SPROUTE

Deploying Cisco Service Provider Network Routing

Student Guide Volume 1

Version 1.2

Part Number: 97-3415-01

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" AND AS SUCH MAY INCLUDE TYPOGRAPHICAL, GRAPHICS, OR FORMATTING ERRORS. 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: Service Provider Routing 9
- Lesson 1: Understanding Service Provider Routing Protocols 11
  - Cisco IP NGN Architecture 13
  - Overview of Routing Protocols 15
  - Interior Gateway Protocols 17
  - Routing Example 18
  - Overview of OSPF 21
  - Link-State Data Structures 23
  - OSPF Area Terminology and Router Types 25
  - OSPF Areas 26
  - OSPF Metric 27
  - Typical OSPF Designs 28
  - Overview of IS-IS 29
  - IS-IS Hierarchical Design 30
  - IS-IS Characteristics 31
  - IS-IS Router and Link Types 33
  - BGP Overview 34
  - BGP Architecture 35
  - BGP Characteristics 36
  - BGP AS Number 39
  - BGP Sessions 40
  - BGP in Customer Connections: Single-Homed Customers 43
  - BGP in Customer Connections: Dual-Attached Customers 44
  - BGP in Customer Connections: Multihomed Customers 45
  - BGP in Customer Connections: Upstream ISP 46
  - BGP in Customer Connections: Transit ISP 47
  - Summary 48
- Lesson 2: Module Summary 49
- Lesson 3: Module Self-Check 51
- Module Self-Check Answers 53
  - Answer Key 53

Module 2: Implement OSPF in the Service Provider Network 55
- Lesson 1: Introducing OSPF Routing 57
  - OSPF in the Cisco IP NGN Architecture 59
  - OSPF and OSPFv3 Key Characteristics 60
  - OSPF Route Entry Creation 61
  - OSPF Data Structures 62
  - Structure of OSPF Network 64
  - Hierarchical Structure of OSPF in a Service Provider Environment 66
  - OSPF LSA Types 68
  - OSPF Stub Areas 70
  - OSPF Not-So-Stubby Areas 71

https://t.me/learningnets

OSPF Operation 73
OSPF Best Path Calculation 75
OSPF Metric 76
Building the LSDB 78
LSA Operation 80
OSPF Link-State Database 83
OSPF Intra-Area Routing 87
OSPF Inter-Area Routing 89
OSPF External Routes 91
OSPF Virtual Link 94
Interpreting OSPF Routes in the Routing Table 95
Calculating Costs for E1 and E2 OSPF Routes 97
OSPF LSDB Overload Protection 98
Summary 99
Lesson 2: Understanding OSPF Operation 101
OSPF Functions 103
OSPF Packet Format 104
OSPF Packet Types 106
OSPF Neighbor States 108
OSPF Link-State Flooding 111
Debug OSPF Packets 113
OSPF Network Types 115
Electing the OSPF DR and BDR 117
OSPF over NBMA Network Types 119
OSPF Adjacency over Metro Ethernet and EoMPLS 121
OSPF Adjacency over Layer 3 MPLS VPN 123
Enabling OSPF on a Link with NBMA Mode 124
Subinterfaces over Multiaccess WAN Networks 126
Enabling OSPF on a Link with Point-to-Point Mode 127
Enabling OSPF on a Link with Point-to-Multipoint Mode 128
Enabling OSPF on a Link with Point-to-Multipoint Nonbroadcast Mode 130
IPv6 Support for OSPF Modes 131
Summary 132
Lesson 3: Implementing OSPF Routing 133
Implement OSPF 135
OSPF Router ID 138
OSPF Passive Interface 140
Verifying Basic OSPF 141
OSPF Virtual Links 145
Configuring OSPF Virtual Links 149
OSPF Cost 150
Cisco Nonstop Forwarding 152
Cisco Nonstop Routing 154
Cisco NSF and NSR for OSPF 155
Graceful Restart for OSPFv3 156
Bidirectional Forwarding Detection 157
Bidirectional Forwarding Detection for OSPF 159
OSPF Authentication Overview 161
Configuring OSPF Authentication 164
Summary 167
Lesson 4: Implementing OSPF Special Area Types 169
OSPF Summarization 171
OSPF Interarea Route Summarization 172
OSPF External Route Summarization 173

https://t.me/learningnets

Default Routes in OSPF 174
OSPF Area Types 176
Configuring OSPF Area Types 178
OSPF Router and LSA Types 179
OSPF Stub Area and Totally Stubby Area 180
Configuring OSPF Stub Areas 182
OSPF Not-So-Stubby Area and Totally Not-So-Stubby Area 183
Configuring OSPF NSSA and Totally NSSA 185
Summary 186
Lesson 5: Module Summary 187
Lesson 6: Module Self-Check 189
Module Self-Check Answers 203
Answer Key 203
Module 3: Implement Integrated IS-IS in the Service Provider Network 207
Lesson 1: Introducing IS-IS Routing 209
IS-IS in the Cisco IP NGN 211
Integrated IS-IS Routing 212
Integrated IS-IS Design Principles 213
Similarities Between IS-IS and OSPF 215
IS-IS Addressing 219
IS-IS Router Types 224
IS-IS Routing Logic 226
Asymmetric IS-IS Routing 228
Symmetric IS-IS Routing 230
IS-IS Packets 231
Integrated IS-IS for IPv6 234
IS-IS Network Types 235
IS-IS Operations 236
IS-IS LSP Flooding 239
IS-IS LSDB Synchronization 240
IS-IS Adjacencies 242
IS-IS Single Topology Restrictions 243
Multitopology IS-IS for IPv6 245
Summary 246
Lesson 2: Implementing Integrated IS-IS Routing 247
Implement OSI Area Routing 248
Implement IS-IS Routing 250
Optimizing the IS-IS Processes 252
Bidirectional Forwarding Detection for IS-IS 254
Nonstop Forwarding for IS-IS 256
IP Route Summarization Configurations in IS-IS Networks 257
Verification of IS-IS 258
Troubleshooting IS-IS Commands 260
Configuring IS-IS to Support IPv6 263
Summary 265
Lesson 3: Module Summary 267
Lesson 4: Module Self-Check 269
Module Self-Check Answers 279
Answer Key 279
Module 4: Implement BGP in the Service Provider Network 281
Lesson 1: Enterprise Connectivity to Service Providers 283
Enterprise Connectivity Requirements 284
Routing Method Selection 285

https://t.me/learningnets

Connection Options....288
Single-Homed Enterprises....289
Dual-Homed Enterprises....290
Multihomed Enterprises....291
Dual-Multihomed Enterprises....292
Summary....293
Lesson 2: Introducing BGP Routing....295
BGP Terminology....296
Autonomous System and BGP....298
BGP Routing Between Autonomous Systems....299
BGP Path Vector Functionality....300
BGP Routing Policies....301
BGP Features....302
BGP Databases....303
BGP Message Types....304
Multiprotocol Extensions for BGP4....306
MPLS VPN Overview....308
Summary....309
Lesson 3: Implementing Basic BGP Routing....311
Planning for BGP....313
Configure Basic EBGP....314
Advertising BGP Networks....316
Configure Basic IBGP....318
Full-Mesh IBGP....320
BGP Support for IPv6....321
Shutting Down a BGP Neighbor....322
BGP Next-Hop Behavior....323
BGP Next Hop Self....324
Cisco IOS XR BGP Configuration Templates....325
BGP Neighbor States....327
BGP Neighbor Authentication....331
Clearing the BGP Session....332
Monitoring BGP Routes....334
BGP Path Selection....335
BGP Route Selection Decision Process....336
BGP Weight Attribute....338
BGP Local Preference Attribute....339
BGP AS Path Attribute....340
BGP Multi-Exit Discriminator Attribute....341
Summary....342
Lesson 4: Module Summary....343
Lesson 5: Module Self-Check....345
Module Self-Check Answers....353
Answer Key....353
Glossary....355

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Course Introduction

Overview

The Deploying Cisco Service Provider Network Routing (SPROUTE) v1.2 is a five-day course that provides service provider professionals with information on the use of advanced routing in implementing scalability for Cisco routers that are connected to LANs and WANs.

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.

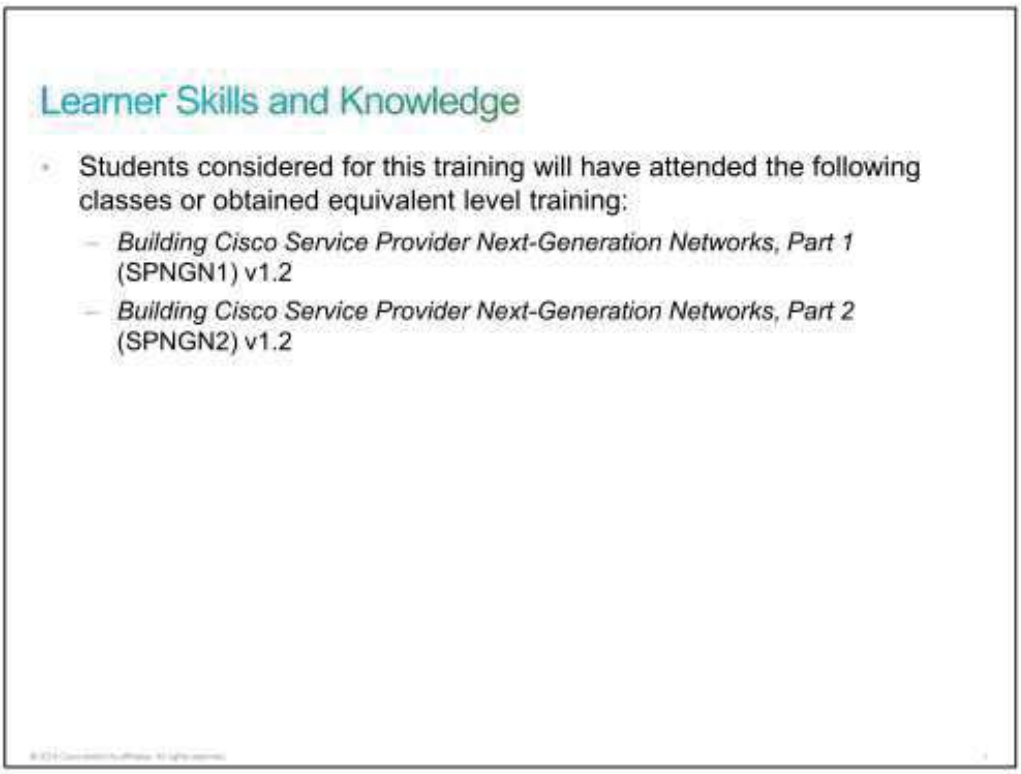

https://t.me/learningnets

Course Goal and Objectives

This topic describes the course goal and objectives.

Upon completing this course, you will be able to meet these objectives:

- Identify the typical routing requirements, and list the routing protocols that are used in service provider networks
- Describe the steps that are needed to implement OSPF in a service provider network
- Describe the importance of the Integrated IS-IS routing protocol for internal routing, and list the steps to follow when you are implementing Integrated IS-IS in a service provider network
- Implement BGP to connect an enterprise to a service provider and to connect a service provider to an upstream service provider
- Describe the tools that are used for routing protocol manipulation, route redistribution, and BGP route selection

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Course Flow

This topic presents the suggested flow of the course materials.

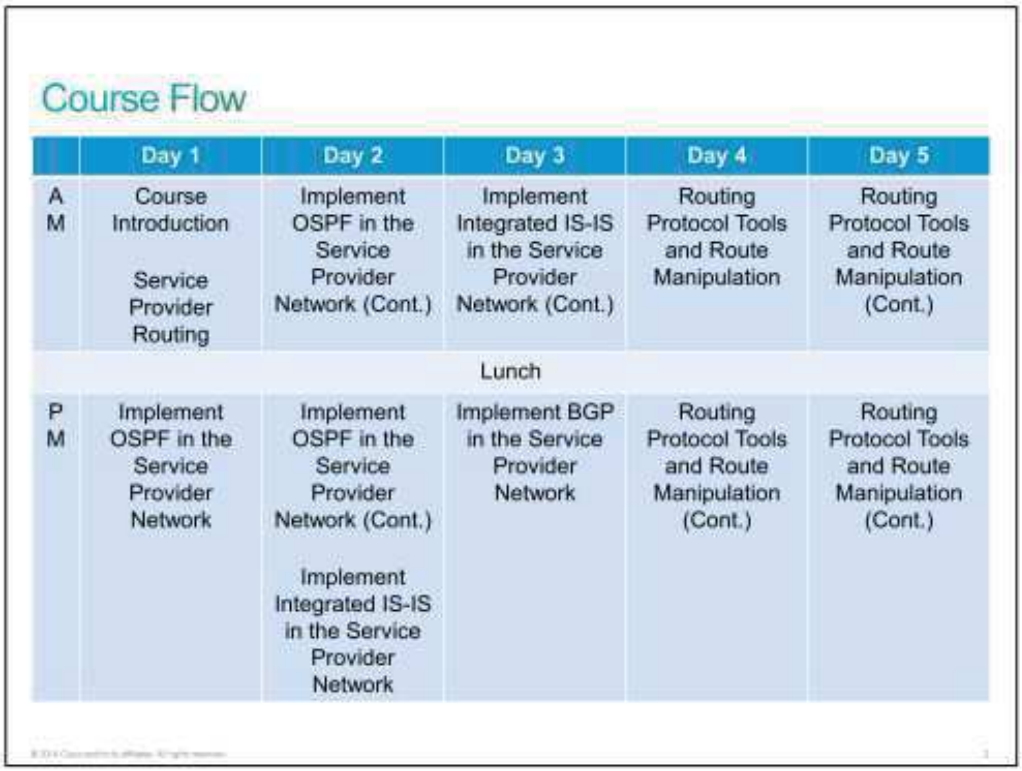

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the lab activities. The exact timing of the subject materials and labs depends on the pace of your specific class.

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

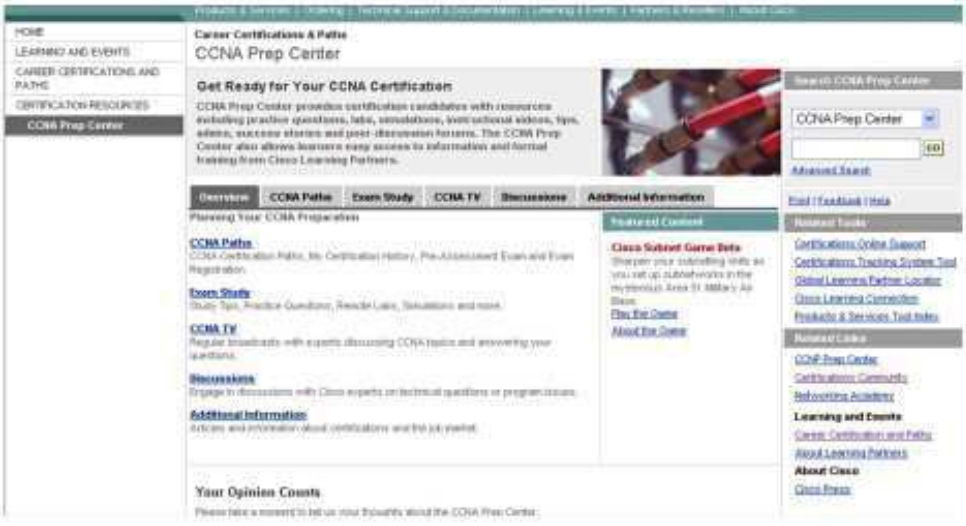

Additional information is available at http://www.cisco.com/go/prepcenter.

Learner Introductions

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

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

Additional References

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

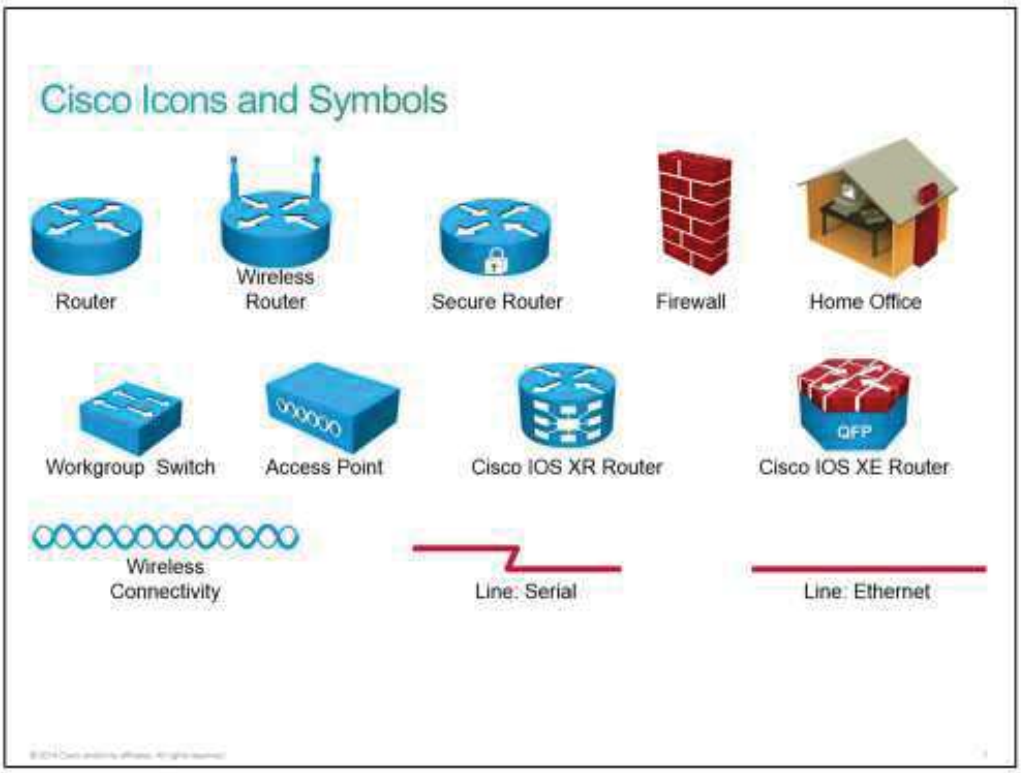

Cisco Glossary of Terms

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

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 1: Service Provider Routing

Introduction

This module identifies the typical routing requirements and lists the routing protocols in service provider networks.

Upon completing this module, you will be able to:

- Describe service provider routing protocols

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Understanding Service Provider Routing Protocols

Overview

This lesson describes the main characteristics of routing protocols that are used in service provider environments. The lesson describes how a service provider ensures IP connectivity to the Internet for end customers and other service providers. The lesson explains the need to exchange Internet routing information via BGP. On the other hand, IGPs are responsible for providing IP connectivity within an AS. The most common IGP protocols in a service provider network are OSPF and IS-IS, which will be briefly described, as well as the BGP routing protocol.

Upon completing this lesson, you will be able to:

- Describe the Cisco IP Next Generation Network architecture
- Describe the characteristics and requirements for routing protocols in service provider environments
- Describe interior gateway protocols
- Provide a routing example
- Describe the characteristics of OSPF in service provider environments
- Describe link-state data structures
- Describe area terminology and router types
- Describe OSPF areas
- Describe OSPF metrics
- Describe typical OSPF designs
- Describe the characteristics of IS-IS in service provider environments

https://t.me/learningnets

- Describe IS-IS hierarchical design
- Describe the characteristics of IS-IS
- Describe IS-IS router and link types
- Describe the characteristics of BGP in service provider environments
- Describe BGP architecture
- Describe the characteristics of BGP
- Describe BGP AS numbers
- Describe how BGP sessions are established
- Describe the use of BGP in a single-homed customer environment
- Describe the use of BGP in a dual-attached customer environment
- Describe the use of BGP in a multihomed customer environment
- Describe the use of BGP in an upstream ISP customer environment
- Describe the use of BGP in a transit ISP customer environment

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Cisco IP NGN Architecture

The Cisco IP Next Generation Network architecture enables service providers to develop fixed and mobile convergence, starting with the transport in the access, aggregation, and core networks. The Cisco IP NGN targets service providers that have an existing centralized wireline services edge network and that are willing to maintain and evolve this network layer.

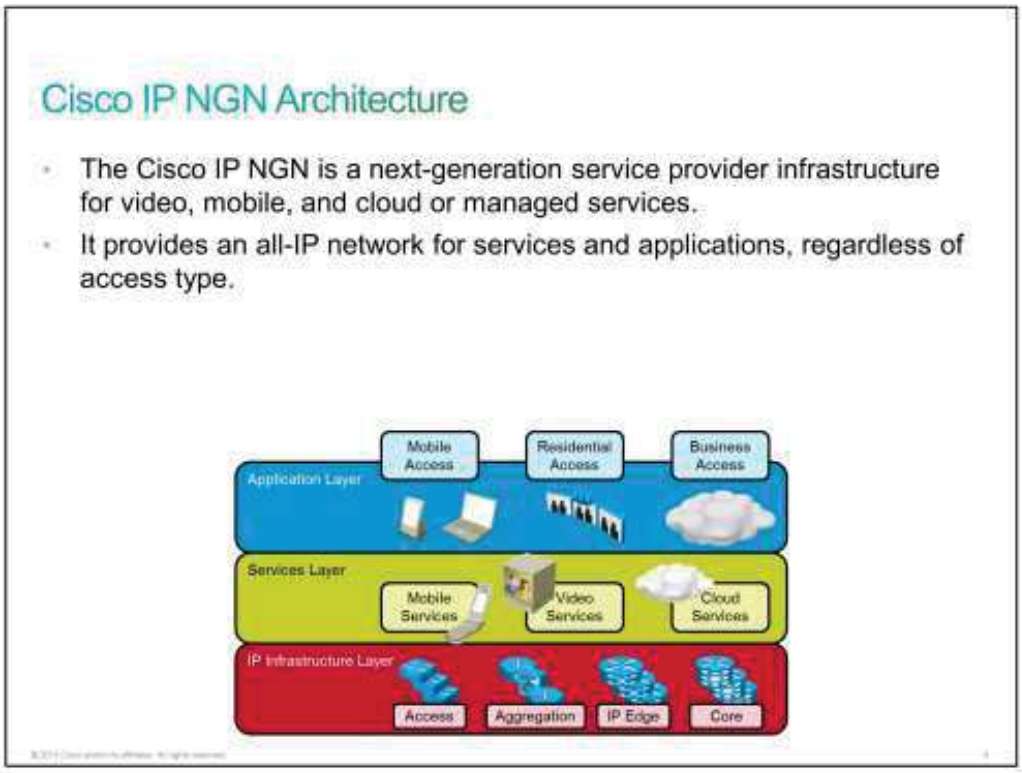

The Cisco IP NGN architecture consists of a flexible, comprehensive, and generic framework that is structured around common layers in the service provider networks: customer premises, access network, aggregation network, edge network, core network, network management, and network admission. The access, aggregation, and core layers are used for the transport of mobile, video, and cloud-managed services.

The purpose of the Cisco IP NGN networks is to provide all-IP transport for services and applications regardless of access type. IP infrastructure, service, and application layers are separated in Cisco IP NGN networks, enabling the addition of new services and applications without changes in the transport network.

https://t.me/learningnets

Cisco IP NGN Architecture (Cont.)

Routing protocols that are used in service provider environments focus on the following:

- *IP infrastructure layer* of the Cisco IP NGN
- Service provider *core* and *edge* devices and customer devices

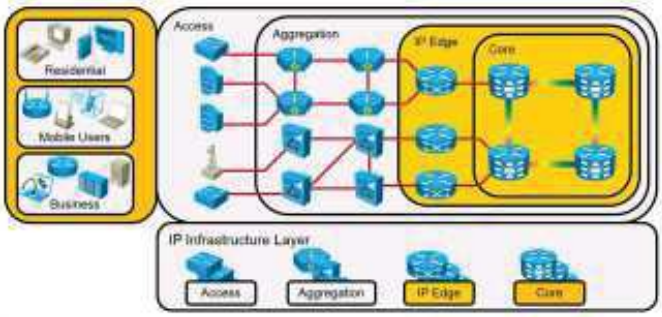

Routing protocols, such as OSPF and IS-IS, are used to carry information about IP networks between routers. Routing protocols are part of the IP infrastructure layer of the Cisco IP NGN, and they are used on service provider core and edge routers as well as customer routers.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

Overview of Routing Protocols

Depending on the type of service provider, there are many different connectivity requirements.

Overview of Routing Protocols

Requirements can be summarized as follows:

- Provide connectivity to the Internet for end customers and subordinate ISPs
- Optionally, provide transit connectivity between service providers (that are Tier 1 ISPs)

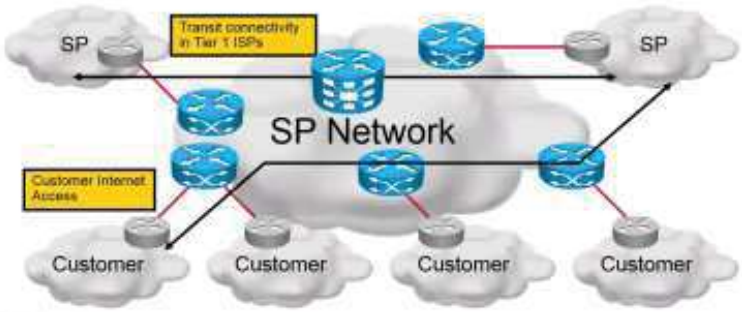

In addition, you must consider local routing within the service provider network to ensure reachability for all local addresses.

Overview of Routing Protocols (Cont.)

IGP and BGP routing:

- **IGP:** exchange local routing information
- **BGP:** exchange external routing information

https://t.me/learningnets

BGP is used to exchange the Internet routing information with other ISPs and the customers who require the information. BGP can be configured to only propagate a default route (for example, for customers), a portion of the complete Internet routing table (for example, for multihomed customers), or the entire Internet routing table (for example, for multihomed customers and subordinate ISPs).

An IGP is used to provide connectivity within an AS. The most important function of an IGP is to provide reachability of BGP neighbors and BGP next-hop addresses.

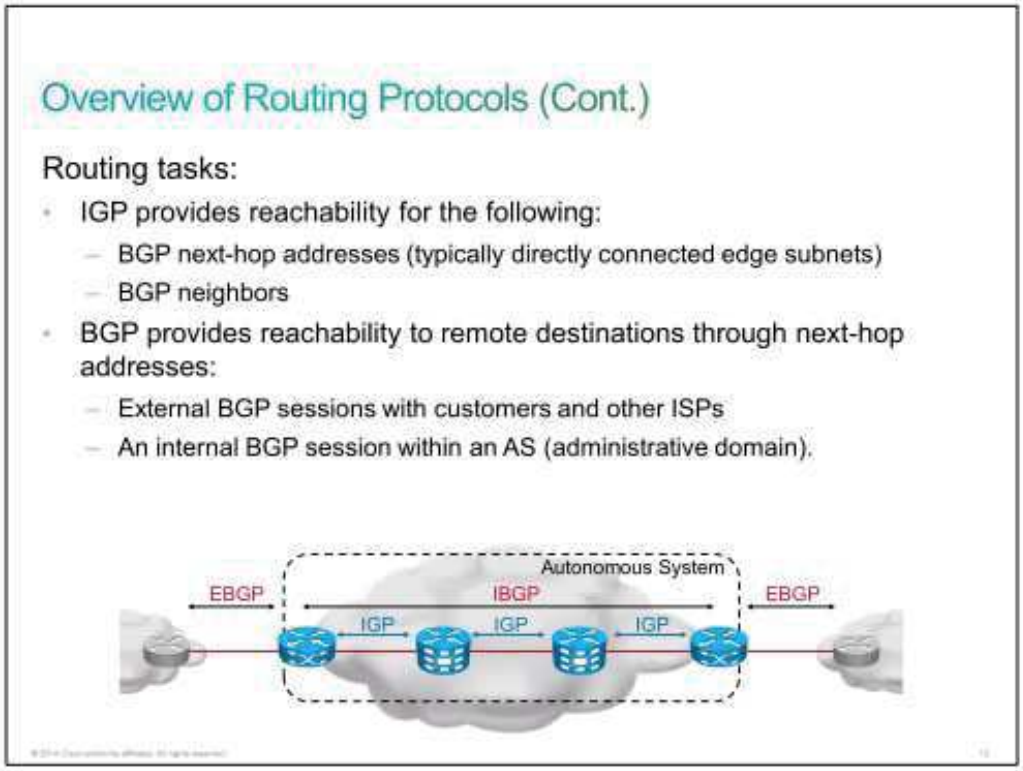

There are two characteristics of BGP that require the assistance of an IGP:

- BGP next-hop addresses do not change as the BGP routes are propagated through an AS.
- Internal BGP neighbors can be several hops away. (External BGP neighbors are typically reachable through a directly connected edge subnet.)

An IGP is therefore needed to propagate the following:

- Next-hop addresses (IP addresses of external BGP neighbors) throughout the AS.
- IP addresses of internal BGP neighbors (typically, loopback addresses).

The simplified illustration shows the three components in a service provider routing environment:

- EGBP sessions with other ASs to exchange the Internet routing information.
- IBGP sessions to carry external routing information across the service provider AS to all routers that require it.
- IGP to provide the reachability of next-hop and neighbor addresses.

https://t.me/learningnets

Interior Gateway Protocols

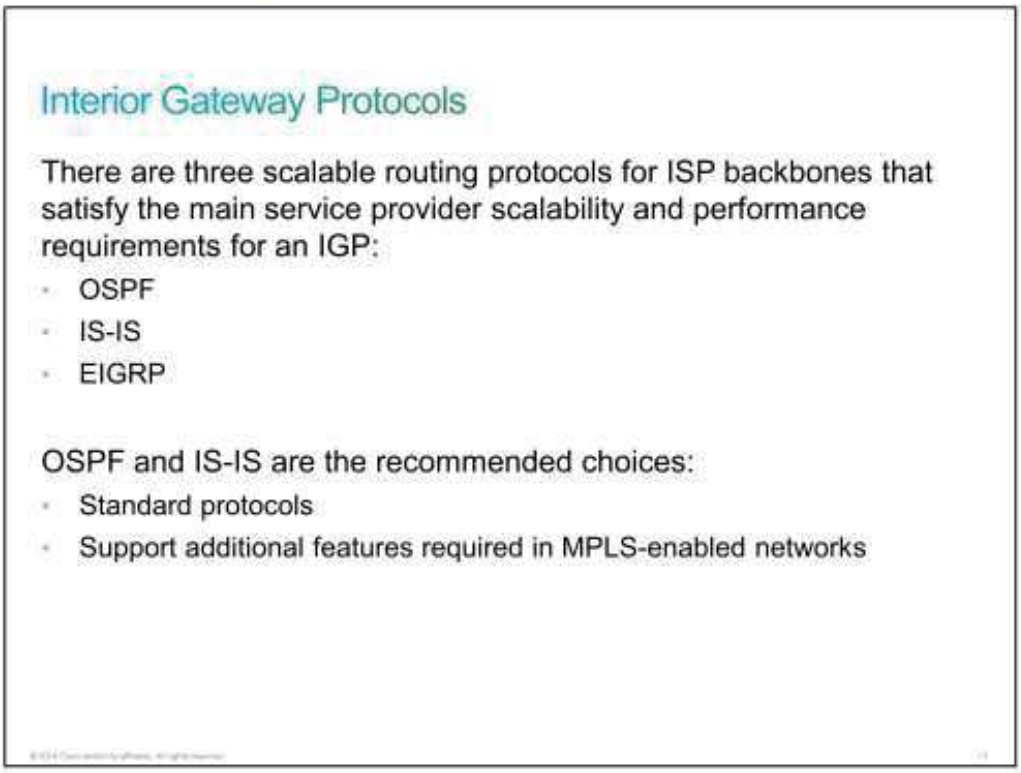

Most service providers today use either OSPF or IS-IS for two reasons:

- EIGRP is Cisco proprietary and may hinder interoperability with devices from other vendors.
- Cisco Multiprotocol Label Switching Traffic Engineering requires the help of a link-state protocol, while EIGRP is a distance-vector protocol.

https://t.me/learningnets

Routing Example

Here you can see some routing examples.

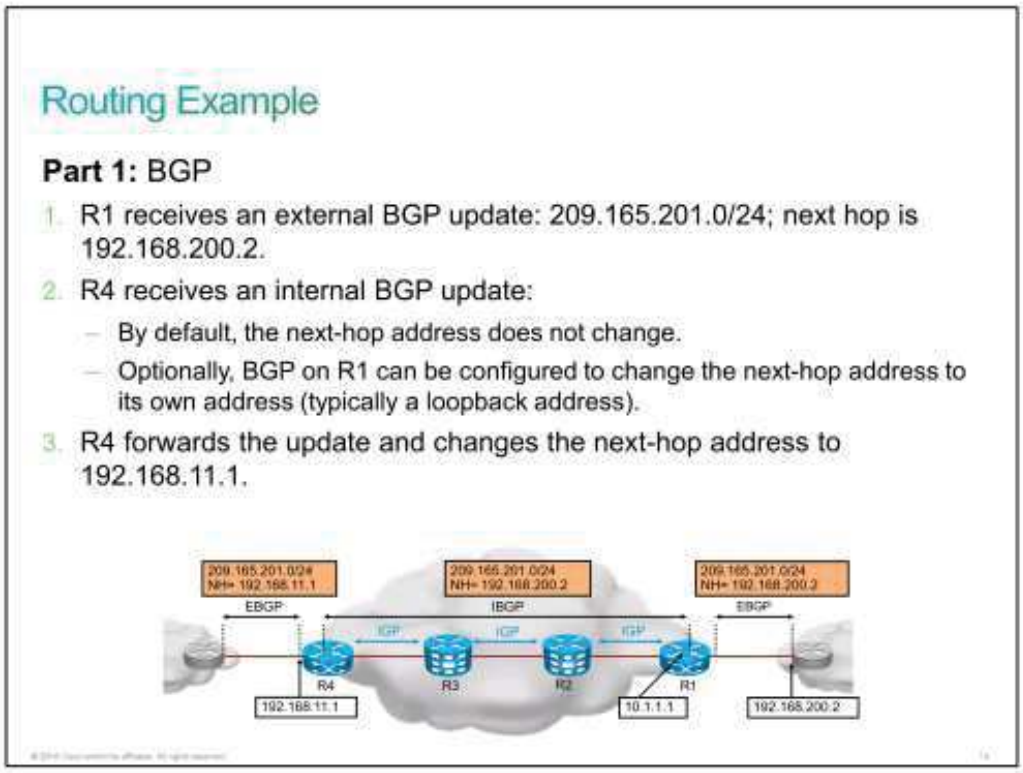

The figure illustrates a sample network where a route is received from an AS over an EBGP session. The route is then forwarded to all other internal routers over an IBGP session. Egress routers will forward the route to other external neighbors.

Note that the next-hop address in the update does not change as the update is forwarded to an internal neighbor. This is the default behavior. Alternatively, if router R1 is configured with the next-hop-self feature, R1 would change the next-hop attribute to its loopback address that is used as the source address for the IBGP session to R4.

After R4 sends the update out to an external neighbor, it will change the next-hop address to its own IP address that is used for the EBGP session with the external neighbor. This is the default behavior.

https://t.me/learningnets

Routing Example (Cont.)

Part 2: IGP

- R1 propagates the BGP next-hop address to all routers in the domain:
  - Edge subnet (192.168.200.0/30) is for the reachability of external BGP next-hop addresses.
  - Loopback address (10.1.1.1/32) is for the reachability of internal BGP neighbors.
- R2 and R3 forward the information:
  - Unchanged (required if the network also uses MPLS-based services such as MPLS VPNs and Cisco MPLS TE).
  - Optionally, summarization can be used within IGP for optimization.

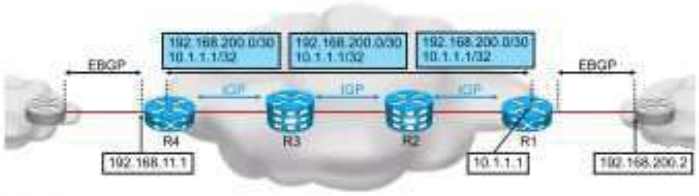

The figure (part 2 of this example) illustrates the required functionality of an IGP in order to support the BGP functionality.

The IGP is propagating two important addresses throughout the AS:

- The IP address of the external neighbor, which is also the next-hop address that is carried within BGP updates coming from the external neighbor. Note that this part is optional if you use the alternative method (using the next-hop-self feature).
- The loopback IP address of the ingress BGP router that is used for all IBGP sessions.

If the same service provider network is also used to provide other services, such as MPLS-based VPNs or Cisco Multiprotocol Label Switching Traffic Engineering, it is important *not* to summarize the next-hop IP addresses in the backbone because this would break MPLS LSPs.

https://t.me/learningnets

Routing Example (Cont.)

Part 3: Routing Table

- End-to-end connectivity is provided through recursive routing table lookups (optimized by Cisco Express Forwarding):
  - BGP for end prefixes
  - IGP for BGP next-hop reachability

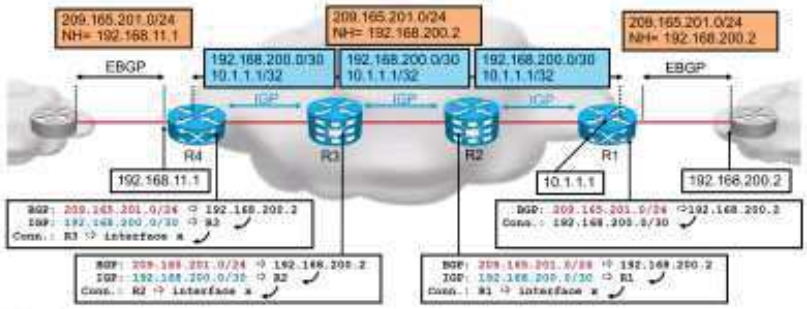

The figure (part 3 of this example) illustrates the final result of BGP and IGP routing information propagation as reflected by the routing table. A recursive set of routes can be observed in the routing tables of all routers:

Note	The router will look up the route to the address with BGP and get a next hop, then look up the route to the next hop, which is learned from the IGP (or static routes).

External BGP routes point to BGP next-hop addresses. Router R1 receives multiple external routes, which all use the same next-hop address that is the interface address of the external BGP peer.

BGP next-hop addresses point to these peers:

- Directly connected external peer on ingress edge routers (router R1 in the example).
- Nonadjacent addresses that are reachable via the IGP (R2, R3, and R4 in the example require reachability to the BGP next-hop address via the IGP). IGP peers are reachable through an attached link.

Note	For performance reasons, routers do not perform recursive lookup when forwarding packets. Cisco Express Forwarding is used to optimize the forwarding table for performance.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Overview of OSPF

The OSPF protocol was developed due to a need to introduce a high-functionality, nonproprietary IGP for the TCP/IP protocol family. The discussion of the creation of a common interoperable IGP for the Internet started in 1988 and did not become formalized until 1991. At that time, the OSPF Working Group requested that OSPF be considered for advancement to Draft Internet Standard.

The OSPF protocol is based on link-state technology, which is a departure from the Bellman-Ford vector-based algorithms that are used in traditional Internet routing protocols such as RIP. OSPF has introduced new concepts such as authentication of routing updates, VLSMs, route summarization, and so on.

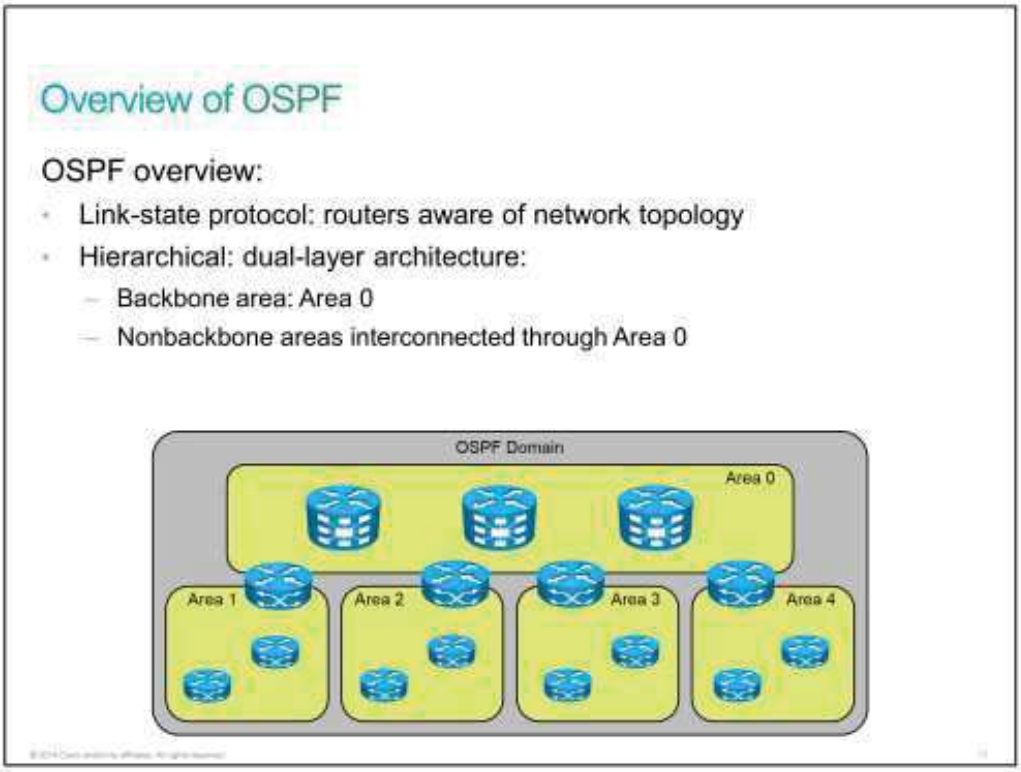

OSPF uses a link-state algorithm in order to build and calculate the shortest path to all known destinations. The algorithm by itself is quite complicated. The following is a high-level, simplified way of looking at the various steps of the algorithm:

- Upon initialization or due to any change in routing information, a router will generate a link-state advertisement. This advertisement will represent the collection of all link states on this router.
- All routers will exchange link states by means of flooding. Each router that receives a link-state update should store a copy in its link-state database and then propagate the update to other routers.
- After the database of each router is completed, the router will calculate a shortest path tree to all destinations. The router uses Dijkstra's algorithm to calculate the shortest path tree. The destinations, the associated cost, and the next hop to reach those destinations will form the IP routing table.
- If no changes in the OSPF network occur, such as the cost of a link or a network being added or deleted, then OSPF should be quiet.

https://t.me/learningnets

OSPF uses flooding to exchange link-state updates between routers. Any change in routing information is flooded to all routers in the network. Areas are introduced to put a boundary on the growth of link-state updates. Flooding and calculation of Dijkstra's algorithm on a router is limited to links within an area. All routers within an area have the exact link-state database. Routers that belong to multiple areas and connect these areas to the backbone area are called ABRs. ABRs must therefore maintain information describing the backbone areas and any other attached areas.

ABRs also provide mechanisms for aggregating routes and cutting down on the unnecessary propagation of subnet information.

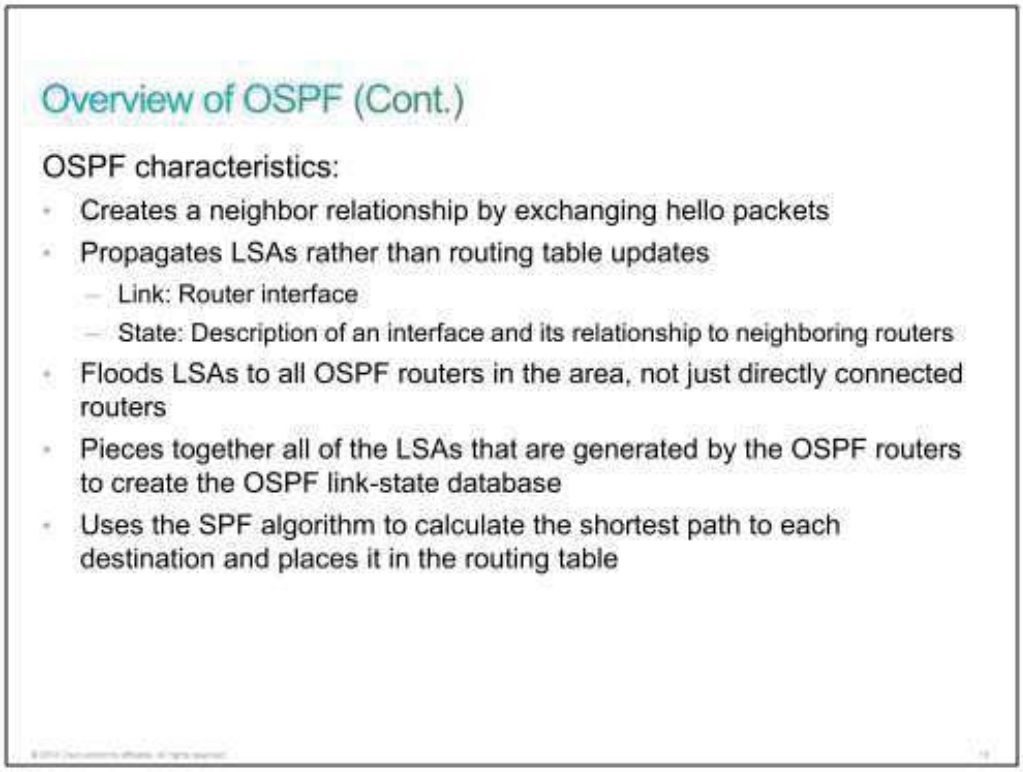

OSPF is a link-state routing protocol. You can think of a link as an interface on a router. The state of the link is a description of this interface and of its relationship to its neighboring routers. A description of the interface would include, for example, the IP address of the interface, the subnet mask, the type of network to which it is connected, the routers that are connected to this network, and so on. The collection of all of these link states forms a link-state database.

A router sends LSA packets to advertise its state periodically (every 30 minutes) and immediately when the router state changes. Information about attached interfaces, metrics that are used, and other variables are included in OSPF LSAs. As OSPF routers accumulate link-state information, they use the SPF algorithm to calculate the shortest path to each node.

A topological (link-state) database is, essentially, an overall picture of networks in relation to routers. The topological database contains the collection of LSAs that are received from all routers in the same area. Because routers within the same area share the same information, they have identical topological databases.

Note	OSPF can operate within a hierarchy. The largest entity within the hierarchy is the AS, which is a collection of networks under a common administration that share a common routing strategy. An AS can be divided into a number of areas, which are groups of contiguous networks and attached hosts.

https://t.me/learningnets

Link-State Data Structures

There are several structures that enable keeping records of link-state information.

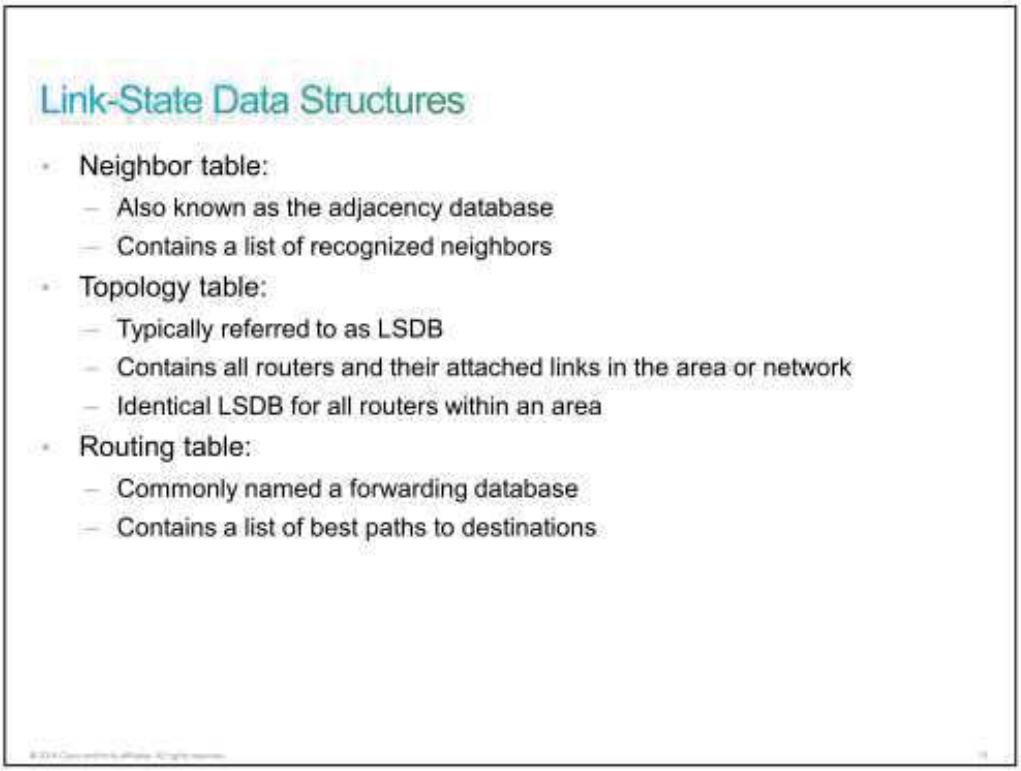

OSPF and IS-IS are classified as link-state routing protocols because of the manner in which they distribute routing information and calculate routes.

Link-state routing protocols collect routing information from all other routers in the network or within a defined area of the network. When link-state routing protocols have collected this information from all routers, each router independently calculates its best paths to all destinations in the network, using Dijkstra's algorithm.

Incorrect information from any particular router is less likely to cause confusion because each router maintains its own view of the network.

For consistent routing decisions to be taken by all of the routers in the network, each router must keep a record of the following information:

- **Its immediate neighbor routers:** If the router loses contact with a neighboring router, within a few seconds, it will invalidate all paths through this router and recalculate its paths through the network. Adjacency information about neighbors is stored in the neighbor table, also known as an adjacency database, in OSPF.
- **All of the other routers in the network, or in its area of the network, and their attached networks:** The router recognizes other routers and networks through LSAs, which are flooded through the network. LSAs are stored in a topology table, also called an LSDB.
- **The best paths to each destination:** Each router independently calculates the best paths to each destination in the network using Dijkstra's algorithm. The best paths are then offered to the routing table or forwarding database. Packets arriving at the router are forwarded based on the information that is held in the routing table.

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Area Terminology and Router Types

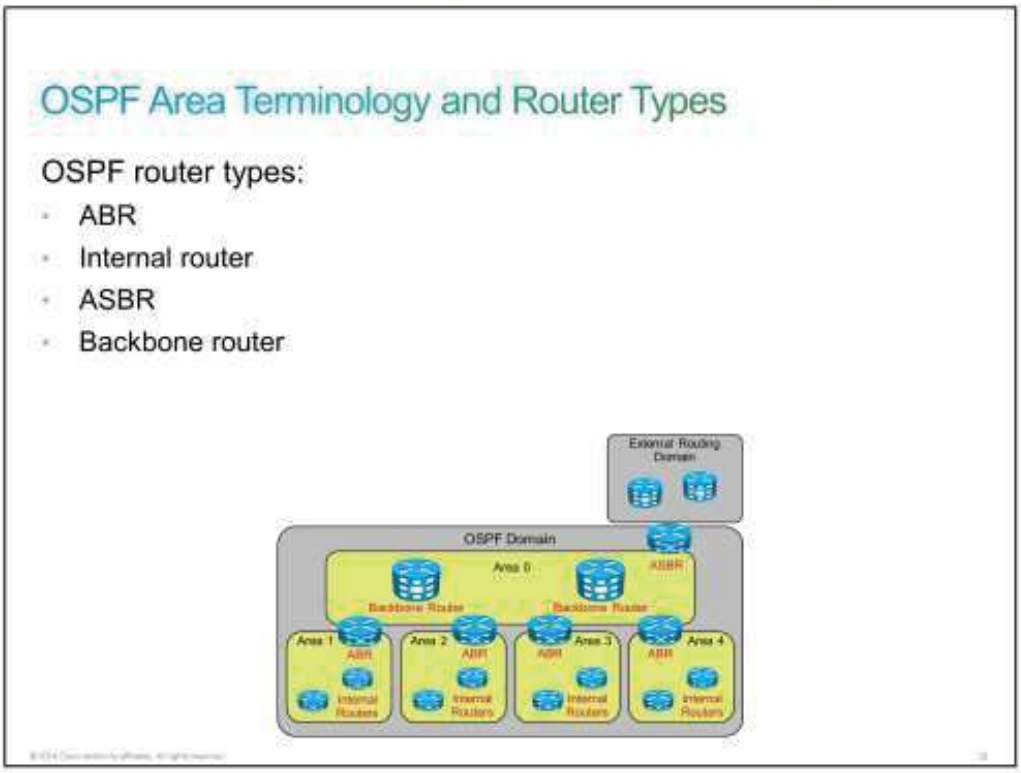

All OSPF areas and routers running an OSPF routing protocol compose the OSPF AS. Routers that make up nonbackbone (normal) areas are known as internal routers, and they have all of the interfaces in one area only. Routers that make up Area 0 are known as backbone routers (internal routers in the backbone). OSPF hierarchical networking defines Area 0 as the core. All other areas connect directly to backbone Area 0. An ABR connects Area 0 to the nonbackbone areas. An OSPF ABR plays a very important role in network design and has interfaces in more than one area. An ABR has the following characteristics:

- It separates LSA flooding zones.
- It becomes the primary point for area address summarization.
- It functions regularly as the source for default routes.
- It maintains the LSDB for each area with which it is connected.

The ideal design is to have each ABR connected to two areas only, the backbone and another area, with three areas being the upper limit.

An ASBR connects any OSPF area to a different routing administration (such as BGP or EIGRP). The ASBR is the point where external routes can be redistributed into OSPF.

https://t.me/learningnets

OSPF Areas

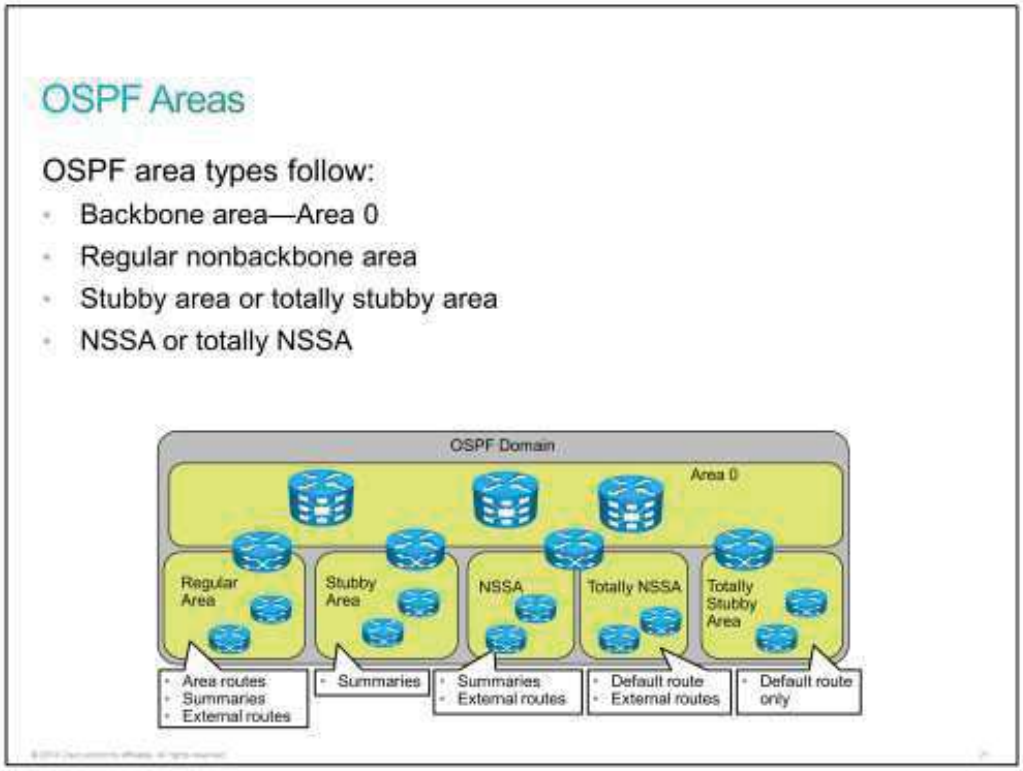

In general, there are six different types of areas:

- Backbone area or Area 0, which typically carries all of the routing information
- Regular nonbackbone areas
- Stubby areas that do not originate or receive any external routes
- Totally stubby areas that do not originate (redistribution from other protocols) or receive any external routes or summaries (only the default route)
- NSSAs that can originate external routes (redistribution from other protocols) but do not receive them from other OSPF areas
- Totally NSSAs that can originate external routes (redistribution from other protocols) but do not receive them from other areas, and the NSSAs do not receive summaries (only the default route)

©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Metric

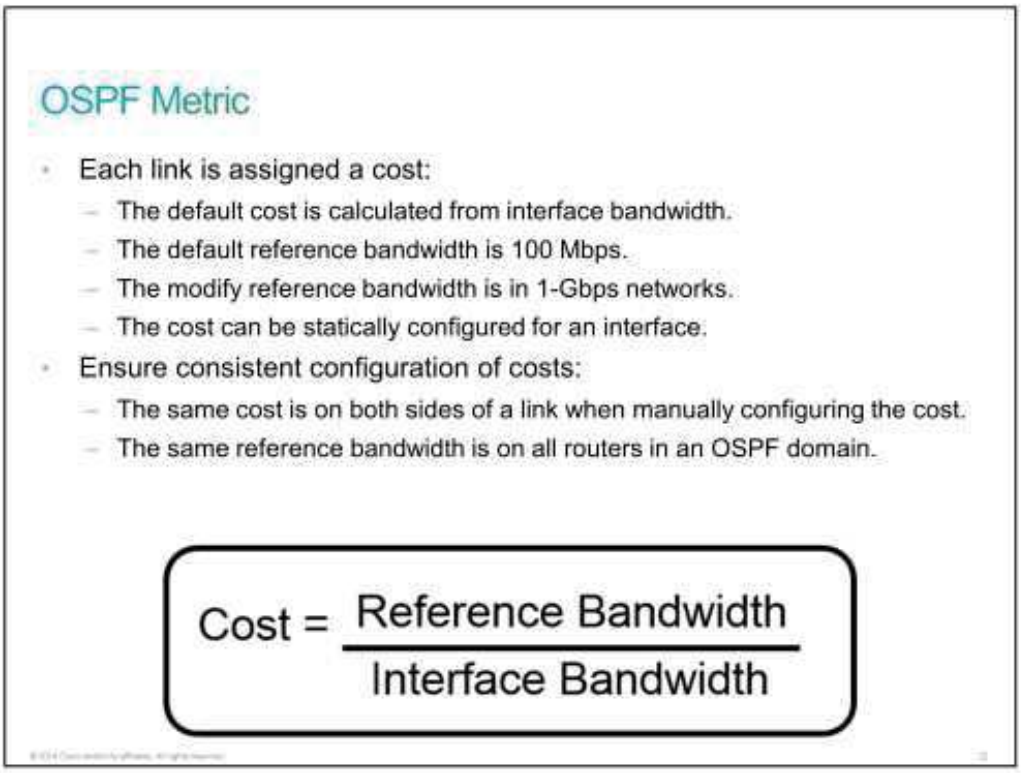

Each link is assigned a cost. The default cost is calculated from the interface bandwidth and the reference bandwidth, which defaults to 100 Mbps. The figure illustrates the formula that is used to calculate costs for individual links based on the default or configured link bandwidth (note that this is not the actual link speed).

The default reference bandwidth is only useful in networks where there are no links faster than 100 Mbps. In faster environments, different links may be given the same cost. For example, 10-Gbps, 1-Gbps, and 100-Mbps links would all be assigned a cost of 1.

Alternatively, cost can be statically configured on a per-link basis. Make sure that you configure costs consistently (the same cost on both routers).

Also make sure that you configure the same reference bandwidth domain-wide.

https://t.me/learningnets

Typical OSPF Designs

Two typical ways of implementing OSPF are the single-area design and multiarea design.

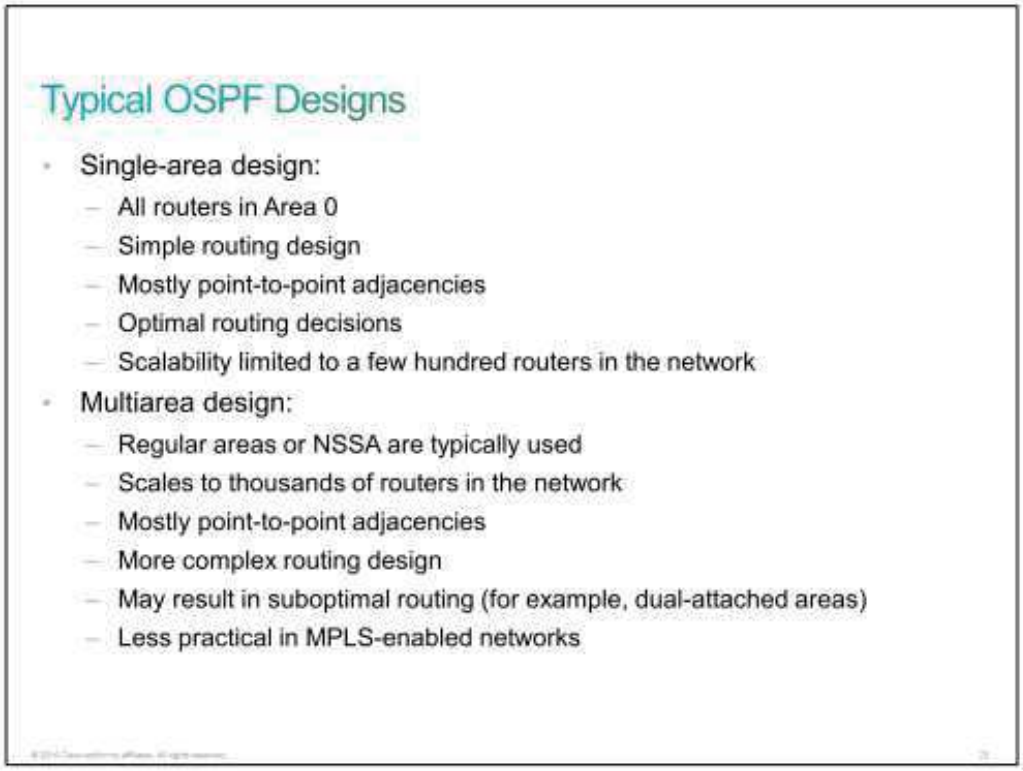

Most modern service provider networks use MPLS with some of the MPLS-based solutions. When implementing MPLS-based VPNs and Cisco Multiprotocol Label Switching Traffic Engineering, it is important to consider the interaction between MPLS Label Distribution Protocol and an IGP. If summarization is used for addresses for which LDP is used to generate LSPs, it will break those LSPs and consequently break MPLS VPNs.

From a design and implementation perspective, it is preferred to implement OSPF using one area (Area 0). The limitation of this approach is the scalability, which is mostly influenced by the number of nodes (routers) in an area.

In large service provider environments, you may be forced to use a hierarchical design. The characteristics and limitations of the hierarchical approach must be considered when designing MPLS solutions.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Overview of IS-IS

IS-IS is a popular IP routing protocol in the ISP industry. The simplicity and stability of IS-IS make it robust in large internetworks. IS-IS is found in large ISPs and in some networks that support OSI protocols.

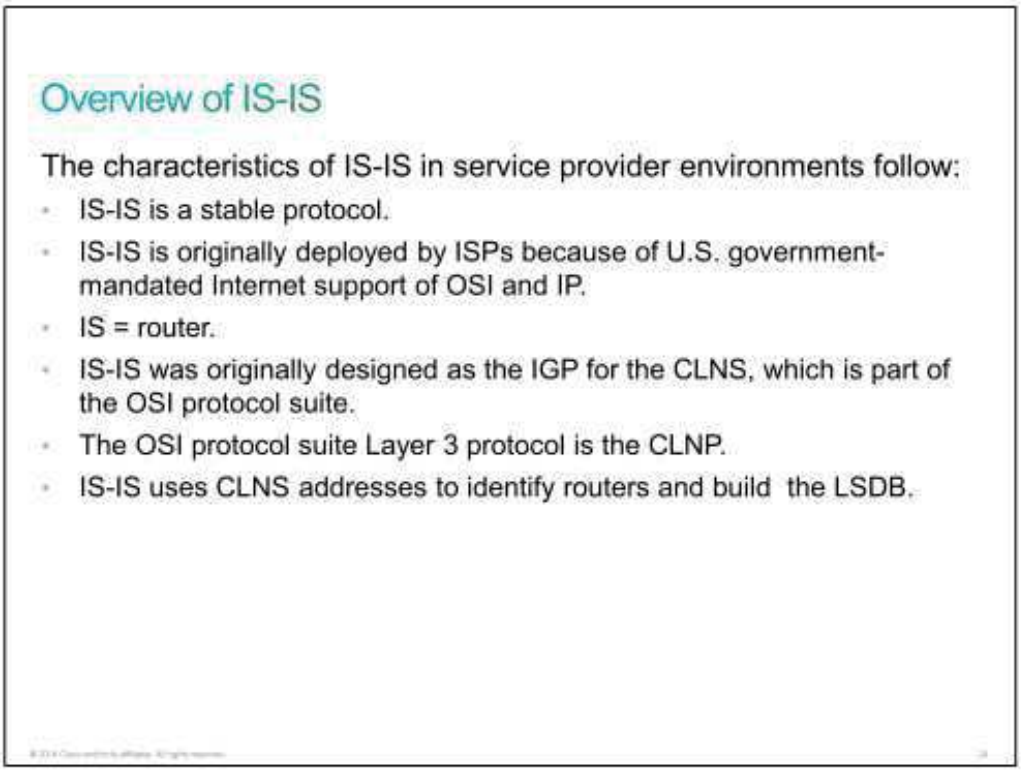

IS-IS development began before the development of OSPF. Large ISPs chose IS-IS because of their unique requirement for scalability, convergence, and stability. The U.S. government also required support for OSI protocols in the early Internet. Although this requirement was later dropped, IS-IS met both constraints.

ISO specifications refer to routers as intermediate systems. Thus, IS-IS is a protocol that allows routers to communicate with other routers. The OSI suite uses CLNS to provide connectionless delivery of data, and the actual Layer 3 protocol is CLNP.

IS-IS uses CLNS addresses to identify the routers and to build the LSDB. IS-IS serves as an IGP for the CLNS. CLNP is the solution for "unreliable" (connectionless) delivery of data, which is similar to IP.

https://t.me/learningnets

IS-IS Hierarchical Design

IS-IS protocol supports a two-level hierarchy for managing and scaling routing in large networks.

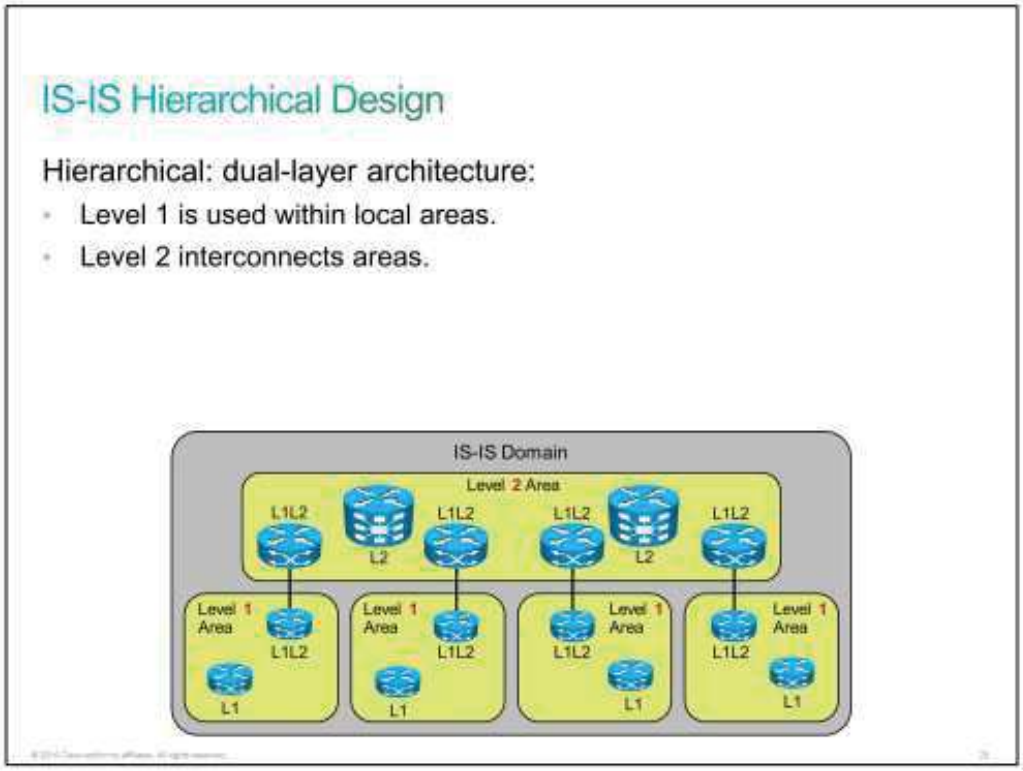

The figure illustrates a hierarchical IS-IS design with these characteristics:

- Level 2 area with Level 2-only routers in the core
- Level 2 area edge with Level 1 and Level 2 routers to connect to other areas using Level 1
- Level 1 areas with Level 1 and Level 2 routers connecting areas to Level 2
- Level 1 areas with Level 1-only routers.

This is a generic representation of what can be implemented using IS-IS. Like OSPF, IS-IS can also be implemented in a simpler fashion with fewer areas or simply one area and one level. Unlike OSPF, all areas do not have to connect to a common backbone area.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

IS-IS Characteristics

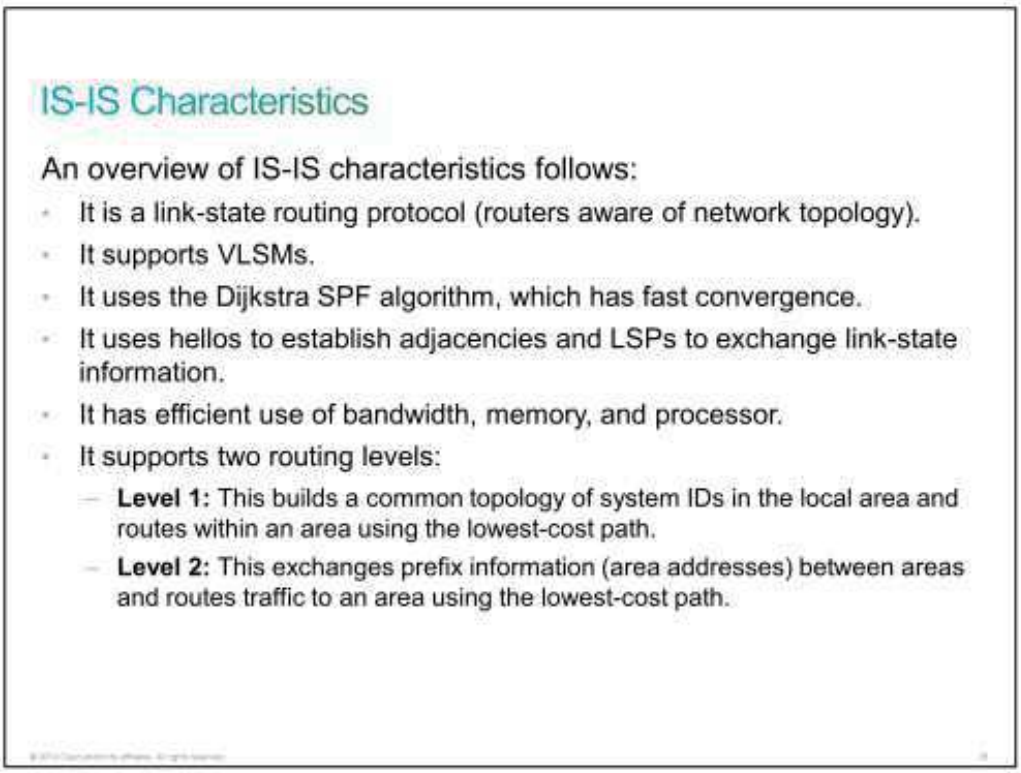

IS-IS is the dynamic link-state routing protocol for the OSI protocol stack. It distributes routing information for routing CLNP data for the ISO CLNS environment. IS-IS operates similarly to OSPF. IS-IS allows the routing domain to be partitioned into areas. IS-IS routers establish adjacencies using a Hello protocol and exchange link-state information using LSPs throughout an area to build the LSDB.

Each router then runs the Dijkstra SPF algorithm against its LSDB to pick the best paths. There is a minimal amount of information that is communicated between areas, which reduces the burden on routers supporting the protocol.

IS-IS routing takes place at two levels within an AS: Level 1 and Level 2.

- Level 1 routing occurs within an IS-IS area. It recognizes the location of the end systems and intermediate systems, and then it builds a routing table to reach each system. All devices in a Level 1 routing area have the same area address. Routing within an area is accomplished by looking at the locally significant address portion (known as the system ID) and choosing the lowest-cost path.
- Level 2 routers learn the locations of Level 1 routing areas and build an interarea routing table. All ISs in a Level 2 routing area use the destination area address to route traffic using the lowest-cost path.

https://t.me/learningnets

IS-IS Characteristics (Cont.)

More IS-IS characteristics follow:

- Each router has topology information for its area.
- IS-IS is part of OSI and was originally used with CLNS only.
- IS-IS still uses CLNS to maintain adjacencies and build an SPF tree.
- Integrated IS-IS can also carry IP routing information in its updates.
- The wide-style metric should be used for large, high-speed service provider networks (24-bit link metric and 32-bit path metric).
- The link cost defaults to 10.
- Each router is identified using a unique NSAP address.

Like OSPF, IS-IS is also a link-state protocol using Dijkstra's algorithm, in which each router has topology information for its area. IS-IS is part of the OSI standard protocol suite and was originally used with CLNS.

Each router is identified using a unique NSAP address, which is part of the CLNS protocol. IS-IS still uses CLNS to maintain adjacencies and build SPF trees, but the integrated version of IS-IS can be used for other protocols, such as IP, and can also have extensions for Cisco Multiprotocol Label Switching Traffic Engineering.

The wide-style metric should be used for large, high-speed service provider networks (24-bit link metric and 32-bit path metric). The link cost defaults to 10 and can be modified to reflect the desired cost. The narrow-style metric can only accommodate 64 metric values, which is typically insufficient in modern networks and may not even be compatible with IS-IS extensions such as those for Cisco MPLS TE.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

IS-IS Router and Link Types

The interaction between routers in an IS-IS domain depends on router types, link types, and area.

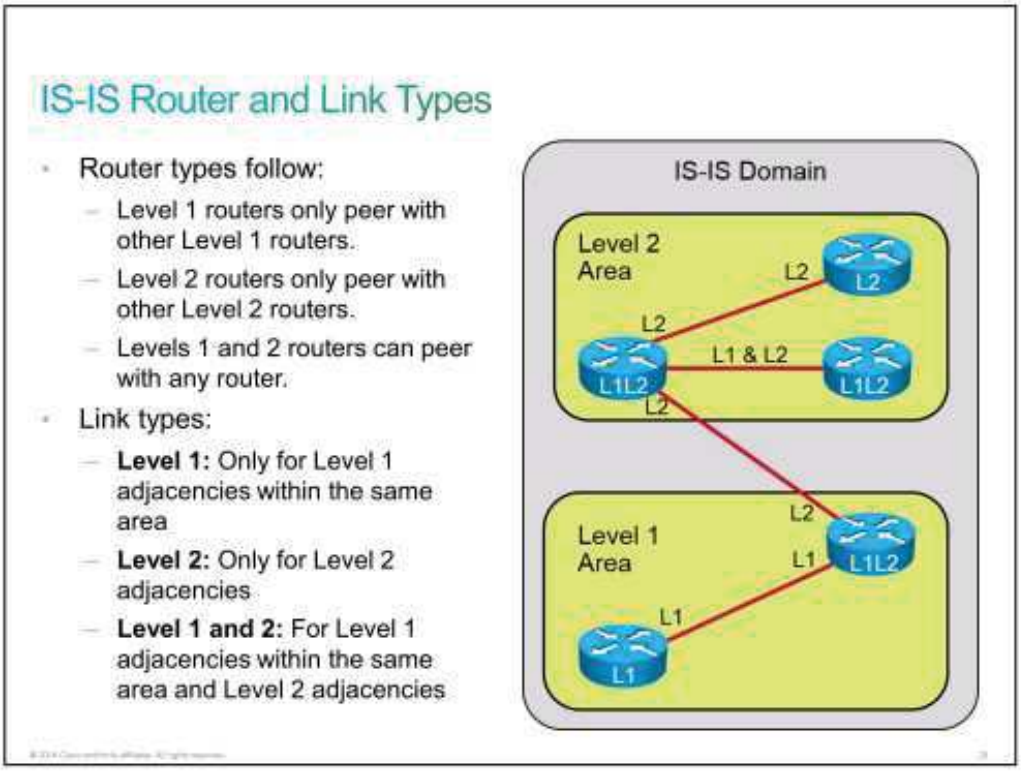

Area

- Level 1 and Level 2 adjacencies can be formed between routers in the same area.
- Only Level 2 adjacencies can be formed between routers in different areas.

https://t.me/learningnets

BGP Overview

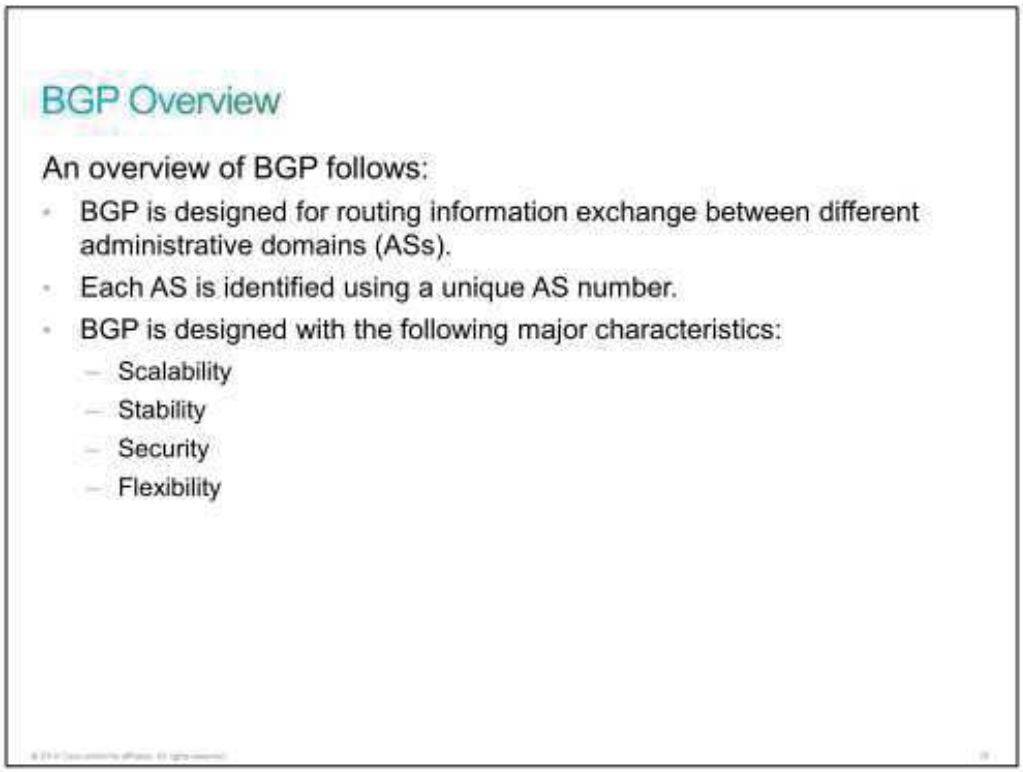

BGP is a distance-vector routing protocol that is designed to meet these criteria:

- **Scalability:** BGP is intended for distributing Internet routing information that constantly grows.
- **Stability:** It is important for BGP to be able to manage constant flapping of routes in an ever-growing Internet, where the likelihood of flapping is also increasing.
- **Security:** Because it is used between an administrative domain and a public environment, BGP must include powerful security mechanisms to protect routers from intrusions from other administrative domains or the Internet in general.
- **Flexibility:** Complex topologies and diverse requirements demand that BGP support advanced mechanisms to implement complex routing policies.

https://t.me/learningnets

BGP Architecture

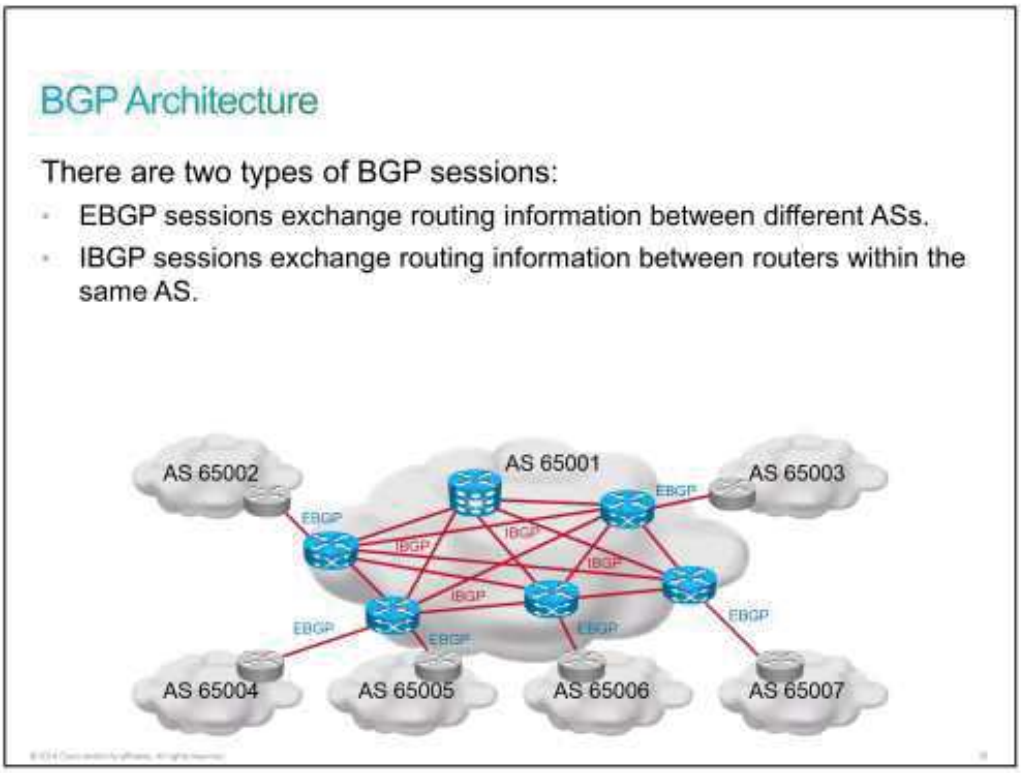

The figure illustrates a general architecture of BGP with these characteristics:

- Each administrative domain is identified using a unique AS number.
- BGP sessions within an AS are called IBGP sessions and differ from EBGP sessions that are used between different ASs.

https://t.me/learningnets

BGP Characteristics

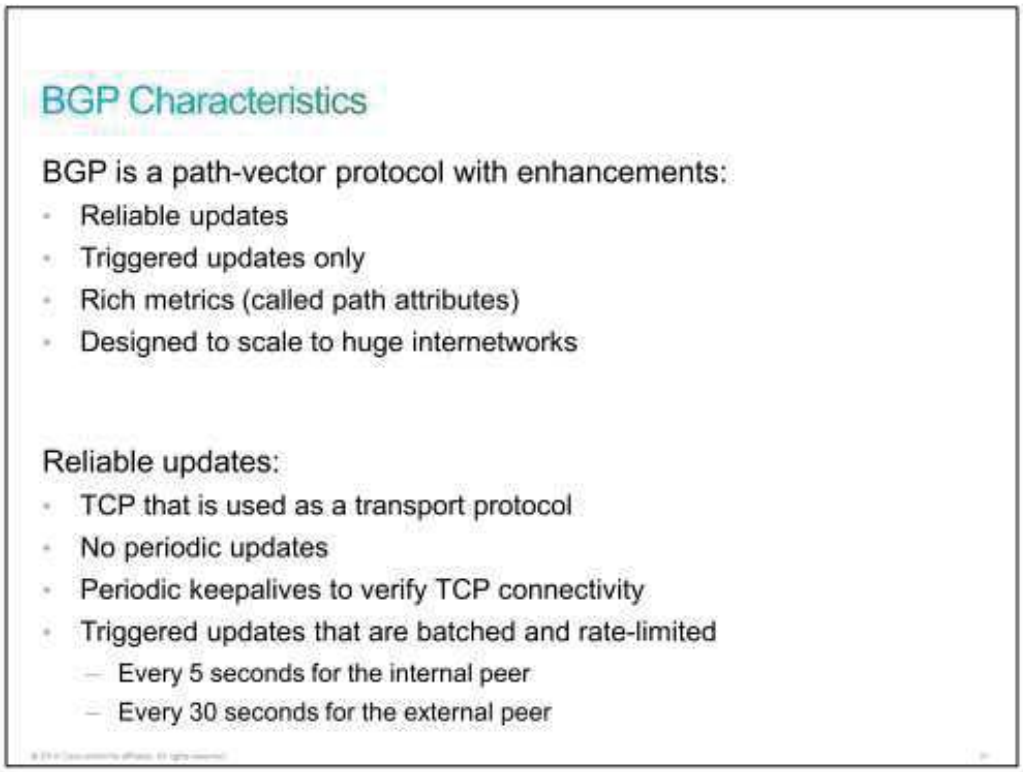

BGP is a path-vector protocol. It is different from distance-vector routing and link-state routing. Each entry in the routing table contains the destination network, the next router, and the path to reach the destination. This means that BGP will announce to its neighbors those IP networks that it can reach itself. The receivers of this information will say, "If that AS can reach those networks, then I can reach them via the AS."

If two different paths are available to reach the same IP subnet, the shortest path is used. This determination requires a mechanism that is capable of measuring the distance. All distance-vector protocols have these mechanisms that are called "metrics." BGP contains a very sophisticated method of computing the shortest path by using attributes that are attached to the reachable IP subnet.

BGP sends routing updates to its neighbors by using a reliable transport. This technique means that the sender of the information always knows that the receiver has actually received the information. As a result, there is no need for periodic updates or routing information refreshes. In BGP, only information that has changed is transmitted.

The reliable information exchange, combined with the batching of routing updates that is also performed by BGP, allows BGP to scale to large, Internet-sized networks.

The reliable transport mechanism that is used by BGP is standard TCP. BGP is an application protocol that uses both the TCP and IP protocols for reliable connections.

Because BGP uses a reliable transport, the sender knows that the receiver has actually received the transmitted information. This capability makes periodic updates unnecessary.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

A router that has received reachability information from a BGP peer must be sure that the peer router is still there. Otherwise, the router could route traffic toward a next-hop router that is no longer available, causing the IP packets to be lost. TCP does not provide a service to signal that the TCP peer has been lost, unless some application data is actually transmitted between the peers. In an idle state, where there is no need for BGP to update its peer, the peer could be unreachable without TCP detecting it. Therefore, BGP takes care of detecting the presence of neighbors by periodically sending small BGP keepalive packets to them. These packets are considered application data by TCP and therefore must be transmitted reliably. According to the BGP specification, the peer router must also reply with a BGP keepalive packet.

When BGP was created, a key design goal was to be able to handle enormous amounts of routing information in large and complex networks. In this environment, many links could go up and down (flapping), causing topology changes that must be considered by the routing protocol. But low convergence time and quick responses to topology changes require fast updates and high CPU power to process both incoming and outgoing updates. The larger the network, the more updates per second that can be expected if an immediate response is required. The presence of too many updates in large networks can jeopardize network scalability.

The designers of BGP decided that scalability was a more important issue than low convergence time, so BGP was designed to batch updates. Any changes that are received within the batch interval time are saved. At the end of the interval, only the remaining result is forwarded in an outgoing update. If a network flaps several times during the batch interval, only the state at the end of the interval is sent in an update. The batching feature avoids an uncontrolled flood of updates all over the Internet because the number of updates is limited by the batching procedure.

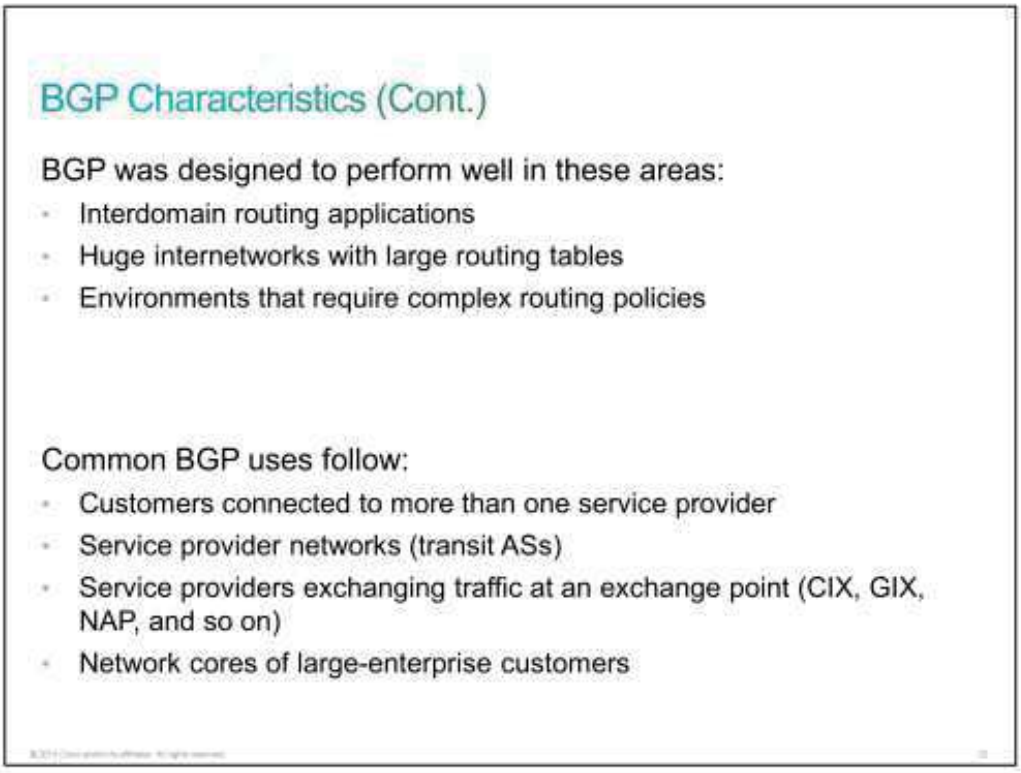

The designers of the BGP protocol succeeded in creating a highly scalable routing protocol that can forward reachability information between ASs (also known as routing domains). The designers had to consider an environment with an enormous number of reachable networks and complex routing policies that were driven by commercial rather than technical considerations.

https://t.me/learningnets

TCP, a well-known and widely proven protocol, was chosen as the transport mechanism. This decision kept BGP simple, but it increased the CPU resource requirements for routers running BGP. The point-to-point nature of TCP also introduces a slight increase in network traffic because any update that should be sent to many receivers has to be multiplied into several copies, which are then transmitted on individual TCP sessions to the receivers.

Whenever there was a design choice between fast convergence and scalability, scalability was the top priority. The batching of updates and the relatively low frequency of keepalive packets are examples of designers placing convergence time second to scalability.

Note	BGP convergence times can be modified with the configuration of nondefault values for BGP scan and advertisement timers.

The figure lists typical scenarios in which BGP is usable. These scenarios include the following:

- Customers connected to more than one service provider.
- ISP networks themselves acting as transit systems and forwarding external traffic.
- Exchange points, which can be defined by the NAP between region and core. International exchange points can be defined by either CIX or GIX points.
- Large enterprises using BGP as their core routing protocol.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP AS Number

An AS number is a number that is assigned by IANA that is used with BGP. The AS number uniquely identifies a network under a single technical administration that has a unique routing policy or is multihomed to the public internet. The AS number is required if you run BGP and peer with your ISP and between ISPs at peering points and Internet exchanges.

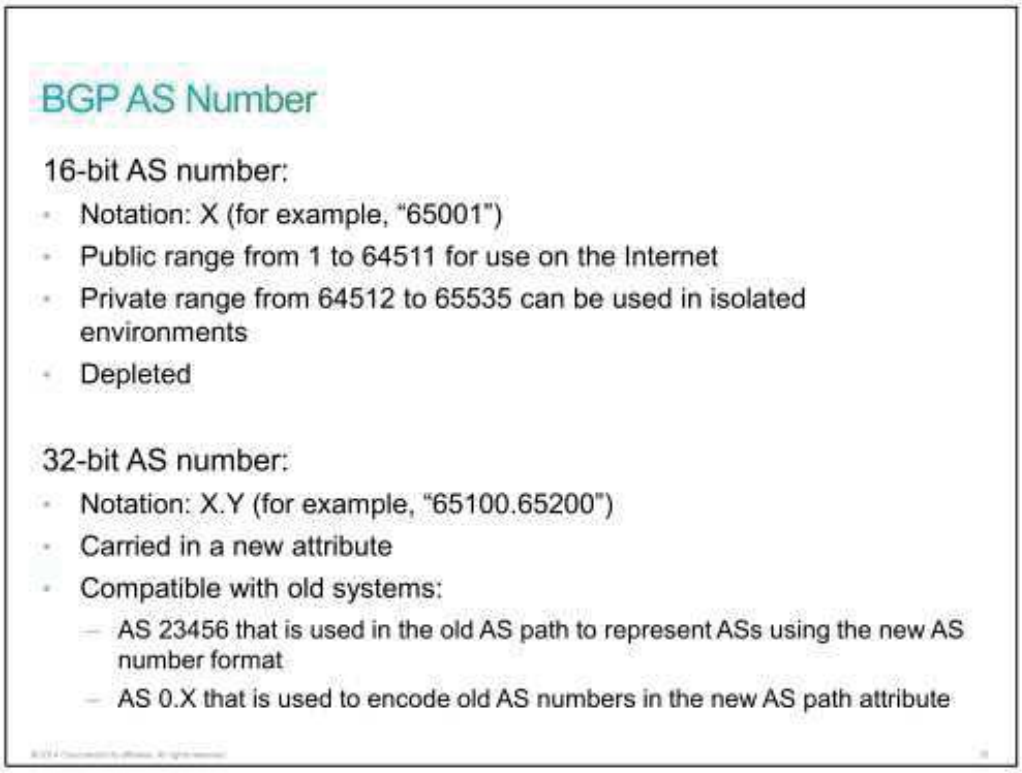

AS numbers come in two forms:

- 16-bit AS numbers have been depleted due to a relatively small number space and the increased demand for companies to be multihomed for increased availability.
- 32-bit AS numbers were introduced to provide more number space while maintaining backward compatibility with 16-bit systems to ease the migration. A 32-bit AS can use two notations: A single 32-bit number (X) or two 16-bit numbers that are joined using a dot (for example, X.Y).

Refer to an online white paper for a detailed explanation of 4-B AS numbers and how they interact with older systems:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6554/ps6599/white_paper_C11_516823.html

https://t.me/learningnets

BGP Sessions

BGP sessions are established over TCP using port number 179. Many session parameters and capabilities are negotiated during session setup by exchanging OPEN messages.

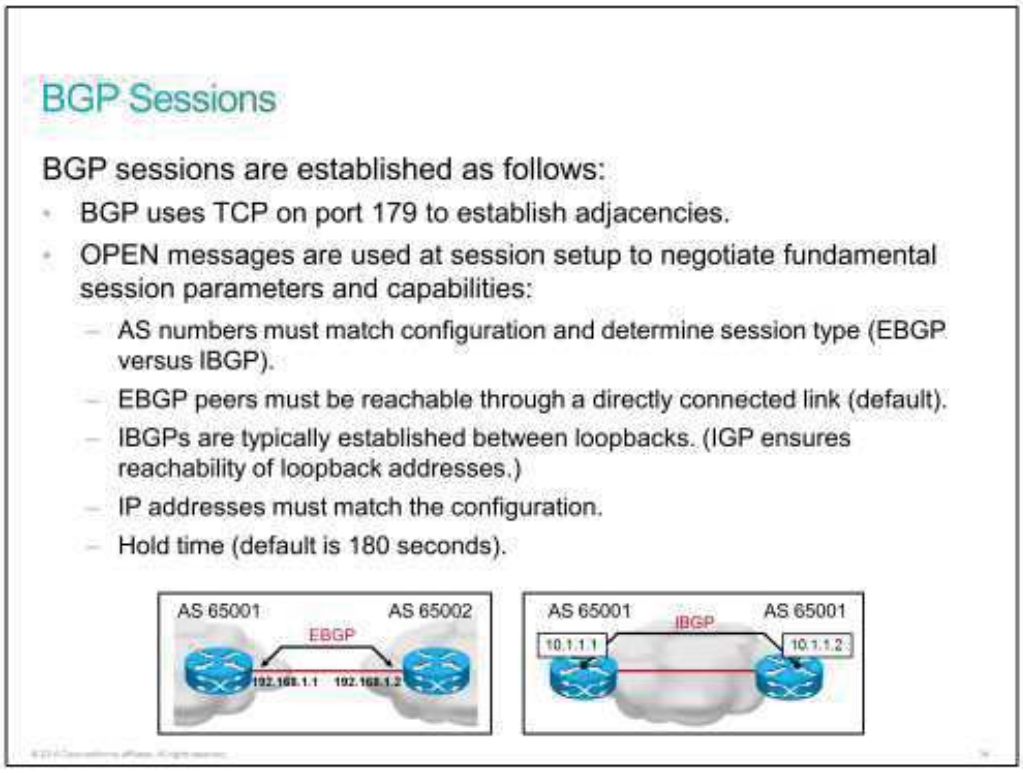

Based on the exchanged AS numbers, both routers will determine if the exchanged numbers match their configuration and select which type the session is (IBGP or EBGP). For EBGP sessions, routers will check to see if the neighbor address is in the routing table as a directly connected address (default requirement). IBGP sessions, on the other hand, can be several hops away, and loopback addresses are typically used to implement IBGP sessions for consistency and stability. The source IP address of a neighbor must also match the configured IP address. Hold-time values are also exchanged, and the lower value is chosen by both routers (keepalive is one-third of hold time).

©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Sessions (Cont.)

EBGP sessions:

- EBGP sessions can form any topology, which is subject to agreements between ASs.
- Received EBGP updates are sent to all other neighbors.
- By default, EBGP neighbors must be directly connected.

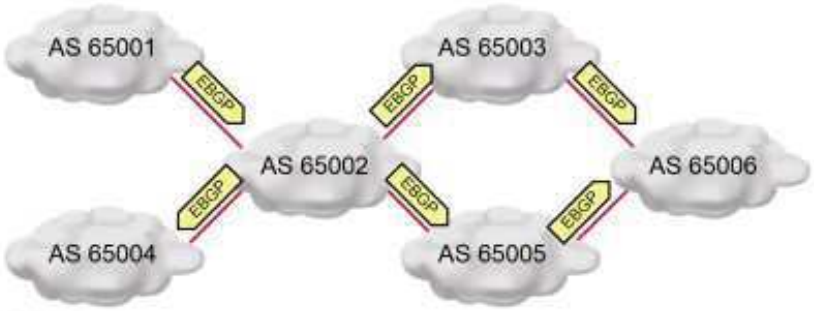

The figure illustrates an arbitrary topology of EBGP sessions between ASs. EBGP has very simple forwarding rules:

- EBGP updates can be sent to all other neighbors (EBGP and IBGP).
- IBGP updates can be sent to EBGP peers.

The AS path is used to prevent updates from looping.

BGP Sessions (Cont.)

IBGP sessions:

- By default, IBGP sessions require a full mesh between all routers within an AS:
  - By default, IBGP updates that are received are not forwarded to other IBGP neighbors.
  - IBGP updates do not scale in large ASs.
- IBGP neighbors can be multiple hops away.

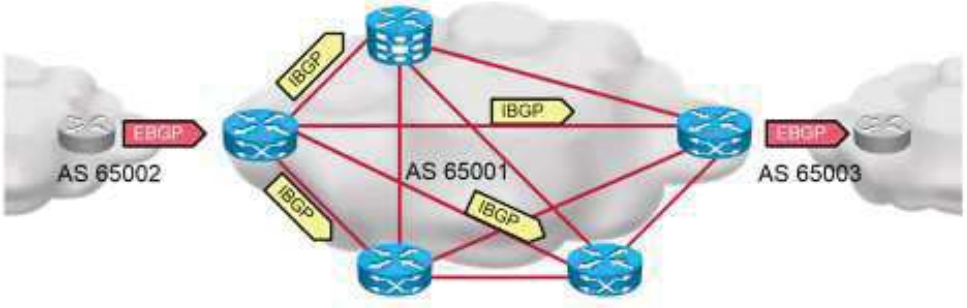

https://t.me/learningnets

IBGP sessions have different forwarding rules, which include a type of split horizon mechanism:

- IBGP updates that are received from a peer can only be forwarded to EBGP peers.
- EBGP updates that are received from a peer can be forwarded to both EBGP and IBGP peers.

In order to ensure that all routers receive all updates, you must configure a full mesh of IBGP sessions between BGP routers in an AS.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP in Customer Connections: Single-Homed Customers

Most single-homed customers (that is, customers that are connected to one ISP using one link) only require static routing because there is no alternative path if the primary path fails (such as a router, link, or ISP failure).

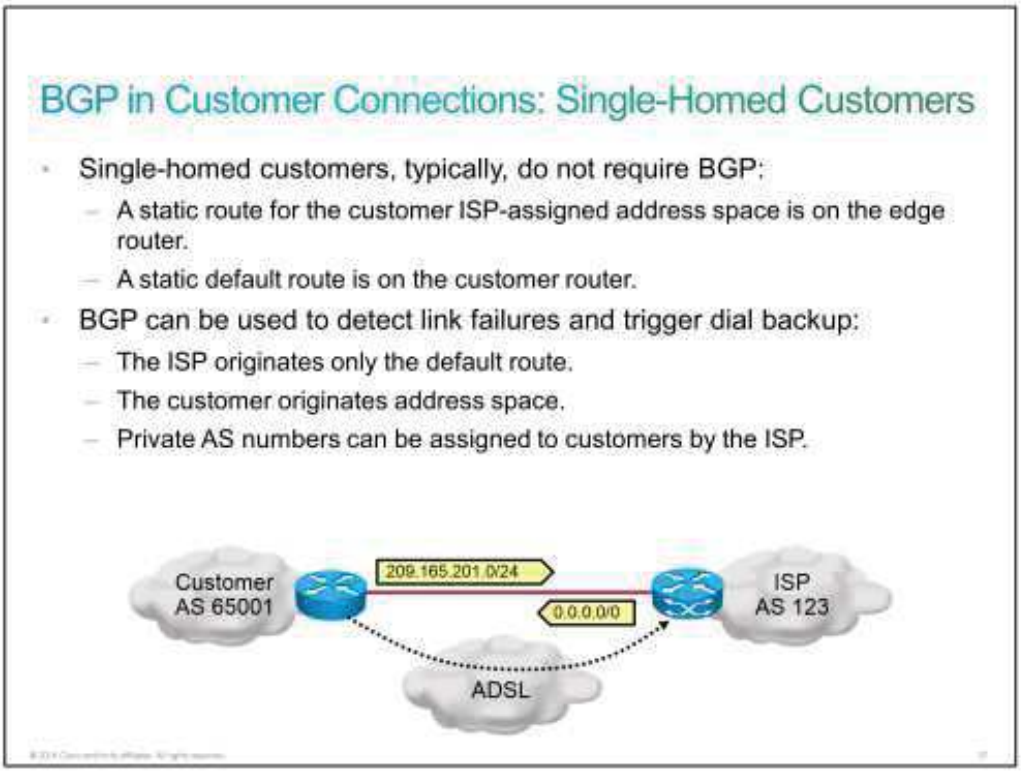

Some single-homed customers may deploy a dial backup solution in which it is beneficial for them to receive notification if their primary path has failed. ISPs can use BGP to send them a default route. If a failure occurs, the BGP session will go down, and the customer can initiate a dial backup connection.

https://t.me/learningnets

BGP in Customer Connections: Dual-Attached Customers

A dual-attached customer (that is, a customer that is connected to the same ISP over two or more links) will preferably require BGP to exchange routing information, enable primary and backup routing or load balancing, and have the ability to detect failed links.

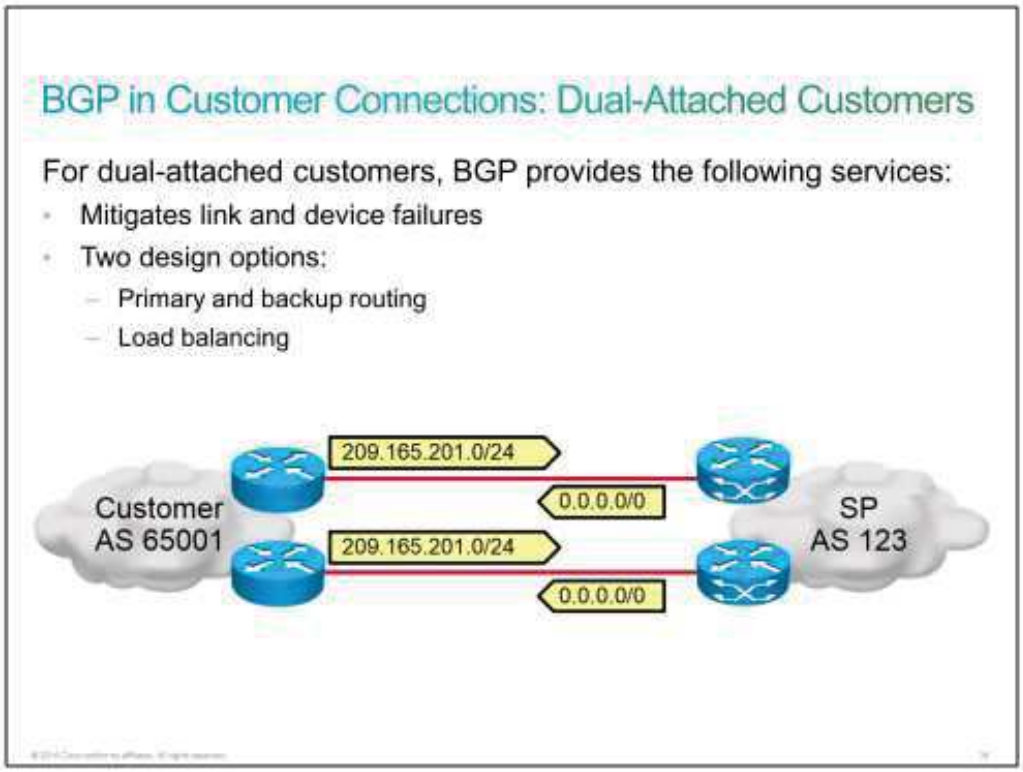

This solution can mitigate router and link failures, but it cannot mitigate ISP failures. Dual-attached customers may use a single router with redundant links toward a service provider.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP in Customer Connections: Multihomed Customers

Multihomed customers (that is, customers that are connected to two or more independent ISPs, using separate links on the same router or on different routers) have the most resilient setup, which can mitigate any single failure of routers, links, or ISPs. These customers will often require complete Internet routing information from all ISPs to give them the most flexibility when implementing load balancing.

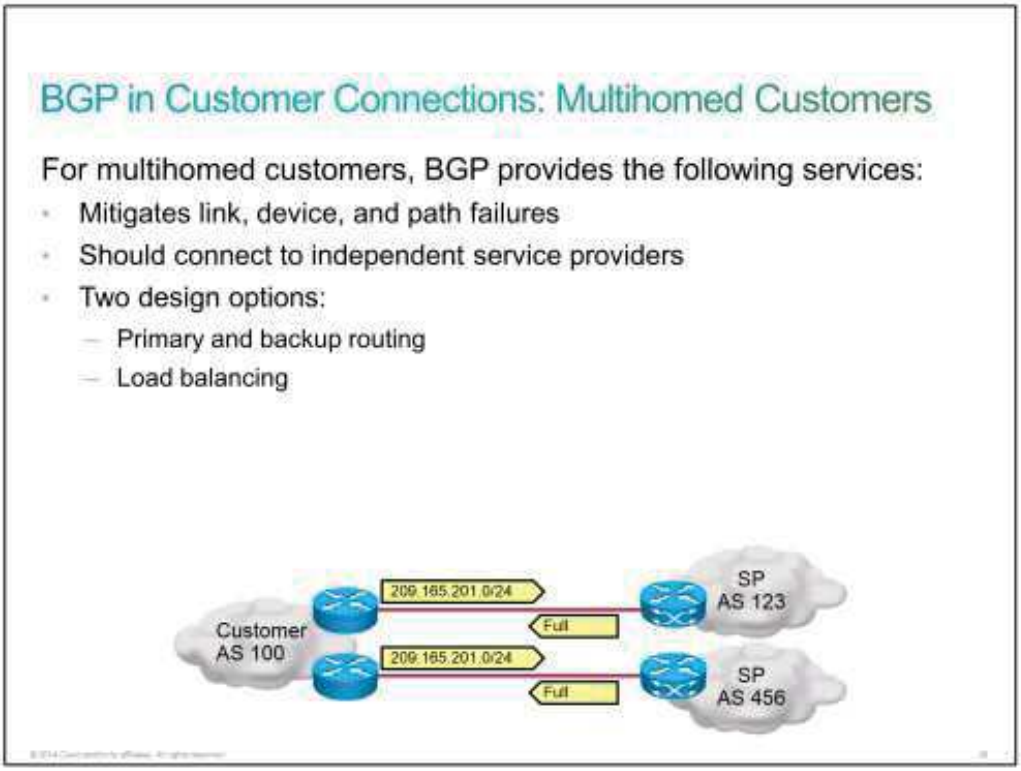

Multihomed customers have the following requirements:

- Public AS number
- Provider-independent address space

https://t.me/learningnets

BGP in Customer Connections: Upstream ISP

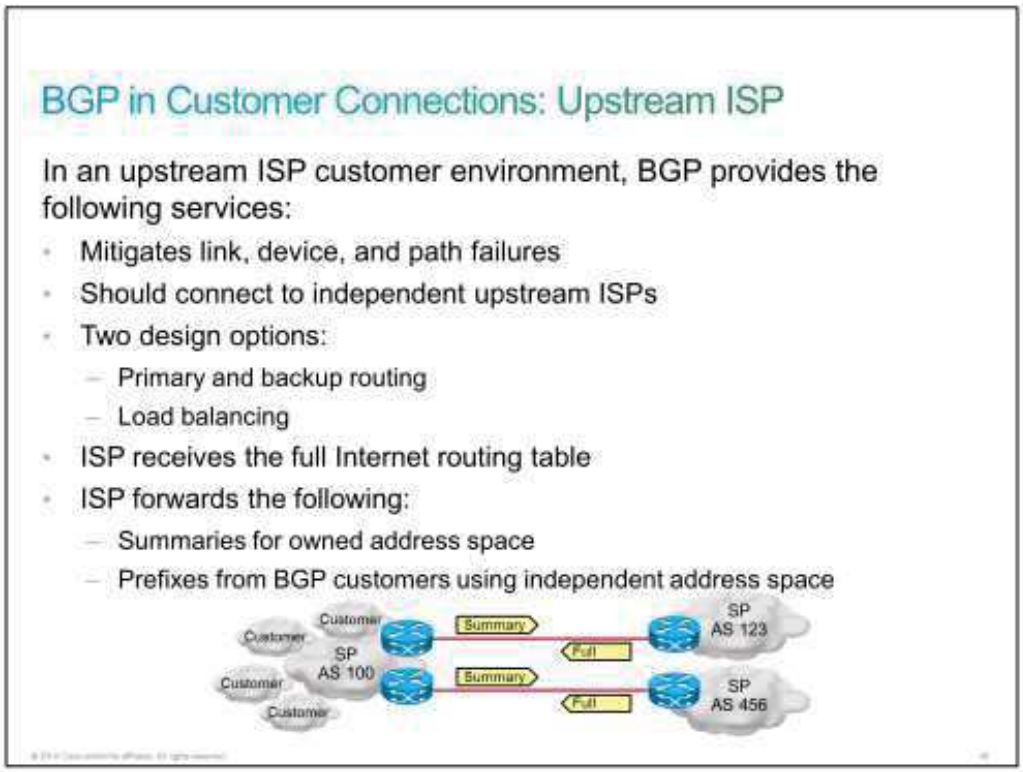

For ISPs, the network is their business, and they will always have multiple links to other ISPs:

- Upstream ISPs, which will provide complete Internet routing information
- Peering ISPs over a local exchange point, which provide routing information for their customers

ISPs will forward a summary for their IP supernet and the individual routes of their BGP-based customers.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP in Customer Connections: Transit ISP

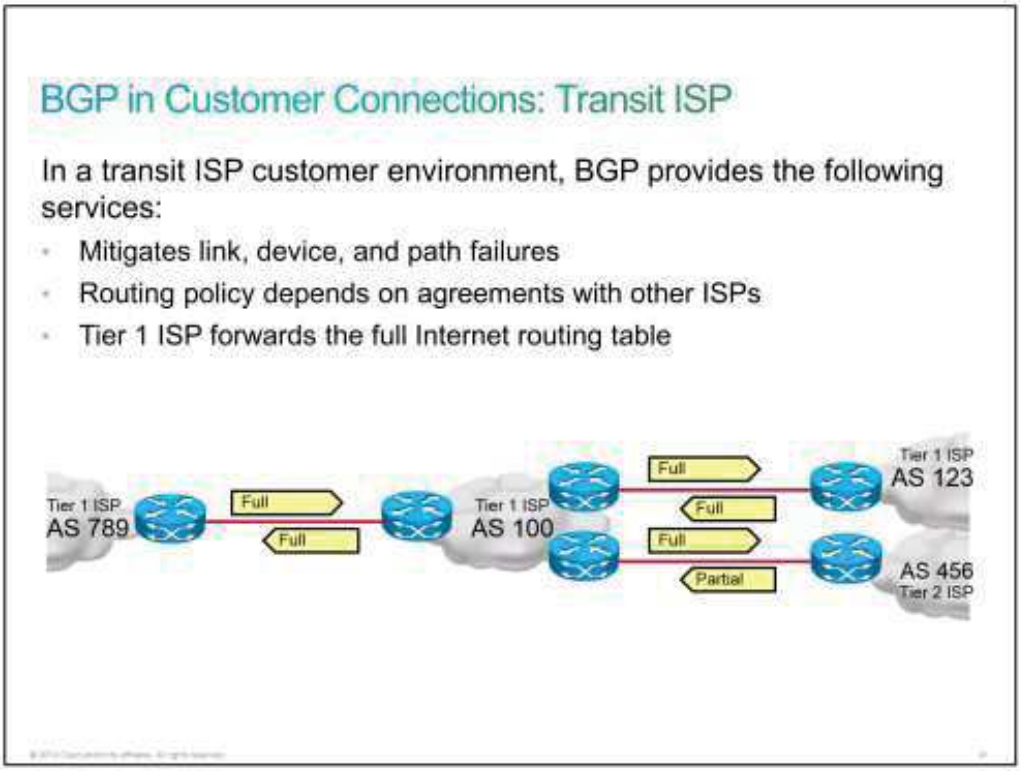

There are many types of transit ISPs, which can be summarized in two major types:

- **Tier 1 ISPs:** Peer with other Tier 1 ISPs to form the backbone of the Internet
- **Tier 2 and Tier 3 ISPs:** Depend on Tier 1 ISPs to reach the rest of the Internet

The relationship between these large ISPs depends on the agreements and charging between the ISPs. The routing policy must be implemented in accordance with the agreements.

In most cases, Tier 1 ISPs will exchange complete Internet routing tables, while Tier 2 and Tier 3 ISPs will only forward the routing information for their subordinate ISPs and end customers.

https://t.me/learningnets

Summary

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

Summary

- Routing protocols are used on the IP infrastructure layer of the Cisco IP NGN.
- BGP is used to exchange external routing information, while IGP is used to exchange local routing information.
- Most service providers use OSPF or IS-IS as IGP.
- IGP propagates BGP next-hop addresses and loopback IP addresses that are used for BGP sessions.
- OPSF uses three data structures: neighbor table, topology table, and routing table.
- The ABR router separates the backbone area from other areas.
- OSPF supports six different types of areas. From a design and implementation perspective, it is preferred to implement OSPF using single-area design.
- OSPF uses link cost as a metric. Link cost is calculated from interface bandwidth.

Summary (Cont.)

- IS-IS is a stable, link-state routing protocol that is often used in service provider environments.
- IS-IS uses hierachical design with Level 1 and Level 2 areas.
- IS-IS is a link-state routing protocol and is similar to OSPF, but it uses CLNS addresses to identify routers and build the LSDB.
- BGP was designed to exchange routing information between autonomous systems.
- BGP uses external and internal BGP sessions to exchange routing information.
- BGP session parameters and capabilities are negotiated during session setup by exchanging OPEN messages.
- BGP is used between service providers and multihomed customers.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Module Summary

Overview

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

This module describes the main characteristics of routing protocols that are used in service provider environments. The module describes how service providers ensure IP connectivity to the Internet, both to end customers and to other service providers. Service providers exchange Internet routing information via BGP, and the OSPF and IS-IS protocols are used for providing IP connectivity within service provider autonomous systems.

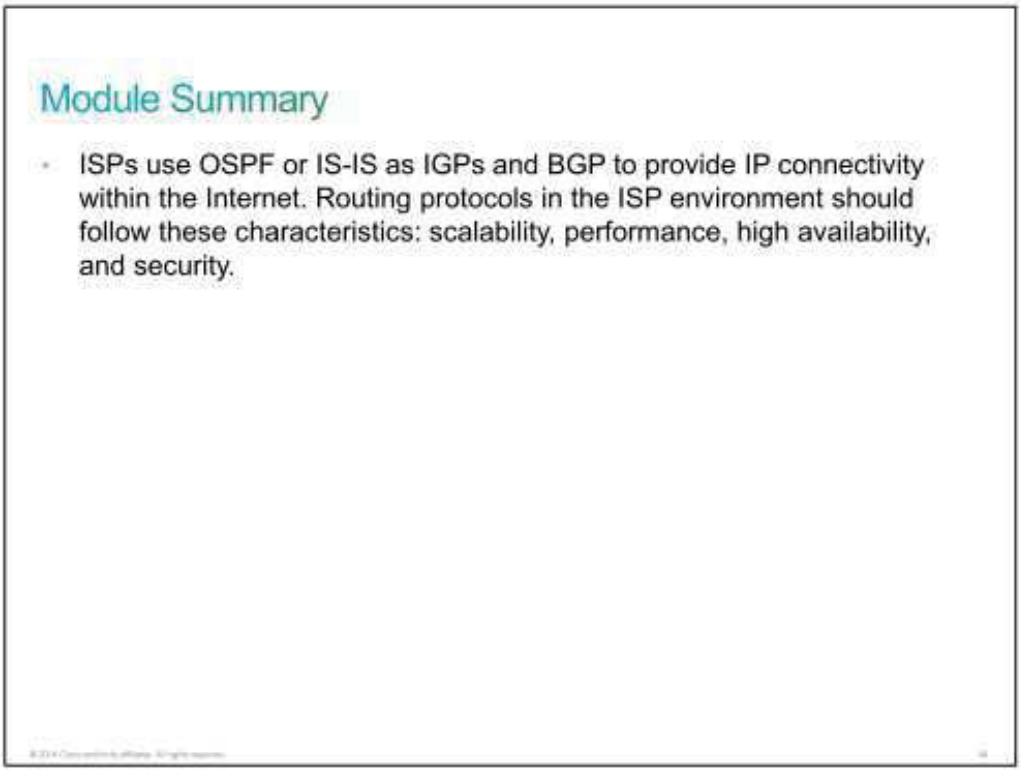

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: 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 a service provider environment, which two statements regarding BGP and IGP are true? (Choose two.) (Source: Understanding Service Provider Routing Protocols)

   A. BGP is used to provide connectivity within an AS.
   B. BGP is used to exchange Internet routing information with other ISPs and customers that require it.
   C. IGP is used to provide reachability of BGP neighbors and BGP next-hop addresses.
   D. IGP is used to exchange external routing information.

2. Which two IGPs are commonly used in service provider backbones? (Choose two.) (Source: Understanding Service Provider Routing Protocols)

   A. RIP
   B. IS-IS
   C. OSPF
   D. EIGRP

3. In OSPF, which area type typically carries all of the routing information? (Source: Understanding Service Provider Routing Protocols)

   A. Stubby area
   B. Totally stubby area
   C. Regular nonbackbone area
   D. Backbone Area 0

https://t.me/learningnets

4. IS-IS and OSPF are distance-vector routing protocols. (Source: Understanding Service Provider Routing Protocols)

   A. True
   B. False

5. Regarding OSPF and IS-IS routing protocols, which two statements about link cost are true? (Choose two.) (Source: Understanding Service Provider Routing Protocols)

   A. IS-IS link cost is calculated based on bandwidth.
   B. IS-IS link cost defaults to 10.
   C. OSPF link cost is calculated based on bandwidth.
   D. OSPF link cost defaults to 100.

6. BGP sessions are established over TCP using port number 179. (Source: Understanding Service Provider Routing Protocols)

   A. True
   B. False

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. B, C
2. B, C
3. D
4. B
5. B, C
6. A

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 2: Implement OSPF in the Service Provider Network

Introduction

This module examines OSPF, which is one of the most commonly used interior gateway protocols in service provider networks. OSPF is an open-standard protocol that is based primarily on RFC 2328, with some enhancements for IPv6 that are based on RFC 2740. OSPF is a complex protocol that is composed of several protocol handshakes, database advertisements, and packet types.

Configuration and verification of OSPF in a Cisco IOS, IOS XE, or IOS XR router is a primary learning objective of this module. The lessons in the module describe steps that are needed to implement OSPF in a service provider network.

Upon completing this module, you will be able to:

- Describe OSPF routing functions and the importance of OSPF in a service provider network
- Describe how OSPF improves packet processes and optimizes routing performance in a service provider network
- Describe implementation steps that are needed to enable OSPF in a service provider network
- Describe OSPF area types and explain the importance of OSPF route summarization in a service provider network

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Introducing OSPF Routing

Overview

OSPF is one of the most commonly used IP routing protocols in networking. It is an open standard that is used by both enterprise and service provider networks.

This lesson introduces each of the major characteristics of the OSPF routing protocol, including a description of link-state routing protocols, the OSPF hierarchical structure, link-state adjacencies, and SPF calculations.

Upon completing this lesson, you will be able to:

- Describe OSPF in the Cisco IP Next Generation Network
- Describe the key characteristics of OSPF and OSPFv3
- Describe how OSPF routes are created
- Describe the data structures that are used by OSPF
- Describe the two-tier hierarchical structure of OSPF, including the characteristics of transit areas and regular areas, and the terminology that is used
- Describe the area hierarchical structure of OSPF in a service provider environment
- Describe OSPF LSA types
- Describe OSPF stub areas
- Describe OSPF NSSAs
- Describe how routers establish OSPF neighbor adjacencies and exchange LSAs
- Describe how OSPF calculates the best path to a network
- Describe the OSPF metric
- Describe how the LSDB is constructed

https://t.me/learningnets

- Describe how LSAs age
- Describe how to interpret content of the OSPF LSDB
- Describe the OSPF LSDB for intra-area routing
- Describe the OSPF LSDB for interarea routing
- Describe the OSPF LSDB for external routes
- Describe the OSPF LSDB for virtual links
- Describe how to interpret the routing table entries for OSPF-learned routes
- Describe how costs are calculated for external routes
- Describe the OSPF LSDB overload protection feature

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF in the Cisco IP NGN Architecture

The OSPF is a popular IP routing protocol in the service provider environment and focuses on routing in the IP infrastructure layer of the Cisco IP Next Generation Network.

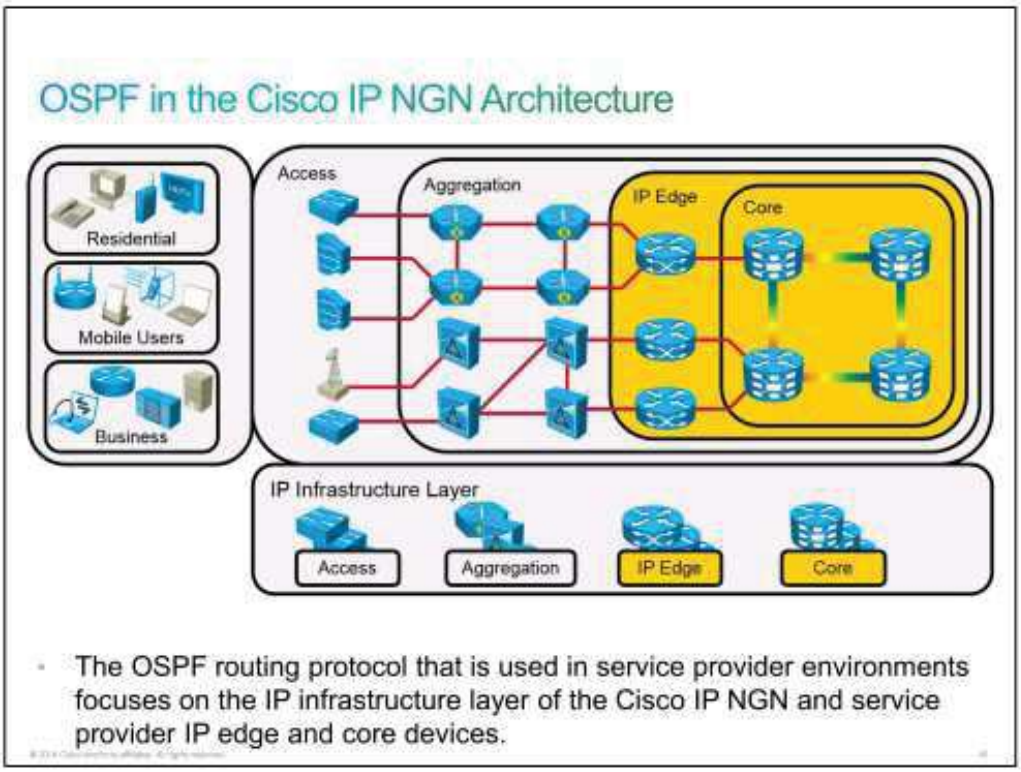

The OSPF protocol can be found on the IP edge and core devices and is responsible for service provider internal routing.

https://t.me/learningnets

OSPF and OSPFv3 Key Characteristics

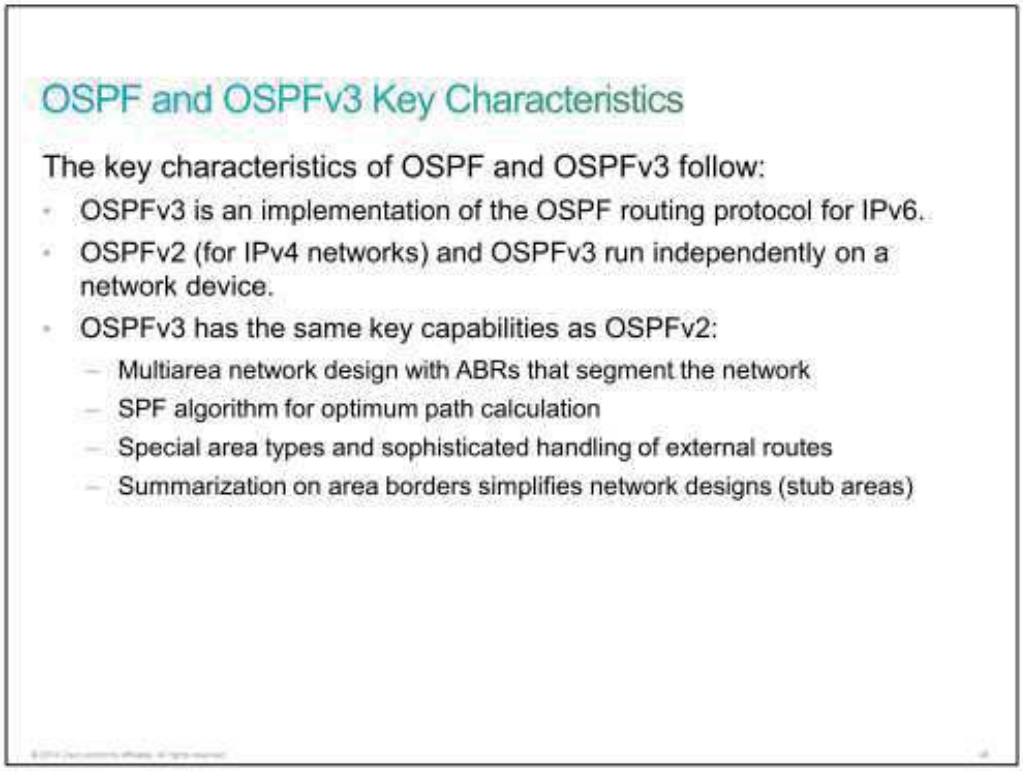

OSPF segments the network into multiple areas that communicate through ABRs. This approach allows greater scalability and relieves the routers from running route calculations for events that are not in their area. Only ABRs need to know the exact topology of all attached areas. These routers send appropriate routes as interarea routes.

OSPF manages external routes differently than internal routes. The routes are propagated across all areas in a special update packet, and they are distinguished in the routing table.

Special area types, such as stub areas and NSSAs, allow for managing external routes and summarization. The core algorithm for best-path calculation is SPF or Dijkstra's algorithm. This algorithm is run every time that there is a topology change in the area. OSPFv3 is a complete rewrite of the OSPF protocol to support IPv6. The foundation remains the same in IPv4 and OSPFv2. OSPFv3 and OSPFv2 run independently on a network device.

The OSPFv3 metric is based on interface costing. The packet types and neighbor discovery mechanisms are the same in OSPFv3 and OSPFv2. OSPFv3 also supports the same interface types, including broadcast, point-to-point, point-to-multipoint, nonbroadcast multiaccess, and virtual links.

https://t.me/learningnets

OSPF Route Entry Creation

The need to overcome the limitations of distance-vector routing protocols led to the development of link-state routing protocols. By using LSAs, each router builds its own view of the network and also maintains a list of neighbors, a list of all routers in the area, and a list of the best paths to each destination. OSPF is classified as a link-state routing protocol because of the manner in which it distributes routing information and calculates routes.

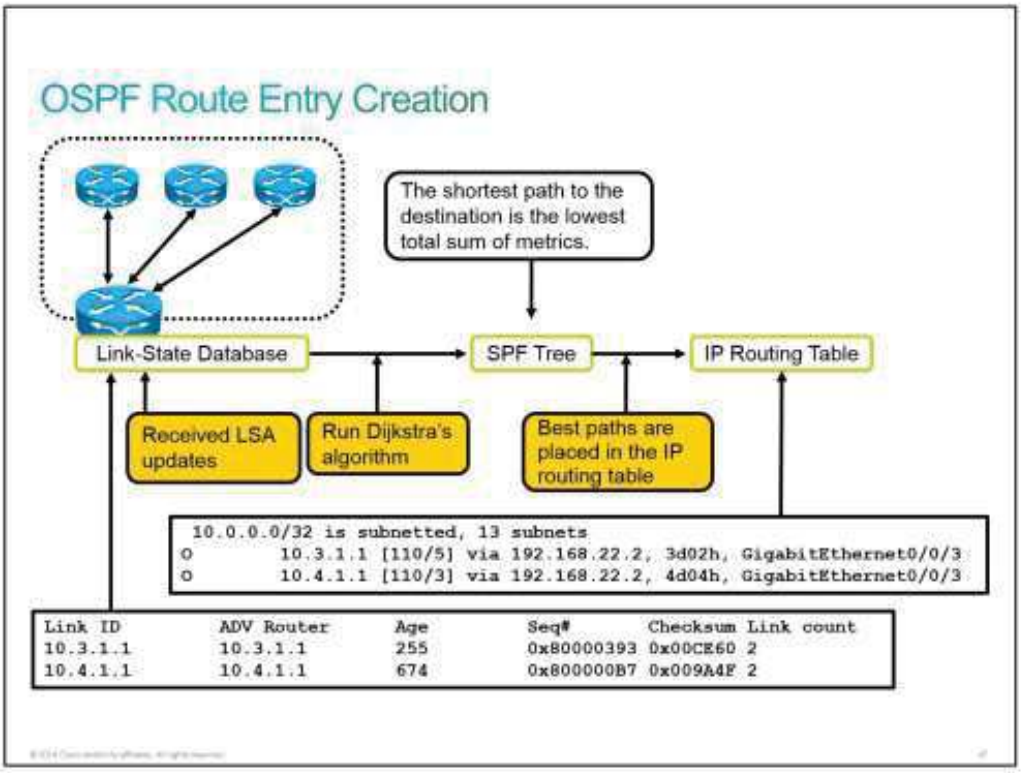

Link-state routing protocols generate routing updates only when a change occurs in the network topology. When a link changes state, the device that detected the change creates an LSA concerning this link.

The LSA is propagated to all neighboring devices using a special multicast address. Each routing device creates a copy of the LSA, updates its LSDB, and forwards the LSA to all neighboring devices within an area. This flooding of the LSA ensures that all routing devices update their databases before updating routing tables to reflect the new topology.

The LSDB is used to calculate the best paths through the network. Link-state routers find the best paths to a destination by applying Dijkstra's algorithm, also known as SPF, against the LSDB to build the SPF tree. The best paths are then selected from the SPF tree and placed in the routing table.

OSPF and IS-IS are examples of link-state routing protocols.

https://t.me/learningnets

OSPF Data Structures

OSPF is a link-state protocol. Routers running OSPF, or link-state routers, recognize more information about the network than their distance-vector counterparts.

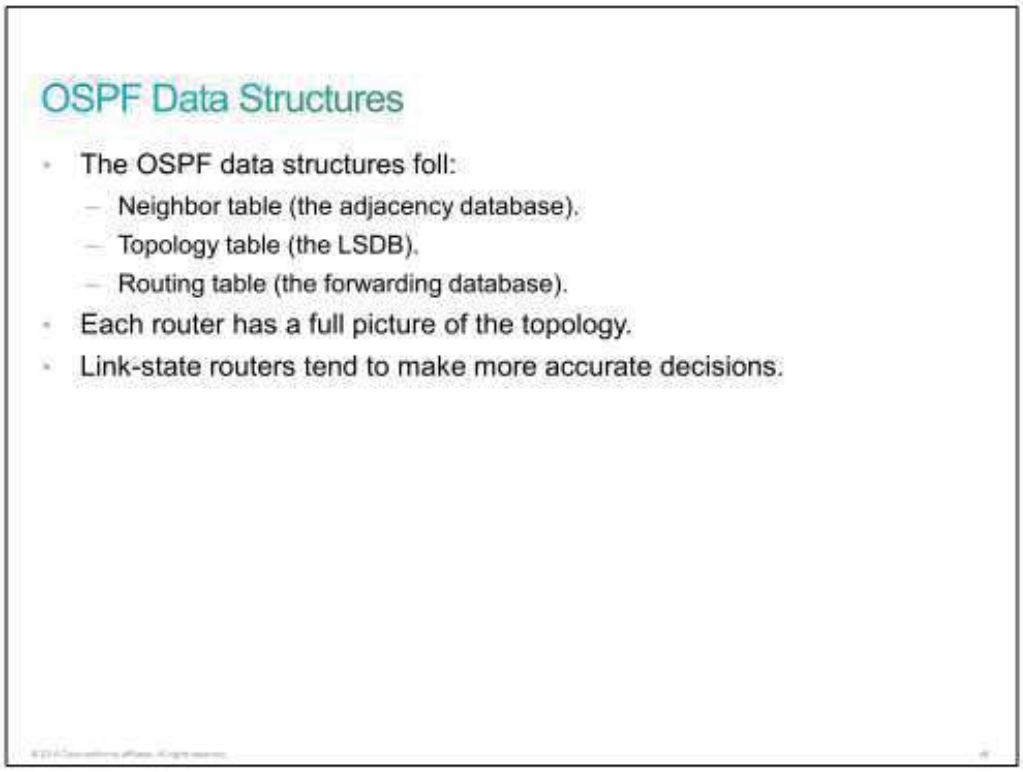

Link-state routing protocols collect routing information from all other routers in the network or from within a defined area of the network. When link-state routing protocols have collected this information from all routers, each router independently calculates its best path to all destinations in the network, using Dijkstra's algorithm.

Incorrect information from any particular router is less likely to cause confusion because each router maintains its own view of the network.

For consistent routing decisions to be taken by all of the routers in the network, each router must keep a record of the following information:

- **Its immediate neighbor routers:** If the router loses contact with a neighboring router, within a few seconds, it will invalidate all paths through this router and recalculate its paths through the network. Adjacency information about neighbors is stored in the neighbor table, also known as an adjacency database, in OSPF.
- **All of the other routers in the network, or in its area of the network, and their attached networks:** The router recognizes other routers and networks through LSAs that are flooded through the network. LSAs are stored in a topology table, also called an LSDB. The LSDB is identical for all OSPF routers in an area. The memory resources that are needed to maintain these tables represent one drawback to link-state protocols.

https://t.me/learningnets

- **The best path to each destination:** Each router independently calculates the best paths to each destination in the network using Dijkstra's algorithm. The best paths are then offered to the routing table or forwarding database. Packets arriving at the router are forwarded based on the information that is held in the routing table. Each router is able to independently select a loop-free and efficient pathway. This benefit overcomes the "routing by rumors" limitation of distance-vector routing.

https://t.me/learningnets

Structure of OSPF Network

In order to understand the structure of an OSPF network, you must be familiar with the two-tier hierarchical structure of OSPF, including the characteristics of transit areas and regular areas, and the terminology that is used.

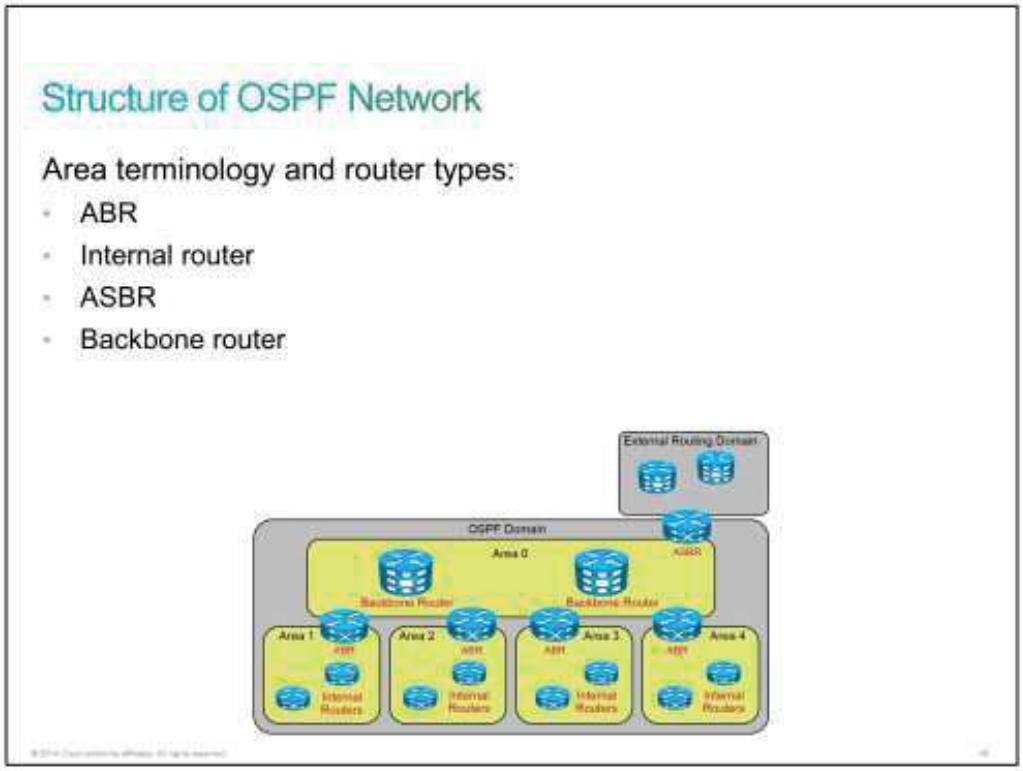

All OSPF areas and routers running the OSPF routing protocol compose the OSPF autonomous system. Routers that make up nonbackbone (normal) areas are known as internal routers and they have all interfaces in one area only. Routers that make up Area 0 are known as backbone routers (internal routers in backbone). OSPF hierarchical networking defines Area 0 as the core. All other areas connect directly to backbone Area 0. An ABR connects Area 0 to the nonbackbone areas. An OSPF ABR plays an important role in network design and has interfaces in more than one area. An ABR has the following characteristics:

- It separates LSA flooding zones.
- It becomes the primary point for area address summarization.
- It functions regularly as the source for default routes.
- It maintains the LSDB for each area with which it is connected.

The ideal design is to have each ABR connected to two areas only, the backbone and another area, with three areas being the upper limit.

An ASBR connects any OSPF area to a different routing administration (such as BGP or EIGRP). The ASBR is the point where external routes can be redistributed into OSPF.

https://t.me/learningnets

To sum up, in order to separate LSA flooding zones, OSPF defines these main router types:

- **Internal router:** This type of router resides inside any area and is a member of one area only.
- **ABR:** This type of router is a member of more than one area. As a border router, it has the ability to control routing traffic from one area to another. Different types of LSAs are exchanged between areas. ABRs can transmit these LSAs, or block them and send default routes instead.
- **ASBR:** This type of router is used to insert external routing information from another non-OSPF autonomous system. ASBR routers generate external LSAs, which can be blocked by ABR routers.

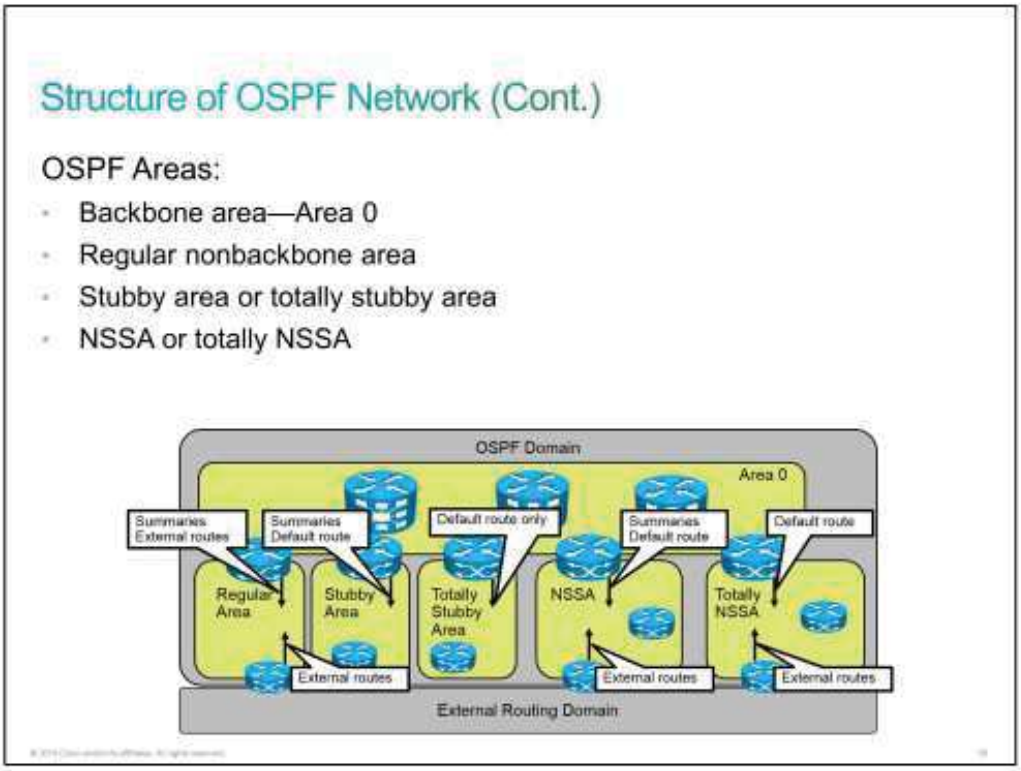

In general, there are six different types of areas:

- Backbone area or Area 0, which typically carries all the routing information.
- Regular nonbackbone areas.
- Stubby areas that do not originate or receive any external routes.
- Totally stubby areas that do not originate (redistribution for other protocols) or receive any external routes or summaries (only the default route).
- NSSAs that can originate external routes (redistribution for other protocols) but do not receive them from other OSPF areas.
- Totally NSSAs that can originate external routes (redistribution for other protocols) but do not receive them from other areas, nor do they receive summaries (only the default route).

https://t.me/learningnets

Hierarchical Structure of OSPF in a Service Provider Environment

One of the most important concepts in OSPF is the existence of hierarchy and areas. OSPF allows collections of contiguous networks to be grouped together. Such a group, together with the routers maintaining interfaces to any of the included networks, is called an area.

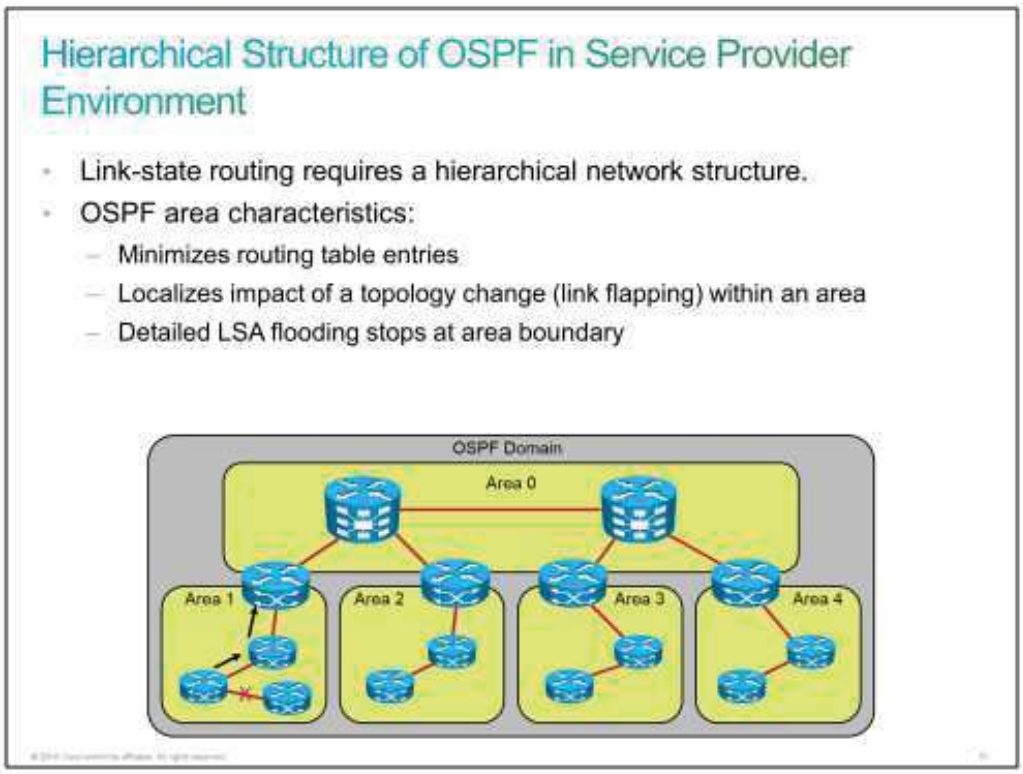

In small networks, the web of router links is not complex and paths to individual destinations are easily deduced. However, in large service provider networks, the web is highly complex and the number of potential paths to each destination is large. Therefore, the Dijkstra calculations that compare all these possible routes can be complex and can take a significant amount of time to complete. All routers must keep a copy of the LSDB—the more OSPF routers, the larger the LSDB. It can be advantageous to have all information in all routers, but this approach does not scale to large network sizes.

OSPF routing protocol reduces the size of the Dijkstra calculations by partitioning the network into areas. The number of routers in an area and the number of LSAs that flood within the area are small, which means that the link-state or topology database for an area is small. Consequently, the Dijkstra calculation is easier and takes less time. Routers inside an area maintain detailed information about the links and only general or summary information about routers and links in other areas.

When a router or link fails, that information is flooded along adjacencies only to the routers in the local area. Routers outside the area do not receive this information. By maintaining a hierarchical structure and limiting the number of routers in an area, an OSPF autonomous system can scale to very large sizes.

OSPF areas require a hierarchical structure, meaning that all areas must connect directly to Area 0, backbone. In the figure, notice that links between Area 1 routers and Area 2, 3, or 4 routers are not allowed.

Traditional design guidelines encouraged usage of OSPF areas. By using multiarea design, you can limit the influence of link flaps (links repeatedly going up and down) and minimize routing table size. You can also take advantage of filtering capabilities on ABRs and contain flooding of LSAs to a small number of routers.

https://t.me/learningnets

On the other hand, modern networks use the opposite approach, called single-area design. You can take advantage of optimal routing because each router sees a complete topology, which is not the case in multiarea design. Currently, unstable links, CPU power, and memory resources are not such concerns as they were in the past. Single-area OSPF design provides support for simplified Cisco MPLS Traffic Engineering implementation.

https://t.me/learningnets

OSPF LSA Types

LSAs are the building blocks of the OSPF LSDB. Individually, they act as database records. In combination, they describe the entire topology of an OSPF network or area.

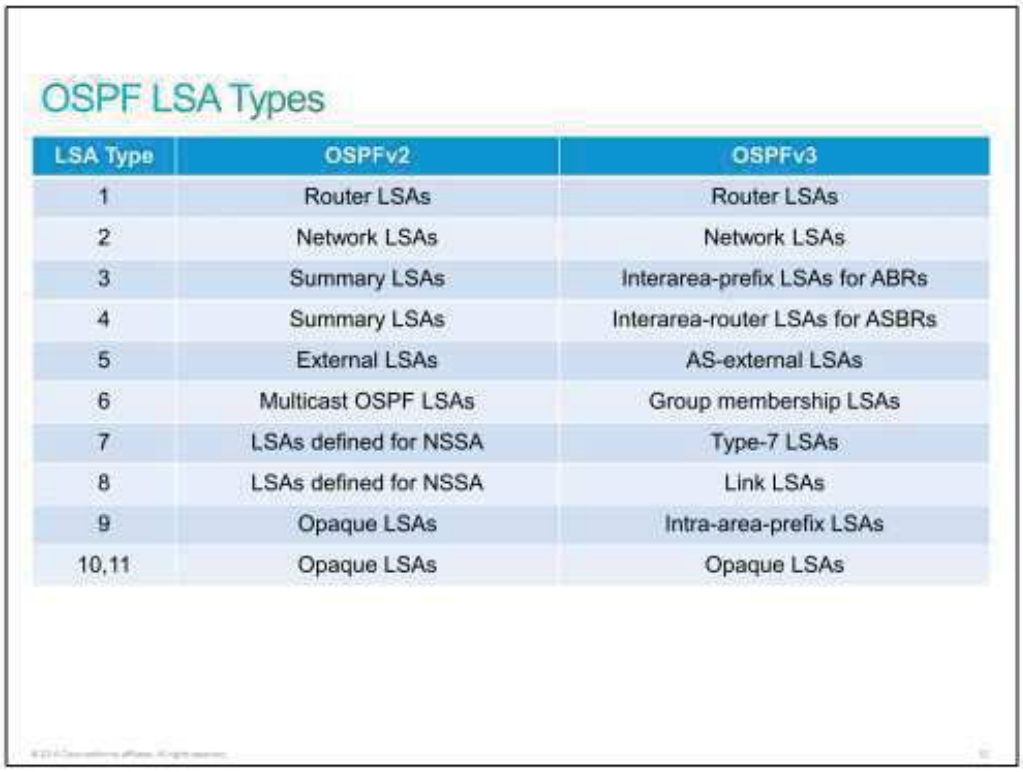

The following are descriptions of each type of LSA. LSA types 1 through 5 and 7 are explained in more detail Most LSA functionality in OSPFv3 is the same as that in OSPFv2, with a few exceptions. In addition, OSPFv3 has two new LSA types to improve OSPF scalability.

Type 1

Every router generates router link advertisements for each area to which it belongs. Router link advertisements describe the state of the router links to the area and are flooded only within that particular area.

In OSPFv3, router interface information may be spread across multiple router LSAs. Receivers must concatenate all router LSAs originated by a given router when running the SPF calculation. In OSPFv3, router LSAs contain no address information. An OSPFv3 router originates a new link LSA for each link to which it is attached. This new link LSA is a type 8 LSA and provides the router link-local address and other addresses on this link.

Type 2

DRs generate network link advertisements for multiaccess networks, which describe the set of routers that are attached to a particular multiaccess network. Network link advertisements are flooded in the area that contains the network.

In OSPFv3, router LSAs and network LSAs do not carry any route information. Route information is carried by the new type 9 intra-area-prefix LSAs.

https://t.me/learningnets

Types 3 and 4

ABRs generate summary link advertisements. Summary link advertisements describe the following interarea routes:

- Type 3 describes routes to networks and aggregates routes.
- Type 4 describes routes to ASBRs.

These LSAs are flooded throughout the backbone area to the other ABRs. The link entries are not flooded into totally stubby areas or NSSAs.

Type 5

ASBRs generate AS external link advertisements. External link advertisements describe routes to destinations that are external to the AS and are flooded everywhere except for stub areas, totally stubby areas, NSSAs, and totally NSSAs.

Type 6

Type 6 LSAs are specialized LSAs that are used in multicast OSPF applications.

Type 7

Type 7 LSAs are used in NSSAs for external routes.

Type 8

Type 8 LSAs are specialized LSAs that are used in internetworking OSPF and BGP. In OSPFv3, link LSAs (type 8) have link-local flooding scope and are never flooded beyond the link with which they are associated. Link LSAs provide the link-local address of the router to all other routers attached to the link, inform other routers that are attached to the link of a list of prefixes to associate with the link, and allow the router to assert a collection of Options bits to associate with the network LSA that will be originated for the link.

Types 9, 10, and 11

The opaque LSAs, types 9, 10, and 11, are designated for future upgrades to OSPF for application-specific purposes. For example, Cisco uses opaque LSAs for MPLS with OSPF. Standard LSDB flooding mechanisms are used for distribution of opaque LSAs. Each of the three types has a different flooding scope.

In OSPFv3, a router can originate multiple intra-area-prefix LSAs (type 9) for each router or transit network, each with a unique link-state ID. The link-state ID for each intra-area-prefix LSA describes its association to either the router LSA or the network LSA and contains prefixes for stub and transit networks.

https://t.me/learningnets

OSPF Stub Areas

To control the advertisement of external routes into an area, OSPF uses stub areas.

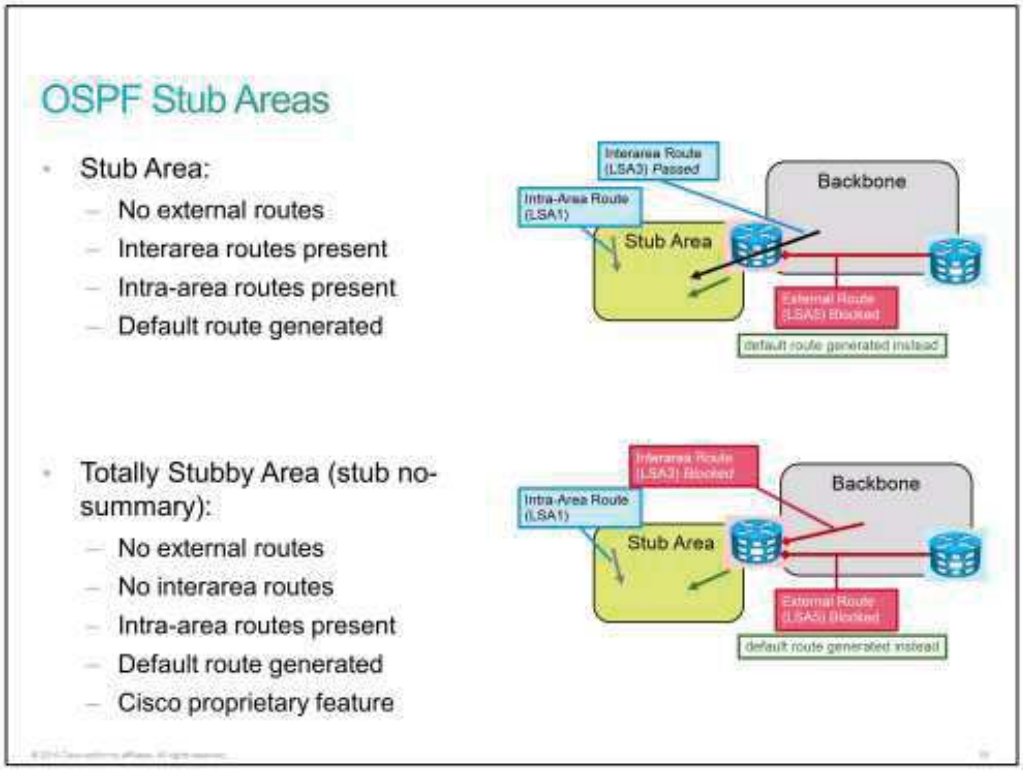

A backbone area is responsible for distributing routing information between multiple areas of an autonomous system. OSPF routing occurring outside of an area is called interarea routing.

The backbone itself has all the properties of an area. It consists of ABRs, routers, and networks only on the backbone. Area 0 is an OSPF backbone area. Any OSPF backbone area has a reserved area ID of 0.0.0.0.

Configuring a stub area reduces the size of the LSDB inside the area, resulting in reduced memory requirements for routers in that area. External network LSAs (type 5), such as those redistributed from other routing protocols into OSPF, are not permitted to flood into a stub area. Routing from these areas to the outside is based on a default route (0.0.0.0/0). If a packet is addressed to a network that is not in the routing table of an internal router, the router automatically forwards the packet to the ABR, which sends a 0.0.0.0/0 LSA. Forwarding the packet to the ABR allows routers within the stub to reduce the size of their routing tables, because a single default route replaces many external routes.

The totally stubby area technique is a Cisco proprietary enhancement that further reduces the number of routes in the routing table. A totally stubby area is a stub area that blocks external type 5 LSAs as well as summary type 3 and type 4 LSAs (interarea routes) from entering the area. Because it blocks these routes, a totally stubby area recognizes only intra-area routes and the default route of 0.0.0.0/0. ABRs inject the default summary link 0.0.0.0/0 into the totally stubby area. Each router picks the closest ABR as a gateway to everything outside the area. Totally stubby areas minimize routing information further than stub areas and increase the stability and scalability of OSPF internetworks. Using totally stubby areas is typically a better solution than using stub areas, as long as the ABR is a Cisco router.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Not-So-Stubby Areas

The OSPF NSSA feature is a nonproprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

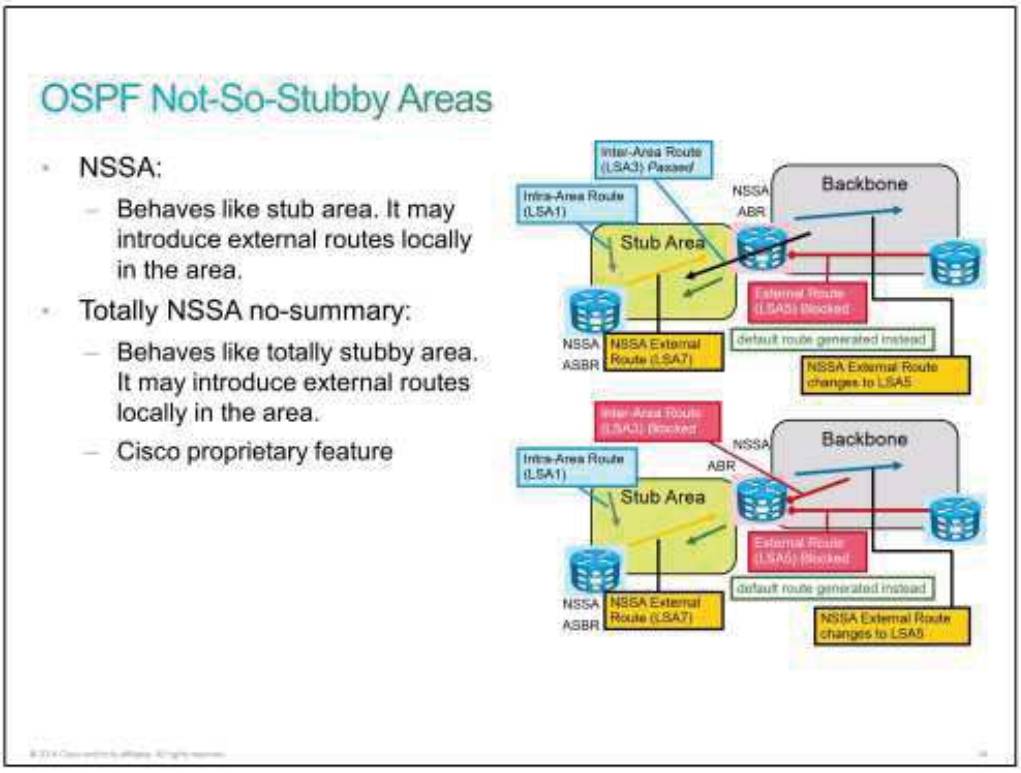

The not-so-stubby areas are characterized as follows:

- **NSSA:** NSSA allows the injection of external routes in a limited fashion into the stub area. Redistribution into an NSSA creates a special type of LSA known as a type 7 LSA, which can exist only in an NSSA. An NSSA ASBR generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain. Type 7 LSAs have a propagate bit in the LSA header to prevent propagation loops between the NSSA and the backbone area. Routes redistributed by an NSSA router internal to the area will have the P bit set. The type 7 LSAs will be converted to type 5 LSAs by the ABRs for the area and propagated to the rest of the network. The P bit is not set only when the NSSA ASBR and NSSA ABR are the same router for the NSSA. Routes redistributed locally on the NSSA ASBR or ABR are injected directly into the backbone as type 5 LSAs and into the NSSA area as type 7 LSAs. To prevent the routes in these type 7 LSAs leaking into the backbone through another ABR for the NSSA area, the P bit is cleared. In this way, the potential for routing loops is reduced. The NSSA retains the other stub area features, and the ABR sends a default route into the NSSA instead of external routes from other ASBRs. The type 7 LSA is described in the routing table as an O N2 or O N1 (N means NSSA). N1 means that the metric is calculated like external type 1. N2 means that the metric is calculated like external type 2. The default is O N2.

Note	The E1 metric adds external and internal costs together to reflect the whole cost to the destination. The E2 metric takes only the external cost, which is reflected in the OSPF cost.

https://t.me/learningnets

- **Totally NSSA:** The OSPF totally NSSA feature is an extension to the NSSA feature like the totally stubby feature is an extension to the stub area feature. It is a Cisco proprietary feature that blocks type 3, 4, and 5 LSAs. A single default route replaces both inbound-external (type 5) LSAs and summary (type 3 and 4) LSAs in the NSSA area totally. The ABRs for the totally NSSA area must be configured to prevent the flooding of summary routes for other areas into the NSSA area. Only ABR routers control the propagation of type 3 LSAs from the backbone.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Operation

A router running a link-state routing protocol must first establish neighbor adjacencies with its neighboring routers. A router achieves this neighbor adjacency by exchanging hello packets with the neighboring routers.

When routers become adjacent, they begin exchanging the link-state information to synchronize the LSDB. Link-state information must be synchronized between routers. Only by reliably flooding link-state information can you ensure that every router in the area or domain has the latest, most accurate view of the network. Only then can the router make reliable routing decisions that are consistent with the decisions of other routers in the network.

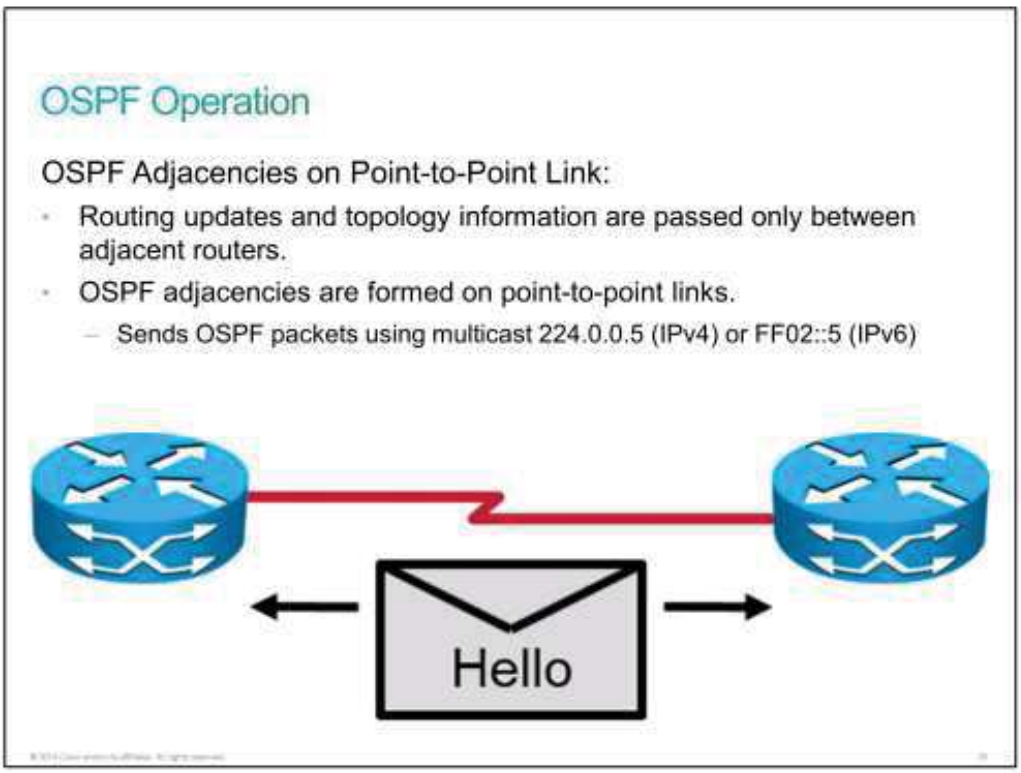

A point-to-point network joins a single pair of routers. A T1 serial line configured with a data link layer protocol, such as PPP or HDLC, is an example of point-to-point network. The two OSPF routers on a point-to-point serial link form a complete adjacency with each other. OSPF routers on the LAN form adjacencies in a different way.

On point-to-point networks, the router dynamically detects its neighboring routers by multicasting its hello packets to all OSPF routers, using the address 224.0.0.5 in OSPFv2 and FF02::5 in OSPFv3. On point-to-point networks, neighboring routers become adjacent whenever they can communicate directly. No DR or BDR election is performed. This is because there can be only two routers on a point-to-point link, so there is no need for a DR or BDR.

The default OSPF hello and dead intervals on point-to-point links are 10 and 40 seconds, respectively.

https://t.me/learningnets

OSPF Operation (Cont.)

OSPF Adjacencies on the LAN Link:

- Forming OSPF adjacencies on LAN links is different than forming them on point-to-point links.
  - Requires DR or BDR election
  - DR and BDR reduce routing update traffic and manage link-state synchronization
  - Sends OSPF packets using multicast 224.0.0.5, 224.0.0.6 (IPv4) or FF02::5, FF02::6 (IPv6)

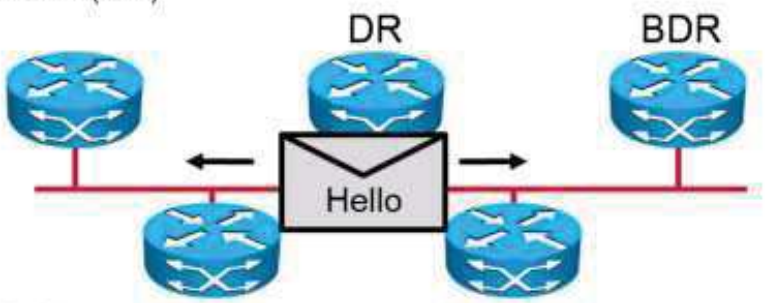

An OSPF router on a multiaccess broadcast network such as Ethernet forms an adjacency with its DR and BDR. A common media segment is the basis for adjacency, such as an Ethernet segment that is connecting two or more routers. When routers first come up on the Ethernet, they perform the hello process and then elect the DR and BDR. The routers then attempt to form adjacencies with the DR and BDR.

The routers on a segment must elect a DR and a BDR to represent the multiaccess broadcast network. The BDR does not perform any DR functions when the DR is operating. Instead, the BDR receives all the information, but the DR performs the LSA forwarding and LSDB synchronization tasks. The BDR performs the DR tasks only if the DR fails. If the DR fails, the BDR automatically becomes the DR, and a new BDR election occurs.

The default OSPF hello and dead intervals on LAN links are 30 and 120 seconds, respectively.

After a DR and BDR have been selected, any router that is added to the network establishes adjacencies with the DR and BDR only. On LAN links, the router dynamically detects its neighboring routers by multicasting its hello packets to all OSPF routers, using the address 224.0.0.5 in OSPFv2 and FF02::5 in OSPFv3. All other routers (not DR nor BDR) are using multicast addresses 224.0.0.6 in OSPFv2 and FF02::6 in OSPFv3 when sending OSPF packets to the DR.

The DR and BDR improve network functioning in the following ways:

- **Reducing routing update traffic:** The DR and BDR act as a central point of contact for link-state information exchange on a multiaccess broadcast network. Therefore, each router must establish a complete adjacency with the DR and the BDR only. Each router, rather than exchanging link-state information with every other router on the segment, sends the link-state information to the DR and BDR only. The DR represents the multiaccess broadcast network in the sense that it sends link-state information from each router to all other routers in the network. This flooding process significantly reduces the router-related traffic on a segment.
- **Managing link-state synchronization:** The DR and BDR ensure that the other routers on the network have the same link-state information about the internetwork. In this way, the DR and BDR reduce the number of routing errors.

https://t.me/learningnets

OSPF Best Path Calculation

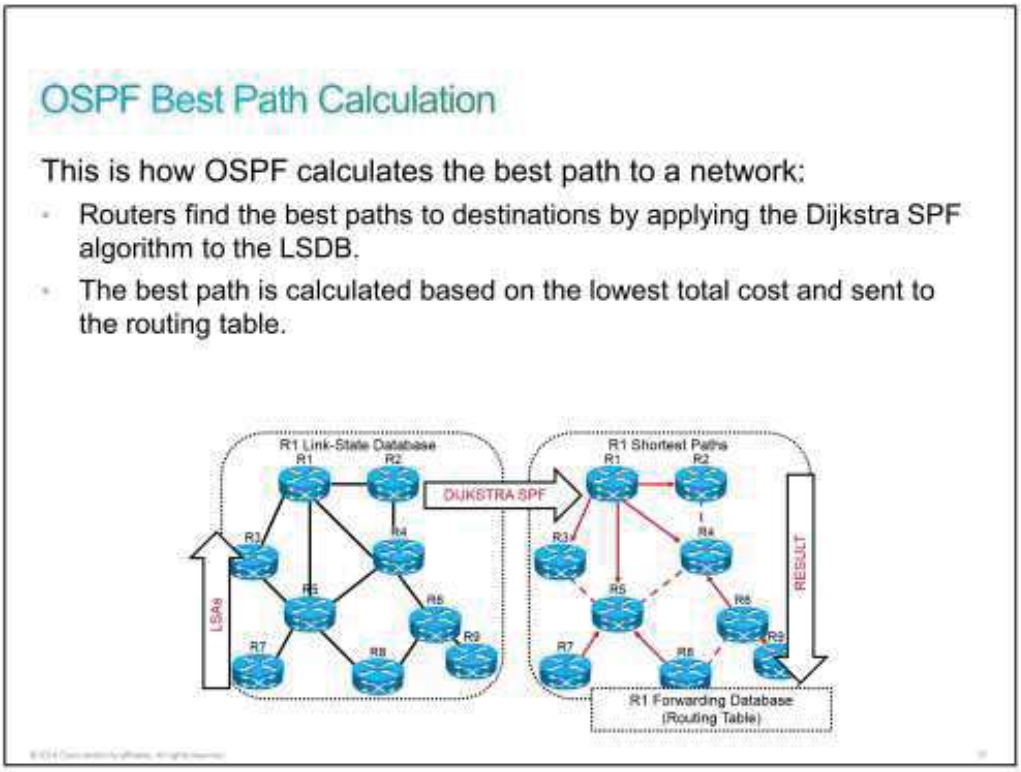

All routers form an adjacency, which is the prerequisite for the LSDB creation process. When they have all formed adjacencies, they start exchanging LSAs. Router R1 has four neighboring routers: R2, R3, R4, and R5. From these routers, R1 receives the LSAs from all other routers in the network. From these LSAs, it can also deduce the links between all routers and draw the web of routers that are depicted in the figure. In this way, an LSDB is created.

Edsger Dijkstra designed a mathematical algorithm for calculating the best paths through complex networks. Link-state routing protocols use Dijkstra's algorithm to calculate the best paths through a network. They assign a cost to each link in the network and place the specific node at the root of a tree, then sum the costs toward each given destination. In this way, they calculate the branches of the tree to determine the best path to each destination. The best path is calculated with respect to the lowest total cost of links to a specific destination and is put in the forwarding database (routing table). For OSPF, the default behavior is that the interface cost is calculated based on its configured bandwidth. You can also manually define an OSPF cost for each interface, which overrides the default cost value.

The figure illustrates an example of a Dijkstra calculation. Each Ethernet link in the figure is assigned an OSPF cost of 10. By summing the costs to each destination, the router can deduce the best path to each destination. The right side of the figure shows the result of the Dijkstra calculation, in which the SPF tree is defining the best paths. From these best paths, which are shown with solid lines, routes to destination networks that are attached to each router are offered to the routing table. For each route, the next-hop address is the appropriate neighboring router (R2, R3, R4, or R5).

https://t.me/learningnets

OSPF Metric

The cost (also called the metric) of an interface in OSPF is an indication of the overhead that is required to send packets across a certain interface. The cost of an interface is inversely proportional to the bandwidth of that interface. A higher bandwidth indicates a lower cost. There is more overhead (higher cost) and time delay in crossing a 56-kb/s serial line than in crossing a 10-Mbps Ethernet line. The formula that is used to calculate the cost is as follows:

Cost = 100,000,000 / bandwidth, in bits per second

For example, it will cost $10^8/10^7 = 10$ to cross a 10-Mb Ethernet line and will cost $10^8/1544000 = 64$ to cross a T1 line. The default reference bandwidth for OSPF is 10^8 bps, or 100 Mb. Increasing the reference bandwidth allows a more granular OSPF design. If changed, it should be changed on all routers in the OSPF domain. The reason that you would change the reference bandwidth is that you may have a link that is faster than 100 Mb in your network. If you have Gigabit networks but are using the default reference bandwidth, then Gigabit links are equal in cost to Fast Ethernet.

By default, the cost of an interface is calculated based on the bandwidth and can be changed by using the OSPF configuration command. If you change the link bandwidth, it will also indirectly change the cost. Only one cost can be assigned per interface, and it is advertised as the link cost in the router link advertisements.

OSPF Metric

These are the main characteristics of the OSPF metric:

- Also called "cost"; COST = 100,000,000 / bandwidth [bps]
- Defined per interface, but may be altered
- Inversely proportional to the bandwidth of that interface

Link Type	Default Cost	Cost with reference 1010 bps
64 kbps serial link	1562	156250
T1 (1.544 Mbps serial link)	64	6476
E1 (2.048 Mbps serial link)	48	4882
Ethernet	10	1000
Fast Ethernet	1	100
Gigabit Ethernet	1	10
10 Gigabit Ethernet	1	1

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Links have different default costs, as follows:

- **A 56-kbps serial link:** The default cost is 1785.
- **A 64-kbps serial link:** The default cost is 1562.
- **T1 (1.544-Mbps serial link):** The default cost is 64.
- **E1 (2.048-Mbps serial link):** The default cost is 48.
- **Ethernet:** The default cost is 10.
- **Fast Ethernet:** The default cost is 1.

The table in the figure also shows the cost of the different link types if the reference bandwidth is changed to 10^{10} bps.

https://t.me/learningnets

Building the LSDB

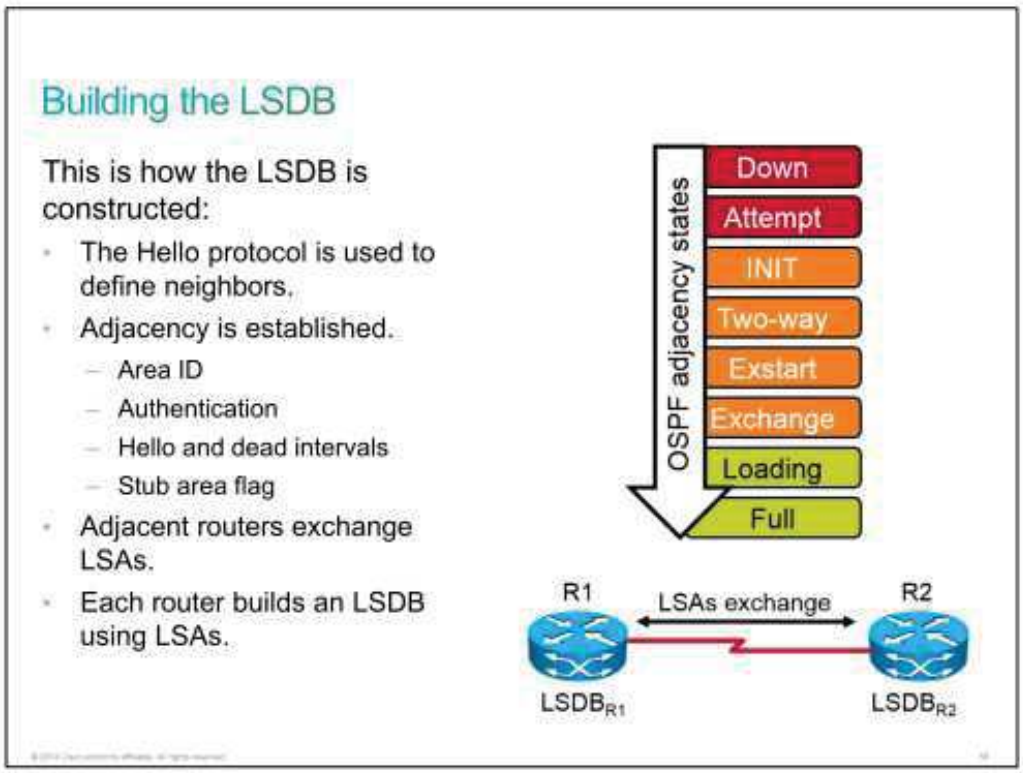

Routers that share a common segment become neighbors on that segment. Neighbors are elected via the Hello protocol. Hello packets are sent periodically out of each interface. Routers become neighbors as soon as they see themselves listed in the neighbor hello packet. Therefore, a two-way communication is guaranteed. Two routers will not become neighbors unless they agree on the following:

- **Area ID:** Two routers having a common segment—their interfaces have to belong to the same area on that segment.
- **Authentication:** OSPF allows for the configuration of a password for a specific area. Routers that want to become neighbors have to exchange the same password on a particular segment.
- **Hello and dead intervals:** OSPF exchanges hello packets on each segment. This is a form of keepalive that is used by routers to acknowledge their existence on a segment and to elect a DR on multiaccess segments. OSPF requires these intervals to be exactly the same between two neighbors.
- **Stub area flag:** Two routers have to agree on the stub area flag in the hello packets to become neighbors. Keep in mind that defining stub areas will affect the neighbor election process.

Adjacency is the next step after the process of becoming neighbors. Adjacent routers are routers that go beyond the simple hello exchange and proceed into the database exchange process. LSU packets are exchanged and acknowledgement is required. An LSU carries one or more LSAs inside it. Selection and processing of LSUs is based on the sequence numbers. At the end, the adjacent routers will have the same link-state database.

https://t.me/learningnets

The following is a brief summary of the states that an interface passes through before becoming adjacent to another router:

- Down
- Attempt
- INIT
- Two-way
- Exstart
- Exchange
- Loading
- Full

https://t.me/learningnets

LSA Operation

Each LSA entry has its own aging timer, LS age, which gives the time, in seconds, since the LSA was originated. The maximum age of the LSA is 3600 seconds, and the refresh time is 1800 seconds. If the LS age reaches 3600 seconds, the LSA must be removed from the database.

After an LSA entry reaches the refresh time, the router that originated the entry sends the LSA (with a higher sequence number in an LSU) to verify that the link is still active. The LSU can contain one or more LSAs. This LSA validation method saves on bandwidth compared with distance vector routers, each of which sends its entire routing table at short intervals.

When each router receives the LSU, it does the following:

- If the LSA does not already exist, the router adds the entry to its LSDB, sends back a LSAck, floods the information to other routers, runs SPF, and updates its routing table.
- If the entry already exists and the received LSA has the same sequence number, the router ignores the LSA entry.
- If the entry already exists but the LSA includes newer information (it has a higher sequence number), the router adds the entry to its LSDB, sends back an LSAck, floods the information to other routers, runs SPF, and updates its routing table.
- If the entry already exists but the LSA includes older information, it sends an LSU to the sender with its newer information.

A combination of the maximum age (max-age) and refresh timers, as well as link-state sequence numbers, helps OSPF maintain a database of only the most recent link-state records.

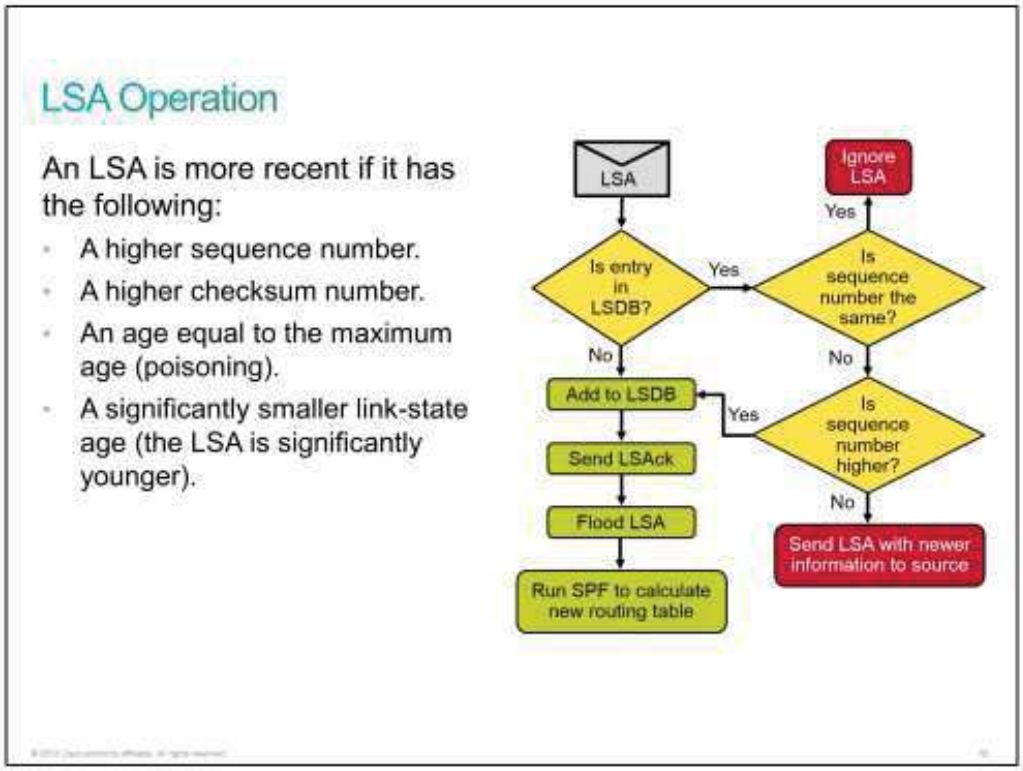

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

LSA Operation (Cont.)

- Each LSA in the LSDB maintains a sequence number:
  - 4-byte number
  - Begins with 0x80000001; ends with 0x7FFFFFFF
- OSPF floods each LSA every 30 minutes.
  - Each time, the sequence number is incremented by one.
  - The LSA with the higher (newer) sequence number is more recent.
  - Ultimately, a sequence number will wrap around to 0x80000001.
  - The existing LSA was prematurely aged to the maximum age (one hour) and flushed:

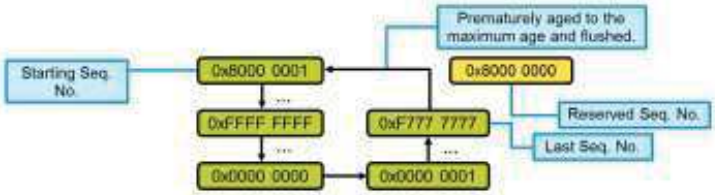

The link-state sequence number field in an LSA header is 32 bits long (4 bytes). Beginning with the left-most bit set, the first legal sequence number is 0x80000001 and the last one is 0x7FFFFFFF. The number is incremented in each new LSA until it reaches 0xFFFFFFFF. The next number is then 0x00000000, followed by 0x00000001, 0x00000002, and so on. 0x8000000 is reserved and never used. The sequence number field is used to detect old or redundant LSAs.

To ensure an accurate database, OSPF floods (refreshes) each LSA every 30 minutes. Each time that a record is flooded, the sequence number is incremented by one. An LSA record will reset its maximum age when it receives a new LSA update. An LSA will never remain in the database longer than the maximum age of 1 hour without a refresh.

It is possible for an LSA to remain in the database for long periods of time, getting refreshed every 30 minutes. At some point, the sequence number will need to wrap back to the starting sequence number. When this occurs, the existing LSA will be prematurely aged out (the maximum age timer is immediately set to 1 hour) and flushed. The LSA will then begin its sequencing at 0x80000001 again.

When a router encounters two instances of an LSA, it must determine which is more recent. The LSA with the higher link-state sequence number is the more recent LSA.

https://t.me/learningnets

LSA Sequence Numbers and Maximum Age

- Every OSPF router announces a router LSA for those interfaces that it owns in that area.
- Router with link ID 10.1.10.1 has been updated three times; the last update was 115 seconds ago.

```
RP/0/RSP0/CPU0:PE1#show ospf database
            OSPF Router with ID (10.1.1.1) (Process ID 1)

                Router Link States (Area 1)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        114         0x80000001 0x0049c1 1

                Router Link States (Area 11)

Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        114         0x80000002 0x004322 1
10.1.10.1       10.1.10.1       115         0x80000003 0x004ddc 2
```

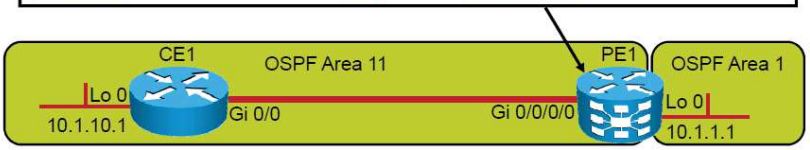

© 2012 Cisco and/or its affiliates. All rights reserved.

The Cisco IOS XR **show ospf database** or Cisco IOS and IOS XE **show ip ospf database** commands display lists of information that are related to the OSPF database for a specific router. The output of the command that is shown in the figure provides an example of how the link-state age and LSA sequence numbers are kept in the database.

Every OSPF router has interfaces in one or more areas and announces a router LSA for those interfaces that it owns in those areas. The link ID is the ID of the router that created the router LSA. The advertising router (shown as "ADV Router" in the output) is the router ID of the OSPF router that announced the router LSA. Generally, the link ID and advertising router for a router LSA are the same.

The highlighted router LSA entry in the OSPF database indicates that the router LSA with link ID 10.1.10.1 has been updated three times (because the sequence number is 0x80000003) and that the last update occurred 115 seconds ago.

Output on the figure is taken from the Cisco IOS XR router. Output from the Cisco IOS and IOS XE routers would be similar to the following:

```
CE1# show ip ospf database
            OSPF Router with ID (10.1.10.1) (Process ID 1)
                Router Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        105         0x80000002 0x004322 1
10.1.10.1       10.1.10.1       105         0x80000003 0x004DDC 2
```


https://t.me/learningnets

OSPF Link-State Database

OSPF maintains a database with only the most recent and up to date LSAs with the use of MaxAge and LSRefreshTime timers, and LS sequence numbers.

Router and Network LSAs

Router and Network LSAs

- One router LSA for every router in an area:
  - Includes a list of directly attached links.
  - Links identified by the IP prefix and link type.
- LSA identified by the router ID of the originating router.
- Floods within its area only; does not cross an ABR.
- Advertised by the DR.
- Floods within its area only; does not cross an ABR.

- One network LSA for each transit broadcast or NBMA network:
  - Includes a list of attached routers on the transit link.
  - Includes a subnet mask of the link.
- Advertised by the DR.
- Floods within its area only; does not cross an ABR.

https://t.me/learningnets

Router and Network LSAs (Cont.)

Router LSAs

Link Type	Description	Link ID
1	Point-to-point connection to another router	Neighboring router ID
2	Connection to a transit network	IP address of DR
3	Connection to a stub network	IP network or subnet number
4	Virtual link	Neighboring router ID

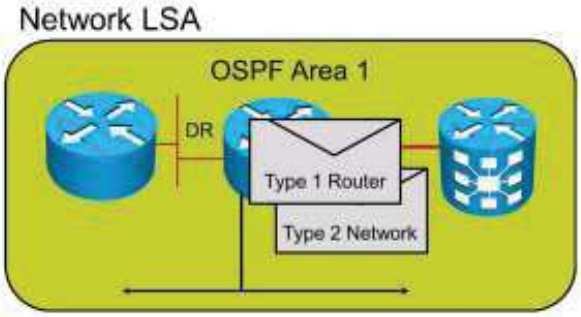

A router advertises a type 1 LSA that floods to all other routers in the area in which it originated. A type 1 LSA describes the collective states of the directly connected links (interfaces) of the router. Each type 1 LSA is identified by the router ID. Each router link is defined as one of four link types: type 1, 2, 3, or 4. The LSA includes a link ID field that identifies, by the network number and mask, the object to which this link connects.

Depending on the type, the link ID has different meanings, as described in the table. A stub network is a dead-end link that has only one attached router. In addition, the type 1 LSA describes whether the router is an ABR or ASBR.

A type 2 LSA is generated for every transit broadcast or NBMA network within an area. A transit network has at least two directly attached OSPF routers. A multiaccess network like Ethernet is an example of a transit network.

The DR of the network is responsible for advertising the network LSA. A type 2 network LSA lists each of the attached routers that make up the transit network, including the DR itself and the subnet mask that is used on the link. The type 2 LSA then floods to all routers within the transit network area. Type 2 LSAs never cross an area boundary. The link-state ID for a network LSA is the IP interface address of the DR that advertises it.

https://t.me/learningnets

Summary LSAs

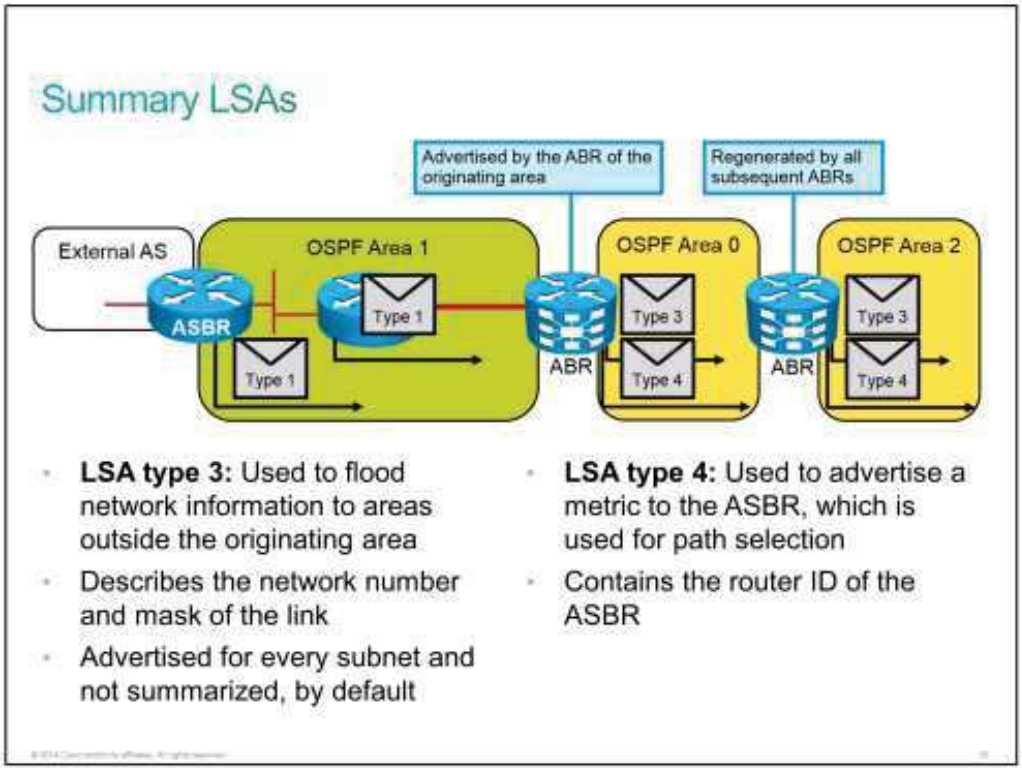

The ABR sends type 3 summary LSAs. Type 3 LSAs advertise any networks that are owned by an area to the rest of the areas in the OSPF AS, as shown in the figure. The link-state ID is set to the network number, and the mask is also advertised.

By default, OSPF does not automatically summarize groups of contiguous subnets, nor does it summarize a network to its classful boundary. The network operator, through configuration commands, must specify how the summarization will occur. By default, a type 3 LSA is advertised into the backbone area for every subnet that is defined in the originating area, which can cause significant flooding problems. Consequently, you should always consider using manual route summarization at the ABR. Summary LSAs are flooded throughout a single area only, but are regenerated by ABRs to flood into other areas. By default, summary LSAs do not contain summarized routes.

A type 4 summary LSA is generated by an ABR only when an ASBR exists within an area. A type 4 LSA identifies the ASBR and provides a route to it. The link-state ID is set to the ASBR router ID. All traffic that is destined to an external AS requires routing table knowledge of the ASBR that originated the external routes.

In the figure, the ASBR sends a type 1 router LSA with a bit (known as the external bit [e bit]) that is set to identify itself as an ASBR. When the ABR (identified with the border bit [b bit] in the router LSA) receives this type 1 LSA, it builds a type 4 LSA and floods it to the backbone, Area 0. Subsequent ABRs regenerate a type 4 LSA to flood into their areas.

https://t.me/learningnets

External LSAs

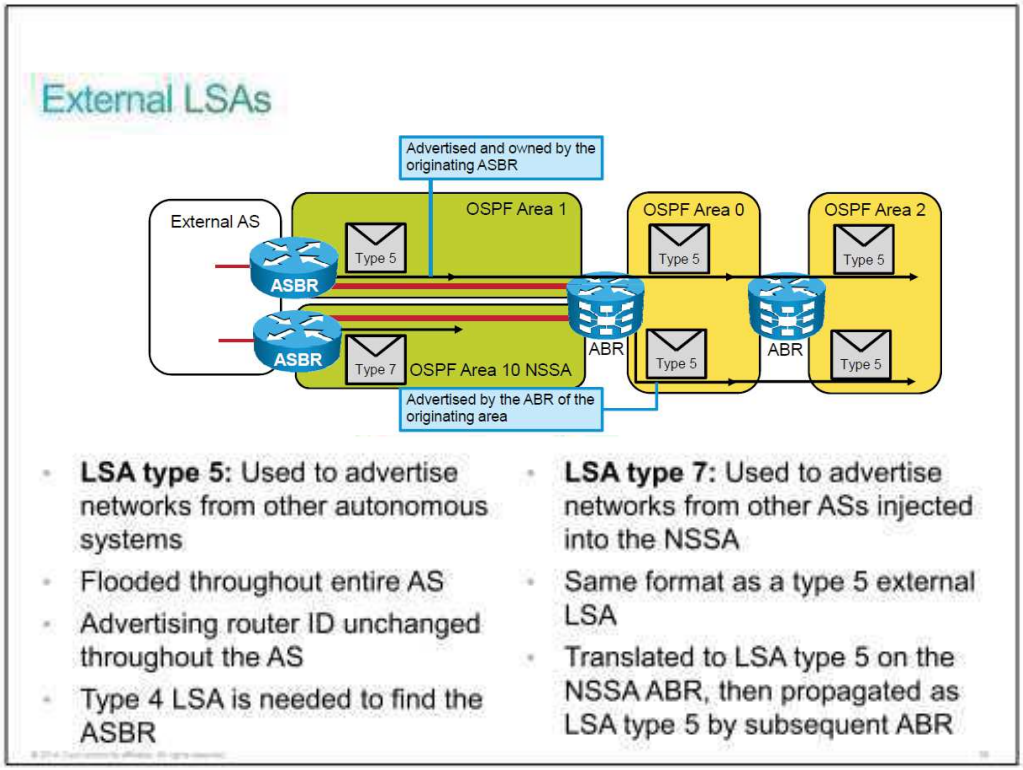

Type 5 external LSAs describe routes to networks outside the OSPF AS. Type 5 LSAs are originated by the ASBR and are flooded to the entire AS. The link-state ID is the external network number. Because of the flooding scope and depending on the number of external networks, the default lack of route summarization can also be a major issue with external LSAs. Therefore, you should always attempt to summarize blocks of external network numbers at the ASBR to reduce flooding problems.

Type 7 external LSAs describe routes to networks outside the OSPF AS. Redistribution from an external AS into an NSSA area creates this special type 7 LSA, which can only exist in an NSSA area. An NSSA ASBR generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain to all areas that can support type 5 LSAs.

Routers that are operating NSSA areas set the n-bit to signify that they can support type 7 NSSAs. These option bits must be checked during neighbor establishment. They must match for an adjacency to form.

The advertising router is set to the router ID of the router that injected the external route into OSPF—a router inside this NSSA area. This address is also set as the "Forward Address" for this prefix—the address that is used to determine the path to take toward this external destination.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Intra-Area Routing

The figure shows the topology that will be used to describe the OSPF LSDB for intra-area routing. Routers are configured for OSPF routing. The Cisco IOS XR **show ospf database** or Cisco IOS and IOS XE **show ip ospf database** commands are used to get information about an OSPF LSDB.

The router link state is type 1 LSA, and the net link state is type 2 LSA.

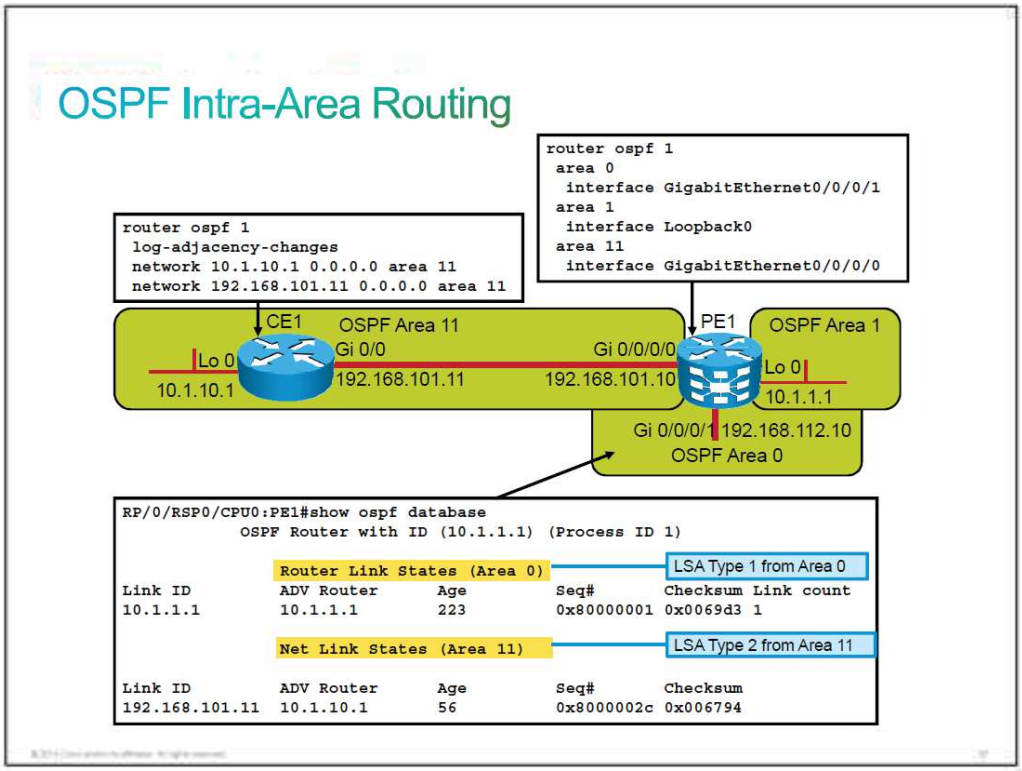

The database columns are as follows:

- **Link ID:** This column identifies each LSA.
- **ADV router:** This column shows the address of the advertising router—the source router of the LSA.
- **Age:** This column shows the maximum age counter in seconds. The maximum configurable age counter is 1 hour, or 3600 seconds.
- **Seq#:** This column shows the sequence number of the LSA. The number begins at 0x80000001 and increases with each update of the LSA.
- **Checksum:** This column shows the checksum of the individual LSA, which can be used to ensure reliable receipt of that LSA.
- **Link count:** This column shows the total number of directly attached links, which is used only on router LSAs. The link count includes all point-to-point, transit, and stub links. Each point-to-point serial link counts as two, and all other links count as one, including Ethernet links.

https://t.me/learningnets

The output in the figure is partial output from an ABR (Cisco IOS XR), router PE1. The complete command output from this router is as follows:

```
RP/0/RSP0/CPU0:PE1# show ospf database
Fri May 26 01:11:49.747 UTC
          OSPF Router with ID (10.1.1.1) (Process ID 1)
              Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        223         0x80000001 0x0069d3 1
              Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        223         0x80000001 0x0080a1
10.1.10.1       10.1.1.1        223         0x80000001 0x0027f0
192.168.101.0   10.1.1.1        223         0x80000001 0x001749
              Router Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        89          0x8000002d 0x00f3e9 1
              Summary Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        223         0x80000001 0x0027f0
192.168.101.0   10.1.1.1        223         0x80000001 0x001749
192.168.112.0   10.1.1.1        223         0x80000001 0x009db7
              Router Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        89          0x8000002e 0x00ed4a 1
10.1.10.1       10.1.10.1       56          0x8000002e 0x00f608 2
              Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.101.11  10.1.10.1       56          0x8000002c 0x006794
              Summary Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        223         0x80000001 0x0080a1
192.168.112.0   10.1.1.1        223         0x80000001 0x009db7
```

The complete output from an internal router (Cisco IOS), router CE1, is as follows:

```
CE1# show ip ospf database
          OSPF Router with ID (10.1.10.1) (Process ID 1)
              Router Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        341         0x80000030 0x00E94C 1
10.1.10.1       10.1.10.1       288         0x80000030 0x00F20A 2
              Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.101.11  10.1.10.1       288         0x8000002E 0x006396
              Summary Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        581         0x80000003 0x007CA3
192.168.112.0   10.1.1.1        582         0x80000003 0x0099B9
```


© 2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Inter-Area Routing

The figure presents the topology that will be used to describe the OSPF LSDB for interarea routing. All routers are configured for OSPF routing. The summary net link states are type 3 LSAs. Because PE1 is the ABR, it has the database for all areas to which it is connected, making it the best place to see the OSPF database. To advertise routes from one area into another, the ABR creates summary links, which you can see using the Cisco IOS XR **show ospf database summary** or Cisco IOS and IOS XE **show ip ospf database summary** commands.

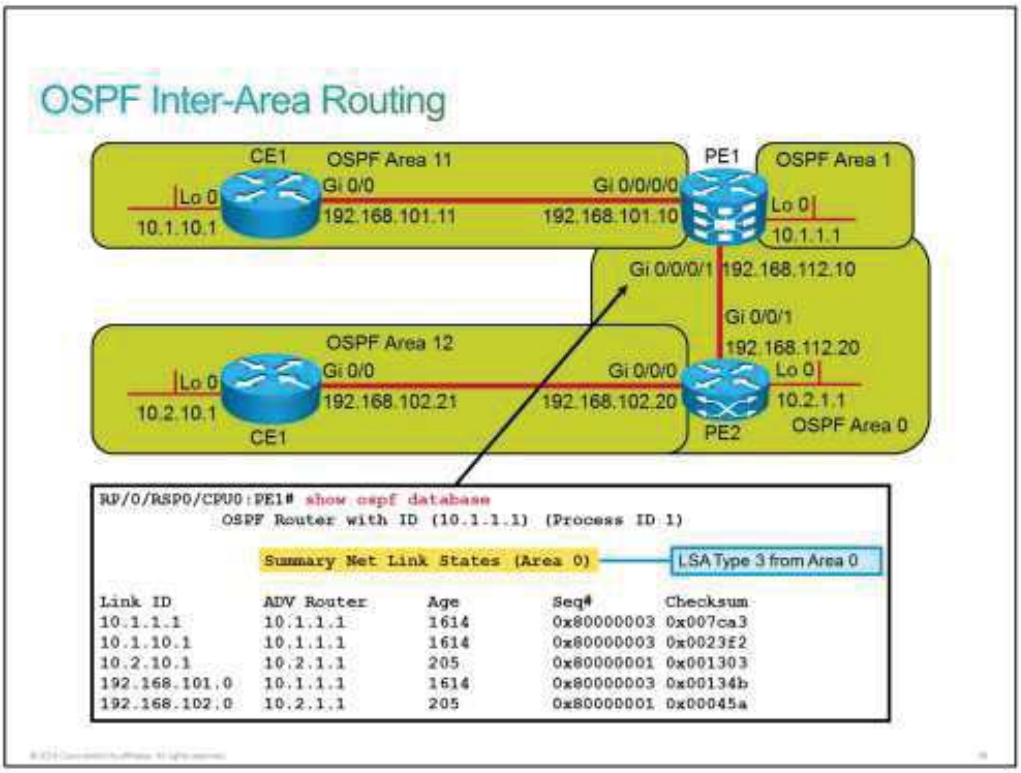

https://t.me/learningnets

The output in the figure is partial output from an ABR (Cisco IOS XR), router PE1. The complete command output from this router is as follows:

```
RP/0/RSP0/CPU0:PE1# show ospf database
Fri May 26 02:48:28.981 UTC
          OSPF Router with ID (10.1.1.1) (Process ID 1)
              Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        729         0x80000004 0x002b22 1
10.2.1.1        10.2.1.1        729         0x80000002 0x00a79b 1
              Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.112.10  10.1.1.1        729         0x80000001 0x0048e5
              Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        147         0x80000004 0x007aa4
10.1.10.1       10.1.1.1        147         0x80000004 0x0021f3
10.2.10.1       10.2.1.1        735         0x80000001 0x001303
192.168.101.0   10.1.1.1        147         0x80000004 0x00114c
192.168.102.0   10.2.1.1        735         0x80000001 0x00045a
              Router Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        1903        0x8000002f 0x00efeb 1
              Summary Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        147         0x80000004 0x0021f3
10.2.10.1       10.1.1.1        728         0x80000001 0x0025f0
192.168.101.0   10.1.1.1        147         0x80000004 0x00114c
192.168.102.0   10.1.1.1        728         0x80000001 0x001648
192.168.112.0   10.1.1.1        147         0x80000004 0x0097ba
              Router Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        1903        0x80000030 0x00e94c 1
10.1.10.1       10.1.10.1       1852        0x80000030 0x00f20a 2
              Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.101.11  10.1.10.1       1852        0x8000002e 0x006396
              Summary Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        147         0x80000004 0x007aa4
10.2.10.1       10.1.1.1        728         0x80000001 0x0025f0
192.168.102.0   10.1.1.1        728         0x80000001 0x001648
192.168.112.0   10.1.1.1        147         0x80000004 0x0097ba
```


©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF External Routes

The figure shows the topology that will be used to describe the OSPF LSDB for external routes. All routers are configured for OSPF routing. The figure illustrates the use of the Cisco IOS XR **show ospf database** or Cisco IOS and IOS XE **show ip ospf database** commands to get information about an OSPF LSDB for external routes.

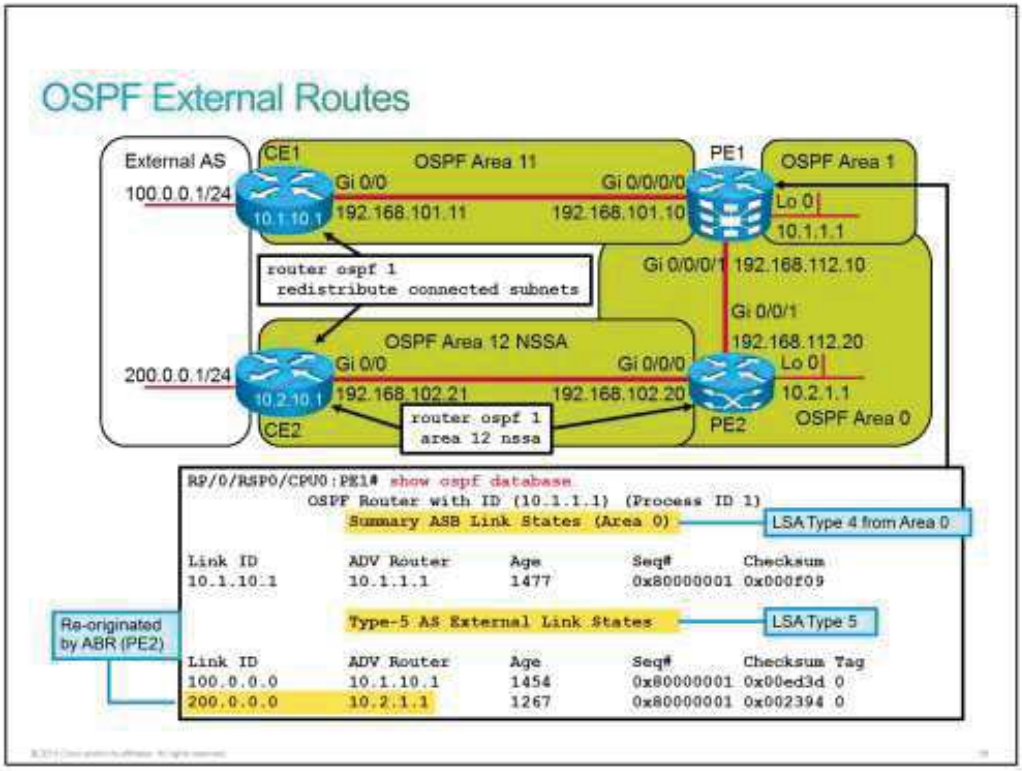

The ASBR (router CE1) creates external LSAs (type 5) to advertise external routes into OSPF. External LSAs are flooded unaltered into all areas. However, the ASBR is not in Area 0. Routers in Area 0 do not know how to reach the ASBR. To advertise the reachability of an ASBR into other areas, the ABR (router PE1) creates ASBR summary LSAs (type 4).

To advertise external routes into an NSSA, the ASBR (router CE2) creates NSSA external LSAs (type 7). The ABR (router PE2) converts type 7 LSAs into type 5 LSAs and propagates the type 5 LSAs into normal areas. The ASBR summary LSAs are not needed in this case, because the ABR originates the external LSA and is reachable within Area 0.

The output in the figure is partial output from the ABR (router PE1). The complete command outputs from the ABRs (PE1 and PE2) are as follows:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1# show ospf database
Fri May 26 03:19:37.433 UTC
          OSPF Router with ID (10.1.1.1) (Process ID 1)
              Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        755         0x80000005 0x002923 1
10.2.1.1        10.2.1.1        1278        0x80000003 0x00ab94 1
              Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.112.10  10.1.1.1        755         0x80000002 0x0046e6
              Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        2015        0x80000004 0x007aa4
10.1.10.1       10.1.1.1        2015        0x80000004 0x0021f3
10.2.10.1       10.2.1.1        625         0x80000002 0x001104
192.168.101.0   10.1.1.1        2015        0x80000004 0x00114c
192.168.102.0   10.2.1.1        625         0x80000002 0x00025b
              Summary ASB Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        1477        0x80000001 0x000f09
              Router Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        1751        0x80000030 0x00edec 1
              Summary Net Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        2015        0x80000004 0x0021f3
10.2.10.1       10.1.1.1        755         0x80000002 0x0023f1
192.168.101.0   10.1.1.1        2015        0x80000004 0x00114c
192.168.102.0   10.1.1.1        755         0x80000002 0x001449
192.168.112.0   10.1.1.1        2015        0x80000004 0x0097ba
              Summary ASB Link States (Area 1)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        1477        0x80000001 0x000f09
10.2.1.1        10.1.1.1        1277        0x80000001 0x0066b9
              Router Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        1751        0x80000031 0x00e74d 1
10.1.10.1       10.1.10.1       1478        0x80000032 0x00f404 2
              Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.101.11  10.1.10.1       1749        0x8000002f 0x006197
              Summary Net Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        2015        0x80000004 0x007aa4
10.2.10.1       10.1.1.1        755         0x80000002 0x0023f1
192.168.102.0   10.1.1.1        755         0x80000002 0x001449
192.168.112.0   10.1.1.1        2015        0x80000004 0x0097ba
              Summary ASB Link States (Area 11)
Link ID         ADV Router      Age         Seq#       Checksum
10.2.1.1        10.1.1.1        1277        0x80000001 0x0066b9
```


https://t.me/learningnets

```
                Type-5 AS External Link States
Link ID         ADV Router      Age         Seq#       Checksum Tag
100.0.0.0       10.1.10.1       1454        0x80000001 0x00ed3d 0
200.0.0.0       10.2.1.1        1267        0x80000001 0x002394 0
PE2#show ip ospf database
          OSPF Router with ID (10.2.1.1) (Process ID 1)
                Router Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.1.1.1        10.1.1.1        1402        0x80000005 0x002923 1
10.2.1.1        10.2.1.1        23          0x80000004 0x00A995 1
                Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.112.10  10.1.1.1        1402        0x80000002 0x0046E6
                Summary Net Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.1.1.1        639         0x80000005 0x0078A5
10.1.10.1       10.1.1.1        639         0x80000005 0x001FF4
10.2.10.1       10.2.1.1        1270        0x80000002 0x001104
192.168.101.0   10.1.1.1        639         0x80000005 0x000F4D
192.168.102.0   10.2.1.1        1270        0x80000002 0x00025B
                Summary ASB Link States (Area 0)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.10.1       10.1.1.1        128         0x80000002 0x000D0A
                Router Link States (Area 12)
Link ID         ADV Router      Age         Seq#       Checksum Link count
10.2.1.1        10.2.1.1        23          0x80000004 0x00E061 1
10.2.10.1       10.2.10.1       112         0x80000007 0x00F410 2
                Net Link States (Area 12)
Link ID         ADV Router      Age         Seq#       Checksum
192.168.102.21  10.2.10.1       112         0x80000004 0x00F916
                Summary Net Link States (Area 12)
Link ID         ADV Router      Age         Seq#       Checksum
10.1.1.1        10.2.1.1        23          0x80000003 0x0024F3
10.1.10.1       10.2.1.1        23          0x80000003 0x00CA43
192.168.101.0   10.2.1.1        23          0x80000003 0x00BA9B
192.168.112.0   10.2.1.1        23          0x80000003 0x003715

                Type-7 AS External Link States (Area 12)
Link ID         ADV Router      Age         Seq#       Checksum Tag
200.0.0.0       10.2.10.1       112         0x80000002 0x004D56 0
                Type-5 AS External Link States
Link ID         ADV Router      Age         Seq#       Checksum Tag
100.0.0.0       10.1.10.1       144         0x80000002 0x00EB3E 0
200.0.0.0       10.2.1.1        23          0x80000002 0x002195 0
```


https://t.me/learningnets

OSPF Virtual Link

The figure presents the topology that will be used to describe the OSPF LSDB when a virtual link is configured. All routers are configured for OSPF routing. The figure illustrates the use of the Cisco IOS XR **show ospf database** or Cisco IOS and IOS XE **show ip ospf database** commands to get information about an OSPF LSDB for virtual links.

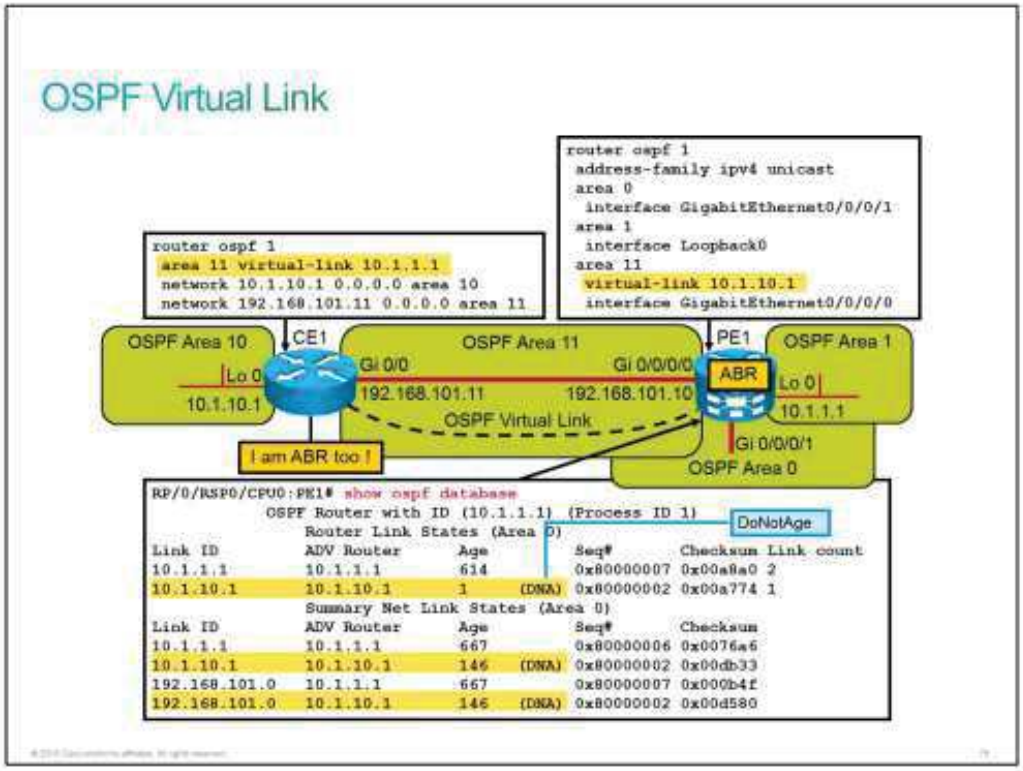

The router link states are type 1 LSAs, and the summary net link states are type 3 LSAs. Both are used to advertise routes from one area into another. Notice that LSAs that are learned through the virtual link have the DNA option. The virtual link is treated like a demand circuit.

Router CE1 considers itself an ABR, because it has a link to Area 0 (the virtual link). As a result, it generates a summary LSA for 10.1.10.1 and 192.168.101.0 into Area 0, which you can see when you issue the **show ospf database** command on PE1.

https://t.me/learningnets

Interpreting OSPF Routes in the Routing Table

Using the example in the figure, you will learn how to interpret the routing table entries for OSPF learned routes.

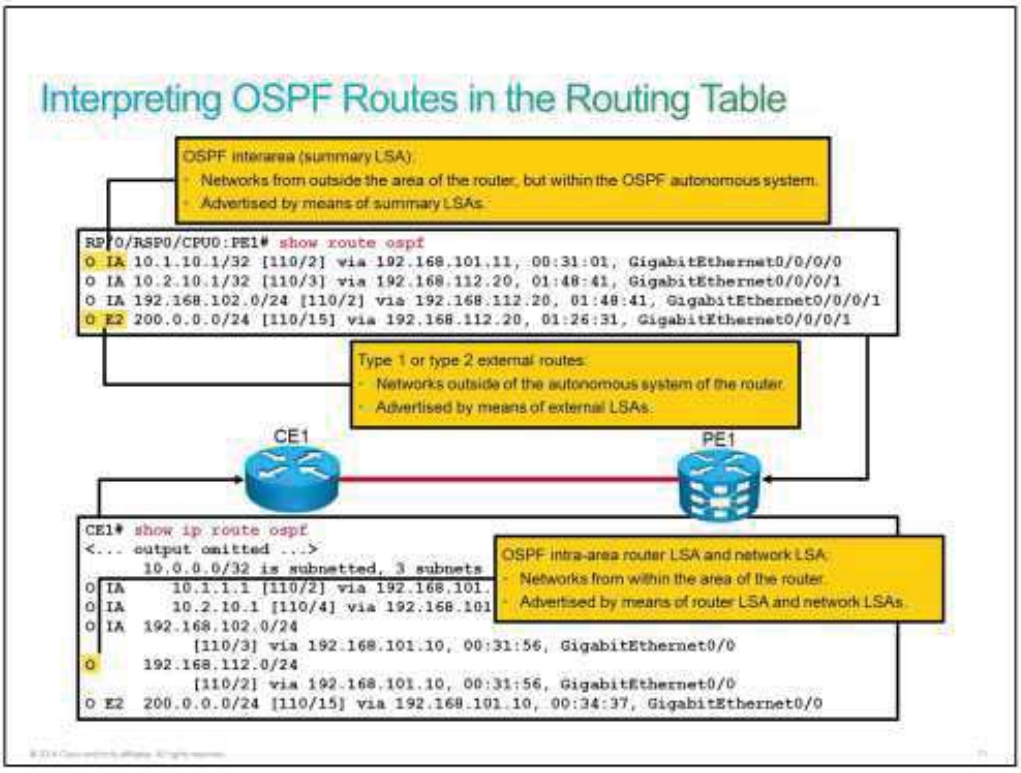

The Cisco IOS XR **show route** and Cisco IOS and IOS XE **show ip route** commands in the figure depict intra-area (O), interarea (O IA) and external type 2 (O E2) routes.

The O E2 entry is an external route from the ASBR, via the ABR. The two numbers in brackets, [110/15], are the administrative distance and the total cost of the route to a specific destination network. In this case, the administrative distance is set to the default of 110 for all OSPF routes, and the total cost of the route has been calculated as 15.

The router and network LSAs describe the details within an area. The routing table reflects this link-state information with a designation of "O," meaning that the route is OSPF intra-area.

When an ABR receives summary LSAs, it adds them to its LSDB and regenerates them into the local area. When an ABR receives external LSAs, it adds them to its LSDB and floods them into the area. The internal routers then assimilate the information into their databases. Summary LSAs appear in the routing table as IA. External LSAs appear in the routing table that is marked as E1 or E2 routes.

The SPF algorithm is then run against the LSDB to build the SPF tree. The SPF tree is used to determine the best paths. The order in which the best paths are calculated is as follows:

- All routers calculate the best paths to destinations within their areas (intra-area) and add these entries to the routing table. These are the type 1 and type 2 LSAs, which are noted in the routing table with a routing designator of "O."
- All routers calculate the best paths to the other areas within the internetwork. These best paths are the interarea route entries, or type 3 and type 4 LSAs, and are noted with a routing designator of O IA.

https://t.me/learningnets

- All routers (except those that are in a form of stub area) calculate the best paths to the external AS (type 5) destinations. These are noted with either an O E1 or an O E2 route designator, depending on the configuration.

At this point, a router can communicate with any network within or outside the OSPF AS.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Calculating Costs for E1 and E2 OSPF Routes

The cost of an external route varies, depending on the external type that is configured on the ASBR. The following external packet types can be configured:

- **E1:** Type O E1 external routes calculate the cost by adding the external cost to the internal cost of each link that the packet crosses. Use this type when there are multiple ASBRs that are advertising an external route to the same AS, to avoid suboptimal routing.
- **E2 (default):** The external cost of O E2 packet routes is always the external cost only. Use this type if only one ASBR is advertising an external route to the AS.

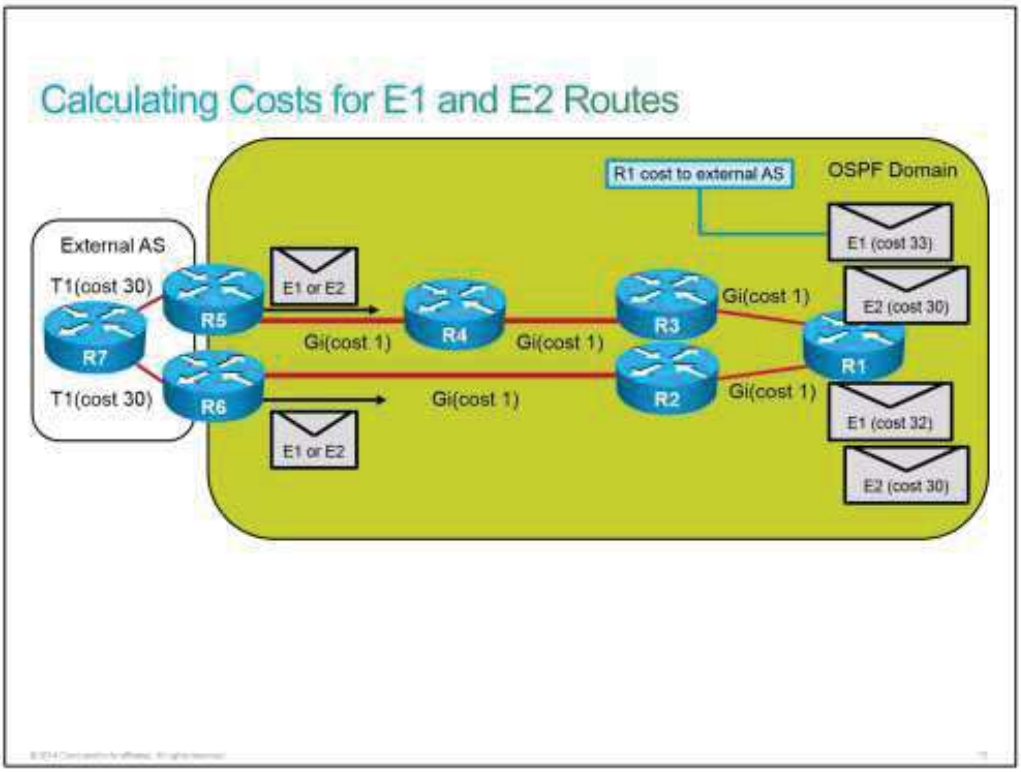

The figure shows the network diagram with two ASBRs (R5 and R6). Both ASBRs are sending external routes into the OSPF AS. External routes can be sent as E1 or E2. R1 receives the same external routes from R2 and R3. In the figure, the path from R1 to R6 is shorter than the path from R1 to R5. If external routes are received as E2 routes (default setting), the cost is the same regardless of the topology in the OSPF domain, which means that there will be suboptimal routing. If external routes are received as E1 routes, the cost is different, because the internal OSPF cost is added to the external cost. The routing is optimal, and the shortest path is selected to the destination.

https://t.me/learningnets

OSPF LSDB Overload Protection

If other routers are misconfigured, causing many prefixes to be redistributed, large numbers of LSAs can be generated. These excessive LSAs can drain local CPU and memory resources. OSPF LSDB overload protection can be configured to protect against this issue by using the Cisco IOS, IOS XE, and IOS XR **max-lsa** router OSPF command.

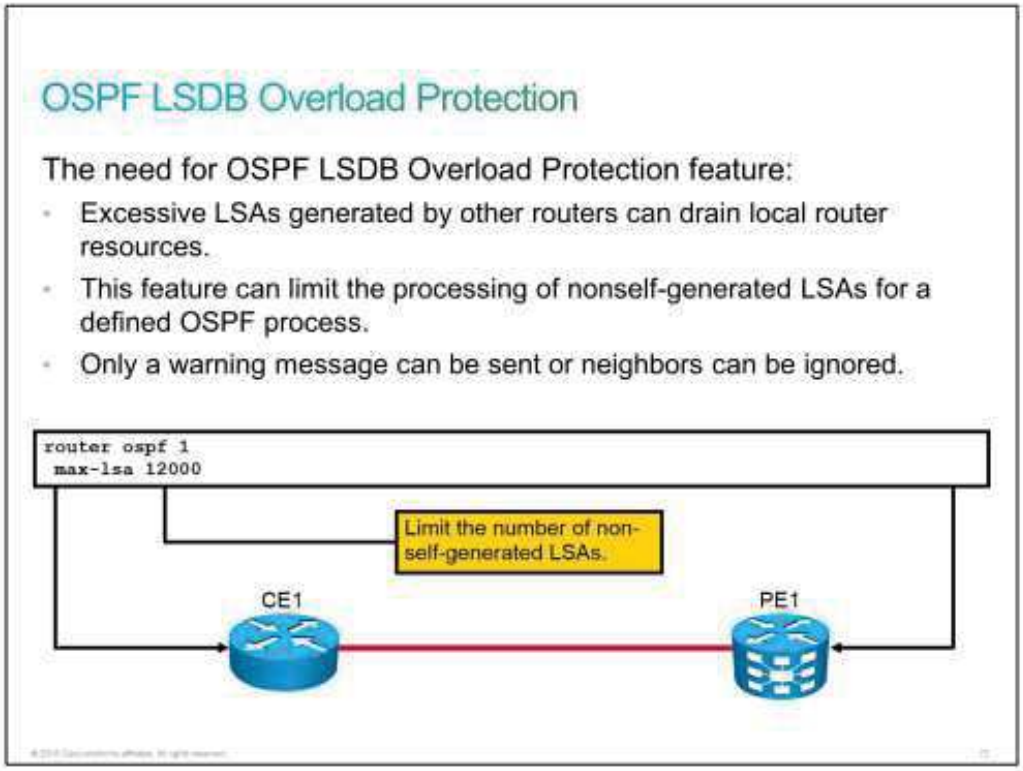

When this feature is enabled, the router keeps count of the number of received (nonself-generated) LSAs that it keeps in its LSDB. An error message is logged when this number reaches a configured threshold number, and a notification is sent when it exceeds the threshold number.

If the LSA count still exceeds the threshold after 1 minute, the OSPF process takes down all adjacencies and clears the OSPF database. This is called the "ignore" state. In the ignore state, no OSPF packets are sent or received by interfaces that belong to that OSPF process.

The OSPF process remains in the ignore state for the time that is defined by the **ignore-time** parameter. The **ignore-count** parameter defines the maximum number of times that the OSPF process can consecutively enter the ignore state before remaining permanently down and requiring manual intervention.

If the OSPF process remains normal for the time that is defined by the **reset-time** parameter, the ignore state counter is reset to 0.

https://t.me/learningnets

Summary

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

Summary

- OSPF uses link-state advertisements to build a topology database, all OSPF routers have the same picture of network topology.
- ASBR routers connects OSPF area to external routing domain.
- OSPF routing protocol reduces the size of the SPF calculations by partitioning the network into multiple areas.
- Hierarchical structure of OSPF minimizes routing table entries, localizes impact of a topology change and stops LSA flooding.
- There are 11 OSPF LSA types. No external routes are propagated into OSPF stub areas. No external and intra-area routes are propagated into OSPF totally stubby areas.
- OSPF routers find the best path to a destination by applying the Dijkstra SPF algorithm to the topology database.
- OSPF cost of an interface is calculated based on the bandwidth and can be changed by using the OSPF configuration command.

Summary (Cont.)

- OSPF uses hello packets to discover neighbors.
- The following are the most commonly used OSPF LSA types: type 1 router, type 2 network, type 3 and 4 summary, type 5 external, and type 7 external.
- OSPF intra-area LSAs are used to describe routes within an area. OSPF inter-area LSAs are used to describe routes from other areas. OSPF external LSAs are used to describe routes from external autonomous systems.
- Intra-area routes are presented in the routing table with „O“, while inter-area routes are presented with „O IA“.
- External routes can be configured to use either E1 or E2 cost type.
- OSPFv2 and OSPFv3 have run independently on a network device.

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Understanding OSPF Operation

Overview

The OSPF protocol has several functions, and five packet types to support them: hello, DBD, link-state request LSR, LSU, and LSAck. These five OSPF packet types enable all OSPF information flow between routers.

The OSPF protocol areas may be made up of different types of network links. You must know which types of network links are being used in order to properly configure OSPF to work with all of them and their different adjacency behaviors, as well as over certain network types. It is important to note that OSPF pays special attention to different network types, such as point-to-point and broadcast networks, and that the OSPF default settings do not always work properly with some network topologies.

This lesson defines packet types and explains where and how these packets interact to build OSPF neighbor adjacencies and maintain the OSPF topology database. The lesson also describes OSPF network types, how the adjacencies are formed for these OSPF network types, and how LSAs are flooded on each.

Upon completing this lesson, you will be able to:

- Describe the functions that are performed by OSPF
- Describe the OSPF packet format
- Describe each of the OSPF packet types
- Describe OSPF neighbor adjacency establishment, LSDB exchange, and synchronization
- Describe link-state flooding and the flooding process steps in a multiaccess network
- Describe how to debug OSPF packets
- Describe OSPF network types
- Describe how the DR and backup designated router BDR are elected
- Describe OSPF in NBMA networks
- Describe OSPF in Metro Ethernet and EoMPLS networks

https://t.me/learningnets

- Describe OSPF in Layer 3 MPLS VPN networks
- Describe how to enable OSPF on a link with NBMA mode
- Describe subinterfaces over multiaccess WAN networks
- Describe how to enable OSPF on a link with point-to-point mode
- Describe how to enable OSPF on a link with point-to-multipoint mode
- Describe how to enable OSPF on a link with point-to-multipoint nonbroadcast mode
- Describe how to configure IPv6 support for OSPF modes

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Functions

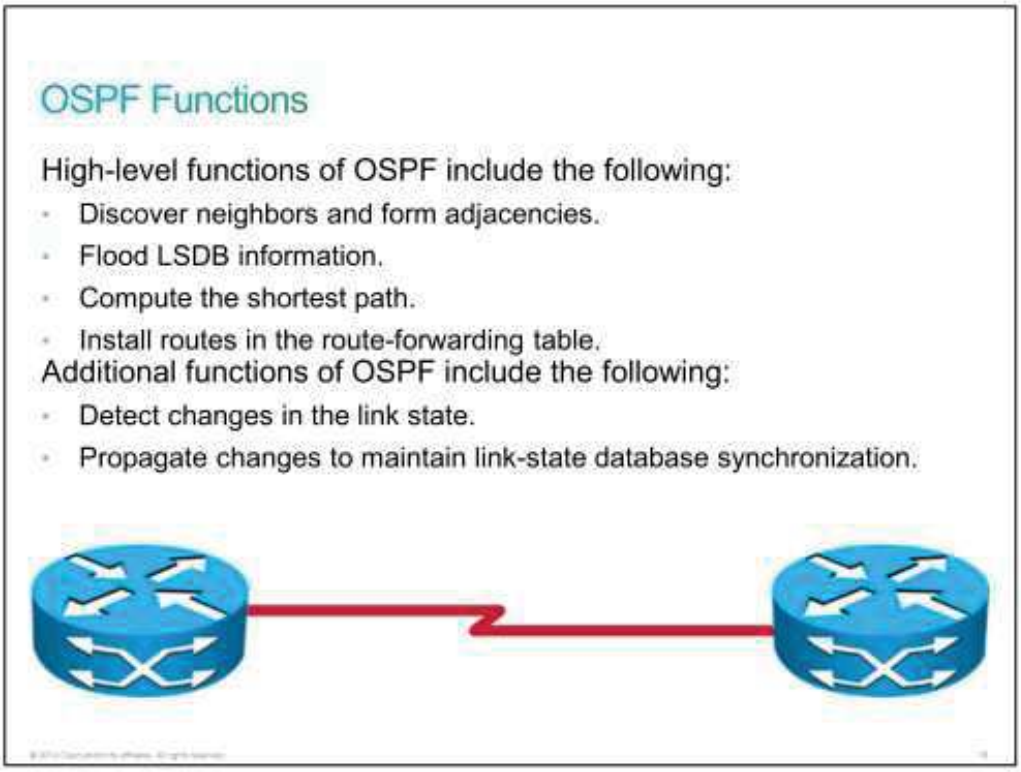

After OSPF initially populates the router forwarding table, the state of the links that are around the OSPF autonomous system may change. OSPF is able to detect these changes and respond by flooding this information throughout the OSPF autonomous system, or at least in the area where the change was detected. The flooding of new information is needed to maintain the LSDBs in all neighboring routers.

For all these functions, several OSPF packet types are involved.

https://t.me/learningnets

OSPF Packet Format

There are five distinct OSPF packet types.

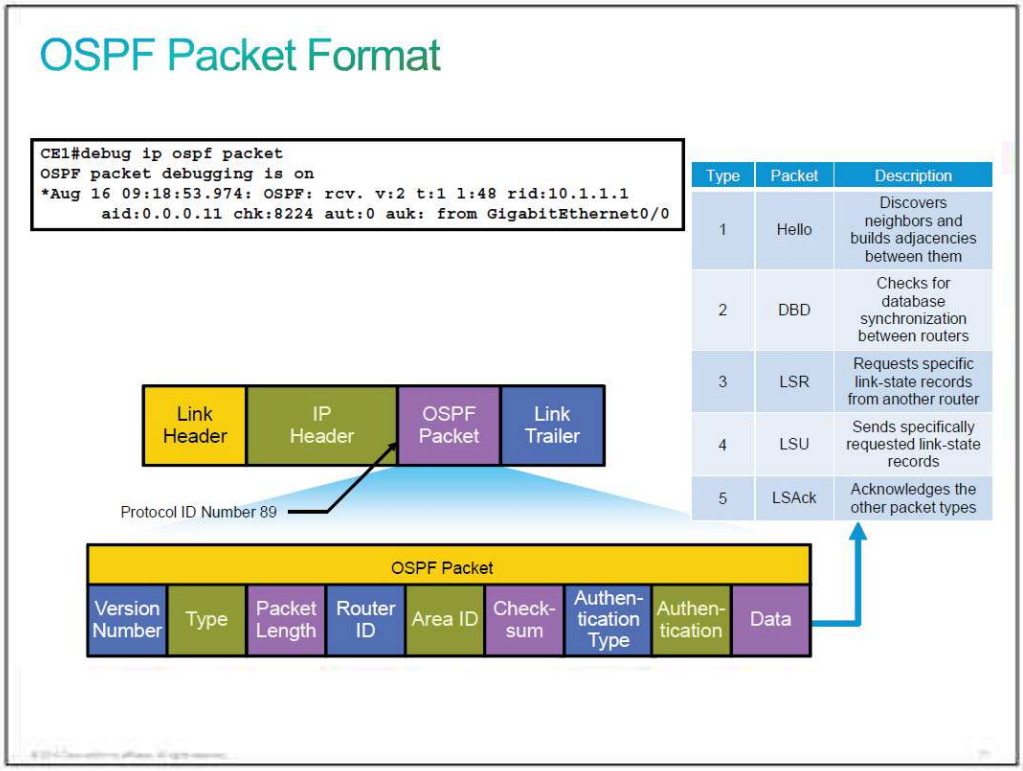

All five OSPF packets are encapsulated directly into an IP payload, as shown in the figure. The OSPF packet does not use TCP or UDP. OSPF requires a reliable packet transport scheme, and because it does not use TCP, it has defined its own acknowledgment routine using an acknowledgment packet (OSPF packet type 5).

©2014 Cisco Systems, Inc.

https://t.me/learningnets

In the IP header, a protocol identifier of 89 defines all OSPF packets. Each of the OSPF packets begins with the same header format. This header has the following fields:

- **Version number:** Version 2 for OSPF with IPv4 and version 3 for OSPF with IPv6.
- **Type:** Differentiates the five OSPF packet types.
- **Packet length:** The length of the OSPF packet in bytes.
- **Router ID:** Defines which router is the source of the packet.
- **Area ID:** Defines the area where the packet originated.
- **Checksum:** Used for packet-header error detection to ensure that the OSPF packet was not corrupted during transmission.
- **Authentication type:** An option in OSPF that describes no authentication, cleartext passwords, or encrypted MD5 formats for router authentication.
- **Authentication:** Used in the authentication scheme.
- **Data:** Each of the five packet types includes different data:
  - **Hello packets:** Contains a list of known neighbors.
  - DBD **packet:** Database descriptor contains a summary of the LSDB, which includes all known router IDs and their last sequence numbers, among several other fields.
  - LSR **packet:** Contains the type of LSA that is needed and the router ID that has the needed LSU.
  - **LSU packet:** Contains the complete LSA entries. Multiple LSA entries can fit in one OSPF update packet.
  - **LSAck packet:** Empty.

https://t.me/learningnets

OSPF Packet Types

Each interface that is participating in OSPF uses the IP multicast address 224.0.0.5 to periodically send hello packets.

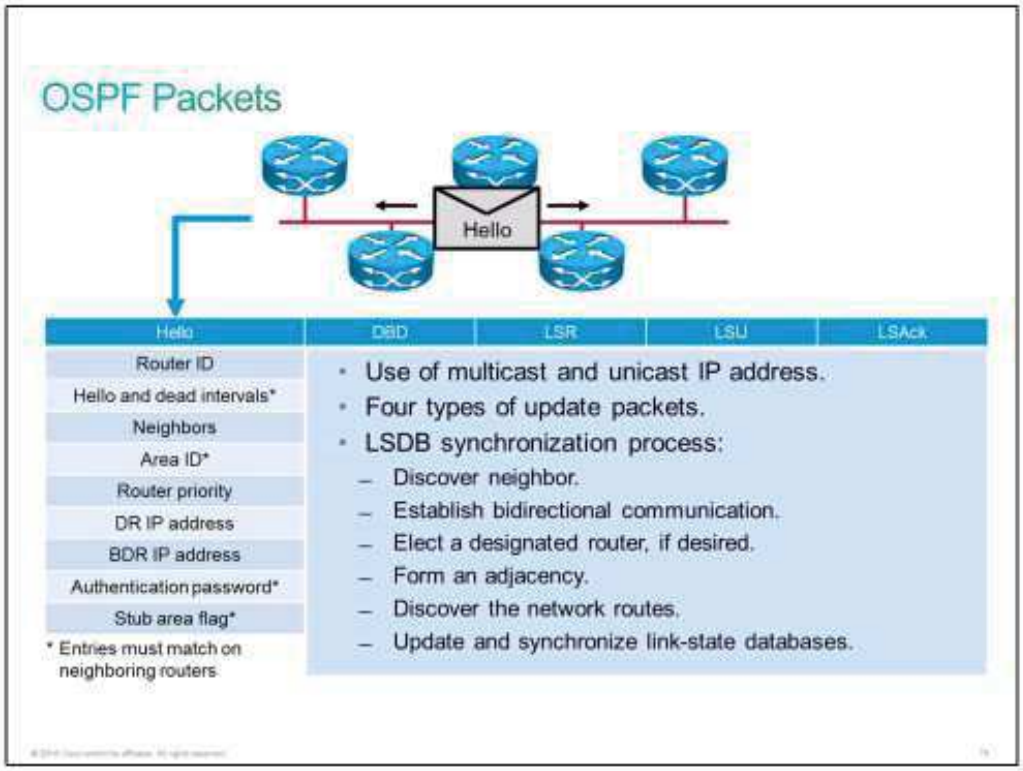

A hello packet contains the following information:

- **Router ID:** The router ID is a 32-bit number that uniquely identifies the router. The highest IP address on an active interface is chosen by default, unless a loopback interface or the router ID is configured. For example, IP address 172.16.12.1 would be chosen over 172.16.1.1. This identification is important in establishing neighbor relationships and coordinating LSU exchanges. Also, the router ID breaks ties during the DRand BDR selection processes if the OSPF priority values are equal.
- **Hello and dead intervals:** The hello interval specifies the frequency, in seconds, with which a router sends hello packets (10 seconds is the default on multiaccess networks). The dead interval is the time in seconds that a router waits to hear from a neighbor before declaring the neighboring router out of service (four times the hello interval, by default). These timers must be the same on neighboring routers; otherwise, an adjacency will not be established.
- **Neighbors:** The Neighbors field lists the adjacent routers with established bidirectional communication. This bidirectional communication is indicated when the router recognizes itself listed in the neighbors field of the hello packet from the neighbor.
- **Area ID:** To communicate, two routers must share a common segment and their interfaces must belong to the same OSPF area on that segment. (They must also share the same subnet and mask.) These routers will all have the same link-state information.
- **Router priority:** The router priority is an 8-bit number that indicates the priority of a router. Priority is used when selecting a DR and BDR.
- **DR and BDR IP addresses:** These are the IP addresses of the DR and BDR for the specific network, if they are known.

https://t.me/learningnets

- **Authentication password:** If router authentication is enabled, two routers must exchange the same password. Authentication is not required, but if it is enabled, all peer routers must have the same password.
- **Stub area flag:** A stub area is a special area. Two routers must agree on the stub area flag in the hello packets. Designating a stub area is a technique that reduces the number of routing updates by replacing many of them with a default route.

The following fields must match when hello packets are exchanged between the neighboring routers: hello and dead intervals, area ID, authentication password, and stub area flag.

After a bidirectional adjacency is formed, OSPF must exchange and synchronize the LSDBs between routers. All routing updates are sent in IP packets (protocol type 89), for which OSPF does not do any fragmentation and reassembly on the IP.

Four types of update packets are used when exchanging and synchronizing LSDBs:

- **Type 2** DBD **packet:** Used to describe the network routes of each neighbor.
- **Type 3** LSR **packet:** After database description packets are exchanged, the routers request missing information by using request packets.
- **Type 4** LSU **packet:** All missing information is sent to the neighbors by sending update packets, which contain different LSAs.
- **Type 5** LSAck **packet:** Every packet is acknowledged, to ensure reliable transport and reliable exchange of information.

Type 4 and type 5 packets are sent to multicast IP addresses, except when retransmitting, when sent across a virtual link, and when sent on nonbroadcast networks. All other packets are sent to unicast IP addresses.

https://t.me/learningnets

OSPF Neighbor States

When OSPF adjacency is formed, a router goes through several state changes before it becomes fully adjacent with its neighbor. The states are Down, Attempt, Init, 2-Way, Exstart, Exchange, Loading, and Full.

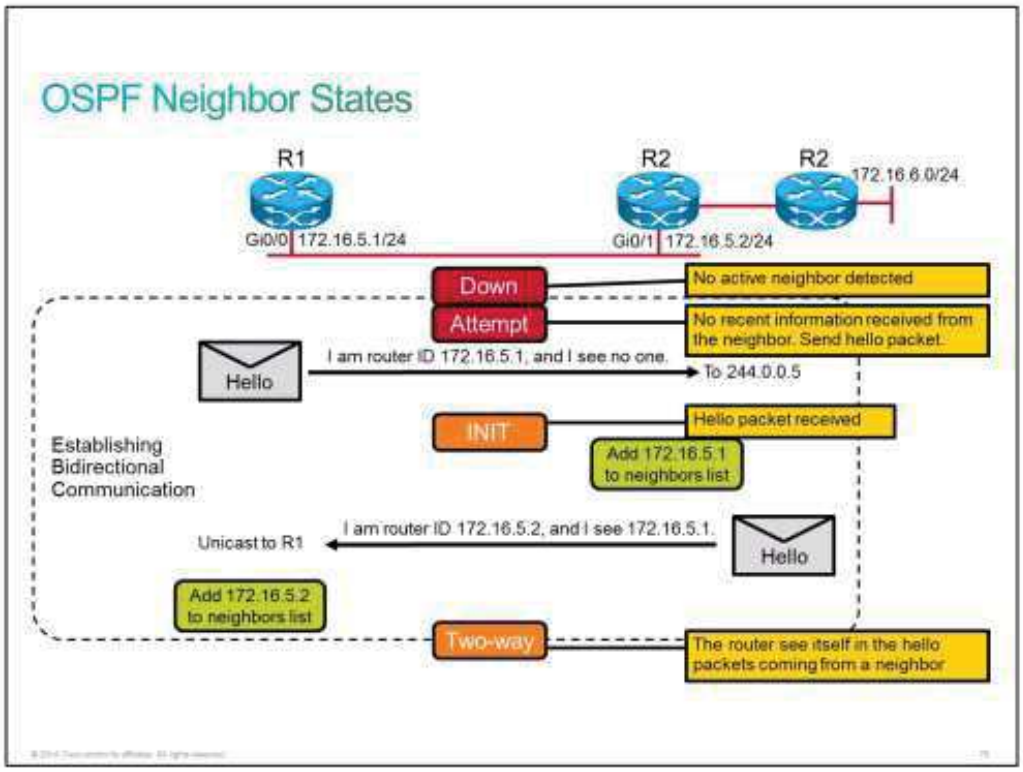

When routers that are running OSPF are initialized, an exchange process using the Hello protocol is the first procedure. The exchange process that happens when routers appear on the network is illustrated in the figure:

- Router R1 is enabled on the LAN and is in a down state, because it has not exchanged information with any other router. It begins by sending a hello packet through each of its interfaces that are participating in OSPF, even though it does not know the identity of the DR or of any other routers. The hello packet is sent out using the multicast address 224.0.0.5.
- All directly connected routers that are running OSPF receive the hello packet from router R1 and add R1 to their lists of neighbors. After adding R1 to the list, other routers are in the INIT state.
- Each router that received the hello packet sends a unicast reply hello packet to R1 with its corresponding information. The neighbor field in the hello packet includes all neighboring routers and R1.
- When R1 receives these hello packets, it adds all the routers that had its router ID in their hello packets to its own neighbor relationship database. After this process, R1 is in the two-way state. At this point, all routers that have each other in their lists of neighbors have established bidirectional communication.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

If the link type is a broadcast network, a DR and BDR must first be selected. The DR forms bidirectional adjacencies with all other routers on the LAN link. This process must occur before the routers can begin exchanging link-state information. Periodically (every 10 seconds on broadcast networks, by default), the routers within a network exchange hello packets to ensure that communication is still working. The hello updates include the DR, BDR, and the list of routers for which hello packets have been received by the router. Remember that "received" means that the receiving router recognizes its own name as one of the entries in the received hello packet.

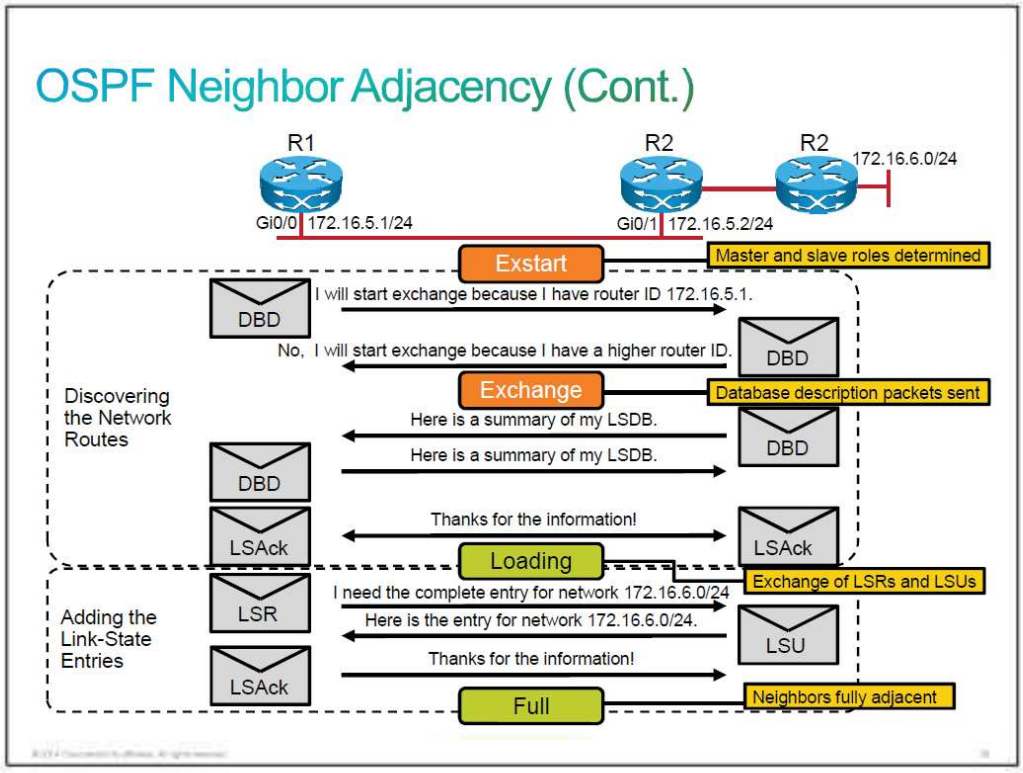

After the DR and BDR are selected, the routers are considered to be in the exstart state. The routers are then ready to discover the link-state information about the internetwork and create their LSDBs. The exchange protocol is used to discover the network routes and it moves all the routers from the exchange state to a full state of communication. The first step in this process is for the DR and BDR to establish adjacencies with each of the other routers. When adjacent routers are in a full state, they do not repeat the exchange protocol unless the full state changes.

Steps for the exchange protocol

1. As shown in the figure, the exchange protocol operates as follows:
   - In the exstart state, the DR and BDR establish adjacencies with each router in the network. During this process, a master-slave relationship is created between each router and its adjacent DR and BDR. The router with the higher router ID acts as the master during the exchange process—router R2 becomes the DR. Only the DR exchanges and synchronizes link-state information with the routers to which it has established adjacencies. Having the DR represent the network in this capacity reduces the amount of routing update traffic.

The master and slave routers exchange one or more DBD packets. The routers are in the exchange state.

A DBD includes information about the LSA entry header that appears in the LSDB of the router. The entries can be about a link or a network. Each LSA entry header includes information about the link-state type, the address of the advertising router, the cost of the link, and the sequence number. The router uses the sequence number to determine the "newness" of the received link-state information.

https://t.me/learningnets

2. When the router receives the DBD, it performs these actions, as shown in the figure:
   - It acknowledges the receipt of the DBD using the LSAck packet.
   - It compares the information that it received with the information that it has. If the DBD has a more up-to-date link-state entry, the router sends an LSR to the other router. When routers start sending LSRs, they are in the loading state.
   - The other router responds with the complete information about the requested entry in an LSU packet. Again, when the router receives an LSU, it sends an LSAck.
3. The router adds the new link-state entries to its LSDB.

When all LSRs are satisfied for a given router, the adjacent routers are considered synchronized. They are in a full state. The routers must be in a full state before they can route traffic. At this point, all the routers in the area should have identical LSDBs.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Link-State Flooding

In a link-state routing environment, it is important for the LSDBs (topology tables) of all routers to stay synchronized. When there is a change in a link state, the routers use a flooding process to notify the other routers in the network of the change.

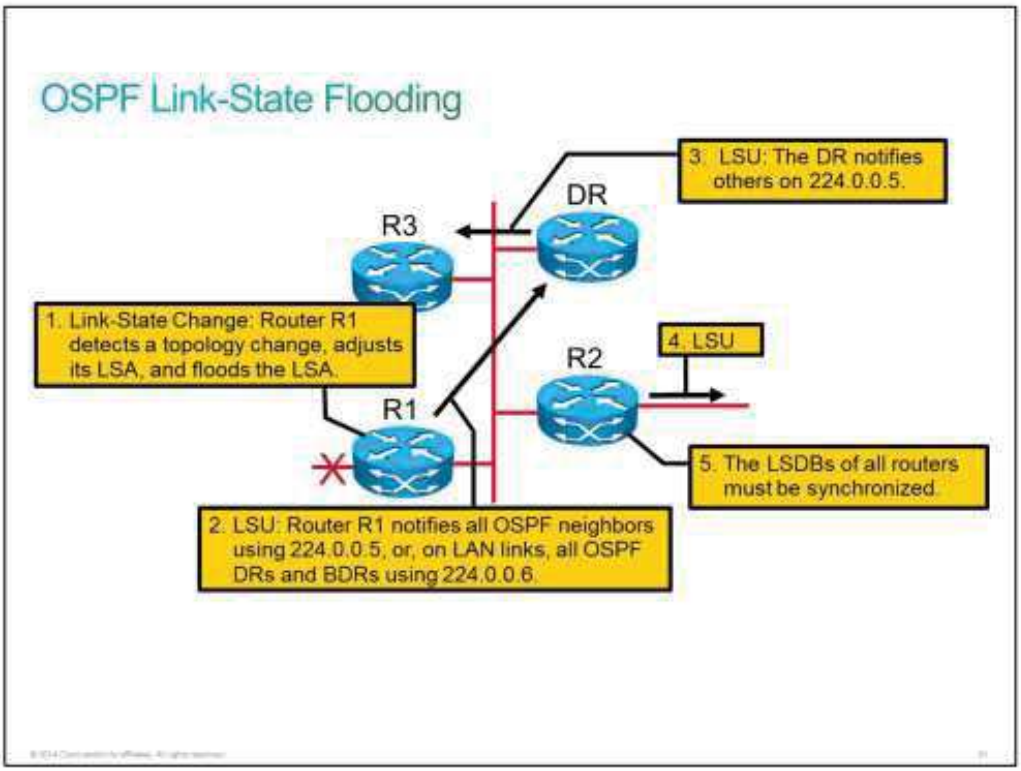

In general, the flooding process steps in a multiaccess network are as follows:

1. A router notices a change in a link state and multicasts an LSU packet to all OSPF neighbors at 224.0.0.5 or to all OSPF DRs and BDRs at 224.0.0.6.
2. The DR acknowledges receipt of the change and floods the LSU to others on the network using the OSPF multicast address 224.0.0.5. After receiving the LSU, each router responds to the DR with an LSAck. To make the flooding procedure reliable, each LSA must be acknowledged separately.
3. If a router is connected to other networks, it floods the LSU to those other networks by forwarding the LSU to the DR of the multiaccess network or to the adjacent router if it is in a point-to-point network.
4. The router updates its LSDB using the LSU that includes the changed LSA. It then recomputes the SPF algorithm against the updated database after a short delay (the SPF delay) and updates the routing table, as necessary.

When is the SPF algorithm run? A change in the topology database is a necessary, but not sufficient, condition for SPF recalculation. These conditions trigger the SPF algorithm:

- The LSA Options field changes.
- The age of the LSA instance is set to MaxAge.
- The length field in the LSA header changes.
- The contents of the LSA (excluding the LSA header) change.

https://t.me/learningnets

An SPF calculation is performed separately for each area in the topology database.

OSPF simplifies the synchronization issue by requiring only adjacent routers to remain synchronized.

Summaries of individual link-state entries, not the complete link-state entries, are sent every 30 minutes to ensure proper LSDB synchronization. Each link-state entry has a timer to determine when the LSA refresh update must be sent.

Each link-state entry also has a maximum age of 60 minutes. If a link-state entry has not been refreshed within 60 minutes, it is removed from the LSDB.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Debug OSPF Packets

To verify that OSPF packets are flowing properly between two routers, use the Cisco IOS XR **debug ospf packet** command or Cisco IOS and IOS XE **debug ip ospf packet** command.

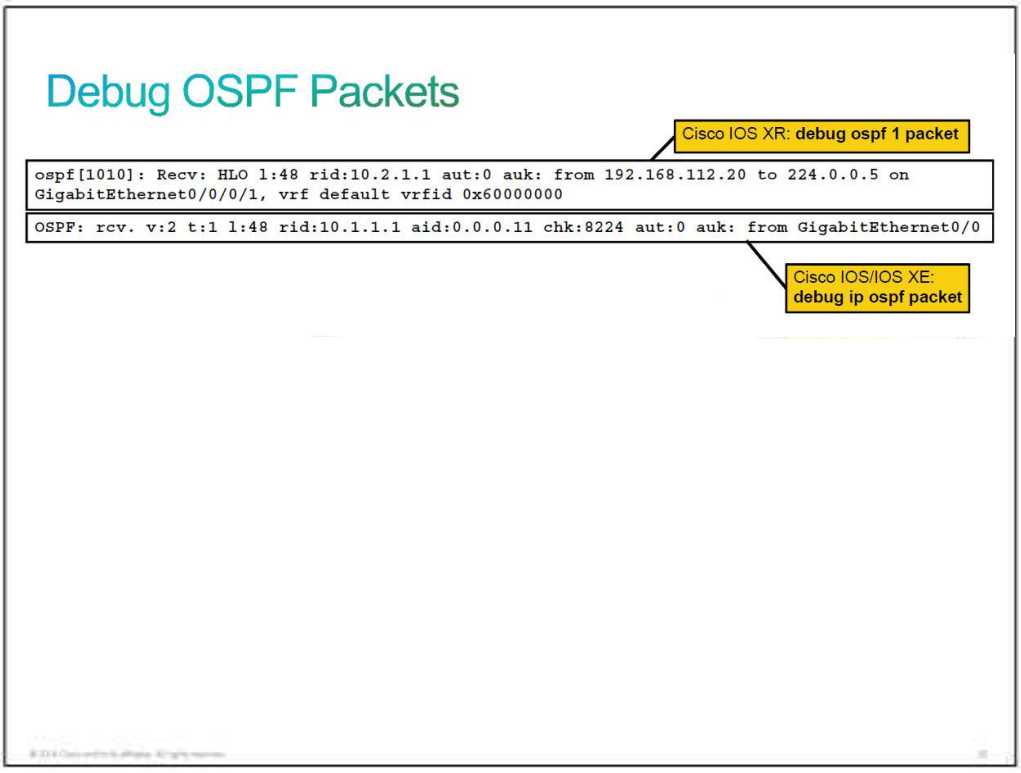

Debug OSPF Packets (Cont.)

Field	Description
OSPF:	OSPF packet
Recv: / rcv.	was received
v:	Provides the version of OSPF
HLO / t:	Specifies the OSPF packet type: 1: hello; 2: DBD; 3: LSR; 4: LSU; 5: LSAck
l:	Specifies the OSPF packet length in bytes
rid:	Provides the OSPF router ID
aid:	Shows the OSPF area ID
chk:	Displays the OSPF checksum
Aut:	Provides the OSPF authentication type: 0: No authentication; 1: Simple password; 2: MD5
auk:	Specifies the OSPF authentication key, if used
keyid	Displays the MD5 key ID; only used for MD5 authentication
seq	Provides the sequence number; only used for MD5 authentication

https://t.me/learningnets

The output of these **debug** commands is shown in the figure. Notice that the output shows the fields in the OSPF header, but they are not described in any detail. After each field, there is a parameter for it. The values are described in the table.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Network Types

OSPF defines distinct types of networks, based on their physical link types. OSPF operation on each type is different, including how adjacencies are established and which configuration is required.

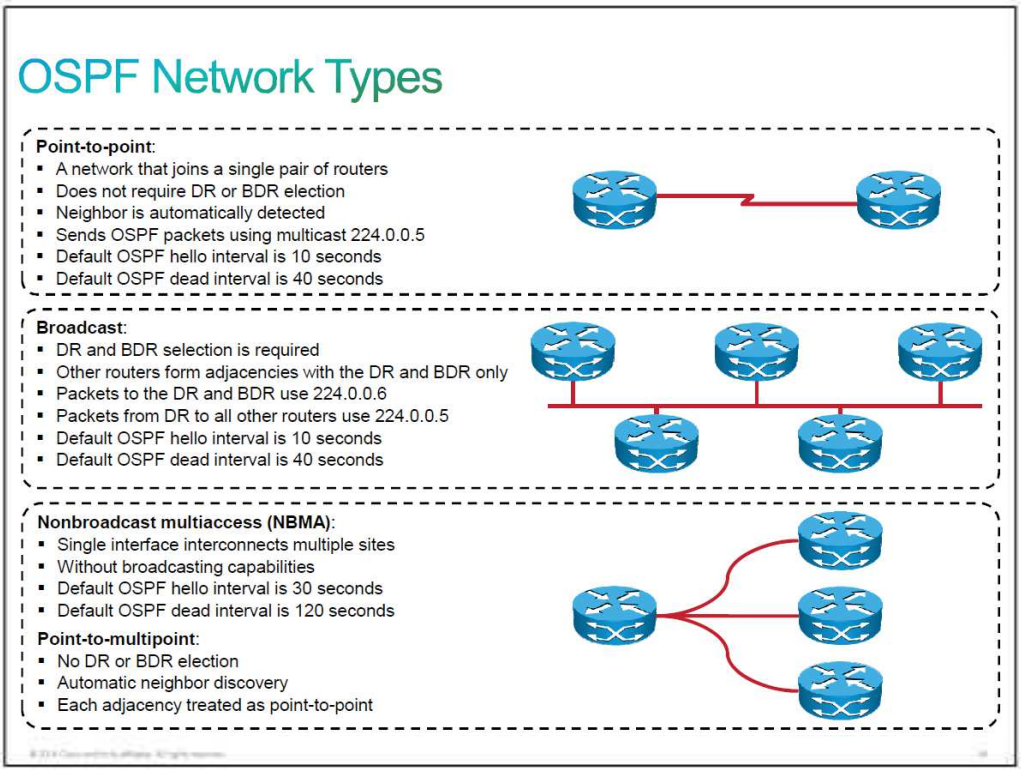

There are three types of networks that are defined by OSPF:

- **Point-to-point:** A network that joins a single pair of routers. A T1 serial line that is configured with a data link layer protocol such as PPP or HDLC is an example of a point-to-point network. On point-to-point networks, the router dynamically detects its neighboring routers by multicasting its hello packets to all OSPF routers, using the address 224.0.0.5. On point-to-point networks, neighboring routers become adjacent whenever they can communicate directly. No DR or BDR election is performed; there can be only two routers on a point-to-point link, so there is no need for a DR or BDR. Usually, the IP source address of an OSPF packet is set to the address of the outgoing interface on the router. The default OSPF hello and dead intervals on point-to-point links are 10 and 40 seconds, respectively.
- **Broadcast:** An OSPF router on a multiaccess broadcast network such as Ethernet forms an adjacency with its DR and BDR. The routers on a segment must elect a DR and a BDR to represent the multiaccess broadcast network. The BDR performs the DR tasks only if the DR fails. The DR and BDR improve network functioning in the following ways:
  - **Reducing routing update traffic:** The DR acts as central point of contact for link-state information exchange. The DR represents the multiaccess broadcast network in the sense that it sends link-state information from each router to all other routers in the network.
  - **Managing link-state synchronization:** The DR ensures that the other routers on the network have the same link-state information about the internetwork.

https://t.me/learningnets

- **NBMA:** A network that interconnects more than two routers but has no broadcast capability. When a single interface interconnects multiple sites over an NBMA network, the nonbroadcast nature of the network can lead to reachability issues. For example, if the NBMA topology is not fully meshed, then a broadcast or multicast that is sent by one router will not reach all the other routers. To implement broadcasting or multicasting in an NBMA network, the router replicates the packets to be broadcast or multicast and sends them individually on each PVC to all destinations. This process is CPU- and bandwidth-intensive. The default OSPF hello and dead intervals on NBMA interfaces are 30 and 120 seconds, respectively.
- **Point-to-multipoint:** The point-to-multipoint mode treats the nonbroadcast network as a collection of point-to-point links. In this environment, the routers automatically identify their neighboring routers but do not elect a DR and BDR. This configuration is typically used with partially meshed networks. Every router in the subnet generates a host (/32) route.

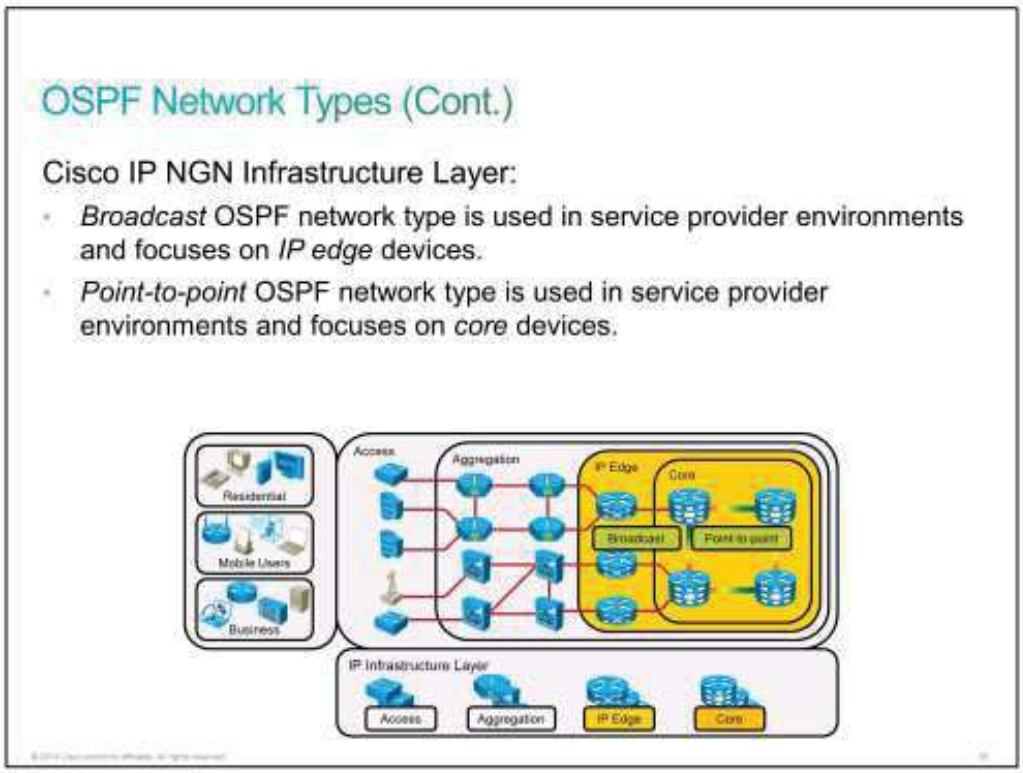

The figure shows usage of the OSPF network types in the service provider IP infrastructure layer of the Cisco IP Next-Generation Network.

https://t.me/learningnets

Electing the OSPF DR and BDR

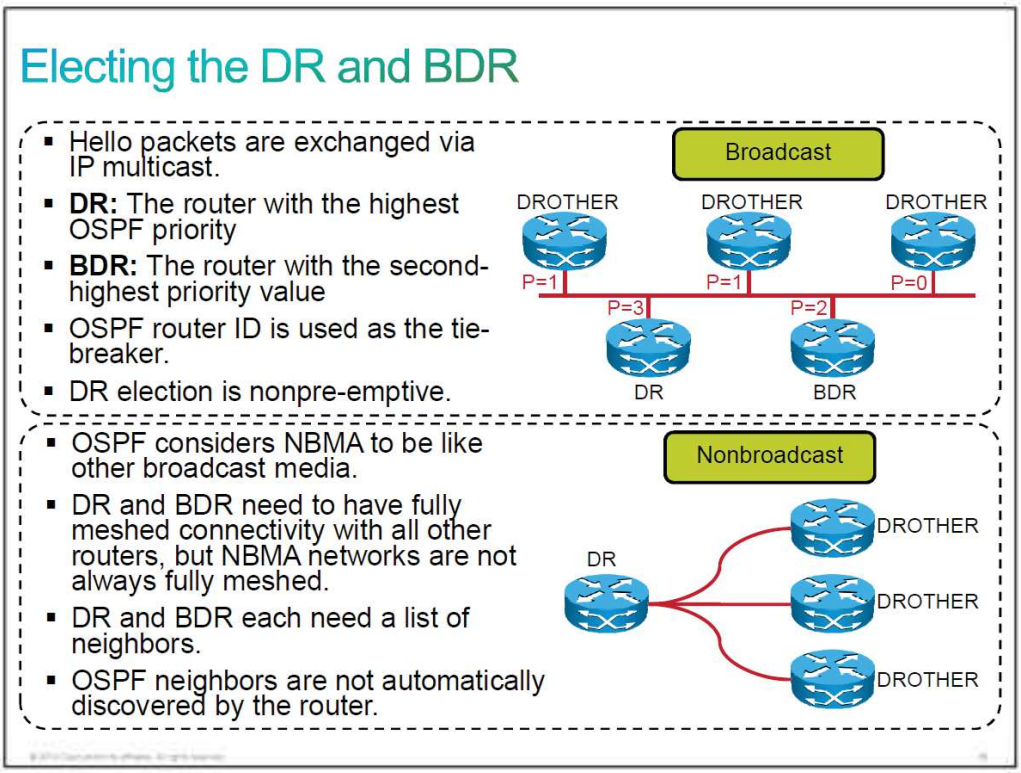

To elect a DR and BDR, the routers view the OSPF priority value of other routers during the hello packet exchange process and then use the following conditions to determine which router to select:

- The router with the highest priority value is the DR.
- The router with the second-highest priority value is the BDR.
- The default for the interface OSPF priority is 1. In the case of a tie, the router ID is used. The router with the highest router ID becomes the DR. The router with the second-highest router ID becomes the BDR.
- A router with a priority that is set to 0 cannot become the DR or BDR. A router that is not the DR or BDR is called a DROTHER.
- If a router with a higher priority value is added to the network, it does not preempt the DR and BDR. The only time that a DR or BDR changes is when one of them is out of service. If the DR is out of service, the BDR becomes the DR, and a new BDR is selected. If the BDR is out of service, a new BDR is elected.

To determine whether the DR is out of service, the BDR uses the wait timer. This timer is a reliability feature. If the BDR does not confirm that the DR is forwarding LSAs before the timer expires, the BDR assumes that the DR is out of service.

The highest IP address on an active interface is normally used as the router ID; however, you can override this selection by configuring an IP address on a loopback interface or using the Cisco IOS, IOS XE, or IOS XR **router-id** router configuration command. The DR operation is at the link level. A DR is selected for every multiaccess broadcast link in the OSPF network.

https://t.me/learningnets

OSPF treats NBMA environments like any other broadcast media environment; however, NBMA clouds are usually built in hub-and-spoke topologies. A hub-and-spoke topology means that the NBMA network is only a partial mesh. In these cases, the physical topology does not provide multiaccess capability on which OSPF relies.

The election of the DR becomes an issue in NBMA topologies because the DR and BDR need to have complete physical connectivity with all routers in the NBMA network. The DR and BDR also need to have a list of all the other routers so that they can establish adjacencies. OSPF cannot automatically build adjacencies with neighboring routers over NBMA interfaces.

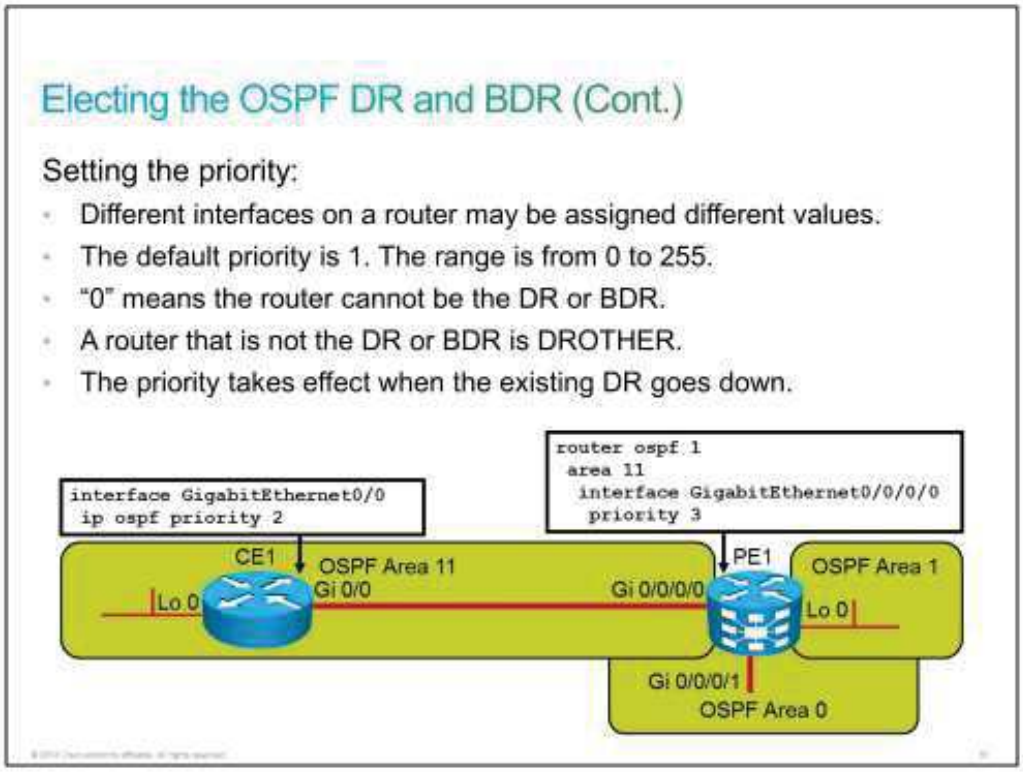

Use the Cisco IOS XR **priority** router command or Cisco IOS and IOS XE **ip ospf priority** interface command to designate which router interfaces on a multiaccess link are the DR and the BDR. The default priority is 1, and the range is from 0 to 255. The interface with the highest priority becomes the DR, and the interface with the second-highest priority becomes the BDR.

Interfaces that are set to a priority of 0 cannot be involved in the DR or BDR election process.

The priority of an interface takes effect only when the existing DR goes down. A DR does not relinquish its status just because a new interface reports a higher priority in its hello packet.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF over NBMA Network Types

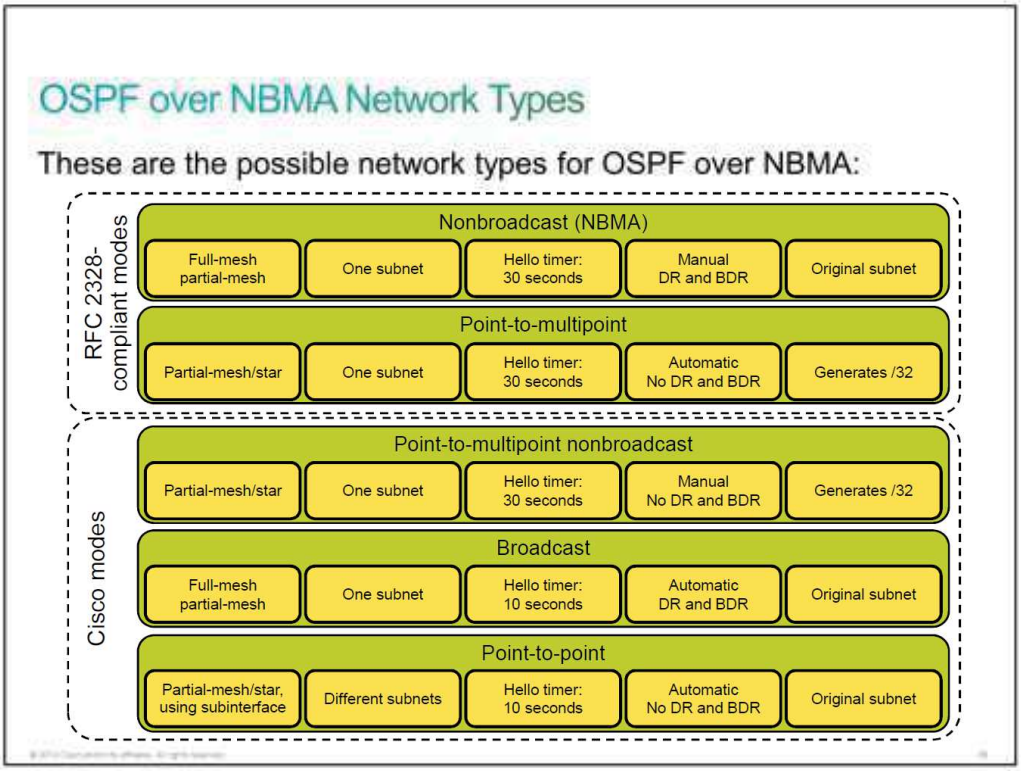

As described in RFC 2328, OSPF runs in one of these two official modes in NBMA topologies:

- **Nonbroadcast multiaccess:** The NBMA mode simulates the operation of OSPF in broadcast networks. Neighbors must be manually configured, and DR and BDR election is required. This configuration is typically used with fully meshed networks.
- **Point-to-multipoint:** The point-to-multipoint mode treats the nonbroadcast network as a collection of point-to-point links. In this environment, the routers automatically identify their neighboring routers but do not elect a DR and BDR. This configuration is typically used with partially meshed networks. Every router in the subnet generates a host (/32) route.

  The choice between NBMA and point-to-multipoint modes determines the way in which the Hello protocol and flooding work over the nonbroadcast network. The main advantage of the point-to-multipoint mode is that it requires less manual configuration, and the main advantage of the nonbroadcast mode is that there is less overhead traffic.

Cisco has defined the following additional modes:

- **Point-to-multipoint nonbroadcast:** Point-to-multipoint nonbroadcast mode is used in place of RFC-compliant, point-to-multipoint mode if multicast and broadcast are not enabled on the virtual circuits, because the router cannot dynamically discover its neighboring routers using hello multicast packets. This mode requires neighbors to be manually configured and does not require DR and BDR election. Every router in the subnet generates a host (/32) route.

https://t.me/learningnets

- **Broadcast:** Makes the WAN interface appear to be a LAN, uses one IP subnet, and OSPF hello packets uses multicast to automatically discover the neighbors. Routers elect the DR and BDR and require a full-mesh or a partial-mesh topology. Broadcast mode is a workaround for statically listing all existing neighboring routers. The interface is set to broadcast and behaves as though the router connects to a LAN. DR and BDR election is still performed; therefore, take special care to ensure either a full-mesh topology or a static election of the DR based on the interface priority.
- **Point-to-point:** Has a different IP subnet on each subinterface. It does not have DR or BDR election. Point-to-point mode is used when only two routers need to form an adjacency on a pair of interfaces. It can be used with either LAN or WAN interfaces.

OSPFv3 networks fall into four types by data link layer protocol:

- **NBMA:** If Frame Relay, ATM, or X.25 is adopted, OSPF defaults the network type to NBMA.
- **Point-to-multipoint:** OSPF does not default the network type for any data link layer protocol to point-to-multipoint.
- **Broadcast:** If Ethernet or FDDI is adopted, OSPF defaults the network type to broadcast.
- **Point-to-point:** If PPP or HDLC is adopted, OSPF defaults the network type to point-to-point.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Adjacency over Metro Ethernet and EoMPLS

EoMPLS is used to deploy Layer 2 MPLS VPNs, in which an MPLS backbone provides a Layer 2 Ethernet port-to-port connection. EoMPLS is also known as a type of Metro Ethernet service.

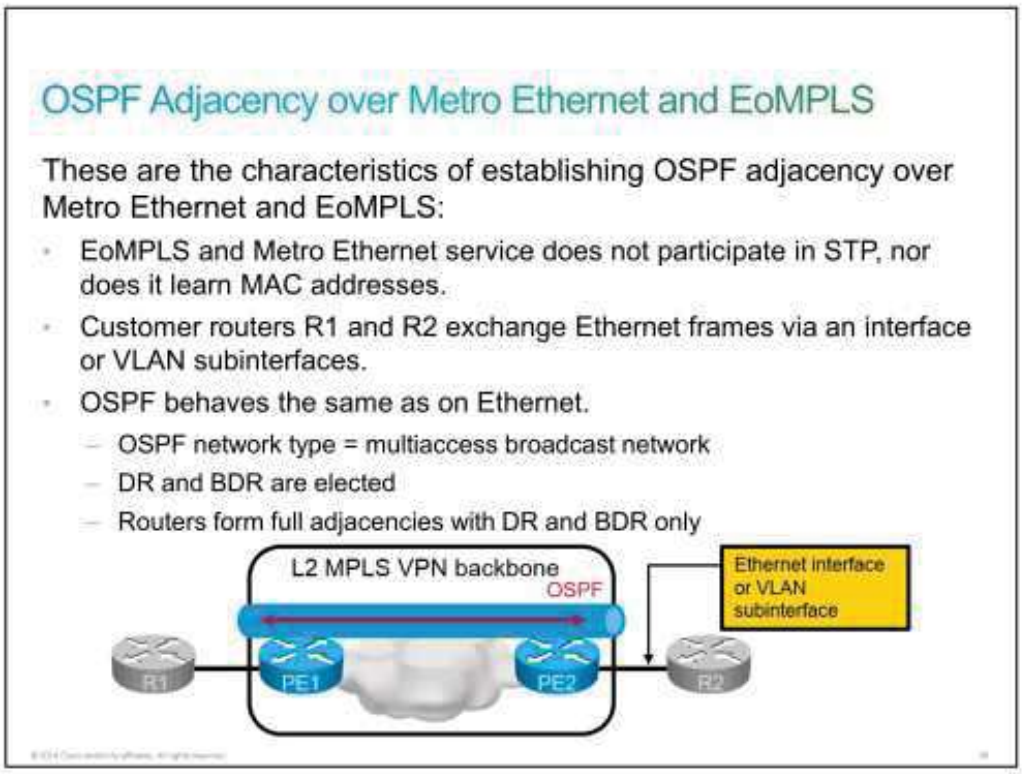

The MPLS backbone of the service provider is used to enable Layer 2 Ethernet connectivity between the customer routers R1 and R2, whether an EoMPLS or Metro Ethernet service is used.

R1 and R2 exchange Ethernet frames. PE router PE1 takes the Ethernet frames that are received from R1 on the link to PE1, encapsulates them into MPLS packets, and forwards them across the backbone to router PE2. PE2 de-encapsulates the MPLS packets and reproduces the Ethernet frames on the link toward R2. EoMPLS and Metro Ethernet typically do not participate in STP and BPDU exchanges, so EoMPLS and Metro Ethernet are transparent to the customer routers.

The Ethernet frames are transparently exchanged across the MPLS backbone. Customer routers can be connected either in a port-to-port or in a VLAN subinterface configuration. In a port-to-port connection, PE routers take whatever Ethernet frame is received and forward the frames across the Layer 2 MPLS VPN backbone. In a VLAN subinterface connection, frames for a particular VLAN—identified with a subinterface in the configuration—are encapsulated and sent across the Layer 2 MPLS VPN backbone.

When deploying OSPF over EoMPLS, there are no changes to the existing OSPF configuration from the customer perspective.

OSPF needs to be enabled, and network commands must include the interfaces that are required by the relevant OSPF area, to start the OSPF properly.

R1 and R2 form a neighbor relationship with each other over the Layer 2 MPLS VPN backbone. From an OSPF perspective, the Layer 2 MPLS VPN backbone, PE1, and PE2 are invisible.

https://t.me/learningnets

A neighbor relationship is established between R1 and R2 directly, and it behaves in the same way as on a regular Ethernet broadcast network.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Adjacency over Layer 3 MPLS VPN

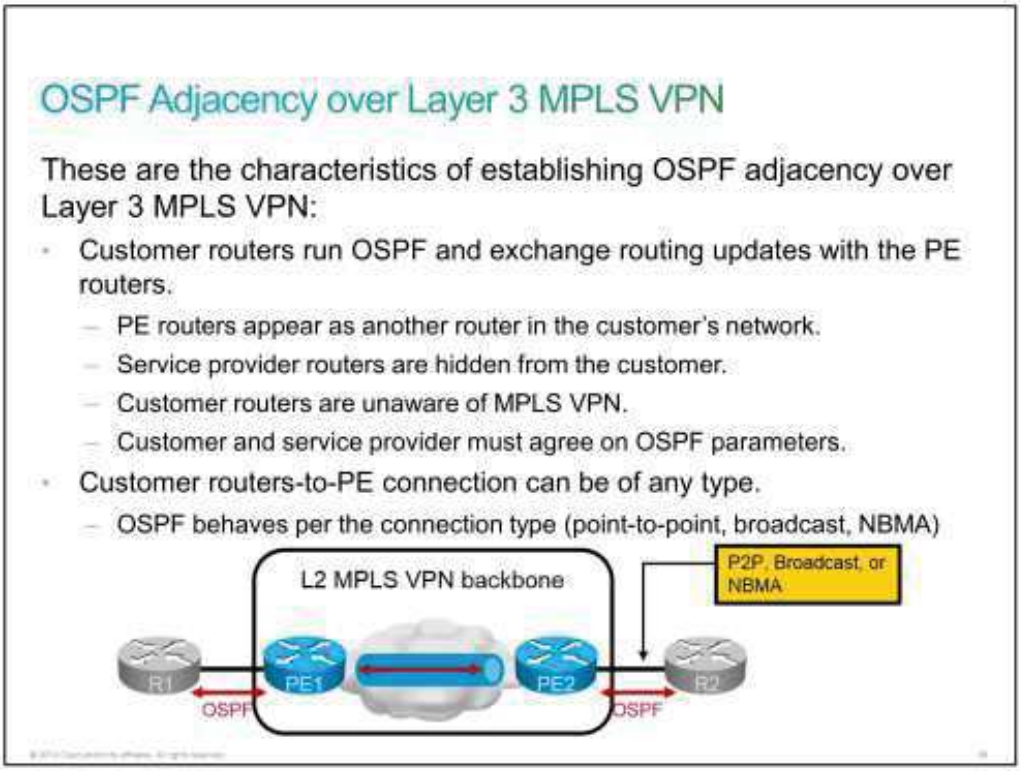

With the Layer 3 MPLS VPN architecture, the ISP provides a peer-to-peer VPN architecture. In this architecture, PE routers participate in customer routing, guaranteeing optimum routing between customer sites. Therefore, the PE routers carry a separate set of routes for each customer, resulting in perfect isolation between the customers. The following applies to the Layer 3 MPLS VPN technology, even when running OSPF as a PE-CE routing protocol:

- The customer routers should not be aware of MPLS VPN; they should run standard IP routing software.
- The P routers must not carry VPN routes for the MPLS VPN solution to be scalable.
- The PE routers must support MPLS VPN services and traditional Internet services.

To OSPF, the Layer 3 MPLS VPN backbone looks like a standard corporate backbone and it runs standard IP routing software. Routing updates are exchanged between the customer routers and the PE routers that appear as normal routers in the customer network. The standard design rules that are used for enterprise Layer 3 MPLS VPN backbones can be applied to the design of the customer network. The P routers are hidden from the customer view, and customer edge routers (CE routers) are unaware of MPLS VPN. Therefore, the internal topology of the Layer 3 MPLS backbone is totally transparent to the customer. The PE routers receive IPv4 routing updates from the CE routers and install them in the appropriate virtual VRF table. This part of the configuration and operation is the responsibility of a service provider.

The PE-CE can have any OSPF network type—point-to-point, broadcast, or even nonbroadcast multiaccess.

The only difference between a PE-CE design and a regular OSPF design is that the customer has to agree with the service provider about the OSPF parameters (AS number, authentication password, and so on); usually, these parameters are governed by the service provider.

https://t.me/learningnets

Enabling OSPF on a Link with NBMA Mode

At the beginning, one or more OSPF routing processes must be enabled on the router, followed by configuration that defines which interfaces are involved in OSPF routing.

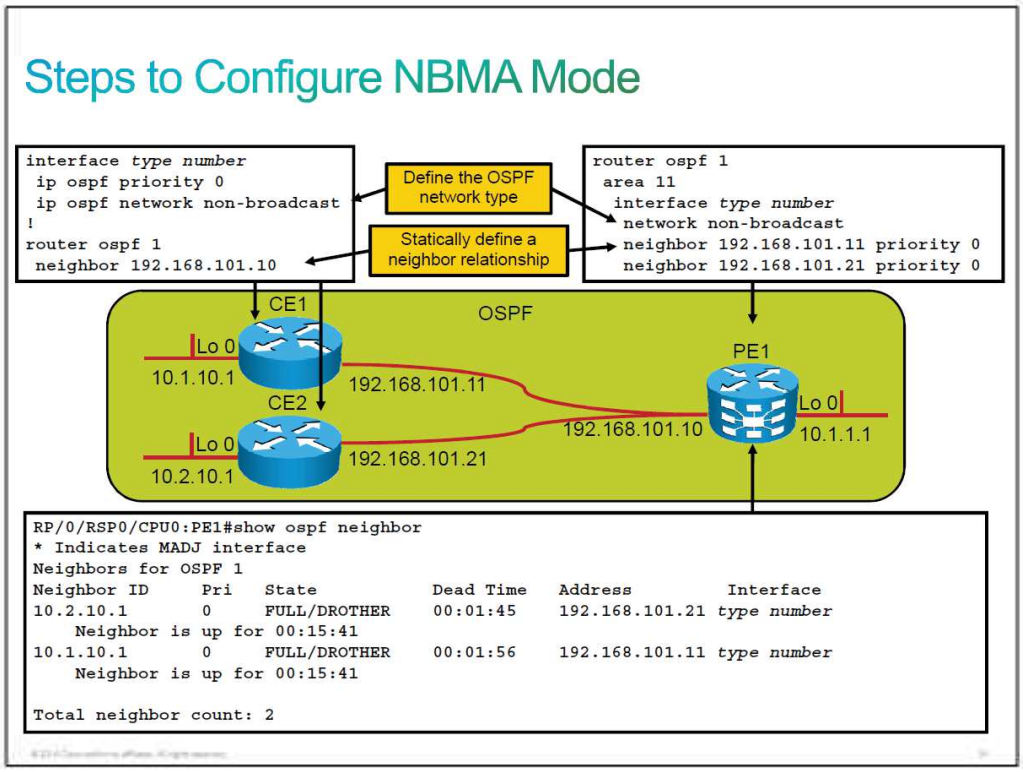

When all the required information is defined and basic OSPF configuration is applied, an implementation plan showing the following tasks is required to configure the NBMA-specific configuration:

- **Manual configuration of the neighbors:** To configure OSPF routers that are interconnecting to nonbroadcast networks, use the Cisco IOS, IOS XE, or IOS XR **neighbor** command in OSPF router configuration mode. The **neighbor** command in the figure shows the configuration when the **priority 0** optional keyword is used. A number indicates the router priority value of the nonbroadcast neighbor that is associated with the specified IP address. The default value is 0. If this value is used, the specified neighbor cannot become a DR or BDR. Note that this configuration is performed on one router for its neighbor. In other words, this example specifies on router PE1 that the router CE1 priority should be 0. Nevertheless, a priority value may already be set on CE1. If so, the priority that is defined locally on CE1 will override the wished priority that is defined on PE1 for CE1, and the **priority** command will have no effect.
- **Definition of the OSPF network type:** To configure the OSPF network type to a type other than the default for a given medium, use the Cisco IOS XR **network** OSPF command or Cisco IOS and IOS XE **ip ospf network** interface command. The default OSPF network type on the Gigabit Ethernet interface is broadcast. The **nonbroadcast** keyword in the figure shows that the network type that is selected is NBMA.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

It is important that both adjacent routers agree on the media type. If one router considers the media as broadcast and the other as nonbroadcast, adjacency may not form because hello intervals are not the same on broadcast and nonbroadcast segments. On broadcast segments, hello messages are sent every 10 seconds. On nonbroadcast segments, hello intervals are sent every 30 seconds. The dead interval, which determines the interval during which at least one hello packet must be received from a neighbor before the router declares the neighbor down, also depends on the hello interval (four times the hello interval by default). Neighboring routers need to have the same interval values to form an adjacency. You can manually adjust the intervals as follows:

- Cisco IOS XR **hello-interval** router OSPF command or Cisco IOS and IOS XE **ip ospf hello-interval** interface command adjusts the OSPF hello interval on this interface.
- Cisco IOS XR **dead-interval** router OSPF command or Cisco IOS and IOS XE **ip ospf dead-interval** interface command adjusts the OSPF dead interval on this interface.

To display OSPF neighbor information on a per-interface basis, you must use the Cisco IOS XR **show ospf neighbor** or Cisco IOS and IOS XE **show ip ospf neighbor** command. The output on the figure shows OSPF neighbors, as seen from Cisco IOS XR router, output that is taken from Cisco IOS and IOS XE router is as follows:

```
CE1# show ip ospf neighbor
Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.1          1   FULL/DR         00:01:41    192.168.101.10  GigabitEthernet0/0
```


https://t.me/learningnets

Subinterfaces over Multiaccess WAN Networks

A physical interface can be split into multiple logical interfaces, called subinterfaces. Subinterfaces were originally created to better handle issues that are caused by split horizon over an NBMA for distance vector-based routing protocols.

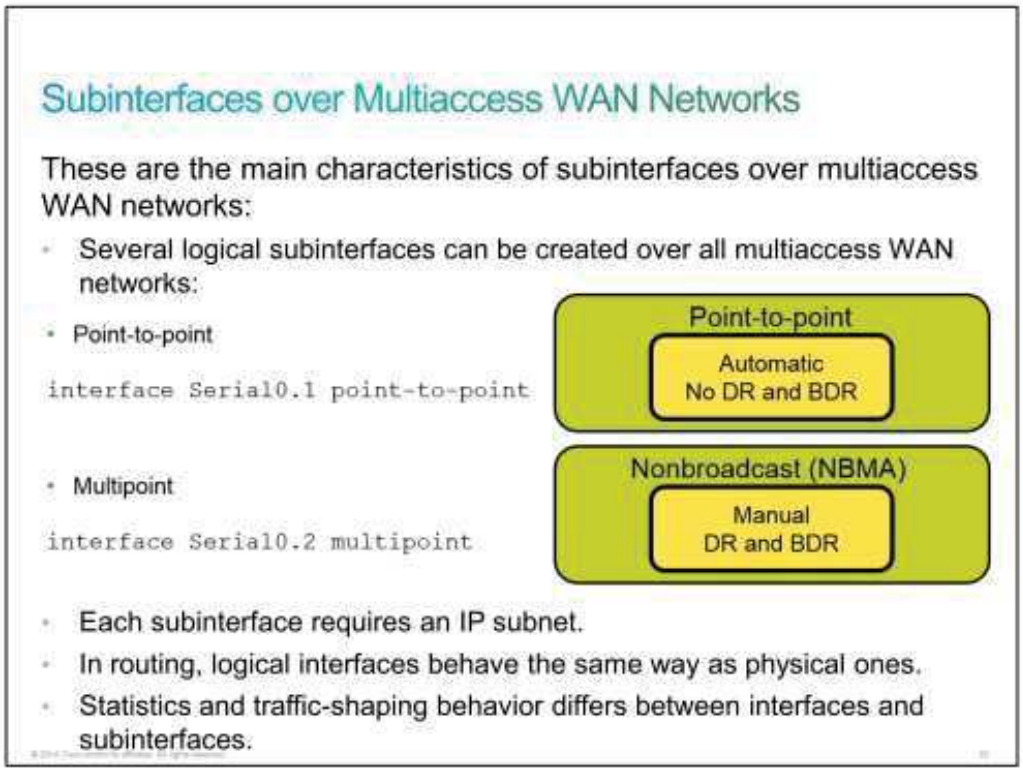

Each subinterface is defined as a point-to-point or a point-to-multipoint interface. During the configuration of subinterfaces, you must choose the Cisco IOS and IOS XE or IOS XR **point-to-point** or **multipoint** keywords. The choice of modes affects the operation of OSPF.

The default OSPF mode on a point-to-point Frame Relay subinterface is the point-to-point mode; the default OSPF mode on a Frame Relay point-to-multipoint subinterface is the nonbroadcast mode. The default OSPF mode on a main Frame Relay interface is also the nonbroadcast mode.

When point-to-point subinterfaces are configured, each PVC gets its own subinterface, and PVCs are treated like point-to-point links. A point-to-point subinterface has the properties of any physical point-to-point interface and requires its own subnet. There is no DR or BDR election process. Neighbor discovery is automatic, so neighbors do not need to be configured. To configure the point-to-point interface, you must select the **point-to-point** keyword during subinterface configuration.

When multipoint subinterfaces are configured, there are multiple PVCs on a single subinterface. Each subinterface requires an IP subnet. Multipoint Frame Relay subinterfaces default to OSPF nonbroadcast mode, which requires neighbors to be statically configured and performs a DR and BDR election. To configure the multipoint interface, you must select the **multipoint** keyword during subinterface configuration.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Enabling OSPF on a Link with Point-to-Point Mode

Networks in point-to-point mode are designed to emulate leased-line networks. This is a Cisco proprietary mode. Because they are treated as point-to-point links, the DR and BDR are not used, and the adjacency forms over the point-to-point network with no DR or BDR election. Only a single subnet is used per point-to-point link.

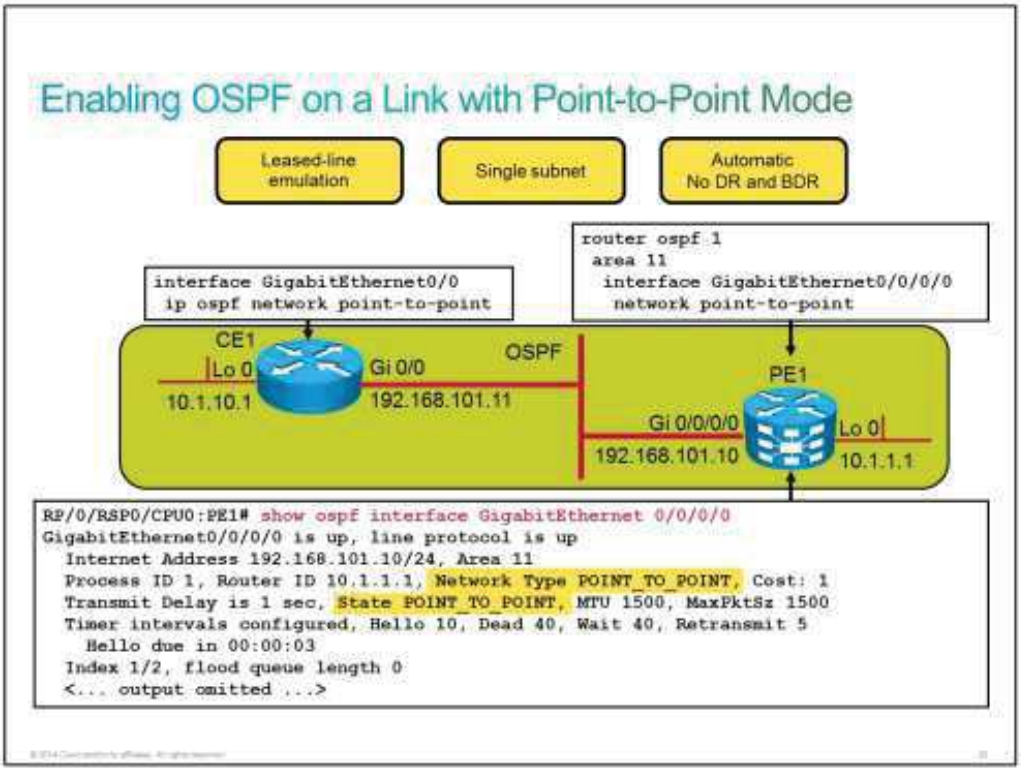

The figure shows partial configuration of CE1 and PE1 in point-to-point mode. This configuration does not require subinterfaces and uses only a single subnet.

The Cisco IOS XR **show ospf interface** command or Cisco IOS and IOS XE **show ip ospf interface** command displays the key OSPF details for each interface. The OSPF network type, area number, cost, and state of the interface are all displayed. The point-to-point and broadcast modes default to a 10-second hello timer. The hello and dead timers on the neighboring interfaces must match for the neighbors to form successful adjacencies. The listed adjacent neighboring routers are all dynamically learned. Manual configuration of neighboring routers is not necessary.

https://t.me/learningnets

Enabling OSPF on a Link with Point-to-Multipoint Mode

Networks in point-to-multipoint mode are designed to work with partial-mesh or star topologies. With RFC 2328-compliant point-to-multipoint mode, OSPF treats all router-to-router connections over the nonbroadcast network as if they were point-to-point links. In point-to-multipoint mode, DRs are not used, and a type 2 LSA is not flooded to adjacent routers. Instead, OSPF point-to-multipoint works by exchanging additional LSUs that are designed to automatically discover neighboring routers and add them to the neighbor table.

In large networks, using point-to-multipoint mode reduces the number of PVCs that are required for complete connectivity, because you are not required to have a full-mesh topology. In addition, not having a full-mesh topology reduces the number of neighbor entries in the neighbor table. Point-to-multipoint mode has the following properties:

- **Does not require a fully meshed network:** This environment allows routing to occur between two routers that are not directly connected but are connected through a router that has virtual circuits to each of the two routers.
- **Does not require a static neighbor configuration:** In nonbroadcast mode, neighboring routers are statically defined to start the DR election process and allow the exchange of routing updates. Because the point-to-multipoint mode treats the network as a collection of point-to-point links, multicast hello packets discover neighboring routers dynamically. Statically configuring neighboring routers is not necessary.
- **Uses one IP subnet:** As in nonbroadcast mode, when you are using point-to-multipoint mode, all routers are on one IP subnet.
- **Duplicates LSA packets:** Also as in nonbroadcast mode, when flooding out a nonbroadcast interface in point-to-multipoint mode, the router must replicate the LSU. The LSU packet is sent to each of the neighboring routers of the interface, as defined in the neighbor table.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

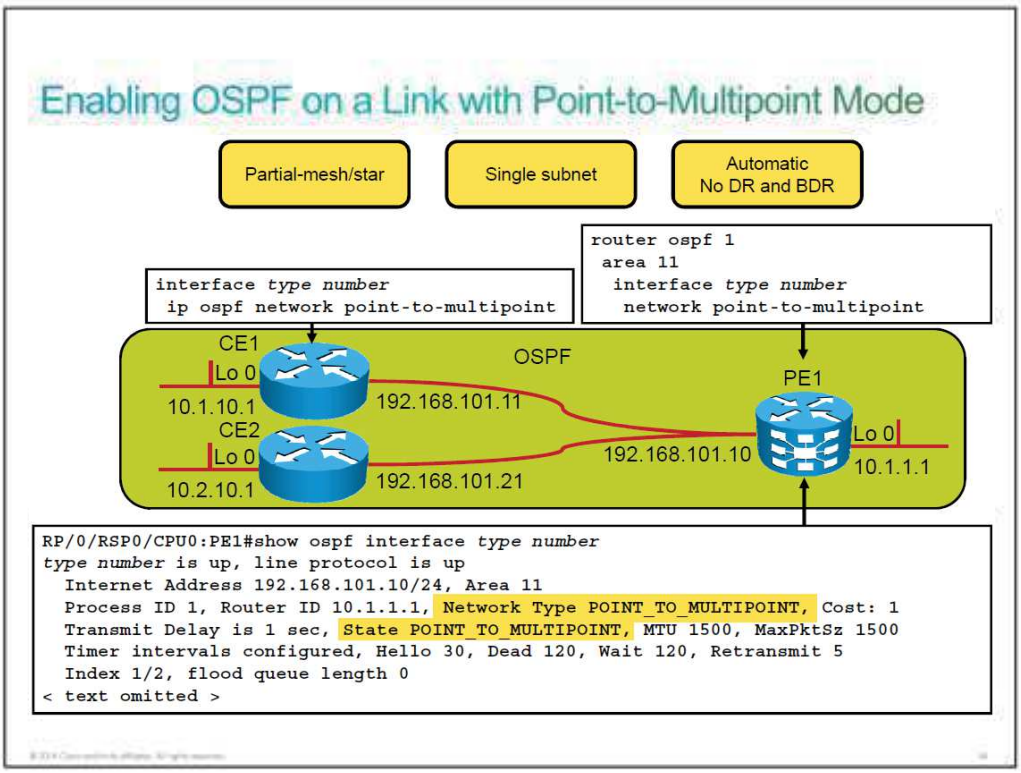

The figure shows partial configurations of CE1 and PE1 in point-to-multipoint mode. This configuration does not require subinterfaces and uses only a single subnet. The CE2 configuration would be very similar to the configuration that is displayed for CE1. In point-to-multipoint mode, a DR or BDR is not used; therefore, DR and BDR election and priorities are not a concern.

The hello interval for a point-to-multipoint interface is 30 seconds, with a dead interval of 120 seconds.

https://t.me/learningnets

Enabling OSPF on a Link with Point-to-Multipoint Nonbroadcast Mode

Cisco defines additional modes for the OSPF neighbor relationship, including point-to-multipoint nonbroadcast. This mode is a Cisco extension of the RFC-compliant point-to-multipoint mode. You must statically define neighbors, and you can modify the cost of the link to the neighboring router to reflect the different bandwidths of each link. The RFC point-to-multipoint mode was developed to support underlying point-to-multipoint virtual circuits that support multicast and broadcast; therefore, this mode allows dynamic neighboring router discovery. For this reason, no DR or BDR is used.

This mode is used in special cases where neighbors cannot be automatically discovered. If, for example, multicast and broadcast are not enabled on the virtual circuits, the RFC-compliant point-to-multipoint mode cannot be used, because the router cannot dynamically discover its neighboring routers using the hello multicast packets; this Cisco mode should be used instead.

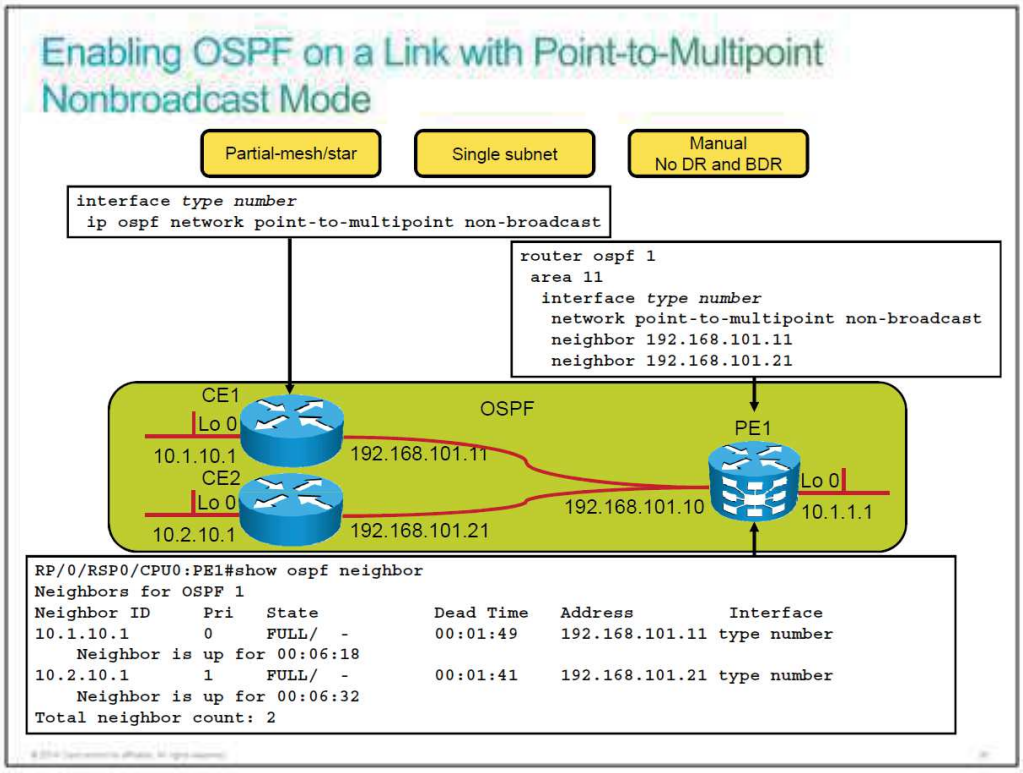

To define point-to-multipoint nonbroadcast mode, you must use the Cisco IOS XR **network point-to-multipoint nonbroadcast** router OSPF command or Cisco IOS and IOS XE IP **OSPF network point-to-multipoint nonbroadcast** interface command.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

IPv6 Support for OSPF Modes

To configure the OSPF network type to a type other than the default for a given medium, use the Cisco IOS XR **network** router OSPF command or Cisco IOS and IOS XE **IPv6 OSPF network** interface command.

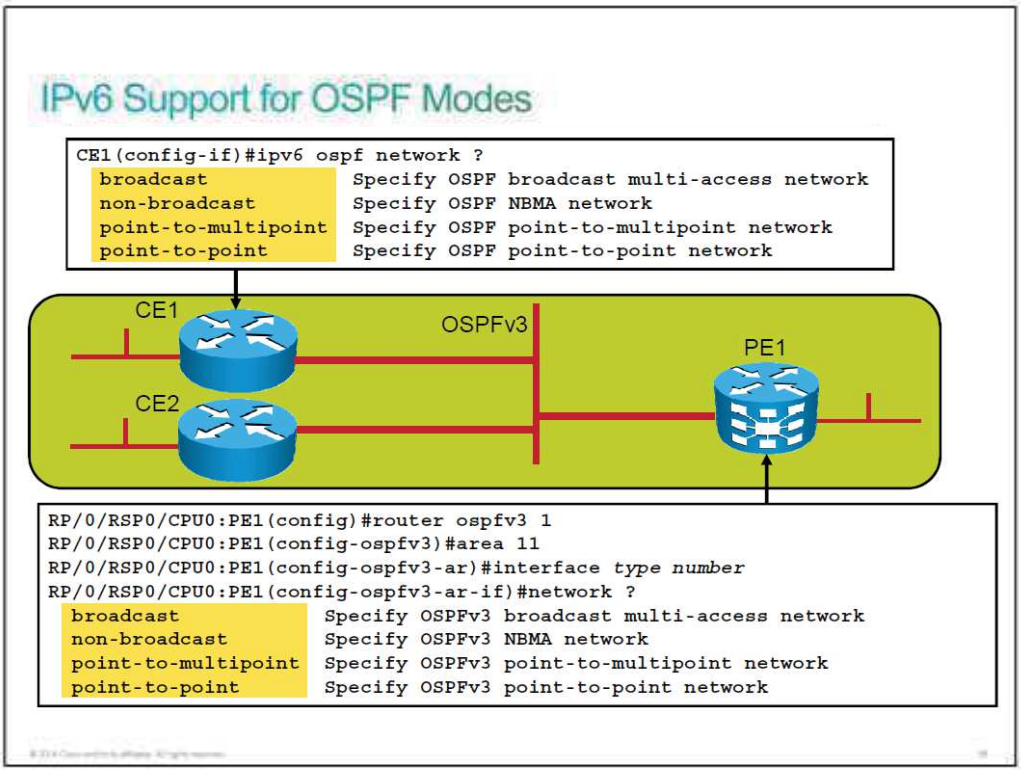

The example in the figure shows only OSPF network type configuration. There must be OSPF version 3 already enabled.

https://t.me/learningnets

Summary

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

Summary

- OSPF uses several types of packets to operate properly.
- OSPF packets are encapsulated directly into an IP payload.
- There are five OSPF packet types: hello, DBD, LSU, LSR, and LSAck.
- The Hello protocol forms logical neighbor adjacency relationships.
- A topology change triggers link-state flooding.
- You can use the **debug ospf packets** Cisco IOS XR command to debug OSPF packets.
- Cisco routers support two RFC-compliant OSPF network types: NBMA and point-to-multipoint; and three Cisco defined OSPF network types: point-to-multipoint nonbroadcast, broadcast, and point-to-point.
- A router with the highest priority becomes a DR. In case of a tie, a router with the highest router ID becomes a DR.

Summary (Cont.)

- It is important that adjacent routers agree on the media type. OSPF network type selection depends on the topology, available subnets, and multicast or broadcast supported.
- OSPF over Metro Ethernet and EoMPLS behaves the same as on the Ethernet.
- OSPF over MPLS VPN behaves as per the connection type between a CE and PE router.
- Both adjacent routers must agree on a the media type. Otherwise, adjacencies may not form due to different hello intervals.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Implementing OSPF Routing

Overview

This lesson discusses the primary configuration commands for a single-area and multiarea OSPF protocol. OSPF design limitations and solutions, such as virtual links, the passive interface, and changing the cost metric, are also described. The lesson concludes with how to secure the OSPF routing protocol by configuring OSPF authentication.

Upon completing this lesson, you will be able to:

- Describe how to implement OSPF in the service provider network
- Describe the importance of the OSPF router ID and how to configure it
- Describe how to configure an OSPF passive interface
- Describe how to verify a basic OSPF configuration
- Describe OSPF virtual links
- Describe how to configure OSPF virtual links
- Describe how to configure OSPF cost
- Describe Cisco Nonstop Forwarding (NSF)
- Describe Cisco Nonstop Routing (NSR)
- Describe NSF and NSR for OSPF
- Describe the OSPFv3 graceful restart feature
- Describe BFD
- Describe how to configure BFD for OSPF
- Describe OSPF authentication in the service provider network

https://t.me/learningnets

- Describe how to configure OSPF authentication

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Implement OSPF

There are two types of routers from the configuration point of view:

- **Routers with single-area configuration:** internal routers, backbone routers, and ASBRs that are residing in one area
- **Routers with a multiarea configuration:** ABRs and ASBRs that are residing in more than one area

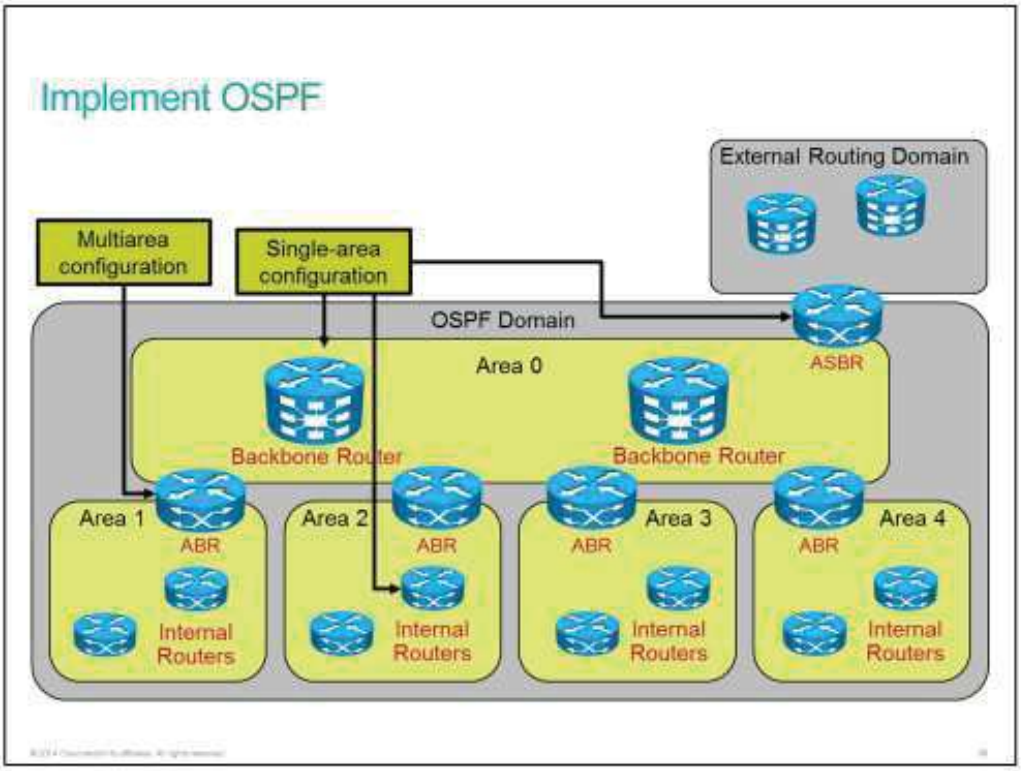

Single-area design in OSPF puts all routers into a single OSPF area. This design results in many updates being processed on every router and in larger routing tables. The OSPF configuration follows a single-area design, in which all the routers are treated as being internal routers to the area and all the interfaces are members of that single area. The best single-area design should use a single area as the backbone area.

Multiarea design is a better solution than single-area design. In a multiarea design, the network is segmented to limit the propagation of LSAs and to make routing tables smaller.

Note	The OSPFv2 Multiarea Adjacency feature allows you to configure a link on the primary interface to enable optimized routing in multiple areas. Each multiarea interface is announced as a point-to-point unnumbered link. The multiarea interface exists as a logical construct over an existing primary interface. The neighbor state on the primary interface is independent of the neighbor state of the multiarea interface. The multiarea interface establishes a neighbor relationship with the corresponding multiarea interface on the neighboring device. You can only configure multiarea adjacency on an interface that has two OSPF speakers. In case of native broadcast networks, the interface must be configured as an OSPF point-to-point type to enable the interface for multiarea adjacency.

https://t.me/learningnets

Implement OSPF (Cont.)

Steps to enable OSPF:

- Create an implementation plan:
  - IP addressing
  - Areas and area types
  - ABRs and ASBRs
- Define summarization and redistribution points.
- Configure OSPF routing processes on every OSPF router.

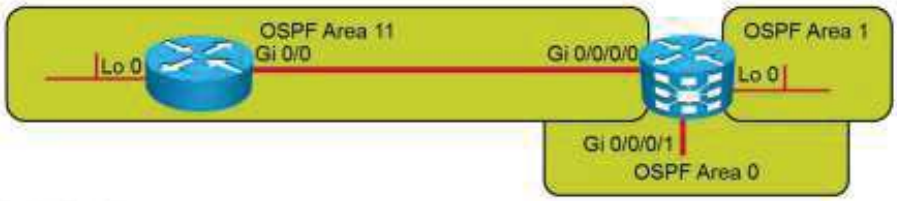

The type of OSPF routing protocol implementation that you should choose depends on your specific needs and topology. When you prepare to deploy OSPF routing in a network, you must first gather the existing state and requirements and then consider different deployment options:

- The IP addressing plan governs how OSPF can be deployed and how well the OSPF deployment might scale. Thus, a detailed IP addressing plan, along with the IP subnetting information, must be created. A solid IP addressing plan should enable usage of OSPF area design and summarization to more easily scale the network, as well as optimize OSPF behavior and propagation of LSAs.
- The network topology consists of links that connect the network equipment (routers, switches, and so on) and belong to different OSPF areas in a multiarea OSPF design. A detailed network topology plan should be presented to assess the OSPF scalability requirements and determine the different OSPF areas, ABRs, and ASBRs as well as summarization and redistribution points.
- An implementation plan must be created before configuring OSPF routing in the network.

After you gather all the required information, you must create an implementation plan that includes the following tasks to perform basic OSPF configuration:

- Define one or more OSPF processes globally on the router.
- Define the interfaces that OSPF will run on.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF for Multiple Areas

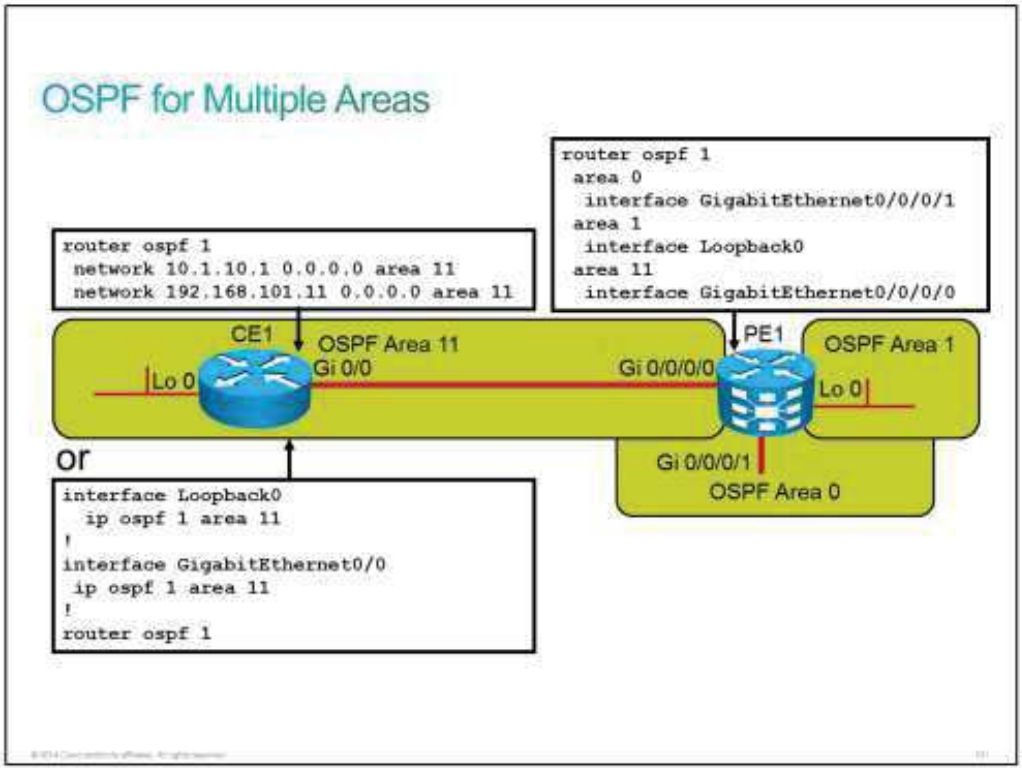

To configure the OSPF process, complete these tasks:

- Enable the OSPF process on the router.
- Identify which interfaces on the router are part of the OSPF process.

To configure an OSPF routing process, use the Cisco IOS, IOS XE, and IOS XR **router ospf** command. The *process-id* parameter inside the **router ospf** command is an internally used identification parameter for the OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process in the router.

To define the interfaces on which OSPF runs and define the area IDs for those interfaces, use the Cisco IOS XR **interface** router OSPF command or Cisco IOS and IOS XE **network area** router OSPF command. The Cisco IOS and IOS XE **wildcard-mask** parameter in the **network area** command determines how to interpret the IP address. The mask has wildcard bits, in which 0 is a match, and 1 indicates that the value is not significant. For example, 0.0.255.255 indicates a match in the first two octets.

Starting with Cisco IOS Release 12.3(11)T (and some specific versions of earlier releases) and all Cisco IOS XE Releases, OSPF can be enabled directly on the interface using the Cisco IOS and IOS XE **ip ospf area** command, which simplifies the configuration of unnumbered interfaces. Because the command is configured explicitly for the interface, it takes precedence over the **network area** command. The command is shown in the figure as well.

The figure shows OSPF configuration on CE1 and PE1 routers. The CE1 router is internal and belongs only to OSPF area 11, and the PE1 router is an ABR and connects two nonbackbone areas (1 and 11) to the backbone area.

https://t.me/learningnets

OSPF Router ID

The OSPF database uses the OSPF router ID to uniquely describe each router in the network. Every router keeps a complete topology database of all routers and links in an area or network; therefore, each router should have a unique router ID.

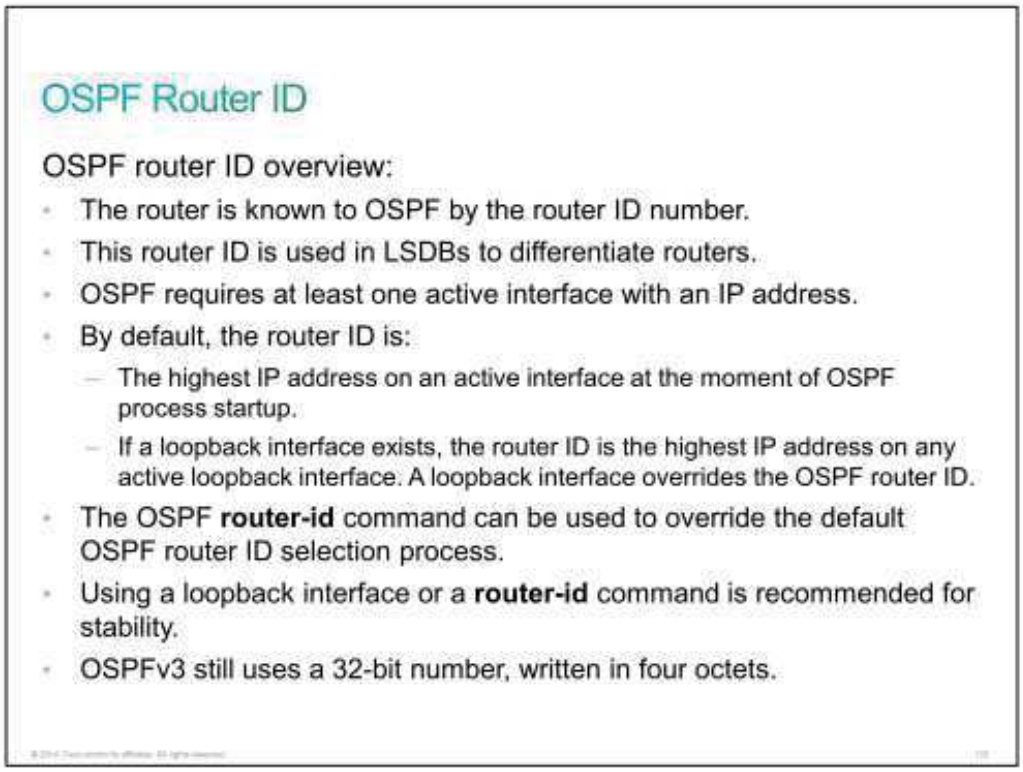

The OSPF routing process chooses a router ID for itself when it starts up. The router ID is a unique IP address that can be assigned in these ways:

- By default, the highest IP address of any active physical interface when OSPF starts is chosen as the router ID. The interface does not have to be part of the OSPF process, but it must be up. There must be at least one IP interface that is up on the router for OSPF to use as a router ID. If no interface is up and available with an IP address when the OSPF process starts, OSPF process will not start.
- If there is a loopback interface, its address will always be preferred as the router ID over a physical interface address, because a loopback interface never goes down. If there is more than one loopback interface, the highest IP address on any active loopback interface becomes the router ID.
- Using the Cisco IOS, IOS XE, or IOS XR **router-id** command is the preferred procedure to set the router ID and is always used in preference to the other two procedures.

When the OSPF router ID is set, it does not change, even if the interface that the router is using for the router ID goes down. The OSPF router ID changes only if the router reloads or if the OSPF routing process restarts.

In OSPF version 3, the OSPF process uses a 32-bit number. This 32-bit number is entered as four octets that are separated by dots and looks like an IP address.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

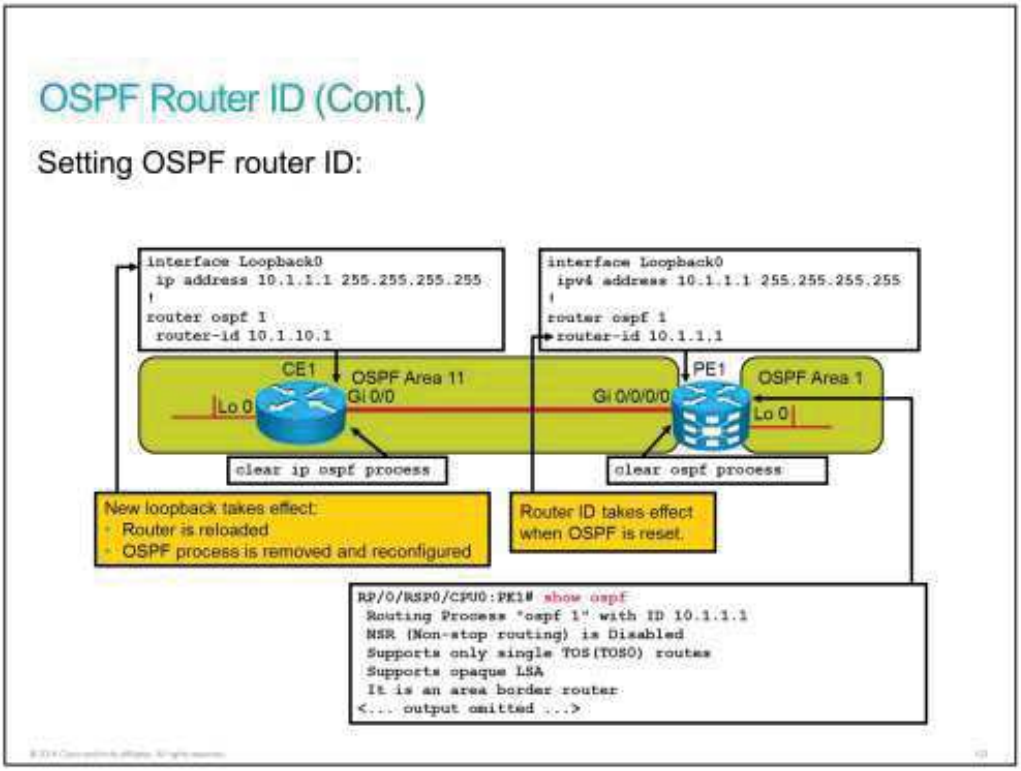

To modify the OSPF router ID to a loopback address, define a loopback interface. Configuring an IP address on a loopback interface overrides the highest IP address that is being used as the router ID. OSPF is more reliable if a loopback interface is configured, because the interface is always active and cannot fail, as opposed to a real interface. For this reason, you should use a loopback address on all key routers. If the loopback address is advertised into OSPF, this address can be pinged for testing purposes. A private IP address can be used to save registered public IP addresses.

If the OSPF process is already running, the router must be reloaded or the OSPF process must be removed and reconfigured before the new loopback address will take effect.

Use the Cisco IOS, IOS XE, and IOS XR **router-id** router OSPF command to ensure that OSPF uses a specific router ID. After the router ID is configured, use the Cisco IOS XR **clear ospf process** command or Cisco IOS and IOS XE **clear ip ospf process** command. This command restarts the OSPF routing process so that it will reselect the new IP address as its router ID. Restarting the OSPF process will temporarily disrupt an operational network. Router IDs must be unique throughout the autonomous system, regardless of how they are configured.

Use the Cisco IOS XR **show ospf** command or Cisco IOS and IOS XE **show ip ospf** command to verify the OSPF router ID. The figure shows partial output from this command on a Cisco IOS XR router.

A useful feature for all routing protocols in Cisco IOS XR is the global **router-id** command. You can override the global **router-id** command in Cisco IOS XR by further configuring a **router-id** command within a given protocol. However, configuring different router IDs per protocol makes management more complicated and provides no gain, so it is not recommended.

https://t.me/learningnets

OSPF Passive Interface

To prevent other routers on a local network from dynamically learning about routes, you can keep routing update messages from being sent through a router interface. In OSPF, the interface address that you specify as passive appears as a stub network in the OSPF domain. OSPF routing information is neither sent nor received through the specified router interface.

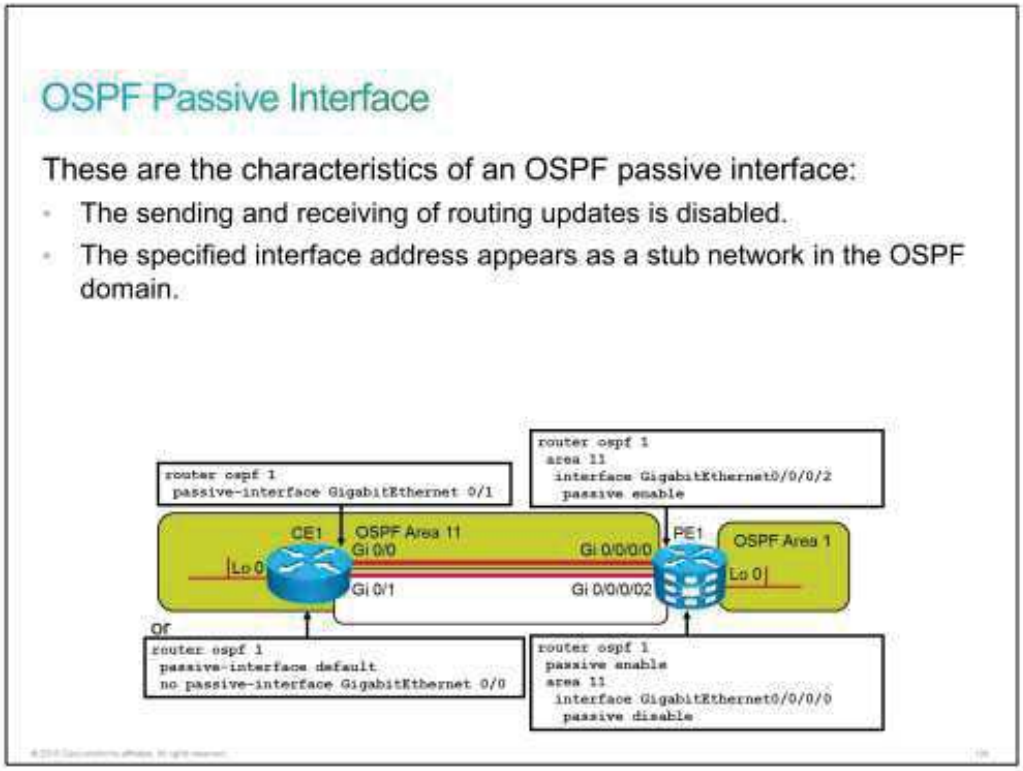

To disable the sending and receiving of OSPF routing updates on an interface, use the Cisco IOS XR **passive enable** router OSPF command or Cisco IOS and IOS XE **passive-interface** router OSPF command.

In service provider and large enterprise networks, many of the distribution routers have more than 200 interfaces. Therefore, many LSAs can be flooded over the domain. The OSPF routing protocol can be configured on all interfaces, and the passive interface feature can be enabled on the interfaces where adjacency is not desired. In some networks, this means the coding of 100 or more passive interface statements. With the default passive interface feature, this problem is solved by allowing all interfaces to be set as passive, by default, using a single Cisco IOS XR **passive enable** router OSPF command or Cisco IOS and IOS XE **passive-interface default** router OSPF command, then configuring individual interfaces where adjacencies are desired, using the Cisco IOS XR **passive disable** router OSPF command or Cisco IOS and IOS XE **no passive-interface** command.

In the figure, routers are configured for the OSPF routing protocol. Between CE1 and PE1 routers, there are two links, but OSPF adjacency should be established only on the first Gigabit Ethernetinterface.

https://t.me/learningnets

Verifying Basic OSPF

After the basic OSPF is configured, you can use several verification commands to verify its operation. One of the first steps is to verify the presence of the OSPF routing process. If the OSPF process has been configured successfully, there must be an OSPF router inside the IP routing table. Additionally, you can verify the OSPF interfaces, routing process neighbors, and database.

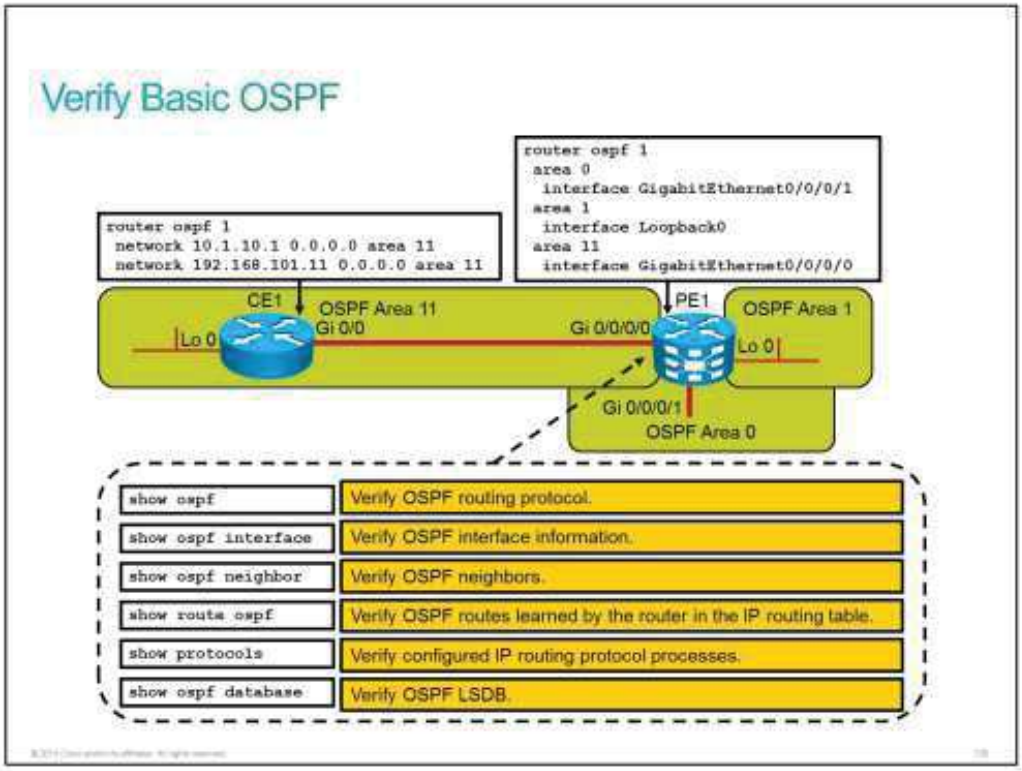

Commands that are shown in the figure are Cisco IOS XR verification commands. Cisco IOS and IOS XE verification commands are as follows: **show ip ospf**, **show ip ospf interface**, **show ip ospf neighbor**, **show ip route ospf**, **show ip protocols**, and **show ip ospf database**.

https://t.me/learningnets

Verify OSPF Routing Protocol

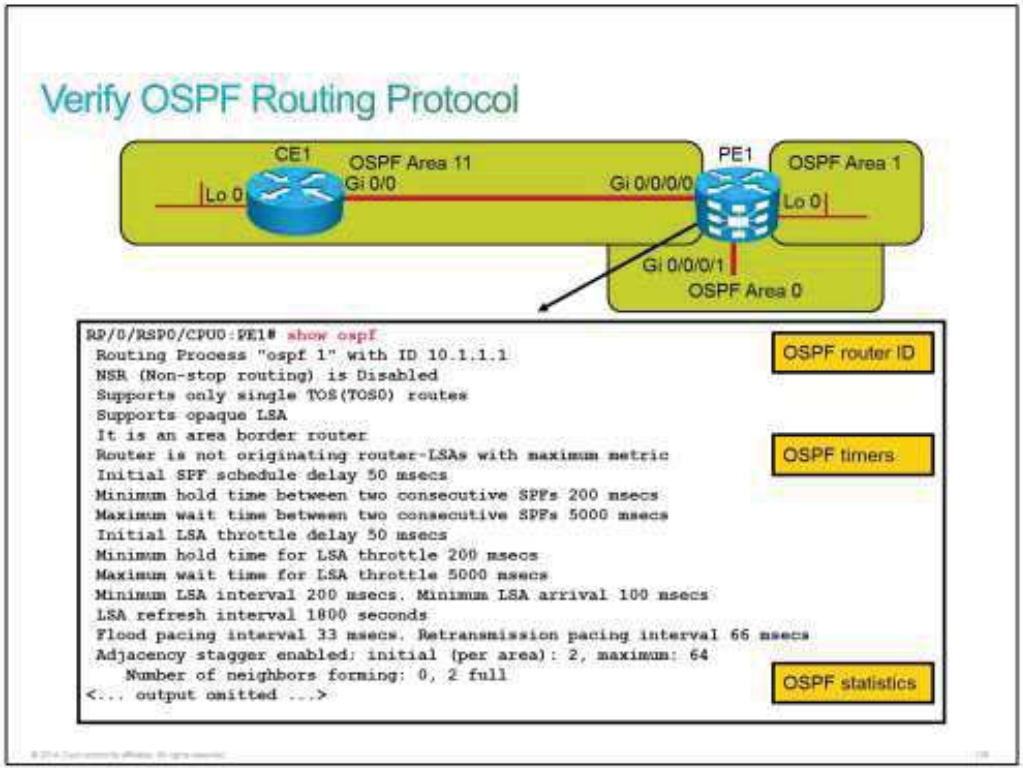

Use the Cisco IOS XR **show ospf** command or Cisco IOS and IOS XE **show ip ospf** command to display general information about OSPF routing processes. The command output shows the OSPF router ID, OSPF timers, the number of times that the SPF algorithm has been executed, and LSA information.

Verify OSPF Interface

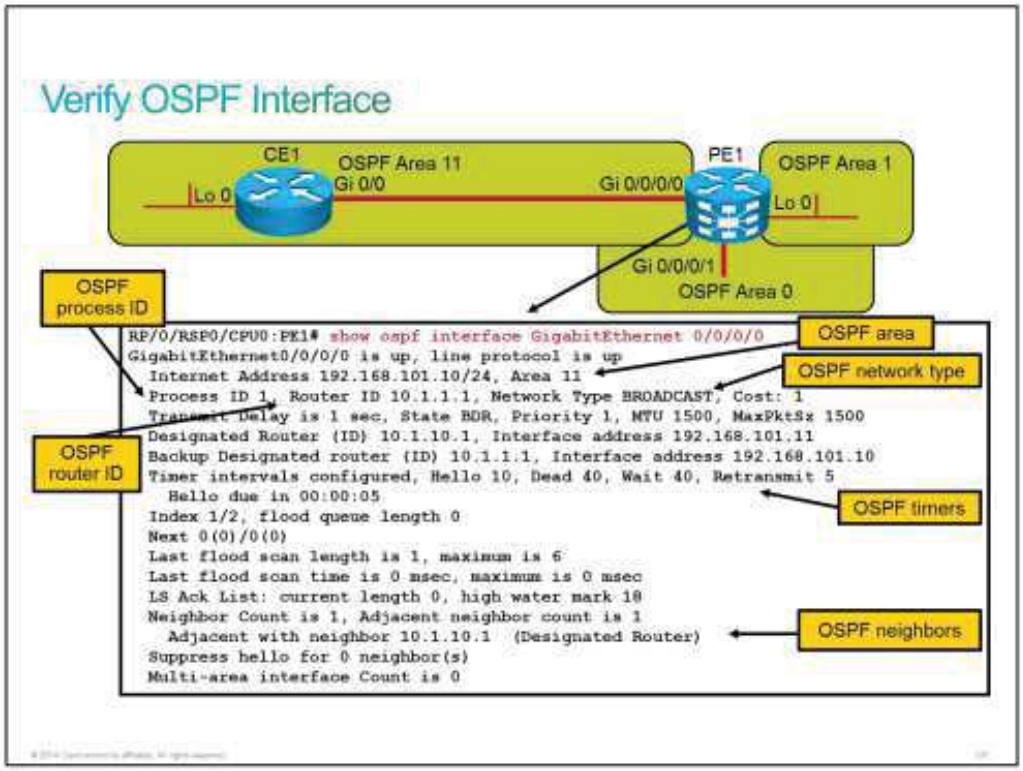

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

The Cisco IOS XR **show ospf interface** command or Cisco IOS and IOS XE **show ip ospf interface** command displays OSPF-related interface information. Use this command to verify that interfaces are configured in the intended areas. In addition, this command displays the timer intervals (including the hello interval) and shows the neighbor adjacencies.

This command also displays other information, such as the OSPF process ID, the OSPF router ID, the OSPF network type, the DR and BDR, timers, and neighbor adjacency.

Verify OSPF Neighbors

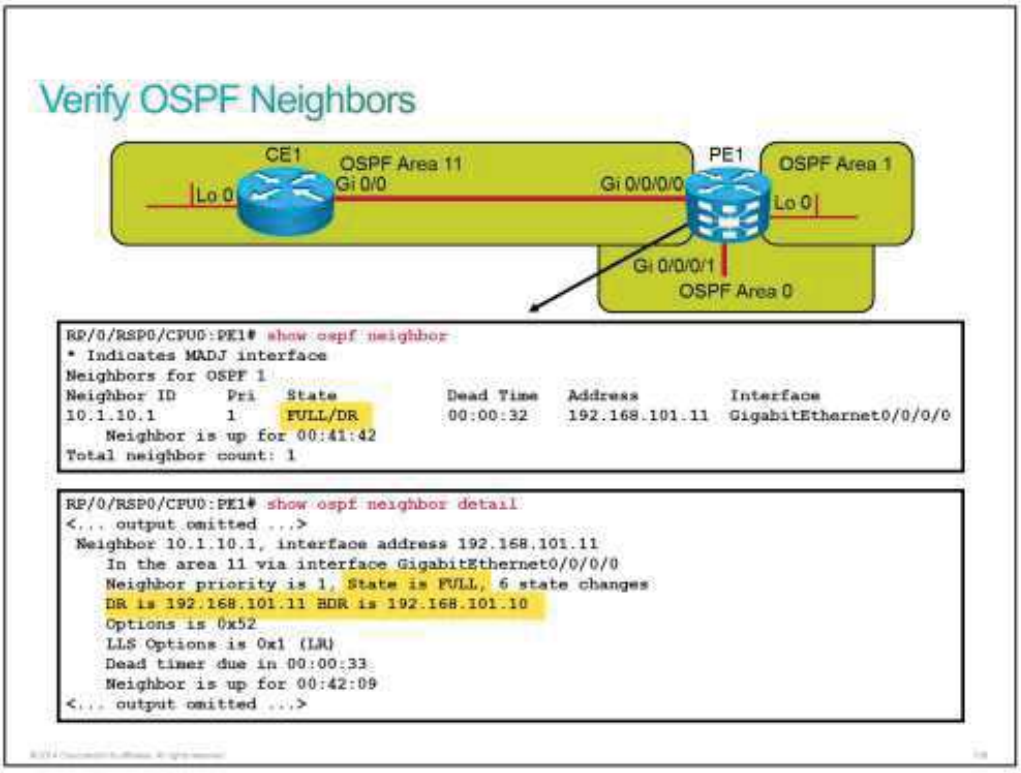

One of the most important OSPF verification and troubleshooting commands is the Cisco IOS XR **show ospf neighbor** command or Cisco IOS and IOS XE **show ip ospf neighbor** command. OSPF does not send or receive updates without having full adjacencies between neighbors. This command displays a list of neighbors, including their OSPF router IDs, their OSPF priorities, their neighbor adjacency states (for example, INIT, exstart, or full), and their dead timers.

https://t.me/learningnets

Verify Routes and Protocols

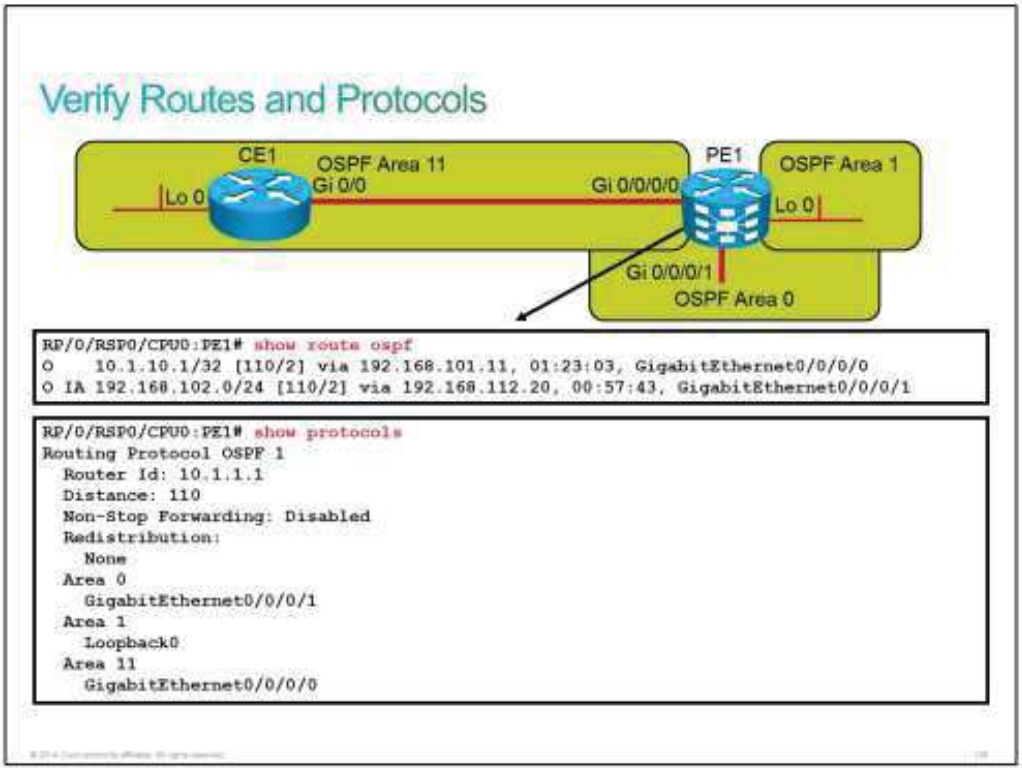

Use the Cisco IOS XR **show route ospf** command or Cisco IOS and IOS XE **show ip route ospf** command to verify the OSPF routes that are known to the router in the IP routing table. This command is one of the best ways to determine connectivity between the local router and the rest of the internetwork. This command also has optional parameters so that you can further specify the information that is to be displayed, including the OSPF process ID.

The "O" code represents OSPF routes, and "IA" means "interarea." In the figure, the 10.1.10.1/32 subnet is recognized on Gigabit Ethernet 0/0/0/0 via neighbor 192.168.101.11. The entry "[110/2]" in the routing table represents the administrative distance that is assigned to OSPF (110) and the total cost of the route to subnet 10.1.10.0 (cost of 2).

Use the Cisco IOS XR **show protocols** command or Cisco IOS and IOS XE **show ip protocols** command to verify the OSPF routing protocol parameters about timers, filters, metrics, networks, and other information for the entire router. The command output in the figure shows that the OSPF routing protocol with process number 1 is configured on the PE1 router. The router ID of the router is 10.1.1.1, and it belongs to Areas 0, 1, and 11.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Virtual Links

OSPF has special restrictions when multiple areas are involved in the OSPF AS. If more than one area is configured, one of these areas must be Area 0. This is called the backbone area. When designing networks, it is good practice to start from the core layer, which becomes Area 0, and then expand into other areas later.

The backbone must be at the center of all other areas, and other areas must be physically connected to the backbone. The main reason is that the OSPF expects all areas to inject routing information into the backbone, which distributes that information into other areas.

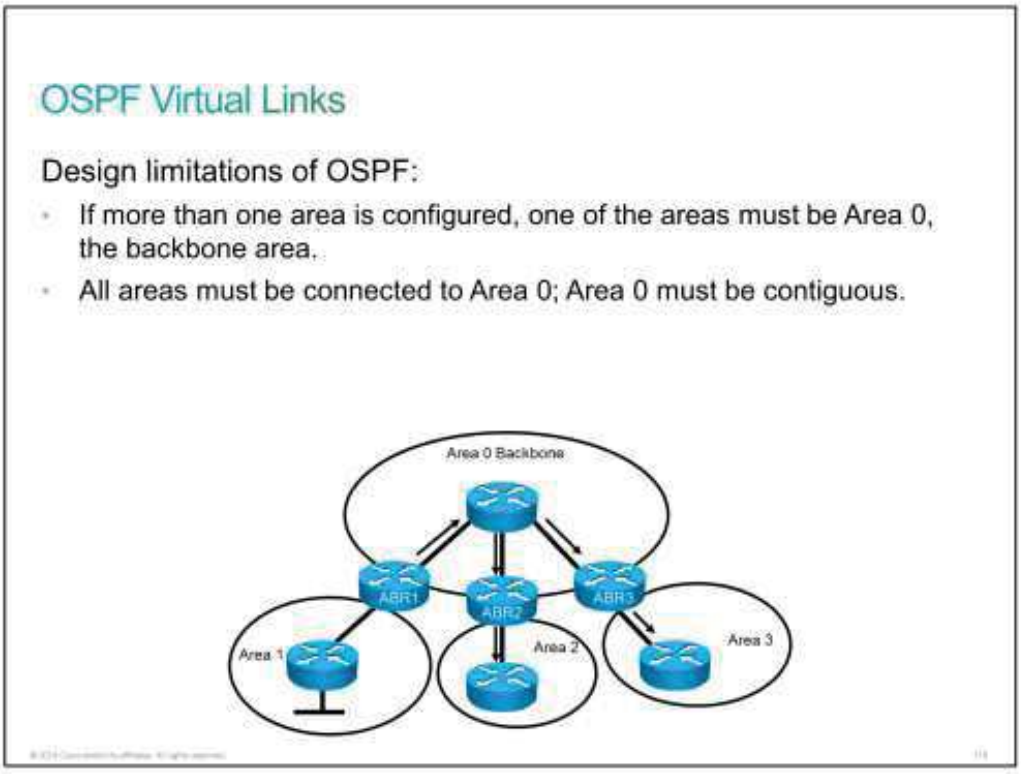

The figure shows that all areas are directly connected to the backbone. In the rare situations in which a new area is introduced that cannot have direct physical access to the backbone, a solution is required. Virtual links must be configured in such cases.

https://t.me/learningnets

OSPF Virtual Links (Cont.)

Virtual links as a solution:

- An extension to the backbone.
- Carried by a nonbackbone area.
- Cannot be created across a stub or NSSA area, or over unnumbered links.
- Virtual links are used for these purposes:
  - Allow areas to connect to the backbone through a nonbackbone area.
  - Repair a discontiguous Area 0 (for example, if two companies merge and have separate backbone areas).

OSPF has a solution to extend the reach of the backbone across other areas. This solution is called the virtual link. It provides an extension to the OSPF backbone and allows a router to connect logically to the backbone, even though there is no direct physical link. Between the two routers that are involved in the creation of the virtual link, there is a nonbackbone area. The routers at each end become part of the backbone, and both act as ABRs.

The virtual link relies on intra-area routing, and its stability is dependent on the stability of the underlying area. Virtual links cannot run through more than one area or over stub areas; they can only run through normal nonbackbone areas. If a virtual link needs to be attached to the backbone across two nonbackbone areas, two virtual links are required—one virtual link to serve each area.

Virtual links are used for two purposes:

- Linking an area that does not have a physical connection to the backbone
- Patching the backbone in case a discontinuity with Area 0 occurs

A good example of when virtual links might be required is when two companies merge that do not have a direct link between their backbone areas. A similar problem occurs when you add a nonbackbone area to an OSPF network, and the new area does not have a direct physical connection to the existing OSPF backbone area, Area 0.

https://t.me/learningnets

OSPF Virtual Links (Cont.)

Example: The role of a virtual link when there is no direct physical connection to Area 0:

- Area 20 is added with no physical access to Area 0.
- A virtual link provides a logical path to the backbone area.
- The OSPF database treats the link between routers ABR1 and ABR2 as a direct link.

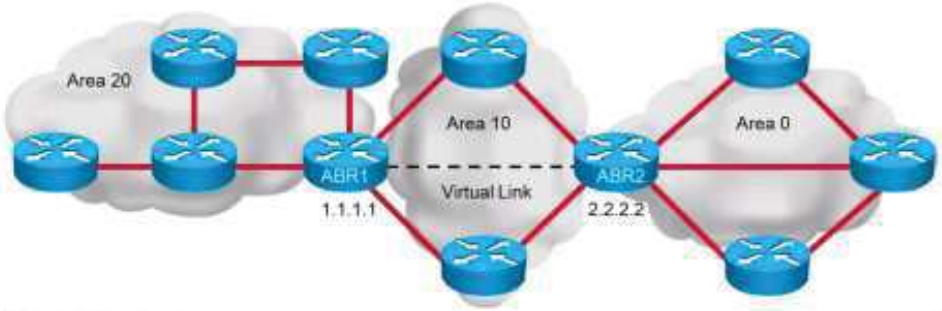

All areas in an OSPF AS must be physically connected to the backbone area (Area 0). When this is not possible, you can use a virtual link to connect to the backbone through a nonbackbone area. The area through which you configure the virtual link is known as a transit area and must have complete routing information.

In the figure, Area 20 is added to the existing OSPF topology. Because it has no direct physical link to the backbone area, Area 0, a virtual link across nonbackbone Area 10 is created. In this case, the virtual link provides a logical path between Area 20 and the backbone area.

The OSPF database treats the virtual link between ABR1 and ABR2 as a direct link. For greater stability, the loopback interface is used as a router ID and virtual links are created using these loopback addresses.

https://t.me/learningnets

OSPF Virtual Links (Cont.)

Example: The role of a virtual link when having a discontiguous Area 0

- Two companies merge without a direct link between them.
- Virtual links are used to connect the discontiguous Area 0s.
- A logical link is built between routers ABR1 and ABR2.
- Virtual links are also recommended for backup or temporary connections.

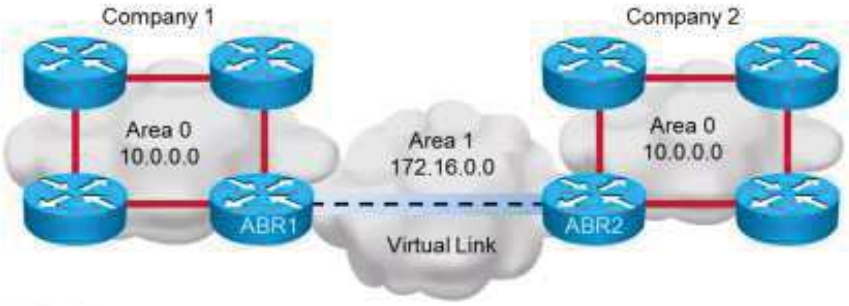

The two-tiered area hierarchy of OSPF requires that all areas are directly connected to the backbone area, Area 0, and that Area 0 is contiguous.

A virtual link is a link that allows discontiguous areas to be connected, or that allows a disconnected area to be connected to Area 0 via a transit area. The OSPF virtual link feature should be used only in very specific cases, for temporary connections or backup after a failure. Virtual links should not be used as a primary backbone design feature.

In the figure, Area 0 is discontiguous, because the two companies have merged and there is no direct link between their backbone areas. A logical link (virtual link) is built between the two ABRs—routers ABR1 and ABR2. This virtual link is similar to a standard OSPF adjacency. However, in a virtual link, the routers do not have to be directly attached to neighboring routers.

The Hello protocol works over virtual links as it does over standard links, in 10-second intervals. However, LSA updates work differently on virtual links. An LSA usually refreshes every 30 minutes. LSAs that are learned through a virtual link have the DNA option set, so that the LSA does not age out. This DNA technique is required to prevent excessive flooding over the virtual link.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Configuring OSPF Virtual Links

This is how you configure the OSPF virtual link:

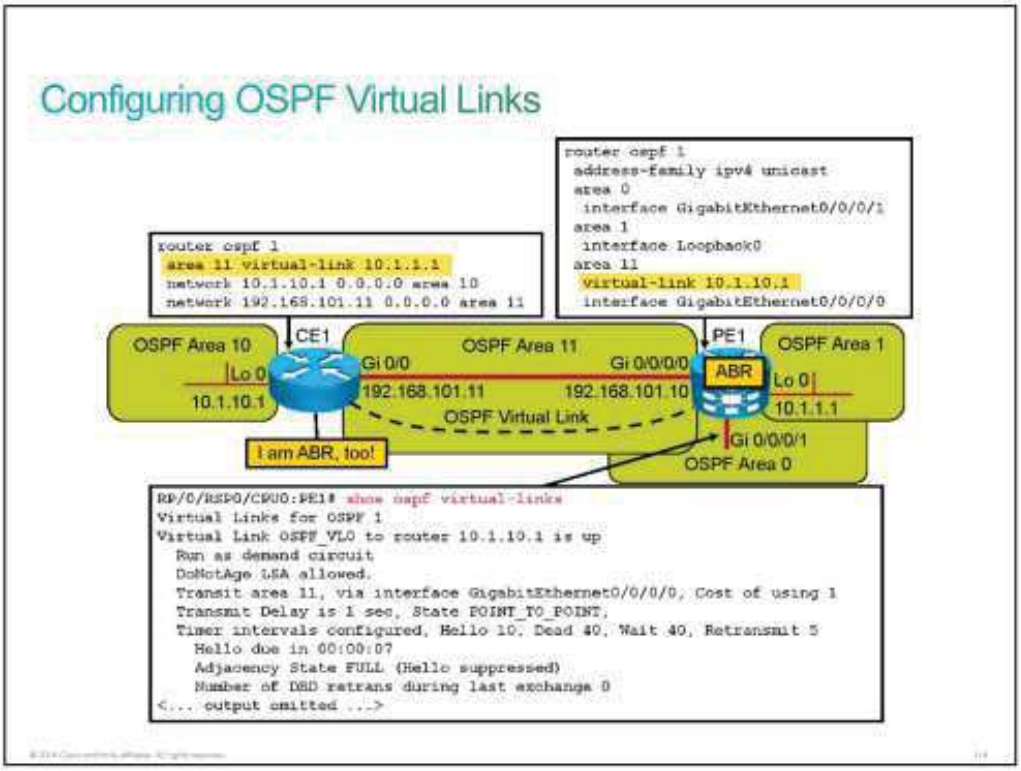

Use the Cisco IOS XR **virtual-link** router OSPF command or Cisco IOS and IOS XE **area virtual-link** router OSPF command to define an OSPF virtual link. The configuration must be done on both sides of the virtual link, and the command on each side must include the router ID of the far-end router. CE1 builds a virtual link to PE1, and PE1 builds a virtual link to CE1. Each router points to the router ID of the other router.

Use the Cisco IOS XR **show ospf virtual-links** command or Cisco IOS and IOS XE **show ip ospf virtual-links** command to verify that the configured virtual link works properly. The figure shows that the virtual link toward the neighboring router with the router ID 10.1.10.1 is up, and Area 11 is used as a transit area. The output also shows that interface GigabitEthernet0/0/0/0 is used to form the virtual link as well as several OSPF timers.

https://t.me/learningnets

OSPF Cost

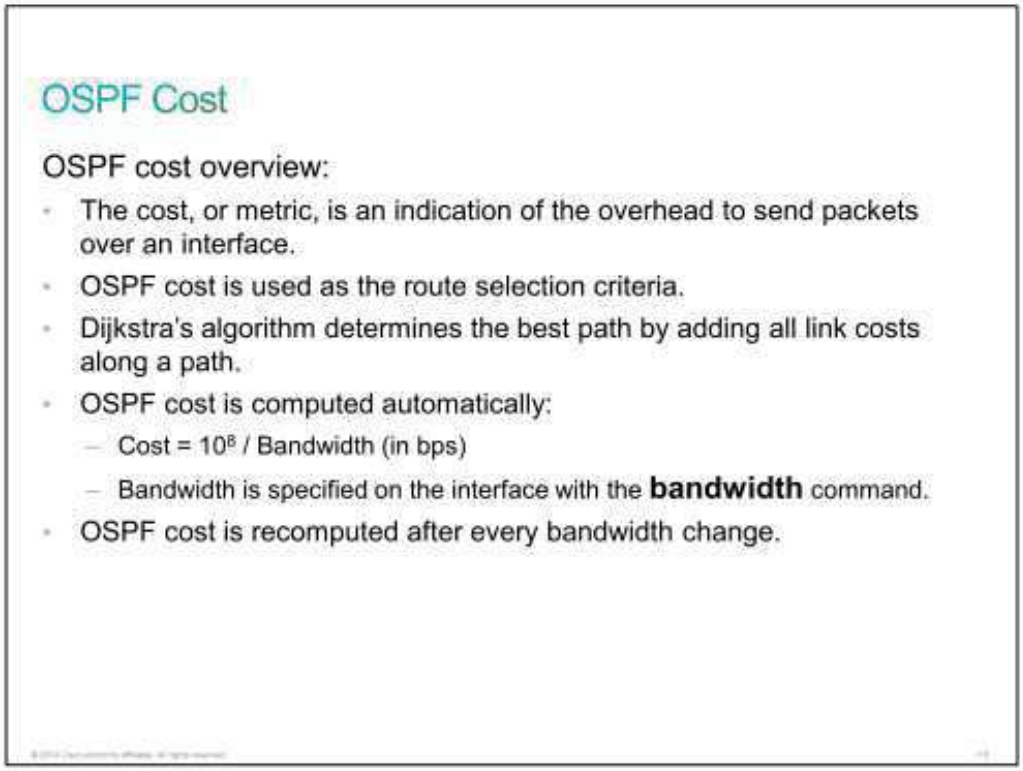

The cost value is a 16-bit positive number between 1 and 65,535, where a lower value is a more desirable metric. For example, a 64-kbps link gets a metric of 1562, while a T1 link gets a metric of 64. Cost is applied on all router link paths, and route decisions are made on the total cost of a path. The metric is only relevant on an outbound path; route decisions are not made for inbound traffic. The OSPF cost is recomputed after every bandwidth change, and Dijkstra's algorithm determines the best path by adding all link costs along a path.

On high-bandwidth links (155 Mbps and more), automatic cost assignment no longer works (it would result in all costs being equal to 1). On these links, OSPF costs must be set manually on each interface.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

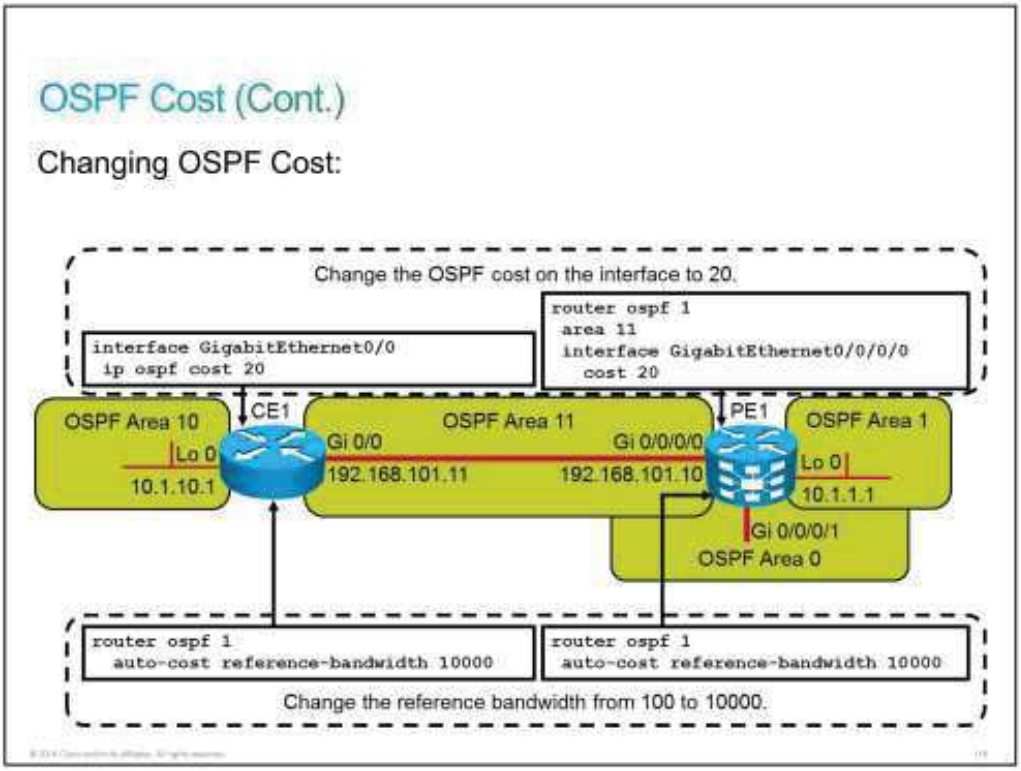

In general, the OSPF cost in Cisco routers is calculated using the following formula: (100 Mbps) / (bandwidth in Mbps).

However, the cost is calculated based on a maximum bandwidth of 100 Mbps, which is a cost of 1. If you have faster interfaces, you may want to recalibrate the cost of 1 to a higher bandwidth.

When you are using the bandwidth of the interface to determine OSPF cost, always remember to use the Cisco IOS, IOS XE, or IOS XR **bandwidth** interface command to accurately define the bandwidth of the interface (in kbps). The **bandwidth** command is the reference for calculating the cost.

To override the automatically calculated default cost, manually define the cost using the Cisco IOS XR **cost** router OSPF command or Cisco IOS and IOS XE **ip ospf cost** interface command. The cost value is an integer from 1 to 65,535. The preferred link is the link with the lowest number.

Interfaces that are faster than 100 Mbps are being used (Gigabit Ethernet and Fast Ethernet). The automatic cost assignment returns the value "1." For all faster links, a new reference value must be configured. Use the Cisco IOS, IOS XE, or IOS XR **auto-cost** router OSPF command to ensure accurate route calculations. The *reference-bandwidth* value in the **auto-cost** command is a reference bandwidth in megabits per second; it ranges from 1 to 4,294,967, with a default value of 100.

https://t.me/learningnets

Cisco Nonstop Forwarding

Cisco Nonstop Forwarding works with the SSO to minimize the amount of time that a network is unavailable to its users following a switchover. The main objective of Cisco NSF is to continue forwarding IP packets following a RP switchover in platforms with dual RPs.

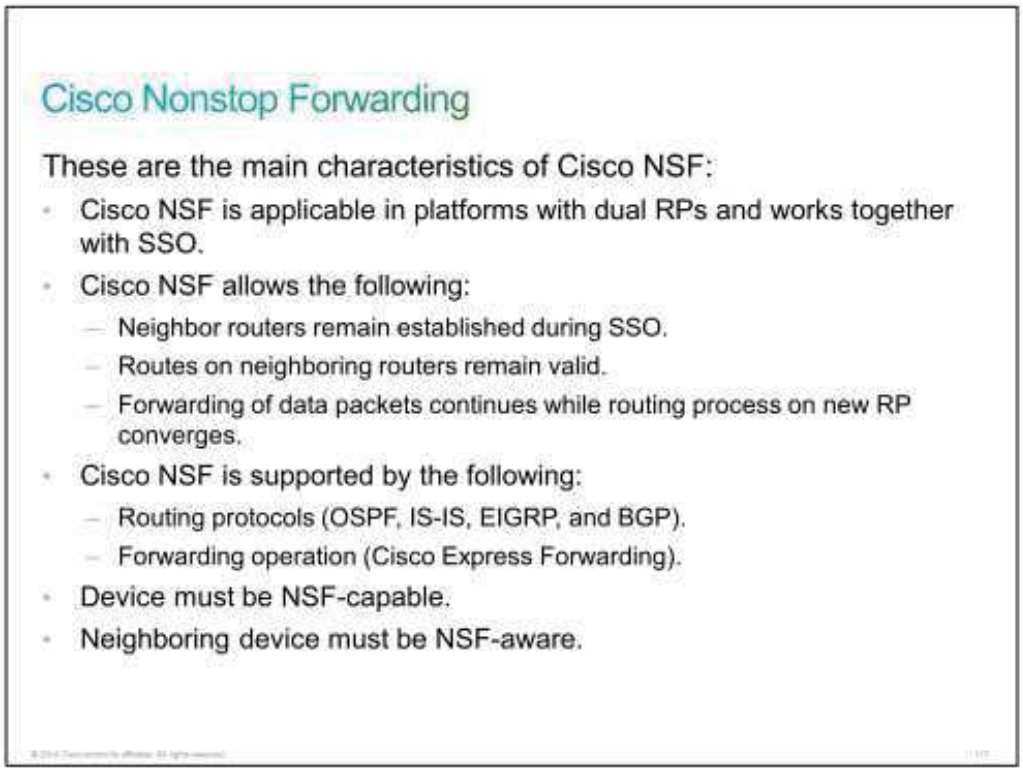

Usually, when a networking device restarts, all routing peers of that device detect that the device went down and then came back up. This transition results in what is called a routing flap, which could spread across multiple routing domains. Routing flaps that are caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. Cisco NSF helps to suppress routing flaps in SSO-enabled devices, thus reducing network instability.

Cisco NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. With Cisco NSF, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards or dual FPs while the standby RP assumes control from the failed active RP during a switchover.

Cisco NSF is supported by the BGP, EIGRP, OSPF, and IS-IS protocols for routing and by Cisco Express Forwarding for forwarding. Of the routing protocols, BGP, EIGRP, OSPF, and IS-IS are enhanced with Cisco NSF capability and awareness, which means that routers running these protocols can detect a switchover and take the necessary actions to continue forwarding network traffic and to recover route information from the peer devices.

A networking device is said to be Cisco NSF-aware if it is running Cisco NSF-compatible software. A device is said to be Cisco NSF-capable if it has been configured to support Cisco NSF; therefore, it would rebuild routing information from Cisco NSF-aware or Cisco NSF-capable neighbors.

https://t.me/learningnets

Cisco NSF depends on Cisco Express Forwarding to continue forwarding packets during switchover while the routing protocols rebuild the RIB tables. When the routing protocols have converged, Cisco Express Forwarding updates the FIB table and removes old route entries. Cisco Express Forwarding, in turn, updates the line cards with the new FIB information.

Cisco NSF uses the IETF standardized graceful restart functionality.

https://t.me/learningnets

Cisco Nonstop Routing

The OSPF support for the NSR with SSO feature enables routers to maintain an OSPF state with other routers and ensure continuous packet forwarding during route processing.

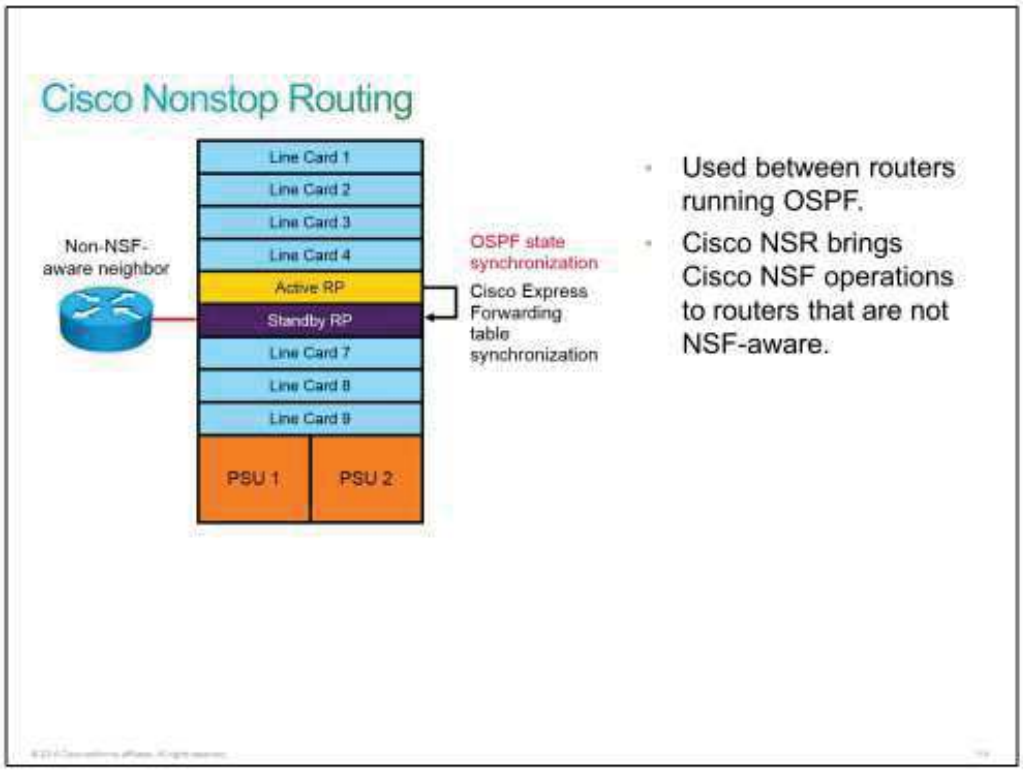

Other routers do not need to be NSF-capable or NSF-aware to benefit from OSPF NSR capabilities.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Cisco NSF and NSR for OSPF

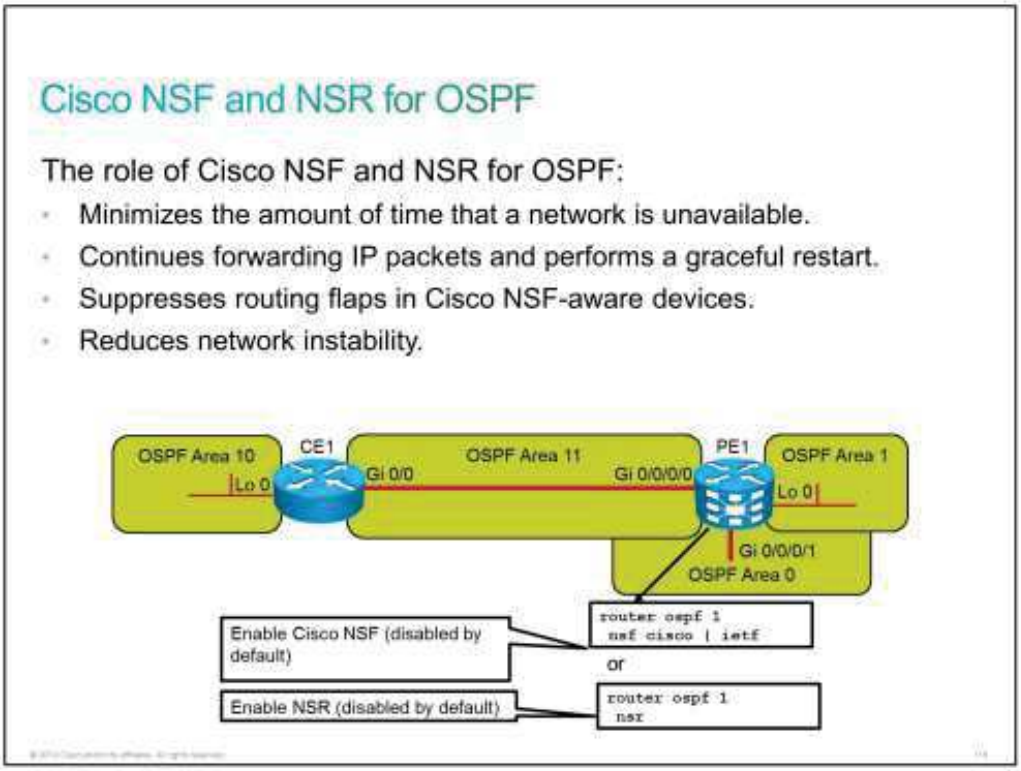

Cisco IOS XR Software, Cisco Nonstop Forwarding for OSPF version 2, allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a failover. With Cisco NSF, peer networking devices do not experience routing flaps. During failover, data traffic is forwarded through intelligent line cards while the standby RP assumes control from the failed RP. The ability of line cards to remain up through a failover and to be kept current with the FIB on the active RP is important to Cisco IOS XR Software NSF operation.

Routing protocols, such as OSPF, run only on the active RP or DRP and receive routing updates from their neighbor routers. When an OSPF Cisco NSF-capable router performs an RP failover, it must perform two tasks to resynchronize its LSDB with its OSPF neighbors. First, it must relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the LSDB for the network.

As quickly as possible after an RP failover, the Cisco NSF-capable router sends an OSPF Cisco NSF signal to neighboring Cisco NSF-aware devices. This signal is in the form of a link-local LSA generated by the failed-over router. Neighbor networking devices recognize this signal as a cue that the neighbor relationship with this router should not be reset. As the Cisco NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

After neighbor relationships are re-established, the NSF-capable router begins to resynchronize its database with all of its NSF-aware neighbors. At this point, the routing information is exchanged between the OSPF neighbors. After this exchange is completed, the NSF-capable device uses the routing information to remove stale routes, update the RIB, and update the FIB with the new forwarding information. OSPF on the router and the OSPF neighbors are now fully converged.

NSR allows an RP failover, process restart, or in-service upgrade to be invisible to peer routers and ensures that there is minimal performance or processing impact. Routing protocol interactions between routers are not impacted by NSR. NSR is built on the warm standby extensions. NSR alleviates the requirement for Cisco NSF and IETF graceful restart protocol extensions.

https://t.me/learningnets

Graceful Restart for OSPFv3

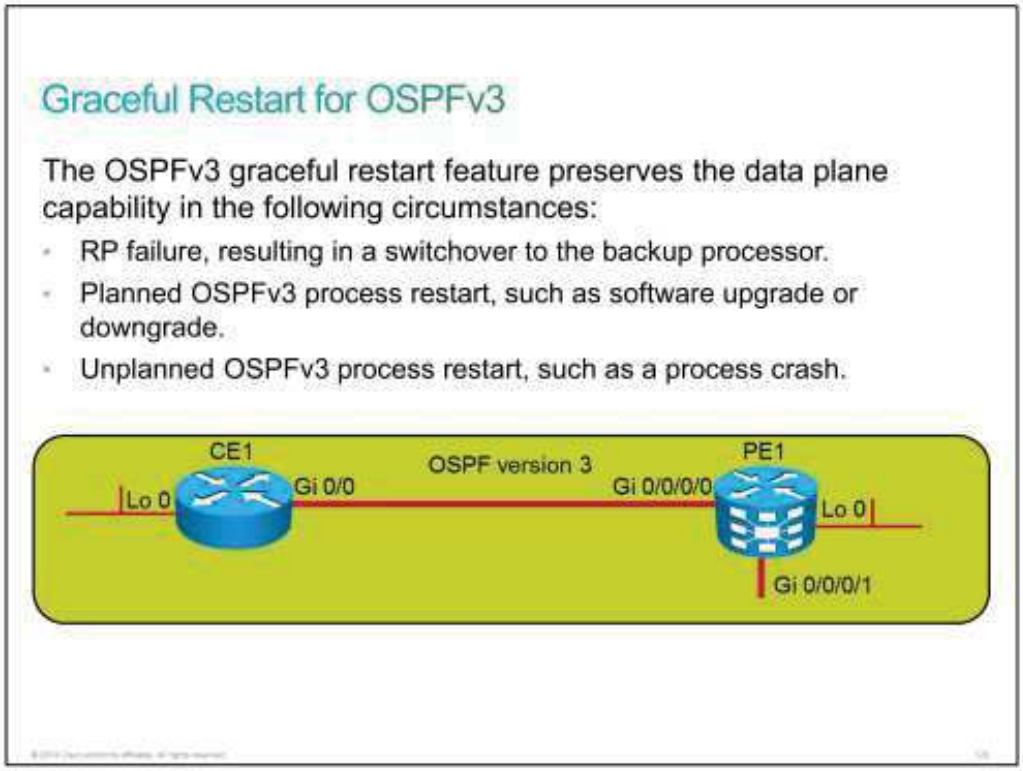

The graceful restart feature supports nonstop data forwarding on established routes while the OSPFv3 routing protocol is restarting. Therefore, this feature enhances high availability of IPv6 forwarding.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Bidirectional Forwarding Detection

In both enterprise and service provider networks, the convergence of business-critical applications onto a common IP infrastructure is becoming more common. Given the criticality of the data, these networks are typically constructed with a high degree of redundancy. While such redundancy is desirable, its effectiveness is dependent upon the ability of individual network devices to quickly detect failures and reroute traffic to an alternate path.

This detection is now typically accomplished via hardware detection mechanisms. However, the signals from these mechanisms are not always conveyed directly to the upper protocol layers. When the hardware mechanisms do not exist (such as Ethernet) or when the signaling does not reach the upper protocol layers, the protocols must rely on their much slower strategies to detect failures. The detection times in existing protocols are typically greater than 1 second and sometimes much longer. For some applications, this is too long to be useful.

BFD provides a low-overhead, short-duration method of detecting failures in the forwarding path between two adjacent routers, including the interfaces, data links, and forwarding planes. BFD is a detection protocol that you enable at the interface and routing protocol levels.

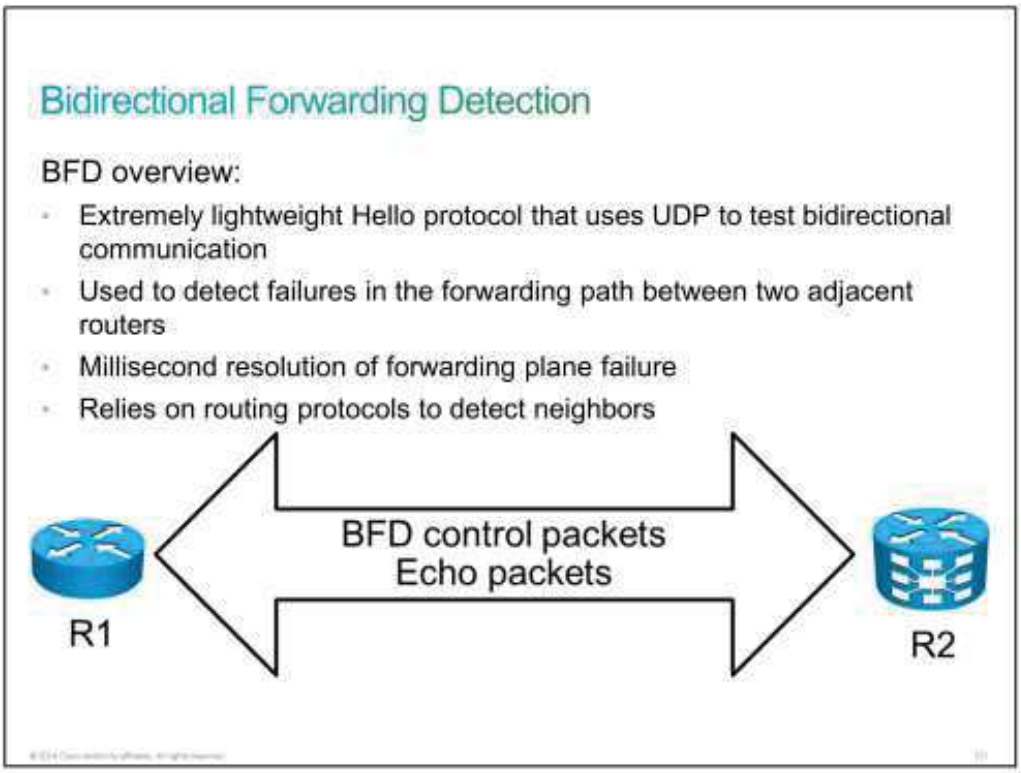

Cisco supports the BFD asynchronous mode, which depends on the sending of BFD control packets between two systems to activate and maintain BFD neighbor sessions between routers. Therefore, in order for a BFD session to be created, you must configure BFD on both BFD peers. Once BFD has been enabled on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers will begin to send BFD control packets to each other at the negotiated interval.

BFD provides fast BFD peer failure detection times independently of all media types, encapsulations, topologies, and the routing protocols BGP, EIGRP, IS-IS, and OSPF. By sending rapid failure detection notices to the routing protocols in the local router to initiate the routing table recalculation process, BFD contributes to greatly reduced overall network convergence time.

https://t.me/learningnets

The BFD protocol has no discovery mechanisms to detect neighbors. It is designed solely as an agent for other applications requiring fast failure detection. When a routing protocol that is configured to use BFD detects a new neighbor, it requests availability tracking from BFD.

BFD can rely on control packets or on echo packets. Echo packets are IP packets addressed to the router itself but sent to the Layer 2 address of the next-hop node. The echo packets thoroughly test the complete bidirectional forwarding path between adjacent routers, as they have to be transmitted by the originating router, propagated to the adjacent router, received by its interface logic, switched by its forwarding engine, and sent back to the originator (because the IP packet is addressed to the router itself).

For example, when R1 sends a BFD echo packet, it sets the destination IP address in the packet to its own interface IP address and the MAC address in the Layer 2 frame header to the neighbor (R2) MAC address. When R2 receives the packet, it performs a Layer 3 lookup and sends the packet toward its final destination (back to R1).

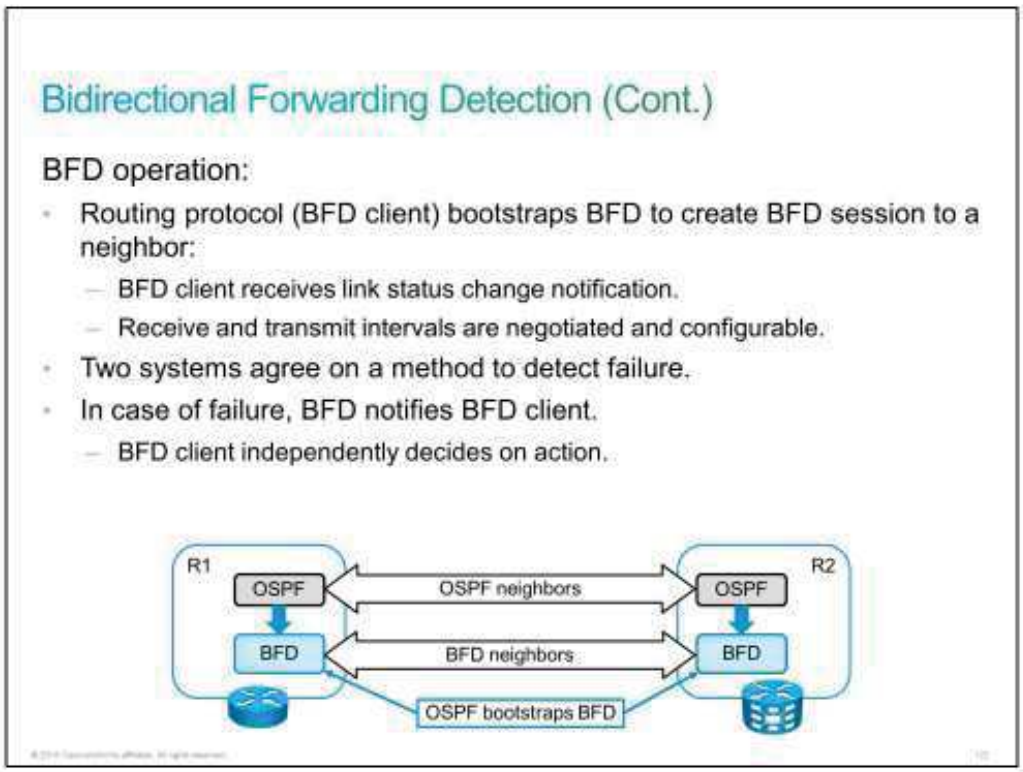

When a routing protocol (OSPF, for example) discovers a neighbor, it sends a request to the local BFD process to initiate a BFD neighbor session with the OSPF neighbor router. Then the BFD neighbor session with the OSPF neighbor router is established. If there is a failure on the link between neighbors, the BFD neighbor session with the OSPF neighbor router is torn down. BFD notifies the local OSPF process that the BFD neighbor is no longer reachable. The local OSPF process tears down the OSPF neighbor relationship. If an alternative path is available the routers will immediately start converging on it.

https://t.me/learningnets

Bidirectional Forwarding Detection for OSPF

First you need to configure BFD on individual interfaces. After you have configured BFD on individual interfaces, you have to tell the routing protocols to use it. In most cases, there is no good reason that the routing protocols would not use BFD whenever possible. The only command that you have to use in these scenarios is the **bfd all-interfaces** Cisco IOS router configuration command.

If you want to be more specific, you could enable BFD on individual interfaces with the specific command for that particular routing protocol. For OSPF, use the **ip ospf bfd** Cisco IOS interface configuration command.

You could also enable BFD on all interfaces with the **bfd all-interfaces** Cisco IOS router configuration command, and disable it on specific interfaces. OSPF uses the **ip ospf bfd disable** Cisco IOS interface configuration command.

You would disable BFD on an interface only if the interface is flaky but you want to retain routing adjacency across short failures. This should not be necessary in most well-designed networks.

To enable BFD and adjust parameters, use the **bfd** Cisco IOS XR router commands.

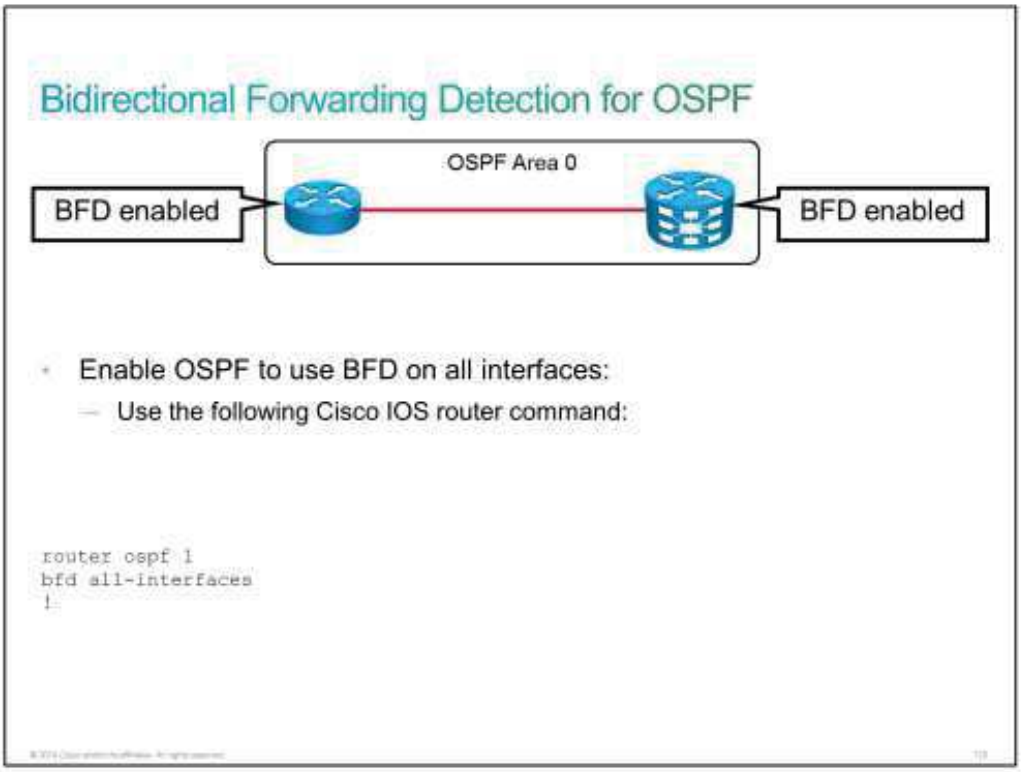

https://t.me/learningnets

Bidirectional Forwarding Detection for OSPF (Cont.)

- Enable OSPF to use BFD on a single interface:
  - Use the following Cisco IOS interface configuration command:

```
interface TenGigabitEthernet3/0/1
ip ospf bfd
bfd interval 100 min_rx 100 multiplier 3
!
```

- To enable BFD and adjust parameters, use the following Cisco IOS XR router commands:

```
router ospf 1
area 0
  interface TenGigE0/1/4/0
    bfd fast-detect
    bfd minimum-interval 100
    bfd multiplier 3
```


©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Authentication Overview

OSPF routing protocol supports authentication of routing updates to prevent attacks to the routing protocol. For example, an attacker might "poison" the routing table of the router by sending a route toward one of the networks, using good cost, and traffic to that network would be diverted to the attacker router. A similar situation applies to an attacker who intercepts a routing update, changes it, and then forwards it. Authentication prevents such attacks by authenticating each routing update. Authentication is accomplished by the exchange of authenticating information that is known to the sending and receiving router only.

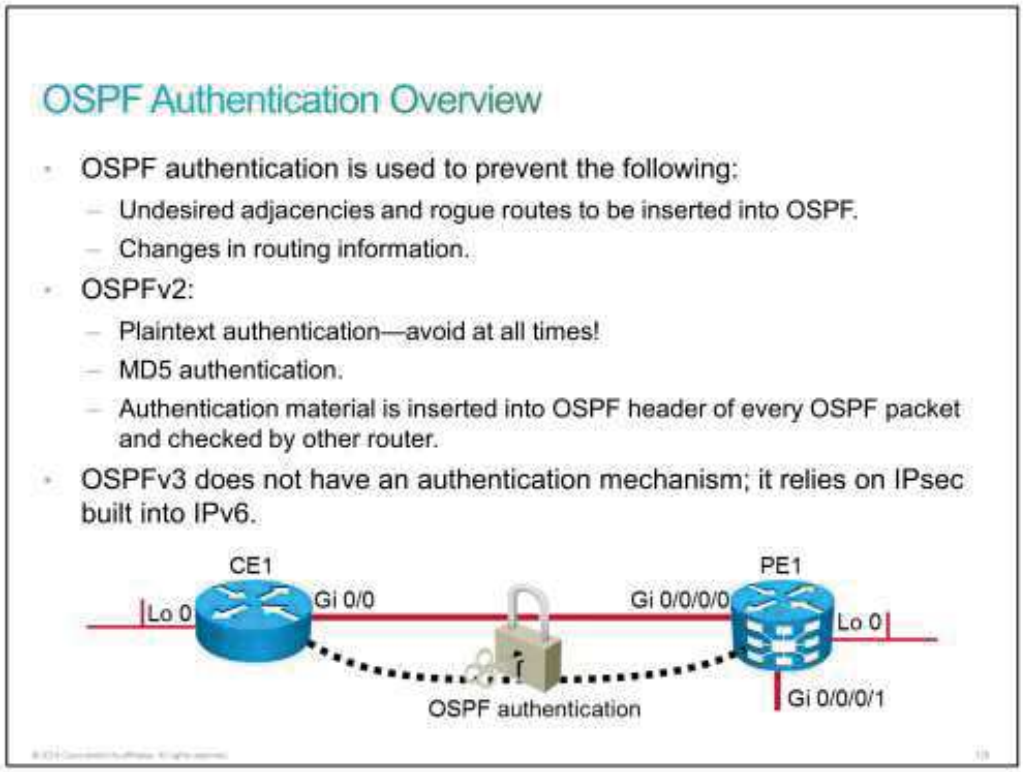

Both OSPFv2 and OSPFv3 support authentication. However, there is significant difference in authentication mechanisms in both protocols. OSPFv2 uses a built-in authentication mechanism and supports plaintext and MD5 authentication. Authentication information (plaintext password or MD5 hash produced from a key and the routing update itself) in OSPFv2 is inserted into the OSPF header and checked by the other router. You should avoid sending an authentication password in plaintext across a line; it is not considered safe.

OSPFv3 does not use a built-in authentication mechanism and relies on IPv6 native security capabilities and native security stack, which uses IPsec.

https://t.me/learningnets

OSPFv2 Authentication

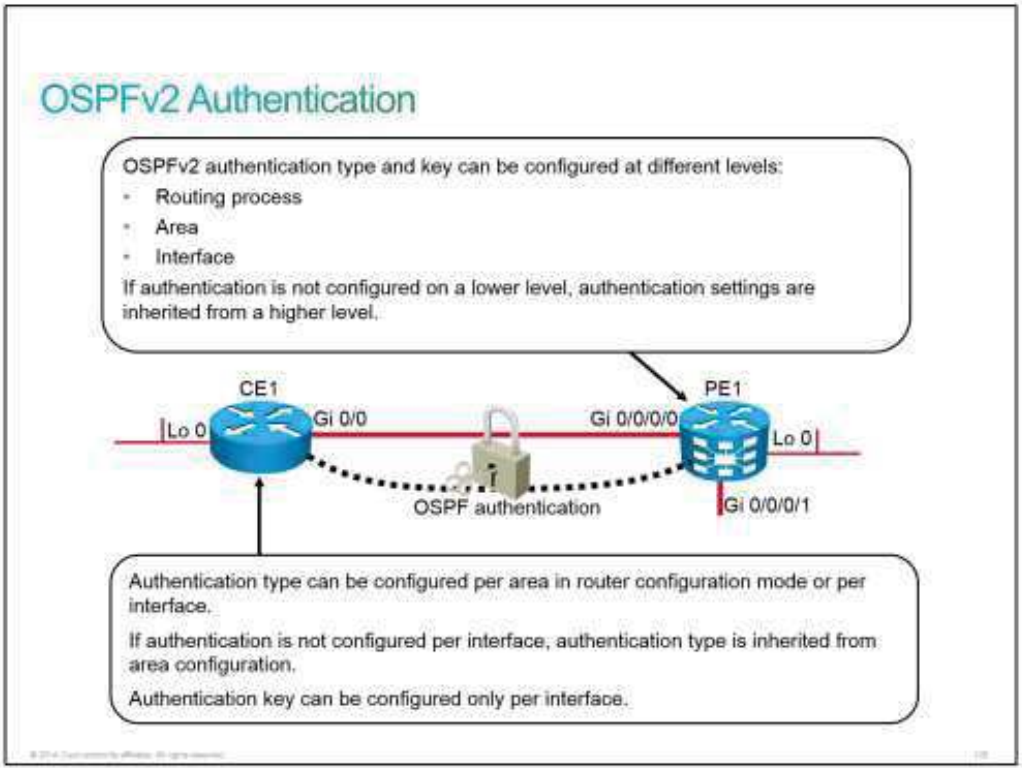

The OSPFv2 authentication type and key on Cisco IOS XR platforms can be configured at different levels. Authentication can be specified for an entire process or area, or on an interface. An interface can be configured for only one type of authentication, not both. Authentication that is configured for an interface overrides authentication that is configured for the area or process. If authentication is not configured on a lower level (for example, on the interface), authentication is inherited from a higher level, for example from an area or process.

If you intend for all interfaces in an area to use the same type of authentication, you can configure fewer commands if you use the **authentication** command in the area configuration submode (and specify the **message-digest** keyword if you want the entire area to use MD5 authentication). This strategy requires fewer commands than specifying the authentication for each interface.

The OSPF authentication type on Cisco IOS and IOS XE platforms can be configured per area in router configuration mode or per interface. However, an authentication key can be specified per interface only. If you intend for all interfaces in an area to use the same type of authentication, you can configure fewer commands if you use the **authentication** command in the area configuration. However, you still have to configure the authentication key on each interface separately.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPFv3 Authentication

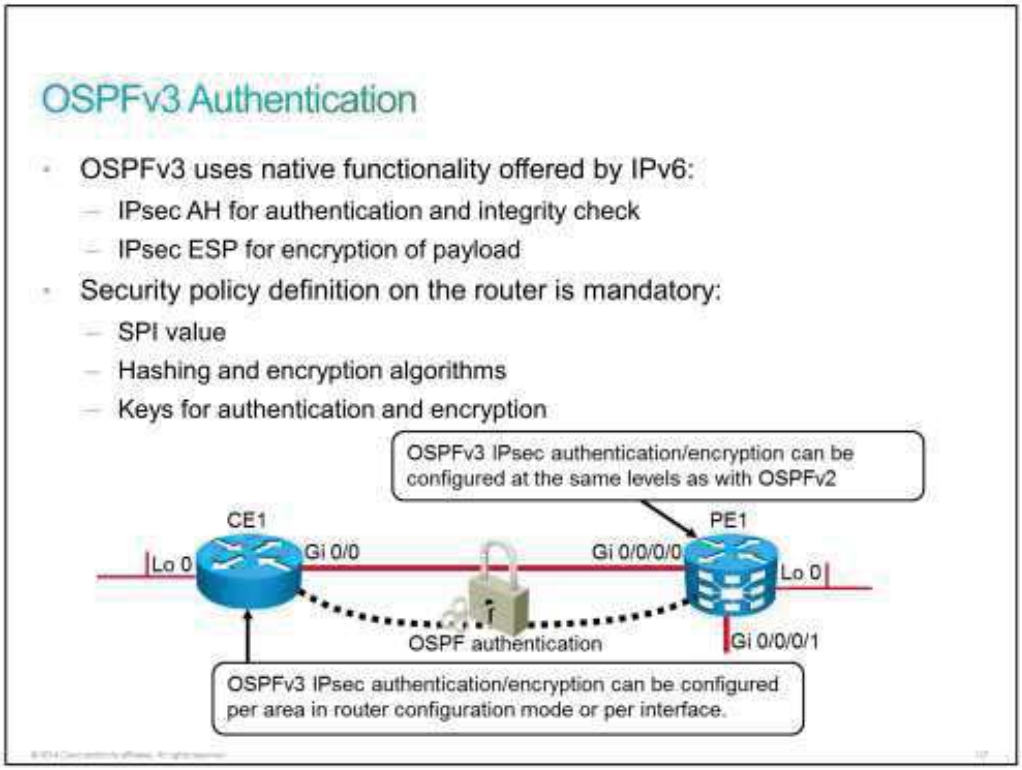

OSPFv3 uses IPv6 native security capabilities and native security stack. Two possible protocols are available:

- AH for authentication and integrity check.
- ESP for encrypting the payload—the routing updates themselves and authentication and integrity check.

Using an IPsec connection for OSPFv3 authentication requires you to define a security policy for every neighbor router. The security policy defines which protocol is used for communication (AH or ESP), hashing and encryption algorithm, keys, and the SPI value.

On Cisco IOS XR platforms, OSPFv3 authentication or encryption can be configured at the same levels as with OSPFv2: OSPF process, area or interface. The concept of inheritance is applied, as with the OSPFv2 authentication.

On Cisco IOS and IOS XE platforms, OSPFv3 authentication or encryption can be configured per area in router configuration mode or per interface.

https://t.me/learningnets

Configuring OSPF Authentication

With the help of the following examples, you will learn how to configure OSPF authentication.

Configuring OSPFv2 Authentication

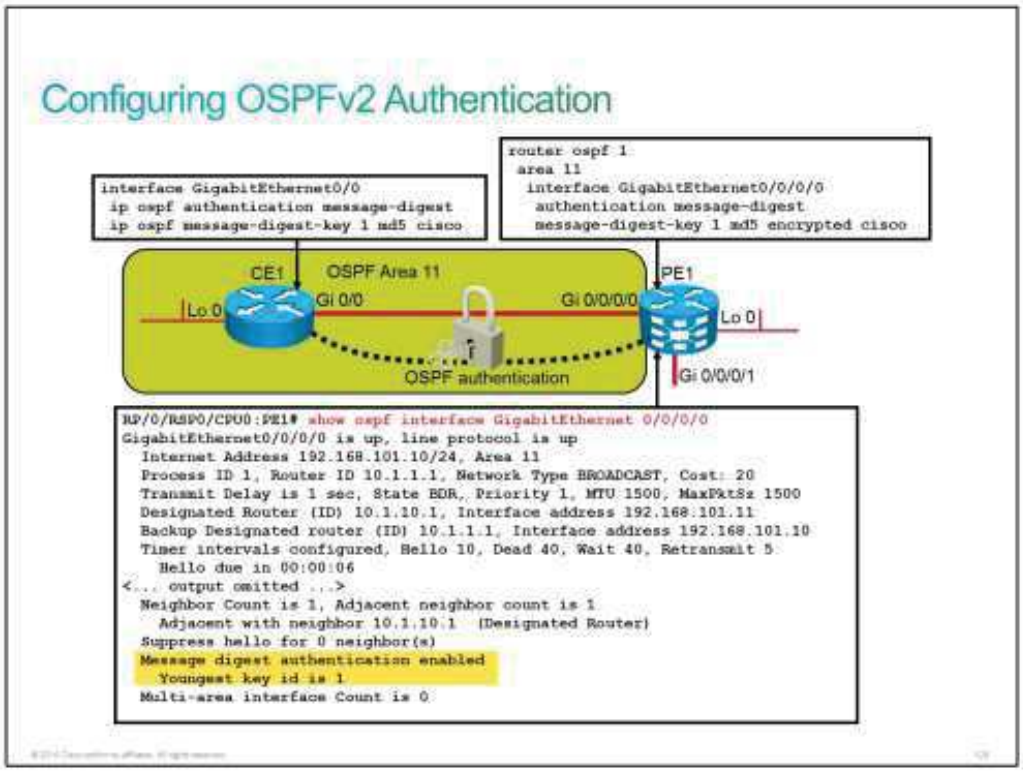

The figure shows a configuration of the OSPFv2 MD5 authentication on a CE1 router on Gigabit Ethernet 0/0 interface and PE1 router on Gigabit Ethernet 0/0/0/0. To enable OSPFv2 authentication on the Cisco IOS XR, use the **authentication message-digest** and **message-digest-key md5** router OSPF commands. To enable OSPFv2 authentication on the Cisco IOS and IOS XE router, use the **ip ospf authentication message-digest** and **ip ospf message-digest-key md5** interface commands. Key-ID in the example is "1" and the actual key is "cisco." Key-ID and key must match on both routers where authentication is being enabled.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

To verify that OSPF authentication is correctly configured and OSPF adjacency is up, use the Cisco IOS XR **show ospf interface** command. In the output, you should see a message that MD5 authentication is enabled and a key different than zero ("0") is used. Similar verification can be performed on the Cisco IOS and IOS XE router by using the **show ip ospf interface** command:

```
CE1# show ip ospf interface GigabitEthernet 0/0
GigabitEthernet0/0 is up, line protocol is up
Internet Address 192.168.101.11/24, Area 11
Process ID 1, Router ID 10.1.10.1, Network Type BROADCAST, Cost: 1
Topology-MTID    Cost    Disabled    Shutdown      Topology Name
      0           1         no          no            Base
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 10.1.10.1, Interface address 192.168.101.11
Backup Designated router (ID) 10.1.1.1, Interface address 192.168.101.10
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  oob-resync timeout 40
  Hello due in 00:00:05
Supports Link-local Signaling (LLS)
Cisco NSF helper support enabled
IETF NSF helper support enabled
Index 1/2, flood queue length 0
Next 0x0(0)/0x0(0)
Last flood scan length is 0, maximum is 5
Last flood scan time is 0 msec, maximum is 4 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.1.1.1  (Backup Designated Router)
Suppress hello for 0 neighbor(s)
Message digest authentication enabled
  Youngest key id is 1
```

Configuring OSPFv3 Authentication

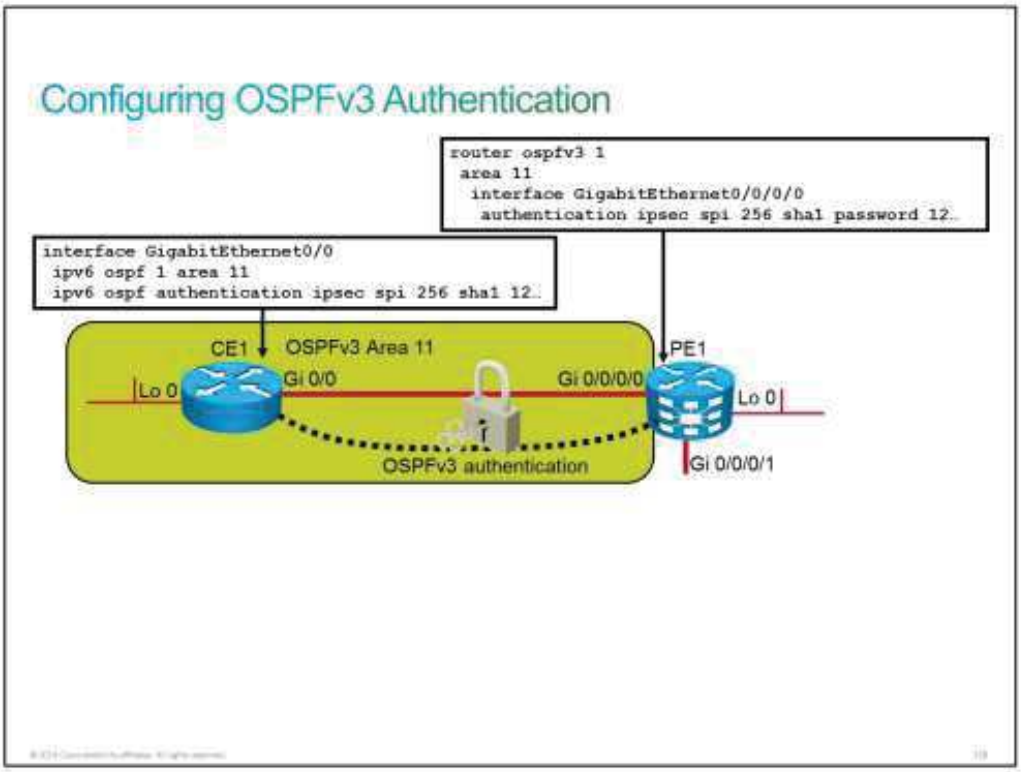

https://t.me/learningnets

The figure shows a configuration of the OSPFv3 SHA authentication on the CE1 router on Gigabit Ethernet 0/0 interface and a PE1 router on Gigabit Ethernet 0/0/0/0. To enable OSPFv3 authentication on the Cisco IOS XR router, use the **authentication ipsec spi** router OSPF command. To enable OSPFv3 authentication on the Cisco IOS and IOS XE router, use the **ipv6 ospf authentication ipsec spi** interface command. The SPI index, hashing method, and a key should match on both routers. The key must be the correct length; the SHA hashing algorithm uses a 160-bit key, so the key must be a 40-hexadecimal number. In the example, the SPI index is 256, and the hashing algorithm is SHA.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

Summary

- Before OSPF is enabled on the service provider network, there must be detailed planning of IP addresses, areas, positioning of ABRs and ASBRs, and redistribution and summarization.
- OSPF requires router ID to uniquely describe each router in a network.
- Designating OSPF interface as passive disables sending and receiving of OSPF packets.
- You can use various **show** commands to verify OSPF.
- A virtual link allows discontiguous Area 0s to be connected, or a disconnected area to be connected to Area 0 via a transit area.

Summary (Cont.)

- You can use the **virtual-link** Cisco IOS XR command to configure OSPF virtual link.
- OSPF cost can be changed by changing bandwidth or cost on the link or by changing reference bandwidth.
- Cisco NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover.
- After an RP failover, the Cisco NSF-capable router sends an OSPF Cisco NSF signal to neighboring Cisco NSF-aware devices to inform the neighbor not to reset the neighbor relationship.
- OSPFv3 graceful restart feature preserves the data plane capability in RP failure or OSPF process restart.

https://t.me/learningnets

Summary (Cont.)

- BFD provides fast peer failure detection times independently of media types, encapsulations, topologies, and the routing protocols.
- You would not use BFD only if an interface is flaky and you want to retain routing adjacency across short failures.
- OSPF supports authentication to prevent undesired adjacencies and changes in routing information.
- To enable OSPFv2 authentication on the Cisco IOS XR, use the **authentication message-digest** and **message-digest-key md5** router OSPF commands.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: Implementing OSPF Special Area Types

Overview

Scalability, improved CPU and memory utilization, and the ability to mix small routers with large routers are all benefits of using proper route summarization techniques. A key feature of the OSPF protocol is the ability to summarize routes at area and AS boundaries. Route summarization is important, because it reduces the amount of OSPF LSA flooding and the sizes of LSDBs and routing tables, which reduce memory and CPU utilization on the routers. The OSPF network can scale to very large sizes, in part because of route summarization.

The OSPF protocol defines several special-case area types, including stub areas, totally stubby areas, and NSSAs. The purpose of all three types of stub areas is to inject default routes into an area so that external and summary LSAs are not flooded. Stub areas are designed to reduce the amount of flooding, the LSDB size, and the routing table size in routers within the area. Network designers should always consider using stub area techniques when building networks. Stub area techniques improve performance in OSPF networks and allow the network to scale to significantly larger sizes.

This lesson defines the different types of route summarization, and discusses OSPF area types and how to configure them.

Upon completing this lesson, you will be able to:

- Describe interarea and external OSPF route summarization
- Describe how OSPF interarea routes are summarized
- Describe how OSPF external routes are summarized
- Describe how default routes are injected into OSPF
- List and describe different OSPF area types
- Describe how to configure different OSPF area types
- Describe OSPF router and LSA types
- Describe OSPF stub area rules

https://t.me/learningnets

- Describe how to configure OSPF stub areas
- Describe OSPF NSSA rules
- Describe how to configure OSPF NSSA and totally NSSA

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Summarization

Route summarization is a key to scalability in OSPF. Route summarization helps solve two major problems: large routing tables and frequent LSA flooding throughout the AS. Every time that a route disappears in one area, routers in other areas get involved in shortest-path calculation. To reduce the size of the area database, you can configure summarization on an area boundary or an AS boundary.

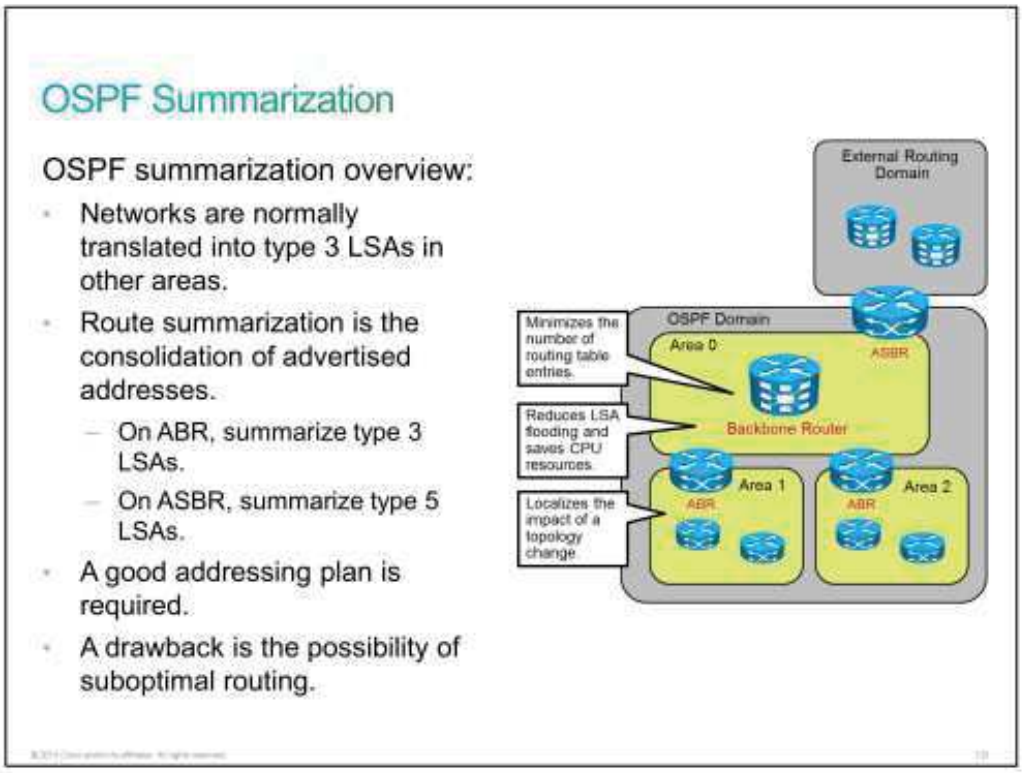

Normally, type 1 and type 2 LSAs are generated inside each area and translated into type 3 LSAs in other areas. With route summarization, the ABRs or ASBRs consolidate multiple routes into a single advertisement. ABRs summarize type 3 LSAs, and ASBRs summarize type 5 LSAs. Instead of advertising many specific prefixes, you need to advertise only one summary prefix.

If the OSPF design includes many ABRs or ASBRs, suboptimal routing is possible. This is one of the drawbacks of summarization.

Route summarization requires a good addressing plan—an assignment of subnets and addresses that is based on the OSPF area structure and lends itself to aggregation at the OSPF area borders.

https://t.me/learningnets

OSPF Interarea Route Summarization

The summarization of internal routes can be done only by ABRs. Without summarization, all the prefixes from an area are passed into the backbone as type 3 interarea routes. When summarization is enabled, the ABR intercepts this process and instead injects a single type 3 LSA, which describes the summary route into the backbone. Multiple routes inside the area are summarized.

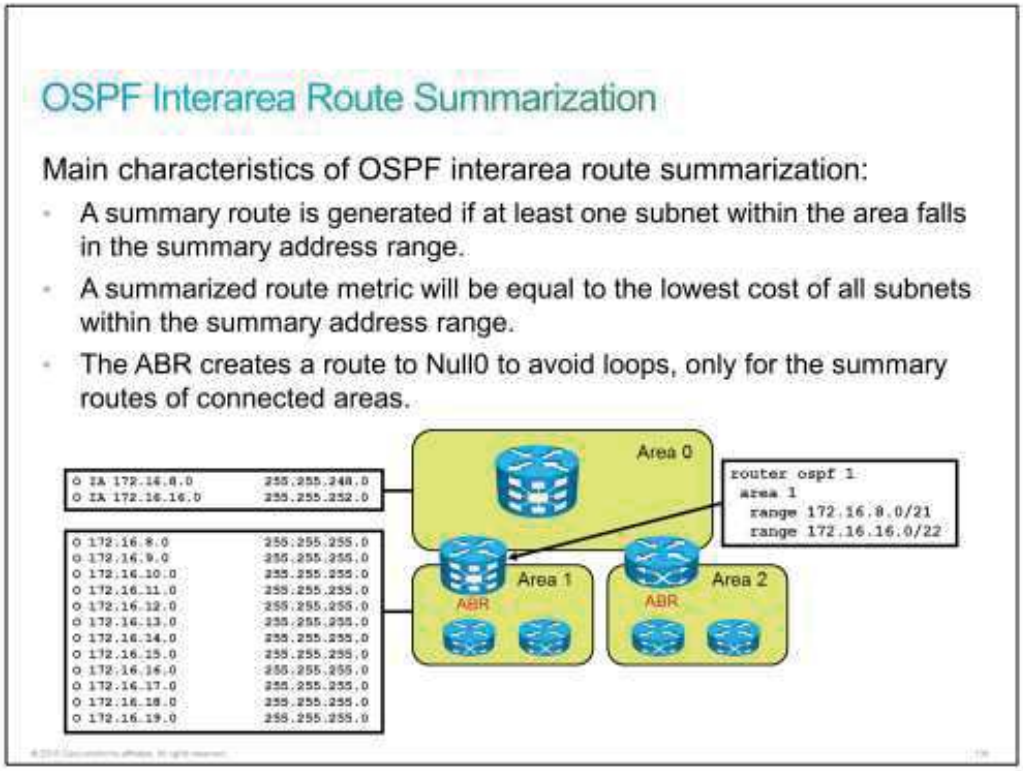

A summary route is generated if at least one subnet within the area falls in the summary address range. The summarized route metric is equal to the lowest cost of all the subnets within the summary address range. Interarea summarization can only be done for the intra-area routes of connected areas, and the ABR creates a route to Null0 to avoid loops in the absence of more specific routes.

Network numbers in areas should be assigned contiguously to ensure that these addresses can be summarized into a minimal number of summary addresses. In the figure, the list of 12 networks in the routing table of the ABR can be summarized into two summary address advertisements. The block of addresses from 172.16.8.0 through 172.16.15.0/24 can be summarized using 172.16.8.0/21, and the block from 172.16.16.0 through 172.16.19.0/24 can be summarized using 172.16.16.0/22.

To consolidate and summarize routes at an area boundary, use the **range** Cisco IOS XR router OSPF command, or **area range** Cisco IOS and IOS XE router OSPF command. The ABR will summarize routes for a specific area before injecting them into a different area via the backbone as type 3 summary LSAs. The OSPFv3 uses same command syntax as OSPFv2.

Cisco IOS, Cisco IOS XE, and Cisco IOS XR Software creates a summary route to the Null0 interface when manual summarization is configured, to prevent routing loops. For example, if the summarizing router receives a packet to an unknown subnet that is part of the summarized range, the packet matches the summary route based on the longest match. The packet is forwarded to the Null0 interface (in other words, it is dropped), which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF External Route Summarization

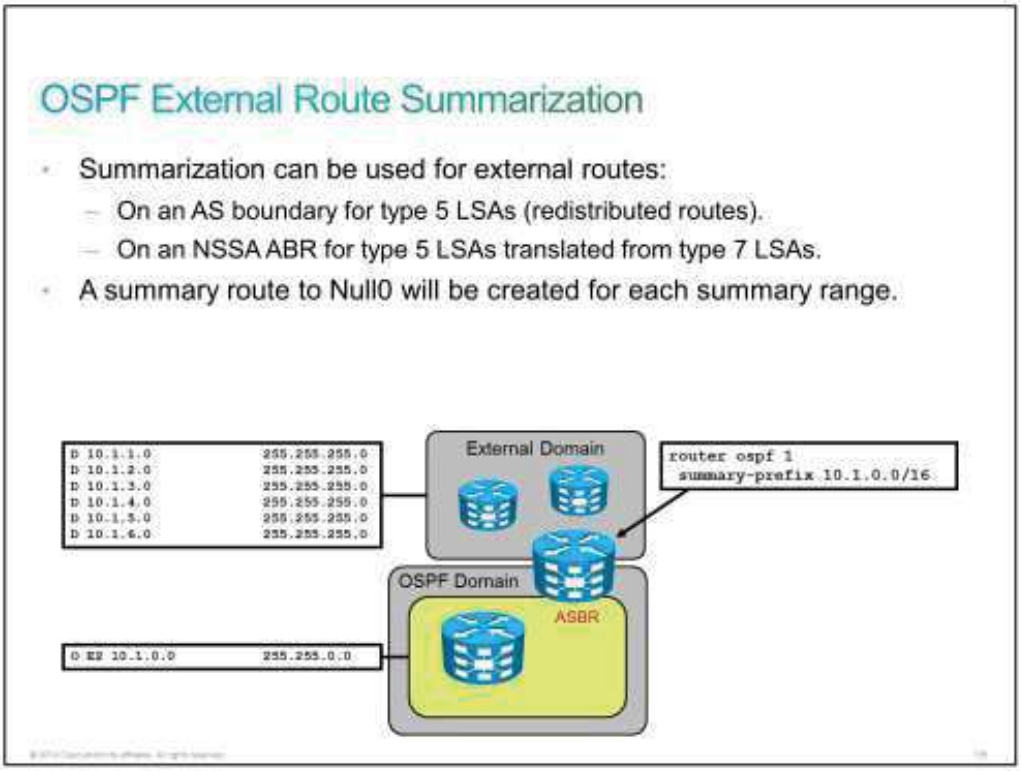

Summarization can also be used for external routes. Each route that is redistributed into OSPF from other protocols is advertised individually with an external LSA. To reduce the size of the OSPF LSDB, you can configure a summary for external routes. Summarization of external routes can be done on the ASBR for type 5 LSAs (redistributed routes) before injecting them into the OSPF domain. Summarization of external routes can be done by the ABR in NSSAs as well, where the ABR router creates type 5 summary routes from type 7 external routes. Without summarization, all the redistributed external prefixes from external autonomous systems are passed into the OSPF area. A summary route to Null0 is created automatically for each summary range.

To create aggregate addresses for OSPF at an AS boundary, use the **summary-prefix** Cisco IOS XR router OSPF command, or **summary-address** Cisco IOS and IOS XE router OSPF command. The OSPF version 3 uses same command syntax as OSPF version 2. The ASBR will summarize external routes before injecting them into the OSPF domain as type 5 external LSAs. The figure depicts route summarization on the ASBR. An external AS that is running EIGRP has its routes redistributed into OSPF.

https://t.me/learningnets

Default Routes in OSPF

To be able to perform routing from an OSPF AS toward external networks, you must either know all the destination networks or create a default route. To learn all the destinations, you can either create a number of static routes or configure redistribution. But the most scalable and optimized way is by using a default route.

The benefits of a default route include the following:

- Smaller routing table.
- Fewer resources and less CPU power needed; no need to recalculate the SPF algorithm if one or more networks fail.

To generate a default external route into an OSPF routing domain, use the **default-information originate** Cisco IOS, Cisco IOS XE, and Cisco IOS XR router OSPF configuration command. The OSPFv3 uses same command syntax as OSPFv2. There are two ways to advertise a default route into a standard area:

- Advertise 0.0.0.0/0 into the OSPF domain when the advertising router already has a default route. Use the **default-information originate** command to allow the ASBR to originate a type 5 default route inside the OSPF AS.

  You can also use different keywords in the configuration command to configure dependency on IP routing table entries.
- Advertise 0.0.0.0/0, regardless of whether the advertising router already has a default route. The second method can be accomplished by adding the keyword **always**. You can also use a route policy to define dependency on any condition inside the route policy.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

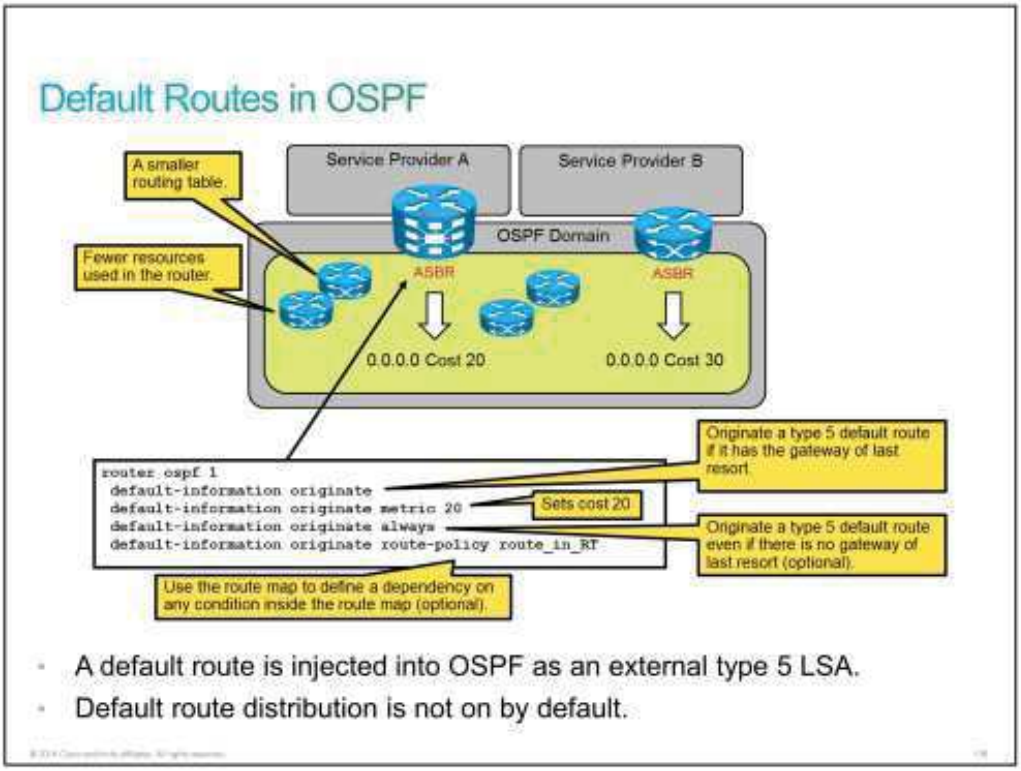

In the figure, an OSPF network is multihomed to two service providers. In this design, service provider A is preferred, and service provider B is used as a backup. To define the priority, the optional **metric** parameter has been used to establish a preference for the default route to service provider A.

https://t.me/learningnets

OSPF Area Types

OSPF is based on a two-level hierarchical area structure. The hierarchy defines backbone and nonbackbone areas. Each area has its own topology, which is invisible from outside the area. A router that belongs to several areas (an ABR) has several topology databases. All areas have to be connected to a backbone area or linked to it with a virtual link. The backbone area has to be contiguous. A nonbackbone area can be discontiguous.

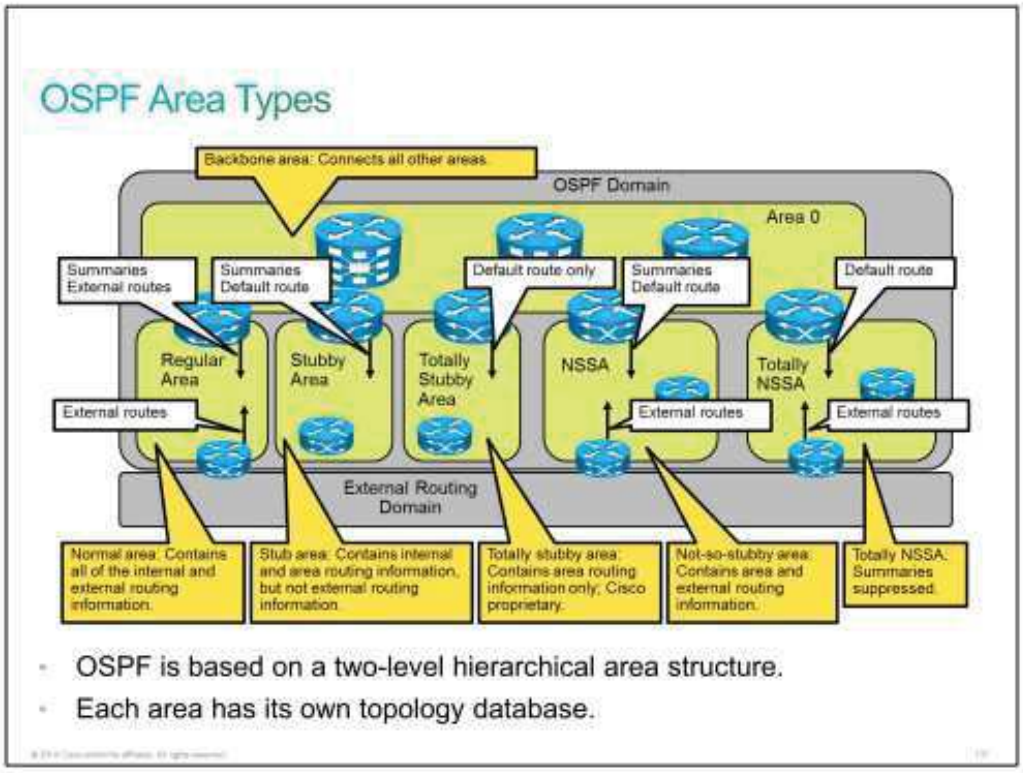

The characteristics that are assigned to an area control the type of route information that it receives. The possible area types are as follows:

- **Normal area:** This default area accepts link updates, route summaries, and external routes.
- **Backbone area (transit area):** The backbone area is the central entity to which all other areas connect. The backbone area is labeled Area 0. All other areas connect to this area to exchange and route information. The OSPF backbone includes all the properties of a standard OSPF area.
- **Stub area:** This area does not accept information about routes that are external to the AS, such as routes from non-OSPF sources. If routers need to route to networks outside the AS, they use a default route, which is noted as 0.0.0.0. Stub areas cannot contain ASBRs, except that ABRs may also be ASBRs.
- **Totally stubby area:** This area does not accept external AS routes or summary routes from other areas that are internal to the AS. If the router needs to send a packet to a network that is external to the area, it sends the packet using a default route. Totally stubby areas cannot contain ASBRs, except that ABRs may also be ASBRs.
- **NSSA:** NSSA is an addendum to the OSPF RFC. This area defines a special type 7 LSA. An NSSA offers benefits that are similar to those of a stub area or totally stubby area. However, NSSAs allow ASBRs, which are prohibited in a stub area.

https://t.me/learningnets

- **Totally NSSA:** This area does not accept external AS routes or summary routes from other areas that are internal to the AS. A totally NSSA contains area routing information and external (LSA 7) routing information, and is Cisco proprietary.

https://t.me/learningnets

Configuring OSPF Area Types

To configure different OSPF areas on the Cisco IOS XR router, use the commands as shown in the figure.

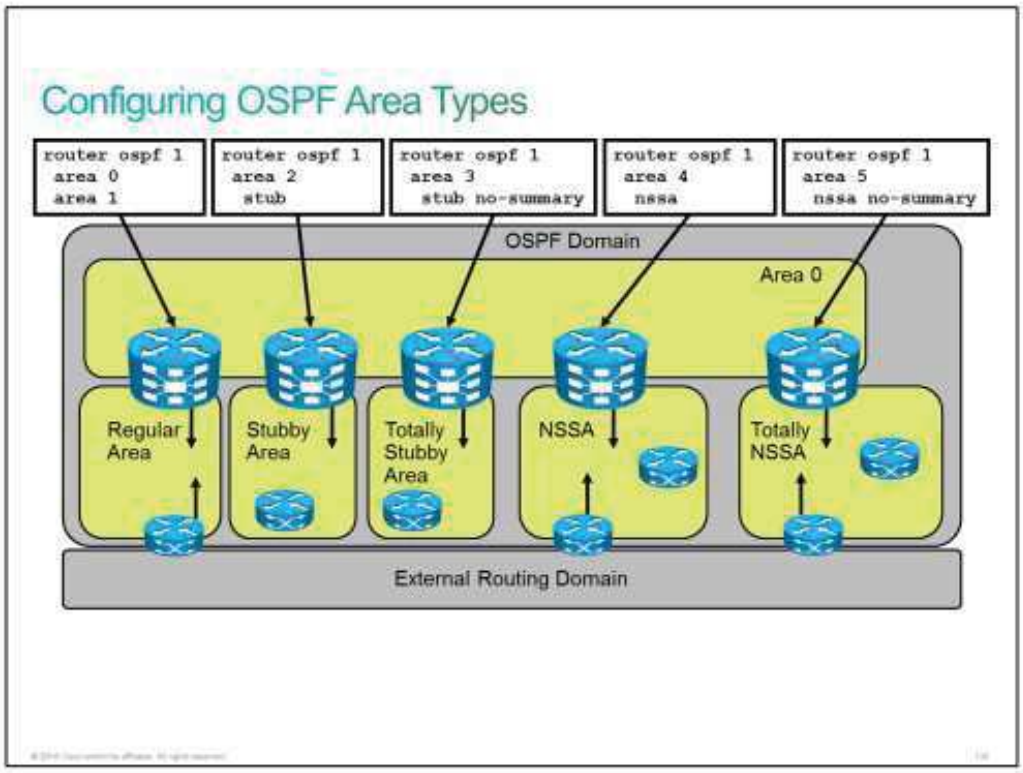

Use the **area 0** Cisco IOS XR router OSPF command, or **network Area 0** Cisco IOS and Cisco IOS XE router OSPF command to enable the backbone area. To configure the nonbackbone area, use the same command with a nonzero value. To enable stubby, totally stubby, NSSA, or totally NSSA areas, use the **stub [no-summary]** or **nssa [no-summary]** Cisco IOS XR router OSPF commands, or **area stub [no-summary]** or **area nssa [no-summary]** Cisco IOS and IOS XE router OSPF command.

https://t.me/learningnets

OSPF Router and LSA Types

A router originates a router-LSA for each area that it belongs to. Such an LSA describes the collected states of the router's links to the area. LSAs are crucial to OSPF and building the topology. LSA type depends on the router type that it originates from.

The OSPF defines three main router types:

- **Internal router:** This type of router resides inside any area and is a member of one area only.
- **ABR:** This type of router is a member of more than one area. As a border router, it has the ability to control routing traffic from one area to another. Different types of LSAs are exchanged between areas. ABRs can transmit these LSAs, or block them and send default routes instead.
- **ASBR:** This type of router is used to insert external routing information from another non-OSPF AS. ASBRs generate external LSAs, which can be blocked by ABRs.

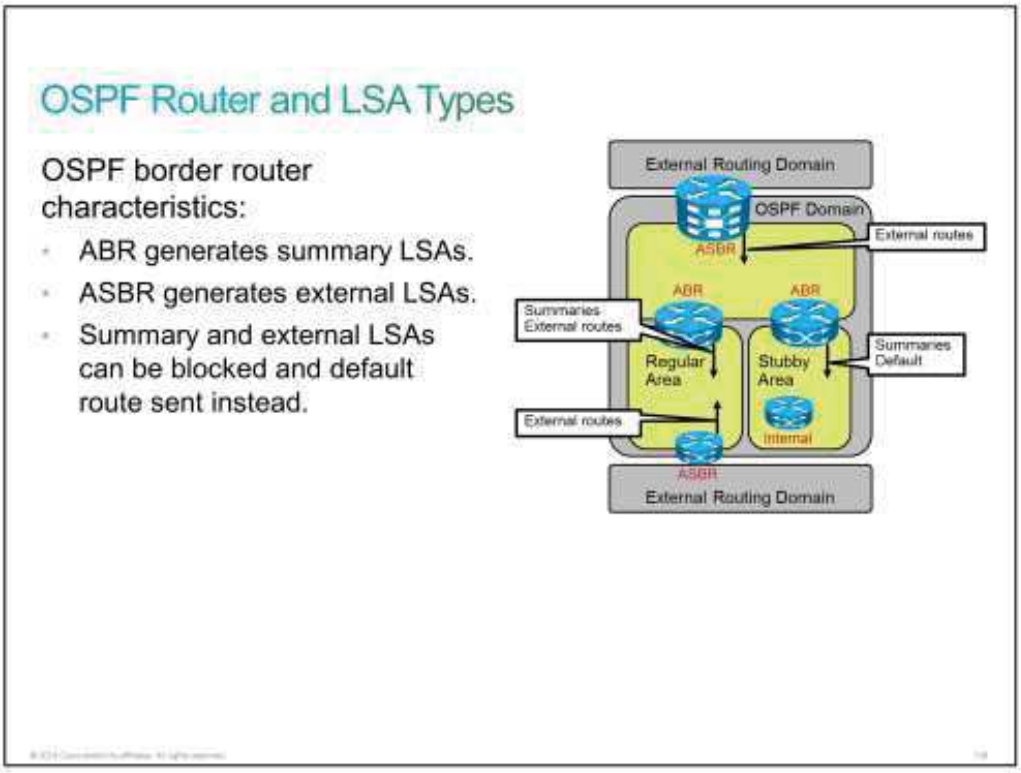

OSPF routers are able to generate and send the following six types of LSAs:

- Router link (type 1 LSA)
- Network link (type 2 LSA)
- Network summary (type 3 LSA)
- ASBR summary (type 4 LSA)
- External (type 5 LSA)
- NSSA external (type 7 LSA)

https://t.me/learningnets

OSPF Stub Area and Totally Stubby Area

Stub and totally stubby areas do not carry any external routes (type 5 LSAs). An area can be qualified as a stub or totally stubby area if it has the following characteristics:

- There is either a single exit point from that area, or, if there are multiple exits, one or more ABRs inject a default into the stub area and suboptimal routing paths are acceptable. Routing to other areas or autonomous systems could take a suboptimal path to reach the destination by exiting the area at a point that is farther from the destination than other exit points.
- All OSPF routers inside the stub area, including ABRs and internal routers, must be configured as stub routers before they can become neighbors and exchange routing information.
- There is no ASBR inside the area.
- The area is not the backbone area, Area 0.
- The area is not needed as a transit area for virtual links. Recall that a virtual link is a link that allows an area to connect to the backbone via a transit area. Virtual links are generally used for temporary connections or backup after a failure; they should not be considered as part of the primary OSPF design.

OSPF Stub Areas

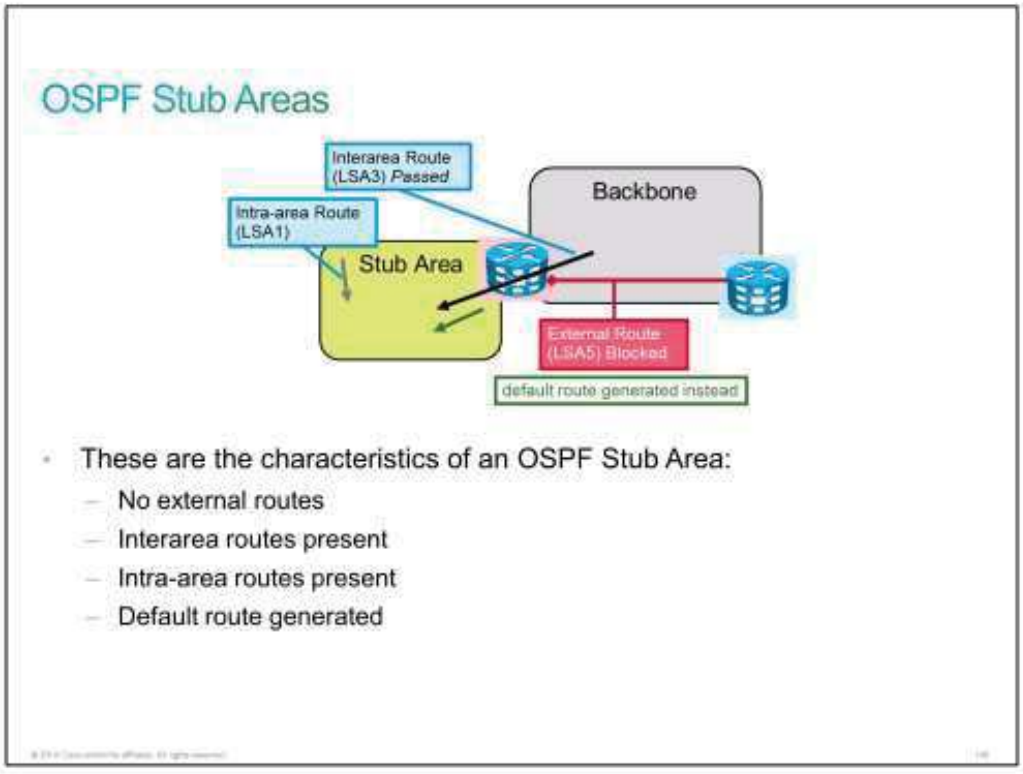

Configuring a *stub area* reduces the size of the LSDB inside the area, resulting in reduced memory requirements for routers in that area. External network LSAs (type 5), such as those that are redistributed from other routing protocols into OSPF, are not permitted to flood into a stub area.

https://t.me/learningnets

Routing from these areas to the outside is based on a default route (0.0.0.0). If a packet is addressed to a network that is not in the routing table of an internal router, the router automatically forwards the packet to the ABR, which sends a 0.0.0.0 LSA. Forwarding the packet to the ABR allows routers within the stub to reduce the size of their routing tables, because a single default route replaces many external routes.

A stub area is typically created when a hub-and-spoke topology is used, with each spoke being a stub area, such as a branch office. In this case, the branch office does not need to know about every network at the headquarters site, because it can use a default route to reach the networks.

OSPF Totally Stubby Area

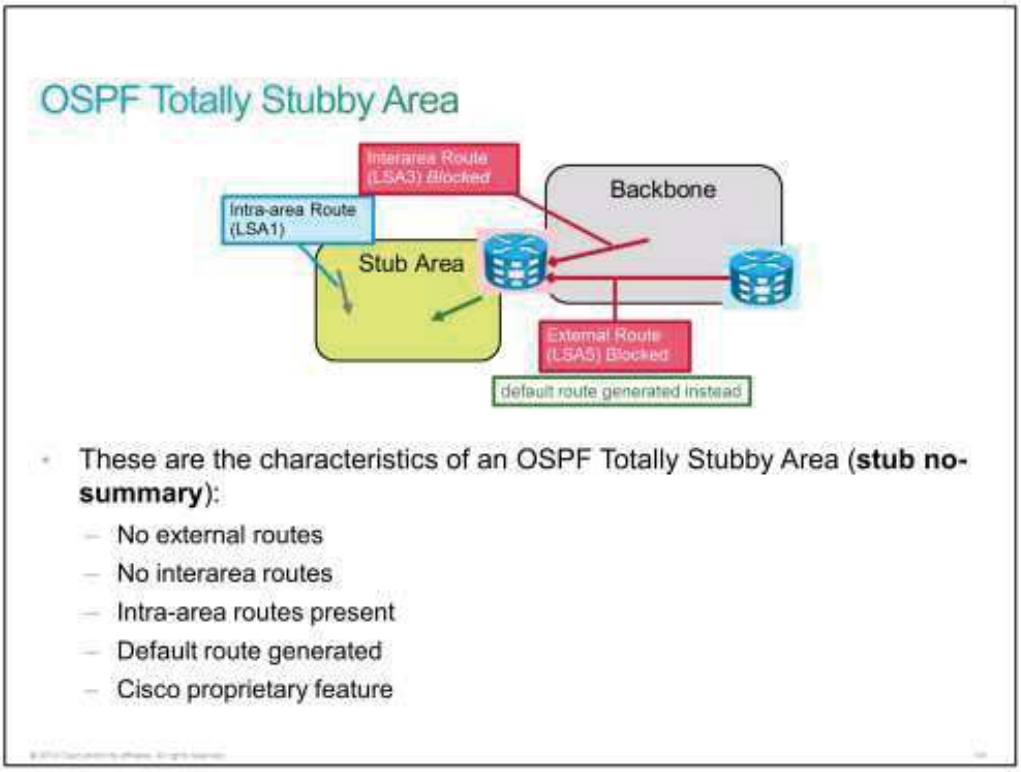

The *totally stubby area* technique is a Cisco proprietary enhancement that further reduces the number of routes in the routing table. A totally stubby area is a stub area that blocks external type 5 LSAs as well as summary type 3 and type 4 LSAs (interarea routes) from entering the area.

Because it blocks these routes, a totally stubby area recognizes only intra-area routes and the default route of 0.0.0.0. ABRs inject the default summary link 0.0.0.0 into the totally stubby area. Each router picks the closest ABR as a gateway to everything outside the area.

Totally stubby areas minimize routing information further than stub areas and increase the stability and scalability of OSPF internetworks. Using totally stubby areas is typically a better solution than using stub areas, as long as the ABR is a Cisco router.

https://t.me/learningnets

Configuring OSPF Stub Areas

When you have gathered all the required information for an implementation plan, you must complete the following tasks to configure OSPF stub areas:

- Configure all routers in the stub area as stub routers.
- Configure the cost for the default route on an ABR (optional).

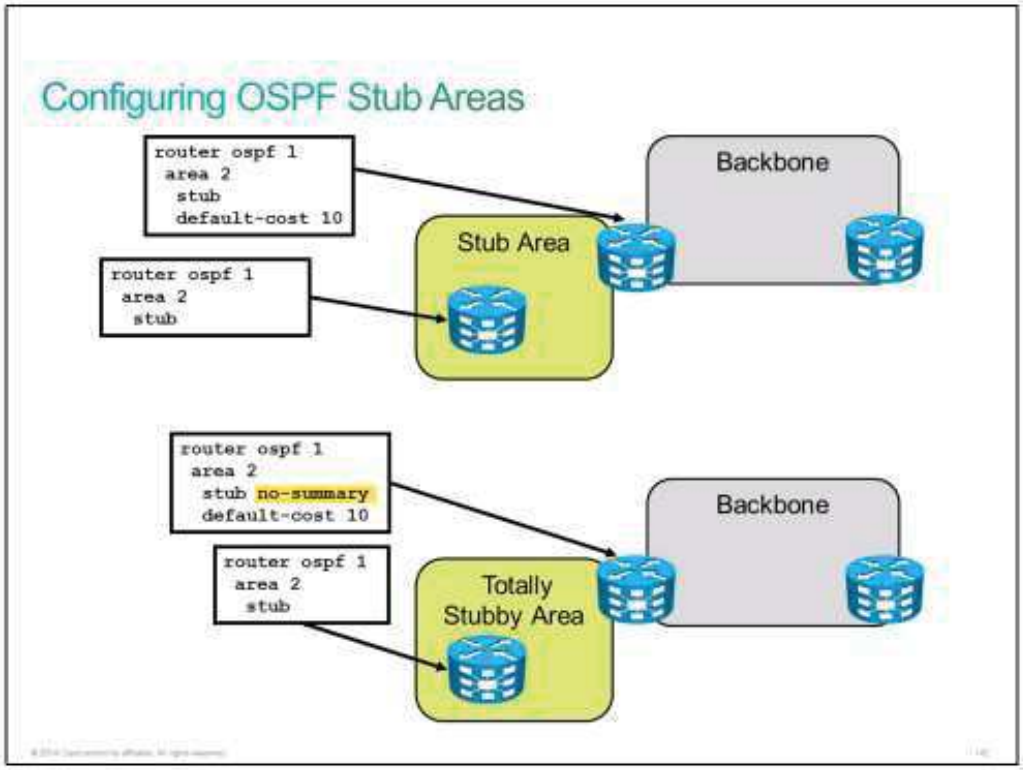

To configure an area as a stub, all routers inside the area must be configured as stub routers. The **stub** Cisco IOS XR router OSPF command, or **area stub** Cisco IOS and IOS XE router OSPF command is used to define an area as a stub area. By default, the ABR will advertise a default route with a cost of 1. You can change the cost of the default route by using the **default-cost** Cisco IOS XR router OSPF command, or **area default-cost** Cisco IOS and IOS XE router OSPF command. You only use this command on ABRs that are attached to stub areas or NSSAs.

The addition of the **no-summary** command on the ABR creates a totally stubby area and prevents all summary LSAs from entering the stub area.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

OSPF Not-So-Stubby Area and Totally Not-So-Stubby Area

The OSPF NSSA feature is described by RFC 3101 and was first introduced in Cisco IOS Software Release 11.2. It is a nonproprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

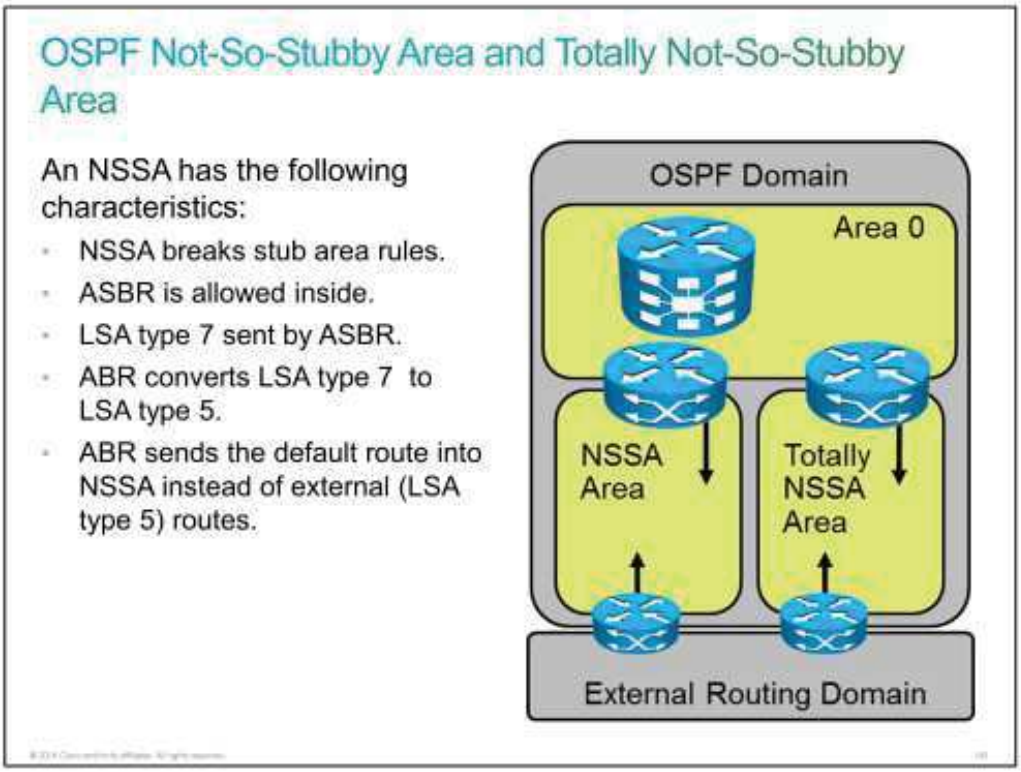

Redistribution into an NSSA creates a special type of LSA known as a type 7 LSA, which can exist only in an NSSA. An NSSA ASBR generates this LSA, and an NSSA ABR translates it into a type 5 LSA, which gets propagated into the OSPF domain. Type 7 LSAs have a propagate bit in the LSA header to prevent propagation loops between the NSSA and the backbone area. The NSSA retains the other stub area features; the ABR sends a default route into the NSSA instead of external routes from other ASBRs.

The type 7 LSA is described in the routing table as an O N2 or O N1 (N means NSSA). N1 means that the metric is calculated like external type 1 (E1); N2 means that the metric is calculated like external type 2 (E2). The default is O N2. The E1 metric adds external and internal costs together to reflect the whole cost to the destination. The E2 metric takes only the external cost, which is reflected in the OSPF cost.

The OSPF totally NSSA feature is an extension to the NSSA feature like the totally stubby feature is an extension to the stub area feature. It is a Cisco proprietary feature that blocks type 3, 4, and 5 LSAs. A single default route replaces both inbound-external (type 5) LSAs and summary (type 3 and 4) LSAs in the totally NSSA. The ABRs for the totally NSSA must be configured to prevent the flooding of summary routes for other areas into the NSSA. Only ABRs control the propagation of type 3 LSAs from the backbone. If an ABR is configured on any other routers in the area, it will have no effect at all.

https://t.me/learningnets

OSPF Not-So-Stubby Area and Totally Not-So-Stubby Area (Cont.)

- NSSA:
  - Behaves like stub area.
  - May introduce external routes locally in the area.
- Totally NSSA (**no-summary**):
  - Behaves like totally stubby area.
  - May introduce external routes locally in the area.
  - Cisco proprietary feature.

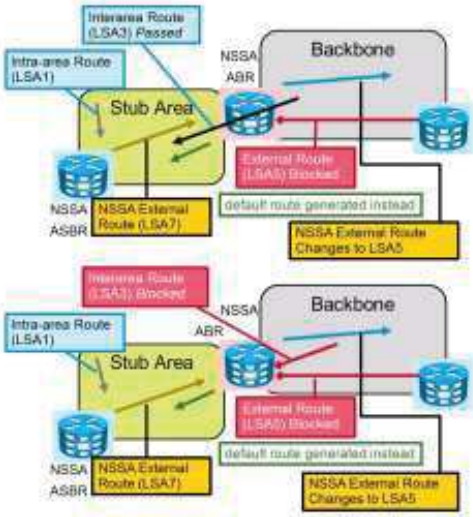

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Configuring OSPF NSSA and Totally NSSA

When you have gathered all the required information for an implementation plan, you must complete the following tasks to configure OSPF NSSA:

- Configure all routers in the NSSA as NSSA routers.
- Configure the cost for the default route on an ABR (optional).

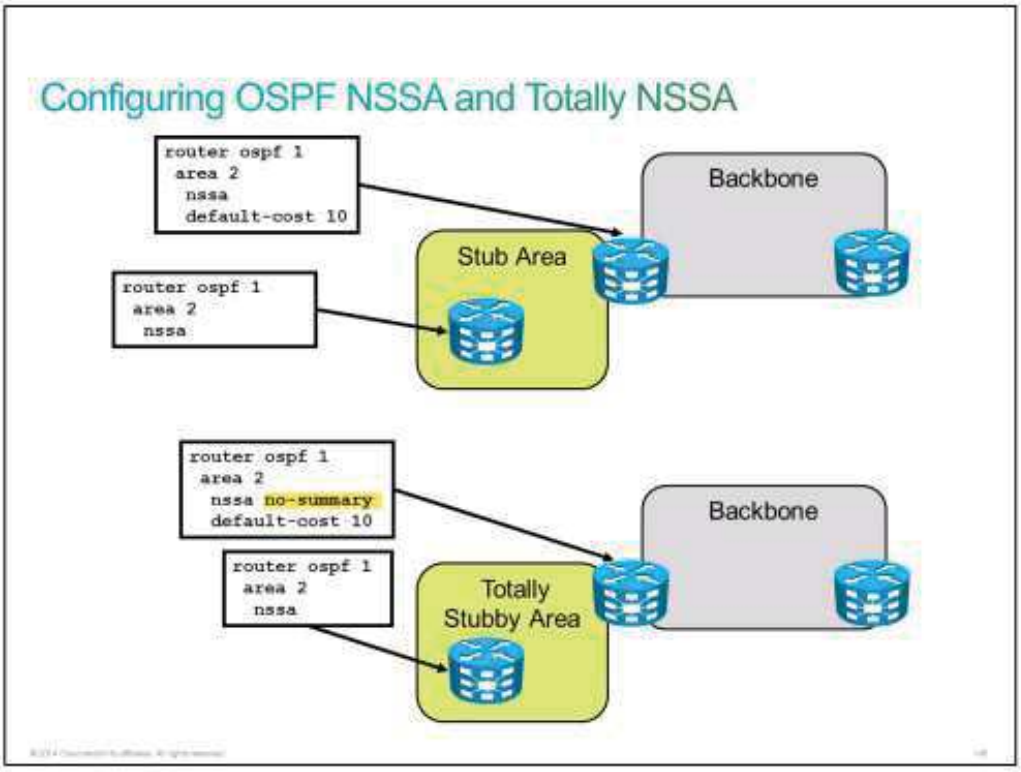

To configure an area as an NSSA, all routers inside the area must be configured as NSSA routers. The **nssa** Cisco IOS XR router OSPF command, or **area nssa** Cisco IOS and IOS XE router OSPF command, is used to define an area as an NSSA. By default, the ABR will advertise a default route with a cost of 1. You can change the cost of the default route by using the **default-cost** Cisco IOS XR router OSPF command, or **area default-cost** Cisco IOS and IOS XE router OSPF command. You only use this command on ABRs that are attached to NSSAs.

The addition of the **no-summary** command on the ABR creates a totally NSSA and prevents all summary LSAs from entering the NSSA.

https://t.me/learningnets

Summary

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

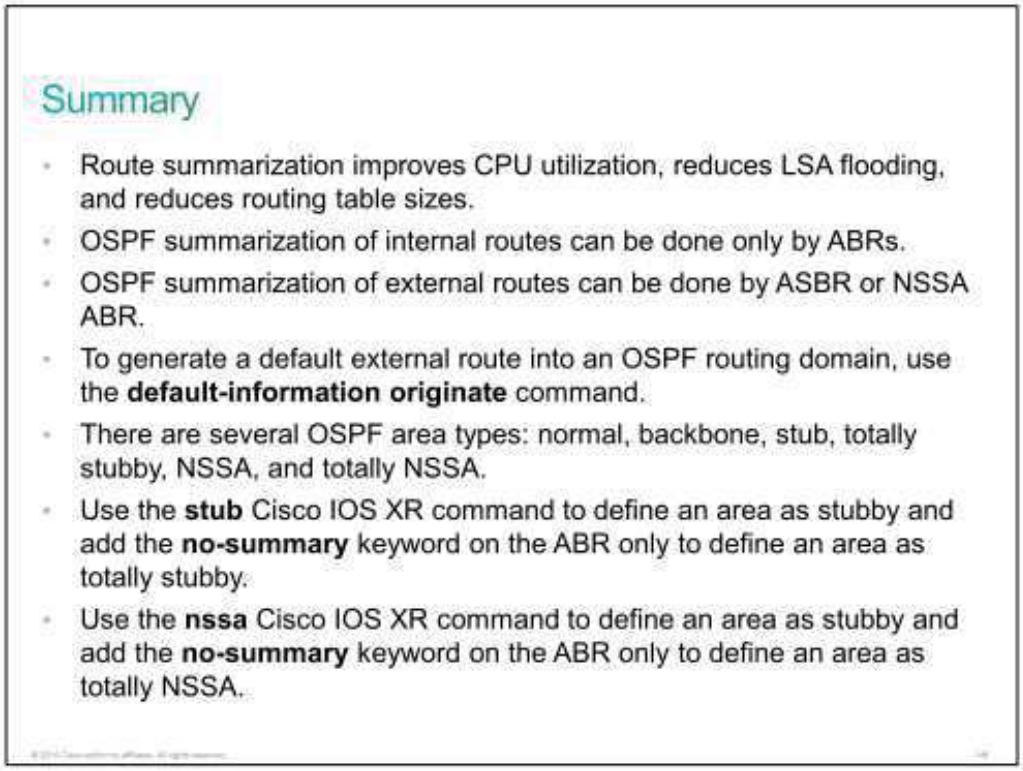

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 5: Module Summary

Overview

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

OSPF is one of the most commonly used interior gateway protocols in service provider networks. OSPF is a complex, open-standard protocol that is composed of several protocol handshakes, database advertisements, and packet types.

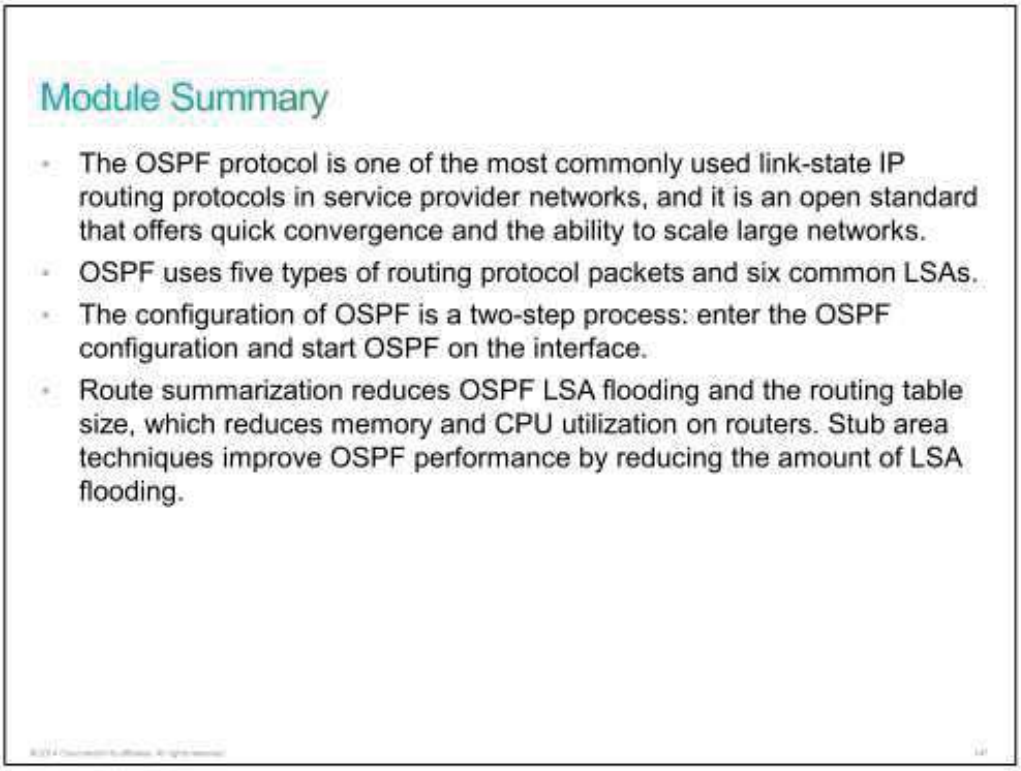

https://t.me/learningnets

 ©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. All of these tables are maintained by a link-state routing protocol except which one? (Source: Introducing OSPF Routing)

   A. Routing
   B. Topology
   C. Update
   D. Neighbor

2. The memory that is needed to maintain tables is one disadvantage of link-state protocols. (Source: Introducing OSPF Routing)

   A. True
   B. False

3. Match each table to its function. (Source: Introducing OSPF Routing)

Neighbor	Stores best paths
Topology	Stores LSAs
Routing	Stores adjacencies

https://t.me/learningnets

4. Which term refers to the router that connects Area 0 to a nonbackbone area? (Source: Introducing OSPF Routing)

   A. Area Boundary Router
   B. Area Border Router
   C. Autonomous System Boundary Router
   D. Backbone router

5. Which OSPF packet helps form neighbor adjacencies? (Source: Introducing OSPF Routing)

   A. Exchange packet
   B. Hello packet
   C. Neighbor discovery packet
   D. Adjacency packet

6. Which criterion does SPF use to determine the best path? (Source: Introducing OSPF Routing)

   A. Lowest delay
   B. Highest bandwidth
   C. Lowest total cost of the route
   D. Total bandwidth of the route

7. Which table is populated as a result of SPF calculations? (Source: Introducing OSPF Routing)

   A. Topology
   B. Routing
   C. Adjacency
   D. Neighbor

8. Cisco recommends no more than _____ area or areas per ABR in addition to Area 0. (Source: Introducing OSPF Routing)

   A. One
   B. Two
   C. Four
   D. Eight

9. An area border router maintains _____. (Source: Introducing OSPF Routing)

   A. A separate database for each area with which it is connected
   B. A single database for all areas
   C. Two databases: one for the backbone and one for all other areas
   D. A separate routing table for each area

10. In a multiarea network, any area can be the backbone area, although it is most often Area 0. (Source: Introducing OSPF Routing)

    A. True
    B. False

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

11. When an OSPF router receives an LSA, it is installed in the _____. (Source: Introducing OSPF Routing)

A. Neighbor table
B. Topology table
C. Routing table
D. Update table

12. An OSPF router receives an LSA and checks the sequence number of the LSA. This number matches the sequence number of an LSA that the receiving router already has. What does the receiving router do with the received LSA? (Source: Introducing OSPF Routing)

A. Ignores the LSA
B. Adds the LSA to the database
C. Sends the newer LSU to the source router
D. Floods the LSA to the other routers

13. An OSPF router receives an LSA. The router checks the sequence number of the LSA and finds that this number is higher than the sequence number that it already has. Which two tasks does the router perform with the LSA? (Choose two.) (Source: Introducing OSPF Routing)

A. Ignores the LSA
B. Adds the LSA to the database
C. Sends the newer LSU to the source router
D. Floods the LSA to the other routers

14. An OSPF router receives an LSA. The router checks the sequence number of the LSA and finds that this number is *lower* than the sequence number that it already has. What does the router do with the LSA? (Source: Introducing OSPF Routing)

A. Ignores the LSA
B. Adds the LSA to the database
C. Sends the newer LSU to the source router
D. Floods the LSA to the other routers

15. Each LSA has its own age timer. By default, how long does an LSA wait before requiring an update? (Source: Introducing OSPF Routing)

A. 30 seconds
B. 1 minute
C. 30 minutes
D. 1 hour

16. OSPF does not require a Hello protocol on point-to-point links because adjacent routers are directly connected. (Source: Understanding OSPF Operation)

A. True
B. False

https://t.me/learningnets

17. When the DR fails, the BDR automatically builds new adjacencies, exchanges databases with other routers, and takes over as the DR. (Source: Understanding OSPF Operation)

A. True
B. False

18. Which destination IP address does OSPF use when advertising to all SPF routers? (Source: Understanding OSPF Operation)

A. 224.0.0.6
B. 224.0.0.5
C. 255.255.255.255
D. The IP address of the output interface

19. The BDR, like the DR, maintains a complete set of adjacencies on a broadcast link. (Source: Understanding OSPF Operation)

A. True
B. False

20. What is the IP protocol number for OSPF packets? (Source: Understanding OSPF Operation)

A. 20
B. 76
C. 86
D. 89

21. All of these are OSPF packet types except which one? (Source: Understanding OSPF Operation)

A. LSU
B. LSR
C. DBD
D. LSAck
E. Hello
F. Query

22. Which multicast address does the OSPF Hello protocol use? (Source: Understanding OSPF Operation)

A. 224.0.0.5
B. 224.0.0.6
C. 224.0.0.7
D. 224.0.0.8

23. The Hello protocol sends periodic messages to ensure that a neighbor relationship is maintained between adjacent routers. (Source: Introducing OSPF Routing)

A. True
B. False

https://t.me/learningnets

24. Place the exchange protocol states in the correct order. (Source: Introducing OSPF Routing)

Down

Init

Two-way

Exstart

Exchange

Loading

Full

25. DBD packets are involved during which two states? (Choose two.) (Source: Understanding OSPF Operation)

A. Exstart
B. Loading
C. Exchange
D. Two-way

26. At which interval does OSPF refresh LSAs? (Source: Understanding OSPF Operation)

A. 10 seconds
B. 30 seconds
C. 30 minutes
D. 1 hour

27. All of these fields are in an OSPF packet header except which one? (Source: Understanding OSPF Operation)

A. Packet length
B. Router ID
C. Authentication type
D. MaxAge time

28. Three routers are connected to an Ethernet LAN. One is a small router that should not take on the role of the DR or BDR. How do you ensure that it never will? (Source: Understanding OSPF Operation)

A. Set the interface priority to 100.
B. Set the interface priority to 0.
C. Leave the interface priority set to 1, and set the priority of the other two routers to 10
D. Use the **no designated-router** command on the Ethernet interface.

https://t.me/learningnets

29. What is the default hello interval for NBMA interfaces? (Source: Understanding OSPF Operation)

A. 10 seconds
B. 30 seconds
C. 60 seconds
D. 120 seconds

30. An OSPF router automatically builds adjacencies with neighboring routers on an NBMA link. (Source: Understanding OSPF Operation)

A. True
B. False

31. Which mode of OSPF operation is RFC-compliant? (Source: Understanding OSPF Operation)

A. Point-to-multipoint nonbroadcast
B. Point-to-multipoint
C. Broadcast
D. Point-to-point

32. Which two statements regarding an OSPF adjacency that is built over a Layer 2 MPLS VPN backbone are true? (Choose two.) (Source: Understanding OSPF Operation)

A. The DR and BDR are elected.
B. The OSPF parameters must be agreed upon with the service provider
C. Provider edge routers appear as additional routers in the customer network.
D. The OSPF network type is multiaccess broadcast.

33. What are the three types of networks that are defined by OSPF? (Choose three.) (Source: Understanding OSPF Operation)

A. Point-to-point
B. Broadcast
C. Point-to-multipoint
D. Point-to-multipoint nonbroadcast
E. Nonbroadcast multiaccess

34. When an OSPF adjacency is built over a Layer 3 MPLS VPN backbone, customer routers are unaware of the MPLS VPN topology. (Source: Understanding OSPF Operation)

A. True
B. False

35. Which two statements regarding OSPF nonbroadcast mode are correct? (Choose two.) (Source: Understanding OSPF Operation)

A. This mode requires manual **neighbor** commands.
B. This mode does not use a DR and BDR.
C. This mode uses a DR and BDR.
D. This mode requires multiple subnets.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

36. When you are using an OSPF **neighbor** Cisco IOS command, you must configure it under an interface. (Source: Understanding OSPF Operation)

A. True
B. False

37. Which OSPF mode requires **neighbor** commands? (Source: Understanding OSPF Operation)

A. Broadcast
B. Point-to-point
C. Point-to-multipoint
D. Point-to-multipoint nonbroadcast

38. Which OSPF **show** command describes a list of OSPF adjacencies? (Source: Implementing OSPF Routing)

A. **show ospf interface**
B. **show ospf**
C. **show route**
D. **show ospf neighbor**

39. All of these techniques are used for router ID selection except which one? (Source: Implementing OSPF Routing)

A. Highest IP address on an interface
B. IP address on a loopback interface
C. Lowest IP address when multiple loopback interfaces are used
D. The **router-id** command

40. When you use the **router-id** command, the router ID immediately changes to the IP address that has been entered. (Source: Implementing OSPF Routing)

A. True
B. False

41. Which Cisco IOS and IOS XE network statement is used to configure OSPF on an interface with an IP address of 172.16.1.1 in Area 0? (Source: Implementing OSPF Routing)

A. **network 172.16.0.0 0.0.0.255 area 0**
B. **network 172.16.1.1 0.0.0.0 area 0**
C. **network 172.16.1.1 255.255.255.255 area 0**
D. **network 172.16.0.0 0.0.255.255 area 0**

42. Only one OSPF process can run on a Cisco router at one time. (Source: Implementing OSPF Routing)

A. True
B. False

https://t.me/learningnets

43. A router has a Fast Ethernet interface with an IP address of 172.16.45.1, a Loopback 0 interface with an IP address of 10.3.3.3, a Loopback 1 interface with an IP address of 10.2.2.2, and a router-id command with an IP address of 10.1.1.1. Which router ID will be selected? (Source: Implementing OSPF Routing)

A. 172.16.45.1
B. 10.3.3.3
C. 10.2.2.2
D. 10.1.1.1

44. The **show ospf neighbor** command shows the FULL state on one of the two neighbors in its table. Which neighbor or neighbors will successfully exchange LSDB information? (Source: Implementing OSPF Routing)

A. A neighbor that is in the FULL state
B. A neighbor that is not in FULL state
C. Any neighbor, regardless of whether or not it is in the FULL state
D. No neighbor, regardless of whether or not it is in the FULL state

45. Which two Cisco IOS XR **show** commands can be used to verify the OSPF router ID of a router? (Choose two.) (Source: Implementing OSPF Routing)

A. **show ospf interface**
B. **show ospf neighbor**
C. **show ospf**
D. **show route**

46. Where will a type 3 LSA be sent? (Source: Implementing OSPF Routing)

A. Only within the area that it originates from
B. Within the area that it originates from plus the backbone area
C. Within the area that it originates from and between all other areas
D. Within the backbone area and between all other areas

47. A network uses Gigabit Ethernet, and you want OSPF to correctly calculate the metric using bandwidth. Which command should you use to ensure that this happens? (Source: Implementing OSPF Routing)

A. **ip ospf cost** on the interface
B. **auto-cost reference-bandwidth** under the OSPF routing process
C. **bandwidth** under the interface
D. **bandwidth** under the OSPF routing process

48. Looking at the routing table, you notice "**[110/55]**". What does this mean? (Source: Implementing OSPF Routing)

A. The O E1 cost is 110, and the O E2 cost is 55.
B. The administrative distance is 110, and the metric is 55.
C. The administrative distance is 55, and the metric is 110.
D. The total cost of the route is 165.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

49. What does it mean if a route in the routing table has an indicator of O? (Source: Implementing OSPF Routing)

A. It is intra-area.
B. It is interarea.
C. It is external.
D. It is a stub.

50. What is the difference between a type 3 LSA and a type 4 LSA? (Source: Implementing OSPF Routing)

A. A type 3 LSA is a summary LSA, and a type 4 LSA is E1.
B. A type 3 LSA is E1, and a type 4 LSA is a summary.
C. A type 3 LSA is a summary for networks, and a type 4 LSA is a summary for ASBRs.
D. A type 3 LSA is a summary for ASBRs, and a type 4 LSA is a summary for networks.

51. When OSPF authentication is configured between two routers, each router has its own unique password. (Source: Implementing OSPF Routing)

A. True
B. False

52. Which three options are used to generate the message digest when OSPF MD5 authentication is configured? (Choose three.) (Source: Implementing OSPF Routing)

A. Packet
B. Sequence number
C. Key ID
D. Key
E. Router ID

53. Which Cisco IOS and IOS XE command is used to specify that OSPF MD5 authentication will be used? (Source: Implementing OSPF Routing)

A. **ip ospf authentication simple**
B. **ip ospf authentication**
C. **ip ospf authentication-key**
D. **ip ospf message-digest-key**
E. **ip ospf authentication message-digest**

54. What are the two reasons why route summarization is important? (Choose two.) (Source: Implementing OSPF Special Area Types)

A. Reduces type 1 LSA flooding
B. Reduces type 3 LSA flooding
C. Reduces the size of the routing table
D. Reduces the size of the neighbor table

https://t.me/learningnets

55. Which two features play a key role in route summarization? (Choose two.) (Source: Implementing OSPF Special Area Types)

A. Contiguous IP addressing
B. Discontiguous IP addressing
C. FLSM
D. VLSM

56. Which Cisco IOS XR command would you use to summarize routes into Area 0 from the ABR? (Source: Implementing OSPF Special Area Types)

A. **summary-address**
B. **range**
C. **network**
D. **summary**

57. Which command would you use to summarize routes into OSPF from the ASBR? (Source: Implementing OSPF Special Area Types)

A. **summary-address**
B. **range**
C. **network**
D. **summary**

58. A default route is identified in the OSPF database as a _____. (Source: Implementing OSPF Special Area Types)

A. Type 1 LSA
B. Type 2 LSA
C. Type 3 LSA
D. Type 4 LSA
E. Type 5 LSA

59. The primary purpose of a default route is to reduce the sizes of the routing table and the LSDB. A default route avoids detailed updating of routes by inserting a single 0.0.0.0 route into the routing table, making this 0.0.0.0 route act as a gateway of last resort. (Source: Implementing OSPF Special Area)

A. True
B. False

60. When should you use the **always** keyword with the **default-information originate** command? (Source: Implementing OSPF Special Area Types)

A. Only by default; configuration not required.
B. When you want to send summarized routes.
C. When your default route is always in the routing table.
D. When you want the default route advertised, even if it is not in the routing table.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

61. A summary LSA (type 3 LSA) is designed to automatically summarize a network into blocks. (Source: Implementing OSPF Special Area Types)

A. True
B. False

62. Route summarization reduces the flooding of which two LSA types? (Choose two.) (Source: Implementing OSPF Special Area Types)

A. Router
B. Network
C. Summary
D. External
E. NSSA

63. You are at the ABR of Area 1 and want to classfully summarize network 172.16.32.0 through 172.16.63.0 into Area 0. What is correct Cisco IOS XR configuration commands that you would use? (Source: Implementing OSPF Special Area Types)

A. area 1 range 172.16.0.0 255.255.0.0
B. area 0 range 172.16.0.0 255.255.0.0
C. area 0 range 172.16.0.0 0.0.255.255
D. area 1 range 172.16.0.0 0.0.255.255

64. You are at the ASBR between an OSPF Area 0 and an EIGRP network. EIGRP routes are being redistributed into OSPF. What is the correct summarization Cisco IOS XR command to summarize the EIGRP block 172.16.32.0 through 172.16.63.0? (Source: Implementing OSPF Special Area Types)

A. summary-address 172.16.32.0 255.255.224.0
B. summary-address 172.16.32.0 0.0.31.255
C. ip summary-address 172.16.32.0 0.0.31.255
D. ip summary-address 172.16.32.0 255.255.224.0
E. ip eigrp 100 summary-address 172.16.32.0 0.0.31.255
F. ip eigrp 100 summary-address 172.16.32.0 255.255.224.0

65. It is important that you always summarize the routes from Area 0 into other areas. Suboptimal path selection can occur if you do not. (Source: Implementing OSPF Special Area Types)

A. True
B. False

66. All of these are permitted in a stub area except which one? (Source: Implementing OSPF Special Area Types)

A. An ABR
B. An ASBR
C. Summary routes
D. Summary LSAs

https://t.me/learningnets

67. Which type of router advertises the default into a stub area? (Source: Implementing OSPF Special Area Types)

A. ASBR
B. Backbone router
C. ABR
D. Internal router

68. What is the meaning of the **no-summary** parameter of the **stub** command? (Source: Implementing OSPF Special Area Types)

A. There is no route summarization in the stub area.
B. No summary LSAs are sent into the stub area.
C. No type 5 LSAs are sent into the stub area.
D. There are no external LSAs in the stub area.

69. The default route has a cost of 1 from stub area ABR if no **default-cost** command is used. (Source: Implementing OSPF Special Area Types)

A. True
B. False

70. A disadvantage of NSSA is that it does not have a totally stubby feature, such as a normal stub area. (Source: Implementing OSPF Special Area Types)

A. True
B. False

71. Which characteristic is not a prerequisite for stub areas? (Source: Implementing OSPF Special Area Types)

A. Virtual links not allowed
B. ASBRs not allowed
C. ABRs not allowed
D. One way in and out of the stub area

72. Stub area design will not improve _____. (Source: Implementing OSPF Special Area Types)

A. CPU utilization on routers in the stub
B. Memory requirements on routers in the stub
C. Ability to reach outside networks
D. LSDB size on routers in the stub

73. A type 7 LSA appears in the routing table as an _____. (Source: Implementing OSPF Special Area Types)

A. O E1 route
B. O E2 route
C. O N2 route
D. O I/A route

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

74. What is the difference between a stub area and a totally stubby area configuration? (Source: Implementing OSPF Special Area Types)

A. **no-summary** option on the ABR
B. **totally-stubby** command on the internal routers
C. **nssa** command on the internal routers
D. **default-cost** command on the ABR

75. A stub area blocks summary LSAs (types 3 and 4 LSAs). (Source: Implementing OSPF Special Area Types)

A. True
B. False

76. Where should you configure the **stub** command when you are configuring a stub area? (Source: Implementing OSPF Special Area Types)

A. On all routers in the area
B. On the ABR
C. On the ASBR
D. On routers that require stub capability within the area

77. In NSSA, the NSSA ABR translates type 7 LSAs into type 5 LSAs. (Source: Implementing OSPF Special Area Types)

A. True
B. False

78. The ABR injects a default route into which three types of areas? (Choose three.) (Source: Implementing OSPF Special Area Types)

A. Stub
B. Totally stubby NSSA
C. Totally stubby
D. Area 0

https://t.me/learningnets

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. C
2. A
3. Routing — Stores best paths
   Topology — Stores LSAs
   Neighbor — Stores adjacencies
4. B
5. B
6. C
7. B
8. B
9. A
10. B
11. B
12. A
13. B, D
14. C
15. C
16. B
17. B
18. B
19. A
20. D
21. F

https://t.me/learningnets

22. A
23. A
24. — Down
— Init
— Two-way
— Exstart
— Exchange
— Loading
— Full
25. A, C
26. C
27. D
28. B
29. B
30. B
31. B
32. A, D
33. A, B, E
34. A
35. A, C
36. B
37. D
38. D
39. C
40. B
41. B
42. B
43. D
44. A
45. A, C
46. D
47. B
48. B
49. A
50. C
51. B
52. A, C, D

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

53. E
54. B, C
55. A, D
56. B
57. A
58. E
59. A
60. D
61. B
62. C, D
63. A
64. A
65. B
66. B
67. C
68. B
69. A
70. B
71. C
72. C
73. C
74. A
75. B
76. A
77. A
78. A, B, C

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 3: Implement Integrated IS-IS in the Service Provider Network

Introduction

The Integrated IS-IS protocol is a part of the OSI suite of protocols. This module describes the importance of the Integrated IS-IS routing protocol for internal service provider routing, and lists the steps to take when you are implementing Integrated IS-IS into the service provider network.

The OSI suite uses CLNS to provide connectionless delivery of data, and the actual Layer 3 protocol is CLNP. CLNP is the solution for unreliable (connectionless) delivery of data, similar to IP. IS-IS uses CLNS addresses to identify the routers and to build the LSDB.

IS-IS operates in strictly CLNS terms; however, Integrated IS-IS supports IP routing as well as CLNS. CLNS addresses are required to configure and troubleshoot IS-IS, even when it is used only for IP.

IS-IS supports the most important characteristics of OSPF and EIGRP, because it supports VLSM and converges quickly. Each protocol has advantages and disadvantages, but this commonality makes any of the three protocols scalable and appropriate for supporting the large-scale networks of today.

Upon completing this module, you will be able to:

- Explain the features, benefits, and operation of the Integrated IS-IS protocol that runs in service provider networks
- Implement Integrated IS-IS in a service provider network

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Introducing IS-IS Routing

Overview

IS-IS is a proven and extensible IP routing protocol that converges quickly and supports VLSM. IS-IS is a public standard, published as ISO 9542 and republished as RFC 995. Integrated IS-IS is specified in RFC 1195 and offers support for IP and OSI protocols. Although not as common, IS-IS is comparable to, and in some cases preferable to, OSPF protocol.

Upon completing this lesson, you will be able to:

- Describe Integrated IS-IS in the Cisco IP NGN
- Describe Integrated IS-IS design principles
- Describe the best practices for Integrated IS-IS designs
- Compare IS-IS and OSPF
- Describe CLNS addressing as it is used in proper IS-IS deployment
- Describe the IS-IS router types used in the IS-IS-enabled network
- Describe the routing logic used in IS-IS networks
- Describe asymmetric routing in IS-IS networks
- Describe symmetric routing in IS-IS networks
- Describe the IS-IS packet structure and the importance of different TLVs
- Describe the TLVs used to support IPv6 in IS-IS networks
- Describe the IS-IS network types
- Describe IS-IS operations in broadcast networks vs. point-to-point networks
- Describe the LSP flooding process in IS-IS networks

https://t.me/learningnets

- Describe the IS-IS link-state database synchronization process
- Describe IS-IS adjacencies
- Describe the design considerations for IS-IS networks running both IPv4 and IPv6
- Describe multitopology IS-IS for IPv6

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

IS-IS in the Cisco IP NGN

IS-IS is a popular IP routing protocol in the service provider industry. The simplicity and stability of IS-IS make it robust in large internetworks. IS-IS is found in large service providers and in some networks that support OSI protocols.

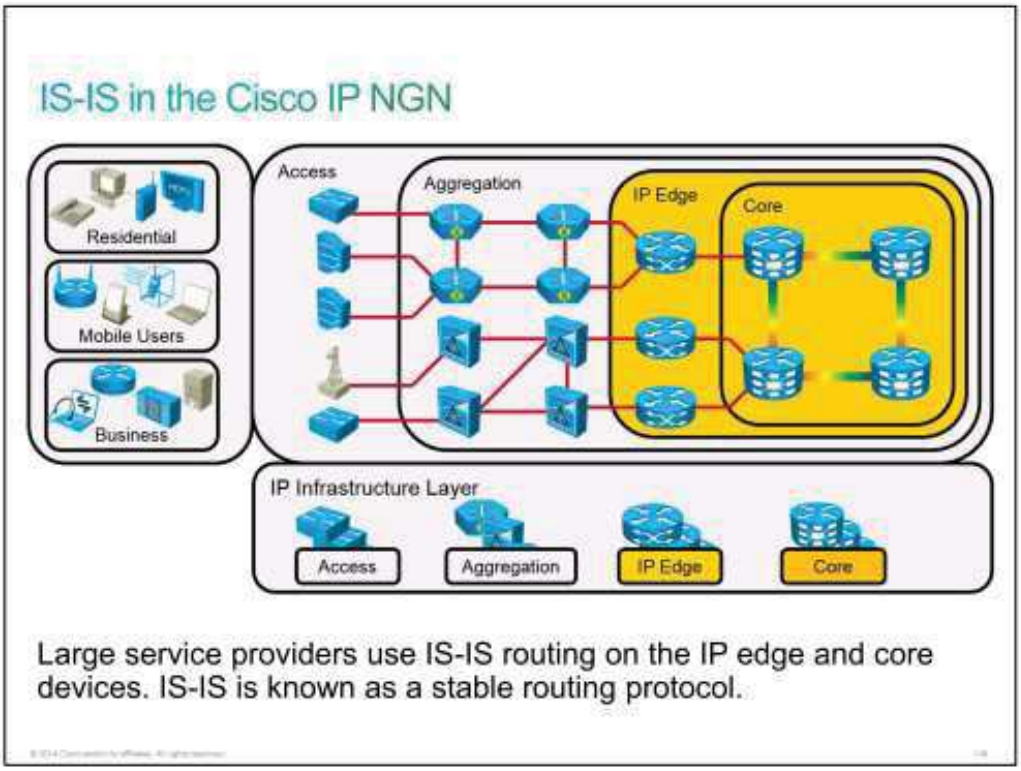

IS-IS development began before the development of OSPF. Large service providers often chose IS-IS as the IGP because of its unique capabilities for scalability, convergence, and stability. The U.S. government also required support for OSI protocols in the early days of the Internet. Although this requirement was later dropped, IS-IS met both constraints.

Later, businesses typically chose OSPF because it was a more widely supported native IP protocol. Today, finding information and expertise for IS-IS is more difficult than for OSPF. Nevertheless, some of the largest networks in the world persist in using IS-IS, which is a tribute to its capabilities.

https://t.me/learningnets

Integrated IS-IS Routing

Integrated IS-IS, or dual IS-IS, is an implementation of the IS-IS protocol for routing multiple network protocols, IP, and CLNS. Integrated IS-IS is specified in RFC 1195 and ISO 10589.

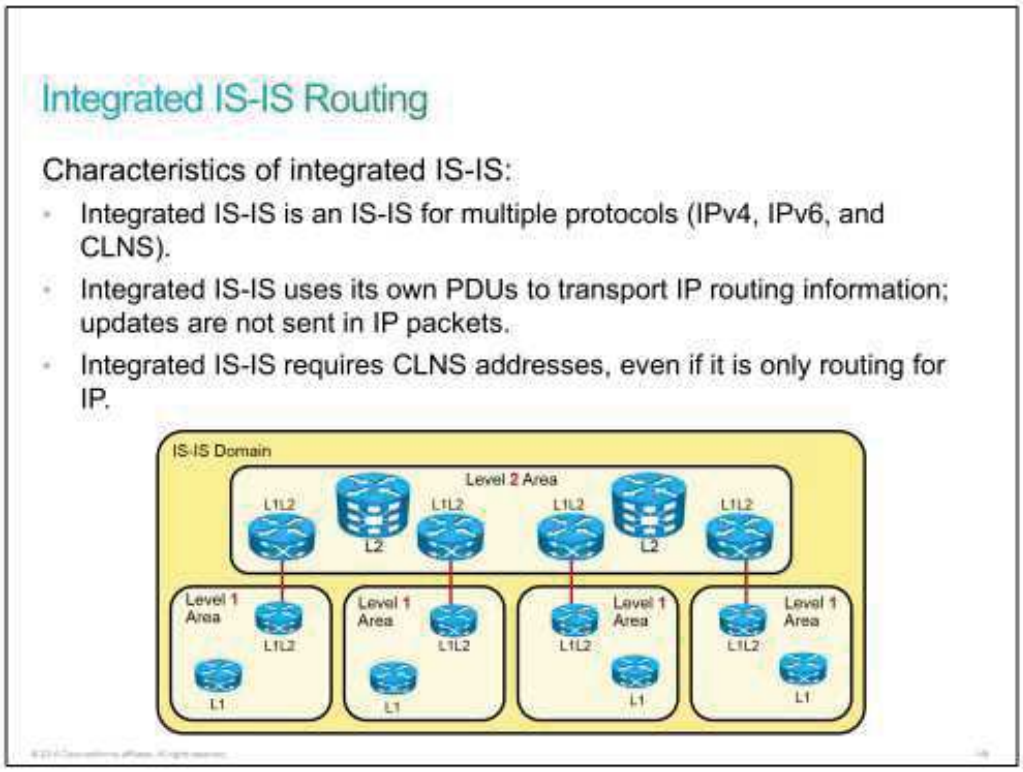

Integrated IS-IS tags CLNP routes with information about IP networks and subnets. As an alternative to OSPF, Integrated IS-IS combines ISO CLNS and IP routing in one protocol. Integrated IS-IS can be used for IP routing, CLNS routing, or for a combination of the two.

Integrated IS-IS uses its own PDUs, including IP reachability information, to transport information between routers. IS-IS information is not carried within a network-layer protocol, but instead it is carried directly within data-link layer frames.

This protocol-independence makes IS-IS easily extensible; there is also a version of Integrated IS-IS that supports IPv6. Because IS-IS uses CLNS addresses to identify the routers and to build the LSDB, an understanding of CLNS addresses is required to configure and troubleshoot IS-IS, even when it is used only for routing IP.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Integrated IS-IS Design Principles

Effective networks are well planned. The first and most important step in building a scalable network is developing a good addressing plan that allows for route summarization. Route summarization is possible only when you are using a hierarchical addressing structure.

Integrated IS-IS Design Principles

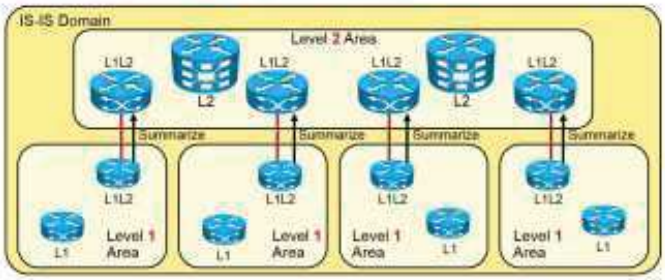

- IP and CLNP addresses must be planned.
- The two-level hierarchy is used for scalability:
  - Limits LSP flooding
  - Provides an opportunity for summarization

Integrated IS-IS Design Principles (Cont.)

Integrated IS-IS characteristics:

- Summarization:
  - Limits update traffic
  - Minimizes router memory and CPU usage
- The default, narrow metrics, are limited to a 6-bit interface and 10-bit path metric.
- Wide metrics allow a 24-bit interface and 32-bit path metric.
- The default metric is 10 on all interfaces.

https://t.me/learningnets

Effective address planning presents opportunities to group devices into areas. Using areas confines the scope of LSP propagation and saves bandwidth. Level 1-2 routers, on the border between a Level 1 area and the Level 2 backbone, are logical places to implement route summarization.

Route summarization saves memory because each intermediate system is no longer responsible for the LSPs of the entire routing domain. Route summarization also saves CPU usage because a smaller routing table is easier to maintain.

One issue with IS-IS is that older implementations, those that are using the narrow metrics, are limited to a maximum interface metric of 63 (6 bits) and a maximum total path metric of 1,023 (10 bits). There is little room to distinguish between paths. The Cisco IOS, IOS XE, and IOS XR Software support wide metrics that allow a 24-bit interface and 32-bit path metrics. The default, however, is still the narrow metrics.

IS-IS, as it is implemented on Cisco routers, does not automatically scale the interface metric. Instead, all IS-IS interfaces have a default metric of 10; this setting can be changed manually. If the default metric is not adjusted on each interface, the IS-IS metric becomes similar to the hop count metric that is used by the RIP.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Similarities Between IS-IS and OSPF

The differences between OSPF and IS-IS are small, but they do exist.

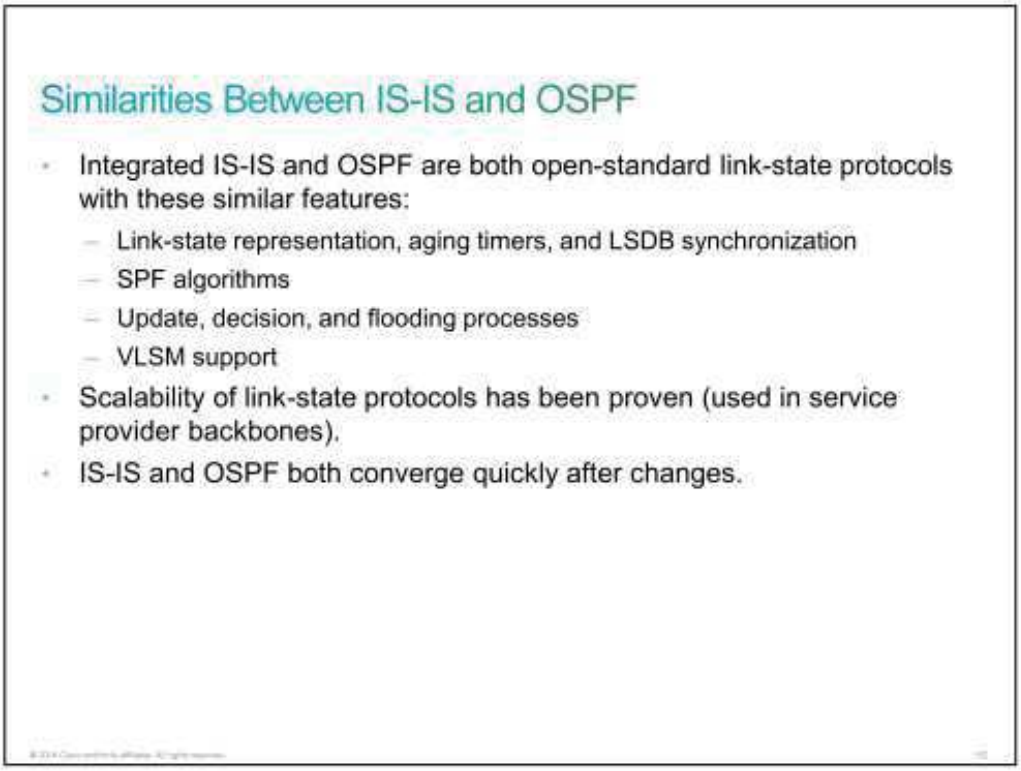

IS-IS and OSPF are more similar than dissimilar. Both routing protocols have the following characteristics:

- They are open-standard link-state routing protocols.
- They support VLSM.
- They use similar mechanisms, such as LSAs, link-state aging timers, and LSDB synchronization, to maintain the health of the LSDB.
- They use the SPF algorithm, with similar update, decision, and flooding processes.
- They are successful in the largest and most demanding deployments (service provider networks).
- They converge quickly after network changes.

https://t.me/learningnets

Integrated IS-IS vs. OSPF: Area Design

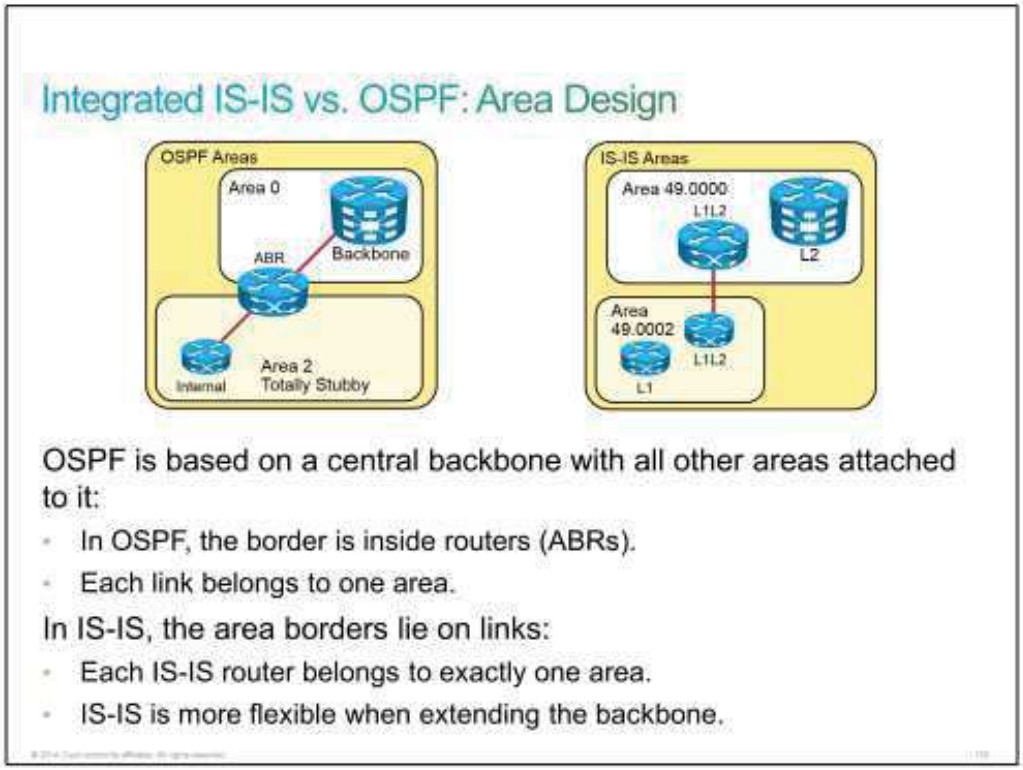

The figure summarizes the differences between OSPF and Integrated IS-IS.

With OSPF, network design is constrained by the fact that OSPF is based on a central backbone, Area 0, with all other areas being physically attached to Area 0. The border between areas is inside the ABRs; each link is in only one area. When you use this type of hierarchical model, a consistent IP addressing structure is necessary to summarize addresses into the backbone. Summarization also reduces the amount of information that is carried in the backbone and advertised across the network.

In comparison, IS-IS has a hierarchy of Level 1 and Level 2 or Level 1-2 routers, and the area borders lie on links. IS-IS permits a more flexible approach to extending the backbone. The backbone can be extended by simply adding more Level 2 and Level 1-2 routers, a less complex process than with OSPF.

https://t.me/learningnets

OSPF and IS-IS Characteristics

OSPF and IS-IS Characteristics

OSPF characteristics:

- Area border inside routers (ABRs)
- Each link is in only one area
- More complex to extend backbone
- Many small LSAs sent
- Runs on top of IP
- Requires IP addresses
- Default metric is scaled by interface bandwidth
- Not easy to extend
- Equipment, personnel, and information more readily available

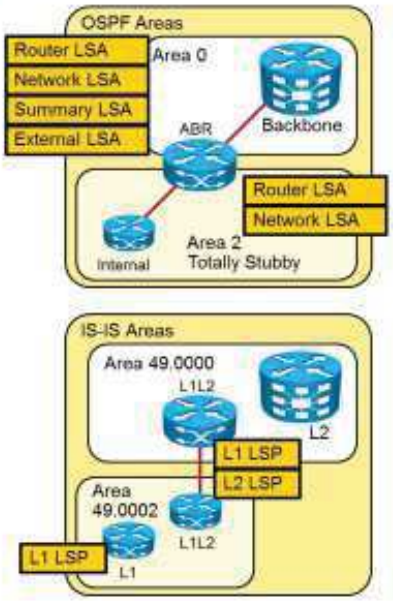

OSPF and IS-IS Characteristics (Cont.)

Integrated IS-IS characteristics:

- Area border on links
- Each router is in only one area
- Simple extension of backbone
- Fewer LSPs sent
- Runs on top of the data link layer
- Requires IP and CLNS addresses
- Default metric is 10 for all interfaces
- Easy to support new protocols with new TLV tuples
- Equipment, personnel, and information not as easily available

OSPF produces many small LSAs. IS-IS updates are grouped by the router and are sent as one LSP. Thus, as network complexity increases, the number of IS-IS updates is not an issue. Each update packet must be routed, though, and routing takes network resources, so more packets represent a larger impact on the network. Because IS-IS uses significantly fewer LSPs, more routers, at least 1000, can reside in a single area, making IS-IS more scalable than OSPF.

https://t.me/learningnets

OSPF runs over IP, whereas IS-IS runs through CLNS.

IS-IS is also more efficient than OSPF in the use of CPU resources and in the way it processes routing updates. Not only are there fewer LSPs to process (LSAs, in OSPF terminology), but also the mechanism by which IS-IS installs and withdraws prefixes is less intensive. IS-IS uses NET addresses that are already summarized.

Both OSPF and IS-IS are link-state protocols so they provide fast convergence. The convergence time depends on a number of factors, such as timers, number of nodes, and the type of router. Based on the default timers, IS-IS detects a failure faster than OSPF does; therefore, convergence occurs more rapidly. If there are many neighboring routers and adjacencies, the convergence time may also depend on the processing power of the router. IS-IS is less CPU-intensive than OSPF.

New features are not easily implemented in OSPF packets; they require the creation of a new LSA. The OSPF description schema is difficult to extend, because of compatibility issues and because it was developed exclusively for IPv4. IS-IS is easy to extend through the TLV mechanism. TLV strings, called tuples, encode all IS-IS updates. IS-IS can easily grow to cover IPv6, or any other protocol, because extending IS-IS consists of simply creating new TLVs.

A company may choose OSPF over IS-IS because OSPF is more optimized and because it was designed exclusively as an IP routing protocol. For example, OSPF defines different area types (normal, stub, and NSSA). The default OSPF metric is related to the interface bandwidth, while IS-IS defaults to a metric of 10 on all interfaces.

If a company does choose OSPF, it will require networking equipment that supports OSPF and network engineers that are familiar with OSPF theory and operation. Finding both equipment and personnel to support an OSPF infrastructure is relatively easy. Furthermore, OSPF documentation is much more readily available than documentation for IS-IS.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

IS-IS Addressing

Unlike IP addresses, CLNS addresses apply to entire nodes and not to interfaces. Because IS-IS was originally designed for CLNS, IS-IS requires CLNS addresses, even if the router is only used for routing IP. CLNS addresses that are used by routers are called NSAPs. One part of an NSAP address is the NSEL byte. When an NSAP is specified with an NSEL of 0, then the NSAP is called the NET.

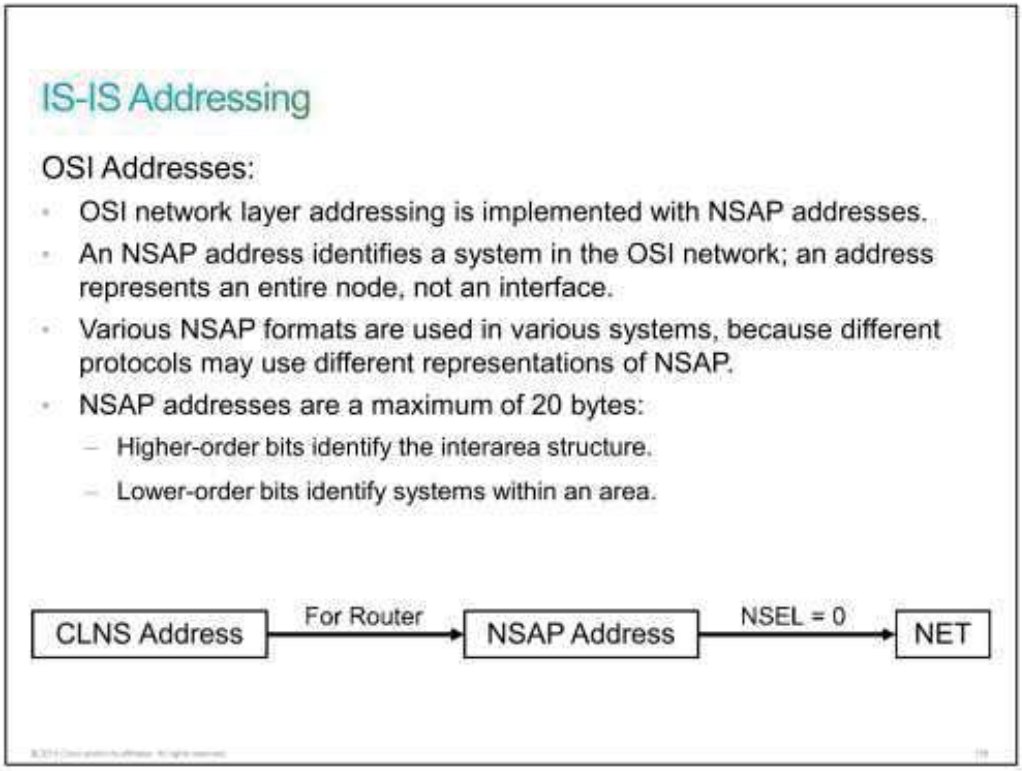

IS-IS LSPs use NSAP addresses to identify the router and build the topology table and the underlying IS-IS routing tree; therefore IS-IS requires NSAP addresses to function properly, even if it is used only for routing IP. NSAP addresses contain the following:

- OSI address of the device
- Link to the higher-layer process

The NSAP address is equivalent to the combination of the IP address and upper-layer protocol in an IP header.

NSAP addresses have a maximum size of 20 bytes. The higher-order bits identify the interarea structure, and the lower-order bits identify unique systems within an area.

There are various NSAP address formats.

https://t.me/learningnets

IS-IS Addressing (Cont.)

Integrated IS-IS NSAP address structure:

- The Cisco implementation of Integrated IS-IS distinguishes only the following three fields in the NSAP address:
  - **Area address:** Variable-length field (1 to 13 octets) composed of the higher-order NSAP octets, excluding system ID and NSEL (typically at least 1 byte, AFI set to 49 plus area ID).
  - **System ID:** Intermediate system identifier in an area; fixed length of six octets in Cisco IOS software.
  - **NSEL:** One octet NSAP selector, service identifier (0 for a router).
- Total length of NSAP is from 8 (minimum) to 20 octets (maximum).

IS-IS Addressing (Cont.)

AFI Value	Address Domain
39	ISO Data Country Code (DCC)
45	E.164
47	ISO 6523 International Code Designator (ICD)
49	Locally administered (private)

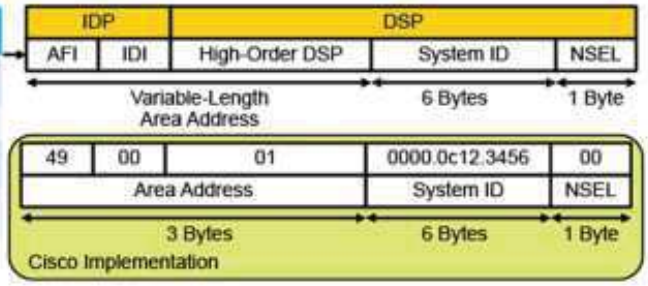

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

The Cisco implementation of Integrated IS-IS divides the NSAP address into three fields: the area address, the system ID, and the NSEL. Cisco routers routing CLNS use addressing that conforms to the ISO 10589 standard. ISO NSAP addresses consist of these elements:

- The AFI and the IDI make up the IDP of the NSAP address. The IDP corresponds roughly to an IP classful major network:
  - The AFI byte specifies the format of the address and the authority that is assigned to that address. Some valid values are shown in the figure.
  - Addresses starting with the AFI value of 49 are private addresses, analogous to RFC 1918 for IP addresses. IS-IS routes these addresses; however, this group of addresses should not be advertised to other CLNS networks because they are ad hoc addresses. Other companies that use a value of 49 may have created different numbering schemes that, when used together, could create confusion.
  - The IDI identifies a subdomain under the AFI. For instance, 47.0005 is assigned to civilian departments of the U.S. government, and 47.0006 is assigned to the U.S. Department of Defense.
- The DSP contributes to routing within an IS-IS routing domain.
  The DSP comprises the HO-DSP, the system ID, and the NSEL:
  - The HO-DSP subdivides the domain into areas. The HO-DSP is approximately the OSI equivalent of a subnet in IP.
  - The system ID identifies an individual OSI device. In OSI, a device has an address, just as it does in DECnet, whereas in IP, each interface has an address.
  - The NSEL identifies a process on the device and corresponds roughly to a port or socket in IP. The NSEL is not used in routing decisions.

The simplest NSAP format, used by most companies that are running IS-IS as their IGP, comprises the following:

- **Area address:** It must be at least 1 byte, separated into two parts:
  - The AFI, set to 49, which signifies that the AFI is locally administered and thus, individual addresses can be assigned by the company
  - The area identifier (ID), the octets of the area address following the AFI.
- **System ID:** Cisco routers that are compliant with the U.S. GOSIP version 2.0 standards require a 6-byte system ID.
- **NSEL:** NSEL must always be set to 0 for a router. NET is an NSAP address with an NSEL of 0.

Routers use the NET to identify themselves in the IS-IS PDUs. For example, you might assign 49.0001.0000.0c12.3456.00, which represents the following:

- AFI of 49
- Area ID of 0001
- System ID of 0000.0c12.3456, the MAC address of a LAN interface
- NSEL of 0

The area address is also referred to as the prefix.

https://t.me/learningnets

IS-IS Addressing (Cont.)

Identifying systems in IS-IS:

- All routers within an area must use the same area address.
- The area address is used in Level 2 routing.
- The system ID identifies the intermediate system.
- The system ID is used in Level 1 routing and must be unique within an area.
- The system ID must be unique within the Level 2 routers that form the routing domain.
- A domain-wide unique system ID is used; it may be a MAC or IP address.

IS-IS Addressing (Cont.)

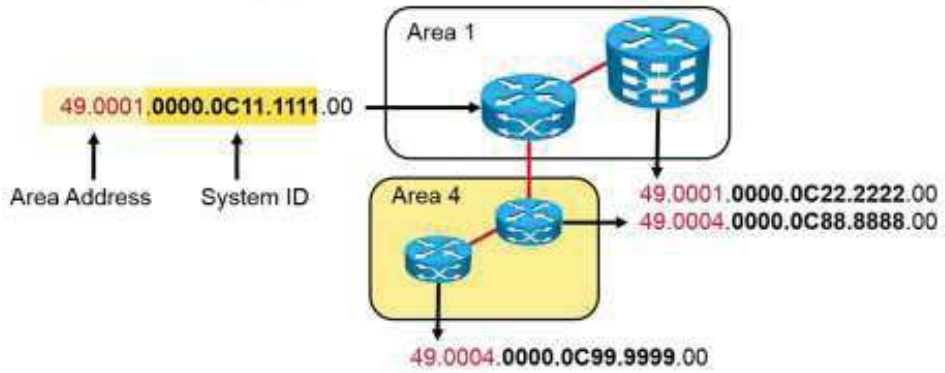

The area address uniquely identifies the routing area, and the system ID identifies each node.

The first part of an NSAP is the area address, and it is associated with the IS-IS routing process. Unlike OSPF, an IS-IS router can be a member of only one area. All routers in an area must use the same area address, which defines the area. The area address is used in Level 2 (interarea) routing.

The 6-byte NSAP system ID must be unique within an area. Using a MAC address from the router, or for Integrated IS-IS, to encode an IP address into the system ID is customary. All of the system IDs in a domain must be of equal length. Cisco enforces this OSI directive by fixing the length of the system ID at 6 bytes.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Level 1 intra-area routing is based on system IDs; therefore, each intermediate system (router) must have a unique system ID within the area.

All Level 2 intermediate systems eventually recognize all other intermediate systems in the Level 2 backbone; therefore, they must also have unique system IDs.

Thus, system IDs should remain unique across the domain. If the system IDs remain unique, there can never be a conflict at Level 1 or Level 2 (if, for example, a device moves into a different area).

https://t.me/learningnets

IS-IS Router Types

The path of connected Level 2 and Level 1-2 routers is called the backbone. All areas and the backbone must be contiguous.

IS-IS Router Types

Level 1, Level 2, and Level 1-2 Routers:

- Level 1 (similar to OSPF internal nonbackbone routers):
  - Intra-area routing enables routers to communicate.
  - Level 1 area is a collection of Level 1 and Level 1-2 routers.
  - Level 1 intermediate system keeps a copy of the Level 1 area LSDB.
- Level 1-2 (similar to OSPF ABR):
  - Intra-area and interarea routing are used.
  - Level 1-2 intermediate system keeps separate Level 1 and Level 2 LSDBs, and advertises the default route to Level 1 routers.
- Level 2 (similar to OSPF backbone routers):
  - Interarea routing is used.
  - The Level 2 (backbone) area is a contiguous set of Level 1-2 and Level 2 routers.
  - The Level 2 intermediate system keeps a copy of the Level 2 area LSDB.

IS-IS Router Types (Cont.)

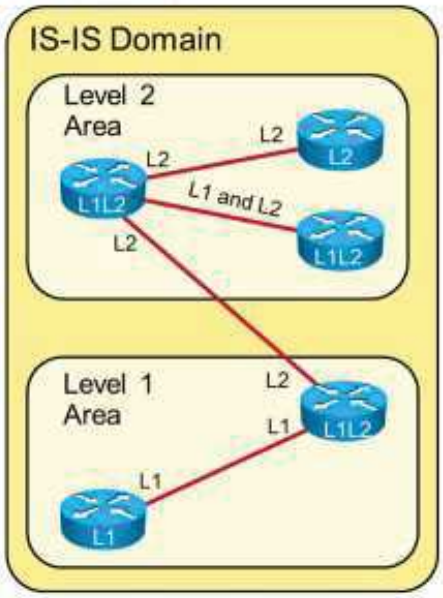

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Recall that IS-IS defines three types of routers:

- **Level 1:** Level 1 routers learn about paths within the areas they connect to (intra-area).
- **Level 2:** Level 2 routers learn about paths between areas (interarea).
- **Level 1-2:** Level 1-2 routers learn about paths both within and between areas. Level 1-2 routers are equivalent to ABRs in OSPF.

Area boundaries fall on the links. Each IS-IS router belongs to exactly one area. Neighboring routers learn whether they are in the same area or different areas, and they negotiate appropriate adjacencies, Level 1, Level 2, or both. Each router keeps a copy of the LSDBs for the levels that are its responsibility.

A Level 1-2 router automatically advertises to all Level 1 routers (within its area) that it is a potential exit point of the area. Level 1 routers default to the nearest attached Level 1-2 router to route outside its own area. You can consider IS-IS areas as "totally stub" areas by default, although optionally leaking Level 2 routes into Level 1is possible.

https://t.me/learningnets

IS-IS Routing Logic

IS-IS routing flows naturally from the OSI address plan, in which areas are identified and unique system IDs are given to each device.

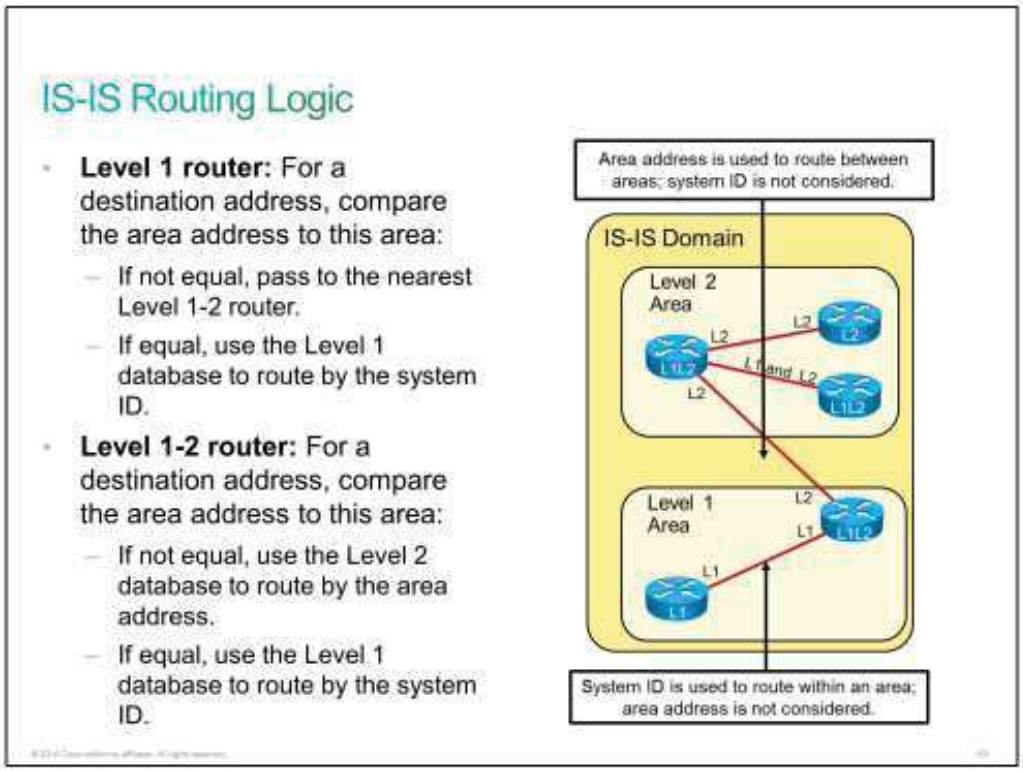

The area address portion of the NSAP address can range from 1 to 13 bytes in length, as specified by the ISO standard. Therefore, an NSAP for an IS-IS network can be as little as 8 bytes in length. The NSAP is usually longer, to permit some granularity in the allocation of areas. The area address prefix is common to all devices in an area and it is unique for each area. Routers are in the same area if they share the same area address.

Routing within an area involves collecting system IDs and adjacencies for all routers in an area, and using Dijkstra's algorithm to compute best paths between the devices. Level 1 routers are aware only of the local area topology. They pass the traffic that is destined to travel outside the area to the closest Level 1-2 router.

Routing between areas is based on the area address. Level 2 routers in different areas exchange area address information and use Dijkstra's algorithm to compute best paths between areas. They pass traffic between areas to the closest Level 1-2 router.

ESH and ISH packets are used for routers (intermediate systems) and end systems to detect each other. When a host is required to send a packet to another host, the packet goes to one of the routers on a network that is directly attached to the host. If the destination host is in the same area, the router searches for the destination system ID and forwards the packet appropriately along the best route.

If the destination address is a host in another area, the Level 1 router sends the packet to the nearest Level 1-2 router. Forwarding through Level 2 routers continues until the packet reaches a Level 2 (or Level 1-2) router in the destination area.

Within the destination area, routers forward the packet along the best path until the destination host is reached.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

Because each router makes its own best-path decisions at every hop along the way, there is a significant chance that paths will not be reciprocal. That is, return traffic can take a different path than the outgoing traffic. For this reason, knowing the traffic patterns within your network and tuning IS-IS for optimal path selection is important, if necessary.

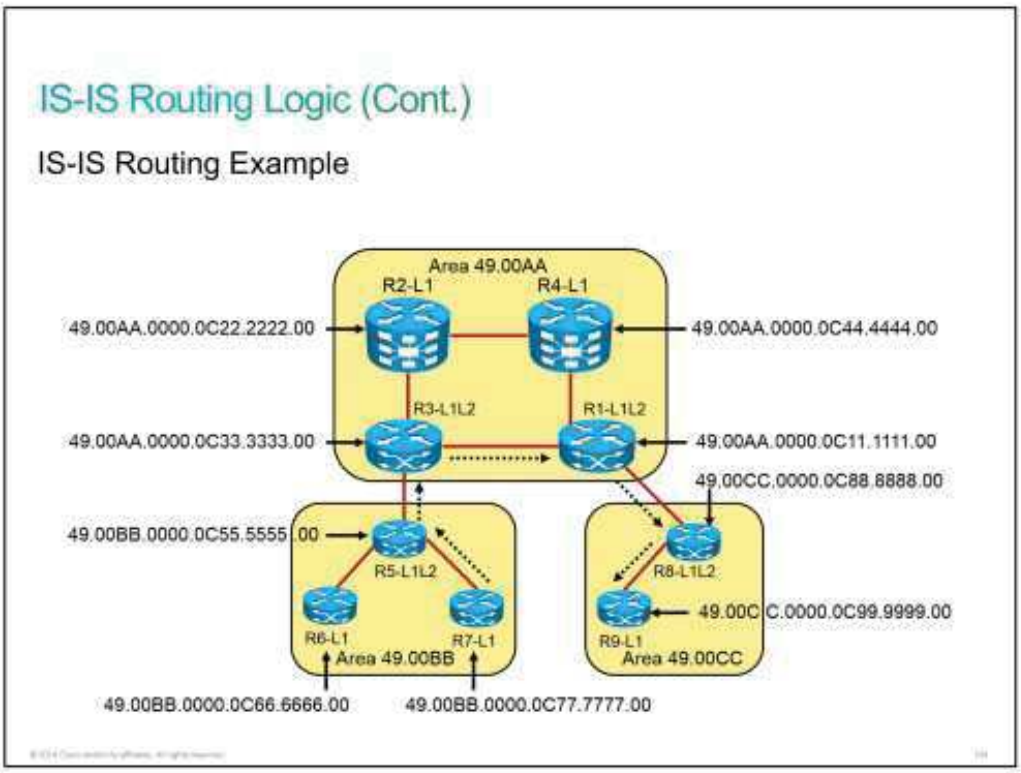

Consider traffic from router R7 to router R9:

1. R7 recognizes that the prefix (49.00CC) of R9 is not the same as the prefix (49.00BB) of R7. R7 therefore passes the traffic to the closest Level 1-2 router, R5. R7 uses its Level 1 topology database to find the best path to R5.
2. R5 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC, which is R3. R5 does not use the destination system ID in this decision.
3. R3 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC, which is R1. R3 does not use the destination system ID in this decision.
4. R1 uses its Level 2 topology database to pick the best next hop to reach the prefix 49.00CC, which is R8. R1 does not use the destination system ID in this decision.
5. R8 recognizes that the prefix (49.00CC) of R9 is the same as the prefix (49.00CC) of R8. R8, therefore, passes the traffic to R9, using its Level 1 topology database to find the best path.

https://t.me/learningnets

Asymmetric IS-IS Routing

Asymmetric routing (packets taking different paths in different directions) is not detrimental to the network; however, this type of routing can make troubleshooting difficult and it is sometimes a symptom of suboptimal design.

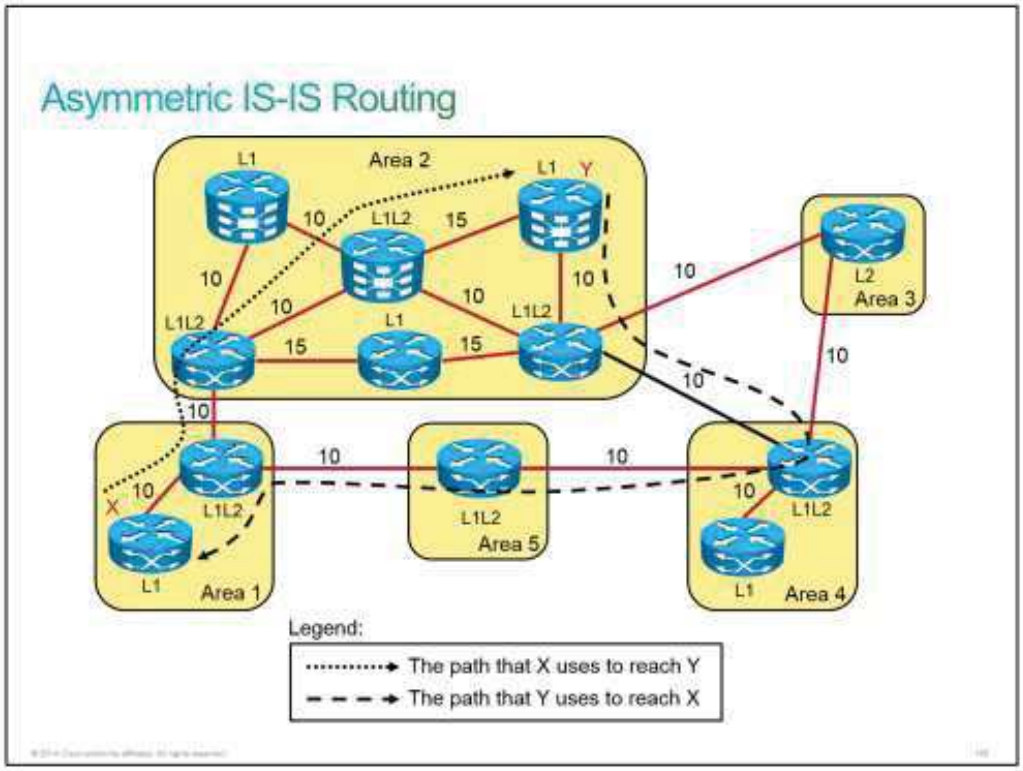

In the figure, Area 1 contains two routers:

- One router borders Area 2 and is a Level 1-2 intermediate system.
- The other router is contained within the area and is a Level 1 router only.

Area 2, however, has many routers:

- A selection of routers is specified as Level 1. The routers route either internally to that area or to the exit points (the Level 1-2 routers).
- The three Level 1-2 routers form a chain across Area 2, linking to the neighbor Areas 1, 3, and 4. Although the middle router of the three Level 1-2 routers does not link directly to another area, the middle router must support Level 2 routing to ensure that the backbone is contiguous. If the middle router fails, the other Level 1-only routers cannot perform the Level 2 function (despite having a physical path across Area 2), and the backbone is broken.

Area 3 contains one router that borders Areas 2 and 4, yet it has no intra-area neighbors and is performing Level 2 functions only. If you add another router to Area 3, the border router reverts to Level 1-2 functions.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

In the figure, symmetric routing does not occur because Level 2 details are hidden from Level 1 routers that recognize only a default route to the nearest Level 1-2 router. Traffic from router X to router Y flows from router X to its closest Level 1-2 router. The Level 1-2 router then forwards the traffic along the shortest path to the destination area (Area 2). When the traffic flows into Area 2, the traffic is routed along the shortest intra-area path to router Y.

Router Y routes return packets to router X via its nearest Level 1-2 router. The Level 1-2 router recognizes the best route to Area 1 via Area 4, based on the lowest-cost Level 2 path.

Because Level 1 and Level 2 computations are separate, the path taken from router Y back to router X is not necessarily the least-cost path from router Y to router X.

https://t.me/learningnets

Symmetric IS-IS Routing

Route leaking is a feature that allows selected Level 2 routes to leak in a controlled manner to Level 1 routers, which helps to avoid asymmetric routing.

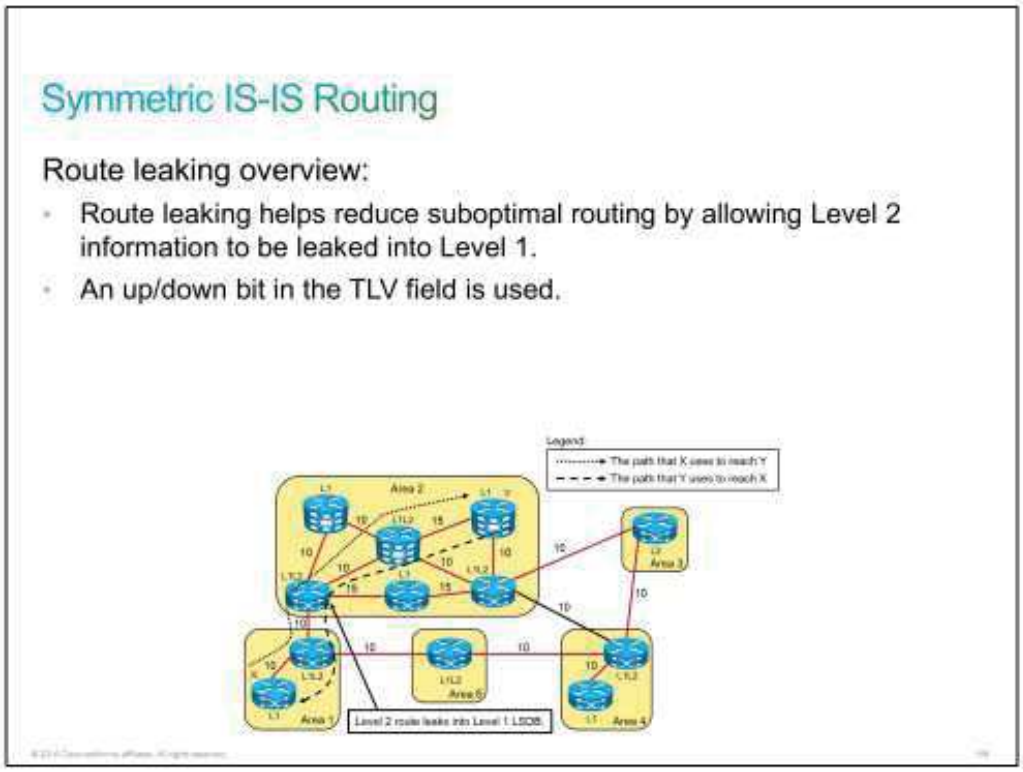

Route leaking helps reduce suboptimal routing by providing a mechanism for leaking, or redistributing, Level 2 information into Level 1 areas. With more detail about interarea routes, a Level 1 router is able to make a better choice about which Level 1-2 router should receive the packet. As shown in the figure, the route to reach router X is leaked from Area 1 into the Level 1 routers in Area 2. Therefore, router Y knows that the best path to reach router X is via the Level 1/Level 2 router that connects directly to Area 1.

Leaked routes are referred to as interarea routes in the routing table and the IS-IS database. When you are viewing the routing table, all the leaked routes are marked with an "ia" designation.

Route leaking is defined in RFC 2966, Domain-wide Prefix Distribution with Two-Level IS-IS, for use with the narrow metric TLV types 128 and 130. The IETF has also defined route leaking for use with the wide metric (using TLV type 135).

To implement route leaking, an up/down bit in the TLV is used to indicate whether or not the route that is identified in the TLV has been leaked. If the up/down bit is set to 0, the route originated within that Level 1 area.

If the up/down bit is set to 1, the route has been redistributed into the area from Level 2. The up/down bit is used to prevent routing loops; a Level 1-2 router does not readvertise, into Level 2, any Level 1 routes that have the up/down bit set to 1.

Route leaking should be planned and deployed carefully to avoid a situation in which any topology change in one area makes it necessary to recompute many routes in all other areas.

https://t.me/learningnets

IS-IS Packets

The OSI stack defines a unit of data as a PDU. OSI recognizes a frame as a data-link PDU, and a packet as a network PDU.

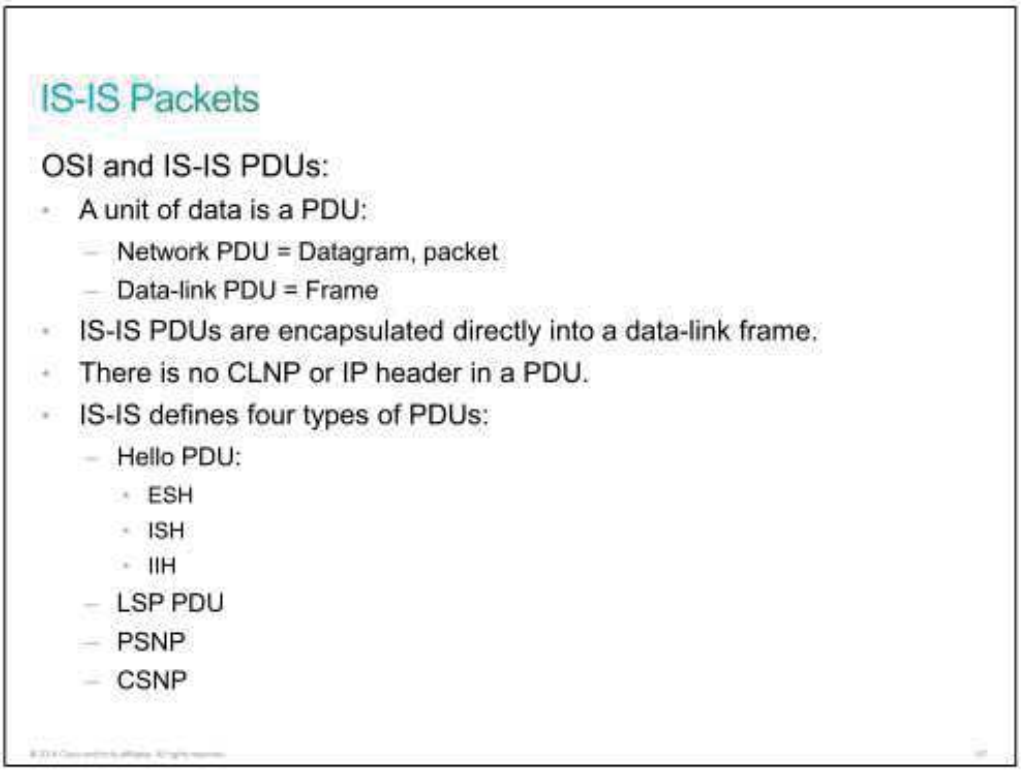

The first eight octets of all IS-IS PDUs are header fields that are common to all PDU types. The TLV information is stored at the very end of the PDU. Different types of PDUs have a set of currently defined TLV codes. Any TLV codes that are not recognized by a router should be ignored and passed through unchanged.

IS-IS PDUs are encapsulated directly into an OSI data-link frame. IS-IS defines four types of PDUs:

- **Hello PDU:** This type of PDU is used to establish and maintain adjacencies:
  - ESH: Announces the presence of an end system. ESHs are sent by all end systems. Intermediate systems (routers) listen for these hellos to discover the end systems.
  - ISH: Announces the presence of an intermediate system (router). ISHs are sent by all intermediate systems. End systems listen for these hellos to discover the intermediate systems.
  - IIH: Enables the intermediate systems to detect IS-IS neighbors and form adjacencies
- **LSP PDU:** Used to distribute link-state information
- **PSNP:** Used to acknowledge and request missing pieces of link-state information
- **CSNP:** Used to describe the complete list of LSPs in the LSDB of a router

https://t.me/learningnets

IS-IS Packets (Cont.)

Link-state packet:

- A router describes itself with an LSP.
- LSPs are sequenced to prevent duplication:
  - LSPs assist with synchronization.
  - Sequence numbers begin at 1.
  - Sequence numbers are increased to indicate the newest LSP.
- LSPs in LSDB have a remaining lifetime:
  - Allow synchronization, decreasing timer
- Each set includes a TLV.

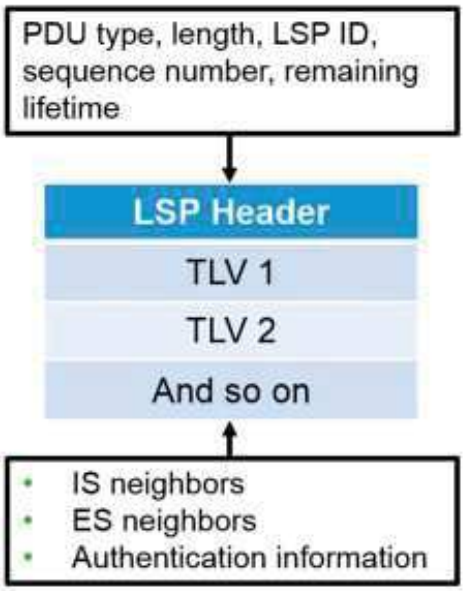

IS-IS Packets (Cont.)

TLV	Type Code	Length Field	Value Variable Length
Area address	1	Area ID length+1	Areas
Intermediate system neighbors	2	Neighbor count+1	IS neighbors
IP internal reachability	128	Number of connected prefixes	Connected IP prefixes: 4-byte metric, 4-byte prefix, 4 byte mask
IP external reachability	130	Number of redistributed prefixes	Redistributed IP prefixes: 4-byte metric, 4-byte prefix, 4 byte mask

In IS-IS, the characteristics of a router are defined by an LSP. The LSP of a router contains an LSP header and TLV fields:

- An LSP header includes these elements:
  - The PDU type and length
  - The LSP ID

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

  - The LSP sequence number, used to identify duplicate LSPs and to ensure that the latest LSP information is stored in the topology table.
  - The remaining lifetime for the LSP, which is used to age out LSPs
- TLV: Some examples of TLV fields include the following:
  - Type Code 1 = area addresses
  - Type Codes 2 and 6 = intermediate system neighbors
  - Type Code 3 = end system neighbors
  - Type Code 10 = authentication information
  - Type Code 128 = IP internal reachability information
  - Type Code 129 = protocols supported
  - Type Code 130 = IP external reachability information
  - Type Code 132 = IP interface addresses

Note	Refer to ISO 10589 and RFC 1195 for more information on the TLVs.

LSPs are given sequence numbers that provide information to the receiving routers:

- Ensures that the latest LSPs are used in their route calculations.
- Avoids entering duplicate LSPs in the topology tables.

If a router reloads, the sequence number is set to 1. The router then receives its previous LSPs from its neighbors. These LSPs have the last valid sequence number before the router reloaded. The router records this number and reissues its own LSPs with the next higher sequence number.

Each LSP has a remaining lifetime that is used by the LSP aging process to ensure the removal of outdated and invalid LSPs from the topology table after a suitable time. This process is known as the count-to-zero operation; twelve hundred seconds is the default start value.

Each LSP includes specific information about networks and stations that are attached to a router. This information is found in multiple TLV fields that follow the common header of the LSP. The TLV structure is a flexible way to add data to the LSP, and an easy mechanism for adding new data fields that may be required in the future.

https://t.me/learningnets

Integrated IS-IS for IPv6

Cisco has added multitopology support to IS-IS to increase flexibility in IS-IS deployment within a dual-stack environment. IS-IS can be deployed using two SPF instances, one for IPv4 and one for IPv6.

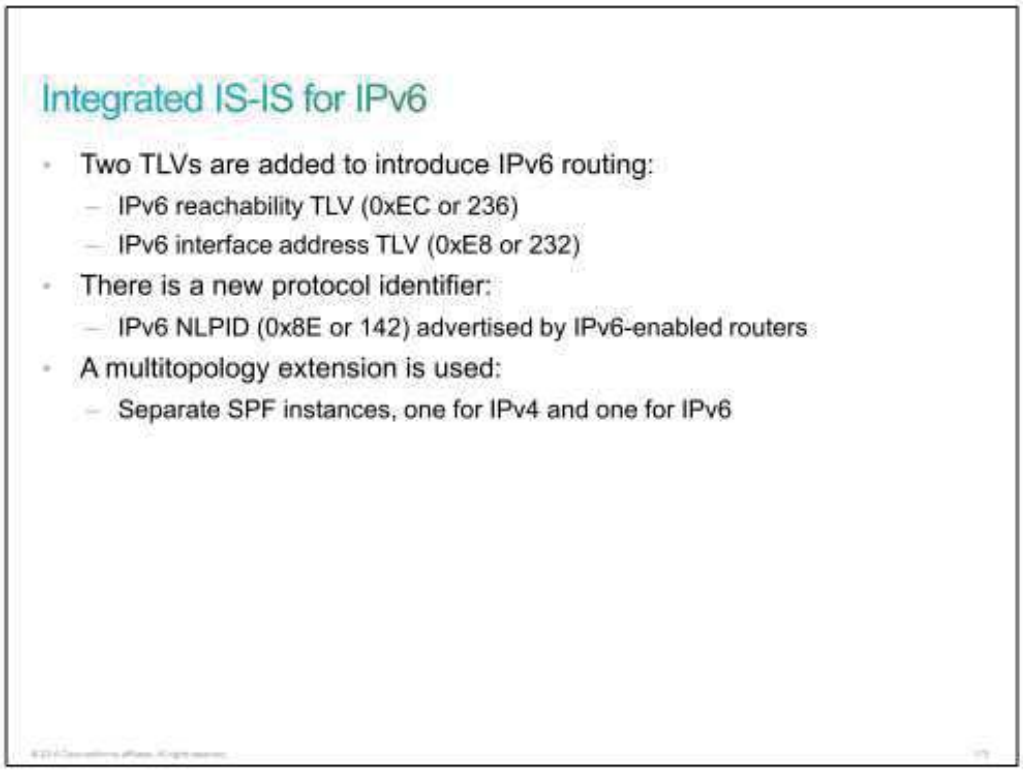

Two TLVs are added in IS-IS for IPv6 support. These two TLVs are used to describe IPv6 reachability and IPv6 interface addresses:

- IPv6 reachability TLV (0xEC or 236):
  - Describes network reachability (routing prefix, metric, options)
  - Equivalent to IPv4 internal and external reachability TLVs (type code 128 and 130)
- IPv6 interface address TLV (0xE8 or 232):
  - Equivalent to IPv4 interface address TLV (type code 132)
  - For hello PDUs, which must contain the link-local address
  - For LSPs, which must only contain the non-link-local address

The protocols-supported TLV (type code 129) lists the supported NLPIDs. All IPv6-enabled IS-IS routers advertise an NLPID value of 0x8E (142). The NLPID of IPv4 is 0xCC (204).

Multitopology IS-IS provides some flexibility when you are transitioning to IPv6. A separate topology is kept for both IPv4 and IPv6 networks; because some links may not be able to carry IPv6, IS-IS specifically keeps track of those links, minimizing the possibility for the traffic to be "black-holed."

Single topology IS-IS, where there is one SPF instance for both IPv4 and IPv6, also remains a possibility that is even easier to administer, but the network must be homogenous. The same links must carry IPv4 and IPv6 simultaneously.

https://t.me/learningnets

IS-IS Network Types

IS-IS network topologies can be divided into two general types: point-to-point and broadcast.

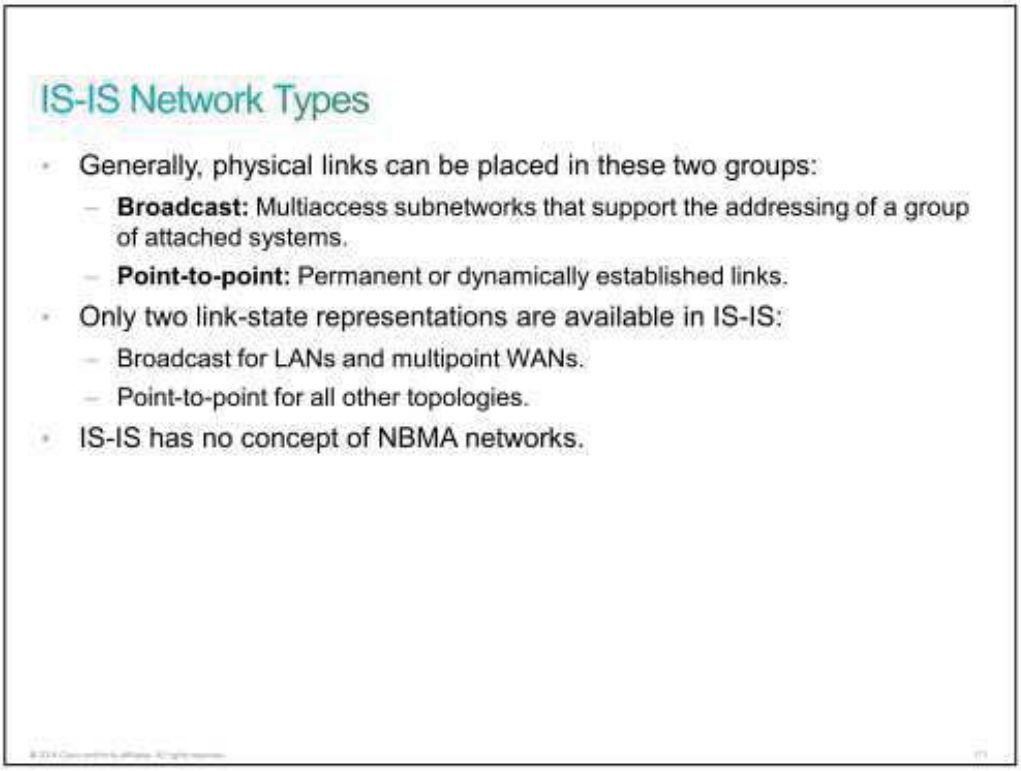

IS-IS supports two media representations for its link states:

- Broadcast for LANs and multipoint WAN links
- Point-to-point for all other media

Note	IS-IS has no concept of NBMA networks. Using point-to-point links, such as point-to-point subinterfaces, over NBMA networks is recommended.

https://t.me/learningnets

IS-IS Operations

Broadcast networks are LAN interfaces or multipoint WAN interfaces. Separate IS-IS adjacencies are established for Level 1 and Level 2. If two neighboring routers in the same area run both Level 1 and Level 2, they establish two adjacencies, one for each level. The router stores the Level 1 and Level 2 adjacencies in separate Level 1 and Level 2 adjacency tables.

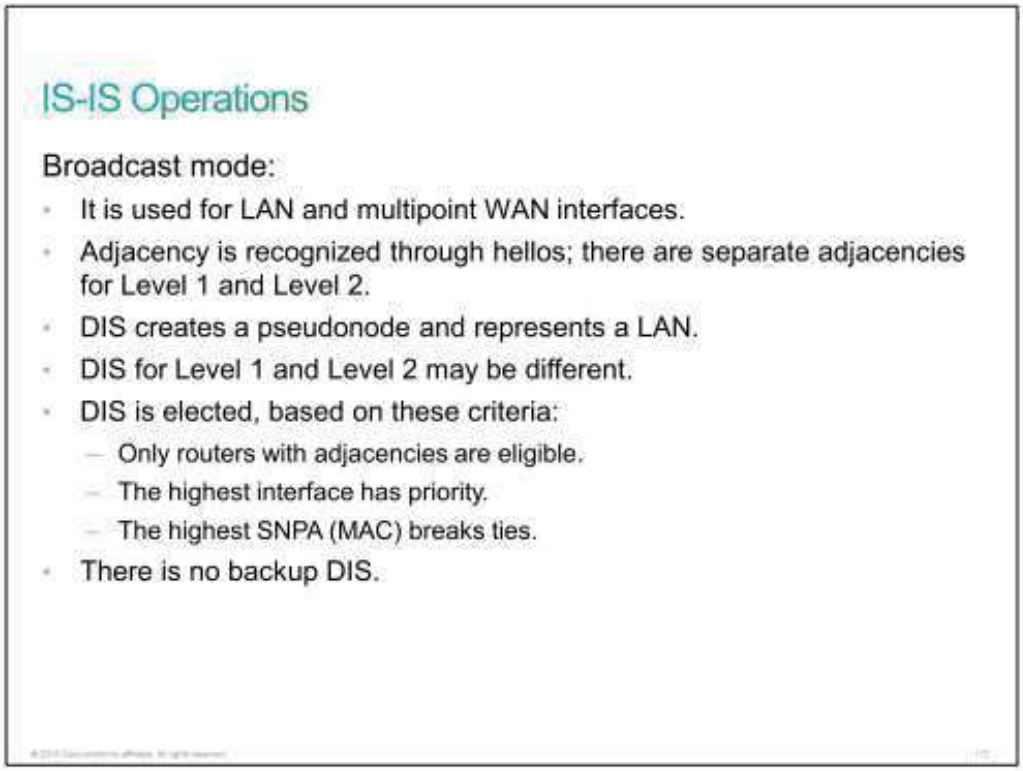

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

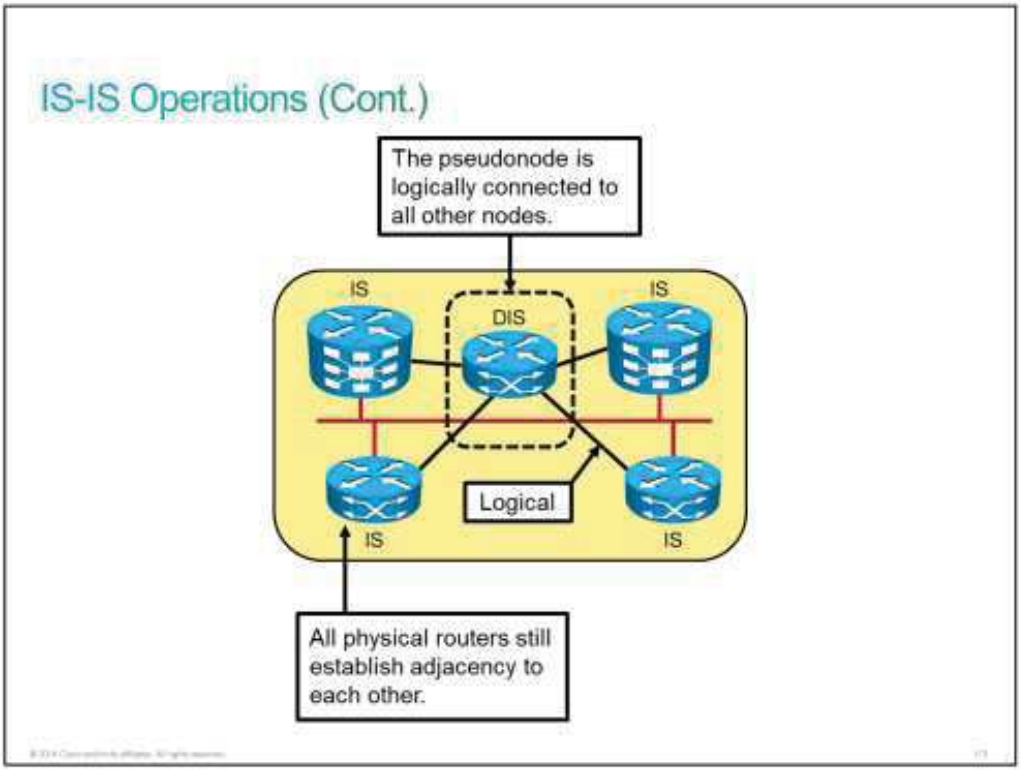

On LANs, routers establish the two adjacencies with specific Layer 1 and Layer 2 IIH PDUs. Routers on a LAN establish adjacencies with all other routers on the LAN.

IIH PDUs announce the area address. Separate IIH packets announce the Level 1 and Level 2 neighbors. Adjacencies are formed and are based on the area address that is communicated in the incoming IIH and the type of router (Level 1 or Level 2). Level 1 routers accept Level 1 IIH PDUs from their own area and establish adjacencies with other routers in their own area. Level 2 routers accept only Level 2 IIH PDUs and establish only Level 2 adjacencies.

Dijkstra's algorithm requires a virtual router (a pseudonode), represented by the DIS, to build a directed graph for broadcast media. The DIS is the router that creates the pseudonode and acts on behalf of the pseudonode. Two major tasks that are performed by the DIS include creating and updating the pseudonode LSP and flooding LSPs over the LAN. The criteria for DIS selection are as follows:

- The highest priority (the priority value is configurable)
- The highest SNPA; on LANs, the SNPA is the MAC address. The SNPA for a WAN interface is the virtual circuit identifier.

An example would be the DLCI on a Frame Relay connection. If the WAN interface is using HDLC encapsulation, the SNPA is simply HDLC.

Cisco router interfaces have a default Level 1 and Level 2 priority of 64. You can configure the priority from 0 to 127. The Level 1 DIS and the Level 2 DIS on a LAN may or may not be the same router because an interface can have different Level 1 and Level 2 priorities.

A selected router is not guaranteed to remain the DIS. Any adjacent intermediate system with a higher priority automatically takes over the DIS role. This behavior is called pre-emptive. Because the IS-IS LSDB is synchronized frequently on a LAN, giving priority to another intermediate system over the DIS is not a significant issue. IS-IS does not use a backup DIS, and routers on a LAN establish adjacencies with both, the DIS and with all other routers.

https://t.me/learningnets

IS-IS Operations (Cont.)

Level 1 and Level 2 LSPs and IIHs:

- The two-level nature of IS-IS requires separate types of LSPs: Level 1 and Level 2 LSPs.
- DIS is a representative of a LAN:
  - DIS sends pseudo-Level 1 and pseudo-Level 2 LSPs for a LAN.
  - There is a separate DIS for Level 1 and Level 2.

	Broadcast	Point-to-Point
Usage	LAN, full-mesh WAN	PPP, HDLC, partial-mesh WAN
Hello timer	3.3 sec for DIS, else 10 sec	10 sec
Adjacencies	n*(n-1)/2	n-1
Uses DIS	Yes	No
LSP and IIH	Sent as multicast	Sent as unicast
IIH type	Level 1 IIH, Level 2 IIH	Point-to-point IIH

Level 1 and Level 2 LSP: IS-IS uses a two-level area hierarchy. The link-state information for these two levels is distributed separately, which results in Level 1 LSPs and Level 2 LSPs. Each intermediate system originates its own LSPs (one for Level 1 and one for Level 2).

On a LAN, one router (the DIS) sends out LSP information on behalf of the LAN. The DIS represents a pseudonode. The DIS sends out the separate Level 1 or Level 2 LSPs for the pseudonode. The Level 1 DIS and the Level 2 DIS on a LAN may or may not be the same router because an interface can have different Level 1 and Level 2 priorities.

LSPs on point-to-point links are sent as unicast, whereas on broadcast media (LANs) LSPs are sent as multicast.

Level 1 and Level 2 IIH: IIHs are used to establish and maintain neighbor adjacency between intermediate systems. The default hello interval is every 10 seconds; however, the hello interval timer is adjustable.

On a LAN, separate Level 1 and Level 2 IIHs are sent periodically as multicasts to a multicast MAC address. Level 1 announcements are sent to the AllL1IS multicast MAC address 0180.C200.0014, and Level 2 announcements are sent to the AllL2IS multicast MAC address 0180.C200.0015.

The default hello interval for the DIS is three times faster (that is, three times smaller) than the interval for the other routers so that DIS failures can be quickly detected.

A neighbor is declared dead if hellos are not received within the hold time. The hold time is calculated as the product of the hello multiplier and hello time. The default hello time is 10 seconds, and the default multiplier is three; therefore, the default hold time is 30 seconds.

Unlike LAN interfaces with separate Level 1 and Level 2 IIHs, point-to-point links have a common point-to-point IIH format that specifies whether the hello relates to Level 1 or Level 2 or both. Point-to-point hellos are sent to the unicast address of the connected router.

The table summarizes the differences between broadcast and point-to-point links.

https://t.me/learningnets

IS-IS LSP Flooding

An IS-IS update process is responsible for flooding the LSPs throughout the IS-IS domain. An LSP is typically flooded to all adjacent neighbors except the neighbor from which it was received. Level 1 LSPs are flooded within their local areas. Level 2 LSPs are flooded throughout the backbone.

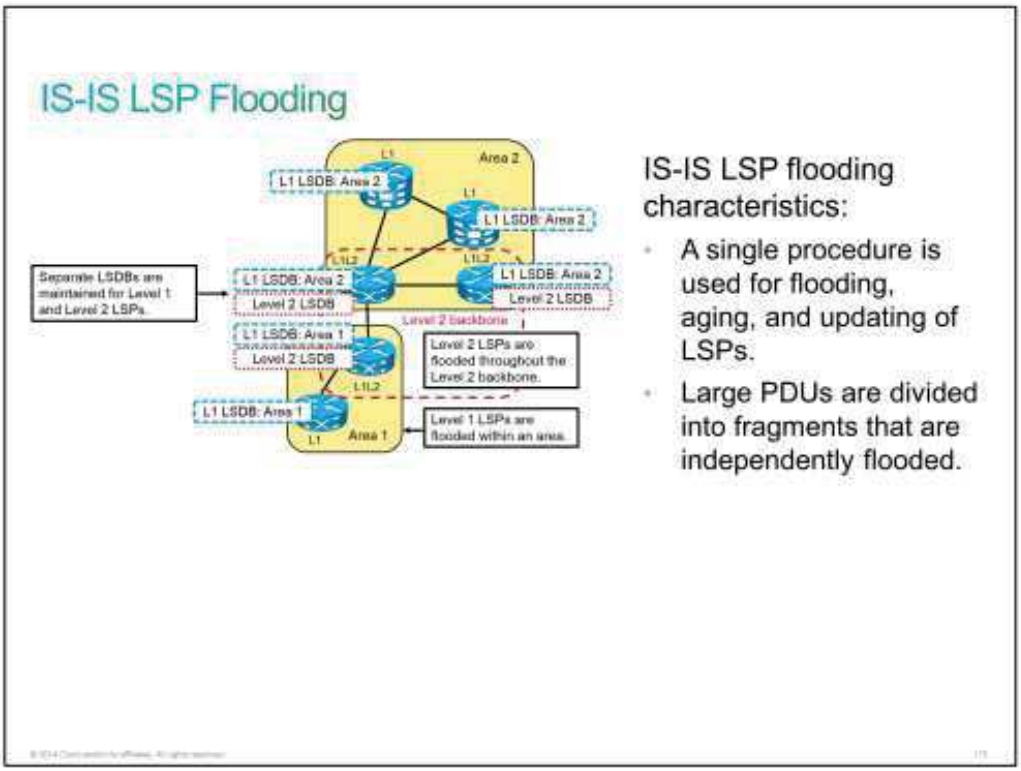

Each intermediate system originates its own LSP (one for Level 1 and one for Level 2). These LSPs are identified by the system ID of the originator and an LSP fragment number starting at 0. If an LSP exceeds the MTU, it is fragmented into several LSPs, numbered 1, 2, 3, and so on.

IS-IS maintains the Level 1 and Level 2 LSPs in separate LSDBs.

When an intermediate system receives an LSP, it examines the checksum and discards any invalid LSPs, flooding them with an expired lifetime age. If the LSP is valid and newer than what is currently in the LSDB, it is retained, acknowledged, and given a lifetime of 1200 seconds.

The age is decremented every second until it reaches 0, at which point the LSP is considered to have expired. When the LSP has expired, it is kept for an additional 60 seconds before it is flooded as an expired LSP.

https://t.me/learningnets

IS-IS LSDB Synchronization

SNPs are used to acknowledge the receipt of LSPs and to maintain LSDB synchronization.

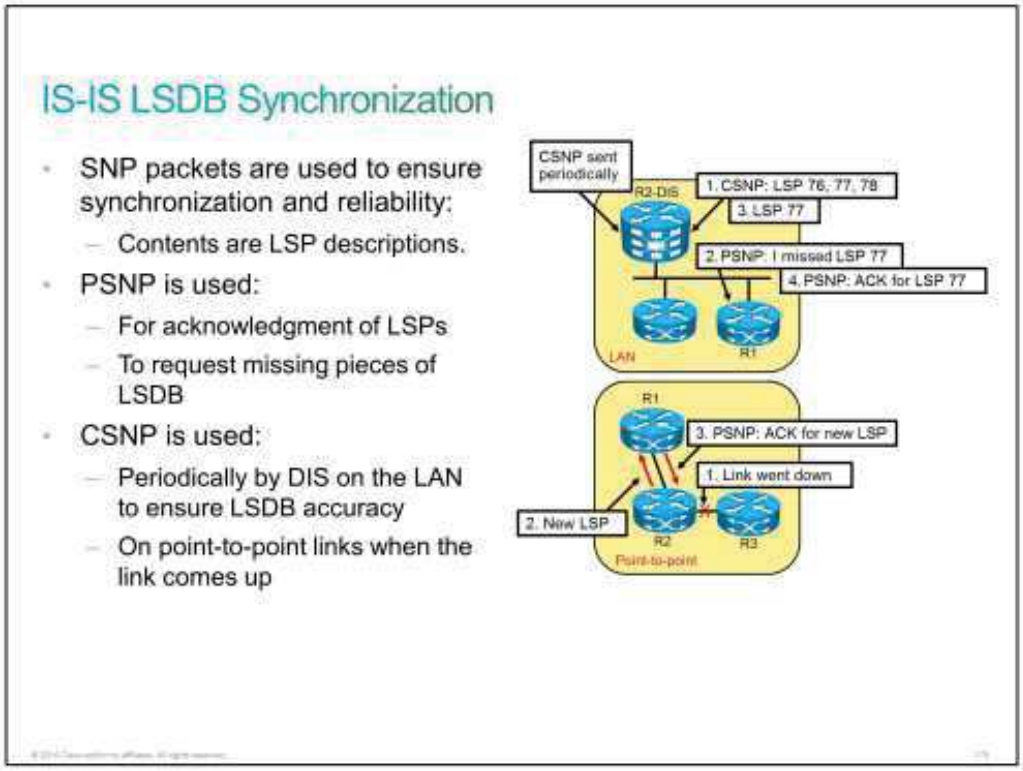

There are two types of SNPs: CSNP and PSNP. The use of SNPs differs between point-to-point and broadcast media. CSNPs and PSNPs share the same format; that is, each carries summarized LSP information. The main difference is that CSNPs contain summaries of all LSPs in the LSDB, while PSNPs contain only a subset of LSP entries.

Separate CSNPs and PSNPs are used for Level 1 and Level 2 adjacencies. Adjacent IS-IS routers exchange CSNPs to compare their LSDB. In broadcast subnetworks, only the DIS transmits CSNPs. All adjacent neighbors compare the LSP summaries received in the CSNP with the contents of their local LSDBs to determine if their LSDBs are synchronized (in other words, if they have the same copies of LSPs as other routers for the appropriate levels and area of routing).

CSNPs are periodically multicast (every 10 seconds) by the DIS on a LAN to ensure LSDB accuracy. If there are too many LSPs to include in one CSNP, the LSPs are sent in ranges. The CSNP header indicates the starting and ending LSP ID in the range. If all LSPs fit in the CSNP, the range is set to default values. In the example, router R1 compares this list of LSPs with its topology table and realizes that it is missing one LSP. Therefore, it sends a PSNP to the DIS (router R2) to request the missing LSP. The DIS reissues only that missing LSP (LSP 77), and router R1 acknowledges it with a PSNP.

Adjacent IS-IS routers use PSNPs to acknowledge the receipt of LSPs and to request transmission of missing or newer LSPs. On point-to-point networks, CSNPs are sent only once when the link comes up to synchronize the LSDBs. After that, LSPs are sent to describe topology changes, and they are acknowledged with a PSNP. The example in the figure shows what happens on a point-to-point link when a link failure is detected. The sequence is as follows:

- A link fails.
- Router R2 notices this failure and issues a new LSP noting the change.

https://t.me/learningnets

- Router R1 receives the LSP, stores it in its topology table, and sends a PSNP back to R2 to acknowledge receipt of the LSP.

https://t.me/learningnets

IS-IS Adjacencies

IIH PDUs announce the area address. On LANs, separate IIH packets announce the Level 1 and Level 2 neighbors.

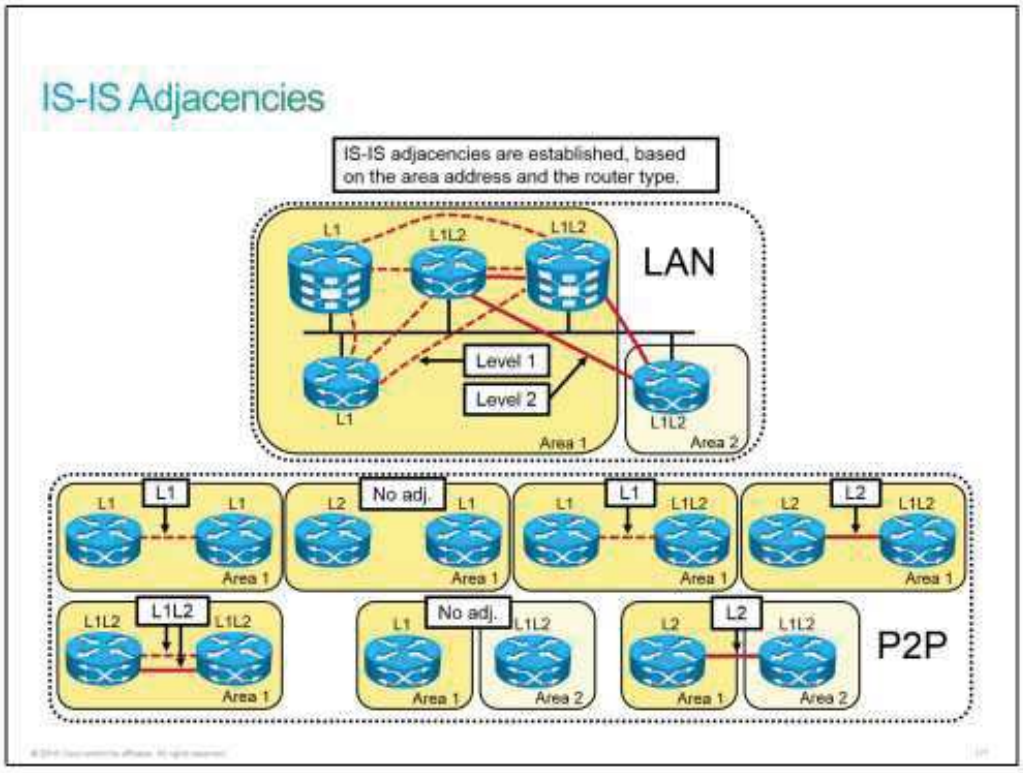

For example, when a LAN has routers from two areas attached, the following processes apply:

- The routers from one area accept Level 1 IIH PDUs only from their own area and therefore establish Level 1 adjacencies only with their own area Level 1 routers.
- The routers from a second area, similarly accept Level 1 IIH PDUs only from their own area.
- The Level 2 routers (or the Level 2 process within any Level 1-2 router) accept only Level 2 IIH PDUs and establish only Level 2 adjacencies.

On point-to-point links, the IIH PDUs are common to both levels but announce the level type and the area address in the hellos as follows:

- Level 1 routers in the same area exchange IIH PDUs that specify Level 1 and establish a Level 1 adjacency.
- Level 2 routers exchange IIH PDUs that specify Level 2 and establish a Level 2 adjacency.
- Two Level 1-2 routers in the same area establish both Level 1 and Level 2 adjacencies and maintain these with a common IIH PDU that specifies the Level 1 and Level 2 information.
- Two Level 1 routers that are physically connected, but that are not in the same area, can exchange IIHs, but they do not establish adjacency because the area addresses do not match.

https://t.me/learningnets

IS-IS Single Topology Restrictions

The original design for Integrated IS-IS defines a single SPF for all routed protocols, which adds the assumption that all interfaces included in the routing decisions run all routed protocols.

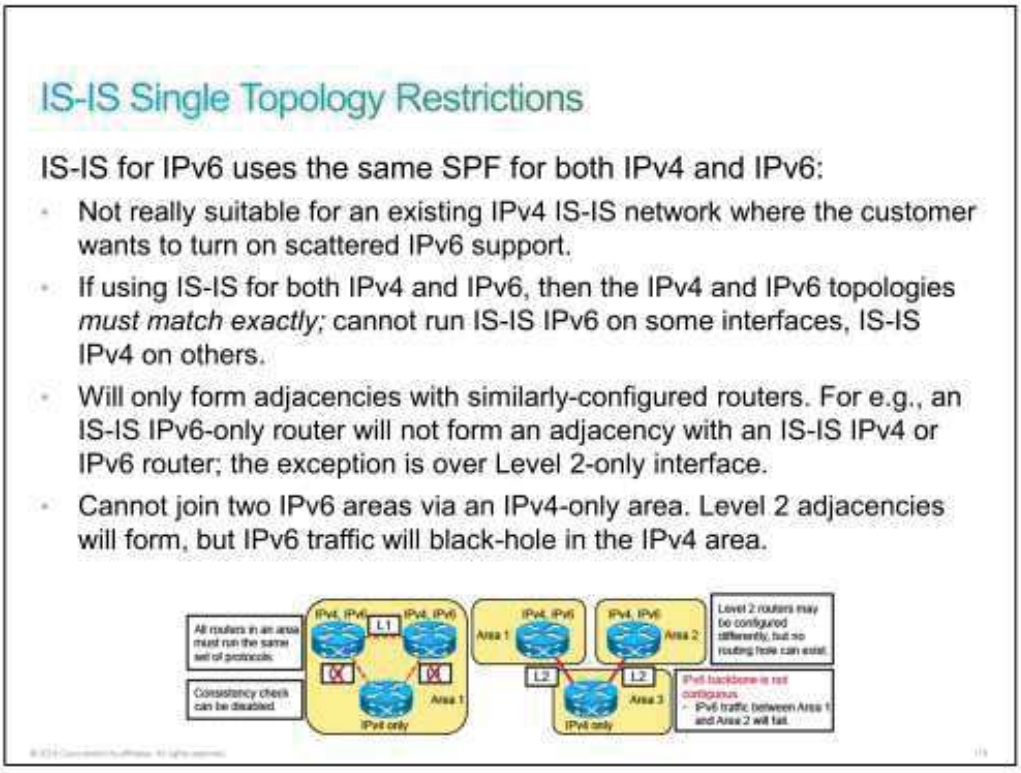

When you migrate from a purely IPv4 environment to a dual-stack environment, a discrepancy in supported protocols would cause adjacencies to fail. The intermediate system performs consistency checks on hello packets and will reject hello packets that do not have the same set of configured address families. For example, a router running IS-IS for both IPv4 and IPv6 will not form an adjacency with a router running IS-IS for IPv4 or IPv6 only. To facilitate a seamless upgrade, the engineer may disable consistency checks during the upgrade to maintain adjacencies active even in a heterogeneous environment.

Suppressing adjacency checking on intra-area links (Layer 1 links) is primarily done during transition from single-topology (IPv4) to multitopology (IPv4 and IPv6) IS-IS networks. Imagine that a service provider is integrating IPv6 into a network and it is not practical to shut down the entire provider router set for a coordinated upgrade. Without disabling adjacency checking—because routers were enabled for IPv6 and IS-IS for IPv6—adjacencies would drop with IPv4-only routers, and IPv4 routing would be severely impacted. With consistency check suppression, IPv6 can be turned up without impacting IPv4 reachability.

https://t.me/learningnets

Use the following commands to disable the adjacency check:

- On the Cisco IOS and Cisco IOS XE router:

```
router isis
 no adjacency-check
```

- On the Cisco IOS XR router:

```
router isis 1
 address-family ipv6 unicast
  adjacency-check disable
```

When single-topology support for IPv6 is employed, either old- or new-style TLVs can be used. However, the TLVs that are utilized to advertise reachability to IPv6 prefixes use wide metrics, so wide metrics should be used within the entire IS-IS domain.

In single-topology IPv6 mode, the configured metric is always the same for both IPv4 and IPv6. The reason for this is that IS-IS establishes routing adjacencies and builds the network topology using CLNS. IPv4 and IPv6 are just routed protocols; for routing information exchange, CLNS is used.

As in any IS-IS network design, Level 2 (backbone) routers must be contiguous. IPv6 adjacency checks are not done on Level 2 links. In the right-side diagram, the Level 2 routers are not contiguous for IPv6; therefore, this is an incorrect network design. The correct example would be when the Level 2 routers are contiguous for both IPv4 and IPv6.

This is called a "routing hole." In the example on the right, adjacencies will be formed across the Layer 2 links between the three areas. However, the IPv6 network is partitioned by the inability of Area 3 to carry IPv6 traffic. Because IS-IS is managing a single topology, the routers will believe that a path for IPv6 exists across Area 3, but all IPv6 traffic sent via that path will fail.

The **adjacency-check** command enables or disables adjacency IPv6 protocol-support checks. If enabled (default), the router will not form an adjacency with a neighbor not supporting IS-IS IPv6.

The **adjacency-check disable** IOS XR command (or **no adjacency-check** IOS/IOS XE command) will allow an adjacency to be formed between an IPv4–IPv6 router and an IPv4 router. This configuration may be convenient when transitioning an IS-IS IPv4 network to IPv4 and IPv6.

Note	The `adjacency-check disable` command suppresses IPv6 checks only. IS-IS IPv4 also checks the protocol support of neighbors and will not allow an adjacency between a router running IS-IS IPv4 and a neighbor not supporting IPv4. For example, IS-IS will never form an adjacency between a router running IPv4 IS-IS only and a router running IPv6 only or between a router running IPv4 and IPv6 IS-IS and a router running IPv6 only.

Note	Also, if the IS-IS router determines that the shortest path to an IPv6 destination lies via a non-IPv6 neighbor, the route to the destination will not be installed in the IPv6 routing table.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Multitopology IS-IS for IPv6

Multitopology IS-IS for IPv6 is available today and is based on the IETF document "M-ISIS: Multi Topology (MT) Routing in IS-IS" (RFC 5120).

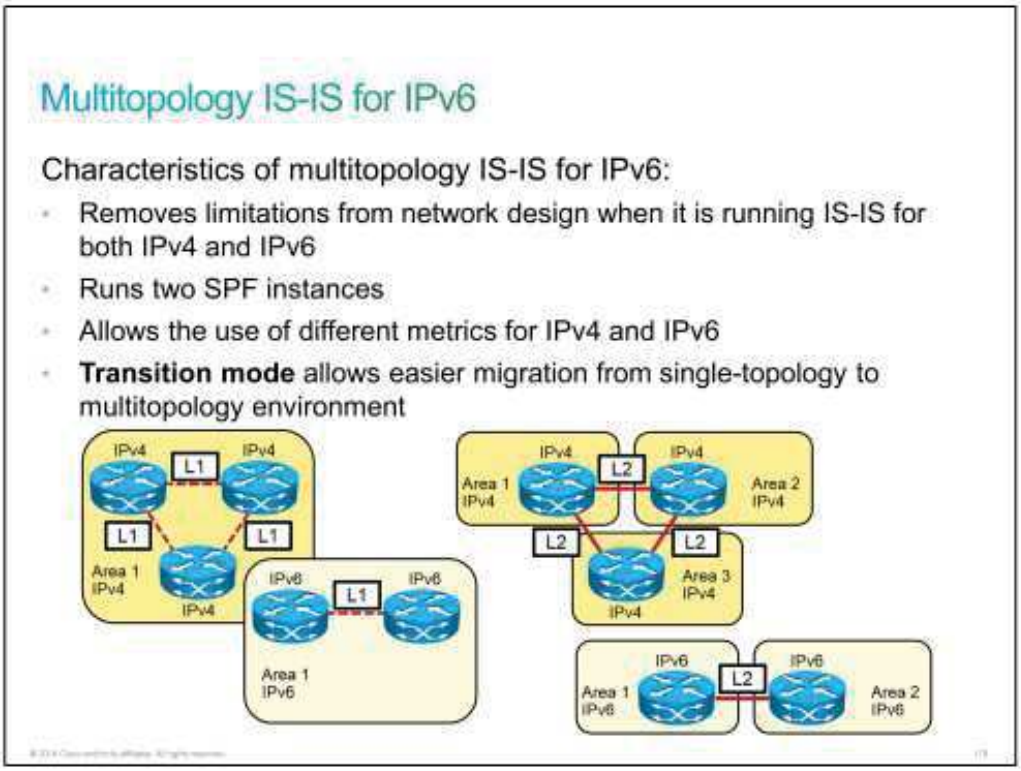

IS-IS multitopology support for IPv6 allows IS-IS to maintain a set of independent topologies within a single area or domain. This mode removes the restriction that all interfaces on which IS-IS is configured must support the identical set of network address families. It also removes the restriction that all routers in the IS-IS area (for Level 1 routing) or domain (for Level 2 routing) must support the identical set of network layer address families. Multiple SPFs are performed, one for each configured topology. Therefore, connectivity existing among a subset of the routers in the area or domain is sufficient for a given network address family to be routable.

When multitopology support for IPv6 is used, use the **metric-style wide** Cisco IOS, Cisco IOS XE, and Cisco IOS XR command to configure IS-IS to use new-style TLVs. TLVs used to advertise IPv6 information in LSPs are defined to use only wide metrics.

All routers in the area or domain must use the same type of IPv6 support, either single-topology or multitopology. A router operating in multitopology mode will not recognize the ability of the single-topology mode router to support IPv6 traffic, which will lead to routing holes in the IPv6 topology. To transition from single-topology support to the more flexible multitopology support, a multitopology transition mode is provided.

The multitopology transition mode allows a network operating in single-topology IS-IS IPv6 support mode to continue to work while upgrading routers to include multitopology IS-IS IPv6 support. While in transition mode, both types of TLVs (single-topology and multitopology) are sent in LSPs for all configured IPv6 addresses, but the router continues to operate in single-topology mode. After all routers in the area or domain have been upgraded to support multitopology IPv6 and are operating in transition mode, transition mode can be removed from the configuration. Once all routers in the area or domain are operating in multitopology IPv6 mode, the topological restrictions of single-topology mode are no longer in effect.

https://t.me/learningnets

Summary

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

Summary

- IS-IS is used in the IP edge and core networks of the Cisco IP NGN.
- Integrated IS-IS is an IS-IS for multiple protocols (IPv4, IPv6, and CLNS)
- IS-IS is a popular, stable, fast-converging IGP routing protocol in the service provider industry, that is positioned to route IPv4, IPv6, or CLNS.
- IS-IS uses similar mechanisms as OSPF.
- An NSAP is the OSI network-layer address. Cisco uses the area address (AFI plus area ID), system ID (6 bytes), and NSEL (0 for router) fields.
- IS-IS defines three types of routers: Level 1, Level 2, and Level 1-2.
- The IS-IS area address is used to route between areas, while the system ID is used to route within an area.
- With IS-IS, asymmetric routing can occur because Level 2 details are hidden from Level 1 routers.
- Route leaking can reduce suboptimal routing by allowing Level 2 routes to be inserted into Level 1.

Summary (Cont.)

- The four types of IS-IS PDUs are hello, LSP, PSNP, and CSNP. LSPs are used by routers to describe their characteristics. LSPs contain a header and TLV fields.
- For IPv6 routing, two TLVs have been added to IS-IS.
- IS-IS recognizes two topology types: point-to-point and broadcast.
- IS-IS broadcast mode is used on LAN and multipoint WAN interfaces.
- Level 1 LSPs are flooded within an area, while Level 2 LSPs are flooded throughout the Level 2 backbone.
- Sequence number PDUs (SNPs) are used to acknowledge the receipt of LSPs and to maintain LSDB synchronization.
- IS-IS adjacencies are established based on the area address and the router type.
- Single-topology IS-IS uses single topology for IPv4 and IPv6.
- Multitopology IS-IS uses separate topologies for IPv4 and IPv6 and supports different topologies for IPv4 and IPv6.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Implementing Integrated IS-IS Routing

Overview

Even when Integrated IS-IS protocol is used to support IP exclusively, network devices must still be configured to use the OSI CLNS protocol. Each IS-IS router requires a NET, and IS-IS packets are directly encapsulated onto the data-link layer instead of traveling inside IP packets.

The commands that are used to configure Integrated IS-IS are slightly different from those of the other IP routing protocols, so understanding how to enable IS-IS processes is important. Additionally, the default settings for IS-IS can result in the inefficient use of router and network resources and suboptimal routing; therefore, a network administrator also needs to know how to effectively tune IS-IS for optimum performance.

This lesson discusses the mechanics of Integrated IS-IS operation in an IP and CLNS environment, and outlines specific commands that are necessary to implement Integrated IS-IS on a Cisco router.

Upon completing this lesson, you will be able to:

- Describe the requirement for CLNS addressing, even when you are using IP in an IS-IS environment
- Describe the configuration process for Integrated IS-IS in an IP environment
- Describe how to optimize the IS-IS process
- Describe how to configure BFD for IS-IS
- Describe how to configure NSF for IS-IS
- Describe how to configure IP route summarization in IS-IS networks
- Describe how to verify IS-IS implementations
- Describe the **show** commands used to troubleshoot IS-IS operations
- Describe how to configure IS-IS to support IPv6

https://t.me/learningnets

Implement OSI Area Routing

A NET address identifies a device (an intermediate system or end system), and not an interface. This is a critical difference between a NET address and an IP address.

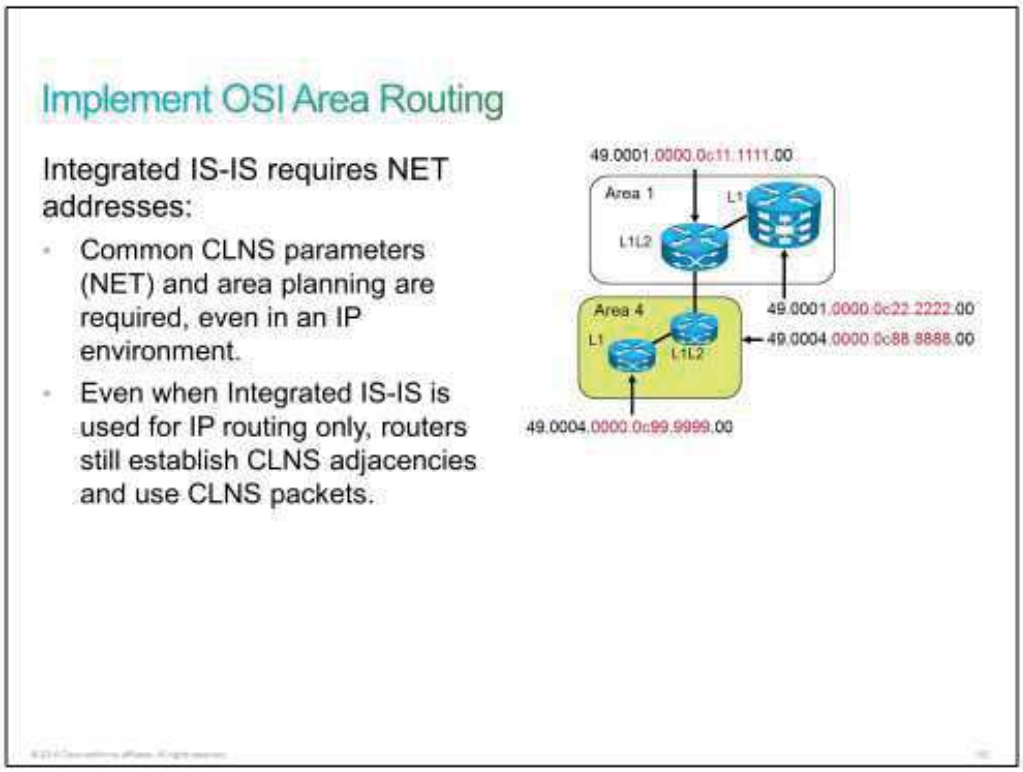

Even if you use Integrated IS-IS only for IP routing, each IS-IS router must have a NET address configured because Integrated IS-IS depends on the support of CLNS routing.

The OSI protocols (hello PDUs) are used to form the neighbor relationship between routers, and the SPF calculations rely on a configured NET address to identify the routers.

A device identifies other devices within its own area based on matching area addresses in their NET. It then knows that it can communicate with these other devices within the same area without using a default route. A default route is injected into the area by the Level 1-2 router. If the area addresses do not match, then the device knows that it must forward that interarea traffic to its nearest Level 1-2 router.

When you are using IS-IS to route IP traffic, IP subnets are treated as leaf objects that are associated with IS-IS areas. When you are routing IP traffic, the router looks up the destination network in its routing table. If the network belongs to a different area, then that interarea traffic is normally forwarded to the nearest Level 1-2 router.

https://t.me/learningnets

Implement OSI Area Routing (Cont.)

Building an IP Routing Table

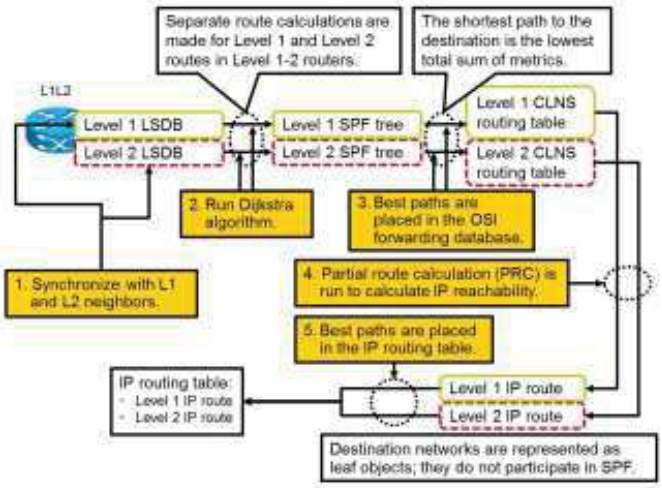

IS-IS uses an OSI forwarding database (routing table) to select the best path to a destination. When the databases are synchronized, routers use the LSDB to calculate the SPF tree to OSI destinations, the NETs. The total of the link metrics along each path determines the shortest path to any given destination.

The IS-IS L1/L2 router synchronizes its LSDB with its IS-IS neighbor (Step 1 in the figure). Level 1 and Level 2 routes have separate LSDBs; therefore, routers may run Dijkstra's algorithm twice (Step 2 in the figure), once for each level, and create separate SPF trees for each level. Routers insert the best paths in the CLNS routing table (the OSI forwarding database) (Step 3 in the figure).

Integrated IS-IS includes IP information in the LSPs, treating it as if it were end-system information, where the ISs are considered as nodes and the IP information is advertised by the ISs as the leaves hanging off the nodes in the shortest path tree. Therefore, updating IP reachability requires only a PRC (Step 4 in the figure); the entire shortest path tree in IS-IS does not need to be recomputed.

The PRC generates best-path choices for IP routes and offers the routes to the IP routing table, where they are accepted, based on normal IP routing table rules (Step 5 in the figure). For example, if more than one routing protocol is running, then the router compares administrative distance. When the IP IS-IS routes are entered in the routing table, they are shown as via Level 1 (i L1) or Level 2 (i L2), as appropriate.

The separation of IP reachability from the core IS-IS network architecture provides Integrated IS-IS better scalability than, for example, OSPF:

- OSPF sends LSAs for individual IP subnets. If an IP subnet fails, the LSA floods through the network and all routers must run a full SPF calculation, which is extremely CPU-intensive.
- Integrated IS-IS builds the SPF tree from CLNS information. If an IP subnet fails, the IS-IS LSP floods through the network, which is the same with OSPF. However, if this is a leaf (stub) IP subnet (that is, if the loss of the subnet does not affect the underlying CLNS architecture), the SPF tree is unaffected; therefore, only a PRC occurs.

https://t.me/learningnets

Implement IS-IS Routing

Four steps are required for the basic setup of IS-IS. Additional commands are available for fine-tuning the configuration.

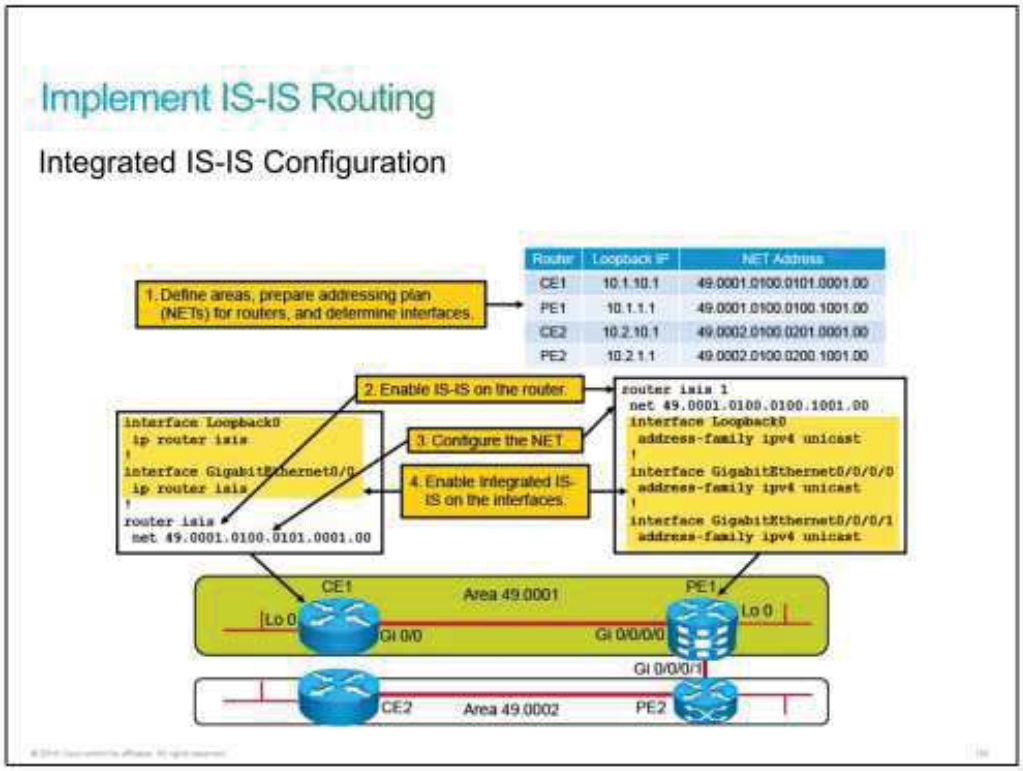

Before you configure Integrated IS-IS, you must map out the areas and plan the addressing. After that is done, you need three commands to enable Integrated IS-IS on a router for IP routing.

- The **router isis** *process-ID* Cisco IOS XR global command and the **router isis** Cisco IOS and Cisco IOS XE global configuration command enable IS-IS as an IP routing protocol. The Cisco IOS or IOS XE command allows you optionally to assign a tag to the process. Just as multiple OSPF processes can be present on the same router, multiple IS-IS processes are possible. The process name is significant only to the local router. If it is omitted, the Cisco IOS and Cisco IOS XE software assume a tag of 0. If more than one IS-IS process is used, then the network plan should indicate which interfaces would participate in which IS-IS process. By default, Cisco IOS, Cisco IOS XE, and Cisco IOS XR software make the router a Level 1-2 router.
- After the IS-IS process is enabled, the router must be identified for IS-IS by assigning a NET to the router with the **net** Cisco IOS, Cisco IOS XE, and Cisco IOS XR router configuration command.
- The final step is to select which interfaces will participate in IS-IS routing. As with other routing protocols, the **address-family ipv4 unicast** Cisco IOS XR command starts IPv4 unicast IS-IS routing on the appropriate interface and the **ip router isis** Cisco IOS and Cisco IOS XE interface configuration command enables IS-IS on the interfaces. Do not forget interfaces to stub IP networks, such as loopback interfaces. If there is more than one IS-IS process, the IS-IS process to which the interface belongs must be specified using the appropriate process name or ID.

You can then use additional commands to fine-tune the IS-IS processes.

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

Implement IS-IS Routing (Cont.)

Define IS-IS area and addressing:

- Area is determined by NET prefix.
  - Assign to support two-level hierarchy
- Addressing
  - IP: Plan to support summarization
  - CLNS: Prefix denotes area. System ID must be unique.

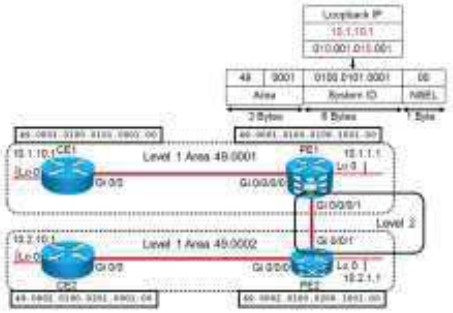

Recall that all intra-area traffic in IS-IS must traverse the Level 2 backbone area. Thus, CLNS addresses must be planned to execute a two-level hierarchy.

You must decide which routers will be backbone (Level 2) routers, which routers will be Level 1-2, and which will be internal area (Level 1) routers. If some routers must do both Level 1 and Level 2 routing, then you should identify the specific interfaces that will participate in each type of routing.

Remember that the CLNS address of a router is called the NET, and it consists of three main parts:

- The prefix, which identifies the area of the router
- The system ID, which uniquely identifies each device. The example in the figure shows how the system ID can be obtained from the router Loopback IP address.
- The NSAP selector NSEL that must be 0

It is not enough to plan the IS-IS area addressing. To have a scalable network, you must also plan IP addressing, and the IP addresses must be planned to allow for summarization of addresses.

Route summarization is the key idea that enables all the benefits of the hierarchical addressing design. Route summarization minimizes routing update traffic and resource utilization.

Be particularly careful when you configure the IP addressing on the router, because troubleshooting IP address misconfigurations with IS-IS is more difficult. The IS-IS neighbor relationships are established over OSI CLNS, not over IP. Because of this approach, two ends of a CLNS adjacency can have IP addresses on different subnets, with no impact to the operation of IS-IS.

https://t.me/learningnets

Optimizing the IS-IS Processes

Optimizing IS-IS will facilitate its smooth functioning and maximize its efficiency. Remember that using the default configuration of IS-IS results in the router having an intermediate system type of Level 1-2. Although this configuration has the advantage of allowing all routers to learn of each other and pass routes without too much administrative oversight, it is not the most efficient way to build an IS-IS network. Routers with the default configuration send out both Level 1 and Level 2 hellos and maintain both Level 1 and Level 2 LSDBs.

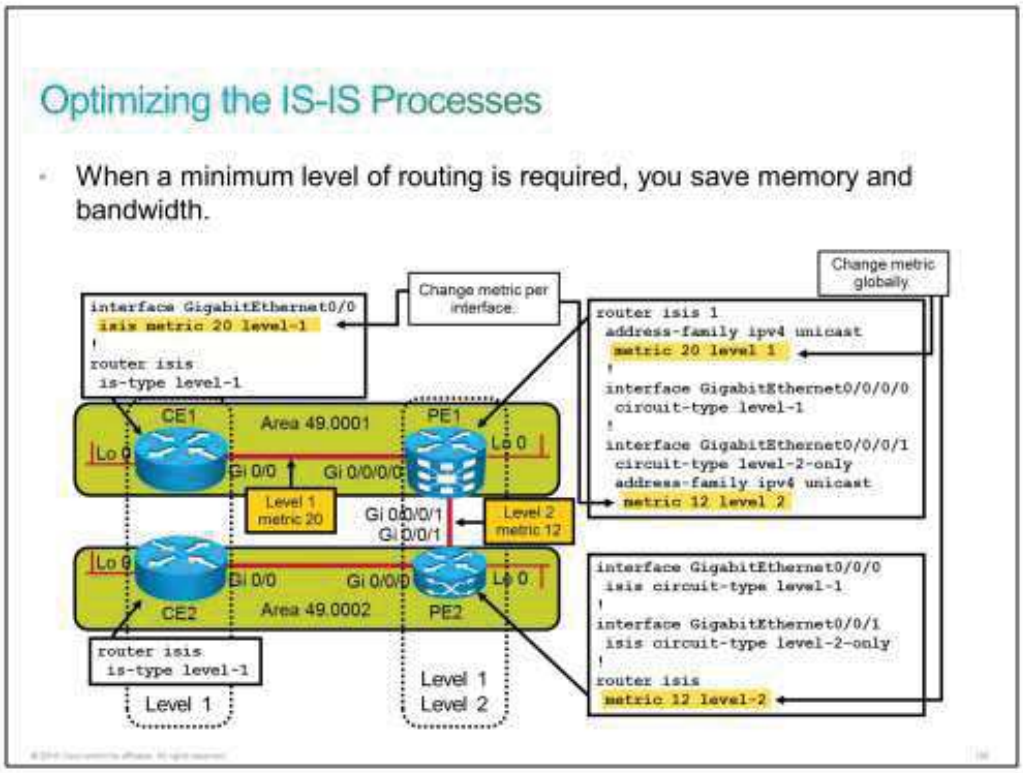

For efficiency, each router should be configured to support the minimum level of routing that is required, enabling the following benefits:

- **Saves memory:** If a router does not need the LSDB for one of the levels, it will not maintain one.
- **Saves bandwidth:** Hellos and LSPs will be sent only for the necessary level.

If a router is to operate only as an internal area router (L1) or a backbone router (L2), then specify this configuration by entering the **is-type** Cisco IOS, Cisco IOS XE, and Cisco IOS XR router configuration command.

Although the router can be a Level 1-2 router, establishing both types of L1 and L2 adjacencies over all interfaces may not be required. If a particular interface has only Level 1 routers connected to it, there is no need for the Level 1-2 router to send Level 2 hellos out that interface. Similarly, if an interface has only Level 2 routers connected to it, there is no need for the Level 1-2 router to send Level 1 hellos out that interface. Doing so would waste bandwidth and router resources by trying to establish adjacencies that do not exist. To make IS-IS more efficient in these types of situations, configure the interface to send only the needed type of hellos. Use the **isis circuit-type** Cisco IOS and Cisco IOS XE interface configuration command, or use the **circuit-type** Cisco IOS XR router interface configuration command.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Unlike most other IP protocols, IS-IS on a Cisco router does not take into account line speed or bandwidth when it sets its link metrics. All interfaces are assigned a metric value of 10 by default. In a network with links of varying types and speeds, this default assignment can result in suboptimal routing. To change the metric value, use the **isis metric** Cisco IOS and Cisco IOS XE interface configuration command, or use the **metric** Cisco IOS XR router interface configuration command. The metric can have different values for Level 1 and Level 2 over the same interface. If the metric value for all interfaces needs to be changed from the default value of 10, then the change needs to be performed one-by-one on all IS-IS interfaces. This change can be time-consuming and error-prone, especially for routers with many IS-IS interfaces. The **metric** Cisco IOS and Cisco IOS XE router configuration command, or **metric** Cisco IOS XR router address family configuration command changes the metric value for all IS-IS interfaces.

In the figure, there are two different areas: area 49.0001 (CE1 and PE1) and area 49.0002 (CE2 and PE2). The CE1 and CE2 routers need to do only Level 1 routing, but PE1 and PE2 routers do Level 1 and Level 2 routing.

https://t.me/learningnets

Bidirectional Forwarding Detection for IS-IS

BFD is a detection protocol, designed to provide fast-forwarding path-failure detection times that verifies the connectivity between two systems. BFD can be configured in two steps. The first step in configuring BFD is setting the baseline parameters for all BFD sessions on an interface.

Bidirectional Forwarding Detection for IS-IS

IS-IS uses BFD on all interfaces:

- Configuration on Cisco IOS:

```
router isis 1
 bfd all-interfaces
```

Bidirectional Forwarding Detection for IS-IS (Cont.)

IS-IS uses BFD on a specific interface only:

- Configuration on Cisco IOS:

```
interface TenGigabitEthernet3/0/1
 ip isis bfd
 bfd interval 100 min_rx 100 multiplier 3
```

- Configuration on Cisco IOS XR:

```
router isis 1
 interface TenGigE0/1/4/0
   bfd fast-detect
   bfd minimum-interval 100
   bfd multiplier 3
```


©2014 Cisco Systems, Inc.

https://t.me/learningnets

The BFD for IS-IS configuration occurs at the interface level, using the following syntax:

```
[no] bfd interval <50-999> min_rx <1-999> multiplier <3-50>
```

Baseline parameters are listed here:

- **interval**: Determines how frequently (in milliseconds) BFD packets will be sent to BFD peers.
- **min_rx**: Determines how frequently (in milliseconds) BFD packets will be expected to be received from BFD peers.
- **multiplier**: Determines the number of consecutive BFD packets that must be missed from a BFD peer before that peer is declared to be unavailable; also informs the higher-layer protocols of the failure.

Once the baseline parameters have been set, individual protocols must be informed that they will be using BFD for failure detection.

https://t.me/learningnets

Nonstop Forwarding for IS-IS

Usually, when a networking device restarts, all routing peers of that device detect that the device went down and then came back up. This transition results in what is called a routing flap. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. NSF allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover.

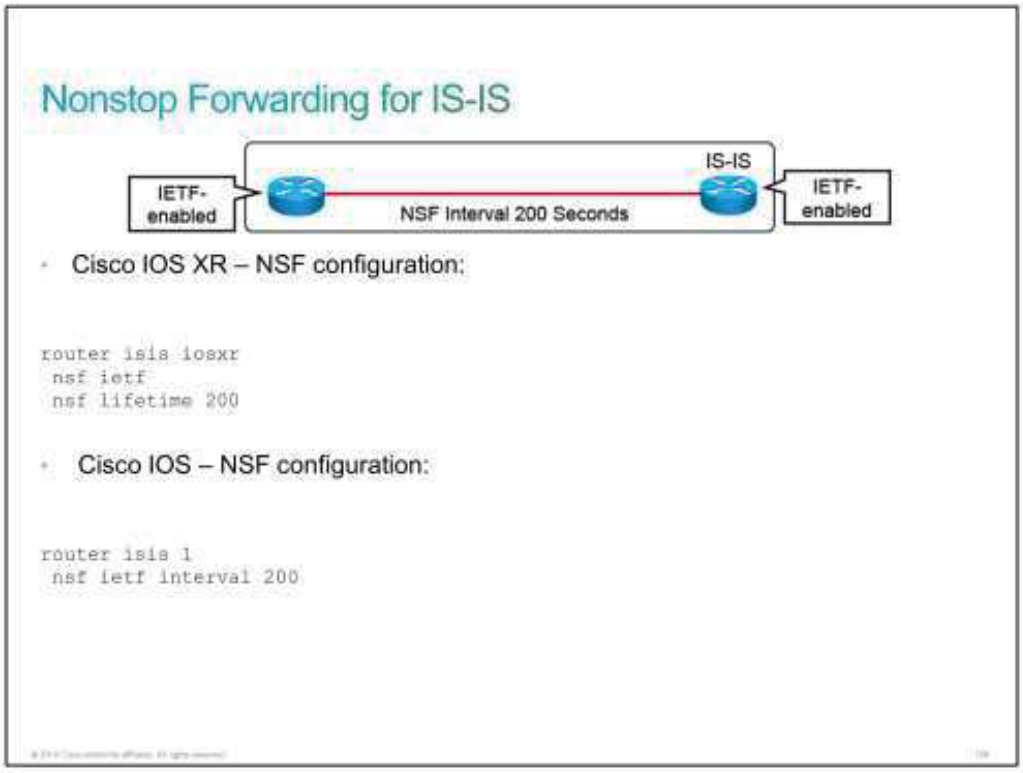

The IS-IS protocol can be configured to use state information that has been synchronized between the active and the standby route processor to recover route information following a switchover instead of information received from peer devices. When an IS-IS NSF-capable router performs an RP switchover, it must perform two tasks in order to resynchronize its Link State Database with its IS-IS neighbors. First, it must relearn the available IS-IS neighbors on the network without causing a reset of the neighbor relationship. Second, it must reacquire the contents of the Link State Database for the network.

The IS-IS NFS feature enables two options: IETF or Cisco Intermediate System-to-Intermediate System. With IETF, neighbor routers provide adjacency and link-state information to help rebuild the routing information following a switchover, whereas the Cisco IS-IS configuration transfers both protocol adjacency and link-state information from the active to the standby RP.

The example shows the enabling of the IETF NSF for IS-IS on the Cisco IOS or IOS XE and Cisco IOS XR routers.

The NSF restart interval has been changed from the 120-second default value to 200 seconds.

https://t.me/learningnets

IP Route Summarization Configurations in IS-IS Networks

Routing protocol scalability is a function of the appropriate use of route summarization.

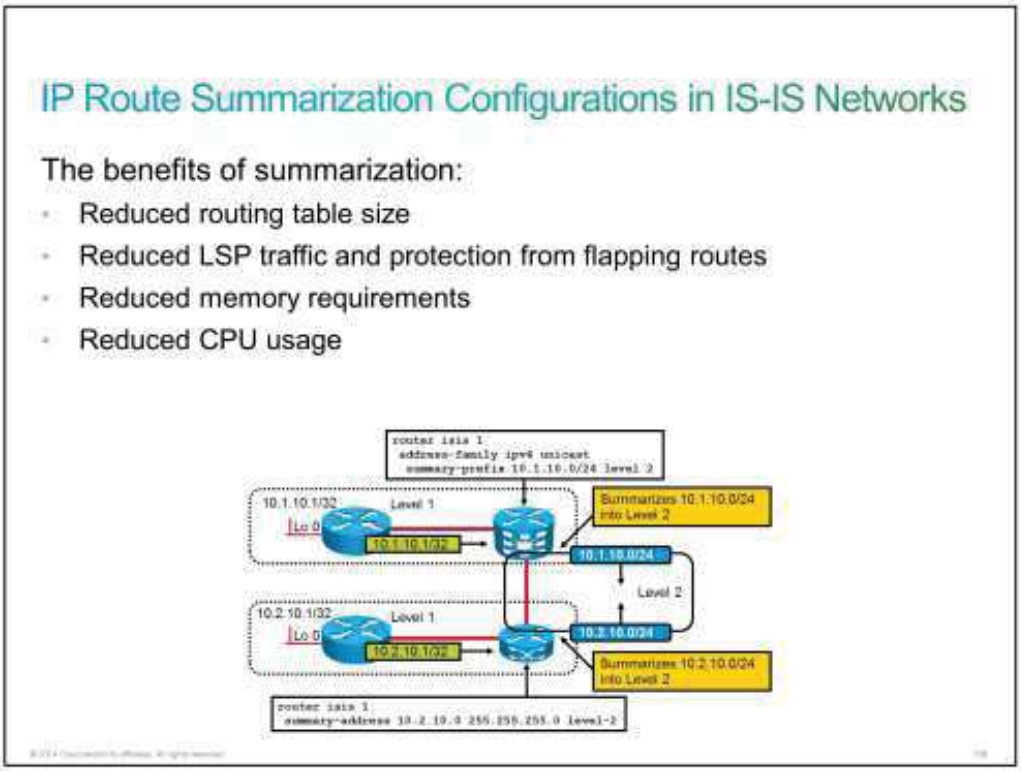

An intermediate system can be configured to aggregate a range of IP addresses into a summary address, using the **summary-address** Cisco IOS and Cisco IOS XE router configuration command, or the **summary-prefix** Cisco IOS XR router address family configuration command as shown in the figure.

This command can be used on any router in an IS-IS network. The router summarizes IP routes into Level 1, Level 2, or both.

https://t.me/learningnets

Verification of IS-IS

To verify the IS-IS configuration and IP functionality of the Integrated IS-IS network, use the **show protocols ipv4** Cisco IOS XR command, or the **show ip protocols** Cisco IOS and Cisco IOS XE command. These commands display the active IP routing protocols, the interfaces on which they are active, and the networks for which they are routing.

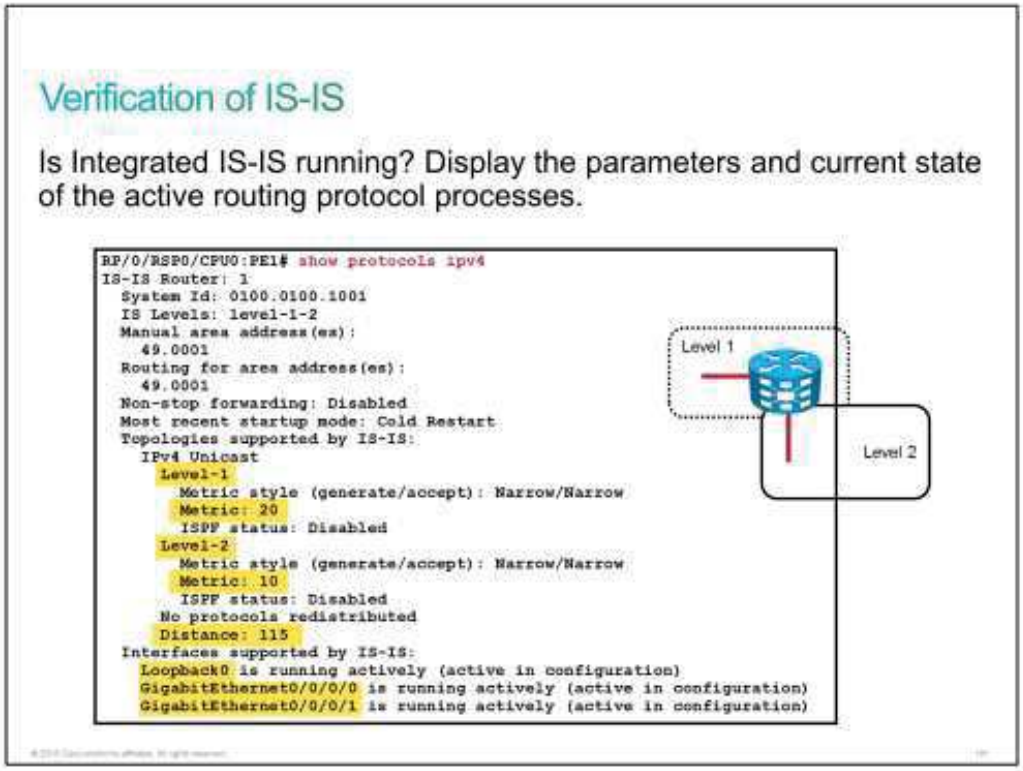

The example shows that interfaces GigabitEthernet 0/0/0/0, GigabitEthernet 0/0/0/1, and Loopback 0 are taking part in Integrated IS-IS, that the metric is different for Level-1 and Level-2 interfaces, that the default metric style is used, and that the default administrative distance of Integrated IS-IS is 115.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

The figure shows an output from a Cisco IOS XR router; similar output from a Cisco IOS and Cisco IOS XE router is displayed here:

```
PE2# show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "isis"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Redistributing: isis
  Address Summarization:
    10.2.10.0/255.255.255.0 into level-2
  Maximum path: 4
  Routing for Networks:
    Loopback0
    GigabitEthernet0/0/0
    GigabitEthernet0/0/1
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.10.1             115      01:59:01
    10.2.10.1             115      00:11:38
    10.1.1.1              115      00:11:38
  Distance: (default is 115)
```

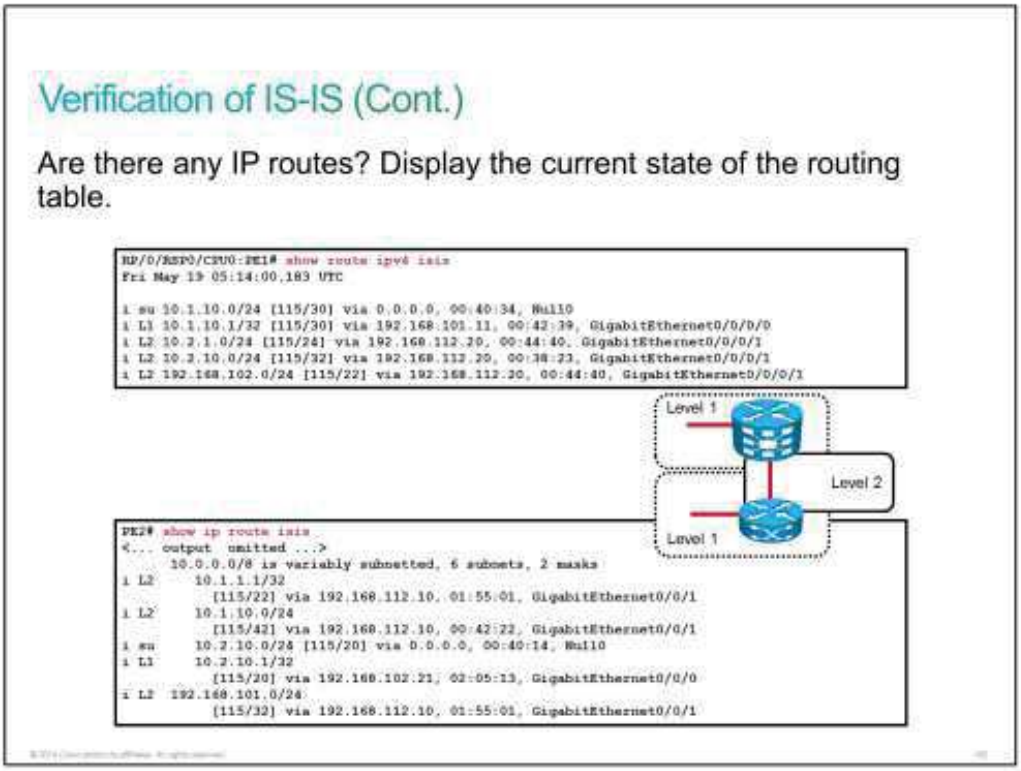

These sample outputs from the **show route ipv4 isis** Cisco IOS XR command, or **show ip route isis** Cisco IOS and Cisco IOS XE command show only the IS-IS routes. One route is from Level 1, as indicated by the **i L1** tag, and the others are from Level 2, as indicated by the **i L2** tag. The route that the router summarizes is indicated by the **i su** tag.

By default, Integrated IS-IS uses an administrative distance of 115. The metric that is shown for each route is taken from the IS-IS cost to the destination.

https://t.me/learningnets

Troubleshooting IS-IS Commands

You can use the **show** commands to verify the router configuration and to troubleshoot the Integrated IS-IS network.

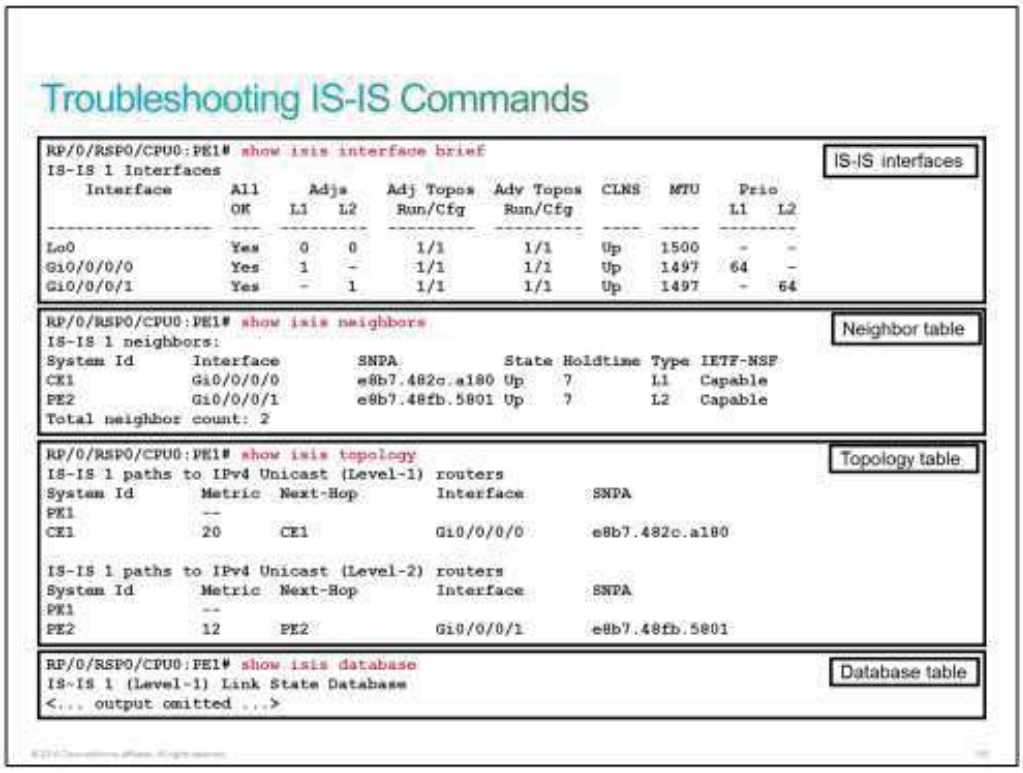

Use the show Commands to Verify the Router Configuration

Command	Description
show isis interface	This Cisco IOS XR command displays the interfaces that are enabled for IS-IS.
show isis neighbors	This Cisco IOS, Cisco IOS XE, and Cisco IOS XR command displays the IS-IS neighbors that are recognized by the system.
show isis topology	This Cisco IOS, Cisco IOS XE, and Cisco IOS XR command displays the Level 1 and Level 2 topology tables, which show the least-cost IS-IS paths to the intermediate systems.
show isis database	This Cisco IOS, Cisco IOS XE, and Cisco IOS XR command displays the contents of the IS-IS LSDB. To force IS-IS to refresh its LSDB and recalculate all routes, issue the **clear isis** Cisco IOS and Cisco IOS XE command (an asterisk (*) can be used to clear all IS-IS processes), or use the **clear isis process** Cisco IOS XR command.
show isis database *LSP_ID* **detail**	This Cisco IOS, Cisco IOS XE, and Cisco IOS XR command displays the detailed contents of each IS-IS LSDB and advertised prefixes for specified routing process.

The examples in the figure show outputs of Cisco IOS XR commands.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

This is an example of the **show isis neighbors** command that was taken from a Cisco IOS and Cisco IOS XE router:

```
PE2# show isis neighbors
System Id      Type Interface   IP Address      State Holdtime Circuit Id
CE2            L1   Gi0/0/0     192.168.102.21  UP    25       PE2.02
PE1            L2   Gi0/0/1     192.168.112.10  UP    26       PE2.03
```

This is an example of the **show isis topology** command that was taken from a Cisco IOS and Cisco IOS XE router:

```
PE2# show isis topology
IS-IS TID 0 paths to level-1 routers
System Id            Metric     Next-Hop            Interface   SNPA
PE2                  --
CE2                  10         CE2                 Gi0/0/0     4055.3986.f968
IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop            Interface   SNPA
PE1                  12         PE1                 Gi0/0/1     4055.392e.c421
PE2                  --
```

This is an example of the **show isis database** command that was taken from a Cisco IOS XR router:

```
RP/0/RSP0/CPU0:PE1# show isis database
Fri May 19 07:48:44.643 UTC
IS-IS 1 (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00           * 0x0000002a   0xb232        545             1/0/0
CE1.00-00             0x00000028   0xa163        570             0/0/0
CE1.01-00             0x00000024   0xd940        1177            0/0/0
Total Level-1 LSP count: 3     Local Level-1 LSP count: 1
IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00           * 0x0000002d   0xc56d        643             0/0/0
PE2.00-00             0x00000030   0x80bd        570             0/0/0
PE2.03-00             0x00000024   0x9be2        980             0/0/0
Total Level-2 LSP count: 3     Local Level-2 LSP count: 1
```

This is an example of the **show isis database** command that was taken from a Cisco IOS and Cisco IOS XE router:

```
PE2# show isis database
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE2.00-00           * 0x0000002D   0xEEFA        1055              1/0/0
PE2.02-00           * 0x00000023   0xC045        480               0/0/0
CE2.00-00             0x00000029   0xAC54        534               0/0/0
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00             0x0000002D   0xC56D        674               0/0/0
PE2.00-00           * 0x00000030   0x80BD        604               0/0/0
PE2.03-00           * 0x00000024   0x9BE2        1014              0/0/0
```

This is an example of the **show isis database** *LSP_ID* **detail** command that was taken from a Cisco IOS XR router:

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1# show isis database CE1.00-00 detail
Mon May 26 14:01:12.828 UTC

IS-IS 1 (Level-1) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
CE1.00-00             0x00000010   0xd114        469             0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  Hostname:     CE1
  IP Address:   10.1.10.1
  Metric: 10         IP 10.1.10.1/32
  Metric: 10         IP 192.168.101.0/24

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
The requested LSP CE1.00-00 was not found in the IS-IS 1 Level-2 LSP Database
```

This is an example of the **show isis database** *LSP_ID* **detail** command that was taken from a Cisco IOS and Cisco IOS XE router:

```
PE2# show isis database PE1.00-00 detail

IS-IS Level-2 LSP PE1.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
PE1.00-00             0x00000017   0xED8C        564               0/0/0
  Area Address: 49.0001
  NLPID:        0xCC 0x8E
  Hostname: PE3
  IP Address:   10.1.1.1
  Metric: 10         IS PE1.01
  Metric: 10         IP 10.1.1.1 255.255.255.255
  Metric: 10         IP 192.168.101.0 255.255.255.0
  Metric: 10         IP 192.168.112.0 255.255.255.0
  Metric: 20         IP-External 10.3.10.0 255.255.255.0
```

Note	The default IS-IS metric for every link is 10. Metric 20 means that the network is not directly connected but is directly connected to the neighbor router.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Configuring IS-IS to Support IPv6

To configure IS-IS to support IPv6 follow these steps:

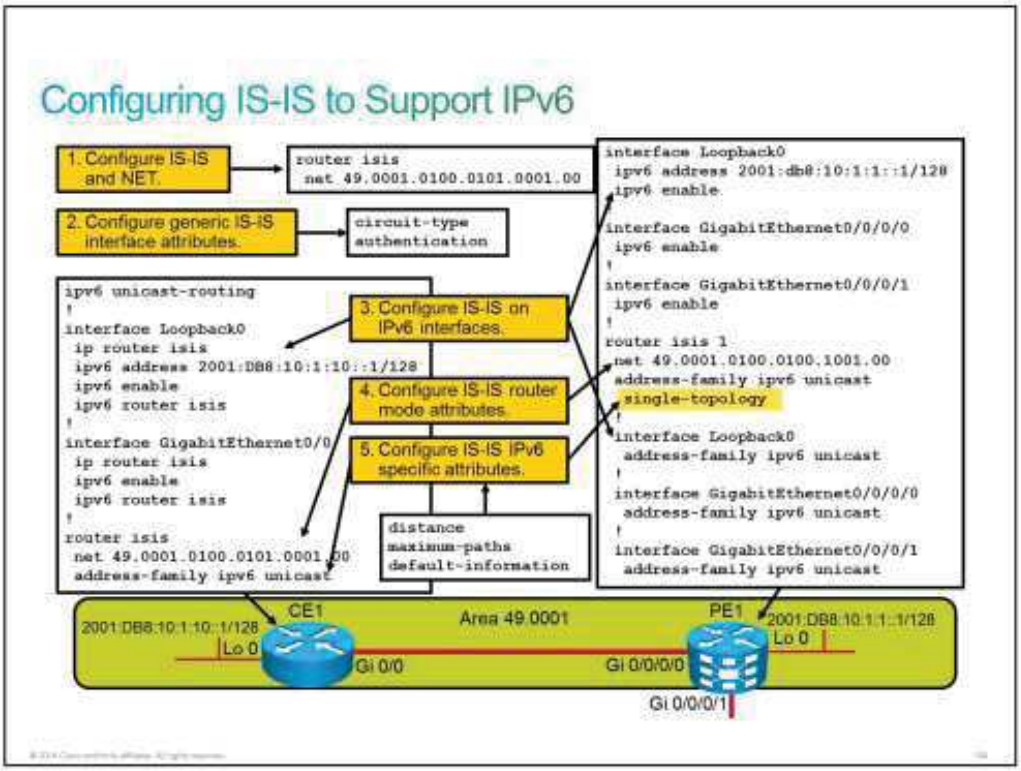

Configuring IS-IS to Support IPv6

Follow these steps to configure IS-IS for IPv6 support:

Step 1 Configure the IS-IS routing process and specify the network entity title.

Step 2 Configure generic IS-IS interface attributes using existing IS-IS commands (circuit-type, authentication, and so on).

Step 3 Configure IS-IS on interfaces. The interfaces must have the IPv6 protocol stack enabled, such as, have an IPv6 address assigned, or autoconfigured.

Step 4 Configure IS-IS router mode attributes. The majority of IS-IS router mode commands are generic and apply to both IPv4 and IPv6. The Cisco IOS XR single-topology router command is valid only in the IPv6 submode. The command instructs IPv6 to use the single topology rather than the default configuration of a separate topology in the multitopology mode.

Step 5 Configure the IS-IS IPv6 specific attributes. IPv6 attributes are configured via the IPv6 address family submode of the router mode.

https://t.me/learningnets

These IS-IS IPv6 Cisco IOS, IOS XE, and IOS XR commands and attributes are used under the address-family IPv6 submode, and are applied to the IPv6 routing table only:

- **distance**: Sets the administrative distance of IS-IS IPv6. The default administrative distance for IS-IS is 115.
- **maximum-paths**: Sets the maximum number of paths allowed for a route learned via IS-IS IPv6. The default number is four.
- **default-information originate**: Configures origination of the IPv6 default route (::/0) by IS-IS. It is used in the same manner as the existing IPv4 **default-information** command.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Summary

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

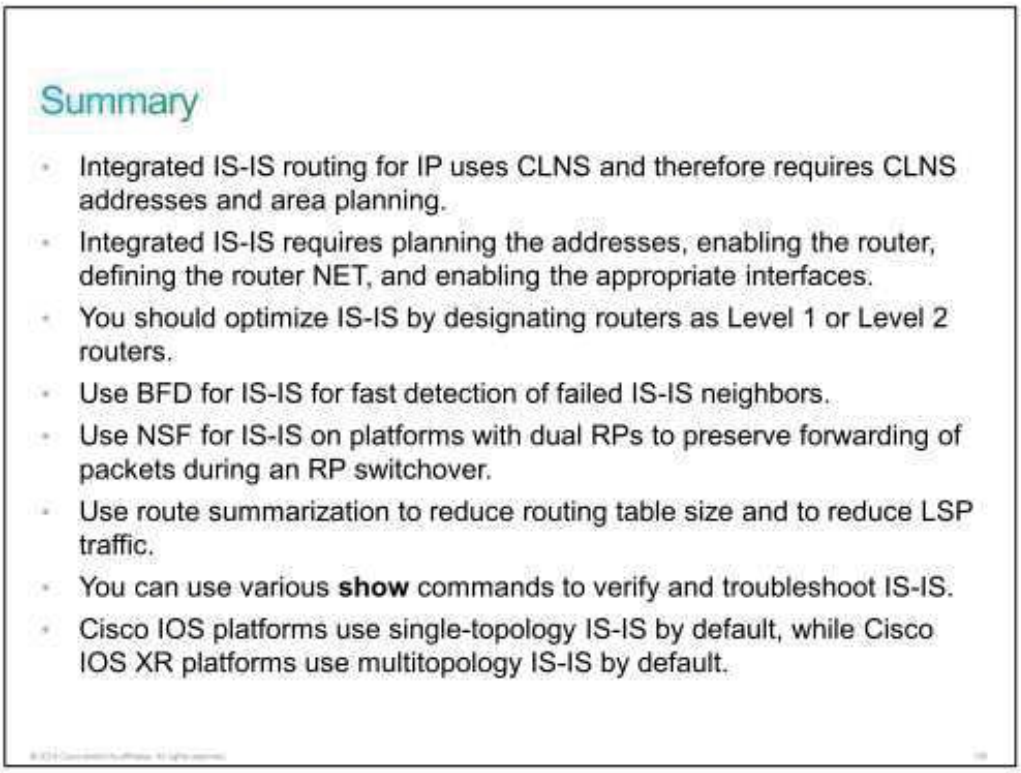

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Module Summary

Overview

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

CLNS provides connectionless delivery of data. As a result, CLNS is the solution for unreliable delivery of data, similar to IP. IS-IS operates in strictly CLNS terms, although Integrated IS-IS supports IP routing as well as CLNS.

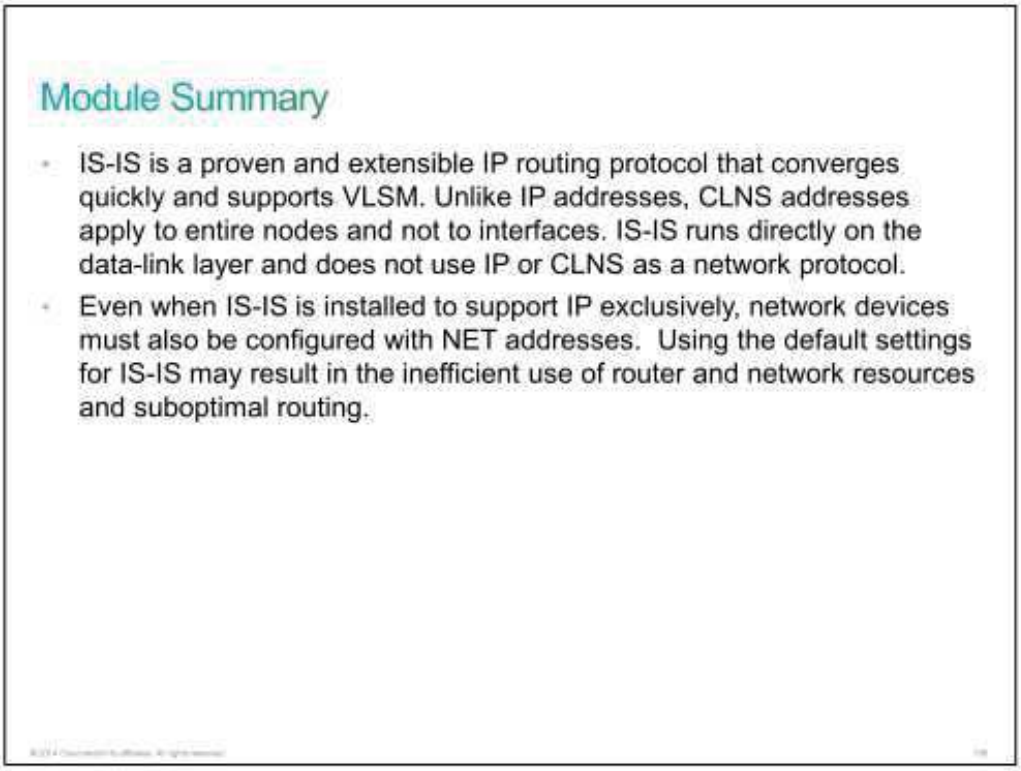

https://t.me/learningnets

 ©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. Which two protocols does Integrated IS-IS support? (Choose two.) (Source: Introducing IS-IS Routing)

   A. IP
   B. IPX
   C. OSI CLNS
   D. AppleTalk

2. Which two options describe IS-IS? (Choose two.) (Source: Introducing IS-IS Routing)

   A. It is an IGP.
   B. It is an EGP.
   C. It is efficient in its use of network resources.
   D. It is an advanced distance vector routing protocol.

3. Which two routing levels does IS-IS support? (Choose two.) (Source: Introducing IS-IS Routing)

   A. Level 1 for IP
   B. Level 1 for interarea routing
   C. Level 2 for interarea routing
   D. Level 2 for CLNS
   E. Level 0 for interdomain routing
   F. Level 1 for intra-area routing
   G. Level 2 for intra-area routing

https://t.me/learningnets

4. Which two options are the responsibility of IS-IS? (Choose two.) (Source: Introducing IS-IS Routing)

A. Level 0 routing
B. Level 1 routing
C. Level 2 routing
D. Level 3 routing

5. Routing between areas is described as _____. (Source: Introducing IS-IS Routing)

A. Level 0 routing
B. Level 1 routing
C. Level 2 routing
D. Level 3 routing

6. Which of the following characteristics can be attributed to IS-IS?

A. link-state protocol
B. fast convergence
C. supports VLSM
D. more extensible
E. documentation and experienced engineers are easy to find
F. most customized to IP
G. metrics scale automatically

7. Level 2 routing is responsible for which task? (Source: Introducing IS-IS Routing)

A. exchanging information about paths between areas
B. building a topology of end systems and intermediate systems in areas
C. using ES-IS to learn prefix information
D. CLNP routing

8. A good IS-IS design features which two properties? (Choose two.) (Source: Introducing IS-IS Routing)

A. CLNS addresses confined to Level 1
B. a summarizable address plan
C. a two-level hierarchy
D. routers with minimal memory and CPU

9. Which two options describe the advantages of IS-IS over OSPF? (Choose two.) (Source: Introducing IS-IS Routing)

A. better support from the IETF
B. more documentation and support
C. ubiquitously implemented
D. support for CLNS
E. more extensible
F. faster convergence

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

10. Put the parts of an NSAP address in the correct order. (Source: Introducing IS-IS Routing)

AFI

IDI

HO-DSP

system ID

NSEL

11. Which component does a NET address identify? (Source: Introducing IS-IS Routing)

A. interface
B. device
C. process
D. protocol

12. Which characteristic describes a NET? (Source: Introducing IS-IS Routing)

A. always has an AFI of 49
B. always has a 3-byte HO-DSP
C. always has an NSEL of 00
D. is never assigned to a router

13. Which prefix does private AFI use? (Source: Introducing IS-IS Routing)

A. 49.0001
B. 37
C. 47
D. 49
E. 37.1921

14. Which term describes an NSAP address? (Source: Introducing IS-IS Routing)

A. NSCAR address
B. CLNS address
C. protocol-specific port
D. replacement for BGP

15. Which two statements about CLNS are true? (Choose two.) (Source: Introducing IS-IS Routing)

A. Within an area, all system IDs must be unique.
B. Within an area, all area addresses (AFI, IDI, HO-DSP) must be identical.
C. Within an area, AFI and IDI must be identical, but HO-DSP may vary.
D. Within an area, all IP addresses must be in the same classful network.

https://t.me/learningnets

16. Which three options are requirements for system ID inside a NET address? (Choose three.) (Source: Introducing IS-IS Routing)

A. must be 8 bytes
B. must be unique in an area
C. must be classless
D. must be unique within Level 2 routers in a routing domain
E. must be 6 bytes on a Cisco router

17. With OSI routing, a Level 1-2 router should compare the destination area address to its own area address, and if they are _____ (Choose two.) (Source: Introducing IS-IS Routing)

A. the same, route at Level 1
B. different, route at Level 1
C. the same, route at Level 2
D. different, route at Level 2

18. Route leaking allows Level 1 information to be leaked into Level 2. (Source: Introducing IS-IS Routing)

A. true
B. false

19. Adjacencies are formed in IS-IS between _____. (Choose two.) (Source: Introducing IS-IS Routing)

A. intermediate systems in the same area
B. pseudonodes
C. two intermediate systems doing routing at the same level (Level 1 or Level 2)
D. interfaces on a broadcast network

20. How are IS-IS PDUs transported? (Source: Introducing IS-IS Routing)

A. within IP packets
B. within CLNS packets
C. directly within frames
D. reliably using TCP

21. What are the four types of IS-IS PDUs? (Choose four.) (Source: Introducing IS-IS Routing)

A. hello
B. SNAP
C. LSP
D. PSNP
E. CSNP
F. IP/IPX
G. OSPF

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

22. What does TLV stand for? (Source: Introducing IS-IS Routing)

A. time, level, value
B. text, level, volume
C. time, length, volume
D. type, length, value

23. Which three topology types are supported by IS-IS? (Choose three.) (Source: Introducing IS-IS Routing)

A. broadcast for LAN links
B. broadcast for multipoint WAN links
C. stub for networks with a single exit
D. point-to-point for LAN links
E. point-to-point for point-to-point WAN links
F. NBMA for WAN links

24. How does IS-IS manage DIS failures? (Source: Introducing IS-IS Routing)

A. The backup intermediate system asserts itself as the new DIS.
B. The failure is recognized quickly and a new DIS is elected.
C. The network reverts to general topology, and every intermediate system forms an adjacency with every other intermediate system.
D. It does not. The DIS must be restored quickly to prevent communication problems.

25. What is the function of CSNPs? (Source: Introducing IS-IS Routing)

A. to request a missing LSP
B. to acknowledge an LSP
C. to provide alerts
D. to maintain LSDB synchronization

26. Which address component is used to identify the router in an IS-IS environment? (Source: Introducing IS-IS Routing)

A. SNPA
B. NRLI
C. NET
D. System number

27. Which two network representations are supported by IS-IS? (Choose two.) (Source: Introducing IS-IS Routing)

A. broadcast
B. NBMA
C. stub
D. point-to-point
E. NSSA

https://t.me/learningnets

28. Which communication method is used by devices in two Level 1 areas? (Source: Introducing IS-IS Routing)

A. using a tunnel
B. on broadcast multiaccess networks
C. through Level 2
D. using PDUs

29. What does the network entity title identify in IP networks? (Source: Introducing IS-IS Routing)

A. a router interface
B. a router process
C. a router
D. a subnet

30. A Level 1-2 intermediate system with the NET 49.000A.0000.0C12.3456.00 receives traffic going to 49.001A.0000.0C78.9AB.00. Which table does it use to route the traffic? (Source: Introducing IS-IS Routing)

A. IS-IS topology
B. Level 1 routing table
C. Level 2 routing table
D. CLNS routing table

31. Which address or addresses does Integrated IS-IS require? (Source: Introducing IS-IS Routing)

A. IP address to use as router ID
B. IP addresses on interfaces
C. CLNS address to identify the device
D. CLNS addresses on interfaces

32. Which two characteristics about Dijkstra calculations in IS-IS are true? (Choose two.) (Source: Introducing IS-IS Routing)

A. perform on separate Level 1 and Level 2 databases
B. perform on TLV
C. use shortest path—lowest sum of link metrics
D. can be disabled

33. Routing from 49.BAD1.1921.6800.0111.00 to 49.BAD7.1921.6800.0112.00 takes place at what level? (Source: Introducing IS-IS Routing)

A. Level 1
B. Level 2
C. Not enough information to determine
D. Routing not possible between these networks

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

34. What is the default IS-IS metric for GigabitEthernet interfaces? (Source: Introducing IS-IS Routing

A. 10
B. 16
C. 83
D. 100

35. IP routing scalability is achieved by _____. (Source: Introducing IS-IS Routing)

A. NET assignment
B. route summarization
C. controlling ES-IS
D. limiting IS-IS resynchronization issues

36. Which Cisco IOS XR command displays the IS-IS routes in the routing table? (Source: Implementing Integrated IS-IS Routing)

A. **show clns route**
B. **show ip neighbor**
C. **show isis route**
D. **show which-route**
E. **show route ipv4 isis**

37. Which Cisco IOS XR command provides a list of all IS-IS areas known to a router? (Source: Implementing Integrated IS-IS Routing)

A. **show clns route**
B. **show ip route**
C. **show which-route**
D. **show isis topology**

38. Which Cisco IOS or IOS XE command reveals the redistribution processes in IS-IS? (Source: Implementing Integrated IS-IS Routing)

A. **show ip protocols**
B. **show clns interface**
C. **show isis route**
D. **show isis database**

39. Which Cisco IOS XR command would you use to view the router interfaces that are currently running IS-IS? (Source: Implementing Integrated IS-IS Routing)

A. **show clns protocol**
B. **show isis interface brief**
C. **show isis route**
D. **show isis database**

https://t.me/learningnets

40. Place the IS-IS configuration steps in the correct order. (Source: Implementing Integrated IS-IS Routing)

define areas and addressing

enable IS-IS on the router.

configure NET

enable IS-IS on the interfaces

41. The router identifies which interfaces participate in IS-IS routing by the _____. (Source: Implementing Integrated IS-IS Routing)

A. **network** command
B. NET that is configured for each interface
C. **ip router isis** command.
D. NET that is configured for each router.

42. IS-IS summarization allows you to _____. (Source: Implementing Integrated IS-IS Routing)

A. summarize the list of NETs
B. summarize the area address
C. summarize a set of IP addresses into a less specific address
D. enumerate the intermediate system and end system neighbors

43. Which Cisco IOS or IOS XE command shows the sources of Integrated IS-IS routing information? (Source: Implementing Integrated IS-IS Routing)

A. **show ip protocols**
B. **show clns protocols**
C. **show ip route isis**
D. **show ip route**

44. What is the default IS-IS routing level of a Cisco router? (Source: Implementing Integrated IS-IS Routing)

A. 0
B. Level 1
C. Level 2
D. 3
E. Level 1-2

45. A NET address is required to configure Integrated IS-IS for routing IP only. (Source: Implementing Integrated IS-IS Routing)

A. true
B. false

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

46. To configure IS-IS on an interface, which Cisco IOS/IOS XE command must be executed from interface configuration mode? (Source: Implementing Integrated IS-IS Routing)

A. **router isis**
B. **isis interface**
C. **ip router isis**
D. **is-type level-1**

47. What is the default IS-IS metric for GigabitEthernet interfaces? (Source: Implementing Integrated IS-IS Routing)

A. 10
B. 16
C. 83
D. 100

https://t.me/learningnets

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

1. A, C
2. A, C
3. C, F
4. B, C
5. C
6. A, B, C, D
7. A
8. B, C
9. D, E
10. — AFI
    — IDI
    — HO-DSP
    — system ID
    — NSEL
11. B
12. C
13. D
14. B
15. A, B
16. B, D, E
17. A, D
18. B
19. A, C

https://t.me/learningnets

20. C
21. A, C, D, E
22. D
23. A, B, E
24. B
25. D
26. C
27. A, D
28. C
29. C
30. C
31. C
32. A, C
33. B
34. A
35. B
36. E
37. D
38. A
39. B
40. — define areas and addressing
    — enable IS-IS on the router.
    — configure NET
    — enable IS-IS on the interfaces
41. C
42. C
43. A
44. E
45. A
46. C
47. A

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module 4: Implement BGP in the Service Provider Network

Introduction

This module describes BGP and its importance to the service provider environment. The primary goal of the module is to describe the BGP implementation steps that connect an enterprise to a service provider, and connect a service provider to upstream service providers. The use of BGP as a routing protocol requires that an administrator understand how to properly configure BGP for scalable internetworking.

Upon completing this module, you will be able to:

- Use different connectivity options to connect an enterprise network to a service provider
- Explain the advantages of using BGP to connect an enterprise network to a service provider
- Implement internal and external BGP sessions to enable service provider routing

https://t.me/learningnets

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 1: Enterprise Connectivity to Service Providers

Overview

To properly plan the enterprise connectivity to a service provider, the designer must understand different connectivity options. Requirements, such as IP addressing, link type and bandwidth availability, routing options, and redundancy should be considered to choose from different connectivity types.

This lesson first describes enterprise connectivity requirements and the routing methods that are available. The lesson concludes with the different types of connectivity that affect routing method selection, link and bandwidth availability, IP addressing, and redundancy.

Upon completing this lesson, you will be able to:

- Describe connectivity requirements when connecting an enterprise network to a service provider
- Describe routing method selection when connecting an enterprise to a service provider
- Describe the different connection options between enterprises and service providers
- Describe single-homed enterprises
- Describe dual-homed enterprises
- Describe multihomed enterprises
- Describe dual-multihomed enterprises

https://t.me/learningnets

Enterprise Connectivity Requirements

Modern corporate networks are connected to the Internet and make use of the Internet for some of their data transport needs. Establishing enterprise connectivity to a service provider requires knowledge and careful planning.

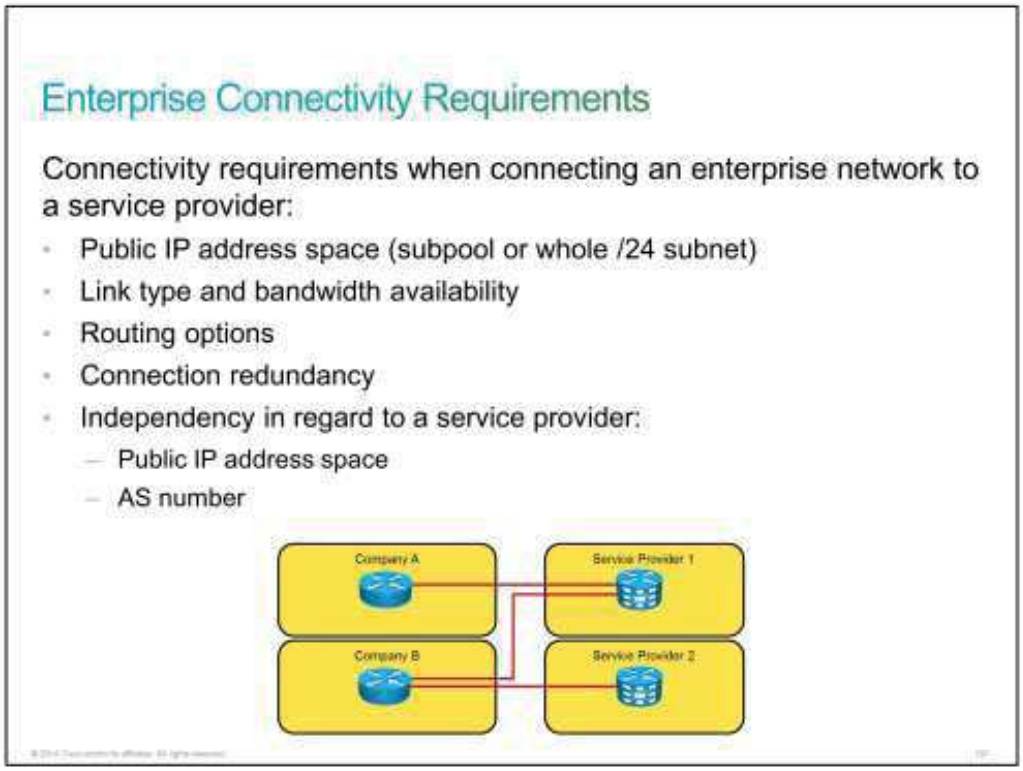

The first parameter to be determined is the number of public IP addresses that are needed. Public IP addresses are used to translate client private addresses for those clients that need to access resources on the Internet. These public IP addresses are also used for those enterprise servers that need to be accessible from the Internet. These servers are either equipped with public addresses or with addresses that are statically translated from private to public addresses.

The link type and speed, which depend on the service provider, are the second parameter to be determined. The link types that are available include leased line, Ethernet over fiber optics or copper, and xDSL. The bandwidth assignment must be done properly to address the connectivity requirements.

Selection of a proper routing protocol is the third parameter. The selection is typically made between static and dynamic routing.

The fourth parameter relates to the issue of connectivity redundancy. An evaluation is needed to determine which kind of redundancy is required for connectivity. Redundancy includes edge router redundancy, link redundancy, and service provider redundancy.

When assessing these parameters, whether an enterprise network needs to be independent of the selected service providers must be taken into account. If independency is required, the public IP address space should not be used from the service provider public address space, but should instead be acquired from the regional Internet authority. Similarly, independency is required for the enterprise network AS. The AS number must be a public assigned number, and not from the private AS number pool that can be assigned by the service provider.

https://t.me/learningnets

Routing Method Selection

Connecting an enterprise network to a service provider requires routing information to be exchanged.

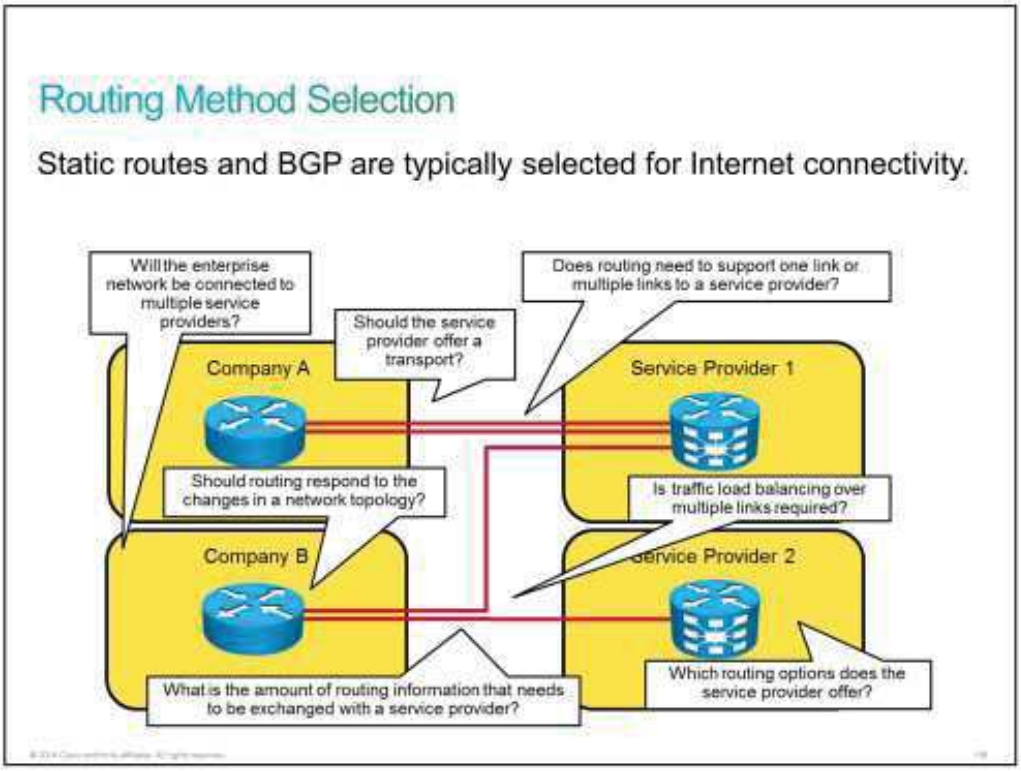

The selection of the routing means depends on the answers to these questions:

- Should the routing methodology respond to the changes in a network topology?
- Does the routing methodology need to support one link, or multiple links, to a service provider?
- Is traffic load balancing over multiple links required?
- Will the enterprise network be connected to multiple service providers?
- Should the service provider offer merely a transport capability that requires connecting different customer locations, perhaps via certain Layer 2 technologies?
- What is the amount of routing information that needs to be exchanged with a service provider?
- Which routing options does the service provider offer?

Routing Method Selection: Static Routes

Static routes are the simplest way of exchanging routing information with a service provider.

https://t.me/learningnets

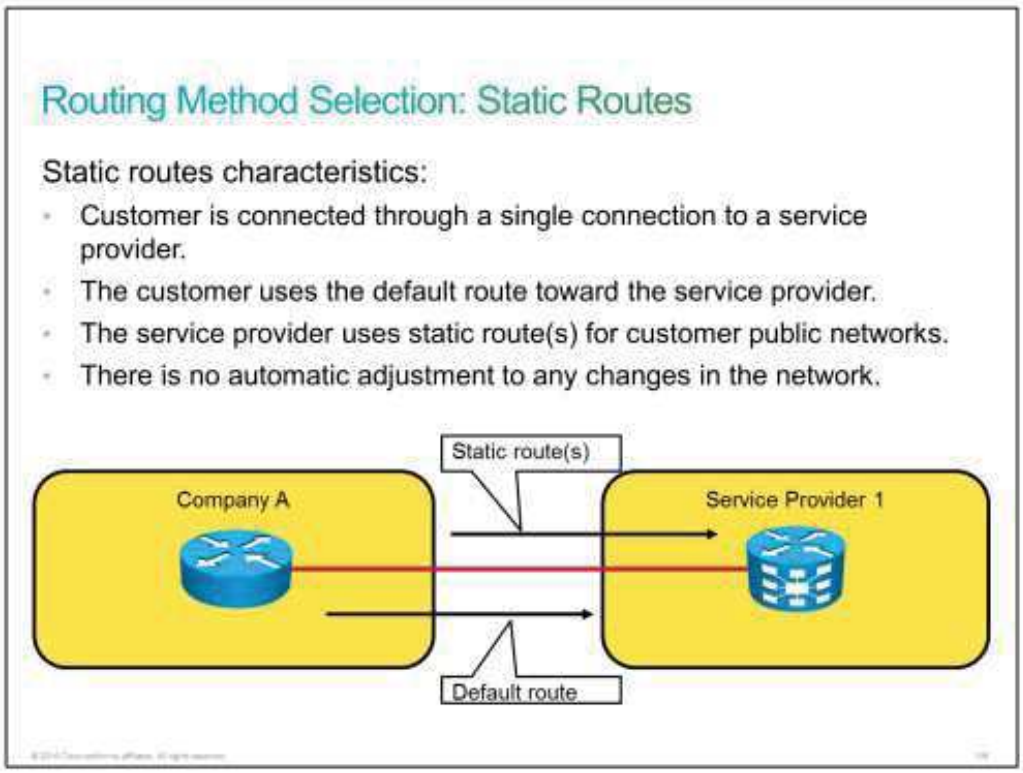

These static routes must agree with the service provider about routing configuration, and they do not change afterward; or, if they do change, these changes should not occur often.

Static routes are typically used when a customer is connected through a single connection to a service provider. The customer can use the default route toward the service provider, and the service provider must deploy a static route or routes that encompass the public networks of the customer. The service provider typically would also redistribute this information into its BGP.

Although static routes are a simple solution, they also have some drawbacks, especially in terms of flexibility and adaptability. For example, if there was a change in a network topology, beyond a directly connected link failure, the static routes would not adapt. If adaptation is required, either the static routes must be combined with an IP SLA functionality that, in turn, would indicate that a static route is down if a certain condition is met, or dynamic routing is used. Note that using an IP SLA cannot completely substitute for dynamic routing, because it cannot react to all changes in the topology in the Internet.

Routing Method Selection: BGP

A typical option that is used to provide dynamic routing exchange when Internet connectivity is deployed is BGP.

https://t.me/learningnets

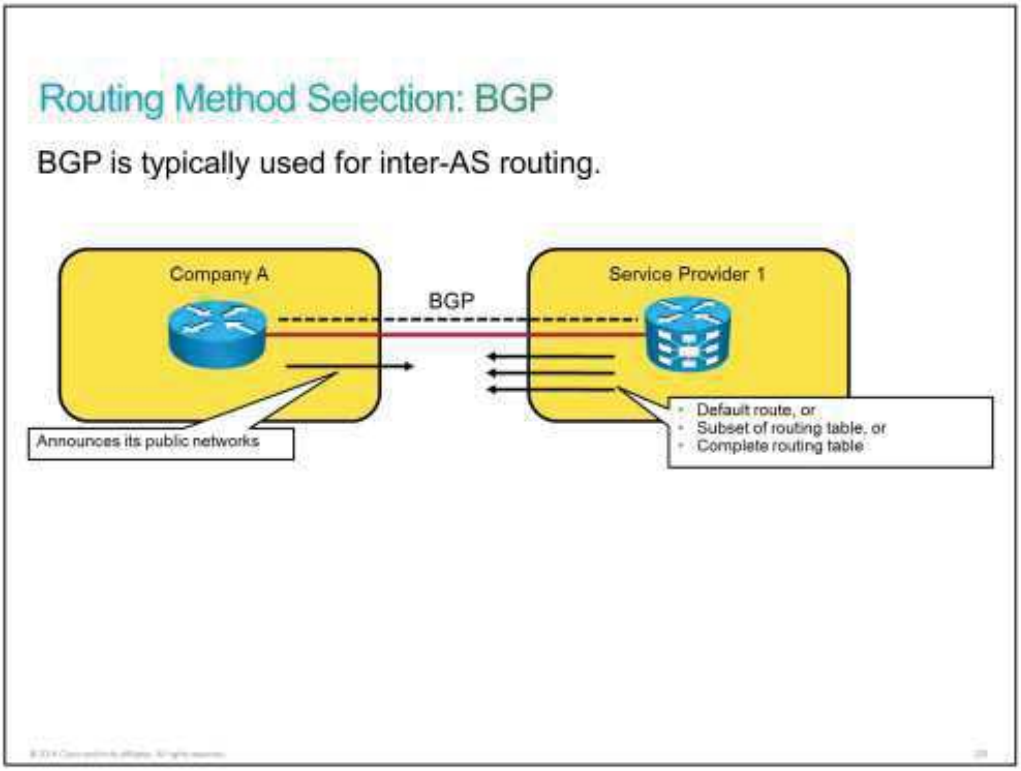

BGP dynamically exchanges routing information, and thus reacts to topology changes, including those changes beyond a link failure.

From the routing perspective, three options can be used:

- A service provider announces a default route only. If two service providers are used, one service provider will be the primary and the second service provider will serve as the backup.
- A service provider announces a default route and a subset of Internet routes, typically from its own public address space. This results in the shortest path being used to the directly connected public networks. For public networks that are not directly connected, one of the service providers is the primary and the second service provider is the backup.
- A service provider announces a complete Internet routing table. The shortest path to any destination will be chosen according to the routing table.

https://t.me/learningnets

Connection Options

When connecting an enterprise network to a service provider, redundancy is a serious concern.

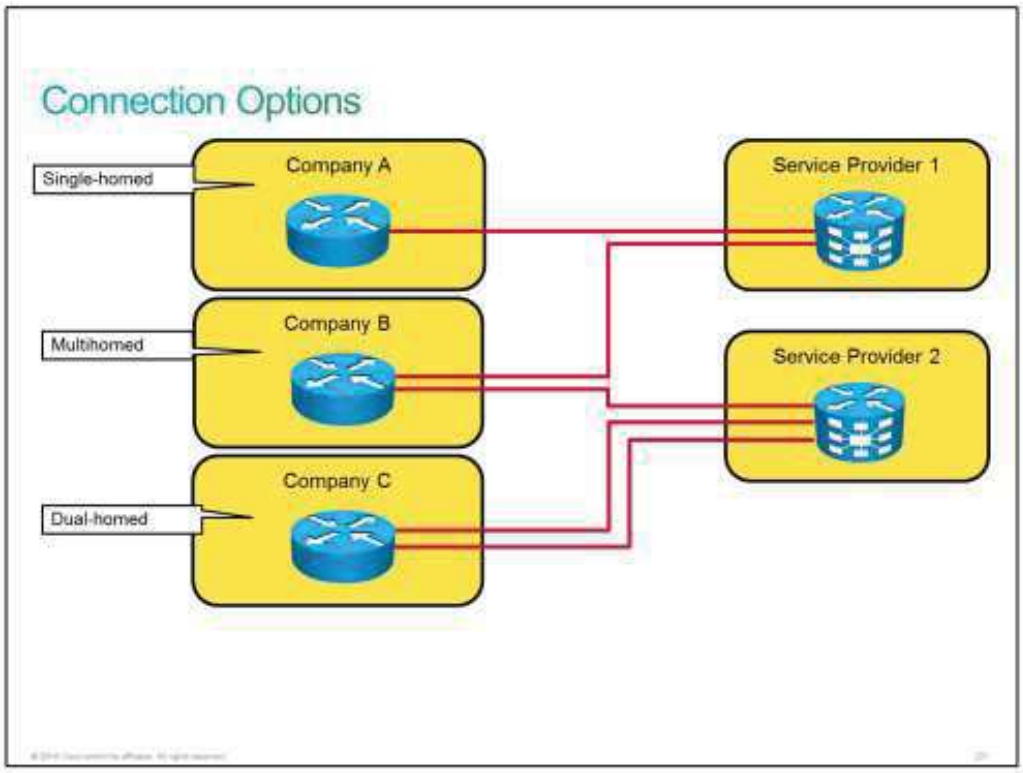

Different aspects can be addressed to achieve connectivity redundancy:

- Deployment of redundant links
- Deployment of redundant equipment
- Use of redundancy within a single router

From the connection perspective, a customer can be connected to a single service provider or to multiple service providers. With a single service provider connection, redundancy can still be achieved if two links toward the same service provider are used effectively, making a customer dual-homed. When no link redundancy is used, the customer is simply single-homed, and in case of a failure within a service provider network, connectivity to the Internet is interrupted.

With multiple service provider connections, redundancy is built into the design, because the customer is multihomed and thus is resistant to a single service provider failure. To enhance the resiliency further, a customer can have two links toward a single service provider, making the solution dual-multihomed.

https://t.me/learningnets

Single-Homed Enterprises

Single-homed service provider connectivity is used in cases when a loss in Internet connectivity or a link failure is not vital or highly problematic to a customer.

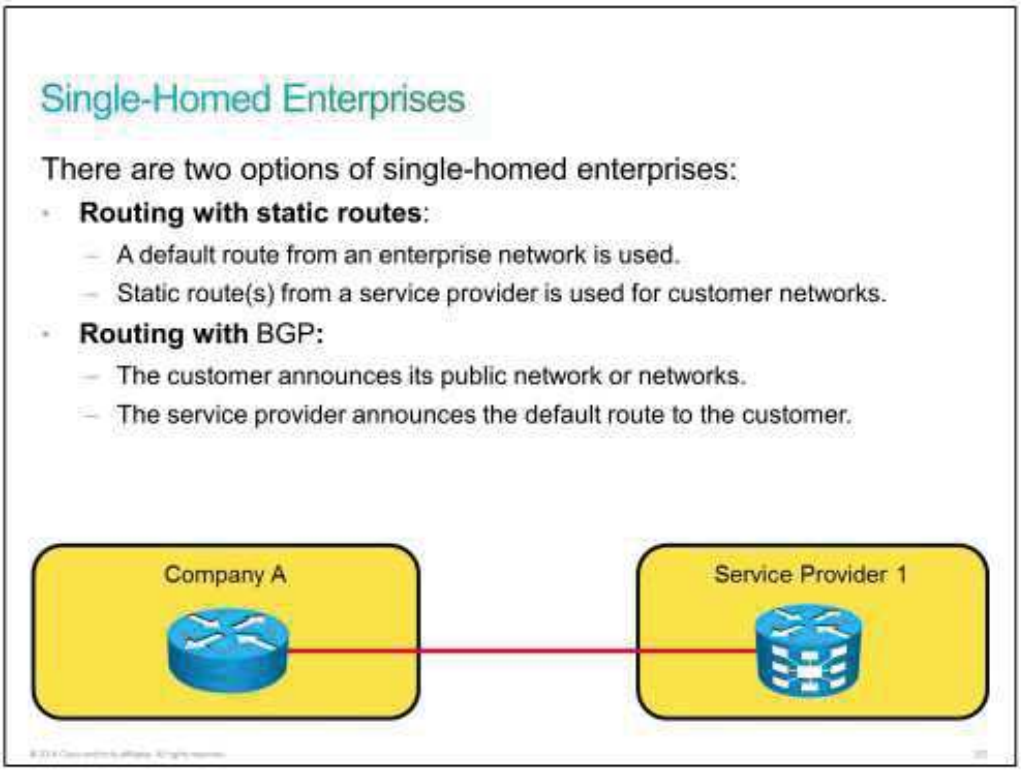

Internet access to a single service provider does not require BGP. Static routes are typically used to manage the routing. If BGP is used, the customer uses it to dynamically announce its public networks to a service provider; the service provider announces only a default route to the customer, because that is sufficient to provide connectivity through a single link to the Internet.

https://t.me/learningnets

Dual-Homed Enterprises

When a customer is connected to a single service provider only, resiliency can still be achieved by deploying a second link to the same service provider. With a second link being used, the routing must be properly configured to allow usage of such a link.

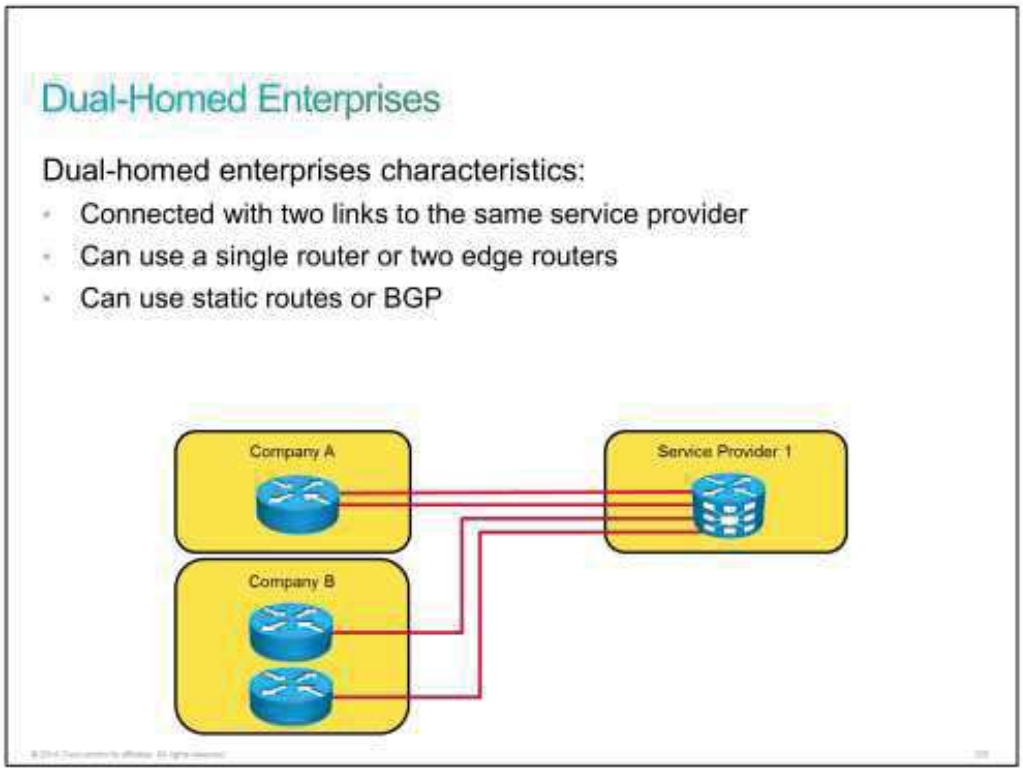

Depending on the SLA that is signed with the service provider, the routing that is deployed could achieve these benefits:

- Primary and backup link functionality, where a single primary link is used to forward and receive traffic to and from the service provider, and the secondary link is used only when the first link fails
- Load sharing between the links, achieved with Cisco Express Forwarding switching

In both cases, the routing that is used can be based on either static or dynamic routes that would typically include BGP.

To enhance the resiliency further, the two links can terminate at separate customer routers.

https://t.me/learningnets

Multihomed Enterprises

Robust resiliency is achieved by connecting to two or more different service providers.

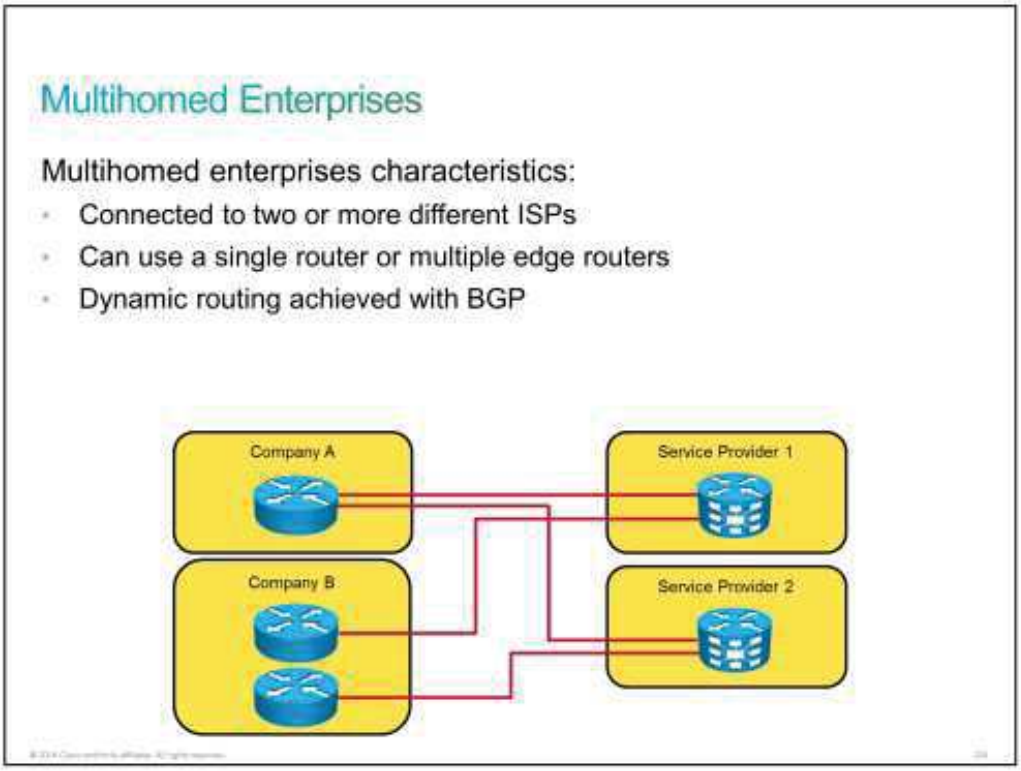

The benefits of connecting to two or more different service providers are significant:

- Greater resistance to a failure than with a directly connected link to a single service provider
- Sharing the load for different destination networks between service providers, based on the network proximity
- Scaling the solution beyond two service providers
- Achieving a solution independent of a service provider. A service provider change requires an update of routing configuration and change of a link. The public IP address space that is used remains the same.

Connections from different service providers can terminate on the same router or different routers to further enhance the resiliency. The routing methodology must be capable of reacting to dynamic changes, and BGP is typically used to achieve this flexibility.

https://t.me/learningnets

Dual-Multihomed Enterprises

Multihoming exists when an AS has more than one connection to the Internet.

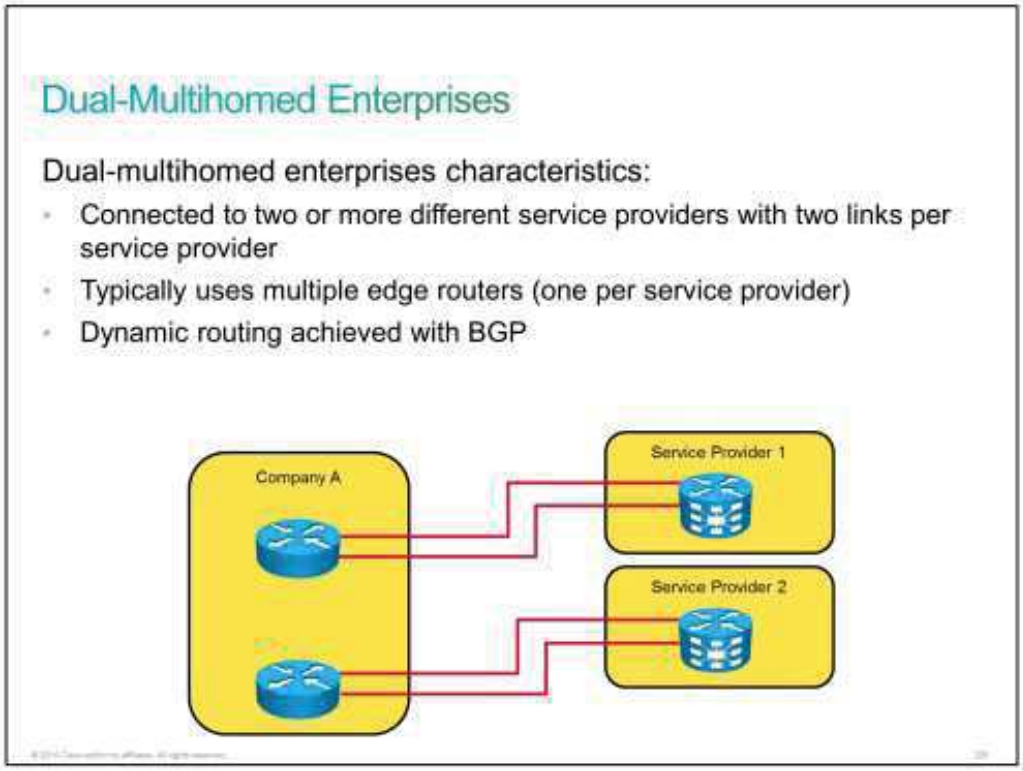

There are two typical reasons for multihoming:

- To increase the reliability of the connection to the Internet. If one connection fails, the other connection remains available.
- To increase the performance of the connection. Paths for better performance can be used to certain destinations.

The benefits of BGP are apparent when an AS has multiple EBGP connections to either a single AS or multiple autonomous systems. Having multiple connections allows an organization to have redundant connections to the Internet, so that if a single path becomes unavailable, connectivity can still be maintained.

https://t.me/learningnets

Summary

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

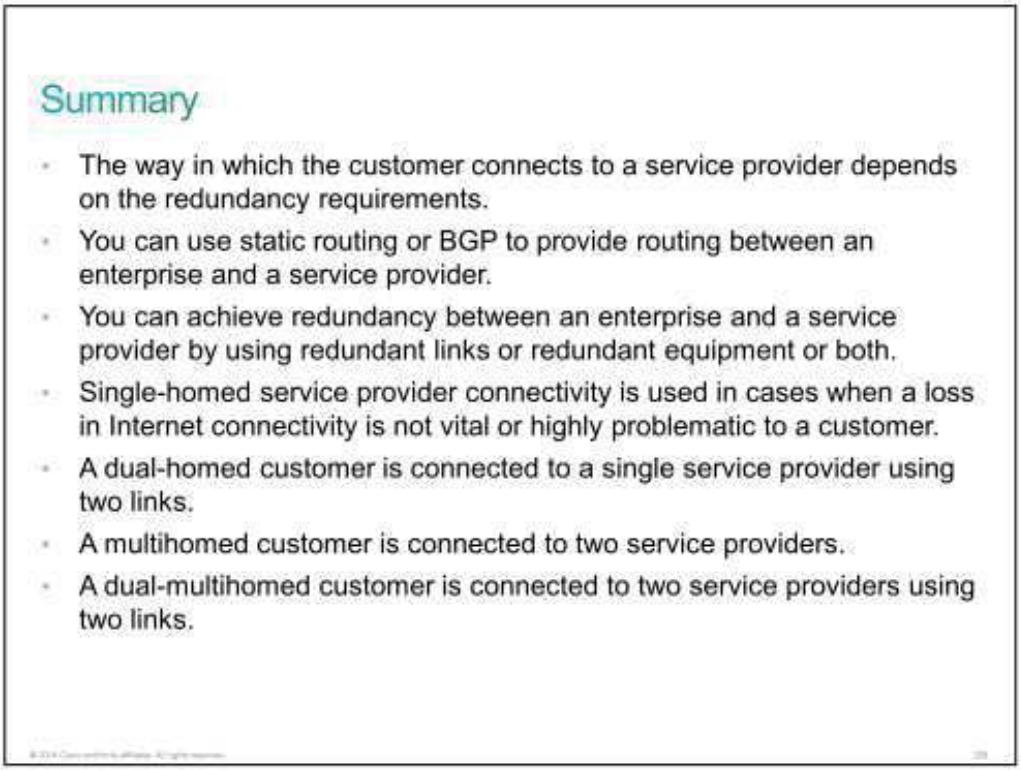

https://t.me/learningnets

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 2: Introducing BGP Routing

Overview

To properly plan enterprise connectivity to a service provider, the designer must understand many aspects of BGP, including the session origination point, the prerequisites for successfully implementing such connectivity, the available routing options, and which routing option should be used in certain cases. The connectivity redundancy considerations must also be taken into account.

BGP is an EGP and this lesson introduces BGP routing. BGP routes between autonomous systems by using path vector attributes upon which routing policy decisions at the AS level are enforced.

Upon completing this lesson, you will be able to:

- Describe BGP terminology
- Describe autonomous systems in BGP networks
- Describe BGP routing between autonomous systems
- Describe BGP path vectors
- Describe BGP routing policies
- Describe BGP features
- Describe the tables used by a router to store BGP information
- Describe the four BGP message types
- Describe the multiprotocol extensions added to BGP to support IPv6
- Describe MPLS VPNs

https://t.me/learningnets

BGP Terminology

The Internet is a collection of autonomous systems that are interconnected to allow communication among these autonomous systems. BGP provides the routing between these autonomous systems.

BGP Terminology

Terms used with BGP:

- **Autonomous system**: A collection of networks under a single administrative domain
- **Interdomain routing:** Routing between the customer and the service provider
- **Internal routing:** Uses IGP (RIP, OSPF, IS-IS, etc.) to exchange routing information inside the AS
- **External routing:** Uses EGP (BGP) to exchange routes between autonomous systems
- Two BGP implementations:
  - IBGP: When BGP is used inside an AS
  - EBGP: When BGP is used between autonomous systems

BGP Terminology (Cont.)

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Enterprises that want to connect to the Internet do so through one or more service providers. If your organization has only one connection to one service provider, then you probably do not need to use BGP; instead, you would use a default route. However, if you have multiple connections to one or multiple service providers, then BGP might be appropriate, because it allows manipulation of path attributes so that the optimal path can be selected.

One way to categorize routing protocols is based on whether the protocols are interior or exterior:

- **IGP** is a routing protocol that exchanges routing information within an AS. RIP, OSPF, and IS-IS are examples of IGPs.
- **EGP** is a routing protocol that exchanges routing information between different autonomous systems. BGP is an example of an EGP.

BGP is an IDRP, also known as an EGP. BGP4 is defined in RFC 4271. As noted in this RFC, the classic definition of an AS is "a set of routers under a single technical administration, using an IGP and common metrics to route packets within the AS, and using an interautonomous system routing protocol (also called an EGP) to determine how to route packets to other autonomous systems."

Autonomous systems can use more than one IGP, potentially with several sets of metrics. From the point of view of BGP, the most important characteristic of an AS is that it appears to other autonomous systems to have a single, coherent interior routing plan and presents a consistent picture of reachable destinations. All parts of an AS must connect to each other.

When BGP is running between routers in different autonomous systems, it is called EBGP. When BGP is running between routers in the same AS, it is called IBGP. BGP allows the path that packets take to be manipulated by the AS. Understanding how BGP works can help to avoid creating problems that result from running BGP for the AS .

https://t.me/learningnets

Autonomous System and BGP

An AS is a collection of networks under a single technical administration. Normally IGPs operate within an AS, while BGP is used between autonomous systems on the Internet.

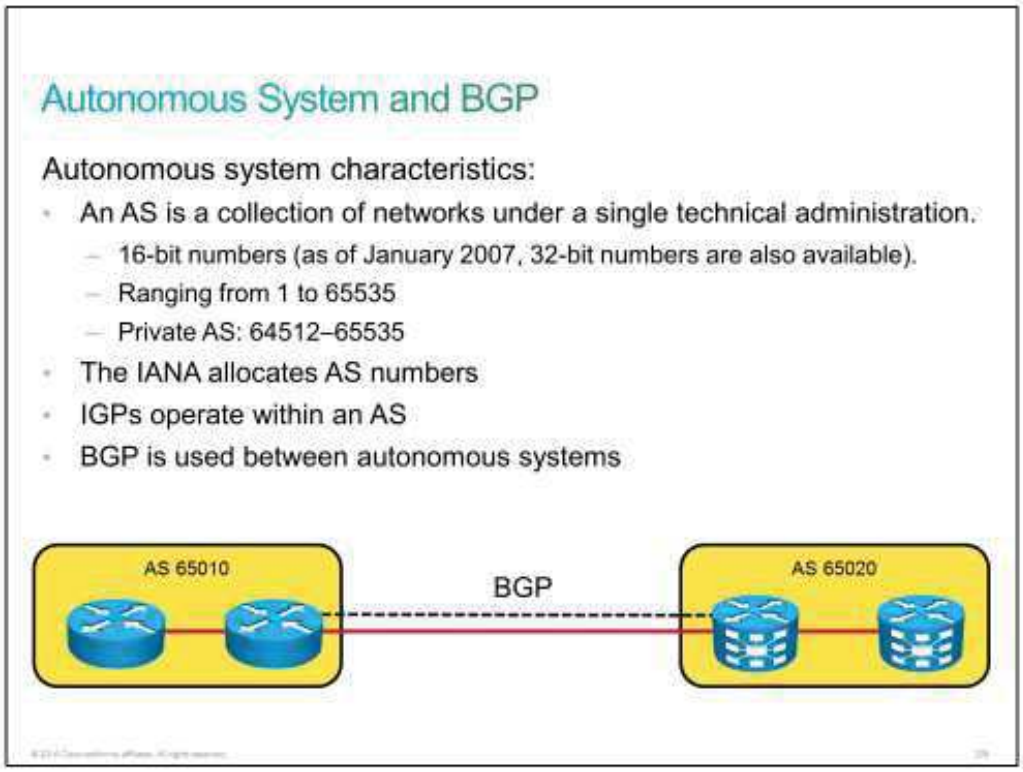

The IANA is the organization that is responsible for allocating AS numbers. Specifically, the ARIN has the jurisdiction to assign numbers for the Americas, the Caribbean, and Africa. RIPE NIC administers AS numbers for Europe, and the APNIC administers the numbers for the Asia Pacific region.

AS numbers traditionally were 16-bit numbers ranging from 1 to 65535. RFC 1930 provides guidelines for the use of AS numbers. A range of AS numbers, 64512 through 65535, is reserved for private use, much like private IP addresses.

As the 16-bit AS range was nearing exhaustion, the IANA extended the AS number registry to a 32-bit range (0 to 4,294,967,295) in 2007. Most networks can still use the "old" AS 16-bit range. Routers that are using the new range need their code to be updated to extend the AS number field size to 32 bits. Routers that are using the old system still use the 16-bit AS number field. When a router that is using the new AS range communicates with a router that is using the old AS range, it uses the reserved AS transition number 23456 in 16-bit form.

For readability purposes, the 32-bit AS range is written in two parts. The old AS ranges from 0.1 to 0.65535; the new AS ranges from 1.0 to 65535.65535.

AS numbers from 65512 to 65535 are still reserved for private usage in the new 16-bit system.

https://t.me/learningnets

BGP Routing Between Autonomous Systems

The main goal of BGP is to provide an interdomain routing system that guarantees a loop-free exchange of routing information between autonomous systems.

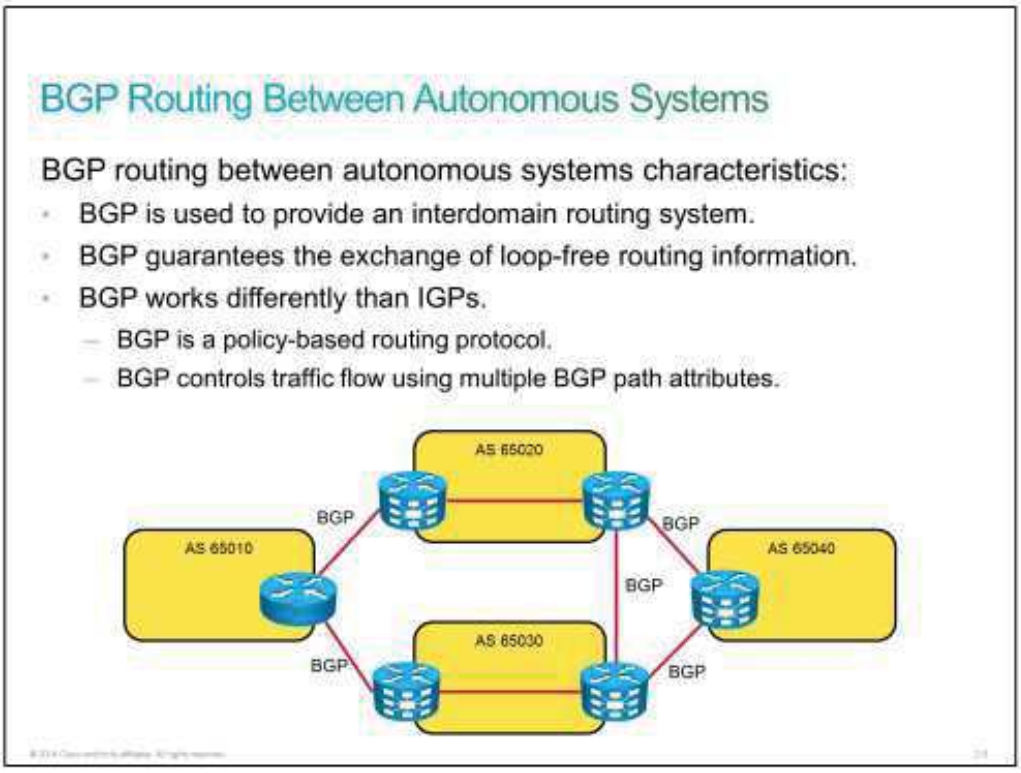

Routers exchange information about paths to destination networks. BGP is a successor of EGP, which was developed to isolate networks from each other as the Internet grew.

BGP works differently than IGPs. An internal routing protocol looks for the quickest path from one point in a network to another, based on certain metrics. In the internal routing protocol, the next hop is the next router; in the BGP, the next hop is the next AS. Rather than treating a router as a single point in the path to any given destination, BGP treats each AS as a single point on the path to the destination. The primary reason that BGP treats an entire autonomous system as a single hop in the AS path is to hide topological details of the AS. No AS can tell what the path through another AS looks like, only that the destination is reachable through that AS. BGP is a PBR protocol that allows an AS to control traffic flow using multiple BGP path attributes. BGP allows a provider to use all its bandwidth by manipulating these path attributes.

https://t.me/learningnets

BGP Path Vector Functionality

Internal routing protocols announce a list of networks and the metrics to get to each network. In contrast, BGP routers exchange network reachability information, called path vectors, which are made up of path attributes (like metrics).

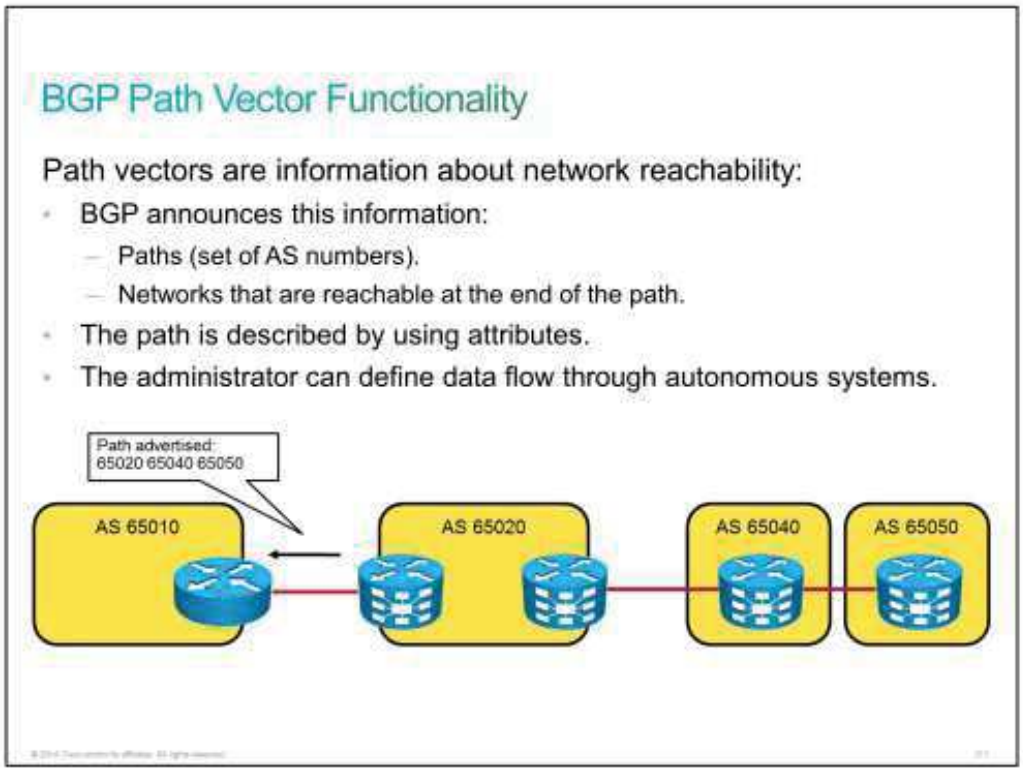

The path vector information includes a list of the complete path of BGP AS numbers (hop-by-hop) that are necessary to reach a destination network and the networks that are reachable at the end of the path. Other attributes include the IP address to get to the next AS (the next-hop attribute), and an indication of how the networks at the end of the path were introduced into BGP (the origin code attribute).

This AS path information is useful to construct a graph of loop-free autonomous systems, and is used to identify routing policies so that restrictions on routing behavior can be enforced, based on the AS path.

The AS path is always loop-free. A router that is running BGP does not accept a routing update that already includes the router AS number in the path list, because the update has already passed through its AS, and accepting it again would result in a routing loop.

An administrator can define policies or rules about how data will flow through the autonomous systems.

https://t.me/learningnets

BGP Routing Policies

BGP allows routing policy decisions at the AS level to be enforced.

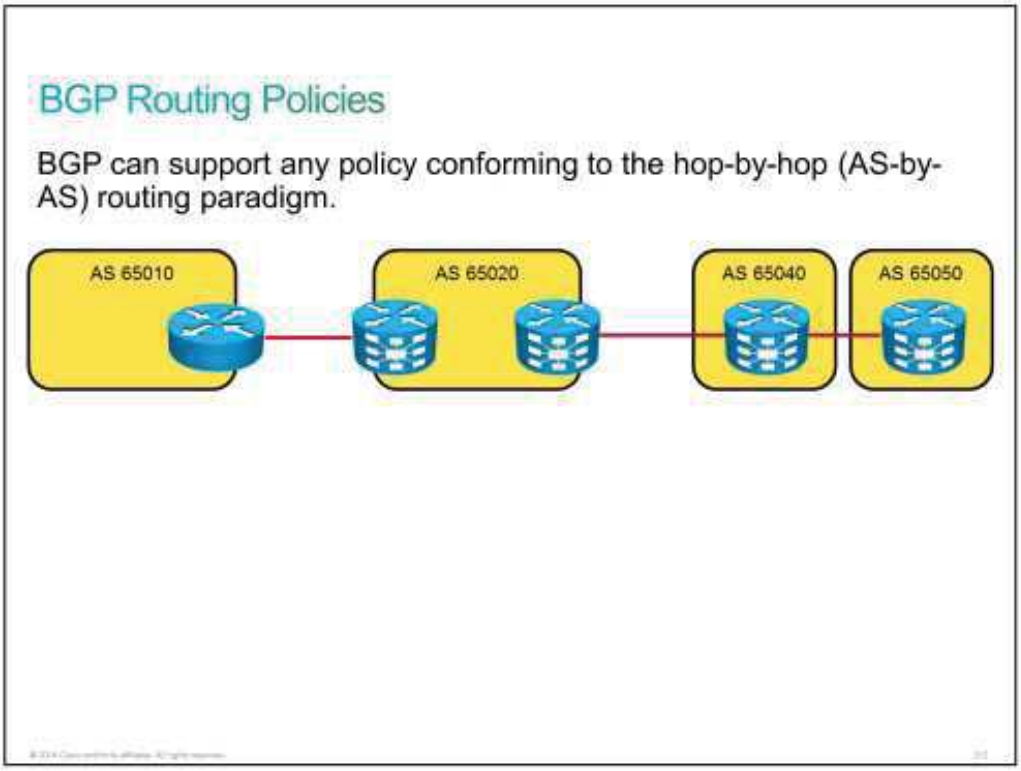

Policies can be implemented for all networks that are owned by an AS, for a certain CIDR block of network numbers (prefixes), or for individual networks or subnetworks.

BGP specifies that a BGP router can advertise to neighboring autonomous systems only those routes that it actually uses. This rule reflects the hop-by-hop routing paradigm that the Internet generally uses.

The hop-by-hop routing paradigm does not support all possible policies. For example, BGP does not enable one AS to send traffic to a neighboring AS, intending that the traffic takes a different route from that taken by traffic that originates in that neighboring AS. In other words, the way that a neighboring AS routes traffic cannot be influenced, but the way that traffic gets to a neighboring AS can be influenced. However, BGP supports any policy that conforms to the hop-by-hop routing paradigm.

Because the Internet currently uses the hop-by-hop routing paradigm only, and because BGP can support any policy that conforms to that paradigm, BGP is highly applicable as an inter-AS routing protocol.

https://t.me/learningnets

BGP Features

BGP is often classified as an advanced distance vector protocol, or even more accurately, as a path vector protocol.

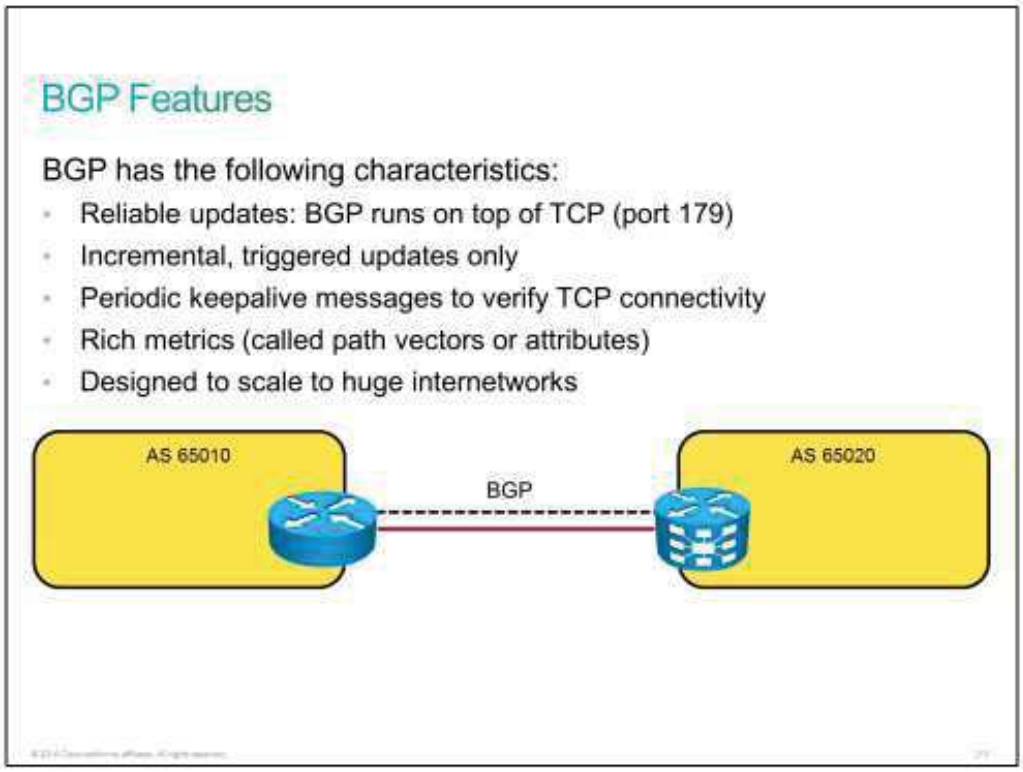

BGP uses TCP as its transport protocol, which provides reliable connection-oriented delivery. BGP assumes that its communication is reliable; therefore, it does not have to implement retransmission or error-recovery mechanisms. BGP uses TCP port 179.

A router that runs the BGP protocol is called a BGP "Speaker." Two routers that are using BGP form a TCP connection with one another and exchange messages to open and confirm the connection parameters. These two BGP routers are called "peer routers," or "neighbors." No routing information is exchanged until the TCP connection has been established.

After the connection is made, BGP peers exchange complete routing tables. However, because the connection is reliable, BGP peers send only changes (incremental, or triggered, updates) after that. Reliable links do not require periodic routing updates; therefore, routers use triggered updates instead. BGP sends keepalive messages.

https://t.me/learningnets

BGP Databases

A router that is running BGP keeps its own tables to store BGP information that it receives from and sends to other routers.

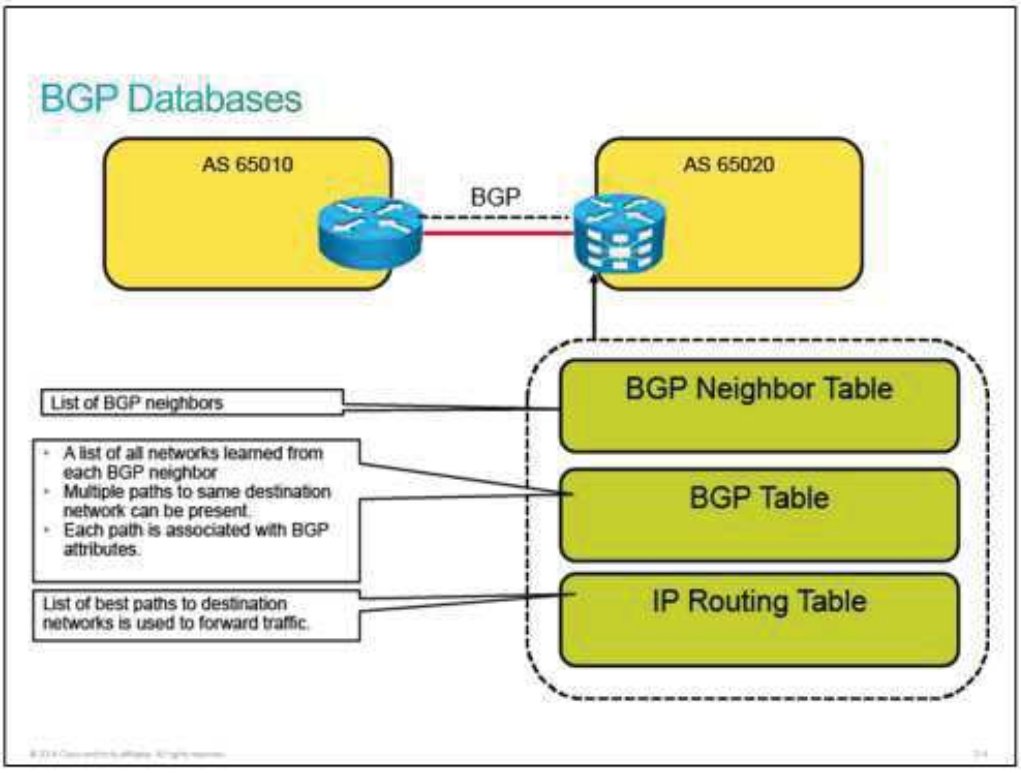

There are three different tables:

- Neighbor table
- BGP table (also called a forwarding database or topology database)
- IP routing table

For BGP to establish an adjacency, it must be explicitly configured for each neighbor. BGP forms a TCP relationship with each of the configured neighbors and keeps track of the state of these relationships by periodically sending a BGP/TCP keepalive message. The BGP sends BGP/TCP keepalives by default every 60 seconds.

After establishing an adjacency, the neighbors exchange the BGP routes that are in their IP routing table. Each router collects these routes from each neighbor that successfully establishes an adjacency, and then places the routes in its BGP forwarding database. All routes that have been learned from each neighbor are placed into the BGP forwarding database. The best routes for each network are selected from the BGP forwarding database, using the BGP route selection process, and are then offered to the IP routing table.

Each router compares the offered BGP routes to any other possible paths to those networks, and the best route—based on administrative distance—is installed in the IP routing table.

EBGP routes (BGP routes that are learned from an external AS) have an administrative distance of 20. IBGP routes (BGP routes that are learned from within the AS) have an administrative distance of 200.

https://t.me/learningnets

BGP Message Types

There are four types of BGP messages: open, keepalive, update, and notification.

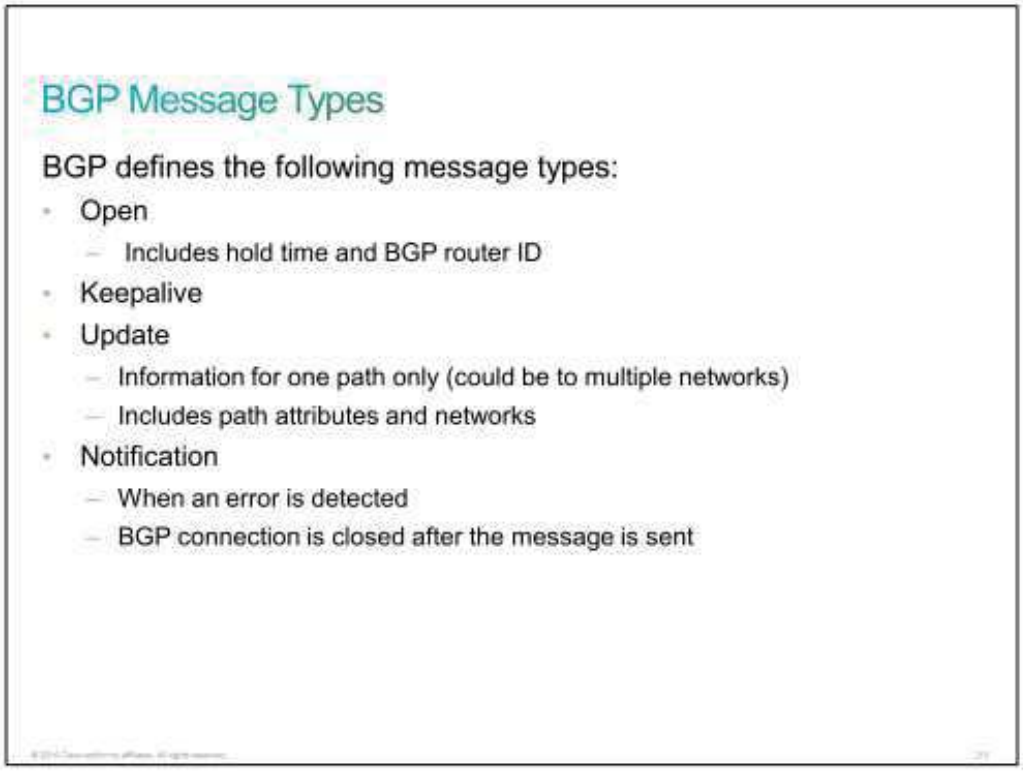

After a TCP connection is established, the first message that is sent by each side is an open message. If the open message is acceptable, the side that receives the message sends a keepalive message confirming the open message. After the receiving side confirms the open message and establishes the BGP connection, the BGP peers can exchange any update, keepalive, and notification messages.

BGP peers initially exchange their full BGP routing tables. Incremental updates are sent only after topology changes in the network occur. BGP peers send keepalive messages to ensure that the connection between the BGP peers still exists, and they send notification packets in response to errors or special conditions.

Here are further details about the different types of BGP messages:

- **Open message**: An open message includes the following information:
  - **Version number:** The suggested version number. The highest common version that both routers support is used. Most BGP implementations today use BGP4.
  - **AS number:** The AS number of the local router. The peer router verifies this information. If it is not the AS number that is expected, the BGP session is ended.
  - **Hold time:** Maximum number of seconds that can elapse between the successive keepalive and update messages from the sender. On receipt of an open message, the router calculates the value of the hold timer by using whichever is smaller: its configured hold time or the hold time that was received in the open message.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

- **BGP router ID:** This 32-bit field indicates the BGP ID of the sender. The BGP ID is an IP address that is assigned to that router, and it is determined at startup. The BGP router ID is chosen in the same way that the OSPF router ID is chosen—it is the highest active IP address on the router, unless a loopback interface with an IP address exists. In this case, the router ID is the highest loopback IP address. The router ID can also be statically configured.
- **Optional parameters:** These parameters are TLV-encoded. An example of an optional parameter is session authentication.

- **Keepalive message:** BGP keepalive messages are exchanged between BGP peers frequently enough to keep the hold timer from expiring. If the negotiated holdtime interval is 0, then periodic keepalive messages are not sent. A keepalive message consists of only a message header.
- **Update message:** A BGP update message has information on one path only; multiple paths require multiple update messages. All the attributes in the update message refer to that path, and the networks are those that can be reached through that path. An update message can include the following fields:
  - **Withdrawn routes:** This list displays IP address prefixes for routes that are withdrawn from service, if any.
  - **Path attributes:** These attributes include the AS path, origin, local preference, and so on (as described later in this module). Each path attribute includes the attribute TLV. The attribute type consists of the attribute flags, followed by the attribute type code.
  - **Network-layer reachability information:** This field contains a list of IP address prefixes that are reachable by this path.
- **Notification message:** A BGP notification message is sent when an error condition is detected; the BGP connection is closed immediately after this is sent. Notification messages include an error code, an error subcode, and data that is related to the error.

https://t.me/learningnets

Multiprotocol Extensions for BGP4

Original BGP4 (RFC 1771) carries IPv4-specific information. Multiprotocol BGP (RFC 2283) uses extensions to add capability to transport routing information for other protocols.

Multiprotocol Extensions for BGP4

Multiprotocol extensions for BGP4 characteristics:

- BGP originally designed for IPv4:
  - Carries IPv4 prefix reachability information
  - Uses IPv4 for transport
- Multiprotocol extensions for BGP4:
  - Enables other protocols besides IPv4
  - New identifier for the address family
  - Most often used in MPLS networks for MPLS VPN
- IPv6-specific extensions:
  - Scoped addresses: NEXT_HOP contains a global IPv6 address and potentially a link-local address
  - NEXT_HOP and NLRI are expressed as IPv6 addresses and prefixes in the multiprotocol attributes

Multiprotocol Extensions for BGP4 (Cont.)

- Still uses TCP for transport:
  - TCP can run over IPv4, transporting IPv6 information
  - TCP can run natively over IPv6

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Multiprotocol extensions for BGP are defined as new attributes. These new attributes define NLRI and a next hop (the next router in the path to the destination) where IPv6 addresses can be used.

BGP4, with multiprotocol extensions, enables the use of many address families. Address families define the type of addresses being carried. The most common address families are IPv4, IPv6, and VPNv4 and VPNv6 for MPLS VPN routes.

An address family is activated within BGP using the Cisco IOS or IOS XE or IOS XR **address-family** router BGP command.

Two new attributes support MP-BGP:

- Multiprotocol reachable NLRI (MP_REACH_NLRI).
- Multiprotocol unreachable NLRI (MP_UNREACH_NLRI).

The MP_REACH_NLRI attribute describes reachable destinations and includes attributes that contain information about network-layer protocol prefixes (specifically IPv6), and the next destination hop-to-reach prefixes.

The IPv6-specific extensions for MP-BGPv4 are listed here:

- **Scoped addresses:** NEXT_HOP contains a global IPv6 address, or potentially a link-local address, when there is link-local reachability with the peer. When link-local addresses are used for peering with a neighbor BGP router, these link-local IPv6 addresses are used as next-hop IP addresses for the routes that are carried by BGP. In the majority of cases, the next-hop IPv6 addresses need to be changed to global IPv6 by attaching a route map or route policy to the neighbor.
- **IPv6 address format:** NEXT_HOP and NLRI are expressed as IPv6 addresses and prefixes in the multiprotocol attributes. You can run MP-BGP over IPv4 or IPv6 transport and can exchange routes for IPv4, IPv6, or both. BGP uses the TCP protocol for peering and this has no relevance to the routes that are carried inside the BGP exchanges. Either IPv4 or IPv6 can be used to transport a TCP connection on the network layer.

https://t.me/learningnets

MPLS VPN Overview

MPLS VPNs are used when a customer has multiple locations that need to be interconnected through a service provider and the customer does not want to use expensive Layer 2 technologies, such as dedicated leased lines.

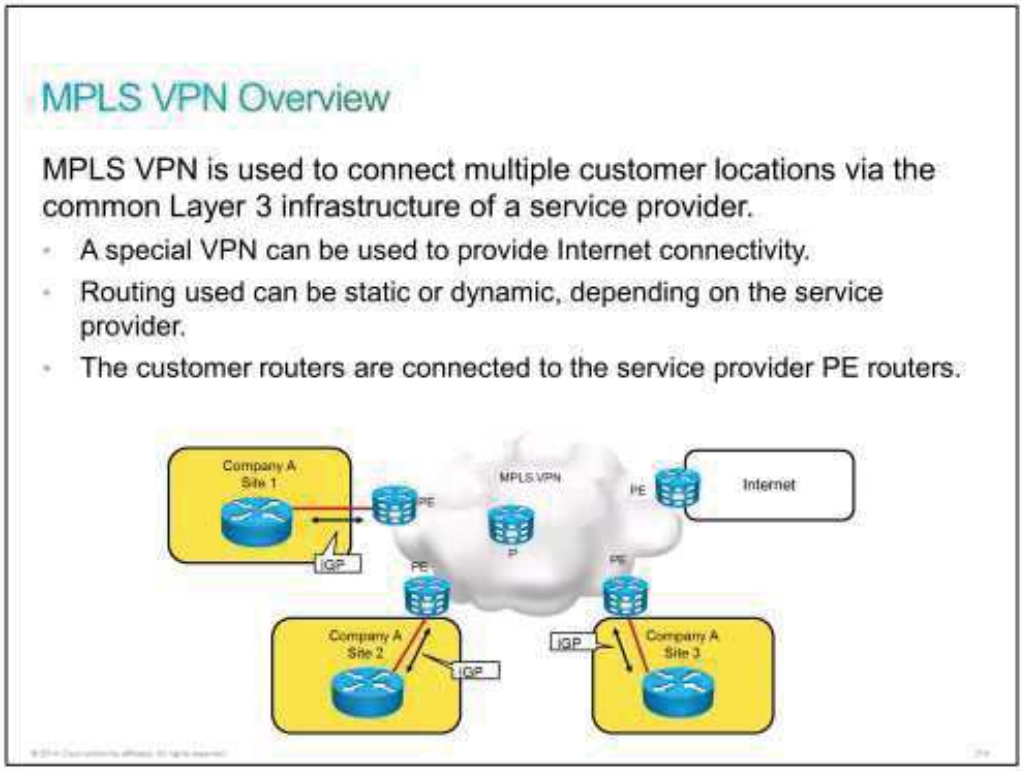

With MPLS VPNs, the service provider uses a common IP-based core network that is enhanced with MPLS technology to provide secure and manageable connectivity for different customers to their geographically diverse sites. In this way, the traffic from different customers can share the same physical wire, but at the same time, traffic is tagged with the labels in such way that the traffic cannot intermix.

When a customer uses MPLS VPN functionality, the routing between the customer and the service provider is required to provide connectivity between the customer locations. Routing options range from static to dynamic and include RIP, EIGRP, OSPF, IS-IS, or even BGP, depending on what the service provider offers. Different locations could use different routing protocols, though this situation typically would not occur.

With an MPLS VPN deployment, the service provider can also offer Internet connectivity through the same MPLS core network, either through a special Internet VPN or through a global routing table. To exchange the Internet routing information, either BGP or a default route is used.

https://t.me/learningnets

Summary

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

Summary

- BGP is the external routing protocol used between autonomous systems.
- An autonomous system is a collection of networks under a single administration and is represented by 16-bit or 32-bit number.
- Forwarding is based on policies and not on best path. BGP routers exchange network reachability information, called path vectors.
- BGP announces a set of AS numbers and networks that are reachable at the end of the path. The path is described by attributes.
- BGP router can advertise to neighboring autonomous systems only those routes that it actually uses.

Summary (Cont.)

- BGP routers establish a TCP session and then exchange routing tables. After that, BGP peers send only incremental triggered updates.
- BGP uses three databases: neighbor table, BGP table, and routing table.
- The four BGP message types are open, keepalive, update, and notification.
- BGP4, with multiprotocol extensions, enables the use of many address families. Address families define the type of addresses being carried.
- MPLS VPNs are used by customers with multiple locations who do not want to use expensive Layer 2 technologies.

https://t.me/learningnets

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 3: Implementing Basic BGP Routing

Overview

This lesson presents the commands and configuration examples that are needed to configure BGP properly. After a successful configuration, BGP will be able to establish a neighbor relationship, set the next-hop address, set the source IP address of a BGP update, and announce networks to other BGP routers.

A router that is running BGP goes through several neighbor states, through which BGP progresses to establish a BGP session. Hints are offered for troubleshooting BGP, because the session can be stuck in the SIA or idle state. EBGP configuration, IBGP configuration, and authentication between neighbors are explained.

BGP is used to perform PBR. To manipulate the best paths that are chosen by BGP, you need to understand the different attributes that BGP uses and how BGP selects the best path that is based on these attributes. This lesson also explains the BGP path selection process and the BGP attributes and their characteristics.

Upon completing this lesson, you will be able to:

- Describe planning for BGP deployments
- Describe the basic configuration steps for EBGP
- Describe how networks are advertised in BGP networks
- Describe the basic configuration steps for IBGP
- Describe full-mesh IBGP networks
- Describe BG support for IPv6
- Describe how to shut down a BGP neighbor
- Describe next-hop behavior in BGP
- Describe the BGP next-hop self feature
- Describe configuration template support in Cisco IOS XR Software

https://t.me/learningnets

- Describe the BGP neighbor states
- Describe BGP neighbor authentication
- Describe clearing the BGP session
- Describe how to monitor BGP routes
- Describe the importance of the BGP path attributes in the path selection
- Describe the decision process of BGP route selection
- Describe the BGP weight attribute
- Describe the BGP local preference attribute
- Describe the BGP AS path attribute
- Describe the BGP multi-exit discriminator attribute

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Planning for BGP

Before initiating BGP configuration, an implementation plan should be prepared.

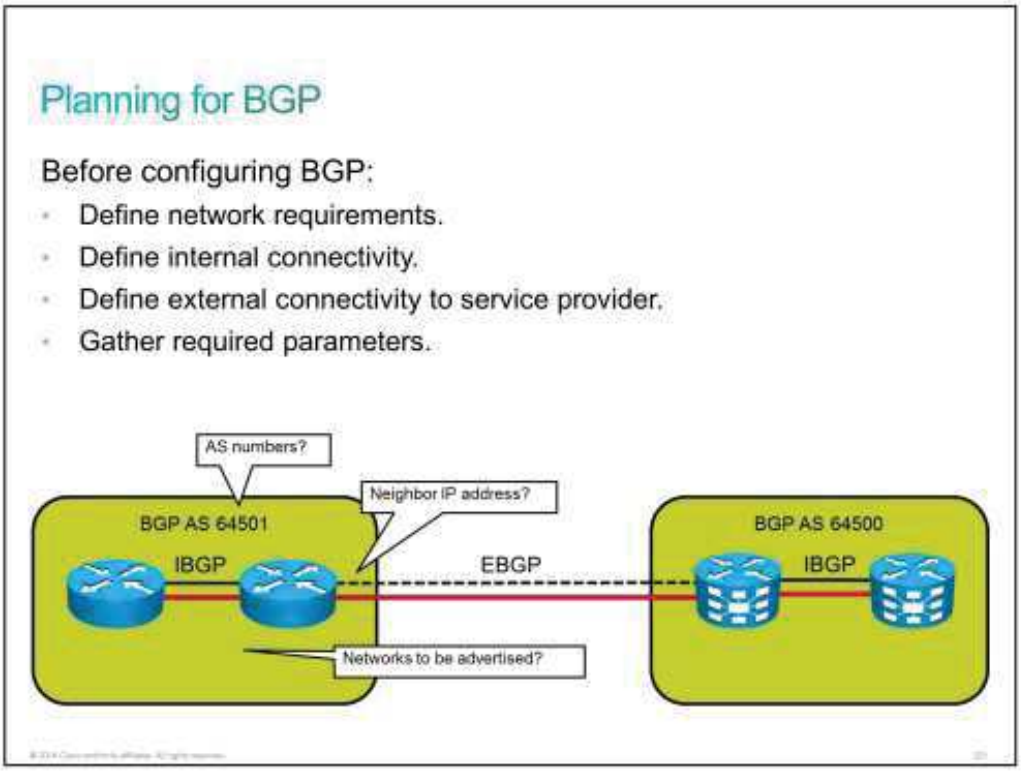

A network administrator must prepare a plane and define the network requirements, including the internal connectivity for IBGP design and configuration, as well as the external connectivity to the service provider for EBGP design and configuration. The next step is to gather all the parameters that are needed to provide enough details for a network operator to start the BGP configuration.

The requirements to configure basic BGP include these details:

- AS numbers (your own, and all remote AS numbers)
- All the neighbors (peers) that are involved in BGP, and IP (IPv4 or IPv6) addressing that is used among the BGP neighbors
- Networks that need to be advertised into BGP

A typical BGP configuration involves configuring BGP between a customer network and a service provider. This process is called EBGP. Many times, IBGP is required, as well as all the collected details for a complete configuration.

https://t.me/learningnets

Configure Basic EBGP

Basic EBGP configuration requires three main steps.

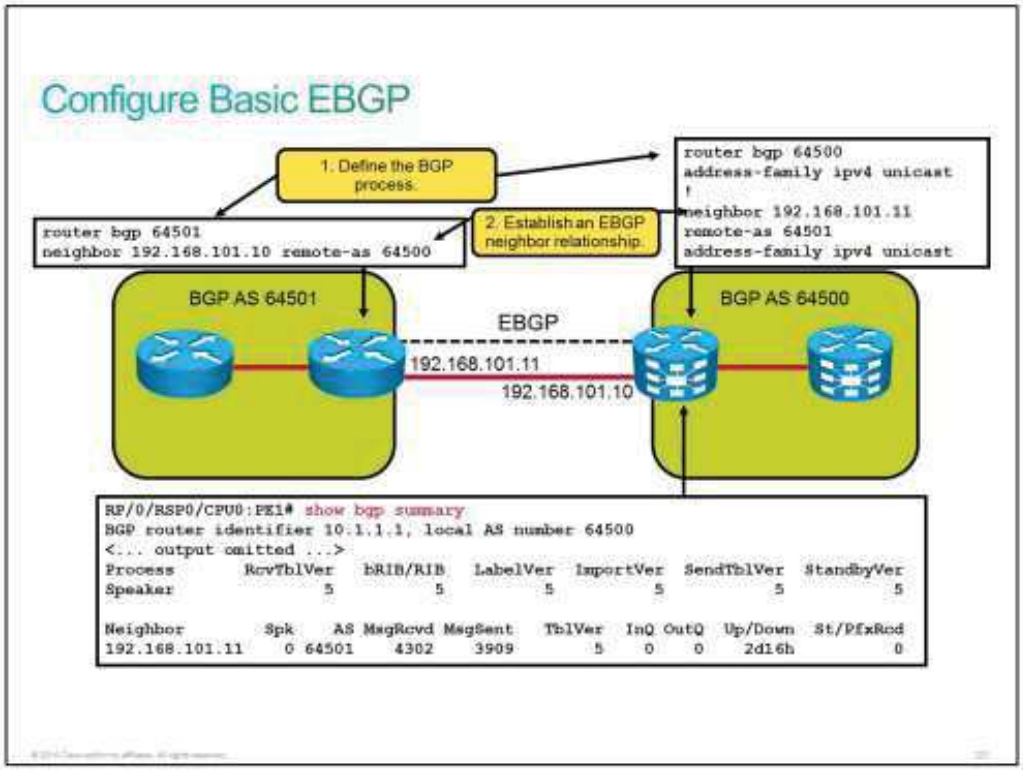

Basic EBGP configuration steps:

- Define the BGP process.
- Establish the neighbor relationship.
- Advertise the networks into BGP.

To perform these steps, information is needed about the neighbors: which AS numbers are used, which IP address is used as the IP address on a remote router (neighbor), and which network will be advertised.

The syntax of the basic BGP configuration commands is similar to the syntax for configuring internal routing protocols. However, there are significant differences in how BGP functions.

Use the Cisco IOS or IOS XE or OS XR **router bgp** *as-number* command to notify the router that any subsequent subcommands belong to this routing process. This command also identifies the local AS in which this router belongs. The router needs to be informed of the AS number so that it can determine whether the BGP neighbors that are to be configured next are IBGP or EBGP neighbors. To establish a connection to another AS, insert the AS number with a Cisco IOS XR **neighbor** router BGP command or Cisco IOS or IOS XE **neighbor remote-as** router BGP command so that the router can properly identify the relationship between the neighboring router and itself.

Only one instance of BGP can be configured on the router at a single time. On the Cisco IOS XR router, enable **address-family** for IPv4 unicast to enable IPv4 unicast capability. The router will start sending and receiving IPv4 routes.

https://t.me/learningnets

The Cisco IOS XR **show bgp summary** or Cisco IOS or IOS XE **show ip bgp summary** command is one way to verify the neighbor relationship. The figure presents the output from a Cisco IOS XR router. Some of the details of this command output are listed:

- **BGP router identifier:** The IP address that all other BGP speakers recognize as representing this router
- **Neighbor:** The IP address that is used in the neighbor statement with which this router has a relationship
- **AS:** The AS number of the listed neighbor
- MsgRcvd: The number of BGP messages that have been received from this neighbor
- MsgSent: The number of BGP messages that are sent to this neighbor
- TblVer: The BGP table version
- InQ: The number of messages waiting to be processed from this neighbor
- OutQ: The number of messages that are queued and waiting to be sent to this neighbor. TCP flow control prevents this router from overwhelming a neighbor with a large update
- **Up/Down:** The length of time that this neighbor has been in the current BGP state (established, active, or idle)
- **State (established, active, idle, open sent, open confirm, or idle [admin]):** The BGP state. A neighbor can be set to administratively shut down (admin state) by using the Cisco IOS XR **shutdown** router BGP command or the Cisco IOS or IOS XE **neighbor shutdown** router BGP command.
- PfxRcd: When the session is in the established state, this value represents the number of BGP network entries that are received from the listed neighbor.

https://t.me/learningnets

Advertising BGP Networks

Two options exist when advertising networks into the BGP.

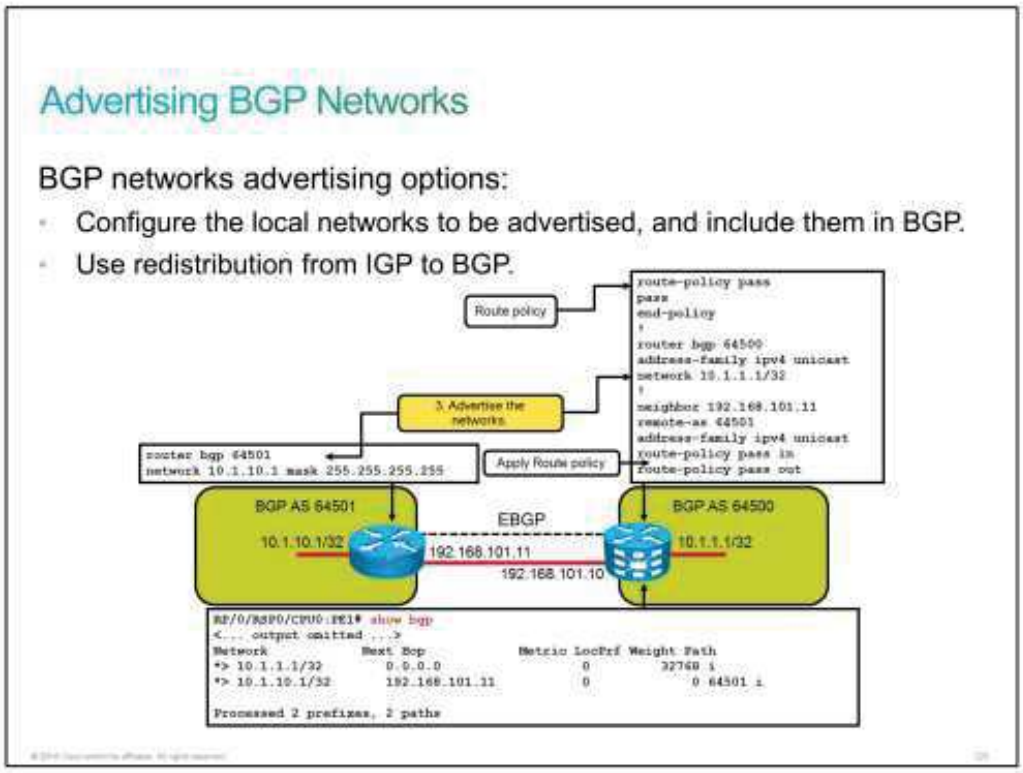

The first option is using the Cisco IOS or IOS XE or IOS XR **network** router BGP command to define the networks that are required. The second option is the redistribution of IGP routes into the BGP routing process.

Use the Cisco IOS or IOS XE or IOS XR **network** router BGP command to permit BGP to advertise a prefix if it is present in the IP routing table. The **network** command determines the networks that the router originates. This concept is different from using the **network** command when configuring IGP. Unlike IGP, the **network** command does not start BGP on specific interfaces; rather, it indicates to BGP which networks it should originate from this router.

The Cisco IOS XR route policy must be applied in order to pass routes in the inbound and outbound directions. Configure the route-policy to pass all routes and use the Cisco IOS XR **route-policy** router BGP command to apply route policy to the inbound and outbound directions.

Use the Cisco IOS XR **show bgp** or Cisco IOS/IOS XE **show ip bgp** command to display the BGP topology database (BGP table). The figure shows a partial sample output of the Cisco IOS XR **show bgp** command; the complete output is shown here:

© 2014 Cisco Systems, Inc.

https://t.me/learningnets

```
RP/0/RSP0/CPU0:PE1# show bgp
Mon Jun 19 22:11:52.515 UTC
BGP router identifier 10.1.1.1, local AS number 64500
BGP generic scan interval 60 secs
BGP table state: Active
Table ID: 0xe0000000   RD version: 7
BGP main routing table version 7
BGP scan interval 60 secs
Status codes: s suppressed, d damped, h history, * valid, > best
              i - internal, r RIB-failure, S stale
Origin codes: i - IGP, e - EGP, ? - incomplete
 Network            Next Hop            Metric LocPrf   Weight Path
*> 10.1.1.1/32        0.0.0.0                  0         32768 i
*> 10.1.10.1/32       192.168.101.11           0             0 64501 i
Processed 2 prefixes, 2 paths
```

The status codes are shown at the beginning of each line of output, and the origin codes are shown at the end of each line. In this output, there is an asterisk (*) in most of the entries in the first column. This asterisk means that the next-hop address is valid. The next-hop address is not always the router that is directly connected to this router. Here are some other options:

- An "**s**," for suppressed, indicates that the specified routes are suppressed (usually because routes have been summarized and only the summary route is being sent).
- A "**d**," for dampening, indicates that the route is being dampened (penalized) for going up and down too often. Although the route might be up right now, it is not advertised until the penalty has expired.
- An "**h**," for history, indicates that the route is unavailable and is probably down; historic information about the route exists, but a best route does not exist.
- An "**r**," for RIB failure, indicates that the route was not installed in the RIB.
- An "**S**," for stale, indicates that the route is stale (this symbol is used in the nonstop forwarding-aware router).

The second column shows ">" when BGP has selected the path as the best path to a network.

The third column is either blank or shows "**i**." If it is blank, BGP learned that route from an external peer. An "**i**" indicates that an IBGP neighbor advertised this path to the router.

The fourth column lists the networks that the router learned.

The Next Hop column lists all the next-hop addresses for each route. This column may contain the entry 0.0.0.0, which signifies that this router is the originator of the route.

The three columns to the left of the Path column list three BGP path attributes that are associated with the path: metric (MED), local preference, and weight.

The column with the Path header may contain a sequence of autonomous systems in the path. From left to right, the first AS that is listed is the adjacent AS, from which this network was learned. The last number (the rightmost AS number) is the originating AS of this network. The AS numbers between these two numbers represent the exact path that a packet takes back to the originating AS. If the path column is blank, the route is from the current AS.

The last column signifies how this route was entered into BGP on the original router. If the last column has an "**i**" in it, the originating router probably used a network statement to introduce this network into BGP. If the character is an "**e**," the originating router learned this network from an EGP, which is the historical predecessor to BGP. A question mark (**?**) signifies that BGP cannot absolutely verify the availability of this network because it is redistributed from an IGP into BGP.

https://t.me/learningnets

Configure Basic IBGP

To establish an IBGP session between routers, use the same neighbor AS number that is configured on the local router.

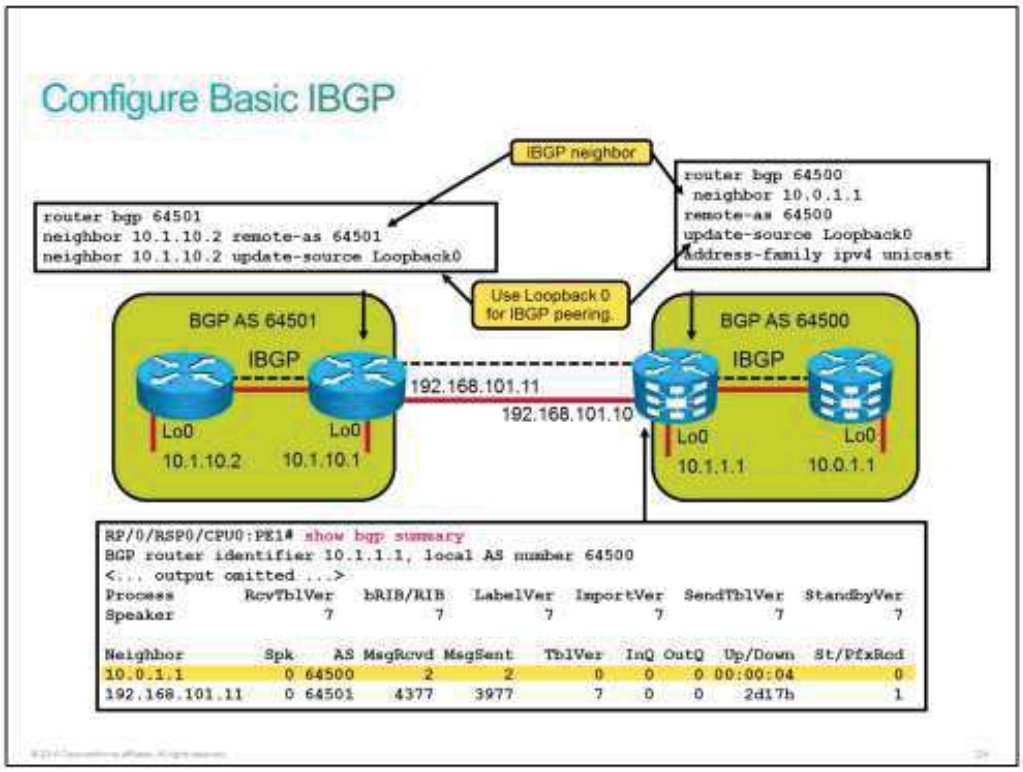

To establish the IBGP session, a loopback interface is used to keep the BGP session stable. As long as the loopback interface is up and reachable, the BGP session will stay up.

The BGP neighbor statement informs the router of the destination IP address for each update packet. The router must decide which IP address to use as the source IP address in the BGP routing update. When a router creates a BGP packet for a neighbor, it checks the routing table for the destination network to reach that neighbor. The IP address of the outbound interface, as the routing table indicates, is used as the source IP address of the BGP packet. When a BGP packet is received for a new BGP session, the source address of the packet is compared to the list of neighbor statements. This source IP address must match the address in the corresponding neighbor statement on the other router. Otherwise, the routers will not be BGP peers because they are not able to establish the BGP session.

Multiple paths can exist to reach each neighbor when peering with IBGP neighboring routers. If the BGP router is using a neighbor address that is assigned to a specific interface on another router, and that interface goes down, the router that is pointing to this address loses its BGP session with that neighbor. If the router peers instead with the loopback interface of the other router, the loopback interface will always be available as long as the router itself does not fail. This peering arrangement adds resiliency to the IBGP sessions because the routers are not tied into a physical interface, which may fail for any number of reasons.

To peer with the loopback of another internal neighbor, the first router would point the neighbor statement to the loopback address of the other internal neighbor. Ensure that both routers have a route to the loopback address of the other neighbor in their routing table. Also ensure that both routers are announcing their loopback addresses into their local routing protocol.

https://t.me/learningnets

The Cisco IOS XR **update-source** router BGP command or the Cisco IOS or IOS XE **neighbor update-source** router BGP command overrides the default source IP address that is used for BGP packets. Telling the router which IP address to use as the source address for all BGP packets is necessary, if a loopback interface is to be used instead of the physical interface.

https://t.me/learningnets

Full-Mesh IBGP

Full-mesh IBGP is required because of the BGP split-horizon rule.

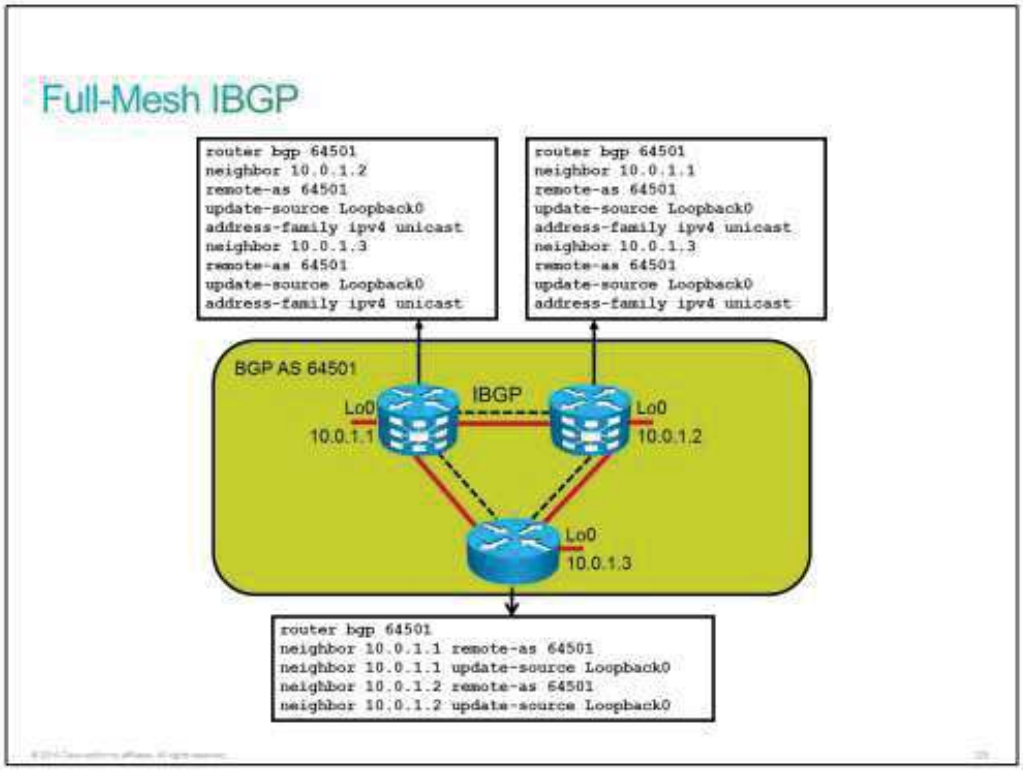

Instead of using full-mesh IBGP, you can also implement BGP route reflector or BGP confederation. These topics are beyond the scope of this course.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Support for IPv6

To exchange IPv4 and IPv6 routes across the BGP session, enable IPv4 and IPv6 unicast address families.

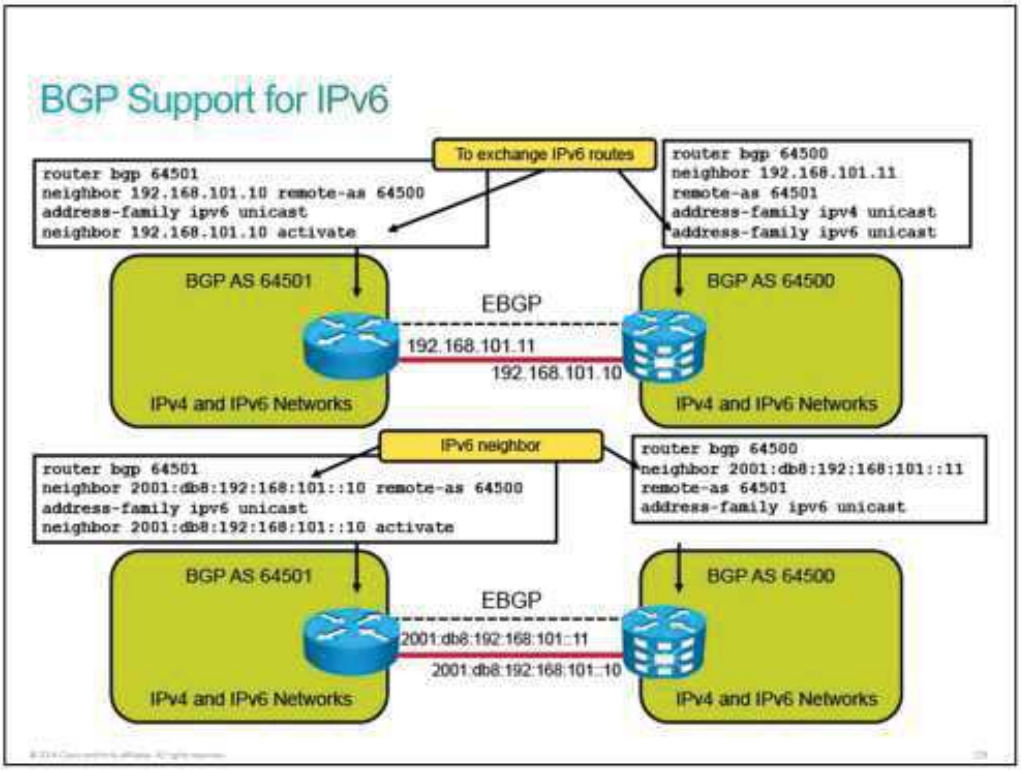

The IPv4 unicast address family is already enabled by default in the Cisco IOS or IOS XE router, but you must enable it in the Cisco IOS XR router.

When a neighbor is reachable on the IPv6 address, use IPv6 instead of the IPv4 neighbor address in the Cisco IOS XR **neighbor** router BGP command or the Cisco IOS or IOS XE **neighbor remote-as** router BGP command. When the BGP session is configured by using IPv6 addresses in the Cisco IOS or IOS XE router, only IPv4 unicast routes are exchanged by default; to enable exchange for IPv6 routes, the IPv6 unicast address family needs to be enabled. On the Cisco IOS XR Software, the exchange of IPv4 prefixes is not supported for IPv6 prefixes. Only IPv6 prefixes can be exchanged between IPv6 prefixes.

https://t.me/learningnets

Shutting Down a BGP Neighbor

If major policy changes are implemented for a neighboring router, and multiple parameters are changed, the neighboring router must be administratively shut down; then the changes must be implemented, and the neighboring router must be brought back up.

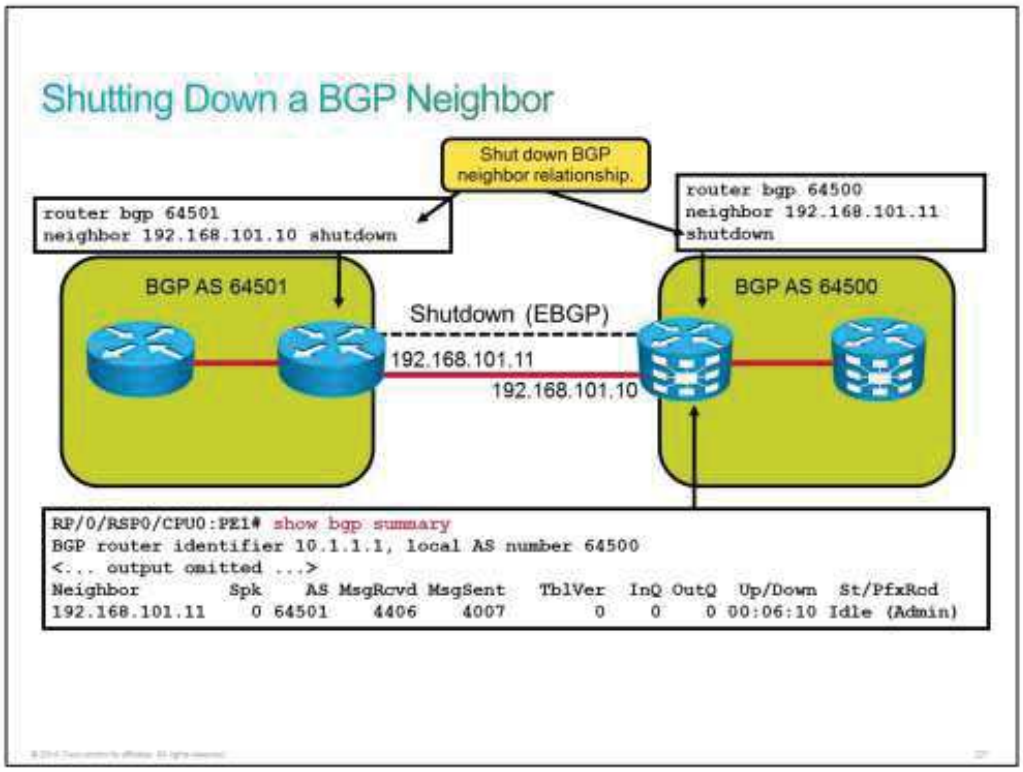

Use the Cisco IOS XR **shutdown** router BGP command or Cisco IOS or IOS XE **neighbor shutdown** router BGP command to administratively shut down a BGP neighbor. You can bring the neighbor back up with the **no** version of this command.

The usage of this command not only terminates the session, but removes all associated routing information as well.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Next-Hop Behavior

The way in which BGP establishes an IBGP relationship is very different from the way that IGPs behave. BGP is an AS-by-AS routing protocol and not a router-by-router routing protocol. The next hop is the IP address that is used to reach the next AS. The method that BGP uses to denote its next-hop address is also very different from the way that an IGP performs the same function.

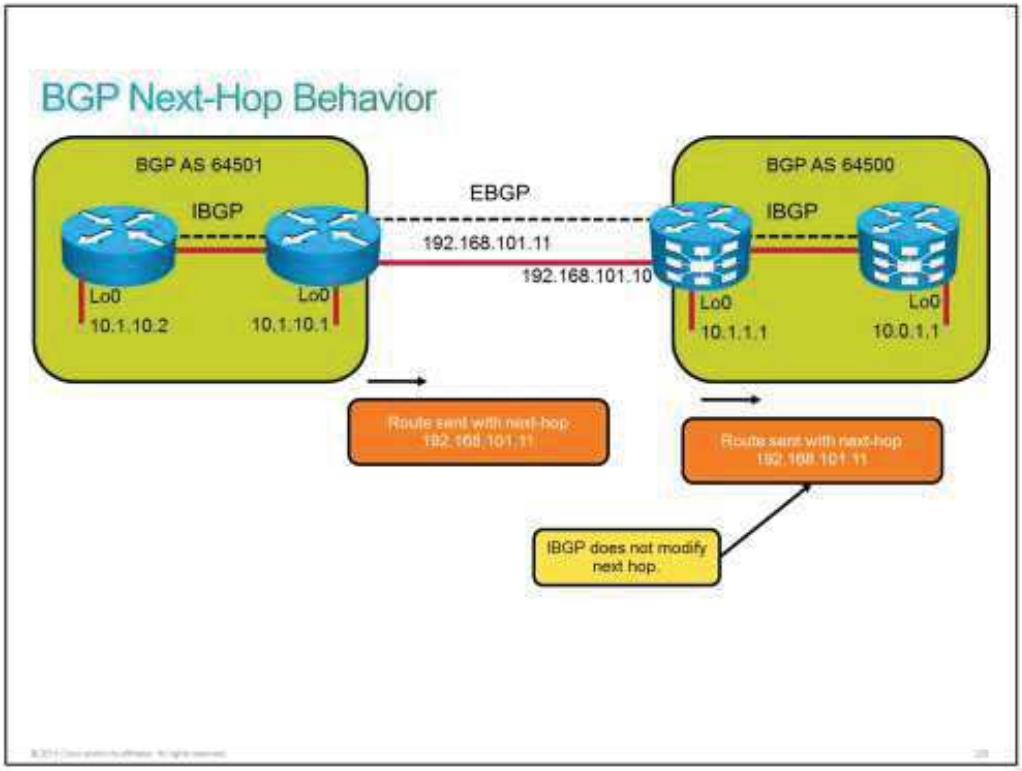

The default next hop is explained here:

- **EBGP:** IP address of the neighbor router that is sending the update
- **IBGP:** IP address that is advertised by EBGP, and should be carried in IBGP

In the figure, the next hop for a route sent via EBGP is set to the IP address of the outgoing interface (192.168.101.11). When the same route is sent to the IBGP neighbor, the next hop is not changed.

https://t.me/learningnets

BGP Next Hop Self

Overriding the default next-hop behavior of a router and forcing it to advertise itself as the next-hop address for routes that are sent to a neighboring router is sometimes necessary.

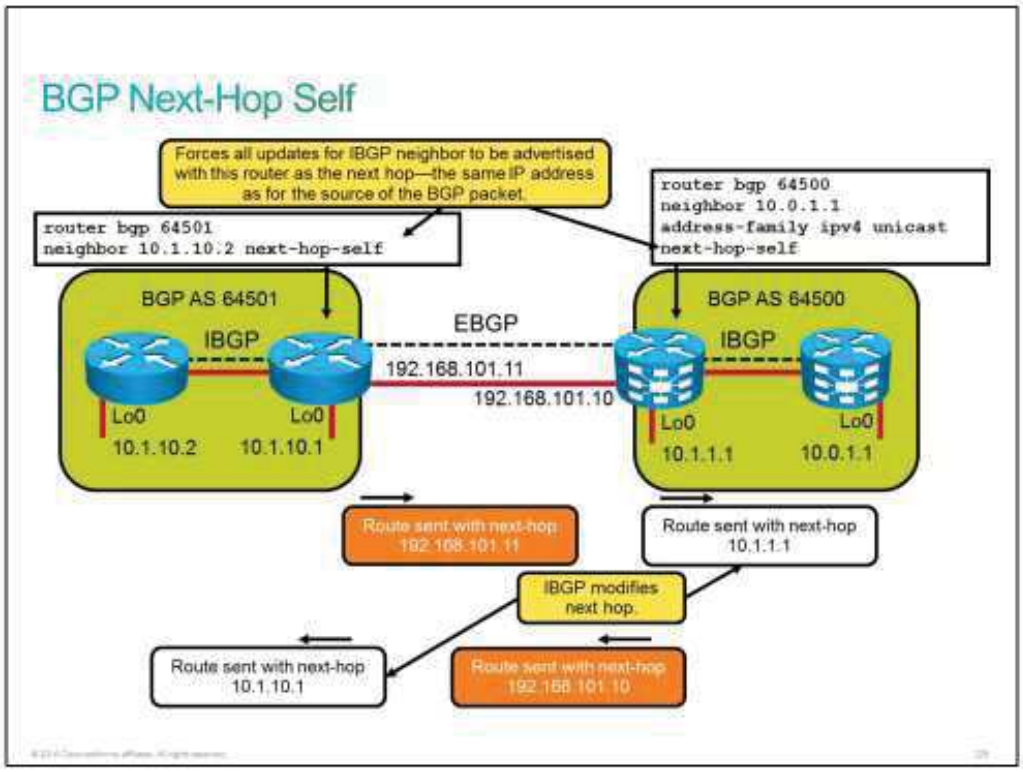

The Cisco IOS XR **next-hop-self** router BGP command or the Cisco IOS or IOS XE **neighbor next-hop-self** router BGP command forces BGP to use its own IP address as the next-hop address for each network that it advertises to its IBGP neighbor, rather than letting the protocol choose the next-hop address to use.

An internal protocol—such as RIP, EIGRP, or OSPF—always uses the source IP address of a routing update as the next-hop address for each network that is placed in the routing table. This command forces BGP to use the source IP address of the update as the next-hop address for each advertised network.

The figure shows the **next-hop-self** set to the IBGP neighbors on the Cisco IOS XR and the Cisco IOS or IOS XE routers. Because the loopback 0 interface is used in the IBGP peering, the update to the IBGP peer is sent with the loopback 0 IP address in the next-hop field.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

Cisco IOS XR BGP Configuration Templates

The **af-group**, **session-group**, and **neighbor-group** configuration commands provide template support for the neighbor configuration in Cisco IOS XR Software.

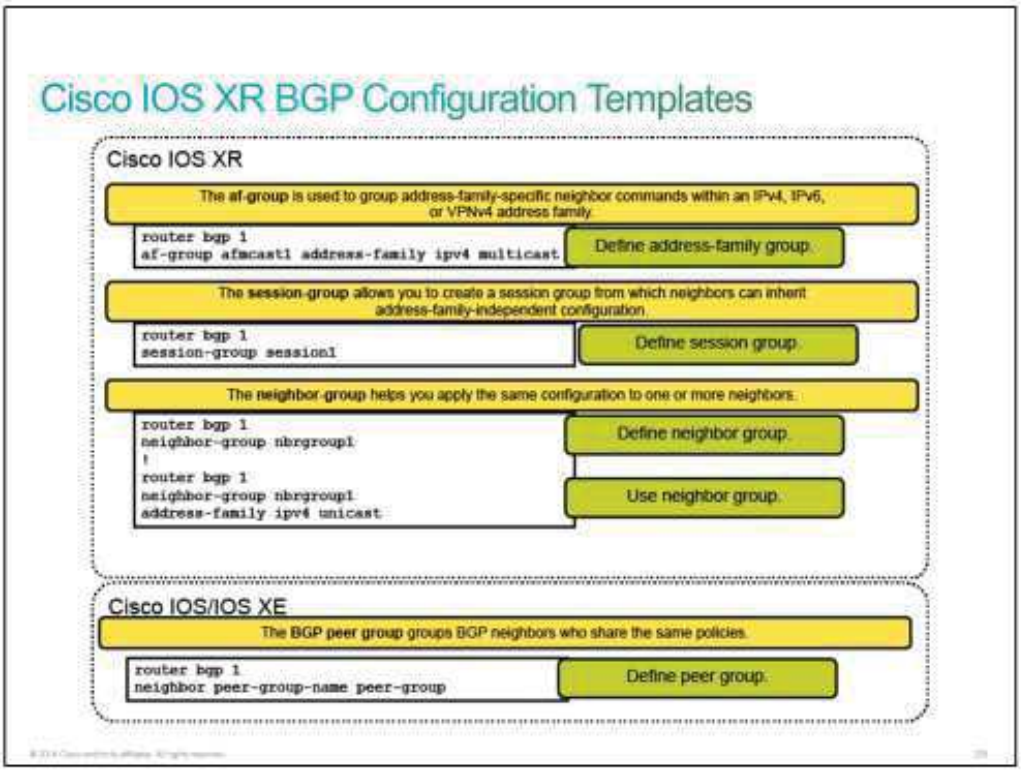

The usage of template configuration commands:

- The **af-group** command is used to group address-family-specific neighbor commands within an IPv4, IPv6, or VPNv4 address family. Neighbors that have the same address-family configuration are able to use the af-group name for their address-family-specific configuration. A neighbor inherits the configuration from an address-family group by way of the **use** command. If a neighbor is configured to use an address-family group, the neighbor (by default) inherits the entire configuration from the address-family group. However, a neighbor does not inherit all of the configuration from the address-family group, if items are explicitly configured for the neighbor. The address-family group configuration is entered under the BGP router configuration mode.
- The **session-group** command allows you to create a session group from which neighbors can inherit address-family-independent configuration. A neighbor inherits the configuration from a session group by way of the **use** command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the entire configuration of the session group. A neighbor does not inherit all of the configuration from a session group if a configuration is done directly on that neighbor.
- The **neighbor-group** command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address-family groups and can comprise the complete configuration for a neighbor. After a neighbor group is configured, a neighbor can inherit the configuration of the group using the **use** command. If a neighbor is configured to use a neighbor group, the neighbor inherits the entire BGP configuration of the neighbor group. However, a neighbor does not inherit all of the configuration from the neighbor group, if items are explicitly configured for the neighbor. In addition, some part of the configuration of the neighbor group could be hidden if a session group or address-family group was also being used.

https://t.me/learningnets

In the Cisco IOS or IOS XE you can group BGP neighbors who share the same policies together in what is called a **BGP peer group**. Instead of configuring each neighbor with the same policy individually, a peer group allows you to group the policies that can be applied to individual peers, thus making an efficient update calculation along with simplified configuration.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Neighbor States

After the TCP handshake is complete, the BGP application tries to set up a session with the neighbor. Several steps must occur for the session to be established.

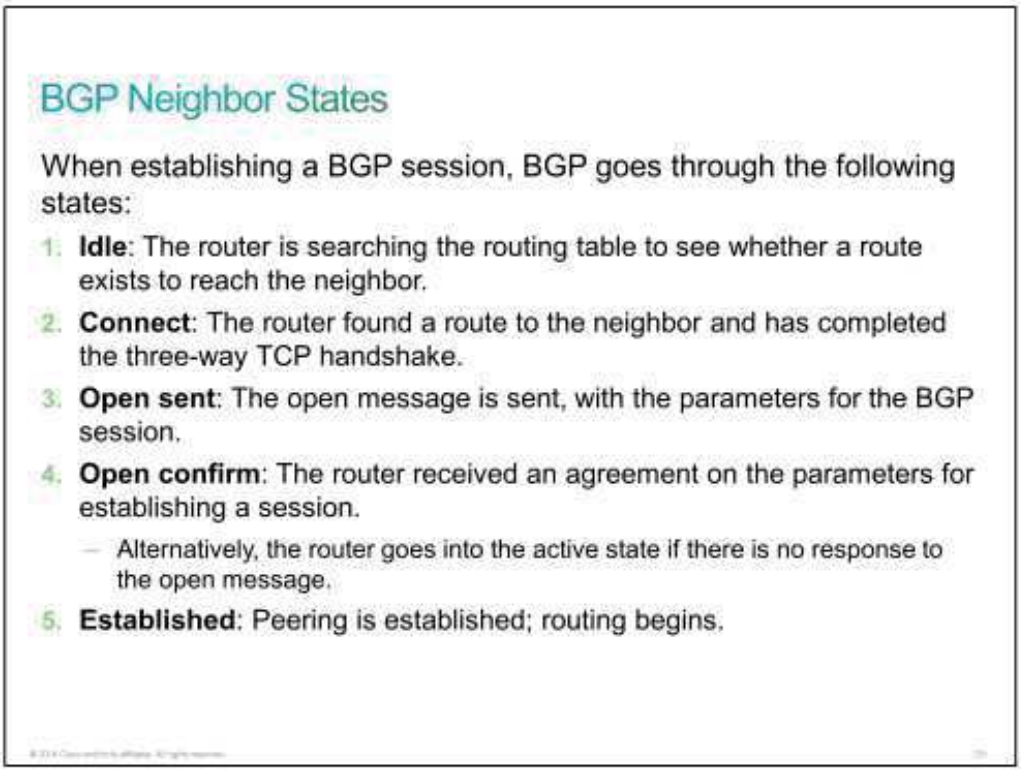

After the **neighbor** command is entered in BGP, BGP takes the IP address that is listed and checks the local routing table for a route to this address. At this point, BGP is in the "idle" state. If BGP does not find a route to the IP address, it stays in the idle state. If it finds a route, it goes to the "connect" state when the TCP handshaking SYN-ACK packet returns.

After the TCP connection is complete, BGP creates a BGP open packet and sends it to the neighbor. Once BGP sends this open packet, the BGP peering session changes to the "open sent" state. If there is no response for five seconds, the state changes to the "active" state.

If a response does come back in a timely manner, BGP goes to the "open confirm" state and starts scanning (evaluating) the routing table for the paths to send to the neighbor. When those paths have been found, BGP then goes to the "established" state and begins routing between the neighbors.

https://t.me/learningnets

BGP Active State

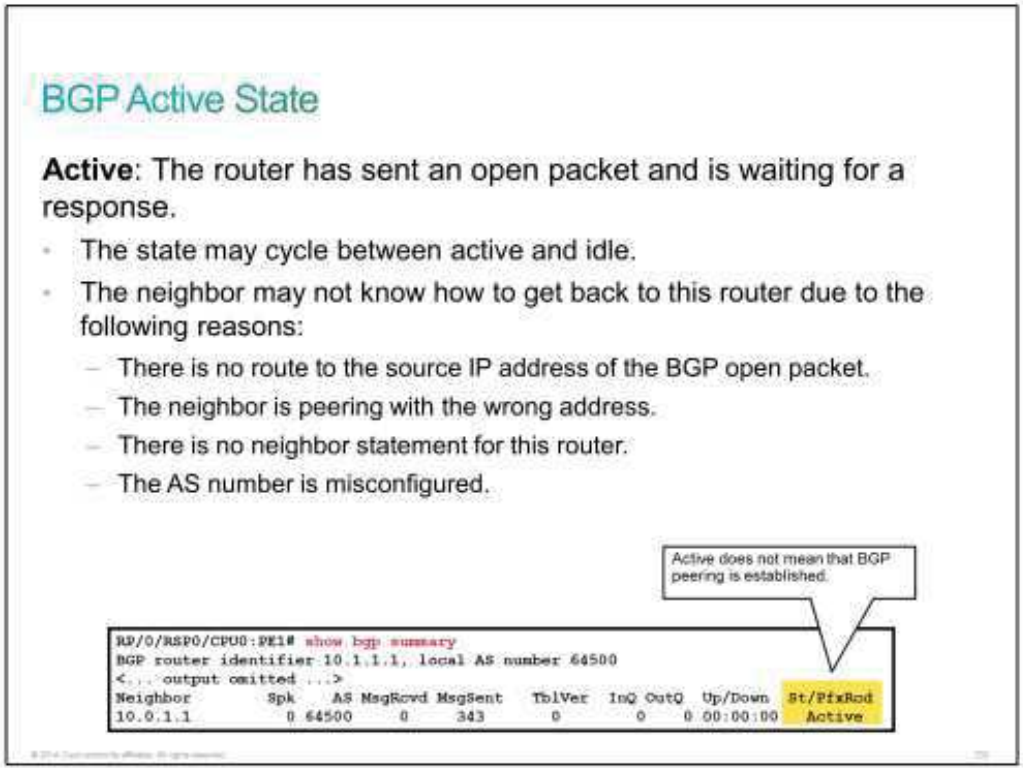

If the router is in the active state, it has found the IP address in the neighbor statement and has created and sent out a BGP open packet. However, the router has not received a response (open confirm packet).

One common problem in this case is that the neighbor may not have a return route to the source IP address. Ensure that the source IP address or network of the packets has been announced to the local routing protocol (IGP).

Another common problem that is associated with the active state occurs when a BGP router attempts to peer with another BGP router that does not have a neighbor statement peering back to the first router, or when the other router is peering with the wrong IP address on the first router. Check to ensure that the other router has a neighbor statement that is peering to the correct address of the router that is in the active state.

If the state toggles between the idle state and the active state, one of the most common problems is AS number misconfiguration.

Note	When an update or a withdraw is not sent to a neighbor, this results in a black-hole routing and consequently into a BGP route inconsistency with a neighbor. To identify such issues, BGP consistency checker can be configured. To configure the consistency checker, use the **bgp consistency-checker {error-message | auto-repair} [interval** *minutes*] router BGP configuration command. Once the process identifies such inconsistency, it reports it with a syslog message and optionally takes action, if the auto-repair keyword is specified.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Established and Idle States

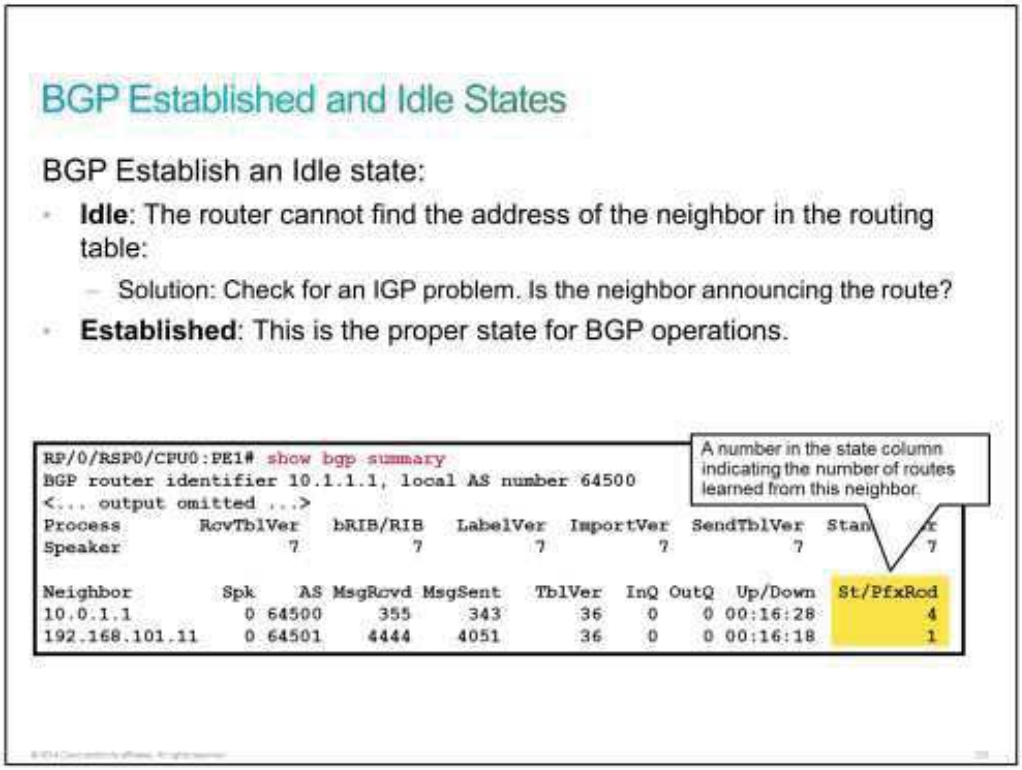


```
RP/0/RSP0/CPU0:PE1# show bgp neighbor
< output omitted >
BGP neighbor is 192.168.101.11
 Remote AS 64501, local AS 64500, external link
 Remote router ID 10.1.10.1
  BGP state = Established, up for 00:16:18
  Last read 00:00:49, Last read before reset 00:16:52
  Hold time is 180, keepalive interval is 60 seconds
  Configured hold time: 180, keepalive: 60, min acceptable hold time: 3
< output omitted >
```

The idle state is an indication that the router does not know how to reach the IP address that is listed in the neighbor statement. The router is idle due to one of these scenarios:

- It is waiting for a static route to that IP address or network to be configured.
- It is waiting for the local routing protocol (IGP) to learn about this network through an advertisement from another router.

The most common reason for a router to enter the idle state is that the neighbor is not announcing the IP address or the network toward which the neighbor statement of the router is pointing. Check these two conditions first to correct this problem:

- Ensure that the neighbor announces the route in its local routing protocol (IGP).
- Verify that an incorrect IP address has not been entered in the neighbor statement.

The established state is the desired state for the neighbor relationship. This state means that both routers have agreed to exchange BGP updates with one another and routing has begun.

https://t.me/learningnets

Use the Cisco IOS XR **show bgp summary** command or the Cisco IOS or IOS XE **show ip bgp summary** command to examine the number of prefixes received. Use the Cisco IOS XR **show bgp neighbor** command or the Cisco IOS or IOS XE **show ip bgp neighbor** command to display information about the BGP connections to neighbors.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Neighbor Authentication

BGP neighbor authentication can be configured on a router so that the router authenticates the source of each routing update packet that it receives. This authentication is accomplished by the exchange of an authentication key (sometimes referred to as a password) that is known to both the sending and the receiving router.

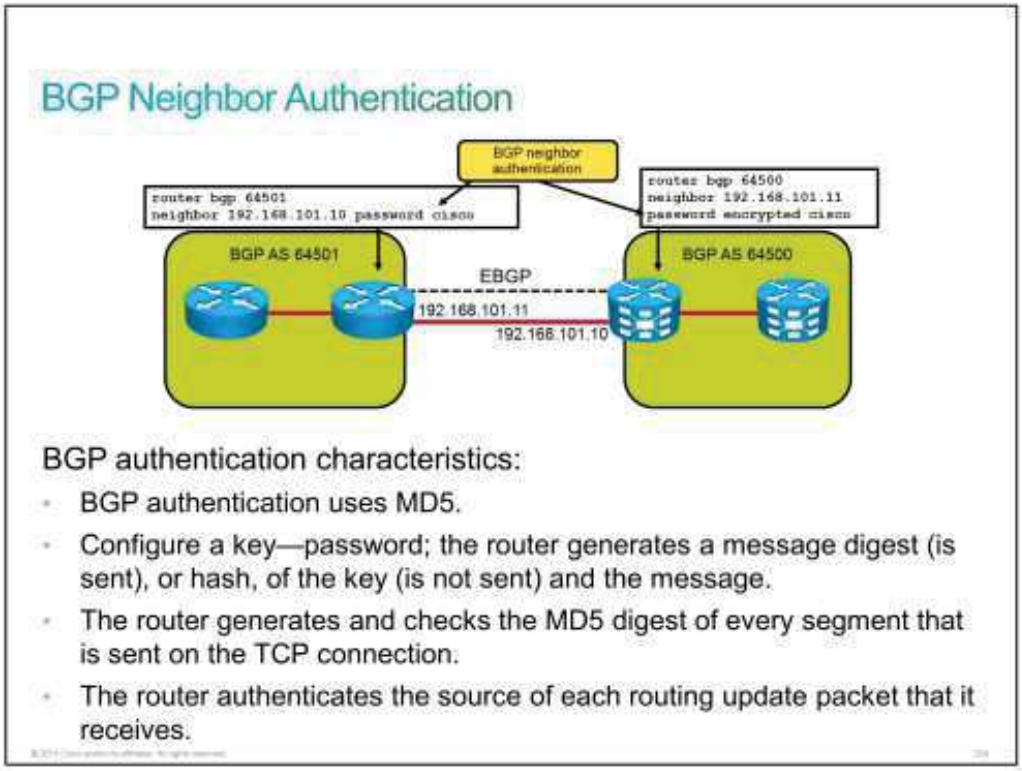

BGP supports MD5 neighbor authentication. MD5 sends a message digest (also called a "hash") that is created using the key and a message. The message digest is then sent instead of the key. The key itself is not sent to prevent it from being read by an eavesdropper on the line while it is being transmitted.

To enable MD5 authentication on a TCP connection between two BGP peers, use the Cisco IOS XR **password** router BGP command or the Cisco IOS or IOS XE **neighbor password** router BGP command.

MD5 authentication can be configured between two BGP peers, which means that each segment that is sent on the TCP connection between the peers is verified. MD5 authentication must be configured with the same password on both BGP peers; otherwise, the connection between the peers will not be made. Configuring MD5 authentication causes the Cisco IOS or IOS XE or IOS XR Software to generate and check the MD5 digest of every segment that is sent on the TCP connection.

https://t.me/learningnets

Clearing the BGP Session

BGP can potentially process huge volumes of routing information. When a policy configuration change occurs, the router cannot go through the huge table of BGP information and recalculate which entry is no longer valid in the local table; also, the router cannot determine which route or routes, already advertised, should be withdrawn from a neighbor.

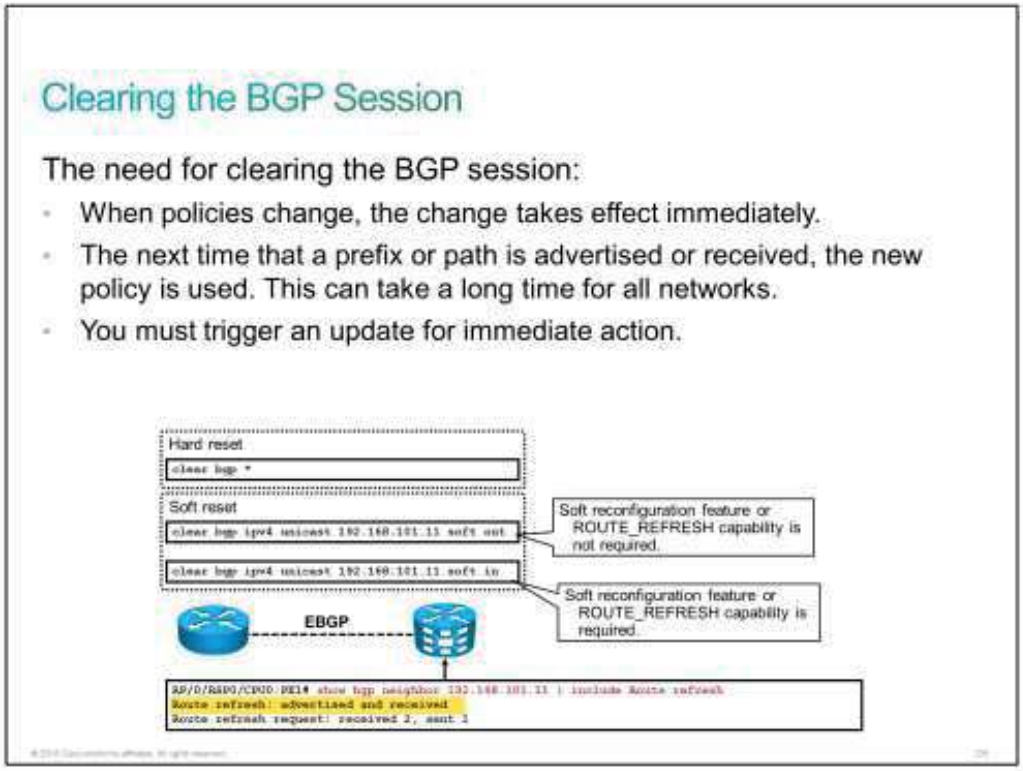

There is an obvious risk that the first configuration change will be immediately followed by a second, which would cause the whole process to start all over again. To avoid such a problem, the Cisco IOS or IOS XE or IOS XR Software applies changes only to those updates that are received or transmitted after the BGP policy configuration change has been performed. The new policy, enforced by the new filters, is applied only on routes that are received or sent after the change.

A network administrator who would like the policy change to be applied on all routes must trigger an update to force the router to let all routes pass through the new filter. If the filter is applied on outgoing information, the router has to resend the BGP table through the new filter. If the filter is applied on incoming information, the router needs its neighbor to resend its BGP table so that it passes through the new filters.

There are three ways to trigger an update:

- **Hard reset**: The Cisco IOS XR **clear bgp** * command or the Cisco IOS or IOS XE **clear ip bgp** * command causes the BGP forwarding table on the router that issued this command to be deleted, and all networks must be relearned from every neighbor. If a router has multiple neighbors, this action is a very dramatic event. This command forces all neighbors to resend their entire tables simultaneously.

https://t.me/learningnets

- **Soft reset**: The Cisco IOS or IOS XE or IOS XR **soft out** option of the previous command causes BGP to do a soft reset for outbound updates. The router that is issuing the BGP soft reset does not reset the BGP session; instead, the router creates a new update and sends the whole table to the specified neighbor. This update includes withdrawal commands for the networks that the other neighbor will not see any more, based on the new outbound policy. The Cisco IOS or IOS XE or IOS XR **soft in** option of the previous command causes BGP to do a soft reset for inbound updates; this option requires the soft reconfiguration feature to be enabled for every neighbor. This feature is very memory-consuming, because the router has to store the BGP routing table received from every neighbor.
- **Route refresh**: The Cisco IOS or IOS XE or IOS XR **soft in** option of the previous command also works without the soft reconfiguration feature. If an inbound soft reset is triggered, BGP sends a REFRESH request to the neighbor, if the neighbor has advertised the ROUTE_REFRESH capability. To determine whether the neighbor has advertised the ROUTE_REFRESH capability, use the Cisco IOS XR **show bgp neighbors** command or the Cisco IOS or IOS XE **show ip bgp neighbors** command.

https://t.me/learningnets

Monitoring BGP Routes

When a BGP session is reset and soft reconfiguration is used, several commands exist to monitor the BGP routes that are received, sent, or filtered.

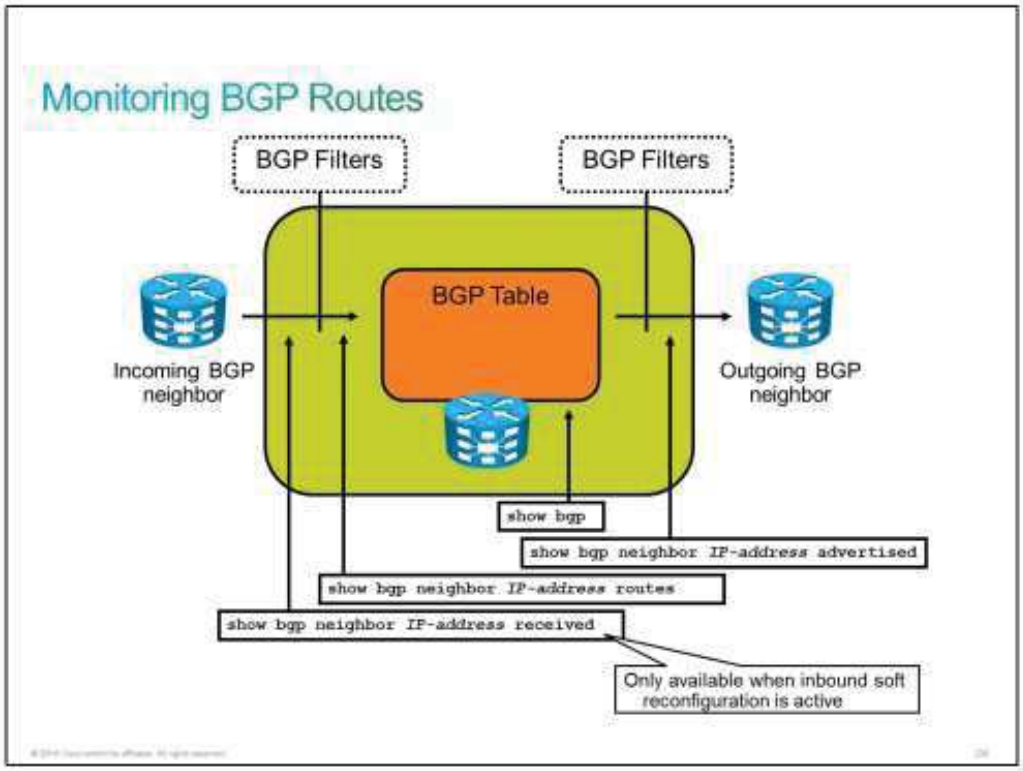

The following commands can be used:

- Use the Cisco IOS XR **show bgp neighbor** *IP-address* **received** command or the Cisco IOS or IOS XE **show ip bgp neighbor** *IP-address* **received** command to list BGP routes that are received from the BGP neighbor, before filters are applied. This command is available only when inbound soft reconfiguration is enabled for the BGP neighbor.
- Use the Cisco IOS XR **show bgp neighbor** *IP-address* **routes** command or the Cisco IOS or IOS XE **show ip bgp neighbor** *IP-address* **routes** command to list BGP routes that are received from the BGP neighbor, after inbound filters are applied.
- Use the Cisco IOS XR **show bgp** command or the Cisco IOS or IOS XE **show ip bgp** command to list all BGP routes that are installed into the BGP routing table.
- Use the Cisco IOS XR **show bgp neighbor** *IP-address* **advertised** command or the Cisco IOS or IOS XE **show ip bgp neighbor** *IP-address* **advertised** command to list BGP routes that are sent to the BGP neighbor.

 © 2014 Cisco Systems, Inc.

https://t.me/learningnets

BGP Path Selection

Routers often have several neighbors and receive routing updates from each neighbor. All routing updates enter the BGP forwarding table; as a result, multiple paths may exist to reach a given network.

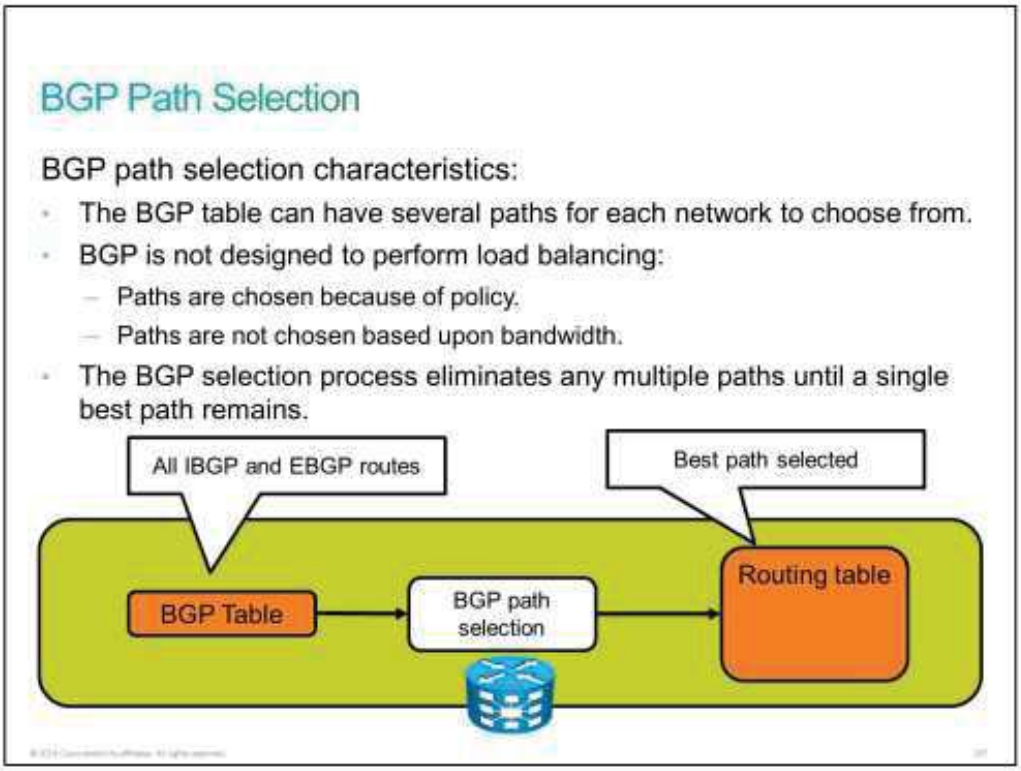

The entire BGP forwarding table can be displayed using the Cisco IOS XR **show bgp** command or the Cisco IOS or IOS XE **show ip bgp** command.

Next, paths for the network are evaluated to determine which path is best. Paths that are not the best are eliminated from the selection criteria but kept in the BGP forwarding table in case the best path becomes inaccessible. If one of the best paths is not accessible, a new best path must be selected.

BGP is not designed to perform load balancing; paths are chosen based on policy, not based on bandwidth. The BGP selection process eliminates any multiple paths until a single best path remains.

The best path is submitted to the routing table manager process and is evaluated against any other routing protocols that can also reach that network. The router usually runs BGP and one of the IGPs—OSPF, IS-IS, and so on—that are sending candidates for the routing table to the routing table manager. The route from the source with the lowest administrative distance is installed in the routing table. The entire routing table can be displayed using the Cisco IOS XR **show route** command or the Cisco IOS or IOS XE **show ip route** command.

https://t.me/learningnets

BGP Route Selection Decision Process

After BGP receives updates about different destinations from different autonomous systems, it chooses the single best path to reach a specific destination.

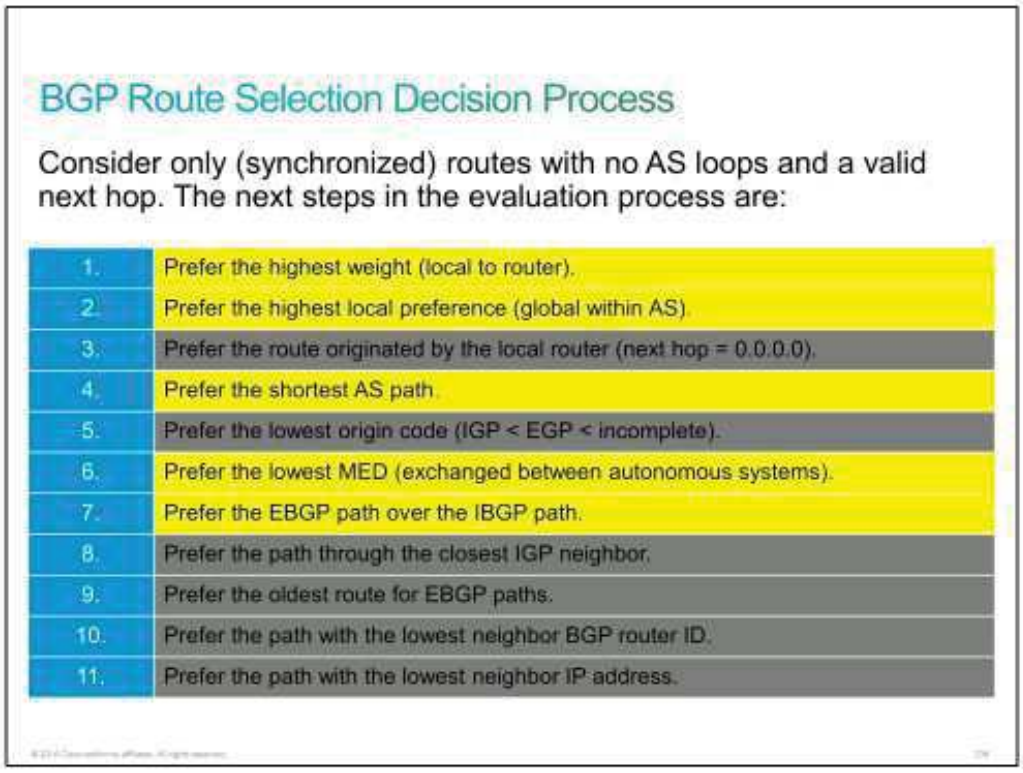

The decision process is based on the BGP attributes. When faced with multiple routes to the same destination, BGP chooses the best route for routing traffic toward the destination. BGP considers only synchronized routes with no AS loops and a valid next hop.

Note	BGP synchronization is disabled by default. With BGP synchronization, enabled routes learned via IBGP must be validated by the interior routing table before they can be advertised to EBGP peers.

The following process summarizes how BGP chooses the best route on a Cisco router:

1. Prefer the route with the highest weight. (Weight is proprietary to Cisco and is local to the router only.)
2. If multiple routes have the same weight, prefer the route with the highest local preference. (The local preference is used within an AS.)
3. If multiple routes have the same local preference, prefer the route that the local router originated. A locally originated route has a next hop of 0.0.0.0 in the BGP table.
4. If none of the routes was locally originated, prefer the route with the shortest AS path.
5. If the AS path length is the same, prefer the lowest origin code (IGP < EGP < incomplete).

https://t.me/learningnets

6. If all origin codes are the same, prefer the path with the lowest MED. (MED is exchanged between autonomous systems.) The MED comparison is made only if the neighboring AS is the same for all routes that are considered, unless the Cisco IOS XR **bgp bestpath med always** command or the Cisco IOS or IOS XE **bgp always-compare-med** command is enabled. The IETF decision regarding BGP MED assigns a value of infinity to the missing MED, making the route that is lacking the MED variable the least preferred. The default behavior of BGP routers is to treat routes without the MED attribute as having a MED of 0, making the route that is lacking the MED variable the most preferred. To configure the router to conform to the IETF standard, use the Cisco IOS XR **bgp bestpath med missing-as-worst** router BGP command or the Cisco IOS or IOS XE **bgp bestpath missing-as-worst** router BGP command.
7. If the routes have the same MED, prefer external paths (External BGP, or EBGP) to internal paths (Internal BGP, or IBGP).
8. If synchronization is disabled and only internal paths remain, prefer the path through the closest IGP neighbor. This step means that the router will prefer the shortest internal path within the AS to reach the destination (the shortest path to the BGP next hop).
9. For EBGP paths, select the oldest route to minimize the effect of routes that are going up and down (flapping).
10. Prefer the route with the lowest neighbor BGP router ID value.
11. If the BGP router IDs are the same, prefer the router with the lowest neighbor IP address.

Only the best path is entered in the routing table and propagated to the BGP neighbors of the router.

https://t.me/learningnets

BGP Weight Attribute

The BGP weight is an attribute that Cisco defines for the path-selection process.

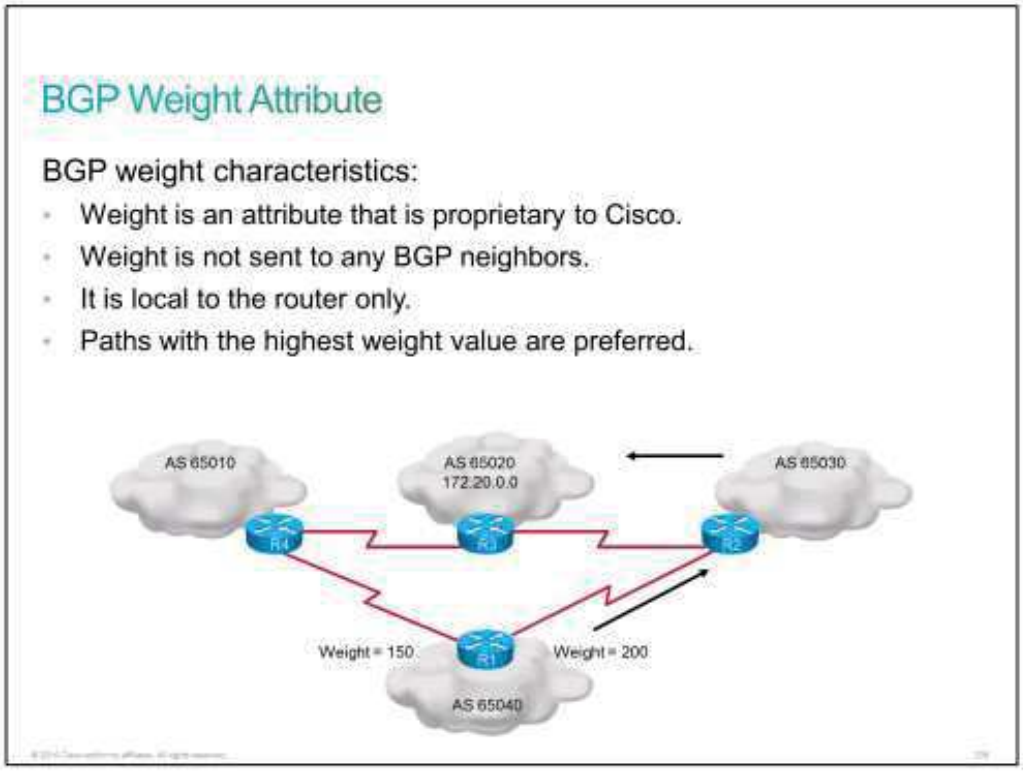

The weight is configured locally on a router and is not propagated to any other routers. This attribute applies when one router is used with multiple exit points out of an AS, as opposed to the local preference attribute that is used when two or more routers provide multiple exit points.

The weight can have a value from 0 to 65535. Paths that the router originates have a weight of 32768 by default, and other paths have a weight of 0 by default.

Routes with a higher weight are preferred when multiple routes exist to the same destination.

In the figure, R2 and R4 learn about network 172.20.0.0 from AS 65020 and propagate the update to R1. R1 has two ways to reach 172.20.0.0, and it must decide which route to take—the path through R2 or the path through R4.

R1 sets the weight of updates that are coming from R2 to 200, and the weight of the updates that are coming from R4 to 150. Because the weight for the route that is pointing to R2 is higher than the weight for the route that is pointing to R4, R1 uses R2 as a next hop to reach 172.20.0.0.

https://t.me/learningnets

BGP Local Preference Attribute

Local preference is a well-known discretionary attribute that provides information to routers in the AS about the path that is preferred for exiting the AS. A path with a higher local preference is preferred.

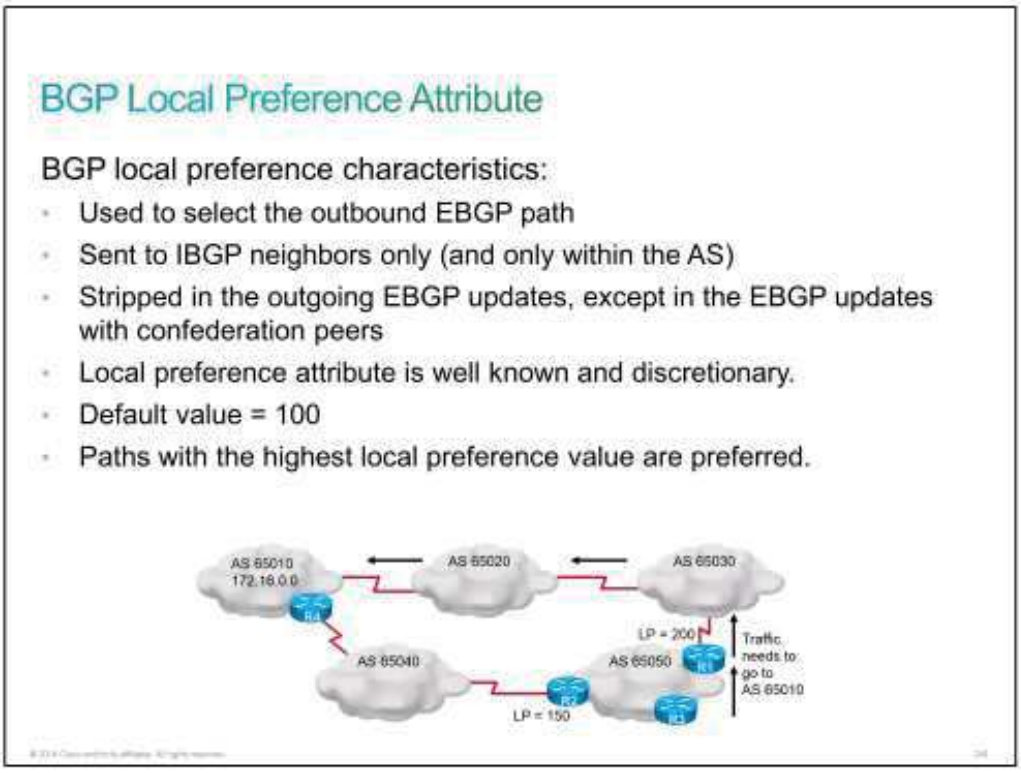

The local preference is an attribute that is configured on a router and exchanged among routers within the same AS only. The default value for local preference on a Cisco router is 100.

To change the default local preference value of 100, use the Cisco IOS or IOS XE or IOS XR **bgp default local-preference** router BGP command.

In the figure, AS 65050 receives updates about network 172.16.0.0 from two neighbors. Each of the networks is advertising a different path to the destination. The local preference on R1 for network 172.16.0.0 is set to 200, and the local preference on R2 for network 172.16.0.0 is set to 150.

Because the local preference information is exchanged within AS 65050, all routers are aware of the exit point for network 172.16.0.0 out of AS 65050. In the figure, R1 is configured with a higher local preference than R2, and all the traffic that is destined for network 172.16.0.0 will be sent to R1 as an exit point from AS 65050.

https://t.me/learningnets

BGP AS Path Attribute

When connections to multiple providers are required, BGP should select the optimum route for traffic to use. The optimum, or best route may not be what the network designer intended, based on design criteria, administrative policies, or corporate mandate.

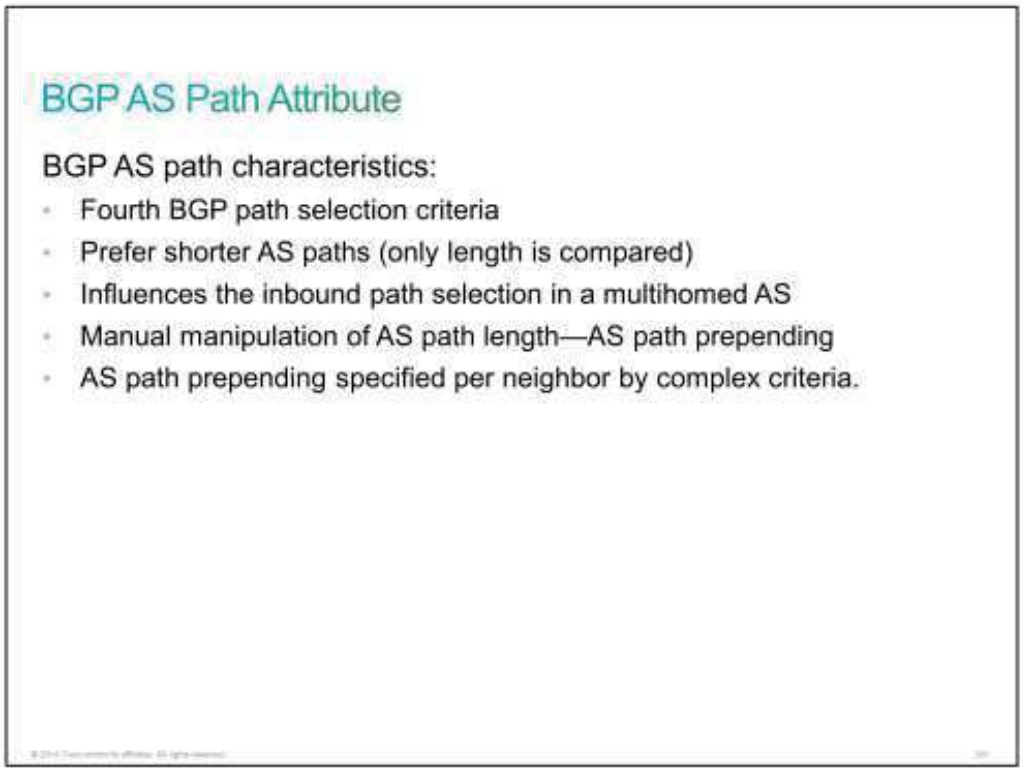

Selecting the appropriate path for outgoing traffic is fairly easy for an AS. Influencing other autonomous systems to select the appropriate path for traffic that is returning to a specific AS is much more complicated. It is unlikely that the operator of an AS can request changes in router configurations in another AS. This limitation makes it virtually impossible to influence another AS to select the desired path that is based on the weight and local preference attributes, because both options would require configuration changes in the neighboring AS.

If no BGP path selection tools are configured on the route to influence the traffic flow, BGP will use the shortest AS path—the fourth option in the path selection process. If the AS path is not manually manipulated by some administrative means, the path that is going over the fewest number of autonomous systems is selected by the router, regardless of available bandwidth.

However, if the AS that is attempting to influence the incoming traffic flow is sending out EBGP updates with a manipulated AS path attribute over that undesired path, the receiver of this update is less likely to select it as the best, because the AS path now appears to be longer. AS path prepending potentially allows the customer to influence the route selection of its service providers. The AS path is extended with multiple copies of the AS number of the sender. There is no exact mechanism to calculate the required prepended AS path length.

The benefit of manipulating AS paths to influence route selection is that the configuration that is needed is done in the AS that is requesting a desired return path.

https://t.me/learningnets

BGP Multi-Exit Discriminator Attribute

The MED attribute, also called the metric, is an optional nontransitive attribute.

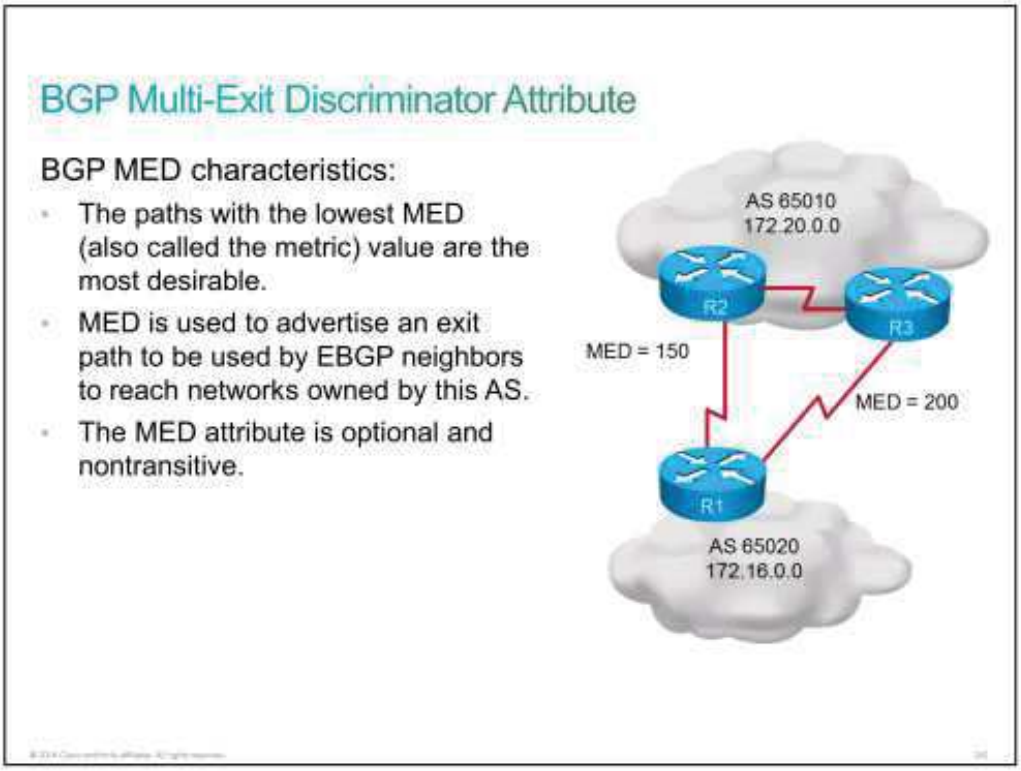

The MED is an indication to EBGP neighbors about the preferred path into an AS. The MED attribute is a dynamic way to influence another AS about which path it should choose to reach a certain route, when multiple entry points into an AS exist. A lower metric is preferred.

Unlike local preference, the MED is exchanged between autonomous systems. The MED is sent to EBGP peers. Those routers propagate the MED within their AS, and the routers within the AS use the MED but do not pass it on to the next AS. When the same update is passed on to another AS, the metric is set back to the default of 0. To change this value, use the Cisco IOS or IOS XE or IOS XR **default-metric** router BGP command. All routes that are advertised to an EBGP neighbor are set to the value that is specified using this command.

MED influences inbound traffic to an AS, and local preference influences outbound traffic from an AS. By default, a router compares the MED attribute only for paths from neighbors in the same way AS. The MED attribute means that BGP is the only protocol that can affect how routes are sent into an AS.

In the figure, the R2 MED attribute is set to 150, and the R3 MED attribute is set to 200. When R1 receives updates from R2 and R3, it picks R2 as the best next hop because its MED of 150 is less than the R3 MED of 200.

https://t.me/learningnets

Summary

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

Summary

- For a BGP configuration, BGP requirements, BGP parameters, and connectivity must be defined.
- Basic EBGP configuration requires three main steps: define the BGP process, establish the neighbor relationship, and advertise the networks into BGP.
- Networks can be advertised into BGP by using the **network** command or by redistribution.
- To exchange IPv6 networks over BGP, you have to activate a neighbor for IPv6 address family.
- You can use the **shutdown** command to manually shut down a BGP neighbor.
- By default, the next hop over EBGP session is the IP address of a router that is sending the update. The next hop over the IBGP session is as advertised by EBGP, and should not change.

Summary (Cont.)

- You can use configuration templates to group configuration that can be applied to several BGP neighbors.
- When establishing a BGP session, the BGP goes through the following states: idle, connect, open sent, open confirm, and established.
- BGP supports MD5 authentication to authenticate each received routing packet.
- You should trigger a BGP update by resetting a BGP session when you change routing policy.
- After BGP receives updates about multiple destinations from different autonomous systems, it follows a multiple-step process for selecting the best route to reach a destination; the best route is a candidate for the routing table.
- Weight, local preference, AS path length, and MED are all attributes that can influence the best route decision.

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Lesson 4: Module Summary

Overview

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

The Internet has proven to be a valuable tool to many companies and service providers, resulting in multiple redundant connections between autonomous systems. The function of BGP is to provide alternatives to using default and static routes to control path selections.

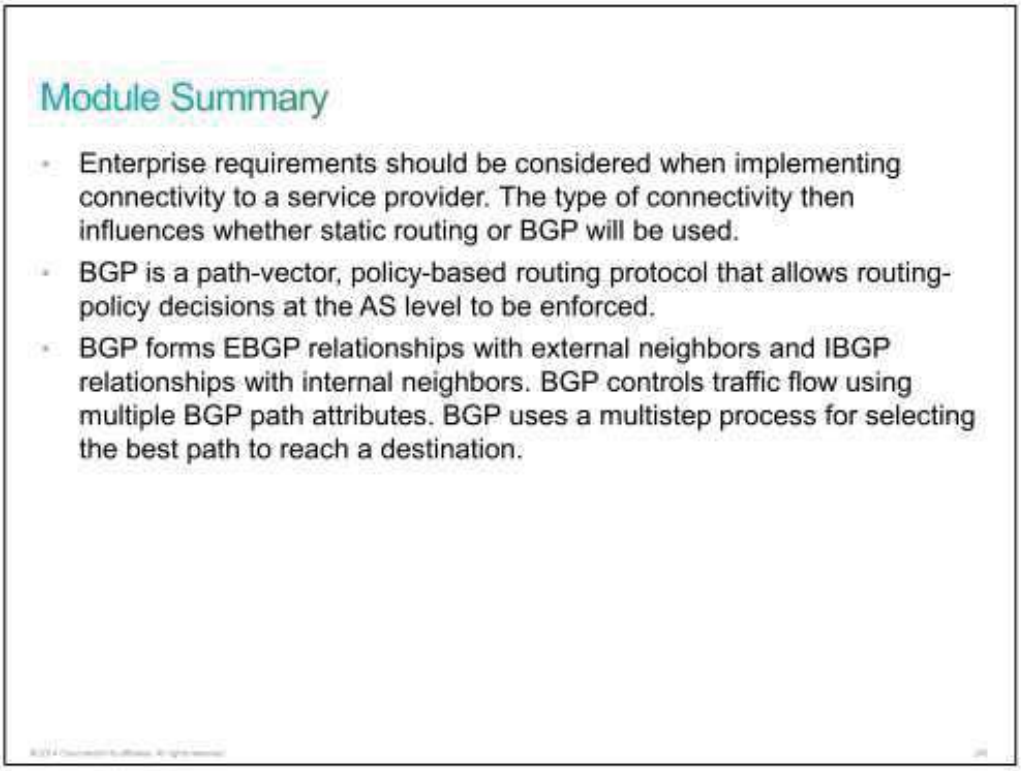

https://t.me/learningnets

 ©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. What is *not* a requirement for enterprise-to-service provider connectivity? (Source: Enterprise Connectivity to Service Providers)

   A. public IP address space
   B. link type and bandwidth availability
   C. AS routing policy
   D. connection redundancy

2. What are three routing update exchange options for enterprise-to-service provider connectivity? (Choose three) (Source: Enterprise Connectivity to Service Providers)

   A. static routes
   B. EIGRP
   C. IS-IS
   D. OSPF
   E. MPLS VPNs
   F. BGP

3. Which one of these statements is not the reason for using BGP as a routing update exchange mechanism? (Source: Enterprise Connectivity to Service Providers)

   A. A customer deploys BGP to announce its public networks.
   B. A BGP is typically used for inter-AS routing.
   C. Customer routers are connected to service provider PE routers.
   D. Customer network implementation requires a complete Internet routing table.

https://t.me/learningnets

4. What are four enterprise-to-service provider connection options? (Choose four) (Source: Enterprise Connectivity to Service Providers)

   A. single-homed
   B. single-attached
   C. dual-homed
   D. dual-attached
   E. multihomed
   F. dual-multihomed

5. What is a characteristic of dual-multihomed service provider connectivity? (Source: Enterprise Connectivity to Service Providers)

   A. a connection to two or more different service providers, with two links per service provider
   B. a connection to multiple service providers, with one link per service provider
   C. the default route pointing to each service provider from an enterprise network
   D. each service provider announcing a default route with a different metric to the enterprise network

6. What are three common ways to perform multihoming? (Choose three.) (Source: Enterprise Connectivity to Service Providers)

   A. Each service provider passes only a default route to the AS.
   B. Each service provider passes a default route and provider-owned specific routes to the AS.
   C. Each service provider passes selected provider-owned routes, but no default routes, to the AS.
   D. Each service provider passes all routes to the AS.

7. Which statement about the AS is true? (Source: Introducing BGP Routing)

   A. The AS is a collection of networks under a single administrative domain.
   B. The AS is a collection of networks that belong to one enterprise network.
   C. The AS requires IGP protocol to exchange routing information between autonomous systems.
   D. EBGP neighbors must be configured within the same AS.

8. What are the two typical reasons for multihoming? (Choose two.) (Source: Introducing BGP Routing)

   A. to increase the reliability of the connection to the Internet
   B. to increase the performance of the connection
   C. to increase the bandwidth of the connection
   D. to simplify the IGP protocol configuration

9. What is a drawback of having all of your connections to a single service provider? (Source: Introducing BGP Routing)

   A. It has redundancy with the multiple connections.
   B. Connectivity issues in that single service provider can cause your autonomous system to lose connectivity to the Internet.
   C. It is not tied into the routing policy of a single connection.
   D. It has more paths to the same networks for better policy manipulation.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

10. Which two conditions are valid reasons to run BGP in an AS? (Choose two.) (Source: Introducing BGP Routing)

A. The AS has only a single connection to another AS.
B. Path and packet flow manipulation is required in this AS.
C. You have a limited understanding of BGP routing and route filtering.
D. The AS is a service provider.

11. Which routing method best describes BGP? (Source: Introducing BGP Routing)

A. distance vector
B. link state
C. path vector
D. hybrid of link state and distance vector

12. What message type is NOT defined by BGP? (Source: Introducing BGP Routing)

A. open
B. keepalive
C. update
D. notification
E. established

13. By default, what are two conditions for routers to be EBGP neighbors? (Choose two.) (Source: Implementing Basic BGP Routing)

A. Routers must be in the same AS.
B. Routers must be in different autonomous systems.
C. Routers are running an IGP between them to establish an adjacency.
D. Routers are directly connected.

14. What does BGP use during the best-path selection process? (Source: Implementing Basic BGP Routing)

A. speed
B. AS routing policy
C. bandwidth plus delay
D. number of routers to reach a destination network

15. Which four parameters are required for a basic BGP configuration? (Choose four) (Source: Implementing Basic BGP Routing)

A. neighbors (peers) that are involved
B. summarization networks
C. AS numbers that are used
D. IP addresses that are used
E. prefix-list filtering
F. networks, which need to be advertised
G. AS-path access list filtering

https://t.me/learningnets

16. Which Cisco IOS XR router BGP command indicates to a BGP router whether an IP address belongs to an IBGP neighbor or to an EBGP neighbor? (Source: Implementing Basic BGP Routing)

A. **shutdown**
B. **update-source Loopback0**
C. **remote-as 65010**
D. **next-hop-self**

17. Which Cisco IOS XR router BGP command sets the source IP address of a BGP update to be the IP address of a specific interface? (Source: Implementing Basic BGP Routing)

A. **shutdown**
B. **update-source Loopback0**
C. **remote-as 65020**
D. **next-hop-self**

18. What is the result of using this command: **router bgp 65010?** (Source: Implementing Basic BGP Routing)

A. The BGP process starts in the router.
B. The BGP process starts on the interface.
C. The neighboring router AS is defined.
D. The router enters into BGP configuration mode, with AS number 65010 used locally.

19. The **network** command that is used in the router BGP process identifies the interfaces from which to advertise BGP updates. (Source: Implementing Basic BGP Routing)

A. true
B. false

20. Which two are NOT valid BGP states in the BGP session establishment? (Choose two) (Source: Implementing Basic BGP Routing)

A. idle
B. connect
C. disconnect
D. open
E. close
F. sent
G. confirm
H. established

21. Which state indicates that the router does not have a path to the neighbor IP address? (Source: Implementing Basic BGP Routing)

A. active
B. idle
C. established
D. open confirm

https://t.me/learningnets

22. Which Cisco IOS or IOS XE **clear ip bgp** command is the least intrusive for resetting a BGP session after changing the outbound policy for neighbor 10.5.5.5? (Source: Implementing Basic BGP Routing)

A. **clear ip bgp ***
B. **clear ip bgp 10.5.5.5 soft out**
C. **clear ip bgp 10.5.5.5**
D. **clear ip bgp 10.5.5.5 soft in**

23. Number the 11 BGP selection criteria in the order in which they are used, from the first to the last, when selecting the BGP path that is submitted to the IP routing table. (Source: Implementing Basic BGP Routing)

prefer the lowest origin code (IGP < EGP < incomplete)	1
prefer the lowest MED	2
prefer the oldest route for EBGP paths	3
prefer the path through the closest IGP neighbor	4
prefer the lowest neighbor IP address	5
prefer the shortest AS path	6
prefer the path with the lowest neighbor BGP router ID	7
prefer the highest weight	8
prefer the highest local preference	9
prefer an EBGP path over an IBGP path	10
prefer the route that was originated by the local router	11

24. Which description applies to the local preference attribute? (Source: Implementing Basic BGP Routing)

A. well-known mandatory
B. well-known discretionary
C. optional transitive
D. optional nontransitive

25. Which description applies to the MED attribute? (Source: Implementing Basic BGP Routing)

A. Well-known mandatory
B. Well-known discretionary
C. Optional transitive
D. Optional nontransitive

https://t.me/learningnets

26. Which description applies to the weight attribute? (Source: Implementing Basic BGP Routing)

A. Weight is a well-known mandatory BGP attribute.
B. Weight is a well-known discretionary BGP attribute.
C. Weight is an optional transitive BGP attribute.
D. Weight is proprietary to Cisco and not advertised to other BGP routers.

27. Which two statements regarding local preference are true? (Choose two.) (Source: Implementing Basic BGP Routing)

A. The higher value for local preference is preferred.
B. Local preference is used only between EBGP neighbors.
C. The lower value for local preference is preferred.
D. Local preference is used only between IBGP neighbors.

28. Which two statements regarding weight are true? (Choose two.) (Source: Implementing Basic BGP Routing)

A. The lower value for weight is preferred.
B. The higher value for weight is preferred.
C. Weight is used only between IBGP neighbors.
D. Weight is used only locally inside the router.

29. Which two statements regarding the MED are true? (Choose two.) (Source: Implementing Basic BGP Routing)

A. The higher value for the MED is preferred.
B. The lower value for the MED is preferred.
C. The MED is exchanged between autonomous systems.
D. The MED is local to an AS.

30. Which two statements regarding the AS path are true? (Choose two.) (Source: Implementing Basic BGP Routing)

A. The shorter AS path is preferred.
B. The longer AS path is preferred.
C. The AS path is prepended and exchanged between autonomous systems.
D. The AS path is local to an AS.

31. Which Cisco IOS or IOS XE or IOS XR router BGP command changes the MED for all routes? (Source: Implementing Basic BGP Routing)

A. **bgp med** *number*
B. **default-metric** *number*
C. **bgp default-metric** *number*
D. **set med** *number*

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

32. The MED is used to decide how to enter an AS from neighboring autonomous systems, when multiple paths exist between two autonomous systems. (Source: Implementing Basic BGP Routing)

A. true
B. false

33. The MED is set inbound when a router is receiving router updates from an EBGP neighbor. (Source: Implementing Basic BGP Routing)

A. true
B. false

https://t.me/learningnets

©2014 Cisco Systems, Inc.

https://t.me/learningnets

Module Self-Check Answers

Answer Key

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

https://t.me/learningnets

prefer the highest local preference — 2

prefer the route that was originated by the local router — 3

prefer the shortest AS path — 4

prefer the lowest origin code (IGP < EGP < incomplete) — 5

prefer the lowest MED — 6

prefer an EBGP path over an IBGP path — 7

prefer the path through the closest IGP neighbor — 8

prefer the oldest route for EBGP paths — 9

prefer the path with the lowest neighbor BGP router ID — 10

prefer the lowest neighbor IP address — 11

24. B
25. D
26. D
27. A, D
28. B, D
29. B, C
30. A, C
31. B
32. A
33. B

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

Glossary

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.

af-group

address family group

AFI

authority and format identifier

APNIC

Asia Pacific Network Information Center

ARIN

American Registry for Internet Numbers. A nonprofit organization established for the purpose of administrating and registrating IP numbers to the geographical areas currently managed by Network Solutions (InterNIC). Those areas include, but are not limited to, North America, South America, South Africa, and the Caribbean.

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.

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.

BDR

backup designated router.

BFD

Bidirectional Forwarding Detection

BFD

Bidirectional Forwarding Detection

https://t.me/learningnets

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

BGP 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.

CE

customer edge. Identifies the network devices, connected to a provider network, that are under the administrative control of the customer.

CIDR

classless interdomain routing. Technique supported by BGP4 and based on route aggregation. CIDR allows routers to group routes together to reduce the quantity of routing information carried by the core routers. With CIDR, several IP networks appear to networks outside the group as a single, larger entity. With CIDR, IP addresses and their subnet masks are written as four octets, separated by periods, followed by a forward slash and a two-digit number that represents the subnet mask.

CLNP

Connectionless Network Protocol

CLNS

Connectionless Network Service

CLNS

Connectionless Network Service

CPU

central processing unit. The hardware within a computer system or smartphone that carries out the instructions of a computer program by performing the basic arithmetical, logical, and input-output operations of the system.

DBD

database description.

DLCI

data-link connection identifier

DNA

DoNotAge

DR

designated router.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

DRP

distributed route processor

DSP

domain-specific part

E1

external type 1

E2

external type 2

EBGP

Exterior Border Gateway Protocol.

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.

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.

ESH

end system hello

Ethernet

Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD and run over a variety of cable types at 10 Mb/s. Ethernet is similar to the IEEE 802.3 series of standards.

FDDI

Fiber Distributed Data Interface. LAN standard, defined by ANSI X3T9.5, specifying a 100-Mb/s token-passing network using fiber-optic cable, with transmission distances of up to 2 km. FDDI uses a dual-ring architecture to provide redundancy.

FIB

forwarding information base

FP

forwarding processors

https://t.me/learningnets

Frame Relay

Industry-standard, packet-switched data link layer protocol that handles multiple virtual circuits between connected devices.

Gigabit Ethernet

Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996.

GOSIP

Government OSI Profile

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.

HO-DSP

high-order domain-specific part

IA

interarea

IANA

Internet Assigned Numbers Authority. Organization operated under the auspices of the ISOC as a part of the IAB. IANA delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains a database of assigned protocol identifiers used in the TCP/IP stack, including autonomous system numbers.

IBGP

Interior Border Gateway Protocol.

IDI

initial domain identifier

IDP

initial domain part

IDRP

Interdomain Routing Protocol

IETF

Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.

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.

https://t.me/learningnets

IIH

IS-IS Hello

INIT state

initial state.

InQ

In queue

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.

IP address

32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address.

IP SLA

(Cisco IOS) IP Service Level Agreement. Cisco IOS mechanism that provides active monitoring of synthetically generated traffic to measure performance statistics for an IP network.

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).

ISH

intermediate system hello

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.

https://t.me/learningnets

ISO

International Organization for Standardization. International organization that is responsible for a wide range of standards, including those relevant to networking. ISO developed the OSI reference model, a popular networking reference model.

ISP

Internet service provider. Company that provides Internet access to other companies and individuals.

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.

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).

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.

LSAck

link-state acknowledgment

LSDB

link-state database.

LSP

Link-State Packet

LSP

label-switched path

LSP

link-state packet

LSP

link-state packet

LSR

link-state request.

LSU

link-state update.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

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.

MD5

Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and Secure Hash Algorithm (SHA) are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. Cisco uses hashes for authentication within the IPsec framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness.

MED

multi-exit discriminator

MP-BGP

multiprotocol BGP

MP-BGPv4

Multiprotocol BGP Version 4

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 preestablished IP routing information.

MsgRcvd

Messages received

MsgSent

Messages sent

MTU

maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.

NBMA

nonbroadcast multiaccess. Term describing a multiaccess network that either does not support broadcasting (such as X.25) or in which broadcasting is not feasible (for example, an SMDS broadcast group or an extended Ethernet that is too large).

NET

network entity title

NSAP

network service access points

https://t.me/learningnets

NSEL

NSAP selector

NSF

Nonstop Forwarding

NSR

Nonstop Routing

NSSA

not-so-stubby area.

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

OutQ

Out queue

P routers

provider routers

PBR

policy-based routing.

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.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

PfxRcd

Prefix received

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.

PRC

partial route calculation

PVC

permanent virtual circuit (or connection). Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection.

RFC

Request for Comments. Document series that is used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some RFCs are humorous or historical. RFCs are available online from numerous sources.

RIB

routing information base

RIP

Routing Information Protocol. A distance-vector routing protocol that uses hop count as a routing metric.

RIPE NIC

Réseaux IP Européens Network Coordination Centre

RP

route processor

SHA

Secure Hash Algorithm.

SIA

stuck-in-active.

SLA

service level agreement.

SNP

Sequence number PDU

https://t.me/learningnets

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.

SPI

security parameter index. This is a number that, together with a destination IP address and security protocol, uniquely identifies a particular security association. When using IKE to establish the security associations, the SPI for each security association is a pseudo-randomly derived number. Without IKE, the SPI is manually specified for each security association.

SSO

Stateful Switchover

SYN-ACK

synchronization-acknowledgment.

syslog

system logging.

TblVer

Table version

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.

TCP/IP

Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks. TCP and IP are the two best-known protocols in the suite.

TLV

type, length, value.

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.

VLSM

variable-length subnet mask. Capability to specify a different subnet mask for the same network number on different subnets. VLSM can help optimize available address space.

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.

 ©2014 Cisco Systems, Inc.

https://t.me/learningnets

VPNv4

Virtual Private Network Version 4

VPNv6

Virtual Private Network Version 6

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.

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.

xDSL

Refers collectively to all types of digital subscriber lines.

https://t.me/learningnets

©2014 Cisco Systems, Inc.

https://t.me/learningnets