



## **CHAPTER CONTENTS**

- 1.1 Objectives**
- 1.2 Preparation**
- 1.3 Background**
- 1.4 Overview of Analysis, Architecture, and Design Processes**
  - 1.4.1 Process Components**
  - 1.4.2 Tactical and Strategic Significance**
  - 1.4.3 Hierarchy and Diversity**
  - 1.4.4 Importance of Network Analysis**
  - 1.4.5 Model for Network Analysis, Architecture, and Design**
- 1.5 A Systems Methodology**
- 1.6 System Description**
- 1.7 Service Description**
- 1.8 Service Characteristics**
  - 1.8.1 Service Levels**
  - 1.8.2 System Components and Network Services**
  - 1.8.3 Service Requests and Requirements**
  - 1.8.4 Service Offerings**
  - 1.8.5 Service Metrics**
- 1.9 Performance Characteristics**
  - 1.9.1 Capacity**
  - 1.9.2 Delay**
  - 1.9.3 RMA**
  - 1.9.4 Performance Envelopes**
- 1.10 Network Supportability**
- 1.11 Conclusion**
- 1.12 Exercises**

# Introduction

I begin this book with a description of the analysis, architecture, and design processes. Many of the concepts and terms used throughout this book are introduced and defined in this chapter. Some of these concepts may be new to you, while others are presented in a different light. Glossaries of terms and acronyms are presented at the end of this book for easy reference.

## 1.1 Objectives

In this chapter I will introduce the fundamental concepts of this book: that the network is part of a system that provides services to its end users; that there are processes for developing an analysis, an architecture, and a design for a network; and that there are ways to characterize a network.

## 1.2 Preparation

In order to understand and apply the concepts in this chapter, you should be familiar with basic networking concepts. This includes the functions and features of the TCP/IP protocol suite, technologies such as the variants of Ethernet, synchronous optical network (SONET), and wave division multiplexing (WDM), and the basics of network routing, security, performance, and management.

## 1.3 Background

Network analysis, architecture, and design have traditionally been considered art, combining an individual's particular rules on evaluating and choosing network technologies; knowledge about how technologies, services, and protocols can be meaningfully combined; experience in what works and what doesn't; along with (often arbitrary) selections of network architectures. However, as with other types of art, success of a particular network design often depends primarily on who is doing

## 4 CHAPTER 1 Introduction

the work, with results that are rarely reproducible. This may have been acceptable in the early days of networking, when networks were more of a hobby than a critical resource and did not directly support revenue generation. Today, however, networks are embedded within our work, home, and outside environments. They are considered “mission-critical”<sup>1</sup> to corporate success and provide near real-time access to information throughout the world. As such, the design of a network must be logical, reproducible, and defensible. This premise is the foundation for this book.

Traditionally, network analysis, architecture, and design have been based on developing and applying a set of rules for the network. In developing a set of rules, an individual may draw from personal experience as well as from general rules such as the 80/20 rule (where 80% of a network’s traffic is local and 20% is remote) or the adage “bridge when you can, route when you must” (bridging being simpler, easier, and cheaper at the time). As we see later in this book, although both of these rules are ancient from the perspective of networking history, they still apply today, albeit in modified form. Such rules were useful when there weren’t many choices in network technologies and services, and when the differences between choices were clearly understood. But times have changed, and our notion of designing networks must adapt to the variety of options now available to us, the variety of services that networks can offer to end users, and the subtle nuances brought about by combining network technologies, techniques, and services.

---

### Example 1.1.

Consider the subtleties in network behavior introduced through the use of virtual private networks, intranets, or VPNs. VPNs are quite useful; however, care must be taken to understand their potential impact on network security, routing, and management. Since VPNs tunnel (encapsulate) and can encrypt traffic flowing across a network, they often require more effort to secure, monitor, and manage. How VPNs impact security, routing, and management will be considered during the architecture process.

---

Network analysis, architecture, and design have traditionally focused on *capacity planning*, which is over-engineering a network to provide an amount of capacity (also known as *bandwidth*) estimated to accommodate most short- and long-term traffic fluctuations over the life cycle of the design. The result is a bandwidth

---

<sup>1</sup>Ambiguous terms such as these will be defined in this chapter.

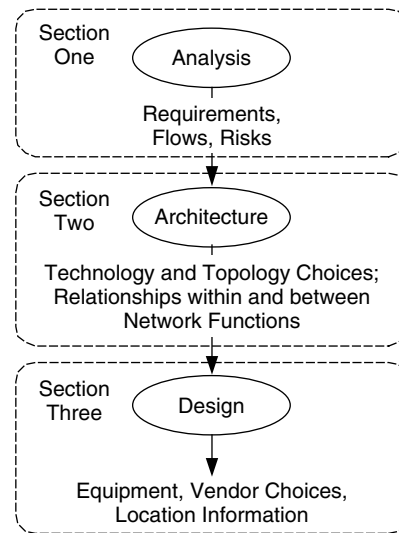
“buffer” that can handle these fluctuations. As network traffic grows over time, this bandwidth buffer is reduced, and users experience problems related to traffic congestion. This is an inefficient use of network resources, wasting money up front in resources that are not used while failing to provide the flexibility needed to adapt to users’ changing traffic requirements. Network bandwidth is only one component of network resources that we must consider. We also need to consider how delay through the network, as well as network reliability, maintainability, and availability (RMA), can be optimized. In today’s evolving networks, delay and reliability can be more important than capacity.

In this book we explore how the analysis, architecture, and design processes have changed and how they continue to change. We discuss how these processes work together in engineering a new or existing network. We approach networks from a different perspective—as a system providing services to its users—and we discuss how networks can be designed to provide many different types of services to users. In taking this approach we emphasize network analysis, which helps us understand what is required of a network in supporting its customers and their applications and devices. As we will see, these processes require an investment in time and effort, but the return on investment is significant. These are powerful tools that can help you build better networks, improving the ability of your organization to get its work done.

This book begins by applying a systems methodology to networking. This methodology is relatively new, and you will learn a number of useful definitions in regard to network analysis, architecture, and design. The rest of this book is logically divided into three sections. The first section covers the analysis process: specifically, how to develop requirements, understand traffic flows, and conduct a risk analysis. The analysis process prepares you for dealing with network architecture, discussed in the second section. Here I describe how to make technology and topology choices for your network, how to understand the relationships among the various functions within your network, and how to use this information to develop an architecture. In the final section the network architecture is used as input for the design process, where location information, equipment, and vendor selections are used to detail the design. Information flows between analysis, architecture, and design processes are presented in Figure 1.1.

Network analysis, architecture, and design will help you identify and apply network services and performance levels needed to satisfy your users. Through these processes you will be able to understand the problems you are trying to address with the new network; determine the service and performance objectives

## 6 CHAPTER 1 Introduction



**FIGURE 1.1** Information Flows Between Network Analysis, Architecture, and Design

needed to tackle these problems; and architect and design the network to provide the desired services and performance levels.

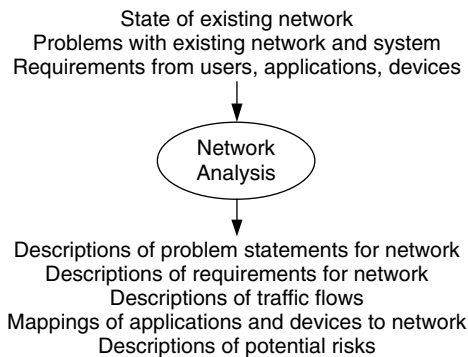
### 1.4 Overview of Analysis, Architecture, and Design Processes

Network analysis, architecture, and design are processes used to produce designs that are logical, reproducible, and defensible. These processes are interconnected, in that the output of one process is used directly as input to the next, thus creating flows of information from analysis to architecture, and from architecture to design.

*Network analysis* entails learning what users, their applications, and devices need from the network (Figure 1.2). It is also about understanding network behavior under various situations. Network analysis also defines, determines, and describes relationships among users, applications, devices, and networks. In the process, network analysis provides the foundation for all the architecture and design decisions to follow. The purpose of network analysis is twofold: first, to listen to users and understand their needs; and second, to understand the system.

In analyzing a network we examine the state of the existing network, including whatever problems it may be having. We develop sets of problem statements

## Overview of Analysis, Architecture, and Design Processes 7



**FIGURE 1.2** Inputs To and Outputs From the Network Analysis Process

and objectives that describe what our target network will be addressing. And we develop requirements and traffic flows, as well as mappings of users, applications, and devices, in support of our problem statements and objectives. As such, network analysis helps us understand what problems we are trying to solve, and in the process we compile information that will be used in developing the architecture and design.

---

### Example 1.2.

The analysis, architecture, and design processes can be applied to any network project, regardless of size or scope. Since we are developing sets of problem statements, objectives, and requirements as input to the analysis process, we can scale the architecture and design to meet the scope of the project. Consider the use of VPNs from Example 1.1. We can develop problem statements, objectives, and requirements for VPNs in an existing network, and develop an analysis, architecture, and design solely around a VPN deployment.

---

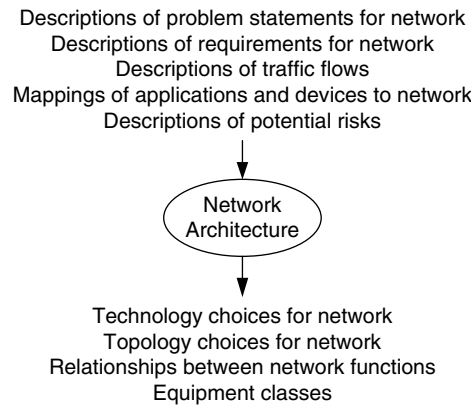
*Network architecture* uses the information from the analysis process to develop a conceptual, high-level, end-to-end structure for the network. In developing the network architecture we make technology and topology choices for the network. We also determine the relationships among the functions of the network (addressing/routing, network management, performance, and security), and how to optimize the architecture across these relationships. There usually is not a single “right” architecture or design for a network; instead there are several that will work, some better than others. The architecture and design processes focus on

## 8 CHAPTER 1 Introduction

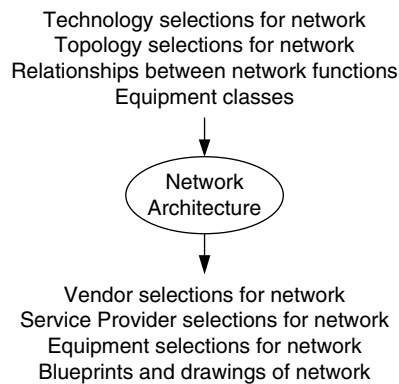
finding those best candidates for architecture and design (optimized across several parameters) for your conditions.

The network architecture process determines sets of technology and topology choices; the classes of equipment needed; and the relationships among network functions (Figure 1.3).

*Network design* provides *physical* detail to the architecture. It is the target of our work, the culmination of analysis and architecture processes. Physical detail includes blueprints and drawings of the network; selections of vendors and service providers; and selections of equipment (including equipment types and configurations) (Figure 1.4).



**FIGURE 1.3** Inputs To and Outputs From the Network Architecture Process



**FIGURE 1.4** Inputs To and Outputs From the Network Design Process



## Overview of Analysis, Architecture, and Design Processes 9

During network design we use an evaluation process to make vendor, service provider, and equipment selections, based on input from the network analysis and architecture. You will learn how to set design goals, such as minimizing network costs or maximizing performance, as well as how to achieve these goals, through mapping network performance and function to your design goals and evaluating your design against its goals to recognize when the design varies significantly from these goals. Network design is also about applying the trade-offs, dependencies, and constraints developed as part of the network architecture. Trade-offs, such as cost versus performance or simplicity versus function, occur throughout the design process, and a large part of network design concerns recognizing such trade-offs (as well as interactions, dependencies, and constraints) and optimizing the design among them. As part of the design process you will also learn how to develop evaluation criteria for your designs.

As we show throughout the remainder of this book, network analysis, architecture, and design combine several things—requirements, traffic flows, architectural and design goals, interactions, trade-offs, dependencies, constraints, and evaluation criteria—to optimize a network’s architecture and design across several parameters. These parameters are chosen and analyzed during the analysis process and prioritized and evaluated during the architecture and design processes. On completion of these processes you should have a thorough understanding of the network and plenty of documentation to take you forward to implementation, testing, and integration.

---

### Example 1.3.

A network’s architecture and design are analogous to the architecture and design of a home. Both the network and home architecture describe the major functional components of each (for the network: network management, addressing and routing, security and privacy, and performance; for the home: plumbing, electrical, HVAC [heating, vacuum, air conditioning], framing) and the relationships among them (for the network: interactions, dependencies, trade-offs, and constraints; for the home: where each component is placed relative to the others). The network and home designs are also similar in that they both provide physical detail to the architecture. For the network this means where major network devices are located; and, for the home, where ducts, outlets, faucets, drains, and so forth are located.

---

### 1.4.1 Process Components


We now add detail to the analysis, architecture, and design processes. Each of these processes has actions that will be taken by project personnel and results or products

## 10 CHAPTER 1 Introduction

of each action. There is also input to begin the process. Thus, the processes are expanded in Figure 1.5.

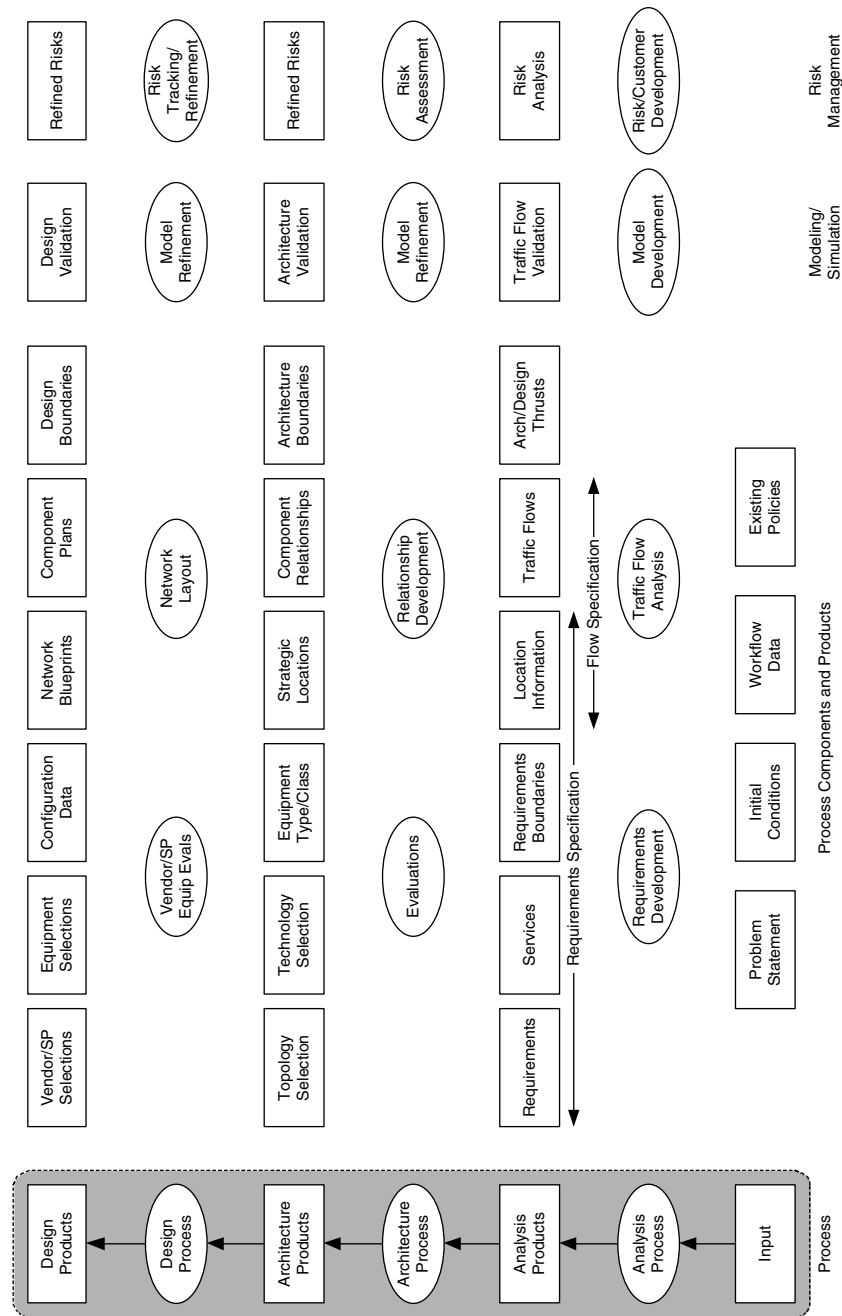
Each of the processes and products has components that describe specific actions or results. The full set of process components is shown in Figure 1.6.

This set of process components represents a complete implementation of network analysis, architecture, and design, and forms the basis for the remainder of this book. Some components, however, may be reduced in importance or removed



**FIGURE 1.5** Processes Shown with Associated Input and Products

### Overview of Analysis, Architecture, and Design Processes 11



**FIGURE 1.6** The Full Set of Process Components

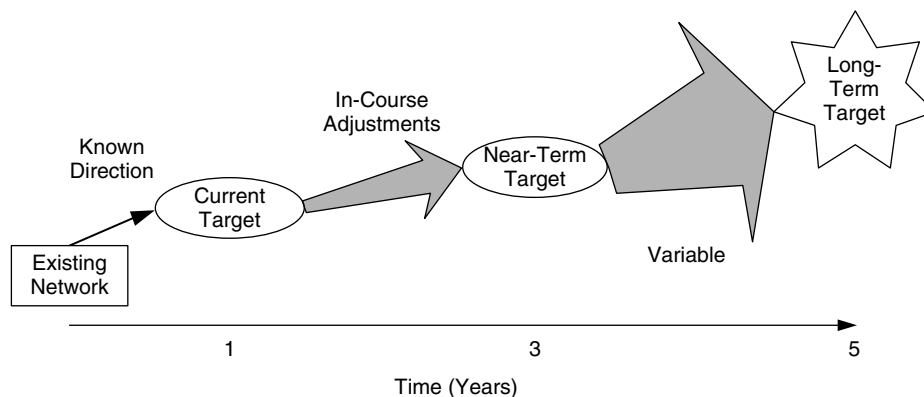
## 12 CHAPTER 1 Introduction

on a per-project basis. Throughout this book we discuss which components are necessary, and which may be reduced in importance.

### 1.4.2 Tactical and Strategic Significance

Network analysis, architecture, and design are part of the engineering process that forms the basis of networking projects. Such projects have immediate, tactical (near-term), and strategic (long-term) significance, and networking projects should consider all of these areas. I recommend that network projects have a plan that includes current, near-term, and long-term targets. While the current target will be a network design, the near-term and long-term targets can be proposed enhancements to the current target, lists of problem statements, objectives, and requirements for near-term and long-term, or all of these. For example, Figure 1.7 shows a one-year/three-year/five-year project plan.

The idea behind this plan is to develop a network design that will be implemented within one year, will prepare us for any changes we might need to make to the network within three years, and will keep us in the direction of what is planned for five years in the future. The long-term (five-year) target is a rough estimate. We will likely not know what new networking technologies, services, or levels of performance will be available five years out, nor will we know how our customers' business plans will change, nor what network problems will arise during that time. But we should have an idea of where we want to be, with the understanding that we may have to make significant changes to the long-term target during those five years. Thus the long-term target is variable.



**FIGURE 1.7** A One-/Three-/Five-Year Project Plan

## Overview of Analysis, Architecture, and Design Processes 13

The current (one-year) target should be well understood and is the focus of our analysis, architecture, and design effort. In between the one-year and five-year targets, the near-term (three-year) target is intended to be somewhat flexible, yet better understood than the long-term target. A significant change in the long-term target (e.g., the result of a planned merger, outsourcing, or major change in a company's business) can be mediated by course corrections in the near-term plan.

Although a one-/three-/five-year plan is shown here, the important concept is to have both tactical and strategic approaches to your plan. Experience shows that one-/three-/five-year plans are very good starting points, but depending on your customers, you may rapidly evolve your network with a six-month/one-year/two-year plan, or take a longer-term view with a one-year/five-year/ten-year plan. I have seen all of these plans work to meet the needs of their customers.

---

### Example 1.4.

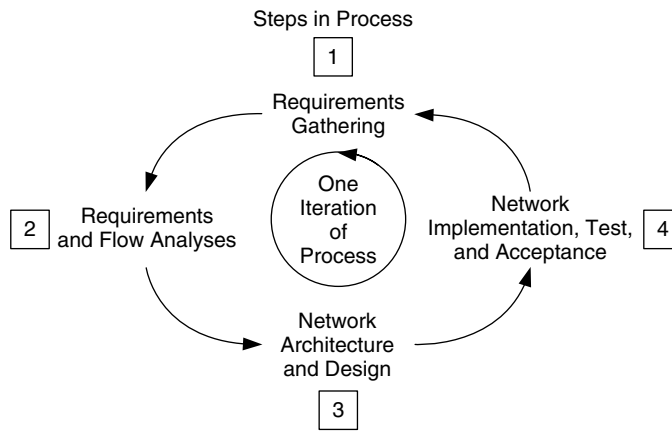
Voice over IP (VoIP) is of interest to many organizations and is an example of a network project that would benefit from tactical and strategic plans. If we apply the one-/three-/five-year plan discussed earlier, the current target (one-year plan) would include the network design for VoIP, based on what is achievable within one year, and the problem statements, objectives, and requirements that result from the requirements analysis process. For example, the current target may be a design that only prepares for VoIP by improving the overall reliability of the network. The near-term target (three-year plan) would conceivably build on the current target to add or expand VoIP to those areas that can support it. The long-term target (five-year plan) would address any major changes that occurred over the previous four years, including advancements in VoIP technology and an assessment whether to continue with VoIP or evolve to new or different technologies.

---

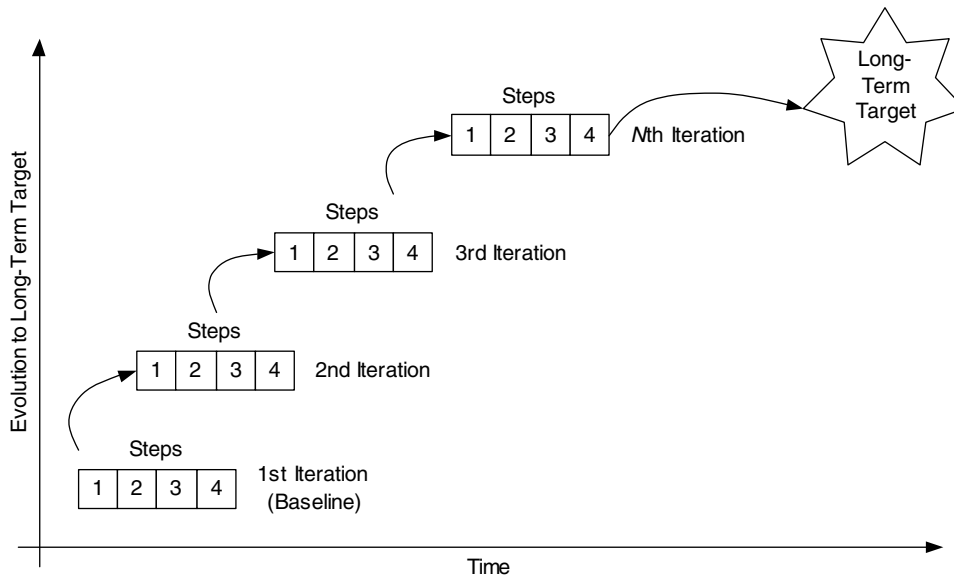
These plans are intended to be iterative and should be regularly reviewed, on the order of twice yearly, once per year, or every two years, depending on your plan. At each iteration the current, near-term, and long-term targets are reviewed and checked against the ongoing sets of problem statements, objectives, and requirements developed during the analysis, architecture, and design processes. One iteration of the cycle (including network implementation, test, and acceptance) is shown in Figure 1.8.

Each iteration is an incremental step toward the near-term and long-term targets, as shown in Figure 1.9. The steps 1, 2, 3, and 4 correspond to the steps in the process shown in Figure 1.8. Thus each iteration is a full cycle as shown above.

14 CHAPTER 1 Introduction



**FIGURE 1.8** The Cyclic and Iterative Nature of Processes



**FIGURE 1.9** Process Iterations Evolve Toward the Long-Term Target

**1.4.3 Hierarchy and Diversity**

All of these processes center around two important characteristics of networks: their levels of hierarchy and diversity. *Hierarchy* is the degree of concentration of networks or traffic flows at interconnection points within the network, as well as

## Overview of Analysis, Architecture, and Design Processes 15

the number of tiers of interconnection points within the network. In general, as networks grow in size and numbers of users, applications, and devices increase, hierarchies provide separation and structure within the network. Hierarchies are important because they help us in determining the sizes of networks, including routing and addressing configurations, and the scaling of network technologies, performance, and service levels. A key concept of this book is understanding these hierarchies, learning how and where they will occur, and learning how to take advantage of them.

Along with hierarchy, there must be some consideration for the degree of *diversity* (a.k.a. redundancy or interconnectivity) in the network design. As hierarchy provides structure in the network, diversity balances this structure by interconnecting the network at different levels in the design to provide greater performance through parts of the network. Diversity is important in that it provides a mechanism to achieve performance within a hierarchical structure. The dynamic between hierarchy and diversity is perhaps one of the most fundamental trade-offs in network architecture and design, and it shows up several times in the analysis, architecture, and design processes.

Hierarchy and diversity may be a bit confusing at this point, but this concept will become clearer as we progress through the book. Hierarchy is fundamental to networking (as it is throughout nature) because it provides a separation of the network into segments. These segments may be separate, smaller networks (subnets) or broadcast domains. Hierarchy is necessary when the amount of traffic on the network grows beyond the capacity of the network or when interactions between devices on the network result in congestion (e.g., broadcast storms).

Figure 1.10 illustrates levels of hierarchy and diversity in a network. This is a typical tree structure for a network, with circles representing networks or routers and lines representing the communications links between networks and/or routers. In this figure there are four levels of hierarchy, from core (or backbone) networks to access networks closest to users. Note that the end points of this tree (commonly referred to as *leaves*; they represent the end networks, devices, or users) all occur at the same level of hierarchy. This does not have to be the case; indeed, in most networks there are leaves at most levels of hierarchy.

An example of adding hierarchy to a network is changing from a flat (bridged or layer 2 switched) structure to a routed structure. This may be done to reduce the size of the broadcast domain or the number of devices reached by a broadcast message. Adding routing to the network breaks a broadcast domain into a number of smaller broadcast domains, and traffic flows are concentrated at routers. Figure 1.11 shows this scenario. Hierarchy is also added to networks when evolving from

16 CHAPTER 1 Introduction




FIGURE 1.10 Hierarchy and Diversity in a Network

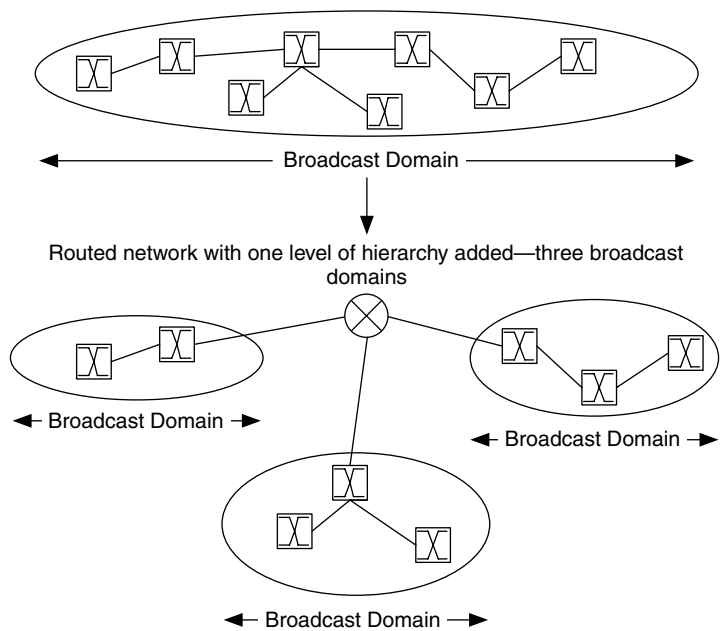


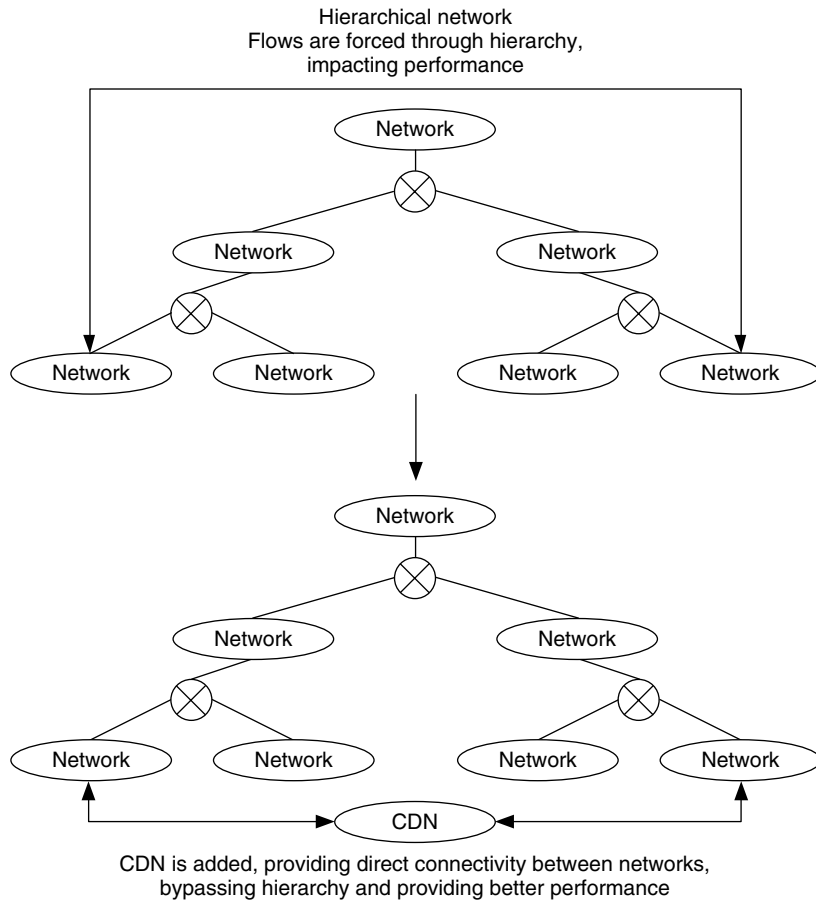
FIGURE 1.11 Hierarchy Added to a Network



Overview of Analysis, Architecture, and Design Processes 17

a single autonomous system (AS) to connecting multiple ASs, as well as when migrating from Interior Gateway Protocols (IGPs) to Exterior Gateway Protocols (EGPs) and to policy-based routing.

A content delivery network (CDN) is an example of adding diversity to a network. A CDN bypasses the core of a network, where congestion is most likely to occur, and directly connects devices or networks lower in the hierarchy (Figure 1.12). This provides better, more predictable performance but can also affect the network hierarchy by modifying its routing behavior.



**FIGURE 1.12** Diversity Added to a Network

## 18 CHAPTER 1 Introduction

### 1.4.4 Importance of Network Analysis

The importance of network analysis is emphasized in this book because experience has shown that networking personnel find it extremely valuable—once they are convinced of its importance. Analysis takes work, and when you know that there will be a payoff, you are more likely to do that work.

In this book you will learn how to gather and analyze network requirements and traffic flows. *Network requirements* are requests for capabilities in the network, usually in terms of performance and function, which are necessary for the success of that network. Network requirements can be gathered and/or derived from customers, applications, and devices. Such requirements are fundamental to a network's architecture and design because they form the basis for customer expectations and satisfaction. Requirements, in conjunction with measurements on the existing network (if there is one), are used to derive traffic flows (sets of network traffic that have some common attributes, such as source/destination address, information type, routing, or other end-to-end information). Analysis of these flows imparts location and directional information onto requirements. This is where performance requirements and architecture start to converge and is often the point in these processes where one can begin to see where “hot spots”—focal points for network performance—will appear in the network. As we will see, evaluating security risks is also part of the analysis process.

Results of the analysis process, the requirements and flow specifications, are then used as input for both network architecture and design. In developing the network architecture, a number of component architectures, targeting particular functions of the network, are evaluated. Desired component architectures are then combined into the reference architecture, which provides a high-level view of your network. This high-level view is then physically detailed during the network design process.

Network analysis is important in that it helps us understand the complexity and nuances of each network and the systems they support. Analysis also provides data upon which various decisions are made, and these data can and should be documented as part of an audit trail for the architecture and design processes. Such data help ensure that the resulting architecture and design are defensible.

#### ***Understanding Network and System Complexity***

In general, networks and the systems they support are becoming increasingly complex. Part of this complexity lies in the sophistication of the capabilities provided by that network. Consider, for example, how services can be incorporated into a current state-of-the-art network. Infrastructure capacity planning, which often

## Overview of Analysis, Architecture, and Design Processes 19

includes traffic over-engineering, may now be expanded to include support for delay-constrained applications and may contain a variety of capacity and delay control mechanisms, such as traffic shaping, quality of service at multiple levels in the network, service-level agreements to couple services to customers, and policies to govern and implement service throughout the network. (Note that *quality of service* refers to determining, setting, and acting on priority levels for traffic flows. A *service-level agreement* is an informal or formal contract between a provider and user that defines the terms of the provider's responsibility to the user and the type and extent of accountability if those responsibilities are not met. Finally, *policies* are high-level statements about how network resources are to be allocated among users.) Analysis of these mechanisms—how they work and interoperate—is covered in detail later in this book.

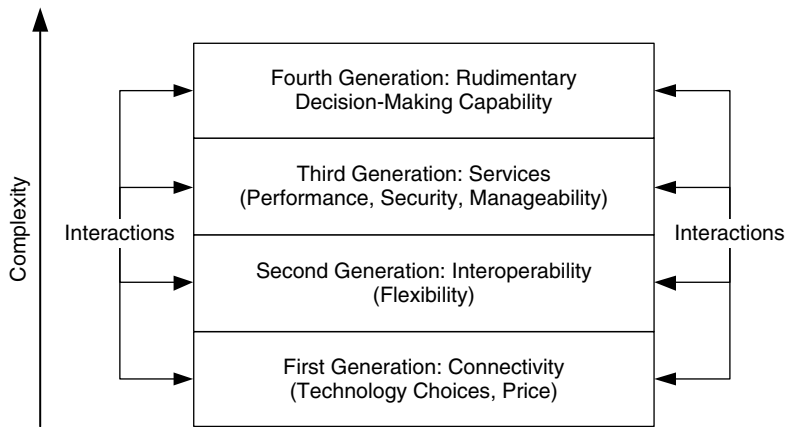
Network and system complexity is nonlinear. Network optimization must consider competing and often conflicting needs. In addition, multiple groups with differing ideas and desires (e.g., users, corporate management, network staff) influence the network design. The network is either designed by committee or through a systematic approach that the groups can agree on.

Networks have evolved to incorporate more sophisticated capabilities. Early (first-generation) networks focused on supporting basic connectivity between devices and on how to scale networks to support growing numbers of users (e.g., segmenting networks using bridges or routers). Second-generation networks focused on interoperability to expand the scope and scale of networks to allow connections among multiple disparate networks. We are currently at the stage in network evolution where service delivery is important to the success of users and their applications. This stage can be considered the third generation of networking. Figure 1.13 illustrates the various generations of networking and their interactions.

We are beginning to see steps toward next-generation capabilities, such as rudimentary decision making within the network. It may be expected that components of the network will evolve to become self-configurable and manageable, especially for those networks that must be configured or administered by end users (e.g., telecommuters, users of mobility/portability services). Indeed, this will become necessary as the complexity and performance of networks increase and as services offered by networks become more sophisticated. Grid networks are a clear step in this direction.


Users, applications, and devices are also evolving more sophisticated capabilities. An example of this is the dynamic between hierarchy and diversity that can be seen in the current Internet. As application and device traffic flows evolve to incorporate information regarding quality, performance, and cost (e.g., real-time

20 CHAPTER 1 Introduction



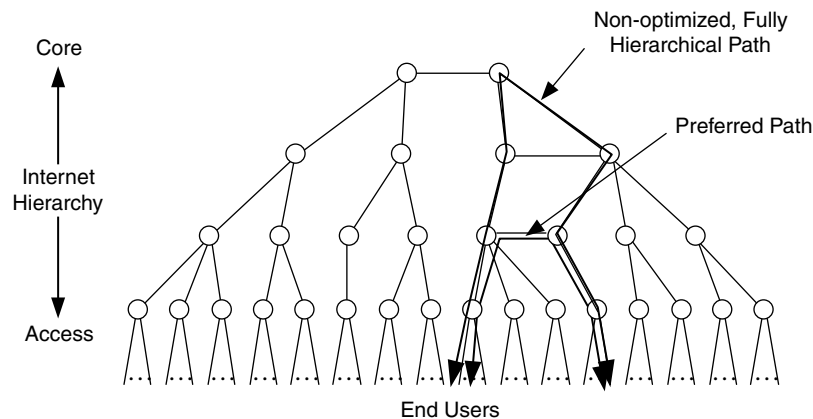
**FIGURE 1.13** Generations of Networking

streaming media), it is expected that these characteristics can be used to ensure paths through the Internet that will support high-performance or high-quality delivery of such traffic. Hierarchy in the Internet often forces traffic flows through nonoptimal paths, crossing several ASs with differing performance and service characteristics, hindering high-performance, high-quality delivery. Figure 1.14 shows a hierarchy of multiple levels of networks from core (or backbone) network providers to access network providers. Traffic flows between end users may travel across several levels of this hierarchy.



**FIGURE 1.14** Hierarchy and Traffic Flow

## Overview of Analysis, Architecture, and Design Processes 21

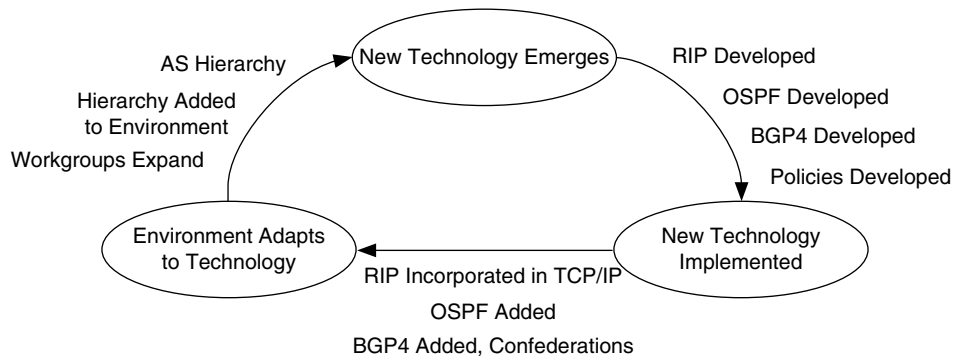


**FIGURE 1.15** Diversity Added to Improve Performance of Select Traffic Flows

To counteract the impact of hierarchy, diversity can be introduced into the Internet at strategic points, providing shortcuts that bypass parts of the Internet. The result is that, for some select flows, paths are optimized for high-performance, high-quality delivery (Figure 1.15). The dynamic between hierarchy and diversity exists in all networks to some degree, and part of the analysis process is determining where and how to apply it. In Figure 1.15 connections are added between networks at the same level of hierarchy, in essence providing a “shortcut” or “bypass” around part of the Internet, resulting in better performance characteristics. This concept of adding diversity to improve performance along select paths can be readily applied to enterprise networks.

Analysis helps us understand how technologies influence networks, users, applications, and devices (and vice versa). This is important for gauging how users of the network will adapt to the network, which affects the overall life cycle of the network. Consider, for example, the evolution of routing protocols, shown in Figure 1.16. Although the Routing Information Protocol (RIP), an Interior Gateway Protocol (IGP) deployed as part of early TCP/IP releases, was simple and easy to use, its limitations were stressed as networks expanded to accommodate larger groups of users (workgroups) and even groups of networks (autonomous systems [ASs]). Routing technology adapted by adding hierarchy to routing, in terms of new IGPs such as Open Shortest Path First (OSPF), as well as in development of Exterior Gateway Protocols (EGPs) such as Border Gateway Protocol (BGP), which can accommodate hierarchy in groups of networks (AS hierarchy). This process continues today as high-level policies are being introduced to control

22 CHAPTER 1 Introduction



**FIGURE 1.16** Routing Evolution

routing at a level above IGP or EGP, and BGP4 introduces hierarchy through grouping BGP4 routers into confederations. We discuss routing protocols in detail in the addressing and routing architecture (see Chapter 6).

Similarly, users, applications, and devices also influence their environment. As new, upgraded, or different users, applications, and devices are added to a network, the requirements on that network may change. The analysis process must examine how high-end computer and applications servers, data storage, analysis and archival systems, and specialized environment-specific devices such as PDAs, video cameras, or medical equipment impact the network.

Finally, the analysis process helps us understand the forces and changes at work within the system (network and its users, applications, and devices). Networks are highly dynamic, changing the rest of the system and being changed by the system. Some of the factors leading to change in the system include usage behavior and patterns; what, how, where, and when each user impacts the system; the emergence of new capabilities—for example, optical switching, faster central processing units (CPUs), and cheaper memory; and changes in scope and scale of the environment, including consolidation and outsourcing.

**Architecture and Design Defensibility**

An important (and often overlooked) part of network analysis is the documentation that provides information about decision making in the network architecture and design processes. During the analysis process we are gathering data that can be used to determine which architectural and design decisions need to be made, details regarding each decision (including reasons for each decision), dependencies between decisions, and any background material used to arrive at these decisions.

## Overview of Analysis, Architecture, and Design Processes 23

Data from network analysis, along with any decisions made during the process, can be documented to form an *audit trail*, that is, the set of documents, data, and decisions, for the architecture and design. Audit trails are useful for describing and defending a particular network architecture or design. An audit trail helps address questions such as “Why did you choose that technology?” “Why doesn’t this new network perform as expected?” or “Why did this network cost so much?” Having documented your analysis of the existing network, the problems to be addressed, requirements for the new network, and all decisions made regarding that network, you will be able to answer questions at any time about your new network.

Decisions made regarding the network architecture and design need to be defensible from several perspectives: technical, in order to be able to address any technical challenges made against your architecture and design; budgetary, to ensure that network costs are within budget or to justify why a budget has been exceeded; schedule, to ensure that time frames for development, installation, testing, and operations are being met; and resources, such as personnel or equipment, to ensure that the customer has everything necessary to build and operate this network.

An audit trail is also useful as a historical document about the network. Over time, after the network is made operational, new network personnel can review this document to understand the logic behind the way the network was designed. Ideally, this document should be periodically reviewed, with new information added regarding changes to the network. Thus, an audit trail becomes a history for that network.

Experience shows that the set of documents, data, and decisions in an audit trail can be vital in making day-to-day tactical design decisions throughout the project. The investment in time at this phase of the project can save large amounts of time and resources later in the project.

The Web is a great tool to use in building an audit trail. Because an audit trail contains information about the old network, the new network, and decisions made about the new network, having this information easily accessible by those who use the network makes a lot of sense. Putting such information on internal Web pages allows easy access by everyone, and changes or additions to the audit trail can be seen immediately. Although there may be some information that your customer might not want everyone to view, such as the results of a risk analysis, most information usually can be accessible to everyone. For information that is restricted (need-to-know), hidden and password-protected Web pages can be used. Of course, when putting sensitive information at a common location, such as a Web site, sufficient security from outside attacks (hackers) should be provided.

## 24 CHAPTER 1 Introduction

An audit trail can be developed using standard word processing and spreadsheet software tools, and software tools specialized for this purpose are available. Problem statements, objectives, requirements, decisions, and all background data are entered into the audit trail, and all information is time stamped. Examples of audit trail information are presented later in this book.

### 1.4.5 Model for Network Analysis, Architecture, and Design

Networking traditionally has had little or no basis in analysis or architectural development, with designers often relying on technologies that are either most familiar to them or that are new, popular, or suggested by vendors and/or consultants. There are serious problems with this traditional approach. In particular, decisions may be made without the due diligence of analysis or architectural development, and such decisions, especially those made during the early phases of the project, are uninformed.

As a result, there may not be an audit trail for the architecture and design; and therefore, the architecture and design may not be defensible. In addition, such an architecture/design may lack consistency in its technological approach. Lacking data from analysis and architecture, we may not have a basis for making technology comparisons and trade-offs. And most importantly, without the proper requirements gathering and analysis, we cannot be sure if our network will meet the needs of its users. Therefore, network analysis, architecture, and design are fundamental to the development of a network.

Network analysis, architecture, and design are similar to other engineering processes in that they address the following areas:

- Defining the problems to be addressed
- Establishing and managing customer expectations
- Monitoring the existing network, system, and its environment
- Analyzing data
- Developing a set of options to solve problems
- Evaluating and optimizing options based on various trade-offs
- Selecting one or more options
- Planning the implementation



## Overview of Analysis, Architecture, and Design Processes 25

Defining the problems to be addressed should entail a quick evaluation of the environment and project—in essence performing a sanity check on the task at hand, as well as determining the size and scope of the problems, determining that you are working on the right problems, and checking the levels of difficulty anticipated in the technologies, potential architectures and designs, administration, management, and politics in that environment. As you size up the problems faced in this project, you should begin to estimate the level of resources needed (e.g., budget, schedule, personnel). You should also develop your own view of the problems affecting that environment. You may find that, from your analysis of the situation, your definition of the problems may differ from the customer's definition. Depending on how far apart your definitions are, you may need to adjust your customer's expectations about the project.

---

### Example 1.5.

Once, in performing an analysis on a customer's metropolitan-area network (MAN), I realized that the problem was not what the customers thought. They thought that the technology chosen at that time, switched multimegabit data service (SMDS), and the routing protocol (OSPF) were not working properly together. However, the problem actually was that the network personnel had forgotten to connect any of their LANs to the MAN. Of course, when they ran tests from one LAN to another, no data were being passed. It was an easy problem to fix, but a lot of work was spent changing the customer's view on the problem and expectations of what needed to be done. The customer originally wanted to change vendors for the routing equipment and replace the SMDS service. Eventually, they were convinced that the equipment and service were fine and that the problem was internal to the organization.

Although SMDS is not widely available anymore, its behavior as a non-broadcast multiple-access (NBMA) technology is similar to other currently available technologies.

---

An early part of every project is determining what your customer's expectations are and adjusting these expectations accordingly. The idea here is not to give customers false expectations or to let them have unrealistic expectations, because this will lead to difficulties later in the project. Instead, the goal is to provide an accurate and realistic view of the technical problems in the network and what it will take to solve them. Customers' expectations will likely focus on budget, schedule, and personnel but may also include their opinions about technologies and methodologies. And, at times, politics become embedded within the project and must be dealt with. The key here is to separate technical from nontechnical issues and to focus on technical and resource issues.

## 26 CHAPTER 1 Introduction

Part of determining customers' expectations means understanding what customers want to do with their network. This may involve understanding the customer's business model and operations. In addition, the customer may expect to have significant input into the development of the network. As you set the customer's expectations, you may need to establish the lines of communication between the network architecture/design group, management, users, and network staff.

Having established what the customer's expectations are, you may need to adjust and manage these expectations. This can be done by identifying trade-offs and options for resources, technologies, architectures, and designs and then presenting and discussing trade-offs and options with your customer. Bringing customers into the process and working with them to make critical decisions about the network will help them become comfortable with the process.

If an existing network is part of this project, monitoring this network, as well as other parts of the system and its environment, can provide valuable information about the current behavior of users, applications, and devices and their requirements for the new network. Monitoring can also validate your and your customer's definitions of the problems with the existing network. When it is possible to monitor the network, you will need to determine what data you want to collect (based on what you want to accomplish with the data), any strategic places in the network where you want to collect this data, and the frequency and duration of data collection.

At this point in the process you should have several sets of information with which you can begin your network analysis. You may have historical data from network management; data captured during monitoring; requirements gathered from users, staff, and management; the customer's definition of the problem; and your definition. All of these data are used in the network analysis, of which there are three parts: requirements or needs analysis, flow analysis, and a risk (security) analysis. Information in these analyses can be placed on the customer's internal Web page, as mentioned earlier, although some information (e.g., the results of the risk analysis) may have to be kept private.

Results of the network analysis are used in the architecture and design processes, where sets of options are developed, including potential architectures, designs, topologies, technologies, hardware, software, protocols, and services.

These sets of options are then evaluated to determine the optimal solutions for the problems. Criteria need to be developed throughout the analysis, architecture, and design processes in order to evaluate these options. Along with these criteria, you will use the results of the network analysis, including requirements, trade-offs, and dependencies between options.

Having selected one or more options, you can complete the network architecture and design and prepare for implementation. At this point you may consider developing a project plan to determine schedule, budget, and resources, as well as major and minor milestones, checkpoints, and reviews.

## 1.5 A Systems Methodology

We begin the network analysis process with a discussion of the systems methodology approach to networking. Applying a systems methodology to network analysis, architecture, and design is a relatively new approach, particularly in the Internet Protocol (IP) world. *Systems methodology* (as applied to networking) means viewing the network that you are architecting and designing, along with a subset of its environment (everything that the network interacts with or impacts), as a system. Associated with this system are sets of services (levels of performance and function) that are offered by the network to the rest of the system. This approach considers the network as part of a larger system, with interactions and dependencies between the network and its users, applications, and devices. As you will see, the systems methodology reveals interesting concepts which are used throughout this book.


One of the fundamental concepts of the systems methodology is that network architectures and designs take into account the services that each network will provide and support. This reflects the growing sophistication of networks, which have evolved from providing basic connectivity and packet-forwarding performance to being a platform for various services. As discussed earlier, we are currently at the stage in network evolution where services are important to the success of many networks (third-generation networks). Some examples of third-generation networks are service-provider networks that support multiple levels of performance and pricing to their customers, content-distribution networks that specialize in high-performance transport, and enterprise networks that incorporate and apply billing and usage models to their customers.

When a network is viewed as part of a system that provides services, the systems methodology works quite well for a variety of networks, from small and simple to large, complex networks. It helps in determining, defining, and describing the important characteristics and capabilities of your network.

## 1.6 System Description

A *system* is a set of components that work together to support or provide connectivity, communications, and services to users of the system. Generically speaking,

## 28 CHAPTER 1 Introduction



**FIGURE 1.17** Generic Components of a System

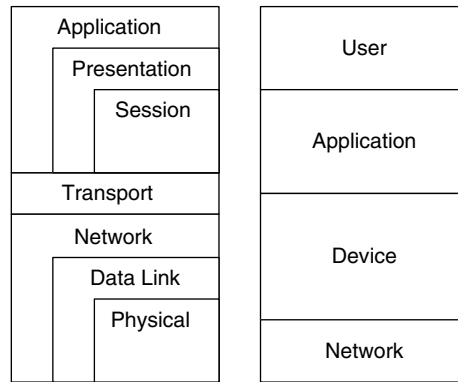
components of the system include users, applications, devices, and networks. Although users of the system can be considered to be outside the system, they also have requirements that include them as part of the system. Throughout this book we include users as part of the system. Figure 1.17 shows how these components are connected within the system.

Figure 1.17 shows the generic components of a system. These components can be subdivided, if necessary, to focus on a particular part of the system. For example, users in a corporate network could be further described as network and computer support personnel, as well as developers and customers of that corporation's product. In a similar sense, applications may be specific to a particular user, customer or group, generic to a customer base, or generic across the entire network.


If we were to compare this view of a system with the open system interconnect (OSI) protocol model, it would look like Figure 1.18. Note that, in this comparison, some of the OSI layers are modified. This is to show that there may be multiple protocol layers operating at one of the system levels. For example, the OSI physical, data link, and network layers may be present at the device level and may also be present multiple times at the network level (e.g., at switches and routers throughout the network).

Figure 1.19 shows that devices can be subdivided by class to show specialized functions, such as storage, computing, or application servers, or an individual device may be subdivided to show its operating system (OS), device drivers, peripheral hardware, or application programming interface (API).

All of these components work together to provide connectivity and communications across the network, among users, applications, and devices. The connectivity and communications can be tailored to meet the specific needs of users and



**FIGURE 1.18** Comparison of OSI Layers to System Levels




**FIGURE 1.19** Device Component Separated into Constituents

applications, such as real-time delivery of voice or streaming media, best-effort delivery of noninteractive data, or reliable delivery of mission-critical data.

The degree of granularity used to describe system components is a trade-off between the amount of detail and accuracy you want in the description and how much time and effort you are willing to put into it. If you are the network architect responsible for a corporate network, you may be able to invest time and resources into developing a detailed description of your system’s components, whereas a consultant or vendor’s design engineer may have little time and resources to spend on such a description. It is important to note, however, that even a small amount of time invested here will pay dividends later in the analysis process.

30 CHAPTER 1 Introduction




**FIGURE 1.20** Traditional View of a System

The traditional view of a system focused on the network providing connectivity between devices (Figure 1.20) and typically did not consider the users or applications.

This traditional view of a system is not complete enough for today’s networks. In particular, we need to include users and their applications in the system description. Experience shows that this degree of descriptiveness in the set (users, applications, devices, and networks) is usually sufficient to provide a complete and accurate description of most general-access systems, yet not so large as to be overwhelming to the network architect. (*General-access* is a term to describe common access of users to applications, computing, and storage resources across a network.) Within this set, users represent the end users, or customers, of the system. These end users are the final recipients of the network services supported by the system.

One reason for identifying components of the system is to understand how these components interface with one another across component boundaries. By defining what the components of the system are, we are also setting what is to be expected across each interface. For example, using the standard set (users, applications, devices, and networks), Figure 1.21 shows potential interfaces.



**FIGURE 1.21** A Generic System with Interfaces Added

Although the description of the system given in Figure 1.21 is usually satisfactory for the start of most network architectures, there may be times when you want to describe more components or have more defined interfaces. For example, the device-network interface may be simple or complex, depending on what you are trying to accomplish. For a network that will be providing simple connectivity, the network-device interface may be a standard LAN interface (e.g., 100BaseT Ethernet) without any prioritization or virtual LAN (VLAN 802.1p/q) tagging. For a network that provides more than simple connectivity, such as quality of service, the device-network interface may be more closely coupled with the device or application. This may be accomplished by using drivers that bypass portions of the protocol stack and APIs that can interpret application performance requirements.

Although the system description is an attempt to identify components across the entire system, we need to recognize that most systems are not completely homogeneous and that components may change in various parts of the system. This usually occurs in parts of the system that perform specific functions, such as a device-specific network (e.g., a video distribution network) or a storage-area network (SAN). For example, although an SAN may be described as the set (users, applications, devices, and networks), users may be other devices in the system, and the only application may be for system storage and archival.

## 1.7 Service Description

The concept of network services in this book builds upon the services work from the Internet Engineering Task Force (IETF). This organization has been developing service descriptions for IP networks. In general, they see network services as sets of network capabilities that can be configured and managed within the network and between networks. We apply this concept to network analysis, architecture, and design, integrating services throughout the entire system. This will help you take advantage of the services concept by analyzing, architecting, and designing based on services, and it will also prepare you for the near future, when services will be configurable and manageable within the network.

*Network services*, or services, are defined here as levels of performance and function in the network. We can look at this from two perspectives: as services being offered by the network to the rest of the system (the devices, applications, and users) or as sets of requirements from the network that are expected by the users, applications, or devices. Levels of performance are described by the performance characteristics capacity, delay, and RMA (reliability, maintainability,

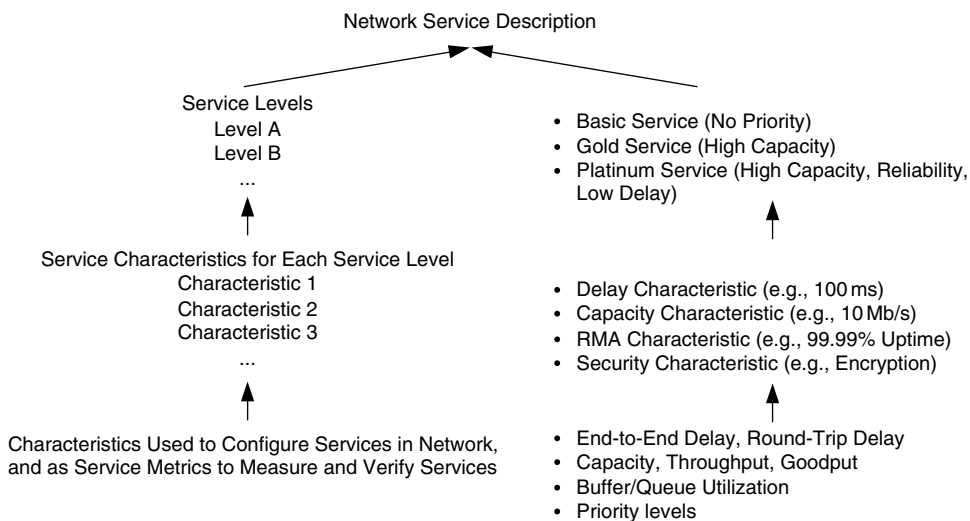
## 32 CHAPTER 1 Introduction

and availability), whereas functions are described as security, accounting, billing, scheduling, and management (and others). This is described in more detail in the next section.

It is important to note that the concept of services used in this book is based on what networks can deliver to the system. Thus, it is not to be confused with services that other parts of the system (e.g., applications) can deliver to each other (e.g., graphics rendering services). When the term *service* is used in this book, it is in reference to network service.

Network services in most of today’s networks are based on best-effort (unpredictable and unreliable) delivery. In addition to best-effort delivery, we examine some new types of services, including high-performance, predictable (stochastic or probabilistic), and guaranteed services. These new services require some different ways of looking at networks, and you will see how to incorporate such services into your architecture and design. We also look at single-tier and multiple-tier performance in the network, and show how to distinguish between them and how they relate to best-effort, predictable, and guaranteed services.

Network services are hierarchical, and individual service characteristics can be grouped together to form higher-level descriptions of a service, as shown in Figure 1.22.



**FIGURE 1.22** Grouping Characteristics into Service Levels and Descriptions



## 1.8 Service Characteristics

One of the goals of network analysis is to be able to characterize services so that they can be designed into the network and purchased from vendors and service providers (e.g., via requests for information [RFI], quote [RFQ], or proposal [RFP], documents used in the procurement process). *Service characteristics* are individual network performance and functional parameters that are used to describe services. These services are offered by the network to the system (the *service offering*) or are requested from the network by users, applications, or devices (the *service request*). Characteristics of services that are requested from the network can also be considered requirements for that network.

Examples of service characteristics range from estimates of capacity requirements based on anecdotal or qualitative information about the network to elaborate listings of various capacity, delay, and RMA requirements, per user, application, and/or device, along with requirements for security, manageability, usability, flexibility, and others.

---

### Example 1.6.

Examples of service characteristics are:

- Defining a security or privacy level for a group of users or an organization
  - Providing 1.5 Mb/s peak capacity to a remote user
  - Guaranteeing a maximum round-trip delay of 100 ms to servers in a server farm
- 

Such requirements are useful in determining the need of the system for services, in providing input to the network architecture and design, and in configuring services in network devices (e.g., routers, switches, device operating systems). Measurements of these characteristics in the network to monitor, verify, and manage services are called *service metrics*. In this book we focus on developing service requirements for the network and using those characteristics to configure, monitor, and verify services within the network.

For services to be useful and effective, they must be described and provisioned end-to-end at all network components between well-defined demarcation points. “End-to-end” does not necessarily mean only from one user’s device to another user’s device. It may be defined between networks, from users to servers, or between specialized devices (Figure 1.23). When services are not provisioned end-to-end, some components may not be capable of supporting them, and thus the services

34 CHAPTER 1 Introduction




FIGURE 1.23 Example Demarcations Points to Describe End-to-End within a Network

will fail. The demarcation points determine where end-to-end is in the network. Determining these demarcation points is an important part of describing a service.

Services also need to be configurable, measurable, and verifiable within the system. This is necessary to ensure that end users, applications, and devices are getting the services they have requested (and possibly have been paying for), and this leads to accounting and billing for system (including network) resources. You will see how service metrics can be used to measure and verify services and their characteristics.

Services are also likely to be hierarchical within the system, with different service types and mechanisms applied at each layer in the hierarchy. For example, Figure 1.24 shows a quality-of-service (QoS) hierarchy that focuses on bulk traffic

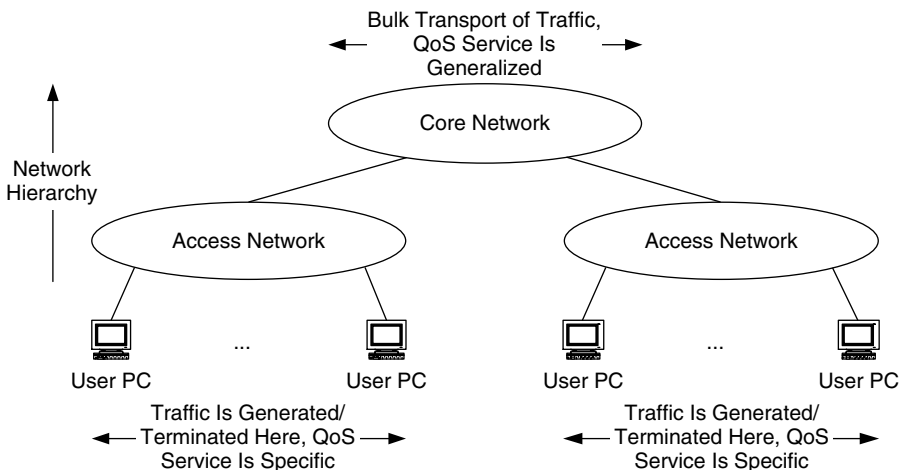


FIGURE 1.24 An Example of Service Hierarchy within a Network


transport in the core of the network while placing specific services at the access network close to the end users, applications, and devices.

### 1.8.1 Service Levels

Service characteristics can be grouped together to form one or more *service levels* for the network. This is done to make service provisioning easier in that you can configure, manage, account, and bill for a group of service characteristics (service level) instead of a number of individual characteristics. For example, a service level (e.g., premium) may combine capacity (e.g., 1.5 Mb/s) and reliability (as 99.99% uptime). Service levels are also helpful in billing and accounting. This is a service-provider view of the network, where services are offered to customers (users) for a fee. This view of networking is becoming more popular in enterprise networks, displacing the view of networks as purely the infrastructure of cost centers.

There are many ways to describe service levels, including frame relay committed information rates (CIRs), which are levels of capacity; classes of service (CoSs), which combine delay and capacity characteristics; and IP types of service (ToSs) and qualities of service (QoSs), which prioritize traffic for traffic conditioning functions, which are described in the performance architecture (see Chapter 8). There can also be combinations of the aforementioned mechanisms, as well as custom service levels, based on groups of individual service characteristics. These combinations depend on which network technology, protocol, mechanism, or combination is providing the service.

In Figure 1.25 service offerings, requests, and metrics are shown applied to the system. In this example a demarcation of services is shown between the device and network components. Depending on the service requirement or characteristic, however, demarcation may also be between the device and application components.



**FIGURE 1.25** Service Requests, Offerings, and Metrics

## 36 CHAPTER 1 Introduction

---

### Example 1.7.

A request came from a customer that each building should have Fast Ethernet (FE) capacity to the rest of the network. As part of the requirements analysis, this request became a requirement for 100 Mb/s peak capacity from the users in each building. This *service request* was then matched in the requirements and design processes by a technology choice that could meet or exceed the request. In this case FE was chosen as the technology, and the service offering was 100 Mb/s to each building. *Service metrics* were then added, consisting of measuring the FE connections from the IP switch or router at each building to the backbone.

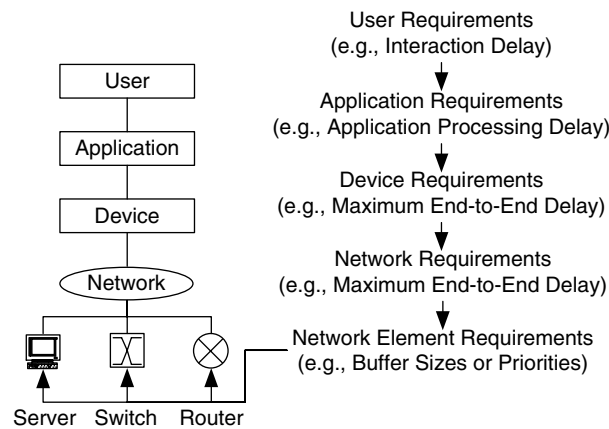
---

Services and service levels can be distinguished by their degrees of predictability or determinism. In the next section we discuss best-effort delivery, which is not predictable or deterministic, as well as predictable and guaranteed services. Services and service levels are also distinguished by their degrees of performance. You will see how the service performance characteristics capacity, delay, and RMA are used to describe services and service levels.

### 1.8.2 System Components and Network Services

Network services are derived from requirements at each of the components in the system. They are end-to-end (between end points that you define) within the system, describing what is expected at each component. Service requirements for the network we are building are derived from each component. There can be user requirements, application requirements, device requirements, and (existing) network requirements. Because we are building the network component, any requirements from the network component come from existing networks that the new network will incorporate or connect to.

Component requirements are added one to another, being refined and expanded as the network comes closer to being realized. User requirements, which are the most subjective and general, are refined and expanded by requirements from the application component, which are in turn refined and expanded by the device and network components. Thus, requirements filter down from user to application to device to network, resulting in a set of specific requirements that can be configured and managed in the network devices themselves. This results in a service offering that is end-to-end, consisting of service characteristics that are configured in each network device in the path (e.g., routers, switches, hubs). As in Figure 1.26, service characteristics are configured and managed within each



**FIGURE 1.26** Requirements Flow Down Components, from User to Network

element and at interfaces between elements. These services are the most specific of all and have the smallest scope (typically a single network device).

Defining network services and service metrics helps keep the system functioning and can provide extra value or convenience to users and their applications. By defining service metrics we are determining what we will be measuring in the network, which will help us in network monitoring and management.

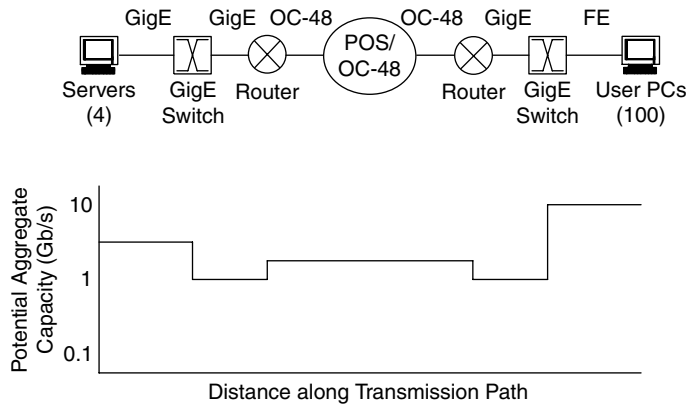
Recall that network services are sets of performance and function, so requirements may also include functions of one of the components. Examples of functions include network monitoring and management, security, and accounting. Services such as these must be considered an integral part of the network architecture and design. In this book, security (and privacy) and network management each have their own architectures. This may seem obvious, but traditionally, services such as security and network management have been afterthoughts in architecture and design, often completely forgotten in the architecture until problems arise.

---

### Example 1.8.

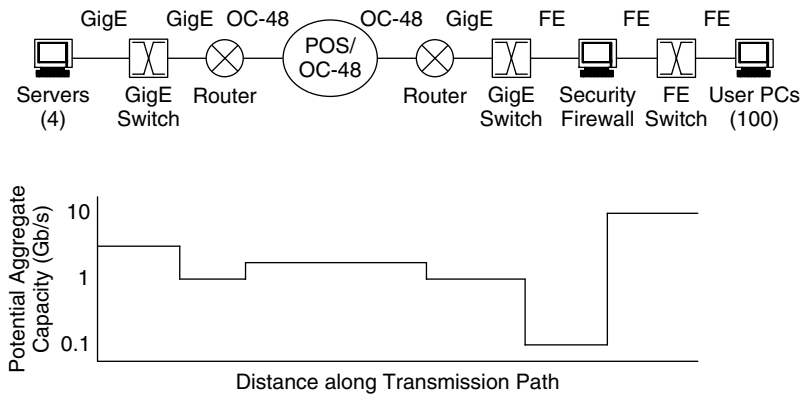
The network path shown in Figure 1.27 was designed to optimize performance between users and their servers. The graph at the bottom of the figure is an estimate of the expected aggregate capacity at each segment of the path. In this network a packet over SONET (POS) link at the OC-48 level (2.544 Gb/s) connects two routers, which then connect to Gigabit Ethernet (GigE) switches.

38 CHAPTER 1 Introduction



**FIGURE 1.27** The Capacity at Each Point in the Transmission Path before the Addition of a Security Firewall

After it was implemented, a security firewall was added at the users’ LAN (with FE interfaces), without it being considered part of the original analysis, architecture, or design. The result was that the firewall changed the capacity characteristics across the path by reducing throughput between the user PCs and the GigE switch, as shown in Figure 1.28.



**FIGURE 1.28** The Capacity at Each Point in the Transmission Path after the Addition of a Security Firewall

One of our architectural and design goals is to identify such performance bottlenecks before the network is implemented. By considering security, network management, services, and routing and addressing in the analysis process, we are

much more likely to understand their behavior and effect on each other and the network. We are therefore able to architect the network to accommodate their requirements and interoperability.

When service characteristics apply to individual network devices, such as routers, switches, data service units (DSUs), and so on, some of these characteristics may be vendor specific. In this book we focus on those characteristics that are part of public standards and not vendor specific.

It is important to note that although standards-based characteristics are “standardized” on the basis of having their descriptions publicly available (e.g., via an IETF RFC), sanctioned by an organization recognized by the networking community, or generally accepted and used as a *de facto standard*, the implementation of characteristics is open to interpretation and often varies across vendors and vendor platforms.

### 1.8.3 Service Requests and Requirements

Service requests and requirements are, in part, distinguished by the degree of predictability needed from the service by the user, application, or device making the request. Based on their predictability, service requests are categorized as best effort, predictable, or guaranteed. Service requests and requirements can also be appropriate for single- or multiple-tier performance for a network.

*Best-effort service* means that there is no control over how the network will satisfy the service request—that there are no guarantees associated with this service. Such requests indicate that the rest of the system (users, applications, and devices) will need to adapt to the state of the network at any given time. Thus, the expected service for such requests will be both unpredictable and unreliable, with variable performance across a range of values (from the network being unavailable to the lowest common denominator of performance across all of the technologies in the end-to-end path). Such service requests either have no specific performance requirements for the network or are based solely on estimates of capacity. When requirements are nonspecific, network performance cannot be tuned to satisfy any particular user, application, or device requirement.

*Guaranteed service* is the opposite of best-effort service. Where best-effort service is unpredictable and unreliable, guaranteed service must be predictable and reliable to such a degree that, when service is not available, the system is held accountable. A guaranteed service implies a contract between the user and provider. For periods when the contract is broken (e.g., when the service is not available), the provider must account for the loss of service and, possibly, appropriately compensate the user.

## 40 CHAPTER 1 Introduction

With best-effort and guaranteed services at opposite ends of the service spectrum, many services fall somewhere between. These are *predictable services*, which require some degree of predictability (more than best effort) yet do not require the accountability of a guaranteed service.

Predictable and guaranteed service requests are based on some a priori knowledge of and control over the state of the system. Such requests may require that the service either operates predictably or is bounded. Therefore, such services must have a clear set of requirements. For the network to provision resources to support a predictable or guaranteed service, the service requirements of that request must be configurable, measurable, and verifiable. This is where service requests, offerings, and metrics are applied.

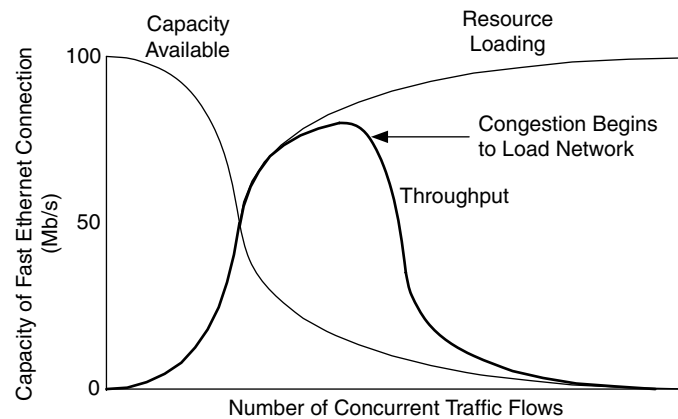
Note that there are times when a service can be best effort, predictable, or guaranteed, depending on how it is interpreted. Therefore, it is important to understand the need for a good set of requirements because these will help determine the types of services to plan for. Also, although the term *predictable* lies in a gray area between best effort and guaranteed, it is the type of service most likely to be served by most performance mechanisms, as we see in Chapter 8.

For example, suppose a device requires capacity (bandwidth) between 4 and 10 Mb/s. There must be a way to communicate this request across the network, a way to measure and/or derive the level of resources needed to support this request, a way to determine whether the required resources are available, and a method to control the information flow and network resources to keep this service between 4 and 10 Mb/s.

Capacity (or bandwidth) is a finite resource within a network. For example, the performance of a 100 Mb/s FE connection between two routers is bounded by that technology. If we were to look at the traffic flows across that 100 Mb/s connection, we would see that, for a common best-effort service, capacity would be distributed across all of the traffic flows. As more flows were added to that connection, the resources would be spread out until, at some point, congestion occurs. Congestion would disrupt the traffic flows across that connection, affecting the protocols and applications for each flow. What is key here is that, in terms of resource allocation, all traffic flows have some access to resources.

This is shown in Figure 1.29. In this figure available capacity (dashed curve) decreases as the number of traffic flows increases. Correspondingly, the loading on the network (solid curve) from all of the traffic flows increases. However, at some point congestion affects the amount of user traffic being carried by the connection, and throughput of the connection (heavy curve) drops. As congestion interferes with the end-to-end transport of traffic, some protocols (e.g., TCP) will retransmit





**FIGURE 1.29** The Performance of a Fast Ethernet Connection under Best-Effort Conditions

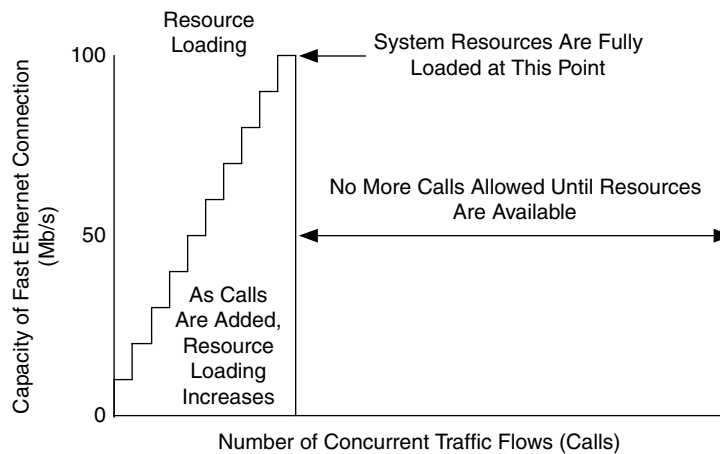
traffic. The difference between the loading and the throughput curves is due to retransmissions. This is undesirable, for while the connection is being loaded, only a percentage of that loading are successfully delivered to destinations. At some point all of the traffic on that connection could be due to retransmissions and throughput would approach zero. This approach is used in best-effort networks.

In contrast, consider a traditional telephony network. Calls are made on this network, and resources are allocated to each call. As more calls are added to the network, at the point where all of the resources have been allocated, additional calls are refused. The exiting calls on the network may suffer no performance degradation, but no new calls are allowed until resources are available. *Call admission control* (CAC) is a mechanism to limit the number of calls on a network, thereby controlling the allocation of resources.

This is shown in Figure 1.30. Individual calls are shown in this figure, and each call is 10 Mb/s for simplicity. As each call is accepted, resources are allocated to it, so the availability drops and loading increases for each call. When the resources are exhausted, no further calls are permitted. Congestion is not a problem for the existing calls, and throughput is maximized. This approach is similar to a guaranteed service.

There is a trade-off between these two approaches to resource allocation in a network. Although a best-effort network allows access to as many traffic flows as possible, performance degradation across all of the traffic flows can occur. Admission control preserves resources for traffic flows that have already been allocated resources but will refuse additional traffic flows when resources are exhausted.

## 42 CHAPTER 1 Introduction



**FIGURE 1.30** The Performance of a Fast Ethernet Connection under CAC

In many networks both approaches (or a hybrid between them) are desired. For example, a voice over IP (VoIP) service, which provides a telephony service across a data network, requires some of the characteristics of CAC while operating over a best-effort network. Such hybrid approaches are discussed in detail in Chapter 8.

Service requests and requirements can also be low or high performance in terms of capacity, delay, and RMA. Low- and high-performance requirements depend on each particular network. A requirement is low or high performance relative to other requirements for that network. *Low performance* is an indicator that the service request or requirement's performance characteristics are less than a performance threshold determined for that network. Likewise, *high performance* is an indicator that the service request or requirement's performance characteristics are greater than a performance threshold determined for that network. Thus, in determining low and high performance for a network, we will develop one or more performance thresholds for that network. Multiple-tier performance indicates that there are multiple tiers of performance for that network. Single-tier performance requirements are roughly equivalent within a network.

Note that low and high performances are not described in terms of best-effort, predictable, or guaranteed service because they are independent of each other. *Best-effort*, *predictable*, and *guaranteed service* refer to the degree of predictability of a request or requirement, whereas *low* and *high performances* refer to a relative performance level for that request or requirement. For example, a network can be

entirely best effort (most current networks are), yet we can often distinguish low- and high-performance requirements for such a network. And when a network has low- and high-performance regions for capacity, delay, and RMA, there may be predictable or guaranteed requirements in either region.

By their nature, each service has its associated set of requirements. These requirements are based on the levels of performance and function desired by the user, application, or device requesting service. Performance requirements are described in terms of capacity, delay, and RMA, whereas functional requirements describe specific functions needed in the service, such as multicast, security, management, or accounting. We use requests for performance and function in developing the network architecture and design—for example, in describing the overall level of performance needed in the network.

As mentioned earlier, service performance requirements (capacity, delay, and RMA) can be grouped together, forming one or more service levels. For example, a service request may couple a specific capacity (e.g., 1.5 Mb/s) with a bound on end-to-end delay (e.g., 40 ms). At times, such service levels can be mapped to well-known service mechanisms such as frame relay CIR or IP ToS or QoS. Thus, service levels are a way to map performance and functional requirements to a well-known or standard network service offering. A properly specified service provides insight into which performance characteristics should be measured in the network to verify service delivery.

#### 1.8.4 Service Offerings

Service requests that are generated by users, applications, or devices are supported by services offered by the network. These service offerings (e.g., via frame relay CIR or IP ToS or QoS, mentioned in the previous section) are the network counterparts to user, application, and device requests for service.

Service offerings map to service requests and thus can also be categorized as best effort, predictable, or guaranteed. Best-effort service offerings are not predictable—they are based on the state of the network at any given time. There is little or no prior knowledge about available performance, and there is no control over the network at any time. Most networks today operate in best-effort mode. A good example of a network that offers best-effort service is the current Internet.

Best-effort service offerings are compatible with best-effort service requests. Neither the service offering nor the request assumes any knowledge about the state of or control over the network. The network offers whatever service is available at that time (typically just available bandwidth), and the rest of the system adapts the flow of information to the available service (e.g., via TCP flow control).

## 44 CHAPTER 1 Introduction

---

### Example 1.9.

An example of a best-effort service request and offering is a file transfer (e.g., using FTP) over the Internet. FTP uses TCP as its transport protocol, which adapts, via a sliding-window flow-control mechanism, to approximate the current state of the network it is operating across. Thus, the service requirement from FTP over TCP is best effort, and the corresponding service offering from the Internet is best effort. The result is that, when the FTP session is active, the performance characteristics of the network (Internet) and flow control (TCP windows) are constantly interacting and adapting, as well as contending with other application sessions for network resources. In addition, as part of TCP's service to the applications it supports, it provides error-free, reliable data transmission.

---

On the other hand, predictable and guaranteed service offerings have some degree of predictability or are bounded. To achieve this, there has to be some knowledge of the network, along with control over the network, in order to meet performance bounds or guarantees. Such services must be measurable and verifiable.

Just because a service is predictable or guaranteed does not necessarily imply that it is also high performance. Take, for example, the telephone network. It offers predictable service but low performance (in terms of capacity). To support voice conversations, this network must be able to support fairly strict delay and delay variation tolerances, even though the capacity per user session (telephone call) is relatively small, or low performance. What is well known from a telephony perspective is somewhat new in the current world of data networking. Support for strict delay and delay variation is one of the more challenging aspects of data network architecture and design.

Predictable and guaranteed service offerings should be compatible with their corresponding service requests. In each case, service performance requirements (capacity, delay, and RMA) in a service request are translated into the corresponding performance characteristics in the service offering.

---

### Example 1.10.

An example of a predictable service request and offering can be seen in a network designed to support real-time streams of telemetry data. An architectural/design goal for a network supporting real-time telemetry is the ability to specify end-to-end delay and have the network satisfy this delay request. A real-time telemetry stream should have an end-to-end delay requirement, and this requirement would form the basis for the service request. For example, this service request may be for an end-to-end delay of 25 ms, with a delay variation

of  $\pm 400 \mu\text{s}$ . This would form the request and the service level (i.e., a QoS level) that needs to be supported by the network. The network would then be architected and designed to support a QoS level of 25 ms end-to-end delay and a delay variation of  $\pm 400 \mu\text{s}$ . Delay and delay variation would then be measured and verified with service metrics in the system, perhaps by using common utilities, such as ping (a common utility for measuring round-trip delay) or TCPdump (a utility for capturing TCP information), or by using a custom application.

---

We use various methods to describe service performance requirements and characteristics within a network, including thresholds, bounds, and guarantees. We also show how to distinguish between high and low performance for each network project.

This approach does not mean that best-effort service is inherently low performance or that predictable or guaranteed services are high performance. Rather, it signifies that predictability in services is an important characteristic and is separate from performance. There are times when a network is best architected for best-effort service, and other times when best-effort, predictable, and guaranteed services are needed. We will see that when predictable or guaranteed services are required in the network, consideration for those requirements tends to drive the architecture and design in one direction, while consideration for best-effort service drives them in another direction. It is the combination of all services that helps make the architecture and design complete.

### 1.8.5 Service Metrics

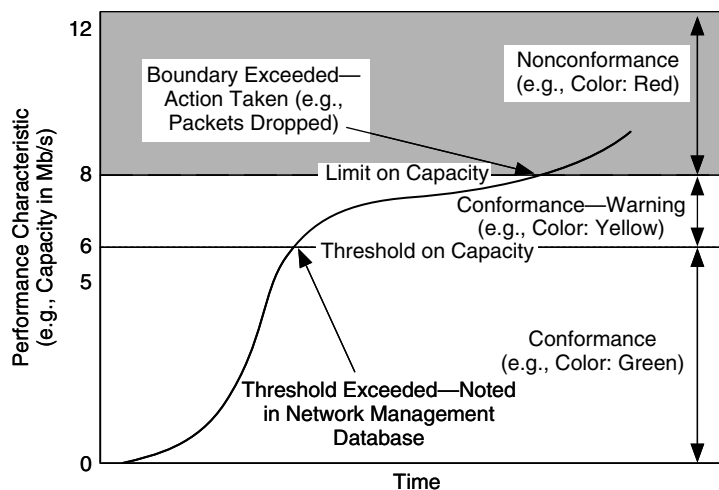
For service performance requirements and characteristics to be useful, they must be configurable, measurable, and verifiable within the system. Therefore, we will describe performance requirements and characteristics in terms of *service metrics*, which are intended to be configurable and measurable.

Because service metrics are meant to be measurable quantities, they can be used to measure thresholds and limits of service. Thresholds and limits are used to distinguish whether performance is in conformance (adheres to) or nonconformance (exceeds) with a service requirement. A *threshold* is a value for a performance characteristic that is a boundary between two regions of conformance and, when crossed in one or both directions, will generate an action. A *limit* is a boundary between conforming and nonconforming regions and is taken as an upper or lower limit for a performance characteristic. Crossing a limit is more serious than crossing a threshold, and the resulting action is usually more serious (e.g., dropping of packets to bring performance back to conformance).

## 46 CHAPTER 1 Introduction

For example, a threshold can be defined to distinguish between low and high performance for a particular service. Both low- and high-performance levels are conforming to the service, and the threshold is used to indicate when the boundary is crossed. This threshold can be measured and monitored in the network, triggering some action (e.g., a flashing red light on an administrator's console) when this threshold is crossed. An example of this might be in measuring the round-trip delay of a path. A threshold of  $N$  ms is applied to this measurement. If the round-trip times exceed  $N$  ms, an alert is generated at a network management station. We discuss this in greater detail in the chapter on network management architecture (Chapter 7).

In a similar fashion, limits can be created with service metrics to provide upper and lower boundaries on a measured quantity. When a limit is crossed, traffic is considered nonconforming (it exceeds the performance requirement), and action is taken to bring the traffic back into conformance (e.g., by delaying or dropping packets). Figure 1.31 shows how limits and thresholds may be applied in the system. In this figure, a threshold of 6 Mb/s is the boundary between low and high performance for a service requirement, and an upper limit of 8 Mb/s is the boundary between conformance and nonconformance for that service. When traffic crosses the 6 Mb/s threshold, a warning is sent to network management (with a color change from green to yellow). These notices can be used to do trend analysis on the network—for example, to determine when capacity needs to be



**FIGURE 1.31** Performance Limits and Thresholds

upgraded. When traffic crosses the 8 Mb/s limit, the network takes action to reduce the capacity used by that traffic flow and an alert is sent to network management (with a color change from yellow to red) until the capacity level drops below 8 Mb/s and is again conforming.

Thresholds and limits are useful applications of service metrics to understand and control performance levels in the network, in support of services.

## 1.9 Performance Characteristics

Services may include one or more of the performance characteristics we have mentioned so far in this chapter: capacity, delay, and RMA. Each characteristic is actually a label for a class of characteristics of that type. For example, the term *capacity* is used as a label for the class of characteristics that involves moving information from place to place, including bandwidth, throughput, goodput, and so forth. Similarly, delay is a label for the class of characteristics that includes end-to-end delay, round-trip delay, and delay variation. RMA is a label for the class of characteristics that includes reliability, maintainability, and availability. Thus, when the terms *capacity*, *delay*, and *RMA* are used in this book, you can use other terms from each class, depending on your network.

There are times when it makes more sense to describe capacity in terms of throughput—for example, when developing requirements for applications. Round-trip delay is commonly used as a measure for delay, although at times delay requirements are expressed in terms of one-way delay.

### 1.9.1 Capacity

*Capacity* is a measure of the system's ability to transfer information (voice, data, video, or combinations of these). Several terms are associated with capacity, such as *bandwidth*, *throughput*, or *goodput*. Although we use the generic term *capacity* throughout this book to reference this class of characteristics, you may choose to use another term in place of or along with capacity.

---

#### Example 1.11.

The bandwidth of a SONET OC-3c link is 155.52 Mb/s, which is three times the bandwidth of an OC-1 link (51.84 Mb/s). This bandwidth does not include data-link, network, or transport-layer protocol (e.g., SONET, IP, or transport control protocol/user

## 48 CHAPTER 1 Introduction

datagram protocol [TCP/UDP]) overhead or, in the case of wide-area networks, the loss in performance due to the bandwidth  $\times$  delay product in the network. When a network or element is performing at its theoretical capacity, it is said to be performing at *line rate*. When an OC-3c circuit was tested, values of realizable capacity (throughput) ranged from approximately 80 to 128 Mb/s (measurements taken at the transport [TCP] layer of the National Research and Education Network [NREN] and Numerical Aerodynamic Simulation [NAS] networks, NASA Ames Research Center, March 1996).

---

### 1.9.2 Delay

*Delay* is a measure of the time difference in the transmission of information across the system. In its most basic sense, delay is the time difference in transmitting a single unit of information (bit, byte, cell, frame, or packet) from source to destination. As with capacity, there are several ways to describe and measure delay. There are also various sources of delay, such as propagation, transmission, queuing, and processing. Delay may be measured in one direction (end-to-end) and both directions (round-trip). Both end-to-end and round-trip delay measurements are useful; however, only round-trip delays can be measured with the use of the practical and universally available utility *ping*.

Another measure of delay incorporates device and application processing, taking into account the time to complete a task. As the size of a task increases, the application processing times (and thus the response time of the system) also increase. This response time, termed here *latency*, may yield important information about the behavior of the application and the network. Latency can also be used to describe the response time of a network device, such as the latency through a switch or router. In this case the processing time is of that switch or router.

Delay variation, which is the change in delay over time, is an important characteristic for applications and traffic flows that require constant delay. For example, real-time and near-real-time applications often require strict delay variation. Delay variation is also known as *jitter*.

Together, delay (end-to-end and round-trip), latency, and delay variation help describe network behavior.

### 1.9.3 RMA

*RMA* refers to reliability, maintainability, and availability. *Reliability* is a statistical indicator of the frequency of failure of the network and its components and represents the unscheduled outages of service. It is important to keep in mind



that only failures that prevent the system from performing its mission, or mission-critical failures (more on this in Chapter 2), are generally considered in this analysis. Failures of components that have no effect on the mission, at least when they fail, are not considered in these calculations. Failure of a standby component needs tending to but is not a mission-critical failure.

Reliability also requires some degree of predictable behavior. For a service to be considered reliable, the delivery of information must occur within well-known time boundaries. When delivery times vary greatly, users lose confidence in the timely delivery of information. In this sense the term *reliability* can be coupled with *confidence* in that it describes how users have confidence that the network and system will meet their requirements.

A parallel can be seen with the airline industry. Passengers (users) of the airline system expect accurate delivery of information (in this case the passengers themselves) to the destination. Losing or misplacing passengers is unacceptable. In addition, predictable delivery is also expected. Passengers expect flights to depart and arrive within reasonable time boundaries. When these boundaries are crossed, passengers are likely to use a different airline or not fly at all. Similarly, when an application is being used, the user expects a reasonable response time from the application, which is dependent on the timely delivery of information across the system.

Along with reliability is maintainability. *Maintainability* is a statistical measure of the time to restore the system to fully operational status after it has experienced a fault. This is generally expressed as a mean-time-to-repair (MTTR). Repairing a system failure consists of several stages: detection; isolation of the failure to a component that can be replaced; the time required to deliver the necessary parts to the location of the failed component (logistics time); and the time to actually replace the component, test it, and restore full service. MTTR usually assumes the logistics time is zero; this is an assumption, which is invalid if a component must be replaced to restore service but takes days to obtain.

To fully describe this performance class, we add availability to reliability and maintainability. *Availability* (also known as operational availability) is the relationship between the frequency of mission-critical failures and the time to restore service. This is defined as the mean time between mission-critical failures (or mean time between failures) divided by the sum of mean time to repair and mean time between mission-critical failures or mean time between failures. These relationships are shown in the following equation, where  $A$  is availability.


$$A = (\text{MTBCF})/(\text{MTBCF} + \text{MTTR}) \text{ or } A = (\text{MTBF})/(\text{MTBF} + \text{MTTR})$$

## 50 CHAPTER 1 Introduction


Capacity, delay, and RMA are dependent on each other. For example, the emphasis of a network design may be to bound delay: A system supporting point-of-sale transactions may need to guarantee delivery of customer information and completion of the transaction within 15 seconds (where the network delay is on the order of 100s of ms); a Web application can have similar requirements. However, in a computation-intensive application we may be able to optimize the system by buffering data during periods of computing. In this case, delay may not be as important as a guarantee of eventual delivery. On the other hand, a system supporting visualization of real-time banking transactions may require a round-trip delay of less than 40 ms, with a delay variation of less than 500  $\mu$ s. If these delay boundaries are exceeded, the visualization task fails for that application, forcing the system to use other techniques.

### 1.9.4 Performance Envelopes

Performance requirements can be combined to describe a performance range for the system. A *performance envelope* is a combination of two or more performance requirements, with thresholds and upper and/or lower limits for each. Within this envelope, levels of application, device, and/or network performance requirements are plotted. Figures 1.32 and 1.33 show two such envelopes. The performance



**FIGURE 1.32** An Example of a 2D Performance Envelope



**FIGURE 1.33** An Example of a 3D Performance Envelope

envelope in Figure 1.32 consists of capacity, in terms of data sizes transferred across the network, and end-to-end delay. In this figure, delay is shown as  $1/\text{delay}$  for consistency.

Figure 1.33 is a 3D performance envelope, showing capacity, delay, and RMA. This envelope also describes two regions of performance, low and high performance, which are functions of the limits and thresholds for capacity, delay, and RMA.

Performance envelopes such as these are useful for visualizing the regions of delay, capacity, and RMA in which the network will be expected to operate based on requirements developed for that network. In Chapter 2 we discuss how requirements are developed for a network.

## 1.10 Network Supportability

The ability of the customer to sustain the required level of performance (that architected and designed into the network) over the entire life cycle of the network is an area of networking that is often neglected. It is a mistake to assume that a successful network architecture and design meet the requirements only on the day it is delivered to the customer and that future requirements are the responsibility of the customer.

Experience indicates operations and support constitute 80% of the life-cycle costs of a system, whereas development, acquisition, and installation represent only

## 52 CHAPTER 1 Introduction

20%. Good network architects/designers take into account the major factors that affect operability and supportability as they make their decisions. Knowledgeable customers insist that they understand the operations and support implications of a network architecture and design. At times, such issues may be of more concern than the feasibility of a new technology.

The postimplementation phases of a network's life cycle can be broken into three elements: operations, maintenance, and human knowledge. The operations element focuses on ensuring that the network and system are properly operated and managed and that any required maintenance actions are identified. The maintenance element focuses on preventive and corrective maintenance and the parts, tools, plans, and procedures for accomplishing these functions. The human knowledge element is the set of documentation, training, and skilled personnel required to operate and maintain the network and system. Design decisions affect each of these factors and have a direct impact on the ability of the customer to sustain the high level of service originally realized upon implementation of the network.

Failure to consider supportability in the analysis, architecture, and design processes has a number of serious consequences. First, a smart customer, when faced with a network architecture/design that obviously cannot be operated or maintained by his or her organization, will reject the network project or refuse to pay for it. Second, a customer who accepts the architecture/design and subsequent implementation will have inadequate resources to respond to network and system outages, experience unacceptable performance after a period of time, and may suffer adverse effects in his or her operation or business (e.g., a loss of their customers or revenue). Other customers will be highly dissatisfied with their network and either require the architect/designer to return and repair the network by providing adequate materials to sustain its required performance level or will prematurely replace it. None of these cases reflects positively on the network architect/designer or implementation team and can lead to finger pointing that can be more painful than any acceptance test.

Key characteristics of a network architecture and design that affect the postimplementation costs include:

- Network and system reliability
- Network and system maintainability
- Training of the operators to stay within operational constraints
- Quality of the staff required to perform maintenance actions

Some examples of key network architecture/design decisions that affect these characteristics include:

- Degree of diversity of critical-path components in network architecture/design
- Quality of network components selected for installation
- Location and accessibility of components requiring frequent maintenance
- Implementation of built-in test equipment and monitoring techniques

Supportability must be considered throughout the life cycle of the network. An accurate assessment of the requirements for continuous service at full performance level must be included in the requirements analysis process, along with a statement of specific, measurable requirements. During the architecture and design processes, trade-offs must take into account the impact of supportability, and the concept of operations must be formulated. Last, during implementation, two major tasks must be accomplished to ensure supportability:

1. Conformance to the network architecture and design must be validated and nonconformance corrected or (at least) documented to ensure that performance is adequate and that maintenance can be performed.
2. Operations and maintenance personnel must understand and be trained in the technologies that are being deployed, including how to operate the network and system properly, when to perform maintenance, and how to most quickly restore service in the event of a fault.

A detailed discussion of how supportability fits into the overall architecture and design processes is provided in Chapter 2.

## 1.11 Conclusion

In this chapter you learned definitions of network analysis, architecture, and design; the importance of network analysis in understanding the system and providing a defensible architecture and design; and the model for the network analysis, architecture, and design processes.

You have also learned that networks are not independent entities but rather a part of the system and that the delivery of network services is a goal of the system. Network services consist of performance and function and are offered to users,

## 54 CHAPTER 1 Introduction

applications, and devices so that they can accomplish their work on the system. In order to architect and design a network to support services, you need to know what they are, how they work together, and how to characterize them. Once you do this, you will have a broad view of what the network will need to support, which you can take to the next levels of detail as you proceed with the network analysis.

By describing the system as a set of components (e.g., user, application, device, network), you can apply interfaces between these components to help understand the relationships, inputs, and outputs between each of the components.

You have also learned about different types of services, from best-effort, unpredictable, and unreliable service to predictable, bounded, and somewhat predictable service, to guaranteed services with accountability.

To go to a level deeper in the discussion about services, we considered the service performance characteristics capacity, delay, and RMA (reliability, maintainability, and availability). These characteristics are useful only if we can measure and verify their values in the system. We discussed these values, as well as service metrics, thresholds, and boundaries. We learned that performance characteristics can be combined into a performance envelope.

Having thought about systems, services, and their characteristics, we are now ready to quantify what we want from our networks. To do this, we first need to gather, analyze, and understand the requirements from the system. This is *requirements analysis*, the next step in the network analysis process.

### 1.12 Exercises

1. In Example 1.3, an analogy was drawn between a network's architecture and design and a home's architecture and design. Provide a similar analogy, using a computer's architecture and design.
2. Hierarchy and interconnectivity are a fundamental trade-off in networks. Given the network hierarchy shown in Figure 1.30, with costs assigned to each link, show how interconnectivity would improve the performance of traffic flowing between Joe's computer and Sandy's computer. Costs are shown as numbers but could represent the capacity of each link or the costs incurred by using each link. What is the total cost of traveling the hierarchy between Joe's computer and Sandy's? In this figure, where would you add a link of cost 15 so that the total cost between Joe's computer and Sandy's is less than it is when you travel the entire hierarchy?
3. In Figure 1.9, connections are added between networks in the Internet to provide a better performing path for select traffic flows. An example of this is a content

delivery network (CDN). What is a CDN? Show how a CDN uses interconnectivity to provide better performance characteristics to its users.

4. In defining where services can be applied in a network, end-to-end is determined by where you want a service to start and stop. For example, if your WAN is supplied by a service provider (e.g., an ATM or frame relay service), you may want to define the end points and characteristics of that service. If you use IP routers at each LAN-WAN interface to that service, describe the following: (1) at which network devices would you define the end points of the service, and (2) what characteristics (service metrics) would you use to measure the service?
5. Service requirements flow from user to application to device to network, becoming more specific along the way. If you were given an application requirement for end-to-end delay (e.g., 100 ms) between an application server on one network and users on another network, for example, how might that translate into delay in the network and devices? What types of service metrics could you use to measure it?
6. For Example 1.5, the delay characteristics for the segments (including the processing at the switches) are as follows: for each GigE segment, 100  $\mu$ s; for the PoS OC-48 segment between routers, 1 ms; for each FE segment, 200  $\mu$ s; and for the security firewall, 5 ms. Draw graphs showing the end-to-end delay performance (in the direction from user PC to server) before and after the security firewall is added.
7. Which of the following applications require best-effort (unpredictable and unreliable), guaranteed (predictable and reliable, with accountability), or predictable service. Give reasons for your choices.
  - High-quality (phone company-grade) voice calls
  - Voice over IP (VoIP) calls
  - File transfers via FTP
  - Audio file downloads
  - A commercial video-on-demand service
  - User access to servers in a corporation
8. Show how performance boundaries and thresholds could be used in the following scenarios.
  - An application has a service requirement for round-trip delay to be less than 100 ms. If delay is greater than 100 ms, notify the network administrator.
  - A user requires capacity of up to 512 Kb/s but may not exceed 1.5 Mb/s. You want to keep track of how much time the user's capacity is between 512 Kb/s and 1.5 Mb/s.