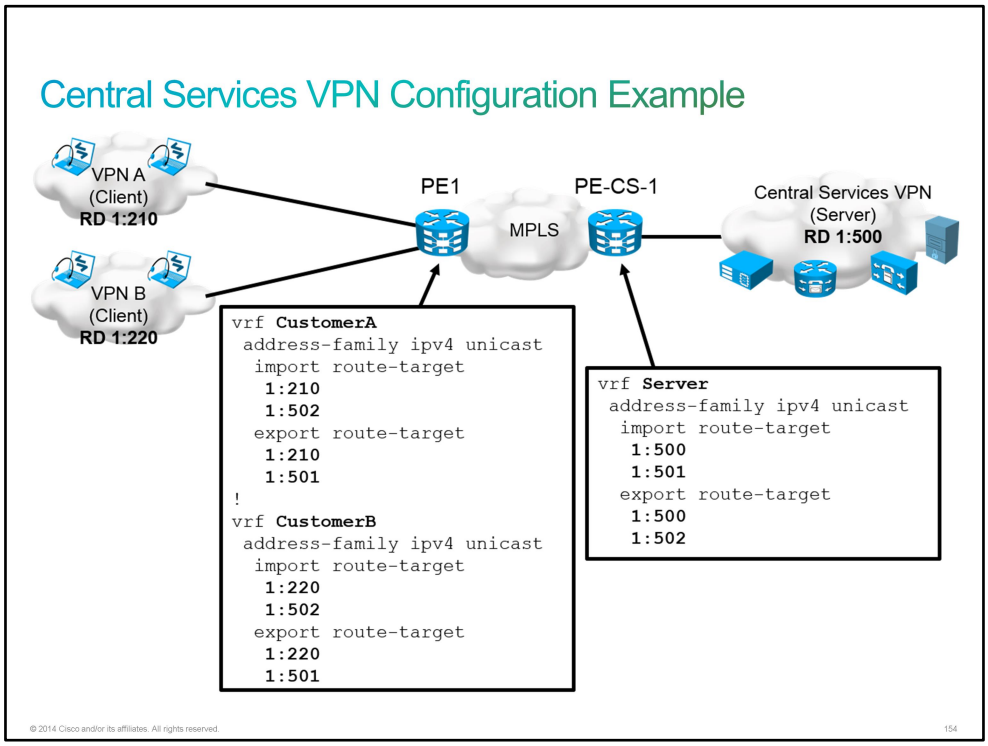


This configuration procedure lists the steps needed to set up a central services VPN. Apart from standard VRF configuration, the route targets in the client and server sites need to be configured in a mirrored fashion.

https://t.me/learningnets

Central Services VPN Configuration Example

The configuration example in the figure shows how to configure a central services VPN in a Cisco IP Next-Generation Network service provider network.



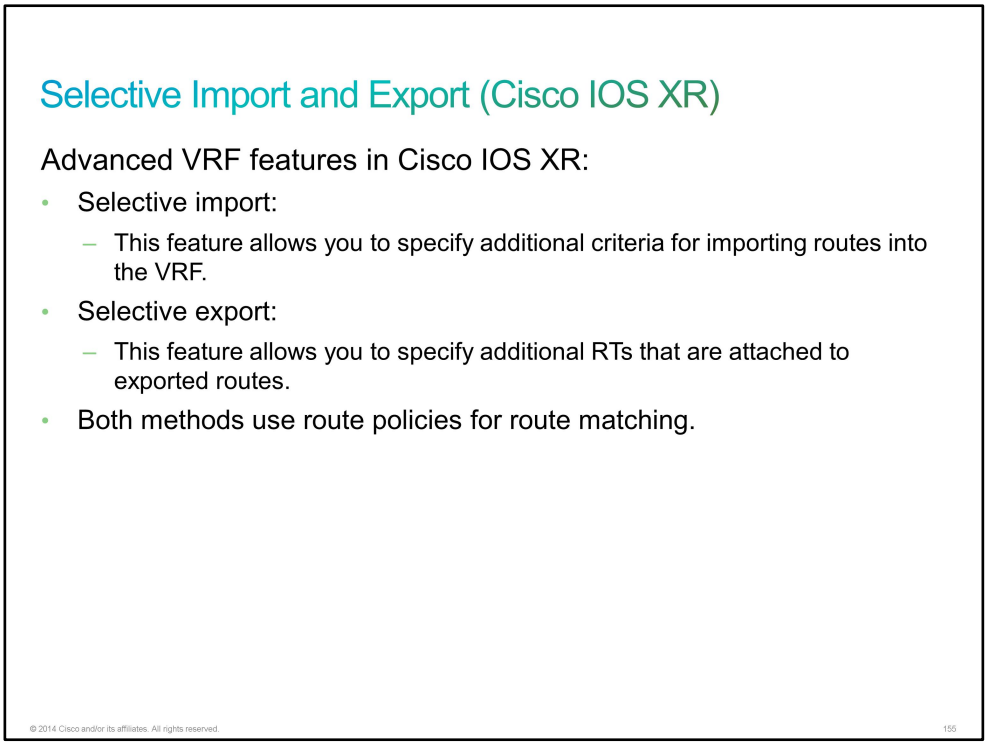
On PE1, two VRF instances are configured for Customer A and Customer B. On the PE-CS-1 router, the VRF server is configured.

Router PE-CS-1 exports routes with RT 1:502 and imports client routes with RT 1:501.

Customer VRF exports routes with RT 1:501 and imports server routes with RT 1:502.

https://t.me/learningnets

Selective Import and Export (Cisco IOS XR)



These VRF features allow you to deploy advanced MPLS VPN topologies or increase the stability of the MPLS VPN backbone:

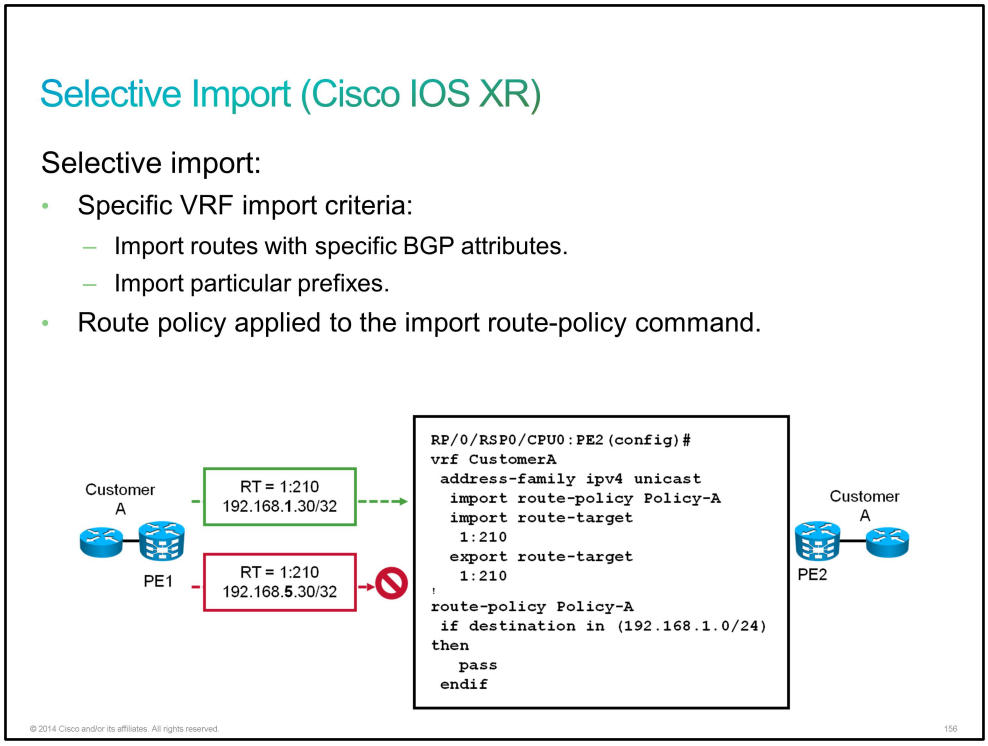
- The selective import feature allows you to select the routes to be imported into a VRF, based on criteria other than the RT of the VRF.
- The selective export feature allows you to attach specific RTs to a subset of routes that are exported from a VRF. By default, the same RTs get attached to all exported routes.

Note	The VRF route limit is also an advanced VRF feature on some platforms that allows you to limit the number of routes that the customer—or other PE routers—can insert in the VRF. This feature prevents undesirable consequences such as configuration errors or DoS attacks.

https://t.me/learningnets

Selective Import (Cisco IOS XR)

Selective route import into a VRF allows you to narrow the route import criteria.



Selective route import uses a route policy that can filter the routes selected by the RT import filter. The routes imported into a VRF are BGP routes, so you can use match conditions in a route policy to match any BGP attribute of a route. These attributes include communities, local preference, MED, AS path, and so on.

The **import route-policy** command is combined with the RT import filter. A route must pass the RT import filter first and then the import route policy. The necessary conditions for a route to be imported into a VRF are as follows:

- At least one of the RTs attached to the route matches one of the import RTs configured in the VRF.
- The route is permitted by the import route policy.

The figure shows an example in which an import route policy is used to match the IPv4 portion of incoming VPNv4 routes and to import into the VRF, only the routes matching a certain prefix.

A configuration similar to this one could be used to accomplish the following:

- Deploy advanced MPLS VPN topologies (for example, a managed router services topology).
- Increase the security of an extranet VPN by allowing only predefined subnetworks to be inserted into a VRF, thus preventing an extranet site from inserting unapproved subnetworks into the extranet.

https://t.me/learningnets

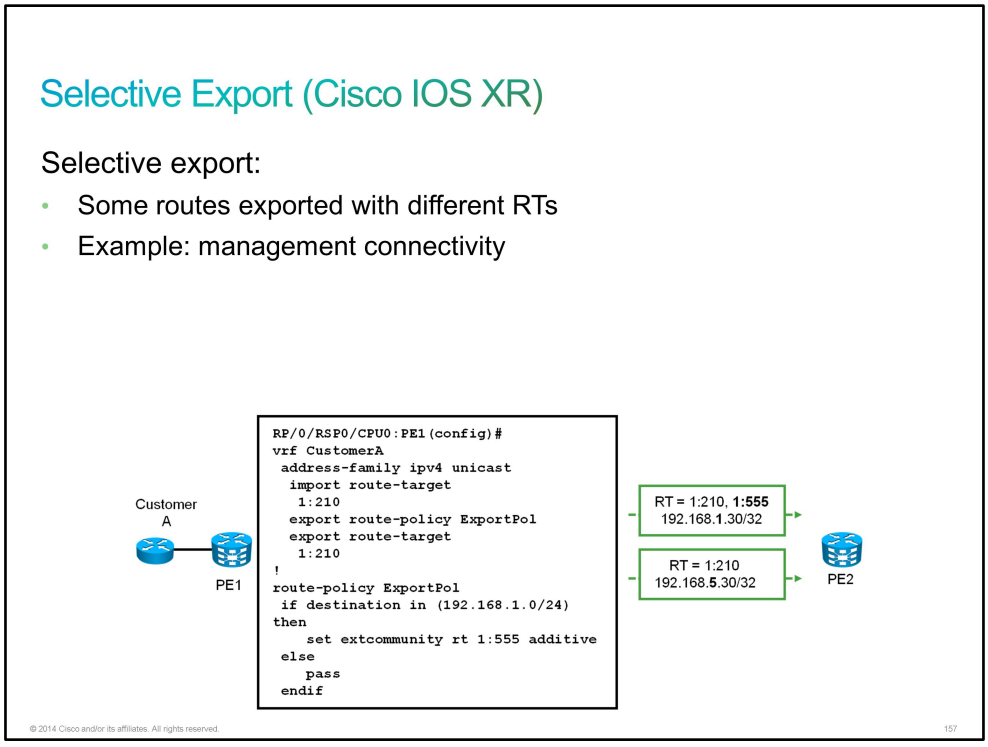
Note	A similar function is usually not needed in an intranet scenario because all customer routers in an intranet are usually under common administration.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Selective Export (Cisco IOS XR)

Some advanced MPLS VPN topologies are easiest to implement if you can attach various RTs to routes exported from the same VRF. This capability allows only a subset of the routes exported from a VRF to be imported into another VRF. Most services in which customer routers need to connect to a common server (for example, network management stations, voice gateways, and application servers) fall into this category.



The **export route-policy** command provides exactly this functionality. A route policy can be specified for each VRF to attach additional RTs to routes exported from that VRF. The **export route-policy** command performs only the attachment of RTs. It does not perform any filtering functions.

Attributes attached to a route with an export route policy are combined with the export RT attributes. If you specify export RTs in a VRF and set RTs with an export route policy, all specified RTs will be attached to the exported route.

Note	The export route policy provides functionality that is almost identical to that of the import route map, but applied to a different VRF. Any requirement that can be implemented with an export route policy can also be implemented with an import route policy. However, the implementation of export maps can be more complicated and difficult to manage.

In the figure, the configuration is implemented with an export route policy.

https://t.me/learningnets

Note	Depending on when you configure the **`export map`** command, you might need to use the **`clear bgp`** command to force the existing BGP session to propagate the extended communities.

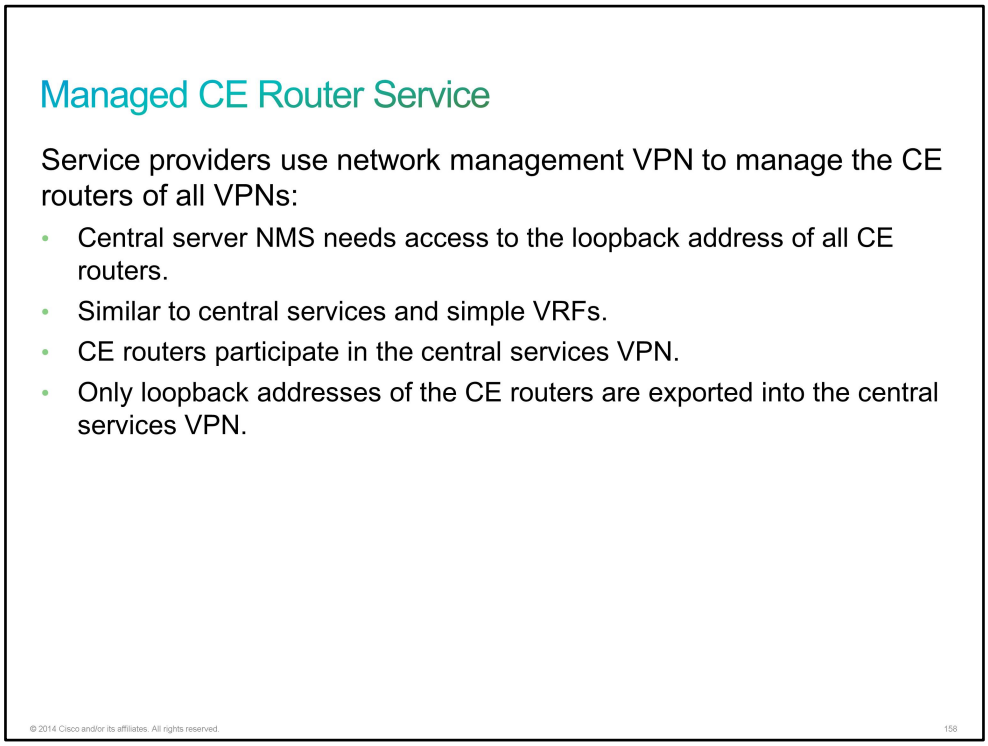
In this example, routes from a certain address block are marked with an additional RT in the originating VRF and are automatically inserted into the receiving VRF, based on their RT.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Managed CE Router Service

If the service provider is managing the customer routers, it is convenient to have a central point that has access to all CE routers but does not have access to the other destinations at the customer sites. This requirement is usually implemented by deploying a separate VPN for management purposes. This VPN needs to see all the loopback interfaces of all the CE routers. All CE routers need to see the network management VPN. The design is similar to that of the central services VPN; the only difference is that you mark only loopback addresses to be imported into the network management VPN.



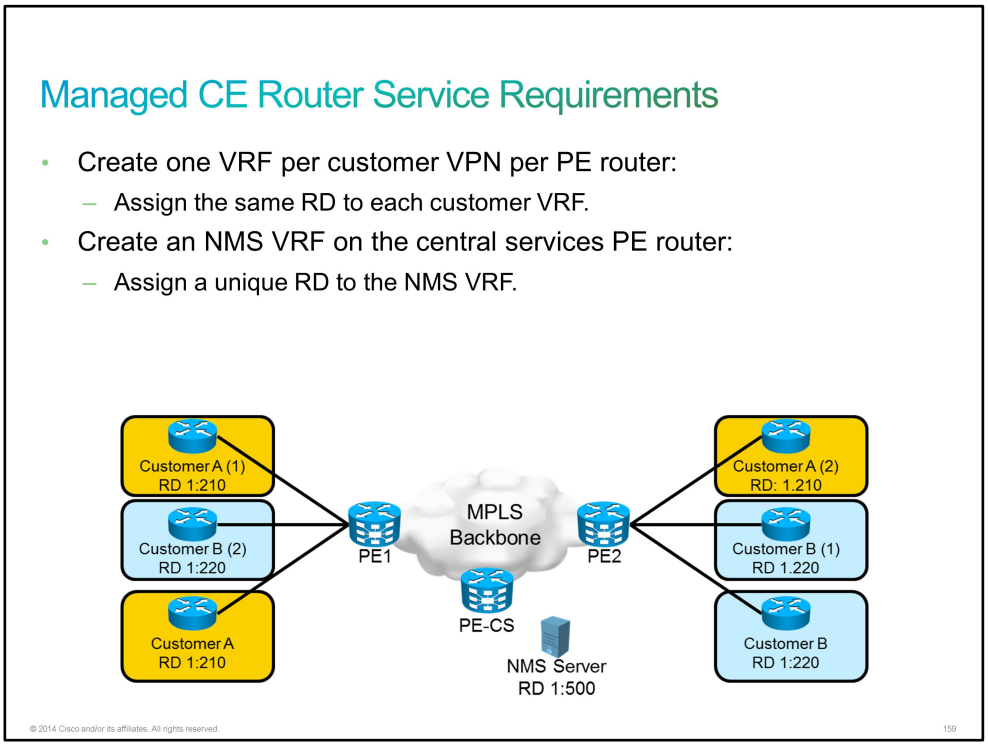
Managed CE router service features include the following:

- The central server NMS needs access to the loopback address of all CE routers.
- It is similar to central services and simple VRFs.
- CE routers participate in the central cervices VPN.
- Only the loopback addresses of the CE routers are exported into the central services VPN.

https://t.me/learningnets

Managed CE Router Service Requirements

To deploy a managed CErouter VPN, you will need to create a VRF for each customer router on every PE. These VRFs can share the same RD, although it is not necessary. The RDs could differ.

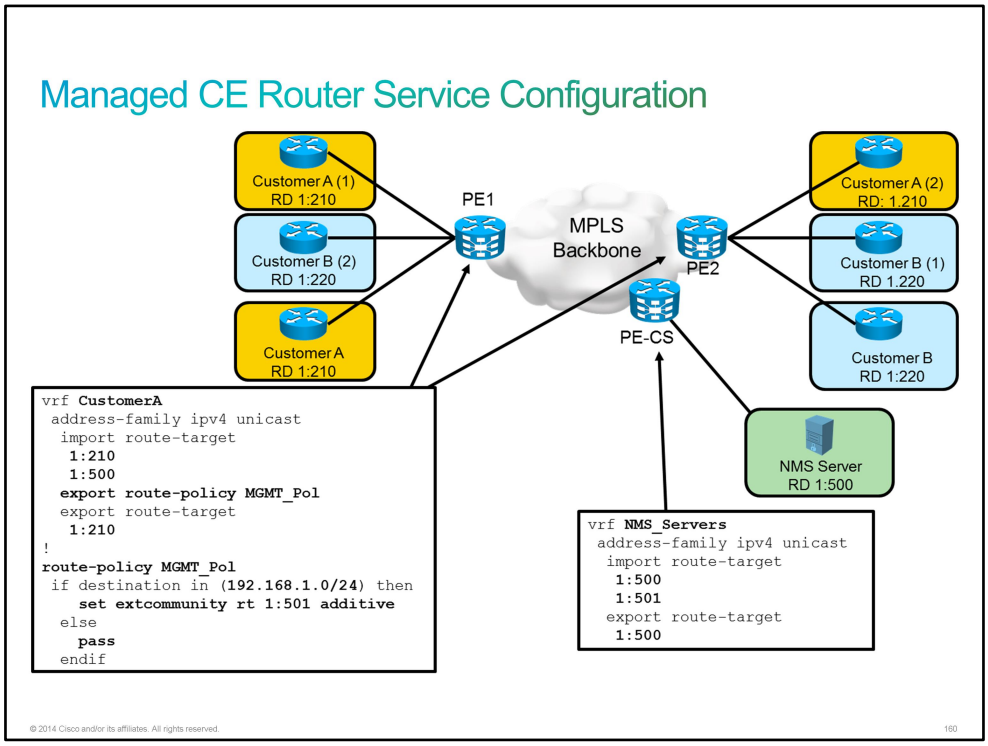


You also need to configure a NMS VRF on the PE router where the network management server is attached. This VRF should have a unique RD.

https://t.me/learningnets

Managed CE Router Service Configuration

You can have a configuration for a customer VRF with differentiated RT export for loopback addresses. The figure shows this example. An export route policy is used to match one part of the IP address space and attach an additional RT to the routes within this address space (CE router loopback addresses).



Note	The routing protocol between PE and CE routers must be secured (with distribute lists or prefix lists) to prevent customers from announcing routes in the address space dedicated to network management; otherwise, customers can gain two-way connectivity to the network management station.

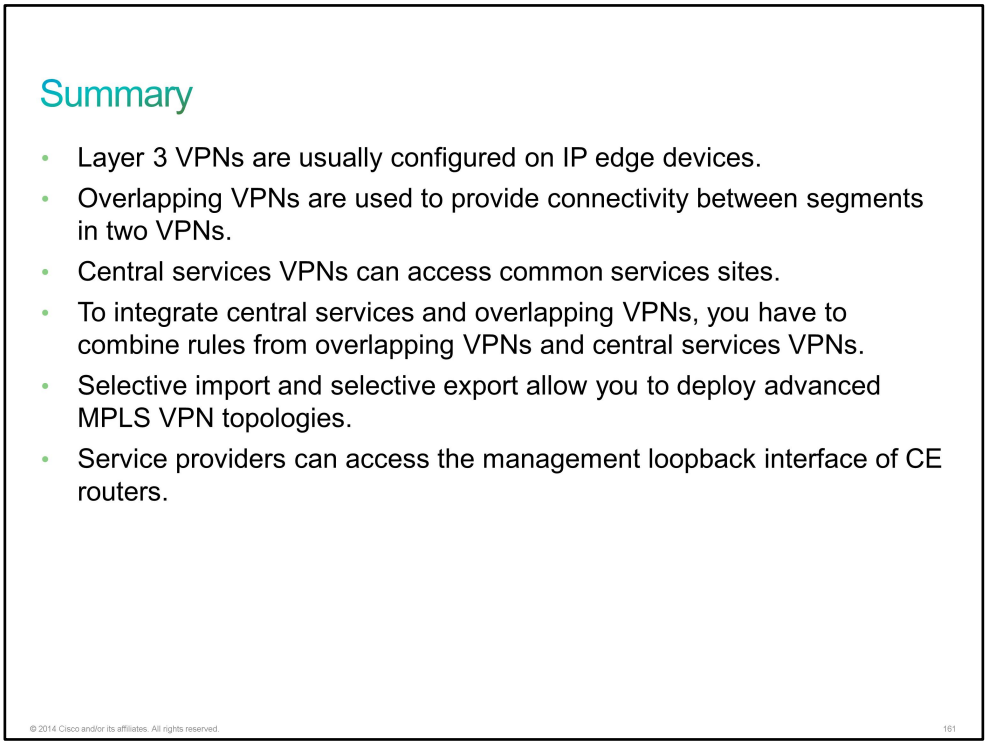
The CE router loopback addresses are then imported into the server VPN based on the additional RT attached to them during the export process.

Note	This design allows client sites to send packets to the network management VPN regardless of the source address. Special precautions should be taken to protect the network management VPN from potential threats and DoS attacks coming from customer sites.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

Lesson 2: Implementing Internet Access in MPLS VPNs

Overview

Integrating Internet access with a MPLS VPN solution is one of the most common service provider business requirements. This lesson provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Upon completing this lesson, you will be able to:

- Describe common customer Internet connectivity scenarios
- Describe classical Internet access
- Describe multisite Internet access
- Describe wholesale Internet access
- Describe the SP shared backbone
- Describe the Internet access design models
- Describe the classic topology of Internet access through global routing
- Describe Internet access through global routing with separate interfaces
- Describe Internet access through global routing with static routing
- Describe Internet access through global routing with BGP
- Describe the VRF-aware tunnel configuration in Cisco IOS XR
- Describe the VRF-aware tunnel configuration in Cisco IOS/XE

https://t.me/learningnets

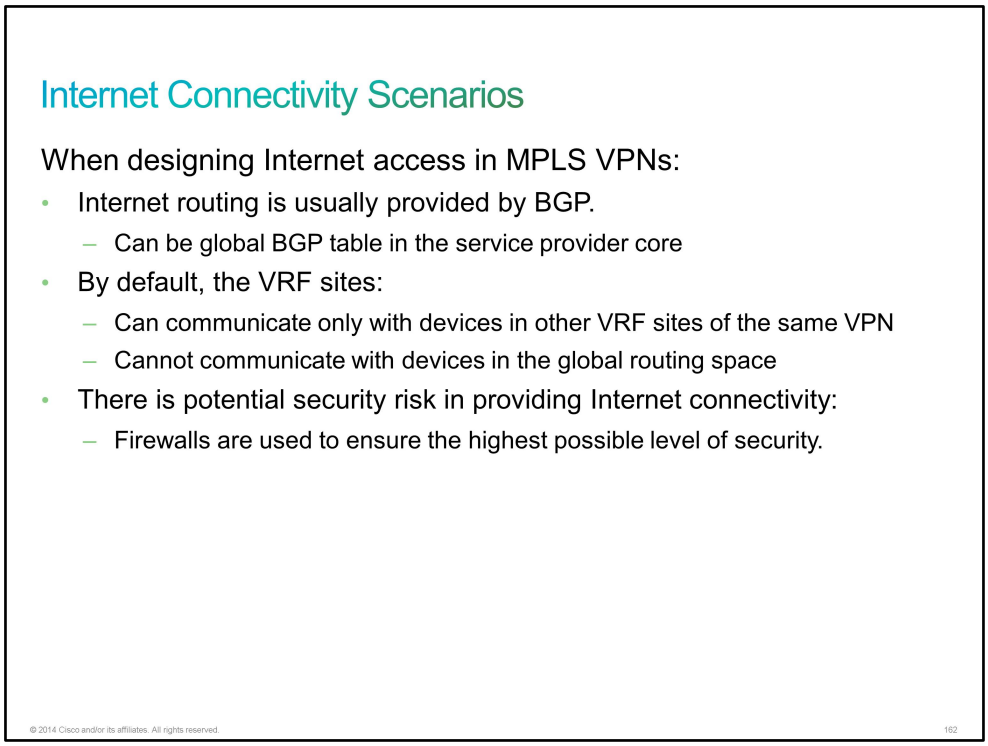
- Describe the multisite Internet access through global routing
- Evaluate Internet access via global routing
- Describe Internet access through a separate VPN
- Describe the Internet access through a separate VPN features
- Describe the separate VPN Internet gateway
- Describe the separate VPN Internet gateway configuration
- Describe the separate VPN Internet access from a central site
- Describe the VRF-aware tunnel configuration in Cisco IOS XR
- Describe the VRF-aware tunnel configuration in Cisco IOS/XE
- Describe redundant Internet access through a separate VPN
- Describe multisite Internet access through a separate VPN
- Describe wholesale Internet access through a separate VPN
- Evaluate Internet access through a separate VPN

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Internet Connectivity Scenarios

Internet routing is usually performed via the BGP table of the MPLS VPN network of the service provider.



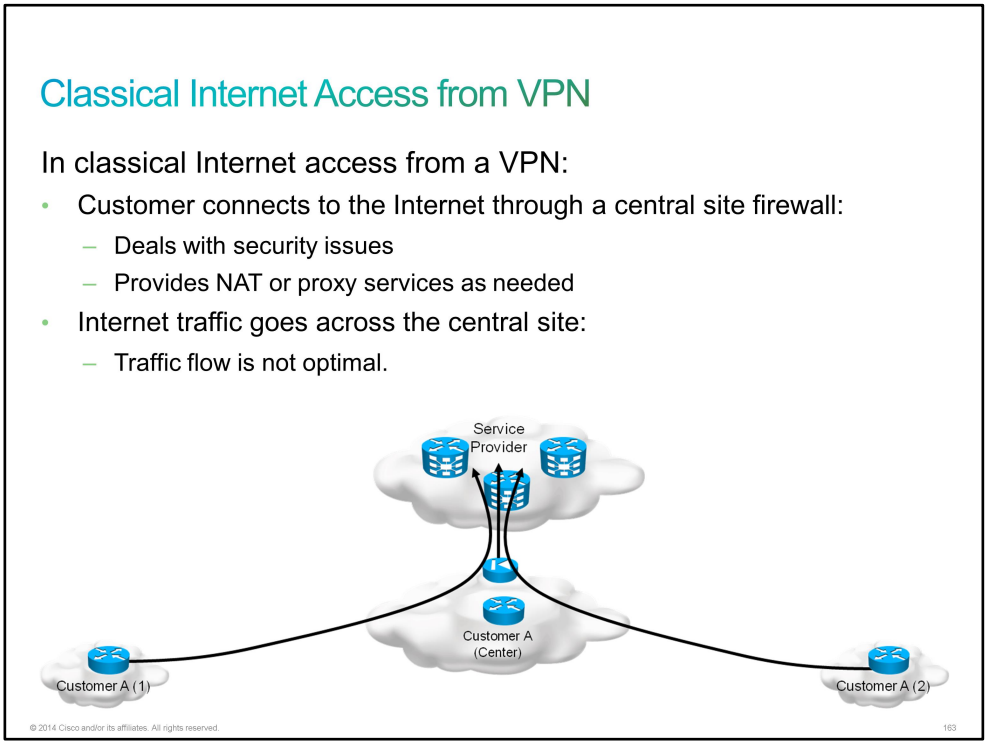
The BGP table is in the global routing space, not in the VRF context. By default, the VRF sites can communicate only with other VRF sites in the same VPN, not with anything in the global routing space. Therefore, something must be done to provide Internet access (global context) to the CE routers (VRF context). The following subtopics explain how to provide Internet access to VRF sites. Internet access is possible only for those customer IP subnets that are not from the private IP addressing space (RFC 1918).

Note	As soon as the VPN has Internet connectivity, a potential security risk exists. Take the proper steps—such as filtering and using a firewall—to ensure the highest possible level of security.

https://t.me/learningnets

Classical Internet Access from VPN

Classical Internet access is implemented through a central firewall that connects the customer network to the Internet in a secure fashion.



The customer network and the Internet are connected only through the firewall. The addressing requirements for this type of connection are simple. The customer is assigned a small block of public address space that is used by the firewall. The customer typically uses private addresses inside the customer network if they are using IPv4 addressing. The firewall performs NAT between the private addresses of the customer and the public addresses that are assigned to the customer by the ISP. Alternatively, the firewall might perform an application-level proxy function that also isolates private and public IP addresses.

If a customer is using IPv6 addressing, there is no need for NAT, but the customer still needs the other functions that a firewall provides.

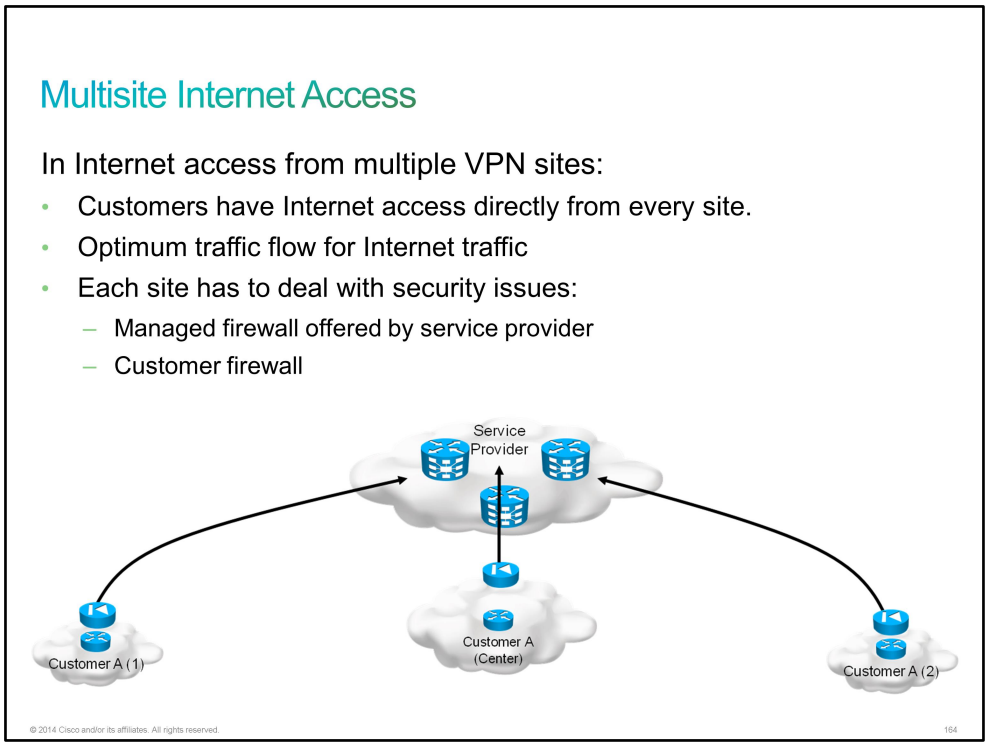
Several benefits are associated with this design. The setup is well known, and the expertise needed to implement it is simple and straightforward. Only one interconnection point between the secure customer network and the Internet needs to be managed.

The major drawback of this design is the traffic flow. All traffic from the customer network to the Internet passes through the central firewall. Although this flow might not be a drawback for smaller customers, it can be a severe limitation for large organizations with many users, especially when they are geographically separated.

https://t.me/learningnets

Multisite Internet Access

Some customers find the traffic flow limitations of the central firewall setup too limiting.



To bypass the limitations of Internet access through a central firewall, some customers use designs in which each customer site has its own independent Internet access.

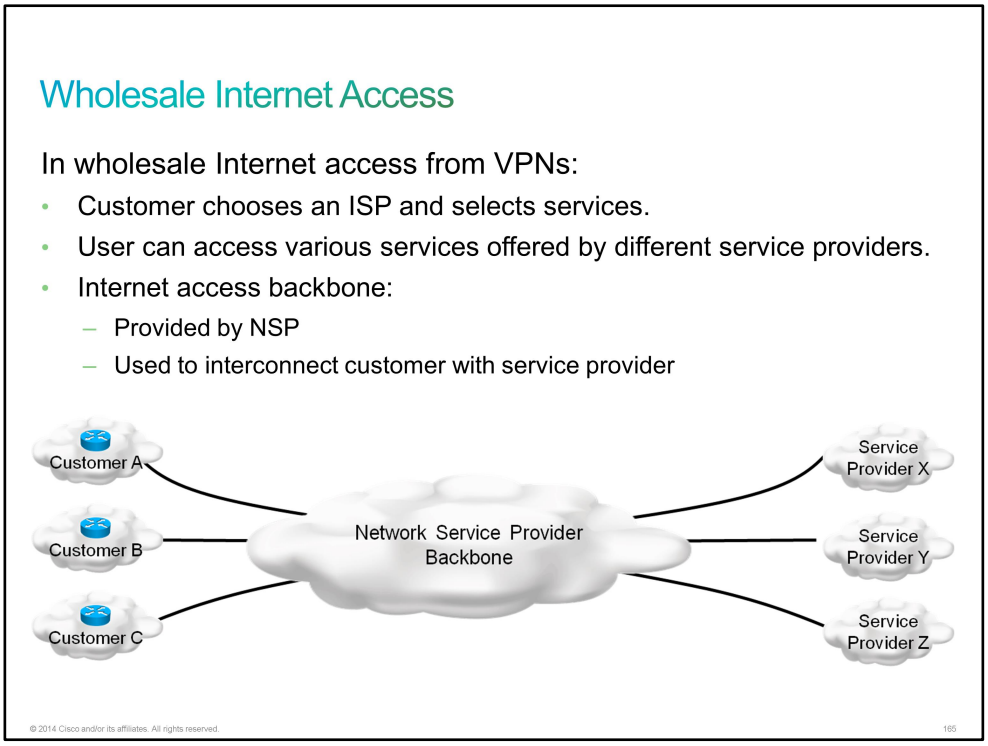
This design solves traffic flow issues, but the associated drawback is higher exposure. Each site needs to be individually secured against unauthorized Internet access, leading to the increased complexity of managing a firewall at every customer site.

To achieve Internet access from every customer site, each CE router must forward VPN traffic toward other customer sites and forward Internet traffic toward Internet destinations. The two traffic types are usually sent over the same physical link, but over different logical links, to minimize costs.

For customers that do not want the complexity of managing their own firewalls, a managed firewall service offered by the service provider can help address the security issues of Internet connectivity.

https://t.me/learningnets

Wholesale Internet Access



A service provider might provide wholesale Internet access from a range of upstream ISPs to satisfy the connectivity and reliability requirements of various customers.

The selection of upstream ISPs and the corresponding configuration processes should therefore be as easy as possible for the service provider. From an Internet perspective, customers A, B, and C are connected to ISP X or ISP Y. The IP address space that the customer uses should be allocated from the block of addresses that is administered by the selected ISP. The service provider that provides wholesale Internet access might need to use a different address for each upstream ISP.

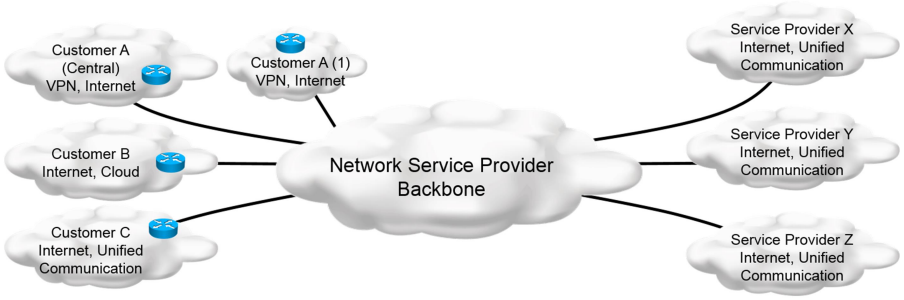
https://t.me/learningnets

Service Provider Shared Backbone

Service Provider Shared Backbone

Multiple providers can offer services over a shared core:

- Internet is traditionally the most popular service.
- Clients access different services from their service providers:
  - Internet, Unified Communications, cloud services, and so on
- Cisco IP NGN support multiple services over a common backbone.

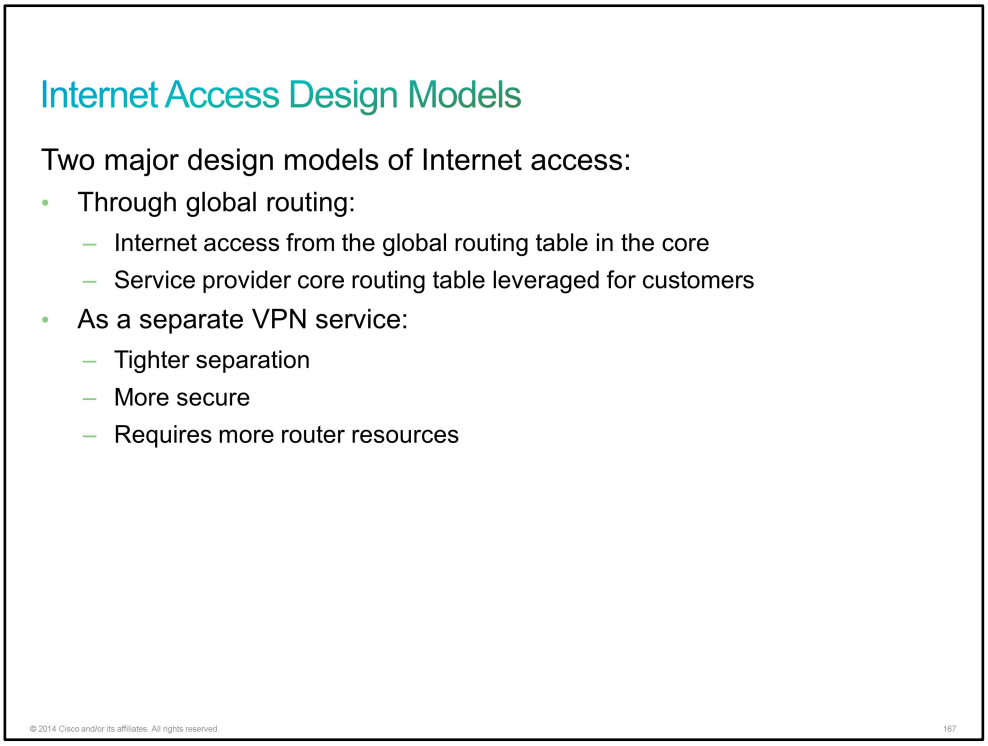


© 2014 Cisco and/or its affiliates. All rights reserved.

Because Internet access is one of the most popular services that service providers offer their customers, many service providers offer Internet access as well as an MPLS VPN service on their shared backbone. Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. A background of common customer Internet connectivity scenarios will help in assessing possible implementations.

https://t.me/learningnets

Internet Access Design Models



Network designers who want to offer Internet access and MPLS VPN services on the same backbone can choose between these two major design models:

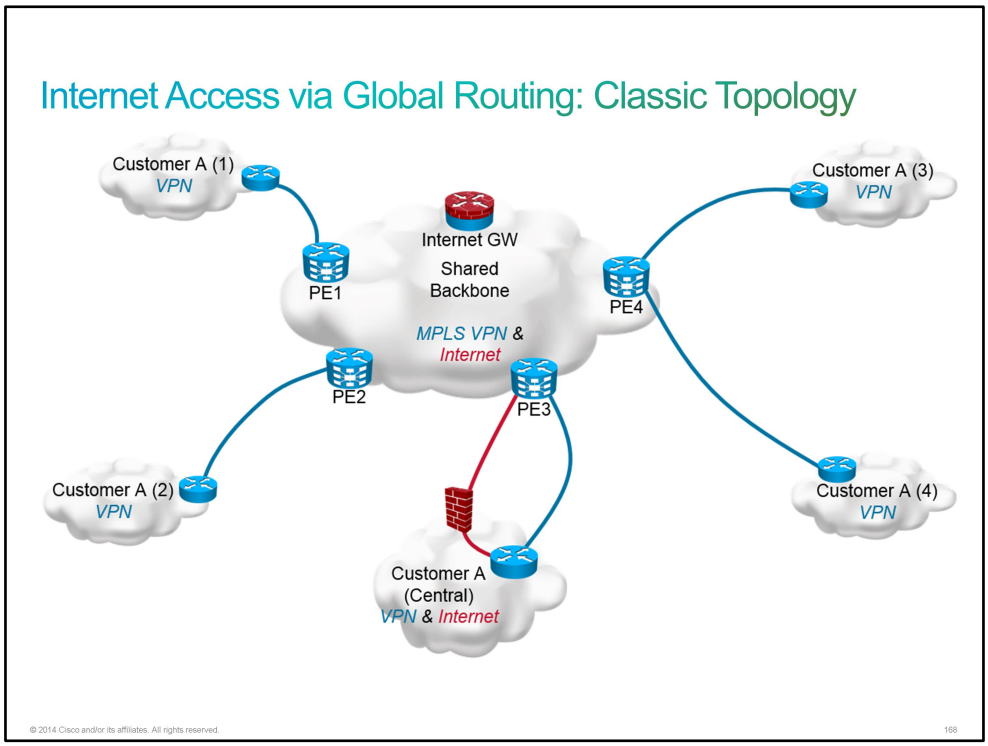
- Internet access that is implemented through global routing on the PE routers and that is not a VPN service
- Internet access that is implemented as a separate VPN in the ISP network

In both cases, security should be the most important concern for customers when they connect to public networks. Customers should isolate private VPNs from Internet traffic, either physically (on a separate interface) or on a subinterface. Appropriate firewall support, either in a dedicated device or integrated in the router Cisco IOS Software, is a necessity. Depending on the network addressing, NAT will be needed for most customers.

https://t.me/learningnets

Internet Access via Global Routing: Classic Topology

The classical Internet access design for a customer is based on a separate Internet access model.



One central customer site has connectivity to the Internet and provides access to the rest of the customer sites. The central site either connects through a firewall or runs the Cisco IOS Firewall Feature Set.

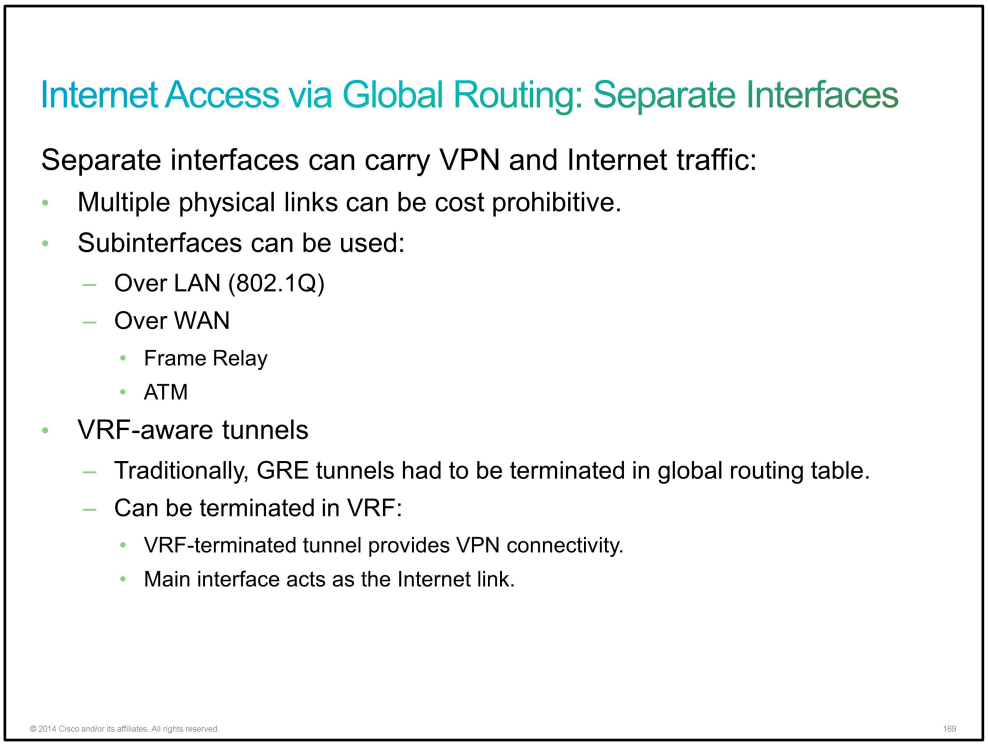
In the shared service provider backbone, the PE routers have full or partial Internet routes to offer the customer. The Internet gateway of the provider is in the same IBGP domain as the provider and PE routers.

This design model can easily map to a customer with an MPLS VPN implementation. In this example, the customer network has been interconnected with an MPLS VPN. A central CE router has Internet connectivity and provides Internet access, through a firewall, for all the sites in the customer network.

This traditional Internet access implementation model provides maximum design flexibility because the Internet access is completely separated from the MPLS VPN services. However, the limitations of traditional IP routing prevent this implementation method from being used for innovative Internet access solutions, such as wholesale Internet access.

https://t.me/learningnets

Internet Access via Global Routing: Separate Interfaces



Instead of separate physical links for VPN and Internet traffic, subinterfaces can be used to create two logical links over a single physical link. Subinterfaces can be configured only on WAN links that use Frame Relay or ATM encapsulation (including xDSL) and on LAN links that use any VLAN encapsulation (802.1Q).

For other encapsulation types, a tunnel interface can be used between the CE router and the PE router. Depending on the router platform and Cisco IOS Software version, VRF-aware tunnels are now supported.

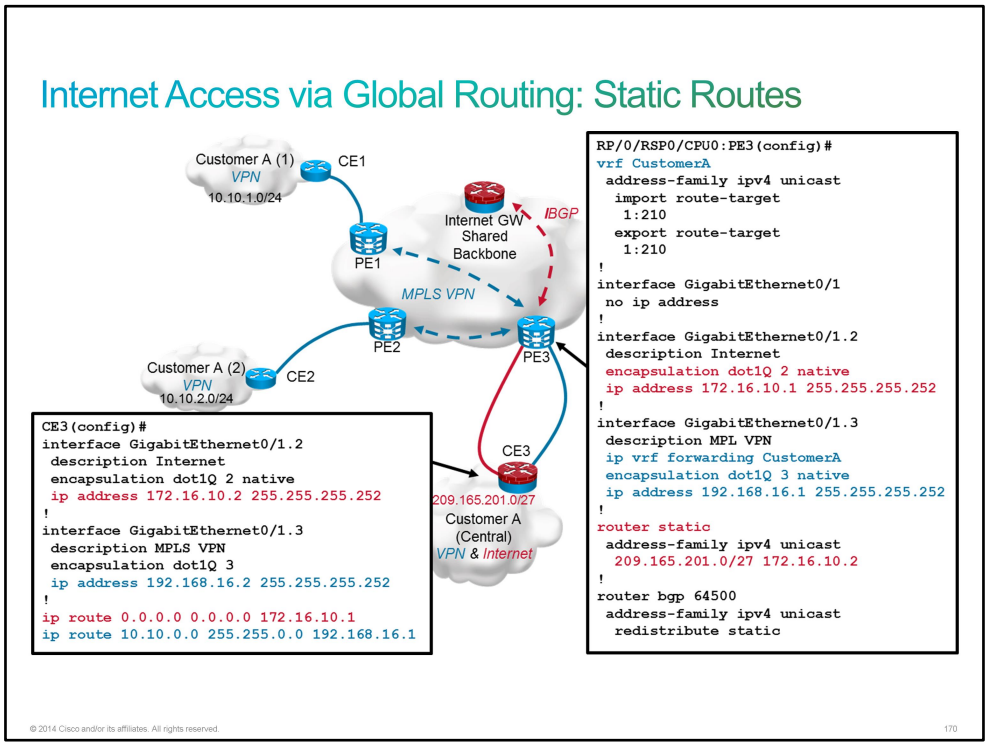
Consider these guidelines when designing VRF-aware tunnels:

- VRF-aware tunnels remove the need for the endpoints of the tunnel to be in a global routing table.
- Without a VRF-aware tunnel, MPLS VPN traffic would need to be tunneled across the Internet interface.

https://t.me/learningnets

Internet Access via Global Routing: Static Routes

Using static routing on the CE and PE routers is the simplest and most common implementation for providing Internet access.



The figure illustrates a configuration that is used to implement Internet access through two 802.1Q subinterfaces with static routes.

In this simple example, the Customer A router does not need to receive full Internet routing. To reach the Internet, the Customer A router just needs a default static route to PE3. PE3 has a route to the Internet through the Internet gateway, as well as a static route for the customer A subnets that point to the Customer A router. The full Internet routing table needs to be present only on the Internet gateway.

The following configuration steps are performed:

1. The customer VRF instance is created for private MPLS VPN.
2. The VPN subinterface is created and associated with the proper VLAN. The subinterface is added to the customer VRF.
3. Another subinterface is created for Internet access.
4. On the PE router, static routes for customer public address space pointing to the customer next hop are configured. These routes are redistributed in BGP routing protocol.

In addition to the PE configuration, the CE router implements a default static route that points to the PE router.

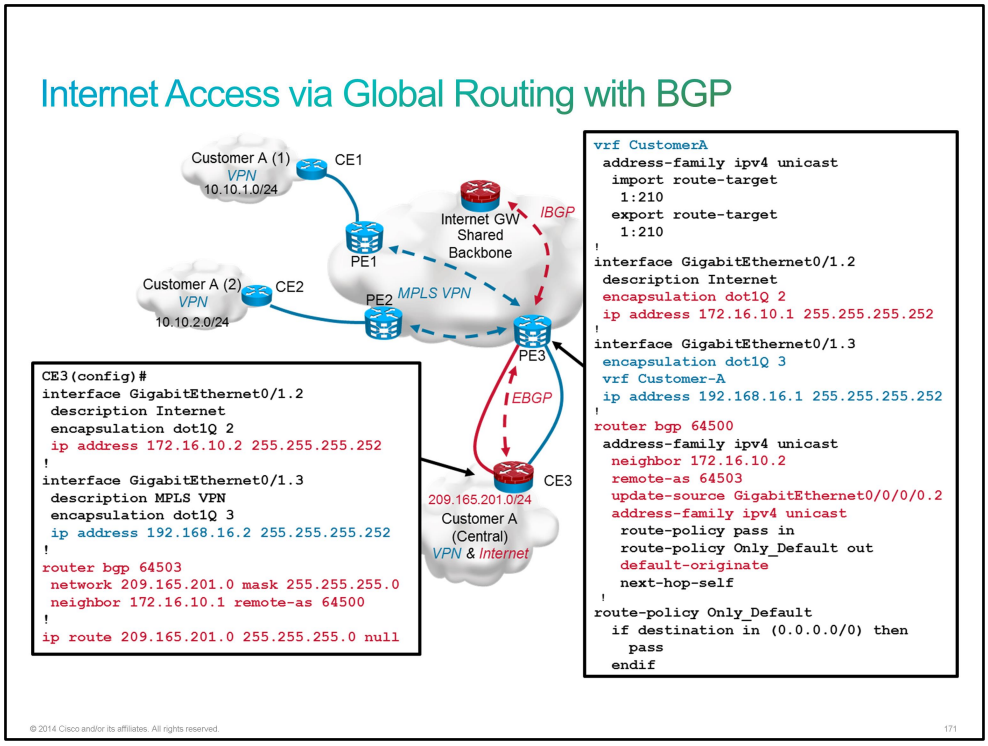
https://t.me/learningnets

Note	On the CE-Central router, distribution of the default route might be needed so that remote sites can also access the Internet. Issues of security and private addresses would need to be resolved.

https://t.me/learningnets

Internet Access via Global Routing with BGP

In this example, Customer A is using a dynamic routing protocol to establish Internet connectivity.



BGP routing protocol is configured on the Customer A router. Customer A advertises its own public address space to EBGP neighbor. Inserting customer public address space into the customer routing table is important. If the network is not in the routing table, the network is not advertised using BGP.

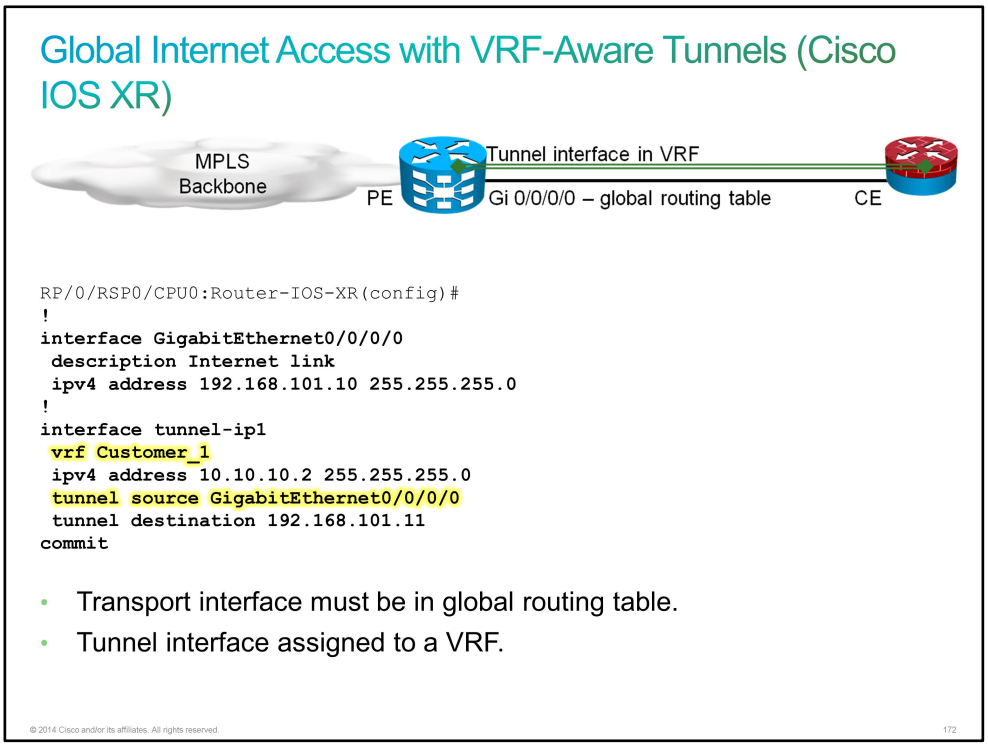
On the PE router, the service provider has to configure the VRF instance and subinterface for the private MPLS VPN of the customer and another subinterface for Internet access.

BGP is used for route exchange between the customer and service provider. The service provider can originate the default route to the customer peer. The service provider can also filter advertised routes and send only the default route to the customer. A route policy is used to filter routes on Cisco IOS XR routers, and prefix lists are used to filter advertised routes on IOS routers.

https://t.me/learningnets

Global Internet Access with VRF-Aware Tunnels (Cisco IOS XR)

Instead of separate links of subinterfaces, you can configure GRE tunnels.



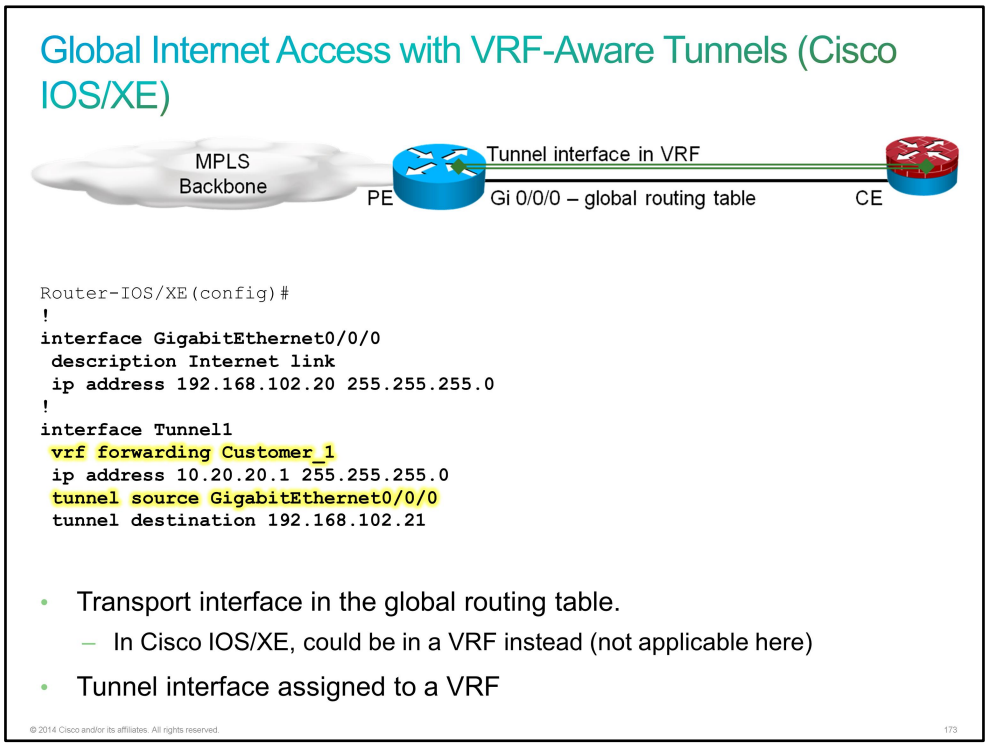
VRF-aware tunnels in Cisco IOS XR can be terminated in a VRF although the tunnel transport goes through the global routing table.

In this example, the physical interface is assigned to the global routing on the PE router. It will be used to carry the Internet traffic. The tunnel-ip1 interface is established between the physical IP address and its directly adjacent IP address on the CE router. The tunnel interface has been assigned to the customer VRF and will ensure the VPN communication.

https://t.me/learningnets

Global Internet Access with VRF-Aware Tunnels (Cisco IOS/XE)

Cisco IOS/XE allows you more flexibility than Cisco IOS XR when deploying VRF-aware tunnels.



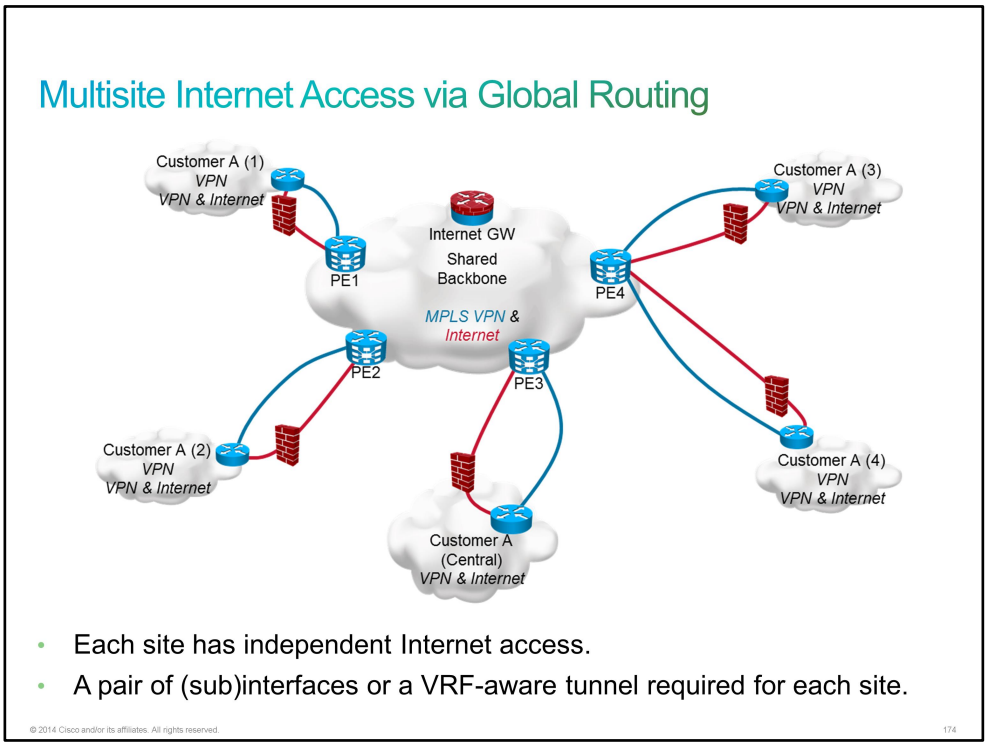
In Cisco IOS/XE the transport interface can be either in the global address space or belong to a VRF. The tunnel interface can be in another VRF.

The additional flexibility is not leveraged in the scenario with Internet access through the global address space. As before, the physical interface is assigned to the global routing. It will be used for Internet traffic. The Tunnel1 interface is established between the physical IP address and its directly adjacent IP address on the CE router. The tunnel interface has been assigned to the customer VRF and will provide connectivity within the VPN.

https://t.me/learningnets

Mutisite Internet Access via Global Routing

Another option is to provide separate Internet access at every customer site.

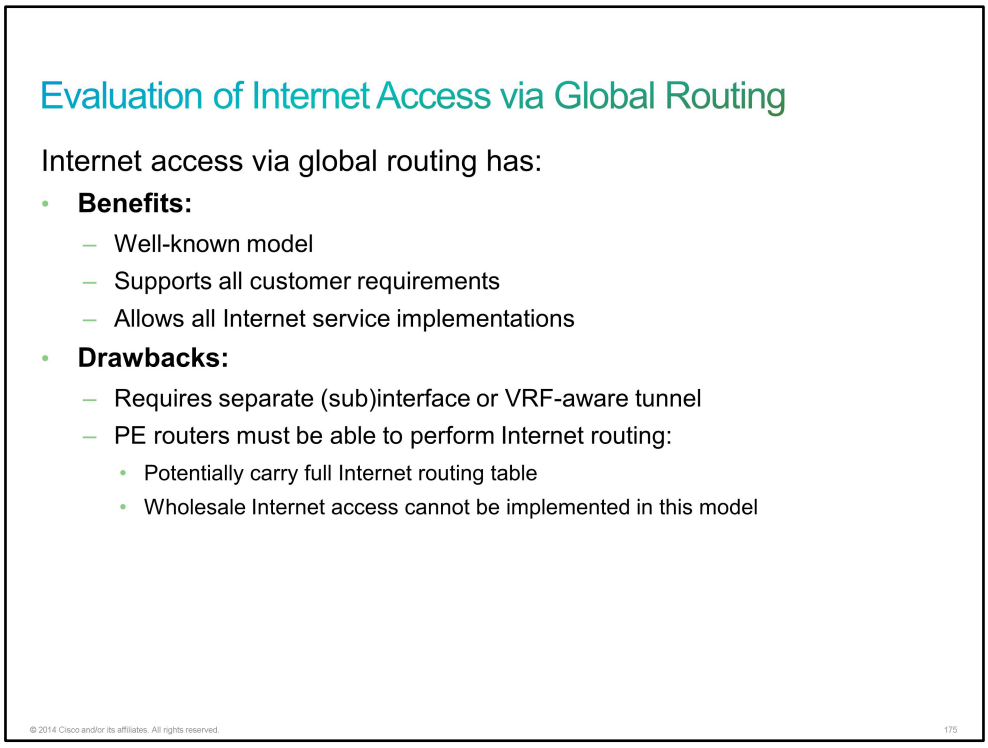


In this case, two physical (or logical) links between every CE router and its PE router would be needed. This design often becomes too complex or too expensive to implement. Issues such as customer route propagation to the Internet and securing access at multiple access points would need to be resolved.

Note	The allowas-in feature might need to be configured on the PE router if the customer is propagating individual site routes to the Internet through BGP.

https://t.me/learningnets

Evaluation of Internet Access via Global Routing



The benefits of a separate Internet access design model are as follows:

- The model is well known and widely understood.
- The model supports all customer requirements, including multihomed customer connectivity with full Internet routing.
- The model allows all Internet service implementations, including BGP sessions with customers.

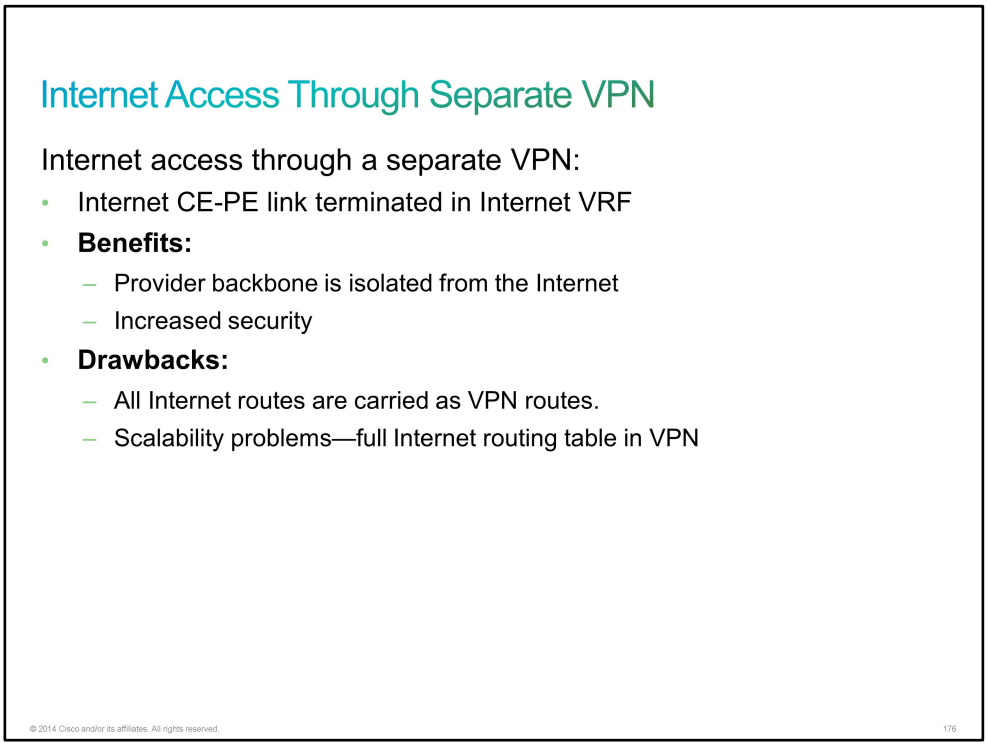
The drawbacks of this model are as follows:

- The model requires two dedicated physical links between the PE and the CE router or specific WAN or LAN encapsulations that might not be suitable for all customers.
- The PE routers must be able to perform hop-by-hop Internet routing and either use the default route to reach the Internet or carry the full Internet routing table.
- Advanced Internet access services (centrally managed firewall service or wholesale Internet access service) cannot be realized with this model.

https://t.me/learningnets

Internet Access Through Separate VPN

For a service provider, implementing Internet access through a separate VPN is similar to offering another managed VPN service.

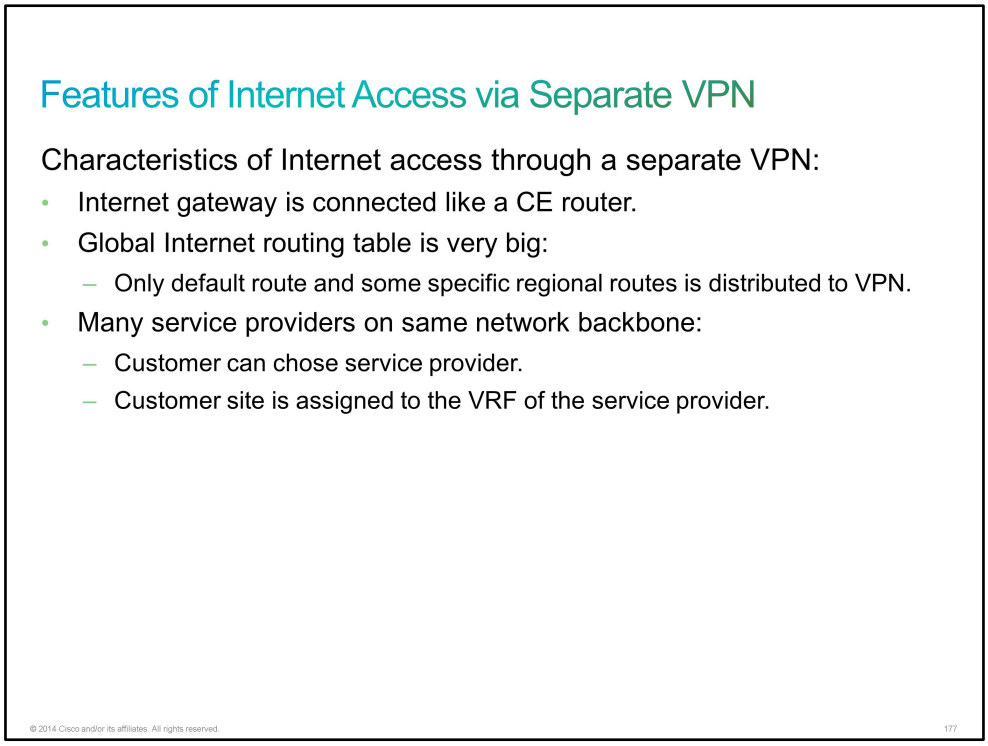


The major benefit of implementing Internet access as a separate VPN is increased isolation between the provider backbone and the Internet—which results in increased security for the provider. The flexibility of MPLS VPN topologies also provides for some innovative design options that allow service providers to offer services that were simply impossible to implement with pure IP routing.

The obvious drawback of running the Internet as a VPN in the MPLS VPN architecture is the scalability of such a solution. The Internet VPN, or possibly multiple Internet VPNs, require considerable router resources to carry the Internet routing table. As a result, scalability constraints are associated with having all the Internet routes inside a single VPN.

https://t.me/learningnets

Features of Internet Access via Separate VPN

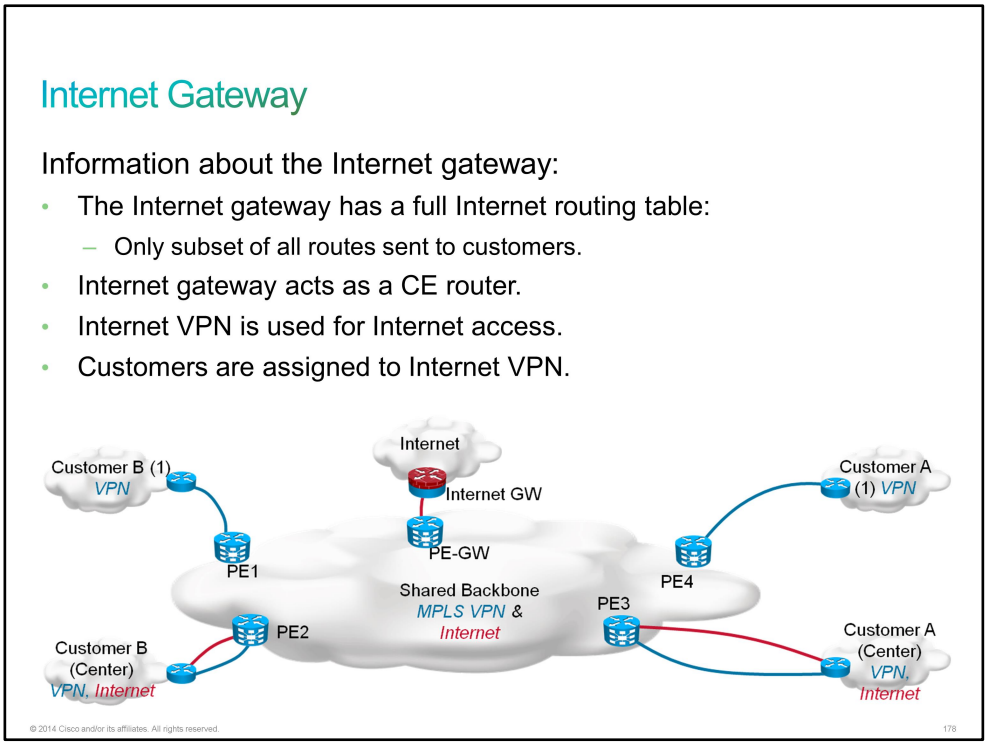


The MPLS VPN architecture can provision a separate VPN to provide Internet access for VPN customers. The service provider defines the Internet VPN and can use different MPLS VPN topologies to implement various types of Internet access. Under this design model, the provider Internet gateways appear as CE routers to the MPLS VPN backbone. Customer Internet access is enabled by using a dual VPN topology that supports both an Internet VPN and a customer VPN across separate customer interfaces.

In this design, the Internet VPN should not contain the full set of global Internet routes because that would make the solution completely nonscalable. The provider Internet gateway routers should announce a default route toward the PE routers. To optimize local routing, the local and regional Internet routes should be inserted in the Internet VPN.

https://t.me/learningnets

Internet Gateway



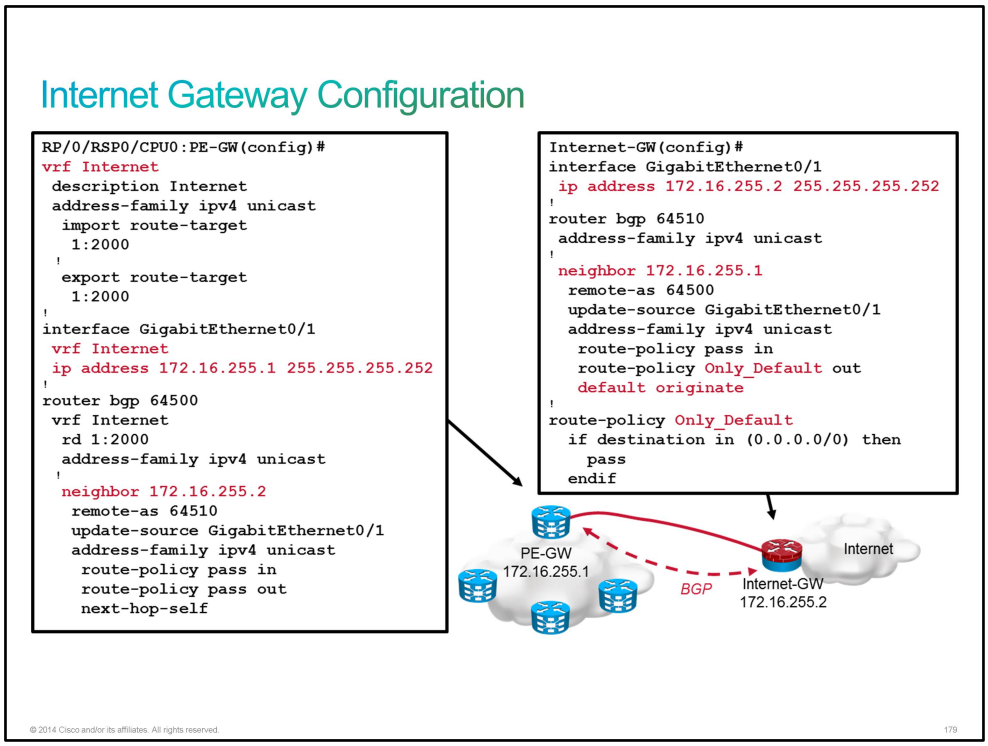
When the service provider implements Internet access as a separate VPN, the Internet backbone is carried on a VPN, which is isolated from the provider backbone. This topology results in increased security for the provider backbone because Internet hosts can reach only PE routers, not the core provider routers. The VPN customers are connected to the Internet simply through an additional VRF instance at the PE.

The Internet gateway acts as a CE router and holds the full BGP routing table. Only the subset of routes and the default route are advertised to the clients assigned to the Internet VPN.

https://t.me/learningnets

Internet Gateway Configuration

The figure shows a sample configuration of the Internet gateway router with default route advertisement.



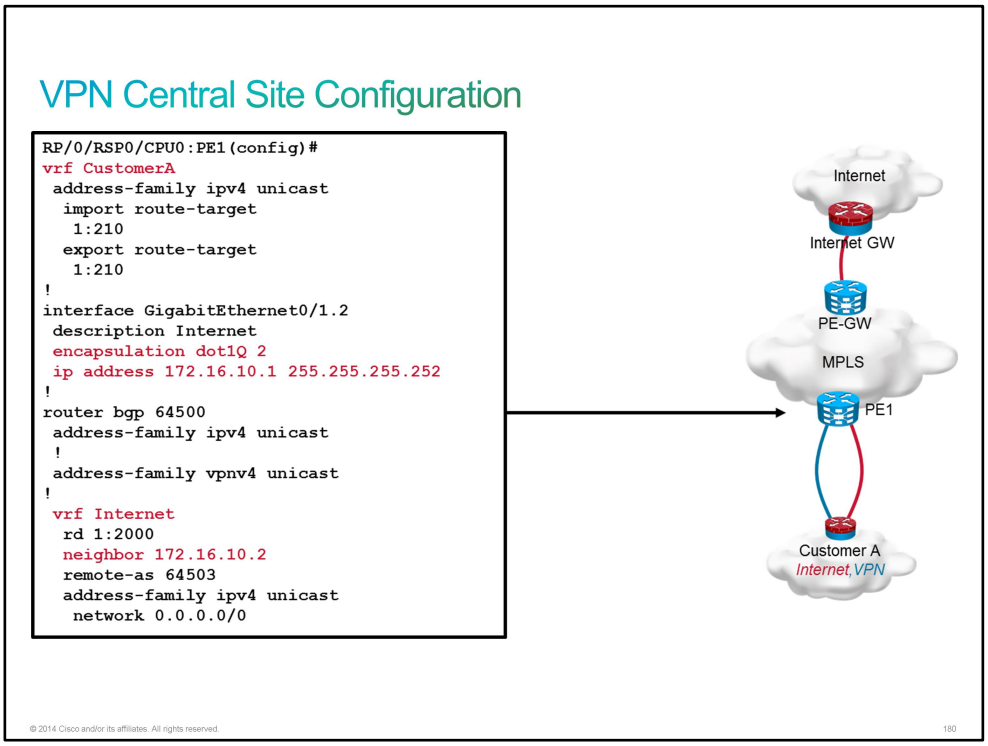
The Internet gateway should be able to route traffic to the Internet. An EBGP session is established with the PE gateway router. Only the default route is advertised to the PE gateway router.

On the PE gateway router, a new VRF instance for Internet is used. The interface facing the Internet gateway is assigned to the Internet VRF. In the BGP process, the Internet VRF has to be enabled and appropriate address families have to be activated (such as IPv4 and IPv6). The PE gateway router distributes the default Internet route among the other PE routers in the MPLS VPN network.

https://t.me/learningnets

VPN Central Site Configuration

The classical Internet access model can easily be implemented with the Internet VPN over the MPLS VPN backbone.



The link between a PE router and the Internet gateway router is assigned to the Internet VRF. The Internet gateway announces a default route to the Internet.

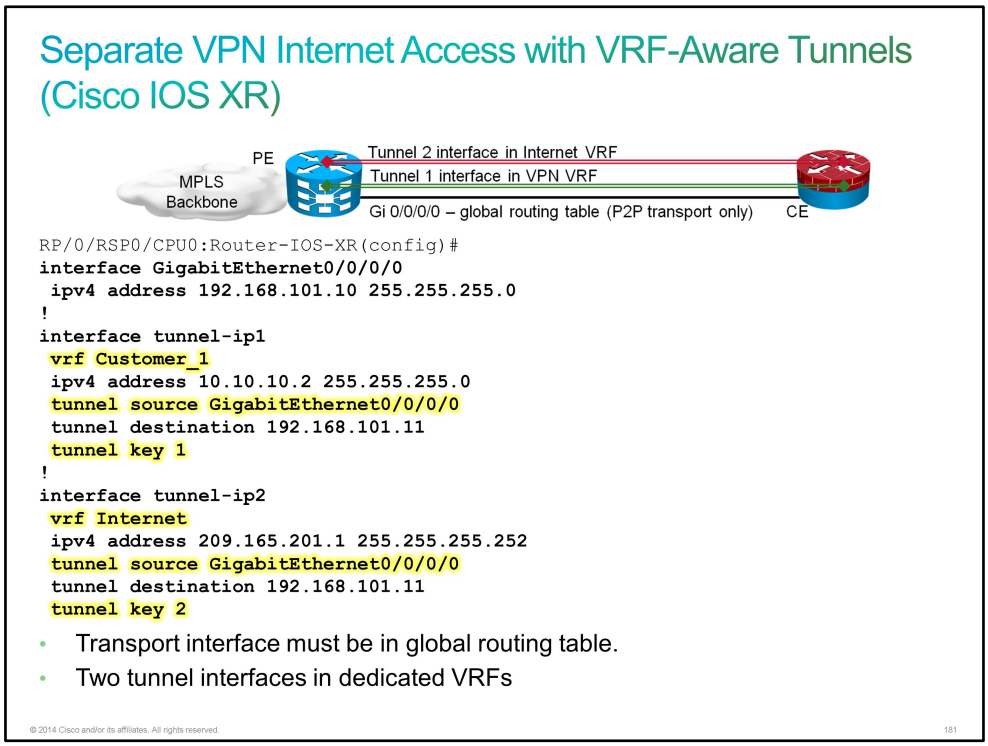
One link between the PE router and each central customer router is assigned to the customer VRF, and one link is assigned to the Internet VRF.

In this example, PE1 connects the Customer A router to both, the Customer A VRF and the Internet VRF.

https://t.me/learningnets

Separate VPN Internet Access with VRF-Aware Tunnels (Cisco IOS XR)

Instead of terminating separate links or subinterfaces in two different VRFs, you can use GRE tunnels.



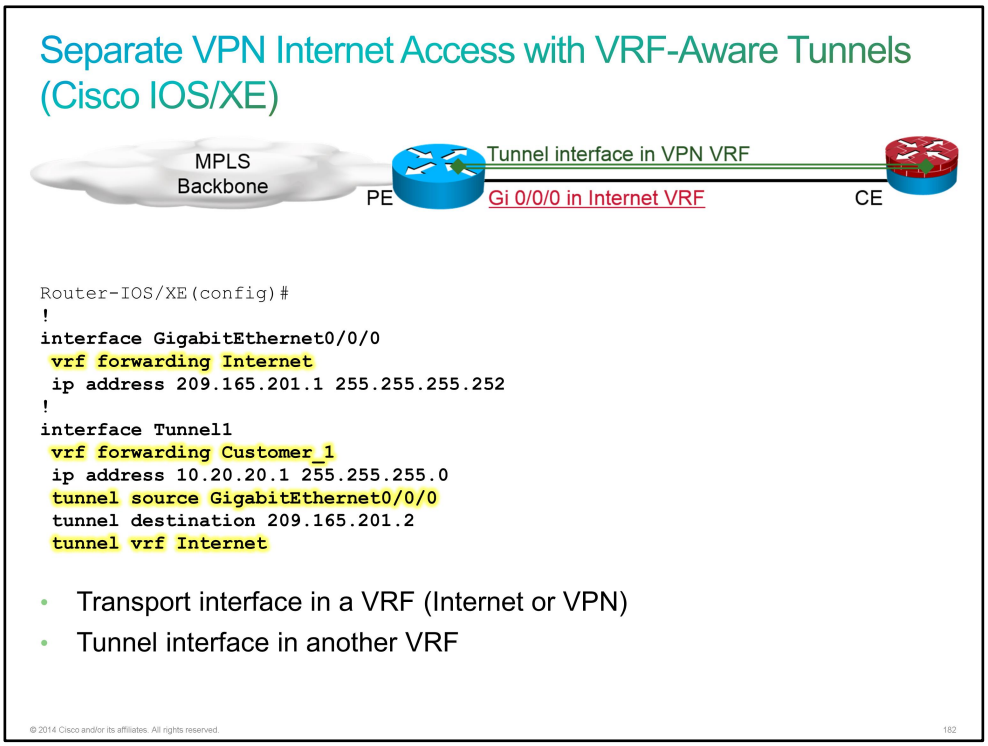
In Cisco IOS XR, you can terminate a tunnel in a VRF but the tunnel transport must go through the global routing space.

To implement Internet access through a separate VPN using VRF-aware tunnels on Cisco IOS XR, two tunnels are needed. The physical interface is assigned to the global routing space and provides transport for both tunnels. One tunnel interface has been assigned to the customer VRF. The other tunnel interface has been assigned to the Internet VRF and will ensure Internet connectivity. GRE tunnel keys, configured with the **tunnel key** command, are used to match the tunnels that are both established between the same source and destination IP address.

https://t.me/learningnets

Separate VPN Internet Access with VRF-Aware Tunnels (Cisco IOS/XE)

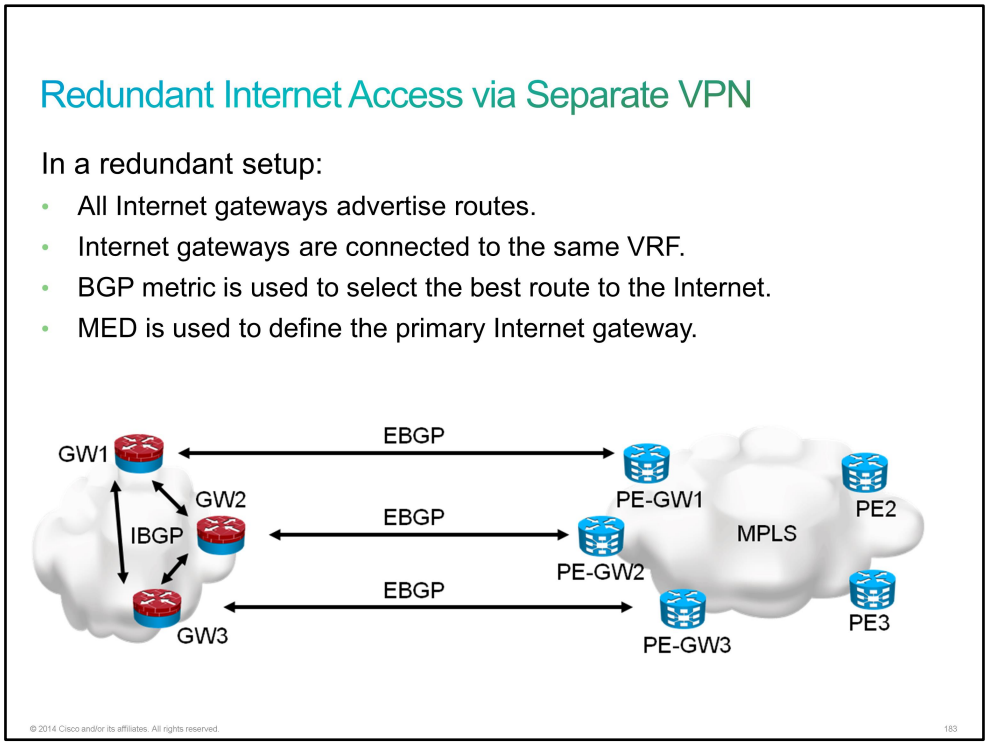
To implement Internet access through a separate VPN on Cisco IOS/XE, only one tunnel is needed.



In this example, the physical interface has been assigned to the Internet VRF and will provide Internet connectivity. The tunnel interface terminates in the customer VRF and will carry VPN traffic.

https://t.me/learningnets

Redundant Internet Access via Separate VPN

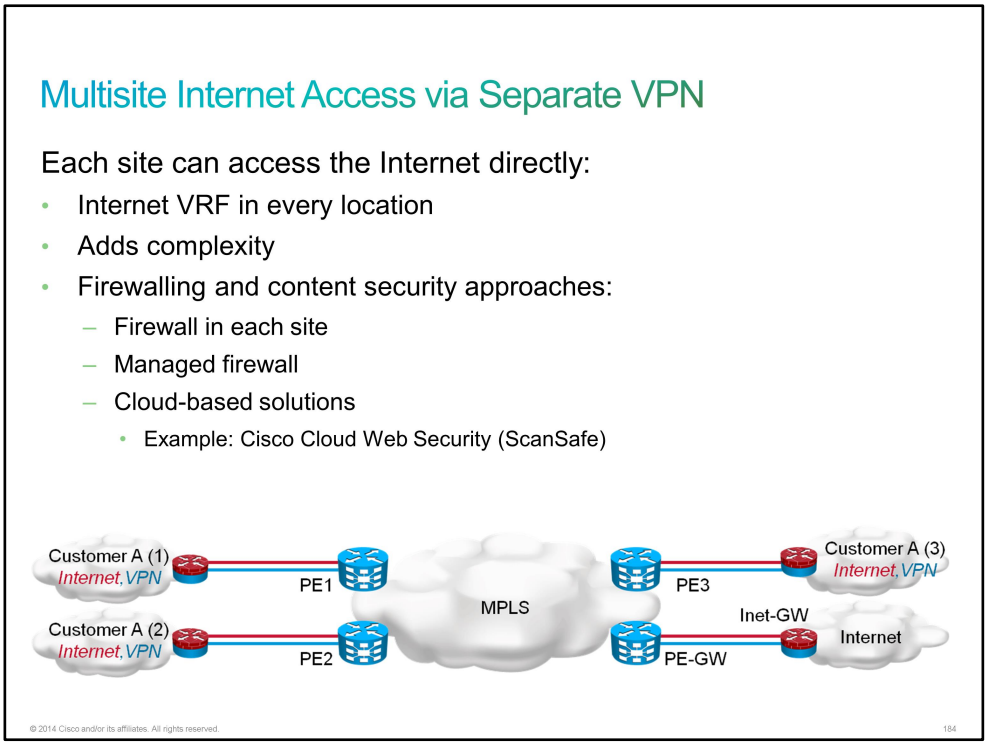


Redundant Internet access is easy to achieve when the Internet service is implemented as a VPN in the MPLS VPN backbone:

- Multiple Internet gateways (acting as CE routers) need to be connected to the MPLS VPN backbone to ensure router and link redundancy.
- All Internet gateways advertise the default route to the PE routers, resulting in routing redundancy.
- The Internet gateways also announce local Internet routes. Because these routes are announced with different BGP attributes—most notably MED—the PE routers select the proper Internet gateway router as the exit point toward those destinations.
- The MED attribute can also be used to indicate the preferred default route to the PE routers. In this setup, one Internet gateway router acts as a primary Internet gateway, and the other Internet gateway router acts as a backup.
- The redundancy that has been established so far covers the path between customer sites and the Internet gateway routers. A failure in the Internet backbone might break the Internet connectivity for the customers if the Internet gateway routers announce the default route unconditionally. Conditional advertisement of the default route is therefore configured on the Internet gateway routers that announce the default route to the PE routers only if the Internet gateway routers can reach an upstream destination.

https://t.me/learningnets

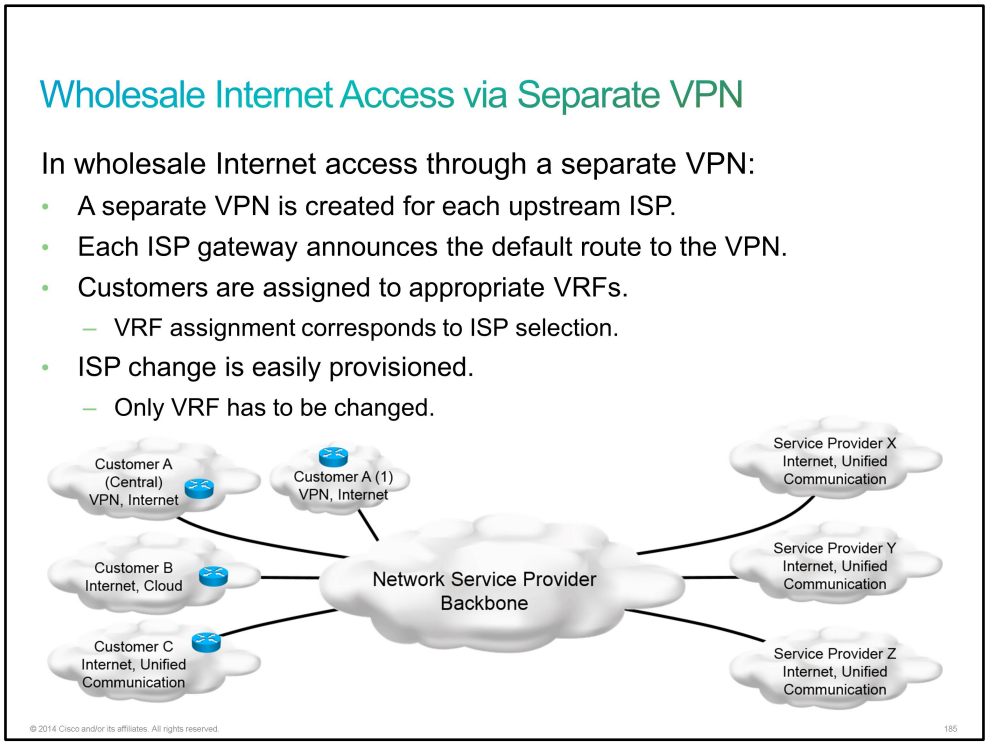
Multisite Internet Access via Separate VPN



Multisite customer Internet access can be implemented by configuring the Internet VRF at every location. This solution adds complexity for the customer because firewall and NAT support might be needed at every site, unless the service provider offers a centrally managed firewall service.

https://t.me/learningnets

Wholesale Internet Access via Separate VPN

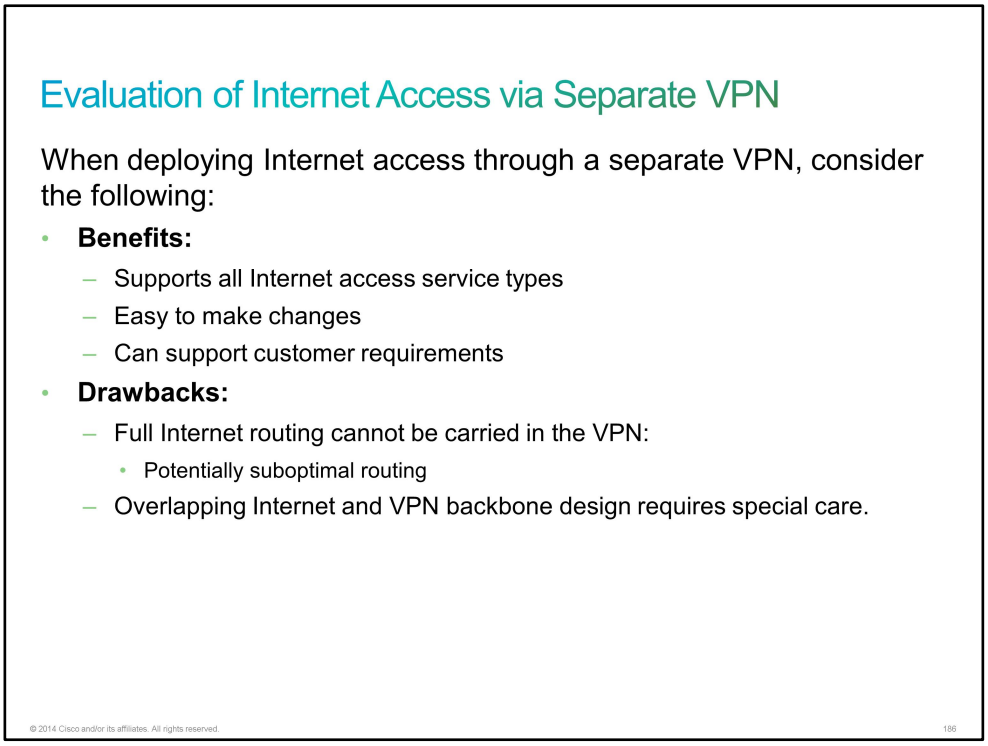


Wholesale Internet access is implemented by creating a separate VPN for every upstream ISP. Acting as a CE router toward the MPLS VPN-based Internet access backbone, the Internet gateway of the upstream ISP announces a default route that is used for routing inside the VPN.

Customers are tied to upstream service providers simply by placing the PE-CE link into the VRF that is associated with the upstream service provider. Changing an ISP becomes as easy as reassigning the interface into a different VRF and managing address-allocation issues.

https://t.me/learningnets

Evaluation of Internet Access via Separate VPN



Internet access that is implemented as a separate VPN has the following benefits:

- This design model supports all Internet access services, ranging from traditional Internet access to innovative services such as wholesale Internet access.
- This design also supports all customer requirements, including full Internet routing on customer routers through an EBGP multihop session with the Internet gateway.

Internet access that is implemented as a separate VPN has the following drawbacks:

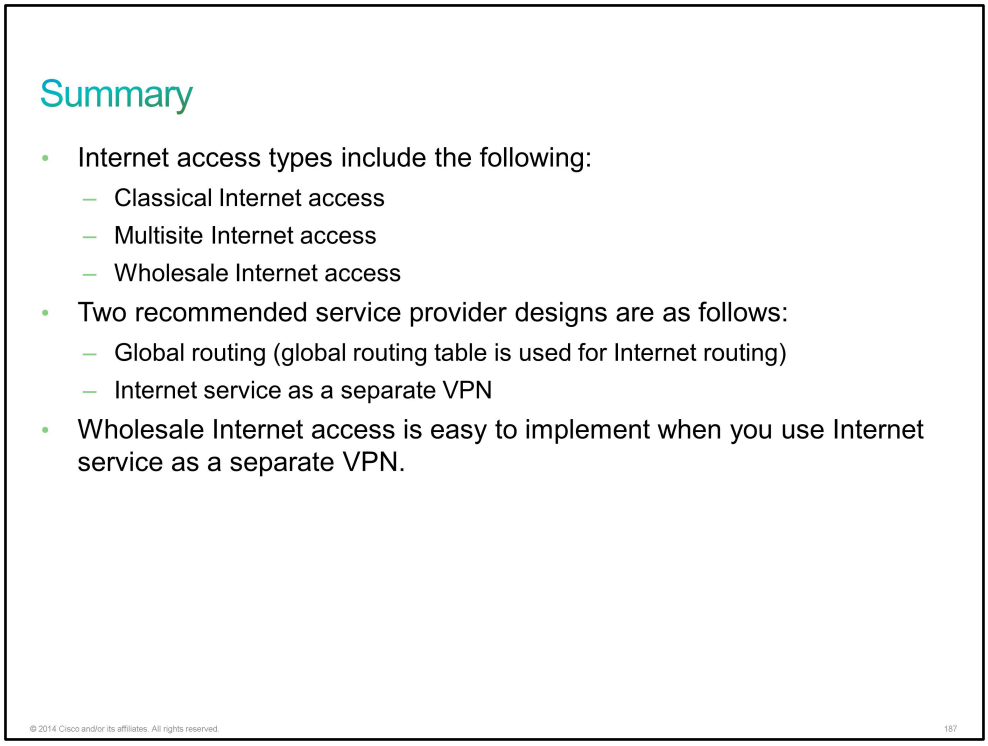
- Full Internet routing cannot be carried inside a VPN; therefore, default routing toward the Internet gateways needs to be used, potentially resulting in suboptimal routing.
- The Internet backbone gateway router is positioned as a CE router connected to the MPLS VPN backbone. If the service provider runs Internet service and MPLS VPN service on the same set of routers, the interconnection between the two services requires special considerations.

  The benefits of the separate VPN design far outweigh the limitations.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

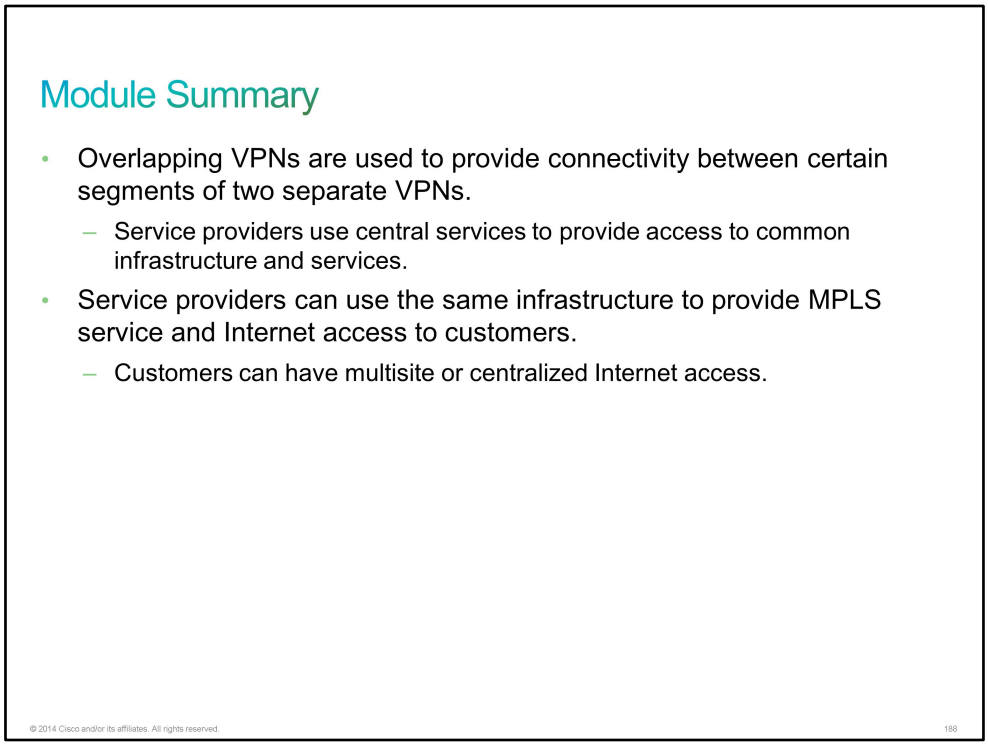
 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Module Summary

Overview

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



https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: 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 the two major design models for implementing Internet access and MPLS VPNs? (Choose two.) (Source: Implementing Internet Access in MPLS VPNs)

   A. using a separate VPN for Internet access
   B. wholesale Internet access
   C. using global routing on PE routers
   D. using default routes

2. What is the recommended implementation option for using global routing to provide Internet access? (Source: Implementing Internet Access in MPLS VPNs)

   A. The separate subinterfaces are placed in separate VRFs.
   B. The Internet (sub)interface is placed in the global routing space.
   C. One interface for the Internet and the MPLS VPN is recommended.
   D. Static routes are always needed for Internet access.

3. Which two are drawbacks of Internet access that is implemented as a separate VPN? (Choose two.) (Source: Implementing Internet Access in MPLS VPNs)

   A. Full Internet routing cannot be carried inside a VRF.
   B. The customer cannot receive full Internet routes.
   C. Default routing toward the Internet gateways need to be used, potentially resulting in suboptimal routing.
   D. The customer must use default routes.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. A, C
2. B
3. A, C

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 4: MPLS IPv6 and Interdomain Solutions

Introduction

This module discusses some advanced configuration features that can help increase the stability of the MPLS VPN backbone. The first lesson discusses various MPLS VPN features that a service provider might use to help meet service requirements and looks at various types of VPN solutions and topologies.

Integrating Internet access with an MPLS VPN solution is one of the most common service provider business requirements. The second lesson provides a good understanding of underlying design issues, several potential design scenarios, and some sample configurations. Various topologies and implementation methods are discussed, along with ways to separate Internet access from VPN services.

Deployments of MPLS have become routine in large-scale global networks, which demand solutions to complex business and network problems. The third lesson describes the two primary components of the Cisco IOS MPLS Inter-Domain Solution: inter-AS and CSC.

Upon completing this module, you will be able to:

- Describe the requirements, usage, and solutions that are associated with overlapping, central services, and managed CE router service VPNs
- Describe common customer Internet connectivity scenarios and identify design models for combining Internet access with MPLS Layer 3 VPN services
- Introduce CSC and inter-AS features with different implementation options

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Deploying IPv6 in an MPLS Environment

Overview

IPv6 is being rolled out on top of or instead of IPv4 due to IPv4 restrictions, most notably the limited global address space. There are multiple methods to deploy IPv6 in MPLS environments. This lesson introduces some concepts that can be used when you implement IPv6 in a service provider environment.

Upon completing this lesson, you will be able to:

- Identify various IPv6 and IPv4 interoperability solutions
- List the various methods that can be used to deploy IPv6 over MPLS
- Describe using IPv6-in-IPv4 tunnels on the CE routers to support IPv6
- Describe using Layer 2 MPLS VPNs to support IPv6
- Describe the 6PE architecture
- Describe the 6PE control plane
- Describe the 6PE data plane
- Describe how to configure 6PE
- Describe how to verify 6VPE
- Describe the 6VPE architecture
- Describe the 6VPE topology
- Describe 6VPE forwarding
- Describe the 6VPE configuration procedure
- Describe the 6VPE configuration
- Describe how to verify 6VPE on a PE router

https://t.me/learningnets

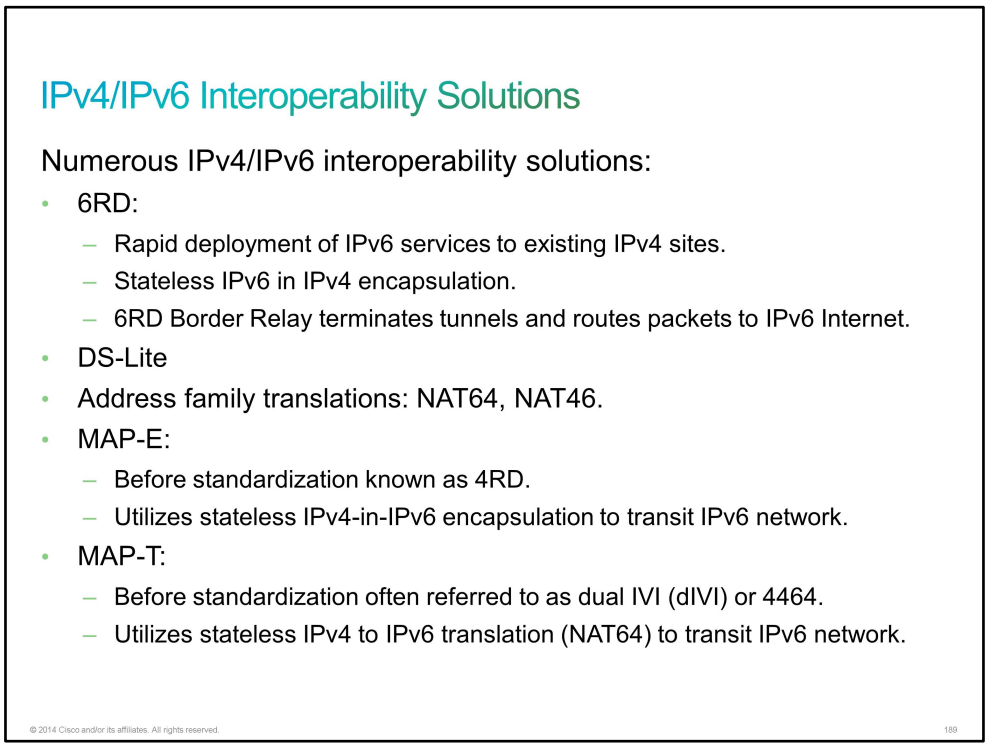
- Describe how to verify 6VPE on a CE router

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

IPv4/IPv6 Interoperability Solutions

A successful IPv6 roll-out depends on the smooth integration and coexistence of IPv6 and IPv4 in service provider environments.



A plethora of technologies are available for service providers to move to IPv6. These technologies use the established techniques of address translation, tunneling, and encapsulation. Some examples include:

- **6RD:** 6RD enables a service provider to rapidly deploy IPv6 services to existing IPv4 sites to which it provides CPE. This approach utilizes stateless IPv6 in IPv4 encapsulation in order to transit IPv4-only network infrastructure. The encapsulation (also known as softwires) must be supported by the CPE, while the 6rd Border Relay must support tunnel termination to route packets to Internet hosts on IPv6. The provider access network continues to be on IPv4, while customers see IPv6 and IPv4 service simultaneously.
- **DS-Lite:** DS-Lite is a combination of tunnel and translation technologies. With DS-Lite, at least part of the service provider network (e.g., access, aggregation network) only supports IPv6 routing. The CPE is provisioned only and natively with IPv6. Any IPv4 traffic on its local LAN is tunneled by the CPE over the IPv6 infrastructure to the gateway. The encapsulation (aka softwires) must be supported by the CPE. The IPv4 address space of the subscriber is a private one. The gateway terminates the tunnel and translates the IPv4 local addressing into globally routable IPv4 (NAT44). If the subscriber network has the capability of using IPv6, the IPv6 traffic is routed natively through the service provider infrastructure. There is a single IPv4 NAT operation applied in the service provider network to the subscriber traffic
- **AFT:** AFT refers to the translation of one IP address from one address family into another IP address of another address family; for instance, from one IPv4 address into an IPv6 address or vice versa. This translation is sometimes denoted as NAT46 (when the initiator is on the IPv4 side) or NAT64 (when the initiator is on the IPv6 side). AFT can be stateful or stateless.

https://t.me/learningnets

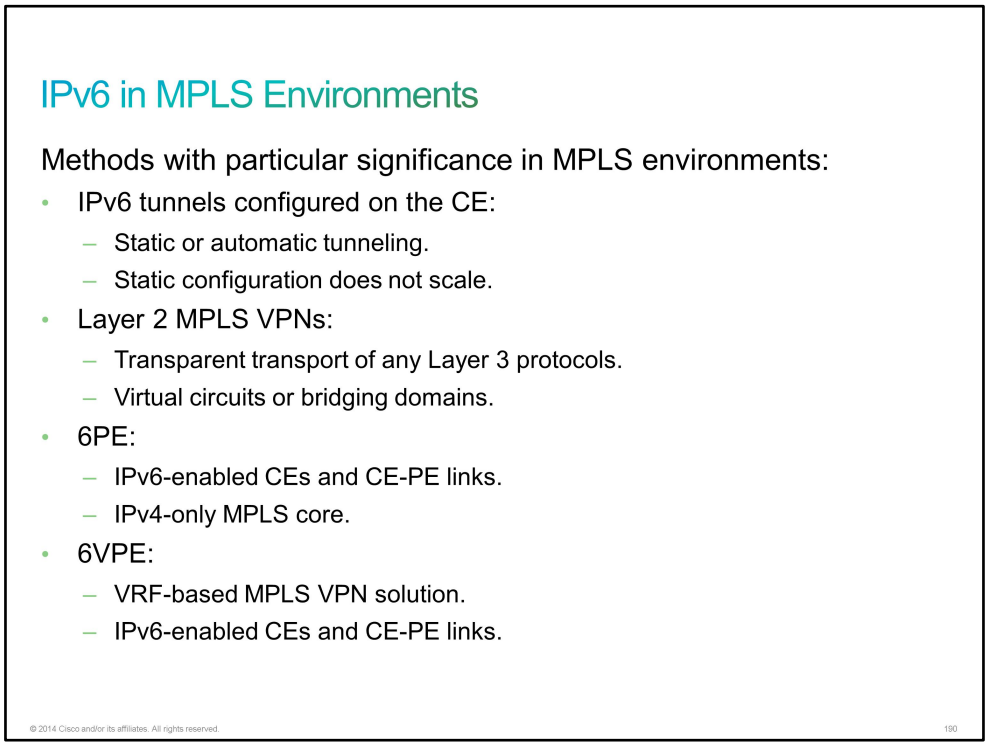
- **MAP-E:** MAP-E is a combination of tunnel and translation technologies. Prior to being standardized, MAP-E was commonly referred to as IPv4 Residual Deployment (4rd). MAP-E enables a service provider to rapidly allow IPv4 services to IPv6 (customer) sites to which it provides CPE. This approach utilizes stateless IPv4-in-IPv6 encapsulation (i.e. tunneling) to transit IPv6-enabled network infrastructure. The encapsulation must be supported by the CPE and MAP-E Gateway/Border Relay, which removes the IPv6 encapsulation from IPv4 packets while forwarding them to the Internet. The provider access network can now be on IPv6, while customers see IPv6 and IPv4 service simultaneously. MAP-E also helps manage IPv4 address exhaustion by keeping the stateful NAT44 on CPE, as usual. MAP-E is attractive to those service providers who have deployed or are planning to deploy IPv6 end-to-end. It is a superior alternative to DS-Lite, given the stateless tunneling.
- **Mapping of Address and Port-Translation Mode:** MAP-T employs a combination of translation technologies. Prior to being standardized, MAP-T was often referred to as dIVI or 4464. enables a service provider to offer IPv4 services to IPv6 enabled (customer) sites to which it provides CPE. This approach utilizes stateless IPv4 to IPv6 translation (i.e. NAT64) to transit IPv6-enabled network infrastructure. The provider access network can now be on IPv6, while customers see IPv6 and IPv4 services simultaneously. MAP-T keeps the stateful NAT44 on CPE, as usual, to handle IPv4 address exhaustion, in addition to stateless NAT64 on CPE and Border Router. MAP-T is attractive to those service providers who have deployed or are planning to deploy IPv6 end-to-end and want to manage IPv4 address exhaustion with utmost predictability. MAP-T is a superior alternative to DS-Lite, given no tunneling and no CGN/44 in the service provider network.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 in MPLS Environments

IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 network.



Depending on the deployment scheme that is chosen and on the current MPLS implementation, enabling IPv6 services might require fewer backbone infrastructure upgrades and less reconfiguration of core routers than enabling IPv4 services, because forwarding is based on labels rather than on the IP header, thus providing a very cost-effective strategy for the deployment of IPv6.

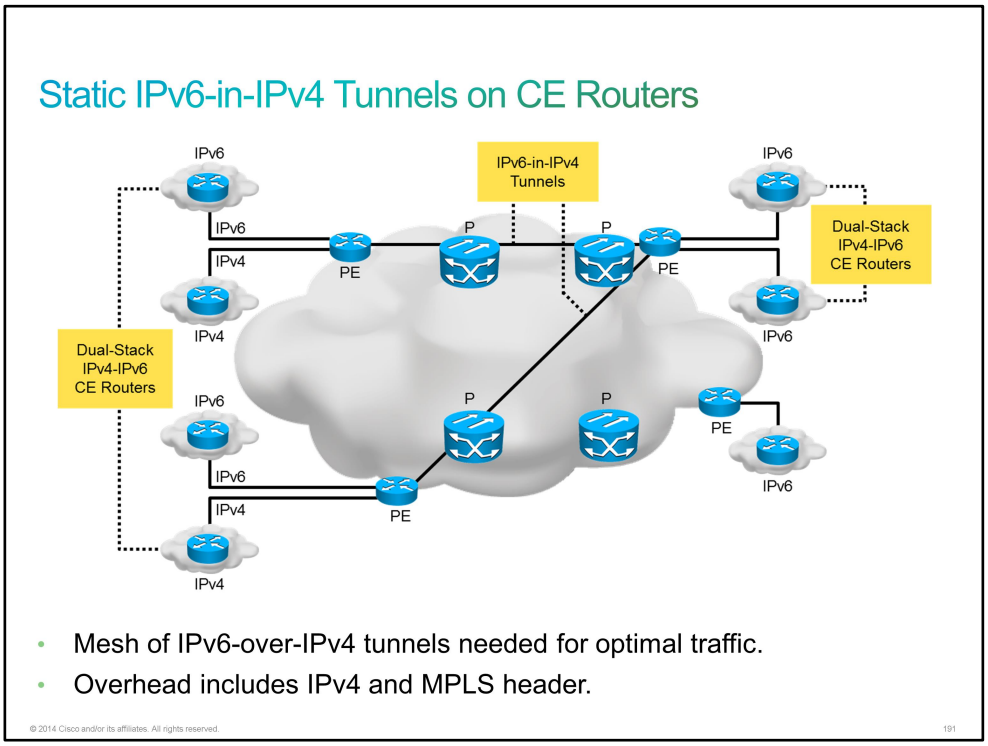
Additionally, the inherent VPN and traffic engineering services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment. A variety of deployment strategies are available or under development:

- **Deploying IPv6 by using tunnels on the CE routers:** This strategy has no impact on and requires no changes to the MPLS provider P routers or PE routers. The strategy uses IPv4 tunnels that are terminated on the CE routers to encapsulate the IPv6 traffic, which then appears as IPv4 traffic within the MPLS provider network.
- **Deploying IPv6 over a Layer 2 MPLS VPN:** This strategy, applicable only on specific Cisco routers (such as the Cisco 12000 and 7600 Series Routers), also requires no changes to the core routing mechanisms, assuming that the provider is already capable of supporting Layer 2 MPLS VPNs.
- **Deploying IPv6 on the PE routers via 6PE:** This strategy requires changes to the PE routers to support a dual-stack implementation, but all the core functions remain IPv4-only, and so the P routers require no changes.

https://t.me/learningnets

Static IPv6-in-IPv4 Tunnels on CE Routers

Using static IPv6-in-IPv4 tunnels on the CE routers is the simplest way to deploy IPv6 over MPLS networks.



This method has no impact on the operation or infrastructure of MPLS and requires no changes to the P routers in the core or to the PE routers that connect to customers.

IPv6 tunnels that are configured on a CE have these characteristics:

- The CE routers are IPv6-aware (dual stack).
- A mesh of IPv6-over-IPv4 tunnels connects CE to CE.
- Overhead includes the IPv4 header and MPLS header.
- MPLS and VPN support native IPv4 and IPv6 tunnels.

Communication between the remote IPv6 domains uses standard tunneling mechanisms, running IPv6 over IPv4 tunnels, as the way that MPLS VPNs support native IPv4 tunnels. The CE routers need to be upgraded to dual stack and configured by using manually configured tunnels or 6to4 tunnels. However, communication with the PE routers is IPv4-only, and the traffic appears to be IPv4 to the MPLS domain. The dual-stack routers use the 6to4 tunnel addresses or an IPv6 prefix that is assigned by a virtual IPv6 ISP (possibly other than the MPLS provider). The figure shows an example for deployment of IPv6 by using tunnels on the CE routers.

https://t.me/learningnets

This alternative is attractive because it has no impact on the MPLS P or PE routers, because it uses IPv4 tunnels to encapsulate IPv6 traffic. After being encapsulated, the traffic appears as IPv4 traffic within the network. However, this option does not permit the delegation of a /48 prefix address from the MPLS service provider address space to customers. Nonetheless, it is probably useful for specific applications, such as adding IPv6 services to an MPLS VPN by configuring the tunnel on the CE. All Cisco IOS Software tunneling mechanisms support this capability.

Pros

From the provider standpoint, using tunnels on CE routers is the simplest way to deploy IPv6 over MPLS networks. This approach has no impact on the operation or infrastructure of MPLS and requires no changes to either the P routers in the core or the PE routers that connect to customers.

Cons

Using tunnels on CE routers presents the same problems as any configured tunnel scheme. As the number of tunnel partners grows, the number of tunnels in an any-to-any mesh increases quickly, making management of the tunnels error-prone and expensive. This problem is not specific to MPLS but rather is a general issue with manually configured tunnels. Automatic tunnels (such as 6to4) can be used to reduce the effort that is required to manage tunnels on the CE, but automatic tunnels have other issues that must be considered.

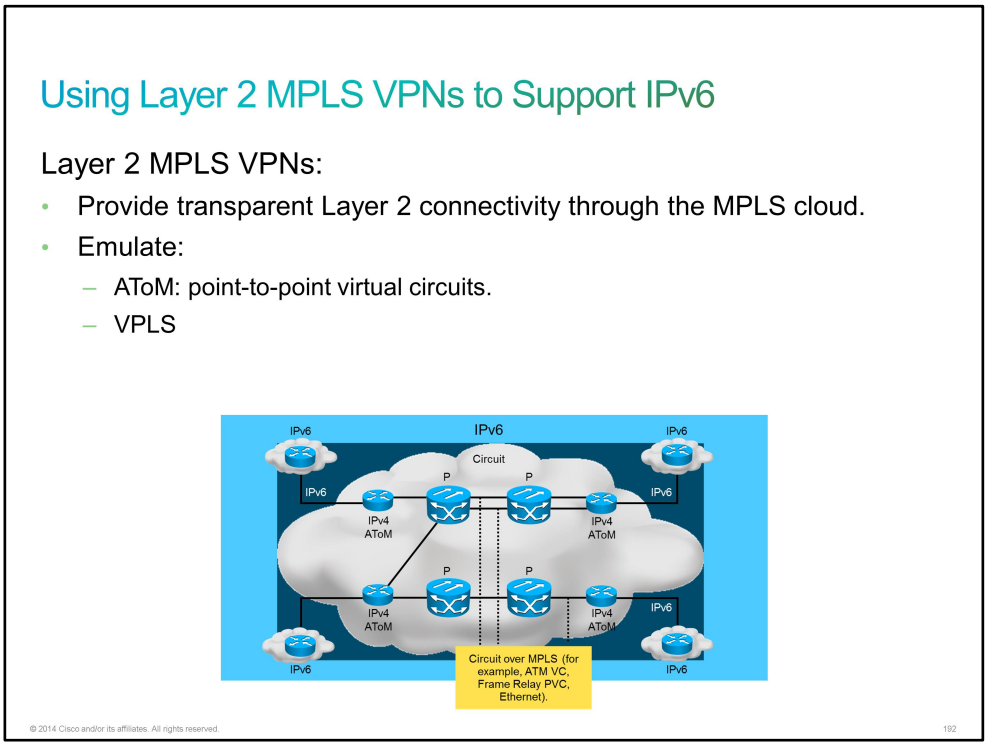
Delegating a global IPv6 prefix for an ISP is also difficult because the ISP does not furnish the IPv6 connectivity; it merely provides transport. Even when the ISP does allocate an address, early traffic levels will be modest, and later traffic might shift from IPv4 to IPv6. Therefore, there is not necessarily a net increase in carried packets. Because the service provider provides transport only over IPv4 over MPLS, this scenario offers little possibility of generating additional revenue for the provider.

From the provider standpoint, using tunnels on CE routers is not a good mechanism to carry IPv6 traffic across an IPv4-based MPLS network. The provider basically puts responsibility for IPv6 transit on the end customers, because the MPLS provider recognizes only IPv4 traffic.

https://t.me/learningnets

Using Layer 2 MPLS VPNs to Support IPv6

Using AToM to deploy IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS.



Circuit transport requires no changes to either the P routers in the core or the PE routers that connect to customers, assuming that the provider has already implemented circuit-over-MPLS technology. This design is a Layer 2 solution.

IPv6 using a Layer 2 MPLS VPN has these characteristics:

- IPv6 is run over pseudowires or in virtual bridge domains.
- Edge MPLS routers need to support a Layer 2 MPLS VPN.
- A mesh of virtual circuits connects PE to PE.
- The PE routers are regular IPv4 routers that are used to aggregate customer IPv6 routers.

Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link, in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled by using a Layer 2 MPLS VPN or EoMPLS, with the IPv6 routers connecting through an ATM or Ethernet interface, respectively. The PE will require an upgrade to a Layer 2 MPLS VPN if it does not already contain that support. The figure shows an example of IPv6 deployment over any circuit transport over MPLS.

The alternative of IPv6 over Layer 2 MPLS VPN is likely to be used by service providers that have ATM or Ethernet links to CE routers. This alternative is fully transparent to users. However, like the previous option, this option does not permit the delegation of a /48 prefix address from the service provider address space to customers. In addition, IPv6 over Layer 2 MPLS VPN does not scale as well as 6PE because service providers need to create circuits through their MPLS networks for their customers.

https://t.me/learningnets

Pros

There are some advantages to using IPv6 over a Layer 2 MPLS VPN:

- Simple solution for service providers with ATM or Ethernet links to CE routers and circuit over MPLS capability.
- Fully transparent IPv6 communication, from the provider perspective.
- No required reconfiguration of the MPLS core.

Cons

There are also some disadvantages to using IPv6 over a Layer 2 MPLS VPN:

- Equivalent to Layer 2 VPN complexity.
- The need to upgrade PE routers that connect to customers to a Layer 2 MPLS VPN.
- The difficulty of delegating a global IPv6 prefix for an ISP.

Description of Pros and Cons

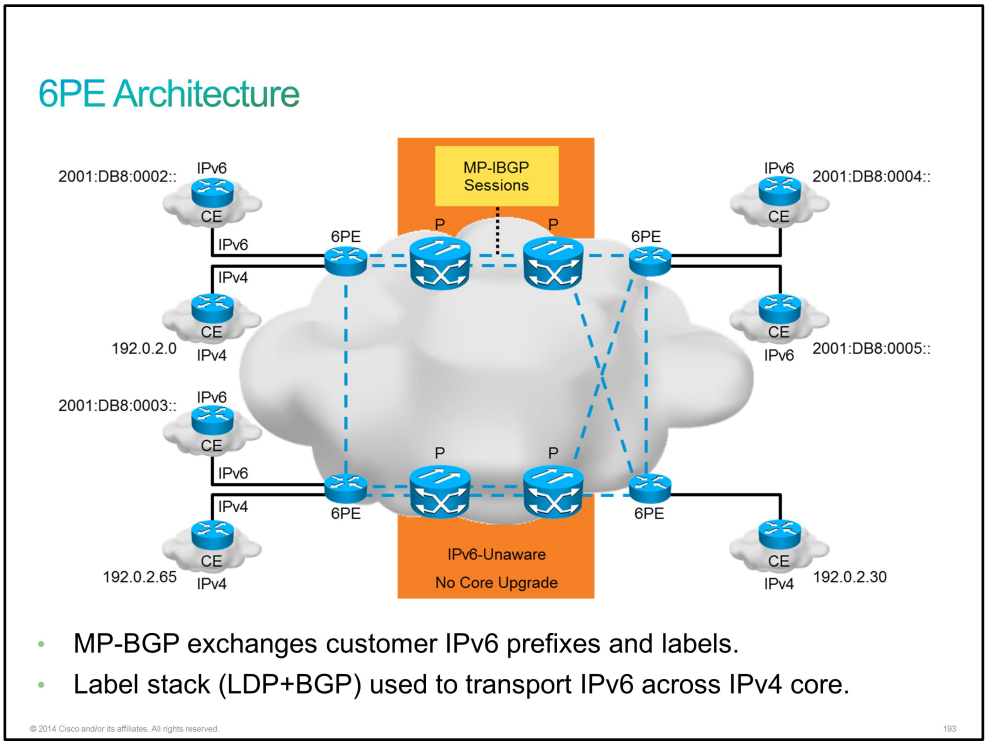
Additional points about IPv6 over Layer 2 MPLS VPN are as follows:

- The use of a Layer 2 MPLS VPN to deploy IPv6 over MPLS networks does not affect the operation or infrastructure of MPLS. Neither the P router software version nor the configuration is changed. However, PE routers that connect to customers must be upgraded to Layer 2 MPLS VPN. Delegating a global IPv6 prefix for an ISP is also difficult.
- Communication between remote IPv6 domains runs native IPv6 protocols over a dedicated link in which the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled via a Layer 2 MPLS VPN, with the IPv6 routers connected through an ATM or Ethernet interface. This potentially heavy tunnel meshing is a drawback of the solution.

https://t.me/learningnets

6PE Architecture

Cisco 6PE6PEenables IPv6 islands to communicate with each other over an MPLS IPv4 core network by using MPLS LSPs. The method relies on BGP extensions in the IPv4 network PE routers (Cisco 6PE routers) to exchange IPv6 reachability information and an MPLS label for each IPv6 address prefix that is announced. Cisco 6PE routers are dual stack (IPv4 and IPv6). This design is a Layer 3 solution.



For the IPv6 transport to be transparent to all but Cisco 6PE routers, you must impose a hierarchy of labels at the Cisco 6PE ingress router. The top label provides connectivity inside the IPv4 MPLS core network. LDP distributes this label. The inner label is used for IPv6 forwarding at each Cisco 6PE egress router. This label is distributed by MP-BGP in the VPN-IPv6 address family.

Cisco 6PE uses MPLS to minimize changes in the ISP core network.

With Cisco 6PE, IPv6 forwarding in the core of the MPLS network is done by label switching, eliminating the need for either IPv6-over-IPv4 tunnels or an additional Layer 2 encapsulation. It allows the appearance of a native IPv6 service to be offered across the network.

Each PE router that must support IPv6 connectivity needs to be upgraded to dual stack (becoming a Cisco 6PE router) and configured to run MPLS on the interfaces that connect to the core. Depending on the site requirements, each router can be configured to forward IPv6 or IPv6 and IPv4 traffic on the interfaces to the CE routers, thus providing the ability to offer only native IPv6 or both IPv6 and native IPv4 services. The Cisco 6PE router exchanges either IPv4 or IPv6 routing information with the CE router through any of the supported routing protocols (depending on the connection). It switches IPv4 and IPv6 traffic over the native IPv4 and IPv6 interfaces that do not run MPLS.

The Cisco 6PE router uses MP-BGP to exchange reachability information with the other 6PE routers in the MPLS domain. It shares a common IPv4 routing protocol (such as OSPFor IS-IS) with the other P and PE devices in the domain. IGP is used for deriving next-hop information.

https://t.me/learningnets

The Cisco 6PE routers encapsulate IPv6 traffic by using two levels of MPLS labels. The top label is distributed by an LDP or TDP that devices in the core use to carry the packet to the egress Cisco 6PE router, using IPv4 routing information. The second or bottom label is associated with the IPv6 prefix of the destination through MP-BGPv4.

Multiprotocol extensions for BGP4 are characterized as follows:

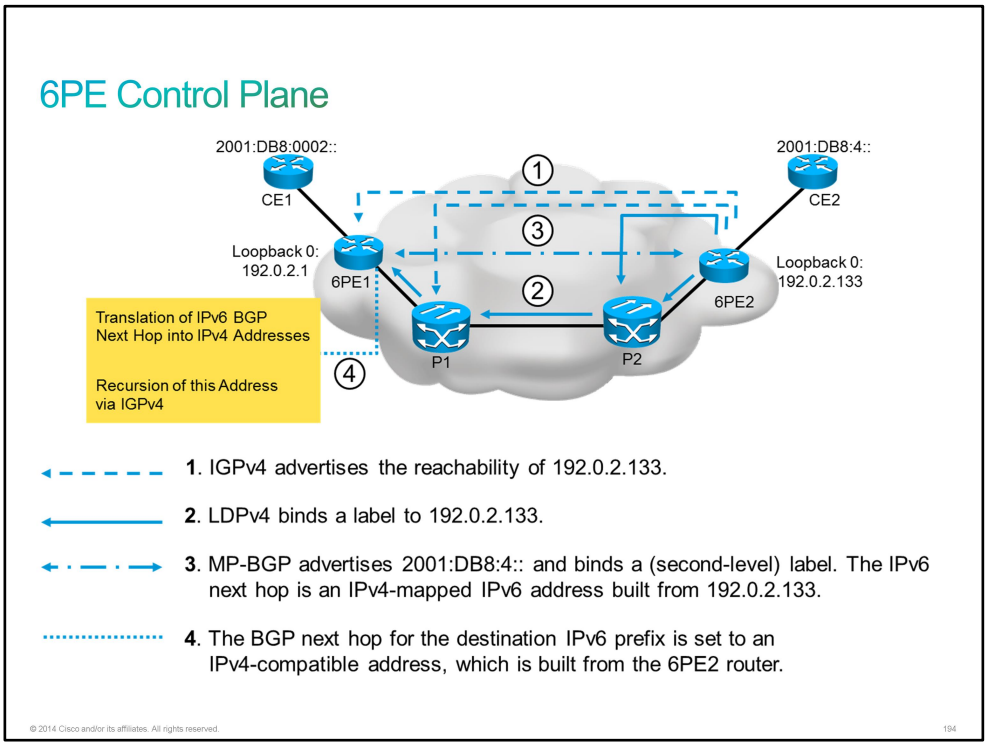
- BGP4 is the primary mechanism to exchange IPv6 reachability information for Cisco 6PE and 6VPE.
- Label binding information is piggybacked with the prefix advertisement in a NLRI attribute.
- BGP protocol extensions, which carry routing information about other protocols, include multicast, MPLS, IPv6, VPN-IPv4, labeled IPv6 unicast (Cisco 6PE), and Cisco 6VPE.
- The exchange of multiprotocol NLRI must be negotiated at the session setup. The exchange uses BGP capability advertisement negotiation procedures.

MP-BGP is the supported EGP for IPv6. MP-BGP extensions for IPv6 support the same features and functionality as IPv4 BGP. IPv6 enhancements to MP-BGP include support for an IPv6 address family, NLRI, and next-hop attributes that use IPv6 addresses.

https://t.me/learningnets

6PE Control Plane

This example describes how IPv6 transport is achieved across an IPv4-based MPLS provider network that is running an IPv4 IGP such as OSPFv2.



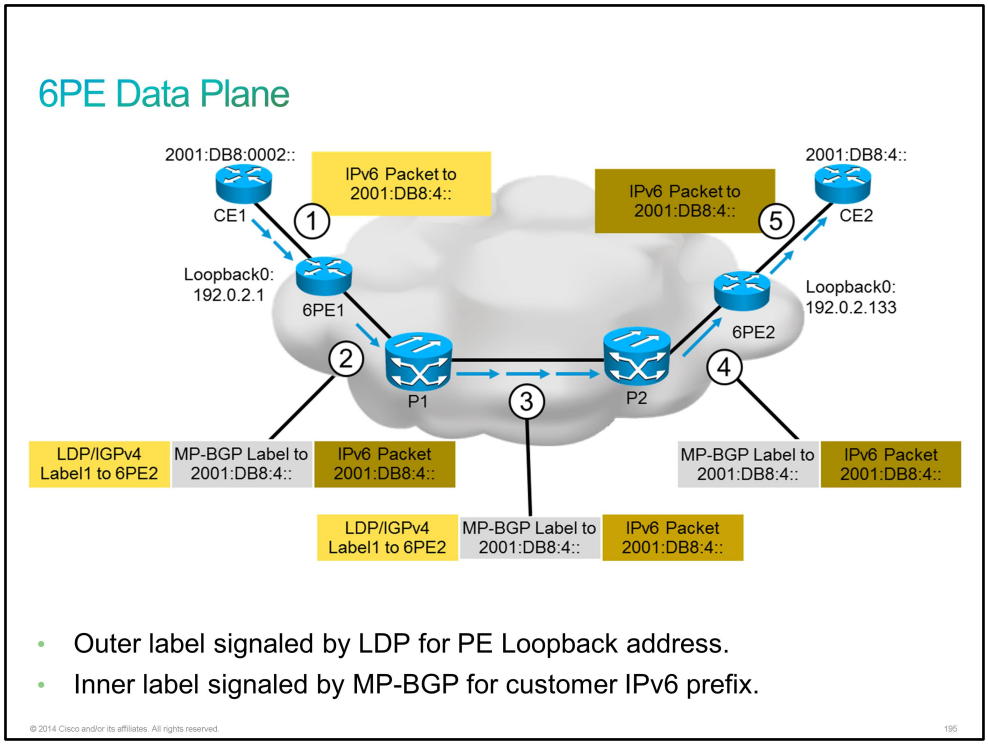
This process describes how the reachability is established between Cisco 6PEs:

1. IGPv4 advertises the reachability of PE routers. The Cisco 6PE routers exchange IPv6 routing information through MP-BGP. The Cisco 6PE routers peer together through MP-BGPv4to exchange IPv6 reachability with the MPLS cloud and to perform IPv6 forwarding. IGP advertises the reachability of the BGP next-hop address; that is, 192.0.2.133.
2. LDPv4 binds labels to create LSPs.
3. The 6PE2 router sends MP-BGP advertisements to the ingress PE router, 6PE1.
4. The BGP next hop for the destination IPv6 prefix is set to an IPv4-compatible address, which is built from the 6PE2 router. The MPLS label binds to the IPv4- or IPv6-mapped address, which is formed from the IPv4 address of 6PE2 (192.0.2.133). Thus, other Cisco 6PE routers (which are part of the mesh edges around the MPLS core) see 6PE2::FFFF:192.0.2.133 as the next-hop IPv6 address when a packet needs to be forwarded to the 2001:DB8:4:: network.

https://t.me/learningnets

6PE Data Plane

6PE packet forwarding is based on label switching using a label stack.



This example describes how routers forward packets in the data plane:

1. An IPv6 packet arrives at the ingress Cisco 6PE router. An IPv6 packet that is destined to a network with prefix 2001:DB8:4:: enters through 6PE1. The MPLS core must be used to carry the packet to its destination, via 6PE2. The routing configuration via IBGP has already advertised the next-hop IPv6 address (::FFFF:192.0.2.133) to 6PE1.
2. The ingress Cisco 6PE router binds MP-BGP (inner) and LDP-IGPv4 (top, or outer) labels to the packet. 6PE1 does a lookup on the IPv6 prefix. Two labels are then created and bound to the packet. MP-BGP binds 2001:DB8:4:: and LDP-IGPv4 binds its label to the IPv4 address of 6PE2 (192.0.2.133). In this way, MP-BGP can carry either IPv4 packets or IPv6 packets from one Cisco 6PE router to another.
3. The IPv6 packet is forwarded to the first P router. MPLS label switching is IPv6-unaware. When P1 receives the MPLS packet, it looks only at the LDP-distributed MPLS label, and a lookup is performed. The result is the imposition of Label2, and the MPLS packet is forwarded to P2. The IPv6 traffic is transparent to the core of the MPLS network.
4. The IPv6 packet is forwarded to the next P router. P2 receives an MPLS packet, which performs an MPLS lookup on Label2 and results in the POP label. As a result, the packet is forwarded to the IPv4 address to which it is bound. The packet still has the inner MPLS label, which is required because a raw IPv6 packet is not allowed within the MPLS core.

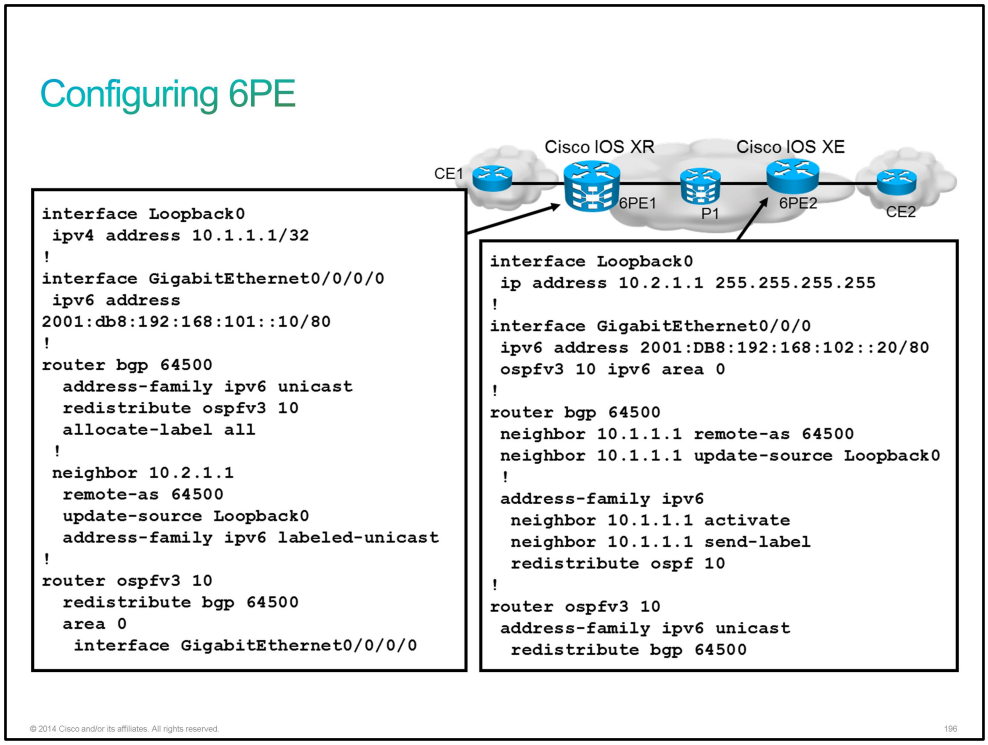
https://t.me/learningnets

5. The egress PE routers remove the MPLS labels and forward the IPv6 packet. The 6PE2 router examines the MPLS packet. The result is that the remaining MPLS label is removed and an IPv6 lookup on the IPv6 destination address is performed, using normal IP destination-based routing. The Cisco 6PE router forwards the IPv6 packet on to the CE device, based on the results of the IPv6 lookup on the IPv6 destination address.

https://t.me/learningnets

Configuring 6PE

In this 6PE scenario, the PE routers run Cisco IOS XR and Cisco IOS XE software.



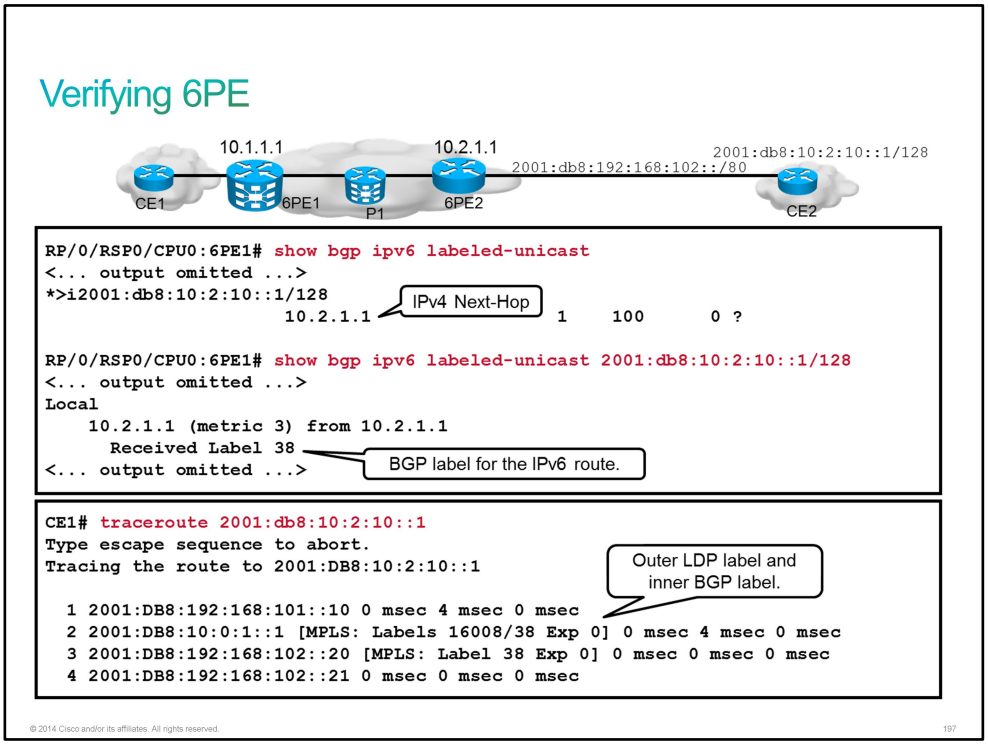
Follow these steps to deploy Cisco 6PE:

1. Configure Cisco 6PE routers as dual-stack routers. The CE-facing interfaces have IPv6 addresses. The core-facing interfaces have IPv4 addresses.
2. Configure Cisco 6PE routers to be part of the IPv4 network and configure MPLS. This configuration is not shown in the example above. The OSPFv3 process 10 on each PE router has been enabled to provide CE-PE IPv6 routing.
3. Configure MP-IBGP. The MP-IBGP session is originated from the Loopback 0 IPv4 address of each PE router.
4. Enable MP-BGP label exchange for IPv6 networks. In Cisco IOS XR you need to enable **address-family ipv6 labeled-unicast** in the BGP neighbor configuration mode. You have to configure label allocation policy. In this example, 6PE1 will allocate labels to all IPv6 routes as a result of the **allocate-label all** command. In Cisco IOS and Cisco IOS XE you need to enable label signaling for a BGP neighbor using the **neighbor** *neighbor-id* **send-label** command.

https://t.me/learningnets

Verifying 6PE

This example illustrates how to verify the 6PE operation.



Once MPLS in functional in the backbone, you can monitor the 6PE operation by examining the IPv6 routes in the PE BGP tables.

In the first output you see that the IPv6 prefixes displayed in the BGP table contain an IPv4 next hop address. This results from the IPv4-only connectivity in the MPLS backbone. If MPLS were not used for traffic forwarding, the IPv6 packets could not be forwarded between the PEs.

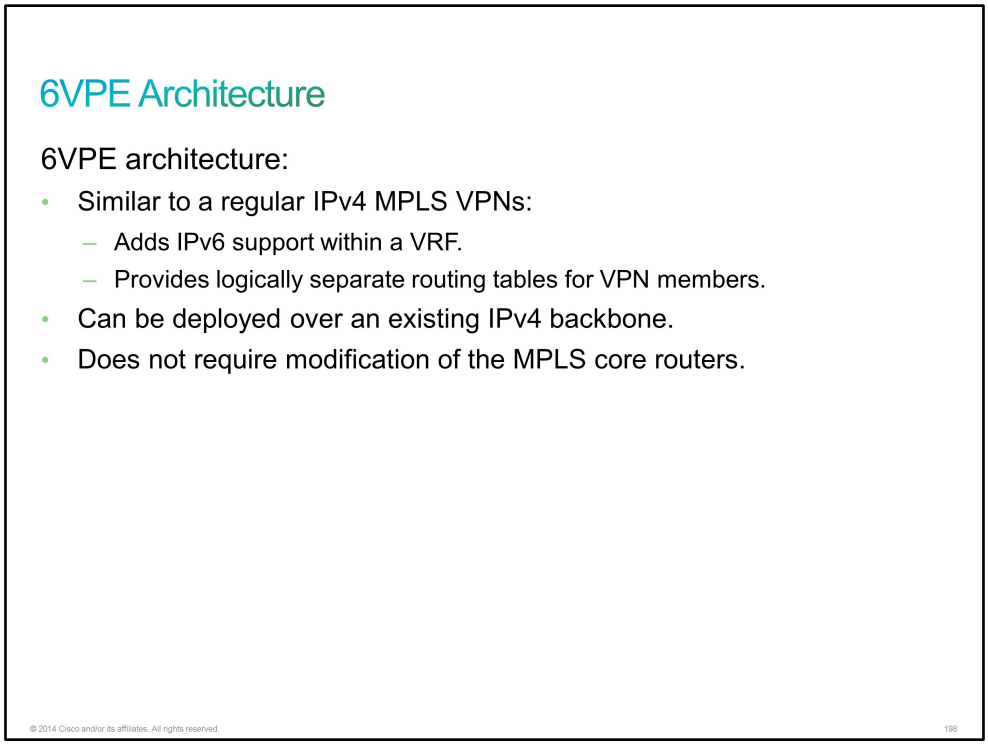
The second output illustrates how to examine a specific IPv6 prefix in the BGP table. It shows a BGP label attached to that prefix. The VPN label, 38 in this example, is used as the inner label when customer traffic is transported through the MPLS backbone.

You can monitor the label swapping operation by using the traceroute utility, as shown in the last output on CE1. When you verify connectivity to CE2, you see that a label stack is used between 6PE1 and the core router. The inner label is the BGP label 38. The outer label is the LDP label that identifies the LSP to 6PE2 Loopback address.

https://t.me/learningnets

6VPE Architecture

IPv6 MPLS VPNs are like IPv4 MPLS VPNs. Customer sites are connected over the shared MPLS backbone so that customers have reachability to their own VPN partner sites only. As with IPv4 MPLS VPNs, this outsourced intranet remains private and is seen by the customer as a closed network that runs on a dedicated network.



The inherent VPN and terminal equipment services that are available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure that supports IPv4 VPNs and MPLS terminal equipment.

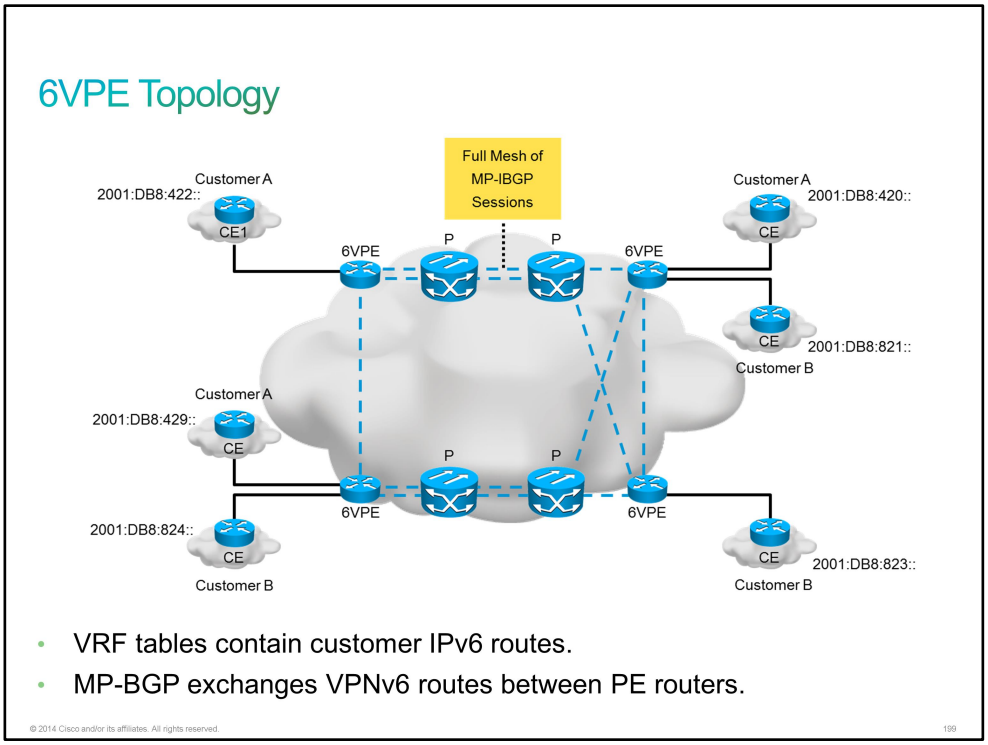
ISPs that offer MPLS VPNs for IPv4 can use Cisco 6VPE to add IPv6 services. This approach has these benefits:

- There is no modification on the MPLS core.
- Cisco 6VPE supports both IPv4 and IPv6 VPNs concurrently on the same interfaces.
- Configuration and operation of IPv6 VPNs are the same as for IPv4 VPNs.

https://t.me/learningnets

6VPE Topology

Service providers that offer MPLS IPv4 VPN services to their customers might look forward to adding IPv6 VPN services to their portfolios. A VPN is said to be an IPv6 VPN when a CE router turns on native IPv6 over an interface or subinterface to the PE router. Cisco 6VPE adds IPv6 VPN capability to Cisco 6PE and enables an ISP to deliver services that are similar to IPv4.



As with IPv4 VPN route distribution, BGP and its extensions are used to distribute routes from an IPv6 VPN site to all other Cisco 6VPE routers that connect to a site of the same IPv6 VPN. Cisco 6PE routers use VRF tables to separately maintain the reachability and forwarding information of each IPv6 VPN (as with IPv4 VPN).

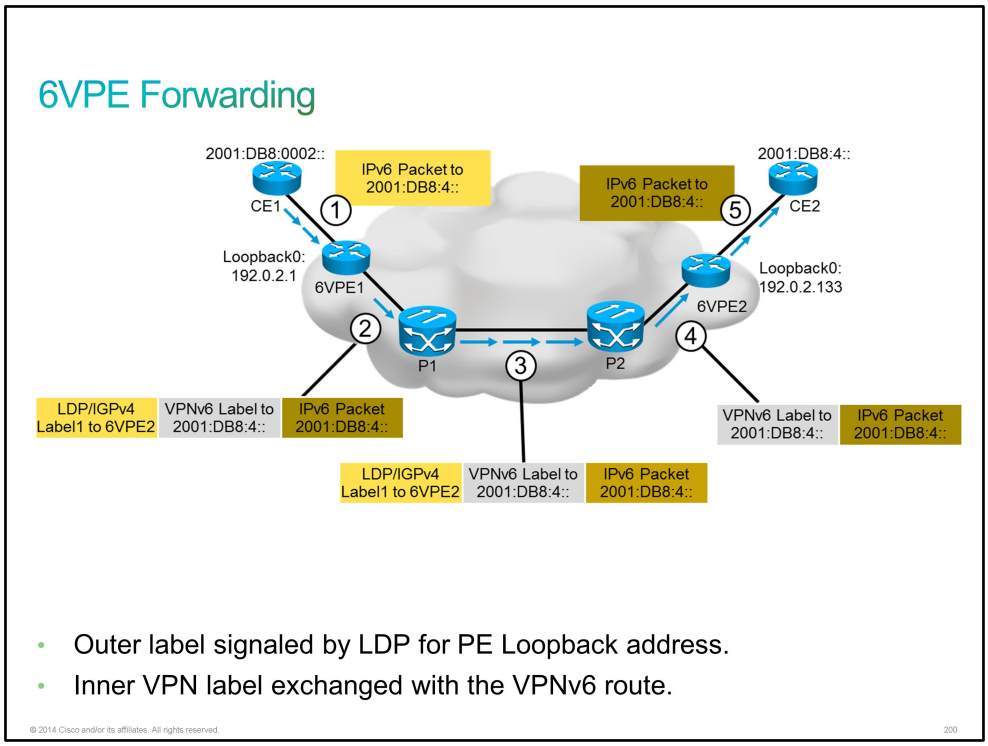
In a BGP/MPLS VPN, MPLS is used to forward packets over the backbone, and MP-BGP is used to distribute VPN routes over the service provider backbone.

MP-BGP allows generic operations over both IPv4 and IPv6. For example, MP-BGP defines an IPv6 VPN address family and describes the route distribution and MPLS tunnel selection.

https://t.me/learningnets

6VPE Forwarding

Packet forwarding in a 6VPE environment follows the same scheme as in IPv4 MPLS VPNs.



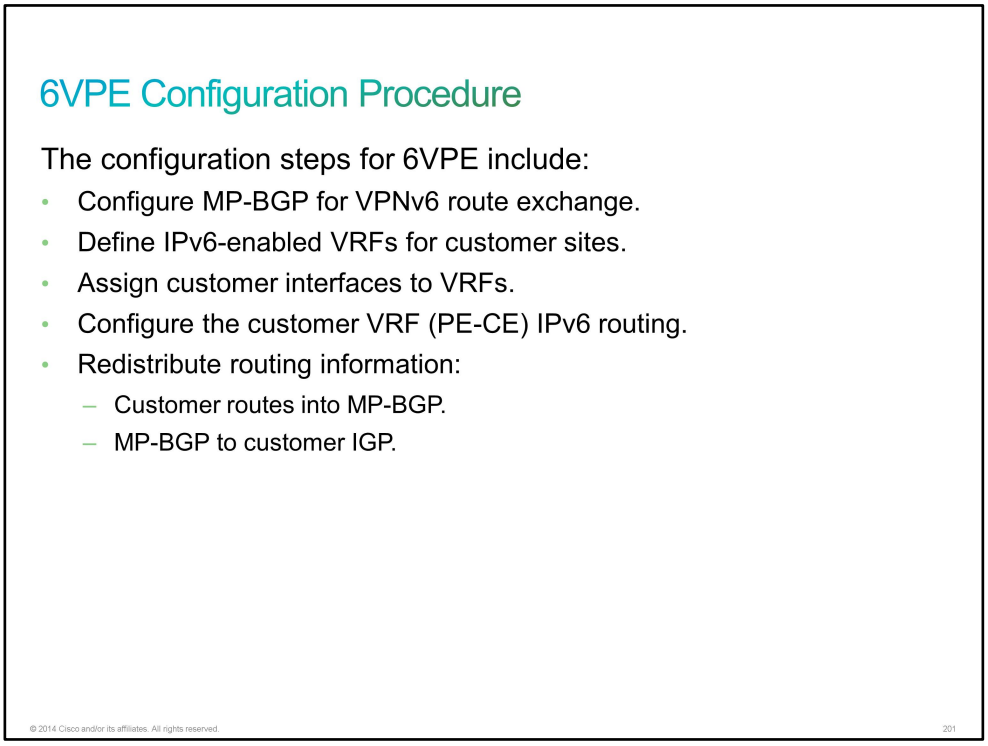
Forwarding in Cisco 6VPE follows these steps:

1. The ingress PE router receives an IPv6 packet and attaches an MPLS label. When a Cisco 6VPE router receives an IPv6 packet from CE1, the router looks up the IPv6 packet destination address in the VRF table. This address enables 6VPE1 to find a VPN IPv6 route. The VPN IPv6 router has an associated MPLS label and an associated BGP next hop. This MPLS label is imposed on the IPv6 packet.
2. The ingress PE router attaches a second label to the IPv6 packet. 6VPE1 pushes another label, which is bound by LDP IGP to the IPv4 address of the BGP next hop, to reach 6VPE2 through the MPLS cloud on the label stack of the labeled IPv6 VPN packet. This top-most imposed label corresponds to the LSP that starts on 6VPE1 and ends on 6VPE2. The bottom label is bound to the IPv6 VPN prefix via BGP. All the P routers in the backbone network switch the VPN packet, based only on the top label in the stack, which points toward the 6VPE2 router. Because of the normal MPLS forwarding rules, the P routers never look beyond the first label and are thus completely unaware of the second label or the IPv6 VPN packet that is carried across the backbone network.
3. The egress router receives the labeled IPv6 packet, removes the labels, and forwards the packet to the CE router. The egress PE router, 6VPE2, receives the labeled IPv6 VPN packet, drops the first label, and performs a lookup on the second label, which uniquely identifies the target VRF table and sometimes even the outgoing interface on 6VPE2. A lookup is performed in the target VRF table, and the IPv6 packet is sent toward the proper CE router in the IPv6 domain or site.

https://t.me/learningnets

6VPE Configuration Procedure

6VPE implementation builds upon a functional MPLS network with IPv4 routing.



These steps need to be completed before configuring the 6VPE feature:

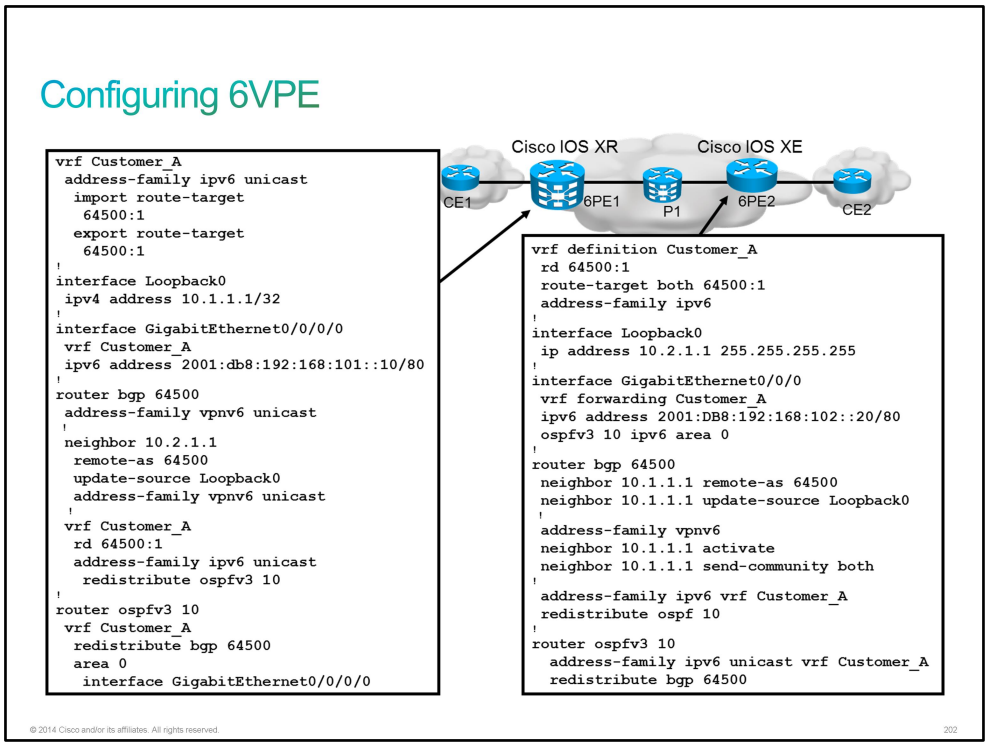
- Loopback interface have been configured with /32 host addresses on the PE routers.
- MPLS LDP has been enabled on the interfaces in the core.
- An IGP has been configured in the backbone.

Once these prerequisites have been met, you can follow the 6VPE configuration procedure.

https://t.me/learningnets

Configuring 6VPE

This example depicts a 6VPE configuration example on Cisco IOS XR and Cisco IOS XE routers.

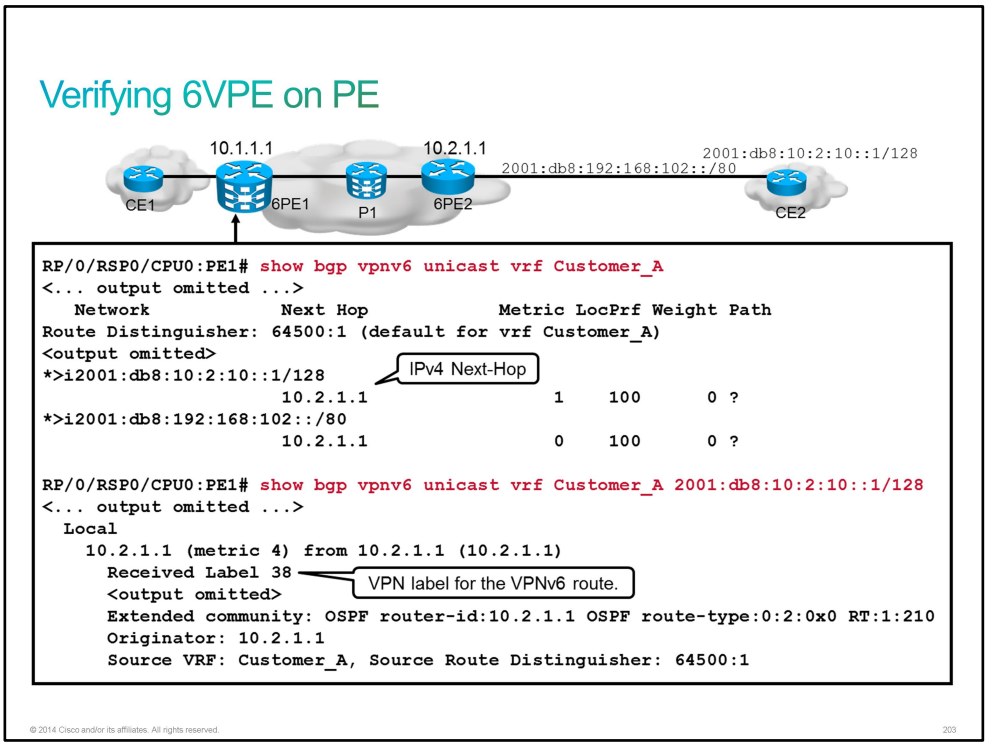


The PE routers have an MP-BGP session for VPNv6 route exchange, VRFs that terminate the IPv6-enabled PE-CE links, and redistribution between an OSPFv3 process that acts the the IPv6 PE-CE routing process and MP-BGP.

https://t.me/learningnets

Verifying 6VPE on PE

This example illustrates how to verify the 6VPE operation on a PE router.



Once MPLS in functional in the backbone, you can monitor the 6VPE operation by examining the VPNv6 routes in the PE BGP tables.

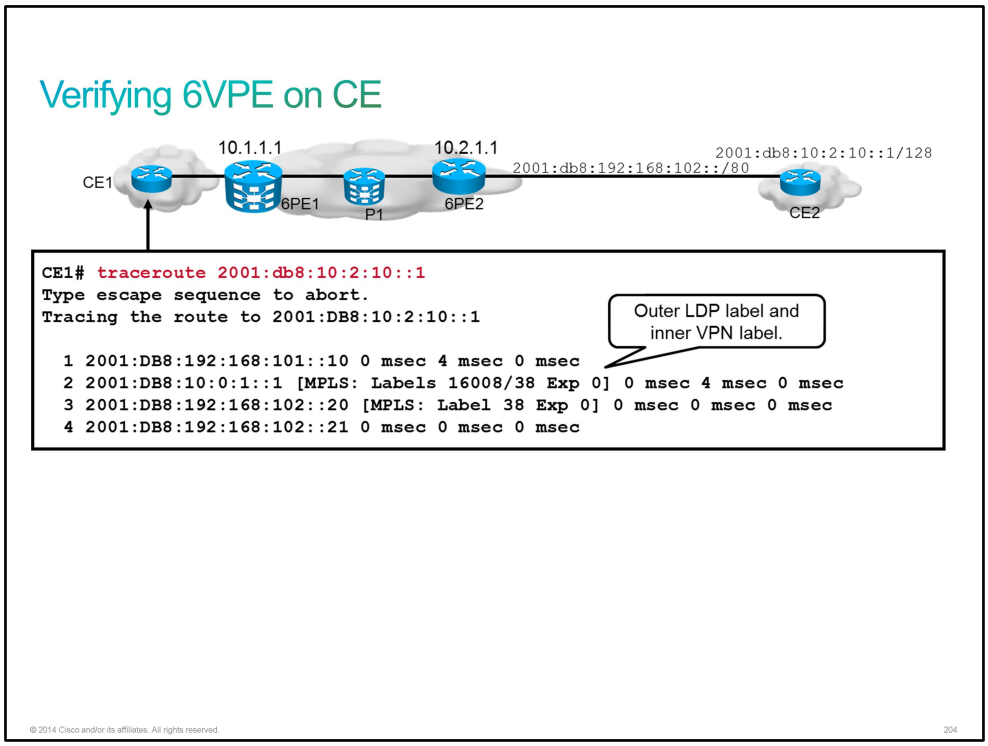
In the first output you see that the IPv6 prefixes displayed in the VPNv6 table contain an IPv4 next hop address. This results from the IPv4-only connectivity in the MPLS backbone. If MPLS were not used for traffic forwarding, the IPv6 packets could not be forwarded between the PEs.

The second output illustrates how to examine a specific IPv6 prefix in the BGP VPNv6 table. It shows a VPN label attached to that prefix. The VPN label, 38 in this example, is used as the inner label when customer traffic is transported through the MPLS backbone.

https://t.me/learningnets

Verifying 6VPE on CE

This example illustrates how to verify the 6VPE operation on a CE router.

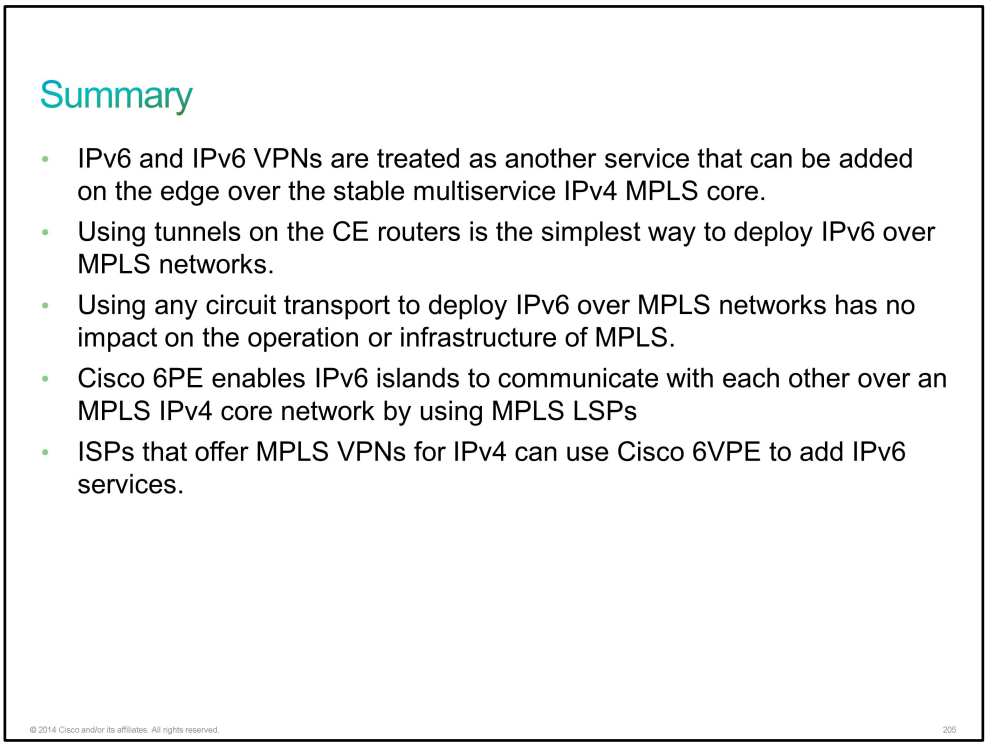


You can monitor the label swapping operation by using the traceroute utility, as shown in the last output on CE1. When you verify connectivity to CE2, you see that a label stack is used between 6PE1 and the core router. The inner label is the VPN label 38. The outer label is the LDP label that identifies the LSP to 6VPE2 Loopback address.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

Lesson 2: Introducing MPLS Interdomain Solutions

Overview

Deployments of MPLS have become routine in large-scale global networks, which demand solutions to address complex business and network problems. There are two primary components of the Cisco IOS Multiprotocol Label Switching Interdomain Solution: the inter-AS and CSC.

Upon completing this lesson, you will be able to:

- Identify two MPLS interdomain solutions (CSC and inter-AS)
- Describe the inter-AS deployment methods
- Describe the inter-AS Back-to-Back VRF method
- Describe the inter-AS Back-to-Back VRF method characteristics
- Describe the inter-AS single-hop MP-EBGP method
- Describe the inter-AS single-hop MP-EBGP method characteristics
- Describe the inter-AS multi-hop MP-EBGP method
- Describe inter-AS multihop MP-EBGP characteristics
- Provide an overview of CSC
- Explain the benefits of CsC
- Explain the operation of CSC
- Describe the CsC implementation options
- Describe the CSC option using IGP and LDP
- Describe the data flow in CSC with IGP/LDP
- Describe the configuration overview of CSC using IGP and LDP

https://t.me/learningnets

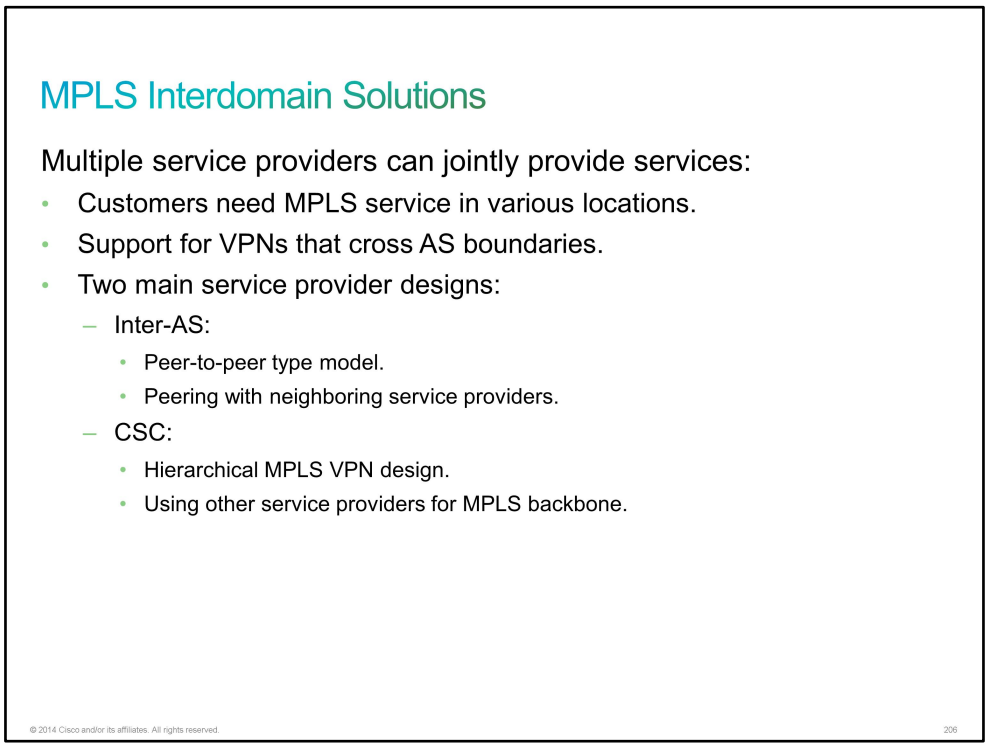
- Describe the configuration of MP-IBGP on CSC-PE for CSC (LDP/IGP)
- Describe the configuration of MPLS LDP in PE VRF for CSC (LDP/IGP)
- Describe the MP-IBGP configuration on PE routers for CSC (IGP/LDP)
- Describe how verify the LDP session between the PE and the CSC-CE
- Describe the configuration of CSC using IGP and LDP
- Describe how to verify the VPN label for the CSC-PE Loopback address on the CSC-CE router
- Describe how to verify the end customer label
- Describe CSC implementation using MP-EBGP
- Describe the data flow in CSC with IGP/LDP
- Describe the configuration of CSC (MP-EBGP) on the CSC-CE routers
- Describe the configuration of CSC (MP-EBGP) on the PE routers
- Describe how to verify MP-EBGP capabilities
- Describe how to verify the BGP and label information on the egress PE
- Describe how to verify the BGP and label information on the ingress PE
- Describe how to verify the BGP and label information on the CSC-CE
- Describe how to verify the BGP and label information on the CSC-PE
- Describe how to verify the CSC functionality using customer traffic

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

MPLS Interdomain Solutions

There are two primary components of the Cisco IOS Multiprotocol Label Switching Interdomain solution: Inter-AS and CSC.



Inter-AS is a peer-to-peer type model that allows the extension of VPNs through multiple provider or multidomain networks. This solution enables service providers to peer up with one another and offer end-to-end VPN connectivity over extended geographical locations for those subscribers who may be out of reach for a single provider.

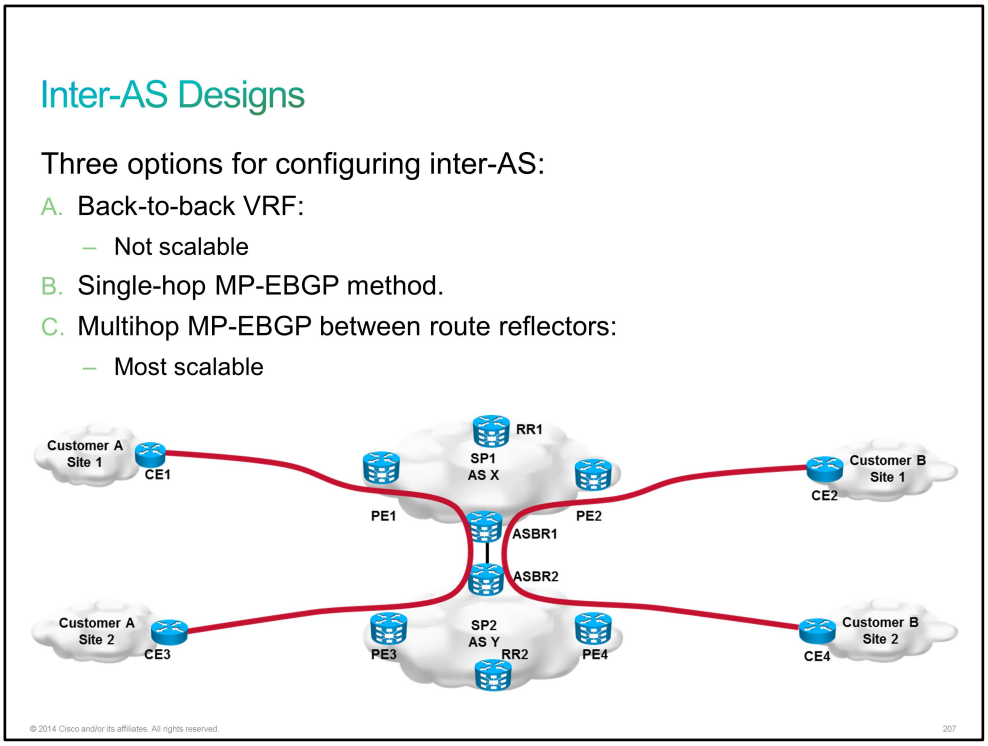
CSC is a hierarchical VPN model that allows small service providers, referred to as customer carriers, to interconnect their IP or MPLS networks over an MPLS backbone. This eliminates the need for customer carriers to build and maintain their own MPLS backbone.

Both inter-AS and CSC can construct scalable networks that help maintain network segmentation based on internal organizational or operational boundaries.

https://t.me/learningnets

Inter-AS Designs

In traditional service provider networks, MPLS VPN was mostly used inside one AS. In this traditional approach all branch offices had to be connected to the same service provider to be able to participate in a MPLS VPN.



Inter-AS introduces techniques to establish MPLS VPNs across multiple autonomous systems. Service providers have to establish interconnection, exchange VPN information, and establish VPNs.

There are three techniques for establishing inter-AS:

A.) Back-to-back VRF:

- The simplest method.
- Not scalable.
- Packets are forwarded as IP packets between the ASBRs.

B.) Single-hop MP-EBGP:

- Medium scalability.
- Some routing overhead in ASBRs.
- MP-EBGP used for VPNv4 prefix exchange between the ASBRs.

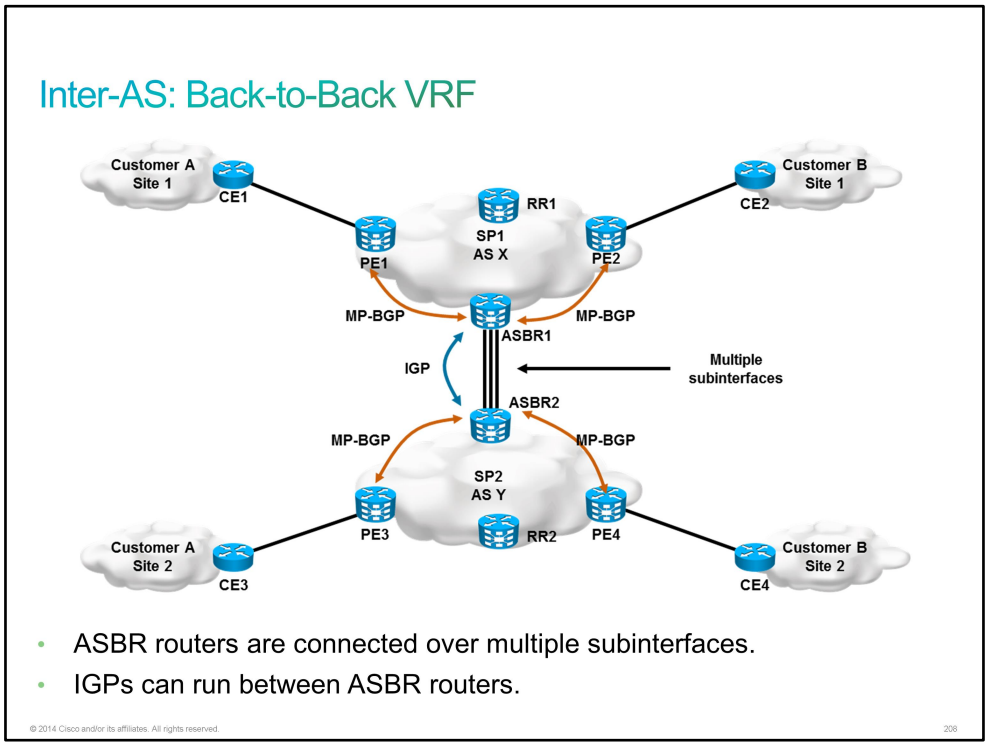
C.) Multihop MP-EBGP between route reflectors:

- Most scalable method.
- ASBRs don’t hold any VPNv4 prefix/label information.
- VPNv4 routes are exchanged between the route reflectors.

https://t.me/learningnets

Inter-AS: Back-to-Back VRF

You can configure inter-AS functionality using different techniques. The first is referred to as the back-to-back VRF method.



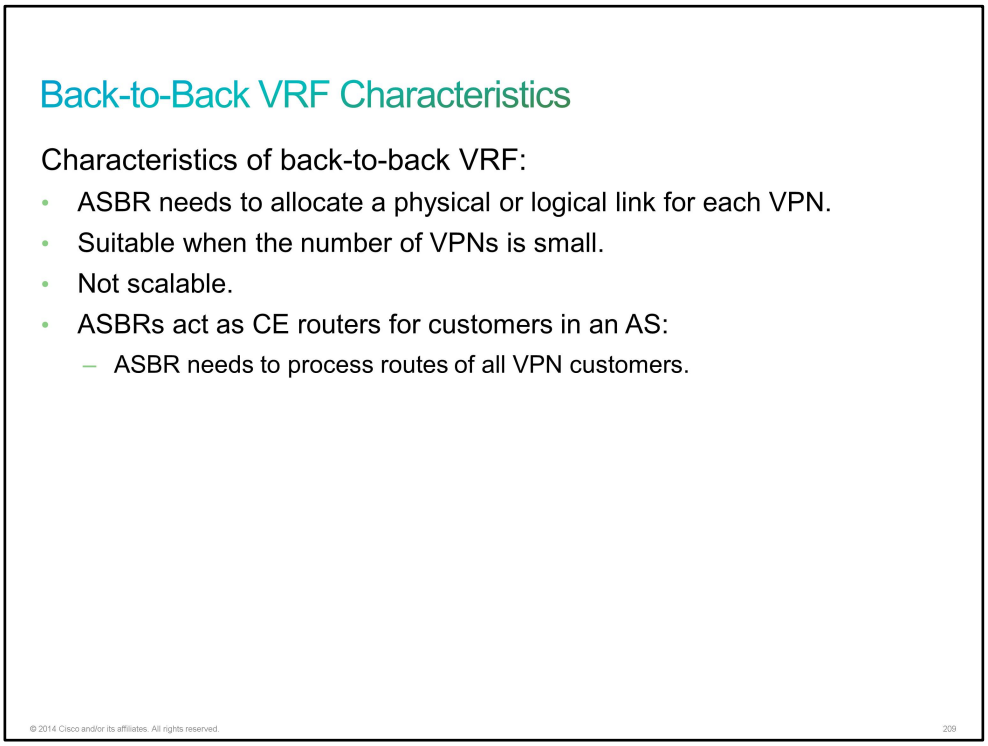
ASBRs are interconnected using multiple interfaces or subinterfaces. Each interface or subinterface is used to carry the traffic of its own VPN.

Each ASBR is acting as a PE router for its customers and a CE router for customers of other service providers. An IGP or BGP are used to exchange customer routing information between ASBRs.

https://t.me/learningnets

Back-to-Back VRF Characteristics

The back-to-back VRF inter-AS method is only suitable when two service providers offer jointly only a small number of MPLS VPNs.

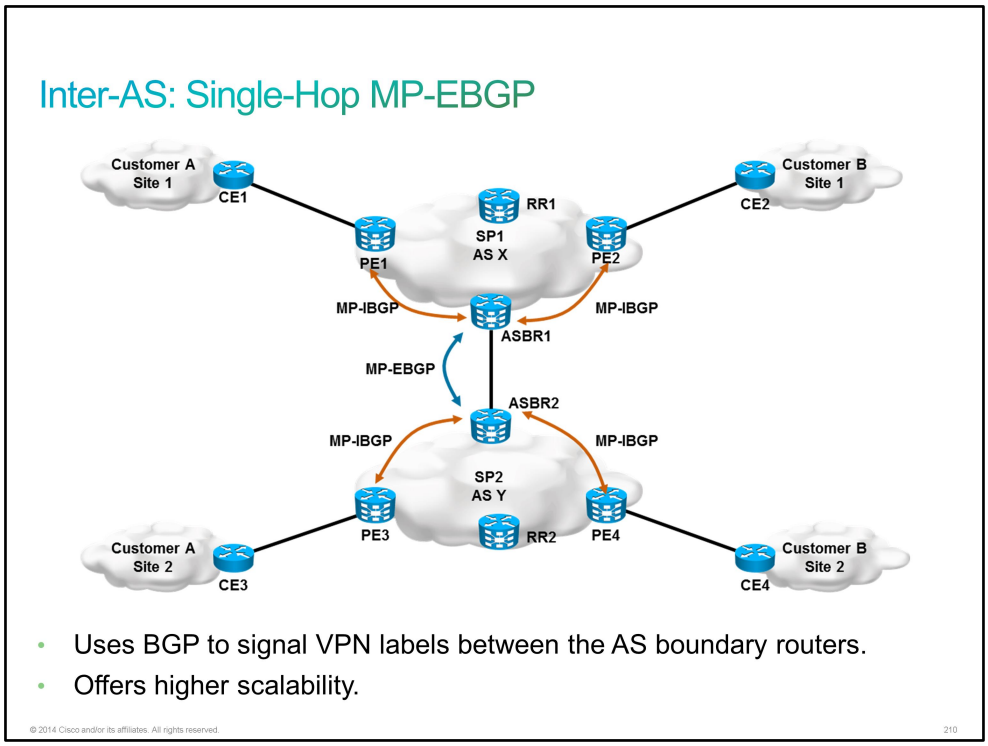


The back-to-back VRF inter-AS method is easy to configure, but offers minimum scalability. Because ASBRs act as CE routers for all VPN customers, these routers have to maintain routes from all customers in their memory.

https://t.me/learningnets

Inter-AS: Single-Hop MP-EBGP

The second method for configuring inter-AS functionality is the single-hop MP-EBGP design.

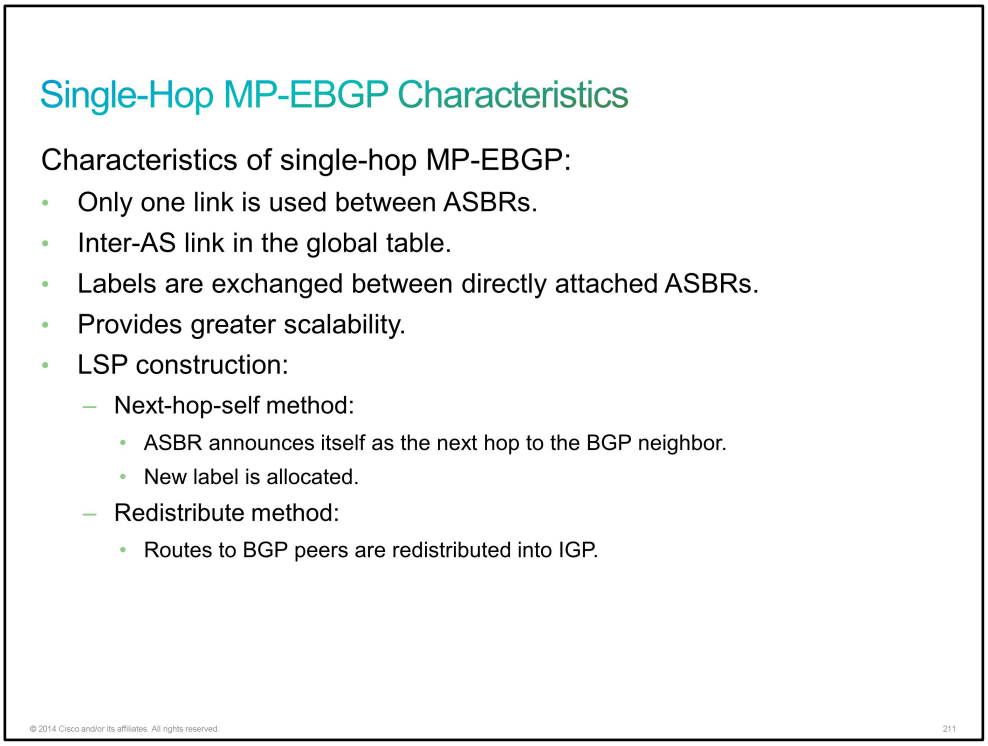


Inside an AS, MPLS and BGP are used to transfer VPN information and construct the LSP. Between autonomous systems, the single-hop MP-EBGP method is used to transfer VPN information and construct the LSPs.

https://t.me/learningnets

Single-Hop MP-EBGP Characteristics

The single-hop MP-EBGP method provides more scalability than the back-to-back VRF method. Only one link is configured between service providers. MP-BGP is used co exchange routing and label information between directly connected routers.

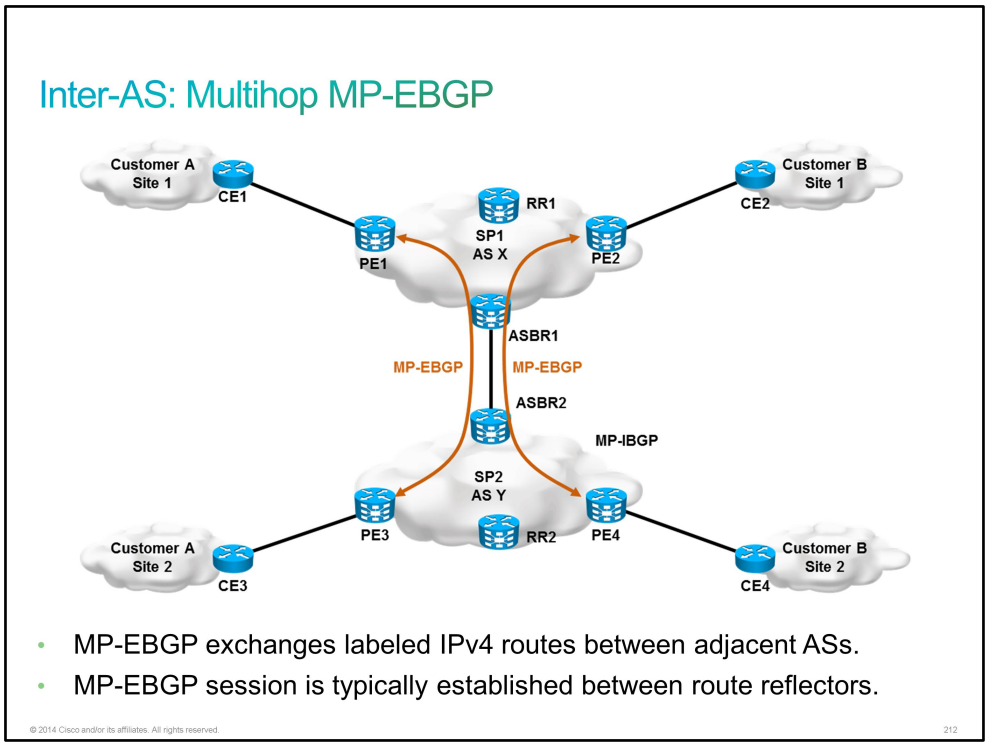


To construct the LSP path, next-hop addresses should be reachable. Routes to BGP peers can be redistributed to the provider IGP, or next-hop-self can be used by having the ASBR replace the next hop with its IP address.

https://t.me/learningnets

Inter-AS: Multihop MP-EBGP

The third method for configuring inter-AS functionality is the multihop MP-EBGP design.

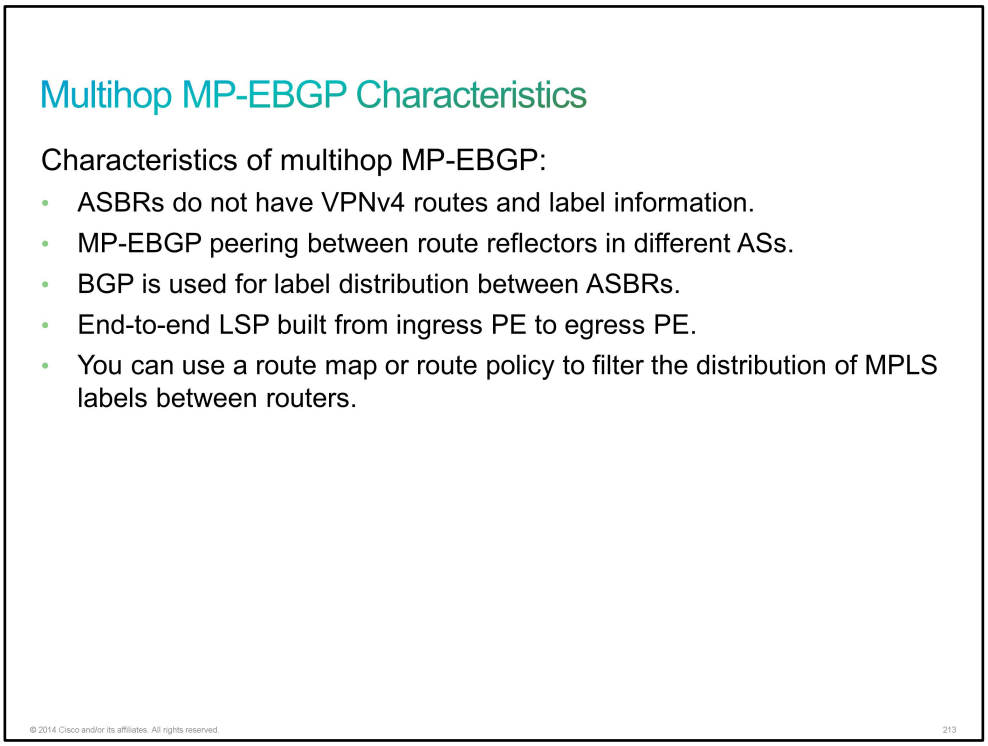


Because BGP only requires the TCP connection to form a BGP neighbor and transfer route information, this third method transfers VPN route information between the source and destination PEs directly over multihop MP-EBGP, and constructs an LSP between the source and destination PEs. If a route reflector exists in the service provider networks, the MP-EBGP session can be established between the route reflectors.

https://t.me/learningnets

Multihop MP-EBGP Characteristics

In the multihop MP-EBGP method, an end-to-end LSP is built from the ingress PE to the egress PE.



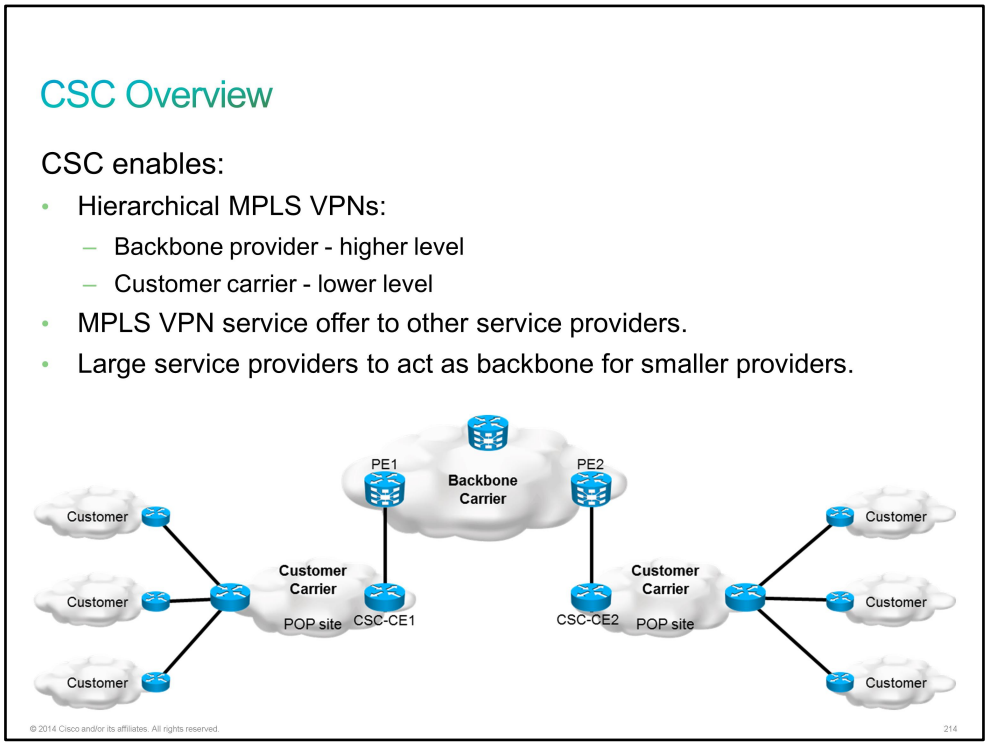
The multihop MP-EBGP method is highly scalable, because there is no route overhead on ASBRs.

You can use a route map or route policy to filter the distribution of MPLS labels between routers.

https://t.me/learningnets

CSC Overview

A carrier network transports traffic between customer sites. Large service providers can interconnect carrier networks of smaller service providers. In this scenario, a smaller service provider acts as a customer for a larger service provider.

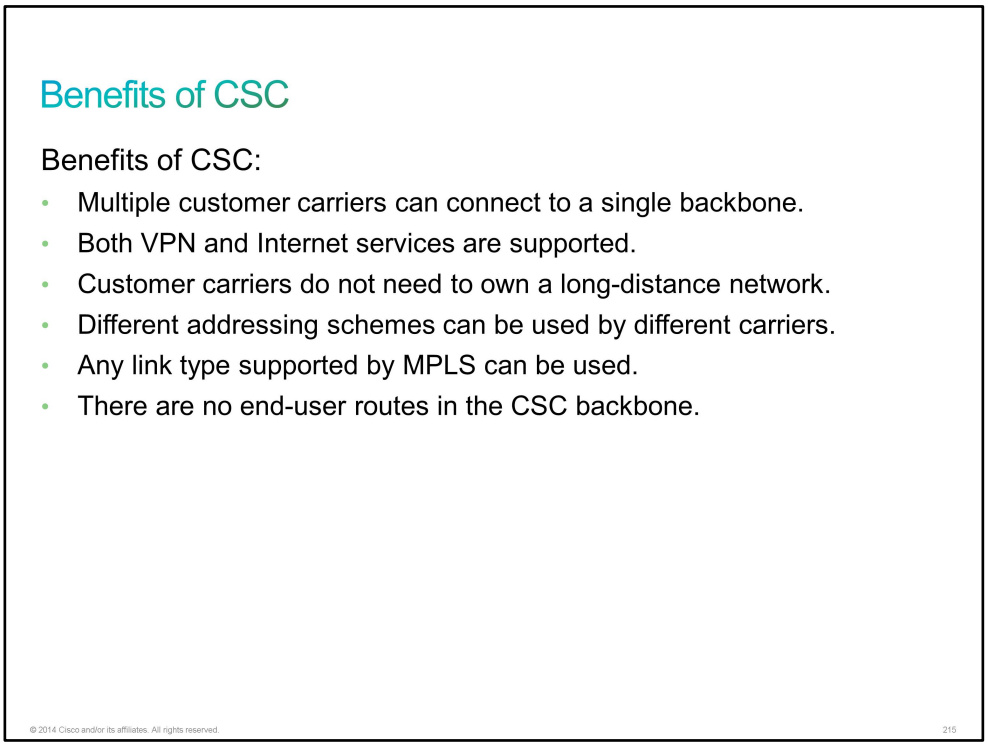


CSC provides MPLS VPN service to other service providers. CSC creates a hierarchical structure with a first-level service provider as the backbone carrier and a second-level service provider as customer carrier. Many customer carrier sites, also called POPsites, can be interconnected using the backbone carrier.

https://t.me/learningnets

Benefits of CSC

CSC paves way for a number of business opportunities.



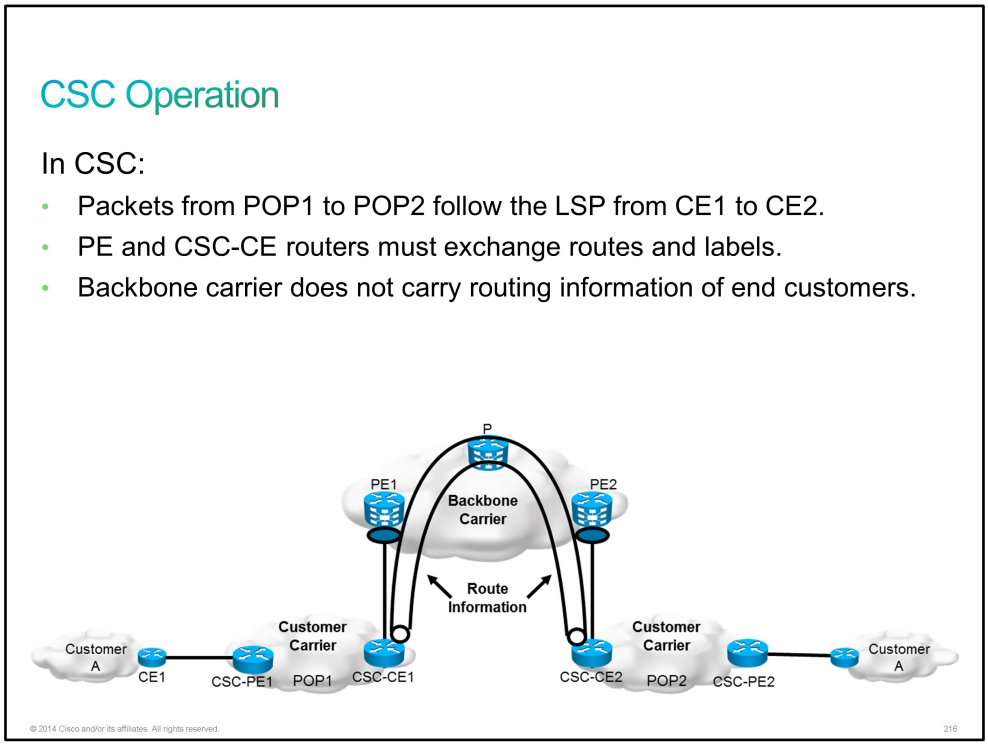
CSC offers a number benefits:

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

https://t.me/learningnets

CSC Operation

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



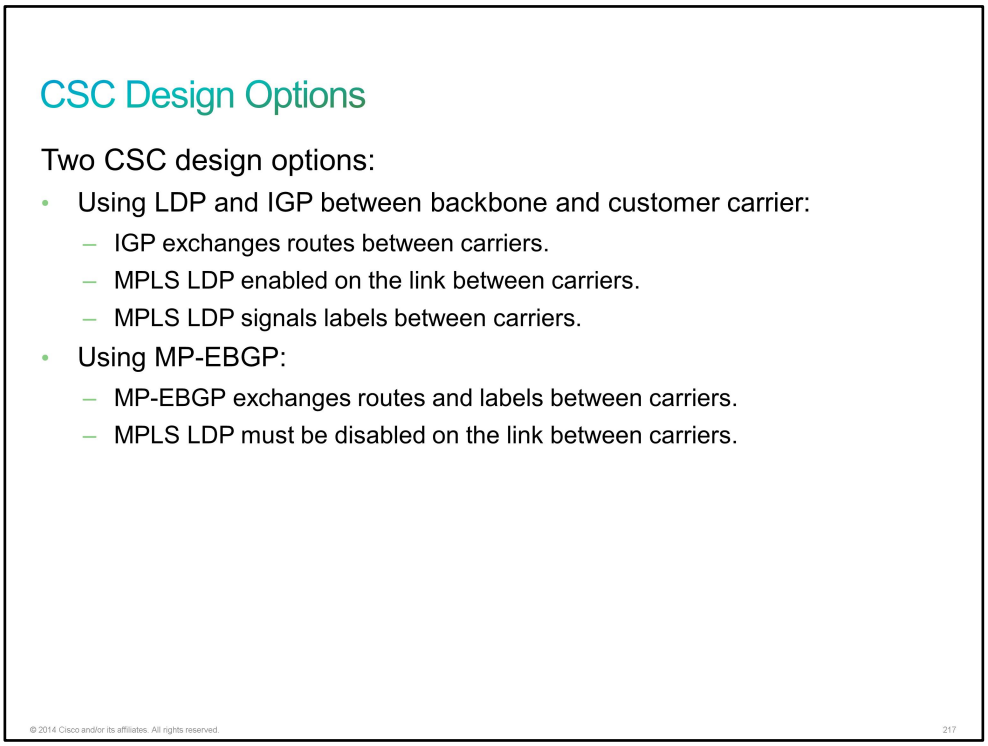
VRF tables are enabled on the CSC PE routers. The label exchange between PE1 and PE2 establishes a LSP from CE1 via the CSC backbone to CE2. Another LSP is also established in the other direction.

The customer carrier can forward the packets between POP1 and POP2 along the LSPs. The CSC backbone does not have to know about end-user sites and their IP addresses.

https://t.me/learningnets

CSC Design Options

To support CSC, the CSC backbone must support MPLS VPNs.



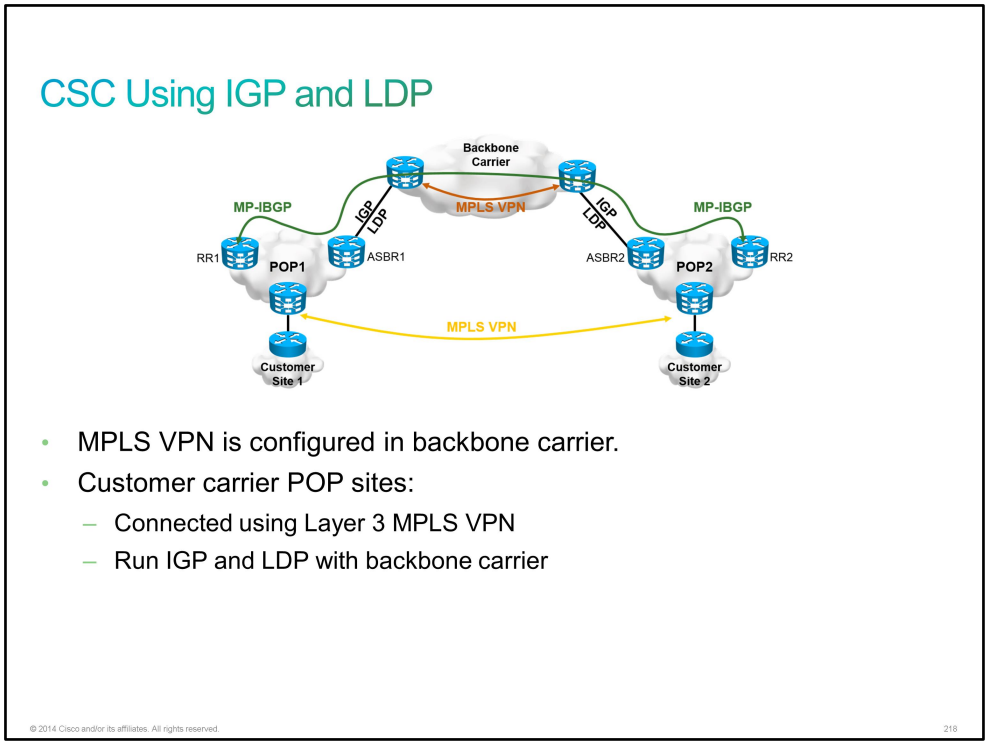
The customer carrier can connect to the backbone carrier in the following ways:

- Using LDP and an IGP:
  - MPLS has to be enabled on the link between the backbone carrier and the customer carrier. LDP is used for label distribution.
  - IGP is used for route exchange.
- Using MP-EBGP:
  - MP-EBGP is used for label and route exchange. LDP must be disabled on the links connecting the customer carrier and the backbone carrier networks, because the two label signaling methods are mutually exclusive.

https://t.me/learningnets

CSC Using IGP and LDP

The first CSC implementation builds on the MPLS and LDP model.

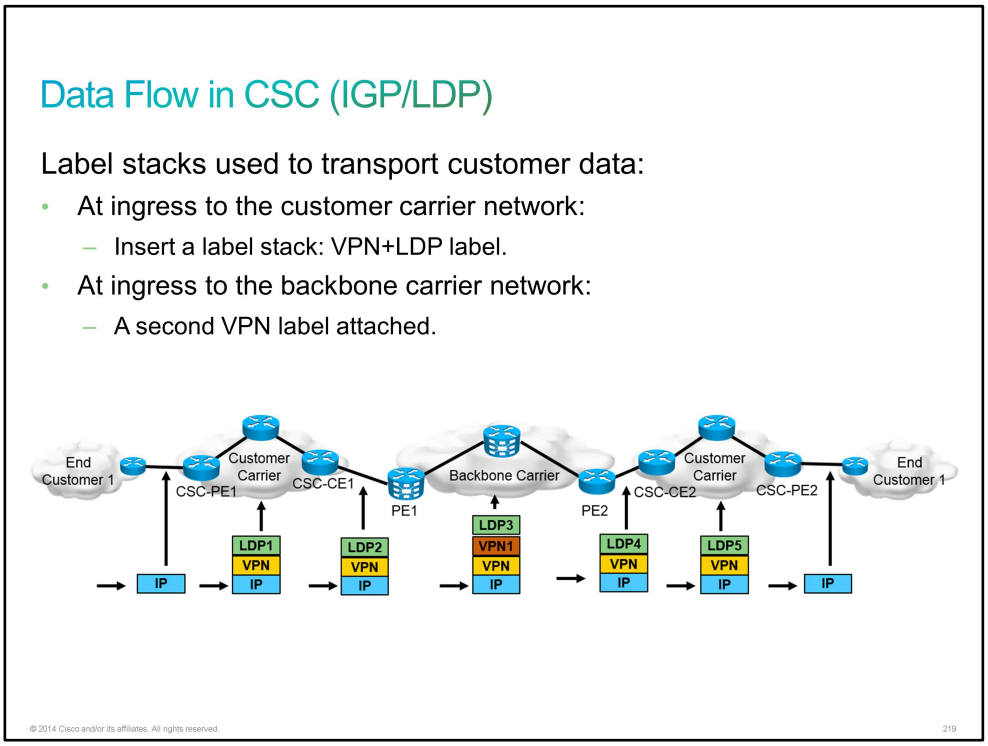


The resulting end-to-end LSP enables the customer carrier to establish a peer relationship between its PE routers that are supporting the end customer. It then enables the customer carrier to use MPLS VPNs to support its end customers over an end-to-end VPN.

https://t.me/learningnets

Data Flow in CSC (IGP/LDP)

In CSC with IGP and LDP, a label stack is used to transport the customer packets through the customer carrier and the backbone carrier networks.



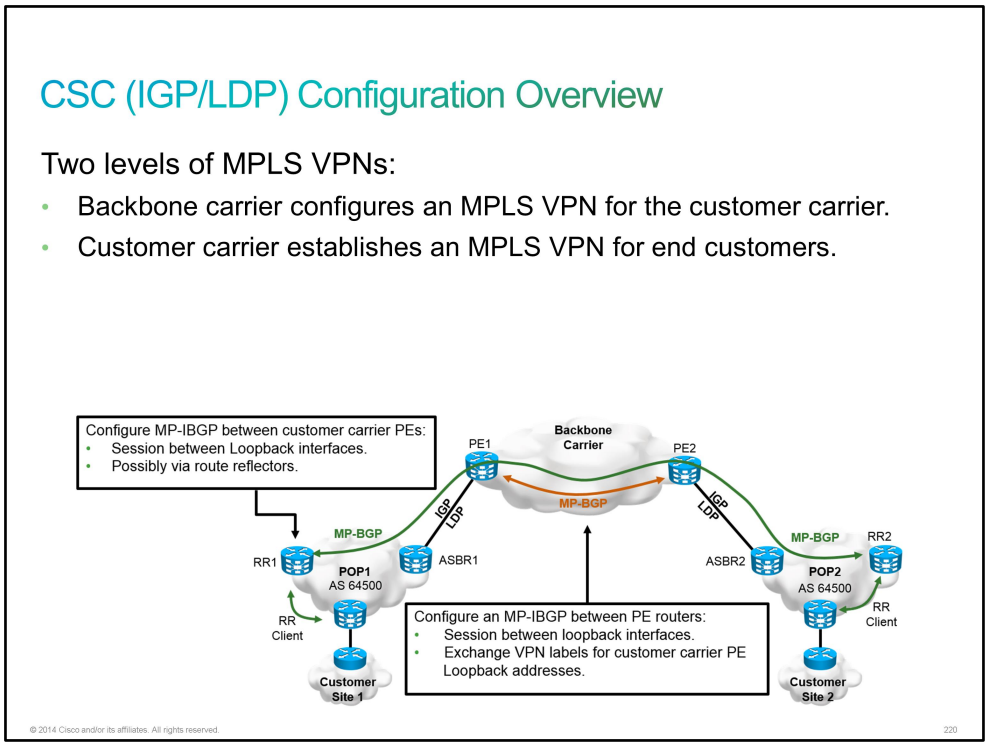
The label stack consists of two labels in the customer carrier network and of three labels in the backbone carrier network.

As illustrated in this example, a customer sends an IP packet. A label stack is inserted, consisting of a VPN label and an LDP label. The VPN label is used to identify packets from the same VPN. The LDP label is attached to packet in order to forward the packet in the customer carrier network. When the packet enters the backbone carrier network, the LDP label is swapped like on any other network link and another VPN label is attached to the packet. This label identifies all packets from this customer carrier.

https://t.me/learningnets

CSC (IGP/LDP) Configuration Overview

CSC can be thought of as an end customer MPLS VPN embedded inside a carrier VPN.

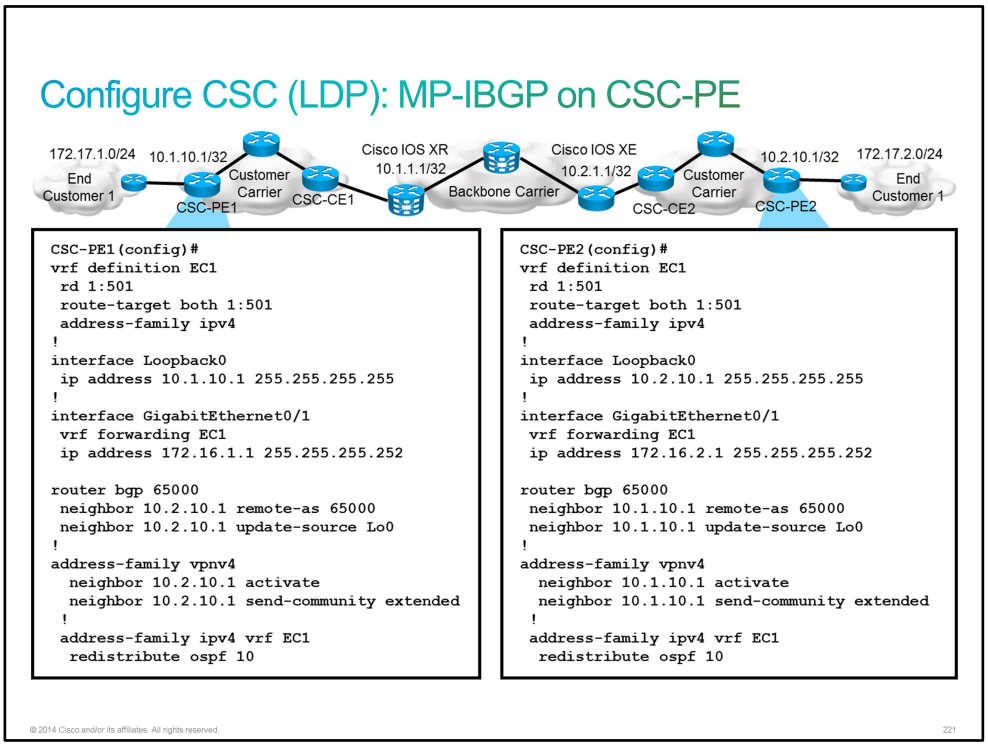


The customer carrier has to configure MP-IBGP sessions between its PE routers, possibly using route reflectors. The backbone carrier has to establish MP-IBGP sessions between its PEs. The MP-IBGP sessions of the customer carrier will be transported inside the VPN established by the backbone carrier.

https://t.me/learningnets

Configure CSC (LDP): MP-IBGP on CSC-PE

This figure illustrates how to configure MP-IBGP on the customer carrier PE routers.



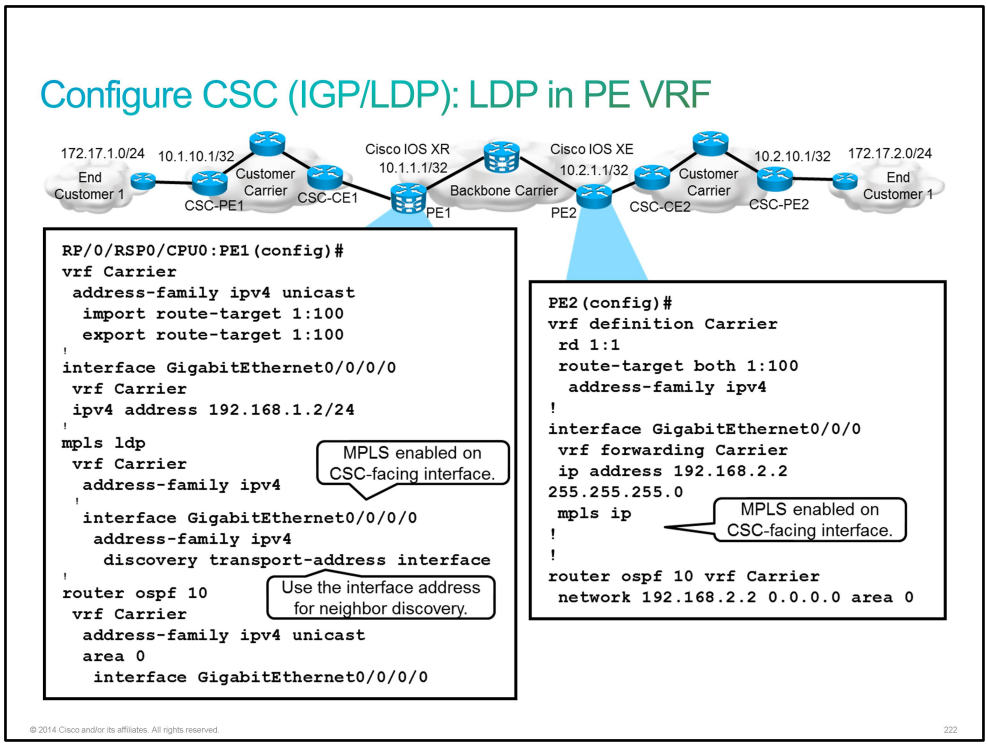
The customer carrier PEs have VRFs for the customer links and MP-IBGP sessions to other customer carrier PEs. This MP-IBGP carries VPNv4/VPNv6 routes with the customer prefixes. The MP-IBGP session has to be established between the Loopback IP addresses (/32 networks) to enable MPLS VPN connectivity between the end customer sites.

Although not shown in this example, MPLS LDP and OSPF are enabled on all internal links within the customer carrier networks, and on the links attached to the PEs.

https://t.me/learningnets

Configure CSC (IGP/LDP): MPLS LDP in PE VRF

This figure illustrates how to enable LDP on the VRF links in Cisco IOS XR and Cisco IOS XR.



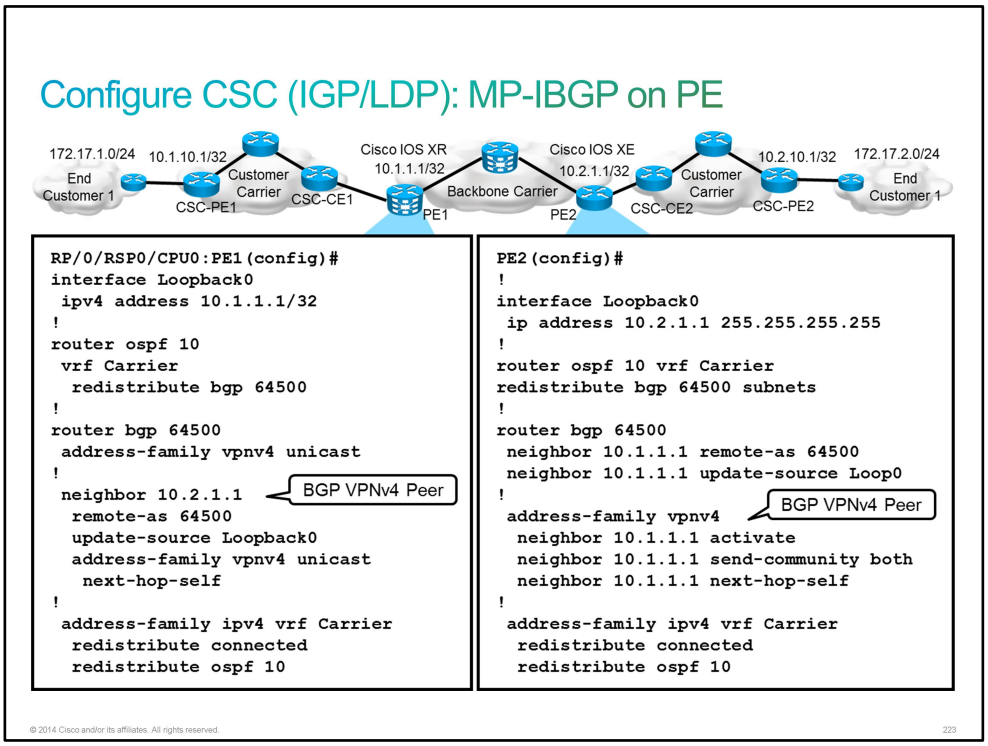
This example illustrates the configuration of PE interfaces connecting to the customer carrier CEs. The interfaces are assigned to carrier VRFs. MPLS LDP and OSPF routing is enabled on them. LDP and IGP will be used to exchange labels for the CSC PE Loopback networks.

To enable MPLS LDP in Cisco IOS XR you need to enter the MPLS LDP VRF sub-configuration mode, choose the interface, and enable MPLS LDP for a specific address family, IPv4 in this example. It is often necessary to set the LDP neighbor discovery transport address to the IP address of the VRF interface. This is done with the **discovery transport-address interface** command, shown in this scenario.

https://t.me/learningnets

Configure CSC (IGP/LDP): MP-IBGP on PE

This figure illustrates how to configure MP-IBGP on PEs running Cisco IOS XR and Cisco IOS XR.



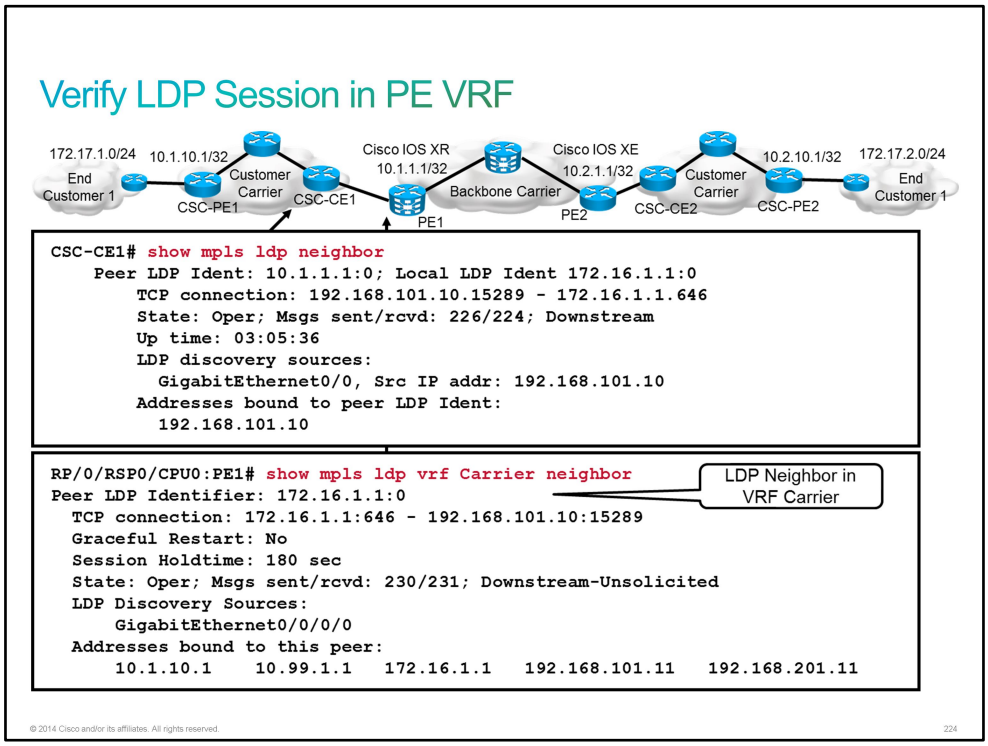
The backbone carrier needs to establish MP-IBGP sessions between the PEs. The full-mesh MP-IBGP requirement is typically resolved by deploying route reflectors.

The MP-IBGP sessions are used to exchange VPNv4/VPNv6 routes, which carry customer carrier prefixes. The only prefixes that have to be exchanged in the CSC solution are the /32 addresses of the CSC-PE Loopback networks.

https://t.me/learningnets

Verify LDP Session in PE VRF

The LDP session between the PE and the CSC-CE is needed to exchange the labels for the CSC-PE Loopback addresses.

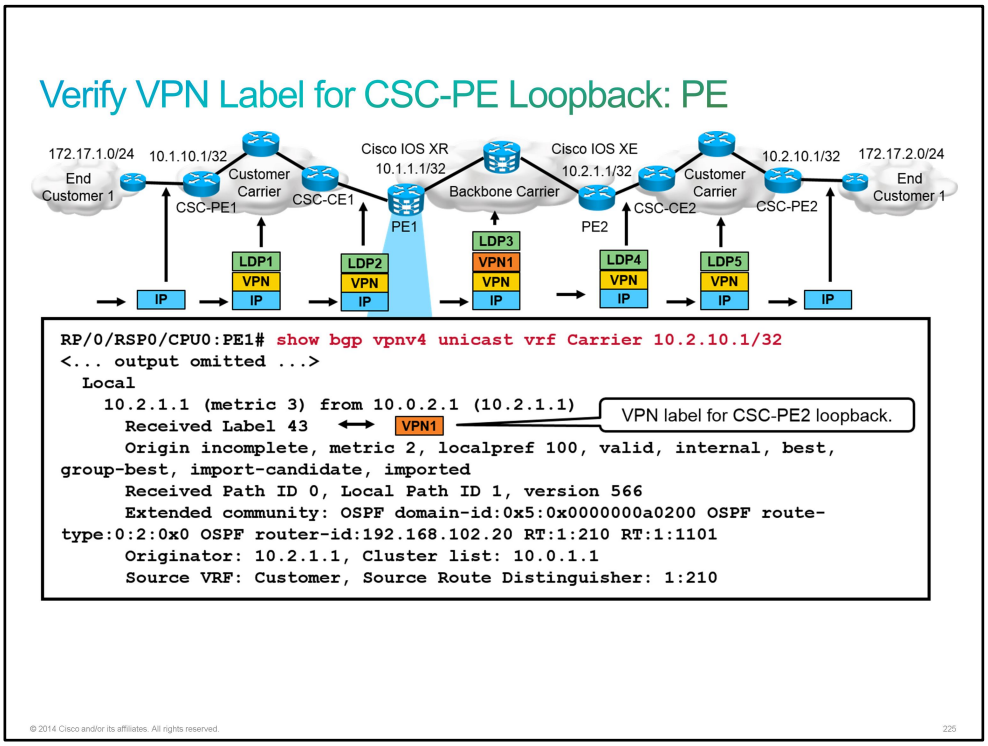


The LDP session is terminated inside the carrier VRF on the backbone carrier PE. You can view the LDP session using the **show mpls ldp neighbor** command. On the PE you need to use the **vrf** option to examine the LDP sessions terminated in a VRF.

https://t.me/learningnets

Verify VPN Label for CSC-PE Loopback: PE

The VPN label exchanged in the backbone carrier network identifies the Loopback IP address of the CSC-PE.



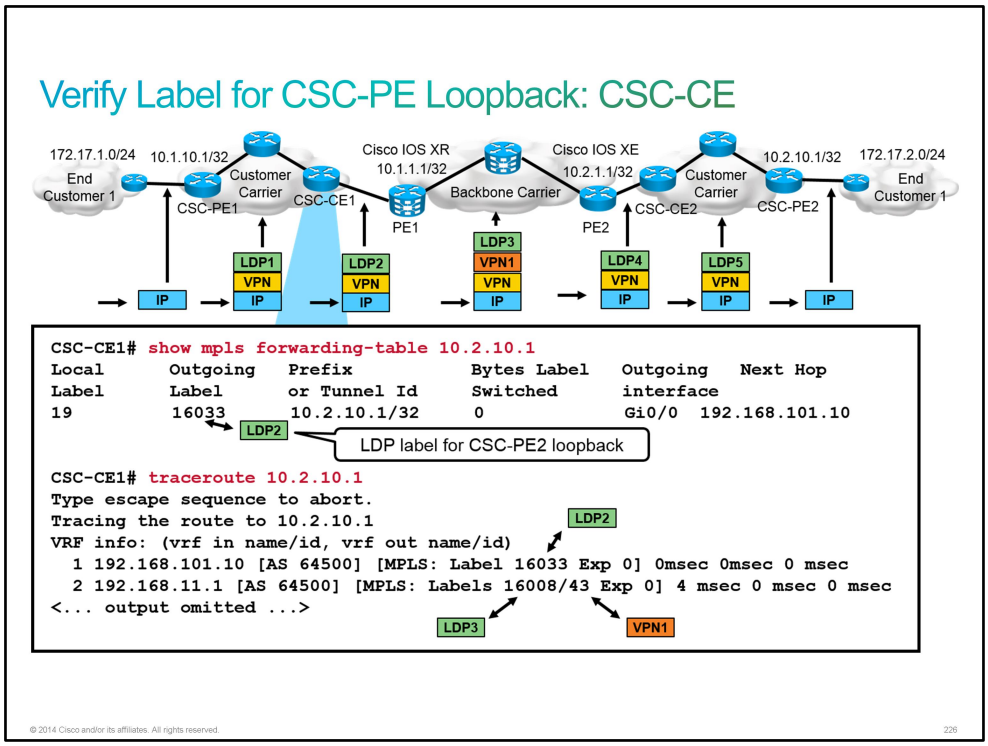
On PE1 you can view the VPN label exchanged in the backbone carrier network for the Loopback IP address of the CSC-PE by examining the appropriate prefix in the BGP table.

This VPN label, depicted as VPN1 in the figure, with the value of 43 in this example, will be used to tag the packets of customer carrier entering the backbone network from CSC-CE1. It will be the second inner-most label on the links inside the backbone, with the exception of the last link due to penultimate hop popping.

https://t.me/learningnets

Verify Label for CSC-PE Loopback: CSC-CE

The labels for the CSC-PE Loopback networks are exchanged via LDP between the PEs and CSC-CEs.



You can examine the label for CSC-PE Loopback network by viewing the MPLS forwarding table on the CSC-CE.

In this example the label received for the CSC-PE2 Loopback address (10.2.10.1/32) via LDP is 16033.

When you use the traceroute utility to examine the LSP to the CSC-PE2 Loopback address, you will see that:

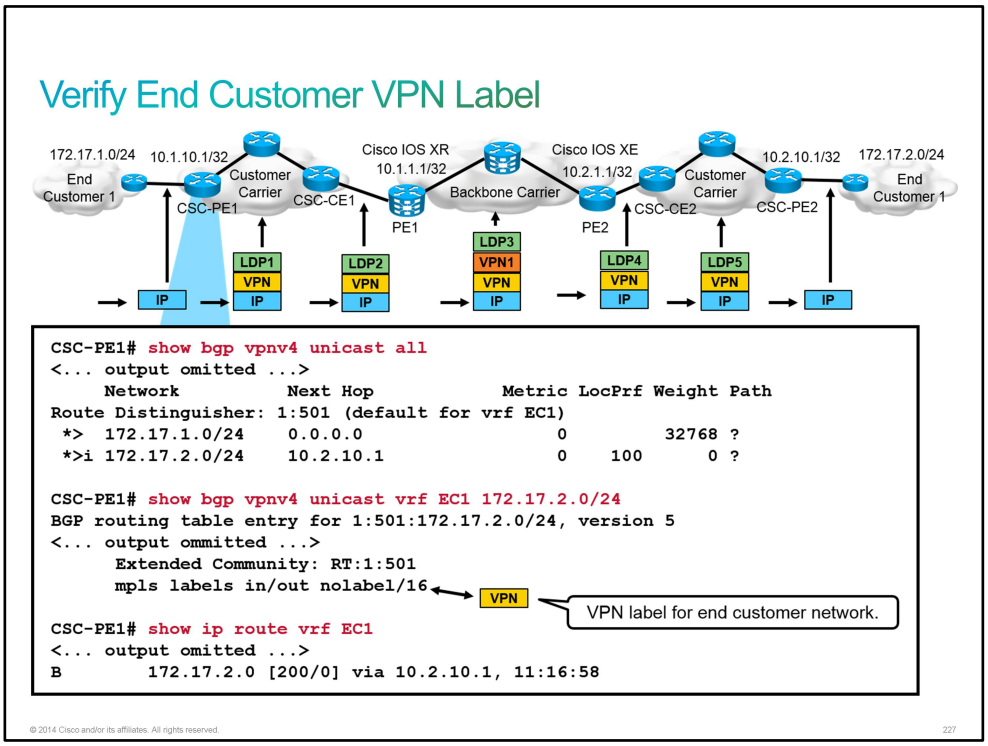
- CSC-CE1 inserts the label 16033 when sending packets over the link towards the PE.
- PE1 inserts the VPN label (43) when transporting the packets through the backbone. This is achieved through label swapping of label 16033 against a label stack. The label stack consists of the LDP label (16008) used on the first link in the backbone network and the VPN label 43. You could see the VPN label 43 in the BGP table on PE1.

Further sections of the LSP are omitted in this example.

https://t.me/learningnets

Verify End Customer VPN Label: CSC-PE

The labels for the end customer networks can be viewed on the CSC-PE routers.



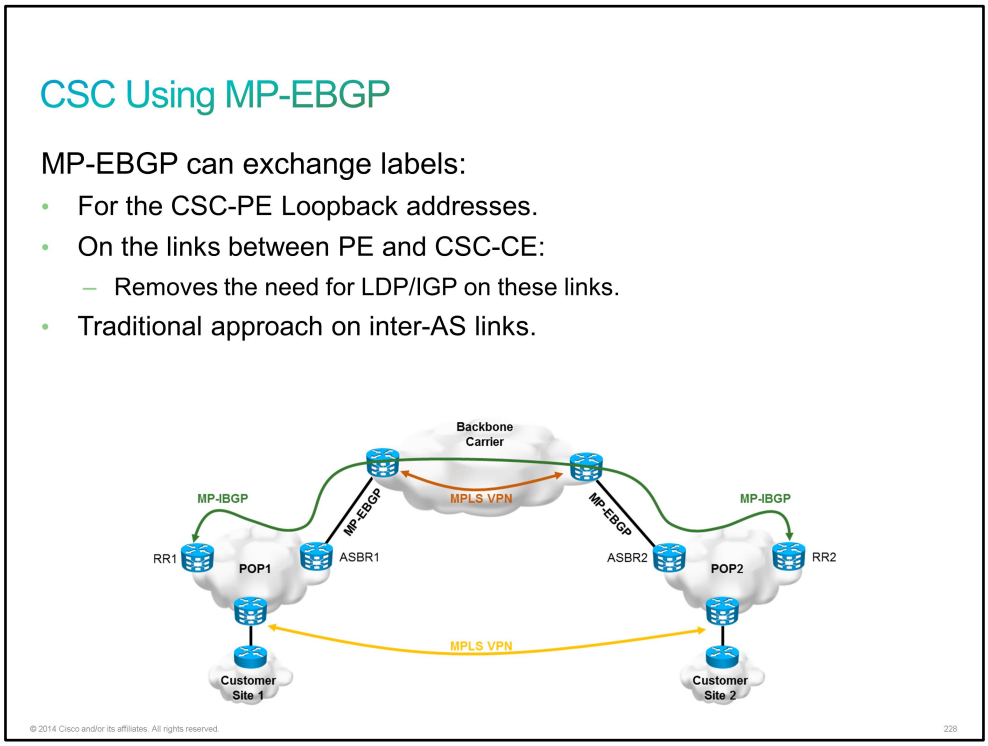
The VPN labels for the end customer networks are exchanged via MP-IBGP established among the CSC-PE routers. You can examine the appropriate prefixes in the BGP table to view the attached label. In this example the VPN label is 16.

You can examine the global routing table and the end customer VRF routing table. The global routing table needs to contain the CSC-PE Loopback addresses, displayed as host /32 routes. The VRF routing table assigned to an end customer will include the end customer networks from the local and remote site.

https://t.me/learningnets

CSC Using MP-EBGP

MP-EBGP offers an alternative method of exchanging labels on the links between the backbone and the customer carriers.



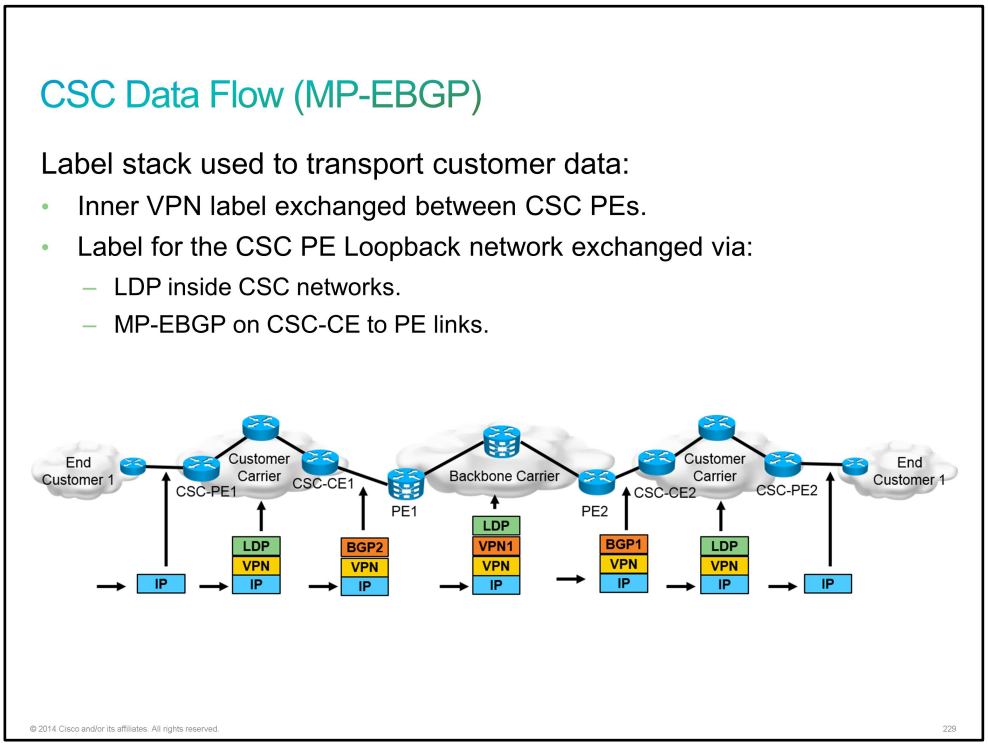
You can use MP-EBGP, instead of LDP, to exchange the labels of the CSC-PE Loopback addresses. This method is often preferred by service providers who traditionally lean towards BGP for inter-domain routing purposes.

MPLS is deployed within the backbone network and within each customer carrier. An MP-IBGP session between the CSC-PE routers provides VPNv4/VPNv6 routing and exchanges VPN labels for the end customer prefixes. An MP-IBGP session between the PE routers provides VPNv4 routing and label exchange for the CSC-PE Loopback addresses. The MP-IBGP full mesh requirement on both levels, backbone and customer carrier, is typically resolved by deploying route reflectors.

https://t.me/learningnets

CSC Data Flow (MP-EBGP)

In CSC using MP-EBGP, a label stack is used to transport the customer packets through the customer carrier and the backbone carrier networks.



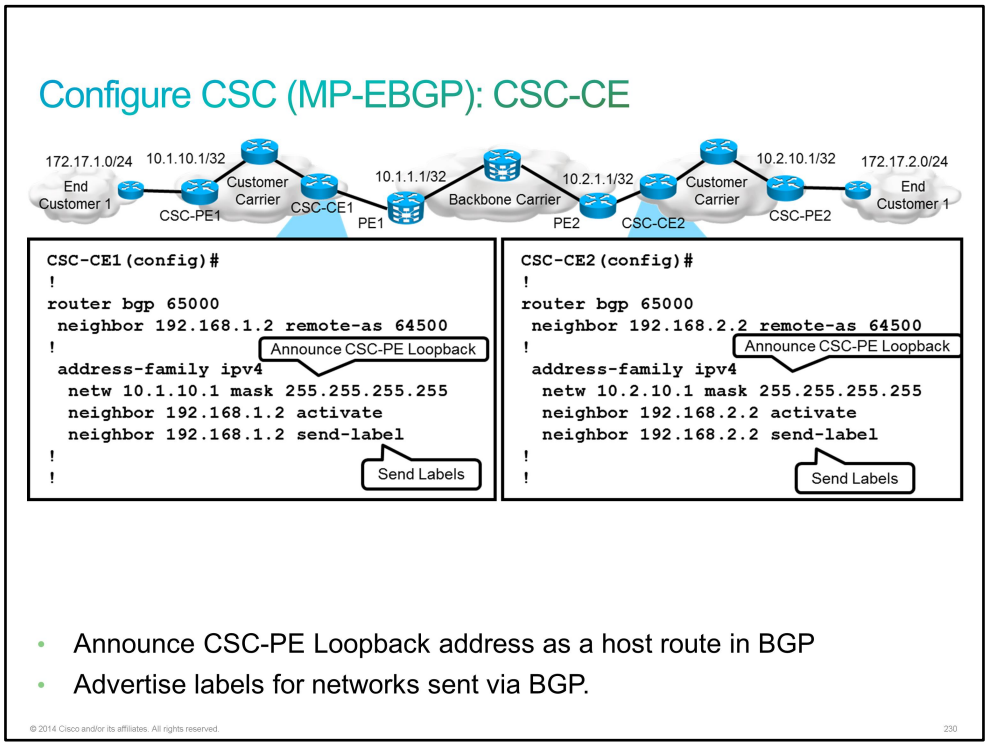
The label stack consists of two labels in the customer carrier network and of three labels in the backbone carrier network.

As illustrated in this example, a customer sends an IP packet. A label stack is attached to the packet. The inner label is a VPN label that identifies packets from the same VPN. The outer label, an LDP label, is attached to packet in order to forward the packet through the customer carrier networks. CSC-CE1 swaps the LDP labels against a MP-EBGP label (BGP2) received from PE1 for the CSC-PE2's Loopback network. PE1 swaps the BGP label against a VPN label of the backbone carrier (VPN1), received via MP-IBGP from PE2, and inserts a third label (LDP) to enable the packet to reach PE2. When the packet reaches PE2, the PE2 swaps the VPN label of the backbone carrier (VPN1) against the BGP label (BGP1) received via MP-EBGP from CSC-CE2. CSC-CE2 swaps this BGP label against an LDP label that identifies the LSP towards CSC-PE2. Due to penultimate hop popping, the packet arrives at CSC-PE2 with a single label: the VPN label. The VPN label identifies the packets from this customer.

https://t.me/learningnets

Configure CSC (MP-EBGP): CSC-CE

The configuration of CSC-CE needs to be adjusted to the CSC deployment method. Instead of MPLS enabled on the link to the PE, MP-EBGP is used to for label signaling.



The CSC-CE routers do not need MPLS LDP or IGP on the links to the PE routers. MP-EBGP is used to signal the labels for the CSC-PE Loopback networks.

To enable MP-EBGP for label signaling, you need configure these two features:

- Announce the CSC-PE Loopback addresses in BGP address-family IPv4. This example assumes that the CSC-CE router has learned the CSC-PE Loopback network via IGP.

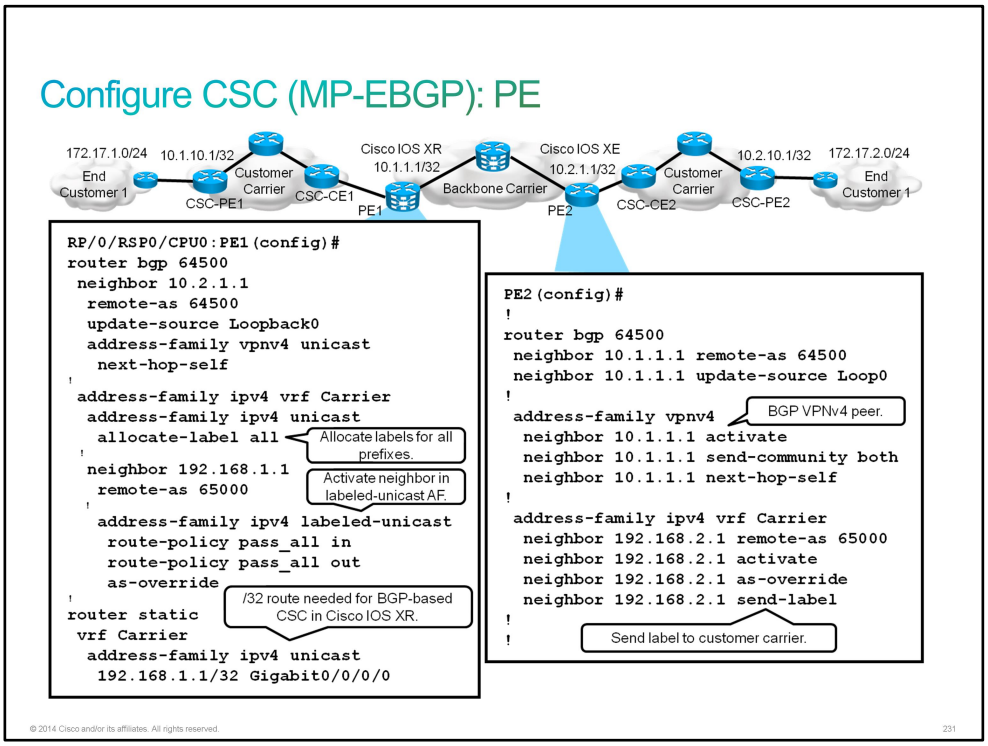
Note	The customer carriers may use alternative variants to configure their routing. For example, internal BGP for IPv4 address family could be used between the CSC routers. With IPv4 IBGP in customer carrier networks, CSC-PE routers could advertise their Loopback networks into BGP. This would remove the need to announce the CSC-PE Loopback networks on the CSC-CE routers.

- Enable label signaling within the MP-EBGP session with the PE router. Use the **neighbor** *neighbor-IP-address* **send-label** command for this purpose.

https://t.me/learningnets

Configure CSC (MP-EBGP): PE

The backbone carrier needs to enable label signaling within the MP-EBGP sessions terminated in the carrier VRFs. The required configuration differs between Cisco IOS XR and Cisco IOS XE.



In Cisco IOS XR Software you need to define the label allocation policy. The label allocation policy is configured in the address-family IPv4/IPv6 unicast submode of the BGP address-family VRF configuration mode. In this example PE1 will allocate labels for all prefixes in the VRF table, as configured with the **allocate-label all** command.

In Cisco IOS XR Software you have to enable the labeled-unicast address family for the MP-EBGP neighbor. You achieve this goal with the **address-family ipv4 labeled-unicast** or **address-family ipv6 labeled-unicast** command. You need to apply a route policy to the MP-EBGP neighbor to allow any updates to be exchanged. The **as-override** command is configured in this example because the customer carrier networks share the same AS number and an update from one customer carrier would never reach another customer carrier with the same AS number.

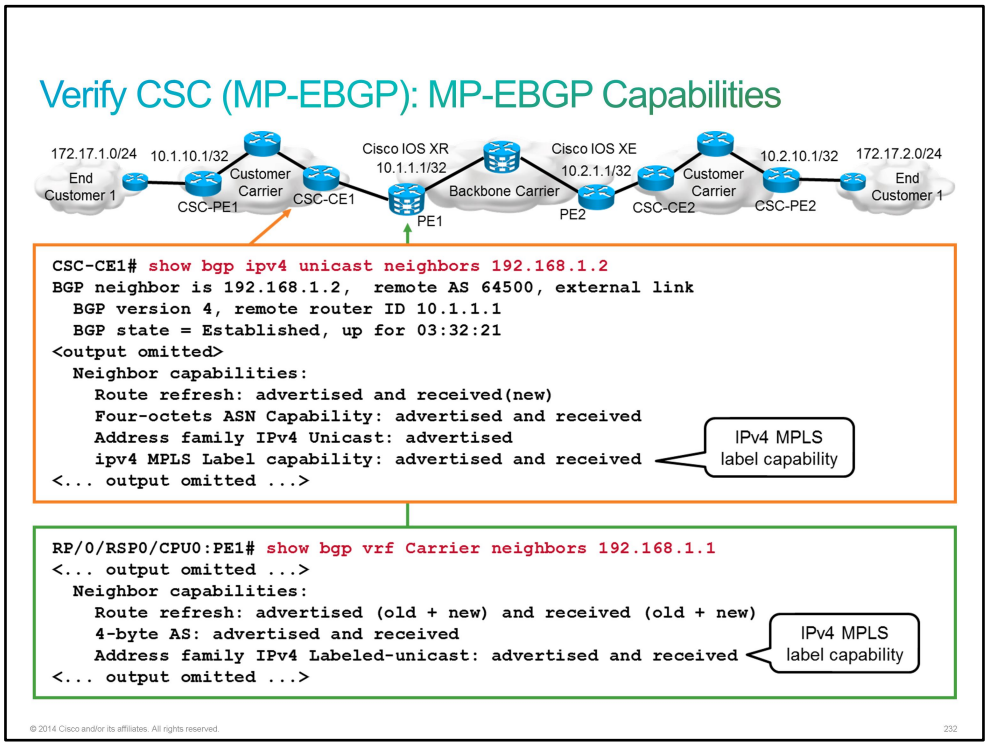
Cisco IOS XR Software requires a /32 host route towards the MP-EBGP neighbor IP address. The required route is configured in the VRF mode of the static routing process.

In Cisco IOS Software, you need to enable LDP signaling within the MP-EBGP session. You configure the feature using the **neighbor** *neighbor-ip-address* **send-label** command under the BGP peer address family configuration mode.

https://t.me/learningnets

Verify CSC (MP-EBGP): MP-EBGP Capabilities

You can start the CSC verification process by checking the capabilities exchanged between the MP-EBGP neighbors.



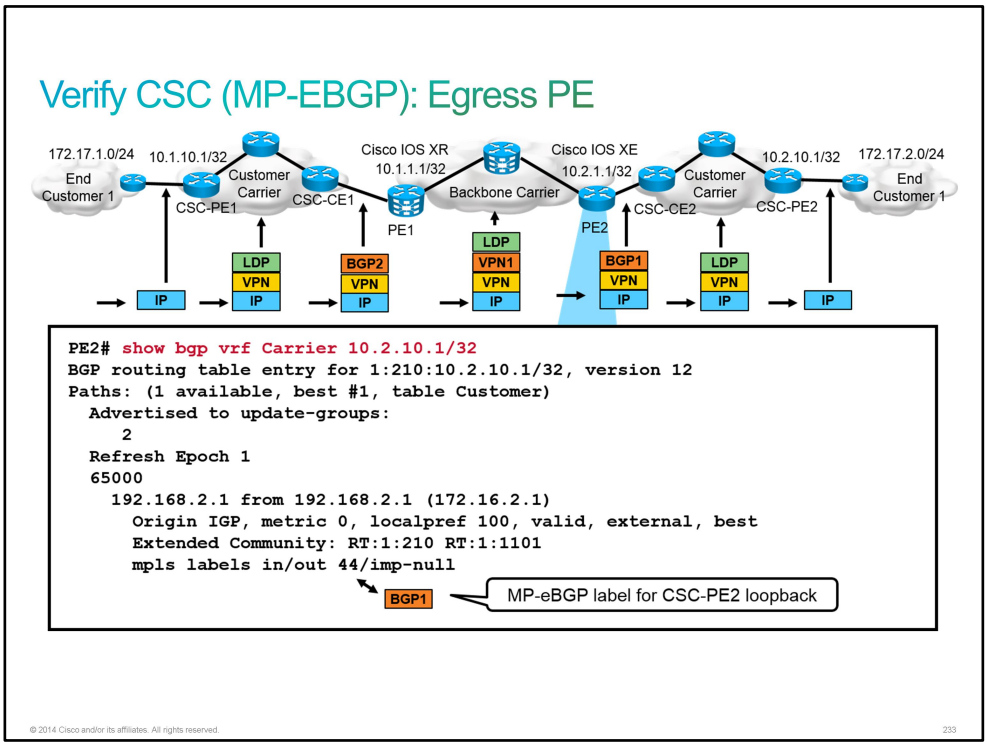
These command outputs describe the capabilities negotiated within the MP-EBGP session between CSC-CE1 (Cisco IOS) and PE1 (Cisco IOS XR).

Both ends report that the MPLS label capability for IPv4 has been successfully negotiated. They are responsible for MPLS label signaling in BGP.

https://t.me/learningnets

Verify CSC (MP-EBGP): Egress PE

MP-EBGP is used to signal the labels for the CSC-PE Loopback networks.



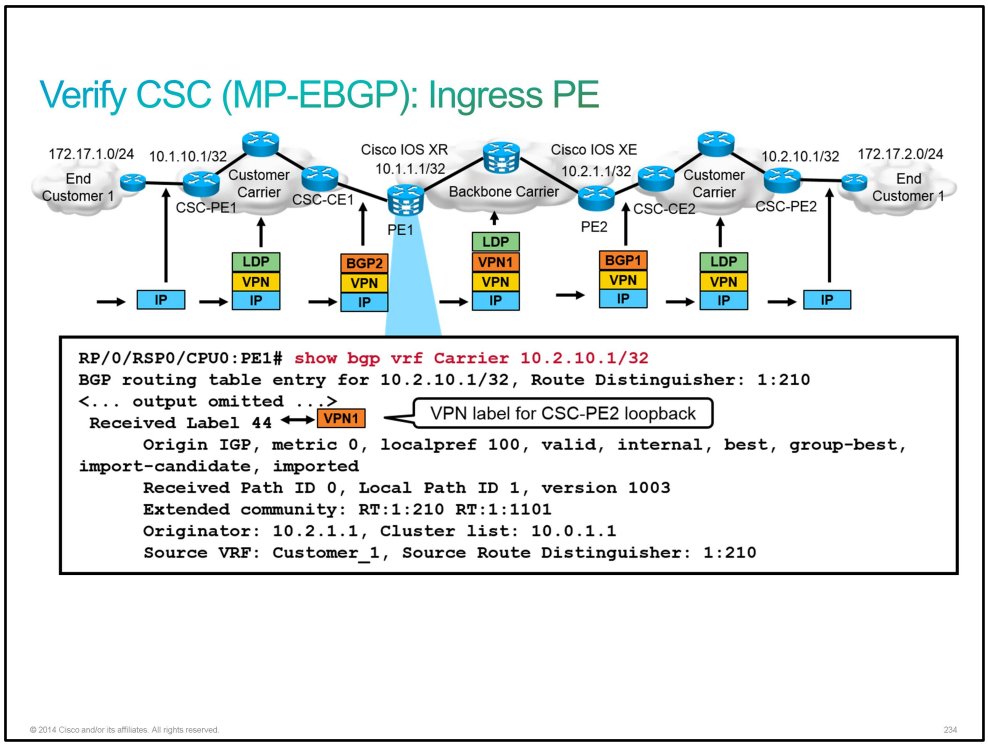
If you want to examine the LSP towards the Loopback address of CSC-PE2, you can view the MPLS label received by PE2 from CSC-CE2. In the data flow diagram this label is depicted as BGP1.

You can display the label by viewing the Loopback network of CSC-PE2 (10.2.10.1/32) in the VRF section of the PE2's BGP table. In this example PE2 has received the MP-EBGP label of 44.

https://t.me/learningnets

Verify CSC (MP-EBGP): Ingress PE

The backbone carrier PEs signal a VPN label for the Loopback network of CSC-PE2.



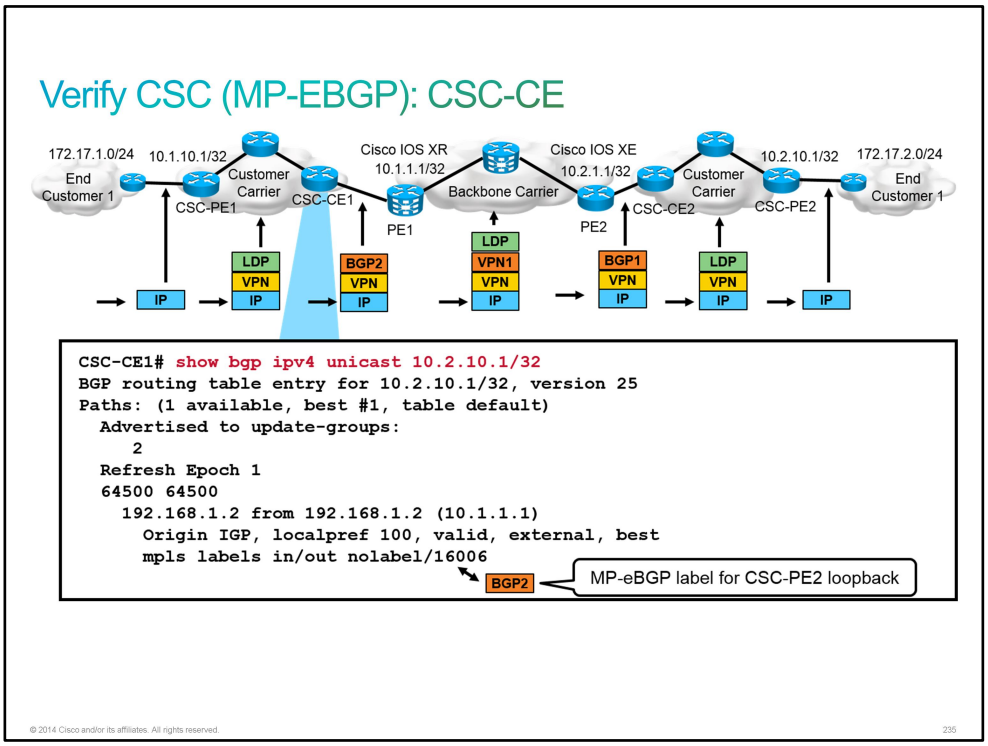
The VPN label exchanged within the backbone carrier network is depicted in the data flow diagram as VPN1.

On PE1, you can display the label by viewing the Loopback network of CSC-PE2 (10.2.10.1/32) in the VRF section of the BGP table. In this example PE1 has received the MP-EBGP label of 44. By default, the VPN label has the same value as the label received by PE2 from CSC-CE2.

https://t.me/learningnets

Verify CSC (MP-EBGP): CSC-CE

PE1 advertises the label to its downstream peer, CSC-CE1. The label is announced via MP-EBGP.



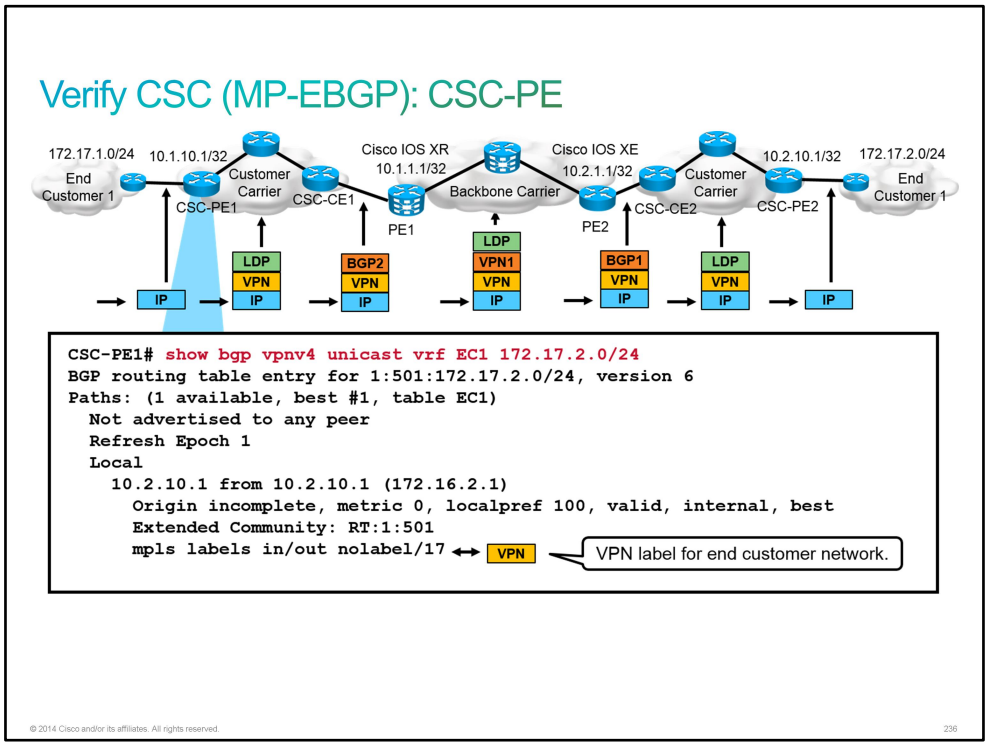
The data flow diagram shows this label as BGP2.

You can display the label by viewing the Loopback network of CSC-PE2 (10.2.10.1/32) in the BGP table of CSC-CE1. In this example the received label has the value of 16006.

https://t.me/learningnets

Verify CSC (MP-EBGP): CSC-PE

You can examine the inner-most VPN label exchanged between the CSC-PE routers for the end customer prefixes.



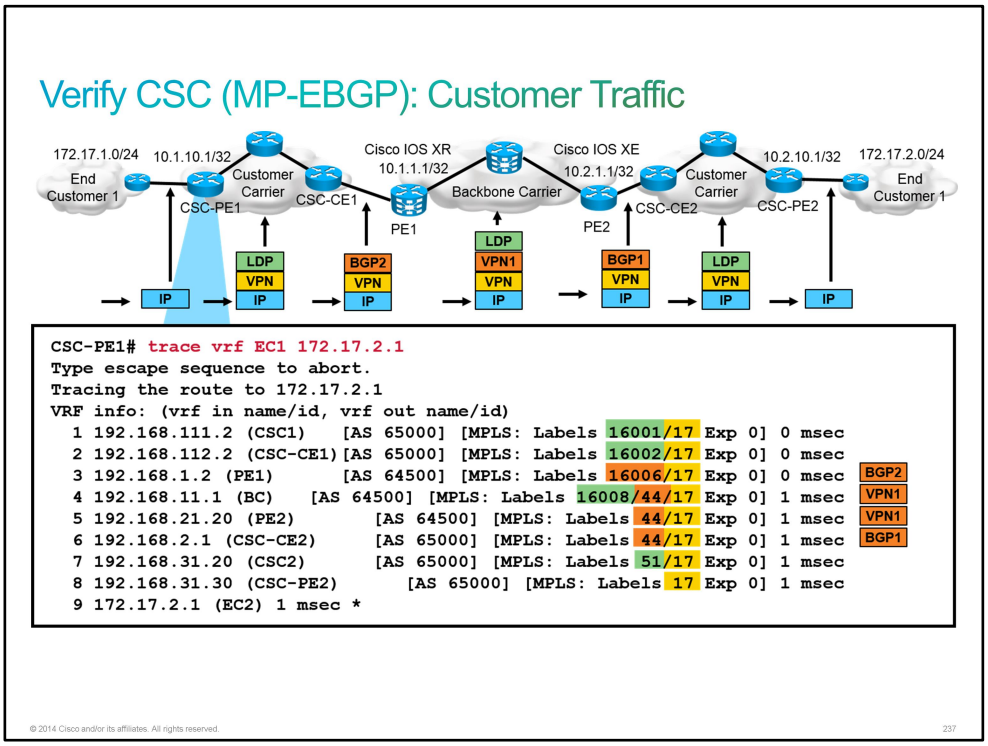
This example illustrates how to examine the end customer prefix 172.17.2.0/24 received by CSC-PE1 from CSC-PE2. The prefix is stored in the VRF section (EC1 VRF) of the BGP table.

The received VPN label is 17.

https://t.me/learningnets

Verify CSC (MP-EBGP): Customer Traffic

You can examine the label stack used on each link by using the traceroute utility from the end customer router or CSC-PE1.



On the link between CSC-PE1 and the internal CSC router, the label stack consists of two labels: the inner most VPN label (17) and the LDP label 16001. The LDP label (16001) is for the Loopback network of CSC-PE2.

The internal CSC router swaps the LDP label 16001 for the Loopback network of CSC-PE2 against 16002. The label stack then consists of two labels: the inner most VPN label (17) and the LDP label 16002.

CSC-CE1 swaps the LDP label against the label received for the Loopback network of CSC-PE2 via MP-EBGP from PE1 (16006). The label stack then consists of two labels: the inner most VPN label (17) and the BGP label 16006.

PE1 swaps the BGP label against the a label stack (VPN label VPN1 + LDP label) that ensures the LSP continuity towards the Loopback network of CSC-PE2. The label stack consists of three labels: the inner most VPN label (17), the VPN label VPN1 (44) and the LDP label 16008.

The internal backbone router is the penultimate hop and pops the LDP label. The label stack consists of two labels: the inner most VPN label (17) and the VPN label VPN1 (44).

PE2 swaps the VPN label for the Loopback network of CSC-PE2 (VPN1) against the BGP label received for that prefix from CSC-CE2. The BGP label happens to have the same value (44). The label stack consists of two labels: the inner most VPN label (17) and the BGP label BGP1 (44).

CSC-CE2 swaps the BGP label against the LDP label for the Loopback network of CSC-PE2. The label stack consists of two labels: the inner most VPN label (17) and the LDP label (51).

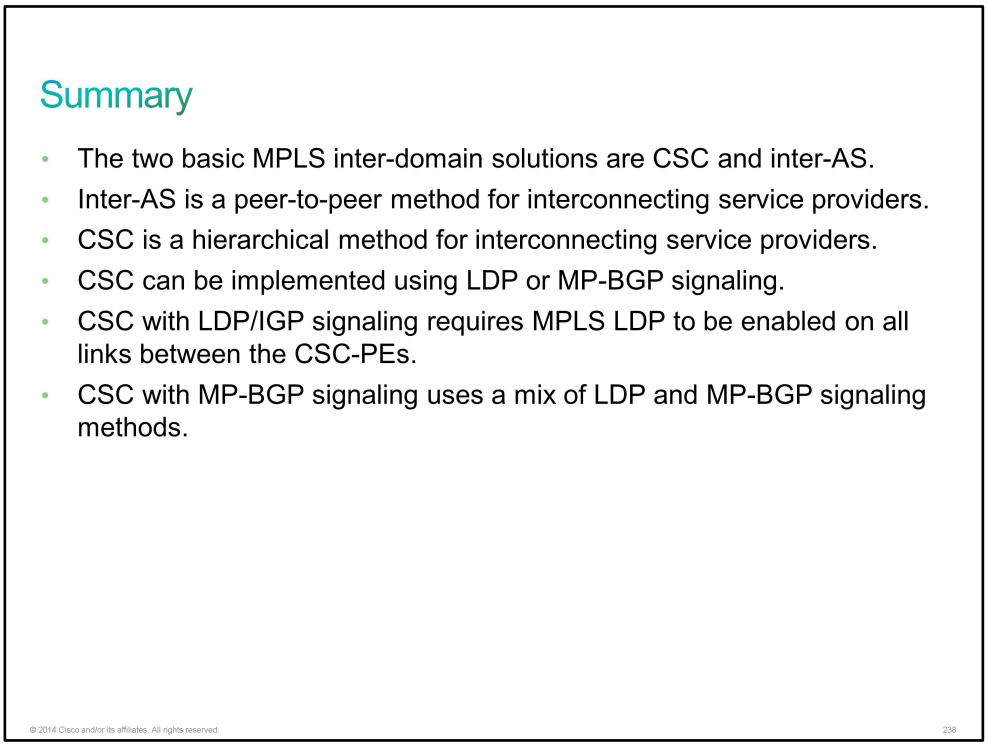
https://t.me/learningnets

The internal customer carrier router is the penultimate hop and pops the LDP label. The packets arrive at CSC-PE2 with only the VPN label (44), which causes the router to forward the packets to the end customer site.

https://t.me/learningnets

Summary

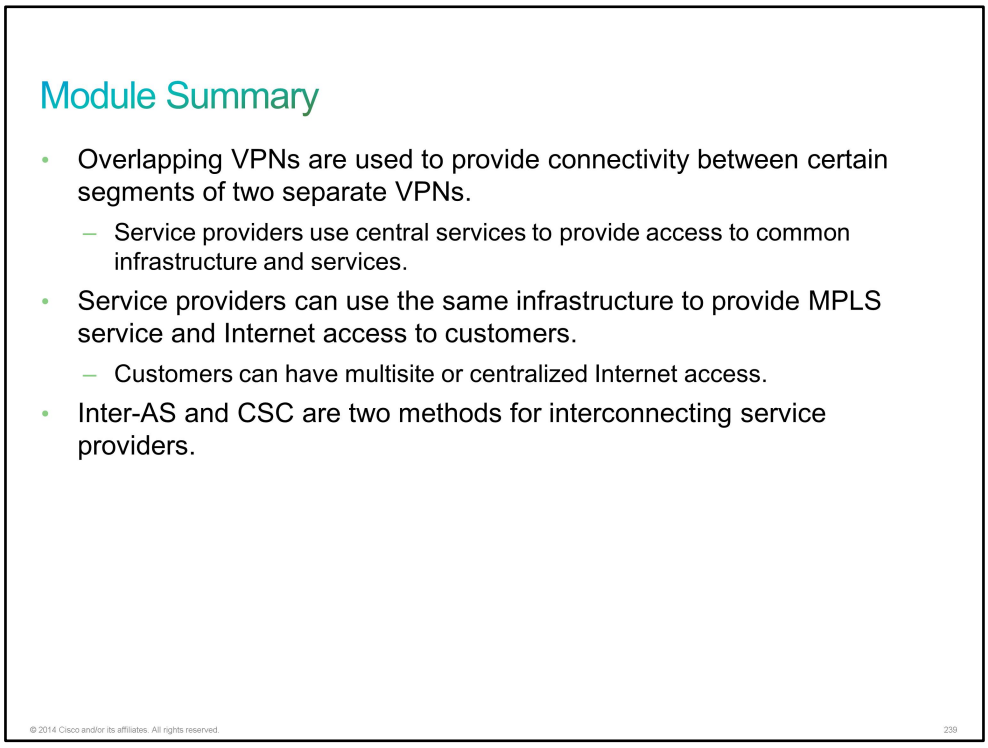
This topic summarizes the primary points that were discussed in this lesson.



https://t.me/learningnets

Lesson 3: Module Summary

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



https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: 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. No tunnel or circuit configuration is required when using Cisco 6PE for IPv6 integration. (Source: Deploying IPv6 in MPLS Environment)

   A. True
   B. False

2. Which next-hop attribute do the 6VPE routers use for vpnv6 routes? (Source: Deploying IPv6 in MPLS Environments)

   A. Global IPv6 address
   B. Link-local IPv6 address
   C. IPv4-compatible IPv6 address
   D. IPv4-mapped IPv6 address
   E. IPv4 address

3. The IPv6 BGP updates in a 6PE environment are referred to as? (Source: Deploying IPv6 in MPLS Environments)

   A. VPNv4 routes
   B. VPNv6 routes
   C. IPv4 routes
   D. IPv6 routes
   E. IPv6-in-IPv4 routes

https://t.me/learningnets

4. Which next-hop attribute do the 6PE routers use for ipv6 routes? (Source: Deploying IPv6 in MPLS Environments)

   A. Global IPv6 address
   B. Link-local IPv6 address
   C. IPv4-compatible IPv6 address
   D. IPv4-mapped IPv6 address
   E. IPv4 address

5. Which two statements about the CSC model are true? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

   A. CSC is a peer-to-peer model.
   B. CSC is a hierarchical model.
   C. Multiple customer carriers can be connected to a single CSC backbone.
   D. End customer routes are inserted in the CSC backbone.

6. Which two protocols can be used to exchange label and route information between the backbone and customer carriers? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

   A. CDP
   B. IGP and LDP
   C. MP-BGP
   D. NTP
   E. L2TPv3

7. Which two statements about the inter-AS model are true? (Choose two.) (Source: Introducing MPLS Interdomain Solutions)

   A. Inter-AS is a peer-to-peer model.
   B. Inter-AS is a hierarchical model.
   C. The back-to-back VRF method is suitable for small environments but is not a scalable solution.
   D. In the multihop EBGP method, the ASBR router contains customer routes.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. A
2. E
3. D
4. E
5. B, C
6. B, C
7. A, C

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 5: Layer 2 VPNs

Introduction

VPWS enables geographically separated sites to be interconnected over a virtual, point-to-point Layer 2 circuit. The virtual connection can cross an MPLS or IP network. VPLS enables remote LAN segments to be linked as a single bridged domain over an MPLS network. The full functions of the traditional LAN, such as MAC address learning, aging, and switching, are emulated across all of the remotely connected LAN segments that are part of a single, bridged domain.

A service provider can offer VPLS service to multiple customers over the MPLS network by defining different bridged domains for different customers. Packets from one bridged domain are never carried over or delivered to another bridged domain, thus ensuring the privacy of the LAN service. VPLS transports Ethernet 802.3, VLAN 802.1Q, and VLAN-in-VLAN (QinQ) traffic across multiple sites that belong to the same Layer 2 broadcast domain. VPLS offers simple VLAN services that include flooding broadcast, multicast, and unknown unicast frames that are received on a bridge. The VPLS solution requires a full mesh of PWs that are established among PE routers. The VPLS implementation is based on LDP-based PW signaling.

Upon completing this module, you will be able to:

- Describe Layer 2 VPNs that are available in the MPLS and IP core
- Implement Any Transport over MPLS
- Implement Virtual Private LAN Service
- Enhance the VPLS scalability

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Introducing Layer 2 VPNs

Overview

There is an ever-increasing demand for the transport of Layer 2 and Layer 3 over a common backbone. This lesson introduces the VPWS and VPLS. VPWS offers point-to-point virtual connections, while VPLS provides LAN-similar multipoint connectivity. A subset of VPWS is AToM. AToM allows an MPLS network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, HDLC, Frame Relay, and ATM.

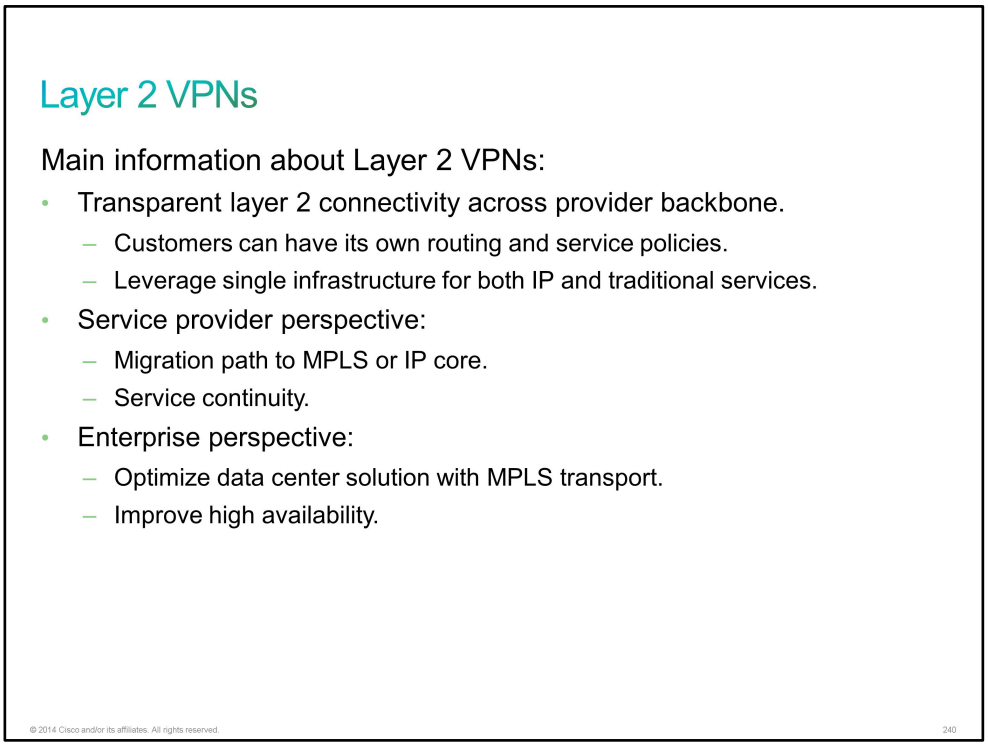
Upon completing this lesson, you will be able to:

- Describe the need for Layer 2 VPNs
- Describe the benefits of Layer 2 VPNs
- Show the different types of Layer 2 VPNs
- Describe the components of a Layer 2 VPN
- Describe the Layer 2 VPN control plane
- Describe the Layer 2 VPN data plane when using an IP core
- Describe the Layer 2 VPN data plane when using an MPLS core
- Describe the two Layer 2 transport services (VPWS and VPLS)
- Describe an AToM example using EoMPLS
- Describe AToM interworking
- Describe the two Layer 2 VPN interworking modes: bridged and routed
- Describe the MEF E-Line, E-LAN and E-Tree services

https://t.me/learningnets

Layer 2 VPNs

Layer 2 VPNs allow both service providers and enterprise to build a single infrastructure for both IP and traditional services. The service providers can migrate their existing ATM and Frame Relay traffic to MPLS or IP core without interrupting the customer service. The enterprises can leverage the extended Layer 2 domain to optimize the data center solution. The MPLS transport enables a host of additional high availability extensions.



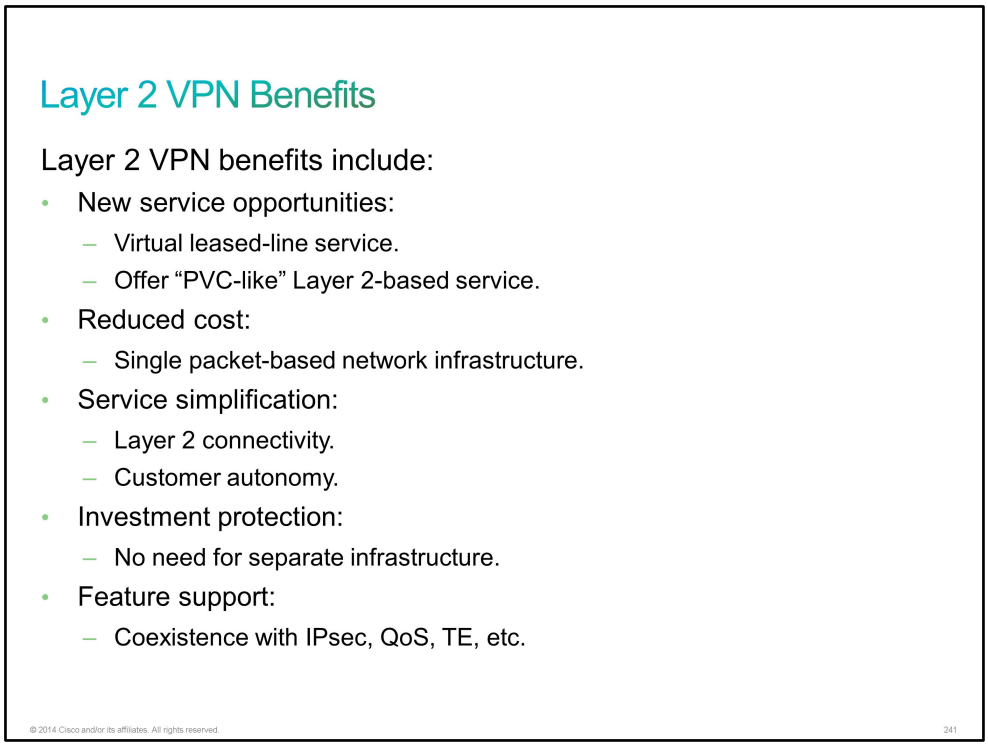
The Layer 2 VPN is a tunneling technology that provides logical separation between the customer and service provider domains, including segmentation of routing, QoS policy, and others.

Layer 2 VPNs also represent an easy-to-implement migration step toward a Layer 3 IP and MPLS VPN system.

https://t.me/learningnets

Layer 2 VPN Benefits

Layer 2 VPNs pave way for a number of business opportunities.



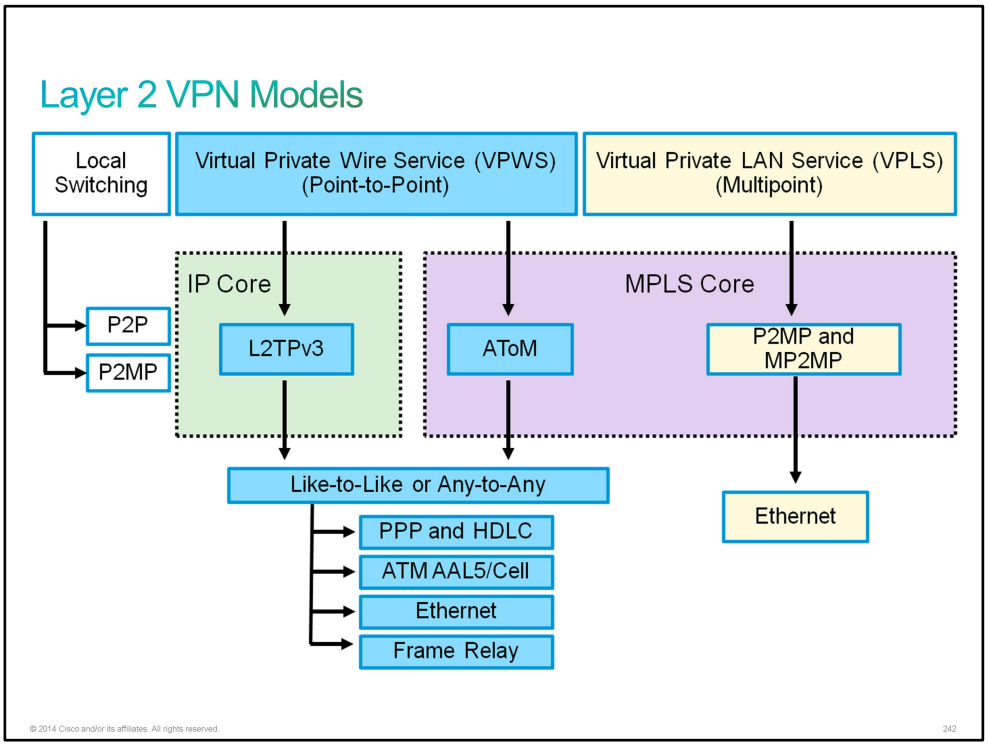
Layer 2 VPN technologies offer a range of benefits to service providers and enterprises:

- **New service opportunities:** The virtual leased lines offer connectivity service that resembles the traditional costly permanent virtual circuits in Frame Relay or ATM environments.
- **Cost savings:** The consolidation of multiple core technologies into a single packet-based network infrastructure lowers the overall cost of system installation and maintenance.
- **Simple connectivity model:** Layer 2 transport provides options for service providers that need to provide Layer 2 connectivity and maintain customer autonomy.
- **Investment protection:** Service providers can extend customer access to existing Layer 2 networks without deploying a new separate infrastructure.
- **Feature support:** Through the use of Cisco IOS and IOS XR features such as IPsec, QoS, and TE, Layer 2 transport can be tailored to meet customer requirements.

https://t.me/learningnets

Layer 2 VPN Models

Layer 2 VPNs are grouped into three main categories: local switching that serves directly connected links, VPWS that provides point-point connectivity and VPLS that enables multipoint connection service.



VPWS can be implemented using L2TPv3, when run in an IP environment, or AToM when deployed in a MPLS core.

VPWS is a point-to-point technology. VPWS supports connections between the same interface types (like-to-like), and between different interface types (any-to-any). The supported attached interface types include Ethernet, Frame Relay, and ATM, including AAL5, PPP, and HDLC.

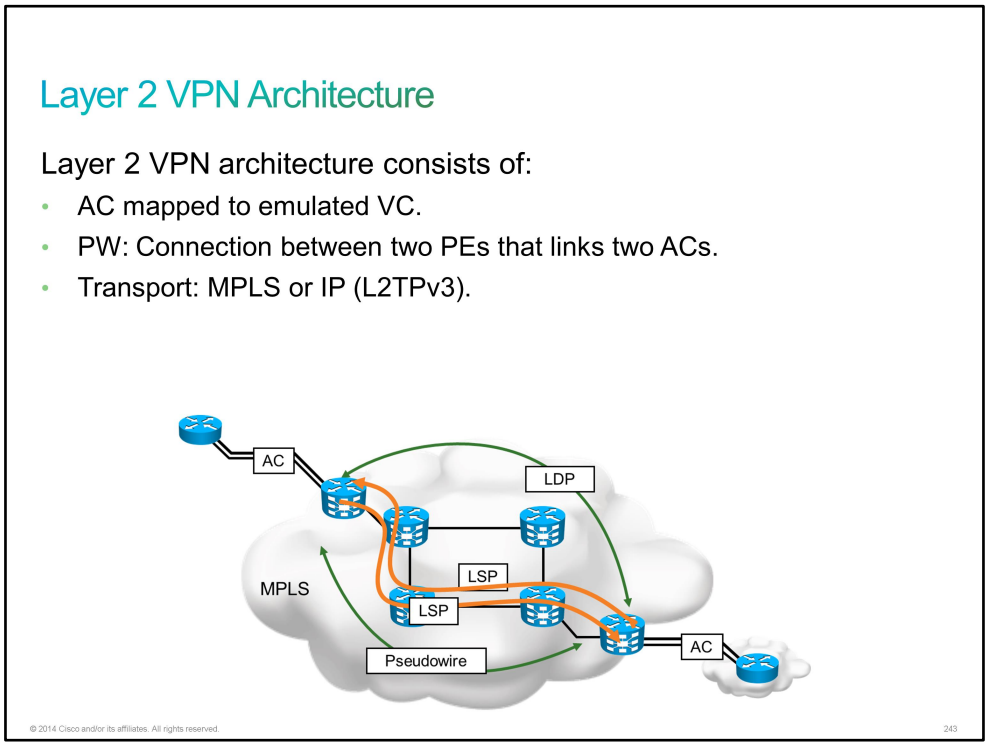
VPLS offers point-to-multipoint and multipoint-to-multipoint connectivity. It leverages MPLS as the transport infrastructure.

IP core transport uses the L2TPv3, and supports only point-to-point connections. The available encapsulations include Ethernet, Frame Relay, and ATM, including AAL5, PPP, and HDLC. They can be linked in any-to-any fashion.

https://t.me/learningnets

Layer 2 VPN Architecture

Layer 2 VPNs generally comprise three elements: attachment circuits, pseudowire(s), and the transport technology.



The Layer 2 VPN solution encompasses three main elements:

- An attachment circuit is the circuit or link directly connected to the PE system that is virtually extended to the other end of the provider cloud. It can represent a Frame Relay DLCI, PPP and HDLC link, ATM VPIand VCI pair, Ethernet port, or VLAN. The attachment circuit is mapped to the emulated VC for transport through the service provider core.
- PW is a point-to-point connection over a packet-switching network. The PW emulates the operation of a "transparent wire" by linking two distant attachment circuits attached to two different PE routers.
- Transport infrastructure is the third element. The transport network can either be MPLS-enabled, and thus capable of supporting any Layer 2 topology (point-to-point, point-to-multipoint, or multipoint-to-multipoint), or IP-based. In the IP core, the L2TP protocol supports point-to-point connections only.

Frames are received on an ingress interface by the ingress PE router. At this point, the frame is a raw Layer 2 frame. In the case of MPLS transport, the ingress PE router encapsulates it into MPLS and tunnels it across the backbone to the egress PE router. The egress PE router decapsulates the packet and reproduces the raw Layer 2 frame on the egress interface.

https://t.me/learningnets

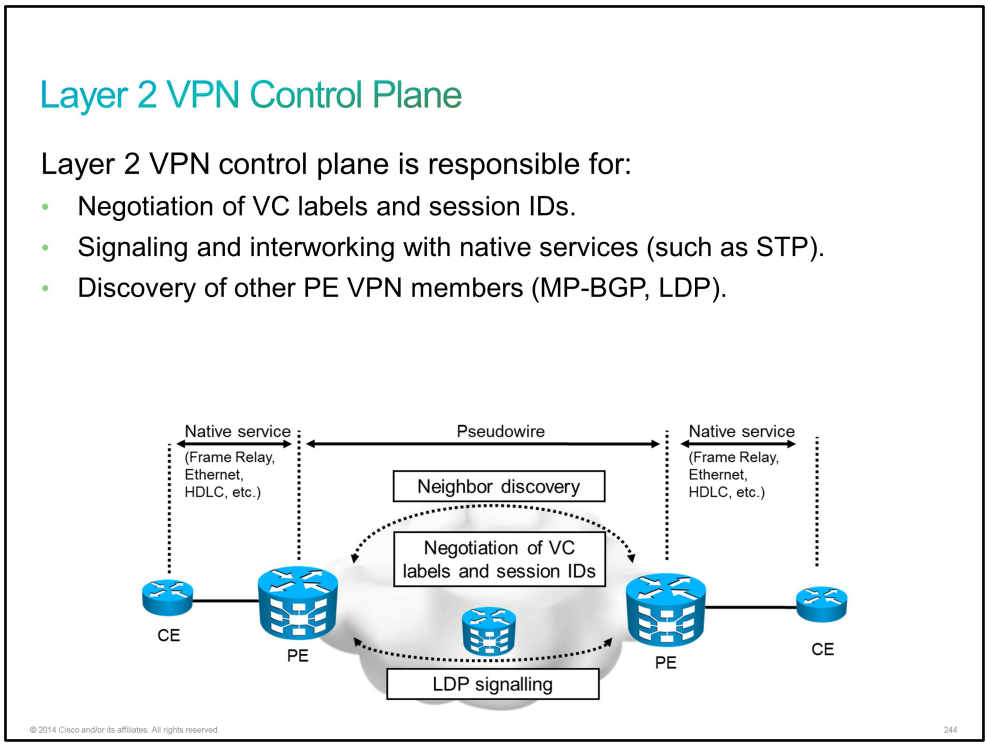
The frames are carried across the MPLS backbone using a label stack of two labels. The top label is used to propagate the packet from the ingress PE router to the correct egress PE router. The second label is used by the egress PE router to forward out the packet on the correct interface. This process is somewhat similar to an MPLS VPN, where the egress PE router uses a VPN label. The top label is called the tunnel label. This name indicates that its use is to tunnel the packet across the MPLS backbone to the egress PE router. The second label is called the VC label. The name indicates that its use is to map the packet to an outgoing VC or link.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Layer 2 VPN Control Plane

A protocol is required between the PE routers so that they can exchange the VC information.



In the case of MPLS transport, the LDP is used for this purpose. A directed multihop LDP session is established between the PE routers. The egress PE router sends an LDP message in which it indicates the label value to use for a VC FEC. That label value is then used by the ingress PE router as the second label in the label stack that is imposed to the frames of the indicated VC FEC.

In the case of IP-based transport, the L2TPv3 session exchanges session parameters and not labels.

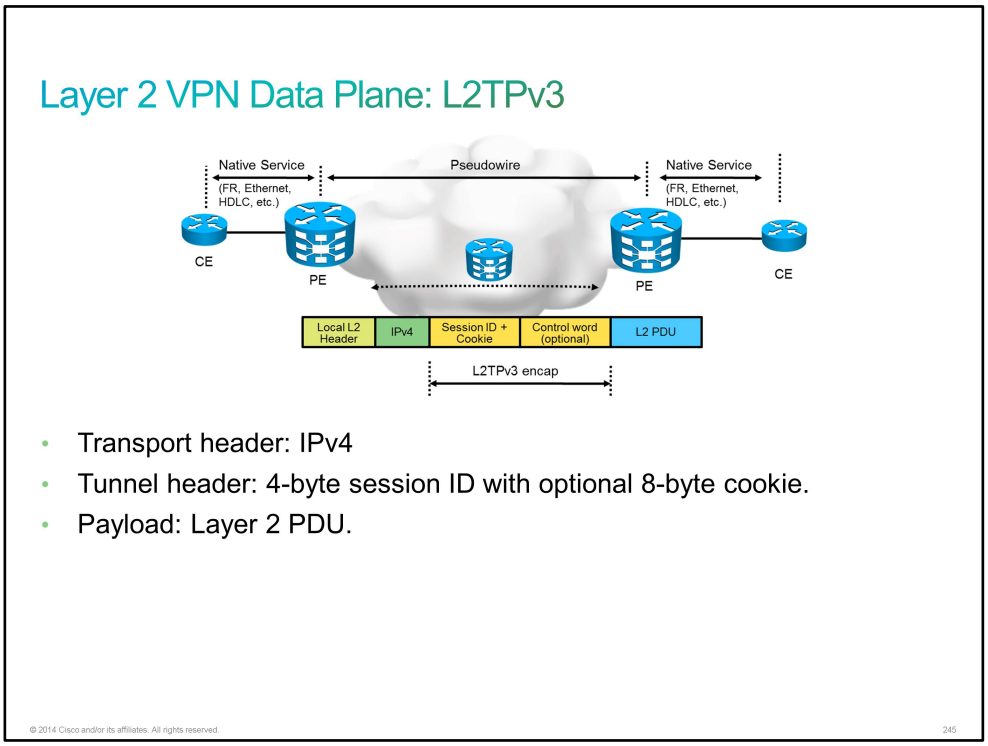
The figure shows a directed multihop LDP session between the ingress and egress PE routers that is used to exchange the VC label. Any ingress-egress PE router pair will need such an LDP session.

The control session is also responsible for providing interworking capabilities with native services, such as Frame RelayLMI. The neighbors can be either statically configured or auto-discovered using MP-BGP extensions or LDP.

https://t.me/learningnets

Layer 2 VPN Data Plane: L2TPv3

L2TPv3 is used to transport Layer 2 frames over pure IP networks.



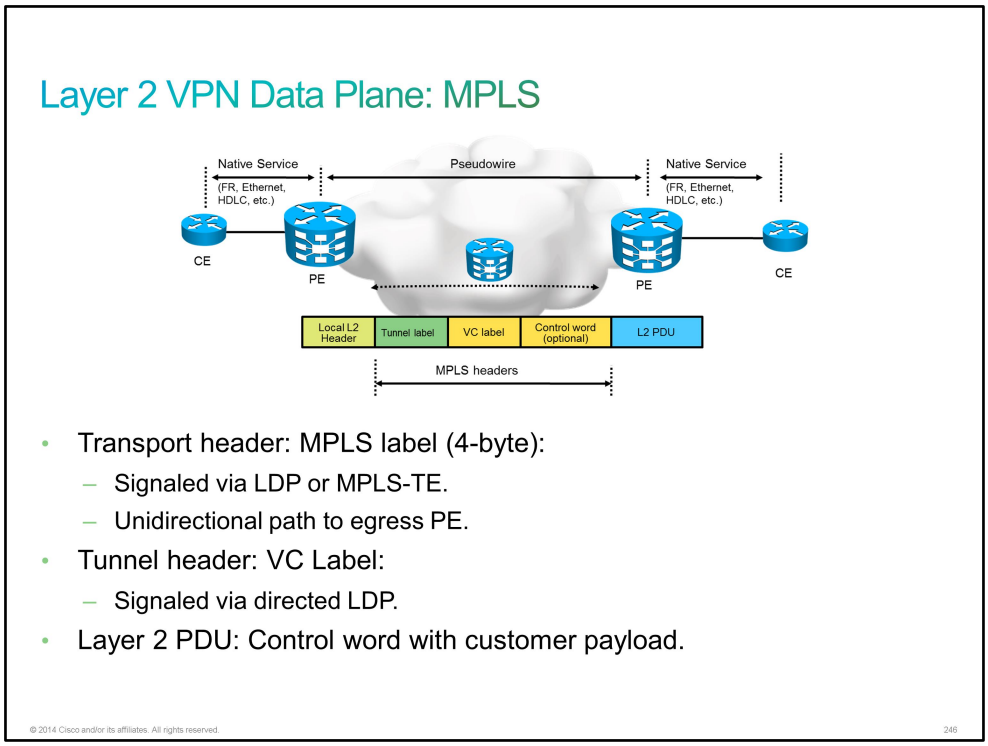
The entire L2TP packet, including payload and L2TP header, is sent within a UDP datagram. Traditionally, L2TP has been used to carry PPP sessions within an L2TP tunnel. L2TPv3 provides additional security features, improved encapsulation, and the ability to carry data links other than just PPP over an IP network (for example, Frame Relay, Ethernet, ATM, and others).

L2TP overhead includes the transport IP header (20 bytes) and an L2TP header of variable length. The only mandatory field in the L2TP header is the session ID (4 bytes). Optional fields are cookie (8 bytes) and control word. The payload of the L2TPv3 packet is the original Layer 2 PDU.

https://t.me/learningnets

Layer 2 VPN Data Plane: MPLS

MPLS transport is based on the same mechanism that you examined for Layer 3 MPLS VPNs.



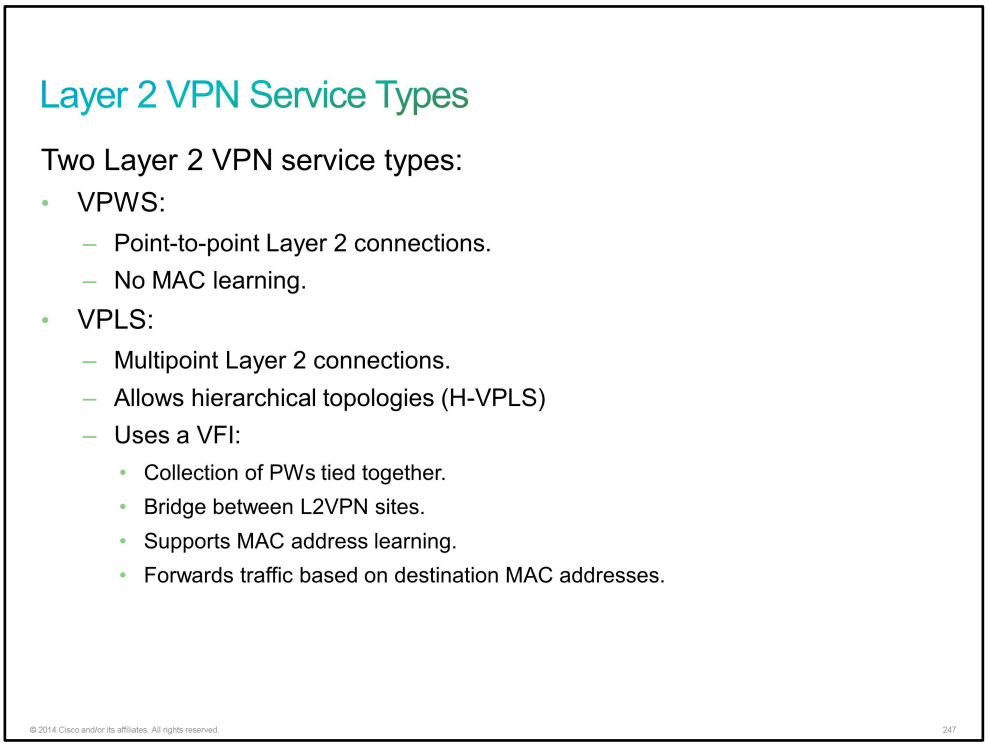
Hop-by-hop LDP signals a unidirectional path to the egress PE router. A directed LDP session exchanges the VC label that serves as the inner label in the MPLS packet. In general, the use of control words is optional, and configurable. Some transport types, such as Frame Relay Layer 2 VPN, require the use of control words.

Depending on the frame encapsulation, some fields of the original frame, such as checksums, may be stripped before MPLS encapsulation.

https://t.me/learningnets

Layer 2 VPN Service Types

There are two major Layer 2 VPN service types.



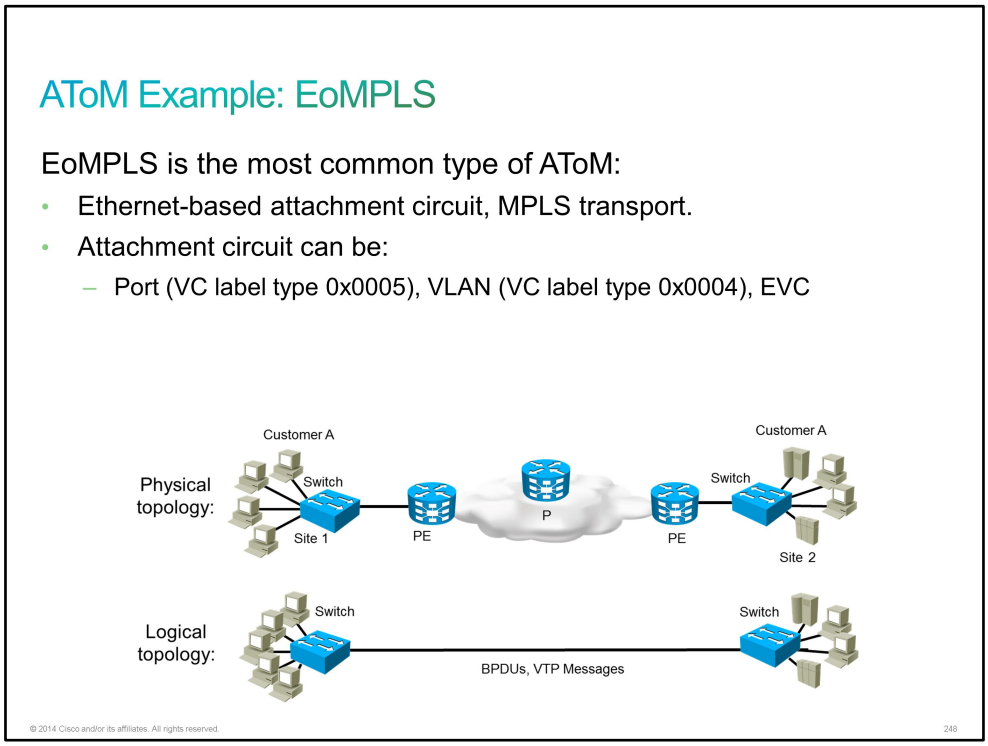
Layer 2 transport services are categorized into two major classes:

- **VPWS:** This service type is used to deploy point-to-point Layer 2 connections. It does not involve MAC learning capabilities. It encompasses two transport methods:
  - L2TPv3
  - AToM. The most common example of AToM is EoMPLS.
- **VPLS.** This technology supports multipoint Layer 2 connections by grouping a collection of PWs terminated on a PE router in a VFI. The VFI represents a virtual extension of the physical circuit attached to the PE system. The VFI resembles a switch that is capable of learning MAC addresses and forwards traffic based on its MAC address table. A VPLS can connect thousands of PEs into a single VLAN and is therefore subject to scalability constraints. To improve the scalability of the solution, H-VPLS topologies enable a two-tier deployment of the PE devices.

https://t.me/learningnets

AToM Example: EoMPLS

EoMPLS is the most ubiquitous example of AToM, which transports Ethernet frames over an MPLS infrastructure.



The virtual link bridging the Ethernet segments at both ends is transparent to SPT BPDUs, VTP packets, and other control messages.

The attachment circuit can be the Ethernet port or 802.1Q subinterface (VLAN). For each attachment circuit, LDP signals a different VC type via the targeted LDP session. VC type 5 is used for Ethernet port mode and VC type 4 is used for Ethernet VLAN mode.

In Ethernet port mode, both ends of PW are connected to Ethernet ports. In this mode, the port is tunneled over PW or, using local switching (also known as attachment circuit-to-attachment circuit cross-connect) switches packets or frames from one attachment circuit to another attached to the same PE node. In Ethernet port mode, the PW is always a type 5 virtual connection. On the ingress PE, the network service provider passes the packets to the PW termination point, adds the MPLS labels to the packets, and sends the packets over the PW. In Ethernet port mode, a VLAN header may or may not be present in the frame. In any case, the PE router carries the frame transparently. This allows an Ethernet trunk to be carried over a single PW.

VLAN mode provides Ethernet VLAN-to-VLAN connectivity. In VLAN mode, each VLAN on a customer-end to provider-end link can be configured as a separate Layer 2 VPN connection, using either virtual connection type 4 or type 5.

Virtual connection type 5 is the default mode. In type 4 virtual connections, on the ingress PE, the VLAN tag maps to a particular PW and the packet is placed on the PW with the VLAN tag untouched.

In Type 5 virtual connections, on the ingress PE that is receiving packets from the CE, the network service provider strips off the customer edge VLAN tag before placing the packets on the PW. On the egress PE, the network service provider pushes the VLAN tag onto the protocol stack before it sends the packet to the CE.

https://t.me/learningnets

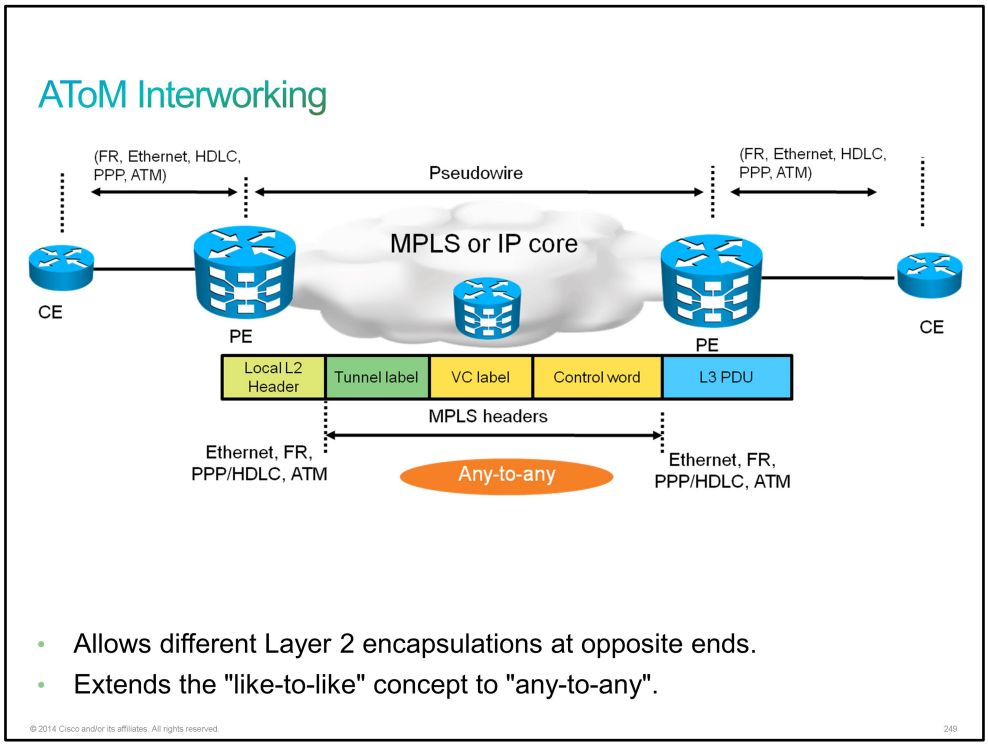
EVCs define a Layer 2 bridging architecture that supports Ethernet services. As an association between two or more user network interfaces, it identifies a point-to-point or multipoint-to-multipoint path within the service provider network. An EVC is a conceptual service pipe within the service provider network.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

AToM Interworking

AToM interworking allows you to provide connectivity between disparate technologies in the Layer 2 VPN sites.



Layer 2 VPN interworking allows you to connect disparate attachment circuits. Cisco routers support these any-to-any combinations:

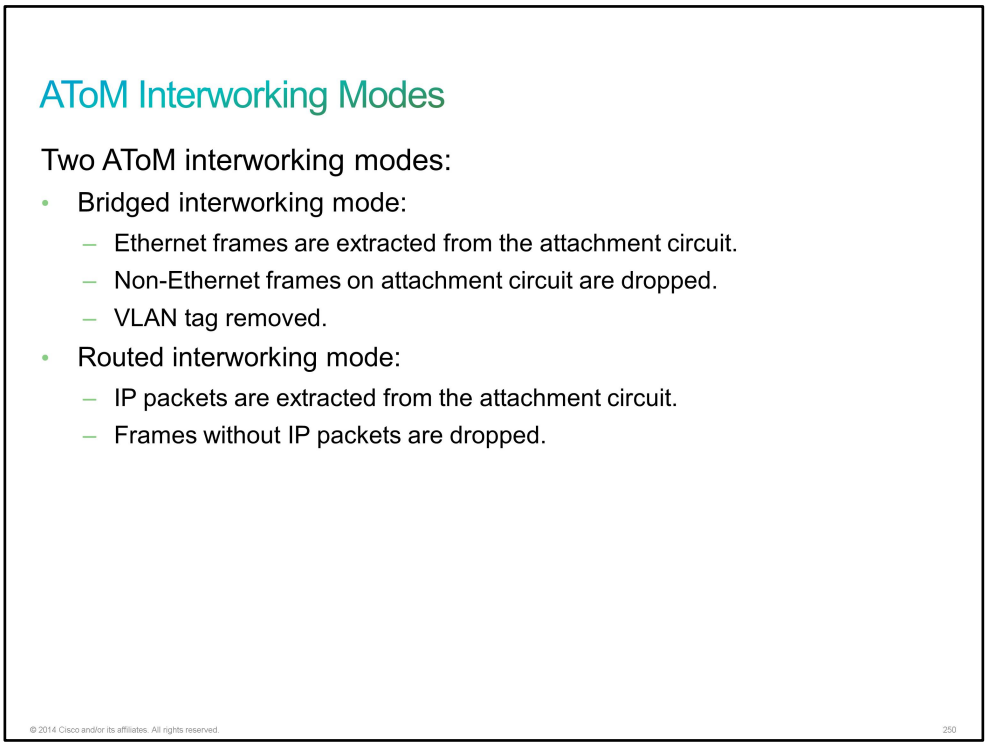
- Ethernet or VLAN to ATM AAL5 interworking.
- Ethernet or VLAN to Frame Relay interworking.
- Ethernet or VLAN to PPP interworking.
- Ethernet to VLAN interworking.
- Frame Relay to ATM AAL5 interworking.
- Frame Relay to PPP interworking.
- Ethernet or VLAN to ATM VPI and VCI interworking.

https://t.me/learningnets

AToM Interworking Modes

The Layer 2 VPN interworking function is implemented in one of two modes:

- Bridged interworking mode.
- Routed interworking mode.



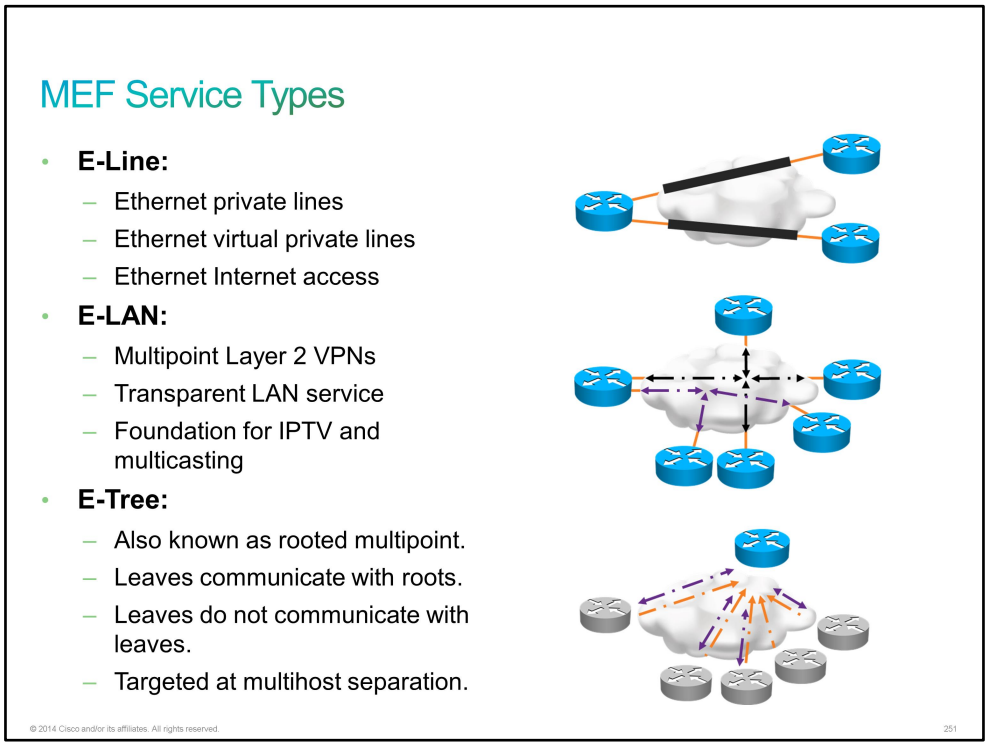
In bridged interworking mode, Ethernet frames are extracted from the attachment circuit and sent over the PW. Attachment circuit frames that are not Ethernet are dropped. In the case of a VLAN, the VLAN tag is removed, leaving an untagged Ethernet frame. This interworking functionality is implemented by configuring the **interworking ethernet** command under the PW class configuration mode.

In routed interworking, IP packets are extracted from the attachment circuit and sent over the PW. Attachment circuit frames are dropped if they do not contain the IPv4 packets. This interworking functionality is implemented by configuring the **interworking ip** command under the PW class configuration mode. The IP interworking mode is used to provide IP connectivity between sites, regardless of the Layer 2 connectivity to these sites. It is different from a Layer 3 VPN because it is point-to-point in nature and the service provider does not maintain any routing information pertaining to customers.

https://t.me/learningnets

MEF Service Types

To bring about the industry acceptance of Ethernet services, it was necessary to clarify and standardize the service range.



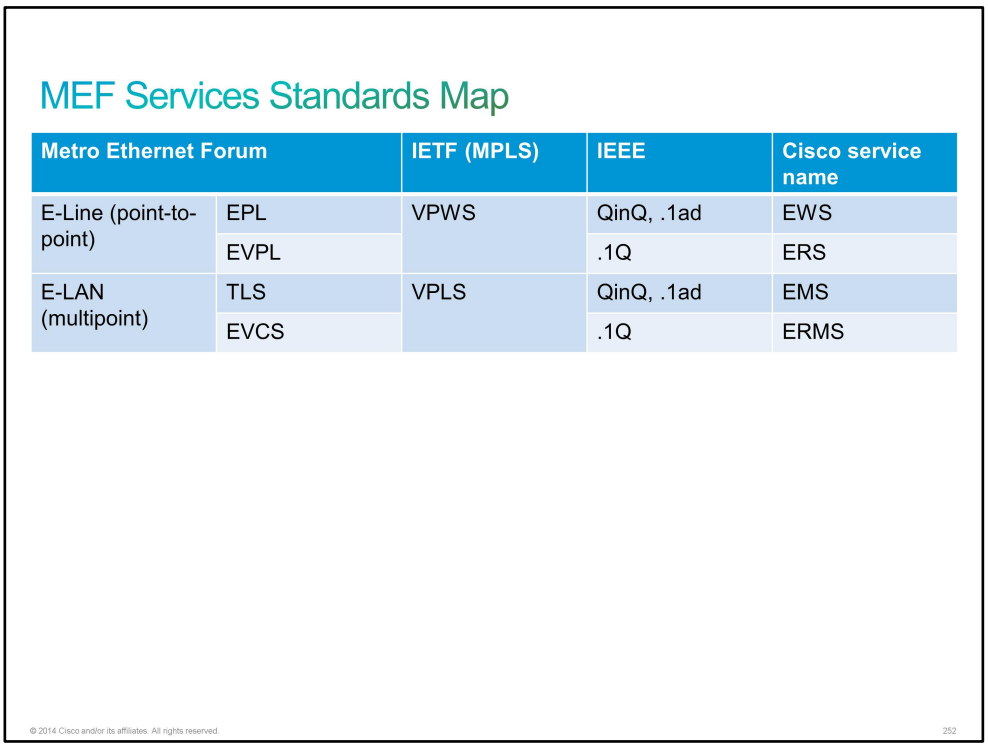
Recognizing this requirement, the industry created the MEF, which played a key role in defining the three major services:

- **E-Line**: A service connecting two customer Ethernet ports over a WAN. This service is further subdivided into EPLs, EVPLs, and EIA.
- **E-LAN**: A multipoint service connecting a set of customer endpoints, giving the appearance to the customer of a bridged Ethernet network connecting the sites. This transparent LAN service is often referred to as a multipoint Layer 2 VPN. It lays the foundation for IPTV and multicasting applications.
- **E-Tree**: A multipoint service connecting one or more roots and a set of leaves, but preventing interleaf communication. This service is also known as rooted multipoint. Specifically, the leaves can communicate with one or more roots, but not with other leaves. The service provides an ideal mechanism for multihost separation. It is considered a major enabler for mobile backhaul and triple-play infrastructure.

All these services provide standard definitions of such characteristics as bandwidth, resilience, and service multiplexing, allowing customers to compare service offerings and facilitating SLAs.

https://t.me/learningnets

MEF Services Standards Map



MEF does not define only the three major categories (E-Line, E-LAN, and E-Tree), but many variants of them. One classification of the variants is based on the EVC mode: port- or VLAN-based.

In addition, the terms E-line, E-LAN, and E-Tree are not the only ones used in the industry. IETF refers to these same services as VPWS and VPLS. The IETF naming is used predominantly throughout this course.

https://t.me/learningnets

Summary

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

Summary

- Layer 2 VPN allows both service providers and enterprise to build a single infrastructure for both IP and traditional services.
- Layer 2 VPNs are grouped into three main categories: local switching that serves directly connected links, methods for transport in the MPLS core, and transport solutions for IP core.
- Pseudowire is a point-to-point connection over a packet-switching network.
- A protocol is required between the PE routers so that they can exchange the VC information.
- L2TPv3 is used to transport Layer 2 frames over pure IP networks.
- A directed LDP session exchanges the VC label that serves as the inner label in the MPLS packet.
- Layer 2 transport services are categorized into two major classes: VPWS and VPLS.

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

Summary (Cont.)

- EoMPLS is the most ubiquitous example of AToM, which transports Ethernet frames over an MPLS infrastructure.
- Layer 2 VPN interworking allows you to connect disparate attachment circuits.
- The L2VPN interworking function is implemented in two modes: Bridged interworking mode or routed interworking mode.
- VPLS creates one or more LANs for each customer who is using the service from the service provider.
- The industry created the Metro Ethernet Forum, which defined the three major services
- Many standardization bodies have issued standards of the Carrier Ethernet architecture, one of them is IEEE

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

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Implementing AToM

Overview

AToM is a subset of VPWS that provides point-to-point virtual connections. AToM allows an MPLS network to provide end-to-end transport for Layer 2 frames and cells. It provides support for Ethernet, PPP, HDLC, Frame Relay, and ATM. This lesson describes AToM characteristics, implementation, and verification.

Upon completing this lesson, you will be able to:

- List the AToM provisioning, forwarding, signaling, and MTU requirements
- Describe the AToM provisioning steps
- Describe the AToM LSP
- Describe the AToM signaling process
- Describe the VCCV OAM feature
- Describe the AToM MTU requirements
- Describe the optional AToM control word field
- Describe the PW stitching concept
- Describe the EoMPLS transport format
- Describe EoMPLS redundancy using PW redundancy
- Describe one-way EoMPLS redundancy
- Describe two-way EoMPLS redundancy
- Provide a high-level overview of the Cisco EVC infrastructure
- Describe the AToM configuration procedure

https://t.me/learningnets

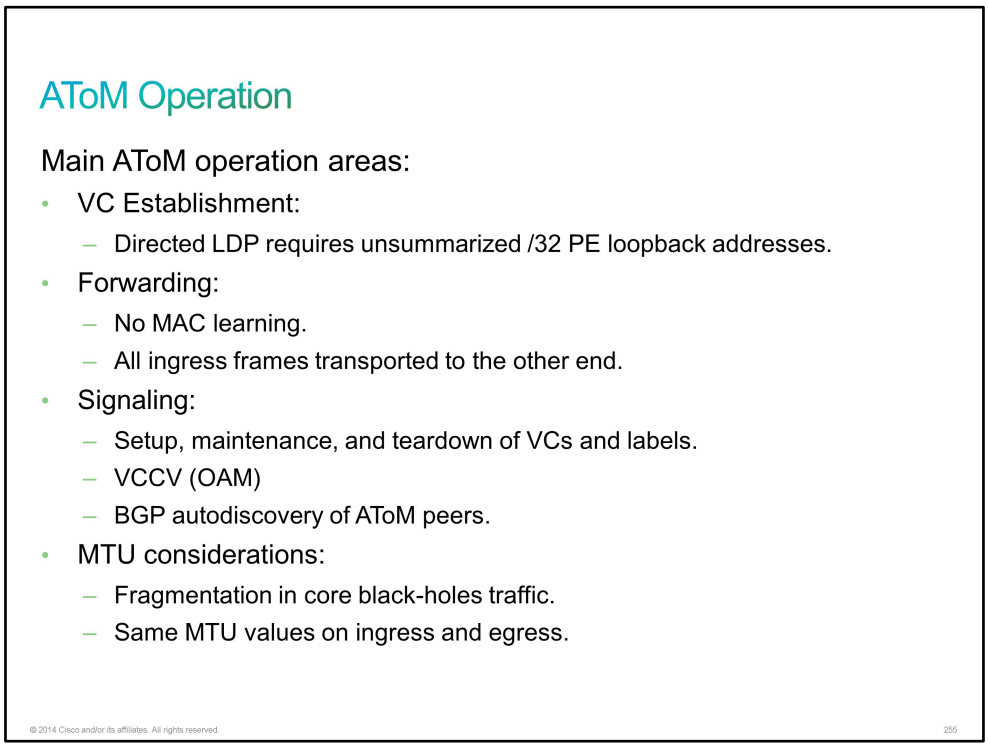
- Describe the EoMPLS configuration that interconnects a VLAN
- Describe the EoMPLS configuration that interconnects untagged frames
- Describe the EoMPLS configuration that interconnects a VLAN and manipulates 802.1Q tagging
- Explain a basic EoMPLS verification method
- Describe how to perform a detailed EoMPLS verification in Cisco IOS XR
- Describe how to perform a detailed EoMPLS verification in Cisco IOS XE
- Describe how to monitor attachment circuits in EoMPLS

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

AToM Operation

When implementing AToM, you need to consider a number of VC creation, forwarding, signaling, and MTU requirements.



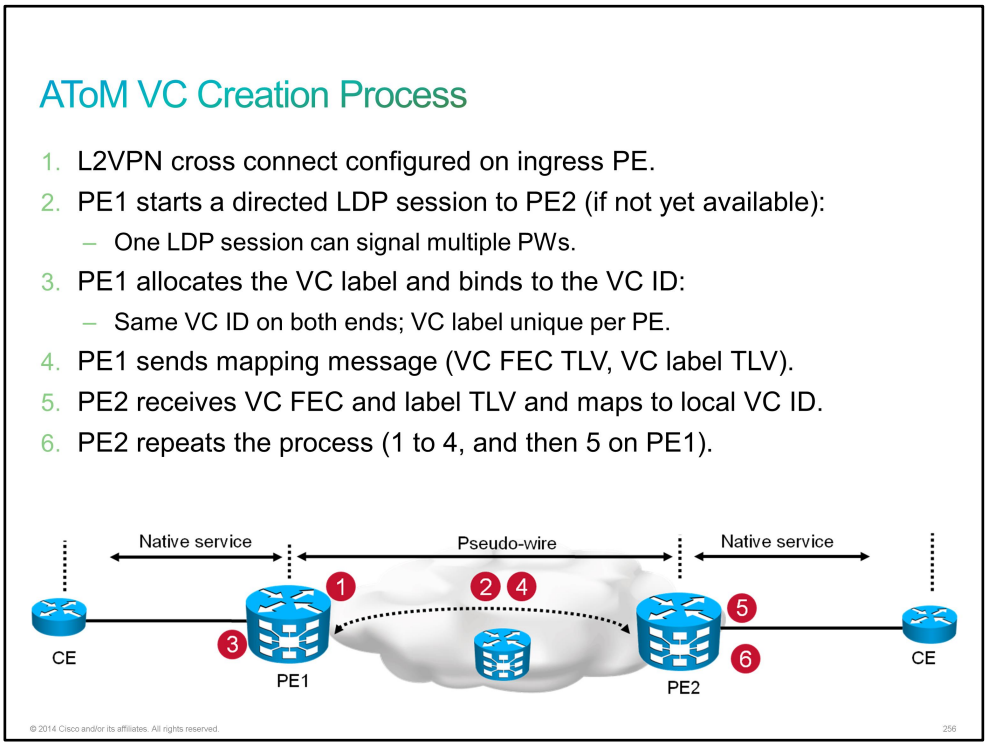
AToM is a subset of Virtual VPWS that enables point-to-point Layer 2 virtual connections over an MPLS infrastructure. Several aspects require special consideration:

- VC creation
- Forwarding
- Signaling
- MTU considerations

https://t.me/learningnets

AToM VC Creation Process

The AToM VC creation process is triggered by an administrator but occurs automatically through LDP signaling.



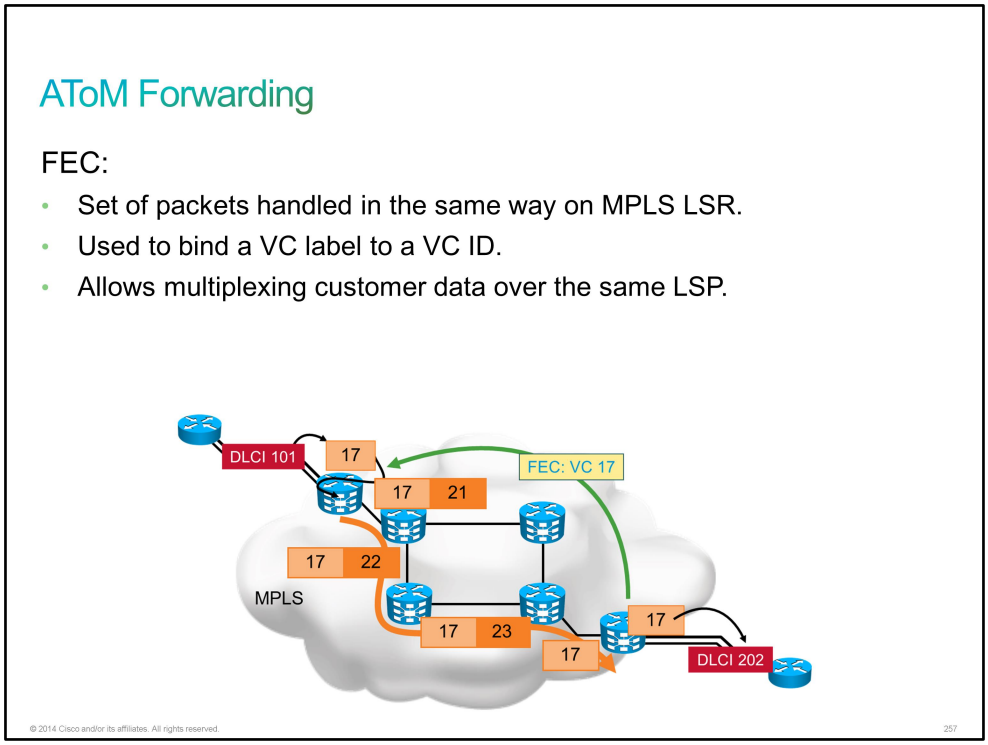
AToM VC creation process occurs in these steps:

1. The **xconnect** command is issued on the ingress PE (port, subinterface, and so on). Alternatively, the PW autodiscovery mechanism can be used to detect the neighbor.
2. PE1 starts a directed LDP session to PE2 if one is not yet available. If a directed session already exists to the destination PE, it is reused for another PW. One LDP session can signal multiple PWs.
3. PE1 allocates a VC label and binds it to the VC ID. The same VC ID value must be configured on both ends. The VC label is unique per PE.
4. PE1 sends a mapping message containing the VC FEC TLV and VC label TLV to the other end.
5. The other end (PE2) receives VC FEC and label TLV and maps it to the locally configured VC ID.
6. The other end (PE2) repeats the process (Steps 1-4), which is then finished on PE1 by receiving the VC FEC and label TLV and mapping it to the locally configured VC ID.

https://t.me/learningnets

AToM Forwarding

AToM packet forwarding follows an AToM LSP.



An IGPis required in the MPLS backbone. All routers in the backbone have routing information about how to send IP packets to each other. LDP is also used between directly connected neighbors. Local labels are assigned to each IGP-derived route. The label values are then propagated to the neighbor across the LDP session.

The IGP, together with the LDP sessions between directly connected neighbors, establishes LSPs from any router inside the backbone to any other router inside the backbone. In the figure, a unidirectional LSP from the upper right to the lower left is established between the ingress and egress PE routers. The tunnel label is used to propagate the packets along the LSP to the correct egress PE router.

The figure also shows the directed multihop LDP session that is used to exchange the VC label between the ingress and egress PE routers. Any ingress-egress PE router pair will need such an LDP session. In this example, the egress PE router allocates the label value 17. The VC label is advertised to the ingress PE router using the directed LDP session between them.

The ingress PE router now forms a label stack. The topmost label, the tunnel label, has the value 21 and is used to guide the packets to the egress PE router. The second label, the VC label, has the value 17 and is used by the egress PE router to propagate out the packets on the correct interface and attachment circuit.

The ingress PE router receives a Frame Relay frame on DLCI 101 on the incoming interface. The DLCI is mapped to the AToM tunnel across the backbone. The Frame Relay frame is therefore encapsulated into MPLS using the label stack with label 21 as the topmost label and label 17 as the second label.

https://t.me/learningnets

The packet is then forwarded along the LSP. The topmost label is used for label swapping in the next hop. The top label is changed to the value 22. In the next hop, label swapping results in label value 23 being the top label. In the router just before the egress router, the incoming label value 23 indicates pop. That label therefore performs PHP. The topmost label is removed, and the packet is propagated to the egress PE router with the label value 17, the VC label, which is now the only label left.

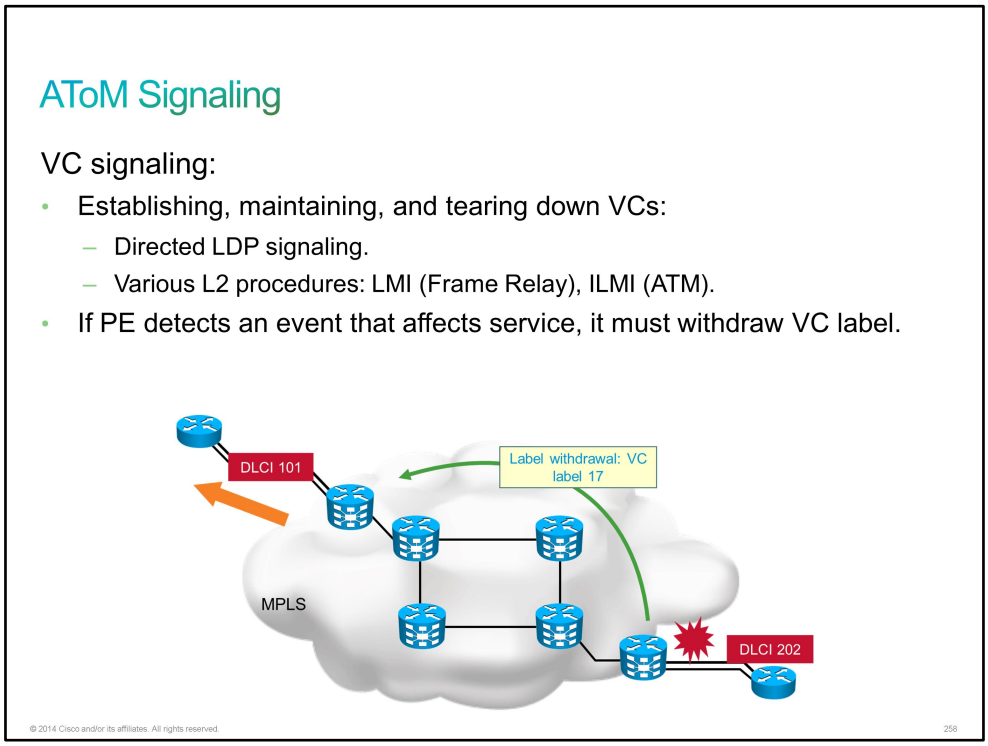
When the PE router receives the packet with label value 17, that label value instructs the PE router to decapsulate the packet and send it out on the associated Frame Relay DLCI. In this case, the DLCI value is 202. The Frame Relay frame is now reconstructed and transmitted.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

AToM Signaling

The signaling of a PW occurs over the targeted LDP sessions and includes the setup, maintenance, and teardown of the VC.



In addition, the PE routers translate the events from the local attachment circuit to the PW and vice versa. This translation depends on the encapsulation of the local attachment circuit.

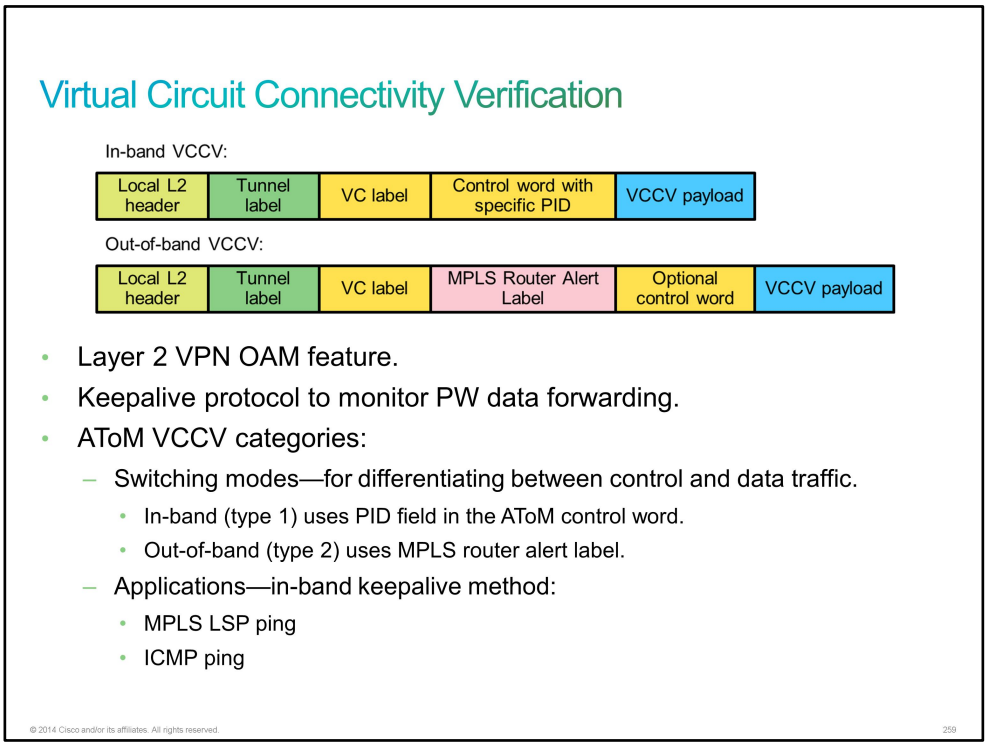
A PE router may provide circuit status signaling on the interface where the customer connects. A PE router that provides Frame Relay services must use LMI procedures with the customer equipment. A PE router that provides ATM services should use ILMI procedures.

An example of the signaling procedure, shown here, is the VC teardown. If a PE router detects a condition that affects normal service, it must withdraw the corresponding VC label.

https://t.me/learningnets

Virtual Circuit Connectivity Verification

The VCCV is a OAM feature for Layer 2 VPNs.



VCCV is Layer 2 VPN OAM feature that allows network operators to run a PE-to-PE keepalive protocol across a specified PW to ensure that the PW data path forwarding does not contain any faults.

When a PW is first signaled using LDP or L2TPv3, a message is sent from the initiating PE to the receiving PE. That message has been extended to include VCCV capability information that indicates to the receiving PE which combinations of control channel and connectivity verification types it is capable of receiving. If the receiving PE agrees to establish the PW, it will return its capabilities in the subsequent signaling message.

When MPLS is used to transport PW packets, VCCV packets are carried over the MPLS LSP. Packets are sent across this channel either as in-band traffic with the data of the PW, or out-of-band.

Two types of VCCV switching modes have been defined to distinguish VCCV packets from regular data traffic:

- **In-band VCCV (type 1):** Control channel type 1 uses a PID field in the AToM control word to identify an AToM VCCV packet. This type of VCCV is used for those PW types that employ the control word.
- **Out-of-band VCCV (type 2):** Control channel type 2 is also referred to as the MPLS router alert label. The MPLS router alert label is carried above the VC label to identify an AToM VCCV packet. Control channel type 2 can be used whether the PW is set up with a control word present or not.

Cisco routers use type 1 switching, if available, when they send MPLS LSP ping packets over an AToM VC control channel. Type 2 switching accommodates those VC types and implementations that do not support or interpret the AToM control word.

https://t.me/learningnets

An AToM VC advertises its AToM VCCV disposition capabilities in both directions—that is, from the originating router (PE1) to the destination router (PE2), and from PE2 to PE1. In some instances, AToM VCs might use different switching types if the two endpoints have different AToM VCCV capabilities. If PE1 supports type 1 and type 2 AToM VCCV switching and PE2 supports only type 2 AToM VCCV switching, there are two consequences:

- LSP ping packets sent from PE1 to PE2 are encapsulated with type 2 switching.
- LSP ping packets sent from PE2 to PE1 use type 1 switching.

You can determine the AToM VCCV capabilities advertised to and received from the peer by entering the **show mpls l2transport binding** command (Cisco IOS Software) or **show l2vpn xconnect all detail** command (Cisco IOS XR Software).

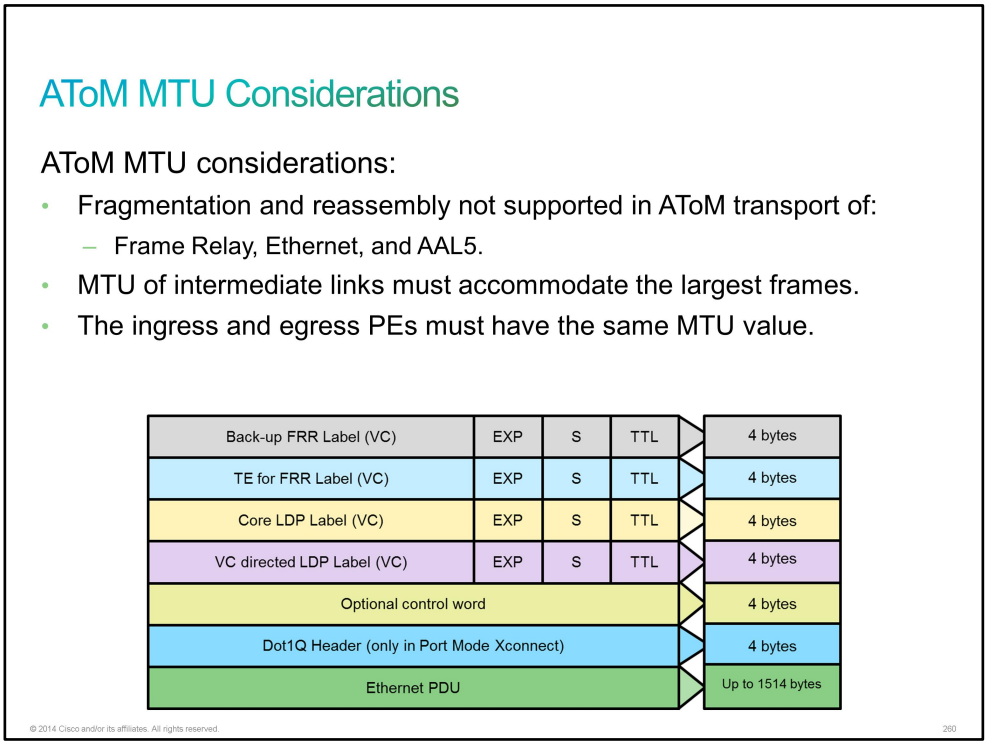
Two verification types can be negotiated by the in-band or out-of-band VCCV control channel in an MPLS environment:

- **ICMP ping:** When this optional connectivity verification mode is used, an ICMP echo packet (ICMPv4 or ICMPv6) achieves connectivity verification.
- **MPLS LSP ping:** This method helps monitor LSPs and quickly isolate MPLS forwarding problems.

https://t.me/learningnets

AToM MTU Considerations

There are several considerations regarding the MTU size in AToM environments.



Unlike IP, most Layer 2 protocols (for example, Frame Relay, Ethernet, and AAL5) do not allow fragmentation of frames. This fact has two implications:

- All intermediate links between the ingress PE router and the egress PE router must be able to carry the largest Layer 2 frame that has been received, including the imposed label stack and the 4-byte control word (if it is used).
- The ingress PE interface and the egress PE interface must have the same MTU value.

Failure to comply with the first rule means that the larger frames, where the label stack and the control word contribute to creating a larger size than can be carried, will be dropped by the backbone.

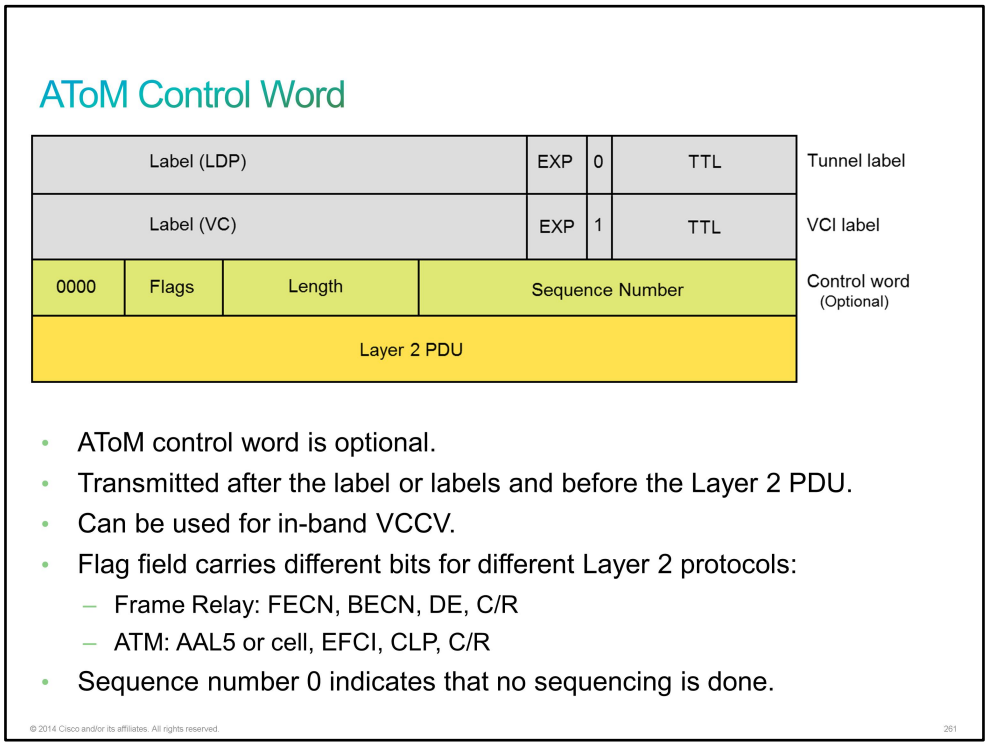
Failure to comply with the second rule means that frames that are forwarded along the LSP will be dropped by the egress PE if the frame size is too large for the egress PE interface.

The figure illustrates the potential overhead imposed by various labels and headers. The MTU size on the core must be able to accommodate the customer MTU size with the added overhead.

https://t.me/learningnets

AToM Control Word

The figure illustrates AToM encapsulation. The topmost label is the tunnel label. This label is followed by the VC label. Following the VC label is the optional control word. Next in the packet is the Layer 2 PDU.



The control word is optional. Some Layer 2 protocols make use of it, while others do not. Both endpoints (ingress PE and egress PE) must agree to either use or not use the control word. It is transmitted after the label stack but before the Layer 2 PDU. It can be used to carry important Layer 2 header information and to guarantee sequenced delivery, if required.

The control word is 32 bits long. It is divided into four fields. The first field is 4 reserved bits, which must always be set to zero. The next field is a 4-bit flag field. The flags have different uses depending on the Layer 2 protocol that is being forwarded. The third field is an 8-bit length field, which is used only if the Layer 2 PDU is shorter than the minimum MPLS packet and padding is required. If no padding is required, the length field is not used. The fourth field is a 16-bit sequence number. Sequence numbering is used only on Layer 2 protocols, and it guarantees ordered delivery. A special value of 0 in the sequence field indicates that there is no guaranteed sequenced delivery.

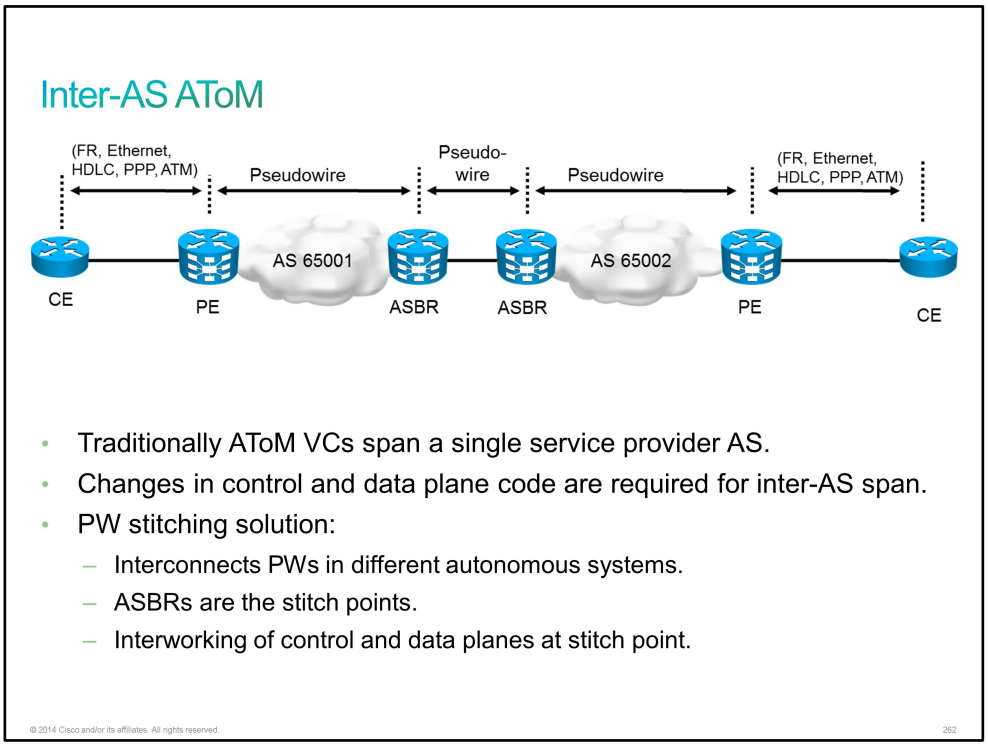
When AToM is used for FRoMPLS, the Frame Relay header is removed and the FECN, BECN, DE, and C/R bits are carried in the control word flag field.

When AToM is used for ATM over MPLS, the first flag in the control word flag field is used to indicate whether it is AAL5 frames or raw ATM cells that are being transported by AToM. The other three flags are used for EFCI, CLP, and C/R. If one of these flags is set in any of the ATM cells that are being transported in the MPLS packet, then the corresponding flag is set in the control word.

https://t.me/learningnets

Inter-AS AToM

Traditionally, AToM enables VC connections through a single AS.



To extend this function to inter-AS deployments, changes in the control and data would be required. PW stitching provides an immediate solution.

PW stitching enables the extension of Layer 2 VPN PWs across an inter-AS boundary or across two separate MPLS networks. Layer 2 VPN PW stitching connects two or more contiguous PW segments to form an end-to-end multihop PW. This end-to-end PW functions as a single point-to-point PW. Layer 2 VPN PW stitching enables the service provider to keep the IP addresses of the edge PE routers private across inter-AS boundaries. Using the IP addresses of the ASBRs, the ASBRs join the PWs of the two domains.

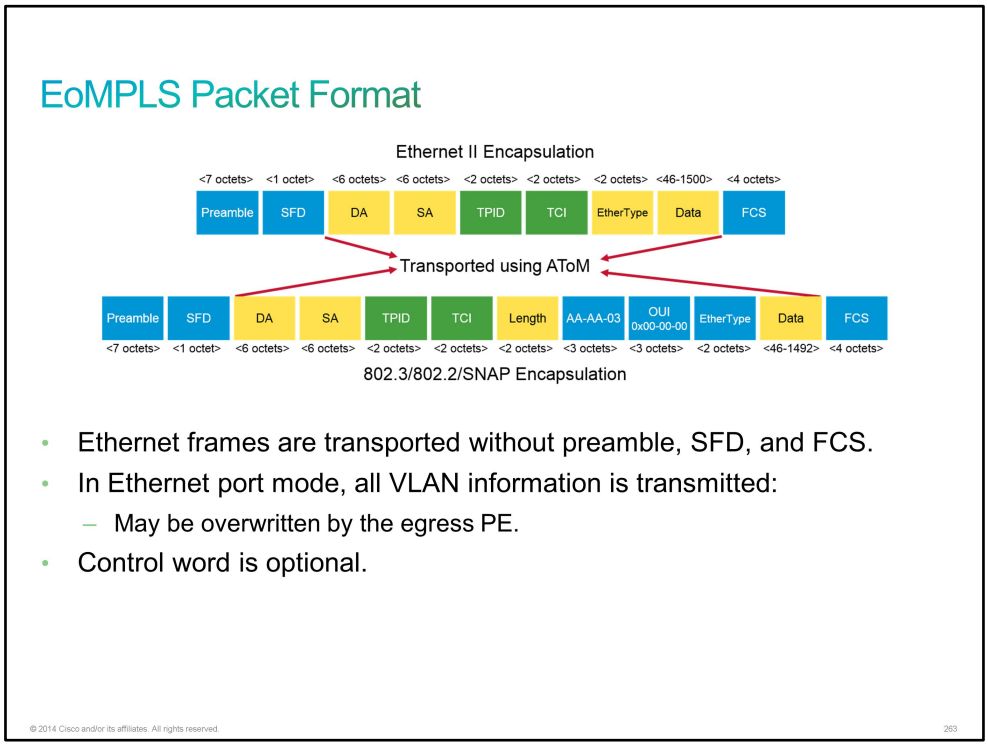
AToM packets forwarded between two PWs are treated the same as any other MPLS packet, with the following exceptions:

- The outgoing VC label replaces the incoming VC label in the packet. New IGP labels and Layer 2 encapsulation are added.
- The incoming VC label TTL field is decremented by one and copied to the outgoing VC label TTL field.
- The incoming VC label EXP value is copied to the outgoing VC label EXP field.
- The outgoing VC label “bottom-of-stack” S bit in the outgoing VC label is set to 1.
- AToM control word processing is not performed at the Layer 2 VPN PW switching aggregation point or ASBR. Sequence numbers are not validated.

https://t.me/learningnets

EoMPLS Packet Format

ATOM encapsulates entire frames but excludes fields whose transmission does not offer any benefits, such as frame synchronization data.



In EoMPLS, the preamble, the SFD and the FCS are excluded from encapsulation.

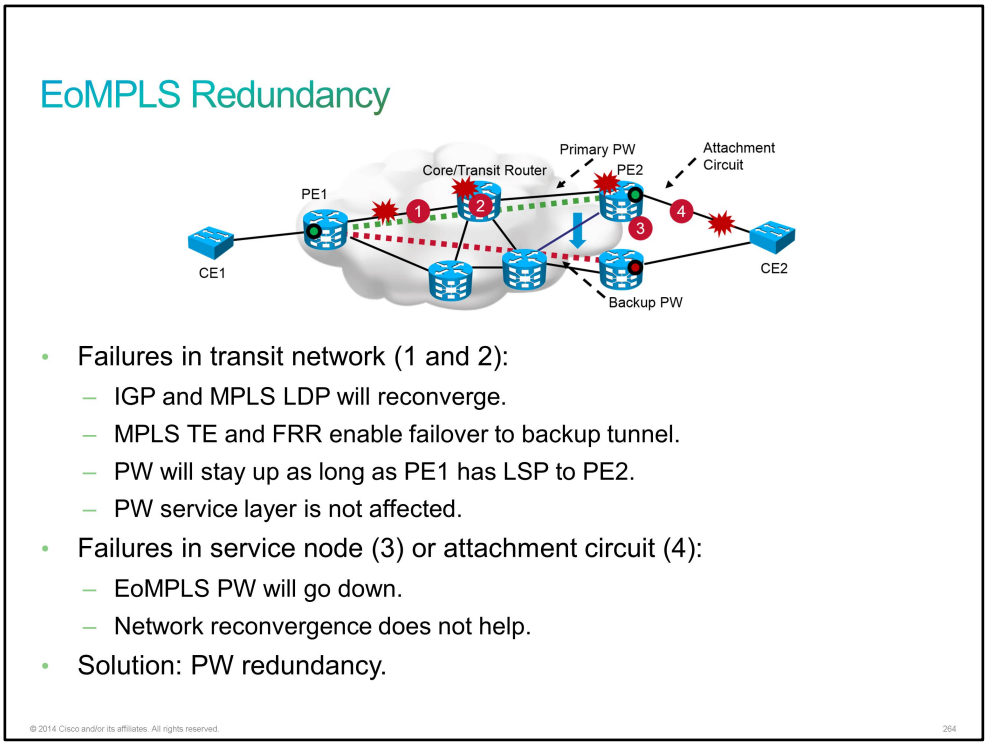
The preamble of an Ethernet frame consists of a 56-bit (7-byte) pattern of alternating 1 and 0 bits, which allows devices on the network to easily detect a new incoming frame. The SFD is designed to break this pattern, and signal the start of the actual frame. The SFD is the 8-bit (1-byte) value marking the end of the preamble of an Ethernet frame. The SFD is immediately followed by the destination MAC address. It has the value 10101011. The FCS provides Layer 2 checksum characters added to a frame for error detection and correction purposes.

In Ethernet port mode, all VLAN information is transmitted. The VLAN tag or tag stack may be overwritten by the egress PE. The egress PE can also manipulate the VLAN tag received over the VC. The control word is optional in EoMPLS.

https://t.me/learningnets

EoMPLS Redundancy

It is common practice to add redundancy to the interconnect between two customer sites to avoid split-subnet scenarios and service interruption. The redundancy solution comprises several building blocks. One block refers to the redundant LAN access and considers issues related to the STP.



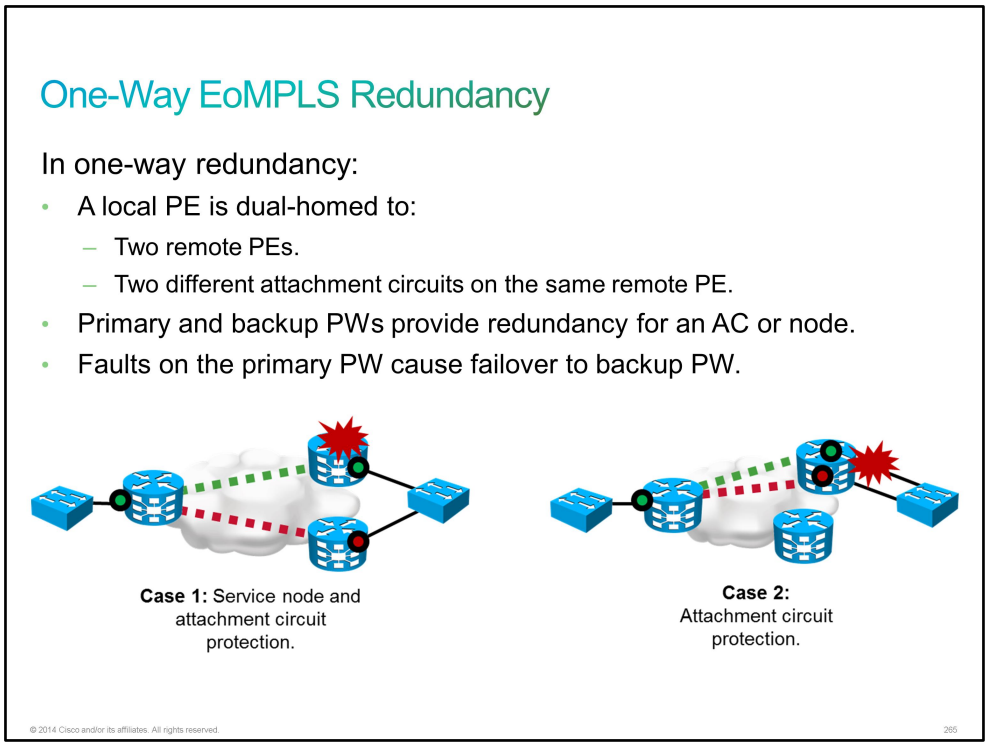
This figure presents a redundant transit network and redundant attachment circuits. In case of failure 1 or 2, as shown above, IGP and MPLS LDP will reconverge. MPLS traffic engineering and FRR will trigger an automatic failover to the backup tunnel. The PW will stay up as long as PE1 has an available LSP path to PE2. This failure scenario does not affect the PW service layer.

In case of failure 3 or 4, as shown above, the PW will go down. Network layer reconvergence will not help re-establish the service. A solution has been developed to address this problem. It is referred to as PW redundancy.

https://t.me/learningnets

One-Way EoMPLS Redundancy

PW redundancy can be implemented using a one-way or two-way method.



In one-way redundancy, the local PE has two PWs to two remote PEs serving the same destination site or to the same egress PE with redundant attachment circuits. One PW is declared as primary, the other as backup. A fault of the primary PW triggers a failover to the backup PW.

https://t.me/learningnets

Two-Way EoMPLS Redundancy

In two-way redundancy, four PWs are used to provide high availability service.

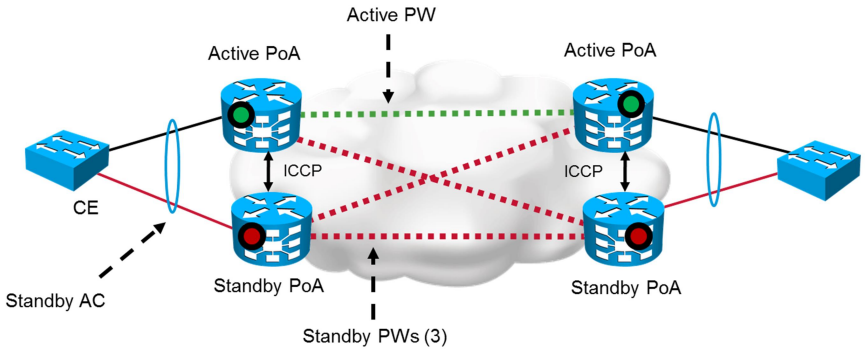
Two-Way EoMPLS Redundancy

Two-way redundancy requires:

- Dual-homing of two local PEs to two remote PEs.
- Four PWs:
  - One primary PW.
  - Three backup PWs.
- MC-LAG:
  - PEs run ICCP.
  - ICCP synchronizes state and forms a redundancy group.

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

Two-Way EoMPLS Redundancy (Cont.)



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

https://t.me/learningnets

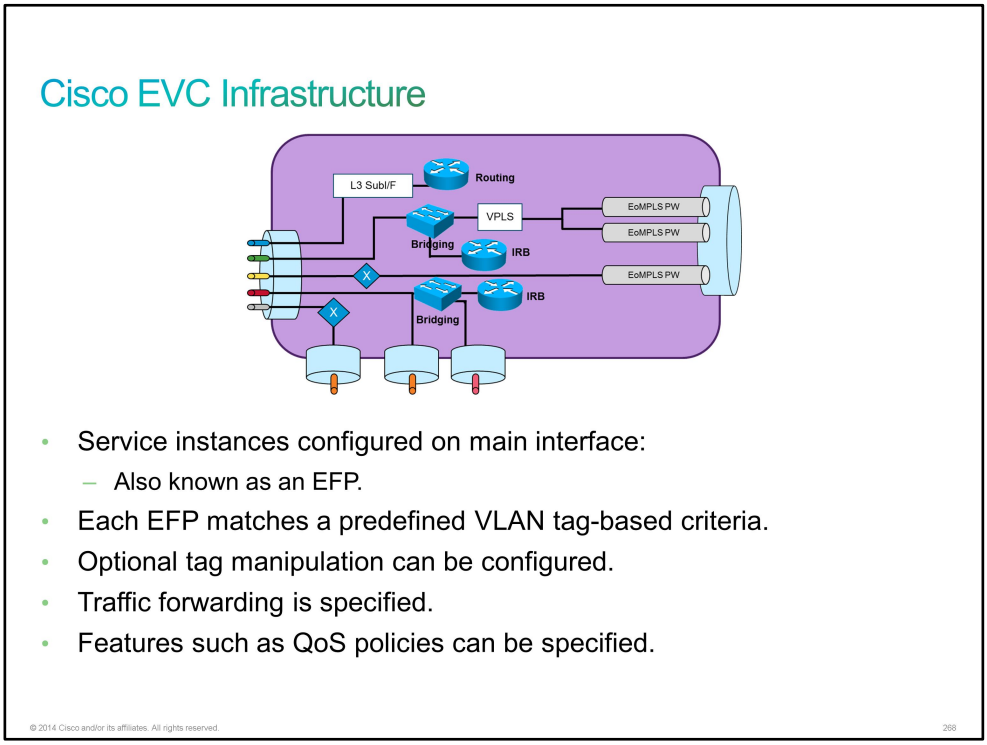
Only one PW is declared as primary. The three remaining PWs are intended for backup. In each site, one attachment circuit is primary, the other backup. To synchronize the redundancy information between the LAN and the PE devices, MC-LAG must be enabled in the network. MC-LAG refers to the PE devices as points of attachment nodes. The points of attachment run ICCP to synchronize state and form a redundancy group.

https://t.me/learningnets

Cisco EVC Infrastructure Overview

Cisco Ethernet Virtual Circuit infrastructure is used to represent Cisco software architecture to address Carrier Ethernet services.

Cisco EVC Infrastructure



Cisco implements EVC using EFP. An EFP is a substream partition of a main interface. On Cisco routers, the EFP is implemented as a Layer 2 subinterface with an encapsulation statement.

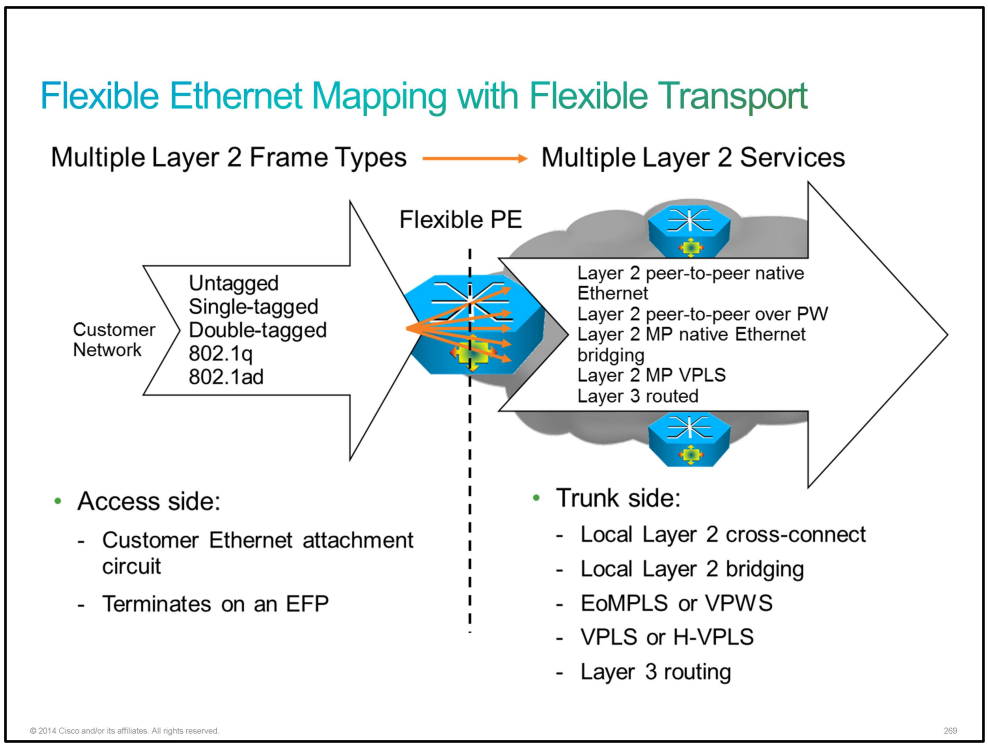
The EVC solution defines these aspects of Ethernet-based attachment circuits and virtual circuits:

- Frame matching based on one or more VLAN tags.
- Optional VLAN tag manipulation.
- Traffic forwarding.
- Additional services, such as QoS policies.

EVC is supported on Cisco IOS and IOS XR Software.

https://t.me/learningnets

Flexible Ethernet Mapping with Flexible Transport



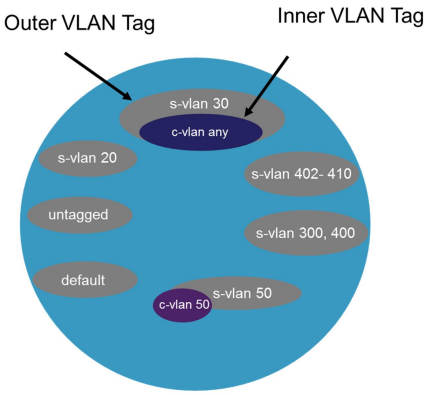
Customer-to-service mapping is illustrated here as a connection between an Ethernet-based data flow on the customer side and a service on the trunk side.

Flexible Ethernet mapping is the ability to process and classify different Ethernet frame types, each with different attributes (EtherTypes, VLAN tags, CoS bits, and so on). Cisco IOS XR Software uses the EFP concept to provide flexible Ethernet mapping. Flexible transport is found on the trunk side. Each Ethernet flow from the customer side is mapped or connected to a service on the trunk side. These service types can be native Ethernet, IP, or IP and MPLS-based, and they form the basis for Layer 2 VPN.

https://t.me/learningnets

Flexible Ethernet Mapping with Flexible Transport (Cont.)

- EFPs enable flexible mapping of frames into Layer 2 services.
- Mapping is based on VLAN tagging:
  - 802.1Q, 802.1ad
  - Single-tag or double-tag
  - Unique or multiple values (ranges or lists)
  - Untagged traffic
  - Unclassified traffic (default)



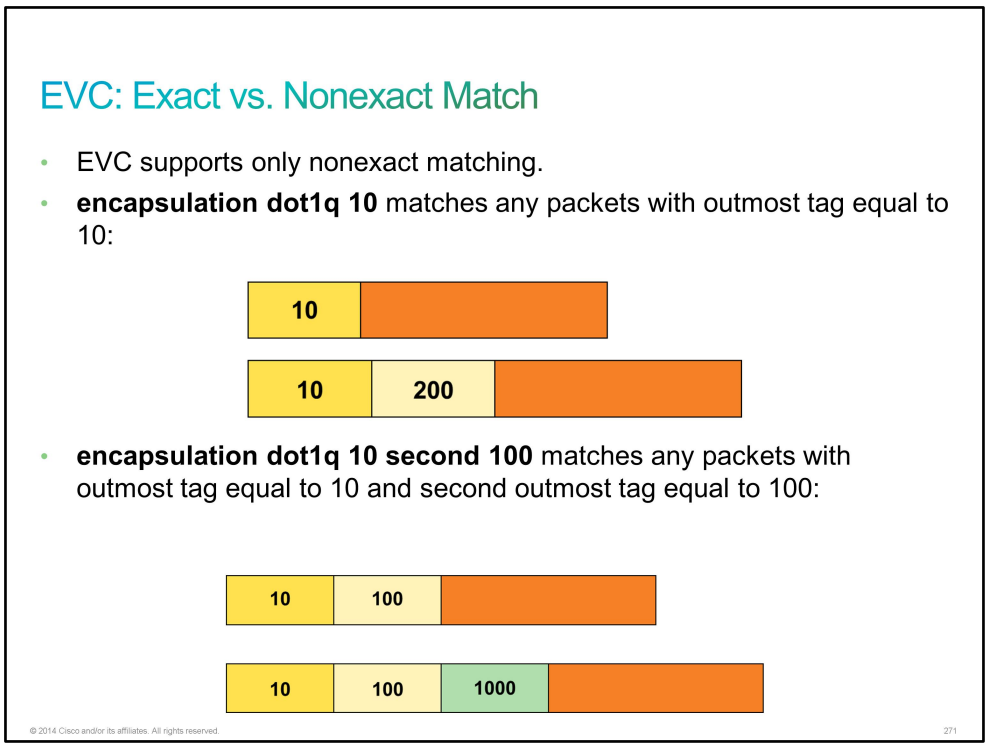
© 2014 Cisco and/or its affiliates. All rights reserved.

One of the EVC advantages is flexible frame matching. Flexible frame matching is a functionality that allows each service instance to match frames with either a unique single VLAN, or a list or range of VLANs. It can also match single- or double-tagged frames, untagged frames, or any frames that are not matched by the specific statement.

Flexible frame matching is the first step when configuring a service instance. Subsequent steps, encapsulation rewrite, and forwarding definition are described in the following pages.

https://t.me/learningnets

EVC: Exact vs. Nonexact Match



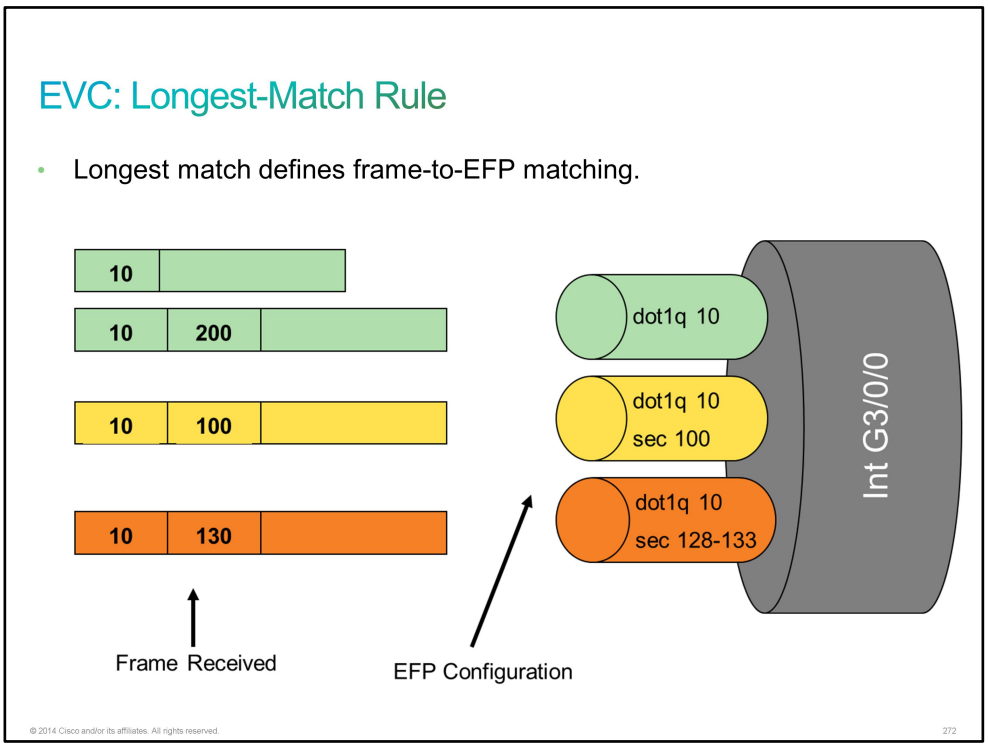
Several subinterfaces with various matching statements can exist on a main interface. The matching logic affects the assignment of a frame to a specific subinterface. The EVC supports only nonexact matching, in which additional inner-VLAN tags are not taken into consideration.

The command **encapsulation dot1q 10** matches any frames with the outmost VLAN tag equal to 10. This includes both 801.1q frames, QinQ frames with any second VLAN tag, as well as frames with more than two VLAN tags, if the outmost tag is 10.

The command **encapsulation dot1q 10 second 100** matches any frames with the outmost VLAN tag equal to 10 and second outmost tag equal to 100. This includes frames that contain more than two VLAN tags, if the first two tags fulfill the defined criteria.

https://t.me/learningnets

EVC: Longest-Match Rule



If several potential matches exist on the same main interfaces, the longest-match rule defines the most specific hit.

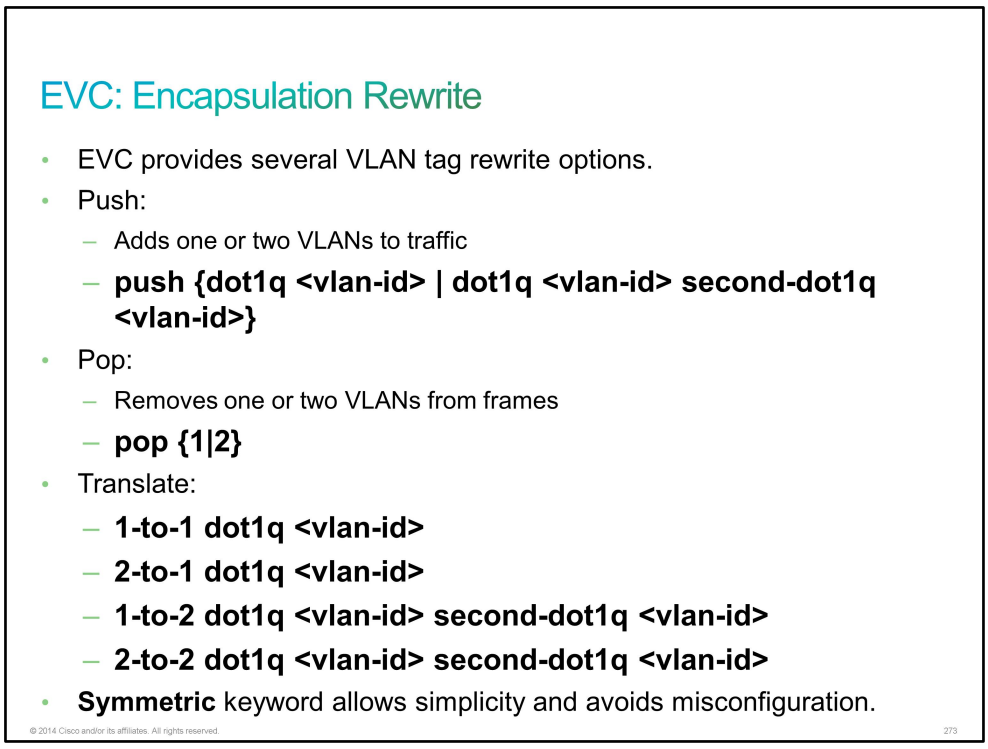
With the three subinterfaces shown above, the first one is the only subinterface that matches the first two frames (single tag 10, and double tag 10-200).

The second frame, with double tag (10-100) is matched by the first and second subinterface. Because the second subinterface defines a more exact match, it will process the second frame.

The third frame, with double tag (10-130) is matched by the first and third subinterface. Because the third subinterface defines a longer match, it will process the third frame.

https://t.me/learningnets

EVC: Encapsulation Rewrite



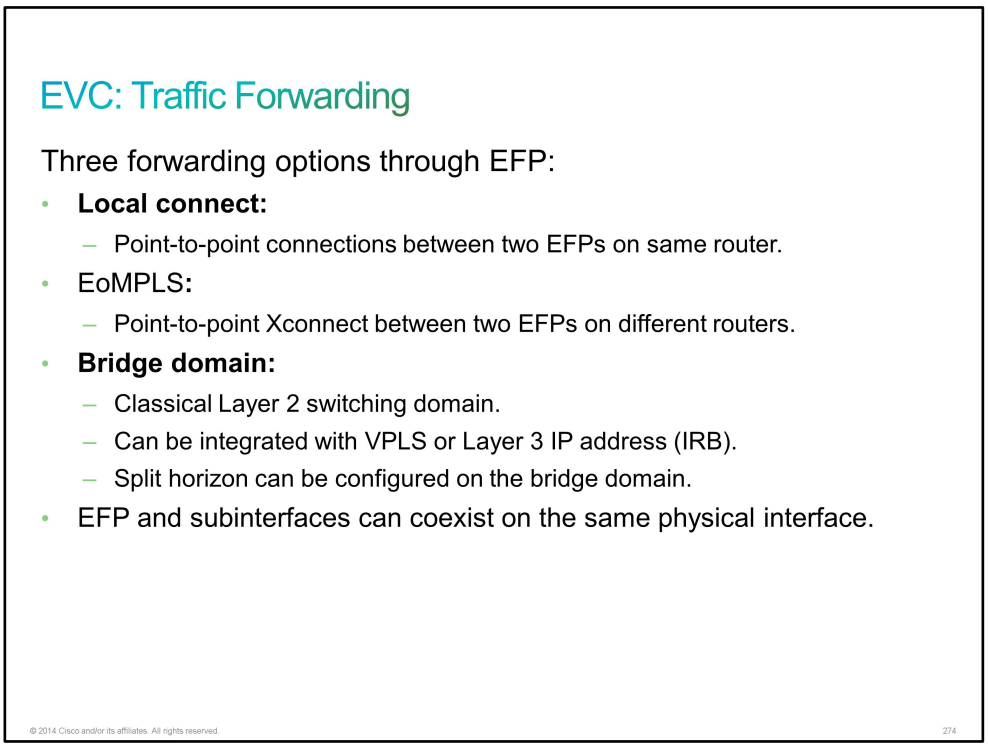
Once a packet is matched and assigned to a given subinterface, its VLAN tags can be modified. You can configure several of these three operations:

- **Push:** This action adds one or two VLANs to traffic using the **push {dot1q <vlan-id> | dot1q <vlan-id> second-dot1q <vlan-id>}** command.
- **Pop:** This action removes one or two VLANs from frames.
- **Translate:** This action allows the replacement of one or two tags by another one or two tags, using one of these commands:
  - **1-to-1 dot1q <vlan-id>**
  - **2-to-1 dot1q <vlan-id>**
  - **1-to-2 dot1q <vlan-id> second-dot1q <vlan-id>**
  - **2-to-2 dot1q <vlan-id> second-dot1q <vlan-id>**

The **symmetric** keyword defines that the pop, push, or translate operation is reversed for egress frames.

https://t.me/learningnets

EVC: Traffic Forwarding



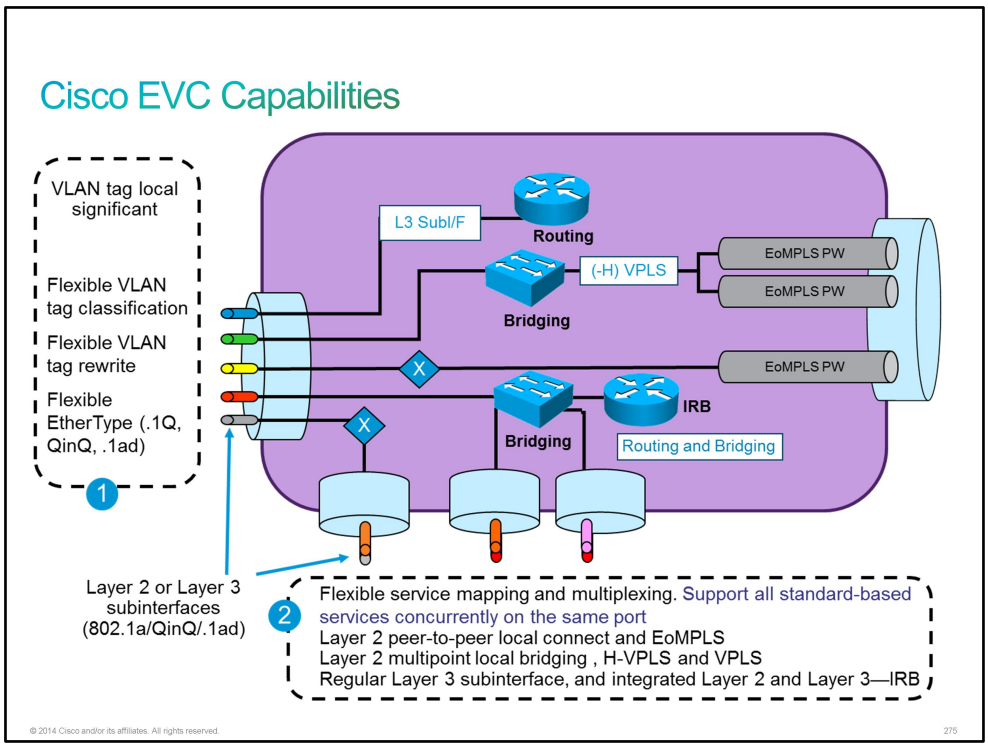
Lastly, the EVC framework is used to define traffic forwarding. Incoming frames can be transmitted to these destination types:

- **Local connect:** This destination represents a point-to-point connection between two EFPs on the same router. No VLANs are required to be defined on the PE device.
- **EoMPLS:** This destination refers to a point-to-point Xconnect between two EFPs on different PE routers. No VLANs are necessary on the PE device.
- **Bridge domain:** This destination describes a VPLS that uses a classical Layer 2 switching domain. The associated SVI can be enabled for a VPLS or Layer 3 IP address (IRB). You can enable or disable the split-horizon rule within the bridge domain.

EFP and subinterfaces can coexist on the same physical interface.

https://t.me/learningnets

Cisco EVC Capabilities

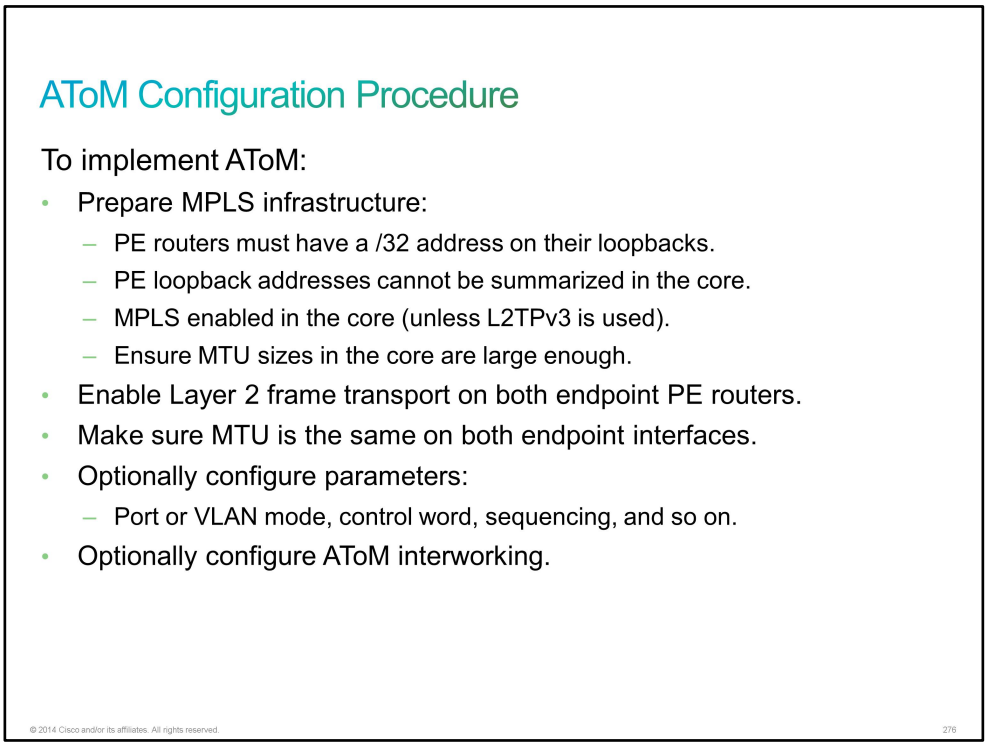


This diagram summarizes the EVC functionality by showing a physical Gigabit Ethernet interface with multiple EFP-based subinterfaces. The subinterfaces can be independently configured with flexible VLAN matching and VLAN tag rewriting. The incoming frames are then forwarded according to the defined parameters: linked to local attachment circuits, bridged or switched over a PW, or routed.

https://t.me/learningnets

AToM Configuration Procedure

AToM configuration builds upon a operational MPLS environment.



To implement AToM, you first need to prepare the MPLS infrastructure. This includes these steps:

- PE routers must have /32 addresses on their loopbacks. This is required for binding the VC label and the transport LDP into a label stack.
- The PE loopback addresses cannot be summarized in the core. Aggregation in the network would break the LSP path.
- MPLS must be enabled in the core unless L2TPv3 is used to provide IP-based transport.
- MTU sizes on the core links must be able to accommodate the customer MTU with the added label overhead.

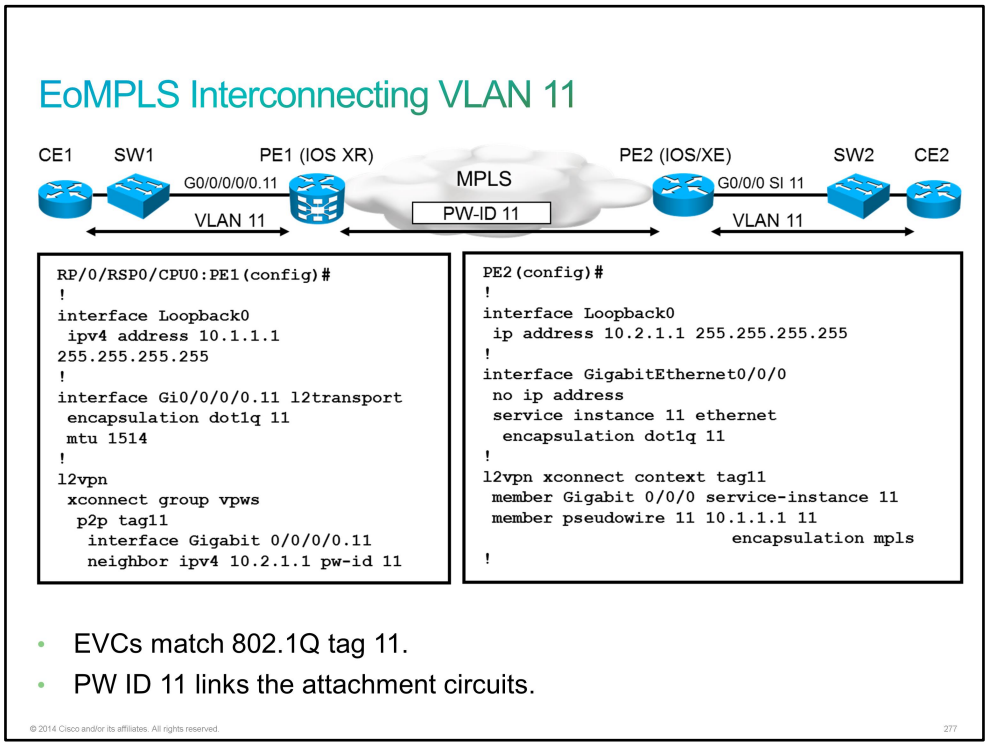
Once the MPLS infrastructure is ready, you will do the following:

- Enable Layer 2 frame transport on both endpoint PE routers.
- Make sure MTU is the same on both endpoint interfaces.
- Optionally configure parameters, such as port or VLAN mode, the control word, sequencing, and so on.
- Optionally configure AToM interworking.

https://t.me/learningnets

EoMPLS Interconnecting VLAN 11

EoMPLS is the most common AToM type. The configuration on Cisco IOS XR platforms differs from the one on Cisco IOS or IOS XE platforms.



On both systems, the PW class can used as a container for optional parameters, such as encapsulation type, the use of control word, transport mode (port or VLAN), sequencing, and others. If you do not configure a PW class, as above, the default encapsulation of MPLS is assumed.

In Cisco IOS XR Software, the attachment circuit must be enabled for Layer 2 transport using the **l2transport** keyword.

On Cisco IOS and IOS XE platforms, the attachment circuit is configured with the **service instance *number* ethernet [name]** interface configuration command. This commands creates an EFP on a Layer 2 interface and enters the service instance configuration mode. You use service instance configuration mode to configure all management and control date plane attributes and parameters that apply to the service instance on a per-interface basis. In this example the service instance matches all frames with 8021.Q tag 11.

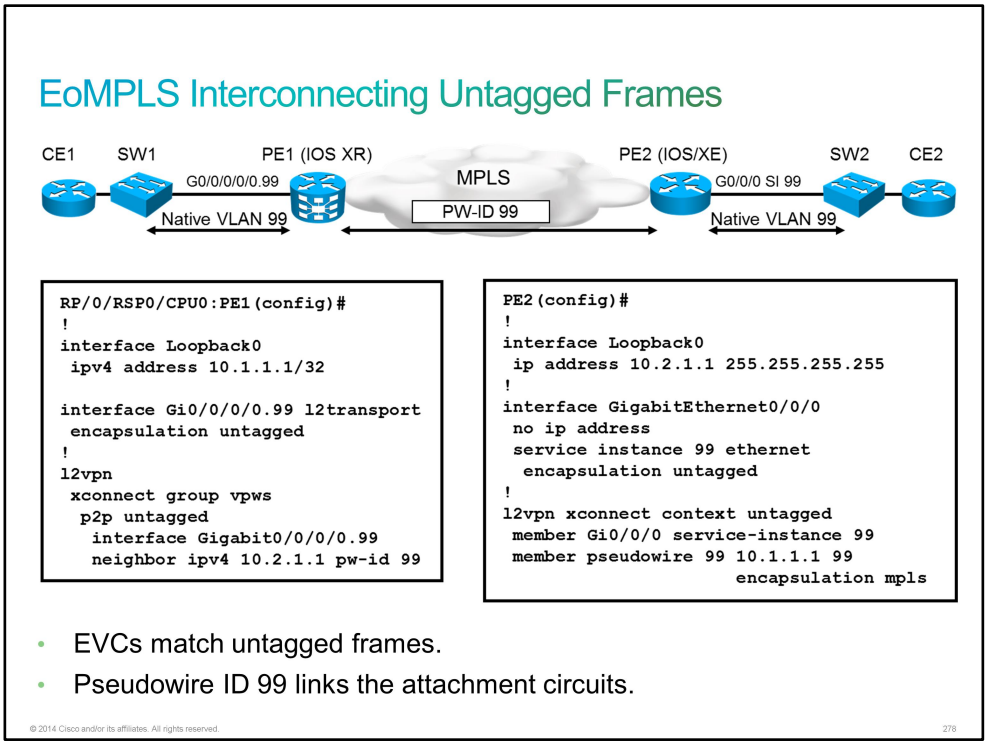
In Cisco IOS XR the EoMPLS service is defined in the l2vpn xconnect group configuration mode. You define a point-to-point service using the **p2p** command. Its elements can include an attachment circuit and a PW. The neighbor command specifies the remote PE address and the VC ID.

In Cisco IOS/XE the EoMPLS service is defined as a l2vpn xconnect context. Its two members are the attachment circuit and the PW. The attachment circuit is the Ethernet service instance 11 on the GigabitEthernet0/0/0 interface. The PW ID 11 connects to the peer (10.1.1.1), uses VC number 11 and encapsulation MPLS.

https://t.me/learningnets

EoMPLS Interconnecting Untagged Frames

Cisco IOS XR offers two ways to configure EoMPLS for untagged frames. An attachment circuit can be a main interface, where the l2transport command is configured under the interface configuration mode, or a subinterface, where the **l2transport** keyword is configured after the subinterface number. In the latter case, depicted in this scenario, you can match the frames using the **encapsulation untagged** command.



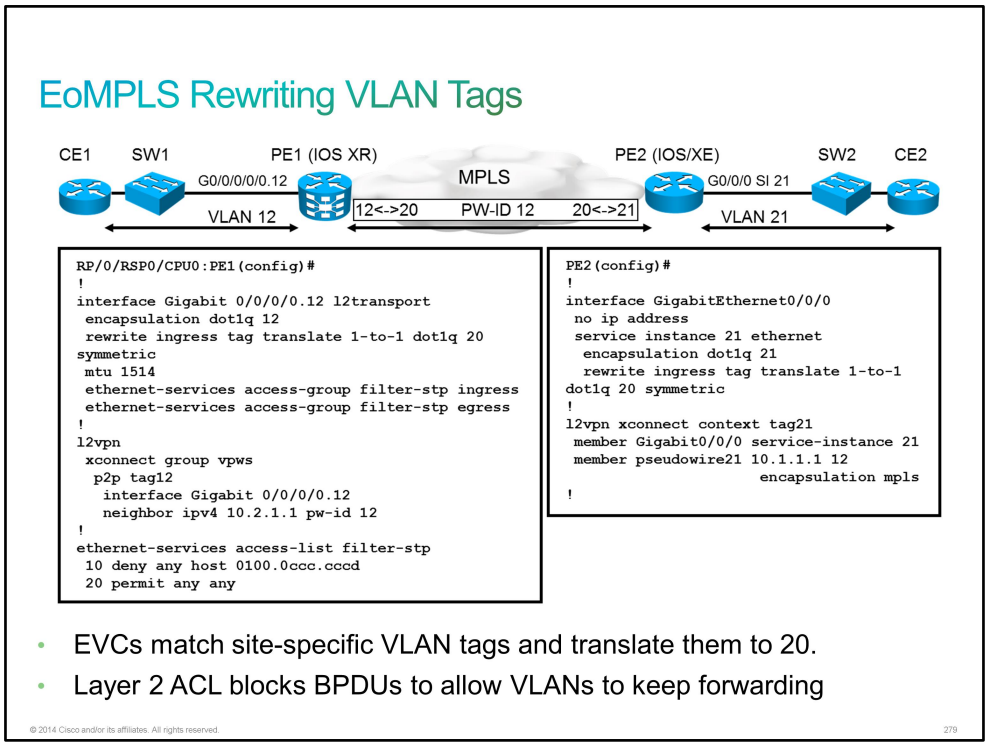
On the Cisco IOS/XE platforms you can use the Ethernet service instances to match specific frames received over an interface. In this example service instance 99 has been configured to match untagged frames.

The EoMPLS service contains a local attachment circuit and a PW to the peer PE. The VC ID must match up on both sides. In this example the VC ID is 99.

https://t.me/learningnets

EoMPLS Rewriting VLAN Tags

This scenario illustrates how to use EoMPLS to interconnect VLAN 12 in one customer site with VLAN 21 in another site. This is achieved with the rewrite feature available in the attachment circuit configuration. The 802.1Q tag 12 is replaced by 802.1Q tag 20 when frames are received by the Cisco IOS XR router and forwarded along the PW. The translation occurs symmetrically for ingress and egress traffic. The Cisco IOS XE router translates 802.1Q tag 21 to 802.1Q tag 20 in symmetric fashion.



The configuration of the Cisco IOS XR router shows a layer 2 ACL that filters STP packets. Customer switches that run PVST+will block switch ports when they receive a BPDU with an insconsistent PVID. This problem occurs when PE devices translate the 802.1Q tags on the PVST+ BPDUs.

When the VLAN tags are rewritten in an EoMPLS service, these syslog messages will be reported by the customer switches:

%SPANTREE-2-RECV_PVID_ERR: Received BPDU with inconsistent peer vlan id 12 on GigabitEthernet0/0 VLAN21.

%SPANTREE-2-BLOCK_PVID_LOCAL: Blocking GigabitEthernet0/0 on VLAN21. Inconsistent local vlan.

This issue is due to the PVID tag that is included with the PVST+ BPDUs. This tag is designed in order to detect misconfigurations and avoid accidental loops. But, in this scenario, it causes each end to be blocked and to not allow traffic to pass.

In order to prevent this problem, you can block the PVST+ BPDUs. This action disables Spanning Tree, and can result in loops if redundant connections are available between the switches. You can block the BPDU in two ways:

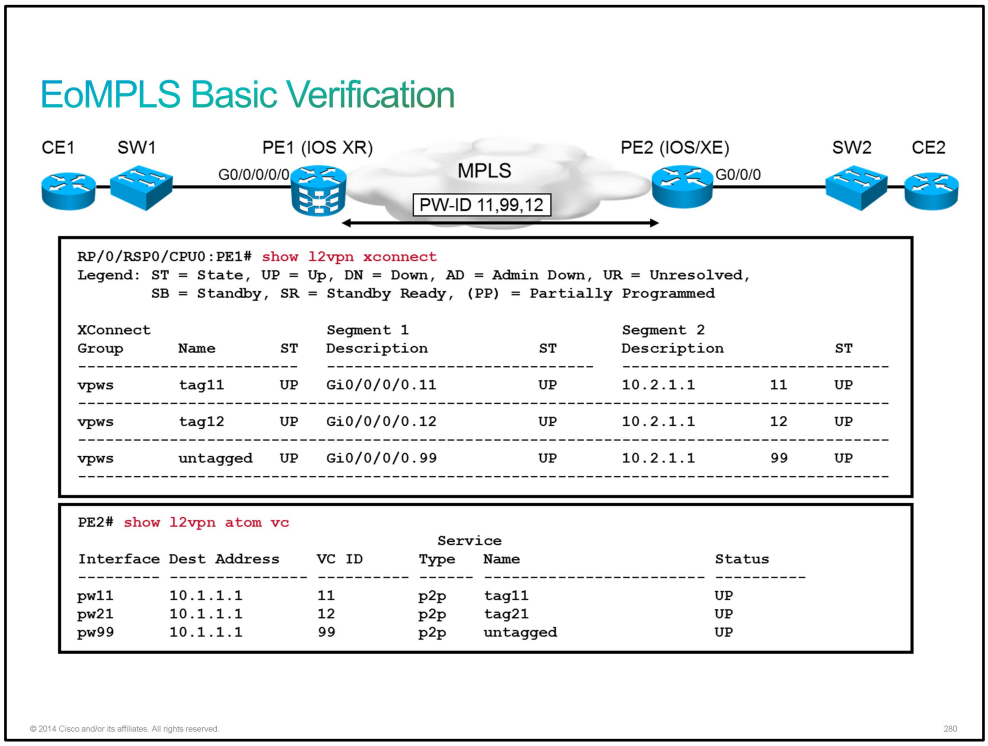
https://t.me/learningnets

- **BPDU filter on switches.** The BPDUs are blocked with the BPDU filter feature on the switches. The BPDU filter blocks BPDUs in both directions, which effectively disables Spanning Tree on the port. The BPDU filter prevents inbound and outbound BPDU. If you enable BPDU filtering on an interface, it is the same as if you disable Spanning Tree on it, which can result in Spanning Tree loops. You can enable the BPDU filters with the **spanning-tree bpdufilter enable imterface** command.
- **Block PVST+ BPDUs on the PE router.** You can filter the PVST+ BPDUs using a layer 2 ethernet-services access-list to deny packets destined to the PVST+ BPDU MAC address. PVST+ BPDU for the non-VLAN 1 (non-native) VLAN are sent to the PVST+ MAC address (also called the SSTP MAC address, 0100.0ccc.cccd), and tagged with a corresponding IEEE 802.1Q VLAN tag. The configuration on the Cisco IOS XR platform illustrates such an ACL.

https://t.me/learningnets

EoMPLS Basic Verification

You can monitor AToM service from a high-level perspective using the **show l2vpn xconnect** and **show l2vpn atom vc** commands in Cisco IOS XR and Cisco IOS/XE, respectively.

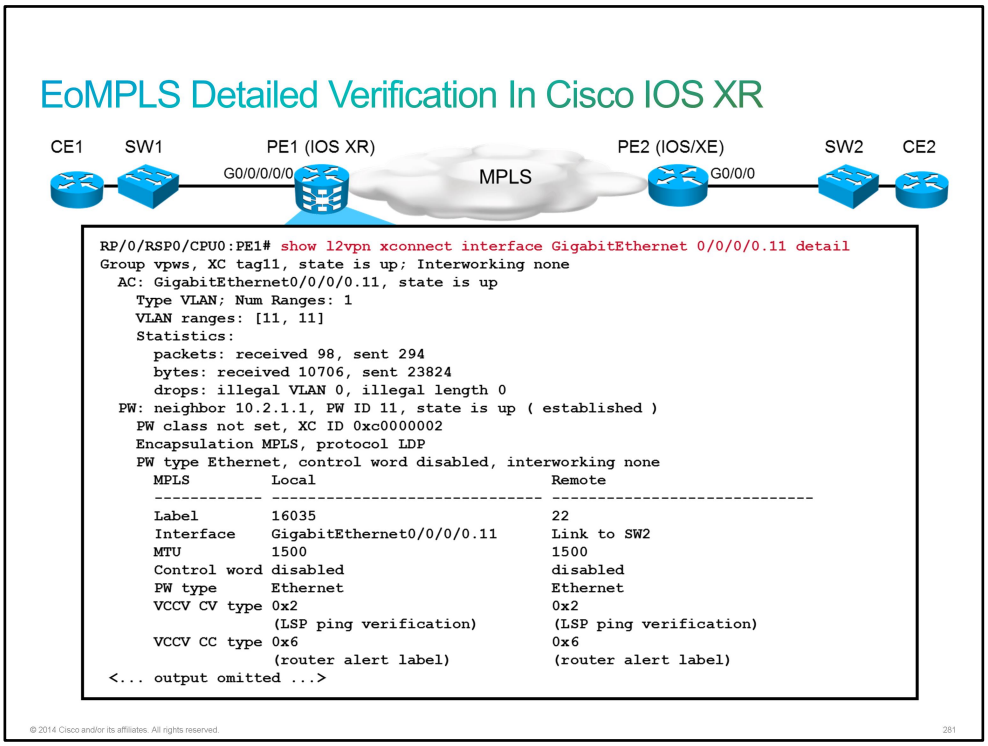


The output provides information about the overall service state, as well as the local attachment point and the PW that constitute the AToM service.

https://t.me/learningnets

EoMPLS Detailed Verification In Cisco IOS XR

You can examine the AToM information in a detailed fashion.

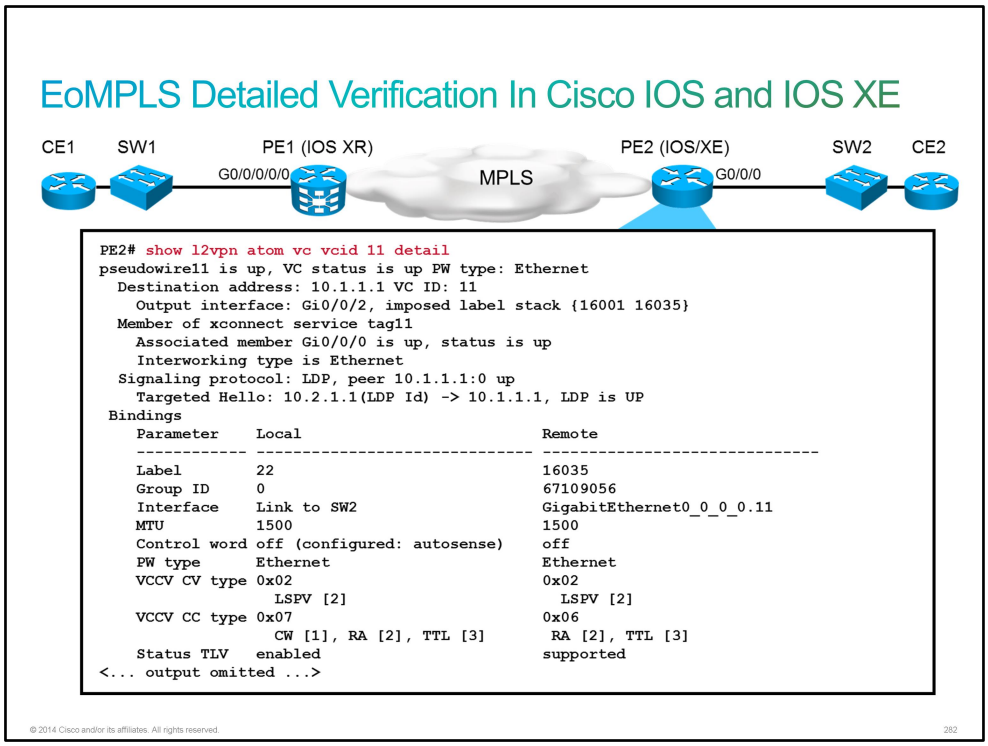


The Cisco IOS XR **show l2vpn xconnect** command with the **detail** keyword provides information about the AToM type, interworking method, processed VLAN, traffic statistics, PW class (which is not set in this example), encapsulation (MPLS in this example), signaling protocol (LDP in this example), and the parameters used locally and remotely. Most of these parameters need to match up for the pseudowire segment to become operational. You can also see the VC label advertised by each end.

https://t.me/learningnets

EoMPLS Detailed Verification In Cisco IOS and IOS XE

Cisco IOS/XE platforms allow you to display detailed AToM information using the **show l2vpn atom** command with the detail keyword.

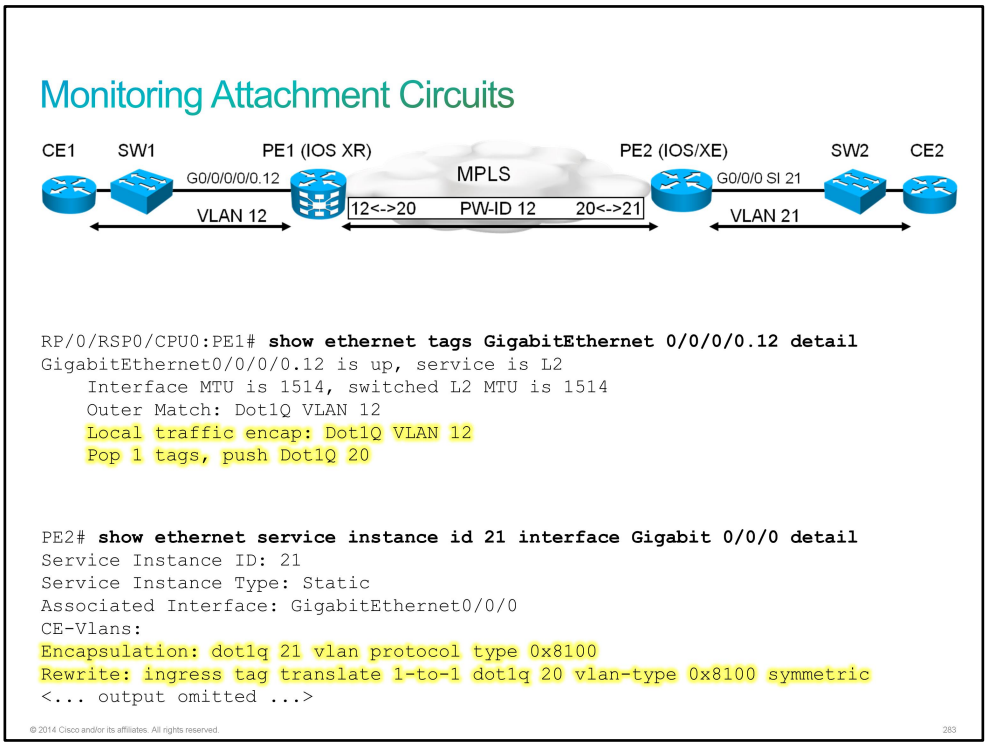


Similarly to the details provided in Cisco IOS XR, you can examine the AToM type, interworking method, traffic statistics, encapsulation (MPLS in this example), signaling protocol (LDP in this example), and the parameters used locally and remotely. You can also see the VC labels advertised by each end of the PW.

https://t.me/learningnets

Monitoring Attachment Circuits

You can monitor the tag matching and manipulation at the attachment circuits.



In Cisco IOS XR you can view the tag matching information when displaying the interface information (not shown above). Further you can use the **show ethernet tags** (Cisco IOS XR) and **show ethernet service instance** (Cisco IOS XE) commands for this purpose. The detailed output displays the outer and inner tag matching, the encapsulation protocol, and the tag manipulation.

In this example, the frames arriving with 802.1Q tag 12 on PE1 are matched. The VLAN tag is replaced by tag 20 in a symmetric fashion. PE2 performs analogous operation for VLAN tag 21 used on the PE-CE link.

https://t.me/learningnets

Summary

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

Summary

- AToM is a subset of Virtual VPWS that enables point-to-point Layer 2 virtual connections over an MPLS infrastructure.
- AToM provisioning occurs in six steps.
- An IGP is required in the MPLS backbone for AToM forwarding.
- The signaling of the VC occurs over the targeted LDP sessions and includes the setup, maintenance, and teardown of the PW.
- VCCV is a Layer 2 VPN OAM feature.
- Unlike IP, most Layer 2 protocols do not allow fragmentation of frames.
- Both endpoints (ingress PE and egress PE) must agree to either use or not use the control word.
- PW stitching enables the extension of Layer 2 VPN PWs across an inter-AS boundary or across two separate MPLS networks.

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

Summary (Cont.)

- AToM is a subset of Virtual VPWS that enables point-to-point Layer 2 virtual connections over an MPLS infrastructure.
- It is common practice to add redundancy to the interconnect between two customer sites to avoid split-subnet scenarios and service interruption.
- Cisco EVC is used to represent Cisco software architecture to address Carrier Ethernet services.
- To implement AToM, you first need to prepare the MPLS infrastructure.

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

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Implementing VPLS

Overview

VPLSenables enterprises to link together their Ethernet-based LANs from multiple sites via the infrastructure provided by their service provider. From the enterprise perspective, the service provider's public network looks like one giant Ethernet LAN. For the service provider, VPLS provides an opportunity to deploy another revenue-generating service on top of their existing network without major capital expenditures. Operators can extend the operational life of equipment in their network.

VPLS uses the provider core to join multiple attachment circuits together to simulate a virtual bridge that connects the multiple attachment circuits together.

Upon completing this lesson, you will be able to:

- Describe the VPLS feature
- Describe the VPLS characteristics
- Describe the VPLS packet forwarding
- Describe the VPLS PE view
- Describe the Virtual Switch Interface and the Virtual Forwarding Instance (VFI)
- Describe the Bridge-Group Virtual Interface (BVI)
- Describe VPLS implementation procedure
- Describe a VPLS configuration example interconnecting VLAN 111
- Describe a VPLS configuration example with VLAN tag rewrite
- Describe a routed VPLS configuration example
- Describe how to verify VPLS operation in Cisco IOS XR
- Describe how to monitor VPLS bridge domain operation in Cisco IOS XE

https://t.me/learningnets

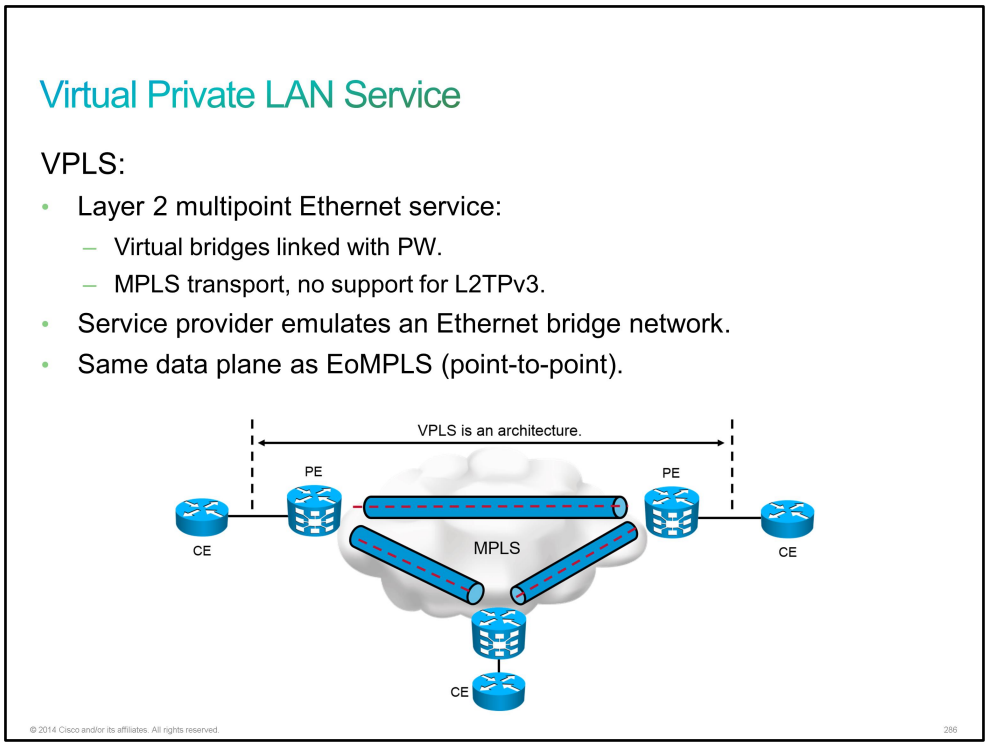
- Describe how to monitor VPLS VFI operation in Cisco IOS XE
- Describe how to monitor attachment circuits in a VPLS

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Virtual Private LAN Service

The purpose of VPLS is to provide a private multipoint LAN-type Ethernet connectivity service.



VPLS emulates a LAN segment over an MPLS backbone across PWs or VCs. VPLS creates one or more LANs for each customer who is using the service from the service provider. Each LAN is completely separate from the other emulated LAN segments. When a customer with different Ethernet sites connects to an MPLS backbone where VPLS is deployed, it appears as if all the sites are interconnected through a virtual Ethernet switch.

For each VPLS, the PE routers are fully meshed with PWs. A PE receiving a frame from another PE can identify which VPLS the frame belongs to, based on a PW label or VC label. As far as each customer is concerned, an Ethernet frame that is sent into the service provider network is delivered to the correct site(s) based on the destination MAC address. It is the task of each PE router to inspect the destination MAC address of each frame arriving from a locally attached site and to forward it to the appropriate destination site. This destination site may be attached to the same PE on a different port or a remote PE. If the destination site is attached to the same PE, the PE locally switches the frame to the correct port. If the destination site is attached to a remote PE, the ingress PE must forward the frame to the appropriate PW to the remote PE. This means that the ingress PE needs to know which egress PE to send the frame to.

There are two ways in which this can be accomplished. One is to have a control plane signaling to carry information about MAC addresses between PEs; another is to have a scheme based on MAC address learning. VPLS takes the latter approach by having each PE take the responsibility for learning which remote PE is associated with a given MAC address. This way, an ingress PE simply needs to identify which frames need to be sent to egress PEs, and egress PEs take care of identifying which local ports to forward the packet to. By inspecting the source MAC address of the frame arriving on a port, whether an actual local port or a PW from a remote PE, and by creating a corresponding entry in the forwarding table, the PE learns where to send future frames with that destination MAC address.

https://t.me/learningnets

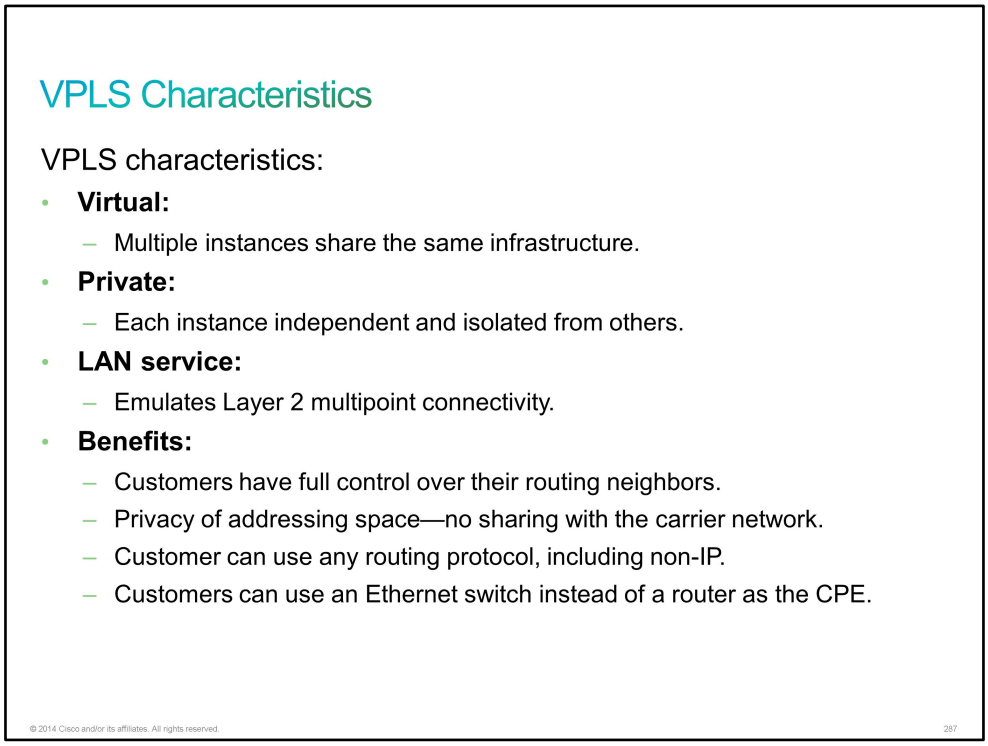
If Ethernet switches are used as CE devices and connected to PE routers, the PEs need to learn the MAC addresses of individual hosts attached to the switches. So, if a host is plugged into the office network served by a switch as a CE, the effect will be felt by all PEs. Thus, for a large deployment, it is better to use routers as CEs than switches.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

VPLS Characteristics

VPLS has become a very attractive technology over the past few years with the advent of MPLS. The reason for this is that some enterprises are very reluctant to relinquish the routing control of their network to the service provider, and they desire Layer 2 VPN services with multipoint connectivity.



VPLS allows service providers to deploy carrier-class service over Ethernet and MPLS-based networks in a reliable and flexible way.

The term implies these characteristics:

- It is “virtual” because multiple instances of this service share the same physical infrastructure.
- It is “private” because each instance of the service is independent from the others.
- It is “LAN service” because it emulates Layer 2 multipoint connectivity between subscribers.

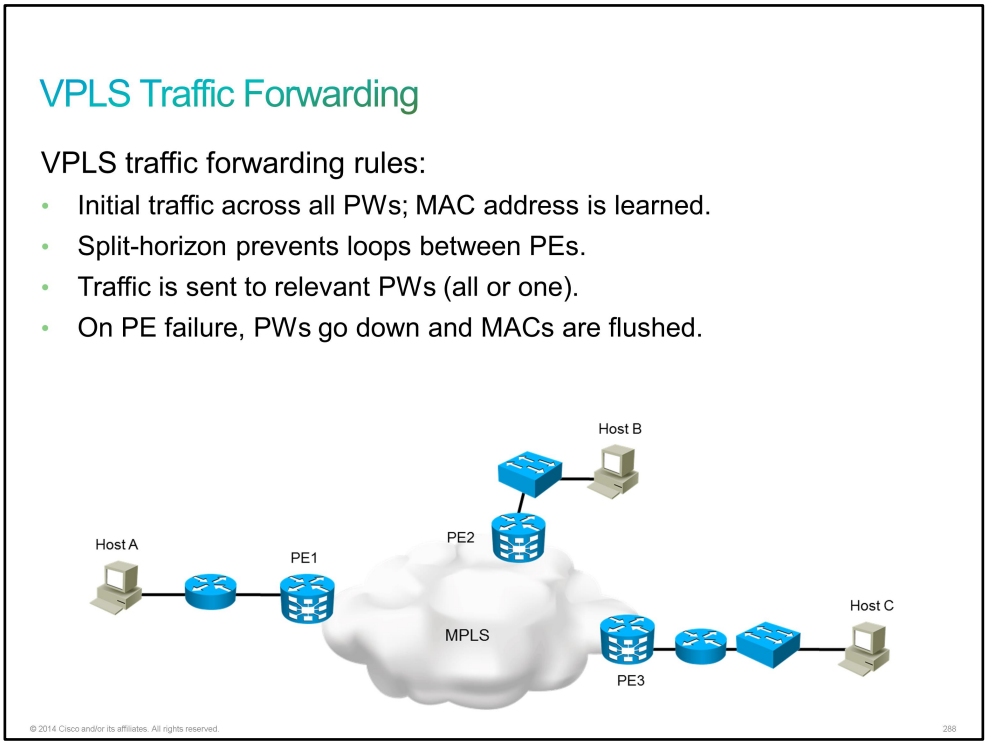
The main VPLS benefits include the following:

- Customers have full operational control over their routing neighbors.
- Privacy of addressing space means that addresses of the carrier infrastructure are completely isolated.
- Customers have a choice of using any routing protocol, including non-IP protocols such as IS-IS.
- Customers can use an Ethernet switch instead of a router as the CPE device.

https://t.me/learningnets

VPLS Traffic Forwarding

The forwarding of unicast frames differs from the forwarding of broadcast and multicast frames.



Forwarding of Unicast Frames

Suppose that the MAC address of host C is C and the MAC address of host B is B, for the customer network in the figure. Suppose host C sends a frame with source MAC address C and destination MAC address B. Suppose that PE3 does not know the location of MAC address B. As a learning bridge would do, PE3 floods the packet on all ports except the port on which it arrived. This means the packet is flooded to the PW to PE2 and the PW to PE1.

PE1 and PE2 know that the packet belongs to the customer VPLS, by virtue of the PW on which the frame arrived. PE1 and PE2 both perform destination MAC address lookup in their VPLS forwarding tables corresponding to this customer. If PE1 does not know the location of MAC address B, it floods the frame on its local ports to the CE. However, it does not flood the frame to any other PE routers. This split-horizon scheme ensures no forwarding loops occur. Similarly, PE2 forwards the frame to the port facing the switch CE.

Receiving frames with MAC address C enables each PE to learn the location of host C. Thus, PE2 and PE1 create an entry in their forwarding tables with an association between MAC address C and their respective PWs to PE3. In this way, all PEs learn the MAC addresses and create an association between MAC addresses and PWs (for remote destinations) in their forwarding tables for that particular VPLS instance.

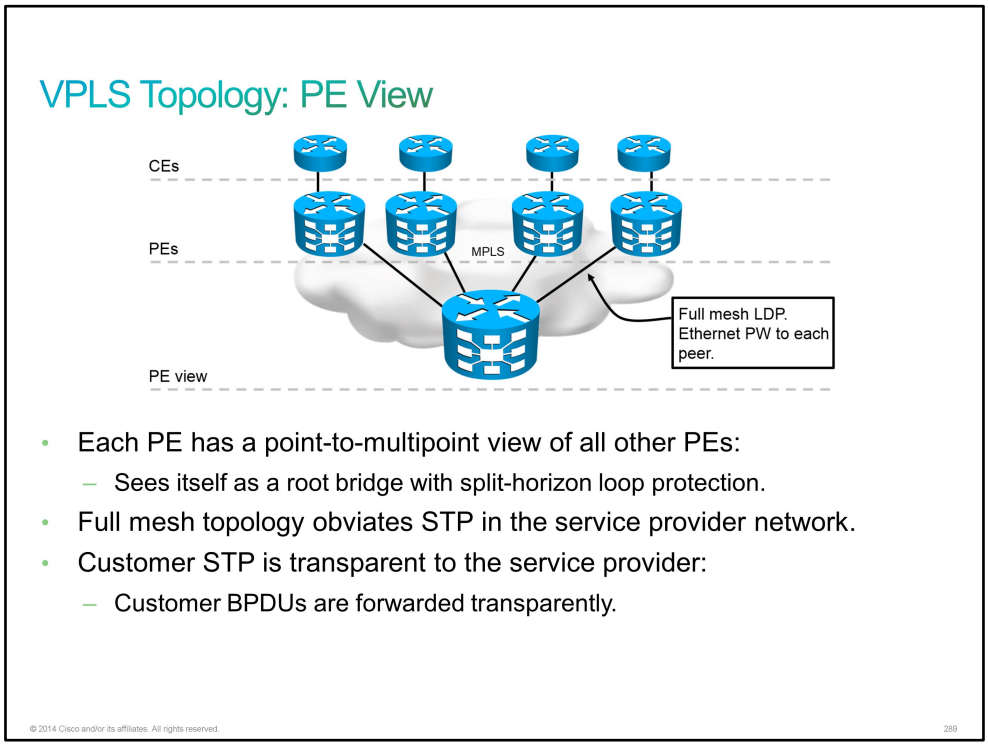
Forwarding of Broadcast and Multicast Frames

Suppose PE3 receives a broadcast frame sent by host C. The frame must be sent to all sites of the customer VPLS. PE3 floods the frame on PWs to PE1 and PE2. In turn, PE1 and PE2 flood the frame to the attached CEs, but due to split horizon, do not send the frames to any PEs. Multicast traffic is also treated this way.

https://t.me/learningnets

VPLS Topology: PE View

The full mesh of PWs in the service provider network allows the implementation of the split-horizon principle that is similar to the split horizon of an internal BGP mesh.



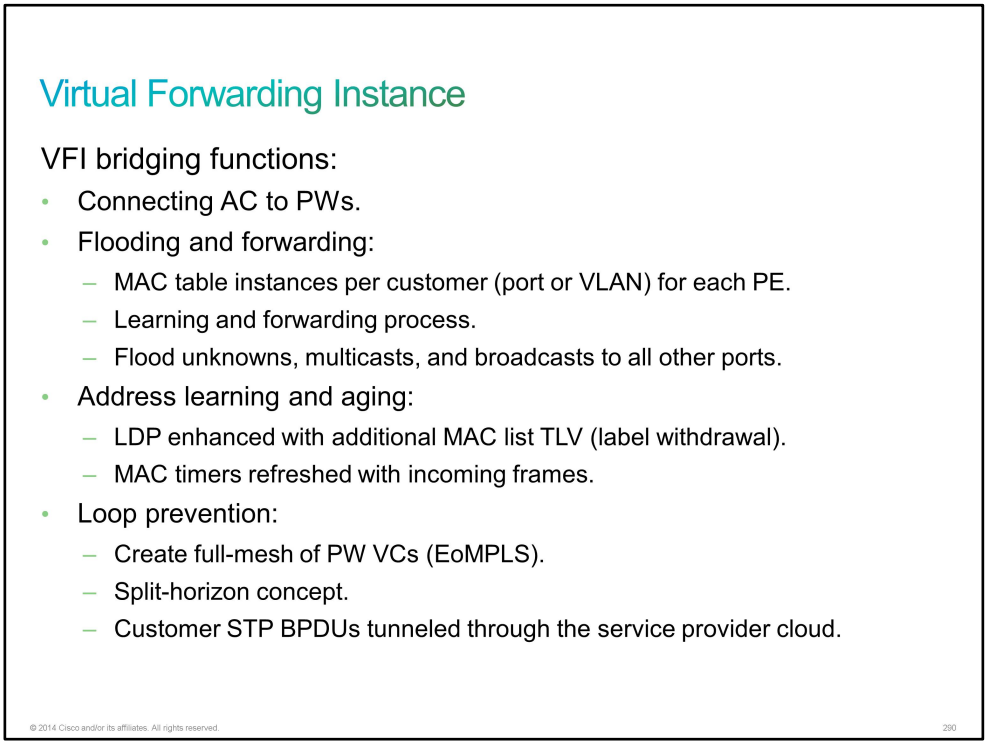
Frames received over a PW are generally not forwarded to another PW. This concept provides loop-free forwarding. Therefore, there is no need for STP in the service provider network. The customer STP BPDUs are transparently tunneled over the PWs to detect loops on the customer layer.

Note	The split-horizon rule must be disabled in certain environments, such as H-VPLS.

https://t.me/learningnets

Virtual Forwarding Instance

The VSI, implemented as a VFI, is a virtual Layer 2 Forwarding entity that defines the VPLS domain membership and resembles virtual switches on PE routers.



A VPLS domain consists of Ethernet interfaces or VLANs that belong to the same (virtual) LAN but are connected to multiple PE devices. The VSI learns remote MAC addresses and is responsible for proper forwarding of the customer traffic to the appropriate end nodes. It is also responsible for guaranteeing that each VPLS domain is loop-free. The VSI is responsible for several functions, namely MAC address management, dynamic learning of MAC addresses on physical ports and VCs, aging of MAC addresses, MAC address withdrawal, flooding, and data forwarding.

All PWs in a VFI are placed by default into the same split-horizon group, which effectively prevents traffic from forwarding to other PWs in the same VFI.

With VPLS Layer 2 VPNs, the customer can connect with a switch or a router.

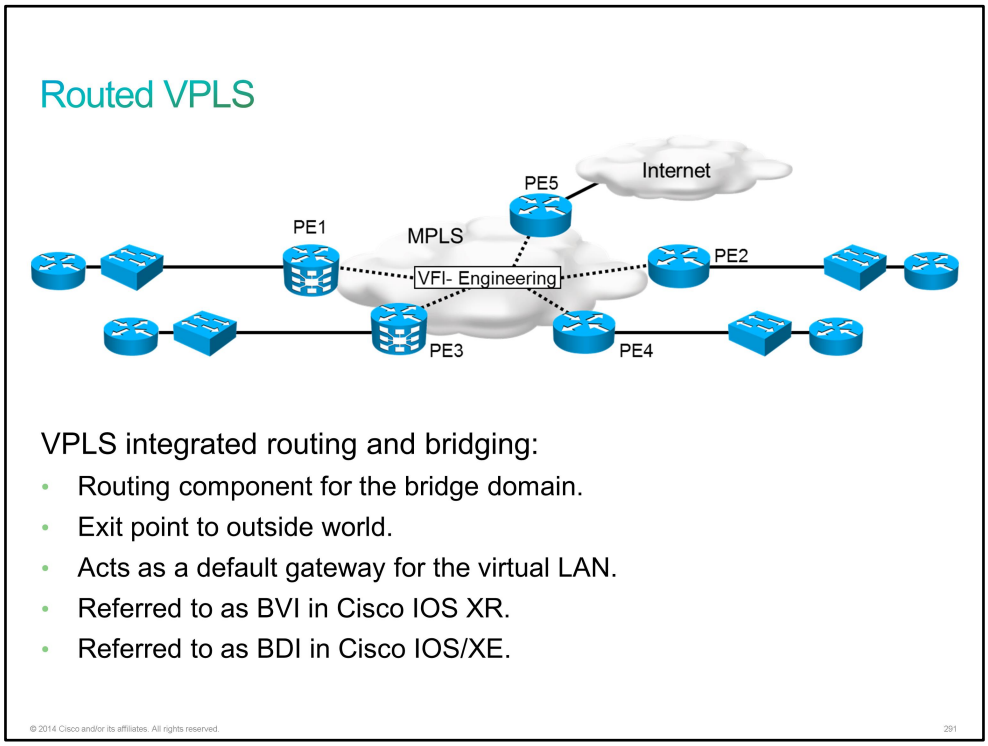
If connecting with a router, the VPLS just looks like a switch to the routing protocols on each side. No MAC learning is done beyond the MAC address of the directly connected CE router interface.

If connecting at Layer 2 with a switch, then MAC learning on the PE of all the MAC addresses on the directly connected customer LAN is possible. This is a situation where MAC limiting or another filtering method would be implemented, or a PBB(802.1ah) can be used.

https://t.me/learningnets

Routed VPLS

Configuration of a bridge-domain creates a Layer 2 domain. In order to exit that Layer 2 domain, connect Layer 3 routers that route between hosts inside the bridge-domain and the outside world.



The routing exchange with the outside world could be provided by a customer device inside a Layer 2 bridge domain.

Alternatively, you can configure a VPLS peer to perform the routing function for the bridge domain. To do this, you need to configure a Layer 3 interface that plugs into a bridge-domain in order to route packets in and out of the bridge-domain.

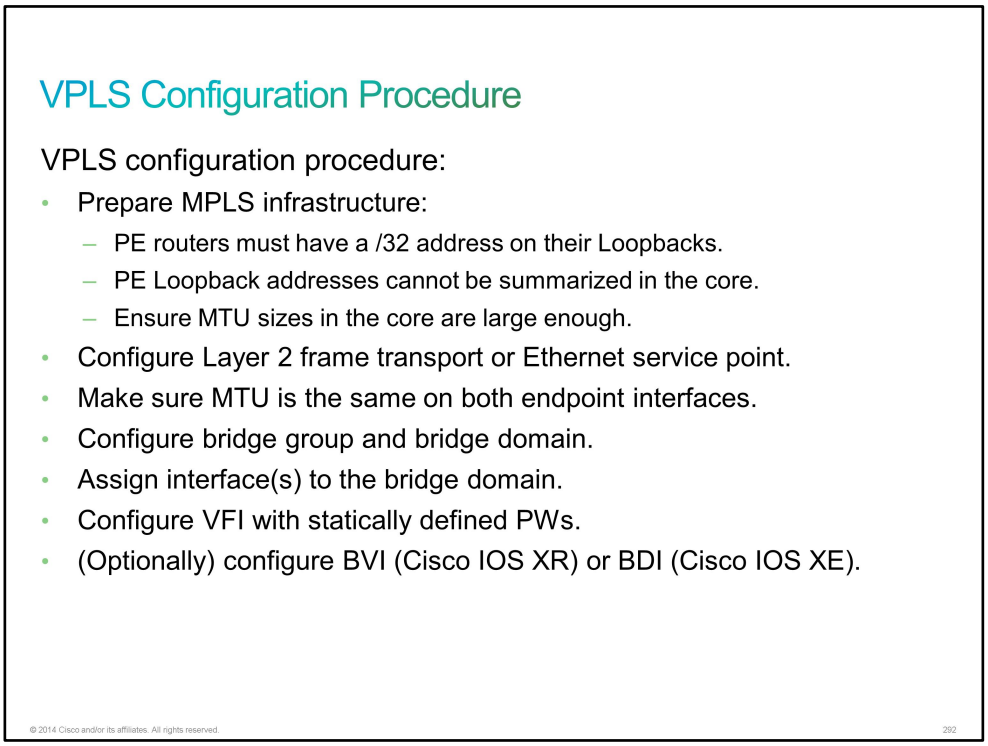
In Cisco IOS XR the routed VPLS feature is implemented using the BVI. In Cisco IOS XE the routing instance for the virtual bridge domain is referred to as the BDI.

VPLS integrated routing and bridging is also known as routed PW and routed VPLS.

https://t.me/learningnets

VPLS Configuration Procedure

You need to prepare the MPLS infrastructure before implementing VPLS.



The VPLS deployment process includes these steps:

- PE routers must have a /32 address on their Loopback interfaces. This is required for binding the VC label and the transport LDP into a label stack.
- PE Loopback addresses cannot be summarized in the core. Aggregation in the network would break the LSP path.
- MPLS must be enabled in the core.
- MTU sizes on the core links must be able to accommodate the customer MTU with the added label overhead.

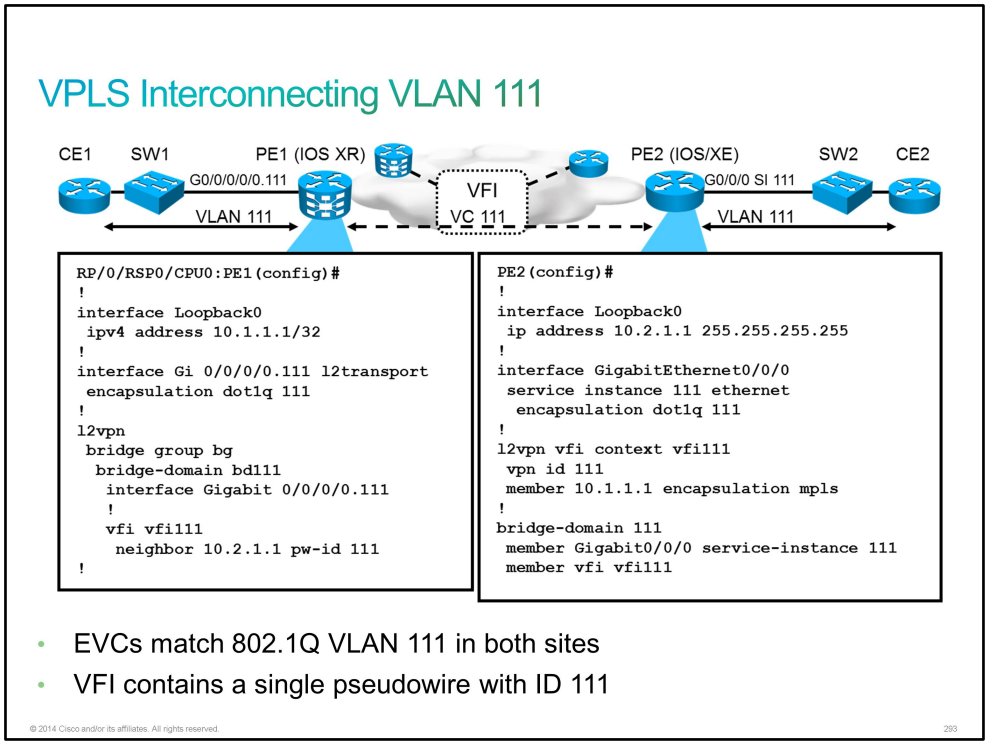
Once the MPLS infrastructure is ready, you will:

- Configure Layer 2 frame transport or Ethernet service instance.
- Make sure the MTU is the same on both endpoint interfaces.
- Configure the bridge group and bridge domain.
- Assign interfaces to bridge domain.
- Configure the VFI with statically defined PWs.
- Optionally configure routing for the bridge domain using BVI (Cisco IOS XR) or BDI (Cisco IOS XE).

https://t.me/learningnets

VPLS Interconnecting VLAN 111

This example shows a basic VPLS implemented on Cisco IOS XR and Cisco IOS XE platforms.



The attachment circuit on the Cisco IOS XR PE is enabled for Layer 2 transport using the **l2transport** keyword. in Cisco IOS XE you configure the attachment circuit using a service instance configured under the main interface.

The Cisco IOS XR VPLS configuration is performed within the Layer 2 VPN (**l2vpn** command), bridge group, and bridge domain configuration mode. The bridge domain contains two main configuration elements: a list of the local attachment circuits, and the VFI that contains the PWs.

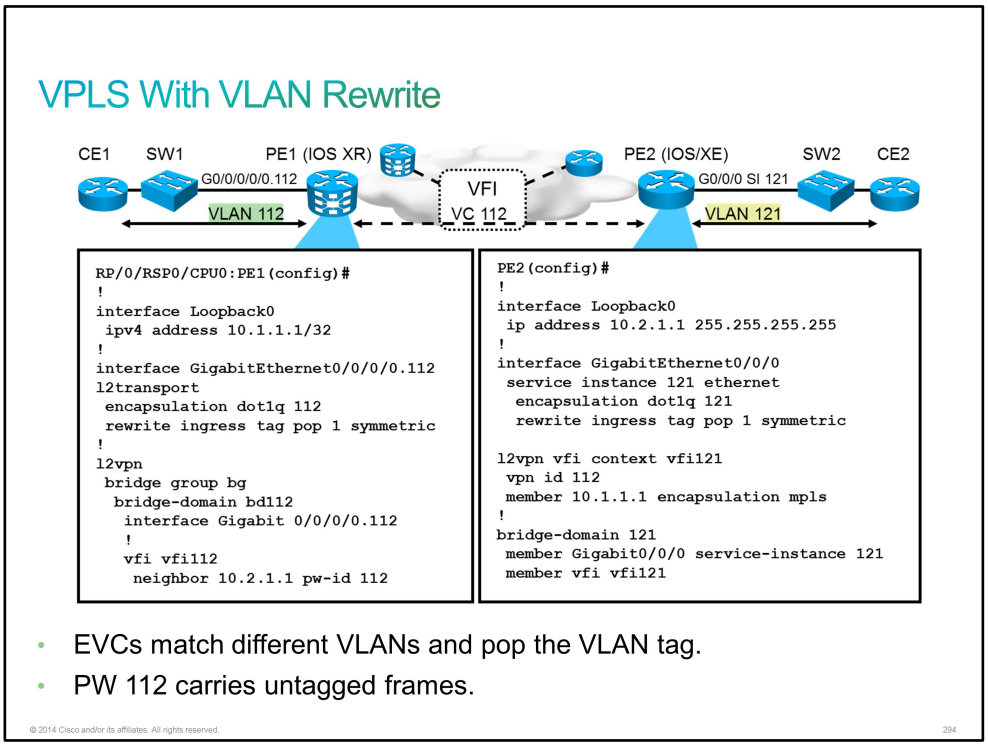
The Cisco IOS XE VPLS configuration consists of two constructs: the VFI and the bridge domain. The VFI is the set of pseudowires configured with the **member** command. The VPN ID configured in the VFI mode must match up with the VC ID configured in Cisco IOS XR with the **neighbor** command. The bridge domain links the attachment circuits with the VFI. The attachment circuit (service instance 111 on GigabitEthernet0/0/0) and the VFI (vfi111) are configured as members of the bridge domain (111).

The names of the bridge group and domain have local significance.

https://t.me/learningnets

VPLS With VLAN Rewrite

This scenario shows two distinct VLANs interconnected via VPLS.



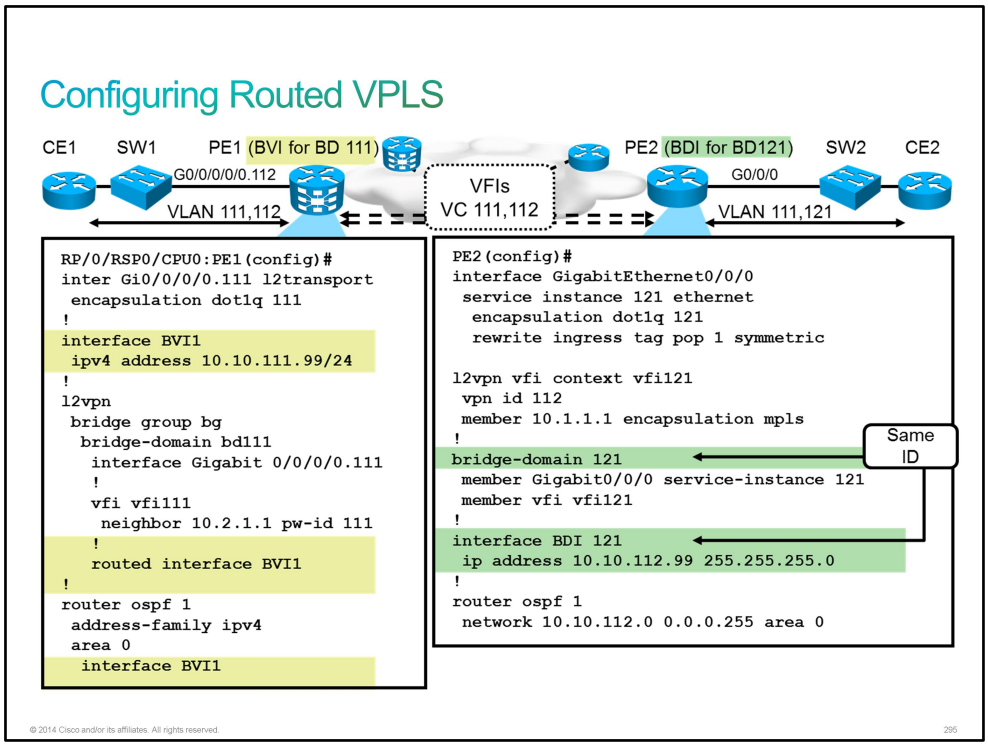
VLAN 112 is connected to PE1 (Cisco IOS XR). VLAN 121 is connected to PE2 (Cisco IOS XE). Each PE has a bridge domain that links the local attachment circuit with the VFI. The VFI is configured with a static configuration of the PW to the peer PE.

Each PE is performing a 802.1Q tag rewrite operation on its local attachment circuit. PE1 pops one 802.1Q tag on its GigabitEthernet0/0/0/0.112 interface. PE2 pops one tag on the service instance 121 of the GigabitEthernt0/0/0 interface. The frames exchanged between the VLAN 112 and 121 are thus forwarded untagged over the PW 112 connecting PE1 and PE2.

https://t.me/learningnets

Configuring Routed VPLS

This example shows how to add a routing component to each bridge domain configured on the PE routers.



PE1 (Cisco IOS XR) acts as the default gateway for the bridge domain that interconnects VLAN 111. PE2 (Cisco IOS XE) acts as the default gateway for the bridge domain that interconnects VLANs 112 and 121.

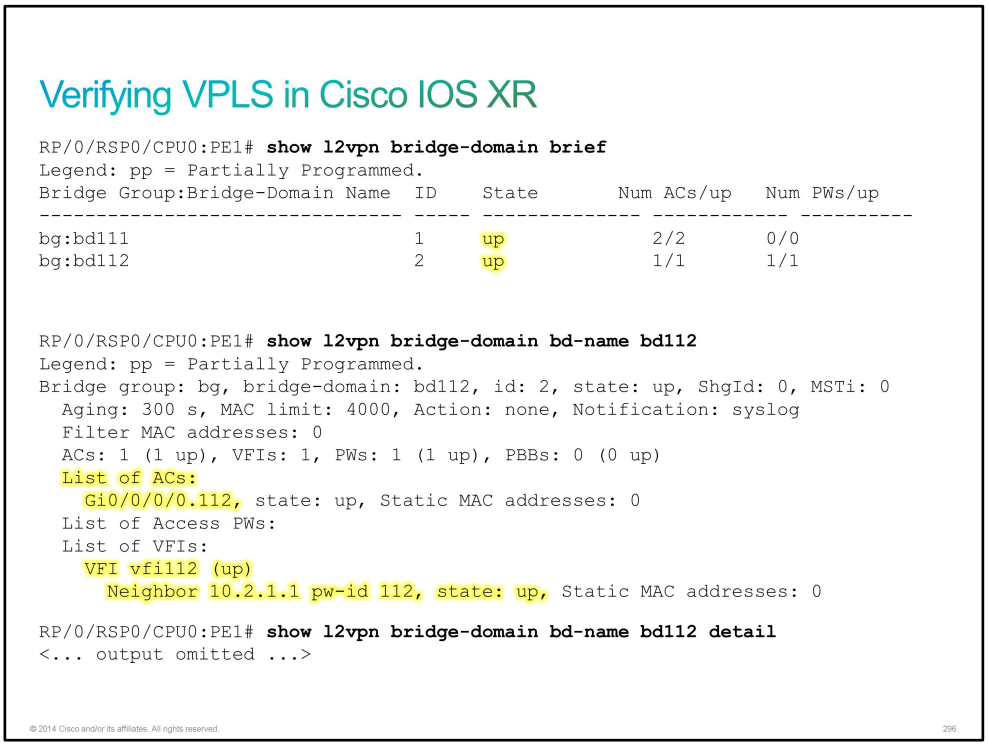
In Cisco IOS XR you need to create a BVI interface, assign it an IP address, and add it to the bridge domain using the **routed interface** command. You will typically include the BVI interface in a routing process of your choice, although static routing is also an option.

In Cisco IOS XE you have to define a BDI with the ID of the bridge domain that the BDI should serve. The BDI requires an IP address and can be included in a routing process.

https://t.me/learningnets

Verifying VPLS in Cisco IOS XR

In Cisco IOS XR you can monitor the VPLS operation using various detail levels. The **show l2vpn bridge-domain** command allows you to choose brief, standard, or detail output.



The brief output provides a summary of the bridge domains and their state. The regular output can be narrowed to a specific bridge domain and displays information about the ACs, VFIs, and PW.

The detail output example for the bridge domain 112 is shown below:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1# show l2vpn bridge-domain bd-name bd112 detail
Legend: pp = Partially Programmed.
Bridge group: bg, bridge-domain: bd112, id: 2, state: up, ShgId: 0, MSTi: 0
  Coupled state: disabled
  MAC learning: enabled
  MAC withdraw: enabled
    MAC withdraw for Access PW: enabled
    MAC withdraw sent on: bridge port up
    MAC withdraw relaying (access to access): disabled
  Flooding:
    Broadcast & Multicast: enabled
    Unknown unicast: enabled
  MAC aging time: 300 s, Type: inactivity
  MAC limit: 4000, Action: none, Notification: syslog
  MAC limit reached: no
  MAC port down flush: enabled
  MAC Secure: disabled, Logging: disabled
  Split Horizon Group: none
  Dynamic ARP Inspection: disabled, Logging: disabled
  IP Source Guard: disabled, Logging: disabled
  DHCPv4 snooping: disabled
  IGMP Snooping: enabled
  IGMP Snooping profile: none
  MLD Snooping profile: none
  Storm Control: disabled
  Bridge MTU: 1500
  MIB cvplsConfigIndex: 3
  Filter MAC addresses:
  P2MP PW: disabled
  Create time: 09/06/2014 10:22:00 (4d11h ago)
  No status change since creation
  ACs: 1 (1 up), VFIs: 1, PWs: 1 (1 up), PBBs: 0 (0 up)
  List of ACs:
    AC: GigabitEthernet0/0/0/0.112, state is up
      Type VLAN; Num Ranges: 1
      VLAN ranges: [112, 112]
      MTU 1500; XC ID 0x1080014; interworking none
      MAC learning: enabled
      Flooding:
        Broadcast & Multicast: enabled
        Unknown unicast: enabled
      MAC aging time: 300 s, Type: inactivity
      MAC limit: 4000, Action: none, Notification: syslog
      MAC limit reached: no
      MAC port down flush: enabled
      MAC Secure: disabled, Logging: disabled
      Split Horizon Group: none
      Dynamic ARP Inspection: disabled, Logging: disabled
      IP Source Guard: disabled, Logging: disabled
      DHCPv4 snooping: disabled
      IGMP Snooping: enabled
      IGMP Snooping profile: none
      MLD Snooping profile: none
      Storm Control: disabled
      Static MAC addresses:
      Statistics:
        packets: received 6360, sent 41561
```


https://t.me/learningnets

```
      bytes: received 461086, sent 3926940
    Storm control drop counters:
      packets: broadcast 0, multicast 0, unknown unicast 0
      bytes: broadcast 0, multicast 0, unknown unicast 0
    Dynamic ARP inspection drop counters:
      packets: 0, bytes: 0
    IP source guard drop counters:
      packets: 0, bytes: 0
List of Access PWs:
List of VFIs:
  VFI vfi112 (up)
    PW: neighbor 10.2.1.1, PW ID 112, state is up ( established )
      PW class not set, XC ID 0xc0000006
      Encapsulation MPLS, protocol LDP
      Source address 10.1.1.1
      PW type Ethernet, control word disabled, interworking none
      Sequencing not set

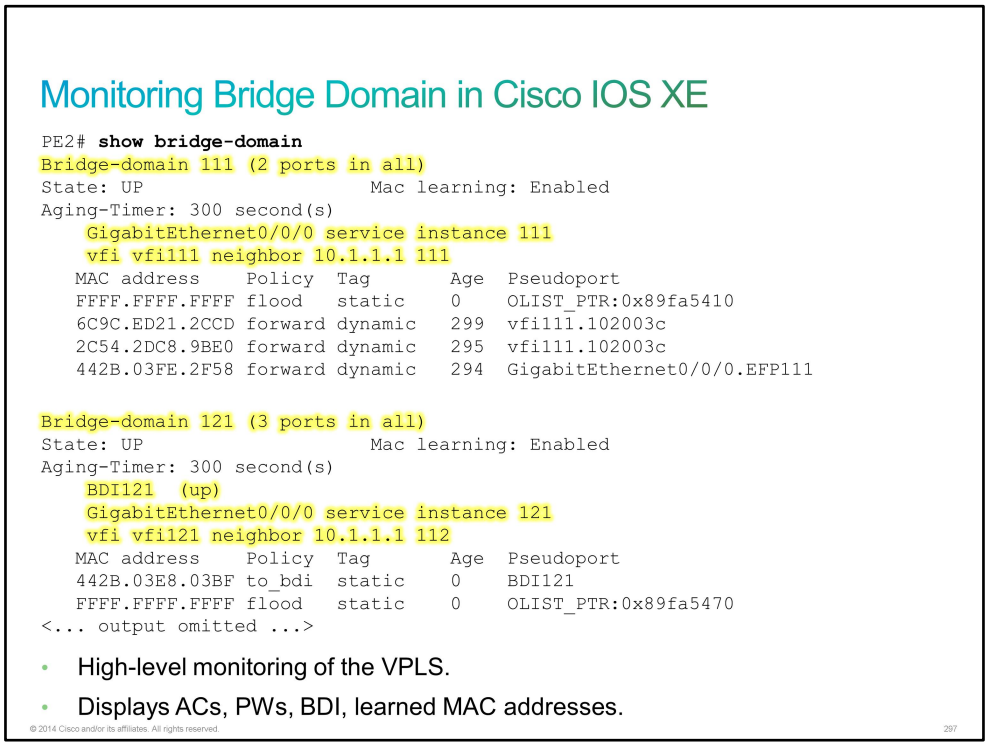
      PW Status TLV in use
        MPLS         Local                          Remote
        ------------ ------------------------------ ------------------------
        Label        16011                          37
        Group ID     0x2                            0x0
        Interface    vfi112                         unknown
        MTU          1500                           1500
        Control word disabled                       disabled
        PW type      Ethernet                       Ethernet
        VCCV CV type 0x2                            0x2
                     (LSP ping verification)        (LSP ping verification)
        VCCV CC type 0x6                            0x6
                     (router alert label)           (router alert label)
                     (TTL expiry)                   (TTL expiry)
        ------------ ------------------------------ ------------------------
      Incoming Status (PW Status TLV):
        Status code: 0x0 (Up) in Notification message
      MIB cpwVcIndex: 3221225478
      Create time: 09/06/2014 10:22:00 (4d11h ago)
      Last time status changed: 09/06/2014 10:28:51 (4d11h ago)
      MAC withdraw messages: sent 0, received 0
      Static MAC addresses:
      Statistics:
        packets: received 41561, sent 6112
        bytes: received 3760696, sent 418630
      Storm control drop counters:
        packets: broadcast 0, multicast 0, unknown unicast 0
        bytes: broadcast 0, multicast 0, unknown unicast 0
    DHCPv4 snooping: disabled
    IGMP Snooping profile: none
    MLD Snooping profile: none
    VFI Statistics:
      drops: illegal VLAN 0, illegal length 0
```

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Monitoring Bridge Domain in Cisco IOS XE

The two main VPLS monitoring methods in Cisco IOS XE include the verification of the bridge domain and of the VFI.



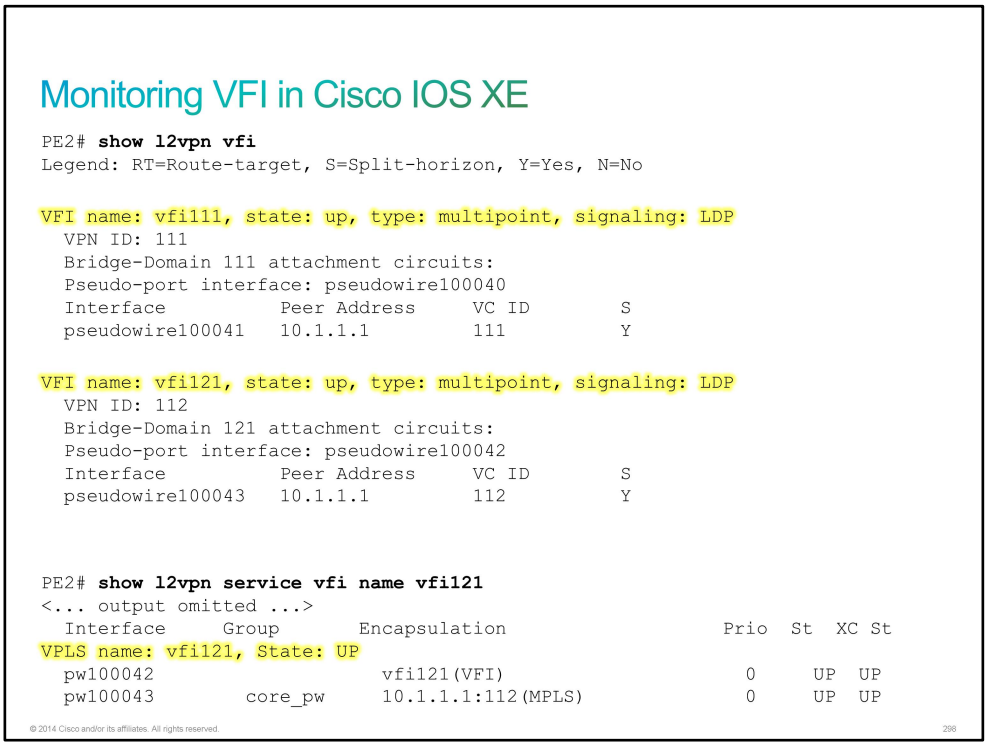
From the high level you can examine the VPLS operation using the **show bridge-domain** command. Optionally you can identify a specific bridge domain. The output includes information about the statistics, attachment circuits, VFI, and MAC addresses learned in the bridge domain.

In this example, bridge domain 121 has a BDI associated with it, which provides routing functionality.

https://t.me/learningnets

Monitoring VFI in Cisco IOS XE

In Cisco IOS XE you can monitor the VFI operation by using the **show l2vpn vfi** or the **show l2vpn service vfi** command. Each command allows you to identify a specific VFI to be examined.

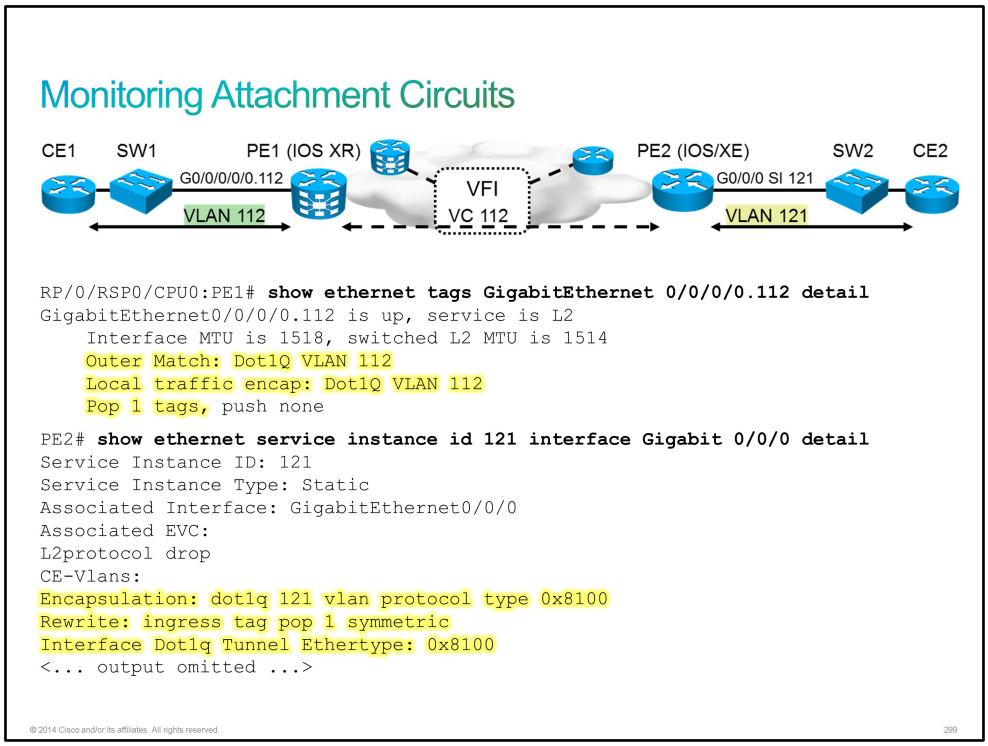


The VFI-related information provided by both commands is similar and reports the state of all PW that build the VFI.

https://t.me/learningnets

Monitoring Attachment Circuits

You can monitor the tag matching and manipulation at the attachment circuits.

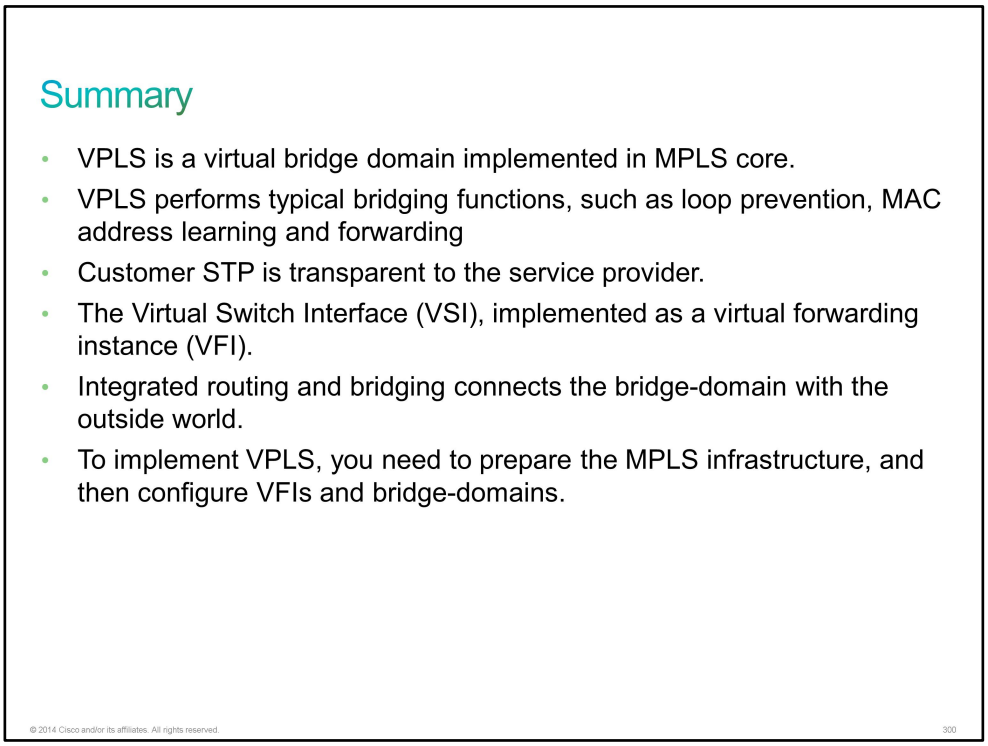


Use the **show ethernet tags** (Cisco IOS XR) and **show ethernet service instance** (Cisco IOS XE) commands for this purpose. The detailed output displays the outer and inner tag matching, the encapsulation protocol, and the tag manipulation.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

Lesson 4: Enhancing Layer 2 VPN Scalability

Overview

Large service provider environments often use highly concentrated points of presence where multiple customer connections are terminated. H-VPLS reduces both signaling and replication overhead by using both full-mesh as well as hub and spoke configurations. It is a widely deployed method of increasing the scalability of the VPLS solutions.

You can further enhance the VPLS scalability by enabling BGP autodiscovery of VPLS peers instead of configuring all pseudowires (PWs) individually.

Upon completing this lesson, you will be able to:

- Describe H-VPLS architectures
- Compare VPLS with H-VPLS
- Describe the H-VPLS edge technologies 802.1ad and 802.1QinQ
- Discuss the use of MPLS in the H-VPLS edge
- Describe the IEEE 802.1ad, 802.1ah, and 802.1ag Carrier Ethernet standards
- Evaluate VPLS and H-VPLS
- Discuss VPLS neighbor discovery and signaling
- Describe the H-VPLS and BGP autodiscovery implementation procedure
- Describe the configuration of switches acting as UPEs in H-VPLS with edge
- Describe an H-VPLS with edge Q-in-Q configuration example for interconnecting VLAN
- Describe an H-VPLS with edge Q-in-Q configuration example for interconnecting VLANs 112 and 121
- Describe the configuration of switches acting as UPEs in H-VPLS with edge 802.1ad
- Describe an H-VPLS with edge 802.1ad configuration example for interconnecting VLAN 111

https://t.me/learningnets

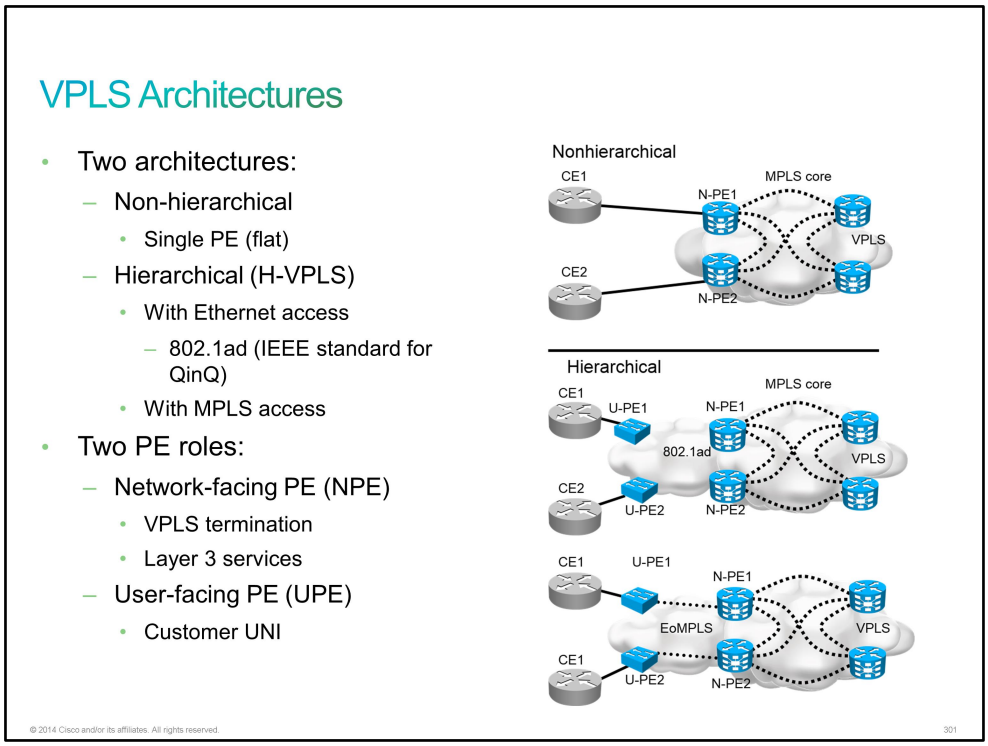
- Describe an H-VPLS with edge 802.1ad configuration example for interconnecting VLANs 112 and 121
- Describe a configuration of H-VPLS with BGP autodiscovery
- Describe how to verify an H-VPLS bridge domain
- Describe how to monitor attachment circuits in an H-VPLS
- Describe how to verify BGP capability exchange for BGP autodiscovery
- Describe how to verify BGP autodiscovery database
- Describe how to verify BGP autodiscovery attributes

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

VPLS Architectures

H-VPLS topologies have been developed to address the scaling limitations within flat VPLS architecture.



VPLS requires a full mesh of tunnel LSPs between all PE routers that participate in the VPLS service. For each VPLS service, n*(n-1)/2 pseudowires must be set up between PEs. This creates signaling overhead. The packet replication requirement deters large-scale deployment. Hierarchical connectivity reduces signaling and replication overhead.

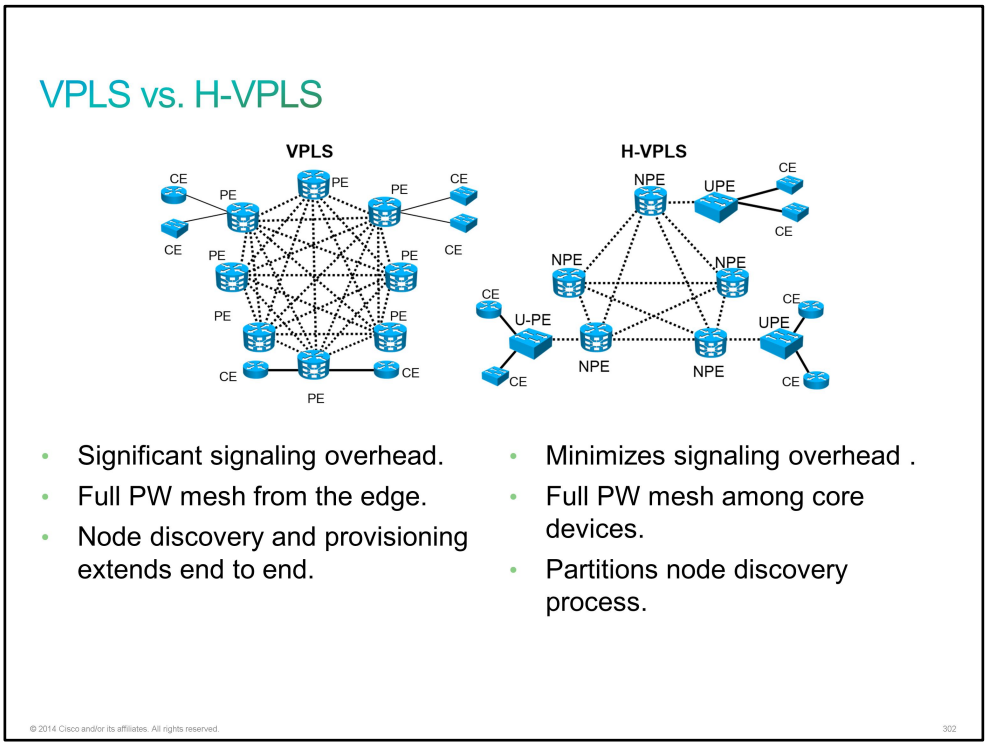
There are two access mechanisms in H-VPLS: Ethernet and MPLS. Ethernet-based access uses the IEEE 802.1ad standard or its predecessor, 802.1 QinQ.

H-VPLS defines user-facing PE (UPE) routers and network-facing PE (NPE) routers.

https://t.me/learningnets

VPLS vs. H-VPLS

H-VPLS removes the scalability constraints by limiting the signaling overhead and partitioning the discovery process.

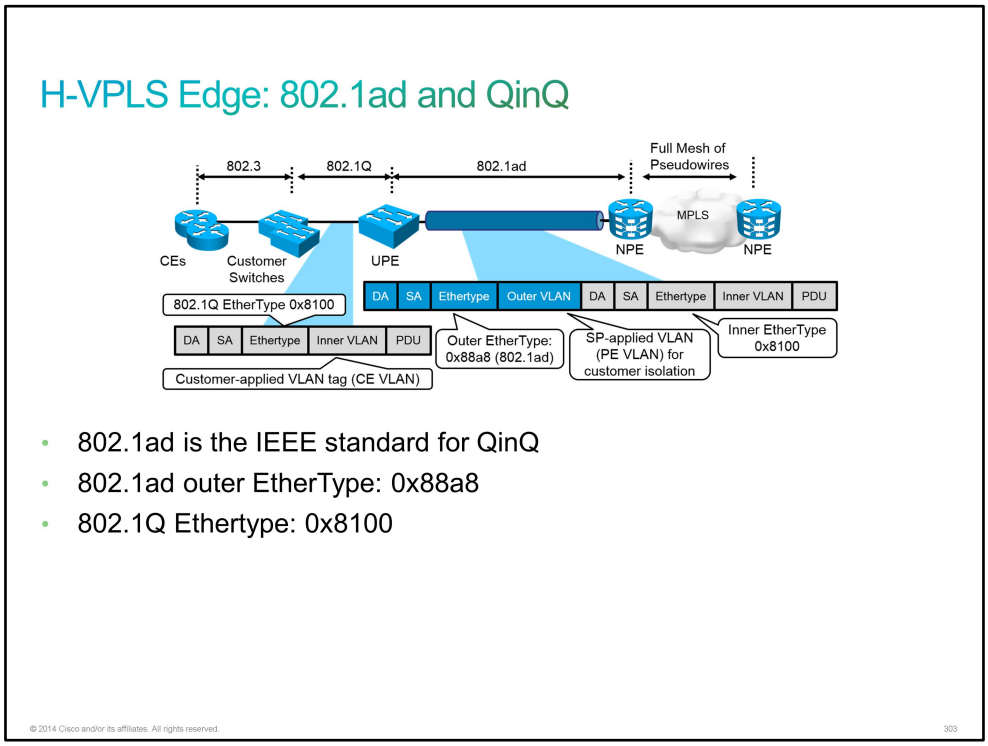


The H-VPLS hierarchy provides the benefits of less signaling in an MPLS core network and less packet replication on NPE routers. The UPE routers have an aggregation role and do some packet replication and MAC address learning.

https://t.me/learningnets

H-VPLS Edge: 802.1ad and QinQ

QinQ and the standards-based IEEE 802.1ad are closely related Ethernet networking technologies. They are also known as provider bridging or stacked VLANs.



QinQ and 802.1ad specify architecture and bridge protocols to provide separate instances of the MAC services to multiple independent users in a manner that does not require cooperation among the users, and requires a minimum of cooperation between the users and the provider. The idea is to provide the possibility for customers to run their own VLANs inside the VLAN of a service provider. This way, the service provider can just configure one VLAN for the customer, and the customer can then treat that VLAN like a trunk.

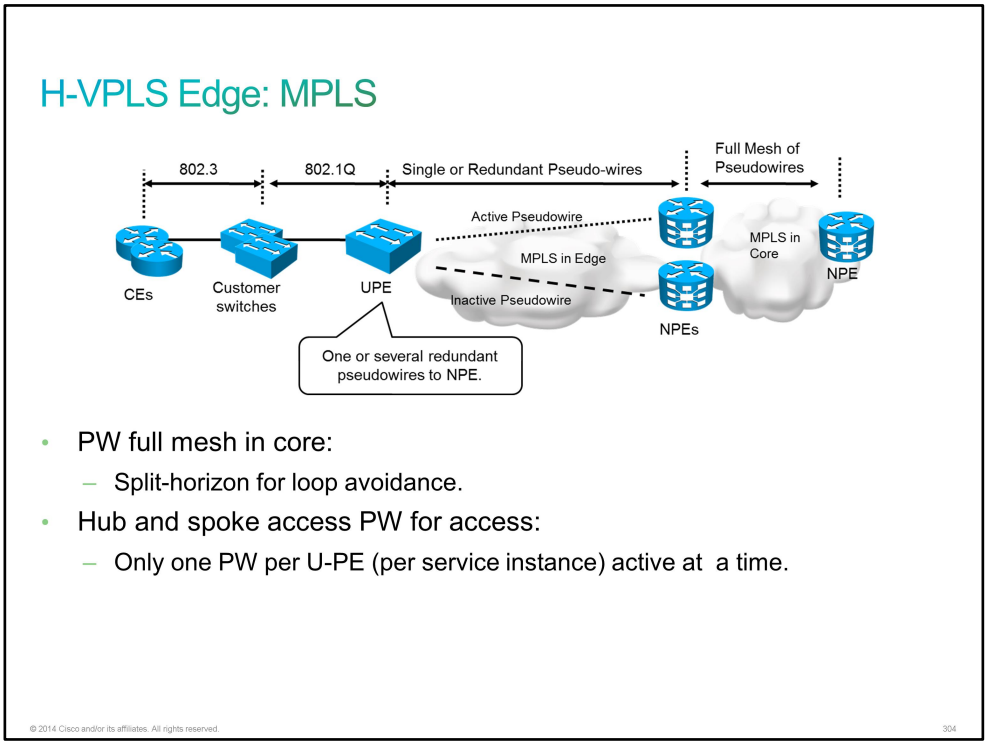
QinQ and 802.1ad have these characteristics:

- The addition of a second tag allows new set of operations. Having multiple tags—the tag stack—allows switches to more easily modify frames. In a tag stack scheme, switches can add (“push”), remove (“pop”) or modify single or multiple tags.
- A multitagged frame not only has multiple VLAN IDs, but has multiple EtherTypes and other VLAN header bit fields.
- A tag stack creates a mechanism for Internet service providers to encapsulate customer single-tagged 802.1Q traffic with a single tag, the final frame being a QinQ or 802.1ad frame. The outer tag is used to identify and segregate traffic from different customers; the inner tag is preserved from the original frame.
- The EtherType values are 0x88a8 (for 802.1ad) and 0x8100 (for QinQ). The EtherType for QinQ is the same as for 802.1Q.

https://t.me/learningnets

H-VPLS Edge: MPLS

The second approach uses MPLS in the edge network, instead of QinQ or 802.1ad.



In the MPLS-based edge, the VPLS core PWs (hub) are augmented with access PWs (spoke) to form a two-tier H-VPLS. The customer sites are connected to UPE devices. The UPE devices have single or redundant connection (PW) to NPE routers. The NPE routers are connected in a basic VPLS full mesh service. For each VPLS service, a single spoke PW is set up between UPE and NPE devices. Unlike traditional PWs that terminate on a physical or logical port, a spoke PW terminates on a VSI on UPE and NPE devices. The UPEs and NPEs treat each spoke connection like an attachment circuit of the VPLS service. The PW label is used to associate the traffic from the spoke PW to a VPLS instance.

UPE Operation

The UPE device supports Layer 2 switching and does all the normal bridging functions of learning and replication on all its ports, including the spoke PW. Packets to an unknown destination are replicated to all ports in the service, including the spoke PW. Once the MAC addresses of CE devices connected to the same UPE device are learned, traffic between them is switched locally, saving the capacity of the spoke PW to NPEs. Similarly, traffic between remotely connected CE devices to different UPE devices is switched directly onto the spoke PW and sent to NPEs over the point-to-point PW. Since the UPE is bridging-capable, only a single PW is required per VPLS instance for any number of access connections in the same VPLS service. This reduces the signaling overhead between UPEs and NPEs.

https://t.me/learningnets

NPE Operation

An NPE device supports all bridging functions for VPLS service and supports the routing and MPLS encapsulation as in basic VPLS. The operation of the NPE is independent of the type of device at the other end of the spoke PW. Thus the NPE will switch traffic between the spoke PW, hub PWs, and attachment circuits once it has learned the MAC addresses.

Dual-Homed UPE

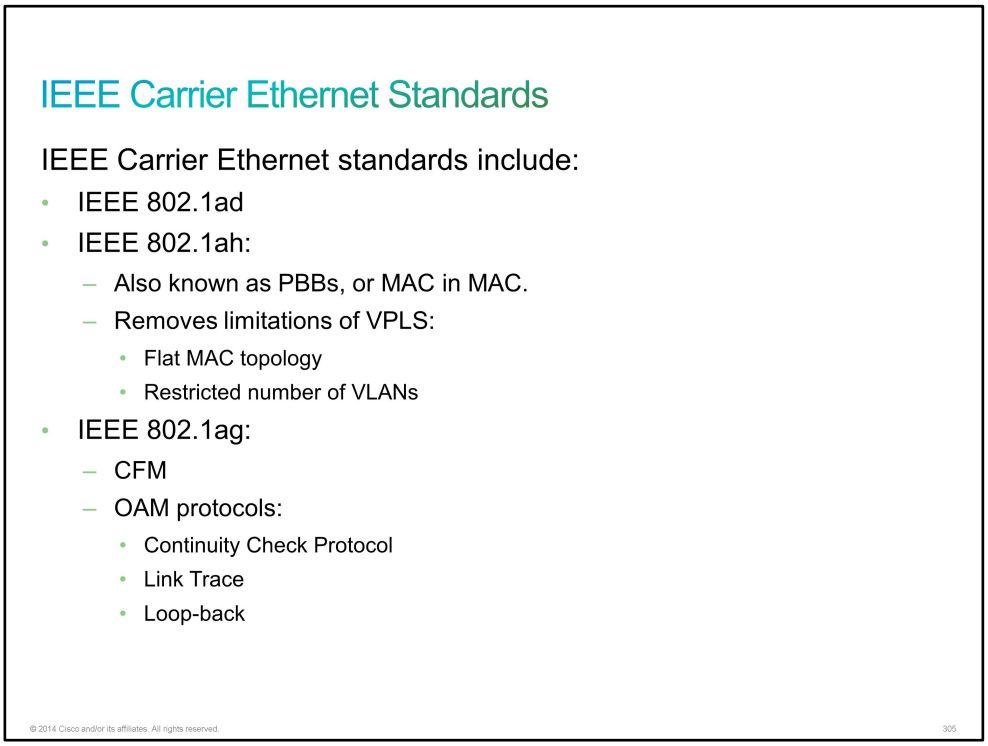
The failure of an NPE router or the connection of the PW can cause the UPE to suffer a total loss of connectivity. To prevent this, redundant connections can be provided. The UPE is dual-homed into two NPE routers. One of the two PWs is designated as primary and the other as standby. The UPEs negotiate PW labels for both PWs, but do not use the standby PW unless the primary PW fails. Spanning-tree instance or manual configuration can be used to designate primary and standby status.

Upon failure of the primary PW, the UPE immediately switches to the standby PW. At this point, the NPE that terminates the standby PW starts learning MAC addresses on the spoke PW. All other NPEs initially continue to send traffic to the initial NPE until they learn that the devices are now connected to a new NPE. To enable a faster unlearning process, the new NPE may send out a flush message using the MAC list TLV (type 0x404) to all NPEs. Upon receiving the message, the NPEs flush the MAC addresses associated with that VPLS instance.

https://t.me/learningnets

IEEE Carrier Ethernet Standards

Many standardization bodies have issued standards that define basic functions and enhancements of the Carrier Ethernet architecture.



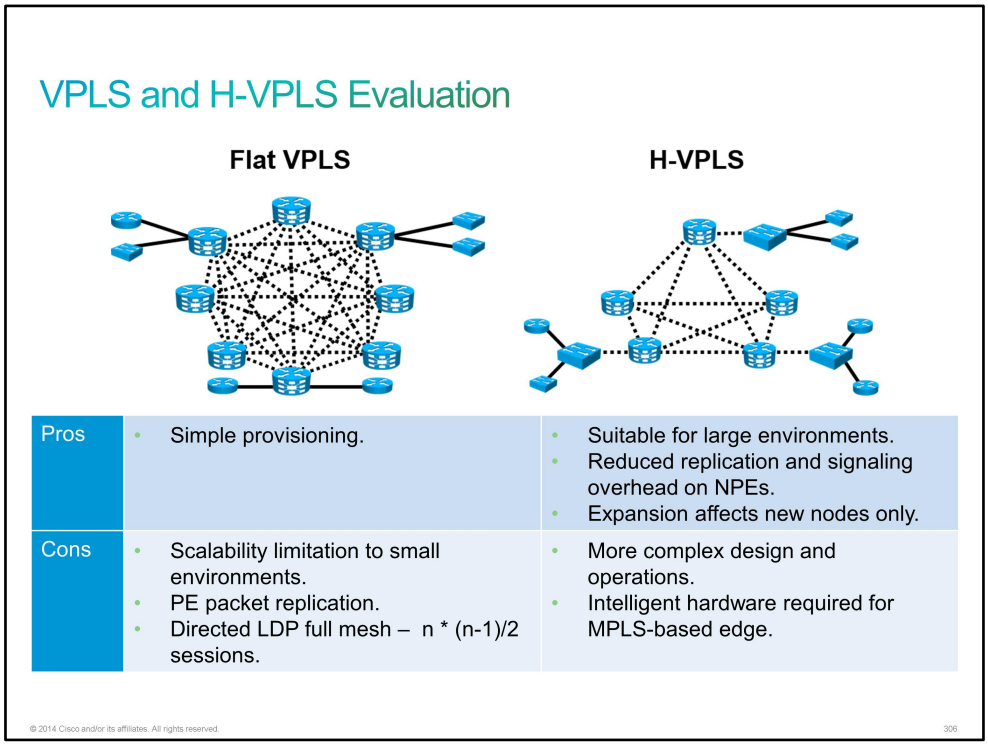
IEEE plays a major role in this field and contributed many important standards, such as the following:

- IEEE 802.1ad, which is the formal name of the 802.1QinQ standard that allows Ethernet frame encapsulation in multiple VLAN tags.
- IEEE 802.1ah, also known as PBBs, or MAC in MAC. This technology addresses the limitations of VPLS, such as flat MAC topology or limited number of VLANs. It scales to large service provider environments.
- IEEE 802.1ag, also referred to as CFM. This set of recommendations defines protocols and practices for Ethernet OAM. Specifically, it comprises three protocols: Continuity Check Protocol, Link Trace, and Loop-back.

https://t.me/learningnets

VPLS and H-VPLS Evaluation

Flat VPLS offers the easiest implementation method. It supports both port and VLAN transport modes, but is limited to small deployments. All packet replication is performed on the PE devices. It results in a directed LDP full mesh of n*(n-1)/2 sessions.

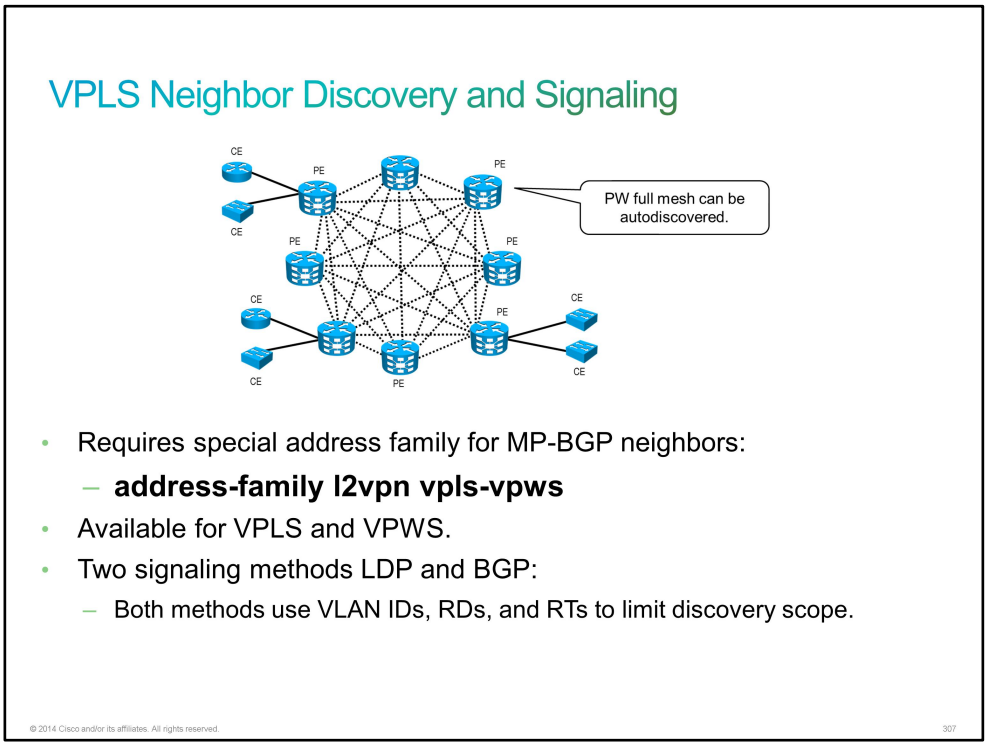


H-VPLS is suitable for larger environments. It places less replication burden and signaling overhead on the NPE routers. It allows the service provider to expand certain network areas without affecting the core and most of the NPE infrastructure. It requires, however, more complex design, provisioning, and operations, and may incur higher cost, if MPLS is deployed in the edge.

https://t.me/learningnets

VPLS Neighbor Discovery and Signaling

The PEs need to identify other members of the same VPLS instance. Once the peers have been identified, each PE needs to set up a PW between themselves and bind the PWs to the particular VSI.



These processes can be automated to remove the need for static configuration of VPLS members. The BGP autodiscovery can be deployed to automatically identify members of the same VPLS.

When BGP autodiscovery is used to find the VPLS peers, you can combine it with one of two PW signaling methods:

- BGP-based signaling
- LDP-based signaling

LDP-Based Signaling

LDP is used for the signaling of PWs that are used to interconnect the VPLS instance of a given customer on the PEs. In order to signal a full mesh of PWs, a full mesh of Targeted LDP (T-LDP) sessions is required between the PEs. In the absence of autodiscovery, these sessions must be manually configured on each PE router. This LDP session is used to communicate the inner label or VC label that must be used for each PW. The network operator assigns a VC ID that is used to identify a particular PW, and it is configured to be the same for a particular VPLS instance on all PE routers.

The LDP session communicates the label value and the FEC element that includes the VC ID, control word bit (indication whether a control word will be used), VC type (Ethernet or VLAN-tagged Ethernet), and interface parameters field (media MTU, and so on.)

https://t.me/learningnets

LDP signals the setup, maintenance, and teardown of a PW between PE devices. Once a PE discovers that other PEs have an association for a particular VPLS instance, the PEs signal using T-LDP to other PEs that a PW is required to be set up between PEs. When a PE device is associated with a particular VSI, LDP transmits a label-mapping message with VC type 0x0005 and a 4-byte VC ID value. If the remote PE has an association with that particular VC ID, it will accept the LDP label-mapping message and respond with its own label-mapping message. Once the two unidirectional VCs are operational, they are combined to form a bidirectional PW.

BGP-Based Autodiscovery and Signaling

The BGP NLRI takes care of autodiscovery and signaling at the same time. The NLRI is generated by a given PE containing the necessary information required by any other PE. These components enable the automatic setting up of a full mesh of PWs for each VPLS.

On each PE, a route distinguisher (RD) and a route target (RT) are configured for each VPLS, like in Layer 3 VPN and Layer 2 VPN. The RT is same for a particular VPLS across all PEs, and is used to identify which VPLS a particular BGP message pertains to. The RD is used to disambiguate routes. On each PE, for each VPLS, an identifier is configured called a VPLS Edge Identifier (VE ID). Each PE involved in a particular VPLS must be configured with a different VE ID. BGP is used to advertise the VE ID to other PEs. This, along with other information in NLRI, is used to calculate the value of the PW label required to reach the advertising PE.

For a given VPLS, a PE requires that each remote PE uses a different PW label to send traffic to that PE. This requirement is done to facilitate the MAC learning process. This way, the receiving PE can learn which PE is associated with the source MAC address of the frame. A PE could send a list of PW labels that are required to reach it, one per remote VPLS edge in that VPLS. However, this would mean that a PE sends a long list of labels if there is a large number of PEs. Instead, the necessary information is carried in BGP NLRI. This allows all remote PEs to calculate the PW label expected by the advertising PE.

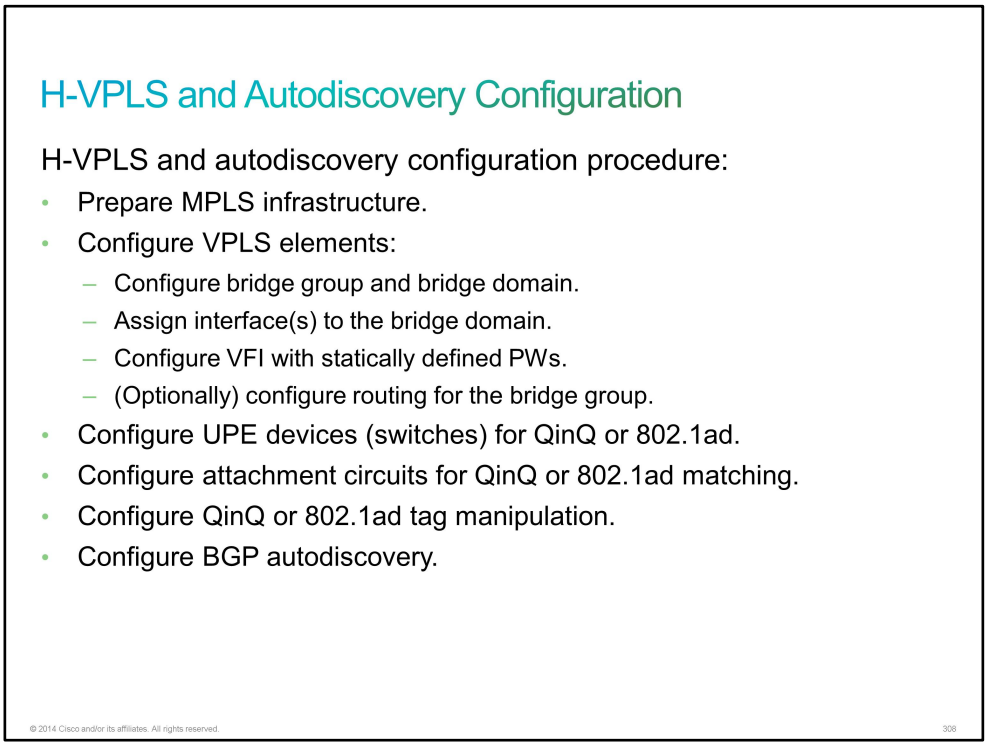
A BGP update message contains the extended community RT (which allows the receiving PE to ascertain which particular VPN the advertisement pertains to), Layer 2 information (control flags to indicate whether a control word is included, the encapsulation type—Ethernet or VLAN-tagged—and MTU), and other BGP attributes such as AS path, origin, and so on.

The NLRI contains the RD, VE ID, label base, VE block offset, and VE block size. The label base, VE block offset and VE block size are the information that the receiving PE requires to calculate the PW label when sending traffic for VPLS to the advertising PE.

https://t.me/learningnets

H-VPLS and Autodiscovery Configuration

Follow this configuration procedure when deploying hierarchical VPLS with QinQ or 802.1ad in the edge and BGP autodiscovery of VPLS peers.



To implement VPLS, you first need to prepare the MPLS infrastructure. This process includes these steps:

- PE routers must have a /32 address on their Loopback interfaces. This is required for binding the VC label and the transport LDP into a label stack.
- PE loopback addresses cannot be summarized in the core. Aggregation in the network would break the LSP path.
- MPLS must be enabled in the core.
- MTU sizes on the core links must be able to accommodate the customer MTU with the added label overhead.

This scenario assumes that a flat VPLS deployment is already in place. This includes the configuration of Layer 2 frame transport or Ethernet service instance. The MTU sizes are the same on all endpoint interfaces. You have configured a bridge group and a bridge domain, and assigned interfaces to bridge domain. Optionally you have also configured routing for the bridge domain.

To enhance the VPLS scalability, you will:

- Configure UPE devices (switches) for QinQ or 802.1ad.
- Configure attachment circuits for QinQ or 802.1ad matching.
- Configure the bridge group and bridge domain.
- Configure QinQ or 802.1ad tag manipulation.

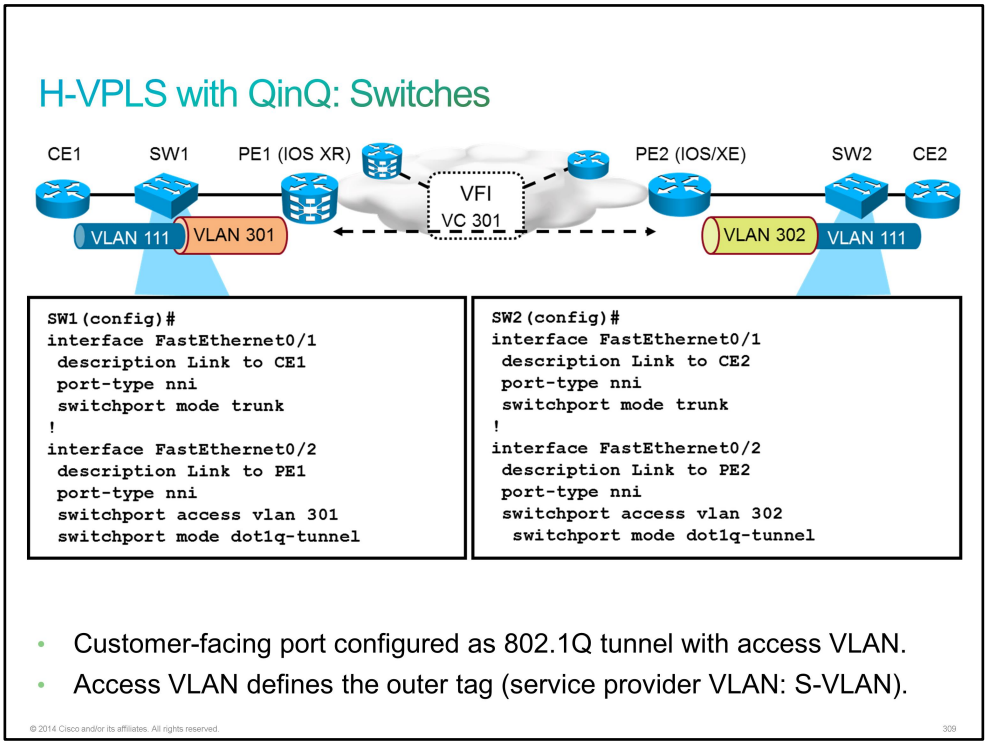
https://t.me/learningnets

- Configure BGP autodiscovery.

https://t.me/learningnets

H-VPLS with QinQ: Switches

This example shows the edge switch configuration in an H-VPLS environment with QinQ.



The customer-facing interfaces are configured as dot1q-tunnel interfaces with an access VLAN.

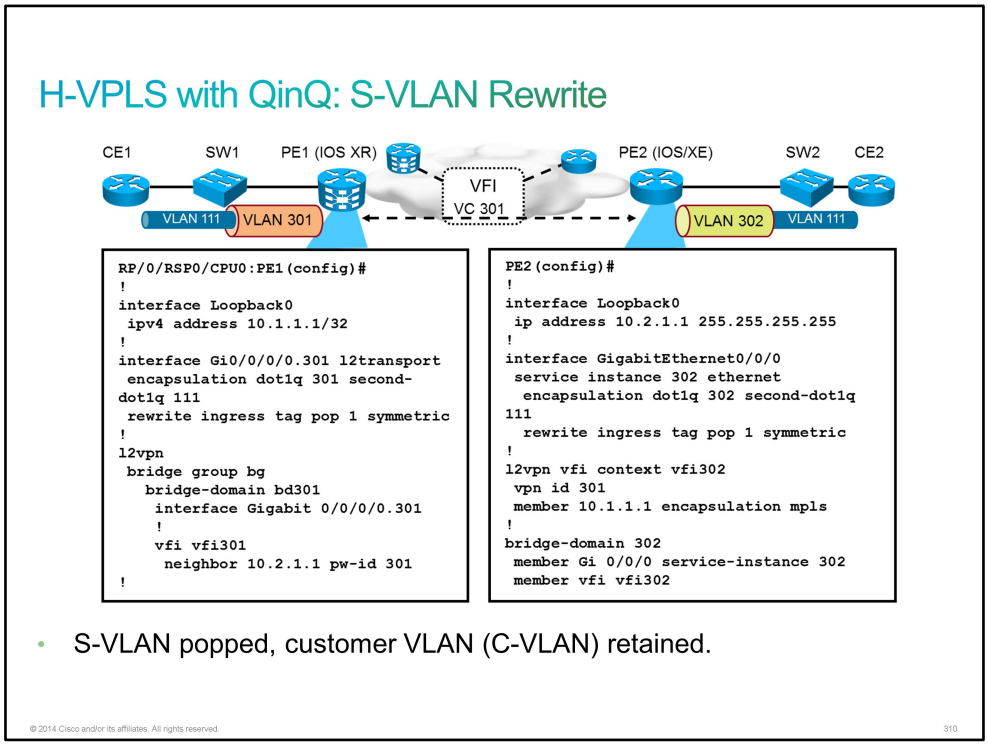
SW1 encapsulates the 802.1Q frames received from CE1 within the provider tag 301. SW2 encapsulates the 802.1Q frames received from CE2 within an provider tag 302. The provider tag appears as the outer tag in a double-tagged Q-in-Q frame.

Although the PE-facing interfaces of the user-facing PE (UPE) switches are configures as normal 802.1Q trunks, they will exchange double-tagged frames with the PE routers.

https://t.me/learningnets

H-VPLS with QinQ: S-VLAN Rewrite

In this scenario the customer VLAN 111 on CE1 is encapsulated within provider VLAN 301 on SW1 and bridged on PE1 and PE2 with the customer VLAN 111 on CE2 encapsulated within provider VLAN 302 on SW2.



The PE router configuration in an H-VPLS scenario resembles basic VPLS deployments. The difference lies with the encapsulation and rewrite settings on the attachment circuits.

The attachment circuit on the Cisco IOS XR PE is enabled for Layer 2 transport using the **l2transport** keyword. The interface GigabitEthernet0/0/0/0.301 terminates double-tagged frames with the outer 802.1Q tag 301 and the inner 802.1Q tag 111. The AC pops the outer tag in a symmetric fashion. The bridge domain will thus transport frames with a single 802.1Q tag (111).

In Cisco IOS XE you configure the attachment circuit using a service instance configured under the main interface. The Ethernet service instance 302 terminates double-tagged frames with the outer 802.1Q tag 302 and the inner 802.1Q tag 111. The AC pops the outer tag in a symmetric fashion. This configuration matches the settings on the PE router running Cisco IOS XR software.

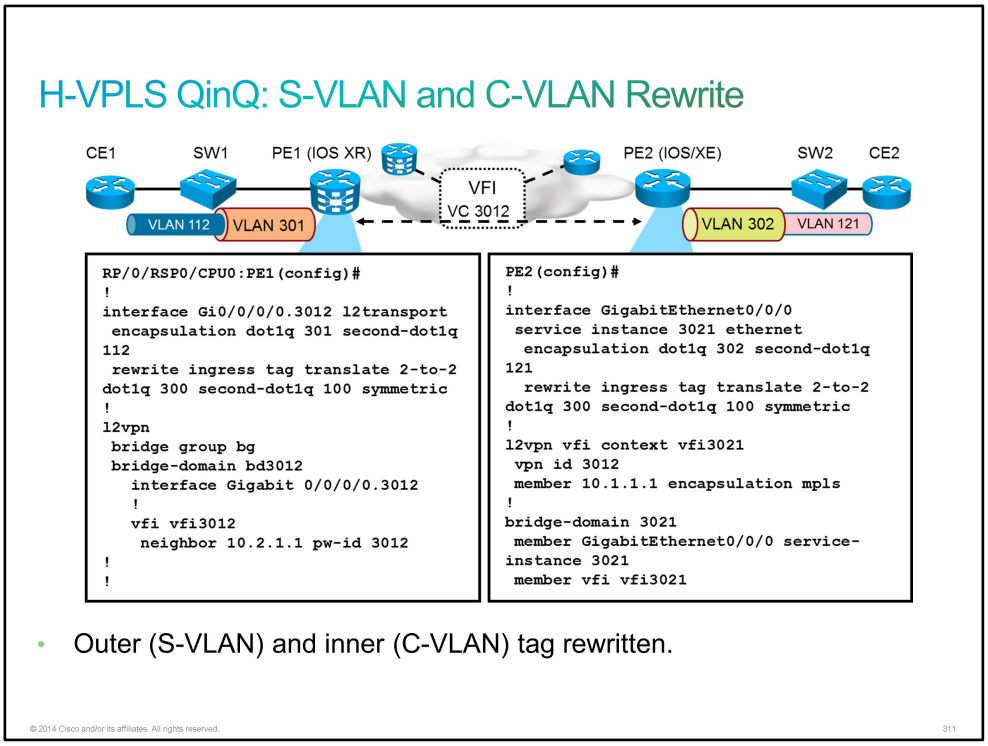
The remaining Cisco IOS XR VPLS configuration is performed within the Layer 2 VPN (**l2vpn** command), bridge group, and bridge domain configuration mode. The bridge domain contains two main configuration elements: a list of the local attachment circuits, and the VFI that contains the PWs.

The remaining Cisco IOS XE VPLS configuration consists of two constructs: the VFI and the bridge domain. The VFI is the set of pseudowires configured with the member command. The VPN ID configured in the VFI mode must match up with the VC ID configured in Cisco IOS XR with the **neighbor** command. The bridge domain links the attachment circuits with the VFI. The attachment circuit (service instance 301 on GigabitEthernet0/0/0) and the VFI (vfi301) are configured as members of the bridge domain (301).

https://t.me/learningnets

H-VPLS QinQ: S-VLAN and C-VLAN Rewrite

In this scenario the customer VLAN 112 on CE1 is encapsulated within provider VLAN 301 on SW1 and bridged on PE1 and PE2 with the customer VLAN 121 on CE2 encapsulated within provider VLAN 302 on SW2.



The attachment circuit on the Cisco IOS XR PE is enabled for Layer 2 transport using the l2transport keyword. The interface GigabitEthernet0/0/0/0.3012 terminates double-tagged frames with the outer 802.1Q tag 301 and the inner 802.1Q tag 112. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. The bridge domain will thus transport double-tagged frames with outer 802.1Q tag 300 and inner 802.1Q tag 100.

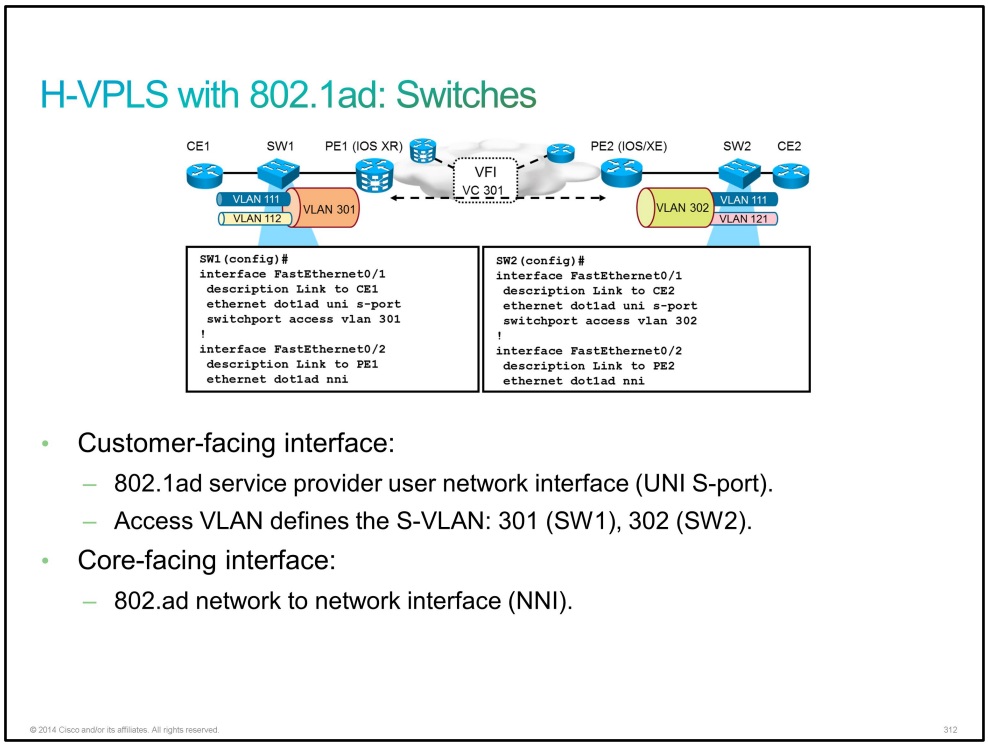
In Cisco IOS XE you configure the attachment circuit using a service instance configured under the main interface. The Ethernet service instance 3021 terminates double-tagged frames with the outer 802.1Q tag 302 and the inner 802.1Q tag 121. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. This configuration matches the settings on the PE router running Cisco IOS XR software.

The remaining configuration is typical for most VPLS deployments.

https://t.me/learningnets

H-VPLS with 802.1ad: Switches

This example shows the edge switch configuration in an H-VPLS environment with 802.1ad.



The customer-facing interfaces are configured as 802.1ad user network interfaces (UNI) of S-port type. They have an access VLAN configured on them that will be used as the outer tag.

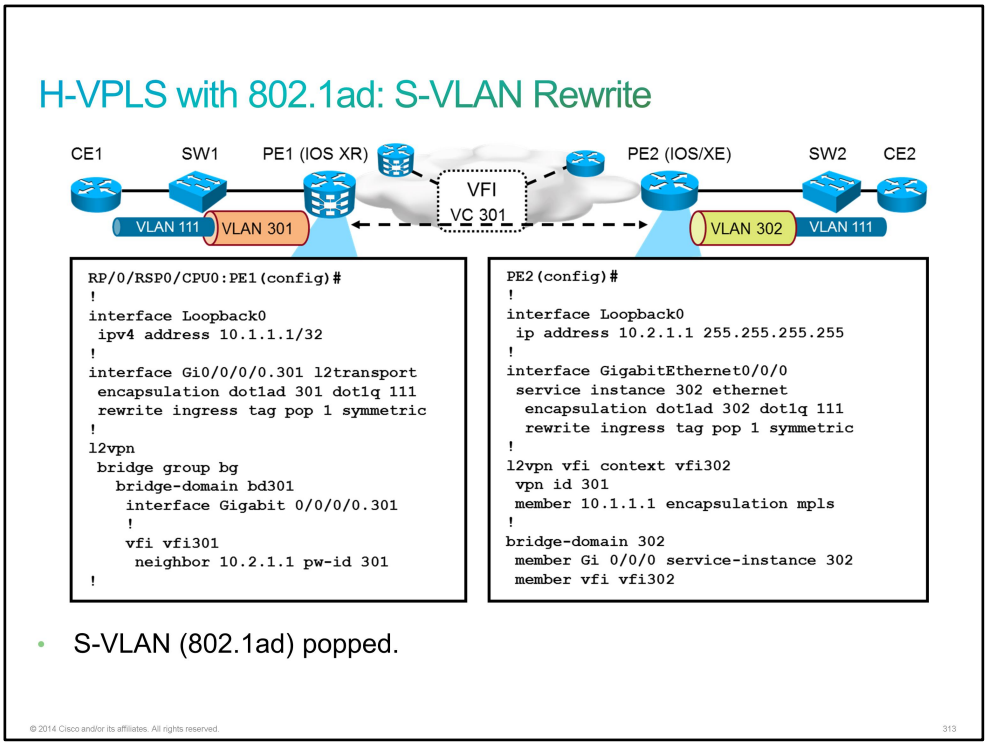
The PE-facing interfaces of the user-facing PE (UPE) switches are configures as 802.1ad network-to-network interfaces (NNI) interfaces. They will exchange double-tagged frames with the PE routers.

SW1 encapsulates the 802.1Q frames received from CE1 within the provider tag 301. SW2 encapsulates the 802.1Q frames received from CE2 within an provider tag 302. The provider tag appears as the outer tag in a double-tagged 802.1ad frame.

https://t.me/learningnets

H-VPLS with 802.1ad: S-VLAN Rewrite

In this scenario the customer VLAN 111 on CE1 is encapsulated within provider VLAN 301 on SW1 and bridged on PE1 and PE2 with the customer VLAN 111 on CE2 encapsulated within provider VLAN 302 on SW2. 802.1ad is used the H-VPLS edge technology.



The attachment circuit on the Cisco IOS XR PE is enabled for Layer 2 transport using the l2transport keyword. The interface GigabitEthernet0/0/0/0.301 terminates double-tagged frames with the outer 802.1ad tag 301 and the inner 802.1Q tag 111. The AC pops the outer tag in a symmetric fashion. The bridge domain will thus transport frames with a single 802.1Q tag (111).

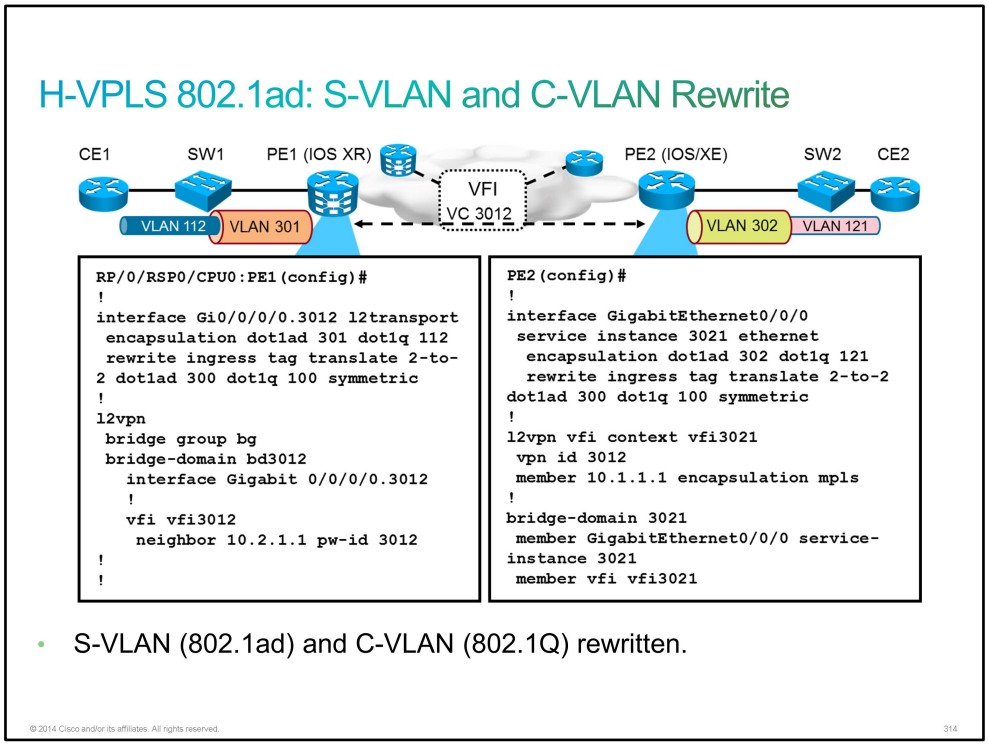
In Cisco IOS XE you configure the attachment circuit using a service instance configured under the main interface. The Ethernet service instance 302 terminates double-tagged frames with the outer 802.1ad tag 302 and the inner 802.1Q tag 111. The AC pops the outer tag in a symmetric fashion. This configuration matches the settings on the PE router running Cisco IOS XR software.

The remaining configuration is typical for most VPLS solutions.

https://t.me/learningnets

H-VPLS 802.1ad: S-VLAN and C-VLAN Rewrite

In this scenario the customer VLAN 112 on CE1 is encapsulated within provider VLAN 301 on SW1 and bridged on PE1 and PE2 with the customer VLAN 121 on CE2 encapsulated within provider VLAN 302 on SW2.



The attachment circuit on the Cisco IOS XR PE is enabled for Layer 2 transport using the l2transport keyword. The interface GigabitEthernet0/0/0/0.3012 terminates double-tagged frames with the outer 802.1ad tag 301 and the inner 802.1ad tag 112. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. The bridge domain will thus transport double-tagged frames with outer 802.1ad tag 300 and inner 802.1Q tag 100.

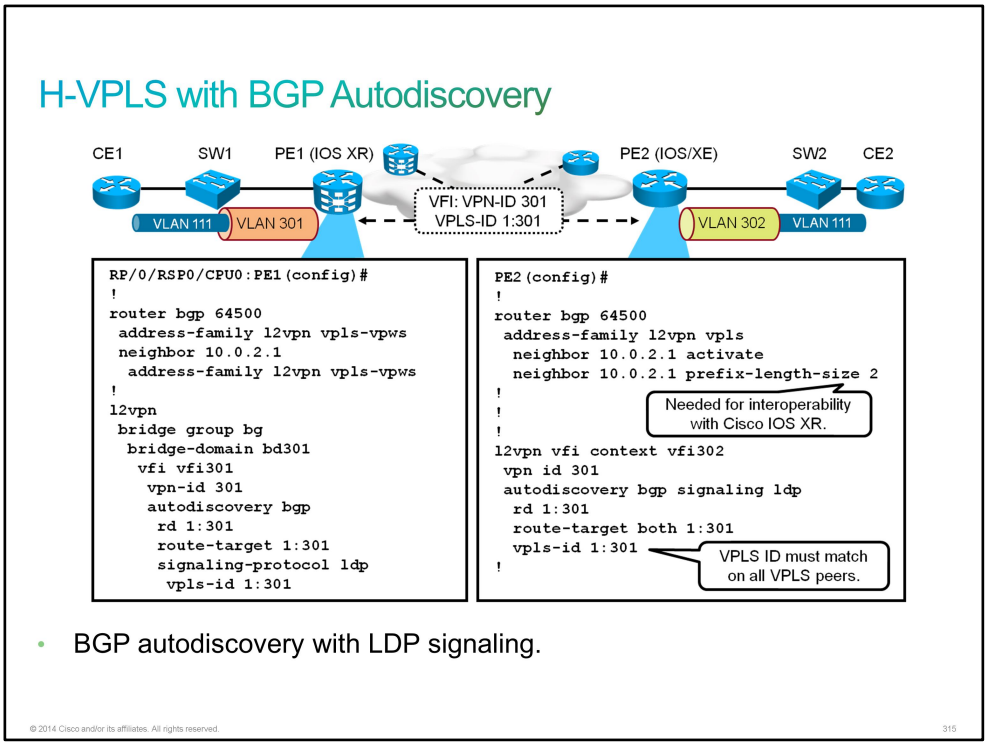
In Cisco IOS XE you configure the attachment circuit using a service instance configured under the main interface. The Ethernet service instance 3021 terminates double-tagged frames with the outer 802.1ad tag 302 and the inner 802.1Q tag 121. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. This configuration matches the settings on the PE router running Cisco IOS XR software.

The remaining configuration is typical for most VPLS deployments.

https://t.me/learningnets

H-VPLS with BGP Autodiscovery

When you add a new VPLS PE to the network, all existing PEs have to be reconfigured in order to have a PW to the new PE because all PEs must be fully meshed. This might become an operational challenge as the number of PEs and bridge-domains increase.



To solve this problem the PEs can discover other PEs automatically through BGP. While there is also a full-mesh requirement for IBGP, it can be lifted by the use of route-reflectors, as in this example. A new PE is typically configured in order to peer with a small number of route-reflectors, all other PEs receive its updates, and the new PE receives the updates from the other PEs.

In order to discover other PEs through BGP, each PE is configured for the vpls-vpws (Cisco IOS XR) and vpls (Cisco IOS XE) address-family and advertises in BGP the bridge-domains in which they want to participate. Once the other PEs that are part of the same bridge-domain are discovered, a PW is established to each of them. BGP is the protocol used for this autodiscovery.

After activating the address family vpls-vpws (Cisco IOS XR) and vpls (Cisco IOS XE) for the session with the route reflector, you need to set the prefix-length-size to 2 to avoid an NLRI mismatch between Cisco IOS XR and Cisco IOS XE. The Cisco IOS XE Software encodes the NLRI length in the first byte in bits format in the BGP Update message. However, the Cisco IOS XR Software interprets the NLRI length in 2 bytes. Therefore, when the BGP neighbor with VPLS-VPWS address family is configured between the IOS and the IOS XR, an NLRI mismatch will occur happen, leading to flapping between neighbors when updates are exchanged.

There are two options for the signaling of the PW to the auto-discovered PEs: BGP and LDP. In this example, you enhance the existing VPLS that bridges the provider VLANs 301 and 302. Instead of a static pseudowire configuration you deploy BGP autodiscovery with LDP signaling.

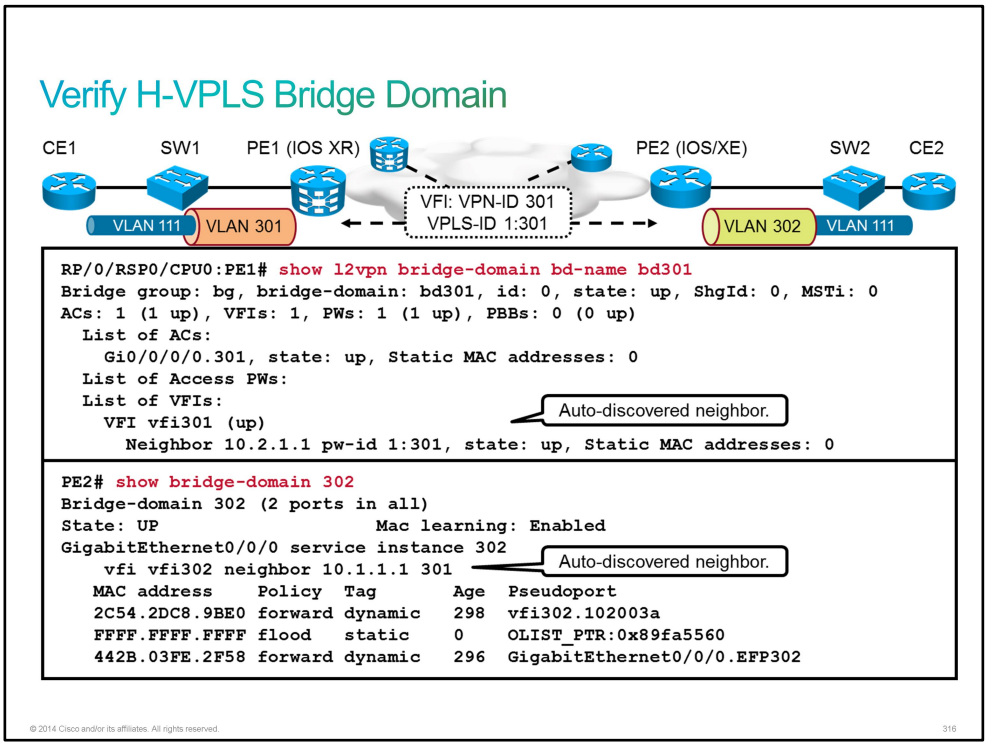
https://t.me/learningnets

You need to configure the VFI for autodiscovery instead of static neighbors. A VPN ID (in this case 301) is mandatory. In the autodiscovery BGP (signaling LDP) configuration mode, you set the route distinguisher, route target, and a VPLS ID. Depending on the platform, the RD may have to be the same on all VPLS peers. The RTs have to match for the updates to be exchanged within a bridge domain. The VPLS IDs need to match up within a VPLS bridge domain. The VPLS ID is typically construed of the BGP AS number and the VPN ID.

https://t.me/learningnets

Verify H-VPLS Bridge Domain

You can verify the bridge domain using the **show l2vpn bridge-domain** (Cisco IOS XR) and **show bridge-domain** (Cisco IOS XE) commands.



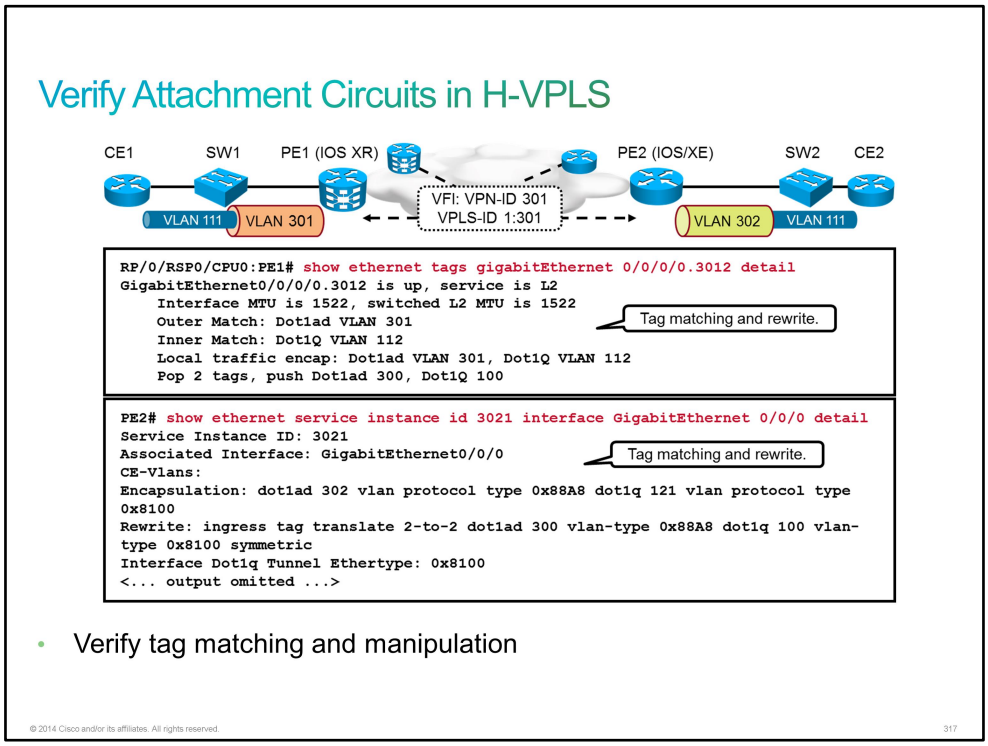
The output will display the state of the bridge domain, its attachment circuits and VFIs. The pseudowires can be configured statically or discovered via BGP. In this case, BGP autodiscovery was used.

The command output in Cisco IOS XE displays the MAC addresses learned within the bridge domain.

https://t.me/learningnets

Verify Attachment Circuits in H-VPLS

To verify the edge technology operation of the H-VPLS solution, you should examine the attachment circuits on the VPLS peers.



To monitor the tag matching and manipulation at the attachment circuits, use the **show ethernet tags** (Cisco IOS XR) and **show ethernet service instance** (Cisco IOS XE) commands. The detailed output displays the outer and inner tag matching, the encapsulation protocol, and the tag manipulation.

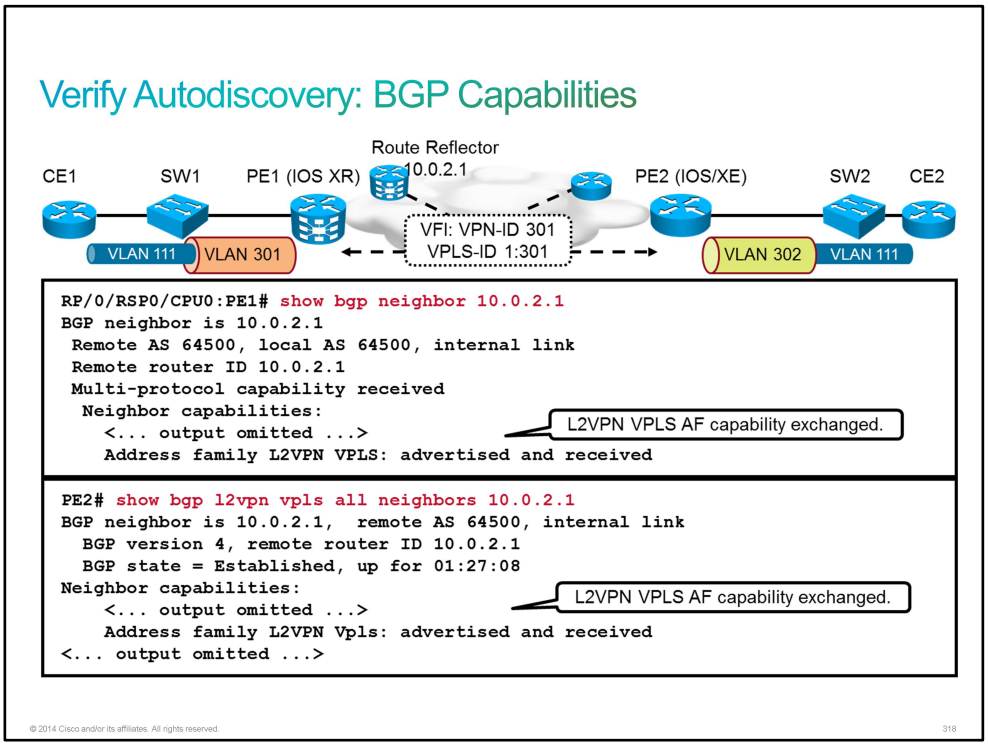
The attachment circuit on the Cisco IOS XR PE has been enabled for Layer 2 transport. The interface GigabitEthernet0/0/0/0.3012 terminates double-tagged frames with the outer 802.1ad tag 301 and the inner 802.1ad tag 112. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. The bridge domain will thus transport double-tagged frames with outer 802.1ad tag 300 and inner 802.1Q tag 100.

In Cisco IOS XE the attachment circuit has been configured using a service instance configured under the main interface. The Ethernet service instance 3021 terminates double-tagged frames with the outer 802.1ad tag 302 and the inner 802.1Q tag 121. The AC translates the outer tag to 300 and inner tag to 100 in a symmetric fashion. This operation matches the tag handling on the PE router running Cisco IOS XR software.

https://t.me/learningnets

Verify Autodiscovery: BGP Capabilities

To verify BGP autodiscovery operation you need to make sure that the Layer 2 VPLS have been successfully negotiated between the BGP peers.



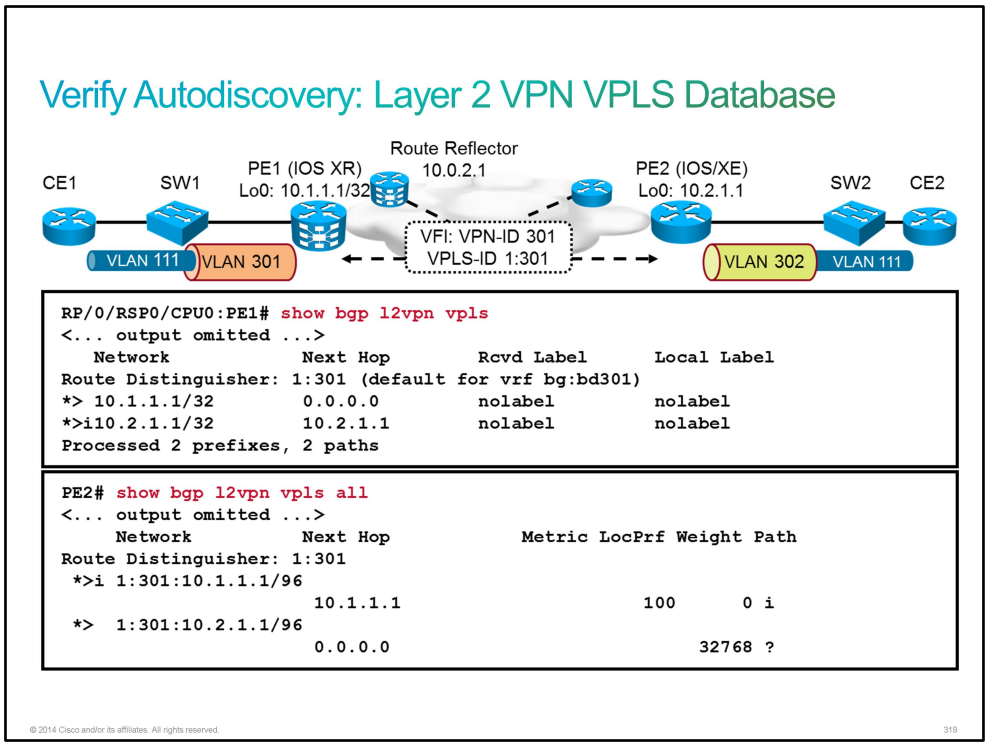
In this case, PE1 and PE2 have an IBGP session with a route reflector (10.0.2.1). The route reflector has also been activated for the address family Layer 2 VPLS exchange.

You see that both VPLS peers, PE1 and PE2, have both advertised and received the address family Layer 2 VPLS to and from the route reflector.

https://t.me/learningnets

Verify Autodiscovery: Layer 2 VPN VPLS Database

To display Layer 2 VPN address family information from the BGP table, use the **show bgp l2vpn vpls** command.



The database contains local generated information and updates received from other VPLS peers within the same bridge domain.

In this scenario you will not see any labels exchanged for the attachment circuits because LDP is used to signal the labels. The task of BGP is to auto-discover the VPLS peers and exchange the Layer 2 VPN attributes.

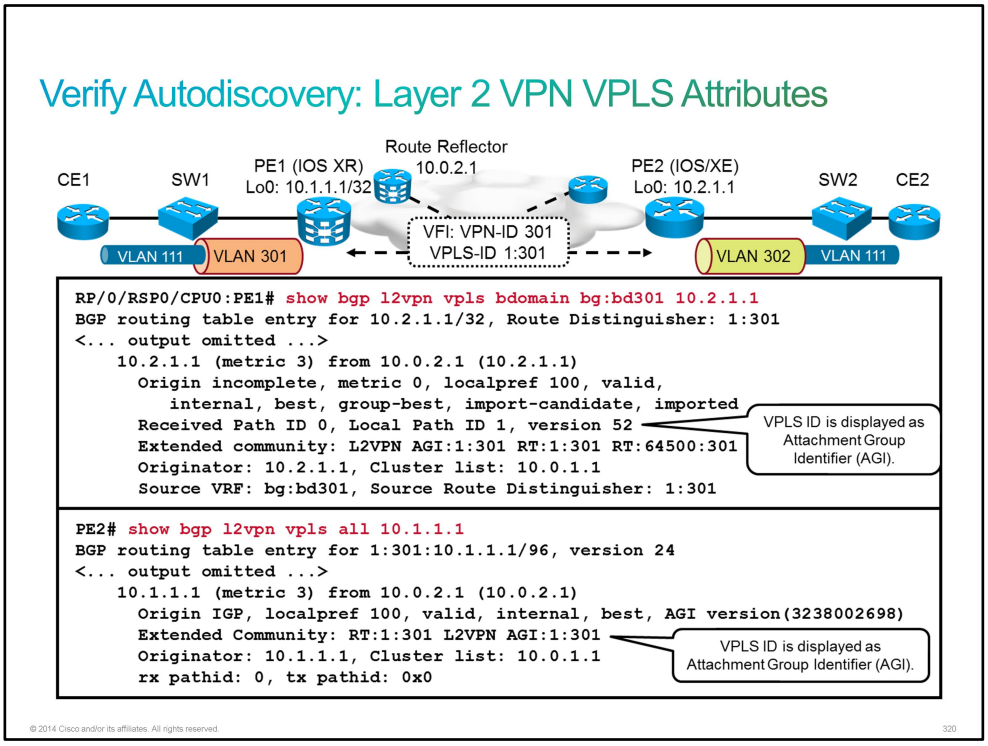
More specifically the responsibilities of BGP can be broken into these tasks:

- Advertise LRI and RTs.
- Learn the RD or RT configured under a VPLS bridge domain.
- Import NLRIs based on RT(s) and pass {VFI_ID, next-hop, layer2info} to the Layer 2 VPN function.
- Provide an API for Layer 2 VPN to retrieve the AS number.

https://t.me/learningnets

Verify Autodiscovery: Layer 2 VPN VPLS Attributes

When you examine the individual entries in the BGP Layer 2 VPN VPLS database you will realize that they do not resemble IP prefixes that you would expect in a traditional BGP table.



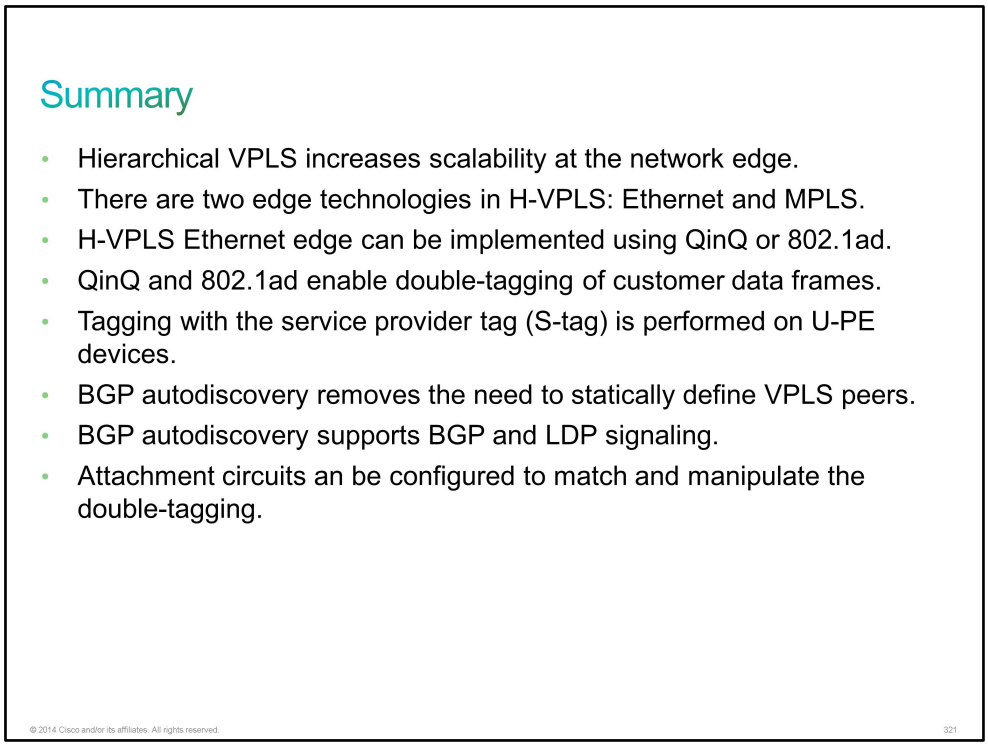
The entries in the BGP Layer 2 VPN VPLS database should be rather thought of as containers for Layer 2 VPN attributes advertised by a given VPLS peer.

These placeholders contain descriptors of bridge domains, such as RTs and Attachment Group Identifiers (AGIs), which are carried within BGP as extended community attributes. The AGIs have been configured within a VFI as VPLS IDs.

https://t.me/learningnets

Summary

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



https://t.me/learningnets

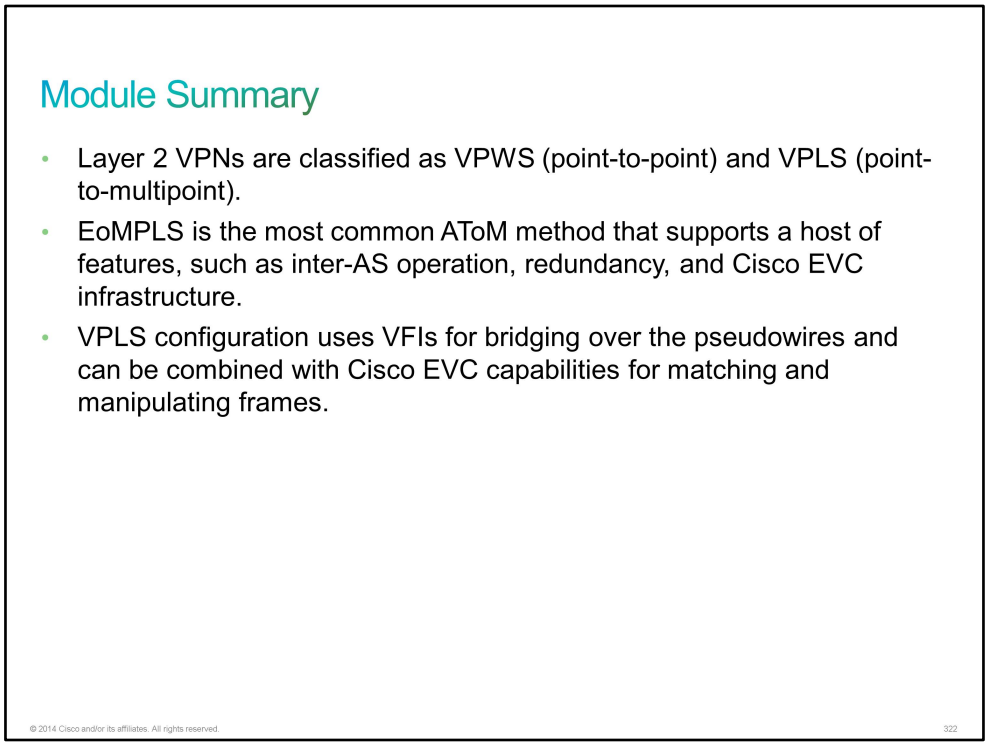
 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 5: Module Summary

Overview

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



Layer 2 VPNs are classified as VPWS (point-to-point) and VPLS (point-to-multipoint). They can be transported over MPLS or IP networks. VPLS can only be carried over MPLS and cannot use L2TPv3, which provides the transport protocol in the pure IP core.

AToM is a subset of VPWS. It is deployed to carry Ethernet, Frame Relay, PPP and HDLC frames, and ATM cells, and it emulates TDM circuits over MPLS. It supports interworking of disparate encapsulations at both ends and allows the connections to span multiple autonomous systems.

https://t.me/learningnets

AToM requires that MPLS and T-LDP sessions are operational, which includes that the loopback addresses on PE routers use a /32 network mask and are not summarized in the core. MTU considerations include the requirement of having identical MTU sizes on both sides and guaranteeing sufficiently large MTU sizes in the core to avoid fragmentation. Fragmentation is not supported for Layer 2 VPNs.

Carrier Ethernet offers a range of enhancements for scalability (802.1ad and 802.1ah), high availability (MST, REP, and G.8032), and OAM (802.1ag and link OAM).

VPLS configuration uses VFIs for bridging over the PWs and can be combined with Cisco EVC capabilities for matching and manipulating frames.

References

For additional information, refer to these resources:

- To learn more about implementing Layer 2 VPN on Cisco IOS XR platforms, refer to *Implementing Virtual Private LAN Services on Cisco IOS XR Software* at this URL: http://www.cisco.com/c/en/us/td/docs/ios_xr_sw/iosxr_r3-7/mpls/configuration/guide/gc37vpls.html.
- To learn more about configuring Layer 2 VPNs and Ethernet services on Cisco IOS XR platforms, refer to *Cisco ASR 9000 Series Aggregation Services Router L2VPN and Ethernet Services Configuration Guide* at this URL: http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r4-1/lxvpn/configuration/guide/lesc41/lesc41ethftrs.html.

 © 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. To establish VPWS in your service provider environment, MPLS must be configured in your IP core network. (Source: Introducing Layer 2 VPNs)

   A. true
   B. false

2. To establish VPLS in your service provider environment, MPLS must be configured in your IP core network. (Source: Introducing Layer 2 VPNs)

   A. true
   B. false

3. Which statement best describes AToM? (Source: Introducing Layer 2 VPNs)

   A. It can be multipoint if multiple attachment circuits are used.
   B. It requires attachment circuit(s).
   C. It requires route targets.
   D. It requires MP-BGP.
   E. It requires route distinguisher(s).

4. Which statement correctly describes pseudowires. (Source: Implementing VPLS)

   A. Pseudowires are used in VPWS only.
   B. VPWS can include one or more pseudowires.
   C. VFI can include one or more pseudowires.
   D. VFI uses point-to-multipoint pseudowires.

https://t.me/learningnets

5. Which statement correctly describes the bridge-domain virtual interface. (Source: Implementing VPLS)

   A. It can be configured in a VPLS only.
   B. It can be configured in a VPWS only.
   C. It can be configured in a VPLS or VPWS.
   D. It is associated with a pseudowire.
   E. It is associated with a VFI.

6. Which statement correctly describes the handling of BPDUs in Layer 2 VPNs? (Source: Implementing VPLS)

   A. VPWS does not transport BPDUs because it is not a virtual bridge.
   B. VPLS does not use BPDUs because it does not provide point-to-point connections.
   C. You will have to filter BPDUs when connecting different VLANs with VPLS.
   D. You will have to filter BPDUs when connecting different VLANs with VPWS.

7. What is required for BGP autodiscovery with LDP signaling? (Source: Enhancing L2VPN scalability)

   A. A route target for the desired address family (IPv4/IPv6)
   B. A pseudowire template that is used for spawning the pseudowires
   C. LDP parameters that are configured on attachment circuits
   D. An L2VPN VPLS address-family

8. What do you need to change on PE routers to migrate from QinQ to 802.1ad in the edge? (Source: Enhancing L2VPN scalability)

   A. Pseudowire configuration
   B. Encapsulation definition on the attachment circuits
   C. Rewrite definition (situational) on the attachment circuits
   D. Encapsulation and rewrite (situational) definition on the attachment circuits
   E. VFI configuration
   F. Bridge domain configuration

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. B
2. A
3. B
4. C
5. A
6. D
7. D
8. D

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Glossary

6PE

Cisco IPv6 Provider Edge Router over MPLS.

6VPE

VPN IPv6 provider edge router

AAL5

ATM adaptation layer 5.

ABR

area border router. Router located on the border of one or more OSPF areas that connects those areas to the backbone network. ABRs are considered members of both the OSPF backbone and the attached areas. They therefore maintain routing tables describing both the backbone topology and the topology of the other areas.

AC

alternating current.

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).

AFT

Address Family Translation.

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.

ASBR

autonomous system boundary router. An ABR located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area.

ASCII

American Standard Code for Information Interchange. 8-bit code for character representation (7 bits plus parity).

ASN

autonomous system number

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.

https://t.me/learningnets

AToM

Any Transport over MPLS

BDI

bridge-domain virtual interface

BGP

Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163.

BGP4

Border Gateway Protocol version 4

BPDU

bridge protocol data unit. Spanning Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network.

BVI

bridge-group virtual interface.

CE

customer edge. Identifies the network devices, connected to a provider network, that are under the administrative control of the customer.

CE bit

congestion experienced bit. Bit in the ECN portion of the DS field that is used with the ECT bit to indicate that endpoints are ECN-capable or that congestion has been experienced.

CE-PE

customer edge-provider edge

CFM

Connectivity Fault Management.

CLI

command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. The UNIX operating system and DOS provide CLIs.

C-network

customer network

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.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

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.

CSC

Carrier Supporting Carrier

DHCP

Dynamic Host Configuration Protocol. Provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them.

DLCI

data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices).

DMVPN

Dynamic Multipoint VPN.

DoS

denial of service. An intentional or unintentional attack on a device that makes the resource unavailable to perform its normal function.

EBGP

Exterior Border Gateway Protocol.

EFP

Ethernet Flow Point

EGP

Exterior Gateway Protocol. Internet protocol for exchanging routing information between autonomous systems. Documented in RFC 904. Not to be confused with the general term *exterior gateway protocol*. EGP is an obsolete protocol that was replaced by BGP.

EIA

Ethernet Internet Access

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.

EoMPLS

Ethernet over MPLS.

https://t.me/learningnets

EPL

Ethernet Private Line

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 Mbps. Ethernet is similar to the IEEE 802.3 series of standards. It is 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.

EVC

Ethernet Virtual Circuit

EVPL

Ethernet Virtual Private Line.

EXP

experimental.

FIB

forwarding information base

Frame Relay

Industry-standard, packet-switched data link layer protocol that handles multiple virtual circuits between connected devices.

FRoMPLS

Frame Replay over MPLS.

GDOI

Group Domain of Interpretation

GET

Group Encrypted Transport.

Gigabit Ethernet

Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.

GRE

Generic Routing Encapsulation. Tunneling protocol that was developed by Cisco and that can encapsulate a variety of protocol packet types inside IP tunnels. This process creates a virtual point-to-point link to Cisco routers at remote points over an IP network.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

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.

H-VPLS

hierarchical VPLS

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.

ICMPv4

Internet Control Message Protocol version 4.

ICMPv6

Internet Control Message Protocol version 6 is the implementation of Internet Control Message Protocol (ICMP) for IPv6. ICMPv5 is defined in RFC 4443.

ID

identifier

IEEE 802.1Q

The IEEE 802.1Q specification establishes a standard method for tagging Ethernet frames with VLAN membership information and defines the operation of VLAN bridges that permit the definition, operation, and administration of VLAN topologies within a bridged LAN infrastructure.

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.

IGPv4

interior gateway protocol version 4.

inter-AS

interautonomous system.

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.

https://t.me/learningnets

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.

IPsec

IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

IPv4

IP version 4. Internet Protocol version 4 is the fourth version in the development of the IP and the first version of the protocol to be widely deployed. Along with IPv6, IPv4 is at the core of standards-based internetworking methods of the Internet. IPv4 is still used to route most traffic across the Internet. IPv4 is a connectionless protocol for use on packet-switched link layer networks (for example, Ethernet). It operates on a best-effort delivery model in that it does not guarantee delivery and does not assure proper sequencing or avoidance of duplicate delivery.

IPv6

IP version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation).

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.

L2TP

Layer 2 Tunneling Protocol. An Internet Engineering Task Force (IETF) standards track protocol defined in RFC 2661 that provides tunneling of PPP. Based upon the best features of L2F and PPTP, L2TP provides an industry-wide interoperable method of implementing VPDN.

L3VPN

Layer 3 VPN

LAN

local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

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).

LFIB

label forwarding information base

LMI

Local Management Interface. Set of enhancements to the basic Frame Relay specification. LMI includes support for a keepalive mechanism, which verifies that data is flowing; a multicast mechanism, which provides the network server with its local DLCI and the multicast DLCI; global addressing, which gives DLCIs global rather than local significance in Frame Relay networks; and a status mechanism, which provides an on-going status report on the DLCIs known to the switch. Known as LMT in ANSI terminology.

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.

LSP

link-state packet

LSR

link-state request.

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.

MED

multi-exit discriminator.

MEF

Metro Ethernet Forum.

mGRE

multipoint Generic Routing Encapsulation.

MP-BGP

Multiprotocol Border Gateway Protocol.

https://t.me/learningnets

MP-BGPv4

MP-BGP version 4.

MP-EBGP

Multiprotocol External Border Gateway Protocol

MP-IBGP

Multiprotocol Interior 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.

MTU

maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.

NAT

Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. Also known as Network Address Translator.

NLRI

Network Layer Reachability Information.

NSSA

not-so-stubby area.

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.

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.

OSPFv2

Open Shortest Path First version 2.

OSPFv3

Open Shortest Path First version 3

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

OUI

Organizational Unique Identifier. Three octets that are assigned by the IEEE in a block of 48-bit LAN addresses.

P router

provider router.

PBB

provider backbone bridge.

PDU

protocol data unit. OSI term for packet.

PE

provider edge. Identifies the network devices, under the administrative control of the provider, that connect to CE devices.

PE-CE

provider edge-customer edge.

PHP

penultimate hop popping.

P-network

provider network

POP

point of presence. In OSS, a physical location where an interexchange carrier installed equipment to interconnect with a local exchange carrier (LEC).

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.

PVID

port VLAN ID.

PVST+

Per VLAN Spanning Tree Plus. Support for dot1q trunks to map multiple spanning trees to a single spanning tree.

PW

pseudowire.

QinQ

queue-in-queue

https://t.me/learningnets

QoS

quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.

RADIUS

Remote Authentication Dial-In User Service. Database for authenticating modem and ISDN connections and for tracking connection time.

RD

route distinguisher.

RIB

routing information base

RIP

Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.

RIPv1

Routing Information Protocol version 1

RIPv2

Routing Information Protocol version 2

RPL

Routing Policy Language

RT

route target.

SLA

service level agreement.

SOO

Site of Origin.

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.

SPT

shortest path tree

SSL

Secure Socket Layer. Encryption technology for the Web used to provide secure transactions, such as the transmission of credit card numbers for e-commerce.

https://t.me/learningnets

SSTP

Shared Spanning Tree Protocol.

STP

Spanning Tree Protocol. Bridge protocol that uses the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops, and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning Tree Protocol standard and the earlier Digital Equipment Corporation Spanning Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version generally is preferred over the Digital version.

SVI

switch virtual interface.

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.

TLV

type, length, value.

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.

VC

virtual circuit. A connection between two endpoints that appears to be a direct connection to the hosts but may actually be switched over several shared physical paths.

VC FEC

virtual circuit forwarding equivalence class.

VCCV

Virtual Circuit Connectivity Verification.

VCI

virtual circuit identifier

VFI

Virtual Forwarding Instance

VFR-Lite

Multi-VRF Customer Edge.

https://t.me/learningnets

VLAN

virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.

VPI

virtual path 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.

VPNv4

virtual private network version 4.

VPNv6

virtual private network version 6.

VPWS

Virtual Private Wire Service.

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.

VRF-Lite

Multi-VRF Customer Edge

VSI

Virtual Switch Interface

VTP

VLAN Trunking Protocol.

WAN

wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets