

ICND2

Interconnecting Cisco Networking Devices, Part 2

Volume 1

Version 2.0

Student Guide

Part Number: Part Number TBD

Do Not Duplicate.
Post beta, not for release.



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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

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

Welcome Students

Note Students, this letter describes important course evaluation access information.

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

Cisco relies on customer feedback to guide business decisions; therefore, your valuable input will help shape future Cisco course curricula, products, and training offerings. Please complete a brief Cisco online course evaluation of your instructor and the course materials in this student kit. On the final day of class, your instructor will provide you with a URL directing you to a short postcourse 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

Do Not Duplicate
Post beta, not for release

The Cisco M-Learning Test and Study App

The Cisco M-Learning Test and Study app is the ideal on-the-go study application for those preparing for Cisco certifications.

Scan the following QR code to get the free Cisco M-Learning Test and Study app along with the 20 free exam questions and free TCP/IP Architecture video.

**Scan the code to get
the iPhone M-Test app**



Do Not Duplicate
Post beta, not for release.

Table of Contents

| | |
|--|------------|
| Course Introduction | 11 |
| Overview | 11 |
| Course Goal and Objectives | 12 |
| Course Flow | 13 |
| Additional References | 14 |
| Your Training Curriculum | 15 |
| Implementing Scalable Medium-Sized Networks | 11 |
| Troubleshooting VLAN Connectivity | 13 |
| VLAN Overview | 13 |
| Creating VLANs | 15 |
| Trunk Operation | 17 |
| Configuring Trunks | 18 |
| Dynamic Trunking Protocol | 110 |
| VLAN Troubleshooting | 112 |
| Trunk Troubleshooting | 115 |
| Summary | 117 |
| Building Redundant Switched Topologies | 119 |
| Issues in Redundant Topologies | 119 |
| Spanning-Tree Operation | 122 |
| Types of Spanning-Tree Protocols | 127 |
| Per VLAN Spanning Tree Plus | 130 |
| Modifying the Bridge ID | 132 |
| Analyzing the STP Topology | 134 |
| Spanning-Tree Failure Consequences | 137 |
| PortFast and BPDU Guard | 140 |
| Summary | 143 |
| Improving Redundant Switched Topologies with EtherChannel | 145 |
| The Need for EtherChannel | 145 |
| Advantages of EtherChannel | 147 |
| EtherChannel Protocols | 149 |
| Configuring EtherChannel | 152 |
| Verifying EtherChannel | 154 |
| Summary | 156 |
| Understanding Layer 3 Redundancy | 157 |
| The Need for Default Gateway Redundancy | 157 |
| Default Gateway Redundancy | 159 |
| HSRP | 161 |
| HSRP Interface Tracking | 163 |
| HSRP Load Balancing | 164 |
| Gateway Load Balancing Protocol | 165 |
| Summary | 167 |
| Module Summary | 169 |

| | |
|---|------------|
| Module Self-Check | 171 |
| <i>Troubleshooting Basic Connectivity</i> | 21 |
| Troubleshooting IPv4 Network Connectivity | 23 |
| Components of Troubleshooting End-to-End Connectivity | 23 |
| Verification of End-to-End Connectivity | 26 |
| Verification of Physical Connectivity Issue | 212 |
| Identification of Current and Desired Path | 217 |
| Default Gateway Issues | 220 |
| Name Resolution Issues | 222 |
| ACL Issues | 225 |
| Summary | 228 |
| Troubleshooting IPv6 Network Connectivity | 229 |
| IPv6 Unicast Addresses | 229 |
| Troubleshooting End-to-End IPv6 Connectivity | 232 |
| Verification of End-to-End IPv6 Connectivity | 234 |
| Identification of Current and Desired IPv6 Path | 239 |
| Default Gateway Issues in IPv6 | 240 |
| Name Resolution Issues in IPv6 | 241 |
| ACL Issues in IPv6 | 243 |
| Summary | 245 |
| Module Summary | 247 |
| Module Self-Check | 249 |
| <i>Implementing an EIGRP-Based Solution</i> | 31 |
| Implementing EIGRP | 33 |
| Dynamic Routing Protocols | 33 |
| Administrative Distance | 37 |
| EIGRP Features | 38 |
| EIGRP Path Selection | 310 |
| EIGRP Metric | 312 |
| EIGRP Configuration | 314 |
| Verification of EIGRP Configuration | 316 |
| Load Balancing with EIGRP | 321 |
| Summary | 322 |
| Troubleshooting EIGRP | 323 |
| Components of Troubleshooting EIGRP | 323 |
| Troubleshooting EIGRP Neighbor Issues | 325 |
| Troubleshooting EIGRP Routing Table Issues | 330 |
| Summary | 334 |
| Implementing EIGRP for IPv6 | 335 |
| EIGRP for IPv6 | 335 |
| EIGRP for IPv6 Commands | 339 |
| EIGRP for IPv6 Configuration Example | 341 |
| Summary | 344 |

| | |
|---|------------|
| Module Summary | 345 |
| Module Self-Check | 347 |
| | 347 |
| <i>Implementing a Scalable, Multiarea Network, OSPF-Based Solution</i> | 41 |
| OSPF Overview | 43 |
| Overview | 43 |
| Link-State Routing Protocol Overview | 44 |
| Link-State Routing Protocol Data Structures | 45 |
| OSPF Metric | 46 |
| Establishing OSPF Neighbor Adjacencies | 47 |
| Building a Link-State Database | 49 |
| OSPF Area Structure | 411 |
| Summary | 414 |
| Multiarea OSPF IPv4 Implementation | 415 |
| Single-Area vs. Multiarea OSPF | 415 |
| Planning for the Implementation of OSPF | 417 |
| Multiarea OSPF Configuration | 418 |
| Multiarea OSPF Verification | 420 |
| Summary | 423 |
| Troubleshooting Multiarea OSPF | 425 |
| OSPF Neighbor States | 425 |
| Components of Troubleshooting OSPF | 428 |
| Troubleshooting OSPF Neighbor Issues | 429 |
| Troubleshooting OSPF Routing Table Issues | 435 |
| Troubleshooting OSPF Path Selection | 438 |
| Summary | 440 |
| Examining OSPFv3 | 441 |
| OSPFv3 Key Characteristics | 441 |
| OSPFv3 Configuration | 444 |
| OSPFv3 Configuration Verification | 446 |
| Summary | 449 |
| Module Summary | 451 |
| References | 451 |
| Module Self-Check | 453 |
| <i>Wide-Area Networks</i> | 51 |
| Understanding WAN Technologies | 53 |
| Overview | 53 |
| Introduction to WAN Technologies | 54 |
| WAN Devices | 56 |
| Serial WAN Cabling | 58 |
| WAN Layer 2 Protocols | 510 |
| WAN Link Options | 512 |
| Summary | 514 |

| | |
|--|------------|
| Configuring Serial Encapsulation | 515 |
| Overview | 515 |
| Serial Communication Links | 517 |
| Configuration of a Serial Interface | 519 |
| HDLC Protocol | 521 |
| Point-to-Point Protocol | 523 |
| PPP Configuration | 525 |
| PPP Authentication: PAP | 527 |
| PPP Authentication: CHAP | 528 |
| Configuring CHAP for PPP Authentication | 529 |
| Verifying CHAP Configuration | 532 |
| Troubleshooting Serial Connections | 534 |
| Summary | 535 |
| Establishing a WAN Connection Using Frame Relay | 537 |
| Understanding Frame Relay | 538 |
| Frame Relay Topologies | 541 |
| Frame Relay Reachability Issues | 542 |
| Frame Relay Signaling | 544 |
| Frame Relay Address Mappings | 545 |
| Configuring Frame Relay | 548 |
| Point-to-Point vs. Multipoint | 550 |
| Configuring Point-to-Point Frame Relay | 551 |
| Configuring Multipoint Frame Relay | 553 |
| Verifying Frame Relay Configuration | 555 |
| Summary | 559 |
| Introducing VPN Solutions | 561 |
| VPNs and Their Benefits | 561 |
| Cisco SSL VPN Solutions | 564 |
| Introducing IPsec | 565 |
| Summary | 567 |
| Configuring GRE Tunnels | 569 |
| GRE Tunnel Overview | 569 |
| GRE Tunnel Configuration | 571 |
| GRE Tunnel Verification | 574 |
| Summary | 575 |
| Module Summary | 577 |
| Module Self-Check | 579 |
| <i>Network Device Management</i> | 61 |
| Configuring Network Devices to Support Network Management Protocols | 63 |
| SNMP Overview | 64 |
| SNMP Versions | 65 |
| Obtaining Data from an SNMP Agent | 66 |
| SNMP Configuration | 68 |
| Syslog Overview | 610 |

| | |
|--|------------|
| Syslog Message Format | 611 |
| Syslog Configuration | 613 |
| NetFlow Overview | 614 |
| NetFlow Architecture | 618 |
| NetFlow Configuration | 619 |
| Summary | 623 |
| Managing Cisco Devices | 625 |
| Router Internal Components Overview | 626 |
| ROM Functions | 628 |
| Stages of the Router Power-On Boot Sequence | 629 |
| Configuration Register | 631 |
| Changing the Configuration Register | 632 |
| Locating Cisco IOS Image Files | 633 |
| Loading Cisco IOS Image Files | 635 |
| Loading Cisco IOS Configuration Files | 638 |
| Cisco IOS Integrated File System and Devices | 640 |
| Managing Cisco IOS Images | 642 |
| Deciphering Cisco IOS Image Filenames | 643 |
| Creating the Cisco IOS Image Backup | 644 |
| Upgrading Cisco IOS Images | 647 |
| Managing Device Configuration Files | 650 |
| Password Recovery | 653 |
| Summary | 656 |
| Licensing | 657 |
| Licensing Overview | 657 |
| Licensing Verification | 662 |
| Permanent License Installation | 663 |
| Evaluation License Installation | 665 |
| Backing up the License | 667 |
| Uninstalling the License | 668 |
| Summary | 670 |
| Module Summary | 671 |
| Module Self-Check | 673 |
| | 673 |
| Glossary | G1 |

Do Not Duplicate.
Post beta, not for release.

Course Introduction

Overview

Interconnecting Cisco Networking Devices, Part 2 (ICND2) v2.0 is an instructor-led course that is presented by Cisco Learning Partners to their end-user customers. This five-day course focuses on using Cisco Catalyst switches and Cisco routers that are connected in LANs and WANs typically found at medium-sized network sites.

Upon completing this training course, you should be able to configure, verify, and troubleshoot the various Cisco networking devices.

Learner Skills and Knowledge

These are the skills and knowledge that learners must possess to benefit fully from the course:

Learner Skills and Knowledge

- Understand network fundamentals
- Implement local-area networks
- Implement Internet connectivity
- Manage network device security
- Implement WAN connectivity
- Implement basic IPv6 connectivity

Course Goal and Objectives

Course Goal

To install, operate, and troubleshoot a medium-sized network, including connecting to a WAN and implementing network security

© 2013 Cisco Systems, Inc.

Do Not Duplicate
Post beta, not for release.

Course Flow

This topic presents the suggested flow of the course materials.

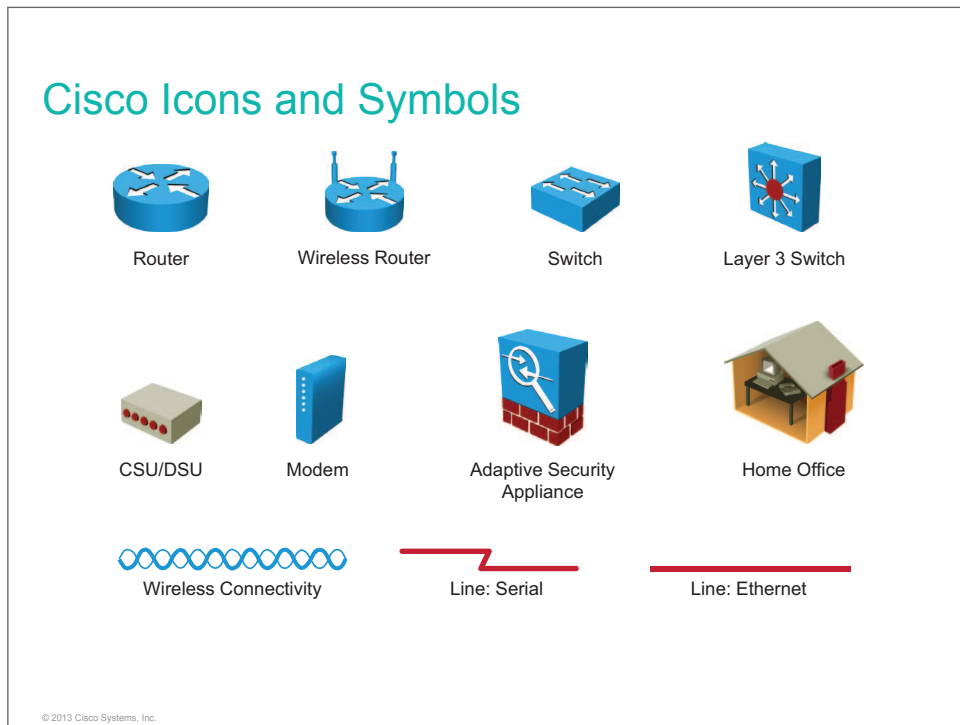
| | AM | PM |
|--------------|---|--|
| Day 1 | Course Intro Review of ICND1 | Implementing Scalable Medium-Sized Networks |
| Day 2 | Implementing Scalable Medium-Sized Networks (Cont.) Troubleshooting Basic Connectivity | Troubleshooting Basic Connectivity (Cont.) Implementing an EIGRP-Based Solution |
| Day 3 | Implementing an EIGRP-Based Solution (Cont.) Implementing a Scalable, Multiarea Network OSPF-Based Solution | Implementing a Scalable, Multiarea Network OSPF-Based Solution (Cont.) |
| Day 4 | Implementing a Scalable, Multiarea Network OSPF-Based | Wide-Area Networks |

© 2013 Cisco Systems, Inc.

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.

Additional References

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



Cisco Glossary of Terms

For additional information on Cisco terminology, refer to the Cisco Internetworking Terms and Acronyms glossary of terms at http://docwiki.cisco.com/wiki/Category:Internetworking_Terms_and_Acronyms%28ITA%29.

Your Training Curriculum

This topic presents the training curriculum for this course.



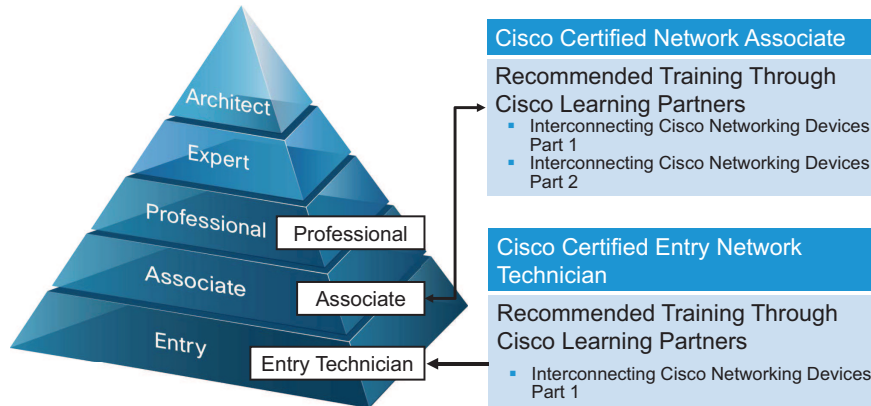
You are encouraged to join the Cisco Certification Community, a discussion forum open to anyone holding a valid Cisco Career Certification (such as Cisco CCIE[®], CCNA[®] Routing and Switching, CCDA[®], CCNP[®], CCIP[®], CCNP[®] Security, and CCNP Voice). It provides a gathering place for Cisco certified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

Training Curriculum

This subtopic presents the training curriculum for this course.

Cisco Career Certifications

Expand Your Professional Options, Advance Your Career.



© 2013 Cisco Systems, Inc.

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

CCNA Prep Center

Save on Cisco CCNA Study

Get the Cisco CCNA Premium Study Bundle and save up to 37 percent.

New to Certifications? What best describes your background?

Start a Networking Career | Military | Student | IT Professional

CERTIFICATIONS HIGHLIGHTS

View Certifications, review our study and practice materials or join a Study Group.

- Introducing the new CCNA Data Center and CCIP Data Center certifications.
- Schedule your CCIE Data Center lab exam now. Avoid the rush.
- Now Available On-demand - The Essentials of CCIE Webinar: Access Now
- CCIE Security written and lab exams v4.0 now available for registration
- Review the updated score reporting process for the CCDE practical exam.
- Retirement of CCIP certification

LEARNING NEWS

Information on the latest happenings on the site.

- What is the IT industry demanding? View this video to find out.
- Congratulations to our new Community Spotlight Award Recipients!
- Check out the Games Arcade, Cisco Learning Network's November Page of the Month
- New Users Getting Started Guide and Video Tour
- Now Available on Demand: Cisco Exam Preparation - Studying for Results

Latest Poll

Is Cisco Learning Network meeting your expectations?

Strongly Agree (46%)
Voters: 60/109

Agree (35%)
Voters: 60/109

Somewhat Agree (10%)
Voters: 20/109

© 2013 Cisco Systems, Inc.

Additional information is available at <http://learningnetwork.cisco.com>.

Learner Introductions

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



© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Implementing Scalable Medium-Sized Networks

This module starts with a review of VLAN and trunk technology. Understanding how VLANs and trunks operate and which protocols are associated with them is important for configuring, verifying, and troubleshooting VLANs and trunks on Cisco access switches. Switched networks introduce redundancy, so an STP loop-avoidance mechanism is needed to prevent undesirable loops. The module also explains EtherChannel technology, which groups several physical interfaces into one logical channel, and the router redundancy process, which solves problems in local networks with redundant topologies.

Objectives

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

- Troubleshoot VLAN connectivity
- Explain how STP works
- Configure link aggregation using EtherChannel
- Describe the purpose of Layer 3 redundancy protocols

Do Not Duplicate.
Post beta, not for release.

Troubleshooting VLAN Connectivity

This lesson addresses some of the common reasons that port connectivity, VLAN configuration, and trunk establishment can fail. It describes the information to look for to identify the source of the problem and determine how to solve it.

Objectives

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


- Describe VLANs
- Create and verify VLANs
- Describe trunks
- Configure and verify trunks
- Describe DTP
- Troubleshoot VLANs
- Troubleshoot trunks

VLAN Overview

This topic describes the basic principles of VLANs.

VLAN Overview

- A VLAN has these characteristics:
 - An independent LAN network
 - A broadcast domain
 - A logical network (subnet)
- VLANs address these needs:
 - Segmentation
 - Security
 - Network flexibility



A VLAN is a group of end stations with a common set of requirements, independent of their physical location. A VLAN has the same attributes as a physical LAN, except that it lets you group end stations even when they are not physically located on the same LAN segment. A VLAN also lets you group ports on a switch so that you can limit unicast, multicast, and broadcast traffic flooding. Flooded traffic that originates from a particular VLAN floods only to the ports belonging to that VLAN. VLAN trunks with IEEE 802.1Q tagging facilitate interswitch communication with multiple VLANs.

A VLAN is a logical broadcast domain that can span multiple physical LAN segments. Within the switched internetwork, VLANs provide segmentation and organizational flexibility. You can design a VLAN structure that lets you group stations that are segmented logically by functions, project teams, and applications without regard to the physical location of the users. Ports in the same VLAN share broadcasts. Ports in different VLANs do not share broadcasts. Containing broadcasts within a VLAN improves the overall performance of the network.

Each VLAN that you configure on the switch implements address learning, forwarding, and filtering decisions and loop-avoidance mechanisms, as if the VLAN were a separate physical bridge. The Cisco Catalyst switch implements VLANs by restricting traffic forwarding to destination ports that are in the same VLAN as the originating ports. When a frame arrives on a switch port, the switch must retransmit the frame only to the ports that belong to the same VLAN. In essence, a VLAN that is operating on a switch limits the transmission of unicast, multicast, and broadcast traffic.

A port normally carries only the traffic for the single VLAN to which it belongs. For a VLAN to span across multiple switches, a trunk is required to connect two switches. A trunk can carry traffic for multiple VLANs.

A VLAN can exist on a single switch or span multiple switches. VLANs can include stations in a single- or multiple-building infrastructures. The process of forwarding network traffic from one VLAN to another VLAN using a router is called *inter-VLAN routing*.

Cisco Catalyst switches have a factory default configuration in which various default VLANs are preconfigured to support various media and protocol types. The default Ethernet VLAN is VLAN 1.

If you want to communicate with the Cisco Catalyst switch remotely for management purposes, the switch must have an IP address. This IP address must be in the management VLAN, which by default is VLAN 1.

Creating VLANs

This topic describes how to create and verify VLANs.

Creating VLANs

```
SwitchX#configure terminal
SwitchX(config)#vlan 2
SwitchX(config-vlan)#name switchlab99
```

- Adds VLAN 2 and names it "switchlab99"

```
SwitchX#configure terminal
SwitchX(config)#interface FastEthernet 0/2
SwitchX(config-if)#switchport access vlan 2
```

- Assigns interface FastEthernet 0/2 to VLAN 2.

© 2013 Cisco Systems, Inc.

The table lists commands to use when adding a VLAN.

| Command | Description |
|------------------------------|---|
| vlan <i>vlan-id</i> | The VID to be added and configured. For <i>vlan-id</i> , the range is 1 to 4094. You can enter a single VID, a series of VIDs separated by commas, or a range of VIDs separated by hyphens. |
| name <i>vlan-name</i> | (Optional) Specifies the VLAN name, an ASCII string from 1 to 32 characters that must be unique within the administrative domain. |

For many Cisco Catalyst switches, you use the **vlan** global configuration command to create a VLAN and enter VLAN configuration mode. Use the **no** form of this command to delete the VLAN. The example in the figure shows how to add VLAN 2 to the VLAN database and how to name it “switchlab99.”

To add a VLAN to the VLAN database, assign a number and name to the VLAN. VLAN 1 is the factory default VLAN. Normal-range VLANs are identified with a number between 1 and 1001.

To add an Ethernet VLAN, you must specify at least a VLAN number. If no name is entered for the VLAN, the default is to append the VLAN number to the command **vlan**. For example, VLAN0004 would be the default name for VLAN 4 if no name is specified.

When an end system is connected to a switch port, it should be associated with a VLAN in accordance with the network design. To associate a device with a VLAN, the switch port to which the device connects is assigned to a single data VLAN and thus becomes an access port. A switch port can become an access port through static or dynamic configuration.

After creating a VLAN, you can manually assign a port or a number of ports to that VLAN. A port can belong to only one VLAN at a time. When you assign a switch port to a VLAN using this method, it is known as a static access port.

On most Cisco Catalyst switches, you configure the VLAN port assignment from interface configuration mode using the **switchport access vlan** command. To configure a bundle of interfaces to a VLAN, use the **interface range** command. Use the **vlan *vlan_number*** command to set static access membership.

Creating VLANs (Cont.)

```
SwitchX#show vlan
VLAN Name                Status    Ports
-----
1  default                 active    Fa0/1
2  switchlab99             active    Fa0/2
3  1002 fddi-default       act/unsup
<output omitted>
```

- Displays information on all configured VLANs.

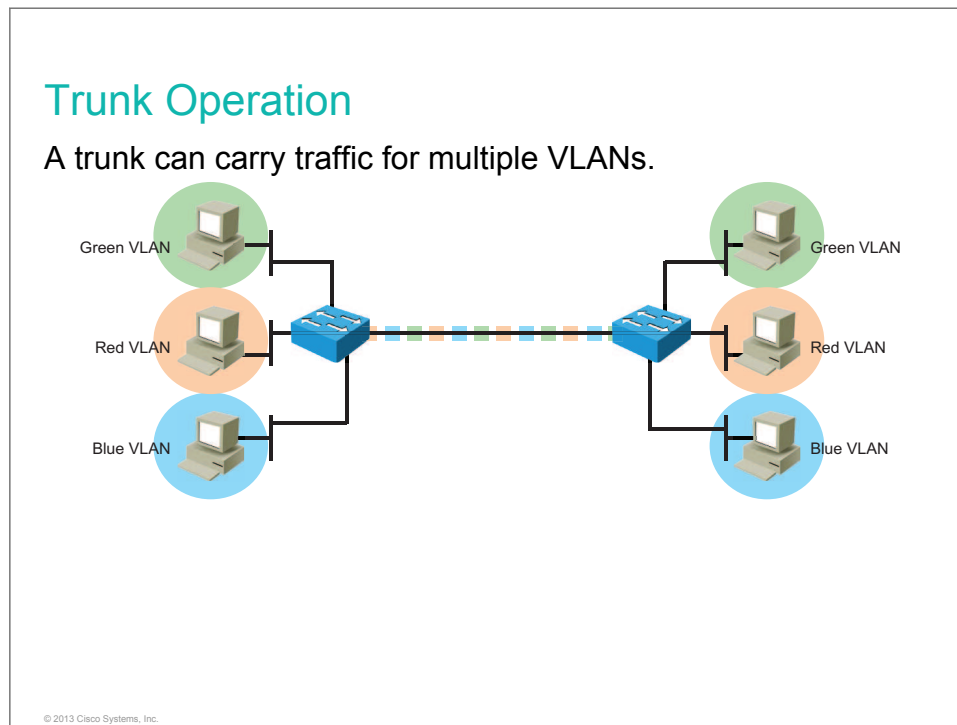
After you configure the VLAN, validate the parameters for that VLAN.

Use the **show vlan** command to display information on all configured VLANs. The command displays configured VLANs, their names, and the ports on the switch that are assigned to each VLAN.

Use the **show vlan id *vlan_number*** or the **show vlan name *vlan-name*** command to display information about a particular VLAN.

Trunk Operation

This topic describes basic principles of trunks.



A port normally carries only the traffic for the single VLAN to which it belongs. For a VLAN to span across multiple switches, a trunk is required to connect two switches. A trunk can carry traffic for multiple VLANs.

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device, such as a router or a switch. Ethernet trunks carry the traffic of multiple VLANs over a single link and allow you to extend the VLANs across an entire network. A trunk does not belong to a specific VLAN—rather, it is a conduit for VLANs between switches and routers.

A special protocol is used to carry multiple VLANs over a single link between two devices. Cisco supports the IEEE 802.1Q trunking protocol. A trunk could also be used between a network device and server or other device that is equipped with an appropriate 802.1Q-capable NIC.

Ethernet trunk interfaces support various trunking modes. You can configure an interface as trunking or nontrunking, or you can have it negotiate trunking with the neighboring interface.

By default on a Cisco catalyst switch, all configured VLANs are carried over a trunk interface. On an 802.1Q trunk port, there is one native VLAN that is untagged (by default, VLAN 1). All other VLANs are tagged with a VID.

When Ethernet frames are placed on a trunk, they need additional information about the VLANs that they belong to. This task is accomplished by using the 802.1Q encapsulation header. IEEE 802.1Q uses an internal tagging mechanism that inserts a 4-byte tag field into the original Ethernet frame between the Source Address and Type or Length fields. Because 802.1Q alters the frame, the trunking device recomputes the FCS on the modified frame. It is the responsibility of the Ethernet switch to look at the 4-byte tag field and determine where to deliver the frame.

Configuring Trunks

This topic describes how to configure and verify IEEE 802.1Q trunks.

Configuring Trunks

- Enter interface configuration mode.
- Configure the Fa0/11 interface as a VLAN trunk.
- The native VLAN is changed to VLAN 99.

```
SwitchX#configure terminal
SwitchX (config)#interface fa0/11
SwitchX (config-if)#switchport mode trunk
SwitchX (config-if)#switchport trunk native vlan 99
```

© 2013 Cisco Systems, Inc.

Configuring Trunk Commands

| Command | Description |
|--|---|
| interface <i>interface interface_number</i> | Enters interface configuration mode for the specified interface. |
| switch mode trunk | Sets the interface to trunk mode. |
| switchport trunk native vlan <i>vlan_number</i> | Sets the native VLAN on the trunk to the specified VLAN number. Traffic from this VLAN is sent untagged. You must ensure that the other end of the trunk link is configured the same way. |

Configuring Trunks (Cont.)

```
SwitchX#show interfaces FastEthernet 0/11 switchport
Name: Fa0/11
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 99
Trunking Native Mode VLAN: 99
<output omitted>
```

- Verifies switchport settings on FastEthernet 0/11

```
SwitchX#show interfaces FastEthernet 0/11 trunk
Port      Mode      Encapsulation  Status Native vlan
Fa0/11    on        802.1q         trunking  99
<output omitted>
```

- Verifies that FastEthernet 0/11 is trunking

© 2013 Cisco Systems, Inc.

To verify a trunk configuration on many Cisco Catalyst switches, use the **show interfaces switchport** and **show interfaces trunk** commands. These two commands display the trunk parameters and VLAN information of the port.

Dynamic Trunking Protocol

This topic explains the purpose of DTP and switchport modes.

Dynamic Trunking Protocol

Switchport mode interactions:

- Manual configuration is recommended.
- Configure the port as trunk or access on both switches.
- The command **nonegotiate** disables negotiation (default).

| | Dynamic Auto | Dynamic Desirable | Trunk | Access |
|-------------------|--------------|-------------------|----------------------|----------------------|
| Dynamic auto | Access | Trunk | Trunk | Access |
| Dynamic desirable | Trunk | Trunk | Trunk | Access |
| Trunk | Trunk | Trunk | Trunk | Limited connectivity |
| Access | Access | Access | Limited connectivity | Access |

© 2013 Cisco Systems, Inc.

Many Cisco Catalyst switches support DTP, which manages automatic trunk negotiation. DTP is a Cisco proprietary protocol. Switches from other vendors do not support DTP. DTP is automatically enabled on a switch port when certain trunking modes are configured on the switch port. DTP manages trunk negotiation only if the port on the other switch is configured in a mode that supports DTP.

You should configure trunk links statically whenever possible. However, Cisco switch ports can run DTP, which can automatically negotiate a trunk link. This protocol can determine an operational trunking mode and protocol on a switch port when it is connected to another device that is also capable of dynamic trunk negotiation.

The default DTP mode is dependent on the Cisco IOS Software version and on the platform. To determine the current DTP mode, issue the command **show dtp interface**.

```
Switch#show dtp interface fa0/1
DTP information for FastEthernet0/1:
  TOS/TAS/TNS:                TRUNK/DESIRABLE/TRUNK
  TOT/TAT/TNT:                802.1Q/802.1Q/802.1Q
  Neighbor address 1:         001646FA9B01
  Neighbor address 2:         000000000000
  Hello timer expiration (sec/state): 17/RUNNING
  Access timer expiration (sec/state) 287/RUNNING
<output omitted>
```

Note A general best practice is to set the interface to **trunk** and **nonegotiate** when a trunk link is required. On links where trunking is not intended, DTP should be turned off.

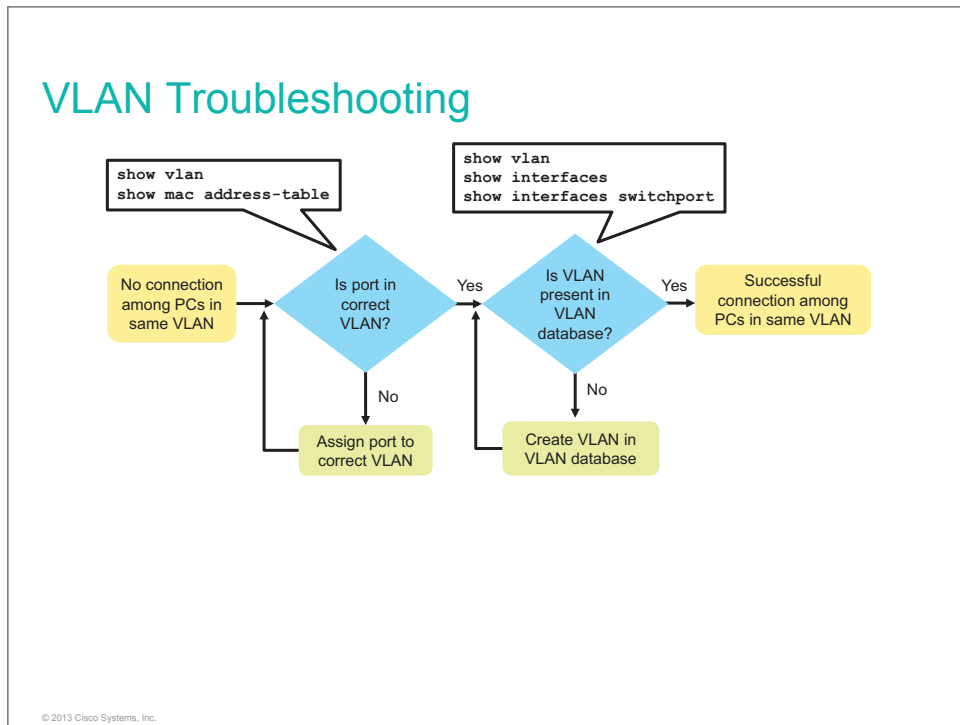
You can configure DTP mode to turn the protocol off or to instruct it to negotiate a trunk link only under certain conditions, as described in the table.

| Mode | Function |
|-------------------|--|
| dynamic auto | Creates the trunk based on the DTP request from the neighboring switch. |
| dynamic desirable | Communicates to the neighboring switch via DTP that the interface is attempting to become a trunk if the neighboring switch interface is able to become a trunk. |
| trunk | Automatically enables trunking regardless of the state of the neighboring switch and regardless of any DTP requests sent from the neighboring switch. |
| access | Trunking not allowed on this port regardless of the state of the neighboring switch interface and regardless of any DTP requests sent from the neighboring switch. |
| nonegotiate | Prevents the interface from generating DTP frames. This command can be used only when the interface switch port mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link. |

The **switchport nonegotiate** interface command specifies that DTP negotiation packets are not sent. The switch does not engage in DTP negotiation on this interface. This command is valid only when the interface switchport mode is access or trunk (configured by using the **switchport mode access** or the **switchport mode trunk** interface configuration command). This command returns an error if you attempt to execute it in dynamic (auto or desirable) mode. Use the **no** form of this command to return to the default setting. When you configure a port with the **switchport nonegotiate** command, the port trunks only if the other end of the link is specifically set to trunk. The **switchport nonegotiate** command does not form a trunk link with ports in either dynamic desirable or dynamic auto mode.

VLAN Troubleshooting

This topic describes how to troubleshoot common VLAN issues.



To troubleshoot VLAN issues when you have no connection between PCs, follow these high-level steps:

1. Use the **show vlan** command to check whether the port belongs to the expected VLAN. If the port is assigned to the wrong VLAN, use the **switchport access vlan** command to correct the VLAN membership.

Use the **show mac address-table** command to check which addresses were learned on a particular port of the switch and to which VLAN that port is assigned.

2. If the VLAN to which the port is assigned is deleted, the port becomes inactive. Use the **show vlan** or **show interfaces switchport** command to verify that the VLAN is present in the VLAN database.

VLAN Troubleshooting (Cont.)

MAC address table verification.

```
SW1#show mac address-table interface FastEthernet 0/1
      Mac Address Table
-----
Vlan    Mac Address      Type        Ports
----    -
10      000c.296a.a21c   DYNAMIC     Fa0/1
10      000f.34f9.9181   DYNAMIC     Fa0/1
Total Mac Addresses for this criterion: 2
```

© 2013 Cisco Systems, Inc.

To display the MAC address table, use the **show mac-address-table** command in privileged EXEC mode. This command displays the MAC address table for the switch. Specific views can be defined by using the optional keywords and arguments. The example shows MAC addresses that were learned on the FastEthernet0/1 interface. It can be seen that MAC address 000c.296a.a21c was learned on the interface FastEthernet0/1 in VLAN 10. If this number is not the expected VLAN number, change the port VLAN membership using the **switchport access vlan** command.

VLAN Troubleshooting (Cont.)

Troubleshooting missing VLANs.

```
SW1#show interfaces FastEthernet 0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 10 (Inactive)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
<output omitted>
```

© 2013 Cisco Systems, Inc.

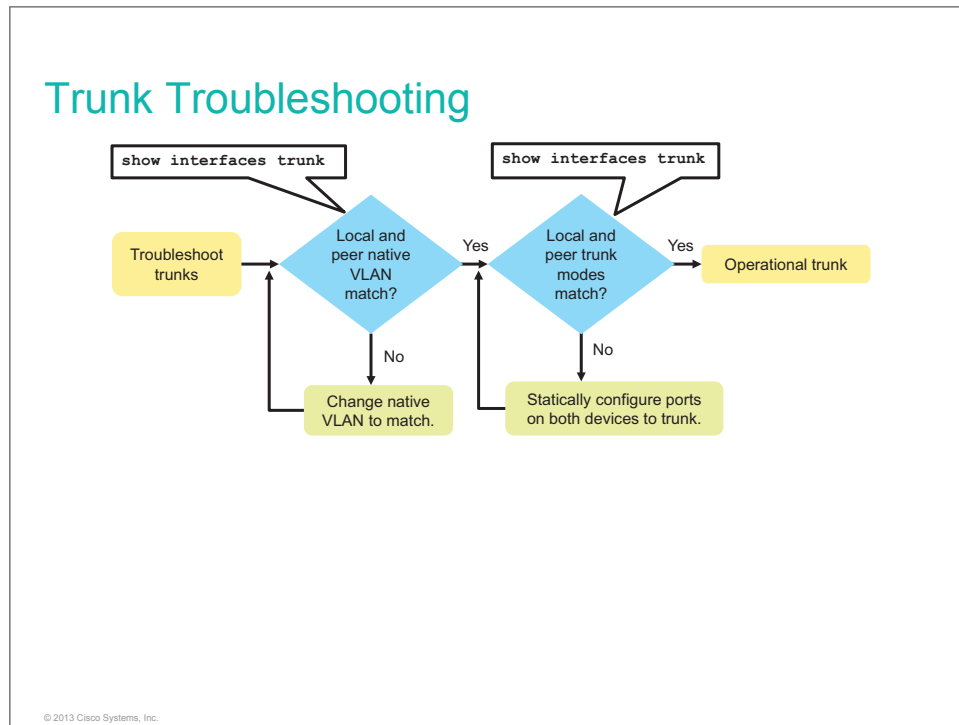
Each port in a switch belongs to a VLAN. If the VLAN to which a port belongs is deleted, the port becomes inactive. All ports belonging to the VLAN that was deleted are unable to communicate with the rest of the network.

Use the command **show interface interface switchport** to check whether the port is inactive. If the port is inactive, it will not be functional until the missing VLAN is created using the **vlan vlan_id** command.

Do Not Duplicate.
Post beta, not for release.

Trunk Troubleshooting

This topic describes how to troubleshoot common trunk issues.



To troubleshoot trunk issues when the trunk is not established or "VLAN leaking" is occurring, follow these high-level steps:

1. Use the **show interfaces trunk** command to check whether a trunk has been established between switches. You should statically configure trunk links whenever possible. However, Cisco Catalyst switch ports, by default, run DTP, which tries to negotiate a trunk link.
2. Use the **show interfaces trunk** command to check whether the local and peer native VLANs match. If the native VLAN does not match on both sides, VLAN leaking occurs.

Trunk Troubleshooting (Cont.)

```
SW1#show interfaces FastEthernet 0/3 trunk
Port      Mode      Encapsulation  Status  Native vlan
Fa0/3     auto      802.1q         not-trunking  2
<output omitted>
```

- Verifies switchport mode, trunk establishment, and the native VLAN on SW1

```
SW2#show interfaces FastEthernet 0/3 trunk
Port      Mode      Encapsulation  Status  Native vlan
Fa0/3     auto      802.1q         not-trunking  1
<output omitted>
```

- Verifies switchport mode, trunk establishment, and the native VLAN on SW2

© 2013 Cisco Systems, Inc.

To display the status of the trunk and native VLAN used on that trunk link, and to verify trunk establishment, use the **show interface trunk** command in privileged EXEC mode. The example shows that the native VLAN on one side of the trunk link was changed to VLAN 2. If one end of the trunk is configured as native VLAN 1 and the other end is configured as native VLAN 2, a frame sent from VLAN 1 on one side is received on VLAN 2 on the other side. VLAN 1 “leaks” into the VLAN 2 segment and this results in connectivity issues. Change the native VLAN to the same VLAN on both sides of the VLAN to avoid this behavior.

Cisco Discovery Protocol notifies you of a native VLAN mismatch on a trunk link with this message:

```
Aug 31 08:34:48.714: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch
discovered on FastEthernet0/3 (2), with SW2 FastEthernet0/3 (1).
```

You should statically configure trunk links whenever possible. Cisco Catalyst switch ports, by default, run DTP, which can determine the operational trunking mode and protocol on a switch port when it is connected to another device that is also capable of dynamic trunk negotiation. If both ends of a trunk are set to dynamic auto trunk mode, a trunk will not be established. The example shows the status of the link as “not-trunking.”

Summary

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

Summary

- A VLAN is a logical broadcast domain that can span multiple physical LAN segments.
- A trunk can carry traffic for multiple VLANs.
- DTP can automatically negotiate a trunk link (not recommended).
- You should verify that the port is in the correct VLAN and that the VLAN is present in the VLAN database.
- You should verify that there is no native VLAN mismatch and that a trunk is established.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Building Redundant Switched Topologies

Most complex networks include redundant devices to avoid single points of failure. Although a redundant topology eliminates some issues, it can introduce other problems. STP is a Layer 2 link management protocol that provides path redundancy while preventing undesirable loops in a switched network.

This lesson identifies the problems that are caused by redundant switched-network topologies and the functions of STP that prevent these problems.

Objectives

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

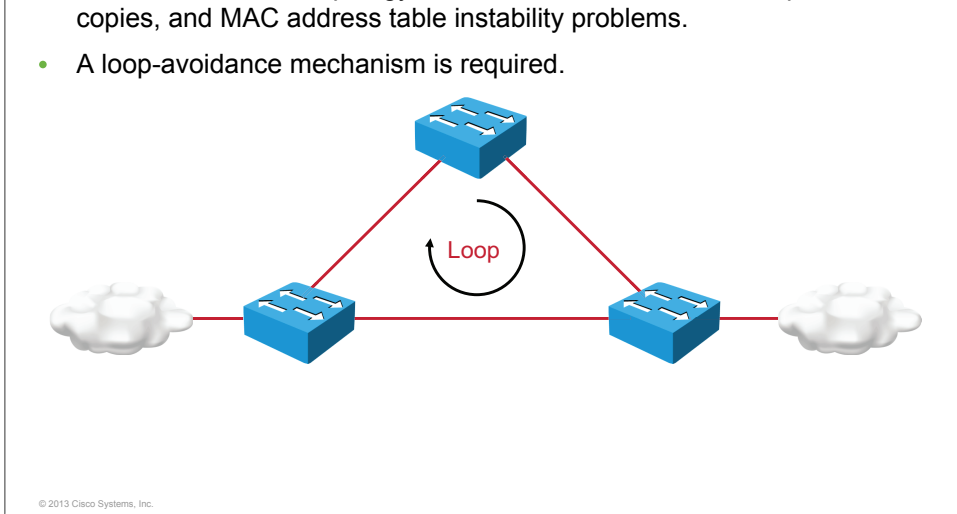
- Describe problems that may arise in redundant switched topologies
- Describe the principles behind STP
- Describe variants of STP and the differences between them
- Explain how PVST+ improves on the concept of STP
- Describe how to make a switch the root bridge
- Describe how to use Cisco IOS commands to analyze the spanning-tree topology and verify the proper operation of STP
- Describe typical symptoms of a major spanning-tree failure and how to recover from that failure
- Demonstrate how to configure and verify PortFast and BPDU guard

Issues in Redundant Topologies

This topic describes how to provide redundant links and devices in switched networks.

Issues in Redundant Topologies

- A redundant topology eliminates single points of failure.
- A redundant switch topology causes broadcast storms, multiple frame copies, and MAC address table instability problems.
- A loop-avoidance mechanism is required.



Redundant designs can eliminate the possibility of a single point of failure causing a loss of function for the entire switched network. However, you must consider some of the problems that redundant designs can cause:

- **Broadcast storms:** Without some loop-avoidance process, each switch floods broadcasts endlessly. This situation is commonly called a broadcast storm.
- **Multiple frame transmission:** Multiple copies of unicast frames may be delivered to destination stations. Many protocols expect to receive only a single copy of each transmission. Multiple copies of the same frame can cause unrecoverable errors.
- **MAC database instability:** Instability in the content of the MAC address table results from copies of the same frame being received on different ports of the switch. Data forwarding can be impaired when the switch consumes the resources that are coping with instability in the MAC address table.

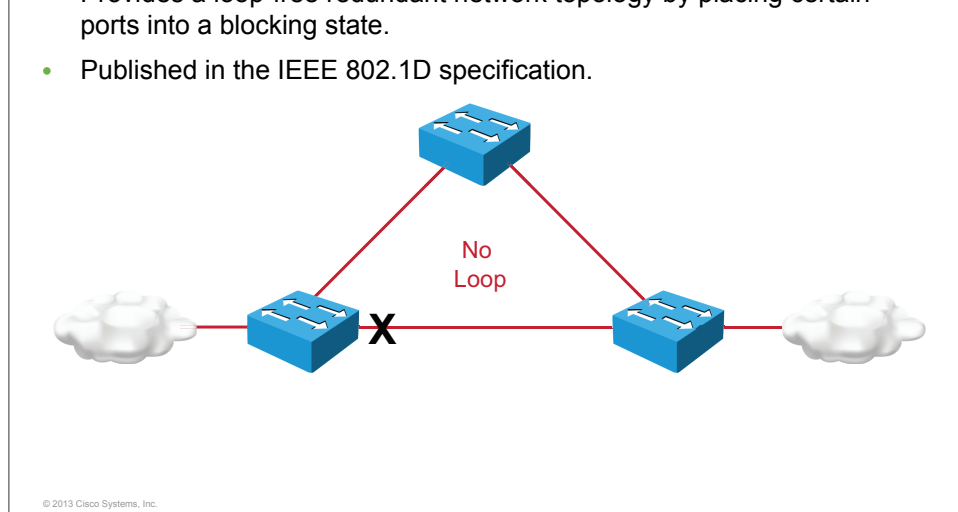
Layer 2 LAN protocols, such as Ethernet, lack a mechanism to recognize and eliminate endlessly looping frames. Some Layer 3 protocols implement a TTL mechanism that limits the number of times that a Layer 3 networking device can retransmit a packet. Lacking such a mechanism, Layer 2 devices continue to retransmit looping traffic indefinitely.

A loop-avoidance mechanism solves these problems. STP was developed to address them.

Issues in Redundant Topologies (Cont.)

Loop resolution with Spanning Tree Protocol:

- Provides a loop-free redundant network topology by placing certain ports into a blocking state.
- Published in the IEEE 802.1D specification.



STP provides loop resolution by managing the physical paths to given network segments. STP allows physical path redundancy while preventing the undesirable effects of active loops in the network. STP is an IEEE committee standard defined as 802.1D.

STP behaves as follows:

- STP uses BPDUs for communication between switches.
- STP forces certain ports into a standby state so that they do not listen to, forward, or flood data frames. The overall effect is that there is only one path to each network segment that is active at any time.
- If there is a problem with connectivity to any of the segments within the network, STP re-establishes connectivity by automatically activating a previously inactive path, if one exists.

Spanning-Tree Operation

This topic describes STP operation using a simple case study.

Spanning-Tree Operation

The spanning-tree algorithm follows these steps:

1. Elects a root bridge
2. Elects a root port for each non-root switch
3. Elects a designated port for each segment
4. Ports transition to forwarding or blocking state

© 2013 Cisco Systems, Inc.

STP and its successor protocols provide loop resolution by managing the physical paths to given network segments. STP allows physical path redundancy while preventing the undesirable effects of active loops in the network. STP forces certain ports into a blocking state. These blocking ports do not forward data frames. The overall effect is that there is only one path to each network segment that is active at any time. If there is a problem with connectivity to any of the segments within the network, STP re-establishes connectivity by automatically activating a previously inactive path, if one exists.

These are the steps of the spanning-tree algorithm:

- 1 Elects a root bridge. The root bridge becomes the switch with the lowest BID. There can be only one root bridge per network. Bridge ID is a combination of bridge priority and the MAC address of the switch. Bridge priority is a number between 0 and 65535, and the default is 32768.
- 2 Elects a root port for each non-root switch, based on the lowest root path cost. The root bridge does not have root ports. Each non-root switch has one root port. The root port shows the direction of the best path to the root bridge.
- 3 Elects a designated port for each segment, based on the lowest root path cost. Each link will have one designated port.
- 4 The root ports and designated ports transition to the forwarding state, and the other ports stay in the blocking state.

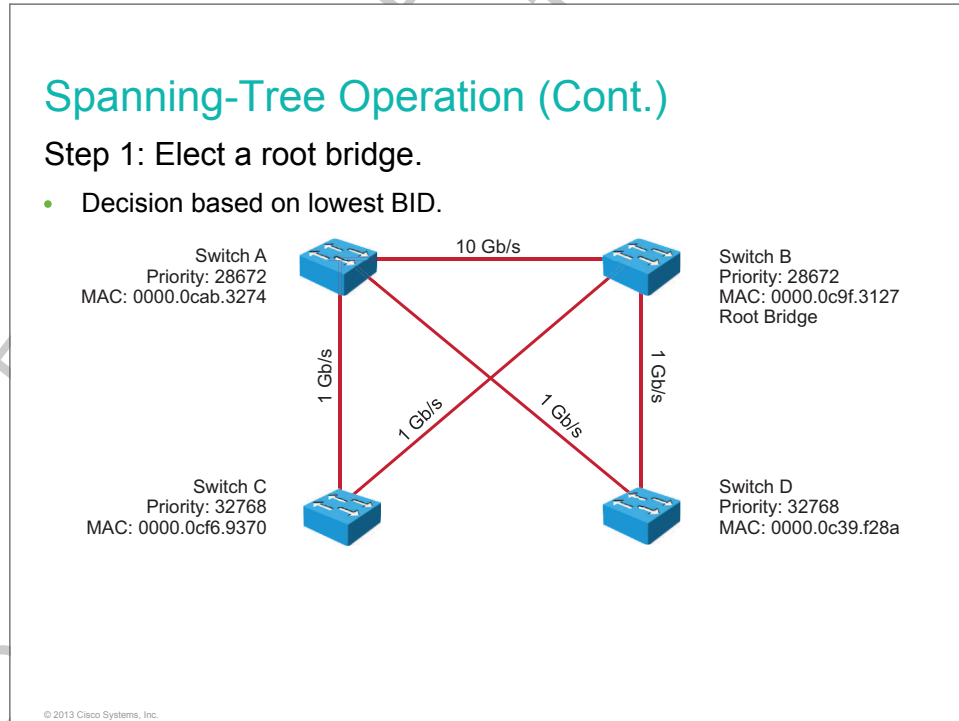
STP path cost depends on the speed of the link. The table shows STP link costs.

| Data rate | STP Cost (802.1D-1998) |
|-----------|------------------------|
| 4 Mb/s | 250 |
| 10 Mb/s | 100 |

| Data rate | STP Cost (802.1D-1998) |
|-----------|------------------------|
| 16 Mb/s | 62 |
| 100 Mb/s | 19 |
| 1 Gb/s | 4 |
| 2 Gb/s | 3 |
| 10 Gb/s | 2 |

STP Port Roles

| Port Role | Description |
|--------------------|--|
| Root port | This port exists on non-root bridges and is the switch port with the best path to the root bridge. Root ports forward traffic toward the root bridge, and the source MAC address of frames received on the root port is capable of populating the MAC table. Only one root port is allowed per bridge. |
| Designated port | This port exists on root and non-root bridges. For root bridges, all switch ports are designated ports. For non-root bridges, a designated port is the switch port that will receive and forward frames toward the root bridge as needed. Only one designated port is allowed per segment. If multiple switches exist on the same segment, an election process determines the designated switch, and the corresponding switch port begins forwarding frames for the segment. Designated ports are capable of populating the MAC table. |
| Nondesignated port | The nondesignated port is a switch port that is not forwarding (blocking) data frames and is not populating the MAC address table with the source addresses of frames that are seen on that segment. |
| Disabled port | The disabled port is a switch port that is shut down. |

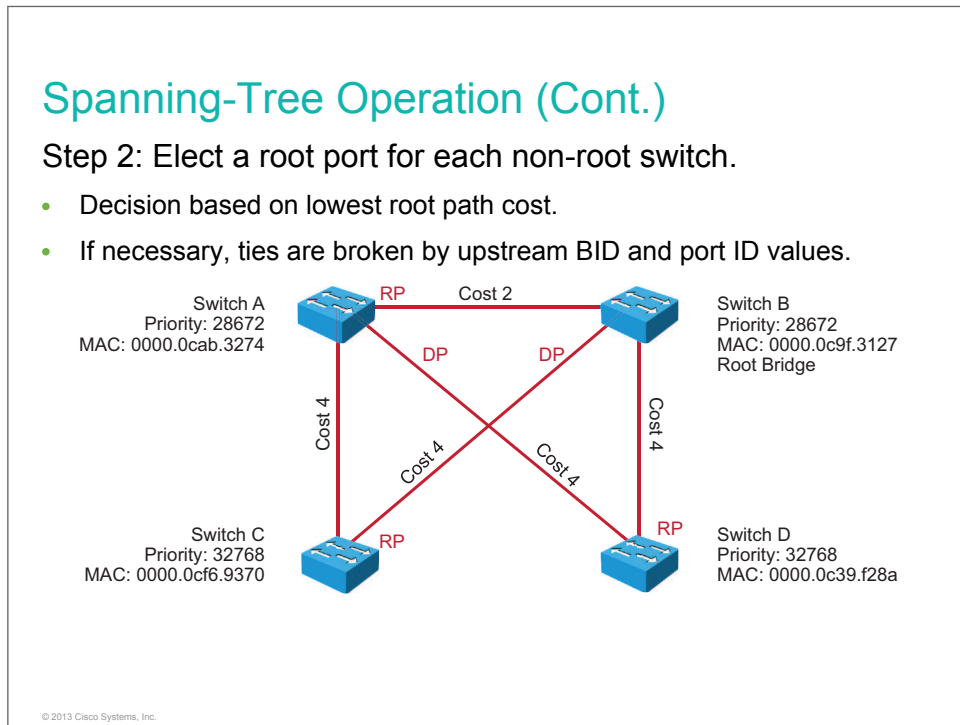


The first step in the spanning-tree algorithm is the election of a root bridge. Initially, all switches assume that they are the root. They start transmitting BPDUs with the Root ID field containing the same value as the Bridge ID field. Thus, each switch essentially claims that it is the root bridge on the network.

As soon as the switches start receiving BPDUs from the other switches, each switch compares the root ID in the received BPDUs against the value that it currently has recorded as the root ID. If the received value is lower than the recorded value (which was originally the BID of that switch), the switch replaces the recorded value with the received value and starts transmitting this value in the Root ID field in its own BPDUs.

Eventually, all switches learn and record the BID of the switch that has the lowest BID, and the switches all transmit this ID in the Root ID field of their BPDUs.

Switch B in the example becomes the root bridge because it has the lowest BID. Switch A and Switch B have the same priority, but Switch B has a lower MAC value.



As soon as a switch recognizes that it is not the root (because it is receiving BPDUs that have a root ID value that is lower than its own BID), it marks the port on which it is receiving those BPDUs as its root port.

BPDUs could be received on multiple ports. In this case, the switch elects the port that has the lowest-cost path to the root as its root port. If two ports have an equal path cost to the root, the switch looks at the BID values in the received BPDUs to make a decision (where the lowest BID is considered best, similar to root bridge election). If the root path cost and the BID in both BPDUs are the same because both ports are connected to the same upstream switch, the switch looks at the Port ID field in the BPDUs and selects its root port based on the lowest value in that field.


The cost associated with each port is, by default, related to its speed (the higher the interface bandwidth, the lower the cost), but the cost can be manually changed.

Switches A, C, and D mark the ports directly connected to Switch B (which is the root bridge) as the root port. These directly connected ports on switches A, C, and D have the lowest cost to the root bridge.

Spanning-Tree Operation (Cont.)

Step 3: Elect a designated port for each segment.

- Decision is based on the lowest root path cost.
- If necessary, ties are broken by upstream BID and port ID.



After electing the root bridge and root ports, the switches determine which switch will become the designated bridge for each Ethernet segment. This process has similarities to the root bridge and root port elections. Each switch connected to a segment sends BPDUs out of the port that is connected to that segment, claiming to be the designated bridge for that segment. At this point, it considers its port to be a designated port.


As soon as a switch starts receiving BPDUs from other switches on that segment, it compares the received values of the root path cost, BID, and port ID fields (in that order) against the values in the BPDUs that it is sending out its own port. The switch stops transmitting BPDUs on the port and marks it as a nondesignated port if the other switch has lower values.

In the example, all ports on the root bridge (Switch B) are designated ports. The ports on Switch A connecting to Switch C and Switch D become designated ports, because they have lower root path costs on each segment.

Spanning-Tree Operation (Cont.)

Step 4: The ports transition to the forwarding or blocking state.

- Root ports and designated ports transition to the forwarding state.
- Other ports stay in the blocking state.



To prevent bridging loops during the time that STP needs to execute its algorithm, all ports start out in the blocking state. As soon as STP marks a port as either a root port or a designated port, the algorithm starts to transition this port to the forwarding state.

Classic (802.1D–1998) and rapid (802.1w and 802.1D–2004) versions of STP both execute the same algorithm in the decision-making process. However, in the transition of a port from the blocking (or discarding, in rapid spanning-tree terms) to the forwarding state, there is a big difference between those two spanning-tree versions. Classic 802.1D would simply take 30 seconds to transition the port to forwarding. The rapid spanning tree algorithm can leverage additional mechanisms to transition the port to forwarding in less than a second.

Although the order of the steps listed in the diagrams suggests that STP goes through them in a coordinated, sequential manner, that is not actually the case. If you look back at the description of each step in the process, you see that each switch is going through these steps in parallel and that it might adapt its selection of root bridge, root ports, and designated ports as new BPDUs are received. As the BPDUs are propagated through the network, all switches eventually have a consistent view of the topology of the network. When this stable state is reached, BPDUs are transmitted only by designated ports.

There are two loops in the sample topology, meaning that two ports should be in the blocking state to break both loops. The port on Switch C that is not directly connected to Switch B (root bridge) is blocked, because it is a nondesignated port. The port on Switch D that is not directly connected to Switch B (root bridge) is also blocked, because it is a nondesignated port.

Types of Spanning-Tree Protocols

This topic describes and compares various types of spanning-tree protocols.

Types of Spanning-Tree Protocols

Spanning-tree standards:

- **IEEE 802.1D:** The legacy standard for bridging and STP
 - **CST:** Assumes one spanning-tree instance for the entire bridged network, regardless of the number of VLANs
- **PVST+:** A Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN configured in the network
- **802.1w (RSTP):** Improves convergence over 1998 STP by adding roles to ports and enhancing BPDU exchanges
- **Rapid PVST+:** A Cisco enhancement of RSTP using PVST+

© 2013 Cisco Systems, Inc.

The STP is a network protocol that ensures a loop-free topology. There are several varieties of spanning-tree protocols:

- STP (IEEE 802.1D) provides a loop-free topology in a network with redundant links.
 - CST assumes one spanning-tree instance for the entire bridged network, regardless of the number of VLANs.
- PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN configured in the network.
- RSTP, or IEEE 802.1w, is an evolution of STP that provides faster convergence of STP. It redefines port roles and link costs.
- Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+. Rapid PVST+ provides a separate instance of 802.1w per VLAN.

Note When Cisco documentation and this course refer to implementing RSTP, they are referring to the Cisco RSTP implementation, Rapid PVST+.

Types of Spanning-Tree Protocols (Cont.)

| Protocol | Standard | Resources Needed | Convergence | Number of Trees |
|-------------|----------|------------------|-------------|--------------------|
| STP | 802.1D | Low | Slow | One |
| PVST+ | Cisco | High | Slow | One for every VLAN |
| RSTP | 802.1w | Medium | Fast | One |
| Rapid PVST+ | Cisco | Very high | Fast | One for every VLAN |

© 2013 Cisco Systems, Inc.

These are characteristics of various spanning-tree protocols:

- STP assumes one 802.1D spanning-tree instance for the entire bridged network, regardless of the number of VLANs. Because there is only one instance, the CPU and memory requirements for this version are lower than for the other protocols. However, because there is only one instance, there is only one root bridge and one tree. Traffic for all VLANs flows over the same path, which can lead to suboptimal traffic flows. Because of the limitations of 802.1D, this version is slow to converge.
- PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN that is configured in the network. The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard. Creating an instance for each VLAN increases the CPU and memory requirements but allows for per-VLAN root bridges. This design allows the STP tree to be optimized for the traffic of each VLAN. Convergence of this version is similar to the convergence of 802.1D. However, convergence is per-VLAN.
- RSTP, or IEEE 802.1w, is an evolution of STP that provides faster STP convergence. This version addresses many convergence issues, but because it still provides a single instance of STP, it does not address the suboptimal traffic flow issues. To support that faster convergence, the CPU usage and memory requirements of this version are slightly higher than those of STP but less than those of RSTP+.
- Rapid PVST+ is a Cisco enhancement of RSTP that uses PVST+. It provides a separate instance of 802.1w per VLAN. The separate instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard. This version addresses both the convergence issues and the suboptimal traffic flow issues. However, this version has the largest CPU and memory requirements.

Types of Spanning Tree Protocols (Cont.)

Default spanning tree configuration for Cisco Catalyst switches:

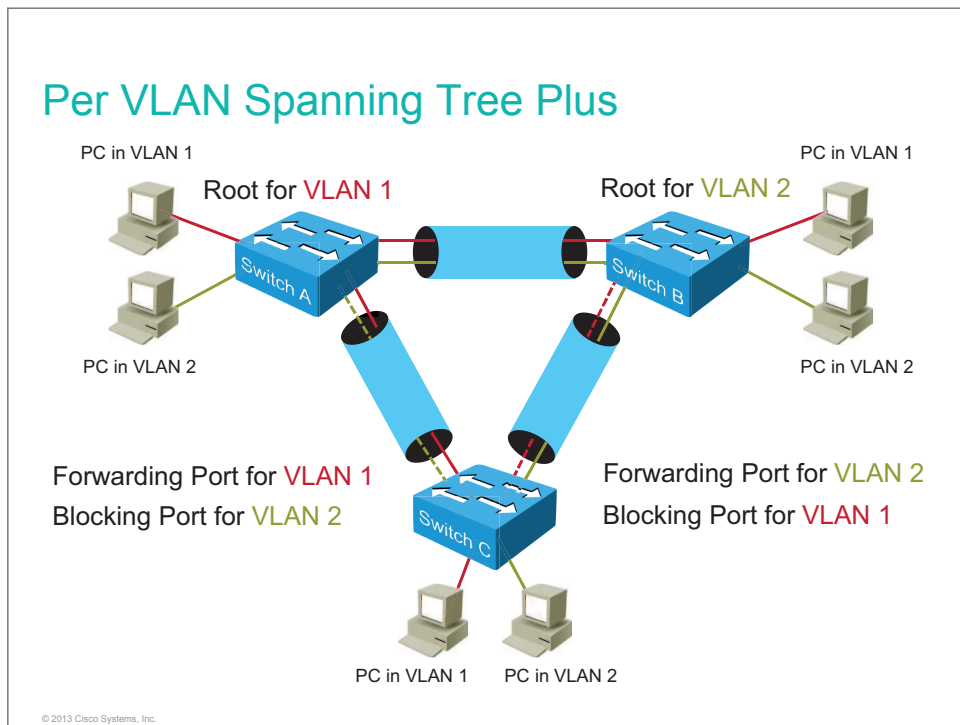
- PVST
- Enabled on all ports
- Slower convergence after topology change than with RSTP.

© 2013 Cisco Systems, Inc.

The default spanning-tree mode for Cisco Catalyst switches is PVST, which is enabled on all ports. PVST has much slower convergence after a topology change than the Rapid PVST.

Per VLAN Spanning Tree Plus

This topic describes the function of the PVST+ protocol.



The 802.1D standard defines a CST that assumes only one spanning-tree instance for the entire switched network, regardless of the number of VLANs. A network running CST has these characteristics:

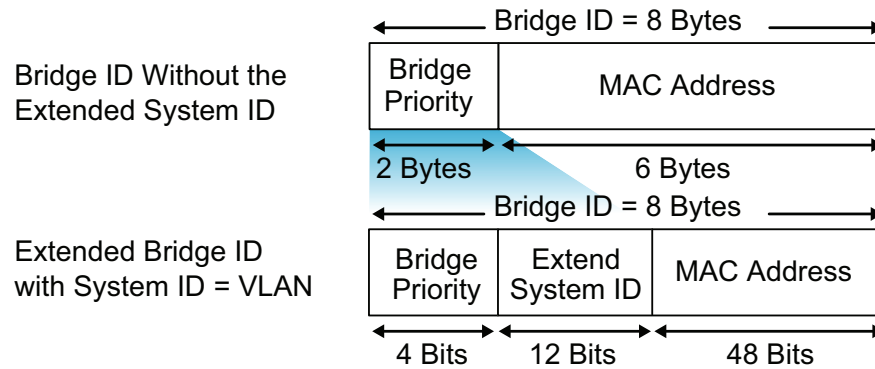
- No load sharing is possible. One uplink must block for all VLANs.
- The CPU is spared. Only one instance of spanning tree must be computed.

PVST+ defines a spanning-tree protocol that has several spanning-tree instances running for the network (one instance of STP per VLAN). Networks running several spanning-tree instances have these characteristics:

- Optimum load sharing can result.
- One spanning-tree instance for each VLAN maintained can mean a considerable waste of CPU cycles for all the switches in the network (in addition to the bandwidth that is used for each instance to send its own BPDUs). This would only be problematic if a large number of VLANs are configured.

Per VLAN Spanning Tree Plus (Cont.)

System ID = VLAN



In a Cisco PVST+ environment, you can tune the spanning-tree parameters so that half the VLANs forward on each uplink trunk. The network must be correctly configured. The configuration must define a different root bridge for each half of the VLANs. Providing different STP root switches per VLAN creates a more redundant network.

Spanning-tree operation requires that each switch has a unique BID. In the original 802.1D standard, the BID was composed of the bridge priority and the MAC address of the switch, and a CST represented all VLANs. PVST+ requires that a separate instance of spanning tree that is run for each VLAN and the BID field must carry VID information. This functionality is accomplished by reusing a portion of the Priority field as the extended system ID to carry a VID.

To accommodate the extended system ID, the original 802.1D 16-bit bridge priority field is split into two fields. The BID includes the following fields:

- **Bridge priority:** A 4-bit field that is still used to carry bridge priority. The priority is conveyed in discrete values in increments of 4096 rather than discrete values in increments of 1, because only the four most significant bits are available from the 16-bit field. In other words, in binary: priority 0 = [0000|<sys-id-ext #>], priority 4096 = [0001|<sys-id-ext #>], and so on. Increments of 1 would be used if the complete 16-bit field was available. The default priority, in accordance with IEEE 802.1D, is 32,768, which is the midrange value.
- **Extended system ID:** A 12-bit field carrying, in this case, the VID for PVST+.
- **MAC address:** A 6-byte field with the MAC address of a single switch.

By virtue of the MAC address, a BID is always unique. When the priority and extended system ID are prepended to the switch MAC address, each VLAN on the switch can be represented by a unique BID.

For example, the VLAN 2 default BID would be 32770 (priority 32768 plus the extended system ID of 2).

If no priority is configured, every switch will have the same default priority and the election of the root for each VLAN is based on the MAC address. This method is a random means of selecting the ideal root bridge. For this reason, it is recommended to assign a lower priority to the switch that should serve as the root bridge.

Modifying the Bridge ID

This topic describes how you can modify the BID of a switch and, with that, the Root Bridge election.

Per VLAN Spanning Tree Plus (Cont.)

```
SW1#show spanning-tree vlan 1
VLAN0001
Spanning tree enabled protocol ieee
Root ID    Priority    28673
           Address    001e.145e.4980
           Cost      19
           Port      3 (FastEthernet0/3)
<output omitted>
```

- SW1 is not the root bridge for VLAN1. This is the switch that is connected to FastEthernet0/3 on SW1.

© 2013 Cisco Systems, Inc.

Per VLAN Spanning Tree Plus (Cont.)

Configure SW1 as the root bridge for VLAN 1.

```
SW1(config)#spanning-tree vlan 1 root primary
```

```
SW1#show spanning-tree vlan 1
VLAN0001
Spanning tree enabled protocol ieee
Root ID    Priority    24577
           Address    001e.147c.6f00
           This bridge is the root
<output omitted>
```

- After modification, SW1 is the root bridge for VLAN1.

© 2013 Cisco Systems, Inc.

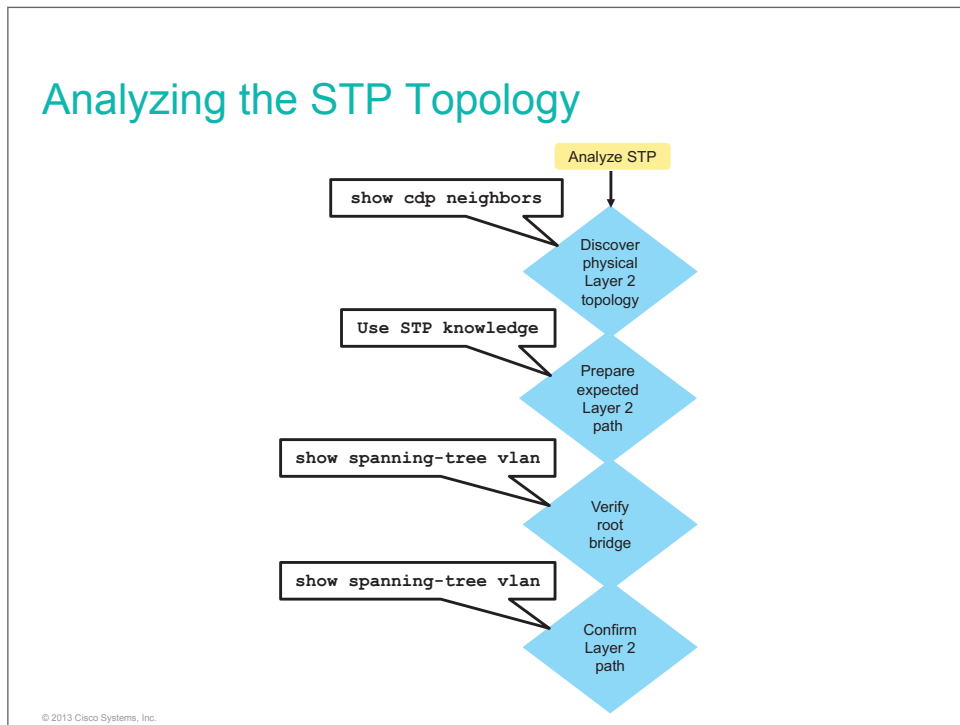
| Command | Description |
|--|---|
| <code>spanning-tree vlan vlan_number root primary</code> | Forces this switch to be the root bridge for the specified VLAN |

The root bridge is elected based on the BID. Since the priority part of the BID is, by default, the same for all switches (32768), the root bridge will be the switch with the lowest MAC address. For load balancing between switches (for example, if you want one switch to be the root bridge for VLAN 1 and the other switch to be the root bridge for VLAN 2), you can modify the priority of the bridge. The easiest way to make a switch the root bridge for a VLAN is the **spanning-tree vlan *vlan_number* root primary** command.

Do Not Duplicate.
Post beta, not for release.

Analyzing the STP Topology

This topic describes how you can analyze the spanning-tree topology and verify the proper operation of STP.




To analyze the STP topology, follow these steps:

1. Discover the physical Layer 2 topology. You could use network documentation, if it exists, or use the **show cdp neighbors** command to discover the physical topology.
2. After you have discovered the physical topology, use your knowledge of STP to determine the expected Layer 2 path. You will need to know which switch is the root bridge.
3. Use the **show spanning-tree vlan** command to determine which switch is the root bridge.
4. Use the **show spanning-tree vlan** command on all switches to find out which ports are in blocking or forwarding state, and thus confirm your expected Layer 2 path.

Analyzing the STP Topology (Cont.)

Verify that the actual STP topology matches the expected topology.



© 2013 Cisco Systems, Inc.

In many networks, the optimal STP topology is determined as part of the network design and then implemented through manipulation of STP priority and cost values. You might run into situations where STP was not considered in the design and implementation, or where it was considered initially, before the network underwent significant growth and change. In such situations, it is important to know how to analyze the actual STP topology in the operational network.

In addition, part of troubleshooting consists of comparing the actual state of the network against the expected state of the network and spotting the differences to gather clues about the problem that you are troubleshooting. You should be able to examine the switches and determine the actual topology, in addition to knowing what the spanning-tree topology is supposed to be.

Analyzing the STP Topology (Cont.)

```
SwitchA#show spanning-tree vlan 100
VLAN0100
Spanning tree enabled protocol ieee
Root ID    Priority    28772
           Address    0000.0c9f.3127
           Cost        2
           Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
Bridge ID   Priority    28772 (priority 28672 sys-id-ext 100)
           Address    0000.0cab.3724
           Hello Time 2 sec  Max Age 20 sec  Forward Delay 15 sec
           Aging Time 300 sec
Interface   Role Sts Cost      Prio.Nbr Type
-----
Gi3/1       Desg FWD 4         128.72  P2p
Gi3/2       Desg FWD 4         128.80  P2p
Te9/1       Root FWD 2         128.88  P2p
```

- The **show spanning-tree** command displays an overview of STP status and topology.

© 2013 Cisco Systems, Inc.

Using the **show spanning-tree** command without specifying any additional options is a good way to get a quick overview of the status of STP for all VLANs that are defined on a switch. If you are interested only in a particular VLAN, you can limit the scope of this command by specifying that VLAN as an option.

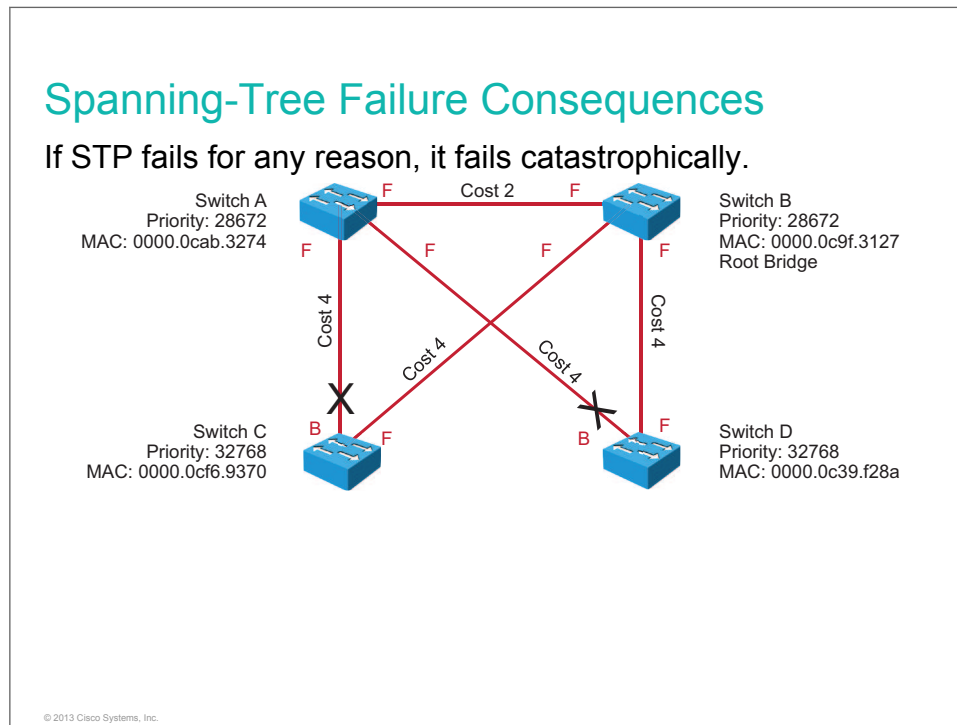
Use the **show spanning-tree vlan *vlan_id*** command to obtain STP information for a particular VLAN. Use this command to get information about the role and status of each port on the switch. The example output on Switch A (root bridge) shows all three ports in the forwarding state (FWD) and the role of the three ports as either designated ports or root ports. Any ports being blocked have status "BLK" in the output.

The output also gives information about the BID of the local switch and the root ID. If Switch A is the root bridge, the Root ID and Bridge ID MAC addresses listed would be the same.

Do Not Duplicate.
Post beta, not for release.

Spanning-Tree Failure Consequences

This topic describes common problems that can result from a redundant topology and describes how to fix them.




With many protocols, a malfunction means that you only lose the functionality that the protocol was providing. For example, if OSPF is malfunctioning on one of your routers, you might lose connectivity to networks that are reachable via that router, but it would generally not affect the rest of your OSPF network. If you still have some way to connect to that router, you can perform your troubleshooting routines to diagnose and fix the problem.

With STP, there are two types of failure. The first is similar to the OSPF problem: STP might erroneously decide to block ports that should have gone into the forwarding state. You might lose connectivity for traffic that would normally pass through this switch, but the rest of the network is unaffected, and you can troubleshoot the switch as long as you still have a way to access it. The second type of failure is much more disruptive. It happens when STP erroneously decides to move one or more ports into the forwarding state.

Spanning-Tree Failure Consequences (Cont.)

What will happen to this network if Switch D erroneously transitions both its ports to the forwarding state?

- Any frame that enters a bridging loop will continue to be forwarded by the switches indefinitely.




Frames that have a destination address recorded in the MAC address table of the switches are simply forwarded to the port that is associated with the MAC address and do not enter a loop. However, any frame that is flooded by a switch—broadcasts, multicasts, and unicasts with an unknown destination MAC address—enters a loop.

Spanning-Tree Failure Consequences (Cont.)

The consequences of STP failure are severe.

- The load on all links in the switched LAN quickly starts increasing.
- Due to the very high load for the CPU, the switch becomes unreachable.



What are the consequences and corresponding symptoms of STP failure?

- The load on all links in the switched LAN quickly starts increasing as more frames enter the loop. This problem is not limited to the links that form the loop but also affects any other links in the switched domain, because the frames are flooded on all links. When the spanning-tree failure is limited to a single VLAN, only links in that VLAN are affected. Switches and trunks that do not carry that VLAN operate normally.
- If the spanning-tree failure has caused more than one bridging loop to exist, traffic increases exponentially, because frames not only start circling but also start getting duplicated. This problem happens because when there are multiple loops, there will be switches that receive a frame on a port and then flood it out on multiple ports, essentially creating a copy of the frame every time they forward it.
- The switches will experience frequent MAC address table changes. This problem happens because frames usually start looping in both directions, causing a switch to see a frame with a certain source MAC address coming in on a port and then see a frame with the same source MAC address coming in on a different port just a fraction of a second later.
- Because of the combination of very high load on all links and the switch CPUs running at maximum load, these devices typically become unreachable, making it nearly impossible to diagnose the problem while it is happening.

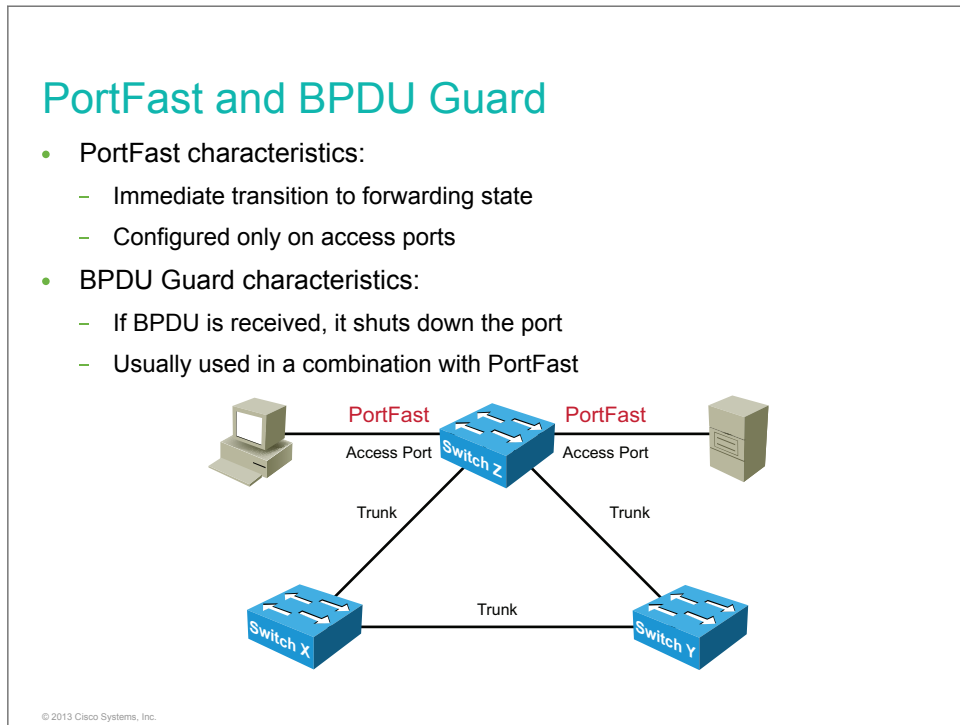
A viable approach is to take over the role of the failing spanning tree by manually removing redundant links in the switched network, either physically or through configuration (if that is still possible), until all loops are eliminated from the topology. When the loops are broken, the traffic and CPU loads should quickly drop to normal levels, and you should regain connectivity to your devices.

Although this intervention restores connectivity to the network, you cannot consider it the end of your troubleshooting process. You have removed all redundancy from your switched network, and you need to restore the redundant links.

Of course, if the underlying cause of the spanning-tree failure has not been fixed, chances are that restoring the redundant links will trigger a new broadcast storm. Before you restore the redundant links, you should spend sufficient time to investigate what happened at the moment that the broadcast storm started. When you eventually start restoring the redundant links, you should carefully monitor the network and have an emergency plan to fall back on if you see a new broadcast storm developing.

PortFast and BPDU Guard

This topic explains why PortFast and BPDU guard are important technologies and how to enable them.



PortFast is a Cisco technology. When a switch port that is configured with PortFast is configured as an access port, that port transitions from the blocking to the forwarding state immediately, bypassing the typical STP listening and learning states. You can use PortFast on access ports that are connected to a single workstation or to a server to allow those devices to connect to the network immediately rather than waiting for spanning tree to converge.

In a valid PortFast configuration, BPDUs should never be received, because that would indicate that another bridge or switch is connected to the port, potentially causing a spanning-tree loop. Cisco switches support a feature called BPDU guard. When it is enabled, BPDU guard puts the port in an error-disabled state (effectively shut down) on receipt of a BPDU.

Note Because the purpose of PortFast is to minimize the time that access ports connecting to user equipment and servers must wait for spanning tree to converge, it should be used only on access ports. If you enable PortFast on a port connecting to another switch, you risk creating a spanning-tree loop. The only exception is when a trunk connects to a non-switch device, typically a server.

PortFast and BPDU Guard (Cont.)

```
SwitchX(config)#interface FastEthernet0/1
SwitchX(config-if)#spanning-tree portfast
SwitchX(config-if)#spanning-tree bpduguard enable
```

- Configures BPDU guard and PortFast on interface FastEthernet0/1

```
SwitchX(config)#spanning-tree portfast bpduguard default
SwitchX(config)#spanning-tree portfast default
```

- Enables PortFast on all nontrunking interfaces and enables BPDU guard globally for all PortFast-enabled ports

© 2013 Cisco Systems, Inc.

The table lists the commands that are used to implement PortFast and BPDU guard.

PortFast and BPDU Guard Commands

| Command | Description |
|---|---|
| spanning-tree portfast | Enables PortFast on a Layer 2 access port and forces it to enter the forwarding state immediately |
| spanning-tree portfast default | Globally enables the PortFast feature on all nontrunking ports. When the PortFast feature is enabled, the port changes from a blocking state to a forwarding state without making the intermediate spanning-tree state changes. |
| spanning-tree bpduguard enable | Enables BPDU guard on a Layer 2 access port |
| spanning-tree portfast bpduguard default | Globally enables the BPDU guard feature |

PortFast and BPDU Guard (Cont.)

```
SwitchX#show running-config interface FastEthernet0/1
Building configuration...
Current configuration : 57 bytes
!
interface FastEthernet0/1
  spanning-tree portfast
  spanning-tree bpduguard enable
end
```

- Verifies that PortFast and BPDU guard have been configured on interface FastEthernet0/1

```
SwitchX#show spanning-tree interface FastEthernet 0/1 portfast
VLAN0010          enabled
```

- Verifies that PortFast is enabled on FastEthernet0/1

© 2013 Cisco Systems, Inc.

The **show running-config interface FastEthernet0/1** command shows the configuration on the FastEthernet 0/1 interface. Using **show spanning-tree interface FastEthernet 0/1 portfast**, you can verify that FastEthernet 0/1 has PortFast enabled.

Summary

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

Summary

- A redundant switch topology causes broadcast storms, multiple frame copies, and MAC address table instability problems.
- STP allows physical path redundancy while preventing the undesirable effects of active loops in the network.
- A root bridge is elected based on the lowest BID.
- There are many STP standards. PVST+ is a Cisco enhancement of STP that provides a separate 802.1D spanning-tree instance for each VLAN configured in the network.
- PVST+ requires that a separate instance of spanning tree is run for each VLAN, and the BID field must carry VID information. The BID includes the bridge priority, extended system ID, and MAC address.

© 2013 Cisco Systems, Inc.

Summary (Cont.)

- PortFast is used on ports connected to a single workstation or server to allow those devices to connect to the network immediately.
- If you enable PortFast on a port connecting to another switch, you risk creating a spanning-tree loop. The BPDU guard feature prevents spanning-tree loops in such cases.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Improving Redundant Switched Topologies with EtherChannel

In hierarchical network design, some links between access and distribution switches may be heavily utilized. The speed of these links can be increased, but only to a certain point. EtherChannel is a technology that allows you to circumvent this restriction by creating logical links made up of several physical links. This lesson describes EtherChannel technology and the various technologies that are available to implement it. You will also learn how to configure EtherChannel and how to verify EtherChannel operations.

Objectives

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


- Describe the idea behind EtherChannel technology
- Describe the advantages of EtherChannel technology
- Identify the two EtherChannel protocols and their modes
- Configure link aggregation using EtherChannel
- Explain what can go wrong with EtherChannel configurations

The Need for EtherChannel

This topic describes the need for EtherChannel technology.

The Need for EtherChannel

- When multiple links aggregate on a switch, congestion occurs.
- One solution is to increase uplink speed, but that solution cannot scale indefinitely.
- Another solution is to multiply uplinks, but loop-prevention mechanisms disable some ports.



© 2013 Cisco Systems, Inc.

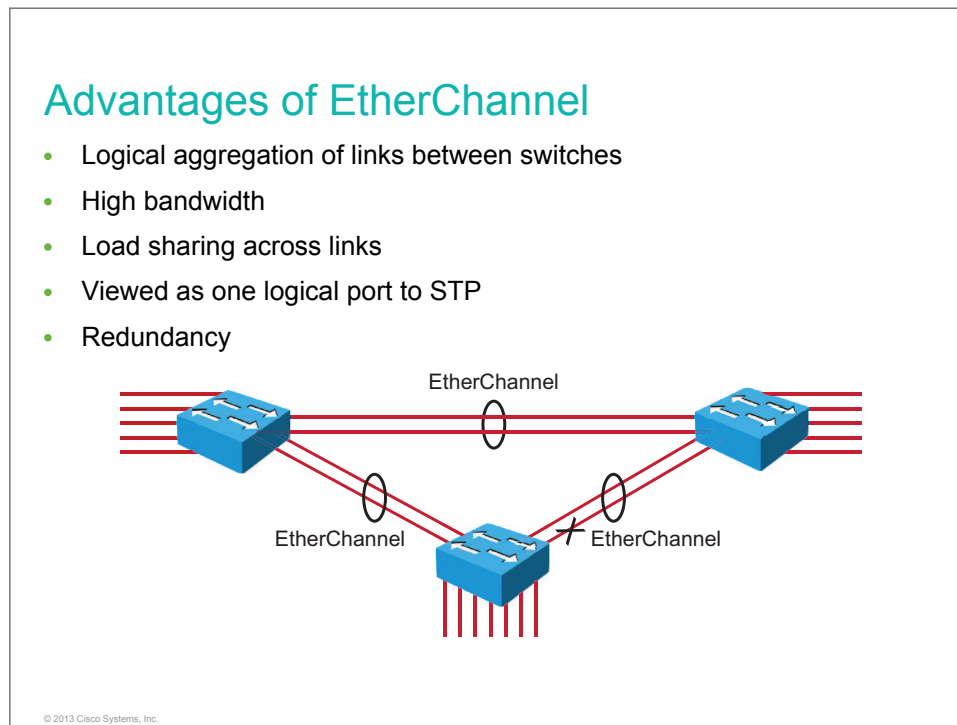
In the figure, traffic coming from several links (usually 100 or 1000 Mb/s) aggregates on the access switch and needs to be sent to distribution switches, at the top of the figure. Because of the aggregation, links with higher bandwidth have to be available between the access and distribution switches.

One solution is to use faster links, such as 10 Gb/s. As the speed increases on the access links, this solution finds its limitation where the fastest possible port is no longer fast enough to aggregate the traffic coming from all access links. The second issue is that faster links are expensive.

Another solution is to multiply the number of physical links between the switches to increase the overall speed of switch-to-switch communication. A downside of this method is that there must be strict consistency in the configuration of each physical link. A second issue is that STP will block one of the links.

Advantages of EtherChannel

This topic describes the advantages for EtherChannel technology.



EtherChannel is a technology that was originally developed by Cisco as a LAN switch-to-switch technique for grouping several Fast Ethernet or Gigabit Ethernet ports into one logical channel.

This technology has many benefits:

- Most configuration tasks can be done on the EtherChannel interface instead of on each individual port, ensuring configuration consistency throughout the links.
- It relies on the existing switch ports. There is no need to upgrade the link to a faster and more expensive connection to have more bandwidth.
- Load balancing is possible between links that are part of the same EtherChannel. Depending on the hardware platform, you can implement one or several methods, such as source MAC-to-destination MAC or source IP-to-destination IP load balancing, across the physical links.
- EtherChannel creates an aggregation that is seen as one logical link. When several EtherChannel bundles exist between two switches, STP may block one of the bundles to prevent redundant links. When STP blocks one of the redundant links, it blocks one EtherChannel, thus blocking all the ports belonging to that EtherChannel link. Where there is only one EtherChannel link, all physical links in the EtherChannel are active because STP sees only one (logical) link.
- EtherChannel provides redundancy, as the overall link is seen as one logical connection, and the loss of one physical link does not create a change in the topology. A spanning-tree recalculation does not need to take place. As long as at least one physical link is present, the EtherChannel is functional, even if its overall throughput decreases.

EtherChannel can be implemented by grouping from two to eight physical links into a logical EtherChannel link. You cannot mix interface types (for example, Fast Ethernet and Gigabit Ethernet) within a single EtherChannel. Keep in mind that the point of EtherChannel is to increase the speed between switches. This concept was extended as the EtherChannel technology became more popular, and some hardware devices other than switches support link aggregation into an EtherChannel link. In any case, EtherChannel creates a one-to-one relationship. You can create an EtherChannel link between two switches or between an EtherChannel-enabled server and a switch, but you cannot send traffic to two different switches through the same EtherChannel link. One EtherChannel link always connects two devices only. The individual EtherChannel group member port configuration must be consistent on both devices. If the physical ports of one side are configured as trunks, the physical ports of the other side must also be configured as trunks. Each EtherChannel has a logical port channel interface. A configuration applied to the port channel interface affects all physical interfaces that are assigned to that interface.

EtherChannel is the Cisco version of port channel technology.

Do Not Duplicate.
Post beta, not for release.

EtherChannel Protocols

This topic describes two protocols for link aggregation. These protocols allow ports with similar characteristics to form a channel through dynamic negotiation with adjoining switches.

EtherChannel Protocols

- Two protocols exist to negotiate EtherChannel creation and maintenance:
 - PAgP is a Cisco proprietary protocol.
 - LACP is an IEEE 802.3ad standard.
- Static EtherChannel can be configured without PAgP or LACP.

© 2013 Cisco Systems, Inc.

PAgP is a Cisco proprietary protocol that aids in the automatic creation of EtherChannel links. When an EtherChannel link is configured using PAgP, PAgP packets are sent between EtherChannel-capable ports to negotiate the forming of a channel. When PAgP identifies matched Ethernet links, it groups the links into an EtherChannel. The EtherChannel is then added to the spanning tree as a single bridge port.


When enabled, PAgP also manages the EtherChannel. PAgP packets are sent every 30 seconds. PAgP checks for configuration consistency and manages link additions and failures between two switches. It ensures that when an EtherChannel is created, all ports have the same type of configuration. In EtherChannel, it is mandatory that all ports have the same speed, duplex setting, and VLAN information. Any port modification after the creation of the channel will also change all the other channel ports.

LACP is part of an IEEE specification (802.3ad) that allows several physical ports to be bundled to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a function similar to PAgP with Cisco EtherChannel. Because LACP is an IEEE standard, you can use it to facilitate EtherChannels in multivendor environments. On Cisco devices, both protocols are supported.

EtherChannel Protocols (Cont.)

PAgP negotiates EtherChannel formation and maintenance.

- **On:** Channel member without negotiation (no protocol).
- PAgP modes:
 - **Desirable:** Actively asking if the other side can or will participate
 - **Auto:** Passively waiting for the other side



| Channel establishment | On | Desirable | Auto |
|-----------------------|-----|-----------|------|
| On | YES | NO | NO |
| Desirable | NO | YES | YES |
| Auto | NO | YES | NO |

© 2013 Cisco Systems, Inc.

PAgP helps create the EtherChannel link by detecting the configuration of each side and making sure that they are compatible so that the EtherChannel link can be enabled when needed. The table shows the settings for PAgP.

| Mode | Purpose |
|----------------|--|
| PAgP auto | This PAgP mode places an interface in a passive negotiating state in which the interface responds to the PAgP packets that it receives but does not initiate PAgP negotiation. |
| PAgP desirable | This PAgP mode places an interface in an active negotiating state in which the interface initiates negotiations with other interfaces by sending PAgP packets. |
| On | This mode forces the interface to channel without PAgP. Interfaces configured in the on mode do not exchange PAgP packets. |


The modes must be compatible on each side. If you configure one side to be in auto mode, it will be placed in a passive state, waiting for the other side to initiate the EtherChannel negotiation. If the other side is also set to auto, the negotiation never starts and the EtherChannel does not form. If you disable all modes by using the **no** command or if no mode is configured, then the interface is placed in the off mode and EtherChannel is disabled.

Note that the on mode manually places the interface in an EtherChannel, without any negotiation. It works only if the other side is also set to on. If the other side is set to negotiate parameters through PAgP, no EtherChannel will form, because the side that is set to on mode will not negotiate.

EtherChannel Protocols (Cont.)

LACP negotiates EtherChannel formation and maintenance.

- **On:** Channel member without negotiation (no protocol).
- LACP modes:
 - **Active:** Actively asking if the other side can or will participate
 - **Passive:** Passively waiting for the other side



| Channel establishment | On | Active | Passive |
|-----------------------|-----|--------|---------|
| On | YES | NO | NO |
| Active | NO | YES | YES |
| Passive | NO | YES | NO |

© 2013 Cisco Systems, Inc.

LACP provides the same negotiation benefits as PAgP. LACP helps create the EtherChannel link by detecting the configuration of each side and making sure that they are compatible, so that the EtherChannel link can be enabled when needed. The table shows the settings for LACP.

| Mode | Purpose |
|--------------|---|
| LACP passive | This LACP mode places a port in a passive negotiating state. In this state, the port responds to the LACP packets that it receives but does not initiate LACP packet negotiation. |
| LACP active | This LACP mode places a port in an active negotiating state. In this state, the port initiates negotiations with other ports by sending LACP packets. |
| On | This mode forces the interface to channel without LACP. Interfaces configured in the on mode do not exchange LACP packets. |

Like PAgP, modes must be compatible on both sides for the EtherChannel link to form. The on mode is mentioned here again because it creates the EtherChannel configuration unconditionally, without PAgP or LACP dynamic negotiation.

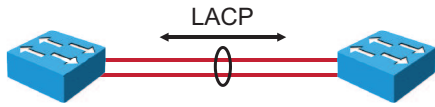
Configuring EtherChannel

This topic describes how to configure EtherChannel.

Configuring EtherChannel

All interfaces within an EtherChannel must have the same configuration:

- Speed and duplex
- Mode (access or trunk)
- Native and allowed VLANs on trunk ports
- Access VLAN on access ports




© 2013 Cisco Systems, Inc.

Follow these guidelines and restrictions when configuring EtherChannel interfaces:

- **EtherChannel support:** All Ethernet interfaces on all modules support EtherChannel (maximum of eight interfaces), with no requirement that interfaces be physically contiguous or on the same module.
- **Speed and duplex:** Configure all interfaces in an EtherChannel to operate at the same speed and in the same duplex mode.
- **VLAN match:** All interfaces in the EtherChannel bundle must be assigned to the same VLAN or be configured as a trunk.
- **Range of VLAN:** An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel.

If you have to change these settings, configure them in EtherChannel interface configuration mode. After you configure the EtherChannel interface, any configuration that you apply to the port channel interface affects individual interfaces as well. The opposite does not apply and will cause interface incompatibility in the EtherChannel.

Configuring EtherChannel (Cont.)



```
SW1(config)#interface range FastEthernet0/1 - 2
SW1(config-if-range)#channel-group 1 mode active
SW1(config-if-range)#exit
SW1(config)#interface port-channel 1
SW1(config-if)#switchport mode trunk
SW1(config-if)#switchport trunk allowed vlan 1,2,20
```

- Creates EtherChannel and configures trunk on SW1

```
SW1(config)#interface range FastEthernet0/4 - 5
SW1(config-if-range)#channel-group 1 mode active
SW1(config-if-range)#exit
SW1(config)#interface port-channel 1
SW1(config-if)#switchport mode trunk
SW1(config-if)#switchport trunk allowed vlan 1,2,20
```

- Creates EtherChannel and configures trunk on SW2

© 2013 Cisco Systems, Inc.

The configuration of an EtherChannel is based on two steps, as described in the table.

| Command | Description |
|--|---|
| interface range <i>interface</i> | Specifies the interfaces that will compose the EtherChannel group. The range keyword allows you to select several interfaces and configure them all together. A good practice is to start by shutting down those interfaces, so that incomplete configuration will not start to create activity on the link. |
| channel-group <i>identifier mode active</i> | Creates the port channel interface, if necessary, and assigns the specified interfaces to it. The identifier specifies a channel group number. |

In the example, FastEthernet0/1 and FastEthernet0/2 are bundled into EtherChannel interface port channel 1. To change Layer 2 settings on the EtherChannel interface, enter EtherChannel interface configuration mode using the **interface port-channel** command, followed by the interface identifier. In the example, the EtherChannel is configured as a trunk interface with allowed VLANs as specified.

Verifying EtherChannel

This topic describes how to verify that EtherChannel is functioning properly.

Verifying EtherChannel

```
SW1#show interface Port-channel1
Port-channel1 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 000f.34f9.9182 (bia 000f.34f9.9182)
  MTU 1500 bytes, BW 200000 Kbit, DLY 100 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
<output omitted>
```

- Verifies interface status

© 2013 Cisco Systems, Inc.

You can use one of several commands to verify an EtherChannel configuration. You can first use the **show interface port-channel** command to display the general status of the EtherChannel interface. In the example, the Port-channel 1 interface is up.

Verifying EtherChannel (Cont.)

```
SW2#show etherchannel summary
Flags: D - down          P - bundled in port-channel
       I - stand-alone  s - suspended
       H - Hot-standby (LACP only)
       R - Layer3       S - Layer2
       U - in use       f - failed to allocate aggregator
       M - not in use, minimum links not met
       u - unsuitable for bundling
       w - waiting to be aggregated
       d - default port
Number of channel-groups in use: 1
Number of aggregators:          1
Group  Port-channel  Protocol    Ports
-----+-----+-----+-----
1      Po1(SU)         LACP       Fa0/1(P)  Fa0/2(P)
```

- Displays a one-line summary per channel group

© 2013 Cisco Systems, Inc.

When several port channel interfaces are configured on the same device, you can use the **show etherchannel summary** command to simply display one line of information per port channel. In this example, the switch has one EtherChannel configured; group 1 uses LACP. The interface bundle consists of the FastEthernet0/1 and FastEthernet0/2 interfaces. You can see that the group is Layer 2 EtherChannel and that it is in use (shown by the letters SU next to the port channel number).

Verifying EtherChannel (Cont.)

```
Switch#show etherchannel Port-channel
Channel-group listing:
-----
Group: 1
-----
Port-channels in the group:
-----
Port-channel: Po1 (Primary Aggregator)
-----
Age of the Port-channel = 4d:01h:29m:00s
<output omitted>
Protocol = LACP
<output omitted>
Ports in the Port-channel:
Index Load Port EC state No of bits
-----+-----+-----+-----+-----
0 00 Fa0/1 Active 4
1 00 Fa0/2 Active 4
Time since last port bundled: 0d:00h:00m:18s Fa0/2
Time since last port Un-bundled: 0d:00h:00m:32s Fa0/2
```

- Displays port channel information

© 2013 Cisco Systems, Inc.

Use the **show etherchannel port-channel** command to display information about the specific port channel interface. In the example, the Port-channel 1 interface consists of two physical interfaces, FastEthernet0/1 and FastEthernet0/2. It uses LACP in active mode. It is properly connected to another switch with a compatible configuration. This is why the port channel is said to be in use.

Note *Load* does not actually indicate the load over an interface. It is meant to be a hexadecimal value that decodes which interface will be chosen for a specific flow of traffic.

Summary

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

Summary

- EtherChannel is a technology that is used to group several ports into one logical channel.
- PAgP and LACP are two protocols for link aggregation. They allow ports with similar characteristics to form a channel through dynamic negotiation with adjoining switches.
- All interfaces within an EtherChannel must have the same configuration of speed and duplex mode, native and allowed VLANs on trunks, and access VLAN on access ports.
- Use the **show etherchannel summary** command to quickly identify EtherChannel groups on the switch.

© 2013 Cisco Systems, Inc.

Understanding Layer 3 Redundancy

This lesson explains solutions to routing problems in a local network with a redundant topology. One solution to these problems is explained through the router redundancy process. This lesson identifies HSRP, VRRP, and GLBP as Layer 3 redundancy protocols. It also describes how to use the **show standby** and **show glbp** commands.

Objectives

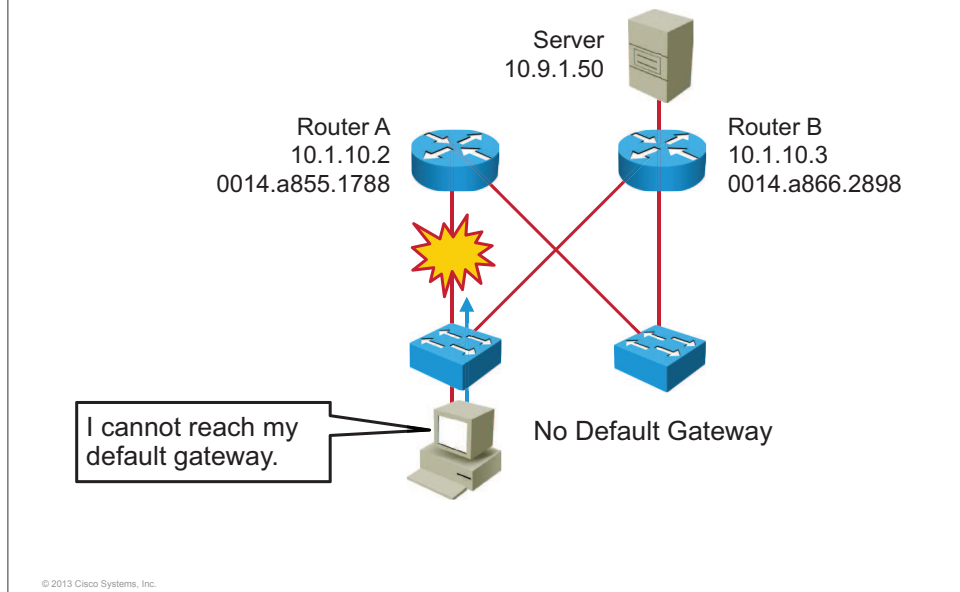
Upon completing this lesson, you will be able to meet these objectives:

- Describe routing issues in connection to redundancy
- Explain the router redundancy process and what happens when a failover occurs
- Identify HSRP and VRRP as Layer 3 redundancy protocols
- Describe the idea behind HSRP interface tracking
- Describe the idea behind HSRP load balancing
- Identify GLBP as a redundancy protocol

The Need for Default Gateway Redundancy

This topic describes routing issues that occur when you use default gateways in a redundant network.

The Need for Default Gateway Redundancy



Each client receives only one default gateway. There is no means by which to configure a secondary gateway, even if a second route exists to carry packets off the local segment.

For example, primary and secondary paths between the access layer equipment and the distribution layer switches provide continuous access in the event of a link failure between those two layers. Primary and secondary paths between the distribution layer switches and the core layer switches provide continuous operation if there is a failure between those two layers.

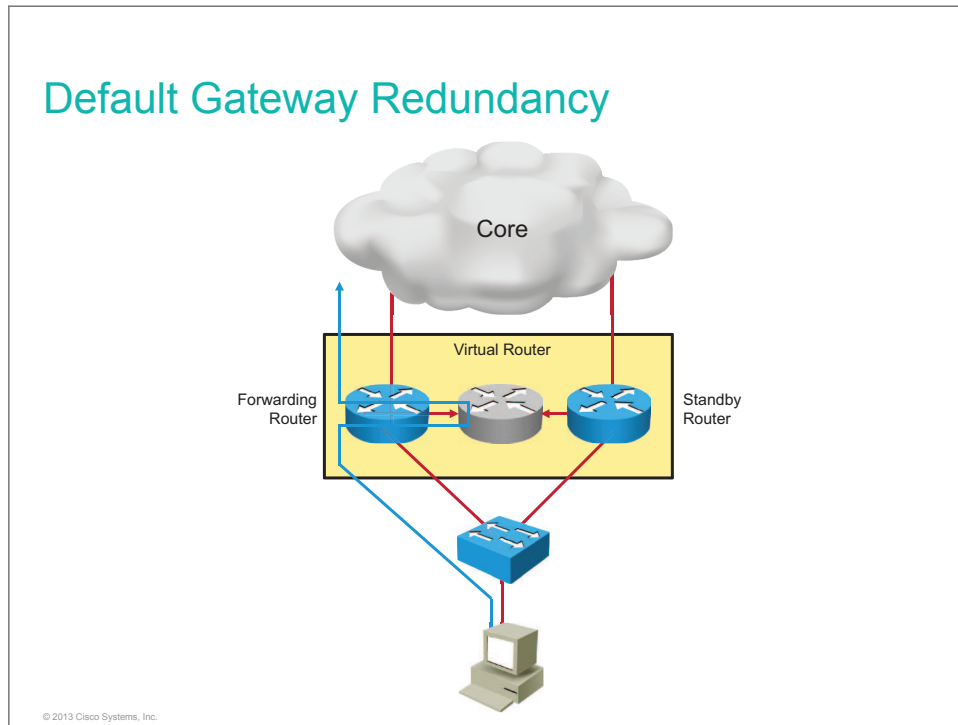
In this example, Router A is responsible for routing packets for Subnet A, and Router B is responsible for routing packets for Subnet B. If Router A becomes unavailable, the routing protocols can quickly and dynamically converge and determine that Router B will now transfer packets that would otherwise have gone through Router A. Most workstations, servers, and printers, however, do not receive this dynamic routing information.

End devices are typically configured with a single default gateway IP address that does not change when the network topology changes. If the router whose IP address is configured as the default gateway fails, the local device is unable to send packets off the local network segment, effectively disconnecting it from the rest of the network. Even if a redundant router exists that could serve as a default gateway for that segment, there is no dynamic method by which these devices can determine the address of a new default gateway.

Even though the example is explained on routers, in modern networks, these routers would actually be Layer 3 switches. These are high-performance devices for routing, but, in contrast to routers, they have many interfaces.

Default Gateway Redundancy

This topic describes how router redundancy works.



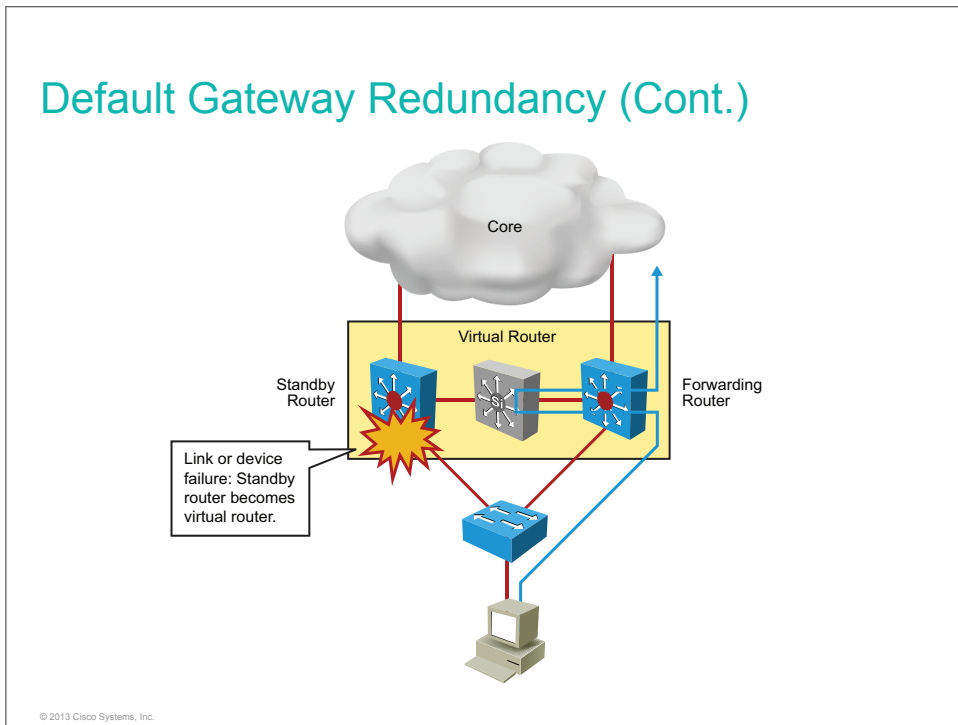
With the type of router redundancy that is shown in the figure, a set of routers works together to present the illusion of a single router to the hosts on the LAN. By sharing an IP (Layer 3) address and a MAC (Layer 2) address, two or more routers can act as a single “virtual” router.

The IP address of the virtual router is configured as the default gateway for the workstations on a specific IP segment. When frames are sent from the workstation to the default gateway, the workstation will use ARP to resolve the MAC address that is associated with the IP address of the default gateway. The ARP resolution returns the MAC address of the virtual router. Frames that are sent to the MAC address of the virtual router can then be physically processed by any active or standby router that is part of that virtual router group.

A protocol is used to identify two or more routers as the devices that are responsible for processing frames that are sent to the MAC or IP address of a single virtual router. Host devices send traffic to the address of the virtual router. The physical router that forwards this traffic is transparent to the end stations.

The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router. The transition from one forwarding router to another is transparent to the end devices.

Default Gateway Redundancy (Cont.)



These are the steps that take place when a router fails:


- 1 The standby router stops seeing hello messages from the forwarding router.
- 2 The standby router assumes the role of the forwarding router.
- 3 Because the new forwarding router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service.

HSRP

This topic describes HSRP.

HSRP

- HSRP defines a group of routers—one active and one standby.
- Virtual IP and MAC addresses are shared between the two routers.
- To verify HSRP state, use the **show standby** command.
- HSRP is Cisco proprietary, and VRRP is a standard protocol.



The diagram illustrates HSRP Group 1. It features three router icons on a yellow background. On the left is a blue router labeled 'Active'. In the center is a grey router labeled 'Virtual'. On the right is a blue router labeled 'Standby'. The entire group is titled 'HSRP Group 1' at the top.

© 2013 Cisco Systems, Inc.

HSRP defines a standby group of routers, with one router designated as the active router. HSRP provides gateway redundancy by sharing IP and MAC addresses between redundant gateways. The protocol consists of virtual MAC and IP addresses that are shared between two routers that belong to the same HSRP group.

HSRP Terminology

| Term | Definition |
|----------------|--|
| Active router | The router that is currently forwarding packets for the virtual router |
| Standby router | The primary backup router |
| Standby group | The set of routers participating in HSRP that jointly emulate a virtual router |

The function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router becomes inoperable.

HSRP is a Cisco proprietary protocol, and VRRP is a standard protocol. Beyond that, the differences between HSRP and VRRP are very slight.

HSRP (Cont.)

Use the **show standby** command to verify the HSRP state.

```
R1#show standby
Vlan1 - Group 1
  State is Active
    2 state changes, last state change 00:00:10
  Virtual IP address is 10.1.1.100
  Active virtual MAC address is 0000.0c07.ac01
  Local virtual MAC address is 0000.0c07.ac01 (vl default)
  Hello time 3 sec, hold time 10 sec
  Next hello sent in 2.800 secs
  Preemption disabled
  Active router is local
  Standby router is unknown
  Priority 100 (default 100)
  Group name is "hsrp-Vl1-1" (default)
```


© 2013 Cisco Systems, Inc.

To display HSRP information, use the **show standby** command in privileged EXEC mode.

The example output shows the HSRP state on the R1 router. The IP of the virtual router is 10.1.1.100 and R1 is actively routing traffic.

HSRP (Cont.)

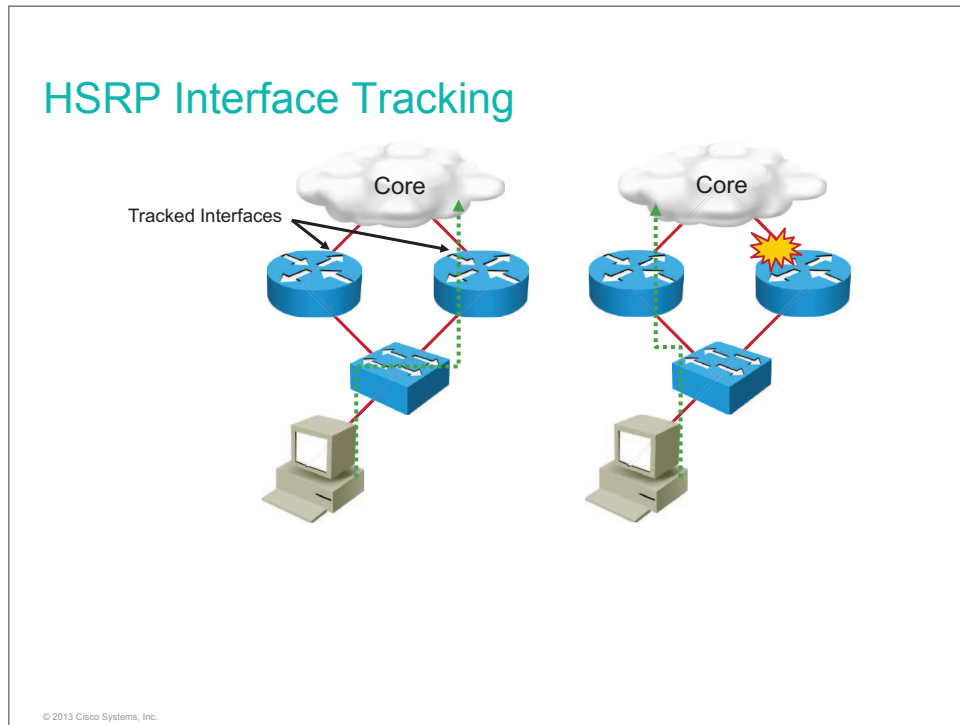
- Active router:
 - Responds to default gateway ARP requests with the virtual router MAC address
 - Assumes active forwarding of packets for the virtual router
 - Sends hello messages
 - Knows the virtual router IP address
- Standby Router
 - Listens for periodic hello messages
 - Assumes active forwarding of packets if it does not hear from active router



© 2013 Cisco Systems, Inc.

HSRP Interface Tracking

This topic describes how interface tracking in HSRP can improve network performance.



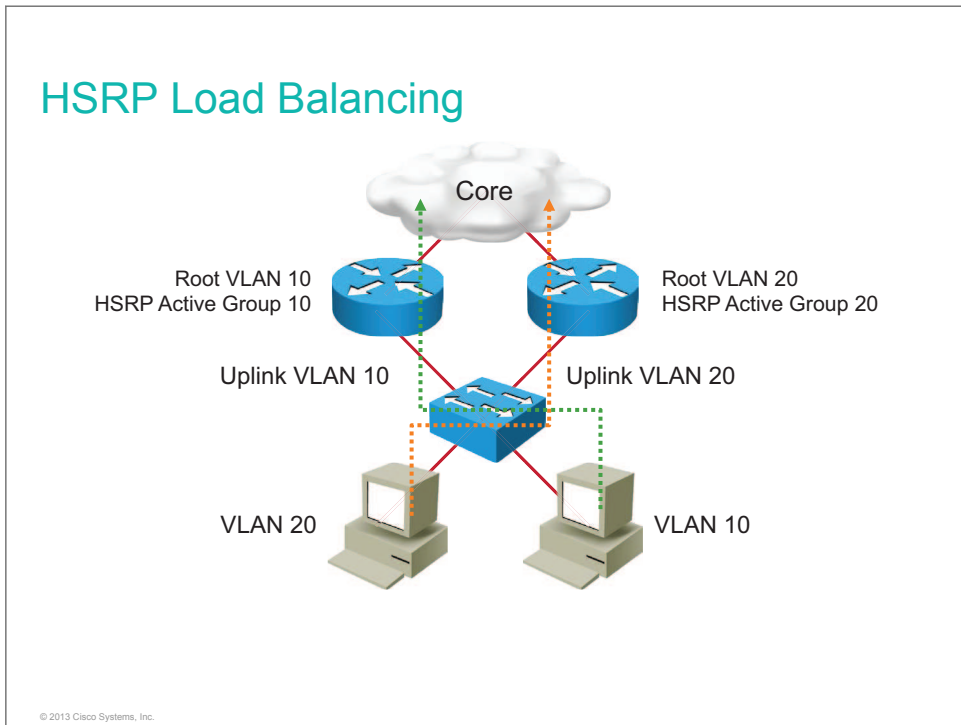
Interface tracking enables the priority of a standby group router to be automatically adjusted based on the availability of the router interfaces. When a tracked interface becomes unavailable, the HSRP tracking feature ensures that a router with an unavailable key interface will relinquish the active router role.

The HSRP group tracks the uplink interfaces. If the uplink on the right switch fails, the router automatically decrements the priority on that interface and sends hello messages with the decremented priority.

Assume that, in the example in the figure, the router on the right is configured with a higher priority, and therefore is handling the traffic toward the core. As soon as the interface through the router on the right fails, the host will be unable to reach the core network. HSRP will make the router on the left the active router.

HSRP Load Balancing

This topic describes how utilizing load balancing in HSRP can improve network performance.



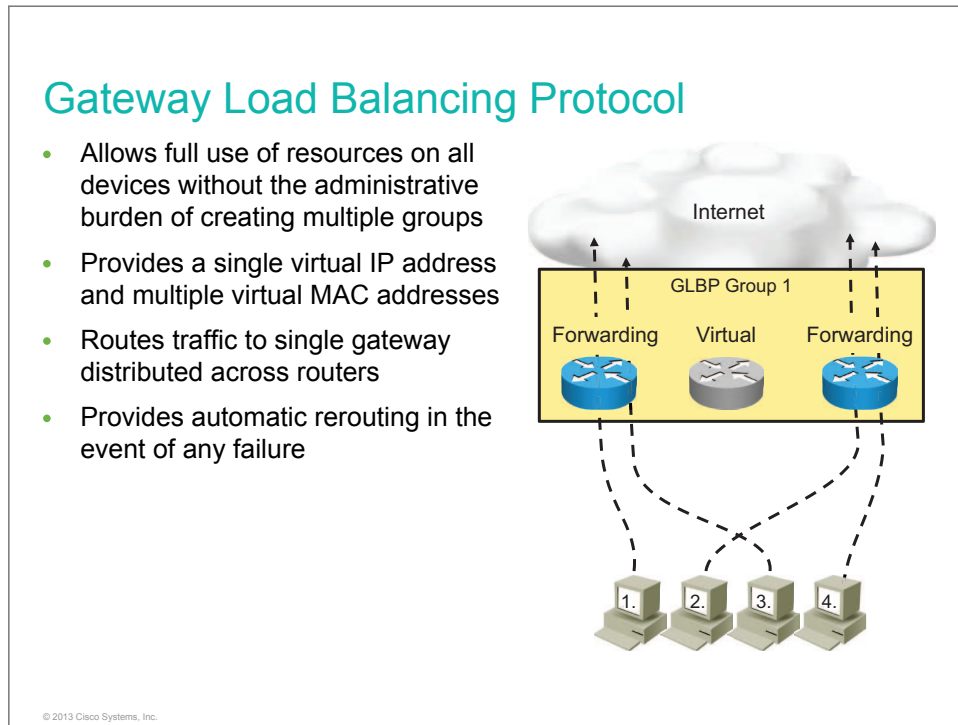
Routers can simultaneously provide redundant backup and perform load sharing across various subnets.

In the figure, two HSRP-enabled routers participate in two separate VLANs. Running HSRP over trunks allows users to configure redundancy among multiple routers.

By configuring HSRP over trunks, you can eliminate situations in which a single point of failure causes traffic interruptions. This feature inherently provides some improvement in overall networking resilience by providing load-balancing and redundancy capabilities between subnets and VLANs.

Gateway Load Balancing Protocol

This topic describes GLBP as a redundancy protocol.



Although HSRP and VRRP provide gateway resiliency, for the standby members of the redundancy group, the upstream bandwidth is not used while the device is in standby mode.

Only the active router in HSRP and VRRP groups forwards traffic for the virtual MAC address. Resources that are associated with the standby router are not fully utilized. You can accomplish some load balancing with these protocols by creating multiple groups and assigning multiple default gateways, but this configuration creates an administrative burden.

GLBP is a Cisco proprietary solution to allow automatic selection and simultaneous use of multiple available gateways, in addition to automatic failover between those gateways. Multiple routers share the load of frames that, from a client perspective, are sent to a single default gateway address.

With GLBP, you can fully utilize resources without the administrative burden of configuring multiple groups and managing multiple default gateway configurations.

Gateway Load Balancing Protocol (Cont.)

```
R1#show glbp
FastEthernet0/1 - Group 1
  State is Active
    1 state change, last state change 00:02:34
  Virtual IP address is 192.168.2.100
  <output omitted>
  Active is local
  Standby is 192.168.2.2, priority 100 (expires in 8.640 sec)
  Priority 100 (default)
  Weighting 100 (default 100), thresholds: lower 1, upper 100
  Load balancing: round-robin
  Group members:
    001e.7aa3.5e71 (192.168.2.1) local
    001e.7aa3.5f31 (192.168.2.2)
  <output omitted>
```

- The **show glbp** command in this example displays information about the status of GLBP group 1.

© 2013 Cisco Systems, Inc.

Gateway Load Balancing Protocol (Cont.)

```
R1#show glbp
<output omitted>
There are 2 forwarders (1 active)
  Forwarder 1
    State is Active
      1 state change, last state change 00:02:23
      MAC address is 0007.b400.0101 (default)
      Owner ID is 001e.7aa3.5e71
      Redirection enabled
      Preemption enabled, min delay 30 sec
      Active is local, weighting 100
  Forwarder 2
    State is Listen
  <output omitted>
```

- The **show glbp** command in this example displays information about the status of GLBP group 1.

© 2013 Cisco Systems, Inc.

To display GLBP information, use the **show glbp** command in privileged EXEC mode.

The example output shows that the virtual router IP address is 192.168.2.100 and that one router is in the Active state and the other router is in the Listen state. "Active" indicates that this router is responsible for responding to ARP requests for the virtual IP address. "Listen" indicates that the router is receiving hello packets and is ready to be activated if the currently active router fails.

Summary

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

Summary

- End devices are typically configured with a single default gateway IP address that does not change when the network topology changes.
- Redundancy protocols provide a mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router.
- HSRP defines a standby group of routers, with one router as the active router. VRRP is standard protocol that provides a similar function.
- GLBP is a Cisco proprietary solution to allow automatic selection and simultaneous use of multiple available gateways in addition to automatic failover between those gateways.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Module Summary

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

Module Summary

- A VLAN is a logical broadcast domain that can span multiple physical LAN segments.
- A loop-avoidance mechanism is required in redundant switch topologies.
- EtherChannel groups several Fast Ethernet or Gigabit Ethernet ports into one logical channel.
- With router redundancy, a set of routers works together to present the illusion of a single virtual router to the hosts on the LAN.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Questions

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

1. Which feature is required for multiple VLANs to span multiple switches? (Source: Troubleshooting VLAN Connectivity)
 - A. a trunk to connect the switches
 - B. a router to connect the switches
 - C. a bridge to connect the switches
 - D. a VLAN configured between the switches
2. What are two reasons for using 802.1Q? (Choose two.) (Source: Troubleshooting VLAN Connectivity)
 - A. to allow switches to share a trunk link with nontrunking clients
 - B. to allow clients to see the 802.1Q header
 - C. to provide inter-VLAN communication over a bridge
 - D. to load-balance traffic between parallel links using STP
 - E. to provide trunking between Cisco switches and other vendor switches
3. Which term commonly describes the endless flooding or looping of frames? (Source: Building Redundant Switched Topologies)
 - A. flood storm
 - B. loop overload
 - C. broadcast storm
 - D. broadcast overload

4. How does STP provide a loop-free network? (Source: Building Redundant Switched Topologies)
- A. by placing all ports in the blocking state
 - B. by placing all bridges in the blocking state
 - C. by placing some ports in the blocking state
 - D. by placing some bridges in the blocking state
5. Which port marks the lowest-cost path from the non-root bridge to the root bridge? (Source: Building Redundant Switched Topologies)
- A. root
 - B. blocking
 - C. designated
 - D. nondesignated
6. Which configuration will actually form an EtherChannel link? (Source: Improving Redundant Switched Topologies with EtherChannel)
- A. Switch A: Auto, Switch B: Auto
 - B. Switch A: Desirable, Switch B: Active
 - C. Switch A: On, Switch B: On
 - D. Switch A: Passive, Switch B: On
7. Which two protocol choices do you have when you are implementing an EtherChannel bundle? (Choose two.) (Source: Improving Redundant Switched Topologies with EtherChannel)
- A. PAgP
 - B. PAgD
 - C. LACP
 - D. LAPD
8. How does STP select the designated port on a segment? (Source: Optimizing STP)
- A. highest-cost path to the closest non-root bridge
 - B. lowest-cost path to the closest non-root bridge
 - C. lowest-cost path to the root bridge
 - D. highest-cost path to the root bridge
9. With STP, what is the state of a nondesignated port? (Source: Optimizing STP)
- A. forwarding
 - B. blocking
 - C. listening
 - D. learning

10. What is the main difference between VRRP and HSRP? (Source: Understanding Layer 3 Redundancy)
- A. HSRP is a standard protocol, while VRRP is Cisco proprietary.
 - B. HSRP is configured at the global level, while VRRP is configured at the interface level.
 - C. VRRP offers Layer 2 first-hop redundancy, while HSRP offers Layer 3 first-hop redundancy.
 - D. VRRP is a standard protocol, while HSRP is Cisco proprietary.

Do Not Duplicate.
Post beta, not for release.

Answer Key

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

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate:
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Troubleshooting Basic Connectivity

This module describes approaches to troubleshooting IP connectivity.

Objectives

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

- Troubleshoot end-to-end connectivity in an IPv4 network
- Troubleshoot connectivity in an IPv6 network

Do Not Duplicate.
Post beta, not for release.

Troubleshooting IPv4 Network Connectivity

Diagnosing and resolving problems is an essential skill of network engineers. A particular problem can be diagnosed and sometimes even solved in many different ways. By using a structured approach to the troubleshooting process, you can greatly reduce the average time it takes to diagnose and solve a problem. This lesson describes various approaches to troubleshooting network connectivity issues.

Objectives

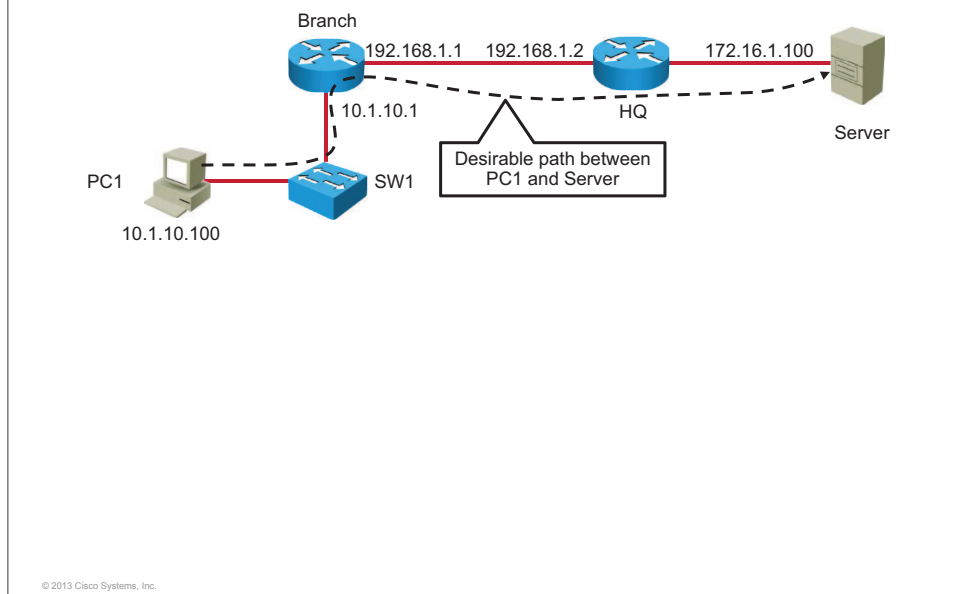
Upon completing this lesson, you will be able to meet these objectives:

- Show various components for troubleshooting IP connectivity
- Describe end-to-end connectivity troubleshooting tools
- Explain how to identify and fix physical connectivity issues
- Show how the current and desired path can be identified
- Describe how a misconfigured gateway affects connectivity
- Explain how a misconfigured name resolution entry affects network connectivity
- Explain how a misconfigured ACL affects network connectivity

Components of Troubleshooting End-to-End Connectivity

This topic describes possible causes of failed IP connectivity.

Components of Troubleshooting End-to-End Connectivity

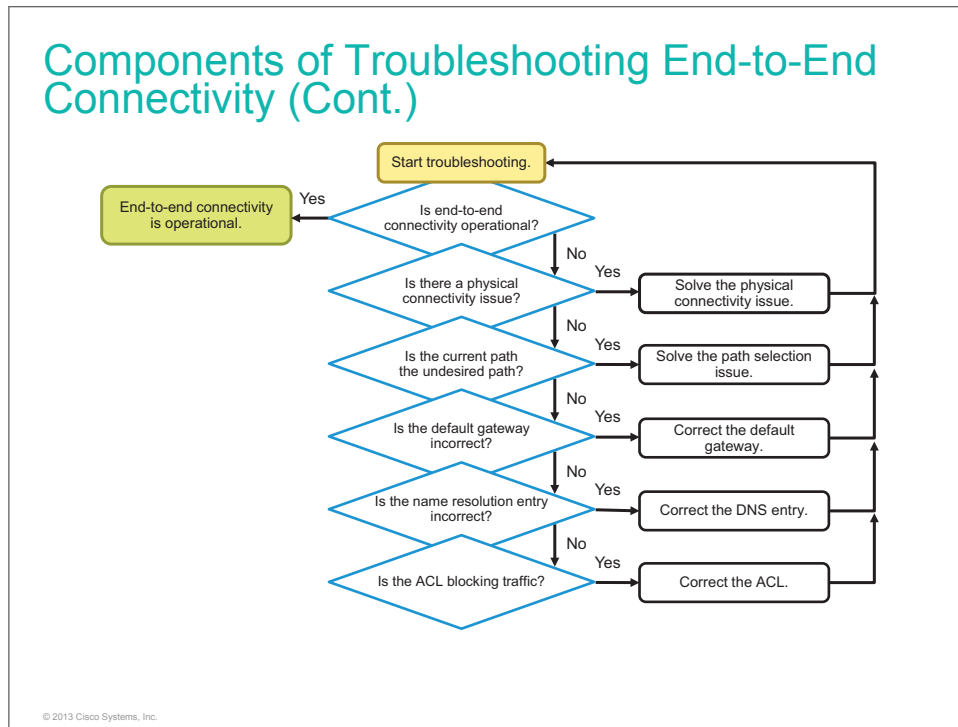


It is impossible to write out a set of troubleshooting procedures that will solve any problem. The troubleshooting process can be guided by structured methods, but the exact steps that are taken at each point along the way cannot be prescribed because they depend on many different factors. Each network is different, each problem is different, and the skill set and experience of each engineer involved in a troubleshooting process are different.

In the scenario that is used in this lesson, the troubleshooting process is described. There is PC1, which wants to access applications on the server. The desirable path is shown in the diagram.

When end-to-end connectivity is not operational, the user will inform the network administrator. The administrator will start the troubleshooting process, as shown in the figure.

Components of Troubleshooting End-to-End Connectivity (Cont.)



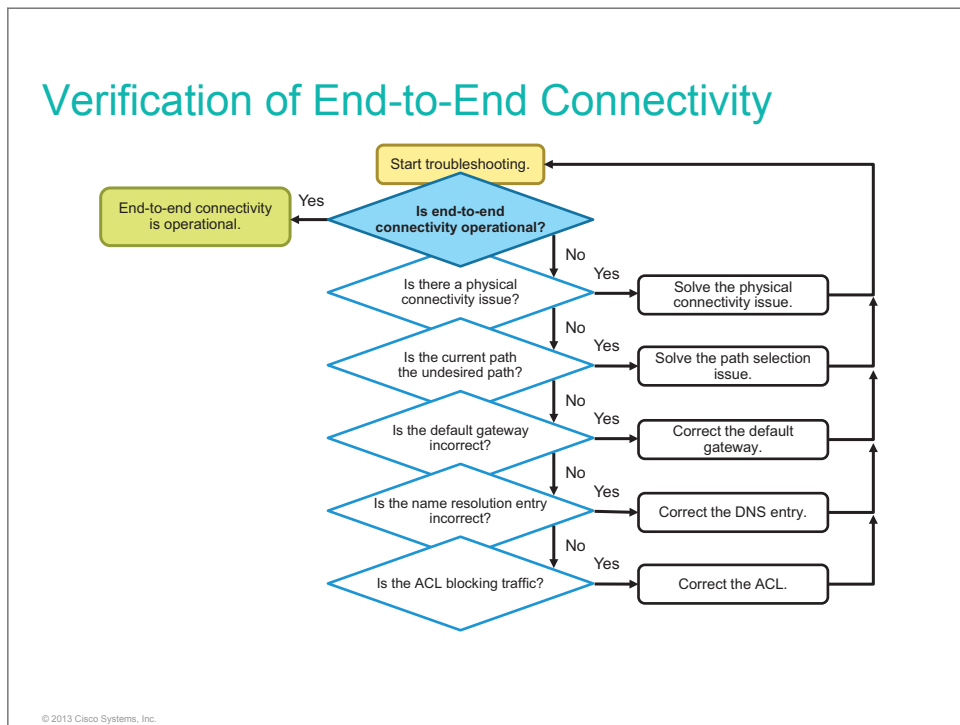
When there is no end-to-end connectivity, these are some items that you should investigate:

- Check the cables, because there might be a faulty cable or interface.
- Make sure that devices are determining the correct path from the source to the destination. Manipulate the routing information, if needed.
- Verify that the default gateway is correct.
- Verify that name resolution settings are correct.
- Verify that there are no ACLs that are blocking traffic.

After every failed troubleshooting step, a solution should be provided to make the step successful. The outcome of this process is operational, end-to-end connectivity.

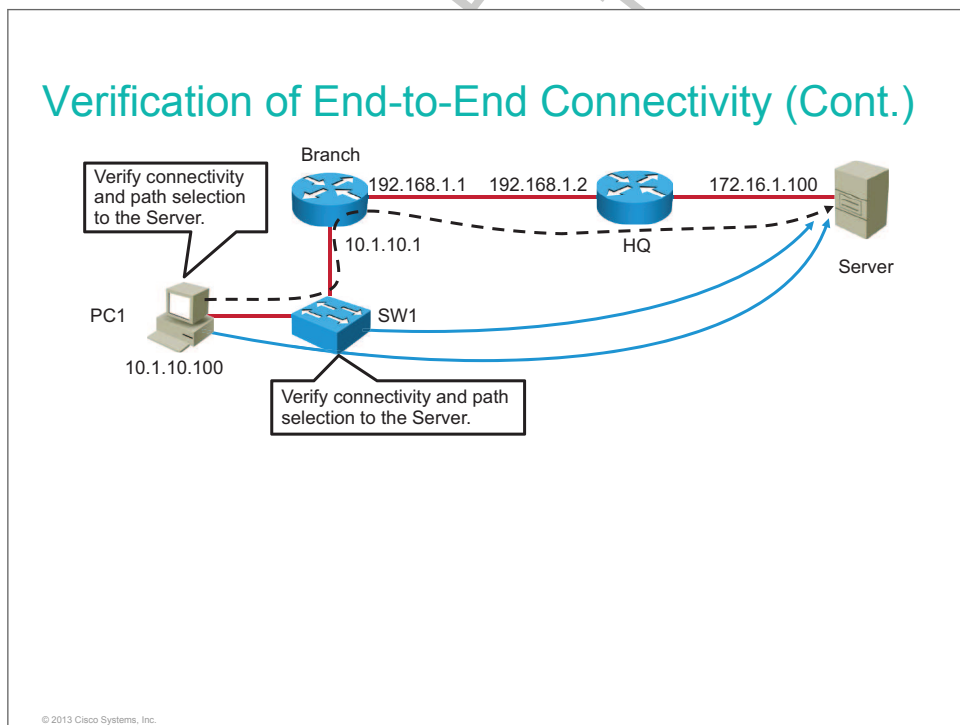
Verification of End-to-End Connectivity

This topic describes usage of end-to-end connectivity verification tools.



The **ping** command is a utility for testing IP connectivity between hosts.

The **traceroute** command is a utility that allows observation of the path between two hosts.



Verification of End-to-End Connectivity (Cont.)

```
C:\Windows\system32>ping 172.16.1.100
Pinging 172.16.1.100 with 32 bytes of data:
Reply from 172.16.1.100: bytes=32 time=8ms TTL=254
Reply from 172.16.1.100: bytes=32 time=1ms TTL=254
Reply from 172.16.1.100: bytes=32 time=1ms TTL=254
Reply from 172.16.1.100: bytes=32 time=1ms TTL=254
Ping statistics for 172.16.1.100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round-trip times in milliseconds:
        Minimum = 1ms, Maximum = 8ms, Average = 2ms
```

- Successful ping from PC1

© 2013 Cisco Systems, Inc.

Verification of End-to-End Connectivity (Cont.)

```
C:\Windows\system32>tracert 172.16.1.100
Tracing route to 172.16.1.100 over a maximum of 30 hops
  0  1 ms  <1 ms  <1 ms  10.1.10.1
  1  10 ms  2 ms  1 ms  192.168.1.2
  2  13 ms  2 ms  1 ms  172.16.1.100
Trace complete.
```

- Successful trace from PC1

© 2013 Cisco Systems, Inc.

Verification of End-to-End Connectivity (Cont.)

```
SW1#ping 172.16.1.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
```

- Successful ping from switch

```
SW1#traceroute 172.16.1.100
Type escape sequence to abort.
Tracing the route to 172.16.1.100
 0 10.1.1.1 0 msec 0 msec 0 msec
 1 192.168.1.2 1 msec 1 msec 1 msec
 2 172.16.1.100 1 msec 1 msec 1 msec
```

- Successful trace from switch

© 2013 Cisco Systems, Inc.

Ping is probably the most widely known connectivity-testing tool in networking and has been part of Cisco IOS Software from the beginning. The **ping** command is a utility for testing IP connectivity between hosts. It sends out requests for responses from a specified host address. The **ping** command uses a Layer 3 protocol that is a part of the TCP/IP suite called **ICMP**, and it uses ICMP echo request and ICMP echo reply packets. If the host at the specified address receives the ICMP echo request, it responds with an ICMP echo reply packet.

Traceroute is a utility that allows observation of the path between two hosts. Use the **traceroute** Cisco IOS command or **tracert** Windows command to observe the path between two hosts. The trace generates a list of hops that are successfully reached along the path. This list provides important verification and troubleshooting information. If the data reaches the destination, the trace lists the interface on every router in the path. If the data fails at some hop along the way, the address of the last router that responded to the trace is known. This address is an indication of where the problem or security restrictions reside.

Verification of End-to-End Connectivity (Cont.)

The **telnet** command can be used to test transport layer connectivity for any port.

```
SW1#telnet 172.16.1.100
Trying 172.16.1.100 ... Open
```

- Use Telnet to connect to the standard Telnet TCP port.

```
SW1#telnet 172.16.1.100 80
Trying 172.16.1.100, 80 ... Open
```

- Using Telnet to connect to TCP port 80 tests availability of the HTTP service.

```
SW1#telnet 172.16.1.100 25
Trying 172.16.1.100, 25 ...
Percent connection refused by remote host
```

- Using Telnet to connect to TCP port 25 tests availability of the SMTP service.

© 2013 Cisco Systems, Inc.

While ping can be used to test network layer connectivity, it can also be used to test transport layer connections from the command line. For instance, consider that you are troubleshooting a problem where someone cannot send email through a particular SMTP server. You ping the server, and it responds. This means that the network layer between you and the server is operational. Now, how do you verify the transport layer? Of course, you could configure a client and start a top-down troubleshooting procedure, but it would be convenient if you could first establish that Layer 4 is operational. So how can you do this?

Although the Telnet server application runs on its own well-known port number 23 and Telnet clients connect to this port by default, you can specify a specific port number on the client and connect to any TCP port that you want to test. At the least, this will show you if the connection is accepted (as indicated by the word “Open” in the output), if the connection is refused, or if it times out. From any of those responses, you can draw further conclusions concerning the connectivity. Certain applications, if they use an ASCII-based session protocol, might even display an application banner, or you might be able to trigger some responses from the server by typing in some keywords. Good examples of these types of protocols are SMTP, FTP, and HTTP.

Verification of End-to-End Connectivity (Cont.)

```
C:\Windows\system32>arp -a
Interface: 10.1.10.100 --- 0xd
Internet Address      Physical Address      Type
10.1.10.1             54-75-d0-8e-9a-d8    dynamic
224.0.0.22           01-00-5e-00-00-16    static
224.0.0.252          01-00-5e-00-00-fc    static
255.255.255.255      ff-ff-ff-ff-ff-ff    static
```

- Host-based tool: **arp**

© 2013 Cisco Systems, Inc.

Verification of End-to-End Connectivity (Cont.)

```
SW1#show mac address-table
      Mac Address Table
-----
Vlan  Mac Address      Type        Ports
----  -
All   0100.0ccc.cccc    STATIC      CPU
All   0100.0ccc.cccd    STATIC      CPU
1     5475.d08e.9ad8    DYNAMIC     Fa0/13
10    000c.29bc.4654    DYNAMIC     Fa0/1
10    000f.34f9.9201    DYNAMIC     Fa0/1
10    5475.d08e.9ad8    DYNAMIC     Fa0/13
Total Mac Addresses for this criterion: 6
```

- Switch tool: **show mac address-table**

© 2013 Cisco Systems, Inc.

The **arp** Windows command displays and modifies entries in the ARP cache that are used to store IP addresses and their resolved Ethernet physical (MAC) addresses. As shown in the figure, the **arp** command lists all devices that are currently in the ARP cache. The information that is displayed for each device includes the IP address, physical (MAC) address, and the type of addressing (static or dynamic).

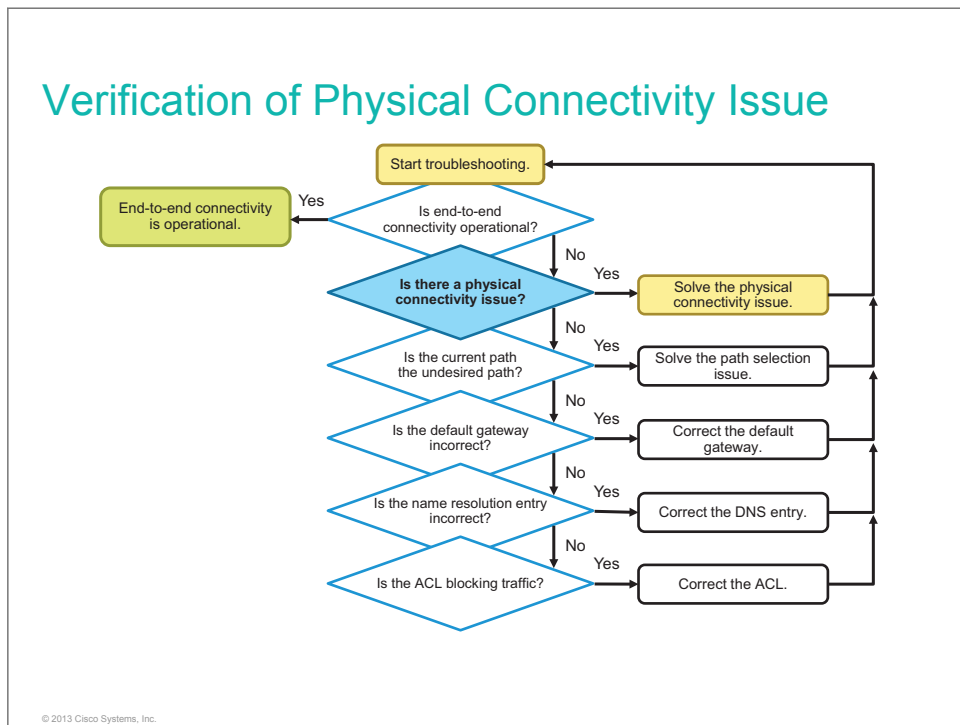
The cache can be cleared by using the **arp -d** command, if you want to repopulate the cache with updated information.

A switch forwards a frame only to the port where the destination is connected. To do this, the switch consults its MAC address table. The MAC address table lists which MAC address is connected to which port. Use the **show mac address-table** command to display the MAC address table on the switch.

Do Not Duplicate.
Post beta, not for release.

Verification of Physical Connectivity Issue


This topic describes how to identify physical connectivity issues and how to fix them.



Inevitably, troubleshooting processes involve a component of hardware troubleshooting. There are three main categories of things that could be the cause of a failure on the network: hardware failures, software failures (bugs), and configuration errors. A fourth category might be performance problems, but performance problems are more of a symptom than the cause of a problem. Having a performance problem means that there is a difference between the expected behavior and the observed behavior. This can mean one of two things: the expectation was not in line with reality, or the system is functioning as it should but the result is not what was expected or promised. In this case, the problem is not technical but organizational in nature and cannot be resolved through technical means. If the system is not functioning as expected, there must be an underlying reason that must be a hardware failure, a software failure, or a configuration error.

In essence, any network device is a specialized computer that, at a minimum, consists of a CPU, RAM, and storage, allowing it to boot and run the operating system and interfaces that allow for the reception and transmission of network traffic. Therefore, after you decide that a problem that you are observing on a given device might be hardware-related, it is worthwhile to at least verify the operation of these generic components. The most commonly used Cisco IOS commands for this purpose are the **show processes cpu**, **show memory**, and **show interface** commands. This topic discusses the **show interface** command.

Verification of Physical Connectivity Issue (Cont.)



```
Branch#show interfaces GigabitEthernet 0/1
GigabitEthernet0/1 is up, line protocol is up
<output omitted>
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
<output omitted>
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
<output omitted>
  0 output errors, 0 collisions, 1 interface resets
  0 unknown protocol drops
```

- Displays drops and errors on physical interface

© 2013 Cisco Systems, Inc.

The interfaces that the traffic passes through are another component that is always worth verifying when you are troubleshooting performance-related issues and you suspect the hardware to be at fault. In most cases, the interfaces are one of the first things that you will verify while tracing the path between devices.

The output of the **show interface** command lists these important statistics that should be checked:

- **Input queue drops:** Input queue drops (and the related ignored and throttle counters) signify that at some point more traffic was delivered to the router than it could process, which does not necessarily indicate a problem because it could be normal during traffic peaks. However, it could be an indication that the CPU cannot process packets in time, so if this number is consistently high, you should try to determine at which moments these counters are increasing and how this relates to CPU usage.
- **Output queue drops:** Output queue drops indicate that packets were dropped due to congestion on the interface. Seeing output drops is normal at any point where the aggregate input traffic is higher than the output traffic. During traffic peaks, packets are dropped if traffic is delivered to the interface faster than it can be sent out. However, although this is considered normal behavior, it leads to packet drops and queuing delays, so applications that are sensitive to packet drops and queuing delays, such as VoIP, might suffer from performance issues. Consistently seeing output drops might indicate that you need to implement an advanced queuing mechanism to provide good QoS to each application.
- **Input errors:** Input errors indicate errors that are experienced during the reception of the frame, such as CRC errors. High numbers of CRC errors could indicate cabling problems, interface hardware problems, or, in an Ethernet-based network, duplex mismatches.
- **Output errors:** Output errors indicate errors, such as collisions, during the transmission of a frame. In most Ethernet-based networks, full-duplex transmission is the norm, and half-duplex transmission is the exception. In full-duplex transmission, operation collisions cannot occur. Therefore, collisions, especially late collisions, often indicate duplex mismatches.

Verification of Physical Connectivity Issue (Cont.)

- A common cause for performance problems in Ethernet-based networks is a duplex or speed mismatch between two ends of a link.
- Duplex configuration guidelines:
 - Point-to-point Ethernet links should always run in full-duplex mode.
 - Half-duplex is not common anymore and mostly encountered if hubs are used.
 - Autonegotiation of speed and duplex is recommended on ports connected to noncritical end points.
 - Manually set the speed and duplex on links between networking devices and ports connected to critical end points.
 - Half-duplex on both ends performs better than a duplex mismatch.

© 2013 Cisco Systems, Inc.

Verification of Physical Connectivity Issue (Cont.)

```
SW1#show interfaces FastEthernet 0/1
FastEthernet0/1 is up, line protocol is up (connected)
Hardware is Fast Ethernet, address is 0017.0e6c.8e81 (bia 0017.0e6c.8e81)
<output omitted>
Full-duplex, 100Mb/s, media type is 10/100BaseTX
<output omitted>
```

- Displays duplex and speed settings

© 2013 Cisco Systems, Inc.

A common cause of interface errors is a mismatched duplex mode between two ends of an Ethernet link. In many Ethernet-based networks, point-to-point connections are the norm, and the use of hubs and the associated half-duplex operation is becoming less common. Most Ethernet links today operate in full-duplex mode, and while collisions were formerly seen as normal for an Ethernet link, collisions today often indicate that duplex negotiation has failed and that the link is not operating in the correct duplex mode.

The IEEE 802.3ab Gigabit Ethernet standard mandates the use of autonegotiation for speed and duplex. In addition, although it is not strictly mandatory, practically all Fast Ethernet NICs also use autonegotiation, by default. The use of autonegotiation for speed and duplex is the current recommended practice.

However, if duplex negotiation fails for some reason, it might be necessary to set the speed and duplex manually on both ends. Typically, this would mean setting the duplex mode to full-duplex on both ends of the connection. However, if even this cannot be made to work, running half-duplex on both ends is always to be preferred over a duplex mismatch.

Autonegotiation of speed and duplex is recommended on ports connected to noncritical end points, such as end users. You should manually set the speed and duplex on links between networking devices and ports connected to critical end points, such as servers.

The table summarizes possible settings of speed and duplex for a connection between a switch port and an end-device NIC. The table gives just a general idea about speed and duplex misconfiguration combinations.

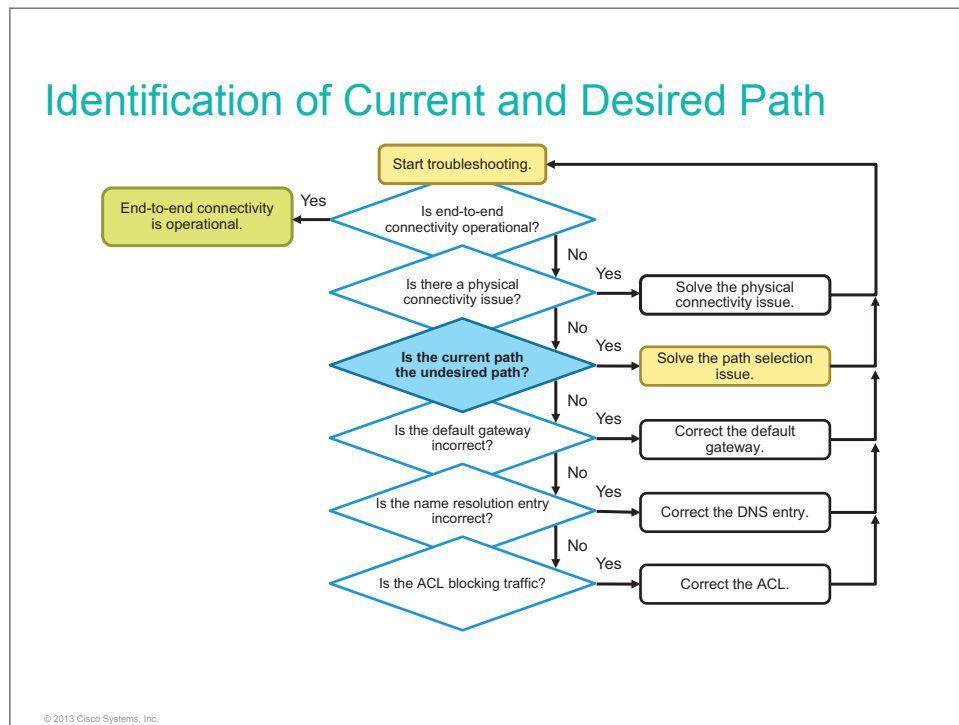
Autonegotiation Valid Configuration (10/100/1000 Mb/s NIC)

| Configuration NIC (Speed/Duplex) | Configuration Switch (Speed/Duplex) | Resulting NIC Speed/Duplex | Resulting Switch Speed/Duplex | Comments |
|----------------------------------|-------------------------------------|----------------------------|-------------------------------|---|
| AUTO | AUTO | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | Assuming maximum capability of Catalyst switch, and NIC is 1000 Mb/s, full duplex |
| 1000 Mb/s, full duplex | AUTO | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | Link is established, but the switch does not see any autonegotiation information from NIC. Since Catalyst switches support only full-duplex operation with 1000 Mb/s, they default to full duplex, and this happens only when operating at 1000 Mb/s. |
| AUTO | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | Assuming maximum capability of NIC is 1000 Mb/s, full duplex |
| 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | 1000 Mb/s, full duplex | Correct manual configuration |
| 100 Mb/s, full duplex | 1000 Mb/s, full duplex | No link | No link | Neither side establishes link due to speed mismatch |
| 100 Mb/s, full duplex | AUTO | 100 Mb/s, full duplex | 100 Mb/s, half-duplex | Duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication |

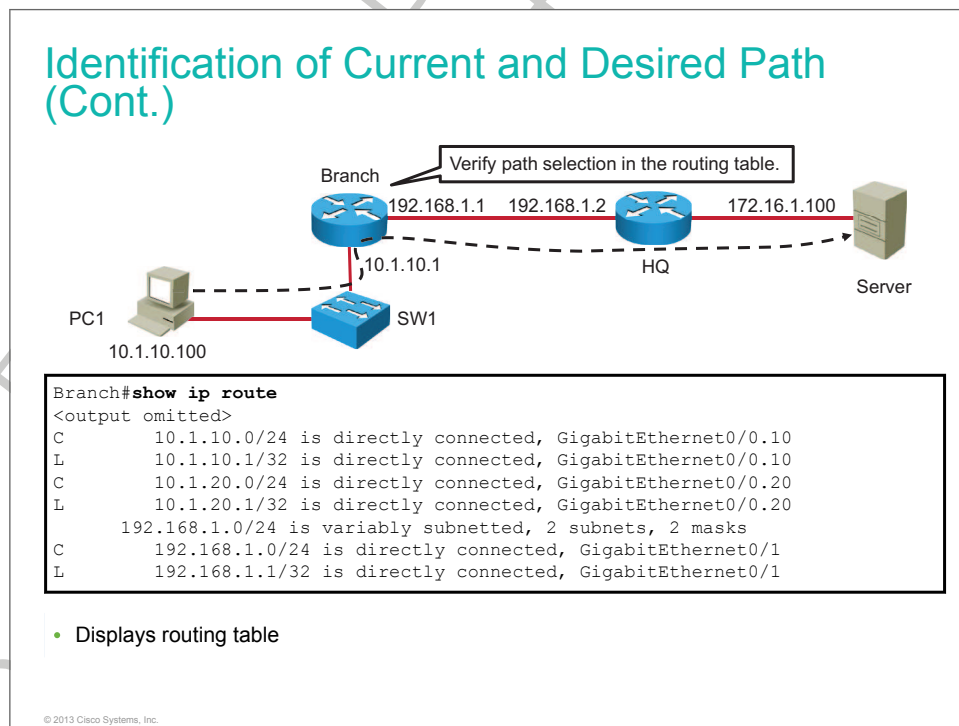
| Configuration NIC (Speed/Duplex) | Configuration Switch (Speed/Duplex) | Resulting NIC Speed/Duplex | Resulting Switch Speed/Duplex | Comments |
|----------------------------------|-------------------------------------|----------------------------|-------------------------------|---|
| AUTO | 100 Mb/s, full duplex | 100 Mb/s, half-duplex | 100 Mb/s, full duplex | Duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication |
| 100 Mb/s, full duplex | 100 Mb/s, full duplex | 100 Mb/s, full duplex | 100 Mb/s, full duplex | Correct manual configuration |
| 100 Mb/s, half-duplex | AUTO | 100 Mb/s, half-duplex | 100 Mb/s, half-duplex | Link is established, but switch does not see any autonegotiation information from NIC and defaults to half-duplex when operating at 10/100 Mb/s |
| 10 Mb/s, half-duplex | AUTO | 10 Mb/s, half-duplex | 10 Mb/s, half-duplex | Link is established, but switch does not see FLP and defaults to 10 Mb/s, half-duplex |
| 10 Mb/s, half-duplex | 100 Mb/s, half-duplex | No link | No link | Neither side establishes link due to speed mismatch |
| AUTO | 100 Mb/s, half-duplex | 100 Mb/s, half-duplex | 100 Mb/s, half-duplex | Link is established, but NIC does not see any autonegotiation information and defaults to 100 Mb/s, half-duplex |
| AUTO | 10 Mb/s, half-duplex | 10 Mb/s, half-duplex | 10 Mb/s, half-duplex | Link is established, but NIC does not see FLP and defaults to 10 Mb/s, half-duplex |

Identification of Current and Desired Path

This topic explains how to identify paths throughout the network.



To troubleshoot Layer 3 connectivity, you need to have a good understanding of the processes that are involved in routing a packet from a host across multiple routers to the final destination.



As you study the network that is shown in the figure, you should ask yourself these questions:

- Which decisions will PC1 make, which information does it need, and which actions will it perform to successfully send a packet that is destined for Server to the first-hop router Branch?

- Which decisions will router Branch make, which information does it need, and which actions will it perform to successfully send the packet from PC1 that is destined for Server to router HQ?

On the router, use the **show ip route** command to examine the routing table. The routing table on the router does not have the route to Server (172.16.1.100).

Identification of Current and Desired Path (Cont.)

- **Directly connected:** Router attaches to this network
- **Local host routes:** Local IP address on the router interface
- **Static routing:** Entered manually by a system administrator
- **Dynamic routing:** Learned by exchange of routing information
- **Default route:** Statically or dynamically learned—used when no explicit route to network is known

```
Router#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override
<output omitted>
```

- Displays routing table codes

© 2013 Cisco Systems, Inc.

The routing tables can be populated by these methods:

- **Directly connected networks:** This entry comes from having router interfaces that are directly attached to network segments. This method is the most certain method of populating a routing table. If the interface fails or is administratively shut down, the entry for this network will be removed from the routing table. The administrative distance is 0 and therefore will pre-empt all other entries for this destination network. Entries with the lowest administrative distance are the best, most-trusted sources.
- **Local host routes:** This entry comes from the local IP address on the router interface. The subnet mask represents the host route.
- **Static routes:** A system administrator manually enters static routes directly into the configuration of a router. The default administrative distance for a static route is 1. Therefore, the static routes will be included in the routing table, unless there is a direct connection to this network. Static routes can be an effective method for small, simple networks that do not change frequently. For larger and unstable networks, the solution with static routes does not scale.
- **Dynamic routes:** The router learns dynamic routes automatically when the routing protocol is configured and a neighbor relationship to other routers is established. The information is responsive to changes in the network and updates constantly. There is, however, always a lag between the time that a network changes and when all of the routers become aware of the change. The time delay for a router to match a network change is called convergence time. A shorter convergence time is better for users of the network. Different routing protocols perform differently in this regard. Larger networks require the dynamic routing method, because there are usually many addresses and constant changes. These changes require updates to routing tables across all routers in the network, or connectivity is lost.

- **Default routes:** A default route is an optional entry that is used when no explicit path to a destination is found in the routing table. The default route can be manually inserted, or it can be populated from a dynamic routing protocol.

The **show ip route** command displays the routing table in a router. The first part of the output explains the codes, presenting the letters and the associated source of the entries in the routing table.

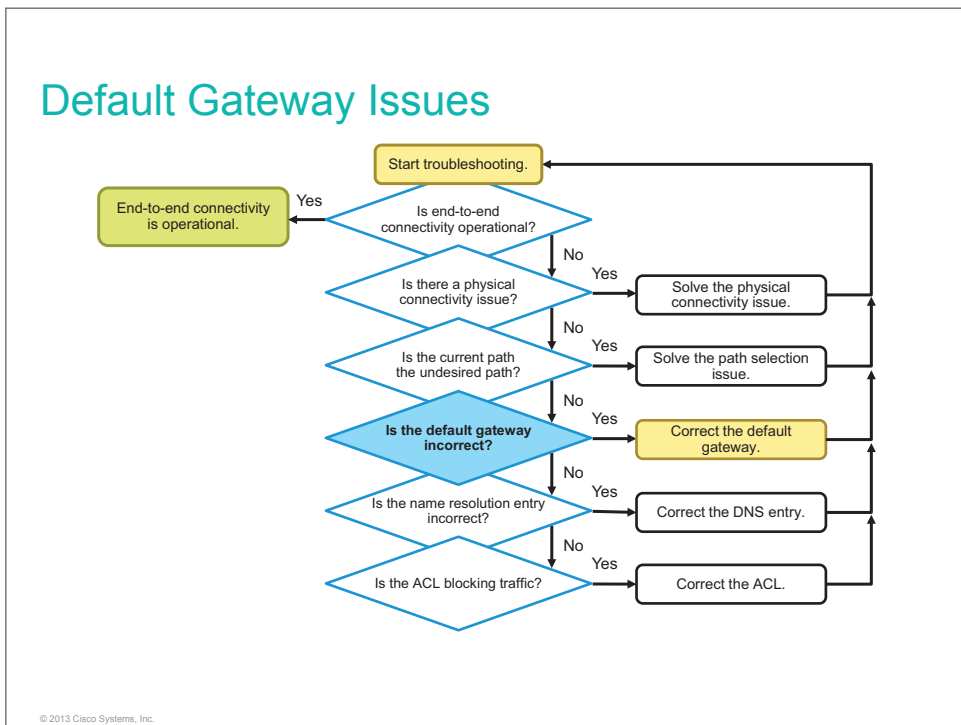
- **L:** Reserved for the local host route
- **C:** Reserved for directly connected networks
- **S:** Reserved for static routes
- **R:** Reserved for RIP
- **O:** Reserved for the OSPF routing protocol
- **D:** Reserved for EIGRP. Letter D stands for DUAL, the update algorithm that is used by EIGRP.

If the destination address in a packet:

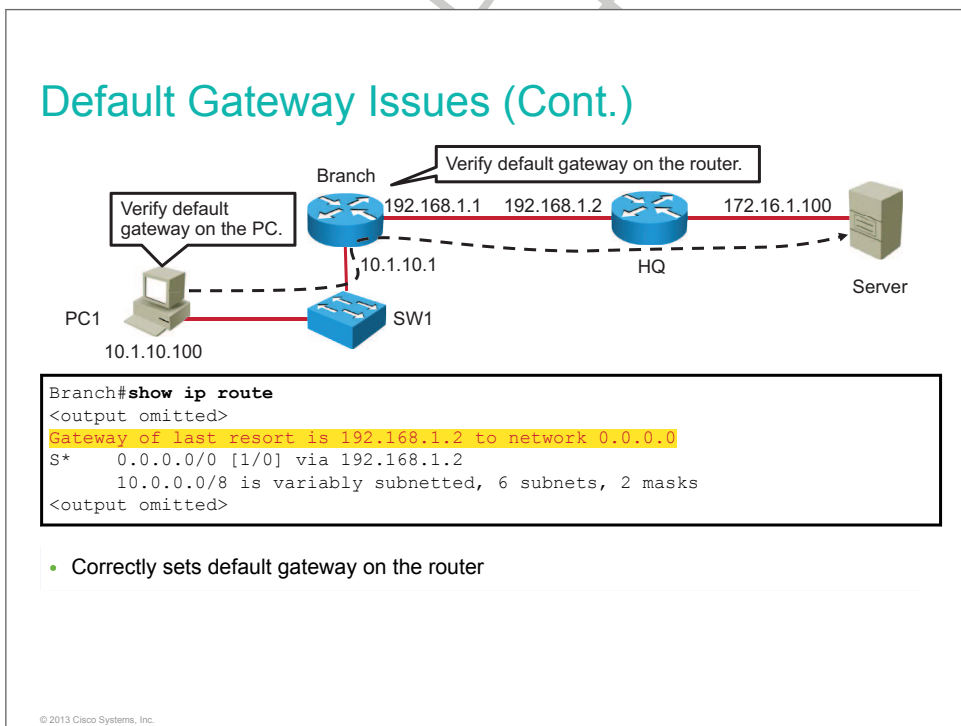
- does not match an entry in the routing table, then the default route is used. If there is not a default route that is configured, the packet is discarded.
- matches a single entry in the routing table, then the packet is forwarded through the interface that is defined in this route.
- matches more than one entry in the routing table and the routing entries have the same prefix (network mask), then the packets for this destination can be distributed among the routes that are defined in the routing table.
- matches more than one entry in the routing table and the routing entries have different prefixes (network masks), then the packets for this destination are forwarded out of the interface that is associated with the route that has the longer prefix match.

Default Gateway Issues

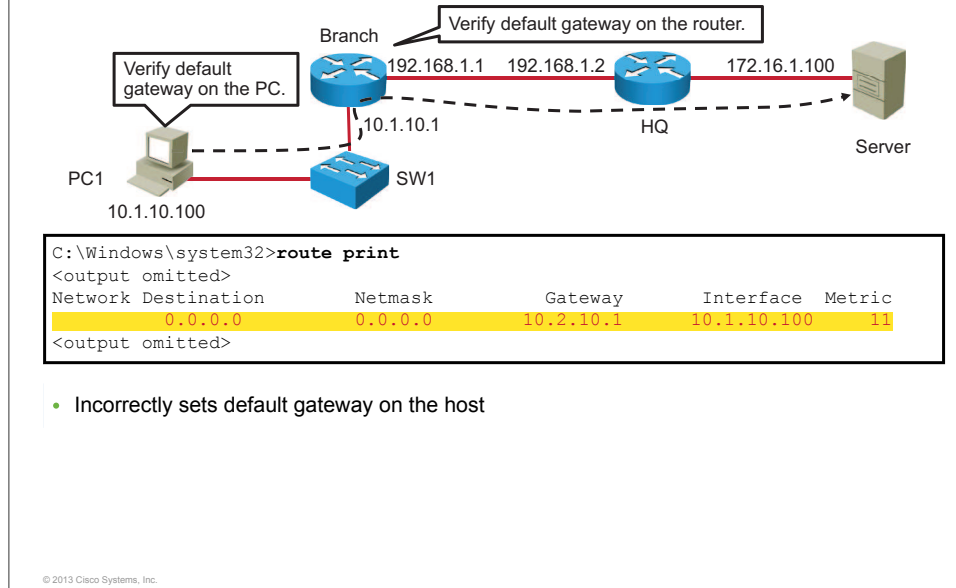
This topic identifies how a misconfigured default gateway will influence network behavior.



In the absence of a detailed route on the router or an incorrect default gateway on the host, communication between two endpoints in different networks will not work.



Default Gateway Issues (Cont.)

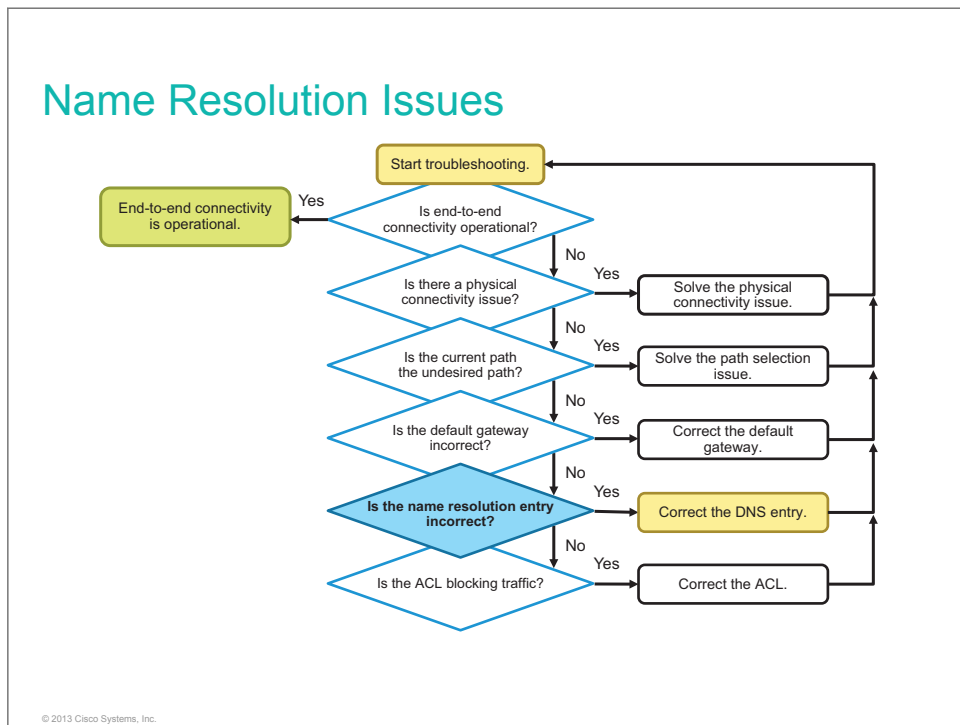


Use the **show ip route** Cisco IOS command and the **route print** Windows command to verify the presence of the default gateway.

In the example, the Branch router has the correct default gateway, which is the IP address of the HQ router. PC1 has the wrong default gateway. PC1 should have the default gateway of Branch router 10.1.10.1.

Name Resolution Issues

This topic identifies how missing name resolution mapping will influence network behavior.



The mapping of computer names to IP addresses can be done in two ways:


- **Static:** The system administrator creates a text file, called a hosts file, and enters each computer name and IP address. The file is then distributed on the network. When a request for a connection to another computer is made, the file is used to resolve the name to the correct IP address. This system works well for simple networks that change infrequently.
- **Dynamic:** The DNS protocol controls the DNS, a distributed database with which you can map host names to IP addresses.

When you configure name resolution on the device, you can substitute the host name for the IP address with all IP commands, such as **ping** or **telnet**.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are formed with periods (.) as the delimiting characters. For example, Cisco is a commercial organization that IP identifies by a .com domain name, so its domain name is cisco.com. A specific device in this domain, for example, the FTP system, is identified as ftp.cisco.com.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names that is mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, specify the name server that is present on your network, and enable the DNS.

Name Resolution Issues (Cont.)



```
172.16.1.100 Server
```


- Enter name for IP mapping in the hosts file on the PC.

© 2013 Cisco Systems, Inc.

The hosts file provides the function of translating human-friendly host names into IP addresses that identify and locate hosts in an IP network. In some operating systems, the content of the hosts file is preferred over other methods, such as the DNS. Unlike the DNS, the hosts file is under the direct control of the local computer administrator.

In a Windows operating system, the file is located at `c:\windows\system32\drivers\etc\hosts`. Other operating systems may have the hosts file in a different location, use a different file, or may not have it at all. Open the hosts file in a text editor such as Notepad.

Name Resolution Issues (Cont.)



```
C:\Windows\system32>ping Server
Pinging Server [172.16.1.100] with 32 bytes of data:
Reply from 172.16.1.100: bytes=32 time=47ms TTL=254
Reply from 172.16.1.100: bytes=32 time=36ms TTL=254
Reply from 172.16.1.100: bytes=32 time=36ms TTL=254
Reply from 172.16.1.100: bytes=32 time=36ms TTL=254
Ping statistics for 172.16.1.100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 36ms, Maximum = 47ms, Average = 38ms
```

- Verifies connectivity of the server, using the **ping** command and the host name as the destination

© 2013 Cisco Systems, Inc.

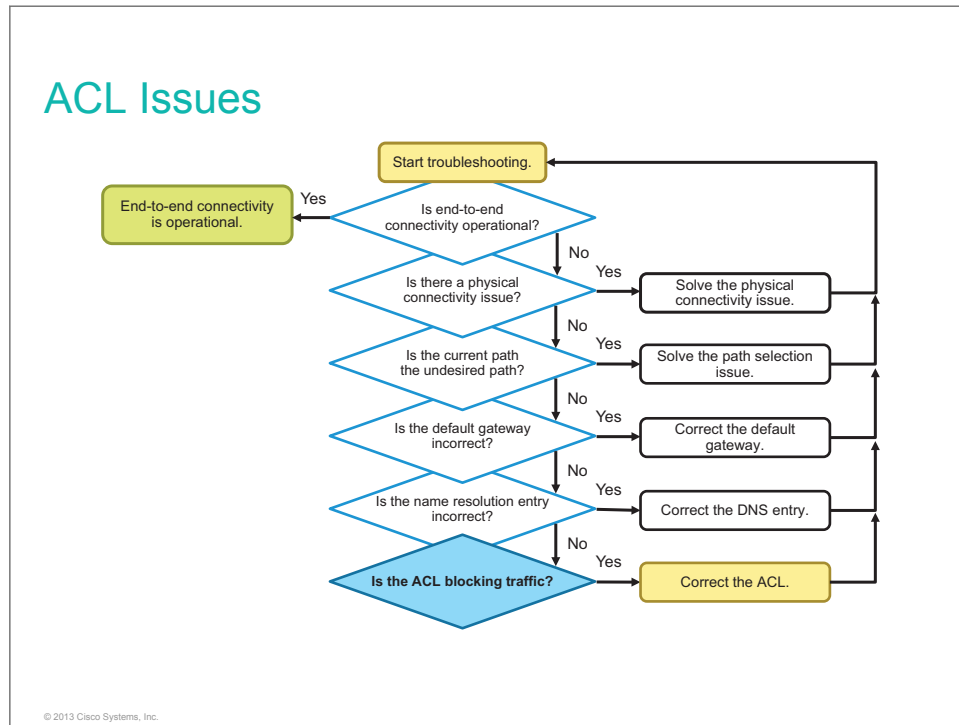
To verify static name resolution, verify connectivity to the server using the host name **Server**. The ping should be successful.

To make a static name resolution entry on a switch or a router, use the **ip host name ip_address** command. For example, if you want to add an entry "Server" that will resolve into IP address 172.16.1.100, the syntax would be **ip host Server 172.16.1.100**.

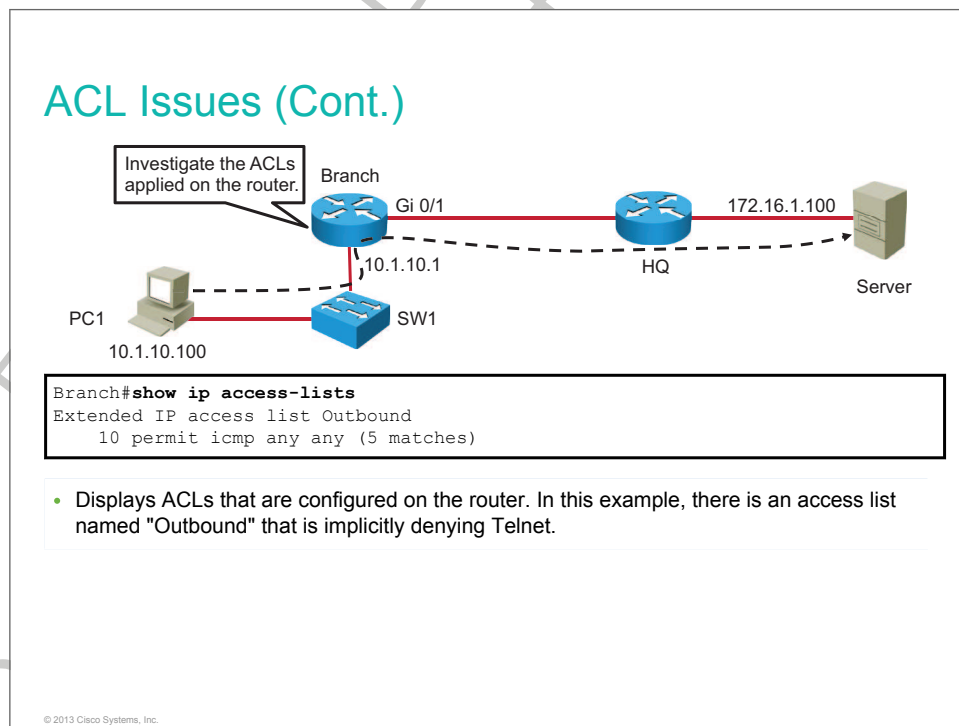
Do Not Duplicate.
Post beta, not for release.

ACL Issues

This topic identifies how ACLs can influence end-to-end connectivity.




On the routers, there may be ACLs configured that prohibit a protocol to pass the interface in the inbound or outbound direction.



Use the **show ip access-lists** command to display the contents of all ACLs. By entering the ACL name or number as an option for this command, you can display a specific ACL.

ACL Issues (Cont.)

PC is unable to use Telnet to connect to the Server. Is there an ACL on the Gi 0/1?



```
Branch#show ip interface GigabitEthernet 0/1 | include access list
Outgoing access list is Outbound
Inbound access list is not set
```

- Displays placement of the ACL on the interface


© 2013 Cisco Systems, Inc.

The **show ip interfaces** command displays IP interface information and indicates whether any IP ACLs are set on the interface.

In the **show ip interfaces GigabitEthernet0/1** command output that is shown, IP ACL Outbound has been configured on the interface as an outbound ACL.

ACL Issues (Cont.)

Correct the ACL so that the router will permit Telnet traffic.




```
Branch(config)#ip access-list extended Outbound
Branch(config-ext-nacl)#permit tcp any any eq 23
```

- Adds the ACL entry to allow Telnet

© 2013 Cisco Systems, Inc.

IP ACL Outbound permits only the ICMP protocol, which is why ping will work. In order to allow Telnet from PC1 to the server, you need to add an entry in IP ACL Outbound to allow TCP protocol and port 23 for Telnet.

ACL Issues (Cont.)



```
Branch#show ip access-lists
Extended IP access list Outbound
 10 permit icmp any any (5 matches)
 20 permit tcp any any eq telnet (17 matches)
```

- Displays the corrected ACLs that are configured on the router

© 2013 Cisco Systems, Inc.

After correcting IP ACL Outbound, a Telnet connection from PC1 to the server will be successful.

Summary

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

Summary

- First test end-to-end connectivity by using the **ping**, **tracert**, or **tracert** commands.
- Isolate physical connectivity issues by examining the output of the **show interface** command.
- Make sure devices are determining the correct path from source to the destination.
- If there is no exact route to the destination, verify the default gateway on the devices.
- Adjust the name resolution entry to represent the current scenario.
- Adjust ACL entries to allow end-to-end connectivity.

© 2013 Cisco Systems, Inc.

Troubleshooting IPv6 Network Connectivity

As with IPv4 networks, problems will arise in IPv6 networks. You will be required to troubleshoot these problems, and you can use the same structured approach as troubleshooting IPv4 networks. However, because there are differences in IPv4 and IPv6 operations, troubleshooting IPv6 networks has its own specifics. For example, instead of verifying ARP entries in IPv4, you have to verify neighbor discovery entries in IPv6. This lesson first describes how to verify end-to-end IPv6 connectivity. The lesson then describes how to verify current paths in networks and how to verify DNS and default gateway settings. The lesson concludes with verification of IPv6 ACLs.

Objectives

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


- Explain types of IPv6 unicast addresses
- Describe possible causes of failed IPv6 connectivity
- Describe usage of IPv6 end-to-end connectivity verification tools
- Explain how to identify IPv6 paths throughout the network
- Describe how to verify that the IPv6 default gateway is correctly set
- Identify how missing IPv6 name resolution mapping will influence network behavior
- Identify how ACLs can influence end-to-end-IPv6 connectivity

IPv6 Unicast Addresses

This topic provides an overview of IPv6 addressing.

IPv6 Unicast Addresses

- **Global:** Starts with 2000:: $/3$ and assigned by IANA



© 2013 Cisco Systems, Inc.

There are several basic types of IPv6 unicast addresses: global, reserved, private (link-local), loopback, and unspecified.

RFC 4291 specifies 2000:: $/3$ to be the global unicast address space that the IANA may allocate to the RIRs. A global unicast address is an IPv6 address from the global unicast prefix. The structure of global unicast addresses enables the aggregation of routing prefixes, which limits the number of routing table entries in the global routing table. Global unicast addresses that are used on links are aggregated upward through organizations and eventually to the ISPs.

Link-local addresses are a new concept of addressing with IP in the network layer. These addresses refer only to a particular physical link. Link-local addresses typically begin with "FE80." The next digits can be defined manually. If you do not define them manually, the interface MAC address is used, based on the EUI-64 format.

As in IPv4, there is a provision for a special loopback IPv6 address for testing. Datagrams that are sent to this address "loop back" to the sending device. However, in IPv6 there is just one address rather than an entire block for this function. The loopback address is 0:0:0:0:0:0:0:1, which is normally expressed as "::1".


In IPv4, an IP address of all zeros has a special meaning. It refers to the host itself and is used when a device does not know its own address. In IPv6, this concept has been formalized, and the all-zero address is named the "unspecified" address. It is typically used in the source field of a datagram that is sent by a device that seeks to have its IP address configured. You can apply address compression to this address—because the address is all zeros, the address becomes simply "::".

The IETF reserved a portion of the IPv6 address space for various uses, both present and future. Reserved addresses represent 1/256th of the total IPv6 address space.

- The lowest address within each subnet prefix (the interface identifier set to all zeros) is reserved as the "subnet-router" anycast address.
- The 128 highest addresses within each $/64$ subnet prefix are reserved to be used as anycast addresses.

IPv6 Unicast Addresses (Cont.)

- **Private:** link-local (starts with FE80::/10)

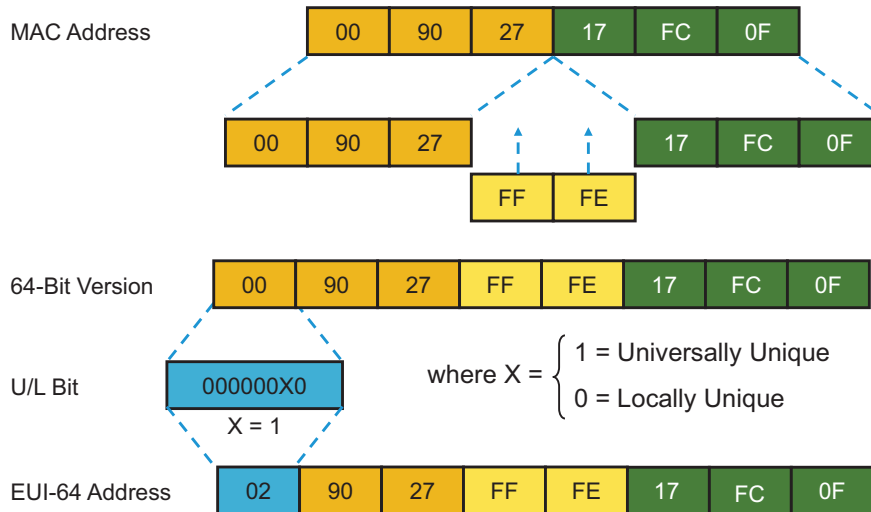


- **Loopback:** (::1)
- **Unspecified:** (::)
- **Reserved:** Used by the IETF

© 2013 Cisco Systems, Inc.

EUI-64 Interface ID Assignment

IPv6 Unicast Addresses (Cont.)

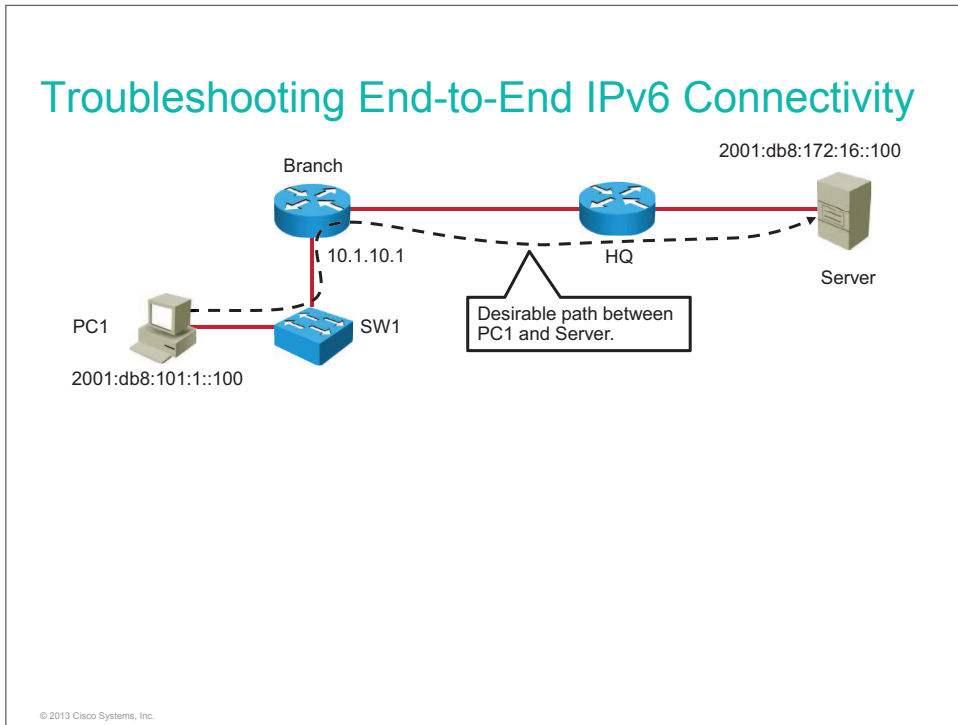


© 2013 Cisco Systems, Inc.

The EUI-64 standard explains how to stretch IEEE 802 MAC addresses from 48 to 64 bits by inserting the 16-bit 0xFFFE in the middle (at the 24th bit) of the MAC address to create a 64-bit, unique interface identifier. In the first byte of the OUI, bit 7 indicates the scope: 0 for global and 1 for local. Because most burned-in addresses are globally scoped, bit 7 will usually be 0. The EUI-64 standard also specifies that the value of the seventh bit be inverted. So, for example, MAC address **00-90-27-17-FC-0F** becomes **02-90-27-17-FC-0F**. The resulting EUI-64 address on network 2001:0DB8:0:1::/64 would be 2001:0DB8:0:1:0290:27FF:FE17:FC0F.

Troubleshooting End-to-End IPv6 Connectivity

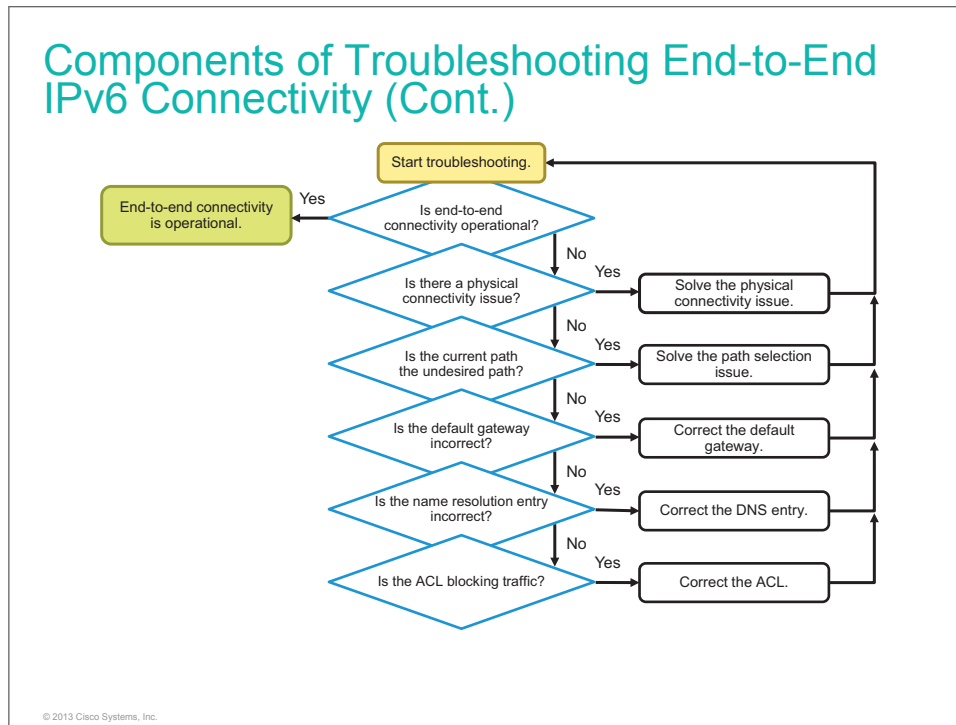
This topic describes possible causes of failed IPv6 connectivity and procedures for how to troubleshoot IPv6 connectivity.



As with troubleshooting IPv4 connectivity, the troubleshooting process for IPv6 can be guided by structured methods. The overall troubleshooting procedure is the same as troubleshooting IPv4, with differences that are related to IPv6 specifics.

In the scenario that is used in this lesson, the troubleshooting process is described. There is a PC1 that wants to access applications on the server. The lesson will focus only on troubleshooting steps that are specific to IPv6. The desirable path is shown in the figure.

Components of Troubleshooting End-to-End IPv6 Connectivity (Cont.)



When end-to-end connectivity is not operational, the user will inform the network administrator. The administrator will start the troubleshooting process, as shown in the figure.

When there is no end-to-end connectivity, these are some items that you would want to investigate:

- If there is an issue with physical connectivity, solve it by adjusting the configuration or changing the hardware.
- Make sure that devices are determining the correct path from the source to the destination. Manipulate the routing information if needed.
- Verify that the default gateway is correct.
- Check if everything is correct regarding the name resolution settings. There should be a name resolution server that is accessible over IPv4 or IPv6.
- Verify that there are no ACLs blocking traffic.

After every failed troubleshooting step, a solution should be provided to make the step successful. The outcome of this process is operational, end-to-end connectivity.

Verification of End-to-End IPv6 Connectivity

This topic describes the usage of end-to-end IPv6 connectivity verification tools.

Verification of End-to-End IPv6 Connectivity

```
C:\Windows\system32>ping 2001:DB8:172:16::100
```

- The ping utility on the PC can be used to test IPv6 connectivity.

```
C:\Windows\system32>tracert 2001:DB8:172:16::100
```

- The traceroute utility on a PC allows observation of the IPv6 path.

© 2013 Cisco Systems, Inc.

The ping utility can be used to test end-to-end IPv6 connectivity by providing the IPv6 address as the destination address. The utility recognizes the IPv6 address when one is provided and uses IPv6 as a protocol to test connectivity.

The ping utility on the PC can be used to test IPv6 connectivity:


```
C:\Windows\system32>ping 2001:DB8:172:16::100
Pinging 2001:db8:172:16::100 with 32 bytes of data:
Reply from 2001:db8:172:16::100: time=19ms
Reply from 2001:db8:172:16::100: time=1ms
Reply from 2001:db8:172:16::100: time=1ms
Reply from 2001:db8:172:16::100: time=1ms
Ping statistics for 2001:db8:172:16::100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 19ms, Average = 5ms
```

Traceroute is a utility that allows observation of the path between two hosts and supports IPv6. Use the **tracert** Cisco IOS command or **tracert** Windows command, followed by the IPv6 destination address, to observe the path between two hosts. The trace generates a list of IPv6 hops that are successfully reached along the path. This list provides important verification and troubleshooting information.

The traceroute utility on the PC allows observation of the IPv6 path:

```
C:\Windows\system32>tracert 2001:DB8:172:16::100
Tracing route to 2001:db8:172:16::100 over a maximum of 30 hops
  0  1 ms  1 ms  <1 ms  2001:db8:101:1::1
  1  10 ms  1 ms  1 ms  2001:db8:172:16::100
Trace complete.
```

Verification of End-to-End IPv6 Connectivity (Cont.)



```
Branch#ping 2001:DB8:172:16::100
```

- The ping utility on the router can be used to test IPv6 connectivity.

```
Branch#traceroute 2001:DB8:172:16::100
```

- Successful trace from a router to verify the IPv6 path.

© 2013 Cisco Systems, Inc.

The ping utility on the router can be used to test IPv6 connectivity:

```
Branch#ping 2001:DB8:172:16::100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:172:16::100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

The traceroute utility on the router allows observation of the IPv6 path:

```
Branch#traceroute 2001:DB8:172:16::100
Type escape sequence to abort.
Tracing the route to 2001:DB8:172:16::100
 1 2001:DB8:209:165::2 0 msec 0 msec 0 msec
```

Verification of End-to-End IPv6 Connectivity (Cont.)

The **telnet** command can be used to test transport layer connectivity for any TCP port over IPv6.

```
C:\Windows\system32>telnet 2001:DB8:172:16::100
HQP#
```

- Use Telnet to connect to the standard Telnet TCP port from a PC.

```
C:\Windows\system32>telnet 2001:DB8:172:16::100 80
HTTP/1.1 400 Bad Request
Date: Wed, 26 Sep 2012 07:27:10 GMT
Server: cisco-IOS
Accept-Ranges: none
400 Bad Request
Connection to host lost.
```

- Use Telnet to connect to TCP port 80, which tests availability of the HTTP service.

© 2013 Cisco Systems, Inc.

Similar to IPv4, you can use Telnet to test end-to-end transport layer connectivity over IPv6 using the **telnet** command from a PC, router, or switch. When you provide the IPv6 destination address, the protocol stack determines that the IPv6 protocol has to be used. If you omit the port number, the client will connect to port 23. You can also specify a specific port number on the client and connect to any TCP port that you want to test.

In the example, you can see two connections from a PC to the Cisco IOS router. The first one connects to port 23 and tests Telnet over IPv6. The second connects to port 80 and tests HTTP over IPv6.

Verification of End-to-End IPv6 Connectivity (Cont.)

```
C:\Windows\system32>netsh interface ipv6 show neighbor
Interface 13: LAB
Internet Address          Physical Address          Type
-----
fe80::9c5a:e957:a865:bde9 00-0c-29-36-fd-f7        Stale
fe80::fa66:f2ff:fe31:7250 f8-66-f2-31-72-50        Reachable (Router)
ff02::2                  33-33-00-00-00-02        Permanent
ff02::16                 33-33-00-00-00-16        Permanent
ff02::1:2                33-33-00-01-00-02        Permanent
ff02::1:3                33-33-00-01-00-03        Permanent
ff02::1:ff05:f9fb        33-33-ff-05-f9-fb        Permanent
ff02::1:ff31:7250        33-33-ff-31-72-50        Permanent
ff02::1:ff65:bde9        33-33-ff-65-bd-e9        Permanent
ff02::1:ff67:bae4        33-33-ff-67-ba-e4        Permanent
```

- Neighbor discovery table on a PC

© 2013 Cisco Systems, Inc.

Verification of End-to-End IPv6 Connectivity (Cont.)

```
Branch#show ipv6 neighbors
IPv6 Address                               Age Link-layer Addr State Interface
FE80::21E:7AFF:FE79:7A81                   8 001e.7a79.7a81 STALE Gi0/1
2001:DB8:101:1:A083:AEE4:E7C5:2CCA        46 000c.2936.fdf7 STALE Gi0/0
2001:DB8:209:165::2                        0 001e.7a79.7a81 REACH Gi0/1
2001:DB8:101:1:C31:CD87:7505:F9FB         0 000c.2952.51fd REACH Gi0/0
```

- Neighbor discovery table on a router

© 2013 Cisco Systems, Inc.

When troubleshooting end-to-end connectivity, it is useful to verify mappings between destination IP addresses and Layer 2 Ethernet addresses on individual segments. In IPv4, this functionality is provided by ARP. In IPv6, the ARP functionality is replaced by the neighbor discovery process and ICMPv6. The neighbor discovery table caches IP addresses and their resolved Ethernet physical (MAC) addresses. As shown in the figure, the **netsh interface ipv6 show neighbor** Windows command lists all devices that are currently in the neighbor discovery table cache. The information that is displayed for each device includes the IP address, physical (MAC) address, and type of addressing. By examining the neighbor discovery table, you can verify that destination IPv6 addresses map to correct Ethernet addresses.

The figure also shows an example of the neighbor discovery table on the Cisco IOS router. The table includes the IPv6 address of the neighbor, age in minutes since the address was confirmed to be reachable, and state. The states are explained in the table:

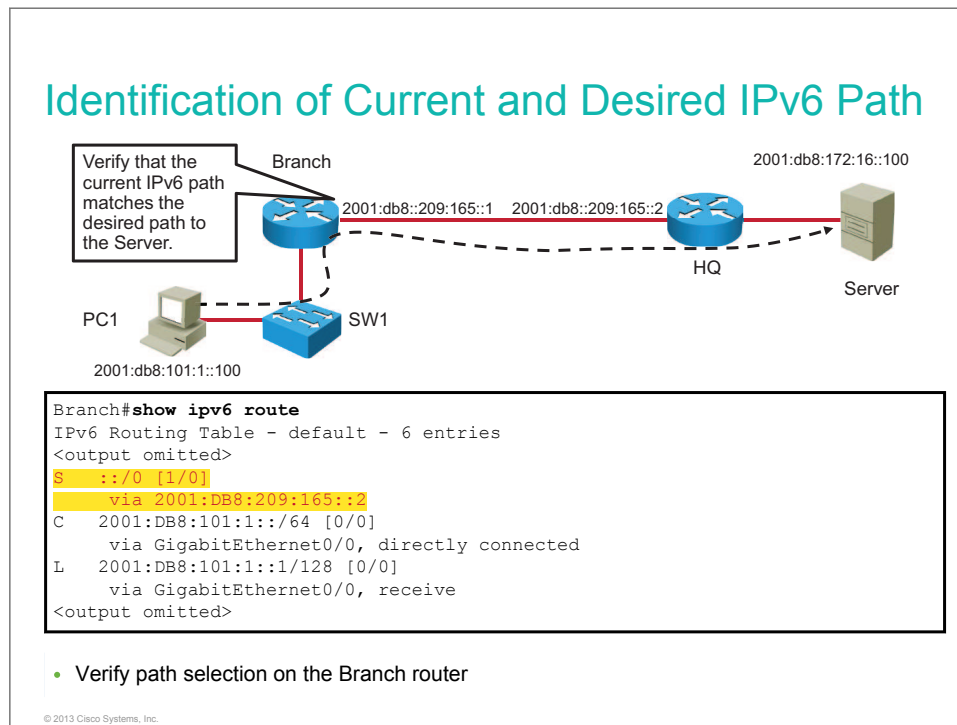
| State | Description |
|--------------------|--|
| INCMP (Incomplete) | Address resolution is being performed on the entry. A neighbor solicitation message has been sent to the solicited-node multicast address of the target, but the corresponding neighbor advertisement message has not yet been received. |
| REACH (Reachable) | Positive confirmation was received within the last ReachableTime milliseconds that the forward path to the neighbor was functioning correctly. While in REACH state, the device takes no special action as packets are sent. |
| STALE | More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. While in STALE state, the device takes no action until a packet is sent. |

| State | Description |
|-------|--|
| DELAY | More than ReachableTime milliseconds have elapsed since the last positive confirmation was received that the forward path was functioning properly. A packet was sent within the last DELAY_FIRST_PROBE_TIME seconds. If no reachability confirmation is received within DELAY_FIRST_PROBE_TIME seconds of entering the DELAY state, send a neighbor solicitation message and change the state to PROBE. |
| PROBE | A reachability confirmation is actively sought by resending neighbor solicitation messages in RetransTimer milliseconds until a reachability confirmation is received. |

Do Not Duplicate.
Post beta, not for release.

Identification of Current and Desired IPv6 Path

This topic describes how to identify the current and desired path in an IPv6 network.



To verify that the current IPv6 path matches the desired path to reach destinations, use the **show ipv6 route** command on a router to examine the routing table.

The routing table on the Branch router in the example has a default route that is configured and will be used to route packets to the server (2001:db8:172:16::100).

Default Gateway Issues in IPv6

In the absence of the default gateway on a host, communication between two endpoints in a different network will not work. This topic describes how to verify that the default gateway is correctly set.

Default Gateway Issues in IPv6

```
C:\Windows\system32>ipconfig
Windows IP Configuration
    Connection-specific DNS Suffix  . : 
    IPv6 Address. . . . .           : 2001:db8:101:1:dd42:a044:fa67:bae4
    Temporary IPv6 Address. . . . . : 2001:db8:101:1:c31:cd87:7505:f9fb
    Link-local IPv6 Address . . . . . : fe80::dd42:a044:fa67:bae4%13
    IPv4 Address. . . . .           : 10.1.1.100
    Subnet Mask . . . . .           : 255.255.255.0
    Default Gateway . . . . .       : fe80::fa66:f2ff:fe31:7250%13
                                   10.1.1.1
```

- Verify the default gateway on a PC

© 2013 Cisco Systems, Inc.

If a PC needs access to other networks in addition to the directly connected network, correct configuration of the default gateway is very important. If a PC has to send a packet to a network that is *not* directly connected, the packet has to be sent to the default gateway, which is the first router on the path to destinations. The default gateway then forwards the packet toward the destination.

Note You will see a percent sign (%), followed by a number, at the end of the IPv6 link-local address and at the end of the default gateway. The number following the percent sign identifies an interface on the PC and is not part of the IPv6 address, and should be ignored when determining the IPv6 address of the default gateway.

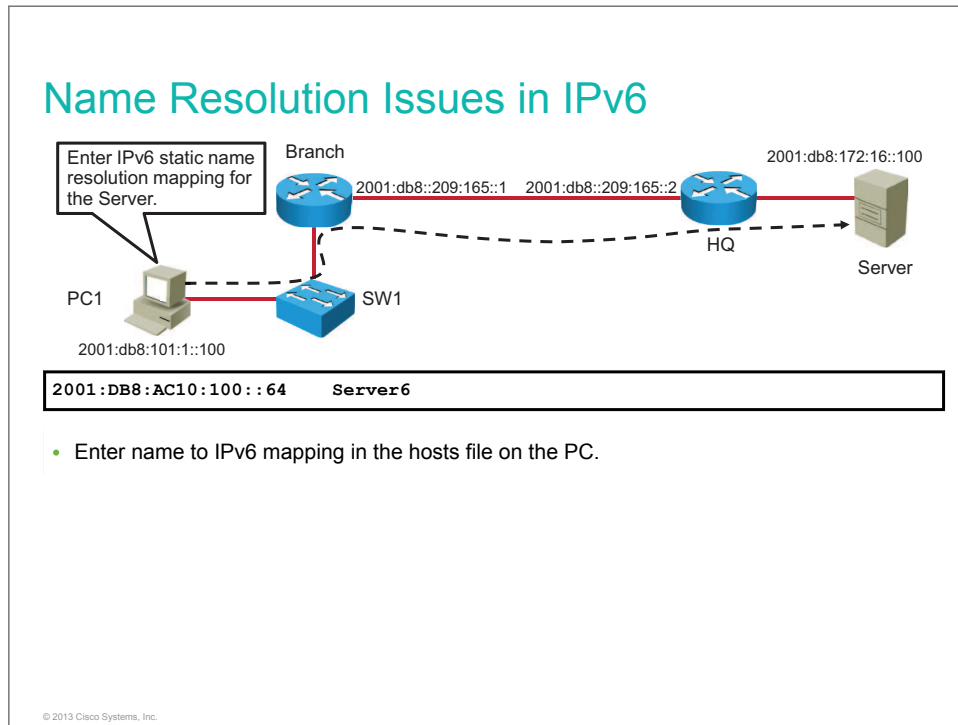
In IPv6, the default gateway can be configured using stateless autoconfiguration or manually. In the case of stateless autoconfiguration, the default gateway is advertised to PCs that are using route advertisements. In IPv6, the IPv6 address that is advertised inside route advertisements as a default gateway is the link-local IPv6 address of a router interface. If the default gateway is configured manually, which is unlikely, the default gateway can be set either to the global IPv6 address or to the link-local IPv6 address.

Note A link-local address is intended only for communications within the segment of a local network or a point-to-point connection that a host is connected to. The link-local IPv6 addresses are assigned with the fe80::/64 prefix.

To verify that a PC has the default gateway set, you can use the **ipconfig** command on a Microsoft Windows PC or the **ifconfig** command on Linux and Mac OS X. In the example, the PC has the IPv6 default gateway set to the link-local address of the Branch router.

Name Resolution Issues in IPv6

This topic identifies how missing IPv6 name resolution mapping influences network behavior.



Because IPv6 networks are long and difficult to remember, DNS is even more important for IPv6 than for IPv4.

The hosts file serves the function of translating human-friendly host names into IPv6 addresses that identify and locate a host in an IPv6 network. In some operating systems, the hosts file content is preferred over other methods, such as the DNS. Unlike the DNS, the hosts file is under the direct control of the local computer administrator.

For a Windows operating system, the file is located at `c:\windows\system32\drivers\etc\hosts`. Other operating systems may have the hosts file in a different location, use a different file, or may not have it at all. Open the hosts file in a text editor such as Notepad.

Name Resolution Issues in IPv6 (Cont.)

Use name resolution mapping with ping command.

```
C:\Windows\system32>ping Server6
Pinging Server [2001:db8:172:16::100] with 32 bytes of data:
Reply from 2001:db8:172:16::100: time=54ms
Reply from 2001:db8:172:16::100: time=46ms
Reply from 2001:db8:172:16::100: time=46ms
Reply from 2001:db8:172:16::100: time=45ms
Ping statistics for 2001:db8:172:16::100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 45ms, Maximum = 54ms, Average = 47ms
```

- Verify connectivity of the server using the ping command and the host name as the destination.

© 2013 Cisco Systems, Inc.

To verify static name resolution, verify connectivity to the server using host name Server6. The ping should be successful.

ACL Issues in IPv6

This topic describes how ACL misconfiguration can cause malfunction of an IPv6 network, and describes the commands that you can use to troubleshoot ACL issues.

ACL Issues in IPv6

Telnet to Server is not working—investigate ACLs applied on the router.

Branch
Gi 0/1

HQ

Server
2001:db8:172:16::100

PC1
2001:db8:101:1::100

```
Branch#show ipv6 access-list  
IPv6 access list Outbound  
  permit icmp any any (44 matches) sequence 10
```

- Display IPv6 ACLs configured on the router.

© 2013 Cisco Systems, Inc.

You can use the **show ipv6 access-list** command to verify whether there are any IPv6 ACLs configured on a router. In the example, there is an ACL named Outbound configured on the router.

ACL Issues in IPv6 (Cont.)

PC is not able to telnet to the Server. Is there an ACL on the Gi 0/1?

Branch
Gi 0/1

HQ

Server
2001:db8:172:16::100

PC1
2001:db8:101:1::100

```
Branch#show ipv6 interface GigabitEthernet0/1 | include access list  
Outbound access list Outbound
```

- Display placement of the ACL on the interface.

© 2013 Cisco Systems, Inc.

Use the **show ipv6 interface** command to verify if an ACL is attached to an interface.

In the example, there is an ACL named Outbound configured on the router. The ACL is applied to the GigabitEthernet0/1 interface in the outbound direction.

ACL Issues in IPv6 (Cont.)

```
Branch(config)#ipv6 access-list extended Outbound
Branch(config-ext-nacl)#permit tcp any any eq 23
```

- Add ACL entry to allow Telnet.

© 2013 Cisco Systems, Inc.

In the example, there is an ACL named Outbound configured on the router. The ACL is applied to the GigabitEthernet0/1 interface in the outbound direction. The ACL permits only ICMP protocol, which is why ping will work. In order to allow Telnet from PC1 to the server, you need to add an entry in the Outbound ACL to allow protocol TCP and port 23 for Telnet.

ACL Issues in IPv6 (Cont.)

```
Branch#show ipv6 access-list
IPv6 access list Outbound
  permit icmp any any (44 matches) sequence 10
  permit tcp any any eq telnet (7 matches) sequence 20
```

- Display corrected ACLs configured on the router.

© 2013 Cisco Systems, Inc.

After correcting the ACL, a Telnet connection from PC1 to the server will be successful.

Summary

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

Summary

- For troubleshooting end-to-end IPv6 connectivity, you can use the same structured approach as for IPv4.
- Use the **ping**, **tracert**, and **telnet** utilities to verify end-to-end IPv6 connectivity.
- Use the **show ipv6 route** command to verify the current IPv6 path on a router.
- The IPv6 gateway on a PC should be set using stateless autoconfiguration or manually.
- Each host should have a DNS server that is configured. The server can be accessed using either IPv4 or IPv6.
- Use the **show ipv6 access-list** and **show ipv6 interfaces** commands to verify whether there are any IPv6 ACLs that are configured to deny traffic.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Module Summary

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

Module Summary

- Physical connectivity must be verified.
- The current path must be equal to the desired path.
- End devices must be configured with the correct gateway.
- End devices must be configured with the correct name resolution entry.
- It is important to ensure that an ACL is not blocking traffic.
- End-to-end connectivity must be verified.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Questions

Use the questions here to review what you learned in this module.

- Which command would you use to check whether there are any input or output errors on a GigabitEthernet0/0 interface? (Source: Troubleshooting IPv4 Network Connectivity)
 - show ip route Gi0/0**
 - show ip interfaces Gi0/0**
 - show interfaces Gi0/0**
 - show mac-address-table**
- Connect the routing table entries to their description. (Source: Troubleshooting IPv6 Network Connectivity)

| | | |
|-----------------------|--------------------------|---|
| A. local-host route | <input type="checkbox"/> | entered manually by a system administrator |
| B. directly connected | <input type="checkbox"/> | statically or dynamically learned and used when no explicit route to network is known |
| C. dynamic routing | <input type="checkbox"/> | learned by exchange of routing information |
| D. default route | <input type="checkbox"/> | local IP address on the router interface |
| E. static routing | <input type="checkbox"/> | router attaches to this network |
- Which command would you use to verify that there is an ACL applied to an interface? (Source: Troubleshooting IPv4 Network Connectivity)
 - show access lists GigabitEthernet 0/1**
 - show access lists**
 - show ip interface GigabitEthernet 0/1**
 - show ip access list**

4. Which command would you use to identify the current IPv6 path on a router? (Source: Troubleshooting IPv6 Network Connectivity)
- A. **show ip route**
 - B. **show ip interfaces brief**
 - C. **show arp**
 - D. **show ipv6 route**

Do Not Duplicate.
Post beta, not for release.

Answer Key

1. C
2.

| | | |
|----|--------------------|---|
| A. | directly connected | router attaches to this network |
| B. | local-host route | local IP address on the router interface |
| C. | static routing | entered manually by a system administrator |
| D. | dynamic routing | learned by exchange of routing information |
| E. | default route | statically or dynamically learned and used when no explicit route to network is known |
3. C
4. D

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Implementing an EIGRP-Based Solution

Describe the components and metrics of EIGRP, how EIGRP selects routes, the creation of an implementation plan, and basic EIGRP configuration.

Objectives

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

- Explain dynamic routing protocols, EIGRP, and the basic configuration of EIGRP
- Troubleshoot EIGRP
- Describe the implementation of EIGRP for IPv6

Do Not Duplicate.
Post beta, not for release.

Implementing EIGRP

EIGRP is an advanced distance vector routing protocol that was developed by Cisco. EIGRP is suited for many different topologies and media. In a well-designed network, EIGRP scales well and provides quick convergence times with minimal overhead. EIGRP is a popular choice for a routing protocol on Cisco devices. This lesson describes how to configure and monitor EIGRP.

Objectives

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

- Describe dynamic routing protocols
- Describe the purpose of administrative distance
- Describe EIGRP features
- Explain how EIGRP chooses the best path
- Describe the EIGRP composite metric
- Configure EIGRP
- Verify EIGRP configuration
- Explain load balancing with EIGRP

Dynamic Routing Protocols

This topic describes dynamic routing protocols.

Dynamic Routing Protocols

A dynamic routing protocol has these purposes:

- The discovery of remote networks
- Maintaining up-to-date routing information
- Choosing the best path to destination networks
- The ability to find a new best path if the current path is no longer available

© 2013 Cisco Systems, Inc.

A routing protocol is a set of processes, algorithms, and messages that is used to exchange routing information and populate the routing table with the choice of best paths for the routing protocol. Routing protocols are a set of rules by which routers dynamically share their routing information. As routers become aware of changes to the networks for which they act as the gateway or changes to links between routers, this information is passed onto other routers. When a router receives information about new or changed routes, it updates its own routing table and, in turn, passes the information to other routers. In this way, all routers have accurate routing tables that are updated dynamically and can learn about routes to remote networks.

All routing protocols have the same purpose: to learn about remote networks and to quickly adapt whenever there is a change in the topology. The method that a routing protocol uses to accomplish this purpose depends upon the algorithm that it uses and the operational characteristics of this protocol. The operations of a dynamic routing protocol vary, depending on the type of routing protocol and on the routing protocol itself.

Although routing protocols provide routers with up-to-date routing tables, there are costs that put additional demands on the memory and processing power of the router. First, the exchange of route information adds overhead that consumes network bandwidth. This overhead can be a problem, particularly for low-bandwidth links between routers. Second, after the router receives the route information, protocols such as EIGRP and OSPF process it extensively to make routing table entries. This means that routers that use these protocols must have sufficient processing capacity to implement the algorithms of the protocol as well as to perform timely packet routing and forwarding.

Dynamic Routing Protocols (Cont.)

Different protocols behave differently:

- IGP versus EGP
- Distance vector vs. link state
- Classless vs. classful

© 2013 Cisco Systems, Inc.

An AS, otherwise known as a routing domain, is a collection of routers under a common administration, such as an internal company network or an ISP network. Because the Internet is based on the AS concept, the following two types of routing protocols are required:

- IGPs: These routing protocols are used to exchange routing information within an AS. EIGRP, IS-IS, and OSPF are examples of IGPs.
- EGPs: These routing protocols are used to route between autonomous systems. BGP is the EGP of choice in networks today.

Within an AS, most IGP routing can be classified as distance vector or link state:

- **Distance vector**: The distance vector routing approach determines the direction (vector) and distance (such as hops) to any link in the internetwork. Some distance vector protocols periodically send complete routing tables to all of the connected neighbors. In large networks, these routing updates can become enormous, causing significant traffic on the links. Distance vector protocols use routers as signposts along the path to the final destination. The only information that a router knows about a remote network is the distance or metric to reach this network and which path or interface to use to get there. Distance vector routing protocols do not have an actual map of the network topology. EIGRP is an example of a distance vector routing protocol.
- **Link state**: The link-state approach, which uses the SPF algorithm, creates an abstract of the exact topology of the entire internetwork, or at least of the partition in which the router is situated. Using the analogy of signposts, using a link-state routing protocol is like having a complete map of the network topology. The signposts along the way from the source to the destination are not necessary because all link-state routers use an identical "map" of the network. A link-state router uses the link-state information to create a topology map and to select the best path to all destination networks in the topology. The OSPF and IS-IS protocols are examples of link-state routing protocols.

Also, there is classful and classless routing:

- **Classful routing:** Classful routing is a consequence of the fact that subnet masks are not advertised in the routing advertisements that most distance vector routing protocols generate. When a classful routing protocol is used, all subnetworks of the same major network (Class A, B, or C) must use the same subnet mask, which is not necessarily a default major class subnet mask. Routers that are running a classful routing protocol perform automatic route summarization across network boundaries. Classful routing is obsolete in networks today.
- **Classless routing:** Classless routing protocols can be considered second-generation protocols because they are designed to address some of the limitations of the earlier classful routing protocols. A serious limitation in a classful network environment is that the subnet mask is not exchanged during the routing update process, thus requiring the same subnet mask to be used on all subnetworks within the same major network. Another limitation of the classful approach is the need to automatically summarize to the classful network number at all major network boundaries. In the classless environment, the summarization process is controlled manually and can usually be invoked at any bit position within the address. Because subnet routes are propagated throughout the routing domain, manual summarization may be required to keep the size of the routing tables manageable. Classless routing protocols include RIPv2, EIGRP, OSPF, and IS-IS.


Do Not Duplicate
Post beta, not for release

Administrative Distance

Administrative Distance

- Multiple routing protocols and static routes can be used at the same time.
- Routers choose the routing source with the lowest administrative distance.

I need to send a packet from network A to network B. Which route is the best?



© 2013 Cisco Systems, Inc.

Multiple routing protocols and static routes may be used at the same time. If there are several sources for routing information, such as specific routing protocols, static routes, and even directly connected networks, an administrative distance value is used to rate the trustworthiness of each routing information source. Cisco IOS Software uses the administrative distance feature to select the best path when it learns about the same destination network from two or more routing sources.

An administrative distance is an integer from 0 to 255. A routing protocol with a lower administrative distance is more trustworthy than one with a higher administrative distance.

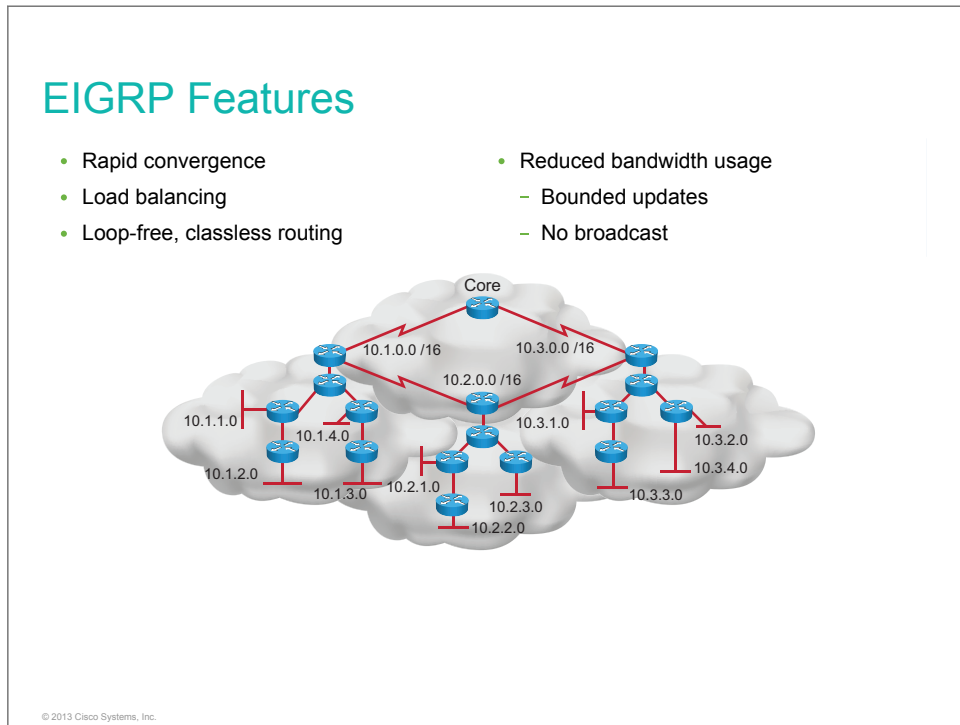
As shown in the figure, the router must deliver a packet from network A to network B. The router must choose between two routes. One is routed by EIGRP, and the other is routed by OSPF. Although the OSPF route seems like the logical choice, given that there are fewer hops to the destination network, the EIGRP route will be identified as more trustworthy and added to the routing table of the router.

The table shows the default administrative distance for selected routing information sources.

| Route Source | Default Distance |
|-------------------------|--|
| Connected interface | 0 |
| Static route address | 1 |
| EIGRP | 90 |
| OSPF | 110 |
| RIP | 120 |
| External EIGRP | 170 |
| Unknown or unbelievable | 255 (will not be used to pass traffic) |

EIGRP Features

This topic describes the features of EIGRP.



EIGRP is a Cisco proprietary routing protocol that combines the advantages of link-state and distance vector routing protocols. EIGRP may act like a link-state routing protocol, because it uses a Hello protocol to discover neighbors and form neighbor relationships, and only partial updates are sent when a change occurs. However, EIGRP is based on the key distance vector routing protocol principle, in which information about the rest of the network is learned from directly connected neighbors. EIGRP is an advanced distance vector routing protocol that includes the following features:

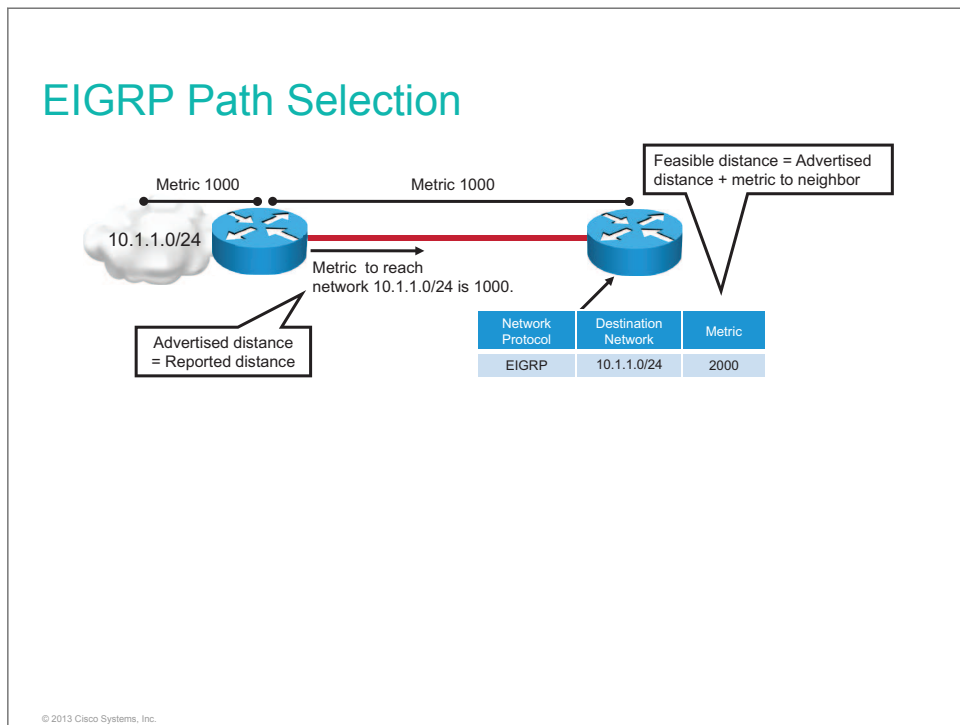
- **Rapid convergence:** EIGRP uses the DUAL to achieve rapid convergence. As the computational engine that runs EIGRP, DUAL resides at the center of the routing protocol, guaranteeing loop-free paths and backup paths throughout the routing domain. A router that uses EIGRP stores all available backup routes for destinations so that it can quickly adapt to alternate routes. If the primary route in the routing table fails, the best backup route is immediately added to the routing table. If no appropriate route or backup route exists in the local routing table, EIGRP queries its neighbors to discover an alternate route.
- **Load balancing:** EIGRP supports unequal metric load balancing as well as equal metric load balancing, which allows administrators to better distribute traffic flow in their networks.
- **Loop-free, classless routing:** Because EIGRP is a classless routing protocol, it advertises a routing mask for each destination network. The routing mask feature enables EIGRP to support discontinuous subnetworks and VLSMs.

- **Reduced bandwidth usage:** EIGRP uses the terms “partial” and “bounded” when referring to its updates. EIGRP does not make periodic updates. The term “partial” means that the update only includes information about the route changes. EIGRP sends these incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. The term “bounded” refers to the propagation of partial updates that are sent only to those routers that the changes affect. By sending only the routing information that is needed and only to those routers that need it, EIGRP minimizes the bandwidth that is required to send EIGRP updates. EIGRP uses multicast and unicast rather than broadcast. Multicast EIGRP packets use the reserved multicast address of 224.0.0.10. As a result, end stations are unaffected by routing updates and requests for topology information.

Do Not Duplicate.
Post beta, not for release.

EIGRP Path Selection

This topic describes EIGRP path calculation.



Each EIGRP router maintains a neighbor table. This table includes a list of directly connected EIGRP routers that have an adjacency with this router. Neighbor relationships are used to track the status of these neighbors. EIGRP uses a lightweight Hello protocol to monitor the connection status with its neighbors.

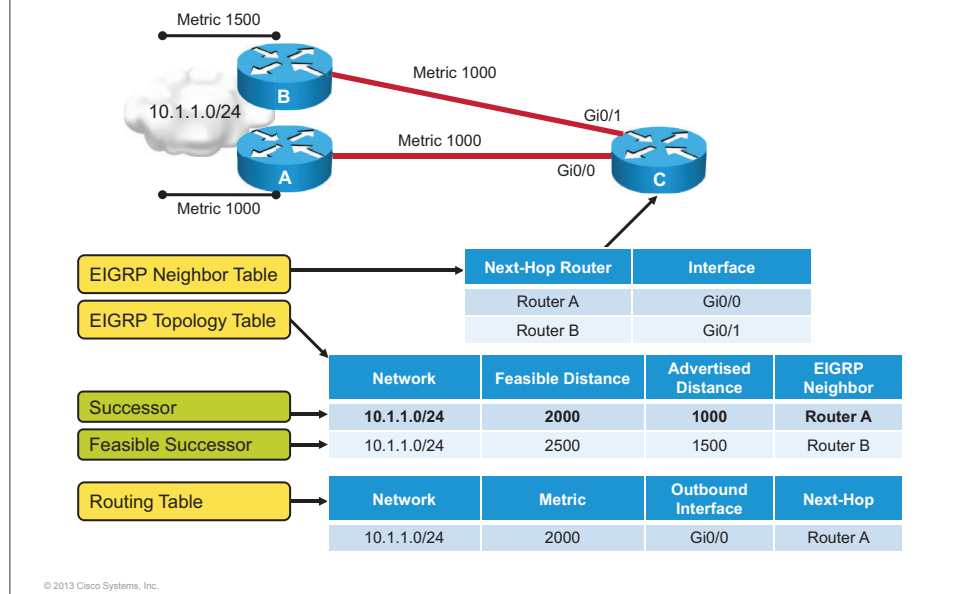
Each EIGRP router maintains a topology table for each routed protocol configuration. The topology table includes route entries for every destination that the router learns from its directly connected EIGRP neighbors. EIGRP chooses the best routes to a destination from the topology table and places these routes in the routing table.

To determine the best route (successor) and any backup routes (feasible successors) to a destination, EIGRP uses the following two parameters:

- **AD:** The EIGRP metric for an EIGRP neighbor to reach a particular network, also sometimes referred to as the *reported distance*.
- **FD:** The AD for a particular network that is learned from an EIGRP neighbor plus the EIGRP metric to reach this neighbor. This sum provides an end-to-end metric from the router to the remote network.

A router compares all **FDs** to reach a specific network and then selects the lowest FD and places it in the routing table. The FD for the chosen route becomes the EIGRP routing metric to reach this network in the routing table.

EIGRP Path Selection (Cont.)



The EIGRP topology database contains all of the routes that are known to each EIGRP neighbor. As shown in the example, routers A and B sent their routing tables to router C, whose table is displayed. Both routers A and B have routes to network 10.1.1.0/24, as well as to other networks that are not shown.

Router C has two entries to reach 10.1.1.0/24 in its topology table. The EIGRP metric for router C to reach both routers A and B is 1000. Add this metric (1000) to the respective AD for each router, and the results represent the FDs that router C must travel to reach network 10.1.1.0/24.

Router C chooses the least FD (2000) and installs it in the IP routing table as the best route to reach 10.1.1.0/24. The route with the least FD that is installed in the routing table is called the “successor route.”

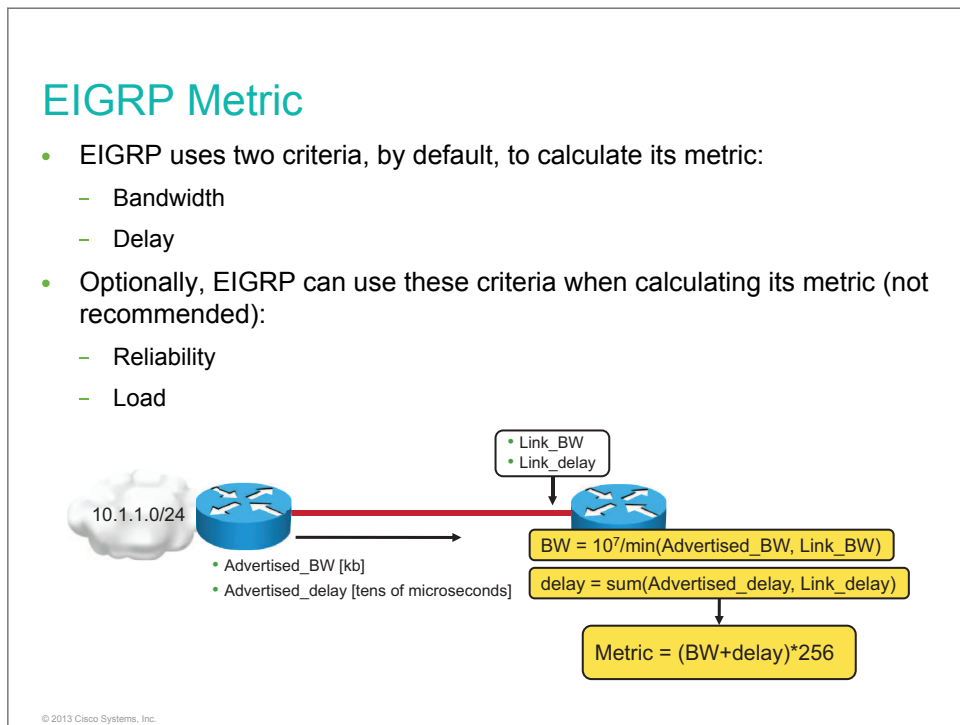
Router C then chooses a backup route to the successor that is called a “feasible successor route,” if one or more feasible successor routes exist. To become a feasible successor, a route must satisfy this feasibility condition: A next-hop router must have an AD that is less than the FD of the current successor route (therefore, the route is tagged as a feasible successor). This rule is used to ensure that the network is loop-free.

If the route via the successor becomes invalid, possibly because of a topology change, or if a neighbor changes the metric, DUAL checks for feasible successors to the destination route. If a feasible successor is found, DUAL uses it, avoiding the need to recompute the route. A route will change from a passive state to an active state if no feasible successor exists, and a recomputation must occur to determine the new successor.

Note In this example, values for the EIGRP metric and for FDs and ADs are optimized for explanation purposes. The real metric values are much larger.

EIGRP Metric

This topic describes the EIGRP composite metric.



The EIGRP metric can be based on several criteria, but EIGRP uses only two of these, by default:

- **Bandwidth:** The smallest bandwidth of all outgoing interfaces between the source and destination, in kilobits.
- **Delay:** The cumulative (sum) of all interface delay along the path, in tens of microseconds.

These criteria can be used, but are not recommended because they typically result in frequent recalculation of the topology table:

- **Reliability:** This value represents the worst reliability between the source and destination, which is based on keepalives.
- **Load:** This value represents the worst load on a link between the source and destination, which is computed based on the packet rate and the configured bandwidth of the interface.

The composite metric formula is used by EIGRP to calculate metric value. The formula consists of values K1 through K5, which are known as EIGRP metric weights. By default, K1 and K3 are set to 1, and K2, K4, and K5 are set to 0. The result is that only the bandwidth and delay values are used in the computation of the default composite metric. The metric calculation method (K values) and the EIGRP AS number must match between EIGRP neighbors.

Although an MTU is exchanged in EIGRP packets between neighbor routers, the MTU is not factored into the EIGRP metric calculation.

EIGRP uses scaled values to determine the total metric: $256 * ([K1 * \text{bandwidth}] + [K2 * \text{bandwidth}] / [256 - \text{Load}] + K3 * \text{Delay}) * (K5 / [\text{Reliability} + K4])$, where if $K5 = 0$, the $(K5 / [\text{Reliability} + K4])$ part is not used (that is, equals 1). Using the default K values, the metric calculation simplifies to $256 * (\text{bandwidth} + \text{delay})$.

EIGRP Metric (Cont.)

```
HQ#show interfaces serial 0/0/0
Serial0/0/0 is up, line protocol is down
  Hardware is GT96K Serial
  Description: Link to Branch
  MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
<output omitted>
```

- Verifies the EIGRP metric values on the Serial 0/0/0 interface of router HQ

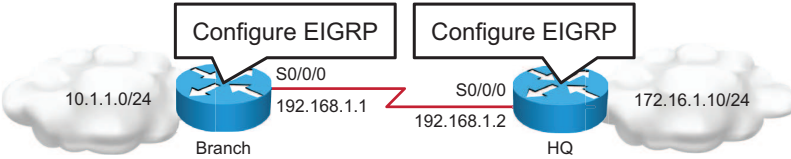
© 2013 Cisco Systems, Inc.

By using the **show interface** command, you can examine the actual values that are used for bandwidth, delay, reliability, and load in the computation of the routing metric. The output in the figure shows the values that are used in the composite metric for the Serial 0/0/0 interface.

EIGRP Configuration

This topic describes how to configure basic EIGRP.

EIGRP Configuration



```
Branch(config)#router eigrp 100
Branch(config-router)#network 10.1.1.0
Branch(config-router)#network 192.168.1.0
```

- Configuration of EIGRP on the Branch router

```
HQ(config)#router eigrp 100
HQ(config-router)#network 172.16.1.0 0.0.0.255
HQ(config-router)#network 192.168.1.0 0.0.0.255
```

- Configuration of EIGRP on the HQ router

© 2013 Cisco Systems, Inc.

| Command | Description |
|---|--|
| <code>router eigrp as_number</code> | Enables the EIGRP routing process for the AS that is specified |
| <code>network network_number wildcard_mask</code> | Associates the network with the EIGRP routing process. Use of the wildcard mask is optional. |

The **router eigrp** global configuration command enables EIGRP.

Use the **router eigrp** and **network** commands to create an EIGRP routing process. Note that EIGRP requires an AS number. The AS parameter is a number between 1 and 65,535 that is chosen by the network administrator.

The **network** command is used in the router configuration mode.

Note The AS number that EIGRP refers to in the parameter can be assigned any 16-bit value. As opposed to OSPF, the AS number in EIGRP must match on all routers that are involved in the same EIGRP process.

The **network** command in EIGRP has the same function as in other IGP routing protocols:

- The **network** command defines a major network number to which the router is directly connected. Any interface on this router that matches the network address in the **network** command will be enabled to send and receive EIGRP updates. The EIGRP routing process looks for interfaces that have an IP address that belongs to the networks that are specified with the **network** command. The EIGRP process begins on these interfaces.
- This network (or subnet) will be included in EIGRP routing updates.

To configure EIGRP to advertise specific subnets only, use the *wildcard-mask* option with the **network** command. You can think of the wildcard mask as the inverse of a subnet mask; for example 255.255.255.0 is 0.0.0.255.


For more details about the **router eigrp** command, check the *Cisco IOS IP Routing: EIGRP Command Reference* at this URL: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_11.html.

For more details about the **network** command, check the *Cisco IOS IP Routing: Protocol-Independent Command Reference* at this URL: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_10.html.

Do Not Duplicate.
Post beta, not for release.

Verification of EIGRP Configuration

Verification of EIGRP Configuration



```
Branch#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H  Address          Interface      Hold    Uptime    SRTT    RTO    Q  Seq
   (sec)              (ms)          (ms)    (sec)    (ms)    (sec)  Cnt Num
0  192.168.1.2      S0/0/0        11     00:17:22  1596    5000   0   3
```

- Verification of EIGRP neighbors on the Branch router. The Branch router has one neighbor. Branch is receiving hello packets from the peer through its Serial 0/0/0 interface.

© 2013 Cisco Systems, Inc.

Verification of EIGRP Configuration (Cont.)



```
HQ#show ip eigrp neighbors
EIGRP-IPv4 Neighbors for AS(100)
H  Address          Interface      Hold    Uptime    SRTT    RTO    Q  Seq
   (sec)              (ms)          (ms)    (sec)    (ms)    (sec)  Cnt Num
0  192.168.1.1      Se0/0/0       13     01:21:35  254     1524   0   4
```


- Verification of EIGRP neighbors on the HQ router. The HQ router has one neighbor. HQ is receiving hello packets from the peer through its Serial 0/0/0 interface.

© 2013 Cisco Systems, Inc.

Use the **show ip eigrp neighbors** command to display the neighbors that EIGRP discovered and to determine when neighbors become active and inactive. The command is also useful for debugging transport problems.

| Field | Description |
|------------|--|
| AS(100) | Process number that is specified with the router command |
| Address | IP address of the EIGRP peer |
| Interface | Interface on which the router is receiving hello packets from the peer |
| Hold (sec) | Length of time (in seconds) that Cisco IOS Software waits to hear from the peer before declaring it down. If the peer is using the default hold time, this number is less than 15. If the peer configures a nondefault hold time, the nondefault hold time is displayed. |
| Uptime | Elapsed time (in hours:minutes:seconds) since the local router first heard from this neighbor |
| Q Cnt | Number of EIGRP packets (update, query, and reply) that the software is waiting to send |
| Seq Num | Sequence number of the last update, query, or reply packet that was received from this neighbor |

Verification of EIGRP Configuration (Cont.)



```
Branch#show ip eigrp interfaces
EIGRP-IPv4 Interfaces for AS(100)
      Xmit Queue  Mean   Pacing Time  Multicast  Pending
Interface  Peers  Un/Reliable  SRTT  Un/Reliable  Flow Timer  Routes
Gi0/0      0      0/0         0      0/1         0           0
S0/0/0     1      0/0        1596   0/1         7984        0
```

- Displays information about interfaces that are configured for EIGRP on the Branch router


© 2013 Cisco Systems, Inc.

Use the **show ip eigrp interfaces** command to determine on which interfaces EIGRP is active and to learn information about EIGRP that relates to those interfaces. If you specify an interface (for example, **show ip eigrp interfaces Fa0/0**), only this interface is displayed. Otherwise, all interfaces on which EIGRP is running are displayed. If you specify AS (for example, **show ip eigrp interfaces 100**), only the routing process for the specified AS is displayed. Otherwise, all EIGRP processes are displayed.

| Field | Description |
|--------------------------------|---|
| Interface | Interface over which EIGRP is configured |
| Peers | Number of directly connected EIGRP neighbors on the interface |
| Xmit Queue Unreliable/Reliable | Number of packets remaining in the Unreliable and Reliable queues |
| Mean SRTT | Average <u>SRTT</u> interval (in milliseconds) for all neighbors on the interface |

| Field | Description |
|---------------------------------|--|
| Pacing Time Unreliable/Reliable | Number of milliseconds to wait after transmitting unreliable and reliable packets |
| Multicast Flow Timer | Number of milliseconds to wait for acknowledgment of a multicast packet by all neighbors before transmitting the next multicast packet |
| Pending Routes | Number of routes in the packets in the transmit queue waiting to be sent |

Verification of EIGRP Configuration (Cont.)




```
Branch#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i -IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
+ - replicated route, % - next hop override
```

- Displays routes on the Branch router. Routes marked with D are those acquired through EIGRP.

© 2013 Cisco Systems, Inc.

Verification of EIGRP Configuration (Cont.)



```
Gateway of last resort is not set
 10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C    10.1.1.0/24 is directly connected, GigabitEthernet0/0
L    10.1.1.1/32 is directly connected, GigabitEthernet0/0
 172.16.0.0/24 is subnetted, 1 subnets
D    172.16.1.0 [90/156160] via 192.168.1.2, 01:12:45, Serial0/0/0
 192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Serial0/0/0
L    192.168.1.1/32 is directly connected, Serial0/0/0
```

- Displays routes on the Branch router. Routes marked with D are those acquired through EIGRP.

© 2013 Cisco Systems, Inc.

The **show ip route** command displays the current entries in the routing table. EIGRP has a default administrative distance of 90 for internal routes and 170 for routes that are imported from an external source, such as default routes. When compared to other IGP's, EIGRP is preferred by Cisco IOS Software because it has the lowest administrative distance.

Verification of EIGRP Configuration (Cont.)

```

Branch#show ip eigrp topology
EIGRP-IPv4 Topology Table for AS(100)/ID(10.1.1.1)
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 28160
   via Connected, Serial0/0/0
P 172.16.1.0/24, 1 successors, FD is 156160
   via 192.168.1.2 (156160/128256), Serial0/0/0
P 10.1.1.0/24, 1 successors, FD is 28160
   via Connected, GigabitEthernet0/0
    
```

- Displays entries in the EIGRP topology table. All routes throughout the EIGRP autonomous system are displayed here.

© 2013 Cisco Systems, Inc.

The **show ip eigrp topology** command displays the EIGRP topology table, the active or passive state of routes, the number of successors, and the FD to the destination. Use the **show ip eigrp topology all-links** command to display all paths, even those that are not feasible.

| Field | Description |
|------------------|--|
| Codes | The state of this topology table entry. Passive and active refer to the EIGRP state with respect to this destination; update, query, and reply refer to the type of packet that is being sent. |
| P – passive | Indicates that no EIGRP computations are being performed for this destination |
| A – active | Indicates that EIGRP computations are being performed for this destination |
| U – update | Indicates that an update packet was sent to this destination |
| Q – query | Indicates that a query packet was sent to this destination |
| R – reply | Indicates that a reply packet was sent to this destination |
| R – reply status | A flag that is set after the software has sent a query and is waiting for a reply |
| 10.1.1.0 | Destination IP network number |
| /24 | Destination subnet mask |
| successors | Number of successors. This number corresponds to the number of next hops in the IP routing table. If "successors" is capitalized, then the route or next hop is in a transition state. |

| Field | Description |
|-------------------|---|
| FD | The FD is the best metric to reach the destination or the best metric that was known when the route went active. This value is used in the feasibility condition check. If the reported distance of the router (the metric after the slash) is less than the FD, the feasibility condition is met and this path is a feasible successor. After the software determines that it has a feasible successor, it does not need to send a query for this destination. |
| replies | The number of replies that are still outstanding (have not been received) with respect to this destination. This information appears only when the destination is in active state. |
| state | The exact EIGRP state that this destination is in. It can be the number 0, 1, 2, or 3. This information appears only when the destination is in the active state. |
| via | The IP address of the peer that told the software about this destination. The first n of these entries, where n is the number of successors, are the current successors. The remaining entries on the list are feasible successors. |
| (156,160/128,256) | The first number is the EIGRP metric that represents the cost, or FD, to the destination. The second number is the EIGRP metric that this peer advertised. |
| Serial0/0/0 | The interface from which this information was learned. |

Do Not Duplicate
Post beta, not for release

Load Balancing with EIGRP

This topic describes how EIGRP utilizes load balancing.

Load Balancing with EIGRP

EIGRP knows two types of load balancing:

- Equal-cost load balancing:
 - By default, up to four routes with a metric equal to the minimum metric are installed in the routing table.
 - The routing table can have up to 16 entries for the same destination.
- Unequal-cost load balancing:
 - By default, it is *not* turned on.
 - Load balancing can be performed through paths that are 128 times worse than the route with the lowest FD.

The diagram shows two scenarios of load balancing between two routers. In the first scenario, labeled 'Equal-Cost Load balancing', two blue routers are connected by two parallel paths, each with a metric of 100. The text 'Equal-Cost Load balancing' is centered between the routers. In the second scenario, labeled 'Unequal-Cost Load balancing', two blue routers are connected by two parallel paths. The top path has a metric of 100, and the bottom path has a metric of 25. The text 'Unequal-Cost Load balancing' is centered between the routers. A small copyright notice '© 2013 Cisco Systems, Inc.' is visible at the bottom left of the diagram area.

Equal-cost load balancing is the ability of a router to distribute traffic over all of its network ports that are the same metric from the destination address. Load balancing increases the use of network segments and increases effective network bandwidth.

For IP, Cisco IOS Software applies load balancing across up to four equal-cost paths, by default. With the **maximum-paths** router configuration command, up to 32 equal-cost routes can be kept in the routing table. If you set the value to 1, you disable load balancing. When a packet is process-switched, load balancing over equal-cost paths occurs on a per-packet basis. When packets are fast-switched, load balancing over equal-cost paths occurs on a per-destination basis.

For more details about the **maximum-paths** command, check the *Cisco IOS IP Routing: EIGRP Command Reference* at this URL: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_09.html.

EIGRP can also balance traffic across multiple routes that have different metrics. This type of balancing is called unequal-cost load balancing.

Summary

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

Summary

- EIGRP is a classless, advanced distance vector routing protocol that runs the DUAL algorithm.
- The composite metric formula is used by EIGRP to calculate metric value; by default, it uses only bandwidth and delay.
- EIGRP is configured on a router through the **router eigrp** and **network** commands.
- There are three tables:
 - The EIGRP neighbor table lists directly connected routers that are running EIGRP.
 - The EIGRP topology table lists all routes that are learned from each EIGRP neighbor.
 - The routing table lists the best routes from the EIGRP topology table and the other routing processes.

© 2013 Cisco Systems, Inc.

Troubleshooting EIGRP

EIGRP is one of the routing protocols that is commonly used in large enterprise networks. Troubleshooting problems related to the exchange of routing information is one of the most essential skills for a network engineer who is involved in the implementation and maintenance of large routed enterprise networks that use EIGRP as the IGP.

This lesson provides the troubleshooting flow and the Cisco IOS commands that you can use to gather information from the EIGRP data structures and routing processes.

Objectives

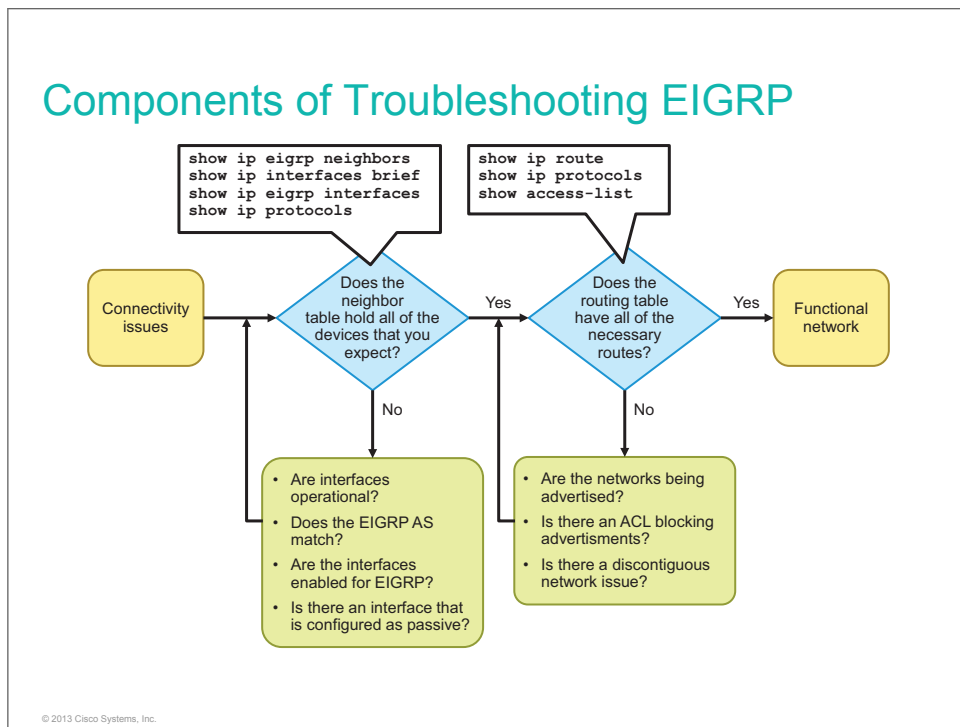
Upon completing this lesson, you will be able to meet these objectives:

- Describe the basic components of troubleshooting a network that is running EIGRP
- Identify and resolve EIGRP neighbor relationship issues
- Identify and resolve EIGRP routing table issues

Components of Troubleshooting EIGRP

This topic describes how to troubleshoot EIGRP.

Components of Troubleshooting EIGRP



After configuring EIGRP, you would first test connectivity to the remote network. If the ping fails, you should check if the router has EIGRP neighbors. Neighbor adjacency might not be up for a number of reasons, such as the following:

- The interface between the devices is down.
- The two routers have mismatching EIGRP autonomous systems.
- Proper interfaces are not enabled for the EIGRP process.
- An interface is configured as passive.

Aside from these issues, there are a number of other, more advanced issues that can cause neighbor relationships to not be formed. Two examples are misconfigured EIGRP authentication or mismatched K values, based on which EIGRP calculates its metric.

If the EIGRP neighbor relationship is up between the two routers, there may be a routing problem. These are some issues that may cause a connectivity problem for EIGRP:

- Proper networks are not being advertised on remote routers.
- An access list is blocking advertisements of remote networks.
- Automatic summary is causing confusion in your discontinuous network.

Troubleshooting EIGRP Neighbor Issues

This topic describes how to troubleshoot EIGRP neighbor issues.

Troubleshooting EIGRP Neighbor Issues

Branch#**show ip route eigrp**
Branch#

- Investigates if there are EIGRP routes in the routing table. There are none in this example.

© 2013 Cisco Systems, Inc.

Troubleshooting EIGRP Neighbor Issues (Cont.)

Branch#**show ip interface brief**

| Interface | IP-Address | OK? | Method | Status | Protocol |
|--------------------|-------------|-----|--------|--------|----------|
| GigabitEthernet0/0 | 10.1.1.1 | YES | manual | up | up |
| Serial0/0/0 | 192.168.1.1 | YES | manual | up | up |

<output omitted>

- Verifies that the protocol and status of the link between neighboring routers is up

© 2013 Cisco Systems, Inc.

The first thing you would normally do if there is no connectivity to the remote network is to issue the **show ip route** command. If there is no path to the remote network, you might have a neighbor issue.

A prerequisite for the neighbor relationship to form between the Branch and the HQ router is OSI Layer 3 connectivity. By investigating the output of the **show ip interface brief** command, you can verify that the status and protocol are both up for the Serial0/0/0 interface that is connected to the Branch router. A ping from Branch to HQ will confirm IP connectivity between the devices.

If the ping is not successful, check the cabling and verify that the interfaces on connected devices are on a common subnet. A log message that states that EIGRP neighbors are "not on common subnet" indicates that there is an incorrect IP address on one of the two EIGRP neighbor interfaces.

Troubleshooting EIGRP Neighbor Issues (Cont.)

```
Branch#show ip protocols
Routing Protocol is "eigrp 1"
<output omitted>
```

```
HQ#show ip protocols
Routing Protocol is "eigrp 2"
<output omitted>
```

- Because the EIGRP autonomous systems do not match, routers will not form a neighbor adjacency.


© 2013 Cisco Systems, Inc.

The command that starts the EIGRP process is followed by a number that is the AS number, **router eigrp AS**. While the process ID has local significance with OSPF, the AS number in EIGRP is important because it must match between the neighbors.

The "Routing for Networks" section of the **show ip protocols** command output indicates which networks have been configured. Any interfaces in those networks participate in EIGRP.

For more details about **show ip protocols** and related commands, check the *Cisco IOS IP Routing: Protocol-Independent Command Reference* at this URL: http://www.cisco.com/en/US/docs/ios/iproute_pi/command/reference/iri_book.html.


Troubleshooting EIGRP Neighbor Issues (Cont.)



```
Branch#show ip eigrp interfaces Serial 0/0/0
EIGRP-IPv4 Interfaces for AS(1)
Interface      Peers    Xmit Queue  Mean    Pacing Time  Multicast  Pending
              Un/Reliable SRTT        Un/Reliable Flow Timer  Routes
S0/0/0         0        0/0         0       0/0          0          0
```

© 2013 Cisco Systems, Inc.

Troubleshooting EIGRP Neighbor Issues (Cont.)



```
HQ#show ip eigrp interfaces Serial 0/0/0
EIGRP-IPv4 Interfaces for AS(1)
Interface      Peers    Xmit Queue  Mean    Pacing Time  Multicast  Pending
              Un/Reliable SRTT        Un/Reliable Flow Timer  Routes
```

- If serial interfaces on both routers are not enabled for the EIGRP process, a neighbor adjacency will not be formed. In this example, HQ does not have Serial 0/0/0 enabled for EIGRP and therefore routers are not becoming EIGRP neighbors.

© 2013 Cisco Systems, Inc.

The network command that is configured under the EIGRP routing process indicates which router interfaces will participate in EIGRP. The **show ip eigrp interfaces interface** command will show you which interfaces are enabled for EIGRP. If connected interfaces are not enabled for EIGRP, then neighbors will not form an adjacency. If an interface is not on the list, this means the router is not using EIGRP through that interface.

| Parameter | Description |
|-------------------------------------|--|
| AS(100) | Autonomous system number that is specified with the router command |
| Interface | Interface over which the EIGRP is configured |
| Peers | Number of directly connected EIGRP neighbors on the interface |
| Xmit Queue Unreliable and Reliable | Number of packets remaining in the Unreliable and Reliable queues |
| Mean SRTT | Average <u>SRTT</u> interval (in milliseconds) for all neighbors on the interface |
| Pacing Time Unreliable and Reliable | Number of milliseconds to wait after transmitting unreliable and reliable packets |
| Multicast Flow Timer | Number of milliseconds to wait for acknowledgment of a multicast packet by all neighbors before transmitting the next multicast packet |
| Pending Routes | Number of routes in the packets in the transmit queue that are waiting to be sent |

Troubleshooting EIGRP Neighbor Issues (Cont.)

```

HQ#show ip protocols
<output omitted>
Routing Protocol is "eigrp 1"
<output omitted>
Routing for Networks:
 172.16.0.0
 192.168.1.0
Passive Interface(s):
 Serial0/0/0
<output omitted>

```

- Because the HQ interface S0/0/0 is configured as a passive neighbor, an adjacency is not formed.

© 2013 Cisco Systems, Inc.

With EIGRP running on a network, the **passive-interface** command stops both outgoing and incoming routing updates because the effect of the command causes the router to stop sending and receiving hello packets over an interface. For this reason, routers will not become neighbors.

To verify whether any interface on a router is configured as passive, use the **show ip protocols** command in privileged mode.

For an example of when you may want to configure an interface as passive towards the LAN, consider a situation in which you want to advertise LAN networks, but you do not want to risk taking hello packets into the LAN.

To configure an interface as a passive interface in EIGRP, you will use the **passive-interface** *interface* command in the EIGRP router configuration mode. To disable the interface as passive, use the **no passive-interface** *interface* command.

Do Not Duplicate:
Post beta, not for release.

Troubleshooting EIGRP Routing Table Issues

This topic describes how to troubleshoot EIGRP routing table issues.

Troubleshooting EIGRP Routing Table Issues

I cannot reach network 172.16.1.0/24.
Are all networks being advertised?

```
graph LR
    Branch[Branch] --- S0_0_0[192.168.1.1 S0/0/0]
    HQ[HQ] --- S0_0_0_2[192.168.1.2 S0/0/0]
    Branch --- N1[10.1.1.0/24]
    HQ --- N2[172.16.1.0/24]
```

```
HQ#show ip protocols
Routing Protocol is "eigrp 1"
<output omitted>
  Routing for Networks:
    192.168.1.0
    10.0.0.0
<output omitted>
```

- The HQ router is missing the network statement for the 172.16.1.0 network.

© 2013 Cisco Systems, Inc.

Branch and HQ have their neighbor adjacency set up, but a ping test from the Branch router to a host in the 172.16.1.0/24 network is not successful. Checking the routing table of the Branch router leads you to the conclusion that there is a route missing to the destination network of 172.16.1.0/24.

Using the **show ip protocols** command on the HQ router will show you that the network 172.16.1.0/24 is not being advertised to EIGRP neighbors:

```
HQ#sh ip protocols | begin Routing for Networks
*** IP Routing is NSF aware ***
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    192.168.1.0 0.0.0.0.255
    10.0.0.0 0.0.0.0.255
<output omitted>
```

Connectivity will be restored if the HQ router is configured to advertise the 172.16.1.0/24 network:

```
HQ(config)#router eigrp 1
HQ(config-router)#network 172.16.1.0
```

You can now again investigate the routing table of the Branch router and confirm that it has a route to the 172.16.1.0/24 network:

```
Branch#show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
```

i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
+ - replicated route, % - next hop override

Gateway of last resort is not set

10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks

C 10.1.1.0/24 is directly connected, GigabitEthernet0/0

L 10.1.1.1/32 is directly connected, GigabitEthernet0/0

172.16.0.0/24 is subnetted, 1 subnets

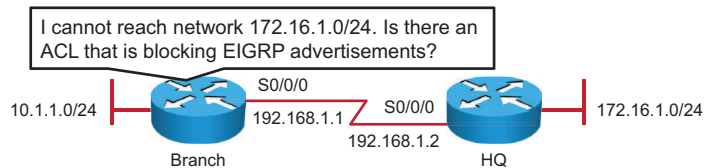
D 172.16.1.0 [90/156160] via 192.168.1.2, 00:28:12, Serial0/0/0

192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks

C 192.168.1.0/24 is directly connected, Serial0/0/0

L 192.168.1.1/32 is directly connected, Serial0/0/0

Troubleshooting EIGRP Routing Table Issues (Cont.)




```
HQ#show ip protocols
<output omitted>
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
<output omitted>
```

- Check if any access lists are applied to the EIGRP network advertisements.

© 2013 Cisco Systems, Inc.

Troubleshooting EIGRP Routing Table Issues (Cont.)



```
Branch#show ip protocols
<output omitted>
Routing Protocol is "eigrp 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
<output omitted>
```


- Check if any access lists are applied to the EIGRP network advertisements.

© 2013 Cisco Systems, Inc.

Access lists provide filtering for different protocols, and these access lists may affect the exchange of the routing protocol messages that are causing routes to be missing from the routing table. The **show ip protocols** command shows whether there are any filter lists that are applied to EIGRP.

Troubleshooting EIGRP Routing Table Issues (Cont.)

EIGRP can be configured to perform automatic summarization on classful boundaries that are causing issues with discontinuous networks.



© 2013 Cisco Systems, Inc.

EIGRP can be configured to automatically summarize routes at classful boundaries. If you have discontinuous networks, automatic summarization will cause routing confusion.

In the figure, router B is not receiving individual routes for the 172.16.1.0/24 and 172.16.2.0/24 subnets. Both router A and router B automatically summarized those subnets to the 172.16.0.0/16 classful boundary when sending EIGRP update packets to router B. The result is that router B has two routes to 172.16.0.0/16 in the routing table, which can result in inaccurate routing and packet loss:

```
RouterB#show ip route
<output omitted>
```

```
Gateway of last resort is not set
```

```
10.0.0.0/24 is subnetted, 2 subnets
C    10.1.1.0 is directly connected, Serial0/2/0
C    10.2.2.0 is directly connected, Serial0/3/0
D    172.16.0.0/16 [90/2172416] via 10.1.1.1, 00:03:51, Serial0/2/0
     [90/2172416] via 10.2.2.3, 00:00:14, Serial0/3/0
```

The behavior of the **auto-summary** command is disabled, by default, on Cisco IOS Software versions 15 or later. Older software generally has automatic summarization enabled, by default.

To disable automatic summarization, enter the **no auto-summary** command in the router eigrp configuration mode:

```
RouterB(config)#router eigrp 1
RouterB(config-if)#no auto-summary
```

Summary

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

Summary

- Use the **show ip route** command to verify routes to remote networks.
- Use the **show ip eigrp neighbors** command to verify the EIGRP neighbor relationship.
- Use the **show ip interface brief** command to verify that the link between devices is operational.
- Use the **show ip eigrp interface *interface*** command to verify that the interface is participating in the EIGRP process.
- Use the **show ip protocols** command to verify that EIGRP AS numbers match, that proper networks are being advertised, that there are no interfaces misconfigured as passive, and that there is no ACL blocking EIGRP advertisements.

© 2013 Cisco Systems, Inc.

Implementing EIGRP for IPv6

Although proprietary to Cisco, EIGRP is widely used. Supporting IPv6 is important for the continued success of EIGRP. This lesson describes Cisco EIGRP support for IPv6, including its operation, configuration, and verification.

Objectives

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

- Describe how EIGRP is used for IPv6
- Explain commands that are used for enabling EIGRP for IPv6
- Present an example configuration of EIGRP for IPv6

EIGRP for IPv6

This topic describes how EIGRP for IPv6 differs from the EIGRP for IPv4.

EIGRP for IPv6

- Easy to configure
- Advanced distance vector mechanism with some features that are common to link-state protocols
- Uses protocol-dependent modules to support multiple protocols
- Supports IPv6 as a separate routing context

© 2013 Cisco Systems, Inc.

Although the configuration and management of EIGRP for IPv4 and EIGRP for IPv6 are similar, they are configured and managed separately.

EIGRP is inherently a multiprotocol routing protocol because it has supported non-IP IPX and AppleTalk for some time. IPv6 support is added as a separate module. IPv6 EIGRP is configured and managed separately from IPv4 EIGRP, but the mechanisms and configuration techniques will be familiar to those who are skilled with EIGRP for IPv4.

For example, both the IPv4 and IPv6 EIGRP implementations include a shutdown feature that allows the routing protocol to be configured but easily disabled. Both use the DUAL to optimize the routing path. Both are scalable to large networks. There are also a few differences in the IPv4 and IPv6 features. For example, in contrast with IPv4 EIGRP, IPv6 EIGRP is configured over a link—there is no network statement as there is for IPv4.

EIGRP for IPv6 (Cont.)

- Neighbor discovery
- Incremental updates
- Fast convergence—DUAL
- Uses multicast for updates
- Composite metric
- Load balancing
- Three tables:
 - Neighbor table
 - Topology table
 - Routing table

© 2013 Cisco Systems, Inc.

The basic components of EIGRP for IPv6 remain the same as the IPv4 version.

EIGRP uses a small hello packet to discover other EIGRP-capable routers on directly attached links and forms durable neighbor relationships. Updates may be acknowledged by using a reliable transport protocol, or they may be unacknowledged—depending on the specific function that is being communicated. The protocol provides the flexibility that is needed to unicast or multicast updates, acknowledged or unacknowledged.

Hello packets and updates are set to the well-known, link-local multicast address FF02::A that Cisco obtained from the IANA. This multicast distribution technique is more efficient than the broadcast mechanism that is used by earlier, more primitive routing protocols, such as RIPv1. EIGRP for IPv4 also uses multicast for update distribution.

EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature minimizes the bandwidth that is required for EIGRP packets.

DUAL, which is an EIGRP algorithm for determining the best path through the network, uses several metrics to select efficient, loop-free paths. When multiple routes to a neighbor exist, DUAL determines which route has the lowest metric (named the FD) and enters this route into the routing table. Other possible routes to this neighbor with larger metrics are received, and DUAL determines the reported distance to this network. The reported distance is defined as the total metric that is advertised by an upstream neighbor for a path to a destination. DUAL compares the reported distance with the FD, and if the reported distance is less than the FD, DUAL considers the route to be a feasible successor and enters the route into the topology table. The feasible successor route that is reported with the lowest metric becomes the successor route to the current route if the current route fails. To avoid routing loops, DUAL ensures that the reported distance is always less than the FD for a neighbor router to reach the destination network; otherwise, the route to the neighbor may loop back through the local router.

When there are no feasible successors to a route that has failed, but there are neighbors advertising the route, a recomputation must occur. This is the process where DUAL determines a new successor. The amount of time that is required to recompute the route affects the convergence time. Recomputation is processor-intensive. It is advantageous to avoid unneeded recomputation. When a topology change occurs, DUAL will test for feasible successors. If there are feasible successors, DUAL will use them in order to avoid unnecessary recomputation.

EIGRP updates contain five metrics: minimum bandwidth, delay, load, reliability, and MTU. Of these five metrics, by default, only minimum bandwidth and delay are used to compute the best path. Unlike most metrics, minimum bandwidth is set to the minimum bandwidth of the entire path, and it does not reflect how many hops or low-bandwidth links are in the path. Delay is a cumulative value that increases by the delay value of each segment in the path.

When a router discovers a new neighbor, it records the neighbor address and interface as an entry in the neighbor table. One neighbor table exists for each protocol-dependent module. 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 informed of the topology change.

The topology table contains all destinations that are advertised by neighboring routers. Each entry in the topology table includes the destination address and a list of neighbors that have advertised the destination. For each neighbor, the entry records the advertised metric, which the neighbor stores in its routing table. An important rule that distance vector protocols must follow is that if the neighbor advertises this destination, the neighbor must use the route to forward packets.

EIGRP for IPv6, like EIGRP for IPv4, is able to do load balancing. This is the capability of a router to distribute traffic over all the router network ports that are the same distance from the destination address. Load balancing increases the utilization of network segments and so increases effective network bandwidth. There are two types of load balancing:

- **Equal-cost path:** Applicable when different paths to a destination network report the same routing metric value
- **Unequal-cost path:** Applicable when different paths to a destination network report different routing metric values

EIGRP for IPv6 Commands

This topic lists and explains commands that are used to deploy EIGRP for IPv6.

EIGRP for IPv6 Commands

```
Router (config)#ipv6 unicast-routing
```

- Globally enables IPv6 routing and must be the first IPv6 command executed on the router

```
Router (config)#ipv6 router eigrp 1
```

- Creates and enters the EIGRP router submode with the AS 1

```
Router (config-rtr)#no shutdown
```

- EIGRP for IPv6 has a shutdown feature. The routing process should be in "no shut" mode in order to start running.

```
Router (config-if)#ipv6 eigrp 1  
Router (config-if)#no shutdown
```

- Configures EIGRP for IPv6 on an interface

© 2013 Cisco Systems, Inc.

| Command | Description |
|------------------------------------|---|
| ipv6 unicast-routing | By default, IPv6 traffic forwarding is disabled. This command enables it. |
| ipv6 router eigrp as-number | To place the router in router configuration mode, create an EIGRP routing process in IPv6, configure this process, and use the ipv6 router eigrp command in the global configuration mode. |
| no shutdown | EIGRP for IPv6 has a shutdown feature. The routing process should be in no shutdown mode in order to start running. Default behavior is different between Cisco IOS Software versions. |
| [no] ipv6 eigrp as-number | To enable EIGRP for IPv6 on a specified interface, use the ipv6 eigrp command in the interface configuration mode. To disable EIGRP for IPv6, use the no form of this command. |

These are some common configuration commands for EIGRP for IPv6. The syntax for these commands is similar, if not identical, to their IPv4 counterparts. For more information, follow this link: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_08.html.

EIGRP for IPv6 has a shutdown feature. The routing process should be in the **no shutdown** mode in order to start running.

EIGRP for IPv6 Commands (Cont.)

```
Router#show ipv6 eigrp topology
```

- Displays entries in the EIGRP IPv6 topology table

```
Router#show ipv6 eigrp neighbors
```

- Displays the neighbors that are discovered by EIGRP for IPv6

```
Router#show ipv6 route eigrp
```

- Shows EIGRP routes in the IPv6 routing table

© 2013 Cisco Systems, Inc.

The three **show** commands listed above have the same role as they have in EIGRP for IPv4.

To display entries in the EIGRP for IPv6 topology table, use the **show ipv6 eigrp topology** command in privileged EXEC mode.

To display the neighbors discovered by EIGRP for IPv6, use the **show ipv6 eigrp neighbors** command.

The **show ipv6 route eigrp** command shows the content of the IPv6 routing table that includes the routes specific to EIGRP.

EIGRP for IPv6 Configuration Example

This topic shows the principle of EIGRP configuration in an example.

EIGRP for IPv6 Configuration Example

```
Branch (config)#ipv6 router eigrp 1
Branch (config-router)#exit
Branch (config)#interface GigabitEthernet0/1
Branch (config-if)#ipv6 eigrp 1
```

- EIGRP for IPv6 configuration on Branch router

© 2013 Cisco Systems, Inc.

The example is a two-router network.

On the Branch router, EIGRP for IPv6 is enabled with AS 1. Notice that the **no shutdown** command is not present under **ipv6 router eigrp 1**, because EIGRP for IPv6 routing processes are turned on, by default, on Cisco IOS Software version 15. After that, EIGRP is enabled on the interface GigabitEthernet 0/1.

EIGRP for IPv6 Configuration Example (Cont.)

```
Branch (config)#ipv6 router eigrp 1
Branch (config)#exit
Branch (config)#interface GigabitEthernet0/0
Branch (config-if)#ipv6 eigrp 1
Branch (config-if)#exit
Branch (config)#interface GigabitEthernet0/1
Branch (config-if)#ipv6 eigrp 1
```

- EIGRP for IPv6 configuration on HQ router

© 2013 Cisco Systems, Inc.

On the HQ router, first, EIGRP for IPv6 is enabled with AS 1. Notice that the **no shutdown** command is not present under **ipv6 router eigrp 1**, because EIGRP for IPv6 routing processes are turned on by default on Cisco IOS Software version 15. After that, interfaces GigabitEthernet0/0 and GigabitEthernet 0/1 are enabled for IPv6 EIGRP.

EIGRP for IPv6 Configuration Example (Cont.)

```
Branch#show ipv6 eigrp interfaces
EIGRP-IPv6 Interfaces for AS(1)
      Xmit Queue PeerQ      Mean Pacing Time Multicast Pending
Interface Peers Un/Reliable Un/Reliable SRTT Un/Reliable Flow Timer Routes
Gi0/1      1      0/0      0/0      9      0/0      50      0
```

- Verifies that the Branch router has GigabitEthernet 0/1 interface

```
Branch#show ipv6 eigrp neighbors
EIGRP-IPv6 Neighbors for AS(1)
H   Address                Interface      Hold  Uptime   SRTT   RTO    Q  Seq
                               (sec)          (ms)                Cnt  Num
0   Link-local address:  Gi0/1         12    00:20:48  9     100    0  2
    FE80::FE99:47FF:FEE5:2671
```

- Verifies EIGRP neighbors

In the first output of the **show ipv6 eigrp interfaces** command, one neighbor is on the GigabitEthernet 0/1 interface, which is the only interface that is included in the EIGRP process.

The second example shows the output of the **show ipv6 eigrp neighbors** command. Details for the neighbor are displayed, such as the link-local address, interface, hold time, and uptime. Notice that the link-local address is used for forming EIGRP neighbor adjacency. In IPv6, neighbor interaction and next-hop addresses are always link-local.

To see more information on these two commands, follow this link: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_13.html.

EIGRP for IPv6 Configuration Example (Cont.)

```
Branch#show ipv6 eigrp topology
EIGRP-IPv6 Topology Table for AS(1)/ID(209.165.201.1)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status
P 2001:DB8:D1A5:C900::/64, 1 successors, FD is 28160
   via Connected, GigabitEthernet0/1
P 2001:DB8:AC10:100::/64, 1 successors, FD is 156160
   via FE80::FE99:47FF:FEE5:2671 (156160/128256), GigabitEthernet0/1
```

- Verifies the EIGRP for IPv6 topology table

© 2013 Cisco Systems, Inc.

EIGRP for IPv6 Configuration Example (Cont.)

```
Branch#show ipv6 route eigrp
IPv6 Routing Table - default - 4 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
D   2001:DB8:AC10:100::/64 [90/156160]
   via FE80::FE99:47FF:FEE5:2671, GigabitEthernet0/1
```

- Verifies IPv6 routes that are acquired via EIGRP

© 2013 Cisco Systems, Inc.

The **show ipv6 eigrp topology** command displays the topology table of EIGRP for IPv6 routes. All the routes are present in the topology table, but only the best routes are in the routing table.

The output of the **show ipv6 route eigrp** command is shown. Here, you are presented with a route that is learned by the EIGRP routing protocol.

To see more information on the command follow this link: http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_13.html.

Summary

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

Summary

- EIGRP has extended its multiprotocol support to IPv6.
- EIGRP for IPv6 is configured per interface on Cisco routers (there is no **network** command).
- You can enable IPv6 routing with the **ipv6 unicast-routing** command.
- EIGRP for IPv6 has a shutdown feature. The routing process should be in **no shutdown** mode in order to start running.

© 2013 Cisco Systems, Inc.

Module Summary

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

Module Summary

- EIGRP is a classless, advanced distance vector routing protocol that runs the DUAL algorithm.
- EIGRP is configured on a router through the **router eigrp** and **network** commands.
- There are three tables:
 - The EIGRP neighbor table lists directly connected routers running EIGRP.
 - The EIGRP topology table lists all routes learned from each EIGRP neighbor.
 - The routing table lists the best routes from the EIGRP topology table and the other routing processes.
- When you suspect an EIGRP issue, first check if the neighbors are up, then start troubleshooting the routing table.
- EIGRP for IPv6 is enabled per interface (there is no **network** command). EIGRP for IPv6 has a shutdown feature. The routing process should be in no shutdown mode in order to start running.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Do Not Duplicate.
Post beta, not for release.

Questions

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.

- Which command correctly specifies that network 10.0.0.0 is directly connected to a router that is running EIGRP? (Source: Implementing EIGRP)
 - Router(config)#**network 10.0.0.0**
 - Router(config)#**router eigrp 10.0.0.0**
 - Router(config-router)#**network 10.0.0.0**
 - Router(config-router)#**router eigrp 10.0.0.0**
- Connect the EIGRP features with their descriptions. (Source: Implementing EIGRP)

| | | |
|----------------------------|--------------------------|---|
| A. load balancing | <input type="checkbox"/> | a direct consequence of using partial updates |
| B. reduced bandwidth usage | <input type="checkbox"/> | routing mask is advertised for each destination network |
| C. classless routing | <input type="checkbox"/> | EIGRP knows two types: equal and unequal |
| D. DUAL | <input type="checkbox"/> | EIGRP algorithm by which EIGRP achieves rapid convergence |
- Which two criteria does EIGRP use, by default, to calculate its metric? (Choose two.) (Source: Implementing EIGRP)
 - bandwidth
 - reliability
 - load
 - MTU
 - delay
- Connect the terms to their descriptions. (Source: Implementing EIGRP)

| | | |
|----------------------------|--------------------------|--|
| A. feasible distance | <input type="checkbox"/> | used to rate the trustworthiness of each routing information source |
| B. administrative distance | <input type="checkbox"/> | the end-to-end metric from the router to that remote network |
| C. autonomous system | <input type="checkbox"/> | describes a set of contiguous routers that run the same routing protocol and share routing information |
| D. advertised distance | <input type="checkbox"/> | the EIGRP metric for an EIGRP neighbor to reach a particular network |
- Which command would you use to investigate which interfaces are enabled for the EIGRP routing process? (Source: Troubleshooting EIGRP)
 - show ip eigrp interfaces**
 - show ip eigrp neighbors**
 - show ip interfaces brief**
 - show eigrp enabled interfaces**

6. In which two ways does the configuration of EIGRP differ on IPv6 from IPv4? (Choose two.)
(Source: Implementing EIGRP for IPv6)
- A. The **network** command is changed into the **ipv6 network** command for EIGRP for IPv6 .
 - B. EIGRP for IPv6 can only be explicitly enabled with the **no shutdown** command. There is no **network** command.
 - C. EIGRP for IPv6 is configured per interface on Cisco routers.
 - D. If you run EIGRP for IPv6, you have to run EIGRP for IPv4; but if you run EIGRP for IPv4, you do not need to run EIGRP for IPv6.

Do Not Duplicate:
Post beta, not for release.

Answer Key

1. C
2. A. DUAL
B. load balancing
C. reduced bandwidth usage
D. classless routing
EIGRP algorithm by which EIGRP achieves rapid convergence
EIGRP knows two types: equal and unequal
a direct consequence of using partial updates
routing mask is advertised for each destination network
3. A, E
4. A. advertised distance
B. feasible distance
C. administrative distance
D. autonomous system
the EIGRP metric for an EIGRP neighbor to reach a particular network
the end-to-end metric from the router to that remote network
used to rate the trustworthiness of each routing information source
describes a set of contiguous routers that run the same routing protocol and share routing information
5. A
6. B, C

Do Not Duplicate, not for release.
Post beta, not for release.

Do Not Duplicate:
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Implementing a Scalable, Multiarea Network, OSPF-Based Solution

This module examines OSPF, which is one of the most commonly used IGPs in IP networking. OSPF is a complex protocol, and therefore configuration and verification of OSPF on a Cisco router is a primary learning objective.

The module discusses the primary configuration commands for a multiarea OSPF and explains the benefits of a multiarea OSPF solution compared to a single-area solution. Specifically, it covers link-state protocols, OSPF components, the OSPF metric, the way in which OSPF operates, and how to configure multiarea OSPF. Several OSPF **show** commands are also described in this module for verification purposes.

Objectives

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

- Describe the basic components and terms of OSPF
- Describe how to implement a multiarea OSPF
- Troubleshoot multiarea OSPF
- Implement OSPF in an IPv6 network

Do Not Duplicate.
Post beta, not for release.

OSPF Overview

Overview

OSPF is a link-state routing protocol that is often used in networks due to scalability, fast convergence, and multivendor environment support. Understanding OSPF operations and OSPF terms is therefore very important for network engineers that would like to design, implement, and troubleshoot scalable networks.

In this lesson, a quick overview of link-state routing protocols, OSPF data structures, and OSPF metrics is done before explaining the functions of OSPF and its packet types. The lesson describes how the link-state database is built using LSAs. OSPF areas are exposed as a structure that gives OSPF scalability, area terminology is explained, and important design limitations are exposed. The lesson also highlights the OSPF neighbor adjacency process.

Objectives

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

- Explain the basic idea behind link-state protocols
- Describe the data structures that are used by link-state routing protocols
- Describe the OSPF metric
- Describe how OSPF neighbor adjacencies are established
- Describe how routers build and synchronize the link-state database
- Describe the two-tier hierarchical structure of OSPF, including the characteristics of transit areas and regular areas as well as the terminology that is used

Link-State Routing Protocol Overview

OSPF is a link-state routing protocol. This topic describes link-state routing protocol basics.

Link-State Routing Protocol Overview

Link-state routing protocols such as OSPF have several advantages when compared to distance vector routing protocols:

- Link-state protocols are more scalable.
- Each router has a full picture of a topology.
- Updates are sent when a topology change occurs and are reflooded periodically.
- Link-state protocols respond quickly to topology changes.
- More information is communicated between routers.

© 2013 Cisco Systems, Inc.

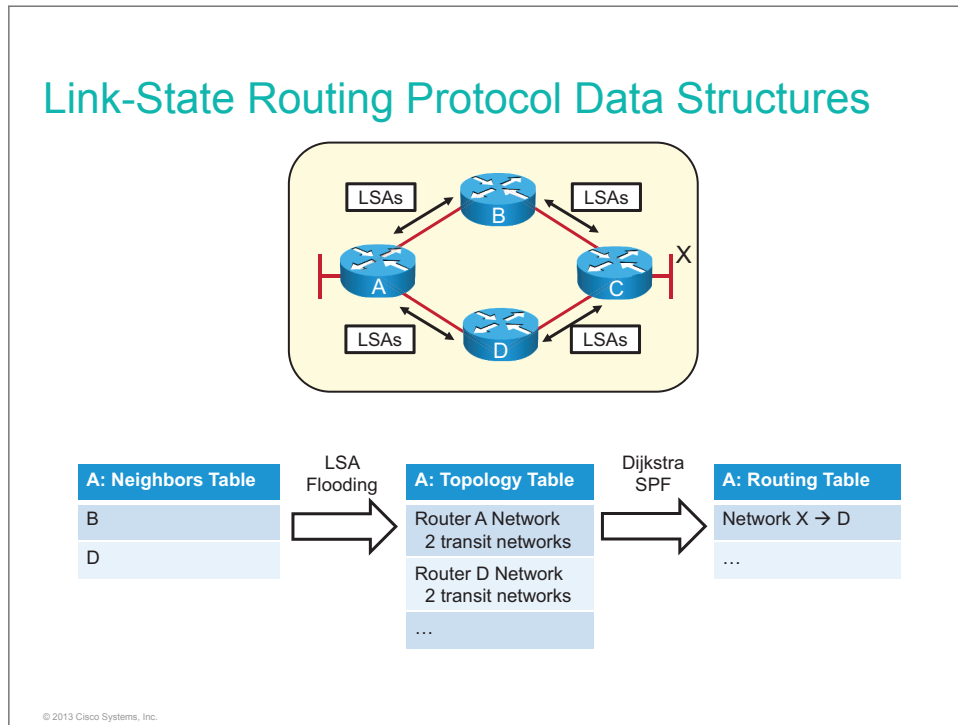
When a failure occurs in a network, routing protocols should detect the failure as soon as possible and find another path across the network. Only link-state protocols support fast convergence with support for scalability and multivendor environments, so they are the only type of IGP that is found in large network environments.

Link-state protocols have the following advantages when compared to distance-vector routing protocols:

- **They are more scalable:** Link-state protocols use a hierarchical design and can scale to very large networks if properly designed.
- **Each router has a full picture of a topology:** Because each router contains full information about all of the routers and links in a network, each router is able to independently select a loop-free and efficient pathway, which is based on cost, to reach every neighbor in the network.
- **Updates are sent when a topology change occurs and are reflooded periodically:** Link-state protocols send updates of a topology change. By using triggered updates, bandwidth is preserved. Additionally, updates are made periodically—by default every 30 minutes.
- **They respond quickly to topology changes:** Link-state protocols establish neighbor relations with adjacent routers. The failure of a neighbor is detected quickly, and this failure is communicated by using triggered updates to all routers in the network. This immediate reporting generally leads to fast convergence times.
- **More information is communicated between routers:** Routers that are running a link-state protocol have a common view on the network. This means that each router has full information about other routers and links between them, including the metric on each link.

Link-State Routing Protocol Data Structures

This topic describes data structures that are used by link-state routing protocols.



A router that is running a link-state routing protocol must first recognize other routers and establish a neighbor adjacency with its neighboring routers. A router achieves this neighbor adjacency by exchanging hello packets with the neighboring routers. After a router establishes a neighbor adjacency by using the hello packets, a neighbor is put into the neighbor database. In the example, router A recognizes routers B and D as neighbors.

After a neighbor relationship is established between routers, the routers synchronize their LSDBs by reliably exchanging LSAs. An LSA describes a router and networks that are connected to this router. LSAs are stored in the LSDB (topology database). By exchanging all LSAs, routers learn the complete topology of the network. Each router should have the same topology database.

After the topology database is built, each router applies the SPF algorithm to the topology map. The SPF algorithm uses the Dijkstra algorithm. The SPF algorithm builds a tree, where the root of the tree is the router itself and leaves are distant networks. The router places itself at the root of a tree and calculates the shortest path to each destination based on the cumulative cost that is required to reach this destination.

The best paths to destinations are then put into the routing table. The routing table includes a destination network and next-hop IP address. In the example, the routing table on router A states that a packet should be sent to router D to reach network X.


Whenever there is a change in a topology, new LSAs are created and sent throughout the network. All routers change their LSDB at the receipt of the new LSA, and the SPF algorithm is run again on the updated LSDB to verify new paths to destinations.

OSPF Metric

This topic describes the OSPF metric.

OSPF Metric

- OSPF uses path cost as a metric.
- By default, the cost is calculated based on the interface bandwidth.
- $\text{Cost} = \text{Reference Bandwidth} / \text{Interface Bandwidth}$, where reference bandwidth is 100 Mb/s.
- Path cost is a cumulated cost of all links on the path to destinations.



© 2013 Cisco Systems, Inc.

A metric is an indication of the overhead that is required to send packets across a certain interface. OSPF uses cost as a metric. Smaller cost indicates a better path than higher cost. The cost of an interface is inversely proportional to the bandwidth of this interface, so a higher bandwidth indicates a lower cost. There is more overhead, higher cost, and more time delays that are involved in crossing a 10-Mb/s Ethernet line than in crossing a 100-Mb/s Ethernet line.

The formula that is used to calculate OSPF cost is $\text{cost} = \text{reference bandwidth} / \text{interface bandwidth (in b/s)}$.

The default reference bandwidth is 10^8 , which is 100,000,000 or the equivalent of the bandwidth of Fast Ethernet. Therefore, the default cost of a 10-Mb/s Ethernet link will be $10^8 / 10^7 = 10$, and the cost of a 100-Mb/s link will be $10^8 / 10^8 = 1$. The problem arises with links that are faster than 100 Mb/s. Because OSPF cost has to be an integer, all links that are faster than Fast Ethernet will have an OSPF cost of 1. In this case, it is required to change OSPF cost on an interface manually or to adjust the reference bandwidth to a higher value.

The cost to reach a distant network from a router is a cumulated cost of all links on the path from the router to the network. In the example, the cost from router A to network X via router B is 30 (10 + 20), and the cost via router D is 20 (10 + 10). The path via router D is better because it has a lower cost.

Establishing OSPF Neighbor Adjacencies

This topic describes how OSPF neighbor adjacencies are established.

Establishing OSPF Neighbor Adjacencies

- OSPF routers first establish adjacencies.
- Hello packets are periodically sent to multicast address 224.0.0.5.
- Routers must agree on certain information (*) inside the hello packet before adjacency can be established.

Router ID
Hello/dead interval *
Neighbors
Area ID *
Router priority
DR IP address
BDR IP address
Authentication data *

Envelope icon
Hello

© 2013 Cisco Systems, Inc.

Neighbor OSPF routers must recognize each other on the network before they can share information because OSPF routing depends on the status of the link between two routers. This process is done using the Hello protocol. The Hello protocol establishes and maintains neighbor relationships by ensuring bidirectional (two-way) communication between neighbors. Bidirectional communication occurs when a router recognizes itself and is listed in the hello packet that is received from a neighbor.

Each interface that is participating in OSPF uses multicast address 224.0.0.5 to periodically send hello packets. A hello packet contains the following information:

- **Router ID:** The router ID is a 32-bit number that uniquely identifies the router. The router ID is, by default, the highest IP address on a loopback interface, if configured. If the router ID is not configured, it is the highest IP address on any interface. The router ID can also be configured manually using the **router-id** command. It is recommended to always use a loopback IP address for the router ID or to set the router ID manually. In this way, the router ID is stable and will not change.
- **Hello and dead intervals:** The hello interval specifies the frequency in seconds at which a router sends hello packets. The default hello interval on multiaccess networks is 10 seconds. The dead interval is the time in seconds that a router waits to hear from a neighbor before declaring the neighboring router out of service. By default, the dead interval is four times the hello interval. These timers must be the same on neighboring routers; otherwise, an adjacency will not be established.
- **Neighbors:** The Neighbors field lists the adjacent routers with established bidirectional communication. This bidirectional communication is indicated when the router recognizes itself and is listed in the Neighbors field of the hello packet from the neighbor.
- **Area ID:** To communicate, two routers must share a common segment and their interfaces must belong to the same OSPF area on this segment. The neighbors must also share the same subnet and mask. These routers will all have the same link-state information.

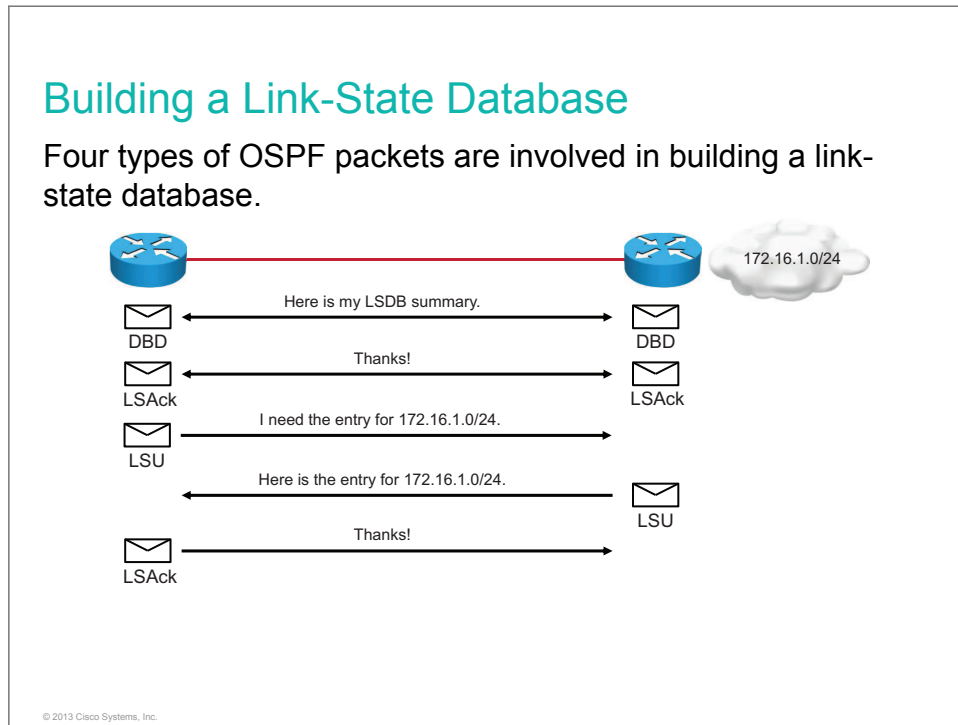
- **Router priority:** The router priority is an 8-bit number that indicates the priority of a router. OSPF uses the priority to select a DR and BDR. In certain types of networks, OSPF elects DRs and BDRs. The DR acts as a hub to reduce traffic between routers.
- **DR and BDR IP addresses:** These are the IP addresses of the DR and BDR for the specific network, if they are known.
- **Authentication data:** If router authentication is enabled, two routers must exchange the same authentication data. Authentication is not required, but if it is enabled, all peer routers must have the same key configured.

Note OSPF DRs and BDRs are not discussed in this course.

Do Not Duplicate.
Post beta, not for release.

Building a Link-State Database

This topic describes how routers build and synchronize an LSDB.



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

- **DBD packet:** A DBD packet is used to describe the network routes of each neighbor.
- **LSR packet:** After DBD packets are exchanged, the routers request missing information by using LSR packets.
- **LSU packet:** All missing information is sent to the neighbors by sending LSU packets that contain different LSAs.
- **LSAck packet:** Every packet receives an LSAck to ensure reliable transport and reliable exchange of information.

When two routers discover each other and establish adjacency using hello packets, they use the exchange protocol to exchange information about the LSAs.

As shown in the figure, the exchange protocol operates as follows:

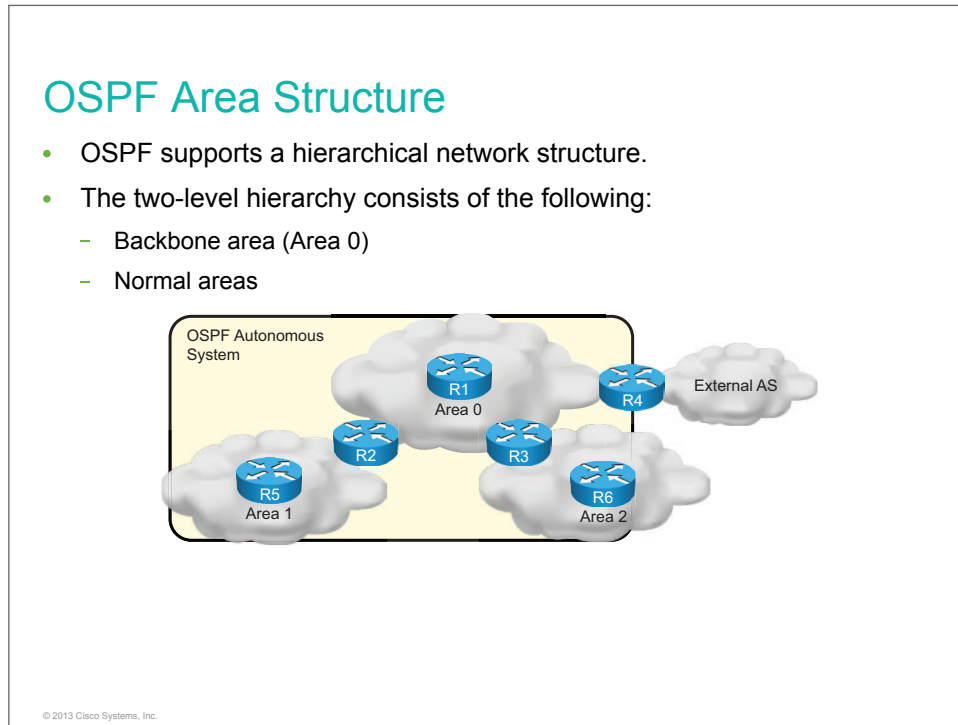
- 1 The routers exchange one or more DBD packets. A DBD includes information about the LSA entry header that appears in the LSDB of the router. Each LSA entry header includes information about the link-state type, the address of the advertising router, the cost of the link, and the sequence number. The router uses the sequence number to determine the “newness” of the received link-state information.
- 2 When the router receives the DBD, it acknowledges the receipt of the DBD that is using the LSAck packet.
- 3 Routers compare the information that they receive with the information that they have. If the received DBD has a more up-to-date, link-state entry, the router sends an LSR to the other router to request the updated link-state entry.
- 4 The other router responds with the complete information about the requested entry in an LSU packet. The other router adds the new link-state entries to its LSDB.

5 Again, when the router receives an LSU, it sends an LSAck.

Do Not Duplicate:
Post beta, not for release.

OSPF Area Structure

This topic describes the two-tier hierarchical structure of OSPF, including the characteristics of transit areas and regular areas as well as the terminology that is used.



In small networks, the web of router links is not complex, and paths to individual destinations are easily deduced. However, in large networks, the web is highly complex, and the number of potential paths to each destination is large. Therefore, the Dijkstra calculations that are comparing all of these possible routes can be very complex and can take a significant amount of time to complete.

Link-state routing protocols usually reduce the size of the Dijkstra calculations by partitioning the network into areas. The number of routers in an area and the number of LSAs that flood within the area are small, which means that the link-state or topology database for an area is small. Consequently, the Dijkstra calculation is easier and takes less time. Routers that are inside an area maintain detailed information about the links and only general or summary information about routers and links in other areas. However, summarization is not done by default; it must be configured. Another advantage of using multiarea OSPF design is that a topology change in an area causes LSA flooding only within the area. SPF recalculations therefore occur only in an area where topology change has happened.

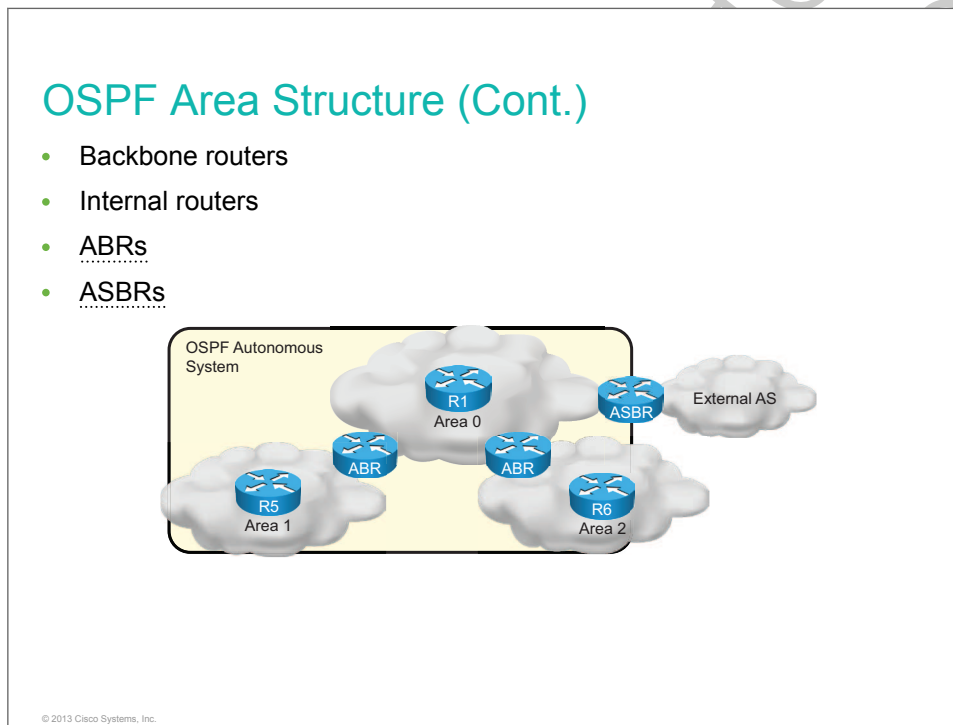
Link-state routing protocols use a two-layer area hierarchy:

- **Backbone or transit area:** The primary function of this OSPF area is to quickly and efficiently move IP packets. Backbone areas interconnect with other OSPF area types. The OSPF hierarchical area structure requires that all areas connect directly to the backbone area. In the figure, links between Area 1 and Area 2 routers are not allowed. Generally, end users are not found within a backbone area, which is also known as OSPF Area 0.

- **Normal or nonbackbone area:** The primary function of this OSPF area is to connect users and resources. Normal areas are usually set up according to functional or geographical groupings. By default, a normal area does not allow traffic from another area to use its links to reach other areas. All traffic from other areas must cross a transit area such as Area 0. Normal areas can be of different types. Normal area types affect the amount of routing information that is propagated into the normal area. For example, instead of propagating all routes from the backbone area into a normal area, you could propagate only a default route.

Note An OSPF area is identified using a 32-bit Area ID. It can be expressed as either a decimal number or dotted decimal. Both formats can be used at the same time. For example, Area 0 and Area 0.0.0.0 are equivalent. The same goes for Area 14 and Area 0.0.0.14. Area 300 would be the same as 0.0.1.44.

In general, any one router should have no more than 60 neighbors. Generally, an area should have no more than 50 routers. Areas with unstable links should be smaller. In general, to maximize stability, one router should not be in more than three areas.



All OSPF areas and routers that are running the OSPF routing protocol comprise the OSPF AS.

Routers that make up Area 0 are known as backbone routers. OSPF hierarchical networking defines Area 0 as the core. All other areas connect directly to backbone Area 0.

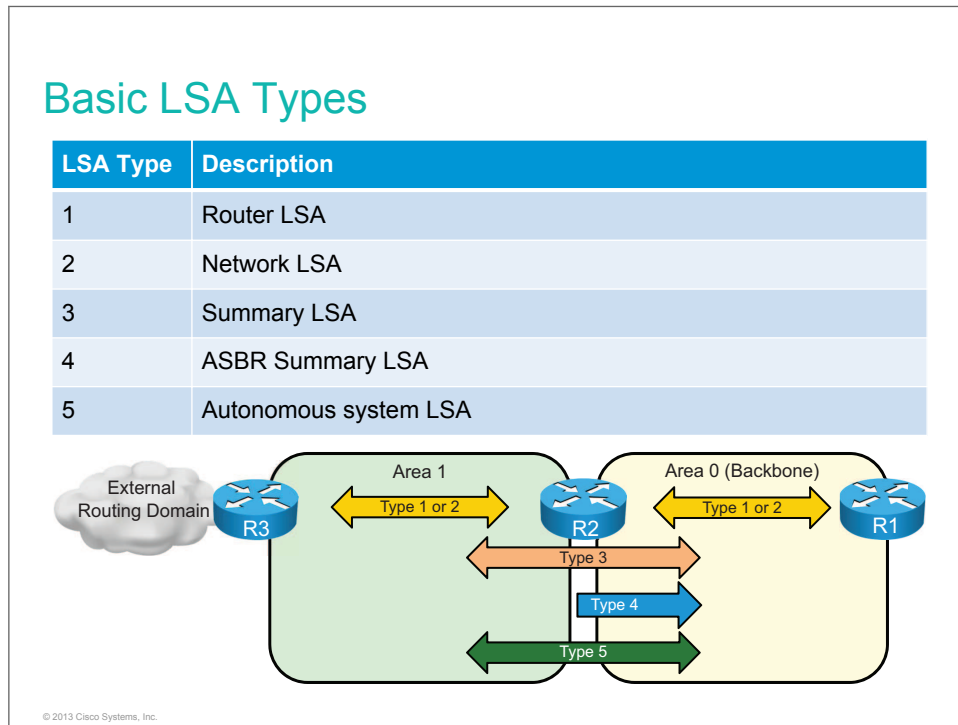
Routers that make up nonbackbone (normal) areas are known as internal routers; they have all interfaces in one area only.

An ABR connects Area 0 to the nonbackbone areas. An OSPF ABR plays a very important role in network design and has interfaces in more than one area. An ABR has the following characteristics:

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

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

An ASBR connects any OSPF area to a different routing administration. The ASBR is the point where external routes can be introduced into the OSPF AS.



LSAs are the building blocks of the OSPF LSDB. Individually, they act as database records; in combination, they describe the entire topology of an OSPF network or area. There are many types of LSAs; some basic ones are described below.

Type 1: Every router generates router link advertisements for each area to which it belongs. Router link advertisements describe the state of the router links to the area and are flooded only within this particular area.

Type 2: Generated by the DR, these are flooded in the area that contains the network.

Type 3: An ABR takes the information that it learned in one area and summarizes it for another area. This summarization is not on by default. This summarization means smaller routing tables of OSPF routers.

Type 4: This informs the rest of the OSPF domain how to get to the ASBR.

Type 5: External link advertisements that are generated by the ASBRs. They get flooded everywhere, except into special areas.

In the example, R2 is an ABR between Areas 0 and 1. Area 0 is usually called the backbone area. R3 acts as the ASBR between the OSPF routing domain and an external domain. LSA types 1 and 2 are flooded between routers within an area. Type 3 and type 5 LSAs are flooded within backbone and standard areas. Type 4 LSAs are injected into the backbone by the ABR because all routers in the OSPF domain need to reach the ASBR (R3).

Summary

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

Summary

- Link-state routing protocols such as OSPF are more scalable and converge faster than distance-vector routing protocols.
- OSPF uses three data structures: neighbors table, topology table, and routing table.
- OSPF uses cost as a metric. The cost of an interface is inversely proportional to the bandwidth of the interface. Smaller cost indicates a better path than higher cost.
- OSPF routers first establish adjacency using hello packets.
- OSPF routers synchronize the LSDB by using the exchange protocol, which utilizes four OSPF packet types.
- OSPF supports a two-tier hierarchical network architecture.

© 2013 Cisco Systems, Inc.

Multiarea OSPF IPv4 Implementation

OSPF routing protocol supports a two-tier hierarchical structure. By utilizing a two-tier or multiarea OSPF design, you can increase the network scalability and reduce the load and utilization on routers due to fewer SPF calculations and smaller routing tables. This lesson first describes differences between single-area and multiarea OSPF design. Then the lesson describes how OSPF implementation should be planned. Finally, the lesson describes how to configure and verify multiarea OSPF design in an IPv4 network.

Objectives

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

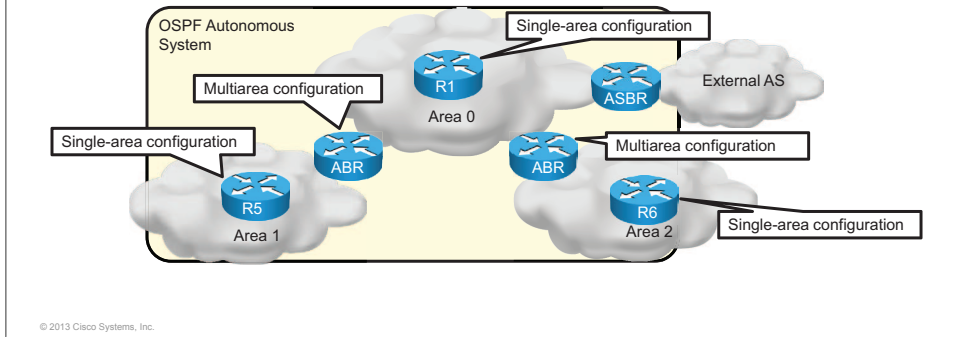
- Explain the difference between single-area and multiarea OSPF
- Describe how OSPF implementation should be planned
- Configure multiarea OSPF
- Verify multiarea OSPF

Single-Area vs. Multiarea OSPF

This topic describes the differences between single-area and multiarea OSPF implementation.

Single-Area vs. Multiarea OSPF

- Single-area OSPF:
 - Many LSAs processed on every router
 - Large routing tables
- Multiarea OSPF:
 - LSA processing confined to an area
 - Smaller routing tables if summarization is used



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

Multiarea design is a better solution than single-area design. In a multiarea design, the network is segmented to limit the propagation of LSAs inside an area and to make routing tables smaller by utilizing summarization. There are two types of routers from the configuration point of view:

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

While multiarea OSPF is scalable and a powerful routing protocol, it requires much knowledge to properly design, implement, or troubleshoot it.

Planning for the Implementation of OSPF

This topic describes how OSPF implementation should be planned.

Planning for the Implementation of OSPF

- Assess the requirements and options:
 - Verify IP addressing.
 - Verify network topology.
- Define ABRs and ASBRs.
- Create an implementation plan.
- Configure OSPF.

© 2013 Cisco Systems, Inc.

You can implement the OSPF routing protocol as a single-area or multiarea OSPF. The type of OSPF implementation that you should choose depends on your specific requirements and existing topology.

- When you prepare to deploy OSPF routing in a network, you must first gather the existing state and requirements. Then, based on this, you should consider single-area or multiarea deployment.
 - The IP addressing plan governs how OSPF can be deployed and how well the OSPF deployment might scale. Thus, a detailed IP addressing plan, along with the IP subnetting information, must be created. A solid IP addressing plan should enable the usage of OSPF multiarea design and summarization to more easily scale the network as well as optimize OSPF behavior and the propagation of LSAs.
 - The network topology consists of links that connect the network equipment and that belong to different OSPF areas in a multiarea OSPF design. Network topology is important to determine primary and backup links. Primary and backup links are defined by the changing OSPF cost on interfaces.
- A detailed network topology plan should be used to determine the different OSPF areas, ABRs, ASBRs, and summarization points, if multiarea OSPF will be used.
- An implementation plan must be created before configuring OSPF routing in the network.
- After the implementation plan, OSPF can actually be implemented. The requirements and existing state will determine the deployment type.

Multiarea OSPF Configuration

This topic shows how to configure a basic multiarea OSPF network.

Multiarea OSPF Configuration

```

R1(config)# interface GigabitEthernet0/0
R1(config-if)# ip address 10.64.0.1 255.255.255.0
R1(config-if)# ip ospf cost 10
R1(config-if)# exit
R1(config)# router ospf 1
R1(config-router)# network 10.64.0.0 0.0.0.255 area 0
    
```

OSPF configuration on R1

© 2013 Cisco Systems, Inc.


The configuration of multiarea OSPF is based on commands that are described in the table.

| Command | Description |
|--|---|
| ip ospf cost <i>cost</i> | Specifies the OSPF cost of sending a packet on an interface. The cost can be a value in the range from 1 to 65,535. |
| router ospf <i>process_id</i> | Configures an OSPF routing process. The <i>process-id</i> parameter is an internally used identification parameter for the OSPF routing process. It is locally assigned and can be any positive integer. A unique value is assigned for each OSPF routing process in the router. |
| network <i>network wildcard_mask area area_id</i> | Defines the interfaces on which OSPF runs and defines the area IDs for those interfaces. The <i>wildcard_mask</i> parameter determines how to interpret the IP address. The mask has wildcard bits in which 0 is a match and 1 indicates that the value is not significant. For example, 0.0.255.255 indicates a match in the first two octets. |

Note OSPF can also be enabled directly on an interface using the **ip ospf area** *area_id* command, which simplifies the configuration of unnumbered interfaces. Because the command is configured explicitly for the interface, it takes precedence over the **network area** command.

The figure shows the OSPF configuration for the Gigabit Ethernet network on the R1 router. Router R1 is in Area 0, router R3 is in Area 1, and router R2 is the ABR between the two areas. R1 is configured for OSPF with a process ID value of 1, and a network statement assigns the GigabitEthernet0/0 interface to Area 0. Additionally, the OSPF cost is changed to 10 on the GigabitEthernet0/0 interface.

Multiarea OSPF Configuration (Cont.)



```
R2(config)# interface GigabitEthernet0/0
R2(config-if)# ip address 10.64.0.2 255.255.255.0
R2(config-if)# exit
R2(config)# interface GigabitEthernet0/1
R2(config-if)# ip address 10.2.1.2 255.255.255.0
R2(config-if)# exit
R2(config)# router ospf 1
R2(config-router)# network 10.64.0.0 0.0.0.255 area 0
R2(config-router)# network 10.2.1.2 0.0.0.0 area 1
```

OSPF configuration on R2

© 2013 Cisco Systems, Inc.

R2 is running OSPF process 1 and has one network statement for Area 0 and one for Area 1. One network statement assigns the GigabitEthernet0/0 interface to Area 0, and the other network statement assigns the GigabitEthernet0/1 interface to Area 1. The network statement for Area 1 specifies the exact IP address of the interface with wildcard mask 0.0.0.0, while the network statement for Area 0 specifies the IP address of the entire network on the interface.

Configuration for R3 is not shown, but it is similar to the configuration on R1. The only difference is that R3 has GigabitEthernet0/1 in Area 1.

Multiarea OSPF Verification

This topic describes how to verify multiarea OSPF.

Multiarea OSPF Verification

Verify OSPF neighbors. R2 has two neighbors in this example:

```
R2# show ip ospf neighbor
Neighbor ID Pri State Dead Time Address Interface
10.64.0.1 1 FULL/DR 00:00:31 10.64.0.1 GigabitEthernet0/0
10.2.1.1 1 FULL/DR 00:00:51 10.2.1.1 GigabitEthernet0/1
```

Verify OSPF-enabled interfaces. R2 has two in this example: GigabitEthernet0/0 in Area 0 and GigabitEthernet 0/1 in Area 1:

```
R2# show ip ospf interface
GigabitEthernet0/0 is up, line protocol is up
Internet Address 10.64.0.2/24, Area 0
Process ID 1, Router ID 10.64.0.2, Network Type BROADCAST, Cost: 1
<output omitted>
GigabitEthernet0/1 is up, line protocol is up
Internet Address 10.2.1.2/24, Area 1
Process ID 1, Router ID 10.64.0.2/24, Network Type BROADCAST, Cost: 1
<output omitted>
```

To verify multiarea OSPF configuration, use the same commands that you would use to verify single-area OSPF. Use the **show ip ospf neighbor** command to verify OSPF neighbors. This command displays the neighbor router ID, neighbor priority, OSPF state, dead timer, neighbor interface IP address, and interface through which the neighbor is accessible. In the example, the R2 router has two neighbors, and each neighbor is reachable through a separate interface.

Use the **show ip ospf interface** command to display OSPF-related information on OSPF-enabled interfaces. This command will reveal useful information such as the OSPF process ID to which the interface is assigned, the area that the interfaces are in, and the cost of the interface.

Multiarea OSPF Verification (Cont.)

Verify for which networks that R2 is routing:

```
R2# show ip protocols
*** IP Routing is NSF aware ***
Routin Protocol is "ospf 1"
<output omitted>
Routing for Networks:
 10.64.0.0 0.0.0.255 area 0
 10.2.1.2 0.0.0.0 area 1
Routing Information Sources:
Gateway      Distance    Last Update
10.64.0.1    110         00:06:07
10.2.1.1     110         00:06:07
<output omitted>
```

© 2013 Cisco Systems, Inc.

Use the **show ip protocols** command to verify the OSPF status. The output of the command will reveal which routing protocols are configured on a router, including routing protocol specifics. In the case of OSPF, you can see the router ID, number of areas in the router, and networks that the router routes for. In the example, the R2 router, which is the ABR, is configured for two areas, with one interface in each area.

Multiarea OSPF Verification (Cont.)

Verify the routing table:

```
R1# show ip route ospf
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override
Gateway of last resort is not set
 10.0.0.0/24 is subnetted, 1 subnets
 O IA   10.2.1.0 [110/11 via 10.64.0.2, 00:46:20, GigabitEthernet0/0]
<output omitted>
```

© 2013 Cisco Systems, Inc.

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

In the example, the 10.2.1.0 subnet is recognized on GigabitEthernet 0/0 via neighbor 10.64.0.2. The “O” code represents OSPF routes, and “IA” indicates interarea, which means that the route originated from another area. Recall that R1 is in Area 0, and the 10.2.1.0 subnet is connected to router R2 in Area 1. The entry “[110/11]” in the routing table represents the administrative distance that is assigned to OSPF (110) and the total cost of the route to subnet 10.2.1.0 (cost of 11).

Do Not Duplicate.
Post beta, not for release.

Summary

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

Summary

- Multiarea OSPF design enables segmentation of a network to limit the propagation of LSAs inside an area and to make routing tables smaller (if utilizing summarization).
- The type of OSPF implementation depends on requirements and existing topology.
- Multiarea OSPF configuration is similar to single-area configuration. Interfaces are assigned to areas using the **network** command.
- Multiarea OSPF verification is similar to single-area verification. Interarea routes are marked with the "IA" code in the routing table.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Troubleshooting Multiarea OSPF

OSPF is a link-state routing protocol and therefore it scales well with a growing network. However, this scalability introduces complexity in design, configuration, and maintenance. This lesson introduces OSPF neighbor states, which are important when troubleshooting OSPF adjacencies. Then the lesson lists some of the common issues surrounding an OSPF network and provides a flow chart approach to troubleshooting these issues. The lesson also describes each issue and provides troubleshooting recommendations.

Objectives

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

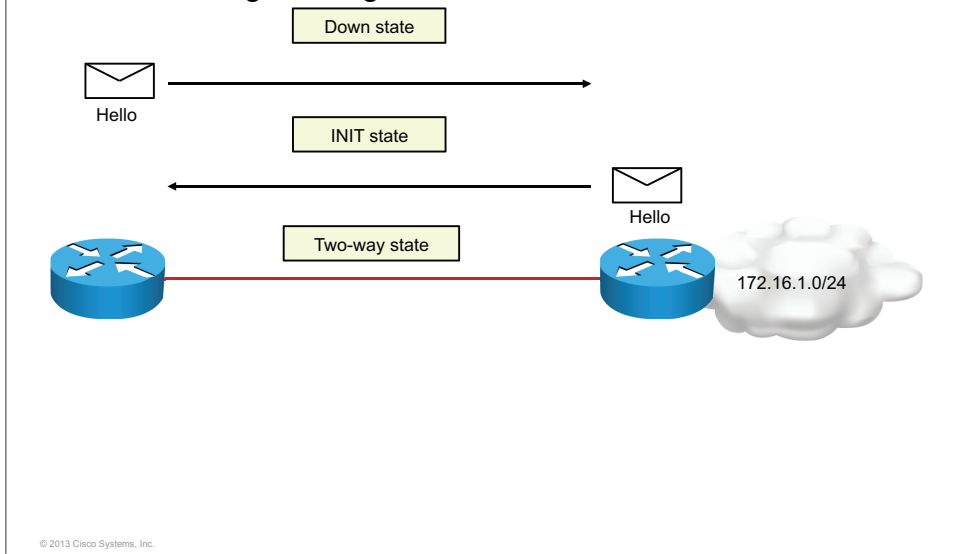
- Describe OSPF neighbor states
- Describe how to troubleshoot OSPF
- Troubleshoot OSPF neighbor issues
- Troubleshoot OSPF routing table issues
- Troubleshoot OSPF path selection issues

OSPF Neighbor States

Before learning how to troubleshoot OSPF, it is important to learn how OSPF routers traverse different OSPF states when adjacencies are being established. This topic describes OSPF neighbor states.

OSPF Neighbor States

OSPF routers go through different OSPF states:



When routers that are running OSPF are initialized, an exchange process using the Hello protocol is the first procedure.

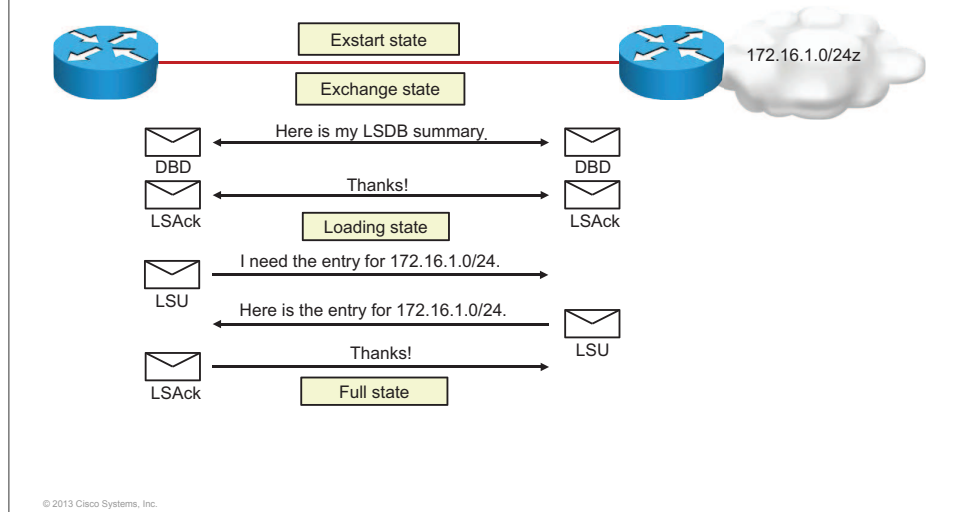
The exchange process that happens when routers appear on the network is illustrated in the figure:

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

If the link type is a broadcast network, such as a LAN link like Ethernet, a DR and BDR must first be selected. The DR acts as a central exchange point for routing information and reduces the amount of routing information that the routers have to exchange. The DR and BDR are selected after routers are in the two-way state. A router with the highest priority will become the DR. If there is a tie, a router with the highest router ID will become the DR. Among the routers on a LAN that were not elected as the DR or BDR, the exchange process stops at this point, and the routers remain in the two-way state.

OSPF Neighbor States (Cont.)

All states except two-way and full are transitory, and routers should not remain in these states for extended periods of time.



After the DR and BDR have been selected (or applied to LAN link types), the routers are considered to be in the exstart state. The routers are then ready to discover the link-state information about the internetwork and create their LSDBs. The exchange protocol is used to discover the network routes, and it brings all of the routers from the exchange state to a full state of communication. The first step in this process is for the DR and BDR to establish adjacencies with each of the other routers.

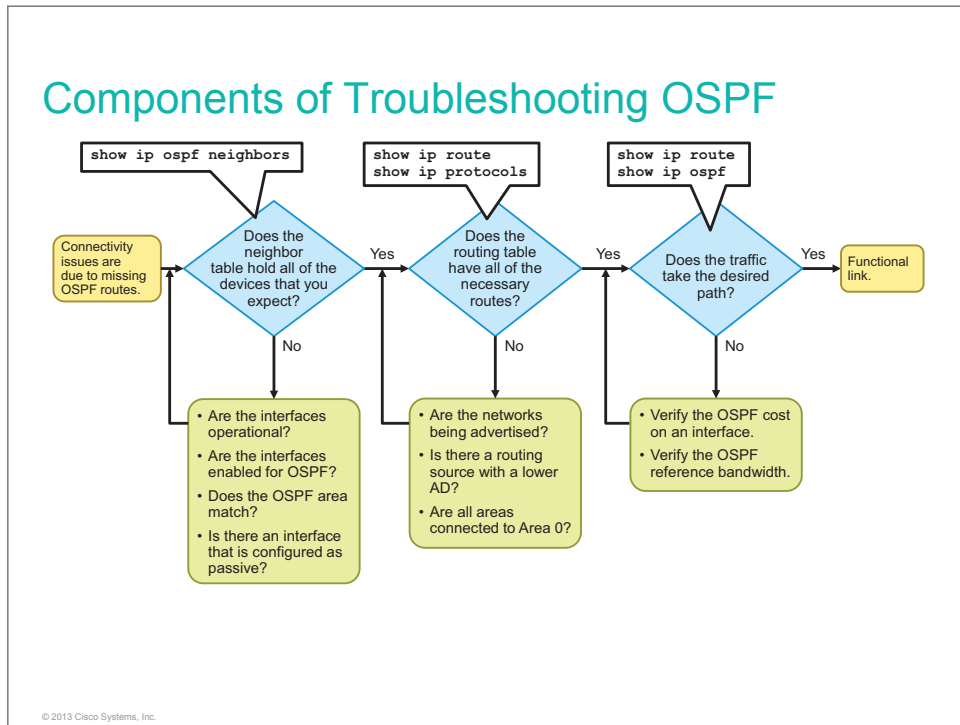
As shown in the figure, the exchange protocol continues as follows:

- 1 In the exstart state, the DR and BDR establish adjacencies with each router in the network. During this process, a master-slave relationship is created between each router and its adjacent DR and BDR. The router with the higher router ID acts as the master during the exchange process. The master-slave election dictates which router will start the exchange of routing information. This step is not shown in the figure.
- 2 The master and slave routers exchange one or more DBD packets. The routers are in the exchange state.
- 3 A router compares the DBD that it received with the LSAs that it has. If the DBD has a more up-to-date link-state entry, the router sends an LSU to the other router. When routers start sending LSAs, they are in the loading state.
- 4 When all LSAs have been satisfied for a given router, the adjacent routers are considered synchronized and are in the full state.

When troubleshooting OSPF neighbors, you should be aware that all states except two-way and full are transitory, and routers should not remain in these states for extended periods of time. However, the most likely problem that you will experience is that you will not see neighbors at all. In this case, verify the OSPF configuration.

Components of Troubleshooting OSPF

This topic describes how to troubleshoot OSPF.



When you are first notified that there are connectivity issues in your network, you should first test connectivity, using the **ping** and **traceroute** commands. If there are connectivity issues and your network uses OSPF as the routing protocol, then follow these high-level steps to troubleshoot it:

- 1 Verify if your router established adjacency with a neighboring router using the **show ip ospf neighbors** command. If adjacency between two routers is not established, the routers cannot exchange routes. If adjacency is not established, you should first verify if the interfaces are operational and enabled for OSPF. If the interfaces are operational and enabled for OSPF, you should also make sure that the interfaces on both routers are configured for the same OSPF area and the interfaces are not configured as passive interfaces.
- 2 If adjacency between two routers is established, but you see no routes in the routing table using the **show ip route** command, you should first verify if there is another routing protocol with lower administrative distance running in the network. In this case, OSPF routes would not be considered to be put into the routing table. If no other routing protocols are configured, verify if all of the required networks are advertised into OSPF. In the case of multiarea OSPF, you should also verify if all regular nonbackbone areas are connected directly to Area 0 or the backbone area. If a regular area is not connected to the backbone area, routers in this area will not be able to send and receive updates to and from other areas.
- 3 If you see all of the required routes in the routing table but the path that traffic takes is not correct, you should verify the OSPF cost on interfaces on the path. You should also be careful in cases where you have interfaces that are faster than 100 Mb/s, because all interfaces above this bandwidth will have the same OSPF cost, by default.

Troubleshooting OSPF Neighbor Issues

This topic describes how to troubleshoot OSPF neighbor issues.

Troubleshooting OSPF Neighbor Issues

HQ is not my OSPF neighbor.
Are my interfaces Layer 2 operational?

OSPF Area 1
10.1.1.0/24

OSPF Area 0
172.16.1.0/24

Branch S0/0/0 192.168.1.1 S0/0/0 192.168.1.2 HQ

```
Branch# show ip interface brief
Interface                IP-Address      OK? Method Status  Protocol
GigabitEthernet0/0      10.1.1.1        YES manual up      up
Serial0/0/0              192.168.1.1     YES manual up      up
<output omitted>
```

Verify if Serial 0/0/0 on the Branch router is Layer 2 operational.

© 2013 Cisco Systems, Inc.

A prerequisite for the neighbor relationship to form between the Branch and Headquarters routers is OSI Layer 3 connectivity. By investigating the **show ip interface brief** output, you can verify that the status and protocol are both up for the Serial0/0/0 interface that is connected to the Branch router. This confirms that the link is operational on Layer 2.

Troubleshooting OSPF Neighbor Issues (Cont.)

HQ is not my OSPF neighbor.
Are my interfaces Layer 3 operational?

OSPF Area 1
10.1.1.0/24

OSPF Area 0
172.16.1.0/24

Branch S0/0/0 192.168.1.1 S0/0/0 192.168.1.2 HQ

```
Branch# ping 192.168.1.2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.2, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Verify Layer 3 connectivity between the Branch and Headquarters routers.

© 2013 Cisco Systems, Inc.

A ping from the Branch to Headquarters routers will confirm IP connectivity between the devices: If the ping is not successful, check the cabling and verify that interfaces on connected devices are operational and that they are on a common subnet with the same subnet mask.

In the example, the Serial0/0/0 interface is enabled on both routers.

Troubleshooting OSPF Neighbor Issues (Cont.)

OSPF Area 1
10.1.1.0/24

OSPF Area 0
172.16.1.0/24

Branch

HQ

Serial0/0/0
192.168.1.1

Serial0/0/0
192.168.1.2


```
Branch# show ip ospf interface
Serial0/0/0 is up, line protocol is up
Internet Address 192.168.1.1/24, Area 0
Process ID 1, Router ID 209.165.201.1, Network Type POINT_TO_POINT, Cost: 64
<output omitted>
```

If the interfaces on both routers are not enabled for OSPF, the adjacency will not form.

© 2013 Cisco Systems, Inc.

The **network** command that is configured under the OSPF routing process indicates which router interfaces will participate in OSPF. You can use the **show ip ospf interface** command to verify which interfaces are enabled for OSPF. The output will also show you which interface is functional and the OSPF-related parameters. If connected interfaces on two routers are not enabled for OSPF, the neighbors will not form an adjacency.

Troubleshooting OSPF Neighbor Issues (Cont.)



```
HQ# show ip ospf interface
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.1.2/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 64
<output omitted>
```

If the interfaces on both routers are not enabled for OSPF, the adjacency will not form.

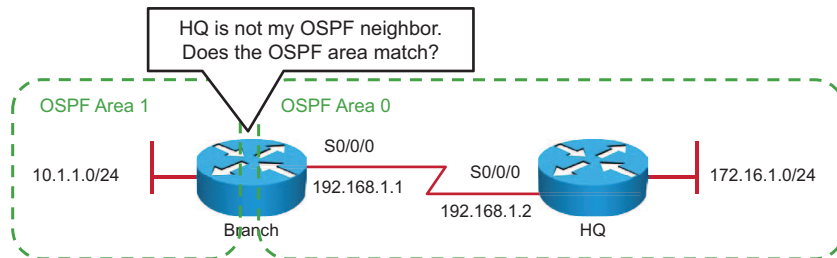
© 2013 Cisco Systems, Inc.

You can also use the **show ip protocols** command to verify which interfaces are configured for OSPF. The output will show you IP addresses or networks that are enabled using the **network** command. If an IP address on an interface falls within a network that has been enabled for OSPF, the interface will be enabled for OSPF. The output of this command will also show you if OSPF is enabled on an interface, using the **ip ospf area** command. The following is an example of the **show ip protocols** command:

```
HQ# show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    172.16.1.0 0.0.0.255 area 0
    192.168.1.0 0.0.0.255 area 0
  Routing on Interfaces Configured Explicitly (Area 0):
    Loopback0
<output omitted>
```

In the example, OSPF is enabled on the Serial0/0/0 interface on both routers.

Troubleshooting OSPF Neighbor Issues (Cont.)

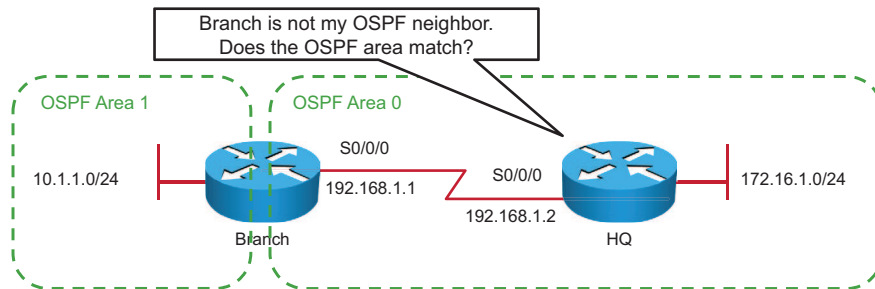


```
Branch# show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
Maximum path: 4
Routing for Networks:
 10.1.1.0 0.0.0.255 area 1
 192.168.1.0 0.0.0.255 area 0
<output omitted>
```

If the OSPF area does not match on both ends, the adjacency will not form.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)



```
HQ# show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
Routing for Networks:
 172.16.1.0 0.0.0.255 area 0
 192.168.1.0 0.0.0.255 area 0
<output omitted>
```

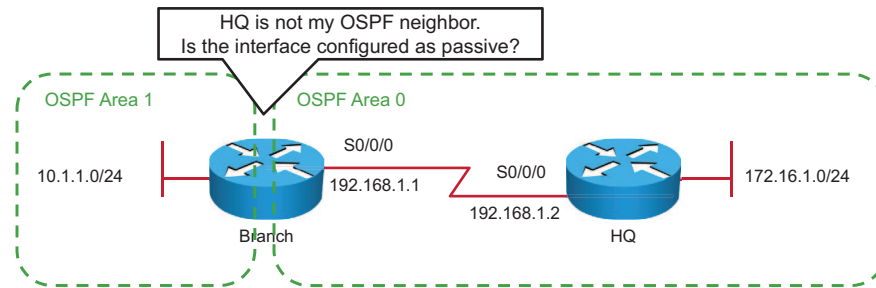
If the OSPF area does not match on both ends, the adjacency will not form.

© 2013 Cisco Systems, Inc.

When you specify networks that will be advertised using OSPF, you have to provide the OSPF area number. The OSPF area numbers on two directly connected interfaces have to be the same, or the adjacency will not form. You can verify the area that an interface has been enabled for by using the **show ip protocols** command.

In the example, OSPF is enabled for the same area on both routers.

Troubleshooting OSPF Neighbor Issues (Cont.)

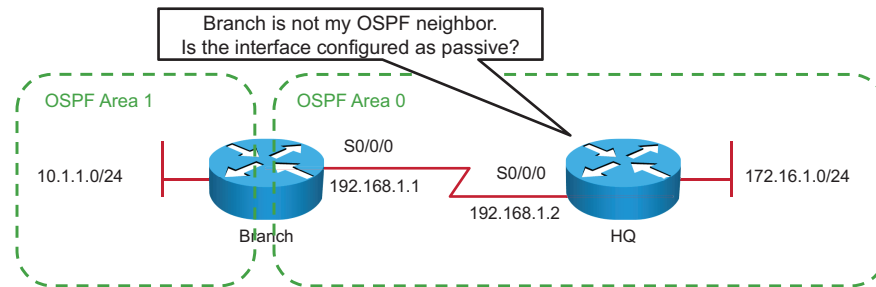


```
Branch# show ip protocols
Routing Protocol is "ospf 1"
<output omitted>
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 1
    192.168.1.0 0.0.0.255 area 0
<output omitted>
```

Check if the interface toward the Headquarters router is configured as passive.

© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)



© 2013 Cisco Systems, Inc.

Troubleshooting OSPF Neighbor Issues (Cont.)

```
HQ# show ip protocols
Routing Protocol is "ospf 1"
<...output omitted...>
Routing for Networks:
 172.16.1.0 0.0.0.255 area 0
 192.168.1.0 0.0.0.255 area 0
Passive Interface(s):
 Serial0/0/0
<output omitted>
```

Headquarters has the interface toward the Branch router configured as passive. This is why the two routers are not forming an adjacency.

With OSPF running on a network, the **passive-interface** command stops both outgoing and incoming routing updates because the effect of the command causes the router to stop sending and receiving hello packets over an interface. For this reason, the routers will not become neighbors.

To verify if any interface on a router is configured as passive, use the **show ip protocols** command in privileged mode.

An example in which you want to configure the interface as passive is handing off a link to a third-party organization that you have no control over (for example, an ISP). In this case, you would need to advertise this particular link through your own network but not allow the third party to receive hellos or send hellos to your device. This would be a security risk.

To configure an interface as a passive interface in OSPF, you will use the **passive-interface interface** command in the OSPF router configuration mode. To disable the interface as passive, use the **no passive-interface interface** command.

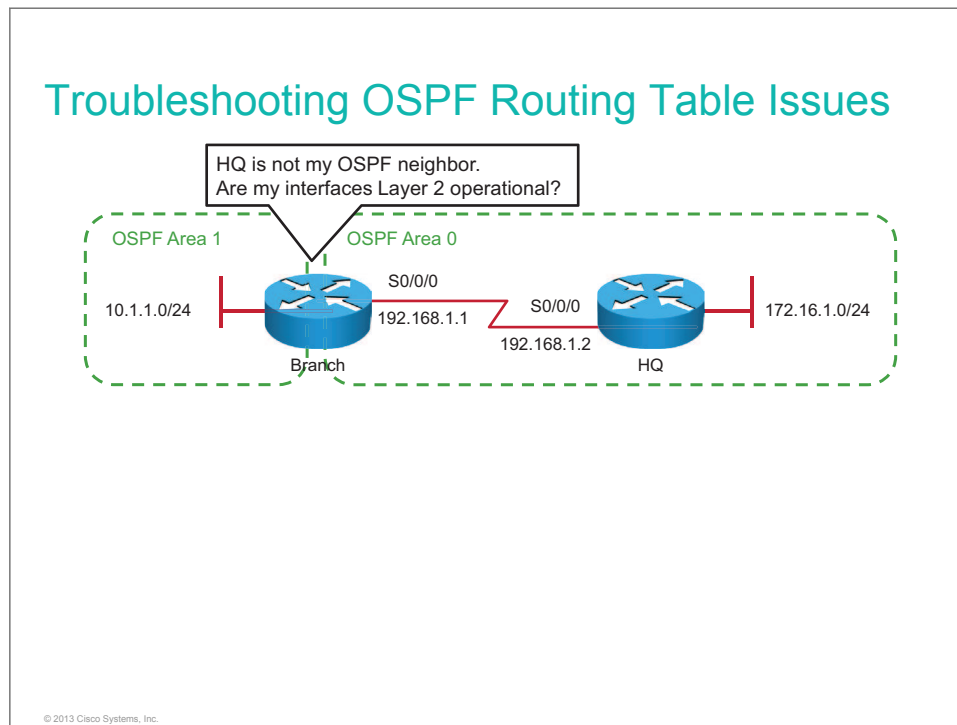
Once you disable the passive interface, the routers should become adjacent as indicated by the **show ip ospf neighbor** command output. Recall that two routers should be in the FULL state in order to exchange LSAs.

```
HQ# show ip ospf neighbor
Neighbor ID    Pri   State           Dead Time   Address        Interface
209.165.201.1  0     FULL/-         00:00:31   192.168.1.1   Serial0/0/0
```

Note Routers will establish the FULL state only with the DR and BDR, while the established state will be two-way with other routers.

Troubleshooting OSPF Routing Table Issues

This topic describes how to troubleshoot OSPF routing table issues.



Troubleshooting OSPF Routing Table Issues (Cont.)

```
HQ# show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Router ID 1.1.1.1
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    192.168.1.0 0.0.0.255 area 0
<output omitted>
```

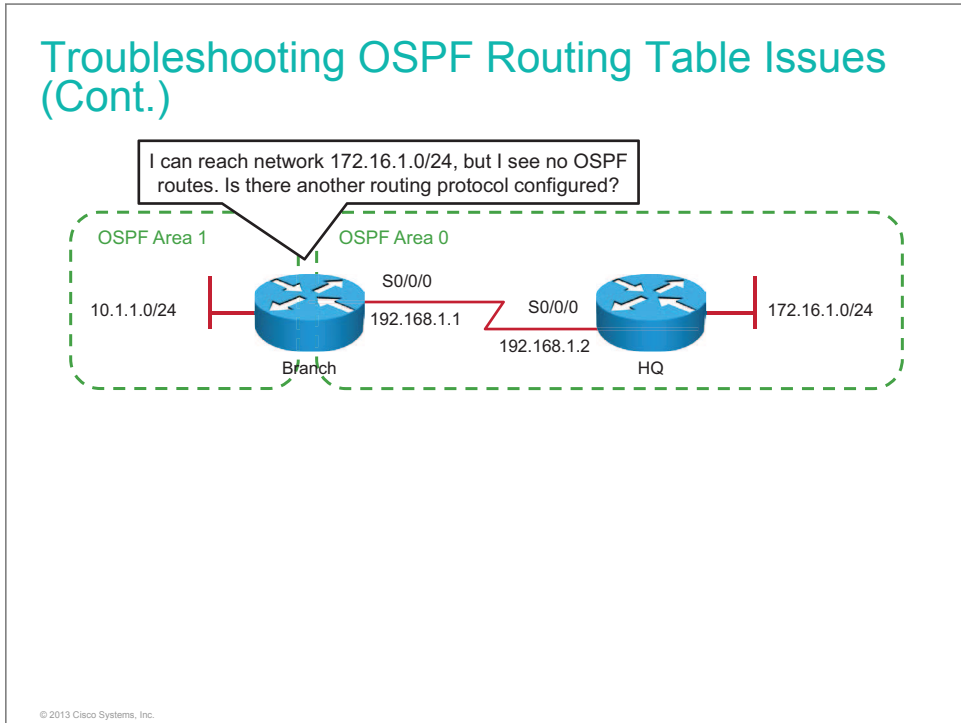
Headquarters has to advertise the network in order to reach it from the Branch router.

© 2013 Cisco Systems, Inc.

The Branch and Headquarters routers have their neighbor adjacency set up, but a ping test from the Branch router to a host in the 172.16.1.0/24 network is not successful. Checking the routing table of the Branch router leads you to the conclusion that there is a route missing to the destination network of 172.16.1.0/24.

You can use the **show ip protocols** command on the Headquarters router to verify if the 172.16.1.0/24 network is being advertised to OSPF neighbors.

In the example, the Headquarters router is configured to advertise the 172.16.1.0/24 network to the neighbor.



Troubleshooting OSPF Routing Table Issues (Cont.)

```
Branch# show ip route 172.16.1.0
Routing entry for 172.16.1.0/24
  Known via "eigrp 1", distance 90, metric 2297856, type internal
  Redistributing via eigrp 1
  Last update from 192.168.1.2 on Serial0/0/0, 00:00:39 ago
  Routing Descriptor Blocks:
  * 192.168.1.2, from 192.168.1.2, 00:00:39 ago, via Serial0/0/0
    Route metric is 2297856, traffic share count is 1
    Total delay is 25000 microseconds, minimum bandwidth is 1544 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 1
```

If several routing protocols are configured on routers, administrative distance will decide which protocol will be used.

© 2013 Cisco Systems, Inc.

When you have more than one routing protocol that is configured in a network, you may receive routing information about a network through an undesired routing protocol. Recall that routing protocol administrative distance influences which routes will be installed in the routing table. Although this does not affect connectivity, you may want to receive all routing information through the same routing protocol for the sake of easier troubleshooting and management. To verify which routing protocols are configured and their administrative distances, use the **show ip protocols** command:


```

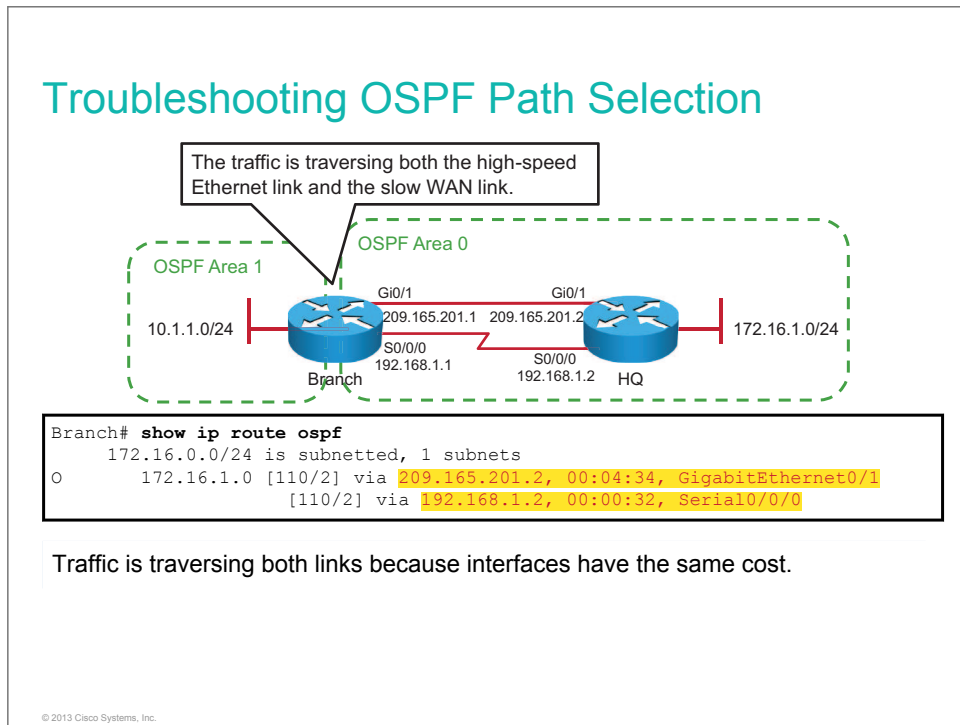
Branch# show ip protocols
Routing Protocol is "ospf 1"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is FILTER_OSPF
  Router ID 209.165.201.1
  It is an area border router
  Number of areas in this router is 2. 2 normal 0 stub 0 nssa
  Maximum path: 4
  Routing for Networks:
    10.1.1.0 0.0.0.255 area 1
    192.168.1.0 0.0.0.255 area 0
  Reference bandwidth unit is 100 mbps
  Routing Information Sources:
    Gateway          Distance      Last Update
    1.1.1.1           80           00:02:37
  Distance: (default is 110)
Routing Protocol is "eigrp 1"
  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 1
  EIGRP NSF-aware route hold timer is 240s
  Automatic network summarization is not in effect
  Maximum path: 4
  Routing for Networks:
    10.0.0.0
    192.168.1.0
  Routing Information Sources:
    Gateway          Distance      Last Update
    (this router)    90           00:12:09
    192.168.1.2      90           00:02:39
  Distance: internal 90 external 170

```

In the example, the route about the 172.16.1.0/24 network has been received through EIGRP and OSPF. However, because EIGRP with an administrative distance of 90 is more trustworthy than OSPF with an administrative distance of 110, the EIGRP route will be installed in the routing table.

Troubleshooting OSPF Path Selection

Incorrect path selection usually does not lead to a loss of connectivity. However, certain links in a network should not be utilized, if possible. This applies, for example, to backup WAN links, which can be charged by the amount of transferred data and can be expensive. This topic describes how to troubleshoot incorrect path selection that is done by OSPF.



When you have redundant paths that are available in a network, you have to make sure that traffic takes the desired path through the network. For example, you could have two locations that are connected via the primary, high-speed link and via the dial-up, low-speed link for backup purposes. In this case, you have to make sure that the backup link is used only when the primary link fails.

In the example, network 172.16.0.0/24 is reachable from the Branch router via the GigabitEthernet0/1 interface and the Serial0/0/0 interface. Because both interfaces have the same OSPF cost, load balancing across both links will be utilized. The reason for the same OSPF cost on both interfaces could be that someone manually changed the cost on the interfaces, or there is incorrect reference bandwidth when managing interfaces that are faster than 100 Mb/s. Recall that the OSPF cost is calculated as interface bandwidth divided by reference bandwidth, which is 100 Mb/s by default. For example, with two interfaces such as 1000 Mb/s and 100 Mb/s, both will have the same OSPF cost with a value of 1. In this case, it is either required to increase the reference bandwidth to 1000 Mb/s or to manually change the OSPF cost on an interface to reflect the actual bandwidth of the interface.

Use the **show ip ospf interface** command to verify the OSPF cost on an interface:

```
Branch# show ip ospf interface
GigabitEthernet0/1 is up, line protocol is up
  Internet Address 209.165.201.2/27, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
<output omitted>
Serial0/0/0 is up, line protocol is up
  Internet Address 192.168.1.2/24, Area 0
  Process ID 1, Router ID 1.1.1.1, Network Type POINT_TO_POINT, Cost: 1
```

Once you increase the OSPF cost on the Serial0/0/0 interfaces on both routers, only the preferred route will be installed in the routing table:

```
Branch(config)# interface Serial0/0/0
Branch(config-if)# ip ospf cost 10

Branch# show ip route ospf
      172.16.0.0/24 is subnetted, 1 subnets
O       172.16.1.0 [110/2] via 209.165.201.2, 00:14:31, GigabitEthernet0/1
```

The reference bandwidth can be changed from the default of 100 Mb/s. You can verify what the reference bandwidth is by using the **show ip ospf** command:

```
Branch# show ip ospf
<output omitted>
Cisco NSF helper support enabled
Reference bandwidth unit is 100 mbps
Area BACKBONE(0)
  Number of interfaces in this area is 4
  Area has no authentication
<output omitted>
```

Changed reference bandwidth means changed cost on the link. Make sure that all of the routers within the OSPF AS have the same reference bandwidth. One way to change it is by using the **auto-cost reference-bandwidth bandwidth_in_Mbits_per_second** command from the router OSPF configuration mode.

Summary

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

Summary

- OSPF routers traverse different OSPF states when adjacencies are being established.
- When troubleshooting connectivity problems due to OSPF, you should first verify OSPF neighbors.
- If OSPF areas on two routers do not match, an adjacency will not form.
- OSPF routing issues could stem from the fact that networks are not advertised, or there could be an ACL that is blocking a routing advertisement.
- When you have OSPF path selection issues, you have to make sure that traffic takes the desired path through the network by manipulating the OSPF cost.

© 2013 Cisco Systems, Inc.

Examining OSPFv3

OSPF is a widely used IGP. Upgrading the protocol to support IPv6 generated a number of significant changes to how the protocol behaves. Understanding the differences between OSPFv2 and OSPFv3 is required for the successful deployment and operation of an IPv6 network that is using OSPF for routing. This lesson describes OSPFv3, which is the IPv6-capable version of the OSPF routing protocol, and describes its operations and implementation.

Objectives

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

- Introduce OSPFv3 and describe how it is different from OSPF for IPv4
- Configure OSPFv3
- Verify the OSPFv3 configuration

OSPFv3 Key Characteristics

This topic describes basic facts about OSPFv3 and compares it to OSPFv2.

OSPFv3 Key Characteristics

- OSPFv3 is an implementation of the OSPF routing protocol for IPv6.
- OSPFv2 (for IPv4) and OSPFv3 (for IPv6) run independently on the router.
- OSPFv3 has the same key capabilities as OSPFv2 for IPv4 networks:
 - Multiarea network design with ABRs that segment the network
 - SPF algorithm for optimal path calculation

© 2013 Cisco Systems, Inc.

OSPFv3 is a complete rewrite of the OSPFv2 protocol to support IPv6. The foundation remains, for the most part, the same as in OSPFv2.

Note OSPFv3 and OSPFv2 run independently on a router.

The OSPFv3 metric is still based on interface costing.

The packet types and neighbor discovery mechanisms are the same in OSPFv3 as they are for OSPFv2.

LSAs are still flooded throughout an OSPF domain, and many of the LSA types are the same, although a few have been renamed or newly created.

Note OSPFv3 is defined in RFC 5340.

OSPFv3 Key Characteristics (Cont.)

- The router ID is still a 32-bit number that is based on the IPv4 address of the router. If there is no IPv4 address that is present on the router, you are prompted to configure it using the **router-id** command.
- Adjacencies and next-hop attributes use link-local addresses.
- IPv6 is used for transport of the LSA.
- OSPFv3 is enabled per link, not per network.
- OSPFv3 communicates using IPv6 multicast addresses.

© 2013 Cisco Systems, Inc.

In OSPFv2, the router ID is derived from the “highest” IPv4 address of an existing router. It is a general practice to set a loopback interface on the router for the purpose of maintaining the router ID or setting it administratively in the routing process configuration.

In OSPFv3, the OSPF process no longer requires an IPv4 address for the router ID, but it does require a 32-bit number to be set. This 32-bit number is entered as four octets that are separated by dots [.] and looks like an IP address. The router ID is set using the **router-id** *router_id* command. If it is not manually set, then the router ID will be the same as the highest configured loopback address on the router. If there is also no loopback configured on the device, then it will use the highest IP address of a physical interface. On devices that do not have IPv4 addresses, you have to set the router ID manually.

OSPFv3 adjacencies use link-local addresses to communicate. Router next-hop attributes are neighboring router link-local addresses. Because link-local addresses have the same prefix, OSPF needs to store the information about the outgoing interface.

OSPFv3 uses IPv6 for transport of LSAs. IPv6 protocol number 89 is used.

OSPFv3 is enabled per link and identifies which networks (prefixes) are attached to this link for determining prefix reachability propagation and the OSPF area.

OSPFv3 takes advantage of IPv6 multicasting by using FF02::5 for all OSPF routers and FF02::6 for the OSPF DR and OSPF BDR.

OSPFv3 Configuration

This topic provides a basic OSPFv3 configuration example.

OSPFv3 Configuration

```
Branch(config)# ipv6 router ospf 99
Branch(config-rtr)# router-id 1.1.1.1
Branch(config-rtr)# interface GigabitEthernet0/1
Branch(config-rtr)# exit
Branch(config-if)# ipv6 address 2001:DB8:D1A5:C900::1/64
Branch(config-if)# ipv6 ospf 99 area 0
```

OSPFv3 configuration on the Branch router

© 2013 Cisco Systems, Inc.

The example shows a multiarea OSPFv3 network of two routers. Both routers are in Area 0. Additionally, the Headquarters router connects to Area 1.

Configuring OSPFv3 Commands

| Command | Description |
|--|--|
| <code>ipv6 router ospf process-id</code> | Enables OSPF for IPv6 router configuration mode. The <i>process-id</i> is an internal identification. It is locally assigned and can be a positive integer from 1 to 65,535. |
| <code>ipv6 ospf process-id area area-id</code> | Enables OSPFv3 on an interface and assigns it to the specified area. |
| <code>router-id router_id</code> | This command is executed in OSPF router configuration mode to statically configure a router ID, which is the name for the router within the OSPFv3 process. |

Note By default, IPv6 traffic forwarding is disabled. Make sure to use the `ipv6 unicast-routing` command to enable it.

OSPFv3 Configuration (Cont.)

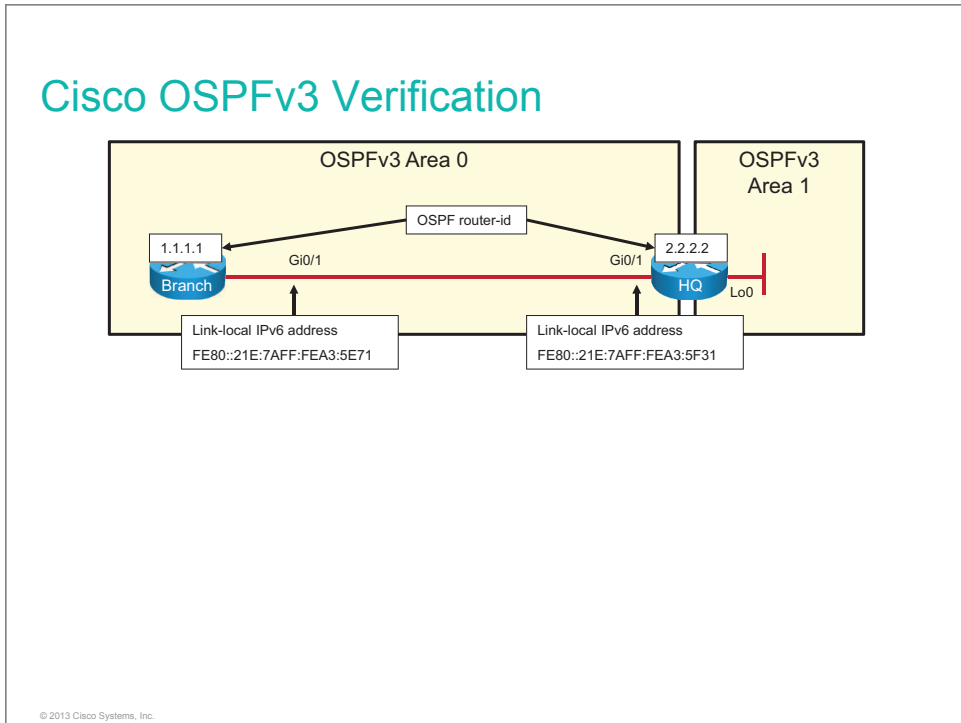
```
HQ(config)# ipv6 router ospf 99
HQ(config-rtr)# router-id 2.2.2.2
HQ(config-rtr)# exit
HQ(config)# interface Loopback0
HQ(config-if)# ipv6 address 2001:DB8:AC10:100::64/64
HQ(config-if)# ipv6 ospf 99 area 0.0.0.1
HQ(config-if)# exit
HQ(config)# interface GigabitEthernet0/1
HQ(config-if)# ipv6 address 2001:DB8:D1A5:C900::2/64
HQ(config-if)# ipv6 ospf 99 area 0.0.0.0
```

OSPFv3 configuration on the Headquarters router

© 2013 Cisco Systems, Inc.

OSPFv3 Configuration Verification

This topic describes commands for troubleshooting OSPFv3.



Cisco OSPFv3 Verification (Cont.)

```
Branch# show ipv6 ospf interface
GigabitEthernet0/1 is up, line protocol is up
  Link Local Address FE80::21E:7AFF:FEA3:5E71, Interface ID 5
  Area 0.0.0.0, Process ID 99, Instance ID 0, Router ID 1.1.1.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.1, local address FE80::21E:7AFF:FEA3:5E71
  Backup Designated router (ID) 2.2.2.2, local address FE80::21E:
7AFF:FEA3:5F31
  <output omitted>
  Adjacent with neighbor 2.2.2.2 (Backup Designated Router)
  <output omitted>
```

OSPFv3 interface information on the Branch router

© 2013 Cisco Systems, Inc.

The **show ipv6 ospf interface** command displays interfaces with enabled OSPFv3. The example shows the OSPFv3 area, process ID, and router ID. Additionally, the output displays the adjacent OSPF neighbor that is connected to the interface.

Cisco OSPFv3 Verification (Cont.)


```
Branch# show ipv6 ospf
Routing Process "ospfv3 99" with ID 1.1.1.1
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
<output omitted>
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Reference bandwidth unit is 100 mbps
Area BACKBONE(0.0.0.0)
  Number of interfaces in this area is 1
<output omitted>
```

OSPFv3 general information on the Branch router

© 2013 Cisco Systems, Inc.

The **show ipv6 ospf** command displays OSPFv3 general information. The example shows the OSPFv3 process ID, router ID, timers, areas that are configured, and reference bandwidth.

Cisco OSPFv3 Verification (Cont.)



```
Branch# show ipv6 ospf neighbor
OSPFv3 Router with ID (1.1.1.1) (Process ID 99)
Neighbor ID  Pri  State           Dead Time   Interface ID  Interface
2.2.2.2      0  FULL/-        00:00:37   6             Gi0/1
```

OSPFv3 neighbor information on the Branch router

© 2013 Cisco Systems, Inc.

The **show ipv6 ospf neighbor** command that is issued on the Branch router shows that it has one OSPFv3 neighbor. This neighbor has router ID 2.2.2.2 and is available through GigabitEthernet0/1 on the Branch router.

Cisco OSPFv3 Verification (Cont.)

```
Branch# show ipv6 route ospf
IPv6 Routing Table - default - 10 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
       IA - ISIS interarea, IS - ISIS summary, D - EIGRP, EX - EIGRP external
       ND - ND Default, NDP - ND Prefix, DCE - Destination, NDR - Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
OI 2001:DB8:AC10:100::64/64 [110/64]
   via FE80::21E:7AFF:FEA3:5F31, GigabitEthernet0/1
```

OSPFv3 route information on the Branch router

Issuing the **show ipv6 route ospf** command on the Branch router returns IPv6 routes that are available to this router. It has one route to the Loopback0 interface on the Headquarters router. This route is marked with "OI," which stands for OSPF Inter-area route. In the IPv6 world, routers use link-local addresses for next-hop communication.

Summary

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

Summary

- OSPFv3 for IPv6 supports the same basic mechanisms that OSPFv2 for IPv4 supports, including the use of areas to provide network segmentation and LSAs to exchange routing updates.
- OSPFv3 is configured per-interface on Cisco routers.
- To verify OSPFv3 configuration, use commands that are similar to those that are used for OSPFv2. Instead of the keyword **ip**, use the keywords **ipv6**: **show ipv6 ospf interface** and **show ipv6 ospf**.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Module Summary

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

Module Summary

- OSPF uses three data structures: neighbors table, topology table, and routing table.
- Multiarea OSPF design enables segmentation of a network to limit the propagation of LSAs inside an area and to make routing tables smaller.
- When troubleshooting connectivity problems due to OSPF, you should first verify OSPF neighbors.
- OSPFv3 for IPv6 supports the same basic mechanisms that OSPFv2 for IPv4 does, including the use of areas to provide network segmentation and LSAs to exchange routing updates.

© 2013 Cisco Systems, Inc.

References

There are no additional references for this module.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Questions

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

1. All of these tables are maintained by a link-state routing protocol except which one? (Source: OSPF Overview)
 - A. routing
 - B. topology
 - C. update
 - D. neighbor
2. Match each table to its function. (Source: OSPF Overview)

| | | |
|-------------|--------------------------|--------------------|
| A. neighbor | <input type="checkbox"/> | stores LSAs |
| B. topology | <input type="checkbox"/> | stores best paths |
| C. routing | <input type="checkbox"/> | stores adjacencies |
3. Which OSPF packet helps form neighbor adjacencies? (Source: OSPF Overview)
 - A. exchange packet
 - B. hello packet
 - C. neighbor discovery packet
 - D. adjacency packet
4. Which criterion does SPF use to determine the best path? (Source: OSPF Overview)
 - A. lowest delay
 - B. highest bandwidth
 - C. lowest total cost of the route
 - D. total bandwidth of the route

5. What are the three possible benefits of multiarea design in OSPF? (Choose three.) (Source: Multiarea OSPF Implementation)
- A. reduced amount of LSA flooding
 - B. reduced number of SPF calculations
 - C. reduced size of the neighbor table
 - D. reduced size of the routing table
6. Which two commands are required for a basic OSPF configuration? (Choose two.) (Source: Multiarea OSPF IPv4 Implementation)
- A. **network ip-address mask area area-id**
 - B. **network ip-address wildcard-mask area area-id**
 - C. **router ospf process-id**
 - D. **ip router ospf**
7. Which OSPF **show** command describes a list of OSPF adjacencies? (Source: Multiarea OSPF IPv4 Implementation)
- A. **show ip ospf interface**
 - B. **show ip ospf**
 - C. **show ip route**
 - D. **show ip ospf neighbor**
8. Which two interfaces will have the lowest OSPF cost by default? (Choose two.) (Source: Troubleshooting Multiarea OSPF)
- A. Fast Ethernet
 - B. T1
 - C. Gigabit Ethernet
 - D. E1
9. Which command is used to enable OSPFv3? (Source: Examining OSPFv3)
- A. (config-if)#**ipv6 ospf process_id area area_id**
 - B. (config-rtr)#**ipv6 ospf process_id area area_id**
 - C. (config-if)#**network ipv6_address wildcard_mask area area_id**
 - D. (config-if)#**network ipv6_address wildcard_mask**
10. Which command is used to display OSPFv3-related interface information? (Source: Examining OSPFv3)
- A. **show ipv6 ospf interface**
 - B. **show ipv6 interface**
 - C. **show ipv6 ospf**
 - D. **show ipv6 interface**

Answer Key

1. C
2. A. routing stores best paths
B. topology stores LSAs
C. neighbor stores adjacencies
3. B
4. C
5. A, B, D
6. B, C
7. D
8. A, C
9. A
10. A

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate:
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Wide-Area Networks

WANs are most often fee-for-service networks, providing the means for users to access resources across a wide geographical area. Some services are considered Layer 2 connections between your remote locations, typically provided by a telco over its WAN switches. Some of these technologies include a serial point-to-point (leased line) connection and Frame Relay connections.

Other connections leverage the Internet infrastructure, a Layer 3 alternative, to interconnect the remote locations of an organization. To provide security across the public Internet, you may implement a VPN solution.

This lesson introduces the components of a VPN solution for WAN connectivity, explains how to configure a PPP connection, and describes Frame Relay operation, configuration, and troubleshooting. The lesson also introduces the GRE tunneling protocol.

Objectives

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

- Explain WAN technologies
- Configure a serial connection
- Describe Frame Relay technology and its basic configuration
- Describe VPN solutions
- Configure GRE tunnels

Do Not Duplicate.
Post beta, not for release.

Understanding WAN Technologies

Overview

As an enterprise grows beyond a single location, it needs to interconnect LANs in various locations using a WAN. There are several technologies that are involved in the functioning of WANs, including hardware devices and software functions. This lesson describes the functions and characteristics of WANs and contrasts them with LANs. The lesson also explores how WANs relate to the OSI reference model in their design and function and which major hardware components are typically used in WAN environments.

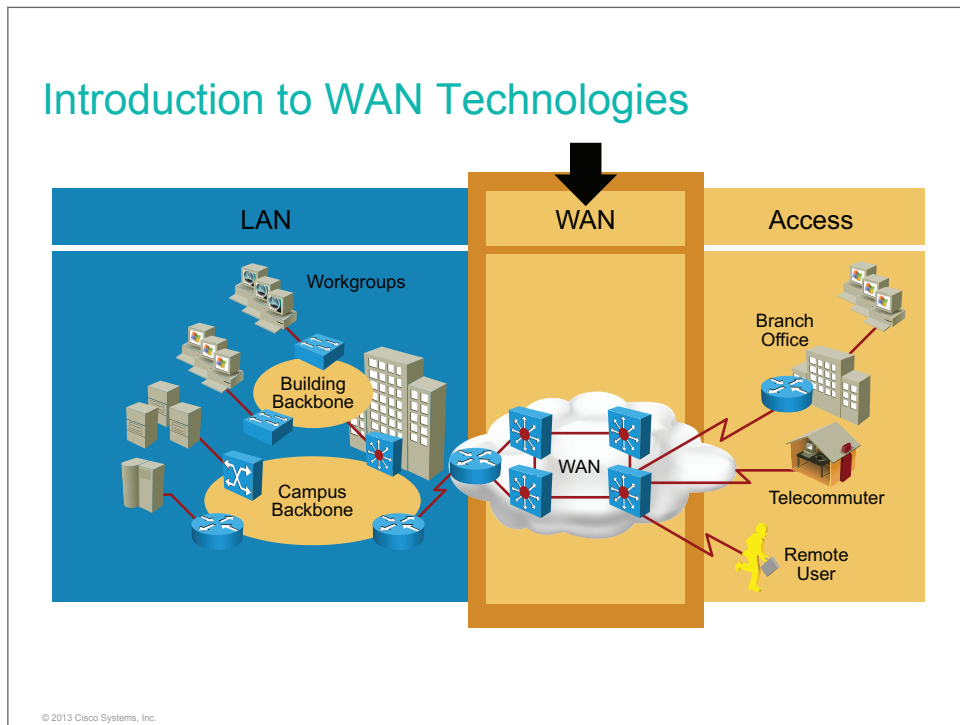
Objectives

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

- Explain WAN technologies
- List the WAN devices and their functions in a WAN environment
- Describe various options for WAN cabling
- List various Layer 2 WAN protocols
- Describe the major WAN communication link options

Introduction to WAN Technologies

A WAN is a data communications network that operates beyond the geographic scope of a LAN. This topic describes the characteristics of a WAN.




WANs use facilities that are provided by a service provider, or carrier, such as a telephone or cable company. They connect the locations of an organization to each other, to locations of other organizations, to external services, and to remote users. WANs carry various traffic types such as voice, data, and video.

These are three major characteristics of WANs:

- WANs generally connect devices that are separated by a broader geographical area than a LAN can serve.
- WANs use the services of carriers such as telcos, cable companies, satellite systems, and network providers.
- WANs use serial connections of various types to provide access to bandwidth over large geographic areas.

Introduction to WAN Technologies (Cont.)



There are several reasons why WANs are necessary in a communications environment.

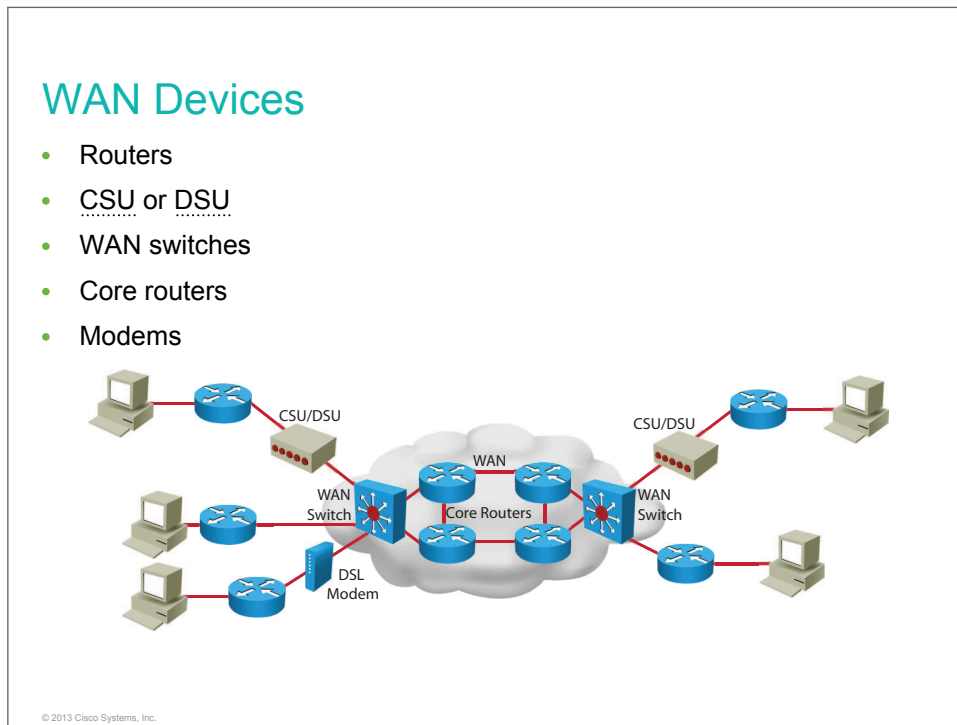
LAN technologies provide speed and cost efficiency for the transmission of data in organizations in relatively small geographic areas. However, there are other business needs that require communication among remote users, including the following:

- People in the regional or branch offices of an organization need to be able to communicate and share data.
- Organizations often want to share information with other organizations across large distances.
- Employees who travel on company business frequently need to access information that resides on their corporate networks.

Because it is not feasible to connect computers across a country or around the world in the same way that computers are connected in a LAN environment with cables, different technologies have evolved to support this need. Increasingly, the Internet is being used as an inexpensive alternative to an enterprise WAN for some applications.

WAN Devices

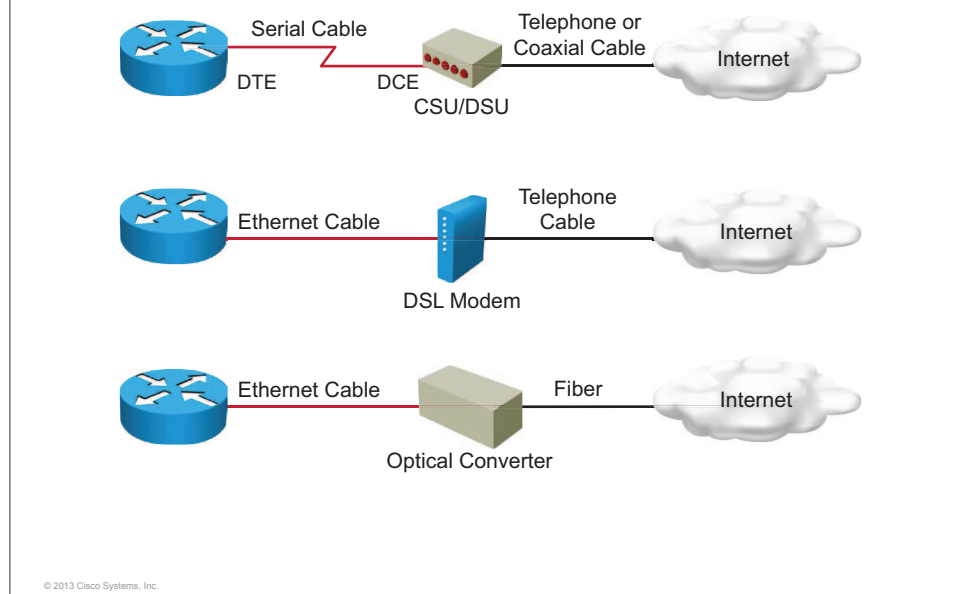
There are several devices that operate at the physical layer in a WAN. This topic describes those devices and their functions in a WAN environment.



WANs use numerous types of devices that are specific to WAN environments, including the following:

- **CSU/DSU:** Digital lines, such as T1 or T3 carrier lines, require a CSU and a DSU. The two are often combined into a single piece of equipment that is called the CSU/DSU. The CSU provides termination for the digital signal and ensures connection integrity through error correction and line monitoring. The DSU converts the T-carrier line frames into frames that the LAN can interpret and vice versa.
- **Modem:** A modem modulates an analog carrier signal to encode digital information and also demodulates the carrier signal to decode the transmitted information.
- **WAN switch:** A WAN switch is a multiport internetworking device that is used in carrier networks.
- **Router:** A router provides internetworking and WAN access interface ports that are used to connect to the service provider network. These interfaces may be serial connections or other WAN interfaces. With some types of WAN interfaces, an external device such as a CSU/DSU or modem (analog, cable, or DSL) is required to connect the router to the local POP of the service provider.
- **Core router:** A core router resides within the middle or backbone of the WAN, rather than at its periphery. To fulfill this role, a router must be able to support multiple telecommunications interfaces of the highest speed in use in the WAN core. It also must be able to forward IP packets at wire speed on all of these interfaces. The router must support the routing protocols that are being used in the core.

WAN Devices (Cont.)



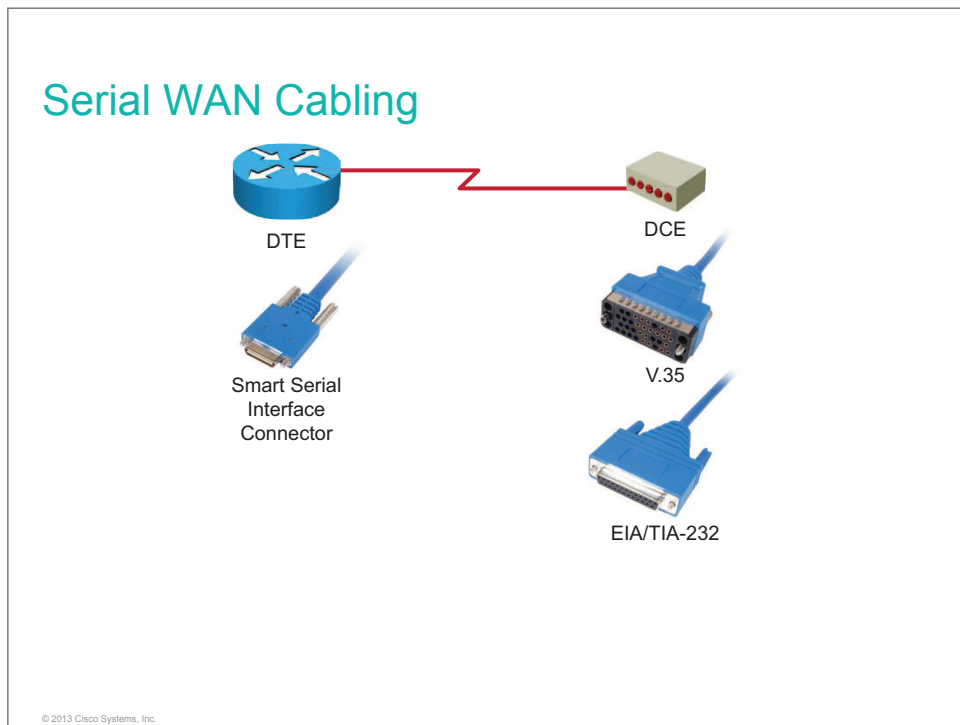
Devices on the subscriber premises are referred to as CPE. The subscriber owns the CPE or leases the CPE from the service provider. A copper or fiber cable connects the CPE to the nearest exchange or CO of the service provider. This cabling is often called the local loop or “last mile.”

These are three of the numerous WAN connection types:

- **CSU/DSU:** The router connects to the CSU/DSU with a serial cable (for example, V.35). The CSU/DSU then connects to the service provider infrastructure using a telephone or coaxial cable (for example, a T1 or E1 line). Devices that put data on the local loop are called DCE (or CSU/DSU, in this example). The customer devices that pass the data to the DCE are called DTE (or the router, in this example). The DCE primarily provides an interface for the DTE into the communication link on the WAN cloud. The CSU/DSU can also be implemented as a module within the router, in which case, you do not need a serial cable.
- **DSL modem:** A router connects to a DSL modem using an Ethernet cable. The modem connects to the service provider network using a telephone cable. The modem can also be implemented as a router module, in which case, you do not need an Ethernet cable.
- **Optical fiber converter:** An optical fiber converter is where a fiber-optic link terminates and where optical signals are converted into electrical signals. The converter can also be implemented as a router or switch module.

Serial WAN Cabling

This topic describes the cabling that is available for serial WAN connections.



WAN physical layer protocols describe how to provide electrical, mechanical, operational, and functional connections for WAN services. The WAN physical layer also determines the interface between the DTE and DCE. The DTE and DCE interfaces on Cisco routers use various physical layer protocols, including the following:

- **V.35:** V.35 is the ITU-T standard for synchronous communications between an NAD and a packet network. Originally specified to support data rates of 48 kb/s, it now supports speeds of up to 2.048 Mb/s using a 34-pin rectangular connector.
- **EIA/TIA-232:** EIA/TIA-232 allows signal speeds of up to 64 kb/s on a 25-pin D-connector over short distances.

These protocols establish the codes and electrical parameters that the devices use to communicate with each other. The method of facilitation that is used by the service provider largely determines the choice of protocol. There are many other protocols in addition to these two protocols.

When you order a cable, you receive a shielded serial transition cable that has the appropriate connector for the standard that you specify. The router end of the shielded serial transition cable has a DB-60 connector that connects to the DB-60 port on a serial WIC. Because five different cable types are supported with this port, the port is sometimes called a five-in-one serial port. The other end of the serial transition cable is available with the connector that is appropriate for the standard that you specify. The documentation for the device to which you want to connect should indicate the standard for this device. To support higher densities in a smaller form factor, Cisco introduced a Smart Serial cable. The serial end of the Smart Serial cable is a 26-pin connector. It is much smaller than the DB-60 connector that is used to connect to a five-in-one serial port. These transition cables support the same five serial standards, are available in either a DTE or DCE configuration, and are used with 2-port serial connections as well as with 2-port asynchronous and synchronous WICs.

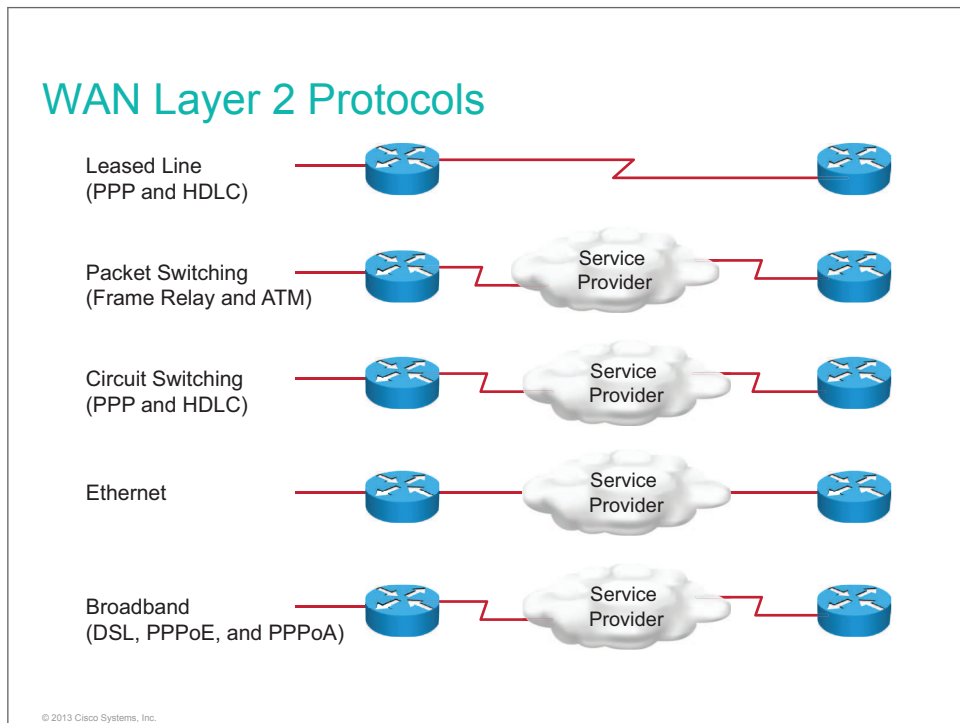
Usually, the cable itself is marked DTE on one end and DCE on the other end.

Your CPE, in this case a router, is the DTE. The data DCE, commonly a modem or a CSU/DSU, is the device that is used to convert the user data from the DTE into a form that is acceptable to the WAN service provider. The synchronous serial port on the router is configured as DTE or DCE, depending on the attached cable, which is ordered as either DTE or DCE to match the router configuration. If the port is configured as DTE (which is the default setting), it will require external clocking from the DCE device.

Do Not Duplicate.
Post beta, not for release.

WAN Layer 2 Protocols

In addition to physical layer devices, WANs require data link layer protocols to establish the link across the communication line from the sending to the receiving device. This topic lists and explains different Layer 2 WAN protocols.



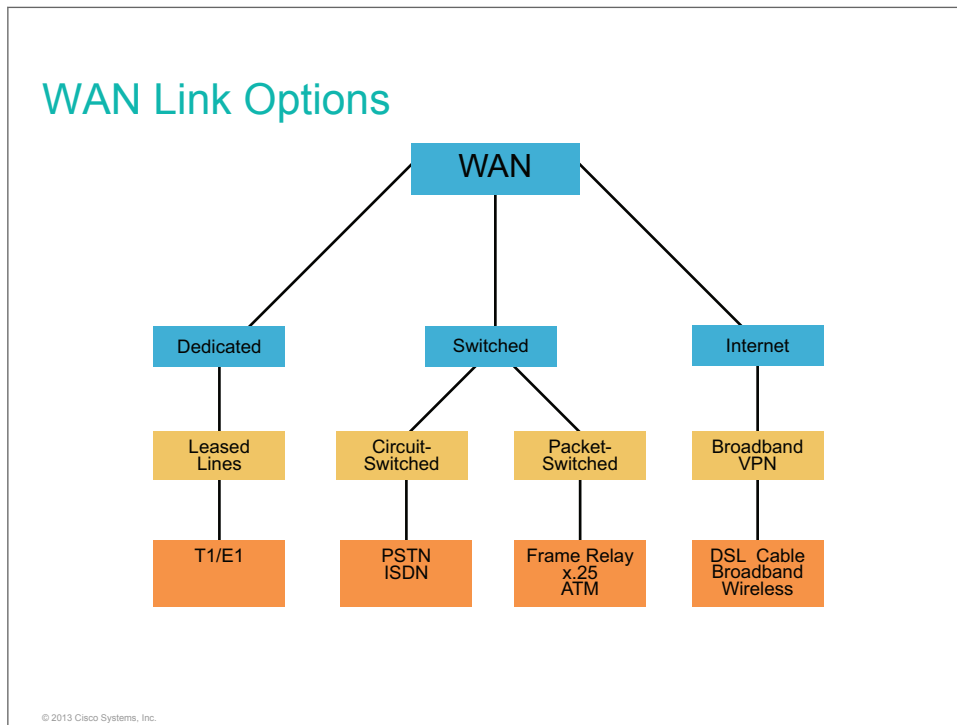
On each WAN connection, data is encapsulated into frames before it crosses the WAN link. To ensure that the correct protocol is used, you must configure the appropriate Layer 2 encapsulation type. The choice of Layer 2 protocol depends on the WAN technology and the communicating equipment. These are typical WAN protocols:

- **HDLC:** HDLC is the Cisco default encapsulation type on point-to-point connections, dedicated links, and circuit-switched connections. You typically use HDLC when two Cisco devices are communicating across a point-to-point connection. HDLC is a bit-oriented, synchronous, data-link layer protocol.
- **PPP:** PPP provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. PPP was designed to work with several network layer protocols, including IP. PPP also has built-in security mechanisms, such as PAP and CHAP.
- **Frame Relay:** This protocol is an industry-standard, switched, data-link layer protocol. It defines how connections between DTE and DCE are maintained for remote terminal access and computer communications in public data networks.
- **ATM:** ATM is the international standard for cell relay in which multiple service types, such as voice, video, and data, are conveyed in fixed-length (53-byte) cells. ATM, a cell-switched technology, uses fixed-length cells that allow processing to occur in hardware, and it reduces transit delays. ATM is designed to take advantage of high-speed transmission media such as T3, E3, and SONET.

- **Ethernet:** The emergence of Ethernet as a viable method of providing both point-to-point and multipoint services has been driven by an abundance of new fiber deployment to business areas. Enterprise customers with years of Ethernet experience in the campus have developed such a comfort level and confidence with Ethernet that they are now asking their service providers for Ethernet as an access option. Ethernet may be the most scalable transport technology ever developed—starting at 10 Mb/s, it has now evolved to 100 Gb/s.
- **Broadband:** Broadband in data communications refers to data transmission in which multiple pieces of data are sent simultaneously to increase the effective rate of transmission, regardless of the actual data rate. In network engineering, this term refers to transmission methods in which two or more signals share a medium, such as the following technologies:
 - **DSL, PPPoE, and PPPoA:** DSL, PPPoE, and PPPoA are a family of technologies that provide digital data transmission over the wires of a local telephone network.
 - **Cable Ethernet:** A cable modem is a type of modem that provides access to a data signal that is sent over the cable TV infrastructure.

WAN Link Options

There are a number of ways in which WANs are accessed, depending on the data transmission requirements for the WAN. This topic describes the major WAN communication link options.



Many options for implementing WAN solutions are currently available. They differ in technology, speed, and cost. Familiarity with these technologies is an important part of network design and evaluation.

WAN connections can be either over a private infrastructure or over a public infrastructure such as the Internet.

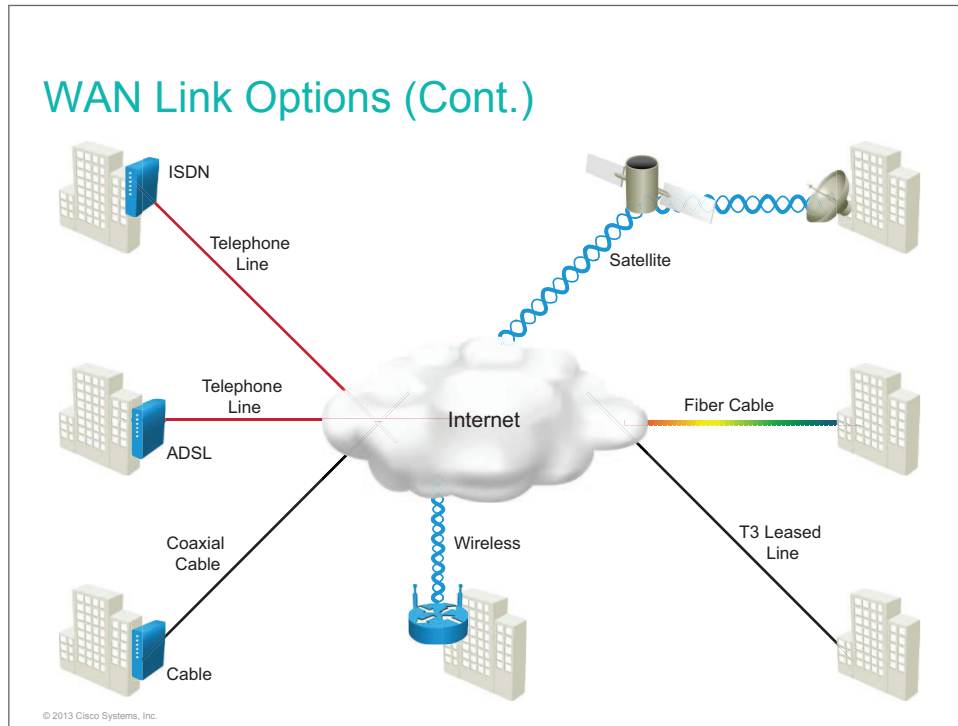
Private WAN Connection Options

Private WAN connections include dedicated and switched communication link options:

- **Dedicated communication links:** When permanent dedicated connections are required, point-to-point lines are used with various capacities that are limited only by the underlying physical facilities and the willingness of users to pay for these dedicated lines. A point-to-point link provides a pre-established WAN communications path from the customer premises through the provider network to a remote destination. Point-to-point lines are usually leased from a carrier and are also called leased lines.
- **Switched communication links:** Switched communication links can be either circuit-switched or packet-switched.
 - **Circuit-switched communication links:** Circuit switching dynamically establishes a dedicated virtual connection for voice or data between a sender and a receiver. Before communication can start, it is necessary to establish the connection through the network of the service provider. Examples of circuit-switched communication links are analog dialup (PSTN) and ISDN.
 - **Packet-switched communication links:** Many WAN users do not make efficient use of the fixed bandwidth that is available with dedicated, switched, or permanent circuits because the data flow fluctuates. Communications providers have data networks that are available to more appropriately service these users. In packet-switched networks, the data is transmitted in labeled frames, cells, or packets. Packet-switched communication links include Frame Relay, ATM, X.25, and Metro Ethernet.

Public WAN Connection Options

Public connections use the global Internet infrastructure. Until recently, the Internet was not a viable networking option for many businesses because of the significant security risks and lack of adequate performance guarantees in an end-to-end Internet connection. With the development of VPN technology, however, the Internet is now an inexpensive and secure option for connecting to teleworkers and remote offices where performance guarantees are not critical. Internet WAN connection links are through broadband services such as DSL, cable modem, and broadband wireless, and they are combined with VPN technology to provide privacy across the Internet. Broadband connection options are typically used to connect telecommuting employees to a corporate site over the Internet.



ISPs use several different WAN technologies to connect their subscribers. The connection type that is used on the local loop, or last mile, may not be the same as the WAN connection type that is employed within the ISP network or between various ISPs.

Each of these technologies provides advantages and disadvantages for the customer. Not all technologies are available at all locations. When a service provider receives data, it must forward this data to other remote sites for final delivery to the recipient. These remote sites connect either to the ISP network or pass from ISP to ISP to the recipient. Long-range communications are usually those connections between ISPs or between branch offices in very large companies.

Summary

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

Summary

- There are three major characteristics of a WAN:
 - Connection of devices that are separated by wide geographical distances
 - Use of the services of carriers such as telcos, cable companies, satellite systems, and network providers
 - Use of serial connections of various types to access bandwidth over large geographic areas
- The major types of devices that are used for WAN access environments include routers, modems (CSUs/DSUs), and other networking devices such as WAN switches.
- WAN physical layer protocols establish the codes and electrical parameters that the devices use to communicate with each other. Choosing a protocol is largely determined by the method of facilitation of the service provider.

© 2013 Cisco Systems, Inc.

Summary (Cont.)

- WANs require data-link layer protocols to establish the link across the communication line from the sending to the receiving device.
- WAN connections can be either over a private infrastructure or over a public infrastructure, such as the Internet. Private WAN connections include both dedicated and switched communication link options.

© 2013 Cisco Systems, Inc.

Configuring Serial Encapsulation

Overview

One of the most common types of WAN connection is the point-to-point connection. Point-to-point connections are used to connect LANs to service provider WANs and to connect LAN segments within an enterprise network. A LAN-to-WAN, point-to-point connection is also referred to as a serial connection or leased-line connection. WAN services are typically leased from a service provider. Some WAN services operate as Layer 2 connections between your remote locations and are typically provided by a telco provider over its WAN switches.

PPP emerged as an encapsulation protocol for transporting IP traffic over point-to-point (leased line) serial connections. PPP encapsulation has been carefully designed to retain compatibility with most commonly used supporting hardware. This lesson describes the operation, configuration, and verification of PPP.

Objectives

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

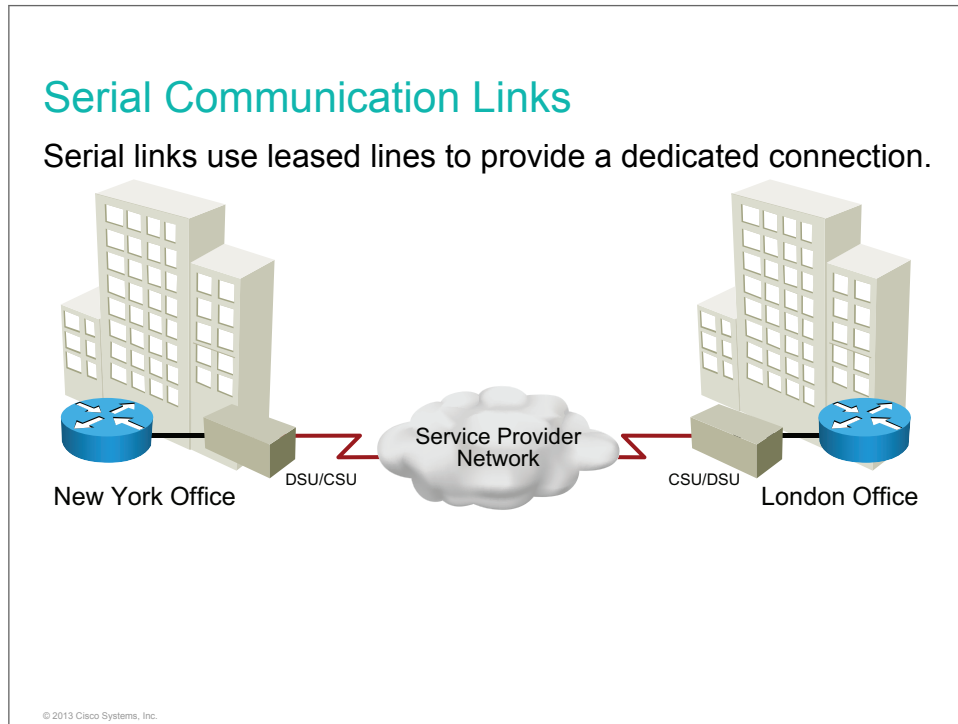
- Explain the idea behind serial links
- Configure a serial interface
- Describe HDLC protocols
- Describe Point-to-Point Protocol
- Configure a serial link with PPP encapsulation
- Describe PAP authentication
- Describe CHAP authentication
- Configure CHAP for PPP
- Configure and verify authentication for PPP

- Troubleshoot serial connections

Do Not Duplicate.
Post beta, not for release.

Serial Communication Links

A point-to-point (or serial) communication link provides a single, established WAN communications path from the customer premises through a carrier network to a remote network. This topic describes the functions of point-to-point technology.



When permanent dedicated connections are required, a point-to-point link is used to provide a pre-established WAN communications path from the customer premises through the provider network to a remote destination. A serial line can connect two geographically distant sites, such as a corporate office in New York and a regional office in London. Point-to-point lines are usually leased from a carrier and are therefore often called leased lines. For a point-to-point line, the carrier dedicates fixed transport capacity and facility hardware to the line that is leased by the customer. The carrier will, however, still use multiplexing technologies within the network.

Leased lines are a frequently used type of WAN access, and they are generally priced based on the bandwidth that is required and the distance between the two connected points.

Point-to-point links are usually more expensive than shared services such as Frame Relay. The cost of leased-line solutions can become significant when they are used to connect many sites over increasing distances. However, there are times when the benefits outweigh the cost of the leased line. The dedicated capacity removes latency or jitter between the endpoints. Constant availability is essential for some applications such as VoIP or video over IP.

A router serial port is required for each leased-line connection. If the underlying network is based on the North American (T-carrier) or European (E-carrier) technologies, the leased line connects to the network of the carrier through a CSU/DSU. The purpose of the CSU/DSU is to provide a clocking signal to the customer equipment interface from the DSU and terminate the channelized transport media of the carrier on the CSU. The CSU also provides diagnostic functions such as a loopback test. Most T1 or E1 TDM interfaces on current routers include approved CSU/DSU capabilities.

Leased lines provide permanent dedicated capacity and are used extensively for building WANs. Leased lines have been the traditional connection of choice but have a number of disadvantages. Leased lines have a fixed capacity. However, WAN traffic is often variable and leaves some of the capacity unused. In addition, each endpoint needs a separate physical interface on the router, which increases equipment costs. Any change to the leased line generally requires a site visit by the carrier personnel.

Serial Communication Links (Cont.)

- Typical WAN speeds for the U.S.:
 - T1 = (1.544 Mb/s)
 - T2 = 4 T1 lines (6 Mb/s)
 - T3 = 28 T1 lines (45 Mb/s)
 - T4 = 168 T1 lines (275 Mb/s)
- Typical WAN speeds for Europe:
 - E1 = (2 Mb/s)
 - E2 = 128 E0 lines (8 Mb/s)
 - E3 = 16 E1 lines (34 Mb/s)
 - E4 = 64 E1 lines (140 Mb/s)


© 2013 Cisco Systems, Inc.

Bandwidth refers to the rate at which data is transferred over the communication link. The underlying carrier technology depends on the bandwidth that is available. There is a difference in bandwidth points between the North American T-carrier specification and the E-carrier system.

Leased lines are available in different capacities and are generally priced based on the bandwidth that is required and the distance between the two connected points.

Configuration of a Serial Interface

Configuration of a Serial Interface




```
RouterA(config)# interface Serial 0/0/0
RouterA(config-if)# clockrate 64000
RouterA(config-if)# bandwidth 64
RouterA(config-if)# no shutdown
```

Configuration of serial interface on Router A

© 2013 Cisco Systems, Inc.

Configuration of a Serial Interface (Cont.)



```
RouterB(config)# interface Serial 0/0/0
RouterB(config-if)# bandwidth 64
RouterB(config-if)# no shutdown
```

Configuration of serial interface on Router B

© 2013 Cisco Systems, Inc.

Configuring Serial Interface Commands

| Command | Description |
|--|--|
| <code>interface Serial interface_number</code> | Enters serial interface configuration mode for the specified interface |
| <code>bandwidth bandwidth</code> | Sets interface bandwidth metric (in kb/s) |

| Command | Description |
|------------------------------------|---|
| <code>clock rate clock_rate</code> | Sets the interface clock rate in b/s—This command used on DCE interfaces only |
| <code>no shutdown</code> | Enables the serial interface |

Serial interfaces require a clock signal to control the timing of the communications. In most environments, a DCE device such as a CSU/DSU will provide the clock. By default, Cisco routers are DTE devices, but they can be configured as DCE devices.

Note The serial cable that is attached determines the DTE or DCE mode of the Cisco router. Choose the cable to match the network requirement.

Each connected serial interface must have an IP address and subnet mask to route IP packets.

Note A common misconception for students that are new to networking and Cisco IOS Software is to assume that the **bandwidth** command changes the physical bandwidth of the link. The **bandwidth** command modifies only the bandwidth metric that is used by routing protocols such as EIGRP and OSPF. Sometimes, a network administrator changes the bandwidth value to have more control over the chosen outgoing interface.

Configuration of a Serial Interface (Cont.)

```
RouterB# show controllers Serial 0/0/0
Interface Serial0/0/0
Hardware is GT96K
DTE V.35idb at 0x4753C1F4, driver data structure at 0x47543900
wic_info 0x47543F2C
Physical Port 1, SCC Num 1
<text omitted>
```

Configuration of serial interface on Router B

© 2013 Cisco Systems, Inc.

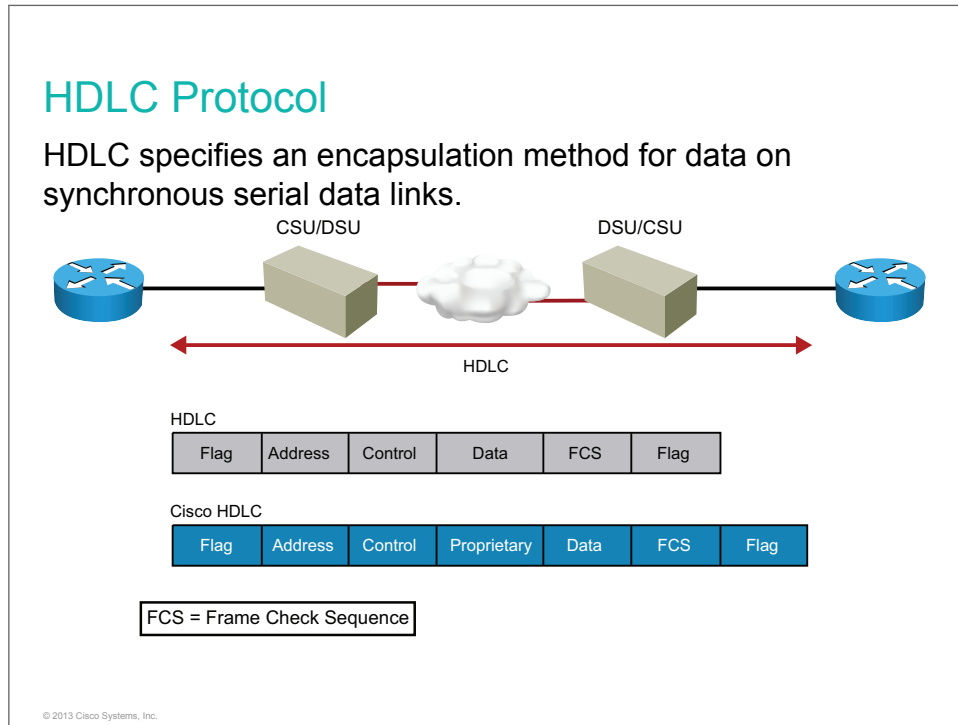
The **show controllers** command displays information about the physical interface itself. This command is useful with serial interfaces to determine the type of cable that is connected without the need to physically inspect the cable itself.

The figure shows a serial interface with an attached DTE cable.

The information that is displayed is determined when the router initially starts and represents only the type of cable that was attached when the router was started. If the cable type is changed after startup, the **show controllers** command will not display the cable type of the new cable.

HDLC Protocol

The HDLC protocol is one of two major data-link protocols that are commonly used with point-to-point WAN connections. This topic describes HDLC.



On each WAN connection, data is encapsulated into frames before crossing the WAN link. To ensure that the correct protocol is used, you need to configure the appropriate Layer 2 encapsulation type. The choice of protocol depends on the WAN technology and the communicating equipment.

The ISO developed HDLC as a synchronous data-link layer, bit-oriented protocol. HDLC uses synchronous serial transmission to provide error-free communication between two points. HDLC defines a Layer 2 framing structure that allows for flow control and error checking by using acknowledgments, control characters, and checksum. Each frame has the same format, whether it is a data frame or a control frame. HDLC may not be compatible, however, between devices from different vendors because of the way that each vendor may have chosen to implement it.

When you want to transmit frames over synchronous or asynchronous links, you must remember that those links have no mechanism to mark the beginnings or ends of the frames. HDLC uses a frame delimiter to mark the beginning and the end of each frame.

There is a Cisco implementation of HDLC that is the default encapsulation for serial lines. Cisco HDLC is very streamlined. There is no windowing or flow control, and only point-to-point connections are allowed. The Cisco HDLC implementation includes proprietary extensions in the data field, as shown in the figure. The extensions allowed multiprotocol support at a time before PPP was specified. Because of the modification, the Cisco HDLC implementation will not interoperate with other HDLC implementations. HDLC encapsulations vary, but PPP should be used when interoperability is required.

HDLC Protocol (Cont.)

```
RouterA# show interfaces Serial 0/0/0
Serial0/0/0 is up, line protocol is up
Hardware is GT96K Serial
Description: Link to HQ
Internet address is 192.168.1.1/24
MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation HDLC, loopback not set
Keepalive set (10 sec)
CRC checking enabled
Last input 00:00:02, output 00:00:05, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: weighted fair
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
Conversations 0/1/256 (active/max active/max total)
Reserved Conversations 0/0 (allocated/max allocated)
Available Bandwidth 1158 kilobits/sec
<output omitted>
```

- Verifies correct configuration of HDLC encapsulation on RouterA Serial 0/0/0 interface. By default, Cisco devices use the Cisco HDLC serial encapsulation method on synchronous serial lines.

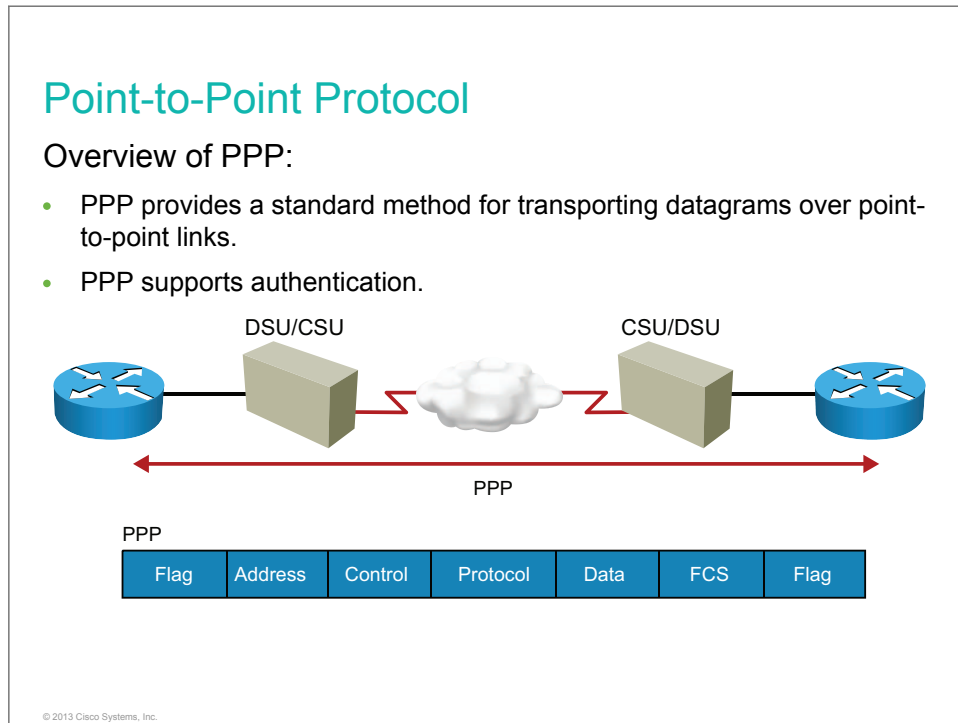
© 2013 Cisco Systems, Inc.

Cisco HDLC is the default encapsulation method that is used by Cisco devices on synchronous serial lines. Cisco HDLC is used as a point-to-point protocol on leased lines between two Cisco devices. If you are connecting to a device that is not a Cisco device, you should use synchronous PPP.

If the default encapsulation method has been changed, use the **encapsulation hdlc** command in privileged EXEC mode to re-enable HDLC.

Point-to-Point Protocol

This topic describes the characteristics of PPP and how it is enabled on a serial interface.



Cisco HDLC is a data-link layer protocol that can be used on leased lines between two Cisco devices. For communicating with a device from another vendor, synchronous PPP is a better option.

PPP originally emerged as an encapsulation protocol for transporting IP traffic over point-to-point links.

PPP provides a standard method for transporting multiprotocol datagrams (packets) over point-to-point links.

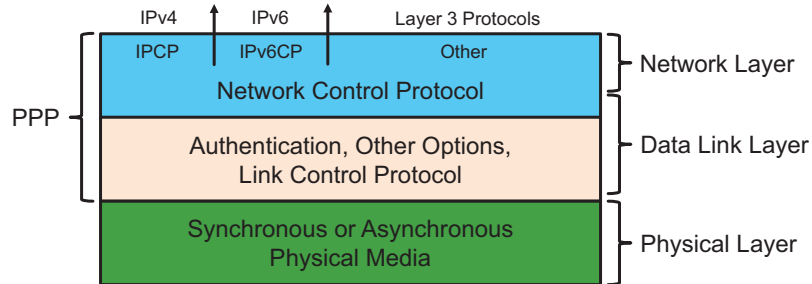
There are many advantages to using PPP, including the fact that it is not proprietary. Moreover, it includes many features that are not available in HDLC, including the following:

- The link-quality management feature monitors the quality of the link. If too many errors are detected, PPP takes down the link.
- PPP supports PAP and CHAP authentication.

Point-to-Point Protocol (Cont.)

PPP is a layered architecture:

- PPP can carry packets from several protocol suites using NCP.
- PPP controls the setup of several link options using LCP.



PPP includes these three main components:

- A method for encapsulating multiprotocol datagrams.
- Extensible LCP to establish, configure, and test the WAN data-link connection.
- A family of NCPs for establishing and configuring different network layer protocols. PPP allows the simultaneous use of multiple network layer protocols.

LCP provides versatility and portability to a wide variety of environments. LCP is used to automatically determine the encapsulation format option, to manage varying limits on sizes of packets, and to detect a loopback link and terminate the link. Other optional facilities that are provided are authentication of the identity of its peer on the link and determination of when a link is functioning correctly or failing.

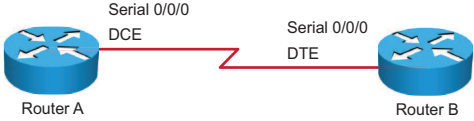
The authentication phase of a PPP session is optional. After the link has been established and the authentication protocol is chosen, the peer can be authenticated. If the authentication option is used, authentication takes place before the network layer protocol configuration phase begins.

The authentication options require the calling side of the link to enter authentication information to help ensure that the user has permission from the network administrator to make the call. Peer routers exchange authentication messages.

PPP Configuration

This topic shows a typical example of a PPP configuration.

PPP Configuration

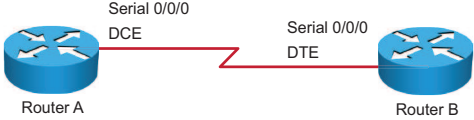


```
RouterA(config)# interface Serial 0/0/0
RouterA(config-if)# ip address 10.0.1.1 255.255.255.0
RouterA(config-if)# encapsulation ppp
RouterA(config-if)# bandwidth 512
RouterA(config-if)# clockrate 64000
RouterA(config-if)# no shutdown
```

PPP configuration on Router A

© 2013 Cisco Systems, Inc.

PPP Configuration (Cont.)



```
RouterB(config)# interface Serial 0/0/0
RouterB(config-if)# ip address 10.0.1.2 255.255.255.0
RouterB(config-if)# encapsulation ppp
RouterB(config-if)# bandwidth 512
RouterB(config-if)# no shutdown
```

PPP configuration on Router B

© 2013 Cisco Systems, Inc.

To set PPP as the encapsulation method to be used by a serial interface, use the **encapsulation ppp** interface configuration command.

The **encapsulation ppp** command has no arguments, but you must first configure the router with an IP routing protocol to use PPP encapsulation. Remember that if you do not configure PPP on a Cisco router, the default encapsulation for serial interfaces is HDLC. In this example, the bandwidth value is set to 512 kb/s.

Configuring PPP Encapsulation Commands

| Command | Description |
|---|---|
| bandwidth <i>bandwidth</i> | Sets bandwidth on the interface. This does not physically change the bandwidth of the interface. It sets a bandwidth metric for routing protocols to use (for example, OSPF) with best path calculations. |
| clockrate <i>clock_rate</i> | On the DCE side, sets the clock rate to a specified value |
| interface Serial <i>interface_number</i> | Enters serial interface configuration mode for the specified interface |
| ip address <i>ip_address subnet_mask</i> | Sets IP address on the interface |
| encapsulation ppp | Sets the interface encapsulation to PPP |
| no shutdown | Brings up the interface |

PPP Configuration (Cont.)

```
RouterA# show interfaces Serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Description: Link to RouterB
  Internet address is 10.0.1.1/24
  MTU 1500 bytes, BW 512 Kbit, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation PPP, LCP Open
  Open: IPCP, CDPCP, loopback not set
  Keepalive set (10 sec)
  CRC checking enabled
  Last input 00:00:36, output 00:00:01, output hang never
  Last clearing of "show interface" counters 00:01:09
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/0 (size/max total/threshold/drops)
    Conversations 0/1/256 (active/max active/max total)
    Reserved Conversations 0/0 (allocated/max allocated)
    Available Bandwidth 384 kilobits/sec
<output omitted>
```

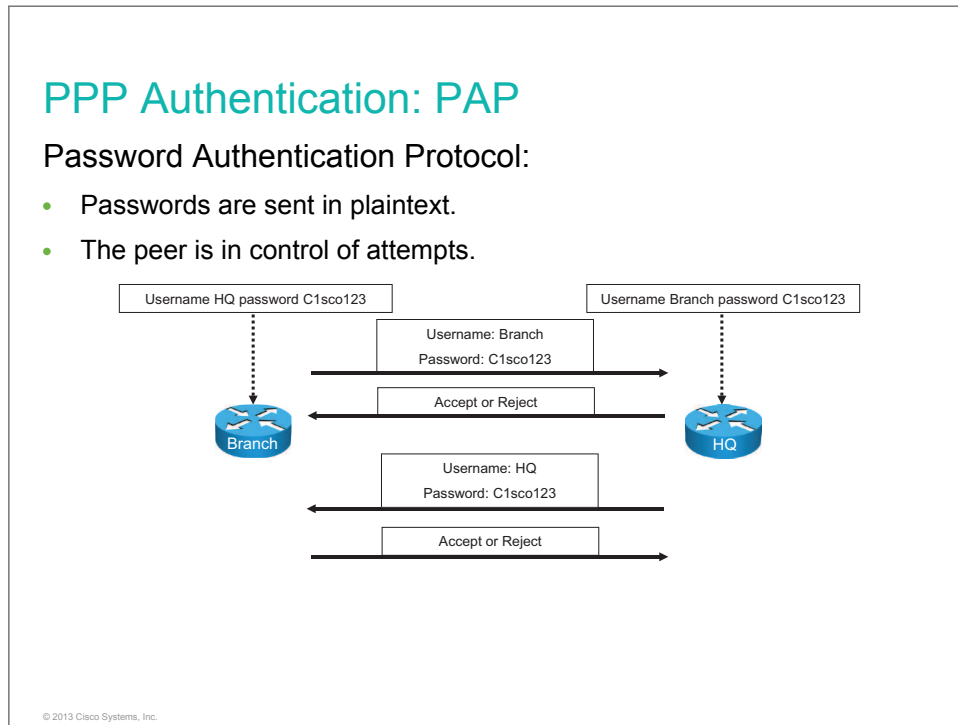
Verifies that proper encapsulation is enabled on the Serial 0/0/0 interface

© 2013 Cisco Systems, Inc.

PPP encapsulation is enabled on the serial interface. Observe that LCP is running and that two protocols have been negotiated: IP and Cisco Discovery Protocol.

PPP Authentication: PAP

This topic describes PAP, one of the two PPP authentication protocols.



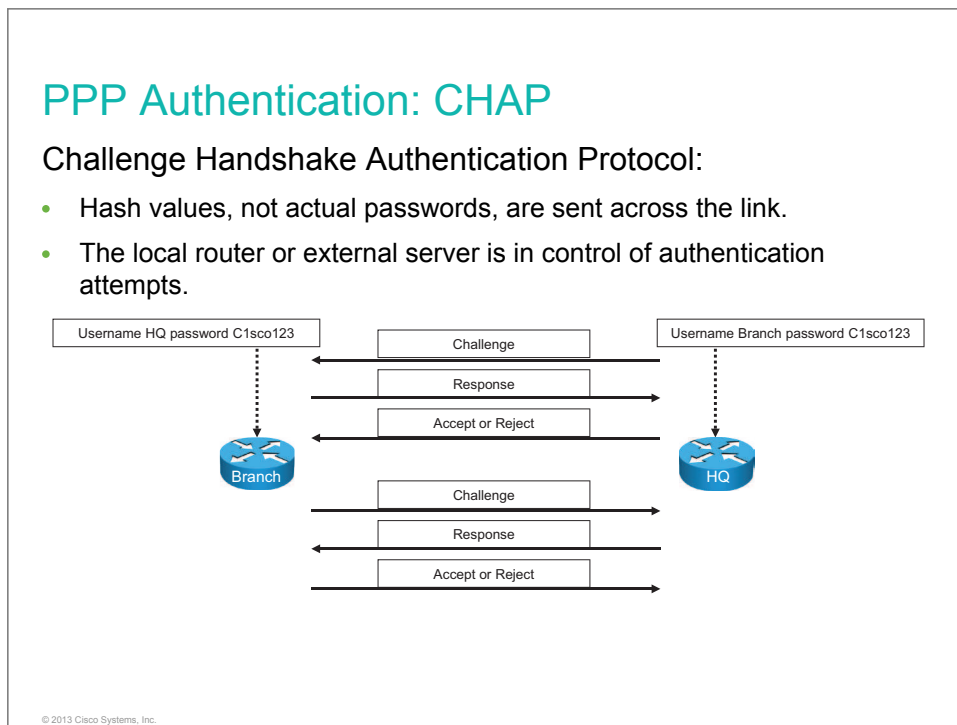
PAP is a two-way handshake that provides a simple method for a remote node to establish its identity. PAP is performed only upon initial link establishment. There is no encryption. The username and password are sent in plaintext. After the PPP link establishment phase is complete, the remote node repeatedly sends a username and password pair to the router until authentication is acknowledged or the connection is terminated.

PAP is not a strong authentication protocol, but it may be adequate in environments that use token-type passwords that change with each authentication. PPP is not secure in most environments. Also, there is no protection from playback or repeated trial-and-error attacks—the remote node is in control of the frequency and timing of the login attempts.

In the example, the Branch router first sends its PAP username and password to the Headquarters router. The Headquarters router evaluates the Branch router credentials against its local database. If the Branch router credentials match, the Headquarters router accepts the connection. If not, the Headquarters router rejects the connection. This is the two-way handshake in which the Branch router authenticates to the Headquarters router. Then the reverse process occurs with the Headquarters router authenticating to the Branch router.

PPP Authentication: CHAP

This topic describes CHAP, one of the two PPP authentication protocols.



CHAP is more secure than PAP. It involves a three-way exchange of a shared secret. Once authentication is established with PAP, it essentially stops working. This leaves the network vulnerable to attack. Unlike PAP, which only authenticates once, CHAP conducts periodic challenges to make sure that the remote node still has a valid password value.

CHAP, which uses a three-way handshake, occurs at the startup of a link and periodically thereafter to verify the identity of the remote node using a three-way handshake.

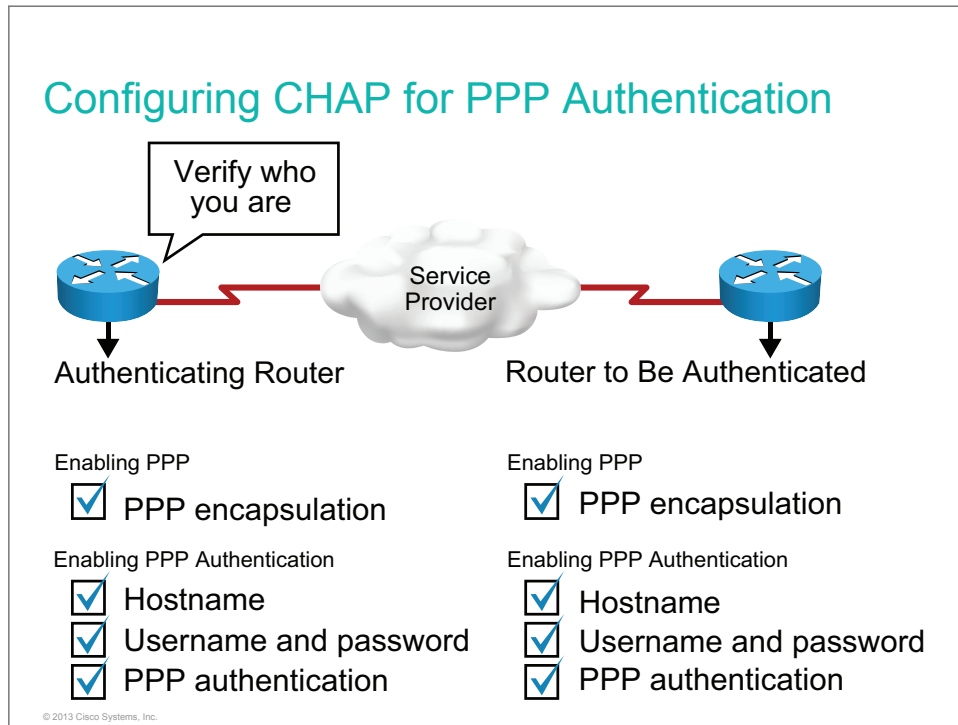
After the PPP link establishment phase is complete, the local router sends a challenge message to the remote node. The remote node responds with a value that is calculated using a one-way hash function, typically MD5, based on the password and challenge message. The local router checks the response against its own calculation of the expected hash value. If the values match, the authentication is acknowledged. Otherwise, the connection is terminated immediately.

CHAP provides protection against playback attack by using a variable challenge value that is unique and unpredictable. Because the challenge is unique and random, the resulting hash value will also be unique and random. The use of repeated challenges is intended to limit exposure to any single attack. The local router or a third-party authentication server is in control of the frequency and timing of the challenges.

In the example, the Headquarters router sends a challenge message to the Branch router. The Branch router responds to the Headquarters router by sending its CHAP username and password. The Headquarters router evaluates the Branch router credentials against its local database. If it matches, it accepts the connection. If not, it rejects the connection. This is a three-way handshake of the Headquarters router authenticating the Branch router. Then a three-way handshake of the Branch router authenticating the Headquarters router follows.

Configuring CHAP for PPP Authentication


This topic shows how authentication can be used with PPP.



To enable PPP encapsulation with CHAP authentication on an interface, complete this checklist:

- Enable PPP encapsulation as the Layer 2 protocol of an interface.
- Enable PPP authentication by completing these steps:
 1. Configure the router host name to identify it.
 2. Configure the username and password to authenticate the PPP peer.
 3. Choose CHAP as the authentication technique.

Configuring CHAP for PPP Authentication (Cont.)



```
Router(config)# hostname RouterA
RouterA(config)# username RouterB password Cisco123
RouterA(config)# interface Serial 0/0/0
RouterA(config-if)# ip address 10.0.1.1 255.255.255.0
RouterA(config-if)# encapsulation ppp
RouterA(config-if)# ppp authentication chap
RouterA(config-if)# clock rate 64000
```

Configuring CHAP authentication on Router A

© 2013 Cisco Systems, Inc.


The host name on one router must match the username that the other router has configured. The passwords must also match. The same password must be configured on both routers. Both routers must be configured for the same PPP authentication type.

The output shows CHAP configuration on Router A.

Configuring CHAP Authentication Commands

| Command | Description |
|---|--|
| hostname <i>hostname</i> | Sets a device host name |
| username <i>username</i> password <i>password</i> | Configures a new user to the device |
| interface <i>interface_name</i> | Enters interface configuration mode for the specified interface |
| ip address <i>ip_address</i> <i>subnet_mask</i> | Sets an IP address on the interface |
| encapsulation ppp | Configures a link with PPP-type encapsulation |
| ppp authentication chap | Enables CHAP authentication on the interface with PPP encapsulation |
| clock rate <i>clock_rate</i> | Configures the clock rate for a hardware connection on a serial interface. A command is used in interface configuration mode and is set on the DCE device. |

Configuring CHAP for PPP Authentication (Cont.)



```
Router(config)# hostname RouterB
RouterB(config)# username RouterA password C1sco123
RouterB(config)# interface Serial 0/0/0
RouterB(config-if)# ip address 10.0.1.2 255.255.255.0
RouterB(config-if)# encapsulation ppp
RouterB(config-if)# ppp authentication chap
```

Configuring CHAP authentication on Router B

© 2013 Cisco Systems, Inc.

The output shows CHAP configuration on Router B.

Note A globally defined user poses security concerns because these credentials can also use Telnet or SSH to connect to the device (if login local is defined under the vty lines). When implementing CHAP, vty lines need to be secured so that an outside source cannot access your device. One simple example of securing access would be using an access list that protects the vty lines.

Verifying CHAP Configuration

This topic shows how to verify the CHAP authentication configuration.

Verifying CHAP Configuration

```
RouterA# show interfaces Serial 0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is GT96K Serial
  Internet address is 10.0.1.1/24
  MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation PPP, LCP Open
  Open: IPCP, CDPCP, loopback not set
  Keepalive set (10 sec)
  CRC checking enabled
  Last input 00:00:21, output 00:00:03, output hang never
  Last clearing of "show interface" counters 00:00:47
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: weighted fair
  Output queue: 0/1000/64/0 (size/max total/threshold/drops)
    Conversations 0/1/256 (active/max active/max total)
    Reserved Conversations 0/0 (allocated/max allocated)
    Available Bandwidth 1158 kilobits/sec
```

Verifies the PPP encapsulation configuration on the Serial 0/0/0 interface and verifies that the connection is still working after configuring authentication

© 2013 Cisco Systems, Inc.

Use the **show interfaces** command to verify configuration. The figure shows that PPP encapsulation has been configured and LCP has established a connection, as indicated by “LCP Open” in the command output.

Verifying CHAP Configuration (Cont.)

```
RouterX# debug ppp authentication
Oct 23 11:08:10.642: %LINK-3-UPDOWN: Interface Serial0/0/0, changed state to
up
Oct 23 11:08:10.642: Se0/0/0 PPP: Using default call direction
Oct 23 11:08:10.642: Se0/0/0 PPP: Treating connection as a dedicated line
Oct 23 11:08:10.642: Se0/0/0 PPP: Session handle[CC000003] Session id[5]
Oct 23 11:08:10.642: Se0/0/0 PPP: Authorization required
Oct 23 11:08:10.674: Se0/0/0 CHAP: O CHALLENGE id 4 len 28 from "RouterX"
Oct 23 11:08:10.718: Se0/0/0 CHAP: I CHALLENGE id 1 len 28 from "RouterY"
Oct 23 11:08:10.718: Se0/0/0 CHAP: Using hostname from unknown source
Oct 23 11:08:10.718: Se0/0/0 CHAP: Using password from AAA
Oct 23 11:08:10.718: Se0/0/0 CHAP: O RESPONSE id 1 len 28 from "RouterX"
Oct 23 11:08:10.722: Se0/0/0 CHAP: I RESPONSE id 4 len 28 from "RouterY"
Oct 23 11:08:10.722: Se0/0/0 PPP: Sent CHAP LOGIN Request
Oct 23 11:08:10.726: Se0/0/0 PPP: Received LOGIN Response PASS
Oct 23 11:08:10.726: Se0/0/0 PPP: Sent LCP AUTHOR Request
Oct 23 11:08:10.726: Se0/0/0 PPP: Sent IPCP AUTHOR Request
Oct 23 11:08:10.726: Se0/0/0 LCP: Received AAA AUTHOR Response PASS
Oct 23 11:08:10.726: Se0/0/0 IPCP: Received AAA AUTHOR Response PASS
```

(Continued in next figure)

© 2013 Cisco Systems, Inc.

Verifying CHAP Configuration (Cont.)

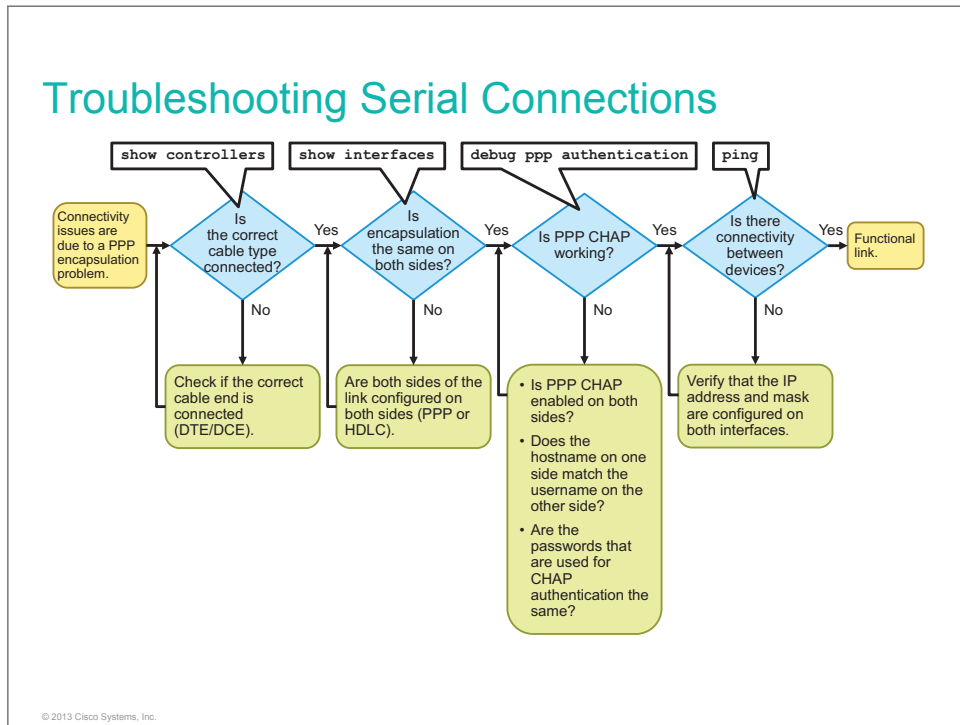
```
Oct 23 11:08:10.726: Se0/0/0 CHAP: O SUCCESS id 4 len 4
Oct 23 11:08:10.742: Se0/0/0 CHAP: I SUCCESS id 1 len 4
Oct 23 11:08:10.746: Se0/0/0 PPP: Sent CDPCP AUTHOR Request
Oct 23 11:08:10.746: Se0/0/0 CDPCP: Received AAA AUTHOR Response PASS
Oct 23 11:08:10.746: Se0/0/0 PPP: Sent IPCP AUTHOR Request-if)#
Oct 23 11:08:11.742: %LINEPROTO-5-UPDOWN: Line protocol on Interface
Serial0/0/0, changed state to up
```

The **debug ppp authentication** command shows the successful CHAP output and verifies PPP authentication.

The figure illustrates the router output that occurs during CHAP authentication. Because two-way authentication is configured (that is, each router authenticates the other), messages appear that reflect both the authenticating process and the process of being authenticated.

Troubleshooting Serial Connections

This topic describes how to troubleshoot PPP encapsulation.



To troubleshoot link issues that are due to misconfigured encapsulation and authentication on the serial interface, follow these high-level steps:

- 1 Check if the correct cable is connected to the device (DTE and DCE).
- 2 Verify that both sides of the serial connection are configured with the same encapsulation—PPP or HDLC. HDLC is the Cisco default.
- 3 If PPP encapsulation is correctly configured on both sides of the link, verify that the CHAP authentication is successful. The CHAP authentication may not be successful when the host name on one side does not match the username on the other side, or the password is not the same on both sides. The password is often misconfigured due to an additional space character at the end of the password string. Also, make sure that the same authentication method (PAP or CHAP) is selected on both sides of the link.
- 4 If CHAP authentication is successful, the link should be operational. If a ping still does not work, verify that the IP address and mask are correctly set on both sides of the link.

Summary

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

Summary

- In addition to the ISO-developed HDLC, there is a Cisco implementation of HDLC, which is the default encapsulation for serial lines on Cisco routers.
- PPP is a common Layer 2 protocol for the WAN. There are two components of PPP: LCP, which negotiates the connection, and NCP, which encapsulates traffic.
- To set PPP as the encapsulation method to be used by a serial interface, use the **encapsulation ppp** interface configuration command.
- You can configure PPP to use PAP or CHAP. PAP sends everything in plaintext. CHAP uses an MD5 hash.
- For CHAP authentication, the remote device must have a corresponding username entry for the local router, with a matching password.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Establishing a WAN Connection Using Frame Relay

Frame Relay is a standardized WAN technology that is a well-proven packet-switching, connection-oriented technology that is used to interconnect remote sites.

This lesson describes the basic functionality of Frame Relay, including topologies, reachability issues, and LMI signaling. The lesson also describes how to configure basic Frame Relay and Frame Relay over point-to-point and multipoint subinterfaces. The lesson concludes with verification of Frame Relay operations.

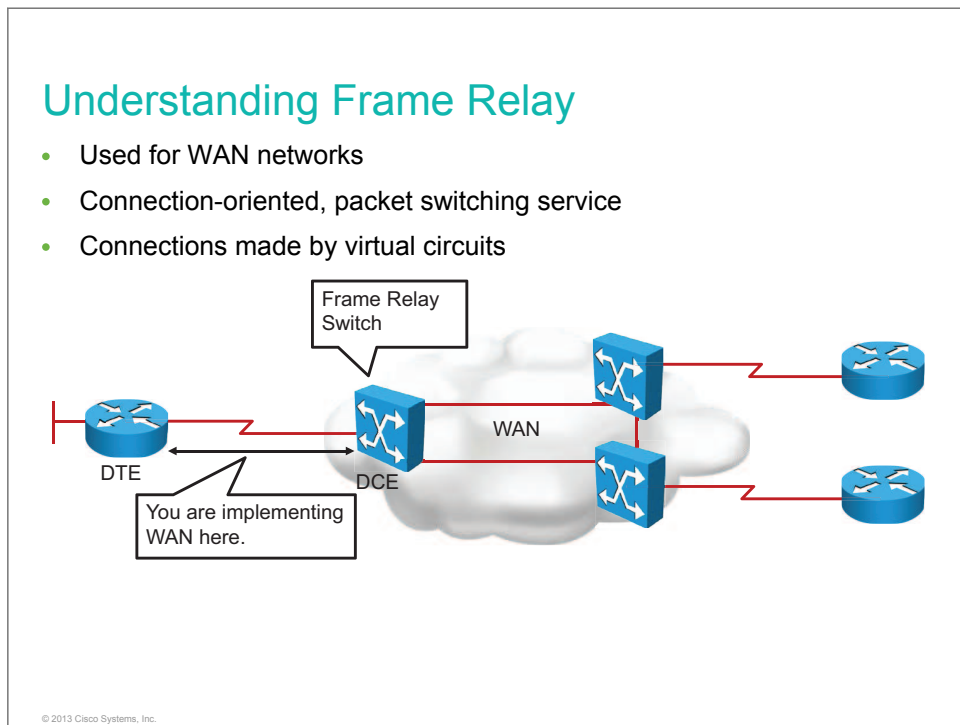
Objectives

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

- Describe Frame Relay technology and its basic configuration
- Explain Frame Relay topologies
- Explain Frame Relay reachability issues
- Explain Frame Relay LMI signaling
- Explain Frame Relay address mappings
- Configure basic Frame Relay
- Explain the difference between point-to-point and multipoint Frame Relay
- Configure point-to-point Frame Relay
- Configure multipoint Frame Relay
- Verify Frame Relay operation

Understanding Frame Relay

This topic describes the basic functionality of Frame Relay.



Frame Relay was originally designed for use across ISDN interfaces. Today, it is used over various other network interfaces as well. Frame Relay is a connection-oriented, data-link technology that is streamlined to provide high performance and efficiency. For error protection, it relies on upper-layer protocols and dependable fiber and digital networks. Frame Relay is an example of a packet-switched technology. Packet-switched networks enable end stations to dynamically share the network media and available bandwidth.

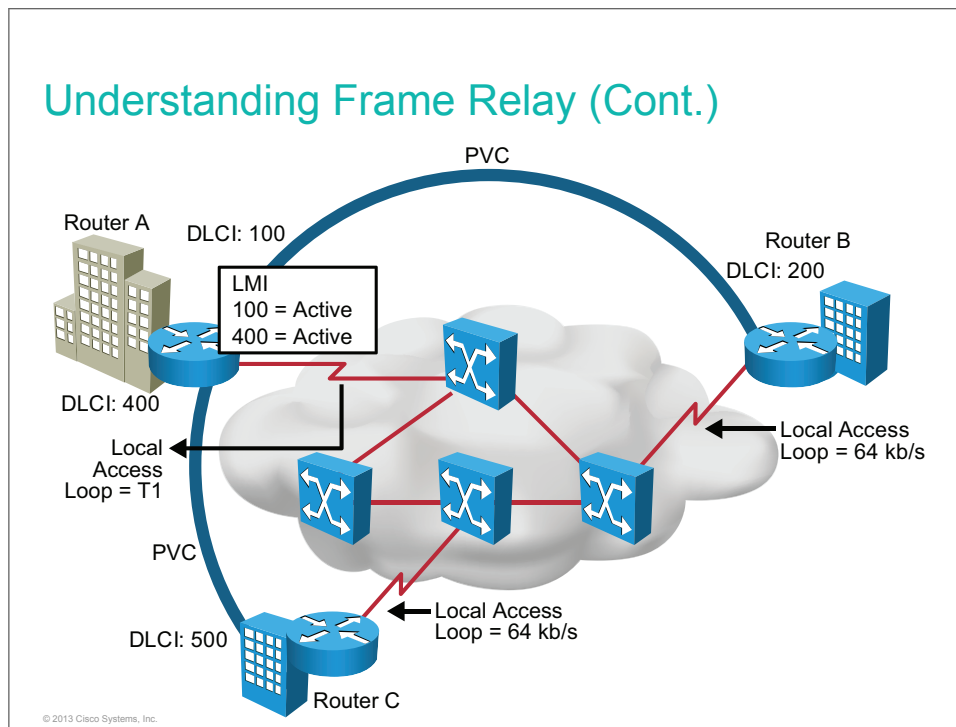
Frame Relay defines the interconnection process between the router and the local access switching equipment of the service provider. It does not define how the data is transmitted within the Frame Relay service provider cloud.

Devices that are attached to a Frame Relay WAN fall into two categories:

- **DTE:** DTE is generally considered to be the terminating equipment for a specific network. DTE devices are typically located on the customer premises and may be owned by the customer. Examples of DTE devices are FRADs, routers, and switches.
- **DCE:** DCE is service provider-owned internetworking devices. The purpose of DCE devices is to provide clocking and switching services in a network and to transmit data through the WAN. In most cases, the switches in a WAN are Frame Relay switches.

Frame Relay provides a means for statistically multiplexing many logical data conversations, referred to as VCs, over a single physical transmission link by assigning connection identifiers to each pair of DTE devices. The service provider switching equipment constructs a switching table that maps the connection identifier to outbound ports. When a frame is received, the switching device analyzes the connection identifier and delivers the frame to the associated outbound port. The complete path to the destination is established before transmission of the first frame.

Understanding Frame Relay (Cont.)



These terms are used frequently in Frame Relay discussions. They may be slightly different from the terms that your Frame Relay service provider uses.

- **Local access rate:** Local access rate is the clock speed (port speed) of the connection (local loop) to the Frame Relay cloud. The local access rate is the rate at which data travels into or out of the network, regardless of other settings.
- **VC:** A VC is a logical circuit, uniquely identified by a DLCI that is created to ensure bidirectional communication from one DTE device to another. A number of VCs can be multiplexed into a single physical circuit for transmission across the network. This capability can often reduce the complexity of the equipment and network that is required to connect multiple DTE devices. A VC can pass through any number of intermediate DCE devices (Frame Relay switches). A VC can be either a PVC or an SVC.
- **PVC:** A PVC provides permanently established connections that are used for frequent and consistent data transfers between DTE devices across the Frame Relay network.
- **DLCI:** Frame Relay VCs are identified by DLCIs. The Frame Relay service providers (for example, telcos) typically assign DLCI values. A DLCI contains a 10-bit number in the address field of the Frame Relay frame header that identifies the VC. DLCIs have local significance because the identifier references the point between the local router and the local Frame Relay switch to which the DLCI is connected. Therefore, devices at opposite ends of a connection can use different DLCI values to refer to the same VC.
- **CIR:** CIR specifies the maximum average data rate that the network undertakes to deliver under normal conditions. When subscribing to a Frame Relay service, you specify the local access rate (for example, 56 kb/s or T1). Typically, you are also asked to specify a CIR for each DLCI. If you send information faster than the CIR on a given DLCI, the network marks some frames with a DE bit. The network does its best to deliver all packets but discards DE packets first if there is congestion. Many inexpensive Frame Relay services are based on a CIR of zero. A CIR of zero means that every frame is a DE frame, and the network throws away any frame when it needs to.

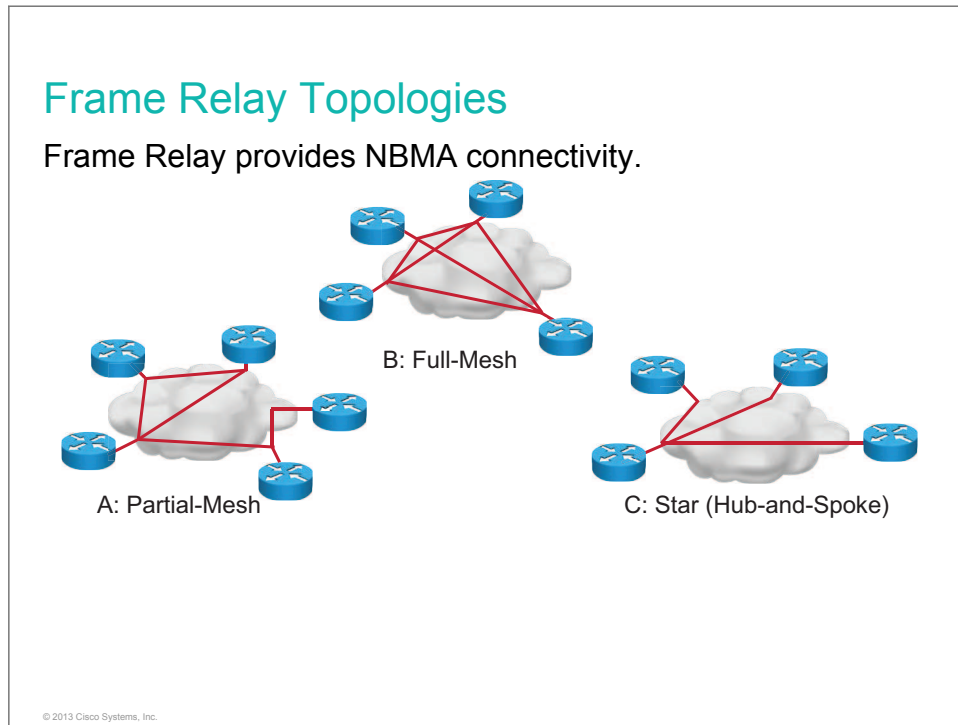
- **Inverse ARP:** Inverse ARP is a method of dynamically associating the network layer address of the remote router with a local DLCI. Inverse ARP allows a router to automatically discover the network address of the remote DTE device that is associated with a VC.
- **LMI:** LMI is a signaling standard between the router (DTE device) and the local Frame Relay switch (DCE device) that is responsible for managing the connection and maintaining status between the router and the Frame Relay switch. Basically, the LMI is a mechanism that provides status information about Frame Relay connections between the router (DTE) and the Frame Relay switch (DCE). Every 10 seconds or so, the end device polls the network, either requesting a dumb sequenced response or channel status information. If the network does not respond with the requested information, the user device may consider the connection to be down.

As shown in the figure, router A has two VCs that are configured on one physical interface. A DLCI of 100 identifies the VC that connects to router B. A DLCI of 400 identifies the VC that connects to router C. At the other end, a different DLCI number is used to identify the VC. The LMI at router A allows the router to learn DLCIs for PVCs that are available from the local Frame Relay switch.

Do Not Duplicate.
Post beta, not for release.

Frame Relay Topologies

This topic describes Frame Relay topologies.



By default, a Frame Relay network provides NBMA connectivity between remote sites. An NBMA environment is treated like other broadcast media environments, such as Ethernet, where all of the routers are on the same subnet.

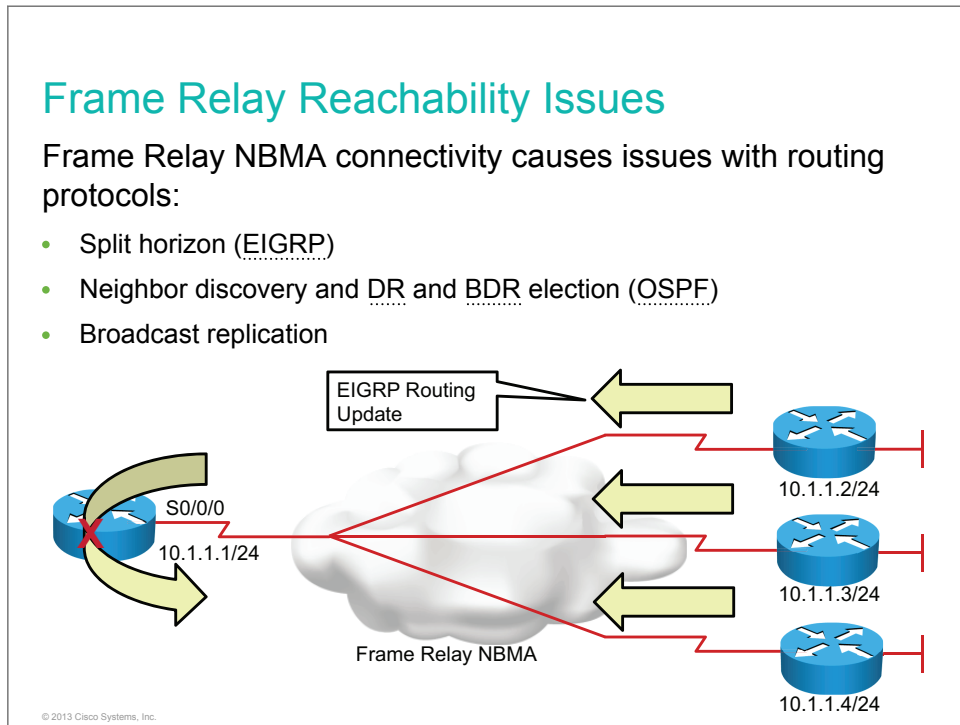
However, to reduce cost, NBMA clouds are usually built in a hub-and-spoke topology. With a hub-and-spoke topology, the physical topology does not provide the multiaccess capabilities that Ethernet does, so each router may not have separate PVCs to reach the other remote routers on the same subnet. Split horizon is one of the main issues that you encounter when Frame Relay is running multiple PVCs over a single interface.

Frame Relay allows you to interconnect your remote sites in various topologies:

- **Hub-and-spoke topology:** Remote sites are connected to a central site that generally provides a service or an application. The hub-and-spoke topology is the most popular Frame Relay network topology. It is the least expensive topology because it requires the fewest PVCs. In the figure, the central router provides a multipoint connection because it uses a single interface to interconnect multiple PVCs.
- **Full-mesh topology:** In this topology, all routers have VCs to all other destinations. A full-mesh topology, although costly, provides direct connections from each site to all other sites and allows for redundancy. When one link goes down, a router can reroute traffic through another site. As the number of nodes in this topology increases, a full-mesh topology can become very expensive. Use the $n(n - 1) / 2$ formula to calculate the total number of links that are required to implement a full-mesh topology, where n is the number of nodes. For example, to fully mesh a network of 10 nodes, 45 links are required: $10(10 - 1) / 2$.
- **Partial-mesh topology:** In a partial-mesh topology, not all sites have direct access to all other sites. Depending on the traffic patterns in your network, you may want to have additional PVCs connect to remote sites that have large data traffic requirements.

Frame Relay Reachability Issues

This topic describes Frame Relay reachability issues and how to solve them.



In any Frame Relay topology, when a single multipoint interface must be used to interconnect multiple sites, reachability issues may result because of the NBMA nature of Frame Relay. The Frame Relay NBMA topology can cause the following problems:

- **Split horizon:** In case of EIGRP and RIP, the split horizon rule reduces routing loops by preventing a routing update that is received on an interface from being forwarded out of the same interface. In a scenario using a hub-and-spoke Frame Relay topology, a remote router (a spoke router) sends an update to the Headquarters router (the hub router) that is connecting multiple PVCs over a single physical interface. The Headquarters router receives the broadcast on its physical interface but cannot forward that routing update through the same interface to other remote (spoke) routers. Split horizon is not a problem if there is a single PVC on a physical interface because this type of connection would be point-to-point.
- **Neighbor discovery and DR and BDR election:** OSPF over NBMA networks works in nonbroadcast network mode, by default, and neighbors are not automatically discovered. You can statically configure neighbors. However, you have to make sure that the hub router becomes a DR. An NBMA network behaves like Ethernet, and on Ethernet, a DR is needed to exchange routing information between all routers on a segment. Therefore, only the hub router can act as a DR because it is the only router that has PVCs with all other routers.

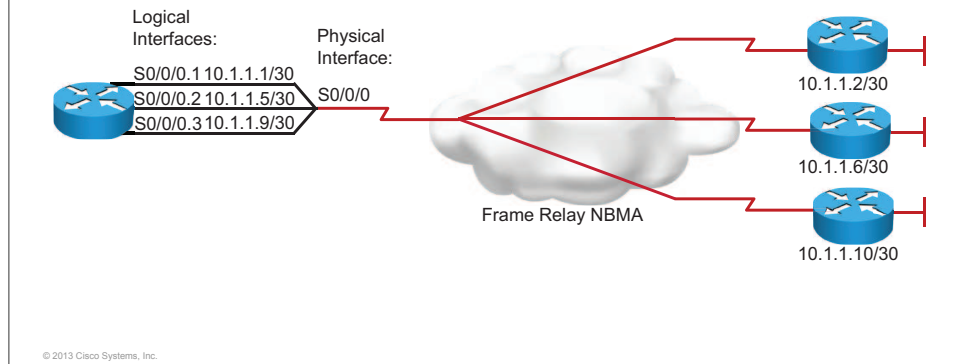
Note When configuring OSPF over Frame Relay, you can also use broadcast and point-to-multipoint OSPF network mode to overcome the split horizon problem.

- **Broadcast replication:** With routers that support multipoint connections over a single interface that terminate many PVCs, the router must replicate broadcast packets, such as routing update broadcasts, on each PVC to the remote routers. These replicated broadcast packets consume bandwidth and cause significant latency variations in user traffic.

Frame Relay Reachability Issues (Cont.)

Subinterfaces are one solution to routing problems in NBMA networks:

- A physical interface simulates multiple, logical, point-to-point interfaces.
- Each subinterface is on a separate IP network.
- Each subinterface is associated with a Frame Relay PVC.

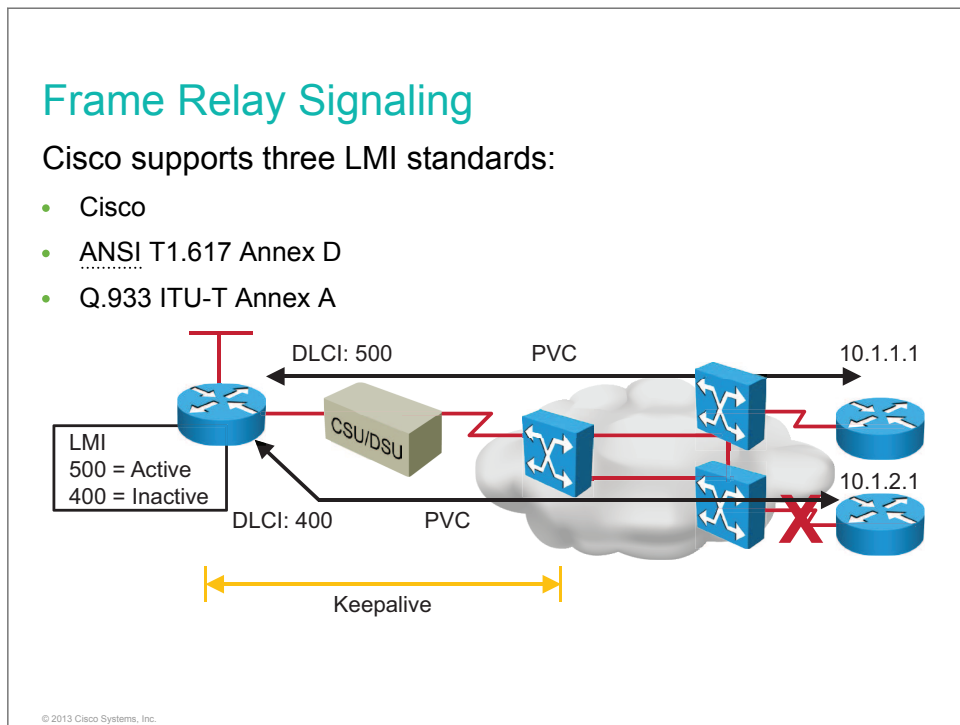


There are several ways to solve the routing update reachability issues:

- One method for solving the reachability issues that are produced by split horizon may be to turn off split horizon. However, disabling split horizon increases the chances of routing loops in your network.
- Another method is to use a fully meshed topology. However, this topology increases costs.
- The last method is to use subinterfaces. To enable forwarding of broadcast routing updates in a hub-and-spoke Frame Relay topology, you can configure the hub router with logically assigned interfaces that are called *subinterfaces*, which are logical subdivisions of a physical interface. In split-horizon routing environments, routing updates that are received on one subinterface can be sent out another subinterface. In subinterface configuration, each VC can be configured as a point-to-point connection that allows each subinterface to act like a leased line. When you use a Frame Relay point-to-point subinterface, each subinterface is on its own subnet.

Frame Relay Signaling

This topic describes Frame Relay LMI signaling.



LMI is a standard for signaling between a router and a Frame Relay switch. LMI is responsible for managing the connection and maintaining the status between the devices.

Although LMI is configurable, the Cisco router tries to autosense which LMI type that the Frame Relay switch is using. The router sends one or more complete LMI status requests to the Frame Relay switch. The Frame Relay switch responds with one or more LMI types, and the router configures itself with the last LMI type that was received. Cisco routers support three LMI types:

- **Cisco:** LMI type that was developed jointly by Cisco, StrataCom, Northern Telecom (Nortel), and Digital Equipment Corporation
- **ANSI:** ANSI T1.617 Annex D
- **Q.933A:** ITU-T Q.933 Annex A

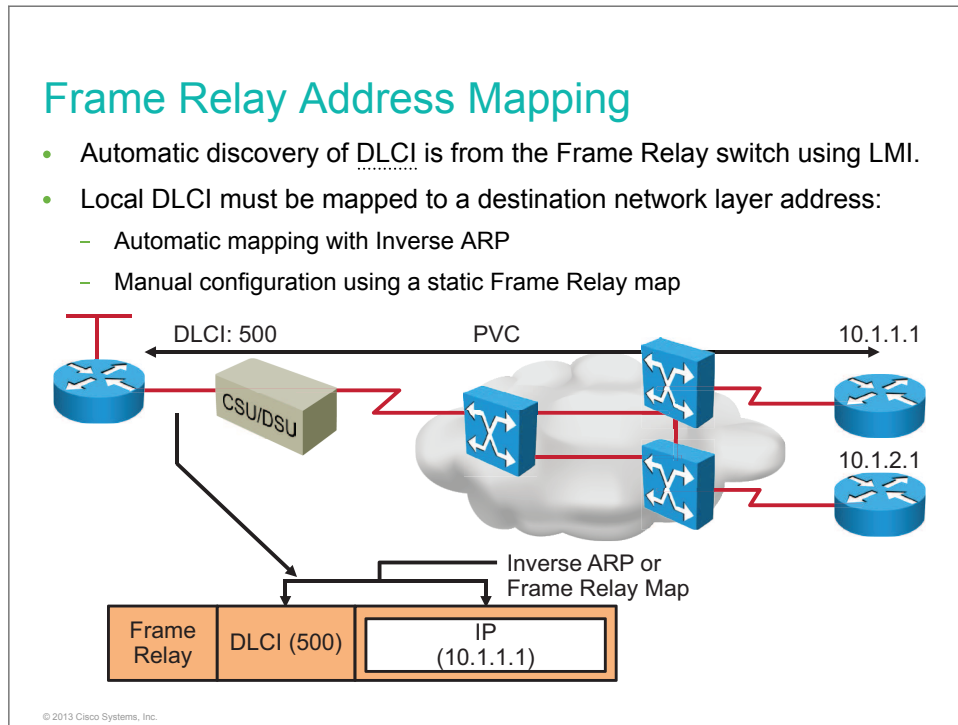
You can also manually configure the appropriate LMI type from the three supported types to ensure correct Frame Relay operation.

When the router receives LMI information, it updates its VC status to one of three states:

- **Active:** Indicates that the VC connection is active and that routers can exchange data over the Frame Relay network
- **Inactive:** Indicates that the local connection to the Frame Relay switch is working, but the remote router connection to the remote Frame Relay switch is not working
- **Deleted:** Indicates that either no LMI is being received from the Frame Relay switch, or there is no service between the router and the local Frame Relay switch

Frame Relay Address Mappings

This topic describes Frame Relay address mappings and Inverse ARP.

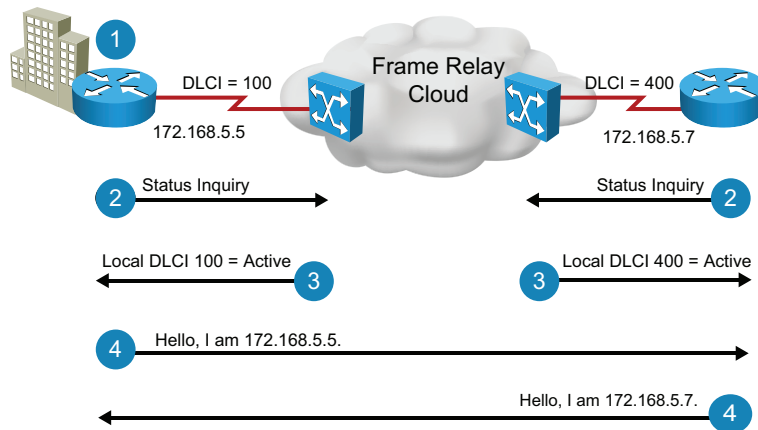


A Frame Relay connection requires that, on a VC, the local DLCI must be mapped to a destination network layer address such as an IP address. Routers can automatically discover their local DLCI from the local Frame Relay switch using the LMI protocol.

On Cisco routers, the local DLCI can be dynamically mapped to the remote router network layer addresses with Inverse ARP. Inverse ARP associates a given DLCI to the next-hop protocol address for a specific connection. Inverse ARP is described in RFC 1293. Instead of using Inverse ARP to automatically map the local DLCIs to the remote router network layer addresses, you can manually configure a static Frame Relay map in the map table.

As shown in the figure, using Inverse ARP, the router on the left can automatically discover the remote router IP address and then map it to the local DLCI. In this case, the local DLCI of 500 is mapped to the 10.1.1.1 IP address. Therefore, when the router must send data to 10.1.1.1, it uses DLCI 500.

Frame Relay Address Mapping (Cont.)

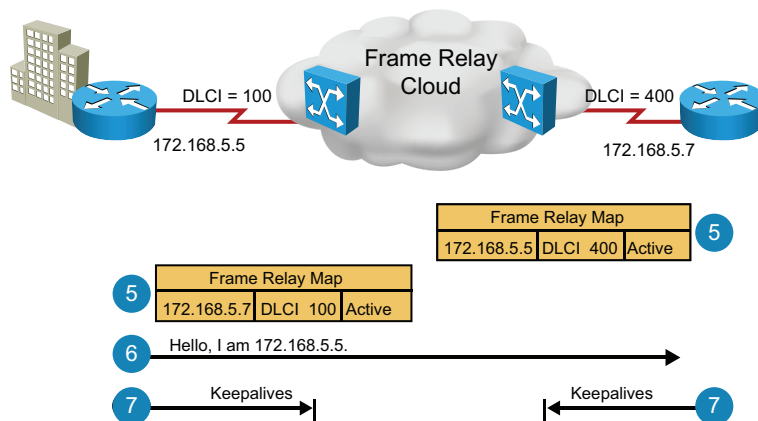


© 2013 Cisco Systems, Inc.

The following is a summary of how Inverse ARP and LMI signaling work with a Frame Relay connection.

- 1 Each router connects to the Frame Relay switch through a CSU/DSU.
- 2 When Frame Relay is configured on an interface, the router sends an LMI status inquiry message to the Frame Relay switch. The message notifies the switch of the router status and asks the switch for the connection status of the router VCs.
- 3 When the Frame Relay switch receives the request, it responds with an LMI status message that includes the local DLCIs of the PVCs to the remote routers to which the local router can send data.
- 4 For each active DLCI, each router sends an Inverse ARP packet.

Frame Relay Address Mapping (Cont.)



© 2013 Cisco Systems, Inc.

- 5 When a router receives an Inverse ARP message, it creates a map entry in its Frame Relay map table that includes the local DLCI and the remote router network layer address. Note that the router DLCI is the local DLCI, not the DLCI that the remote router is using. Any of the three connection states can appear in the Frame Relay map table.

Note If Inverse ARP is not working or the remote router does not support Inverse ARP, you must manually configure static Frame Relay maps, which map the local DLCIs to the remote network layer addresses.

- 6 Every 10 seconds, the router exchanges LMI information with the switch (keepalives).
- 7 The router changes the status of each DLCI to active, inactive, or deleted, based on the LMI response from the Frame Relay switch.

Frame Relay Address Mapping (Cont.)

Configure a static Frame Relay map in these situations:

- A Frame Relay peer does not support Inverse ARP.
- You want to control broadcast traffic across a PVC.
- You want to have different Frame Relay encapsulations across PVCs.

© 2013 Cisco Systems, Inc.

When the remote router does not support Inverse ARP, the Frame Relay peers have different Frame Relay encapsulation types. Or, when you want to control broadcast and multicast traffic over the PVC, you must statically map the local DLCI to the remote router network layer address. These static Frame Relay map entries are referred to as static maps.

Configuring Frame Relay

This topic describes how to configure Frame Relay.

Configuring Frame Relay

This router uses Inverse ARP.

DLCI=100
10.1.1.1/24

WAN

This router uses a static map.

DLCI=200
10.1.1.2/24

```
HQ(router)# interface Serial0/0/0
HQ(router-if)# ip address 10.1.1.1 255.255.255.0
HQ(router-if)# encapsulation frame-relay
HQ(router-if)# bandwidth 64
```

Configuration on the Headquarters router

© 2013 Cisco Systems, Inc.

Configuring Frame Relay (Cont.)

This router uses Inverse ARP.

DLCI=100
10.1.1.1/24

WAN

This router uses a static map.

DLCI=200
10.1.1.2/24

```
Branch(router)# interface Serial0/0/0
Branch(router-if)# ip address 10.1.1.2 255.255.255.0
Branch(router-if)# encapsulation frame-relay
Branch(router-if)# bandwidth 64
```

Configuration on the Branch router

© 2013 Cisco Systems, Inc.

A basic Frame Relay configuration assumes that you want to configure Frame Relay on one or more physical interfaces and that the routers support LMI and Inverse ARP. Inverse ARP is enabled by default on Cisco routers.

| Command | Description |
|---|---|
| interface <i>interface</i> | Enters interface configuration mode for the specified interface. |
| ip address <i>ip_address subnet_mask</i> | Sets the IP address on the interface. |
| encapsulation frame-relay [cisco ietf] | Configures Frame Relay encapsulation on an interface. Optionally, you can specify Frame Relay encapsulation type. You can choose between "cisco" and "ietf," where "cisco" is the default and "ietf" is an IETF standard. |
| bandwidth <i>bandwidth</i> | Sets the bandwidth metric on an interface. This is used by some routing protocols (for example, OSPF) for calculating the best path to the destination. |

In the example, both routers are configured with Frame Relay on the Serial0/0/0 interface. The Headquarters and Branch routers use Inverse ARP to learn the IP address of the host on the other side.

Point-to-Point vs. Multipoint

This topic explains the differences between point-to-point and multipoint Frame Relay.

Point-to-Point vs. Multipoint

Two types of subinterfaces:

- **Point-to-point:**
 - Subinterfaces act like leased lines.
 - Each point-to-point subinterface requires its own subnet.
 - Point-to-point is applicable to hub-and-spoke topologies.
- **Multipoint:**
 - Subinterfaces act like NBMA networks, so they do not resolve split-horizon issues.
 - Multipoint can save address space because it uses a single subnet.
 - Multipoint is applicable to partial-mesh and full-mesh topologies.

© 2013 Cisco Systems, Inc.

You can configure subinterfaces in either of two modes:

- **Point-to-point:** A single, point-to-point subinterface is used to establish one PVC connection to another physical interface or subinterface on a remote router. In this case, each pair of the point-to-point routers is on its own subnet, and each point-to-point subinterface has a single DLCI. In a point-to-point environment, because each subinterface acts like a point-to-point interface, update traffic is not subject to the split-horizon rule.
- **Multipoint:** A single, multipoint subinterface is used to establish multiple PVC connections to multiple physical interfaces or subinterfaces on remote routers. In this case, all of the participating interfaces are in the same subnet. In this environment, because the subinterface acts like a regular NBMA Frame Relay interface, update traffic is subject to the split-horizon rule. This rule is a method of preventing routing loops in distance-vector routing protocols (for example, EIGRP). The mechanism prohibits a router from advertising a route back onto the interface from which it was learned. When a device participating in route advertisements receives an update from an interface, the device will forward updates through all interfaces except the interface upon which it received the update.

Configuring Point-to-Point Frame Relay

This topic shows how to configure point-to-point Frame Relay.

Configuring Point-to-Point Frame Relay

```

HQ(router)# interface Serial0/0/0
HQ(router-if)# no ip address
HQ(router-if)# encapsulation frame-relay
HQ(router-if)# interface Serial0/0/0.110 point-to-point
HQ(router-subif)# ip address 10.1.1.1 255.255.255.252
HQ(router-subif)# bandwidth 64
HQ(router-subif)# frame-relay interface-dlci 110
HQ(router-subif)# interface Serial0/0/0.120 point-to-point
HQ(router-subif)# ip address 10.1.1.5 255.255.255.252
HQ(router-subif)# bandwidth 64
HQ(router-subif)# frame-relay interface-dlci 120
    
```

Configuration of point-to-point subinterfaces on the Headquarters router


© 2013 Cisco Systems, Inc.

When configuring Frame Relay point-to-point subinterfaces, first enable Frame Relay encapsulation on the physical interface. Then create a point-to-point subinterface and assign an IP address, bandwidth, and DLCI to it. On point-to-point subinterfaces, you do not need to use the Frame Relay map command to perform static address mapping because it is always assumed that the end point of the point-to-point connection automatically resides on the same subnet as the start point. It is also not required to enable or disable Inverse ARP because there is only a single remote destination on a point-to-point PVC, and discovery is not necessary.

| Command | Description |
|---|---|
| encapsulation frame-relay | Configures Frame Relay encapsulation on an interface. |
| bandwidth <i>bandwidth</i> | Sets the bandwidth on an interface. |
| interface <i>interface.subinterface</i> point-to-point | Creates a point-to-point subinterface and enters subinterface configuration mode. |
| frame-relay interface-dlci <i>dlci</i> | Assigns a DLCI to a subinterface. This command is not needed if you use static mappings between IP addresses and DLCIs. |

In the example, the Headquarters router is configured with two point-to-point subinterfaces. One subinterface uses DLCI 110 to reach the Branch 1 router, and the other subinterface uses DLCI 120 to reach the Branch 2 router.

Configuring Point-to-Point Frame Relay (Cont.)




```
BR1(router)# interface Serial0/0/0
BR1(router-if)# no ip address
BR1(router-if)# encapsulation frame-relay
BR1(router-if)# interface Serial0/0/0.110 point-to-point
BR1(router-subif)# ip address 10.1.1.2 255.255.255.0
BR1(router-subif)# bandwidth 64
BR1(router-subif)# frame-relay interface-dlci 110
```

Configuration of Branch 1 for point-to-point Frame Relay

© 2013 Cisco Systems, Inc.

The Branch 1 router has one subinterface to reach the Headquarters router. It uses DLCI 110 to reach the Headquarters router.

Configuring Point-to-Point Frame Relay (Cont.)



```
BR2(router)# interface Serial0/0/0
BR2(router-if)# no ip address
BR2(router-if)# encapsulation frame-relay
BR2(router-if)# interface Serial0/0/0.210 point-to-point
BR2(router-subif)# ip address 10.1.1.6 255.255.255.0
BR2(router-subif)# bandwidth 64
BR2(router-subif)# frame-relay interface-dlci 210
```

Configuration of Branch 2 for point-to-point Frame Relay

© 2013 Cisco Systems, Inc.

The Branch 2 router has one subinterface to reach the Headquarters router. It uses DLCI 210 to reach the Headquarters router.

Configuring Multipoint Frame Relay

This topic shows how to configure multipoint Frame Relay.

Configuring Multipoint Frame Relay

```

HQ(router)# interface Serial0/0/0
HQ(router-if)# no ip address
HQ(router-if)# encapsulation frame-relay
HQ(router-if)# interface Serial0/0/0.1 multipoint
HQ(router-subif)# ip address 10.1.1.1 255.255.255.0
HQ(router-subif)# bandwidth 64
HQ(router-subif)# frame-relay map ip 10.1.1.2 110 broadcast
HQ(router-subif)# frame-relay map ip 10.1.1.3 120 broadcast
    
```

Configuration of multipoint subinterfaces on the Headquarters router


© 2013 Cisco Systems, Inc.

When configuring Frame Relay multipoint subinterfaces, first enable Frame Relay encapsulation on the physical interface. Then create a multipoint subinterface and assign an IP address and bandwidth. Use static mappings between IP addresses and DLCIs using the **frame-relay map** command.

| Command | Description |
|--|--|
| encapsulation frame-relay | Configures Frame Relay encapsulation on an interface. |
| interface interface.subinterface point-to-point | Creates a point-to-point subinterface and enters subinterface configuration mode. |
| bandwidth bandwidth | Sets the bandwidth on an interface. This is used by some routing protocols (for example, OSPF) to use when calculating the best path through the network. |
| frame-relay interface-dlci dlci | Assigns a DLCI to a subinterface. This command is not needed if you use static mappings between IP addresses and DLCIs. |
| frame-relay map ip ip_address dlci [broadcast] | Defines mapping between the destination protocol address and the DLCI. An optional broadcast command forwards broadcasts when multicast is not enabled. |

In the example, the Headquarters router is configured with two point-to-point subinterfaces. One subinterface uses DLCI 110 to reach the Branch 1 router, and the other subinterface uses DLCI 120 to reach the Branch 2 router.

Configuring Multipoint Frame Relay (Cont.)




```
BR1 (router)# interface Serial0/0/0
BR1 (router-if)# encapsulation frame-relay
BR1 (router-if)# ip address 10.1.1.2 255.255.255.252
BR1 (router-if)# bandwidth 64
BR1 (router-if)# frame-relay map ip 10.1.1.1 110
```

Configuration of multipoint subinterfaces on the Branch 1 router

© 2013 Cisco Systems, Inc.

The Branch 1 router has one interface to reach the Headquarters router. It uses DLCI 110 to reach the Headquarters router.

Configuring Multipoint Frame Relay (Cont.)



```
BR2 (router)# interface Serial0/0/0
BR2 (router-if)# encapsulation frame-relay
BR2 (router-if)# ip address 10.1.1.6 255.255.255.252
BR2 (router-if)# bandwidth 64
BR2 (router-if)# frame-relay map ip 10.1.1.1 210
```

Configuration of multipoint subinterfaces on the Branch 2 router

© 2013 Cisco Systems, Inc.

The Branch 2 router has one interface to reach the Headquarters router. It uses DLCI 210 to reach the Headquarters router.

Verifying Frame Relay Configuration

This topic describes the Frame Relay **show** commands that you can use to verify that Frame Relay is running as intended.

Verifying Frame Relay Configuration

```
Branch# show interfaces Serial0/0/0
Serial0/0/0 is up, line protocol is up
  Hardware is WIC MBRD Serial
  Internet address is 192.168.1.1/24
  MTU 1500 bytes, BW 1544 Kbit/sec, DLY 20000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation FRAME-RELAY, loopback not set
  Keepalive set (10 sec)
  LMI enq sent 630, LMI stat recvd 616, LMI upd recvd 0, DTE LMI up
  LMI enq recvd 15, LMI stat sent 0, LMI upd sent 0
  LMI DLCI 1023 LMI type is CISCO frame relay DTE
  Broadcast queue 0/64, broadcasts sent/dropped 9/0, interface broadcasts 0
  Last input 00:00:04, output 00:00:04, output hang never
  Last clearing of "show interface" counters 01:45:04
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: weighted fair

<output omitted>
```

Displays interface status, information, and counters

© 2013 Cisco Systems, Inc.

The **show interfaces** command displays information about the encapsulation and Layer 1 and Layer 2 status. Verify that the encapsulation is set to Frame Relay.

The command also displays information about the LMI type and the LMI DLCI.

The output also displays the Frame Relay DTE or DCE type. Normally, the router will be the DTE. However, a Cisco router can be configured as the Frame Relay switch. In this case, the type will be DCE.

The command output represents the basic Frame Relay configuration example.

Verifying Frame Relay Configuration (Cont.)

```
Branch# show frame-relay lmi
LMI Statistics for interface Serial0/0/0 (Frame Relay DTE) LMI TYPE = CISCO
Invalid Unnumbered info 0          Invalid Prot Disc 0
Invalid dummy Call Ref 0          Invalid Msg Type 0
Invalid Status Message 0          Invalid Lock Shift 0
Invalid Information ID 0          Invalid Report IE Len 0
Invalid Report Request 0          Invalid Keep IE Len 0
Num Status Enq. Sent 834          Num Status msgs Rcvd 820
Num Update Status Rcvd 0          Num Status Timeouts 14
Last Full Status Req 00:00:21     Last Full Status Rcvd 00:00:21
```

Displays LMI statistics

© 2013 Cisco Systems, Inc.

Use the **show frame-relay lmi** command to display LMI traffic statistics. For example, this command shows the number of status messages that are exchanged between the local router and the local Frame Relay switch. This command output helps isolate the problem to a Frame Relay communications issue between the carrier switch and your router.

The figure displays sample output that shows the number of status messages that are exchanged between the local router and the local Frame Relay switch.

The table describes a few of the fields in the **show frame-relay lmi** output.

show frame-relay lmi Output Fields

| Field | Description |
|----------------------|--|
| LMI Type | Signaling or LMI specification. The options are Cisco, ANSI, or ITU-T. |
| Num Status Enq. Sent | Number of LMI status inquiry messages that were sent. |
| Num Status Msgs Rcvd | Number of LMI status messages that were received. |

Verifying Frame Relay Configuration (Cont.)

```
Branch# show frame-relay pvc
PVC Statistics for interface Serial0/0/0 (Frame Relay DTE)
      Active      Inactive      Deleted      Static
Local          1             0             0             0
Switched       0             0             0             0
Unused         0             0             0             0
DLCI = 120, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0/0/0
input pkts 18          output pkts 18          in bytes 962
out bytes 962          dropped pkts 0          in pkts dropped 0
out pkts dropped 0    out bytes dropped 0
in FECN pkts 0        in BECN pkts 0         out FECN pkts 0
out BECN pkts 0        in DE pkts 0           out DE pkts 0
out bcast pkts 13     out bcast bytes 442
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
pvc create time 02:32:29, last time pvc status changed 02:32:29
```

Displays PVC statistics

© 2013 Cisco Systems, Inc.

Use the **show frame-relay pvc** command to display the status of each configured PVC as well as traffic statistics.

The table describes a few fields of the **show frame-relay pvc** command output.

show frame-relay pvc Output Fields

| Field | Description |
|------------|--|
| DLCI | One of the DLCI numbers for the PVC. |
| DLCI USAGE | Lists "SWITCHED" when the router is used as a switch or "LOCAL" when the router or access server is used as a DTE device. |
| PVC STATUS | Status of the PVC. The DCE device reports the status, and the DTE device receives the status. When you disable the LMI mechanism on the interface by using the no keepalive command, the PVC status is STATIC. Otherwise, the PVC status is exchanged using the LMI protocol. <ul style="list-style-type: none"> STATIC: LMI is disabled on the interface. ACTIVE: The PVC is operational and can transmit packets. INACTIVE: The PVC is configured but is down. DELETED: The PVC is not present (DTE device only), which means that no status is received from the LMI protocol. |
| INTERFACE | Specific interface that is associated with this DLCI. |

Verifying Frame Relay Configuration (Cont.)

```
Branch# show frame-relay map
Serial0/0/0 (up): ip 192.168.1.2 dlci 120(0x78,0x1C80), dynamic,
broadcast,
CISCO, status defined, active
```

Displays Frame Relay map entries

Use the **show frame-relay map** command to display the current map entries and information about the connections.

The following information explains the **show frame-relay map** output that appears in the example.

- The “120” output is the local DLCI number in decimal.
- The “0x78” output is the hexadecimal conversion of the DLCI number ($0 * 78 = 120$ decimal).
- The “0x1C80” output is the value as it would appear “on the wire” because of the way that the DLCI bits are spread out in the address field of the Frame Relay frame.
- The “192.168.1.2” output is the remote router IP address (a dynamic entry that is learned via the Inverse ARP process).
- Broadcast and multicast are enabled on the PVC.
- The PVC status is active.

Summary

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

Summary

- Frame Relay is a packet-switched, connection-oriented, data-link technology.
- LMI is a signaling standard between the router and the Frame Relay switch. LMI is responsible for managing the connection and maintaining the status between the devices.
- In Frame Relay, a local DLCI must be mapped to a remote destination IP address, either manually or using Inverse ARP.
- To configure basic Frame Relay, set the encapsulation type to Frame Relay on an interface.
- It is recommended that you use Frame Relay point-to-point subinterfaces to solve routing protocol issues.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Introducing VPN Solutions

Cisco VPN solutions provide an Internet-based WAN infrastructure for connecting branch offices, home offices, business partner sites, and remote telecommuters to all or portions of a company network. With cost-effective, high-bandwidth Internet connectivity that is secured by encrypted VPN tunnels, you can reduce WAN bandwidth costs while increasing connectivity speeds. This lesson describes the benefits of VPN implementation.

Objectives

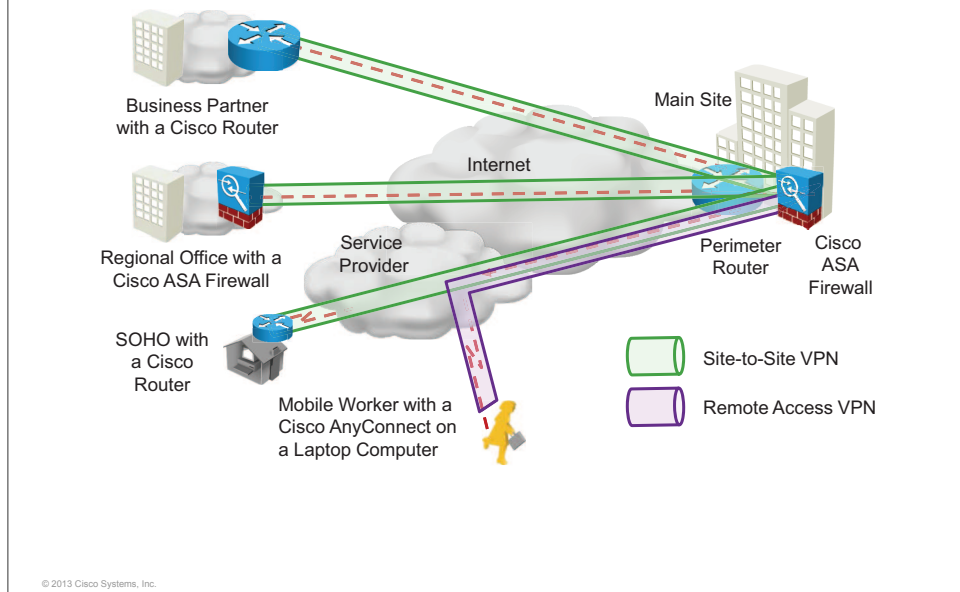
Upon completing this lesson, you will be able to meet these objectives:

- Describe the purpose of VPNs
- List the two Cisco SSL VPN Solutions
- Identify the role of IPsec

VPNs and Their Benefits

This topic describes remote connection options including the VPN solution and its benefits.

VPNs and Their Benefits



Organizations need secure, reliable, and cost-effective networks to connect corporate headquarters, branch offices, and suppliers. With the growing number of teleworkers, enterprises have an increasing need for secure, reliable, and cost-effective ways to connect to people working in SOHOs and other remote locations with resources on corporate sites.

The figure illustrates the remote connection topologies that modern networks use to connect remote locations. In some cases, the remote locations only connect to the headquarters location, while in other cases, remote locations connect to multiple sites. The regional (branch) office in the figure connects to the headquarters and partner sites, while the mobile worker has a single connection to the headquarters.

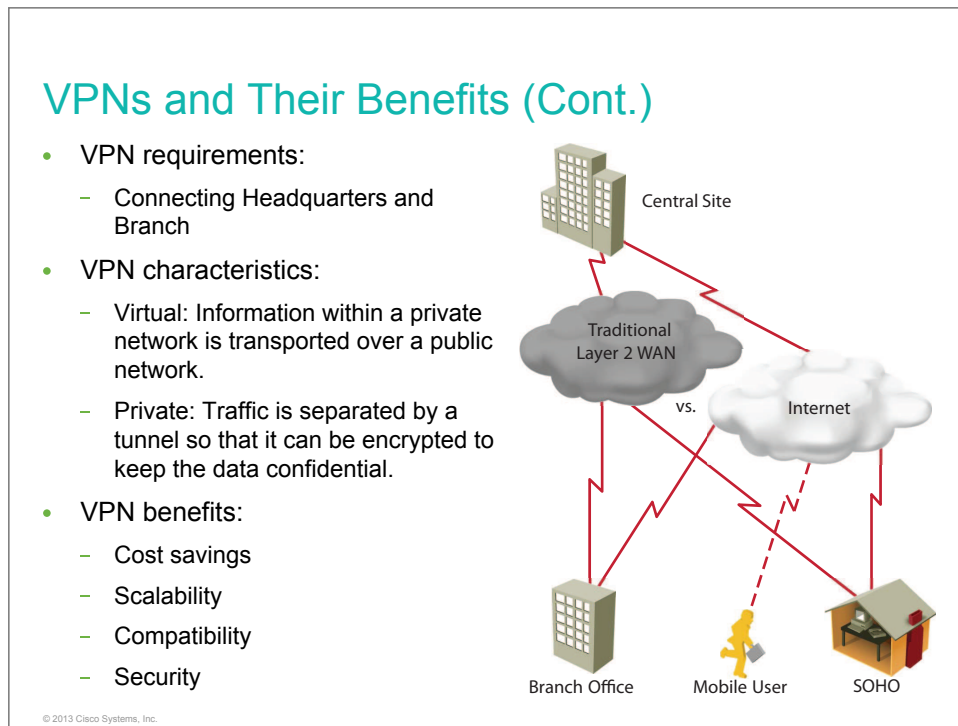
VPNs offer flexible and scalable connectivity. Site-to-site connections can provide a secure, fast, and reliable remote connection. This is the most common option for teleworkers, combined with remote access over broadband, to establish a secure VPN over the public Internet.

With a VPN, the information from a private network is transported over a public network, such as the Internet, to form a virtual network instead of using a dedicated Layer 2 connection. To remain private, the traffic can be encrypted to keep the data confidential. For our purposes, a VPN will be defined as an encrypted connection between private networks over a public network, usually the Internet.

There are two types of VPN networks:

- **Site-to-site VPN:** This is an extension of a classic WAN network. End hosts send and receive traffic through a VPN device, which could be a router or Cisco ASA. This device is responsible for encapsulating and encrypting outbound traffic for all of the traffic from a particular site and sending it through a VPN tunnel over the Internet to a peer VPN device on the target site. Upon receipt, the peer VPN gateway strips the headers, decrypts the content, if encrypted, and relays the packet toward the target host that is inside its private network.
- **Remote-access VPN:** This can support the needs of telecommuters, mobile users, and extranet, consumer-to-business traffic. In a remote-access VPN, each host typically has Cisco AnyConnect VPN Client software that is installed. Whenever the host tries to send any traffic, the Cisco AnyConnect VPN Client software encapsulates and encrypts (if encryption is used) this traffic before sending it over the Internet to the VPN gateway at the edge of the target network. Upon receipt, the VPN gateway behaves as it does for site-to-site VPNs.

Cisco ASA is a network security device that also acts as a VPN end point (concentrator).



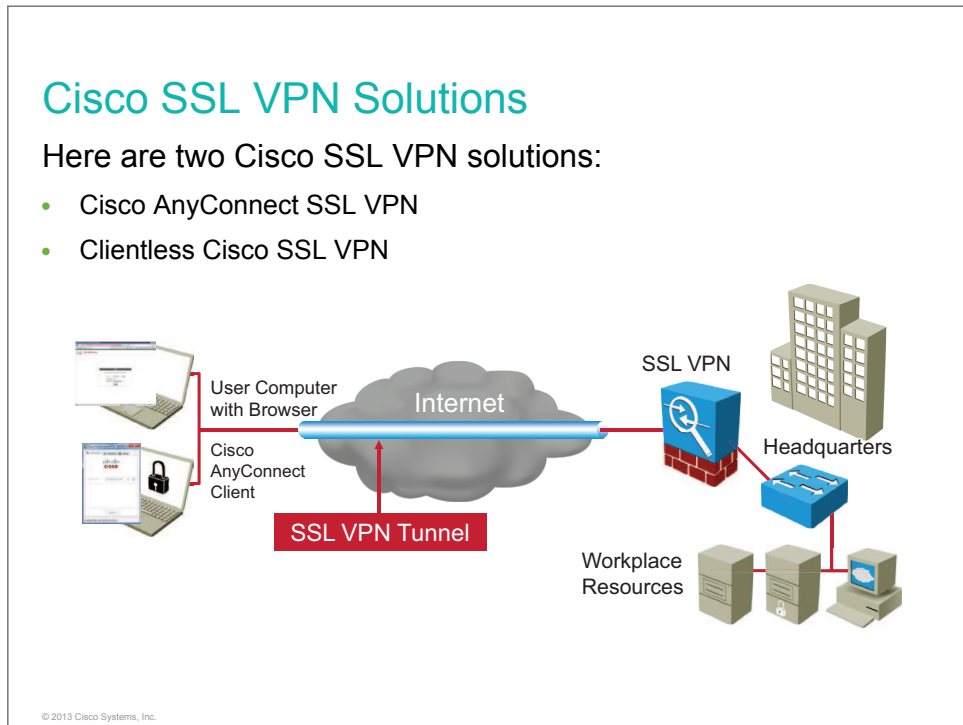
A VPN is an encrypted connection between private networks over a public network such as the Internet. The “V” stands for “virtual” and the “N” stands for “network.” The information from a private network is securely transported over a public network, the Internet, to form a virtual network. The “P” stands for “private.” This means VPN data is separated from the rest of the traffic. VPN traffic can be either encrypted or unencrypted. Instead of using a dedicated Layer 2 connection such as a leased line, a VPN uses virtual connections that are routed through the Internet from the private network of the company to the remote site or employee host.

These are the benefits of a VPN:

- **Cost savings:** VPNs enable organizations to use cost-effective, third-party Internet transport to connect remote offices and remote users to the main corporate site, therefore eliminating expensive, dedicated WAN links. Furthermore, with the advent of cost-effective, high-bandwidth technologies such as DSL, organizations can use VPNs to reduce their connectivity costs while simultaneously increasing remote connection bandwidth.
- **Scalability:** VPNs enable corporations to use the Internet infrastructure within ISPs and devices, which makes it easy to add new users. Therefore, corporations are able to add large amounts of capacity without adding significant infrastructure.
- **Compatibility with broadband technology:** VPNs allow mobile workers, telecommuters, and people who want to extend their work day to take advantage of high-speed, broadband connectivity, such as DSL and cable, to gain access to their corporate networks, providing workers significant flexibility and efficiency. Furthermore, high-speed, broadband connections provide a cost-effective solution for connecting remote offices.
- **Security:** VPNs can provide the highest level of security by using advanced encryption and authentication protocols that protect data from unauthorized access. Two options here are IPsec and SSL.

Cisco SSL VPN Solutions

This topic describes the two Cisco SSL VPN solutions.



Cisco IOS SSL VPN is a technology that provides remote access connectivity from almost any Internet-enabled location using a web browser and its native SSL encryption, or using the Cisco AnyConnect VPN Client that is installed on a PC.

Introducing IPsec

This topic describes the IPsec framework, which provides the security for a VPN.

IPsec Characteristics

- IPsec acts at the network layer, protecting and authenticating IP packets.
- IPsec is a framework of open standards that is algorithm-independent.
- IPsec services provide four critical functions:
 - Confidentiality
 - Data integrity
 - Authentication
 - Anti-replay protection

© 2013 Cisco Systems, Inc.

IPsec is an [IETF](#) standard that defines how a VPN can be configured using the IP addressing protocol. IPsec is not bound to any specific encryption, authentication, security algorithms, or keying technology. IPsec is a framework of open standards that spells out the rules for secure communications. IPsec relies on existing algorithms to implement the encryption, authentication, and key exchange. By not binding IPsec to specific algorithms, IPsec allows newer and better algorithms to be implemented without patching the existing IPsec standards. IPsec provides data confidentiality, data integrity, and origin authentication between participating peers at the IP layer. IPsec secures a path between a pair of gateways, a pair of hosts, or a gateway and a host.

IPsec works at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers). As a result, IPsec can protect virtually all application traffic because the protection can be implemented from Layers 4 through 7. All implementations of IPsec have a plaintext Layer 3 header, so there are no issues with routing. IPsec functions over all Layer 2 protocols such as Ethernet, [ATM](#), or Frame Relay.

IPsec security services provide four critical functions:

- **Confidentiality (encryption):** The sender can encrypt the packets before transmitting them across a network. By doing so, no one can eavesdrop on the communication. If the communication is intercepted, it cannot be read.
- **Data integrity:** The receiver can verify that the data was transmitted through the Internet without being changed or altered in any way. IPsec ensures data integrity by using checksums, which is a simple redundancy check.
- **Authentication:** Authentication ensures that the connection is made with the desired communication partner. The receiver can authenticate the source of the packet by guaranteeing and certifying the source of the information. IPsec uses [IKE](#) to authenticate users and devices that can carry out communication independently. IKE uses several types of authentication including username and password, one-time password, biometrics, [PSKs](#), and digital certificates.

- **Anti-replay protection:** Anti-replay protection verifies that each packet is unique and not duplicated. IPsec packets are protected by comparing the sequence number of the received packets with a sliding window on the destination host or security gateway. A packet that has a sequence number that is before the sliding window is considered either late or a duplicate packet. Late and duplicate packets are dropped.

Do Not Duplicate.
Post beta, not for release.

Summary

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

Summary

- Organizations implement VPNs because they are less expensive, easier to scale than traditional WANs, and can provide security.
- A site-to-site VPN is an extension of a classic WAN network.
- Remote access VPNs can support the needs of telecommuters, mobile users, and extranet, consumer-to-business traffic.
- IPsec protects and authenticates IP packets and is a framework of open standards that is algorithm-independent.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Configuring GRE Tunnels

GRE is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE enables network expansion across a single-protocol backbone environment.

Objectives

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

- Describe GRE tunneling
- Configure a GRE tunnel
- Verify a GRE tunnel

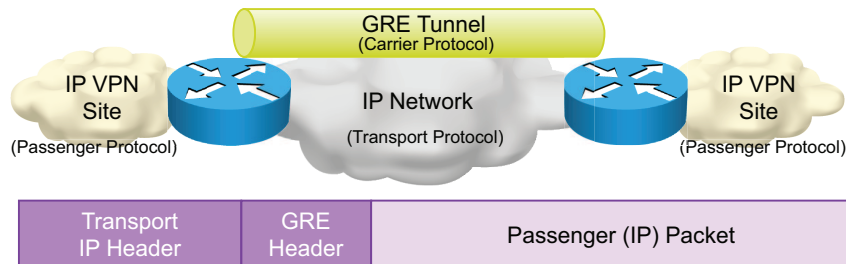
GRE Tunnel Overview

This topic describes the purpose of implementing a GRE tunnel.

GRE Tunnel Overview

GRE = Generic Routing Encapsulation:

- One of many tunneling protocols
- IP protocol 47: defines GRE packets
- Allows routing information to be passed between connected networks
- No encryption



© 2013 Cisco Systems, Inc.

GRE, developed by Cisco, is one of many tunneling protocols. It is a simple, general-purpose protocol that is designed to manage the transportation of multiprotocol and IP multicast traffic between two or more sites that may only have IP connectivity. GRE is designed to encapsulate arbitrary types of network layer packets inside arbitrary types of network layer packets, as defined in RFC 1701, *Generic Routing Encapsulation (GRE)*; RFC 1702, *Generic Routing Encapsulation over IPv4 Networks*; and RFC 2784, *Generic Routing Encapsulation (GRE)*.

A tunnel interface supports a header for each of the following:

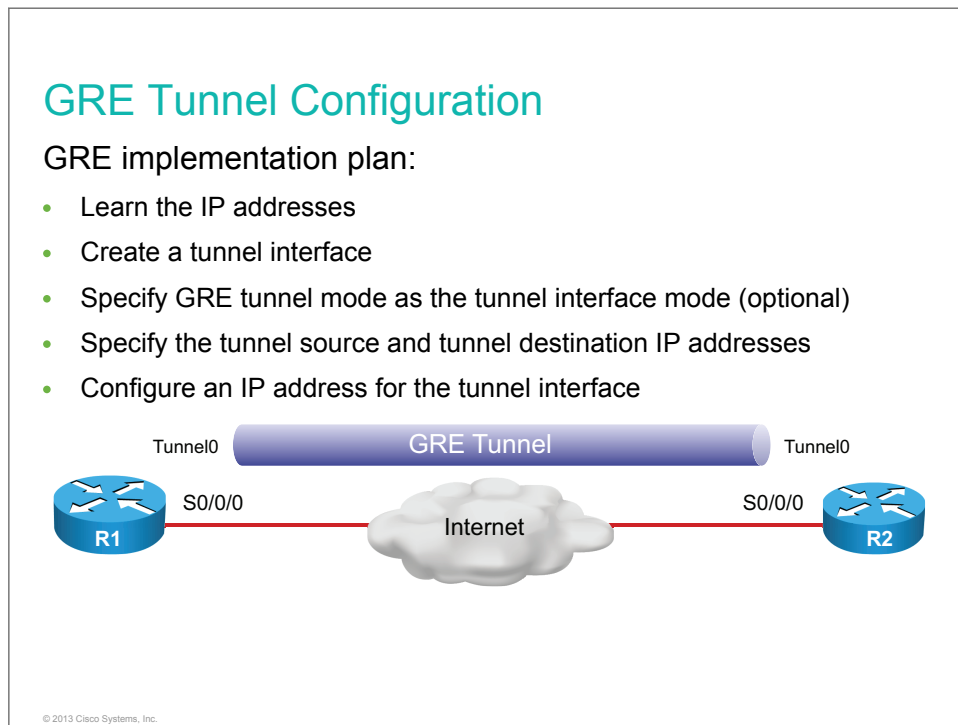
- A passenger protocol or encapsulated protocol, such as IPv4 or IPv6. This protocol is the one that is being encapsulated.
- A carrier or encapsulation protocol (GRE, in this case).
- A transport delivery protocol, such as IP, which is the protocol that carries the encapsulated protocol.

GRE has these characteristics:

- GRE encapsulation uses a protocol-type field in the GRE header to support the encapsulation of any OSI Layer 3 protocol.
- GRE itself is stateless. It does not include any flow-control mechanisms, by default.
- GRE does not include any strong security mechanisms to protect its payload.
- The GRE header, together with the tunneling IP header, creates at least 24 bytes of additional overhead for tunneled packets.

GRE Tunnel Configuration

This topic describes how to configure a GRE tunnel.



After you have assessed your requirements, you can create an implementation plan. To implement a GRE tunnel, you must do the following:

- Learn the IP addresses.
- Create a tunnel interface.
- Optionally, specify GRE tunnel mode as the tunnel interface mode. (GRE tunnel mode is the default tunnel interface mode of Cisco IOS Software.)
- Specify the tunnel source IP address.
- Specify the tunnel destination IP address.
- Configure an IP address for the tunnel interface.

GRE Tunnel Configuration (Cont.)



```
Branch(config)# interface Tunnel0
Branch(config-if)# tunnel mode gre ip
Branch(config-if)# ip address 192.168.2.1 255.255.255.0
Branch(config-if)# tunnel source 209.165.201.1
Branch(config-if)# tunnel destination 209.165.202.130
```

Configuration of GRE tunnel on the Branch router

© 2013 Cisco Systems, Inc.

GRE Tunnel Configuration (Cont.)



```
HQ(config)# interface Tunnel0
HQ(config-if)# tunnel mode gre ip
HQ(config-if)# ip address 192.168.2.2 255.255.255.0
HQ(config-if)# tunnel source 209.165.202.130
HQ(config-if)# tunnel destination 209.165.201.1
```

Configuration of GRE tunnel on the Headquarters router

© 2013 Cisco Systems, Inc.

The sample configuration illustrates a basic GRE tunnel configuration on the Branch and Headquarters routers. The minimum configuration requires specification of the tunnel source and destination addresses. You must also configure an IP subnet to provide IP connectivity across the tunnel link.

The figure shows a tunnel configuration for the Branch and Headquarters routers. Both tunnel interfaces have the tunnel source set as the local GigabitEthernet 0/1 interface and the tunnel destination set as the peer router GigabitEthernet0/1 interface. The IP address is assigned to the tunnel interfaces on both routers.

| Command | Description |
|---|--|
| tunnel mode gre ip | Specifies GRE tunnel mode as the tunnel interface mode in interface tunnel configuration mode. This is the default tunnel mode on Cisco routers, so, in fact, it is not necessary to enter this command. |
| tunnel source <i>ip_address</i> | Specifies the tunnel source IP address in interface tunnel configuration mode. |
| tunnel destination <i>ip_address</i> | Specifies the tunnel destination IP address in interface tunnel configuration mode. |
| ip address <i>ip_address mask</i> | Specifies the IP address of the tunnel interface. |

Do Not Duplicate:
Post beta, not for release.

GRE Tunnel Verification

This topic describes how to verify a GRE tunnel.

GRE Tunnel Verification

```
Branch# show ip interface brief | include Tunnel
Tunnel0          192.168.2.1      YES manual up      up
```

Verifies that the tunnel interface is up

```
Branch# show interface Tunnel 0
Tunnel0 is up, line protocol is up
Hardware is Tunnel
Internet address is 192.168.2.1/24
MTU 17916 bytes, BW 100 Kbit/sec, DLY 50000 usec,
    reliability 255/255, txload 1/255, rxload 1/255
Encapsulation TUNNEL, loopback not set
Keepalive not set
Tunnel source 209.165.201.1, destination 209.165.202.130
Tunnel protocol/transport GRE/IP
<output omitted>
```

Verifies that the tunnel interface is up and shows tunnel IPs, source and destination IPs, and tunnel protocol

© 2013 Cisco Systems, Inc.

To determine whether the tunnel interface is up or down, use the **show ip interface brief** command. You can verify the state of a GRE tunnel by using the **show interface tunnel** command. The line protocol on a GRE tunnel interface is up as long as there is a route to the tunnel destination.

GRE Tunnel Verification (Cont.)

```
Branch# show ip route
<output omitted>
    192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.2.0/24 is directly connected, Tunnel0
L    192.168.2.1/32 is directly connected, Tunnel0
    209.165.201.0/24 is variably subnetted, 2 subnets, 2 masks
C    209.165.201.0/27 is directly connected, GigabitEthernet0/1
L    209.165.201.1/32 is directly connected, GigabitEthernet0/1
```

Verifies the tunnel route between the Branch and Headquarters routers

© 2013 Cisco Systems, Inc.

By issuing the **show ip route** command, you can identify the route between the Branch and Headquarters routers. Because a tunnel is established between the two routers, the path is seen as directly connected.

Summary

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

Summary

- GRE is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside IP tunnels.
- You must configure a tunnel source and tunnel destination to establish a GRE tunnel as the IP address of the tunnel itself.
- You should verify that the tunnel interface is up after configuring it.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Module Summary

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

Module Summary

- A WAN can be interconnected over a private infrastructure or over a public infrastructure such as the Internet.
- PPP is a common Layer 2 protocol for the WAN. There are two components of PPP: LCP, which negotiates the connection, and NCP, which encapsulates traffic.
- Frame Relay is a packet-switched, connection-oriented, data-link technology.
- Organizations implement VPNs because they are less expensive and easier to scale than traditional WANs, while still offering mechanisms for secure communication.
- GRE is a tunneling protocol that can encapsulate a wide variety of protocol packet types inside of IP tunnels, but it does not provide encryption.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Questions

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

1. Which two statements about WANs are true? (Choose two.) (Source: Configuring Serial Encapsulation)
 - A. WANs generally connect devices that are located over a broader geographical area.
 - B. WANs generally connect devices that are close to each other.
 - C. WAN stands for World Around Networks.
 - D. WANs use connections of various types to provide access to bandwidth over large geographical areas.
2. Which feature does PPP use to encapsulate multiple protocols? (Source: Configuring Serial Encapsulation)
 - A. NCP
 - B. LCP
 - C. IPCP
 - D. IPXP
3. What is the purpose of LCP? (Source: Configuring Serial Encapsulation)
 - A. to perform authentication
 - B. to negotiate control options
 - C. to encapsulate multiple protocols
 - D. to specify asynchronous versus synchronous

4. Which two statements best describe CHAP? (Choose two.) (Source: Configuring Serial Encapsulation)
- A. CHAP is performed periodically.
 - B. CHAP uses a two-way handshake.
 - C. CHAP uses a three-way handshake.
 - D. CHAP uses a two-way hash function.
 - E. CHAP passwords are sent in plaintext.
5. With CHAP, how does a remote node respond to a challenge message? (Source: Configuring Serial Encapsulation)
- A. with a hash value
 - B. with a return challenge
 - C. with a plaintext password
 - D. with an encrypted password
6. Match each Frame Relay operation component with its definition. (Source: Establishing a WAN Connection Using Frame Relay)
- | | | |
|----------------------|--------------------------|--|
| A. SVC | <input type="checkbox"/> | method of dynamically associating a remote network layer address with a local DLCI |
| B. Inverse ARP | <input type="checkbox"/> | VC that is dynamically established on demand and is torn down when transmission is complete |
| C. LMI | <input type="checkbox"/> | maximum average data rate |
| D. CIR | <input type="checkbox"/> | signaling standard between the router device and the Frame Relay switch that is responsible for managing the connection and maintaining status between the devices |
| E. local access rate | <input type="checkbox"/> | clock speed of the connection to the Frame Relay cloud |
7. What identifies the logical circuit between the router and the local Frame Relay switch? (Source: Establishing a WAN Connection Using Frame Relay)
- A. DLCI
 - B. LMI signal
 - C. FECN packet
 - D. BECN packet
8. Which characteristic of Frame Relay can cause reachability issues when a single interface is used to interconnect multiple sites? (Source: Establishing a WAN Connection Using Frame Relay)
- A. intermittent
 - B. point-to-point
 - C. error-correcting
 - D. NBMA

9. Which VC status on a Cisco router indicates that the local connection to the Frame Relay switch is working, but the remote router connection to the Frame Relay switch is not working? (Source: Establishing a WAN Connection Using Frame Relay)
- A. LMI state
 - B. active state
 - C. deleted state
 - D. inactive state
10. What are two types of VPNs? (Choose two.) (Source: Introducing VPN Solutions)
- A. remote-access
 - B. remote-to-site
 - C. remote-to-remote
 - D. site-to-site
11. Which two of these options are advantages of VPNs over traditional Layer 2 WANs? (Choose two.) (Source: Introducing VPN Solutions)
- A. are less expensive
 - B. provide scalability
 - C. require less processing for routers
 - D. require less knowledge to install and maintain
12. Which command is used to specify GRE tunnel mode as the tunnel interface mode? (Source: Configuring GRE Tunnels)
- A. **tunnel mode ip gre**
 - B. **tunnel mode gre ip**
 - C. **tunnel gre ip**
 - D. **tunnel gre**
13. Which two commands could you use to verify that a GRE tunnel is up? (Choose two.) (Source: Configuring GRE Tunnels)
- A. **show ip interface brief**
 - B. **show interface tunnel**
 - C. **show gre tunnel**
 - D. **show tunnel interface**

Answer Key

1. A, D
2. A
3. B
4. A, C
5. A
6. A. local access rate
B. SVC
C. CIR
D. LMI
E. Inverse ARP
clock speed of the connection to the Frame Relay cloud
VC that is dynamically established on demand and is torn down when transmission is complete
maximum average data rate
signaling standard between the router device and the Frame Relay switch that is responsible for managing the connection and maintaining status between the devices
method of dynamically associating a remote network layer address with a local DLCI
7. A
8. D
9. D
10. A, D
11. A, B
12. B
13. A, B

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Network Device Management

Network staff is responsible for managing each device on the network according to best industry practices and for reducing device downtime. This module describes commands and processes to determine network operational status, gather information about remote devices, and manage Cisco IOS Software images, configuration files, and devices on a network. The module also explains how to enable Cisco IOS Software feature sets by obtaining and validating a Cisco software license.

Objectives

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

- Describe how network devices can be managed and monitored
- Describe the management of Cisco devices
- Explain licensing

Do Not Duplicate.
Post beta, not for release.

Configuring Network Devices to Support Network Management Protocols

This lesson provides an overview of some of the tools for monitoring and troubleshooting Cisco devices. SNMP basic components and operation are explained. A graphing tool can use SNMP to periodically poll an SNMP agent (for example, a router) and graph the values. Syslog is another useful protocol that allows a network device to send event notification messages to a centralized logging server. The lesson ends with a basic explanation of the NetFlow protocol. The NetFlow protocol creates an environment and the tools to understand network traffic and how it is flowing.

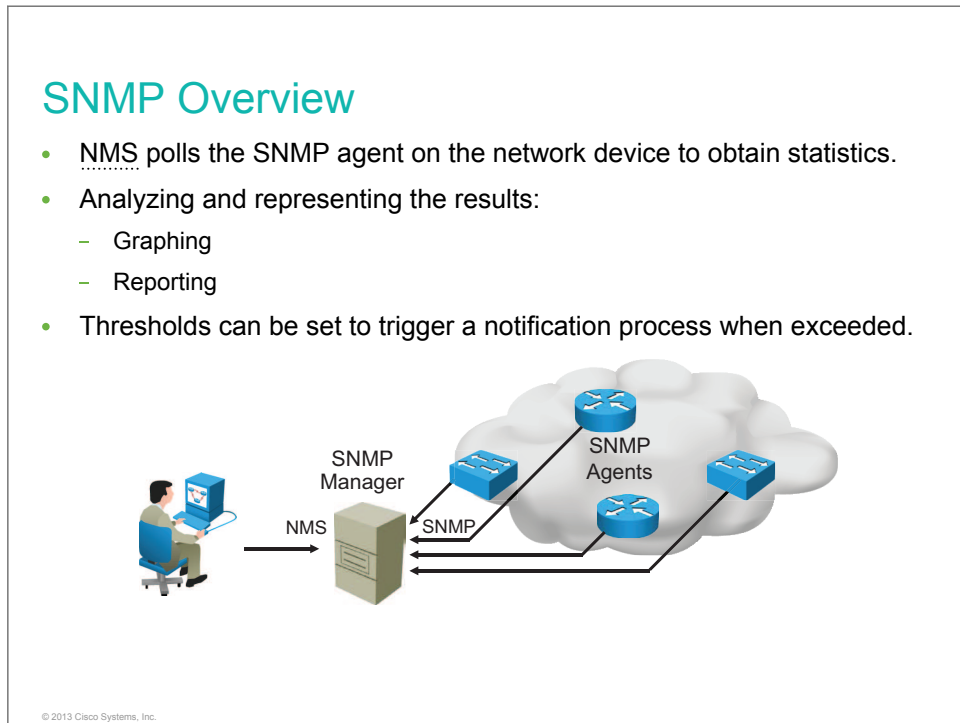
Objectives

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

- Explain why SNMP is used
- Compare and describe different SNMP versions
- Describe how to obtain data from an SNMP agent
- Configure a Cisco device for SNMP access
- Explain why syslog is used
- Describe the format of syslog messages
- Configure syslog on a Cisco device
- Describe the purpose of NetFlow
- Describe the NetFlow architecture
- Configure and verify NetFlow on a Cisco device

SNMP Overview

This topic describes the purpose of SNMP.



SNMP is an application layer protocol that provides a message format for communication between managers and agents. The SNMP system consists of an SNMP manager, an SNMP agent, and an MIB. The SNMP manager can be part of an NMS such as Cisco Prime Infrastructure. The agent and MIB reside on the network device. To configure SNMP on the device, you define the relationship between the manager and the agent.

Routers and other network devices keep statistics about the information of their processes and interfaces locally. SNMP on a device runs a special process that is called an agent. This agent can be queried, using SNMP, to obtain the values of statistics or parameters. By periodically querying or “polling” the SNMP agent on a device, statistics can be gathered and collected over time by an NMS. This data can then be processed and analyzed in various ways. Averages, minimums, or maximums can be calculated, the data can be graphed, or thresholds can be set to trigger a notification process when they are exceeded. The NMS polls devices periodically to obtain the values of the MIB objects that it is set up to collect.

The SNMP agent contains MIB variables, and the SNMP manager can request or change the values. A manager can get a value from an agent or store a value in the agent. The agent gathers data from the MIB, the repository for information about device parameters, and network data. The agent can also respond to manager requests to get or set data.

SNMP Versions

This topic compares different SNMP versions.

| SNMP Version | Security | Bulk Retrieval Mechanism |
|--------------|---|--------------------------|
| SNMPv1 | Plaintext authentication with community strings | No |
| SNMPv2c | Plaintext authentication with community strings | Yes |
| SNMPv3 | Strong authentication, confidentiality, and integrity | Yes |

© 2013 Cisco Systems, Inc.

SNMP is implemented in versions 1, 2c, and 3.

Both SNMPv1 and SNMPv2c use a plaintext, community-based form of security to authenticate communication between managers and agents.

SNMPv2c introduced a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism supports the retrieval of tables and large quantities of information, minimizing the number of round trips that are required.


SNMPv3 primarily introduced security features such as confidentiality, integrity, and authentication of messages between managers and agents.

Obtaining Data from an SNMP Agent

This topic describes how to obtain data from an SNMP agent.

Obtaining Data from an SNMP Agent

An SNMP graphing tool periodically polls an SNMP agent (for example, a router) and graphs obtained values:



© 2013 Cisco Systems, Inc.

Because the CPU is one of the key resources, it should be measured continuously. You should gather CPU statistics at the NMS and graph the statistics. You should observe CPU utilization for a long time and define which values are still acceptable. You can set thresholds to these values so that when CPU utilization exceeds this threshold, notifications are sent.

The figure shows a 5-minute CPU utilization graph of a router that was taken using the open source tool called Cacti. The SNMP graphing tool periodically polls the SNMP agent, in this case a router, and graphs obtained values.

Obtaining Data from an SNMP Agent (Cont.)

- MIB is a collection of information that is organized hierarchically.
- **OIDs** uniquely identify managed objects in an MIB.
 - A 5-minute, exponentially moving average of the CPU busy percentage: 1.3.6.1.4.1.9.2.1.58.0

The figure illustrates the process of obtaining data from an SNMP agent. It shows a MIB tree structure on the left, a detailed view of an SNMP object on the right, and a terminal window at the bottom. The terminal window shows the command `snmpget -v2c -c community 10.250.250.14 1.3.6.1.4.1.9.2.1.58.0` and its output: `SNMPv2-SMI::enterprises.9.2.1.58.0 = INTEGER: 11`. Arrows point from labels 'version', 'community', 'IP address', and 'OID number' to the corresponding parts of the command and output. A label 'obtained CPU value' points to the output value '11'.

An MIB is a collection of information that is organized hierarchically. The information is then accessed using a protocol such as SNMP. OIDs uniquely identify managed objects in an MIB hierarchy. OIDs can be depicted as a tree, the levels of which are assigned by different organizations. Top-level MIB OIDs belong to various standard organizations. Vendors define private branches, including managed objects for their own products.

OIDs belonging to Cisco, as shown in the figure, are numbered as follows: .iso (1), .org (3), .dod (6), .internet (1), .private (4), .enterprises (1), and .cisco (9).

The example in the figure shows the output of an SNMP application called "snmpget," which is issued on the network management station. Using the snmpget application, you can manually obtain values for the 5-minute exponentially moving average of the CPU busy percentage, as the output in the figure shows. You must specify the SNMP version, the correct community, the IP address of the network device that you want to query, and the OID number.

SNMP community strings authenticate access to MIB objects and function as embedded passwords. In order for the NMS to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the network device.

SNMP Configuration

This topic describes how to configure a Cisco device for SNMP access.

SNMP Configuration

- Enable SNMP read-write access to the router
- Configure SNMP contact
- Configure SNMP location

Get some useful information from the router via SNMP

© 2013 Cisco Systems, Inc.

To implement SNMP access to the router, you must do the following:

- On the router, configure a community access string with a read-write privilege to permit access to the SNMP.
- Set the system contact and location of the SNMP agent on the router.

SNMP Configuration (Cont.)

SNMP viewer installed

```
R1(config)# snmp-server community Cisco RW
R1(config)# snmp-server location San Jose
R1(config)# snmp-server contact Joe Summer
```

SNMP configuration on R1

© 2013 Cisco Systems, Inc.

Configuration of SNMP is based on the steps described in the table.

| Command | Description |
|---|--|
| <code>snmp-server community string [ro rw]</code> | Defines the community access string with a read-only or read-write privilege |
| <code>snmp-server contact contact_name</code> | Sets the system contact string |
| <code>snmp-server location location</code> | Sets the system location string |

Note The first `snmp-server` command that you issue enables SNMP on the device.

A community string authenticates access to MIB objects and can have one of these attributes:

- **Read-only:** Gives read access to authorized management stations to all objects in the MIB except the community strings but does not allow write access.
- **Read-write:** Gives read and write access to authorized management stations to all objects in the MIB but does not allow access to the community strings.

The sample configuration illustrates a basic SNMP configuration on router R1.

The community access string, "Cisco," is configured to permit read-write SNMP access to router R1. This means that the NMS will be able to retrieve and modify MIB objects from router R1. NMS would retrieve MIB objects, for example, for generating graphs of CPU usage. Any NMS that tries to obtain SNMP information from router R1 must have the community specified as "Cisco."

System contact and the location of the SNMP agent is also set on the router so that these descriptions can be accessed through the configuration file. Configuring the basic information is recommended because it may be useful when troubleshooting your configuration.

Syslog Overview

This topic describes the purpose of syslog.

Syslog Overview

- Syslog is a protocol that allows a network device to send event notification messages across IP networks to event message collectors.
- A device can be configured so that it generates a syslog message and forwards it to various destinations:
 - logging buffer
 - console line
 - terminal lines
 - syslog server

© 2013 Cisco Systems, Inc.

Syslog is a protocol that allows a machine to send event notification messages across IP networks to event message collectors. By default, a network device sends the output from system messages and **debug** privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a syslog server, depending on your configuration. The process also sends messages to the console. Logging services provide a means to gather logging information for monitoring and troubleshooting, to select the type of logging information that is captured, and to specify the destinations of captured syslog messages.

You can set the severity level of the messages to control the type of messages that are displayed on the consoles and each of the destinations. You can time-stamp log messages or set the syslog source address to enhance real-time debugging and management.

You can access logged system messages by using the device CLI or by saving them to a correctly configured syslog server. The switch or router software saves syslog messages in an internal buffer.

You can remotely monitor system messages by viewing the logs on a syslog server or by accessing the device through Telnet, SSH, or through the console port.

Syslog Message Format

This topic describes the format of syslog messages.

Syslog Message Format

The general format of syslog messages is generated by the syslog process on Cisco IOS Software:

```
seq no:timestamp: %facility-severity-MNEMONIC:description
```

An example of a syslog message is informing the administrator that FastEthernet 0/22 came up.

```
*Apr 22 11:05:55.423: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/22, changed state to up
```

© 2013 Cisco Systems, Inc.

By default, the general format of syslog messages that are generated by the syslog process on the Cisco IOS Software is as follows:

```
seq no:timestamp: %facility-severity-MNEMONIC:description
```

The items that are contained in the Cisco IOS Software syslog message are explained in the following table.

| Item | Explanation |
|--------------------|--|
| seq no | Stamps log messages with a sequence number only if the service sequence-numbers global configuration command is configured |
| timestamp | Date and time of the message or event, which appears only if the service timestamps log [datetime log] global configuration command is configured |
| facility | The facility to which the message refers (for example, SNMP, system, and so on) |
| severity | Single-digit code from 0 to 7 that is the severity of the message |
| MNEMONIC | Text string that uniquely describes the message |
| description | Text string containing detailed information about the event being reported |

The eight message severity levels from the most severe level to the least severe level are explained in the table.

| Severity Level | Explanation |
|-----------------------------------|----------------------------------|
| Emergency (severity 0) | System is unusable |
| Alert (severity 1) | Immediate action needed |
| Critical (severity 2) | Critical condition |
| Error (severity 3) | Error condition |
| Warning (severity 4) | Warning condition |
| Notification (severity 5) | Normal but significant condition |
| Informational (severity 6) | Informational message |
| Debugging (severity 7) | Debugging message |

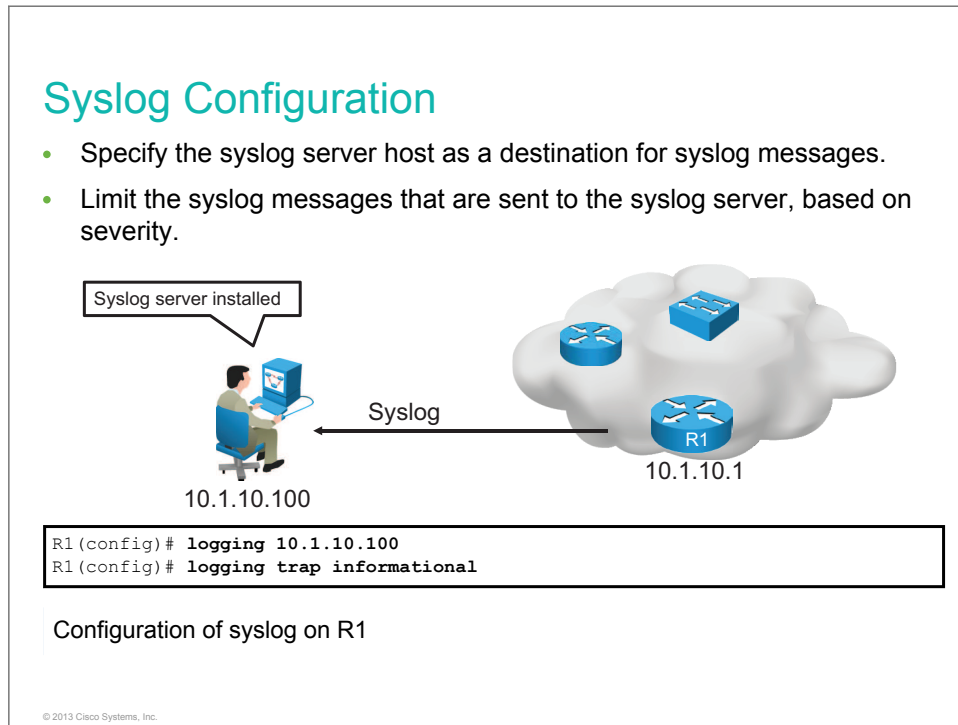
If severity level 0 is configured, this means only Emergency-level messages will be displayed. For example, if severity level 4 is configured, all messages with severity levels up to 4 will be displayed (Emergency, Alert, Critical, Error, and Warning).

The highest severity level is level 7, which is the Debugging-level message. Much information can be displayed at this level, and it can even hamper the performance of your network. Use it with caution.

Do Not Duplicate. Post beta, not for release.

Syslog Configuration

This topic describes how to configure a Cisco device to send logs to the server.



To implement a syslog configuration, specify a syslog server host as a destination for syslog messages and limit the syslog messages that are sent to the syslog server, based on the severity.

Configuration of syslog is based on commands that are described in the table.

| Command | Description |
|--|--|
| logging {hostname ip-address} | Identifies a syslog server host to receive logging messages |
| logging trap severity | Limits the syslog messages that are sent to the syslog server, based on severity |

The figure shows configurations for logging syslog messages to a syslog server with IP address 10.1.10.100, where you can observe syslog messages.

The **logging** command identifies a syslog server host to receive logging messages. By issuing this command more than once, you build a list of syslog servers that receive logging messages. You can limit the syslog messages that are sent to the syslog server, based on severity, using the **logging trap** command.

NetFlow Overview

This topic describes the purpose of NetFlow.

NetFlow Overview

- NetFlow is an application for collecting IP traffic information.
- Reports from NetFlow are like a phone bill.
- NetFlow enables the following:
 - Measuring who uses network resources
 - Accounting and charging for resource utilization
 - Using the measured information to do effective network planning
 - Using the measured information to customize applications and services

© 2013 Cisco Systems, Inc.

NetFlow is a Cisco IOS application that provides statistics on packets flowing through the routing devices in the network. It is emerging as a primary network accounting and security technology.

Cisco IOS NetFlow efficiently provides a key set of services for IP applications, including network traffic accounting, usage-based network billing, network planning, security, DoS monitoring capabilities, and network monitoring.

NetFlow technology creates an environment in which you have the tools to understand how network traffic is flowing. To use an analogy from the telephone industry, NetFlow is like a phone bill, from which you can learn who is talking to whom, how frequently, how long, and so on.

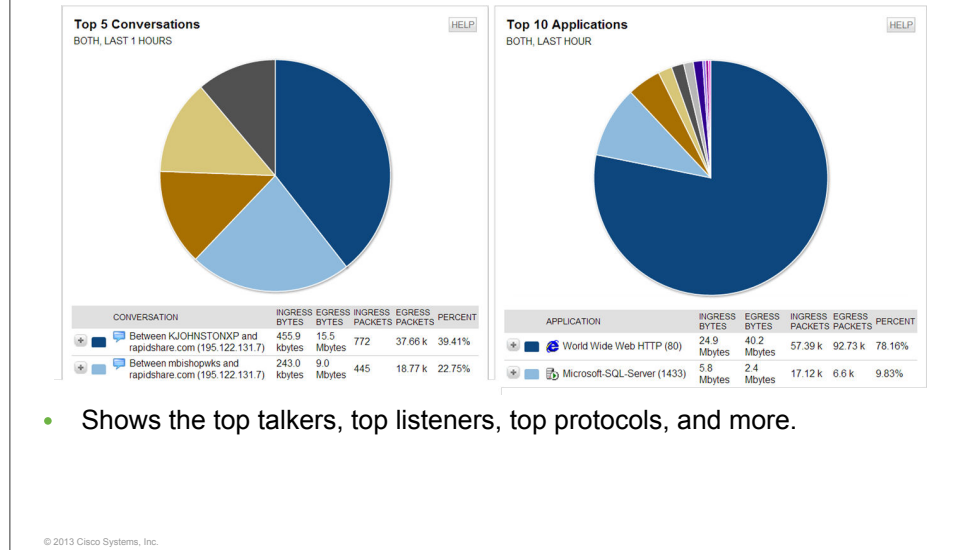
NetFlow is completely transparent to the existing network, including end stations and application software, and network devices such as LAN switches. Also, NetFlow capture and export are performed independently on each internetworking device. NetFlow does not need to be operational on each router in the network.

NetFlow is important for service providers and enterprise customers because it helps address four key requirements:

- Efficiently measuring who is using which network resources for which purpose
- Accounting and charging back according to the resource utilization level
- Using the measured information to do more effective network planning so that resource allocation and deployment are well-aligned with customer requirements
- Using the information to better structure and customize the set of available applications and services to meet user needs and customer service requirements

NetFlow Overview (Cont.)

Example of analysis on a NetFlow collector:



- Shows the top talkers, top listeners, top protocols, and more.

There are several NetFlow collectors and analyzers available on the market. These tools enable you to analyze the traffic on your network by showing the top talkers, top listeners, top protocols, and more. You can see the types of traffic (web, mail, FTP, peer-to-peer, and so on) that are on the network, as well as which devices are sending and receiving most of the traffic. Collecting data provides you with forensic-level telemetry on top talkers, top hosts, and top listeners. And because data is preserved over time, you will be able to analyze network usage trends.

Based on the usage of NetFlow analyzers, you will be able to identify the following:

- The major users of the network
- The websites that are routinely visited and what is downloaded
- Who is generating the most traffic
- If you have enough bandwidth to support mission-critical activity
- Who is using excessive bandwidth

NetFlow Overview (Cont.)

- NetFlow components:
 - NetFlow-enabled network devices
 - NetFlow collector
- NetFlow devices generate NetFlow records that are exported and then collected by a NetFlow collector. Cisco Network Analysis Module is an example of a NetFlow collector. It also processes NetFlow data and provides the results through its GUI.



© 2013 Cisco Systems, Inc.

NetFlow components include the following:


- Network devices that are configured for NetFlow
- NetFlow Collector, which receives NetFlow information from network devices

Network devices collect IP traffic statistics on interfaces where NetFlow is configured. Network devices then export these statistics as NetFlow records to a central server that runs NetFlow Collector software. The Collector also performs traffic analysis.

NetFlow Overview (Cont.)

Cisco defines a flow as a unidirectional sequence of packets with seven common values:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Layer 3 protocol type
- ToS
- Input logical interface



© 2013 Cisco Systems, Inc.

A NetFlow network flow is defined as a unidirectional stream of packets between a given source and destination. The source and destination are each defined by a network layer IP address and transport layer source and destination port numbers.

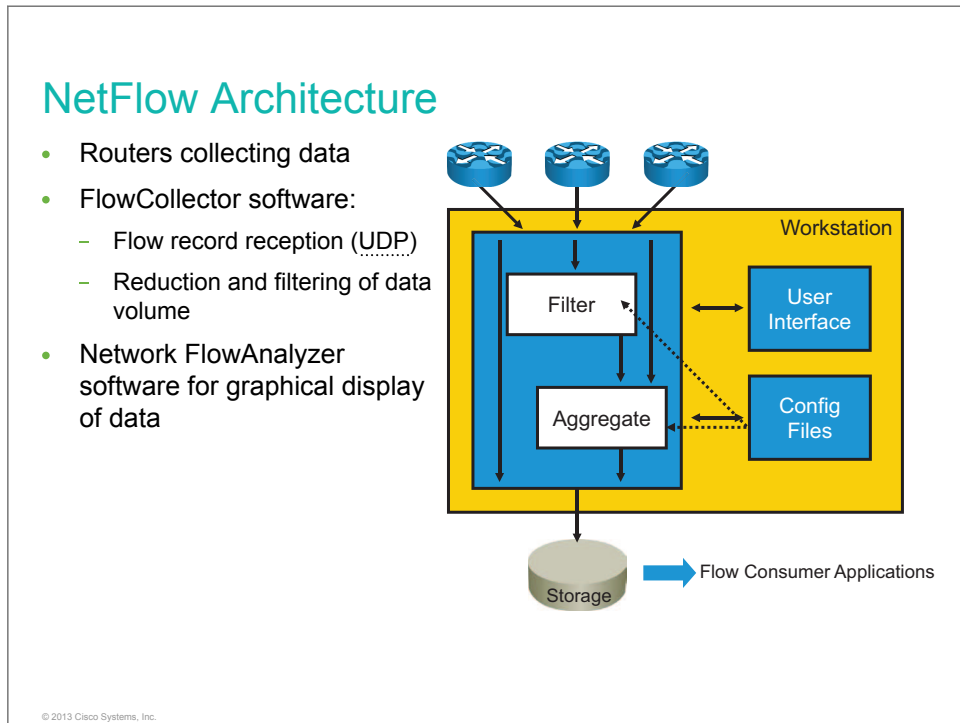
A flow is defined by the combination of seven key fields, which define a unique flow. If a packet has one key field that is different from another packet, it is considered to belong to another flow. Flows are stored in the NetFlow cache.

Note ToS is a field in the IP header. It is used as a mechanism to mark IP packets with different priorities in order to receive different treatment in terms of throughput, reliability, and latency.

Do Not Duplicate:
Post beta, not for release.

NetFlow Architecture

This topic describes NetFlow architecture.



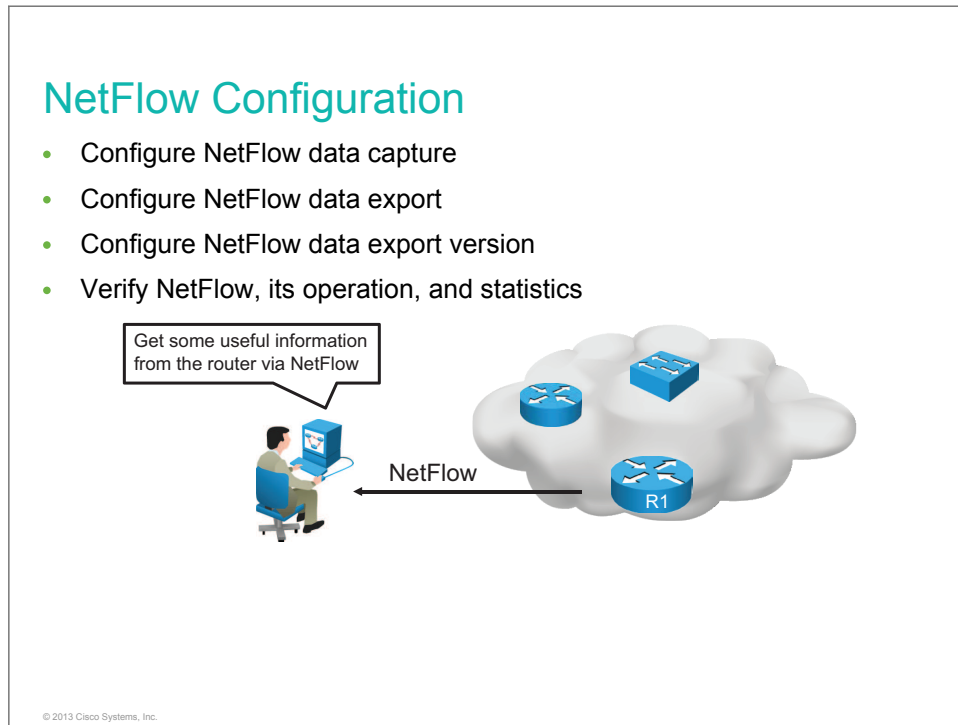
The purpose of FlowCollector is to grab flow export data from multiple routers, filter and aggregate the data according to the policies of the customer (that is, keep the data that you want and discard the rest), and store this summarized or aggregated data instead of raw flow data to minimize disk space consumption. The data is written to a disk at specified intervals in flat files. The customer may run multiple collection schemes or threads concurrently; for example, customers who want data both for planning and for billing may want to store different “cuts” of the data to support these applications via different aggregation schemes.

FlowCollector passively listens to specified UDP ports to receive and process exported NetFlow datagrams. The FlowCollector application provides a high-performance, easy-to-use solution that scales to accommodate consumption of NetFlow export data from multiple devices in order to support key flow-consumer applications, including accounting, billing, and network planning and monitoring.

FlowAnalyzer provides the means to do near real-time visualization and analysis of recorded and aggregated flow data. You can specify the router, the aggregation scheme, and the time interval in which you wish to view and then retrieve the relevant data. You can then sort and visualize the data in a manner that makes sense for the users (with bar charts, pie charts, or histograms of the sorted reports). The data can then be exported to spreadsheets such as Excel for more detailed analysis, trending, reporting, and so on.

NetFlow Configuration

This topic describes how to configure and verify NetFlow on a Cisco device.




To implement NetFlow on a router, you must do the following:

- **Configure NetFlow data capture:** NetFlow can capture data from ingress (incoming) and egress (outgoing) packets.
- **Configure NetFlow data export:** You need to specify the IP address or hostname of the NetFlow collector and the UDP port that the NetFlow collector listens to.
- **Configure NetFlow data export version:** You can specify the version of NetFlow data export format.
- **Verify NetFlow, its operation, and statistics:** After you have configured NetFlow, you can analyze the exported data on a workstation running an application such as NetFlow Collection Engine or using several **show** commands on the router itself.

Note

Be aware that NetFlow consumes additional memory. If you have memory constraints, you can preset the size of the NetFlow cache so that it contains a smaller number of entries. The default cache size depends on the platform.

NetFlow Configuration (Cont.)



```
R1(config)# interface GigabitEthernet0/1
R1(config-if)# ip flow ingress
R1(config-if)# ip flow egress
R1(config-if)# exit
R1(config)# ip flow-export destination 10.1.10.100 9996
R1(config)# ip flow-export version 9
```

Configuration of NetFlow on router R1

© 2013 Cisco Systems, Inc.

Configuration of NetFlow is based on the steps that are described in the table.

| Command | Description |
|--|--|
| ip flow {ingress egress} | Enables NetFlow on the interface. Captures traffic that is being received or being transmitted by the interface. |
| ip flow-export destination <i>ip-address udp-port</i> | IP address of the workstation to which you want to send the NetFlow information and the number of the UDP port on which the workstation is listening for this input. UDP port 9996 is commonly used for NetFlow. |
| ip flow-export version <i>version</i> | Specifies the version format that the export packet uses. |

The figure shows configurations for NetFlow data capture and data export to NetFlow collector with IP address 10.1.10.100, where you can analyze the exported data.

Traffic that is being either received or transmitted by the GigabitEthernet 0/0 interface is captured using the **ip flow** command. Captured NetFlow information is then sent to the collector with IP address 10.1.10.100 on UDP port 9996. The **ip flow-export version** command specifies that the export packet uses the version 9 format.

Note NetFlow exports data in UDP in one of five formats (1, 5, 7, 8, and 9). Version 9 is the most versatile data export format but is not backward compatible with version 8 or version 5.

NetFlow Configuration (Cont.)

```
R1# show ip interface GigabitEthernet0/1
<output omitted>
Input features: Ingress-NetFlow, MCI Check
Output features: Access List, Post-Ingress-NetFlow, Egress-NetFlow
```

- Displays if NetFlow is enabled on an interface.

```
R1# show ip flow export
Flow export v9 is enabled for main cache
Export source and destination details :
VRF ID : Default
Destination(1) 10.1.10.100 (9996)
Version 9 flow records
43 flows exported in 15 udp datagrams
```

- Displays the status and the statistics for NetFlow data export

© 2013 Cisco Systems, Inc.

You can use the **show ip interface** command to verify if NetFlow is enabled on an interface. In the example, NetFlow is enabled in the ingress and egress directions on the GigabitEthernet0/1 interface.

Use the **show ip flow export** command to verify the status and statistics for NetFlow accounting data export. In the example, the configured destination for NetFlow export is 10.1.10.100 using UDP port 9996. The version of the configured flow export is 9.

NetFlow Configuration (Cont.)

```
Branch# show ip cache flow
<output omitted>
IP Flow Switching Cache, 278544 bytes
 2 active, 4094 inactive, 31 added
6374 age polls, 0 flow alloc failures
Active flows timeout in 30 minutes
Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 34056 bytes
 2 active, 1022 inactive, 31 added to flow
 0 alloc failures, 0 force free
 1 chunk, 0 chunks added
last clearing of statistics 00:49:48
Protocol    Total    Flows    Packets    Bytes    Packets    Active(Sec)  Idle(Sec)
-----    -
            Flows   /Sec     /Flow     /Pkt     /Sec     /Flow       /Flow
TCP-Telnet  19      0.0      19        58       0.1      6.5        11.7
TCP-WWW    14      0.0       8        202      0.0      0.0        1.5
TCP-other   2       0.0      19        98       0.0      2.2        8.9
<output omitted>
SrcIf      SrcIPAddress  DstIf      DstIPAddress  Pr SrcP DstP  Pkts
Gi0/1     172.16.1.100  Gi0/0.10   10.1.10.100   01 0000 0000 1341
```

- Displays a summary of the NetFlow accounting statistics

© 2013 Cisco Systems, Inc.

You can use the **show ip cache flow** command to display a summary of NetFlow statistics on a Cisco IOS router. Using this command, you can see which protocols use the highest volume of traffic and between which hosts this traffic flows.

The example shows a summary of NetFlow statistics on the Branch router.

Do Not Duplicate:
Post beta, not for release.

Summary

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

Summary

- NMS polls the SNMP agent on a network device to obtain statistics.
- Use the **snmp-server community** command to configure SNMP access to the router.
- Syslog is a protocol that allows a network device to send event notifications to a syslog server.
- Use the **logging** command to identify a syslog server host to receive logging messages.
- NetFlow provides statistics on packets flowing through the routing devices in the network.
- The configuration part of NetFlow consists of configuring data capture and configuring data export.

© 2013 Cisco Systems, Inc.

Do Not Duplicate.
Post beta, not for release.

Managing Cisco Devices

When a Cisco router boots, it performs a series of steps in a particular order. At several points during the process, the router makes a decision about the next step to take. Knowledge of the boot sequence can be of great help when troubleshooting a Cisco router and also when adjusting its configuration. Carefully managing Cisco IOS images and configuration files reduces device downtime and maintains best practices. Cisco IOS image files contain the Cisco IOS Software that is required for a Cisco device to operate. The device configuration files contain a set of user-defined configuration commands that customize the functionality of a Cisco device.

This lesson describes the steps in the router boot sequence and the procedures and commands that are required to manage Cisco IOS images, configuration files, and devices on the network.

Objectives

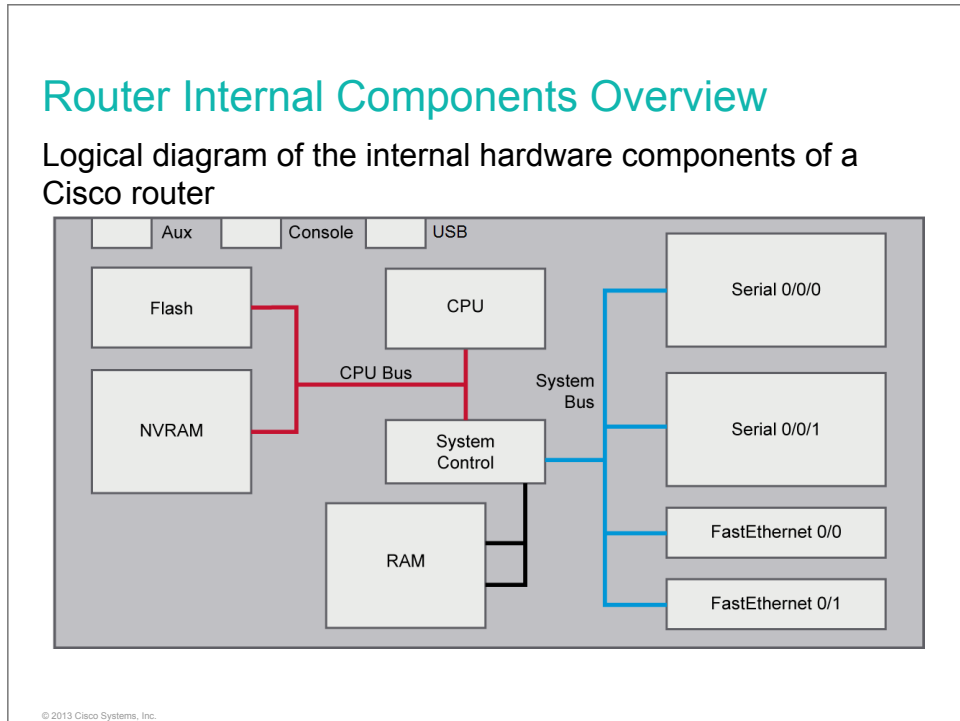
Upon completing this lesson, you will be able to meet these objectives:

- Describe the major internal components of a Cisco router
- Describe the functions of ROM in a Cisco router
- Describe the sequence of events that occur during a router bootup
- Describe how to display the boot information in the configuration register
- Describe how to change the boot information in the configuration register
- Describe the process of locating Cisco IOS images
- Describe the process of loading Cisco IOS images
- Describe the process of loading Cisco IOS configuration files
- Describe the file systems that are used by a Cisco router
- Describe why it is important to create a backup of Cisco IOS images and configuration files
- Describe how to decipher Cisco IOS image filenames
- Describe how to create a backup of a Cisco IOS image to a TFTP server
- Describe how to upgrade a Cisco IOS router from a TFTP server

- Describe the configuration files and their location
- Describe how to perform a password recovery on a Cisco router

Router Internal Components Overview

The major internal components of a Cisco router include the CPU, interfaces, RAM, ROM, flash memory, and NVRAM. This topic describes these major internal components.



A router is a computer, similar to a PC. Routers have many of the same hardware and software components that are found in other computers, including the following:

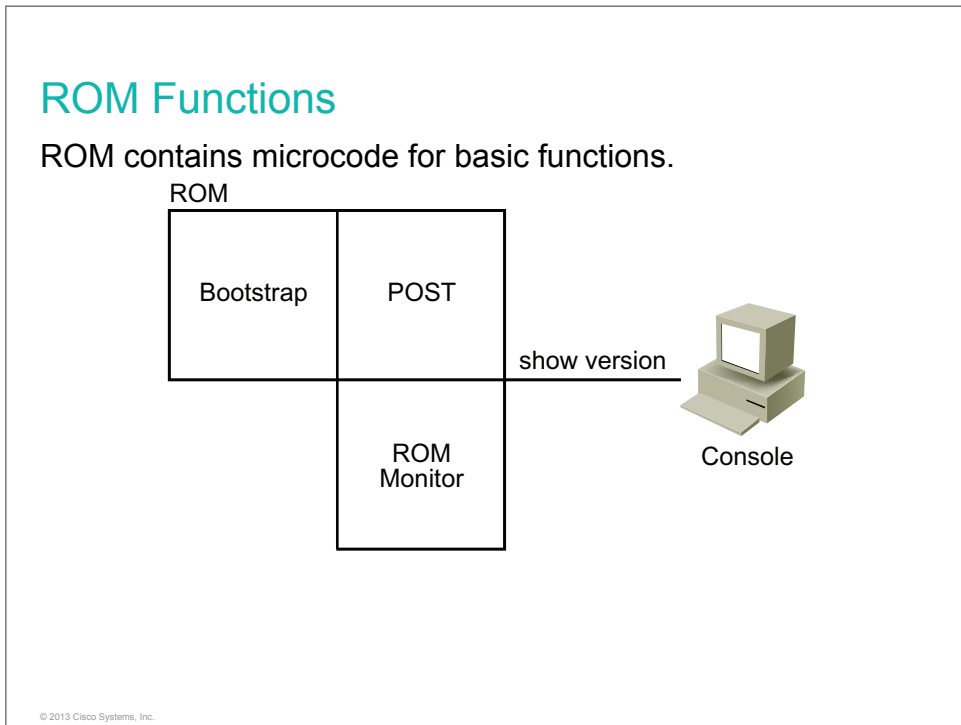
- **CPU:** The CPU executes operating system instructions such as system initialization, routing functions, and switching functions.
- **RAM:** RAM stores the instructions and data that the CPU needs to execute. This read/write memory contains the software and data structures that allow the router to function. RAM is volatile memory and loses its content when the router is powered down or restarted. However, the router also contains permanent storage areas such as ROM, flash, and NVRAM. RAM is used to store the following components:
 - **Operating system:** Cisco IOS Software is copied into RAM during bootup.
 - **Running configuration file:** This file stores the configuration commands that Cisco IOS Software is currently using on the router. With few exceptions, all commands that are configured on the router are stored in the running configuration file, which is also known as “running-config.”
 - **IP routing table:** This file stores information about directly connected and remote networks. It is used to determine the best path to forward the packet.
 - **ARP cache:** ARP cache contains the IPv4 address to MAC address mappings, such as the ARP cache on a PC. The ARP cache is used on routers that have LAN interfaces such as Ethernet interfaces.

- **Packet buffer:** Packets are temporarily stored in a buffer when they are received on an interface or before they exit an interface.
- **ROM:** ROM is a form of permanent storage. This type of memory contains microcode for basic functions to start and maintain the router. ROM contains the ROM monitor, which is used for router disaster recovery functions such as password recovery. ROM is nonvolatile; it maintains the memory contents even when the power is turned off.
- **Flash memory:** Flash memory is nonvolatile computer memory that can be electrically stored and erased. Flash is used as permanent storage for the operating system. In most models of Cisco routers, Cisco IOS Software is permanently stored in flash memory and copied into RAM during the bootup process, where the CPU then executes it. Some older models of Cisco routers run Cisco IOS Software directly from flash. Flash consists of SIMMs or PCMCIA cards that can be upgraded to increase the amount of flash memory. Flash memory does not lose its contents when the router loses power or is restarted.
- **NVRAM:** NVRAM does not lose its information when the power is turned off. Cisco IOS Software uses NVRAM as permanent storage for the startup configuration file (startup-config). All configuration changes are stored in the running configuration file in RAM, and with few exceptions, Cisco IOS Software implements them immediately. To save these changes in case the router is restarted or loses power, the running configuration must be copied to NVRAM, where it is stored as the startup configuration file.
- **Configuration register:** The configuration register is used to control how the router boots. The configuration register value is stored in NVRAM.
- **Interfaces:** Interfaces are the physical connections to the external world for the router and include the following types, among others:
 - Ethernet, Fast Ethernet, and Gigabit Ethernet
 - Asynchronous and synchronous serial
 - USB interface, which can be used to add a USB flash drive to a router
 - Console and auxiliary ports. A console can have an RJ-45 or mini-USB connector.

Although there are several different types and models of routers, every router has the same general hardware components. Depending on the model, these components are located in different places inside the router.

ROM Functions

This topic describes the functions of ROM in a Cisco router.



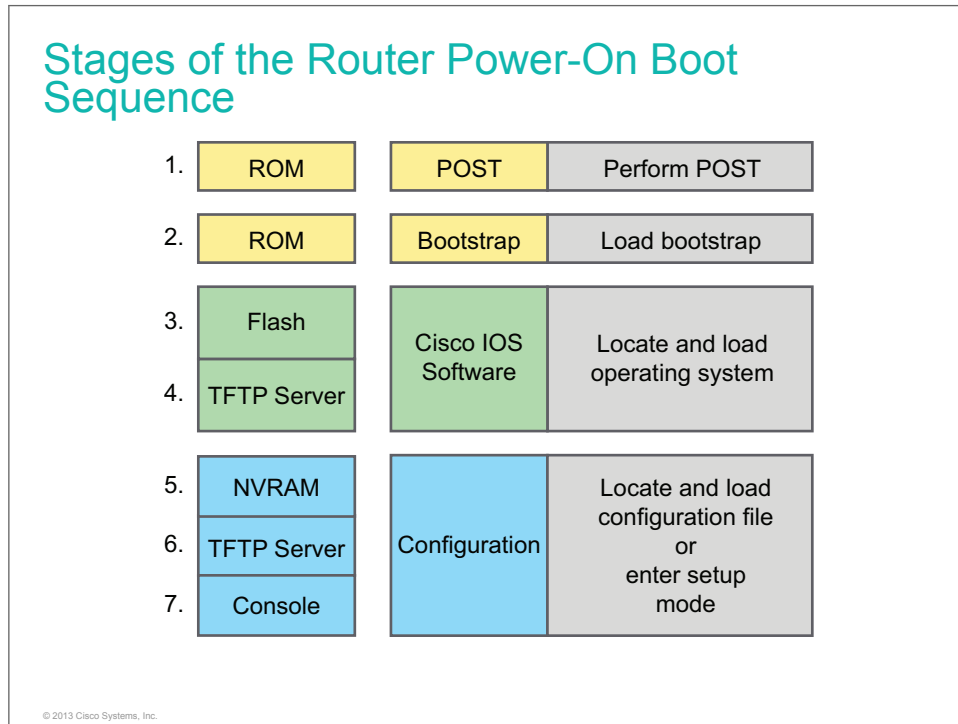
There are three major areas of microcode that are generally contained in ROM:

- **Bootstrap code:** The bootstrap code is used to bring up the router during initialization. It reads the configuration register to determine how to boot and then, if instructed to do so, loads the Cisco IOS Software.
- **POST:** POST is the microcode that is used to test the basic functionality of the router hardware and determine which components are present.
- **ROM monitor:** This area includes a low-level operating system that is normally used for manufacturing, testing, troubleshooting, and password recovery. In ROM monitor mode, the router has no routing or IP capabilities.

Note Depending on the specific Cisco router platform, the components that are listed may be stored in flash memory or in bootstrap memory to allow for field upgrade to later versions.

Stages of the Router Power-On Boot Sequence

When a router boots, it performs a series of steps: performing tests, finding and loading the Cisco IOS Software, finding and loading configurations, and running the Cisco IOS Software. This topic describes the sequence of events that occurs during a router bootstrap.



The sequence of events that occurs during the power-up (boot) of a router is important. Knowledge of this sequence helps to accomplish operational tasks and troubleshoot router problems.

When power is initially applied to a router, the events occur in the following order:

- 1 Perform POST:** This event is a series of hardware tests that verifies if all components of the Cisco router are functional. During this test, the router also determines which hardware is present. POST executes from microcode that is resident in the system ROM.
- 2 Load and run bootstrap code:** Bootstrap code is used to perform subsequent events such as locating the Cisco IOS Software, loading it into RAM, and then running it. When the Cisco IOS Software is loaded and running, the bootstrap code is not used until the next time that the router is reloaded or power-cycled.
- 3 Find the Cisco IOS Software:** The bootstrap code determines the location of the Cisco IOS Software that is to be run. Normally, the Cisco IOS Software image is located in the flash memory, but it can also be stored in other places such as a TFTP server. The configuration register and configuration file determine where the Cisco IOS Software images are located and which image file to use. If a complete Cisco IOS image cannot be located, a scaled-down version of the Cisco IOS Software is copied from ROM into RAM. This version of the Cisco IOS Software is used to help diagnose any problems and can be used to load a complete version of the Cisco IOS Software into RAM.
- 4 Load the Cisco IOS Software:** After the bootstrap code has found the correct image, it then loads that image into RAM and starts the Cisco IOS Software. Some older routers do not load the Cisco IOS Software image into RAM but execute it directly from flash memory instead.
- 5 Find the configuration:** After the Cisco IOS Software is loaded, the bootstrap program searches for the startup configuration file (startup-config) in NVRAM.

- 6 **Load the configuration:** If a startup configuration file is found in NVRAM, the Cisco IOS Software loads it into RAM as the running configuration and executes the commands in the file, one line at a time. The running-config file contains interface addresses, starts routing processes, configures router passwords, and defines other characteristics of the router. If no configuration exists, the router will enter the setup utility or attempt an autoinstall to look for a configuration file from a TFTP server.
- 7 **Run the configured Cisco IOS Software:** When the prompt is displayed, the router is running the Cisco IOS Software with the current running configuration file. The network administrator can now begin using Cisco IOS commands on this router.

Do Not Duplicate.
Post beta, not for release.

Configuration Register

This topic describes how to display the boot information in the configuration register.

Configuration Register

- The configuration register is a 16-bit number that affects router behavior.
- The least-significant 4 bits of the configuration register are called the boot field.
- The boot field in the configuration register specifies how the router locates Cisco IOS Software.

© 2013 Cisco Systems, Inc.

A router has a 16-bit configuration register in NVRAM. Each bit has value 1 (on or set) or value 0 (off or clear), and each bit setting affects the router behavior upon the next reload power cycle.

You can use the 16-bit configuration register to do the following:

- Force the router to boot into the ROM monitor.
- Select a boot source and default boot filename.
- Control broadcast addresses.
- Recover a lost password.
- Change the console line speed.

The lowest 4 bits (the rightmost hexadecimal number) are called the boot field and specify how a router locates the Cisco IOS image file.

Changing the Configuration Register

This topic describes how to change the boot information in the configuration register.

Changing Configuration Register

```
Branch# show version
<output omitted>
Configuration register is 0x2102
```

First, verify the current configuration register value.

```
Branch# configure terminal
Branch(config)# config-register 0x2101
Branch(config)# exit
Branch# copy running-config startup-config
```

Set the configuration register value.

```
Branch# show version
<output omitted>
Configuration register is 0x2102 (will be 0x2101 at next reload)
```

Verify the new configuration register value.

© 2013 Cisco Systems, Inc.

Before altering the configuration register, you should determine how the router is currently loading the software image. The **show version** command will display the current configuration register value. The last line of the display contains the configuration register value.

You can change the default configuration register setting with the **config-register** global configuration command. The configuration register is a 16-bit register. The lowest 4 bits of the configuration register (bits 3, 2, 1, and 0) form the boot field. A hexadecimal number is used as the argument to set the value of the configuration register. The default value of the configuration register is 0x2102.

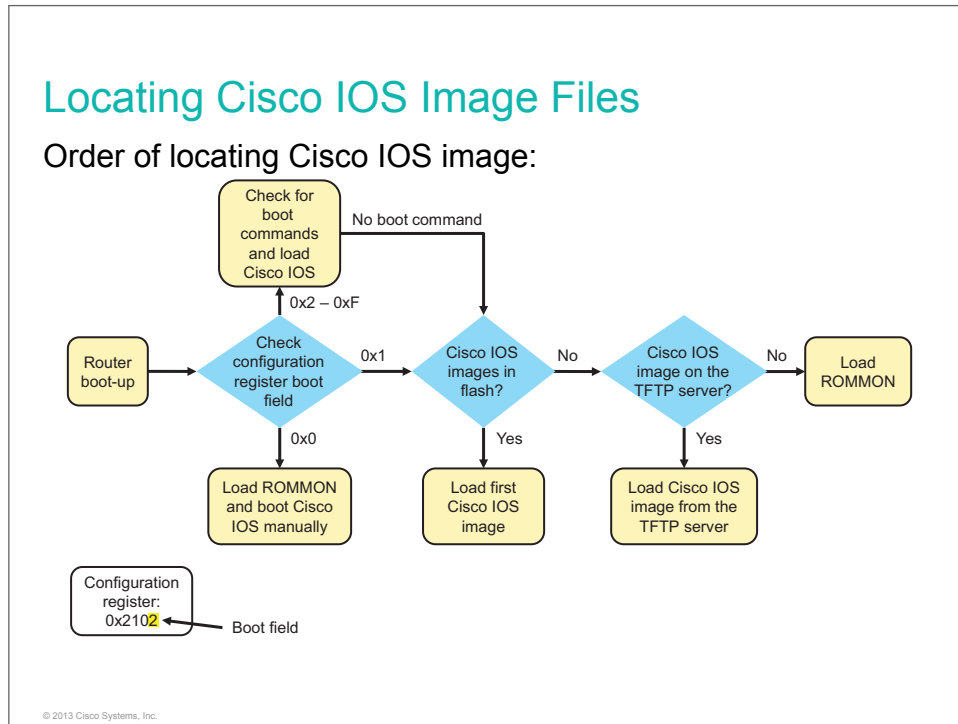
The **show version** command is used to verify changes in the configuration register setting. The new configuration register value takes effect when the router reloads.

In this example, the **show version** command indicates that the configuration register setting of 0x2101 will be used during the next router reload. The new configuration register value will cause the router to boot the first Cisco IOS image that is found in Flash.

When using the **config-register** command, all 16 bits of the configuration register are set. Be careful to modify only the bits that you are trying to change, such as the boot field, and leave the other bits as they are. Remember that the other configuration register bits perform functions that include the selection of the console b/s rate and whether to use the saved configuration in NVRAM.

Locating Cisco IOS Image Files

When a Cisco router boots, it searches for the Cisco IOS image in a specific sequence. It searches for the location that is specified in the configuration register, flash memory, a TFTP server, and ROM. This topic describes the process of locating the Cisco IOS image.



The bootstrap code is responsible for locating the Cisco IOS Software. It searches for the Cisco IOS image according to the following sequence:

1. The bootstrap code checks the boot field of the configuration register. The configuration register has several uses, such as for telling the router how to boot up or for password recovery. For example, the factory default setting for the configuration register is 0x2102. This value indicates that the router attempts to load a Cisco IOS Software image from flash memory and loads the startup configuration file from NVRAM. The boot field is the lowest 4 bits of the configuration register and is used to specify how the router boots. These bits can point to flash memory for the Cisco IOS image, the startup configuration file (if one exists) for commands that tell the router how to boot, or a remote TFTP server. Alternatively, these bits can specify that no Cisco IOS image is to be loaded and to start a Cisco ROM monitor. The configuration register bits perform other functions as well, such as selection of the console b/s rate and whether to use the saved configuration file (startup-config) in NVRAM. It is possible to change the configuration register and, therefore, change where the router looks for the Cisco IOS image and the startup configuration file during the bootup process.

For example, a configuration register value of 0x2102 (the "0x" indicates that the digits that follow are in hexadecimal notation) has a boot field value of 0x2. The right-most digit in the register value is 2 and represents the lowest 4 bits of the register. The table indicates how different values of the boot field affect the location of the Cisco IOS image.

| Configuration Register Boot Field Value | Meaning |
|---|--|
| 0x0 | At the next power cycle or reload, the router boots to the ROM monitor. |
| 0x1 | The router boots the first image in flash memory as a system image. |
| 0x2 to 0xF | At the next power cycle or reload, the router sequentially processes each boot system command in global configuration mode. |

- If the boot field value of the configuration register is from 0x2 to 0xF, the bootstrap code parses the startup configuration file in NVRAM for the **boot system** commands that specify the name and location of the Cisco IOS Software image to load. Several **boot system** commands can be entered in sequence to provide a fault-tolerant boot plan.

The **boot system** command is a global configuration command that allows you to specify the source for the Cisco IOS Software image to load. Some of the syntax options that are available include the following:

- This example boots the system boot image file named *igs-bpx-1* from the flash memory device:

```
Branch(config)# boot system flash:c2900-universalk9-mz.SPA.152-4.M1.bin
```

- This example illustrates a configuration that specifies a TFTP server as a source of a Cisco IOS image, with a ROM monitor as the backup:

```
Branch(config)#boot system tftp://c2900-universalk9-mz.SPA.152-4.M1.bin
Branch(config)#boot system rom
```

- If there are no **boot system** commands in the configuration, the router defaults to loading the first valid Cisco IOS image in flash memory and running it.
- If no valid Cisco IOS image is found in flash memory, the router attempts to boot from a network TFTP server using the boot field value as part of the Cisco IOS image filename.
- After six unsuccessful attempts at locating a TFTP server, the router will load the ROM monitor.

Note The procedure of locating the Cisco IOS image depends on the Cisco router platform and default configuration register values. The procedure that was just described applies to the Cisco 3900, 2900, and 1900 Series Integrated Services Routers.

Loading Cisco IOS Image Files

This topic describes the process of loading the Cisco IOS image.

Loading Cisco IOS Image Files

```
System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 2011 by cisco Systems, Inc.
Total memory size = 512 MB - On-board = 512 MB, DIMM0 = 0 MB
CISCO2901/K9 platform with 524288 Kbytes of main memory
Main memory is configured to 72/-1(On-board/DIMM0) bit mode with ECC enabled
Readonly ROMMON initialized
program load complete, entry point: 0x80803000, size: 0x1b340
program load complete, entry point: 0x80803000, size: 0x1b340
IOS Image Load Test

Digitally Signed Release Software
program load complete, entry point: 0x81000000, size: 0x5d433c0
Self decompressing the image:
#####
##### [OK]
<output omitted>
```

- The Cisco IOS image file is decompressed and stored to RAM. The output shows the boot process on a router.

© 2013 Cisco Systems, Inc.

When the router locates a valid Cisco IOS image file in flash memory, the Cisco IOS image is normally loaded into RAM to run. If the image needs to be loaded from flash memory into RAM, it must first be decompressed. After the file is decompressed into RAM, it is started. When the Cisco IOS Software begins to load, you may see a string of pound signs (#), as shown in the figure, while the image decompresses.

Loading Cisco IOS Image Files (Cont.)

```
Branch# show version
Cisco IOS Software, C2900 Software (C2900-UNIVERSALK9-M), Version 15.2(4)M1,
RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2012 by Cisco Systems, Inc.
Compiled Thu 26-Jul-12 20:54 by prod_rel_team
ROM: System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1)
Branch uptime is 39 minutes
System returned to ROM by reload at 11:39:24 UTC Tue Nov 20 2012
System image file is "flash0:c2900-universalk9-mz.SPA.152-4.M1.bin"
Last reload type: Normal Reload
Last reload reason: Reload Command
<output omitted>
```

(Continued in next figure)

© 2013 Cisco Systems, Inc.

Loading Cisco IOS Image Files (Cont.)

```
Cisco CISCO2901/K9 (revision 1.0) with 483328K/40960K bytes of memory.
Processor board ID FCZ1642C5XJ
 2 Gigabit Ethernet interfaces
 1 Serial(sync/async) interface
1 terminal line
DRAM configuration is 64 bits wide with parity enabled.
255K bytes of non-volatile configuration memory.
250880K bytes of ATA System CompactFlash 0 (Read/Write)
<output omitted>
Configuration register is 0x2102
```

- Displays information about the currently loaded software, as well as hardware and device information

The **show version** command can be used to help verify and troubleshoot some of the basic hardware and software components of the router. The **show version** command displays information about the version of the Cisco IOS Software that is currently running on the router, the version of the bootstrap program, and information about the hardware configuration, including the amount of system memory.

The output from the **show version** command includes the following:

- **Cisco IOS version**

```
Cisco IOS Software, C2900 Software (C2900-UNIVERSALK9-M), Version 15.2(4)M1,
RELEASE SOFTWARE (fc1)
```

This line from the example output shows the version of the Cisco IOS Software in RAM that the router is using.

- **ROM bootstrap program**

```
ROM: System Bootstrap, Version 15.0(1r)M15, RELEASE SOFTWARE (fc1)
```

This line from the example output shows the version of the system bootstrap software that is stored in ROM and was initially used to boot up the router.

- **Location of Cisco IOS image**

```
System image file is "flash0:c2900-universalk9-mz.SPA.152-4.M1.bin"
```

This line from the example output shows where the Cisco IOS image is located and loaded as well as its complete filename.

- **Interfaces**

```
2 Gigabit Ethernet interfaces
```

```
1 Serial (sync/async) interface
```

This section of the output displays the physical interfaces on the router. In this example, the Cisco 2901 router has two GigabitEthernet interfaces and one serial interface.

- **Amount of NVRAM**

```
255 KB of NVRAM
```

This line from the example output shows the amount of NVRAM on the router.

- **Amount of Flash**

```
250,880 KB of ATA System CompactFlash 0 (Read/Write)
```

This line from the example output shows the amount of flash memory on the router.

- **Configuration register**

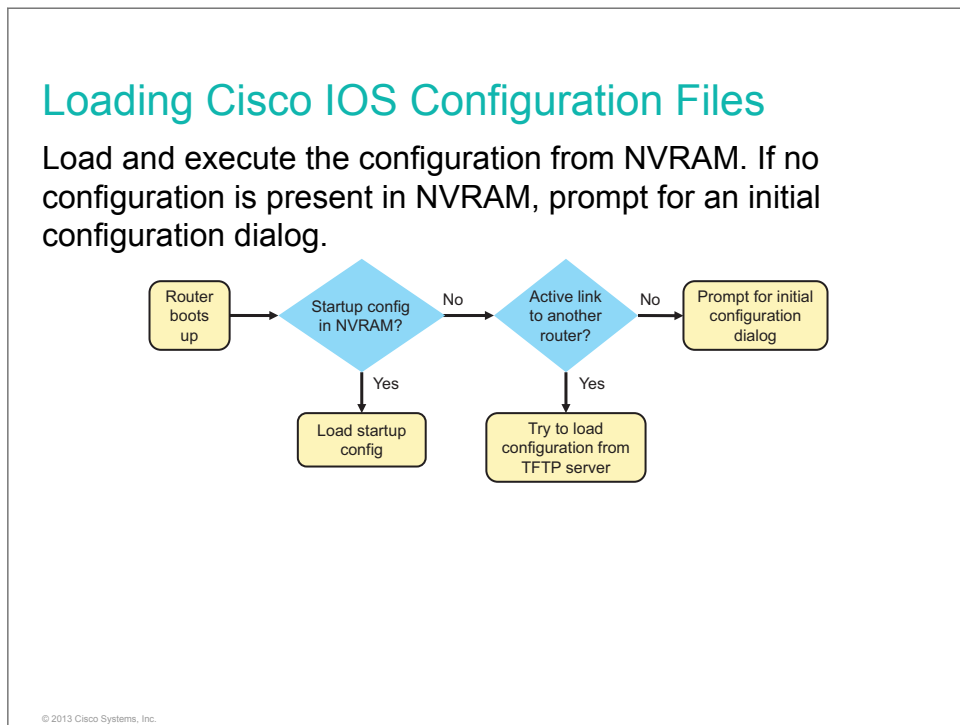
```
Configuration register is 0x2102
```

The last line of the **show version** command displays the current configured value of the software configuration register in hexadecimal format. This value indicates that the router will attempt to load a Cisco IOS Software image from flash memory and load the startup configuration file from NVRAM.

Do Not Duplicate
Post beta, not for release.

Loading Cisco IOS Configuration Files

This topic describes the process of loading the Cisco IOS configuration files.



After the Cisco IOS Software image is loaded and started, the router must be configured to be useful. If there is an existing saved configuration file (startup-config) in NVRAM, it is executed. If there is no saved configuration file in NVRAM, the router either begins autoinstall or enters the setup utility.

If the startup configuration file does not exist in NVRAM, the router may search for a TFTP server. If the router detects that it has an active link to another configured router, it sends a broadcast searching for a configuration file across the active link. This condition will cause the router to pause, but you will eventually see a console message such as the following:

```
%Error opening tftp://255.255.255.255/network-config(Timed out)
%Error opening tftp://255.255.255.255/cisconet.cfg (Timed out)
```

The setup utility prompts the user at the console for specific configuration information to create a basic initial configuration on the router, as shown in this example:

```
<output omitted>
If you require further assistance please contact us by sending email to
export@cisco.com.
Cisco CISCO2901/K9 (revision 1.0) with 483328K/40960K bytes of memory.
Processor board ID FCZ1642C5XJ
2 Gigabit Ethernet interfaces
1 Serial(sync/async) interface
1 terminal line
DRAM configuration is 64 bits wide with parity enabled.
255K bytes of non-volatile configuration memory.
250880K bytes of ATA System CompactFlash 0 (Read/Write)
--- System Configuration Dialog ---
Would you like to enter the initial configuration dialog? [yes/no]:
```


Loading Cisco IOS Configuration Files (Cont.)

```
Branch# show running-config
Building configuration...
Current configuration : 1318 bytes
!
! Last configuration change at 13:11:38 UTC Tue Nov 20 2012
! NVRAM config last updated at 13:11:38 UTC Tue Nov 20 2012
! NVRAM config last updated at 13:11:38 UTC Tue Nov 20 2012
version 15.2
<output omitted>
```

- Displays the current configuration

```
Branch# show startup-config
Using 1318 out of 262136 bytes
!
! Last configuration change at 13:11:38 UTC Tue Nov 20 2012
! NVRAM config last updated at 13:11:38 UTC Tue Nov 20 2012
! NVRAM config last updated at 13:11:38 UTC Tue Nov 20 2012
version 15.2
<output omitted>
```

- Displays the saved configuration

© 2013 Cisco Systems, Inc.

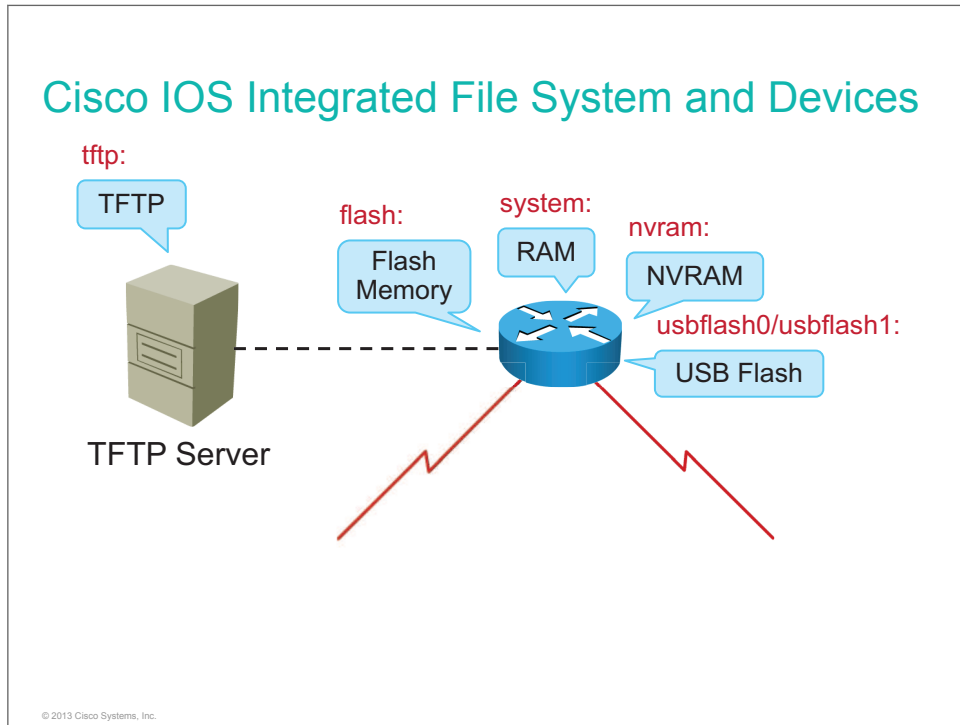
The **show running-config** and **show startup-config** commands are among the most common Cisco IOS Software EXEC commands because they allow you to see the current running configuration in RAM on the router or the startup configuration commands in the startup configuration file in NVRAM that the router will use at the next restart.

If the words "Current configuration" are displayed, the active running configuration from RAM is being displayed.

If there is a message at the top indicating how much nonvolatile memory is being used ("Using 1318 out of 262,136 B" in this example), the startup configuration file from NVRAM is being displayed.

Cisco IOS Integrated File System and Devices

The Cisco IFS feature provides a single interface to all of the file systems that a router uses. This topic describes the file systems that are used by a Cisco router.



The availability of the network can be at risk if the configuration of a router or the operating system is compromised. Attackers who gain access to infrastructure devices can alter or delete configuration files. They can also upload incompatible Cisco IOS images or delete the Cisco IOS image. The changes are invoked automatically or invoked when the device is rebooted.

To mitigate against these problems, you have to be able to save, back up, and restore configurations and Cisco IOS images.

Cisco IOS devices provide a feature that is called the Cisco IFS. This system allows you to create, navigate, and manipulate directories on a Cisco device. The directories that are available depend on the platform. The Cisco IFS feature provides a single interface to all of the file systems that a Cisco router uses, including the following:

- Flash memory file systems
- NFSs: TFTP, RCP, and FTP
- Any other endpoint for reading or writing data (such as NVRAM, the running configuration in RAM, and so on)

An important feature of the Cisco IFS is the use of the URL convention to specify files on network devices and the network. The URL prefix specifies the file system.

Cisco IOS Integrated File System and Devices (Cont.)

```
Branch# show file systems
File Systems:
  Size (b)      Free (b)      Type  Flags  Prefixes
  -           -           -     -      -
  -           -           opaque rw    archive:
  -           -           opaque rw    system:
  -           -           opaque rw    tmpsys:
  -           -           opaque rw    null:
  -           -           network rw    tftp:
* 256610304    153710592    disk  rw    flash0: flash:#
  -           -           disk  rw    flash1:
  262136      255626      nvram rw    nvram:
  -           -           opaque wo   syslog:
  -           -           opaque rw    xmodem:
  -           -           opaque rw    ymodem:
  -           -           network rw    rcp:
  -           -           network rw    http:
  -           -           network rw    ftp:
  -           -           network rw    scp:
  -           -           opaque ro   tar:
  -           -           network rw    https:
  -           -           opaque ro   cns:
```

- Lists all of the available file systems

© 2013 Cisco Systems, Inc.

This figure displays the output of the **show file systems** command, which lists all of the available file systems on a Cisco 2901 router. This command provides insightful information such as the amount of available and free memory and the type of file system and its permissions. Permissions include read only (as indicated by the “ro” flag), write only (as indicated by the “wo” flag), and read and write (as indicated by the “rw” flag).

The FFS has an asterisk preceding it, which indicates the current default file system. The bootable Cisco IOS Software is located in flash memory, so the pound symbol (#) that is appended to the flash listing indicates a bootable disk.

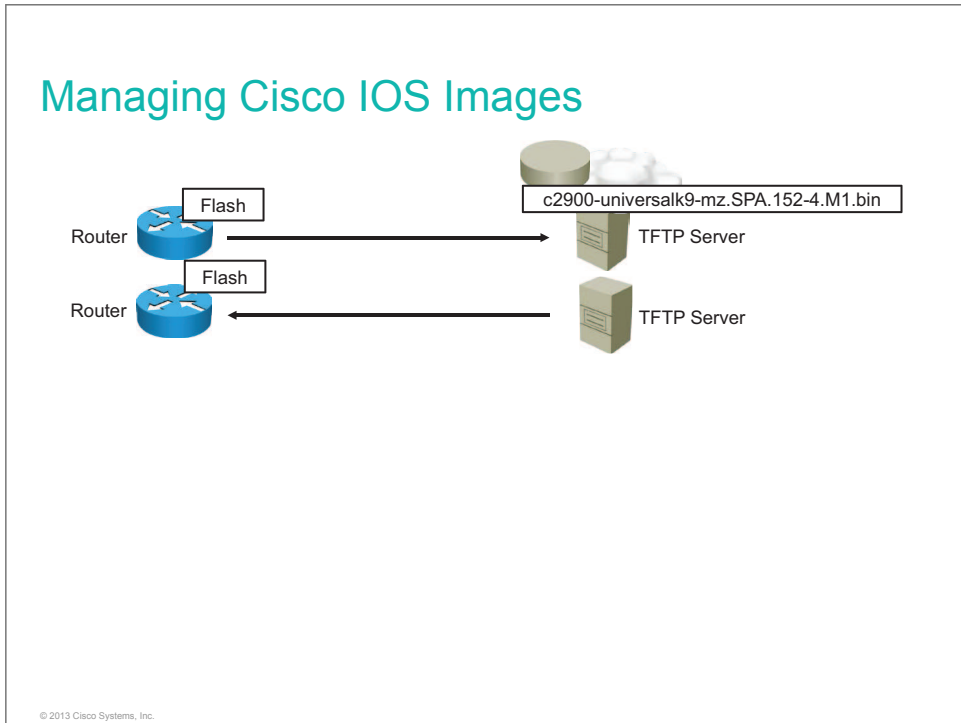
This table contains some commonly used URL prefixes for Cisco network devices.

| Prefix | Description |
|----------------------|--|
| flash: | Flash memory. This prefix is available on all platforms. For platforms that do not have a device named Flash, the flash: prefix is aliased to slot0. Therefore, the flash: prefix can be used to refer to the main flash memory storage area on all platforms. |
| ftp: | FTP network server |
| http: | HTTP network server |
| nvram: | NVRAM |
| rcp: | RCP network server |
| system: | Contains the system memory, including the current running configuration |
| tftp: | TFTP network server |
| usbflash0, usbflash1 | USB flash |

An important feature of the Cisco IFS is the use of the URL convention to specify files on network devices and the network. The URL prefix specifies the file system.

Managing Cisco IOS Images

As a network grows, storage of Cisco IOS Software images and configuration files on a central TFTP server enables control of the number and revision level of Cisco IOS images and configuration files that must be maintained. This topic describes why it is important to create a backup of Cisco IOS images and configuration files.

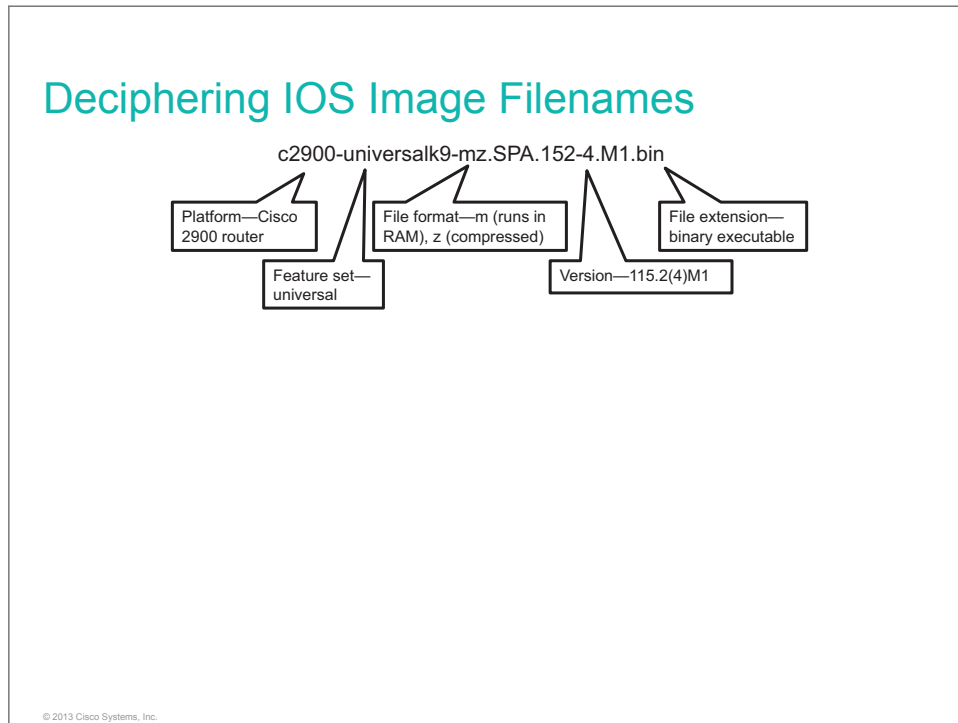


Production internetworks usually span wide areas and contain multiple routers. For any network, it is prudent to retain a backup copy of the Cisco IOS Software image in case the system image in the router becomes corrupted or accidentally erased.

Widely distributed routers need a source or backup location for Cisco IOS Software images. Using a network TFTP server allows image and configuration uploads and downloads over the network. The network TFTP server can be another router, a workstation, or a host system.

Deciphering Cisco IOS Image Filenames

This topic describes how to decipher Cisco IOS image filenames.



Before upgrading a Cisco IOS router, it is necessary to select a Cisco IOS image with the correct feature set and version. The Cisco IOS image file is based on a special naming convention. The name for the Cisco IOS image file contains multiple parts, each with a specific meaning. It is important that you understand this naming convention when upgrading and selecting Cisco IOS Software.

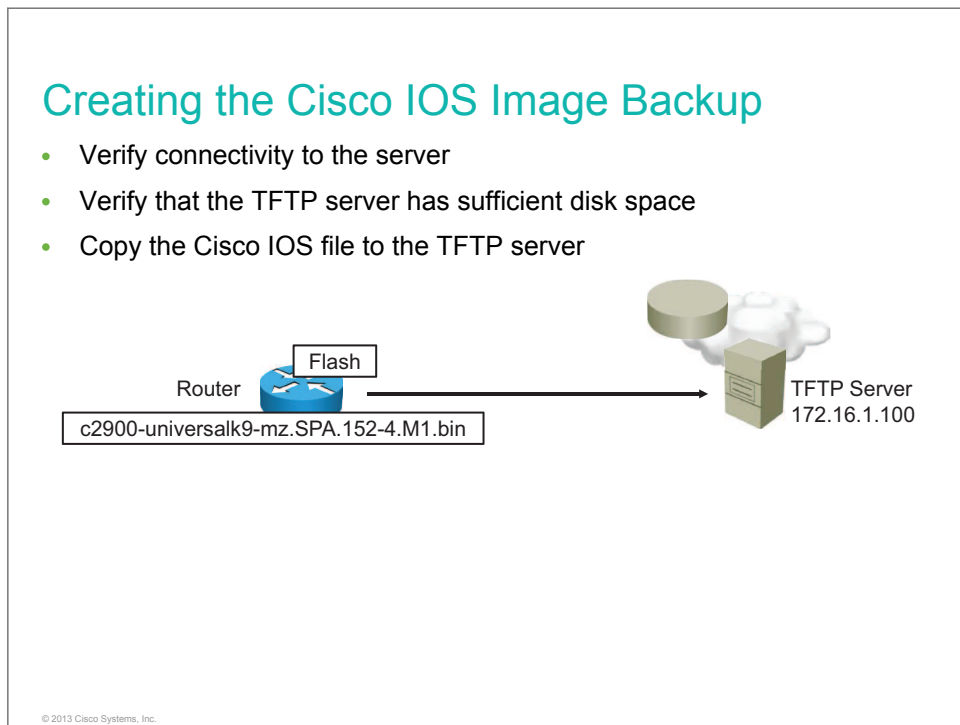
For example, the filename in this figure is explained as follows:

- The first part (`c2900`) identifies the platform on which the image runs. In this example, the platform is a Cisco 2900 Series Integrated Services Router.
- The second part (`universal`) specifies the feature set. In this case, "universal" refers to the universal, single image set that includes the IP base, security, unified communications, and data feature sets. Each router is activated for an IP Base feature set. However, for other feature sets, software activation is needed.
- The third part (`mz`) indicates where the image runs and if the file is compressed. In this example, "mz" indicates that the file runs from RAM and is compressed.
- The fourth part (`15.2(4)M1`) is the version number.
- The final part (`.bin`) is the file extension. This extension indicates that this file is a binary executable file.

Note The Cisco IOS Software naming conventions, field meaning, image content, and other details are subject to change.

Creating the Cisco IOS Image Backup

This topic describes how to create a backup of a Cisco IOS image to a TFTP server.



To maintain network operations with minimum down time, it is necessary to have procedures in place for backing up Cisco IOS images. In this way, you can quickly copy an image back to a router in case of a corrupted or erased image on the router.

Follow these steps to create a backup of Cisco IOS images to the TFTP server:

- Make sure that there is access to the network TFTP server. You can ping the TFTP server to test connectivity.
- Verify that the TFTP server has sufficient disk space to accommodate the Cisco IOS Software image. Use the **show flash0:** command on the router to determine the size of the Cisco IOS image file.
- Copy the image to the TFTP server using the **copy** command.

Creating the Cisco IOS Image Backup (Cont.)

```
Branch# ping 172.16.1.100
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.1.100, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

Verify connectivity to the server.

```
Branch# show flash0:
-#- --length-- -----date/time----- path
1      97794040 Nov 30 1983 00:00:00 +00:00 c2900-universalk9-mz.SPA.
152-4.M1.bin
<output omitted>
```

Verify Cisco IOS image size.

© 2013 Cisco Systems, Inc.

Creating the Cisco IOS Image Backup (Cont.)

```
Branch# copy flash0: tftp:
Source filename []? c2900-universalk9-mz.SPA.152-4.M1.bin
Address or name of remote host []? 172.16.1.100
Destination filename []? c2900-universalk9-mz.SPA.152-4.M1.bin
!!!!!!!!!!!!!!!!!!!!!!
<output omitted>
97794040 bytes copied in 363.468 secs (269058 bytes/sec)
```

Copy image to the TFTP server

© 2013 Cisco Systems, Inc.

Before creating an image backup, verify connectivity to the TFTP server. You can do this by pinging the TFTP server from the router. In the example, the TFTP server is accessible from the router.

You should then make sure that you have sufficient disk space on the TFTP server to accommodate the Cisco IOS Software image. You can use the **show flash** command to verify the Cisco IOS Software image file size. The file in the example is 97,794,040 B (93 MB).

In this example, you will create a backup of the current image file on the router (c2900-universalk9-mz.SPA.152-4.M1.bin) to the TFTP server at 172.16.1.100.

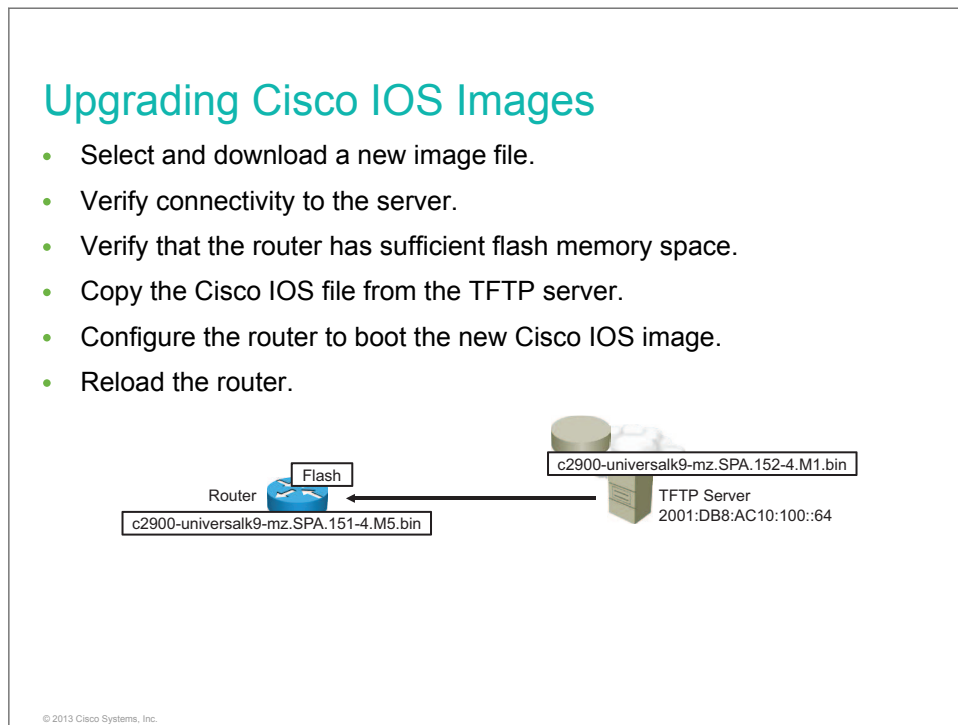
Finally, copy the Cisco IOS Software image file to the server, using the **copy** command. After you issue the command with specified source and destination URLs, you will be prompted for the source file name, IP address of the remote host, and destination filename. After you enter all of this information, transfer of the file will occur. The table describes the command.

| Command | Description |
|---|---|
| copy <i>source-url destination-url</i> | Copies any file from a source to a destination. The exact format of the source and destination URLs varies according to the file or directory location. |

Do Not Duplicate.
Post beta, not for release.

Upgrading Cisco IOS Images

This topic describes how to upgrade a Cisco IOS router from a TFTP server.



Cisco constantly releases new Cisco IOS Software versions to resolve caveats (software defects) and provide new features. If you decide to upgrade the software on the Cisco router, follow these steps:

- Select a Cisco IOS image file that meets your requirements in terms of platform, features, and software defects. Download the file from www.cisco.com and transfer it to the TFTP server.
- Make sure that there is access to the network TFTP server. Ping the TFTP server to test connectivity.
- Make sure that there is sufficient flash memory space on the router that is being upgraded. You can verify the amount of free flash space by using the **show flash0:** command. Compare the free flash space with the new image file size.
- Copy the Cisco IOS image file from the TFTP server to the router using the **copy** command.
- When the image is saved on the router flash memory, you have to instruct the router to load the new image during the bootup. Save the configuration.
- Finally, reload the router in order to boot the new image.

Upgrading Cisco IOS Images (Cont.)

```
Branch# ping 2001:DB8:AC10:100::64
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:AC10:100::64, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

Verify connectivity to the server

```
Branch# show flash0:
-#- --length-- -----date/time----- path
<output omitted>
6      3000320 Nov 20 2012 10:03:30 +00:00 cpexpress.tar
7          1038 Nov 20 2012 10:03:36 +00:00 home.shtml
153710592 bytes available (102899712 bytes used)
```

Verify free flash memory space

© 2013 Cisco Systems, Inc.

After you download the correct Cisco IOS image file and transfer it to the TFTP server, you should verify connectivity to the TFTP server by pinging the TFTP server from the router. In the example, the TFTP server is accessible from the router.

You should then make sure that you have sufficient disk space in flash memory to accommodate the Cisco IOS image. You can use the **show flash** command to verify free flash memory space. Free flash space in the example is 153,710,592 B.

In this example, you will load the new image file (c2900-universalk9-mz.SPA.152-4.M1.bin) from the TFTP server at 2001:DB8:AC10:100::64 to the router. The example uses IPv6 as the transport protocol to show that TFTP can also be used across IPv6 networks.

Upgrading Cisco IOS Image (Cont.)

```
Branch# copy tftp: flash0:
Address or name of remote host []? 2001:DB8:AC10:100::64
Source filename []? c2900-universalk9-mz.SPA.152-4.M1.bin
Destination filename []? c2900-universalk9-mz.SPA.152-4.M1.bin
Accessing tftp://2001:DB8:AC10:100::64/c2900-universalk9-mz.SPA.
152-4.M1.bin...
Loading c2900-universalk9-mz.SPA.152-4.M1.bin from 2001:DB8:AC10:100::64 (via
GigabitEthernet0/0): !!!!!!!!!!!!!!!!!!!!!!!
<output omitted>
[OK - 97794040 bytes]
97794040 bytes copied in 368.128 secs (265652 bytes/sec)
```

Copy the image from the TFTP server.

```
Branch# configure terminal
Branch(config)# boot system flash0://c2900-universalk9-mz.SPA.152-4.M1.bin
Branch# copy running-config startup-config
Branch# reload
```

Set the image to boot and reload the router

© 2013 Cisco Systems, Inc.

Copy the Cisco IOS image file from the server to the router flash memory using the **copy** command. After you issue the command with specified source and destination URLs, you will be prompted for the IP address of the remote host, source file name, and destination file name. After you enter all of the required information, transfer of the file will begin.

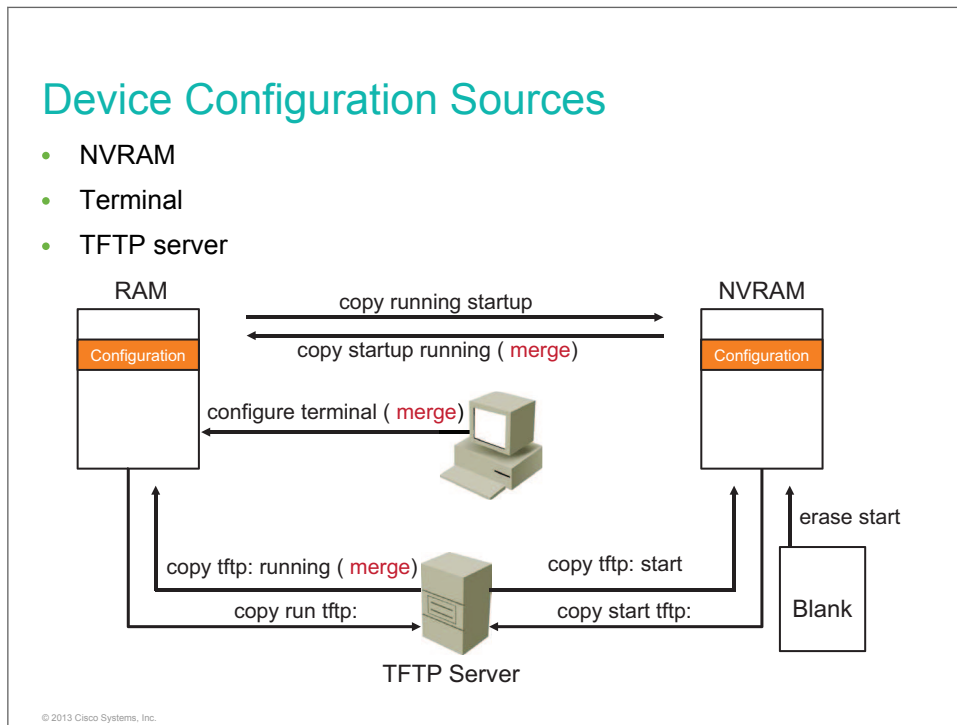
When the image file has been copied to the router, you have to instruct the router to boot the new image file. Use the **boot system** command to instruct the router to boot the specific file. Recall that the boot field in the configuration register has to be set to 0x2 to 0xF in order for the router to check the **boot** commands. Save the configuration.

Reload the router to boot the router with the new image. When the router has booted, you can verify if the new image has loaded using the **show version** command.

| Command | Description |
|------------------------|---|
| boot system url | Specify the system image that the router loads at startup |

Managing Device Configuration Files

Device configuration files contain a set of user-defined configuration commands that customize the functionality of a Cisco device. This topic describes the configuration files and their location.



Configuration files of a Cisco router are stored in the following locations:

- The running configuration is stored in RAM.
- The startup configuration is stored in NVRAM.

You can copy configuration files from the router to a file server using FTP or TFTP. For example, you can copy configuration files to back up a current configuration file to a server before changing its contents, therefore allowing the original configuration file to be restored from the server. The protocol that is used depends on which type of server is used.

You can copy configuration files from a server to the running configuration in RAM or to the startup configuration file in NVRAM of the router for one of the following reasons:

- To restore a backed-up configuration file.
- To use the configuration file for another router. For example, you may add another router to the network and want it to have a similar configuration as the original router. By copying the file to the network server and making the changes to reflect the configuration requirements of the new router, you can save time by not recreating the entire file.
- To load the same configuration commands onto all of the routers in the network so that all of the routers have similar configurations.
- To use the configuration file for another router. For example, you may add another router.

For example, in the **copy running-config tftp** command, the running configuration in RAM is copied to a TFTP server.

Use the **copy running-config startup-config** command after a configuration change is made in RAM and must be saved to the startup configuration file in NVRAM. Similarly, copy the startup configuration file in NVRAM back into RAM with the **copy startup running:** command. Notice that you can abbreviate the commands.

Similar commands exist for copying between a TFTP server and either NVRAM or RAM.

The following examples show common **copy** command usage. The examples list two methods to accomplish the same tasks. The first example is a simple syntax, and the second example provides a more explicit syntax.

- Copy the running configuration from RAM to the startup configuration in NVRAM, overwriting the existing file:

```
R2# copy running-config startup-config
R2# copy system:running-config nvram:startup-config
```

- Copy the running configuration from RAM to a remote location, overwriting the existing file:

```
R2# copy running-config tftp
R2# copy system:running-config tftp
```

- Copy a configuration from a remote source to the running configuration, merging the new content with the existing file:

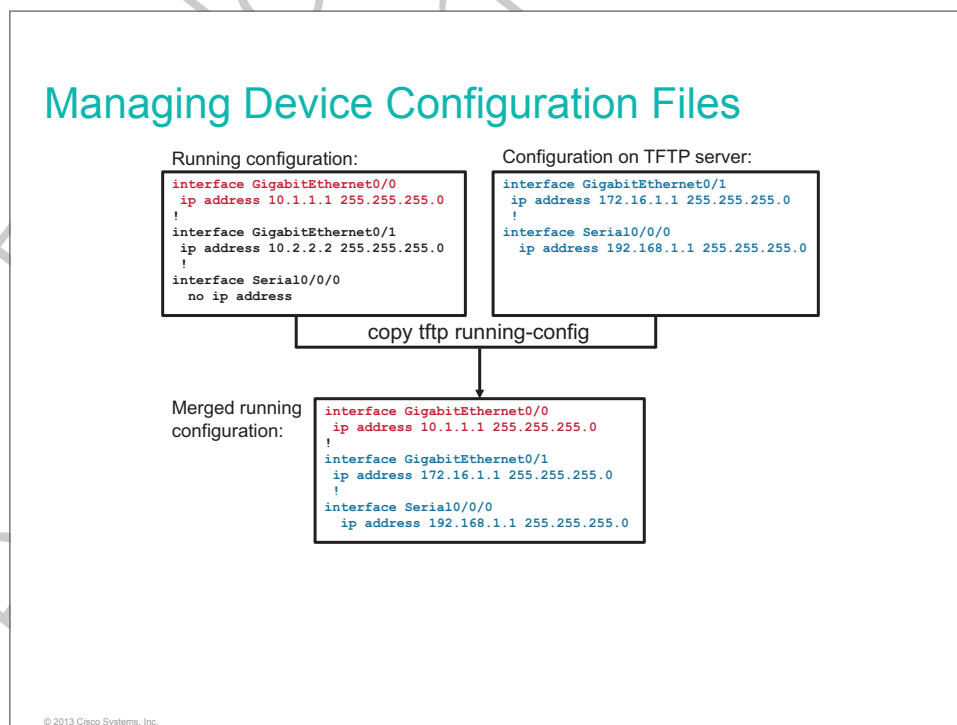
```
R2# copy tftp running-config
R2# copy tftp system:running-config
```

- Copy a configuration from a remote source to the startup configuration, overwriting the existing file:

```
R2# copy tftp startup-config
R2# copy tftp nvram:startup-config
```

Use the **configure terminal** command to interactively create configurations in RAM from the console or remote terminal.

Use the **erase startup-config** command to delete the saved startup configuration file in NVRAM.



This figure shows an example of how to use the **copy tftp running-config** command to merge the running configuration in RAM with a saved configuration file on a TFTP server.

Note When a configuration is copied into RAM from any source, the configuration merges with or overlays any existing configuration in RAM, rather than overwriting it. New configuration parameters are added, and changes to existing parameters overwrite the old parameters. Configuration commands that exist in RAM for which there are no corresponding commands in NVRAM remain unaffected. Copying the running configuration from RAM into the startup configuration file in NVRAM will overwrite the startup configuration file in NVRAM.

Managing Device Configuration Files (Cont.)

```
Branch# copy running-config tftp
Address or name of remote host []? 172.16.1.100
Destination filename [running-config]? config.cfg
.!!
1684 bytes copied in 13.300 secs (129 bytes/sec)
```

Upload and save the current configuration to a TFTP server

```
Branch# copy tftp running-config
Address or name of remote host []? 2001:DB8:AC10:100::64
Source filename []? config.cfg
Destination filename [running-config]?
Accessing tftp://2001:DB8:AC10:100::64/config.cfg...
Loading config.cfg from 2001:DB8:AC10:100::64 (via GigabitEthernet0/0): !
[OK - 1684/3072 bytes]

1684 bytes copied in 17.692 secs (99 bytes/sec)
```

Merge a configuration file from the TFTP server with the running configuration of the RAM.

© 2013 Cisco Systems, Inc.

You can use the TFTP servers to store configurations in a central place, allowing centralized management and updating. Regardless of the size of the network, there should always be a copy of the current running configuration online as a backup.

The **copy running-config tftp** command allows the current configuration to be uploaded and saved to a TFTP server. The IP address or name of the TFTP server and the destination filename must be supplied. A series of exclamation marks in the display shows the progress of the upload.

The **copy tftp running-config** command downloads a configuration file from the TFTP server to the running configuration of the RAM. Again, the address or name of the TFTP server and the source and destination filename must be supplied. In the example, IPv6 is used as a transport protocol. In this case, because you are copying the file to the running configuration, the destination filename should be running-config. This process is a merge process, not an overwrite process.

Password Recovery

An enable password controls access to the router privileged EXEC mode. If the password is mistyped or forgotten, access to the router privileged EXEC mode is not possible, and you must perform a password recovery procedure. This topic describes how to perform a password recovery on a Cisco router.

Password Recovery

The password recovery procedure differs for different router and switch platforms.

- 1 Switch off the router.
- 2 Switch on the router. Press **Break** to enter ROM monitor mode.
- 3 Once the router is on ROM monitor mode, set the configuration register to 0x2142.

```
rommon 1> confreg 0x2142
```
- 4 Reset the router.

```
rommon 1> reset
```
- 5 Enter privileged EXEC mode.

```
Router> enable
```

© 2013 Cisco Systems, Inc.

You can use the configuration register to perform the password recovery procedure. You have to set the configuration register value to a value that will instruct the router to ignore the startup configuration, which includes the forgotten enable password. Because a user cannot enter privileged EXEC mode in order to change the configuration register, the register has to be changed in ROM monitor. In order to enter ROM monitor, reboot the router and press **Break** to interrupt the boot process to get into ROM monitor.

Follow these steps to perform password recovery:

- Either switch off or shut down the router.
- Switch on the router. Press **Break** to interrupt the boot process to get into ROM monitor.
- When the router is in ROM monitor mode, set the configuration register to 0x2142. The hexadecimal number of 4 will instruct the router to ignore the startup configuration at the next reload.
- Reset the router. The router reboots but ignores the saved configuration. Do not enter the interactive setup dialog.
- Enter privileged EXEC mode. You should be able to do so because the saved configuration is ignored and the empty configuration without the enable password is loaded.

Password Recovery (Cont.)

6 Copy "startup-config" to "running-config."

```
Router# copy startup-config running-config
```

7 Bring up interfaces.

```
Router(config-if)# no shutdown
```

8 Enter global configuration mode and change the enable password.

```
Router# configure terminal
Router(config)# enable secret newpassword
```

9 Change the configuration register back to the initial value.

```
Router(config)# config-register 0x2102
```

© 2013 Cisco Systems, Inc.

Password Recovery (Cont.)

1 Copy "running-config" to "startup-config"

0

```
Router# copy running-config startup-config
```

© 2013 Cisco Systems, Inc.

- Copy "startup-config" to "running-config" in order to load the saved configuration. After this step, all interfaces might be disabled. You have to enable the desired interfaces using the **no shutdown** command.
- Because the startup configuration merged into the running configuration, the interfaces will be shut down. Bring up the appropriate interfaces using the **no shutdown** command.
- Enter global configuration mode and change the enable password to a new value. Do not forget or mistype the password this time.

- Change the configuration register back to the initial value. Change the value to the previously recorded value or to 0x2102. This will instruct the router to not ignore the startup configuration at the next reload.
- Copy "running-config" to "startup-config" in order to save changes regarding the new enable password and the configuration register value.

Note The password recovery procedure differs for different router and switch platforms. Refer to www.cisco.com for the password recovery procedure for different platforms.

Do Not Duplicate.
Post beta, not for release.

Summary

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

Summary

- The major components of a router are CPU, RAM, flash memory, ROM, NVRAM, and interfaces.
- A router first performs a POST test when booting.
- When a router boots, it searches for the Cisco IOS image in a specific sequence.
- When a router locates a valid Cisco IOS image in flash memory, the Cisco IOS image is loaded into RAM to run.
- After a router loads the Cisco IOS image, the router loads startup-config (if any startup-config is present on the router).
- The configuration register is a 16-bit number that affects router behavior, including locating a Cisco IOS image.
- You can use a TFTP server to store router configurations in a central place.

© 2013 Cisco Systems, Inc.

Licensing

This lesson explains the universality of Cisco IOS images and the concept behind licensing. The Cisco IOS Software Activation feature is an orchestrated collection of processes and components to activate Cisco IOS Software feature sets by obtaining and validating Cisco software licenses. With the Cisco IOS Software Activation feature, you can enable licensed features and register licenses.

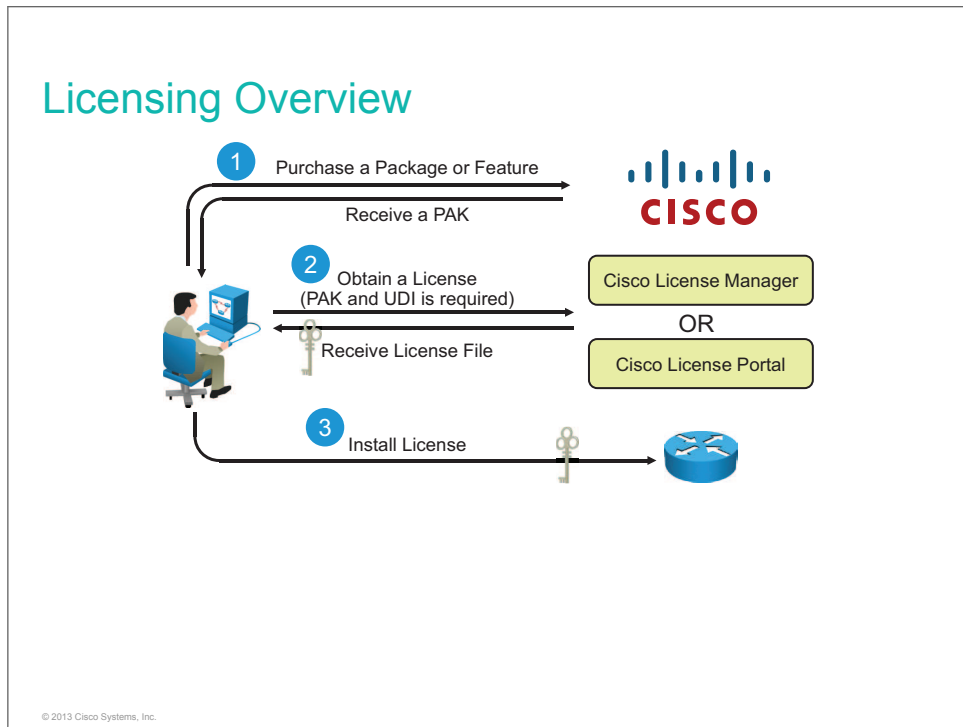
Objectives

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

- Explain the idea behind Cisco IOS image licensing
- Explain how a current version of the license can be identified
- Explain how to install a permanent license
- Explain how to install an evaluation license
- Explain how to back up a license
- Explain how to uninstall a license

Licensing Overview

This topic describes the universality of Cisco IOS images and the idea behind licensing.



When you order a new router, it is shipped preinstalled with the software images and the corresponding permanent licenses for the packages and features that you specified.

Note Use the Cisco IOS **show license** command to determine the licenses that are activated on your system.

Your router comes with an evaluation license, also known as a temporary license, for most packages and features that are supported on your router. If you want to try a new software package or feature, activate the evaluation license for this package or feature. If you want to permanently activate a software package or feature on your router, you must get a new software license.

Software Claim Certificates are used for licenses that require software activation. The claim certificate provides the PAK for your license and important information regarding the Cisco EULA. In most cases, Cisco or your Cisco partner will have already activated the licenses that were ordered at the time of purchase, and no Software Claim Certificate is provided.

Complete the following steps to permanently activate a software package or feature on the router:

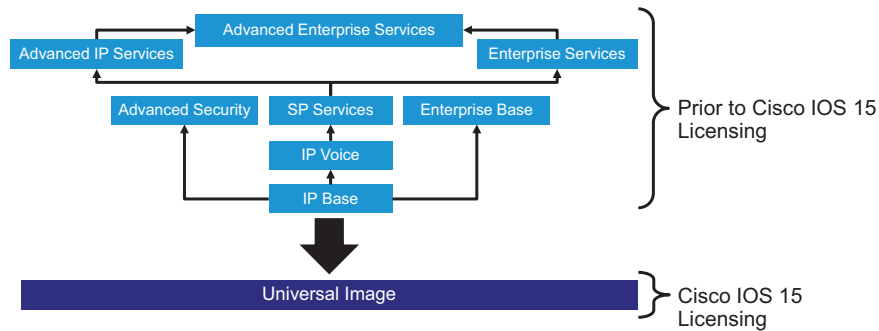
- 1 Purchase the software package or feature that you want to install. You receive a PAK with your purchase.
- 2 Get the license file using one of these options:
 - **Cisco License Manager**, which is a free software application available at <http://www.cisco.com/go/clm>.
 - **Cisco License Registration Portal**, the web-based portal for obtaining and registering individual software licenses, available at <http://www.cisco.com/go/license>.
- 3 Use the Cisco IOS CLI to install and manage licenses.

To obtain the license, you might also need the UDI, which has two main components: the PID and the serial number. The following example shows the output from the **show license udi** command that reveals the product ID and serial number of the router:

```
Router# show license udi
Device#   PID                               SN                               UDI
```

Licensing Overview (Cont.)

- Prior to Cisco IOS version 15.0, a software image was selected based on the required feature set of the customer.
- There were 8 software packages (images) that satisfied requirements in different service categories.



Prior to the Cisco IOS version 15.0 release, a software image was selected based on the required feature set of the customer. There were eight software packages (images) that satisfied requirements in different service categories. Cisco IOS Software Packaging consists of eight packages for Cisco routers.

Five packages are designed to satisfy requirements in four typical service categories:

| Software Image/Package | Features |
|------------------------|---|
| IP Base | IP data is the entry-level Cisco IOS Software Image |
| IP Voice | Converged voice and data, <u>VoIP</u> , <u>VoFR</u> , and IP telephony |
| Advanced Security | Security and VPN features including Cisco IOS Firewall, <u>IPS</u> , <u>IPSec</u> , <u>3DES</u> , and VPN |
| SP Services | Adds <u>SSH/SSL</u> , <u>ATM</u> , <u>VoATM</u> , and <u>MPLS</u> to IP voice |
| Enterprise Base | Enterprise protocols: multiprotocol support |

Three additional premium packages offer new Cisco IOS Software feature combinations that address more complex network requirements. All features merge in the Advanced Enterprise Services package that integrates support for all routing protocols with voice, security, and VPN capabilities:

| Software Image/Package | Features |
|------------------------|--|
| Advanced IP Services | Full Cisco IOS Software features |
| Enterprise Services | Enterprise base, full IBM support, and service provider services |
| SP Services | MPLS, ATM, and VoATM |

Feature inheritance is another powerful aspect of Cisco IOS Software Packaging. After a feature is introduced, it is included in the more comprehensive packages. Feature inheritance facilitates migration by clarifying the feature content of the different packages and how they relate to each other.

Licensing Overview (Cont.)

- Since the introduction of Cisco IOS 15.0 Software, the universal image contains all packages and features in *one* image.
- Multiple technology package licenses can be installed and activated on the Cisco 1900, 2900, and 3900 Series Integrated Services Router platforms.
- Individual features can be enabled or disabled by license keys.

| Technology Package License | Features |
|----------------------------|---|
| IP Base | Entry-level Cisco IOS functionality |
| DATA | MPLS, ATM, and multiprotocol support |
| Unified Communications | VoIP and IP telephony |
| Security | Cisco IOS Firewall, IPS, IPSEC, 3DES, and VPN |

© 2013 Cisco Systems, Inc.

Beginning with the Cisco 1900, 2900, and 3900 Series Integrated Services Routers, Cisco has revised the licensing model of Cisco IOS Software. Routers come with IP Base installed, and additional feature pack licenses can be installed as bolt-on additions to expand the feature set of the device.

The Cisco IOS universal image contains all packages and features in one image. The universal image on the Cisco 1900, 2900, and 3900 routers are a superset of Cisco IOS simplified technology packages. Each package is a grouping of technology-specific features. Multiple technology package licenses can be installed and activated on the Cisco 1900, 2900, and 3900 platforms.

Note Use the **show license feature** command to view the technology package licenses and feature licenses that are supported on your router.

Premium features beyond what is included in the default IP Base package are generally grouped into three major Technology Package Licenses: Data, Security, and Unified Communications. These three packages represent the vast majority of features that are available in Cisco IOS Software.

The following table lists the technology package licenses that are supported on Cisco ISR G2 platforms (Cisco 1900, 2900, and 3900 routers).

| Technology Package License | Features |
|-------------------------------|---|
| ipbasek9 (IP Base) | Entry-level Cisco IOS functionality |
| datak9 (DATA) | MPLS, ATM, and multiprotocol support |
| uck9 (Unified Communications) | VoIP and IP telephony |
| securityk9 (Security) | Cisco IOS Firewall, IPS, IPSEC, 3DES, and VPN |

Note The IP Base license is a prerequisite for installing the Data, Security, and Unified Communications license.

Do Not Duplicate.
Post beta, not for release.

Licensing Verification

This topic describes how to verify an installed license through CLI.

Licensing Verification

```
Router# show license
Index 1 Feature: ipbasek9
  Period left: Life time
  License Type: Permanent
  License State: Active, In Use
  License Count: Non-Counted
  License Priority: Medium
Index 2 Feature: securityk9
  Period left: Not Activated
  Period Used: 0 minute 0 second
  License Type: EvalRightToUse
  License State: Not in Use, EULA not accepted
  License Count: Non-Counted
  License Priority: None
<output omitted>
```

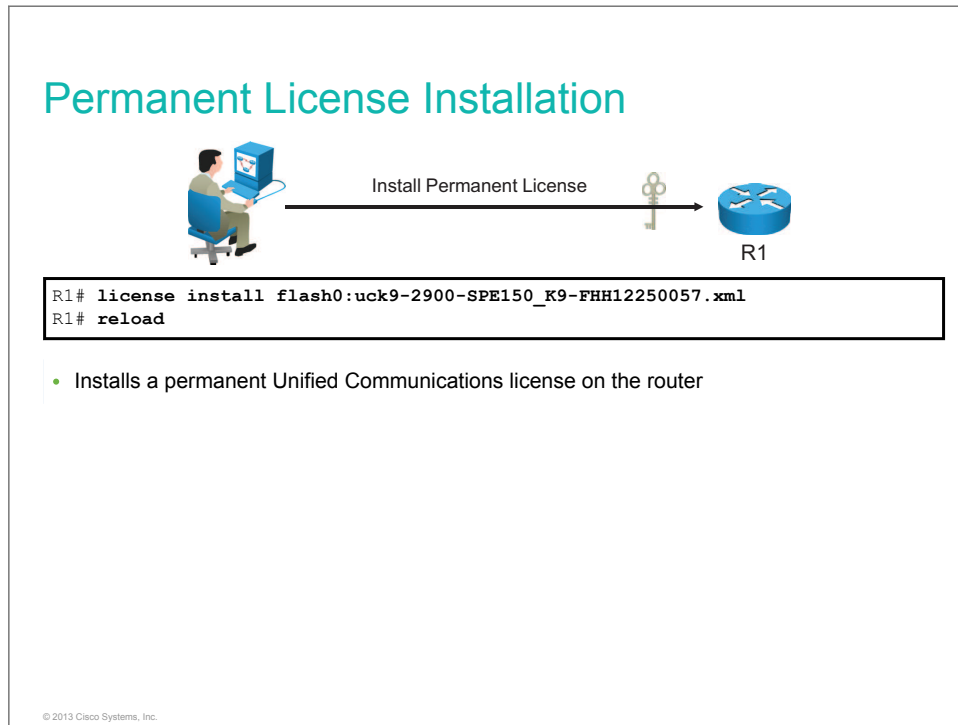
- Displays information about all Cisco IOS Software licenses

© 2013 Cisco Systems, Inc.

Use the **show license** command in privileged EXEC mode to see information about Cisco IOS Software licenses.

Permanent License Installation

This topic describes how to install a permanent license using the CLI.



Installing the permanent license is based on the steps that are described in the table.

| Command | Description |
|--|--|
| <code>license install stored-location-url</code> | Installs a license file. |
| <code>reload</code> | Reloads the router. A reload is not required if an evaluation license is active. A reload is required to activate a technology package license if an evaluation license is not active. |

The figure shows the configuration for installing the permanent Unified Communications license on the router. It is assumed that you obtained the license file from Cisco and stored it on the Flash of the router.

Permanent licenses are perpetual (that is, no usage period is associated with them). Once permanent licenses are installed, they provide all of the permissions that are needed to access features in the software image.

Note Cisco manufacturing preinstalls the appropriate permanent license on the ordered device for the purchased feature set. No customer interaction with the Cisco IOS Software Activation processes is required to enable a license on new hardware.

Use the **license install** command to install the permanent license:

```
R1# license install flash0:uck9-C2900-SPE150_K9-FHH12250057.xml
Installing licenses from "uck9-C2900-SPE150_K9-FHH12250057.xml"
Installing...Feature:uck9...Successful:Supported
1/1 licenses were successfully installed
0/1 licenses were existing licenses
0/1 licenses were failed to install
upt-3945-1#
```

```
*Jul  7 17:24:57.391: %LICENSE-6-INSTALL: Feature uck9 1.0 was installed in this
device.
UDI=C3900-SPE150/K9:FHH12250057; StoreIndex=15:Primary License Storage
*Jul  7 17:24:57.615: %IOS_LICENSE_IMAGE_APPLICATION-6-LICENSE_LEVEL: Module name =
c2900
Next reboot level = uck9 and License = uck9
```

Reload the router after the license is successfully installed using the **reload** command. Use the **show version** command after the router is reloaded to verify that the license has been installed.

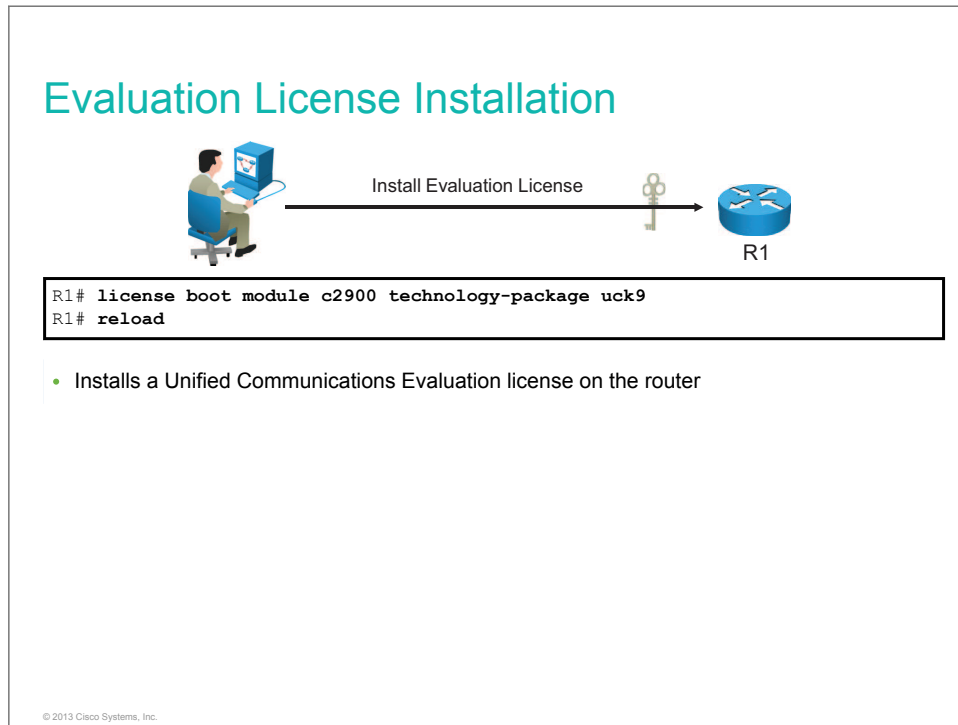
Note A reload is not required if an evaluation license is active. A reload is required to activate a Technology Package License if an Evaluation license is not active.

```
R1# show version
<output omitted>
License Info:
License UDI:
-----
Device#      PID                      SN
-----
*0           C3900-SPE150/K9         FHH12250057
Technology Package License Information for Module:'c2900'
-----
Technology   Technology-package      Technology-package
Current      Type                    Next reboot
-----
ipbase       ipbasek9               Permanent            ipbasek9
security     None                   None                 None
uc           uck9                   Permanent            uck9
data         None                   None                 None
Configuration register is 0x0
```

Do Not Duplicate. Post beta, not for release.

Evaluation License Installation

This topic describes how to install an Evaluation license using the CLI.



Note Starting with Cisco IOS Releases 15.0(1)M6, 15.1(1)T4, 15.1(2)T4, 15.1(3)T2, and 15.1(4)M, Evaluation licenses are replaced with Evaluation Right To Use licenses. Follow the steps that are detailed in the table to activate an Evaluation Right to Use license.

Activating the Evaluation license is based on the steps that are described in the table.

| Command | Description |
|---|---|
| license boot module <i>module-name</i> technology-package <i>package-name</i> | Enables the Evaluation license. Use the ? command with the module command to see the module name for your router and with the technology-package command to see the software packages and features that are supported on your router. |
| reload | Reloads the router. A reload is required to activate the software package. |

The figure shows the configuration for activating a Unified Communications Evaluation license on the router. Evaluation licenses are temporary, and you use them to evaluate a feature set on new hardware. Temporary licenses are limited to a specific usage period (for example, 60 days).

Use the **license boot module** command to enable the Evaluation license:

```
R1(config)#license boot module c2900 technology-package uck9
PLEASE READ THE FOLLOWING TERMS CAREFULLY.  INSTALLING THE LICENSE OR
LICENSE KEY PROVIDED FOR ANY CISCO PRODUCT FEATURE OR USING SUCH
PRODUCT FEATURE CONSTITUTES YOUR FULL ACCEPTANCE OF THE FOLLOWING
TERMS.  YOU MUST NOT PROCEED FURTHER IF YOU ARE NOT WILLING TO BE BOUND
BY ALL THE TERMS SET FORTH HEREIN.
Use of this product feature requires an additional license from Cisco,
```

together with an additional payment. You may use this product feature on an evaluation basis, **without payment to Cisco, for 60 days**. Your use of the product, including during the 60-day evaluation period, is subject to the Cisco End User License Agreement at <http://www.cisco.com/en/US/docs/general/warranty/English/EU1KEN.html>. If you use the product feature beyond the 60-day evaluation period, you must submit the appropriate payment to Cisco for the license. After the 60-day evaluation period, your use of the product features will be governed solely by the Cisco End User License Agreement (link above), together with any supplements relating to such product features. The above applies even if the evaluation license is not automatically terminated and you do not receive any notice of the expiration of the evaluation period. It is your responsibility to determine when the evaluation period is complete, and you are required to make payment to Cisco for your use of the product features beyond the evaluation period. Your acceptance of this agreement for the software features on one product shall be deemed your acceptance with respect to all such software on all Cisco products you purchase, which includes the same software. (The foregoing notwithstanding, you must purchase a license for each software feature you use past the 60-day evaluation period, so that if you enable a software feature on 1000 devices, you must purchase 1000 licenses for use past the 60-day evaluation period.) Activation of the software command line interface will be evidence of your acceptance of this agreement.

ACCEPT? [yes/no]: **yes**

% use 'write' command to make license boot config take effect on next boot

Nov 27 08:44:14.395: %IOS_LICENSE_IMAGE_APPLICATION-6-LICENSE_LEVEL: Module name=c2900 Next reboot level = uck9 and License = uck9

Nov 27 08:44:15.023: %LICENSE-6-EULA_ACCEPTED: EULA for feature uck9 1.0 has been accepted. UDI=CISCO2901/K9:FCZ1642C5XD; StoreIndex=1:Built-In License Storage

Reload the router after the license is successfully installed using the **reload** command. Use the **show license** command after the router is reloaded to verify that the license has been installed.

R1# **show license**

```
Index 1 Feature: ipbasek9
  Period left: Life time
  License Type: Permanent
  License State: Active, In Use
  License Count: Non-Counted
  License Priority: Medium
Index 2 Feature: securityk9
  Period left: Not Activated
  Period Used: 0 minute 0 second
  License Type: EvalRightToUse
  License State: Not in Use, EULA not accepted
  License Count: Non-Counted
  License Priority: None
Index 3 Feature: uck9
  Period left: 8 weeks 3 days
  Period Used: 9 minutes 30 seconds
  License Type: EvalRightToUse
  License State: Active, In Use
  License Count: Non-Counted
  License Priority:
```

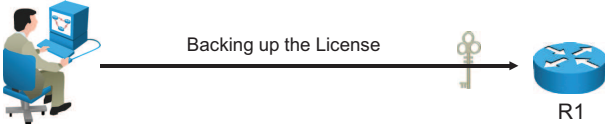
Low

<output omitted>

Backing up the License

This topic describes how to back up the license.

Backing up the License



```
R1# license save flash:all_licenses.lic
```

- Saves the license to the Flash of the router

© 2013 Cisco Systems, Inc.

Saving or backing up the license is based on the steps that are described in the table.

| Command | Description |
|---|---|
| <code>license save file-sys://lic-location</code> | Saves copies of all licenses in a device. <i>lic-location</i> : The license storage location can be a directory or a URL that points to a file system. Use the ? command to see the storage locations that are supported by your device. |

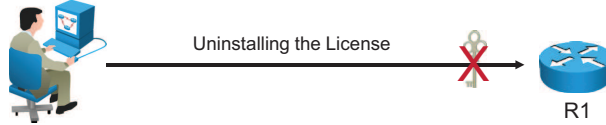
The figure shows the configuration for backing up the license on the router. Saved licenses are restored by using the **license install** command.

```
Router# license save flash:all_licenses.lic
license lines saved ..... to flash:all_licenses.lic
```

Uninstalling the License

This topic describes how to uninstall the license.

Uninstalling the License



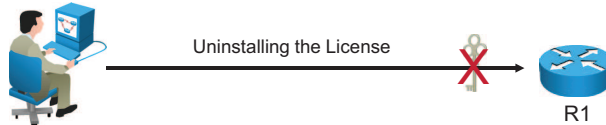
To clear an active permanent license from the router, perform the following tasks:

```
Router(config)# license boot module c3900 technology-package uck9 disable
Router(config)# exit
Router# reload
```

- Disable the technology package

© 2013 Cisco Systems, Inc.

Uninstalling the License (Cont.)



```
Router# license clear uck9
Router# configure terminal
Router(config)# no license boot module c3900 technology uck9 disable
Router(config)# exit
Router# reload
```

- Clear the license

© 2013 Cisco Systems, Inc.

Note Some licenses, such as built-in licenses, cannot be cleared. Only licenses that have been added by using the **license install** command are removed. Evaluation licenses are not removed.

To clear an active permanent license from the Cisco 3900, 2900, and 1900 routers, perform the following tasks:

- Disable the technology package
- Clear the license

Clearing the active permanent license is based on the steps that are described in the table.

| Command | Description |
|---|--|
| license boot module <i>module-name</i> technology-package <i>package-name</i> disable | Disables the active license. |
| reload | Reloads the router. A reload is required to make the software package inactive. |
| license clear <i>feature-name</i> | Clears the Technology Package License from license storage. |
| no license boot module <i>module-name</i> technology-package <i>package-name</i> disable | Clears the license boot module <i>module-name</i> technology-package <i>package-name</i> disable command that is used for disabling the active license. |
| reload | Reloads the router. A reload is required to make the software package inactive. |

The figure shows the configuration for how to clear an active permanent license from the Cisco 3900, 2900, and 1900 routers. First, you need to disable the technology package and then clear the license. Each of these two steps require reload of the router.

The following example shows how to clear an active license on a Cisco 3900 router:

```
Router(config)# license boot module c3900 technology-package uck9 disable
% use 'write' command to make license boot config take effect on next boot
Router(config)# exit
Router# reload
Router# license clear uck9
*Jul 7 00:34:23.691: %SYS-5-CONFIG_I: Configured from console by consoleclear uck9
Feature: uck9
  1 License Type: Permanent
    License State: Active, Not in Use
    License Addition: Exclusive
    License Count: Non-Counted
    Comment:
    Store Index: 15
    Store Name: Primary License Storage
Are you sure you want to clear? (yes/[no]): yes
upt-3945-1#
*Jul 7 00:34:31.223: %LICENSE-6-REMOVE: Feature uck9 1.0 was removed from this device.
UDI=C3900-SPE150/K9:FHH12250057; StoreIndex=15:Primary License Storage
Router#
Router# configure terminal
Router(config)# no license boot module c3900 technology uck9 disable
Router(config)# exit
Router# reload
```

Summary

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

Summary

- Obtain the license using Cisco License Manager or the Cisco License Registration Portal and use Cisco IOS commands to install the license.
- Use the **show license** command in privileged EXEC mode to see information about Cisco IOS Software licenses.
- Use the **license install** command to install the permanent license.
- Use the **license save** command to back up the license.
- Use the **license clear** command to remove the license.

© 2013 Cisco Systems, Inc.

Module Summary

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

Module Summary

- NetFlow provides statistics on packets flowing through the routing devices in the network, while SNMP provides a lot more statistics from networking devices.
- To maintain network operations with minimum down time, it is necessary to have procedures in place for backing up Cisco IOS images.
- The universal images on the Cisco 1900, 2900, and 3900 Series Integrated Services Routers are a superset of Cisco IOS simplified technology packages; each package is a grouping of technology-specific features.

Do Not Duplicate.
Post beta, not for release.

Module Self-Check

Do Not Duplicate.
Post beta, not for release.

Questions

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

1. The SNMP system consists of which three components? (Choose three.) (Source: Configuring Network Devices to Support Network Management Protocols)
 - A. SNMP manager
 - B. SNMP agent
 - C. SNMP trap
 - D. MIB
 - E. threshold
2. A router is configured with the **snmp-server community Cisco RO** command. An NMS is trying to access this router via SNMP. Which kind of access does the NMS have? (Source: Configuring Network Devices to Support Network Management Protocols)
 - A. The NMS can only graph obtained results.
 - B. The NMS can graph obtained results and change the host name of the router.
 - C. The NMS can only change the host name of the router.
3. You want to configure the router with IP address 10.1.1.1 to send syslog messages of all severities, including debugging, to the syslog server with IP address 192.168.1.240. Which configuration is correct? (Source: Configuring Network Devices to Support Network Management Protocols)
 - A. **logging 10.1.1.1**
logging trap debugging
 - B. **logging 192.168.1.240**
logging trap debugging
 - C. **logging 192.168.1.240**
logging debugging
 - D. **logging 10.1.1.1**
logging debugging
4. Which management protocol could you use to discover which host generates the highest volume of traffic? (Source: Configuring Network Devices to Support Network Management Protocols)
 - A. SNMP
 - B. syslog
 - C. NetFlow
5. Which stage during a Cisco router bootup process occurs last? (Source: Managing Cisco Devices)
 - A. POST
 - B. find and load Cisco IOS Software
 - C. find and load bootstrap
 - D. find and load configuration

6. Which stage of a Cisco router bootup process verifies that all router components are operational? (Source: Managing Cisco Devices)
- A. POST
 - B. find Cisco IOS Software
 - C. find bootstrap
 - D. find configuration
7. Which Cisco router component is used primarily to store the startup-config file? (Source: Managing Cisco Devices)
- A. RAM
 - B. ROM
 - C. NVRAM
 - D. flash memory
 - E. configuration register
8. Which one of the following is a low-level operating system that is normally used for manufacturing, testing, and troubleshooting? (Source: Managing Cisco Devices)
- A. POST
 - B. bootstrap
 - C. mini-Cisco IOS Software
 - D. ROM monitor
9. What happens if the router cannot find a valid startup configuration file in NVRAM during router bootup? (Source: Managing Cisco Devices)
- A. The router enters the setup mode.
 - B. The router attempts to restart.
 - C. The router runs the ROM monitor.
 - D. The router performs a shutdown.
10. Which Cisco IOS command is used to download a copy of the Cisco IOS image file from a TFTP server? (Source: Managing Cisco Devices)
- A. **copy ios tftp**
 - B. **copy tftp flash**
 - C. **copy flash tftp**
 - D. **backup flash tftp**
11. Which Cisco IOS command displays the amount of memory that is available where the Cisco IOS image is stored on your router? (Source: Managing Cisco Devices)
- A. **show flash**
 - B. **show nvram**
 - C. **show memory**
 - D. **show running-config**

12. Where is the running configuration of the router usually stored? (Source: Managing Cisco Devices)
- A. BIOS
 - B. RAM
 - C. NVRAM
 - D. bootflash
13. Which technology package license provides basic Cisco IOS functionality for the router? (Source: Licensing)
- A. ipbase
 - B. data
 - C. unified communications
 - D. security
14. Which command is used to install a permanent license? (Source: Licensing)
- A. **install license**
 - B. **license boot module**
 - C. **license install**
 - D. **license load**

Do Not Duplicate.
Post beta, not for release.

Answer Key

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

Do Not Duplicate:
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Do Not Duplicate:
Post beta, not for release.

Do Not Duplicate.
Post beta, not for release.

Glossary

| Term | Definition |
|-----------|---|
| ABR | area border router. Router located on the border of one or more OSPF areas that connects those areas to the backbone network. ABRs are considered members of both the OSPF backbone and the attached areas. They therefore maintain routing tables describing both the backbone topology and the topology of the other areas. |
| ACL | access control list. A list kept by routers to control access to or from the router for a number of services (for example, to prevent packets with a certain IP address from leaving a particular interface on the router). |
| AD | advertised distance |
| algorithm | Well-defined rule or process for arriving at a solution to a problem. In networking, algorithms commonly are used to determine the best route for traffic from a particular source to a particular destination. |
| ANSI | American National Standards Institute. A voluntary organization composed of corporate, government, and other members that coordinates standards-related activities, approves U.S. national standards, and develops positions for the United States in international standards organizations. ANSI helps develop international and U.S. standards relating to, among other things, communications and networking. ANSI is a member of the IEC and the ISO. |
| ARP | Address Resolution Protocol. Internet protocol that is used to map an IP address to a MAC address. Defined in RFC 826. |
| AS | autonomous system. A collection of networks under a common administration sharing a common routing strategy. Autonomous systems are subdivided by areas. An autonomous system must be assigned a unique 16-bit number by the IANA. |
| ASBR | autonomous system boundary router. An ABR located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area. |
| ASCII | American Standard Code for Information Interchange. 8-bit code for character representation (7 bits plus parity). |
| ATM | Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as E3, SONET, and T3. |
| b/s | bits per second |
| b/s | bits per second. |
| BDR | backup designated router. |

| Term | Definition |
|--------------------------|--|
| BGP | Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by RFC 1163. |
| BID | bridge ID. |
| blocking | In a switching system, a condition in which no paths are available to complete a circuit. The term also is used to describe a situation in which one activity cannot begin until another is completed. |
| BPDU | bridge protocol data unit. Spanning Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network. |
| broadcast | Data packet that are sent to all nodes on a network. Broadcasts are identified by a broadcast address. |
| CHAP | Challenge Handshake Authentication Protocol. Security feature supported on lines using PPP encapsulation that prevents unauthorized access. CHAP does not itself prevent unauthorized access, but merely identifies the remote end. The router or access server then determines whether that user is allowed access. |
| CIR | committed information rate. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. CIR, measured in bits per second, is one of the key negotiated tariff metrics. |
| Cisco ASA | Cisco Adaptive Security Appliance |
| Cisco Discovery Protocol | Media- and protocol-independent device-discovery protocol. Using Cisco Discovery Protocol, a device can advertise its existence to other devices and receive information about other devices on the same LAN or on the remote side of a WAN. Cisco Discovery Protocol runs on all media that support SNAP, including LANs, Frame Relay, and ATM media. |
| Cisco ISR G2 | Cisco Integrated Services Routers Generation 2 |
| CO | central office. The local telephone company office to which all local loops in a given area connect and in which circuit switching of subscriber lines occurs. |
| CPE | customer premises equipment. Terminating equipment, such as terminals, telephones, and modems, supplied by the telephone company, installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment residing on the customer site. |
| CRC | cyclic redundancy check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node. |
| CST | Common Spanning Tree. |

| Term | Definition |
|-------------|--|
| CSU | channel service unit. Digital interface device that connects end-user equipment to the local digital telephone loop. Often referred to together with DSU, as <i>CSU/DSU</i> . |
| DBD | database description. |
| DCE | data communications equipment (EIA expansion). |
| DE | discard eligible. If the network is congested, DE traffic can be dropped to ensure the delivery of higher priority traffic. |
| DLCI | data-link connection identifier. Value that specifies a PVC or an SVC in a Frame Relay network. In the basic Frame Relay specification, DLCIs are locally significant (connected devices might use different values to specify the same connection). In the LMI extended specification, DLCIs are globally significant (DLCIs specify individual end devices). |
| DNS | Domain Name System. System used on the Internet for translating names of network nodes into addresses. |
| DoS | denial of service. |
| DR | designated router. |
| DSU | data service unit. Device used in digital transmission that adapts the physical interface on a DTE device to a transmission facility, such as T1 or E1. The DSU also is responsible for such functions as signal timing. Often referred to together with CSU, as <i>CSU/DSU</i> . |
| DTE | data terminal equipment. Device at the user end of a user-network interface that serves as a data source, destination, or both. DTE connects to a data network through a DCE device (for example, a modem) and typically uses clocking signals generated by the DCE. DTE includes such devices as computers, protocol translators, and multiplexers. |
| DTP | Dynamic Trunking Protocol. |
| DUAL | Diffusing Update Algorithm. Convergence algorithm used in EIGRP that provides loop-free operation at every instant throughout a route computation. Allows routers involved in a topology change to synchronize at the same time, while not involving routers that are unaffected by the change. |
| EGP | Exterior Gateway Protocol. Internet protocol for exchanging routing information between autonomous systems. Documented in RFC 904. Not to be confused with the general term <i>exterior gateway protocol</i> . EGP is an obsolete protocol that was replaced by BGP. |
| EIA/TIA-232 | Common physical layer interface standard, developed by EIA and TIA, that supports unbalanced circuits at signal speeds of up to 64 kb/s. Closely resembles the V.24 specification. Formerly called <i>RS-232</i> . |

| Term | Definition |
|---------------|---|
| EIGRP | Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency, and combines the advantages of link-state protocols with those of distance vector protocols. |
| endpoint | H.323 terminal or gateway. An endpoint can call and be called. It generates and terminates the information stream. |
| EtherChannel | Developed and copyrighted by Cisco Systems. Logical aggregation of multiple Ethernet interfaces used to form a single higher bandwidth routing or bridging endpoint. |
| EUI | extended universal identifier |
| EULA | End User License Agreement. |
| Fast Ethernet | Any of a number of 100-Mb/s Ethernet specifications. Fast Ethernet offers a speed increase 10 times that of the 10BaseT Ethernet specification while preserving such qualities as frame format, MAC mechanisms, and MTU. Such similarities allow the use of existing 10BaseT applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification. |
| FCS | frame check sequence. Extra characters added to a frame for error control purposes. Used in HDLC, Frame Relay, and other data link layer protocols. |
| FD | feasible distance |
| FFS | flash file system |
| flooding | Traffic passing technique used by switches and bridges in which traffic received on an interface is sent out all the interfaces of that device except the interface on which the information was received originally. |
| FLP | Fast Link Pulse. A type of link pulse that encodes information used in autonegotiation. |
| forwarding | Process of sending a frame toward its ultimate destination by way of an internetworking device. |
| FRAD | Frame Relay access device. Any network device that provides a connection between a LAN and a Frame Relay WAN. |
| frame | Logical grouping of information sent as a data link layer unit over a transmission medium. Often refers to the header and the trailer, used for synchronization and error control, that surround the user data contained in the unit. The terms <i>cell</i> , <i>datagram</i> , <i>message</i> , <i>packet</i> , and <i>segment</i> also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles. |
| FTP | File Transfer Protocol. Application protocol, part of the TCP/IP protocol stack, used for transferring files between network nodes. FTP is defined in RFC 959. |

| Term | Definition |
|------------------|--|
| full duplex | Capability for simultaneous data transmission between a sending station and a receiving station. |
| gateway | In the IP community, an older term referring to a routing device. Today, the term <i>router</i> is used to describe nodes that perform this function, and <i>gateway</i> refers to a special-purpose device that performs an application-layer conversion of information from one protocol stack to another. |
| Gigabit Ethernet | Standard for a high-speed Ethernet, approved by the IEEE (Institute of Electrical and Electronics Engineers) 802.3z standards committee in 1996. |
| GLBP | Gateway Load Balancing Protocol. |
| GRE | Generic Routing Encapsulation. Tunneling protocol that was developed by Cisco and that can encapsulate a variety of protocol packet types inside IP tunnels. This process creates a virtual point-to-point link to Cisco routers at remote points over an IP network. |
| half duplex | Capability for data transmission in only one direction at a time between a sending station and a receiving station. BSC is an example of a half-duplex protocol. |
| HDLC | high-level data link control. Bit-oriented synchronous data link layer protocol developed by ISO. Derived from SDLC, HDLC specifies a data encapsulation method on synchronous serial links using frame characters and checksums. |
| host | Computer system on a network. Similar to node, except that host usually implies a computer system, whereas node generally applies to any networked system, including access servers and routers. |
| HSRP | Hot Standby Router Protocol. Provides high network availability and transparent network topology changes. HSRP creates a hot standby router group with a lead router that services all packets sent to the hot standby address. The lead router is monitored by other routers in the group. If it fails, one of the standby routers inherits both the lead position and the hot standby address. |
| IANA | Internet Assigned Numbers Authority. Organization operated under the auspices of the ISOC as a part of the IAB. IANA delegates authority for IP address-space allocation and domain-name assignment to the InterNIC and other organizations. IANA also maintains a database of assigned protocol identifiers used in the TCP/IP stack, including autonomous system numbers. |
| ICMP | Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792. |
| IEEE | Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today. |

| Term | Definition |
|--------------|---|
| IEEE 802.1Q | Networking standard that supports VLANs (virtual LANs) on an Ethernet network. The standard has a system for VLAN tagging for frames on the Ethernet medium and procedures to be used by switches in handling those frames. |
| IEEE 802.1Q | The IEEE 802.1Q specification establishes a standard method for tagging Ethernet frames with VLAN membership information and defines the operation of VLAN bridges that permit the definition, operation, and administration of VLAN topologies within a bridged LAN infrastructure. |
| IEEE 802.3 | IEEE LAN protocol that specifies an implementation of the physical layer and the MAC sublayer of the data link layer. IEEE 802.3 uses CSMA/CD access at a variety of speeds over a variety of physical media. Extensions to the IEEE 802.3 standard specify implementations for Fast Ethernet. Physical variations of the original IEEE 802.3 specification include 10Base2, 10Base5, 10BaseF, 10BaseT, and 10Broad36. Physical variations for Fast Ethernet include 100BaseT, 100BaseT4, and 100BaseX. |
| IETF | Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC. |
| IFS | IOS File System feature, which provides a single interface to all of the file systems that a router uses. |
| IGP | interior gateway protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP. |
| IKE | Internet Key Exchange. IKE establishes a shared security policy and authenticates keys for services (such as IPsec) that require keys. Before any Ipse traffic can be passed, each router/firewall/host must verify the identity of its peer. This can be done by manually entering pre-shared keys into both hosts or by a CA service. |
| INIT state | initial state. |
| internetwork | Collection of networks interconnected by routers and other devices that functions (generally) as a single network. Sometimes called an internet, which is not to be confused with the Internet. |
| Inverse ARP | Inverse Address Resolution Protocol. Method of building dynamic routes in a network. Allows an access server to discover the network address of a device associated with a virtual circuit. |
| IP | Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security. Defined in RFC 791. |

| Term | Definition |
|------------|---|
| IP address | 32-bit address assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. CIDR provides a new way of representing IP addresses and subnet masks. Also called an Internet address. |
| IPS | intrusion prevention system |
| IPsec | IP Security. A framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPsec provides these security services at the IP layer. IPsec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPsec. IPsec can protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. |
| IPv4 | IP version 4 |
| IPv6 | IP version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (next generation). |
| IPX | Internetwork Packet Exchange. NetWare network layer (Layer 3) protocol used for transferring data from servers to workstations. IPX is similar to IP and XNS. |
| ISDN | Integrated Services Digital Network. Communication protocol offered by telephone companies that permits telephone networks to carry data, voice, and other source traffic. |
| IS-IS | Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, whereby ISs (routers) exchange routing information based on a single metric to determine network topology. |
| kb/s | kilobits per second. A bit rate expressed in thousands of bits per second. |
| LACP | Link Aggregation Control Protocol. |
| LAN | local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies. |
| LCP | link control protocol. Protocol that establishes, configures, and tests data-link connections for use by PPP. |

| Term | Definition |
|-----------------|--|
| LMI | Local Management Interface. Set of enhancements to the basic Frame Relay specification. LMI includes support for a keepalive mechanism, which verifies that data is flowing; a multicast mechanism, which provides the network server with its local DLCI and the multicast DLCI; global addressing, which gives DLCIs global rather than local significance in Frame Relay networks; and a status mechanism, which provides an on-going status report on the DLCIs known to the switch. Known as LMT in ANSI terminology. |
| LMI | Local Management Interface. Set of enhancements to the basic Frame Relay specification. LMI includes support for a keepalive mechanism, which verifies that data is flowing; a multicast mechanism, which provides the network server with its local DLCI and the multicast DLCI; global addressing, which gives DLCIs global rather than local significance in Frame Relay networks; and a status mechanism, which provides an on-going status report on the DLCIs known to the switch. Known as LMT in ANSI terminology. |
| load balancing | In routing, the capability of a router to distribute traffic over all its network ports that are the same distance from the destination address. Good load-balancing algorithms use both line speed and reliability information. Load balancing increases the use of network segments, thus increasing effective network bandwidth. |
| logical channel | Nondedicated, packet-switched communications path between two or more network nodes. Packet switching allows many logical channels to exist simultaneously on a single physical channel. |
| loop | Route where packets never reach their destination, but simply cycle repeatedly through a constant series of network nodes. |
| LSA | link-state advertisement. Broadcast packet used by link-state protocols that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables. Sometimes called an LSP. |
| LSAck | link-state acknowledgement. |
| LSDB | link-state database. |
| LSR | link-state request. |
| LSU | link-state update. |
| MAC | Media Access Control. Lower of the two sublayers of the data link layer defined by the IEEE. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used. |
| MAC address | Standardized data link layer address that is required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are 6 bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address. |

| Term | Definition |
|-------------------|--|
| MAC address | Standardized data link layer address that is required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are 6 bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address. |
| MAC address table | A switch will forward the frame only to the port where the destination is connected. But the switch needs to know to which port is the destination device connected. MAC address table lists which MAC address is connected to which port. MAC table is populated from previous communication done by the devices that communicate between themselves. Initially (for example, after a restart) the switch doesn't have this information, so it has to send frames out of all its ports. |
| Mb/s | megabits per second. A bit rate expressed in millions of binary bits per second. |
| MD5 | Message Digest 5. A one-way hashing algorithm that produces a 128-bit hash. Both MD5 and Secure Hash Algorithm (SHA) are variations on MD4 and are designed to strengthen the security of the MD4 hashing algorithm. Cisco uses hashes for authentication within the IPsec framework. Also used for message authentication in SNMP v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness. |
| MIB | Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as SNMP or CMIP. The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches. |
| MPLS | Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information. |
| MTU | maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle. |
| multicast | Single packets copied by the network and sent to a specific subset of network addresses. These addresses are specified in the Destination Address field. |
| multicast address | Single address that refers to multiple network devices. Synonymous with group address. |
| NAD | network access device. |
| NBMA | nonbroadcast multiaccess. Term describing a multiaccess network that either does not support broadcasting (such as X.25) or in which broadcasting is not feasible (for example, an SMDS broadcast group or an extended Ethernet that is too large). |

| Term | Definition |
|---------|---|
| NCP | Network Control Protocol. Series of protocols for establishing and configuring different network layer protocols, such as for AppleTalk over PPP. |
| NCP | Network Control Protocol. Series of protocols for establishing and configuring different network layer protocols, such as for AppleTalk over PPP. |
| NetFlow | A feature of some routers that allows them to categorize incoming packets into flows. Because packets in a flow often can be treated in the same way, this classification can be used to bypass some of the work of the router and accelerate its switching operation. |
| NFS | Network File System. As commonly used, a distributed file system protocol suite developed by Sun Microsystems that allows remote file access across a network. In actuality, NFS is simply one protocol in the suite. NFS protocols include NFS, RPC, XDR, and others. These protocols are part of a larger architecture that Sun refers to as ONC. |
| NIC | network interface card. Board that provides network communication capabilities to and from a computer system. Also called an adapter. |
| NMS | network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer, such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources. |
| NVRAM | nonvolatile RAM. RAM that retains its contents when a unit is powered off. |
| OID | object identifier. Values are defined in specific MIB modules. The Event MIB allows a user or an NMS to watch over specified objects and to set event triggers based on existence, threshold, and boolean tests. An event occurs when a trigger is fired; this means that a specified test on an object returns a value of true. To create a trigger, a user or an NMS configures a trigger entry in the mteTriggerTable of the Event MIB. This trigger entry specifies the OID of the object to be watched. For each trigger entry type, corresponding tables (existence, threshold, and boolean tables) are populated with the information required for carrying out the test. The MIB can be configured so that when triggers are activated (fired) either an SNMP Set is performed, a notification is sent out to the interested host, or both. |
| OSI | Open Systems Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability. |
| OSPF | Open Shortest Path First. Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing. OSPF was derived from an early version of the IS-IS protocol. |
| OUI | Organizational Unique Identifier. Three octets that are assigned by the IEEE in a block of 48-bit LAN addresses. |

| Term | Definition |
|----------|---|
| packet | Logical grouping of information that includes a header containing control information and (usually) user data. Packets most often are used to refer to network layer units of data. The terms <i>datagram</i> , <i>frame</i> , <i>message</i> , and <i>segment</i> also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles. |
| PAGP | Port Aggregation Protocol. |
| PAK | Product Authorization Key. |
| PAP | Password Authentication Protocol. Authentication protocol that allows PPP peers to authenticate one another. The remote router attempting to connect to the local router is required to send an authentication request. Unlike CHAP, PAP passes the password and the host name or username in the clear (unencrypted). PAP does not itself prevent unauthorized access but merely identifies the remote end. The router or access server then determines whether that user is allowed access. PAP is supported only on PPP lines. |
| PCMCIA | Personal Computer Memory Card International Association. Standard used for credit-card-sized computer peripherals. Type 1 devices are very thin memory cards, Type 2 devices include most modems and interfaces, and Type 3 devices are used for disk drives and thicker components. |
| PID | product identifier |
| POP | point of presence. In OSS, a physical location where an interexchange carrier installed equipment to interconnect with a local exchange carrier (LEC). |
| POST | power-on self test. Set of hardware diagnostics that runs on a hardware device when this device is powered on. |
| PPP | Point-to-Point Protocol. Successor to SLIP that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and ARA. PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: LCP and NCP. |
| PPPoA | PPP over ATM. |
| PPPoE | PPP over Ethernet. |
| protocol | Formal description of a set of rules and conventions that govern how devices on a network exchange information. |
| PSK | pre-shared key. |
| PSTN | public switched telephone network. General term referring to the variety of telephone networks and services in place worldwide. Sometimes called POTS. |

| Term | Definition |
|-------------|--|
| PVC | permanent virtual circuit (or connection). Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. In ATM terminology, called a permanent virtual connection. |
| PVST+ | Per-VLAN Spanning Tree Plus. Support for dot1q trunks to map multiple spanning trees to a single spanning tree. |
| QoS | quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability. |
| queue | In routing, a backlog of packets waiting to be forwarded over a router interface. |
| RAM | random-access memory. Volatile memory that can be read and written by a microprocessor. |
| rcp | remote copy protocol. Protocol that allows users to copy files to and from a file system residing on a remote host or server on the network. The rcp protocol uses TCP to ensure the reliable delivery of data. |
| redundancy | In internetworking, the duplication of devices, services, or connections so that, in the event of a failure, the redundant devices, services, or connections can perform the work of those that failed. |
| RFC | Request For Comments. Document series used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some are humorous or historical. RFCs are available online from numerous sources. |
| RIP | Routing Information Protocol. IGP supplied with UNIX BSD systems. The most common IGP in the Internet. RIP uses hop count as a routing metric. |
| RIR | regional Internet registry |
| root bridge | Exchanges topology information with designated bridges in a spanning-tree implementation to notify all other bridges in the network when topology changes are required. This prevents loops and provides a measure of defense against link failure. |
| router | Network layer device that uses one or more metrics to determine the optimal path along which network traffic should be forwarded. Routers forward packets from one network to another based on network layer information. Occasionally called a gateway (although this definition of gateway is becoming increasingly outdated). |
| RSTP | Rapid Spanning Tree Protocol. |
| SMTP | Simple Mail Transfer Protocol. Internet protocol providing email services. |

| Term | Definition |
|--------|---|
| SNMP | Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security. |
| SOHO | small office, home office. Networking solutions and access technologies for offices that are not directly connected to large corporate networks. |
| SPF | shortest path first. Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms. Sometimes called Dijkstra's algorithm. |
| SRTT | smoothed round-trip time |
| SSH | Secure Shell Protocol. Protocol that provides a secure remote connection to a router through a TCP application. |
| SSL | Secure Socket Layer. Encryption technology for the Web used to provide secure transactions, such as the transmission of credit card numbers for e-commerce. |
| STP | Spanning Tree Protocol. Bridge protocol that uses the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops, and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning Tree Protocol standard and the earlier Digital Equipment Corporation Spanning Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version generally is preferred over the Digital version. |
| SVC | switched virtual circuit. Virtual circuit that is dynamically established on demand and is torn down when transmission is complete. SVCs are used in situations where data transmission is sporadic. Called a switched virtual connection in ATM terminology. |
| syslog | system logging |
| T1 | Digital WAN carrier facility. T1 transmits DS-1-formatted data at 1.544 Mb/s through the telephone-switching network, using AMI or B8ZS coding. |
| T3 | Digital WAN carrier facility. T3 transmits DS-3-formatted data at 44.736 Mb/s through the telephone switching network. |
| TDM | time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots. Bandwidth is allocated to each channel regardless of whether the station has data to transmit. |
| telco | Abbreviation for telephone company. |

| Term | Definition |
|----------|---|
| Telnet | Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log into remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854. |
| Telnet | Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log into remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854. |
| TFTP | Trivial File Transfer Protocol. Simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, username and password). |
| topology | Physical arrangement of network nodes and media within an enterprise networking structure. |
| ToS | type of service |
| trunk | Physical and logical connection between two switches across which network traffic travels. A backbone is composed of a number of trunks. |
| TTL | Time to Live. |
| UDI | universal device identifier. |
| UDP | User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768. |
| unicast | Message sent to a single network destination. |
| V.35 | ITU-T standard describing a synchronous, physical layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and in Europe, and is recommended for speeds up to 48 kb/s. |
| VC | virtual circuit. |
| VID | VLAN ID. The identification of the VLAN, which is used by the standard IEEE 802.1Q. Being 12 bits, it allows for the identification of 4096 VLANs. |
| VLAN | virtual LAN. Group of devices on one or more LANs that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible. |
| VLSM | variable-length subnet mask. Capability to specify a different subnet mask for the same network number on different subnets. VLSM can help optimize available address space. |

| Term | Definition |
|--------|---|
| VoATM | Voice over ATM. Voice over ATM enables a router to carry voice traffic (for example, telephone calls and faxes) over an ATM network. When sending voice traffic over ATM, the voice traffic is encapsulated using a special AAL5 encapsulation for multiplexed voice. |
| VoFR | Voice over Frame Relay. VoFR enables a router to carry voice traffic (for example, telephone calls and faxes) over a Frame Relay network. When sending voice traffic over Frame Relay, the voice traffic is segmented and encapsulated for transit across the Frame Relay network using FRF.12 encapsulation. |
| VoIP | Voice over IP. The capability to carry normal telephony-style voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls and faxes) over an IP network. In VoIP, the DSP segments the voice signal into frames, which then are coupled in groups of two and stored in voice packets. These voice packets are transported using IP in compliance with ITU-T specification H.323. |
| VPN | virtual private network. Enables IP traffic to travel securely over a public TCP/IP network by encrypting all traffic from one network to another. A VPN uses tunneling to encrypt all information at the IP level. |
| VRRP | Virtual Router Redundancy Protocol. |
| WAN | wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs. |
| WAN | wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs. |
| WIC | WAN interface card. Connects the system to the WAN link service provider. |
| X.25 | ITU-T standard that defines how connections between DTE and DCE are maintained for remote terminal access and computer communications in PDNs. X.25 specifies LAPB, a data link layer protocol, and PLP, a network layer protocol. Frame Relay has to some degree superseded X.25. |
| 3DES | Triple Data Encryption Standard. |
| 802.1Q | Networking standard that supports VLANs (virtual LANs) on an Ethernet network. The standard has a system for VLAN tagging for frames on the Ethernet medium and procedures to be used by switches in handling those frames. |

Do Not Duplicate.
Post beta, not for release.