

Introduction to Switched Networks

Objectives

Upon completion of this chapter, you will be able to answer the following questions:

- How do you describe the convergence of data, voice, and video in the context of switched networks?
- How do you describe a switched network in a small- to medium-sized business?
- How do you explain the process of frame forwarding in a switched network?
- How do you compare a collision domain to a broadcast domain?

Key Terms

This chapter uses the following key terms. You can find the definitions in the Glossary.

Converged network page 3

Call control page 6

Voice messaging page 6

Mobility page 6

Automated attendant page 6

Cisco Borderless Network page 6

Hierarchical page 8

Modularity page 8

Resiliency page 8

Flexibility page 8

Access page 8

Distribution page 8

Core page 8

Form factor page 11

Fixed configuration switch page 14

Modular configuration switch page 14

Stackable configuration switch page 15

Traffic flow analysis page 16

Multilayer switch page 16

Port density page 17

Small form-factor pluggable (SFP)
page 18

Forwarding rate page 19

Power over Ethernet (PoE) page 19

Frame forwarding page 23

Ingress port page 24

Egress port page 25

MAC address table page 25

Store-and-forward switching page 29

Cut-through switching page 29

Fragment free switching page 31

Collision domain page 32

Broadcast domain page 32

Introduction (1.0.1.1)

Modern networks continue to evolve to keep pace with the changing way that organizations carry out their daily business. Users now expect instant access to company resources from anywhere and at any time. These resources not only include traditional data but also video and voice. There is also an increasing need for collaboration technologies that allow real-time sharing of resources between multiple remote individuals as though they were at the same physical location.

Different devices must seamlessly work together to provide a fast, secure, and reliable connection between hosts. LAN switches provide the connection point for end users into the enterprise network and are also primarily responsible for the control of information within the LAN environment. Routers facilitate the movement of information between LANs and are generally unaware of individual hosts. All advanced services depend on the availability of a robust routing and switching infrastructure on which they can build. This infrastructure must be carefully designed, deployed, and managed to provide a necessary stable platform.

This chapter begins an examination of the flow of traffic in a modern network. It examines some of the current network design models and the way that LAN switches build forwarding tables and use the MAC address information to efficiently switch data between hosts.



Class Activity 1.0.1.2: Sent or Received Instructions

Individually, or in groups (per the instructor's decision), discuss various ways that hosts send and receive data, voice, and streaming video.

Develop a matrix (table) listing network data types that can be sent and received. Provide five examples.

Note

For an example of the matrix, see the document prepared for this modeling activity.

Save your work in either hard- or soft-copy format. Be prepared to discuss your matrix and statements in a class discussion.

LAN Design (1.1)

In this section, you will explore the design of local-area networks. The Cisco Borderless Network architecture for delivery of services and applications provides a setting for the exploration of switched network design. And you will learn how the fundamental core-distribution-access model applies to switched networks.

Converged Networks (1.1.1)

Converged networks were cutting edge ten years ago, but now they are standard fare for switched environments. The integration of voice, video, and data on a switched infrastructure provides a seamless experience for users. IP phones and video devices are fully integrated into the data network.

Growing Complexity of Networks (1.1.1.1)

Our digital world is changing. The ability to access the Internet and the corporate network is no longer confined to physical offices, geographical locations, or time zones. In today's globalized workplace, employees can access resources from anywhere in the world, and information must be available at any time and on any device. These requirements drive the need to build next-generation networks that are secure, reliable, and highly available.

Data networks originally served the purpose of transporting data between workstations and servers. As networks became more reliable, voice and video traffic was integrated with data traffic, creating a converged network. A converged network is one where data, voice, and video are integrated. Next-generation converged networks must not only support current expectations and equipment but must also be able to integrate legacy platforms.

Legacy Equipment

Legacy equipment can hinder convergence. Figure 1-1 illustrates legacy telephone equipment. A business site can contain equipment that supports both legacy PBX telephone systems and IP-based phones. This sort of equipment is rapidly migrating toward IP-based phone switches.



Figure 1-1 Legacy Components

Advanced Technology

Although converged networks have existed for some time now, they were initially only feasible in large enterprise organizations because of the network infrastructure and complex management requirements. There were high network costs associated with convergence because more expensive switch hardware was required to support the additional bandwidth. Converged networks also required extensive management in relation to QoS, because voice and video data traffic needed to be classified and prioritized on the network. Few individuals had the expertise in voice, video, and data networks to make convergence feasible and functional.

Over time, convergence has become easier to implement and manage, and less expensive to purchase. Figure 1-2 illustrates some of the newer platforms for converged networks that help to provide access to the network anytime, anywhere, and on any device.



Figure 1-2 Converged Network Components

Elements of a Converged Network (1.1.1.2)

To support collaboration, business networks employ converged solutions using voice systems, IP phones, voice gateways, video support, and videoconferencing, as illustrated in Figure 1-3.

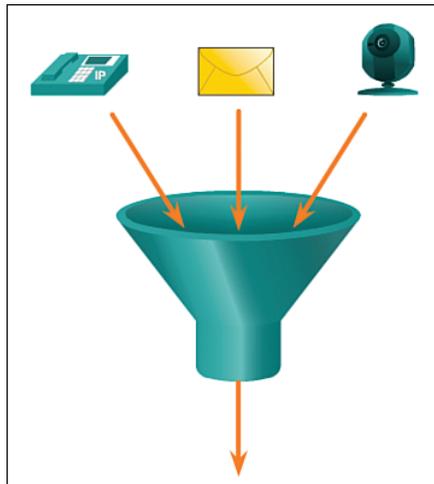


Figure 1-3 Many Types of Traffic on One Network

Including data services, a converged network with collaboration support can include features such as the following:

- **Call control:** Telephone call processing, caller ID, call transfer, hold, and conference
- **Voice messaging:** Voicemail
- **Mobility:** Receive important calls wherever you are
- **Automated attendant:** Serve customers faster by routing calls directly to the right department or individual

One of the primary benefits of transitioning to the converged network is that there is just one physical network to install and manage. This results in substantial savings over the installation and management of separate voice, video, and data networks. Such a converged network solution integrates IT management so that any moves, additions, and changes are completed with an intuitive management interface. A converged network solution also provides PC softphone application support, as well as point-to-point video, so that users can enjoy personal communications with the same ease of administration and use as a voice call.

The convergence of services onto the network has resulted in an evolution in networks from a traditional data transport role to a superhighway for data, voice, and video communication. This one physical network must be properly designed and implemented to allow the reliable handling of the various types of information that it must carry. A structured design is required to allow management of this complex environment.

Video

Video 1.1.1.2: Observing Spanning Tree Protocol Operation

Go to the online course and play the video in the second graphic to view a few of the collaboration services in action.

Cisco Borderless Network (1.1.1.3)

With the increasing demands of the converged network, the network must be developed with an architectural approach that embeds intelligence, simplifies operations, and is scalable to meet future demands. One of the more recent developments in network design is the *Cisco Borderless Network*.

The Cisco Borderless Network is a network architecture combining innovation and design that allows organizations to support a borderless network that can connect anyone, anywhere, anytime, on any device—securely, reliably, and seamlessly. This architecture is designed to address IT and business challenges, such as supporting the converged network and changing work patterns.

The Cisco Borderless Network provides the framework to unify wired and wireless access, including policy, access control, and performance management across many different device types. Using this architecture, the borderless network is built on a hierarchical infrastructure of hardware that is scalable and resilient, as shown in Figure 1-4. By combining this hardware infrastructure with policy-based software solutions, the Cisco Borderless Network provides two primary sets of services: network services and user and endpoint services, all managed by an integrated management solution. It enables different network elements to work together and allows users to access resources from any place at any time, while providing optimization, scalability, and security.

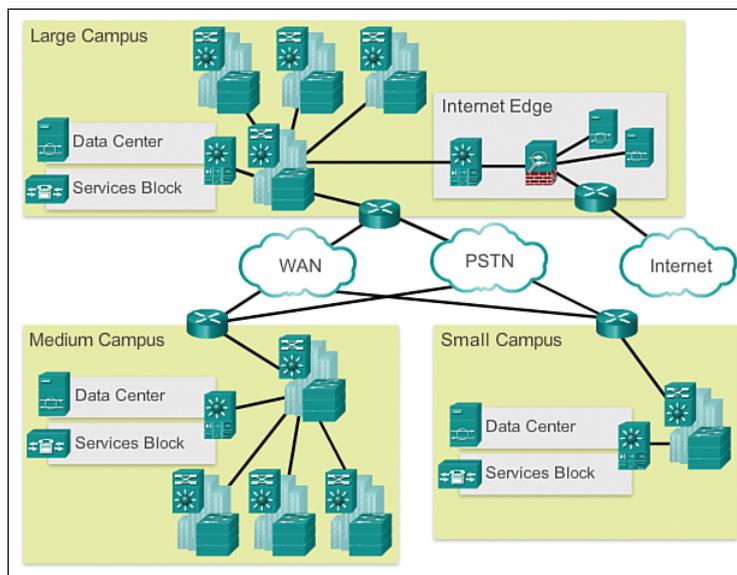


Figure 1-4 Cisco Borderless Network

Video

Video 1.1.1.3: Observing Spanning Tree Protocol Operation

Go to the online course and play the video in the second graphic to learn more about the evolution of the borderless network.

Hierarchy in the Borderless Switched Network (1.1.1.4)

Creating a borderless switched network requires that sound network design principles are used to ensure maximum availability, flexibility, security, and manageability. The borderless switched network must deliver on current requirements and future

required services and technologies. Borderless switched network design guidelines are built upon the following principles:

- **Hierarchical:** Facilitates understanding the role of each device at every tier; simplifies deployment, operation, and management; and reduces fault domains at every tier
- **Modularity:** Allows seamless network expansion and integrated service enablement on an on-demand basis
- **Resiliency:** Satisfies user expectations for keeping the network always on
- **Flexibility:** Allows intelligent traffic load sharing by using all network resources

These are not independent principles. Understanding how each principle fits in the context of the others is critical. Designing a borderless switched network in a hierarchical fashion creates a foundation that allows network designers to overlay security, mobility, and unified communication features. Two time-tested and proven hierarchical design frameworks for campus networks are the three-tier layer model, as shown in Figure 1-5, and the two-tier layer model, as shown in Figure 1-6.

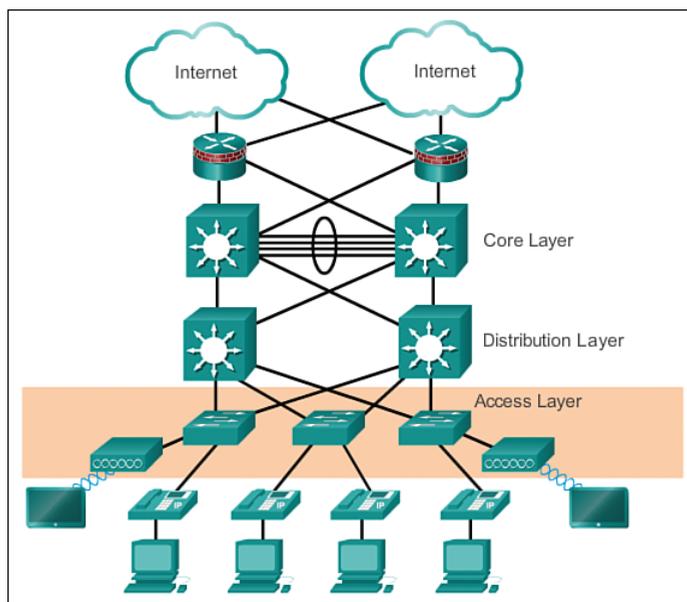


Figure 1-5 Access Layer

The three critical layers within these tiered designs are the *access*, *distribution*, and *core* layers. Each layer can be seen as a well-defined, structured module with specific roles and functions in the campus network. Introducing modularity into the campus

hierarchical design further ensures that the campus network remains resilient and flexible enough to provide critical network services. Modularity also helps to allow for growth and changes that occur over time.

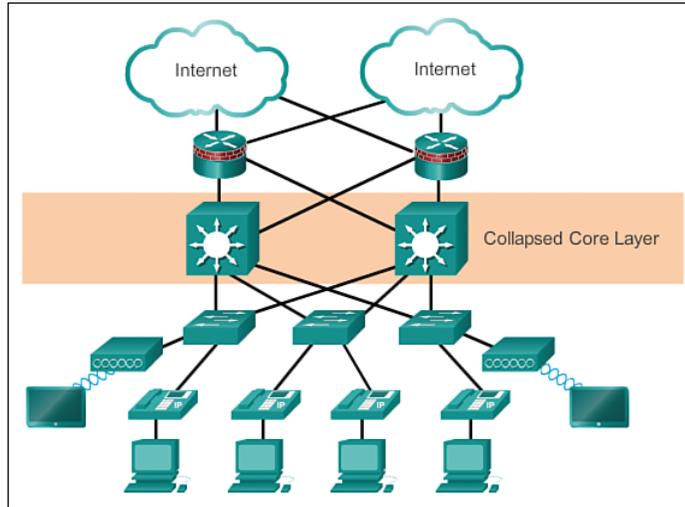


Figure 1-6 Collapsed Core

Access, Distribution, and Core Layers (1.1.1.5)

The access-distribution-core hierarchical network model is the most referenced network model in computer networking. It is simple, but it carries the rudimentary information necessary to convey networking concepts in context.

Access Layer

The access layer represents the network edge, where traffic enters or exits the campus network. Traditionally, the primary function of an access layer switch is to provide network access to the user. Access layer switches connect to distribution layer switches, which implement network foundation technologies such as routing, quality of service, and security.

To meet network application and end-user demand, the next-generation switching platforms now provide more converged, integrated, and intelligent services to various types of endpoints at the network edge. Building intelligence into access layer switches allows applications to operate on the network more efficiently and securely.

Distribution Layer

The distribution layer interfaces between the access layer and the core layer to provide many important functions, including

- Aggregating large-scale wiring closet networks
- Aggregating Layer 2 broadcast domains and Layer 3 routing boundaries
- Providing intelligent switching, routing, and network access policy functions to access the rest of the network
- Providing high availability through redundant distribution layer switches to the end user and equal-cost paths to the core
- Providing differentiated services to various classes of service applications at the edge of the network

Core Layer

The core layer is the network backbone. It connects several layers of the campus network. The core layer serves as the aggregator for all the other campus blocks and ties the campus together with the rest of the network. The primary purpose of the core layer is to provide fault isolation and high-speed backbone connectivity.

Figure 1-7 shows a three-tier campus network design for organizations where the access, distribution, and core are each separate layers. To build a simplified, scalable, cost-effective, and efficient physical cable layout design, the recommendation is to build an extended-star physical network topology from a centralized building location to all other buildings on the same campus.

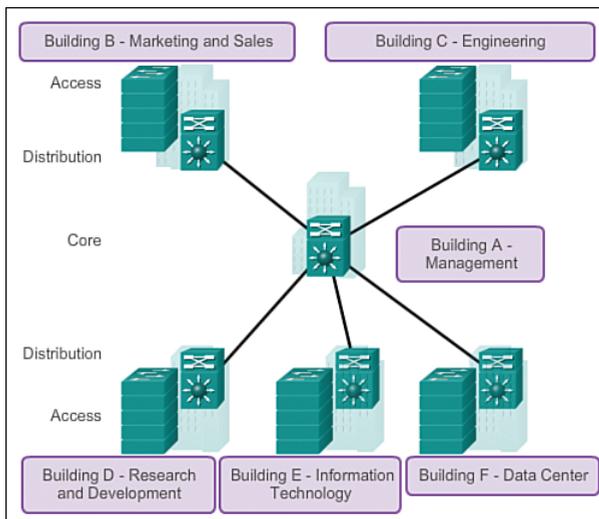


Figure 1-7 Three-Tier Campus Network Design

In some cases where extensive physical or network scalability does not exist, maintaining separate distribution and core layers is not required. In smaller campus locations where there are fewer users accessing the network or in campus sites consisting of a single building, separate core and distribution layers might not be needed. In this scenario, the recommendation is the alternate two-tier campus network design, also known as the collapsed core network design.

Figure 1-8 shows a two-tier campus network design example for an enterprise campus where the distribution and core layers are collapsed into a single layer.

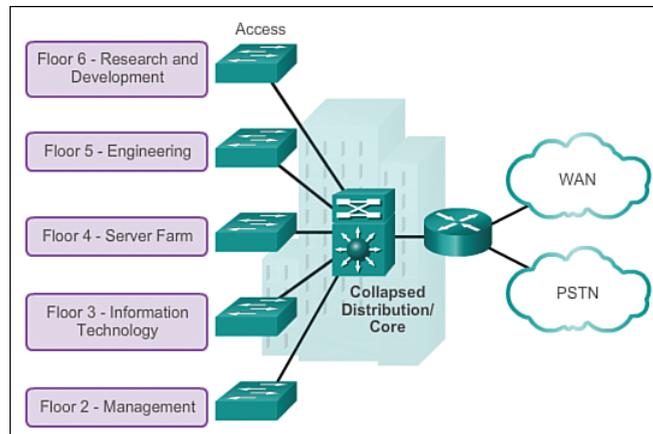


Figure 1-8 Two-Tier Campus Network Design

Interactive
Graphic

Activity 1.1.1.6: Identify Switched Network Terminology

Go to the online course to perform this practice activity.

Switched Networks (1.1.2)

In this topic, you will learn about the various types of switches and their *form factors*. A discussion of multilayer switching will put in context our exploration of access layer switches, which are the focus of this course.

Role of Switched Networks (1.1.2.1)

The role of switched networks has evolved dramatically in the last two decades. It was not long ago that flat Layer 2 switched networks were the norm. Flat Layer 2 data networks relied on the basic properties of Ethernet and the widespread use of hub repeaters to propagate LAN traffic throughout an organization. As shown in Figure 1-9, networks have fundamentally changed to switched LANs in a hierarchical network.

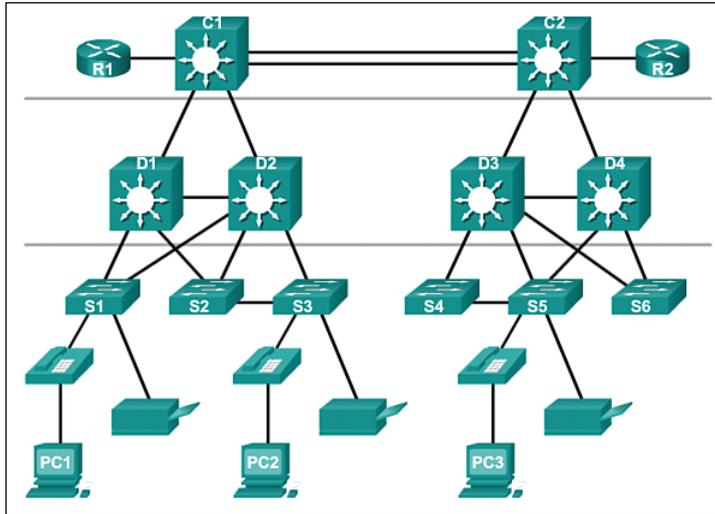


Figure 1-9 Hierarchical Networks

A switched LAN allows more flexibility, traffic management, and additional features, such as

- Quality of service
- Additional security
- Support for wireless networking and connectivity
- Support for new technologies, such as IP telephony and mobility services
- Layer 3 functionality

Figure 1-10 shows the hierarchical design used in the borderless switched network.

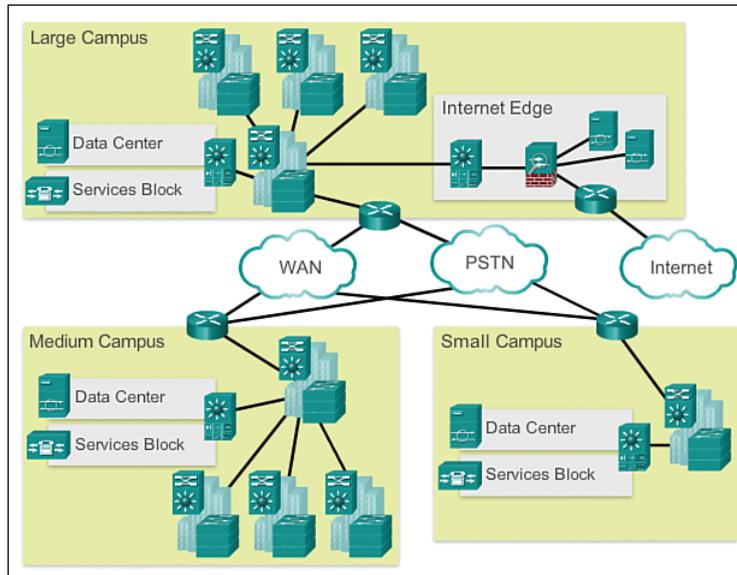


Figure 1-10 Borderless Switched Network

Form Factors (1.1.2.2)

There are various types of switches used in business networks. It is important to deploy the appropriate types of switches based on network requirements. Here are some common business considerations when selecting switch equipment:

- **Cost:** The cost of a switch will depend on the number and speed of the interfaces, supported features, and expansion capability.
- **Port Density:** Network switches must support the appropriate number of devices on the network.
- **Power:** It is now common to power access points, IP phones, and even compact switches using Power over Ethernet (PoE). In addition to PoE considerations, some chassis-based switches support redundant power supplies. PoE will be explored in Section 1.1.3.3.
- **Reliability:** The switch should provide continuous access to the network.
- **Port Speed:** The speed of the network connection is of primary concern to end users.
- **Frame Buffers:** The ability of the switch to store frames is important in a network where there might be congested ports to servers or other areas of the network.

- **Scalability:** The number of users on a network typically grows over time; therefore, the switch should provide the opportunity for growth.

When selecting the type of switch, the network designer must choose between a fixed or a modular configuration, and stackable or nonstackable. Another consideration is the thickness of the switch, which is expressed in number of rack units. This is important for switches that are mounted in a rack. For example, the *fixed configuration switches* shown in Figure 1-11 are all 1 rack unit (1U). These options are sometimes referred to as switch form factors.



Figure 1-11 Fixed Configuration Switches

Fixed Configuration Switches

Fixed configuration switches do not support features or options beyond those that originally came with the switch. The particular model determines the features and options available; features and options are limited to those that originally come with the switch. For example, a 24-port gigabit fixed switch cannot support additional ports. There are typically different configuration choices that vary in how many and what types of ports are included with a fixed configuration switch.

Modular Configuration Switches

Modular configuration switches offer more flexibility in their configuration. Modular configuration switches typically come with different-sized chassis that allow for the installation of different numbers of modular line cards, as shown in Figure 1-12. The line cards actually contain the ports. The line card fits into the switch chassis the way that expansion cards fit into a PC. The larger the chassis, the more modules it can support. There can be many different chassis sizes to choose from. A modular switch with a single 24-port line card could have an additional 24-port line card added to bring the total number of ports up to 48.



Figure 1-12 Modular Switches

Stackable Configuration Switches

Stackable configuration switches can be interconnected using a special cable that provides high-bandwidth throughput between the switches, as shown in Figure 1-13. Cisco StackWise technology allows the interconnection of up to nine switches. Switches can be stacked one on top of the other with cables connecting the switches in a daisy-chain fashion. The stacked switches effectively operate as a single larger switch. Stackable switches are desirable where fault tolerance and bandwidth availability are critical and a modular switch is too costly to implement. Using cross-connected connections, the network can recover quickly if a single switch fails. Stackable switches use a special port for interconnections. Many Cisco stackable switches also support StackPower technology, which enables power sharing among stack members.



Figure 1-13 Stackable Switches

Traffic Flow (1.1.2.3)

To select the appropriate switch for a network, you need to have specifications that detail the target traffic flows. Companies need a network that can meet evolving requirements. A business might start with a few PCs interconnected so that they can share data. As the business adds more employees, devices—such as PCs, printers, and

servers—are added to the network. Accompanying the new devices is an increase in network traffic. Some companies also rely on converged VoIP phone systems, which add more traffic.

To select the appropriate switches, it is important to perform and record traffic flow analyses regularly. *Traffic flow analysis* is the process of measuring the bandwidth usage on a network and then analyzing the data for performance tuning, capacity planning, and making hardware improvement decisions. Analyzing the various traffic sources and their impact on the network allows you to more accurately tune and upgrade the network to achieve the best possible performance.

There are many ways to monitor traffic flow on a network. Individual switch ports can be manually monitored to record bandwidth utilization over time. Traffic flow analysis tools can automatically record traffic flow data in a database and perform an associated trend analysis. While the software is collecting data, you can see how every interface is performing at any given point in time on the network. This gives the network administrator a visual means of identifying traffic flow patterns.

Multilayer Switching (1.1.2.4)

Multilayer switches are typically deployed in the core and distribution layers of an organization's switched network. Multilayer switches are characterized by their ability to build a routing table, support a few routing protocols, and forward IP packets at a rate close to that of Layer 2 forwarding. Multilayer switches often support specialized hardware, such as application-specific integrated circuits (ASIC). ASICs, along with dedicated software data structures, can streamline the forwarding of IP packets independent of the CPU.

There is a trend in networking toward a pure Layer 3 switched environment. When switches were first used in networks, none of them supported routing; now, almost all switches support routing. It is likely that soon all switches will incorporate a route processor because the cost of doing so is decreasing relative to other constraints. Eventually the term *multilayer switch* will be redundant.

The Catalyst 2960 switches shown in Figure 1-14 illustrate the migration to a pure Layer 3 environment. With IOS Releases prior to 15.x, these switches supported only one active switched virtual interface (SVI). With IOS Release 15.x, these switches now support multiple active SVIs, as well as support for static routes! This means that the switch can be remotely accessed through multiple IP addresses on distinct networks.



Figure 1-14 Cisco Catalyst 2960 Series Switches

Packet Tracer
Activity

Packet Tracer Activity 1.1.2.5: Comparing 2960 and 3560 Switches

In this activity, you will use various commands to examine three different switching topologies and compare the similarities and differences between the 2960 and 3560 switches. You will also compare the routing table of a 1941 router with a 3560 switch.

Switch Features (1.1.3)

Relative to routers, the features associated with a switch or a product line of switches vary dramatically. It is important for a switch administrator to understand the features available so that well-informed switch-purchasing decisions are made for an organization.

Port Density (1.1.3.1)

The *port density* of a switch refers to the number of ports available on a single switch. Figure 1-15 shows the port density of three different switches.

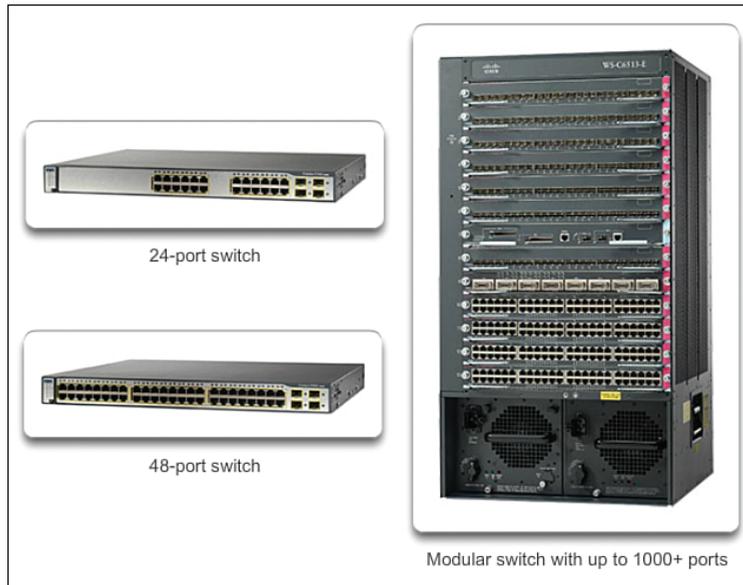


Figure 1-15 Port Densities

Fixed configuration switches typically support up to 48 ports on a single device. They have options for up to four additional ports for *small form-factor pluggable (SFP)* devices. High port densities allow for better use of limited space and power. If there are two switches that each contain 24 ports, they would be able to support up to 46 devices, because at least one port per switch is lost with the connection of each switch to the rest of the network. In addition, two power outlets are required. Alternatively, if there is a single 48-port switch, 47 devices can be supported, with only one port used to connect the switch to the rest of the network and only one power outlet needed to accommodate the single switch.

Modular switches can support very high port densities through the addition of multiple switch port line cards. For example, some Catalyst 6500 switches can support in excess of 1000 switch ports.

Large enterprise networks that support many thousands of network devices require high-density, modular switches to make the best use of space and power. Without using a high-density modular switch, the network would need many fixed configuration switches to accommodate the number of devices that need network access. This approach can consume many power outlets and a lot of closet space.

The network designer must also consider the issue of uplink bottlenecks. For example, to achieve target performance, a series of fixed configuration switches might require many ports for bandwidth aggregation between switches. With a single modular switch, bandwidth aggregation is less of an issue, because the backplane of the chassis can provide the necessary bandwidth to accommodate the devices connected to the switch port line cards.

Forwarding Rates (1.1.3.2)

Forwarding rates define the processing capabilities of a switch by rating how much data the switch can process per second. Switch product lines are classified by forwarding rates, as shown in Figure 1-16. Entry-level switches have lower forwarding rates than enterprise-level switches. Forwarding rates are important to consider when selecting a switch. If the switch forwarding rate is too low, it cannot accommodate full wire-speed communication across all of its switch ports. Wire speed is the data rate that each Ethernet port on the switch is capable of attaining. Data rates can be 100 Mb/s, 1 Gb/s, 10 Gb/s, or 100 Gb/s.

For example, a typical 48-port Gigabit Ethernet switch operating at full wire speed generates 48 Gb/s of traffic. If the switch only supports a forwarding rate of 32 Gb/s, it cannot run at full wire speed across all ports simultaneously. Fortunately, access layer switches typically do not need to operate at full wire speed, because they are physically limited by their uplinks to the distribution layer. This means that less expensive, lower-performing switches can be used at the access layer, and more expensive, higher-performing switches can be used at the distribution and core layers, where the forwarding rate has a greater impact on network performance.

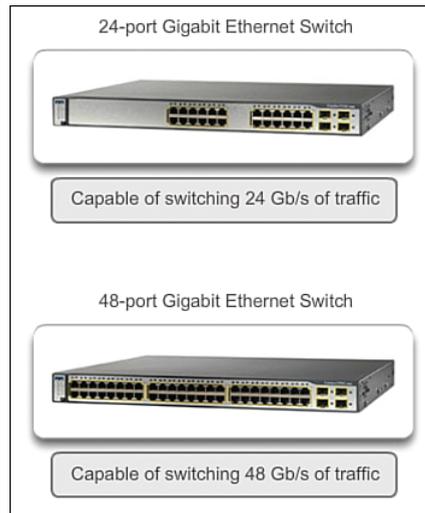


Figure 1-16 Forwarding Rate

Power over Ethernet (1.1.3.3)

Power over Ethernet (PoE) allows the switch to deliver power to a device over the existing Ethernet cabling. This feature can be used by IP phones and some wireless access points. The highlighted devices in Figure 1-17 have PoE ports.

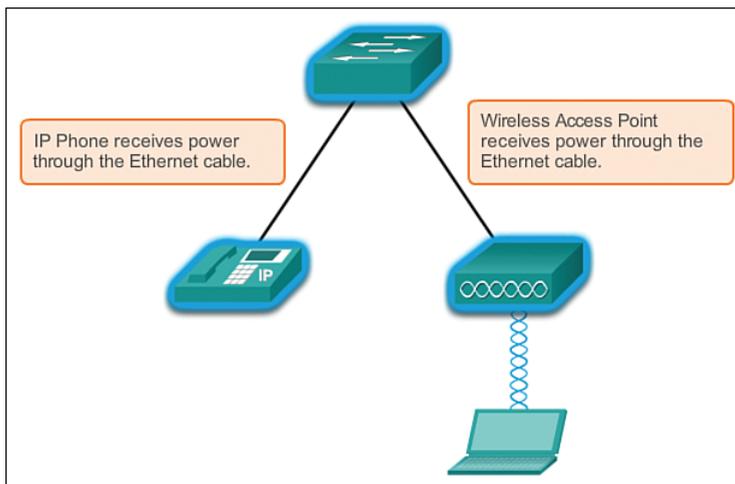


Figure 1-17 Power over Ethernet

PoE allows more flexibility when installing wireless access points and IP phones, allowing them to be installed anywhere that there is an Ethernet cable. A network administrator should ensure that the PoE features are required, because switches that support PoE are expensive.

The relatively new Cisco Catalyst 2960-C and 3560-C Series compact switches support PoE pass-through. PoE pass-through allows a network administrator to power PoE devices connected to the switch, as well as the switch itself, by drawing power from certain upstream switches. The highlighted switch in Figure 1-18 represents a Cisco Catalyst 2960-C.

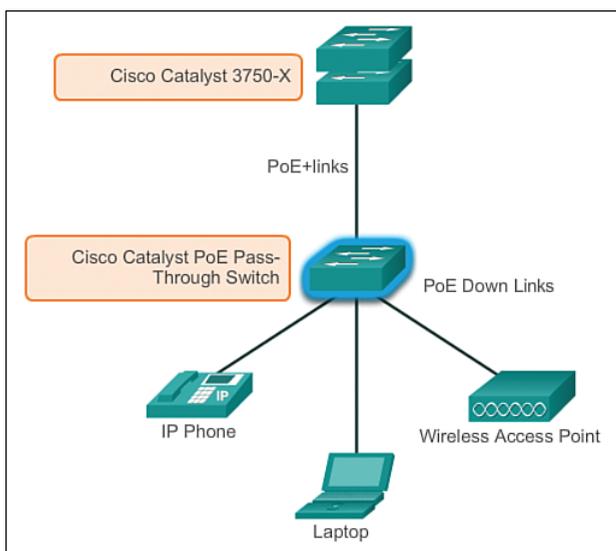


Figure 1-18 PoE Pass-Through

Cisco Catalyst Switch Breakdown (1.1.3.4)

While switches can be categorized in various ways, Cisco Catalyst switches are usually described in terms of the core-distribution-access hierarchy, as shown in Figure 1-19. The core and distribution layers often include the same types of switches, depending on the size of the network. Similarly, the distribution and access layers often include the same types of switches.



Figure 1-19 Switches in the Hierarchical Design Model

In general, the core and distribution layers incorporate four types of switches:

- **Cisco Catalyst 6500 Series Switches:** These switches scale to 4-terabit capacity with the Virtual Switching System, with up to 160 gigabits per slot; the switches are 100 Gigabit Ethernet ready, and support enhanced security, manageability, and wireless control.
- **Cisco Catalyst 4500E Series Switches:** These switches support modularity, offering 1.6-terabits-per-second capacity with the Virtual Switching System; these switches offer high availability bolstered by Control Plane Policing (CPP), and are ideal for collapsed distribution-access and small- to medium-distribution deployments.
- **Cisco Catalyst 4500-X Series Switches:** These switches are fixed aggregation switches for space-constrained environments, in a 1 RU form factor, and operate at 1.6-terabits-per-second capacity.

- **Cisco Catalyst 3750-X Series Switches:** These switches are stackable fixed-configuration switches for smaller, restrictive deployments, with advanced Layer 3 and Layer 2 switching and security services, and support for Gigabit and 10 Gigabit Ethernet aggregation, including comprehensive support for Borderless Networks services.

The distribution and access layers typically incorporate the following types of switches:

- **Cisco Catalyst 4500E Series Switches:** These switches come with high capacity (848 gigabits) and density (240 full Power Over Ethernet Plus ports), with 60 Watt Universal Power Over Ethernet to power a large range of devices, and high availability with Stateful Switchover (SSO).
- **Cisco Catalyst 3750-X Series Switches:** These switches are stackable fixed-configuration switches, with StackWise Plus and StackPower for high availability and operational efficiency, service and network modules for service upgrades, and full Power Over Ethernet Plus and comprehensive Borderless Networks services.
- **Cisco Catalyst 3560-X Series Switches:** These switches are fixed-configuration switches for campus and branch deployments, with high-availability and advanced security features, service and network modules for service upgrades, and full Power Over Ethernet Plus and comprehensive Borderless Networks services.
- **Cisco Catalyst 3560 and 3560-C Series Compact Switches:** These are sleek, quiet switches that deliver comprehensive access services outside the wiring closet and support for Power Over Ethernet Plus, Cisco EnergyWise, and advanced QoS, as well as providing a unique PoE pass-through capability that eliminates the need for power outlets.

The access layer normally incorporates the following types of switches:

- **Cisco Catalyst 2960 Series Switches:** These are stackable fixed-configuration Layer 2 switches that are a cost-effective solution for mid-sized organizations and branch offices, and provide full Power Over Ethernet Plus and baseline Borderless Networks services.
- **Cisco Catalyst 2960 and 2960-C Series Compact Switches:** These are sleek, quiet switches that deliver baseline access services outside the wiring closet, with support for Power Over Ethernet Plus, Cisco EnergyWise, and advanced QoS, and provide unique PoE pass-through capability that eliminates the need for power outlets.

With such a wide selection of switches to choose from in the Catalyst product line, an organization can carefully determine the ideal combination to meet the needs of the employees and the customers.

**Interactive
Graphic****Activity 1.1.3.5: Identify Switch Hardware**

Go to the online course to perform this practice activity.

**Lab 1.1.3.6: Selecting Switch Hardware**

In this lab, you will complete the following objectives:

- Part 1: Explore Cisco Switch Products
 - Part 2: Select an Access Layer Switch
 - Part 3: Select a Distribution/Core Layer Switch
-

The Switched Environment (1.2)

In this section you learn about *frame forwarding* of LAN switches and the role of broadcast domains and collision domains in a switched environment.

Frame Forwarding (1.2.1)

Computer networking is enabled by switching. Often people make the mistake of thinking that switching is specific to LANs. In reality, switching is a generic concept that applies to any networking device with interfaces on it. Switching in a generic sense refers only to the use of some sort of table to instruct a networking device what port to use to send out a packet based on the port in which the packet entered, coupled with specific information embedded in the packet. It really is up to your imagination what a generic switch might use to switch packets; it comes down to the set of rules used to build the table.

Switching as a General Concept in Networking and Telecommunications (1.2.1.1)

The concept of switching and forwarding frames is universal in networking and telecommunications. Various types of switches are used in LANs, WANs, and the public

switched telephone network (PSTN). The fundamental concept of switching refers to a device making a decision based on two criteria:

- *Ingress port*
- Some sort of address embedded in the frames or packets processed by the device

The decision on how a switch forwards traffic is made in relation to the flow of that traffic. The term *ingress* is used to describe where a frame enters the device on a port. The term *egress* is used to describe frames leaving the device from a particular port.

When a LAN switch makes a decision, it is based on the ingress port and the destination address of the message.

A LAN switch maintains a table that it uses to determine how to forward traffic through the switch. In Table 1-1, you see the information that a generic LAN switch might use to forward Ethernet frames.

Table 1-1 Generic LAN Switch

Port Table	
Destination Address	Port
EE	1
AA	2
BA	3
EA	4
AC	5
AB	6

With Table 1-1, the following conclusions can be made:

- If a message enters port 1 and has a destination address of EA, the switch forwards the traffic out port 4.
- If a message enters port 5 and has a destination address of EE, the switch forwards the traffic out port 1.
- If a message enters port 3 and has a destination address of AB, the switch forwards the traffic out port 6.

The only intelligence of the LAN switch is its ability to use its table to forward traffic based on the ingress port and the destination address of a message. With a LAN switch, there is only one master switching table that describes a strict association between addresses and ports; therefore, a message with a given destination address always exits the same *egress port*, regardless of the ingress port it enters.

Cisco LAN switches forward Ethernet frames based on the destination MAC address of the frames.

Dynamically Populating a Switch MAC Address Table (1.2.1.2)

Switches use MAC addresses to direct network communications through the switch to the appropriate port toward the destination. A switch is made up of integrated circuits and the accompanying software that controls the data paths through the switch. For a switch to know which port to use to transmit a frame, it must first learn which devices exist on each port. As the switch learns the relationship of ports to devices, it builds a table called a MAC address or content addressable memory (CAM) table. CAM is a special type of memory used in high-speed searching applications.

LAN switches determine how to handle incoming data frames by maintaining the *MAC address table*. A switch builds its MAC address table by recording the MAC address of each device connected to each of its ports. The switch uses the information in the MAC address table to send frames destined for a specific device out the port that has been assigned to that device.

A switch populates the MAC address table based on source MAC addresses. When a switch receives an incoming frame with a destination MAC address that is not found in the MAC address table, the switch forwards the frame out of all ports (flooding) except for the ingress port of the frame. When the destination device responds, the switch adds the source MAC address of the frame and the port where the frame was received to the MAC address table. In networks with multiple interconnected switches, the MAC address table contains multiple MAC addresses for a single port connected to the other switches.

The following steps describe the process of building the MAC address table:

1. The switch receives a frame from PC 1 on Port 1 in Figure 1-20.

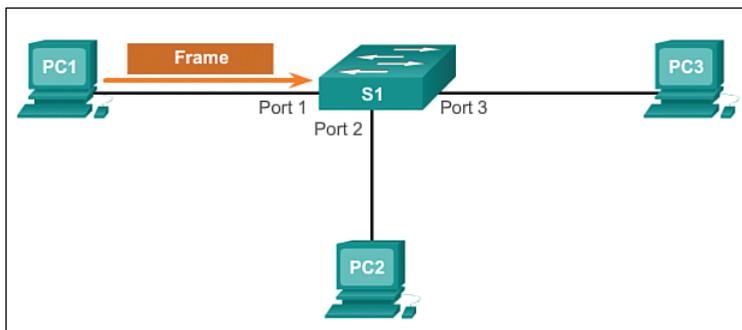


Figure 1-20 Receipt of a Frame

2. The switch examines the source MAC address and compares it to the MAC address table.
 - If the address is not in the MAC address table, it associates the source MAC address of PC 1 with the ingress port (Port 1) in the MAC address table, as shown in Figure 1-21.
 - If the MAC address table already has an entry for that source address, it resets the aging timer. An entry for a MAC address is typically kept for five minutes.

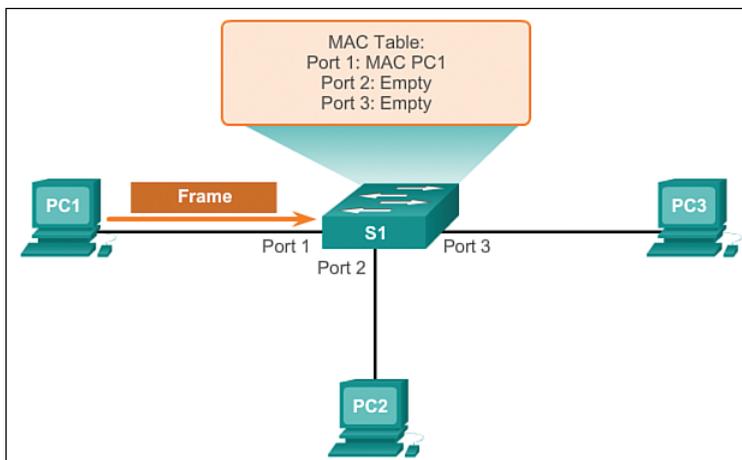


Figure 1-21 Parse Source MAC Address Against MAC Address Table

3. After the switch has recorded the source address information, the switch examines the destination MAC address.
 - If the destination address is not in the MAC table or if it's a broadcast MAC address, as indicated by all Fs, the switch floods the frame to all ports except the ingress port, as shown in Figure 1-22.

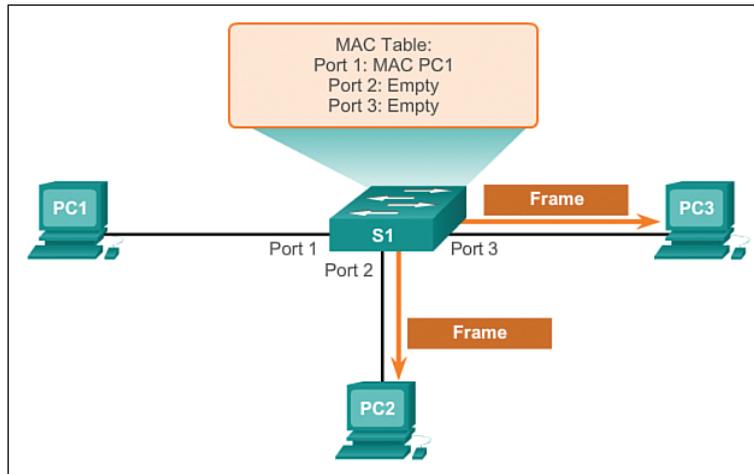


Figure 1-22 Switch Forwards Frame Out All Ports If Destination Is Not in MAC Address Table

4. The destination device (PC 3) replies to the frame with a unicast frame addressed to PC 1, as shown in Figure 1-23.

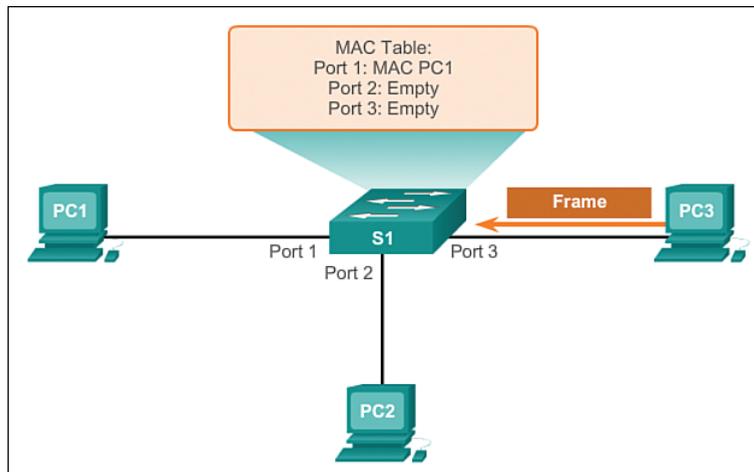


Figure 1-23 Frame Recipient Replies with Unicast Frame

5. The switch enters the source MAC address of PC 3 and the port number of the ingress port into the address table. The destination address of the frame and its associated egress port are found in the MAC address table, as shown in Figure 1-24.

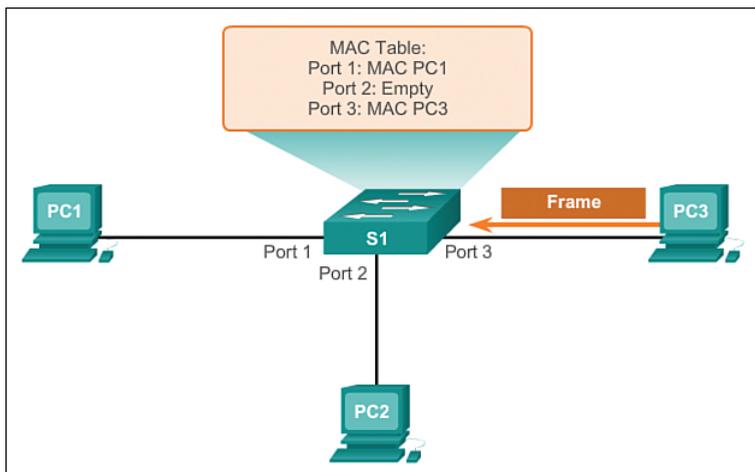


Figure 1-24 Switch Populates MAC Address Table with New Source MAC Address

6. The switch can now forward frames between these source and destination devices without flooding, because it has entries in the address table that identify the associated ports, as shown in Figure 1-25.

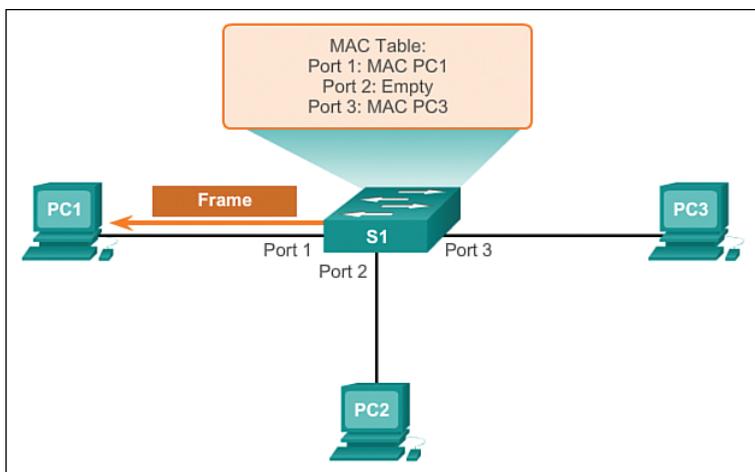


Figure 1-25 Switch Forwards Frame Out All Ports Associated with Original Sender

Switch Forwarding Methods (1.2.1.3)

As networks grew and enterprises began to experience slower network performance, Ethernet bridges (early versions of a switch) were added to networks to limit the size of the collision domains. In the 1990s, advancements in integrated circuit technologies allowed LAN switches to replace Ethernet bridges. These LAN switches were able to move the Layer 2 forwarding decisions from software to application-specific

integrated circuits (ASIC). ASICs reduce the packet-handling time within the device and allow the device to handle an increased number of ports without degrading performance. This method of forwarding data frames at Layer 2 was referred to as *store-and-forward switching*. This term distinguished it from *cut-through switching*. These switching methods are explored in this section.

The store-and-forward method makes a forwarding decision on a frame after it has received the entire frame and checked the frame for errors using a mathematical error-checking mechanism known as a cyclic redundancy check (CRC). The CRC was invented by Wesley Peterson at IBM in 1961.

By contrast, the cut-through frame forwarding method begins the forwarding process after the destination MAC address of an incoming frame and the egress port have been determined.

Store-and-Forward Switching (1.2.1.4)

Store-and-forward switching has two primary characteristics that distinguish it from cut-through: error checking and automatic buffering.

Error Checking

A switch using store-and-forward switching performs an error check on an incoming frame. After receiving the entire frame on the ingress port, as shown in Figure 1-26, the switch compares the frame check sequence (FCS) value in the last field of the datagram against its own FCS calculations. The FCS is an error-checking process that helps to ensure that the frame is free of physical and data-link errors. If the frame is error-free, the switch forwards the frame. Otherwise, the frame is dropped.

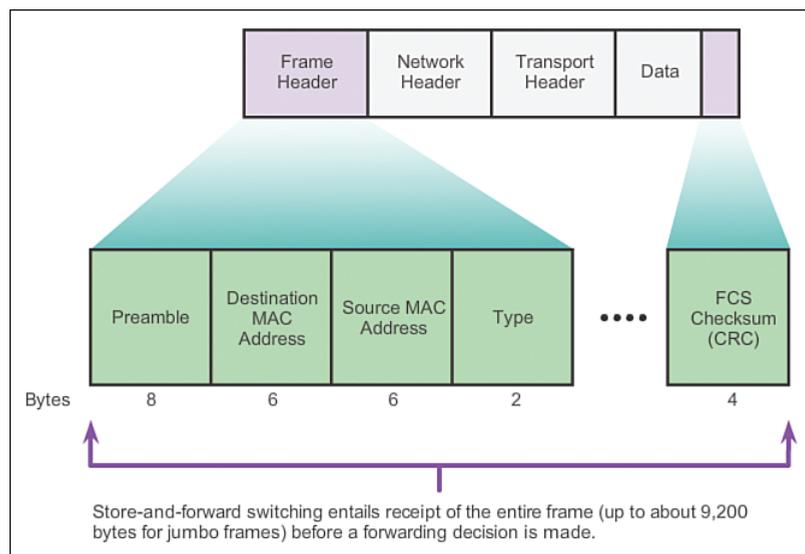


Figure 1-26 Store-and-Forward Switching

Automatic Buffering

The ingress port buffering process used by store-and-forward switches provides the flexibility to support any mix of Ethernet speeds. For example, handling an incoming frame traveling into a 100-Mb/s Ethernet port that must be sent out a 1-Gb/s interface would require using the store-and-forward method. With any mismatch in speeds between the ingress and egress ports, the switch stores the entire frame in a buffer, computes the FCS check, forwards it to the egress port buffer, and then sends it.

A store-and-forward switch drops frames that do not pass the FCS check and therefore does not forward invalid frames. By contrast, a cut-through switch can forward invalid frames because no FCS check is performed.

Cut-Through Switching (1.2.1.5)

An advantage to cut-through switching is the ability of the switch to start forwarding a frame earlier than store-and-forward switching. There are two primary characteristics of cut-through switching: rapid frame forwarding and fragment free.

Rapid Frame Forwarding

As indicated in Figure 1-27, a switch using the cut-through method can make a forwarding decision as soon as it has looked up the destination MAC address of the frame in its MAC address table. The switch does not have to wait for the rest of the frame to enter the ingress port before making its forwarding decision.

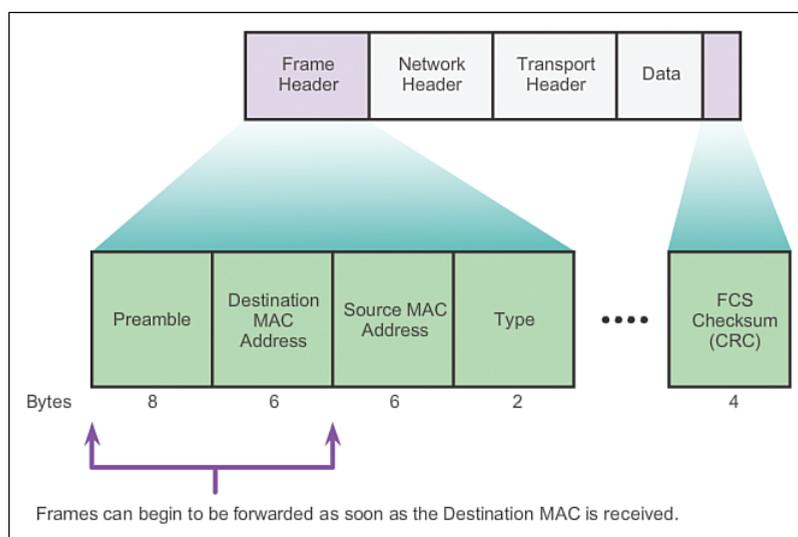


Figure 1-27 Cut-Through Switching

With today's MAC controllers and ASICs, a switch using the cut-through method can quickly decide whether it needs to examine a larger portion of a frame's headers for additional filtering purposes. For example, the switch can analyze past the first 14 bytes (the source MAC address, the destination MAC address, and the EtherType fields) and examine an additional 40 bytes to perform more sophisticated functions relative to IPv4 Layers 3 and 4.

The cut-through switching method does not drop most invalid frames. Frames with errors are forwarded to other segments of the network. If there is a high error rate (invalid frames) in the network, cut-through switching can have a negative impact on bandwidth, thus clogging bandwidth with damaged and invalid frames.

Fragment Free

Fragment free switching is a modified form of cut-through switching in which the switch waits for the collision window (64 bytes) to pass before forwarding the frame. This means that each frame will be checked into the data field to make sure that no fragmentation has occurred. Fragment free mode provides better error checking than cut-through, with practically no increase in latency.

The lower latency speed of cut-through switching makes it more appropriate for extremely demanding, high-performance computing (HPC) applications that require process-to-process latencies of 10 microseconds or less.

Interactive Graphic

Activity 1.2.1.6: Frame Forwarding Methods

Go to the online course to perform this practice activity.

Interactive Graphic

Activity 1.2.1.7: Switch It!

Go to the online course to perform this practice activity.

Switching Domains (1.2.2)

Access switches determine collision domains. Routers and multilayer switches determine broadcast domains. However, VLANs coincide with broadcast domains in a switched environment, so access switches also contribute to the determination of broadcast domains. In this topic, you will explore the relationship between collision domains and broadcast domains.

Collision Domains (1.2.2.1)

In hub-based Ethernet segments, network devices compete for the medium, because devices must take turns when transmitting. The network segments that share the same bandwidth between devices are known as *collision domains*, because when two or more devices within that segment try to communicate at the same time, collisions can occur.

It is possible, however, to use a switch device, operating at the OSI data link layer, to divide a network into segments and reduce the number of devices that compete for bandwidth. When a switch is used, each port represents a new segment. Each new segment is a new collision domain. More bandwidth is available to the devices on the segment, and collisions in one collision domain do not interfere with the other segments. This is also known as microsegmentation.

As shown in Figure 1-28, each switch port connects to a single PC or server, and each switch port represents a separate collision domain.

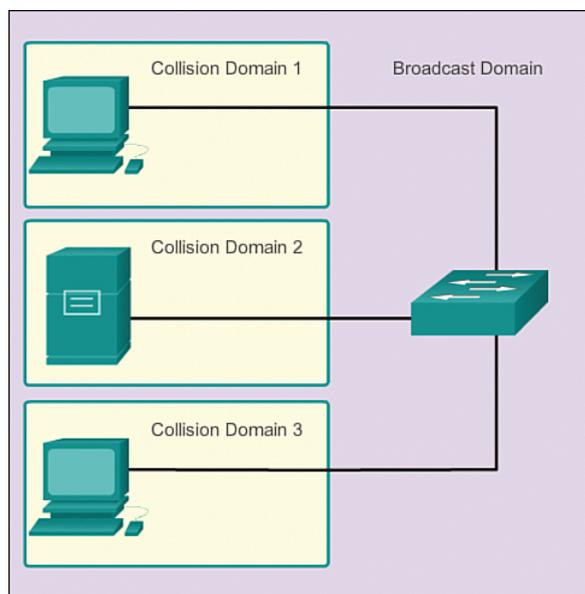


Figure 1-28 Collision Domains and Broadcast Domains

Broadcast Domains (1.2.2.2)

Although switches filter most frames based on MAC addresses, they do not filter broadcast frames. For other devices on the LAN to receive broadcast frames, switches must flood these frames out all ports except the one on which the broadcast was received. A collection of interconnected switches forms a single *broadcast domain*.

Only a network layer device, such as a router, can divide a Layer 2 broadcast domain. Routers are used to segment both collision and broadcast domains.

When a device sends a Layer 2 broadcast, the destination MAC address in the frame is set to all binary 1s. A frame with a destination MAC address of all binary 1s, or all Fs in hexadecimal, is received by all devices in the broadcast domain.

The Layer 2 broadcast domain is referred to as the MAC broadcast domain. The MAC broadcast domain consists of all devices on the LAN that receive broadcast frames from a host.

Video**Video 1.2.2.2: Broadcast Domains I**

Go to the online course and view the first half of the animation.

When a switch receives a broadcast frame, it forwards the frame out each of its ports, except the ingress port where the broadcast frame was received. Each device connected to the switch receives a copy of the broadcast frame and processes it. Broadcasts are sometimes necessary for initially locating other devices and network services, but they also reduce network efficiency. Network bandwidth is used to propagate the broadcast traffic. Too many broadcasts and a heavy traffic load on a network can result in congestion: a slowdown in the network performance.

When two switches are connected together, the broadcast domain is increased.

Video**Video 1.2.2.2: Broadcast Domains II**

Go to the online course and view the second half of the animation.

In this case, a broadcast frame is forwarded to all connected ports on switch S1. Switch S1 is connected to switch S2. The frame is then also propagated to all devices connected to switch S2.

Alleviating Network Congestion (1.2.2.3)

LAN switches have special characteristics that make them effective at alleviating network congestion. First, they allow the segmentation of a LAN into separate collision domains. Each port of the switch represents a separate collision domain and provides the full bandwidth to the device or devices that are connected to that port. Second, they provide full-duplex communication between devices. A full-duplex connection can carry transmitted and received signals at the same time. Full-duplex connections have dramatically increased LAN network performance and are required for 1-Gb/s Ethernet speeds and higher.

Switches interconnect LAN segments (collision domains), use a table of MAC addresses to determine the segment to which the frame is to be sent, and can lessen or eliminate collisions entirely. Following are some important characteristics of switches that contribute to alleviating network congestion:

- **High port density:** Switches have high port densities: 24- and 48-port switches are often just 1 rack unit (1.75 inches) in height and operate at speeds of 100 Mb/s, 1 Gb/s, and 10 Gb/s. Large enterprise switches can support many hundreds of ports.
- **Large frame buffers:** The ability to store more received frames before having to start dropping them is useful, particularly when there might be congested ports to servers or other parts of the network.
- **Port speed:** Depending on the cost of a switch, it might be possible to support a mixture of speeds. Ports of 100 Mb/s, and 1 or 10 Gb/s, are common (100 Gb/s is also possible).
- **Fast internal switching:** Having fast internal forwarding capabilities allows high performance. The method that is used can be a fast internal bus or shared memory, which affects the overall performance of the switch.
- **Low per-port cost:** Switches provide high port density at a lower cost. For this reason, LAN switches can accommodate network designs featuring fewer users per segment, therefore increasing the average available bandwidth per user.

**Interactive
Graphic****Activity 1.2.2.4: Circle the Domain**

Go to the online course to perform this practice activity.

Summary (1.3)



Class Activity 1.3.1.1: It's Network Access Time

Use Packet Tracer for this activity. Internet connectivity is not required in this design. Work with a classmate to create two network designs to accommodate the following scenarios:

Scenario 1: Classroom Design (LAN)

- 15 student end devices represented by one or two PCs
- One instructor end device, preferably represented by a server
- Stream video presentations over a LAN connection

Scenario 2: Administrative Design (WAN)

- All requirements as listed in Scenario 1
- Access to and from a remote administrative server for video presentations and pushed updates for network application software

Both the LAN and WAN designs should fit on one Packet Tracer file screen. All intermediary devices should be labeled with the switch model (or name) and the router model (or name).

Save your work and be ready to justify your device decisions and layout to your instructor and the class.

Interactive Graphic

Activity 1.3.1.2: Basic Switch Configurations

Go to the online course to use the Syntax Checker to perform basic switch configurations.

Packet Tracer Activity

Packet Tracer Activity 1.3.1.3: Skills Integration Challenge

As a recently hired LAN technician, your network manager has asked you to demonstrate your ability to configure a small LAN. Your tasks include configuring initial settings on two switches using the Cisco IOS and configuring IP address parameters on host devices to provide end-to-end connectivity. You are to use two switches and two hosts/PCs on a cabled and powered network.

We have seen that the trend in networks is toward convergence using a single set of wires and devices to handle voice, video, and data transmission. In addition, there has been a dramatic shift in the way businesses operate. No longer are employees constrained to physical offices or by geographic boundaries. Resources must now be seamlessly available anytime and anywhere. The Cisco Borderless Network architecture enables different elements, from access switches to wireless access points, to work together and allow users to access resources from any place at any time.

The traditional three-layer hierarchical design model divides the network into core, distribution, and access layers, and allows each portion of the network to be optimized for specific functionality. It provides modularity, resiliency, and flexibility, which provide a foundation that allows network designers to overlay security, mobility, and unified communication features. In some networks, having a separate core and distribution layer is not required. In these networks, the functionality of the core layer and the distribution layer is often collapsed together.

There are various types of switches used in business networks. It is important to deploy the appropriate types of switches based on network requirements. When selecting the type of switch, the network designer must choose between a fixed or modular configuration, and stackable or nonstackable. Another consideration is the thickness of the switch, which is expressed in number of rack units. A network administrator might choose to implement a multilayer switch. Multilayer switches are characterized by their ability to build a routing table, support a few routing protocols, and forward IP packets at a rate close to that of Layer 2 forwarding. Other switch features that should be considered include port density, forwarding rates, power capabilities (such as PoE), and scalability features.

Cisco LAN switches use ASICs to forward frames based on the destination MAC address. Before this can be accomplished, the switch must first use the source MAC address of incoming frames to build a MAC address table in content-addressable memory (CAM). If the destination MAC address is contained in this table, the frame is forwarded only to the specific destination port. In cases where the destination MAC address is not found in the MAC address table, the frames are flooded out all ports, except the one on which the frame was received.

Switches use either store-and-forward or cut-through switching. Store-and-forward reads the entire frame into a buffer and checks the CRC before forwarding the frame. Cut-through switching only reads the first portion of the frame and starts forwarding it as soon as the destination address is read. Although this is extremely fast, no error checking is done on the frame before forwarding.

Every port on a switch forms a separate collision domain, allowing extremely high-speed, full-duplex communication. Switch ports do not block broadcasts, and connecting switches together can extend the size of the broadcast domain, often resulting in degraded network performance.

Practice

The following activities provide practice with the topics introduced in this chapter. The Labs and Class Activities are available in the companion *Switched Networks Lab Manual* (ISBN 978-1-58713-372-5). The Packet Tracer Activities PKA files are found in the online course.



Class Activities

- Class Activity 1.0.1.2: Sent or Received Instructions
- Class Activity 1.3.1.1: It's Network Access Time



Labs

- Lab 1.1.3.6: Selecting Switch Hardware



Packet Tracer Activities

- Packet Tracer Activity 1.1.2.5: Comparing 2960 and 3560 Switches

Check Your Understanding Questions

Complete all the review questions listed here to test your understanding of the topics and concepts in this chapter. The appendix “Answers to ‘Check Your Understanding’ Questions” lists the answers.

1. What are the layers of the switch hierarchical design model? (Choose three.)
 - A. Access
 - B. Data link
 - C. Core
 - D. Network access
 - E. Enterprise
 - F. Distribution

2. Which of the following characteristics describe a converged network? (Choose two.)
 - A. Support of voice and video, both using the same switch
 - B. Separate wiring infrastructure for voice and video traffic
 - C. Affordability for small and medium businesses
 - D. Cheaper equipment cost

3. When an appropriate switch form factor for a network is being determined, what should be selected when fault tolerance and bandwidth availability are desired but the budget is limited?
 - A. Stackable switch
 - B. Nonstackable switch
 - C. Fixed configuration switch
 - D. Modular switch

4. Which cost-effective physical network topology design is recommended when building a three-tier campus network that connects three buildings?
 - A. Bus
 - B. Mesh
 - C. Extended star
 - D. Dual-ring

5. When the appropriate switch form factor for a network is being determined, what type of switch should be selected when future expansion is important and cost is not a limiting factor?
 - A. Stackable switch
 - B. 1-rack-unit switch
 - C. Fixed configuration switch
 - D. Modular switch

6. Fill in the blank. The technology that allows a switch to deliver power to a device like an IP phone or an access point through the data cable is known as _____.

-
7. Which of the following statements about Layer 2 Ethernet switches are true? (Choose two.)
- A. Layer 2 switches prevent broadcasts.
 - B. Layer 2 switches have multiple collision domains.
 - C. Layer 2 switches route traffic between different networks.
 - D. Layer 2 switches decrease the number of broadcast domains.
 - E. Layer 2 switches can send traffic based on the destination address.
8. A network administrator is researching enterprise-level switches to upgrade the network infrastructure. Which switching feature defines the overall amount of data that the switch can process each second?
- A. Forwarding rate
 - B. Wire speed
 - C. PoE
 - D. Port density
9. Which option best describes a switching method?
- A. Cut-through: makes a forwarding decision after receiving the entire frame
 - B. Store-and-forward: forwards the frame immediately after examining its destination MAC address
 - C. Cut-through: provides the flexibility to support any mix of Ethernet speeds
 - D. Store-and-forward: ensures that the frame is free of physical and data-link errors
10. Which service is provided by an automated attendant feature on a converged network?
- A. Point-to-point video
 - B. Call routing
 - C. IT management interface
 - D. Videoconferencing
11. A medium-sized company wants to add IP phones to its network. Should it consider buying a switch that supports PoE?
- A. Yes, because PoE increases port density.
 - B. Yes, because PoE provides more flexibility in placing IP phones.
 - C. No, because PoE has no effect on the use of VoIP devices on a network.
 - D. Yes, because PoE adds Layer 3 functionality to a switch.

12. Which switching mode describes a switch that transfers a frame as soon as the destination MAC address is read?
- A. Fragment free
 - B. Cut-through
 - C. Store-and-forward
 - D. Latency forwarding