Networking Computers with Switches

Switches are an important component in any wired network. At one point, switches were optional in most, but they have steadily replaced Ethernet hubs because of their extra capabilities. Additionally, advanced switches can perform functions of a router.

In this chapter, you'll learn the details of how a switch works and its benefits over a hub. You'll also learn about the differences between managed and unmanaged switches, between layer 2 and layer 3 switches, and how a switch can be used to create VLANs. Last, you'll learn some basics about switch speeds and switch security.

- **Connecting multiple computers**
- **Understanding physical ports**
- **Comparing hubs and switches**
- **Comparing managed and unmanaged switches**
- **Exploring switch speeds**
- **Understanding security options**

Connecting Multiple Computers

Chapter 2 introduced basic connectivity with hubs, switches, and routers. As a reminder, *switches* (or *hubs*) connect computers in a network. In larger organizations, routers connect multiple networks into a local area network (LAN).

Consider Figure 8.1. It shows three separate subnetworks of 192.168.1.0/24, 192.168.5.0/24, and 192.168.7.0/24. Each of the subnetworks has a central switch connecting the devices. A router connects the three.

The /24 represents CIDR notation. It indicates the subnet mask has the first 24 bits set to a 1. In other words, the subnet mask is 255.255.255.0.

Packets sent by computers on this network go through a switch first. The switch learns which computers are connected to which port. It uses this knowledge to determine the path for every packet it receives. In contrast, routers move packets between the subnetworks.

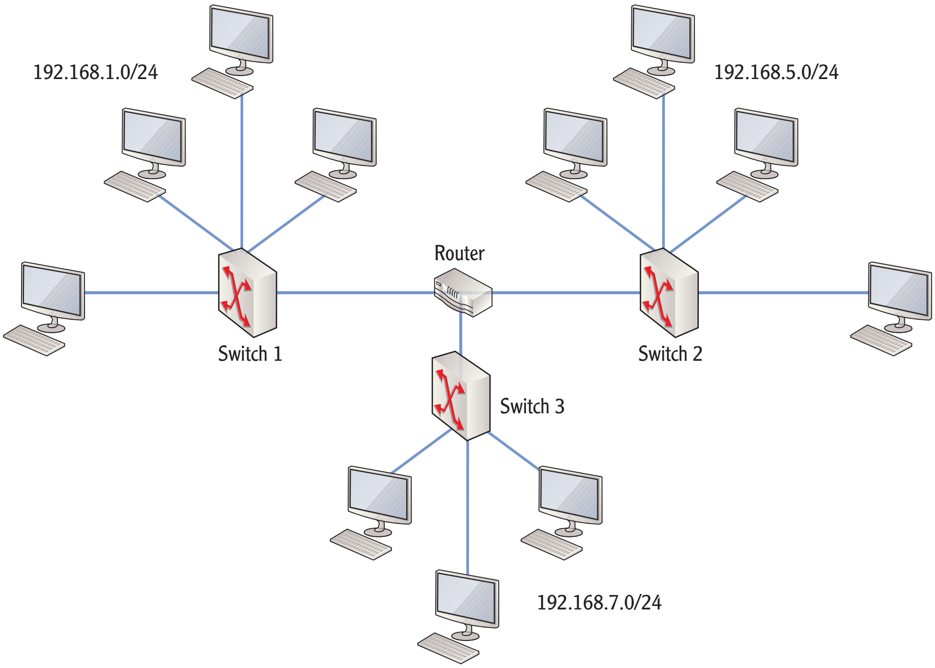

FIGURE 8.1 Using switches to connect computers

Networks, Subnets, and Subnetworks

The terms *networks*, *subnetworks*, and *subnets* can easily be confusing. The terms are sometimes mixed together, and it's worth identifying the differences. In general, a network is two or more computers or other network devices connected together. When they are connected, they can share data and resources with each other.

Both a subnet and a subnetwork are a group of computers with the same network ID. Additionally, routers separate subnetworks and subnets from each other. However, there is a subtle difference between these two.

A subnet is a network that started as a classful network and was divided into multiple subnets. For example, you can divide a single Class C network of 192.168.1.0/24 into four subnetworks of 192.168.1.0/26, 192.168.1.64/26, 192.168.1.128/26, and 192.168.1.192/26. (The "Subnetting IPv4 Addresses" section of Chapter 5 explained subnetting.)

(Continues)

Networks, Subnets, and Subnetworks (Continued)

Subnetworks use different classful IP ranges without subnetting them. For example, you can create one subnetwork with an address in 192.168.1.0/24, and a different subnetwork could have an address range in 192.168.5.0/24. This creates two subnetworks without subnetting a classful IP address. This is actually very common in private networks. There are more than enough classful IP address ranges for even very large organizations to use without subnettting.

Unfortunately, not everyone uses these terms in the same way. You'll often hear technicians call a subnetwork a *subnet* or simply a *network*. Some insist that they are all called *networks*. Some insist that a subnet is created only when a classful IP address range has been subnetted, and the rest are networks.

The debate will continue. However, the convention I'm using in this chapter is to separate subnetworks, subnets, and networks. If I called them all *networks*, then it would be easy for you to become confused between a local area network (composed of multiple subnetworks) and a network on one side of the router (which is only one subnetwork).

Although the terminology can be tricky, the primary message should still be clear. Switches connect and track computers within a subnetwork. Routers connect and track subnetworks.

Switches track the location of the computers on networks. Routers track networks or subnetworks, not computers.

If the destination computer is on the same subnetwork, the switch forwards the packet to the destination computer. If the destination computer is on a different subnetwork, the switch forwards the packet to the router for routing to the correct subnetwork.

Notice that the switch is the central device for each subnetwork. It could be a hub, but for several reasons, switches have replaced hubs in many networks. Although most networks have switches connected as shown in Figure 8.1, many network line drawings omit the icon of the switch. For example, Figure 8.2 shows a network line drawing with the switches omitted.

Even though Figure 8.2 doesn't show the switches for each subnetwork, the computers have to be connected to a central device. They wouldn't all be connected directly to the router.

Understanding Physical Ports

Twisted-pair cable uses RJ-45 connectors. Each end of the twisted-pair cable has an RJ-45 connector, and the connector plugs into the port.

Switches have physical ports where physical cables plug in. For example, if the network is using twisted-pair cable, the switch will have RJ-45 ports that accept RJ-45 connectors. If the cable is fiber optic, the switch has physical ports that accept the fiber-optic connectors.

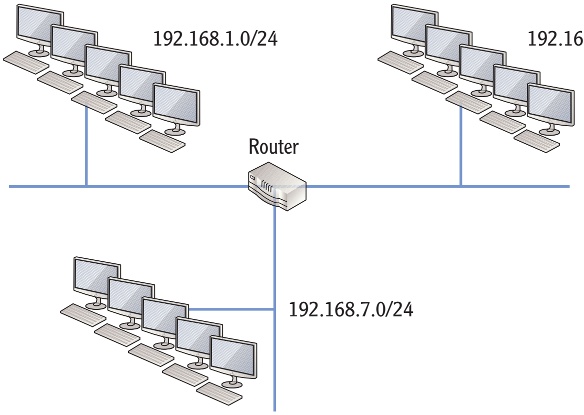

FIGURE 8.2 Using switches to connect computers

Uplinks connect two switches together or connect the switch to a router.

Most switches that connect end user computers have RJ-45 ports since twisted pair is the most commonly used media. Some switches include fiber-optic ports for uplinks.

Physical ports and logical ports are not the same. A physical port is something you can touch and accepts a cable. A logical port is simply a number that is embedded in a packet. For example, HTTP uses a logical default port of 80. When the packet reaches the destination computer, the logical port identifies the service or application that will process the data.

Chapters 3 and 4 covered logical ports. These are numbers embedded in data packets.

Identifying the Number and Type of Ports

The number of ports on a switch or hub varies according to the physical size of the device. Hubs are less expensive than switches and commonly have between 4 and 24 ports. This is usually enough for a small office/home office (SOHO) network. If selecting a device for a small business with 8 or 10 users, a 24-port device may be a reasonable investment, allowing for future growth.

Switches typically have between 8 to 64 ports. You can purchase switches in two separate designs:

Form Factor Switch This has a set number of ports built into the switch, and the number of ports can't be changed. Form factor switches can have any number of ports, but 48 is the maximum for most form factor switches. These switches are great where simplicity is required.

Modular Switch A modular switch starts with few to zero ports and can expand to hundreds of ports. You can then add plug-in modules to add ports. This is similar to a computer that can accept additional memory modules. For example, the computer may start with 1 GB RAM, but you can add RAM when your needs change. Similarly, you can buy a modular switch with a module that includes eight ports but then add modules to increase the number of available ports.

You can add different types of modules to a modular switch. This includes modules for typical RJ-45 ports, wireless services, video services, and more.

Selecting a switch design is based on several factors. For example, you will want to ensure you have enough ports for your immediate needs while also considering future growth requirements. Most modular switches require programming, which adds administrative overhead, while many form factor switches work right out of the box.

Identifying Ports in Drawings

When switches are included in network drawings or connection maps, the ports are usually labeled. This allows technicians to identify what port goes to what system. Switch ports are commonly labeled with E, F, or Gi followed by a number. For example, the following conventions are common:

E The first 10 Mbps port is labeled as E0. This indicates Ethernet port 0. Some manufacturers represent the first Ethernet port on a modular switch as E0/0, which represents the first port on the first module.

F The first 100 Mbps port is typically labeled as F0 or F0/0, a Fast Ethernet port. Compare this to the second port on the first module, which is labeled as F0/1. Fast Ethernet ports can also be labeled as Fa instead of just F.

Gi The first 1000 Mbps port is labeled as Gi0/0, a gigabit port. Similarly, the first port on the second module is Gi1/0.

Many computer technologies use zero-based numbering where the first item is a 0 instead of a 1. In a switch, the first port is 0 instead of 1.

As an example, consider Figure 8.3. Based on how the ports are labeled, you can gain additional information about the switches. The ports on switch 1 are labeled as E0 through E4, so it is a 10 Mbps switch. The ports on switch 2 are labeled as F0 through F4, so it is a 100 Mbps switch. The ports on switch 1 are labeled as G0 through G3, so it is a 1000 Mbps switch.

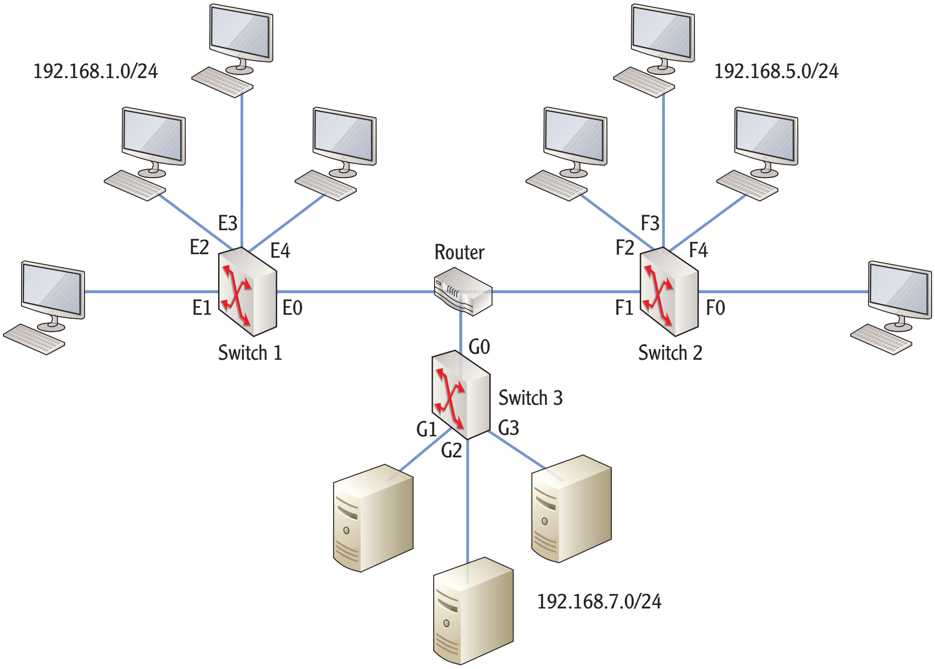

FIGURE 8.3 Identifying switch ports

Comparing Hubs and Switches

Chapter 2 introduced hubs, switches, and the concepts of broadcast and collision domains. Before going too far, it's worth repeating and expanding some of the key material.

A hub is a layer 1 device that connects multiple network devices. When using a hub, bandwidth decreases as you add more devices to the network since all devices share bandwidth equally. Any data sent into one port of a hub goes out all other ports. With a switch, each port is separated from each other. The ports do not share bandwidth with one another.

Hubs create a single collision domain and a single broadcast domain. Switches create multiple collision domains but share a common broadcast domain. Routers create separate collision and broadcast domains.

Understanding Collision Domains

A collision domain is a group of devices on the same segment that are subject to collisions. A hub creates a single collision domain. A switch creates multiple collision domains.

Ethernet Carrier Sense Multiple Access/Collision Detection (CSMA/CD) helps each device determine when they can send data across the network. The devices listen for traffic, and if they don't hear any traffic, they are free to send data. However, just as two people can start talking at the same time, two computers can start sending data at the same time.

Consider Figure 8.4. This shows several computers connected on the same segment. PC-1 and PC-4 are both sending traffic at the same time, and a collision occurs.

Collisions on the network degrade the network performance. Both computers must resend data each time a collision occurs.

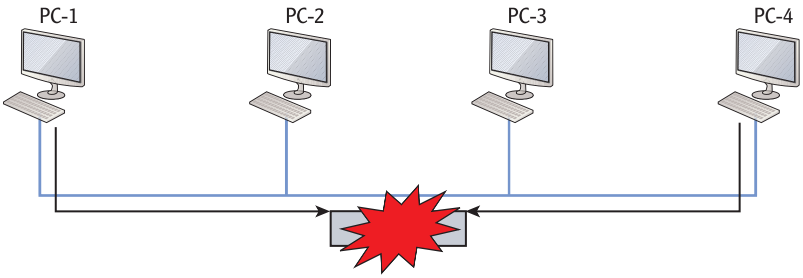

FIGURE 8.4 Collisions on a collision domain

CSMA/CD has a recovery mechanism for retransmission of lost data after a collision. First, a jamming signal is transmitted on the segment letting all devices know there has been a collision. Other devices postpone data transmission until the devices that had the collision resend their data. Once this process is complete, the network is reopened for business, and other devices may transmit their data.

If you think this creates a lot of network traffic, you are correct. Add collisions, and network performance will rapidly decrease.

You can reduce the number of collisions by increasing the number of collision domains. Since a hub has a single collision domain and a switch creates a separate collision domain for each port, you reduce collisions by replacing hubs with switches.

A hub has no intelligence. All traffic received on one port is flooded to all other ports.

Identifying a Collision Domain with a Hub

A hub connects multiple computers into a single collision domain. In other words, all devices connected with a hub contend for equal access to the same segment.

Figure 8.5 shows four computers connected with a hub. This is logically the same as Figure 8.4 shown earlier. Each computer uses CSMA/CD to listen before transmitting and can cause a collision if it sends data at the same time as another computer.

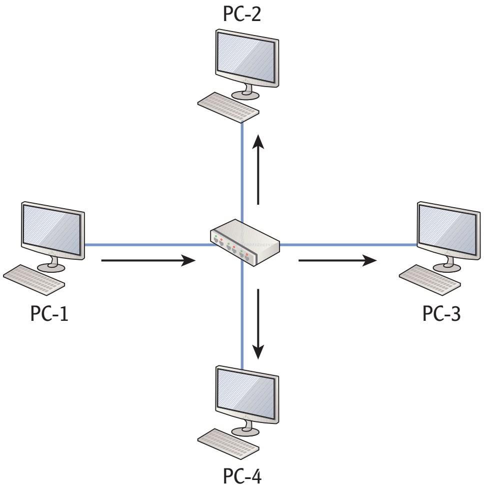

FIGURE 8.5 A single collision domain created by a hub

The hub acts as the central point in the network, and any traffic destined for another device will go to the hub first. In Figure 8.5, PC-1 sends data to PC-3. However, the hub forwards the packet to all computers connected to it. PC-3 will process the packet while PC-2 and PC-4 will discard it. However, this packet can cause a collision if any other computer sends data at the same time.

Identifying Collision Domains with a Switch

In contrast, Figure 8.6 shows how a switch creates multiple collision domains. Unicast traffic sent from the source computer is passed only to the destination computer. The switch creates an internal connection between PC-1 and PC-3. PC-2 and PC-4 don't receive the data.

If PC-2 sends data to PC-4 at the same time PC-1 is sending data to PC-3, it doesn't cause a collision. Instead, the switch makes an internal connection between PC-2 and PC-4.

A logical question is, "How does the switch track the computers connected to the ports?" That's a great question. The short answer is that a switch maps physical ports to computers' media access control (MAC) addresses.

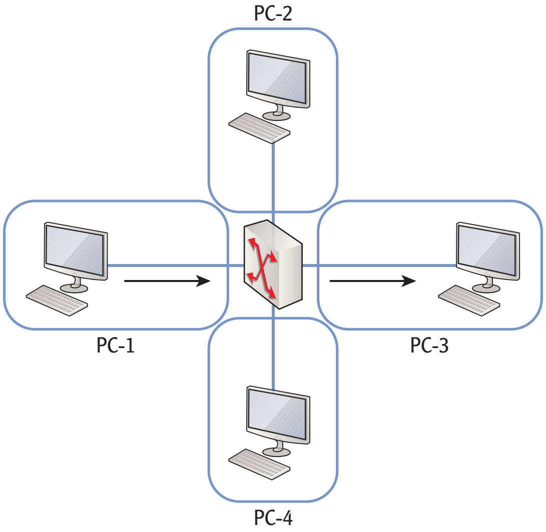

FIGURE 8.6 Multiple collision domains created by a switch

Mapping Ports to MAC Addresses

Chapter 3 introduced *MAC addresses*. As a reminder, a MAC address is a 48-bit address expressed in a hexadecimal format. For example, a MAC address looks like 00-23-5A-33-C4-CA.

Every network interface card (NIC) has a MAC address assigned to it. MAC addresses are typically burned into the card and unchangeable, though some NICs allow you to modify the MAC. Additionally, when a computer sends data to another computer, it always includes both its own IP address and its own MAC address as part of the source information.

A simple switch starts with very little knowledge when it's turned on. It knows what ports it has, but it does not know which computers are connected to which ports. However, as traffic is sent through the switch, it learns. It populates an internal *MAC address table* with the MAC addresses of each computer and maps them to the port to which they're connected.

Consider Figure 8.7. It shows a four-port switch with a computer connected to each port, and it shows their MAC addresses. Imagine this switch is just turned on. When PC-1 sends data to PC-3, the switch doesn't know what port PC-3 is on, so it sends the data to all ports. However, the packet from PC-1 includes the MAC address of PC-1. The switch silently says "gotcha" and starts populating the MAC table by logging port number F0 with the MAC address of PC-1.

Four bits are used to represent each hexadecimal character. Chapter 6 covered hexadecimal numbers in more depth.

Advanced switches (managed switches) are configurable. You can configure them with the MAC addresses of connected computers.

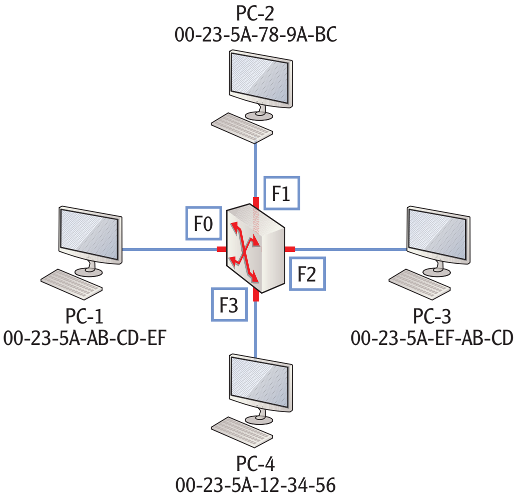

MAC Table

Port	MAC
F0	00-23-5A-AB-CD-EF
F1	00-23-5A-78-9A-BC
F2	00-23-5A-EF-AB-CD
F3	00-23-5A-12-34-56

FIGURE 8.7 **Mapping ports to MAC addresses**

When PC-3 answers, it includes the destination MAC address of PC-1 and the source address of PC-3. Again, the switch silently says "gotcha" and logs the MAC address of PC-3 with port F2 in the MAC address table. Since the switch knows that PC-1 is connected to port F0 (based on the MAC address), it internally switches the data from PC-3 to PC-1 on port F0. In a very short period, the switch will learn the MAC addresses of each computer along with their associated ports.

When the switch maps the MAC addresses to ports as traffic passes, the table is updated dynamically. However, an administrator can configure a managed switch with specific MAC addresses as static entries. Dynamic entries can be overwritten as time passes, but static entries remain.

Static entries are used for port security. Port security is explained later in this chapter.

Managed switches have configurable ports. You cannot configure ports on an unmanaged switch.

Comparing Managed and Unmanaged Switches

Switches can be configurable or nonconfigurable. If the switch must be customized in some way, then the administrator must be able to configure it. This requires additional knowledge on the part of the administrator and extra time. In other words, a configurable switch has administrative overhead.

A configurable switch is a *managed switch*. A nonconfigurable switch is an *unmanaged switch*. Determining what switch you need requires some basic knowledge of these two types of switches.

Understanding Unmanaged Switches

An unmanaged switch is just like a hub with respect to administrative overhead. There isn't any. You take the switch out of the box and plug it in, and it works.

The switch will monitor the traffic from each of the ports and build the MAC address table. As mentioned, the MAC address table maps the MAC addresses of the connected computers to their respective ports.

Even though an unmanaged switch doesn't require any administration, it does provide performance benefits over the simple hub. It still creates separate collision domains and increases performance on the network.

Unmanaged switches operate at layer 2 of the OSI Model.

Chapter 3 presented the OSI Model. Layer 2 is the Data Link layer.

Understanding Managed Switches

In contrast to an unmanaged switch, a managed switch can be configured. Managed switches are commonly managed using protocols such as Telnet or Secure Shell (SSH), and administrators can monitor and configure the switch remotely. As a reminder, SSH encrypts the traffic so that it can't be read if intercepted by a protocol analyzer or sniffer, while Telnet transmits in clear text.

Chapter 4 discussed PuTTY, a common application that uses SSH to administer managed switches from a remote location.

Some of the management tasks that an administrator can perform are as follows:

- Configure static entries in the MAC table
- Configure duplex settings (half-duplex or full-duplex) on ports
- Monitor performance of the switch using the Simple Network Management Protocol (SNMP)
- Configure the switch to send alerts called traps with SNMP when certain events occur
- Create a virtual LAN (VLAN)
- Configure port mirroring

Port mirroring sends a copy of all traffic on the switch to a single port. Administrators can capture traffic on this port for monitoring with a packet sniffer.

Although managed switches provide many more capabilities, they are also more expensive. Before spending the extra money for a managed switch, you should first ensure that your administrative staff can support them. If not, you may have a shiny new electronic toy with expensive extra features that no one knows how to use.

Managed switches can operate at layer 2 or layer 3 of the OSI Model.

Layer 2 is the Data Link layer. Layer 3 is the Network layer.

Comparing Layer 2 and Layer 3 Switches

A layer 2 switch has the primary purpose of segmenting collision domains at layer 2 of the OSI Model. Each port is segmented from the others. Layer 2 switches are hardware based, which makes them extremely fast. They use the integrated circuitry on the main board (the hardware) to move data between ports at lightning speed.

As you'd expect, a *layer 3* switch operates at layer 3 of the OSI Model. It includes standard switching functionality, but also contains routing capability to route layer 3 traffic just as if it were a router. Although the router is a great layer 3 device, it can be slow, because additional procesing of the packets must take place by the integrated software. A hardware-based switch is quicker.

Routers route traffic on layer 3. A layer 3 switch can route traffic on layer 3 like a router.

Only managed switches can be configured to route traffic on layer 3 like a router. Additionally, managed switches can be configured to create virtual local area networks. Unmanaged switches will work on layer 2 only.

As mentioned previously, a regular switch (a layer 2 switch) creates separate collision domains. It does pass broadcasts, so broadcast traffic goes to all ports on the layer 2 switch. In other words, a layer 2 switch does not create separate broadcast domains.

In contrast, a router does not pass broadcasts. The router creates separate broadcast domains. However, a layer 3 switch acts like a router and creates separate broadcast domains. You can use this to ensure traffic is routed to only certain ports on a switch without replacing the switch with a router.

Using a Managed Switch to Create a VLAN

A *virtual LAN (VLAN)* is like a LAN inside a LAN. However, just as the name implies, it is created virtually, not with extra physical hardware.

The benefits of creating a VLAN include the following:

- Improved LAN security, because broadcast traffic is limited to specific ports
- The ability to group workstations or servers based on needs, not physical location
- Improved network performance for each separate broadcast domain

Imagine this scenario. You are the administrator for an organization that has several departments including the sales and finance departments. Recently financial data was leaked. The source of the leak is unknown, but security is being tightened everywhere financial data flows. Traffic from the computers in the financial department needs to be isolated with the least cost.

Many other security measures could be taken, but this explanation is focused only on the creation of VLANs.

Figure 8.8 shows the configuration before creating the VLAN. Although the switch is actually hosting many more computers, the figure shows only four computers for simplicity's sake. There shouldn't be any surprises in this diagram to you except that it's drawn with the switch on one side and the computers on the other.

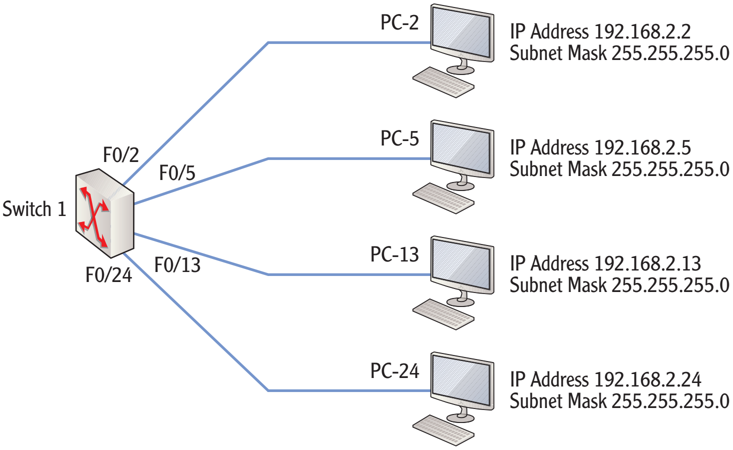

FIGURE 8.8 Connecting computers with a switch

The goal is to create two separate VLANs with this switch—one for sales and one for finance. VLANs need VLAN identifiers (VLAN IDs) and VLAN names, and Table 8.1 shows the VLAN configuration. The sales department will be on VLAN ID 2 using ports F0/0 through F0/11 and a network ID of 192.168.2.0/24. The finance department will be on VLAN ID 4 using ports F0/12 through F0/24 with a network ID of 192.168.4.0/24.

Network IDs are derived from the IP address and subnet mask. Chapter 5 explained how to calculate the network ID.

TABLE 8.1 VLAN configuration

VLAN ID	VLAN name	Port range	Subnet range	Subnet ID
2	Sales	F0/0-F0/11	192.168.2.0/24	192.168.2.0
4	Finance	F0/12-F0/24	192.168.4.0/24	192.168.4.0

Figure 8.9 shows this VLAN configuration. Even though the switch is shown twice, there is only one physical switch. However, it is using specific ports outlined in Table 8.1 to create the two VLANs.

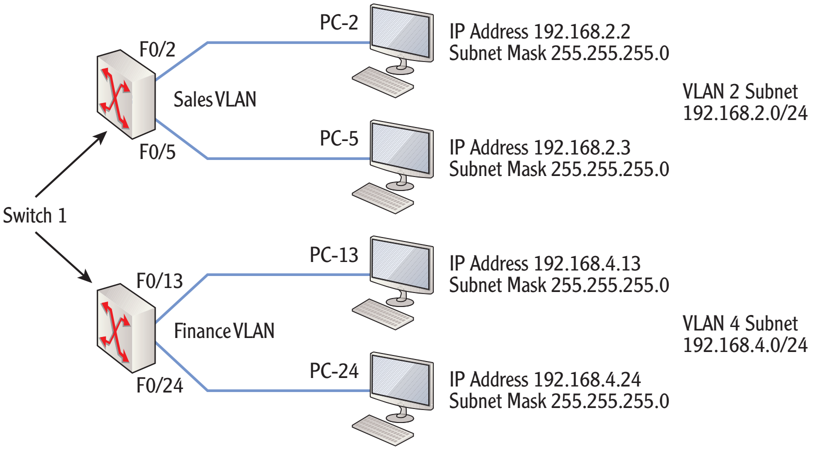

FIGURE 8.9 Creating VLANs with a switch

You would ensure that the financial computers are connected only to switch ports F0/12 through F0/24 and their IP addresses are changed so that they have a network ID of 192.168.4.0/24. Additionally, you need to ensure that all the sales computers are connected only to ports F0/0 through F0/12.

Although this example shows the basics of creating a VLAN, VLANs can be much more complex. For example, VLANs can be created to span multiple switches.

Consider Figure 8.10. This shows four departments connected with five switches. Each switch is dedicated to a specific department, but as the company grows, more salespeople are added than the office space can support. Some salespeople are sitting in the office space where the HR switch is connected. However, you can create VLANs so that these salespeople are virtually connected to the sales switch.

The following are some basic points to remember with VLANs:

- A VLAN must have at least two ports before traffic can flow, but it can have more.
    - You can create 24 two-port VLANs on a 48-port switch.
    - You can create two 24-port VLANs on a 48-port switch.
    - You can create any mixture as long as each VLAN has at least two ports.
- All ports don't have to be used in a VLAN.
- VLANs can span multiple switches.

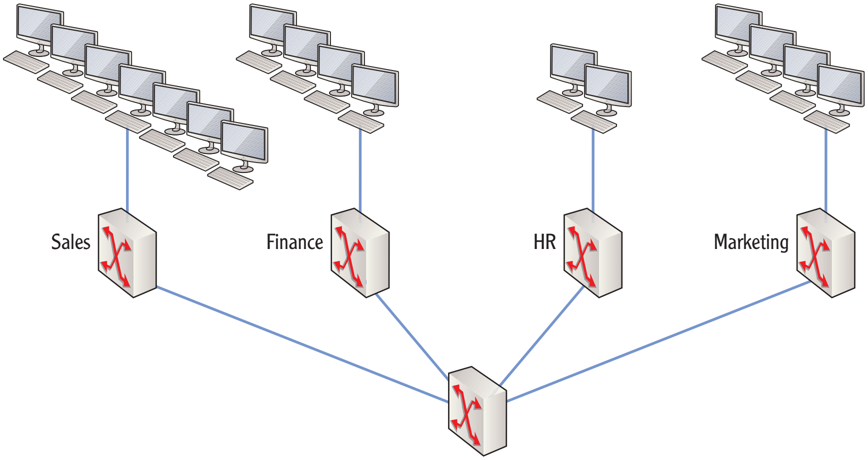

FIGURE 8.10 Multiple switches connecting different departments

Exploring Switch Speeds

A major consideration when managing and purchasing a switch is the transmission speed. This indicates the bandwidth of the switch, or how much data it can process at a time. Speeds are rated in megabits per second (Mbps) or gigabits per second (Gbps), per physical port.

You'll often hear the terms *bandwidth* and *speeds* used interchangeably. You can compare this to traffic on a highway. A single-lane road may be able to handle 100 cars an hour. However, if traffic increases and you need to handle 1,000 cars an hour, you can widen the road and add more lanes. The first 100 cars may or may not reach their destination any faster than before the road was widened, but as traffic increases, more cars will be able to get there faster because of the extra bandwidth of the road. Similarly, you can increase the amount of traffic a network can handle by increasing the bandwidth of the devices on the network.

The IEEE 802.3 standard identifies several basic speeds you're likely to see in networks today. Although more speeds are possible, Table 8.2 shows common speeds used with twisted-pair cables.

TABLE 8.2 Ethernet speeds

Protocol	Speed	Comments
IEEE 802.3	10 Mbps	10 million bits per second
IEEE 802.3u	100 Mbps	100 million bits per second

(Continues)

TABLE 8.2 *(Continued)*

Protocol	Speed	Comments
IEEE 802.3z	1000 Mbps	1000 million bits per second (1 gigabit)
IEEE 802.3an	10 Gbps	10 gigabits per second

There are three important speeds to pay attention to when looking at a switch:

- Transmission speed
- Uplink speed
- Backplane speed

These speeds are discussed in detail in the following sections.

Many switches use autosense to detect the speed of connected devices. For example, a switch rated at 100/1000 can operate at either 100 Mbps or 1000 Mbps on each port.

Identifying Transmission Speeds

Switches are commonly represented with port speeds such as 10/100 Mbps or 100/1000 Mbps. The 10/100 means that a port may operate at either 10 Mbps or 100 Mbps, and the 100/1000 means that a port can operate at either 100 Mbps or 1000 Mbps.

The limiting factors are the capabilities of the end devices and the cable grade that is used. In other words, CAT 5 twisted-pair cable can't be used for 1000 Mbps, though CAT 5E can.

Chapter 7 presented different categories of twisted-pair cables.

Similarly, if a computer has a 10 Mbps NIC, the switch can send data to the NIC only at 10 Mbps no matter how fast the switch is.

High-speed switches are available. Of course, they are more expensive. If you have a group of users who need to share large files, stream audio and video, or use Voice over IP (VoIP), it's worth getting the high-speed switches. If you do, you also need to ensure that the connecting cable and individual NICs meet the speed requirements.

If you are using a managed switch, you can manually configure individual ports for speed. Some ports could be set at 10 Mbps, some at 100 Mbps, and some at 1000 Mbps, as long as the switch supports all the speeds. Additionally, you can configure the ports individually for half-duplex or full-duplex.

Most network hardware available today uses full-duplex. However, you can downgrade a port to half-duplex for compatibility with legacy hardware.

Autosense for speed between the PC and the port is a common option with many switches. In other words, you don't have to set the speed, but the switch automatically determines the best settings for optimal speed.

Understanding the Uplink Port

An *uplink port* is a special port on a switch used to connect the switch to another switch or to another device. In contrast, other ports on the switch are called *access links*.

Uplink ports offer scalability by allowing you to add switches to the network in a daisy chain. You can also use the uplink port to connect the switch to a router for access to other subnets.

Uplink Ports and Crossover Cables

An uplink port is wired so that a straight-through cable can be used to connect it to other switches or routers. However, if you're using a regular port, it's not wired to connect two switches, and you'll need to use a crossover cable.

Many new switches can automatically sense whether a crossover or straight-through connection is needed and configure the switch internally using Auto-MDIX. Other devices use the MDI/MDI-X button that you can toggle to change the port from straight-through to crossover.

The uplink port may be labeled as an uplink port. However, on some switches the port is not labeled but instead shares the capability with another port. For example, on some smaller switches, one of the ports may include a push button labeled as MDI/MDI-X (for medium dependent interface/medium dependent interface crossover). When MDI is selected, it works as a regular port. When MDI-X is selected, it works as an uplink port. If the MDI/MDI-X button is not available, you may need to use a crossover cable.

Chapter 7 identified the differences between crossover cables and straight-through cables.

Many switches offer the ability to bundle access links together to act as a single link between switches. If your switch has a bundling option, you can configure multiple ports together as a single link. This uses the *link aggregation control protocol (LACP)*.

LACP defined in IEEE 802.3ax forms a single logical channel between devices with multiple physical links. For example, with LACP enabled, you could bundle five 100 Mbps ports as one logical link. This gives an effective throughput in this trunk of 500 Mbps. If full-duplex is used (and it normally is), your effective data throughput is 1 Gbps. Table 8.3 shows some possible speeds when bundling ports into a trunk.

Chapter 3 covered half-duplex and full-duplex. Most cables and network devices use full-duplex by default, but managed switches can be configured to match the existing hardware.

TABLE 8.3 Trunking throughput speeds

No. of ports	Port speed	Duplex setting	Effective throughput at uplink
One	10	Half	10 Mbps
One	100	Full	200 Mbps
Five	100	Half	500 Mbps
Five	100	Full	1000 Mbps (Gigabit)

Identifying Backplane Speed

Another speed to consider with switches is the backplane speed. The backplane speed is the internal speed of the switch. The faster this speed is, the better the overall performance of the switch.

Backplane speed applies only to modular switches, not form factor switches. It measures how fast data is transferred between modules in the switch.

Depending on the manufacturer, backplane speed may be measured at a couple different points. The first would be the speed on the chassis where the modules plug in. This is sometimes referred to as the speed between application-specific integrated circuits (ASICs), or the ASICs speed. This is similar in concept to the bus speed on a computer.

The second backplane speed measurement is between ports on the different blades on the same chassis. This is slightly different from the ASICs speed and is sometimes referred to as the *port-to-port speed*.

Understanding Security Options

Security is never something that is done once and it's over. It's an ongoing process.

Security is required in any network. Any time data needs to be kept secret, security is required. In addition, a company without any proprietary data quickly becomes a company without revenue.

The basic security principle of security in-depth dictates that multiple layers of security are required. Additionally, security needs to be regularly reviewed and updated. Some of the common security steps you can take with switches are as follows:

- Keep network hardware protected with physical security (in a locked room, in a locked rack).
- Change default passwords on managed switches to a complex password.

- Never use blank passwords.
- Use a secure protocol (such as SSH) to remotely manage switches.
- Use SNMP version 3 (instead of just SNMP or SNMP version 2) for best security.
- Consider port security.
- Consider hardware redundancy for maximum availability of network resources.

These last two items (port security and hardware redundancy) are explained in more depth in the following two sections.

Understanding Port Security

Port security helps you restrict what devices can connect to ports on a switch. The danger is that if someone can walk into your organization and simply plug a computer into an RJ-45 jack in the wall, they can access your network. That's disconcerting to both administrators and organization executives.

One method of port security is to configure each port with the MAC address of a specific computer. Only that computer can connect. If a device with a different MAC attempts a connection, the switch refuses the connection. If you have five computers, this won't take much time. However, if you have 500 computers, it can be quite time-consuming and tedious.

Only managed switches can use port security. Unmanaged switches can't be configured for specific MACs.

An alternative is to configure the port to remember which MAC addresses it learned and to set a threshold for the maximum number of addresses allowed. For example, you may set the threshold at one or two addresses, and any MAC learned after that would trip an alarm. The alarm may notify an administrator using an SNMP trap (or error message) or may shut down communication with the port.

Another element of port security is ensuring that unused ports are not enabled. For example, you may have a 48-port switch that is cabled to 48 RJ-45 wall jacks in your organization. However, you are currently using only 40 of the RJ-45 jacks. The other jacks don't have computers attached. The switch ports where these jacks are connected should be disabled. This prevents someone from coming in, plugging in a computer to an empty jack, and accessing your network.

Fault tolerance means that a fault can occur and a system can tolerate it. Fault tolerance can be implemented at the disk level, the server level, the site level, and more.

Planning Hardware Redundancy

The switch in a typical network configuration presents itself as a single point of failure. If the switch fails, *poof*, all the computers connected through the switch lose connectivity to network resources. However, you can build in fault tolerance by adding hardware redundancy.

Hardware redundancy simply means that additional components are added to ensure that the failure of one component doesn't result in a complete failure. For example, many modular switches come with more than one power supply. If one power supply fails, the switch can continue to operate. In some switches, the power supplies are completely redundant, meaning that the switch will continue to operate with no loss of capability. In other switches, a failure of one power supply may affect only some of the ports on the switch but not all of them.

It's also possible to configure a fail-over state in a trunked environment with LACP. In other words, if one or more ports fail in the combined trunk, the capability can be switched to another port.

Obviously, adding redundant capabilities costs more money. The majority of the time, the redundant component is not actively utilized but is instead just there in case a failure occurs. When considering hardware redundancy for switches, you need to consider how critical connectivity is for the devices.

It may be that a failure of a particular switch results in immediate loss of critical resources and lost revenue. It's worthwhile adding additional redundancy for this switch. On the other hand, the loss of a switch may not be critical. You may be able to replace it within a day without any impact on the business's bottom line. In this case, the added cost of redundancy is not necessary.

The Essentials and Beyond

In this chapter, you learned about many of the capabilities and inner workings of switches. You learned how the switch uses the MAC addresses to create a MAC address table. It then creates multiple collision domains by sending data only to the destination device based on the MAC address instead of to all devices connected to the switch. You also learned how some switches are managed and can be configured, while other switches simply work by plugging them in. Some switches operate as layer 2 and layer 3 switches, and you can create VLANs within a managed switch. Port security can control what devices can connect to a switch.

Additional Exercises

- Draw a diagram of computers in your network up to the nearest router. Identify the different broadcast domains and the different collision domains.
- Locate a switch in your network. Count the number of normal ports it uses.
- Locate a switch in your network. Identify the uplink port. How is it labeled?
- Locate an active switch in your network. Identify the physical security used to protect it.

(Continues)

The Essentials and Beyond *(Continued)*

To compare your answers to the author's, please visit **www.sybex.com/go/networkingessentials**.

Review Questions

1. A ______________ switch is expandable. You can add ports by adding components.
2. You are looking at a drawing of a 100 Mbps switch and want to identify what device is connected to the first port. How will the port be labeled?

   A. E0
   B. F0
   C. F1
   D. Gi1

3. True or false. Layer 2 switches create separate broadcast domains.
4. Your network includes computers connected via hubs. You want to reduce the number of collisions to the least number possible. What should you do?

   A. You should replace the hubs with bridges.
   B. You should replace the hubs with managed hubs.
   C. You should replace the hubs with switches.
   D. You should replace the hubs with firewalls.

5. True or false. A managed switch requires less administrative overhead than an unmanaged switch.
6. A layer 3 switch functions just like a ______________.
7. What does a switch maintain to track the location of computers?

   A. A MAC address table
   B. A routing table
   C. A layer 3 table
   D. A managed table

8. A switch has 48 ports. How many VLANs can you create with it?

   A. 2
   B. 12
   C. 24
   D. 48

9. A 100 Mbps switch is configured to combine five ports using LACP with full-duplex. What is the effective throughput at the uplink?

   A. 100 Mbps
   B. 500 Mbps
   C. 600 Mbps
   D. 1000 Mbps

10. You want to ensure that only known computers can connect to a switch. What should you implement?