Chapter 1

Network Design and Concepts

THE CCIE QUALIFICATION EXAM TOPICS COVERED IN THIS CHAPTER INCLUDE THE FOLLOWING:

- ✓ General routing protocol concepts
- ✓ Know the OSI reference model
- ✓ Understanding hierarchical topologies
- ✓ Designing scalable networks
- ✓ Increasing fault tolerance
- ✓ Understand binary, decimal, and hexadecimal conversion

When designing networks, it is extremely important to start with a good, solid network topology. The thought you put into this initial step will determine how well your network design will perform in the future.

In this chapter, we'll discuss network topology designs that help you optimize network features. We'll teach you how to design a hierarchical topology using the Cisco three-layer model and show you how to build an internetwork that is scalable, manageable, and cost effective with improved performance. This chapter will also teach you how to build fault-tolerant internetworks and how to perform load balancing for both LANs and WANs. We will also review some of the fundamentals of networking.

Although some of the topics in this chapter may seem to be extremely basic to most readers, they are topics that will be covered in the CCIE qualification examination. Likewise, it is extremely important to have a solid foundation of the networking environment prior to discussing the higher-level technologies. For this reason, we most certainly must discuss these topics.

General Routing Concepts

First we'll discuss routing within the network environment. Keep in mind the difference between switching and routing. Switching occurs at the OSI Data-Link layer (layer 2), and routing occurs at the OSI Network layer (layer 3). Switches forward *frames* based on MAC addresses, and routers forward *packets* based on a logical layer 3 address, such as an IP address. With that said, we should note that the single topic of routing in itself is multifaceted. We'll discuss aspects such as static versus dynamic routing and briefly cover interior and exterior routing protocols. Finally, we will discuss distance vector versus link-state routing protocols.

Static vs. Dynamic Routing

Routing protocols are dynamic, meaning that they can make forwarding decisions based on changes to network topology, whereas static routes are manually configured on the routers and can only make a static forwarding decision. However, if a static route's next-hop goes down (drops out of the routing table), that static route will be removed from the routing table.

Static routing is efficient in a hub-and-spoke network, which has no redundant paths forming any type of a mesh or partial mesh topology. Dynamic routing protocols can determine the best routes to a destination network automatically. Dynamic routing protocols use metrics to make their routing decisions. Also, because dynamic routing automatically adds and deletes routes as the network topology changes, the network administration is greatly simplified.

Interior vs. Exterior Routing Protocols

Routing protocols can be divided into two separate categories: interior gateway protocols (IGPs) and exterior gateway protocols (EGPs). IGPs are used to perform routing within a single autonomous system (AS) or a single administrative network domain. EGPs are used to communicate, or route traffic, between separate autonomous systems (separate administrative network domains). Border Gateway Protocol (BGP) is the main EGP in use today, and most other routing protocols are IGPs (RTMP, RIP, IGRP, EIGRP, OSPF, IS-IS, NLSP).

Distance Vector vs. Link-State Routing Protocols

The early routing protocols were all distance vector, based on the Bellman-Ford routing algorithm. They advertise routes to destination networks based on the distance and direction required to reach the destination network. The most popular vector or metric is hop count. All routers that are configured with hopcount-based distance vector routing protocol must advertise their full routing tables to each of their neighbor routers at specified intervals (for example, every 30 seconds, 60 seconds, 90 seconds, etc.). This advertisement of the full routing table is accomplished by broadcasting the packets out every interface in the router, which results in a lot of broadcast in the network using up valuable wide area network (WAN) bandwidth.

The following are some examples of distance vector routing protocols:

- Routing Table Maintenance Protocol (RTMP)—AppleTalk
- Routing Information Protocol (RIPv1 and RIPv2)
- Interior Gateway Routing Protocol (IGRP)
- DEC DNA Phase IV

To correct some of the inefficiencies of the distance vector routing protocols, link-state routing protocols were developed. Link-state routing protocols are based on the shortest path first (SPF also known as Djikstra) algorithm. Routers running a link-state routing protocol compile information about themselves (their IP addresses, directly connected links, and the up or down status of those links). This compiled information is sent to every other router in the network, and then each router independently calculates the best path to each destination network and maintains a map or topology of the entire network. Routers will transmit a partial or incremental routing update only when and if a directly connected link to a router changes its up or down status, thus the issue with all the broadcast traffic in the network is avoided.

The following are some examples of link-state routing protocols:

- Open Shortest Path First (OSPF)
- Intermediate System-to-Intermediate System (IS-IS)
- NetWare Link Services Protocol (NLSP)
- DECnet Phase V

Now that you understand general routing protocol concepts, it is imperative that you understand the networking environment. The most important aspect of the networking environment is the OSI model, because it provides a reference by which all networking functions can be understood.

OSI Reference Model

The topic of the Open System Interconnection (OSI) model is covered in just about every networking course and book, so we will provide a quick review of the seven-layered model. The OSI model was developed by the International Organization for Standardization (ISO) in 1984 to describe the flow of data on a network. The seven layers (from bottom to top) are as follows: Physical, Data-Link, Network, Transport, Session, Presentation, and Application. You can remember the layers with the statement "Please Do Not Throw Sausage Pizza Away."

Each layer is self-contained, meaning that each layer can be implemented independently. If you run IP (layer 3), it can ride on top of various layer 2 protocols, such as Ethernet, Frame Relay, High-Level Data Link Control (HDLC), Point-to-Point Protocol (PPP), and so on. In the following sections, we'll review each layer.

Physical Layer (OSI Layer 1)

The Physical layer describes the transportation of the bits over the physical media (metallic cable, fiber, air waves). Keep in mind that the bits (1s and 0s) can be represented by either digital or analog signals. It also defines the signaling specifications, cable types, and interfaces as well as voltage levels, data rates, and maximum transmission distances. Repeaters and hubs operate in the Physical layer because they do not delineate broadcast or collision domains.

The following are some examples of the Physical layer standards:

- RJ-45
- EIA/TIA-232
- V.35

Data-Link Layer (OSI Layer 2)

The Data-Link layer provides reliable transport of the bits across the Physical layer. It formats the bits into frames for transmission. The Data-Link layer provides sequencing of frames, flow control, synchronization, and physical addressing. Bridges and layer 2 switches operate in the

Data-Link layer and make forwarding decisions based on a MAC address. The following are some examples of the Data-Link layer standards:

- Frame Relay
- Asynchronous transfer mode (ATM)
- High-level Data Link Control (HDLC)
- Ethernet
- Integrated Services Digital Network (ISDN)
- Point-to-Point Protocol (PPP)

And remember that, especially when discussing Ethernet, the Data-Link layer is divided into two sublayers for a local area network (LAN). The upper sublayer is the Logical Link Control (LLC) sublayer, which manages communications between devices. The lower sublayer is the Media Access Control (MAC) sublayer, which manages the protocol access to the physical media. Devices that operate in this layer can utilize a unique physical MAC address.

Network Layer (OSI Layer 3)

The Network layer is responsible for data forwarding and methods to determine the best forwarding path to a destination. At this layer, data is forwarded in units called packets. This layer specifies routing protocols, logical network addressing, and packet fragmentation. Routers and layer 3 switches operate at this layer and make forwarding decision based on a layer 3 address, such as an IP address. Routers and layer 3 switches both define collision and broadcast domains. As a CCIE, this is where you will really prove your worth!

The following are some examples of the Network layer standards:

- Internet Protocol (IP)
- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)
- Internetwork Packet Exchange (IPX)

Transport Layer (OSI Layer 4)

The Transport layer provides transport of the data from the upper layers. It can provide end-to-end error checking and recovery, multiplexing, virtual circuit management, and flow control. Messages are assigned a sequence number at the transmission end. At the receiving end the packets, or segments, which are also called Transport layer protocol data units (PDUs), are reassembled and checked for errors. And when TCP is used on the Transport layer, additional functionality such as flow control is included. Flow control, which is also known as windowing, manages the data flow to ensure that the transmitting device does not send more data than the receiving device can handle. The following are some examples of the Transport layer standards:

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

- Sequenced Packet Exchange (SPX)
- Real-Time Transport Protocol (RTP)

Real-Time Transport Protocol (RTP) is actually a special Transport layer protocol that rides on top of UDP, but it is considered a Transport layer protocol.

Session Layer (OSI Layer 5)

The Session layer provides a control structure for communication between various applications. It establishes, manages, and terminates communication connections called *sessions*. The management of these sessions involves the synchronization of dialog control by using checkpoints in the data stream.

The following are some examples of the Session layer standards:

- NetBIOS
- Session Control Protocol (SCP)
- Real-Time Control Protocol (RTCP)

Presentation Layer (OSI Layer 6)

The Presentation layer is responsible for code conversion functions. These functions ensure that data sent from one application on one device is readable by an application on another device. This is accomplished through converting the character representation formats, *data compression*, and encryption. Voice encoding schemes are also specified at this layer.

The following are some examples of the Presentation layer standards:

- ASCII
- Moving Picture Experts Group (MPEG)
- Graphics Interchange Format (GIF)
- Voice CODECs (e.g., G.711, G.729a, G.726, G.728)

Application Layer (OSI Layer 7)

The Application layer provides the user and operating system with access to the network services. This layer interacts with the software applications by identifying communication resources, determining network availability, and distributing information services. It also provides synchronization between the peer applications that reside on separate systems.

The following are some examples of the Application layer standards:

- Telnet
- File Transfer Protocol (FTP)

- Simple Mail Transfer Protocol (SMTP)
- Simple Network Management Protocol (SNMP)

Now that you understand the seven layers of the OSI model, you can conceptually rationalize the flow of data across a network environment. Next, we'll turn our attention to the many aspects of a good, reliable network design. An architect must design a solid foundation on which to build a skyscraper. A network must also be designed with a good solid foundation; if not, it will surely crumble in a matter of time.

Hierarchical Topologies

Hierarchy helps us to understand where things belong, how things fit together, and what functions go where. It brings order and understanding to otherwise complex models. If you want a pay raise, hierarchy dictates that you ask your boss, not your subordinate. The boss is the person whose role it is to grant (or deny) your request.

Hierarchy has many of the same benefits in network design that it does in other areas. When a hierarchical model is used to design a network, the network is more predictable. Using a hierarchical model, you can define the levels at which certain functions should be performed. For example, you would ask your boss, not your subordinate, for a raise because of their relative positions in the business hierarchy. Likewise, you can use tools such as access lists at certain levels in hierarchical networks and avoid them at others.

Let's face it: large networks can be extremely complicated, with multiple protocols, detailed configurations, and diverse technologies. Hierarchy helps you to summarize a complex collection of details into an understandable model. Then, as specific configurations are needed, the model dictates the appropriate manner in which to apply them.

Benefits of Hierarchical Topologies

Hierarchy can be applied to network topology in many ways, and Cisco has long encouraged using the hierarchical approach when designing the network topology. The benefits of hierarchy to network topology include improvements in the following areas:

- Scalability
- Manageability
- Performance
- Cost

We'll look at each of these in a bit more depth.

Scalability

Hierarchical networks, which are easier to scale than other models (such as a flat network model), are actually composed of many individual modules, each with a specific position within the hierarchy. Because their design is modular, expansion can often be as simple as adding new

modules into the overall internetwork. A flat network model does not lend itself to future physical growth or additional segmentation of business functionality.

Manageability

Hierarchical networks are easier to manage than other types of networks because they are easier to troubleshoot. For example, anyone familiar with Ethernet will know what great fun it is to troubleshoot 10Base2 (a coaxial network infamous for poor troubleshooting avenues). If the network is down, where do you begin (assuming you lack sophisticated diagnostic tools)? You'll need more cable when installing 10BaseT, but the cost is almost always justified because troubleshooting a star network is so much easier than troubleshooting a bus network. Hierarchical networks offer similar advantages in troubleshooting. It is much simpler to isolate problems within a hierarchy than in other models, such as meshed networks. In a flat network, you have to search for issues within a single, large network, and in a hierarchical network, the large network is segmented into separate partitions, making it easier to isolate the issue.

Performance

Improvements to network performance may well justify hierarchical network design. Networks that use hierarchical design can take advantage of advanced routing features such as *route summarization,* which results in smaller routing tables and faster convergence in large networks. Fully meshed networks require larger routing tables and converge slower because of the greater number of possible paths.

Cost

In the end, overall cost is often the driving force when building networks. Due to the properties we just discussed, hierarchical networks generally require fewer administrator hours to maintain and can make more efficient use of hardware and other resources. You can anticipate hardware needs more readily than in nonhierarchical networks, which will be explained more in the next section. In addition, you can more accurately purchase and share WAN bandwidth between layers of hierarchy.

As you can see, there are many benefits of a hierarchical approach to network design as it pertains to scalability, manageability, performance, and cost. Now we'll mix things up a little bit by discussing the Cisco three-layer hierarchical model.

The Three-Layer Hierarchical Model

Just when you finally memorized all the aspects of the OSI reference model, Cisco created its own hierarchical model that you now need to learn. This model is used to help you design a scalable, reliable, cost-effective hierarchical internetwork. Cisco defines three layers of hierarchy. The Cisco three-layer hierarchical model refers to a conceptual guideline to follow in your network design; it does not refer to the network data flow, as the OSI model does. The three layers are as follows:

- Core
- Distribution
- Access

Each layer has specific responsibilities. Remember, however, that the three layers are logical and not necessarily physical. Three layers do not necessarily mean three separate devices. In the OSI model, another logical hierarchy, the 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 you build physical implementations of hierarchical networks, you may have many devices in a single layer or you might have a single device performing functions at two layers. The definition of the layers is logical, not physical.

The *core layer* provides high-speed transport between locations. It provides low latency, high availability, and redundancy. No compression, access lists, or encryption should be performed at this layer because they will introduce latency.

The *distribution layer* provides route policies/filtering. It typically provides access lists, distribution lists, route summarization, security policies, compression, and encryption. Layer 3 switches can be implemented in this layer.

The *access layer* provides local and remote access to the network. This layer provides shared and switched layer 2 bandwidth, MAC filtering, and virtual local area network (VLAN) segmentation. Typically, remote access servers and virtual private network (VPN) concentrators are also implemented at this layer.

Figure 1.1 depicts a network in which the three-layer hierarchy model has been implemented.

FIGURE 1.1 A three-layered hierarchical network

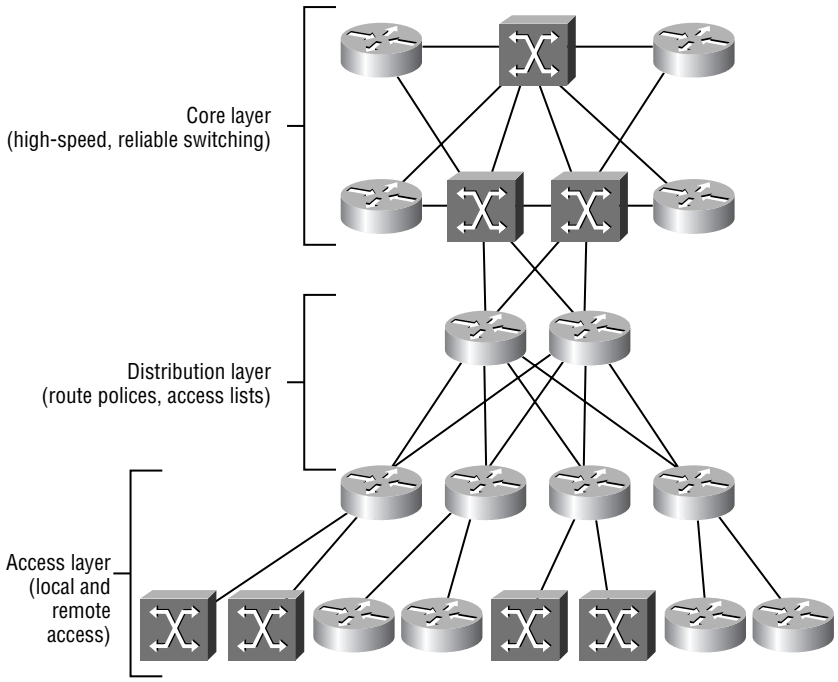

You should not add new routers below the access layer. To do so would expand the diameter of the network, which breaks the predictability of the topology. If you need to add new routers to support additional workgroups, they should communicate through the distribution layer and thus be peers (instead of subordinates) to the other access layer routers.

As already noted, having three separate levels does not have to imply the use of three separate routers. It could be fewer, or it could be more. Remember, this is a logical approach, and in some networks, two layers may be combined. Therefore, as in smaller networks, the core and distribution layers will probably be combined.

At this point, we have covered hierarchical topologies as they pertain to network design as well as good design characteristics to implement into your network design. But when you design a network, you also have a responsibility to take into account factors such as reliability and availability, efficiency, adaptability, and security. These are all criteria that will add value to your network design and ensure the longevity of the network.

Scalable Internetworks

Increasing demands for connectivity both in businesses and at home has led to extraordinary growth of today's internetworks, and because of this growth, it's important for internetworks to be scalable.

A scalable internetwork is continually growing, so it has to be flexible and it has to be easy to add components to it. An ideal design is based on the hierarchical model to simplify its management and to permit well-planned growth that honors the network's requirements.

Following are mandatory characteristics of a scalable internetwork:

- It's reliable and available.
- It's responsive.
- It's efficient.
- It's adaptable.
- It's easily accessible while being secure.

But that's enough of an introduction! In the following sections, we'll discuss these topics and help you understand the role each of these functions plays in the overall network design.

Reliability and Availability

When it comes to reliability, the internetwork's core layer is the most critical. The Cisco definition of *reliable* is "an internetwork that can respond quickly to changes in the network topology and accommodate failures by rerouting traffic."

The following list includes some Cisco internetwork operating system (IOS) features that serve to provide reliability and availability:

Reachability Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and NetWare Link Services Protocol (NLSP) use expanded metrics that can go beyond the hop count limitations of distance vector routing algorithms. These routing protocols (OSPF, EIGRP, and NLSP) analyze a combination of factors to establish the real cost of a path to a network, making Cisco routers capable of supporting very large internetworks.

Convergence Scalable routing protocols can converge quickly because of each router's complete understanding of the internetwork and ability to quickly detect problems.

Alternate paths routing Because OSPF and EIGRP build a complete map of the internetwork, a router can dynamically reroute traffic to an alternate path if a problem occurs.

Load balancing Through static routes and dynamic routing protocols (such as EIGRP, OSPF, RIP, etc.), the Cisco IOS is able to perform *load balancing*. This allows for redundant links and for more bandwidth to be available to locations needing more than just one link. For example, if two T1 WAN links were installed between buildings, the actual bandwidth between them would reach approximately 3 megabits per second (Mbps). In addition, this helps convergence time.

Tunneling Running a *tunneling protocol* affords the ability to communicate across WAN links previously unreachable. For example, if you have a WAN link that supports only TCP/IP and you want to manage a Novell NetWare server that supports only IPX across it, you could tunnel IPX packets inside IP packets and achieve your goal. However, remember that this causes overhead on the router.

Dial backup You can configure dial backup links for redundancy on your WAN links. This can also be configured as bandwidth on demand, providing extra bandwidth whenever a link becomes saturated, enhancing the link's reliability and availability as well as its performance.

Responsiveness

Because the network administrator is responsible for ensuring that users don't experience delays in responsiveness as the internetwork grows, they must be keenly aware of the latency factor that each piece of equipment (routers, switches, and bridges) contributes to the internetwork.

The Cisco IOS provides mitigation for the latency needs of each protocol running on your internetwork with various queuing techniques. We will discuss the various queuing and quality of service (QoS) techniques in Chapter 18.

Efficiency

The task of creating smoothly running, efficient LANs and internetworks is obviously important, but optimizing the bandwidth on a WAN can be difficult. The best way to reduce the bandwidth usage is to reduce the amount of update traffic on the LAN that will be sent over your WAN.

The following Cisco IOS features are available to help reduce bandwidth usage:

Access control lists (ACLs) These are used to permit or deny certain types of traffic from entering or exiting a specific router interface. ACLs can stop basic traffic, broadcasts, and protocol updates from saturating a particular link. TCP/IP, IPX, and AppleTalk can all be filtered extensively. ACLs can also be implemented to control access to devices in the network.

Snapshot routing Commonly used for ISDN connections when running distance vector protocols, *snapshot routing* allows routers to exchange full distance vector routing information at an interval defined by the administrator.

Compression over WANs The Cisco IOS supports TCP/IP header and data compression to reduce the amount of traffic crossing a WAN link. You can configure link compression so that header and data information are compressed into packets. This is accomplished by the Cisco IOS prior to sending the frame across the WAN. Cisco IOS also provides PPP compression and header compression for the RTP protocol.

Dial-on-demand routing (DDR) This allows wide area links to be used selectively. With it, the administrator can define "interesting" traffic on the router and initiate point-to-point WAN links based on that traffic. Interesting traffic is defined by access lists, so there's a great deal of flexibility afforded the administrator. For instance, an expensive ISDN connection to the Internet could be initiated to retrieve e-mail but not retrieve a web resource. DDR is an effective tool if WAN access is charged according to a quantified time interval, and it's best to use it in situations in which WAN access is infrequent.

Reduction in routing tables entries By using route summarization and incremental updates, you can reduce the number of router processing cycles by reducing the entries in a routing table. Route summarization, which summarize all the routes advertised into one entry, occurs at major network boundaries. Incremental updates save bandwidth by sending only topology changes instead of the entire routing table when transmitting updates.

Adaptability

Another important goal for an administrator is to design an internetwork that responds well to change. To achieve this, internetworks need to be able to do the following:

- Pass both routable and nonroutable network protocols. For example, TCP/IP is routable, and Microsoft's NetBIOS Enhanced User Interface (NetBEUI) is not routable, only bridgeable.
- Create islands of networks using different protocols. This allows you to add protocols used by the network islands to core layer routers or to use tunneling in the backbone to connect the islands, thus eliminating the necessity of adding unwanted protocols to the core backbone.
- Balance multiple protocols in a network. Each protocol has different requirements, and the internetwork must be able to accommodate the specific issues of each.

The Cisco IOS also has many features that contribute to network adaptability:

EIGRP Cisco's proprietary EIGRP allows you to use multiple protocols within one routing algorithm. EIGRP supports IP, IPX, and AppleTalk. It also allows unequal cost load balancing, which we will discuss in Chapter 9, "IP Interior Gateway Protocols."

Redistribution This allows you to exchange routing information between networks that use different routing protocols. For example, you can update a routing table from a network running IGRP on a router participating in a RIP network.

Bridging By using source-route bridging and integrated routing and bridging, you can integrate your older networks and protocols that do not support routing into the new internetwork.

Accessible but Secure

Access layer routers must be both accessed and used to connect to a variety of WAN services while maintaining security to keep hackers out.

The following Cisco IOS features support these requirements:

Dedicated and switched WAN support You can create a direct connection with Cisco routers using basic or digital services (a T1, for example). Cisco routers also support many different switched services such as Frame Relay, Switched Multimegabit Data Services (SMDS), X.25, and ATM to give you options to meet cost, location, and traffic requirements.

Exterior Gateway Protocol support Both exterior gateway protocol (EGP) and Border Gateway Protocol (BGP) are supported by the Cisco IOS. BGP, discussed in detail later in this book, is used mostly by Internet Service Providers (ISPs) and has mostly replaced EGP.

Access control lists These are used to prevent specific kinds of traffic from either entering or leaving a Cisco router.

Authentication protocols Cisco supports both Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) for providing authentication on WAN connections using Point-to-Point Protocol (PPP).

Hopefully, you now have a better understanding of what is required for a good solid foundation in your network design, as well as some of the additional factors—such as reliability and availability, efficiency, adaptability, and security—that must also be taken into account in a good network design. Next we'll discuss fault tolerance, or assuring that a host is able to find a path to the internetwork.

Fault-Tolerant Topologies

Some networks are more important than others. Of course, *all* networks are important, right? In some situations, however, network availability (or the lack thereof) can be much more costly.

When you're designing networks, the use of many features can significantly increase fault tolerance and decrease the possibility of network outages. Perhaps you have worked with servers that included disk mirroring or Redundant Array of Inexpensive Disks (RAID) or even redundant servers. These all use one form of protection. In this section, we will discuss techniques that ensure that, first, hosts can find a path to the internetwork, and second, once they find a path, the path actually works! Remember to always keep the three-layer hierarchical model in mind and that to properly implement fault tolerance, you must consider fault tolerance at each layer of the network.

Redundant LAN Configurations

It does little good to install routers at the access level if the workstations cannot find and use them. This leads us to the issue of investigating how different workstations find routers that lead to the internetwork and how you can help those workstations find redundant paths out of the LAN. We'll consider this problem for the following protocols:

- Transmission Control Protocol/Internet Protocol (TCP/IP)
- Internetwork Packet Exchange (IPX)
- AppleTalk

Transmission Control Protocol/Internet Protocol (TCP/IP)

Most network administrators have configured a default router (or default gateway) on a host when setting up TCP/IP. This, along with configuring the IP address, the subnet mask, and perhaps the DNS server, is standard when setting up any TCP/IP device.

In Figure 1.2, Workstation A and Workstation B are assigned the default gateway of 172.16.10.1. Server A, which has two NICs and is acting as a router between the two NICs, also gets a default gateway of 172.16.10.1. Workstation C, however, must be assigned a default gateway of 172.16.20.1, which is the first (and only, in this diagram) router that it sees. It cannot contact the router at 172.16.10.1 directly because they are on separate data-link networks! Once the administrator has this all mapped out, they can configure all the devices with TCP/IP information.

Some implementations of TCP/IP will allow for multiple default gateways; others provide for the workstation to listen to routing updates to learn of routers. Either method will provide the client with redundant paths out if the primary router should fail and should be considered when available. Unfortunately, the most common method of default router configuration is to statically assign the default router at the client. This means that should the router fail, there are two options: either fix the router or reconfigure the workstation. Hardly fault tolerant! We will look at two Cisco solutions to this problem: Hot Standby Router Protocol (HSRP) and proxy Address Resolution Protocol (ARP) support.

FIGURE 1.2 A sample internetwork

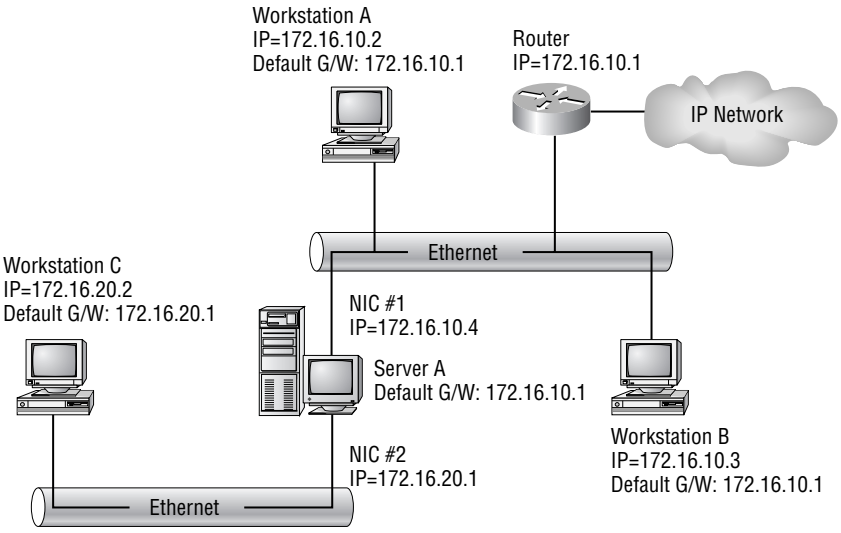

Default Gateway and HSRP

Hot Standby Router Protocol (HSRP) can allow IP devices to keep working through their default router even when that router fails. It does this by creating what Cisco calls a virtual router, or *phantom router,* on the network. This phantom router does not exist physically, but it does have a Media Access Control (MAC) address and an IP address. Workstations are configured to use the phantom router's IP address as a default gateway. The phantom addresses are actually passed among the physical routers participating in HSRP. If the physical router hosting the phantom router's MAC and IP addresses fails, another physical router automatically answers to the phantom's MAC and IP addresses and accepts the traffic. The workstations need never be aware that the hardware they are talking to has changed, and the MAC and IP addresses they have been using continue to function as if nothing had ever happened.

We will discuss HSRP and the various ways of configuring this functionality in Chapter 22, "IP Services."

Default Gateway and Proxy ARP

You can configure some IP stacks to take advantage of proxy ARP. Under normal circumstances, workstations will use the Address Resolution Protocol (ARP) to find the hardware addresses that are on their local network. When using proxy ARP, however, these workstations will send out ARP requests for *every* IP device that they want to communicate with, regardless of whether or not it is on their local network. Any router that is hearing this request and that is able to reach the desired IP address can respond to the ARP with its own MAC address. From the workstation's view, it looks like the whole world is one big LAN. The routers take care of the details of reaching remote segments. Proxy ARP is now enabled by default in all Cisco routers.

The end result is that workstations can dynamically locate redundant paths out of the LAN. When the proxy ARP request (which is a broadcast) is sent out, a response can come from any router able to reach the required destination, and thus if one router fails, the workstation can immediately begin to communicate with the internetwork through any other available routers. Understand, however, that overhead will result on any router performing proxy ARP. Also, be aware that proxy ARP is a security concern because any device can respond to a proxy ARP, which could enable that traffic to be intercepted by unauthorized devices.

To configure workstations to run proxy ARP, simply set the default gateway of the workstations to their own IP address. Once you have reconfigured your default gateway to the IP address of the workstation, try pinging a remote device. Turn on `debug ip packet` on the router and see what happens.

Internetwork Packet Exchange (IPX)

Internetwork Packet Exchange (IPX) is dynamic in the assignment of the address and default router. IPX clients can issue a "find network number" request, and any router that can provide access to the requested network answers it. If that particular router goes away, the client will automatically reissue the request. If there is a different path out, the new router will answer the client request and the client can then take advantage of the new path. Once again, it's completely dynamic.

What this means is that at the access layer, anytime you provide two paths out, AppleTalk and IPX clients will automatically find them and use them, and that increases the fault tolerance of your network. As we mentioned, if the clients cannot find paths out, the internetwork is not much use to them. IP clients are typically more challenging because they generally are not as dynamic at finding paths out as IPX or AppleTalk clients.

AppleTalk

Have you ever wondered why you don't have to play these little gateway games with AppleTalk? The reason is that both addressing and default router configuration are dynamic with this protocol. With AppleTalk, the workstations actually listen to the Routing Table Maintenance Protocol (RTMP) routing updates. They don't build routing tables as routers do, but they do pay attention to the source AppleTalk address of the update. They then use that address as their default gateway. RTMP updates are broadcast every 10 seconds, which means that if you lose your default router on a network, workstations will take a maximum of 10 seconds to learn any redundant router address.

Redundant WAN Connections

As you have just seen, you can provide redundancy in the links between clients and servers on the LAN using several techniques. Now we will look at ways to provide redundancy inside the WAN.

Consider the network illustrated in Figure 1.3. This is a full mesh network, in which every node has a direct link to every other node. For fault tolerance, this is great, but it is far from efficient and does not scale well. Also, it has departed from the hierarchical topology we looked at earlier. There is a solution, however, that will preserve hierarchy while providing redundancy

in the WAN. A full mesh is commonly used within a layer (especially the core layer); do not implement a full mesh between layers.

FIGURE 1.3 A full mesh network

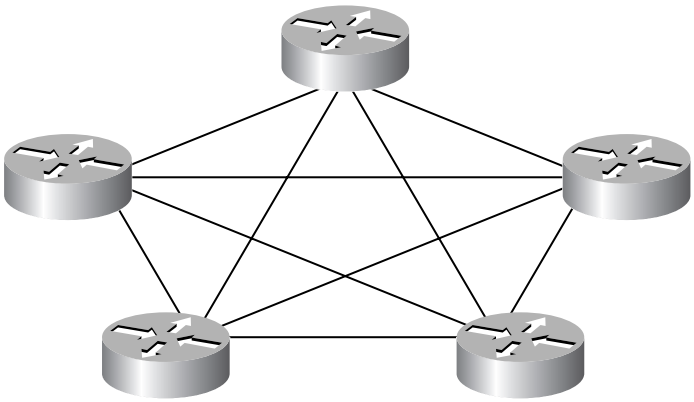

Partial Mesh Topology

We have implemented a partial mesh in the network shown in Figure 1.4. Notice that we have preserved our hierarchy, yet each node has a redundant link to the layer above it. This design provides all the advantages of hierarchical design, is scalable, and can take advantage of load balancing.

FIGURE 1.4 Redundant hierarchical network

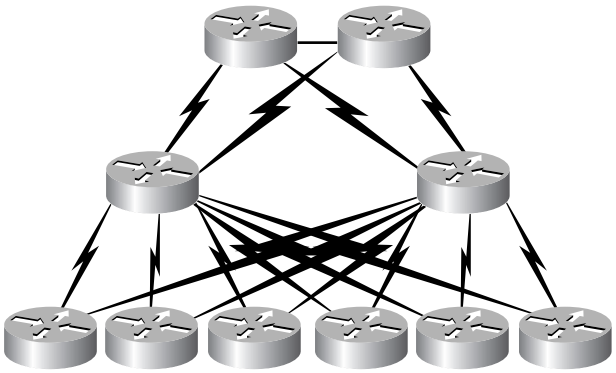

You can add the additional WAN connections in several ways. You could add them in identical pairs—that is, you could install two T1 lines rather than one. This provides the ultimate in redundancy. If one T1 fails, another is waiting to go. From a cost perspective, however, this

can be similar to buying two new cars just in case one gets a flat tire. True, you will probably never have to walk to work, but that security will certainly cost you.

An alternative to identical connections to the next layer is using links that are not the same, that is, perhaps a T1 and a 56Kbps backup line. Should the primary line fail, internetwork connectivity can be preserved, although generally at a reduced level. Once again, cost will most likely determine the capacity of the backup line.

Cisco has a solution that is similar to this second example, in which the two connections are not the same. In this case, the second, or backup, line is not even running until the primary line fails! We will look at this solution next.

Dial-on-Demand Routing (DDR) Backup

Not all redundant links have to be dedicated lines. In many cases, an ISDN Basic Rate Interface (BRI) is used to back up a dedicated leased line. This can be a great advantage because you will probably not want to bring the ISDN up unless the primary line fails (or becomes overloaded). Cisco's DDR allows this configuration. The ISDN line can be configured to become active only when the primary line either fails or is under heavy load. Of course, should the primary line fail and you have to depend on your backup, you will likely not have your normal bandwidth available. You will, however, likely be paying significantly less than you would to have a pair of dedicated lines.

We will discuss the various configurations and functionality of DDR in Chapter 15, "Integrated Services Digital Network (ISDN) and Dial Backup."

As you just saw, you can provide redundancy inside the WAN environment by using both redundant links or circuits and DDR. Both of these functions come with a financial cost to you, so you will most likely want to assure the best possible return on investment (ROI). To do so, you will want to use all the bandwidth that you are paying for, which brings us to the topic of load balancing.

Performance: Load Balancing

Redundant links are not cheap to operate, but they are called for in some situations. If you are going to pay for redundant links, you would likely want to use both lines when they are both available, and that involves load balancing.

A good design rule is to keep bandwidth consistent within a layer of hierarchy whenever possible and to use technologies such as DDR when it's not possible to purchase equal links.

In the following sections, we'll discuss various methods of utilizing load balancing in the WAN environment with three different protocols: Internet Protocol (IP), Internetwork Packet Exchange (IPX), and AppleTalk.

Load Balancing with Internet Protocol (IP)

With most IP routing protocols, load balancing is automatic. Dynamic routing protocols are supposed to find the redundant paths, and dynamic IP routing protocols will use both available paths. This is not, however, always be a good thing.

Some routing protocols (for example, those that use hop count) could see these two paths and load-balance across them just fine until the 56Kbps line is full. At that point, the traffic is equally balanced. These protocols, however, are not smart enough to realize that more than 90 percent of the total bandwidth is going unused on the T1! Once any link is operating at capacity, these routing protocols are not capable of sending additional traffic across links still not at capacity because they do not understand capacity as a metric. This problem is called pinhole congestion, and you can avoid it by using advanced routing protocols such as EIGRP.

Load Balancing with Internetwork Packet Exchange (IPX)

By default, IPX will not load-balance across multiple links; however, Cisco provides a way to enable IPX load balancing. You can use the `ipx maximum-paths` command, which specifies a number of links to load-balance across. We will discuss this command when we cover the IPX protocol later on in this book.

Load Balancing with AppleTalk

AppleTalk, like IPX, considers only one path to a remote network. You can set the number of parallel routing paths that can be used by AppleTalk by using the `appletalk maximum-paths` command. Remember to set this on all your routers, not just one router.

We've been discussing various methods of utilizing load balancing in the WAN environment. Fault tolerance is also needed in the multiaccess or LAN environment.

Fault Tolerance on Multiaccess Segments

LAN segments are very reliable when compared with their wide area counterparts. However, failure occurs on LANs as well, making fault tolerance an important issue. The most well-known fault tolerance mechanisms on a LAN are the dual rings encountered in Fiber Distributed Data Interface (FDDI) networks, where upon primary ring failure, a secondary backup fiber link will automatically take over, and the Spanning-Tree Protocol (STP) encountered in Ethernet networks, where upon a path failure, a redundant path can automatically take over.

In this section, you will also learn about fault tolerance methods that occur at the Data-Link and Network layers of the OSI reference model. We'll look at the following two types of fault tolerance:

- Internet Control Message Protocol (ICMP) redirects
- Proxy Address Resolution Protocol (proxy ARP)

Internet Control Message Protocol (ICMP) Redirects

The Internet Control Message Protocol (ICMP) is used primarily for error handling and testing; however, one of the processes it uses can provide fault tolerance. Here's how this type of fault tolerance works: When a packet is received by a router on a particular interface and the destination for the packet is outside that interface, an ICMP redirect tells the sender of that packet

to send the packet to a different router. The sender should then update its local routing table so that further packets are sent directly to the correct address.

Proxy ARP

Proxy Address Resolution Protocol (proxy ARP) is a variation on the ARP protocol in which an intermediate device, such as a router, sends an ARP response on behalf of an end node to the requesting host. Proxy ARP has been defined and referenced in many Requests for Comments (RFCs). This technology once had a strong following, and one of the benefits is that it can reduce bandwidth usage on slow-speed WAN links. As networks grew, however, proxy ARP did not scale with them.

Enabling Proxy ARP on Cisco Routers

By default, proxy ARP is enabled on Cisco routers, as displayed in the results of the following `show ip interface` command:

```
RouterA#show ip interface ethernet 0
Ethernet0 is up, line protocol is up
  Internet address is 10.1.0.1/16
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.9
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Null turbo vector
  IP multicast fast switching is disabled
  IP multicast distributed fast switching is disabled
  Router Discovery is disabled
  IP output packet accounting is disabled
```

```
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  Web Cache Redirect is disabled
  BGP Policy Mapping is disabled
RouterA#
```

Disabling Proxy ARP on Cisco Routers

To disable Proxy ARP on a Cisco router, use the `no ip proxy-arp` command, as follows:

```
RouterA(config)#interface ethernet 0
RouterA(config-if)#no ip proxy-arp
RouterA(config-if)#^Z
RouterA#
```

Most Cisco routers default to a four-hour ARP time-out.

The Advantages and Disadvantages of Proxy ARP

Proxy ARP has the following advantages:

- Simple configuration, no need to configure clients with gateway
- Load balancing, although it's somewhat random
- Immediate fault tolerance, if addresses have not been recently contacted

Using proxy ARP involves the following disadvantages:

- Creation of a lot of broadcast traffic
- Waiting for ARP cache on the workstation to time out in event of failure
- Lack of control over which router is primary and which is secondary

Proxy ARP does provide some fault tolerance on a multiaccess segment, but it does not give the level of control that most administrators want. A more robust and flexible method is needed. In response to this need, Cisco developed the Hot Standby Router Protocol (HSRP). We will discuss HSRP in Chapter 22.

Now let's shift gears and discuss the basic aspect of networking, We're of course referring to the 1s and 0s and how they represent data information.

Binary, Decimal, Hexadecimal Conversion

As a CCIE candidate, you must know how to perform conversions between decimal, binary, and hexadecimal numbers. This is the aspect of networking that is most important for you to thoroughly understand. The digital world comprises bits of information represented by a series of 1s and 0s, and if you are going to exist in this digital world, you must understand how the 1s and 0s are able to represent bits of information. A numbering system that uses 1s and 0s is known as a binary system, meaning that two states exist: a 1 or a 0. Now to make it easier to understand 1s and 0s, we represent them as either decimal or hexadecimal values. Because of this, it is imperative that you know how to convert binary to decimal and then to hexadecimal.

The decimal numeric system is based on 10 digits, 0 through 9. After 9, you use 2 digits starting at 10 and then cycle the rightmost digit from 0 to 9 again. At this point, the left digit is incremented to the next higher value; thus, after 19 comes 20.

The binary number system uses two digits, a 0 and a 1. Binary numbers are the language of computer systems. IP addresses and MAC addresses are represented by binary numbers. Each digit (0 or 1) is a *bit,* thus the binary number 10110110 is a binary number with 8 bits. An IP (version 4) address comprises 32 bits, and a MAC address comprises 48 bits. IP addresses are usually displayed in a dotted decimal format, meaning that each 8-bit number is represented as a decimal number. Likewise, a MAC address is displayed in hexadecimal format, meaning that it takes two hexadecimal characters to represent the original 8-bit number.

Binary to Decimal Conversion

To convert a binary number to decimal, multiply each digit by a power of 2. Each digit, or bit, has a decimal equivalent of from 1 to 128 (reading right to left), which is based on the location of the bit in the binary number. This is similar to converting decimal numbers based on 1s, 10s, 100s, and so on. In decimal format, the number 111 is 100+10+1. In binary format, the number 11111111 is the sum of 128+64+32+16+8+4+2+1, or 255. For the binary 8 bits 10110110, the result of the conversion is 128+0+32+16+0+4+2+0=182.

Figure 1.5 an *X* is used to represent the 8 bits within the binary word. Notice that, starting from the right to left, each digit is multiplied by 2 (1, 2, 4, 8, 16, 32, 64, and 128). The number is the value that the digit position holds. A 1 means that you use that digit position value, and a 0 means that you do not use that digit position value. So, if you add up the digit values with a 1, you will get 128+32+16+4+2=182.

FIGURE 1.5 Binary to decimal conversion

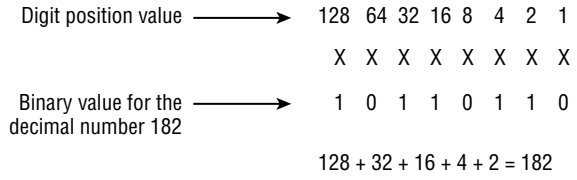

To convert decimal to binary, you simply need to figure out which digit position values are needed, so that when the values are added up, they equal the required decimal number.

Memorize 1, 2, 4, 8, 16, 32, 64, and 128. Use this as you read a binary number from right to left. It should be helpful in converting faster.

Binary to Hexadecimal Conversion

To convert a binary number to its hexadecimal equivalent, you must first divide the 8 bits of the binary word into two *nibbles* of 4 bits each. This is because with hexadecimal representation, there is only 16 possible values for the 4 bits to represent (0 through 15), thus each of the four digit positions represents a value the same as it does in the binary world. In hexadecimal format, the number 1111 is the sum of 8+4+2+1, or 15. For the binary 4 bits 0110, the result of the conversion is 0+4+2+0=6. So, the binary 8 bits of 01100111 become two nibbles of 0110 and 0111. The first nibble, 0110, is equal to 6, and the second nibble, 0111, is equal to 7 (0+4+2+1=7), thus the hexadecimal equivalent of the 8 binary bits is 67.

This all works well until the total value reaches from 10 through 15 because you only have one digit to represent the nibble (4-bit value). When this happens, you can perform the conversion by switching to letters after the number 9, thus the number 10 equals *A,* and the number 11 equals *B,* and so on until the highest possible number, 15, which equals *F.*

In Figure 1.6, we show the decimal numbers 0 through 15 with their binary and hexadecimal equivalents.

FIGURE 1.6 Decimal to binary to hexadecimal conversion

Decimal	Binary	Hex
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

A decimal number of **182** equals a binary number of **10110110**. This binary number can be broken into two nibbles: 1011 and 0110.

Here's how to convert the binary into hexadecimal:
1011=11=**B**
0110=6=**6**

Therefore, the decimal number 182 equals the hexadecimal value **B6**.

The hexadecimal format is used with MAC addressing, IPX addressing, and token ring addressing. It is common to represent a hexadecimal number with 0x before the number so that it is not confused with a decimal number. The hexadecimal number 16 is written as 0x10, not 10. It is also common to use the term *hex* when referring to hexadecimal.

We will discuss how to convert binary to decimal and to hexadecimal further in Chapter 8, "IP Addressing, and Subnetting," when we actually discuss the IP version 4 (IPv4) as well as the new IP version 6 (IPv6) standard.

Summary

A good network topology design makes it much easier to design and deploy network applications that run over the top of the network. Cisco recommends using a hierarchical design, which offers many benefits, including predictability, scalability, efficiency, cost control, and security.

Further, Cisco recommends that medium to large businesses use a three-layered hierarchical model consisting of core, distribution, and access layers. Each layer has clearly defined functions, and once the network is established, it can scale significantly before it needs to be reengineered.

Topologies that enhance network fault tolerance are often required. IPX and AppleTalk dynamically find their gateways to the internetwork, but IP features such as HSRP and proxy ARP can improve fault tolerance in workstation-to-router communication. Redundant WAN links can provide additional fault tolerance and can be used inside hierarchical designs. Technologies such as DDR provide for backup links. When redundant links are used, design consideration should be given to load balancing. You should also identify and avoid issues such as pinhole congestion.

In this chapter, we discussed decimal to binary to hexadecimal conversion. We also reviewed the basic features of the Cisco internetwork operating system (IOS). It is important that you have a firm understanding of the basics offered in this chapter before you move on to the other chapters in this book.

Exam Essentials

Understand the difference between static routing and dynamic routing. Static routing is efficient in a hub-and-spoke network that has no redundant paths forming any type of a mesh or partial mesh topology. However, static routes must be manually configured. Dynamic routing protocols can determine the best routes to a destination network automatically. Dynamic routing protocols use metrics to make their routing decisions.

Know the seven layers of the OSI model. The seven layers (from bottom to top) are as follows: Physical, Data-Link, Network, Transport, Session, Presentation, and Application. You can remember the layers with the statement "Please Do Not Throw Sausage Pizza Away."

Know the benefits of the hierarchical network model. The benefits of hierarchy to network topology include improvements to scalability, manageability, performance and cost.

Know the three layers of the Cisco three-layer hierarchical model Cisco defines three layers of hierarchy: core, distribution, and access.

Understand binary to decimal conversion. To convert a binary number to decimal, multiply each digit by a power of 2. Each digit, or bit, has a decimal equivalent of from 1 to 128 (reading right to left), which is based on the location of the bit of the binary number. This is similar to decimal numbers based on 1s, 10s, 100s, and so on. In decimal format, the number 111 is 100+10+1. In binary format, the number 11111111 is the sum of 128+64+32+16+8+4+2+1, or 255.

Key Terms

Before you take the exam, be certain you are familiar with the following terms:

access layer
bit
core layer
data compression
distribution layer
frames
hierarchy
load balancing
nibbles
packets
phantom router
route summarization
sessions
snapshot routing
tunneling protocol

Review Questions

1. Which of the following are advantages of hierarchical design?
   - **A.** Fault tolerance
   - **B.** Scalability
   - **C.** Ease of management
   - **D.** Predictability
   - **E.** All the above
2. Which of the following are layers in Cisco's three-layer hierarchical design? (Choose all that apply.)
   - **A.** Backbone
   - **B.** Core
   - **C.** End node
   - **D.** Access
   - **E.** Distribution
3. Which of the following should be included at the core layer? (Choose all that apply.)
   - **A.** Packet filtering
   - **B.** Firewalling
   - **C.** Fast throughput
   - **D.** Fault tolerance
   - **E.** Additional devices
4. How many layers of hierarchy should you add below the access layer?
   - **A.** None
   - **B.** One
   - **C.** Two
   - **D.** Three
   - **E.** Four
5. Which of the following are recommended at the distribution layer? (Choose all that apply.)
   - **A.** Packet filtering
   - **B.** Access lists
   - **C.** Queuing
   - **D.** Redundant WAN connections
   - **E.** Firewalls

6. Which of the following protocols allow for dynamic location of default routers? (Choose all that apply.)
   - **A.** IP
   - **B.** IPX
   - **C.** AppleTalk
   - **D.** NetBEUI

7. Which of the following methods will allow IP workstations to locate routers dynamically? (Choose all that apply.)
   - **A.** HSRP
   - **B.** Workstation listening to routing protocols
   - **C.** Proxy ARP
   - **D.** RTMP

8. You need to add a new site to your hierarchical network. Which of the following are possible places to connect the new site into your existing network? (Choose all that apply.)
   - **A.** Access layer
   - **B.** Distribution layer
   - **C.** Core layer
   - **D.** Backbone

9. What is the hexadecimal equivalent of decimal 182?
   - **A.** 67
   - **B.** 92
   - **C.** B6
   - **D.** 1A

10. You have a T1 link from an access layer router to a distribution layer router, and you have a BRI DDR connection to another distribution layer router. The DDR is configured to run in case of failure. Which of the following do you have?
    - **A.** Proxy ARP
    - **B.** Fault tolerance
    - **C.** Load balancing
    - **D.** HSRP
    - **E.** None of the above

Answers to Review Questions

1. E. Hierarchical design provides fault tolerance, scalability, manageability, and predictability.
2. B, D, E. The three layers of the Cisco hierarchical model are core, distribution, and access.
3. C, D. The core layer should not do anything to hinder packet flow through the network. It should provide fast throughput and fault tolerance.
4. A. No network should be installed below the access layer.
5. A, B, C, D, E. Access lists, queuing, redundancy, and firewalls should be used in the distribution layer.
6. A, B, C. IP, IPX and Appletalk protocols all allow clients to locate gateways dynamically. NetBEUI is not a routable protocol.
7. A, B, C. HSRP, dynamic routing protocols, and proxy ARP can be used to allow IP hosts to find and use alternate default gateways.
8. B, C. If the new site might be a future hub, it can be connected directly to the core. It should never be connected through the access layer.
9. C. Decimal 182 equals the binary equivalent of 10110110, which becomes the two nibbles 1011 and 0110. So, 1011=B, and 0110=6.
10. B. ISDN BRI can provide backup in case of failure.