ROUTE

Implementing Cisco IP Routing

Volume 1

Version 1.0

Student Guide

Text Part Number: 97-2814-01

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
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.**

DISCLAIMER WARRANTY: THIS CONTENT IS BEING PROVIDED “AS IS.” CISCO MAKES AND YOU RECEIVE NO WARRANTIES IN CONNECTION WITH THE CONTENT PROVIDED HEREUNDER, EXPRESS, IMPLIED, STATUTORY OR IN ANY OTHER PROVISION OF THIS CONTENT OR COMMUNICATION BETWEEN CISCO AND YOU. CISCO SPECIFICALLY DISCLAIMS ALL IMPLIED WARRANTIES, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE. This learning product may contain early release content, and while Cisco believes it to be accurate, it falls subject to the disclaimer above.

Students, this letter describes important course evaluation access information!

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

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

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

Sincerely,

Cisco Systems Learning

Table of Contents

Volume 1

Course Introduction

Overview

Implementing Cisco IP Routing (ROUTE) v1.0 is an instructor-led training program presented by Cisco training partners to their end customers. This five-day course is designed to help students prepare for Cisco CCNP® certification. The ROUTE course is a component of the CCNP curriculum.

The ROUTE course is designed to provide professionals of medium to large network sites with information on the use of advanced routing in implementing scalability for Cisco routers that are connected to LANs and WANs. The goal is to train professionals to dramatically increase the number of routers and sites using these techniques instead of redesigning the network when additional sites or wiring configurations are added. The ROUTE training reinforces the instruction by providing students with hands-on labs to ensure they thoroughly understand how to implement advanced routing within their networks.

Student Skills and Knowledge

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

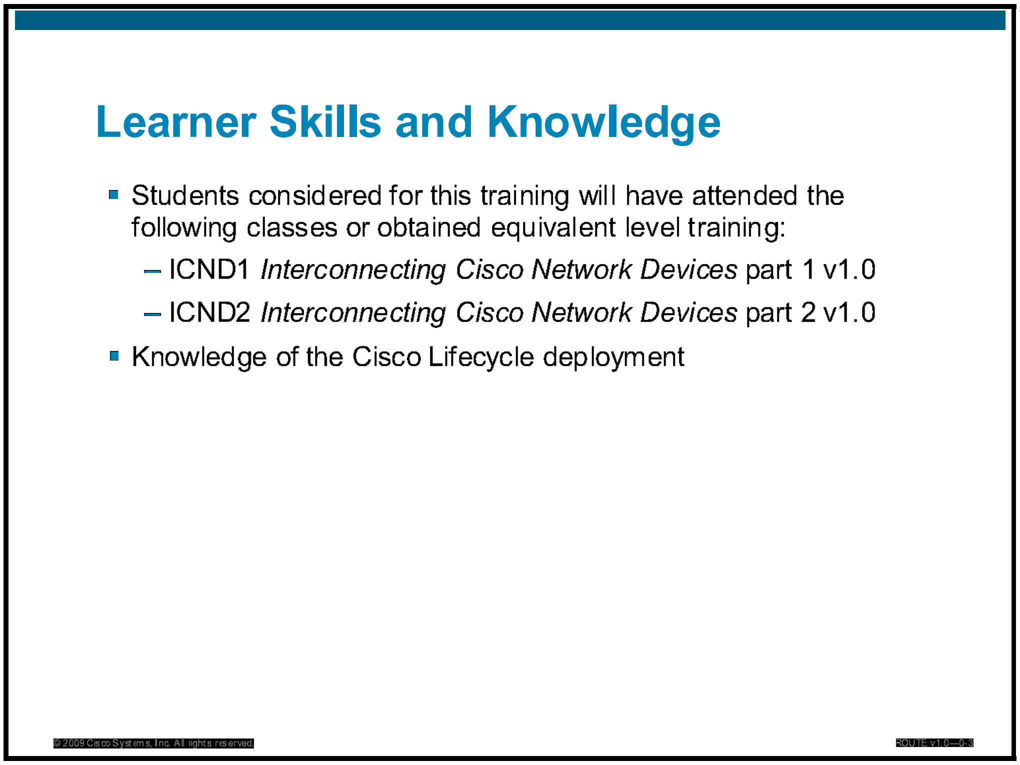

Learner Skills and Knowledge

- Students considered for this training will have attended the following classes or obtained equivalent level training:
  - ICND1 *Interconnecting Cisco Network Devices* part 1 v1.0
  - ICND2 *Interconnecting Cisco Network Devices* part 2 v1.0
- Knowledge of the Cisco Lifecycle deployment

Course Goal and Objectives

This topic describes the course goal and objectives.

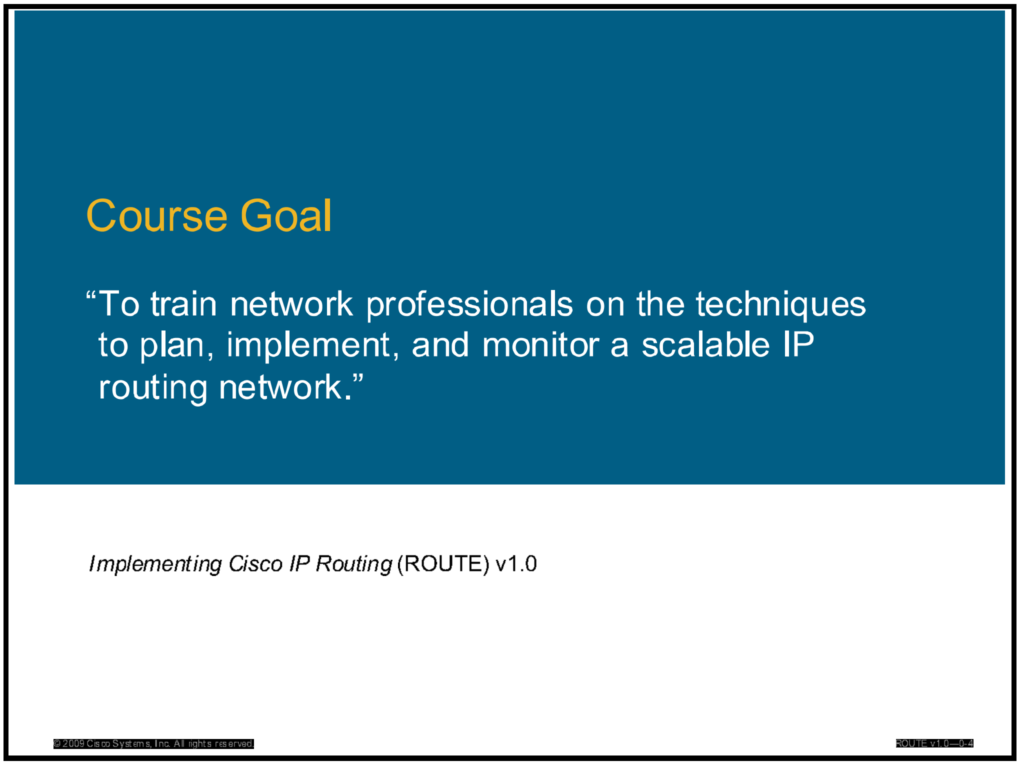

Course Goal

"To train network professionals on the techniques to plan, implement, and monitor a scalable IP routing network."

Implementing Cisco IP Routing (ROUTE) v1.0

ROUTE v1.0—0-4

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

- Plan Routing Services to meet requirements
- Implement an EIGRP-based solution
- Implement a Scalable Multiarea Network OSPF-based solution
- Implement an IPv4-based redistribution solution
- Implement Path Control
- Implement and verify a Layer 3 solution using BGP to connect an enterprise network to an Internet service provider

Course Flow

This topic presents the suggested flow of the course materials.

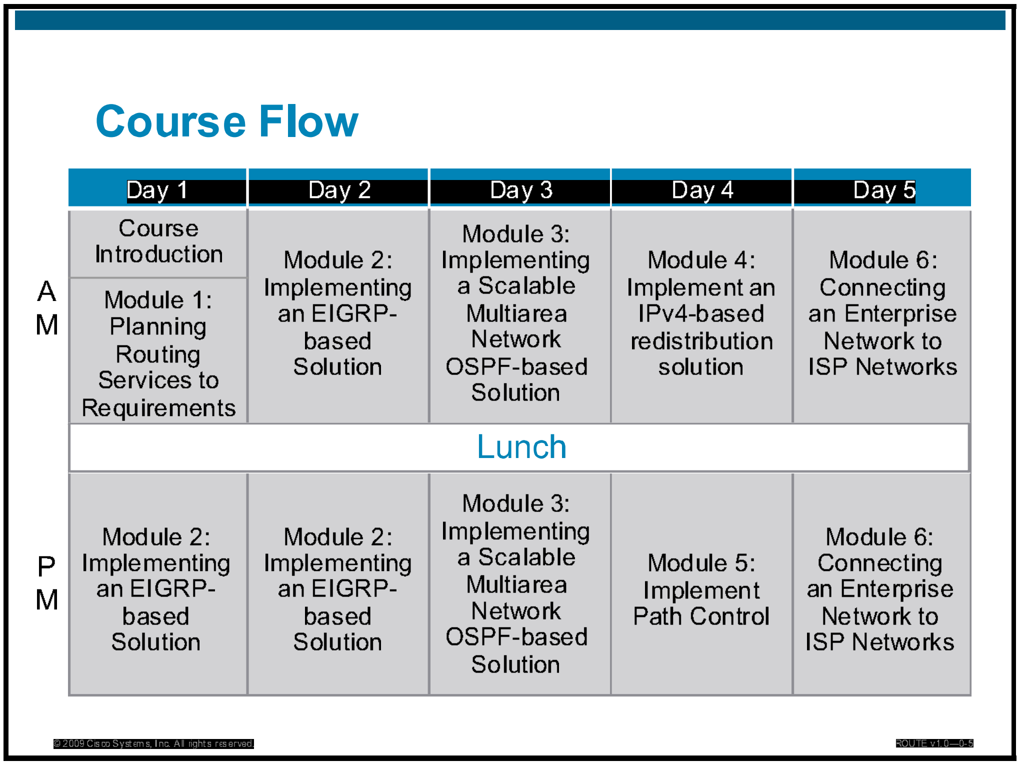

Course Flow

	Day 1	Day 2	Day 3	Day 4	Day 5
AM	Course Introduction Module 1: Planning Routing Services to Requirements	Module 2: Implementing an EIGRP-based Solution	Module 3: Implementing a Scalable Multiarea Network OSPF-based Solution	Module 4: Implement an IPv4-based redistribution solution	Module 6: Connecting an Enterprise Network to ISP Networks
	Lunch				
PM	Module 2: Implementing an EIGRP-based Solution	Module 2: Implementing an EIGRP-based Solution	Module 3: Implementing a Scalable Multiarea Network OSPF-based Solution	Module 5: Implement Path Control	Module 6: Connecting an Enterprise Network to ISP Networks

ROUTE v1.0—0-5

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.

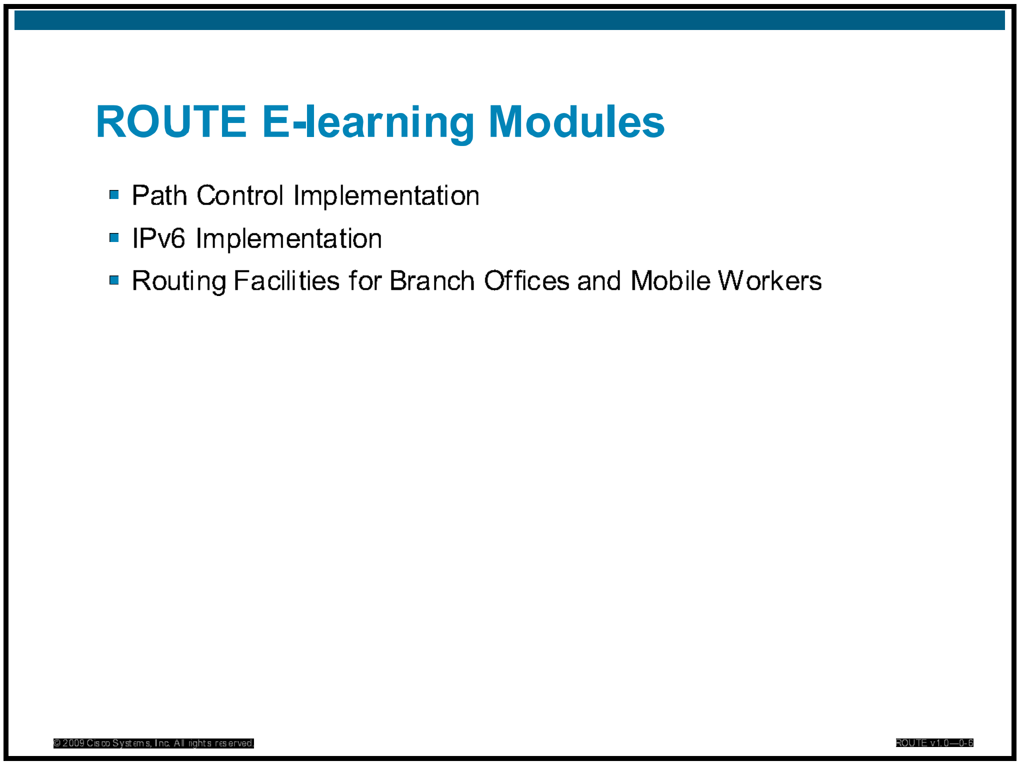

ROUTE E-learning Modules

- Path Control Implementation
- IPv6 Implementation
- Routing Facilities for Branch Offices and Mobile Workers

ROUTE v1.0—0-6

Implementing Cisco IP Routing (ROUTE) v1.0 training has three e-learning modules, which also include information required to pass the *Cisco 642-902 ROUTE certification exam*. The following modules provided:

- "Implementing Path Control"
- "Implementing IPv6"
- "Routing Facilities for Branch Offices and Mobile Workers"

The ELT content is supplied on a CD that is given out to each student, along with the other course materials.

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.

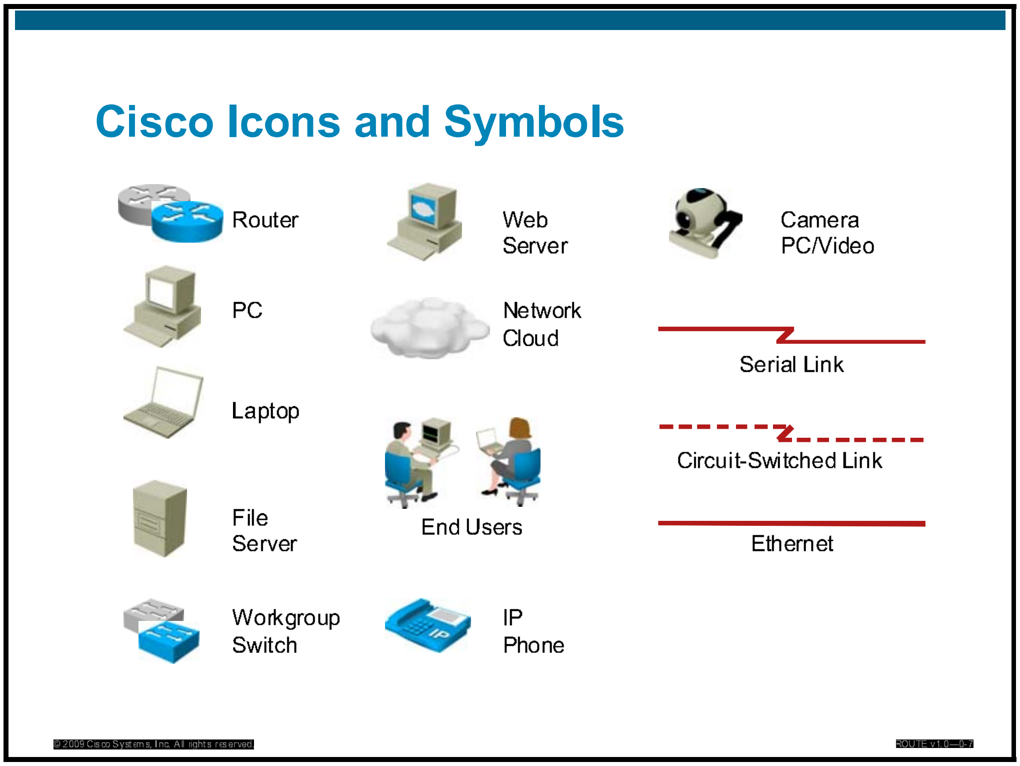

Cisco Icons and Symbols

ROUTE v1.0—0-7

Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the *Cisco Internetworking Terms and Acronyms* glossary of terms at http://www.cisco.com/en/US/docs/internetworking/terms_acronyms/CISCO12.html.

Your Training Curriculum

This topic presents the training curriculum for this course.

Cisco Career Certifications

Cisco Certifications

www.cisco.com/go/certifications

ROUTE v1.0—0-8

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®, CCVP™, 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 www.cisco.com/go/certifications.

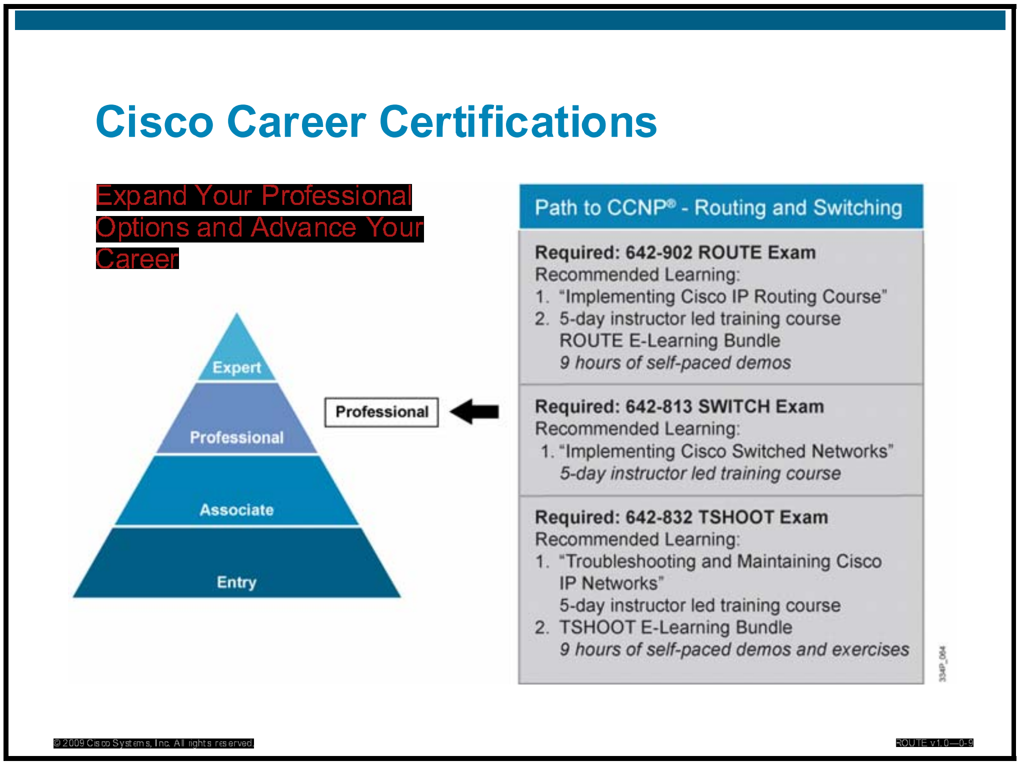

Cisco Career Certifications

Expand Your Professional Options and Advance Your Career

Path to CCNP® - Routing and Switching

Required: 642-902 ROUTE Exam
Recommended Learning:
1. "Implementing Cisco IP Routing Course"
2. 5-day instructor led training course
   ROUTE E-Learning Bundle
   9 hours of self-paced demos

Required: 642-813 SWITCH Exam
Recommended Learning:
1. "Implementing Cisco Switched Networks"
   5-day instructor led training course

Required: 642-832 TSHOOT Exam
Recommended Learning:
1. "Troubleshooting and Maintaining Cisco IP Networks"
   5-day instructor led training course
2. TSHOOT E-Learning Bundle
   9 hours of self-paced demos and exercises

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—0-9

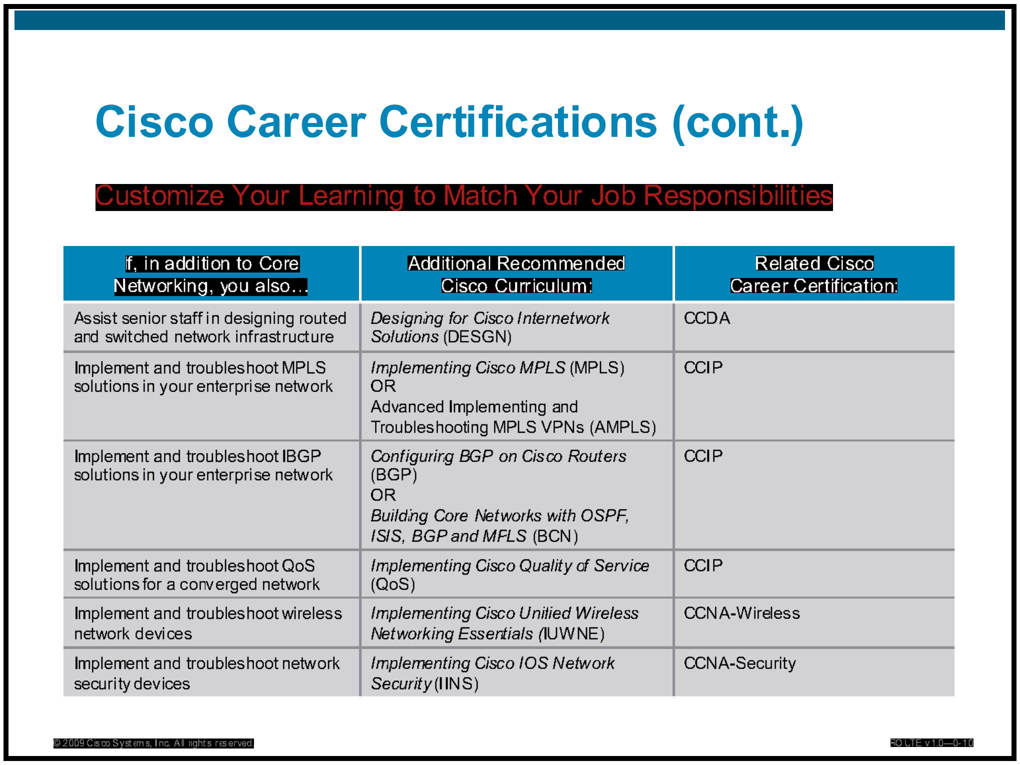

Cisco Career Certifications (cont.)

Customize Your Learning to Match Your Job Responsibilities

f, in addition to Core Networking, you also…	Additional Recommended Cisco Curriculum:	Related Cisco Career Certification:
Assist senior staff in designing routed and switched network infrastructure	*Designing for Cisco Internetwork Solutions* (DESGN)	CCDA
Implement and troubleshoot MPLS solutions in your enterprise network	*Implementing Cisco MPLS* (MPLS) OR Advanced Implementing and Troubleshooting MPLS VPNs (AMPLS)	CCIP
Implement and troubleshoot IBGP solutions in your enterprise network	*Configuring BGP on Cisco Routers* (BGP) OR *Building Core Networks with OSPF, ISIS, BGP and MFLS* (BCN)	CCIP
Implement and troubleshoot QoS solutions for a converged network	*Implementing Cisco Quality of Service* (QoS)	CCIP
Implement and troubleshoot wireless network devices	*Implementing Cisco Unified Wireless Networking Essentials* (IUWNE)	CCNA-Wireless
Implement and troubleshoot network security devices	*Implementing Cisco IOS Network Security* (IINS)	CCNA-Security

© 2009 Cisco Systems, Inc. All rights reserved.

SOLTE v1.0—0-10

General Administration

This topic presents the general administration for this course.

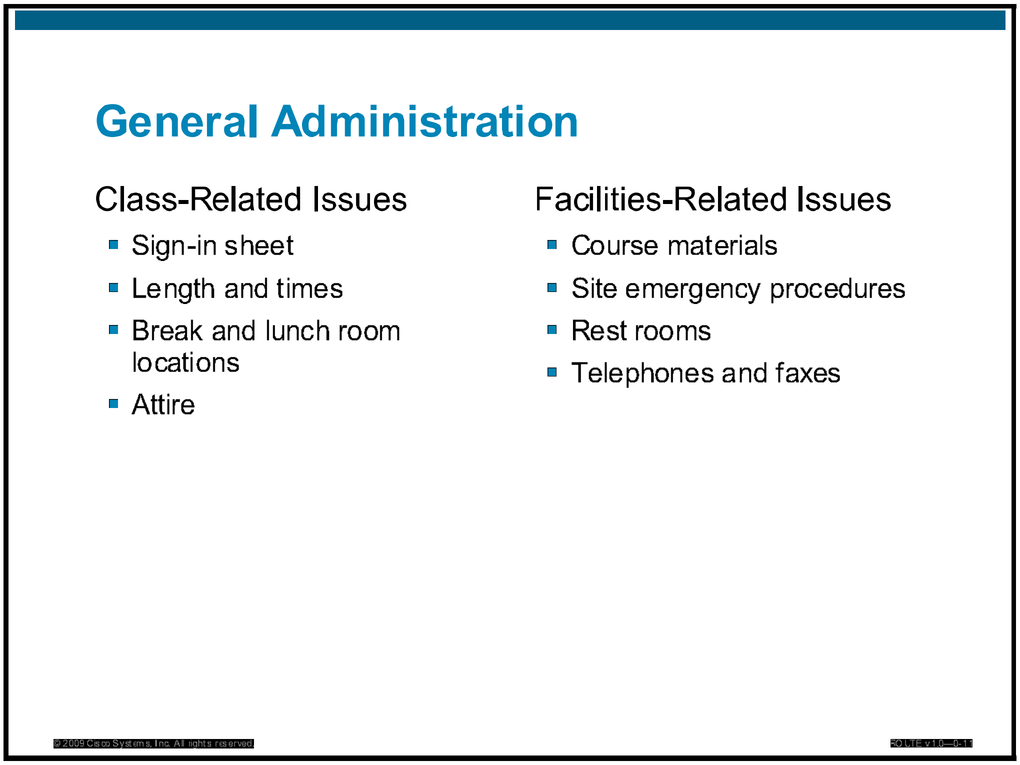

General Administration

Class-Related Issues

- Sign-in sheet
- Length and times
- Break and lunch room locations
- Attire

Facilities-Related Issues

- Course materials
- Site emergency procedures
- Rest rooms
- Telephones and faxes

The instructor will discuss the following administrative issues so that you know exactly what to expect from the class:

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

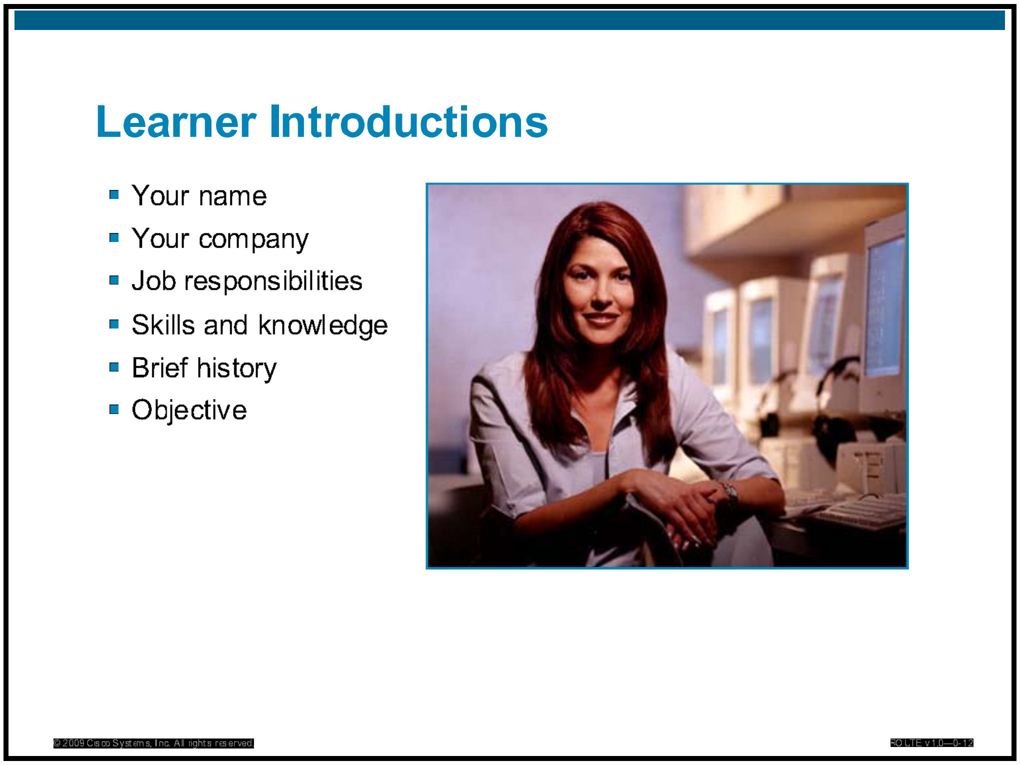

Learner Introductions

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

Prepare to share this information:

- Your name
- Your company
- Your job responsibilities
- The prerequisite skills that you have
- A profile of your experience
- What you would like to learn from this course

Module 1

Planning Routing Services to Requirements

Overview

The convergence of voice, video, and data has not only changed the conceptual network models but has also affected the way that networks support services and applications. Correct information must be identified and collected to use in the implementation plan.

This module describes Cisco conceptual models and architectures for converged networks, as well as how to build an implementation plan.

Module Objectives

Upon completing this module, you will be able to describe the converged network requirements of various network and networked applications within Cisco network architectures, including the creation of the implementation plan. This ability includes being able to meet these objectives:

- Identify the distinctive business and technical requirements of complex enterprise networks (compared to the simpler networks of CCNA)
- Assess a provided network design and to select the proper tools and resources and planning the work
- Assess a provided network design to create an implementation plan
- Review ICND2 skills and knowledge
- Discuss lab results to assess skills needed for implementing complex networks

Lesson 1

Assessing Complex Enterprise Network Requirements

Overview

This lesson starts by introducing Cisco Enterprise Architectures and describing how they align with the traditional three-layer hierarchical network model. The lesson examines the Cisco Enterprise Composite Network Model and discusses the traffic patterns in converged networks. It also introduces the Cisco vision of the future of the Intelligent Information Network (IIN) and the Cisco Service-Oriented Network Architecture (Cisco SONA). The lesson concludes with a discussion of where routing protocols fit into these models.

Objectives

Upon completing this lesson, you will be able to describe the converged network requirements of various networks and network applications within Cisco network architectures. This will result in the ability to identify the information that you must collect when filling out an implementation plan. This ability includes being able to meet these objectives:

- Define Cisco network models.
- Understand traffic conditions in a converged network.
- Understand the Cisco SONA framework.
- Understand routing and routing protocols.

Defining Cisco Network Models

This topic describes Cisco network models, starting with the Cisco Enterprise Architectures and how they map to a traditional three-layer hierarchical network model.

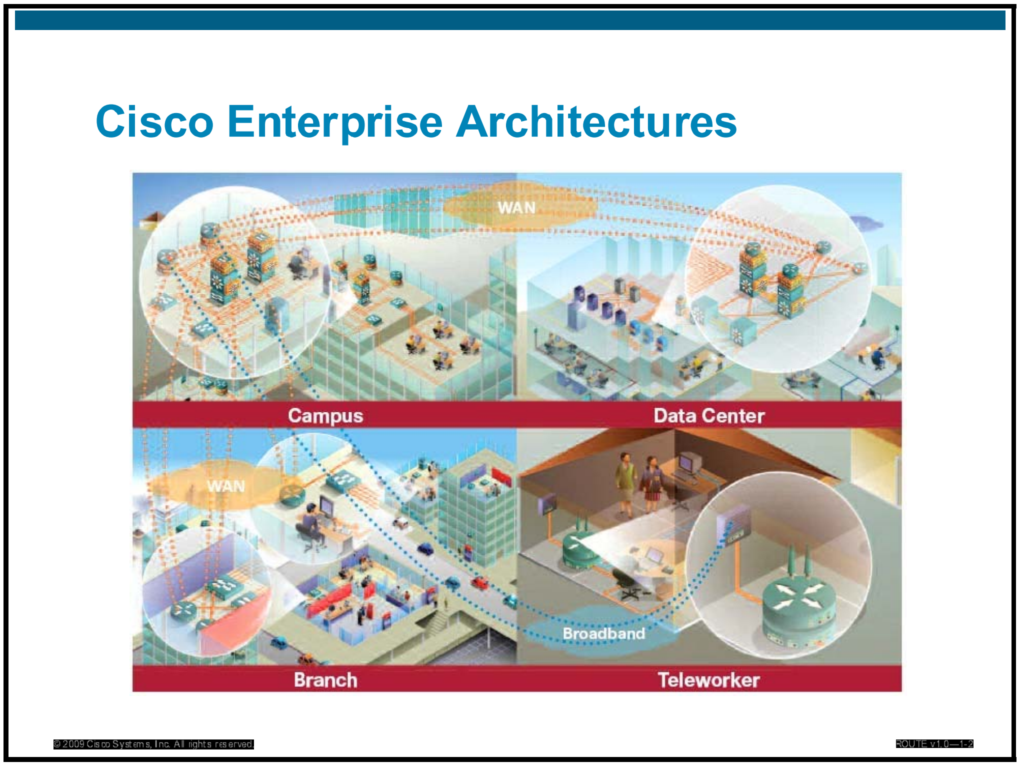

Cisco Enterprise Architectures

ROUTE v1.0—1-2

Cisco provides an enterprise-wide systems architecture that helps companies protect, optimize, and grow the infrastructure that supports their business processes. The architecture provides for the integration of the entire network—campuses, data centers, WANs, branches, and teleworkers—offering staff secure access to tools, processes, and services.

The Cisco Enterprise Campus Architecture combines intelligent switching and routing of core infrastructure with tightly integrated productivity-enhancing technologies, including IP communications, mobility, and advanced security. The architecture provides the enterprise with high availability through a resilient multilayer design, redundant hardware and software features, and automatic procedures for reconfiguring network paths when failures occur.

Multicast provides optimized bandwidth consumption and quality of service (QoS) prevents oversubscription, ensuring that real-time traffic, such as voice and video, and critical data are not dropped or delayed.

Integrated security protects against and mitigates the impact of worms, viruses, and other attacks on the network, even at the port level. The Cisco enterprise-wide architecture extends support for standards, such as 802.1x and Extensible Authentication Protocol (EAP). It also provides the flexibility to add IP Security (IPsec) and Multiprotocol Label Switching (MPLS), virtual private networks (VPNs), identity and access management, and VLANs to compartmentalize access. These features help improve performance and security and decrease costs.

The Cisco Enterprise Data Center Architecture is a cohesive, adaptive network architecture. It supports the requirements for consolidation, business continuance, and security, while enabling emerging service-oriented architectures, virtualization, and on-demand computing. IT staff can easily provide departmental staff, suppliers, or customers with secure access to applications and resources. This simplifies and streamlines management, significantly reducing overhead. Redundant data centers provide backup using synchronous and asynchronous data and application replication. The network and devices offer server and application load balancing to maximize performance. This solution allows the enterprise to scale without major changes to the infrastructure.

The Cisco Enterprise Branch Architecture allows enterprises to extend head-office applications and services, such as security, IP communications, and advanced application performance to thousands of remote locations and users or to a small group of branches. Cisco integrates security, switching, network analysis, caching, and converged voice and video services into a series of integrated services routers in the branch, so that the enterprises can deploy new services when they are ready, without buying new equipment. This solution provides secure access to voice, mission-critical data, and video applications anywhere and anytime. Advanced network routing, VPNs, redundant WAN links, application content caching, and local IP telephony call processing provide a robust architecture with high levels of resilience for all the branch offices. An optimized network leverages the WAN and LAN to reduce traffic and save bandwidth and operational expenses. The enterprise can easily support branch offices with the ability to centrally configure, monitor, and manage devices located at remote sites, including tools such as AutoQoS, which proactively resolve congestion and bandwidth issues before they affect network performance.

The Cisco Enterprise Teleworker Architecture allows enterprises to securely deliver voice and data services to small remote offices and home offices using a standard broadband access service. This ability provides a business resiliency solution for the enterprise and a flexible work environment for employees. Centralized management minimizes IT support costs, while robust integrated security mitigates the unique security challenges of this environment. Integrated security and identity-based networking services enable the enterprise to extend campus security policies to the teleworker. Staff can securely log in to the network over an "always-on" VPN and gain access to authorized applications and services from a single cost-effective platform. Productivity can further be enhanced by adding an IP phone, providing cost-effective access to a centralized IP communications system with voice and unified messaging services.

The Cisco Enterprise WAN Architecture offers the convergence of voice, video, and data services over a single IP communications network, which enables the enterprise to cost-effectively span large geographic areas. QoS, granular service levels, and comprehensive encryption options help ensure the secure delivery of high-quality corporate voice, video, and data resources to all corporate sites, enabling staff to work productively and efficiently wherever they are located. Security is provided with multiservice VPNs (IPsec and MPLS) in Layer 2 or Layer 3 WANs, hub-and-spoke topologies, or full-mesh topologies.

Cisco Hierarchical Network Model

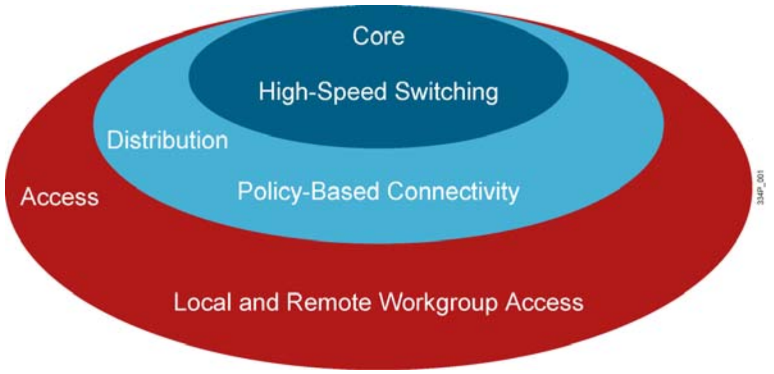

Core

High-Speed Switching

Distribution

Policy-Based Connectivity

Access

Local and Remote Workgroup Access

ROUTE v1.0—1-3

Traditionally, the three-layer hierarchical model has been used in network designs. This model provides a modular framework that allows flexibility in network design and facilitates implementation and troubleshooting. The hierarchical model divides networks or their modular blocks into access, distribution, and core layers with these features:

- **Access layer**: This layer is used to grant user access to network devices. In a network campus, the access layer generally incorporates switched LAN devices with ports that provide connectivity to workstations and servers. In a WAN environment, the access layer for remote sites or teleworkers may provide access to the corporate network across WAN technology.
- **Distribution layer:** This layer aggregates the wiring closets and uses switches to segment workgroups and isolate network problems in a campus environment. Similarly, the distribution layer aggregates WAN connections at the edge of the campus and provides policy-based connectivity.
- **Core layer (also referred to as the backbone):** This layer is a high-speed backbone and is designed to switch packets as fast as possible. Because the core is critical for connectivity, it must provide a high level of availability and adapt to changes very quickly.

Note The hierarchical model can be applied to any network type: LAN, WAN, wireless LAN (WLAN), metropolitan-area networks (MAN), VPN, or any modular block of the Cisco networking model.

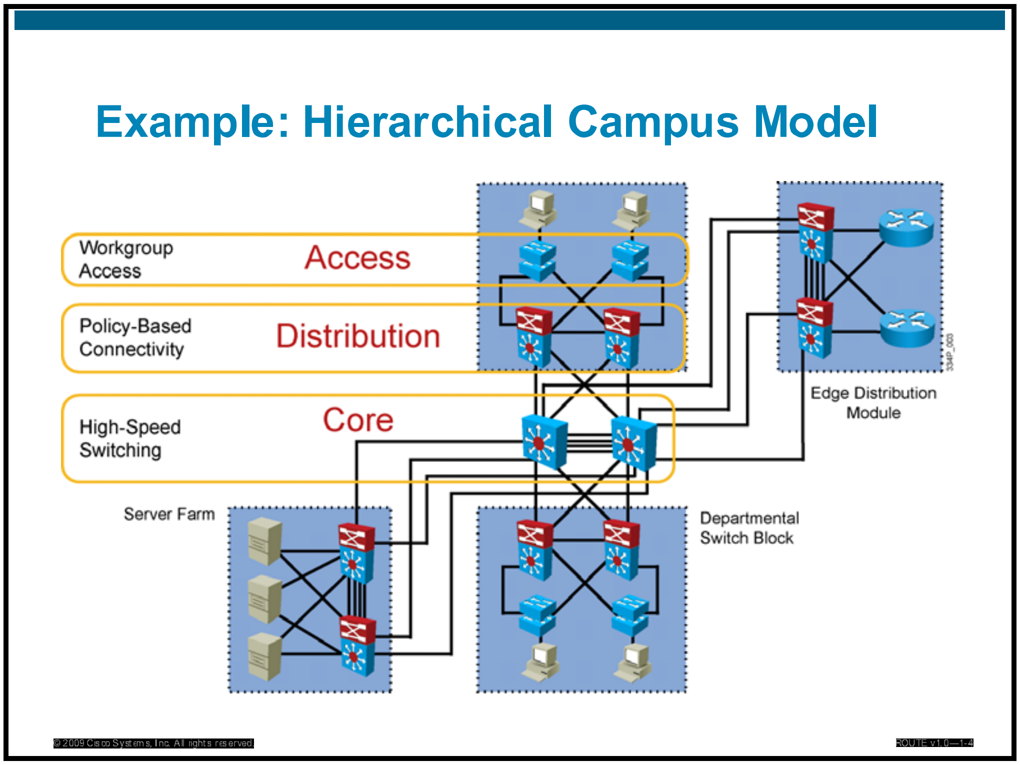

Example: Hierarchical Campus Model

ROUTE v1.0—1-4

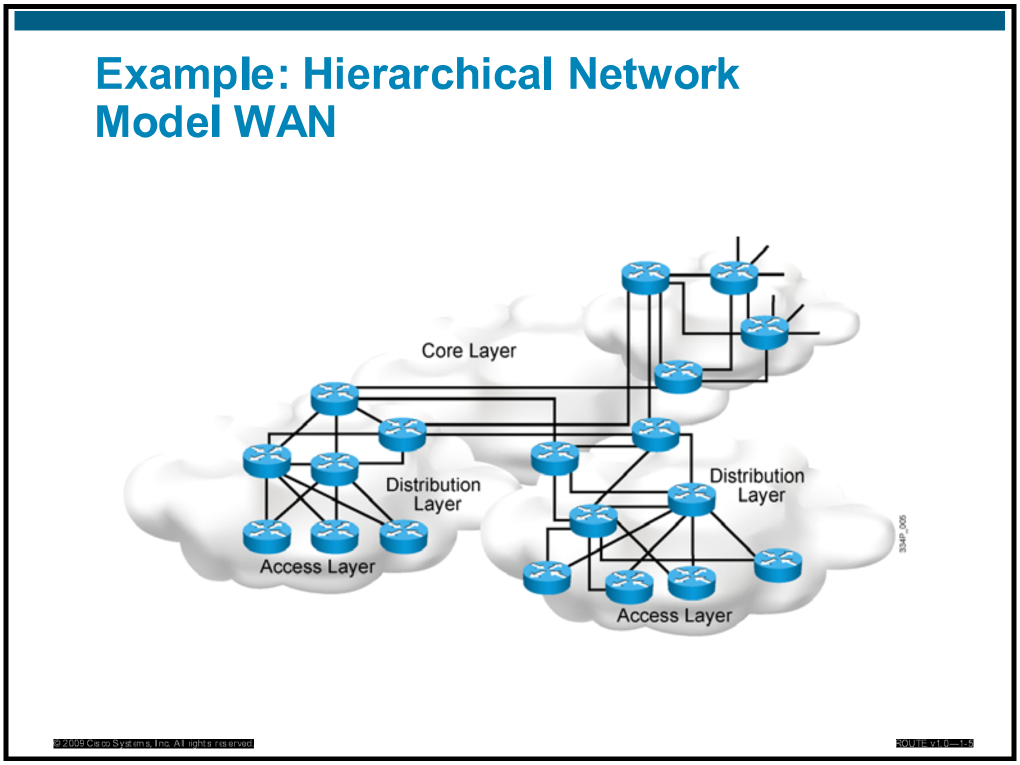

Example: Hierarchical Network Model WAN

ROUTE v1.0—1-5

For example, the hierarchical model can be applied specifically to the enterprise campus.

The hierarchical model can also be applied to the enterprise WAN. Obviously, another model is required to break down and analyze an existing modern enterprise network or plan a new one.

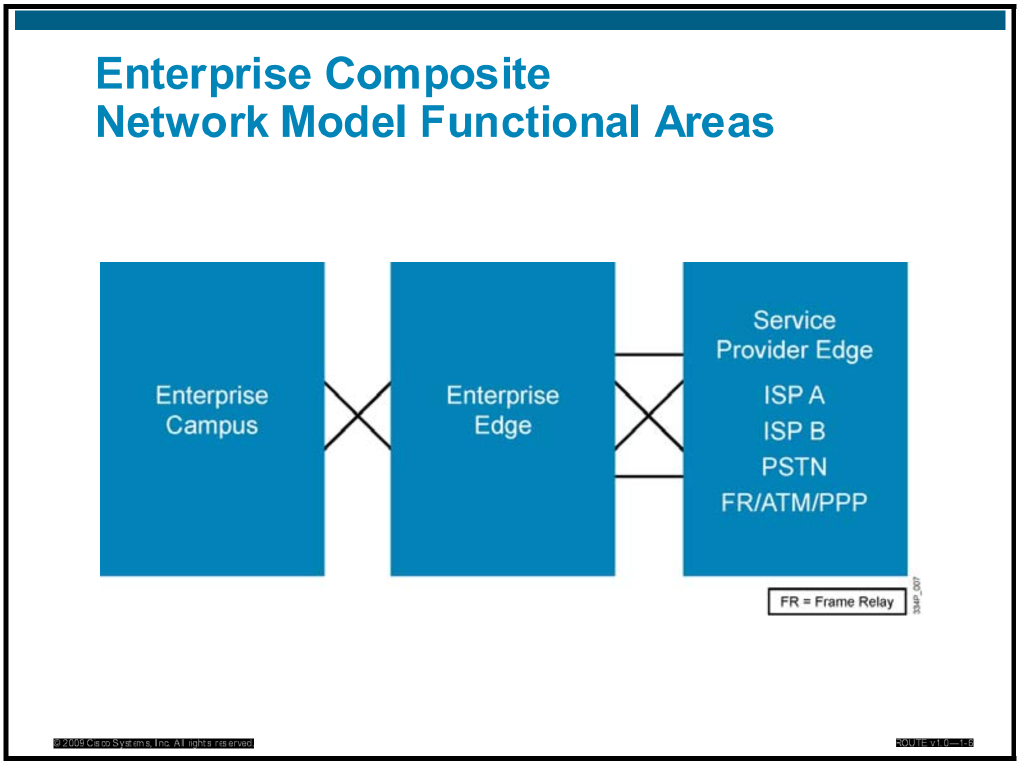

Enterprise Composite Network Model Functional Areas

Enterprise Campus

Enterprise Edge

Service Provider Edge

ISP A

ISP B

PSTN

FR/ATM/PPP

FR = Frame Relay

ROUTE v1.0—1-6

Since intelligent network service security has become of critical importance to all network planning and implementation, Cisco has developed a set of best practices for security called the Cisco SAFE Blueprint. SAFE helps network designers and administrators properly deploy security solutions to support network solutions and the existing network infrastructure.

SAFE includes the Enterprise Composite Network Model, which can be used by network professionals to describe and analyze any modern enterprise network.

Three functional areas are defined by the model:

- **Enterprise Campus**: This functional area contains the modules required to build a hierarchical, highly robust campus network. Access, distribution, and core principles are applied to these modules.
- **Enterprise Edge**: This functional area aggregates connectivity from the various elements at the edge of the enterprise network. It provides a description of connectivity to remote locations, the Internet, and remote users.
- **Service Provider Edge:** This area provides a description of connectivity to service providers such as Internet service providers (ISPs), WAN providers, and the public switched telephone network (PSTN).

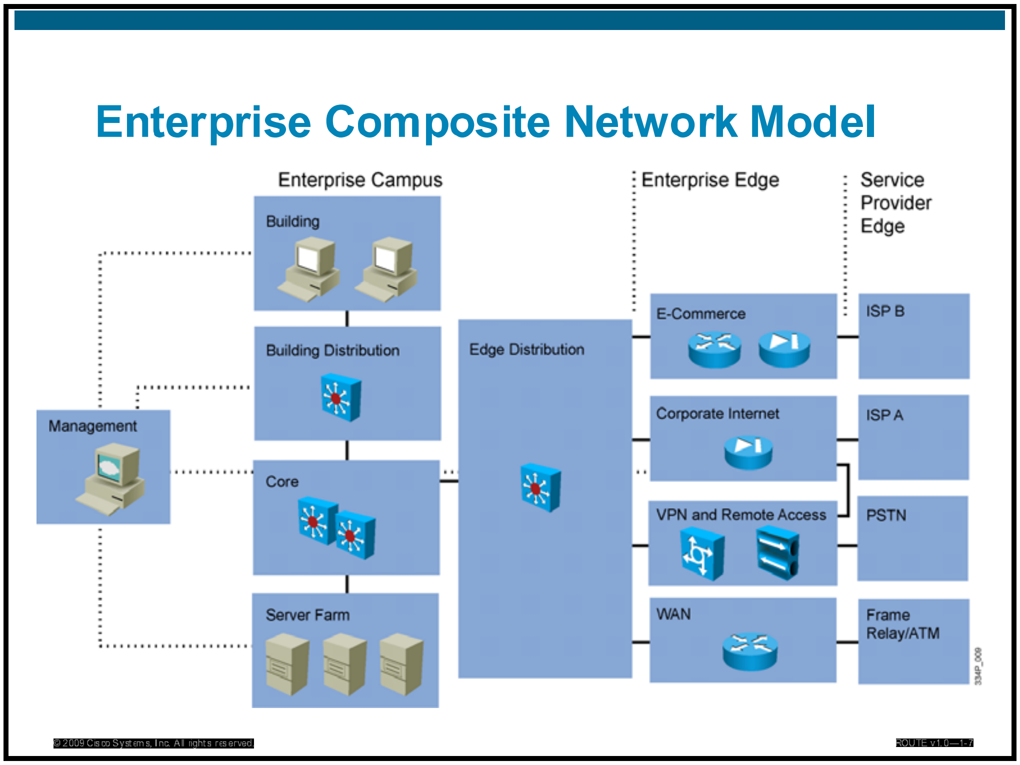

Enterprise Composite Network Model

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-7

Various modules form an integrated converged network that supports business processes.

As shown in the figure, the campus comprises six modules:

- Building, with access switches and end devices (PCs and IP phones)
- Building distribution, with distribution multilayer switches
- Core, sometimes called the backbone
- Edge distribution, which concentrates all branches and teleworkers accessing the campus via WAN or Internet
- Server farm, which represents the data center
- Management, which represents the network management functionality

Additional modules in the other functional areas represent e-commerce functionality, corporate Internet connections, remote access and VPN connections, and traditional WAN (Frame Relay, ATM, and leased lines with PPP) connections.

Traffic Conditions in a Converged Network

This topic describes the traffic types and requirements in converged networks.

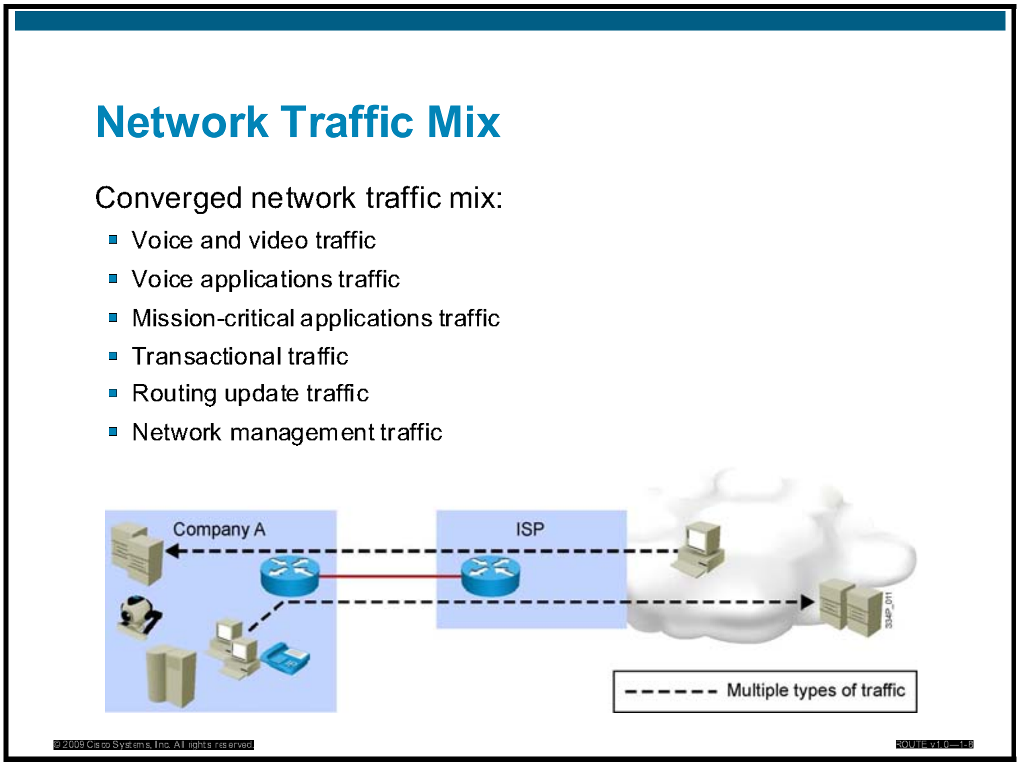

Network Traffic Mix

Converged network traffic mix:

- Voice and video traffic
- Voice applications traffic
- Mission-critical applications traffic
- Transactional traffic
- Routing update traffic
- Network management traffic

Company A
ISP
Multiple types of traffic
334P_011

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-8

Converged networks with integrated voice, video, and data contain various traffic patterns:

- Voice and video traffic, for example, IP telephony, and video broadcast and conferencing
- Voice applications traffic, generated by voice-related applications (such as contact centers)
- Mission-critical traffic, generated, for example, by stock exchange applications
- Transactional traffic, generated by e-commerce applications
- Routing update traffic, from routing protocols like Routing Information Protocol (RIP), Open Shortest Path First Protocol (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS) Protocol, and Border Gateway Protocol (BGP)
- Network management traffic

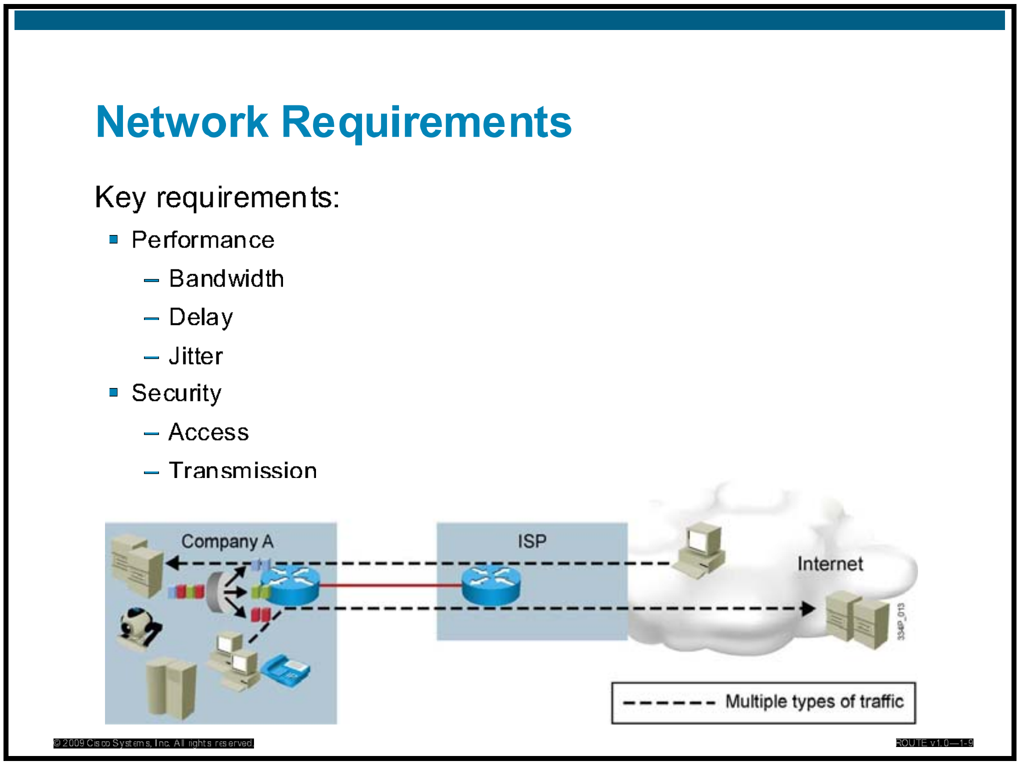

Network Requirements

Key requirements:

- Performance
  - Bandwidth
  - Delay
  - Jitter
- Security
  - Access
  - Transmission

The diversity of the traffic mix poses stringent requirements on the network in terms of performance and security. The requirements significantly differ, depending on the traffic type. For example, voice and video require constant bandwidth and low delay and jitter, while transactional traffic requires high reliability and security with relatively low bandwidth. Video traffic is frequently carried as IP multicast traffic. Also, voice applications such as IP telephony require high reliability and availability, because the user expectations for a "dial tone" in the IP network are exactly the same as in traditional phone network. To meet the traffic requirements in the network, for example, voice and video traffic must be treated differently from other traffic, such as web-based traffic. QoS mechanisms are mandatory in converged networks.

Security is a key issue not only for fixed networks but also wireless mobility, for which access to the network is possible from virtually anywhere. Several security strategies, such as device hardening with strict access control and authentication, intrusion protection, intrusion detection, traffic protection with encryption, and so on can minimize or even eliminate network security threats.

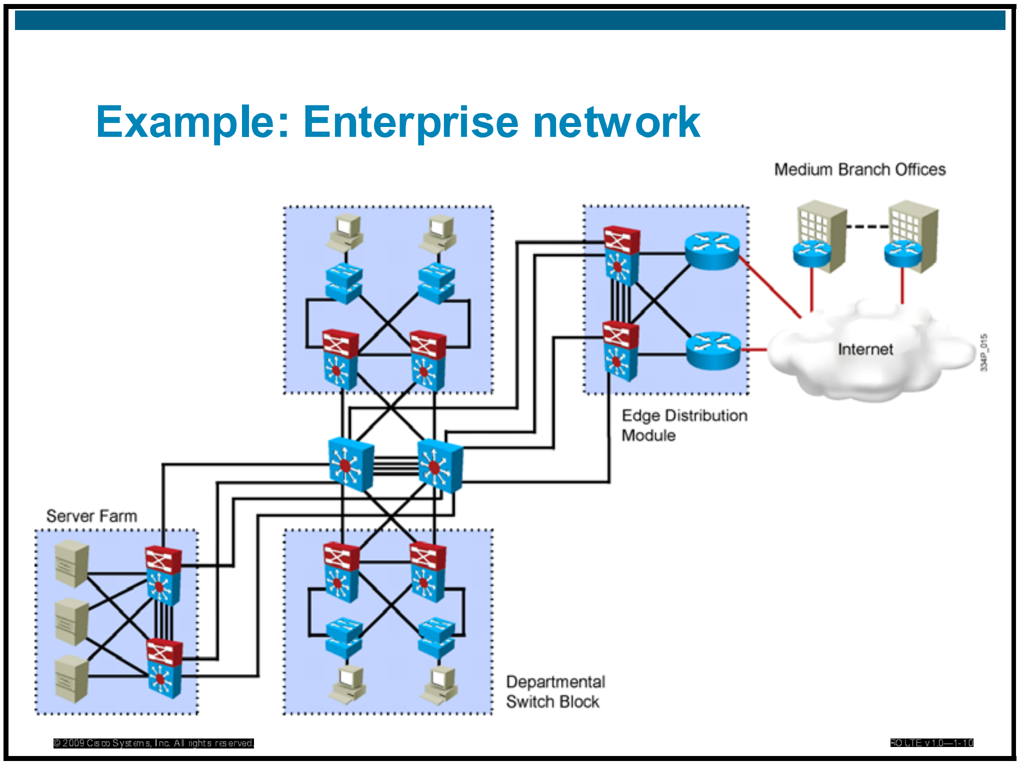

Example: Enterprise network

ROUTE v1.0—1-10

The figure above shows a hierarchical enterprise network with some remote offices. In such environments, many different traffic types exist.

IP telephony is used as well as video applications, which add a lot of delay to time-sensitive (VoIP) and bandwidth-consuming (video) traffic streams.

Server farms contain storage for mission-critical data and e-commerce applications generating transactional traffic. Traffic toward the server farm requires fast transport and bandwidth guarantees.

You must have remote and Internet-connectivity Layer 3 devices sending updates in order to efficiently route traffic.

Network design and configuration must be able to provide guaranteed services for all this traffic and satisfy the requirements for performance and security.

Cisco SONA Framework

This topic describes Cisco SONA, which guides an evolution of enterprise networks toward an IIN. The IIN and its features are also described in this section.

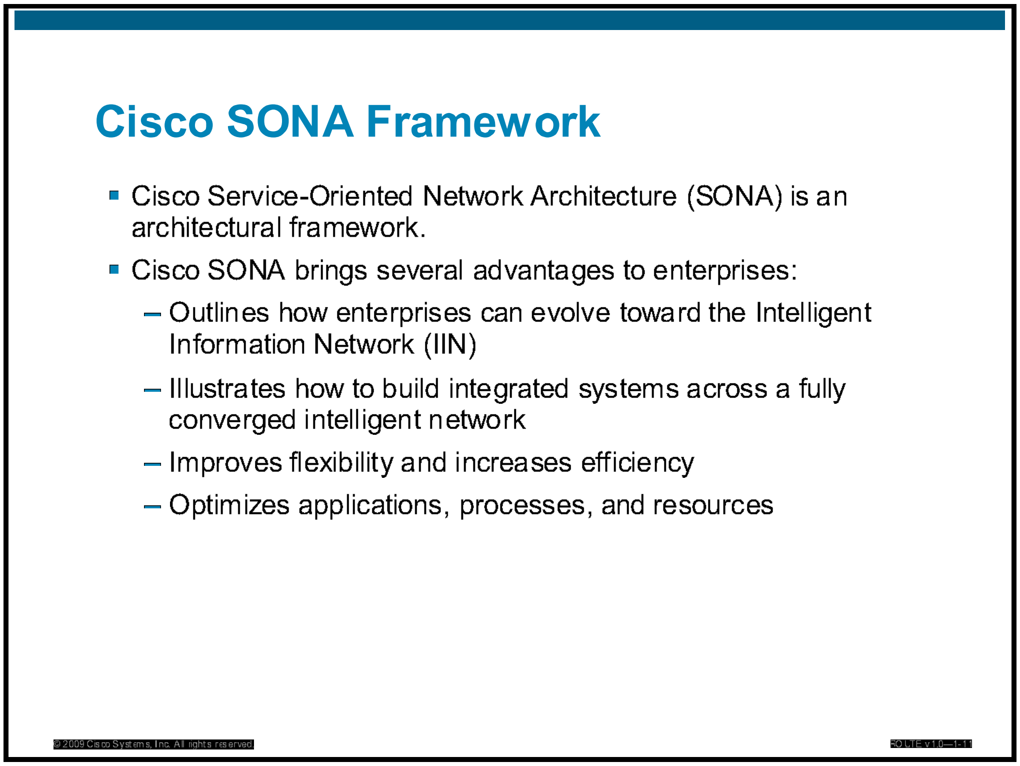

Cisco SONA Framework

- Cisco Service-Oriented Network Architecture (SONA) is an architectural framework.
- Cisco SONA brings several advantages to enterprises:
  - Outlines how enterprises can evolve toward the Intelligent Information Network (IIN)
  - Illustrates how to build integrated systems across a fully converged intelligent network
  - Improves flexibility and increases efficiency
  - Optimizes applications, processes, and resources

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—1-11

Cisco is helping organizations address new IT challenges, such as the deployment of service-oriented architectures, web services, and virtualization. Cisco SONA is an architectural framework that guides the evolution of enterprise networks to an IIN. The Cisco SONA framework provides several advantages to enterprises:

- Outlines the path toward the IIN
- Illustrates how to build integrated systems across a fully converged IIN
- Improves flexibility and increases efficiency, which results in optimized applications, processes, and resources

Cisco SONA uses the extensive product line services, proven architectures, and the experience of Cisco and its partners to help enterprises achieve their business goals.

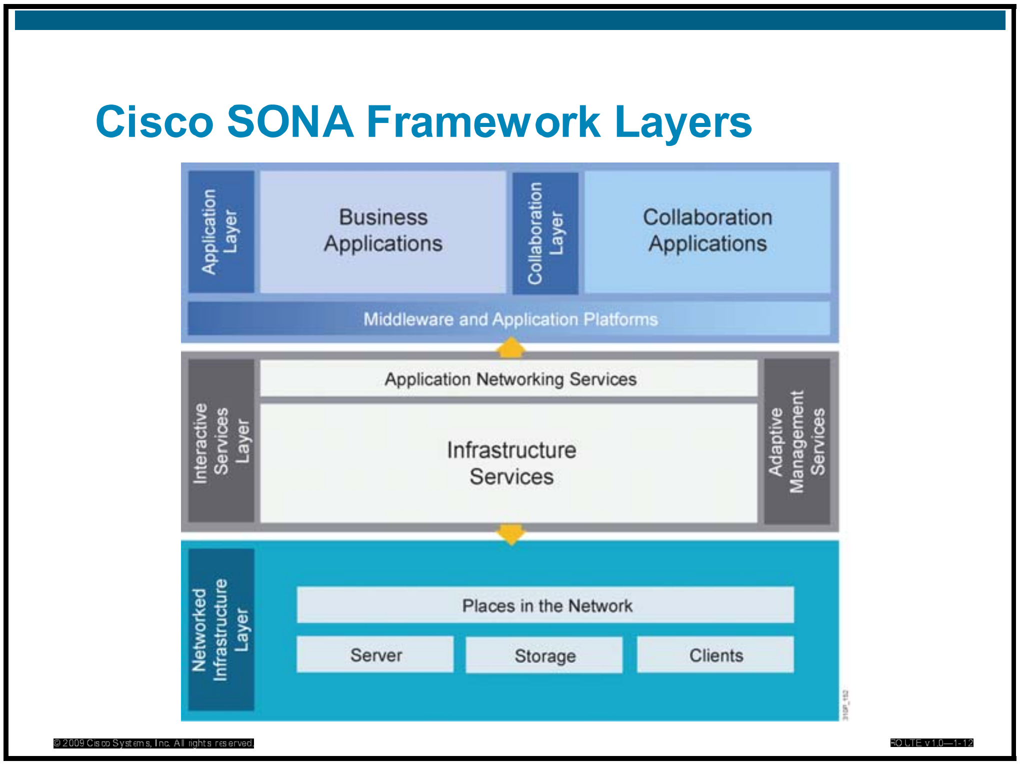

Cisco SONA Framework Layers

Application Layer

Business Applications

Collaboration Layer

Collaboration Applications

Middleware and Application Platforms

Interactive Services Layer

Application Networking Services

Infrastructure Services

Adaptive Management Services

Networked Infrastructure Layer

Places in the Network

Server

Storage

Clients

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-12

The Cisco SONA framework shows how integrated systems can both allow for a dynamic, flexible architecture and provide for operational efficiency through standardization and virtualization. It centers on the concept that the network is the common element that connects and enables all components of the IT infrastructure. Cisco SONA outlines these three layers of the IIN:

- **Networked infrastructure layer:** This layer is where all of the IT resources are interconnected across a converged network foundation. The IT resources include servers, storage, and clients. The network infrastructure layer represents how these resources exist in different places in the network, including the campus, branch, data center, WAN and MAN, and teleworker. The objective for customers in this layer is to have “anywhere and anytime” connectivity.
- **Interactive services layer**: This layer enables efficient allocation of resources to applications and business processes delivered through the networked infrastructure. This layer comprises these services:
  - Voice and collaboration services
  - Mobility services
  - Security and identity services
  - Storage services
  - Computer services
  - Application networking services
  - Network infrastructure virtualization
  - Services management
  - Adaptive management services

- **Application layer:** This layer includes business applications and collaboration applications. The objective for customers in this layer is to meet business requirements and achieve efficiencies by leveraging the interactive services layer.

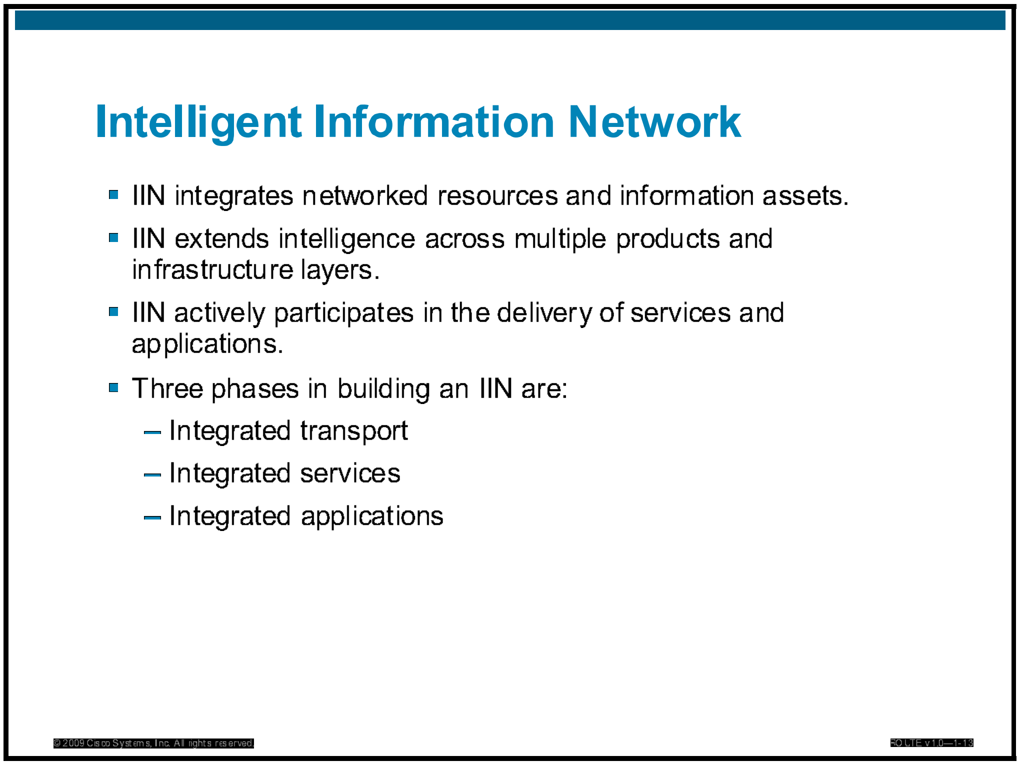

Intelligent Information Network

- IIN integrates networked resources and information assets.
- IIN extends intelligence across multiple products and infrastructure layers.
- IIN actively participates in the delivery of services and applications.
- Three phases in building an IIN are:
  - Integrated transport
  - Integrated services
  - Integrated applications

ROUTE v1.0—1-13

The Cisco vision of the future of the Intelligent Information Network (IIN) encompasses these features:

- **Integration of networked resources and information assets that have been largely unlinked**: Modern converged networks with integrated voice, video, and data require that IT departments more closely link the IT infrastructure with the network.
- **Intelligence across multiple products and infrastructure layers**: The intelligence built into each component of the network is extended network-wide and applies end to end.
- **Active participation of the network in the delivery of services and applications**: With added intelligence, the IIN makes it possible for the network to actively manage, monitor, and optimize service and application delivery across the entire IT environment.

With the listed features, the IIN offers much more than basic connectivity, bandwidth for users, and access to applications. The IIN offers end-to-end functionality and centralized, unified control that promotes true business transparency and agility.

The IIN technology vision offers an evolutionary approach that consists of three phases in which functionality can be added to the infrastructure as required:

- **Integrated transport:** Everything, including data, voice, and video, is consolidated onto an IP network for secure network convergence. By integrating data, voice, and video transport into a single, standards-based, modular network, organizations can simplify network management and generate enterprise-wide efficiencies. Network convergence also lays the foundation for a new class of IP-enabled applications delivered through Cisco Unified Communications solutions.

- **Integrated services:** Once the network infrastructure has been converged, IT resources can be pooled and shared or "virtualized" to flexibly address the changing needs of the organization. Integrated services help to unify common elements, such as storage and data center server capacity. By extending virtualization capabilities to encompass server, storage, and network elements, an organization can transparently use all of its resources more efficiently. Business continuity is also enhanced because shared resources across the IIN provide services in the event of a local systems failure.
- **Integrated applications:** With Cisco Application-Oriented Networking (AON) technology, Cisco has entered the third phase of building the IIN. This phase focuses on making the network "application-aware" so that it can optimize application performance and more efficiently deliver networked applications to users. In addition to capabilities such as content caching, load balancing, and application-level security, Cisco AON Software makes it possible for the network to simplify the application infrastructure by integrating intelligent application message handling, optimization, and security into the existing network.

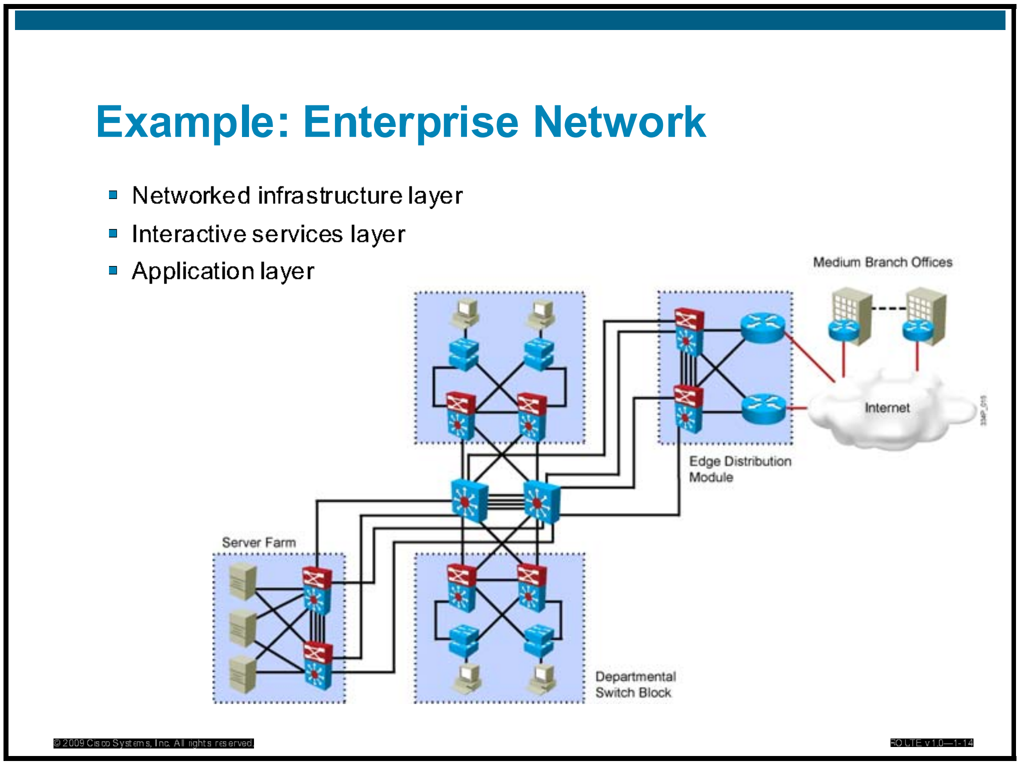

Example: Enterprise Network

- Networked infrastructure layer
- Interactive services layer
- Application layer

ROUTE v1.0—1-14

The figure above shows a hierarchical enterprise network with some remote offices. Segmentation can be done to three basic layers of SONA:

- Networked infrastructure layer
- Interactive services layer
- Application layer

The networked infrastructure layer represents the physical infrastructure—the combination of network, servers, clients, and storage hardware that is deployed throughout an enterprise network.

The interactive services layer represents the network-based functionality by making resources available to applications and business processes. Application delivery, real-time communication, management, mobility, security, transport, and virtualization are parts of the interactive services layer.

The application layer represents the enterprise software that addresses the needs of organizational processes and data flow, often in a large, distributed environment.

Routing and Routing Protocols

This topic describes routing and routing protocols.

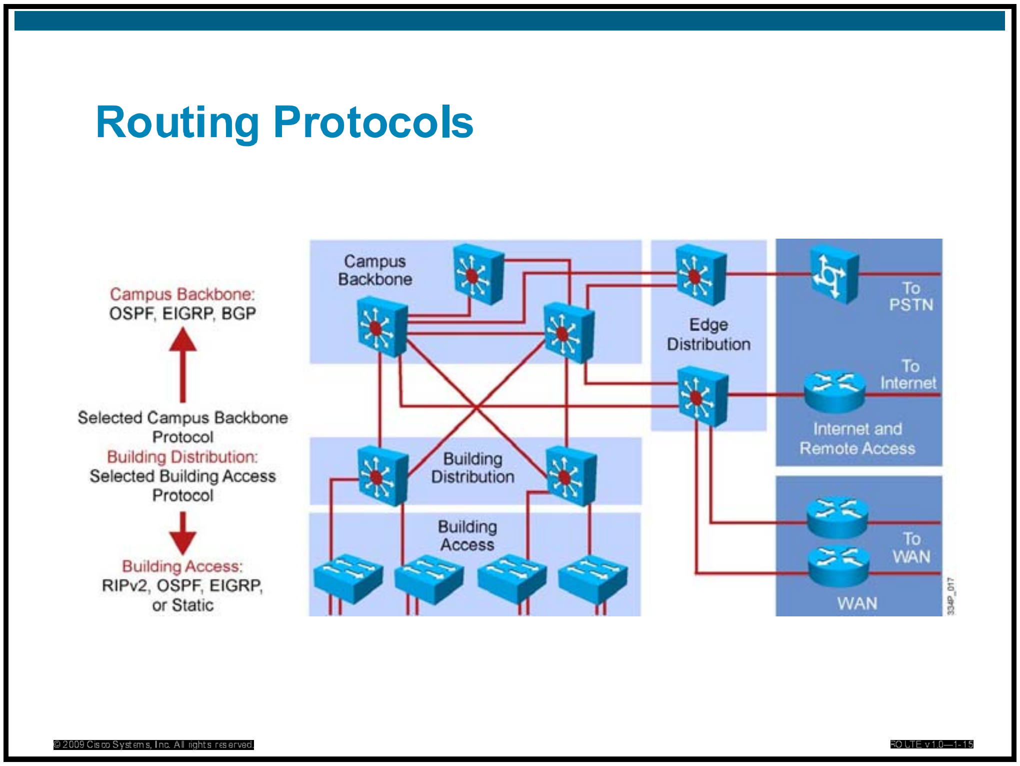

Routing Protocols

Campus Backbone: OSPF, EIGRP, BGP

Selected Campus Backbone Protocol

Building Distribution: Selected Building Access Protocol

Building Access: RIPv2, OSPF, EIGRP, or Static

ROUTE v1.0—1-15

To review, the focus of this course is on selecting, planning, implementing, tuning, and troubleshooting IP advanced routing protocols. This is a technical course at the level of Cisco CCNP®.

All of the models and tools described previously are important in the initial part of the process of selecting and planning.

The best practice is to use one IP routing protocol throughout the enterprise if possible. In many cases, this practice is not possible, which will be discussed in detail in another module. For example, Border Gateway Protocol (BGP) will be a factor in the Corporate Internet and E-Commerce modules if multihomed to ISPs is implemented. You will usually use static routes for remote access and VPN users. Therefore, you are likely to have to deal with multiple routing protocols.

The Enterprise Composite Network Model can assist in determining where each routing protocol is implemented, where the boundaries are, and how traffic flows are managed.

It is obvious that advanced IP routing protocols must be implemented in all core networks to support high availability requirements. Less advanced routing protocols (such as RIP) and static routes may exist at the access and distribution levels within modules.

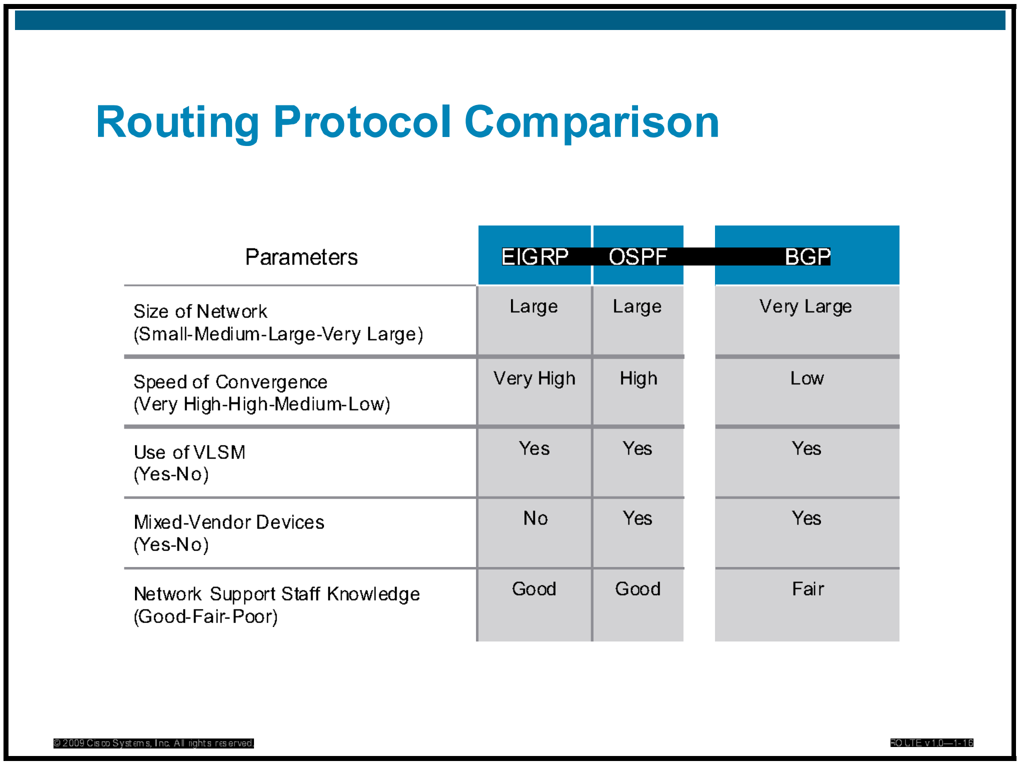

Routing Protocol Comparison

Parameters	EIGRP	OSPF	BGP
Size of Network (Small-Medium-Large-Very Large)	Large	Large	Very Large
Speed of Convergence (Very High-High-Medium-Low)	Very High	High	Low
Use of VLSM (Yes-No)	Yes	Yes	Yes
Mixed-Vendor Devices (Yes-No)	No	Yes	Yes
Network Support Staff Knowledge (Good-Fair-Poor)	Good	Good	Fair

This table provides a simple comparison of three IP routing protocols. The remainder of this course consists of technical details for each of these.

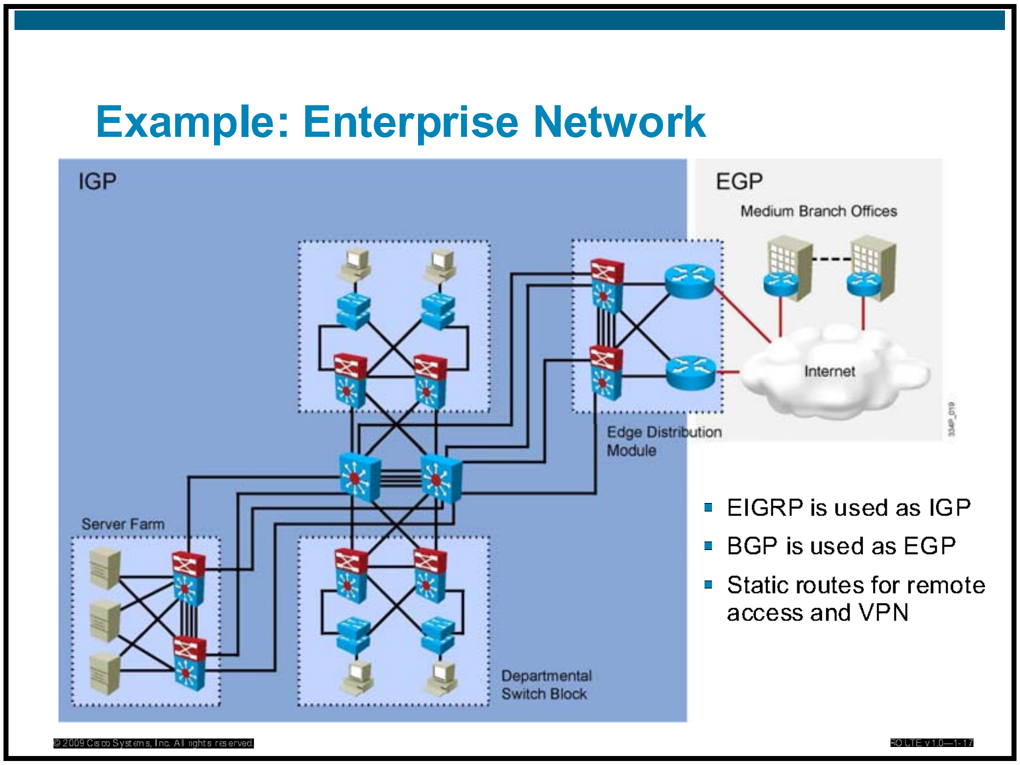

Example: Enterprise Network

- EIGRP is used as IGP
- BGP is used as EGP
- Static routes for remote access and VPN

ROUTE v1.0—1-17

Based on a best practice, one IP routing protocol has been selected throughout the whole enterprise network in the figure above. Enterprise networks usually employ an Interior Gateway Protocol (IGP) such as RIP, EIGRP, or OSPF for the exchange of routing information within their networks. EIGRP has been used in the example above, as it has very fast convergence and supports a large network size. The network in the figure above has Internet connectivity in which multihoming with multiple routers has been implemented. For such interautonomous system connectivity, an Exterior Gateway Protocol (EGP) is used. BGP is an example of an EGP protocol, is selected above. It supports very large networks and has excellent traffic policy options. Besides advanced IP routing protocols supporting high availability requirements, static routes exist at the access and distribution levels for remote and VPN access.

Summary

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

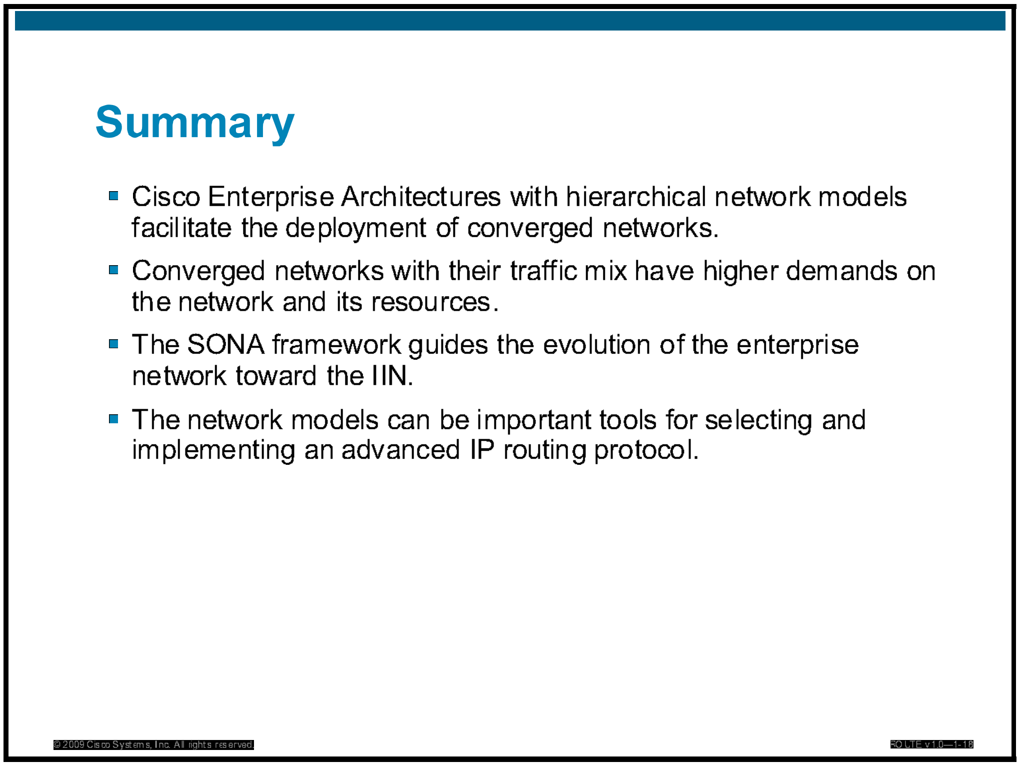

Summary

- Cisco Enterprise Architectures with hierarchical network models facilitate the deployment of converged networks.
- Converged networks with their traffic mix have higher demands on the network and its resources.
- The SONA framework guides the evolution of the enterprise network toward the IIN.
- The network models can be important tools for selecting and implementing an advanced IP routing protocol.

ROUTE v1.0—1-18

Lesson 2

Creating an Implementation Plan and Documenting the Implementation

Overview

An implementation plan and its documentation are a result of good processes and procedures during network design, implementation, and performance testing. This lesson assesses a provided network design, identifies network requirements, creates an implementation plan, and provides guidelines for creating the documentation. To create an implementation plan, you must have detailed network information, tools, resources, and a work plan. By selecting the proper tools and resources, as well as a plan of work, you can make implementation of the network is faster, more cost-effective, and capable of meeting high industry standards.

Objectives

Upon completion of this lesson, you will be able to describe the requirements of the enterprise network, implementation plan, and documentation of the implementation process, as well as describe their results. These abilities include being able to meet these objectives:

- Create an implementation plan
- Define the implementation plan tasks
- Develop the implementation plan documentation
- Create an implementation plan example

Creating an Implementation Plan

This topic describes the steps required to create a typical implementation plan, the types of information it contains and types of tasks within it.

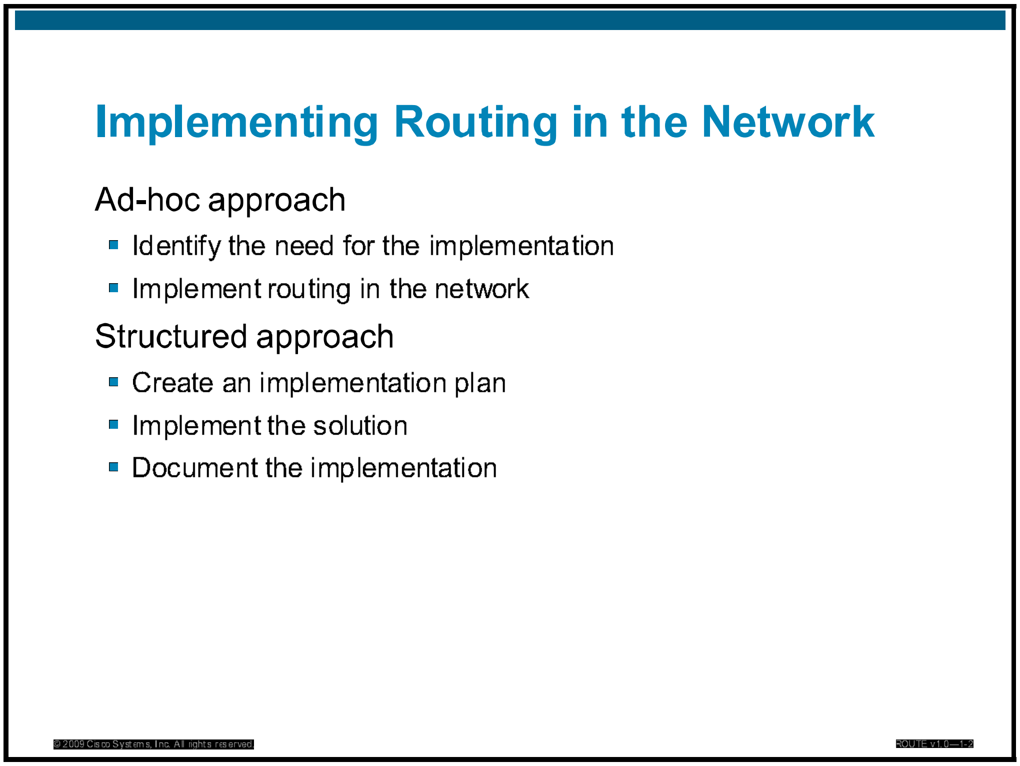

Implementing Routing in the Network

Ad-hoc approach

- Identify the need for the implementation
- Implement routing in the network

Structured approach

- Create an implementation plan
- Implement the solution
- Document the implementation

ROUTE v1.0—1-2

How is the implementation of routing protocols performed? For any other process, the following options exist:

- Ad-hoc approach
- Structured approach

In an ad-hoc approach, the network engineer identifies the need for routing protocol implementation and implements the solution without planning any of the tasks. If the size of the network is increasing, new equipment and remote offices are added to its administration. Many activities such as connectivity, routing, and security, are required. The network engineer can simply examine and configure the required functionalities as they arrive. Scalability issues, suboptimal routing, and security issues are more likely to occur with this approach. A good implementation plan is required to avoid such difficulties.

In a structured approach, the network engineer identifies the need for a routing protocol implementation and starts with planning first. Based on the existing topology, the engineer reviews all new changes, taking into account many aspects of the implementation. The engineer defines a new topology, including an IP addressing plan, scalability issues, link utilization, remote network connectivity, and other network parameters. The engineer does not review the technical aspect of the implementation only. The implementation plan must meet technical and business requirements as well. The engineer writes all details into the implementation plan documentation prior to the implementation. After the successful implementation he creates good documentation. This documentation includes the implementation plan itself, along with tools, resources, and implementation results.

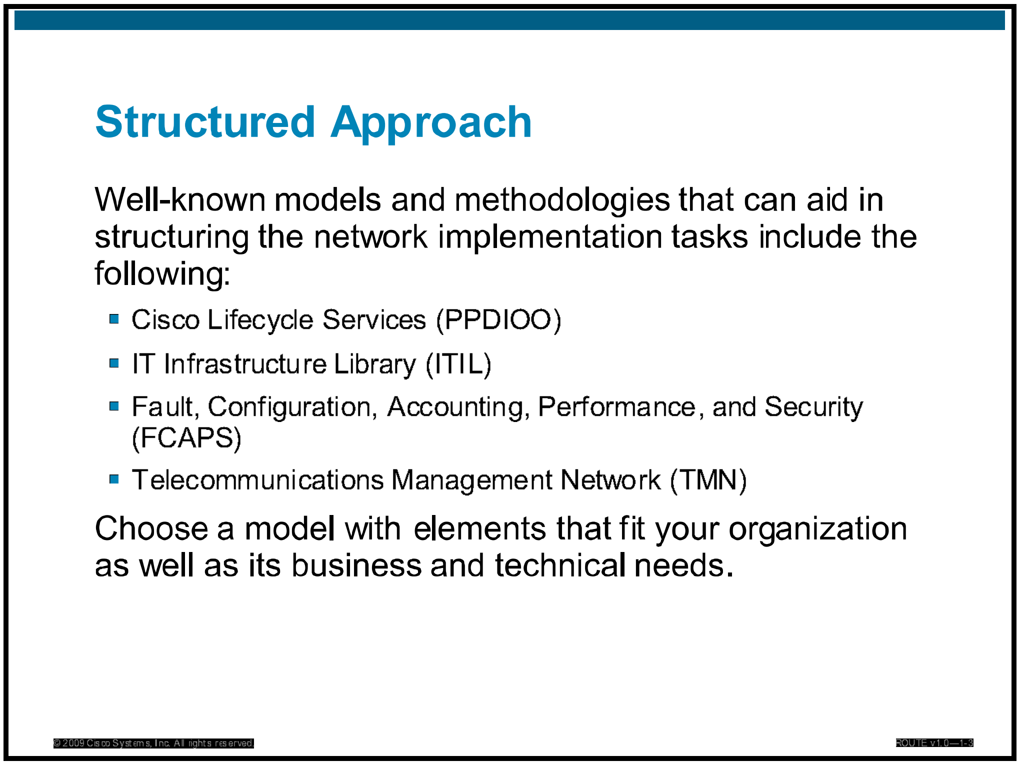

Structured Approach

Well-known models and methodologies that can aid in structuring the network implementation tasks include the following:

- Cisco Lifecycle Services (PPDIOO)
- IT Infrastructure Library (ITIL)
- Fault, Configuration, Accounting, Performance, and Security (FCAPS)
- Telecommunications Management Network (TMN)

Choose a model with elements that fit your organization as well as its business and technical needs.

The implementation plan is part of the well-known models and methodologies of every IT company, which can help structure the network implementation task. These methodologies are generic models that categorize the lifecycle approach of each process and help provide high-quality IT services.

For many models and methodologies related to network implementation, network implementation with an implementation plan is just one of the building blocks. The following models are a few good examples:

- The Cisco Lifecycle Services approach defines the minimum set of activities needed, by technology and by network complexity, to help customers successfully deploy and operate Cisco technologies and optimize their performance throughout the lifecycle of the network. This approach is referred to as the PPDIOO model based on the six phases in the network lifecycle: Prepare, Plan, Design, Implement, Operate, and Optimize. The implementation plan is part of the Design phase and the implementation itself is part of the Implement phase.
- IT Infrastructure Library (ITIL) is a framework of best practices for IT service management that provides high quality IT services. IT services are aligned with business requirements and processes in IT. The implementation plan and implementation are part of ITIL best practices.
- The Fault, Configuration, Accounting, Performance, and Security (FCAPS) model was created by the International Organization for Standardization (ISO). It defines five categories as the minimum necessary for successful network management: configuration management, fault management, accounting management, performance management and security management. Both the implementation plan and implementation are in the configuration management category.

- The Telecommunications Management Network (TMN) model is a protocol model similar to the FCAPS model and defines a framework for the management of telecommunication networks. The Telecommunications Standardization Sector (ITU-T) took the main aspects of the FCAPS model, refined them, and created a framework for which the implementation plan and implementation itself are one of the building blocks.

Note Each organization is different and has different requirements. Choose the model and elements that fit your organization and its business and technical needs.

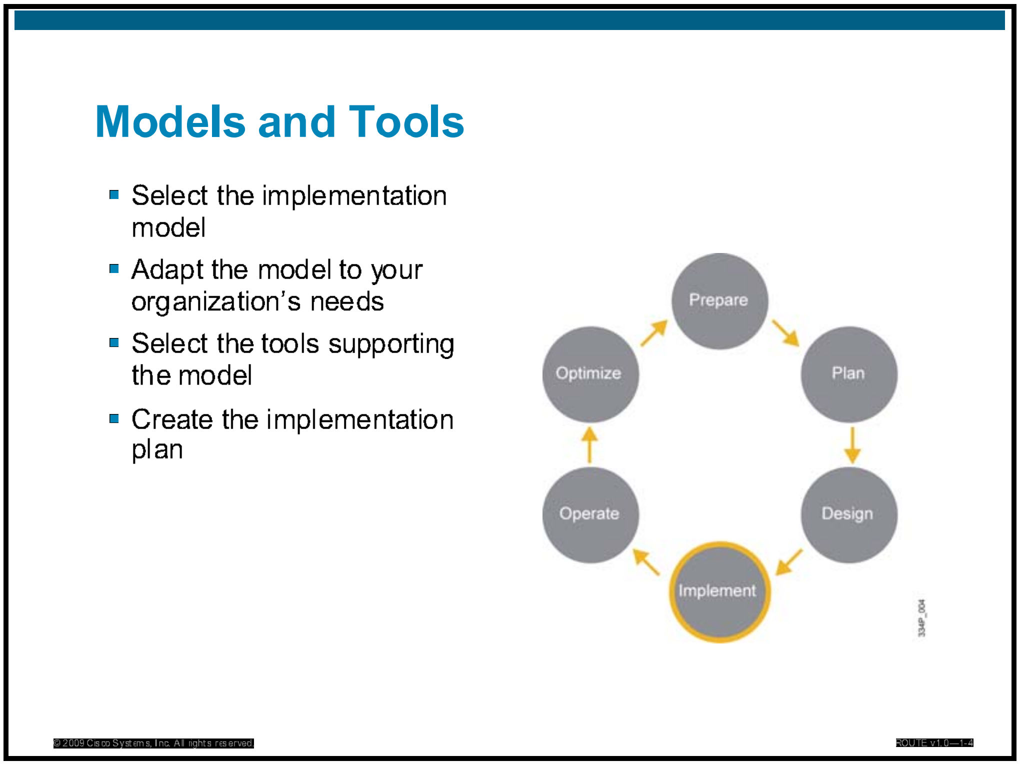

Models and Tools

- Select the implementation model
- Adapt the model to your organization’s needs
- Select the tools supporting the model
- Create the implementation plan

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-4

After you decide on a structured approach, you must choose a model and methodology.. You may combine different models to adapt the solution to fit requirements. The Cisco Lifecycle Services approach is a step-by-step approach for successfully deploying technology solutions; it will be used as an example throughout the course.

Once you have selected an implementation model, you must adapt it to the needs of your organization. If you choose service components and define processes and procedures properly when creating your implementation plan, you can produce a successful implementation.

You must select cost-effective tools to successfully deploy and optimize Cisco technologies.

Once you collect the requirements, models, and tools. you must create the implementation plan. Then you can successfully implement the solution.

Implementation Plan Tasks

This topic describes the types of tasks that are detailed in a typical implementation plan.

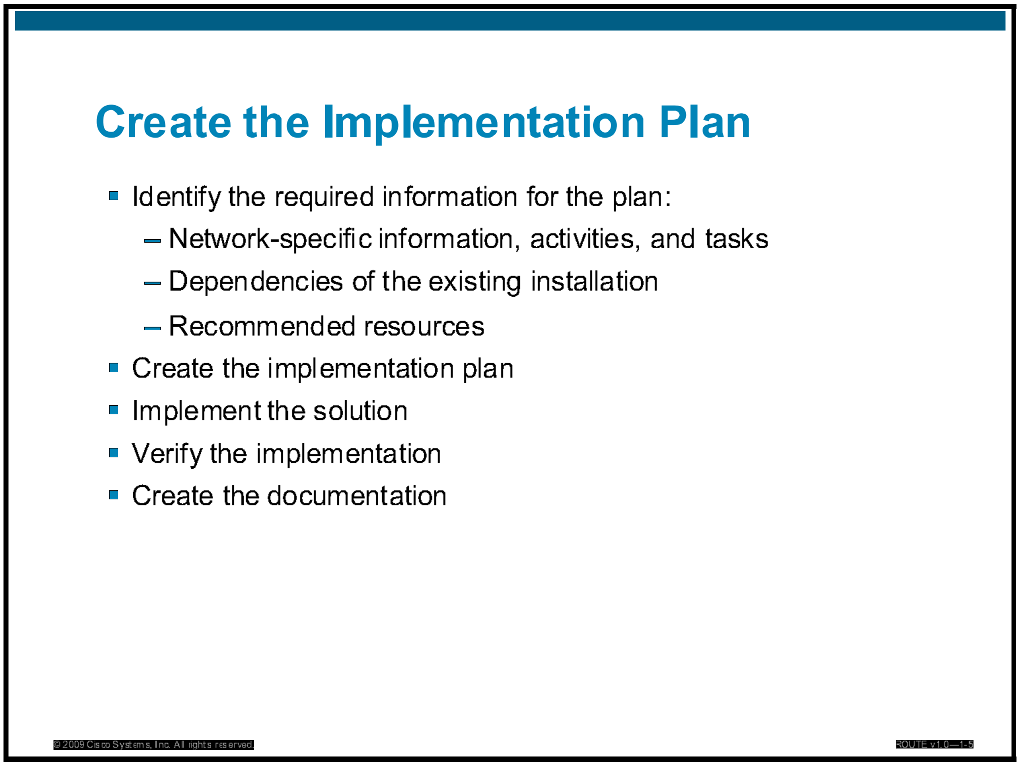

Create the Implementation Plan

- Identify the required information for the plan:
  - Network-specific information, activities, and tasks
  - Dependencies of the existing installation
  - Recommended resources
- Create the implementation plan
- Implement the solution
- Verify the implementation
- Create the documentation

ROUTE v1.0—1-5

During the design process you must clearly define the model and network and business requirements before you can create an implementation plan.

Prior to developing the implementation plan, you must identify the following required information:

- Network-specific information, activities, and tasks associated with implementation plan development
- Dependencies of your implementation plan development on other service components
- Recommended resources to accomplish the activities and tasks associated with implementation plan development

The next logical step is to develop the implementation plan, followed by implementation, verification, and creation of good documentation.

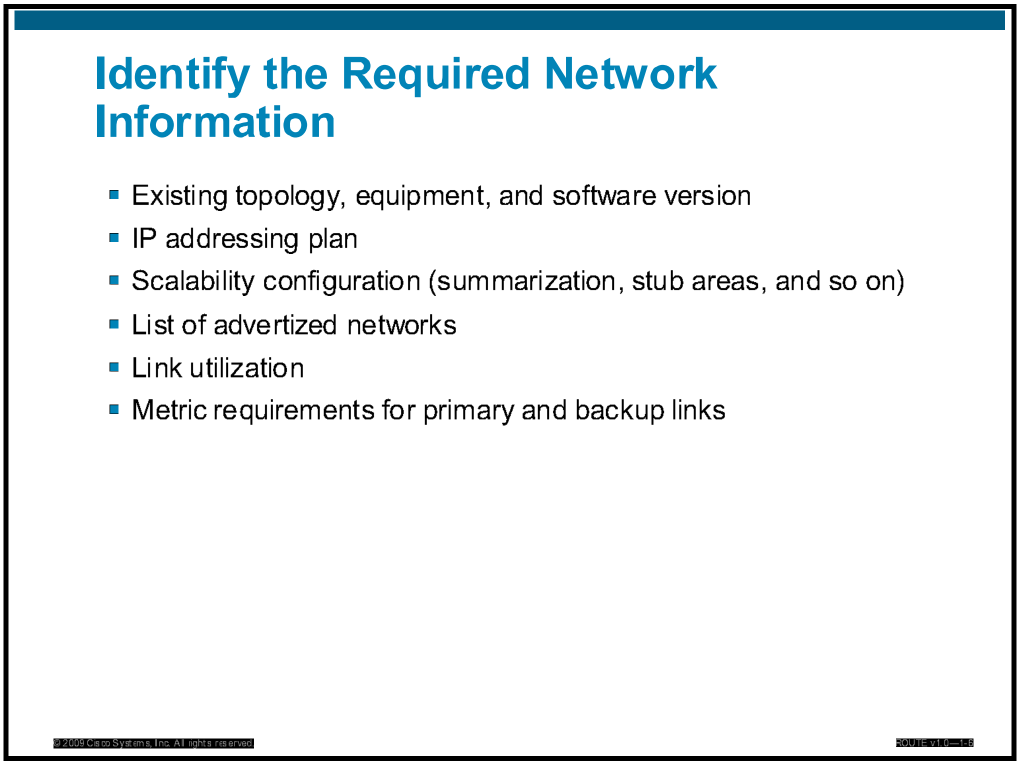

Identify the Required Network Information

- Existing topology, equipment, and software version
- IP addressing plan
- Scalability configuration (summarization, stub areas, and so on)
- List of advertized networks
- Link utilization
- Metric requirements for primary and backup links

ROUTE v1.0—1-6

One of your most important tasks is to identify network-specific information, because the implementation must support the topology and its requirements. You will likely need to change the existing network installation in order to have a successful implementation. The following network-specific information is required:

- Existing topology, equipment, software version
- IP addressing plan
- Scalability requirements (summarization, stub areas, and so on)
- List of advertised networks
- Link utilization
- Metric requirements for primary and backup links

Based on the information above, the network engineer can decide about the tasks required. The existing network may not require topology, IP addressing, or any other changes. In a best case scenario, when the network is built following the Cisco recommended design, the new implementation is just an addition to the existing network.

Identify Other Requirements

- Site specific implementation requirements
- Dependencies on existing installation
- Configuration and verification commands
- Implementation schedule and resources
- Tools

Command
#debug ip eigrp
(config-router)#eigrp stub
(config-if)#ip summary-address eigrp 1 10.x.0.0 255.255.0.0
(config-if)#ip summary-address eigrp 1 0.0.0.0 0.0.0.0
(config-router)#network 10.x.0.0 0.0.255.255
(config-router)#no auto-summary
(config)#router eigrp 1

Command List

Documentation

Backups
Schedule
Resources

CLI and GUI management

Online Resources

In addition to network information, there are many other requirements for the successful creation of an implementation plan and for implementation:

- Site-specific implementation requirements
- Dependencies on the existing installation
- Configuration and verification commands
- Implementation schedule and resources
- Tools

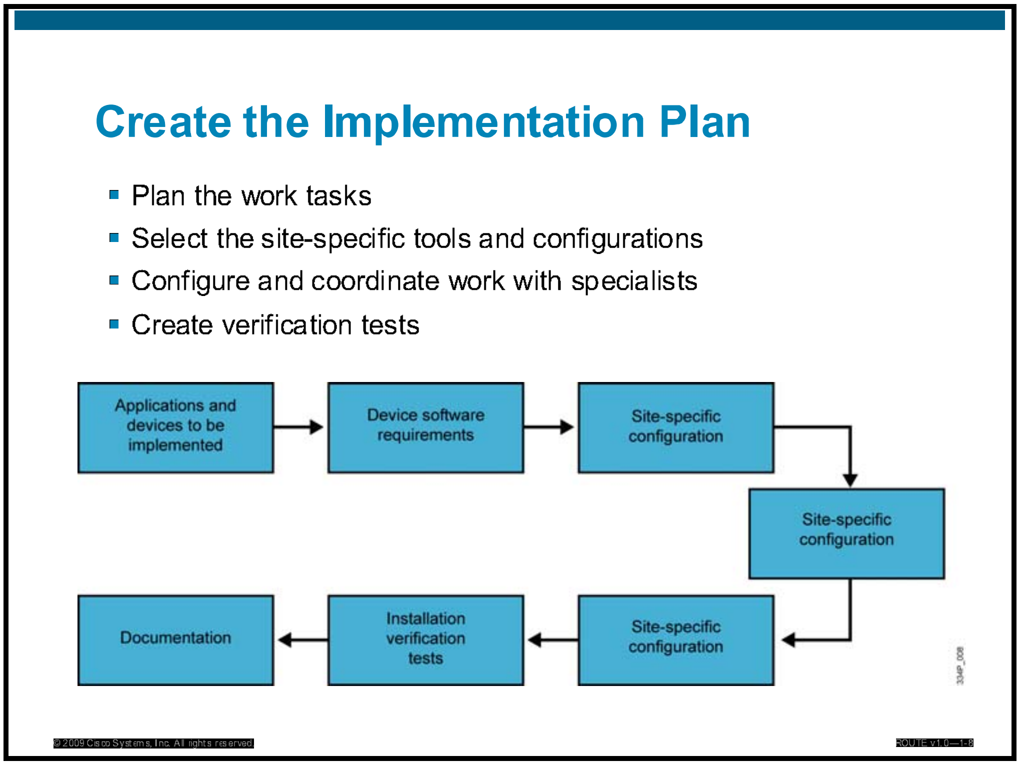

Create the Implementation Plan

- Plan the work tasks
- Select the site-specific tools and configurations
- Configure and coordinate work with specialists
- Create verification tests

Applications and devices to be implemented → Device software requirements → Site-specific configuration → Site-specific configuration → Site-specific configuration → Installation verification tests → Documentation

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—1-8

After you gather network requirements, existing documentation, implementation schedule options, identified implementation risks, management plan, and roles and responsibilities implementation, you can create the plan.

You must define and document the following tasks to create a site-specific implementation plan:

- Identify applications and devices to be implemented
- Identify installation tasks and checklists
- Create site-specific configurations
- Define site-specific installation tasks and checklists
- Define device configuration and software requirements
- Create installation verification tests

Implementation Plan Documentation

This topic describes the types of implementation information that should be documented and how to document them.

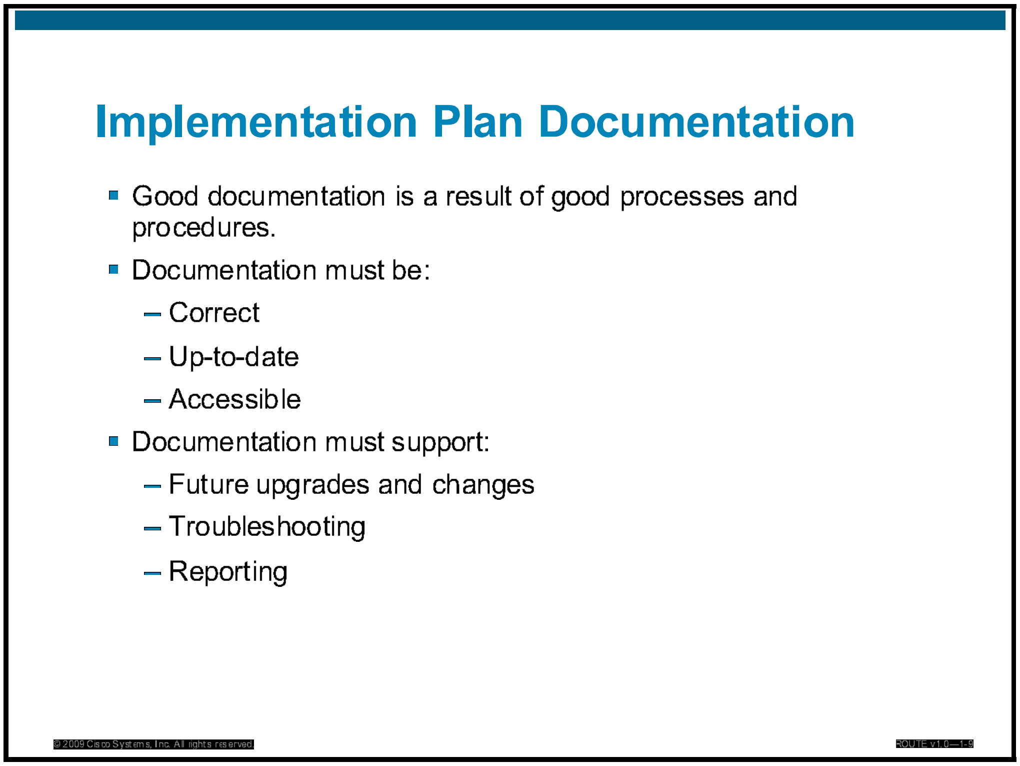

Implementation Plan Documentation

- Good documentation is a result of good processes and procedures.
- Documentation must be:
  - Correct
  - Up-to-date
  - Accessible
- Documentation must support:
  - Future upgrades and changes
  - Troubleshooting
  - Reporting

ROUTE v1.0—1-9

Creating and documenting an excellent implementation plan according to a well-known model and methodology is the first step for good documentation.

Documentation must be correct and up-to-date as you will use it during the implementation process and during verification at the end. At the end of the implementation, you will add all verification steps and results to produce documentation that is useful for future processes. The documentation will provide the last known good status of the network and, together with all the details inside, will make it easy to create an implementation plan for future changes and upgrades in the network.

At the same time, the documentation must be accessible. This is one of the requirements for a successful troubleshooting session. The documentation contains all of the information about the equipment, configuration, and known issues, as well as baseline, verification tasks, and their results. Having good documentation available to a troubleshooting engineer at any time is essential to ensuring efficient troubleshooting.

If the IT department needs to create a report, the documentation can support the information in the report, because it contains all the tasks performed, the schedule, and the resources involved.

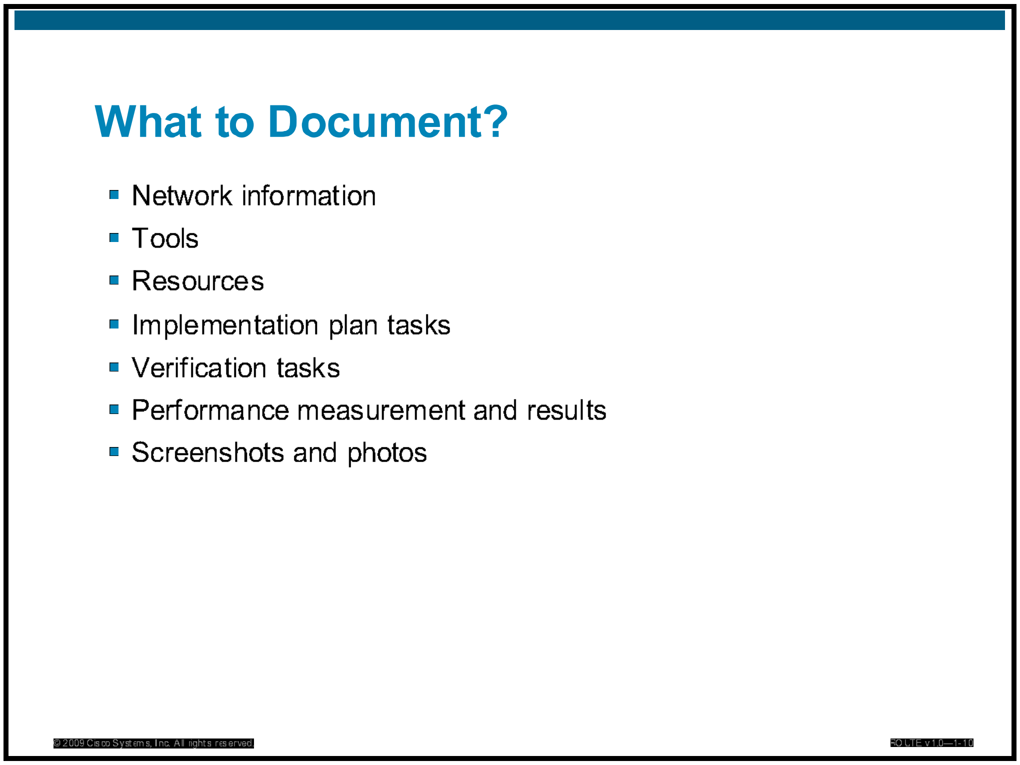

What to Document?

- Network information
- Tools
- Resources
- Implementation plan tasks
- Verification tasks
- Performance measurement and results
- Screenshots and photos

ROUTE v1.0—1-10

The documentation consists of:

- Network information
- Tools
- Resources
- Implementation plan tasks
- Verification tasks
- Performance measurement and results
- Screenshots and photos

Each part of the documentation presents its own phase of the network lifecycle implementation and verification process. The documentation creation process cannot be finished in one step; it is not finished until the end of the project. The process starts with the implementation plan, which describes all the tasks needed and ends with the verification steps.

The typical process when creating documentation includes creating a template and adding information to it during every step of the implementation process. Finally, several verification steps are required to verify that the information is correct. If a standard company template does not exist, then convert the documentation to the company standards and standard models and methodologies. At the end of the project, safely store the document, as it can be used at any time to review the network and determine when changes are required.

Implementation Plan Example

This topic describes how to assess a network design and create an implementation plan.

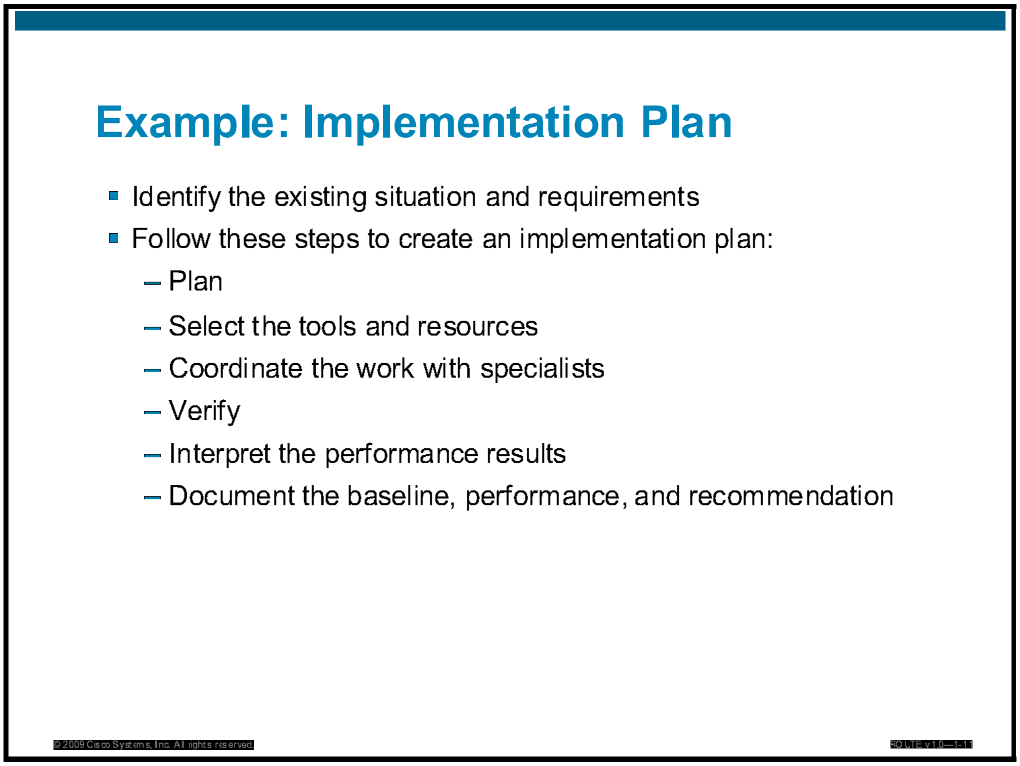

Example: Implementation Plan

- Identify the existing situation and requirements
- Follow these steps to create an implementation plan:
  - Plan
  - Select the tools and resources
  - Coordinate the work with specialists
  - Verify
  - Interpret the performance results
  - Document the baseline, performance, and recommendation

ROUTE v1.0—1-11

In order to create an implementation plan, you must define the existing situation and requirements correctly. Review the given network, select tools and resources, and create the implementation tasks required.

The following steps are required during the creation of an implementation plan:

- Plan
- Select the tools and resources
- Coordinate work with specialists
- Verify
- Interpret performance results
- Document baseline, performance, and recommendations

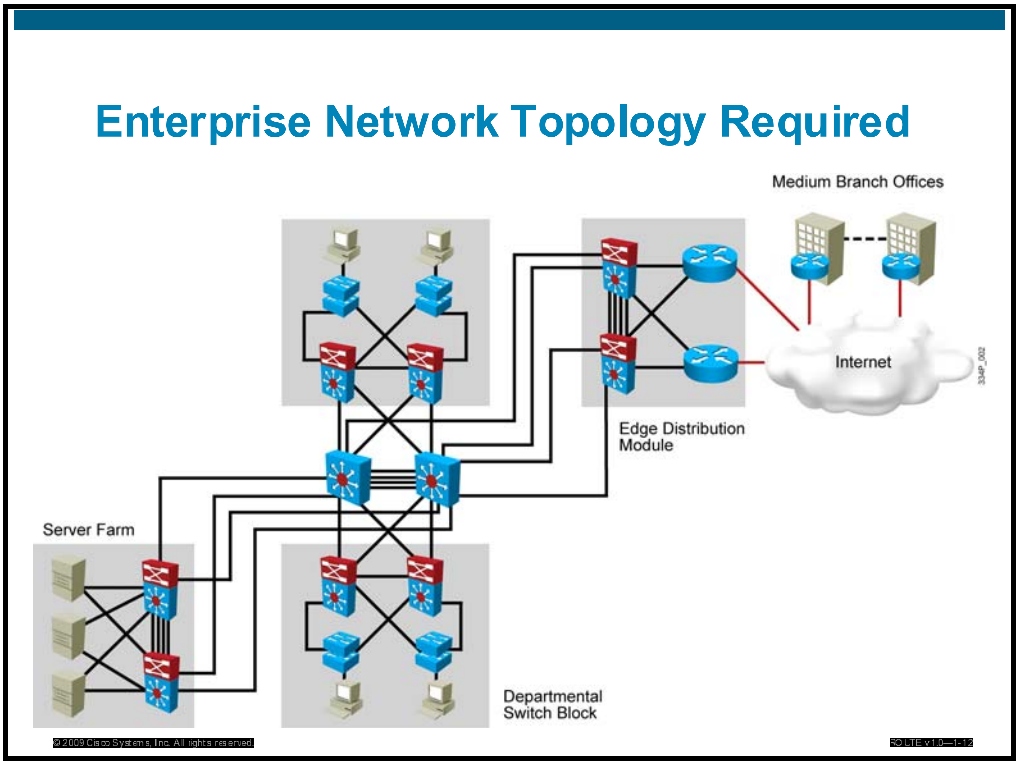

Enterprise Network Topology Required

The figure presents an enterprise network in which a hierarchical design is applied. The company would like to implement a scalable solution with a routing protocol that provides fast convergence. For optimal routing and packet forwarding, hierarchical addressing with summarization is required. Users require high-speed access to the server farm with redundant connectivity for protection. The company has many remote offices and a redundant connection to the Internet is required to provide the remote offices with nonstop access to its server farm. For remote offices, a secure connection must be implemented to prevent unauthorized persons from accessing data.

Network professionals must review the existing topology and other network information needed to implement a new solution. Network professionals must take all requirements into account and create a complete implementation plan. They must document an implementation plan, along with the results of the verification tests.

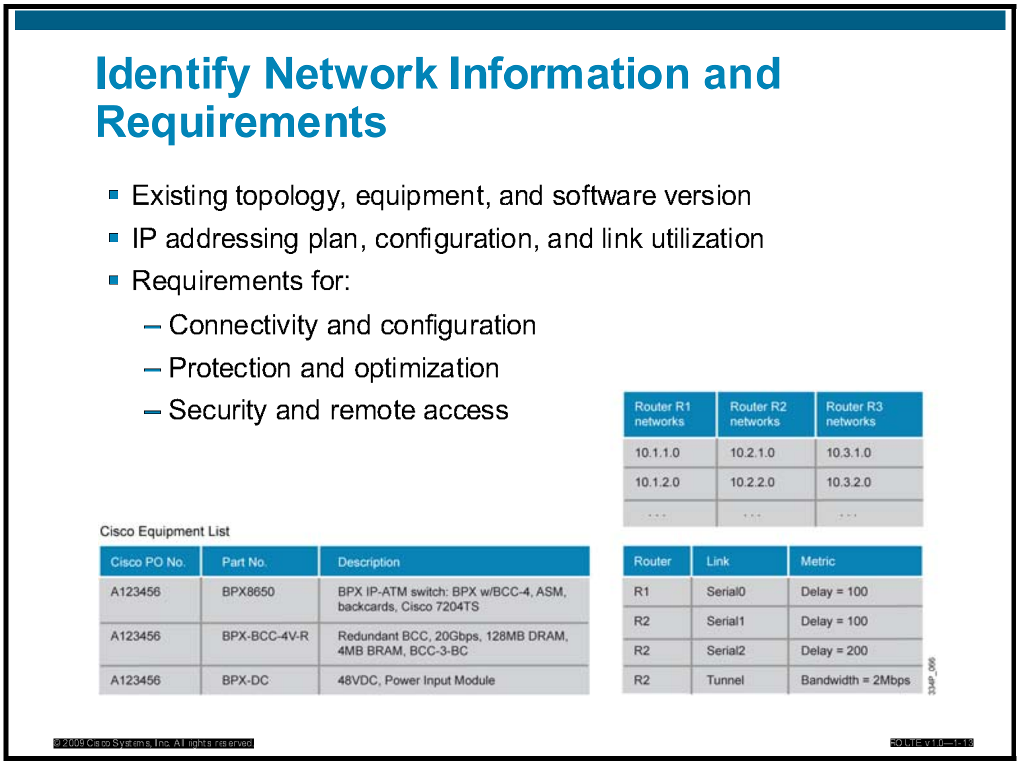

Identify Network Information and Requirements

- Existing topology, equipment, and software version
- IP addressing plan, configuration, and link utilization
- Requirements for:
  - Connectivity and configuration
  - Protection and optimization
  - Security and remote access

Router R1 networks	Router R2 networks	Router R3 networks
10.1.1.0	10.2.1.0	10.3.1.0
10.1.2.0	10.2.2.0	10.3.2.0
. . .	. . .	. . .

Cisco Equipment List

Cisco PO No.	Part No.	Description
A123456	BPX8650	BPX IP-ATM switch: BPX w/BCC-4, ASM, backcards, Cisco 7204TS
A123456	BPX-BCC-4V-R	Redundant BCC, 20Gbps, 128MB DRAM, 4MB BRAM, BCC-3-BC
A123456	BPX-DC	48VDC, Power Input Module

Router	Link	Metric
R1	Serial0	Delay = 100
R2	Serial1	Delay = 100
R2	Serial2	Delay = 200
R2	Tunnel	Bandwidth = 2Mbps

The first step before creating an implementation plan is to gather existing information about the network and all the requirements.

The existing topology provides redundant connectivity among all of the network devices. Internet connectivity is dual-homed, which provides redundant access to the remote sites as well as World Wide Web resources. The equipment can provide all of the functionalities required, but the software version of the operation system must be upgraded.

The networking equipment has existing IP addressing that need to be changed in order to ensure optimal routing and forwarding of packets as well as summarization. Requirements for server farm access and remote office connectivity do not include changes in the Quality of Service (QoS) configuration. The server farm hosts the company's critical applications. Aside from Voice over IP (VoIP), these applications require preferred treatment. Open Shortest Path First (OSPF) is configured in the network. This configuration must be changed, as a faster convergence time is required. The EIGRP routing protocol is a better selection.

Security configuration is required to provide secure access to internal resources and to provide remote office connectivity. Existing security is sufficient and no changes are needed.

After identification of network information, document all details and requirements, including:

- A list of equipment, topology (physical and logical), and design documents
- The current and required software versions
- The current configuration and documentation, such as for IP addressing, summarization, routing information, QoS, and security
- Site requirement specifications, including IP addressing, required software, topology changes, routing protocol requirements, QoS, and security.

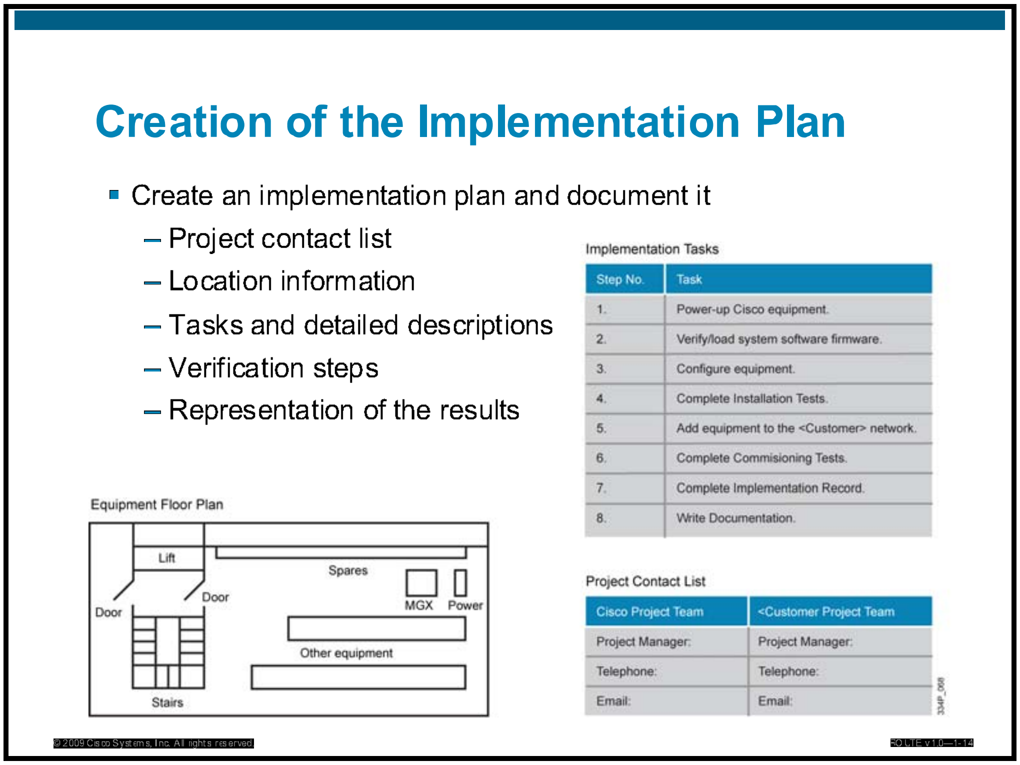

Creation of the Implementation Plan

- Create an implementation plan and document it
  - Project contact list
  - Location information
  - Tasks and detailed descriptions
  - Verification steps
  - Representation of the results

Implementation Tasks

Step No.	Task
1.	Power-up Cisco equipment.
2.	Verify/load system software firmware.
3.	Configure equipment.
4.	Complete Installation Tests.
5.	Add equipment to the <Customer> network.
6.	Complete Commisioning Tests.
7.	Complete Implementation Record.
8.	Write Documentation.

Equipment Floor Plan

Project Contact List

Cisco Project Team	<Customer Project Team
Project Manager:	Project Manager:
Telephone:	Telephone:
Email:	Email:

You must identify the current status of the network and current network requirements before creating the first part of documentation. You must then obtain the following information:

- Project contact list and statements of work
- Location information and means of accessing to the premises
- Tools and resources
- Assumptions
- Tasks and detailed descriptions
- Network Staging Plan

The following examples show the typical content of an implementation plan and a description of each section.

The project contact list introduces all of the people involved and their commitments.

Project Contact List

Cisco Project Team	<Customer> Project Team
Project Manager:	Project Manager:
Telephone:	Telephone:
Email:	Email:
Project Engineer:	Project Engineer:
Telephone:	Telephone:
Email:	Email:
Design Engineer:	Design Engineer:
Telephone:	Telephone:
Email:	Email:
Account Manager:	Account Manager:
Telephone:	Telephone:
Email:	Email:
Systems Engineer:	Systems Engineer:
Telephone:	Telephone:
Email:	Email:

Location information and access details of the premises define where the equipment is located and how to reach it.

Equipment Floor Plan

Location	Details
Floor	
Room	
Suite	
Position	
Rack No.	

A tools description provides a list of tools that the implementation engineer will require to carry out the work detailed in this document.

Tools Required

Item No.	Item
1.	PC with a VT100 emulator, 10BaseT interface, FTP Server and TFTP Client applications
2.	Console port cable DB9-RJ45/DB25
3.	10BaseT Ethernet cable

The implementation task list must provide a breakdown of the implementation process, followed by a detailed description of each activity. The output of each activity should be indicated on the implementation record.

Implementation Tasks

Step No.	Task
1.	Connect to the router
2.	Verify the current installation and create a backup file
3.	Change the Cisco IOS version on all devices
4.	Update the IP address configuration on distribution routers
5.	Configure EIGRP
6.	Verify the configuration and record the results

After the implementation plan is completed successfully, documentation must be created with all of the details, verification steps, and results.

Summary

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

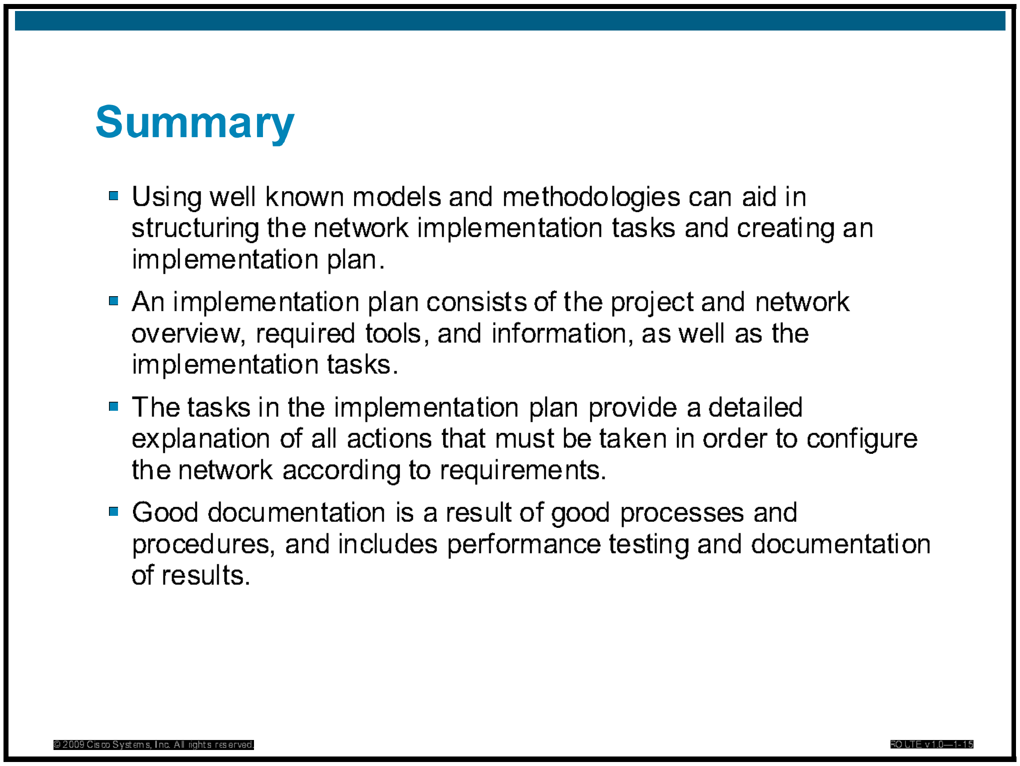

Summary

- Using well known models and methodologies can aid in structuring the network implementation tasks and creating an implementation plan.
- An implementation plan consists of the project and network overview, required tools, and information, as well as the implementation tasks.
- The tasks in the implementation plan provide a detailed explanation of all actions that must be taken in order to configure the network according to requirements.
- Good documentation is a result of good processes and procedures, and includes performance testing and documentation of results.

ROUTE v1.0—1-15

Lesson 3

Lab 1-1 Debrief

Overview

In Lab 1-1, students create the implementation plan for routing protocol selection and implementation. The first part of the lab is focused on gathering requirements and required data. After a student successfully surveys the existing topology and gathers all of the data, the student must create an implementation plan, and then perform implementation and verification. The student must then document the project.

After students complete the lab, the instructor will lead a discussion about lab topology, tasks, verification, and checkpoints. The instructor will also provide a sample solution and different alternatives. Students will present their implementation plans and solutions.

Objectives

Upon completing this lesson, you will be able to explain the gathering of network requirements and required data. You will be able to create an implementation plan, verify it, and document the whole process. This ability includes being able to meet these objectives:

- Complete the Lab Overview and Verification
- Describe a Sample Solution and Alternatives

Lab Overview and Verification

This topic describes lab topology and key checkpoints used to create a solution and to start with the verification.

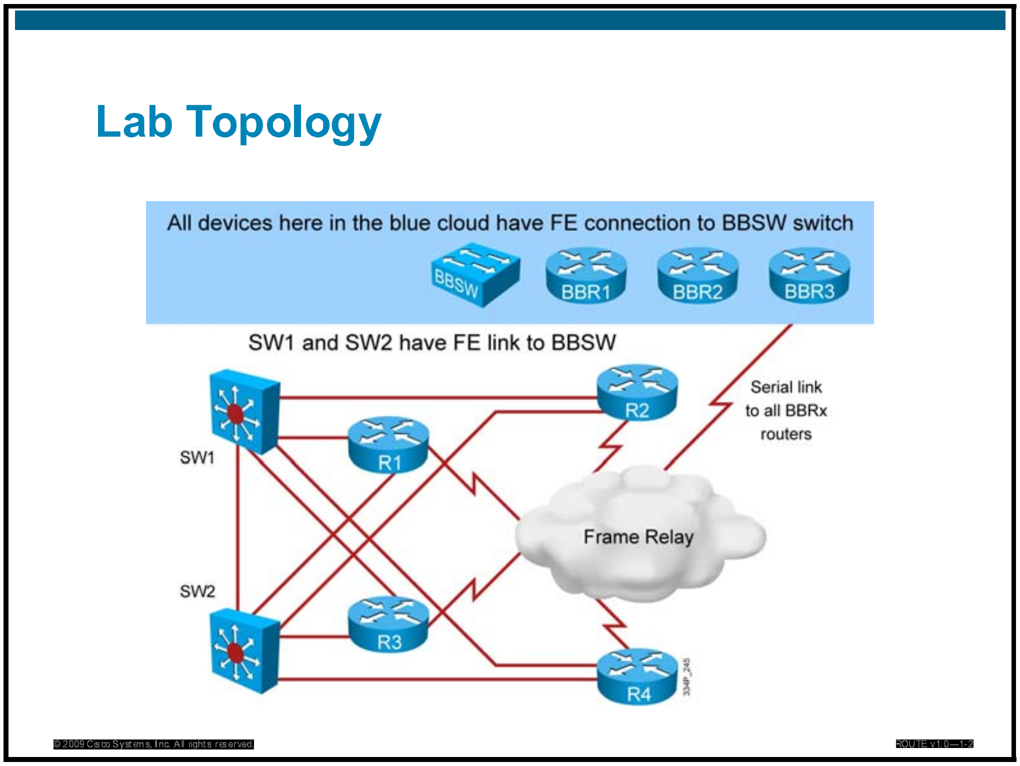

Lab Topology

All devices here in the blue cloud have FE connection to BBSW switch

SW1 and SW2 have FE link to BBSW

Serial link to all BBRx routers

ROUTE v1.0—1-2

The figure above presents the physical lab topology used for all labs in the *Planning Routing Services* course. The topology uses four pod routers, two switches, and backbone equipment. A physical lab is not needed for this lab, as the implementation plan is the theoretical part of each implementation. Implementation and verification make up the practical part, but will be practiced throughout the whole course.

Based on the topology, students will create requirements, gather all of the data, and create implementation plans. Finally, they will describe and document the verification process.

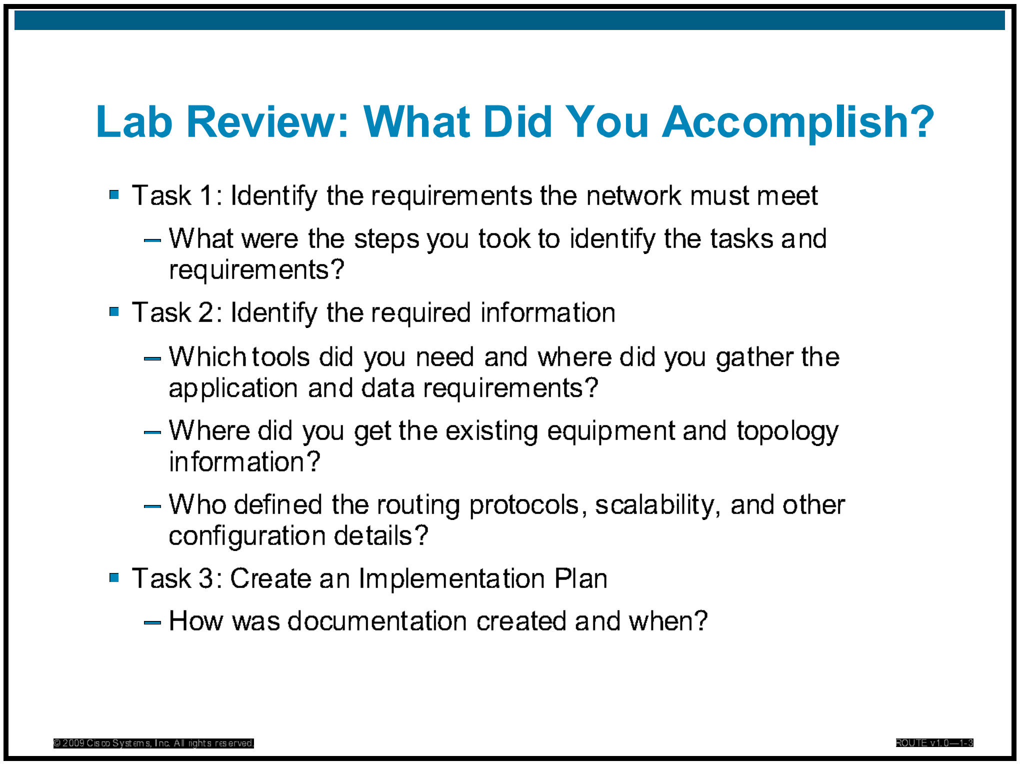

Lab Review: What Did You Accomplish?

- Task 1: Identify the requirements the network must meet
  - What were the steps you took to identify the tasks and requirements?
- Task 2: Identify the required information
  - Which tools did you need and where did you gather the application and data requirements?
  - Where did you get the existing equipment and topology information?
  - Who defined the routing protocols, scalability, and other configuration details?
- Task 3: Create an Implementation Plan
  - How was documentation created and when?

In the first task, you must identify the requirements the network has to meet to establish a foundation for the implementation plan. There are two common approaches to this. You can either define the requirements based on company needs, or gather the requirements from the network administrator. In the first approach, you must select the correct tools to be able to analyze the network and define the requirements in order to start gathering data and creating a good implementation plan. If you choose the second method, you can speed up the process and get the real requirements from the person who knows all the details of the network.

In the second task, the requirements are defined, but the real data is missing. Again, you can do some research and collect all the necessary data using the different tools. As an option, two network administrators can provide all of the data. Then you can select the routing protocols, scalability options, and define other configuration details.

With all of the details, you can create a good implementation plan, then implement, verify, and document the process.

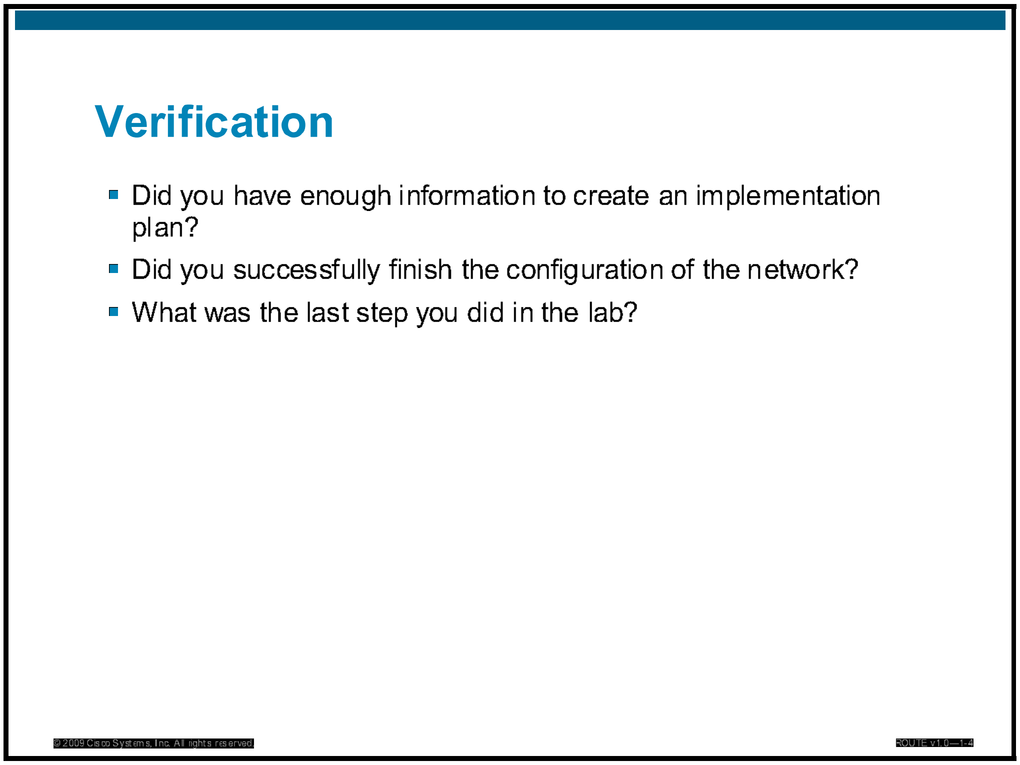

Verification

- Did you have enough information to create an implementation plan?
- Did you successfully finish the configuration of the network?
- What was the last step you did in the lab?

ROUTE v1.0—1-4

A common approach to verifying the implementation process for a routing protocol is to follow the following steps:

- Evaluate if enough information was gathered in order to create a good implementation plan.
- Verify that the routing protocol configuration is successful.
- Create the documentation, which includes all the requirements, required data, implementation and verification steps, as well as implmentation results.

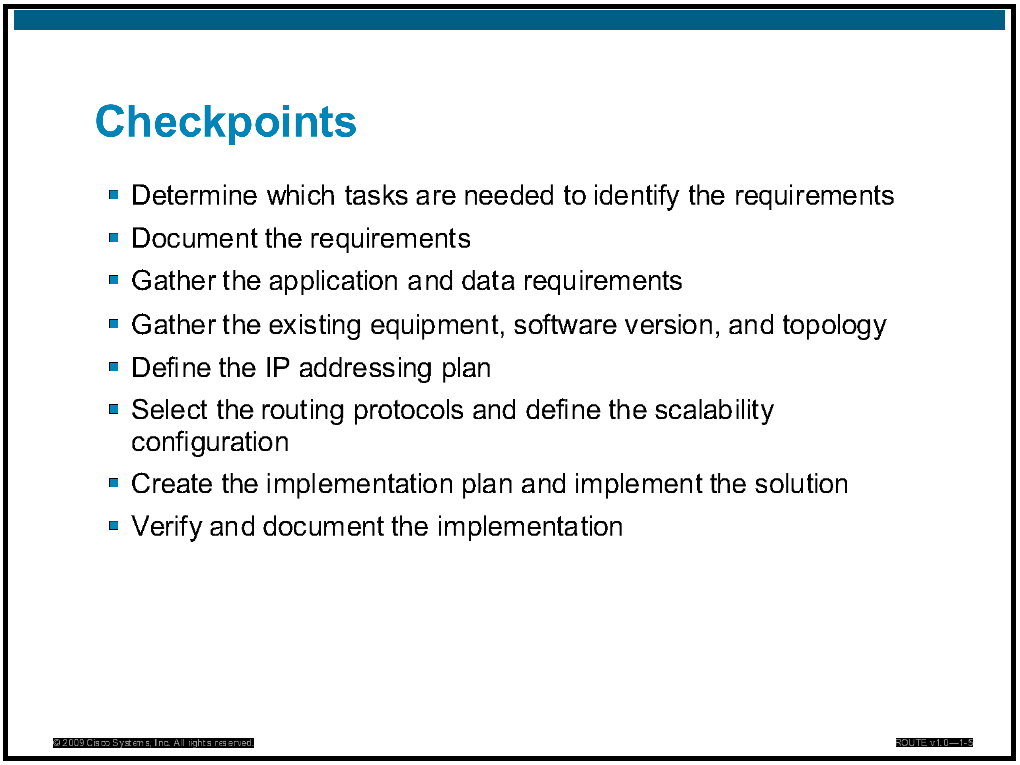

Checkpoints

- Determine which tasks are needed to identify the requirements
- Document the requirements
- Gather the application and data requirements
- Gather the existing equipment, software version, and topology
- Define the IP addressing plan
- Select the routing protocols and define the scalability configuration
- Create the implementation plan and implement the solution
- Verify and document the implementation

During the configuration and verification phase, a network operator can deal with several checkpoints. After completing all configuration tasks, the network operator can complete implementation of a routing protocol or perform additional verification and troubleshooting, as needed.

Optionally, the network operator can check the creation of the implementation plan in different stages using checkpoints verifying each stage.

With different checkpoints, the network operator can verify for proper configuration. The following checkpoints are used for verification:

- Determine which tasks are needed to identify the requirements
- Document the requirements
- Gather the application and data requirements
- Gather the existing equipment, software version, and topology
- Define the IP addressing plan
- Select routing protocols and define scalability configuration
- Create an implementation plan and implement the solution
- Verify and document the implementation

Sample Solutions and Alternatives

This topic describes sample solution and other alternatives.

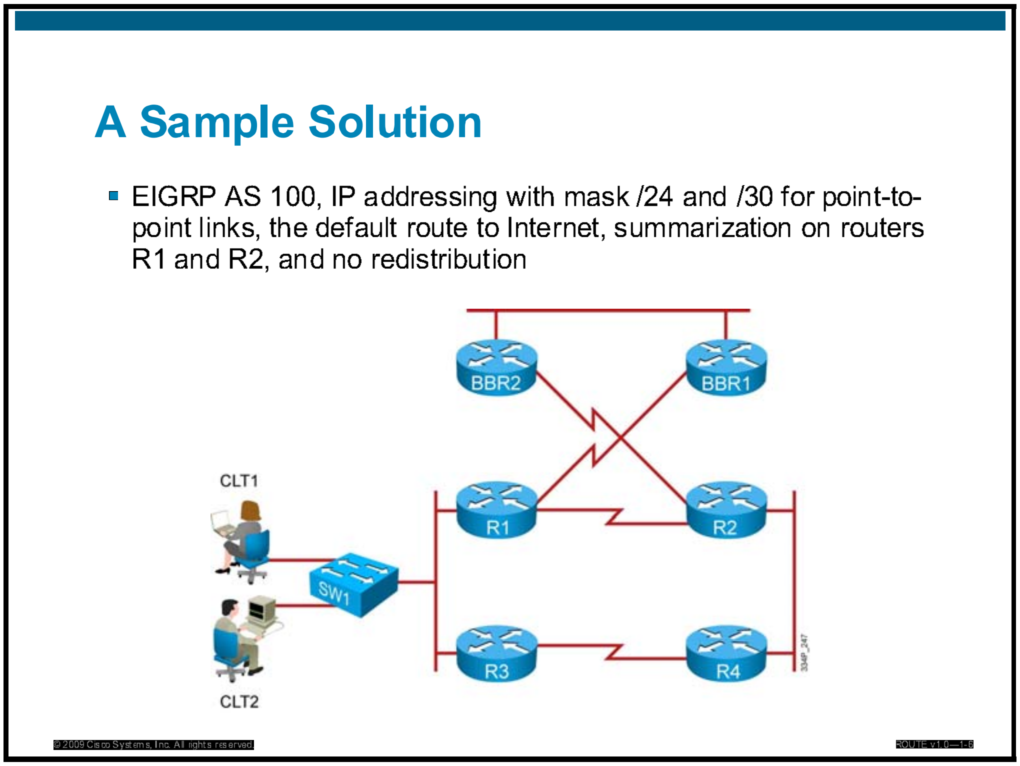

A Sample Solution

- EIGRP AS 100, IP addressing with mask /24 and /30 for point-to-point links, the default route to Internet, summarization on routers R1 and R2, and no redistribution

ROUTE v1.0—1-6

A sample solution includes the implementation details and the details for each task of the implementation plan. Different solutions are possible and the figure points out a few details of a successful configuration.

A proper implementation of the routing protocol might include the following attributes:

- Implementation of EIGRP AS 100
- IP addressing with mask /24 and /30 for point-to-point links
- An Internet gateway that uses the default route announced by routers BBR1 and BBR2
- Implementation of summarization on routers R1 and R2; because only one routing protocol is used, there is no need for redistribution between different routing protocols.

Note As the purpose of lab 1-1 is for you to create an implementation plan, there is no single solution to the lab. The solution presented here is just a sample that satisfies the lab requirements.

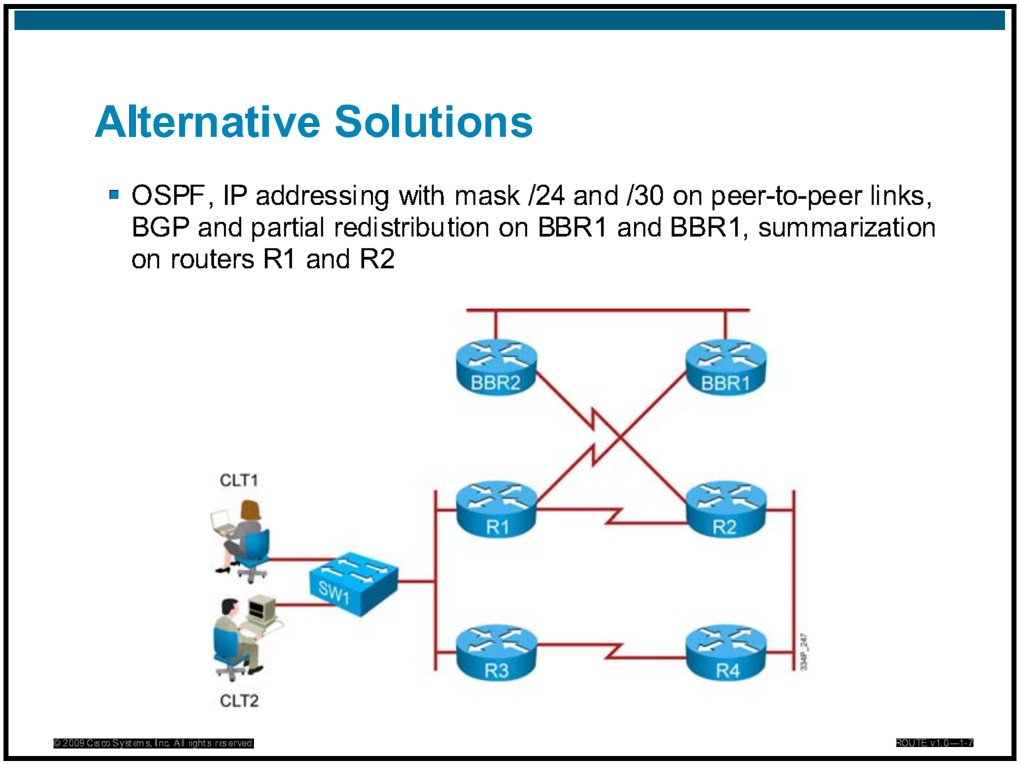

Alternative Solutions

- OSPF, IP addressing with mask /24 and /30 on peer-to-peer links, BGP and partial redistribution on BBR1 and BBR1, summarization on routers R1 and R2

ROUTE v1.0—1-7

You can achieve the same or similar results by using different configuration steps and a different routing protocol.

Instead of EIGRP, you can use the OSPF routing protocol. If you use multihoming and have your own BGP autonomous system number and public IP address space, you can run BGP on routers BBR1 and BBR2 instead of using the default route to the Internet. If you use more than one routing process, you may need to use redistribution or the default route.

Note As the purpose of lab 1-1 is for you to create an implementation plan, there is no single solution to the lab. The alternative solution presented here is just a sample that satisfies the lab requirements.

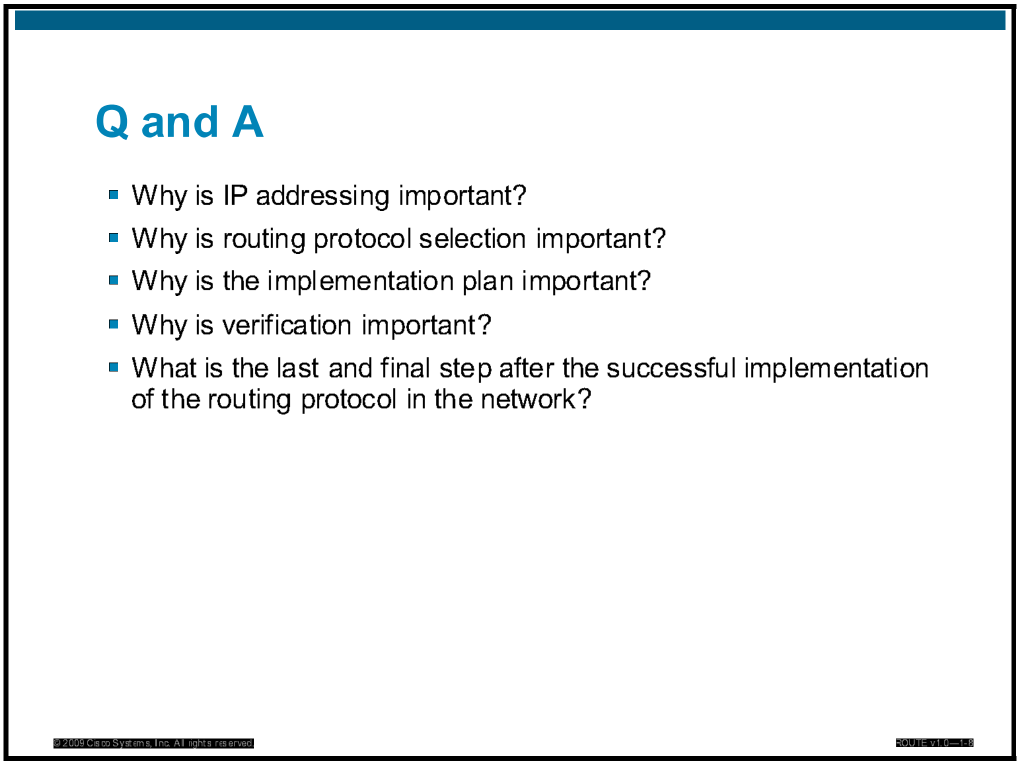

Q and A

- Why is IP addressing important?
- Why is routing protocol selection important?
- Why is the implementation plan important?
- Why is verification important?
- What is the last and final step after the successful implementation of the routing protocol in the network?

IP addressing is important because a good addressing plan makes redistribution and summarization possible. It also helps when a scalable solution is required.

Routing protocol selection is important, because different organizations require different convergence speeds, levels of scalability, and levels of interoperability. EIGRP has a faster convergence speed than OSPF but OSPF might be more scalable in some cases.

An implementation plan is needed in order to correctly implement the proper configuration. Sometimes the steps must be implemented in a specific order. Sometimes the number of steps is so high that without an implementation plan, it is likely that some details might be omitted accidentally.

Verification follows implementation. It is important because it proves our concept and the configuration steps used.

The final step before handover is the creation of documentation. Good documentation is required in order to implement and verify the network. It also helps later when upgrading and troubleshooting takes place.

Summary

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

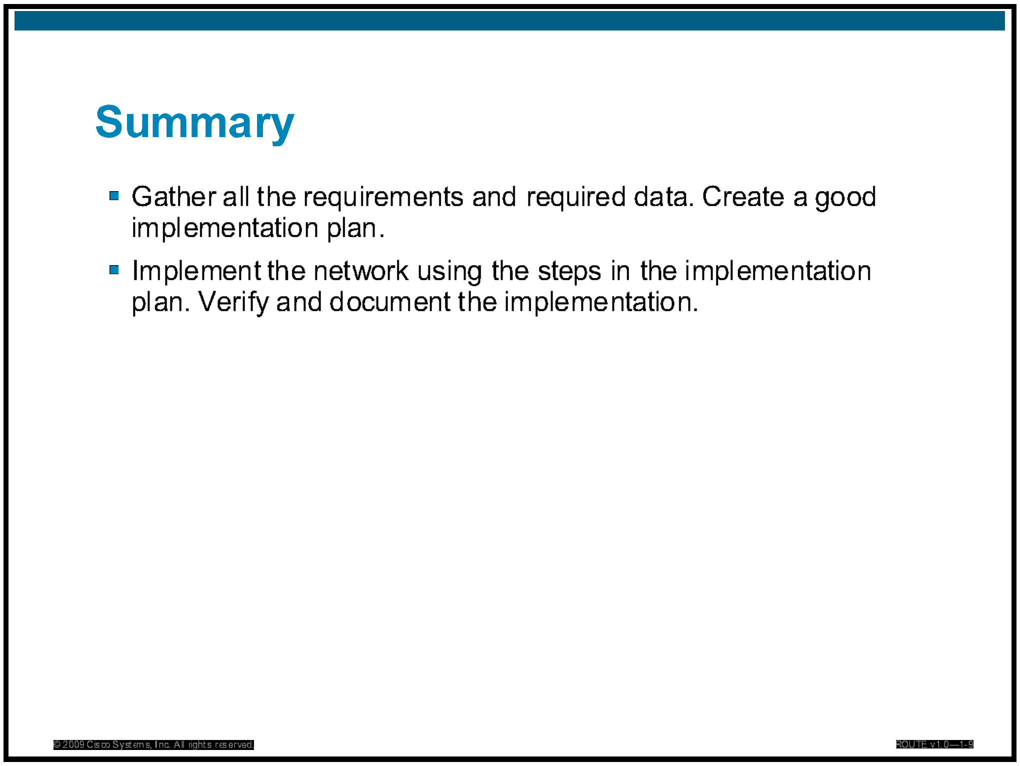

Summary

- Gather all the requirements and required data. Create a good implementation plan.
- Implement the network using the steps in the implementation plan. Verify and document the implementation.

Module Summary

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

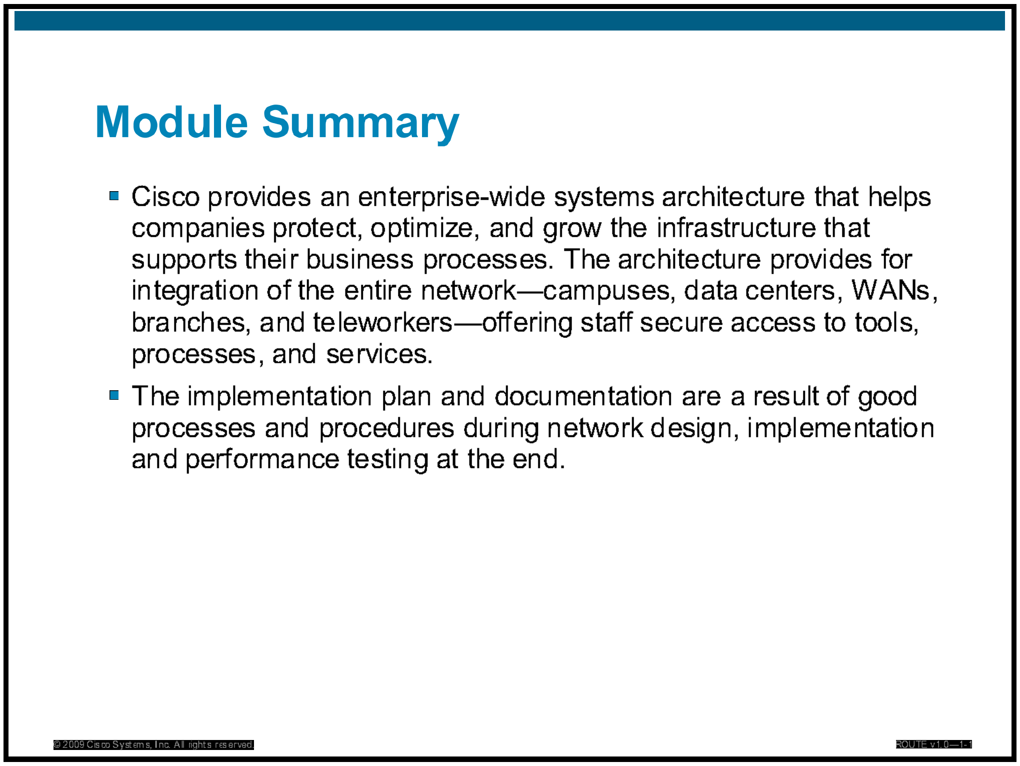

Module Summary

- Cisco provides an enterprise-wide systems architecture that helps companies protect, optimize, and grow the infrastructure that supports their business processes. The architecture provides for integration of the entire network—campuses, data centers, WANs, branches, and teleworkers—offering staff secure access to tools, processes, and services.
- The implementation plan and documentation are a result of good processes and procedures during network design, implementation and performance testing at the end.

This module describes the Cisco conceptual models and architectures for converged networks. It examines the three tiers of the hierarchical network model in detail, the traffic conditions in a converged network, and the use of routing protocols.

Additionally, it describes the creation of an implementation plan, for which the requirements and the required data provide a baseline to define all of the tasks required to produce a successful implementation. Verification after the implementation proves the concept, and documentation is created to finish the implementation process.

.

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.

Q1) Which three layers are parts of the Cisco hierarchical network model? (Choose three.) (Source: Assessing Complex Enterprise Network Requirements)

A) Core
B) Distribution
C) Redistribution
D) Access
E) Workgroup

Q2) What is SAFE? (Source: Assessing Complex Enterprise Network Requirements)

A) security protocol
B) blueprint for network designers and administrators of best practices for the proper deployment of security solutions
C) routing protocol authentication
D) Cisco hierarchical network model white paper

Q3) What are two key network requirements? (Choose two.) (Source: Assessing Complex Enterprise Network Requirements)

A) performance
B) security
C) connectivity
D) convergence speed

Q4) Which advantage is not a SONA advantage for enterprises? (Source: Assessing Complex Enterprise Network Requirements)

A) outlines how enterprises can evolve toward the IIN
B) illustrates how to build integrated systems across a fully converged IIN
C) improves flexibility and increases efficiency, which results in optimized applications, processes, and resources
D) uses the limited product line services

Q5) What are three SONA framework layers? (Source: Assessing Complex Enterprise Network Requirements)

____

____

____

Q6) Which routing protocol supports a very high convergence speed and the use of VLSM? (Source: Assessing Complex Enterprise Network Requirements)

A) EIGRP
B) BGP
C) OSPF
D) RIP

Q7) What are three main steps for a structured approach to implement routing in a network? (Choose three.) (Source: Creating an Implementation Plan and Documenting the Implementation)

A) select the tools used for implementation
B) create an implementation plan
C) implement the solution
D) document the implementation

Q8) What is the name of the Cisco model and methodology that describes a structured approach to network implementation? (Source: Creating an Implementation Plan and Documenting the Implementation)

A) Cisco Lifecycle Services
B) Cisco ITIL
C) Cisco FCAPS
D) Cisco TMN

Q9) Which three items must be identified prior to the creation of an implementation plan? (Choose three.) (Source: Creating an Implementation Plan and Documenting the Implementation)

A) network-specific information, activities and tasks associated with the implementation plan development
B) dependencies your implementation plan development has on other service components
C) recommended resources to accomplish the activities and task associated with implementation plan development
D) implementation plan and verification tasks

Q10) What information must you know in order to create an implementation plan for EIGRP routing protocol? (Choose three.) (Source: Creating an Implementation Plan and Documenting the Implementation)

A) existing topology, equipment, software version
B) IP addressing plan and scalability requirements
C) tools needed to evaluate application requirements
D) list of advertized networks and metrics

Q11) When changing the software version of the existing network infrastructure, Layer 3 devices are part of the implementation plan for routing protocols. (Source: Creating an Implementation Plan and Documenting the Implementation)

A) true
B) false

Q12) What is not part of the implementation plan documentation for configuring a routing protocol in an enterprise network? (Source: Creating an Implementation Plan and Documenting the Implementation)

A) tools
B) future upgrade tasks
C) implementation plan tasks
D) verification tasks
E) performance measurement and results

Q13) Interpretation of performance results must be done prior to the verification steps within the implementation plan for routing protocols. (Source: Creating an Implementation Plan and Documenting the Implementation)

A) true
B) false

Module Self-Check Answer Key

Q1) A, B, D

Q2) B

Q3) A, B

Q4) D

Q5) Networked infrastructure layer, interactive services layer, application layer

Q6) A

Q7) B, C, D

Q8) A

Q9) A, B, C

Q10) A, B, D

Q11) A

Q12) B

Q13) B

Module 2

Implementing an EIGRP-Based Solution

Overview

In routing environments, Enhanced Interior Gateway Routing Protocol (EIGRP) offers benefits and features over historical distance-vector routing protocols such as Routing Information Protocol version 1 (RIPv1). These benefits include rapid convergence, lower bandwidth utilization, and multiple routed protocol support besides IP.

This module describes how EIGRP works and how to implement and verify EIGRP operations. It also explores advanced topics like route summarization, load balancing, EIGRP bandwidth usage, and authentication. The module concludes with a discussion of EIGRP issues and problems as well as how to correct them.

Module Objectives

Upon completing this module, you will be able to implement and verify EIGRP operations. This ability includes being able to meet these objectives:

- Identify the technologies, components, and metrics of EIGRP needed to implement routing in diverse, large-scale internetworks based on requirements.
- Configure EIGRP according to a given implementation plan and set of requirements.
- Discuss the lab results for configuring and verifying EIGRP operations.
- Configure and verify EIGRP over circuit emulation, MPLS VPNs, and Frame Relay for operational performance.
- Discuss the lab results for configuring and verifying EIGRP circuit emulation, MPLS VPNs, and Frame Relay operations.
- Configure and verify EIGRP authentication for operational performance.
- Discuss the lab results for configuring and verifying EIGRP authentication.
- Implement and verify the advanced EIGRP features in an Enterprise Network.
- Discuss the lab results for implementing and verifying EIGRP operations.

Lesson 1

Planning Routing Implementations with EIGRP

Overview

To select the appropriate routing protocols for an internetwork, you must understand the key features and terminology that are necessary to evaluate a given protocol against other choices. Routing protocols are distinguished from one another by the way that each selects the best pathway and the way that each calculates the routing protocol metric. Knowing the correct commands to use when you configure Enhanced Interior Gateway Routing Protocol (EIGRP) helps to ensure that the migration to this routing protocol is smooth and quick.

This lesson reviews the benefits of EIGRP and discusses the key capabilities that distinguish EIGRP from other routing protocols, including the four underlying technologies within EIGRP. The three tables that EIGRP uses in the path selection process are described, and EIGRP metric calculation is explored in detail. An implementation plan is described as the first step in configuring EIGRP, followed by basic EIGRP configuration.

Objectives

Upon completing this lesson, you will be able to describe the components and metrics of EIGRP, how EIGRP selects routes between routers in diverse, large-scale internetworks, the implementation plan creation process, and basic EIGRP configuration. This ability includes being able to meet these objectives:

- List EIGRP capabilities and attributes.
- Define EIGRP operation and metrics.
- Plan and Document EIGRP.
- Implement Basic EIGRP.

EIGRP Capabilities and Attributes

Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet masking (VLSM), partial updates, and support for multiple network layer protocols. This topic describes these capabilities.

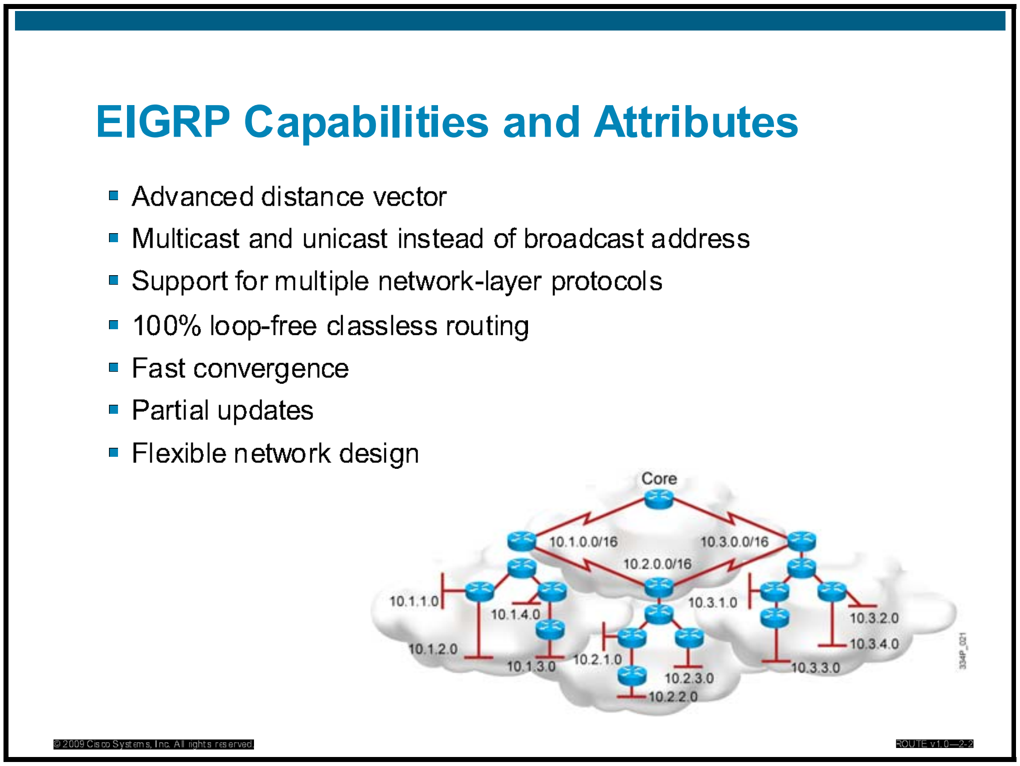

EIGRP Capabilities and Attributes

- Advanced distance vector
- Multicast and unicast instead of broadcast address
- Support for multiple network-layer protocols
- 100% loop-free classless routing
- Fast convergence
- Partial updates
- Flexible network design

ROUTE v1.0—2-2

EIGRP is a Cisco proprietary protocol that combines the advantages of link-state and distance vector routing protocols. EIGRP has its roots as a distance vector routing protocol and is predictable in its behavior. EIGRP is easy to configure and is adaptable to a wide variety of network topologies. The addition of several link-state features, such as dynamic neighbor discovery, makes EIGRP an advanced distance vector protocol. EIGRP is an *enhanced* IGRP because of its rapid convergence and the guarantee of a loop-free topology at all times. A hybrid protocol, EIGRP uses the Diffusing Update Algorithm (DUAL) and includes the following key features:

- **Fast convergence:** A router running EIGRP stores all of its neighbors' routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found.

- **Partial updates:** EIGRP does not send periodic updates. Instead, it sends partial triggered updates; these are sent only when the path or the metric changes for a route and contain information about only the changed routes. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. Because of these two capabilities, EIGRP consumes significantly less bandwidth. This behavior is different from link-state protocols, in which an update is transmitted to all link-state routers within an area.

- **Multiple network-layer protocol support:** EIGRP supports multiple network-layer protocols (for example IP) by using protocol-dependent modules. These modules are responsible for protocol requirements specific to the network layer. The rapid convergence and sophisticated metric structure of EIGRP offers superior performance and stability.

Note This course covers only the TCP/IP implementation of EIGRP.

- **Multicast and unicast:** EIGRP uses multicast and unicast, rather than broadcast. The multicast address used for EIGRP is 224.0.0.10.

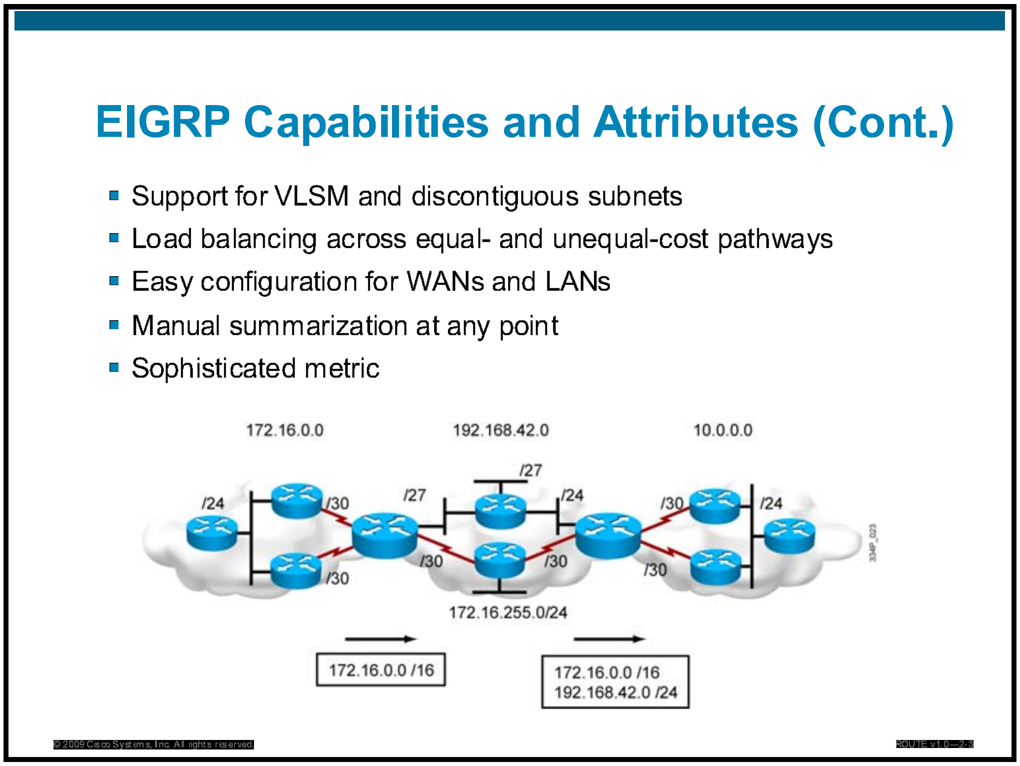

EIGRP Capabilities and Attributes (Cont.)

- Support for VLSM and discontiguous subnets
- Load balancing across equal- and unequal-cost pathways
- Easy configuration for WANs and LANs
- Manual summarization at any point
- Sophisticated metric

172.16.0.0
192.168.42.0
10.0.0.0
/27
/24
/30
/27
/24
/30
/24
/30
/30
/30
/30
/30
172.16.255.0/24
172.16.0.0 /16
172.16.0.0 /16
192.168.42.0 /24
334P_023

ROUTE v1.0—2-3

EIGRP features also include:

- **VLSM support**: EIGRP is a classless routing protocol, which means that it advertises a subnet mask for each destination network. This feature enables EIGRP to support discontinuous subnetworks and VLSM. With EIGRP, routes are automatically summarized at the major network number boundary, but EIGRP can be configured to summarize on any bit boundary on any router interface.

- **Seamless connectivity across all data-link layer protocols and topologies:** EIGRP does not require special configuration to work across any Layer 2 protocols. Other routing protocols, such as Open Shortest Path First (OSPF), use different configurations for different Layer 2 protocols, such as Ethernet and Frame Relay. EIGRP operates effectively in both LAN and WAN environments. WAN support for dedicated point-to-point links and nonbroadcast multiaccess (NBMA) topologies is standard for EIGRP. EIGRP accommodates differences in media types and speeds when neighbor adjacencies form across WAN links and can be configured to limit the amount of bandwidth that the protocol uses on WAN links.

- **Sophisticated metric:** EIGRP represents metric values in a 32-bit format to provide enough granularity. EIGRP supports unequal metric load balancing, which allows administrators to distribute traffic flow more efficiently in their networks.

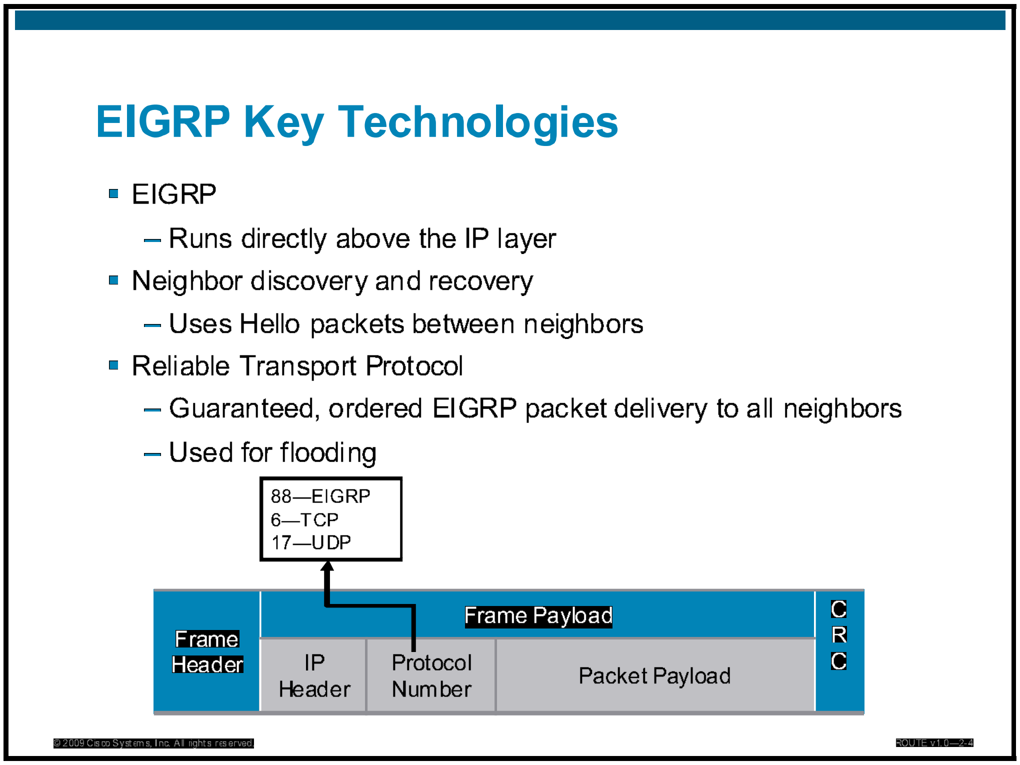

EIGRP Key Technologies

- EIGRP
  - Runs directly above the IP layer
- Neighbor discovery and recovery
  - Uses Hello packets between neighbors
- Reliable Transport Protocol
  - Guaranteed, ordered EIGRP packet delivery to all neighbors
  - Used for flooding

88—EIGRP
6—TCP
17—UDP

Frame Header

Frame Payload

IP Header

Protocol Number

Packet Payload

CRC

ROUTE v1.0—2-4

EIGRP runs directly above the IP layer (protocol number 88) and employs four key technologies that combine to differentiate it from other routing technologies: neighbor discovery and recovery, reliable transport protocol, Diffusing Update Algorithm (DUAL) finite-state machines, and protocol-dependent modules.

- **Neighbor discovery and recovery mechanism:** This mechanism enables routers to learn dynamically about other routers on their directly-attached networks. Routers must also discover when their neighbors become unreachable or inoperative. This process is achieved with low overhead by periodically sending small hello packets. As long as a router receives hello packets from a neighboring router, it assumes that the neighbor is functioning and the two can exchange routing information.

- **Reliable Transport Protocol:** This protocol is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast or unicast packets. For efficiency, only certain EIGRP packets are transmitted reliably.

  For example, on a multiaccess network that has multicast capabilities, such as Ethernet, it is not necessary to send hello packets reliably to all neighbors individually, so EIGRP sends a single multicast hello packet containing an indicator that informs the receivers that the packet need not be acknowledged. Other types of packets, such as updates, contain an indicator in the packet that acknowledgment is required. Reliable transport protocol contains a provision for sending multicast packets quickly, even when unacknowledged packets are pending. This provision helps ensure that convergence time remains low in the presence of varying link speeds.

- **DUAL** enables EIGRP routers to find out whether a path to the destination network is loop-free. DUAL allows a router running EIGRP to find alternate paths based on updates received from other routers..

- **Protocol-dependent modules** are responsible for network layer protocol-specific requirements. The IP-EIGRP module is responsible for sending and receiving EIGRP packets, which are encapsulated in IP as well as for parsing EIGRP packets and informing DUAL of the new information that has been received. IP-EIGRP asks DUAL to make routing decisions, to put results in the IP routing table and to redistribute routes learned by other IP routing protocols.

EIGRP Operation and Metric

EIGRP uses the neighbor table to list adjacent routers. The topology table lists all the learned routes to each destination, while the routing table contains the best route to each destination. This best route is called the successor route. A feasible successor route is a backup route to a destination, which is kept in the topology table. This topic describes how EIGRP uses these tables and routes in its operation.

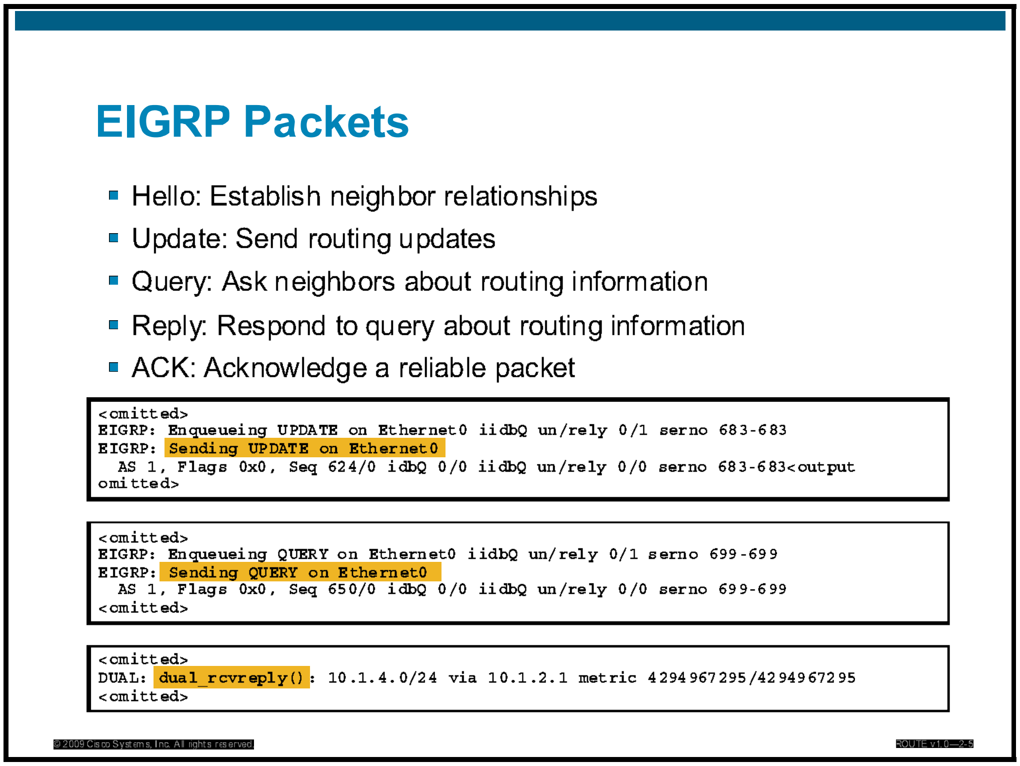

EIGRP Packets

- Hello: Establish neighbor relationships
- Update: Send routing updates
- Query: Ask neighbors about routing information
- Reply: Respond to query about routing information
- ACK: Acknowledge a reliable packet

```
<omitted>
EIGRP: Enqueueing UPDATE on Ethernet0 iidbQ un/rely 0/1 serno 683-683
EIGRP: Sending UPDATE on Ethernet0
  AS 1, Flags 0x0, Seq 624/0 idbQ 0/0 iidbQ un/rely 0/0 serno 683-683<output
omitted>
```

```
<omitted>
EIGRP: Enqueueing QUERY on Ethernet0 iidbQ un/rely 0/1 serno 699-699
EIGRP: Sending QUERY on Ethernet0
  AS 1, Flags 0x0, Seq 650/0 idbQ 0/0 iidbQ un/rely 0/0 serno 699-699
<omitted>
```

```
<omitted>
DUAL: dual_rcvreply(): 10.1.4.0/24 via 10.1.2.1 metric 4294967295/4294967295
<omitted>
```

ROUTE v1.0—2-5

EIGRP uses the following five generic packet types:

- **Hello:** Routers use hello packets for neighbor discovery. The packets are sent as multicasts and do not require acknowledgments.
- **Update:** Update packets contain route change information. They are sent reliably to the affected routers only. These updates can be unicast to a specific router or multicast to multiple attached routers.
- **Query:** When a router performs a route computation and does not have a feasible successor, it sends a reliable query packet to its neighbors to determine if they have a feasible successor for the destination. Queries are normally multicast but can be retransmitted as unicast packets in certain cases.
- **Reply:** A router sends a reply packet in response to a query packet. Replies are unicast reliably to the originator of the query.
- **ACK:** The acknowledgment (ACK) packet acknowledges update, query, and reply packets. ACK packets are unicast hello packets and contain a nonzero acknowledgment number.

Initial Route Discovery

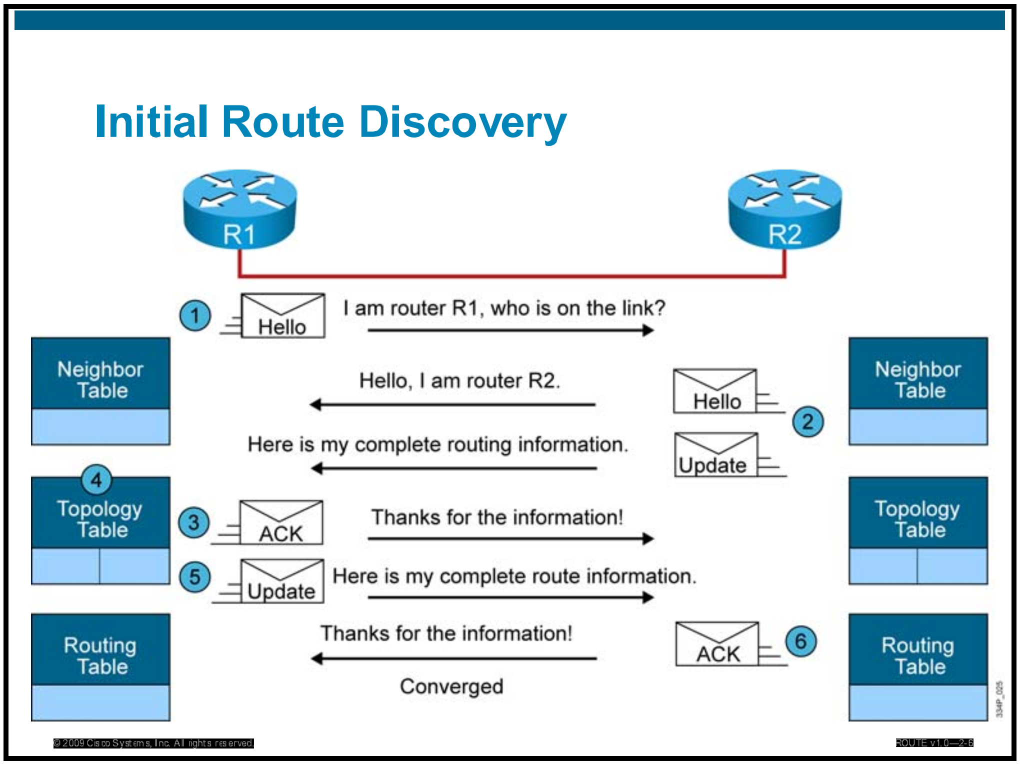
R1
R2
1
Hello
I am router R1, who is on the link?
Neighbor Table
Hello, I am router R2.
Hello
2
Neighbor Table
Here is my complete routing information.
Update
4
Topology Table
3
ACK
Thanks for the information!
Topology Table
5
Update
Here is my complete route information.
Routing Table
Thanks for the information!
ACK
6
Routing Table
Converged
334P_025

ROUTE v1.0—2-6

The process to establish and discover neighbor routes occurs simultaneously with EIGRP. A high-level description of the process follows, using the topology in the figure as an example:

1. A new router (router R1) comes up on the link and sends a hello packet through all of its EIGRP-configured interfaces.

2. Routers receiving the hello packet (router R2) on one interface reply with update packets that contain all the routes they have in their routing tables, except those learned through that interface (split horizon). Router R2 sends an update packet to router R1, but a neighbor relationship is not established until router R2 sends a hello packet to router R1. The update packet from router R2 has the initialization bit set, indicating that this is the initialization process. The update packet includes information about the routes that the neighbor (router R2) is aware of, including the metric that the neighbor is advertising for each destination.

3. After both routers have exchanged hellos and the neighbor adjacency is established, router R1 replies to router R2 with an ACK packet, indicating that it received the update information.

4. Router R1 assimilates all update packets in its topology table. The topology table includes all destinations advertised by neighboring (adjacent) routers. It lists each destination, all the neighbors that can reach the destination, and their associated metric.

5. Router R1 then sends an update packet to router R2.

6. Upon receiving the update packet, router R2 sends an ACK packet to router R1.

After routers R1 and R2 successfully receive the update packets from each other, they are ready to update their routing tables with the successor routes from the topology table.

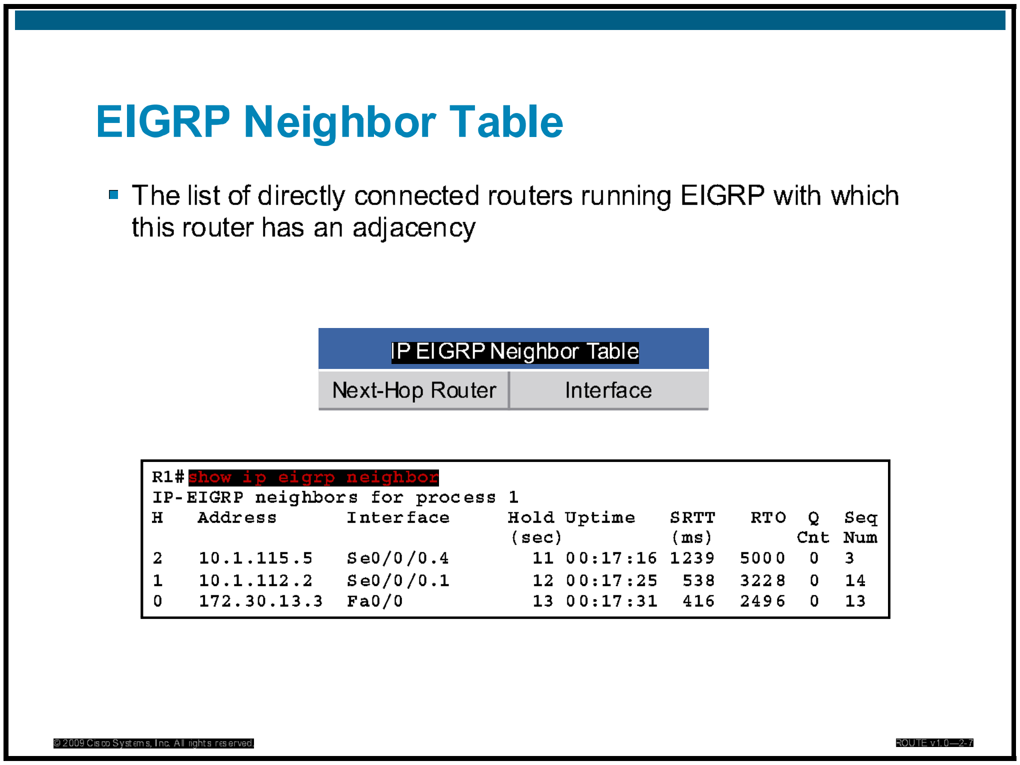

EIGRP Neighbor Table

- The list of directly connected routers running EIGRP with which this router has an adjacency

IP EIGRP Neighbor Table	
Next-Hop Router	Interface

```
R1#show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H   Address        Interface      Hold Uptime   SRTT   RTO  Q   Seq
                                  (sec)         (ms)        Cnt Num
2   10.1.115.5     Se0/0/0.4        11 00:17:16 1239   5000  0   3
1   10.1.112.2     Se0/0/0.1        12 00:17:25  538   3228  0   14
0   172.30.13.3    Fa0/0            13 00:17:31  416   2496  0   13
```

When a router discovers and forms an adjacency with a new neighbor, it records the neighbor's address and the interface through which it can be reached as an entry in the neighbor table. One neighbor table exists for each protocol-dependent module. The EIGRP neighbor table is comparable to the adjacencies database that link-state routing protocols use and serves the same purpose: to ensure bidirectional communication between each of the directly connected neighbors.

When a neighbor sends a hello packet, it advertises a hold time, which is the amount of time that a router treats a neighbor as reachable and operational. If a hello packet is not received within the hold time, the hold time expires and DUAL is aware of the topology change.

The neighbor-table entry also includes information required by the reliable transport protocol. Sequence numbers are employed to match acknowledgments with data packets, and the last sequence number received from the neighbor is recorded, so that out-of-order packets can be detected. A transmission list is used to queue packets for possible retransmission on a per-neighbor basis. Round-trip timers are kept in the neighbor-table entry to estimate an optimal retransmission interval.

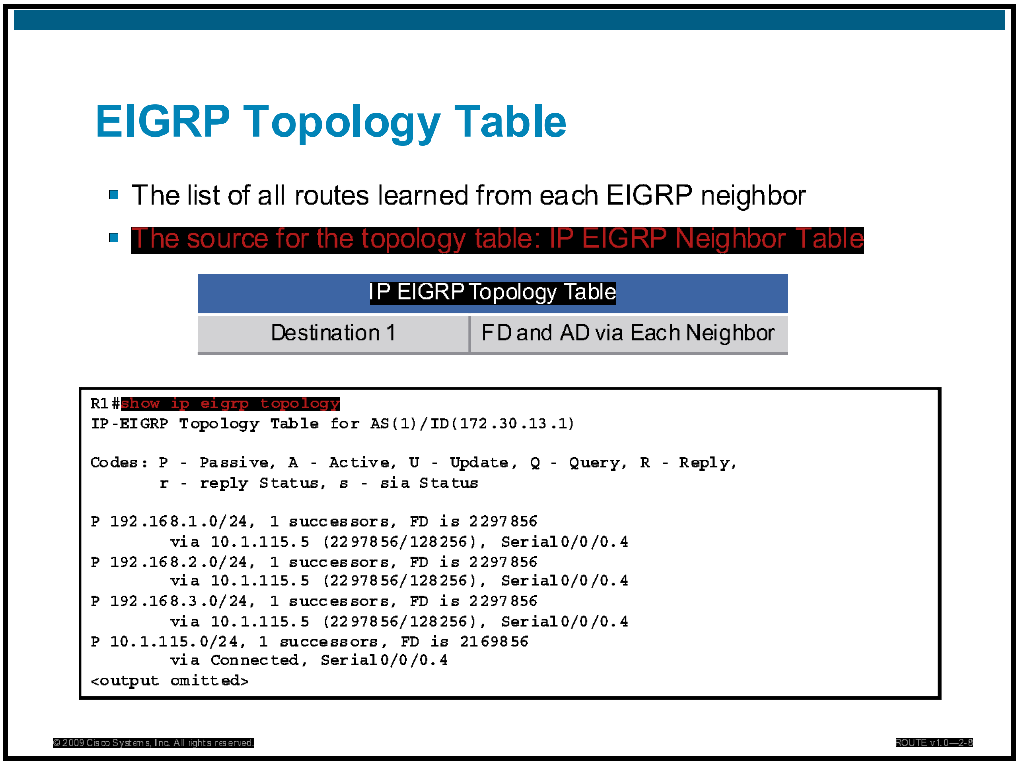

EIGRP Topology Table

- The list of all routes learned from each EIGRP neighbor
- The source for the topology table: IP EIGRP Neighbor Table

IP EIGRP Topology Table	
Destination 1	FD and AD via Each Neighbor

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(1)/ID(172.30.13.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 192.168.1.0/24, 1 successors, FD is 2297856
        via 10.1.115.5 (2297856/128256), Serial0/0/0.4
P 192.168.2.0/24, 1 successors, FD is 2297856
        via 10.1.115.5 (2297856/128256), Serial0/0/0.4
P 192.168.3.0/24, 1 successors, FD is 2297856
        via 10.1.115.5 (2297856/128256), Serial0/0/0.4
P 10.1.115.0/24, 1 successors, FD is 2169856
        via Connected, Serial0/0/0.4
<output omitted>
```

When the router dynamically discovers a new neighbor, it sends an update about the routes it knows to its new neighbor and receives the same from the new neighbor. These updates populate the topology table. The topology table contains all destinations advertised by neighboring routers. It is important to note that if a neighbor is advertising a destination, it must be using that route to forward packets; this rule must be strictly followed by all distance vector protocols.

The topology table also maintains the metric that each neighbor advertises for each destination, the advertised distance (AD), and the metric that this router would use to reach the destination via that neighbor, the feasible distance (FD). The FD is the cost for this router to reach the neighbor for this destination, plus the neighbor's metric to reach the destination.

The topology table is updated when a directly connected route or interface changes or when a neighboring router reports a change to a route.

A topology-table entry for a destination can be in one of two states: active or passive. A destination is in the *passive* state when the router is not performing a recomputation; it is in the *active* state when the router is performing a recomputation. If feasible successors are always available, a destination never has to go into the active state, thereby avoiding a recomputation. The desired state is the passive state.

A recomputation occurs when a destination has no feasible successors. The router initiates the recomputation by sending a query packet to each of its neighboring routers. If the neighboring router has a route for the destination, it sends a reply packet; if it does not have a route, it sends a query packet to its neighbors. In this case, the route is also in the active state in the neighboring router. While a destination is in the active state, a router cannot change the destination's routing table information. After a router has received a reply from each neighboring router, the topology-table entry for the destination returns to the passive state, and the router can select a successor.

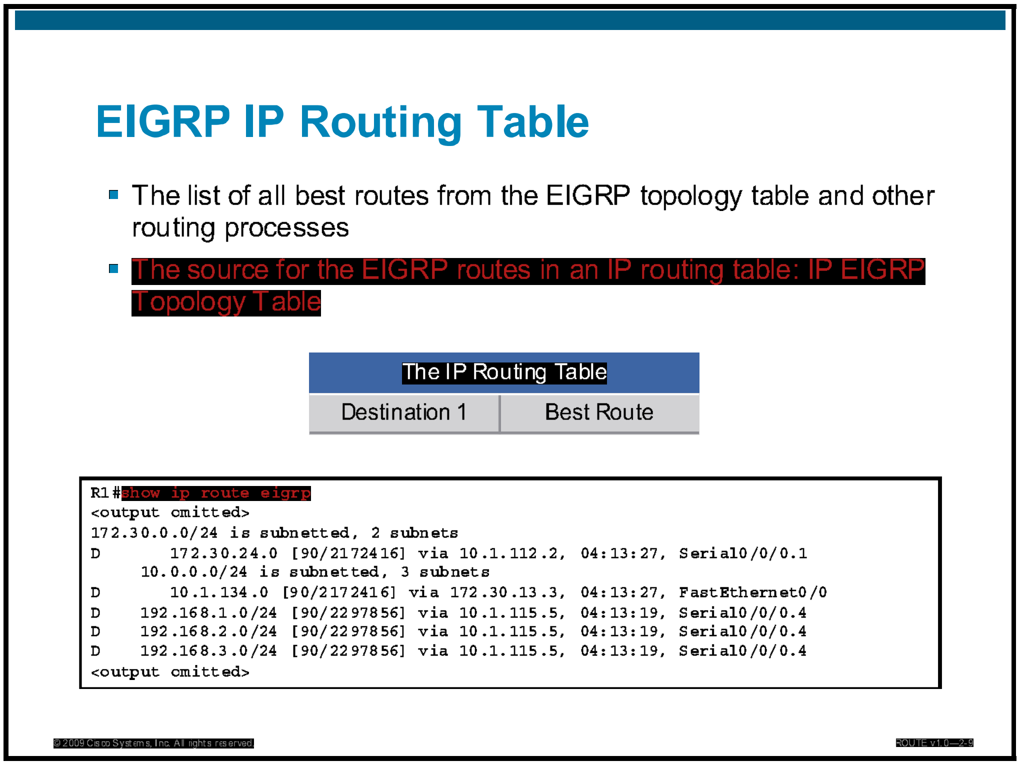

EIGRP IP Routing Table

- The list of all best routes from the EIGRP topology table and other routing processes
- The source for the EIGRP routes in an IP routing table: IP EIGRP Topology Table

The IP Routing Table	
Destination 1	Best Route

```
R1#show ip route eigrp
<output omitted>
172.30.0.0/24 is subnetted, 2 subnets
D       172.30.24.0 [90/2172416] via 10.1.112.2, 04:13:27, Serial0/0/0.1
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.134.0 [90/2172416] via 172.30.13.3, 04:13:27, FastEthernet0/0
D    192.168.1.0/24 [90/2297856] via 10.1.115.5, 04:13:19, Serial0/0/0.4
D    192.168.2.0/24 [90/2297856] via 10.1.115.5, 04:13:19, Serial0/0/0.4
D    192.168.3.0/24 [90/2297856] via 10.1.115.5, 04:13:19, Serial0/0/0.4
<output omitted>
```

ROUTE v1.0—2-9

A router compares all FDs to reach a specific network, then selects the route with the lowest FD and places it in the IP routing table; this is called the successor route. The FD for the chosen route becomes the EIGRP routing metric to reach that network in the routing table.

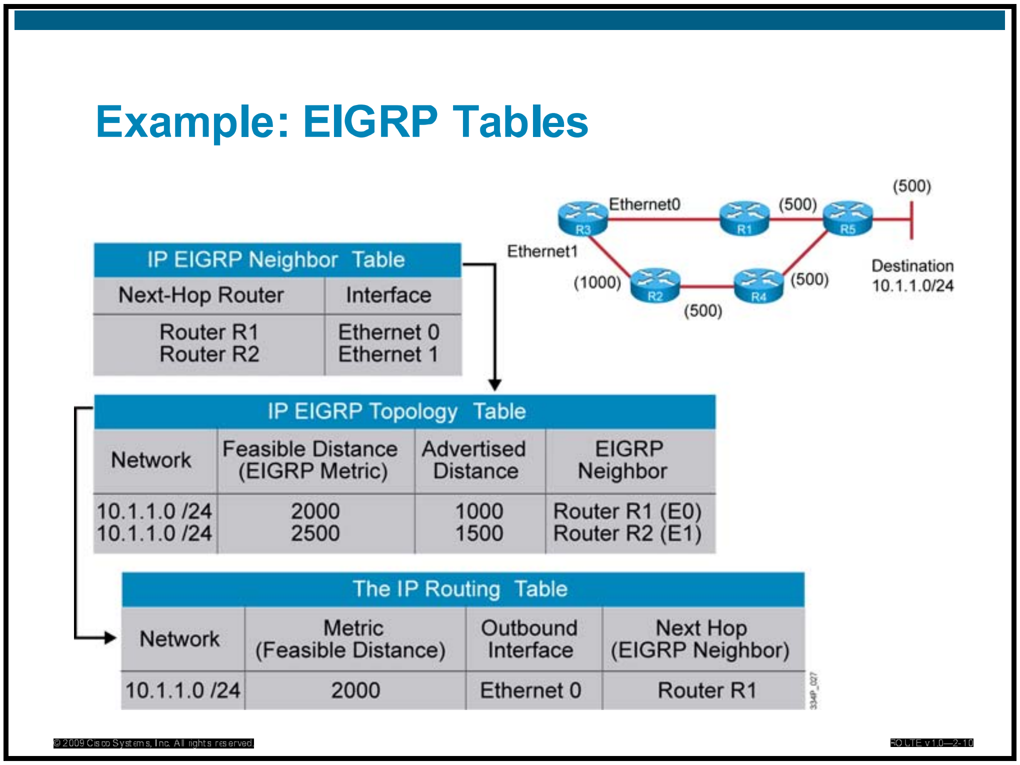

Example: EIGRP Tables

IP EIGRP Neighbor Table

Next-Hop Router	Interface
Router R1 Router R2	Ethernet 0 Ethernet 1

IP EIGRP Topology Table

Network	Feasible Distance (EIGRP Metric)	Advertised Distance	EIGRP Neighbor
10.1.1.0 /24	2000	1000	Router R1 (E0)
10.1.1.0 /24	2500	1500	Router R2 (E1)

The IP Routing Table

Network	Metric (Feasible Distance)	Outbound Interface	Next Hop (EIGRP Neighbor)
10.1.1.0 /24	2000	Ethernet 0	Router R1

ROUTE v1.0—2-10

Example: EIGRP Tables

The network shown illustrates the EIGRP tables; router R3's tables are displayed. Routers R1 and R2 have established a neighbor relationship with router R3 and have sent their routing tables to router R3. Both routers R1 and R2 have paths to network 10.1.1.0/24, among many others that are not shown.

The routing table on router R1 has an EIGRP metric of 1000 for 10.1.1.0/24, so router R1 advertises 10.1.1.0/24 to router R3 with a metric of 1000. Router R3 installs the route to 10.1.1.0/24 via router R1 in its EIGRP topology table with an advertised distance of 1000.

Router R2 has network 10.1.1.0/24 with a metric of 1500 in its IP routing table, so router R2 advertises 10.1.1.0/24 to router R3 with an advertised distance of 1500. Router R3 places the route to 10.1.1.0/24 network via router R2 in the EIGRP topology table with an advertised distance of 1500.

Therefore, router R3 has two entries to reach 10.1.1.0/24 in its topology table. The EIGRP metric for router R3 to reach both routers R1 and R2 is 1000. This cost (1000) is added to the respective advertised distance from each router, resulting in the feasible distances from router R3 to reach network 10.1.1.0/24 shown in the figure.

Router R3 chooses the least-cost feasible distance, which is 2000, via router R1, and installs it in the IP routing table as the best route to reach 10.1.1.0/24. The EIGRP metric in the routing table is equal to the feasible distance from the EIGRP topology table. Router R1 is the successor for the route to 10.1.1.0/24.

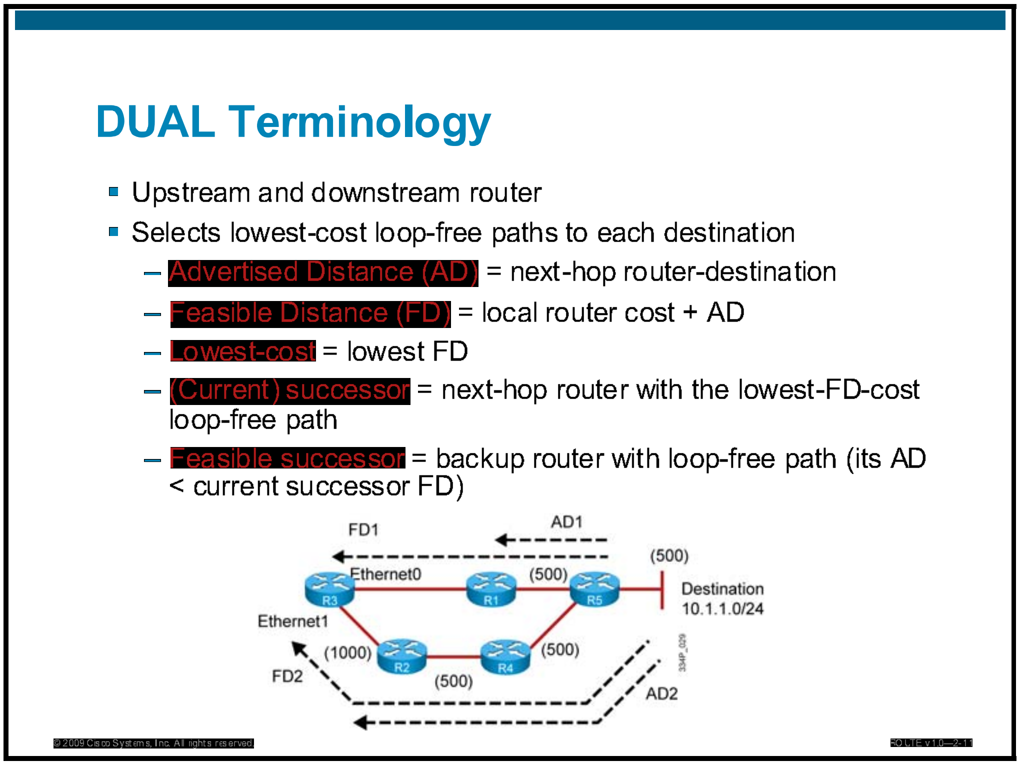

DUAL Terminology

- Upstream and downstream router
- Selects lowest-cost loop-free paths to each destination
  - Advertised Distance (AD) = next-hop router-destination
  - Feasible Distance (FD) = local router cost + AD
  - Lowest-cost = lowest FD
  - (Current) successor = next-hop router with the lowest-FD-cost loop-free path
  - Feasible successor = backup router with loop-free path (its AD < current successor FD)

ROUTE v1.0—2-11

DUAL uses the distance information, known as a metric or cost, to select efficient, loop-free paths.

The lowest-cost route is calculated by adding the cost between the next-hop router and the destination—referred to as the advertised distance (AD)—to the cost between the local router and the next-hop router. The sum of these costs is referred to as the feasible distance (FD).

A successor, also called a current successor, is a neighboring router that has a least-cost path to a destination (the lowest FD) that is guaranteed not to be part of a routing loop; successors are used for forwarding packets. Multiple successors can exist if they have the same FD. By default, up to four successors can be added to the routing table (the router can be configured to accept up to six per destination).

As well as keeping least-cost paths, DUAL keeps backup paths to each destination. The next-hop router for a backup path is called the feasible successor. To qualify as a feasible successor, a next-hop router must have an AD less than the FD of the current successor route.

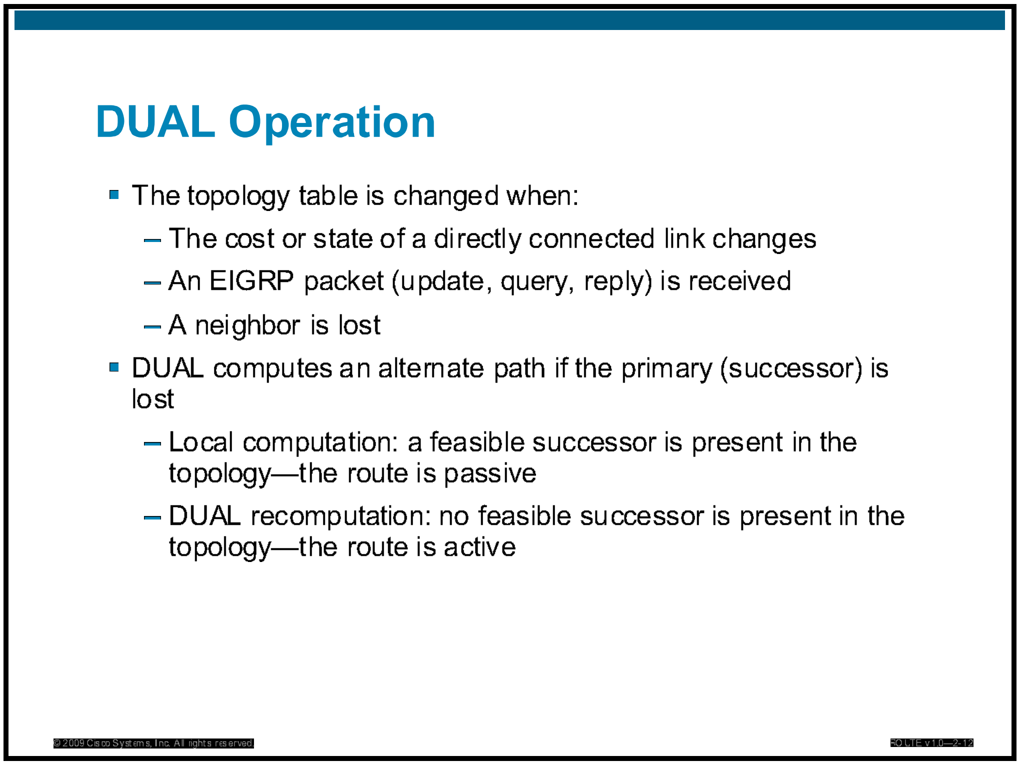

DUAL Operation

- The topology table is changed when:
  - The cost or state of a directly connected link changes
  - An EIGRP packet (update, query, reply) is received
  - A neighbor is lost
- DUAL computes an alternate path if the primary (successor) is lost
  - Local computation: a feasible successor is present in the topology—the route is passive
  - DUAL recomputation: no feasible successor is present in the topology—the route is active

ROUTE v1.0—2-12

If the route via the successor becomes invalid (because of a topology change) or if a neighbor changes the metric, DUAL checks for a feasible successor to the destination route. If one is found, DUAL uses it, thereby avoiding the need to recompute the route. If no suitable feasible successor exists, a recomputation must occur to determine the new successor. Although a recomputation is not processor-intensive, it does affect convergence time, so it is best to avoid any unnecessary recomputations.

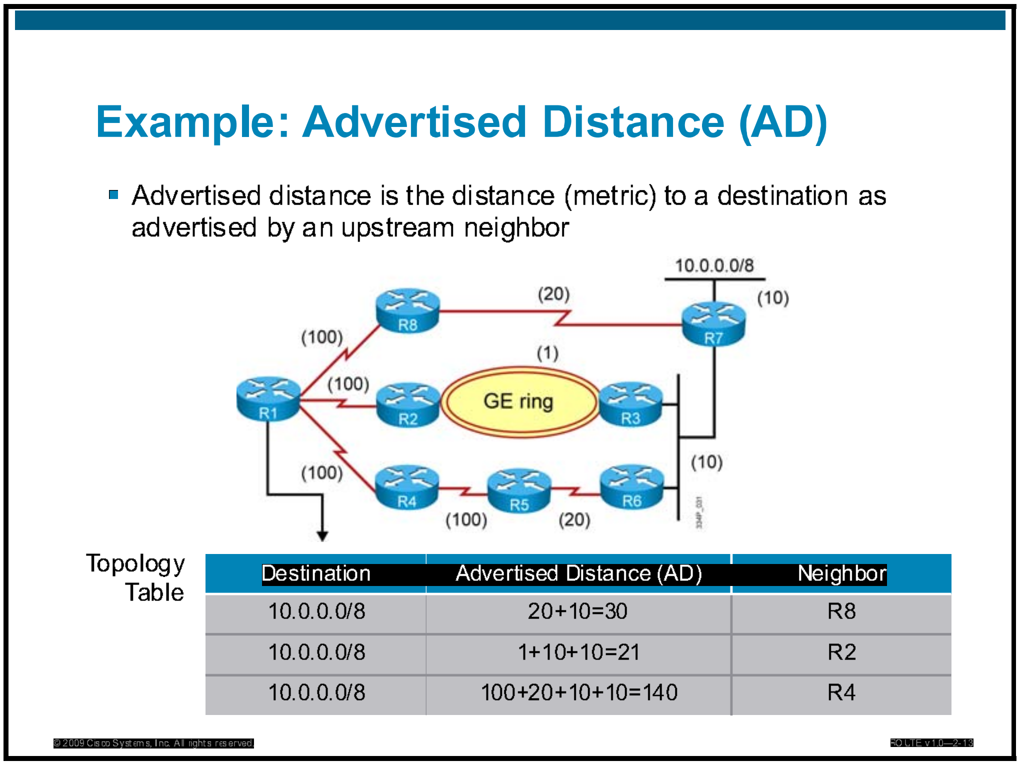

Example: Advertised Distance (AD)

- Advertised distance is the distance (metric) to a destination as advertised by an upstream neighbor

Topology Table

Destination	Advertised Distance (AD)	Neighbor
10.0.0.0/8	20+10=30	R8
10.0.0.0/8	1+10+10=21	R2
10.0.0.0/8	100+20+10+10=140	R4

ROUTE v1.0—2-13

Example: Advertised Distance (AD)

The figure above shows an example of how the AD is calculated. Router R1 has several options available to reach network 10.0.0.0/8. Routers R2, R4, and R8 each send an update to router R1. Each update contains an AD, which is the cost that router is advertising to reach network 10.0.0.0/8.

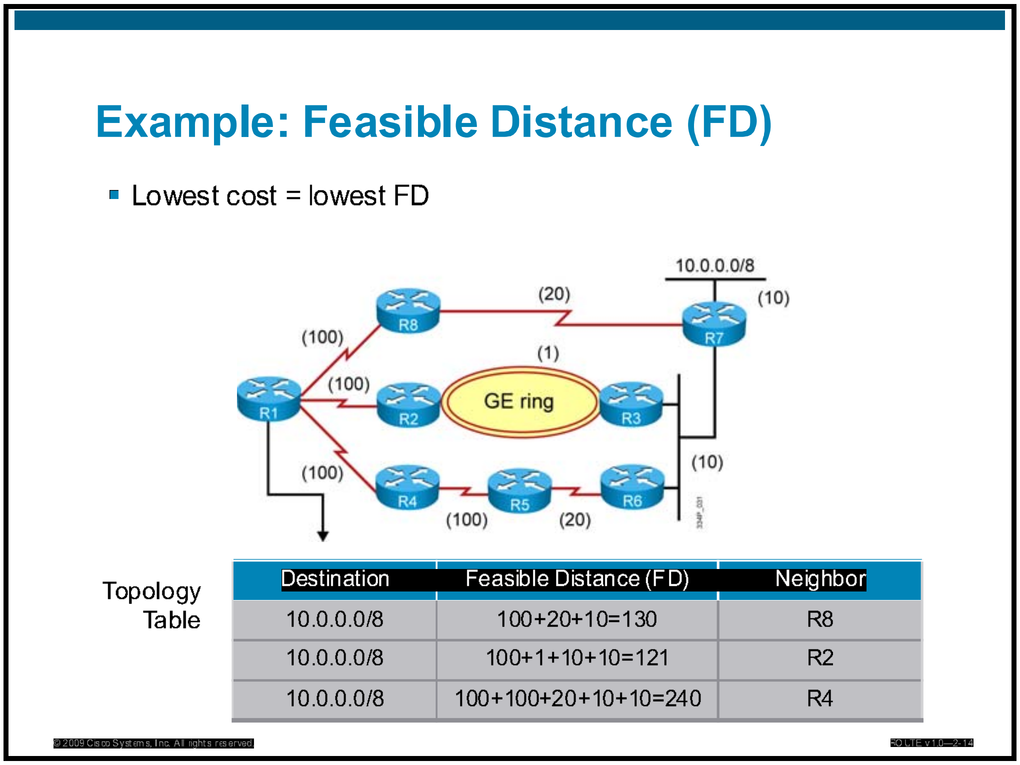

Example: Feasible Distance (FD)

- Lowest cost = lowest FD

Topology Table

Destination	Feasible Distance (FD)	Neighbor
10.0.0.0/8	100+20+10=130	R8
10.0.0.0/8	100+1+10+10=121	R2
10.0.0.0/8	100+100+20+10+10=240	R4

ROUTE v1.0—2-14

Example: Feasible Distance (FD)

The figure above shows an example of how the FD is calculated. Router R1 has several options available to reach network 10.0.0.0/8. Each update from the three neighbors has a different AD. By adding the cost of the local link to routers R2, R4, and R8 to the AD of each path, router R1 calculates the FD for each path to network 10.0.0.0/8.

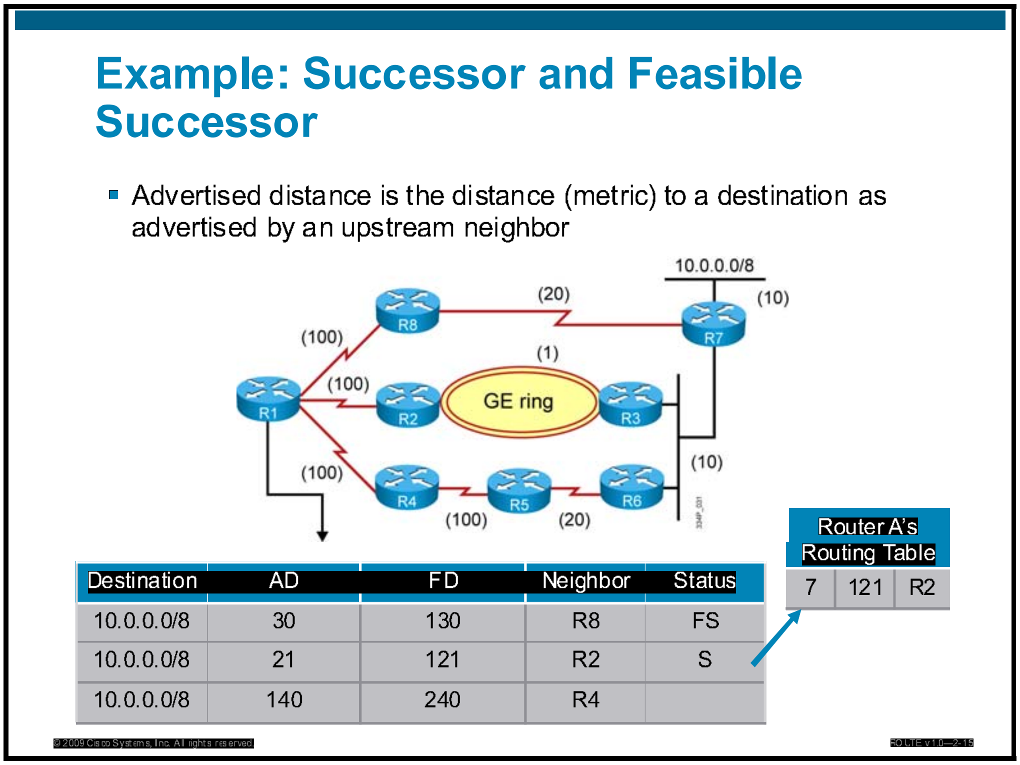

Example: Successor and Feasible Successor

- Advertised distance is the distance (metric) to a destination as advertised by an upstream neighbor

Destination	AD	FD	Neighbor	Status
10.0.0.0/8	30	130	R8	FS
10.0.0.0/8	21	121	R2	S
10.0.0.0/8	140	240	R4	

Router A's Routing Table

7	121	R2

ROUTE v1.0—2-15

Example: Successor and Feasible Successor

The figure above shows the successor and feasible successor on router R1 to network 10.0.0.0/8. Three paths exist for network 10.0.0.0/8. The FD and AD values are calculated for all three paths—the three candidates for the routing table. The candidate with the lowest FD value becomes the successor. If the AD for one of the remaining two candidates is lower than the FD on the successor route, then this candidate becomes a feasible successor. The route via router R2 becomes the successor. The route via router R8 becomes the feasible successor. Only the successor route goes into the routing table.

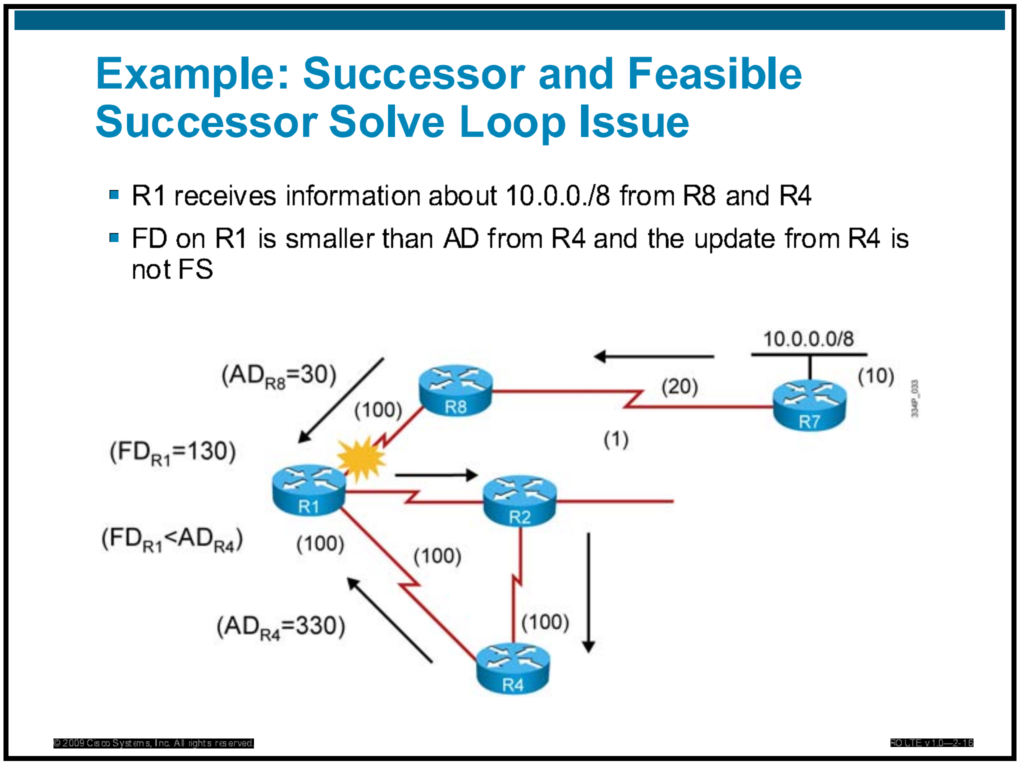

Example: Successor and Feasible Successor Solve Loop Issue

- R1 receives information about 10.0.0./8 from R8 and R4
- FD on R1 is smaller than AD from R4 and the update from R4 is not FS

ROUTE v1.0—2-16

The scenario on this slide shows how the DUAL algorithm solves the EIGRP routing loop issue. Router R8 with an AD of 30 sends the routing update about network 10.0.0.0/8. Router R1 receives an update, calculates the FD value (130), and sends an update to both neighbors. Routers R1, R2, and R4 are in a loop, and the picture above shows the update sent to router R2, which comes back to router R1. The update travelled via routers R2 and R4, and the AD of the update received on router R1 is 330. The value is higher than the FD value (130) calculated from the original update received from router R8. Because the FD of the update coming from R8 is smaller than the AD in the update coming from router R4, the route from the second update does not become feasible successor. This way DUAL solves the routing loop issue.

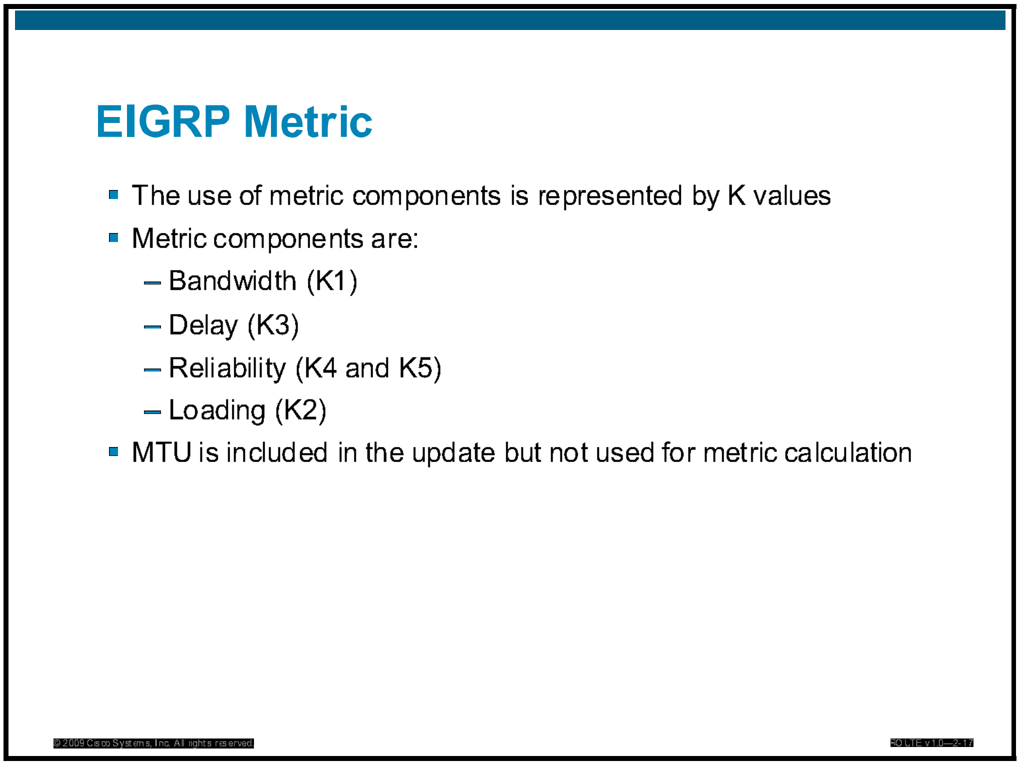

EIGRP Metric

- The use of metric components is represented by K values
- Metric components are:
  - Bandwidth (K1)
  - Delay (K3)
  - Reliability (K4 and K5)
  - Loading (K2)
- MTU is included in the update but not used for metric calculation

ROUTE v1.0—2-17

EIGRP uses the composite metric to determine the best path. The metric can be based on five criteria, but EIGRP uses only two of these criteria by default:

- **Bandwidth:** The smallest bandwidth between source and destination
- **Delay:** The cumulative interface delay along the path

The following criteria can be used, but are not recommended, because using them typically results in frequent recalculation of the topology table:

- **Reliability:** This value represents the worst reliability between source and destination, based on keepalives.
- **Loading:** This value represents the worst load on a link between the source and destination, computed based on the packet rate and the configured bandwidth of the interface.
- **Maximum Transmission Unit (MTU):** This represents the smallest MTU in the path. MTU is included in the EIGRP routing update but is not actually used in the metric calculation.

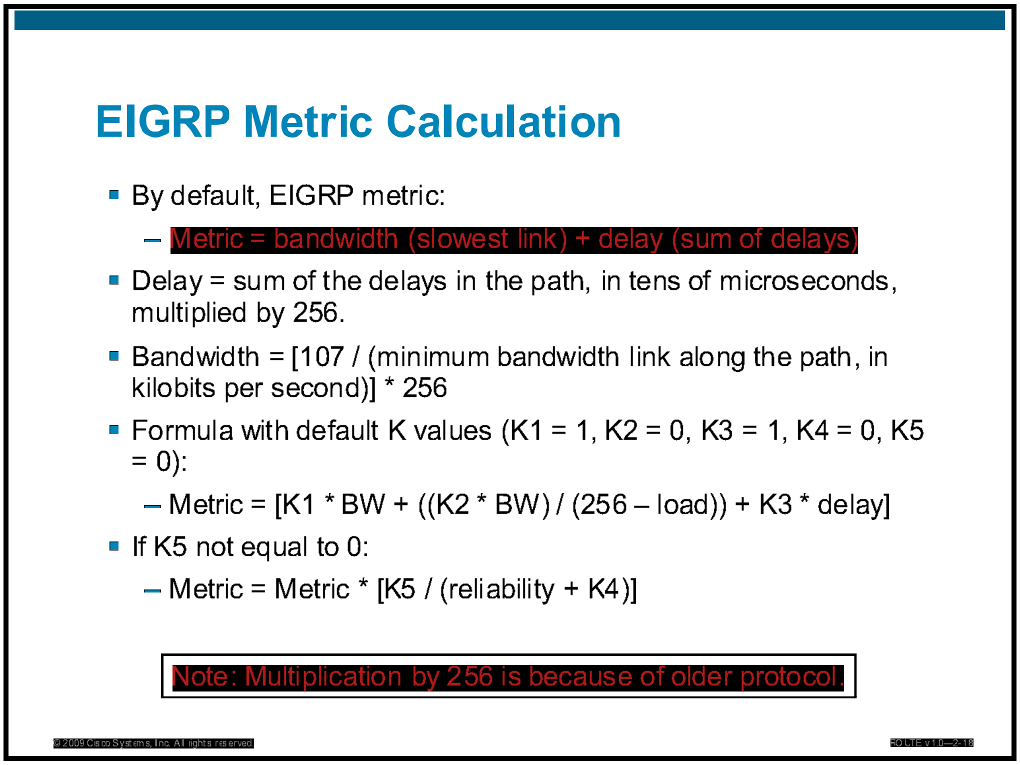

EIGRP Metric Calculation

- By default, EIGRP metric:
  - Metric = bandwidth (slowest link) + delay (sum of delays)
- Delay = sum of the delays in the path, in tens of microseconds, multiplied by 256.
- Bandwidth = [107 / (minimum bandwidth link along the path, in kilobits per second)] * 256
- Formula with default K values (K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5 = 0):
  - Metric = [K1 * BW + ((K2 * BW) / (256 – load)) + K3 * delay]
- If K5 not equal to 0:
  - Metric = Metric * [K5 / (reliability + K4)]

Note: Multiplication by 256 is because of older protocol.

ROUTE v1.0—2-18

EIGRP calculates the metric by adding together weighted values of different variables of the link to the network in question. The default constant weight values are K1 = K3 = 1, and K2 = K4 = K5 = 0.

In EIGRP metric calculations, when K5 is 0 (the default), variables (bandwidth, bandwidth divided by load, and delay) are weighted with the constants K1, K2, and K3. The following is the formula used:

- metric = (K1 * bandwidth)+ [(K2 * bandwidth) / (256 – load)] + (K3 * delay)

If these K values are equal to their defaults, the formula becomes the following:

- metric = (1 * bandwidth) + [(0 * bandwidth) / (256 – load)] + (1 * delay)
- metric = bandwidth + delay

If K5 is not equal to 0, the following additional operation is performed:

- metric = metric * [K5 / (reliability + K4)]

K values are carried in EIGRP hello packets. Mismatched K values can cause a neighbor to be reset (Only K1 and K3 are used, by default, in metric compilation.) These K values should be modified only after careful planning; changing these values can prevent your network from converging and is generally not recommended.

The format of the delay and bandwidth values used for EIGRP metric calculations is different from those displayed by the **show interface** command. The EIGRP delay value is the sum of the delays in the path, in tens of microseconds, multiplied by 256. The **show interface** command displays delay in microseconds. The EIGRP bandwidth is calculated using the minimum bandwidth link along the path, in kilobits per second. The value 10^7 is divided by this value, and then the result is multiplied by 256, because of the older protocol.

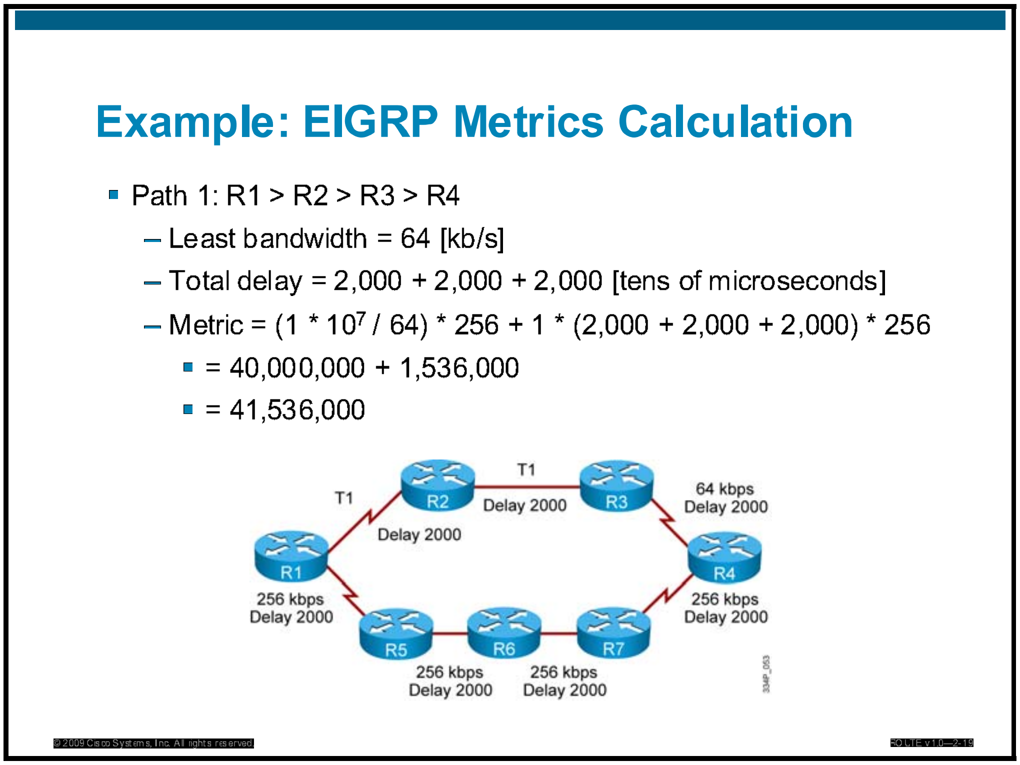

Example: EIGRP Metrics Calculation

- Path 1: R1 > R2 > R3 > R4
  - Least bandwidth = 64 [kb/s]
  - Total delay = 2,000 + 2,000 + 2,000 [tens of microseconds]
  - Metric = (1 * 10^7 / 64) * 256 + 1 * (2,000 + 2,000 + 2,000) * 256
    - = 40,000,000 + 1,536,000
    - = 41,536,000

T1
R2
Delay 2000
T1
R3
Delay 2000
64 kbps
Delay 2000
R1
R4
256 kbps
Delay 2000
256 kbps
Delay 2000
R5
R6
R7
256 kbps
Delay 2000
256 kbps
Delay 2000
334P_053

ROUTE v1.0—2-19

Example: EIGRP Metrics Calculation

Router R1 has two paths to reach networks behind router R4. The bandwidths (in kb/s) and the delays (in tens of microseconds) of the various links are also shown in the figure.

The least bandwidth along the top path (R1 > R2 > R3 > R4) is 64 kb/s. The EIGRP bandwidth calculation for this path is as follows:

- Bandwidth = (10^7/ least bandwidth in kb/s) * 256
- Bandwidth = (10,000,000 / 64) * 256 = 156,250 * 256 = 40,000,000

The delay through the top path is as follows:

- Delay = [(delay R1 → R2) + (delay R2 → R3) + (delay R3 → R4)] * 256
- Delay = [2000 + 2000 + 2000] * 256
- Delay = 1,536,000

Therefore, the EIGRP metric calculation for the top path is as follows:

- Metric = bandwidth + delay
- Metric = 40,000,000 + 1,536,000
- Metric = 41,536,000

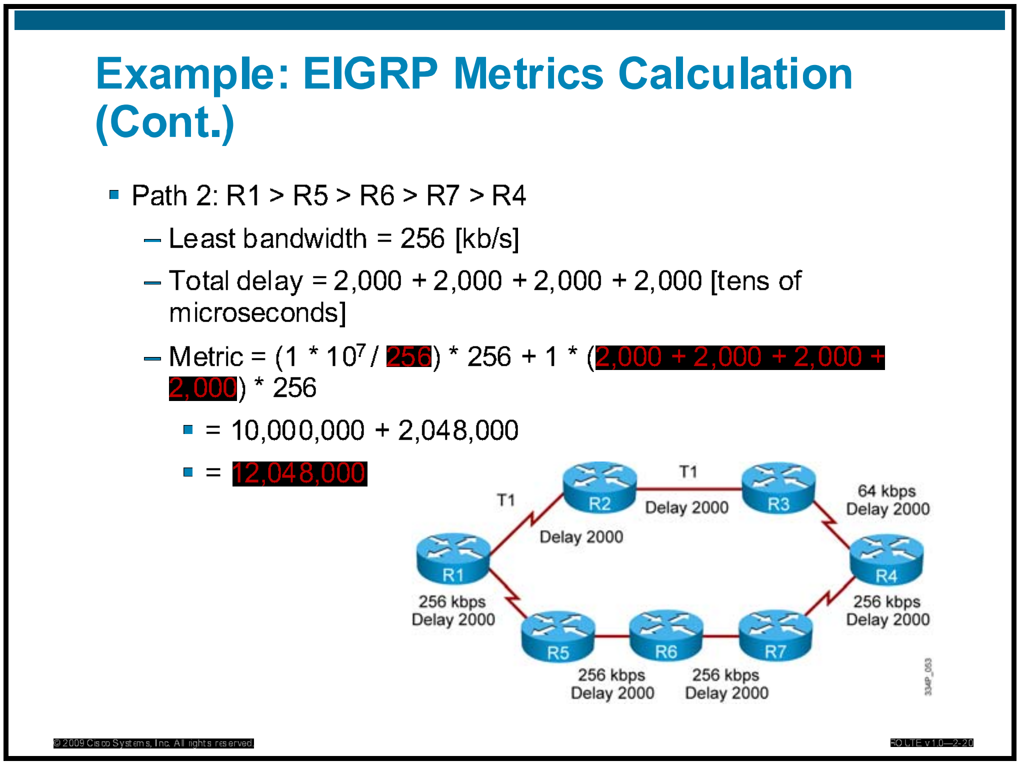

Example: EIGRP Metrics Calculation (Cont.)

- Path 2: R1 > R5 > R6 > R7 > R4
  - Least bandwidth = 256 [kb/s]
  - Total delay = 2,000 + 2,000 + 2,000 + 2,000 [tens of microseconds]
  - Metric = (1 * 10^7 / 256) * 256 + 1 * (2,000 + 2,000 + 2,000 + 2,000) * 256
    - = 10,000,000 + 2,048,000
    - = 12,048,000

ROUTE v1.0—2-20

The least bandwidth along the lower path (R1 → R5 → R6 → R7 → R4) is 256 kb/s. The EIGRP bandwidth calculation for this path is as follows:

- Bandwidth = (10^7 / least bandwidth in kb/s) * 256
- Bandwidth = (10,000,000 / 256) * 256 = 10,000,000

The delay through the lower path is as follows:

- Delay = [(delay R1 → R5) + (delay R5 → R6) + (delay R6 → R7) + (delay R7 → R4)] * 256
- Delay = [2000 + 2000 + 2000 + 2000] * 256
- Delay = 2,048,000

Therefore, the EIGRP metric calculation for the lower path is as follows:

- Metric = bandwidth + delay
- Metric = 10,000,000 + 2,048,000
- Metric = 12,048,000

Router R1 therefore chooses the lower path, with a metric of 12,048,000 over the top path, with a metric of 41,536,000. Router R1 installs the lower path, with a next-hop router of R5 and a metric of 12,048,000, in the IP routing table.

The bottleneck along the top path, the 64-kb/s link, can explain why the router takes the lower path. This slow link means that the rate of transfer to Router R4 can be at a maximum of 64 kb/s. Along the lower path, the lowest speed is 256 kb/s, meaning the throughput rate can be as high as that speed. Therefore, the lower path represents a better choice, for example, for moving large files quickly.

Planning and Documenting for EIGRP

This topic describes how to plan, implement, and document the EIGRP deployment.

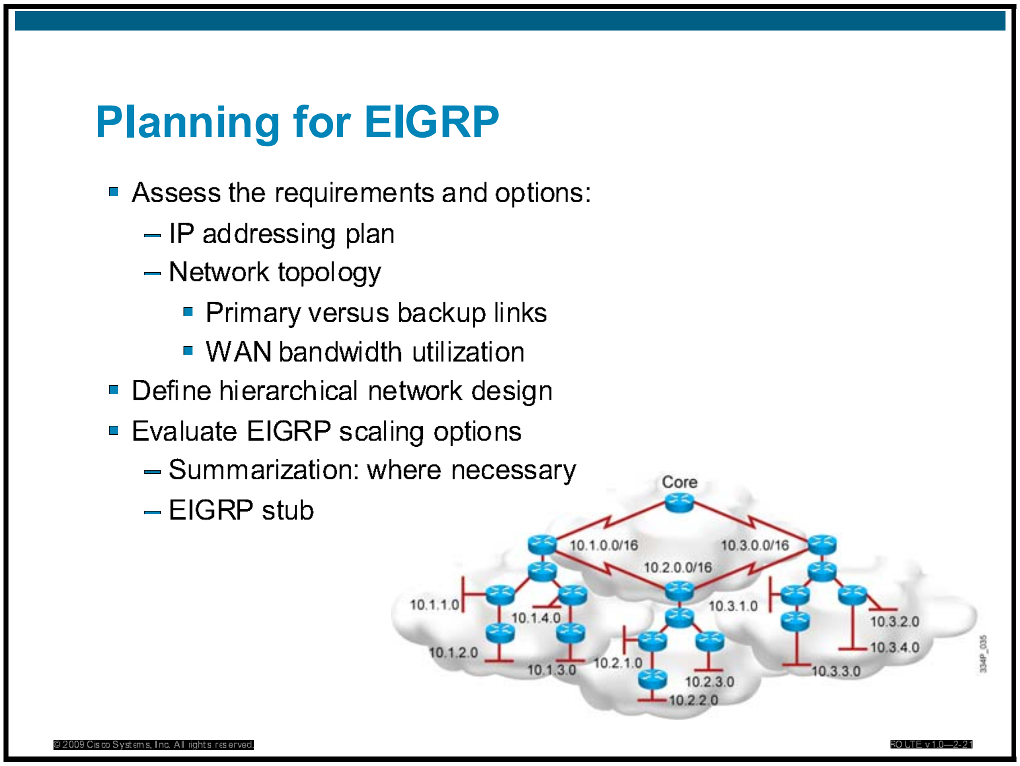

Planning for EIGRP

- Assess the requirements and options:
  - IP addressing plan
  - Network topology
    - Primary versus backup links
    - WAN bandwidth utilization
- Define hierarchical network design
- Evaluate EIGRP scaling options
  - Summarization: where necessary
  - EIGRP stub

ROUTE v1.0—2-21

The EIGRP routing protocol implementation depends on specific needs and topologies. When preparing to deploy EIGRP routing in a network, the existing state and requirements first need to be gathered and different deployment options considered:

- The IP addressing plan determines how EIGRP can be deployed and how well the EIGRP deployment might scale. Thus, a detailed IP addressing plan along with IP subnetting information must be collected. A solid IP addressing plan should enable the usage of EIGRP summarization, making it easier to scale the network and optimize EIGRP behavior.

- A network topology consists of links connecting the network equipment (routers, switches, and so on). A detailed network topology plan should be presented, in order to assess EIGRP scalability requirements and determine which EIGRP features might be required (for example EIGRP stub routing).

- EIGRP can be used to employ traffic engineering, which helps with efficient bandwidth utilization and enables the administrator to have control over the traffic patterns. By changing the interface metrics, EIGRP traffic engineering can be deployed to improve bandwidth utilization.

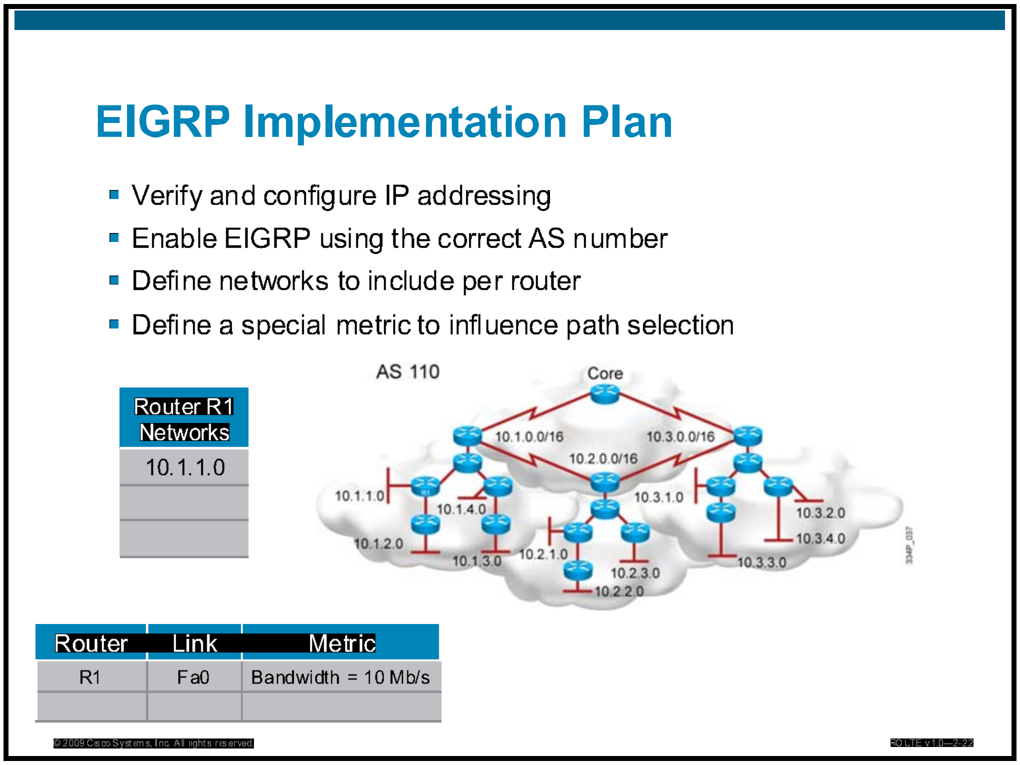

EIGRP Implementation Plan

- Verify and configure IP addressing
- Enable EIGRP using the correct AS number
- Define networks to include per router
- Define a special metric to influence path selection

Router R1 Networks
10.1.1.0

Router	Link	Metric
R1	Fa0	Bandwidth = 10 Mb/s

Once you have assessed the requirements, you can create the implementation plan. The information necessary to implement EIGRP routing is:

- The IP addresses (or, more precisely, the networks) that need to be included and advertised by EIGRP
- The correct autonomous system (AS) number used to enable the EIGRP process, which must be the same on the routers in the EIGRP domain.
- A list of routers where EIGRP must be enabled, along with the connected networks that need to be advertised (per individual router)
- A listing in the table of any specific metric that needs to be applied to certain interfaces in order to deploy EIGRP traffic engineering, along with the interface where the metric needs to be applied

When an implementation plan is created, a list of tasks for each router in the network must be defined:

- Enable the EIGRP routing protocol.
- Configure the proper network statements based on the information collected.

You can also apply the metric to proper interfaces, if you wish.

After implementation, you should confirm that EIGRP is deployed properly on each router:

- Verify the setup of the EIGRP neighbor relationship or relationships.
- Verify that the EIGRP topology table is populated with the necessary information.
- Verify that the IP routing table is populated with the necessary information.
- Verify that there is connectivity in the network.
- Verify that EIGRP behaves as expected when the topology changes (test link failure and router failure events).

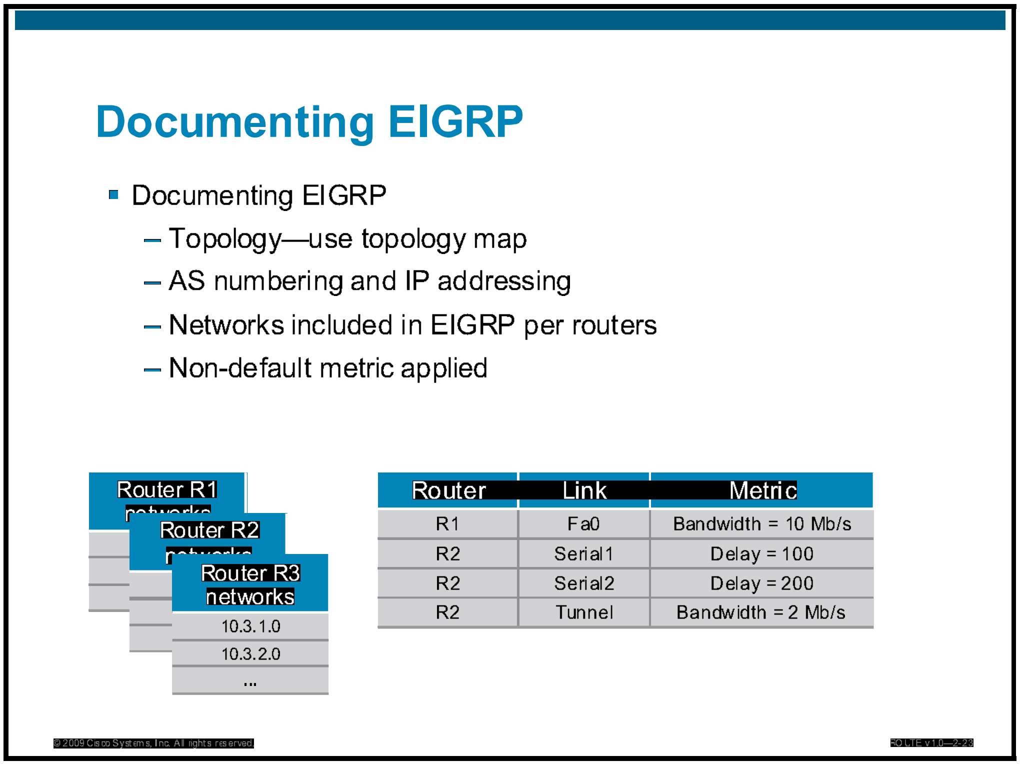

Documenting EIGRP

- Documenting EIGRP
  - Topology—use topology map
  - AS numbering and IP addressing
  - Networks included in EIGRP per routers
  - Non-default metric applied

Router R1 networks

Router R2 networks

Router R3 networks

10.3.1.0

10.3.2.0

...

Router	Link	Metric
R1	Fa0	Bandwidth = 10 Mb/s
R2	Serial1	Delay = 100
R2	Serial2	Delay = 200
R2	Tunnel	Bandwidth = 2 Mb/s

After a successful EIGRP deployment, you should document the solution in order to keep the information about the deployment available for future reference. The implementation plan itself is only half of the information. In order to complete the documentation, you must include the verification process and its results, as well.

Implementing Basic EIGRP

This topic describes how to plan and implement the basic EIGRP configuration.

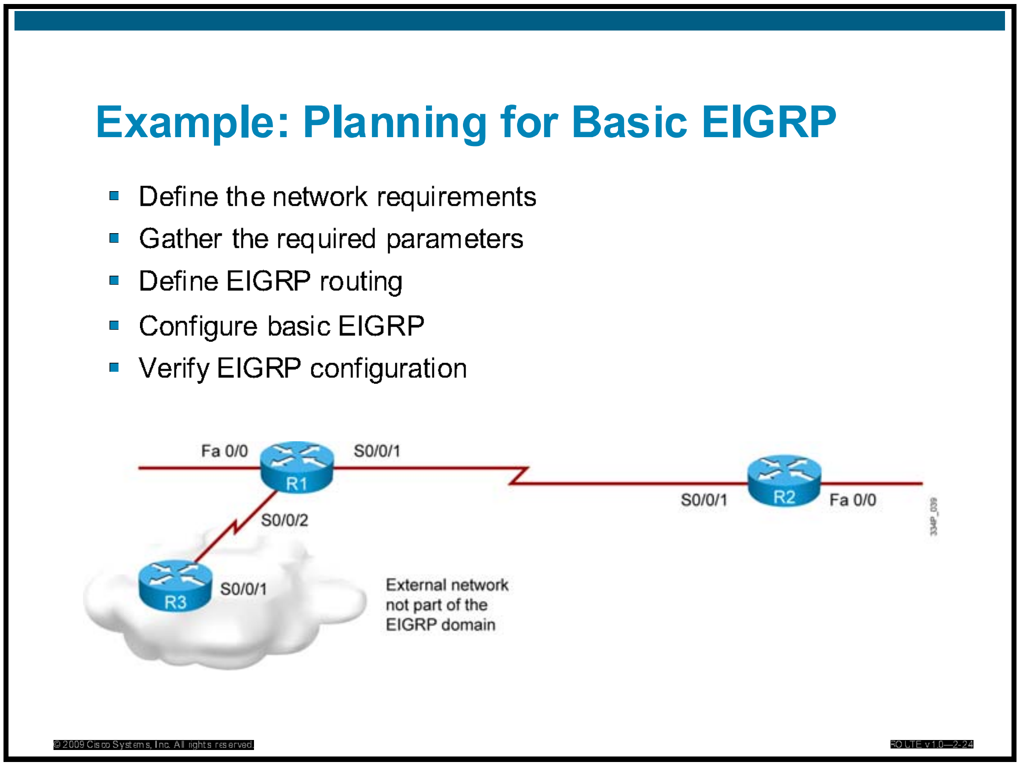

Example: Planning for Basic EIGRP

- Define the network requirements
- Gather the required parameters
- Define EIGRP routing
- Configure basic EIGRP
- Verify EIGRP configuration

ROUTE v1.0—2-24

For the example in the figure above, prepare an implementation plan to configure basic EIGRP and proceed with the configuration.

When you plan for the basic EIGRP configuration, you must ensure that your plan includes the:

- Define the network requirements
- Gather the required parameters
- Define the EIGRP routing
- Configure the basic EIGRP
- Verify the EIGRP configuration

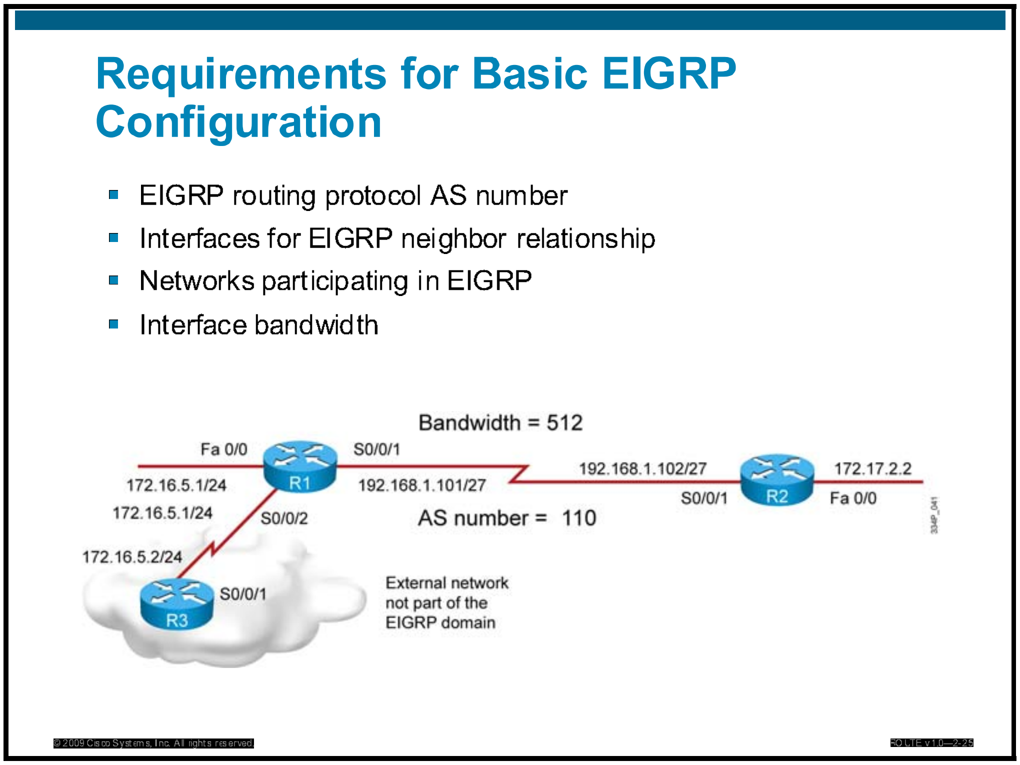

Requirements for Basic EIGRP Configuration

- EIGRP routing protocol AS number
- Interfaces for EIGRP neighbor relationship
- Networks participating in EIGRP
- Interface bandwidth

ROUTE v1.0—2-25

The network in the figure consists of three routers. Routers R1 and R2 are in the same EIGRP autonomous system (AS). Router R3 represents an external network that is not part of the EIGRP AS 110. Requirements for the basic EIGRP configuration are:

- **EIGRP routing protocol AS number**: Routers in the same EIGRP domain must have the same AS number, as each EIGRP process must be started with the same AS number. The AS number 110 is used in our example.
- **Interfaces for EIGRP neighbor relationship**: Interfaces included in the EIGRP routing protocol will exchange routing updates and other packets between their neighbors. You must define the interfaces to show which networks are part of the EIGRP process. Both routers (R1 and R2) have one serial and one Fast Ethernet interface included in the EIGRP process. IP addressing is defined as listed in the figure.
- **Networks participating in EIGRP**: EIGRP routers must advertise their local networks to all neighbors. All the interfaces and their networks must be defined. Both routers (R1 and R2) are advertising directly to connected networks that are part of the EIGRP domain.
- **Interface Bandwidth**: Interface bandwidth is changing the metric of the link. In order to influence path selection, interface bandwidth must be defined properly. The real bandwidth on the serial link between routers R1 and R2 is 512 kb/s and the proper configuration must be applied in order to reflect the real bandwidth. This will result in proper selection of preferred routes in the EIGRP process.

Note Default bandwidth for a serial interface is 1544 kb/s.

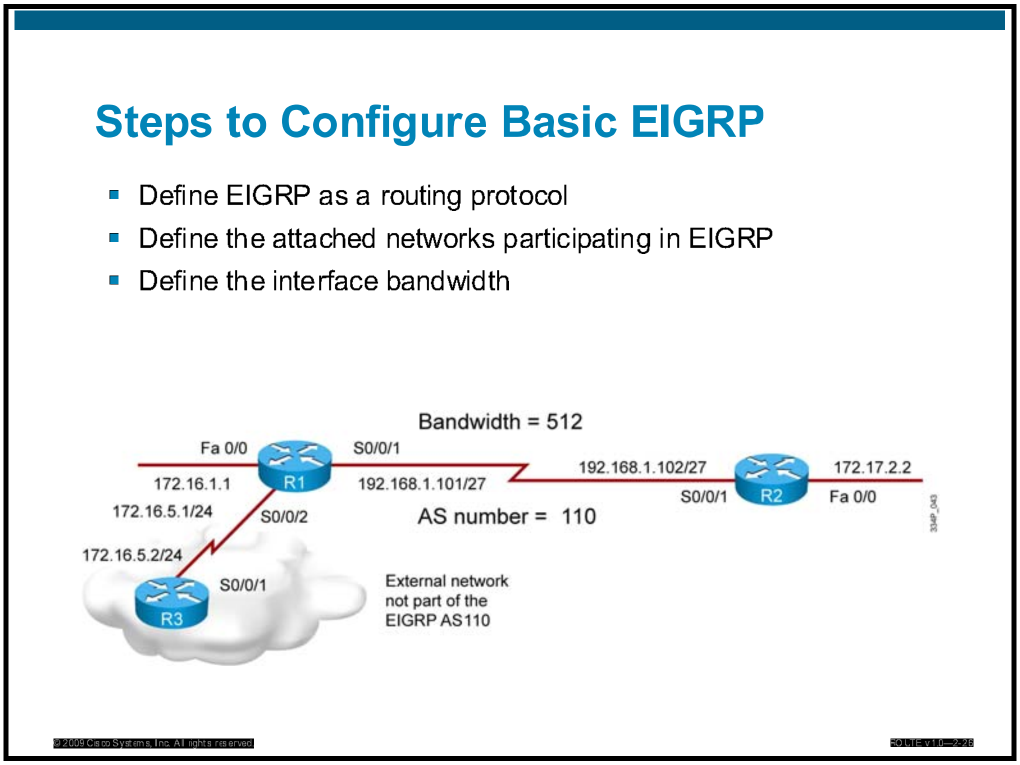

Steps to Configure Basic EIGRP

- Define EIGRP as a routing protocol
- Define the attached networks participating in EIGRP
- Define the interface bandwidth

ROUTE v1.0—2-26

Once you have defined all of the required information, an implementation plan showing the following tasks is required to configure a basic EIGRP configuration:

- Define EIGRP as a routing protocol
- Define the attached networks participating in EIGRP
- Define the interface bandwidth

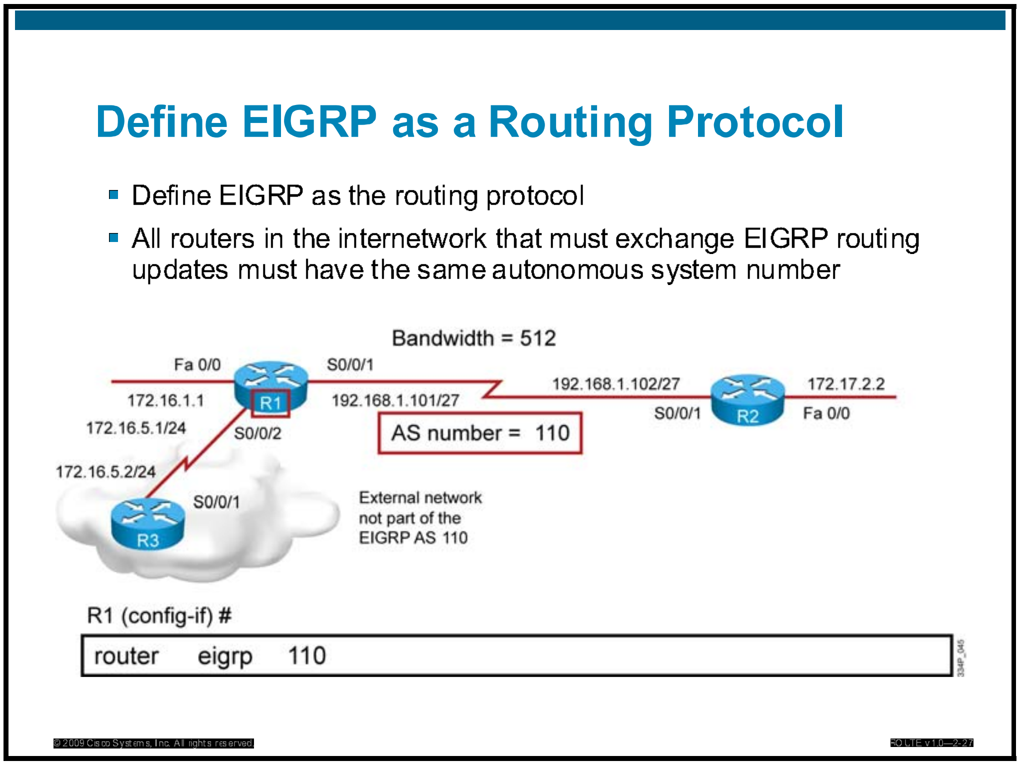

Define EIGRP as a Routing Protocol

- Define EIGRP as the routing protocol
- All routers in the internetwork that must exchange EIGRP routing updates must have the same autonomous system number

R1 (config-if) #

```
router    eigrp    110
```

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-27

During the first step of a basic EIGRP configuration, you must define EIGRP as a routing protocol. You must specify an AS number that identifies the routes to the other EIGRP routers. Be aware that all routers in the same EIGRP domain must have the same AS number.

Use the **router eigrp** ***110*** *c*ommand to configure the EIGRP routing protocol and add any subsequent subcommands belonging to this routing process. This command also identifies the local AS to which this router belongs. AS 110 is used as an example. The command enters router configuration mode.

Note You can configure more than one EIGRP autonomous system on the same router, but you should configure only one EIGRP autonomous system in any single autonomous system.

For more details about the **router eigrp** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

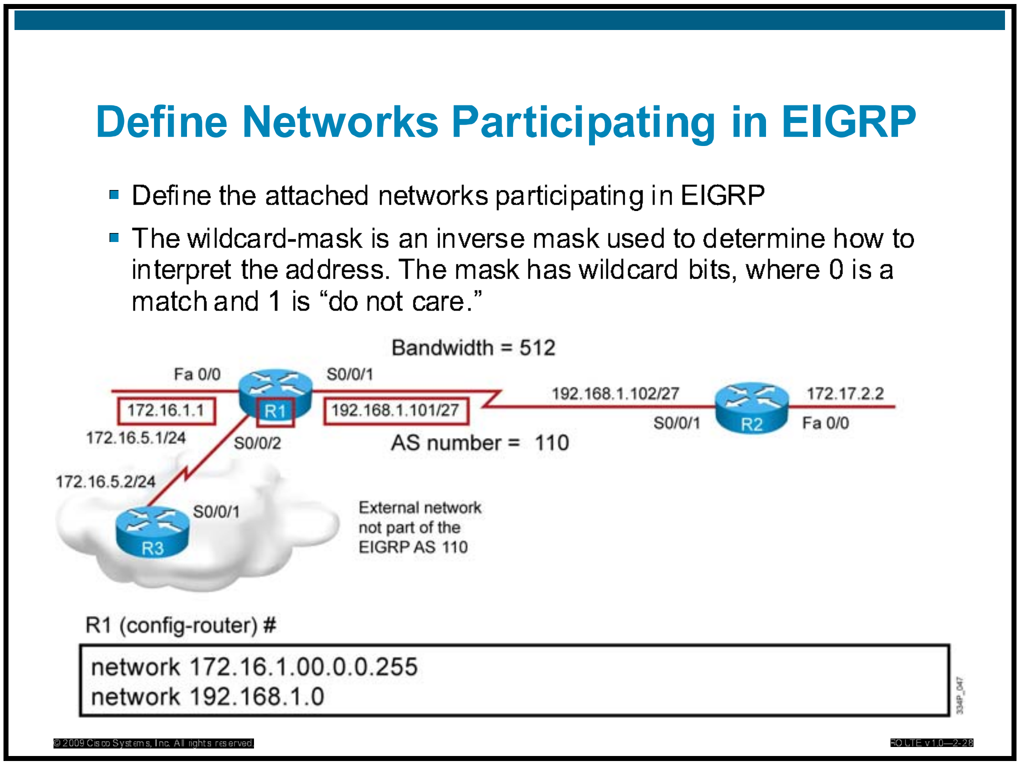

Define Networks Participating in EIGRP

- Define the attached networks participating in EIGRP
- The wildcard-mask is an inverse mask used to determine how to interpret the address. The mask has wildcard bits, where 0 is a match and 1 is “do not care.”

R1 (config-router) #

```
network 172.16.1.00.0.0.255
network 192.168.1.0
```

In order to start sending and receiving EIGRP routing updates, networks of directly connected interfaces must be defined. Only the network statements for interfaces where the router will send and receive updates need to be configured.

Use the **network** ***172.16.1.0 0.0.0.255*** command in router configuration mode to specify the network for an EIGRP routing process.

When the **network** command is configured for an EIGRP routing process, the router matches one or more local interfaces. The **network** command matches only local interfaces that are configured with addresses that are within the same subnet as the address that has been configured with the **network** command. The router then establishes neighbors through the matched interfaces. There is no limit to the number of network statements (**network** commands) that you can configure on a router.

Note The wildcard mask in the network command is optional. It is an inverse mask used to determine how to interpret the network number. The mask has wildcard bits, where 0 indicates a match and 1 indicates the bits which are not relevant. For example, 0.0.255.255 indicates a match in the first two octets.

For more details about the **network** (EIGRP) command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

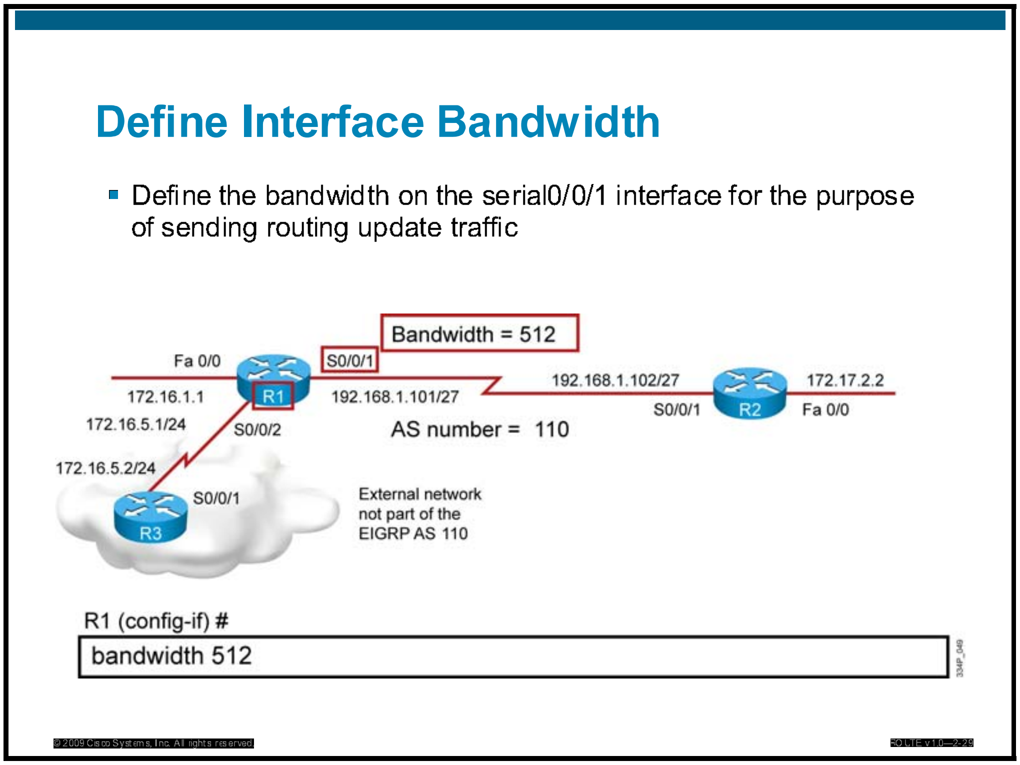

Define Interface Bandwidth

- Define the bandwidth on the serial0/0/1 interface for the purpose of sending routing update traffic

R1 (config-if) #

```
bandwidth 512
```

ROUTE v1.0—2-29

EIGRP uses the minimum path bandwidth to determine a routing metric. The TCP protocol adjusts the initial retransmission parameters based on the apparent bandwidth of the outgoing interface.

Use the **bandwidth** ***512*** command in interface configuration mode to specify or change the informational value used for an EIGRP routing process. If you do not change the bandwidth for the interfaces, EIGRP assumes that the default bandwidth on the serial link is the T1. If the link is actually slower, the router might not be able to converge or routing updates might be lost.

The **bandwidth** command sets an informational parameter only; you cannot adjust the actual bandwidth of an interface with this command. For some media, such as Ethernet, the bandwidth is fixed; for other media, such as serial lines, you can change the actual bandwidth by adjusting hardware. For both classes of media, you can use the **bandwidth** configuration command to communicate the current bandwidth to the higher-level protocols.

Note At higher bandwidths, the value you configure with the bandwidth command is not what is displayed by the **show interface** command. The value shown is used in EIGRP updates and computing the load.

For generic serial interfaces such as PPP or High-Level Data Link Control (HDLC), set the bandwidth to the line speed. For Frame Relay on point-to-point interfaces, set the bandwidth to the committed information rate (CIR). For Frame Relay multipoint connections, set the bandwidth to the sum of all CIRs, or, if the permanent virtual circuits (PVCs) have different CIRs, then set the bandwidth to the lowest CIR multiplied by the number of PVCs on the multipoint connection.

For more details about the **bandwidth** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

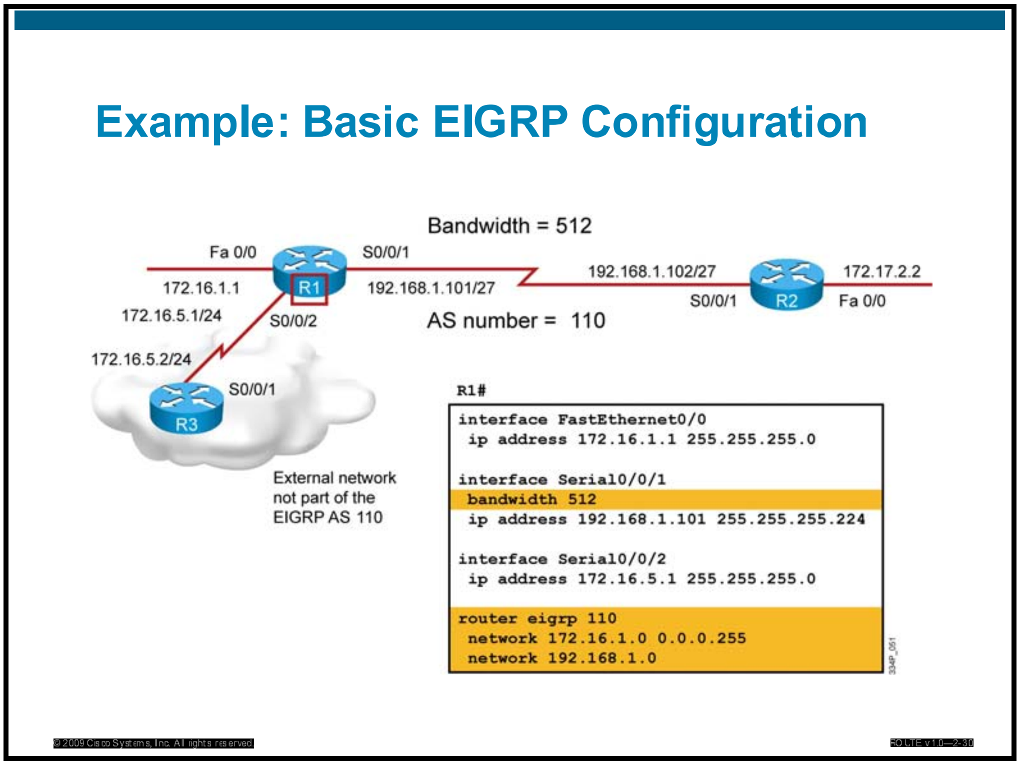

Example: Basic EIGRP Configuration

```
R1#
interface FastEthernet0/0
 ip address 172.16.1.1 255.255.255.0

interface Serial0/0/1
 bandwidth 512
 ip address 192.168.1.101 255.255.255.224

interface Serial0/0/2
 ip address 172.16.5.1 255.255.255.0

router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

Example: Basic EIGRP Configuration

On router R1, EIGRP is enabled in autonomous system 110. The **network 172.16.1.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R1 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.16.1.0/24 subnet will participate in EIGRP. However, the full class B network 172.16.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the Serial 0/0/1 interface, and allows router R1 to advertise this network.

If you do not use the optional wildcard mask, the EIGRP process assumes that all directly connected networks that are part of the overall major network will participate in the EIGRP routing process, and EIGRP will attempt to establish EIGRP neighbor relationships from each interface that is part of that Class A, B, or C major network.

Use the optional wildcard mask to identify a specific IP address, subnet, or network. The router interprets the network number using the wildcard mask to determine which connected networks will participate in the EIGRP routing process. If specifying an interface address, use the mask 0.0.0.0 to match all four octets of the address. An address and wildcard mask combination of 0.0.0.0 255.255.255.255 matches all interfaces on the router.

In the example above, router R1 is connected to router R3, which is external to AS 110. Network 172.16.5.0 is used on the link between router R1 and R3, but the statement for network 172.16.1.0 on router R1 is using a wildcard mask and router R1 does not try to form an adjacency with the router R3. Without the wildcard mask, router R1 would send EIGRP packets to the external network (toward router R3), which would waste bandwidth and CPU cycles and provide unnecessary information to the external network. The wildcard mask in the example above tells EIGRP to establish a relationship with EIGRP routers from interfaces that are part of subnet 172.16.1.0/24 only.

Note The configuration of router R2 is identical to that of R1 described here. The only difference in terms of the EIGRP configuration is that the advertised network of the Fast Ethernet interface is different. The Fast Ethernet interface of router R2 is using a different subnet on the link.

Summary

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

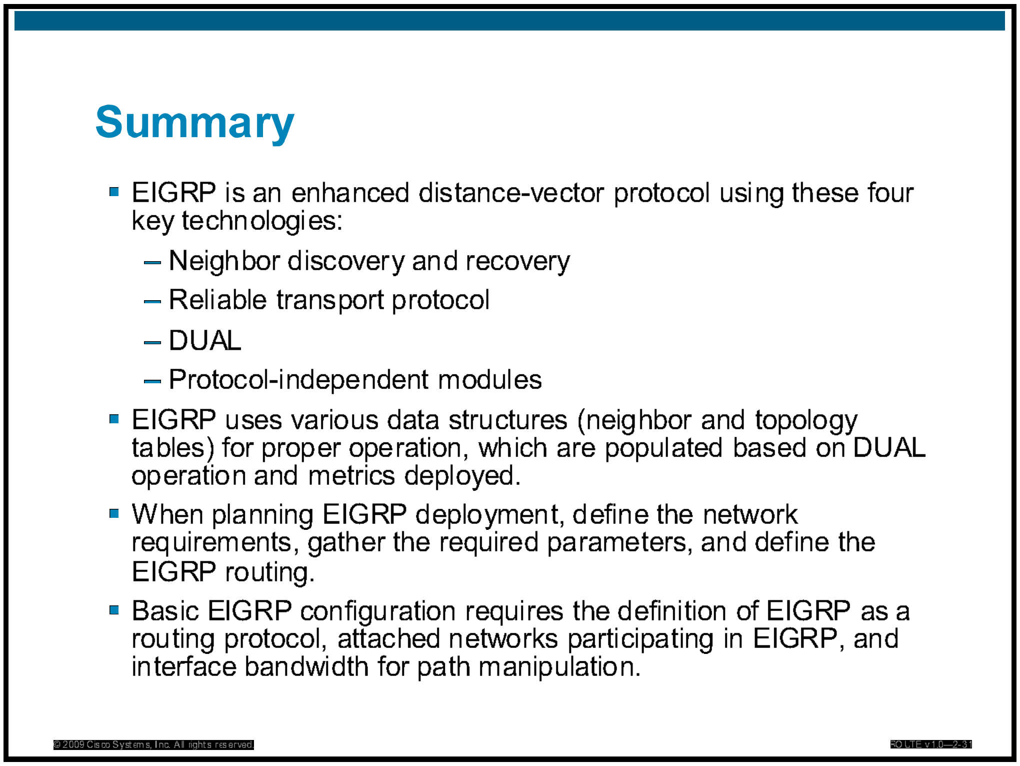

Summary

- EIGRP is an enhanced distance-vector protocol using these four key technologies:
  - Neighbor discovery and recovery
  - Reliable transport protocol
  - DUAL
  - Protocol-independent modules
- EIGRP uses various data structures (neighbor and topology tables) for proper operation, which are populated based on DUAL operation and metrics deployed.
- When planning EIGRP deployment, define the network requirements, gather the required parameters, and define the EIGRP routing.
- Basic EIGRP configuration requires the definition of EIGRP as a routing protocol, attached networks participating in EIGRP, and interface bandwidth for path manipulation.

Lesson 2

Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture

Overview

To assist in verification, this lesson introduces various Cisco IOS software **show** commands and defines the key fields in each. For a scalable Enhanced Interior Gateway Routing Protocol (EIGRP) network, configuring manual route summarization at key points on the internetwork is vital when implementing an optimized network configuration.

Knowing the correct commands to use when you configure EIGRP helps to ensure that migration to this routing protocol is smooth and quick. Understanding which **show** command to use when verifying the EIGRP configuration saves valuable time. This lesson also provides advanced configuration options for EIGRP, including route summarization, passive interfaces, and default network origination.

Objectives

Upon completing this lesson, you will be able to describe how to verify and implement EIGRP routing. This ability includes being able to meet these objectives:

- Verify EIGRP routes for IPv4.
- Verify EIGRP operation for IPv4.
- Use the **passive-interface** command with EIGRP.
- Advertise an IP default network in EIGRP.
- Determine summary boundaries.
- Utilize manual route summarization.

Verify EIGRP Routes for IPv4

This topic describes how EIGRP configuration and operation can be verified using the appropriate commands.

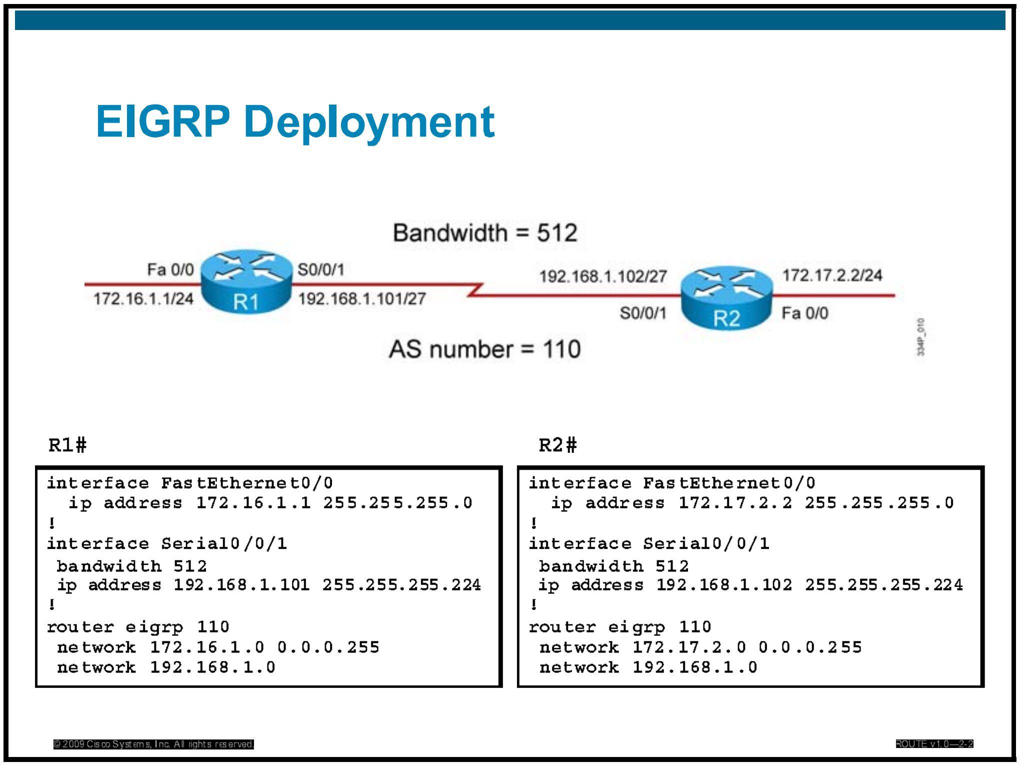

EIGRP Deployment

R1#

```
interface FastEthernet0/0
  ip address 172.16.1.1 255.255.255.0
!
interface Serial0/0/1
 bandwidth 512
 ip address 192.168.1.101 255.255.255.224
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

R2#

```
interface FastEthernet0/0
  ip address 172.17.2.2 255.255.255.0
!
interface Serial0/0/1
 bandwidth 512
 ip address 192.168.1.102 255.255.255.224
!
router eigrp 110
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
```

ROUTE v1.0—2-2

Router R1 has EIGRP enabled in autonomous system 110. The **network 172.16.1.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R1 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.16.1.0/24 subnet will participate in EIGRP. Note, however, the full class B network 1 72.16.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the Serial 0/0/1 interface and allows router R1 to advertise this network.

Router R2 has EIGRP enabled in autonomous system 110. The **network 172.17.2.0 0.0.0.255** command starts EIGRP on the Fast Ethernet 0/0 interface and allows router R2 to advertise this network. With the wildcard mask used, this command specifies that only interfaces on the 172.17.2.0/24 subnet will participate in EIGRP. Note, however, the full class B network 172.17.0.0 will be advertised, because EIGRP automatically summarizes routes on the major network boundary by default. The **network 192.168.1.0** command starts EIGRP on the serial 0/0/1 interface and allows router R2 to advertise this network.

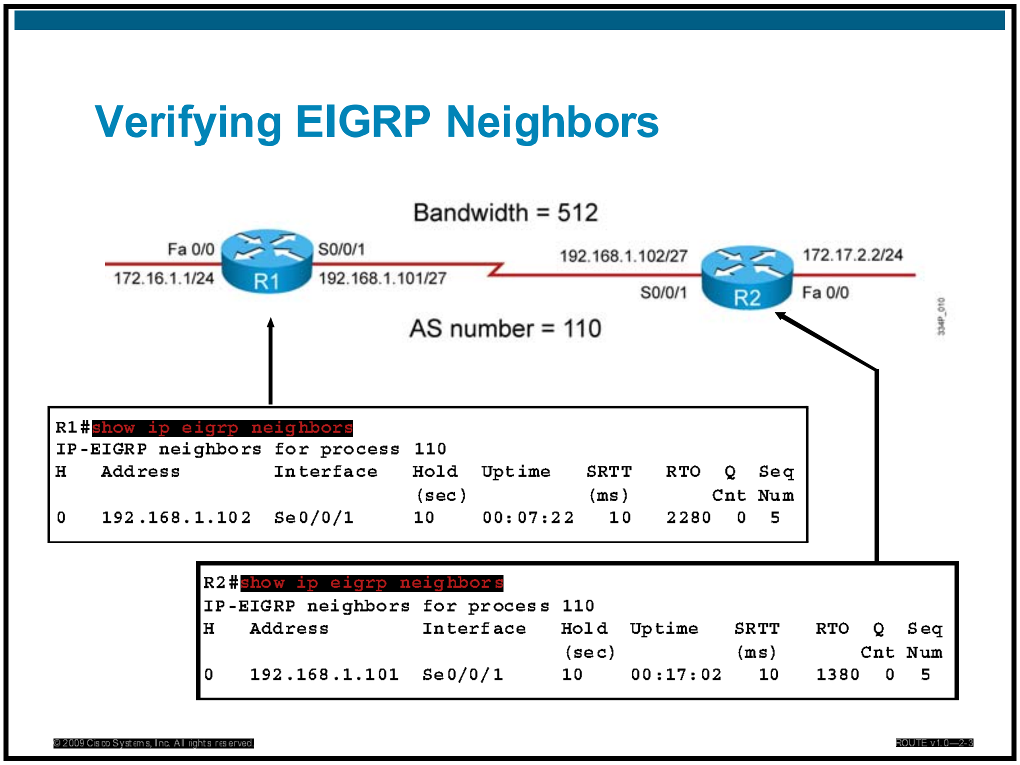

Verifying EIGRP Neighbors

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface    Hold  Uptime     SRTT    RTO  Q   Seq
                                 (sec)            (ms)         Cnt Num
0   192.168.1.102   Se0/0/1      10    00:07:22    10     2280  0   5
```

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface    Hold  Uptime     SRTT    RTO  Q   Seq
                                 (sec)            (ms)         Cnt Num
0   192.168.1.101   Se0/0/1      10    00:17:02    10     1380  0   5
```

ROUTE v1.0—2-3

Both routers R1 and R2 are configured with the EIGRP routing protocol in autonomous system (AS) 110 and are advertising their networks to the neighbors. Before updates are sent, EIGRP is building the EIGRP neighbor table. The EIGRP neighbor table displays the neighbors discovered by EIGRP, including the IP address and interface on which each neighbor is reachable.

The EIGRP neighbor table can be displayed using the **show ip eigrp neighbors** command. Use this command to determine when neighbors become active or inactive. You can also use it for debugging certain types of transport problems.

The outputs of the command in the figure list currently used neighbor relationships—router R1 has formed an adjacency with router R2 and vice versa.

For more details about the **show ip eigrp neighbors** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

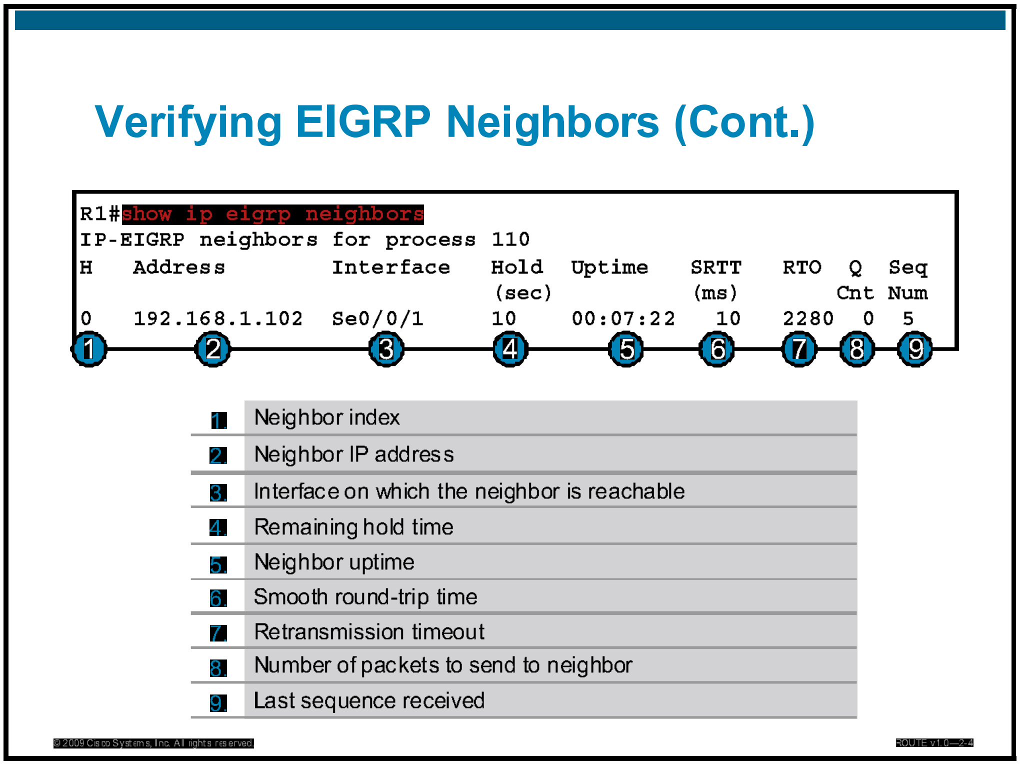

Verifying EIGRP Neighbors (Cont.)

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT    RTO  Q   Seq
                                  (sec)           (ms)         Cnt Num
0   192.168.1.102   Se0/0/1       10    00:07:22   10     2280  0   5
```

1	Neighbor index
2	Neighbor IP address
3	Interface on which the neighbor is reachable
4	Remaining hold time
5	Neighbor uptime
6	Smooth round-trip time
7	Retransmission timeout
8	Number of packets to send to neighbor
9	Last sequence received

ROUTE v1.0—2-4

This output of the **show ip eigrp neighbors** command includes the following key elements:

1. **H (handle):** This column lists the order in which a peering session was established with the specified neighbor. The order is specified with sequential numbering starting with 0.

2. **Address:** This column contains the IP address of the EIGRP peer.

3. **Interface:** This column contains the interface on which the router is receiving hello packets from the peer.

4. **Hold Time:** This column contains the length of time (in seconds) that the Cisco IOS software will wait to hear from the peer before declaring it down. If the peer is using the default hold time, this number will be less than 15. If the peer is configured with a non-default hold time, the non-default hold time will be displayed. Originally, the expected packet was a hello packet, but with current Cisco IOS software releases, any EIGRP packet received after the first hello from that neighbor resets the timer.

5. **Uptime:** This is the elapsed time (in hours: minutes: seconds) since the local router first heard from this neighbor.

6. **Smooth Round Trip Timer (SRTT):** The smooth round-trip time is the number of milliseconds required for an EIGRP packet to be sent to this neighbor and for the local router to receive an acknowledgment of that packet. This timer is used to determine the retransmit interval, also known as the retransmit timeout (RTO).

7. **Retransmission timeout (RTO)**: This is the amount of time (in milliseconds) the software waits before resending a packet from the retransmission queue to a neighbor.

8. **Queue count:** This is the number of EIGRP packets (update, query, and reply) that the software is waiting to send. If this value is consistently higher than 0, a congestion problem might exist. A 0 indicates that no EIGRP packets are in the queue.

9. **Seq Num:** This is the sequence number of the last update, query, or reply packet that was received from this neighbor.

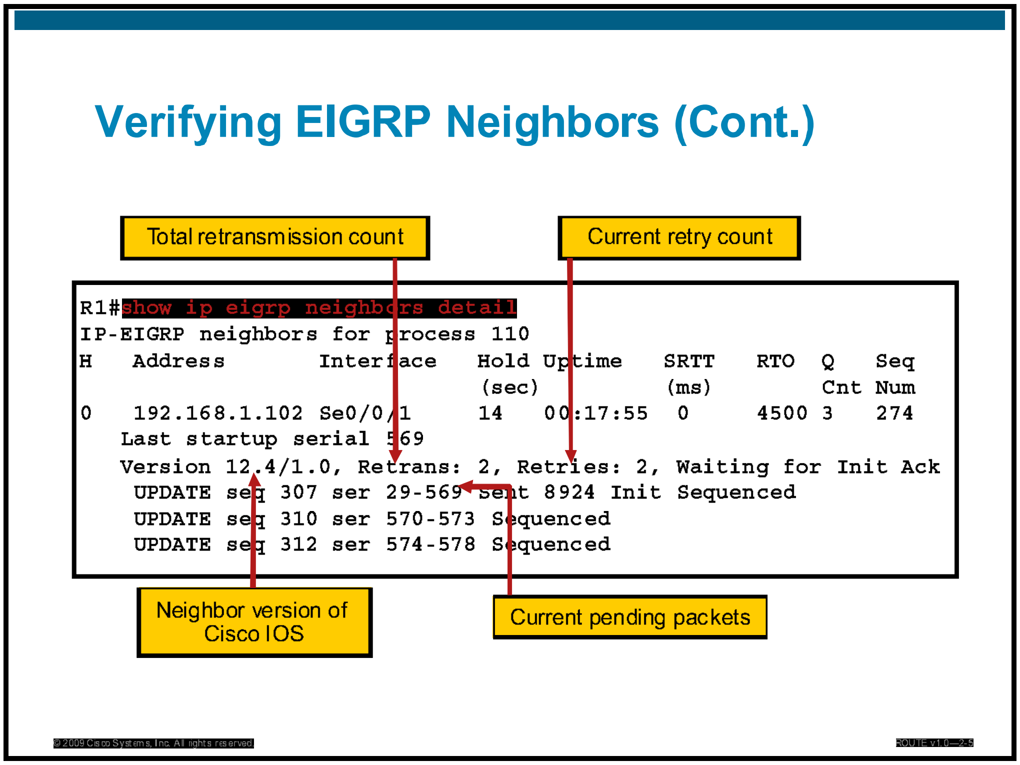

Verifying EIGRP Neighbors (Cont.)

Total retransmission count

Current retry count

```
R1#show ip eigrp neighbors detail
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold Uptime    SRTT    RTO  Q    Seq
                                   (sec)          (ms)         Cnt  Num
0   192.168.1.102  Se0/0/1        14    00:17:55  0       4500 3    274
   Last startup serial 569
   Version 12.4/1.0, Retrans: 2, Retries: 2, Waiting for Init Ack
    UPDATE seq 307 ser 29-569 Sent 8924 Init Sequenced
    UPDATE seq 310 ser 570-573 Sequenced
    UPDATE seq 312 ser 574-578 Sequenced
```

Neighbor version of Cisco IOS

Current pending packets

ROUTE v1.0—2-5

Detailed neighbor information can be examined with the **show ip eigrp neighbor detail** command. The command also reveals how many times a retransmission has occurred, the current retry count (router R1 in the figure has the value 2 for the retry count), the packets that are currently waiting to be sent (you can see that router R1 has three updates waiting to be sent), and the Cisco IOS version on the neighboring router

For more details about the **show ip eigrp neighbors detail** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

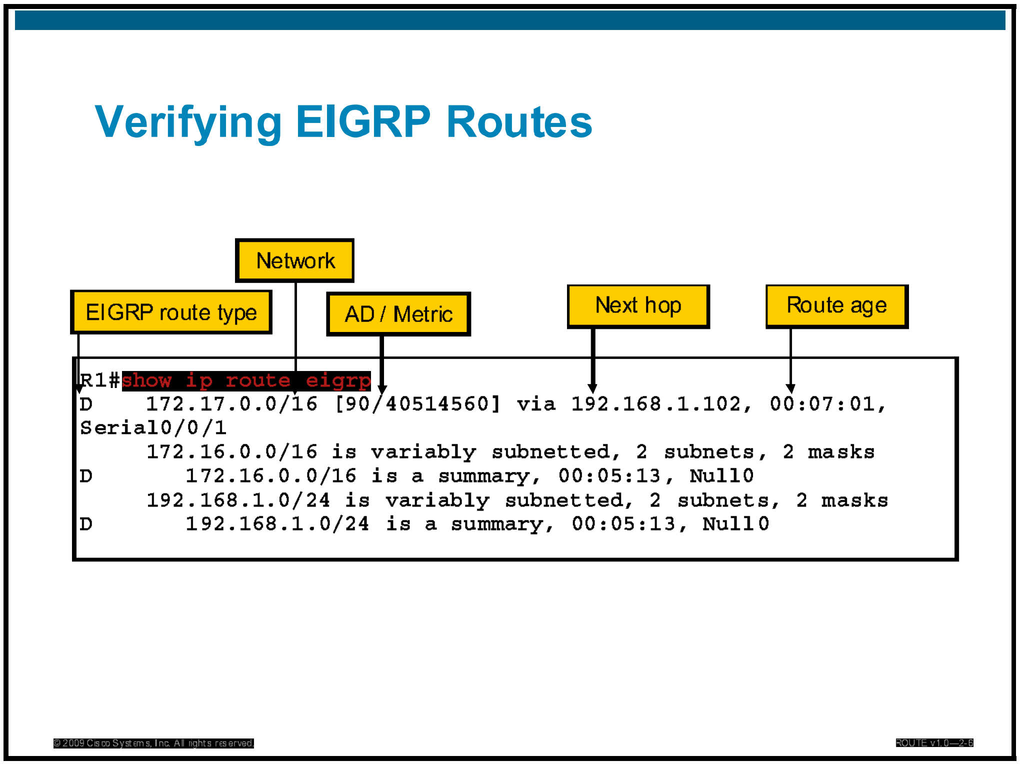

Verifying EIGRP Routes

```
R1#show ip route eigrp
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:07:01,
Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:05:13, Null0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
D       192.168.1.0/24 is a summary, 00:05:13, Null0
```

ROUTE v1.0—2-6

If you run the **show ip route** command, the output will contain all of the routes in the routing table. To verify only EIGRP routes for any neighbors the router recognizes, use the **show ip route eigrp** command.

EIGRP supports several route types: EIGRP routes from the local AS (D), EIGRP routes from the external AS (EX), and summary routes. EIGRP routes in the router's R1 routing table above are identified with a D in the left column; any external EIGRP routes (from outside of this autonomous system) would be identified with an EX.

After the network number, there is a field that looks similar to [90/40514560]. The numbers may be different from the one in the example. The first number, 90 in the example above, is the administrative distance. It is used to select the best path when a router learns two or more routes from different routing sources. For example, consider that this router also uses Routing Information Protocol (RIP), and RIP has a route to network 172.17.0.0 that is three hops away. The router, without administrative distance, cannot compare the three hops for RIP to an EIGRP metric of 40514560. The router does not know the bandwidth associated with the hops and EIGRP does not use a hop count as a metric. To avoid problems like this, Cisco established an administrative distance value for each routing protocol; the lower the value, the more preferred the route is. By default, EIGRP internal routes have an administrative distance of 90 and RIP has an administrative distance of 120. Because EIGRP has a metric based upon bandwidth and delays, it is preferred over the RIP hop count. As a result, in this example, the EIGRP route is installed in the routing table. The second number in the brackets is the EIGRP metric. Recall that the default EIGRP metric is the least-cost bandwidth plus the accumulated delays. The EIGRP metric for a certain network is the same as its feasible distance (FD) in the EIGRP topology table.

The next field, "via 192.168.1.102" in the example above, identifies the address of the next-hop router to which this router passes the packets for the destination network 172.17.0.0/16. The next-hop address in the routing table is the same as the successor in the EIGRP topology table.

Each route also has a time associated with it: the length of time, perhaps days or months, since EIGRP last advertised this network to this router. EIGRP does not refresh routes periodically; it resends the routing information only when neighbor adjacencies change.

The next field in the output is the interface (serial 0/0/1 in this case) from which packets for 172.17.0.0 are sent.

Notice that the routing table includes routes to null0 for the advertised routes. The Cisco IOS software automatically inserts these routes in the table. They are called summary routes. Null0 is a directly connected, software-only interface. The use of the null0 interface prevents the router from trying to forward traffic to other routers in search of a more precise, longer match. For example, if the router in the figure receives a packet to an unknown subnet that is part of the summarized range (such as 172.16.3.5) 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, or sent to the bit bucket) which prevents the router from forwarding the packet to a default route and possibly creating a routing loop.

For more details about the **show ip route** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

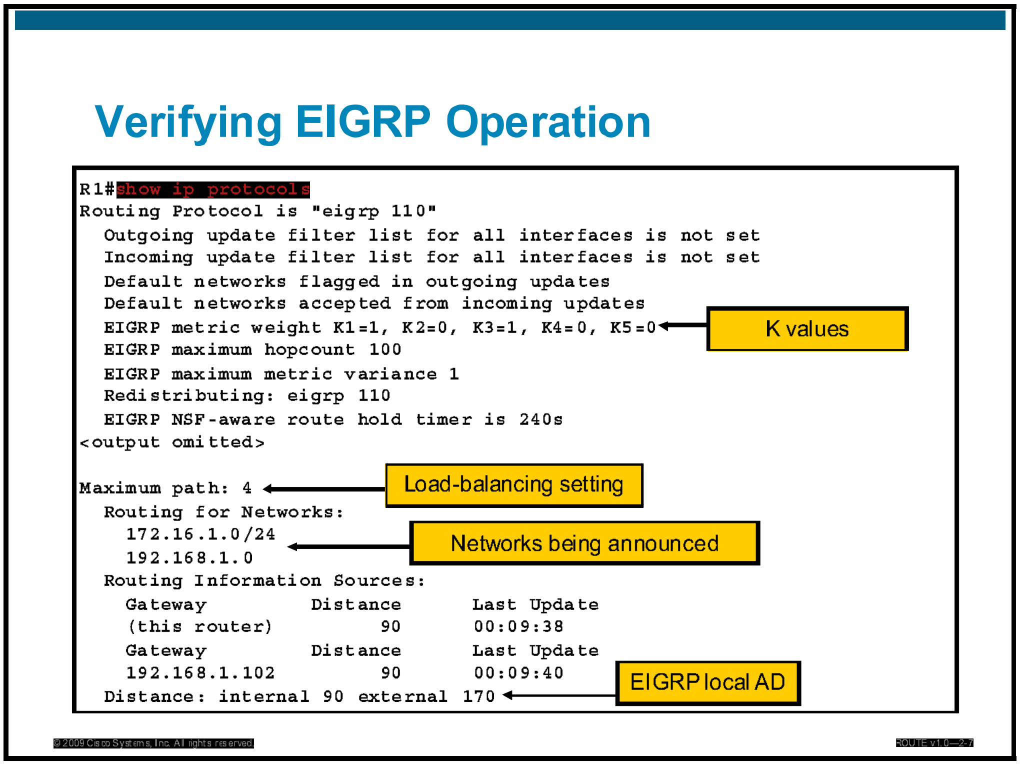

Verifying EIGRP Operation

```
R1#show ip protocols
Routing Protocol is "eigrp 110"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Default networks flagged in outgoing updates
  Default networks accepted from incoming updates
  EIGRP metric weight K1=1, K2=0, K3=1, K4=0, K5=0
  EIGRP maximum hopcount 100
  EIGRP maximum metric variance 1
  Redistributing: eigrp 110
  EIGRP NSF-aware route hold timer is 240s
<output omitted>

Maximum path: 4
  Routing for Networks:
    172.16.1.0/24
    192.168.1.0
  Routing Information Sources:
    Gateway         Distance      Last Update
    (this router)         90      00:09:38
    Gateway         Distance      Last Update
    192.168.1.102         90      00:09:40
  Distance: internal 90 external 170
```

K values

Load-balancing setting

Networks being announced

EIGRP local AD

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-7

To display the parameters and current state of the active routing protocol process, use the **show ip protocols** command.

The sample output in the example above shows that EIGRP process 110 is running. The command output displays any filtering of routing that is occurring on EIGRP outbound or inbound updates. It also identifies if EIGRP is generating a default network or receiving a default network in EIGRP updates.

The command output provides information about additional default settings for EIGRP, such as default K values, hop count, and variance.

Note Because the routers must have identical K values for EIGRP to establish an adjacency, you should run the **show ip protocols** command to determine the current K value setting before attempting an adjacency.

This sample output also indicates that automatic summarization is enabled (this is the default setting) and that the router is allowed to load-balance over a maximum of four paths. (The Cisco IOS software allows configuration of up to six paths for equal-cost load balancing, using the **maximum-path** configuration command.)

The networks that the router is routing are also displayed. As shown in the figure, the format of the output varies, depending on the use of the wildcard mask in the **network** command. If a wildcard mask is used, the network address is displayed with a prefix length. If a wildcard mask is not used, the Class A, B, or C major network is displayed.

The routing information sources section of this command output identifies all other routers that have an EIGRP neighbor relationship with this router.

The **show ip protocols** command output also provides the two administrative distances. First, an administrative distance of 90 applies to networks from other routers inside the autonomous system; these are considered internal networks. Second, an administrative distance of 170 applies to networks introduced to EIGRP for this autonomous system through redistribution; these are called external networks. The source of the external routes is not inside the autonomous system.

For more details about the **show ip protocols** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

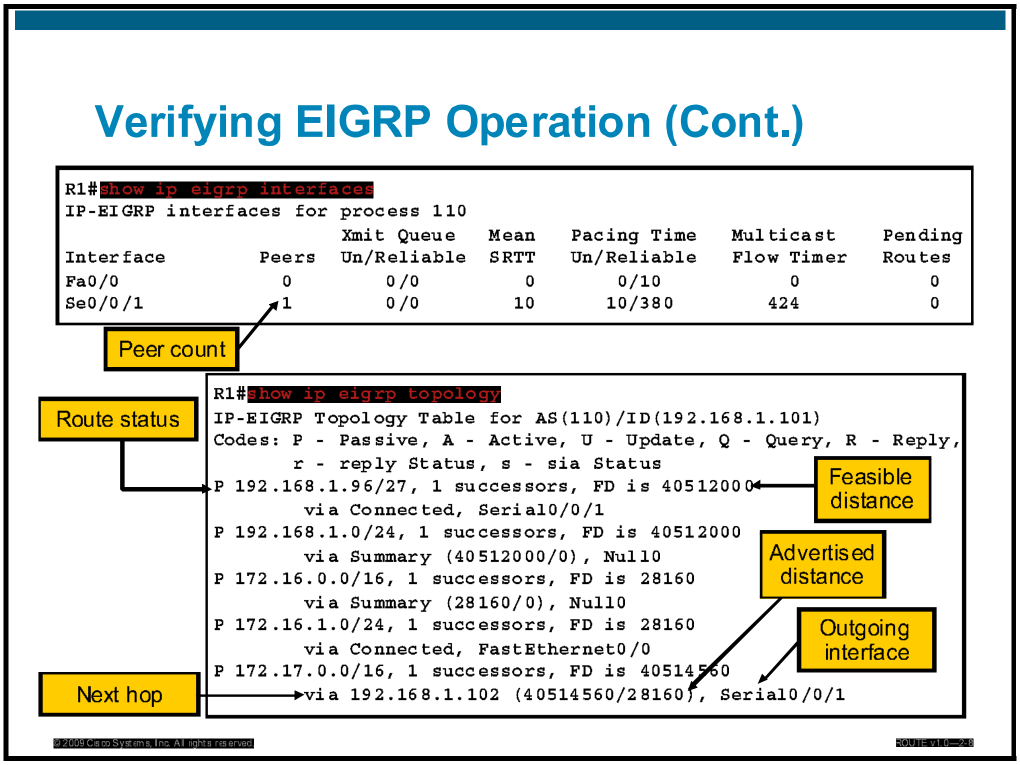

Verifying EIGRP Operation (Cont.)

```
R1#show ip eigrp interfaces
IP-EIGRP interfaces for process 110
                        Xmit Queue   Mean   Pacing Time   Multicast    Pending
Interface        Peers  Un/Reliable  SRTT   Un/Reliable   Flow Timer   Routes
Fa0/0              0        0/0        0       0/10           0           0
Se0/0/1            1        0/0       10      10/380        424           0
```

```
R1#show ip eigrp topology
IP-EIGRP Topology Table for AS(110)/ID(192.168.1.101)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 192.168.1.96/27, 1 successors, FD is 40512000
        via Connected, Serial0/0/1
P 192.168.1.0/24, 1 successors, FD is 40512000
        via Summary (40512000/0), Null0
P 172.16.0.0/16, 1 successors, FD is 28160
        via Summary (28160/0), Null0
P 172.16.1.0/24, 1 successors, FD is 28160
        via Connected, FastEthernet0/0
P 172.17.0.0/16, 1 successors, FD is 40514560
        via 192.168.1.102 (40514560/28160), Serial0/0/1
```

The **show ip eigrp interfaces** command displays information about interfaces configured for EIGRP. Use this command to determine which interfaces EIGRP is active on, and to learn information about EIGRP for interfaces. As shown in this sample output, the following key elements are included in the output:

- **Interface:** Interface over which EIGRP is configured
- **Peers:** Number of directly connected EIGRP neighbors
- **Xmit Queue Un/Reliable:** Number of packets remaining in the Unreliable and Reliable transmit queues
- **Mean SRTT:** Mean smooth round-trip time (SRTT) interval (in seconds).
- **Pacing Time Un/Reliable:** Pacing time used to determine when EIGRP packets should be sent out of the interface (unreliable and reliable packets)
- **Multicast Flow Timer:** Maximum number of seconds in which the router will send multicast EIGRP packets
- **Pending Routes:** Number of routes in the packets sitting in the transmit queue waiting to be sent

To verify EIGRP operations further, you can use the **show ip eigrp topology** command. Use this command to determine the Diffusing Update Algorithm (DUAL) states and to debug any possible DUAL problems. If this command is used without any keywords or arguments, only routes that are feasible successors are displayed. The sample output above shows that Router R1 has an ID of 192.168.1.101 and resides in the autonomous system 110 (the EIGRP ID is the highest IP address on an active interface for this router). The command output lists the networks known by this router through the EIGRP routing process. The codes in the command output showing the state of this topology table entry are defined as follows:

- **Passive (P):** This network is available and installation can occur in the routing table. Passive is the correct state for a stable network. Passive state is an indication that no EIGRP computations are being performed for this destination.
- **Active (A):** This network is currently unavailable and installation cannot occur in the routing table. A network in an active state has outstanding queries. Active state is an indication that EIGRP computations are being performed for this destination.
- **Update (U):** This code applies if a network is being updated (an update packet is being sent to this destination). This code also applies if the router is waiting for an acknowledgment for an update packet.
- **Query (Q):** This code applies if there is an outstanding query packet for this network and the network is not in the active state. The code indicates that a query packet was sent to this destination. This code also applies if the router is waiting for an acknowledgment for a query packet.
- **Reply (R):** This code applies if the router is generating a reply for this network or is waiting for an acknowledgment for a reply packet. This code indicates that a reply packet was sent to this destination.
- **Stuck-in-active (SIA) status:** This code signifies an EIGRP convergence problem for the network with which the route is associated.

The number of successors available for a route is indicated in the command output, as well. In this example, all networks have one successor. If there were equal-cost paths to the same network, a maximum of six paths would be shown. The number of successors corresponds to the number of best routes with equal cost.

For each network, the FD is displayed, followed by the next-hop address and then a field similar to "(40514560/28160)" in the figure. The first number in this field is the FD for that network through this next-hop router. The second number is the advertised distance (AD) from the next-hop router to the destination network.

For more details about the **show ip eigrp interfaces** and **show ip eigrp topology** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

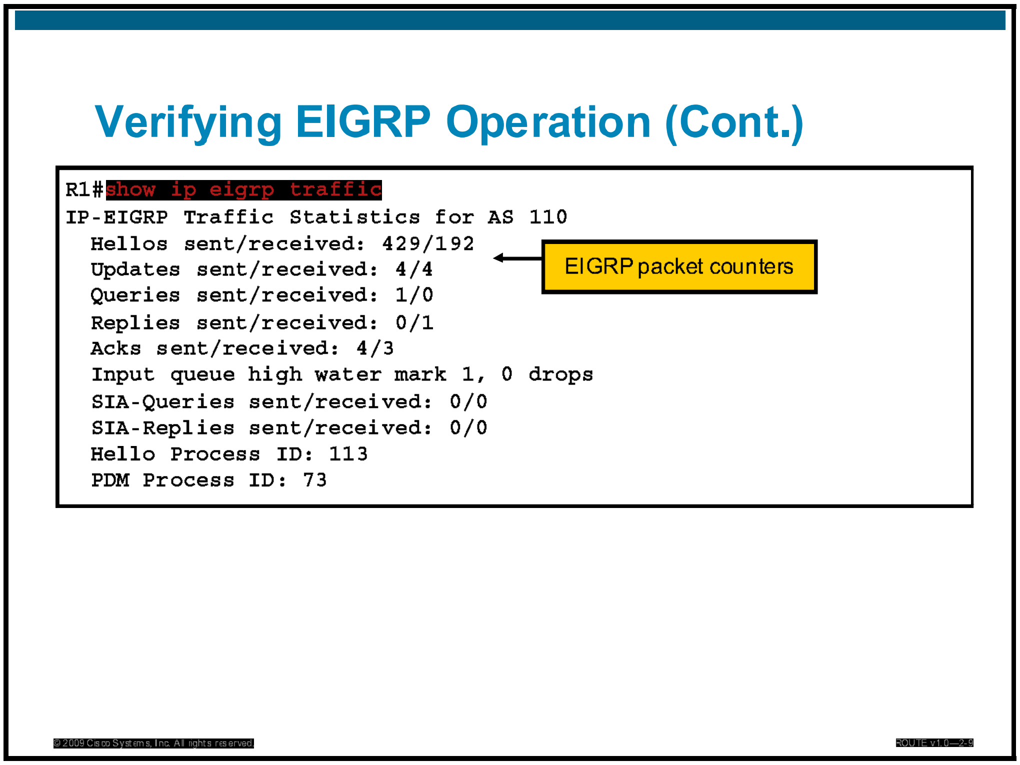

Verifying EIGRP Operation (Cont.)

```
R1#show ip eigrp traffic
IP-EIGRP Traffic Statistics for AS 110
  Hellos sent/received: 429/192
  Updates sent/received: 4/4
  Queries sent/received: 1/0
  Replies sent/received: 0/1
  Acks sent/received: 4/3
  Input queue high water mark 1, 0 drops
  SIA-Queries sent/received: 0/0
  SIA-Replies sent/received: 0/0
  Hello Process ID: 113
  PDM Process ID: 73
```

EIGRP packet counters

ROUTE v1.0—2-9

To examine the number of various EIGRP packets sent and received, use the **show ip eigrp traffic** command, as illustrated in the figure.

You can see that router R1 has sent 429 and received 192 hello messages, sent 4 and received 4 update messages, sent 1 query message and received 0 query messages, sent 0 reply messages and received 1 reply message, and sent 4 ACK messages and received 3 ACK messages.

For more details about the **show ip eigrp traffic** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Using the Passive-Interface Command with EIGRP

This topic describes how to control routing updates using the **passive-interface** command.

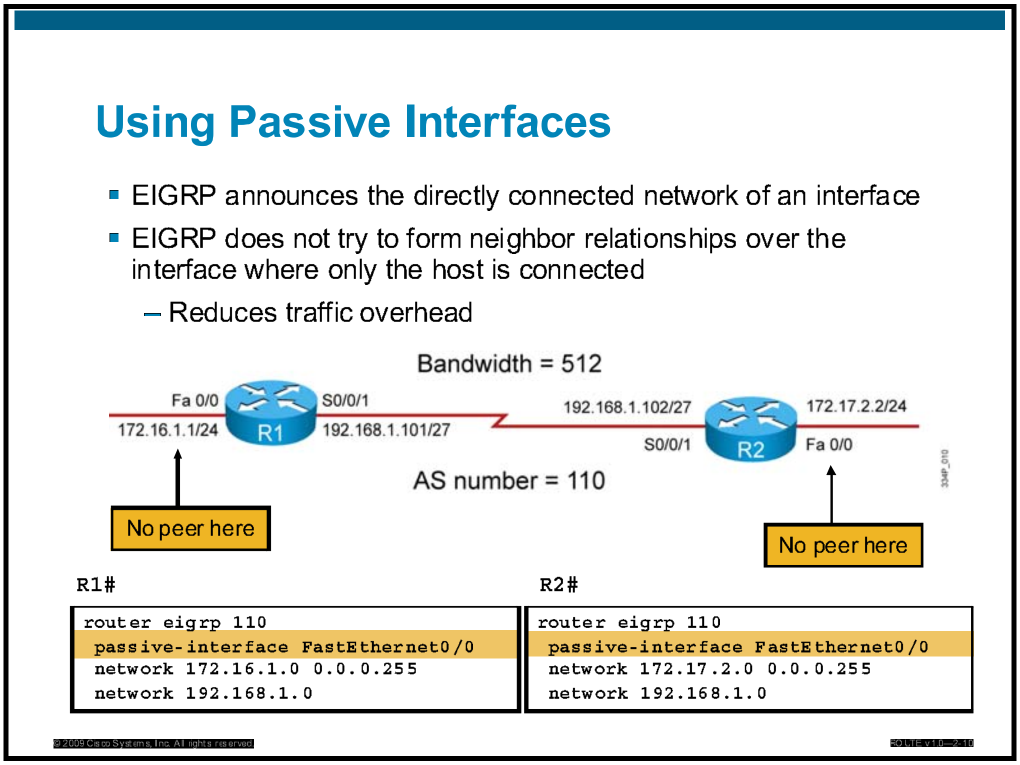

Using Passive Interfaces

- EIGRP announces the directly connected network of an interface
- EIGRP does not try to form neighbor relationships over the interface where only the host is connected
  - Reduces traffic overhead

R1#

```
router eigrp 110
 passive-interface FastEthernet0/0
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

R2#

```
router eigrp 110
 passive-interface FastEthernet0/0
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
```

ROUTE v1.0—2-10

Routers R1 and R2 have no neighbors available over the FastEtehrent0/0 interfaces, thus there is no need to try to establish adjacency over the interfaces. Moreover, the packets sent are overhead to the link bandwidth and also consume CPU resources of the router. In order to stop sending hello packets over the interface without neighbors, use the **passive-interface** command on the specified interface. In the above example, the **passive-interface** command is used in both routers for the Fast Ethernet interface 0/0. EIGRP will not bring up adjacencies on a passive interface, regardless of whether the **neighbor** command is configured.

Note Configuring the **passive-interface** command suppresses all incoming and outgoing routing updates and hello messages.

The **passive-interface** command has the following properties:

- Prevents a neighbor relationship from being established over the passive interface
- Stops routing updates from being received or sent over the passive interface
- Allows a subnet on the passive interface to be announced in an EIGRP process

For more details about the **passive-interface** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

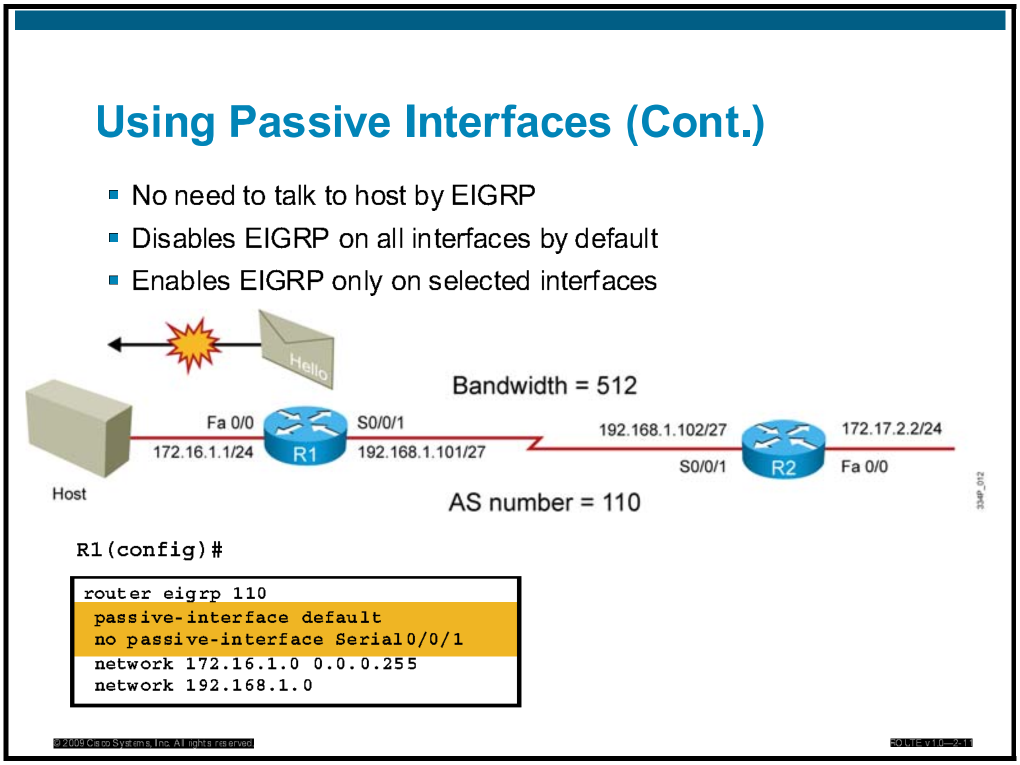

Using Passive Interfaces (Cont.)

- No need to talk to host by EIGRP
- Disables EIGRP on all interfaces by default
- Enables EIGRP only on selected interfaces

```
R1(config)#
```

```
router eigrp 110
 passive-interface default
 no passive-interface Serial0/0/1
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

ROUTE v1.0—2-11

Within ISPs and large enterprise networks, distribution routers may have more than 100 interfaces, so manual configuration of the **passive-interface** command on interfaces where adjacency is not desired may create a problem. In some networks, this means entering 100 or more passive interface statements.

With the default passive interface feature, this issue is solved by allowing all interfaces to be set as passive by default using a single **passive-interface default** command. Where adjacencies are desired, the individual interfaces are configured using the **no passive-interface** command.

In the example above, routers R1 and R2 are configured with the **passive-interface default** command and all interfaces are refusing the establishment of EIGRP adjacency by default. The Serial 0/0/1 interface on each router is then configured to allow EIGRP adjacency as neighbors are expected. The **passive-interface** command is disabled for these interfaces..

For more details about the **passive-interface default** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

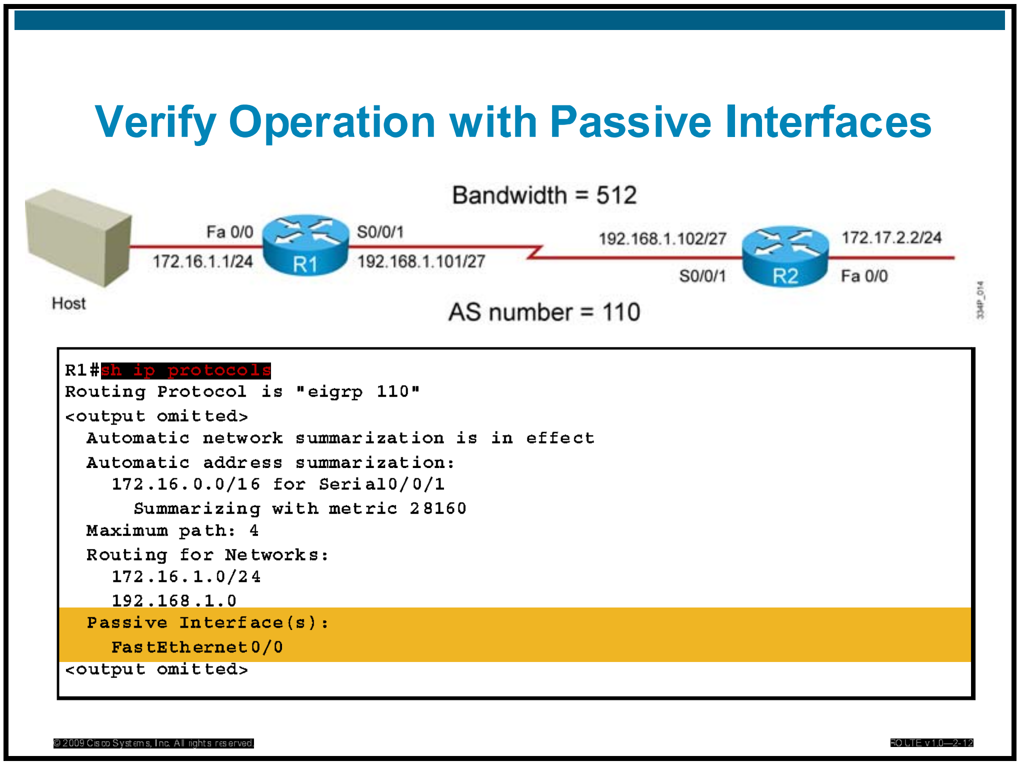

Verify Operation with Passive Interfaces

```
R1#sh ip protocols
Routing Protocol is "eigrp 110"
<output omitted>
  Automatic network summarization is in effect
  Automatic address summarization:
    172.16.0.0/16 for Serial0/0/1
      Summarizing with metric 28160
  Maximum path: 4
  Routing for Networks:
    172.16.1.0/24
    192.168.1.0
  Passive Interface(s):
    FastEthernet0/0
<output omitted>
```

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-12

The main questions to ask when verifying operation with passive interfaces are:

- Do we see all the neighbors?
- Which interfaces in the routing process are passive?

To see all of the EIGRP neighbors available, use the **show ip eigrp neighbors** command.

To see the passive interfaces in the routing protocol, use **show ip protocols** command. The command output above for router R1 shows that interface FastEthernet 0/0 is defined as a passive interface.

For more details about the **show ip protocols** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

Advertising an IP Default Network in EIGRP

This topic describes the **ip default-network** command that is used to configure the last resort gateway or default route.

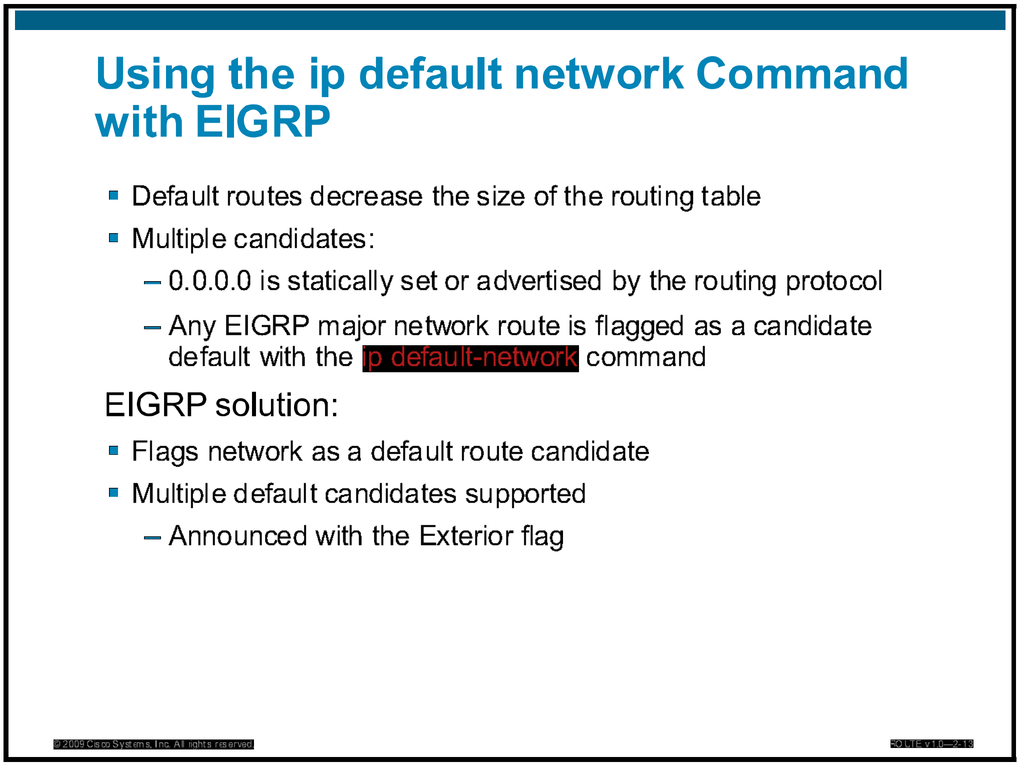

Using the ip default network Command with EIGRP

- Default routes decrease the size of the routing table
- Multiple candidates:
  - 0.0.0.0 is statically set or advertised by the routing protocol
  - Any EIGRP major network route is flagged as a candidate default with the ip default-network command

EIGRP solution:

- Flags network as a default route candidate
- Multiple default candidates supported
  - Announced with the Exterior flag

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-13

The major purpose of using default routes is to decrease the size of the routing table. This especially applies to stub networks or networks at the access layer (generally, it applies to all networks that are on lower hierarchical layers).

The router, before installation of a default route, first collects default route candidates.

- The candidate can be a statically configured default route with the following command: **ip route** *0.0.0.0 0.0.0.0 next-hcp* | *inteıface*. In this command, *inteıface* is an outgoing interface through which all packets with unknown destinations will be forwarded, and *next-hcp* is an IP address to which packets with unknown destinations will be forwarded.
- Any major network residing in the local routing table can become a candidate to use the **ip default-network** command. The command is also used to attach an Exterior flag to any major EIGRP or IGRP route, thus making it a candidate for a default route.

Note In EIGRP, no default routes can be directly injected (as in the OSPF environment with the default-information originate command).

The router examines all the default candidates and selects the best one based on the administrative distance and route metric.

When selected, the router sets the gateway of last resort to the next hop of the selected candidate. This does not apply when the best candidate happens to be one of the directly connected routes.

Note Any route residing in a routing table can be marked as a default candidate even if the route is not an EIGRP route.

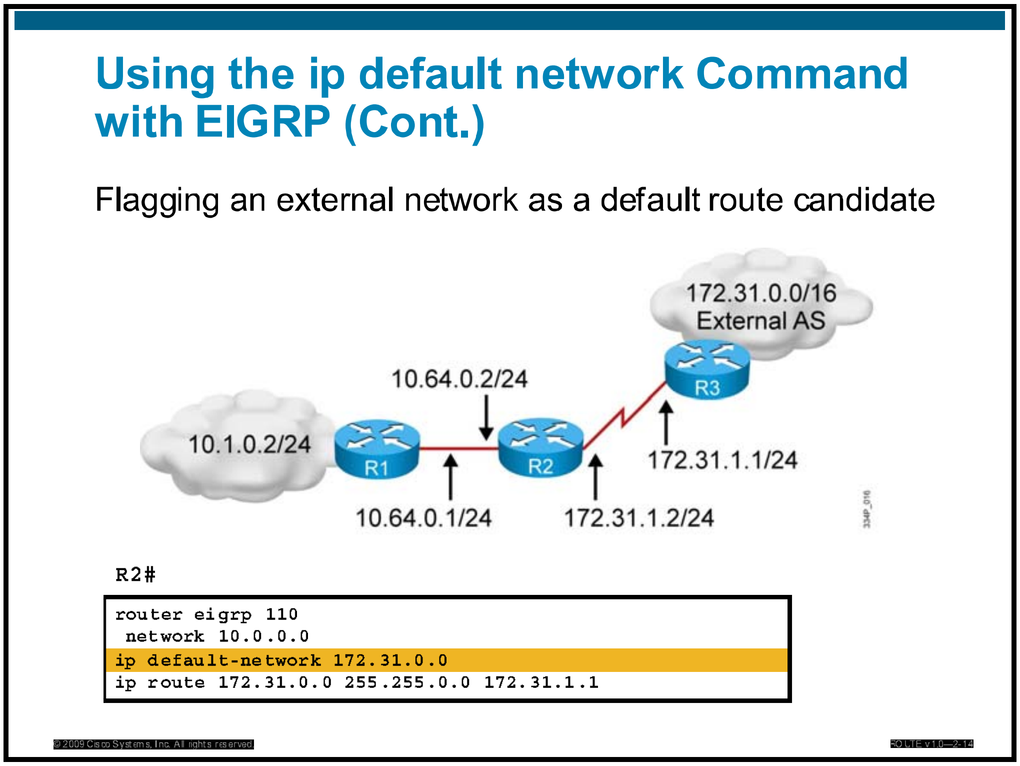

Using the ip default network Command with EIGRP (Cont.)

Flagging an external network as a default route candidate

R2#

```
router eigrp 110
 network 10.0.0.0
ip default-network 172.31.0.0
ip route 172.31.0.0 255.255.0.0 172.31.1.1
```

ROUTE v1.0—2-14

The EIGRP default route can be created with the **ip default-network** command. A router configured with this command considers the network listed in the command as the last-resort gateway that it will announce to other routers.

The network specified by this command must be reachable by the router that uses this command before it announces it as a candidate default route to other EIGRP routers. The network specified by this command must also be passed to other EIGRP routers so that those routers can use this network as their default network and set the gateway of last resort to this default network. This means that the network must either be an EIGRP-derived network in the routing table or be generated using a static route, which has been redistributed into EIGRP.

Note Multiple default networks can be configured; downstream routers use the EIGRP metric to determine the best default route.

Router R2 is has access to the external network 172.31.0.0/16 via its serial interface. The static route is configured in order to provide reachability, as routers R2 and R3 are not exchanging routing updates. Router R2 is configured with the 172.31.0.0 network as a candidate default network using the **ip default-network 172.31.0.0** command. This network is passed to router R1.

For more details about the **ip default-network** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

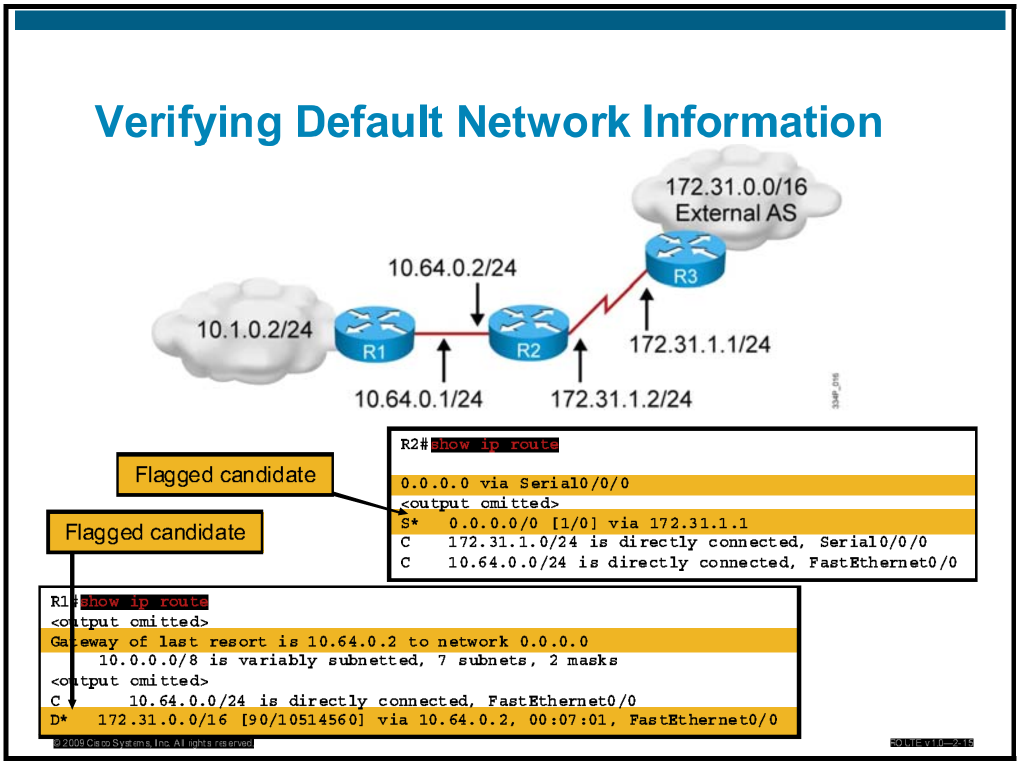

Verifying Default Network Information

```
R2#show ip route

0.0.0.0 via Serial0/0/0
<output omitted>
S*    0.0.0.0/0 [1/0] via 172.31.1.1
C     172.31.1.0/24 is directly connected, Serial0/0/0
C     10.64.0.0/24 is directly connected, FastEthernet0/0
```

```
R1#show ip route
<output omitted>
Gateway of last resort is 10.64.0.2 to network 0.0.0.0
     10.0.0.0/8 is variably subnetted, 7 subnets, 2 masks
<output omitted>
C       10.64.0.0/24 is directly connected, FastEthernet0/0
D*   172.31.0.0/16 [90/10514560] via 10.64.0.2, 00:07:01, FastEthernet0/0
```

ROUTE v1.0—2-15

Routers R1 and R2 are configured for EIGRP and router R2 has a configuration using the **ip default-network** command. In order to verify the processing of the default routes and default candidates, you must look at the IP routing tables on both of the routers.

The 172.31.0.0 network is passed from router R2 to router R1. The **ip default-network** command does not benefit router R2 directly. On router R1, the EIGRP-learned 172.31.0.0 network is flagged as a candidate default network (as indicated by the * in the routing table). Router R1 also sets the gateway of last resort to 10.64.0.2 (toward router R2) to reach the default network of 172.31.0.0.

Note When you configure the **ip default-network** command, a static route (the **ip route** command) is generated in the router's configuration; however, the Cisco IOS software does not display a message to indicate that this has been done. The entry appears as a static route in the routing table of the router in which the command is configured, as can be seen in router R1's configuration and routing table in the figure. This can be confusing if you want to remove the default network. The configuration must be removed with the **no ip route** command.

Determining Summary Boundaries

This topic explains why administrators may need to use manual route summarization over default automatic route summarization.

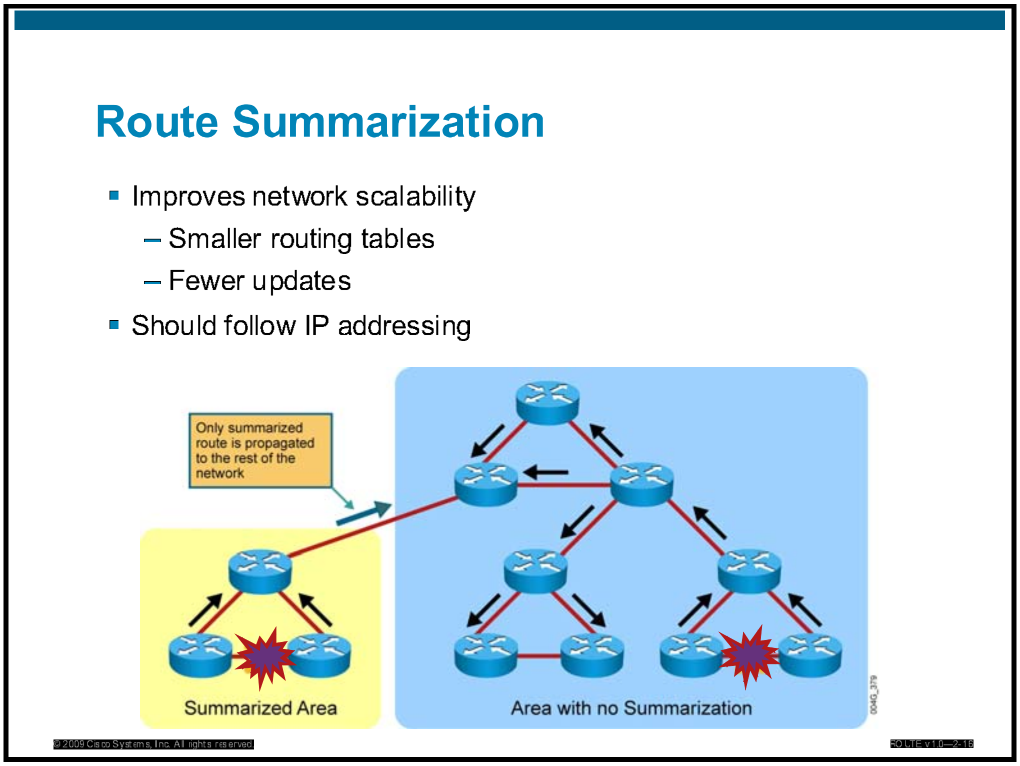

Route Summarization

- Improves network scalability
  - Smaller routing tables
  - Fewer updates
- Should follow IP addressing

Only summarized route is propagated to the rest of the network
Summarized Area
Area with no Summarization
004G_379

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-16

To reduce routing overhead and improve stability and scalability of routing, you can use route aggregation (summarization). However, to implement route aggregation, you must divide the network into contiguous IP address areas. This requires you to have a solid understanding of IP address assignment on route aggregation and hierarchical routing.

The purpose of route summarization is to squeeze several subnets into one aggregate entry that covers all of them. Summarization results in smaller routing tables and smaller updates. Consequently, it also results in less routing traffic and lower CPU utilization (minor changes in the network go unnoticed).

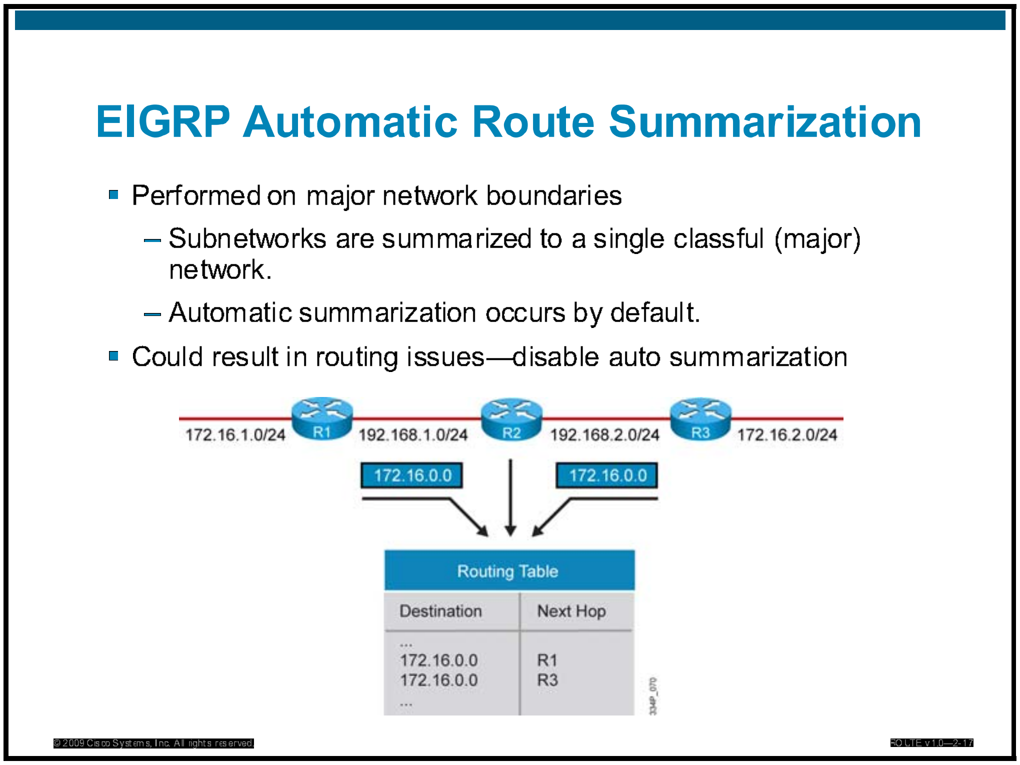

EIGRP Automatic Route Summarization

- Performed on major network boundaries
  - Subnetworks are summarized to a single classful (major) network.
  - Automatic summarization occurs by default.
- Could result in routing issues—disable auto summarization

Routing Table

Destination	Next Hop
...	
172.16.0.0	R1
172.16.0.0	R3
...	

ROUTE v1.0—2-17

Some EIGRP features (such as automatic route summarization routes at major network boundaries) are characteristics of distance vector operation. Traditional distance vector protocols, which are classful routing protocols, cannot assume the mask for networks that are not directly connected, because routing updates do not exchange masks.

EIGRP automatically summarizes routes at the classful boundary. In some cases, you may not want automatic summarization to occur. For example, if you have discontiguous networks, you need to disable automatic summarization to minimize router confusion.

Note Automatic summarization is enabled by default for EIGRP.

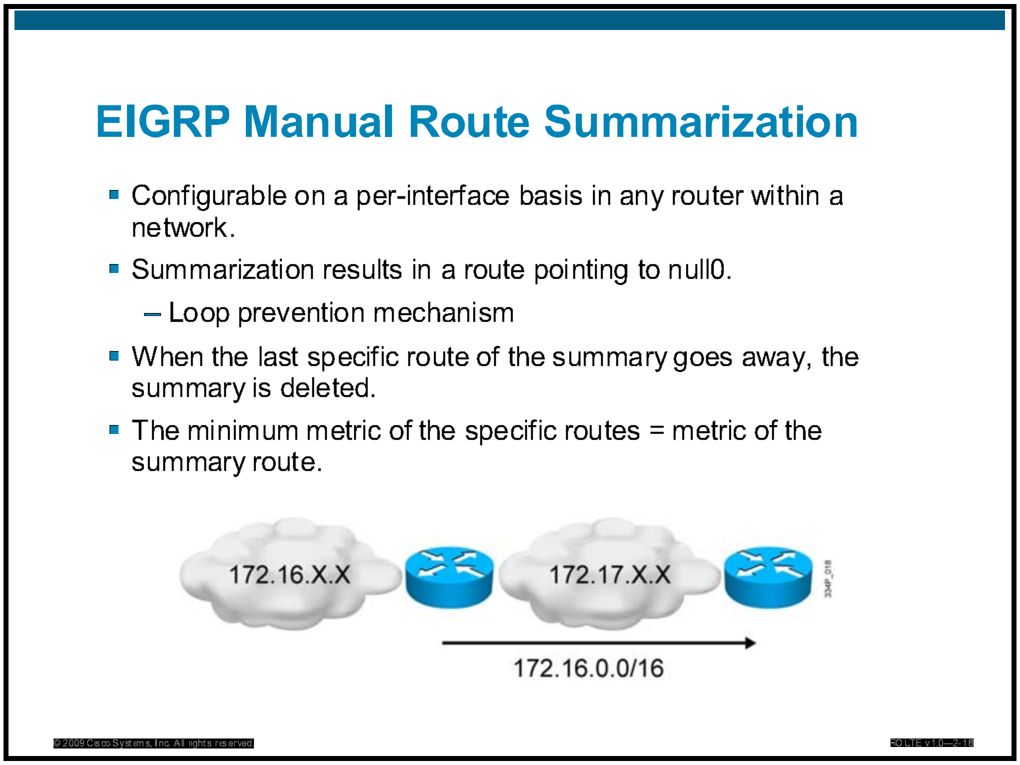

EIGRP Manual Route Summarization

- Configurable on a per-interface basis in any router within a network.
- Summarization results in a route pointing to null0.
  - Loop prevention mechanism
- When the last specific route of the summary goes away, the summary is deleted.
- The minimum metric of the specific routes = metric of the summary route.

A drawback to using distance vector protocols is that you cannot create summary routes at arbitrary boundaries within a major network. The ability to summarize routes is desirable, because it allows you to keep smaller routing tables. EIGRP allows administrators to disable automatic summarization and create one or more summary routes within the network on any bit boundary, as long as a more specific route exists in the routing table. When the last specific route of the summary goes away, the summary is deleted from the routing table.

The minimum metric of the specific routes is used as the metric of the summary route.

Recall that Cisco IOS software automatically inserts summary routes to interface null0 in the routing table for automatically summarized routes, to prevent routing loops. For the same reason, Cisco IOS software also creates a summary route to interface null0 when manual summarization is configured. 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.

For manual summarization to be effective, blocks of contiguous addresses (subnets) must come together at a common router so that the router can advertise a single summary route. The number of subnets that can be represented by a summary route is directly related to the difference in the number of bits between the subnet mask and the summary mask. The formula 2^n, where n equals the difference in the number of bits between the summary and subnet mask, indicates how many subnets can be represented by a single summary route. For example, if the summary mask contains three fewer bits than the subnet mask, eight ($2^3 = 8$) subnets can be aggregated into one advertisement.

For example, if network 10.0.0.0 is divided into /24 subnets and is summarized to the summarization block 10.1.8.0/21, the difference between the /24 networks and the /21 summarizations is 3 bits; therefore, $2^3 = 8$ subnets can be aggregated. The summarized subnets range from 10.1.8.0/24 through 10.1.15.0/24.

When configuring summary routes, the administrator needs to specify the IP address of the summary route and the summary mask. The Cisco IOS software for EIGRP handles many of the details that surround proper implementation, including details about metrics, loop prevention, and removal of the summary routes from the routing table if none of the more specific routes are valid.

Utilizing Manual Route Summarization

This topic explains why administrators may need to use manual route summarization over default automatic route summarization.

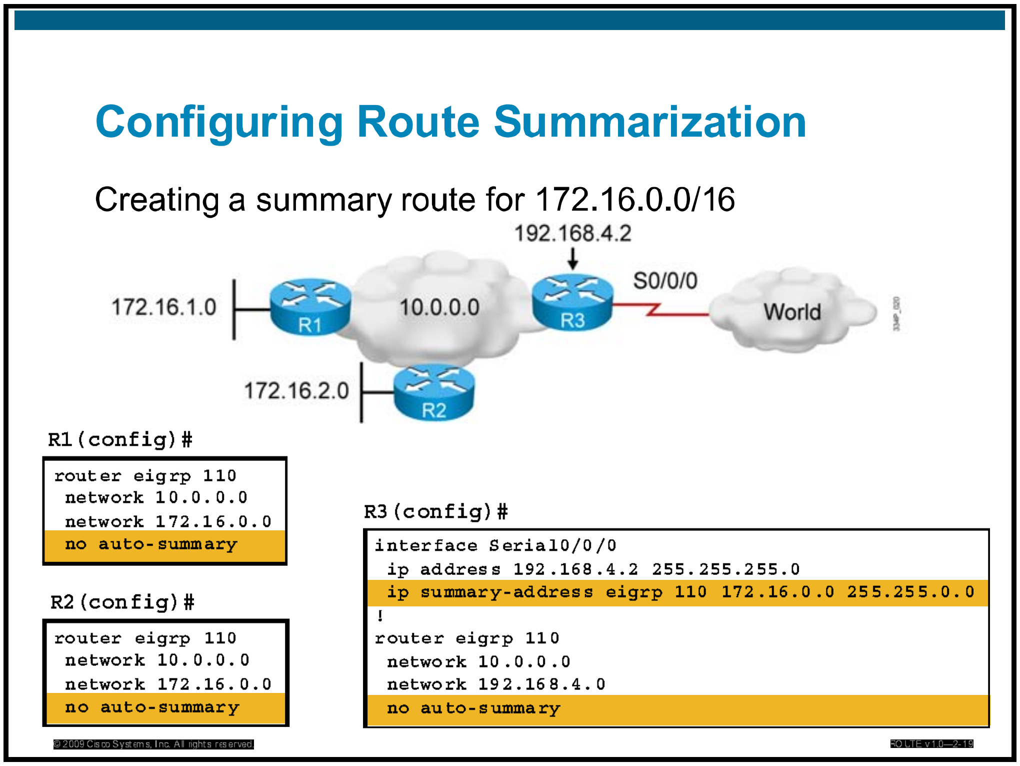

Configuring Route Summarization

Creating a summary route for 172.16.0.0/16

R1(config)#

```
router eigrp 110
 network 10.0.0.0
 network 172.16.0.0
 no auto-summary
```

R2(config)#

```
router eigrp 110
 network 10.0.0.0
 network 172.16.0.0
 no auto-summary
```

R3(config)#

```
interface Serial0/0/0
 ip address 192.168.4.2 255.255.255.0
 ip summary-address eigrp 110 172.16.0.0 255.255.0.0
!
router eigrp 110
 network 10.0.0.0
 network 192.168.4.0
 no auto-summary
```

EIGRP automatically summarizes routes at the classful boundary. In some cases, you may not want automatic summarization to occur. For example, if you have discontiguous networks, you need to disable automatic summarization to minimize router confusion. To disable automatic summarization, use the **no auto-summary** EIGRP router configuration command.

Note EIGRP router does not perform automatic summarization of networks in which it does not participate.

A discontiguous network 172.16.0.0 is used in the network on specific router R1 and R2 interfaces. On routers R1 and R2, automatic summarization has been disabled, so the 172.16.1.0 and 172.16.2.0 subnets are advertised into network 10.0.0.0. The routing tables of routers in the 10.0.0.0 network, including router R3, include these discontiguous subnets.

An EIGRP router automatically summarizes routes only for networks to which it is attached. If a network is not automatically summarized at the major network boundary, as is the case in this example on routers R1 and R2 because autosummarization is turned off, all the subnet routes are carried into the router R3 routing table. Router R3 will not automatically summarize the 172.16.1.0 and 172.16.2.0 subnets, because it does not own the 172.16.0.0 network. Therefore, router R3 will send routes to the 172.16.1.0 and 172.16.2.0 subnets to the WAN. If you configure a summary route on the router R3 interface serial 0/0/0, as shown in the figure, only one route will be sent on the WAN. This route will represent all subnets that belong to network 172.16.0.0.

To configure manual route summarization on router R3, you must select the interface to propagate the summary route. Interface Serial0/0/0 is used in the example above. When configuring the summary route, you should use the **ip summary-address eigrp** command. You must specify the EIGRP routing protocol, AS number, and the summary address and the mask of the routes.

Note For manual route summarization, the summary route is advertised only if a component (a more specific entry) of the summary route is present in the routing table.

For more details about the **auto-summary** and **ip summary-address eigrp** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

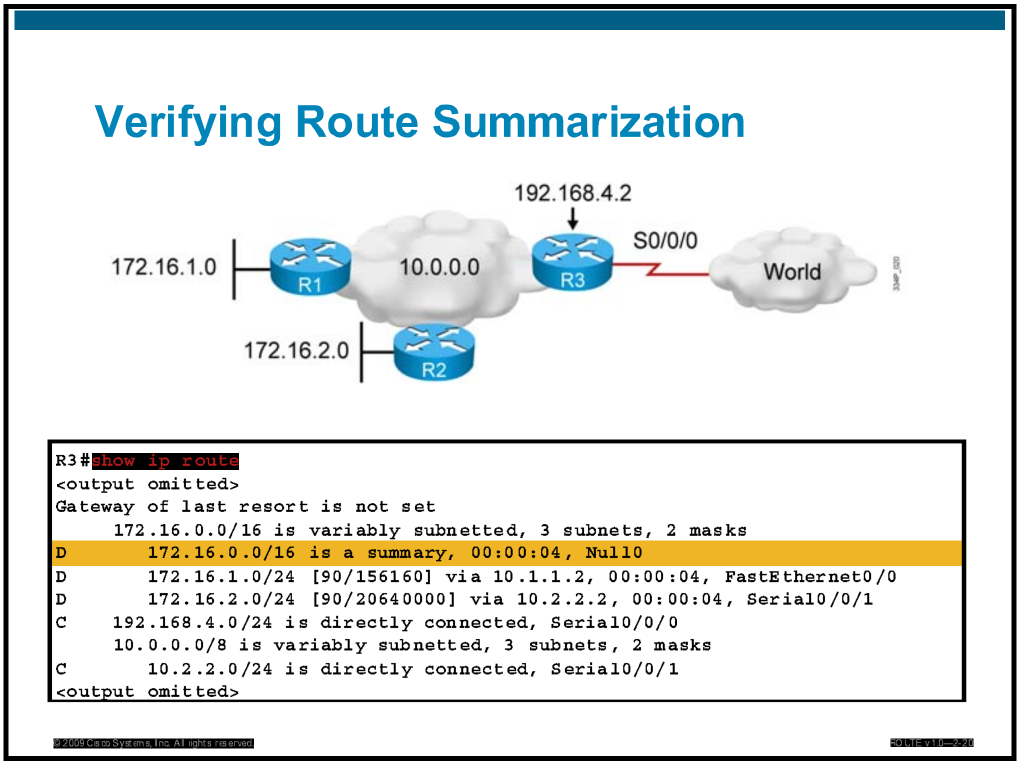

Verifying Route Summarization

```
R3#show ip route
<output omitted>
Gateway of last resort is not set
     172.16.0.0/16 is variably subnetted, 3 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:00:04, Null0
D       172.16.1.0/24 [90/156160] via 10.1.1.2, 00:00:04, FastEthernet0/0
D       172.16.2.0/24 [90/20640000] via 10.2.2.2, 00:00:04, Serial0/0/1
C    192.168.4.0/24 is directly connected, Serial0/0/0
     10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C       10.2.2.0/24 is directly connected, Serial0/0/1
<output omitted>
```

In order to verify that summarization is configured correctly, you must look at the IP routing table. Router R3 was configured for summarization and the routing table on router R3 is presented in the figure above. Router R3 has both 172.16.1.0 and 172.16.2.0, the discontiguous subnets, in its routing table. Because of the summarization, only network 172.16.0.0 is advertised out of the serial0/0/0 interface, though.

The summary route pointing to the Null0 interface prevents routing loops. This approach is based on the assumption that the router doing summarization has more information on the subnets covered by the summary route besides the summary route itself.

Summary

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

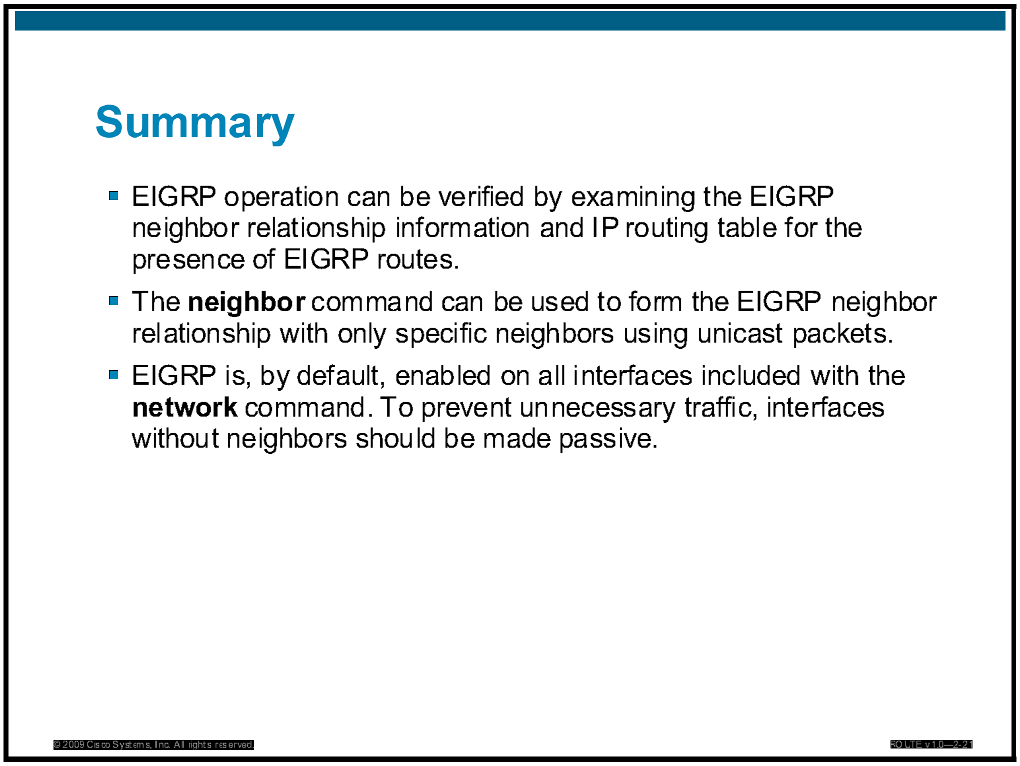

Summary

- EIGRP operation can be verified by examining the EIGRP neighbor relationship information and IP routing table for the presence of EIGRP routes.
- The **neighbor** command can be used to form the EIGRP neighbor relationship with only specific neighbors using unicast packets.
- EIGRP is, by default, enabled on all interfaces included with the **network** command. To prevent unnecessary traffic, interfaces without neighbors should be made passive.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-21

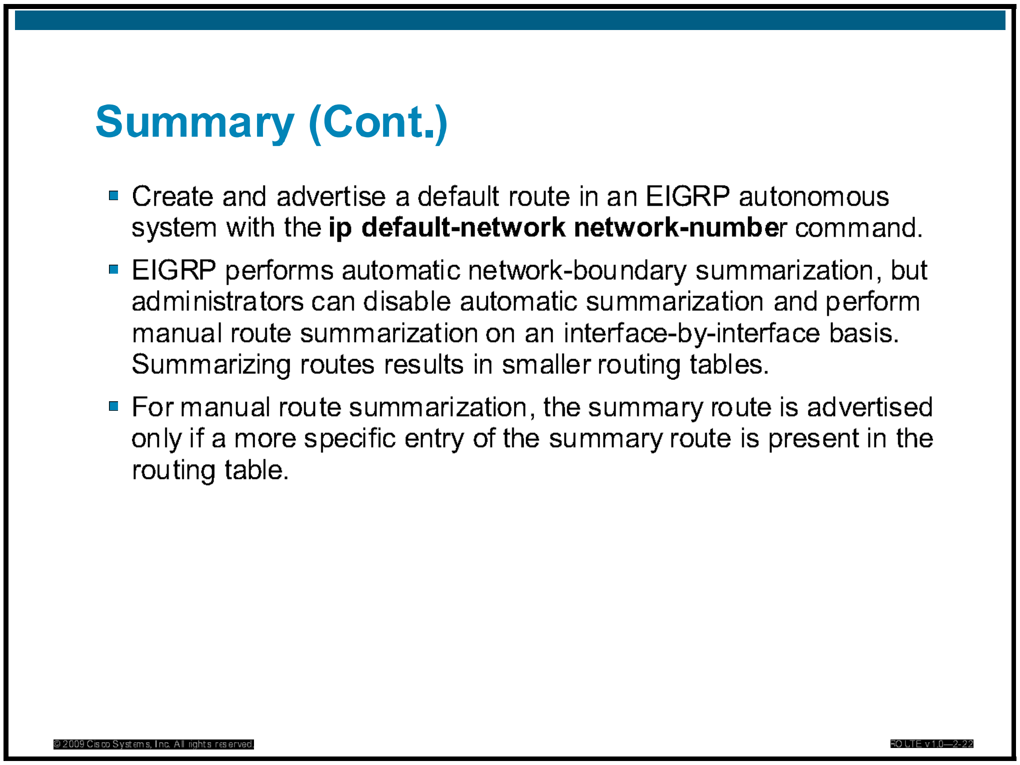

Summary (Cont.)

- Create and advertise a default route in an EIGRP autonomous system with the **ip default-network network-numbe**r command.
- EIGRP performs automatic network-boundary summarization, but administrators can disable automatic summarization and perform manual route summarization on an interface-by-interface basis. Summarizing routes results in smaller routing tables.
- For manual route summarization, the summary route is advertised only if a more specific entry of the summary route is present in the routing table.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-22

Lesson 3

Lab 2-1 Debrief

Overview

In Lab 2-1, students configure and verify EIGRP operations. First a student configures basic EIGRP and advertises all the specific subnets used in the network. Next, the student defines EIGRP path selection so that the primary path is preferred and the second path remains as a backup.

Because EIGRP uses a lot of bandwidth and CPU resources, the student must optimize EIGRP operation. The student must also configure EIGRP operation in scalable way, in which summarization is configured in order to improve convergence time and add stability.

Objectives

Upon completing this lesson, you will be able to explain how to configure and verify EIGRP operations. This ability includes being able to meet these objectives:

- Complete the lab overview and verification
- Describe a sample solution and alternatives

Lab Overview and Verification

This topic describes lab topology, as well as the key checkpoints used to create a solution and start with the verification.

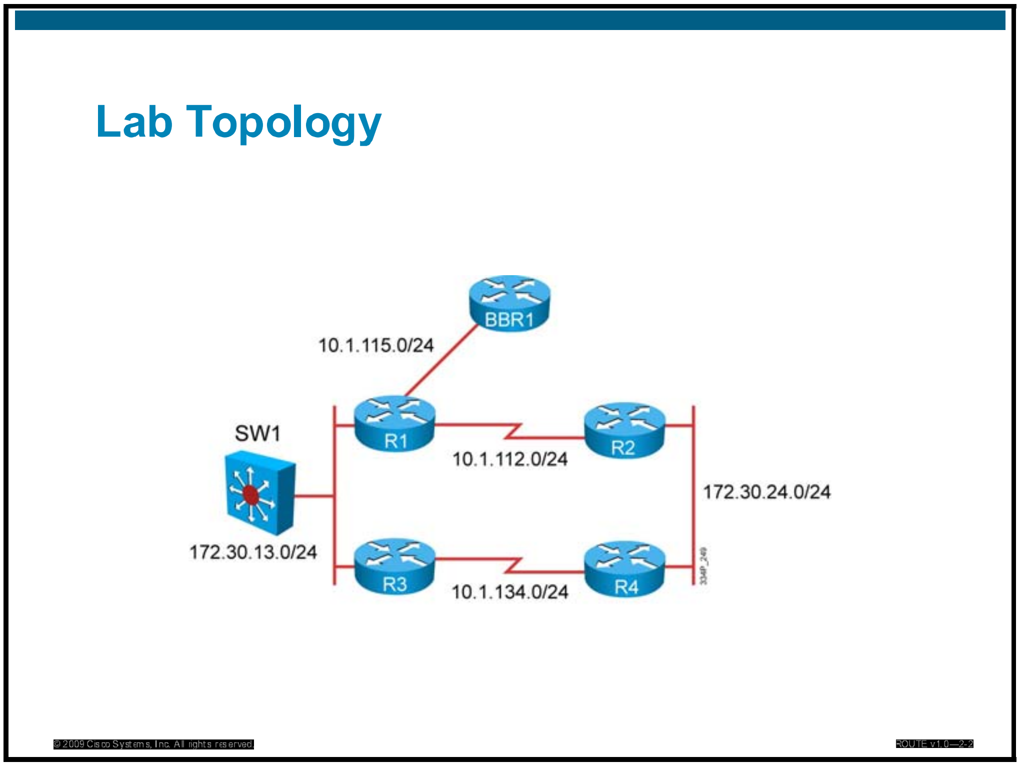

Lab Topology

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-2

The figure above presents the physical lab topology used for Lab 2-1: configure and verify EIGRP operations. The topology uses four pod routers, one backbone router, and one pod switch.

Based on the topology, students will identify the required parameters and configure a basic EIGRP routing protocol in order to establish Layer 3 reachability in the network, as well as influence EIGRP path selection, optimize EIGRP operation, and provide a scalable solution.

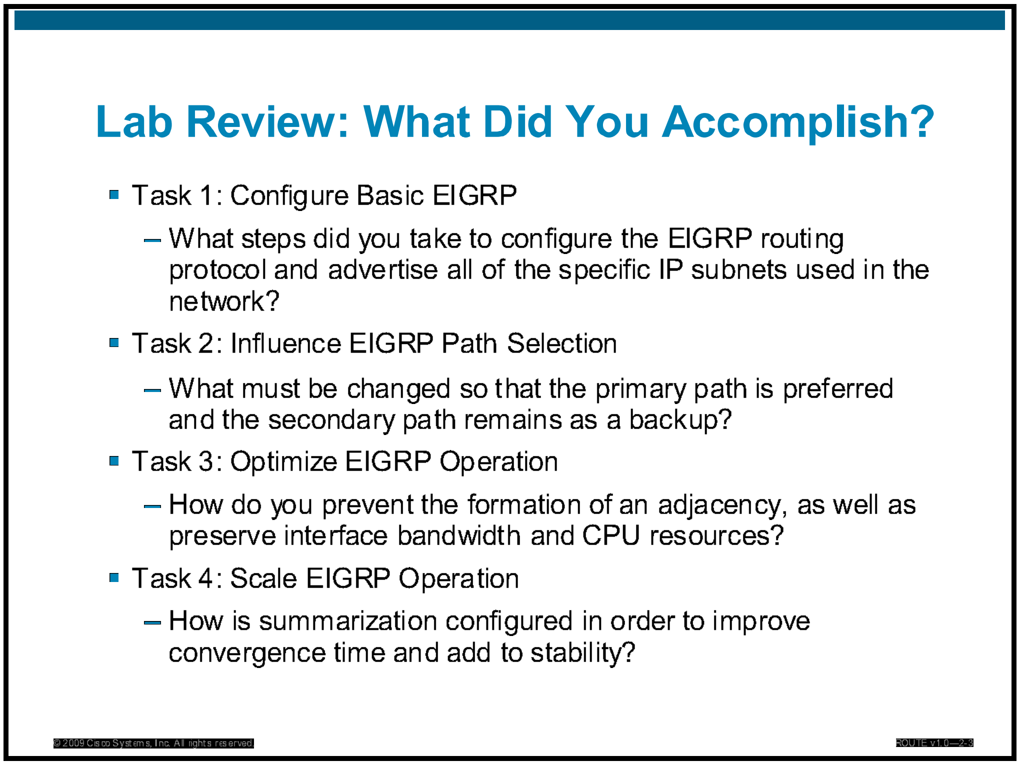

Lab Review: What Did You Accomplish?

- Task 1: Configure Basic EIGRP
  - What steps did you take to configure the EIGRP routing protocol and advertise all of the specific IP subnets used in the network?
- Task 2: Influence EIGRP Path Selection
  - What must be changed so that the primary path is preferred and the secondary path remains as a backup?
- Task 3: Optimize EIGRP Operation
  - How do you prevent the formation of an adjacency, as well as preserve interface bandwidth and CPU resources?
- Task 4: Scale EIGRP Operation
  - How is summarization configured in order to improve convergence time and add to stability?

In the first task, you configured basic EIGRP routing. All routers are configured for EIGRP routing protocol according to the implementation plan.

In the second task, you influenced the EIGRP path selection. You implemented redundancy in the network and the routers are able to select the primary path, while the backup path remains in the routing table. You must also change the metric in order to influence path selection in EIGRP routing protocol.

In the third task, you optimized the EIGRP operation. Routing update suppression prevented the formation of EIGRP adjacencies. At the same time, CPU resources were preserved without the use of filtering. The **passive-interface** command is the solution to optimize the EIGRP operation in this step.

In the fourth task, you configured scaling options of the EIGRP operation. Summarization was configured to summarize many subnets into one summary route, which was sent to the adjacent routers. Only one summary route is sent instead of many more-specific subnets.

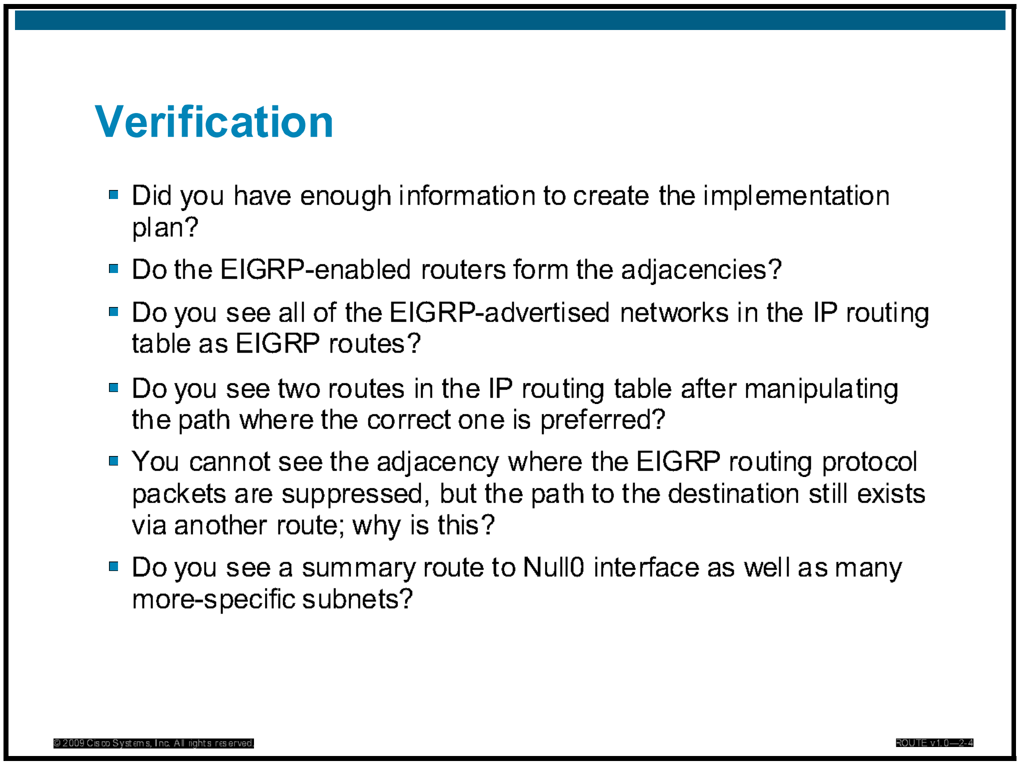

Verification

- Did you have enough information to create the implementation plan?
- Do the EIGRP-enabled routers form the adjacencies?
- Do you see all of the EIGRP-advertised networks in the IP routing table as EIGRP routes?
- Do you see two routes in the IP routing table after manipulating the path where the correct one is preferred?
- You cannot see the adjacency where the EIGRP routing protocol packets are suppressed, but the path to the destination still exists via another route; why is this?
- Do you see a summary route to Null0 interface as well as many more-specific subnets?

ROUTE v1.0—2-4

A common approach to verifying the implementation process for configuring EIGRP operations is as follows:

- In order to create the implementation plan, sufficient information must be gathered.
- After a successful EIGRP configuration, all neighboring routers running EIGRP form an adjacency.
- Adjacent routers start exchanging routing protocol information and EIGRP routes populate the IP routing table.
- When a redundant path exists, one of them becomes primary path. Path manipulation results in the desired path being the primary and redundant paths being the backups.
- You cannot see the adjacency where EIGRP routing protocol packets are suppressed, but the path to the destination still exists via another route.
- The router performing the summarization includes a summary route to the Null0 interface, as well as routes to more-specific subnets. The remaining EIGRP neighbors only have a summary route.

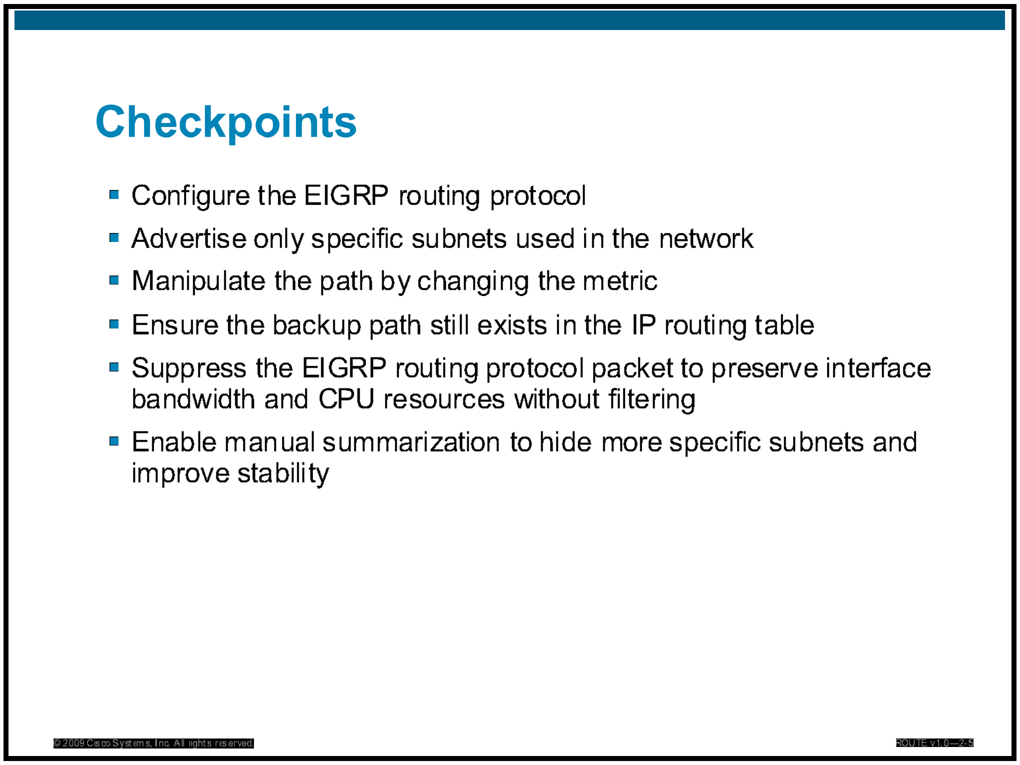

Checkpoints

- Configure the EIGRP routing protocol
- Advertise only specific subnets used in the network
- Manipulate the path by changing the metric
- Ensure the backup path still exists in the IP routing table
- Suppress the EIGRP routing protocol packet to preserve interface bandwidth and CPU resources without filtering
- Enable manual summarization to hide more specific subnets and improve stability

ROUTE v1.0—2-5

During the configuration and verification phase, a network operator can deal with several checkpoints. After completing all the configuration tasks, the network operator may have successfully configured EIGRP operations, or may need to do additional verification and troubleshooting.

With different checkpoints, the network operator can verify for proper configuration. The following checkpoints are used for verification:

- Configure the EIGRP routing protocol.
- Advertise only specific subnets used in the network.
- Manipulate the path by changing the metric.
- Ensure that the backup path still exists in the IP routing table.
- Suppress the EIGRP routing protocol packet to preserve interface bandwidth and CPU resources without filtering.
- Enable manual summarization to hide more specific subnets and improve stability.

Sample Solutions and Alternatives

This topic describes a sample solution and other alternatives.

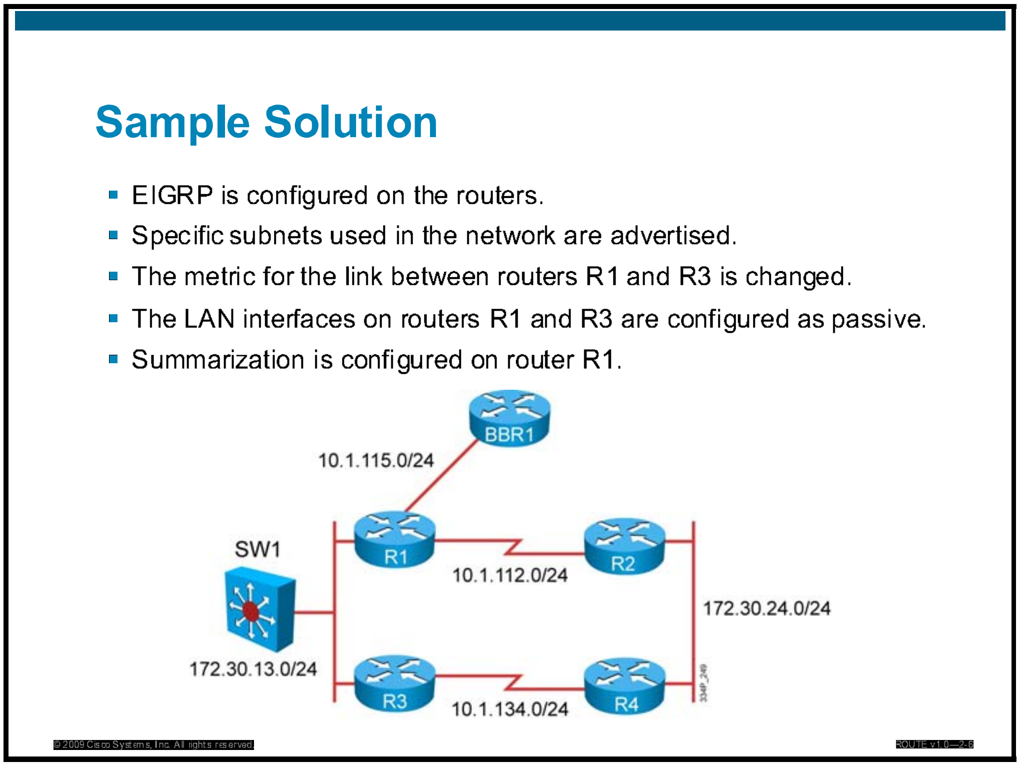

Sample Solution

- EIGRP is configured on the routers.
- Specific subnets used in the network are advertised.
- The metric for the link between routers R1 and R3 is changed.
- The LAN interfaces on routers R1 and R3 are configured as passive.
- Summarization is configured on router R1.

A sample solution includes implementation details and details for each task of the implementation plan. Different solutions are possible and the figure above shows a few details of a successful configuration.

The proper implementation of route redistribution between multiple IP routing protocols includes the following checkpoints:

- EIGRP routing protocol is configured on the pod routers.
- Specific subnets used in the network are advertised.
- The metric for the link between routers R1 and R3 is changed in order to manipulate the primary and backup path.
- The LAN interface on routers R1 and R3 is configured as passive to suppress EIGRP routing packets and preserve link bandwidth and CPU resources.
- Summarization is configured on router R1 for a scalable EIGRP implementation that results in the summary route being advertised only to other neighbors.

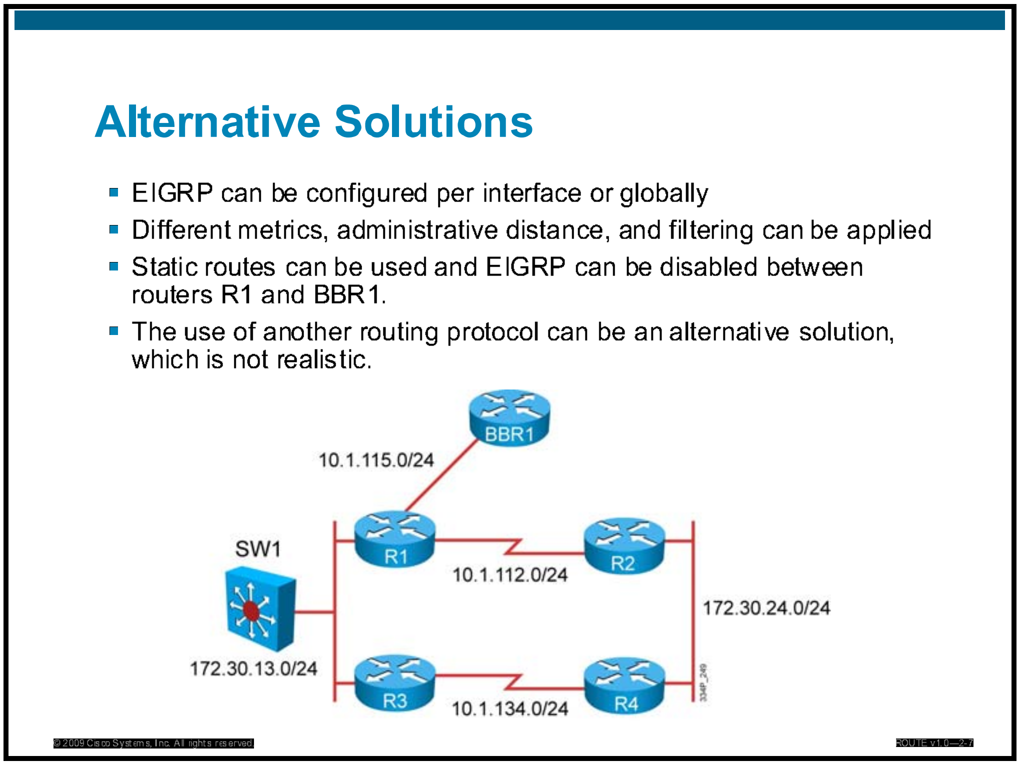

Alternative Solutions

- EIGRP can be configured per interface or globally
- Different metrics, administrative distance, and filtering can be applied
- Static routes can be used and EIGRP can be disabled between routers R1 and BBR1.
- The use of another routing protocol can be an alternative solution, which is not realistic.

ROUTE v1.0—2-7

Different metrics, administrative distance, and filtering can be applied to change the behavior of the EIGRP routing protocol. They can also be applied to manipulate the path when there is a redundancy in the network or to preserve CPU resources.

You can use static routes, as well, but if all of the routers are configured with static routes, the solution will not be scalable. One of the options is to disable EIGRP between routers R1 and BBR1 and configure the default route or several static routes pointing toward router BBR1.

The use of another routing protocol can be an alternative solution, which is not realistic as changing the routing protocol is not the case during fine tuning of the existing protocol.

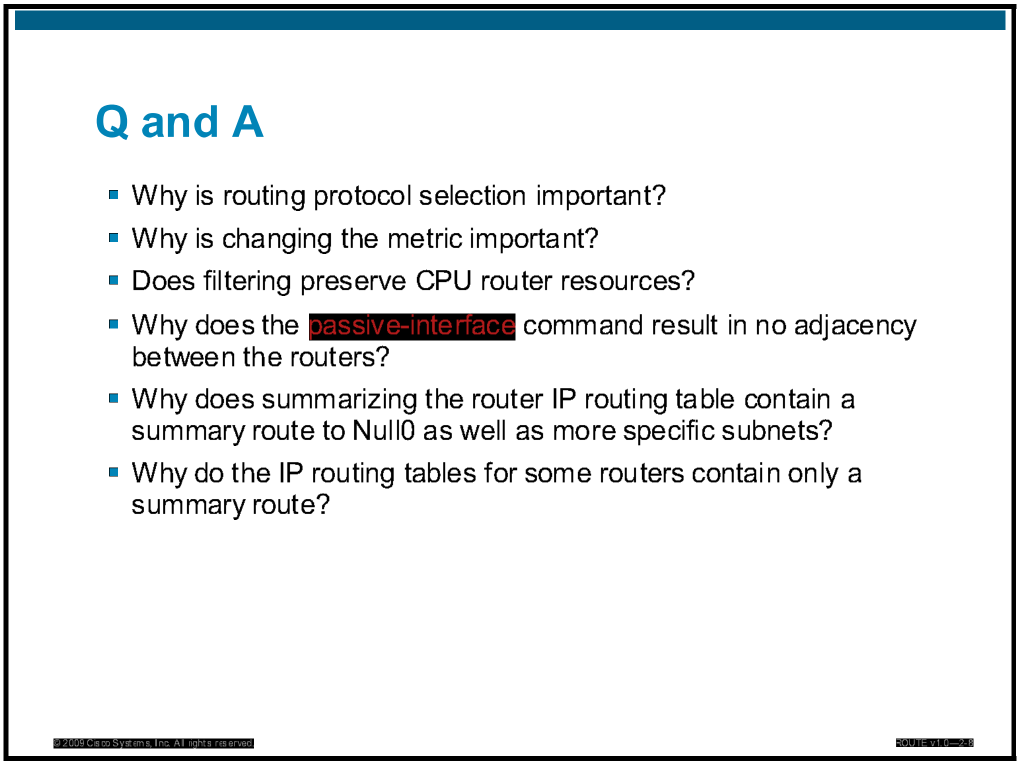

Q and A

- Why is routing protocol selection important?
- Why is changing the metric important?
- Does filtering preserve CPU router resources?
- Why does the passive-interface command result in no adjacency between the routers?
- Why does summarizing the router IP routing table contain a summary route to Null0 as well as more specific subnets?
- Why do the IP routing tables for some routers contain only a summary route?

ROUTE v1.0—2-8

A routing protocol exchanges routing updates and populates the IP routing table, which is used for destination-based forwarding. Different routing protocols process routing updates in different ways.

The metric defines the importance and the quality of the routes in the routing protocol. By manipulating the administrative distance and metric value, you can implement path manipulation, as well.

Filtering does not preserve CPU resources.

The **passive-interface** command suppresses routing protocol packets, preventing routers from forming adjacencies.

The summary route is advertised and the neighboring routers send packets to the router that is summarizing the subnets in the routing table. If one of the more-specific subnets is lost, the router still sends the summary route to its neighbors, but the destination is not reachable and packets must be dropped (sent to the Null0 interface).

A router summarizing the subnets contains the summary route as well as more-specific subnets. Because the idea of summarization is to decrease the sizes of the routing tables for neighboring routers, only a summary route is sent. This is enough to preserve the connectivity in the network.

Summary

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

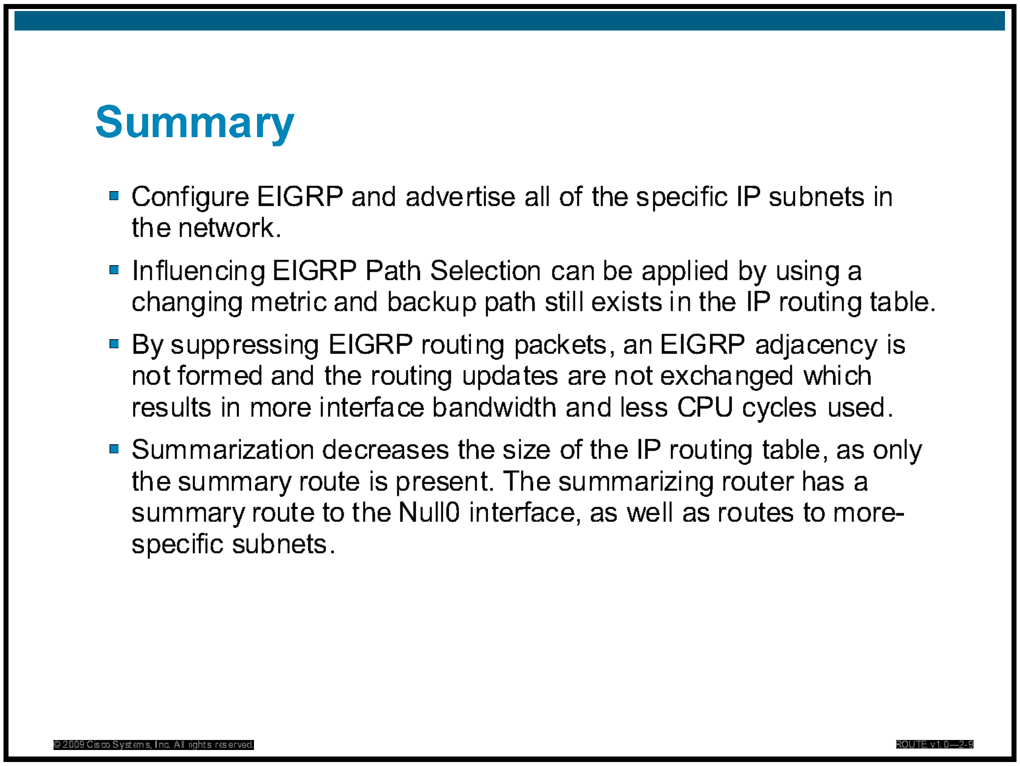

Summary

- Configure EIGRP and advertise all of the specific IP subnets in the network.
- Influencing EIGRP Path Selection can be applied by using a changing metric and backup path still exists in the IP routing table.
- By suppressing EIGRP routing packets, an EIGRP adjacency is not formed and the routing updates are not exchanged which results in more interface bandwidth and less CPU cycles used.
- Summarization decreases the size of the IP routing table, as only the summary route is present. The summarizing router has a summary route to the Null0 interface, as well as routes to more-specific subnets.

ROUTE v1.0—2-9

Lesson 4

Configuring and Verifying EIGRP for the Enterprise WAN Architecture

Overview

EIGRP can operate over various underlying network technologies—Ethernet over Multiprotocol Label Switching (EoMPLS), MPLS virtual private network (MPLS VPN), and physical Frame Relay, as well as multipoint and point-to-point Frame Relay subinterfaces. Load balancing across multiple links is a valuable option for efficient bandwidth utilization. If you limit the amount of bandwidth that Enhanced Interior Gateway Routing Protocol (EIGRP) uses across these WAN links, you can provide user traffic with better access to the WAN links.

This lesson provides insight into EIGRP deployment over various WAN technologies, as well as advanced configuration options for EIGRP load balancing and limitation of EIGRP bandwidth utilization on WAN links.

Objectives

Upon completing this lesson, you will be able to describe, recognize, and deploy EIGRP over various WAN technologies and be able to scale the deployment with load balancing and proper bandwidth utilization. This ability includes being able to meet these objectives:

- Configure and verify EIGRP over Frame Relay and on a physical interface.
- Configure and verify EIGRP over multipoint subinterfaces.
- Configure and verify EIGRP over point-to-point subinterfaces.
- Implement load balancing across equal-metric paths.
- Implement load balancing across unequal-metric aths.
- Determine EIGRP bandwidth use across WAN links.
- Implement EIGRP over Layer 2 and Layer 3 MPLS VPN.

EIGRP Over Frame Relay and on a Physical Interface

This topic describes how EIGRP can be deployed using Frame Relay physical interfaces.

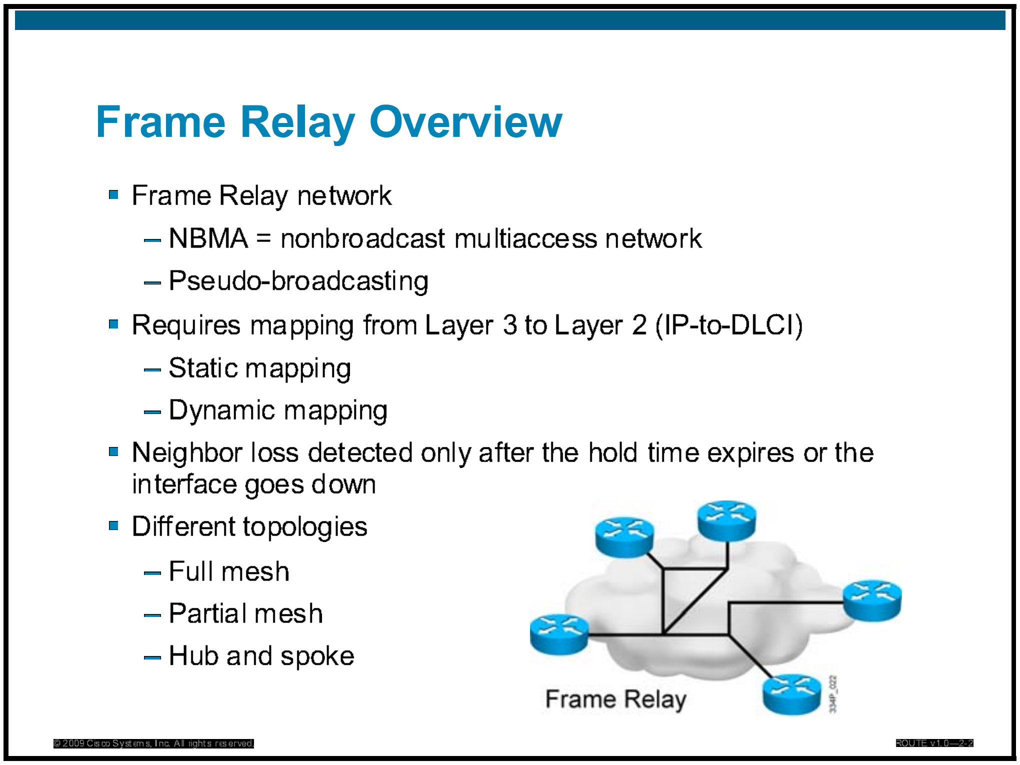

Frame Relay Overview

- Frame Relay network
  - NBMA = nonbroadcast multiaccess network
  - Pseudo-broadcasting
- Requires mapping from Layer 3 to Layer 2 (IP-to-DLCI)
  - Static mapping
  - Dynamic mapping
- Neighbor loss detected only after the hold time expires or the interface goes down
- Different topologies
  - Full mesh
  - Partial mesh
  - Hub and spoke

Frame Relay
334P_022

ROUTE v1.0—2-2

Frame Relay is a switched WAN technology for which virtual circuits are created through the network. In order to provide IP layer connectivity mapping between IP addresses and data-link connection identifiers (DLCIs) must be deployed—either dynamically or statically.

Usually, switched WAN networks do not support broadcasting capability equivalent to LAN broadcasting. To emulate the LAN broadcasting capability that is required by IP routing protocols (for example, to send hello or update packets to all neighbors reachable over an IP subnet), Cisco IOS Software implements pseudo-broadcasting. This is when the Layer 2 code in Cisco IOS Software creates several copies of the same broadcast or multicast packet—one for each neighbor reachable through the WAN media.

In environments where a single router has a large number of neighbors reachable through a single WAN interface, pseudo-broadcasting must be tightly controlled, as it can use a large amount of CPU time and WAN bandwidth. Pseudo-broadcasting is controlled with the broadcast option specified on static maps in a Frame Relay configuration. Pseudo-broadcasting cannot be controlled for neighbors reachable through dynamic maps created via Inverse Address Resolution Protocol (Inverse ARP) on a Frame Relay, as dynamic maps always allow pseudo-broadcasting. To control pseudo-broadcasting in Frame Relay, you must define the manual static maps and disable Inverse ARP.

Neighbor loss is detected only after the hold time expires or the interface goes down. It is important to know that an interface is up as long as at least one DLCI is alive.

The different topologies for Frame Relay are full mesh, partial mesh, and hub and spoke.

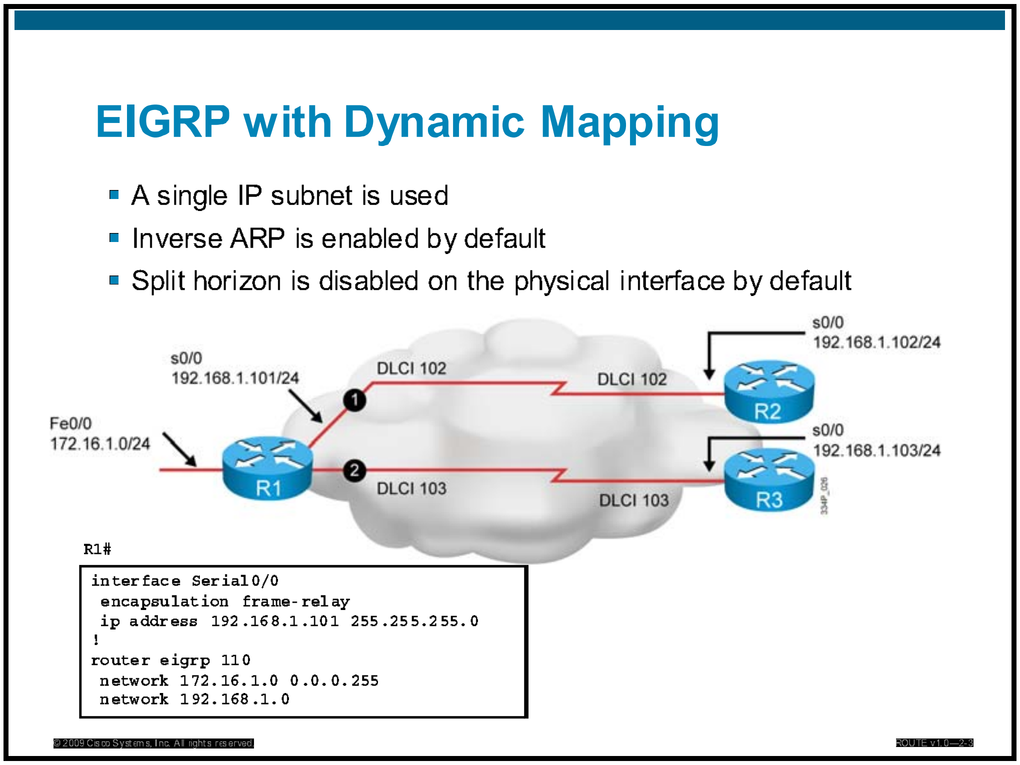

EIGRP with Dynamic Mapping

- A single IP subnet is used
- Inverse ARP is enabled by default
- Split horizon is disabled on the physical interface by default

```
R1#
interface Serial0/0
 encapsulation frame-relay
 ip address 192.168.1.101 255.255.255.0
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

To deploy EIGRP over a physical interface using dynamic mapping, thus relying on Inverse ARP, no changes are needed to the basic configuration. The EIGRP process is enabled using the required autonomous system (AS) number (110 in the example on the figure). Proper interfaces and networks can also be included in the topology by specifying the **network** command under the EIGRP routing process.

The split-horizon behavior is disabled by default on the physical interface. Therefore, routers R2 and R3 can provide connectivity between their connected networks. Inverse ARP does not provide dynamic mapping for the communication between routers R2 and R3; this must be configured manually.

EIGRP with Dynamic Mapping (Cont.)

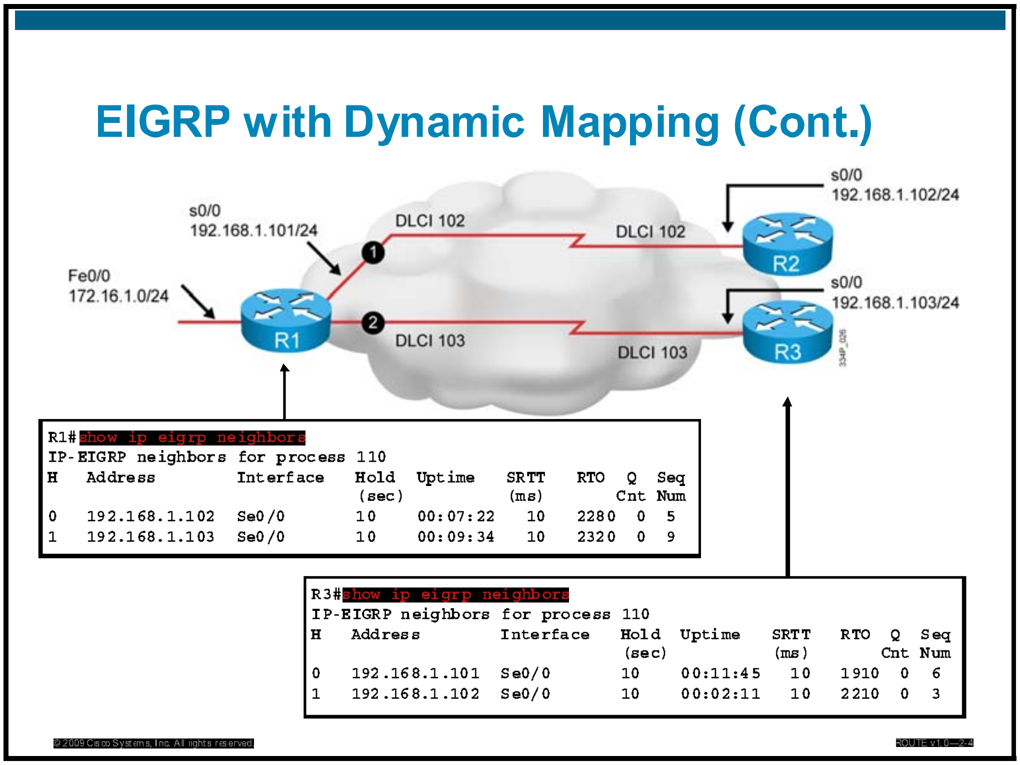

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT    RTO  Q   Seq
                                  (sec)            (ms)        Cnt  Num
0   192.168.1.102   Se0/0         10    00:07:22    10     2280  0   5
1   192.168.1.103   Se0/0         10    00:09:34    10     2320  0   9
```

```
R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT    RTO  Q   Seq
                                  (sec)            (ms)        Cnt  Num
0   192.168.1.101   Se0/0         10    00:11:45    10     1910  0   6
1   192.168.1.102   Se0/0         10    00:02:11    10     2210  0   3
```

The sample outputs of the **show ip eigrp** neighbors command on this slide show the neighbors of routers R1 and R3. Router R1 forms an adjacency with router R2 and another with router R3 over the serial0/0 physical interface. Likewise, routers R2 and R3 form adjacencies with router R1. Apart from that, they can also form an EIGRP adjacency to each other if the IP-to-DLCI mapping for that connectivity is also provided. On the sample output for router R3, it is apparent that router R3 has adjacencies to routers R1 and R2.

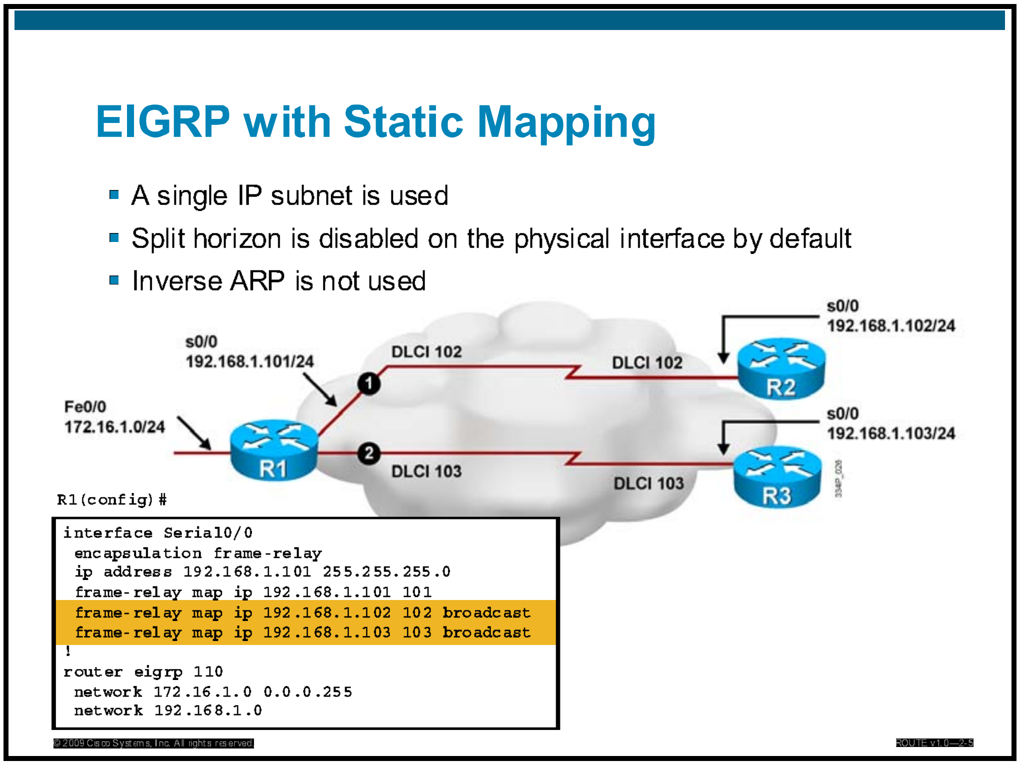

EIGRP with Static Mapping

- A single IP subnet is used
- Split horizon is disabled on the physical interface by default
- Inverse ARP is not used

```
R1(config)#
interface Serial0/0
 encapsulation frame-relay
 ip address 192.168.1.101 255.255.255.0
 frame-relay map ip 192.168.1.101 101
 frame-relay map ip 192.168.1.102 102 broadcast
 frame-relay map ip 192.168.1.103 103 broadcast
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

ROUTE v1.0—2-5

To deploy EIGRP over a physical interface on router R1 using static mapping, thus disabling the Inverse ARP, no changes are needed to the basic configuration of EIGRP. The EIGRP process is enabled using the required AS number (110 in the example in the figure above). Proper interfaces and networks are included in the topology by specifying the **network** command under the EIGRP routing process. In addition, manual IP-to-DLCI mapping statements on the serial 0/0 interface are necessary on all three routers: R1, R2, and R3.

The split-horizon behavior is disabled by default on the physical interface, thus routers R2 and R3 can provide connectivity between their connected networks, as well.

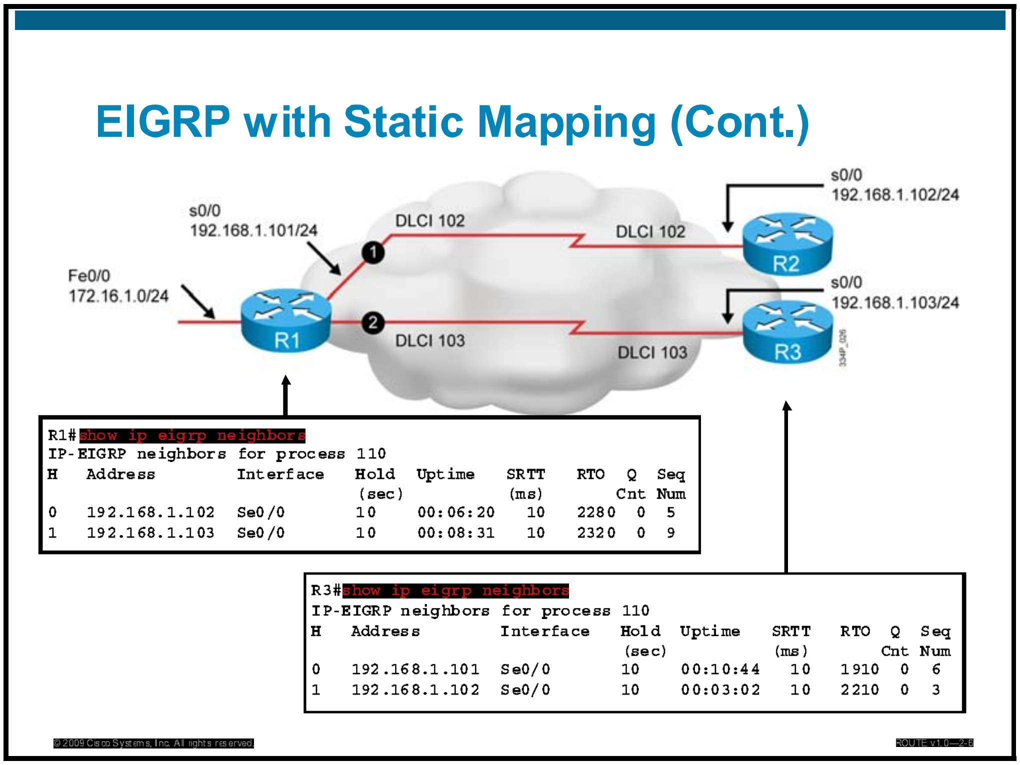

EIGRP with Static Mapping (Cont.)

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT     RTO   Q   Seq
                                  (sec)            (ms)          Cnt  Num
0   192.168.1.102   Se0/0         10    00:06:20    10      2280  0   5
1   192.168.1.103   Se0/0         10    00:08:31    10      2320  0   9
```

```
R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT     RTO   Q   Seq
                                  (sec)            (ms)          Cnt  Num
0   192.168.1.101   Se0/0         10    00:10:44    10      1910  0   6
1   192.168.1.102   Se0/0         10    00:03:02    10      2210  0   3
```

ROUTE v1.0—2-6

The figure above shows the adjacency formed between router R1 and routers R2 and R3 over the serial0/0 physical interface. The adjacency formed using static mapping is the same as the one formed using dynamic mapping. The same applies to routers R2 and R3, which form the adjacency with router R1. Routers R2 and R3 can also form an EIGRP adjacency to each other if the IP-to-DLCI mapping for that connectivity is provided. The sample output from the **show ip eigrp neighbors** command in the figure above shows the neighbors on router R1 and R3. It is apparent that router R1 has two neighbors (routers R2 and R3), and router R3 also has two neighbors (routers R1 and R2).

EIGRP over Multipoint Subinterfaces

This topic describes how EIGRP can be deployed using Frame Relay multipoint subinterfaces.

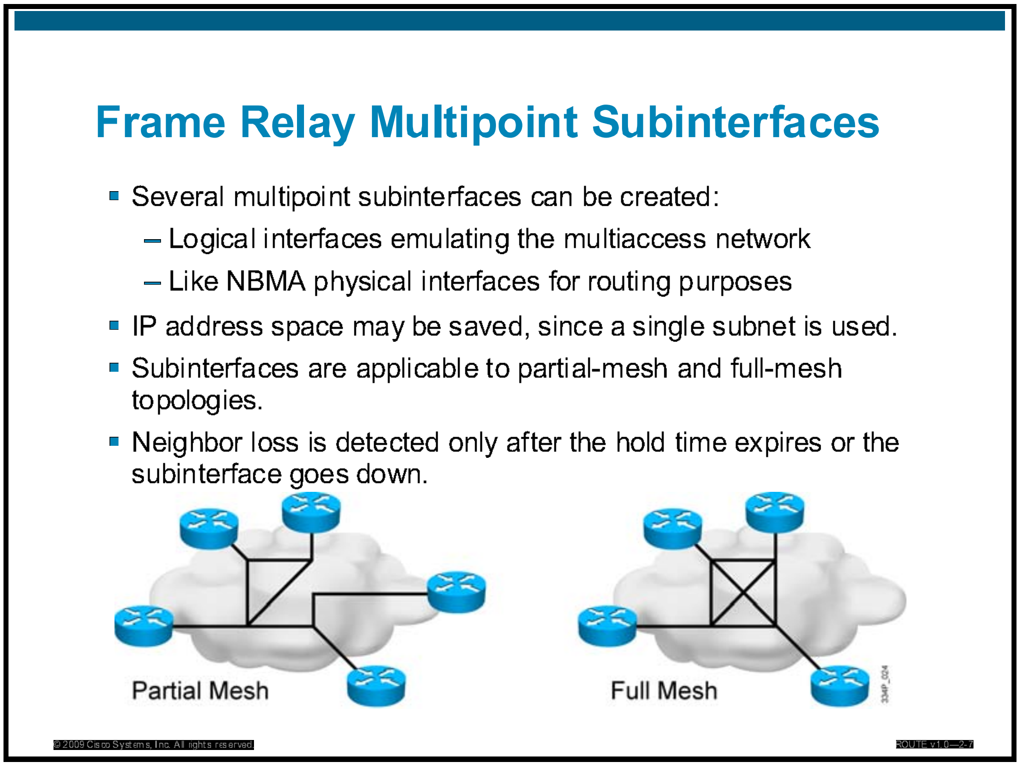

Frame Relay Multipoint Subinterfaces

- Several multipoint subinterfaces can be created:
  - Logical interfaces emulating the multiaccess network
  - Like NBMA physical interfaces for routing purposes
- IP address space may be saved, since a single subnet is used.
- Subinterfaces are applicable to partial-mesh and full-mesh topologies.
- Neighbor loss is detected only after the hold time expires or the subinterface goes down.

ROUTE v1.0—2-7

Several subinterfaces can be created over Frame Relay interfaces. They are logical interfaces emulating a multi-access network and provide the routing equivalent to nonbroadcast multiaccess (NBMA) physical interfaces. As with NBMA physical interfaces, a single subnet is used—preserving the IP address space.

EIGRP neighbor loss detection is particularly slow on multipoint subinterfaces configured over low-speed WAN links. This is because the default values of the EIGRP timers on these interfaces are 60 seconds for the hello timer and 180 seconds for the hold timer. In the worst case, neighbor loss detection can take up to 3 minutes.

Frame Relay multipoint is applicable to partial-mesh and full-mesh topologies. Partial-mesh Frame Relay networks must deal with the possibility of a split horizon, which prevents routing updates from being retransmitted on the same interface on which they were received.

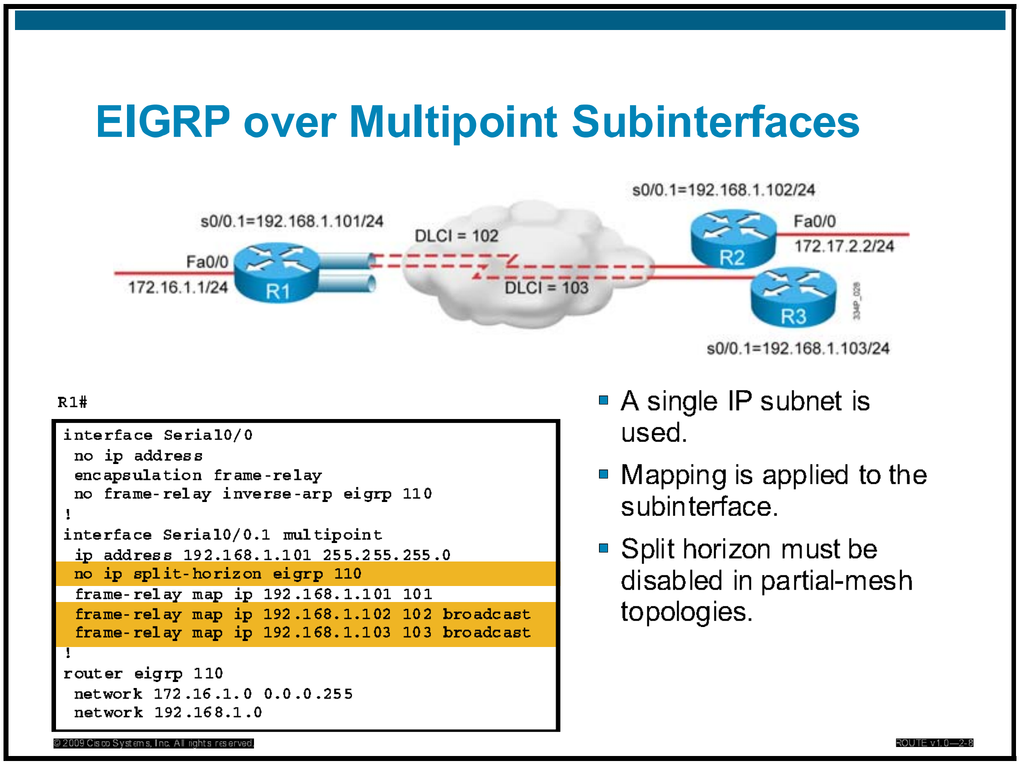

EIGRP over Multipoint Subinterfaces

R1#

```
interface Serial0/0
 no ip address
 encapsulation frame-relay
 no frame-relay inverse-arp eigrp 110
!
interface Serial0/0.1 multipoint
 ip address 192.168.1.101 255.255.255.0
 no ip split-horizon eigrp 110
 frame-relay map ip 192.168.1.101 101
 frame-relay map ip 192.168.1.102 102 broadcast
 frame-relay map ip 192.168.1.103 103 broadcast
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

- A single IP subnet is used.
- Mapping is applied to the subinterface.
- Split horizon must be disabled in partial-mesh topologies.

ROUTE v1.0—2-8

To use the multipoint behavior of a subinterface, you must add the **multipoint** keyword at the end of the **interface** command when creating the subinterface. With Frame Relay, the mapping on the multipoint subinterfaces is created by using the proper local DLCI value:

- By specifying the proper local DLCI value and relying on the Inverse ARP
- With manual IP-to-DLCI mapping

EIGRP is configured with no changes to the basic deployment. To enable the EIGRP process, use the required AS number (110 in the example on the figure) and include the proper interfaces and networks in the topology via the **network** command under the EIGRP routing process. Additionally, you can configure manual IP-to-DLCI mapping statements (via the **frame-relay map** command with the **broadcast** keyword on the serial 0/0 multipoint subinterfaces on routers R1, R2, and R3), in order to define the mapping between a destination protocol address and the DLCI used to connect to the destination address.

If routers R2 and R3 need to provide connectivity between their connected networks, you must disable the EIGRP split horizon on the multipoint subinterface of router R1.

The router R1 configuration includes the **frame-relay map** command to its own IP address on the multipoint serial subinterface in order to ping the local address for router R1 from router R1 itself.

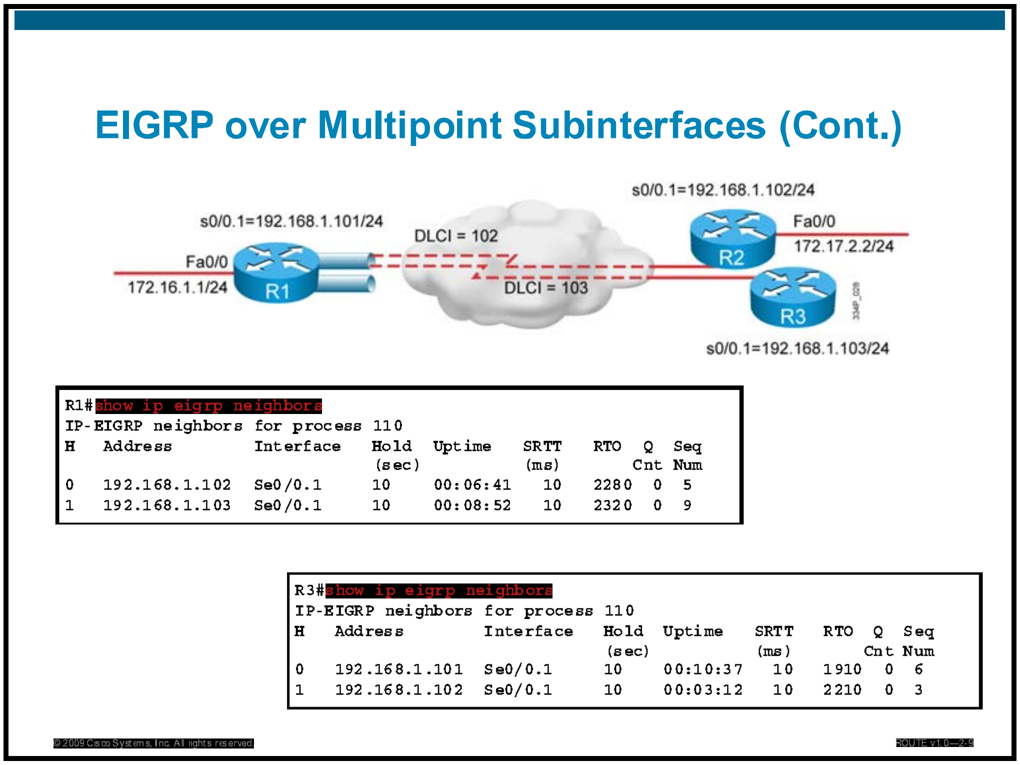

EIGRP over Multipoint Subinterfaces (Cont.)

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface    Hold  Uptime    SRTT   RTO  Q  Seq
                                 (sec)           (ms)       Cnt Num
0   192.168.1.102   Se0/0.1      10    00:06:41    10   2280  0  5
1   192.168.1.103   Se0/0.1      10    00:08:52    10   2320  0  9
```

```
R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface    Hold  Uptime    SRTT   RTO  Q  Seq
                                 (sec)           (ms)       Cnt Num
0   192.168.1.101   Se0/0.1      10    00:10:37    10   1910  0  6
1   192.168.1.102   Se0/0.1      10    00:03:12    10   2210  0  3
```

ROUTE v1.0—2-9

If you want to verify the operation of the EIGRP routing protocol over the Frame Relay multipoint subinterface, you can use the **show ip eigrp neighbors** command. The figure above shows two sample outputs for routers R1 and R3. Router R1 forms the adjacency with routers R2 and R3 over the serial0/0.1 multipoint subinterface. This is done in the same way routers R2 and R3 form the adjacency with router R1 and between each other if IP-to-DLCI mapping for that connectivity is provided.

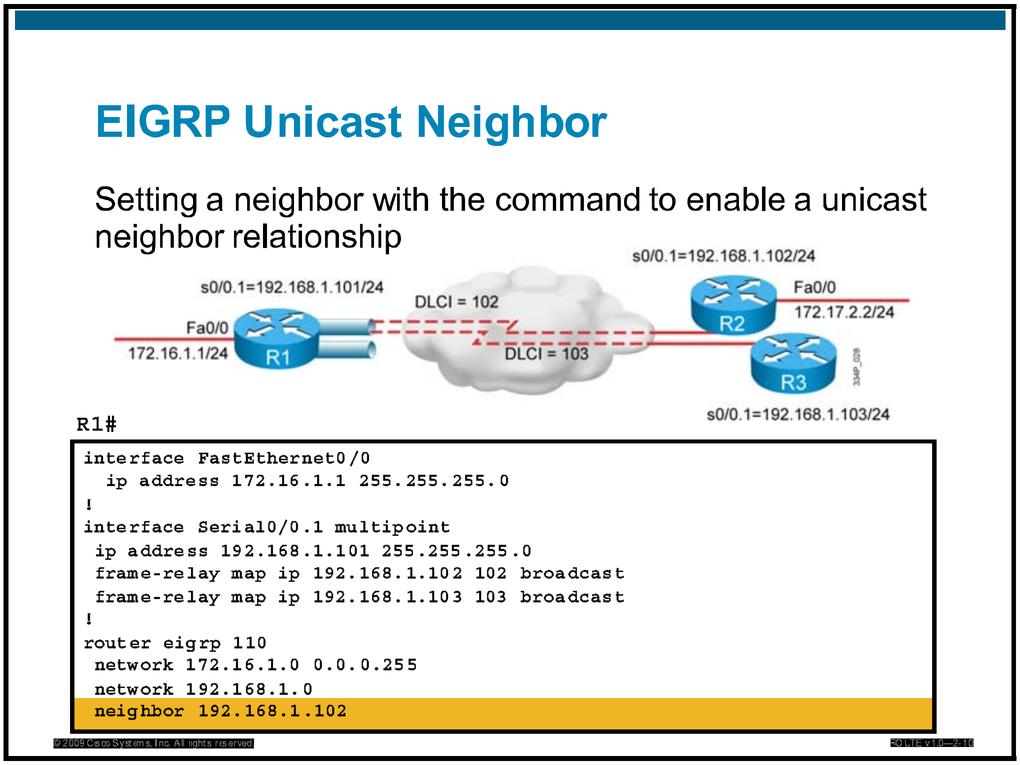

EIGRP Unicast Neighbor

Setting a neighbor with the command to enable a unicast neighbor relationship

R1#

```
interface FastEthernet0/0
  ip address 172.16.1.1 255.255.255.0
!
interface Serial0/0.1 multipoint
 ip address 192.168.1.101 255.255.255.0
 frame-relay map ip 192.168.1.102 102 broadcast
 frame-relay map ip 192.168.1.103 103 broadcast
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
 neighbor 192.168.1.102
```

ROUTE v1.0—2-10

The **neighbor** command is used in EIGRP in order to define a neighboring router with which to exchange routing information. Instead of using multicast packets, EIGRP exchanges routing information with the neighbors in the form of unicast packets whenever the **neighbor** command is configured for an interface. EIGRP stops processing all multicast packets that come inbound on that interface. At the same time, EIGRP stops sending multicast packets on that interface. Multiple neighbor statements can be used to establish peering sessions with specific EIGRP neighbors. The interface through which EIGRP will exchange routing updates must be specified in the neighbor statement. The interfaces through which two EIGRP neighbors exchange routing updates must be configured with IP addresses from the same network.

Note EIGRP neighbor adjacencies cannot be established or maintained over an interface that is configured as passive.

Router R1 is configured with the **neighbor** command for router R2.

For more details about the **neighbor** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

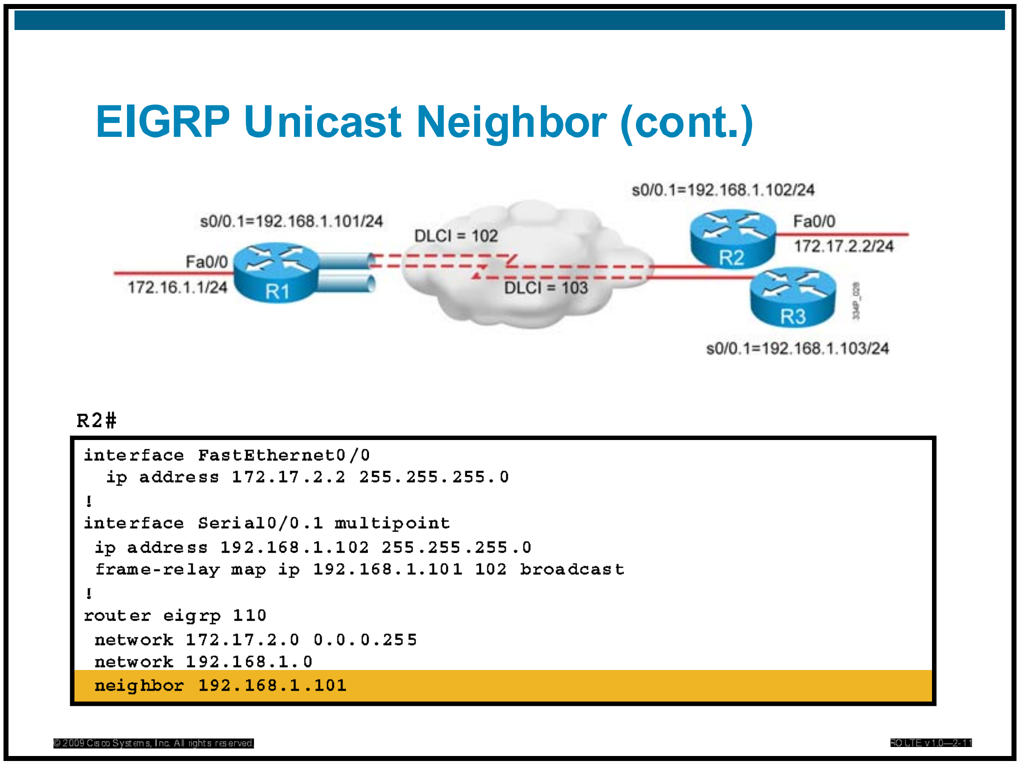

EIGRP Unicast Neighbor (cont.)

R2#

```
interface FastEthernet0/0
  ip address 172.17.2.2 255.255.255.0
!
interface Serial0/0.1 multipoint
 ip address 192.168.1.102 255.255.255.0
 frame-relay map ip 192.168.1.101 102 broadcast
!
router eigrp 110
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
 neighbor 192.168.1.101
```

Router R1 is configured with the **neighbor** command and will not accept multicast packets anymore. In order to establish an adjacency, router R2 must be configured as well. In the example above, router R2 is configured with the **neighbor** command for router R1, which enables the use of unicast packets accepted by router R1.

Router R3 is not configured with the **neighbor** command.

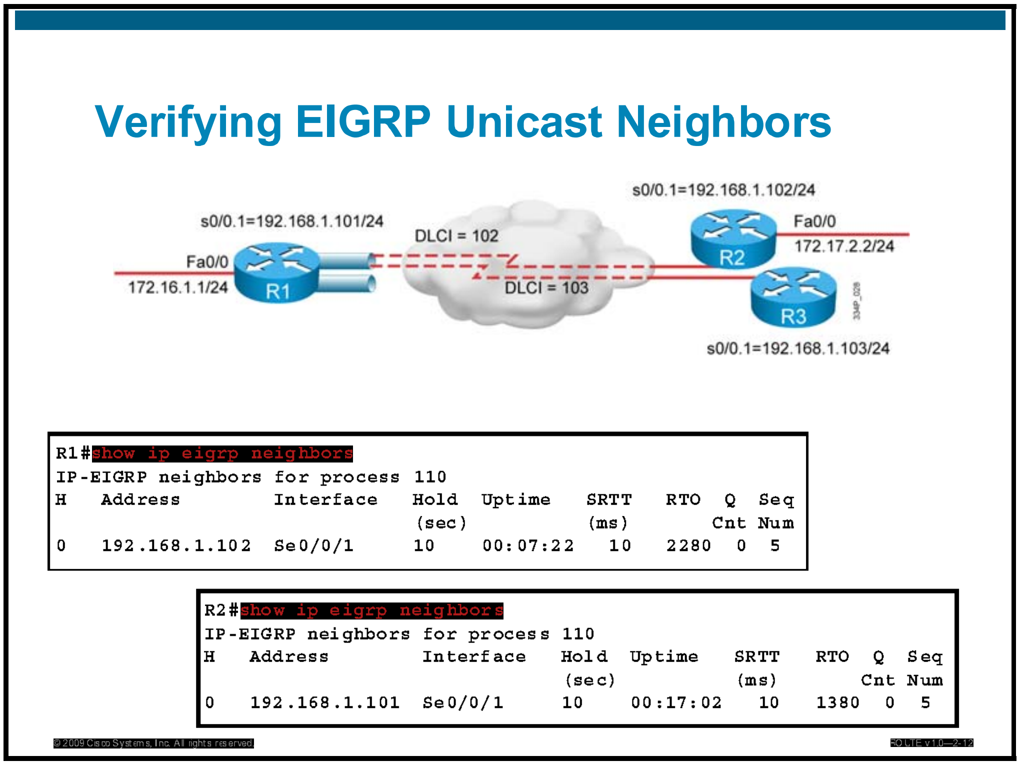

Verifying EIGRP Unicast Neighbors

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q  Seq
                                  (sec)           (ms)       Cnt Num
0   192.168.1.102   Se0/0/1       10    00:07:22   10    2280  0   5
```

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q  Seq
                                  (sec)           (ms)       Cnt Num
0   192.168.1.101   Se0/0/1       10    00:17:02   10    1380  0   5
```

ROUTE v1.0—2-12

To verify the configuration, use the **show ip eigrp neighbors** command. As you can see in the output above, routers R1 and R2 have formed the neighbor relationship. The sample output does not show that the **neighbor** command was used on both routers. It only indicates that a neighbor relationship was established, which is a proof that the configuration was successfully completed.

Router R3 is not using the **neighbor** command and no neighbor relationship was established, because routers R1 and R2 are not accepting multicast packets.

For more details about the **show ip eigrp neighbors** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

EIGRP over Point-to-Point Subinterfaces

This topic describes how EIGRP can be deployed using Frame Relay point-to-point subinterfaces.

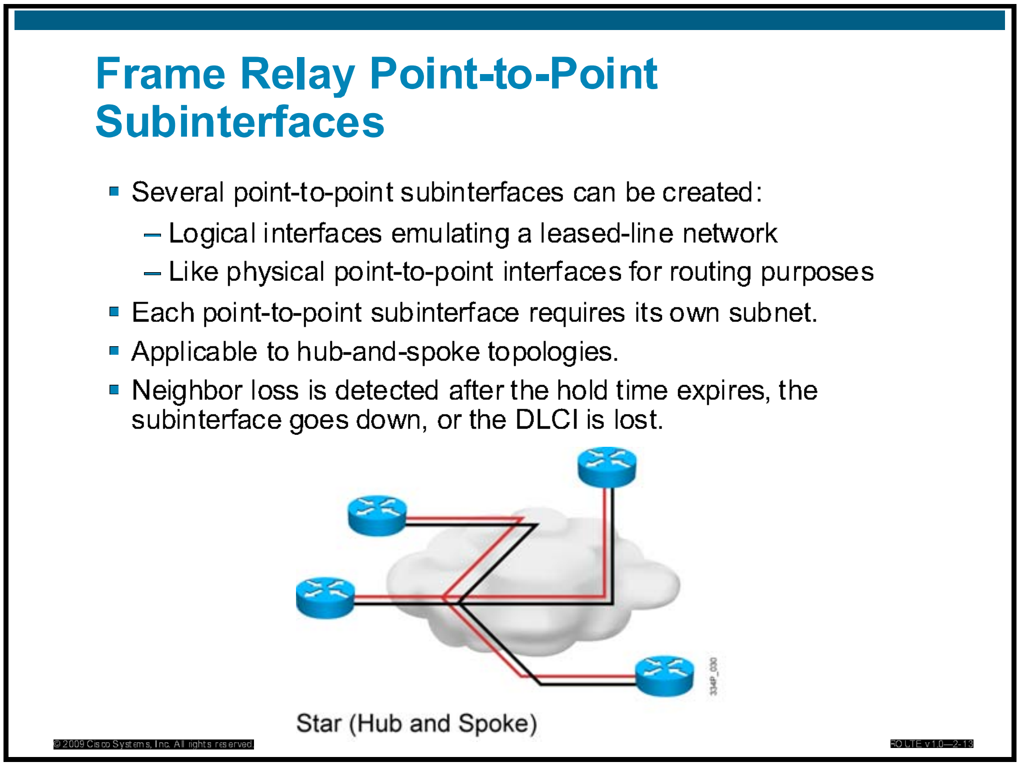

Frame Relay Point-to-Point Subinterfaces

- Several point-to-point subinterfaces can be created:
  - Logical interfaces emulating a leased-line network
  - Like physical point-to-point interfaces for routing purposes
- Each point-to-point subinterface requires its own subnet.
- Applicable to hub-and-spoke topologies.
- Neighbor loss is detected after the hold time expires, the subinterface goes down, or the DLCI is lost.

Star (Hub and Spoke)

ROUTE v1.0—2-13

Several point-to-point subinterfaces can be created over Frame Relay interfaces. They are logical interfaces emulating a leased line network and provide a routing equivalent to point-to-point physical interfaces. As with physical point-to-point interfaces, each interface requires its own subnet. Frame Relay point-to point is applicable to hub-and-spoke topologies.

EIGRP neighbor loss detection is quite fast on point-to-point subinterfaces for several reasons:

- The default values of the EIGRP hello timer and the EIGRP hold timer are identical to the values used on point-to-point links (5 seconds for the hello timer and 15 seconds for the hold timer). In the worst case, the neighbor loss is detected within 15 seconds.
- On Frame Relay networks, the subinterface is declared down if the DLCI attached to the interface is lost and neighbor loss detection is immediate. For multipoint subinterfaces, all of the permanent virtual circuits (PVCs) attached to it must be lost for the interface to be declared down.

Note Neighbor loss detection due to DLCI loss only works if the Frame Relay network supports end-to-end Integrated Local Management Interface (ILMI) signaling. On some Frame Relay networks, one end of the connections might fail (for example, due to a router failure), but the DLCI will still be declared operational at the other end of the connection.

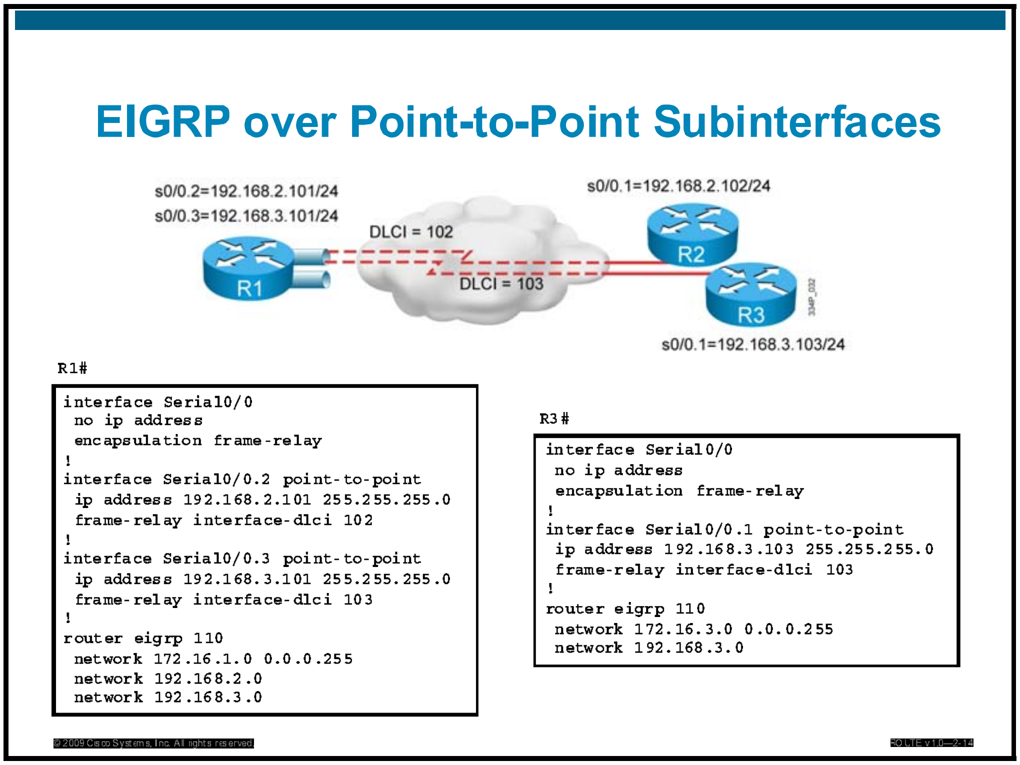

EIGRP over Point-to-Point Subinterfaces

R1#

```
interface Serial0/0
 no ip address
 encapsulation frame-relay
!
interface Serial0/0.2 point-to-point
 ip address 192.168.2.101 255.255.255.0
 frame-relay interface-dlci 102
!
interface Serial0/0.3 point-to-point
 ip address 192.168.3.101 255.255.255.0
 frame-relay interface-dlci 103
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.2.0
 network 192.168.3.0
```

R3#

```
interface Serial0/0
 no ip address
 encapsulation frame-relay
!
interface Serial0/0.1 point-to-point
 ip address 192.168.3.103 255.255.255.0
 frame-relay interface-dlci 103
!
router eigrp 110
 network 172.16.3.0 0.0.0.255
 network 192.168.3.0
```

ROUTE v1.0—2-14

To enable subinterfaces for point-to-point, you need to create them using the **point-to-point** keyword at the end of the **interface** command. With Frame Relay, the mapping of the point-to-point subinterfaces is created by specifying the proper local DLCI value.

EIGRP is configured with no changes to the basic deployment. To enable the EIGRP process, use the required AS number (110 in the example on the figure) and include the proper interfaces and networks in the topology via the **network** command under the EIGRP routing process. Additionally, you can configure manual IP-to-DLCI mapping statements (via the **frame-relay map** command with the **broadcast** keyword on the serial 0/0 multipoint subinterfaces on routers R1, R2, and R3).

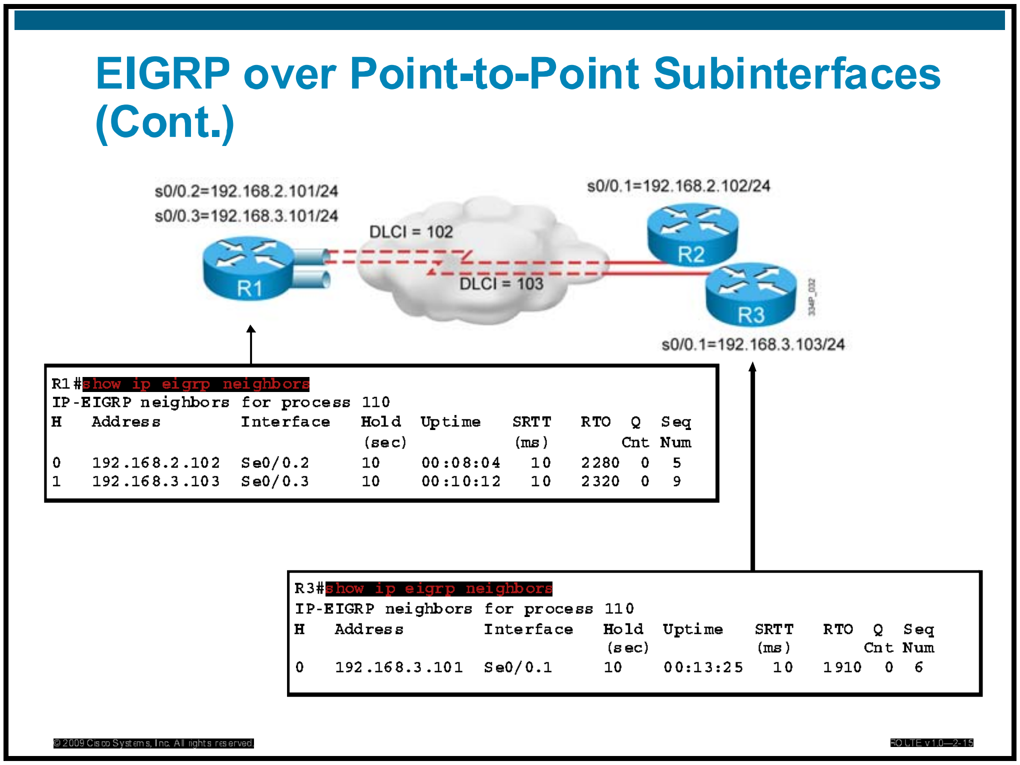

EIGRP over Point-to-Point Subinterfaces (Cont.)

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT    RTO  Q   Seq
                                  (sec)            (ms)        Cnt  Num
0   192.168.2.102   Se0/0.2       10    00:08:04    10     2280  0   5
1   192.168.3.103   Se0/0.3       10    00:10:12    10     2320  0   9
```

```
R3#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime     SRTT    RTO  Q   Seq
                                  (sec)            (ms)        Cnt  Num
0   192.168.3.101   Se0/0.1       10    00:13:25    10     1910  0   6
```

ROUTE v1.0—2-15

The **show ip eigrp neighbors** command can be used to verify the operation of the EIGRP routing protocol over the Frame Relay point-to point subinterface. The figure above shows two sample outputs for router R1 and R3. Router R1 forms an adjacency with router R2 over the serial0/0.2 point-to-point interface and with router R3 over the serial0/0.3 point-to-point subinterface. Likewise, routers R2 and R3 form the adjacency with router R1. On the figure above it is apparent that router R3 has one neighbor over the serial0/0.1 point-to-point subinterface.

Load Balancing Across Equal-Metric Paths

This topic explains how EIGRP performs load balancing across equal-metric paths and describes how to change the default configuration.

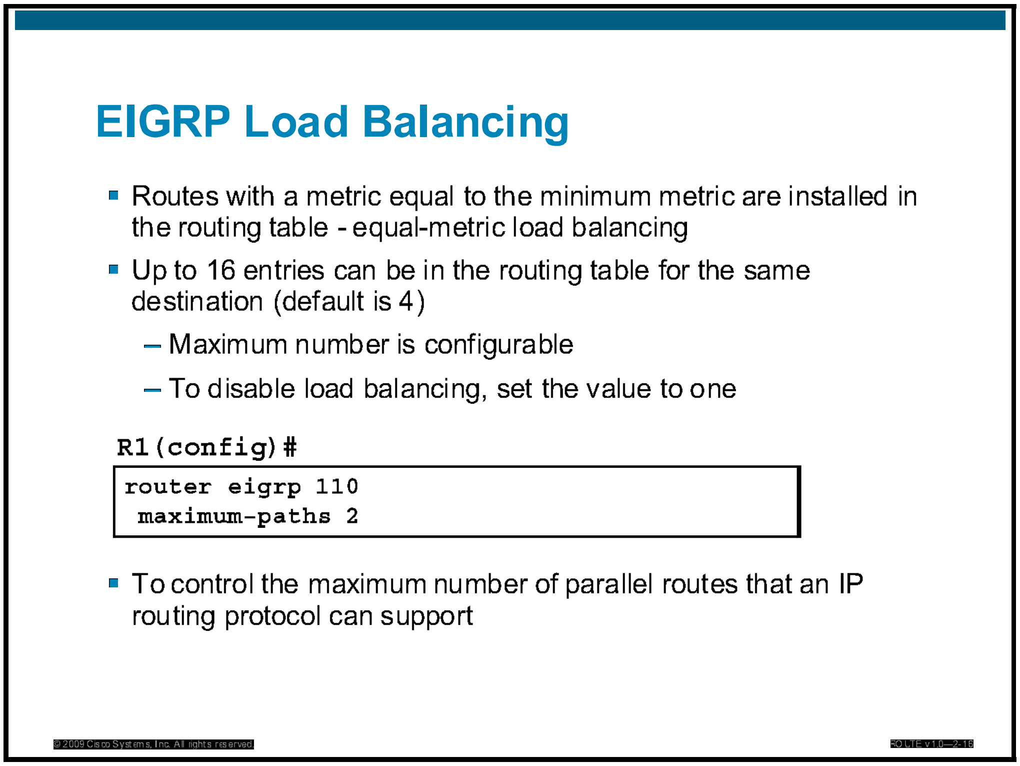

EIGRP Load Balancing

- Routes with a metric equal to the minimum metric are installed in the routing table - equal-metric load balancing
- Up to 16 entries can be in the routing table for the same destination (default is 4)
  - Maximum number is configurable
  - To disable load balancing, set the value to one

R1(config)#

```
router eigrp 110
 maximum-paths 2
```

- To control the maximum number of parallel routes that an IP routing protocol can support

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-16

Equal-metric load balancing is a router's capability to distribute traffic over all of its network ports that have the same metric to the destination address. Load balancing increases the use of network segments and increases the effective network bandwidth.

For IP, the Cisco IOS Software applies load balancing between a maximum of four equal-metric paths by default. You can configure the maximum number of parallel routes that an IP routing protocol can support using the **maximum-paths** router configuration command. Up to 16 equally good routes can be kept in the routing table.

Note Setting the **maximum-paths** value to 1 disables load balancing.

When a packet is process-switched, load balancing over equal-metric paths occurs on a per-packet basis. When packets are fast-switched, load balancing over equal-metric paths occurs on a per-destination basis. (Thus, if you are testing load balancing, do not ping to or from routers with fast-switching interfaces, because the packets generated locally by this router are process-switched rather than fast-switched, and the ping might produce confusing results.)

For more details about the **maximum-paths** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

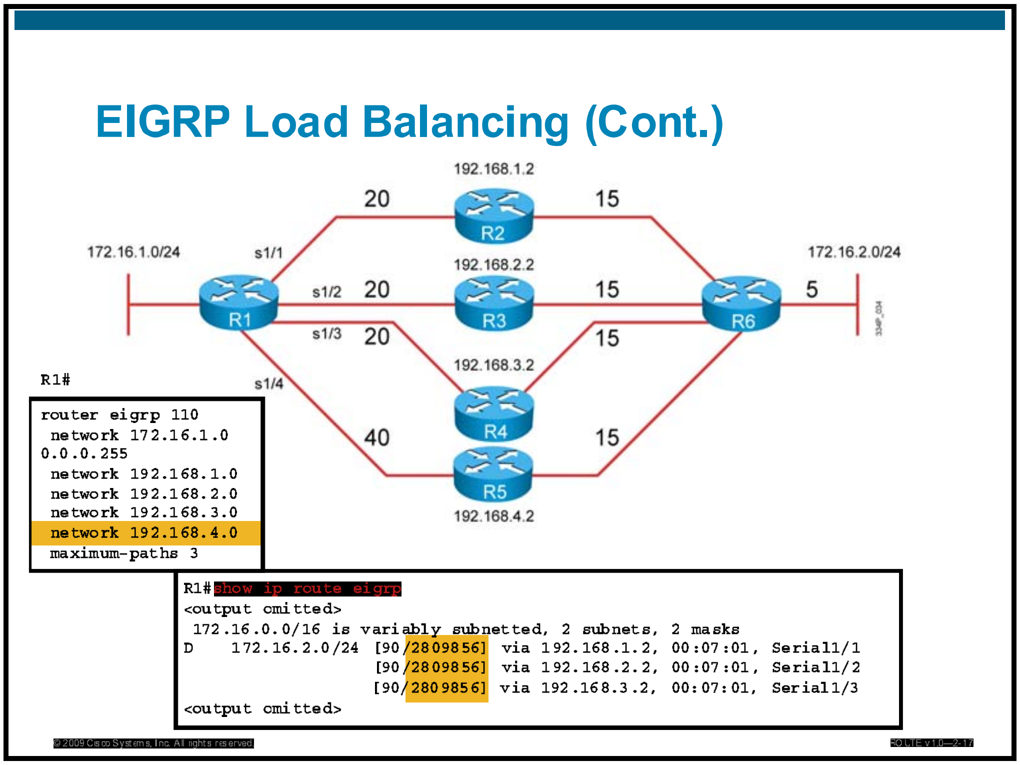

EIGRP Load Balancing (Cont.)

```
R1#
router eigrp 110
 network 172.16.1.0
0.0.0.255
 network 192.168.1.0
 network 192.168.2.0
 network 192.168.3.0
 network 192.168.4.0
 maximum-paths 3
```

```
R1#show ip route eigrp
<output omitted>
 172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.2.0/24  [90/2809856]  via 192.168.1.2, 00:07:01, Serial1/1
                    [90/2809856]  via 192.168.2.2, 00:07:01, Serial1/2
                    [90/2809856]  via 192.168.3.2, 00:07:01, Serial1/3
<output omitted>
```

ROUTE v1.0—2-17

The configuration example above shows the use of the **maximum-paths** command, which is applied under the EIGRP routing process. Router R1 is configured to support up to three equal-metric paths. If the **maximum-paths** command is not used, EIGRP can, by default, use four equal-metric paths.

The sample output above shows router R1's routing table, where the three paths have the same metric (cost). All three paths through routers R2, R3, and R4 are used to reach the destination network 172.16.2.0 behind router R6. The path through router R5 is not used, because the metric is too big. Even if the metric is the same as the others, the path will not be used, as 3 is the maximum number set by the **maximum-paths** command.

Load Balancing Across Unequal-Metric Paths

This topic explains how EIGRP performs load balancing across unequal-metric paths and describes how to configure load balancing.

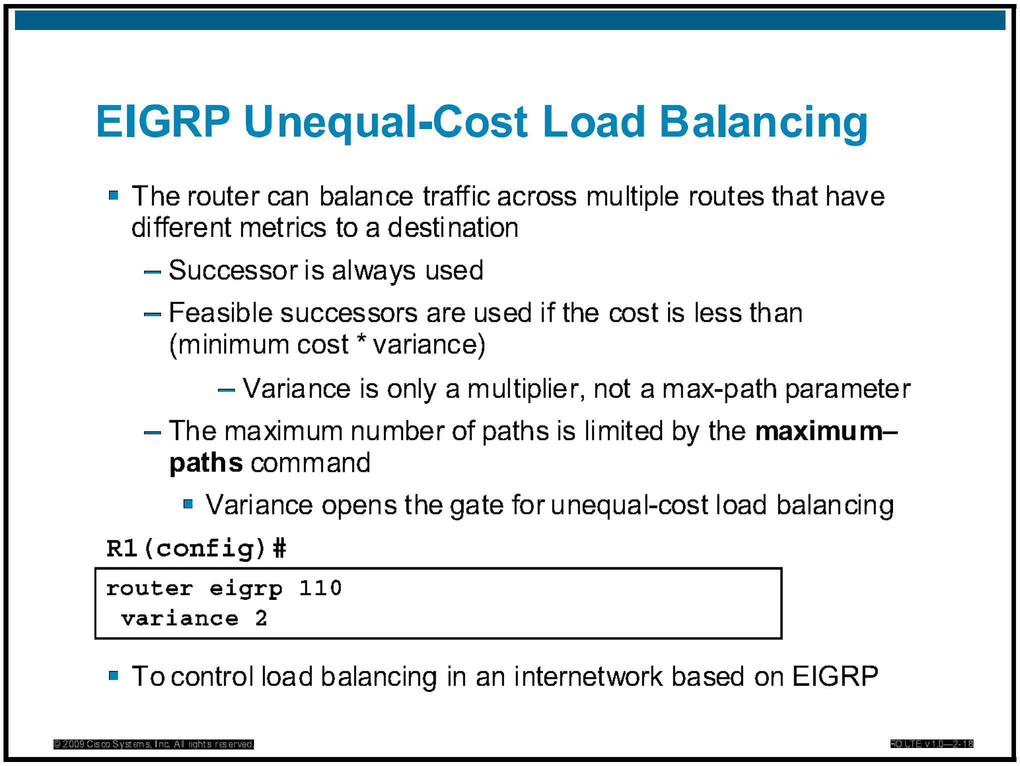

EIGRP Unequal-Cost Load Balancing

- The router can balance traffic across multiple routes that have different metrics to a destination
  - Successor is always used
  - Feasible successors are used if the cost is less than (minimum cost * variance)
    - Variance is only a multiplier, not a max-path parameter
  - The maximum number of paths is limited by the **maximum–paths** command
    - Variance opens the gate for unequal-cost load balancing

R1(config)#

```
router eigrp 110
 variance 2
```

- To control load balancing in an internetwork based on EIGRP

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-18

EIGRP can also balance traffic across multiple routes that have different metrics, which is called unequal-metric load balancing. The degree to which EIGRP performs load balancing is controlled by the **variance** (EIGRP) command. Setting a variance value (1-128) enables EIGRP to install multiple loop-free routes with unequal cost in a local routing table. EIGRP will always install a successor into the local routing table. Additional feasible successors are candidates for the local routing table. Additional entries through EIGRP must meet two criteria to be installed in the local routing table:

- The route must be loop-free. This condition is satisfied when the advertised distance (AD) is less than the total distance, or when the route is a feasible successor.
- The metric of the route must be lower than the metric of the best route (the successor), multiplied by the variance configured on the router.

The default value for the **variance** (EIGRP) command is 1, which indicates equal-cost load balancing—only routes with the same metric as the successor are installed in the local routing table. The variance command is not limiting the maximum number of paths. It is the multiplier that defines the range of metric values that are accepted for load balancing by the EIGRP process. If the variance is set to 2, any EIGRP-learned route with a metric less than 2 times the successor metric will be installed in the local routing table. If the **variance** command allows EIGRP to use 9 paths and the **maximum-path** command sets the maximum paths value to 3, only the first 3 paths of the 9 paths will show up in the IP routing table—the **maximum-path** command will limit the maximum number.

Note EIGRP does not load-share between multiple routes; it only installs the routes in the local routing table. Then the local routing table enables switching hardware or software to load-share between the multiple paths.

For more details about the EIGRP **variance** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

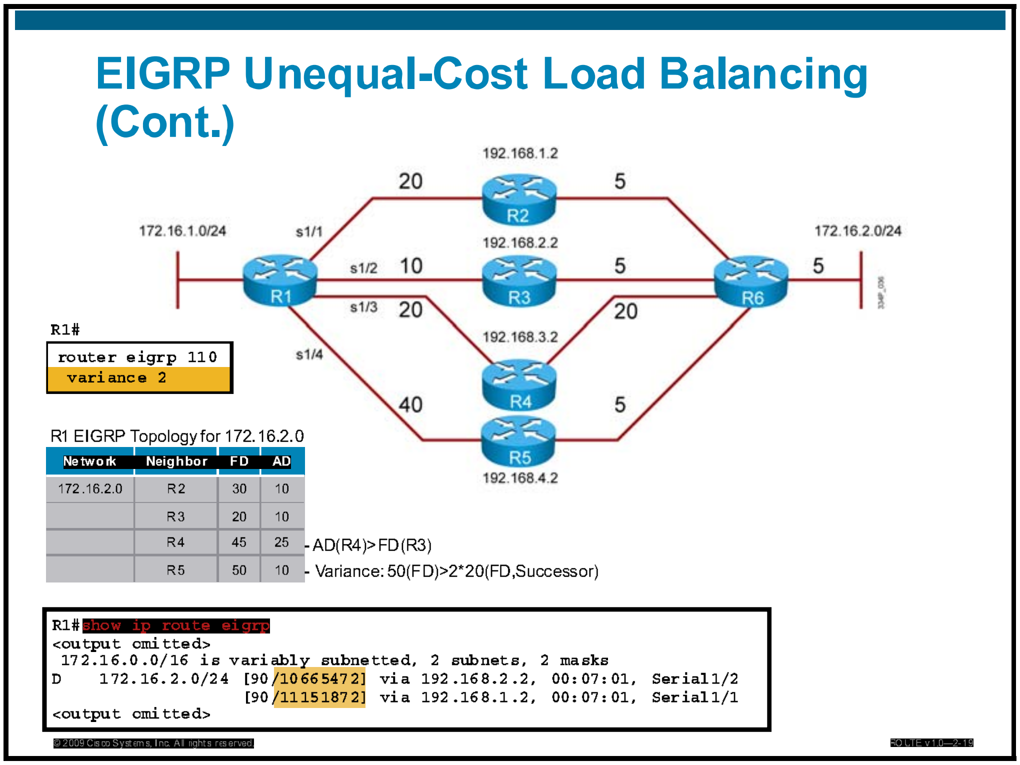

EIGRP Unequal-Cost Load Balancing (Cont.)

```
R1#
router eigrp 110
 variance 2
```

R1 EIGRP Topology for 172.16.2.0

Network	Neighbor	FD	AD
172.16.2.0	R2	30	10
	R3	20	10
	R4	45	25
	R5	50	10

- AD(R4)>FD(R3)
- Variance: 50(FD)>2*20(FD,Successor)

```
R1#show ip route eigrp
<output omitted>
 172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D    172.16.2.0/24 [90/10665472] via 192.168.2.2, 00:07:01, Serial1/2
                   [90/11151872] via 192.168.1.2, 00:07:01, Serial1/1
<output omitted>
```

Router R1 in the example above is configured with a variance of 2, and the range of metric values, which are the feasible distances (FDs) to network 172.16.2.0/24, is from 20 to 50 (generic values, in order to make computation easier). This range of values determines the feasibility of a potential route.

A route is feasible if the next router in the path is closer to the destination than the current router and if the metric of the alternate path is within the variance. Load balancing can use only feasible paths, which are included inside the local routing table. The two feasibility conditions are:

- The local best metric (the current FD) must be greater than the best metric (AD) learned from the next router. In other words, the next router in the path must be closer to the destination than the current router, which prevents routing loops.
- The variance multiplied by the local best metric (the current FD) must be greater than the metric through the next router (the alternative FD). This condition is true if the metric of the alternate path is within the variance.

If both of these conditions are met, the route is called feasible and can be added to the routing table.

The example below shows four paths from router R1 to network 172.16.2.0/24 with the following metrics:

- Path 1: 30 (via R2)
- Path 2: 20 (via R3)
- Path 3: 45 (via R4)
- Path 4: 50 (via R5)

All the paths have different metrics. By default, router R1 only places path 2, via R3, in the routing table, because it is the least-cost path—a successor route with the lowest FD. For it to load balance over paths 1 and 2, a **variance** command must be applied to router R1 to change the value in such way that path 1 (with a greater metric than the successor route) can be used as well. The variance value must be set such that the result of multiplying the variance and successor FD will be greater than the second FD candidate. In our example, the variance value is set to 2, which produces a result of 40 (20 * 2 = 40) and path 1 through router R2, becomes the second route in the local routing table (the FD of 30 is less than 40).

Router R4 is not considered for load balancing with this variance, because the FD through router R4 is more than twice the FD for the successor (router R3). Router R4 will never be a feasible successor no matter what the variance is. This is because router R4's advertised distance of 25 is greater than router R3's FD of 20; therefore, to avoid a potential routing loop, router R4 is not considered a feasible successor.

Router R5 is not considered for load balancing with this variance, because the FD through router R5 is more than twice the FD for the successor (router R3). In this example, however, router R5 will be a feasible successor no matter what the variance is. This is because router R5's advertised distance of 10 is lower than router R3's FD of 20.

Load balancing is proportional to the bandwidth. Routes via R2 and R3 in the picture above are used for load balancing. The FD of the route via R2 equals 30. The FD of the route via R3 equals 20. The total FD is 50 and the ratio of traffic between the two paths (via R2 : via R3) is 3/5 : 2/5.

EIGRP Bandwidth use Across WAN Links

This topic explains how EIGRP can utilize bandwidth across WAN links in an environment where the default bandwidth usage, of up to 50%, may not be optimal.

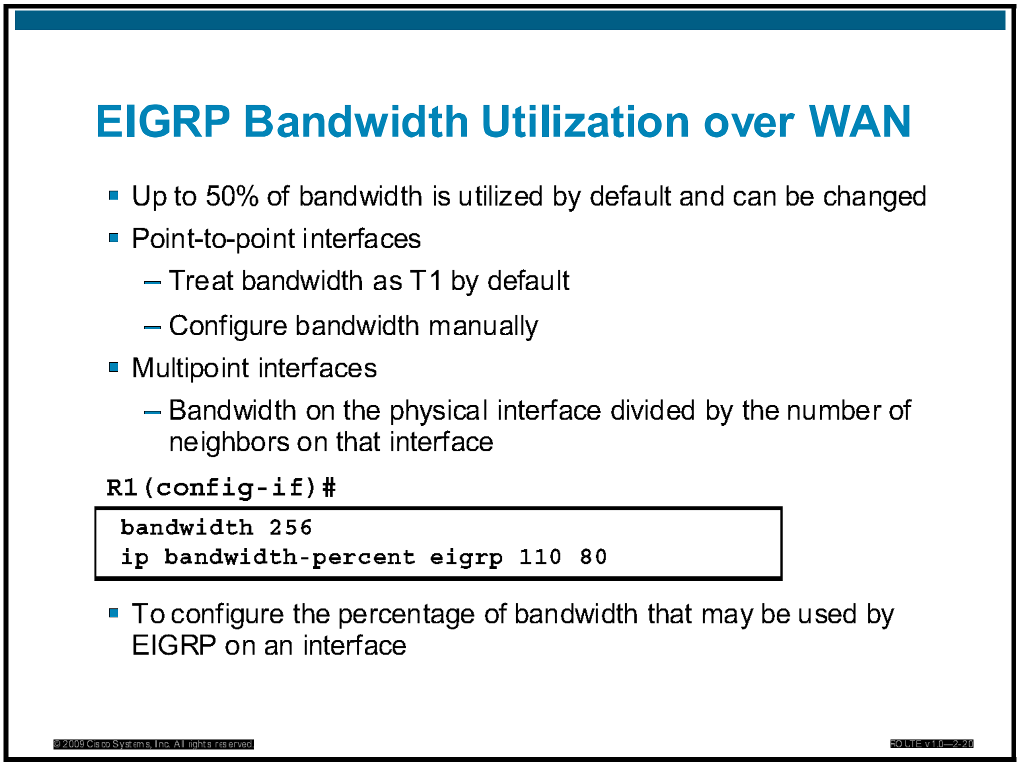

EIGRP Bandwidth Utilization over WAN

- Up to 50% of bandwidth is utilized by default and can be changed
- Point-to-point interfaces
  - Treat bandwidth as T1 by default
  - Configure bandwidth manually
- Multipoint interfaces
  - Bandwidth on the physical interface divided by the number of neighbors on that interface

R1(config-if)#

```
bandwidth 256
ip bandwidth-percent eigrp 110 80
```

- To configure the percentage of bandwidth that may be used by EIGRP on an interface

EIGRP operates efficiently in WAN environments. It is scalable on both point-to-point links and on multipoint NBMA links.

Because of the inherent differences in the operational characteristics of WAN links, the default configuration parameters may not be the best option for all WAN links. A solid understanding of EIGRP operation, coupled with knowledge of available link speeds, can yield an efficient, reliable, and scalable router configuration.

By default, EIGRP may use up to 50 percent of the bandwidth of an interface or subinterface. When calculating how much bandwidth to use, EIGRP uses either the bandwidth of the link set by the **bandwidth** command, or the link's default bandwidth if none is configured. This percentage can be changed on a per-interface basis by using the **ip bandwidth-percent eigrp** interface configuration command.

Cisco IOS Software assumes that the point-to-point Frame Relay subinterfaces (such as all serial interfaces) operate at full T1 link speed. In many implementations however, only fractional T1 speeds are available. Therefore, when configuring these subinterfaces, set the bandwidth to match the contracted committed information rate (CIR).

When configuring multipoint interfaces, especially for Frame Relay (but also for ATM and ISDN PRI), it is important to understand that all neighbors share the bandwidth equally. That is, EIGRP uses the bandwidth command on the physical interface divided by the number of Frame Relay neighbors connected on that physical interface to calculate the bandwidth assigned to each neighbor. The EIGRP configuration should reflect the correct percentage of the actual available bandwidth on the line.

In the figure above, a sample configuration is used, in which an EIGRP process with an AS number 110 can get 80% of the bandwidth configured on that interface. This percentage can be greater than 100, which may be useful if the bandwidth is configured artificially low for routing-policy reasons.

For more details about the **ip bandwidth-percent eigrp** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

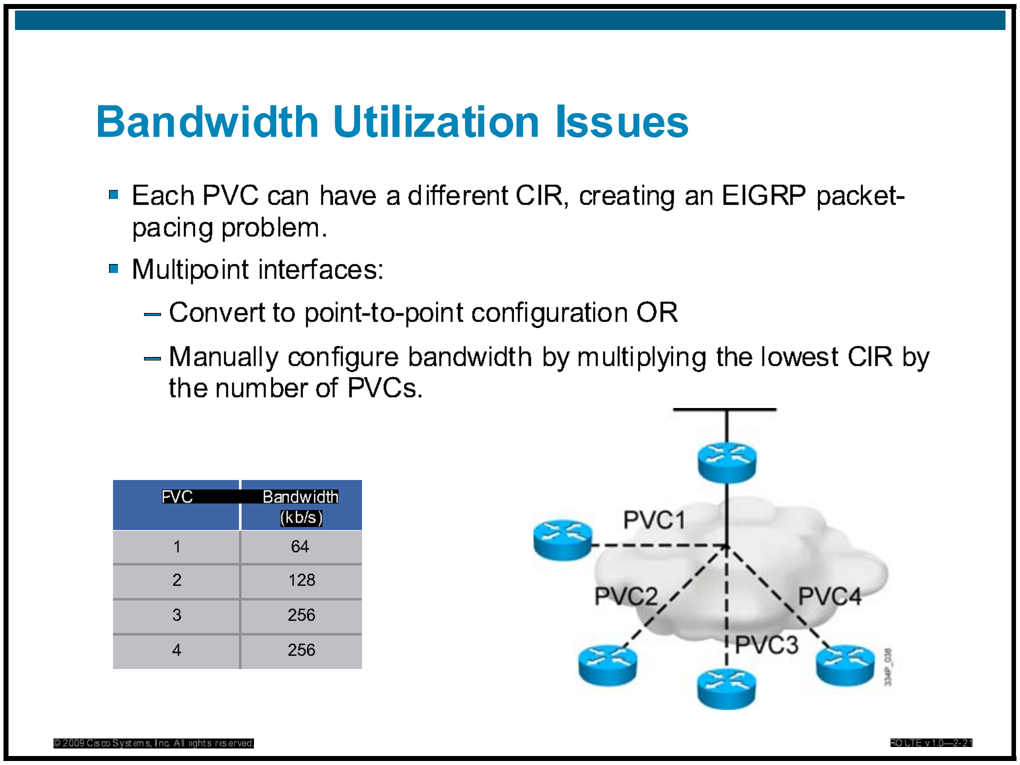

Bandwidth Utilization Issues

- Each PVC can have a different CIR, creating an EIGRP packet-pacing problem.
- Multipoint interfaces:
  - Convert to point-to-point configuration OR
  - Manually configure bandwidth by multiplying the lowest CIR by the number of PVCs.

PVC	Bandwidth (kb/s)
1	64
2	128
3	256
4	256

Each installation has a unique topology and requires a unique configuration. Since CIR values differ, a subinterface will often requires a hybrid configuration that blends the characteristics of point-to-point circuits with multipoint circuits.

When configuring multipoint interfaces, configure the bandwidth to represent the minimum CIR multiplied by the number of circuits. This approach may not fully use the higher-speed circuits, but it ensures that the circuits with the lowest CIR values are not overdriven. If the topology has a small number of very low-speed circuits, these interfaces are defined typically as point-to-point, so that their bandwidth can be set to match the provisioned CIR.

In the figure above, 64k b/s is the smallest amount of bandwidth among all PVCs in the multipoint configuration and this should be taken into account for bandwidth calculation.

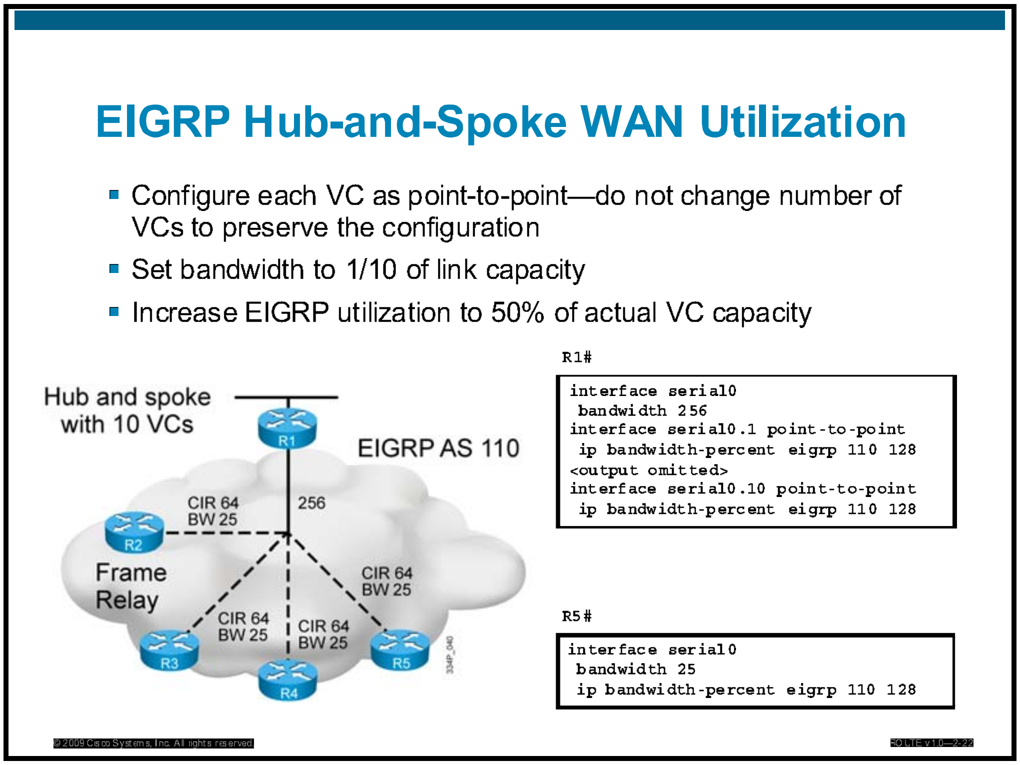

EIGRP Hub-and-Spoke WAN Utilization

- Configure each VC as point-to-point—do not change number of VCs to preserve the configuration
- Set bandwidth to 1/10 of link capacity
- Increase EIGRP utilization to 50% of actual VC capacity

R1#

```
interface serial0
 bandwidth 256
interface serial0.1 point-to-point
 ip bandwidth-percent eigrp 110 128
<output omitted>
interface serial0.10 point-to-point
 ip bandwidth-percent eigrp 110 128
```

R5#

```
interface serial0
 bandwidth 25
 ip bandwidth-percent eigrp 110 128
```

ROUTE v1.0—2-22

The example above shows a hub-and-spoke topology with ten virtual circuits to the ten remotes sites. Only four of the ten remote sites are shown in the figure above. The configurations for routers R1 and R5, using EIGRP AS 110, are also shown in the figure.

The circuits are provisioned as 64-kb/s links, but there is insufficient bandwidth at the interface to support this allocation. For example, if the hub tries to communicate to all remote sites at the same time, the bandwidth that is required exceeds the available link speed of 256 kb/s for the hub (10 times the CIR of 64 kb/s equals 640 kb/s).

The configuration for router R1 shows that the **bandwidth 256** command has been applied to the main interface (serial0) and has set the bandwidth to 256 kb/s. Because 10 virtual circuits (VCs) are configured, each of them automatically gets 10% of the main interface bandwidth. Therefore, the **bandwidth** command is not needed on each subinterface.

The **ip bandwidth-percent eigrp 110 128** command sets the maximum percentage of an interface's bandwidth that EIGRP can use to 128% for EIGRP AS 110. In a point-to-point topology, all virtual circuits are treated equally; the subinterfaces are assigned a bandwidth equal to one-tenth of the available link speed (25 kb/s). Based on the **ip bandwidth-percent eigrp 110 128** command applied on each interface and subinterface, the EIGRP allocation percentage is raised to 128 percent of the specified bandwidth in an attempt to ensure that the EIGRP packets are delivered through the Frame Relay network. This adjustment causes the EIGRP packets to receive 32 kb/s of the provisioned 64 kb/s on each circuit (128% of the 25kb/s equals 32 kb/s). This extra configuration restores the 50-50 ratio that was changed when the bandwidth was set to an artificially low value. In order to ensure that this calculation is correct, the number of VCs must not be changed.

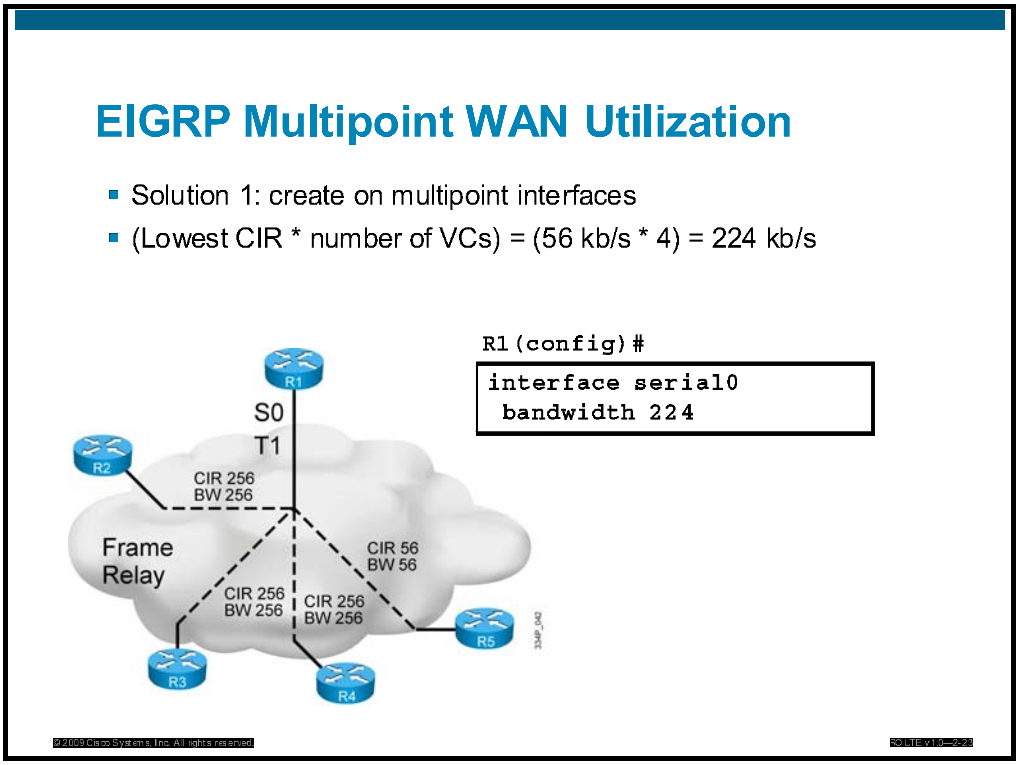

EIGRP Multipoint WAN Utilization

- Solution 1: create on multipoint interfaces
- (Lowest CIR * number of VCs) = (56 kb/s * 4) = 224 kb/s

R1(config)#

```
interface serial0
 bandwidth 224
```

ROUTE v1.0—2-23

In a multipoint topology, where virtual circuits may not have equal bandwidth, the rule is to use the lowest CIR and multiply it by the number of VCs to get the bandwidth that should be set on the interface. With this configuration, the smallest and the slowest link will not suffer with higher speed links.

In the figure above, the example shows four VCs with different CIRs configured. The CIR toward routers R2, R3, and R4 is configured to 256 kb/s and the CIR toward router R5 is configured to 56 kb/s. Thus, on router R1, the bandwidth is set to 224 kb/s, which is 4 times the lowest CIR in the system (56 kb/s).

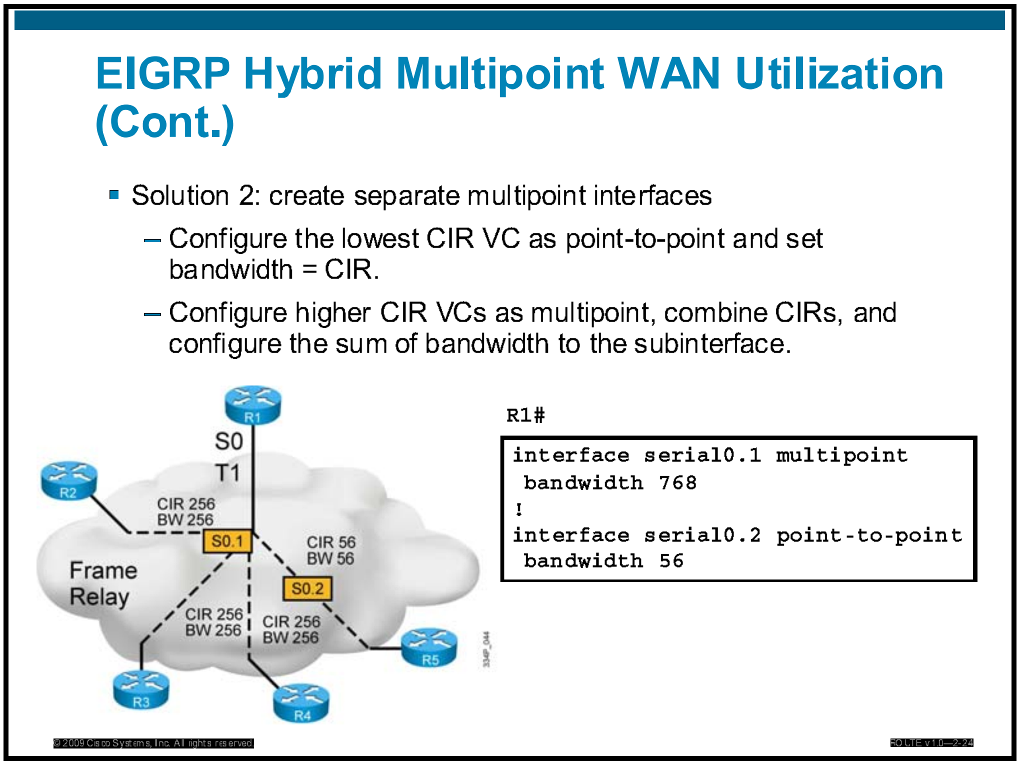

EIGRP Hybrid Multipoint WAN Utilization (Cont.)

- Solution 2: create separate multipoint interfaces
  - Configure the lowest CIR VC as point-to-point and set bandwidth = CIR.
  - Configure higher CIR VCs as multipoint, combine CIRs, and configure the sum of bandwidth to the subinterface.

```
R1#
interface serial0.1 multipoint
 bandwidth 768
!
interface serial0.2 point-to-point
 bandwidth 56
```

ROUTE v1.0—2-24

In a hybrid multipoint and point-to-point topology, you should create separate point-to-point VCs for the lowest CIR VCs. All other VCs, which have a higher and possibly equal CIR, should be configured as one multipoint VC. For this multipoint VC, the CIR must be a combination of all individual CIRs.

The figure above shows an example in which one VC is a low-speed circuit with 56kb/s of bandwidth and the other three VCs have bandwidth of 256 kb/s. The preferred configuration on router R1 shows the low-speed circuit configured as point-to-point with the bandwidth set to the CIR value. The remaining circuits are designated as a multipoint subinterface, and their CIRs are added together to set the bandwidth for the subinterface.

In multipoint interfaces, the bandwidth is shared equally among all circuits. In this case, the bandwidth is set to 768 kb/s, which is the sum of the three CIRs (3 * 256 kb/s = 768 kb/s). Each link will be allocated one-third of this bandwidth, resulting in 256 kb/s each.

EIGRP over EoMPLS and Metro Ethernet

This topic provides an overview of EoMPLS technology and explains how EIGRP can be deployed in an EoMPLS environment.

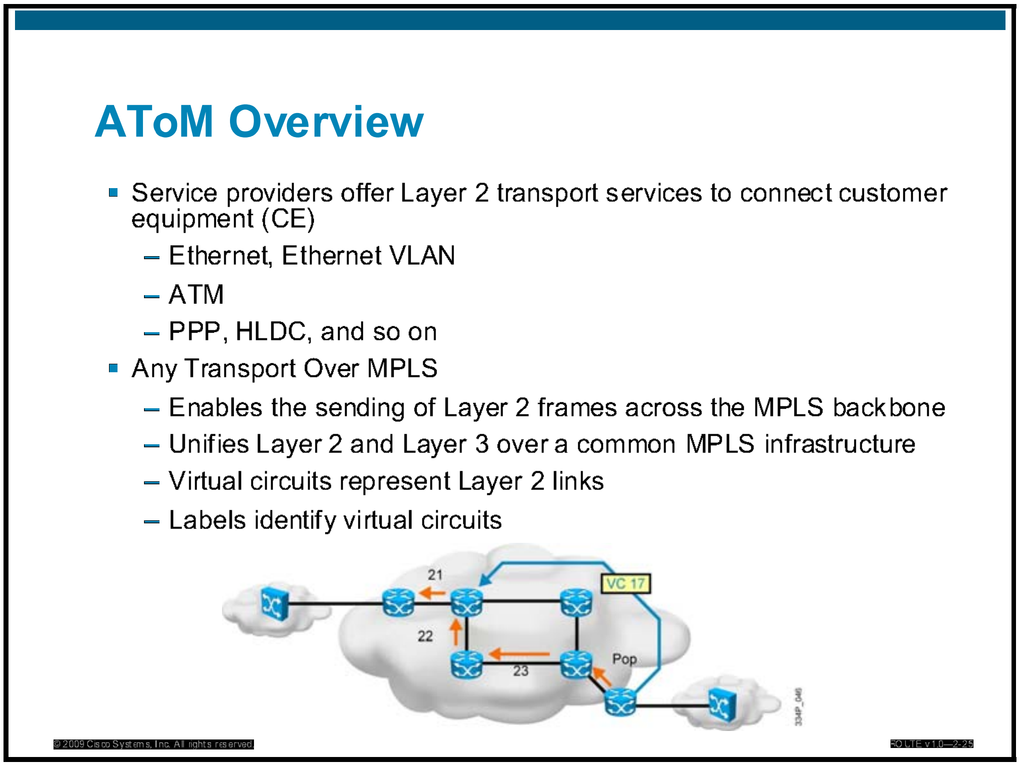

AToM Overview

- Service providers offer Layer 2 transport services to connect customer equipment (CE)
  - Ethernet, Ethernet VLAN
  - ATM
  - PPP, HLDC, and so on
- Any Transport Over MPLS
  - Enables the sending of Layer 2 frames across the MPLS backbone
  - Unifies Layer 2 and Layer 3 over a common MPLS infrastructure
  - Virtual circuits represent Layer 2 links
  - Labels identify virtual circuits

ROUTE v1.0—2-25

Many ISPs currently offer Layer 2 transport services to their customers. These services are offered over a circuit-based infrastructure to build Layer 2 virtual private networks (VPNs).

Initially, VPNs were built using leased lines. Later, service providers offered Layer 2 VPNs based on point-to-point data link layer connectivity, using ATM or Frame Relay virtual circuits. Customers built their own Layer 3 networks to accommodate IP traffic. As a result, separate networks exist for Layer 2 and Layer 3 traffic. However, maintaining separate networks for Layer 2 VPNs and Internet traffic is difficult and costly. Therefore, ISPs want a single IP-based network to provide both Layer 2 and Layer 3 services.

MPLS VPN was introduced to meet the requirement for a unified network for Layer 3 VPN services. However, some customers still wanted Layer 2 connections. This can be Ethernet VLAN extensions across a metropolitan area or ATM services. Any Transport over MPLS (AToM) was introduced to facilitate Layer 2 connectivity across an MPLS backbone.

AToM benefits ISPs that offer Layer 2 connectivity to customers with traditional offerings such as ATM, Frame Relay, and serial/Point-to-Point Protocol (PPP) services. Additionally, it benefits service providers specializing in Ethernet connectivity in metropolitan areas. Services for Layer 2 VPNs also appeal to service providers' enterprise customers, who may already run many of these networks and want just point-to-point connectivity.

Note	For more information about MPLS, please attend the *Implementing MPLS VPN* Cisco course.

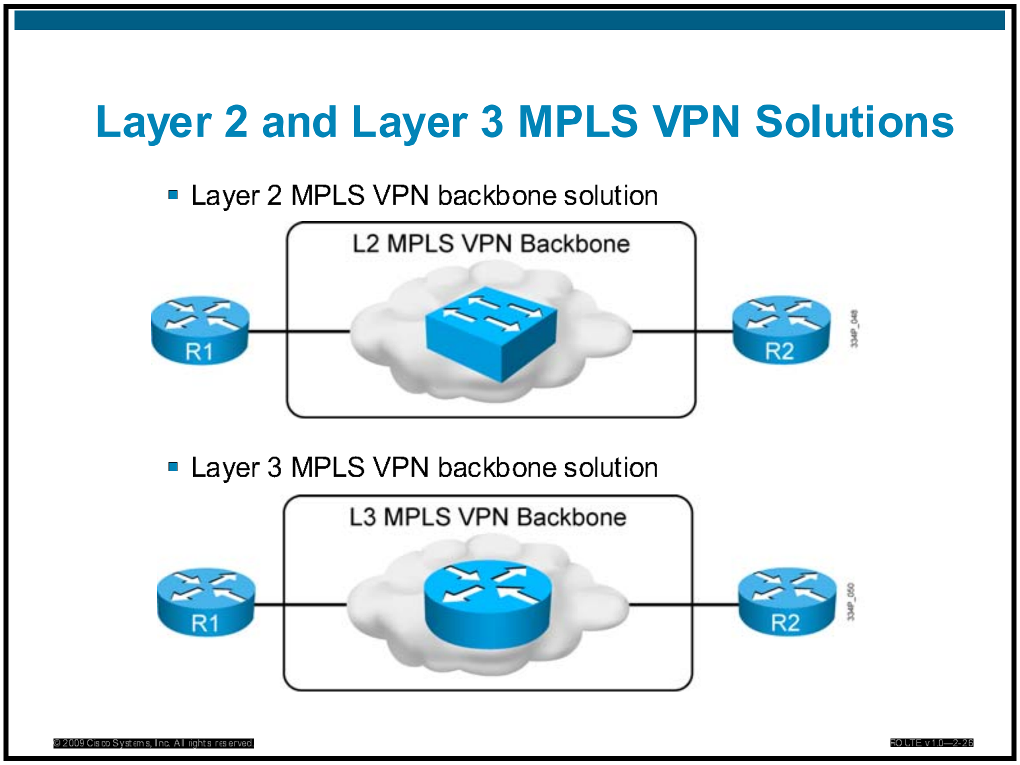

Layer 2 and Layer 3 MPLS VPN Solutions

- Layer 2 MPLS VPN backbone solution

- Layer 3 MPLS VPN backbone solution

ROUTE v1.0—2-26

The figure above presents the basic difference between a Layer 2 MPLS VPN and a Layer 3 MPLS VPN backbone solution. Customer routers (R1 and R2 on this slide) are connected across the MPLS VPN backbone, and it is important to define the difference.

The Layer 2 MPLS VPN backbone solution is providing the Layer 2 service across the backbone, where routers R1 and R2 are connected together directly using the same IP subnet. This slide presents the connectivity through the backbone as a switch.

The Layer 3 MPLS VPN backbone solution is providing the Layer 3 service across the backbone, where routers R1 and R2 are connected to ISP edge routers. A separate IP subnet is used on each side. This slide presents the connectivity through the backbone as a router.

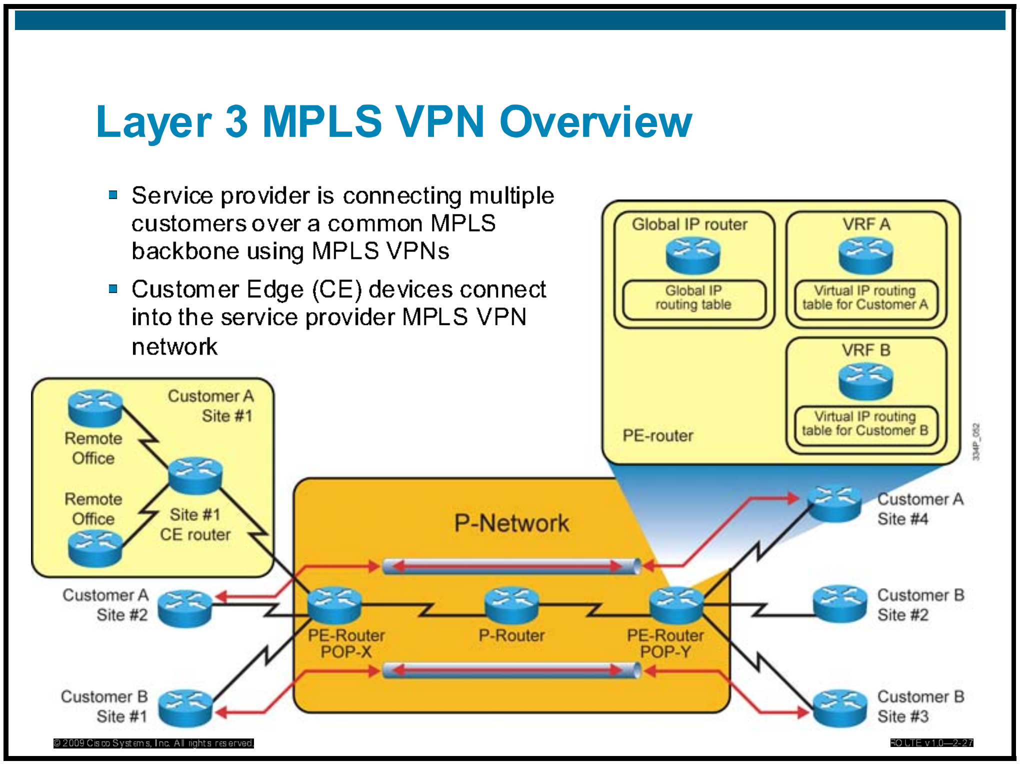

Layer 3 MPLS VPN Overview

- Service provider is connecting multiple customers over a common MPLS backbone using MPLS VPNs
- Customer Edge (CE) devices connect into the service provider MPLS VPN network

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-27

The MPLS VPN architecture provides the ISPs with a peer-to-peer VPN architecture that combines the best feature of an overlay VPN (support for overlapping customer address spaces) with the best features of peer-to-peer VPNs:

- PE routers participate in customer routing, guaranteeing optimum routing between customer sites.
- PE routers carry a separate set of routes for each customer, resulting in perfect isolation between the customers.

The MPLS VPN terminology divides the overall network into the customer-controlled part (C-network) and the provider-controlled part (P-network). Contiguous portions of the C-network are called sites and are linked with the P-network via Customer Edge routers (CE-routers). The CE-routers are connected to the provider edge routers (PE-routers), which serve as the edge devices of the provider network. The core devices in the provider network (P-routers) provide the transit transport across the provider backbone and do not carry customer routes.

The architecture of a PE-router in an MPLS VPN is very similar to the architecture of a point of presence (POP) in the dedicated PE-router peer-to-peer model; the only difference is that the whole architecture is condensed into one physical device. Each customer is assigned an independent routing table (virtual routing table) that corresponds to the dedicated PE-router in the traditional peer-to-peer model. Routing across the provider backbone is performed by another routing process that uses global IP routing table, corresponding to the intra-POP P-router in a traditional peer-to-peer model.

The MPLS VPN backbone provides a Layer 3 backbone in which the CE routers see PE routers as additional customer routers in the path. The PE routers maintain separate routing tables for each customer.

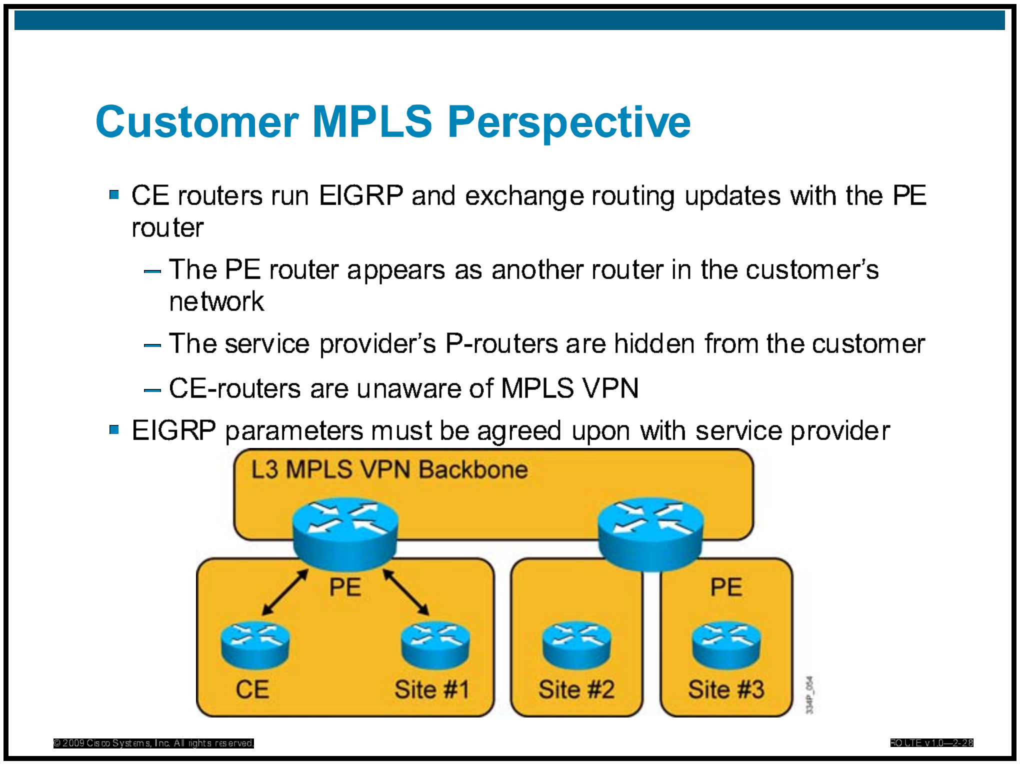

Customer MPLS Perspective

- CE routers run EIGRP and exchange routing updates with the PE router
  - The PE router appears as another router in the customer's network
  - The service provider's P-routers are hidden from the customer
  - CE-routers are unaware of MPLS VPN
- EIGRP parameters must be agreed upon with service provider

L3 MPLS VPN Backbone
PE
PE
CE
Site #1
Site #2
Site #3
334P_054

ROUTE v1.0—2-28

The MPLS VPN technology has the following routing requirements:

- The customer routers should not be aware of MPLS VPN. They should run standard IP routing software.
- The provider core routers (P-routers) must not carry VPN routes, in order to ensure the MPLS VPN solution is scalable.
- The provider edge routers (PE-routers) must support MPLS VPN services and traditional Internet services.

The MPLS VPN backbone looks like a standard corporate backbone to the CE-routers. The CE-routers run standard IP routing software and exchange routing updates with the PE-routers that appear to them as normal routers in customer's network. The standard design rules that are used for enterprise MPLS VPN backbones can be applied to the design of the customer's network. The P-routers are hidden from the customer's view and CE-routers are unaware of the MPLS VPN. The internal topology of the MPLS backbone is therefore transparent to the customer.

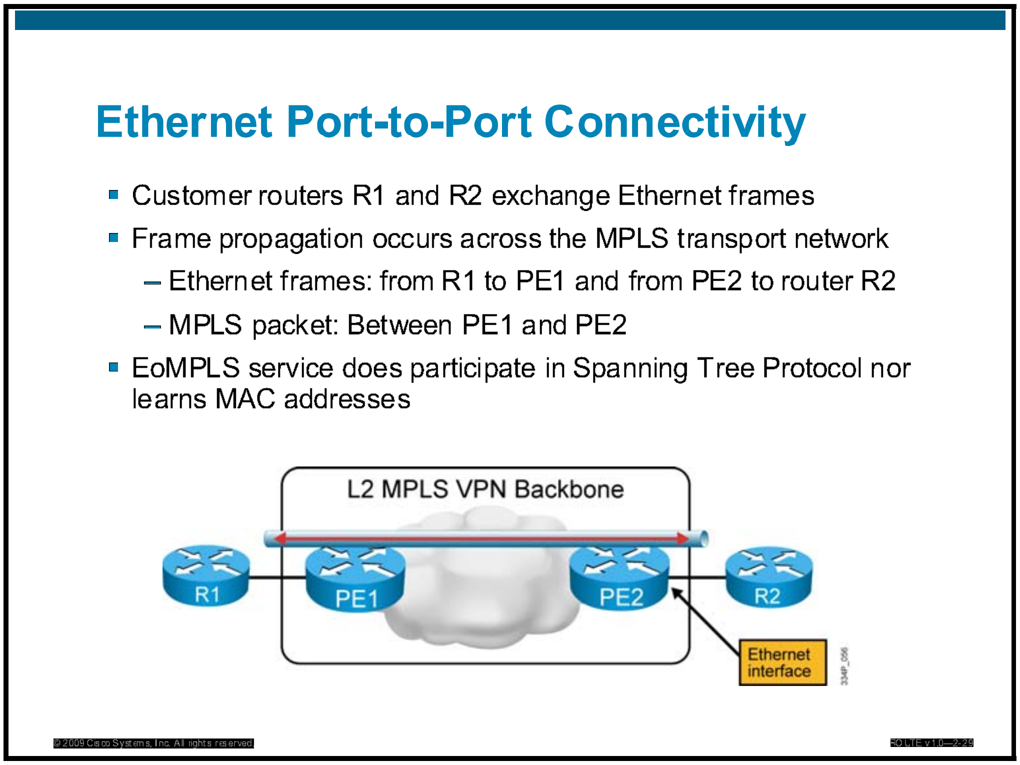

Ethernet Port-to-Port Connectivity

- Customer routers R1 and R2 exchange Ethernet frames
- Frame propagation occurs across the MPLS transport network
  - Ethernet frames: from R1 to PE1 and from PE2 to router R2
  - MPLS packet: Between PE1 and PE2
- EoMPLS service does participate in Spanning Tree Protocol nor learns MAC addresses

ROUTE v1.0—2-29

An MPLS backbone provides a Layer 2 Ethernet port-to-port connection between the two customer routers R1 and R2.

Routers R1 and R2 are exchanging Ethernet frames. The Provider Edge 1 (PE1) router takes whatever Ethernet frame it receives from router R1 on the link to PE1, encapsulates it into an MPLS packet, and forwards it across the backbone to the PE2 router. PE2 de-encapsulates the MPLS packet and reproduces the Ethernet frame on its link towards router R2.

The AToM feature does not include any MAC layer address learning and filtering. That means that routers PE1 and PE2 do not filter any frames based on those addresses.

Nor does the AToM feature use STP (Spanning Tree Protocol). Bridge protocol data units (BPDUs) are propagated transparently and not processed. The LAN loop detection must be performed by other functions or avoided by design.

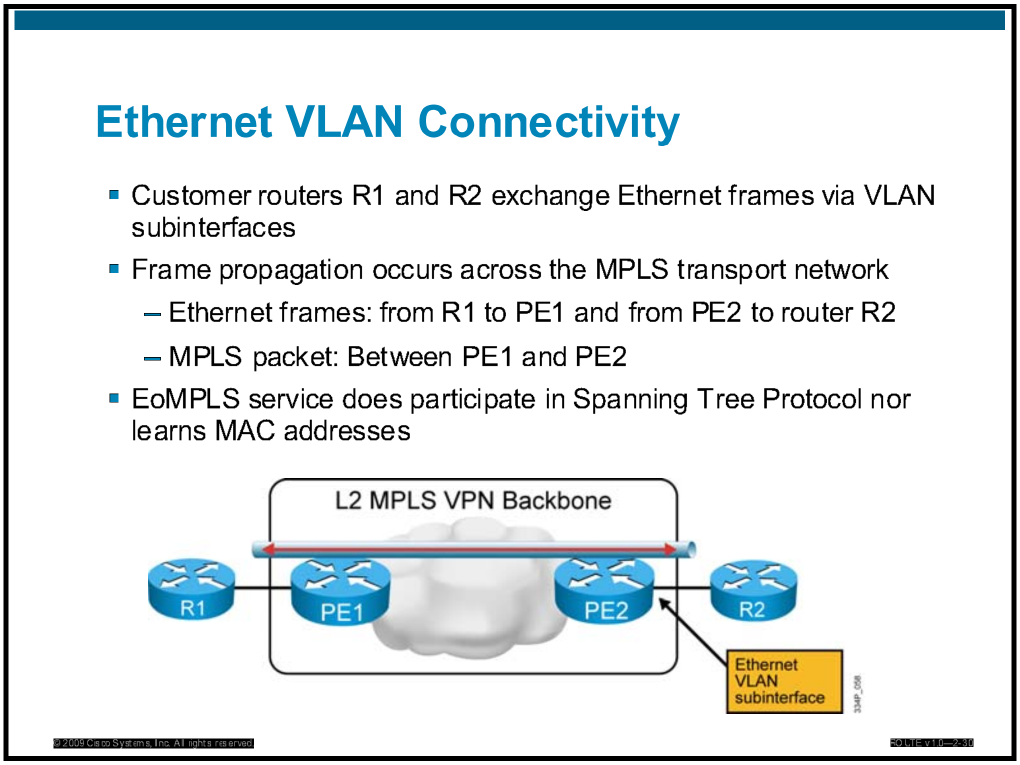

Ethernet VLAN Connectivity

- Customer routers R1 and R2 exchange Ethernet frames via VLAN subinterfaces
- Frame propagation occurs across the MPLS transport network
  - Ethernet frames: from R1 to PE1 and from PE2 to router R2
  - MPLS packet: Between PE1 and PE2
- EoMPLS service does participate in Spanning Tree Protocol nor learns MAC addresses

ROUTE v1.0—2-30

The two customer routers R1 and R2 are connected to the MPLS edge routers PE1 and PE2 via VLAN subinterfaces.

The interface encapsulation between routers R1 and R2 and the PE routers supports VLANs. Different subinterfaces in the PE routers are used to connect to different VLANs. The PE1 subinterface to the VLAN where router R1 is connected is used for AToM forwarding. When an Ethernet frame arrives on the specific VLAN subinterface, it is encapsulated into MPLS and forwarded across the backbone to router PE2. Router PE2 de-encapsulates the packet and reproduces the Ethernet frame on the outgoing subinterface toward router R2.

Although the learning of MAC addresses and STP are not features of AToM, combining AToM with a LAN switch allows the service provider to utilize those missing features.

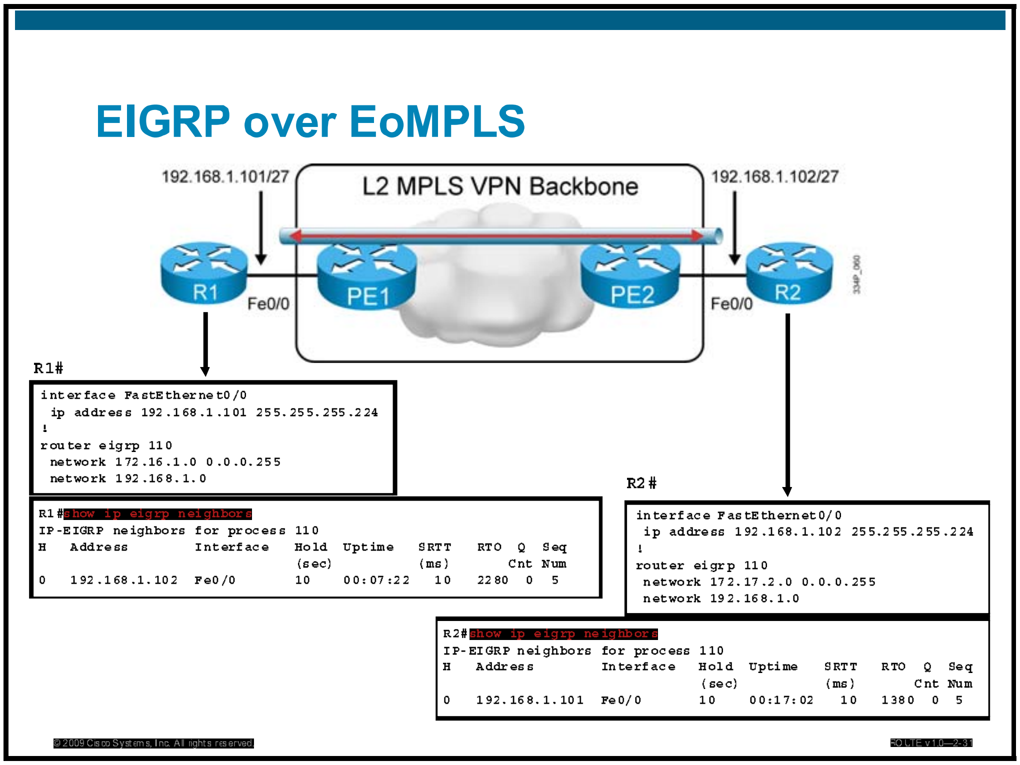

EIGRP over EoMPLS

R1#

```
interface FastEthernet0/0
 ip address 192.168.1.101 255.255.255.224
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q  Seq
                                  (sec)           (ms)        Cnt Num
0   192.168.1.102   Fe0/0         10    00:07:22   10    2280  0   5
```

R2#

```
interface FastEthernet0/0
 ip address 192.168.1.102 255.255.255.224
!
router eigrp 110
 network 172.17.2.0 0.0.0.255
 network 192.168.1.0
```

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q  Seq
                                  (sec)           (ms)        Cnt Num
0   192.168.1.101   Fe0/0         10    00:17:02   10    1380  0   5
```

ROUTE v1.0—2-31

In the example in the slide, it is assumed that the MPLS network is configured properly and only the EIGRP configuration is observed.

When deploying EIGRP over EoMPLS, there are no changes to the EIGRP configuration from the customer perspective. EIGRP needs to be enabled with the correct autonomous system (AS) number (the same on both routers R1 and R2). In addition, the **network** commands must include all of the interfaces required in the EIGRP process. This applies to the link toward routers PE1 and PE2 as well as routers R1 and R2, which will form the neighbor relationship to each other over the MPLS backbone. From the EIGRP perspective, the MPLS backbone and routers PE1 and PE2 are not visible. A neighbor relationship is established directly between routers R1 and R2 as it is visible from the **show ip eigrp neighbors** command output.

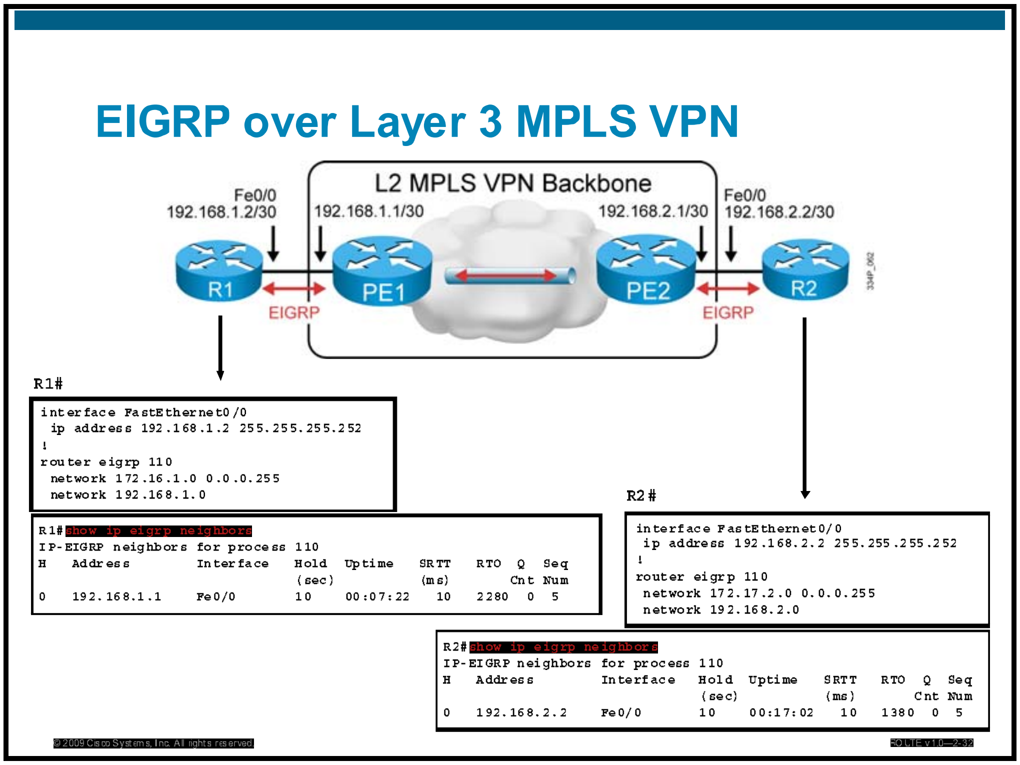

EIGRP over Layer 3 MPLS VPN

```
R1#
interface FastEthernet0/0
 ip address 192.168.1.2 255.255.255.252
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
```

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q   Seq
                                  (sec)           (ms)        Cnt Num
0   192.168.1.1     Fe0/0         10    00:07:22   10    2280  0   5
```

```
R2#
interface FastEthernet0/0
 ip address 192.168.2.2 255.255.255.252
!
router eigrp 110
 network 172.17.2.0 0.0.0.255
 network 192.168.2.0
```

```
R2#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address         Interface     Hold  Uptime    SRTT   RTO  Q   Seq
                                  (sec)           (ms)        Cnt Num
0   192.168.2.2     Fe0/0         10    00:17:02   10    1380  0   5
```

ROUTE v1.0—2-32

Routers R1 and R2 are deployed with EIGRP as if there were a corporate core network between them. EIGRP is enabled on the proper interfaces using the **network** command. The only difference is that the customer has to agree with the service provider regarding the EIGRP parameters (such as the AS number, authentication password, and so on), as these parameters are often governed by the service provider.

The PE routers receive IPv4 routing updates from the CE routers and install these updates in the appropriate virtual routing and forwarding (VRF) table. This part of the configuration and operation is the responsibility of a service provider.

Summary

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

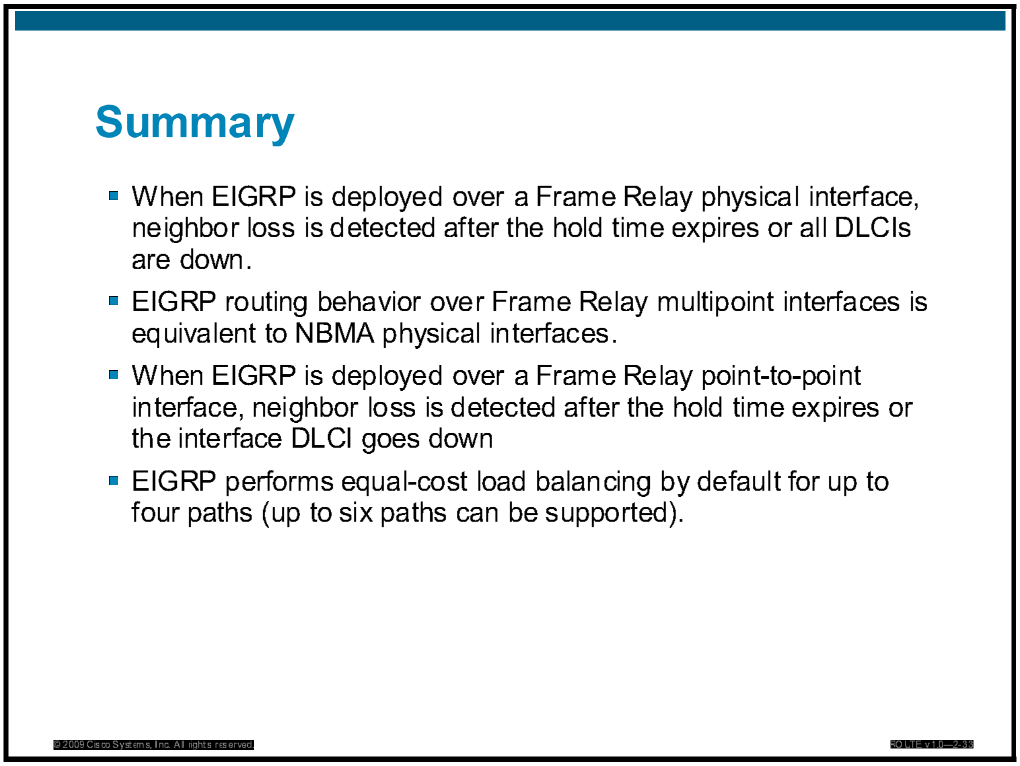

Summary

- When EIGRP is deployed over a Frame Relay physical interface, neighbor loss is detected after the hold time expires or all DLCIs are down.
- EIGRP routing behavior over Frame Relay multipoint interfaces is equivalent to NBMA physical interfaces.
- When EIGRP is deployed over a Frame Relay point-to-point interface, neighbor loss is detected after the hold time expires or the interface DLCI goes down
- EIGRP performs equal-cost load balancing by default for up to four paths (up to six paths can be supported).

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-33

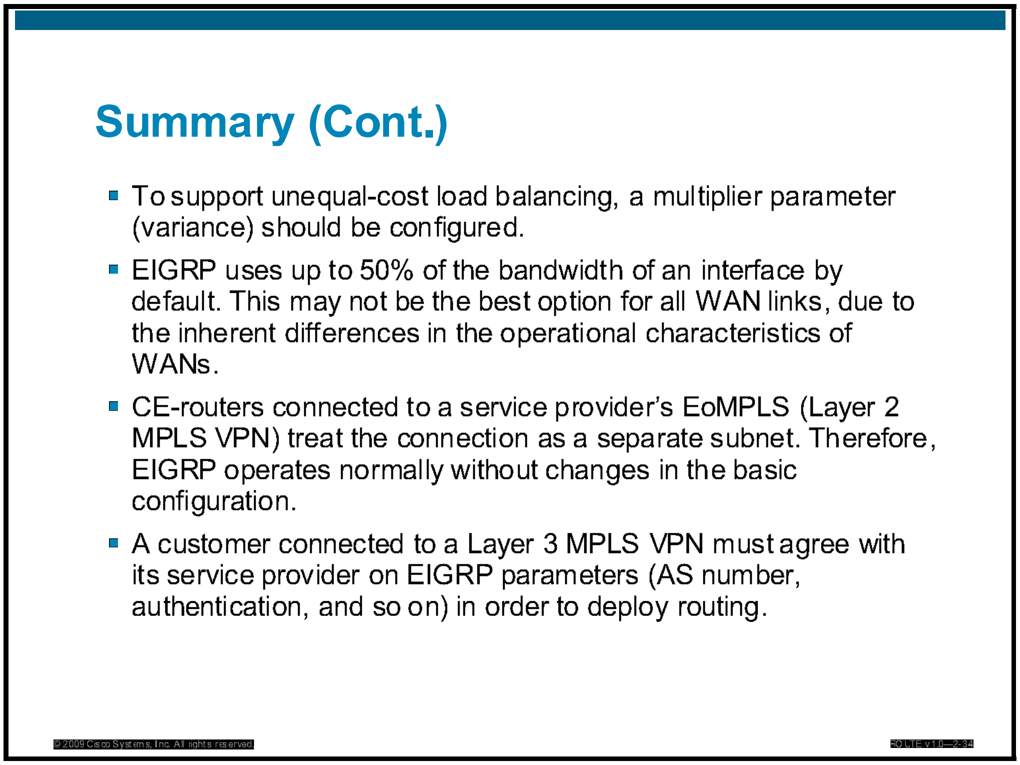

Summary (Cont.)

- To support unequal-cost load balancing, a multiplier parameter (variance) should be configured.
- EIGRP uses up to 50% of the bandwidth of an interface by default. This may not be the best option for all WAN links, due to the inherent differences in the operational characteristics of WANs.
- CE-routers connected to a service provider's EoMPLS (Layer 2 MPLS VPN) treat the connection as a separate subnet. Therefore, EIGRP operates normally without changes in the basic configuration.
- A customer connected to a Layer 3 MPLS VPN must agree with its service provider on EIGRP parameters (AS number, authentication, and so on) in order to deploy routing.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-34

Lesson 5

Lab 2-2 Debrief

Overview

In Lab 2-2, students configure and verify EIGRP circuit emulation and Frame Relay operations. First, they configure EIGRP on point-to-point interfaces, then multipoint interfaces. After that, they adjust EIGRP over multipoint WAN interface and configure EIGRP unequal-cost path load balancing.

Objectives

Upon completing this lesson, you will be able to configure and verify EIGRP circuit emulation and Frame Relay operations. This ability includes being able to meet these objectives:

- Complete the lab overview and verification
- Describe a sample solution and alternatives

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and to start with the verification.

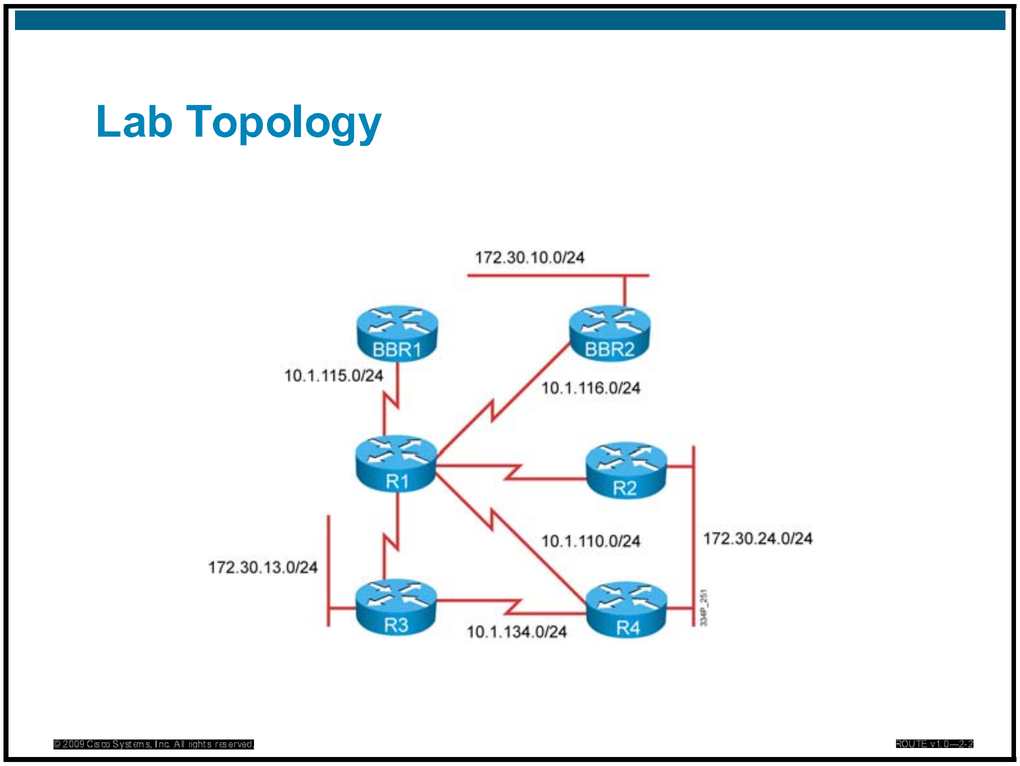

Lab Topology

ROUTE v1.0—2-2

The figure above presents the physical lab topology used for Lab 2-2: configuring and verifying EIGRP circuit emulation and Frame Relay operations. The topology uses four pod routers and two backbone routers. All routers participate in the EIGRP routing protocol.

Based on the topology, students will identify the required parameters and EIGRP configuration over point-to-point and multipoint WAN interfaces.

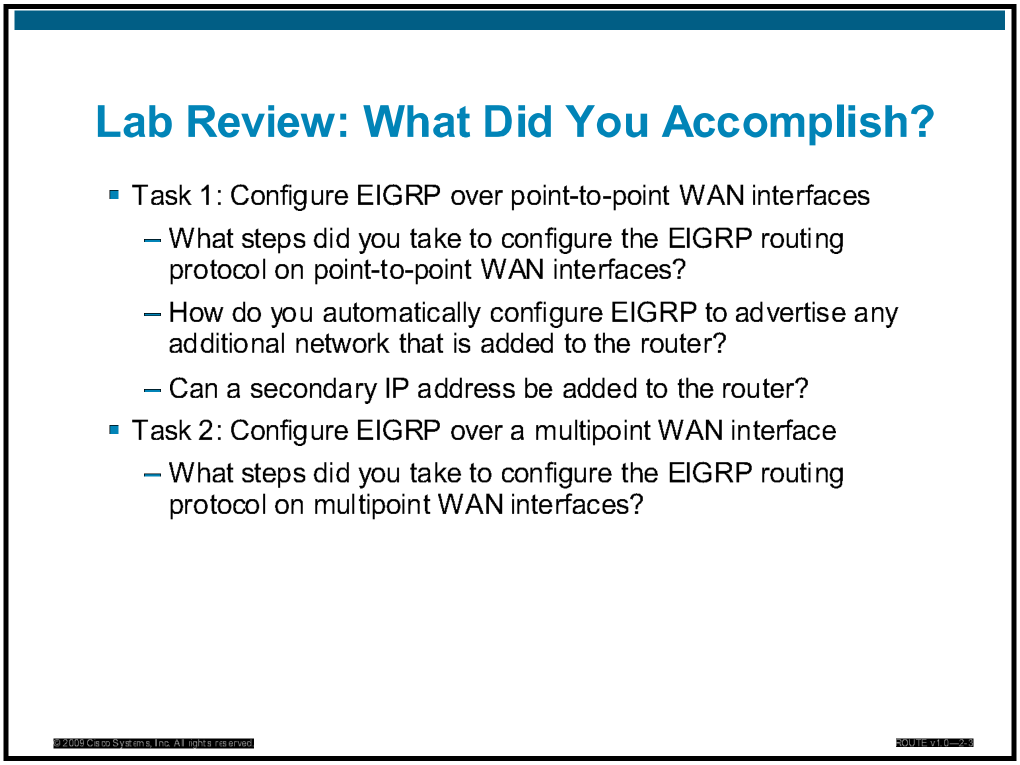

Lab Review: What Did You Accomplish?

- Task 1: Configure EIGRP over point-to-point WAN interfaces
  - What steps did you take to configure the EIGRP routing protocol on point-to-point WAN interfaces?
  - How do you automatically configure EIGRP to advertise any additional network that is added to the router?
  - Can a secondary IP address be added to the router?
- Task 2: Configure EIGRP over a multipoint WAN interface
  - What steps did you take to configure the EIGRP routing protocol on multipoint WAN interfaces?

ROUTE v1.0—2-3

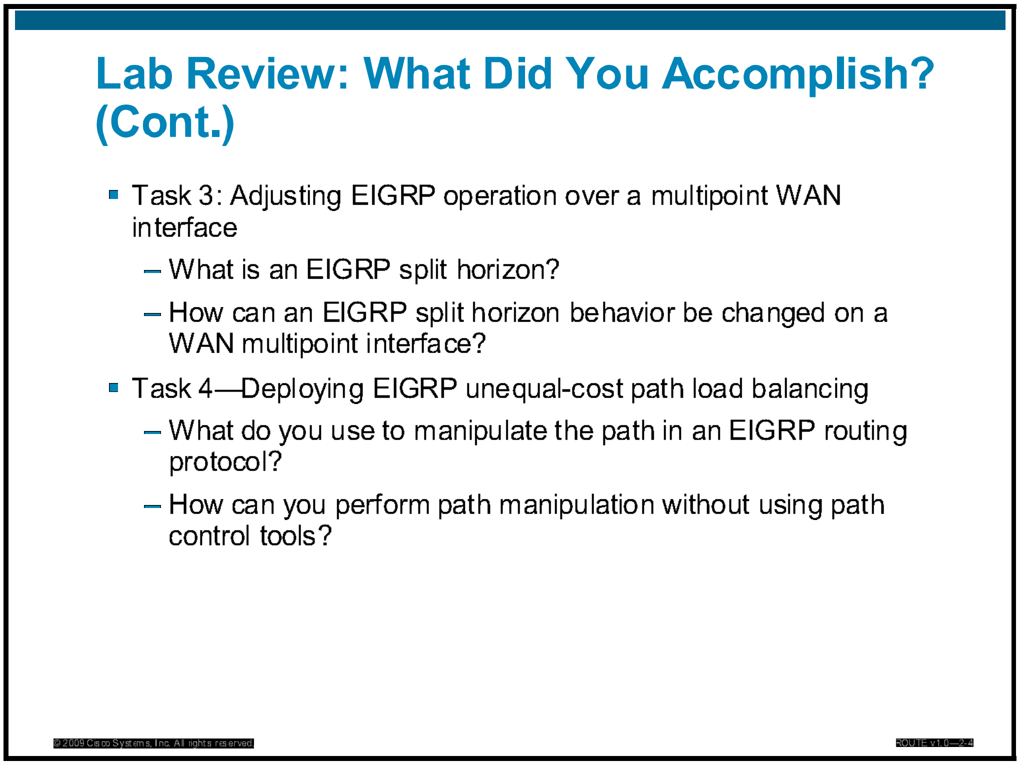

Lab Review: What Did You Accomplish? (Cont.)

- Task 3: Adjusting EIGRP operation over a multipoint WAN interface
  - What is an EIGRP split horizon?
  - How can an EIGRP split horizon behavior be changed on a WAN multipoint interface?
- Task 4—Deploying EIGRP unequal-cost path load balancing
  - What do you use to manipulate the path in an EIGRP routing protocol?
  - How can you perform path manipulation without using path control tools?

ROUTE v1.0—2-4

In the first task, you configured EIGRP over point-to-point WAN interfaces. The EIGRP configuration automatically advertises any additional network that is added to the router. You also added a secondary IP address to the router.

In the second task, you configured EIGRP over a multipoint WAN interface. Again, the EIGRP configuration automatically advertises any additional network that is added to the router.

In the third task, you configured additional EIGRP functionalities in order to adjust the EIGRP operation over a multipoint WAN interface. Split horizon was disabled.

In the fourth task, you deployed EIGRP unequal-cost path load balancing. In order to manipulate the path, you changed the metric in EIGRP.

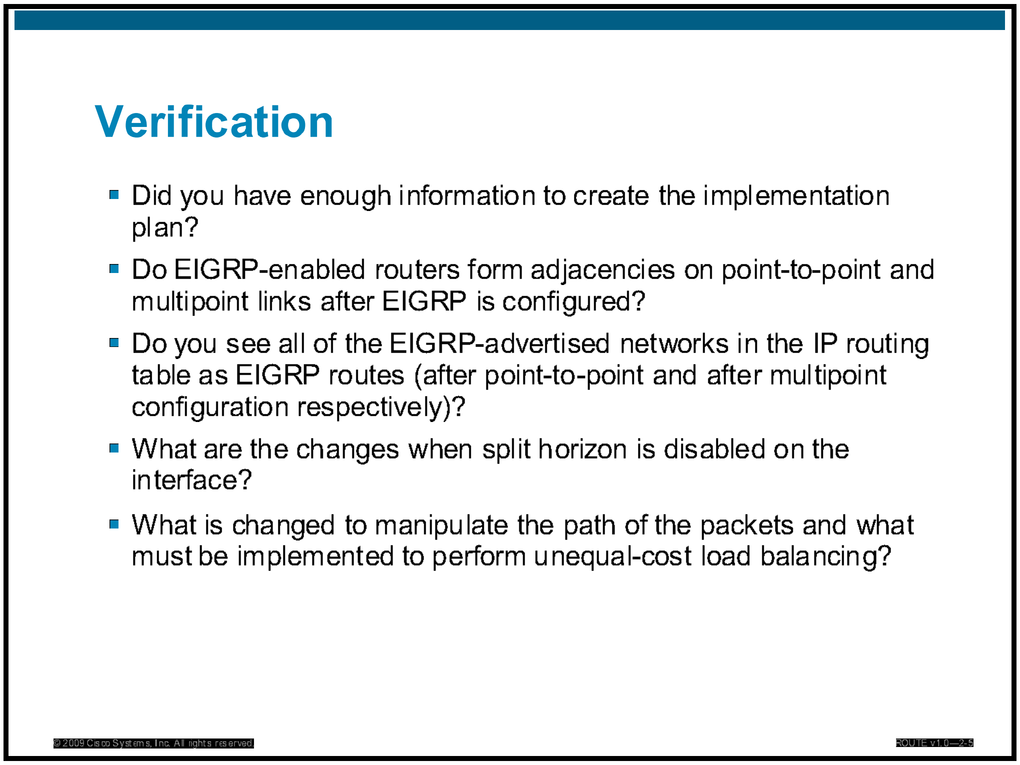

Verification

- Did you have enough information to create the implementation plan?
- Do EIGRP-enabled routers form adjacencies on point-to-point and multipoint links after EIGRP is configured?
- Do you see all of the EIGRP-advertised networks in the IP routing table as EIGRP routes (after point-to-point and after multipoint configuration respectively)?
- What are the changes when split horizon is disabled on the interface?
- What is changed to manipulate the path of the packets and what must be implemented to perform unequal-cost load balancing?

A common approach to verifying the implementation process for a routing protocol is to answer the following questions:

- Did you have sufficient information to create the implementation plan?
- Do EIGRP-enabled routers form adjacencies on point-to-point and multipoint links after EIGRP is configured?
- Do you see all of the EIGRP-advertised networks in the IP routing table as EIGRP routes (after point-to-point and multipoint configuration)?
- What are the changes when split horizon behavior is disabled on the interface?
- What is changed to manipulate the path of the packets and what must be implemented to perform unequal-cost load balancing?

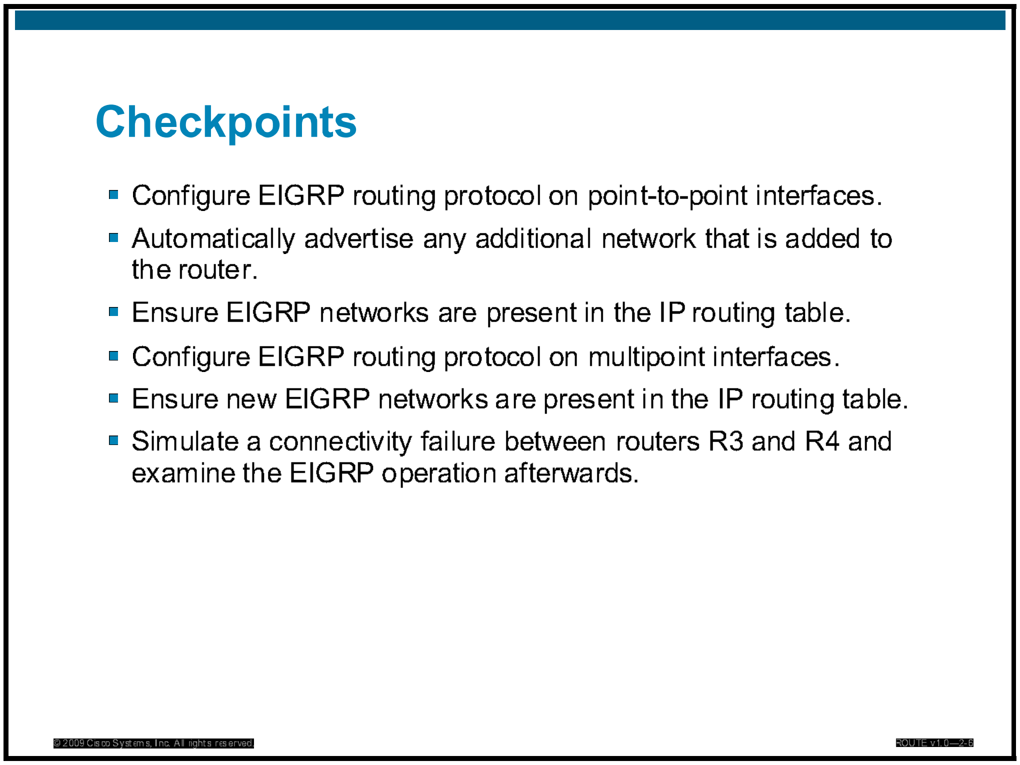

Checkpoints

- Configure EIGRP routing protocol on point-to-point interfaces.
- Automatically advertise any additional network that is added to the router.
- Ensure EIGRP networks are present in the IP routing table.
- Configure EIGRP routing protocol on multipoint interfaces.
- Ensure new EIGRP networks are present in the IP routing table.
- Simulate a connectivity failure between routers R3 and R4 and examine the EIGRP operation afterwards.

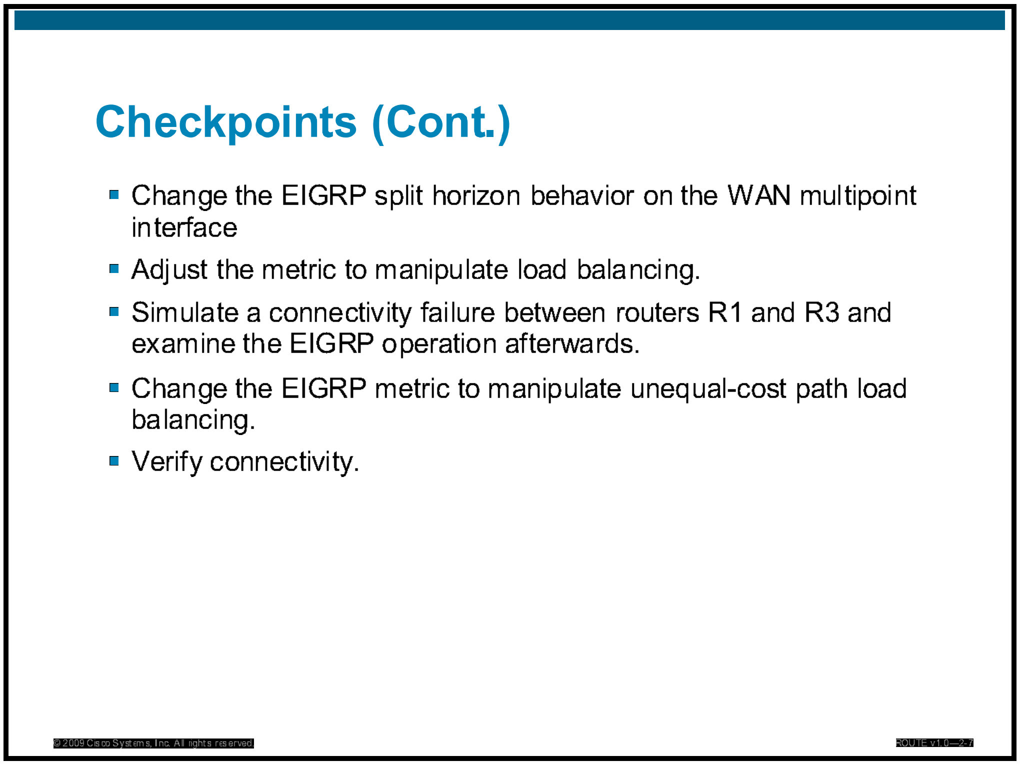

Checkpoints (Cont.)

- Change the EIGRP split horizon behavior on the WAN multipoint interface
- Adjust the metric to manipulate load balancing.
- Simulate a connectivity failure between routers R1 and R3 and examine the EIGRP operation afterwards.
- Change the EIGRP metric to manipulate unequal-cost path load balancing.
- Verify connectivity.

During the configuration and verification phase, the network operator deals with several checkpoints. After completing all configuration tasks, the network operator has either finished the lab successfully or must perform additional verification and troubleshooting.

With different checkpoints, the network operator can verify for proper configuration. The following checkpoints are used for verification:

- Configure the EIGRP routing protocol on point-to-point interfaces.
- Automatically advertise any additional network that is added to the router.
- Ensure that EIGRP networks are present in the IP routing table.
- Configure the EIGRP routing protocol on multipoint interfaces.
- Ensure new EIGRP networks are present in the IP routing table.
- Simulate a connectivity failure between routers R3 and R4 and examine the EIGRP operation afterwards.
- Change the EIGRP split horizon behavior on the WAN multipoint interface.
- Adjust the metric to manipulate load balancing.
- Simulate a connectivity failure between routers R1 and R3 and examine the EIGRP operation afterwards.
- Change the EIGRP metric to manipulate unequal-cost path load balancing.
- Verify connectivity.

Sample Solutions and Alternatives

This topic describes sample solution and other alternatives.

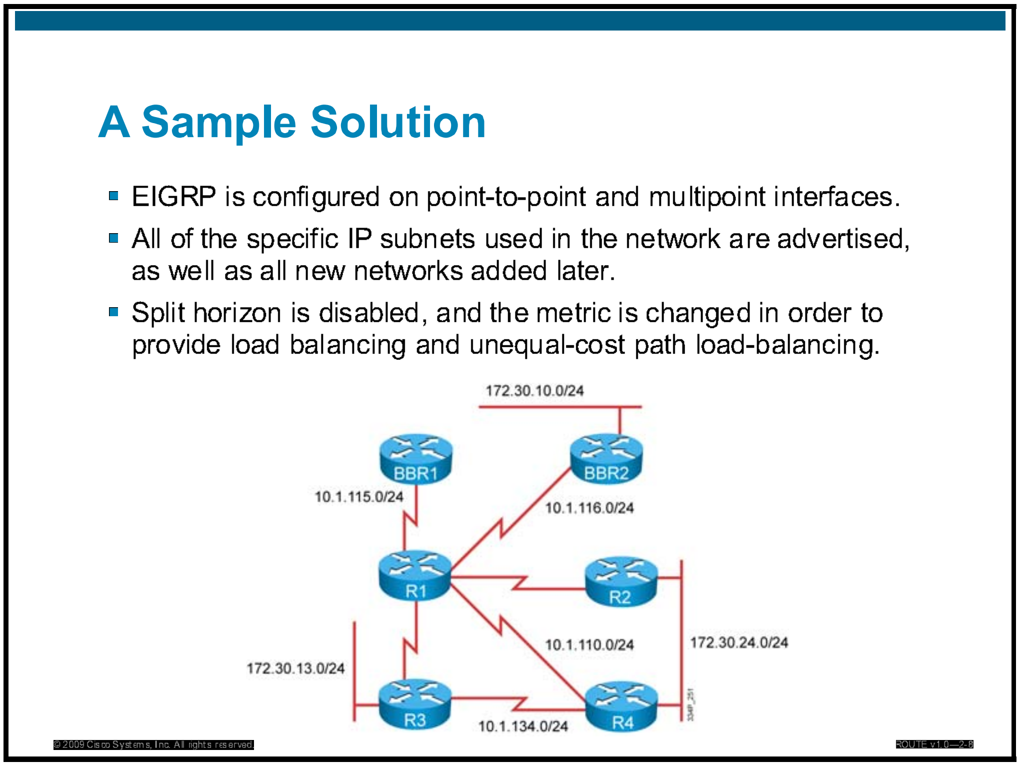

A Sample Solution

- EIGRP is configured on point-to-point and multipoint interfaces.
- All of the specific IP subnets used in the network are advertised, as well as all new networks added later.
- Split horizon is disabled, and the metric is changed in order to provide load balancing and unequal-cost path load-balancing.

ROUTE v1.0—2-8

The sample solution includes implementation details and details for each task of the implementation plan. Different solutions are possible and the figure shows a few details of a successful configuration.

Proper implementation includes the following details:

- EIGRP is configured on point-to-point and multipoint interfaces.
- All of the specific IP subnets used in the network (including all networks added later) are advertised.
- Split horizon behavior is disabled and the metric is changed in order provide load balancing and unequal-cost path load balancing.

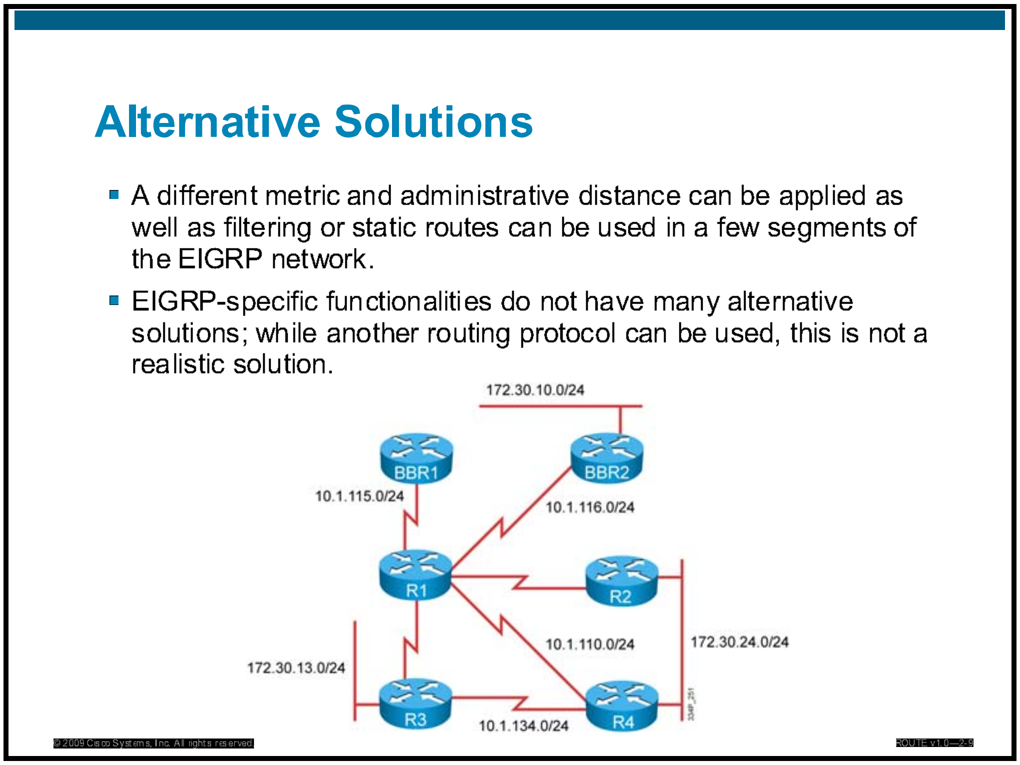

Alternative Solutions

- A different metric and administrative distance can be applied as well as filtering or static routes can be used in a few segments of the EIGRP network.
- EIGRP-specific functionalities do not have many alternative solutions; while another routing protocol can be used, this is not a realistic solution.

ROUTE v1.0—2-9

A different metric and administrative distance can be applied, as well as filtering or static routes can be used in few segments of the EIGRP network to provide a similar solution. However, EIGRP-specific functionalities do not have many alternative solutions. To implement a similar solution, another routing protocol can be used, which is not a realistic solution as changing the routing protocol is not the case during fine tuning of the existing protocol.

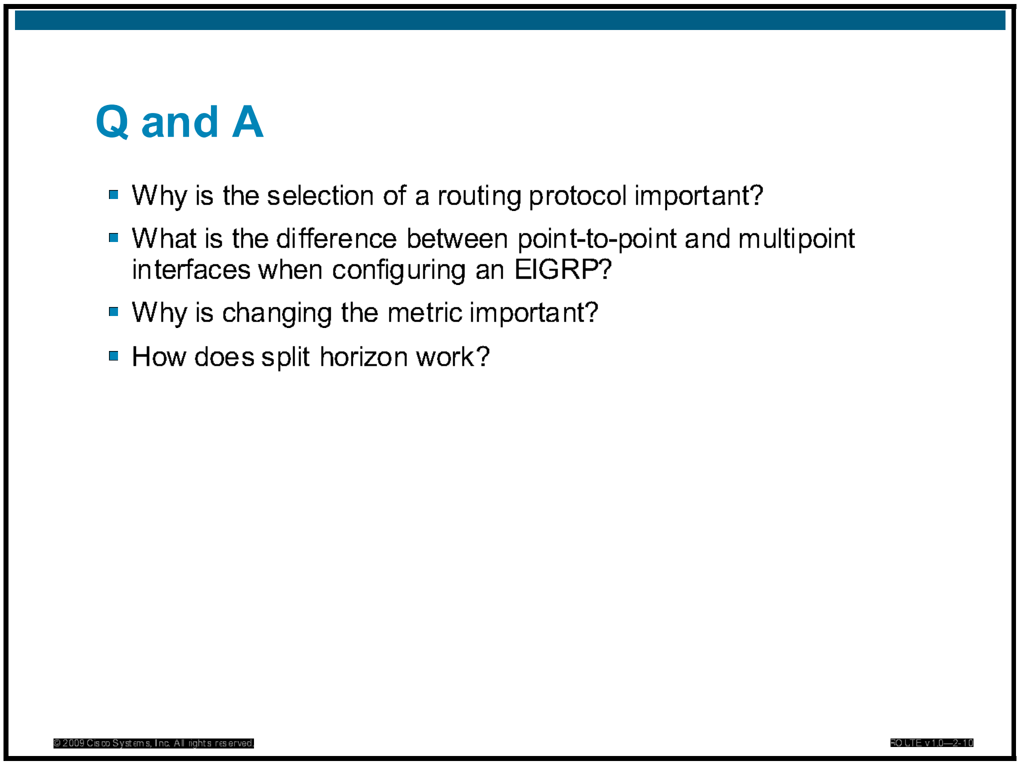

Q and A

- Why is the selection of a routing protocol important?
- What is the difference between point-to-point and multipoint interfaces when configuring an EIGRP?
- Why is changing the metric important?
- How does split horizon work?

ROUTE v1.0—2-10

The routing protocol, with its metric and administrative distance, exchanges routing updates. It also populates the IP routing table, which is used for destination-based forwarding. Different routing protocols process routing updates in different ways.

When configuring EIGRP on point-to-point interfaces, just two routers are connected to the link. They exchange routing updates. If the interface is a multipoint interface, routing updates received from one neighbor must be sent to another neighbor through the same interface.

The metric provides the importance or the quality of the routes within a routing protocol. By manipulating the administrative distance and metric value, you can implement path manipulation, as well.

Because of the split horizon behavior, the routing updates are not sent back to the interface from which they were received. This breaks the exchange of EIGRP routing updates.

Summary

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

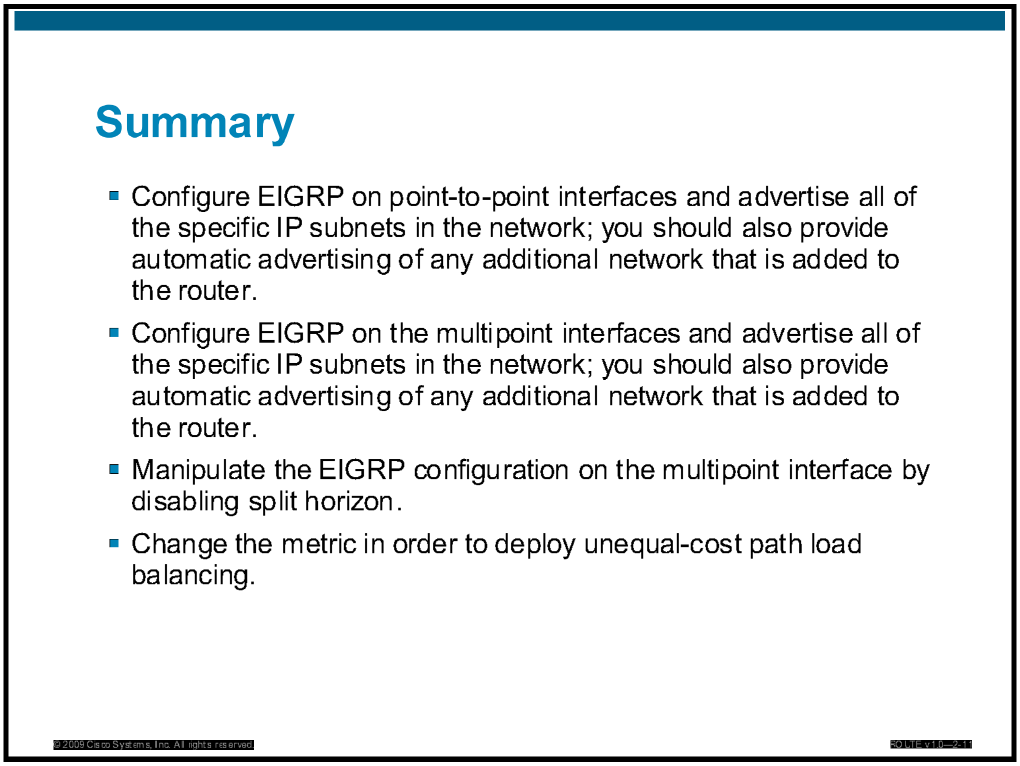

Summary

- Configure EIGRP on point-to-point interfaces and advertise all of the specific IP subnets in the network; you should also provide automatic advertising of any additional network that is added to the router.
- Configure EIGRP on the multipoint interfaces and advertise all of the specific IP subnets in the network; you should also provide automatic advertising of any additional network that is added to the router.
- Manipulate the EIGRP configuration on the multipoint interface by disabling split horizon.
- Change the metric in order to deploy unequal-cost path load balancing.

Lesson 6

Implementing and Verifying EIGRP Authentication

Overview

You can prevent your router from receiving false route updates by configuring neighbor router authentication. Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor authentication (also called neighbor router authentication or route authentication) can be configured in such a way that routers can participate in routing based on predefined passwords.

This lesson describes EIGRP Message Digest 5 (MD5) authentication and how to configure and verify it.

Objectives

Upon completing this lesson, you will be able to implement and verify authentication in an EIGRP network. This ability includes being able to meet these objectives:

- Determine router authentication for EIGRP.
- Determine MD5 authentication for EIGRP.
- Implement MD5 authentication for EIGRP.
- Verify MD5 authentication for EIGRP.

Router Authentication for EIGRP

This topic describes the router authentication used by routing protocols.

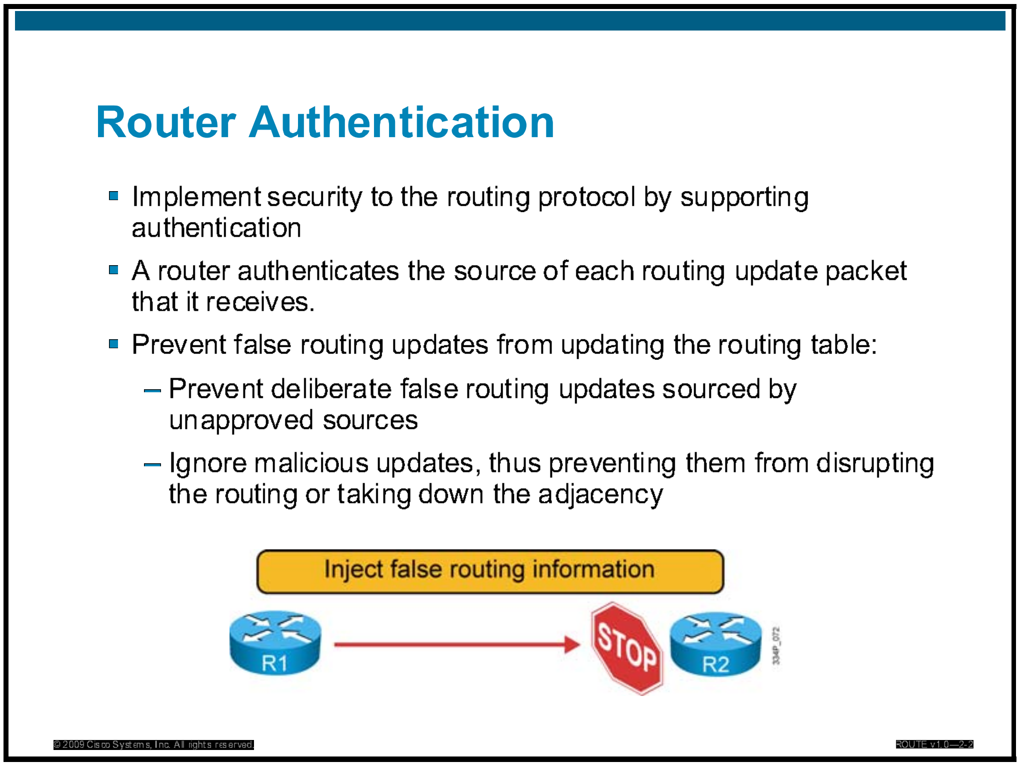

Router Authentication

- Implement security to the routing protocol by supporting authentication
- A router authenticates the source of each routing update packet that it receives.
- Prevent false routing updates from updating the routing table:
  - Prevent deliberate false routing updates sourced by unapproved sources
  - Ignore malicious updates, thus preventing them from disrupting the routing or taking down the adjacency

ROUTE v1.0—2-2

Neighbor router authentication (also called route authentication) can be configured in such way that routers only participate in routing based on predefined passwords.

Without neighbor authentication, unauthorized or deliberately malicious routing updates can compromise the security of your network traffic. A security compromise may occur if any unfriendly party interferes with your network. For example, an unauthorized router might launch a fictitious routing update to convince your router to send traffic to an incorrect destination.

When neighbor authentication has been configured on routers, those routers authenticate the source of each routing update packet they receive. They do so by exchanging an authentication key that is known to both the sending and the receiving router.

By default, no authentication is used for routing protocol packets.

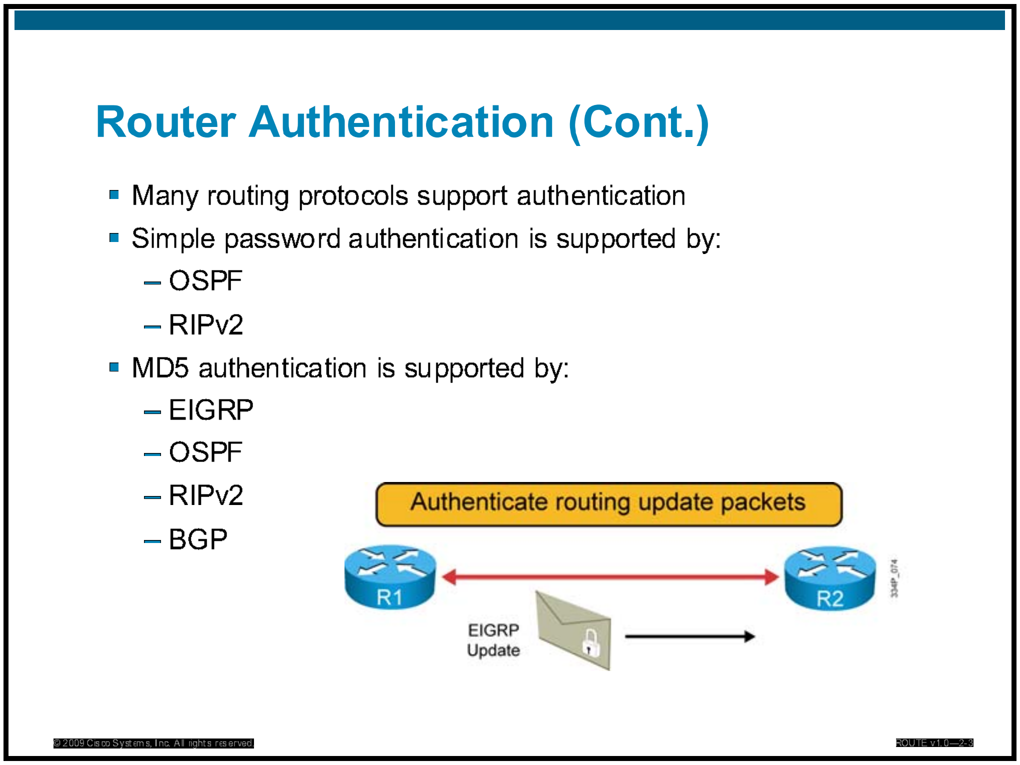

Router Authentication (Cont.)

- Many routing protocols support authentication
- Simple password authentication is supported by:
  - OSPF
  - RIPv2
- MD5 authentication is supported by:
  - EIGRP
  - OSPF
  - RIPv2
  - BGP

Authenticate routing update packets

ROUTE v1.0—2-3

There are two types of authentication: simple password authentication (also called plaintext authentication) and Message Digest 5 (MD5) authentication.

Simple password authentication is supported by Open Shortest Path First (OSPF) and Routing Information Protocol version 2 (RIPv2). MD5 authentication is supported by OSPF, RIPv2, Border Gateway Protocol (BGP), and EIGRP.

Note Authentication for EIGRP, OSPF, and BGP is covered in this course.

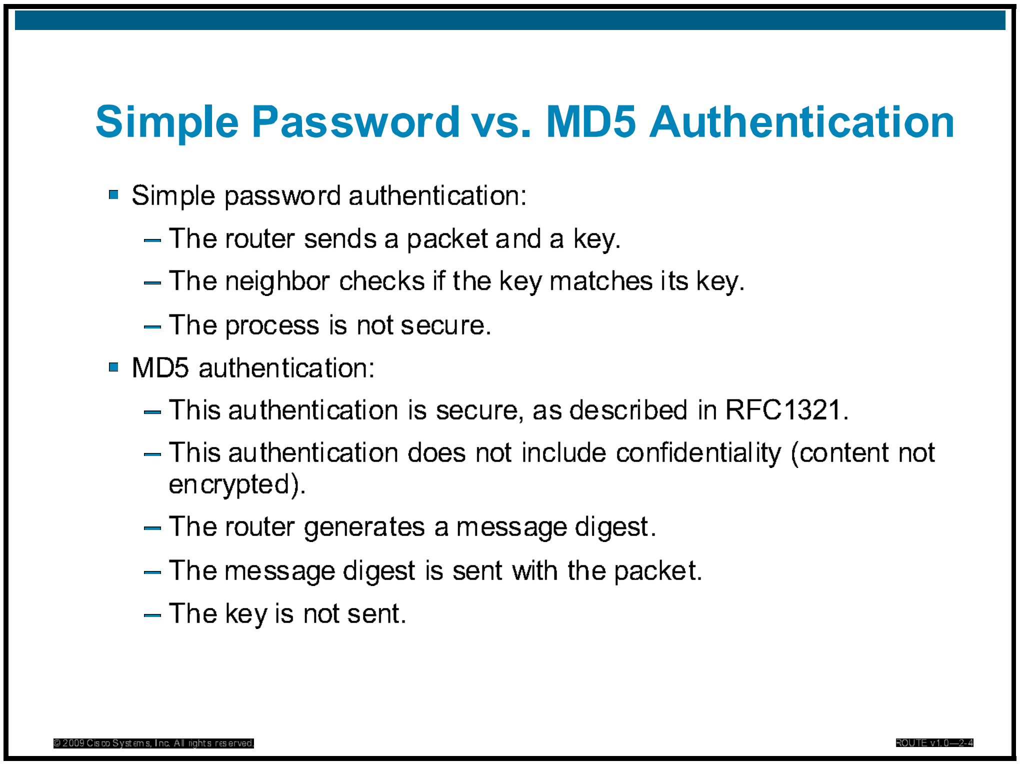

Simple Password vs. MD5 Authentication

- Simple password authentication:
  - The router sends a packet and a key.
  - The neighbor checks if the key matches its key.
  - The process is not secure.
- MD5 authentication:
  - This authentication is secure, as described in RFC1321.
  - This authentication does not include confidentiality (content not encrypted).
  - The router generates a message digest.
  - The message digest is sent with the packet.
  - The key is not sent.

ROUTE v1.0—2-4

The behavior of simple password authentication is the same as that of MD5 authentication, except that MD5 sends a message digest instead of the authenticating key itself. MD5 creates the message digest using the key and a message, but the key itself is not sent, which prevents it from being read while it is being transmitted. Simple password authentication sends the authenticating key itself over the wire.

Note Note that simple password authentication is not recommended for use as part of your security strategy, because it is vulnerable to passive attacks. Anybody with a link analyzer could easily view the password on the wire. The primary use of simple password authentication is to avoid accidental changes to the routing infrastructure. Using MD5 authentication, however, is a recommended security practice.

With simple password authentication, a password (key) is configured on a router, and each participating neighbor router must be configured with the same key. When a packet is sent over the wire, the password, in plaintext, is sent along with it.

MD5 authentication is a cryptographic authentication. A key (password) and key ID are configured on each router. The router uses an algorithm to generate a message digest (also called a hash) from the key and key ID, and appends the message digest to the packet. Unlike simple authentication, the key is not exchanged over the wire; the message digest is sent instead of the key, which ensures that no one can eavesdrop on the line and learn keys during transmission. Although MD5 authentication provides authenticity, it does not provide confidentiality. The content of the routing update is not encrypted.

Note As with all keys, passwords, and other security secrets, it is imperative that you closely guard authenticating keys used in neighbor authentication. In order to obtain the security benefits of this feature, you must keep all authenticating keys confidential. Also, when performing router management tasks via Simple Network Management Protocol (SNMP), do not ignore the risk associated with sending keys using non-encrypted SNMP.

MD5 Authentication for EIGRP

This topic describes MD5 authentication as it is used in EIGRP.

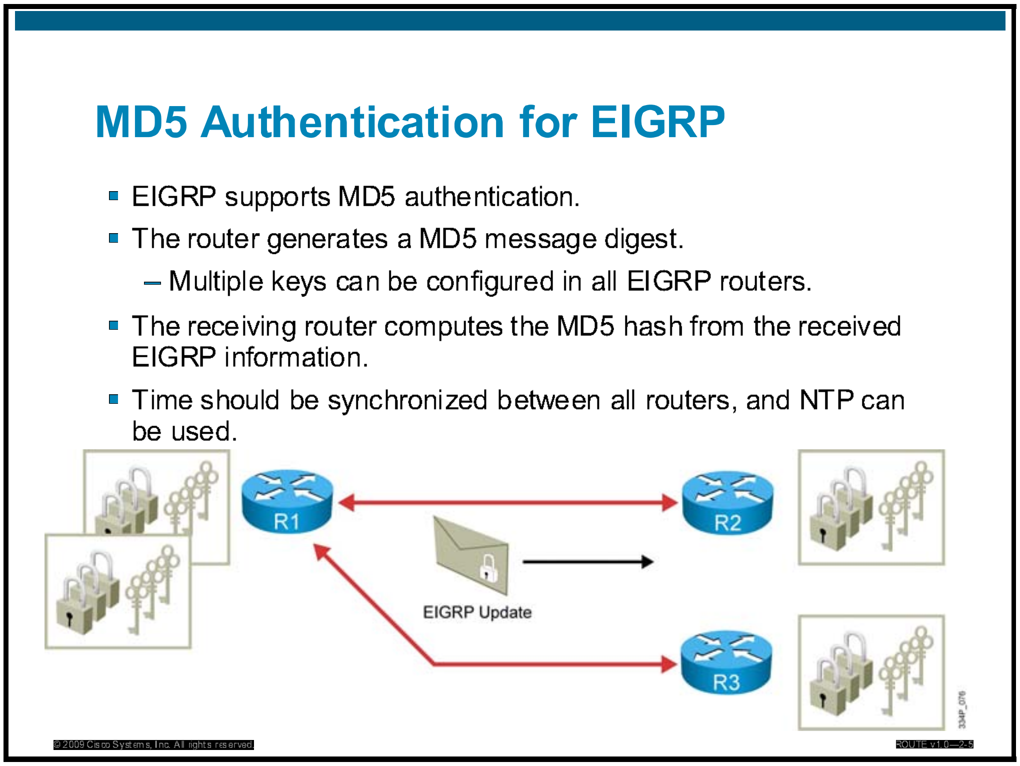

MD5 Authentication for EIGRP

- EIGRP supports MD5 authentication.
- The router generates a MD5 message digest.
  - Multiple keys can be configured in all EIGRP routers.
- The receiving router computes the MD5 hash from the received EIGRP information.
- Time should be synchronized between all routers, and NTP can be used.

EIGRP supports MD5 authentication to prevent the introduction of unauthorized or false routing messages from unapproved sources. EIGRP neighbor authentication (also called neighbor router authentication or route authentication) can be configured in such way that routers can participate in routing based on predefined passwords. By default, no authentication is used for EIGRP packets. EIGRP must be configured to use MD5 authentication.

When neighbor authentication has been configured on a router, the router authenticates the source of each routing update packet that it receives. In order to start using EIGRP MD5 authentication, you must configure an authenticating key (sometimes referred to as a password) and a key identifier on both the sending router and the receiving router. Each EIGRP router takes the key and key ID and generates a message digest that is appended to each routing update and sent to the neighbor. The receiving router computes the MD5 hash from the received EIGRP information. If the hash matches the value received, the packet is accepted. If it does not, the packet is silently dropped.

Each key has its own key ID, which is stored locally. The combination of the key ID and the interface associated with the message uniquely identifies the authentication algorithm and the MD5 authentication key in use. You can increase the security of EIGRP MD5 authentication by making frequent key changes. The definition of multiple keys is supported, which can be changed based on time that is defined in the configuration. Transitioning between the keys is implemented in a way that provides non-disruptive exchange of EIGRP routing updates. The key changes must be well-planned and supported by the time synchronization between the routers. The key rollover works only if the times on the adjacent routers are synchronized. You can use several mechanisms for time synchronization. Network Time Protocol (NTP) is the most commonly used time synchronization mechanism to ensure that the correct time is used by all the participating EIGRP routers using the key rollover mechanism.

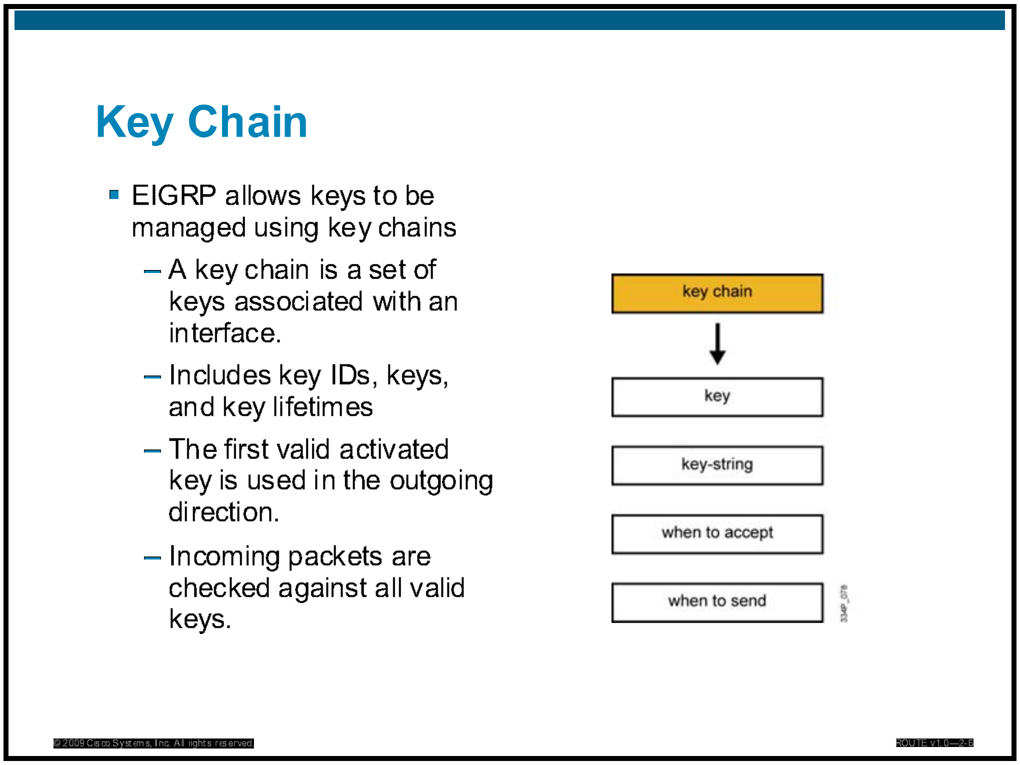

Key Chain

- EIGRP allows keys to be managed using key chains
  - A key chain is a set of keys associated with an interface.
  - Includes key IDs, keys, and key lifetimes
  - The first valid activated key is used in the outgoing direction.
  - Incoming packets are checked against all valid keys.

key chain

key

key-string

when to accept

when to send

ROUTE v1.0—2-6

EIGRP allows keys to be managed using key chains. Each key definition within the key chain can specify a time duration for when that key will be activated (its lifetime). Then, during a given key's lifetime, routing update packets are sent with this activated key. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest. It then uses the first valid key it encounters.

Keys cannot be used at times when they are not activated. Therefore, for a given key chain, you should ensure that key activation times overlap. This will help you avoid any period of time for which no key is activated. If a period occurs during when no key is activated, neighbor authentication cannot occur, and therefore routing updates will fail.

Note that the router needs to know the correct time to be able to rotate through keys in a way that is synchronized with the other participating routers, so that all routers are using the same key at the same moment. Refer to the Network Time Protocol (NTP) and calendar commands in the “Performing Basic System Management” chapter of the *Cisco IOS Configuration Fundamentals Configuration Guide* for information about configuring the time on your router.

Implementing MD5 Authentication for EIGRP

This topic describes how to configure MD5 authentication for EIGRP.

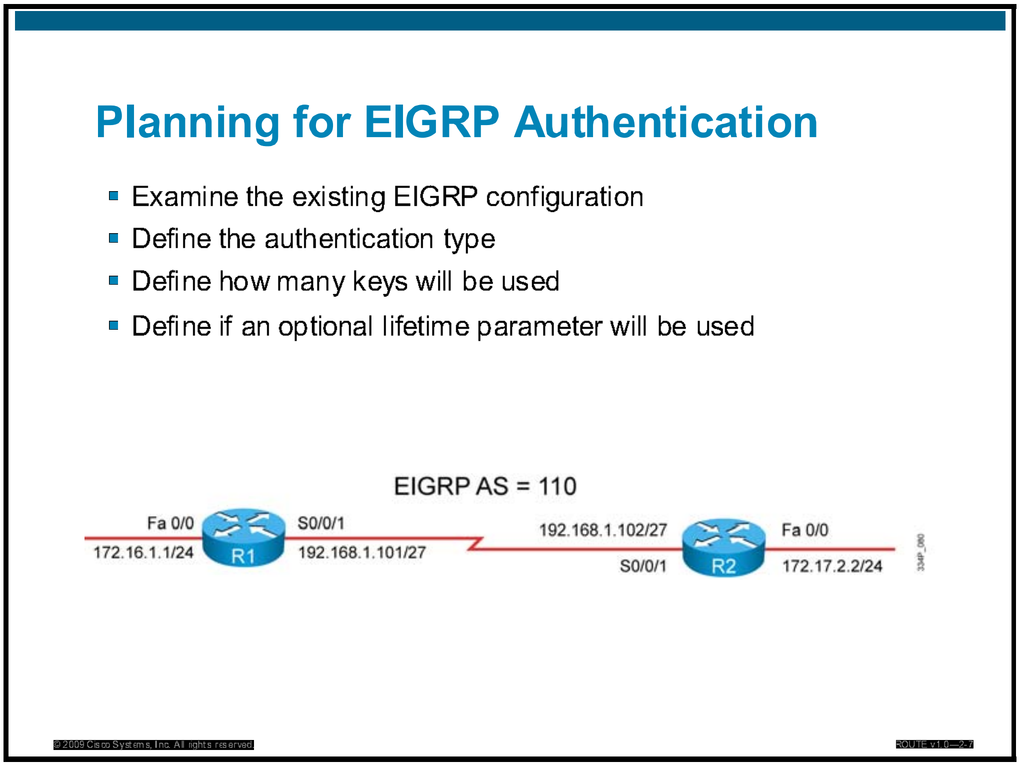

Planning for EIGRP Authentication

- Examine the existing EIGRP configuration
- Define the authentication type
- Define how many keys will be used
- Define if an optional lifetime parameter will be used

ROUTE v1.0—2-7

Before configuring authentication for EIGRP, a network administrator must examine the existing EIGRP configuration and define the network requirements. The existing EIGRP configuration defines which autonomous system (AS) number is used for EIGRP and which routers and interfaces participate in the EIGRP configuration. The network requirements for EIGRP authentication define which parameters must be gathered and include the definition of the authentication type, number of keys used in the EIGRP network, and optional lifetime parameters used.

The next step is to gather of all of the parameters needed to provide enough details for the network operator to start setting up EIGRP authentication.

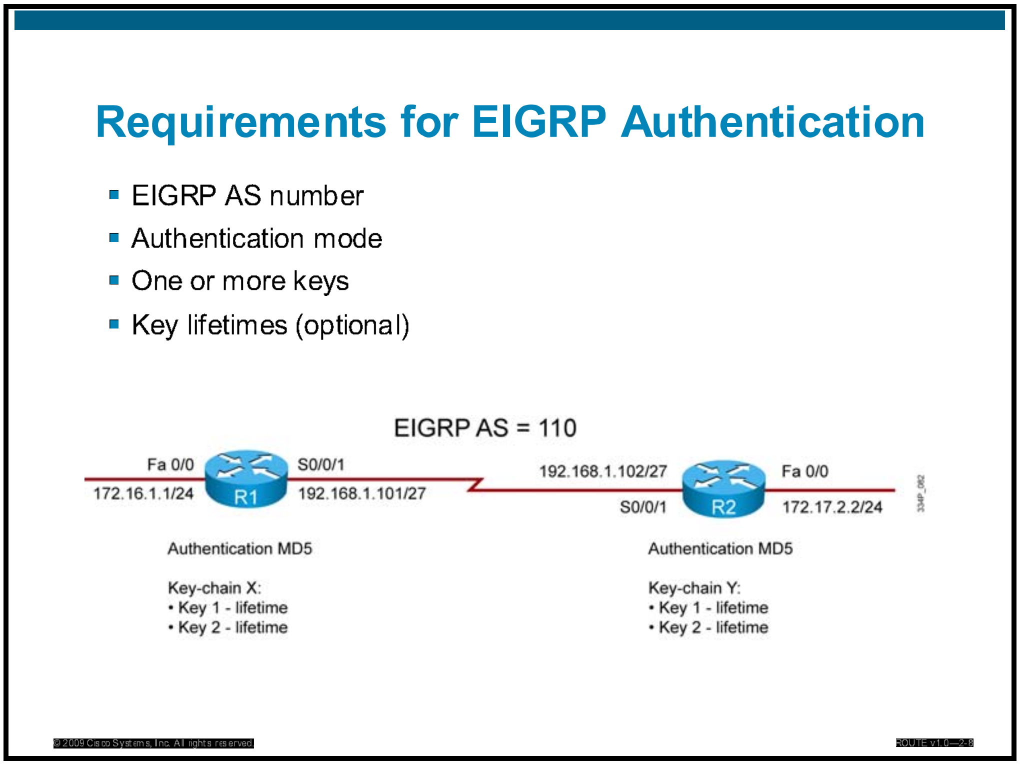

Requirements for EIGRP Authentication

- EIGRP AS number
- Authentication mode
- One or more keys
- Key lifetimes (optional)

ROUTE v1.0—2-8

Requirements to configure EIGRP authentication include the following elements:

- The EIGRP AS number
- The authentication mode selected
- The definition of one or more keys
- The lifetime of each defined key

The EIGRP AS number must be defined, as authentication is locked to the EIGRP process. The authentication mode used is MD5. The definition of the keys is a process by which the network administrator and network designer must develop a security plan. They are not trying to encrypt the packet, but authenticate the source of the EIGRP routing updates. Using more keys and changing them can reduce the potential risk. When more keys are used, it makes sense to change the keys based on predefined time intervals. A key’s lifetime can be defined. It is optional and requires an NTP server in order to synchronize the clocks of all routers running EIGRP.

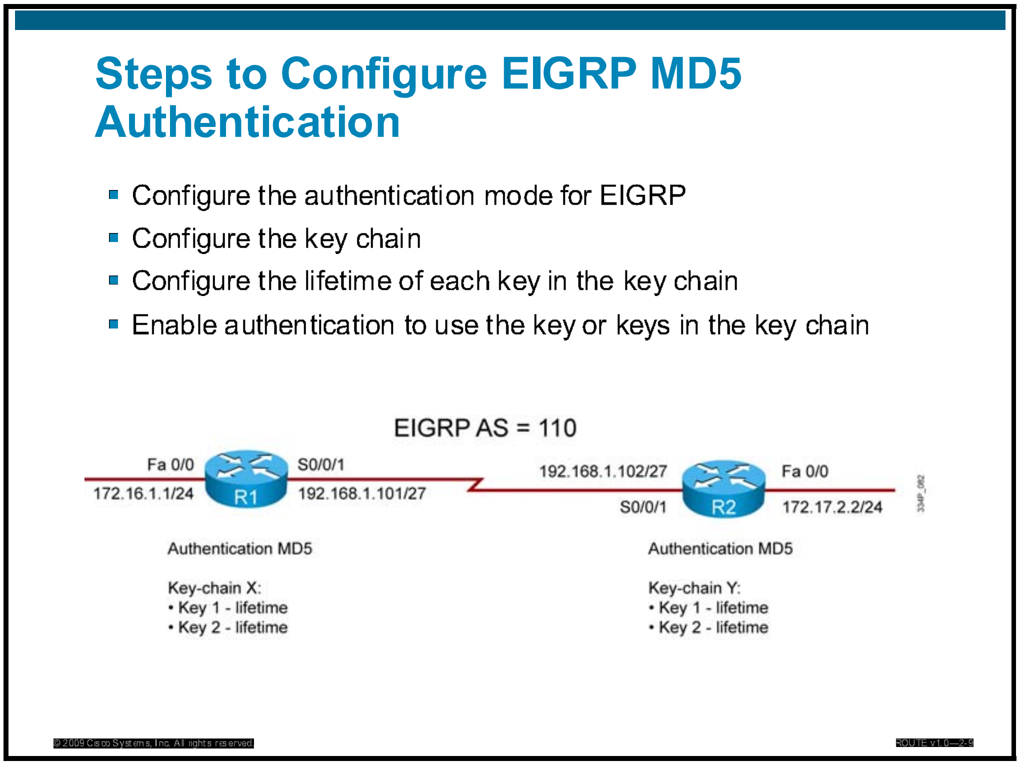

Steps to Configure EIGRP MD5 Authentication

- Configure the authentication mode for EIGRP
- Configure the key chain
- Configure the lifetime of each key in the key chain
- Enable authentication to use the key or keys in the key chain

EIGRP AS = 110

Fa 0/0
172.16.1.1/24
R1
S0/0/1
192.168.1.101/27

192.168.1.102/27
S0/0/1
R2
Fa 0/0
172.17.2.2/24

Authentication MD5

Key-chain X:
• Key 1 - lifetime
• Key 2 - lifetime

Authentication MD5

Key-chain Y:
• Key 1 - lifetime
• Key 2 - lifetime

ROUTE v1.0—2-9

EIGRP MD5 authentication configuration steps:

Step 1 Configure the authentication mode for EIGRP.

Step 2 Configure the key chain.

Step 3 Optional: Configure the lifetime parameters for the keys.

Step 4 Enable authentication to use the key or keys in the key chain.

To complete these steps, you must have information about the existing EIGRP process, definitions of all the keys, and any lifetime parameters to be used in the configuration.

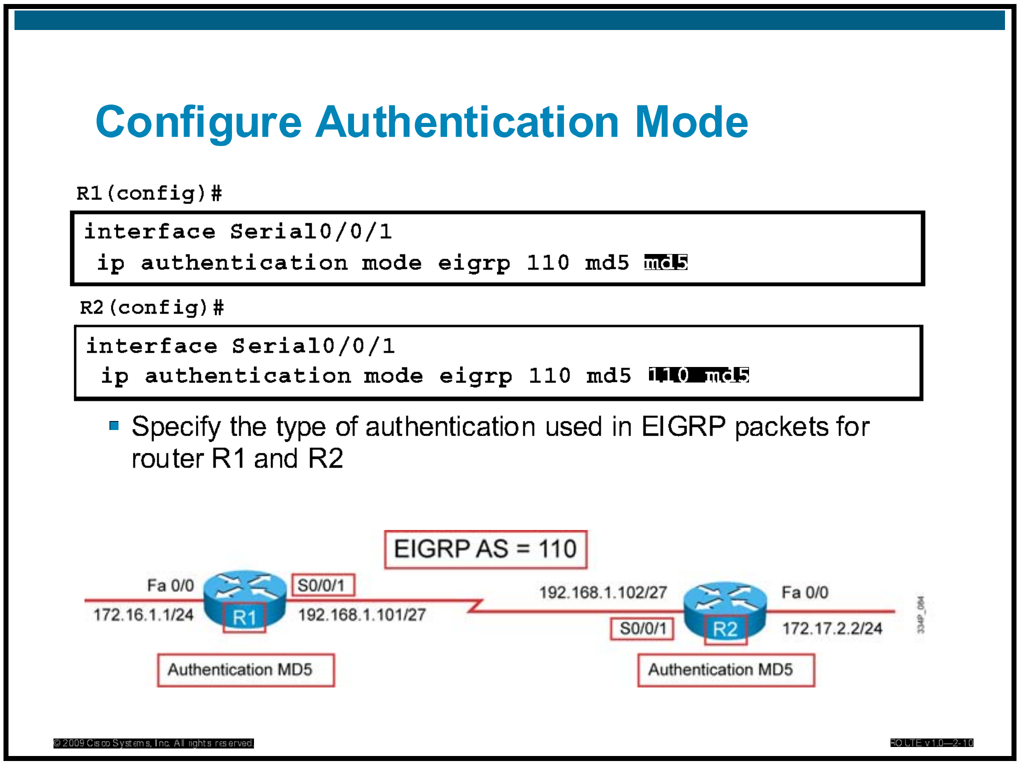

Configure Authentication Mode

```
R1(config)#
interface Serial0/0/1
 ip authentication mode eigrp 110 md5 md5
```

```
R2(config)#
interface Serial0/0/1
 ip authentication mode eigrp 110 md5 110 md5
```

- Specify the type of authentication used in EIGRP packets for router R1 and R2

ROUTE v1.0—2-10

You must configure authentication between any two neighbors exchanging EIGRP routing updates. You must also use the correct AS number and enable the authentication configuration on all interfaces between the neighbors across the EIGRP domain.

The example in this slide shows two routers, R1 and R2, that are running the EIGRP process with AS number 110. They are connected via the serial 0/0/1 interfaces and configuration must be applied to both interfaces involved. To configure MD5 authentication for EIGRP, complete the following steps:

Step 1 Enter the configuration mode for the interface on which you want to enable authentication.

Step 2 Specify the type of the authentication using the **ip authentication mode eigrp 110 md5** interface command.

The only authentication type available is MD5, and AS number 110 is used in the example on this slide, where both routers R1 and R2 are configured for MD5 authentication.

When authentication is configured, an MD5 keyed digest is added to each EIGRP packet in the specified AS.

For more details about the **ip authentication mode eigrp** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

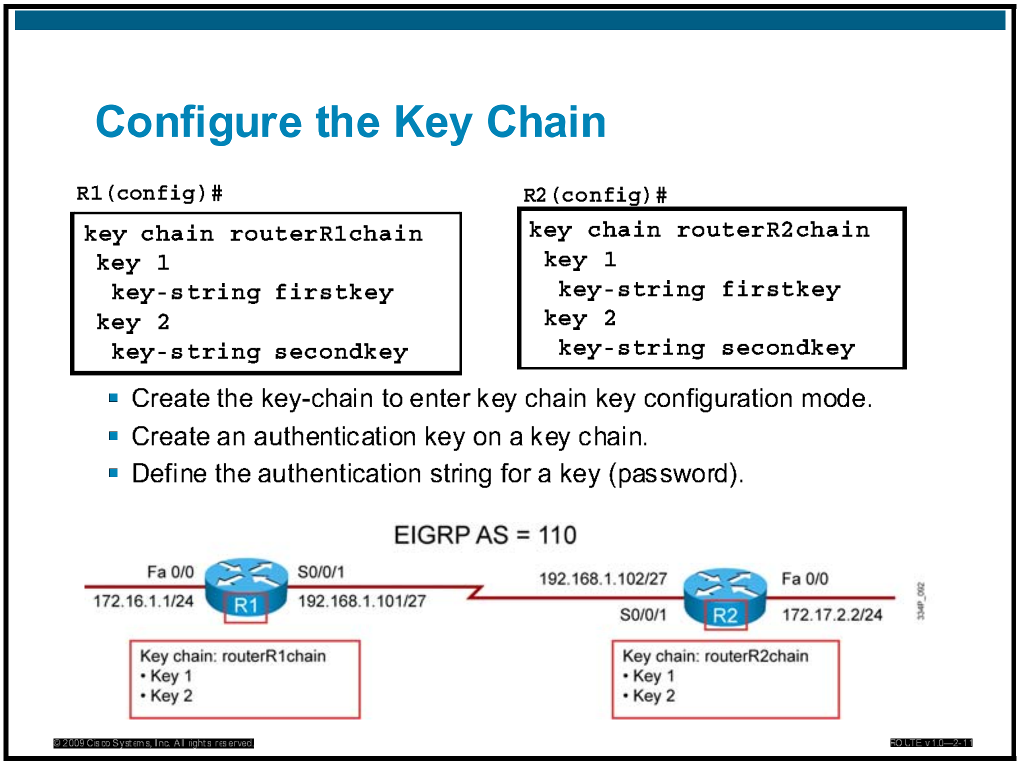

Configure the Key Chain

R1(config)#

```
key chain routerR1chain
 key 1
  key-string firstkey
 key 2
  key-string secondkey
```

R2(config)#

```
key chain routerR2chain
 key 1
  key-string firstkey
 key 2
  key-string secondkey
```

- Create the key-chain to enter key chain key configuration mode.
- Create an authentication key on a key chain.
- Define the authentication string for a key (password).

EIGRP AS = 110
Fa 0/0
172.16.1.1/24
R1
S0/0/1
192.168.1.101/27
192.168.1.102/27
S0/0/1
R2
Fa 0/0
172.17.2.2/24

Key chain: routerR1chain
- Key 1
- Key 2

Key chain: routerR2chain
- Key 1
- Key 2

ROUTE v1.0—2-11

When authentication is enabled on the interface, a group of authentication keys must be defined. The **key chain** global configuration command is used to define all of the keys used for EIGRP MD5 authentication. Once you are in the key chain configuration mode, use the **key** command to identify the key in the key chain. Each key is defined by the number, which defines the key ID. When the key command is used, the configuration enters the key chain key configuration mode, where the **key-string** *authentication-key* configuration command must be used to specify the authentication string (or password). The key ID and authentication string must be the same on all neighboring routers.

In the example on this slide, each interface is configured to enter its own key chain. The key chain configured on router R1 is “routerR1chain” and key-chain on router R2 is “routerR2chain”. Inside each of the key chains, two keys are defined, key 1 and key 2.

Note	The key ID number of an authentication key on a key chain can range from 0 to 2147483647, and there is no need for the numbers to be consecutive.

Key 1 includes the authentication string “firstkey” and key 2 includes the authentication string “secondkey”.

Note	The authentication string used to authenticate sent and received EIGRP packets can contain from 1 to 80 uppercase and lowercase alphanumeric characters; the one exception is that the except that the first character cannot be a number.

For more details about the **key chain**, **key**, and **key-string (authentication)** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

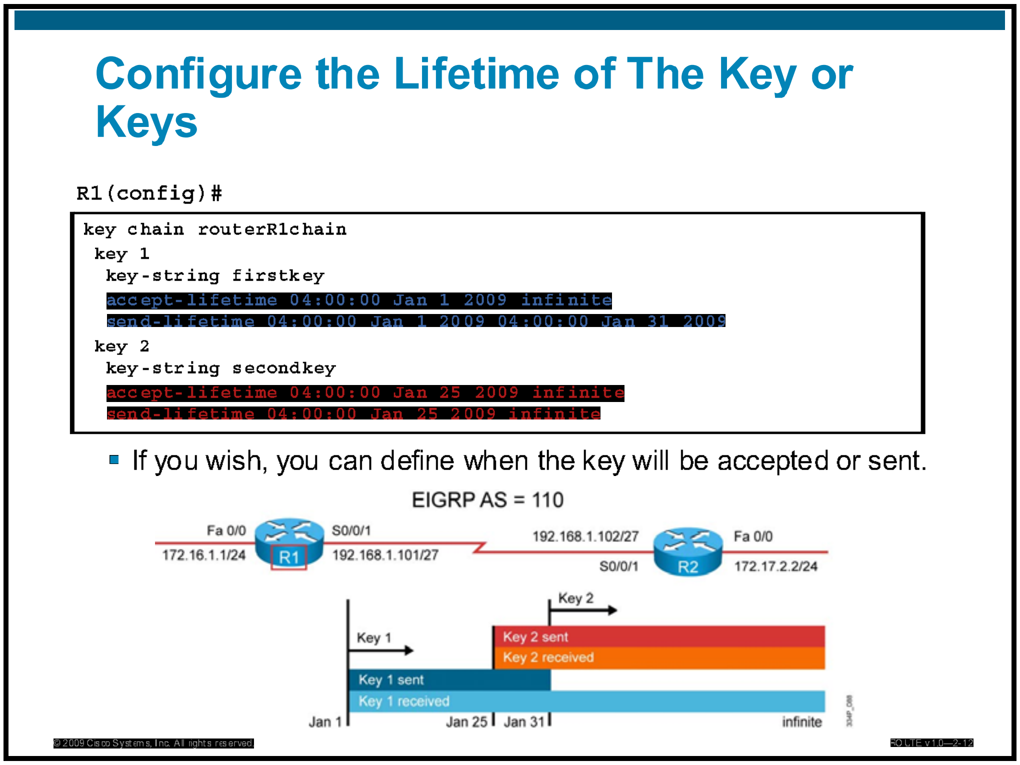

Configure the Lifetime of The Key or Keys

```
R1(config)#
```

```
key chain routerR1chain
 key 1
  key-string firstkey
  accept-lifetime 04:00:00 Jan 1 2009 infinite
  send-lifetime 04:00:00 Jan 1 2009 04:00:00 Jan 31 2009
 key 2
  key-string secondkey
  accept-lifetime 04:00:00 Jan 25 2009 infinite
  send-lifetime 04:00:00 Jan 25 2009 infinite
```

- If you wish, you can define when the key will be accepted or sent.

© 2009 Cisco Systems, Inc. All rights reserved.

ROUTE v1.0—2-12

Once you define more keys and ensure each key includes an authentication string, you can then configure a key lifetime, if you wish. Two commands can be used to define the lifetime of the key. The **accept-lifetime** command in the key chain key configuration mode is used to set the time-period during which the authentication key on a key chain is received as valid. The **send-lifetime** command in the key chain key configuration mode is used to set the time-period during which an authentication key on a key chain is valid to be sent.

In the example on this slide, router R1 is configured with the key chain **routerR1chain** and two keys are inside the key chain. Each key has an authentication string and lifetime specified. The network administrator wants to change the keys on all the routers in the network on a regular basis to improve the security, initially every month. One week is enough time to change the keys on all the routers, so the validity of key 2 is configured 1 week before the expiration of key 1 to ensure that all the routers in the network will accept the configuration.

The authentication string for key 1 is set to **firstkey** and the **accept-lifetime** and **send-lifetime** commands have been used to specify the validity of key 1. This key is acceptable for use on packets received by router R1 from January 1, 2009 onwards, as specified in the **accept-lifetime 04:00:00 Jan 1 2009 infinite** command. However, the **send-lifetime 04:00:00 Jan 1 2009 04:00:00 Jan 31 2009** command specifies that this key was valid for use when sending packets from January 1, 2009 until January 31, 2009. It will no longer be valid for use in sending packets after 4:00 a.m. on January 31, 2009.

The authentication string for key 2 is set to **secondkey.** The **accept-lifetime** and **send-lifetime** commands have been used to specify the validity of key 2, as well as for key 1. Key 2 is acceptable for use on packets received by router R1 from January 25, 2009 onwards, as specified in the **accept-lifetime 04:00:00 Jan 25 2009 infinite** command. This key can also be used when sending packets from January 25, 2009 onwards, as specified in the **send-lifetime 04:00:00 Jan 25 2009 infinite** command.

When more than one key is configured, key 1 is used first until its lifetime expires. Then the next key is used. The bar chart in this slide presents the result of the configuration. From

Starting January 1, 2009, key 1 was used for sent packets as well as for received packets. Starting January 25, 2009, key 2 was valid for sent and received packets, but it was not actually used, since key 1 was still valid. Starting January 31, 2009, key 1 was no longer valid. Immediately key 2 started to be used for all sent and received packets. From January 25, 2009 to January 31, 2009, router R1 will accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1 or a key ID equal to 2. All other MD5 packets will be dropped.

The syntax of the start time in the **accept-lifetime** and **send-lifetime** commands can be either of the following:

- *hh:mm:ss month date year*
- *hh:mm:ss date month year*

When using the **infinite** parameter to configure the lifetime, the key is valid for use on received packets from the start-time value and never expires. The default start time and the earliest acceptable date is January 1, 1993.

If an end time is specified, the key is valid for use on the received packets from the start-time value until the end-time value. The syntax is the same as the start-time value. The end-time value must be after the start-time value. The default end time is **infinite**.

The same or a similar configuration can be applied to router R2 as well. The optional lifetime parameters can be the same or different. It is important that at least one key is valid to send and at least one key is valid as a receive key. Of course both keys must be the same, otherwise it is very likely that the authentication will fail.

For more details about the **accept-lifetime** and **send-lifetime** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

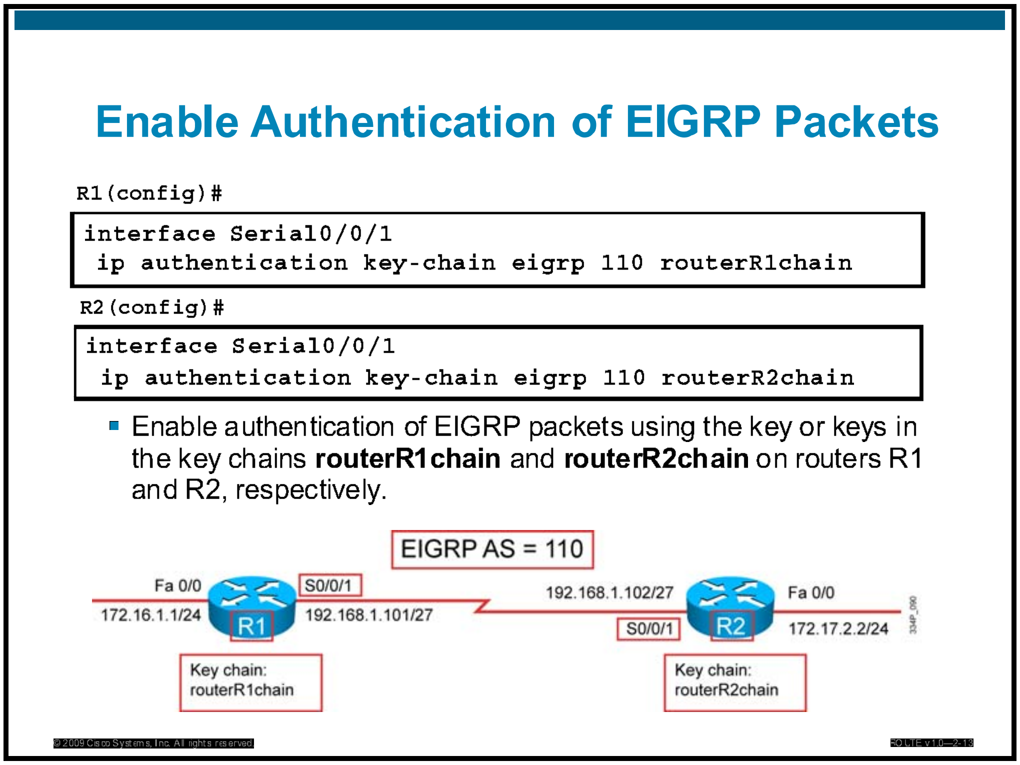

Enable Authentication of EIGRP Packets

R1(config)#

```
interface Serial0/0/1
 ip authentication key-chain eigrp 110 routerR1chain
```

R2(config)#

```
interface Serial0/0/1
 ip authentication key-chain eigrp 110 routerR2chain
```

- Enable authentication of EIGRP packets using the key or keys in the key chains **routerR1chain** and **routerR2chain** on routers R1 and R2, respectively.

ROUTE v1.0—2-13

When an authentication type is selected and a key chain is configured, then authentication of EIGRP packets must be enabled on all interfaces participating in the EIGRP domain as well. Authentication is enabled using the **ip authentication key-chain eigrp** interface command.

In the example on this slide, each interface is configured to use keys in its local key chain. Router R1 is configured to use the **routerR1chain** key chain and router R2 is configured to use the **routerR2chain** key chain. The key chain contains the list of all the available keys and must be configured separately to specify all of the keys required. The name of the key chain is of local significance and can be different on the two neighboring routers.

Note	The name of the authentication key chain is a user-defined string.

For more details about the **ip authentication key-chain eigrp** command, please check the Cisco IOS IP Routing Protocols Command Reference on the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

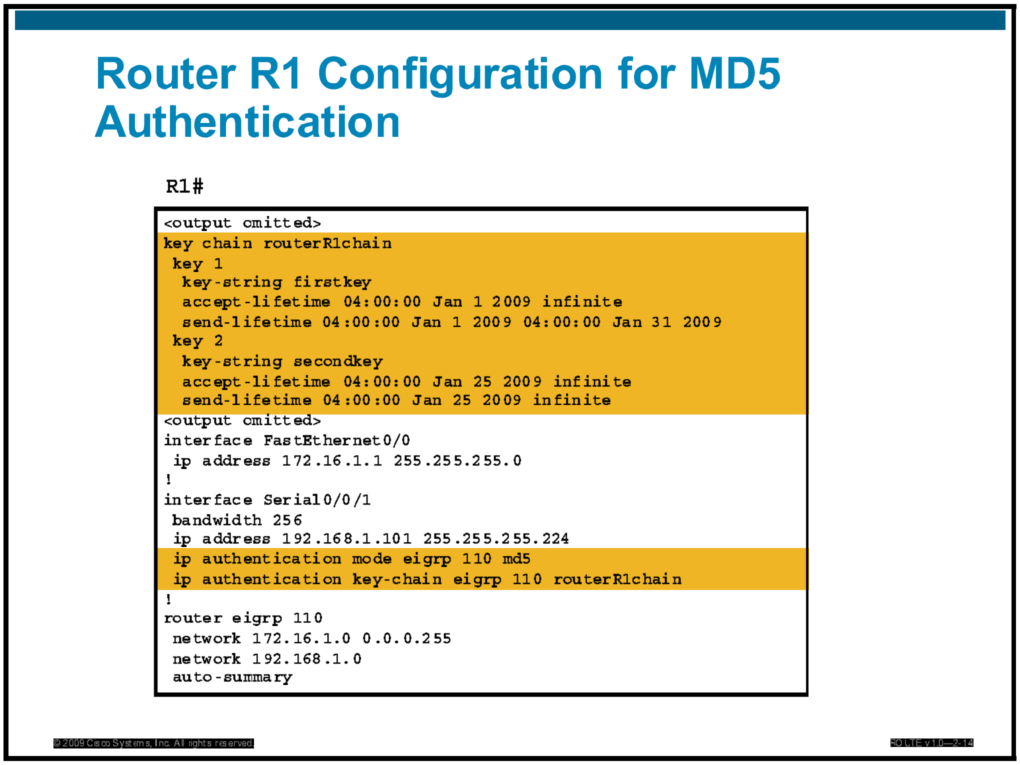

Router R1 Configuration for MD5 Authentication

R1#

```
<output omitted>
key chain routerR1chain
 key 1
  key-string firstkey
  accept-lifetime 04:00:00 Jan 1 2009 infinite
  send-lifetime 04:00:00 Jan 1 2009 04:00:00 Jan 31 2009
 key 2
  key-string secondkey
  accept-lifetime 04:00:00 Jan 25 2009 infinite
  send-lifetime 04:00:00 Jan 25 2009 infinite
<output omitted>
interface FastEthernet0/0
 ip address 172.16.1.1 255.255.255.0
!
interface Serial0/0/1
 bandwidth 256
 ip address 192.168.1.101 255.255.255.224
 ip authentication mode eigrp 110 md5
 ip authentication key-chain eigrp 110 routerR1chain
!
router eigrp 110
 network 172.16.1.0 0.0.0.255
 network 192.168.1.0
 auto-summary
```

ROUTE v1.0—2-14

The configuration of router R1 on this slide shows that MD5 authentication is configured on the serial 0/0/1 interface with the **ip authentication mode eigrp 110 md5** command. The **ip authentication key-chain eigrp 110 routerR1chain** command specifies that the key chain **routerR1chain** is to be used.

The **key chain routerR1chain** command indicates to enter configuration mode for the **routerR1chain** key chain. Two keys are defined. Key 1 is set to **firstkey** with the **key-string firstkey** command. This key is acceptable for use on packets received by router R1 from January 1, 2009 onward, as specified in the **accept-lifetime 04:00:00 Jan 1 2009 infinite** command. In contrast, the **send-lifetime 04:00:00 Jan 1 2009 04:00:00 Jan 31 2009** command specifies that this key is valid for use when sending packets from January 2, 2009 to January 31, 2009.

Key 2 is set to **secondkey** with the **key-string secondkey** command. This key is acceptable for use on packets received by router R1 from January 25, 2009 onward, as specified in the **accept-lifetime 04:00:00 Jan 25 2009 infinite** command. This key can also be used when sending packets from January 25, 2009 onward, as specified in the **send-lifetime 04:00:00 Jan 25 2009 infinite** command.

Router R1 will therefore accept and attempt to verify the MD5 digest of any EIGRP packets with a key ID equal to 1. Router R1 will also accept a packet with a key ID equal to 2. All other MD5 packets will be dropped. Router R1 will send all EIGRP packets using key 2, because key 1 is no longer valid for use when sending packets.

MD5 EIGRP authentication configuration on router R2 is very similar. The lifetimes of the keys are typically the same, and the key chain names can be different (it is of local significance in the router). The IP addresses on the interfaces are different and set according to the IP addressing scheme of router R2. In the **ip authentication key-chain eigrp** command, the configuration must refer to its local key chain name and networks advertised under the EIGRP 110 routing process must reflect the real networks used.

Verifying MD5 Authentication for EIGRP

This topic describes how to verify MD5 authentication for EIGRP.

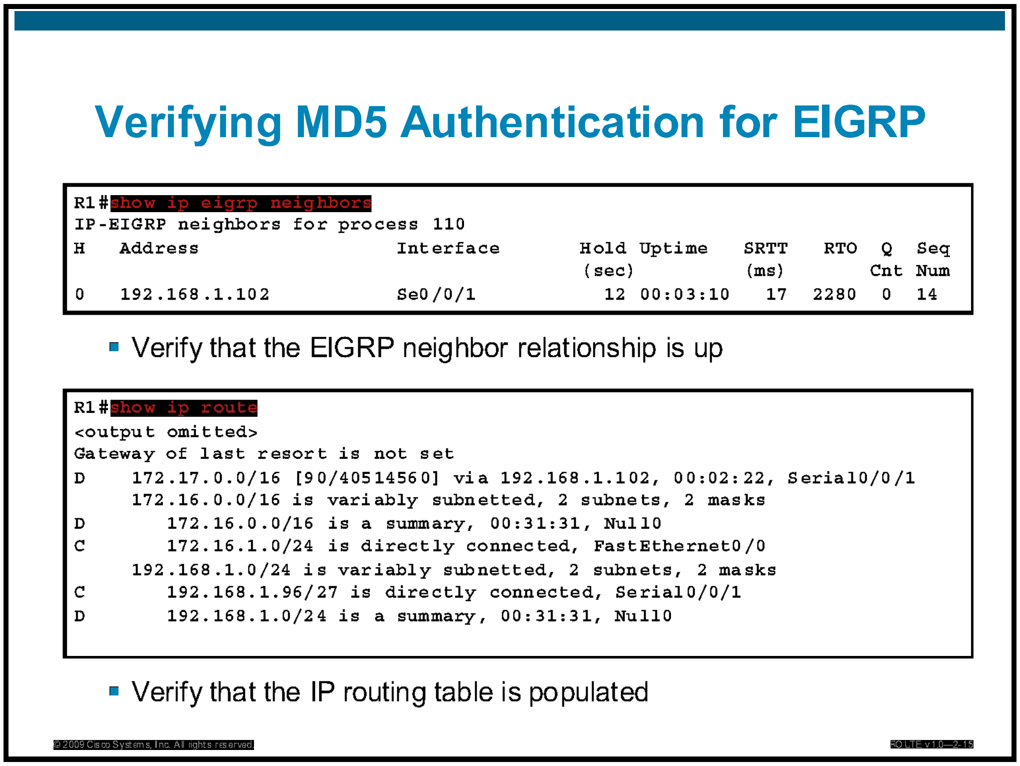

Verifying MD5 Authentication for EIGRP

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
H   Address                 Interface        Hold Uptime    SRTT   RTO  Q  Seq
                                             (sec)          (ms)       Cnt Num
0   192.168.1.102           Se0/0/1            12 00:03:10   17   2280  0  14
```

- Verify that the EIGRP neighbor relationship is up

```
R1#show ip route
<output omitted>
Gateway of last resort is not set
D    172.17.0.0/16 [90/40514560] via 192.168.1.102, 00:02:22, Serial0/0/1
     172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
D       172.16.0.0/16 is a summary, 00:31:31, Null0
C       172.16.1.0/24 is directly connected, FastEthernet0/0
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C       192.168.1.96/27 is directly connected, Serial0/0/1
D       192.168.1.0/24 is a summary, 00:31:31, Null0
```

- Verify that the IP routing table is populated

ROUTE v1.0—2-15

If authentication is not successful, then routers will not process EIGRP packets and will not form neighbor relationships. Also, routers will not build the EIGRP tables and populate the IP routing table with EIGRP routes.

The output on the figure on this slide shows two commands that can be used to verify the EIGRP neighbors and IP routing table.

The **show ip eigrp neighbors** verification command shows the EIGRP neighbors table on router R1, which indicates that the two routers have successfully formed an EIGRP adjacency.

The **show ip route** verification command on router R1 shows that the 172.17.0.0 network has been learned via EIGRP over the serial connection, which proves that the authentication must have been successful.

For more details about the **show ip eigrp neighbors** and **show ip route** commands, please check the Cisco IOS IP Routing Protocols Command Reference via the following link: http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

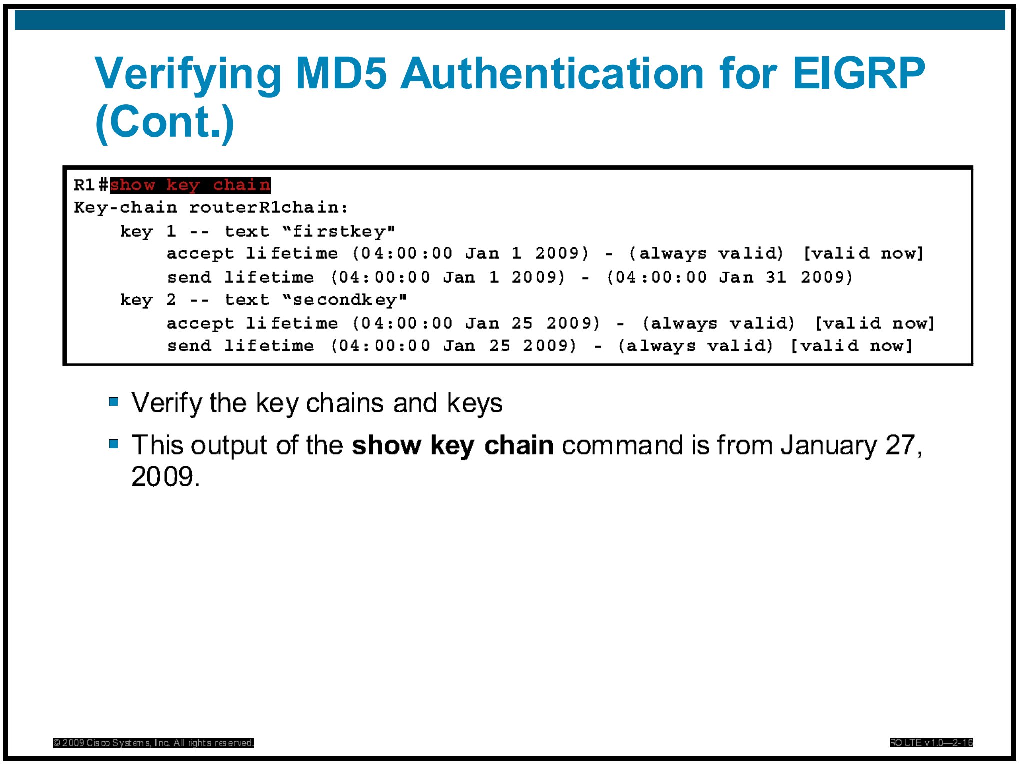

Verifying MD5 Authentication for EIGRP (Cont.)

```
R1#show key chain
Key-chain routerR1chain:
    key 1 -- text "firstkey"
        accept lifetime (04:00:00 Jan 1 2009) - (always valid) [valid now]
        send lifetime (04:00:00 Jan 1 2009) - (04:00:00 Jan 31 2009)
    key 2 -- text "secondkey"
        accept lifetime (04:00:00 Jan 25 2009) - (always valid) [valid now]
        send lifetime (04:00:00 Jan 25 2009) - (always valid) [valid now]
```

- Verify the key chains and keys
- This output of the **show key chain** command is from January 27, 2009.

You can use the **show key chain** verification command to see the key chain, key string, and the lifetime of the keys configured under the key chain. The keys must be the same on both neighbors and the lifetime must be set properly.

The output of the **show key chain** command on router R1 is shown on this slide. Key chain **routerR1chain** and both key 1 (with authentication string **firstkey)** and key 2 (with authentication string **secondkey)** are shown in the sample output. Under each key, the lifetime of the key is shown as well. You can verify the configuration by checking that the output is the same from the neighboring router (router R2 in our topology).

For more details about the **show key chain** command, please check the Cisco IOS IP routing protocols command reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

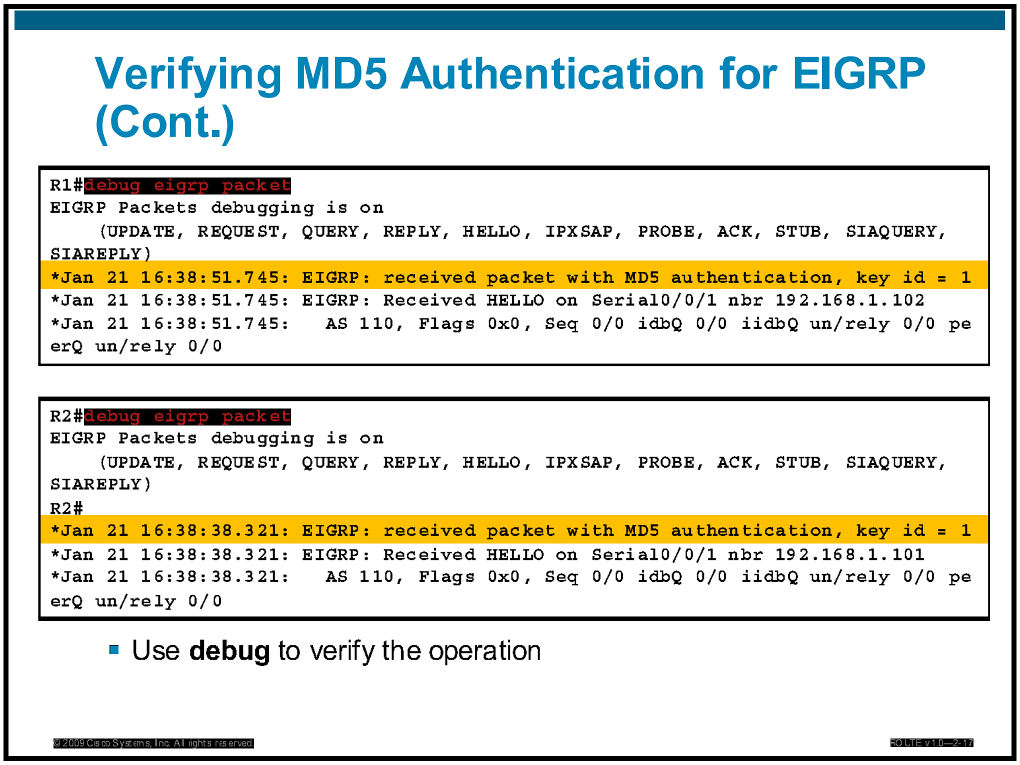

Verifying MD5 Authentication for EIGRP (Cont.)

```
R1#debug eigrp packet
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)
*Jan 21 16:38:51.745: EIGRP: received packet with MD5 authentication, key id = 1
*Jan 21 16:38:51.745: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.102
*Jan 21 16:38:51.745:   AS 110, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 pe
erQ un/rely 0/0
```

```
R2#debug eigrp packet
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)
R2#
*Jan 21 16:38:38.321: EIGRP: received packet with MD5 authentication, key id = 1
*Jan 21 16:38:38.321: EIGRP: Received HELLO on Serial0/0/1 nbr 192.168.1.101
*Jan 21 16:38:38.321:   AS 110, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0 pe
erQ un/rely 0/0
```

- Use **debug** to verify the operation

ROUTE v1.0—2-17

Use the **debug eigrp packet** command to display general debugging information. If a communication session is closing when it should not be, an end-to-end connection problem could be the cause. The **debug eigrp packet** command is useful for analyzing the messages traveling between the local and remote hosts, including authentication messages.

The sample output on this slide shows successful MD5 authentication. The output of the **debug eigrp packet** command on router R1 shows that router R1 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 1, from router R2.

Similarly, the output of the **debug eigrp packet** command on router R2 shows that router R2 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from router R1.

For more details about the **debug eigrp packet** command, please check the Cisco IOS debug command reference via the following link: http://www.cisco.com/en/US/docs/ios/debug/command/reference/db_book.html

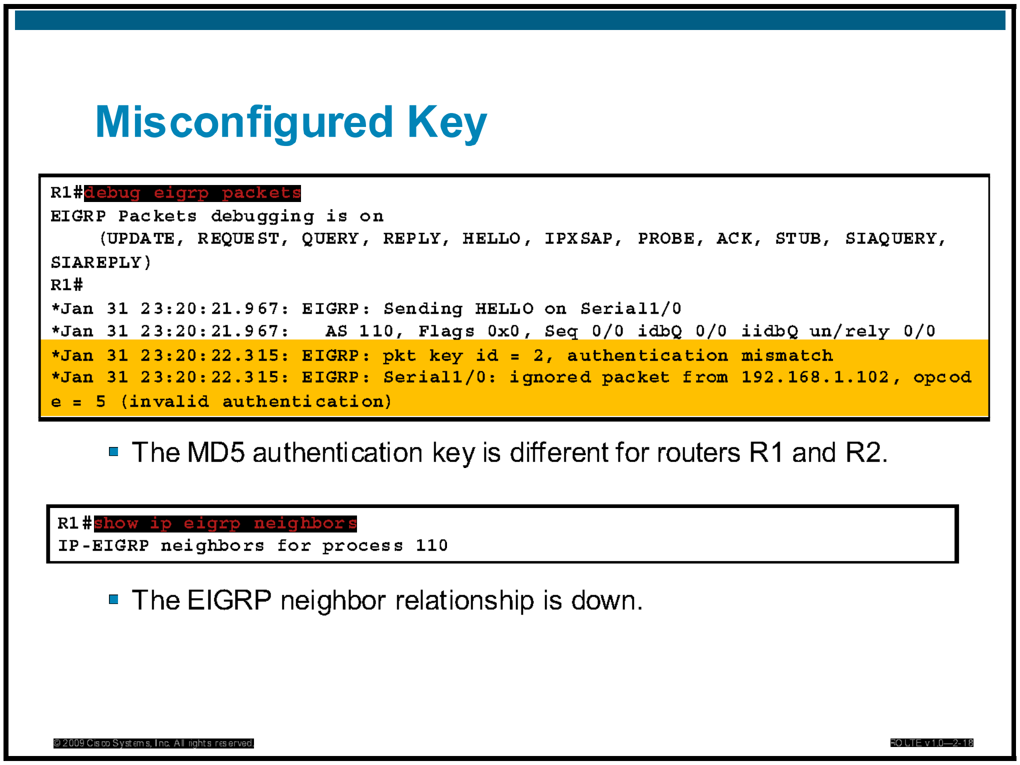

Misconfigured Key

```
R1#debug eigrp packets
EIGRP Packets debugging is on
    (UPDATE, REQUEST, QUERY, REPLY, HELLO, IPXSAP, PROBE, ACK, STUB, SIAQUERY,
SIAREPLY)
R1#
*Jan 31 23:20:21.967: EIGRP: Sending HELLO on Serial1/0
*Jan 31 23:20:21.967:   AS 110, Flags 0x0, Seq 0/0 idbQ 0/0 iidbQ un/rely 0/0
*Jan 31 23:20:22.315: EIGRP: pkt key id = 2, authentication mismatch
*Jan 31 23:20:22.315: EIGRP: Serial1/0: ignored packet from 192.168.1.102, opcod
e = 5 (invalid authentication)
```

- The MD5 authentication key is different for routers R1 and R2.

```
R1#show ip eigrp neighbors
IP-EIGRP neighbors for process 110
```

- The EIGRP neighbor relationship is down.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-18

The sample output on this slide shows the MD5 authentication problems. The output of the **debug eigrp packet** command on router R1 shows that router R1 is receiving EIGRP packets with MD5 authentication, with a key ID equal to 2, from router R2, but there is an authentication mismatch. The EIGRP packets from router R2 are ignored and the neighbor relationship is declared to be down. The output of the **show ip eigrp neighbors** command confirms that router R1 does not have any EIGRP neighbors.

The two routers keep trying to reestablish their neighbor relationship using key 2. Because of the different key strings used by each router in this scenario, router R1 will authenticate hello messages sent by router R2 using the key string **secondkey**. However, when router R1 sends a hello message back to router R2 using a different key string, an authentication mismatch occurs.

Summary

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

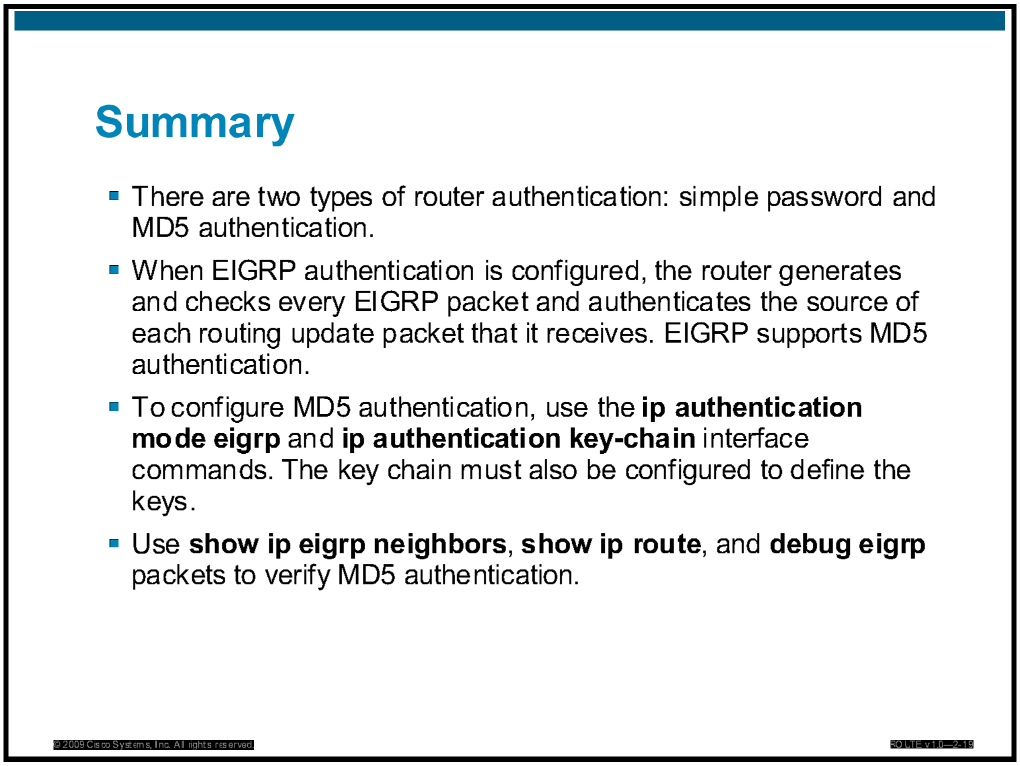

Summary

- There are two types of router authentication: simple password and MD5 authentication.
- When EIGRP authentication is configured, the router generates and checks every EIGRP packet and authenticates the source of each routing update packet that it receives. EIGRP supports MD5 authentication.
- To configure MD5 authentication, use the **ip authentication mode eigrp** and **ip authentication key-chain** interface commands. The key chain must also be configured to define the keys.
- Use **show ip eigrp neighbors**, **show ip route**, and **debug eigrp** packets to verify MD5 authentication.

Lesson 7

Lab 2-3 Debrief

Overview

In Lab 2-3, students configure and verify EIGRP authentication. They also implement EIGRP authentication over LAN and WAN interfaces.

Objectives

Upon completing this lesson, you will be able to configure and verify EIGRP authentication. This ability includes being able to meet these objectives:

- Complete the lab overview and verification
- Describe a sample solution and alternatives

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and to start with the verification.

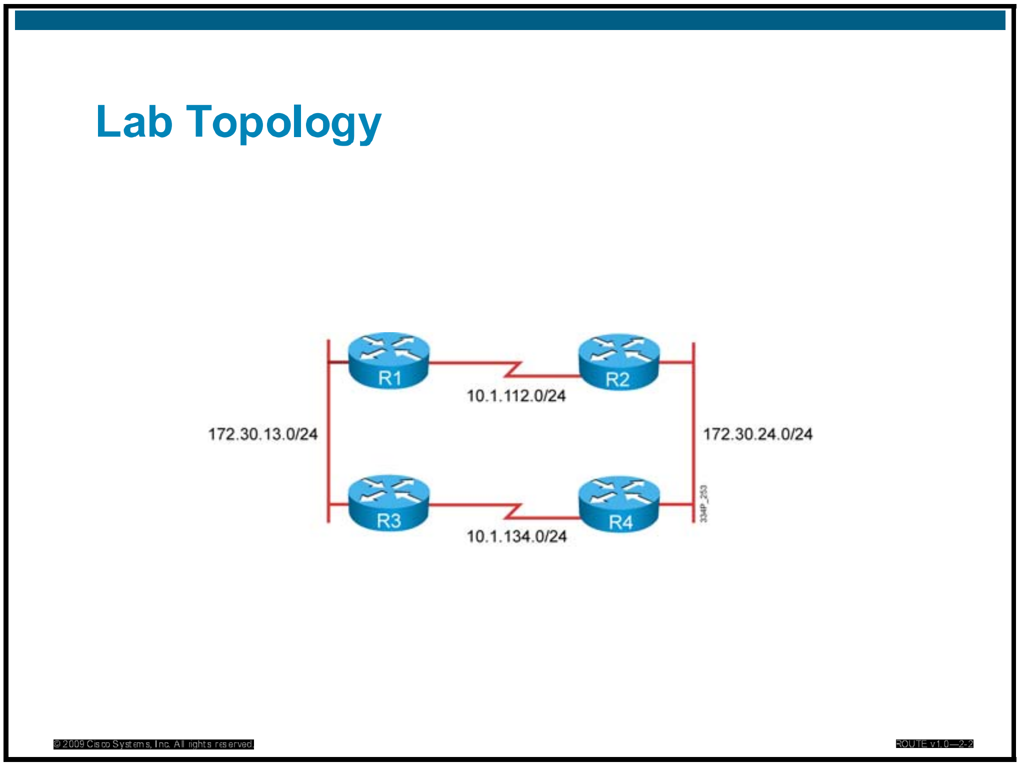

Lab Topology

ROUTE v1.0—2-2

The figure above presents the physical lab topology used for the Lab 2-3: Configure and Verify EIGRP Authentication. The topology uses four pod routers. All routers participate in the EIGRP routing protocol.

Based on the topology, students will identify the required parameters for configuring EIGRP authentication on LAN and WAN interfaces.

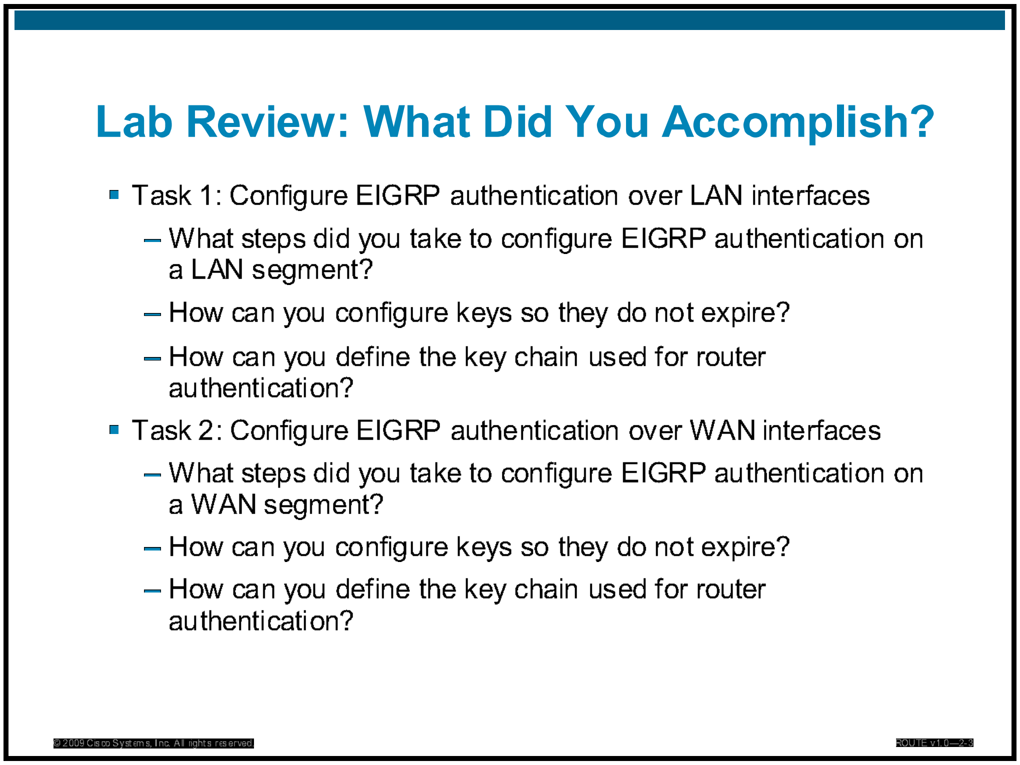

Lab Review: What Did You Accomplish?

- Task 1: Configure EIGRP authentication over LAN interfaces
  - What steps did you take to configure EIGRP authentication on a LAN segment?
  - How can you configure keys so they do not expire?
  - How can you define the key chain used for router authentication?
- Task 2: Configure EIGRP authentication over WAN interfaces
  - What steps did you take to configure EIGRP authentication on a WAN segment?
  - How can you configure keys so they do not expire?
  - How can you define the key chain used for router authentication?

In the first task, you configured EIGRP authentication over LAN interfaces. You configured a key chain with a key that never expires. You enabled secure authentication and used the defined key chain to provide security when exchanging EIGRP packets.

In the second task, you configured EIGRP authentication over WAN interfaces. You configured a second key chain with a key that never expires. And again, you enabled secure authentication using the defined key chain.

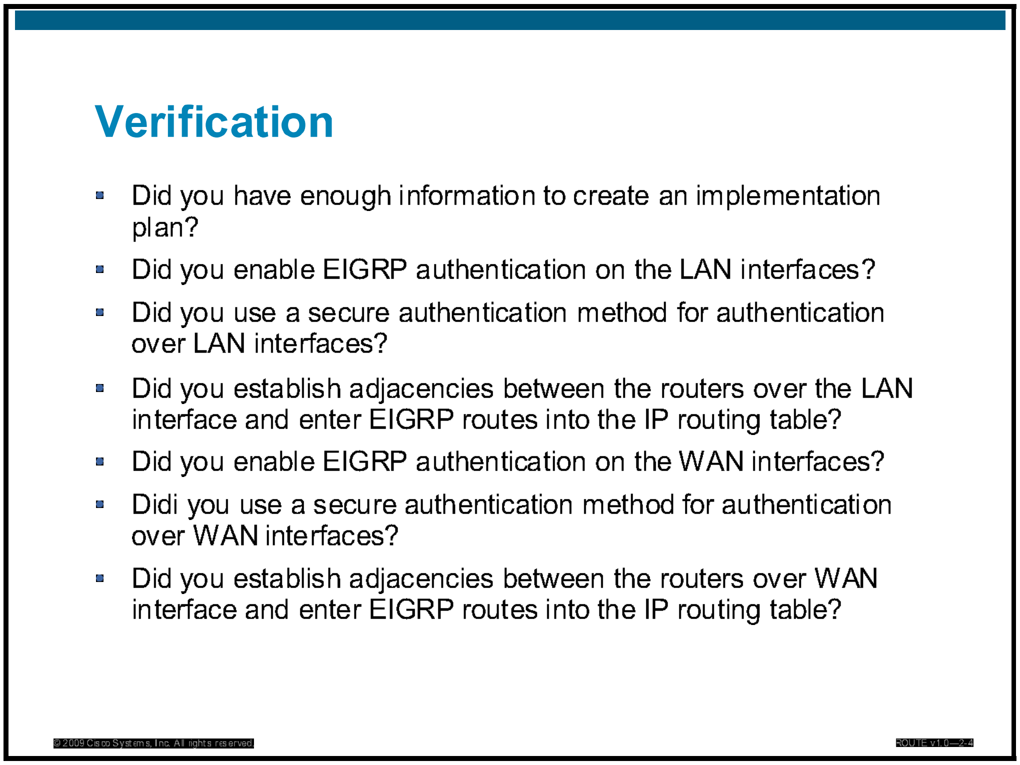

Verification

- Did you have enough information to create an implementation plan?
- Did you enable EIGRP authentication on the LAN interfaces?
- Did you use a secure authentication method for authentication over LAN interfaces?
- Did you establish adjacencies between the routers over the LAN interface and enter EIGRP routes into the IP routing table?
- Did you enable EIGRP authentication on the WAN interfaces?
- Didi you use a secure authentication method for authentication over WAN interfaces?
- Did you establish adjacencies between the routers over WAN interface and enter EIGRP routes into the IP routing table?

ROUTE v1.0—2-4

A common approach to verifying the implementation process for a routing protocol is to answer the following questions:

- Did you have enough information to create an implementation plan?
- Did you enable EIGRP authentication on the LAN interfaces?
- Did you use a secure authentication method for authentication over LAN interfaces?
- Did you establish adjacencies between the routers over the LAN interface and enter EIGRP routes into the IP routing table?
- Did you enable EIGRP authentication on the WAN interfaces?
- Did you use a secure authentication method for authentication over WAN interfaces?
- Did you establish adjacencies between the routers over WAN interface and enter EIGRP routes into the IP routing table?

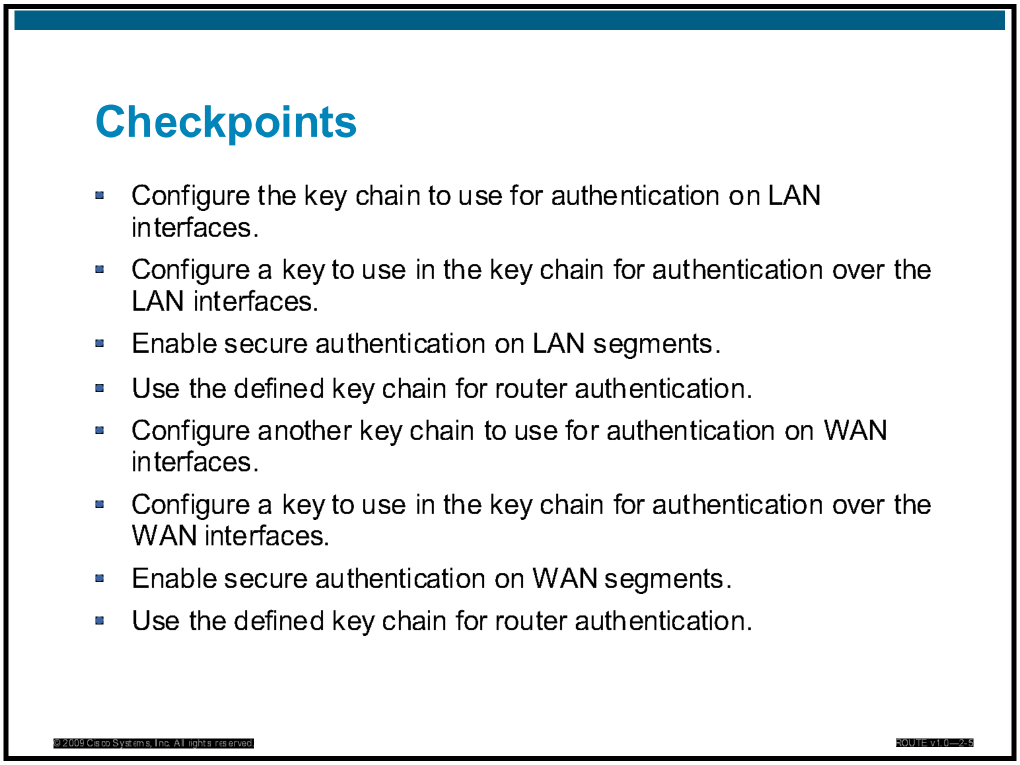

Checkpoints

- Configure the key chain to use for authentication on LAN interfaces.
- Configure a key to use in the key chain for authentication over the LAN interfaces.
- Enable secure authentication on LAN segments.
- Use the defined key chain for router authentication.
- Configure another key chain to use for authentication on WAN interfaces.
- Configure a key to use in the key chain for authentication over the WAN interfaces.
- Enable secure authentication on WAN segments.
- Use the defined key chain for router authentication.

During the configuration and verification phase, the network operator can deal with several checkpoints. After completing all configuration tasks, the network operator has either finished the lab successfully or must perform additional verification and troubleshooting.

With different checkpoints, the network operator can verify for proper configuration. The following checkpoints are used for verification:

- Configure the key chain to use for authentication on LAN interfaces.
- Configure a key to use in the key chain for authentication over the LAN interfaces.
- Enable secure authentication on LAN segments.
- Use the defined key chain for router authentication.
- Configure another key chain to use for authentication on WAN interfaces.
- Configure a key to use in the key chain for authentication over the WAN interfaces.
- Enable secure authentication on WAN segments.
- Use the defined key chain for router authentication.

Sample Solutions and Alternatives

This topic describes a sample solution and other alternatives.

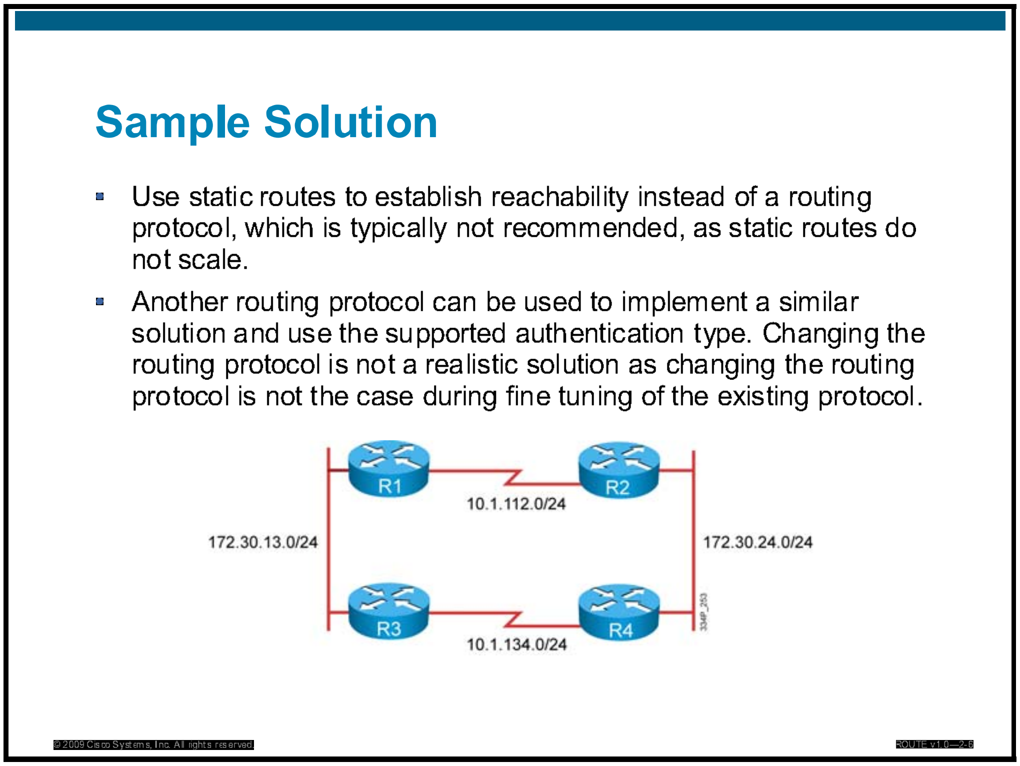

Sample Solution

- Use static routes to establish reachability instead of a routing protocol, which is typically not recommended, as static routes do not scale.
- Another routing protocol can be used to implement a similar solution and use the supported authentication type. Changing the routing protocol is not a realistic solution as changing the routing protocol is not the case during fine tuning of the existing protocol.

ROUTE v1.0—2-6

The sample solution includes implementation details and details for each task of the implementation plan. Different solutions are possible and this slide shows a few details of a successful configuration.

Proper implementation includes the following items:

- Configure a key chain with the key used for LAN authentication and use it for router authentication on LAN segments.
- Configure another key chain with the key used for WAN authentication and use it for router authentication on WAN segments.

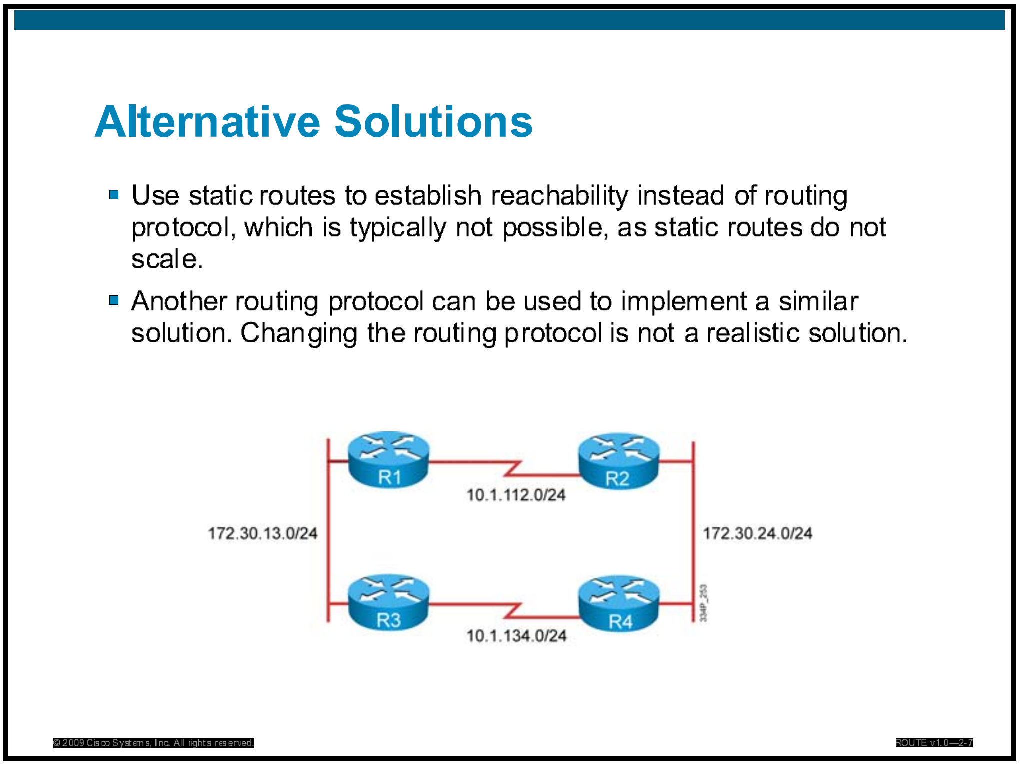

Alternative Solutions

- Use static routes to establish reachability instead of routing protocol, which is typically not possible, as static routes do not scale.
- Another routing protocol can be used to implement a similar solution. Changing the routing protocol is not a realistic solution.

ROUTE v1.0—2-7

Use static routes to establish reachability instead of a routing protocol, which is typically not recommended, as static routes do not scale.

Another routing protocol can be used to implement a similar solution and use the supported authentication type. Changing the routing protocol is not a realistic solution as changing the routing protocol is not the case during fine tuning of the existing protocol.

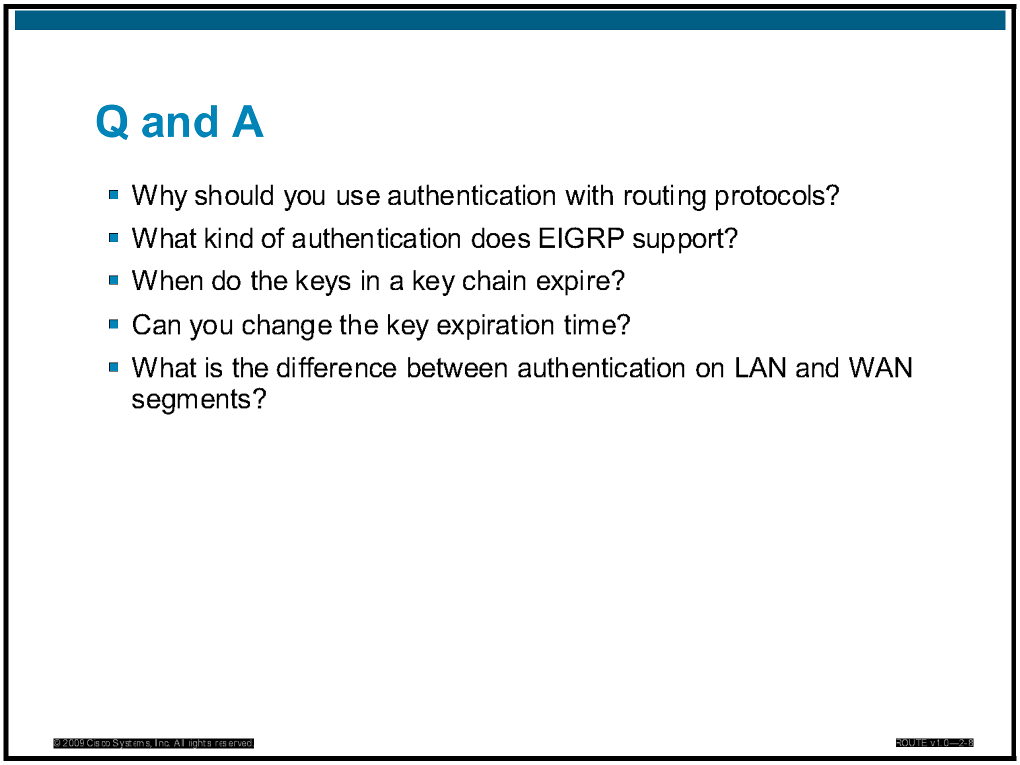

Q and A

- Why should you use authentication with routing protocols?
- What kind of authentication does EIGRP support?
- When do the keys in a key chain expire?
- Can you change the key expiration time?
- What is the difference between authentication on LAN and WAN segments?

ROUTE v1.0—2-8

Authentication provides additional security in networks by verifying the source and destination of each routing update. Only routers with the correct authentication configured can exchange routing protocol packets.

EIGRP supports MD5 authentication.

The expiration time can be configured for each key and the key can be configured not to expire.

There is no difference between LAN and WAN authentication. In both cases, you must enable MD5 authentication, configure a key chain, and use the keys in the key chain properly.

Summary

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

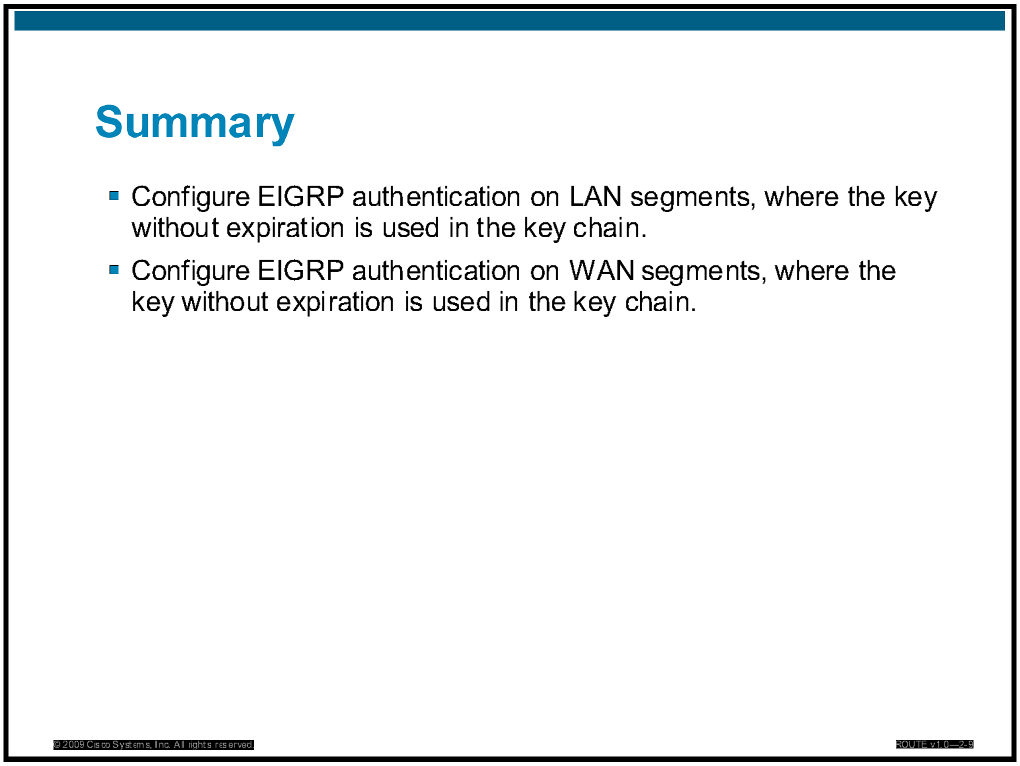

Summary

- Configure EIGRP authentication on LAN segments, where the key without expiration is used in the key chain.
- Configure EIGRP authentication on WAN segments, where the key without expiration is used in the key chain.

Lesson 8

Advanced EIGRP Features in an Enterprise Network

Overview

Network administrators benefit from understanding how to configure Enhanced Interior Gateway Routing Protocol (EIGRP) to prevent common routing problems that hinder network scalability. For example, you can implement EIGRP stub routers to limit the EIGRP query range, making EIGRP more scalable with fewer complications.

EIGRP is a scalable routing protocol, which ensures that as a network grows larger, it operates efficiently and adjusts rapidly to changes. This lesson describes advanced EIGRP features and practical EIGRP-specific design and configuration techniques to implement an effective, scalable network.

Objectives

Upon completing this lesson, you will be able to implement advanced EIGRP features in an Enterprise Network by applying the planned implementation processes using correct Cisco IOS commands and applications. You will also be able to verify that the configuration was correctly implemented. This ability includes being able to meet these objectives:

- Define scalability in large networks.
- Understand EIGRP queries.
- Define SIA connections in EIGRP.
- Understand EIGRP stub routers.

Scalability in Large Networks

This topic explains factors affecting scalability in large internetworks.

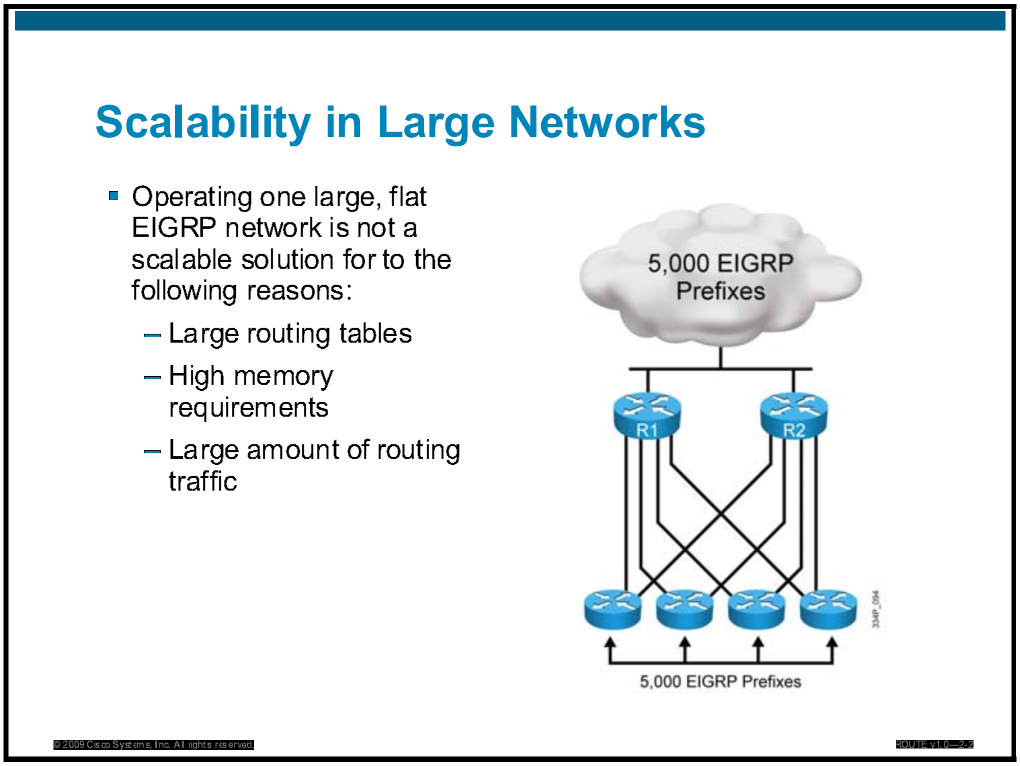

Scalability in Large Networks

- Operating one large, flat EIGRP network is not a scalable solution for to the following reasons:
  - Large routing tables
  - High memory requirements
  - Large amount of routing traffic

ROUTE v1.0—2-2

Large, flat EIGRP networks are normally not scalable for two main reasons:

- High memory demands can lead to problems. The problems result from having a large topology table, having a large number of routes in a routing table, and, in some environments (such as a concentration of routers at a central site), from having a large number of neighbors in an adjacency table.
- High bandwidth demands can also create problems, resulting from the exchange of routing updates. The sending of queries and replies in one large EIGRP domain (which may include links with low bandwidth and links with a significant number of transmission errors) often results in a large amount of routing traffic, which consequently results in even more traffic and congestion.

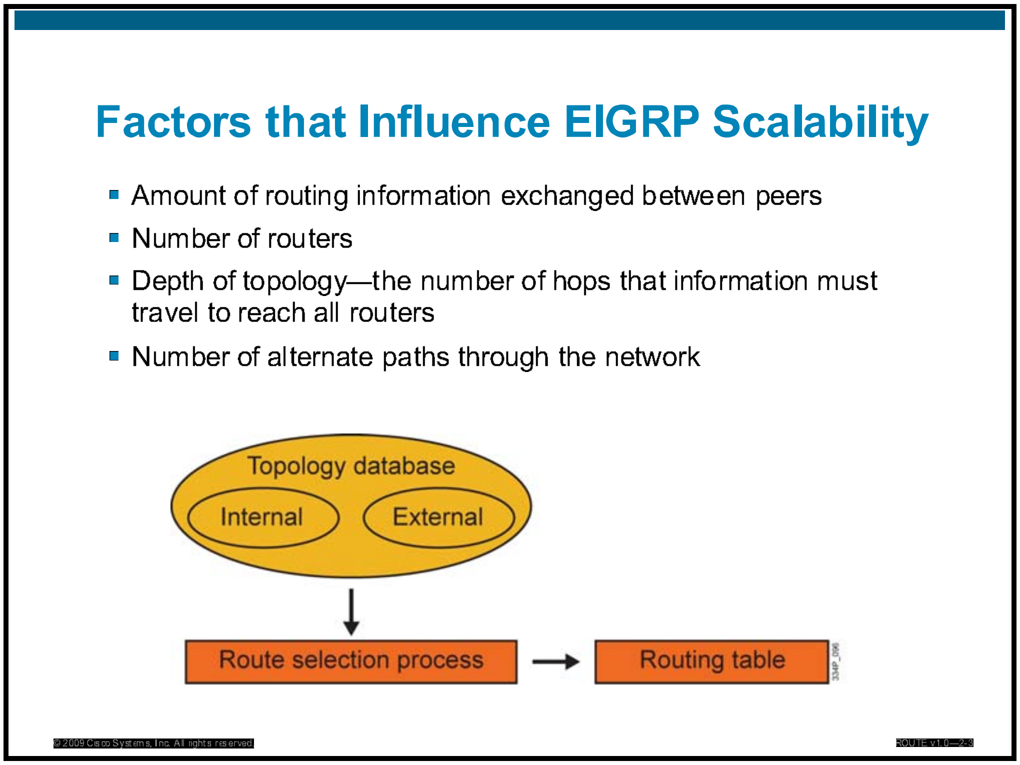

Factors that Influence EIGRP Scalability

- Amount of routing information exchanged between peers
- Number of routers
- Depth of topology—the number of hops that information must travel to reach all routers
- Number of alternate paths through the network

The following are some of the factors that affect network scalability:

- **Amount of information exchanged between neighbors:** If more routing information than is necessary is exchanged between EIGRP neighbors, the routers have to work harder both at neighbor startup and when reacting to changes in the network. Route summarization is needed to improve the convergence time.

- **Number of routers:** When a topology change occurs in the network, EIGRP resource consumption is directly related to the number of routers that are involved in the change.

- **Depth of the topology:** The topology depth can affect the convergence time. Depth refers to the number of hops that information must travel to reach all routers. A multinational network without route summarization is an example of a network with large depth and, therefore, higher convergence times. A three-tiered network design (as described in Module 1) is highly recommended for all IP routing environments. There should never be more than seven hops between any two routing devices on an internetwork. The propagation delay and query process across multiple hops when changes occur may slow network convergence.

- **Number of alternate paths through the network:** A network should provide alternate paths to avoid single points of failure. However, too much complexity (too many alternate paths) can also lead to EIGRP convergence problems, because the EIGRP routing process needs to explore all possible paths for lost routes (using queries). This complexity creates the ideal condition for a router to become stuck-in-active (SIA) as it awaits a response to queries that are being propagated through many alternate paths.

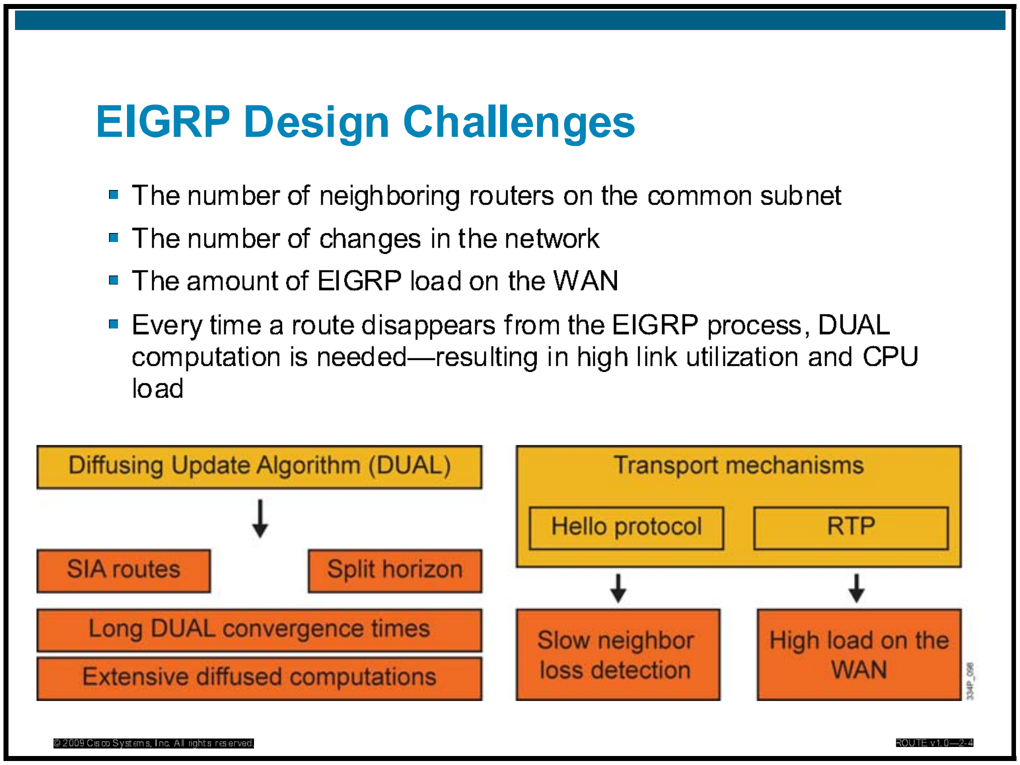

EIGRP Design Challenges

- The number of neighboring routers on the common subnet
- The number of changes in the network
- The amount of EIGRP load on the WAN
- Every time a route disappears from the EIGRP process, DUAL computation is needed—resulting in high link utilization and CPU load

Diffusing Update Algorithm (DUAL)

SIA routes

Split horizon

Long DUAL convergence times

Extensive diffused computations

Transport mechanisms

Hello protocol

RTP

Slow neighbor loss detection

High load on the WAN

ROUTE v1.0—2-4

When you implement EIGRP as the routing protocol, you need to address some design challenges. The three major factors are:

- The size of the topology and routing tables, including the number of neighboring routers on the common subnet
- The number of changes in the network
- The amount of EIGRP load on the WAN

These three factors mainly dictate the EIGRP design and introduce the need for query boundaries (using summarization, redistribution, and so on). Without any boundaries, queries are propagated throughout the EIGRP domain and, often, all of the routers get involved in a Diffusing Update Algorithm (DUAL) computation. This EIGRP feature not only places additional bandwidth demands on links in the network, but also results in high CPU utilization. Frequent DUAL computations have an effect on all tables, which are maintained by the routers—from EIGRP structures to various caches built during the forwarding process.

EIGRP Queries

This topic explains how EIGRP uses queries to converge rapidly when a route is lost.

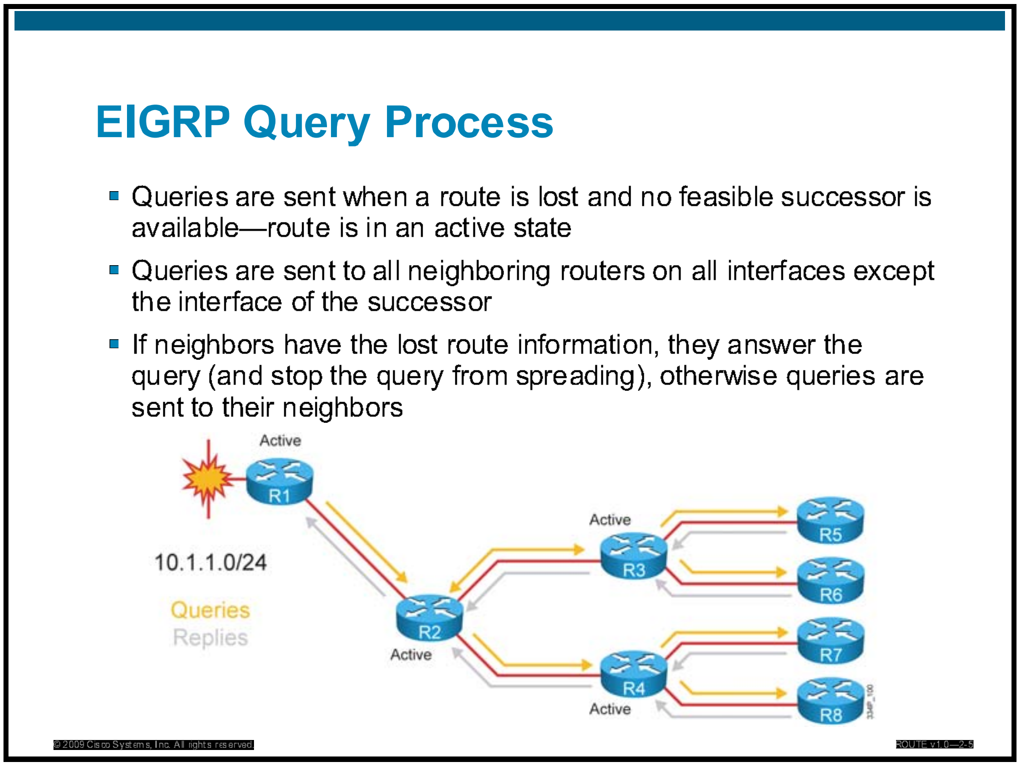

EIGRP Query Process

- Queries are sent when a route is lost and no feasible successor is available—route is in an active state
- Queries are sent to all neighboring routers on all interfaces except the interface of the successor
- If neighbors have the lost route information, they answer the query (and stop the query from spreading), otherwise queries are sent to their neighbors

ROUTE v1.0—2-5

As an advanced distance vector protocol, EIGRP relies on neighboring routers to provide routing information. Recall that when a router loses a route and does not have a feasible successor in its topology table, it looks for an alternative path to the destination. This is known as *going active* on a route; a route is considered passive when a router is not performing recompilation on that route.

When the route is lost, the router sends query packets to all neighbors on interfaces other than the one used to reach the previous successor (split horizon behavior). These packets query if each of the neighbors has a route to the given destination. If a router has an alternate route, it answers the query and does not propagate it further. If a neighbor does not have an alternate route, it queries each of its own neighbors for an alternate path. The queries then propagate through the network, creating an expanding tree of queries. When a router answers a query, it stops the spread of the query through that branch of the network.

The figure in the slide presents a network example for which a single lost route might result in an enormous number of queries sent throughout the EIGRP domain. When the route to network 10.1.1.0 on router R1 is lost, router R1 sends a query to all neighboring routers and to all interfaces except the interface of the successor (split horizon). The query is propagated to router R2. Since it has no information about the lost route, router R2 cascades the query to its neighbors, which cascade it to their neighbors, and so on. Each query requires a reply from the neighbor and the amount of traffic increases. The network topology, in the figure in the slide, shows that there is no redundant path to network 10.1.1.0 available. The EIGRP query propagation process is far from efficient. Many queries are sent and each query is followed by a reply. Several solutions exist to optimize the query propagation process and to limit the amount of unnecessary EIGRP load on the links. The solutions that can be used are summarization, redistribution, and the EIGRP stub routing feature.

SIA Connections in EIGRP

Stuck-in-active (SIA) routes can be some of the most challenging problems to resolve in an EIGRP network. This topic explains why SIA connections occur.

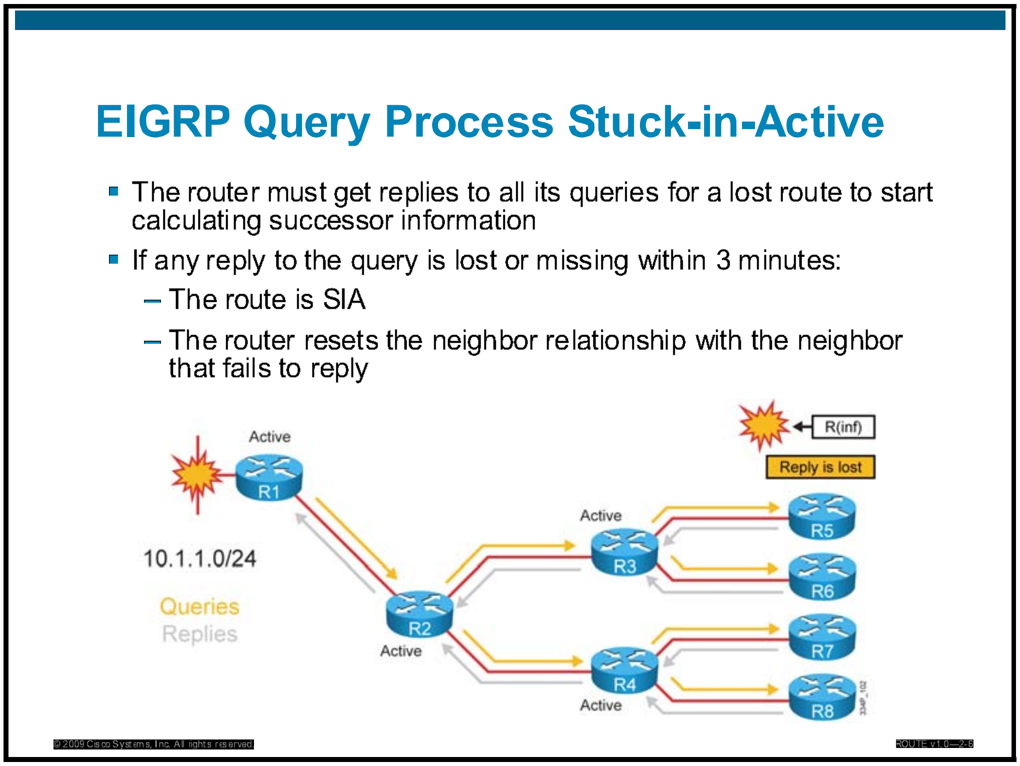

EIGRP Query Process Stuck-in-Active

- The router must get replies to all its queries for a lost route to start calculating successor information
- If any reply to the query is lost or missing within 3 minutes:
  - The route is SIA
  - The router resets the neighbor relationship with the neighbor that fails to reply

ROUTE v1.0—2-6

EIGRP uses a reliable multicast approach to search for an alternate to a lost route; therefore, it is imperative that EIGRP receive a reply for each query it generates in the network.

Once a route goes active and the query sequence is initiated, it can only come out of the active state and transition to the passive state when it receives a reply for every generated query. If the router does not receive a reply to all the outstanding queries within three minutes (the default time), the route goes into the SIA state.

When the route goes into the SIA state, the querying router resets the neighbor relationship to the neighbor that failed to reply. This setting causes the router to go active on all routes known through the lost neighbor, and to re-advertise all the routes that it knows about to the lost neighbor.

The most common reasons for SIA routes are as follows:

- The router being queried is too busy to answer the query because of high CPU usage or memory problems and cannot allocate the memory to process the query or build the reply packet.
- The link between the two routers is not good; therefore, some packets are lost between the routers. While the receiving router receives enough packets to maintain the neighbor relationship, the router does not receive all of the queries or replies.
- A failure causes traffic on a link to flow in one direction only—this is called a unidirectional link.

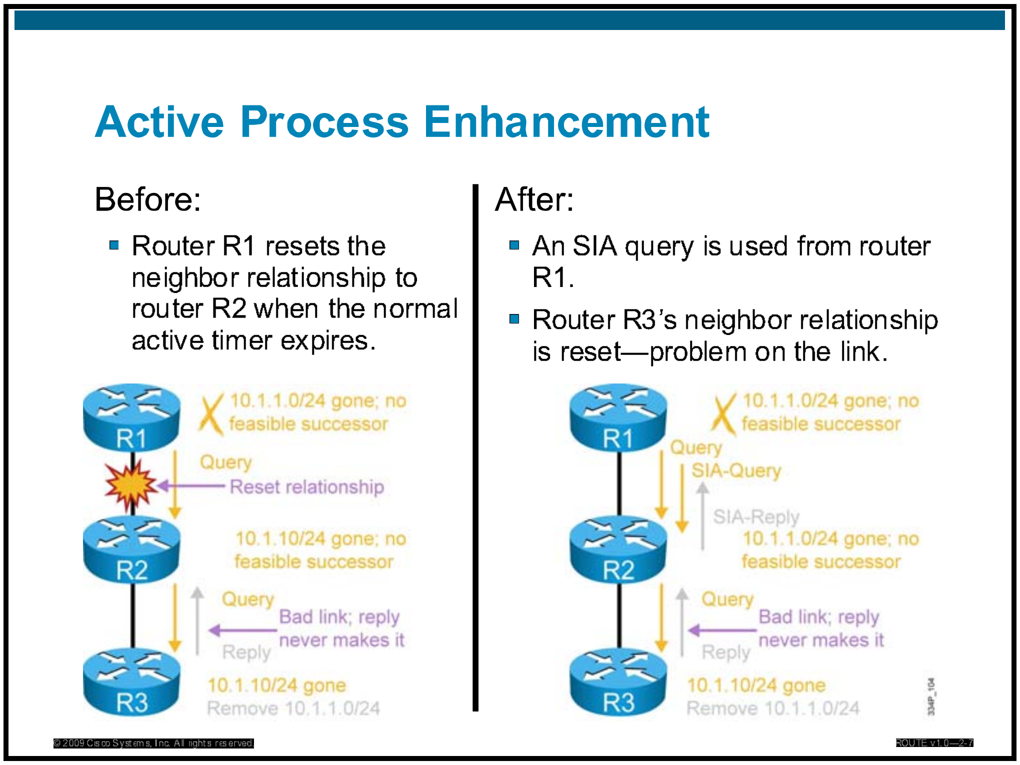

Active Process Enhancement

Before:

- Router R1 resets the neighbor relationship to router R2 when the normal active timer expires.

After:

- An SIA query is used from router R1.
- Router R3's neighbor relationship is reset—problem on the link.

R1
10.1.1.0/24 gone; no feasible successor
Query
Reset relationship
R2
10.1.10/24 gone; no feasible successor
Query
Bad link; reply never makes it
Reply
R3
10.1.10/24 gone
Remove 10.1.1.0/24
R1
10.1.1.0/24 gone; no feasible successor
Query
SIA-Query
SIA-Reply
R2
10.1.1.0/24 gone; no feasible successor
Query
Bad link; reply never makes it
Reply
R3
10.1.10/24 gone
Remove 10.1.1.0/24
334P_104

ROUTE v1.0—2-7

A route becomes active when it goes down or its metric become worse and there are no feasible successors. It sends a query to all its neighbors asking for a new path to the lost route. The process requires replies from all of the neighbors. If replies are lost, then two new messages are required.

SIA query and SIA reply are two new additions to the type, length, value (TLV) triplets in the EIGRP packet header. These packets are generated automatically since Cisco IOS Release 12.1(5) with the Active Process Enhancement feature. This feature enables an EIGRP router to monitor the search progression for a successor route and ensures that the neighbor is still reachable. The result is improved network reliability by reducing the unintended termination of neighbor adjacency.

Before a SIA query and SIA reply were available, the following would occur:

1. Router R1 sends a query for network 10.1.1.0/24 to router R2.

2. Router R2 has no entry for this network, so it queries router R3. If problems exist between router R2 and router R3, the reply packet from router R3 to router R2 may be delayed or lost.

3. Router R1 has no visibility of the downstream progress and assumes that no response indicates problems with router R2. After router R1's three-minute active timer expires, the neighbor relationship with router R2 is reset, along with all known routes from router R2.

With the Active Process Enhancement feature, the events take a different course.

1. Router R1 queries downstream router R2 (with an SIA query) at the midway point of the active timer (one and a half minutes by default) about the status of the route.

2. Router R2 responds (with an SIA reply) that it is searching for a replacement route.

3. Upon receiving this SIA reply response packet, router R1 validates the status of router R2 and does not terminate the neighbor relationship.

4. Meanwhile, router R2 will send up to three SIA queries to router R3. If they go unanswered, router R2 will terminate the neighbor relationship with router R3. Router R2 will then update router R1 with an SIA reply indicating that the network 10.1.1.0/24 is unreachable.

5. Routers R1 and R2 will remove the active route from their topology tables. The neighbor relationship between routers R1 and R2 remains intact.

EIGRP Stub Routers

The stability of large-scale EIGRP networks is often dependent on the range of queries through the network. This topic explains how to mark the spokes of a large network as stubs to reduce the number of EIGRP queries and thus improve network scaling.

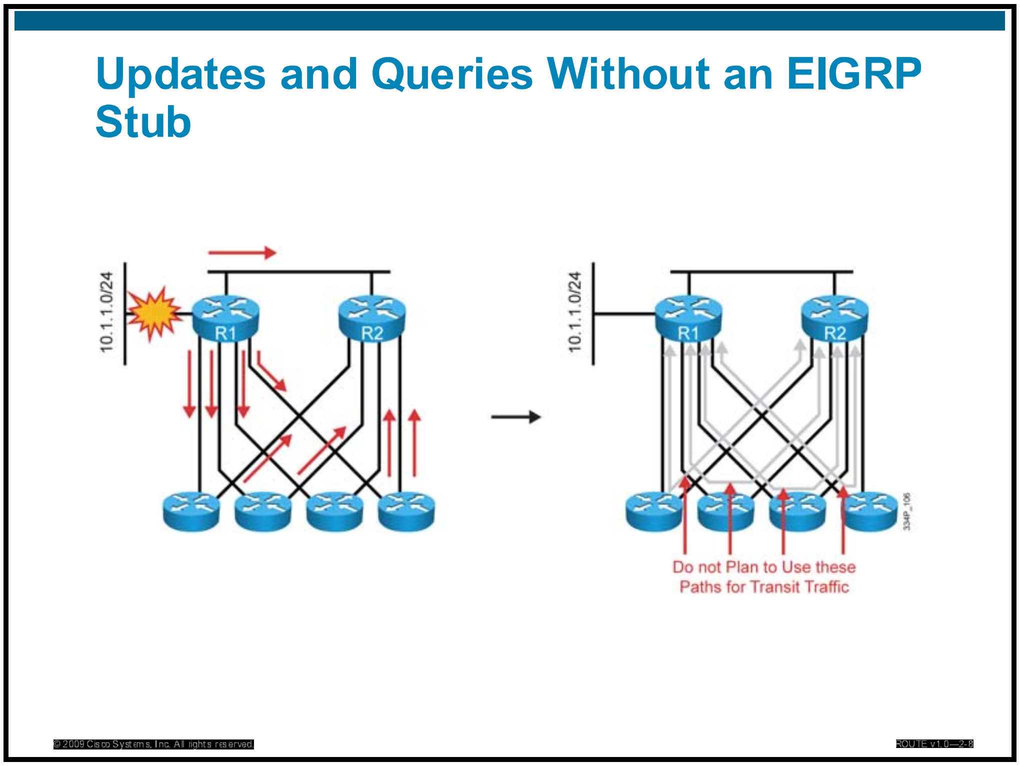

Updates and Queries Without an EIGRP Stub

Do not Plan to Use these Paths for Transit Traffic

ROUTE v1.0—2-8

When a router running EIGRP loses its connection to a network, it first searches for alternate loop-free paths. If it finds none, it then sends queries to each of its neighbors, looking for an alternate path. If the neighbor does not have another path to this destination, it replies with infinity. After receiving all replies, router R1 then removes all references to this route from its local tables. In large hub-and-spoke networks, the hub routers have to build queries and process replies from each of the spokes; this limits scaling.

Without the stub feature, a hub router will send a query to the spoke routers if a route is lost somewhere in the network. If there is a communication problem over the WAN link between the hub router and the spoke router, replies may not be received for all queries (this is known as SIA), and the network may become unstable.

By default (when a router is not stub-enabled), queries for network 10.1.1.0/24 are sent to the remote routers, thus unnecessarily utilizing the bandwidth and possibly invoking routes that are stuck-in-active. Each of the remote sites also sends a query towards router R2, with router R2 receiving five queries that it must process and answer. If these spokes are remote sites, they typically have two connections for redundancy. Router R1 should never use the spokes as a path to anything reachable through router R2, so there is no reason to learn about, or query for, routes through these spokes.

Hub-and-spoke network topologies commonly use stub routing. If a true stub network is required, the hub router can be configured to send a default route to the spoke routers. This approach is the simplest and conserves the most bandwidth and memory on the spoke routers.

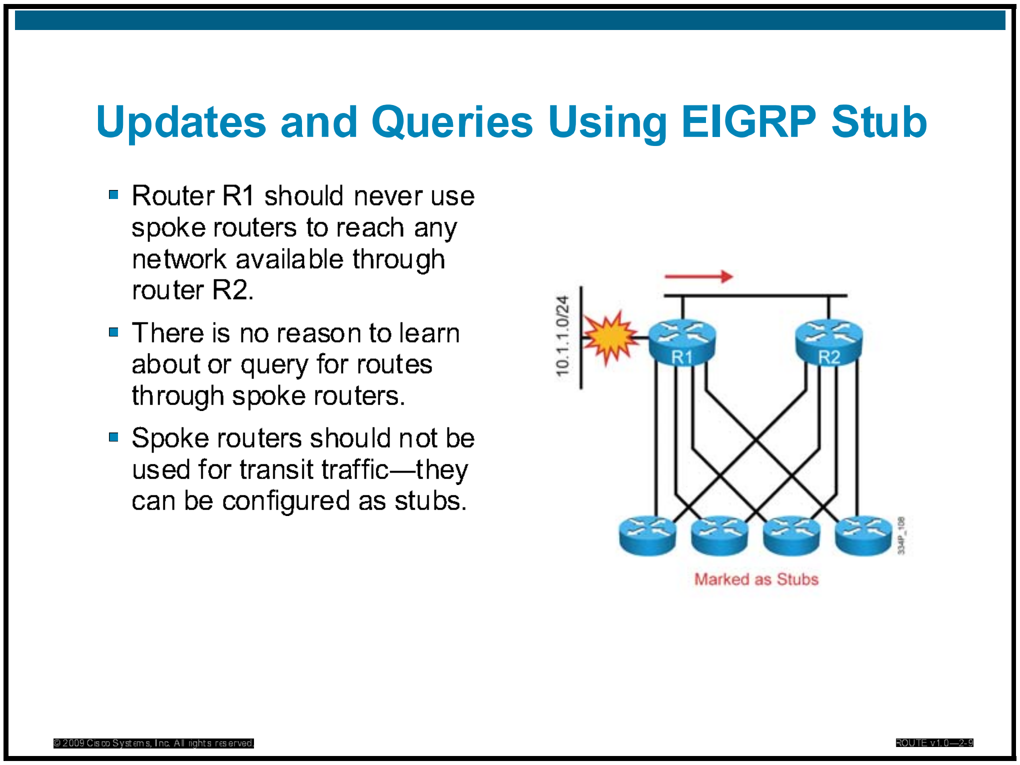

Updates and Queries Using EIGRP Stub

- Router R1 should never use spoke routers to reach any network available through router R2.
- There is no reason to learn about or query for routes through spoke routers.
- Spoke routers should not be used for transit traffic—they can be configured as stubs.

Marked as Stubs

The EIGRP stub routing feature allows a network administrator to prevent the sending of queries to the spoke router under any condition. Remote sites allow the hub (regional office) sites to immediately answer queries without propagating the queries to the remote sites. This saves CPU cycles and bandwidth. It also lessens the convergence time, even when the remote sites are dual-homed to two or more hub (regional) sites.

The figure in the slide shows that spoke routers are configured as stubs to signal to routers R1 and R2 that the paths through the spoke routers should not be used for transit traffic. Router R1 is not sending queries for network 10.1.1.0/24 to stubs, reducing the total number of queries and total bandwidth used. Marking the remote routers as stubs also reduces the complexity of the topology.

It is highly recommended that you use both EIGRP route summarization and EIGRP stub features to provide the best scalability.

Note The EIGRP stub routing feature does not automatically enable route summarization on the hub router. In most cases, the network administrator should configure route summarization on the hub routers.

Note Although EIGRP is a classless routing protocol, it behaves in a classful way by default. For example, one default behavior of EIGRP is to have automatic summarization turned on. When you configure the hub router to send a default route to the remote router, ensure that the **ip classless** command is on the remote router. By default, the **ip classless** command is enabled in all Cisco IOS images that support the EIGRP stub routing feature.

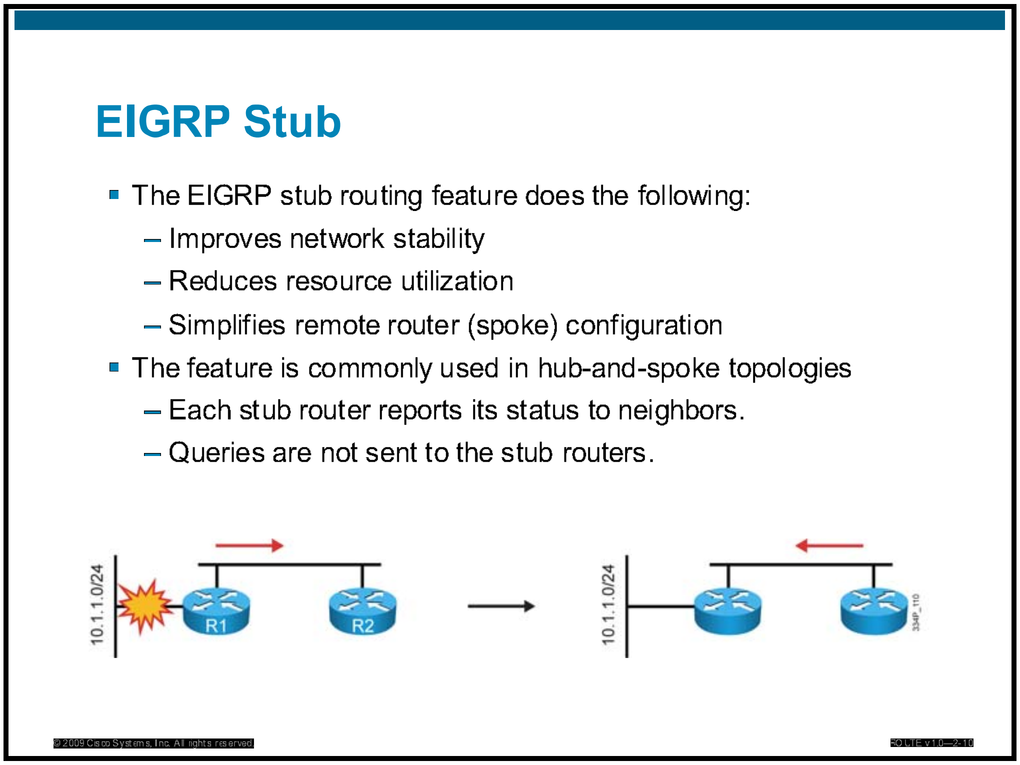

EIGRP Stub

- The EIGRP stub routing feature does the following:
  - Improves network stability
  - Reduces resource utilization
  - Simplifies remote router (spoke) configuration
- The feature is commonly used in hub-and-spoke topologies
  - Each stub router reports its status to neighbors.
  - Queries are not sent to the stub routers.

10.1.1.0/24
R1
R2
10.1.1.0/24
334P_110

ROUTE v1.0—2-10

The EIGRP stub feature was first introduced in Cisco IOS Release 12.0(7)T.

Only the remote routers are configured as stubs. A stub router sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status does not query the stub router for any routes. Therefore, a router that has a stub peer does not query that peer; instead, hub routers connected to the stub router answer the query on behalf of the stub router. The stub routing feature does not prevent routes from being advertised to the remote router.

The EIGRP stub routing feature also simplifies the configuration and maintenance of hub-and-spoke networks, improves network stability, and reduces resource utilization. When stub routing is enabled in dual-homed remote configurations, you do not have to configure filtering on remote routers to prevent them from appearing as transit paths to the hub routers.

Caution EIGRP stub routing should be used on stub routers only. A stub router is defined as a router that is connected to the network core or hub layer and through which core transit traffic should not flow. A stub router should only have hub routers for EIGRP neighbors; ignoring this restriction may cause undesirable behavior.

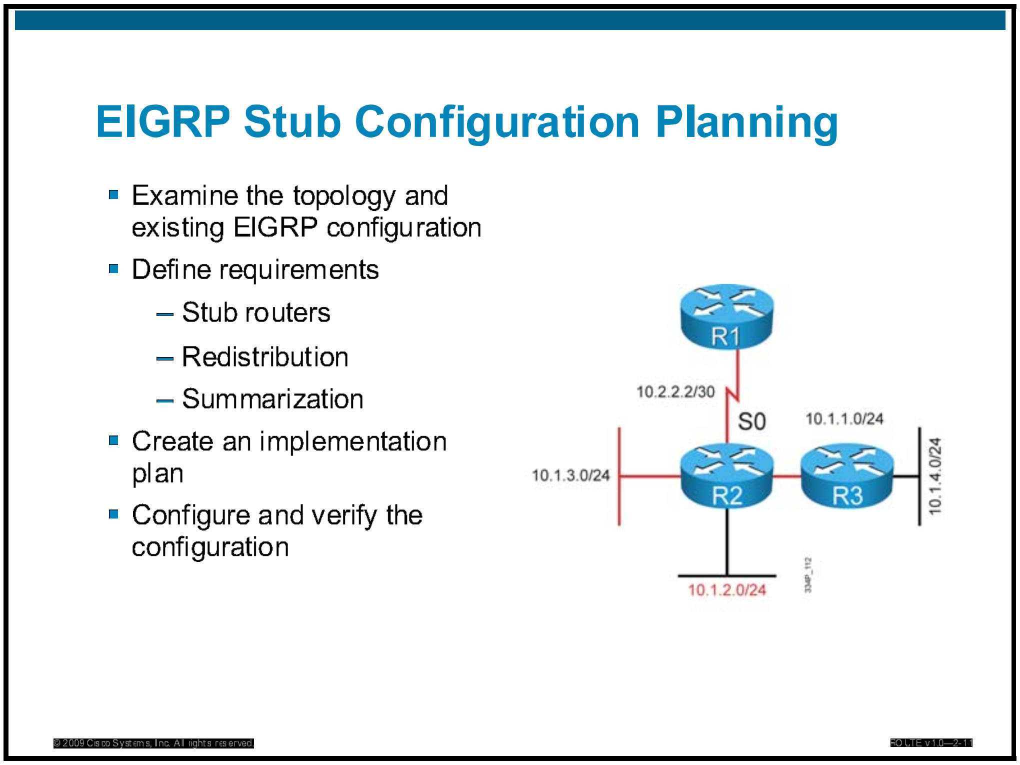

EIGRP Stub Configuration Planning

- Examine the topology and existing EIGRP configuration
- Define requirements
  - Stub routers
  - Redistribution
  - Summarization
- Create an implementation plan
- Configure and verify the configuration

ROUTE v1.0—2-11

When you configure EIGRP stub behavior on stub routers, you should examine the existing topology and configuration and follow the design. The design is based on the topology and requirements. Stub routers, together with the redistribution and summarization, limit the query range in the topology. The next step is to create the implementation plan, then configure the EIGRP stub functionality on all required routers in the EIGRP domain. When you configure EIGRP stub routers, you optimize query and reply processing and you can verify that the configuration and design are both correct.

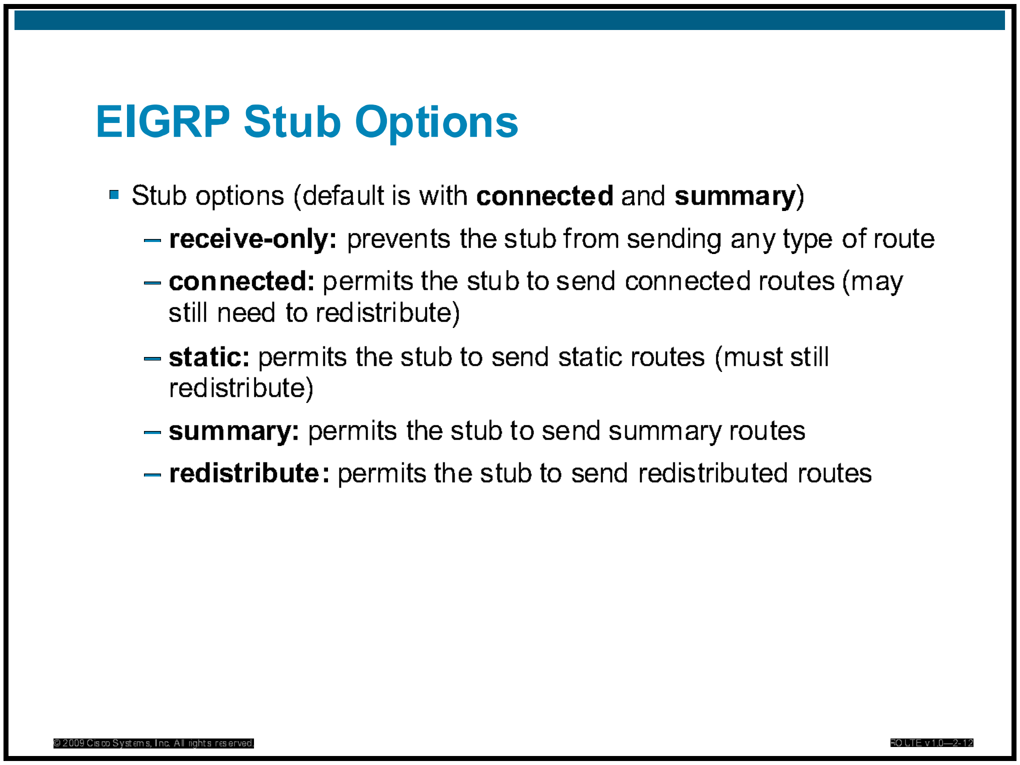

EIGRP Stub Options

- Stub options (default is with **connected** and **summary**)
  - **receive-only:** prevents the stub from sending any type of route
  - **connected:** permits the stub to send connected routes (may still need to redistribute)
  - **static:** permits the stub to send static routes (must still redistribute)
  - **summary:** permits the stub to send summary routes
  - **redistribute:** permits the stub to send redistributed routes

A router configured as a stub shares information about connected and summary routes with all neighboring routers by default.

The **receive-only** option restricts the router from sharing any of its routes with any other router within an EIGRP autonomous system (AS). This option does not permit any other option to be specified, because it prevents any type of route from being sent. The other options (**connected**, **static**, and **summary**) cannot be used with the **receive-only** option. Use this option if there is a single interface on the router.

The **connected** option permits the EIGRP stub routing feature to send connected routes. If a network command does not include the connected routes, it might be necessary to redistribute the connected routes with the **redistribute connected** command under the EIGRP process. This option is enabled by default and is the most widely practical stub option.

The **static** option permits the EIGRP stub routing feature to send static routes. You still need to redistribute static routes with the **redistribute static** command.

The **summary** option permits the EIGRP stub routing feature to send summary routes. You can either create summary routes manually, or create them automatically by enabling **auto-summary** at a major network border router. The **summary** option is enabled by default.

The **redistribute** option permits the EIGRP stub routing feature to send redistributed routes. You still need to redistribute routes with the **redistribute** command.

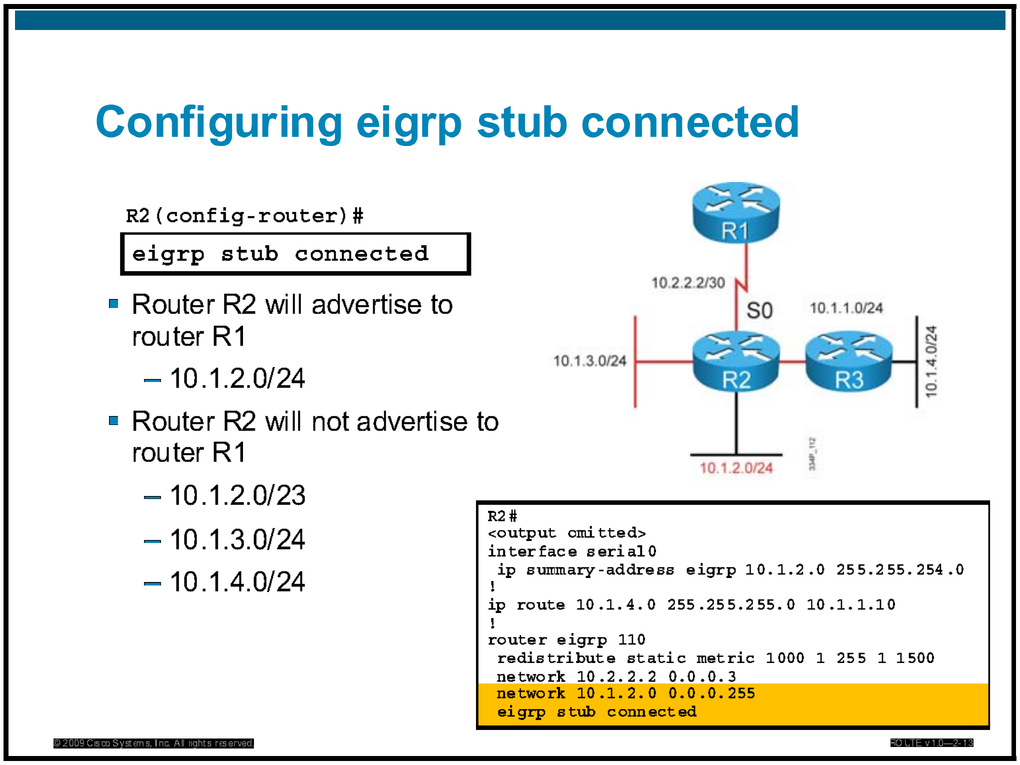

Configuring eigrp stub connected

```
R2(config-router)#
eigrp stub connected
```

- Router R2 will advertise to router R1
  - 10.1.2.0/24
- Router R2 will not advertise to router R1
  - 10.1.2.0/23
  - 10.1.3.0/24
  - 10.1.4.0/24

```
R2#
<output omitted>
interface serial0
 ip summary-address eigrp 10.1.2.0 255.255.254.0
!
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
router eigrp 110
 redistribute static metric 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.3
 network 10.1.2.0 0.0.0.255
 eigrp stub connected
```

You can use the **eigrp stub** router configuration mode command to configure EIGRP stub functionality on the routers. You can modify the **eigrp stub** command with several options to optimize the exchange of routes based on the topology and requirements.

Note The **eigrp stub** command options can be used in any combination except for the **receive-only** keyword, which prevents that router in advertising networks.

The figure in the slide shows the **eigrp stub connected** command on Router R2. Because the **connected** keyword is used, router R2 advertises the connected networks to its neighbors. Router R2, in the example in the slide, advertises the 10.1.2.0/24 route only. Network 10.1.2.0/24 is connected and, at the same time, is covered by the **network** statement under EIGRP AS 110 routing process. Although 10.1.3.0/24 is also a connected network, it is not advertised to router R1 because it is not advertised in a **network** command, and connected routes are not redistributed. The same applies to the static route for network 10.1.4.0/24. This route is not included in the EIGRP routing updates, because the EIGRP stub functionality for static routes is not specified under the EIGRP routing process. The **eigrp stub connected** command is only used as an example.

For more details about **eigrp stub** command, please check the Cisco IOS IP Routing Protocols Command Reference via the following link:
http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html

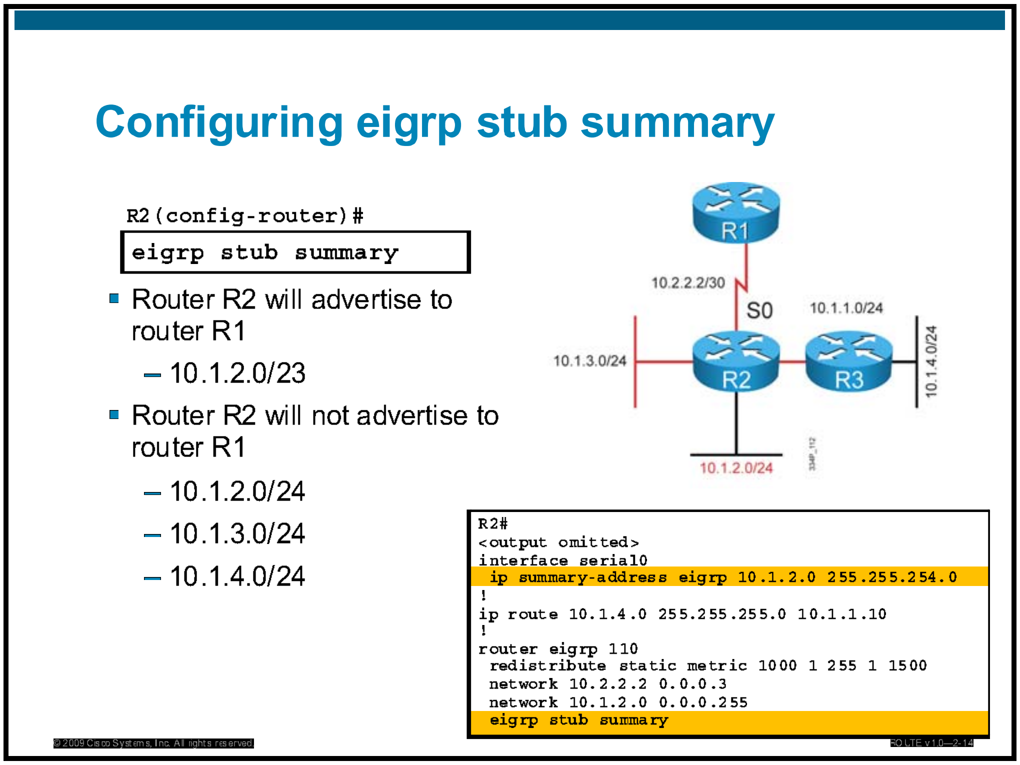

Configuring eigrp stub summary

```
R2(config-router)#
eigrp stub summary
```

- Router R2 will advertise to router R1
  - 10.1.2.0/23
- Router R2 will not advertise to router R1
  - 10.1.2.0/24
  - 10.1.3.0/24
  - 10.1.4.0/24

```
R2#
<output omitted>
interface serial0
 ip summary-address eigrp 10.1.2.0 255.255.254.0
!
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
router eigrp 110
 redistribute static metric 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.3
 network 10.1.2.0 0.0.0.255
 eigrp stub summary
```

ROUTE v1.0—2-14

The figure in the slide shows that the **eigrp stub summary** command is used on router R2. Because the **summary** keyword is used, router R2 will advertise summary routes only to its neighbors. Router R2 will only advertise 10.1.2.0/23, the summary route that is configured on the router R2. No other routes are advertised, because the **eigrp stub summary** command is the only **eigrp stub** command used under the EIGRP AS 110 routing process in the example in the slide.

Note Summary routes can be created manually with the **summary-address** command or automatically at a major network border router with the **auto-summary** command enabled, which is enabled by default.

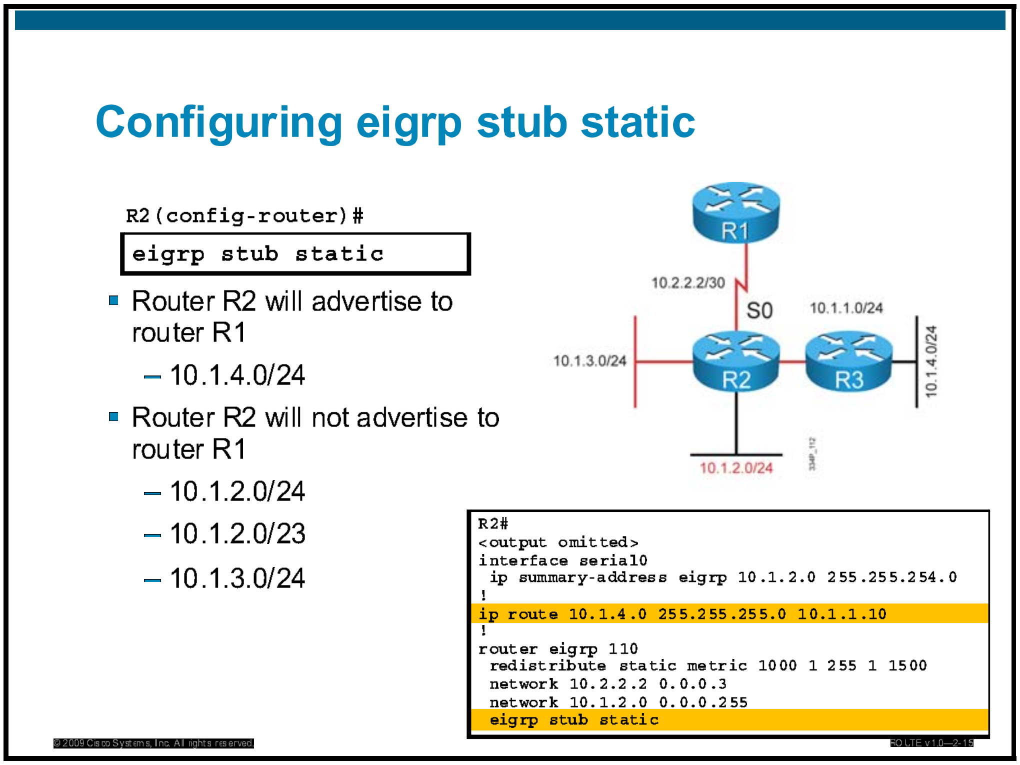

Configuring eigrp stub static

```
R2(config-router)#
eigrp stub static
```

- Router R2 will advertise to router R1
  - 10.1.4.0/24
- Router R2 will not advertise to router R1
  - 10.1.2.0/24
  - 10.1.2.0/23
  - 10.1.3.0/24

```
R2#
<output omitted>
interface serial0
 ip summary-address eigrp 10.1.2.0 255.255.254.0
!
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
router eigrp 110
 redistribute static metric 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.3
 network 10.1.2.0 0.0.0.255
 eigrp stub static
```

ROUTE v1.0—2-15

The figure in the slide shows that the **eigrp stub static** command is used on router R2. Because of the **static** keyword, router R2 will advertise static routes only to its neighbors. Router R2 will only advertise 10.1.4.0/24, the static route that is configured on the router R2. It will not advertise any other routes, because the **eigrp stub static** command is the only **eigrp stub** command used under the EIGRP AS 110 routing process in the example in the slide.

Note Without the configuration of the **eigrp stub static** option, EIGRP will not send any static routes. This includes internal static routes that normally would be automatically redistributed. It will still be necessary to redistribute static routes with the **redistribute static** command.

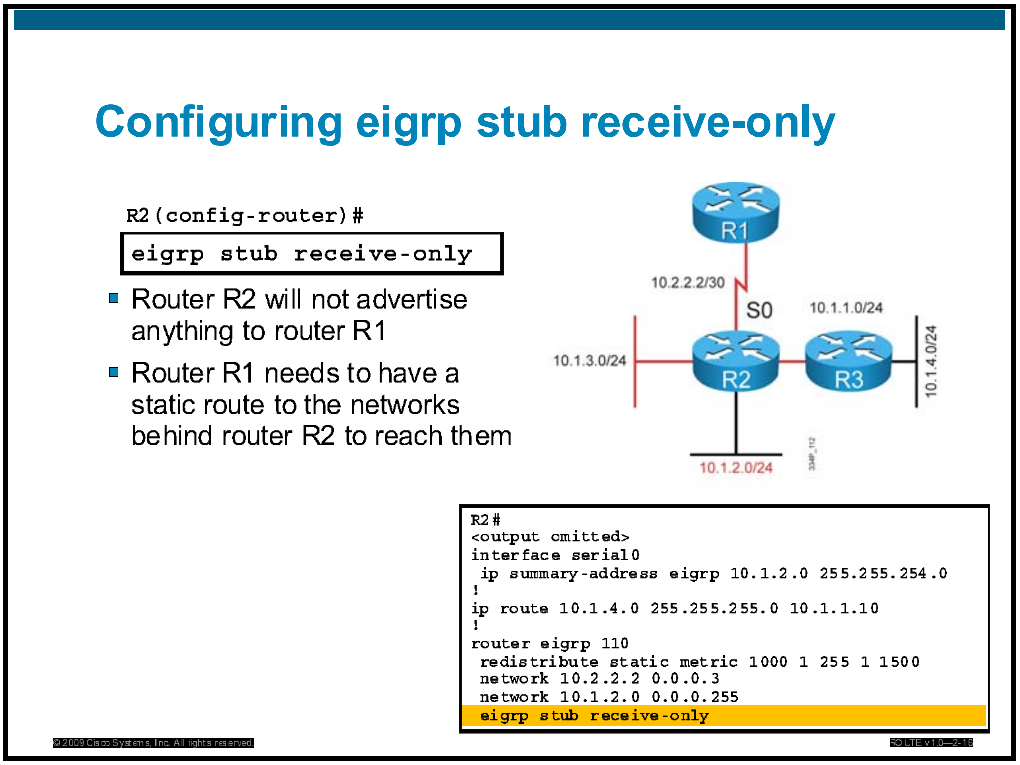

Configuring eigrp stub receive-only

```
R2(config-router)#
eigrp stub receive-only
```

- Router R2 will not advertise anything to router R1
- Router R1 needs to have a static route to the networks behind router R2 to reach them

```
R2#
<output omitted>
interface serial0
 ip summary-address eigrp 10.1.2.0 255.255.254.0
!
ip route 10.1.4.0 255.255.255.0 10.1.1.10
!
router eigrp 110
 redistribute static metric 1000 1 255 1 1500
 network 10.2.2.2 0.0.0.3
 network 10.1.2.0 0.0.0.255
 eigrp stub receive-only
```

ROUTE v1.0—2-16

The figure in the slide shows that the **eigrp stub receive-only** command is used on router R2. Because the **receive-only** keyword is used, router R2 will not advertise anything to its neighbors. Router R2, in the example in the slide, requires that static routes be configured in order to reach the networks behind router R2.

Note The **receive-only** keyword restricts the router from sharing any of its routes with any other router in the same EIGRP autonomous system and cannot be combined with any other option within the **eigrp stub** command.

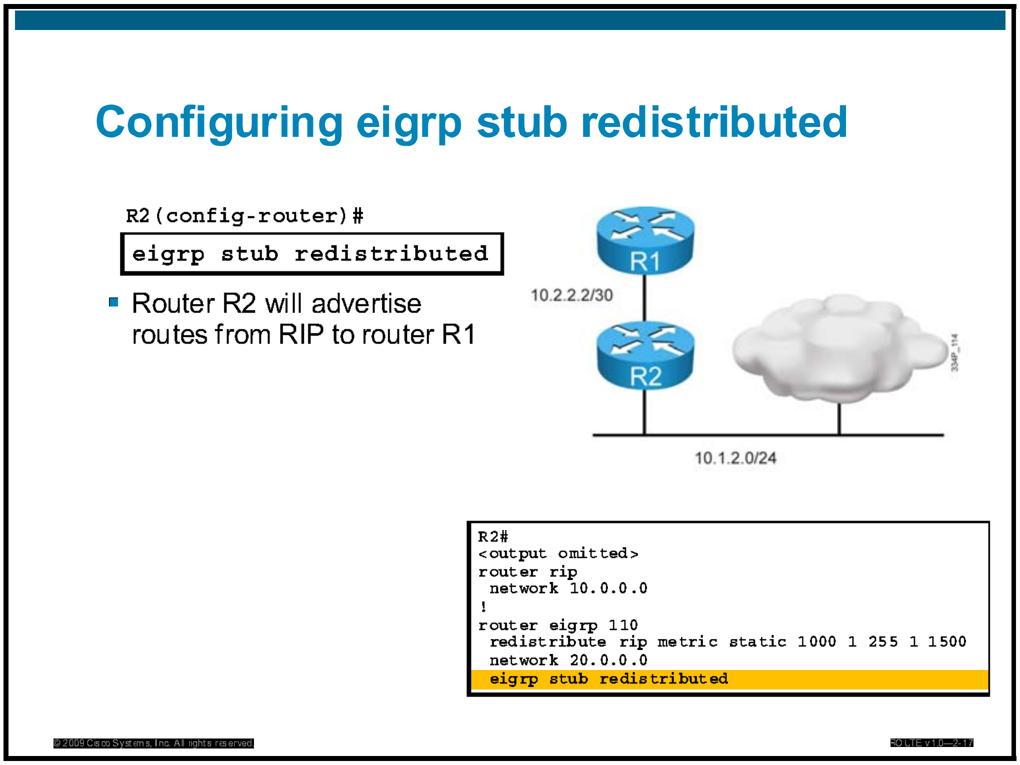

Configuring eigrp stub redistributed

```
R2(config-router)#
eigrp stub redistributed
```

- Router R2 will advertise routes from RIP to router R1

```
R2#
<output omitted>
router rip
 network 10.0.0.0
!
router eigrp 110
 redistribute rip metric static 1000 1 255 1 1500
 network 20.0.0.0
 eigrp stub redistributed
```

From the figure in the slide, it is apparent that router R2 is running the EIGRP AS 110 and RIP routing processes. At the same time, it is configured with the **eigrp stub redistributed** command within the EIGRP AS 110 routing process. Because the **redistributed** keyword is used, router R2 will advertise all RIP routes that are redistributed from the RIP routing protocol into the EIGRP AS 110 process. The sample configuration shows that the **redistribute** command is used under the EIGRP AS 110 process. This is required in order to include redistributed networks within the EIGRP advertisements.

Summary

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

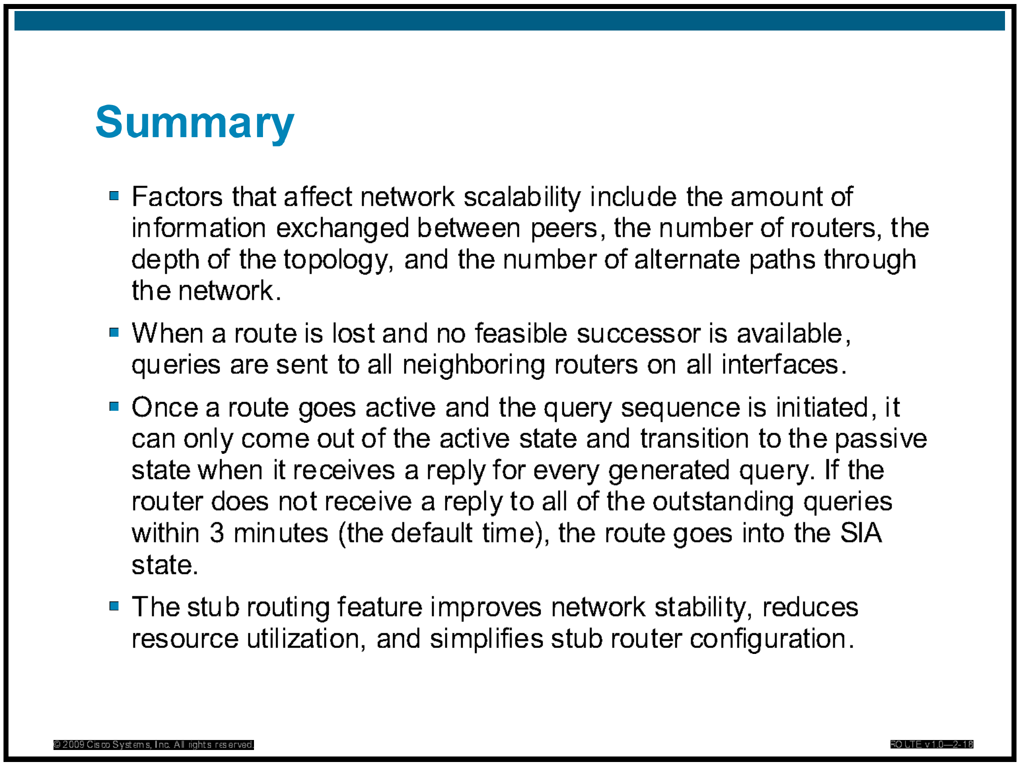

Summary

- Factors that affect network scalability include the amount of information exchanged between peers, the number of routers, the depth of the topology, and the number of alternate paths through the network.
- When a route is lost and no feasible successor is available, queries are sent to all neighboring routers on all interfaces.
- Once a route goes active and the query sequence is initiated, it can only come out of the active state and transition to the passive state when it receives a reply for every generated query. If the router does not receive a reply to all of the outstanding queries within 3 minutes (the default time), the route goes into the SIA state.
- The stub routing feature improves network stability, reduces resource utilization, and simplifies stub router configuration.

Lesson 9

Lab 2-4 Debrief

Overview

In Lab 2-4, students implement and troubleshoot EIGRP operations. They also solve EIGRP adjacency and limited connectivity issues.

Objectives

Upon completing this lesson, you will be able to implement and troubleshoot EIGRP operations to solve EIGRP adjacency and limited connectivity issues. This ability includes being able to meet these objectives:

- Complete the lab overview and verification
- Presents instructions to troubleshoot EIGRP implementation

Lab Overview and Verification

This topic describes the lab topology and key checkpoints used to create a solution and start with the verification.

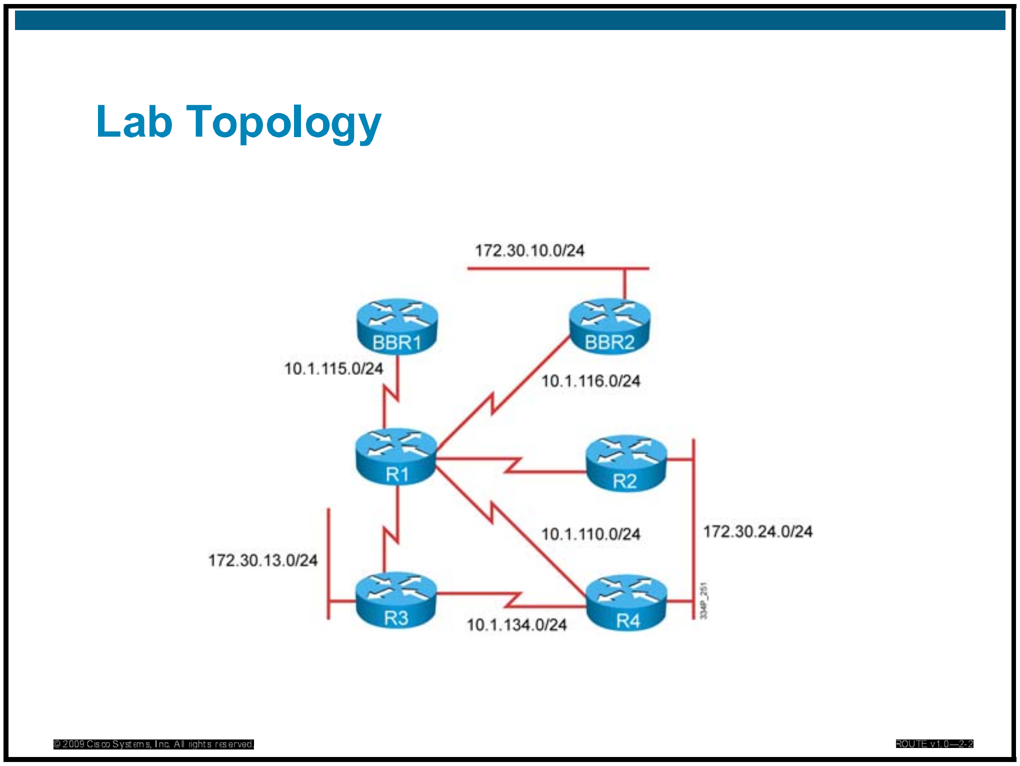

Lab Topology

ROUTE v1.0—2-2

The figure in the slide presents the physical lab topology used for the Lab 2-4: Implement and Troubleshoot EIGRP Operations. The topology uses four pod routers and two backbone routers. All routers are participating in the EIGRP routing protocol.

The configuration is broken in order to prepare the lab for students, who will go through troubleshooting steps.

Based on the topology, students will identify required parameters to implement and troubleshoot EIGRP operations.

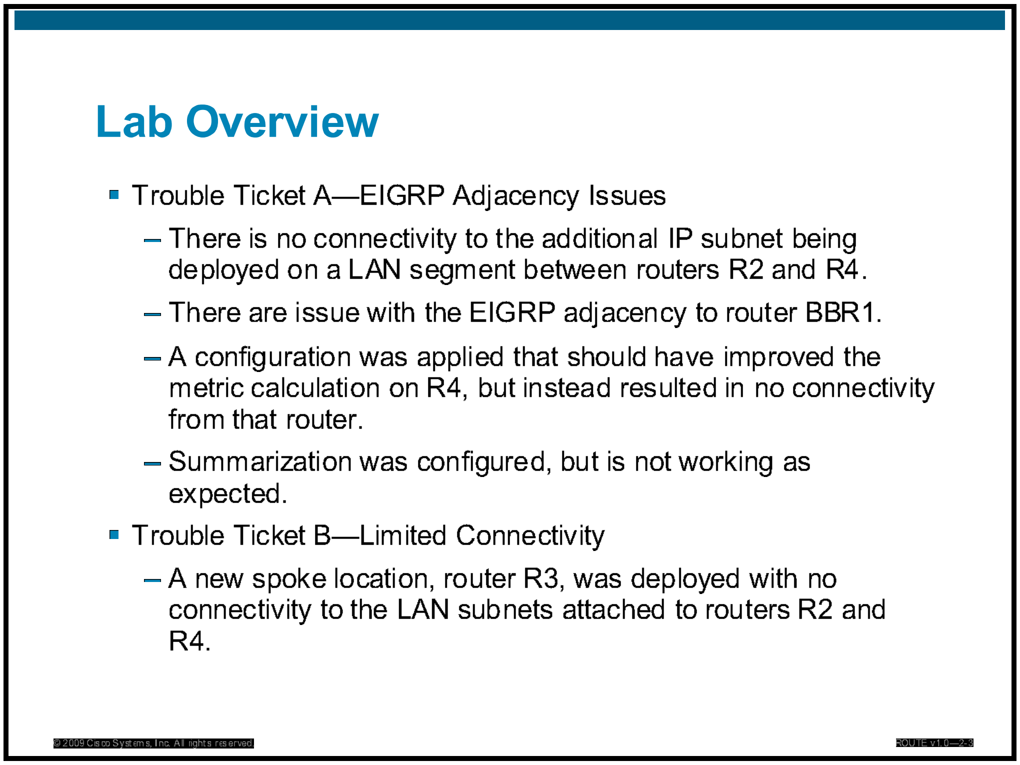

Lab Overview

- Trouble Ticket A—EIGRP Adjacency Issues
  - There is no connectivity to the additional IP subnet being deployed on a LAN segment between routers R2 and R4.
  - There are issue with the EIGRP adjacency to router BBR1.
  - A configuration was applied that should have improved the metric calculation on R4, but instead resulted in no connectivity from that router.
  - Summarization was configured, but is not working as expected.
- Trouble Ticket B—Limited Connectivity
  - A new spoke location, router R3, was deployed with no connectivity to the LAN subnets attached to routers R2 and R4.

The lab consists of two trouble tickets:

- **Trouble Ticket A—EIGRP Adjacency Issues:** There is no connectivity to an additional IP subnet being deployed on a LAN segment between routers R2 and R4. The next problem is an issue with EIGRP adjacency to router BBR1. A configuration that should have improved the metric calculation on router R4 instead caused there to be no connectivity from that router. Finally, summarization is configured, but is not working as expected.
- **Trouble Ticket B—Limited Connectivity:** A newly deployed spoke location, router R3, has no connectivity to the LAN subnets attached to routers R2 and R4.

Instructions

This subtopic presents instructions to troubleshoot EIGRP implementation.

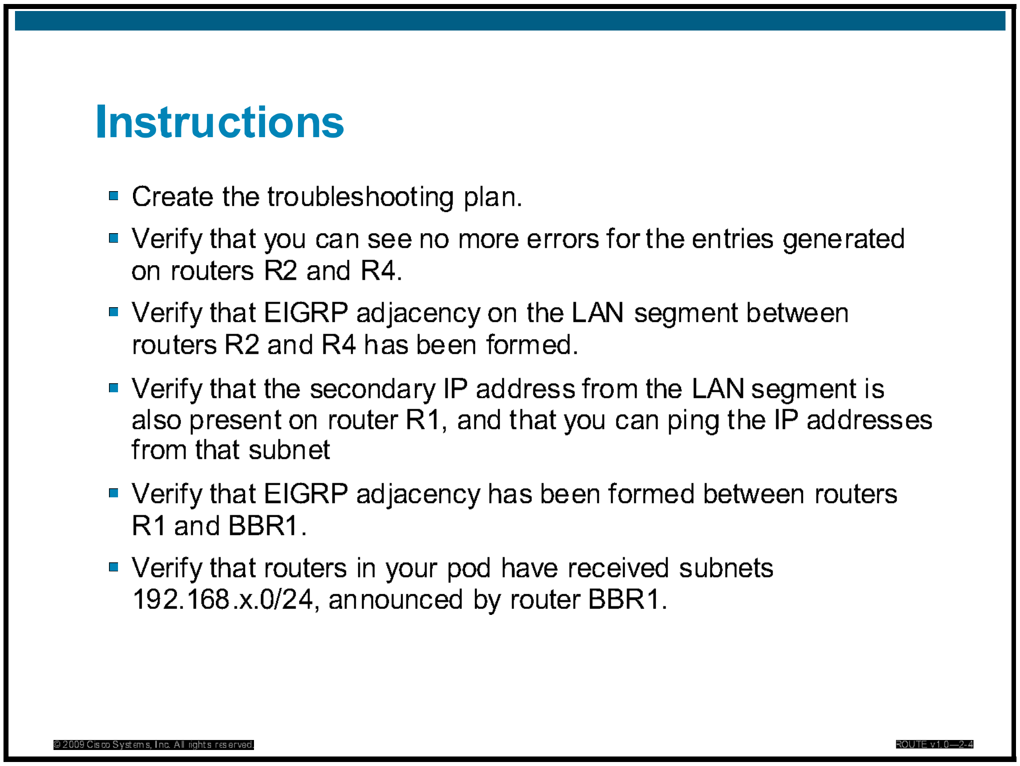

Instructions

- Create the troubleshooting plan.
- Verify that you can see no more errors for the entries generated on routers R2 and R4.
- Verify that EIGRP adjacency on the LAN segment between routers R2 and R4 has been formed.
- Verify that the secondary IP address from the LAN segment is also present on router R1, and that you can ping the IP addresses from that subnet
- Verify that EIGRP adjacency has been formed between routers R1 and BBR1.
- Verify that routers in your pod have received subnets 192.168.x.0/24, announced by router BBR1.

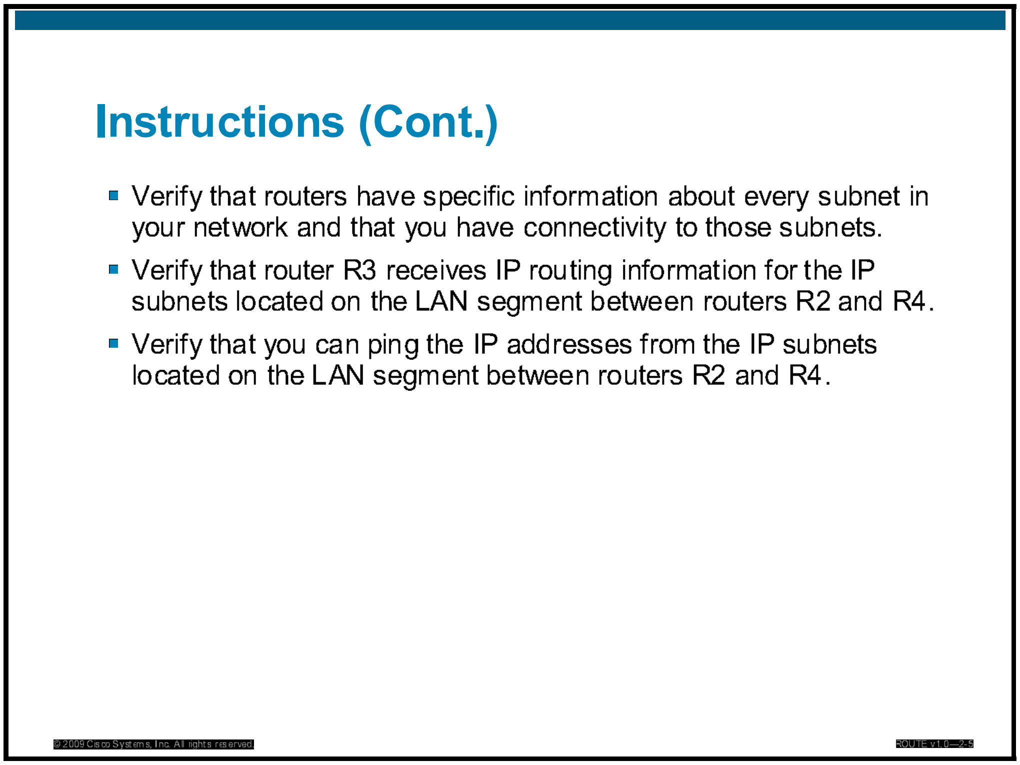

Instructions (Cont.)

- Verify that routers have specific information about every subnet in your network and that you have connectivity to those subnets.
- Verify that router R3 receives IP routing information for the IP subnets located on the LAN segment between routers R2 and R4.
- Verify that you can ping the IP addresses from the IP subnets located on the LAN segment between routers R2 and R4.

A common approach to verifying the implementation process for a routing protocol is to follow these instructions:

- Create the troubleshooting plan.
- Verify that you can see no more errors for the entries generated on routers R2 and R4.
- Verify that EIGRP adjacency on the LAN segment between routers R2 and R4 has been formed.
- Verify that the secondary IP address from the LAN segment is present on router R1 and that you can ping the IP addresses from that subnet.
- Verify that EIGRP adjacency has been formed between routers R1 and BBR1.
- Verify that the routers in your pod have received the subnets 192.168.x.0/24 announced by router BBR1.
- Verify that routers have specific information about every subnet in your network and that you have connectivity to those networks.
- Verify that router R3 receives IP routing information for the IP subnets located on the LAN segment between routers R2 and R4.
- Verify that you can ping the IP addresses from the IP subnets located on the LAN segment between routers R2 and R4.

Summary

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

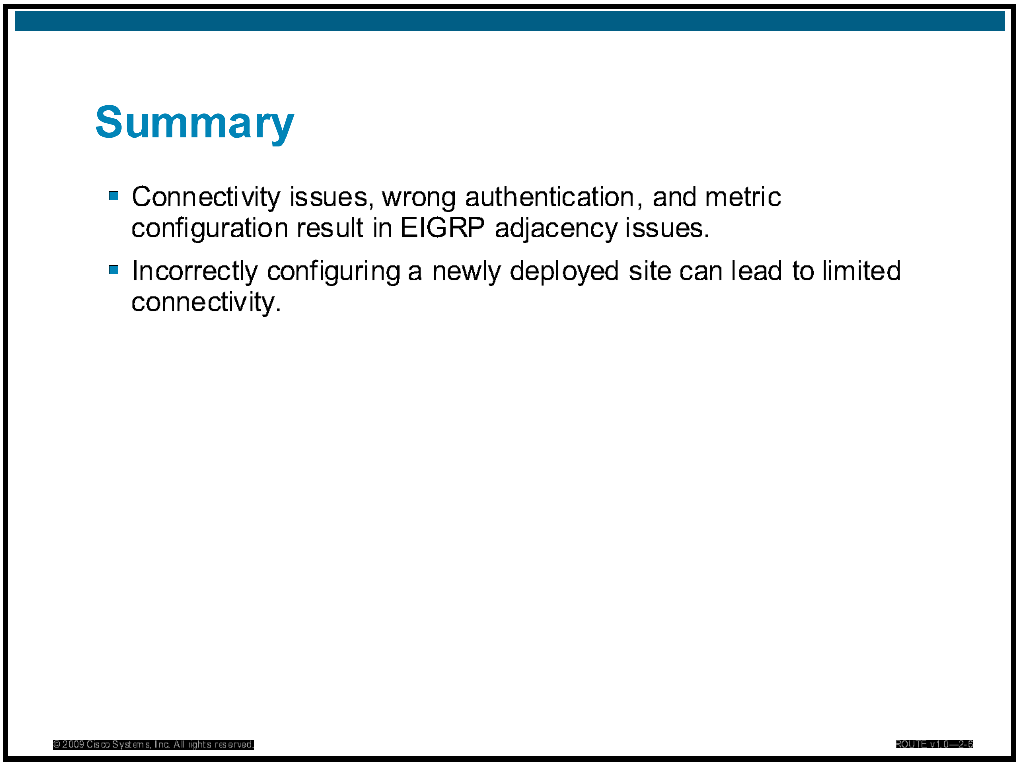

Summary

- Connectivity issues, wrong authentication, and metric configuration result in EIGRP adjacency issues.
- Incorrectly configuring a newly deployed site can lead to limited connectivity.

ROUTE v1.0—2-6

Module Summary

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

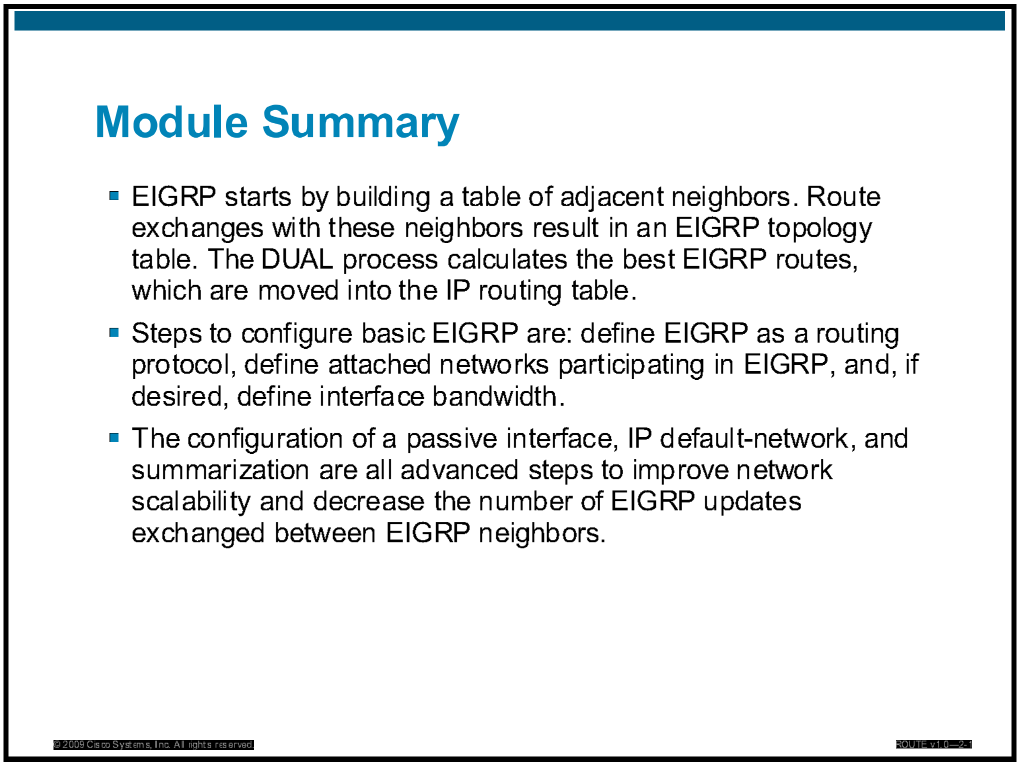

Module Summary

- EIGRP starts by building a table of adjacent neighbors. Route exchanges with these neighbors result in an EIGRP topology table. The DUAL process calculates the best EIGRP routes, which are moved into the IP routing table.
- Steps to configure basic EIGRP are: define EIGRP as a routing protocol, define attached networks participating in EIGRP, and, if desired, define interface bandwidth.
- The configuration of a passive interface, IP default-network, and summarization are all advanced steps to improve network scalability and decrease the number of EIGRP updates exchanged between EIGRP neighbors.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-1

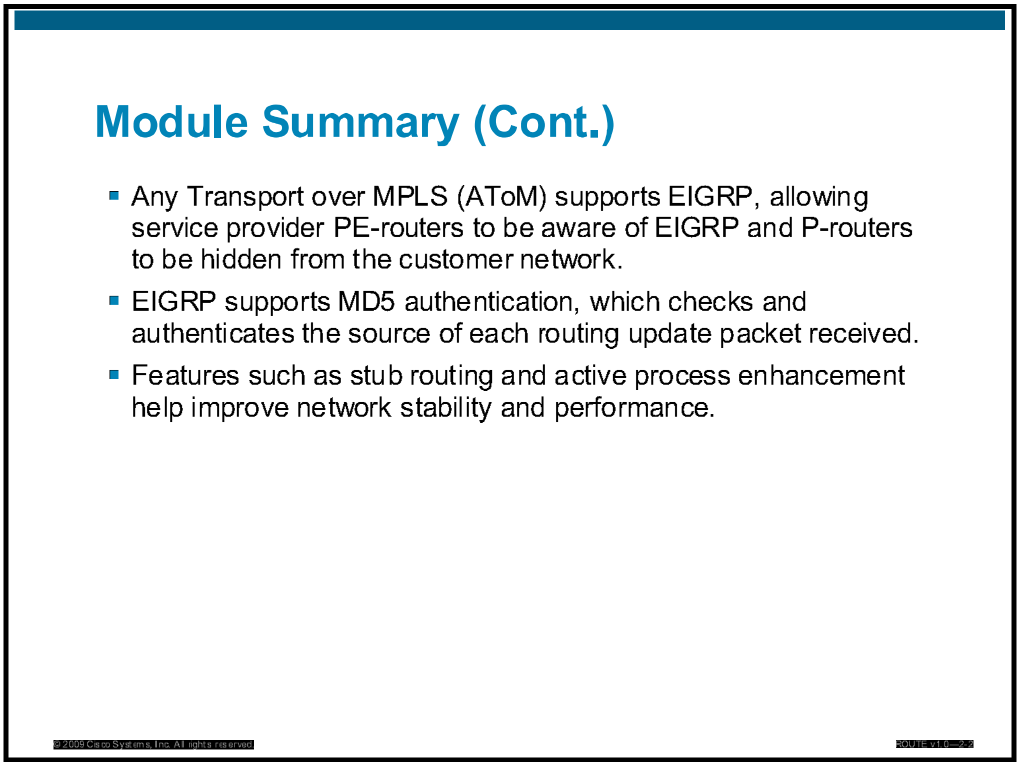

Module Summary (Cont.)

- Any Transport over MPLS (AToM) supports EIGRP, allowing service provider PE-routers to be aware of EIGRP and P-routers to be hidden from the customer network.
- EIGRP supports MD5 authentication, which checks and authenticates the source of each routing update packet received.
- Features such as stub routing and active process enhancement help improve network stability and performance.

© 2009 Cisco Systems, Inc. All rights reserved. ROUTE v1.0—2-2

Configuring Enhanced Interior Gateway Routing Protocol (EIGRP) for your routing environment enables you to achieve benefits such as rapid convergence, lower bandwidth utilization, and multiple routed protocol support. Using EIGRP ensures that as a network grows larger, it will still operate efficiently and adjust to changes rapidly.

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.

Q1) Which three features are benefits of EIGRP? (Choose three.) (Source: Planning Routing Implementations with EIGRP)

A) fast convergence
B) support for VLSM and discontiguous subnets
C) same metric algorithm as OSPF
D) manual route summarization at any point in the network

Q2) What is listed in the EIGRP topology table? (Source: Planning Routing Implementations with EIGRP)

A) directly connected routers that have formed an EIGRP adjacency
B) best routes to a destination network
C) all routes learned from each EIGRP neighbor
D) all EIGRP neighbors in the EIGRP domain

Q3) What is listed in the EIGRP neighbor table? (Source: Planning Routing Implementations with EIGRP)

A) directly connected routers that have formed EIGRP adjacencies
B) best routes to a destination network
C) all routes learned from each EIGRP neighbor
D) all EIGRP neighbors in the EIGRP domain

Q4) What is listed in the IP routing table? (Source: Planning Routing Implementations with EIGRP)

A) directly connected routers that have formed EIGRP adjacencies
B) best routes to a destination network
C) all routes learned from each EIGRP neighbor
D) all EIGRP neighbors in the EIGRP domain

Q5) Which two statements are true of the EIGRP metric calculation? (Choose two.) (Source: Planning Routing Implementations with EIGRP)

A) The following are the default K values: K1 = 1, K2 = 1, K3 = 0, K4=0, K5 = 0.
B) To convert an IGRP metric to an EIGRP metric, multiply the IGRP metric by 256.
C) To convert an EIGRP metric to an IGRP metric, multiply the EIGRP metric by 256.
D) The following are the default K values: K1 = 1, K2 = 0, K3 = 1, K4 = 0, K5=0.

Q6) Which three characteristics are key features of EIGRP? (Choose three.) (Source: Planning Routing Implementations with EIGRP)

A) fast convergence
B) partial updates
C) support for multiple Layer 3 protocols
D) backward compatibility with RIP

Q7) Which of these are key EIGRP technologies? (Choose three.) (Source: Planning Routing Implementations with EIGRP)

A) RTP
B) protocol-dependent modules
C) protocol-independent modules
D) DUAL
E) RMTP

Q8) Which two characteristics are features of EIGRP? (Choose two.) (Source: Planning Routing Implementations with EIGRP)

A) support for load balancing across unequal-cost paths
B) manual summarization at any point on the internetwork
C) provisioning of highly structured area design requirements
D) automatic redistribution of static routes

Q9) Which type of database is a list of all EIGRP adjacencies? (Source: Planning Routing Implementations with EIGRP)

A) EIGRP topology table
B) EIGRP neighbor table
C) IP routing table
D) IP EIGRP adjacency table

Q10) Which type of database contains a list of the best EIGRP routes to reach a destination? (Source: Planning Routing Implementations with EIGRP)

A) EIGRP topology table
B) EIGRP neighbor table
C) IP routing table
D) IP EIGRP adjacency table

Q11) Which type of database contains a list of all possible EIGRP routes to reach a destination? (Source: Planning Routing Implementations with EIGRP)

A) EIGRP topology table
B) EIGRP neighbor table
C) IP routing table
D) IP EIGRP adjacency table

Q12) Which five criteria may be considered by EIGRP when calculating the metric? (Choose five.) (Source: Planning Routing Implementations with EIGRP)

A) MTU
B) bandwidth
C) cost
D) delay
E) load
F) hop count
G) reliability

Q13) Which packet type establishes neighbor relationships? (Source: Planning Routing Implementations with EIGRP)

A) ACK
B) hello
C) query
D) reply
E) update

Q14) Which packet type is responsible for sending routing advertisements? (Source: Planning Routing Implementations with EIGRP)

A) ACK
B) hello
C) query
D) reply
E) update

Q15) What should a network engineer do before configuring EIGRP in the network? (Choose three.) (Source: Planning Routing Implementations with EIGRP)

A) assess the requirements
B) assess the existing configuration and topology
C) create the documentation
D) verify the EIGRP neighbors
E) create an implementation plan

Q16) Which parameters are included in an EIGRP implementation plan? (Choose two.) (Source: Planning Routing Implementations with EIGRP)

A) IP addressing
B) EIGRP AS number
C) the difference in the metric between EIGRP and IGRP
D) feasible distance of feasible successor
E) feasible distance of successor

Q17) Which packet type is used to ask neighbors about routing information? (Source: Planning Routing Implementations with EIGRP)

A) ACK
B) hello
C) query
D) reply
E) update

Q18) DUAL selects as the successor for a specific destination network the next-hop router with which of these? (Source: Planning Routing Implementations with EIGRP)

A) highest FD
B) lowest FD
C) highest AD
D) lowest AD

Q19) What is the formula for selecting a feasible successor? (Source: Planning Routing Implementations with EIGRP)

A) The AD of the current successor route is less than the FD of the feasible successor route.
B) The FD of the current successor route is less than the AD of the feasible successor route.
C) The FD of the feasible successor route is less than the AD of the current successor route.
D) The AD of the feasible successor route is less than the FD of the current successor route.

Q20) What does EIGRP do when a successor fails and there are no feasible successors, but there are alternate paths available? (Source: Planning Routing Implementations with EIGRP)

A) It immediately uses the alternate pathway with the lowest FD and sends queries and updates to ensure that this pathway is loop-free.
B) It automatically uses the alternate pathway with the lowest FD.
C) It sends queries to see if the alternate paths are still viable. When a loop-free path is found, the path is installed in the routing table.
D) It removes the network from the routing table and waits for the periodic update from EIGRP neighbors to see if an alternate route exists.

Q21) Which two conditions signify the active state for EIGRP? (Choose two.) (Source: Planning Routing Implementations with EIGRP)

A) The route can be used and is stable.
B) The route cannot be used.
C) EIGRP queries are outstanding and the router is waiting for EIGRP replies.
D) This is the best route with the lowest FD.

Q22) Test your understanding of EIGRP by matching terms with statements. Write the number of the statement in front of the term that the statement describes. (Source: Planning Routing Implementations with EIGRP)

Term

_____ 1. successor

_____ 2. feasible successor

_____ 3. hello

_____ 4. topology table

_____ 5. IP

_____ 6. update

_____ 7. routing table

_____ 8. DUAL

Statement

1. a network protocol that EIGRP supports
2. a database that contains successor and feasible successor information
3. a database that includes administrative distance
4. a neighbor router that has the best path to a destination
5. a neighbor router that has the best alternate loop-free path to a destination
6. an algorithm used by EIGRP to ensure fast convergence
7. a multicast packet used to discover neighbors
8. a packet sent by EIGRP routers when a new neighbor is discovered or a change occurs

Q23) What is the purpose of the **network** command for EIGRP? (Source: Planning Routing Implementations with EIGRP)

A) to determine which router interfaces participate in EIGRP and which networks the router advertises
B) to specify the AS number to which the router belongs
C) to define the EIGRP neighbors
D) to tell EIGRP which networks to advertise, those directly connected and those learned through EIGRP

Q24) Which command creates a default route for EIGRP? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **ip default-network** *network-number*
B) **ip route 0.0.0.0 0.0.0.0 outbound-interface**
C) **ip route 0.0.0.0 255.0.0.0 outbound-interface**
D) **ip route 0.0.0.0 255.255.255.255 outbound-interface**

Q25) Which command displays an indication if a network is SIA? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **show ip route**
B) **show ip protocol**
C) **show ip eigrp topology**

Q26) What is the correct **network** command to allow updates to propagate only out of interfaces that are part of subnet 10.1.0.0/16? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **network 10.1.0.0 mask 255.255.0.0**
B) **network 10.1.0.0 mask 0.0.255.255**
C) **network 10.1.0.0 255.255.0.0**
D) **network 10.1.0.0 0.0.255.255**

Q27) Which three of these are true of configuring the **ip default-network** command for EIGRP? (Choose three.) (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) The network must be reachable by the router using this command.
B) The command will set the gateway of last resort to 0.0.0.0 on the router issuing this command.
C) The network must be advertised to other neighbors as an EIGRP route.
D) The network will be flagged by other EIGRP routers as a candidate default route.

Q28) What does the passive state in the EIGRP topology table signify? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) There are outstanding queries for this network.
B) The network is unreachable.
C) The network is up and operational, and this state signifies normal conditions.
D) A feasible successor has been selected.

Q29) Which command indicates the number of EIGRP peer routers on an interface? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **show ip eigrp interfaces**
B) **show ip eigrp neighbors**
C) **show ip route**
D) **show ip eigrp topology**

Q30) Which command is used to disable automatic EIGRP network-boundary summarization, and where is it applied? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **no boundary-summarization** at the interface level
B) **no auto-summary** under the routing process
C) **no auto-summary** at the interface level
D) **no boundary-summarization** under the routing process

Q31) Which command is used to configure manual summarization of all the subnets in network 10.1.32.0/21 for EIGRP in AS 101? (Source: Implementing and Verifying Basic EIGRP for the Enterprise LAN Architecture)

A) **ip summary-address eigrp 101 10.1.32.0 255.255.248.0**
B) **ip eigrp 101 summary-address 10.1.32.0 255.255.240**.0
C) **ip summary-address eigrp 101 10.1.32.0 255.255.240.0**
D) **ip eigrp 101 summary-address 10.1.32.0 255.255.248.0**

Q32) By default, how many equal-cost paths to the same destination network can EIGRP place in the routing table? (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) one
B) two
C) four
D) six

Q33) Between headquarters and remote site A, there are two dedicated serial PPP connections, one at 64 kb/s, and the other at 128 kb/s. What is the appropriate variance to allow for unequal-cost load balancing across these links? (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) 1
B) 2
C) 3
D) 4

Q34) What is the default bandwidth percentage that EIGRP uses on WAN links? (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) 25 percent
B) 50 percent
C) 75 percent
D) 100 percent

Q35) Which two statements best describes AToM? (Choose two.) (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) AToM stands for Any Transport over MPLS
B) AToM unifies Layer 2 and Layer 3 over a common MPLS infrastructure.
C) AToM must be enabled for EIGRP in an MPLS environment.
D) Authentication is required in an AToM design.

Q36) Which router does not participate in customer routing? (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) C router
B) CE router
C) PE router
D) P router

Q37) What are three requirements for MPLS VPN technology? (Choose three.) (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) CE routers should not be aware of MPLS VPN.
B) The service provider's P routers must be hidden from the customer.
C) C routers must be directly connected to PE routers.
D) Each PE router must appear as another router in the customer's network.

Q38) You do not need to change the basic configuration when you deploy EIGRP over a physical interface using dynamic mapping, thus relying on Inverse ARP. ((Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) true
B) false

Q39) Which two topologies use EIGRP over Frame Relay multipoint subinterfaces? (Choose two.) (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) point-to-point
B) partial-mesh
C) full-mesh
D) hub-and-spoke

Q40) What are two main reasons for the relatively fast EIGRP neighbor loss detection on point-to-point subinterfaces? (Choose two.) (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) There is a small default EIGRP hello and hold timer, which is identical to the value used on point-to-point links.
B) On Frame Relay networks, the subinterface is declared down if the DLCI attached to the interface is lost.
C) Neighbors send immediate EIGRP update packets to inform each other of neighbor loss.
D) The EIGRP process is checking for neighbors every 5 seconds

Q41) Which topologies use EIGRP over Frame Relay multipoint subinterfaces? (Choose two.) (Source: Configuring and Verifying EIGRP for the Enterprise WAN Architecture)

A) point-to-point
B) partial-mesh
C) full-mesh
D) hub-and-spoke

Q42) Which authentication does EIGRP support? (Source: Implementing and Verifying EIGRP Authentication)

A) MD5
B) MD5 and simple password
C) simple password
D) none

Q43) When EIGRP authentication is configured between two routers, each router has its own unique password. (Source: Implementing and Verifying EIGRP Authentication)

A) true
B) false

Q44) Which three of these are used to generate the message digest when EIGRP MD5 authentication is configured? (Choose three.) (Source: Implementing and Verifying EIGRP Authentication)

A) packet
B) sequence number
C) key ID
D) key
E) router ID

Q45) What does the **accept-lifetime 04:00:00 Jan 1 2006 infinite** command do? (Source: Implementing and Verifying EIGRP Authentication)

A) specifies that a key is acceptable for use on received packets from the first of January 2006 onward
B) specifies that a key is acceptable for use on sent packets from the first of January 2006 onward
C) specifies that a key is acceptable for use on received packets until the first of January 2006
D) specifies that a key is acceptable for use on sent packets until the first of January 2006

Q46) Which command specifies that EIGRP MD5 authentication in AS 100 be used? (Source: Implementing and Verifying EIGRP Authentication)

A) **ip authentication mode eigrp 100 md5**
B) **ip eigrp 100 authentication mode md5**
C) **ip authentication-key eigrp 100**
D) **ip message-digest-key eigrp 100**
E) **ip eigrp 100 authentication message-digest**

Q47) Which command is used to troubleshoot EIGRP authentication? (Source: Implementing and Verifying EIGRP Authentication)

A) **debug ip eigrp adj**
B) **debug ip eigrp packets**
C) **debug eigrp packets**
D) **debug ip eigrp adjacency events**
E) **debug eigrp adj**

Q48) When a router gets a query from a neighboring router that is not a successor for the network listed in the query, and that network is in a passive state on this router, what does the router do? (Source: Advanced EIGRP Features in an Enterprise Network)

A) The router replies that the destination is unreachable.
B) The router attempts to find a new successor. If successful, it replies with new information. If unsuccessful, it marks the destination as unreachable and queries all neighboring routers except the previous successor.
C) The router replies with the current successor information.
D) The router marks the destination as unreachable and queries all neighboring routers except the previous successor.

Q49) Which three of these factors affect network scalability? (Choose three.) (Source: Advanced EIGRP Features in an Enterprise Network)

A) number of alternate paths through the network
B) amount of information exchanged between neighbors
C) the amount of different AS numbers used in the network
D) depth of the topology

Q50) Which three statements about implementing EIGRP stub routers are true? (Choose three.) (Source: Advanced EIGRP Features in an Enterprise Network)

A) Stub routing is commonly used on networks with hub-and-spoke topologies.
B) The EIGRP stub feature should be configured only on remote spoke routers.
C) EIGRP stub routers can and should be used as transit points to other parts of the network and other autonomous systems.
D) Queries are not propagated to EIGRP stub routers; EIGRP updates are sent to stub routers, or a default route is passed.

Q51) How long does a querying router wait to reset a neighbor that fails to reply to a query? (Source: Advanced EIGRP Features in an Enterprise Network)

A) 15 seconds
B) 40 seconds
C) 1 minute
D) 3 minutes

Q52) Which command configures an EIGRP stub router to not send any routing updates? (Source: Advanced EIGRP Features in an Enterprise Network)

A) **eigrp stub**
B) **eigrp stub receive-only**
C) **eigrp stub no-send**
D) **eigrp stub none**

Module Self-Check Answer Key

Q1)	A, B, D
Q2)	C
Q3)	A
Q4)	B
Q5)	B, D
Q6)	A, B, C
Q7)	A, B, D
Q8)	A, B
Q9)	B
Q10)	C
Q11)	A
Q12)	A, B, D, E, G
Q13)	B
Q14)	E
Q15)	A, B, E
Q16)	A, B
Q17)	C
Q18)	B
Q19)	D
Q20)	C
Q21)	B, C
Q22)	A=4, B=5, C=7, D=2, E=1, F=8, G=3, H=6
Q23)	A
Q24)	A
Q25)	C
Q26)	D
Q27)	A, C, D
Q28)	C
Q29)	A
Q30)	B
Q31)	A
Q32)	C
Q33)	B
Q34)	B
Q35)	A,B

Q36)	D
Q37)	A, B, D
Q38)	A
Q39)	B, C
Q40)	A, B
Q41)	B, C
Q42)	A
Q43)	B
Q44)	A, C, D
Q45)	A
Q46)	A
Q47)	C
Q48)	C
Q49)	A, B, D
Q50)	A, B, D
Q51)	D
Q52)	B