General Concepts

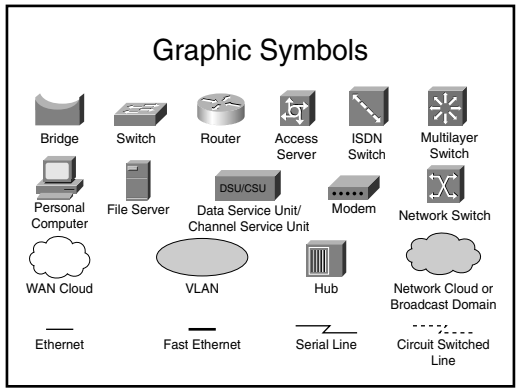

Defining Networks

Several different types of users access the network from many locations:

- **Main office**—Most corporate information is located here. Everyone is connected to the LAN.
- **Branch office**—Remote sites with a separate LAN access the main office through the WAN.
- **Private residences**—Many employees work out of their homes, which become part of the network.
- **Other sites**—Mobile users can connect from virtually anywhere.

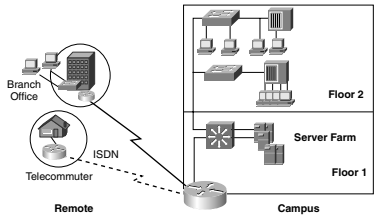

Hierarchical Model

Cisco uses a hierarchical network model. High traffic loads create a need for efficient routing and switching techniques.

Defining a Network's Key Points

Cisco uses a hierarchical network model. The three layers are the access layer, the distribution layer, and the core layer:

- **Access layer**—Provides user connectivity to the network.
- **Distribution layer**—Responsible for routing, filtering, and WAN access.
- **Core layer**—Responsible for fast-switching services.

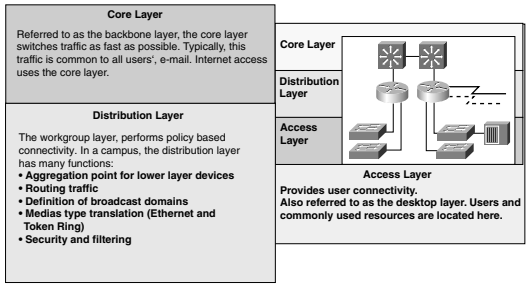

OSI Model

The *OSI model* is a standardized framework for network functions and schemes. It breaks down otherwise complex network interactions into simple elements, allowing developers to modularize design efforts. This method allows many independent developers to work on separate network functions that can be applied in a "plug-and-play" manner.

OSI Model

Application	User interface	Telnet HTTP
Presentation	Encryption and other processing	ASCII JPEG
Session	Manages multiple applications	Operating systems Scheduling

OSI Model *(Continued)*

Transport	Provides reliable or unreliable delivery and some error correction	TCP UDP SPX
Network	Provides logical addressing used by routers	IP IPX
Data link	Creates frames from bits of data Uses MAC addresses to access endpoints Provides error detection but no correction	802.3 802.2 HDLC
Physical	Specifies voltage, wire speed, and pinout cables	EIA/TIA V.35

Protocol data units (PDUs) are used to communicate between layers.
Encapsulation is the method of adding headers and trailers as data moves down the stack. The receiving device strips the header, which contains directions for that layer (de-encapsulation).

OSI Model Summary

- The OSI model provides a standardized way to create and implement network standards and schemes.
- The OSI model allows plug-and-play applications, simplified building blocks, and modularized development.
- The OSI model has seven layers. Mnemonics are useful for remembering the layers and their functions (such as Pick Don's Nose Then Spit Potatoes Afterward).
- Encapsulation is the process of adding layer-specific instructions (for the receiving device) as headers and trailers.
- De-encapsulation is the reverse process of encapsulation.

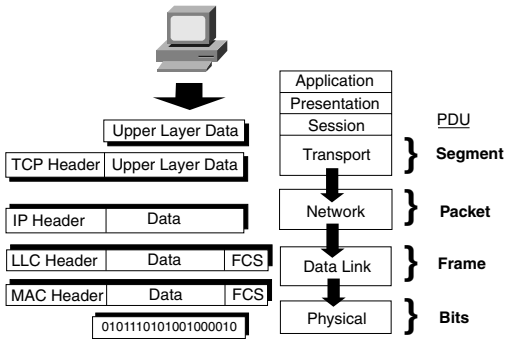

Lower (Data Link) Layers

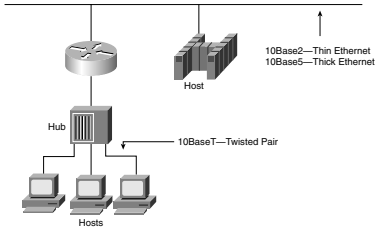

Physical layer functions are as follows:

- Media type
- Connector type
- Signaling type

The physical layer specifies

- Voltage levels
- Data rates
- Maximum transmission rates and distances
- Physical connectors and pinouts

Type	Name	Distance	Carrier
10Base2	Thinnet	Up to 185 meters	Coaxial
10Base5	Thicknet	500 meters	Coaxial
10BaseT	Ethernet signals	100 meters	Twisted pair

Collision/Broadcast Domains

All stations on an Ethernet segment are connected to the same segment. Therefore, all signals are received by all devices. When devices send signals at the same time, a collision occurs. A scheme is needed to detect and compensate for collisions.

- **Collision domain**—A group of devices connected to the same physical medium so that if two devices access the medium at the same time, a collision results. This is a Layer 1 domain.
- **Broadcast domain**—A group of devices on the network that receive one anothers' broadcast messages. This is a Layer 2 domain.

- **Ethernet hubs**—Devices that allow the concentration of many devices into a single segment. They have the following characteristics:
  - — Physical layer devices.
  - — Do not manipulate or view traffic.
  - — Do not create separate collision domains.
  - — Use carrier sense multiple access collision detect (CSMA/CD). When a collision occurs, both stations resend the signal after a random period. Collisions increase with the number of stations.
  - — Regenerate the signal, allowing traffic to travel longer distances.

Data Link Layer Functions

- Perform physical addressing.
- Provide support for connection-oriented and connectionless services.
- Provide for frame sequencing and flow control.

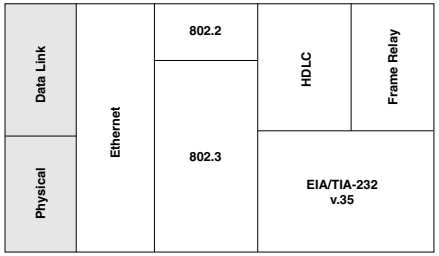

Two sublayers perform the data link functions:

Media Access Control (MAC) sublayer (802.3)—Responsible for how data is sent over the wire. The MAC address is a 48-bit address expressed as 12 hex digits. MAC defines the following:

- Physical addressing
- Network topology
- Line discipline
- Error notification
- Orderly delivery of frames
- Optional flow control

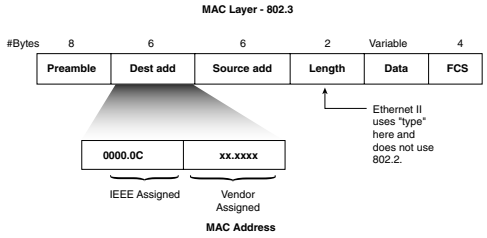

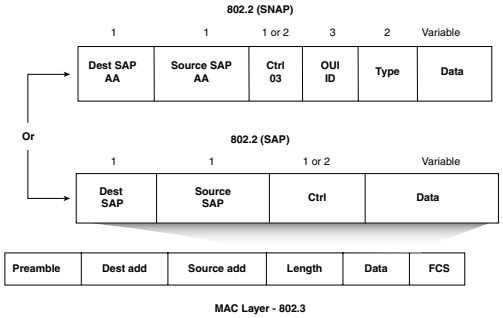

Logical Link Control (LLC) sublayer (802.2)—Responsible for identifying and encapsulating different protocol types. There are two types of LLC frames: Service Access Point (SAP) and Subnetwork Access Protocol (SNAP).

Data Link Layer Devices

Bridges and Layer 2 switches function at the data link layer. Hardware ASICs allow switches to operate at gigabit speeds, whereas bridges make decisions based on software rules, which takes much longer. When a bridge or switch receives a frame, it processes the frame as follows:

- If the destination device is on the same segment as the originating frame, the bridge blocks the frame from going out other ports. This is known as *filtering*.
- If the destination device is on a different segment than the originating frame, the bridge forwards the frame to the appropriate segment.

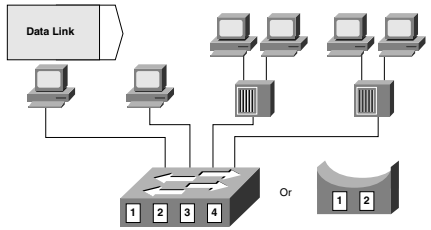

- If the destination device is unknown to the bridge, the bridge forwards the frame to all segments except the one on which it was received. This is called *flooding*.

The purpose of Layer 2 Ethernet devices is to reduce collisions. (Other Layer 2 types are discussed later.) They have the following characteristics:

- Each segment defines a collision domain.
- All devices connected to the same bridge or switch belong to the same broadcast domain.

Network Layer Functions

Network traffic must often span devices that are not locally attached or that belong to separate broadcast domains. Two pieces of information are needed to do this:

- A logical address associated with the source and destination stations
- A path through the network to reach the desired destinations

Router Operation at the Network Layer

Routers operate by gathering and trading data on different networks and selecting the best path to those networks. Routing tables contain the following information:

- **Network addresses**—32-bit addresses.
- **Interface**—The port used to reach a given destination.
- **Metric**—Criteria used to influence path selection when multiple paths exist. Metrics include hops, time, and speed.

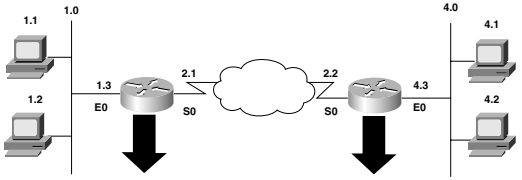

Routing Table

NET	INT	METRIC
1	E0	0
2	S0	0
4	S0	1

Routing Table

NET	INT	METRIC
1	S0	1
2	S0	0
4	E0	0

Transport Layer Functions

A logical connection (session) must be established to connect two devices in a network. The transport layer

- Allows end stations to multiplex multiple upper-layer segments into the same data streams
- Provides reliable data transport (guaranteed delivery) between end stations (on request)

Lower Layers Summary

- The physical layer specifies the media type, connectors, signaling, voltage level, data rates, and distances required to interconnect network devices.
- Hubs allow several end stations to communicate as if they were on the same segment.
- A collision occurs when two stations transmit at the same time.
- Hubs have a single collision domain and a broadcast domain.
- The data link layer determines how data is transported.
- Bridges and Layer 2 switches function at the data link layer.
- All devices connected to a bridge or Layer 2 switch belong to the same broadcast domain.
- All devices connected to a single segment of a Bridge or Layer 2 switch belong to the same collision domain.
- The network layer defines how to transport traffic between devices that are not locally attached.
- The transport layer defines session setup rules between two end stations.
- Routers use routing tables to navigate paths to distant networks.

Assembling and Cabling Cisco Devices

LAN Specifications and Connections

The term *Ethernet* encompasses several LAN implementations. Physical layer implementations vary, and all support various cabling structures. There are three main categories:

- **Ethernet (DIX) and IEEE 802.3**—Operate at 10 Mbps over coaxial cable, UTP, or fiber.
- **100 Mbps Ethernet (Fast Ethernet IEEE 802.3u)**—Operates over UTP or fiber.
- **1000 Mbps Ethernet**—Gigabit Ethernet that operates over fiber.

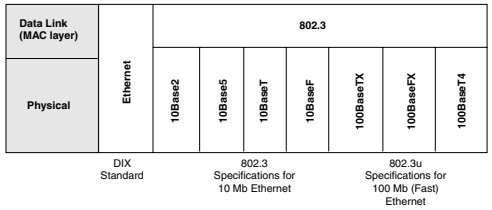

Fast Ethernet can be used throughout the campus environment. The following table gives examples of each campus layer.

	Ethernet 10BaseT Position	Fast Ethernet Position
Access layer	Provides connectivity between the end-user device and the access switch	Gives high-performance PCs and workstations 100 Mbps access to the server.
Distribution layer	Not typically used at this layer	Provides connectivity between access and distribution layers. Provides connectivity from the distribution to core layers. Provides connectivity from the server block to the core layer.
Core layer	Not typically used at this layer	Provides interswitch connectivity.

The following table compares cable and connector specifications. Fast Ethernet requires unshielded twisted-pair (UTP) Category 5 cabling.

	10Base5	10BaseT	100BaseTX	100BaseFX
Medium	50-ohm coaxial (thick)	EIA/TIA Category 3, 4, 5, UTP 2 pair	EIA/TIA Category 5 UTP 2 pair	62.5/125 micron multimode fiber
Maximum segment length	500 meters	100 meters	100 meters	400 meters
Topology	Bus	Star	Star	Point-to-point
Connector	AUI	ISO 8877 (RJ-45)	ISO 8877 (RJ-45)	Duplex media interface connector (MIC) ST

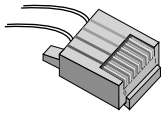

The RJ-45 Connector

Pin	Wire Pair T is Tip R is Ring
1	Pair 2 T2
2	Pair 2 R2
3	Pair 3 T3
4	Pair 1 R1
5	Pair 1 T1
6	Pair 3 R3
7	Pair 4 T4
8	Pair 4 R4

An RJ-45 connector is used with UTP cabling. The two types of connections are straight-through and crossover.

Straight-through cables are typically used to connect different devices, such as switch-to-router connections.

Crossover cables are typically used to connect similar devices, such as switch-to-switch connections. The primary exception to this rule is switch-to-hub connections, which use a crossover cable.

Some device ports are marked with an X. In general, use a straight-through cable when only one of the ports is marked.

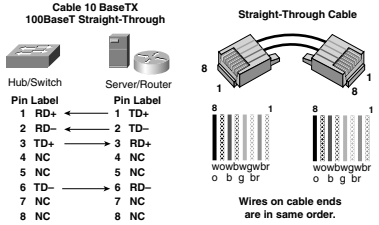

LAN Specifications and Connections Summary

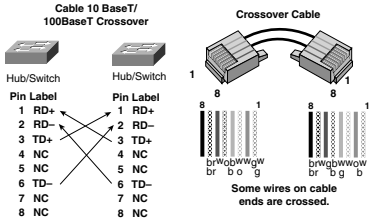

- Ethernet has several LAN specifications, including IEEE 802.3 (10 Mbps), IEEE 802.3u (100 Mbps), and Gigabit Ethernet (1000 Mbps).
- UTP Category 5 is required for Fast Ethernet.
- Straight-through cables are typically used to connect different device types, such as a router and a switch. The exception is a switch-to-hub connection, which requires a crossover cable.
- Crossover cables are typically used to connect similar devices, such as a switch and a switch.

WAN Specifications and Connections

There are several ways to carry traffic across the WAN. The implementation depends on distance, speed, and the type of service required. The speeds of connections vary from 56 Kbps to T1/E1 (1.544/2.048 Mbps). WANs use serial communication for long-distance communication. Cisco routers use a proprietary 60-pin connector. The network end of the cable must match the service hardware.

Cabling Routers for Serial Connectors

When cabling routers, you need to determine whether you need a data terminal equipment (DTE) connector or a data circuit-terminating equipment (DCE) connector:

- **DTE**—The endpoint of the user's device on the WAN link.
- **DCE**—The point where responsibility for delivery data passes into the hands of the SP. The DCE provides clocking and is responsible for forwarding traffic.

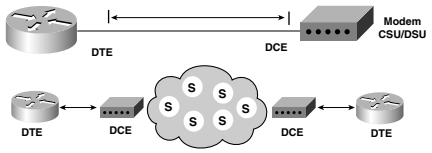

If you connect routers back-to-back, one of the routers will be a DTE, and the other will be a DCE.

Router Ports

Routers can have fixed or modular ports:

- **Fixed ports**—Each port has a port type and number (such as "Ethernet 0").
- **Modular ports**—Each port has a port type, slot number, and port number (such as "serial 1/0").

Configuring Devices

You must establish a connection through a console port in order to configure a Cisco device. Some devices use a rollover cable to connect a console port to a PC. To set up the connection, do the following:

1. Cable the device using a rollover cable. You might need an adapter for the PC.
2. Configure the terminal emulation application with the following COM port settings: 9600 bps, 8 data bits, no parity, 1 stop bit, and no flow control.

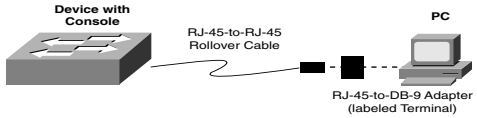

WAN Specifications and Connections Summary

- WANs use serial transmission for long-distance communication.
- Cisco routers use a proprietary 60-pin connector on serial ports.
- A DTE/DCE is the point where the service provider assumes for the WAN. A DCE provides clocking.
- Routers have either fixed or modular ports. The syntax you use to configure each interface depends on the type of port.
- Rollover cables are used to set up a console connection.

Operating and Configuring a Cisco IOS Device

Basic Operation of Cisco IOS Software

Cisco IOS software enables network services in switches and routers. Cisco IOS Software provides the following features:

- Network protocols and functions
- Connectivity
- Security
- Scalability
- Reliability
- Management

The Cisco IOS command-line interface (CLI) can be accessed through a console connection, modem connection, or Telnet session. These connections are called EXEC sessions.

Starting a Switch

When a Catalyst switch is started for the first time, a default configuration is loaded. Three main operations are performed during normal startup:

- A power-on self-test (POST) checks the hardware.
- A startup routine initiates the operating system.
- Software configuration settings are loaded.

Initial Startup Procedure

1. Before you start the switch, verify the following:
   - All network cable connections are secure.
   - A terminal is connected to the console port.
   - A terminal application is selected.
2. Attach the switch to the power source to start the switch (there is no on/off switch).
3. Observe the boot sequence.

LEDs on the front panel of the switch provide information on switch status during startup, normal operation, and fault conditions. Pressing the mode button (shown in the figure) toggles through the LED display modes, which include the following:

- Port status
- BW utilization
- Full-duplex support

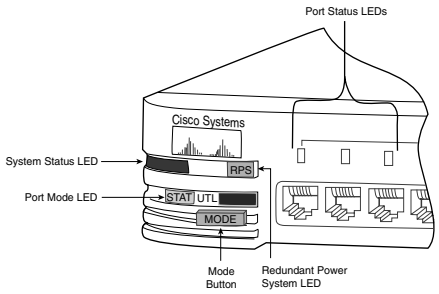

The following table details switch LED status indicators.

Catalyst Switch LED Keys

LED	Status
System LED	**Green**—System is powered and operational. **Amber**—System malfunction.
Redundant power supply	**Green**—Redundant power supply is operational. **Amber—**Redundant power supply is installed but not operational. **Flashing amber**—The internal power supply and redundant power supply have power, and the internal power supply is powering the switch.
Port status (STAT LED on)	**Green**—Link is present. **Flashing green**—Activity. **Alternating green and amber**—Link fault. **Amber**—Port is not forwarding.
Bandwidth utilization (UTL LED on)	**One to eight LEDs on**—0.1 to less than 6 Mbps. **Nine to 16 LEDs on**—6 to less than 120 Mbps. **17 to 24 LEDs on**—120 to 280 Mbps.
Full-duplex (FDUP LED on)	**Green**—Ports are configured in full-duplex mode. **Off**—Ports are half-duplex.

Getting Help

Several commands built into the IOS software provide help when you're entering configuration commands:

- **?**—Displays a list of commonly used commands.
- **More**—Appears at the bottom of the screen when more information exists. Display the next screen by pressing the Spacebar. Display the next line by pressing the Return key. Press any other key to return to the user-mode prompt.
- **s?**—Lists all commands that start with s.
- **show ?**—Lists all variants of the **show** command.
- **show running-configuration**—Displays the currently active configuration in memory, including any changes made in the session that have not yet been saved.

- **show config**—Displays the last saved configuration.
- **show version**—Displays information about the system hardware and software.
- **show interfaces**—Displays information on connections and ports that connect with other devices.

Starting a Switch Summary

- The Catalyst status LEDs are generally green when the switch is functioning and amber when there is a malfunction.
- Port LEDs are green during the POST. The power LED remains green when the test is complete. All other LEDs go off after the test completes unless there is a malfunction.
- After a successful POST, the Menu Console logon screen appears. From here, you can enter three different modes: menu (M), command-line (K), or IP configuration (I).
- The CLI has several help commands, including **?** and **show**.

Starting a Router

When a Cisco router is started for the first time, it does not have an initial configuration. The router prompts the user for a minimum of details. This basic setup is not intended for entering complex configurations or protocol features. The **setup** command gives you the following options:

- Go to the EXEC prompt without saving the created configuration
- Go back to the beginning of setup without saving the created configuration
- Accept the created configuration, save it to NVRAM, and exit to EXEC mode

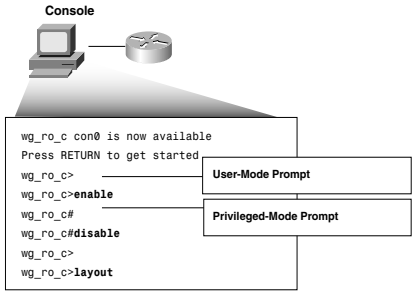

Default answers appear in square brackets ([ ]). You can accept the defaults by pressing the Return key. At the first setup prompt, you can enter **no** to discontinue setup. You can abort the setup process at any time by pressing Ctrl-C.

Access Levels

User EXEC level provides a limited number of basic commands.

Privileged EXEC (enable mode) level gives you access to all router commands. This level can be password-protected. The **enable** command gives you access to this mode. (**disable** takes you back to user mode.)

Console Error Messages

When you enter an incorrect command, you receive one of the following messages:

Error Message	Meaning	How to Get Help
% Ambiguous command: show con	Not enough characters were entered to define a specific command.	Reenter the command followed by a question mark (?) with no space between the command and the question mark.
% Incomplete command	Keywords or values are missing.	Reenter the command followed by a question mark with a space between the command and the question mark.
% Invalid input detected at caret marker	The command was entered incorrectly. The caret marks the point of the error.	Enter a question mark to display all the commands or parameters that are available in this mode.

History Buffer

The command history lets you review previously entered commands. This buffer defaults to ten lines, but you can configure it to a maximum of 256 lines using the **history size** command:

- **terminal history size** *lines*—Sets the session command buffer size
- **history size** *line*—Sets the buffer size permanently
- **show history**—Shows the command buffer contents

CLI Editing Sequences

The Cisco IOS Software gives you shortcuts to speed the editing process.

Command	Action
Ctrl-A	Moves the cursor to the beginning of the line
Ctrl-E	Moves the cursor to the end of the line
Esc-B	Moves the cursor back one word

Command	Action
Esc-F	Moves the cursor forward one character
Ctrl-B	Moves the cursor back one character
Ctrl-F	Moves the cursor forward one word
Ctrl-D	Deletes a single character
Backspace	Removes one character to the left of the cursor
Ctrl-R	Redisplays a line
Ctrl-U	Erases a line
Ctrl-W	Erases a word
Ctrl-Z	Ends configuration mode and returns to EXEC mode
Tab	Completes a partially entered (unambiguous) command
Ctrl-P or up arrow	Recalls commands, beginning with the most recent
Ctrl-N or down arrow	Returns the more recent commands in the buffer

Starting a Router Summary

- The startup configuration routine option appears when no valid configuration exists in NVRAM.
- You can access the setup configuration dialog by entering the **setup** command in privileged mode.
- The **?** command displays the available commands in a given mode.
- The enhanced editing mode includes a set of keyboard functions to simplify using the CLI.
- The command history feature lets you see a list of previously entered commands.

Configuring the Router

From privileged EXEC mode, the **configure terminal** command provides access to global configuration mode. From global configuration mode, you can access specific configuration modes, such as the following:

- **Interface**—Configures operations on a per-interface basis
- **Subinterface**—Configures multiple virtual interfaces
- **Controller**—Supports commands that configure controllers (such as E1 and T1)
- **Line**—Configures the operation of a terminal line
- **Router**—Configures IP routing protocols
- **IPX-router**—Configures the Novell network layer protocol

Assigning a Router Name Example

The **hostname** command can name a router:

```
>enable
#configure terminal
(config)#hostname Router
Router(config)
```

Configuring a Serial Interface Example

```
Router#configure terminal
Router(config)#interface s1
Router(config-if)#clock rate 64000
Router(config-if)#bandwidth 64
Router#show interface serial 1
```

Notes:

- Unambiguous abbreviations of commands are allowed.
- Abbreviations of delimiters are not allowed. For example, a clock rate of 64,000 cannot be abbreviated to 64.
- The **bandwidth** command overrides the default bandwidth (1.544 Mbps). The bandwidth entered has no effect on the line's actual speed.

Major Command/Subcommand Relationship

Commands that indicate a process or interface that will be configured are called *major commands*. Major commands cause the CLI to enter a specific configuration mode.

Major commands have no effect unless they are immediately followed by a subcommand that supplies the configuration entry.

```
Router(config)
  #interface serial 0
Router(config-if)
  #shutdown

Router(config)
  #router rip
Router(config-router)
  #network 10.0.0.0
```

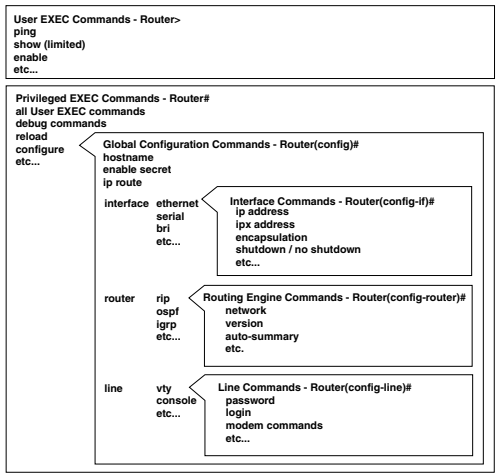

Configuring Router Password Examples

```
Router(config)#line
  console 0
Router(config-line)
  #login
Router(config-line)
  #password homer

Router(config)#line
  vty 0 4
Router(config-line)
  #login
Router(config-line)
  #password bart
```

The numbers 0 to 4 in the **line vty** command specify the number of Telnet sessions allowed in the router. You can also set up a different password for each line by using the **line vty port number** command.

```
Router(config)#enable password apu
Router(config)#enable secret flanders
Router(config)#service password-encryption
```

The **no enable** command disables the privileged EXEC mode password.

The **no enable secret** command disables the encrypted password.

Note: When the enable secret password is set, it is used instead of the enable password.

Configuring the Router Summary

- Entering the **configure terminal** command from enable mode places you in global configuration mode. From this mode, you have access to the interface, subinterface, controller, line, router, and IPX-router configuration modes.
- You must save your running configuration to NVRAM with the **copy running-config startup-config** command. Failing to save your configuration to NVRAM causes your configurations to be lost if your router is reloaded.
- Router security is achieved by password-protecting various access modes.
- Interface type and numbers must be defined when the **interface** command is used.
- Use the **show interface** command to verify configuration changes.

Managing Your Network Environment

Discovering Neighbors with CDP

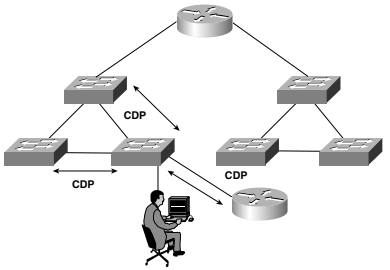

CDP is a proprietary tool that enables access to protocol and address information on directly connected devices. CDP runs over the data link layer, allowing different network-layer protocols (such as IP and IPX) to learn about each other. CDP runs over all LANs, Frame Relay, ATM, and other WANs employing SNAP encapsulation. CDP starts up by default on bootup and sends updates every 60 seconds.

- **show cdp**—Allows you to view CDP output.
- **cdp enable**—Enables CDP on an interface. **no cdp enable** disables.
- **cdp run**—Allows other CDP devices to get information about your device.
- **no cdp run**—Prevents other CDP devices from getting information about your device.
- **show cdp neighbors**—Displays the CDP updates received on the local interfaces.
- **show cdp neighbors detail**—Displays updates received on the local interfaces. This command displays the same information as the **show cdp entry** * command.
- **show cdp entry**—Displays information about neighboring devices.
- **show cdp traffic**—Displays information about interface traffic.
- **show cdp interface**—Displays information about interface status and configuration.

Discovering Neighbors with CDP Summary

- CDP gathers information on directly connected devices.
- CDP passes packets of information between neighboring devices.
- The **show cdp neighbors** command yields the following information for adjacent devices: attached interfaces, hardware platform, and remote port ID.
- The **show cdp entry** * command yields some Layer 3 protocol information (such as IP addresses).

Getting Information About Remote Devices

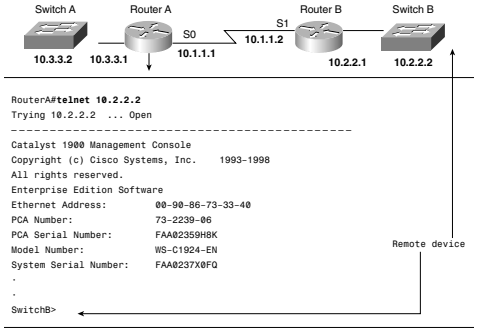

Telnet is an underlying TCP/IP protocol for accessing remote computers. It allows connections and remote console sessions from one device to one or more other remote devices.

Telnet Procedure

To establish a Telnet session, use the **telnet** or **connect** commands. A router's IP address and host name can be used as delimiters.

```
RouterA#telnet 10.2.2.2
RouterB#connect RouterA
RouterA#show sessions
```

Note: **show sessions** displays a list of connected hosts.

Suspending and Resuming Sessions

Press Ctrl-Shift-6 and then press x to suspend the current session.

Press Enter or enter **resume** to resume the last active session.

resume session # reconnects you to a specific session. The **show session** command finds the session number.

Ping/Trace

You can verify connectivity using the **ping** command. In addition to confirming connectivity, **ping** tells you the minimum, average, and maximum times for packets making the roundtrip to the target system and back. You can assess the path's reliability using this command:

```
Router#ping 10.1.1.10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.1.10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/4 ms
```

The **trace** command can be used to view the actual routes that packets take between devices:

```
Router#trace 10.1.1.10

Type escape sequence to abort.
Tracing the route to 10.1.1.10

4 msec 4 msec 4 msec
Router#
```

Getting Information About Remote Devices Summary

- Telnet allows remote connections to distant devices.
- You open a Telnet session by entering the **telnet** or **connect** command, followed by the target device's IP address or host name.
- The **show sessions** command displays a list of connected hosts, their IP addresses, their byte counts, the idle time, and the session name.
- Use the **show user** command to list all active Telnet sessions.
- To reestablish a suspended Telnet session, press the Enter key, use the **resume** command (for the most recent session), or use the **resume session number** command. (Use **show session** to get session numbers.)
- The **ping** and **trace** commands can be used to obtain information about network devices and to check for connectivity.

Router Boot Sequence and Verification

When a router boots up, it goes through the following sequence:

1. The router checks its hardware with a power-on self-test (POST).
2. The router loads a bootstrap code.
3. The Cisco IOS Software is located and loaded using the information in the bootstrap code.
4. The configuration is located and loaded.

When this sequence is complete, the router is ready for normal operation.

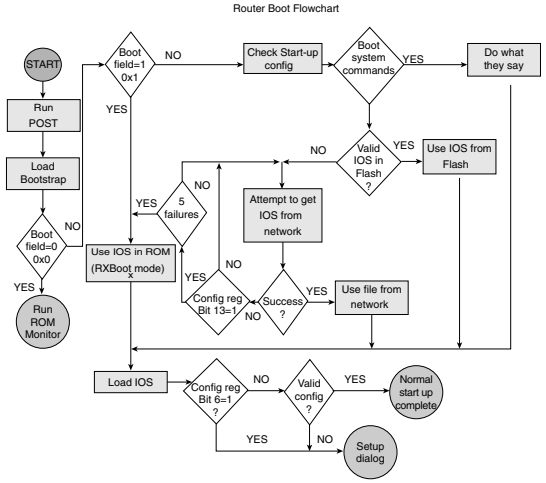

Router Components

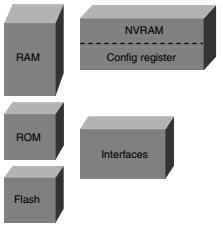

The major router components are as follows:

- **RAM (random-access memory)**—Contains key Cisco IOS Software and data structures.
- **ROM (read-only memory)**—Contains startup microcode.
- **Flash memory**—Flash contains the Cisco IOS Software image. Some routers run the Cisco IOS image directly from Flash and do not need to transfer it to RAM.
- **NVRAM (nonvolatile RAM)**—Stores the configuration. Uses a battery when power is removed.
- **Config reg**—Controls the boot-up method.
- **Interfaces**—Physical connections can include Token Ring, FDDI, and so on.

Altering the Configuration Register

Before changing the configuration register, use the **show version** command to determine the current setting. The last line contains the register value. Changing the value changes the location of the Cisco IOS load. The **reload** command must be used for the new configuration to be set. The register value is checked only during the boot process.

Configuration Register Boot Field Value	Meaning
0x0	Use ROM monitor mode (manually boot using the **b** command).
0x1	Automatically boot from ROM (provides a Cisco IOS subset).
0x2 to 0xF	Examine NVRAM for boot system commands (0x2 is the default if the router has Flash).

When the Cisco IOS Software is loaded, the router must be configured. Configurations in NVRAM are executed. If one does not exist in NVRAM, the router initiates an auto-install or setup utility. The auto-install routine downloads the config file from a TFTP server.

There are other sources of configurations. The Cisco IOS **copy** commands are used to move configurations from one component or device to another. The syntax is **copy object <source> <dest>**. Here's an example:

```
copy running-config
  startup-config
```

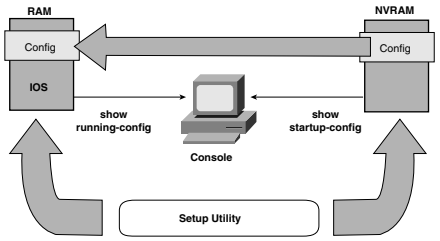

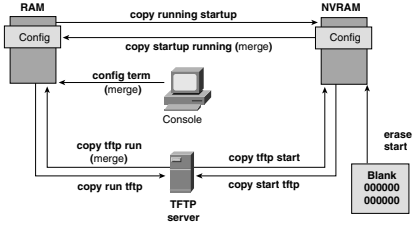

Note: When a configuration is copied into RAM, it merges with the existing configuration in RAM. It does not overwrite the existing configuration.

The **show running-config** and **show startup-config** commands are useful troubleshooting aids. These commands allow you to view the current configuration in RAM or the startup configuration commands in NVRAM.

In NVRAM:

```
wg_rp_c#show startup-config
Using 1359 out of 32762 bytes
!
version 12.0
!
        --More--
```

You know that you are looking at the startup configuration file when you see a message at the top telling you how much nonvolatile memory has been used.

In RAM:

```
wg_ro_c#show running-config
Building configuration...
Current configuration:
!
version 12.0
!
        --More--
```

You know that you are looking at the current configuration file when you see the words "Current configuration" at the top of the display.

Key Feature of IFS

The Cisco IOS File System (IFS) feature provides an interface to the router file systems. The universal resource locator (URL) convention allows you to specify files on network devices. Here are the URL prefixes for Cisco network devices:

- **Bootflash**—Boot Flash memory
- **Flash**—Available on all platforms
- **Flh**—Flash load helper log files
- **ftp**—File Transfer Protocol network server
- **nvram**—NVRAM
- **rcp**—Remote copy protocol network server
- **slot0**—First PCMCIA Flash memory card
- **slot1**—Second PCMCIA Flash memory card
- **System**—Contains the system memory and the running configuration
- **tftp**—Trivial File Transfer Protocol (TFTP) network server

How to Manage Cisco IOS Images

It is always prudent to retain a backup copy of your Cisco IOS Software image in case your router software becomes corrupted. Here's a Cisco IOS upgrade example:

```
wg_ro_a#show flash
wg_ro_a#copy flash tftp
wg_ro_a#copy tftp flash
```

When using the **copy flash** command, you must enter the IP address of the remote host and the name of the source and destination system image file. The router prompts you for this information. If no free Flash memory space is available, or if the Flash memory has never been written to, the erase routine is required.

Router Boot Sequence and Verification Summary

- The major components of the router are RAM, ROM, Flash memory, NVRAM, the configuration register, and the interfaces.
- The four major areas of microcode contained in ROM are bootstrap code, POST code, ROM monitor, and a mini Cisco IOS Software.
- The router configuration can come from NVRAM, a terminal, or a TFTP server.
- You can back up your software image on the network server by using the **copy flash [location]** command.

Catalyst Switch Operations

Basic Layer 2 Switching (Bridging) Functions

Ethernet switching operates at OSI Layer 2, creating dedicated network segments and interconnecting segments. Layer 2 switches have three main functions:

- **MAC address learning**—A Layer 2 switch learns the MAC addresses of devices attached to each of its ports. The addresses are stored in a bridge forwarding database.
- **Forwarding and filtering**—Switches determine which port a frame must be sent out to reach its destination. If the address is known, the frame is sent only on that port; if the address is unknown, the frame is flooded to all ports except the one from which it originated.
- **Loop avoidance**—When the switched network has redundant loops, the switch can prevent duplicate frames from traveling over multiple paths.

Bridging and Switching Comparison

Bridging	Switching
Software-based	Hardware- (ASIC) based
One spanning tree instance per bridge	Many spanning tree instances per switch
Usually up to 16 ports per bridge	More ports on a switch

Frame Transmission Modes

There are three primary frame-switching modes:

- **Cut-through**—The switch checks the destination address and immediately begins forwarding the frame. This can decrease latency.
- **Store and forward**—The switch waits to receive the entire frame before forwarding. The entire frame is read, and a cyclic redundancy check (CRC) is performed. If the CRC is bad, the frame is discarded. Latency increases as a function of frame length.
- **Fragment-free (modified cut-through)**—The switch reads the first 64 bytes before forwarding the frame. 64 bytes is the minimum number of bytes necessary to detect and filter out collision frames. This is the default mode for Catalyst 1900.

How Switches Learn Addresses

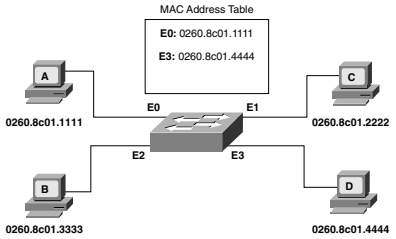

A switch uses its bridge forwarding table (called a MAC table in Catalyst) address table when forwarding frames to devices. With an empty bridge forwarding table, the switch must flood frames to all ports other than the one it arrived on. This is the least-efficient way to transmit data.

Initially, the switch MAC address table is empty. Then Station A with the MAC address sends a frame to station C. When the switch receives this frame, it does the following:

- Because the MAC table is empty, the switch must flood the frame to all other ports (except E0, the frame origin).
- The switch notes the source address of the originating device and associates it with port E0 in its MAC address table entry. Note that the table uses the source address to populate the table, not the destination address.

The switch continues to learn addresses in this manner, continually updating the table. As the MAC table becomes more complete, the switching becomes more efficient, because frames are filtered to specific ports rather than being flooded out all ports.

Broadcast and Multicast Frames

Broadcast and multicast frames are flooded to all ports other than the originating port. Broadcast and multicast addresses never appear as a frame's source address, so the switch does not learn these addresses.

Basic Layer 2 Switching (Bridging) Functions Summary

- Ethernet switches are Layer 2 devices that increase a network's available bandwidth by creating separate network segments.
- Switches have three modes of frame transmission:
  - **Cut-through**—Only the destination address is checked before the frame is forwarded.
  - **Store and forward**—The entire frame is checked before being forwarded.
  - **Fragment-free**—Only the first 64 bytes are checked before forwarding.
- Switches learn, store, and use MAC addresses to determine where a frame should be transmitted.
- A frame is forwarded to a specific port only when the destination address is known. Otherwise, it is flooded out all ports other than the one it was received on.

Redundant Topology Overview

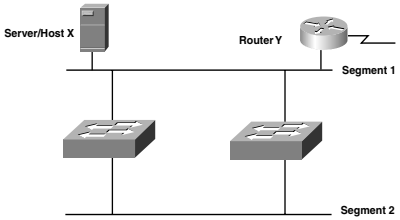

A *redundant topology* has multiple connections to switches or other devices. Redundancy ensures that a single point of failure will not cause the entire switched network to fail. However, redundancy can cause problems in a network, including broadcast storms, multiple copies of frames, and MAC address table instability.

Broadcast Storms

The flooding of broadcast frames can cause a broadcast storm (indefinite flooding of frames) unless there is a mechanism in place to prevent it.

An example of a broadcast storm is shown in the figure and is described here:

1. Host X sends a broadcast frame, which is received by switch A.
2. Switch A checks the destination and floods it to the bottom Ethernet link, segment 2.
3. Switch B receives the frame on the bottom port and transmits a copy to the top segment.
4. Because the original frame arrives at switch B through the top segment, switch B transmits the frame a second time. The frame now travels continuously in both directions.

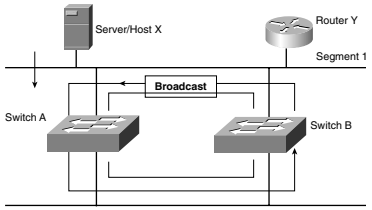

Multiple Frame Transmissions

Most protocols cannot correctly handle duplicate transmissions. Protocols that use sequence numbering assume that the sequence has recycled. Other protocols process the duplicate frame with unpredictable results. Multiple frame transmissions occur as follows:

1. Host X sends a frame to Router Y. One copy is received over the direct Ethernet connection, segment 1. Switch A also receives a copy.
2. Switch A checks the destination address. If the switch does not find an entry in the MAC address table for Router Y, it floods the frame on all ports except the originating port.
3. Switch B receives the frame on segment 2. Switch B then forwards the frame to segment 1.

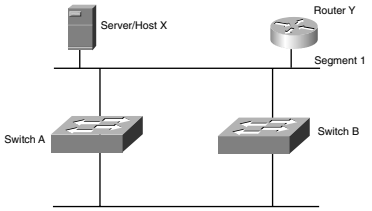

Note: Router Y has now received two copies of the same frame.

Database Instability

Database instability occurs when a switch receives the same frame on different ports. The following example shows how this occurs:

1. Host X sends a frame to Router Y. When the frame arrives at switch A and switch B, they both learn the MAC address for host X and associate it with 0.
2. The frame is flooded out port 1 of each switch (assuming that Router Y's address is unknown).

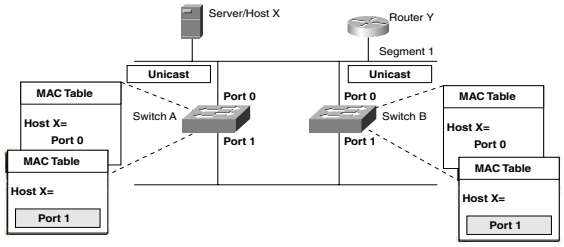

3. Switch A and switch B receive the frame on port 1 and incorrectly associate host X's MAC address with that port.
4. This process repeats indefinitely.

Multiple Loops

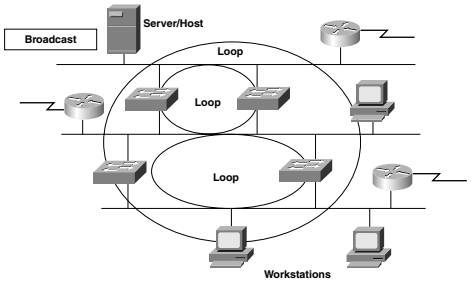

Multiple loops can occur in large switched networks. When multiple loops are present, a broadcast storm clogs the network with useless traffic. Packet switching is adversely affected in this case and might not work at all. Layer 2 cannot prevent or correct broadcast storms.

Redundant Topology Summary

- A broadcast storm occurs when broadcast messages propagate endlessly throughout a switched network.
- Multiple transmissions of the same message cause errors in most protocols.
- A switch's MAC address table becomes unstable when the switch receives the same frame on different ports.
- Layer 2 devices cannot recognize or correct looping traffic without help.

Spanning-Tree Protocol

The *Spanning-Tree Protocol* prevents looping traffic in a redundant switched network by blocking traffic on the redundant links. If the main link goes down, the spanning tree activates the standby path. Spanning-Tree Protocol operation is transparent to end stations.

The Spanning-Tree Protocol was developed by DEC and was revised in the IEEE 802.1d specification. The two algorithms are incompatible. Catalyst switches use the IEEE 802.1d Spanning-Tree Protocol.

Spanning Tree Operation

Link Speed	Cost (Reratify IEEE Spec)	Cost (Previous IEEE Spec)
10 Gbps	2	1
1 Gbps	4	1
100 Mbps	19	10
10 Mbps	100	100

Spanning-Tree Protocol assigns roles to switches and ports so that there is only one path through the switch network at any given time. This is accomplished by assigning a single root bridge, root ports for nonroot bridges, and a single designated port for each network segment. On the root bridge, all ports are designated ports.

On the root bridge, all ports are set to the forwarding state. For the nonroot bridge, the root port is set to the forwarding state. (The switch might also have designated ports set to the forwarding state). The port with the lowest-cost path to the root bridge is chosen as the root port.

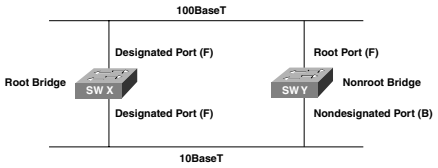

One designated port is assigned on each segment. The bridge with the lowest-cost path to the root bridge is the designated port.

Nondesignated ports are set to the blocking state (which does not forward any traffic).

Selecting the Root Bridge

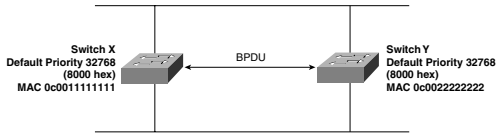

Switches running the Spanning-Tree Protocol exchange information at regular intervals using a frame called the *bridge protocol data unit* (BPDU). Each bridge has a unique bridge ID. The bridge ID contains the bridge MAC address and a priority number. The midrange value of 32768 is the default priority. The bridge with the lowest bridge ID is selected as the root bridge. When switches have the same priority, the one with the lowest MAC address is the root bridge. In the figure, Switch X is the root bridge.

Port States

Frames take a finite amount of time to travel or propagate through the network. This delay is known as *propagation delay.* When a link goes down, spanning tree activates previously blocked links. This information is sent throughout the network, but not all switches receive this information at the same time. To prevent temporary loops, switches wait until the entire network is updated before they set any ports to the forwarding state. Each switch port in a network running the Spanning-Tree Protocol is in one of the following states:

- Blocking
- Listening
- Learning
- Forwarding

The forward delay is the time it takes for a port to go to a higher state. It usually takes 50 seconds for a port to go from the blocking state to the forwarding state (20 max age + 15 listening + 15 learning), but the timers can be adjusted.

Spanning-Tree Recalculation

When a link fails, the network topology must change. Connectivity is reestablished by placing key blocked ports in the forwarding state.

1 Edit not clear in illustration below

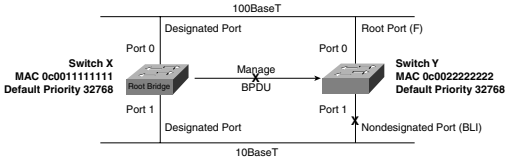

In the figure, if switch X fails, switch Y does not receive the BPDU. If the BPDU is not received before the max age timer expires, spanning tree begins recalculating the network. In the figure, switch Y is now the root bridge. If switch X comes back up, spanning tree recalculates the network, and switch X is again the root bridge.

Time to Converge

A network is said to have converged when all ports in a switched network are in either the blocked or forwarding state after a topology change.

Spanning-Tree Protocol Summary

- The Spanning-Tree Protocol prevents loops in a redundant network.
- Spanning-Tree Protocol assigns a root bridge, root ports for nonroot bridges, and designated port segments. In a converged network, ports are either in forwarding or blocking state.
- BPDUs are exchanged every two seconds. The bridge ID is made up of the MAC address and priority. The bridge with the lowest bridge ID is the root bridge.
- The four port states are blocking, listening, learning, and forwarding.
- When a link fails, spanning tree adjusts the network topology to ensure connectivity.

Configuring the Catalyst 1900 Switch

An IP address must be assigned to a switch to use Telnet or Simple Network Management Protocol (SNMP).

A 32-bit subnet mask denotes which bits in the IP address correspond to the host and network portions of the address.

The default gateway is used when the switch must send traffic to a different IP network. The default gateway is a Layer 3 device (router) that can access other networks.

Configuring the IP Address

Before configuring the switch, you must identify the IP address, subnet mask, and default gateway on the switch:

```
RouterA(config)#ip address 10.1.5.22 255.255.255.0
RouterA (config)#ip default-gateway 10.1.5.44
```

Use the **no ip address** command to reset the IP address to the factory default of 0.0.0.0. Use the **no ip default-gateway** command to delete a configured default gateway and set the gateway address to the default value of 0.0.0.0.

The IP address, subnet mask, and default gateway settings can be viewed with the **show ip** command.

Duplexing

Duplexing is a mode of communication in which both ends can send and receive information. With full duplex, bidirectional communication can occur at the same time. Half duplex is also bidirectional, but signals can flow in only one direction at a time.

Half duplex:

- CSMA/CD susceptible to collisions
- Multipoint attachments
- Can connect with both half-duplex and full-duplex devices
- Efficiency is typically rated at 50 to 60 percent
- Nodes sharing their connection to a switch port must be in half-duplex mode

Full duplex:

- Can send and receive data at the same time
- Collision-free
- Point-to-point connection only
- Uses a dedicated switched port with separate circuits
- Efficiency is rated at 100 percent in both directions
- Both ends must be configured to run in full-duplex mode

Duplex Interface Configuration

The Catalyst 1900 can autonegotiate the duplex connection. This mode is enabled when both speed and duplex flags are set to auto. The **show interfaces** command shows the current settings.

```
duplex {auto | full | full-flow-control | half}
```

- **duplex auto**—Autonegotiation of duplex mode
- **duplex full-flow-control**—Full-duplex mode with flow control

Managing MAC Addresses

MAC address tables contain three types of addresses:

- Dynamic addresses are learned by the switch and then are dropped when they are not in use.
- Permanent and static addresses are assigned by an administrator.

MAC Address Configuration

The **mac-address-table** global configuration command is used to associate a MAC address with a particular switched port interface. The syntax for the **mac-address-table** command is

```
mac-address-table {permanent, restricted static} {mac-address type module/
port
  (src-if-list)}
```

You verify the MAC address table settings using the **show mac-address-table** command.

Note: The Catalyst 1900 can store a maximum of 1024 MAC addresses in its MAC address table. After the table is full, it floods all new addresses until one of the existing entries gets aged out.

- **mac-address-table permanent**—Sets a permanent MAC address
- **no mac-address-table permanent**—Deletes a permanent MAC address
- **mac-address-table restricted static**—Sets a restricted static address to an interface
- **no mac-address-table restricted static**—Deletes a restricted static address
- **Mac-address-table src-if-list**—Sets a restricted address to a port

Port References (Catalyst 1900)

Different commands refer to the same ports in different ways:

- The **show running config** output refers to e0/1 as interface Ethernet 0/1.
- The **show spantree output** refers to e0/1 as port Ethernet 0/1.
- The **show vlan-membership** output refers to e0/1 as port 1.

Port Security

The *port security* feature restricts the number of MAC addresses used on a switch or restricts the use of a port to a specified group of users. The number of devices on a secured port can range from one to 132. The MAC addresses are assigned either automatically or by the administrator (assigned statically).

Address violations occur when a secured port receives a source address already assigned to another secured port or when a port exceeds its address table size limit. When a violation occurs, the action can be suspended, ignored, or disabled.

A suspended port is reenabled when a valid address is received. A disabled port must be reenabled manually. If the action is ignored, the switch port remains enabled.

Here is the procedure for configuring the IP address:

```
RouterA(config)#interface e0/1
RouterA(config-if)#port secure max-mac-count 1
RouterA(config-if)#exit
RouterAhost#show mac-address-table security
RouterA(config-if)#exit
RouterA(config)#address-violation ignore
```

The **no port secure** command disables addressing security and sets the maximum number of addresses on the interface to the default (132).

The **show** command yields a list of enabled ports and their security statuses.

The action for an address violation can be suspend, disable, or ignore.

Use the **no address-violation** command to set the switch to its default value (suspend).

Configuring the Catalyst 1900 Switch Summary

- To configure global switch parameters (switch, host name, or IP address), use the **config term** command. To configure a particular port, use the **interface** command while in global configuration mode.
- MAC address tables can be dynamic, permanent, or static.
- Switches are assigned IP addresses for network management purposes.
- A default gateway is used to reach a network that has a different IP address.
- Use the various **show** commands to verify switch configuration.

VLANs

VLAN Operation Overview

The *virtual LAN* (VLAN) allows you to group physically separate users into the same broadcast domain. The use of VLANs improves security, segmentation, and flexibility. The use of VLANs also decreases the cost of arranging users, because no extra cabling is required.

VLAN Characteristics

VLANs allow an administrator to define user groups logically rather than by their physical locations. For example, you can arrange user groups such as accounting, engineering, and finance rather than grouping everyone on the first floor, everyone on the second floor, and so on.

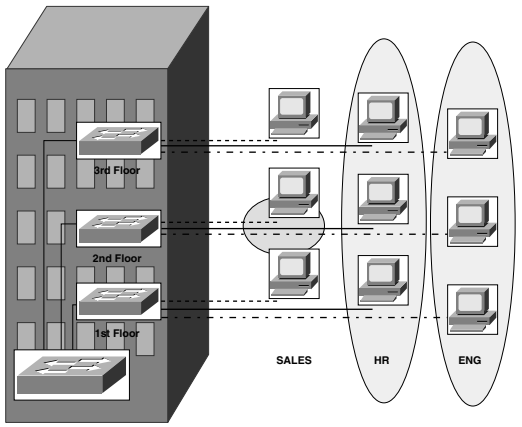

- VLANs define broadcast domains that can span multiple LAN segments.
- VLAN segmentation is not bound by the physical location of users.
- Each switch port can be assigned to only one VLAN.
- Ports not assigned to the same VLAN do not share broadcasts, improving network performance.
- A VLAN can exist on one switch or on multiple switches.
- VLANs can connect across wide-area networks (WANs).

The figure shows a VLAN design. VLANs are defined by user functions rather than locations.

VLAN Operation

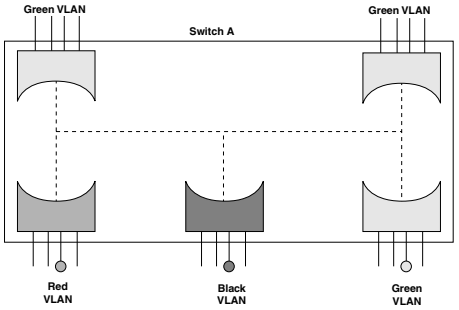

Each VLAN on a switch behaves as if it were a separate physical bridge. The switch forwards packets (including unicasts, multicasts, and broadcasts) only to ports assigned to the same VLAN from which it originated. This reduces on network traffic.

VLANs require a trunk to span multiple switches. Each trunk can carry traffic for multiple VLANs.

VLAN Assignment

A port can be assigned (configured) to a given VLAN. VLAN membership can be designated as either static or dynamic:

- **Static assignment—** The VLAN port is statically configured by an administrator.
- **Dynamic assignment—**The switch uses a VMPS (VLAN Membership Policy Server). The VMPS is a database that maps MAC addresses to VLANs. A port can belong to only one VLAN at a time. Multiple hosts can exist on a single port only if they are all assigned to the same VLAN.

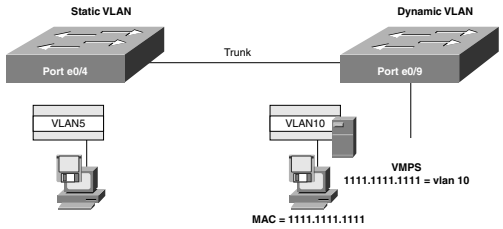

Inter-Switch Link

Inter-Switch Link (ISL) is a Cisco-proprietary protocol designed to carry VLAN traffic between switches. ISL provides point-to-point links in full-duplex or half-duplex mode. ISL is performed with ASICs, which operate at wire speeds and let VLANs span the backbone.

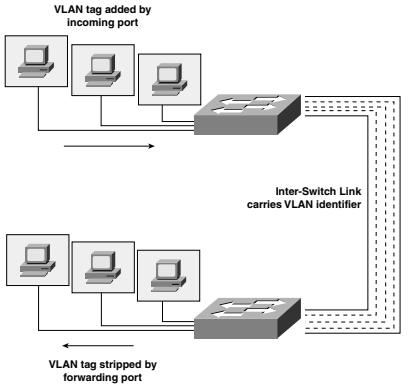

ISL Tagging

ISL frame tagging multiplexes VLAN traffic onto a single physical path. It is used for connections between switches, routers, and network interface cards. A non-ISL-capable device treats ISL-encapsulated Ethernet frames as protocol errors if the frame size exceeds the maximum transmission unit (MTU). ISL tagging is a protocol-independent function that occurs at OSI Layer 2. ISL can maintain redundant links and can load-balance traffic.

ISL Encapsulation

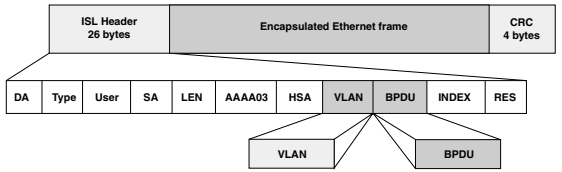

ISL-enabled ports encapsulate each frame with a 26-byte ISL header and a 4-byte CRC. ASICs allow this to occur at wire speed (low latency). The number of VLANs supported depends on the switch. The Catalyst 1900 supports 64 VLANs with a separate spanning-tree instance for each VLAN.

VLAN Operation Summary

- A VLAN is a broadcast domain that can span multiple physical LAN segments.
- VLANs improve performance, flexibility, and security by restricting broadcasts.
- VLANs only forward data to ports assigned to the same VLAN.
- VLAN ports can be assigned either statically or dynamically.
- ISL is a Cisco-proprietary protocol used to share and manage VLAN information across switches.
- ISL trunks encapsulate frames with an ISL header CRC.

Configuring a VLAN

VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that maintains VLAN configuration consistency throughout a common administrative domain by managing VLAN additions, deletions, and name changes across multiple switches. VTP server updates are propagated to all connected switches in the network, which reduces the need for manual configuration (promotes scaling) and minimizes the risk of errors caused by duplicate names or incorrect VLAN types.

VTP operates in server, client, or transparent mode. The default is server mode. VLAN updates are not propagated over the network until a management domain name is specified or learned.

VTP Example

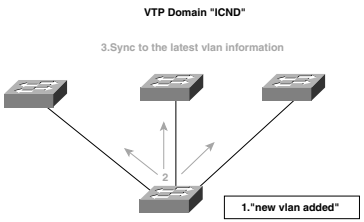

The VTP server notifies all switches in its domain that a new VLAN, named ICND, has been added. The server advertises VLAN configuration information to maintain domain consistency.

How VTP Works

Whenever a change to a VLAN occurs, the VTP server increments its configuration revision number and then advertises the new revision throughout the domain. When a switch receives the advertisement, it overwrites its configuration with the new information if the new revision number is higher than the one it already has.

VTP Advertisements

VTP advertisements are flooded over the factory default VLAN (VLAN1) every five minutes or whenever there is a change. The **delete vtp** command resets the configuration number.

VTP Modes

VTP operates in server, client, or transparent mode. The default is server mode. VLAN configurations are not advertised until a management domain name is specified or learned.

Server Mode	Client Mode	Transparent Mode
Sends and forwards VTP advertisements.	Sends and forwards VTP advertisements.	Forwards VTP advertisements.
Syncs VLAN configuration information with other switches.	Syncs VLAN configuration information with other switches.	Does *not* sync VLAN configuration information with other switches.
Configurations are saved in NVRAM.	Configurations are *not* saved in NVRAM.	Configurations are saved in NVRAM.
Switch can create VLANs.	Cannot create VLANs.	Switch can create VLANs.
Switch can modify VLANs.	Switch cannot modify VLANs.	Switch can modify VLANs.
Switch can delete VLANs.	Cannot delete VLANs.	Switch can delete VLANs.

VTP Pruning

VTP pruning improves bandwidth by keeping unnecessary traffic from flooding the entire domain.

By default, a trunk carries traffic for all VLANs in the VTP management domain.

With VTP pruning enabled, updated traffic from station A is not forwarded to switches

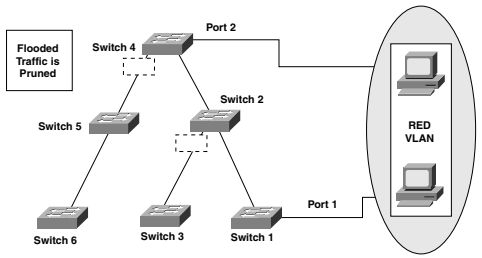

3, 5, and 6, because traffic for the red VLAN has been pruned on the links indicated on switches 2 and 4.

Here is the **vtp** command:

```
vtp [server | transparent] [domain domain-name] [trap {enable | disable}]
  [password password] [pruning {enable | disable}]
```

- *domain-name* can be specified or learned.
- **vtp trap** generates NMP messages.
- *password* can be set for the VTP management domain. The password entered should be the same for all switches in the domain.
- **pruning** propagates the change throughout the domain.

VTP trunk Command

The **trunk** command sets a Fast Ethernet port to trunk mode. This command turns trunking on or off and sets the negotiation state:

```
trunk [on | off | desirable | auto | nonegotiate]
```

- **desirable**—The port turns on trunking if the connected device is in the On, Desirable, or Auto state.
- **auto**—Enables trunking if the connected device is set to On or Desirable.
- nonegotiate—The port is set to the permanent ISL trunk.

Here is the procedure for configuring VTP:

```
RouterA(config)#vtp transparent domain springfield trap enable password
  cisco pruning enable
RouterA(config)#int fa0/26
RouterA(config-if)#trunk on desirable
RouterA(config-if)#exit
RouterA(config)#address-violation {s | d | i}
RouterA(config)#exit
RouterA#show vtp
RouterA#show trunk A
```

On the Catalyst 1900, the two Fast Ethernet ports are interfaces fa0/26 and fa0/27.

Here is the procedure for configuring a VLAN:

```
RouterA#config t
RouterA(config)#vlan 7 name springfield
RouterA(config)#int fa0/26
RouterA(config-if)#vlan-membership static 7
RouterA(config-if)#exit
RouterA(config)#exit
RouterA#show vlan7
RouterA#show vlan-membership
RouterA#show spantree 1
```

Configuring a VLAN Summary

- VTP advertises and synchronizes VLAN configuration information.
- The three VTP modes are server (the default), client, and transparent.
- VTP messages include a configuration revision number. When a switch receives a higher configuration number, it overwrites its configuration with the newly advertised one.
- VTP pruning restricts flooded traffic to some trunk lines.
- VLAN 1 is the default VLAN configuration on the Catalyst 1900 switch.
- To configure a VLAN, you must enable VTP, enable trunking, create a VLAN, and assign that VLAN to a port.

TCP/IP Overview

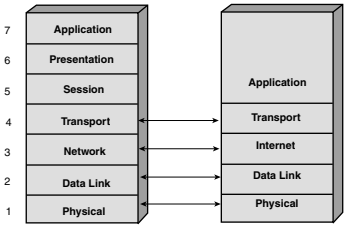

The *Transmission Control Protocol/Internet Protocol* (TCP/IP) suite of protocols is used to communicate across any set of interconnected networks. These protocols, initially developed by Defense Advanced Research Projects Agency (DARPA), are well-suited for communication across both LANs and WANs.

The protocol suite includes Layer 3 and 4 specifications, as well as specifications for higher-layer applications such as e-mail and file transfer.

The TCP/IP protocol stack closely follows the OSI Reference Model. All standard Layer 1 and 2 protocols are supported (called the network interface layer in TCP/IP).

TCP/IP Datagrams

TCP/IP information is sent through datagrams. One message can be broken up into a series of datagrams that must be reassembled at the destination. Three layers are associated with the TCP/IP protocol stack:

- **Application layer**—Specifications exist for e-mail, file transfer, remote login, and other applications. Network management is also supported.
- **Transport layer**—Transport services allow multiple upper-layer applications to use the same data stream. TCP and UDP protocols at this layer provide the following functions:
  - — Flow control (through windowing)
  - — Reliability (through sequence numbers and acknowledgments)

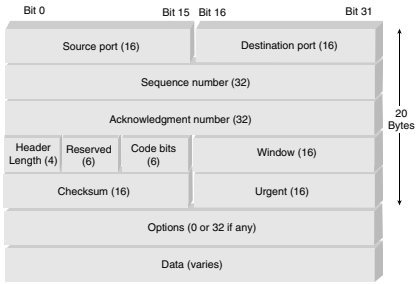

- **Internet layer**—Several protocols operate at the TCP/IP Internet layer:
  - — IP provides connectionless, best-effort routing of datagrams.
  - — ICMP provides control and messaging capabilities.
  - — ARP determines the data link layer address for known IP addresses.
  - — RARP determines network addresses when data link layer addresses are known.

TCP

TCP is a connection-oriented, reliable protocol that breaks messages into segments and reassembles them at the destination station (resending anything not received). TCP also provides a virtual circuit between applications.

UDP

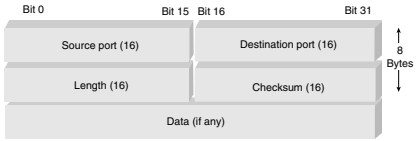

UDP is a connectionless, unreliable protocol used for applications that provide their own error recovery process. It trades reliability for speed. UDP is simple and efficient but unreliable. UDP does not check for segment delivery.

Connection-Oriented Services

A connection-oriented service establishes and maintains a connection during a transmission. The service first establishes a connection and then sends data. After the data transfer is complete, the session is torn down.

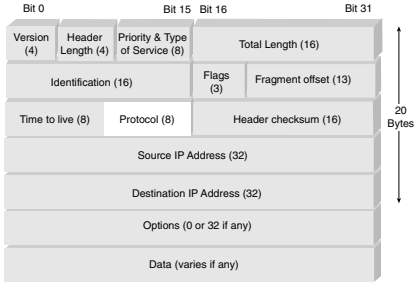

Port Numbers

Both TCP and UDP can send data from multiple upper-layer applications on the same datagram. Port (or socket) numbers are used to keep track of different conversations crossing the network at any given time. Well-known port numbers are controlled by the Internet Assigned Numbers Authority (IANA). For example, Telnet is always defined by port 23. Applications that do not use well-known port numbers have them randomly assigned from a specific range.

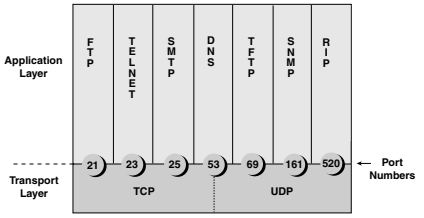

Port Number Ranges

- Numbers below 1024 are considered well-known ports.
- Numbers above 1024 are dynamically assigned ports.
- Vendor-specific applications have reserved ports (usually above 1024).

How TCP Connections Are Established

End stations use control bits called SYN (for synchronize) and Initial Sequence Numbers (ISN) to synchronize during connection establishment.

Three-Way Handshake

The synchronization requires each side to send its own initial sequence number and to receive a confirmation of it in acknowledgment (ACK) from the other side.

1. Host A sends a SYN segment with sequence number 100.
2. Host B sends an ACK and confirms the SYN it received. Host B also sends a SYN. The ACK field in host B now expects to hear sequence 101.
3. Host A sends an ACK verifying the SYN and passes data.

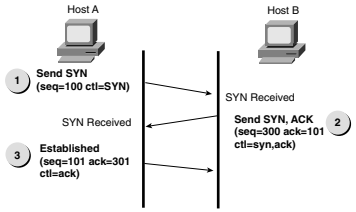

TCP Windowing

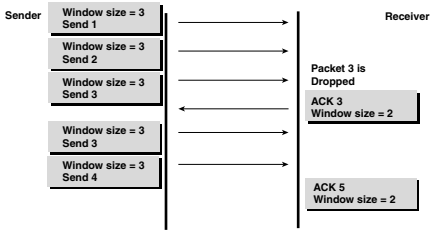

Windowing ensures that one side of a connection is not overwhelmed with data that it cannot process. The window size from one end station tells the other side of the connection how much it can accept at one time. With a window size of 1, each segment must be acknowledged before another segment is sent. This is the least-efficient use of bandwidth.

1. The sender sends three packets before expecting an ACK.
2. The receiver can handle only a window size of 2. So it drops packet 3, specifies 3 as the next packet, and specifies a window size of 2.
3. The sender sends the next two packets but still specifies its window size of 3.
4. The receiver replies by requesting packet 5 and specifying a window size of 2.

TCP Sequence and Acknowledgment Numbers

TCP uses forward reference acknowledgments. Each datagram is numbered so that at the receiving end TCP reassembles the segments into a complete message. If a segment is not acknowledged within a given time period, it is resent.

IP

IP provides connectionless, best-effort delivery routing of datagrams. The protocol field in the header determines the Layer 4 protocol being used (usually TCP or UDP).

Other Internet Layer Protocols

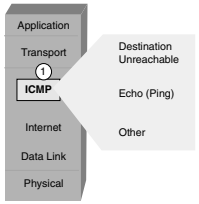

ICMP, ARP, and RARP are three protocols used by the Internet layer to IP. The *Internet Control Message Protocol* (ICMP) is used to send error and control messages. Messages such as destination unreachable, time exceeded, subnet mask request, echo, and others are used by ICMP.

Address Resolution Protocol (ARP) maps a known IP address to a MAC sublayer address. An ARP cache table is checked when looking for a destination address. If the address is not in the table, ARP sends a broadcast looking for the destination station.

Reverse ARP

Reverse Address Resolution Protocol (RARP) maps a known MAC address to an IP address. Dynamic Host Configuration Protocol (DHCP) is a modern implementation of RARP.

TCP/IP Overview Summary

- The TCP/IP protocol suite includes Layer 3 and 4 specifications.
- UDP is connectionless (no acknowledgments). No software checking for segment delivery is done at this layer.
- TCP is a reliable connection-oriented protocol. Data is divided into segments, which are reassembled at the destination. Missing segments are resent.
- Both TCP and UDP use port (or socket) numbers to pass information to the upper layers. A socket is an IP address in conjunction with a port number.
- The three-way handshake is a synchronization process. Sequence numbers and ACK are used to establish connections.

TCP/IP Address Overview

In a TCP/IP environment, each node must have a unique 32-bit logical IP address. Each IP datagram includes the source and destination IP addresses in the header.

Host and Network Address

Each company listed on the Internet is viewed as a single network. This network must be reached before a host within that company can be contacted. A two-part addressing scheme allows the IP address to identify both the network and the host.

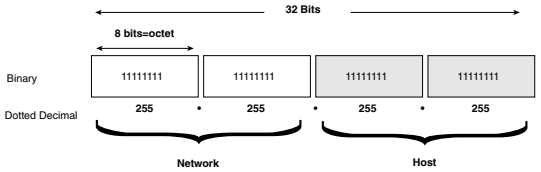

- All the endpoints within a network share a network number.
- The remaining bits identify each host within that network.

IP Address Classes

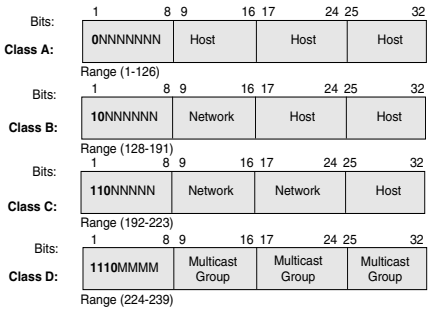

There are five classes of IP: Classes A through E. Classes A, B, and C are the most common. Class A has 8 network bits and 24 host bits. (So there are few Class A networks, but each has many hosts.) Class C addresses allow for many more networks, each with fewer hosts. This scheme was based on the assumption that there would be more small networks than large networks in the world.

Note: The address range for all five classes is shown in the figure. Class D is used for multicast purposes, and Class E addresses are used for research.

Class C Address Breakdown

Number of Bits	Subnet Mask	Subnets	Hosts
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2

The example in the figure shows networks A and B connected by a router. Network B has a Class A address (10.0.0.0). The routing table contains entries for network addresses (not hosts within that network). In the example, 172.16.0.0 and 10.0.0.0 refer to the wires at each end of the router. Network 10.0.0.0 is a special case of Class A networks. It is typically used in private networks.

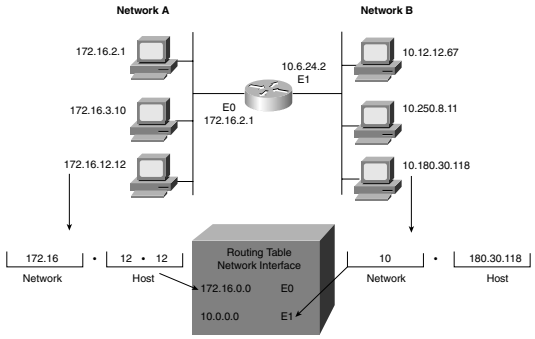

TCP/IP Address Summary

- In a TCP/IP environment, each end station has a 32-bit logical IP address that has a network and host portion.
- The address format is known as dotted-decimal notation. The range is 0.0.0.0 to 255.255.255.255.
- Five address classes are suited to different types of users.
- The total number of available hosts on a network can be derived by using the formula $2^n - 2$, where n is the number of bits in the host portion.

Implementing Subnet Planning

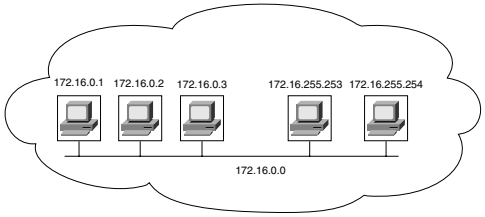

Without subnets, the organization operates as a single network. These flat topologies result in short routing tables but, as the network grows, the use of bandwidth becomes inefficient. (All systems on the network receive all the broadcasts on the network.) Network addressing can be made more efficient by breaking the addresses into smaller segments, or subnets. Subnetting provides additional structure to an addressing scheme without altering the addresses.

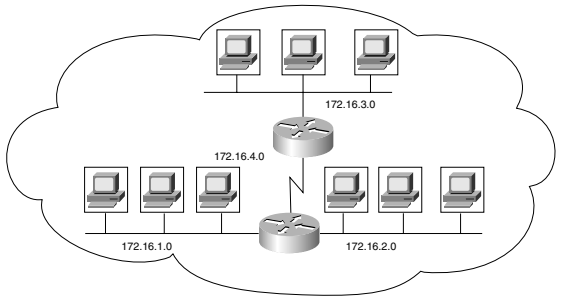

In the figure, the network address 172.16.0.0 is subdivided into four subnets: 172.16.1.0, 172.16.2.0, 172.16.3.0, and 172.16.4.0. If traffic were evenly distributed to each end station, the use of subnetting would reduce the overall traffic seen by each end station by 75 percent.

Subnet Mask

A *subnet mask* is a 32-bit value written as four octets. In the subnet mask, each bit determines how the corresponding bit in the IP address should be interpreted (network, subnet, or host). The subnet mask bits are coded as follows:

- Binary 1 for the network bits
- Binary 1 for the subnet bits
- Binary 0 for the host bits

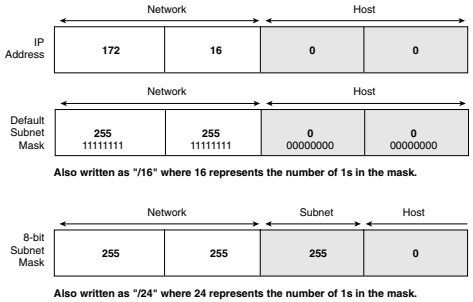

Although dotted decimal is the most common format, the subnet can be represented in several ways:

- **Dotted decimal**—172.16.0.0 255.255.0.0
- **Bit count**—172.16.0.0/16
- **Hexadecimal**—172.16.0.0 0xFFFF0000

The **ip netmask-format** command can be used to specify the format of network masks for the current session. Dotted decimal is the default.

Default Subnet Masks

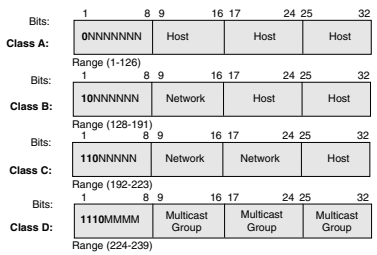

Each address class has a default subnet mask. The default subnet masks only the network portion of the address, the effect of which is no subnetting. With each bit of subnetting beyond the default, you can create $2^n - 2$ subnets. These examples show the effect of adding subnet bits.

Address	Subnet Address	Number of Subnets	Comments
10.5.22.5/8	255.0.0.0	0	This is the default Class A subnet address. The mask includes only the network portion of the address and provides no additional subnets.
10.5.22.5/16	255.255.0.0	254	This Class A subnet address has 16 bits of subnetting, but only the bits in the second octet (those beyond the default) contribute to the subnetting.
155.13.22.11/16	255.255.0.0	0	In this case, 16 bits are also used for subnetting, but because the default for a Class B address is 16 bits, no additional subnets are created.
155.13.10.11/26	255.255.255.192	1022	In this case, there is a total of 26 bits of subnetting, but the Class B address can use only 10 of them to create subnets. The result is the creation of 1024 subnets ($2^{10} - 2$).

How Routers Use Subnet Masks

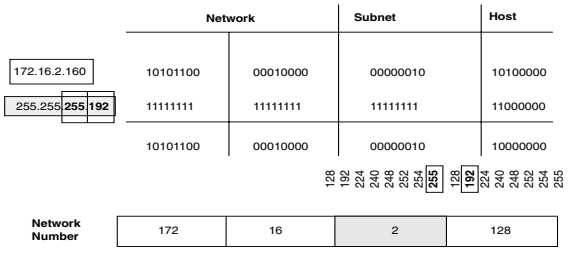

To determine an address's subnet, a router performs a logical AND operation with the IP address and subnet mask. Recall that the host portion of the subnet mask is all 0s. The result of this operation is that the host portion of the address is removed, and the router bases its decision on only the network portion of the address.

In the figure, the host bits are removed, and the network portion of the address is revealed. In this case, a 10-bit subnet address is used, and the network (subnet) number 172.16.2.128 is extracted.

Broadcast Addresses

Broadcast messages are sent to every host on the network. There are three kinds of broadcasts:

- **Directed broadcasts**—You can broadcast to all hosts within a subnet and to all subnets within a network. (170.34.2.255 sends a broadcast to all hosts in the 170.34.2.0 subnet.)
- **Flooded broadcasts (255.255.255.255)**—Local broadcasts within a subnet.
- You can also broadcast messages to all hosts on all subnets within a single network. (170.34.255.255 sends a broadcast to all subnets in the 170.34.0.0 network.)

Identifying Subnet Addresses

Given an IP address and subnet mask, you can identify the subnet address, broadcast address, first usable address, and last usable address using this method:

	174	24	4	176	
174.24.4.176	10101110	00011000	00000100	10110000	Host
255.255.255.192	11111111	11111111	11111111	11000000	Mask
174.24.4.128	10101110	00011000	00000100	10000000	Subnet
174.24.4.191	10101110	00011000	00000100	10111111	Broadcast
174.24.4.129	10101110	00011000	00000100	10000001	First
174.24.4.190	10101110	00011000	00000100	10111110	Last

1. Write down the 32-bit address. Directly below that, write down the subnet mask.
2. Draw a vertical line just after the last 1 bit in the subnet mask.
3. Copy the portion of the IP address to the left of the line. Place all 0s for the remaining free spaces to the right. This is the subnet number.
4. Copy the portion of the IP address to the left of the line. Place all 1s for the remaining free spaces to the right. This is the broadcast address.
5. Copy the portion of the IP address to the left of the line. Place all 0s in the remaining free spaces until you reach the last free space. Place a 1 in that free space. This is your first usable address.
6. Copy the portion of the IP address to the left of the line. Place all 1s in the remaining free spaces until you reach the last free space. Place a 0 in that free space. This is your last usable address.

How to Implement Subnet Planning

Subnetting decisions should always be based on growth estimates rather than current needs.

To plan a subnet, follow these steps:

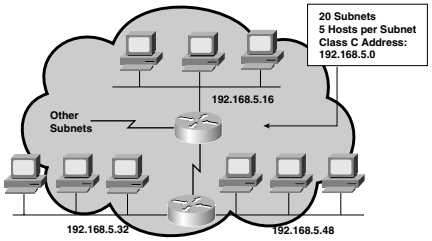

1. Determine the number of subnets and hosts for each subnet required.
2. The address class you are assigned and the number of subnets required determine the number of subnetting bits used. For example, with a Class C address and a need for 20 subnets, you will have a 29-bit mask (255.255.255.248). This allows for the Class C default 24-bit mask and 5 bits required for 20 subnets. (The formula $2^n - 2$ yields only 14 subnets for 4 bits, so 5 bits must be used.)
3. The remaining bits in the last octet are used for the host field. In this case, each subnet has $2^3 - 2$, or 6 hosts.
4. The final host addresses are a combination of the network/subnet plus each host value. The hosts on the 192.168.5.32 subnet would be addressed as 192.168.5.33, 192.168.5.34, 192.168.5.35, and so forth.

Implementing Subnet Planning Summary

- Breaking up networks into smaller segments (or subnets) improves network efficiency and conserves IP addresses.
- A 32-bit subnet mask determines the boundary between the subnet host portions of the IP address using 1s and 0s.
- A subnet defines a broadcast domain in a routed network.
- Cisco IOS Software supports directed, local network, and subnet broadcasts.
- Subnet planning should be based on future growth predictions rather than current needs.

Configuring IP Addresses

An IP address must be assigned to a switch if you plan to use SNMP or connect to the switch through a Web browser or Telnet. If the switch needs to send traffic to a different IP network, the traffic is routed to a default gateway.

Here's the procedure for configuring a switch IP address:

```
SwitchA>enable
SwitchA#config term
SwitchA(config)#ip address 10.2.5.10 255.255.255.0
SwitchA(config)#ip default-gateway 10.2.5.2
SwitchA(config)#exit
SwitchA#show ip
```

The **no ip address** command resets the address to the default (0.0.0.0).

Each unique IP address can have a host name associated with it. A maximum of six IP addresses can be specified as named servers. *Domain Name System* (DNS) is a system used to translate names into addresses. If a system sees an address it does not recognize, it refers to DNS, which is enabled by default with a server address of 255.255.255.255. The **ip domain-lookup** and the **no ip domain-lookup** commands turn DNS on and off, respectively.

Router IP Host Names

When names are used to route traffic, they must be translated into addresses. Routers must be able to associate host names with IP addresses to communicate with other IP devices. The **ip host** command manually assigns host names to addresses.

Configuring IP Addresses Summary

- The **ip address** command sets the IP address and subnet mask.
- The **ip name-server** command defines which hosts can provide the name service.
- DNS translates node names into addresses.
- The **show hosts** command displays host names and addresses.

Configuring Network Routing

VLANs create Layer 2 segments. End stations in different segments (broadcast domains) cannot communicate with each other without the use of a Layer 3 device, such as a router. Each VLAN must have a separate physical connection on the router, or trunking must be enabled on a single physical connection for inter-VLAN routing to work.

The figure depicts a router attached to a switch. The end stations in the two VLANs communicate by sending packets to the router, which forwards them to the other VLAN. This setup is referred to as a *router on a stick*.

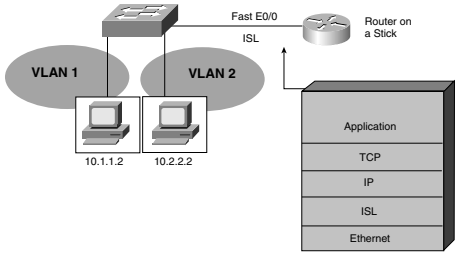

Dividing Physical Interfaces into Subinterfaces

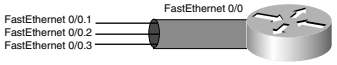

ISL trunking requires the use of subinterfaces. A *subinterface* is a logical, addressable interface on the router's physical Fast Ethernet port. Several subinterfaces can be on a single port (one per VLAN). The **encapsulation isl** *domain* command (in subinterface configuration mode) enables ISL. The *domain* parameter refers to the VLAN domain number.

In the figure, the FastEthernet 0 interface is divided into multiple subinterfaces (FastEthernet 0.1, FastEthernet 0.2, and so on).

Configuring Network Routing Summary

- "Router on a stick" is a router attached only to a switch. The router receives packets from one VLAN and forwards them to another VLAN.
- A subinterface is required to support ISL trunking.
- To configure a router on a stick, enable ISL on the switch port connected to the router, enable ISL encapsulation on the router's FastEthernet subinterface, and assign a network layer address to each subinterface.

Determining IP Routes

Routing Overview

Routing is the process of getting packets and messages from one location to another.

Key Information a Router Needs

The router needs the following key information:

- **Destination address**—The destination (typically an IP address) of the information being sent
- **Sources of information**—Where the information came from (typically an IP address)
- **Possible routes**—Likely routes to get from source to destination
- **Best route**—The best path to the intended destination
- **Status of routes**—Known paths to the most current destinations

A router is constantly learning about routes in the network and storing this information in its routing table. The router uses its table to make forwarding decisions. The router learns about routes in one of two ways:

- Manually (routing information is entered by the network administrator)
- Dynamically (routing processes running in the network)

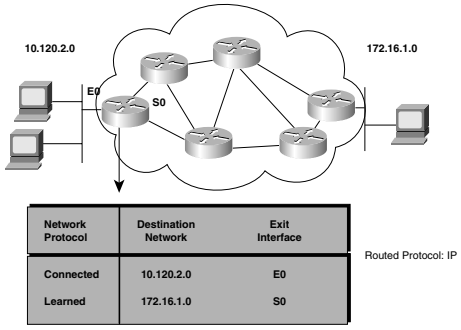

Identifying Static and Dynamic Routes

A router uses static or dynamic routes when forwarding packets:

- *Static routes* are manually entered by the network administrator. These routes must be manually updated whenever there is a topology change.
- *Dynamic routes* are learned by the router. Unlike static routes, topology changes are learned without administrative intervention and are automatically propagated throughout the network.

Examining Static Routes

Static routes specify the path packets take, allowing precise control over a network's routing behavior. Static routes are sometimes used to define a *gateway of last resort*. This is

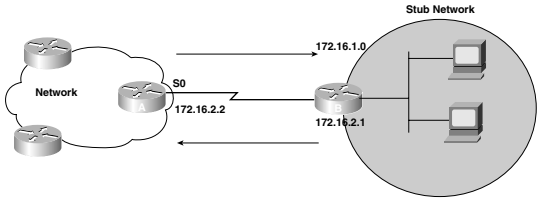

where a packet is routed if no other suitable path can be found. Static routes are also used when routing to a stub network. A *stub network* is a network accessed by a single route. Often, static routes are the only way on to or off of a stub network. Static routes are also used for security reasons or when the network is small.

Examining Static Route Configuration

The **ip route** command configures a static route in global configuration mode. This command manually sets the routing table. This table entry will not accept dynamic changes as long as the path is active.

```
ip route network [mask] {address | interface} [distance] [permanent]
```

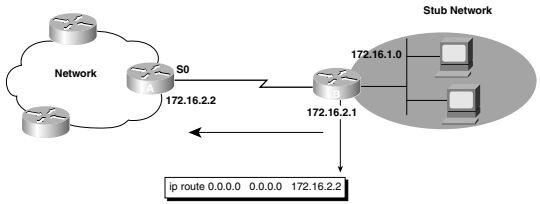

- *address*—IP address of the next-hop router.
- *interface*—Interface to the destination network. Must be a point-to-point interface.
- *distance* (Optional)—Defines the administrative distance.
- *permanent* (Optional)—Specifies that the route will not be removed, even if the interface shuts down.

A default route is a special type of static route. Use a default route when the route is not known or when storing the needed information is unfeasible.

Routing Summary

- Routing is the process of sending packets from one location to another. A router needs to know the destination address, source address, initial possible routes, and best path to route packets.
- Routing information is stored in the router's routing table.
- Static routes are user-defined, whereas dynamic routes are learned by the router running a routing protocol.
- Use the **ip route** command to configure a static route.
- A default route is used for situations in which the route is not known or when it is unfeasible for the routing table to store sufficient information about the route.

Dynamic Routing Overview

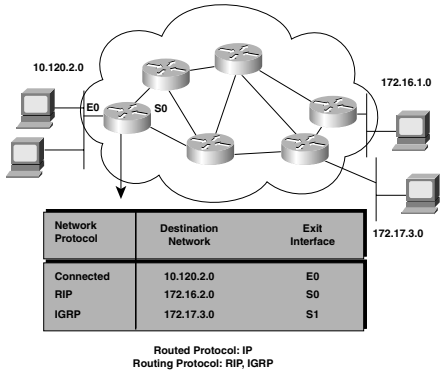

Routing protocols are used to determine paths between routers and to maintain routing tables. Dynamic routing uses routing protocols to disseminate knowledge throughout the network. A routing protocol defines communication rules and interprets network layer address information. Routing protocols describe the following:

- Routing update methods
- Information contained in updates
- When updates are sent
- Paths to other routers

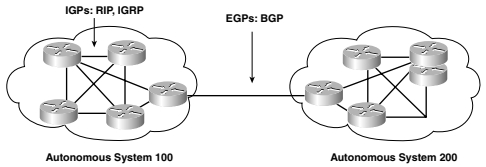

Autonomous Systems

An *autonomous system* (AS) is a group of networks under a common administrative domain. *Interior Gateway Protocols* (IGPs), such as RIP and IGRP, exchange routing information within an autonomous system. *Exterior Gateway Protocols* (EGPs) are used to connect autonomous systems. Border Gateway Protocol (BGP) is an example of an EGP.

Ranking Routes with Administrative Distance

Several routing protocols can be used simultaneously in the same network. When there is more than a single source of routing information, an administrative distance value is used to rate the trustworthiness of each routing information source. The administrative distance metric is an integer from 0 to 255. In general, a route with a lower number is considered more trustworthy and is more likely to be used.

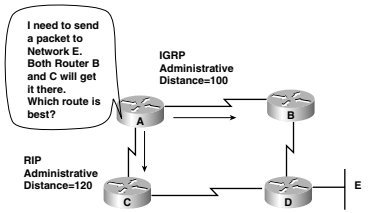

Default Distance Values

Route Source	Default Distance
Connected interface	0
Static route address	1
EIGRP	90
IGRP	100
OSPF	110
RIP	120
Ext. EIGRP	170
Unknown	255

Routing Protocol Classes

There are three basic routing protocol classes:

- **Distance vector**—Uses the direction (vector) and distance to other routers as metrics. RIP and IGRP are both distance vector protocols.
- **Link state**—Also called shortest path first, this protocol re-creates the topology of the entire network.
- **Balanced hybrid**—Combines the link-state and distance vector algorithms.

How Distance Vector Protocols Route Information

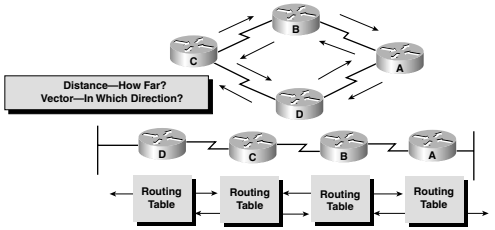

Routers using distance vector-based routing share routing table information with each other. This method of updating is called *routing by rumor.* Each router receives updates from its direct neighbor. In the figure, router B shares information with routers A and C. Router C shares routing information with routers B and D. In this case, the routing information is distance vector metrics (such as number of hops). Each router increments the metrics as they are passed on (incrementing hop count, for example).

Distance accumulation is a method that keeps track of the routing distance between any two points in the network, but the routers do not have an exact topology of an internetwork.

How Information Is Discovered with Distance Vectors

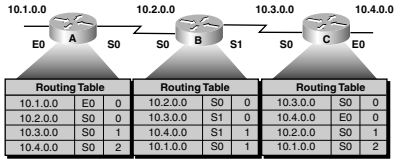

Network discovery is the process of learning about destinations that are not directly connected. As the network discovery proceeds, routers accumulate metrics and learn the best paths to various destinations. In the figure, each directly connected network has a distance of 0. Router A learns about other networks based on information it receives from Router B. Router A increments the distance metric for any route learned by Router B. For example, router B knows about the networks to which Router C is directly connected. Router B then shares this information with Router A, which increments the distance to these networks by 1.

Examining Distance Vector Routing Metrics

Distance vector routing protocols use routing algorithms to determine the best route. These algorithms generate a metric value for each path through the network. The smaller the metric, the better the path. Metrics can be calculated based on one or more characteristics of a path.

Commonly Used Metrics

- **Hop count**—Number of times a packet goes through a router.
- **Ticks**—Delay on a data link using IBM PC clock ticks (approximately 55 ms).
- **Cost**—An arbitrary value based on a network administrator-determined value. Usually bandwidth, cost in dollars, or time.
- **Bandwidth**—A link's data capacity.
- **Delay**—Time required to reach the destination.
- **Load**—Network activity.
- **Reliability**—Bit error rate of each network link.
- **MTU (maximum transmission unit)**—The maximum message length allowed on the path.

Link-State Routing

The *link-state-based routing algorithm* (also known as shortest-path first [SPF]) maintains a database of topology information. Unlike the distance vector algorithm, link-state routing maintains full knowledge of distant routers and how they interconnect. Link-state routing provides better scaling than distance vector routing for the following reasons:

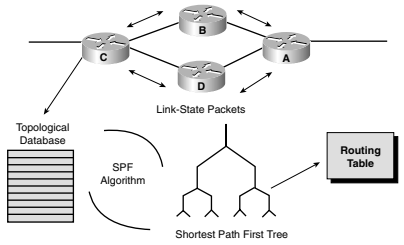

- Link state sends only topology changes. Distance vector sends complete routing tables.
- Link-state updates are sent less often than distance vector.
- Link state is hierarchical, which limits the scope of route changes.
- Link state supports classless addressing and summarization.

Balanced Hybrid Routing

Balanced hybrid routing combines aspects of both distance vector and link-state protocols. Balanced hybrid routing uses distance vector with more accurate metrics, but unlike distance vector routing protocols, it updates only when there is a topology change. Balanced hybrid routing provides faster convergence while limiting the use of resources such as bandwidth, memory, and processor overhead. The Cisco EIGRP is an example of a balanced hybrid protocol.

Dynamic Routing Summary

- Routing protocols use the network layer address to forward packets to the destination network.
- An AS is a collection of networks under a common administrative domain.
- More than one routing protocol can be used at the same time. Administrative distance is used to rate the trustworthiness of each information source.
- Distance vector, link state, and balanced hybrid are the most common IGPs.
- Distance vector-based algorithms send copies of routing tables. As network discovery proceeds, routers accumulate metric information used to determine the best path to distant networks.
- A link-state routing algorithm, also known as SPF, sends network topology information rather than metrics.
- Balanced hybrid routing combines aspects of both distance-vector and link state-protocols.

Distance Vector Routing

Any topology change in a network running a distance vector protocol triggers an update in the routing tables. The topology updates follow the same step-by-step process as the initial network discovery.

The process of updating routing tables is as follows: A router compares the information contained in the update to its current table. If the update contains information about a better route to some destination (a smaller total metric), the router updates its own routing table. As always, the distance metric is incremented 1 greater than the neighbor it received the update from. During updates, the router sends its entire routing table to each of its adjacent neighbors. The table includes the total path cost (defined by its metric) and the logical address of the first router on the path to each destination network. In the figure, Router B is one unit of cost from Router A. Therefore, it adds 1 to all costs reported by Router A.

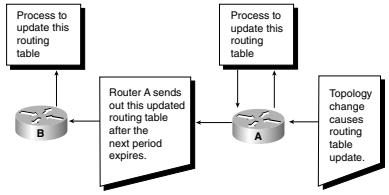

How Routing Loops Occur in Distance Vector

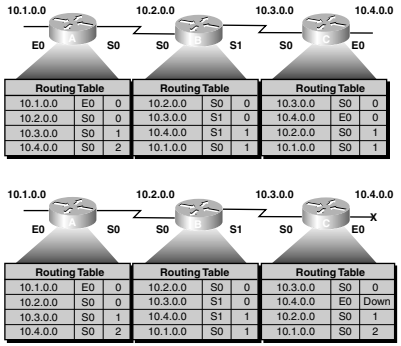

During updates, routing loops can occur if the network has inconsistent routing entries. Slow convergence on a new configuration is one cause of this phenomenon. The network is converged when all routers have consistent routing tables. The figure illustrates how a routing loop occurs. Before a network failure, all routers have correct tables. The figure uses hop count as a cost metric, so the cost of each link is 1. Router C is directly connected to network 10.4.0.0 with a distance of 0. Router A's path to network 10.4.0.0 is through Router B, with a hop count of 2.

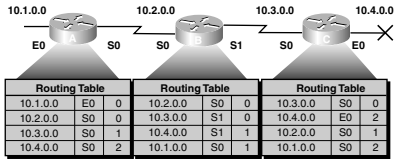

If network 10.4.0.0 fails, Router C detects the failure and stops routing packets to that network. At this point, Routers A and B do not know of the failure. Router A's table still shows a valid path to 10.4.0.0 through Router B. If Router B sends out its normal update to A and C, Router C sees a valid path to 10.4.0.0 through Router B and updates its routing table to reflect a path to network 10.4.0.0 with a hop count of 2. (Remember, it must add 1 to the hop count.) Now router C sends an update to Router B, which then updates Router A. Router A detects the modified distance vector to network 10.4.0.0 and recalculates its own distance vector to 10.4.0.0 as 4. With each update, the incorrect information continues to bounce between the routers. Without some mechanism to prevent this, the updates continue. This condition, called *count to infinity,* continuously loops packets around the network.

Some distance vector routing algorithms have a Time-To-Live (TTL) value, but the routing loop problem can occur with this limit. You can avoid this by defining infinity as some maximum number. When you set this value, the routing loop continues until the metric exceeds the maximum allowed value. After the metric value exceeds the maximum, the network is unreachable.

Eliminating Routing Loops Through Split Horizon

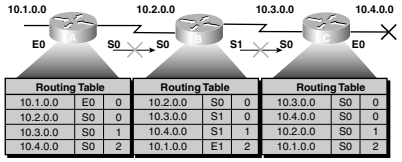

Split horizon is one way to eliminate routing loops and speed up convergence. The idea behind split horizon is that it is never useful to send information about a route back in the direction from which the update came. If the router has no valid alternative path to the network, it is considered inaccessible. Split horizon also eliminates unnecessary routing updates, thus speeding convergence.

Route Poisoning

Route poisoning, which is part of split horizon, also eliminates routing loops caused by inconsistent updates. Route poisoning basically locks the table (using hold-down timers) until the network has converged.

Example of Route Poisoning

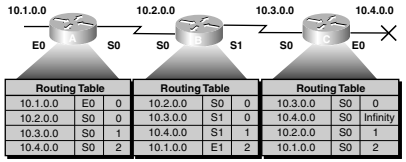

The figure provides the following example: When network 10.4.0.0 goes down, Router C "poisons" its link to network 10.4.0.0 with an infinite cost (marked as unreachable).

Router C is no longer susceptible to incorrect updates about network 10.4.0.0 coming from neighboring routers that might claim to have a valid alternative path. After the hold-down timer expires (which is just longer than the time to convergence), Router C begins accepting updates again.

Poison Reverse

When Router B sees the metric to 10.4.0.0 jump to infinity, it sends a return message (overriding split horizon) called a *poison reverse* back to Router C, stating that network 10.4.0.0 is inaccessible. This message ensures that all routers on that segment have received information about the poisoned route.

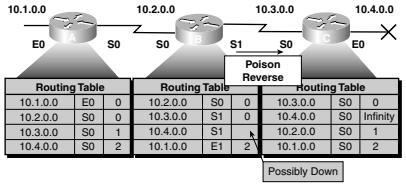

Avoiding Routing Loops with Triggered Updates

A triggered update is sent immediately in response to a change in the network. The router detecting the change immediately sends an update message to adjacent routers, which then generate their own triggered updates. This continues until the network converges. There are two problems with triggered updates:

- The update message can be dropped or corrupted.
- The updates do not happen instantly. It is possible that a router issued a regular update before receiving the triggered update. If this happens, the bad route can be reinserted into a router that received the triggered update.

The solution is as follows: Hold-down timers dictate that when a route is invalid, no new route with a same or worse metric will be accepted for the same destination for a certain period of time. This allows the triggered update to propagate throughout the network.

Characteristics of Hold-Down Timers

- They are used to prevent regular update messages from inappropriately reinstating a route that might have gone bad.
- Hold-down timers force routers to hold any changes for a period of time.
- The hold-down period should be calculated to be just greater than the amount of time it takes for updates to converge.

Hold-Down Implementation Process

1. When a router receives an update that a network is down, the router marks the route as inaccessible and starts a hold-down timer.

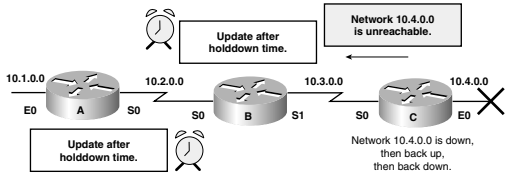

2. If an update is received from a neighboring router with a better metric, the router removes the timer and uses the new metric.
3. If an update with a poorer metric is received before the hold-down timer expires, the update is ignored.
4. During the hold-down period, routes appear in the routing table as "possibly down."

Distance Vector Routing Summary

- Distance vector routing protocols maintain routing information by updating routing tables with neighboring routing tables.
- Defining a max count prevents infinite loops.
- Split horizon solves routing loops by preventing routing updates from being sent back in the same direction from which they came.
- Route poisoning sets downed routes to infinity to make that route unreachable.
- A triggered update is sent immediately in response to a change. Each router receiving a triggered update sends its own until the network converges.
- Hold-down timers prevent regular update messages from reinstating failed routes.
- More than one loop-preventing solution can be implemented on networks that have multiple routes.

Enabling RIP

To enable a dynamic routing protocol, you must do the following:

- Select a routing protocol (such as RIP or IGRP).
- Assign IP network numbers.
- Assign network/subnet addresses and the appropriate subnet mask to interfaces.

The **network** command starts up the routing protocol. The **network** command also specifies a directly connected network and advertises that network.

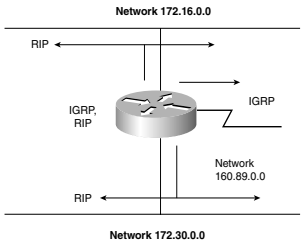

RIP

Cisco devices support two versions of *Routing Information Protocol* (RIP): RIP Version 1 (RFC 1058) and an enhanced version, RIP Version 2, a classless routing protocol.

Here are key RIP characteristics:

- RIP is a distance vector routing protocol.
- Hop count is used as the metric for path selection (the maximum is 15).
- Routing updates broadcast every 30 seconds (the default).
- RIP can load-balance over six equal-cost paths (the default is 4).
- Only one network mask can be used for each classful network (RIPv1).
- RIPv2 permits variable-length subnet masks on the internetwork.
- RIPv2 performs triggered updates.

A *classless routing protocol* allows routers to summarize information about several routes in order to cut down on the quantity of information carried by the core routers. With classless IP configured, packets received with an unknown subnet of a directly attached network are sent to the next hop on the default route. With *classless interdomain routing* (CIDR), several IP networks appear to networks outside the group as a single, larger entity.

Defining Paths

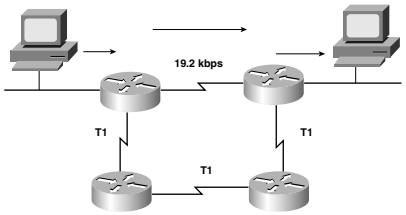

Load balancing occurs when a router has several equal-cost paths to the same destination. *Load sharing* is when a router has several unequal-cost paths to the same destination. If a router has unequal-cost paths to the destination, it does not load-balance unless the **variance** command is used. RIPv1 doesn't support unequal load balancing. Load balancing can be disabled by setting the maximum number of paths to 1.

Here's the procedure for configuring RIP:

```
RouterA>enable
RouterA#config term
RouterA(config)#router rip
RouterA(config-router)#network 10.3.2.0
RouterA(config-router)#exit
RouterA(config)#exit
RouterA#show ip protocols
```

- **show ip protocols**—Shows whether a router is delivering bad routing information
- **show ip route**—Shows RIP routing tables
- **debug ip rip**—Displays RIP routing updates (**no debug all** disables)

Note: The **network** command specifies the autonomous system and starts up the routing protocol in the specified network. The **network** command also allows the router to advertise that network.

Enabling RIP Summary

- To configure a dynamic routing protocol, select a protocol, assign a network number, and assign network addresses for each interface.
- RIP, a distance vector routing protocol, uses hop count as a route selection metric. RIP can load-balance across equal-cost paths.
- The **ip classless** command prevents the router from dropping packets destined for unknown subnets of directly connected networks.

Enabling Interior Gateway Routing Protocol

Interior Gateway Routing Protocol (IGRP) is a distance vector routing protocol with sophisticated metrics and improved scalability, allowing better routing in larger-sized networks. IGRP uses delay and bandwidth as routing metrics (reliability and load are optional). IGRP has a default maximum hop count of 100 hops. (RIP's maximum is 15 hops.) This can be reconfigured to 255 hops.

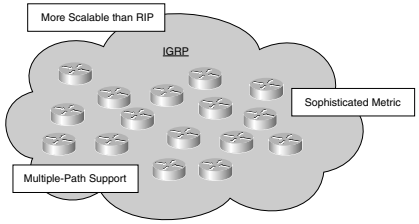

IGRP can load-balance across six nonequal paths, increasing the available bandwidth and providing route redundancy.

IGRP Metrics

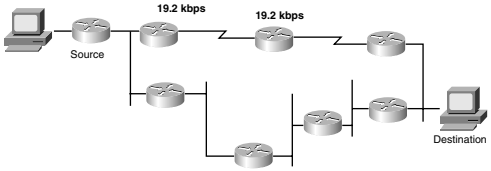

IGRP achieves greater route selection accuracy than RIP with the use of a composite metric. The path that has the smallest metric value is the best route. IGRP's metric includes the following components: bandwidth delay, reliability (based on keepalives), loading (bits per second), and maximum transmission unit (MTU). You can significantly affect performance by adjusting IGRP metric values.

Paths of Different Metrics

IGRP allows load balancing over unequal paths. If two unequal paths are used and one path is four times better than the other, the better path will be used four times as often. The **variance** command specifies the metric range allowed for load balancing across multiple paths.

Here's the procedure for configuring IGRP:

```
RouterA>enable
RouterA#config term
RouterA(config)#router igrp 100
RouterA(config-router)#network 170.8.0.0
RouterA(config-router)#variance 1
RouterA(config-router)#traffic share balanced
RouterA(config-router)#exit
RouterA(config)#exit
RouterA#show ip protocols
```

The syntax for the router IGRP command includes the AS number. All routers within an autonomous system must use the same system number.

- The default value of **variance** is 1 (equal-cost load balancing).
- The **traffic share balanced** command distributes traffic proportionally to the metrics' ratios.
- Use **show ip protocols** to verify the IGRP protocol configuration.
- Use **show ip route** to display the contents of the IP routing tables.

Enabling IGRP Summary

- IGRP has increased scalability and a more sophisticated routing metric than RIP. IGRP can load balance over unequal paths.
- IGRP's routing metric is a composite of bandwidth, delay, reliability, load, and MTU.
- The **debug ip igrp** configuration commands display routing and transaction information for troubleshooting purposes.

Access Lists and Their Applications

As a network grows, it becomes more important to manage the increased traffic going across the network. Access lists help limit traffic by filtering traffic based on packet characteristics. Access lists define a set of rules used by routers to identify particular types of traffic. Access lists can be used to filter both incoming and outgoing traffic on a router's interface. An access list applied to a router specifies rules for only traffic going through the router. Traffic originating from a router is not affected by that router's access lists. (It *is* subject to access lists within other routers as it passes through them.)

Packet Filtering

Access lists can be configured to permit or deny incoming and outgoing packets on an interface. By following a set of conventions, the network administrator can exercise greater control over network traffic by restricting network use by certain users or devices.

Applications of an IP Access List

To establish an access list, you must define a sequential list of permit and deny conditions that apply IP addresses or IP protocols. Access lists filter only traffic going through the router; they do not filter traffic originated from the router. Access lists can also filter Telnet traffic in to or out of the router's vty ports.

Access List Type		Number Range/Identifier
IP	Standard Extended Named	1-99 100-199 Name (Cisco IOS 11.2 and later)
IPX	Standard Extended SAP filters Named	800-899 900-999 1000-1099 Name (Cisco IOS 11.2 F and later)

Other Access List Uses

- Access lists allow finer granularity of control when you're defining priority and custom queues.
- Access lists can be used to identify "interesting traffic," which triggers dialing in dial-on-demand routing (DDR).
- Access lists filter and in some cases alter the attributes within a routing protocol update (route maps).

Types of Access Lists

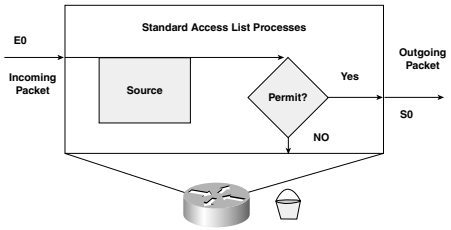

There are two general types of access lists:

- Standard access lists check the source address of packets. Standard access lists permit or deny output for an entire protocol suite based on the source network/subnet/host IP address.
- Extended IP access lists check both source and destination packet addresses. Extended lists specify protocols, port numbers, and other parameters, giving administrators more flexibility and control.

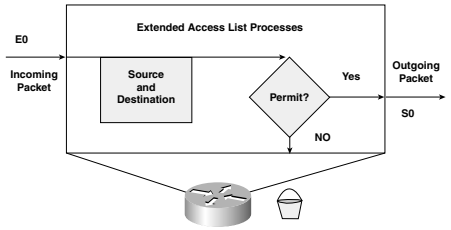

Standard	Extended
Filter based on source	Filter based on source and destination
Permit or deny the entire TCP/IP protocol suite	Specify a specific IP protocol and port number
Range: 1 to 99	Range: 100 to 199

Access List Process Options

- **Inbound access lists**—Incoming packets are processed prior to being sent to the outbound interface. If the packet is to be discarded, this method reduces overhead (no routing table lookups). If the packet is permitted, it's processed in the normal way.
- **Outbound access lists**—Outgoing packets are processed by the router first and then are tested against the access list criteria.

Permit or Deny Process

Access list statements are operated on one at a time from top to bottom. After a packet header match is found, the packet is operated on (permitted or denied), and the rest of the statements are skipped.

If no match is found, the packet is tested against the next statement until a match is found or the end of the list is reached. An implicit **deny** statement is present at the end of the list. (All remaining packets are dropped.)

Unless there is at least one **permit** statement in an access list, all traffic is blocked.

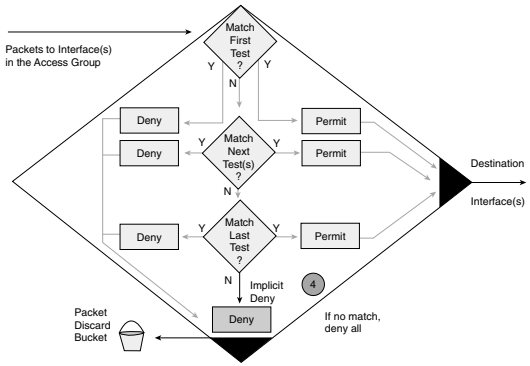

Guidelines for Implementing Access Lists

- Be sure to use the correct numbers for the type of list and protocols you want to filter.
- You can use only one access list per protocol, direction, and interface. A single interface can have one access list per protocol.
- Put more-specific statements before more-general ones. Frequently occurring conditions should be placed before less-frequent conditions.
- Additions are always put at the end of the access list. You cannot selectively add or remove statements in the middle of an access list.
- Without an explicit **permit any** statement at the end of a list, all packets not matched by other statements are discarded. Every access list should include at least one **permit** statement.
- An interface with an empty access list applied to it allows (permits) all traffic. Create your statements before applying the list to an interface.
- Access lists filter only traffic going through the router.

Protocol Access List Identifiers

The access list number entered by the administrator determines how the router handles the access list. The arguments in the statement follow the number. The types of conditions allowed depend on the type of list (defined by the access list number). Conditions for an access list vary by protocol. You can have several different access lists for any given protocol, but only one protocol is allowed on any access list (one protocol per direction per interface).

TCP/IP Packet Tests

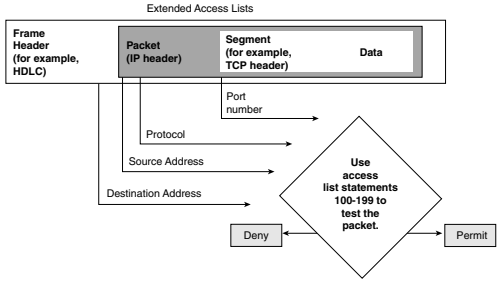

For TCP/IP packets, access lists check the packet and upper-layer headers for different items (this depends on the type of access list, standard or extended). Standard access lists are assigned a number in the range 1 to 99. Extended access lists use the range 100 to 199. As soon as a packet is checked for a match with the **access list** statement, it is either permitted to an interface or discarded.

Wildcard Masking

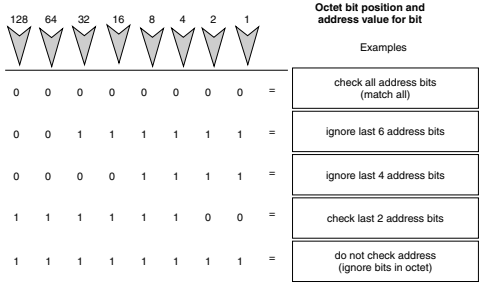

It is not always necessary to check bytes within an address. Wildcard masking identifies which bits should be checked or ignored. Administrators can use this tool to select one or more IP addresses for filtering. Wildcard masking is exactly opposite of subnet masking.

- A wildcard mask bit 0 means check the corresponding bit value.
- A mask bit 1 means do not check (ignore) that corresponding bit value.

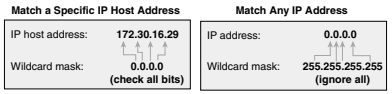

To specify an IP host address within a **permit** or **deny** statement, enter the full address followed by a mask of all 0s (0.0.0.0).

To specify that all destination addresses will be permitted in an access list, enter 0.0.0.0 as the address, followed by a mask of all 1s (255.255.255).

Abbreviated Commands in Wildcard Masking

You can use abbreviations rather than typing an entire wildcard mask:

- **Checking all addresses**—To match a specific address, use **host**. For example, **172.30.16.29 0.0.0.0** can be written as **host 172.30.16.29**.
- **Ignoring all addresses**—Use the word **any** to specify all addresses. For example, **0.0.0.0 255.255.255.255** can be written as **any**.

Access Lists and Their Applications Summary

- Access lists filter packets as they pass through the router.
- The two general types of access lists are standard and extended. Standard lists filter based on only the source address, and extended lists filter based on source and destination addresses, as well as specific protocols and numbers.
- Access lists can be set to either inbound or outbound. For inbound access lists, the packets are processed first and then routed to an outbound interface (assuming that the filter passes them). In outbound access lists, the packets are sent to the interface and then routed.
- If a packet meets a **permit** statement's criteria, it is passed to the next statement. If a packet meets a **deny** statement's criteria, it is immediately discarded.
- More-restrictive statements should be at the top of the list.
- Only one access list per interface, per protocol, per direction is allowed.
- Every access list should have at least one **permit** statement.
- For IP, standard access lists use the number range 1 to 99, and extended access lists use 100 to 199. For IPX, standard access lists use 800 to 899, and extended access lists use 900 to 999.
- Wildcard masking is used to filter single IP addresses or blocks of addresses.

Access List Configuration

Principles of Configuring Access Lists

Access lists are processed from top to bottom, making statement ordering critical to efficient operation. Always place specific and frequent statements at the beginning of an access list. Named access lists allow the removal of individual statements (but no reordering). To reorder statements, you must remove and re-create the whole list with the proper statement ordering. Use a text editor to create lists. Remember that all access lists end with an implicit **deny all** statement.

Access List Syntax

The syntax for a standard and extended IP access lists is **access-list access-list-number {permit | deny} source [mask]**.

```
access-list access-list-number {permit | deny} protocol source
    source-wildcard
  [operator port] destination destination-wildcard [operator port]
  [established] [log]
```

operator port can be less than, greater than, equal to, or not equal to a port number. **established** (used for inbound TCP only) allows only established connections to pass packets. **log** sends a logging message to the console.

After the statements are added, they are applied to an access group using the following syntax:

```
ip access-group access-list-number {in | out}
```

Here is the procedure for configuring extended IP access lists:

```
RouterA>enable
RouterA#access-list 101 deny tcp 172.16.4.0 0.0.0.255 72.16.3.0 0.0.0.255 eq
21
RouterA#config term
RouterA(config)#interface ethernet 0
RouterA(config-if)#access group 101 in
RouterA(config)#exit
RouterA#show ip interface
```

Named Access Lists

Named IP access addresses (Cisco IOS Release 11.2 and later) allow alphanumeric strings as identifiers rather than numbers. Named access lists can be standard or extended. Named IP access lists also allow you to delete individual statements from an access list.

You should use named access lists when more than 99 standard or extended access lists are configured on any router. Duplicate names are not allowed on any one router. (You can use the same name on two different routers.)

Guidelines for Placing Access Lists

Extended access lists can block traffic from leaving the source. They should be as close as possible to the source of the traffic to be denied. Standard access lists block traffic at the destination. They should be as close as possible to the destination of the traffic to be denied.

Virtual Terminal Access Lists

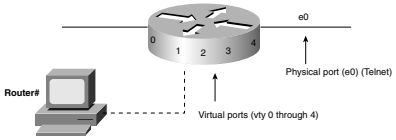

In addition to physical ports, devices also have virtual ports (called virtual terminal lines). There are five such virtual terminal lines, numbered vty 0 through vty 4. Standard and extended access lists do not prevent router-initiated Telnet sessions.

Virtual terminal access lists can block vty access to the router or block access to other routers on allowed vty sessions. Restrictions on vty access should include all virtual ports, because users can connect through any vty port. The syntax for a vty access list is **line vty {vty# | vty-range}**.

After you add the vty statements, you assign them to the router with the following command:

```
access-class access-list-number {in | out}
```

Specifying **in** prevents incoming Telnet connections, and **out** prevents Telnet connections to other routers from the vty ports.

Access List Configuration Summary

- Here are some general guidelines for configuring access lists:
  - All access lists end with an implicit **deny**.
  - More-specific tests should precede more-general tests.
  - Frequently used tests should precede infrequent tests.
- Standard access lists filter based on source addresses only.
- Extended access lists filter based on source and destination addresses, protocols, and ports.
- The **access-list** command assigns statements to a list. The **access-group** command assigns an access list to an interface.
- Named access lists allow you to identify access lists with alphanumeric strings rather than numbers. You can delete entries from a named access list.
- Extended access lists should be close to the source of the traffic to be denied.
- Standard access lists should be close to the destination.
- Access lists can be used to control virtual terminal (vty) access to or from a router.
- The **line vty** and **access-class** commands are used to configure and set vty access lists.

IPX Routing Overview

Cisco routers are compatible with NetWare (Novell) networks. They have the following features:

- Interface support, including native ISDN and ATM
- IPX access filters for several protocols (IPX, RIP, SAP, NCP, and NetBIOS)
- Support for EIGRP and NLSP
- Serverless LAN support
- Dial-on-demand routing (DDR)

Key Novell NetWare Features

- Novell IPX addresses use 32 bits for the network number and 48 bits for the node number, expressed as a hexadecimal number.
- The MAC address of an interface is used as the node number.
- Multiple logical networks can be configured on a single interface, but each network must use a different encapsulation type.

- The default routing protocol is IPX RIP.
- NetWare servers and routers announce their services to clients using SAP broadcasts. One type of SAP broadcast is *Get Nearest Server* (GNS).

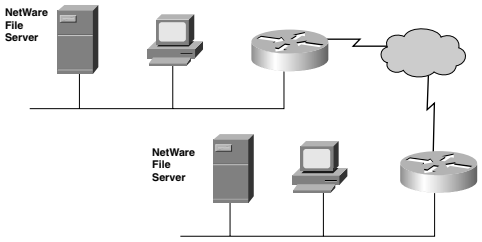

NetWare Protocol Stack

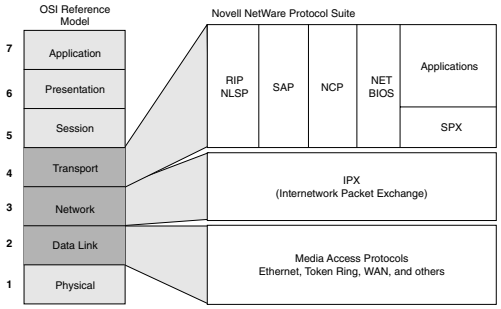

Novell IPX/SPX is a proprietary suite of protocols based on the Xerox Network Systems (XNS) protocol suite. All common media access protocols are supported on the NetWare protocol stack. IPX is a Layer 3 connectionless protocol.

Novell IPX Addressing

Novell IPX addresses have two parts: a network number and a node number. IPX has the following characteristics:

- The network number can be up to eight hexadecimal digits long. (The leading 0s are not shown.)
- A network number is assigned to servers and routers by the network administrator.
- Clients dynamically learn the network address upon startup.

IPX Ethernet Frame Structures

Cisco routers support all four framing variations allowed by NetWare. Each encapsulation type has a specific use:

- **Ethernet_802.3 (raw Ethernet)**—The default for NetWare 2.0 to 3.11
- **Ethernet_802.2**—The default for NetWare 3.12 and later
- **Ethernet_II**—Used with TCP/IP and DECnet
- **Ethernet_SNAP**—Used with TCP/IP and AppleTalk

Multiple encapsulations can be used on a single interface as long as multiple network numbers have also been assigned. Clients and servers with different framing types cannot communicate directly with each other.

Cisco Encapsulation Types

In the figure, the Novell framing type is matched to its Cisco equivalent. The default encapsulation type is used if one is not specified. All devices must use the same encapsulation type if they are to communicate directly.

	Novell IPX Name	Cisco IOS Name
Ethernet	Ethernet_802.3	novell-ether
	Ethernet_802.2	sap
	Ethernet_II	arpa
	Ethernet_SNAP	snap
Token Ring	Token-Ring	sap
	Token-Ring_SNAP	snap
FDDI	FDDI_SNAP	snap
	FDDI_802.2	sap
	FDDI_Raw	novell-fddi

Specify encapsulation when you configure IPX networks

IPX RIP

IPX RIP (a distance-vector routing protocol) uses ticks and hop count as metrics. RIP checks ticks first and then uses hop count if two or more paths have the same tick value. IPX routers broadcast copies of their routing tables in the same manner that IP routers do. The split-horizon algorithm prevents routing loops in IPX networks.

Service Advertising Protocol

All NetWare servers advertise their service types and service addresses. NetWare uses SAP broadcasts to dynamically announce, locate, add, and remove services on the network.

Router SAP Tables

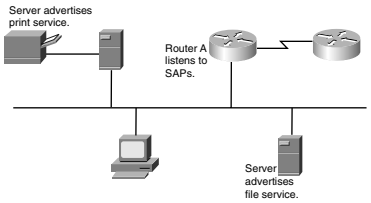

Rather than forward SAP broadcasts (which would add significant traffic to the network), routers build and send SAP tables. The tables are sent every 60 seconds by default.

Initiating a Connection to a NetWare Server

When a client powers up, it broadcasts a GNS SAP query. All local NetWare file servers respond with a SAP reply. The client can then log into the target server. If a Cisco router receives a GNS query, it does not respond unless no NetWare servers are on the network.

IPX Routing Summary

- Novell NetWare supports access lists and filters, scalable routing protocols, serverless LAN, dial-on demand routing, and rich management control.
- RIP, SAP, IPX, SPX, and NLSP are included in the NetWare protocol stack.
- A Novell IPX address is made up of an administrator-assigned network number and a node number (usually derived from an interface MAC address).
- NetWare servers broadcast information tables every 60 seconds. Routers use this information to build SAP tables, which they share with other routers.
- Novell IPX RIP is a distance vector routing protocol. It uses hop count and ticks for a metric.
- Network devices advertise their services using SAP broadcasts.
- NetWare clients use GNS and SAP queries to locate network services.

IPX Filtering

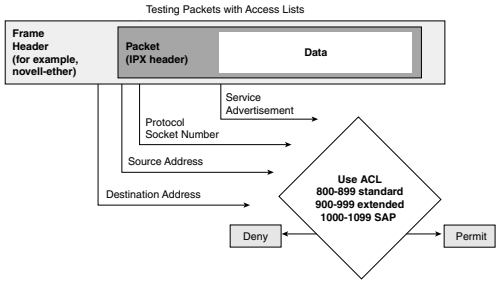

Novell IPX uses access lists to filter packets. Packets can be filtered using standard, extended, or SAP access lists. Standard access lists (numbered from 800 to 899) use destination and source IPX addresses to filter packets. Extended access lists (numbered from 900 to 999) use them to filter packets. SAP filter access lists use service advertisement numbers to filter packets and are numbered from 1000 to 1099. Packets are permitted or denied based on the criteria specified in the access list statements. Wildcard masks are used to specify individual addresses or blocks of addresses.

IPX Standard Access Lists

The syntax for standard IPX access lists is as follows:

```
access-list access-list-number {deny | permit} source-network
  [.source-node [source-node-mask]] destination-network
  [.dest-node [dest-node-mask]]
```

source-network.network-node and *destination-network.network-node* denote the IPX source and destination addresses, respectively. (–1 equals any network.)

After entering the statements, you apply the access list to the interface using the following command:

```
ipx access-group access-list-number [in | out]
```

Note: IPX access lists default to outbound filters.

IPX Extended Access Lists

Use the following syntax for extended IPX access lists:

```
access-list access-list-number {deny | permit} protocol
  [source-network]
  [[.source-node] source-node-mask] |
  [.source-node source-network-mask. source-node-mask]]
  [source-socket] [dest.network] [[[.dest-node]dest-node-mask] |
  [.dest-nodedest-network-mask. destnodemask]] [des-socket] [log]
```

protocol can be a name or number (decimal) of an IPX protocol type. The **log** parameter logs IPX access control list violations to a syslog server whenever a packet matches a particular access list entry.

Procedure for Configuring Extended IP Access Lists

SAP Operation

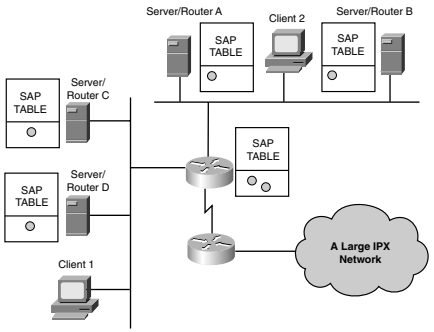

To minimize overhead traffic, routers synchronize the list of available services by forwarding SAP tables rather than forwarding every SAP broadcast. The router advertises its SAP table every 60 seconds by default.

SAP Filter Types

IPX input SAP filters limit the number of services entered into the SAP table. The propagated SAP updates contain a subset of all services.

IPX output SAP filters limit the number of services propagated from the table. The propagated SAP updates contain a subset of all the known services.

Be sure that all clients will see all advertisements required for their application processing. Always place SAP filters as close as possible to the source of the SAP information. This is the most efficient use of bandwidth.

Examining SAP Filter Configuration

Use the following syntax for SAP filter definition statements:

```
access-list access-list-number {deny | permit} network [.node]
  [network-mask.node-mask] [service-type [server-name]]
```

SAP filters use the range 1000 to 1999. Each SAP service type is identified by a hexadecimal number (a 0 matches all services).

After entering the definition statements, you activate the SAP filter with the following command:

```
ipx [input | output]-sap-filter access-list-number
```

IPX Filtering Summary

- Novell addressing is used to create standard, extended, and SAP filter access lists.
- Standard access lists permit or deny information based on source and destination.
- Extended access lists filter on protocol type, source network/node, destination network/node, IPX protocol, and IPX socket number.
- Routers build SAP tables based on SAP server advertisements. Routers forward their tables every 60 seconds by default.
- SAP filters can be placed on a router for both incoming and outgoing traffic.
- SAP filters control the propagation of SAP messages.

Configuring IPX Routing

You must do the following to configure Novell IPX as a routing protocol:

- Start the IPX routing process.
- Enable load sharing to balance packets across multiple routes and links.
- Assign unique network numbers to each interface. (Multiple network numbers can be assigned to an interface for different encapsulation types.)
- Change the encapsulation type, if required.

IPX Configuration Commands

The following commands are used when configuring IPX routing:

- **ipx routing [node]**—Enables IPX routing and SAP services.
- **ipx maximum-paths [paths]**—Enables load sharing. The default is 1, and the maximum is 64.
- **ipx network network [encapsulation encapsulation type]**—Enables IPX routing on a particular interface.

Subinterfaces

Subinterfaces allow a single physical interface to support multiple logical subinterfaces or networks. Novell IPX subinterfaces have the following characteristics:

- Each subinterface must have a distinct network number and encapsulation type.
- All clients and servers using the same network number must also have the same encapsulation type.
- An interface can have multiple encapsulation types only if multiple network numbers are assigned to that interface.

You can assign multiple network numbers to a single interface using subinterfaces or by assigning primary and secondary networks on the interface. In the figure, the router is configured with two subinterfaces, allowing it to communicate with two servers using different encapsulation types. The two servers cannot communicate with each other directly, because they use different encapsulation types (the router can provide connectivity).

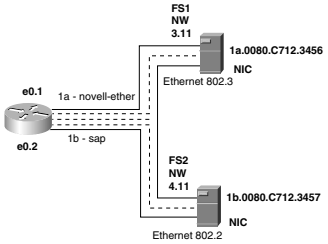

Here's the procedure for configuring IPX:

```
RouterA>enable
RouterA#config term
RouterA(config)#ipx routing
RouterA(config)#ipx maximum-paths 2
RouterA(config)#interface ethernet 0.1
RouterA(config-if)#ipx network 9e encapsulation novell-ether
RouterA(config-if)#exit
RouterA(config)#exit
RouterA#show ipx interface
```

Troubleshooting IPX Routing

The following commands can be used to troubleshoot IPX routing:

- **debug ipx routing activity**—Displays information about IPX routing update packets
- **debug ipx sap**—Displays information about IPX SAP packets
- **ping ipx**—Checks IPX host reachability and network connectivity

Configuring IPX Routing Summary

- The three major commands used to configure IPX routing are **ipx routing**, **ipx maximum-paths**, and **ipx network**.
- Subinterfaces allow a single interface to support multiple logical networks and enable multiple encapsulations per interface.

WAN Concepts and Terminology

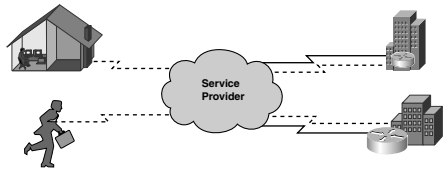

Wide-area networks (WANs) connect networks, users, and services across broad geographic areas. Companies use WANs to connect company sites for information exchange.

Three WAN Connection Types

WAN services are generally leased from service providers on a subscription basis. There are three main types of WAN connections (services):

- **Leased-line**—Provides a preestablished connection through the service provider's network (WAN) to a remote network. Leased lines provide a reserved connection for the client but are costly. Leased-line connections are typically synchronous serial connections with speeds up to 45 Mbps (E3).

- **Circuit-switched**—Provides a dedicated circuit path between sender and receiver for the duration of the "call." Circuit switching is used for basic telephone service or Integrated Services Digital Network (ISDN). Circuit-switched connections are best for clients that require only sporadic WAN usage.

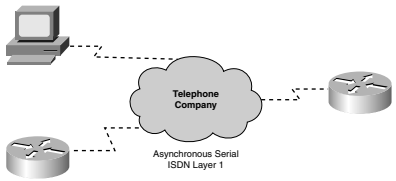

- **Packet-switched**—Devices transport packets using virtual circuits (VCs) that provide end-to-end connectivity.

Programmed switching devices provide physical connections. Packet headers are used to identify the destination. Packet switching offers leased-line-type services over shared lines, but at a much lower cost. Packet-switched networks typically use serial connections with speeds ranging from 56 Kbps to E3.

Other WAN Connections

- **Digital Subscriber Line (DSL)**—Delivers high-bandwidth connections over existing copper telephone lines. There are four varieties of DSL:
  - — Asymmetric digital subscriber line (ADSL)
  - — High-data-rate digital subscriber line (HDSL)
  - — Single-line digital subscriber line (SDSL)
  - — Very-high-data-rate digital subscriber line (VDSL)

  DSL does not use the entire bandwidth available on the twisted pair, leaving room for a voice channel.
- **Cable**—Uses a coaxial cable to transport the data.

WAN Cabling

The router end of the cable connects to the DB-60 port on a serial WAN interface card (using a DB-60 connector). The connector on the other end of the serial cable is specified according to the standard used.

The ports on either end of a WAN connection are specified as DTE (data terminal equipment) or data communications equipment (DCE). DCE converts user data into the service provider's preferred format. The port configured as DTE requires external clocking from the CSU/DSU or another DCE device.

Physical Parameters for WAN Connections

A WAN service provider assigns your organization the parameters required for making the WAN link connection.

- Customer premises equipment (CPE) is located on the subscriber's premises. It includes both equipment owned by the subscriber and devices leased by the service provider.
- Demarcation, or demarc, marks the point where CPE ends and the local loop begins. Usually, it is located in the telecommunications closet.
- Local loop, or "last-mile," is the cabling from the demarc into the WAN service provider's central office (CO).

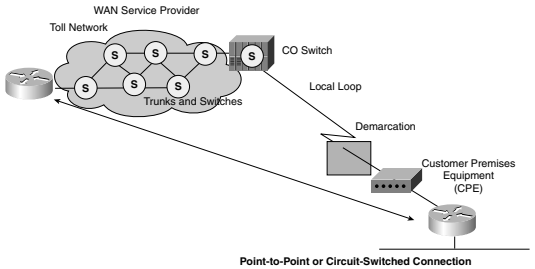

- The central office is a switching facility that provides a point of presence for WAN service. The central office is the entry point to the WAN cloud and the exit point from the WAN for called devices.
- A switching point for calls.
- The toll network is a collection of trunks inside the WAN cloud.

Layer 2 Encapsulation Protocols

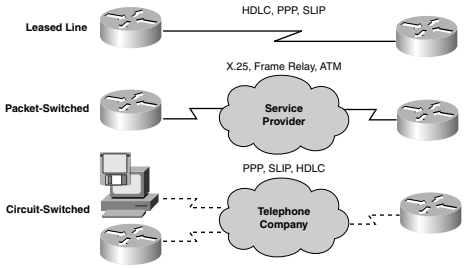

High-level data link control (HDLC) is the default encapsulation type on point-to-point dedicated links and circuit-switched connections. HDLC should be used for communication between Cisco devices.

Point-to-Point Protocol (PPP) provides connections between devices over several types of physical interfaces, such as asynchronous serial, HSSI, ISDN, and synchronous. PPP works with several network layer protocols, including IP and IPX. PPP uses PAP and CHAP for basic security.

X.25/Link Access Procedure, Balanced (LAPB) defines connections between DTE and DCE for remote terminal access. LAPB is a data link layer protocol specified by X.25.

Frame Relay is the industry-standard switched data link layer protocol. Frame Relay (based on X.25) can handle multiple virtual circuits.

Asynchronous Transfer Mode (ATM) is the international standard for cell relay using fixed-length (53-byte) cells for multiple service types. Fixed-length cells allow hardware processing, which greatly reduces transit delays. ATM takes advantage of high-speed transmission media, such as E3, T3, and SONET.

WAN Concepts and Terminology Summary

- WANs connect devices across broad geographic regions. Companies use WANs to connect various sites.
- Leased-line or point-to-point connections provide a dedicated connection.
- Circuit-switched connections provide a dedicated circuit path for the duration of the call. Circuit switching is best for sporadic WAN usage.
- Packet-switched connections use virtual circuits to provide end-to-end connectivity.
- The five serial standards supported by Cisco devices are EIA/TIA-232, EIA/TIA-449, V.35, X.21, and EIA/TIA-530.
- Typical WAN protocols include HDLC, PPP, SLIP, and ATM.

Configuring HDLC and PPP Encapsulation

HDLC is a data link protocol used on synchronous serial data links. HDLC cannot support multiple protocols on a single link because it lacks a mechanism to indicate which protocol it is carrying.

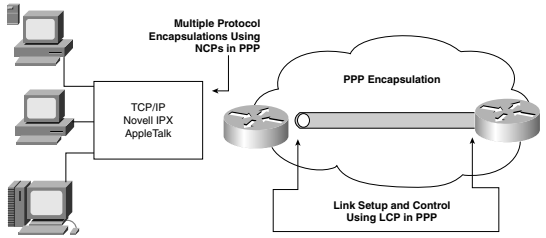

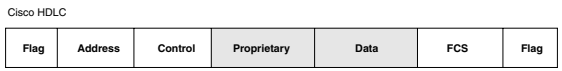

The Cisco version of HDLC uses a proprietary field that acts as a protocol field. This field makes it possible for a single serial link to accommodate multiple network-layer protocols. Cisco's HDLC is a point-to-point protocol that can be used on leased lines between two Cisco devices. PPP should be used when communicating with non-Cisco devices.

To change the encapsulation back to HDLC from some other protocol, use the following command from interface configuration mode:

```
Router(config-if)#encapsulation hdlc
```

PPP Encapsulation

PPP uses a Network Control Protocol (NCP) component to encapsulate multiple protocols and uses Link Control Protocol (LCP) to set up and negotiate control options on the data link.

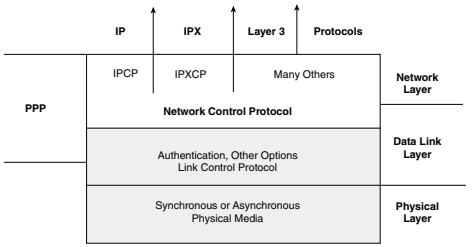

PPP Configuration Options

Cisco routers using PPP encapsulation include the LCP options shown in the following table.

Feature	How It Operates	Protocol
Authentication	Requires a password; performs challenge handshake	PAP CHAP
Compression	Compresses data at source; reproduces data at destination	Stacker or protocol
Error detection	Monitors data dropped on link; avoids frame looping	Magic Number
Multilink	Load balancing across multiple links	Multilink Protocol (MP)

- Authentication options are Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP).
- Compression options increase the effective throughput on PPP connections.
- For error detection, the quality and magic number options help ensure a reliable, loop-free data link.
- Multilink is available in Cisco IOS Release 11.1 and later. It improves throughput and reduces latency between peer routers.
- PPP callback is available in Cisco IOS Release 11.1. It offers enhanced security. After making the initial DDR call, the router requests that it be called back and then terminates its call.

Establishing a PPP Session

The three phases of PPP session establishment are link establishment, authentication, and network protocol.

- **Link establishment—**Each PPP device sends LCP packets to configure and test the data link. Options such as maximum receive unit, compression, and link authentication are negotiated here. Default values are assumed when no figures are present.
- **Authentication (optional)—**After the link is established, the peer can be authenticated.
- **Network layer protocol—**NCP packets are used to select and configure network layer protocols. After they are configured, the network layer protocols can begin sending datagrams over the link.

PPP Authentication Protocols

PPP Authentication Protocol is a simple two-way handshake that's used to establish a remote node's identity. It takes place after the PPP link is established. The remote node repeatedly sends its username and password to the router until authentication is acknowledged or the connection is terminated.

CHAP is a three-way handshake that takes place at link startup and periodically throughout the session to verify the remote node's identity. After the PPP link is established, the local router sends a challenge message to the remote node. The remote node responds with a calculated value (typically, an MD5 function is used). The local router checks the response against its own calculated value. If the values match, the authentication is acknowledged. Otherwise, the connection is terminated immediately.

How Secure Is PAP/CHAP?

With PAP, passwords are sent across the link without encryption or protection against trial-and-error attacks. This level of security is usually sufficient for token-type passwords that change with each authentication.

CHAP uses unpredictable challenge values, which limit exposure to attacks. Local router or authentication servers (TACACS) control the challenges' frequency and timing.

PPP Encapsulation and Authentication Overview

You must do the following before enabling PAP or CHAP:

- Enable PPP protocol encapsulation on each router.
- Assign a host name to each router.
- Define a remote username and password for each router to accept the authentication process.

Here's a CHAP configuration example:

```
RouterA>enable
RouterA#configterm
RouterA(config)#hostname flanders
RouterA(config)#username ned password maude
RouterA(config)#interface serial 0
RouterA(config-if)#encapsulation ppp
```

```
RouterA(config-if)#ppp authentication chap
RouterA(config-if)#exit
RouterA(config)#exit
RouterA#show interface s0
```

password must be the same for both routers using CHAP.

To encrypt passwords, enter the **service password-encryption** command while in global configuration mode.

Configuring HDLC and PPP Encapsulation Summary

- HDLC is the default protocol on serial data links for Cisco devices. Cisco's proprietary HDLC supports multiprotocol environments.
- PPP encapsulates Layer 3 data over point-to-point links.
- LCP options for PPP define authentication, passwords and challenge handshakes, compression, error detection, and multilink parameters.
- The three PPP session establishment phases are link establishment, authentication, and network layer configuration.
- PPP authentication includes PAP, a simple two-way handshake conducted only upon initial establishment, and CHAP, a three-way password-based handshake done at link establishment and periodically throughout the session.

ISDN BRI Concepts

Integrated Services Digital Network (ISDN) is a collection of standards that define an integrated voice/data architecture over the Public Switched Telephone Network (PSTN). ISDN standards define the hardware and call setup schemes. ISDN provides the following benefits:

- **Multiple traffic feeds**—Voice, video, telex, and packet-switched data are all available over ISDN.
- **Fast call setup**—ISDN uses out-of-band (D, or delta channel) signaling for call setup. ISDN calls can often be set up and completed in less than one second.
- **Fast bearer (B) channel services (64 kbps per channel)**—With multiple B channels (two B channels with BRI), ISDN offers 128 kbps. Leased lines usually provide only 56 kbps in North America.

ISDN Standard Access Methods

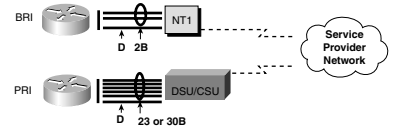

Channel	Capacity	Mostly Used For
B	64 kbps	Circuit-switched data (HDLC, PPP)
D	16/64 kbps	Signaling information (LAPD)

With BRI, there are two bearer (B) channels (6 kbps each) and one delta (D) channel (16 kbps). (BRI is sometimes written as 2B+D.) The B channels are used for digitized voice and high-speed data transport. The D channel is used for signaling. The D channel can also be used for low-rate packet data (such as alarms). D channel traffic is transported using the LAPD data link layer protocol.

In North America and Japan, Primary Rate Interface (PRI) has 23 B channels and 1 D channel (all channels are 64 kbps). In Europe, PRI has 30 B channels and 1 D channel.

ISDN Call Setup

The D channel initiates the call by establishing a path between switches and passing the called number. Local switches use the SS7 signaling protocol to complete the path and pass the called number to the terminating ISDN switch. When the destination receives the setup information, it uses the D channel to signal to the ISDN switch that is available. The B channel is now connected end-to-end and can carry conversation or data.

ISDN Functions

Customer premises equipment (CPE) connects to the ISDN switch. The ISDN standards define functions (devices) that act as transition points between reference-point interfaces. With BRI, you must determine whether you need a transition device (NT1) between the router and the service provider's ISDN switch. Connectors labeled as BRI U have a built-in NT1. Connectors marked BRI S/T require an external NT1.

Caution: Insert the cable running from an ISDN BRI port in only an ISDN jack or switch. ISDN BRI uses voltages that can seriously damage non-ISDN devices.

ISDN Device Types and Reference Points

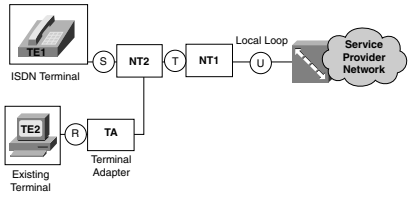

- **TE1 (terminal endpoint 1)**—Devices have a native ISDN interface.
- **NT2 (network termination 2)**—Aggregates and switches all ISDN lines at the customer service site using a customer switching device.
- **NT1 (network termination 1)**—Converts BRI signals into a form used by the ISDN digital line. An NT1 terminates the local loop.
- **TE2 (terminal endpoint 2)**—A device that requires a TA.
- **TA (terminal adapter)**—Converts EIA/TIA-232, V.35, and other signals into BRI signals.
- **R**—The connection point between a non-ISDN-compatible device and a terminal adapter.
- **S**—The connection point into the customer switching device (NT2). Enables calls between customer equipment.
- **T**—The outbound connection from the NT2 to the ISDN network. This reference point is electrically identical to the S interface.
- **U**—The connection point between NT1 and the ISDN network.

Determining the Router ISDN Interface

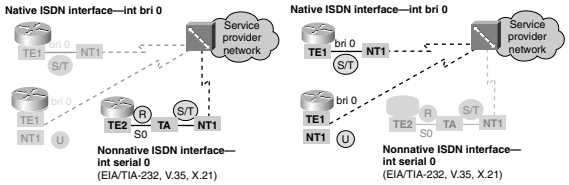

Cisco routers might not have a native ISDN terminal, and those with terminals might not have the same reference point. To avoid damaging equipment, you need to evaluate each router carefully.

Connectors labeled BRI have a native ISDN interface built in. (Your router is a TE1.) A router might also have a built-in NT1 (BRI U interface). If your router interface is labeled BRI, you must use an external TA device.

Warning: Never connect a router with a U interface to an NT1. It will most likely ruin the interface!

ISDN Switch Types

Service providers use several different types of switches for their ISDN services. Before connecting a router to an ISDN service, you must be aware of the switch types used at the central office. You must specify this information during router configuration to allow ISDN service.

Service Provider Identifiers

Service Provider Identifiers (SPIDs) can be assigned by your service provider to identify your switch at the central office. Your switch must be identified before a connection can be made (during call setup).

The syntax for configuring a SPID on your switch is as follows:

```
isdn [spid1 | spid2] spid-number [ldn]
```

[spid1 | spid2] specifies SPID as either the first or second B channel, *spid-number* is the number assigned by the ISDN service provider, and *ldn* is an optional local dial number.

ISDN BRI Concepts Summary

- ISDN standards define a digital architecture for integrated voice/data capability using the public switched telephone network. ISDN provides multiple user-traffic feeds, fast call setup, and fast data transfer rates.
- ISDN protocols include the E-series protocol for the telephone network and ISDN; the I-series protocol for ISDN concepts, aspects, and interfaces; and the Q-series protocol for switching and signaling.
- ISDN BRI has 2 64 kbps B channels and 1 16 kbps D channel.
- ISDN PRI has 23 B channels and 1 D channel.
- Reference points define connection points between functions.
- SPIDs are required before you can access the ISDN network.

Dial-on-Demand Overview

Dial-on-demand routing (DDR) refers to an as-needed connection service over the PSTN. DDR is typically used for low-volume, periodic connections. It can offer substantial savings over traditional WAN connections. DDR is well-suited for telecommuters, satellite offices sending sales transactions or order requests, and automated customer order systems.

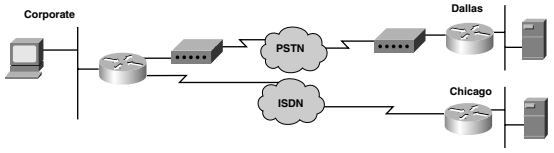

DDR Process

DDR uses the concept of "interesting traffic" to determine when a dialup connection should be made. Interesting traffic is defined in a router lookup table. When traffic is defined as interesting, the router locates the next-hop router and any dialing instructions (located in a dialer map). If the link is not already up, the router initiates a connection, and traffic is sent. As soon as a link is enabled, the router transmits both interesting and uninteresting traffic. The call is terminated if no more additional interesting traffic is sent within a specified time period.

Defining Static Routes for DDR

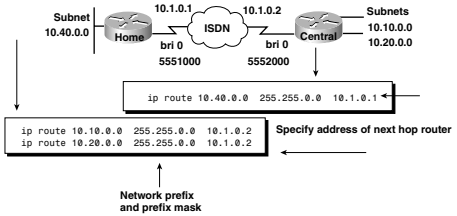

Static routes are a necessity for DDR, because you want to maintain exact control over which routes are used to reach each destination. These routes must be manually configured on all participating routers, because static routes have no routing updates. To manually configure a route, use the following command:

```
ip route [network prefix] [prefix mask] {address | interface}
  [distance [permanent]]
```

network prefix is the address of the destination network, *address* | *interface* are the address and interface of the next-hop router, and **permanent** sets the static condition.

Specifying Interesting Traffic for DDR

Interesting packets are determined by the network administrator. They can be defined by protocol type, source address, or destination host. Use the following command to define interesting packets:

```
dialer-list dialer-group protocol protocol-name {permit | deny |
  list access-list-number}
```

dialer-group maps the dialer list to an interface, and **list** *access-list-number* assigns an access list to the dialer group.

Other important DDR commands are **dialer-group**, which links interesting traffic created in the **dialer-list** command to the interface, and **dialer-map**, which defines one or more dial-on-demand numbers.

Dial-on-Demand Summary

- DDR refers to dynamic connections made over dialup facilities on an as-needed basis.
- DDR is best-suited for low-volume, periodic connections.
- To configure legacy DDR, define static routes (**ip route** command), specify interesting traffic (**dialer-list** command), and configure the dialer information (**dialer-group** command).
- All participating routers must have static routes defined to reach the remote networks.

Frame Relay Overview

Frame Relay is a connection-oriented Layer 2 protocol that allows several data connections (called virtual circuits) to be multiplexed onto a single physical link. Frame Relay relies on upper-layer protocols for error correction. Frame Relay specifies only the connection between a router and a service provider's local access switching equipment. The data transmission within the service provider's Frame Relay cloud is not specified.

A connection identifier is used to map packets to outbound ports on the service provider's switch. When the switch receives a frame, a lookup table is used to map the frame to the correct outbound port. The entire path to the destination is determined before the frame is sent.

Frame Relay Stack

Most Frame Relay functions exist at the lower two layers of the OSI Reference Model. Frame Relay is supported on the same physical serial connections that support point-to-point connections. Cisco routers support the following serial connections: EIA/TIA-232, EIA/TIA-449, V.35, X.21, EIA/TIA-530.

Upper-layer information (such as IP data) is encapsulated by Frame Relay and is transmitted over the link.

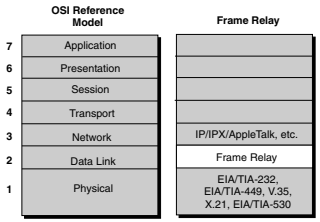

Frame Relay Terms

- **BECN (Backward Explicit Congestion Notification)**—A message sent to a source router when a Frame Relay switch recognizes congestion in the network. A BECN message requests a reduced data transmission rate.
- **CIR (Committed Information Rate)**—The minimum guaranteed data transfer rate agreed to by the Frame Relay switch.
- **DLCI (Data Link Connection Identifier)**—Identifies the logical circuit between the router and the Frame Relay switch.
- **FECN (Forward Explicit Congestion Notification)**—A message sent to a destination device when a Frame Relay switch senses congestion in the network.
- **Inverse ARP**—Routers use Inverse ARP to discover the network address of a device associated with a VC.
- **LMI (Local Management Interface)**—A signaling standard used to manage the connection between the router and the Frame Relay switch. LMIs track and manage keepalive mechanisms, multicast messages, and status. LMI can be configured (in Cisco IOS 11.2 and later), but routers can autosense LMI types by sending a status request to the Frame Relay switch. The router configures itself to match the LMI type response. The three types of LMIs supported by Cisco Frame Relay switches are Cisco (developed by Cisco, StrataCom, Northern Telecom, and DEC), ansi Annex D (ANSI standard T1.617), and q933a (ITU-T Q.933 Annex A).
- **VC (virtual circuit)**—A logical circuit between two network devices. A VC can be permanent (PVC) or switched (SVC). PVCs save bandwidth (there is no circuit establishment or teardown) but can be expensive. SVCs are established on-demand and are torn down when transmission is complete. VC status can be active, inactive, or deleted.

Dynamic Mapping with Inverse ARP

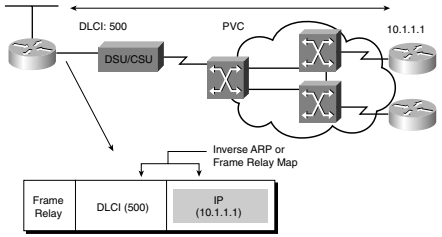

To correctly route packets, each DLCI must be mapped to a next-hop address. These addresses can be dynamically mapped using Inverse ARP or can be manually configured. After the address is mapped, it is stored in the router's Frame Relay map table.

LMI Signaling Process

1. The router connects to a Frame Relay switch through a channel service unit/data service unit (CSU/DSU).
2. The router sends a VC status inquiry to the Frame Relay switch.
3. The switch responds with a status message that includes the DLCI's information for the usable PVCs.
4. The router advertises itself by sending an Inverse ARP to each active DLCI.
5. The routers create map entries with the local DLCI and network-layer address of the remote routers. Static maps must be configured if Inverse ARP is not supported.
6. Inverse ARP messages are sent every 60 seconds.
7. LMI information is exchanged every 10 seconds.

Frame Relay Overview Summary

- Frame Relay is a connection-oriented Layer 2 protocol that allows several data connections (VCs) to be multiplexed onto a single physical link.
- Cisco routers support Frame Relay on the following types of serial connections: EIA/TIA-232, EIA/TIA-449, V.35, X.21, and EIA/TIA-530.
- Local DLCI addresses can be dynamically mapped using Inverse ARP or manually configured using static Frame Relay maps.

- Local Management Interface (LMI) signaling is used by Frame Relay switches to manage connections and maintain status between the devices. The supported LMI types are cisco, ansi, and q933a.

Configuring Frame Relay

The three commands used to configure basic Frame Relay on a router select the Frame Relay encapsulation type, establish the LMI connection, and enable Inverse ARP:

```
encapsulation frame-relay [cisco | ietf]
frame-relay lmi-type {ansi | cisco | q933i}
frame-relay inverse-arp [protocol] [dlci]
```

Here's the procedure for configuring basic Frame Relay:

```
RouterA>enable
RouterA#config term
RouterA(config)#int ser 1
RouterA(config-if)#address 10.16.0.1  255.255.255.255
RouterA(config-if)#encapsulation frame-relay cisco
RouterA(config-if)#frame-relay lmi-type   cisco
RouterA(config-if)#bandwidth 64
RouterA(config-if)#frame-relay inverse-arp ip 16
RouterA(config-if)#exit
RouterA(config)#exit
RouterA#
```

Verifying Frame Relay Operations

The following commands verify and display Frame Relay information:

- **show interface**—Displays Layer 1 and Layer 2 status, DLCI information, and the LMI DLCIs used for the local management interface.
- **show frame-relay lmi**—Displays LMI traffic statistics (LMI type, status messages sent, and invalid LMI messages).
- **show frame-relay pvc**—Displays the status of all configured connections, traffic statistics, and BECN FECN packets received by the router.
- **show frame-relay map**—Displays the current map entries for static and dynamic routes. The **frame-relay-inarp** command clears all dynamic entries.

Static Frame Relay Map Configuration

A router's address-to-DLCI table can be defined statically when inverse ARP is not supported. These static maps can also be used to control broadcasts. To statically configure the map table, use the following command:

```
frame-relay map protocol protocol-address dlci [broadcast]
  [ietf | cisco| payload-compression packet-by-packet]
```

protocol specifies bridging or logical link control, *protocol-address* is the network layer and address of the destination device, *dlci* is the local dlci, *broadcast* is an optional parameter used to control broadcasts and multicasts over the VC, and **payload compression** is an optional Cisco proprietary compression method.

Frame Relay Topology

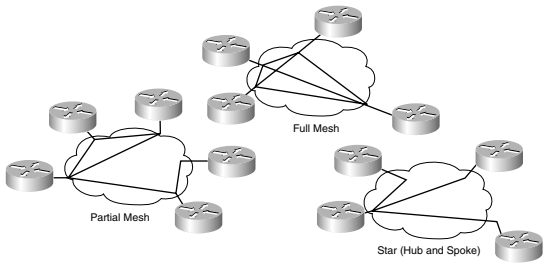

Frame Relay is a nonbroadcast multiaccess (NBMA) connection scheme. This means that although Frame Relay interfaces support multipoint connections by default, broadcast routing updates are not forwarded to remote sites. Frame Relay networks can be designed using star, full-mesh, and partial-mesh topologies.

A *star topology,* also known as a hub-and-spoke configuration, is the common network topology. Remote sites are connected to a central site, which usually provides a service. Star topologies require the fewest PVCs, making them relatively inexpensive. The hub router provides a multipoint connection using a single interface to interconnect multiple PVCs.

In a *full-mesh topology,* all routers have virtual circuits to all other destinations. Although it is expensive, this method provides redundancy, because all sites are connected to all other sites. Full-mesh networks become very expensive as the number of nodes increases. The number of links required in a full-mesh topology with n nodes is $(n – (n – 1)) / 2$.

In a *partial-mesh topology,* not all sites have direct access to all other sites. Connections usually depend on the traffic patterns within the network.

Configuring Frame Relay Summary

- The **frame-relay**, **frame-relay lmi-type**, and **frame-relay inverse-arp** commands are used to configure Frame Relay.
- The **frame-relay map** command is used to configure static address-to-DLCI tables.
- The three WAN topologies used to interconnect remote sites are star, partial-mesh, and full-mesh.

Configuring Frame Relay Subinterfaces

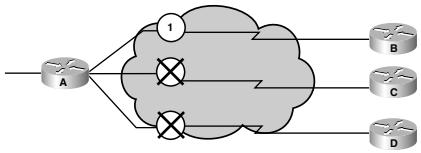

Frame Relay provides NBMA connectivity between sites, which means that although remote locations can reach each other, routing update broadcasts are not forwarded to all locations. Frame Relay networks use split horizon to prevent routing loops. With split horizon activated, if a remote router receives an update on an interface that has multiple PVCs, the router cannot forward that broadcast to routers on other PVCs on the same interface.

Routing Update Replication

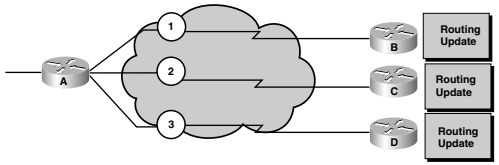

When several DLCIs terminate in a single router, that router must replicate all routing updates and service advertisements on each DLCI. These updates consume bandwidth, cause latency variations, and consume interface buffers, which leads to higher packet rate loss.

You should consider broadcast traffic and virtual circuit placement when designing Frame Relay networks to avoid negatively affecting critical user data.

Resolving Reachability Issues in Frame Relay

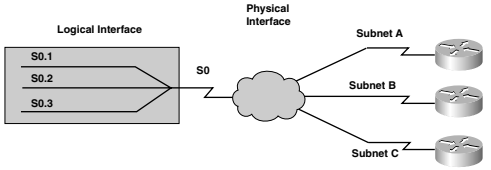

You can resolve reachability issues by configuring subinterfaces on the router. These logically assigned interfaces allow the router to forward broadcast updates in a Frame Relay network. Subinterfaces are logical subdivisions of a physical interface. Routing updates received on one subinterface can be sent out another subinterface without violating split horizon rules. If you configure virtual circuits as point-to-point connections, the subinterface acts similar to a leased line.

Subinterface Configuration

Subinterfaces can be configured as either point-to-point or multipoint. With point-to-point, one PVC connection is established with another physical interface or subinterface on a remote router using a single subinterface. With multipoint, multiple PVC connections are established with multiple physical interfaces or subinterfaces on remote routers on a single subinterface. All interfaces involved use the same subnet, and each interface has its own local DLCI.

To configure a subinterface, use the following command:

```
frame-relay interface-dlci dlci-number
```

To select a subinterface, use the following command:

```
interface serial number.subinterface-number {multipoint |
  point-to-point}
```

The range of subinterface numbers is to 4294967293. The number that precedes the period (.) must match the physical interface number to which this subinterface belongs.

To configure the local DLCI on the subinterface, use the following command:

```
frame-relay interface-dlci dlci-number
```

The *dlci-number* defines the local DLCI number being linked to the subinterface. This is the only way to link an LMI-derived PVC to a subinterface. (LMI does not know about subinterfaces.)

Here is the procedure for configuring basic Frame Relay:

```
RouterA>enable
RouterA#config term
RouterA(config)#int ser0
RouterA(config-if)#no ip address
RouterA(config-if)#encapsulation frame-relay cisco
RouterA(config-if)#interface serial0.3 point-to-point
RouterA(config-if)#frame-relay inverse-dlci 120
RouterA(config-if)#exit
RouterA(config)#exit
RouterA#
```

The **frame-relay inverse-dlci 120** command is required for all point-to-point configurations and multipoint subinterfaces for which Inverse ARP is enabled. Do not use this command on physical interfaces.

Configuring Frame Relay Subinterfaces Summary

- Split horizon does not allow routing updates received on one interface to be forwarded out the same interface.
- Routing updates received on one subinterface can be sent out another subinterface configured on the same physical interface.
- Virtual circuits can be configured as point-to-point connections, allowing subinterfaces to act like leased lines.
- Subinterfaces can be configured to support point-to-multipoint or point-to-point connection types.