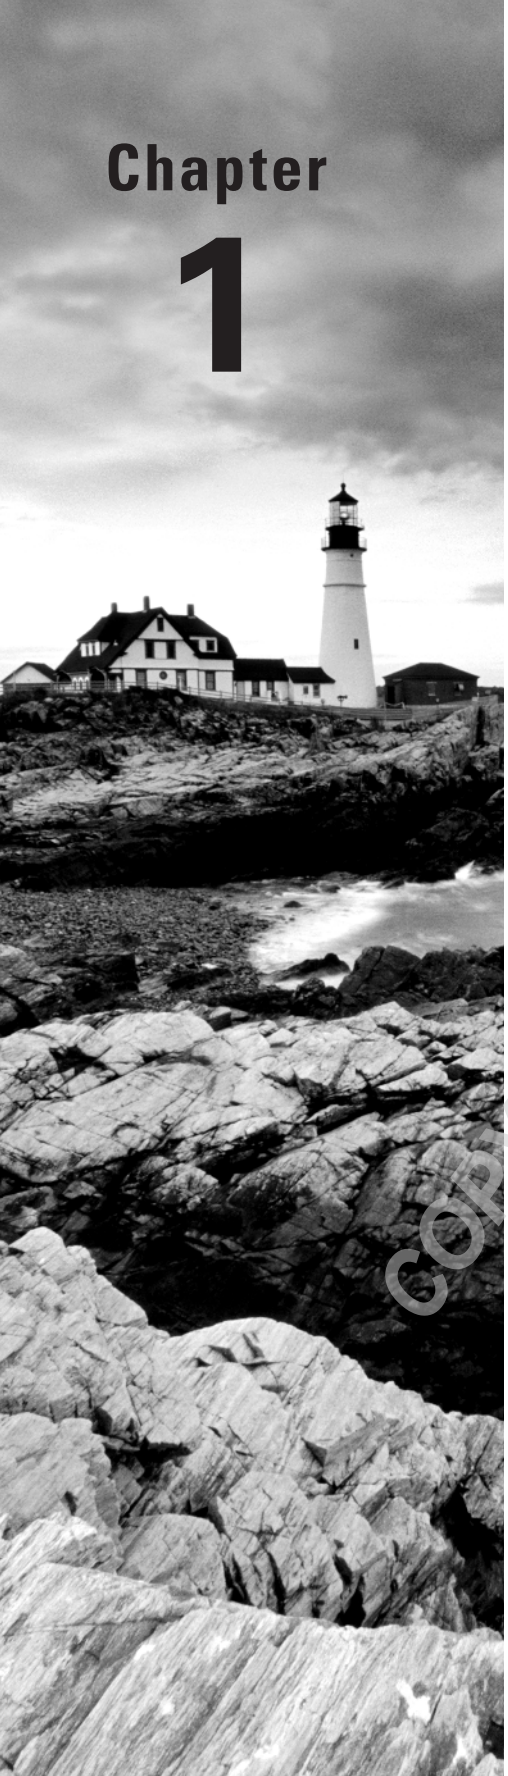

Chapter 1

Network Fundamentals

WE'LL COVER THE FOLLOWING CCNA EXAM TOPICS IN THIS CHAPTER:

1.0 Network Fundamentals

✓ **1.1 Explain the role and function of network components**

- 1.1.a Routers
- 1.1.b L2 and L3 switches
- 1.1.c Next-generation firewalls and IPS

✓ **1.2 Describe characteristics of network topology architectures**

- 1.2.a 2 tier
- 1.2.b 3 tier
- 1.2.c Spine-leaf
- 1.2.d WAN
- 1.2.e Small office/home office (SOHO)

✓ **1.3 Compare physical interface and cabling types**

- 1.3.a Single-mode fiber, multimode fiber, copper
- 1.3.b Connections (Ethernet shared media and point-to-point)
- 1.3.c Concepts of PoE

COPYRIGHTED MATERIAL

This chapter is really an internetworking review, focusing on how to connect networks together using Cisco routers and switches. As a heads up, I've written it with the assumption that you have a bit of basic networking knowledge and/or have read the first chapter in the CCNA Part 1: Understanding Cisco Networking Technologies study guide.

That said, there isn't a whole lot of new material here, but even if you're a seasoned network professional, you should still read through *all* chapters to make sure you get how the objectives are currently covered.

Let's start by defining exactly what an internetwork is: You create an internetwork when you connect two or more networks via a router and configure a logical network addressing scheme with protocols like IP or IPv6.

We'll move on to covering network components like routers and switches and defining what exactly creates a Small Office Home Office Network (SOHO). After that, we'll touch on Next Generation Firewalls (NGFWs) and network architect models, and then I'll guide you through an overview of Ethernet and the wiring used in Local Area Networks (LANs) and Wide Area Networks (WANs).

To find your included bonus material, as well as Todd Lammle videos, practice questions & hands-on labs, please see `www.lammle.com/ccna`

Network Components

So why is it so important to learn Cisco internetworking anyway?

Networks and networking have grown exponentially over the past 20 years, and understandably so. They've had to evolve at light speed just to keep up with huge increases in basic, mission-critical user needs from simply sharing data and printers to bigger burdens like multimedia remote presentations, conferencing, and the like. Unless everyone who needs to share network resources is located in the same office space, the challenge is to connect relevant networks so all users can share the wealth of whatever services and resources they need, on site or remotely.

Figure 1.1 shows a basic *local area network (LAN)* connected via a *hub*, which is basically an antiquated device, which connects wires together and is typically used in SOHO networks.

Keep in mind that a simple SOHO network like this one would be considered one collision domain and one broadcast domain.

FIGURE 1.1 A very basic SOHO network

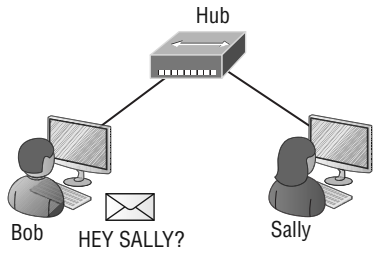

Things really can't get much simpler than this. And yes, though you can still find this configuration in some home networks, even many of those as well as the smallest business networks are more complicated today.

Routers, Switches, and Oh So SOHO!

Figure 1.2 shows a network that's been segmented with a switch, making each network segment that connects to the switch its own separate collision domain. Doing this results in a lot less chaos!

FIGURE 1.2 A switch can break up collision domains.

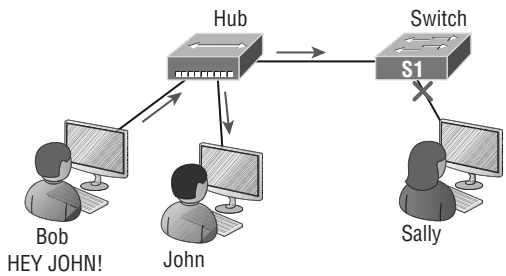

So, this is a great start, but I really want you to note that this network is still just one, single broadcast domain. This means we've really only reduced our PC's chaos—not eliminated it.

For example, if there's some sort of vital announcement that everyone in our network neighborhood needs to hear about, it will definitely still get loud! You can see that the hub used in Figure 1.2 just extended the one collision domain from the switch port. The result is that John received the data from Bob but, happily, Sally did not, which is good because Bob intended to talk with John directly. If he had needed to send a broadcast instead, everyone, including Sally, would have received it, causing unnecessary congestion.

Here's a list of some of the things that commonly cause LAN traffic congestion:

- Too many hosts in a collision or broadcast domain
- Broadcast storms
- Too much multicast traffic
- Low bandwidth
- Adding hubs for connectivity to the network
- A bunch of ARP broadcasts

Take another look at Figure 1.2 and make sure you see that I extended the main hub from Figure 1.1 to a switch in Figure 1.2. I did that because hubs don't segment a network; they just connect network segments. Basically, it's an inexpensive way to connect a couple of PCs, which can work for really simple home use and troubleshooting, but that's about it!

As our community grows, we'll need to add more streets along with traffic control and even some basic security. We'll get this done by adding routers because these convenient devices are used to connect networks and route packets of data from one network to another. Cisco became the de facto standard for routers because of its unparalleled selection of high-quality router products and fantastic service. So never forget that by default, routers are basically employed to efficiently break up a *broadcast domain*—the set of all devices on a network segment, which are allowed to "hear" all broadcasts sent out on that specific segment.

Figure 1.3 depicts a router in our growing network, creating an internetwork and breaking up broadcast domains.

FIGURE 1.3 Routers create an internetwork.

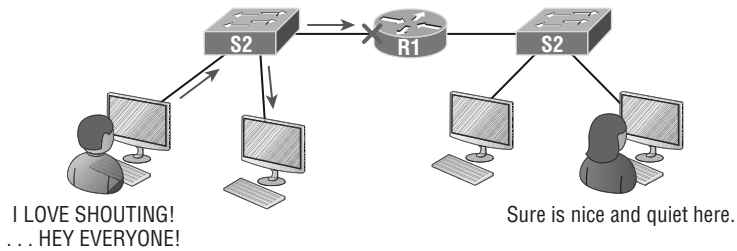

The network in Figure 1.3 is actually a pretty cool little network. Each host is connected to its own collision domain because of the switch, and the router has created two broadcast domains. So now Sally is happily living in peace in a completely different neighborhood, no longer subjected to Bob's incessant shouting! If Bob wants to talk with Sally, he has to send a packet with a destination address using her IP address—he cannot broadcast for her!

But there's more… Routers provide connections to *wide area network (WAN)* services as well via a serial interface for WAN connections—specifically, a V.35 physical interface on a Cisco router.

Let me make sure you understand why breaking up a broadcast domain is so important. When a host or server sends a network broadcast, every device on the network must read and process that broadcast—unless you have a router. When the router's interface receives this broadcast, it can respond by basically saying, "no thanks," and discard the broadcast without forwarding it on to other networks. Even though routers are known for breaking up broadcast domains by default, it's important to remember that they break up collision domains as well.

There are two advantages to using routers in your network:

- They don't forward broadcasts by default.
- They can filter the network based on layer 3 (Network layer) information such as an IP address.

Conversely, we don't use layer 2 switches to create internetworks because they don't break up broadcast domains by default. Instead, they're employed to add functionality to a network LAN. The main purpose of these switches is to make a LAN work better—to optimize its performance—providing more bandwidth for the LAN's users. Also, these switches don't forward packets to other networks like routers do. Instead, they only "switch" frames from one port to another within the switched network. And don't worry, even though you're probably thinking, "Wait—what are frames and packets?" I promise to completely fill you in later in this chapter. For now, think of a packet as a package containing data.

Okay, so by default, switches break up collision domains, but what are these things? A *Collision domain* is an Ethernet term used to describe a network scenario in which one device sends a packet out on a network segment and every other device on that same segment is forced to pay attention to it no matter what. This isn't efficient because if a different device tries to transmit at the same time, a collision will occur, requiring both devices to retransmit, one at a time—not good! And this happens a lot in a hub environment, where each host segment connects to a hub that represents only one collision domain and a single broadcast domain. By contrast, each and every port on a switch represents its own collision domain, allowing network traffic to flow much more smoothly.

Layer 2 switching is considered hardware-based bridging because it uses specialized hardware called an *application-specific integrated circuit (ASIC)*. ASICs can run up to high gigabit speeds with very low latency rates.

Latency is the time measured from when a frame enters a port to when it exits a port.

Switches read each frame as it passes through the network. The layer 2 device then puts the source hardware address in a filter table and keeps track of which port the frame was received on. This information (logged in the bridge's or switch's filter table) is what helps the machine determine the location of the specific sending device.

Figure 1.4 shows a switch in an internetwork and how John is sending packets to the Internet. Sally doesn't hear his frames because she's in a different collision domain. The destination frame goes directly to the default gateway router, so Sally doesn't even see John's traffic.

FIGURE 1.4 Switches work at layer 2.

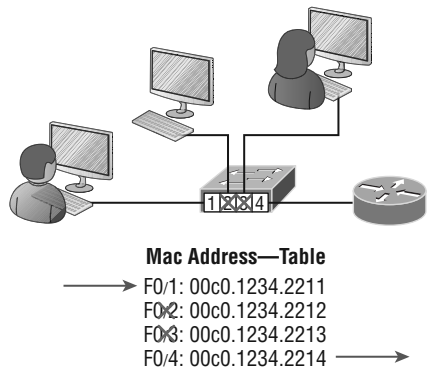

The real estate business is all about location, location, location, and it's the same way for layer 2 and layer 3 devices. Although both need to be able to negotiate the network, it's crucial to remember that they're concerned with very different parts of it. Primarily, layer 3 machines, like routers, need to locate specific networks, whereas layer 2 machines like switches and bridges need to eventually locate specific devices. So, networks are to routers as individual devices are to switches and bridges. And routing tables that "map" the internetwork are for routers as filter tables that "map" individual devices are for switches and bridges.

After a filter table is built on the layer 2 device, it will forward frames only to the segment where the destination hardware address is located. If the destination device is on the same segment as the source host, the layer 2 device will block the frame from going to any other segments. If the destination is on a different segment, the frame can be transmitted only to that segment. This is called *transparent bridging*.

When a switch interface receives a frame with a destination hardware address that isn't found in the device's filter table, it will forward the frame to all connected segments. If the unknown device that was sent the "mystery frame" replies to this forwarding action, the switch updates its filter table regarding that device's location. But in the event the destination address of the transmitting frame is a broadcast address, the switch will forward all broadcasts to every connected segment by default.

All devices that the broadcast is forwarded to are considered to be in the same broadcast domain. This can be a problem because layer 2 devices propagate layer 2 broadcast storms that can seriously choke performance, and the only way to stop a broadcast storm from propagating through an internetwork is with a layer 3 device—a router!

Next-Generation Firewalls and IPS

Today's networks definitely need security, and as our network grows, we'll need to increase protection for it. Just like we'd add locks to our doors and windows, and then maybe a fence, then even a bigger fence—with a locked gate and even some barbed wire to top it off… We can go on and on here. You get the picture.

There are new devices out that are actually seriously solid firewalls. I'll mention Next Generations Firewalls (NGFW) providing full layer-7 inspection, as though it's just a bump in the wire (meaning little delay), which is mostly true. However, it's totally true that every company, including Cisco, markets their devices like this.

Figure 1.5 pictures devices in a small network and how a basic firewall or NGFW can be placed to provide security in your network.

FIGURE 1.5 Physical components of a network

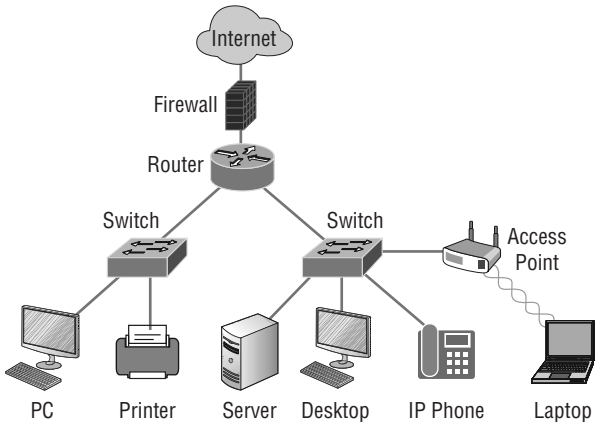

Firewall and NGFW design can be pretty complicated, but we don't need to get into the weeds here. Since this is a Cisco book, we're going to stick with Cisco technologies for our firewall. Cisco has a Next Generation Firewall (NGFW) called Firepower that they acquired from a company called SourceFire in 2013.

Now understand this is going to be a brief introduction to NGFW and Intrusion Prevention Systems (IPS). Why? Because I have a two-book series on CCNP Security Securing Cisco Network Firewalls (SCNF) that covers this topic in depth with well over 1500 pages of information! That's a lot of firewall info. We'll just get our feet wet for now.

So let's start by defining a NGFW and what it has to do with Intrusion Prevention Systems (IPS). NGFWs are considered a third-generation firewall technology that provides full packet reassembly and deep-packet inspect up to and through layer 7.

NGFW's are popular because they permit application visibility and control (AVC) as well as offer intrusion prevention system (IPS) policies, which help us look for attacks on known client vulnerabilities.

And no one said this technology is cheap. For example, the newer firewalls can provide SSL decryption, which sounds simple, but there's a catch. To be able provide that kind of

shield at close to wire speed you've got to have hardware encryption acceleration capability, which will cost you plenty!

The NGFWs today have everything but the kitchen sink in their code just to stay competitive, and this causes all sorts of issues for the manufacturers when struggling to keep up with the market.

Here's a taste… All NGFWs must, at a minimum, include the following:

- Be router and switch compatible (L2/L3)
- They must add packet filtering with IPS and Malware inspection capability
- Provide Network Address Translation (NAT)
- Permit stateful inspection
- Permit Virtual Private Networks (VPNs)
- Provide URL and Application filtering
- Implement QoS
- Third-party integration
- Support for REST API

That's not a short list, and the items in it are all absolutely required because NGFWs must pack a powerful security punch in order to lock our modern networks down tight!

Figure 1.6 shows a Cisco Firepower NGFW blocking an attacker trying to exploit a vulnerability on my network, which IPS stopped dead. The red line on top indicates all the attacks and the blue line is for the data. Wow—that's a lot of attacks!

FIGURE 1.6 NGFW can stop attacks in real time.

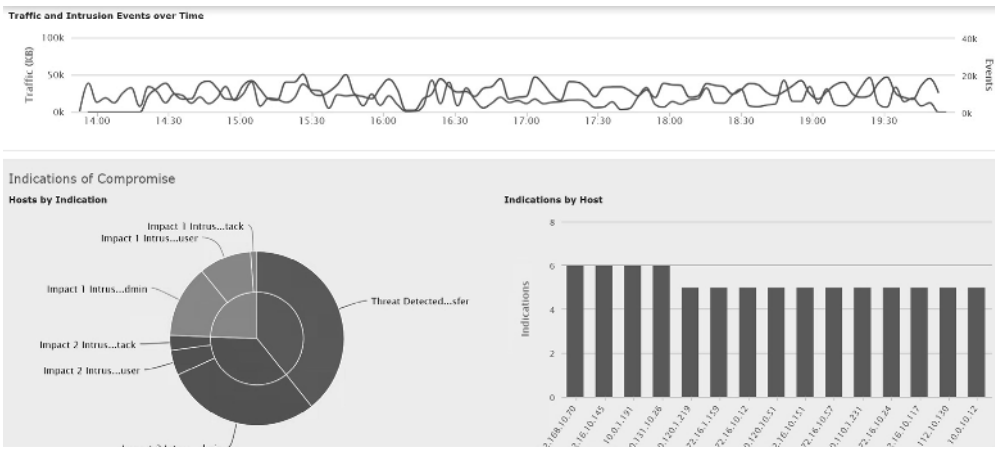

Now, let's dig into those attacks. Well see how the Cisco Firepower NGFW came to the rescue by blocking all those attacks via my IPS policy, as shown in Figure 1.6! Figure 1.7 displays some of the events that were caught by Cisco Firepower.

FIGURE 1.7 Cisco IPS policy to the rescue!

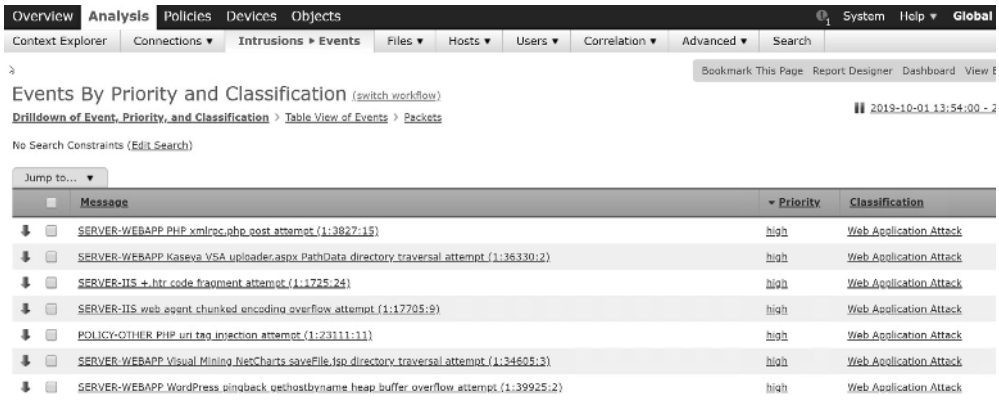

And just so you know, some of those attacks shown in Figure 1.7 were some really serious ones! Figure 1.8 displays all the actual packets that were dropped.

FIGURE 1.8 Cisco Firepower IPS policy dropped the bad guys packets!

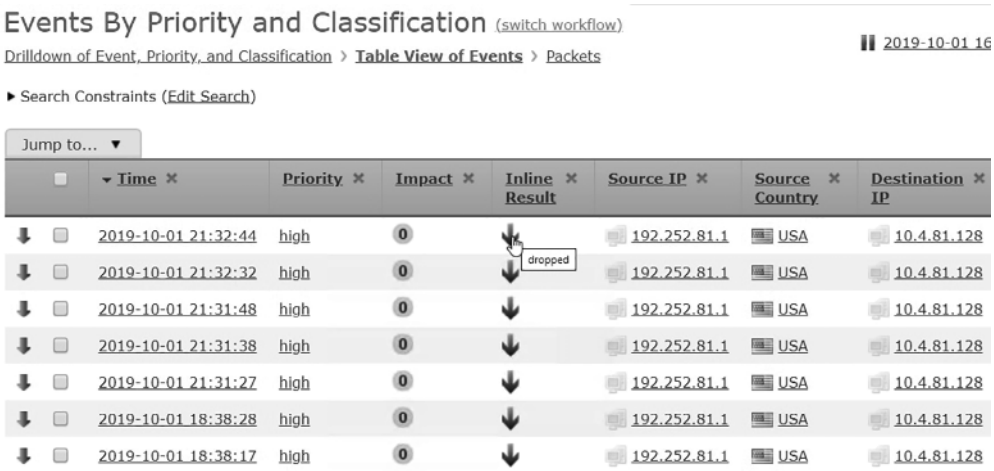

NGFWs perform a deeper inspection compared to your traditional firewall. For instance, the stateful ASA that Cisco is moving away from is being replaced with new Firepower Threat Defense (FTD) devices, which are true NGFW devices.

Network Topology Architectures

Most of us were exposed to hierarchy early in life, and anyone with older siblings learned what it was like to be at the bottom of it. Regardless of where you first discovered the concept of hierarchy, most of us experience it in many aspects of our lives. It's *hierarchy* that helps us understand where things belong, how things fit together, and what functions go where. It brings order to otherwise complex models. If you want a pay raise, for instance, hierarchy dictates that you ask your boss, not your subordinate, because that's the person whose role it is to grant or deny your request. So basically, understanding hierarchy helps us discern where we should go to get what we need.

Hierarchy offers a lot of the same benefits in network design that it does in life. When used properly, it makes networks more predictable and helps us to define which areas should perform certain functions. For example, you can use tools like access lists at certain levels within hierarchical networks and avoid them at others.

Large networks can be extremely complicated, involving multiple protocols, detailed configurations, and diverse technologies. Hierarchy helps us summarize a complex collection of details into an understandable model, bringing order from the chaos. Then, as specific configurations are needed, the model dictates the correct way to apply them.

The Cisco Three-Layer Hierarchical Model (3-Tier)

The Cisco hierarchical model can help you design, implement, and maintain a scalable, reliable, cost-effective hierarchical internetwork.

Cisco defines three layers of hierarchy, as shown in Figure 1.9, each with specific functions, and it's referred to as a 3-tier network architecture.

Each layer has specific responsibilities. Keep in mind that the three layers are logical, so they aren't necessarily physical devices. Consider the OSI model, another logical hierarchy. Its seven layers describe functions but not necessarily protocols, right? Sometimes a protocol maps to more than one layer of the OSI model, and sometimes multiple protocols communicate within a single layer.

In the same way, when we build physical implementations of hierarchical networks, we may have many devices in a single layer, or there may be a single device performing functions at two layers. Just remember that the definition of the layers is logical, not physical!

Let's take a closer look at each of the layers now.

FIGURE 1.9 The Cisco hierarchical model

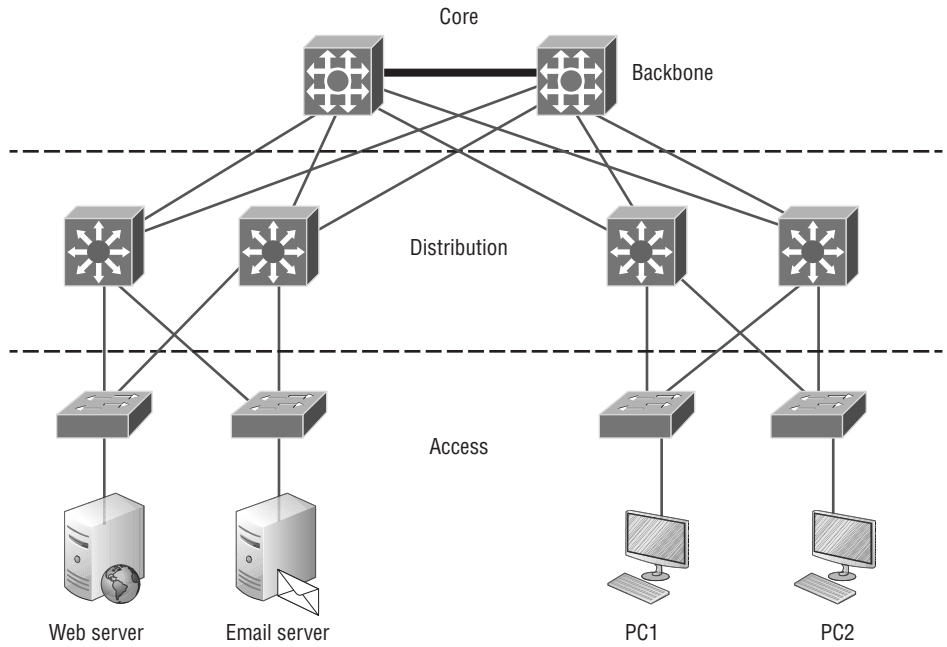

The Core Layer

The *core layer* is literally the core of the network. At the top of the hierarchy, this layer is responsible for transporting large amounts of traffic both reliably and quickly. The prime purpose of the network's core layer is to switch traffic as fast as possible. The traffic transported across the core is common to a majority of users, but user data is processed at the distribution layer, which forwards the requests to the core if needed.

If there's a failure in the core, *every single user* can be affected! This is why fault tolerance at this layer is so important. The core is likely to see large volumes of traffic, so speed and latency are driving concerns here. Given the function of the core, some vital design specifics come into view. Let's start with things we don't want to happen here:

- Never do anything to slow down traffic. This includes making sure you don't use access lists, perform routing between virtual local area networks, or implement packet filtering.
- Don't support workgroup access here.
- Avoid expanding the core, e.g., adding routers as the internetwork grows. If performance becomes an issue in the core, go with upgrades over expansion.

Here's a list of goals we want to achieve as we design the core:

- Design the core for high reliability. Consider data-link technologies that facilitate both speed and redundancy, like Gigabit Ethernet with redundant links or even 10 Gigabit Ethernet.
- Design with speed in mind. The core should have very little latency.
- Select routing protocols with lower convergence times. Fast and redundant data-link connectivity is no help if your routing tables are shot!

The Distribution Layer

The *distribution layer* is sometimes referred to as the workgroup or aggregation layer and is the communication point between the access layer and the core. The primary functions of the distribution layer are to provide routing, filtering, and WAN access and to determine how packets can access the core, if needed. The distribution layer must determine the fastest way that network service requests are handled—for instance, how a file request is forwarded to a server. After the distribution layer determines the best path, it forwards the request to the core layer if necessary. The core layer then quickly transports the request to the correct service.

The distribution layer is where we implement policies for the network because we have a lot of flexibility in defining network operation here. There are several things that should generally be handled at the distribution layer:

- Routing
- Implementing tools (like access lists), packet filtering, and queuing
- Implementing security and network policies, including address translation and firewalls
- Redistributing between routing protocols, including static routing
- Routing between VLANs and other workgroup support functions
- Defining broadcast and multicast domains

At the distribution layer, it's key to avoid anything limited to functions exclusively belonging to one of the other layers!

The Access Layer

The *access layer* controls user and workgroup access to internetwork resources and is sometimes referred to as the *desktop layer*. The network resources most users need are available locally because the distribution layer handles any traffic for remote services.

Here are some of the tasks the access layer carries out:

- Continued (from distribution layer) use of access control and policies
- Creation of separate collision domains (microsegmentation/switches)
- Workgroup connectivity into the distribution layer
- Device connectivity

- Resiliency and security services
- Advanced technology capabilities (voice/video, etc.)
- QoS Marking

Technologies like Gigabit or Fast Ethernet switching are frequently seen in the access layer as well.

I can't stress this enough—just because there are three separate layers does not imply three separate devices! There could be fewer or there could be more. After all, this is a *layered* approach.

Collapsed Core (2-Tier)

The Collapsed core design is also referred to as 2-tier because it's only 2-layers. But in concept, it's like the 3-tier only less expensive and geared for smaller companies. The design is meant to maximize performance and user availability to the network, while still allowing for design scalability over time.

In a 2-tier, the distribution is merged with the core layer, as shown in Figure 1.10.

FIGURE 1.10 Real-life collapsed core (2-tier) Image

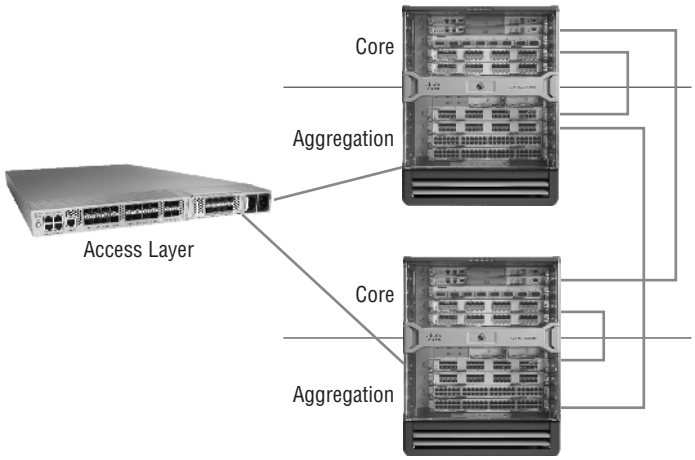

Here you see the Core and Distribution (also called Aggregation) are both running on the same large enterprise switch. The Access layer switches connect into the enterprise switch, only in the defined Aggregation ports.

This design is much more economical and it's still very functional in a campus environment, where your network many not really grow significantly larger over time. It's known as a "collapsed core" and refers to a design in which the distribution layer and core layer functions are implemented by a single device.

The big reason the collapsed core design exists is for reducing network costs, while maintaining most of the benefits of the three-tier hierarchical model.

Spine-Leaf

I've been writing about Cisco's three-tier network design for a really long time, and I just did again. But today's data centers demand a new design, and one was finally created that works really well called a *leaf-and-spine* topology. This design is still pretty old as of this writing, it's just not decades old!

Here's how it works: Your typical data center has racks filled with servers. In the leaf-and-spine design, there are switches found at the top and end of each rack that connect to these servers, with a server connecting into each switch for redundancy.

People refer to this as a top-of-rack (ToR) design because the switches physically reside at the top of a rack. Figure 1.11 pictures a top of rack network design.

FIGURE 1.11 Top of Rack Network Design

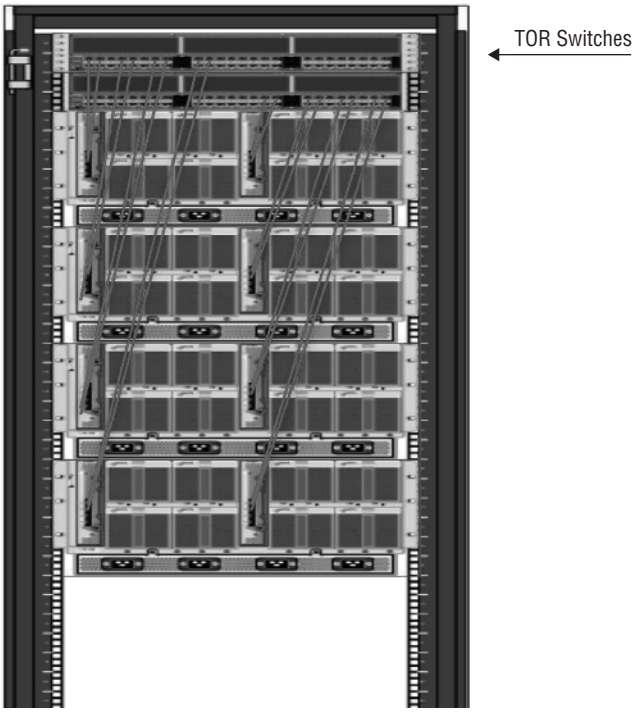

These ToR switches act as the leaves within the leaf-and-spine topology. The ports in the leaf switches connect to a node, e.g., a server in the rack, a firewall, a load-balancing appliance, or a router leaving the data center, as well as to the spine switch. Check out Figure 1.12.

FIGURE 1.12 Spine-leaf design

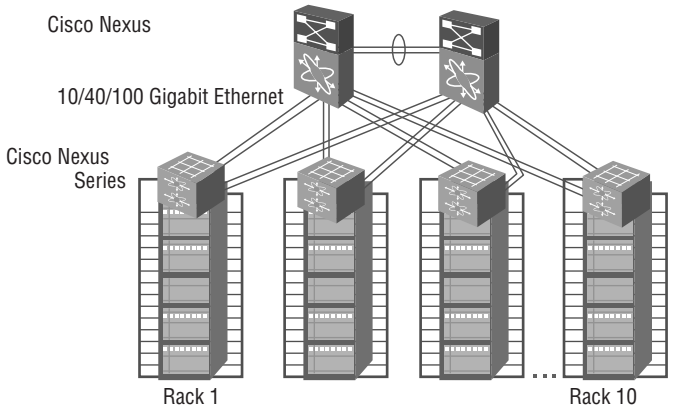

You can see each leaf switch connecting to every spine switch, which is great because it means that we no longer need a gazillion connections between switches. Keep in mind that the spine only connects to leaf devices, not to servers or end devices.

And interestingly enough, when you connect your ToR data center switches in a leaf-and-spine topology, all of your switches are the same distance away from one another (single switch hop).

WAN

Let's begin WAN basics by asking, what's the difference between a wide area network (WAN) and a local area network (LAN)? Clearly there's the distance factor, but modern wireless LANs can cover some serious turf, so there's more to it than that. What about bandwidth? Here again, really big pipes can be had for a price in many places, so that's not it either. So what's the answer we're looking for?

A major distinction between a WAN and a LAN is that while you generally own a LAN infrastructure, you usually lease a WAN infrastructure from a service provider. Modern technologies sometimes blur this characteristic somewhat, but this factor still fits neatly into the context of Cisco's exam objectives.

There are several reasons why WANs are necessary in corporate environments today. LAN technologies provide pretty solid speeds—10/25/40/100Gbps is now common—and they're definitely pricey. The thing is, these solutions really only work well in relatively small geographic areas. We still need WANs in a communications environment because some business needs require connections to remote sites for many reasons:

- People in the regional or branch offices of an organization need to be able to communicate and share data.
- Organizations often want to share information with other organizations across large distances.
- Employees who travel on company business frequently need to access information that resides on their corporate networks.

Here are three major characteristics of WANs:

- WANs generally connect devices that are separated by a broader geographic area than a LAN can serve.
- WANs use the services of carriers like telcos, cable companies, satellite systems, and network providers.
- WANs use serial connections of various types to provide access to bandwidth over large geographic areas.

The first key to understanding WAN technologies is to be familiar with the different WAN topologies, terms, and connection types commonly used by service providers to join our LAN networks together.

Defining WAN Terms

Before you run out and order a WAN service type from a provider, you really need to understand the following terms that service providers typically use. Here they are in Figure 1.13:

FIGURE 1.13 WAN terms

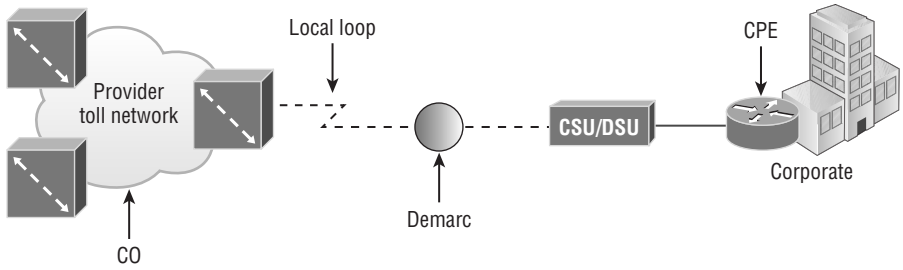

Customer premises equipment (CPE) *Customer premises equipment (CPE)* is equipment that's typically owned by the subscriber and located on the subscriber's premises.

CSU/DSU A Channel Service Unit/Data Service Unit (CSU/DSU) is a device that's used to connect a DTE to a digital circuit like a T1/T3 line. A device is considered DTE if it's either a source or destination for digital data—for example, PCs, servers, and routers. In Figure 1.13, the router is considered DTE because it's passing data to the CSU/DSU, which will forward the data to the service provider. Although the CSU/DSU connects to the service provider's infrastructure using a telephone or coaxial cable like a T1 or E1 line, it connects to the router with a serial cable. The most important aspect to remember for the CCNA objectives is the CSU/DSU provides clocking of the line to the router.

Demarcation point The *demarcation point* (demarc for short) is the precise spot where the service provider's responsibility ends and the CPE begins. It's generally a device in a telecommunications closet owned and installed by the telecommunications company (telco). It's your responsibility to cable (extended demarc) from this box to the CPE, which is usually a connection to a CSU/DSU.

Local loop The *local loop* connects the demarc to the closest switching office, referred to as the central office.

Central office (CO) This point connects the customer's network to the provider's switching network. Make a mental note that a *central office (CO)* is sometimes also referred to as a *point of presence (POP)*.

Toll network The *toll network* is a trunk line inside a WAN provider's network. This network is a collection of switches and facilities owned by the Internet service provider (ISP).

Optical fiber converters Even though I'm not employing this device in Figure 1.13, optical fiber converters are used where a fiber-optic link terminates to convert optical signals into electrical signals and vice versa. You can also implement the converter as a router or switch module.

Make sure you're comfortable with these terms, what they represent, and where they're located, as shown in Figure 1.13, because they're key to understanding WAN technologies.

WAN Connection Bandwidth

Next, I want you to know these basic but very important bandwidth terms used when referring to WAN connections:

Digital Signal 0 (DS0) This is the basic digital signaling rate of 64 Kbps, equivalent to one channel. Europe uses the E0 and Japan uses the J0 to reference the same channel speed. Typical to T-carrier transmission, this is the generic term used by several multiplexed digital carrier systems and is also the smallest-capacity digital circuit. 1 DS0 = 1 voice/data line.

T1 Also referred to as a DS1, a T1 comprises 24 DS0 circuits bundled together for a total bandwidth of 1.544 Mbps.

E1 This is the European equivalent of a T1 and comprises 30 DS0 circuits bundled together for a bandwidth of 2.048 Mbps.

T3 Referred to as a DS3, a T3 comprises 28 DS1s bundled together, or 672 DS0s, for a bandwidth of 44.736 Mbps.

OC-3 Optical Carrier (OC) 3 uses fiber and is made up of three DS3s bundled together. It's made up of 2,016 DS0s and avails a total bandwidth of 155.52 Mbps.

OC-12 Optical Carrier 12 is made up of four OC-3s bundled together and contains 8,064 DS0s for a total bandwidth of 622.08 Mbps.

OC-48 Optical Carrier 48 is made up of four OC-12s bundled together and contains 32,256 DS0s for a total bandwidth of 2488.32 Mbps.

Physical Interfaces and Cables

Ethernet was first implemented by a group called DIX, Digital, Intel, and Xerox. They created and implemented the first Ethernet LAN specification, which the IEEE used to create the IEEE 802.3 committee. This was a 10 Mbps network that ran on coax, then eventually twisted-pair and fiber physical media.

The IEEE extended the 802.3 committee to three new committees known as 802.3u (Fast Ethernet), 802.3ab (Gigabit Ethernet on category 5), and then finally 802.3ae (10 Gbps over fiber and coax). There are more standards evolving almost daily, like 100 Gbps Ethernet (802.3ba).

When designing your LAN, it's really important to understand the different types of Ethernet media available to you. Sure, it would be great to run TenGigabit Ethernet to each desktop and 100 Gbps between switches, but I dare you to justify the cost of that network! However, if you mix and match the different types of Ethernet media methods currently available, you can come up with a cost-effective network solution that still works really great.

The *EIA/TIA* (Electronic Industries Alliance and the newer Telecommunications Industry Association) is the standards body that creates the Physical layer specifications for Ethernet. The EIA/TIA specifies that Ethernet use a *registered jack (RJ) connector* on *unshielded twisted-pair (UTP)* cabling (RJ45). But the industry is moving toward simply calling this an 8-pin modular connector.

Every Ethernet cable type that's specified by the EIA/TIA has inherent attenuation, which is defined as the loss of signal strength as it travels the length of a cable and is measured in decibels (dB). The cabling used in corporate and home markets is measured in categories. A higher-quality cable will have a higher-rated category and lower attenuation. For example, category 6 is better than category 5 because category 6 cables have cables have more wire twists per foot and therefore less crosstalk. Crosstalk is the unwanted signal interference from adjacent pairs in the cable.

Here's a list of some of the most common IEEE Ethernet standards, starting with 10 Mbps Ethernet:

10Base-T (IEEE 802.3) 10 Mbps using category 3 unshielded twisted pair (UTP) wiring for runs up to 100 meters. Unlike with the 10Base-2 and 10Base-5 networks, each device must connect into a hub or switch, and you can have only one host per segment or wire. It uses an RJ45 connector (8-pin modular connector) with a physical star topology and a logical bus.

100Base-TX (IEEE 802.3u) 100Base-TX, most commonly known as Fast Ethernet, uses EIA/TIA category 5, 5E, or 6 UTP two-pair wiring. One user per segment and up to 100 meters long, it uses an RJ45 connector with a physical star topology and a logical bus.

100Base-FX (IEEE 802.3u) Uses fiber cabling 62.5/125-micron multimode fiber. Point-to-point topology and up to 412 meters long. It uses ST and SC connectors, which are media-interface connectors.

1000Base-CX (IEEE 802.3z) Copper twisted-pair, called twinax, is a balanced coaxial pair that can run only up to 25 meters and uses a special 9-pin connector known as the High Speed Serial Data Connector (HSSDC). This is used in Cisco's Data Center technologies.

1000Base-T (IEEE 802.3ab) Category 5, four-pair UTP wiring up to 100 meters long and up to 1 Gbps.

1000Base-SX (IEEE 802.3z) This is the implementation of 1 Gigabit Ethernet, running over multimode fiber-optic cable instead of copper twisted-pair cable, using short wavelength laser. Multimode fiber (MMF) using 62.5- and 50-micron core and uses an 850 nanometer (nm) laser that can go up to 220 meters with 62.5-micron, 550 meters with 50-micron.

1000Base-LX (IEEE 802.3z) Single-mode fiber that uses a 9-micron core and 1300 nm laser that can go from 3 kilometers up to 10 kilometers.

1000Base-ZX (Cisco standard) 1000BaseZX, or 1000Base-ZX, is a Cisco specified standard for Gigabit Ethernet communication. 1000BaseZX operates on ordinary single-mode fiber-optic links with spans up to 43.5 miles (70 km).

10GBase-T (802.3.an) 10GBase-T is a standard proposed by the IEEE 802.3an committee to provide 10 Gbps connections over conventional UTP cables, (category 5e, 6, or 7 cables). 10GBase-T allows the conventional RJ45 used for Ethernet LANs and can support signal transmission at the full 100-meter distance specified for LAN wiring.

Ethernet Cabling

A discussion about Ethernet cabling is an important one, especially if you are planning on taking the Cisco exams. You need to really understand the following three types of cables:

- Straight-through cable
- Crossover cable
- Rolled cable

We will look at each in the following sections, but first, let's take a look at the most common Ethernet cable used today: the category 5 Enhanced Unshielded Twisted Pair (UTP), shown in Figure 1.14:

FIGURE 1.14 Category 5 Enhanced UTP cable

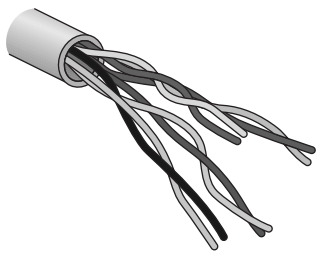

Category 5 Enhanced UTP cable can handle speeds up to a gigabit with a distance of up to 100 meters. Typically we'd use this cable for 100 Mbps and category 6 for a gigabit, but the category 5 Enhanced is rated for gigabit speeds and category 6 is rated for 10 Gbps!

Straight-Through Cable

The *straight-through cable* is used to connect the following devices:

- Host to switch or hub
- Router to switch or hub

Four wires are used in straight-through cable to connect Ethernet devices. It's relatively simple to create this type, and Figure 1.15 shows the four wires used in a straight-through Ethernet cable.

FIGURE 1.15 Straight-through Ethernet cable

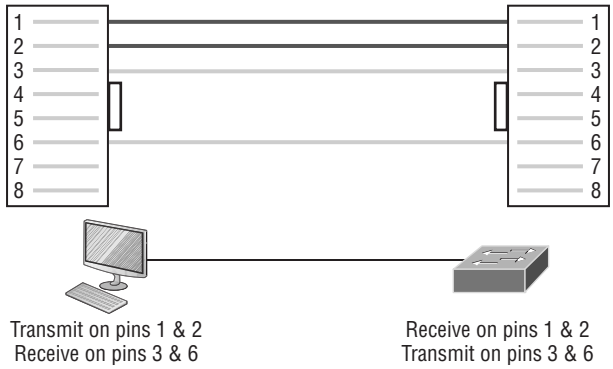

Notice that only pins 1, 2, 3, and 6 are used. Just connect 1 to 1, 2 to 2, 3 to 3, and 6 to 6 and you'll be up and networking stat. Just remember that this would be a 10/100 Mbps Ethernet-only cable and wouldn't work with gigabit, voice, or other LAN or WAN technology.

Crossover Cable

The *crossover cable* can be used to connect these devices:

- Switch to switch
- Hub to hub
- Host to host
- Hub to switch
- Router direct to host
- Router to router

The same four wires used in the straight-through cable are used in this cable; we just connect different pins together. Figure 1.16 shows how the four wires are used in a cross-over Ethernet cable.

FIGURE 1.16 Crossover Ethernet cable

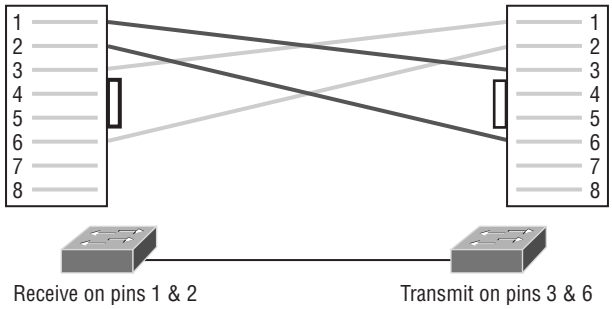

Notice here that instead of connecting 1 to 1, 2 to 2, and so on, here we connect pins 1 to 3 and 2 to 6 on each side of the cable. Figure 1.17 shows some typical uses of straight-through and crossover cables:

FIGURE 1.17 Typical uses for straight-through and cross-over Ethernet cables

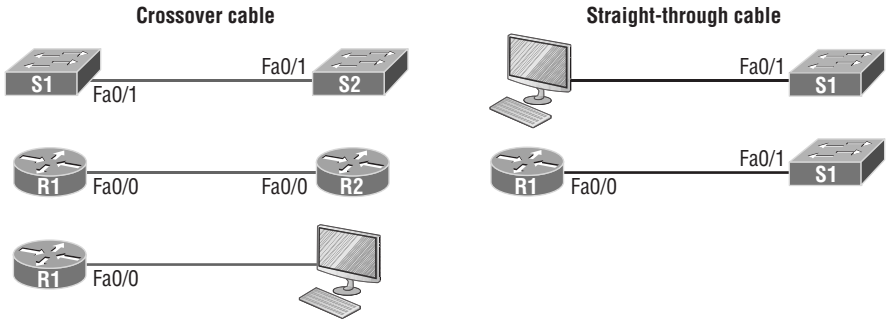

The crossover examples in Figure 1.17 are switch port to switch port, router Ethernet port to router Ethernet port, and router Ethernet port to PC Ethernet port. For the straight-through examples I used PC Ethernet to switch port and router Ethernet port to switch port.

It's very possible to connect a straight-through cable between two switches, and it will start working because of autodetect mechanisms called auto-mdix. But be advised that the CCNA objectives don't typically consider autodetect mechanisms valid between devices!

UTP Gigabit Wiring (1000Base-T)

In the previous examples of 10Base-T and 100Base-T UTP wiring, only two wire pairs were used. That's definitely not good enough for Gigabit UTP transmission!

The 1000Base-T UTP wiring pictured in Figure 1.18 requires four wire pairs and uses more advanced electronics so that each and every pair in the cable can transmit simultaneously. Even so, gigabit wiring is almost identical to my earlier 10/100 example, except that we'll use the other two pairs in the cable too.

FIGURE 1.18 UTP Gigabit crossover Ethernet cable

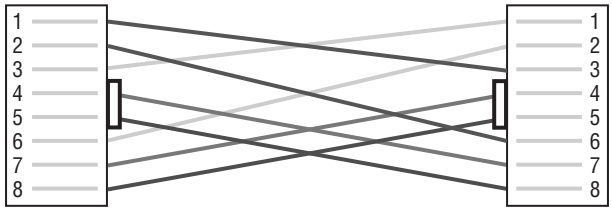

For a straight-through cable it's still 1 to 1, 2 to 2, and so on up to pin 8. And in creating the gigabit crossover cable, you'd still cross 1 to 3 and 2 to 6, but you would add 4 to 7 and 5 to 8—pretty straightforward!

Fiber Optic

Fiber-optic cabling has been around for a long time and provides some solid advantages. The cable allows for very fast transmission of data, is made of glass (even plastic), is very thin, and works as a waveguide to transmit light between two ends of the fiber. Fiber optics has been used to go very long distances, as in intercontinental connections, but it's becoming more and more popular in Ethernet LAN networks due to the fast speeds available. Also, unlike UTP, it's immune to interference like crosstalk.

The main components of this cable are the core and the cladding. The core will hold the light, and the cladding confines the light to the core. Remember, the tighter the cladding, the smaller the core; the smaller the core, the less light sent through it, but it can go faster and farther.

Figure 1.19 pictures a 9-micron core, which is very small and can be measured against a human hair—50 microns!

FIGURE 1.19 Typical fiber cable. Dimensions are in um (10^{-6} meters). Not to scale.

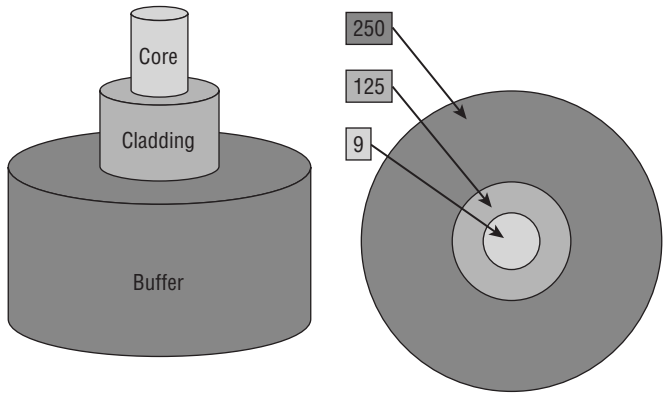

So here, the cladding is 125 microns, which is actually a fiber standard that allows manufacturers to make connectors for all fiber cables. The last piece of this cable is the buffer, which is there to protect the delicate glass.

There are two major types of fiber optics: single-mode and multimode. Figure 1.20 shows the differences between multimode and single-mode fibers.

FIGURE 1.20 Multimode and single-mode fibers

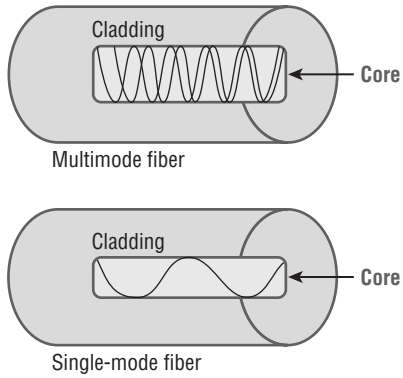

Single-mode is more expensive, has a tighter cladding, and can go much farther distances than multimode. The difference comes in the tightness of the cladding, which as I mentioned, makes a smaller core. It also means that only one mode of light will propagate down the fiber. Multimode is looser and has a larger core so it allows multiple light particles to travel down the glass. These particles have to be put back together at the receiving end. This means the distance is less than that with single-mode fiber, which allows only very few light particles to travel through the fiber.

There are about 70 different connectors for fiber, and Cisco uses a few of them.

Power over Ethernet (802.3af, 802.3at)

Power over Ethernet (PoE and PoE+) technology describes a system for transmitting electrical power, along with data, to remote devices over standard twisted-pair cable in an Ethernet network. This technology is useful for powering IP phones (Voice over IP, or VoIP), wireless LAN access points, network cameras, remote network switches, embedded computers, and other appliances. These are all situations where it would be inconvenient, expensive, and possibly not even feasible to supply power separately. A big reason for this is because the main wiring must be installed by qualified, licensed electricians in order to meet legal and/or insurance mandates.

The IEEE has created a standard for PoE called 802.3af. For PoE+ it's referred to as 802.3at. These standards describe precisely how a powered device is detected and also defines two methods of delivering Power over Ethernet to a given powered device. Keep in mind that PoE+ standard, 802.3at, delivers more power than 802.3af, which is compatible with Gigabit Ethernet with four-wire pairs at 30w.

This process happens one of two ways: either by receiving the power from an Ethernet port on a switch (or other capable device) or via a power injector. And you can't use both approaches to get the job done. And be careful here because doing this wrong can lead to serious trouble! Be sure before connecting.

Figure 1.18 gives you an example of a Cisco Next Generation Firewall (NGFW). It has eight ports that can be routed or switched ports, and ports 7 & 8 are listed as PoE ports at 0.6A. This is an excellent new NGFW that I personally use in my office.

FIGURE 1.21 NGFW ports provide PoE

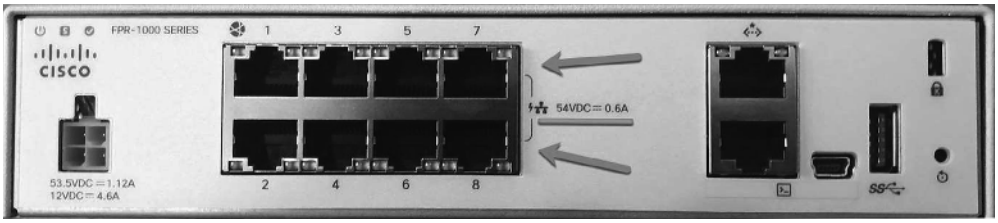

Remember, if you don't have a switch with PoE, you can use a power injector.

Be really careful when using an external power injector! Take your time and make darn sure the power injector provides the voltage level for which your device was manufactured.

Summary

I started this chapter by defining exactly what an internetwork, which was defined as: You create an internetwork when you connect two or more networks via a router and configure a logical network addressing scheme with protocols like IP or IPv6.

I then moved on to covering network components like routers, switches and defining what exactly creates a Small Office Home Office Network (SOHO).

After that, I touched on Next Generation Firewalls (NGFWs), and Cisco's design architecture of 3-tier and 2-tier, and spine-leaf was covered as well.

Finally, a thorough overview of Ethernet and the wiring used in Local Area Networks (LANs) and Wide Area Networks (WANs) was discussed.

Exam Essentials

Differentiate between a Switch and a Router Switches operate at layer 2 of the OSI model and only read frame hardware addresses in a frame to make a switching decision. Routers read to layer 3 and use routed (logical) address to make forwarding decision on a packet

Understand the term SOHO SOHO means small office, home office, and is small network connecting a user or small handful of users to the internet and office resources such as servers and printers. Usually just one router and a switch or two, plus a firewall.

Define 3-Tier Architecture The Cisco hierarchical model can help you design, implement, and maintain a scalable, reliable, cost-effective hierarchical internetwork. Cisco defines three layers of hierarchy, the core, distribution, and access, each with specific functions, and it's referred to as a 3-tier network architecture.

Define 2-Tier Acthitecture 2-tier architecture is also referred to as the collapsed core design because it's only 2-layers. But in concept, it's like the 3-tier only less expensive and geared for smaller companies. The design is meant to maximize performance and user availability to the network, while still allowing for design scalability over time. In a 2-tier, the distribution layer is merged with the core layer.

Define spine-leaf Also referred to as leaf-and-spine topology, in the leaf-and-spine design, there are switches found at the top of each rack that connect to the servers in the rack, with a server connecting into each switch for redundancy. People refer to this as a top-of-rack (ToR) design because the switches physically reside at the top of a rack.

Review Questions

You can find the answers to these questions in the Appendix.

1. Which one of the following is true about the Cisco core layer in the three-tier design?
   - **A.** Never do anything to slow down traffic. This includes making sure you don't use access lists, perform routing between virtual local area networks, or implement packet filtering.
   - **B.** It's best to support workgroup access here.
   - **C.** Expanding the core, e.g., adding routers as the internetwork grows, is highly recommended as a first step in expansion.
   - **D.** All cables from the Core must connect to the TOR.
2. Which one of the following best describes a SOHO network?
   - **A.** It uses ff:ff:ff:ff:ff:ff as a layer 2 unicast address, which makes it more efficient in a small network
   - **B.** It uses UDP as the Transport layer protocol exclusively, which saves bandwidth in a small network.
   - **C.** A single or small group of users connecting to a switch, with a router providing a connection to the internet
   - **D.** SOHO is the network cabling used from the access layer to the TOR
3. Which two of the following describe the access layer in the three-tier network design? (Choose two.)
   - **A.** Microsegmentation
   - **B.** Broadcast control
   - **C.** PoE
   - **D.** Connections to TOR
4. Which fiber type is a Cisco standard and has a distance of over 40 miles?
   - **A.** 1000Base-SX
   - **B.** 1000Base-LX
   - **C.** 1000Base-ZX
   - **D.** 10GBase-T
5. What is the speed of a T3?
   - **A.** 1.544Mbps
   - **B.** 2.0Mbps
   - **C.** 100Mpbs
   - **D.** 44.736Mbps

6. Which of the following is *not* provided by today's NGFWs?
   - **A.** IPS Inspection
   - **B.** Layer 2 deep packet inspection
   - **C.** Application Visibility and Control (AVC)
   - **D.** Network Address Translation (NAT)

7. Which of the following is the standard for PoE+?
   - **A.** 802.3P
   - **B.** 802.3af
   - **C.** 802.3at
   - **D.** 802.3v6

8. Which of the following defines a two-tier design?
   - **A.** The access layer connects to the distribution layer, and the 2-tiers then connect to the core layer.
   - **B.** In a two-tier design, the distribution layer is merged with the core layer.
   - **C.** It's best to support workgroup access in the two-tier layer
   - **D.** All cables from the core must connect to the two-tier TOR

9. What is the speed of the 802.3.an standard?
   - **A.** 100Mbps
   - **B.** 1Gbps
   - **C.** 10Gbps
   - **D.** 100Gbps

10. In a spine-leaf design, which is true?
    - **A.** The switches are found at the top of each rack that connect to the servers in the rack.
    - **B.** The distribution layer is merged with the core layer.
    - **C.** The access layer connects to the distribution layer, and the two-tiers then connect to the core layer.
    - **D.** All cables from the core must connect to the spine, which connects to the leaf device.