Nothing in this world can take the place of persistence. Talent will not; nothing is more common than unsuccessful people with talent. Genius will not; unrewarded genius is almost a proverb. Education will not; the world is full of educated derelicts. Persistence and determination alone are omnipotent. The slogan "Press On" has solved and always will solve the problems of the human race.

—Calvin Coolidge

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

- List internetwork design goals
- Identify key requirements of internetwork design
- Describe a methodology for internetwork design

CHAPTER 1

Internetwork Design Overview

If you have ever been tasked with a network design project, whether large or small, undoubtedly you have had to make several difficult design decisions. You may have even found several "right answers" to the problem, but then needed to select the best "right answer." What, then, is the best "right answer?" It depends on how well you and your team know your customer and their requirements. The truth is, you can approach a network design in several ways. However, most network designs (the successful ones, at least) follow some fundamental guidelines.

This chapter focuses on the goals of internetwork design, including the technical and business trade-offs you must understand prior to making design choices. Sometimes, you and your team may be diverted from the topic into business issues that may not necessarily be technically oriented. As a network designer, you must remember to keep the team focused and gather the relevant information to make the design meet the customer's goals. This chapter provides you with an internetwork design methodology road map for use when approaching an internetwork design to keep your project on track.

Internetwork Design Goals

An internetwork can be generally defined as two or more local-area networks (LANs) interconnected by one or more Layer 3 devices (ordinarily routers). An internetwork may be contained within a single building, or may span the globe. Although there are specific differences between LANs and wide-area networks (WANs), in general terminology (and in this book) the word *network* refers to either type.

The first step in designing an internetwork is to establish and document the goals of the design. These goals will be particular to each organization or situation. Certain requirements tend to always show up in any good network design, however. They are as follows:

- Functionality
- Scalability
- Adaptability
- Manageability
- Cost effectiveness

Functionality

First and foremost, you cannot design a network without first fully knowing what you are trying to accomplish. Gathering all the requirements is often a very difficult task, but when the network is deployed, it must work as designed. There is absolutely no room for negotiation here. The network must enable users to meet their individual job requirements in such a way that the overall business requirements of the organization are met. The network must provide end-to-end application availability at some specified level of service (defined by management as the optimal compromise between functionality and cost). An MCI executive once told me that, before you embark on a new challenge, you must know the answer to the question, “How do you measure your success?” I urge you to ask yourself this same question of your network design on a regular basis, because it leads to an overall understanding of the tasks you are trying to achieve and keeps you focused on success.

Scalability

The network must be able to grow as the organization grows, and as more of the organization is included in the network. The initial design must be scalable across several orders of magnitude of network growth.

Take company XYZ, for example. They have been acquiring new companies at a rate of three to five per year. As a measure of your success, the network infrastructure must scale in a modular fashion so that new acquisitions “snap in” to their existing infrastructure. If scalability is not present in company XYZ’s network, the network absorption of this new company will undoubtedly fail or will become a management nightmare. I have seen many networks in the past which, due to poor planning for scalability, had become a jumbled “spaghetti net.” It is vital, therefore, to ensure that a design will scale in the future, even if large-scale expansion is not necessarily required today.

Adaptability

The network should be designed with an eye toward future technologies and should not include any design elements that would limit adoption of new technologies as they become available. There may be trade-offs between this and cost effectiveness throughout a network design/implementation. For example, Voice over IP (VoIP) and multicast are new technologies rapidly being adopted in many internetworks. Network designs should certainly be able to support these technologies without requiring a “forklift upgrade.” This is done by provisioning hardware and software that has future-proofed options for expansion and upgradability.

Manageability

The network should be designed to facilitate proactive network monitoring and management, to ensure ongoing stability of operation and availability of resources (see the sidebar later in this chapter, "Total Cost of Ownership"). You should consider a network management strategy as carefully as you consider the network design. What this means is that the network must work as designed, but it also must be *supportable*. If a highly complex design is delivered to the network-management team, it may require an excessive amount of your time and support to work with network-operations personnel. The key is to remember that another organization or individual may support the network you designed. Your reputation as a designer is on the line and, as the adage says, "Perception is reality."

Cost Effectiveness

The benefits of the network to the organization, however quantified, must equal or outweigh the costs. The cost of implementing the network design must be within agreed-upon budgetary constraints. For example, a design requirement may implicitly exceed the financial commitment of the customer. If that is the case, you must identify this and find an alternative solution, if one exists. If an alternative does not exist, this must be communicated back to the customer with his options so that he can make an informed decision.

Key Design Issues and Requirements

The cost trade-offs involved in internetwork design can be viewed in two ways:

- The costs could be categorized according to the technology, whether it be WAN or LAN design.
- The costs could be divided between one-time, fixed costs, and costs that recur on a regular (often monthly) basis.

In the WAN environment, the fixed costs typically are for equipment purchases, such as modems, channel service unit/data service units (CSU/DSUs), and router interfaces, as well as for any circuit-provisioning costs and network-management tools and platforms such as NetView and HP OpenView. The recurring costs are the monthly circuit fees from the service provider. Additionally, recurring costs are those for support and maintenance of the WAN, including any network management center personnel.

In the LAN environment, the fixed costs again include equipment purchases—such as routers, switches, and hubs—and the purchase and installation costs of the physical cabling for the network. The recurring costs are the salaries of administration staff who attend to daily network operations. Day-to-day operational tasks include user moves, adds, and changes related to end stations, along with the growth, maintenance, and troubleshooting of the network infrastructure.

The real cost of ownership of a network is most often overlooked in the preliminary phases of network design. Typically, recurring costs such as cost of circuits and staff tend to predominate in LAN and WAN environments. When weighing trade-offs between cost and benefit, you should first consider changes that offer the potential for reducing recurring costs. As previously stated, however, neglecting the real cost of ownership is a trap that many network designers fall into.

Total Cost of Ownership

According to a survey conducted by The Gartner Group (described further at www.cisco.com/warp/public/779/smbiz/solutions/t1.shtml), the miscalculation of labor costs involved in setting up and managing WANs has left many organizations underestimating network Total Cost of Ownership (TCO) by half. The study suggests labor accounts for a whopping 43% of some networks' TCO, with the rest of the money covering such items as training, end-user downtime, disaster prevention, and information recovery.

Design Methodology

The flowchart in Figure 1-1 outlines a simple methodology that you can use in your network design. Notice that the first three steps to be completed are one-time and sequential. The initial steps of designing the network topology and of devising addressing and naming conventions should be completed early, and should not require major revision later. These foundational steps are very important for the next steps to successfully occur. The next three steps are a recurring loop that never goes back beyond the fourth step. Keeping router Internetworking Operating System (IOS) code current through online bug searches and tools available on Cisco Connection Online (CCO), www.cisco.com, should be viewed as an ongoing task. Much like code maintenance, capacity planning and maintenance should also be an area of ongoing attention.

Figure 1-1 *Design Methodology*

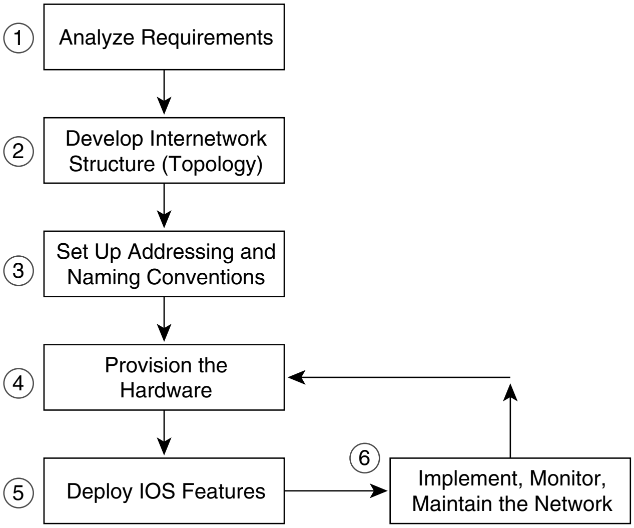
1
Analyze Requirements
2
Develop Internetwork Structure (Topology)
3
Set Up Addressing and Naming Conventions
4
Provision the Hardware
5
Deploy IOS Features
6
Implement, Monitor, Maintain the Network

Step 1: Analyze Requirements

Step 1 in the design methodology is to analyze the requirements of the network and its users. Network user needs constantly change, both in response to changing business conditions and in response to changes in technology itself. As more voice and video-based network applications (for example, Cisco IP/TV) become available, for example, the pressure to increase network bandwidth intensifies. Needs analysis includes not only the business case (cost/benefit analysis) for adopting such bandwidth-hungry applications, but also a detailed analysis of the procedures and costs required to upgrade the network to provide the needed bandwidth. Similar examples might include extending the corporate network to include regional and international offices, or integrating a telephone system into what was previously a data-only network.

Step 2: Develop the Internetwork Structure

Step 2 in the design methodology is to develop the overall network topology using a three-tiered hierarchical model. In this model, the network is divided into core, distribution, and access tiers. These tiers describe a set of discrete functions performed at each tier, as well as a network topology typically associated with each tier. Each tier within the topology has its own function. By keeping the tiers separate, the hierarchical design method produces a highly flexible and scalable network. Figure 1-2 shows an internetwork structure that follows the hierarchical model of network design.

Figure 1-2 *Develop the Internetwork Structure*

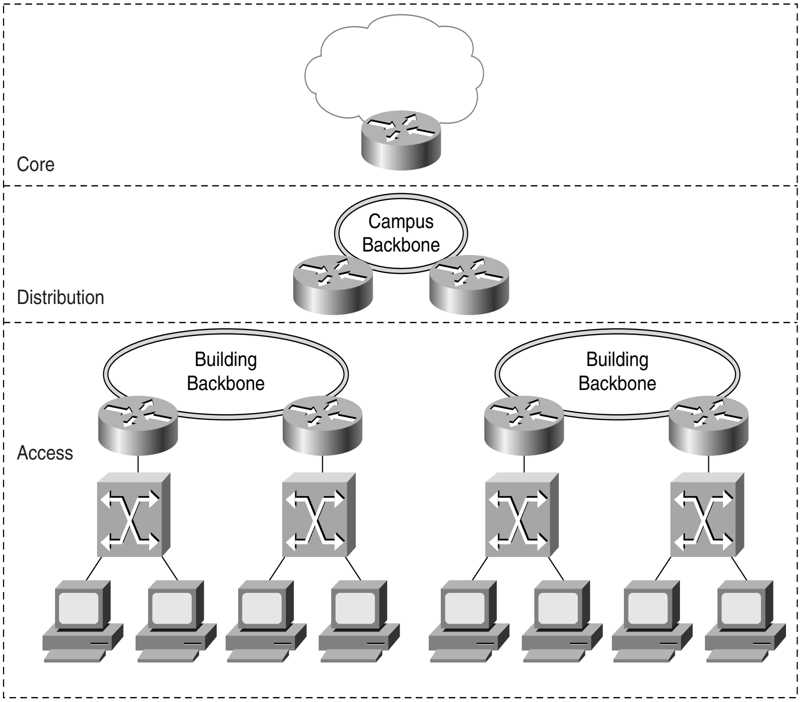

In the hierarchical model, the core tier primarily provides the wide-area links between geographically remote sites, tying a number of campus networks together in a corporate or enterprise WAN. Hosts are rarely in the core tier; core services are typically leased from a telecom service provider (for example, T1/T3, Frame Relay, SMDS, and ATM). In fact, placing hosts in the core is not generally recommended.

NOTE The use of the word *tier* here is unrelated to the seven-layer OSI reference model. The three-tiered model describes a network design methodology, not a protocol stack. It is worth noting, however, that a boundary between topological tiers is usually created with an OSI reference model Layer 3 device (for example, a router).

The distribution tier generally refers to the distribution of network services to multiple LANs within a campus network environment. This tier is where the campus backbone network is found and is typically based on FDDI, Fast or Gigabit Ethernet, or sometimes campus ATM. In Figure 1-2, the distribution tier is implemented as a FDDI ring. This tier is often where network policy is implemented as well (that is, security, access lists, and so on). In many corporate networks, for example, you'll find that the distribution tier is commonly where company-wide servers (such as e-mail, Internet proxy, firewalls, demilitarized zones [DMZs], and so on) are placed.

The access tier is usually a LAN or a group of LANs, typically Ethernet or Token Ring, that provides users with first-line access to network services. The access tier is where almost all hosts are attached to the network, including servers of all kinds as well as user workstations.

Step 3: Set Up Addressing and Naming Conventions

Step 3 in the design methodology is to develop the overall addressing scheme by assigning blocks of addresses to portions of the network, thus simplifying address administration and producing a more scalable internetwork. In the TCP/IP example in Figure 1-3, IP address 10.0.0.0 with 16 bits of subnetting is used throughout the organization. This campus has been allocated a contiguous block of 254 of these addresses. The address is then further allocated, with each building receiving approximately 16 contiguous subnets.

Figure 1-3 *Set Up Addressing and Naming Conventions*

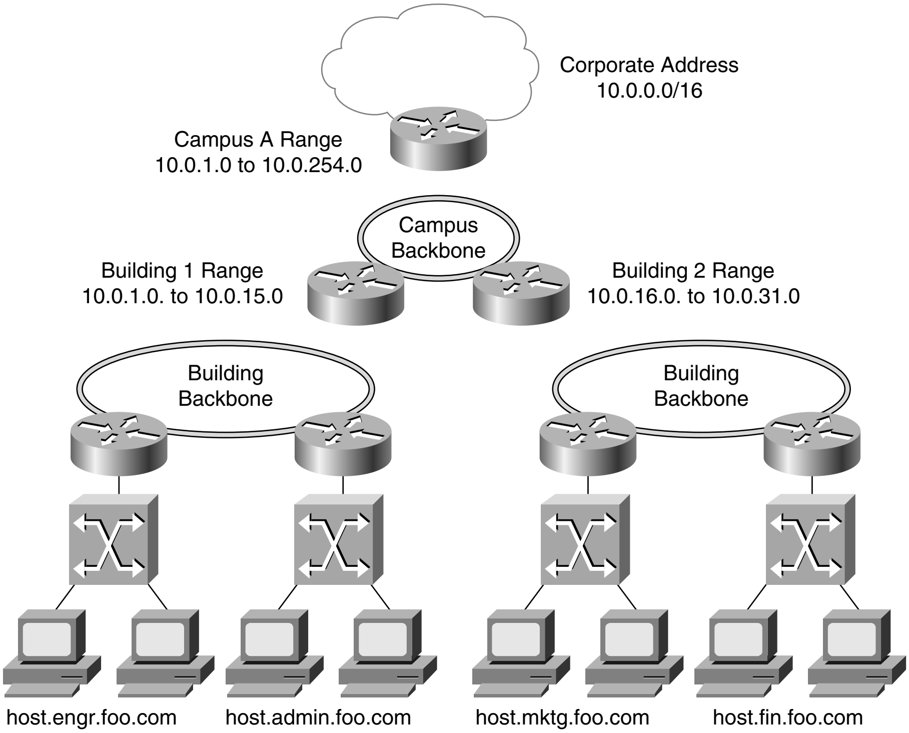

If the routing protocols used in the internetwork support variable-length subnet masking (VLSM), you can deploy a true hierarchical addressing scheme. In a generic TCP/IP example, an 8-bit subnet mask can be in use at the core tier, a 16-bit mask can be in use at the distribution tier, and a 24-bit mask can be applied to the access tier. Careful allocation of addresses in a hierarchical design can result in efficient summarization of routes in the routing tables.

Consider 10.0.0.0 as a sample network number that you are assigned. Your network will consist of 50 regions. An example of an 8-bit subnet mask will be chosen at the core for summarization. The subnet and subnet mask at the core location would be 10.0.0.0 255.0.0.0. The result is that all the specific route entries at the access and distribution tiers will not be seen to the outside world. After you do this, the core basically says to any external network, "If your IP packet destination is '10.anything,' come to me." At the distribution tier, a 16-bit mask is used to identify each region without allowing all the specific routes at the access layer to be propagated to other distribution layer routers. A typical subnet and mask combination at the distribution tier would be 10.1.0.0 255.255.0.0.

After you do this, the distribution tier router basically says to any external network, "If your IP packet destination is '10.1.anything,' come to me." Finally, the access tier uses a 24-bit mask. A typical subnet and mask combination at the access tier would be 10.1.1.0 255.255.255.0. As you migrate down the hierarchical model, your subnet masks become more granular (more specific), and, as you move back to the top, the masks become less specific. By using hierarchical addressing to fit your hierarchical model, the result is *scalability* and *stability*. Networks following this model can scale to thousands of nodes and be extremely stable.

The naming scheme is also designed in a systematic way, with common prefixes used for naming components within an organization. The systematic naming convention also makes the network more scalable and easier to manage. In Figure 1-3, each site has a named prefix that identifies the domain to which it belongs. All admin hosts contain admin.foo.com in their name, for example. Generally, onto this name you will prepend a name such as tiffany.admin.foo.com for a specific user's workstation (Tiffany's workstation) within the hostadmin domain. This proves invaluable for documentation and subsequent troubleshooting and also should be applied to other areas where naming is important, such as router names and interface descriptions. In the future, as more sites are added, you have room to select some unused addresses within a preallocated range. Also, the names become predictable and organized for optimal support.

Step 4: Provision the Hardware

Step 4 in the methodology is to use vendor documentation to select the LAN and WAN hardware components, to implement the internetwork design. LAN devices include router models, switch models, cabling systems, and backbone connections. WAN devices include modems, CSU/DSUs, and remote access servers. The selection process typically includes consideration of the function and features of the particular devices, including their expandability and management capabilities. The initial cost of the equipment is always part of the decision process.

An example of a hierarchical hardware design uses a Cisco 7200 class router at the core tier, a Cisco 3600 class router at the distribution tier, and a Cisco 2600 class router at the access tier. An excellent tool available on CCO can assist you in router product selection. You can find this tool at www.cisco.com/public/reseller_mktplace/select.html. Note also that the IOS version may be impacted by the type of hardware selected; certain hardware features require specific levels of software. An example would be the new NPE-300 CPU for the 7200 platform. This CPU requires a minimum IOS of 12.0 to operate and may impact (in a positive or negative way) your ultimate hardware decision.

Step 5: Deploy Cisco IOS Software Features

Step 5 in the methodology is to deploy Cisco IOS features as appropriate. Many of these features, such as access lists, proxy features, traffic shaping, queuing, QoS, and compression, are available to support bandwidth management. Other features provide support for security, general network management, tariff management, and bandwidth management.

Initial deployment of Cisco IOS features is directly related to the applications in use on host machines. Consideration for protocol types is critical. You might consider asking yourself some of the following questions:

- Are the protocols routable, like IP and IPX?
- Are the protocols nonroutable, like NetBIOS and SNA, DEC MOP (Maintenance Operation Protocol) or LAT (local-area transport) traffic?

Such planning will ensure that NetBIOS doesn't get added as an afterthought—for example, when users of Windows machines accept the defaults. Also, you need to understand which mechanisms are in place to do remote management and router discovery. If so, where are the network-management systems (NMSs) located, and how often will they poll the network?

Additional deployment of features is based on the three-tier hierarchical model. The access tier interconnect devices often employ static definitions, proxy services, and first-cut filtering. The distribution tier typically uses features providing compression, congestion control, and security. Finally, the core tier generally focuses on QoS (Quality of Service) features (for example, WRED—Weighted Random Early Detection). For more information on WRED, check out www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/wred.htm.

Step 6: Implement, Monitor, and Manage the Network

The last step in the methodology is to deploy the network. If possible, model and test the new design in a lab environment prior to full deployment. A very good way to do this is to use a modeling tool called Netsys. Netsys enables you to see the network impact of proposed configuration changes before they are made. A good general-reference Web page for Netsys can be found at www.cisco.com/warp/public/458/54.html. Later on, during actual deployment, you should use a phased approach to reduce the impact on the user communities. Perform operational data gathering continuously so that planning for higher-bandwidth applications can occur proactively. A number of SNMP and RMON tools that allow for proactive network management are available today. Deploy new hardware and Cisco IOS software features as required to support new applications. Remember to have solid test and contingency plans before you implement any new design!

Summary

A successful network design requires persistence. As a designer, you are responsible for making your network design a success. This chapter covered a number of issues that you should consider when approaching any internetwork design and outlined an internetwork design methodology to keep you focused on providing a fully functional network design solution. The better you understand your design goals and requirements, the more successful your design will be.

Chapter Review Questions

1. List and describe five goals of a sound network design.
2. What is the key trade-off that must be made in most network designs?
3. Describe the steps that need to take place in a good network design methodology.
4. At which tier of the hierarchical model is QoS usually deployed?
5. Why is scalability so important in good network design?
6. Why is adaptability so important in good network design?
7. At which tier of the hierarchical model are firewalls most often deployed?