DESGN

Designing for Cisco Internetwork Solutions

Version 1.1

Student Guide

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Table of Contents

Volume 1

DESGN

Course Introduction

Overview

Designing for Cisco Internetwork Solutions (DESGN) v1.1 will enable you to gather internetworking requirements, identify solutions, and design the network infrastructure and elements to ensure the basic functionality of the proposed solutions. The purpose of this five-day course is to provide you with the knowledge and skills to achieve associate level competency in network design. The course is the first design course in a curriculum supporting the Cisco network design certification track. The course focuses on the technology and methods that are currently available.

Course Outline

The Course Introduction includes these topics:

- Course Objectives
- Cisco Certifications
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Flow Diagram
- Icons and Symbols
- Learner Introductions
- Case Study and Simulations
- Course Evaluations

Course Objectives

This topic lists the course objectives.

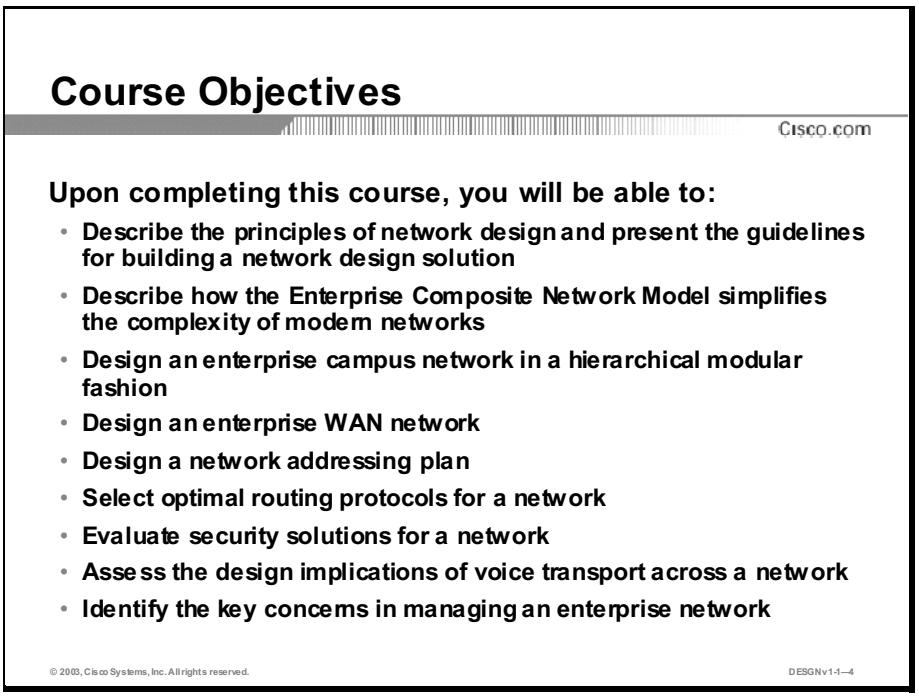

Upon completing this course, you will be able to:

- Describe the principles of the network design and present the guidelines for building a network design solution
- Describe how the Enterprise Composite Network Model simplifies the complexity of modern networks
- Design an enterprise campus network in a hierarchical modular fashion
- Design an enterprise WAN network
- Design a network addressing plan
- Select optimal routing protocols for a network
- Evaluate security solutions for a network
- Assess the design implications of voice transport across a network
- Identify the key concerns in managing an enterprise network

Cisco Certifications

This topic discusses Cisco certification and provides information on resources available for additional details.

Cisco provides three levels of general certifications for IT professionals with several different paths (or tracks) and designations. In addition, Cisco provides a variety of Cisco Qualified Specialist focused certifications to show knowledge in specific technologies, solutions, or job roles. For details, go to http://www.cisco.com/go/certifications.

Learner Skills and Knowledge

This topic lists the course prerequisites.

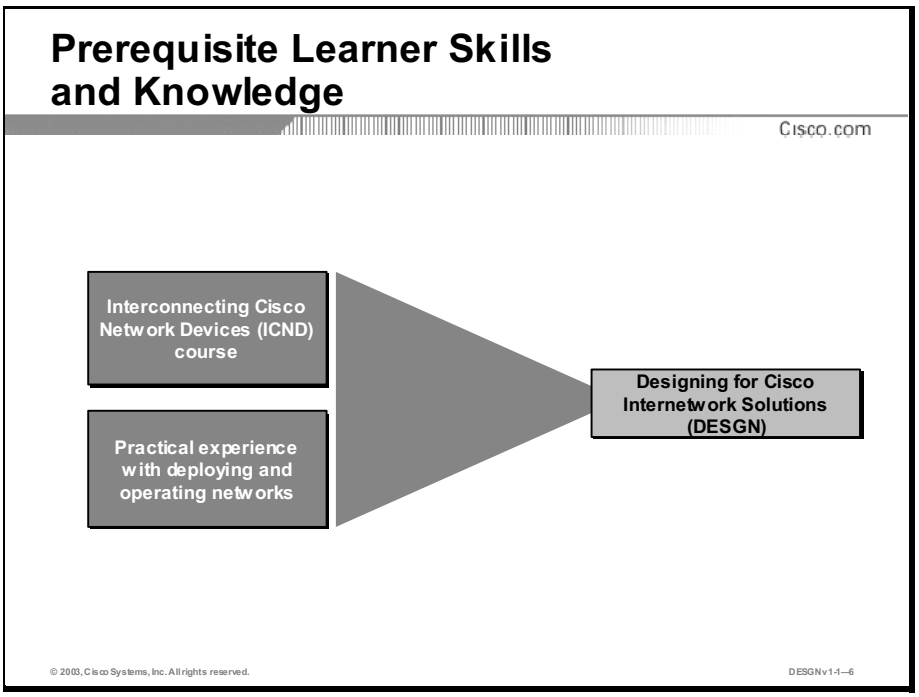

To benefit fully from this course, you must have these prerequisite skills and knowledge:

- *Interconnecting Cisco Network Devices* (ICND) course or hold Cisco CCNA® certification
- Practical experience with deploying and operating networks based on Cisco network devices and Cisco IOS software

The learners with Cisco CCNP® or equivalent level of knowledge and experience will have the advantage of being able to participate more actively in classroom discussions.

Learner Responsibilities

This topic discusses the responsibilities of the learners.

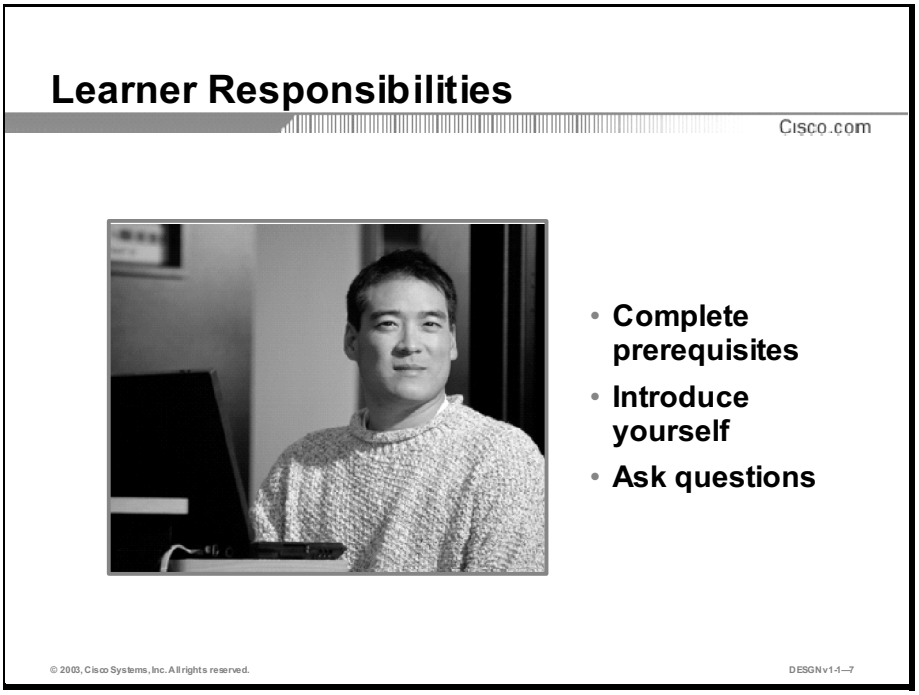

To take full advantage of the information presented in this course, you must have completed the prerequisite requirements.

In class, you are expected to participate in all lesson exercises and assessments.

In addition, you are encouraged to ask any questions relevant to the course materials.

If you have pertinent information or questions concerning future Cisco product releases and product features, please discuss these topics during breaks or after class. The instructor will answer your questions or direct you to an appropriate information source.

General Administration

This topic lists the administrative issues for the course.

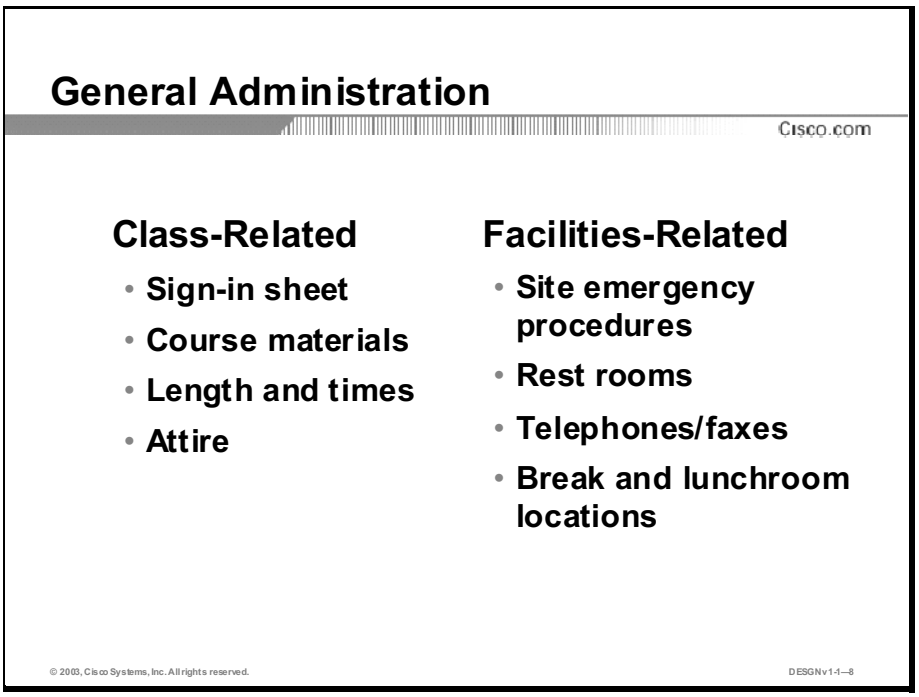

The instructor will discuss the administrative issues noted here so you know exactly what to expect from the class.

- Sign-in process
- Starting and anticipated ending times of each class day
- Class breaks and lunch facilities
- Appropriate attire during class
- Materials you can expect to receive during class
- What to do in the event of an emergency
- Location of the rest rooms
- How to send and receive telephone and fax messages

Course Flow Diagram

This topic covers the suggested flow of the course materials.

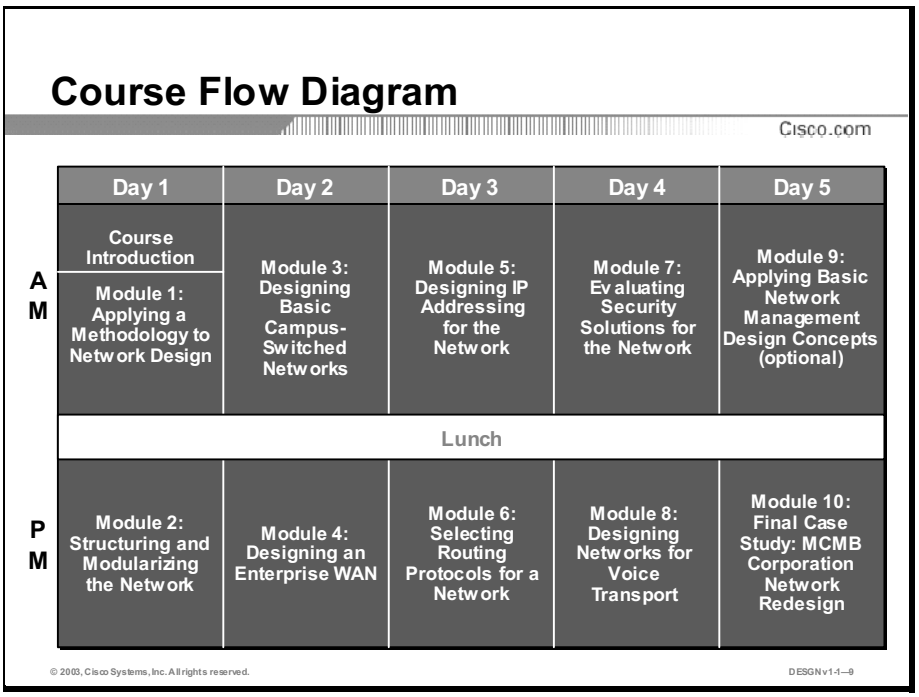

The schedule reflects the recommended structure for this course. This structure allows enough time for the instructor to present the course information and for you to work through the laboratory exercises. The exact timing of the subject materials and exercises depends on the pace of your specific class.

The instructor may introduce the Final Case Study at the beginning of the course, giving you the opportunity to do the ongoing exercises after each module or for the homework. The final discussion about possible designs will occur at the end of the last day.

Note: The “Applying Basic Network Management Design Concepts”, module will be presented in the course, depending on the availability of time.

Icons and Symbols

This topic shows the Cisco icons and symbols used in this course.

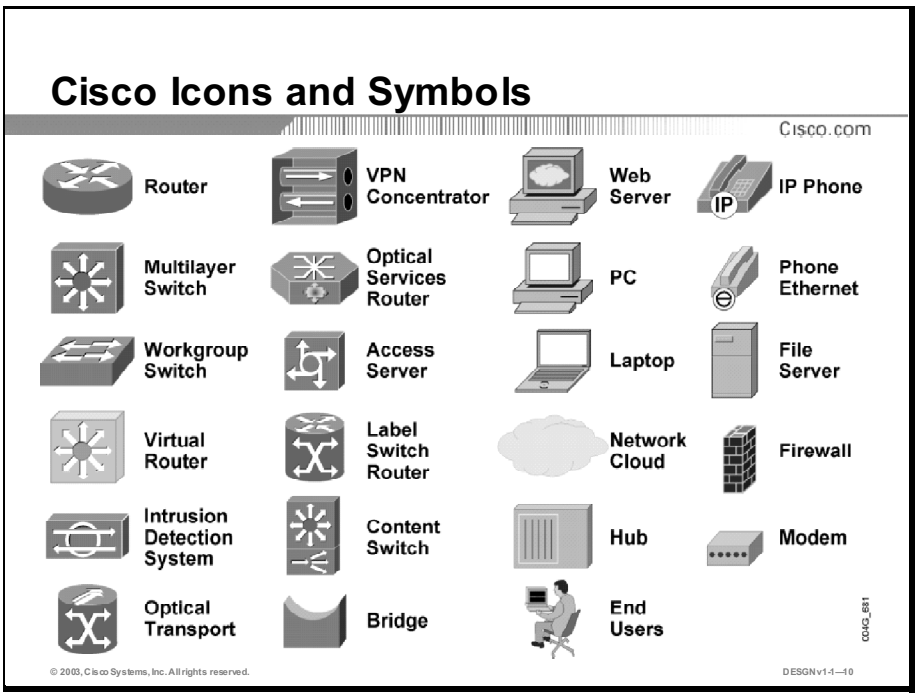

Learner Introductions

This is the point in the course where you introduce yourself.

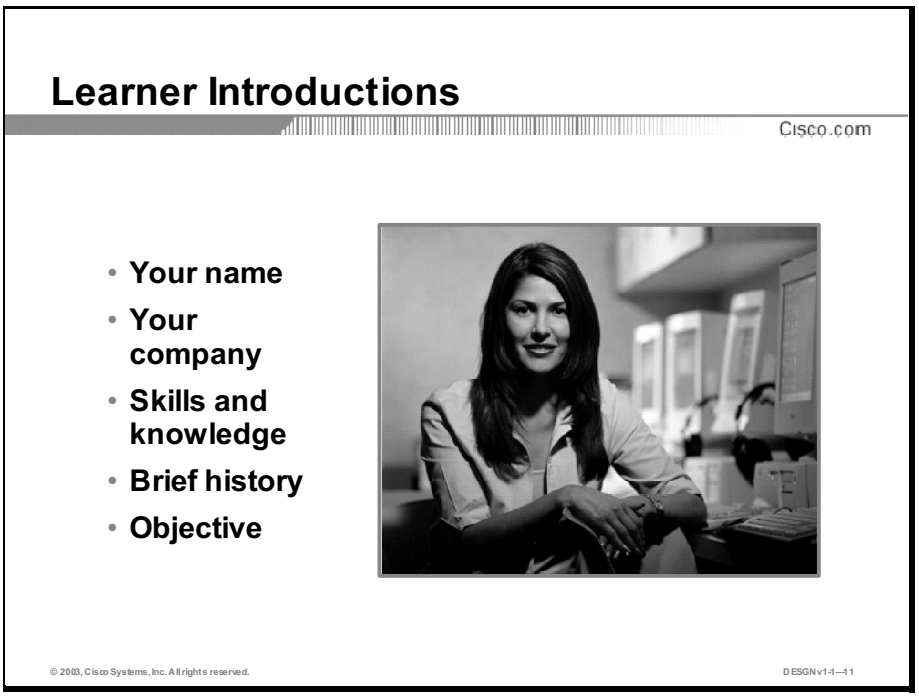

Prepare to share the following information:

- Your name
- Your company
- If you have most or all of the prerequisite skills
- A profile of your experience
- What you would like to learn from this course

Case Study and Simulations

The case study and simulations encourage you to use knowledge obtained in the course. The purpose of the case study and simulations is to provide practical application of the information you learn in this course.

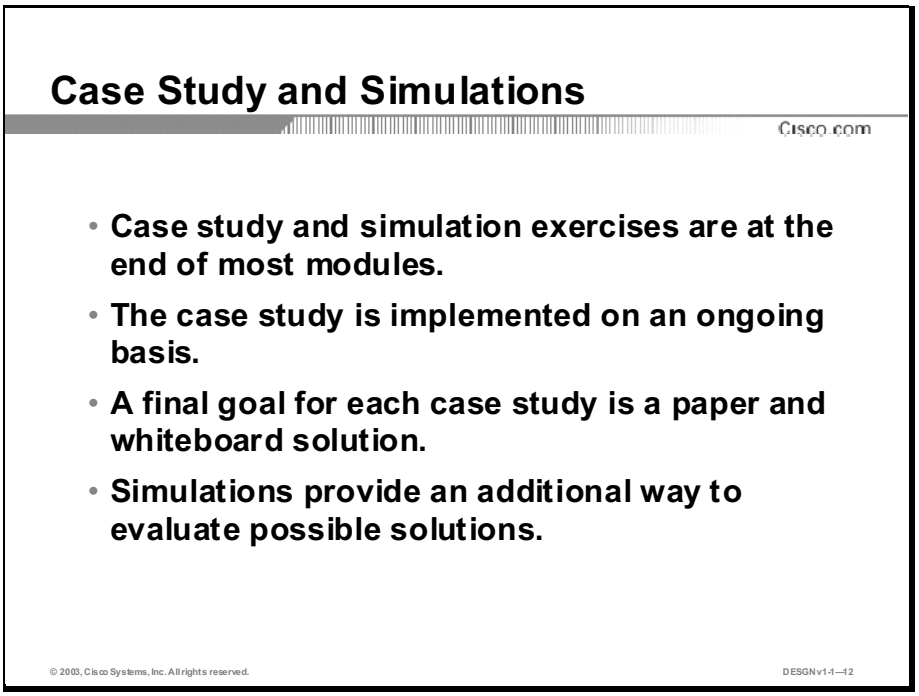

The case study is implemented on an ongoing basis, starting with an initial problem (in the first module) and later focusing on the topics covered by respective modules. The case study does not appear in all course modules but covers most of the design processes and tasks that you must perform in real-life situations. The case study will be completed on paper and may be presented using a whiteboard.

The simulations are used to evaluate selected problems and are presented in some modules only. They can be performed on paper only or through the customized application.

The detailed distribution of the case study tasks and simulations is explained in the *Specific Instructions* paragraph.

Objectives

After completing the case study and simulations in *Designing for Cisco Internetwork Solutions,* you will be able to:

- Identify missing information in the provided scenario and outline major design tasks
- Evaluate the effects of new applications on WAN links
- Plan the capacity of WAN links based on the simulation results
- Propose a campus design for a given network scenario
- Explain the benefits of using switched LAN solution compared to shared LAN solution
- Explain the differences of using Layer 2 (L2) and Layer 3 (L3) switching solutions in the campus
- Select an effective WAN transport
- Select an appropriate WAN backup strategy
- Propose an IP addressing plan for a given network scenario
- Select a suitable routing protocol for the network
- Explain the importance of fast convergence in the network
- Explain the effect of transporting voice traffic across the data networks

Disclaimers

Network design and architecture is both an art and a science. Some of the design processes are well established and based on explicit data. The numerous architectural combinations available to a designer may result in different designs. Each design choice depends on numerous parameters, ranging from technical factors to business requirements. In the case study and associated design tasks, only a few of the possible parameters are given. The result is appropriate solutions for each task of the case study.

The multiple solutions are not a problem. They reflect the art and science of network design. In real-world network design, there are few operational networks that are exactly the same.

For each task of the case study, a solution is provided that is associated with assumptions and reasoning. There is no claim that the provided solution is the best or the only solution. Your solution may be more appropriate for the assumptions that you made. The provided solution offers you a way to compare and contrast your solution with other possibilities.

Case Study Guidelines

Whether you are working individually or as part of a group, follow these guidelines for each case study:

- Use the scenarios, information, and parameters provided at each task of an ongoing case study. If there are ambiguities, make reasonable assumptions and proceed. For all the tasks, use the initial customer scenario and build on the solutions you developed so far.
- You may use any and all documentation, books, white papers, and so on.
- In each task of the case study, you act as a network design consultant. Make creative proposals to help the enterprise accomplish its goals. When your ideas differ from the provided solutions, justify your ideas.
- Use any design strategies that you feel are appropriate.
- Use any internetworking technologies that you feel are appropriate.
- A final goal for each case study is a paper and whiteboard solution. You do not need to provide the specific product names.

Specific Instructions

The specific instructions for the case study and simulations are listed in this section. Each task is explained briefly with the purpose of the simulation described. The case studies and simulations always appear in their respective module.

Module 1: Applying a Methodology to Network Design

Case Study: Network Upgrade

The purpose of the initial task of the case study is to identify the missing information in the provided scenario, using the design principles presented, and to outline major design tasks that result from the stated network requirements.

Simulation: New Applications

The purpose of this simulation is to demonstrate the use of simulation as a supplementing design tool that allows the designer to obtain results without actually performing prototype installation and operation.

Module 2: Structuring and Modularizing the Network

Case Study: Designing a Network Hierarchy

This purpose of this case study exercise is to apply the Enterprise Composite Network Model to the requirements of the company and develop a basic network hierarchy.

Module 3: Designing Basic Campus Switched Networks

Case Study: Campus Design

The result of this task is your proposal for a campus switching network design, based on the requirements from the initial scenario.

Simulation 1: Shared vs. Switched LAN

One of the requirements in the initial scenario was the proof of the benefits of LAN switching as compared with a shared LAN approach. In this simulation, the benefits are shown.

Simulation 2: Layer 2 vs. Multilayer Switching

One of the typical design issues of the campus design is the data link layer versus multilayer option. The purpose of this simulation is to demonstrate the benefits and drawbacks of each solution.

Module 4: Designing an Enterprise WAN

Case Study: WAN Upgrade and Backup

The task in this stage of the case study is to select WAN links and WAN backup based on the initial requirements and the simulation results modeling new applications.

Module 5: Designing IP Addressing for the Network

Case Study: Network Addressing Plan

In this task, you will propose an IP addressing scheme that is hierarchical and scalable.

Module 6: Selecting Routing Protocols for a Network

Case Study: Routing Protocol Selection

The purpose of this task is to select a routing protocol and to outline its deployment, based on the problems explained in the initial scenario and taking into account the routing design approach to scalable networks.

Simulation: Network Convergence

In this module, the simulation is used to examine the importance of fast network convergence after certain failures. Both physical design and logical design, including the routing, can contribute to faster convergence.

Module 7: Evaluating Security Solutions for the Network

Case Study: Designing Network Security

The purpose of this task is to apply Cisco's Security Architecture for Enterprise Blueprint (SAFE) Blueprint to your earlier design based on the Enterprise Composite Network Model and the requirements of the enterprise.

Module 8: Designing Networks for Voice Transport

Simulation: Voice Transport over IP Network

The purpose of this simulation is to show the effect of mixing data traffic with voice traffic on the same medium and to focus on the issues associated with the voice transport over IP networks.

Module 10: Final Case Study: MCMB Corporation Network Redesign

The purpose of this Case Study is to use all of the principles taught in this course.

Course Evaluations

Cisco relies on customer feedback to make improvements and guide business decisions. Your valuable input will help shape future Cisco learning products and program offerings.

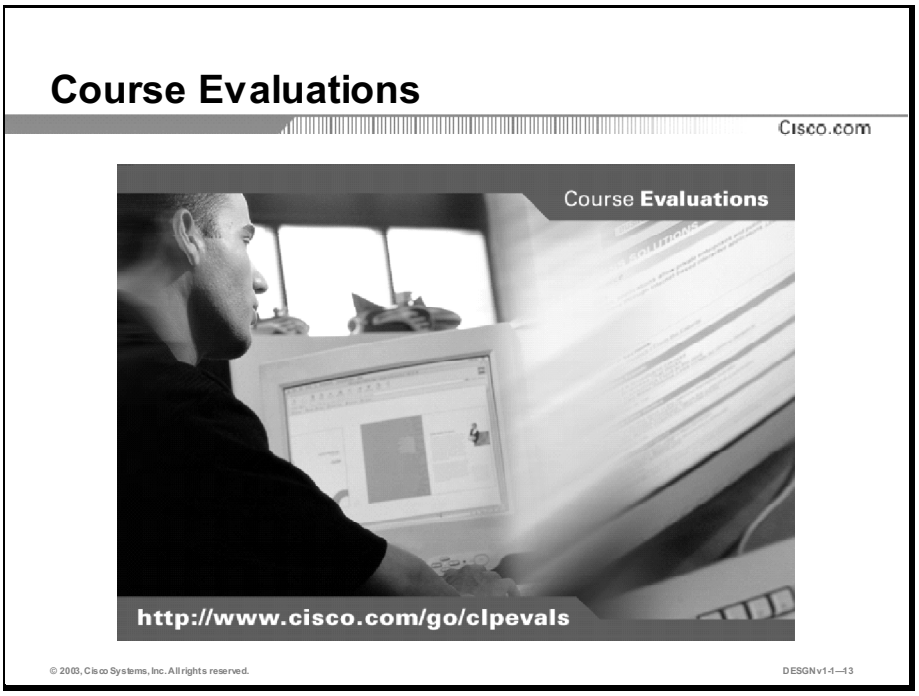

On the first and final days of class, your instructor will provide the following information needed to fill out the evaluation:

- Course acronym *(printed on student kit side label)* _____
- Course version number *(printed on student kit side label)* _____
- Cisco Learning Partner ID # _____
- Instructor ID # _____
- Course ID # *(for courses registered in Cisco Learning Locator)* _____

Please use this information to complete a brief (approximately 10 minutes) online evaluation concerning your instructor and the course materials in the student kit. To access the evaluation, go to http://www.cisco.com/go/clpevals.

After the completed survey has been submitted, you will be able to access links to a variety of Cisco resources, including information on the Cisco Career Certification programs and future Cisco Networkers events.

If you encounter any difficulties accessing the course evaluation URL or submitting your evaluation, please contact Cisco via email at clpevals_support@external.cisco.com.

Module 1

Applying a Methodology to Network Design

Overview

A network design must meet the requirements of the organization it supports. As a network designer, you should understand the needs of the organization and follow a methodology that helps match needs to the network implementation. This module introduces principles and guidelines for building an effective network design.

Module Objectives

Upon completing this module, you will be able to describe a network design methodology and the procedure used to implement a network design.

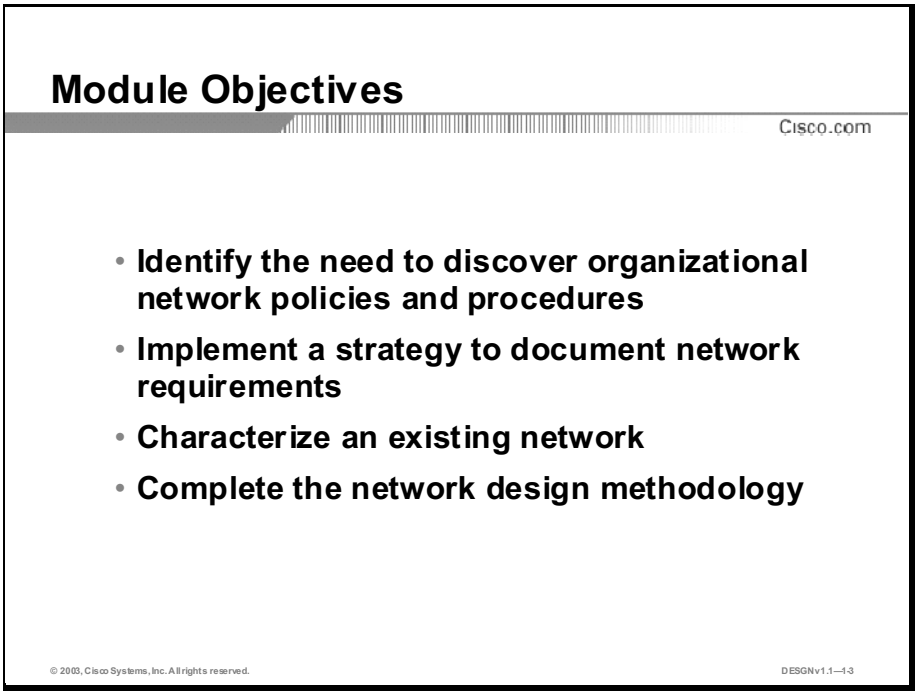

Module Outline

The outline lists the components of this module.

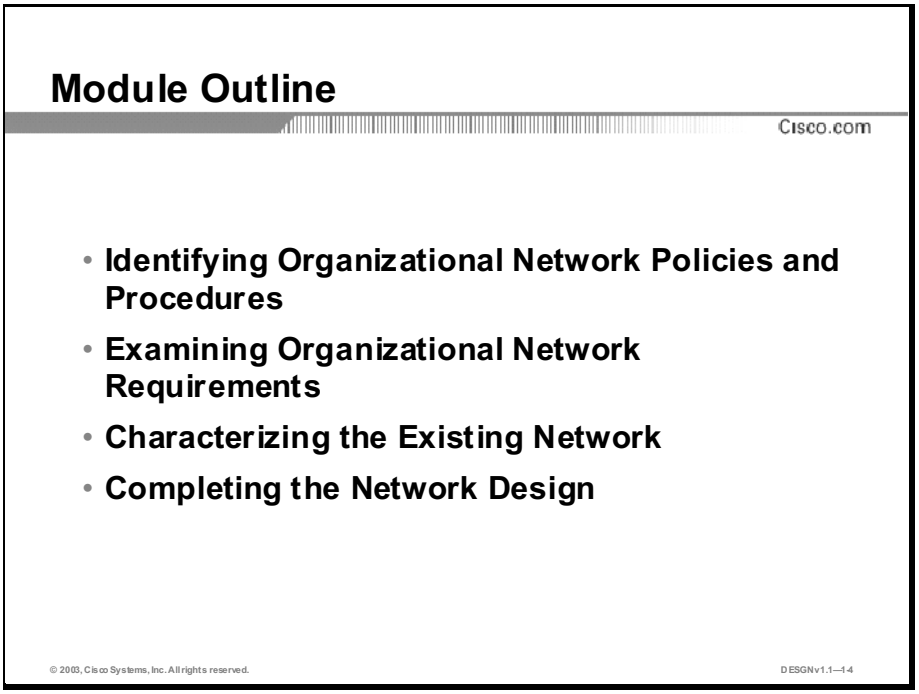

Identifying Organizational Network Policies and Procedures

Overview

Network designers will identify organizational processes and map them to the network infrastructure. This introductory lesson presents a network organizational model, its supporting infrastructure, and deployment scenarios. This lesson also examines the impact of organizational policies and procedures and describes how networking assists in achieving organizational goals.

Relevance

You must understand an organization's procedures before you can determine the enterprise network requirements.

Objectives

Upon completing this lesson, you will be able to identify the need to discover organizational network policies and procedures. This includes being able to meet these objectives:

- Identify the network organizational model components
- Identify the network architecture components
- Identify the network organizational policies
- Describe the network organizational procedures
- Describe selected deployment scenarios for an organization's network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with basic networking concepts and technologies

Outline

The outline lists the topics included in this lesson.

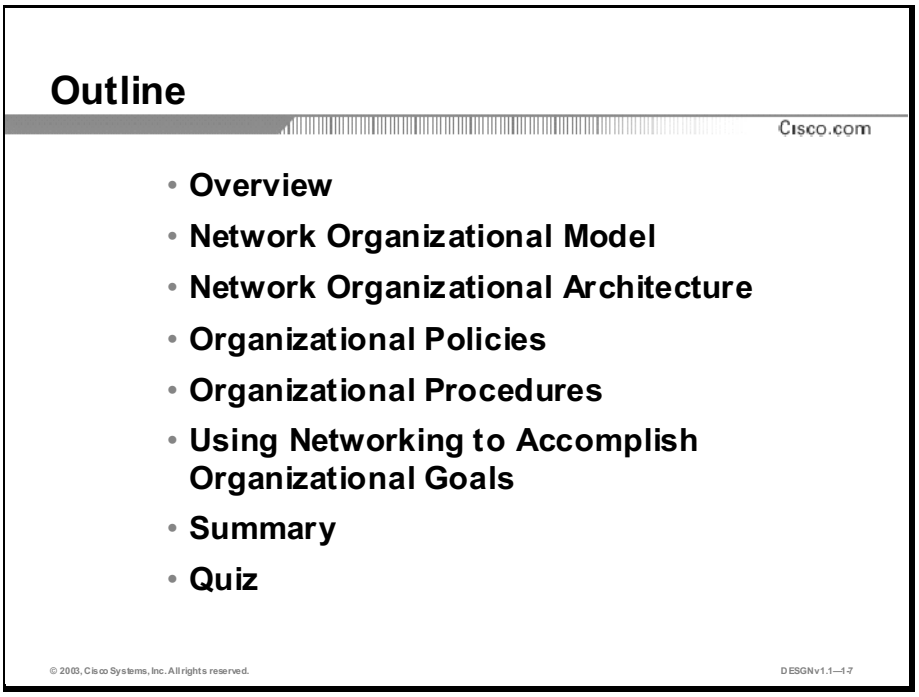

Network Organizational Model

Modern networks are evolving to provide integration within the organization as well as to provide access to the outside world. This topic identifies the network organizational model components.

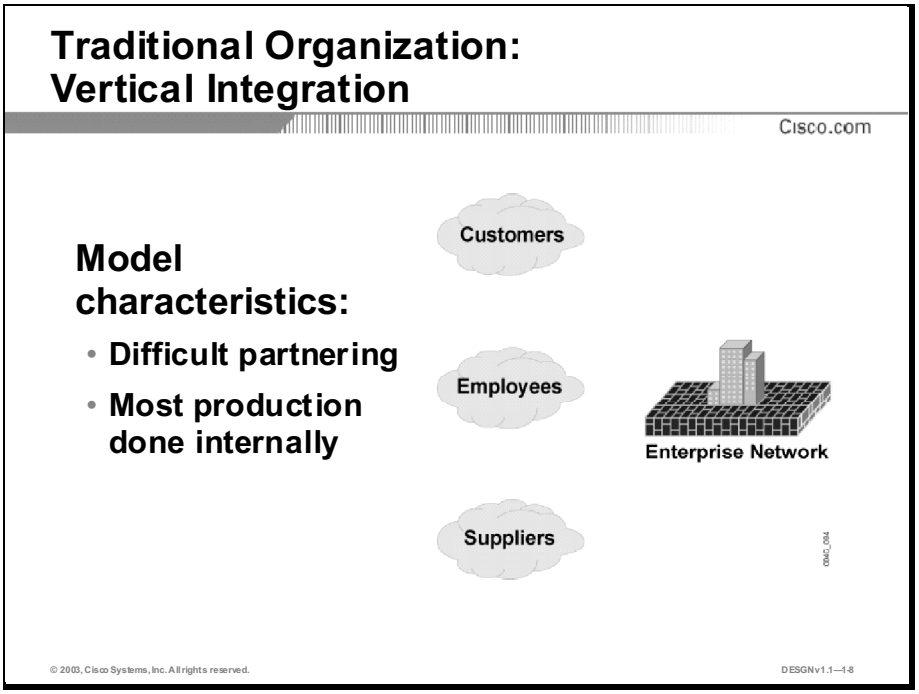

A traditional corporation is often built on a model that presents a closed infrastructure and provides limited external integration. The vertical integration model of intraorganizational communication often presents network connections externally, without much effort spent on providing network services. Only very limited access is offered to external users.

Internetworking introduces a more responsive model of organizational communication. The Internet presents an open market with a wide availability of goods, resources, and people. Within such a market, Internet applications offer the potential to increase service availability dramatically.

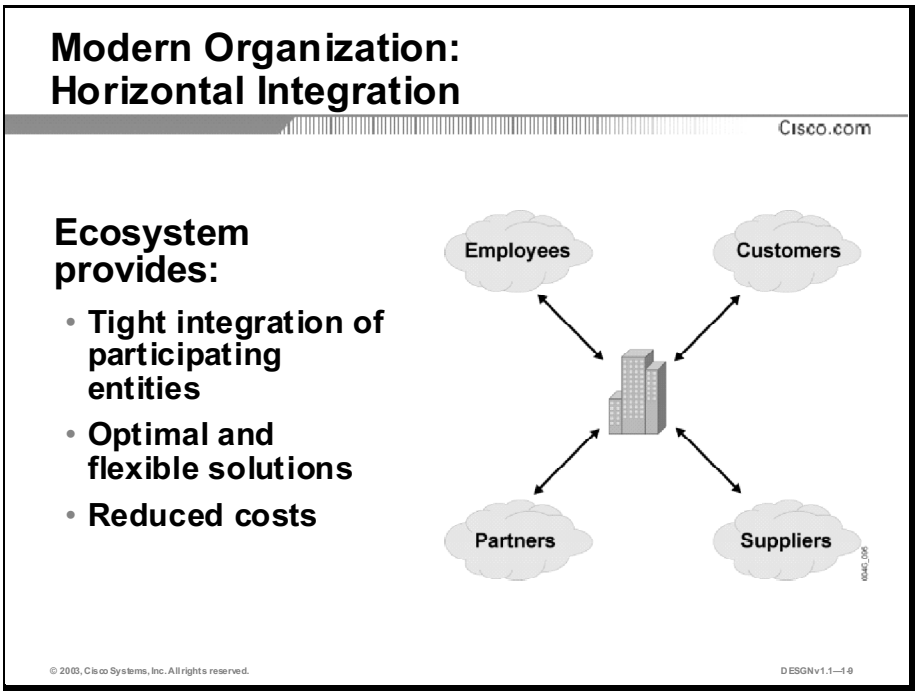

With the introduction of the Internet, such transaction costs may drop to be equal to, or lower than, the existing internal transaction costs. Horizontal integration to take advantage of the reduced costs may make sense.

Horizontal integration involves focus on core competencies and partnership with others to provide supporting activities. This approach forms the basis for the network organizational model. This model is based on three core assumptions:

- The relationships that a company maintains with its key partners can be as much of a competitive differentiator as its core products or services.
- The manner in which a company shares information and systems is a critical element in the strength of its relationships.
- Being connected is no longer adequate. Organizational relationships and the communications that support them must exist in a networked fabric.

The horizontal network organizational model builds on a system that integrates all the participating entities into an organizational ecosystem. Organizations can create ecosystems either internally within an organization, externally with partners and suppliers, or both.

A modern network organizational model is formed by these entities:

- **Employees:** For organizations to function most effectively, information must be readily available to employees. Intranet applications provide the backbone for immediate access to current information and services.
- **Customers:** Using online support services, customers are provided more convenient services faster, and any customer, large or small, located virtually anywhere, has access to these services. In addition, online support services cost less than traditional services.
- **Partners:** Successful partnerships leverage the resources of each partner.

- **Suppliers:** The purchasing function (ordering, delivery, and billing) can be both time and labor intensive and expensive. Organizations can leverage their networks to create links to their suppliers, resulting in less costly, just-in-time transactions.

The leaders in any area of endeavor will be those who employ network technology effectively to reach the goals of improved productivity, reduced time to market, greater revenue, lower expenses, and stronger customer relationships. An organization's network must be designed and deployed to support the organization's requirements effectively. A networked organization views the network as a means to support applications that generate revenue, reduce costs, and improve customer and supplier relationships.

Example: Automobile Manufacturing Ecosystem

An example of how a network can improve an ecosystem is the modern automobile manufacturer. Automobile manufacturers contract with partners who specialize in manufacturing particular components and technologies. For example, a partner who has the expertise in manufacturing engines produces the engine. With the cooperation of the component partners, the automobile manufacturer assembles the completed automobile. When all the partners are online, the transactions have minimal cost with just-in-time manufacturing. These relationships and the ability to share information save time and money for the automobile manufacturer and its partners.

Network Organizational Architecture

The modern organizational model is built around a modular architecture. This architecture supports applications built on common network solutions, using shared network services over a scaled network infrastructure. This topic identifies the network architecture components.

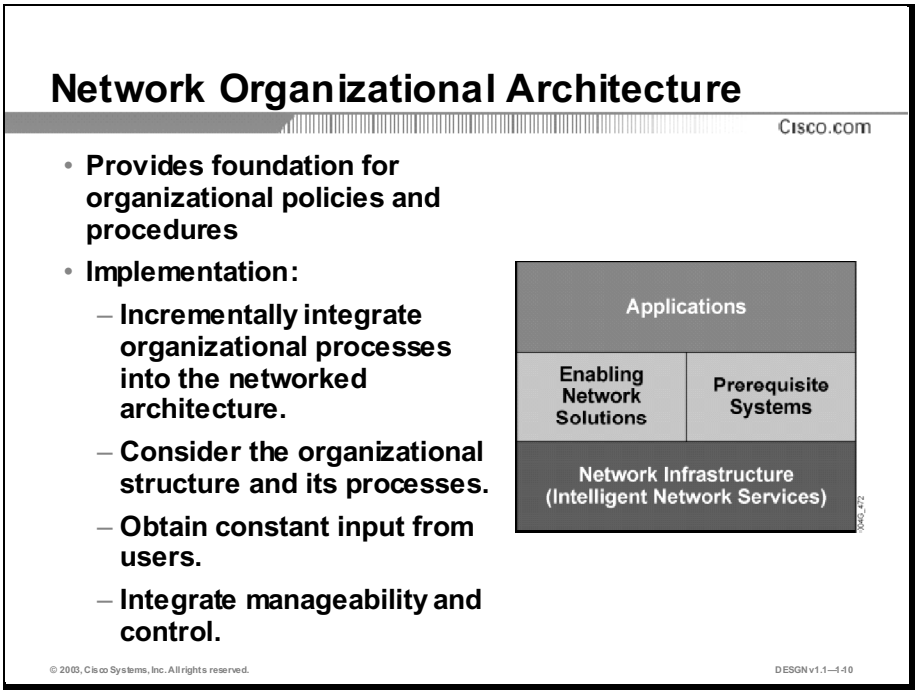

Modern organizations require a flexible, scalable, and robust infrastructure that is built on a network architecture structured for current and future organizational growth. Modern organizations are able to streamline operations for two major reasons:

- They use technology aligned with their organizational needs.
- They establish a technology foundation that allows them to build critical applications more easily.

Network Architecture Components

The network organizational architecture is divided into layers, making the organizational processes easier to implement or expand. Each layer of the architecture has special tasks that contribute to the success of the organizational processes:

- **Applications:** Applications address organizational goals directly. They offer a discrete set of functions, accessed via the network, for authorized users such as employees, customers, suppliers, or partners.
- **Enabling network solutions:** Solutions such as voice transport or content networking make modern networks more versatile so that they can provide a greater range of functions and better support application requirements.
- **Prerequisite systems:** Prerequisite systems are combinations of structured data and business logic, sometimes wrapped in an application that exposes information as requested or directed.

- **Network infrastructure:** Network infrastructure includes network platforms and links, coupled with intelligent network services. It provides quality of service in a highly available, managed, and secure network.

Implementing a Network Organizational Model

An implementation path toward the network organizational model should be incremental and logical, starting small and growing while success builds. For example, many companies select customer support as a critical area because they foresee the potential of the network to help develop closer relationships with their customers.

The network architecture of an organization must reflect its logical structure and the processes conducted within the organization and its ecosystem.

Input from users ensures that applications will be usable and that the resulting architecture will bring the anticipated benefits to the organization.

Open access to information, resources, and services through a network organizational environment sets new standards for relationships with customers, clients, partners, suppliers, and employees. Therefore, control and manageability are desired on all levels of the network architecture.

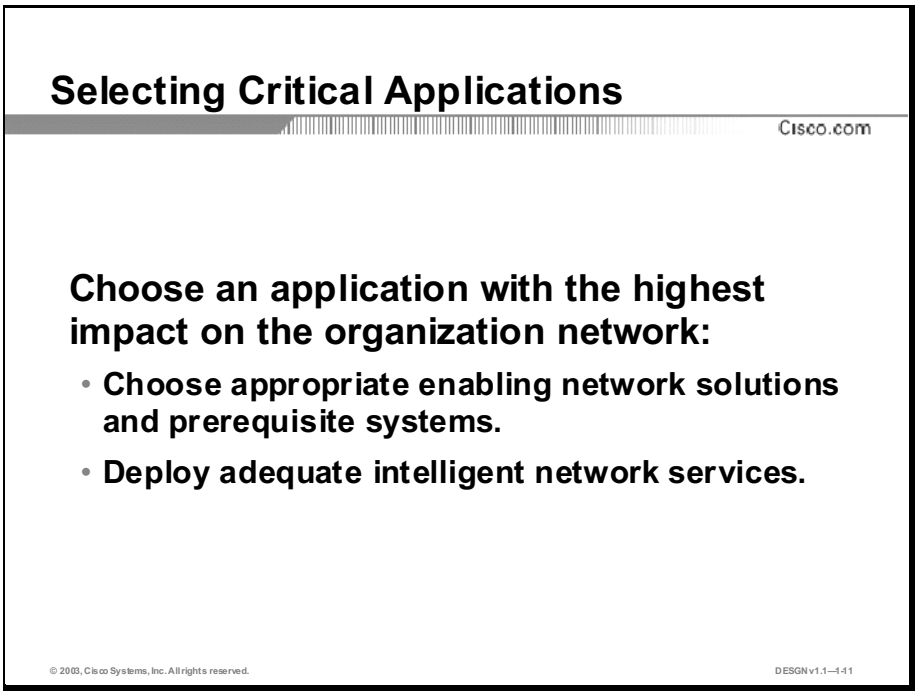

An organization should begin the network organizational model implementation by selecting an application that has the greatest impact on organizational processes, keeping in mind that the model is not about incremental improvements in existing tasks. The approach looks for new ways of sharing information, tools, and systems to build stronger organizational relationships.

Given the process requirements, the design should include the necessary network solutions, prerequisite systems, and intelligent network services. For example, if the process requires secure search of repository data, the application must provide the appropriate database, search engine, and security.

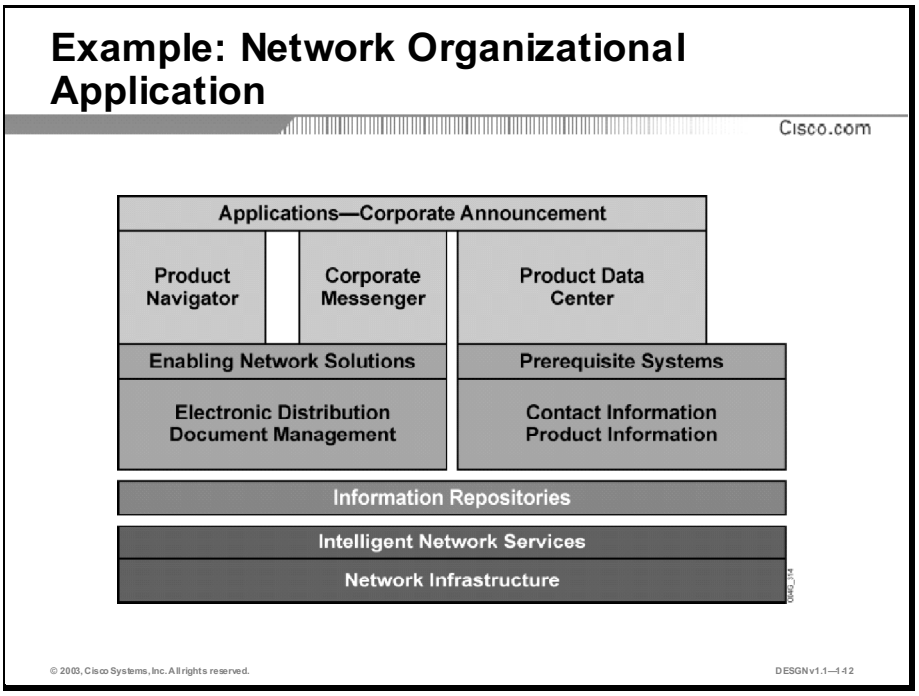

The example presents an organizational application, which introduces mechanisms for faster and easier announcement of new or modified products to corporate users, partners, or customers.

The figure illustrates three different business applications. The applications are simple to use, and users obtain all necessary data directly from each application.

- **Product navigator:** The product navigator application helps customers find information about a product, based on their needs.
- **Corporate messenger:** If customers want to receive all announcements for a given product, such as a new model, they can use the messaging application. This application sends information automatically to customers without manual intervention from the corporate administrator.
- **Product data center:** The product data center presents an application to retrieve product information from the central storage in the format that the customer demands. For example, the customer may choose to display the page on the web or download it in Adobe Acrobat format.

Enabling network solutions and prerequisite systems are needed for the network organizational application to work. The application on the upper layer uses the lower-level solutions to access the information requested by the user, faster and more easily. For example, when a user requests an information sheet for a product, the application accesses the repository and, with the help of the document management technology, searches for the requested document.

Organizational Policies

Every organization implements specific policies to achieve its goals. You should be familiar with an organization's policies before designing an enterprise network. This topic identifies the network organizational policies.

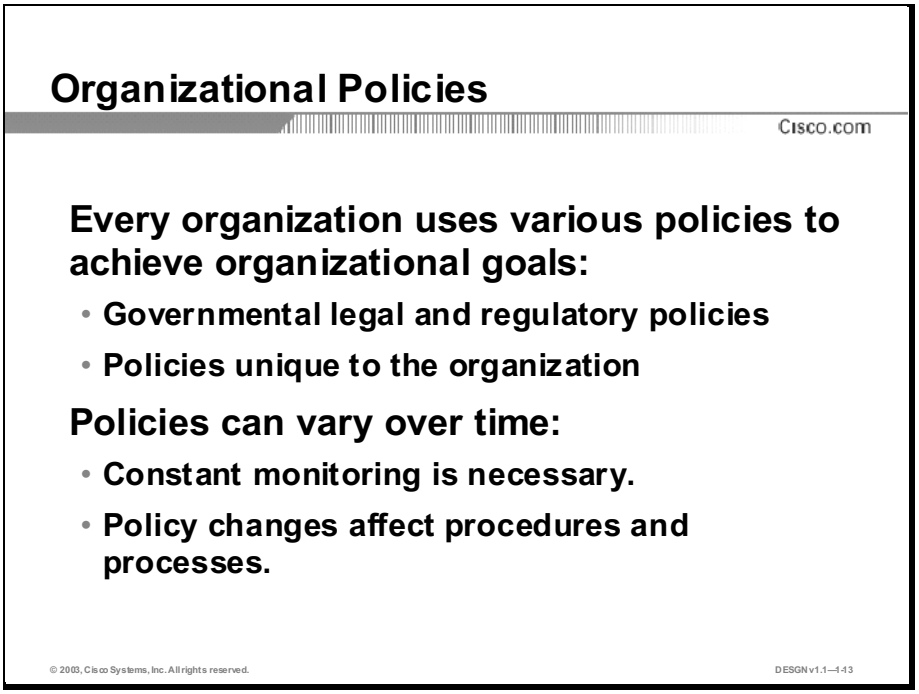

Policies are the rules and guidelines that an organization follows to achieve its goals. Policies are developed, implemented, and maintained at all levels of an organization. Internal documents explain the policies, procedures, and standards of an organization's operations. Organizational policies are divided into:

- Governmental legal and regulatory policies
- Policies specific to the organization

Governmental Legal and Regulatory Policies

Organizations implement policies in response to legal, regulatory, or contractual requirements. In the United States, generally accepted accounting principles (GAAP) drive many accounting policies. In other geographies, available network equipment choices may be limited by governmental Post, Telephone, and Telegraph (PTT) organizations.

Policies Unique to the Organization

Unique policies require an understanding of the organization's goals, mission, and desired outcomes. Understanding specific policies may also require sensitivity to an organization's risk tolerance.

Examples often include policies relevant to building a network infrastructure that address technology, vendor orientation, and preference policies. If an organization already uses a certain type of equipment, the network maintenance organization understands the equipment operation and is therefore less inclined to change the vendors. For example, the organization may use equipment that runs a standard routing protocol. The organization's common policy for Layer 3 (L3) equipment is that the network must support the selected routing protocol. Equipment that violates this policy should not be included in the network design.

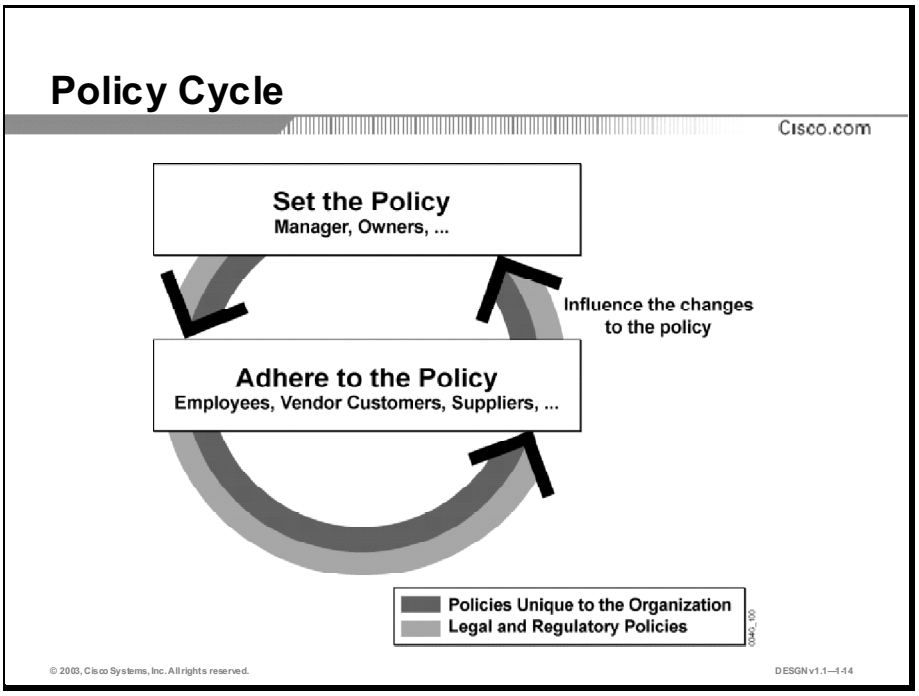

Management typically sets the policy and monitors its implementation. Management measures the subsequent impact on the organization, and redefines the policy as necessary to align with organizational goals.

The typical goals of organizational policy are to direct and align the effort of the organization to achieve a common goal. Therefore, employees, partners, customers, and suppliers must comply with the policies to maintain business relationships.

Policies within the organization may vary over time, and this impacts the organizational workflows and outcomes. Evolving market demands, regulatory changes, or organizational operations may drive these policy changes. Small companies may change their internal policies to meet the requirements of partner companies' network specifications, or to accept new solutions that do not align with current policy.

Example: Policies for Partnering

Company B decides to use a new application with a partnering company A, but finds that the application is not compliant with the policy of company A. The management of company A may decide to redefine the partnership with company B because of the noncompliance in policies.

Example: Inter-Organization Communication Policies

Organization X requires secure communication with organization Y. However, neither communication nor network security exists within organization X. This situation requires organization X to develop a new policy that defines how secure communication with organization Y will be achieved.

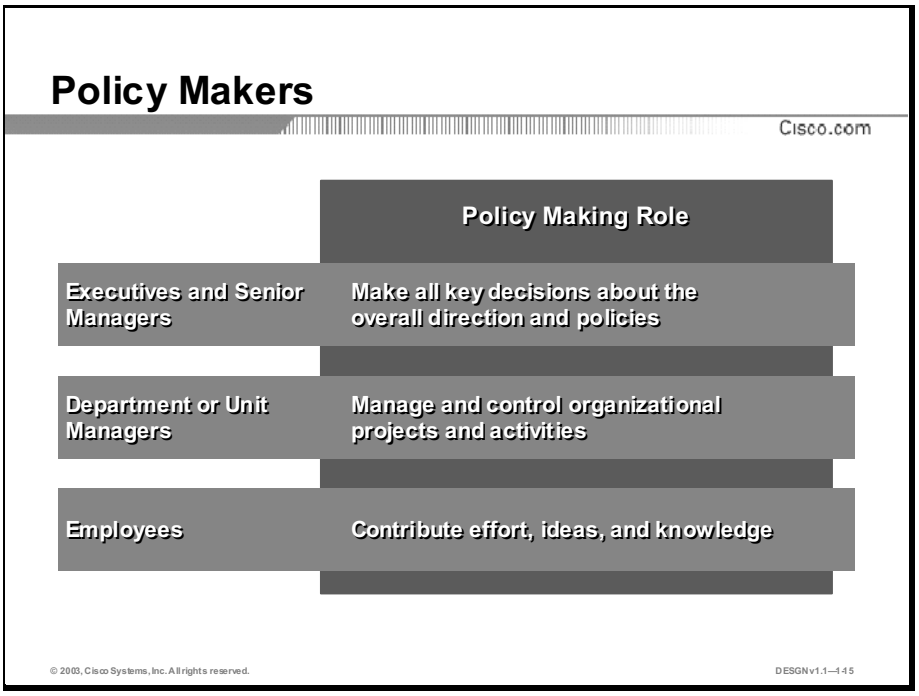

Organizations are divided into levels, each with its own responsibilities and policies. For example, there may be three core levels within an organization:

- **Executives and senior managers:** Make major strategic business decisions from the top of an organizational hierarchy.
- **Departmental or unit managers:** Assign tasks to nonmanagerial employees. In many companies, one manager may oversee a number of subordinate managers within a department, who in turn provide direction to nonmanagerial employees. However, within a small company, a single manager may control all activities within a single department.
- **Employees:** Use their knowledge, effort, and ideas to contribute to the completion of the tasks assigned by their supervisors. Employees are often specialists. For example, workers on an automobile assembly line are specialists in the specific parts of the car they assemble. Specialization enables an organization to develop and implement complex and highly demanding tasks.

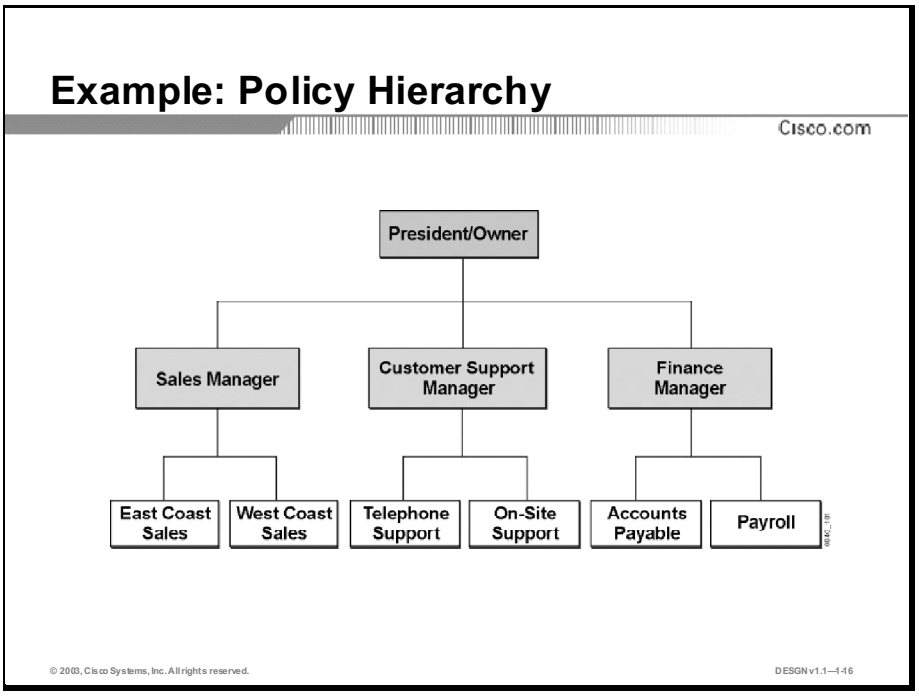

The organizational chart in the figure illustrates a hierarchical structure for a small- to mid-sized organization.

The president or owner is at the top of the hierarchy. The president and managers decide which projects they will undertake, identify the benefits the projects will bring to the organization, and address the obstacles they need to overcome for the projects to succeed.

The figure presents three upper managers who oversee departments:

- **Sales manager:** Manages and controls the east and west coast sales departments
- **Customer support manager:** Controls the customer support department and any other department involved in customer relations (for example, a customer call center)
- **Finance manager:** Controls the finance department, which is responsible for carrying out financial transactions

Other departments, such as information technology (IT) and manufacturing, are not represented.

Organizational Procedures

Organizations define procedures that support their organizational policies. Organizational procedures significantly influence network design. Each department has its own functions and tasks. The number and functions of departments may vary depending on the size and type of organization. In a successful company, all departments need to work together to achieve the best results. Every department has an assigned role in the organizational procedures. Well-structured organizations can react more quickly and compete more effectively in a rapidly changing environment. This topic describes the network organizational procedures.

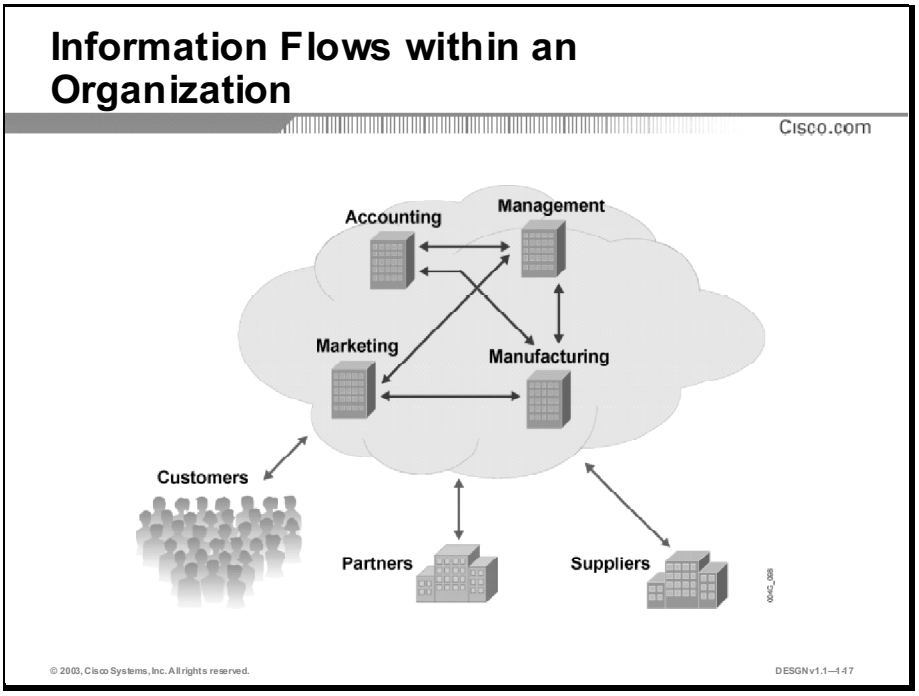

The figure illustrates a sample organizational structure and the procedural interrelations among its components.

Typically, departments such as manufacturing, accounting, marketing, and management are interconnected internally through information flows that help to achieve organizational goals. External links to partners, customers, and suppliers also exist. Most organizational procedures affect a number of departments as well as external entities.

As the network designer, you must determine the major processes and procedures for each department and the communications to external partners, suppliers, or customers. For example, suppliers communicate more often to manufacturing than to marketing, and partners often relate more directly to marketing and management than to manufacturing, although the flow of communications may vary.

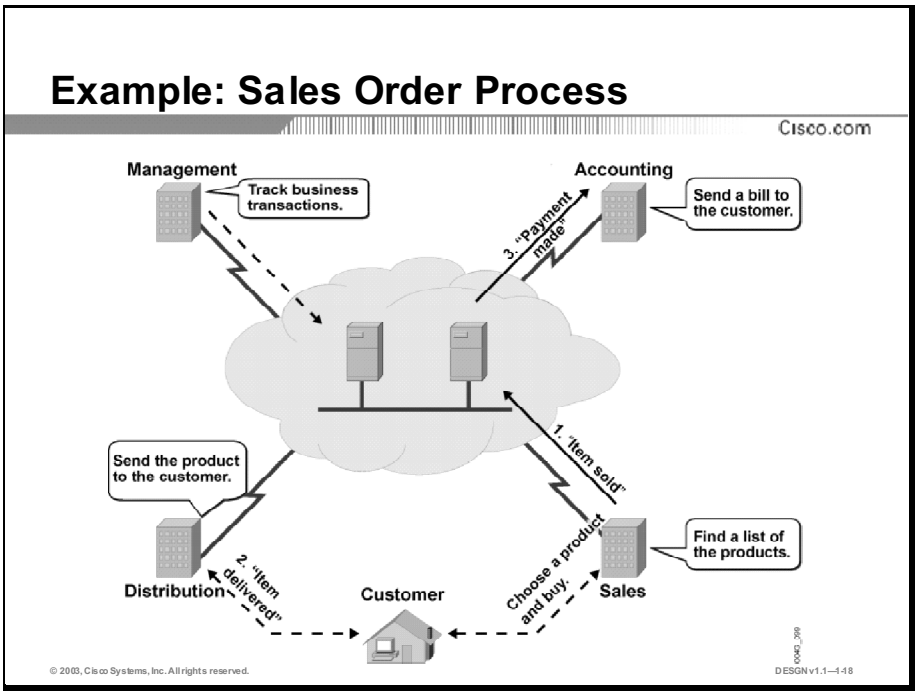

The figure illustrates the communications that occur when a customer buys a product from an organization and how information technology can enhance the relationships and procedures within the organization. The underlying information technology provides the services for all departments. Data is available to all involved participants, without having to enter that data more than once.

In the example, the customer chooses one of the products from a list based on sales and marketing data available in the system. When the selection is completed, the application stores the collected data from the customer, requests the delivery, and initiates the payment process.

The distribution division uses the order information to ensure the delivery of the product. At the same time, the application informs the remaining departments about the status of the purchase. The application updates the database information and proceeds with additional tasks.

The accounting department controls the payment process and generates the invoice.

Using Networking to Accomplish Organizational Goals

The integration of the Internet with the organization offers new opportunities for improvements to network organizational processes and procedures while challenging traditional policies. Few modern companies can survive without carefully established partnerships that require an open, yet controlled, access. This topic describes selected deployment scenarios for an organization's network.

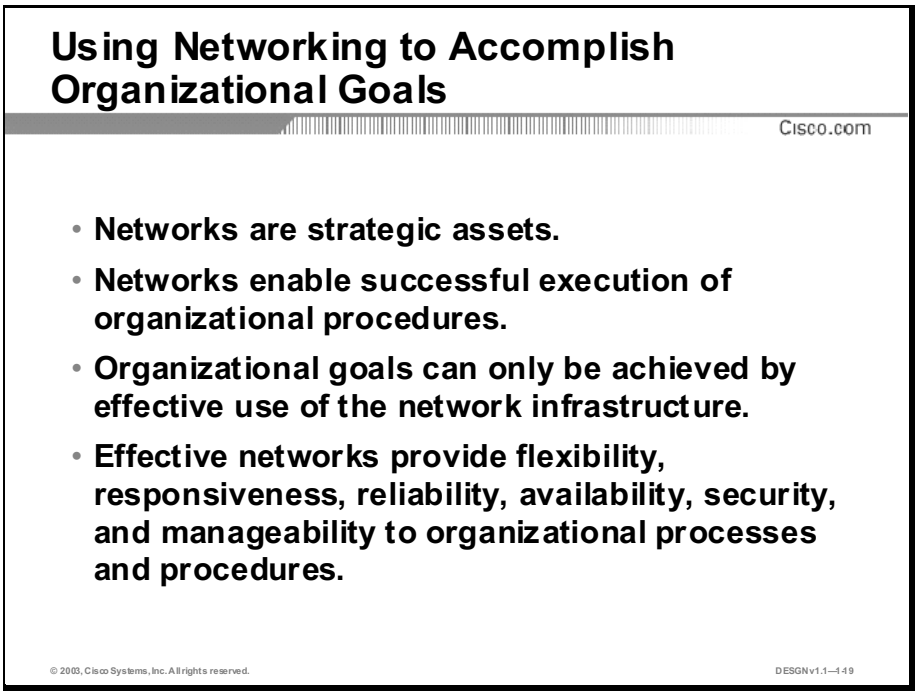

Many traditional networks are focused on internal connectivity only, rather than on reliability and manageability. A modern network based on the network organizational model overcomes these issues, creating a more flexible, more responsive, and highly available open network.

Network technology based on the network organizational model can dramatically increase efficiency, productivity, and customer satisfaction.

Example: Internet Financial Transactions

In the past, the customer had to be physically present to carry out complex financial transactions with a bank. The Internet, a globally enabled network, supports most financial transactions from the home or from the office.

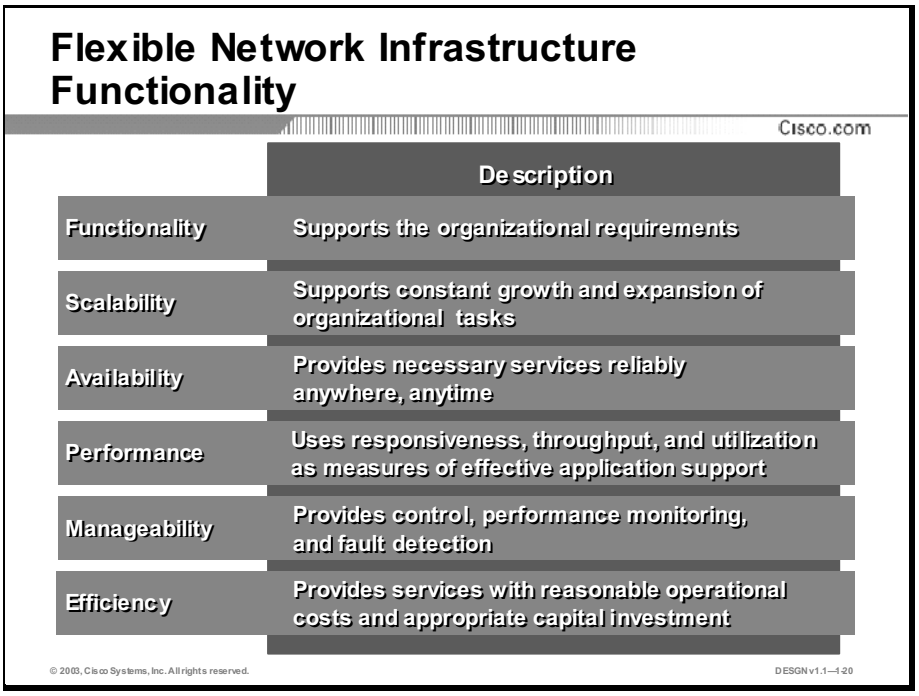

Business networks allow organizations to stay current on the latest industry trends and developments, resulting in a competitive advantage. A flexible network infrastructure that can scale to constant change is a key to successful deployment of network applications that support organizational processes.

Goals cannot be achieved without a full understanding and integration of these networking requirements:

- **Functionality:** Organizational applications require a fully functional network to support the goals defined by the organization. The network must be able to support the applications that are required to conduct the organizational processes.
- **Scalability:** Constant growth and expansion of the organizational processes require the network infrastructure to support the same scalability as the organization, with easy and inexpensive investment in the infrastructure.
- **Availability:** Critical business applications require networks to provide services on a 24-hours-a-day basis, thus to be highly available. All components of the network infrastructure must provide redundancy and resiliency.
- **Performance:** Organizational applications require a certain level of performance from the network. Networks must be able to identify users and applications and provide the responsiveness and throughput required while maintaining economic levels of utilization.
- **Manageability:** Management systems play an important role in today's organizational processes. They improve control, capacity management, performance monitoring, and fault detection for professionals who are not necessarily technical experts in all the applicable disciplines. For example, a management system may provide case-tracking or statistical analyses of critical business events. As a critical asset, organizations must proactively manage the network.
- **Efficiency:** Efficiency of the network infrastructure in the organization provides optimum results with appropriate investment. The approach, such as converging voice and data on a single integrated network, can lead to enhanced efficiency and cost-effectiveness.

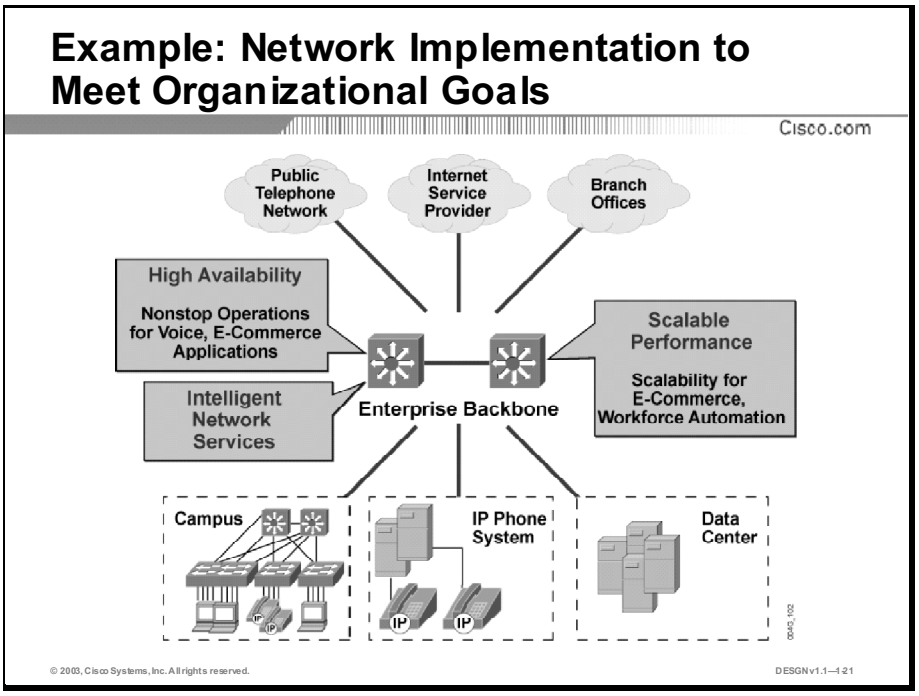

The figure illustrates a network infrastructure, with an emphasis on the features needed to accomplish specific organizational goals.

A high-speed connection and fast network devices in the enterprise network backbone (supporting a network organizational model) can achieve the required performance. The devices in this network design are based on versatile, scalable platforms, which can provide further functionality with reduced additional investment. Such scalability may provide sufficient capacity to run an organizational application such as e-commerce without immediately replacing the existing infrastructure. This network configuration also may support the high availability, reliability, and manageability that an e-commerce application requires to be available 24 hours a day to support users all over the world.

The example network transports data and other information, such as voice, thereby increasing efficiency while reducing installation, ongoing operations, and management costs.

Summary

This topic summarizes the key points discussed in this lesson.

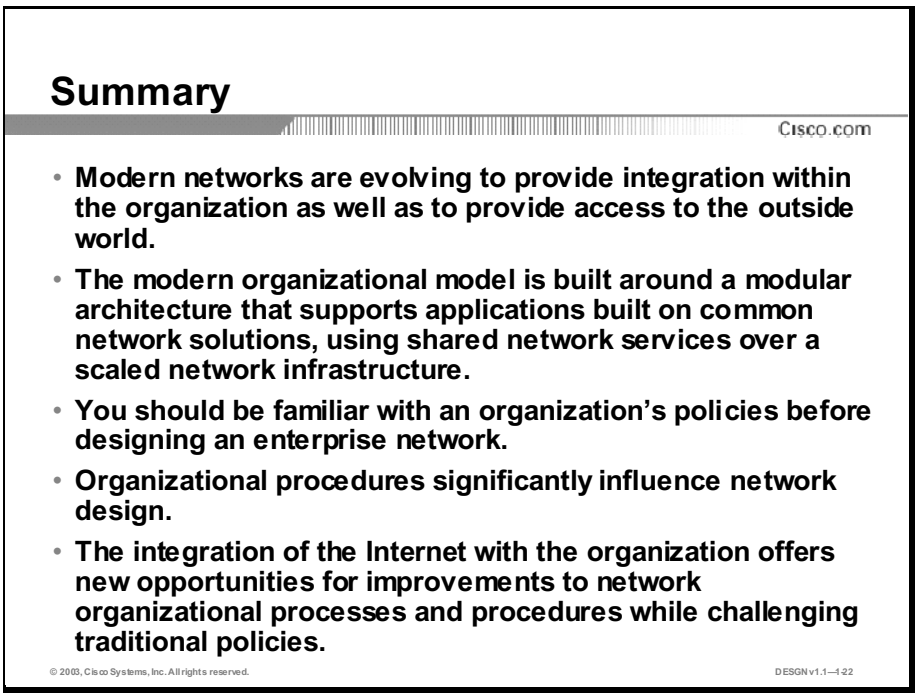

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three features identify the traditional business model? (Choose three.)

A) vertical business organization

B) high efficiency

C) isolated information systems

D) long production cycles

E) instant access to relevant information

Q2) Which three options are the components of a network architecture? (Choose three.)

A) product navigator

B) applications

C) intelligent network services

D) corporate messenger

E) training

F) enabling network solutions

Q3) Which two sets of policies affect the achievement of organizational goals? (Choose two.)

A) external policies

B) organization-specific policies

C) employment policies

D) governmental policies

E) networked architecture policies

Q4) To allow a company to react more rapidly and compete more efficiently, the IT infrastructure should reflect the organizational ____.

A) structure

B) procedures

C) department

D) management hierarchy

Q5) Which three network features are needed to achieve organizational goals? (Choose three.)

A) scalability

B) high availability

C) self-adaptability

D) performance

E) adherence to business standards

Quiz Answer Key

Q1) A, C, D

Relates to: Network Organizational Model

Q2) B, C, F

Relates to: Network Organizational Architecture

Q3) B, D

Relates to: Organizational Policies

Q4) A

Relates to: Organizational Procedures

Q5) A, B, D

Relates to: Using Networking to Accomplish Organizational Goals

Examining Organizational Network Requirements

Overview

The network design methodology is derived from the planning, design, implementation, operation, and optimization (PDIOO) methodology that reflects a network's life cycle. The design cycle begins by identifying organizational requirements and concludes with the network design verification. The final step is to validate the design through a prototype or pilot project. To design a network that meets the stated requirements, you must identify the organizational goals, organizational constraints, technical goals, and technical constraints.

This lesson describes the design process and the processes of determining what applications and network services already exist and which are planned, along with associated organizational and technical goals and constraints.

Relevance

This lesson provides the learner with the guidelines for extracting customer requirements for the intended network. It also serves as a template for a designer.

Objectives

Upon completing this lesson, you will be able to implement a strategy to document network requirements. This includes being able to meet these objectives:

- Describe the role of design as an integral phase in a network life cycle
- Identify the steps in the network design methodology
- Identify a process for gathering enterprise network design requirements
- Identify the planned applications and network services in an intended network
- Assess organizational goals for the network design project
- Identify organizational constraints that affect the network design
- Assess technical goals for the network design project
- Identify technical constraints that affect the network design

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with basic networking concepts and technologies

Outline

The outline lists the topics included in this lesson.

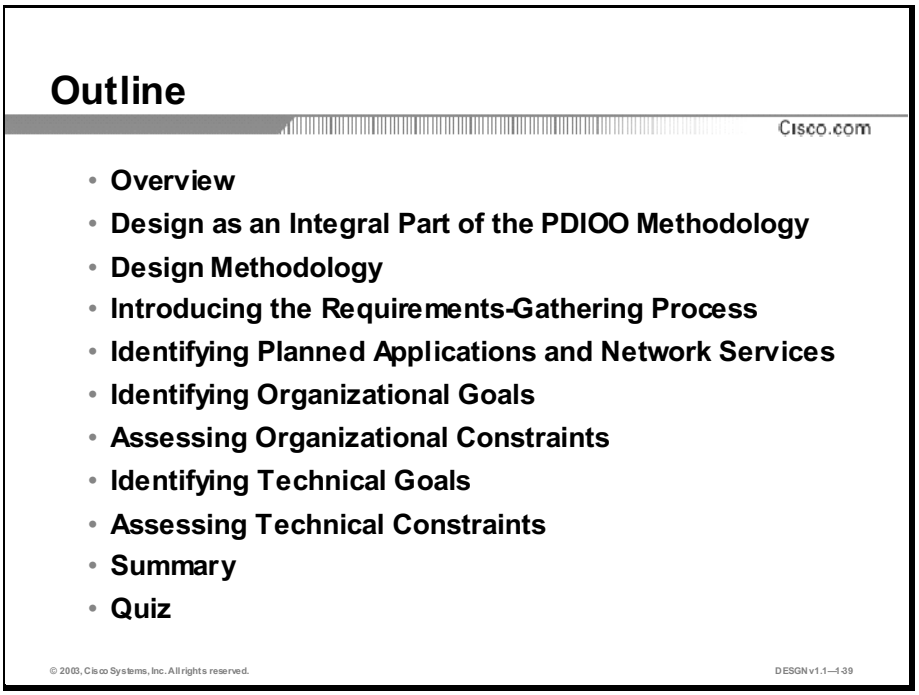

Design as an Integral Part of the PDIOO Methodology

The planning, design, implementation, operation, and optimization (PDIOO) methodology reflects the phases of a standard network's life cycle. These phases are encountered in the operation of every network. There is a close relationship between design and the other PDIOO phases. While design is one of the five PDIOO phases, design elements are present in all the other phases, and the other phases influence design decisions. This topic describes the role of design as an integral phase in a network life cycle.

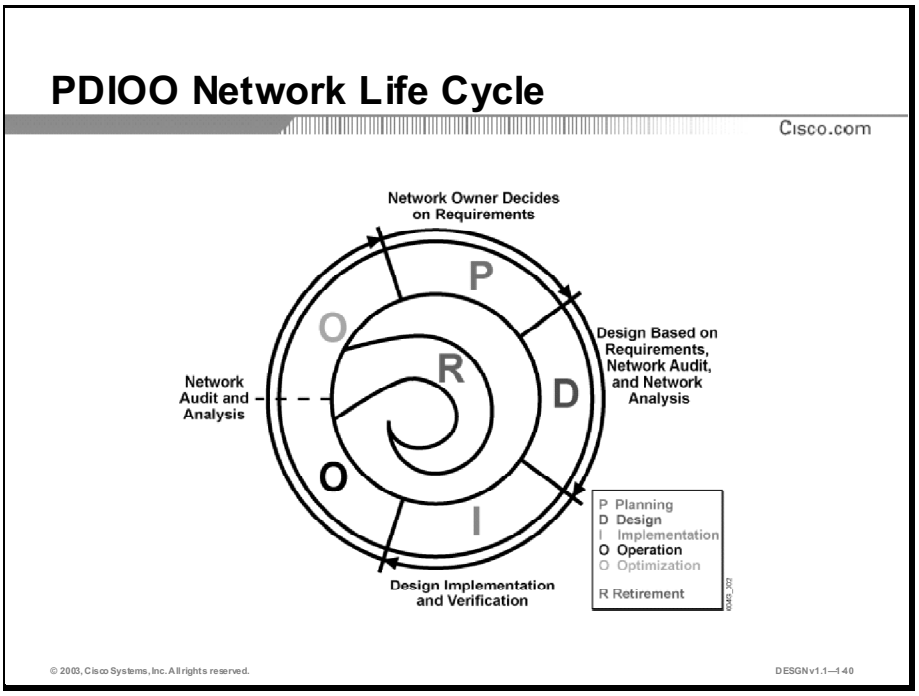

The PDIOO life cycle phases are closely related, and include:

- **Planning:** The network requirements are identified based on goals, facilities, user needs, and so on.
- **Design:** The initial requirements derived in the planning phase drive the activities of the network design specialists. The network design specification produced provides the basis for the implementation activities.
- **Implementation:** After the design has been approved, implementation (and verification) begins. The network is built according to the design specifications.
- **Operation:** Operation is the final test of the appropriateness of the design. The fault detection, correction, and performance monitoring that occur in daily operations provide initial data for the optimization phase.
- **Optimization:** The optimization phase involves proactive management of the network. The goal of proactive management is to identify and resolve issues before they affect the organization. Reactive fault detection and correction (troubleshooting) is needed when proactive management cannot predict and mitigate failures. In the PDIOO process, the optimization phase may prompt a network redesign if too many network problems and errors arise or if performance does not meet expectations.

- **Retirement:** Retirement is not a defined part of the PDIOO process, but is a natural part of the network life cycle. When the network is out of date, the organization can take it out of production. If the features permit, the retired equipment may be redeployed.

Design Methodology

A methodology is a set of procedures, or a documented process, that you can use, without having to re-create the path from the start to the completion of the effort. This topic identifies the steps in the network design methodology.

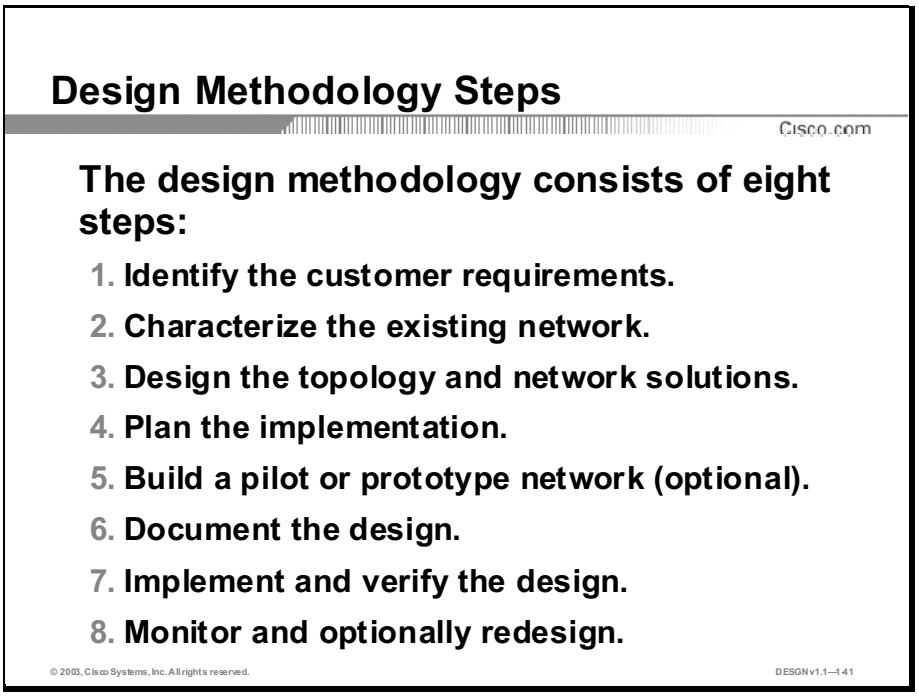

A design methodology:

- Ensures that no step is missed when the process is followed.
- Encourages consistency in the creative process, enabling network designers to set appropriate deadlines and to maintain both customer and manager satisfaction.
- Lets customers and managers validate that there has been thought given to how to meet their requirements.
- Provides a framework for the deliverables of the design process.

The design methodology consists of multiple steps. Some steps are intrinsic to the design phase, while others are related to other PDIOO phases. There are seven mandatory steps and one optional step:

Step 1 **Identify customer requirements:** In this step, key decision makers identify the initial requirements. This is typically done within the PDIOO planning phase.

Step 2 **Characterize the existing network:** Characterization of the existing network includes the network audit and the network analysis. During the network audit, the existing network is thoroughly checked for integrity and quality. During the network analysis, network behavior (traffic, congestion, and so on) is analyzed.

Note:	Complete this step only when an existing network is in place.

Step 3 **Design the topology and network solutions:** In this step, you complete the actual design. Decisions on network infrastructure, intelligent network services, and network solutions (Voice over IP [VoIP], content networking, and so on) are made.

Step 4 **Plan the implementation:** In this step, you prepare the implementation procedures to expedite and clarify the actual implementation. Cost assessment is undertaken at this time.

Step 5 **Build a pilot or prototype network:** In this optional step, you can build a pilot or prototype network to verify the design late in the PDIOO design phase or early in the PDIOO implementation phase.

Step 6 **Document the design:** In this step, you write the actual design documents.

Step 7 **Implement and verify the design:** In this step, by building a network, you verify the actual implementation of the design. This step maps directly to the implementation phase of the PDIOO methodology.

Step 8 **Monitor and optionally redesign:** After the network is built, it is put into operation. Network operators should constantly monitor and check the network for errors. This step is a part of the operation and optimization phases of the PDIOO methodology.

Note: Implement a pilot (prototype) network as often as possible to identify and correct problems that might otherwise lead to redesign later.

Introducing the Requirements-Gathering Process

The design requirements gathering process is composed of six steps as milestones for the designer. You will discuss these steps or milestones with the staff to determine and gather the necessary data and documentation. This topic identifies a process for gathering enterprise network design requirements.

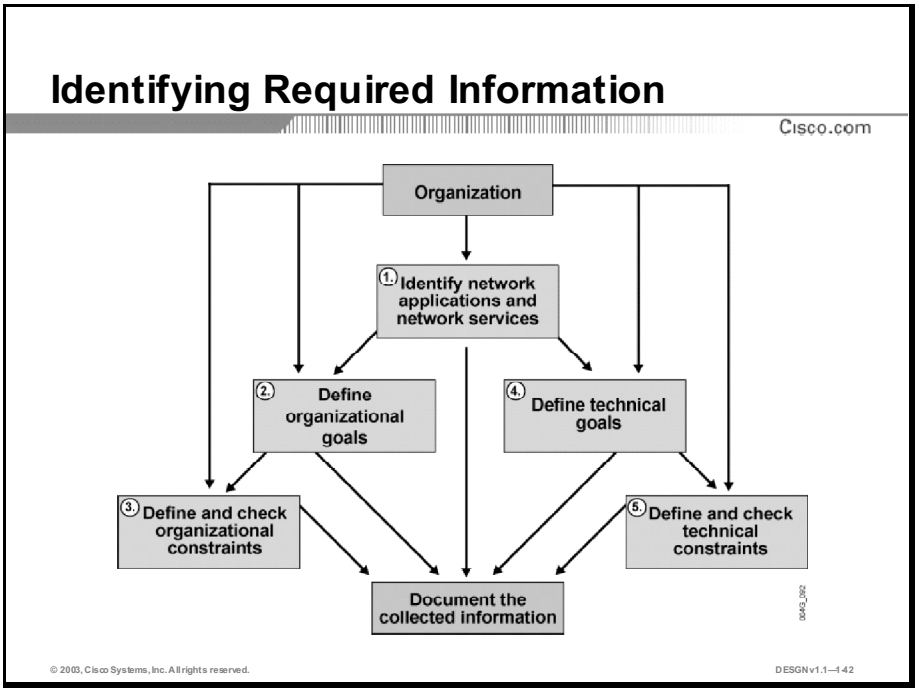

The process is not unidirectional. You might return to a step and make additional inquiries about issues as they arise during the design process.

Step 1 Identify network applications and network services.

Step 2 Define organizational goals.

Step 3 Define and check on all the possible organizational constraints.

Step 4 Define the technical goals.

Step 5 Define and check on all possible technical constraints to consider.

When you complete the data-gathering steps, you will be ready to interpret and analyze the data and develop a design proposal.

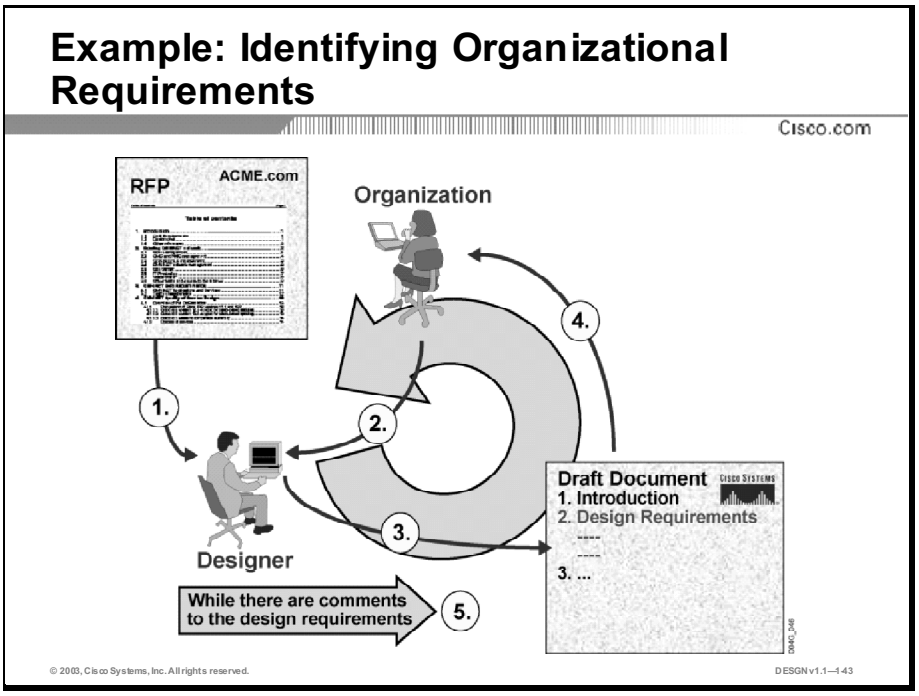

This figure illustrates an iterative approach to developing the final design requirements document.

Step 1 Extract the raw customer requirements (RFP, RFI).

Step 2 Query the customer for raw requirements (verbal description).

Step 3 Produce a draft document that describes the design requirements.

Step 4 Verify the design requirements with the customer for approval.

Step 5 Revise the document as necessary to eliminate errors and omissions.

Identifying Planned Applications and Network Services

A key critical step in data gathering is to determine what applications are planned for use and how important the applications are. The use of a table helps to organize and categorize the solutions for the applications and services planned. This topic helps you identify the planned applications and network services in an intended network.

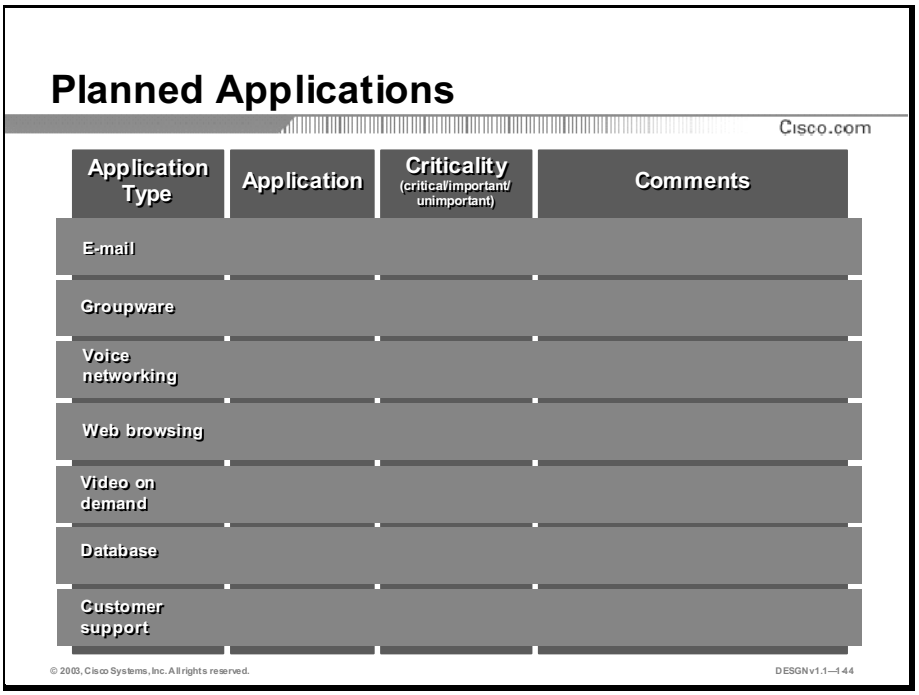

A planned applications table lists:

- Planned application types such as e-mail, groupware (tools that aid group work), voice networking, web browsing, video on demand (VoD), database, file sharing and transfer, and computer-aided manufacturing
- Concrete applications that will be used (for example, Microsoft Internet Explorer, Lotus Notes)
- Importance of certain applications denoted with critical, important, and unimportant keywords
- Additional comments taken in the data gathering process

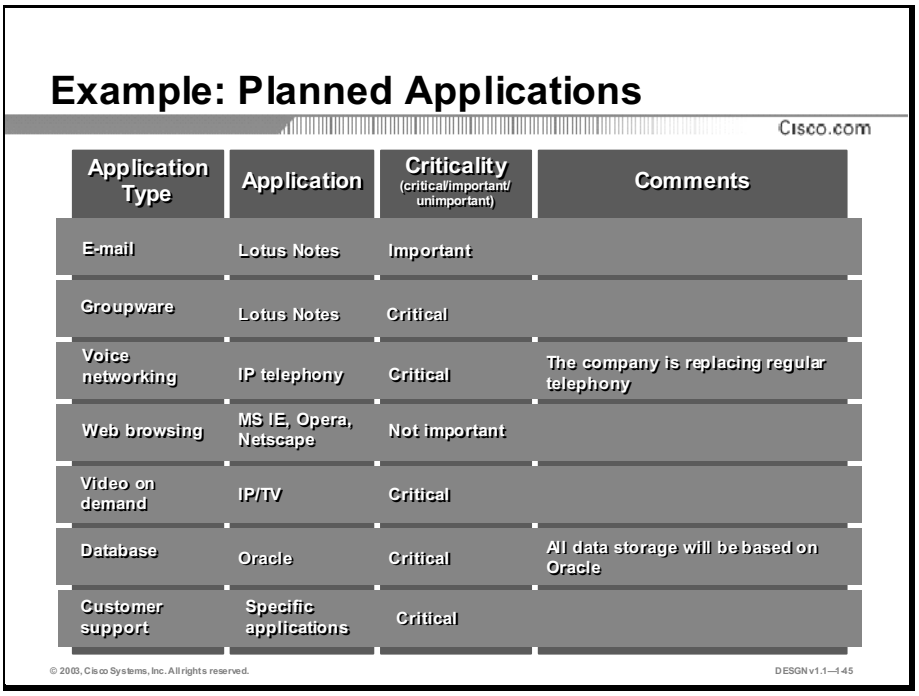

The figure shows gathered data about the planned applications for Corporation X.

Planned Applications Worksheet

Use the table as a template to identify and evaluate planned applications in the course case study and for future design efforts. Remember to include your applications in the Application Type column.

Application Type	Application	Criticality	Comments
E-mail			
Groupware			
Voice networking			
Web browsing			
Video on demand			
Database			
Customer support applications			

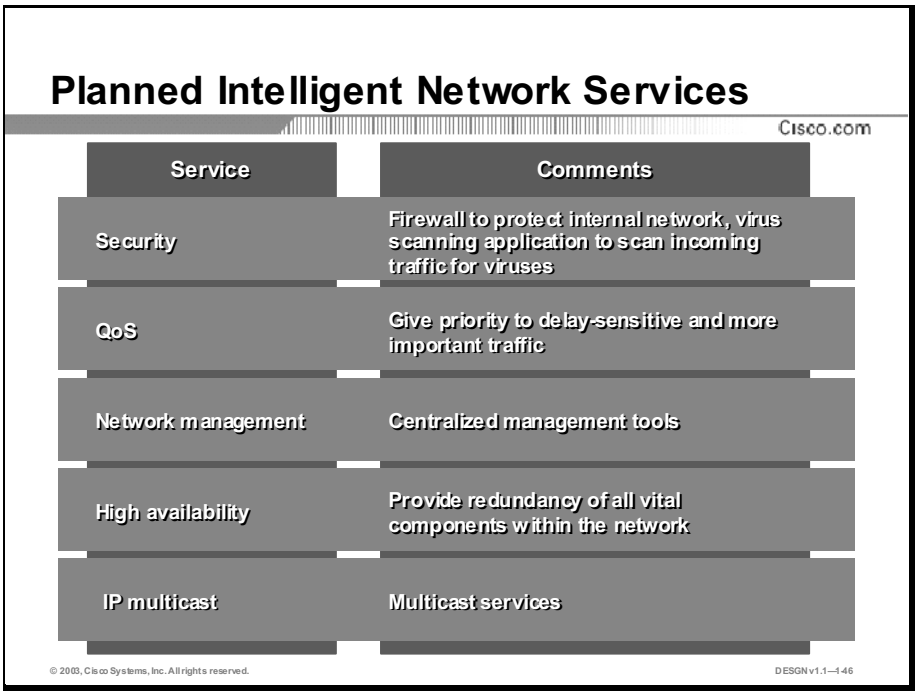

The planned intelligent network services table is similar to the planned applications table. It lists:

- Intelligent network services planned for the network
- Additional comments about those services

Intelligent network services include security, quality of service (QoS), network management, high availability, and IP multicast. Software distribution, backup, directory services, host naming, user authentication, and authorization are examples of general services and solutions deployed to support the applications in a typical organization.

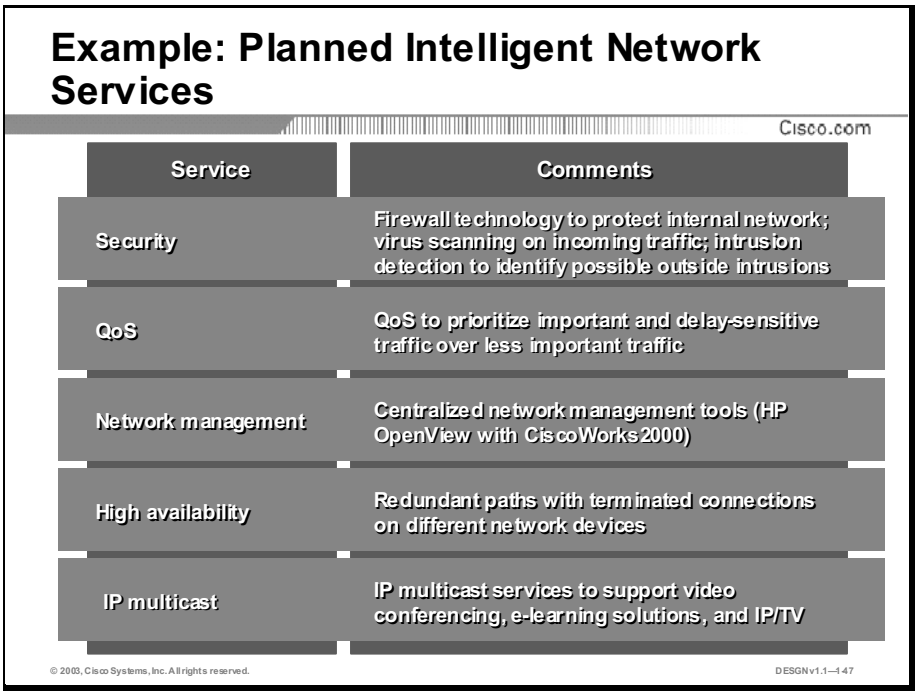

The figure shows data gathered for the intelligent network services planned for Corporation X.

Planned Intelligent Network Services Worksheet

Use the table as a template to identify and evaluate planned intelligent network services in the course case study and for future design efforts. Remember to include your services in the Service column.

Service	Comments
Security	
QoS	
Network management	
High availability	
IP multicast	

Identifying Organizational Goals

An effective network solution will often impact organizational processes. Every design project should consider the organizational goals to be achieved. This topic helps you assess the organizational goals for a network design project.

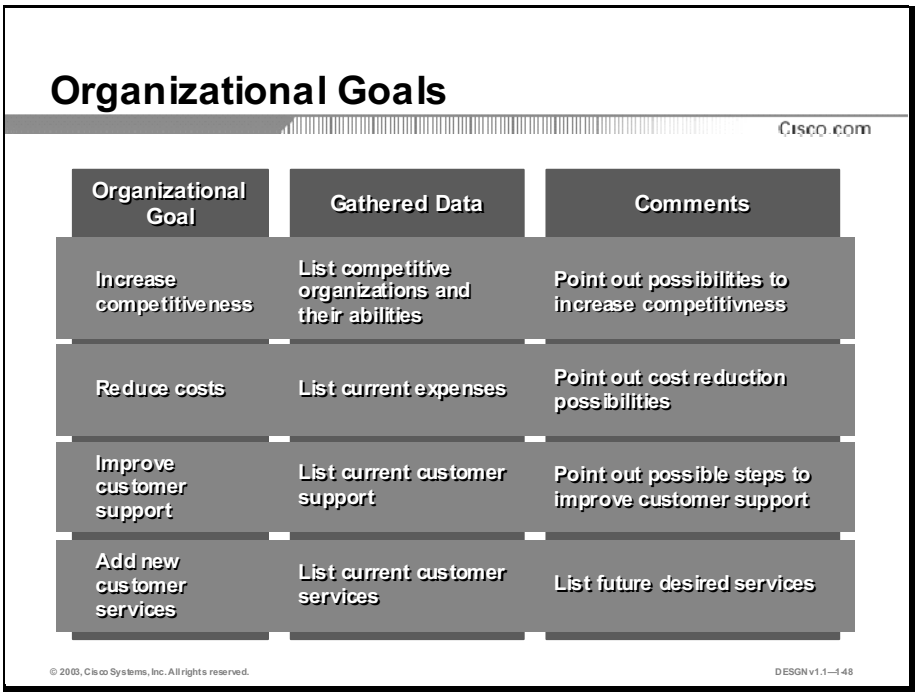

Network designers are often eager to start with the analysis of the technical goals before considering the organizational goals and constraints. However, detailed attention to organizational goals and constraints is important for the success of the project. You will determine the organization's expectations about the outcomes that the design will bring to the organization.

Preliminary research on the organization's activities, products, services, market, suppliers, and competitive advantages, in addition to the organization's structure, enhances the process of positioning the technologies and products to be used in the network.

To identify the organizational goals, you will:

- Identify and discuss the customer requirements, the goals to be achieved, and the purpose of the new network.
- Determine the criteria for success.
- Understand the consequences of failure.

Organizational goals differ from organization to organization. Basic goals that every commercial organization is eager to attain include:

- Increase the generated revenue and the profitability of the organization's operation. A new design should reduce costs in certain segments and propel growth in others. Discuss expectations on how the new network will influence revenues and profits with the customer.

- Improve data availability and interdepartmental communications inside an organization to shorten development cycles and enhance productivity.
- Facilitate customer support and additional customer services that can expedite reaction to customer needs and deliver better customer satisfaction.
- Open the organization information infrastructure to all key constituencies (prospects, investors, customers, partners, suppliers, and employees). Build relationships and information accessibility to a new level as a basis for the network organizational model.

Goals are similar among governmental, charitable, religious, and educational organizations. Most of these entities focus on maximizing the available resources to attain the organization's goals and objectives. In not-for-profit organizations, the key measures are typically stated in terms of cost containment, service quality, service expansion, and resource deployment.

Design requirements are typically stated in a request for proposal (RFP) or request for information (RFI) document from the enterprise. The first step in the design process should be to document the design requirements and gain enterprise verification and approval.

The table of organizational goals and gathered data will help you assess the importance of each organizational goal:

- **Increase competitiveness:** List competitive organizations along with their advantages and weaknesses. Note possible improvements that may increase competitiveness or effectiveness.
- **Reduce costs:** Reducing operational costs can result in increased profitability (the revenue increase is not necessary) or increased services with the same revenue. List current expenses to help determine where to reduce costs.
- **Improve customer support:** Customer support services help gain competitive advantage. List current customer support services with comments about possible and desired improvements.
- **Add new customer services:** List current customer services, and note future and desired (requested) services.

The figure provides an example of data gathered about the organizational goals of Corporation X.

Organizational Goals Worksheet

Use the table as a template to identify organizational goals in the course case study and for future design efforts.

Goal	Data	Comments

Assessing Organizational Constraints

When assessing organizational goals, it is important to analyze any organizational constraints that may affect the network design. This topic helps you identify organizational constraints that affect the network design.

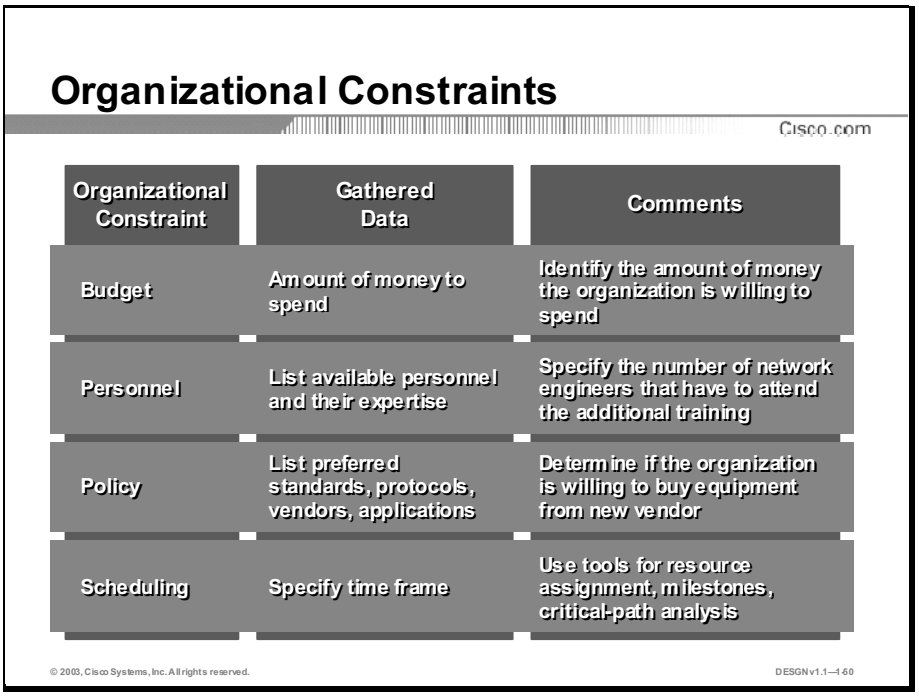

The figure presents a sample table of typical organizational restraints:

- **Budget:** Reduced budgets or limited resources often force network designers to implement an affordable network design that compromises availability, manageability, performance, and scalability. The budget includes all equipment purchases, software licenses, maintenance agreements, staff training, and so on. You must know the budget available to invest on a solid design. It is also useful to know the areas in which to compromise network performance to meet budget requirements.
- **Personnel:** Availability of trained personnel within the organization may be a design consideration. Organizations may not have either trained personnel or enough personnel. Additional constraints may be imposed if the organization is outsourcing network management. Trained technicians should verify that all elements are working in concert and are recognized on the network. You should consider implementation and operation staff. You must know the number and availability of operations personnel, their expertise, and possible training requirements.
- **Policies:** Organizations have different policies regarding protocols, standards, vendors, and applications. You must understand these policies to design the network successfully. Determine customer policies related to single or multivendor platforms.
- **Scheduling:** The new network design is often driven by the introduction of new network applications. The implementation time frames for new applications are often tightly connected and therefore influence the time available for network design. You should discuss the project time frame with the organization's executive management and gain approval.

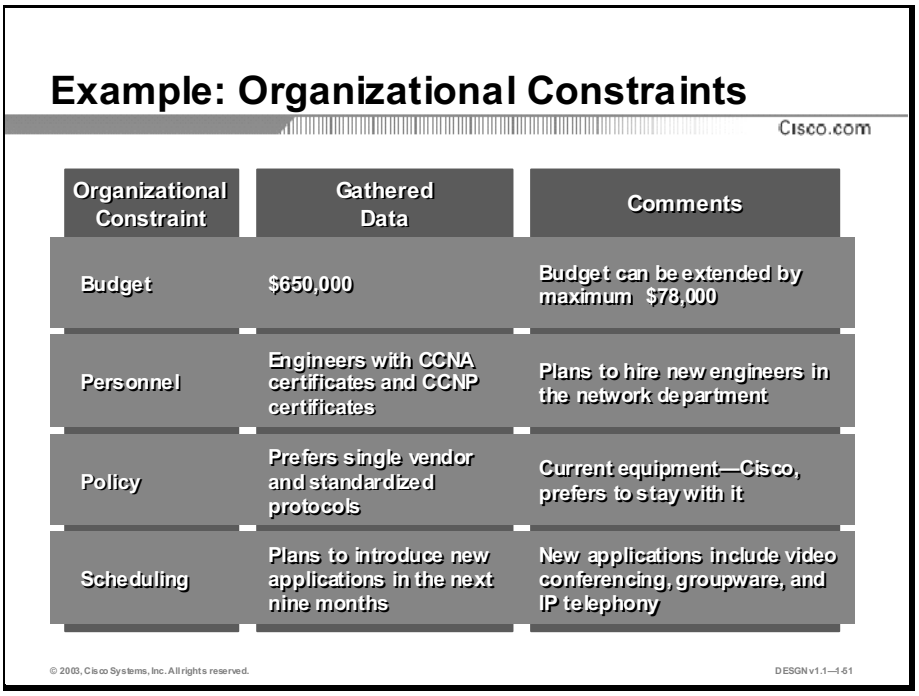

The figure presents gathered organizational constraints and accompanying data for Corporation X.

Organizational Constraints Worksheet

Use the table as a template to help determine organizational constraints in the course case study and for future design efforts. Remember to include your constraints in the Constraint column.

Constraint	Data	Comments
Budget		
Personnel		
Policy		
Scheduling		

Identifying Technical Goals

As the organization network grows, so does dependency on the network and the applications that utilize it. Network-accessible organization data and mission-critical applications that are essential to the organizational operations of the organization depend on network availability. This topic helps you assess technical goals for the network design project.

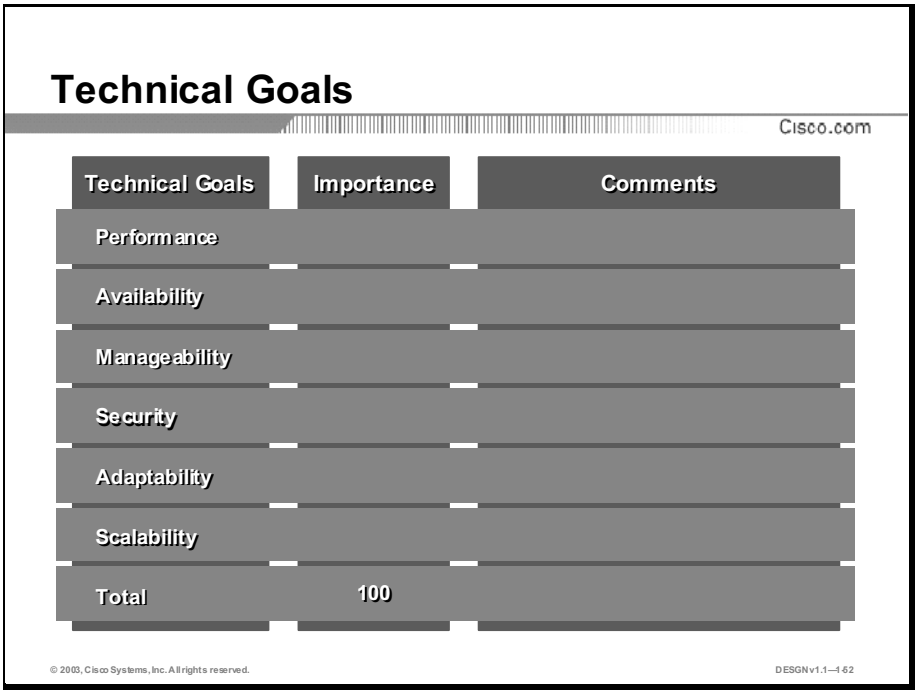

This list describes common technical goals:

- **Improve performance of the network:** Responsiveness and throughput degrade as the number of users and applications increase. The first goal of network redesign is to increase performance by upgrading link speed, partitioning the network, or both.
- **Improve security and reliability of mission-critical applications and data:** Increased threats from both inside and outside the enterprise network require the most up-to-date security rules and technologies to avoid disruptions of network operation.
- **Decrease the expected downtime and related expenses:** When a network failure occurs, downtime must be minimal, and the network must respond quickly to minimize related costs.
- **Modernize outdated technologies:** New network technologies and applications demand continued modernization of equipment and technologies.
- **Improve scalability of the network:** Networks must allow for upgrades and future growth.
- **Simplify the network management:** Simplify network management for better understanding and use.

Use a table to help identify technical goals. Different goals have different levels of importance, which each organization should determine. The sum of the individual percentages should be 100, thus providing direction when choosing equipment, protocols, features, and so on.

Typical technical goals include performance, availability, manageability, security, adaptability, and scalability.

Note: Performance is a general term that includes responsiveness, throughput, and resource utilization. Networked application users and their managers are usually most sensitive to responsiveness issues; speed is critical. Network managers often look to throughput as a measure of effectiveness. Executives with capital budget responsibility often evaluate resource utilization as a measure of economic efficiency. In presenting performance information, carefully consider the audience.

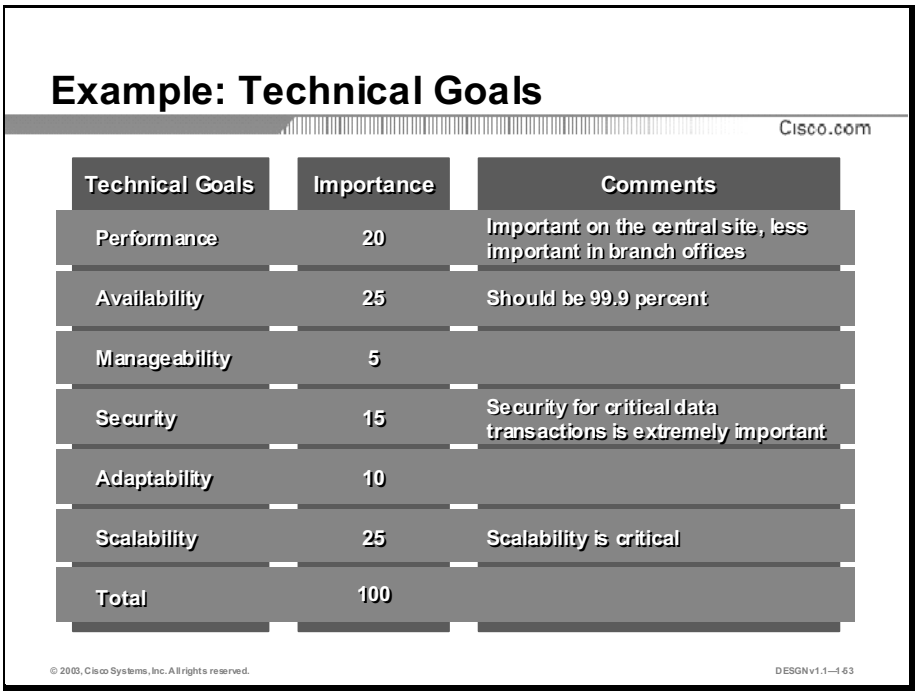

The figure displays desired technical goals gathered for Corporation X, along with their importance rating and additional comments. Specific technical goals are rated in importance on a scale from 1 to 100, with the sum totaling 100.

In the example, Corporation X gives great importance to availability, scalability, and performance. These requirements suggest the need for redundant equipment, redundant paths, high-speed links, and so on.

Technical Goals Worksheet

Use the table as a template to identify and evaluate technical goals in the course case study and for future design efforts. Remember to include your goals in the Technical Goals column.

Technical Goals	Importance	Comments
Performance		
Availability		
Manageability		
Security		
Adaptability		
Scalability		

Assessing Technical Constraints

Network designers may face various technical constraints during the design process. Good network design addresses constraints by identifying possible tradeoffs. This topic helps you identify technical constraints that affect the network design.

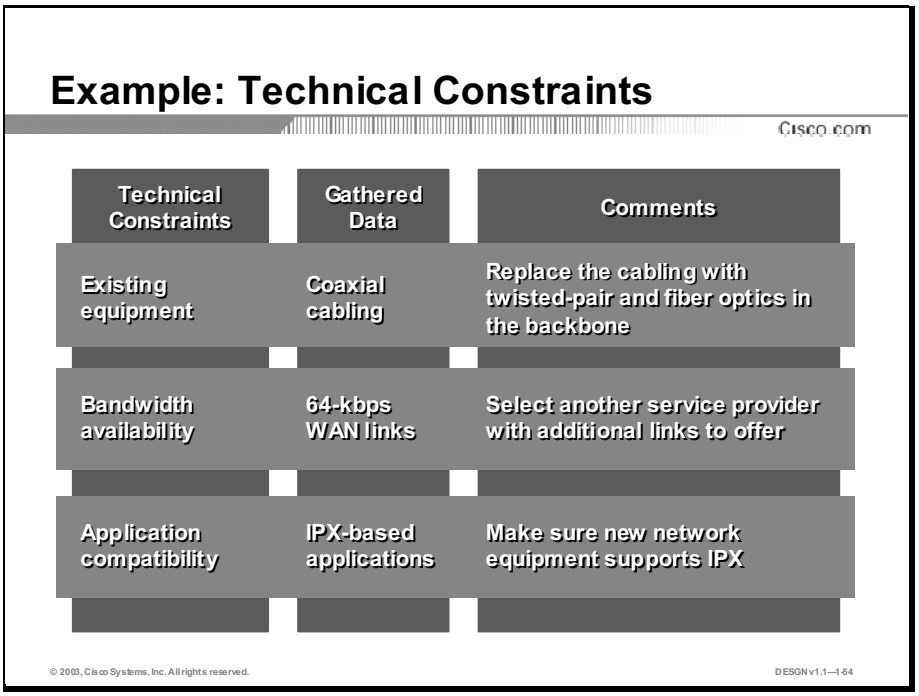

Using a table can facilitate the process of gathering technical constraints. Identify the technical constraints (for example, existing equipment, bandwidth availability, application compatibility) and then note the current situation and the necessary changes required to mitigate a certain constraint.

Consider the following technical constraints:

- **Existing equipment:** The network design process is usually progressive; legacy equipment must coexist with new equipment.
- **Bandwidth availability:** Insufficient bandwidth in parts of the network, where the bandwidth cannot be increased due to technical constraints, must be resolved by other means.
- **Application compatibility:** If new applications are not being introduced at the same time as the new network, the design must provide compatibility with old applications.

Example: Technical Constraints

The figure presents technical constraints for Corporation X. The designer notes that coaxial cabling still exists and suggests replacing the cabling with twisted-pair and fiber optics. The bandwidth availability indicates that a change of service providers may be warranted. Application compatibility suggests that the designer should take care when choosing equipment.

Technical Constraints Worksheet

Use the table as a template to identify and evaluate technical constraints in the course case study and for future design efforts.

Technical Constraints	Gathered Data	Comments
Existing equipment		
Bandwidth availability		
Application compatibility		

Summary

This topic summarizes the key points discussed in this lesson.

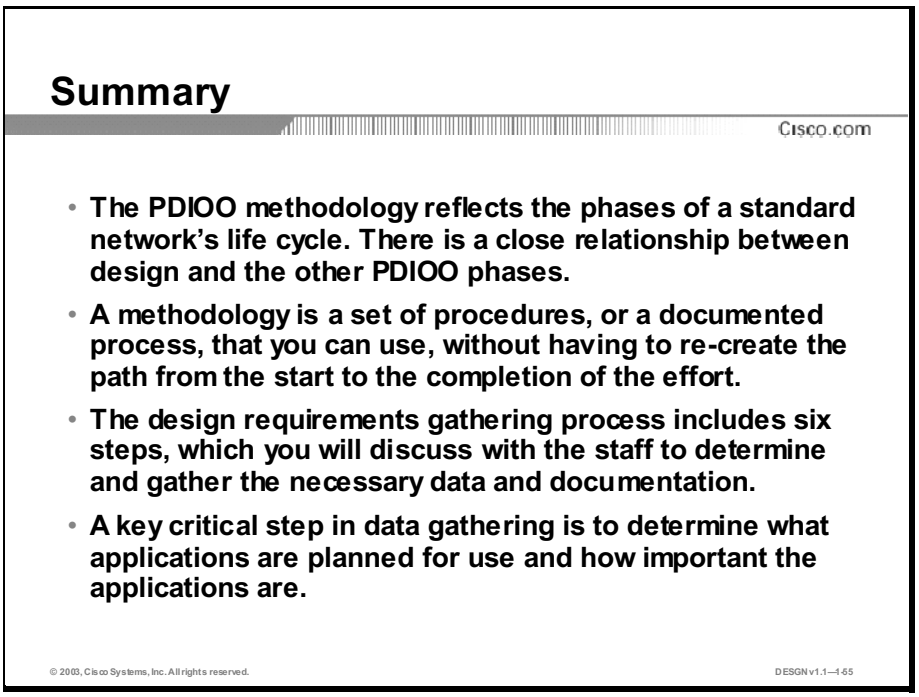

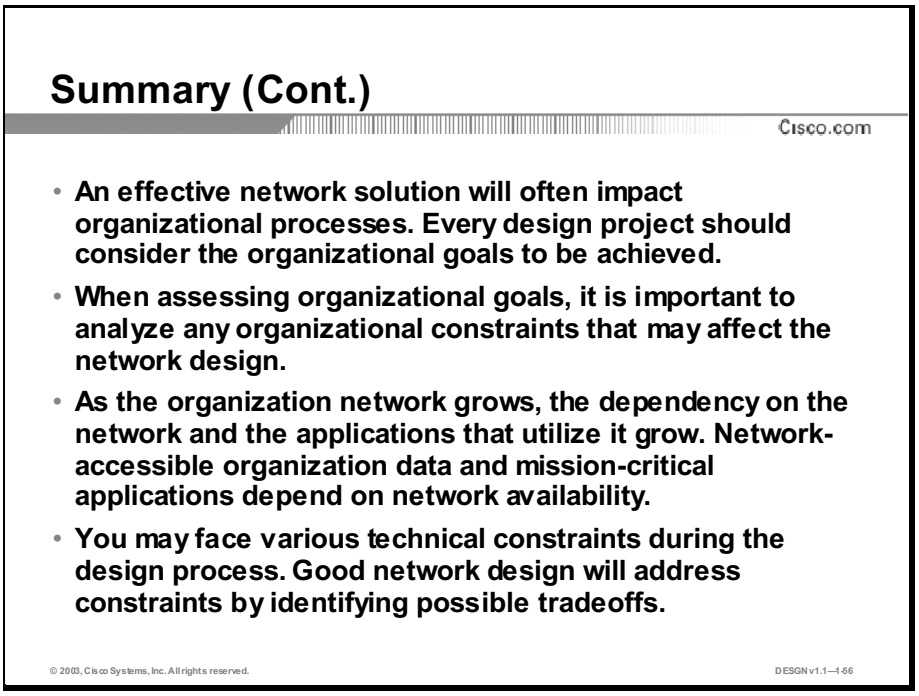

References

For additional information, refer to this resource:

- Oppenheimer, P. *Top-Down Network Design: A Systems Analysis Approach to Enterprise Network Design*. Indianapolis, Indiana: Macmillan Technical Publishing—Cisco Press; 1999.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

The existing network is outdated and the implemented redundancy is unsatisfactory. The WAN connections have become a bottleneck.

Q1) The planning phase requirements have a direct influence on the _____ phase decisions.

A) retirement

B) design

C) optimization

D) test

Q2) The _____ phase is important for design verification.

A) design

B) planning

C) building

D) implementation

Q3) Match each design methodology step in the lettered list with the correct procedure description in the numbered list.

A) extracting customer requirements

B) characterizing the existing network

C) designing topology and network solutions

D) building a pilot network

E) planning the implementation

F) documenting the design

G) implementing and verifying the design

H) monitoring and optional redesigning

_____ 1. build a prototype network

_____ 2. build a network, verify design

_____ 3. verify the health of the network

_____ 4. develop design documents

_____ 5. perform network audit and analysis

_____ 6. identify initial design requirements

_____ 7. plan implementation steps, assess costs

_____ 8. create design and documentation, plan implementation

Q4) Which three pieces of information must you identify before the network design can commence? (Choose three.)

A) organizational and technical goals

B) technical constraints

C) existing and new network applications

D) expected revenue growth

E) ROI

Q5) Select two typically planned applications for an organization. (Choose two.)

A) e-mail

B) groupware

C) video on demand

D) IP telephony

E) QoS

Q6) Which three items are considered organizational goals? (Choose three.)

A) increase competitiveness

B) reduce costs

C) determine budget

D) improve customer support

E) improve QoS

Q7) What is considered to be an organizational constraint?

A) budget

B) planned applications

C) technical goals

D) legacy equipment

Corporation X is planning to introduce e-learning for its employees. Videoconferencing will be its next step in facilitating organizational meetings. The company is looking for an alternative telephony service to reduce their operational costs.

Q8) Which three are considered to be technical goals? (Choose three.)

A) high security

B) ease of management

C) available budget

D) higher reliability

E) increased revenue

F) facilitated customer support

Q9) Which three are considered to be technical constraints? (Choose three.)

A) available budget

B) existing equipment

C) company policy

D) bandwidth availability

E) application compatibility

F) scheduling

Quiz Answer Key

Q1) B

Relates to: Design as an Integral Part of the PDIOO Methodology

Q2) D

Relates to: Design as an Integral Part of the PDIOO Methodology

Q3) 1=D, 2=G, 3=B, 4=F, 5=H, 6=A 7=E, 8=C

Relates to: Design Methodology

Q4) A, B, C

Relates to: Introducing the Requirements-Gathering Process

Q5) C, D

Relates to: Identifying Planned Applications and Network Services

Q6) A, B, D

Relates to: Identifying Organizational Goals

Q7) A

Relates to: Assessing Organizational Constraints

Q8) A, B, D

Relates to: Identifying Technical Goals

Q9) B, D, E

Relates to: Assessing Technical Constraints

Characterizing the Existing Network

Overview

In many organizations, a network already exists and the new design requires restructuring and upgrading the existing network. When a network already exists, you should examine the existing network. You will map the existing network topology and audit and measure network traffic using a variety of available tools. Then you will analyze the data. The lesson concludes with guidelines for creating a summary report that describes the network health, which is essential for any successful redesign solution.

Relevance

Characterizing an existing network will provide the information you need to redesign the network.

Objectives

Upon completing this lesson, you will be able to characterize an existing network. This includes being able to meet these objectives:

- Identify the information you need to collect about the existing network infrastructure
- Identify major features of the existing network
- List the tools that help in auditing the existing network
- Identify the existing network traffic and applications
- List the tools for monitoring and analyzing network traffic
- Analyze the existing network health based on audit and traffic measurements

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Understanding of network monitoring and troubleshooting concepts and tools

Outline

The outline lists the topics included in this lesson.

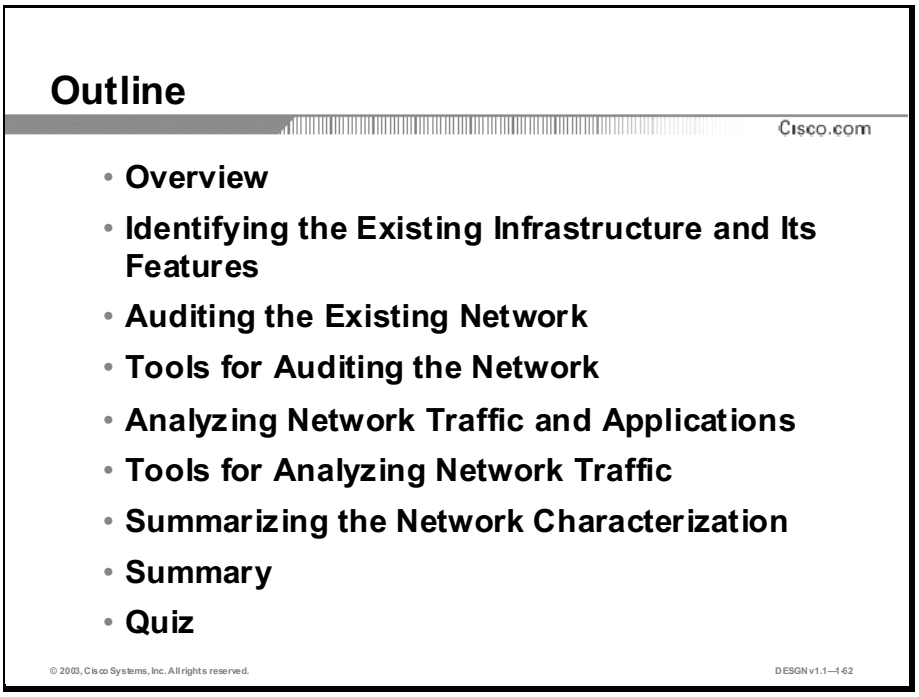

Identifying the Existing Infrastructure and Its Features

The first step in characterizing an existing network is to gather as much information about the network as possible. Organization input, a network audit, and traffic analysis provide the key information you need. This topic identifies the information you need to collect about the existing network infrastructure.

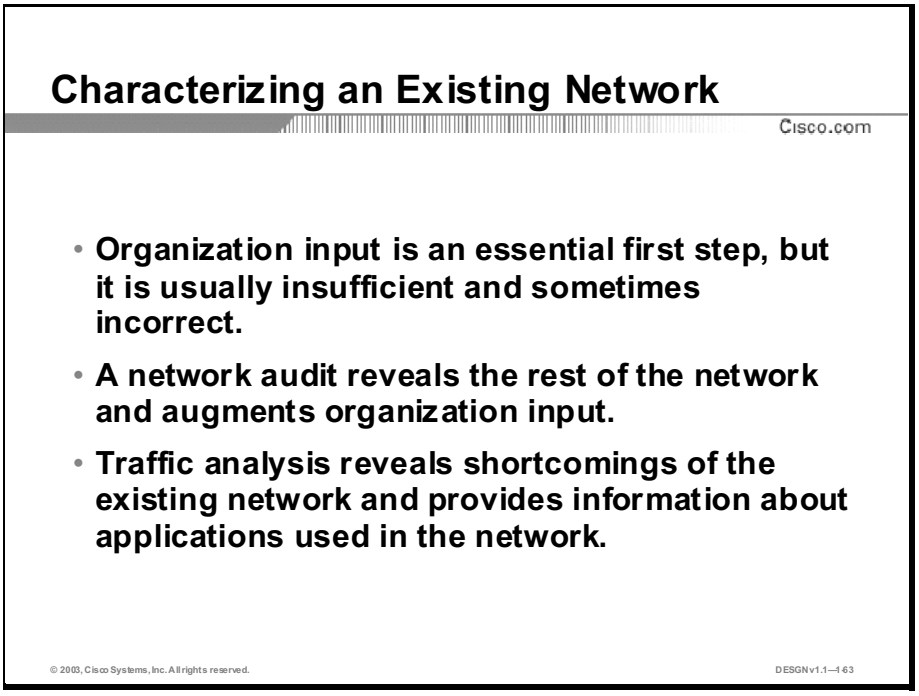

The characterization of a network is typically based on these steps:

Step 1 Review existing documentation about the network, and use verbal input from the organization to gain a first impression about the network.

Step 2 Perform a network audit that adds detail to the description of the network.

Step 3 Use traffic analysis information to augment organizational input when describing the applications and protocols used in the network.

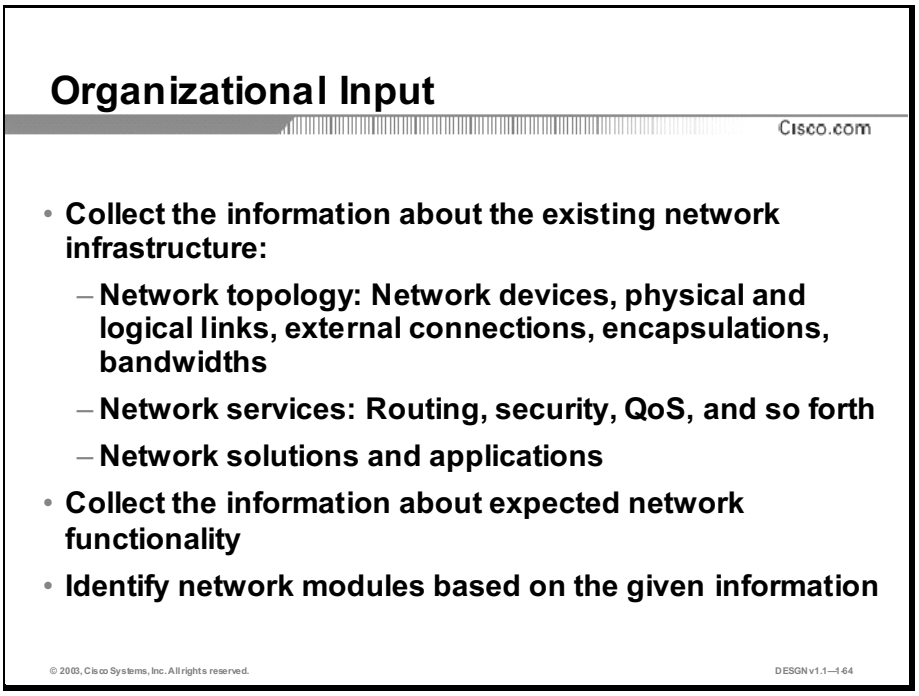

Use initial data to draft a document that describes the existing network infrastructure. This document should cover these topics:

- Existing network infrastructure
- Expected network functionality
- Network topology
- Network services
- Network solutions
- Network applications

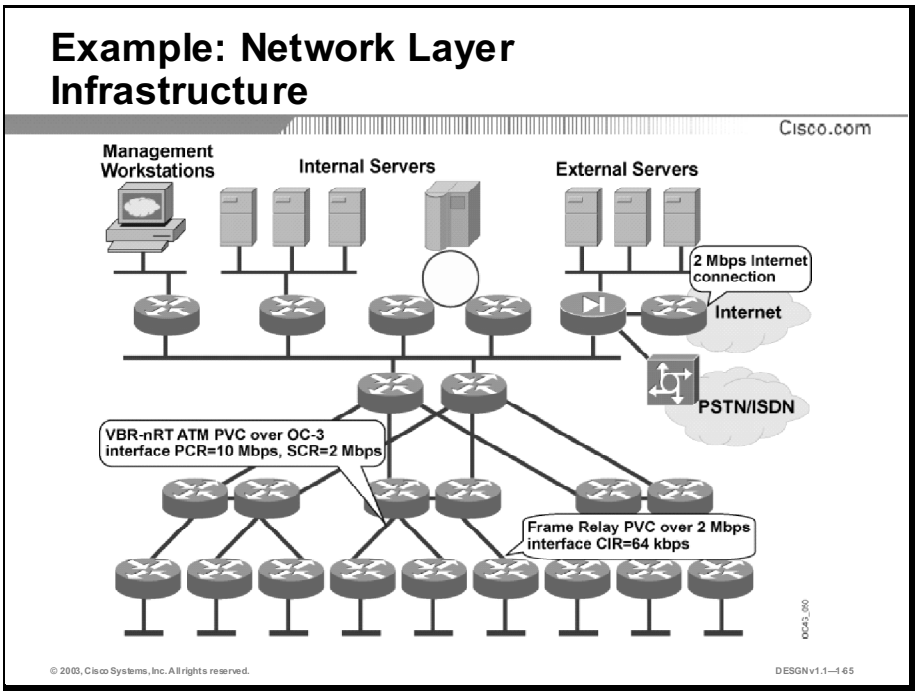

The figure presents the OSI reference model network layer (Layer 3) topology of an example wide-area network (WAN).

Based on the OSI reference model network layer topology, two redesign issues are raised:

- Routing design:
  - The topology of the entire network is relevant (OSI Layers 1, 2, and 3).
- Firewall design:
  - The topology around the Internet connection is relevant.
  - The internal servers and applications that should be accessible from the Internet are relevant.
  - The requirements for outbound connections are relevant.

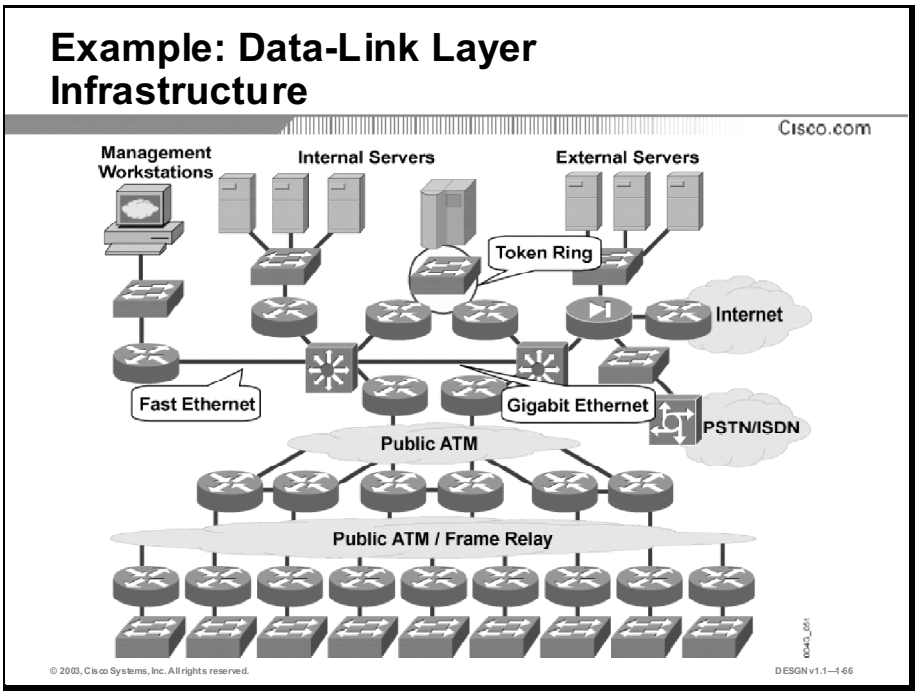

The figure illustrates the OSI data link layer (Layer 2) topology of an example network.

The OSI Layer 2 map reveals more network devices. For some designs, it is important to include the description of OSI Layer 2/1 topology. The figure reveals LAN devices in addition to the interfaces connected to a public WAN. However, the map hides the logical (Layer 3) links between the routers.

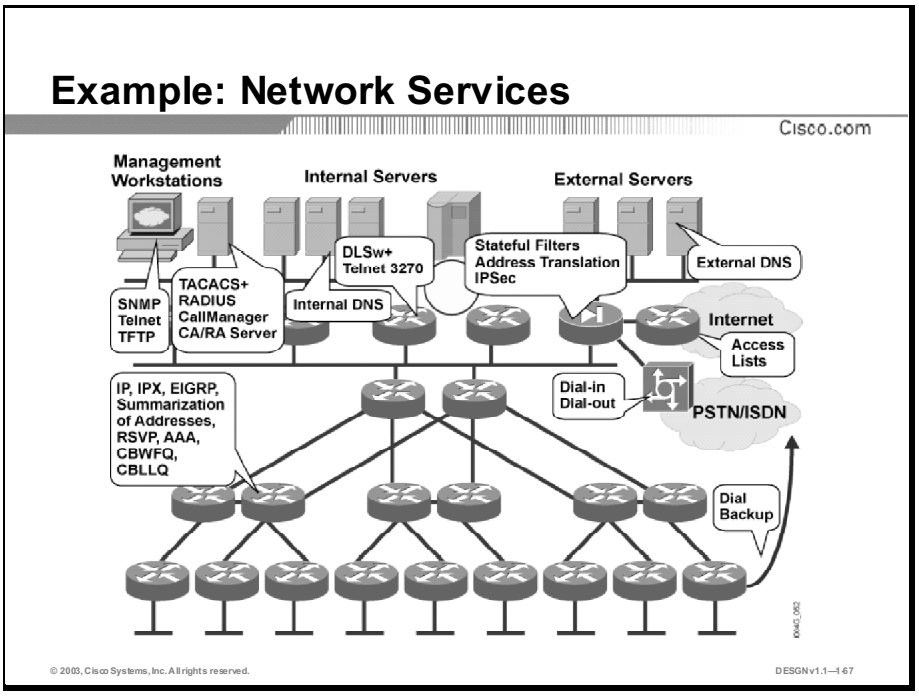

The figure illustrates the network services in an example network.

The detailed list of network services that is used in the network describes the protocols and major flows through the network. This information provides a necessary foundation for the development of a good design.

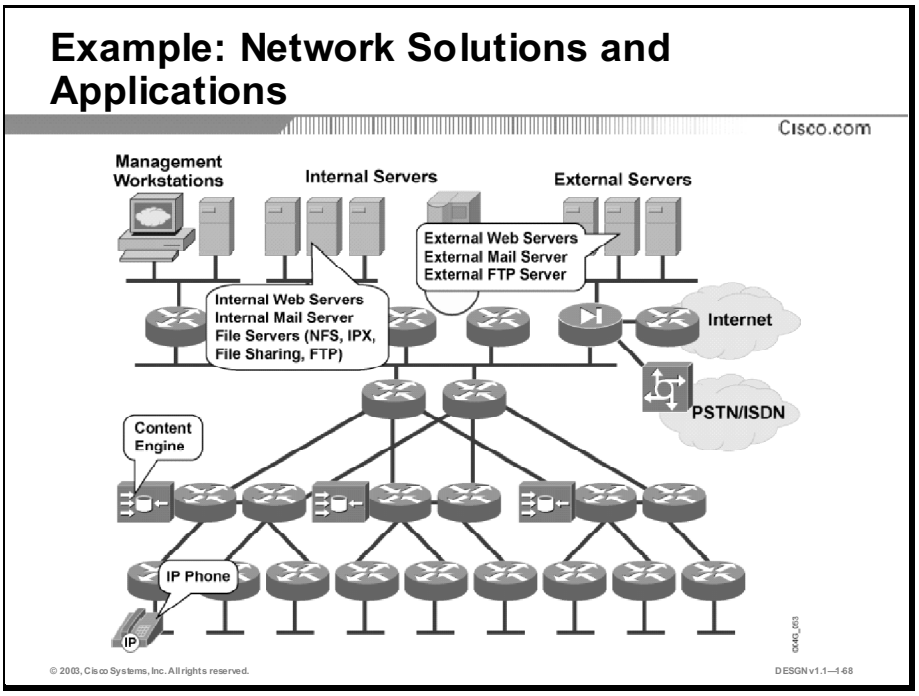

The figure illustrates the names of the services and the applications in the example network.

You should collect detailed descriptions of the protocols that these applications and services use.

Auditing the Existing Network

The auditing process adds detail to the initial network documentation that you created from existing documentation and organizational input. This topic helps you identify the major features of the existing network.

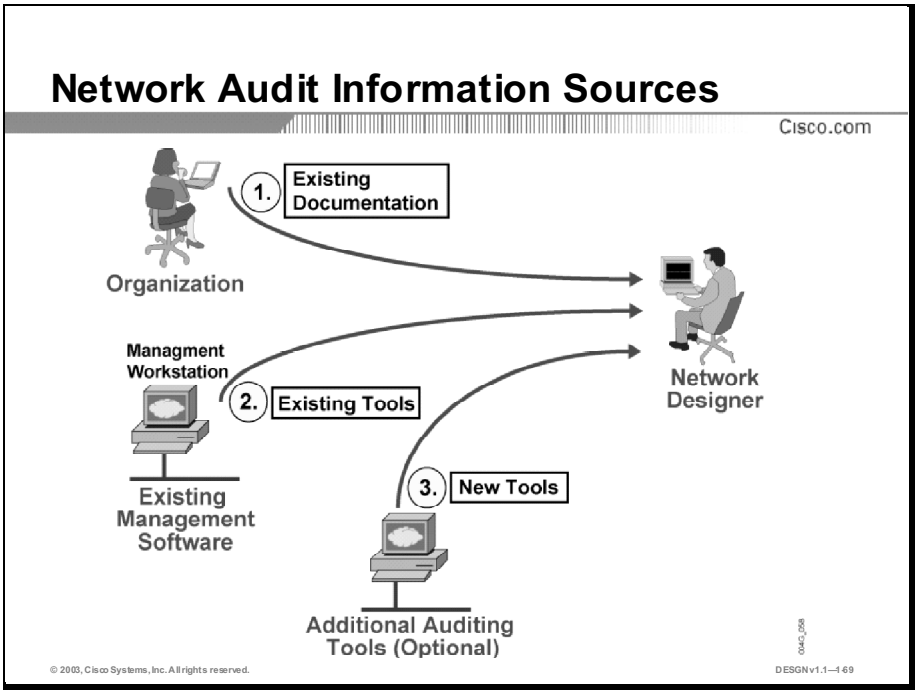

The figure illustrates three different sources of information that you can use in the auditing process.

The auditing process starts by consolidating existing information about the network. You can gather up-to-date information from the existing management software. If the organization has insufficient tools, you may choose to introduce additional software tools temporarily, or even permanently, if they prove useful.

The audit provides details such as the following:

- List of network devices
- Hardware specifications and software versions of network devices
- Configurations of network devices
- Output of various auditing tools to verify and augment the existing documentation
- Link, CPU, and memory uses of network devices
- Unused ports, modules, and slots in network devices

The audit process balances both detail and effort to produce as much information as needed or possible. It should not require that a large set of CPU-heavy auditing tools be purchased and installed in the organizational network to collect configurations of network devices.

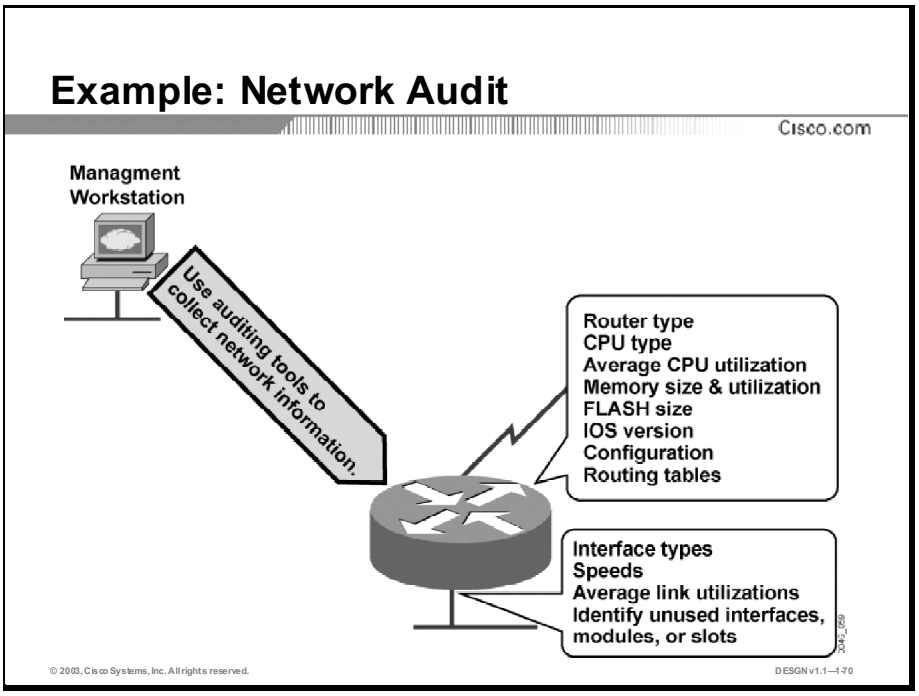

You typically perform the auditing process from a central location in a secure environment that is allowed to access all network devices.

The figure illustrates how either a manual or automated auditing process collects information from the network management workstation. The auditing process should collect all information relevant to the redesign. Use the same process for all devices in the network affected by the design.

Tools for Auditing the Network

You can manually audit a small network, and use tools to audit a large network. This topic lists the tools that help in auditing the existing network.

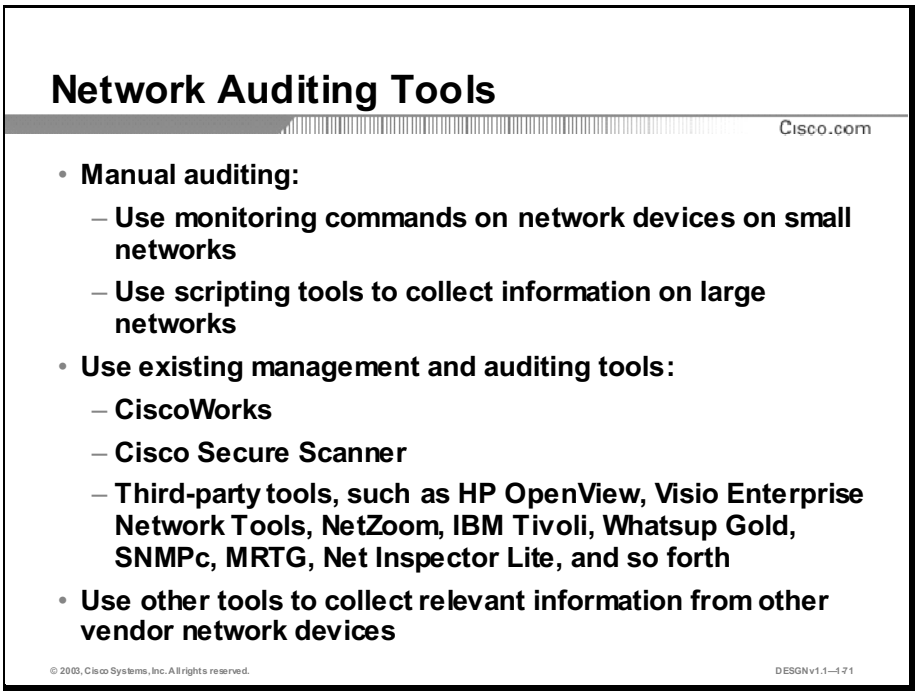

You can audit a small network without special tools. You can use monitoring commands on a small number of network devices to collect the relevant information. You can partially automate the approach by using scripting tools to execute the monitoring commands automatically.

In large networks, a manual auditing approach is too time consuming and unreliable. You can use these special tools to collect the relevant information from the network devices:

- CiscoWorks to map a network and collect different types of information (for example, network topology, hardware and software versions, configurations, and so on)
- Cisco Secure Scanner to find security vulnerabilities
- Third-party tools (for example, HP OpenView, Visio Enterprise Network Tools, NetZoom, IBM Tivoli, Whatsup Gold, SNMPc, MRTG, Net Inspector Lite, and so on)

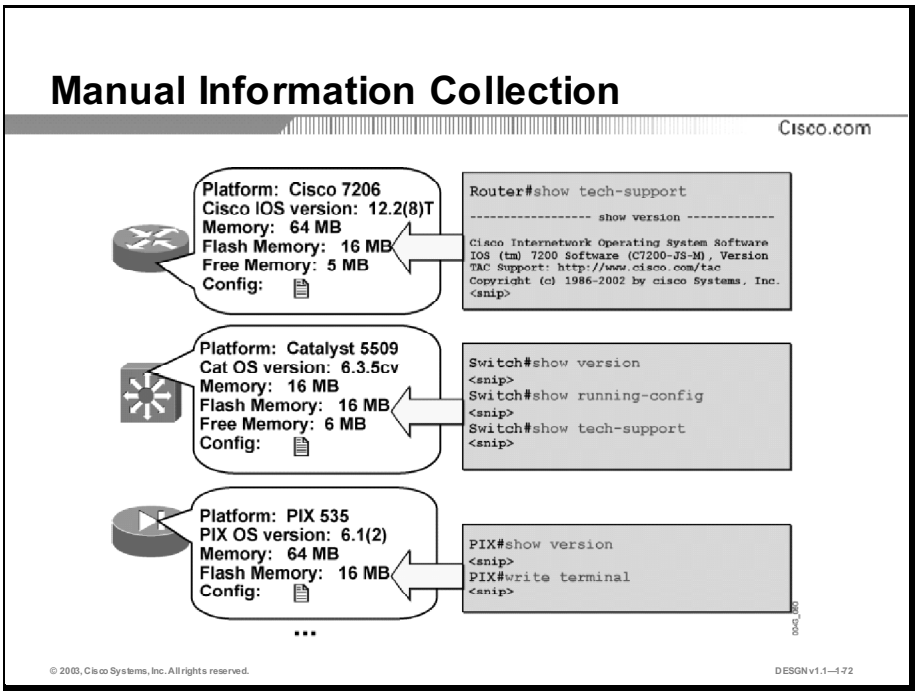

You can perform the auditing process manually by using various monitoring commands. Use this option only with relatively small networks.

The figure illustrates three different types of network devices, the information to collect, and commands used to obtain the information:

- On Cisco routers running Cisco IOS software, the **show tech-support** command displays all information about the router. Use the **show processes cpu** command to determine CPU use and the **show processes memory** command to view memory usage.
- On Cisco switches running Catalyst software, the most useful commands vary, depending on the version of the software (for example, use **show version**, **show running-config,** or **show tech-support**, if available).
- On Cisco Secure PIX firewalls, a printout of the configuration is usually needed (for example, use **show version** and **write terminal**).

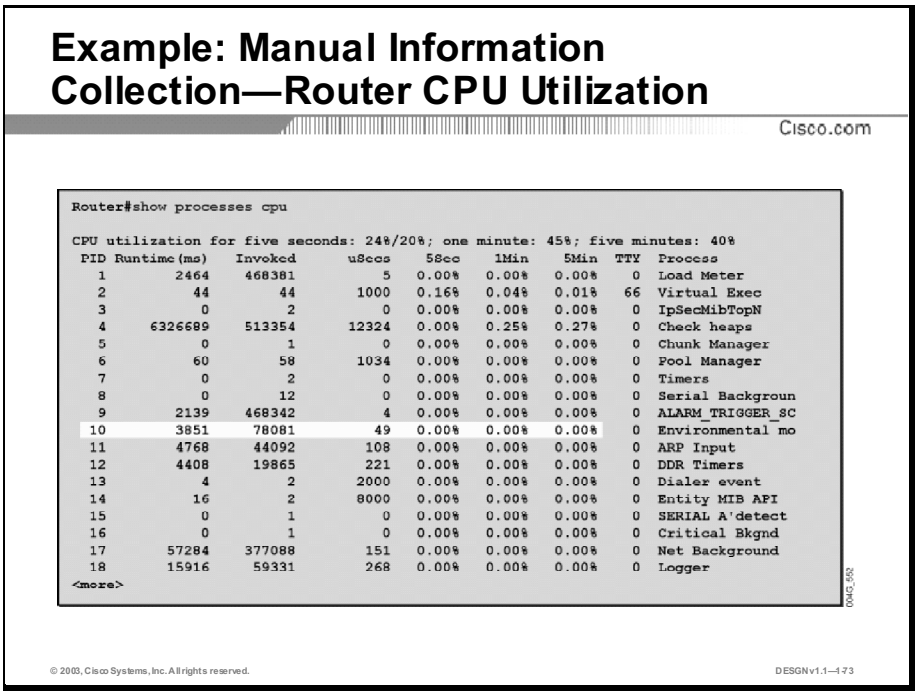

The figure illustrates the printout from the **show processes cpu** command.

The printout provides information about the network device CPU utilization, which is useful in describing the network health.

The table describes the fields displayed in the **show processes cpu** command output.

Field	Description
CPU utilization for five seconds	CPU utilization for the last five seconds. The first number indicates the total; the second number indicates the percent of CPU time spent at the interrupt level.
one minute	CPU utilization for the last minute
five minutes	CPU utilization for the last five minutes
PID	Process ID
Runtime (ms)	CPU time the process has used, expressed in milliseconds
Invoked	Number of times the process has been invoked
uSecs	Microseconds of CPU time for each process invocation
5Sec	CPU utilization by task in the last five seconds
1Min	CPU utilization by task in the last minute
5Min	CPU utilization by task in the last five minutes
TTY	Terminal that controls the process
Process	Name of the process

Example: Manual Information Collection—Router Memory Utilization

Cisco.com

```
Router#show process memory
Total: 26859400, Used: 8974380, Free: 17885020
  PID TTY  Allocated      Freed    Holding    Getbufs    Retbufs Process
    0   0      88464       1848    6169940          0          0 *Init*
    0   0        428    1987364        428          0          0 *Sched*
    0   0  116119836  105508736     487908     373944      55296 *Dead*
    1   0        284        284       3868          0          0 Load Meter
    2  66       5340       1080      17128          0          0 Virtual Exec
    3   0        668        284       7252          0          0 IpSecMibTopN
    4   0          0          0       6868          0          0 Check heaps
    5   0         96          0       6964          0          0 Chunk Manager
    6   0      17420     231276       6964       5388     254912 Pool Manager
    7   0        284        284       6868          0          0 Timers
    8   0        284        284       6868          0          0 Serial Backgroun
    9   0          0          0       6868          0          0 ALARM_TRIGGER_SC
   10   0        284        284       6868          0          0 Environmental mo
   11   0        316    3799360       7184          0          0 ARP Input
   12   0    2547784    1033916       7372       6804          0 DDR Timers
   13   0        284        284      12868          0          0 Dialer event
   14   0      10744       2284      15328          0          0 Entity MIB API
   15   0         96          0       6964          0          0 SERIAL A'detect
   16   0         96          0       6964          0          0 Critical Bkgnd
   17   0      23412       2632      15404          0          0 Net Background
<more>
```

0043_503

The figure illustrates the printout from the **show processes memory** command.

The table describes the fields in the **show processes memory** command output.

Field	Description
Total	Total amount of memory held
Used	Total amount of used memory
Free	Total amount of free memory
PID	Process ID
TTY	Terminal that controls the process
Allocated	Bytes of memory allocated by the process
Freed	Bytes of memory freed by the process, regardless of who originally allocated it
Holding	Amount of memory currently allocated to the process
Getbufs	Number of times the process has requested a packet buffer
Retbufs	Number of times the process has relinquished a packet buffer
Process	Process name
Total	Total amount of memory held by all processes

Analyzing Network Traffic and Applications

Traffic analysis verifies the set of applications and protocols used in the network and determines the traffic patterns of the applications. This topic helps you identify the existing network traffic and applications.

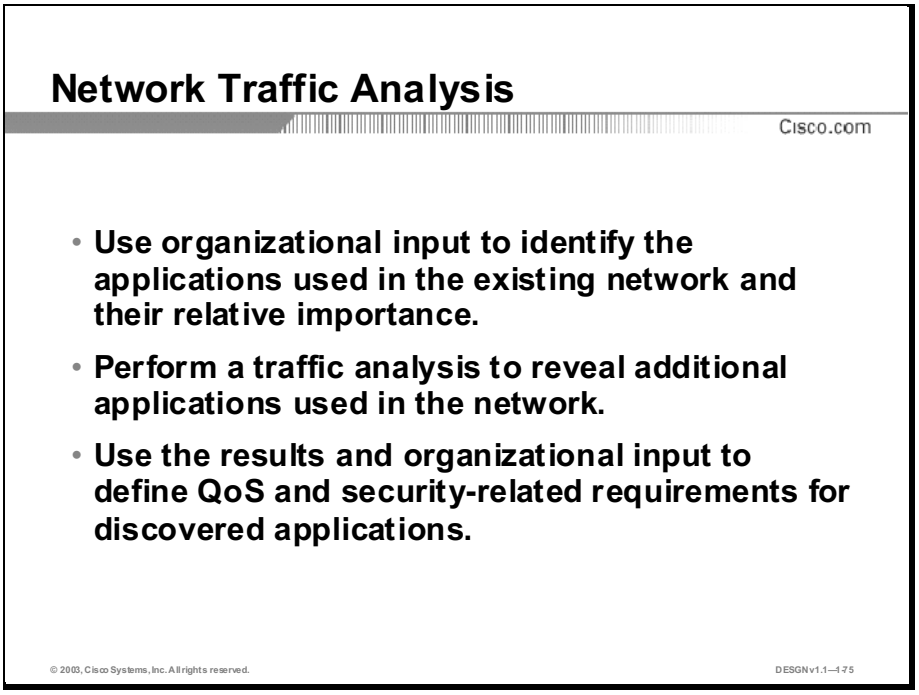

An organization can identify current and planned applications. However, a traffic analysis is important to reveal additional applications that can affect network performance.

You can perform a traffic analysis by using services available in Cisco routers such as NetFlow or by using dedicated hardware- or software-based analyzers that can be moved around the network.

Describe each application in terms of:

- Importance
- QoS requirements
- Security requirements
- Scope (describing in which network modules an application or protocol is used)

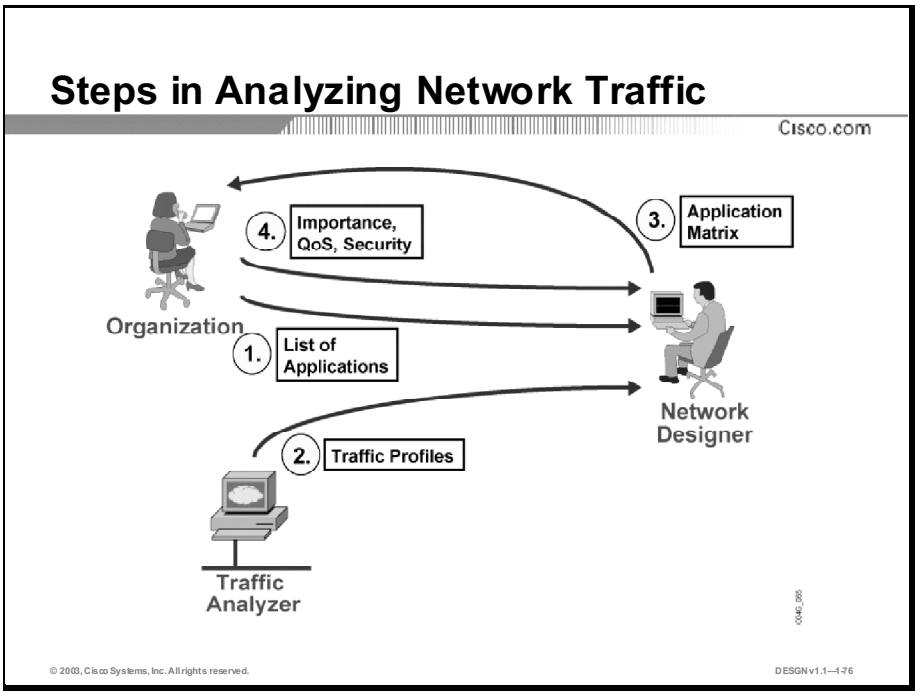

Use an interactive approach to create a list of applications and protocols used in the network.

Evaluate the results of the traffic analysis using the following steps:

Step 1 Use direct organizational input.

Step 2 Verify the organization's list of applications by using a traffic analyzer.

Step 3 Present the organization with the extended list of applications.

Step 4 Generate the final list of applications with their requirements (importance, QoS, security).

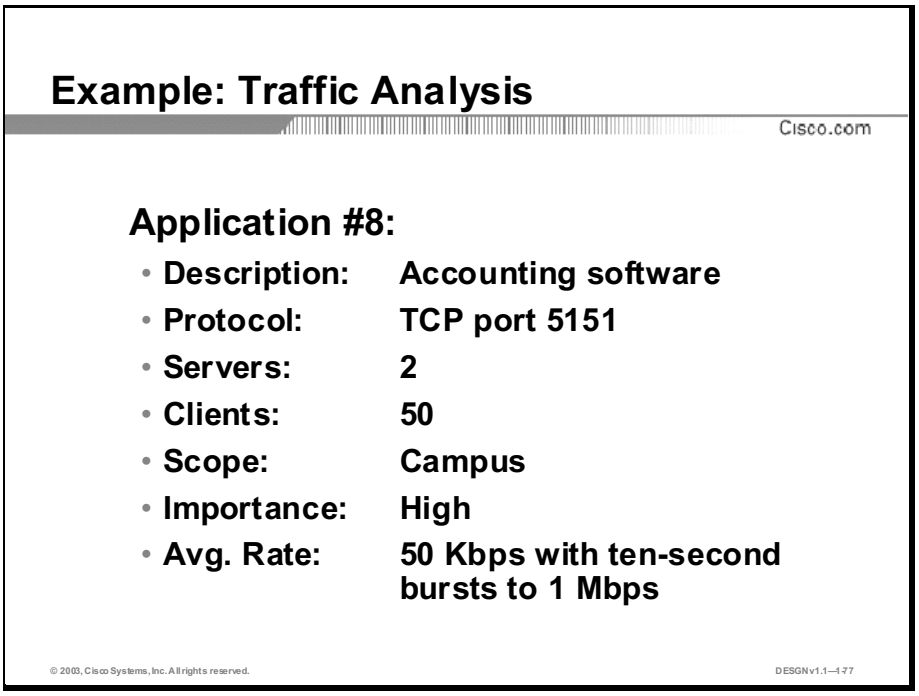

The figure illustrates a fictitious application (number 8 on the list) with information that may or may not be relevant for characterizing the existing network.

Assuming that the requirement concerns QoS in a classical WAN network module, the information is relevant because it describes the:

- Application (TCP port 5151), which is needed to perform classification
- Importance, which is needed when evaluating how much bandwidth should be allocated to this application
- Current bandwidth consumption according to the present QoS implementation

However, if there is a concern about a secure and resilient connection to the Internet, this information might not be relevant.

Tools for Analyzing Network Traffic

Tools used for traffic analysis range from manual identification of applications using IOS software commands in combination with network-based application recognition (NBAR) or NetFlow, to those where dedicated software- or hardware-based analyzers capture live packets. This topic lists the tools for monitoring and analyzing network traffic.

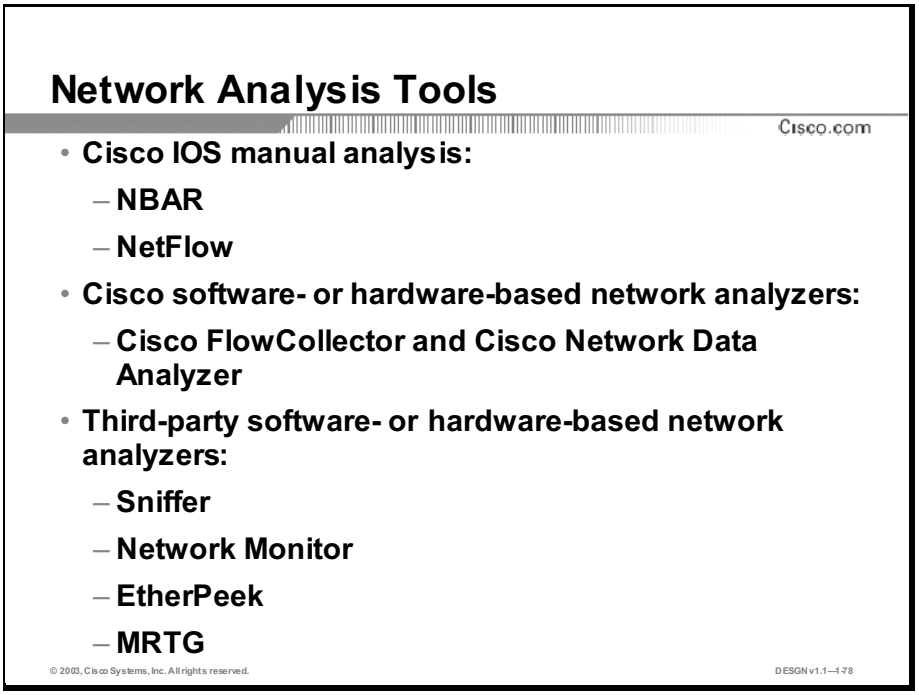

The tool options for traffic analysis are:

- Use NBAR to identify well-known applications and protocols in the network.
- IOS NetFlow technology is an integral part of IOS software that collects and measures data as it enters specific routers or switch interfaces. NetFlow identifies lesser-known applications as it gathers the information of every single flow. You can manually collect the information using the IOS software **show ip cache flow** command.
- Cisco FlowCollector and Data Analyzer allow automatic information gathering of every single flow in the network segment.
- Use third-party hardware- or software-based products to analyze traffic in different subnets of the network. An example of a third-party software-based analyzer is Multi Router Traffic Grapher (MRTG).

Example: NBAR Printout

Cisco.com

```
Router#show ip nbar protocol-discovery

FastEthernet0/0.2
                         Input                     Output
  Protocol               Packet Count              Packet Count
                         Byte Count                Byte Count
                         30 second bit rate (bps)  30 second bit rate(bps)
  ---------------------- ------------------------  ------------------------
  http                   46384                     79364
                         5073520                   64042528
                         305                       1655
  secure-http            2762                      2886
                         429195                    1486350
                         0                         0
  snmp                   143                       10676
                         17573                     1679322
                         0                         0
  telnet                 1272                      12147
                         122284                    988834
                         0                         0
  ntp                    5383                      0
                         624428                    0
                         0                         0
  dns                    305                       235
                         31573                     55690
                         50                        120
    23412        2632      15404          0          0 Net Background
<more>
```


The figure presents an example printout of the IOS NBAR **show ip nbar protocol-discovery** command. This command illustrates the statistics gathered with the NBAR Protocol Discovery feature. It displays statistics for all interfaces on which protocol discovery is currently enabled. The default output of this command includes average 30 seconds bit rate (in bits per second), input byte count, input packet count, and protocol name.

Protocol discovery provides an easy way to discover application protocols transiting an interface. The protocol discovery feature identifies any protocol traffic that NBAR supports. Use protocol discovery to monitor both input and output traffic.

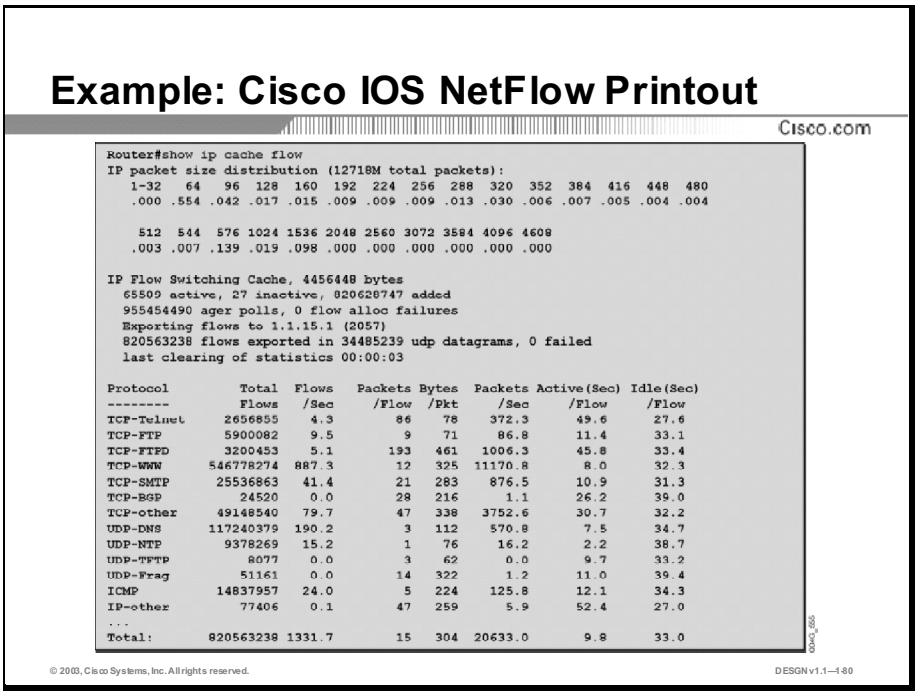

The figure presents an example printout of the IOS NetFlow feature using the **show ip cache flow** command. By analyzing NetFlow data, you can identify the cause of congestion, determine the class of service (CoS) for each user and application, and identify the source and destination network for the traffic. NetFlow provides extremely granular and accurate traffic measurements and high-level aggregated traffic collection.

The table describes the fields displayed in the **show ip cache flow** command output.

Field	Description
Bytes	Number of bytes of memory used by the NetFlow cache
active	Number of active flows in the NetFlow cache at the time this command was entered
inactive	Number of flow buffers that are allocated in the NetFlow cache
added	Number of flows created since the start of the summary period
ager polls	Number of times the NetFlow code looked at the cache to expire entries
flow alloc failures	Number of times the NetFlow code tried to allocate a flow but could not
exporting flows	IP address and User Datagram Protocol (UDP) port number of the workstation to which flows are exported
Flows exported	Total number of flows exported and the total number of UDP datagrams
Failed	Number of flows that could not be exported by the router
last clearing of statistics	Standard time output (hh:mm:ss) since the **clear ip flow stats** command was executed
Protocol	IP protocol and the “well-known” port number as described in RFC 1340

Field	Description
Total Flows	Number of flows for this protocol since the last time statistics were cleared
Flows/Sec	Average number of flows for this protocol seen per second
Packets/Flow	Average number of packets observed for the flows seen for this protocol
Bytes/Pkt	Average number of bytes observed for the packets seen for this protocol
Packets/Sec	Average number of packets for this protocol per second
Active(Sec)/Flow	Sum of all the seconds from the first packet to the last packet of an expired flow
Idle(Sec)/Flow	Sum of all the seconds from the last packet seen in each nonexpired flow

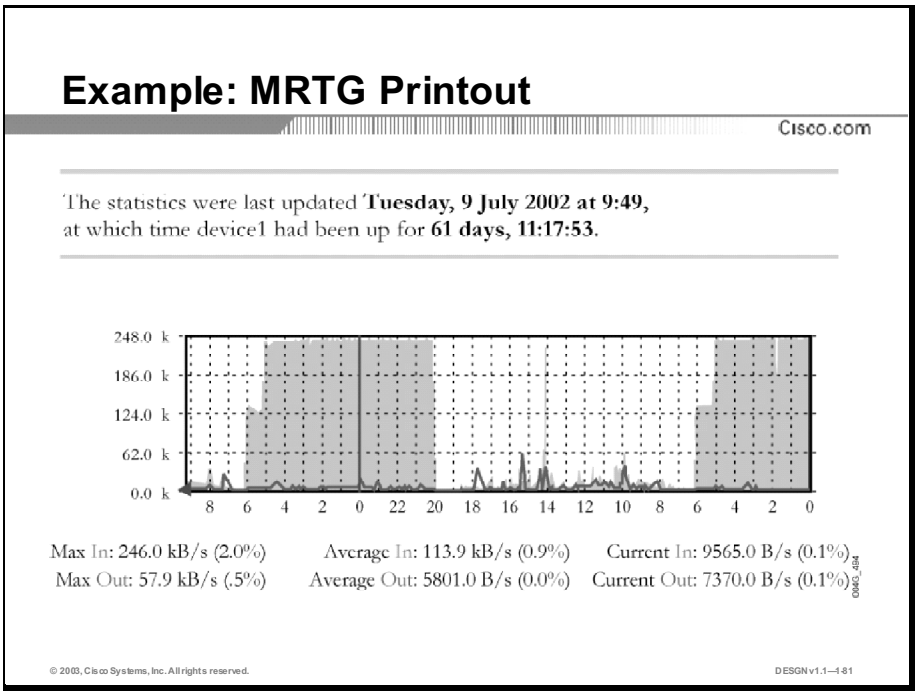

The figure presents an example printout from the MRTG tool that illustrates the daily throughput on the link to the Internet.

Network Health Analysis

You can use this checklist to help determine the health status of the existing network:

- No shared Ethernet segments are saturated (no more than 40 percent network utilization).
- No WAN links are saturated (no more than 70 percent network utilization).
- The response time is generally less than 100 milliseconds ($1/10^{th}$ of a second).
- No segments have more than 20 percent broadcasts/multicasts.
- No segments have more than one cyclic redundancy check (CRC) error per million bytes of data.
- On the Ethernet segments, less than 0.1 percent of the packets result in collisions.
- The Cisco routers are not over utilized (five-minute CPU utilization no more than 75 percent).
- The number of output queue drops has not exceeded 100 in an hour on any Cisco router.
- The number of input queue drops has not exceeded 50 in an hour on any Cisco router.
- The number of buffer misses has not exceeded 25 in an hour on any Cisco router.
- The number of ignored packets has not exceeded 10 in an hour on any interface on a Cisco router.

Summarizing the Network Characterization

Network characterization results in a summary report describing the health of the network. This topic analyzes the existing network health based on audit and traffic measurements.

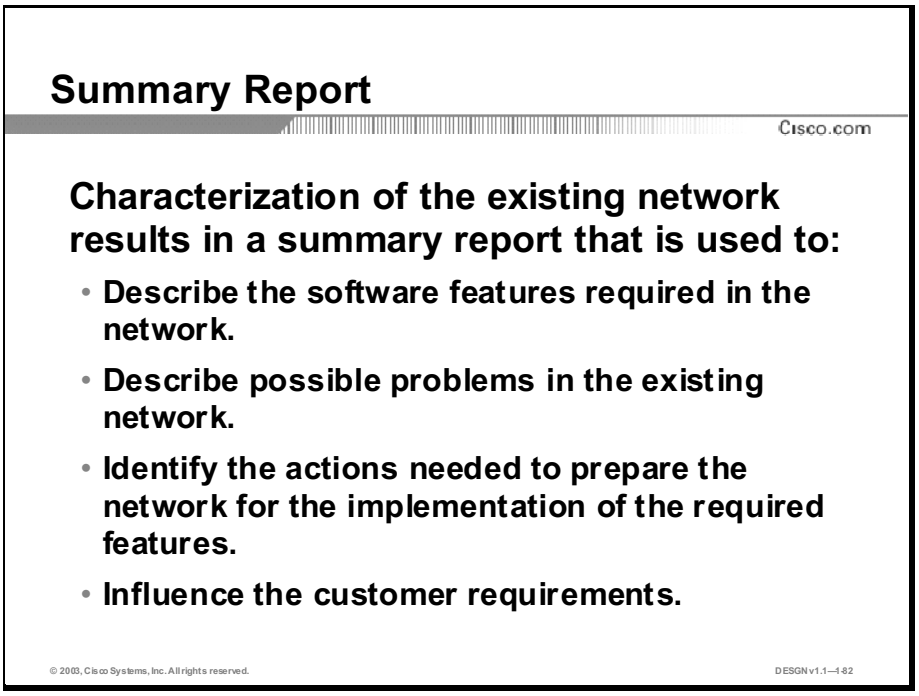

Organization input, network audit, and traffic analysis should provide enough information to identify possible problems in the existing network. You must convert the collected information into a concise summary report that identifies possible drawbacks of the network. With this information, you can propose hardware and software upgrades to support the network requirements or influence the organizational requirements.

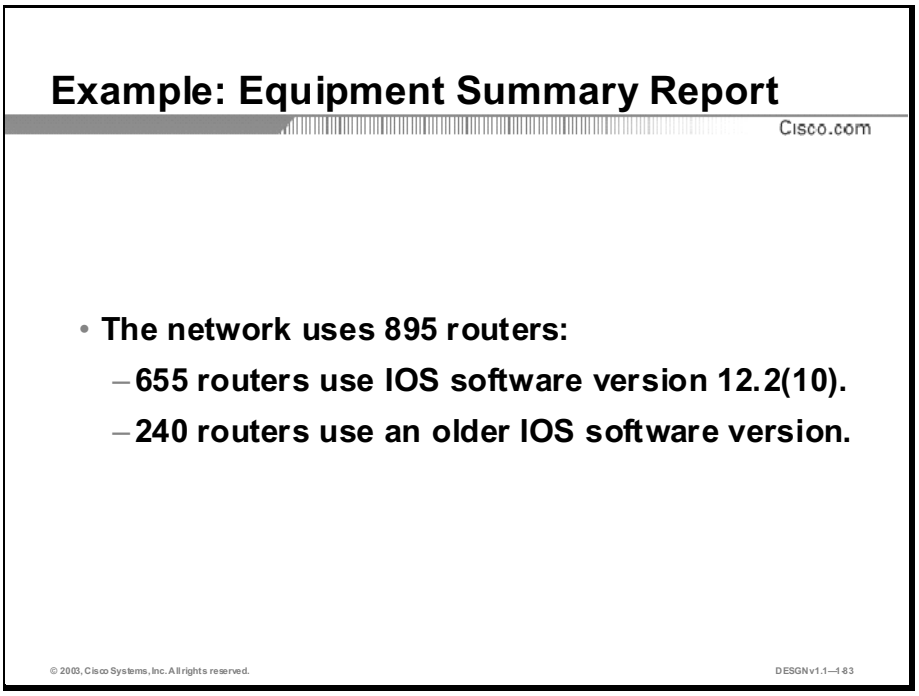

In the figure, the summary report identifies the number of routers with a common IOS software version.

Example: Summary Report Problem Statement

Cisco.com

- **Requirement: Queuing in the WAN**
- **Identified problem:**
  - **Existing IOS software version does not support new queuing technologies.**
  - **15 out of 19 routers with older IOS software are in the WAN.**
  - **12 out of 15 routers do not have enough memory to upgrade to IOS software version 12.2.**
  - **5 out of 15 routers do not have enough FLASH memory to upgrade to IOS software version 12.2.**

The summary report conclusions should identify shortcomings of the existing infrastructure.

In the example, queuing is required. However, the existing IOS software version does not support the needed queuing features. In addition, some routers do not have enough RAM or Flash memory for an upgrade.

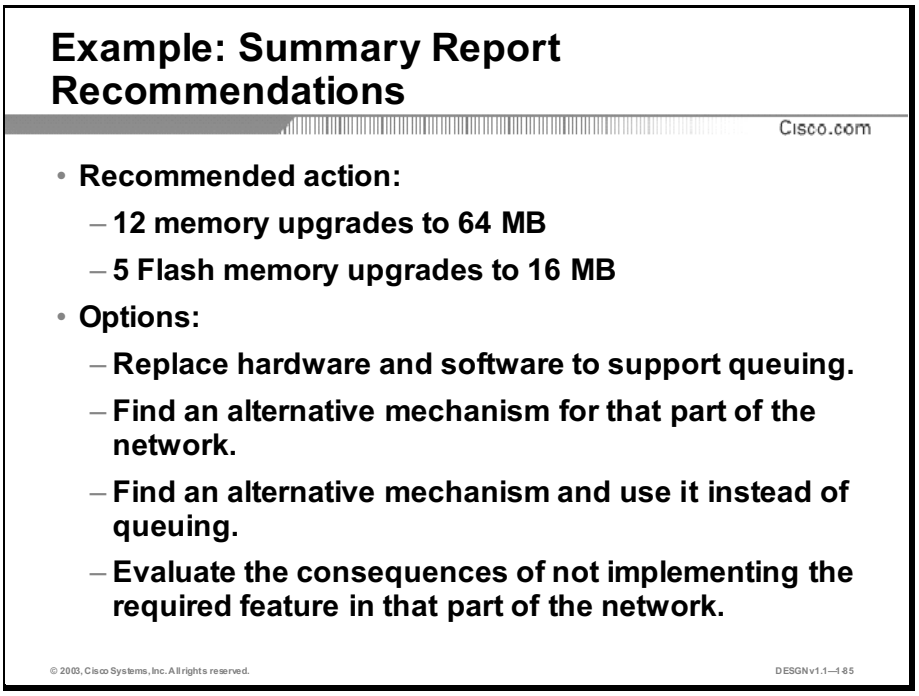

The figure presents example recommendations in a summary report.

Summary report recommendations relate to the existing network and to the organizational requirements. Use the summary report to recommend hardware and software to support the required feature.

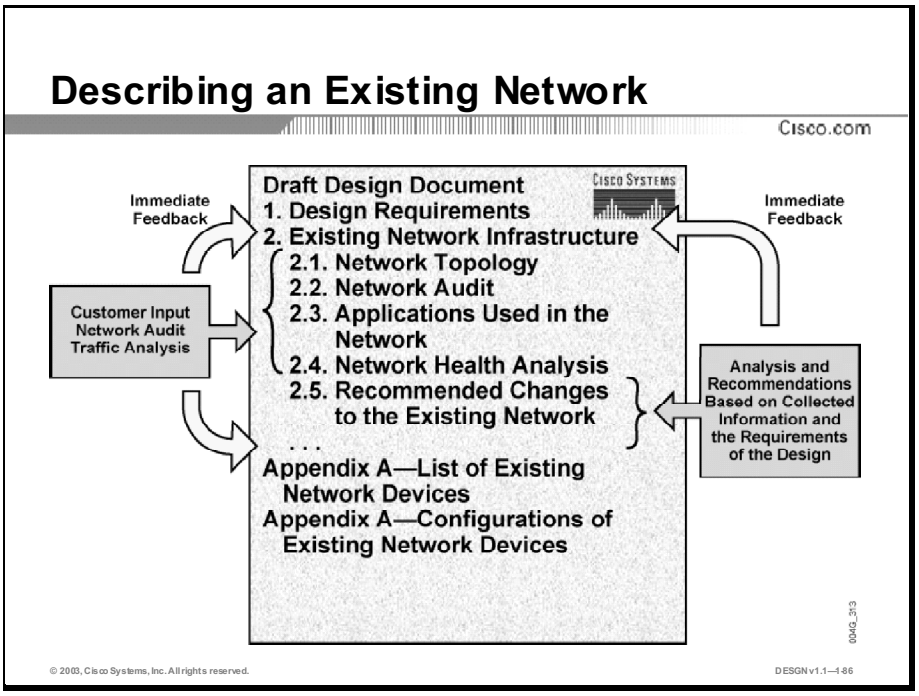

After thorough examination of the existing network, you create a draft design document. The figure illustrates the not yet fully developed index of the draft design document, including the topic that describes the existing network. The *Design Requirements* and *Existing Network Infrastructure* sections are closely related. The examination of the existing network may demonstrate the need for changes to the *Design Requirements*. Data from both sections has a direct influence on the design of the network.

Typical existing network draft documentation should include these items:

- Logical (Layer 3) topology map or maps (with the topology divided into network modules if the network is too large to fit into one topology map)
- Physical (Layer 1) topology map or maps
- Network audit results (types of traffic in a network, traffic congestion points, suboptimal traffic paths, and so on)
- Summary describing the major network services used in the existing network; for example, Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), IP Security (IPSec); you can attach configurations of all network devices as either a separate document or an appendix to the design document.
- Summary description of applications and overlay services used in the network
- Summary describing issues that may impact the design or established design requirements
- List of existing network devices, with the platform and software versions
- Configurations of existing network devices

Summary

This topic summarizes the key points discussed in this lesson.

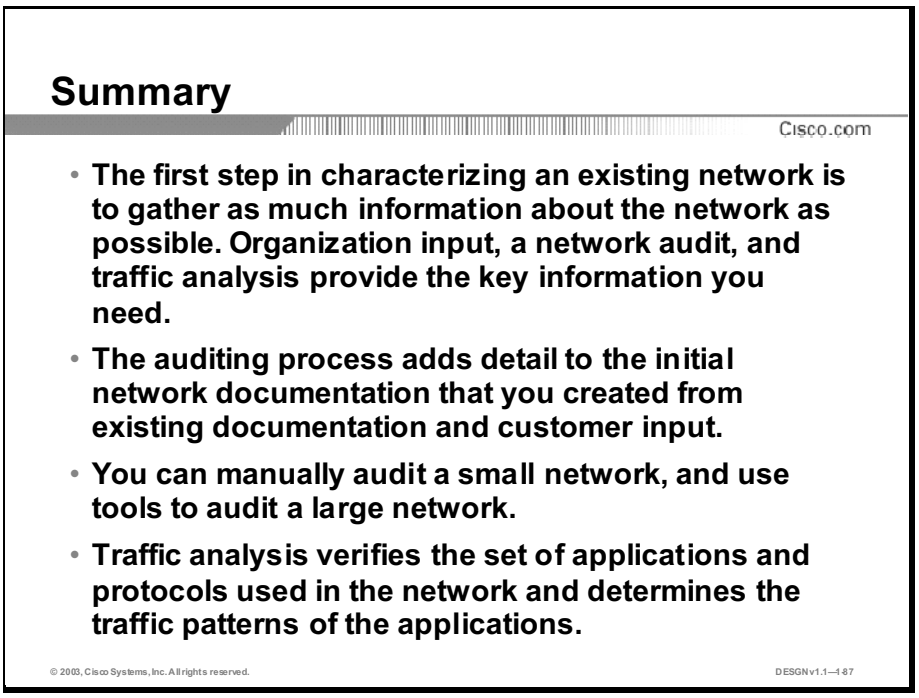

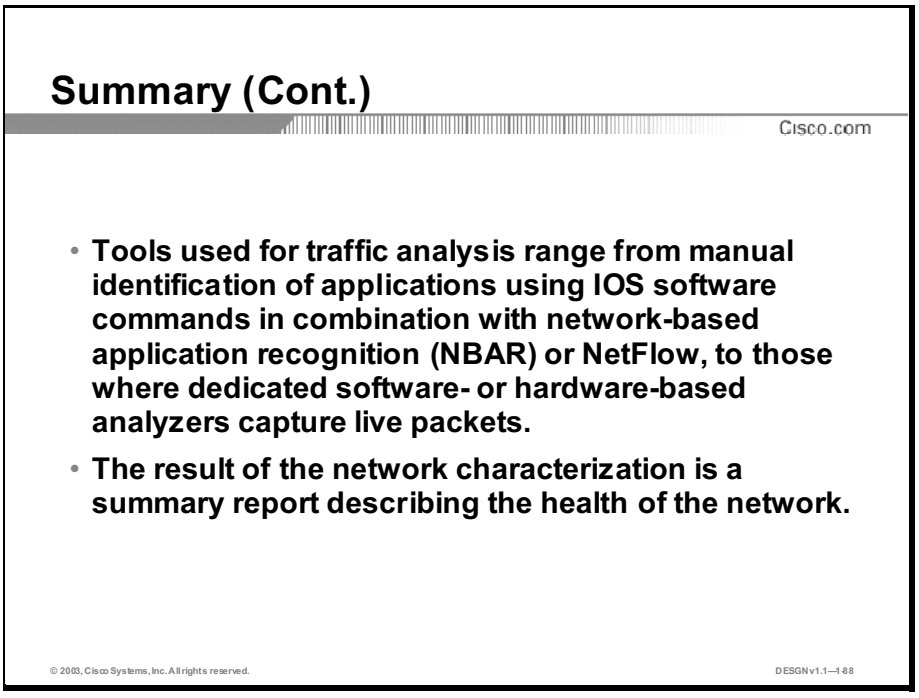

References

For additional information, refer to this resource:

- Oppenheimer, P. *Top-Down Network Design: A Systems Analysis Approach to Enterprise Network Design*. Indianapolis, Indiana: Macmillan Technical Publishing—Cisco Press; 1999.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three tasks are part of the characterization process of an existing network? (Choose three.)

A) collecting information using the existing documentation and direct organizational input

B) using tools to analyze network traffic

C) using design tools to create a framework for the design

D) using tools for automated auditing of the network

E) identifying the organization's business objectives

Q2) The network design should define an effective Layer 2 and Layer 3 topology in the Enterprise Campus. Which information should you collect during a network audit to help you define the Layer 2 and Layer 3 topology? (Choose three.)

A) routing tables to determine suboptimal packet paths

B) geographical locations of each site

C) link use to determine overloaded links

D) external connections to Internet service providers and partner networks

E) configuration of network devices

Q3) Which command can you use to determine the average CPU use on a Cisco router?

A) command **`show processes memory`**

B) command **`show processes cpu`**

C) command **`show cpu utilization`**

D) command **`show cpu`**

Q4) Which three parameters can you identify using a traffic analyzer? (Choose three.)

A) protocol specification (IP protocol ID, TCP/User Datagram Protocol (UDP) port number)

B) average bit rate and packet rate

C) QoS requirements

D) importance of the application

E) devices that use the application and their addresses

Q5) Which three commands reveal information about individual applications, protocols, or flows? (Choose three.)

A) command **`show processes memory`**

B) command **`show ip nbar protocol-discovery`**

C) command **`show ip interface`**

D) command **`show ip cache flow`**

E) command **`show processes cpu`**

Q6) Fill in the missing item. The network health analysis produces a report that is _____.

A) based on the organizational requirements

B) based on the existing network and expected functionality

C) used to sell more boxes

D) stored in a Microsoft Word document

Q7) Select the best answer. What input is used to create the documentation of an existing network?

A) existing documentation and the organizational input

B) auditing and analytical tools

C) organizational input, auditing, and analytical tools

D) monitoring commands on routers and switches

Quiz Answer Key

Q1) A, B, D

Relates to: Identifying the Existing Infrastructure and Its Features

Q2) A, C, E

Relates to: Auditing the Existing Network

Q3) B

Relates to: Tools for Auditing the Network

Q4) A, B, E

Relates to: Analyzing Network Traffic and Applications

Q5) B, C, D

Relates to: Tools for Analyzing Network Traffic

Q6) B

Relates to: Summarizing the Network Characterization

Q7) C

Relates to: Summarizing the Network Characterization

Completing the Network Design

Overview

After you gather organizational requirements, and document and audit the existing network, you are ready to design a network solution, plan the implementation, and (optionally) build a network prototype. You should document each process during the implementation to make it easier to create the final design document.

The lesson begins with an explanation of how to assess the scope of the design project and how to complete the list of requirements. After gathering all customer requirements and inputting them into decision tables, you will identify and gather missing information and assess the scope of the project again to ensure a comprehensive understanding of the network needs.

Relevance

This lesson provides the learner with the knowledge needed to apply an appropriate network design methodology using a modular top-down approach.

Objectives

Upon completing this lesson, you will be able to implement a network design. This includes being able to meet these objectives:

- Describe the top-down network design approach
- Discuss the role of decision tables in network design
- Assess the scope of a network design project
- Describe the structured process of designing a network solution
- Describe the components of a network implementation plan
- Explain when and how to build a prototype or pilot network
- Describe each component of a network design document

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Understanding of network design methodologies and concepts

Outline

The outline lists the topics included in this lesson.

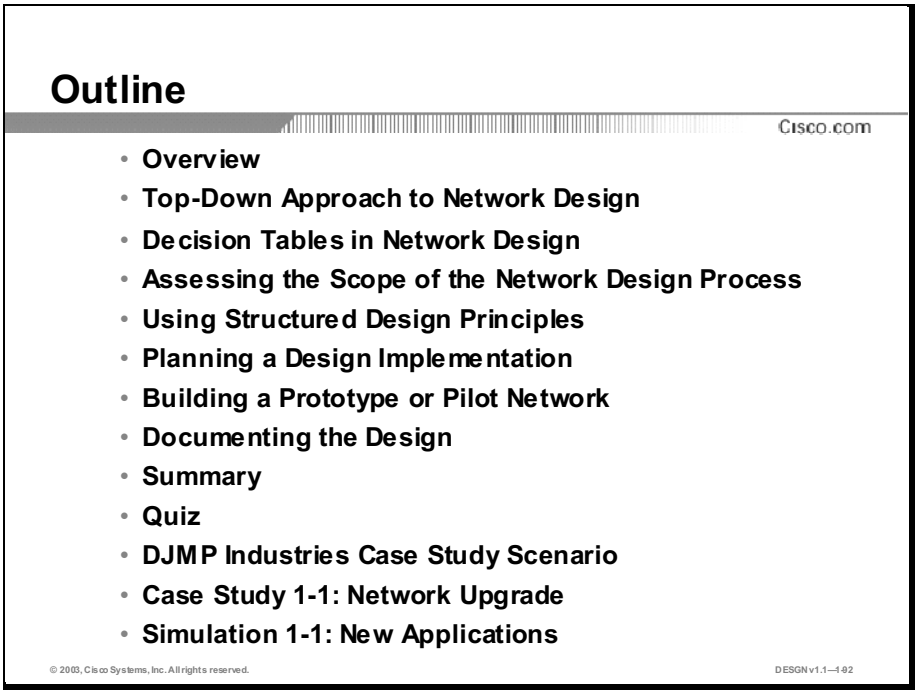

Top-Down Approach to Network Design

Designing an enterprise network is a complex project. Top-down design facilitates the process by dividing it into smaller, more manageable steps. This topic describes the top-down network design approach.

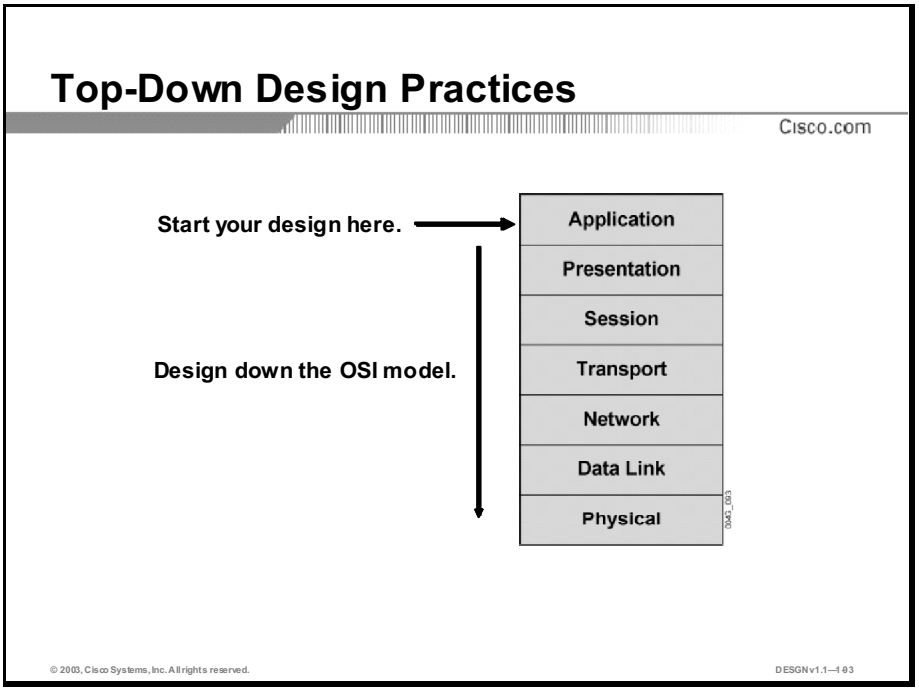

Top-down design clarifies the design goals and initiates the design from the perspective of the required applications and network solutions (IP telephony, content networking, and so on). The top-down approach adapts the physical infrastructure to the needs of the network solution.

When you deploy a bottom-up approach by selecting network devices and technologies first, the network may not meet the needs of the organization. With a bottom-up or connect-the-dots approach, the risk of redesigning the network is very high.

To complete a top-down design:

- Analyze the organization's requirements and applications.
- Complete the design from the top of the Open System Interconnection (OSI) reference model to the bottom:
  - Define requirements at the upper OSI layers (applications, presentation, and session).
  - Specify the infrastructure required in the lower OSI layers (transport, network, data link, and physical).
- Gather additional data on the network that may influence the logical and physical design and adapt the design to the new data as required.

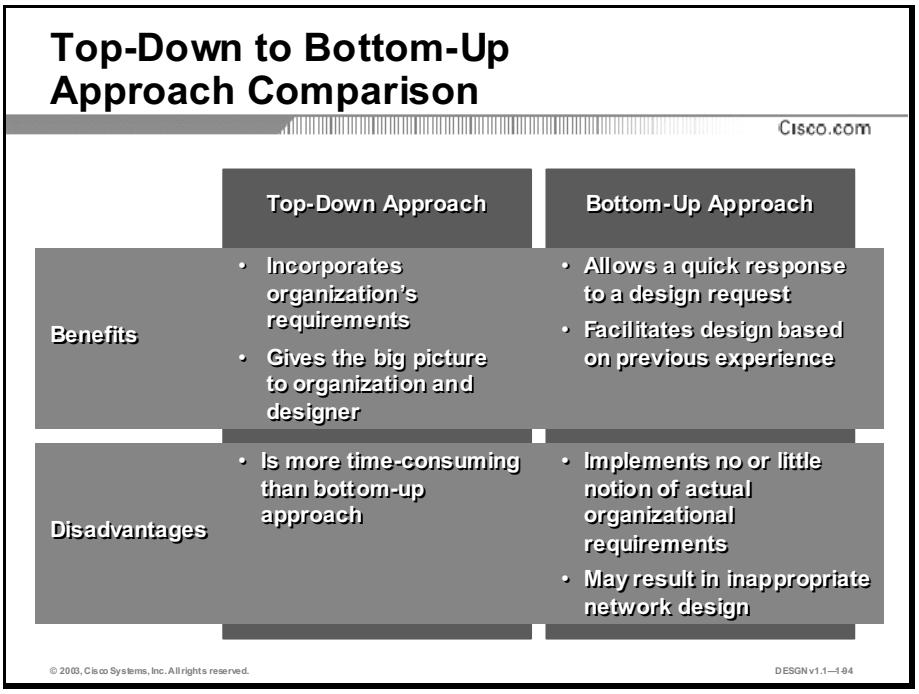

Benefits of the top-down approach over a bottom-up approach include:

- Incorporates the customer organization's requirements
- Provides the "big picture" of the desired network to the customer and the designer
- Provides a design that is appropriate for current requirements and future developments

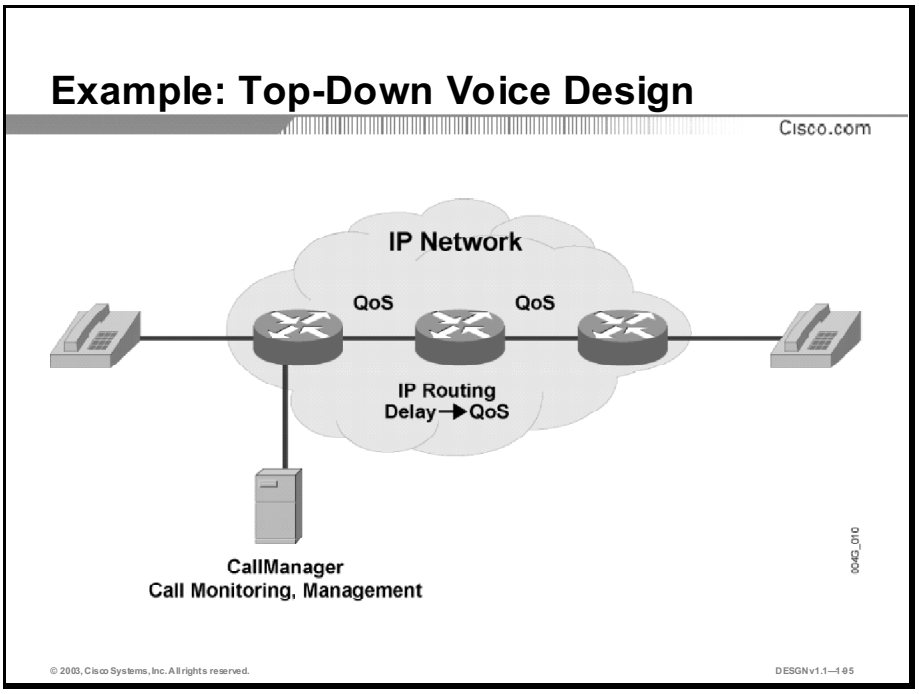

The figure illustrates the basics of the top-down approach when designing an IP telephony network solution from the perspective of the designer.

The organization requires a data network capable of supporting IP telephony, thus avoiding the costs of having two separate networks. The organization needs Voice over IP (VoIP) as the driving application. IP routing and QoS are needed at the transport and network layers. The CallManager addresses the application need for call routing.

The resulting network design includes IP-enabled routers (and other devices not shown in the figure) so that IP routing takes place in the network. To implement IP telephony, the delay in the IP network is managed with specific QoS mechanisms.

A Cisco CallManager is placed inside the network to manage and monitor IP telephone calls.

Note: The CallManager is a server-based application that establishes and maintains signaling and control for IP telephone sessions.

Decision Tables in Network Design

Decision tables provide a simple and systematic summary of different parameters and states of a network used in decision making. Decision tables facilitate the selection of the most appropriate option from many possibilities. This topic discusses the role of decision tables in network design.

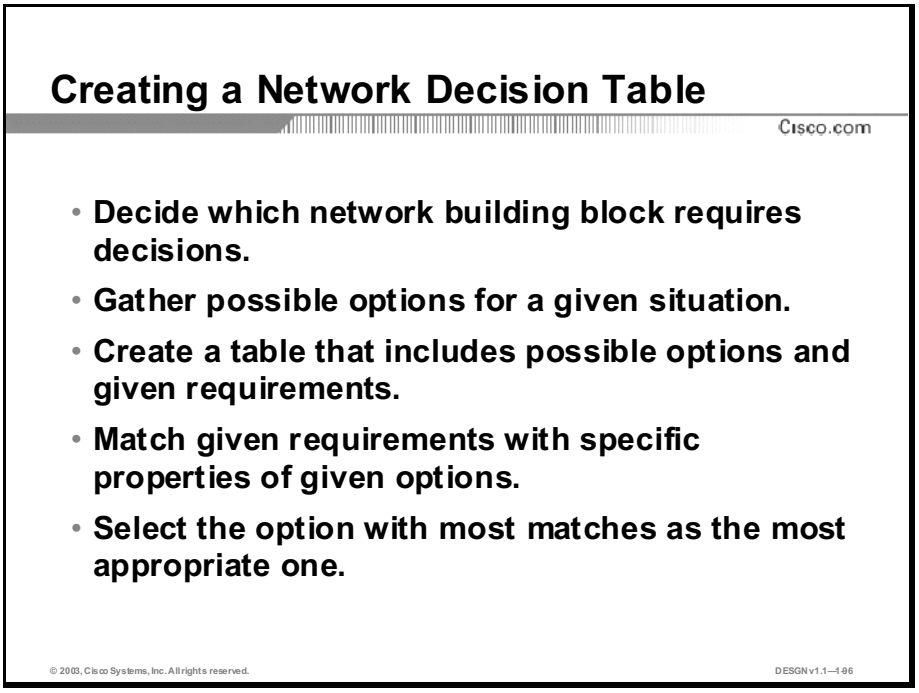

Decision tables provide the means for making a decision when multiple solutions exist for a given problem:

- Simplify the selection of the most appropriate option. Decide on the type of specific network building blocks (topology, routing protocols, security, and so forth).

Basic guidelines for creating a network design decision table include:

Step 1 Decide which network building block (physical topology, routing protocol, security implementation, and so on) requires decisions.

Step 2 Collect possible options for a given situation. Be certain to include all options to obtain maximum value from the decision table. A thorough survey of the existing state of technology and considerable knowledge are needed to include all options.

Step 3 Create a table to include possible options and given requirements. Add the parameters or properties of specific options.

Step 4 Match the given requirements with the specific properties of the given options.

Step 5 Select the most appropriate option, that is, the option with the most matches, when all requirements are equally treated. However, if some requirements are considered more important than others, you can implement a system of weights, where you assign each of the requirements a weight proportional to its importance in decision making.

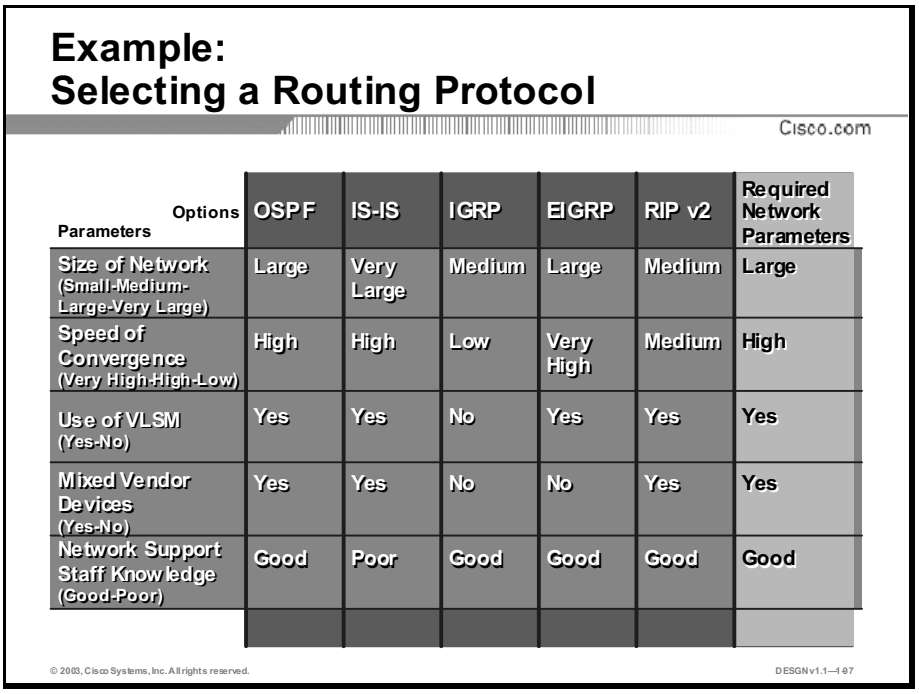

The figure provides an example decision table for selecting a routing protocol based on multiple criteria. Several routing protocols are considered as possible options based on their compliance to five different required parameters. The protocol you choose should meet these requirements:

- It should support a large network (up to 100 or more routers). Routing Information Protocol Version 2 (RIPv2) and Interior Gateway Routing Protocol (IGRP) protocols do not meet this requirement.
- It must have a high speed of convergence. This precludes RIPv2 and IGRP.
- The use of variable-length subnet masking (VLSM) is required. IGRP does not support VLSM.
- It must support Cisco and other vendors' equipment. Enhanced Interior Gateway Routing Protocol (EIGRP) and IGRP are Cisco proprietary protocols, so they do not support mixed vendor environments.
- Network staff should have a good knowledge of the chosen protocol, enabling them to troubleshoot the network. Most network administrators have only a basic knowledge of Intermediate System-to-Intermediate System (IS-IS), as it is not a widely used protocol.

Note: All requirements are of the same importance in this example, so no weights are used.

Based on the stated requirements, OSPF would be the routing protocol of choice.

Decision Table Worksheet

The table is provided as a template. Fill in the parameters on each row, the available options for each column, and make your own decision table.

Parameter						Required Network Parameters

Assessing the Scope of the Network Design Process

In assessing the scope of a network design, you must determine if the design is for a new network or is a modification of the entire network, a single segment or module, a set of LANs, a WAN, or a remote-access network. In another review of the design, the network designer determines if the design addresses a single function or all of the OSI model layers. This topic assesses the scope of a network design project.

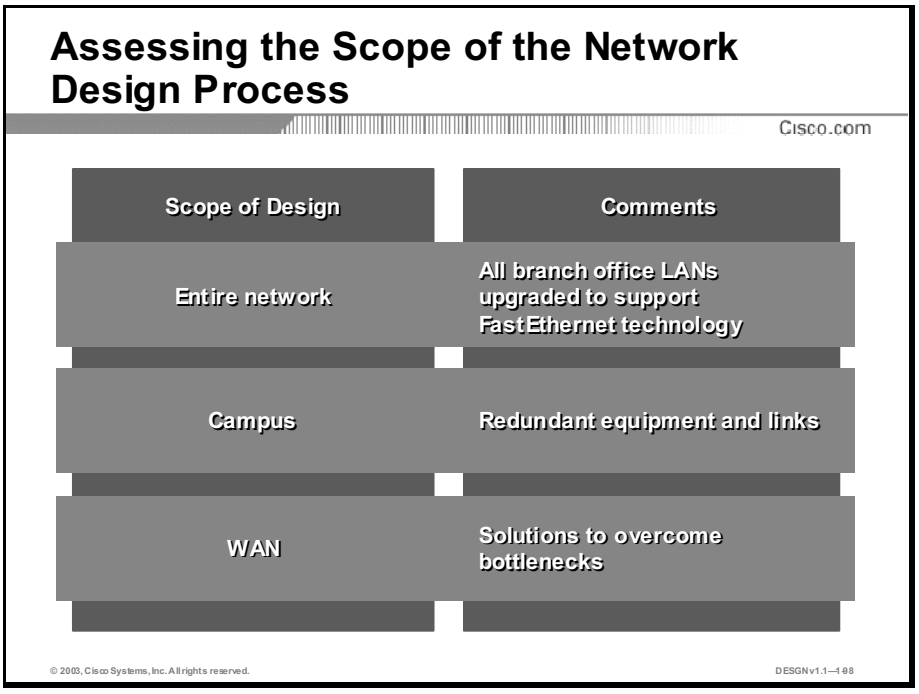

You must determine if the design:

- Is an upgrade or modification of an existing network, or if it is for a new network
- Applies to the whole network or only to a certain segment or module such as the campus or WAN

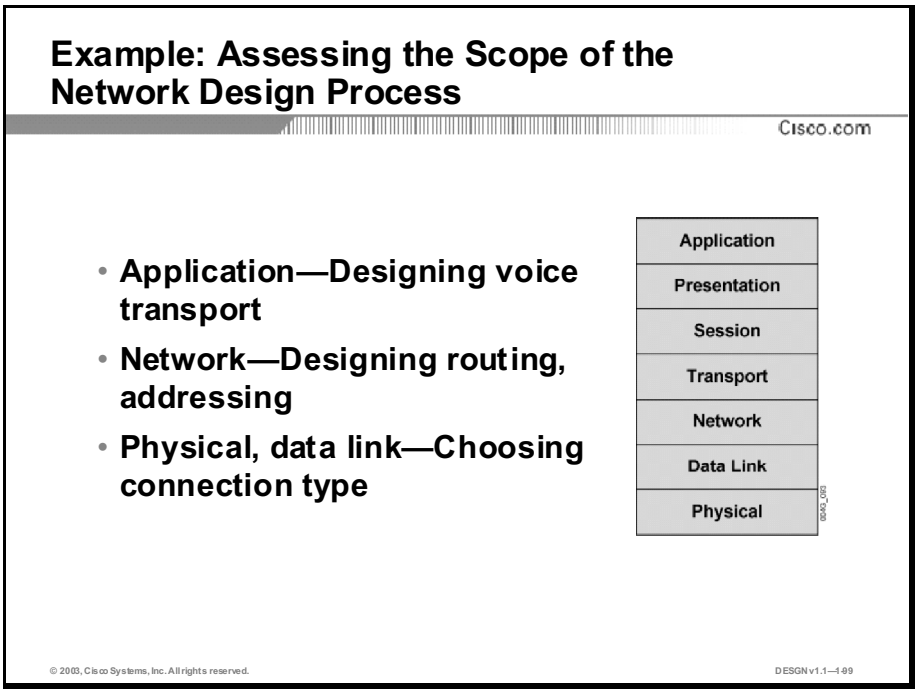

The OSI reference model is important in the design phase. The network designer should review the project scope from the protocol layer perspective and decide whether the design is needed at the network layer or if other layers are involved:

- The application layer includes the design of voice transport, as well as other applications.
- The network layer includes the design of routing and addressing.
- The physical and data link layers include decisions about the connection types and technologies to be used, for example, Gigabit Ethernet, ATM, and Frame Relay.

Example: Assessing the Scope of the Network

Corporation X Network Design Scope Assessment

Scope of Design	Comments
Entire network	The central office needs a backbone redesign. All branch office LANs will be upgraded to Fast Ethernet technology.
Network layer	Introduction of private IP addresses requires a new addressing plan. Certain LANs need segmentation. Routing must support the new addressing plan and provide greater reliability and redundancy.
Data link layer	The central office backbone and some branch offices require redundancy with redundant equipment and redundant links.

Design Scope Worksheet

Use the table as a template to identify the scope of the network design in the course case study and for future design efforts.

Scope of Design	Comments

Using Structured Design Principles

The output of the design should be a model of the complete system. To achieve this, the top-down approach is highly recommended. Do not focus on the network components, technologies, or protocols. Instead, focus on a systematic process that uses the business goals, technical objectives, and existing and future network services and applications. This systematic approach requires structured design practices, such as logical, physical, and functional models. This topic describes the structured process of designing a network solution.

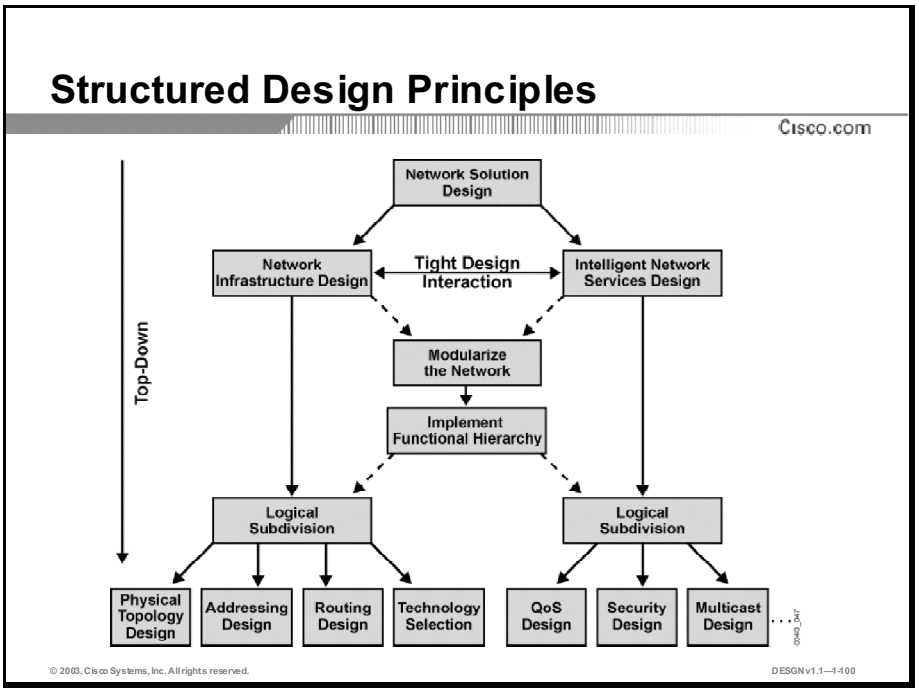

Network infrastructure and intelligent network services design are tightly connected, both bound to the same logical, physical, and functional models. These elements are logically subdivided, in both the network infrastructure and intelligent network services components. Use the top-down approach during all design phases.

Structured design practices focus on dividing the design task into related, less complex components.

- First, identify the logical connectivity requirements of the applications, with a focus on the necessary network solutions and the supporting network services. Examples of network solutions include voice, content networking, and storage networking. The network services required include availability, management, security, QoS, and IP multicast.
- Split the network functionally to develop the network infrastructure and hierarchy requirements. In this course, the Cisco Architecture for Voice, Video and Integrated Data (AVVID) provides a consistent infrastructure.
- Design each structured element separately in relation to other elements. Network infrastructure and intelligent network services design are tightly connected, both bound to the same logical, physical, and functional models.

Logical Structure

After you identify the connectivity requirements, you can work in more detail on each functional module. Each task involved in designing the network infrastructure and intelligent network services is a logical structure, which must be designed separately but in close relation with other structures. Your goal is one homogenous network.

Some logical structures are more closely related than others. Network infrastructure elements are more closely related to each other than to intelligent network services; for example, physical topology and addressing design are very closely related, while addressing and QoS are not.

Physical Structure

Several approaches to physically structuring a functional module exist. Designers have used the three-layer hierarchical structure—core, distribution, and access—for nearly a decade. In this approach, three separate yet related physical structures are developed instead of a single, large network. Hierarchical physical structuring gives the designer meaningful and functionally homogenous elements within each module. Selecting the functionality and required technologies is easier when applied to separate structured network elements than when it is applied to the complex network.

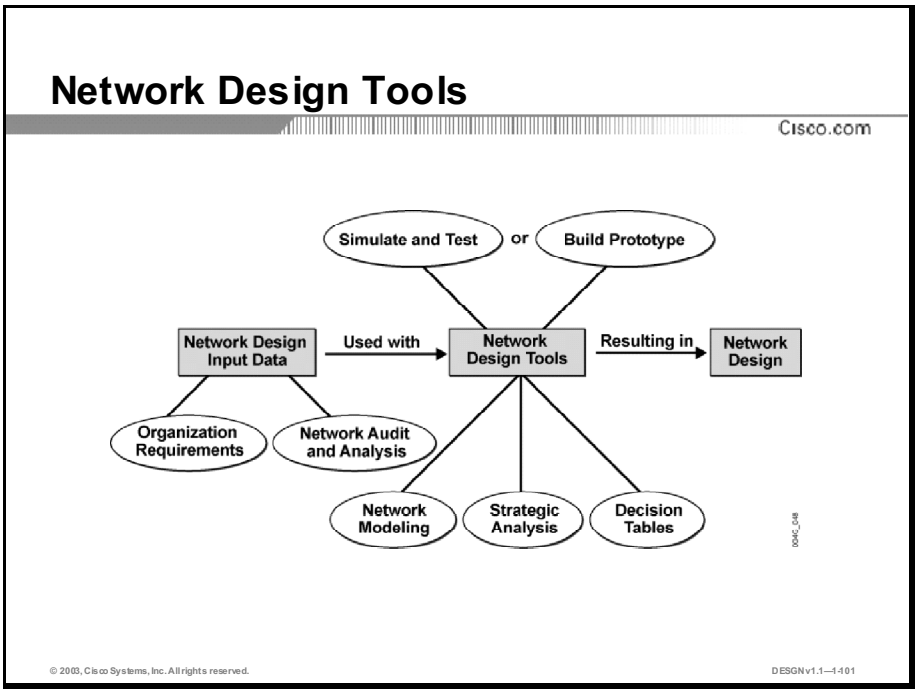

Several types of tools are available to help you design a complex modern network:

- **Network modeling tools:** Network modeling tools are helpful when the requirements are many and large. These tools model both simple and complex networks. The program processes the information provided and returns a proposed configuration. You can then modify the configuration and reprocess it to add redundant links, support additional sites, and so on.
- **Strategic analysis tools:** Strategic analysis or what-if tools help designers and others involved in the design (engineers, technologists, and business and marketing professionals) develop network and service plans, including detailed technical and business analysis. These tools attempt to calculate the effects of specific network components through simulated scenarios.
- **Decision tables:** Decision tables are manual tools for choosing specific characteristics of a network from multiple options, based on required parameters.
- **Simulation and verification tools or services:** These tools or services are used to verify the acquired design, lessening the need for a pilot network implementation.

To verify a network design produced with the help of network modeling tools, strategic analysis tools, and decision tables, use simulation and test tools or build a pilot or prototype network. The pilot or prototype network also serves to confirm the appropriateness of the design implementation plan.

Planning a Design Implementation

When the design is complete, you are ready to document the implementation and migration in as much detail as possible. This topic describes the components of a network implementation plan.

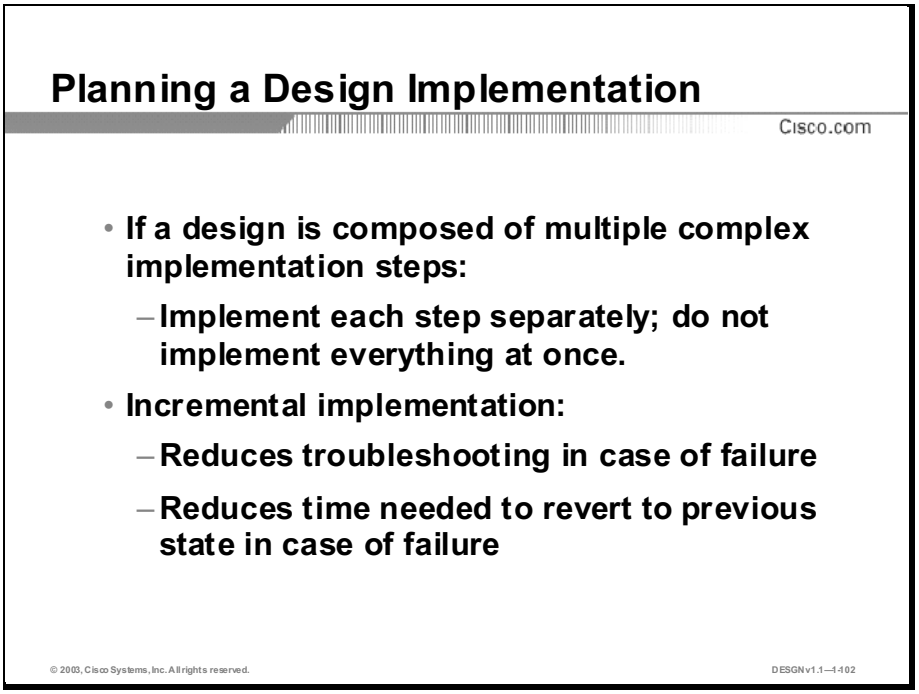

The more detailed the design documentation, the less knowledgeable the network engineer needs to be to implement the design. Very complex implementation steps usually require that the designer carries out the implementation, whereas other staff members can complete well-documented detailed implementation steps without the direct involvement of the designer.

When implementing a design you must consider the possibility of a failure, even after a successful pilot or prototype network test. You need a process test at every step and a procedure to revert to the original setup in case there is a problem.

Implementation steps and estimated times should be listed in a table.

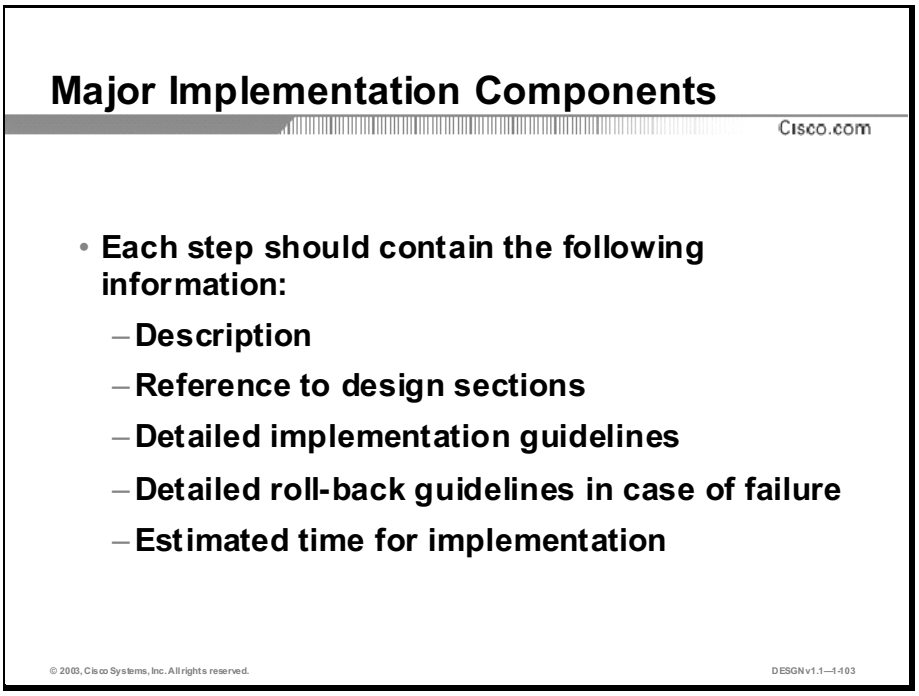

Implementation of a network design consists of several phases (install hardware, configure systems, launch into production, and so on). Each phase consists of several steps, and each step should contain:

- Description of the step
- Reference to design documents
- Detailed implementation guidelines
- Detailed roll-back guidelines in case of failure
- Estimated time needed for implementation

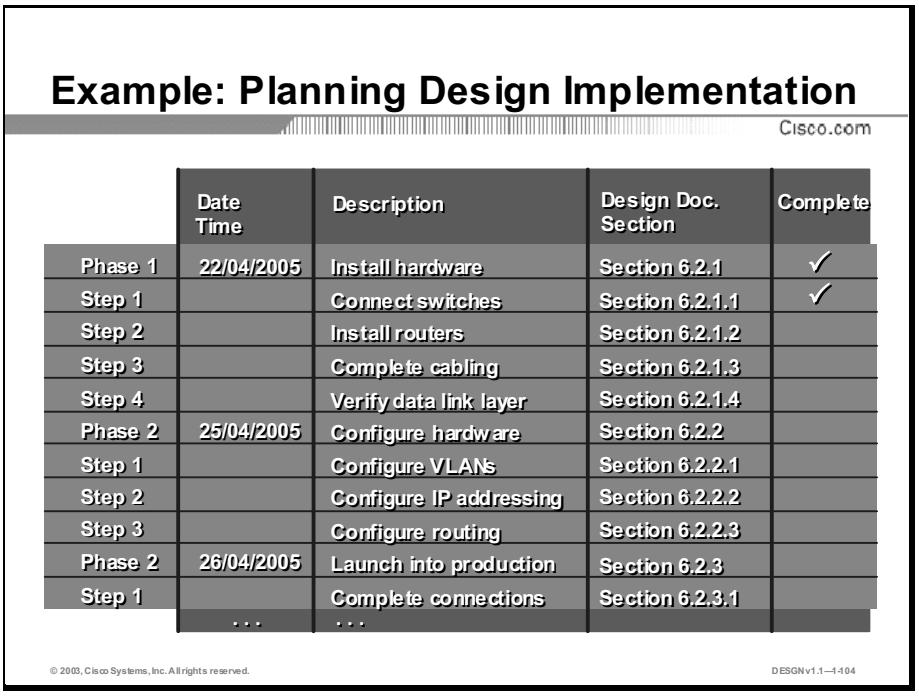

The figure provides a summary of all steps needed to complete the design implementation. Each step is briefly described, with references to help an implementor locate further details. The detailed descriptions should refer to the section of the design document that describes precisely what needs to be accomplished.

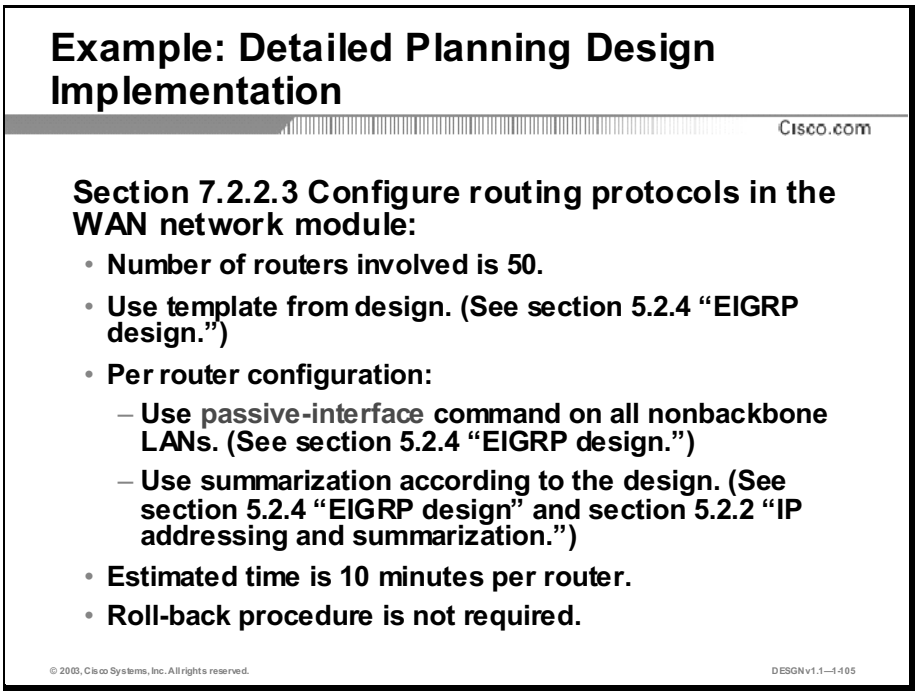

The figure provides a detailed description of one implementation step. It describes the configuration of EIGRP on 50 routers in the network and lists the two major components of the step (per router configuration procedure).

Note:	The reference to the design is used to provide the details about EIGRP implementation.

Building a Prototype or Pilot Network

After a design is complete, you must verify it. You can test the design in an existing or live network (pilot) or in a prototype network that will not affect the existing network. This topic explains when and how to build a prototype or pilot network.

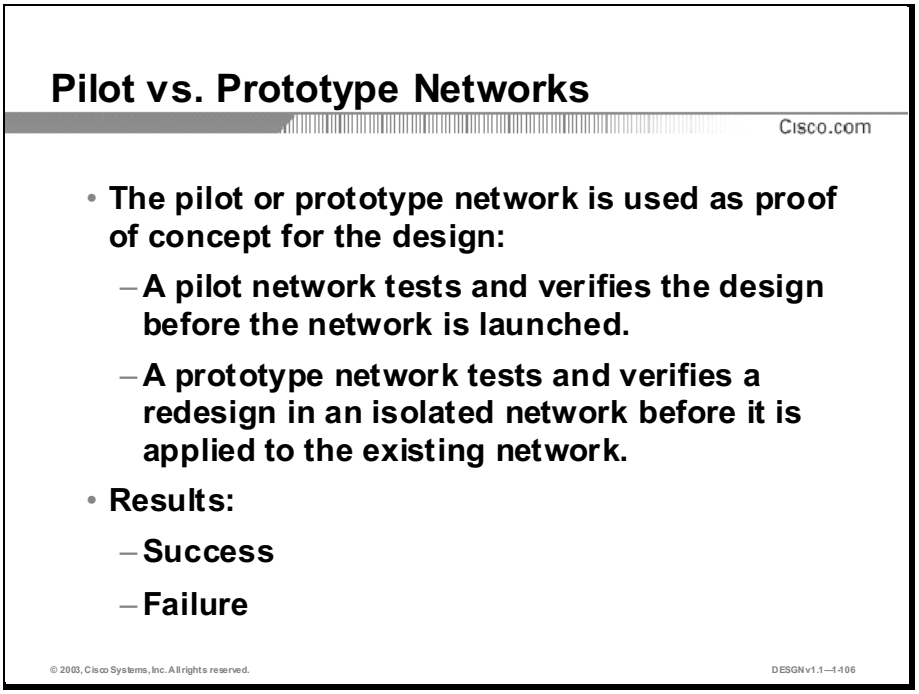

A pilot network is usually set up when the design is used to create a completely new network or to add to an existing network. A prototype network is usually used to verify designs that are implemented on an existing network infrastructure.

A successful design implementation in either a pilot or a prototype network serves as a proof of concept, in preparation for full implementation.

A prototype or pilot implementation can have one of two results:

- **Success:** This result is usually enough to prove the concept of the design.
- **Failure:** This result is usually used to correct the design and repeat the prototype or pilot phase. In the case of small deviations, the design can immediately be corrected and tested in the prototype or pilot network.

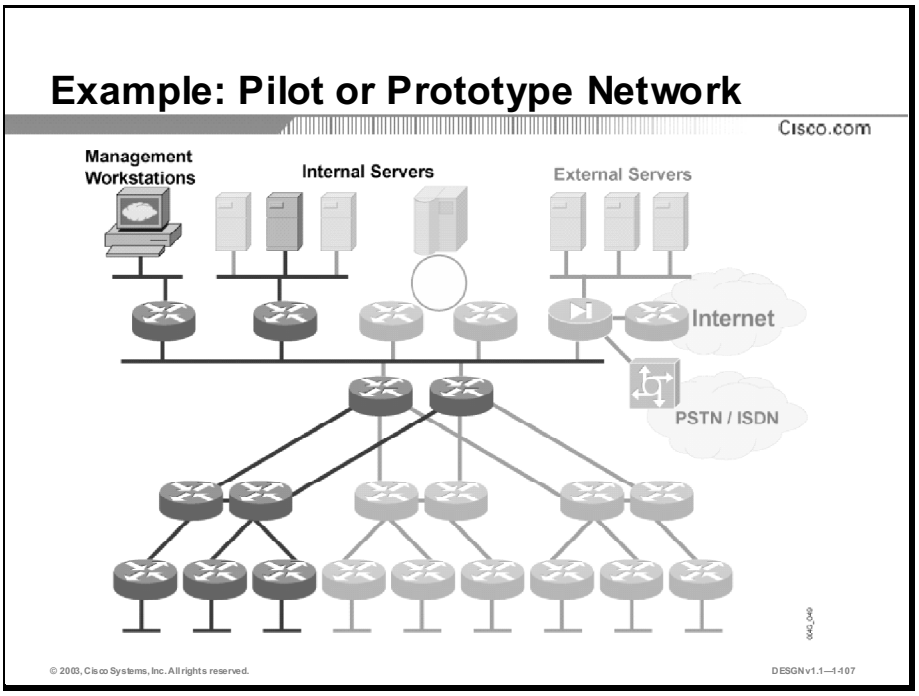

The figure presents an example topology subset of an existing network or a planned network. The highlighted areas indicate the parts of the network involved in the redesign.

The example represents two verification options:

- **Topology of the existing network and the emphasized part of the topology, replicated in a prototype network:** The design is then tested in the prototype network.
- **Topology of the new network as proposed by the design:** The emphasized part of the topology is implemented first in the pilot phase where the design is verified.

Documenting the Design

A design document lists the design requirements, documents the existing network, documents the network design, identifies the proof of concept strategy, and details an implementation plan. This topic describes each component of a network design document.

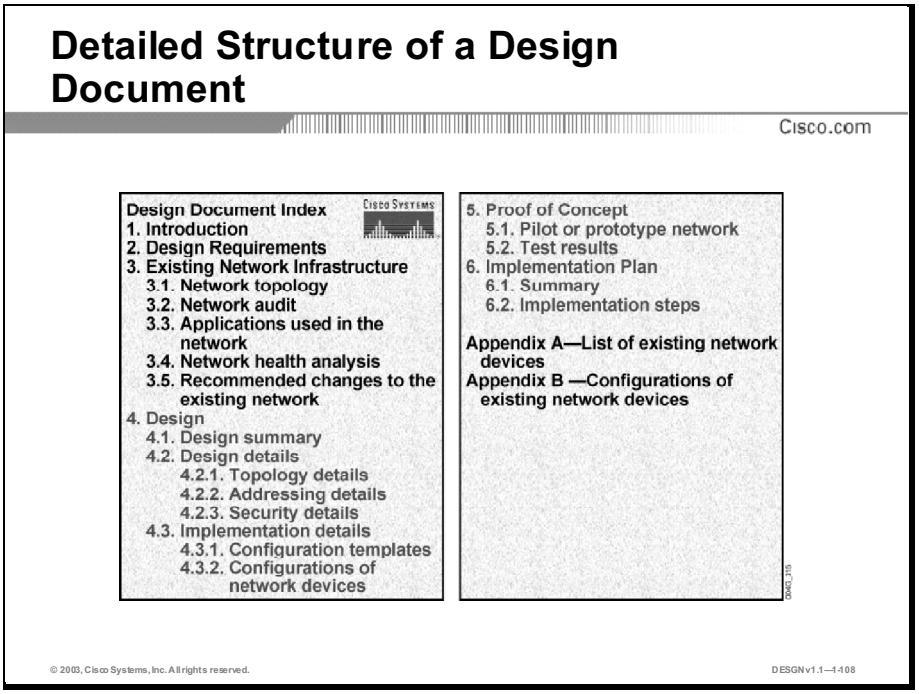

The final design document structure should include the following sections:

- **Introduction:** This section presents the main reasons leading to the network design or redesign.
- **Design Requirements:** This section identifies organization requirements and design goals to fulfill.
- **Existing Network Infrastructure:** This section describes an existing network and is used only for a network redesign.
- **Design:** This section identifies design and implementation. The *design details* describe the topology, addressing, and security. Include *implementation details* such as configuration templates and exact configurations of network devices.
- **Proof of Concept:** This section describes pilot or prototype network verification.
- **Implementation Plan:** This section provides the details that technical staff need to carry out implementation as quickly and smoothly as possible, without requiring the presence of the designer.
- **Appendixes:** These include lists and optional configurations of existing network devices (on network redesign).

Summary

This topic summarizes the key points discussed in this lesson.

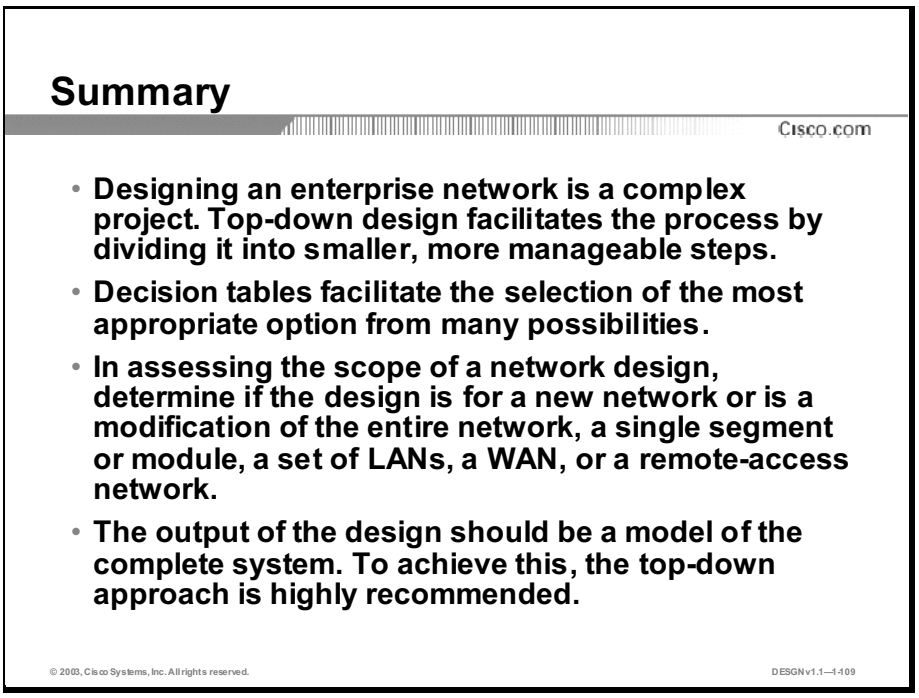

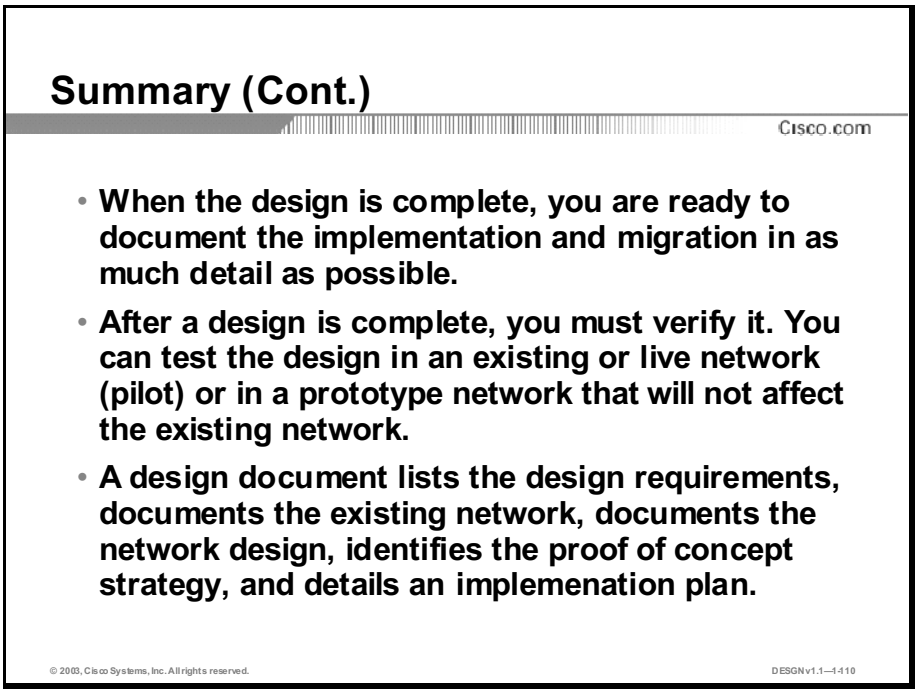

References

For additional information, refer to this resource:

- Oppenheimer, P. *Top-Down Network Design: A Systems Analysis Approach to Enterprise Network Design*. Indianapolis, Indiana: Macmillan Technical Publishing—Cisco Press; 1999.

Next Steps

For the associated case study and exercise, refer to the following sections that follow the Quiz:

- DJMP Industries Case Study Scenario
- Case Study 1: Network Upgrade
- Simulation 1: New Applications

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Starting at the first step, number each step below according to the sequence of the top-down design stages:

_____ 1. define upper OSI layers

_____ 2. analyze customer requirements

_____ 3. choose underlying technology

_____ 4. gather additional information during design

Q2) Select the statement that best describes the role of decision tables in the design process.

A) Decision tables introduce the concept of modularity into the design process.

B) Decision tables enable the designer to create a network design based only on data gathered from the network analysis.

C) Decision tables provide an accurate summary of the design, facilitating the decision on the basis of appropriateness.

D) Decision tables provide a means for decision making when multiple solutions exist for a given network issue.

Q3) Select four determinants of the design project scope. (Choose four.)

A) WAN upgrade

B) campus upgrade

C) network layer redundancy

D) data link layer redundancy

E) application layer redundancy

F) network redesign

Q4) A new bank network is being designed. Connections to bank subsidiaries are required to be redundant to reduce the possibility of connection outages. The importance of secure transactions was emphasized throughout the initial requirement documents and verbal communications with the customer. Number the design issues in the correct order according to their importance under these particular circumstances.

_____ 1. physical topology design

_____ 2. addressing design

_____ 3. security design

_____ 4. modularizing the network

Q5) Choose the two correct types of tools that should be used during the network design process. (Choose two.)

A) network modeling tools

B) network management tools

C) simulate and test tools

D) network implementation tools

Q6) Which three of these items should be present in an implementation plan? (Choose three.)

A) implementation description and references to the design document

B) old and new configurations of the network devices

C) roll-back procedure

D) time, date, and duration (optional cost)

E) application profile

Q7) A design that describes the introduction of IPSec encryption and authentication is required for an existing network's classical WAN module. Which approach would you use to verify the design?

A) pilot network

B) prototype network

C) live network

D) cable network

Q8) What three components does the design document usually include? (Choose three.)

A) design

B) existing cabling

C) design requirements

D) summary of L2 devices

E) implementation plan

Quiz Answer Key

Q1) Correct order: 2, 1, 4, 3

Relates to: Top-Down Approach to Network Design

Q2) D

Relates to: Decision Tables in Network Design

Q3) A, C, D, F

Relates to: Assessing the Scope of the Network Design Process

Q4) Correct order: 2, 3, 1, 4

Relates to: Using Structured Design Principles

Q5) A, C

Relates to: Using Structured Design Principles

Q6) A, C, D

Relates to: Planning a Design Implementation

Q7) B

Relates to: Building a Prototype or Pilot Network

Q8) A, C, E

Relates to: Documenting the Design

DJMP Industries Case Study Scenario

This case study analyzes the network infrastructure of DJMP Industries, a fictitious manufacturer of portable speed bumps. The company has provided you with a short description of the current situation and its plans. It is your job, as a network designer, to identify all of the company's requirements and data that will allow you to provide an effective solution.

Company Facts

DJMP Industries, a manufacturer of portable speed bumps, is an international company with headquarters in San Jose, California. The company is one of the leading suppliers of portable speed bumps in the world. The demand for the company's flagship product is constantly increasing and the company faces the need for tighter integration of its customers, partners, and suppliers into its information infrastructure.

The company's San Jose headquarters site consists of two buildings: the central building and building A. There are approximately 200 employees located at headquarters. DJMP Industries has three regional offices in the United States (Boston, Denver, and Houston) with 35, 50, and 50 employees in each office, respectively. Assume that the number of workstations is equal to the number of employees. Each regional office has a few smaller remote offices with up to 5 employees. Currently, there are 8 remote offices that are connected on-demand, using ISDN to a nearby regional office with the following distribution: Boston—2, Houston—3, Denver—3. Research and Development (R&D) is located in Houston while Engineering, Production, and Manufacturing are in Denver. Production is highly automated.

Current Situation

The current situation does not provide for future growth, and the company is seeking a scalable solution to replace and upgrade the existing infrastructure, especially communications. The most urgent task is to solidify the internal network infrastructure that connects the company offices with the headquarters. Because the company is adding new intranet and Internet-based applications to support its expanding business, the requirement for strategic communications infrastructure seems even more urgent.

The applications that the company is currently running include a set of custom-developed applications that run on IP and use a proprietary protocol running on top of TCP. The routing protocol used is Routing Information Protocol Version 1 (RIPv1), and the IP addressing is flat with no hierarchy.

The users frequently experienced slow response times. The company performed a monitoring and analysis of the traffic on links to regional offices and to the Internet. The analysis showed the average 24-hour link utilization as listed in the 24-Hour WAN-Link Utilization table (cells show percentage utilization).

24-Hour WAN-Link Utilization

From/To	San Jose Headquarters	Boston	Denver	Houston	Internet
San Jose Headquarters		32	45	42	10
Boston	25				
Denver	30				
Houston	32				
Internet	25				

Currently, the entire campus network at the headquarters is a shared Ethernet LAN, and servers are present in both buildings (central and A). There is severe congestion of the LAN, especially at peak hours. The placement of the servers presents additional problems. Although the WAN links are terminated at the central building, regional offices also access the servers in building A. A central firewall located in the central building of the campus provides Internet connectivity. The figure illustrates the core of DJMP's internal network (for clarity, the small remote offices are not shown).

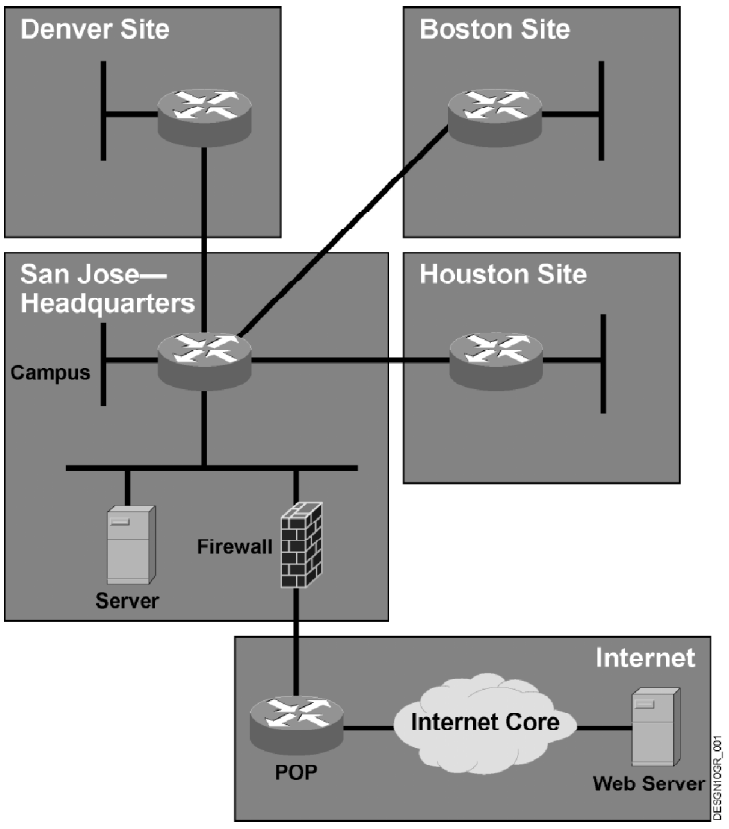

Core network of the DJMP Industries

Plans and Requirements

The company is extending its worldwide presence and will soon open two international offices, one in Europe (London) and one in Asia (Singapore), each with approximately 10 employees. The company is considering using the Internet as a connectivity option, with VPNs terminated at headquarters. In addition, the company would like to find a solution that would lower the cost of its international voice calls and is seriously considering a Voice over IP (VoIP) solution, initially for its international offices.

For its communications infrastructure, the company is seeking a solution for the headquarters campus. The company plans to restructure its campus LAN and is considering switched solutions, along with proper placement of servers. Because the effects of introducing the switched LAN are unknown, the company wants some proof of the technology's viability. Because of some bad experiences with network outages in the past, the LAN solution must be highly redundant.

With the expansion and modernization of its IT infrastructure, the company will introduce several new intranet and Internet-based applications based on the following Internet protocols: web browsing (HTTP), e-mail, Telnet, and FTP. The existing legacy applications will remain in use at least for the next year and a half; the traffic produced by the legacy applications is not expected to decrease. The IT department performed a survey of typical application usage patterns, and an external consultant identified five distinct profiles of the typical users who will use the new applications: engineer, researcher, e-commerce user, administrator, and salesperson. The table illustrates the estimated applications mix for a typical user, along with the expected intensity of usage (light, heavy).

Typical Users and Applications Mix

User	Web Browsing	E-Mail	Telnet	File Transfer
Engineer	Light	Light	Light	Light
Researcher	Heavy	Light		
E-Commerce User	Heavy			
Administrator	Light	Heavy	Light	
Salesperson	Light	Light		

The table describes the users based on location and the types of applications they use.

Distribution of Users of New Applications

Location	Engineer	Researcher	E-Commerce User	Administrator	Salesperson
San Jose Headquarters			10	25	
Boston			20		10
Denver	20	5		10	
Houston		25		10	

The introduction of new applications will result in an additional load on the company's links to regional offices. The tighter integration and growth of remote offices expected in the future will even further increase the traffic load on the WAN links. The company would like to upgrade the WAN infrastructure to provide sufficient bandwidth between the regional offices and headquarters and, at the same time, find a solution for better convergence during the network failures. Routing Information Protocol (RIP), as its current routing protocol, has proven to be inadequate. The company is aware of the drawbacks of its current IP addressing scheme and is seeking a better solution.

Case Study 1: Network Upgrade

Complete this case study to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the Course Introduction module
- DJMP Industries Case Study Scenario, presented at the end of Module 1 (Applying a Methodology to Network Design)
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the design methodology implementation details. Upon completing this case study, you will be able to meet these objectives:

- Document the company's requirements
- Document the existing network
- Identify the missing information
- Outline the major design areas for a given case

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, present in the Course Introduction module
- Case study solutions, presented in Appendix B, Case Study Solutions

Exercise Procedure

Complete these steps:

Step 1 Read the DJMP Industries Case Study Scenario completely before commencing the exercise. Allow 10 to 15 minutes for reading.

Step 2 Discuss the scenario with your group. Allow 10 minutes for a discussion.

Step 3 Document any information that you think is missing from the scenario and that you consider necessary for the design. List these items, and provide a brief comment for each. Use the Missing Items table.

Missing Items

Step 4 Outline the major design areas that you need to address in designing the solution for the given customer scenario. List the tasks, and provide a brief comment for each. Use the Major Design Tasks table.

Major Design Tasks

Step 5 Keep in mind that you just purchased an extremely powerful network simulation tool, and decide where it can help you in the design decisions. List some possible scenarios where you could evaluate the effects of the new design by using the simulation tool. Use the Possible Simulation Scenarios table.

Possible Simulation Scenarios

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class and you have justified any major deviations from the case study solution.

Simulation 1: New Applications

Complete this exercise to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the Course Introduction module
- DJMP Industries Case Study Scenario, presented at the end of Module 1 (Applying a Methodology to Network Design)
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the design methodology implementation details. Upon completing this simulation, you will be able to meet these objectives:

- Evaluate the effects of new applications on the WAN links
- Plan the capacity of the WAN links based on the simulation results

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, present in the Course Introduction module
- Case study solutions, presented in Appendix B, Case Study Solutions

Exercise Procedure

Read the New Applications scenario below, and answer the questions that appear in the text. Discuss possible answers, and describe your considerations in the classroom.

New Applications Scenario

The company has provided you with information about its existing network, planned applications, and number of users. In addition, you determined that all existing WAN links run at 64 kbps. You used a simulation tool to simulate the existing load on the WAN links. Afterward, you simulated the additional load imposed by the new applications and you graphed the results.

Initial Traffic

According to both the data that the customer provided to you (existing load on the links) and the topology of the existing network (as shown in the DJMP Industries Case Study Scenario), the simulation indicated that some links were saturated. The simulation focused on the 30-minute interval. The loaded links are marked with yellow (dashed) lines. The green (solid) links are not saturated. The threshold for considering the link to be loaded is set at 30 percent, and, for the heavily loaded link, the threshold is 60 percent.

The figure describes the company's network:

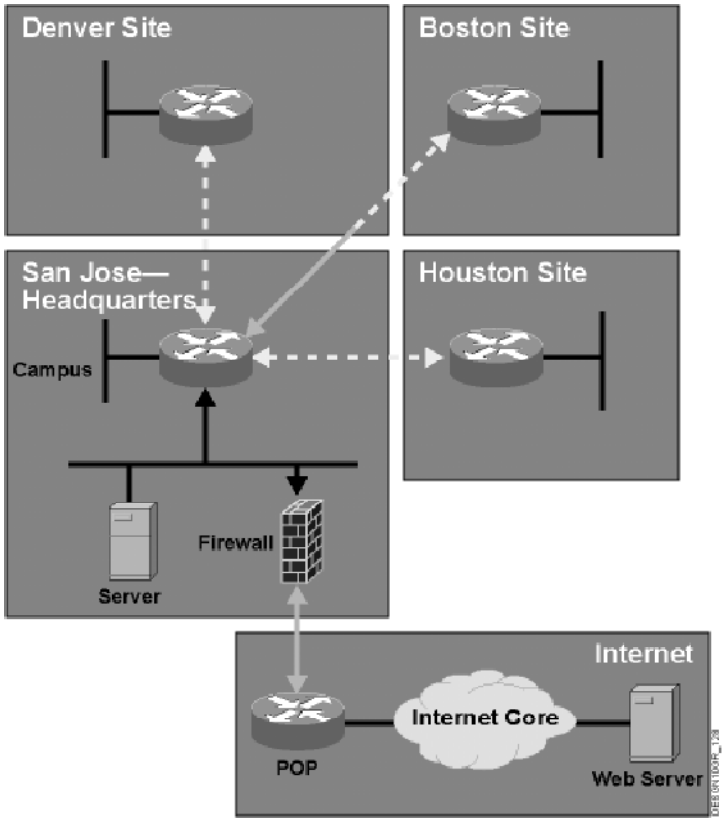

The graphs show the results of simulating 30 minutes of the existing traffic on a certain WAN link (both directions) and on the link to the Internet POP.

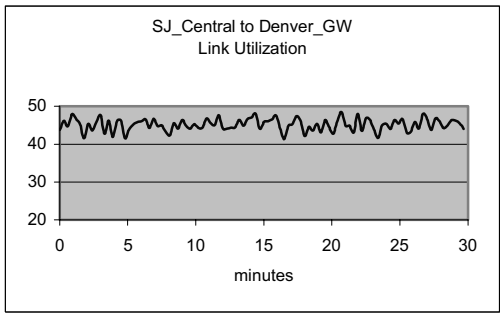

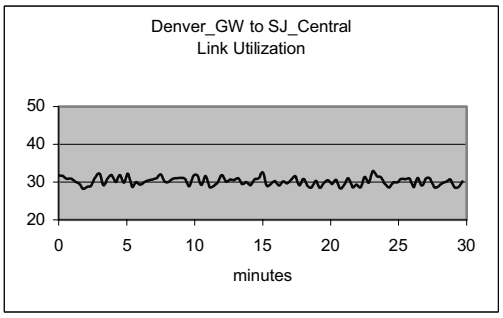

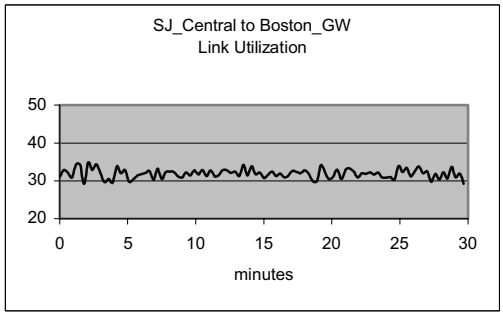

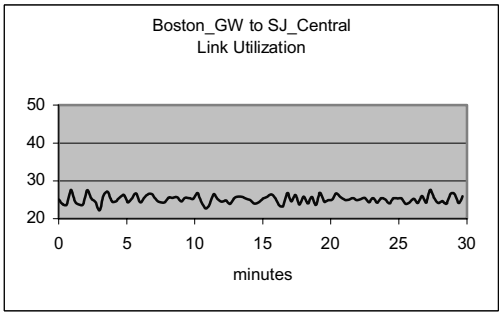

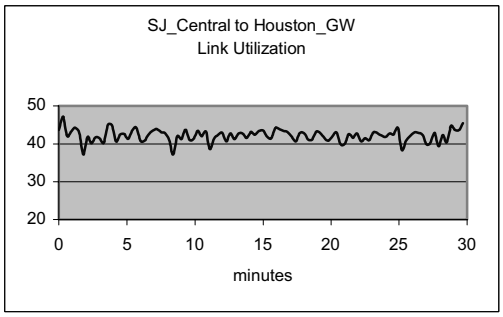

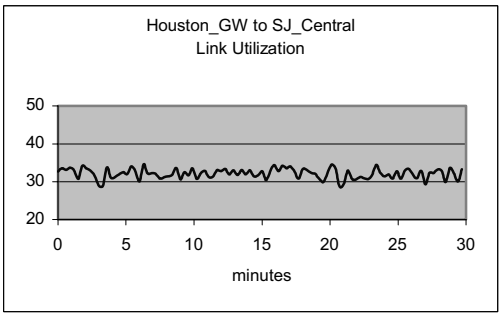

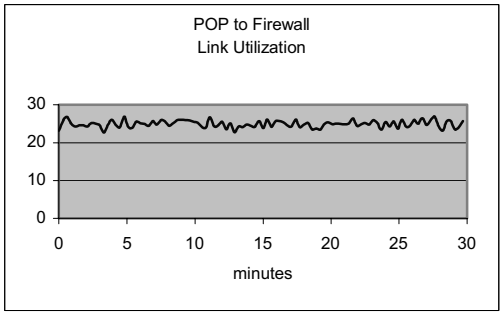

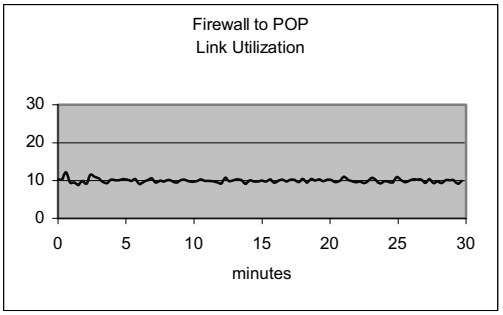

Q1) Observe the directions in which the load was higher. What can you determine from the results?

_____
_____
_____
_____
_____

Loaded Network—New Applications Introduced

You simulated the effect of new applications on the same topology with the same bandwidths of the WAN links. The load imposed with new applications and their respective users was determined from the data the customer supplied. You performed the simulation (30-minute interval), and the resulting graphs are shown in the figures. Observe carefully the results for all four simulated links (to all regional offices and to Internet POP).

Note: In a real-life situation, the observed interval should be longer and should include the peak-hour traffic that gives the most relevant results. However, because of the simulation tool's limitations (long calculation times), the observed interval is set to 30 minutes.

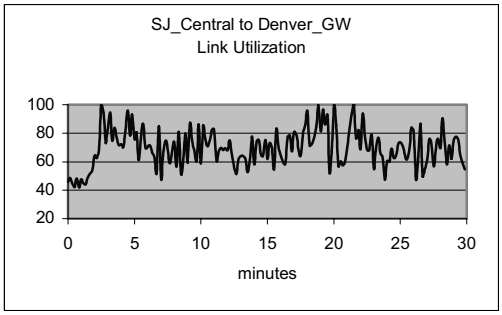

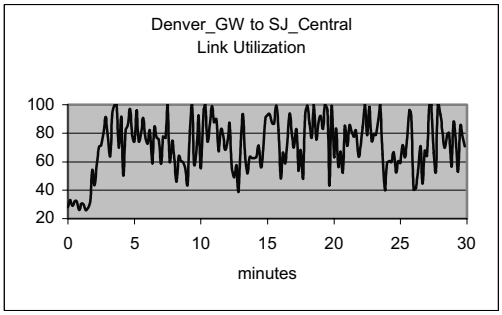

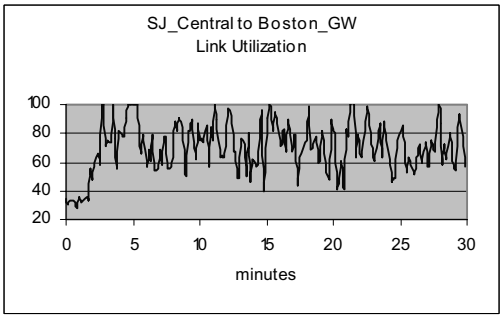

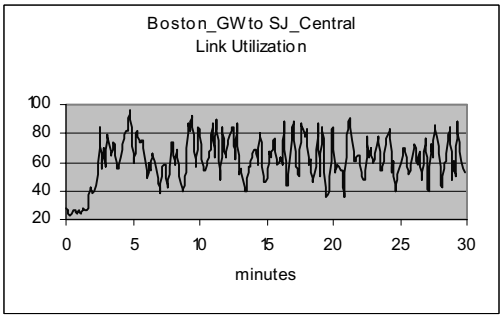

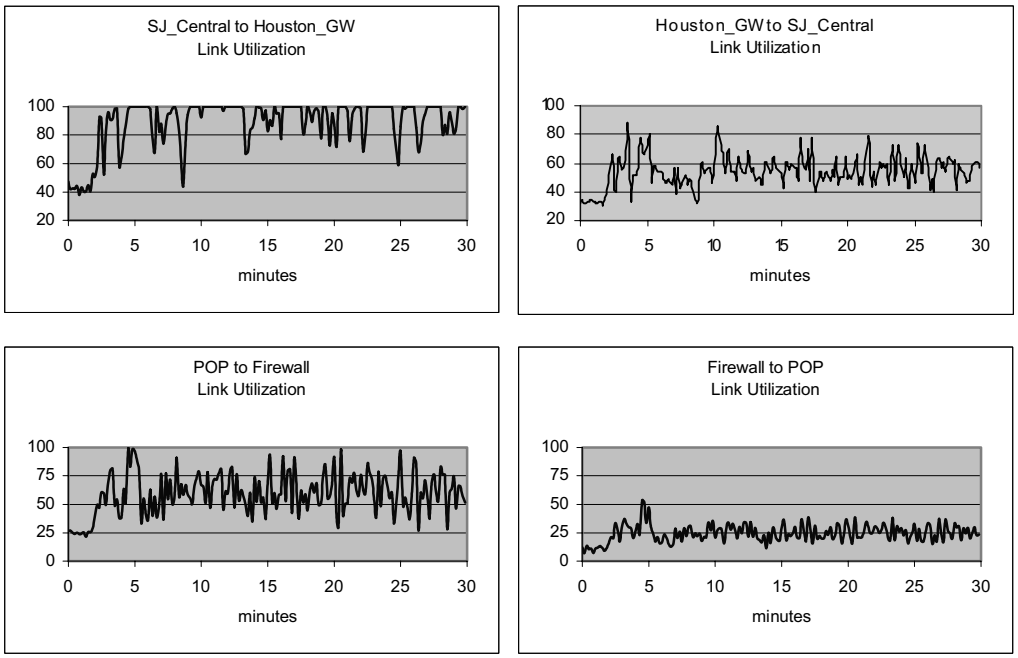

Q2) What can you determine from the results? Compare the planned number of users and applications for each of the regional offices. In which direction are the links saturated?

Q3) When you compare the results from the initial traffic simulation with the results from the simulations of the new applications, you observe that the traffic from Denver to San Jose Headquarters is now also significant. Why?

You observed the result of web (HTTP) traffic on the Houston–San Jose link. The average response times per web page were compared for two categories, light HTTP (smaller, less complex web pages) and heavy HTTP (larger web pages). The graphs are illustrated in these figures.

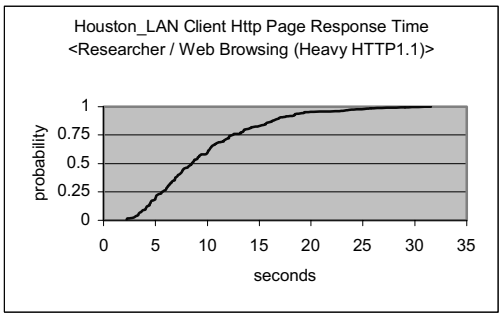

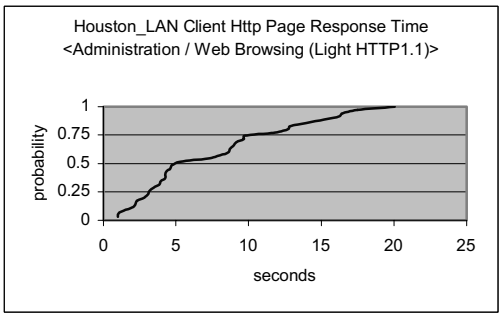

Q4) What can you determine from the graphs?

_____
_____
_____
_____
_____

Increased Link Speed: Houston–San Jose

You decided to increase the link speed on the Houston–San Jose connection to 128 kbps. You repeated the simulation and observed the link utilization along with the web page response times. The results are listed in these graphs.

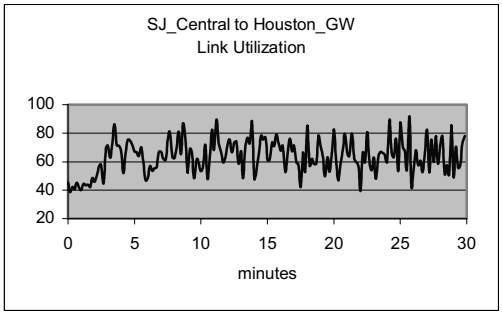

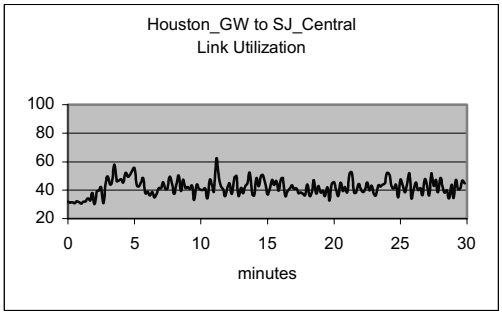

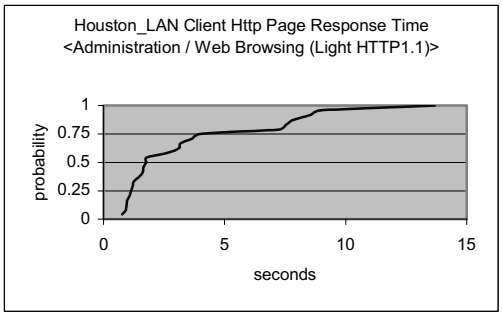

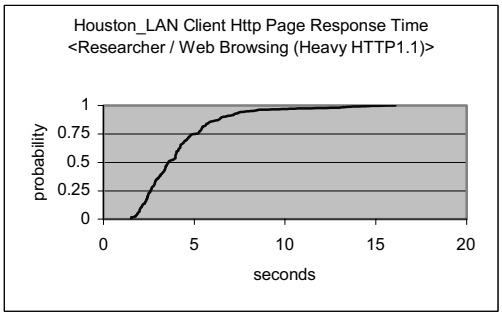

Q5) What do you observe from the graphs?

Module 2

Structuring and Modularizing the Network

Overview

Cisco has developed a blueprint that you can use to simplify the complexity of modern networks. This blueprint, called the Enterprise Composite Network Model, is a modular hierarchical approach to network design. The Enterprise Composite Network Model helps you to view the network based on the functional, logical, and physical components, thereby reducing complexity.

This module describes a basic network hierarchy. It then expands the basic network hierarchy to a more comprehensive, modular, hierarchical network of functional areas for use in designing modern networks.

Module Objectives

Upon completing this module, you will be able to explain the Enterprise Composite Network Model and describe its components.

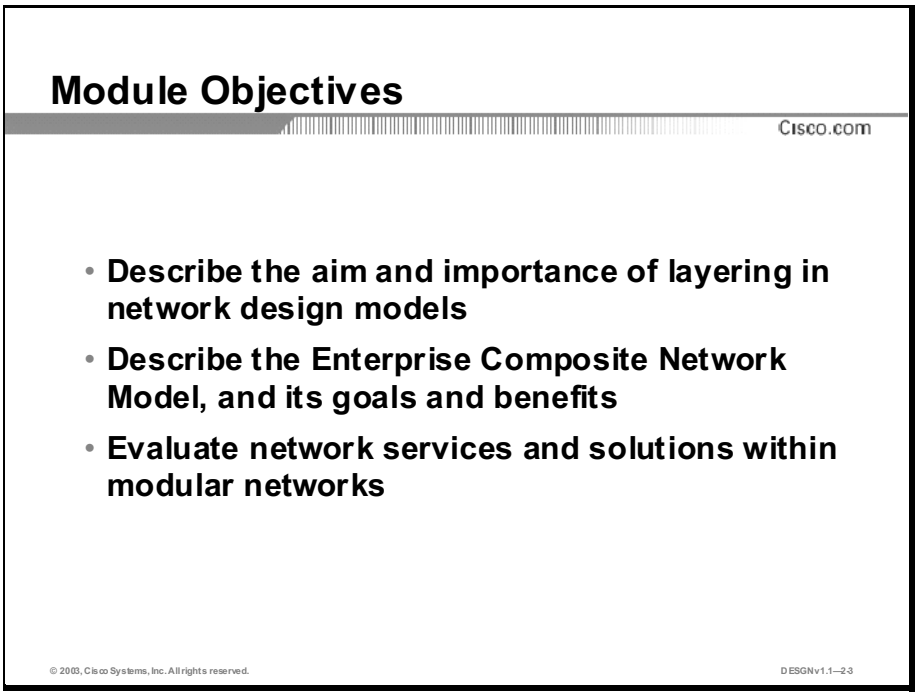

Module Outline

The outline lists the components of this module.

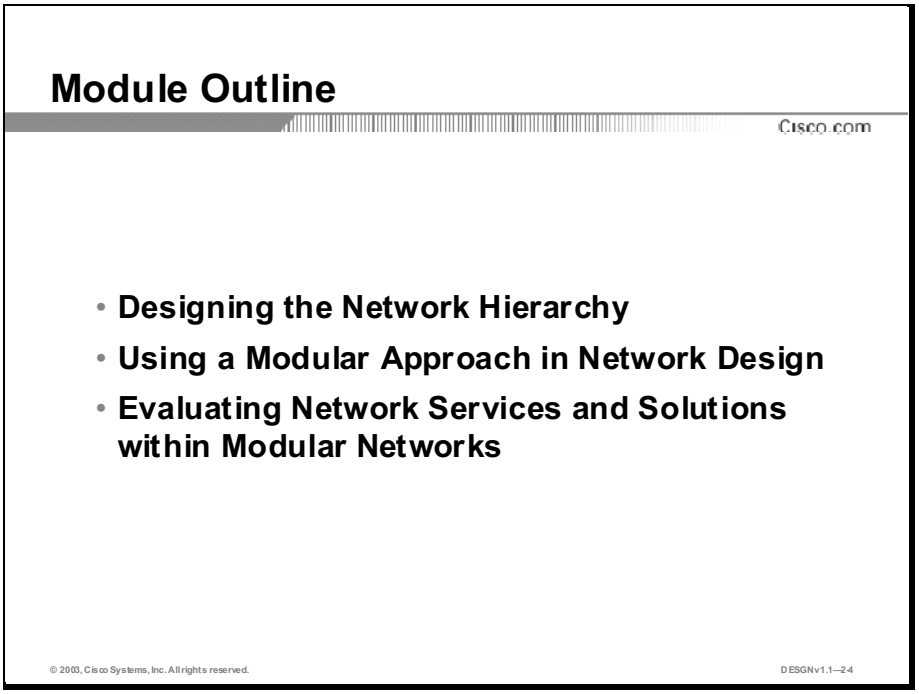

Designing the Network Hierarchy

Overview

The hierarchical network structure is composed of the access, distribution, and core layers. Each layer has its own functions used to develop a hierarchical network design. Historically used in the design of enterprise LAN and WAN data networks, a hierarchical model applies equally within the functional modules of the Enterprise Composite Network Model, discussed in the Using a Modular Approach in Network Design lesson.

Relevance

The hierarchical model serves as the predecessor to the Enterprise Composite Network Model. By understanding its functions, you will better understand today's Enterprise Composite Network Model. You will also understand how to design its use of functional areas containing hierarchical modularity for designing complex networks.

Objectives

Upon completing this lesson, you will be able to describe the aim and importance of layering in network design models. This includes being able to meet these objectives:

- Describe each layer in the hierarchical network model
- Describe the role and functions of the network access layer
- Describe the role and functions of the network distribution layer
- Describe the role and functions of the network core layer

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in Cisco IOS software.

Outline

The outline lists the topics included in this lesson.

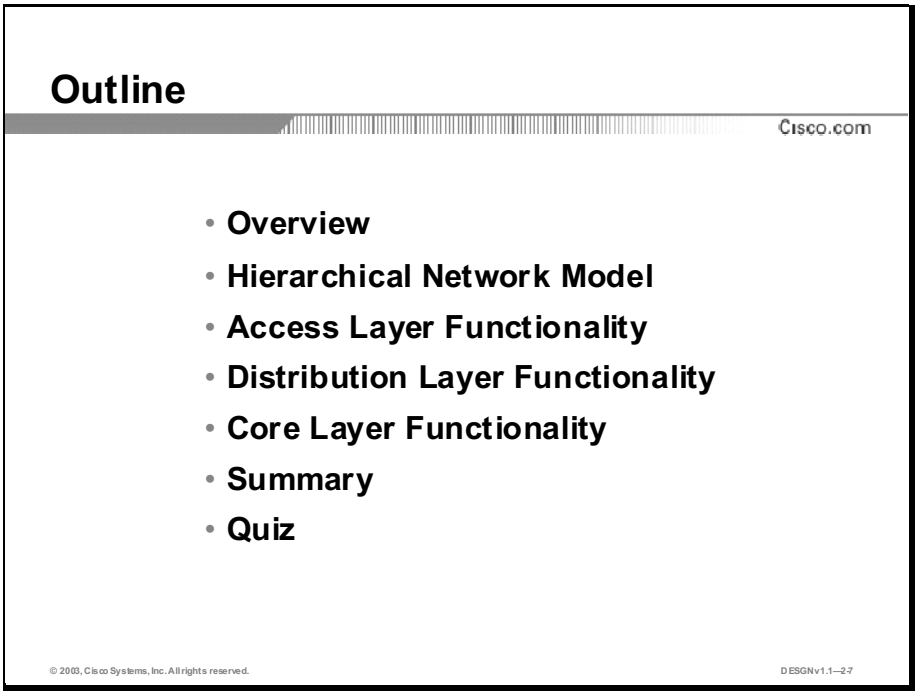

Hierarchical Network Model

The hierarchical network model provides a modular view of a network, making it easier to design a network. This topic discusses the significance of hierarchical layering used in the Enterprise Composite Network Model.

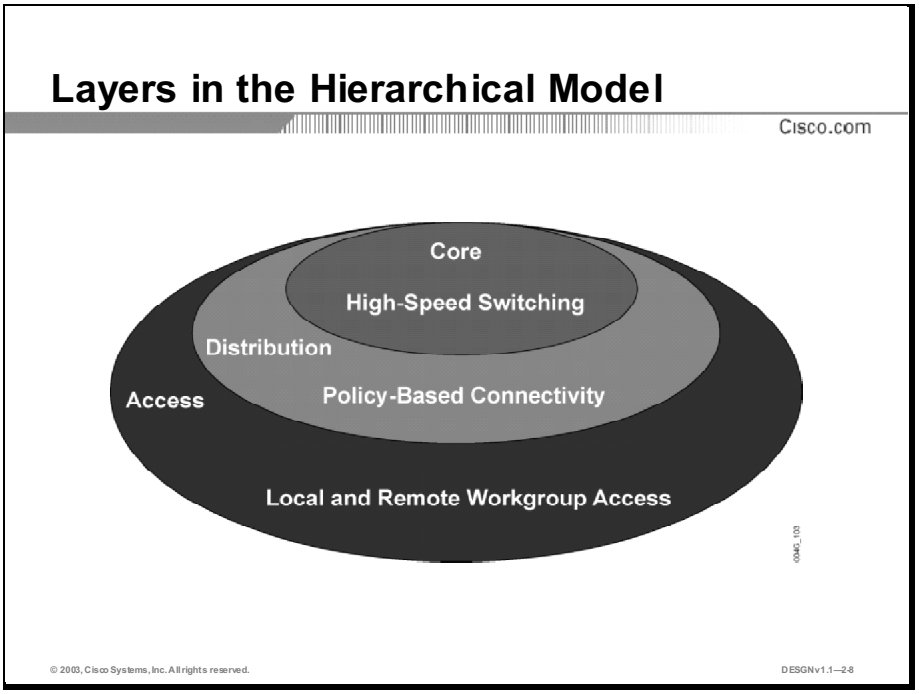

Modern networks are extremely complex. Organizations continue to increase the requirements for bandwidth, reliability, and functionality from their networks. If you are a network designer, you are probably concerned with how you will keep pace with the constant changes in the internetworking industry.

You can easily break networks down into smaller functional components because networks have natural physical, logical, and functional boundaries. The hierarchical model provides a tool to ensure that the network design is scalable, reliable, available, responsive, efficient, adaptable, flexible, accessible but secure, and manageable.

The hierarchical model includes three discrete layers of internetworking:

- **Access layer:** Provides local and remote workgroup and user access to the network
- **Distribution layer:** Provides policy-based connectivity
- **Core (backbone) layer:** Provides high-speed transport to satisfy the connectivity and transport needs of the distribution layer devices

Each layer in the model focuses on specific functions to help you choose the right systems and features. The devices of all three hierarchical layers must implement the physical and data link layers of the OSI reference model to achieve basic connectivity within the network. Network designers typically implement the layers as distinct physical entities, but that is not always necessary. You can implement each layer within a network device, represented by physical media. You can even omit a particular layer, but you should maintain hierarchy.

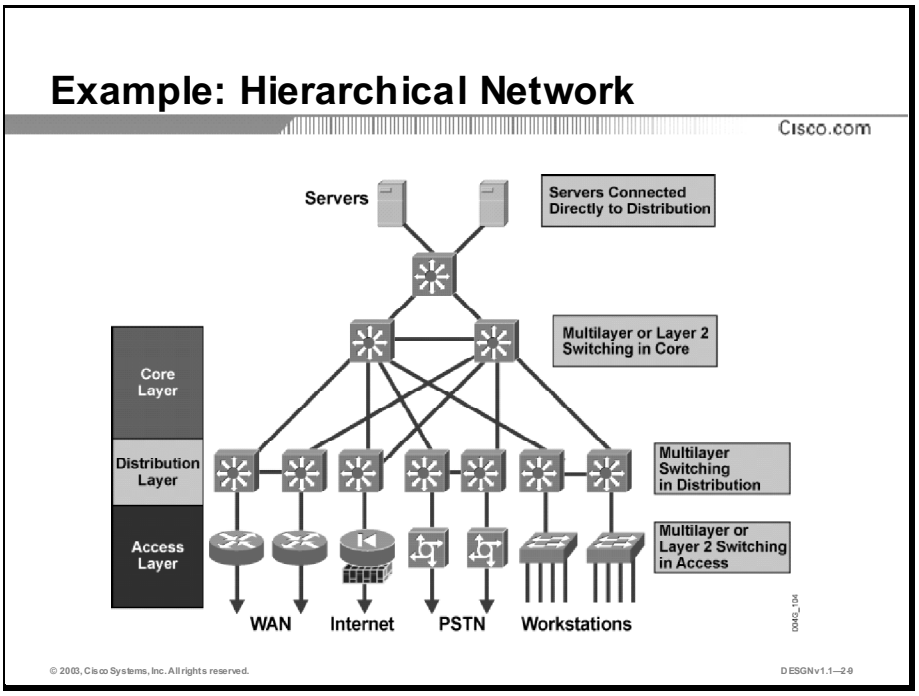

The three-layer hierarchical model specifies this general mapping scheme:

- **Access layer:** Access layer devices control traffic by localizing service requests to the access media. Access layer devices must also provide connectivity without compromising network integrity. For example, the devices at the access layer must detect whether a telecommuter dialing in is legitimate, yet must also require minimal telecommuter authentication steps.
- **Distribution layer:** Distribution layer devices control access to resources that are available at the core layer and must, therefore, make efficient use of bandwidth. In addition, a distribution layer device must address the quality of service (QoS) needs for different protocols by implementing policy-based traffic control to isolate backbone and local environments. Policy-based traffic control enables you to prioritize traffic to ensure the best performance for the most time-critical and time-dependent applications.
- **Core layer:** Core layer devices provide services that optimize communication transport within the network. In addition, core layer devices are expected to provide maximum availability and reliability. Core layer devices should be able to maintain connectivity when the circuits connecting them fail. A fault-tolerant network design ensures that failures do not have a major impact on network connectivity.

Access Layer Functionality

The purpose of the access layer is to grant user access to network resources. This topic describes the role and functions of the access layer.

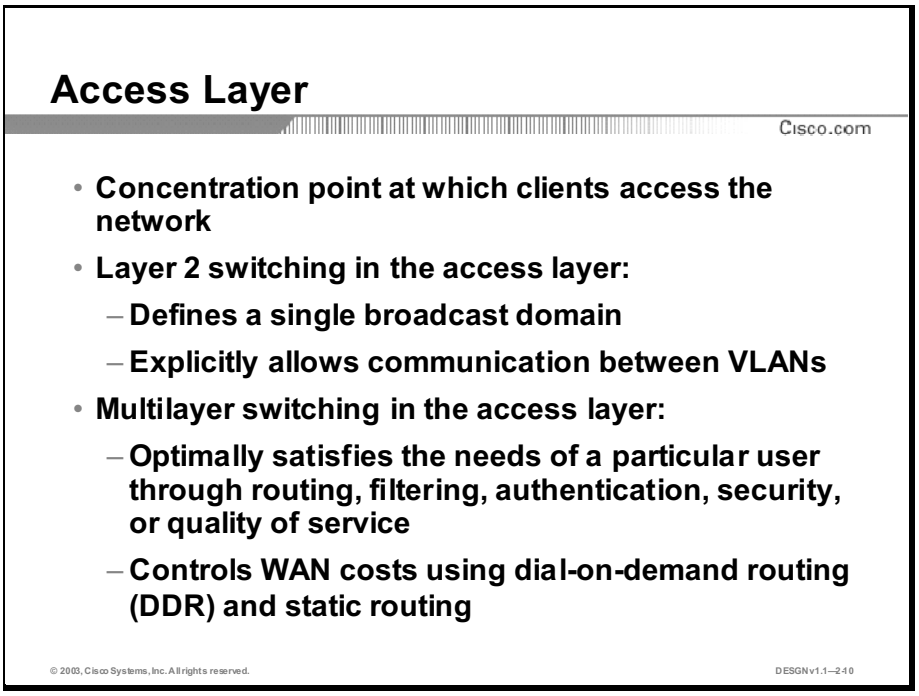

The access layer has these characteristics:

- In the campus environment, the access layer incorporates shared, switched, or subnetted LAN access devices with ports available to workstations and servers.
- In the WAN environment, the access layer provides sites with access to the corporate network via a wide-area technology, such as Frame Relay, ISDN, leased lines, digital subscriber line (DSL) over traditional telephone copper lines, or coaxial cable.
- The access is granted only to authenticated users or devices.

You can provide access to end users as part of two different scenarios:

- **Using Layer 2 switching (campus):** The access layer aggregates end user-switched 10/100 ports and provides Fast Ethernet, Fast EtherChannel, and Gigabit Ethernet uplinks to the distribution layer. Applying more than one VLAN will satisfy the connectivity requirements and reduce the size of the broadcast domains, each with its own IP subnet. Multiple VLANs allow each VLAN to support its own spanning tree and alternate paths in case of failure. Transportation between the access layer switches and the distribution layer switches is based on a Layer 2 trunking system, either Inter-Switch Link (ISL) or 802.1Q. A multilayer distribution switch can also provide the inter-VLAN communication for the access layer.
- **Using multilayer switching (WAN):** Access routing provides access to remote office environments using wide-area technologies combined with features such as route propagation, packet filtering, authentication, and so on. In a dial-up connection environment, you can use dial-on-demand routing (DDR) and static routing to control costs.

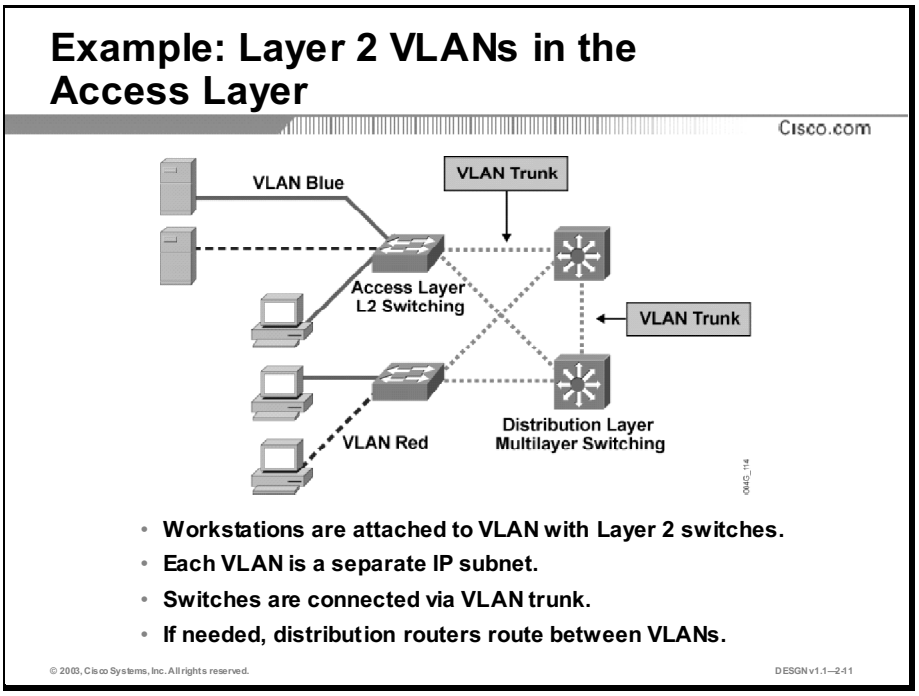

In the example shown in the figure, the access layer aggregates end user-switched 10/100 ports and provides uplinks to the distribution layer. Two small VLANs satisfy the connectivity requirements and reduce the size of the broadcast domains (each of them has its own IP subnet). The method of transportation that is deployed between the access layer switches and the distribution layer switches is based on a Layer 2 trunking system, such as ISL or 802.1Q. All switches are aware of the VLAN presence, making it possible for each of them to select a primary path (per VLAN Spanning Tree) and to find an alternative in case of various failures.

Spanning tree features such as UplinkFast and PortFast further improve the design. UplinkFast unblocks the blocked uplink port on a switch and transitions it to the forwarding state immediately, without transitioning the port through the listening and learning states. PortFast causes a port attaching the workstations to enter the spanning-tree forwarding state immediately, bypassing the listening and learning states.

Note:	When connectivity requirements exist between separate VLANs, a distribution multilayer switch or router may act as a VLAN transit node.

Distribution Layer Functionality

The distribution layer aggregates the wiring closets and uses a combination of Layer 2 and multilayer switching to segment workgroups and isolate network problems, preventing them from impacting the core layer. It is also used to enforce policy within the network. This topic describes the role and functions of the distribution layer.

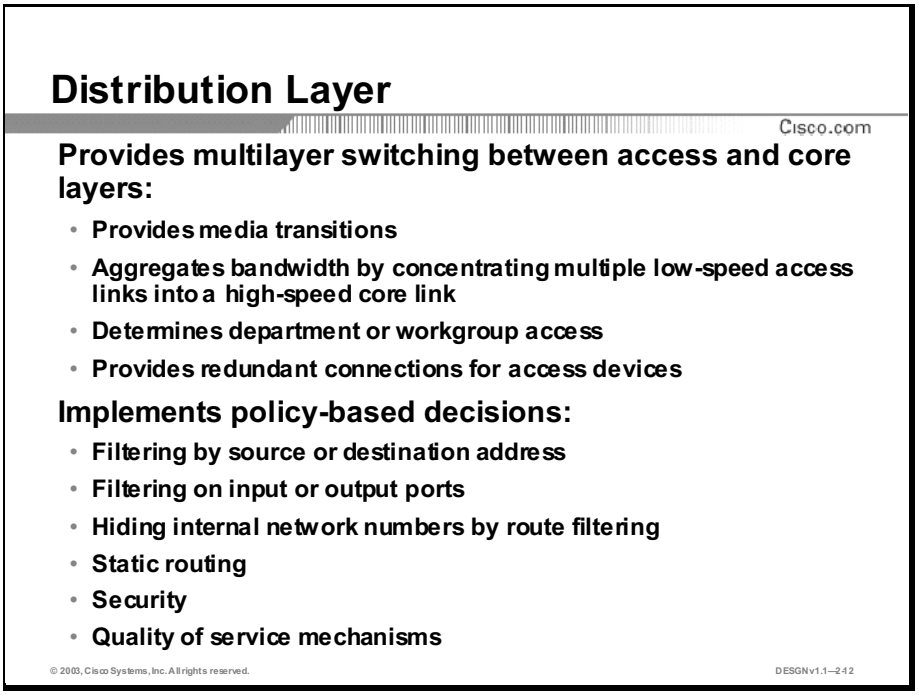

The distribution layer represents a routing boundary between the access and core layers and is the place where routing and packet manipulation are performed. The distribution layer connects network services to the access layer and implements policies regarding security, traffic loading, and routing.

The distribution layer is often the layer that delineates broadcast domains, although you can define broadcast domains at the access layer.

The distribution layer allows the core layer to connect diverse sites while maintaining high performance. To maintain good performance in the core, the distribution layer can redistribute between bandwidth-intensive access layer routing protocols and optimized core routing protocols.

To further improve routing protocol performance, the distribution layer can summarize routes from the access layer. For some networks, the distribution layer offers a default route to access layer routers and runs dynamic routing protocols when communicating with core routers.

In short, the distribution layer is the layer that provides policy-based connectivity. In terms of IP routing, it represents a redistribution point between routing domains or the demarcation between static and dynamic routing protocols. It can also be the point at which tasks such as controlled routing decisions and filtering take place.

Example: Distribution Layer in the Routed Campus Network

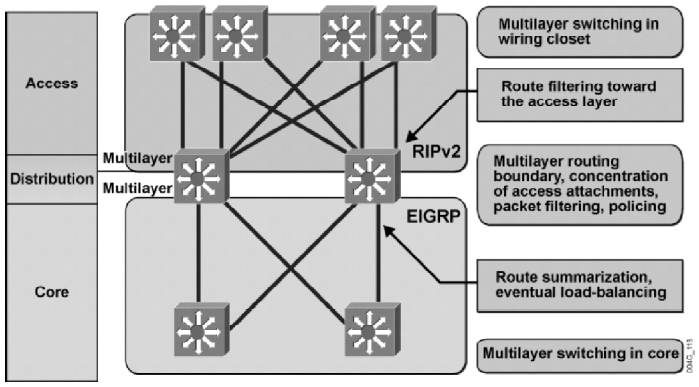

The figure shows the switching types and positioning of various features around the distribution layer in an example network. The distribution layer has highly redundant connectivity, both toward the access layer and toward the core layer. The distribution layer in a routed campus network has these characteristics:

- Multilayer switching is used toward the access layer.
- Multilayer switching is performed in the distribution layer and extended toward the core layer.
- Two-way route redistribution is used to exchange the routes between the routing processes.
- Route filtering is configured on interfaces toward the access layer.
- Route summarization is configured on interfaces toward the core layer.

Core Layer Functionality

The core layer is a high-speed backbone, which is designed to switch packets as fast as possible. This topic describes the role and functions of the core layer.

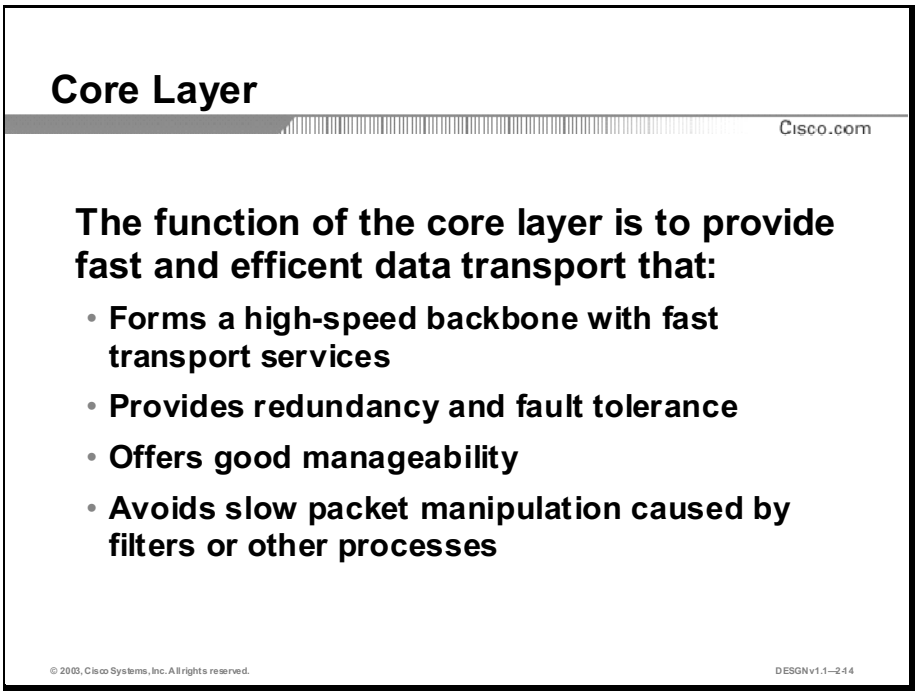

Because the core is critical for connectivity, it must provide a high level of redundancy and must adapt to changes very quickly. A full mesh is strongly suggested. A well-connected partial mesh with multiple paths from each device is effectively a design requirement. The core layer should not perform any packet manipulation, such as checking access lists and filtering, which would slow down the switching of packets.

Core devices are most reliable when they can accommodate failures by rerouting traffic and respond quickly to changes in the network topology. The core devices must be able to implement scalable protocols and technologies, alternate paths, and load balancing.

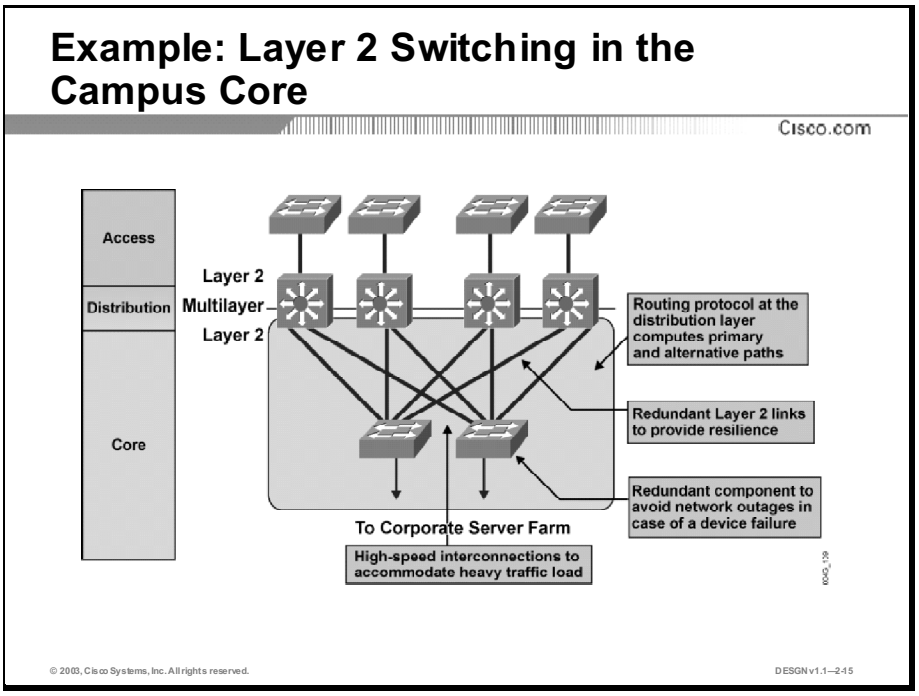

The figure shows Layer 2 switching in the campus core of an example network. A typical packet flow between access sites would follow these steps:

Step 1 A packet is Layer 2 switched toward the distribution switch.

Step 2 The distribution switch performs multilayer switching toward a core interface.

Step 3 The packet is Layer 2 switched across the LAN core.

Step 4 The receiving distribution switch performs multilayer switching toward an access LAN.

Step 5 The packet is Layer 2 switched across access LAN to the destination host.

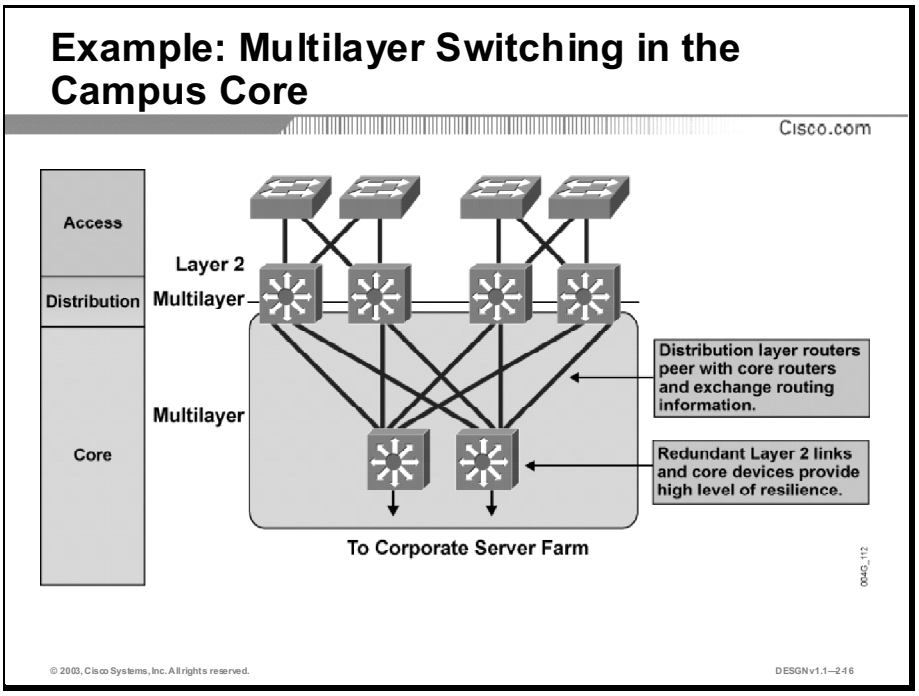

The figure shows multilayer switching in the campus core of an example network. A typical packet flow between access sites would follow these steps:

Step 1 A packet is Layer 2 switched toward the distribution switch.

Step 2 The distribution switch performs multilayer switching toward a core interface.

Step 3 The packet is multilayer-switched across the LAN core.

Step 4 The receiving distribution switch performs multilayer switching toward an access LAN.

Step 5 The packet is Layer 2 switched across the access LAN to the destination host.

Because core devices accommodate failures by rerouting traffic and respond quickly to changes in the network topology, and because there is no cost in performance for routing in the core, most implementations have multilayer switching in the core layer. The core layer can then more readily implement scalable protocols and technologies, alternate paths, and load balancing.

Summary

This topic summarizes the key points discussed in this lesson.

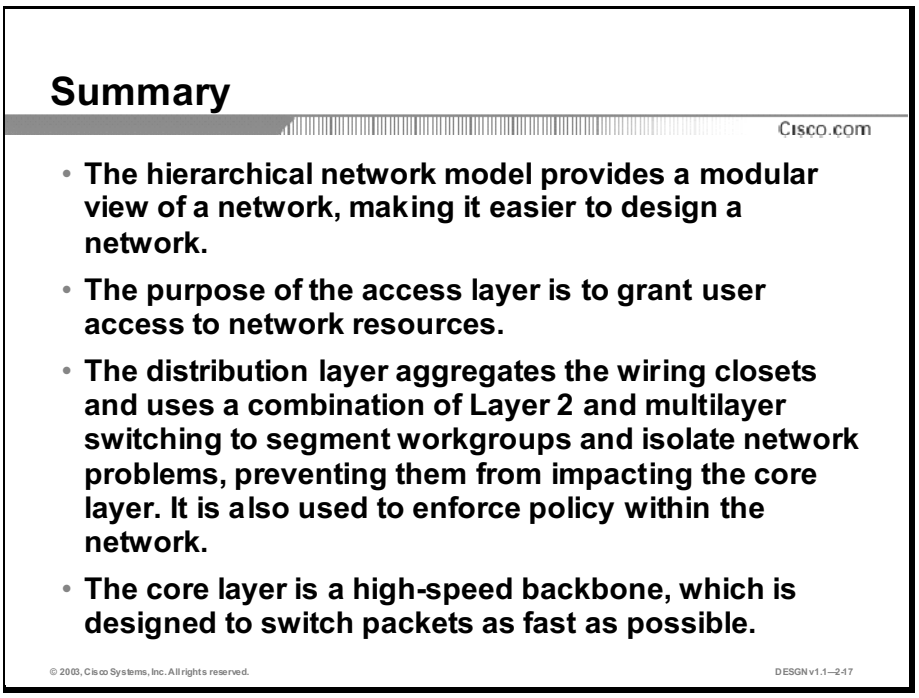

References

For additional information, refer to these resources:

- Oppenheimer, P. *Top-Down Network Design: A Systems Analysis Approach to Enterprise Network Design*. Indianapolis, Indiana: Macmillan Technical Publishing—Cisco Press; 1999.
- "Internetworking Design Basics" chapter of *Cisco Internetwork Design Guide* at http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

This figure presents a sample hierarchically structured network. Some of the devices are marked with letters.

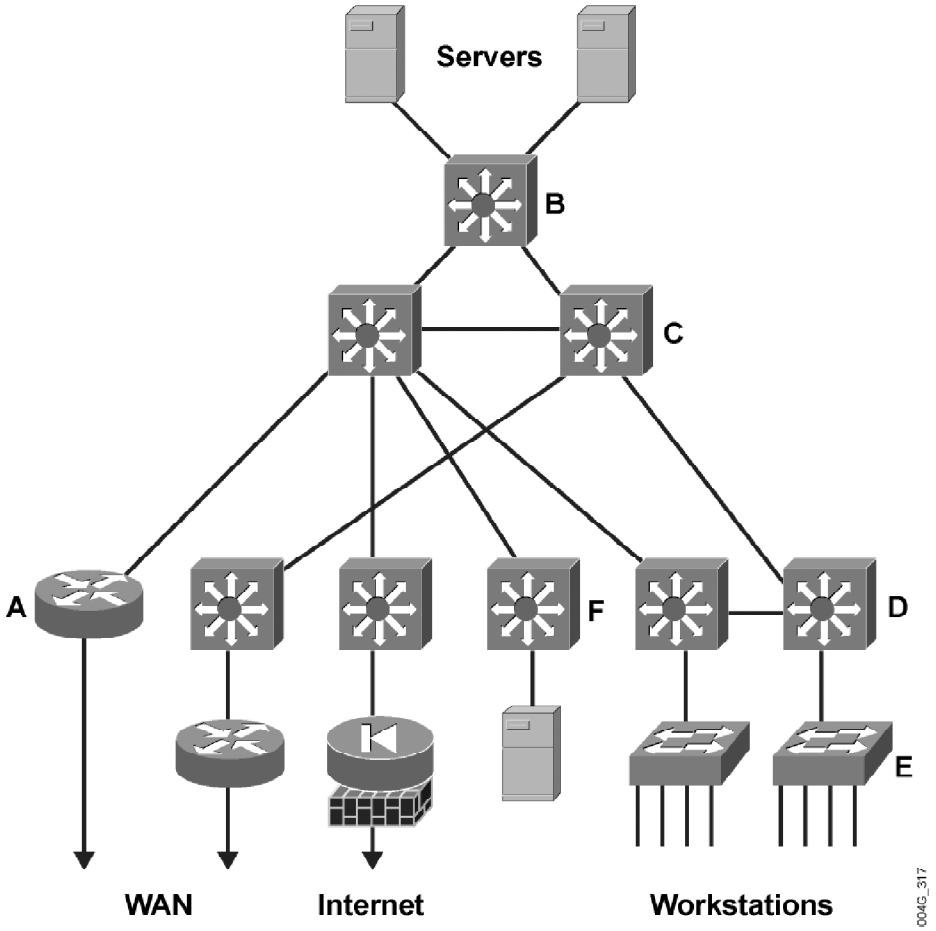

Q1) How would you map the marked devices to the access, distribution, and core layers in this figure?

Layer	Device (letter)

This figure represents a sample network. Devices are grouped within boxes identified by a letter.

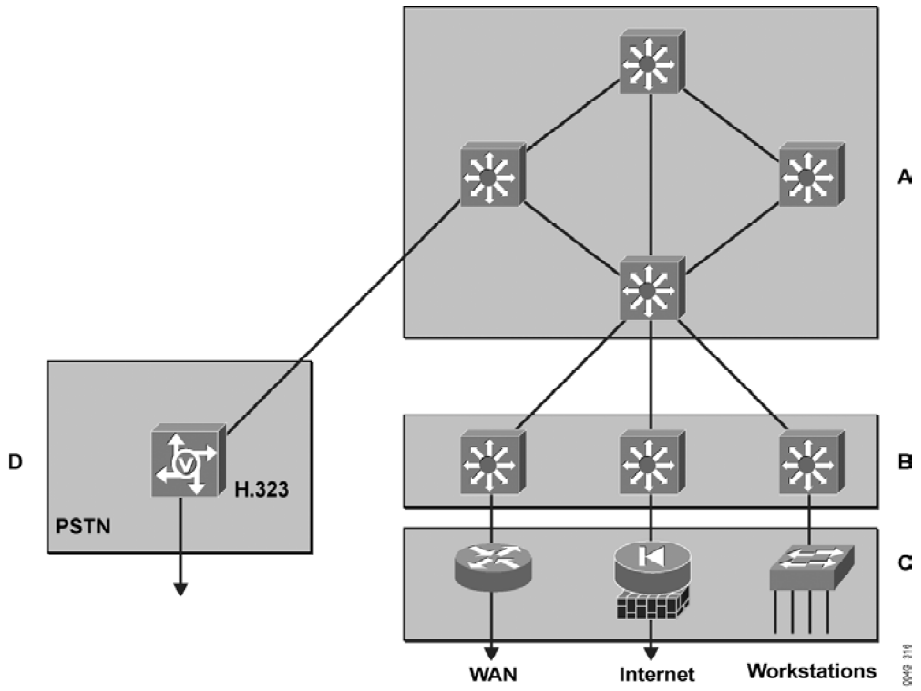

Q2) Identify the position of the access layer in the network diagram by selecting two identified boxes. (Choose two.)

A) A

B) B

C) C

D) D

Q3) Which three statements describe the key features of the distribution layer? (Choose three.)

A) The distribution layer aggregates access layer links.

B) The distribution layer represents a routing boundary between the access and core layers.

C) The distribution layer provides policy-based connectivity.

D) The distribution layer provides connectivity to the network.

E) The distribution layer concentrates user access by providing fast throughput.

F) The distribution layer focuses on fast packet switching.

Q4) What are the three roles of the core layer in a LAN design? (Choose three.)

A) provides high-speed data transport

B) performs packet filtering

C) serves as a fast convergent infrastructure with a high level of redundancy

D) avoids data manipulation

E) performs mainly policy-based decisions

F) provides access to the network

Quiz Answer Key

Q1) Access=A, B, E, F
Distribution=A, B, D, F
Core=C

Relates to: Hierarchical Network Model

Q2) C, D

Relates to: Access Layer Functionality

Q3) A, B, C

Relates to: Distribution Layer Functionality

Q4) A, C, D

Relates to: Core Layer Functionality

Using a Modular Approach in Network Design

Overview

The Enterprise Composite Network Model provides a blueprint or framework for designing networks. The modularity built into the model allows flexibility in network design and facilitates implementation and troubleshooting. This lesson describes the basic functional areas and modules of an enterprise network and focuses on the design considerations within a network module and between the modules.

Relevance

A modular network design model provides flexibility to develop an effective enterprise network design.

Objectives

Upon completing this lesson, you will be able to describe the Enterprise Composite Network Model and its goals and benefits. This includes being able to meet these objectives:

- Describe the modules of the Enterprise Composite Network Model
- Describe the components and functions of Enterprise Campus modules
- Describe the components and functions of Enterprise Edge modules
- Explain the components and functions of Service Provider Edge modules

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in IOS software

Outline

The outline lists the topics included in this lesson.

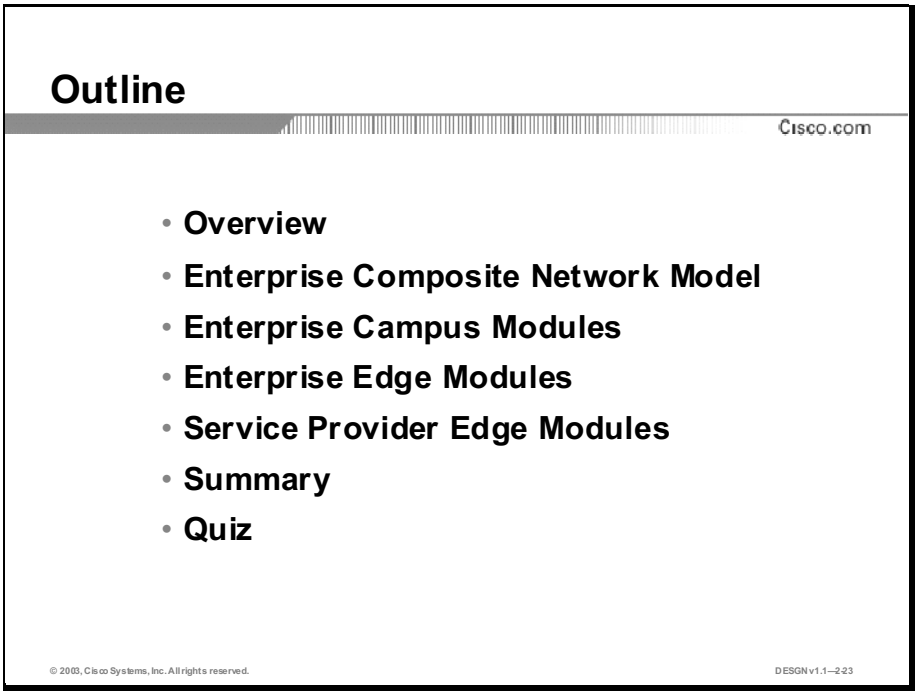

Enterprise Composite Network Model

To scale the hierarchical model (access, distribution, and core) Cisco developed the Enterprise Composite Network Model, which reduces the enterprise network into further physical, logical, and functional boundaries. Hierarchy is embedded as required into each module of the Enterprise Composite Network Model. This topic describes the Enterprise Composite Network Model and its goals and benefits.

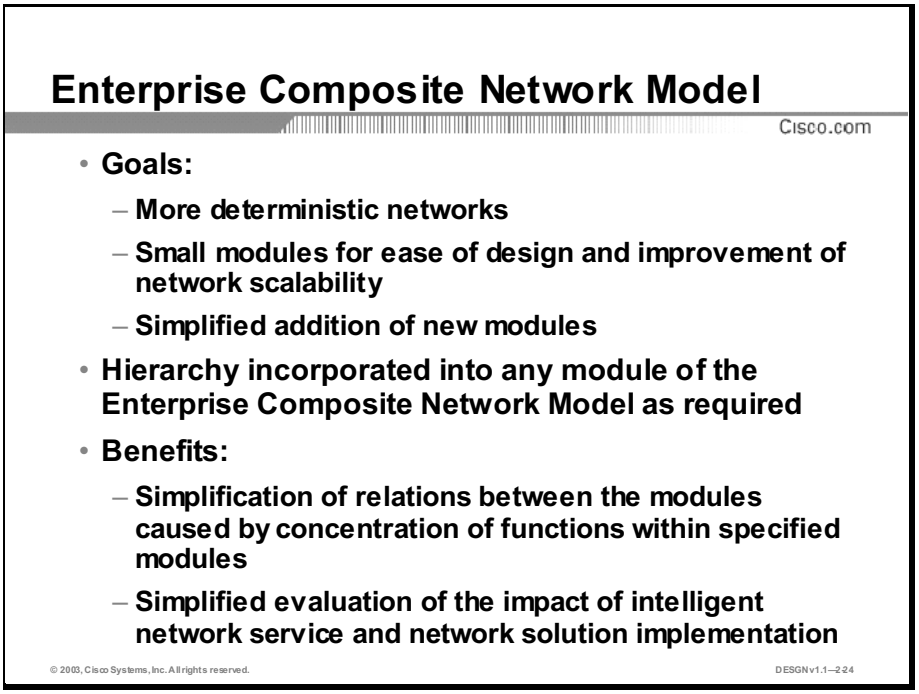

These are the goals of the Enterprise Composite Network Model:

- A deterministic network with clearly defined boundaries between modules. The model has clear demarcation points. A network designer knows exactly what traffic is allowed into and out of these demarcation points.
- Ease of design and increased network scalability achieved by using a “divide and conquer” method of problem solving.
- Simplified scalability. Adding a building to the campus, a remote office to a WAN, or servers to the server farm becomes a simpler task.

The Enterprise Composite Network Model enables network designers to concentrate on each module and on the relationships between the modules. There are clear boundaries with well-defined physical and logical points of entry that provide clear locations for policy enforcement.

This model provides additional integrity in network design. This integrity allows the designer to evaluate any network solution (for example, IP telephony, Content Networking [CN], or Storage Networking) and any intelligent network service (for example, security, quality of service [QoS], or network management) with respect to each network module and in relation to the overall network infrastructure.

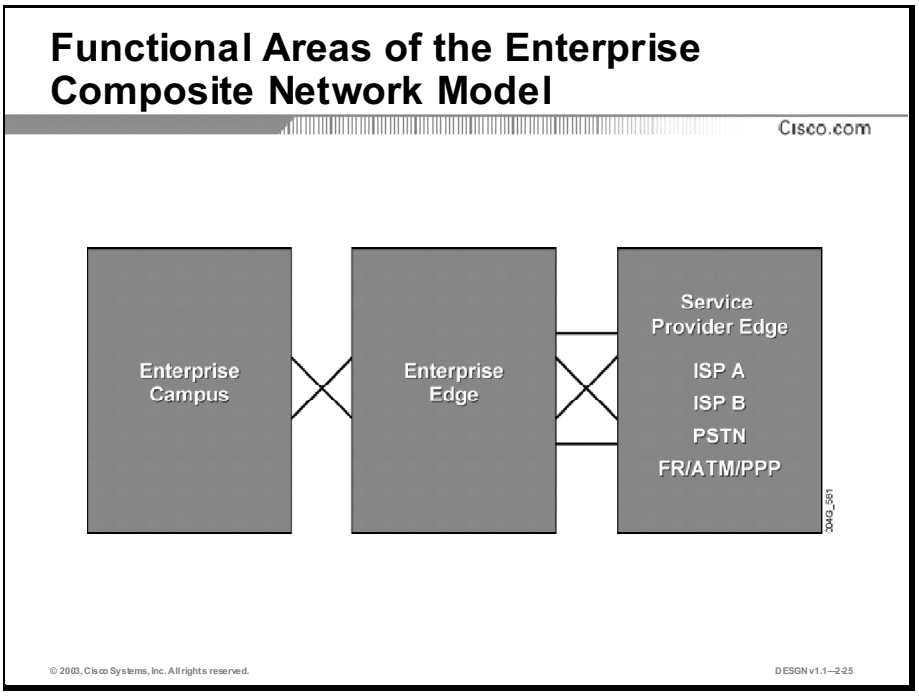

The Enterprise Composite Network Model introduces high-level modularity into the network structure. The basic idea is to divide the entire network into functional areas containing network modules. These internal network modules still maintain the hierarchical concept of the core, distribution, and access layers as needed.

Note:	Access, distribution, and core layers can appear in any functional area of the Enterprise Composite Network Model.

The Enterprise Composite Network Model has three major functional areas:

- **Enterprise Campus:** Includes the modules required to build a highly robust campus network that provides reliability, availability, scalability, and flexibility. This area contains all the network elements for independent operation within one campus location. No remote connections or Internet access are provided in this functional area. An enterprise may have more than one campus.
- **Enterprise Edge:** Aggregates the connectivity from the elements at the edge of the enterprise campus network. The Enterprise Edge functional area filters traffic from the edge modules and routes it into the Enterprise Campus functional area. The Enterprise Edge functional area contains all the network elements for efficient and secure communication between an Enterprise Campus and remote locations, business partners, mobile users, and the Internet.
- **Service Provider Edge:** The modules in this functional area are not implemented in an enterprise. The Service Provider Edge modules enable communication with other networks using WAN technologies and Internet service providers (ISPs).

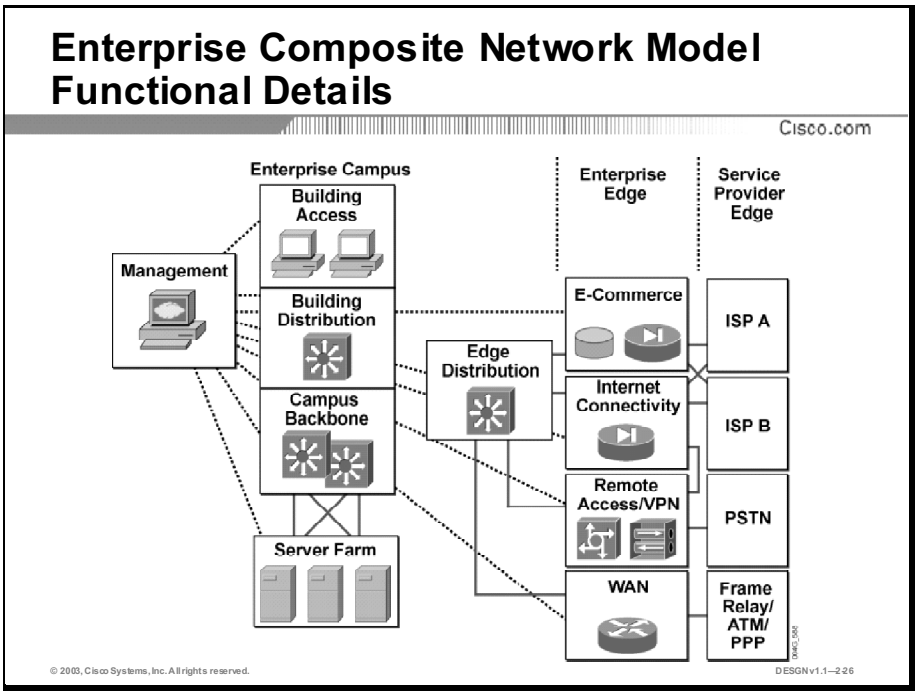

The second layer of modularity represents a more granular view of the modules within each functional area. Each module performs specific roles in the network and has specific requirements, but the size of the module is not meant to reflect its scale in a real network. For example, the Edge Distribution module in the Enterprise Campus represents the aggregation devices for remote sites to access Enterprise Campus devices.

The Enterprise Composite Network Model allows network designers to focus only on a selected module and its functions. Designers can describe each network solution and intelligent network service on a per-module basis but validate each as part of the complete enterprise network design.

To achieve scalability, you can add modules. In addition, the modules may have submodules.

Guidelines

Use these guidelines to create an enterprise network:

- Divide the network into the Enterprise Campus and Enterprise Edge functional areas, where the Enterprise Campus functional area includes all devices and connections within one location, and the Enterprise Edge functional area covers all communications with remote locations and the Internet.
- Define clear boundaries between the Enterprise Campus and Enterprise Edge networks.

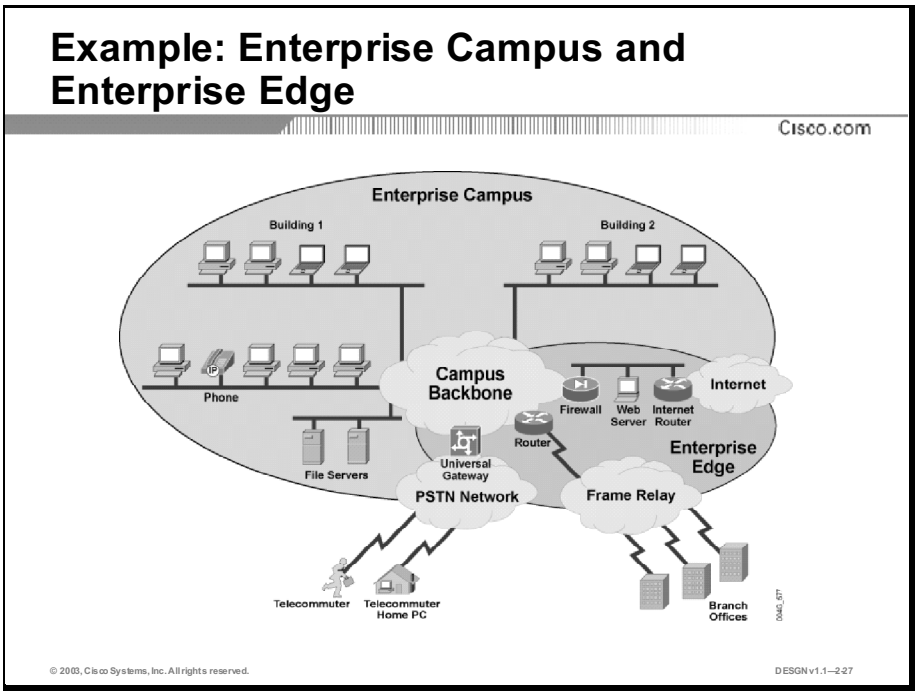

The figure shows the first step when designing an example enterprise network. The network is divided into the Enterprise Campus functional area and the Enterprise Edge functional area.

Enterprise Campus Modules

The Enterprise Campus functional area includes the Campus Infrastructure module, the Network Management module, the Server Farm module, and the Edge Distribution module. This topic describes the Enterprise Campus functional area and each of its component modules.

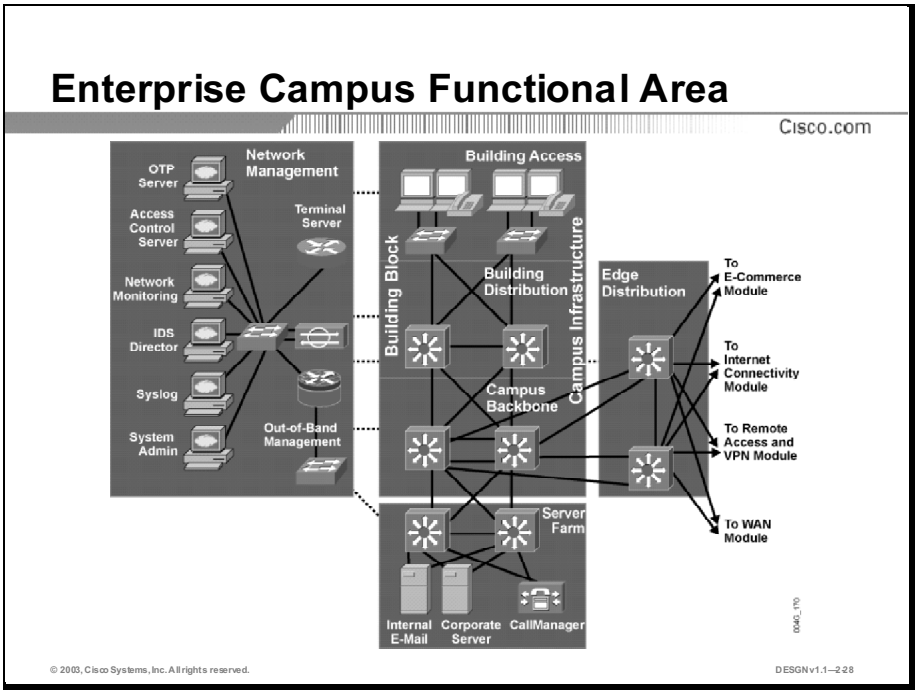

The Campus Infrastructure module includes three submodules:

- The Building Access submodule, located within a campus building, aggregates end users from different workgroups and provides uplinks to the Building Distribution module.
- The Building Distribution submodule aggregates the wiring closets within a building and provides connectivity to the Campus Backbone submodule.
- The Campus Backbone submodule is the core layer of the Campus Infrastructure module. This submodule interconnects the Building Distribution submodules with the Server Farm, Network Management, and Edge Distribution modules. In the Enterprise Campus functional area, the buildings and different parts of the campus connect together across a high-performance, switched backbone called the Campus Backbone submodule.

Network redundancy and high availability are provided at each layer of the Campus Infrastructure module.

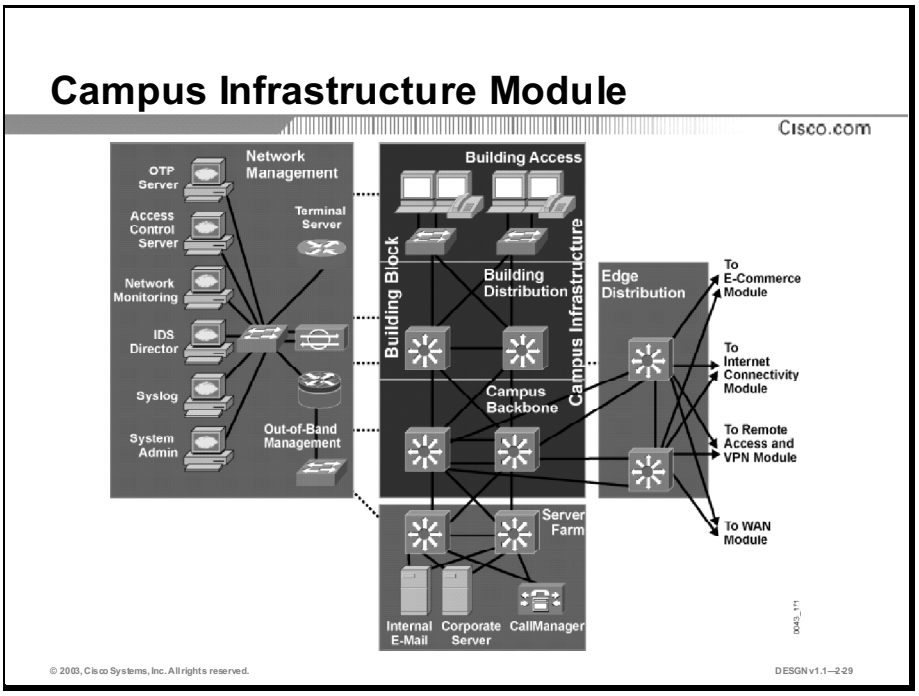

The Campus Infrastructure module interconnects users within a campus with the Server Farm and Edge Distribution modules. This module is composed of submodules (Building Access, Building Distribution, and Campus Backbone). Each building contains a Building Access and a Building Distribution submodule. To scale from a building model to the campus infrastructure, a Campus Backbone submodule between buildings is added. The campus backbone also provides Campus Infrastructure module connectivity to the Edge Distribution and Server Farm modules.

A high-capacity, centralized Server Farm module provides internal server resources to users. The Network Management module supports security, monitoring, logging, troubleshooting, and other common management features from end to end. The Edge Distribution module provides connectivity between the Enterprise Campus and the Enterprise Edge functional areas.

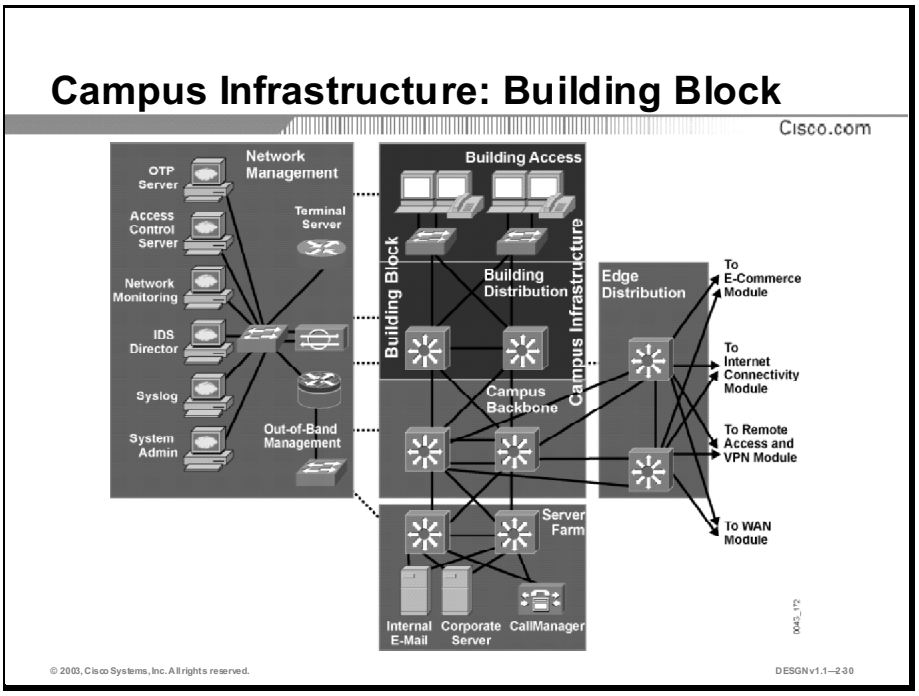

The Campus Infrastructure design consists of a number of buildings (building blocks) connected across a Campus Backbone submodule. Networking in each building relies on the Building Access and Building Distribution submodules. In the most general model, Layer 2 (L2) data link switching is used in the Building Access submodule and multilayer switching in the Building Distribution submodule.

The building block has two submodules:

- **Building Access:** Contains end-user workstations, IP Phones, network printers, and access switches that connect devices to the Building Distribution submodule. The Building Access submodule performs important services such as broadcast suppression, protocol filtering, network access, and QoS marking.
- **Building Distribution:** Provides aggregation of access networks using multilayer switching. Distribution performs routing, QoS, and access control. Requests for data flow into these switches and into the Campus Backbone submodule. Responses follow the identical path in reverse. Redundancy and load balancing with the Building Access and Campus Backbone submodules are recommended.

In the example shown in the figure, each Building Distribution submodule has two equal-cost paths into the Campus Backbone submodule. This provides fast failure recovery, because each distribution switch maintains two equal-cost paths in the routing table to every destination network. When one connection to the Campus Backbone submodule fails, all routes switch over immediately to the remaining path in about one second after the link failure is detected.

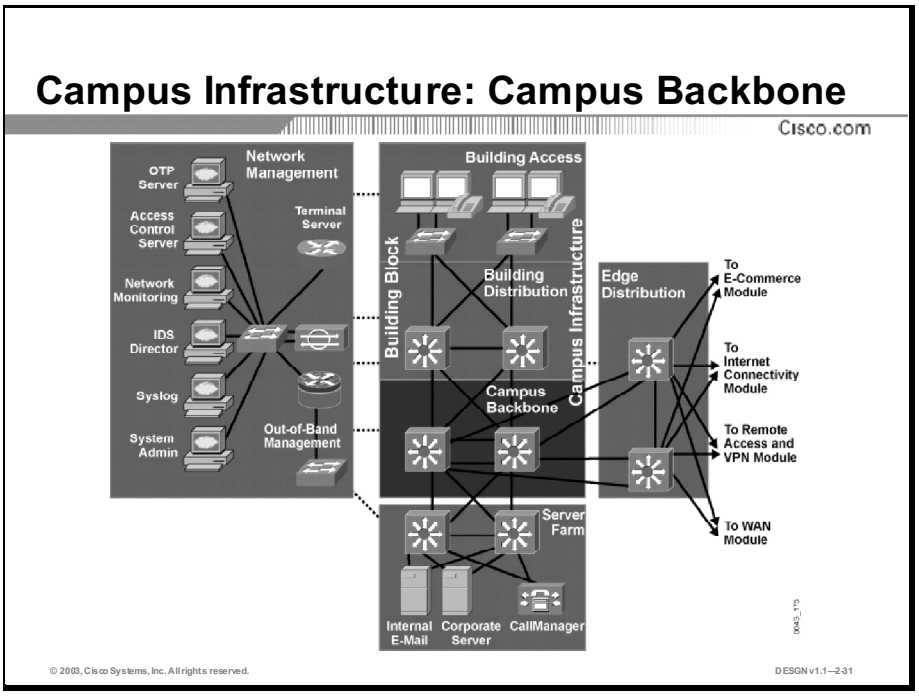

The Campus Backbone submodule provides redundant and fast-converging connectivity between buildings, as well as with the Server Farm and Edge Distribution modules. It routes and switches traffic as fast as possible from one module to another. In general, this module uses multilayer switches for high-throughput functions with added routing, QoS, and security features.

Note: The access, distribution, and core layers can appear in any functional area or module of the Enterprise Composite Network Model.

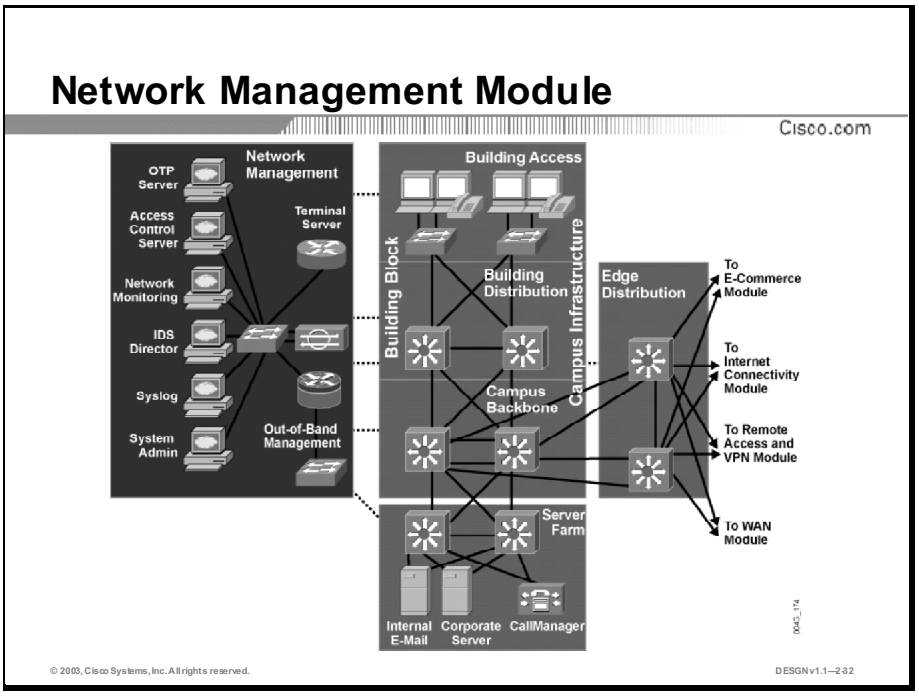

The Network Management module performs intrusion detection, system logging, TACACS+/RADIUS and One Time Password (OTP) authentication, in addition to network monitoring and general configuration management functions. For management purposes, an out-of-band connection (network on which no production traffic travels) to all network components is recommended. For locations where an out-of-band network is impossible, the Network Management module uses the production network. The Network Management module provides configuration management for nearly all devices in the network through the use of two primary technologies:

- IOS routers act as terminal servers and provide a dedicated management network segment. The routers provide a reverse-Telnet function to the console ports on the Cisco devices throughout the enterprise.
- More extensive management features (software changes, content updates, log and alarm aggregation, and Simple Network Management Protocol [SNMP] management) are provided through the dedicated management network segment.

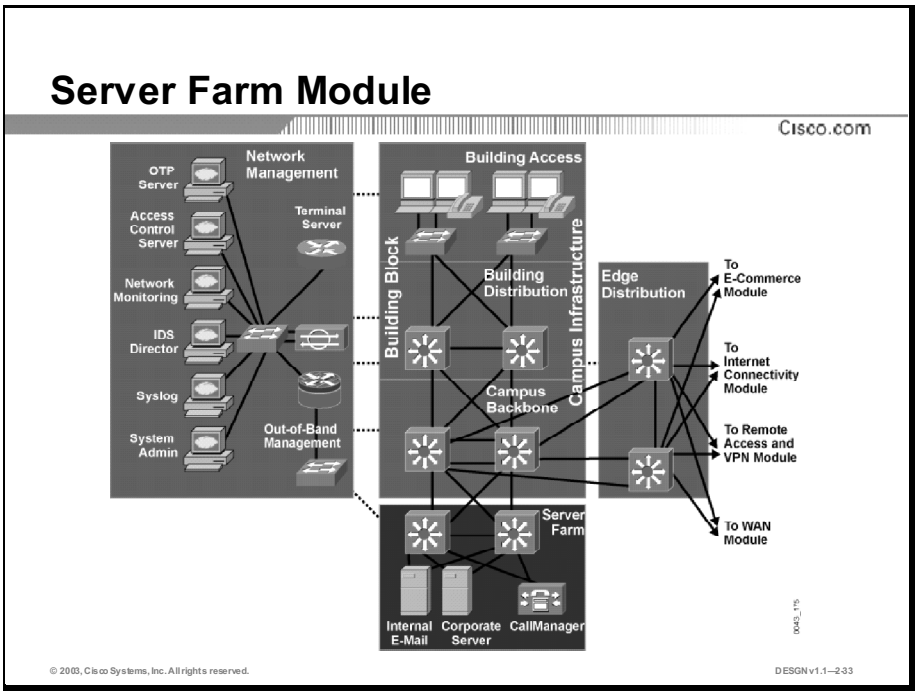

The Server Farm module contains internal e-mail and corporate servers providing application, file, print, e-mail, and Domain Name System (DNS) services to internal users. Because access to these servers is vital, they are connected to two different switches in the figure, enabling full redundancy and load sharing. Moreover, the Server Farm module switches are cross-connected with Campus Backbone switches, enabling high reliability and availability of all servers in the Server Farm module.

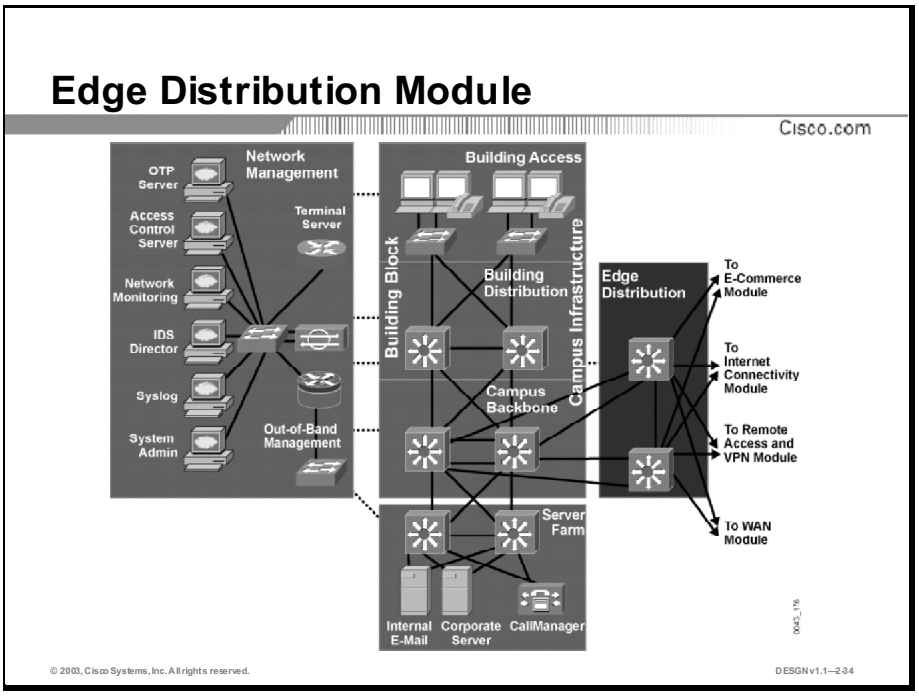

The Edge Distribution module aggregates the connectivity from the various elements at the Enterprise Edge and routes the traffic into the Campus Backbone submodule. Its structure is similar to the Building Distribution submodule. Both modules use access control to filter traffic, although the Edge Distribution module can rely on the entire Enterprise Edge functional area devices to perform additional security functions. Both modules use multilayer switching to achieve high performance, but the Edge Distribution module can add more security functions because the performance requirements are usually not as great.

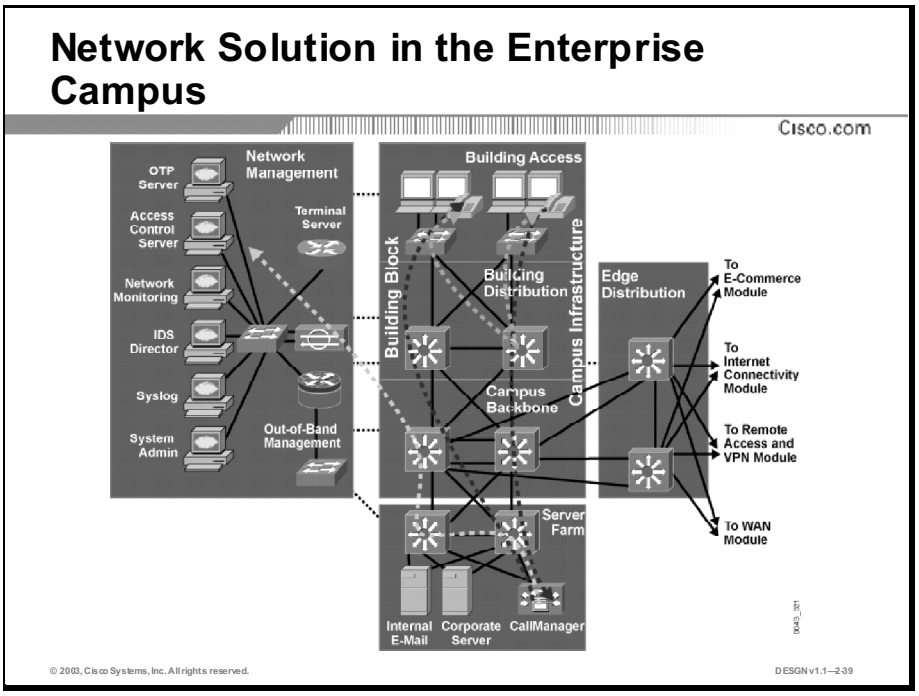

The figure shows how a network solution (IP telephony) uses all the modules of the Enterprise Campus.

Step 1 The control session between the calling IP Phone and the CallManager is established.

Step 2 The CallManager finds the called party and informs both parties of the new call.

Step 3 The voice session is established between the IP Phones.

Step 4 A Call Detail Record (CDR) is added in the CallManager or network accounting system.

Guidelines

Use these guidelines to create an Enterprise Campus network design:

Step 1 Select modules within the campus that act as buildings, with Building Access and Building Distribution submodules.

Step 2 Determine the locations and the number of Building Access switches and their uplinks to Building Distribution switches.

Step 3 Select the appropriate Building Distribution switches, taking into account the number of Building Access switches and end users. Use at least two Building Distribution switches for redundancy.

Step 4 Consider two uplink connections from each Building Access switch to two Building Distribution switches.

Step 5 Determine where servers are or will be located and design the Server Farm module with at least two distribution switches, which connect all servers in a full redundant mode.

Step 6 Design the Network Management module with the out-of-band (reverse Telnet, special Ethernet) connections to all critical devices in the network.

Step 7 Determine the module within the campus that acts as an interface between the Enterprise Campus and Enterprise Edge functional area. This is the Edge Distribution module. Use distribution switches appropriate to the size, volume, and complexity of the network.

Step 8 Implement the Edge Distribution module in a redundant manner.

Step 9 Design the Campus Backbone submodule of the Campus Infrastructure module, consisting of at least two switches and taking into account the expected volume of traffic between modules.

Step 10 Interconnect all modules of the Enterprise Campus with the Campus Backbone submodule of the Campus Infrastructure module in a redundant manner.

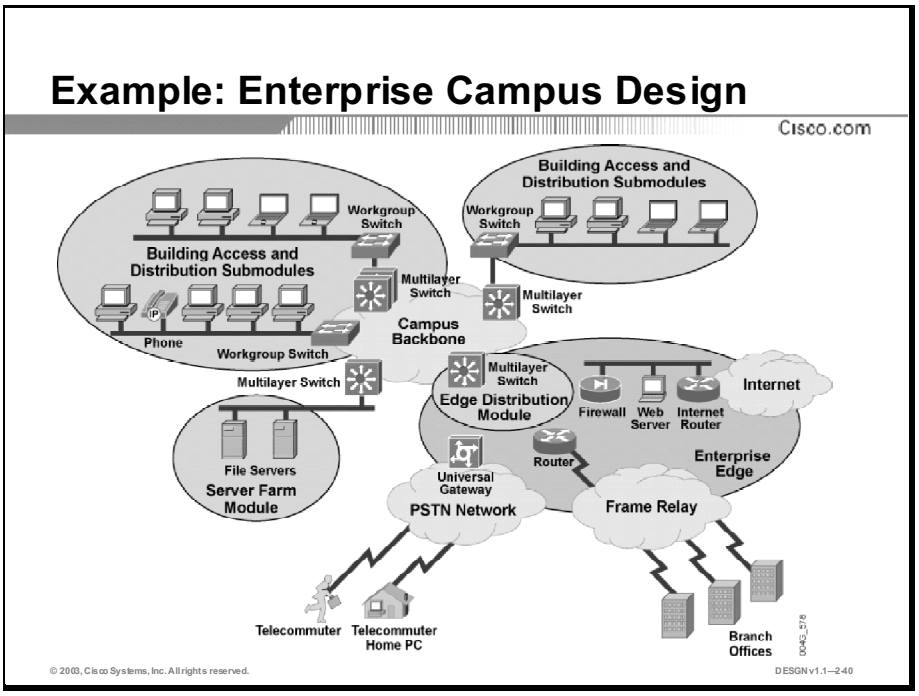

The example shows how the Enterprise Campus is divided into easy manageable building blocks: the Server Farm module, the Edge Distribution module, and the Campus Backbone submodule. Note how the Enterprise Edge is reachable only through the Edge Distribution module.

Enterprise Edge Modules

The Enterprise Edge functional area includes the E-Commerce module, the Internet Connectivity module, the Remote Access and VPN module, and the WAN module. This topic describes the components of the Enterprise Edge and explains the importance of each module.

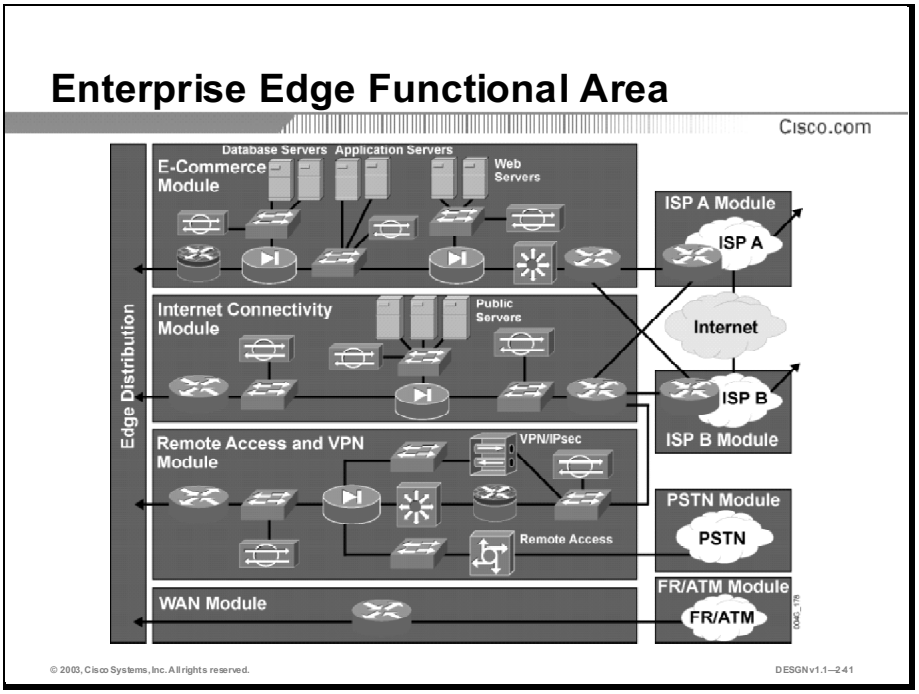

Each module connects to the Edge Distribution module, which provides connectivity between the Enterprise Edge and the Enterprise Campus functional areas. The Enterprise Edge modules use different services and WAN technologies, which are typically available from service providers.

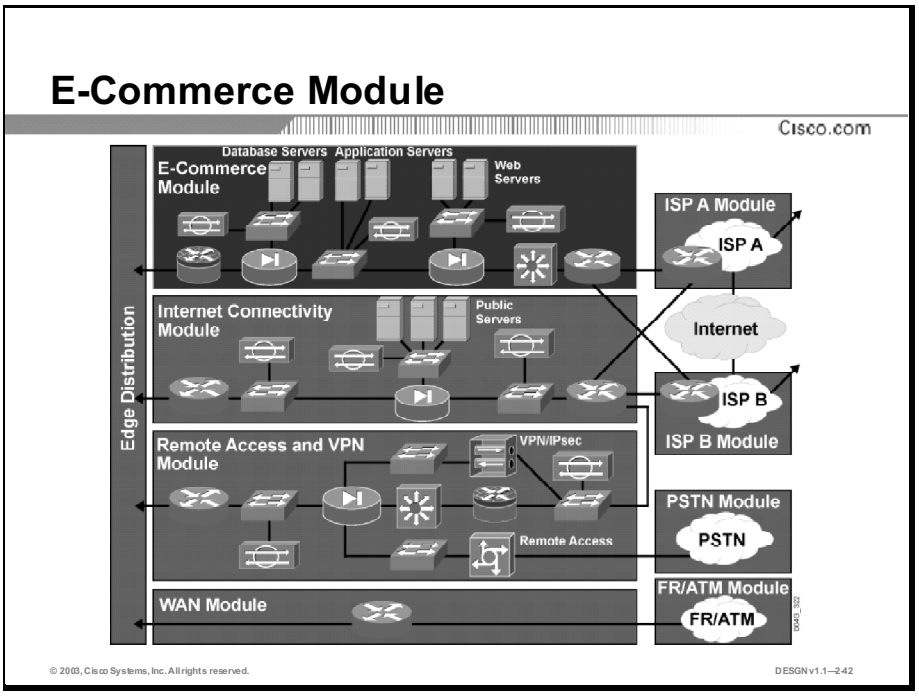

The E-Commerce module enables enterprises to successfully deploy e-commerce applications and take advantage of the opportunities provided by the Internet. The majority of traffic is initiated outside the enterprise network. All e-commerce transactions pass through a series of intelligent services that provide scalability, security, and high availability within the overall e-commerce network design. To build a successful e-commerce solution, enterprises need these network devices:

- **Web servers:** Act as the primary user interface for e-commerce navigation
- **Application servers:** Host the various applications
- **Database servers:** Contain the application and transaction information that is the heart of the e-commerce business implementation
- **Firewall or firewall routers:** Governs communication between and provides security between the various users of the system
- **Network Intrusion Detection System (NIDS) appliances:** Provide monitoring of key network segments in the module to detect and respond to attacks against the network
- **Multilayer switch with Intrusion Detection System (IDS) modules:** Provide traffic transport and integrated security monitoring

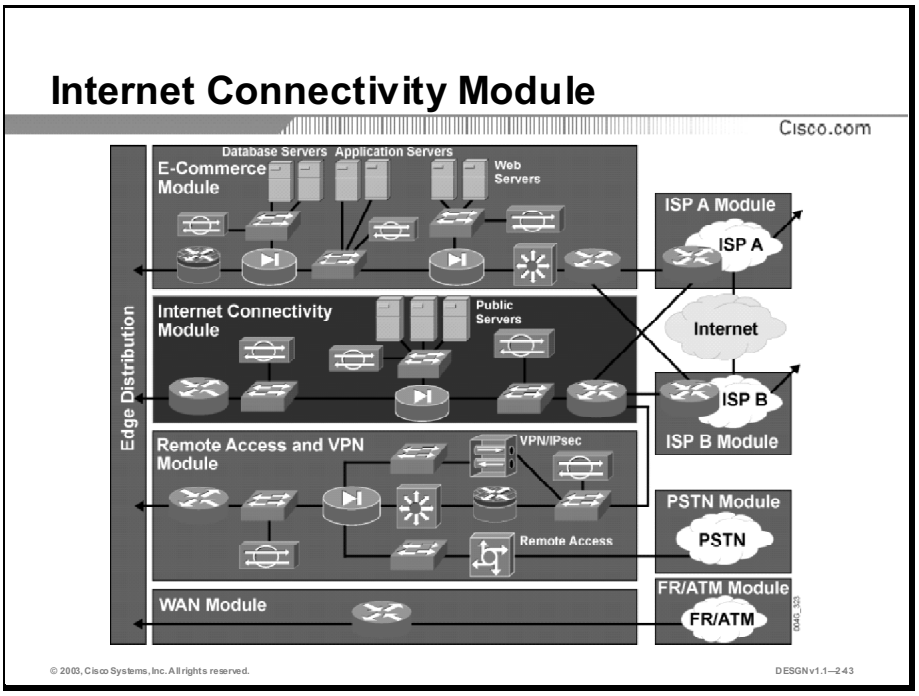

The Internet Connectivity module provides internal users with connectivity to Internet services so that Internet users can access the information on the public servers, including HTTP, FTP, Simple Mail Transfer Protocol (SMTP), and DNS. In addition, this module accepts Virtual Private Network (VPN) traffic from remote users and remote sites and forwards it to the Remote Access and VPN module, where VPN termination takes place. The Internet Connectivity module does not serve e-commerce applications. Major components of the Internet Connectivity module are:

- **SMTP mail servers:** Act as a relay between the Internet and the intranet mail servers.
- **DNS servers:** Serve as the authoritative external DNS server for the enterprise; relay internal requests to the Internet.
- **Public servers (FTP and HTTP):** Provide public information about the organization. Each server on the public services segment has host intrusion detection software to monitor against any rogue activity at the operating system level, in addition to activity in common server applications, including HTTP, FTP, and SMTP.
- **Firewall or firewall routers:** Provides network-level protection of resources, stateful filtering of traffic, and VPN termination for remote sites and users.
- **L2 switches:** Ensure that data from managed devices can cross directly to the IOS firewall only.
- **Edge routers:** Provide basic filtering and multilayer connectivity to the Internet.

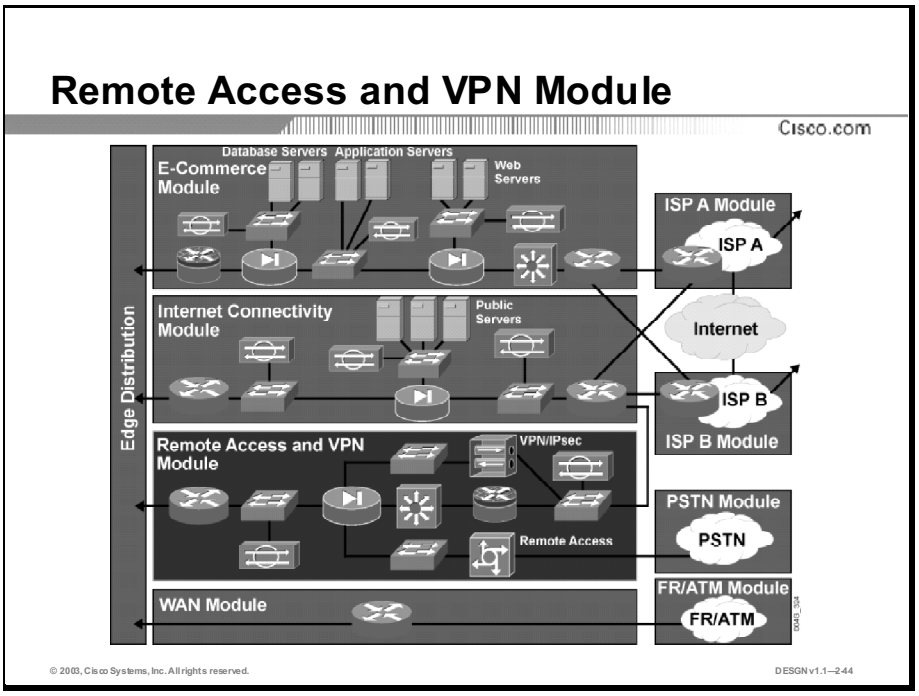

The Remote Access and VPN module terminates VPN traffic, forwarded by the Internet Connectivity module, from remote users and remote sites. It also initiates VPN connections to remote sites using the Internet Connectivity module. Furthermore, the module terminates dial-in connections received through the Public Switched Telephone Network (PSTN) and, after successful authentication, grants dial-in users access to the network. Major components of the Remote Access and VPN module are:

- **Dial-in access concentrators:** Terminate dial-in connections and authenticate individual users
- **VPN concentrators:** Terminate IP Security (IPSec) tunnels and authenticate individual remote users
- **Firewalls:** Provide network-level protection of resources and stateful filtering of traffic; provide differentiated security for remote-access users; authenticate trusted remote sites and provide connectivity using IPSec tunnels
- **L2 switches:** Provide L2 connectivity for devices
- **NIDS appliances:** Provide Layer 4 to Layer 7 monitoring of key network segments in the module

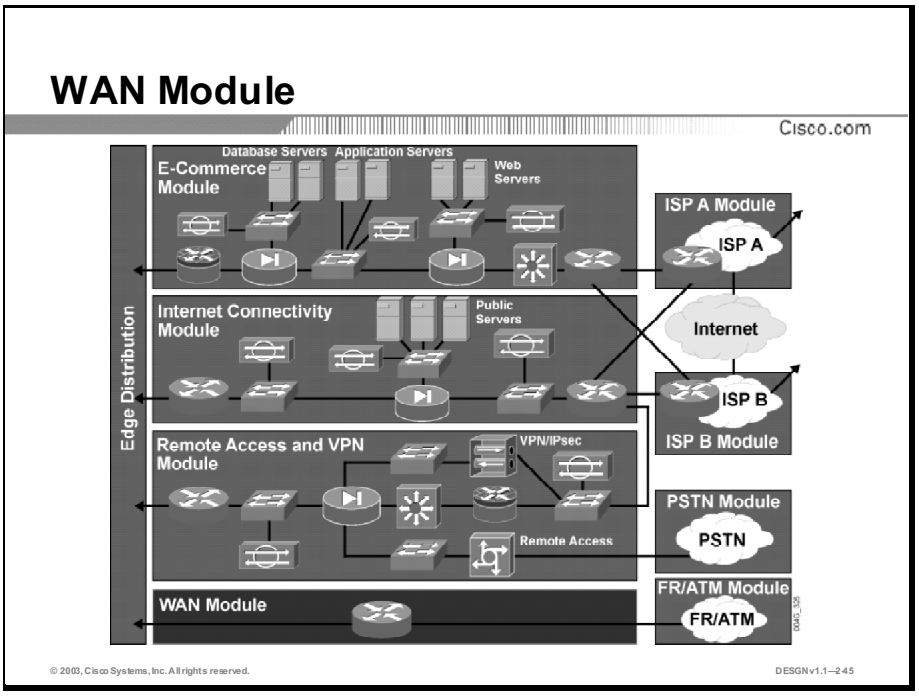

The WAN module uses different WAN technologies for routing traffic between remote sites and the central site. In addition to traditional media (leased lines) and circuit switched data link technologies (Frame Relay and ATM), the WAN module can use more recent WAN physical layer technologies, including SONET/Synchronous Digital Hierarchy (SDH), cable, digital subscriber line (DSL), and wireless. All Cisco devices supporting these WAN technologies, in addition to routing, access control, and QoS mechanisms, can be used in this module. While security is not as critical when all links are enterprise-owned, you should consider security in the network design.

Guidelines

Use these guidelines to create the Enterprise Edge functional area:

Step 1 Determine which part of the edge is used exclusively for permanent connections to remote locations (branch offices) and assign it to the WAN module. All WAN devices supporting Frame Relay, ATM, cable, leased lines, SONET/SDH, and so on, are located here.

Step 2 Determine the connections from the corporate network into the Internet and assign them to the Internet Connectivity module. The Internet Connectivity module should have security to prevent any unauthorized access from the Internet to the internal network. The public web servers are in this module or the E-Commerce module.

Step 3 Design the Remote Access and VPN module if the enterprise requires VPN connections or dial-in for accessing the internal network from the outside world. Implement a security policy. Users should not be able to access the internal network directly without authentication and authorization. The VPN sessions use the connectivity from the Internet Connectivity module.

Step 4 Create the E-Commerce module (for business-to-business or business-to-customer scenarios) when customers or partners require Internet access to business

applications and database servers. Deploy a high-security policy that allows customers access to predefined servers and services and restricts all other operations.

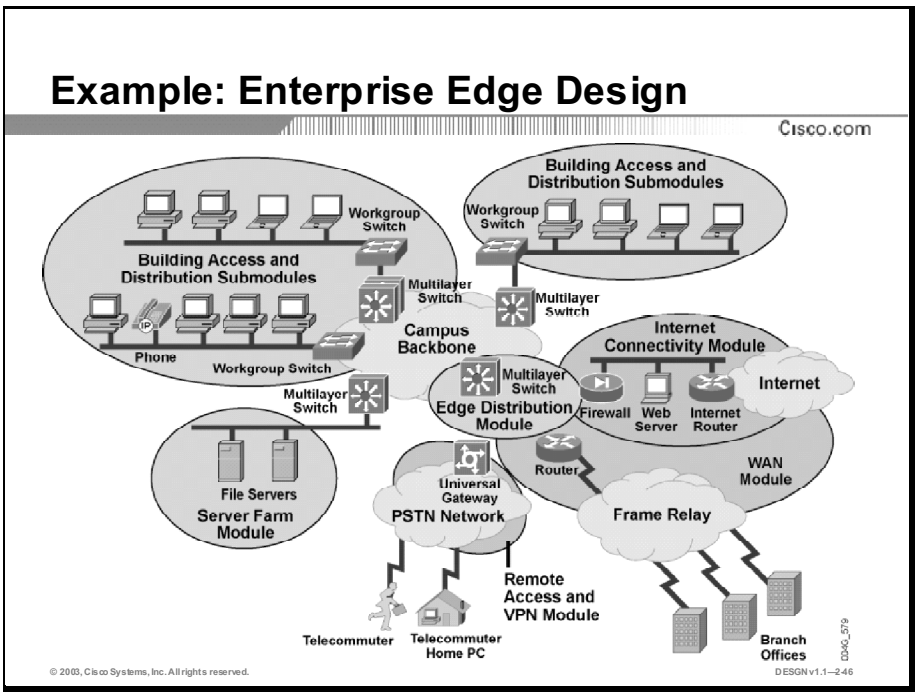

The figure shows how the Enterprise Edge functional area is divided into easily manageable modules. The Enterprise Edge functional area may include the E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules. The modularity and independence among modules is noticeable.

Service Provider Edge Modules

The Service Provider Edge functional area is located on a service provider network, and enables communications with the outside world. This topic describes the Service Provider Edge functional area and its component modules.

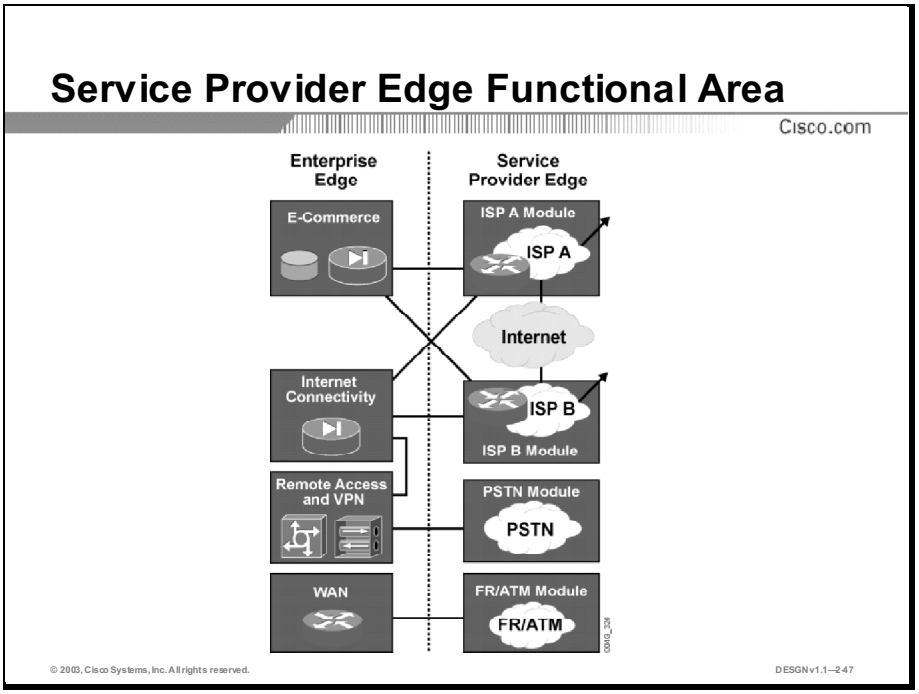

The last functional area in the Enterprise Composite Network Model is the Service Provider Edge. The modules in this area are not implemented by the enterprise itself, but they are necessary to enable communication with other networks and most often use different WAN technologies and ISPs.

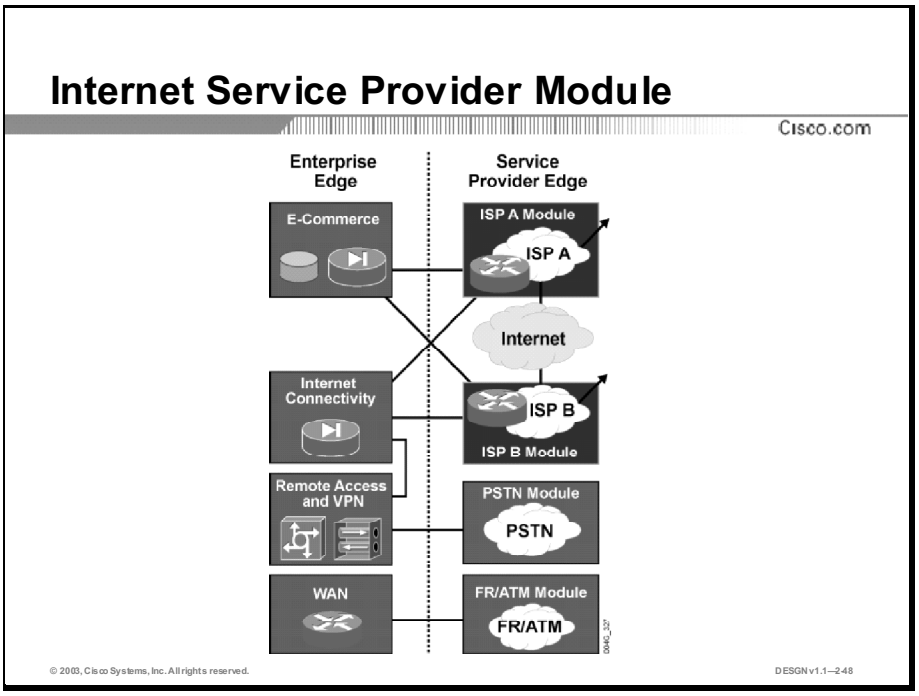

The Internet Service Provider module enables enterprise IP connectivity to the Internet. This service is essential for enabling Enterprise Edge services, such as the E-Commerce, Internet Connectivity, and Remote Access and VPN modules. To provide redundant connections to the Internet, enterprises connect to two or more ISPs. Physical connection between the ISP and the enterprise can come from any of the WAN technologies.

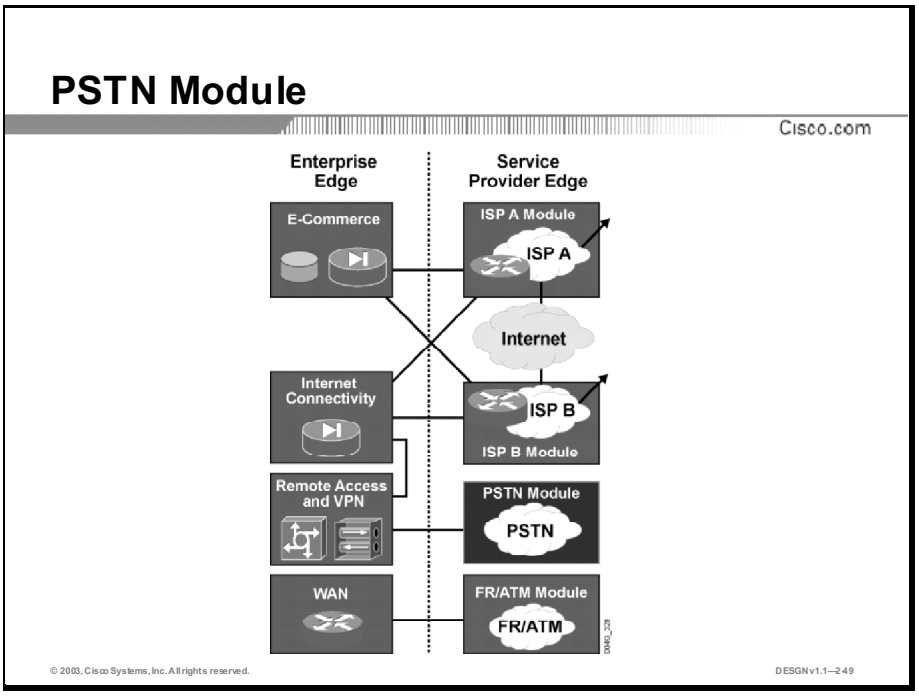

The PSTN module represents the dial-up infrastructure for accessing the enterprise network using ISDN, analog, and wireless telephony (cellular) technologies. Enterprises can also use it to back up existing WAN links. WAN backup connections are generally established on demand and are disconnected after idle timeout.

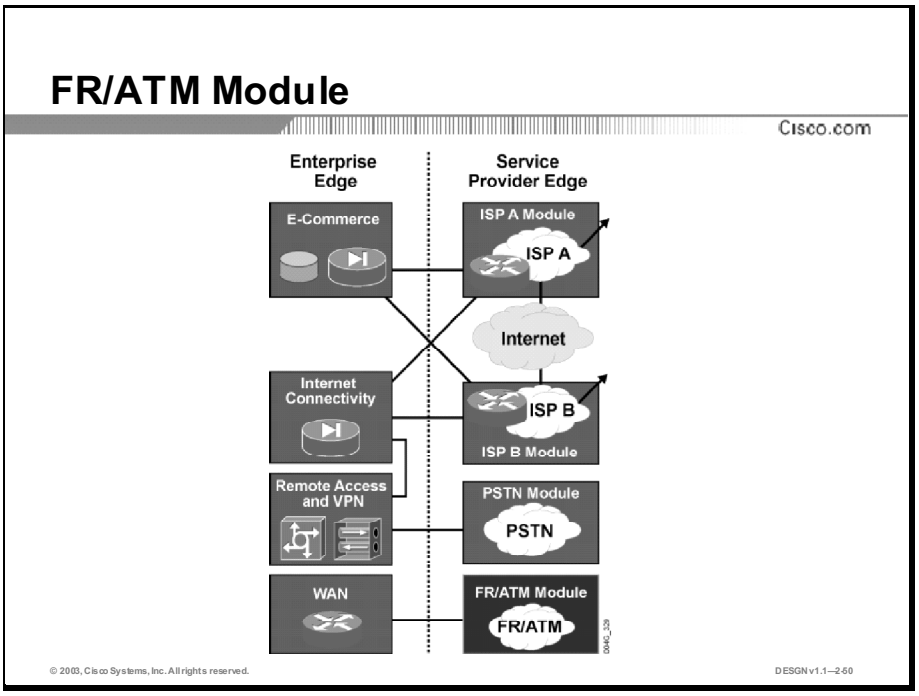

The FR/ATM module includes all WAN technologies for permanent connectivity with remote locations. Traditional Frame Relay and ATM are still the most frequently used; however, many modern technologies can fit into the same module:

- Frame Relay is a connection-oriented, packet-switching technology designed to transmit data traffic efficiently at data rates of up to those provided by E3 and T3 connections. Its ability to connect multiple remote sites across a single physical connection reduces the number of point-to-point physical connections required to link sites together.
- ATM is an alternative to Frame Relay and supports higher speeds. It is a high-performance, cell-oriented switching and multiplexing technology for carrying different types of traffic.
- Leased lines provide the simplest permanent point-to-point connection between two remote locations. The carrier company reserves point-to-point links for the company's private use. Because the connection is exclusive, the carrier (service provider) can assure a given level of quality. The fee for the connection is a fixed monthly rate.
- SONET and SDH are standards for transmission over optical networks. Europe uses SDH; its equivalent, SONET, is used in North America.
- Cable technology uses existing coaxial cable television (CATV) cables. This technology, coupled with cable modems, provides much greater bandwidth than telephone lines and can provide extremely fast access to the Internet or enterprise network.
- DSL is a modem technology that uses existing twisted-pair telephone lines to transport high-bandwidth data, such as voice, data, and video. DSL is sometimes referred to as "last-mile" technology because it is used only for connections from a telephone switching station (service provider) to a home or office, not between switching stations. DSL is used mostly by telecommuters to access enterprise networks. However, many more companies are migrating from traditional Frame Relay to DSL technology using VPNs because of its cost-efficiency.
- Wireless technology is another modern technology for interconnecting remote LANs. The point-to-point signal transmissions take place through the air over a terrestrial radio or microwave platform rather than through copper or fiber cables. Fixed wireless does not

require satellite feeds or local phone service. An advantage of fixed wireless is its ability to connect with users in remote areas without the need for laying down new cables. However, this technology is limited to shorter distances and may be degraded by weather conditions.

Guidelines

Use these guidelines for the proper use of Service Provider Edge modules:

- When connecting to the Internet, consider redundant connection to two service providers or two different connections to a single service provider.
- For mobile and remote users that require direct dial-up access, choose ISDN, analog modem, or cellular wireless. These technologies provide speeds ranging from 9600 bps to 128 kbps.
- For higher speeds (greater than or equal to T3/E3 rates), use ATM over SONET/SDH links. ATM may be available at rates down to T1/E1. It may also be possible to use Packet over SONET (POS) when an enterprise can acquire access to “dark fiber.”
- When connecting to remote locations via the WAN module with speeds of less than or equal to T3/E3 rates, take into consideration leased-line and Frame Relay connections. Leased lines are typically cheaper for shorter distances (a few miles or kilometers); however, with greater distance, Frame Relay becomes competitive. (In the United States, Frame Relay is typically less expensive, regardless of circuit distance.) Not all service providers offer Frame Relay above T1/E1 rates.
- Investigate if DSL is a possible solution for a concentration of remote locations because DSL brings considerable savings compared to the traditional WAN links.
- Use cable connections only for telecommuters since cable does not provide any throughput guarantee.

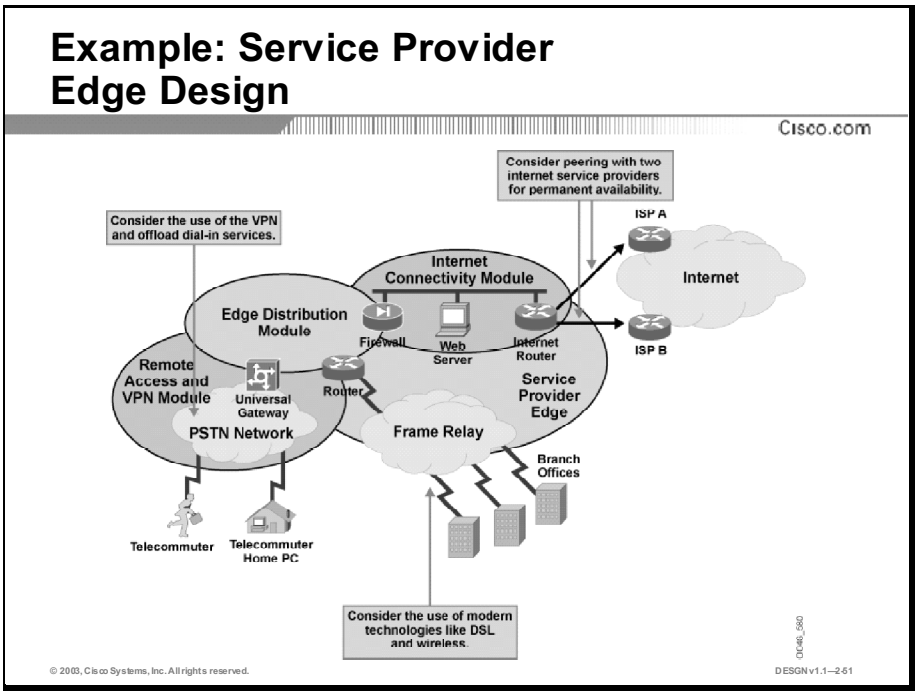

The example shows that enterprises can use different WAN technologies to provide connectivity to a campus network. Enterprises use PSTN for telecommuters and Frame Relay is for concentration of remote locations. Internet access is accomplished with redundant connections to two ISPs.

Summary

This topic summarizes the key points discussed in this lesson.

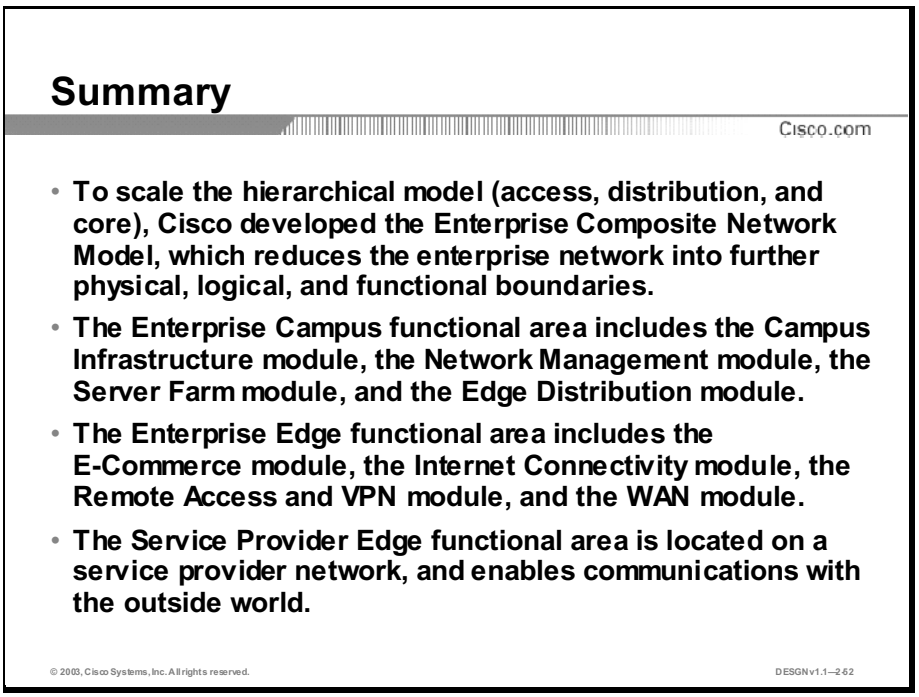

References

For additional information, refer to these resources:

- *SAFE White Paper: A Security Blueprint for Enterprise Networks,* http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm
- *SAFE: Extending the Security Blueprint to Small, Midsize, and Remote-User Networks,* http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safes_wp.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three modules are parts of the Enterprise Campus functional area? (Choose three.)

A) Network Management module

B) Edge Distribution module

C) Internet Connectivity module

D) E-Commerce module

E) Server Farm module

Q2) Which module or submodule in the Enterprise Campus functional area interconnects the building block with the Server Farm module and Edge Distribution module?

A) Network Management module

B) Server Distribution module

C) Building Distribution submodule

D) Enterprise Edge functional area

E) Enterprise Distribution module

F) Campus Backbone submodule

Q3) Which type of server is typically located in the Internet Connectivity module?

A) Internet

B) public

C) private

D) corporate

Q4) The E-Commerce module uses three types of servers—_____, _____ and, _____—for web communication with users who are running applications and storing data. (Choose three.)

A) database

B) web

C) private

D) application

E) public

F) Internet

Q5) Which two modules are connected to the Remote Access and VPN module? (Choose two.)

A) Service Provider module

B) PSTN module

C) Server Farm module

D) WAN module

E) Internet Connectivity module

Quiz Answer Key

Q1) A, B, E

Relates to: Enterprise Composite Network Model

Q2) F

Relates to: Enterprise Campus Modules

Q3) B

Relates to: Enterprise Edge Modules

Q4) A, B, D

Relates to: Enterprise Edge Modules

Q5) B, E

Relates to: Service Provider Edge Modules

Evaluating Network Services and Solutions within Modular Networks

Overview

Businesses operating large enterprise networks seek an enterprise-wide infrastructure to provide a solid foundation for emerging application solutions such as IP telephony, content delivery, and storage networking.

In this lesson, the intelligent network services (for example, security and high availability) and network solutions such as voice transport or Content Networking within and between modules are presented with respect to the modules forming the Enterprise Composite Network Model. The main focus of the lesson is to explain the relationship between network modules and to describe how concentration on the functions of each module simplifies network design and deployment.

Relevance

Intelligent network services and network solutions are key technologies that you will implement on the network.

Objectives

Upon completing this lesson, you will be able to evaluate network services and solutions within modular networks. This includes being able to meet these objectives:

- Explain the role of intelligent network services in an enterprise network design
- Describe the security design issues within an enterprise network
- Describe the high-availability considerations within an enterprise network
- Explain the role of network solutions in an enterprise network design
- Describe the IP telephony considerations within an enterprise network
- Describe the Content Networking functions within an enterprise network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in IOS software
- Familiarity with network modularity concepts, particularly with the Enterprise Composite Network Model

Outline

The outline lists the topics included in this lesson.

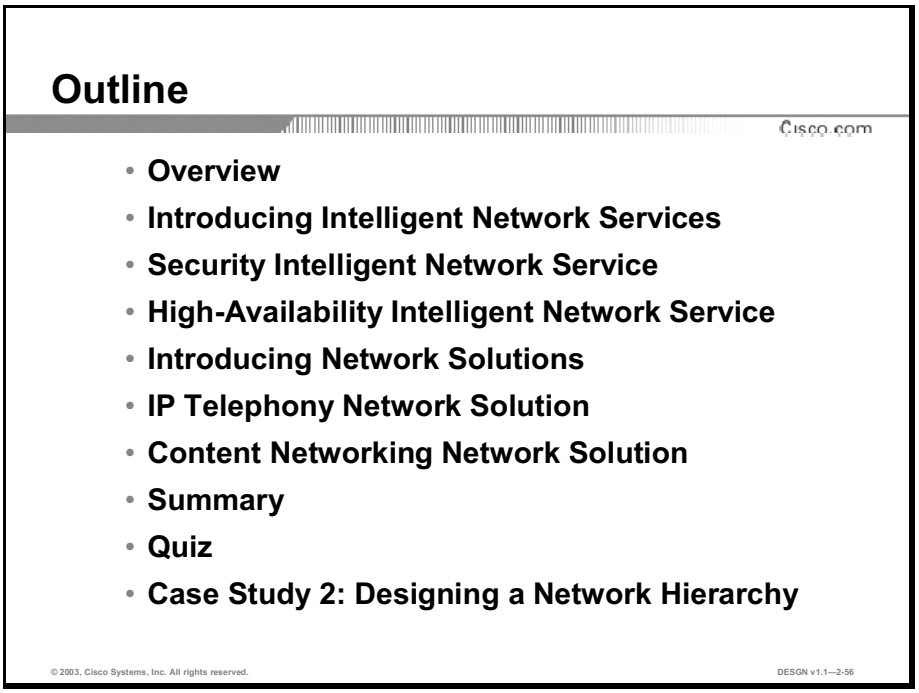

Introducing Intelligent Network Services

Beyond just moving a datagram between two points, intelligent network services essentially add intelligence to the network infrastructure. Intelligent network services support application awareness within the network. This topic describes intelligent network services.

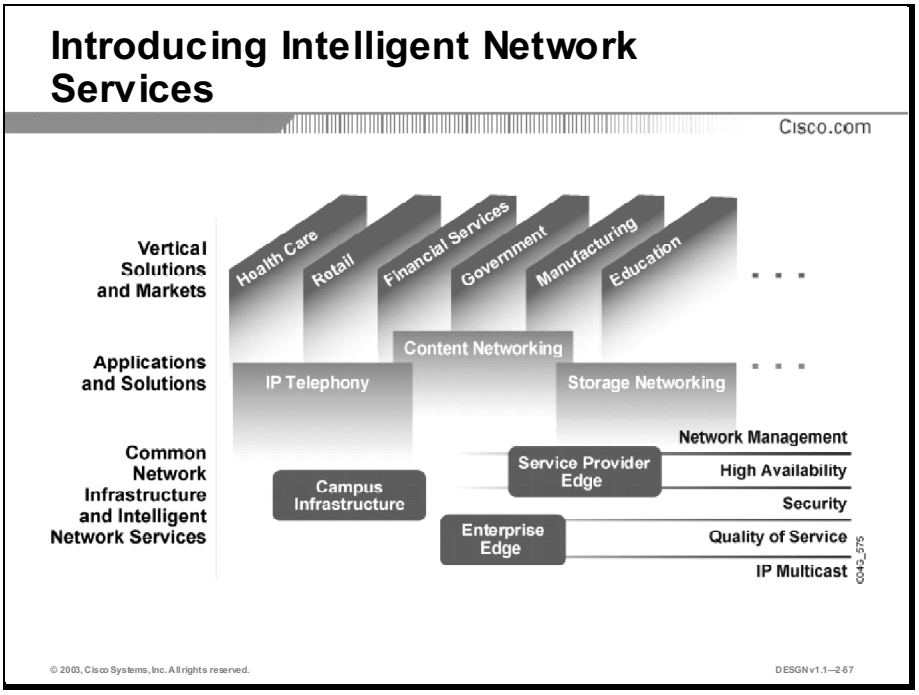

In IP, the forwarding service assumed that end nodes in the network were intelligent and that the network core was “dumb.” With advances in networking software and hardware, the network is able to offer increasingly rich, intelligent mechanisms for forwarding information.

Through intelligent network classification, the network can identify traffic based on application content and context. Advanced network services act on classified traffic to regulate performance, ensure security, facilitate delivery, and improve manageability.

An intelligent network service is a supporting, but not an ultimate solution. For example, implementing QoS is not an ultimate solution within the network. It is necessary to enable other solutions and applications such as IP telephony. Therefore, QoS is not a solution; it is an intelligent network service. However, voice communication is an ultimate goal; therefore, it is a network solution.

Here are examples of intelligent network services:

- **Network management:** Includes LAN management for advanced management of multilayer switches; routed WAN management for monitoring, traffic management, and access control to administer the routed infrastructure of multiservice networks; service management for managing and monitoring service level agreements; and VPN/security management for optimizing VPN performance and security administration.
- **Security:** Ensures the security of the network and applications through authentication, encryption, and failover. Security features include stateful, application-based filtering, defense against network attacks, per-user authentication and authorization, and real-time alerts.

- **High availability:** Ensures end-to-end availability for services, clients, and sessions. Implementation includes reliable, fault-tolerant network devices to automatically identify and overcome failures, and resilient network technologies.
- **QoS:** Manages the delay, delay variation (jitter), bandwidth availability, and packet loss parameters on a network to meet the diverse needs of voice, video, and data applications. QoS features provide value-added functionality such as network-based application recognition (NBAR) for classifying traffic on an applications basis, a Service Assurance Agent (SAA) for end-to-end QoS measurements, Resource Reservation Protocol (RSVP) signaling for admission control and reservation of resources, and a variety of configurable queue insertion and servicing disciplines.
- **IP multicasting:** Provides bandwidth-conserving technology that reduces network traffic by delivering a single stream of information intended for many corporate recipients and homes through the transport network. Multicasting enables distribution of videoconferencing, corporate communications, distance learning, distribution of software, and other applications. Multicast packets are replicated only as necessary in the network by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols resulting in the most efficient delivery of data to multiple receivers.

Note: In this course, two examples of intelligent network services (security and high availability) are introduced giving an overview of how to implement a sample intelligent network service on a network infrastructure.

To support network solutions efficiently, you should deploy the underlying intelligent network services on a module-by-module basis, as required by a network solution. You can replicate these design elements to other modules of the enterprise network as the network changes. Thus, modularization to small subsets of the overall network simplifies the network design and often reduces the cost and complexity of the network.

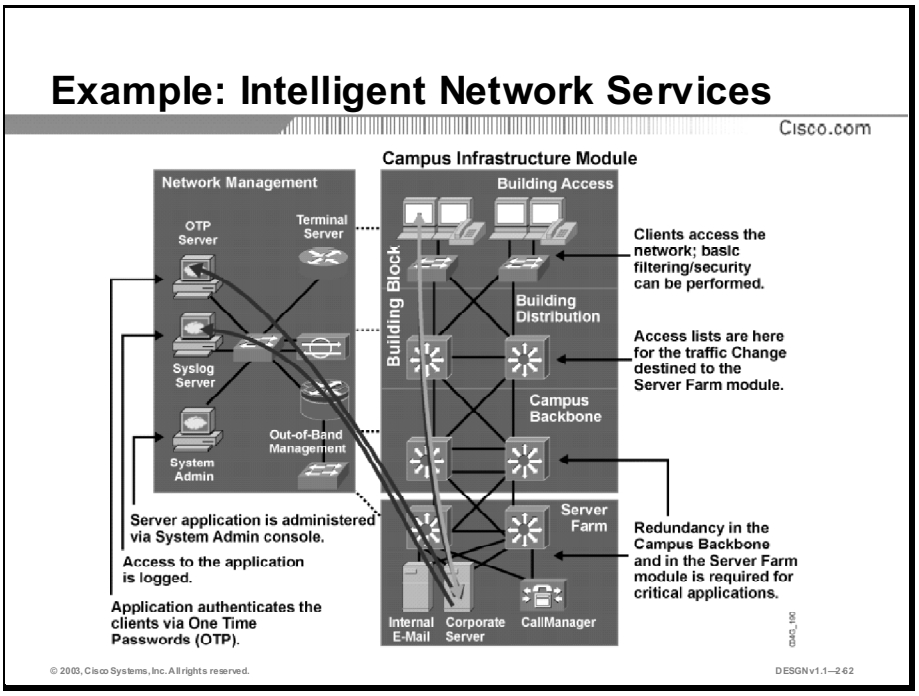

The figure shows a network solution that utilizes all modules of the Enterprise Campus functional area and requires different intelligent network services for reliable and secure data transfer.

A client-server application might have the following requirements within an enterprise campus network:

- At the Building Access submodule, the client application is granted controlled access to the network.
- At the Building Distribution submodule, the packet filtering ensures that the application data is forwarded to the Campus Backbone submodule to reach the Server Farm module.
- The Campus Infrastructure module provides duplicate network links with fast convergence.
- In the Server Farm module, the application requires highly redundant connections to file servers.
- The Network Management module grants access to the Server Farm module (using One Time Passwords (OTPs) and monitors and logs successful and unsuccessful access to the application.

Security Intelligent Network Service

Security is an intelligent network service that increases the integrity of the network by protecting network resources and users from internal and external threats. Without a full understanding of the threats involved, network security deployments tend to be incorrectly configured, too focused on security devices, or lacking appropriate threat response options. This topic describes security as an intelligent network service.

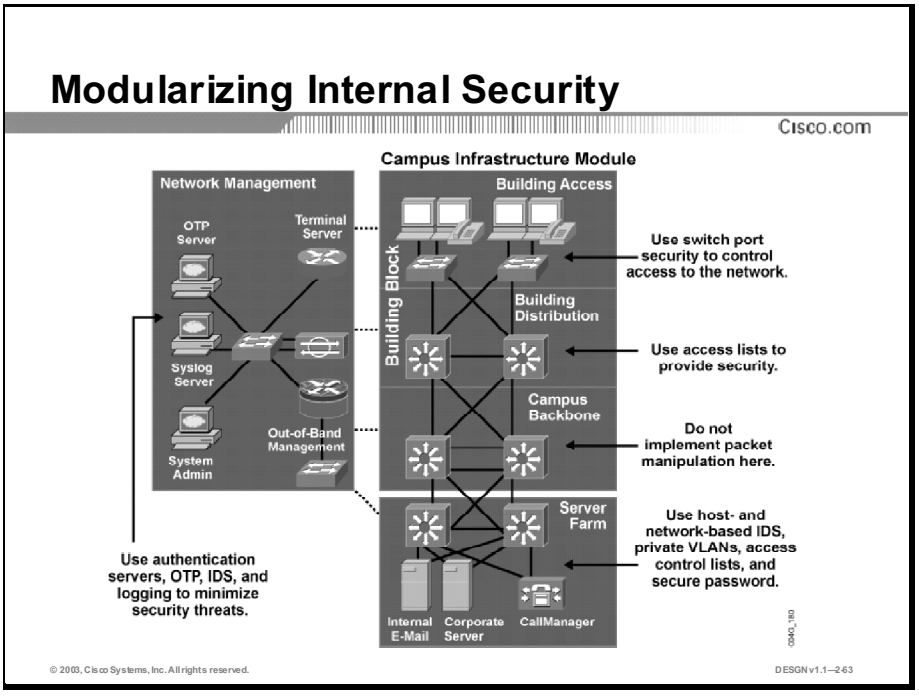

You can evaluate and apply security on a module-by-module basis within the Enterprise Composite Network Model. Some security considerations for each module are:

- The Campus Backbone submodule in the Campus Infrastructure module switches packets as quickly as possible. It should not perform any security functions, because these would slow down packet switching.
- The Building Distribution submodule performs packet filtering, to keep unnecessary traffic from the Campus Backbone submodule. Packet filtering at the Building Distribution submodule is a security function because it prevents some undesired access to other modules. Given that switches in this submodule are usually Layer 3-aware multilayer switches, the Building Distribution submodule is often the first location that can filter based on network layer information.
- At the Building Access submodule, access is controlled at the port level with respect to the data link layer information (for example, MAC addresses).
- The Server Farm module provides application services to end users and devices. Given the high degree of access that most employees have to these servers, they often become the primary target of internally originated attacks. Use host- and network-based IDSs, private VLANs, and access control to provide a much more comprehensive response to attacks. Onboard IDSs within multilayer switches can inspect traffic flows on the Server Farm module.

- The Network Management module securely manages all devices and hosts within the enterprise architecture. Syslog provides important information regarding security violations and configuration changes by logging security-related events (authentication, and so on). An authentication, authorization, and accounting (AAA) security server can work in combination with the OTP server to provide a very high level of security to all local and remote users. AAA and OTP authentication reduces the likelihood of a successful password attack.

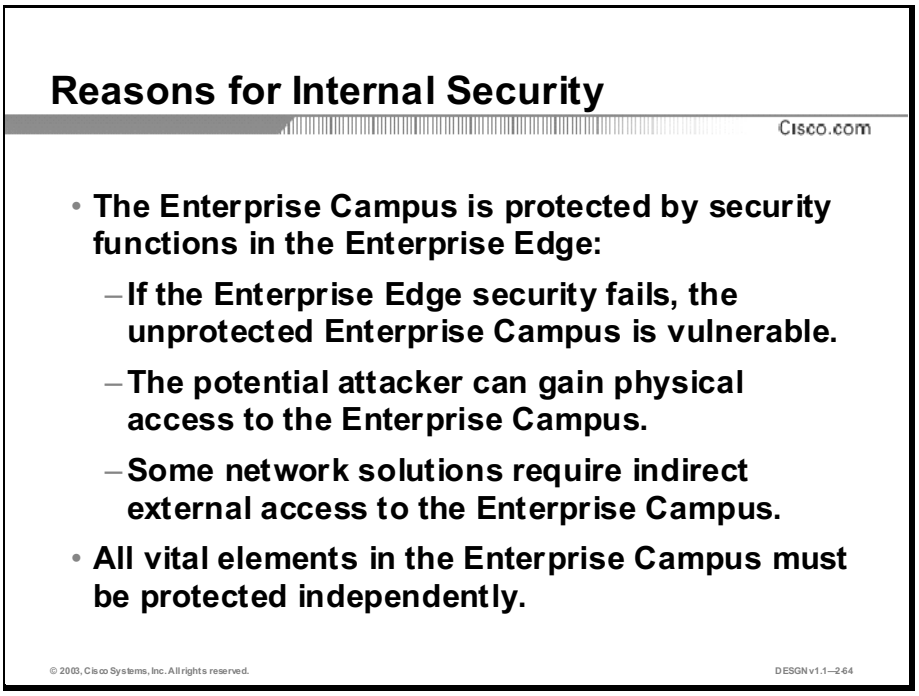

Several reasons exist for strong protection of the internal Enterprise Campus functional area, including security functions in each individual element of the Enterprise Campus:

- Relying on the security established at the Enterprise Edge fails as soon as security in the Enterprise Edge is compromised. Several layers of security increase the protection of the Enterprise Campus, where, usually, the most strategic assets reside.
- The potential attacker can gain physical access to devices in the Enterprise Campus. Relying on physical security is not enough.
- Very often external access does not stop at the Enterprise Edge. Applications require at least an indirect access to the Enterprise Campus resources, requiring strong security.

Designing Security

Security solutions must be designed in a layered and independent way:

- Establish several layers of protection.
- Make sure that security functions at one layer or in one network module do not rely on the security function in other layers or modules.

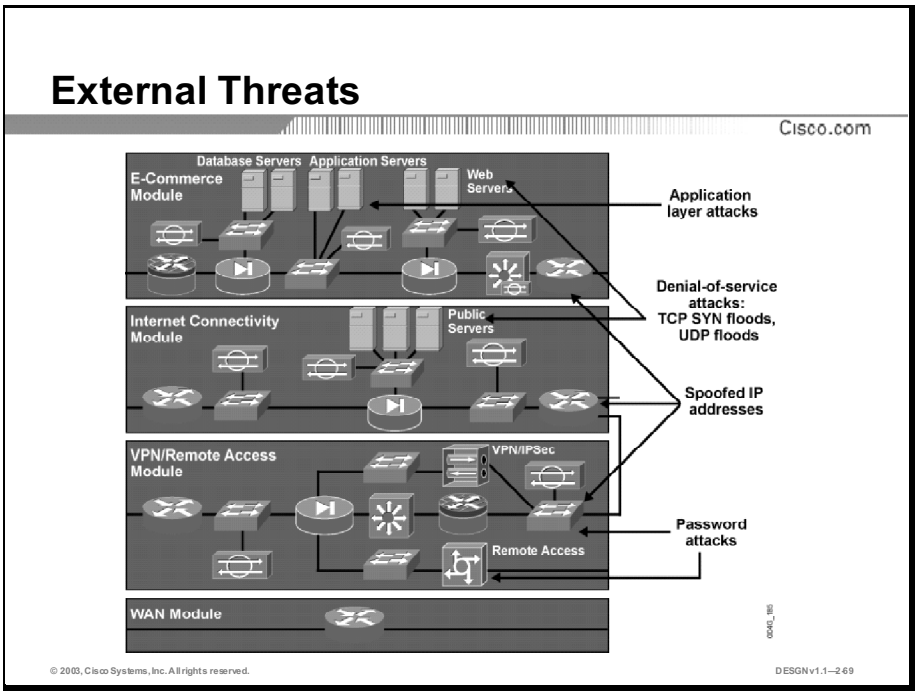

When designing security in an enterprise network, you will identify the Enterprise Edge as the first front to stop potential attacks from the outside. The Enterprise Edge is like a wall with small doors and strong “guards” that efficiently control any access.

Here are four methods of attack on an enterprise network from the outside:

- IP spoofing
- Password attacks
- Denial of service attacks
- Application layer attacks

Because of the complexity of network applications, access control must be extremely granular and flexible, yet still provide strong security. You should balance the ease of use of network applications and resources against the security measures imposed on the network users.

High-Availability Intelligent Network Service

Enterprise networks carry mission-critical information. Organizations need to protect the integrity of this information with internetworking platforms that offer a sufficient level of resilience and high availability. This topic describes high availability as an intelligent network service

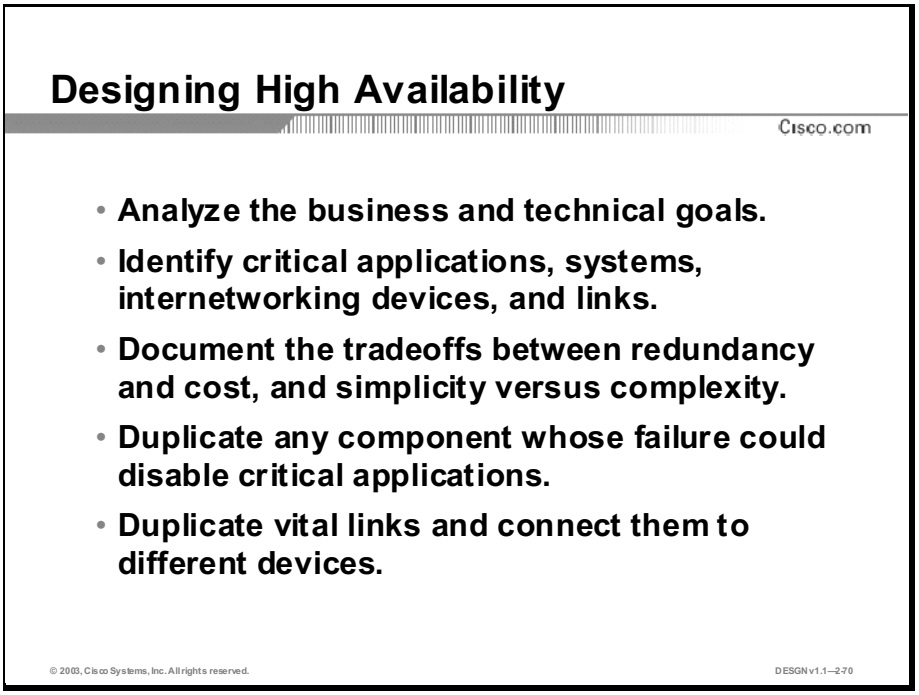

Redundant network designs meet requirements for network high availability by duplicating network links and interconnectivity devices. Redundancy eliminates the possibility of having a single point of failure on the network. The goal is to duplicate required components whose failure could disable critical applications.

Because redundancy is expensive to deploy and maintain, you should implement redundant topologies with care. Be sure to select a level of redundancy that matches the requirements for availability and affordability.

Before you select redundant design solutions, first analyze the business and technical goals to establish the required availability. Make sure that you can identify critical applications, systems, internetworking devices, and links. Analyze the tolerance for risk and the consequences of not implementing redundancy. Discuss the tradeoffs of redundancy versus cost and simplicity versus complexity. Redundancy adds complexity to the network topology and to network addressing and routing.

These redundancies exist in networking:

- Device redundancy, including card and port redundancy
- Redundant physical connections to workstations and servers
- Route redundancy
- Link redundancy

Note: High availability is not ensured end-to-end simply by making the backbone redundant. If communication on a local segment is disrupted for any reason, that information will not reach the backbone. In other words, high availability from end to end is only possible when redundancy is deployed throughout the network.

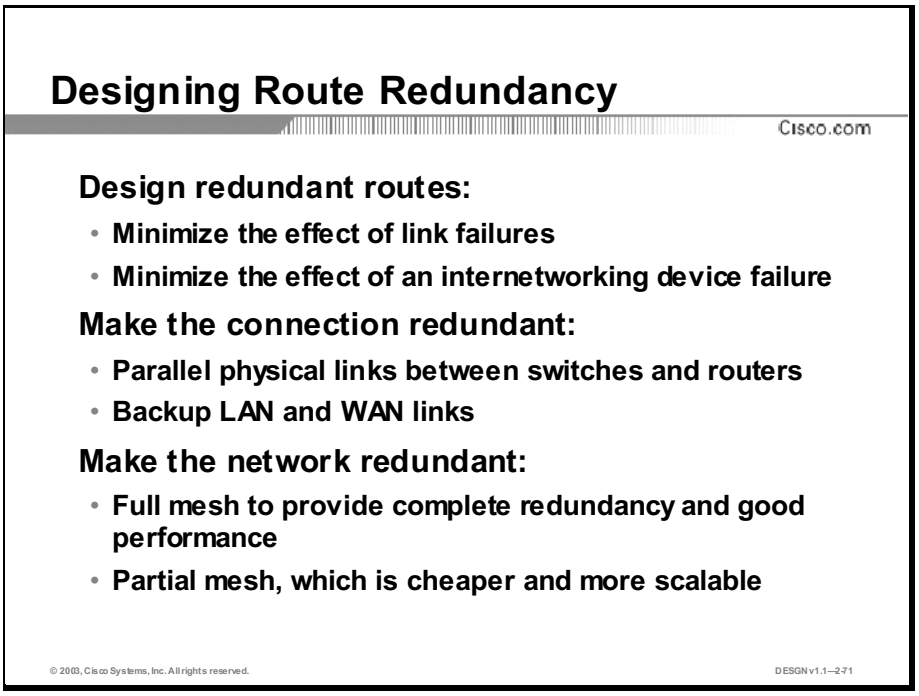

The common approach in designing route redundancy is to implement partial redundancy (a partial mesh instead of a full mesh, and backup links to the alternative concentrator), protecting only the most vital points of the network, such as the links between the layers and concentration devices.

Design Consideration: Full-Mesh vs. Partial-Mesh Redundancy

The full-mesh design forms any-to-any connectivity and is ideal for connecting a reasonably small number of devices together. However, as the network topology grows, the number of links required to maintain a full mesh increases dramatically.

The partial-mesh network is similar to the full-mesh network with some of the trunks removed. The partial-mesh backbone is appropriate for a campus network, where the traffic predominantly goes into one centralized Server Farm module.

Example: Campus Infrastructure Redundancy

Cisco.com

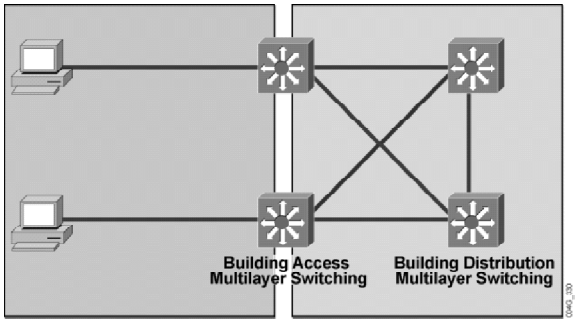

- **The Building Access network is partially meshed with the Building Distribution switches.**
- **The Building Access switch has a chance to recover from a link or Building Distribution switch failure.**

When multilayer switching is deployed in the Building Access submodule, Layer 2 (data link switching) is normally used between the Building Access switch and the workstations, and multilayer switching is used between the Building Access and Building Distribution switches.

The multilayer switches select the primary and backup path between the Building Access and Building Distribution submodules based on the cost of the link as computed by the routing protocol algorithm. The best path is placed in the forwarding table, and, in the case of equal paths, load sharing takes place.

Example: Enterprise Edge Redundancy

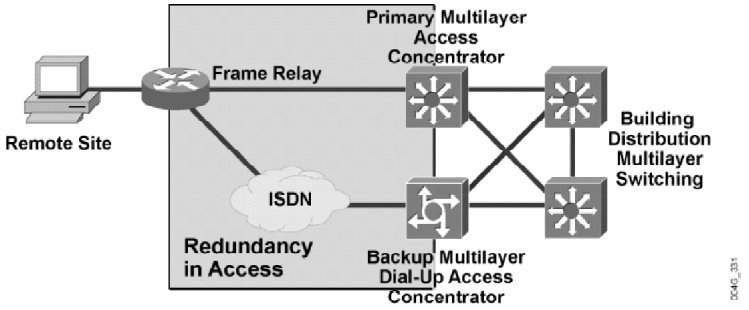

- **The remote site establishes a backup connection via the ISDN backup interface.**

Because WAN links are often critical pieces of the internetwork, redundant media are often deployed in WAN environments. You can provision backup links so that they are always on or only become active when a primary link goes down or becomes congested.

The backup links can use different technologies, as in the above scenario, where a Frame Relay circuit is used in parallel with a backup ISDN circuit. The primary requirement is sufficient capacity to meet critical needs.

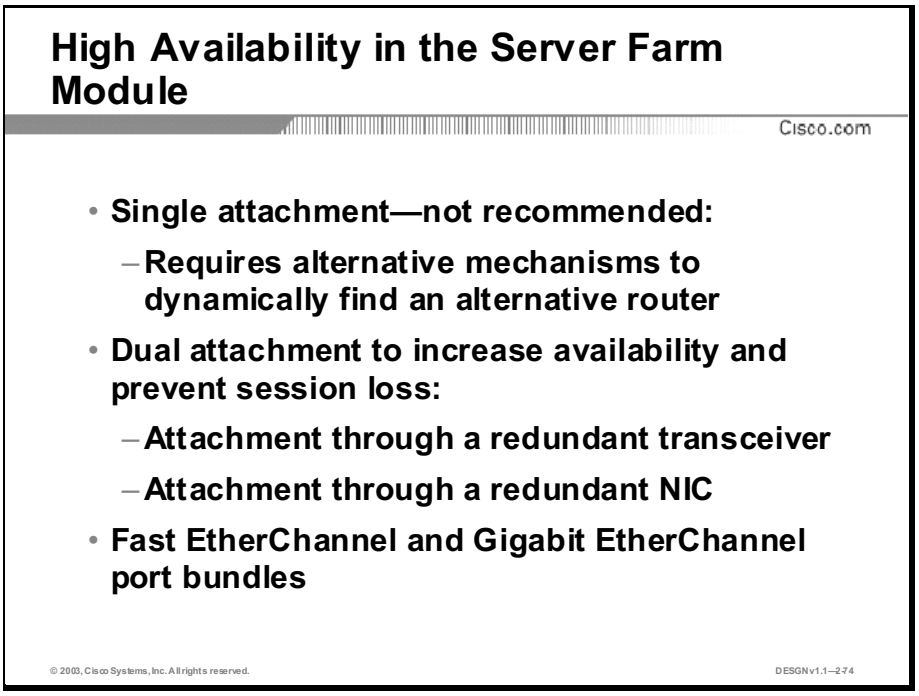

Because of the high cost, most network designers do not often implement a completely redundant network. Instead, network designers implement partially redundant internetworks. When network designers think about improving reliability on critical workstations and servers, the solution often depends on the workstation hardware and operating system software in use. These are some common attachment methods:

- **Single attachment:** In this case, a workstation needs to find dynamically an alternative router (Routing Information Protocol [RIP], Address Resolution Protocol [ARP], Router Discovery Protocol [RDP], or Hot Standby Router Protocol [HSRP]).
- **Attachment through a redundant transceiver:** Physical redundancy with a redundant transceiver attachment is suitable in environments where the workstation hardware or software does not support redundant attachment options.
- **Attachment through a redundant network interface card (NIC):** Some environments (for example, most UNIX servers) support redundant attachment through dual NICs (primary and backup) that the device driver presents as a single interface to the operating system.
- **Fast EtherChannel and Gigabit EtherChannel port bundles:** Use port bundles to group multiple Fast or Gigabit Ethernet ports into a single logical transmission path between switch and router, host, or other switch (Spanning Tree Protocol [STP] regards a channel as one link.). The switch distributes frames across the ports in an EtherChannel according to the source and destination MAC addresses. If a port within an EtherChannel fails, traffic previously carried over the failed port switches to the remaining ports within the EtherChannel.

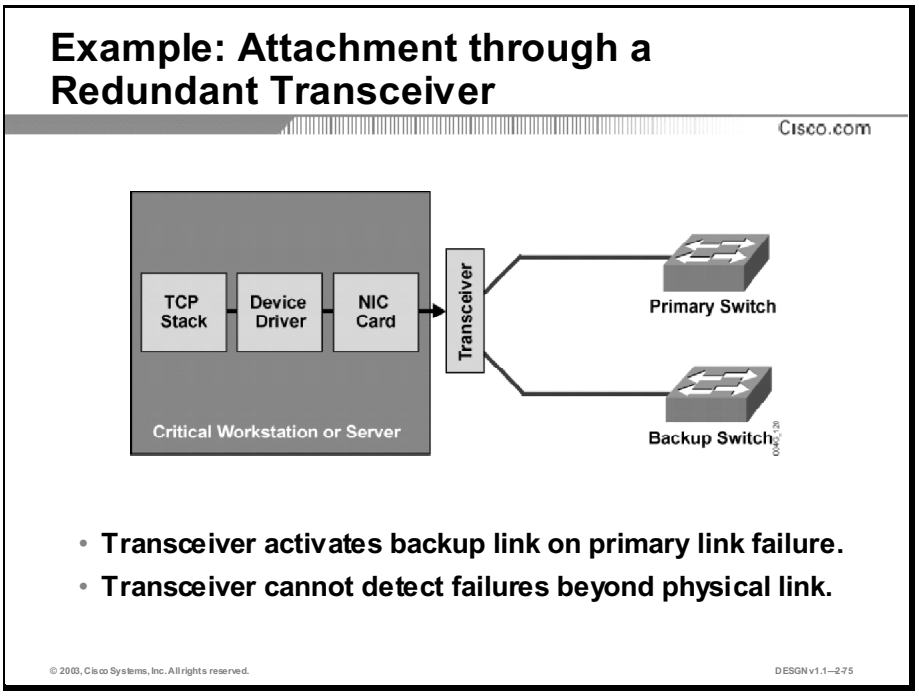

The figure shows server-to-switch connections implemented with a redundant transceiver. The redundant transceiver has two uplink ports that are usually connected to two access switches. The transceiver activates the backup port after it detects a link failure (carrier loss) on the primary port.

Note:	The redundant transceiver can detect only physical layer failures. It cannot detect failures inside the switch or failures beyond the first switch.

Note:	This type of redundancy is most often implemented on servers.

Example: Attachment through a Redundant NIC

Cisco.com

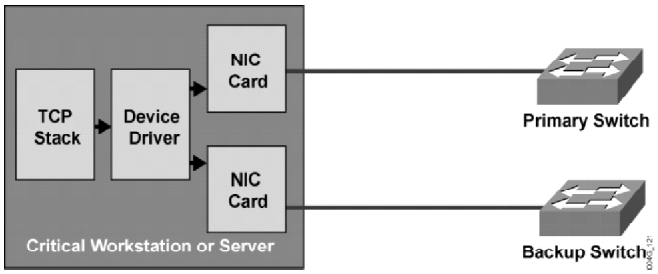

- **Device driver presents two NIC cards as a single logical interface.**
- **This setup uses one MAC address on both interfaces.**
- **Backup card is activated when the primary link is gone.**

The figure presents the case where redundancy is provided by installing an additional interface card in the server. The device driver presents configured NIC cards as a single interface (one IP address) to the operating system. The two NICs might use a common MAC address, or they might use two distinct MAC addresses and rely on gratuitous ARP to provide proper IP-to-MAC address mapping on switches when the backup interface card is activated.

Note: The workstation sends Gratuitous ARP messages to update the ARP and the forwarding tables on attached neighboring nodes (for example, Layer 2 switches).

Introducing Network Solutions

The network infrastructure and the intelligent network services provide the platform on which network designers implement network solutions (the term “overlay solutions” is sometimes used). You evaluate all the aspects of a network solution because modern networks must support a wide range of solutions. This topic describes network solutions.

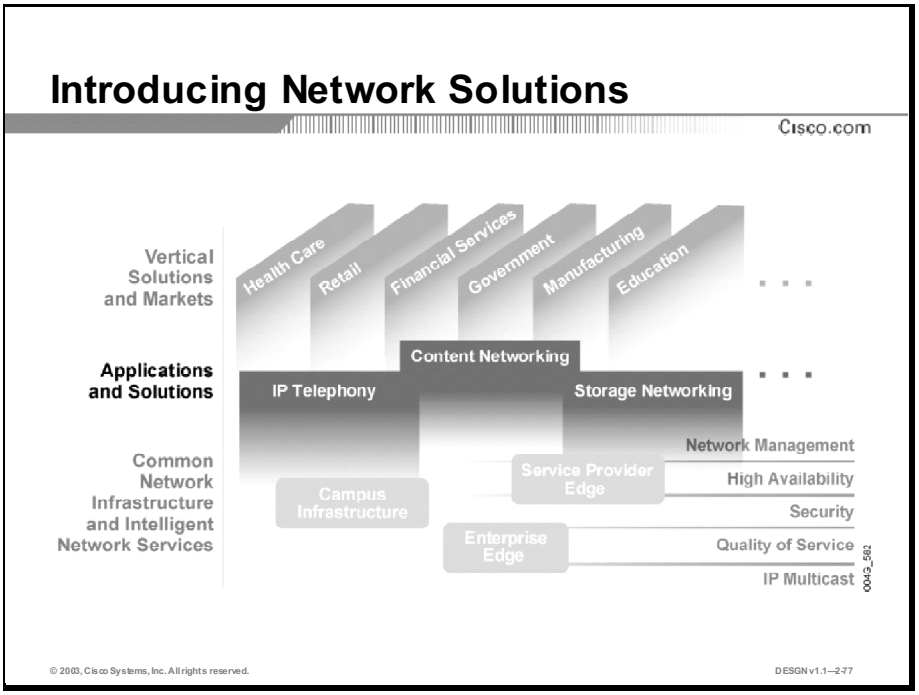

Examples of network solutions are:

- **IP telephony:** The convergence of voice, video, and data on a single IP network is changing the way enterprises communicate. Voice and video can now be transported as high-priority data, lowering network costs and optimizing business communications.
- **Content Networking:** Through internal developments and acquisition of breakthrough content-networking technologies, Cisco has built a comprehensive architecture for optimizing Web site performance and content delivery. This architecture comprises the five essential technology building blocks that provide the foundation for all of Cisco’s existing and future content-networking solutions.
- **Storage Networking:** Driven by workforce collaboration, e-commerce, and e-learning, storage networking has emerged as an important networked application. Cisco Storage Networking provides high capacity, low latency networking for disaster recovery, data replication, and storage consolidation.

Note: This course introduces two examples of network solutions (IP telephony and content networking) to provide an overview of how to implement a sample network solution network infrastructure.

Most network solutions span the entire enterprise network, but their requirements vary from network module to module. Modularity in network design allows you to create design elements to replicate as the network grows. As each element in the network design requires change, the cost and complexity of making the upgrade is confined to a small subset of the overall network.

You must focus on the specific functions that a network solution within a specific network module requires to simplify the design problem. Conversely, focusing on an individual module and its specific functions means that you must specify the required interfaces to other modules.

From the user perspective, the entire enterprise network is an integrated platform on which the network solution is implemented. The deployed solution must appear as a single solution to users.

IP Telephony Network Solution

IP telephony is a network solution that relies on the infrastructure and intelligent network services. To ensure successful implementation of a voice solution, you must consider the enterprise infrastructure and its configuration. This topic describes IP telephony as a network solution.

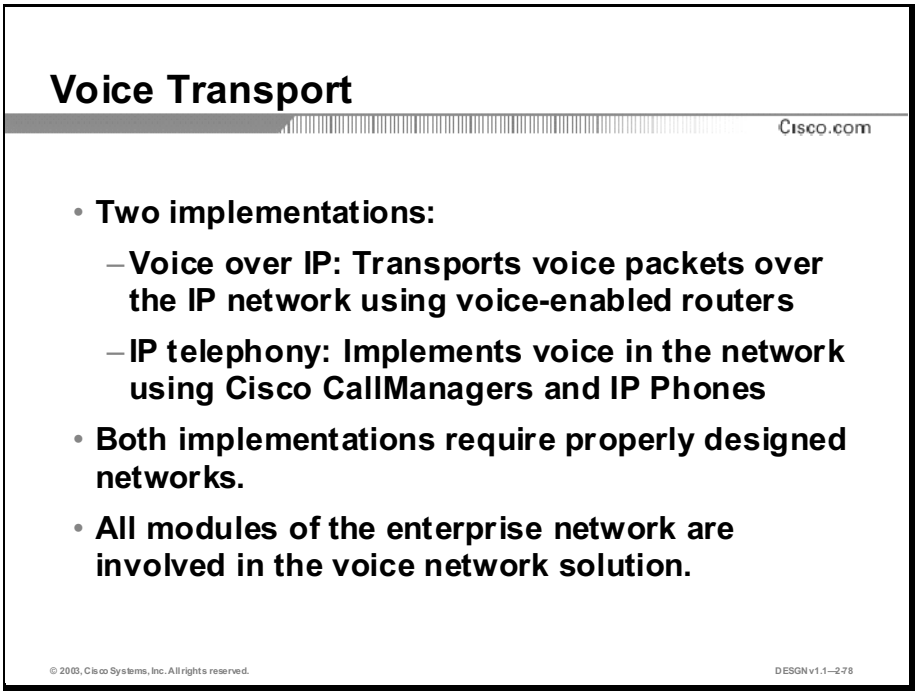

Voice is a very general term, divided into two implementations:

- **Voice over IP (VoIP):** VoIP allows voice-enabled routers to convert analog voice into IP packets or packetized digital voice channels and route those packets between locations. Users do not often notice that VoIP is implemented in the network. They use their traditional telephones connected to a PBX. However, the PBX is not connected to the Public Switched Telephone Network (PSTN) or to another PBX but to a voice-enabled router that is an entry point to VoIP. Voice-enabled routers can also terminate IP telephones using session initiation protocol (SIP) for call control and signaling.
- **IP telephony:** IP telephony is a voice implementation where traditional telephones are replaced with IP Phones and a server, Cisco CallManager, for call control and signaling. The IP Phone itself performs voice-to-IP conversion. Connection to the PSTN requires a voice-enabled router or other gateway in the Enterprise Edge functional area, where calls are forwarded to the PSTN.

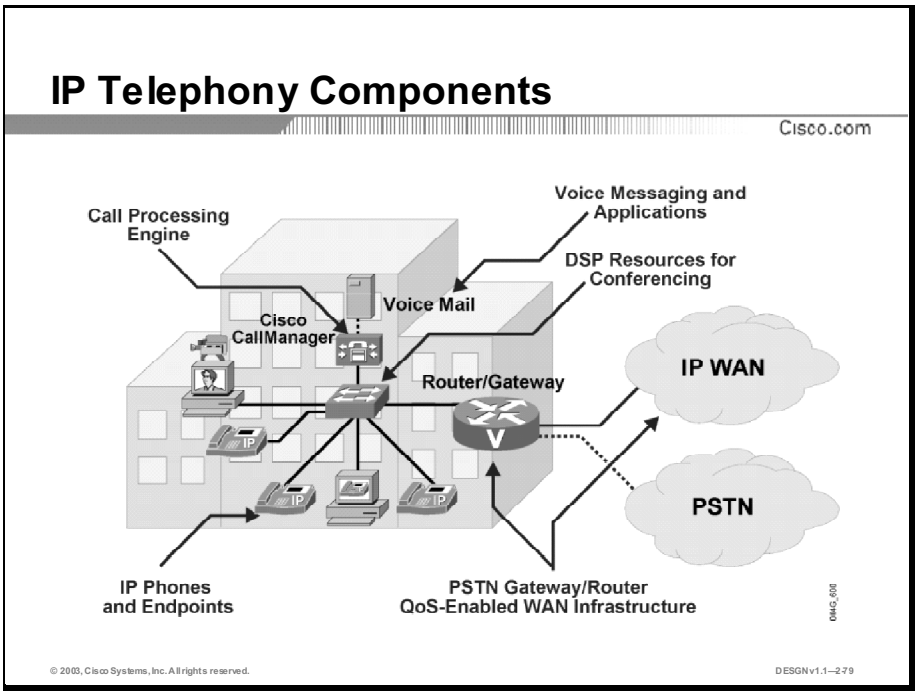

There are four main voice-specific components of the IP telephony network:

- **IP Phones:** IP Phones support calls in an IP telephony network. They perform voice-to-IP (and vice versa) coding and compression using special hardware. IP Phones offer services such as user directory lookups and Internet access for stock quotes. The telephones are active network devices and require power for their operation. Typically a network connection or an external power supply provides the power.
- **Switches with inline power:** Switches with inline power enable a modular wiring-closet infrastructure to provide centralized power for Cisco IP telephony networks. These switches are similar to traditional switches, with an option to provide power to the LAN ports where IP Phones are connected. In addition, they perform some basic QoS mechanisms, such as packet classification, which is a baseline for prioritizing voice through the network.
- **Call processing manager:** The CallManager provides central call control and configuration management for IP Phones. CallManager provides the core functionality to initialize IP telephony devices and perform call setup and routing of calls throughout the network. CallManager supports clustering, which provides a distributed scalable and highly available IP telephony model.
- **Voice gateway:** Voice gateways, also called voice-enabled routers or switches, provide voice services such as voice-to-IP coding and compression, PSTN access, IP packet routing, backup call processing, and voice services. Backup call processing allows voice gateways to take over call processing in case the primary call processing manager fails. Typically, voice gateways support a subset of the call processing functionality supported by CallManager.

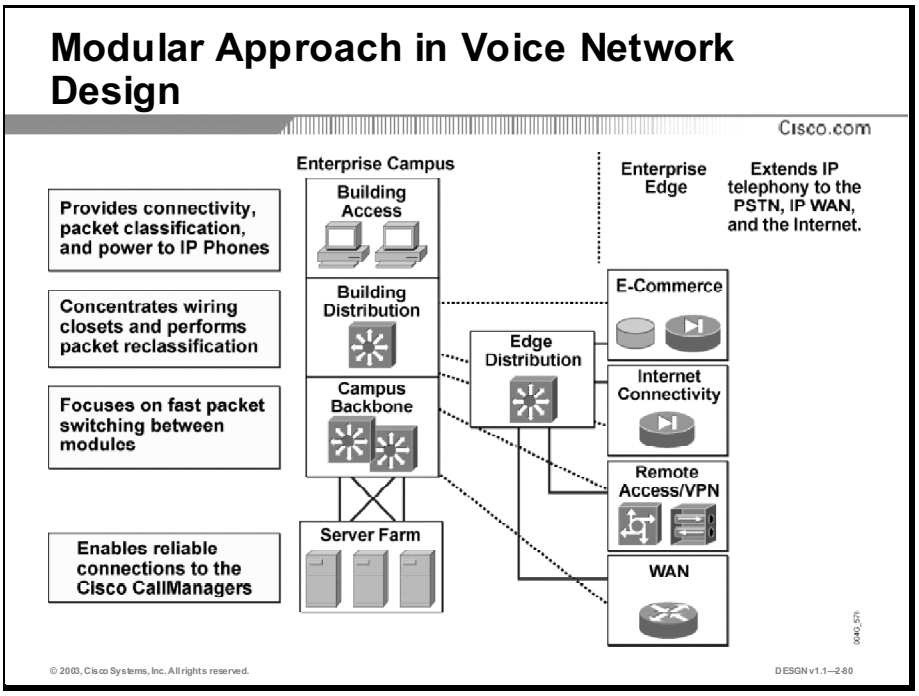

Enterprises that want to deploy new network services such as IP telephony on their networks face new design challenges. They need to deploy a delay-sensitive overlay service from end to end, through all enterprise network modules.

To simplify design, implementation, and troubleshooting, the modular approach is recommended. You should evaluate each module and submodule of the enterprise network before proceeding:

- **Building Access submodule:** IP Phones and user computers attach to L2 switches. Switches provide power and packet classification, which is essential for proper voice packet manipulation through the network.
- **Building Distribution submodule:** This submodule performs packet reclassification if the Building Access submodule is unable to perform packet classification. It concentrates Building Access switches (wiring closets) and provides redundant uplinks to the Campus Backbone submodule.
- **Campus Backbone submodule:** The Campus Backbone submodule forms the core of the network. All enterprise network modules are attached to it, and, therefore, virtually all traffic between application servers and clients traverses the Campus Backbone submodule. With the advent of wire-speed L3 gigabit switching devices, LAN backbones have migrated to switched gigabit architectures, which combine all the benefits of routing with wire-speed packet forwarding.
- **Server Farm module:** This module includes multilayer switches and CallManagers. Because CallManagers are the heart of IP telephony, redundant links and redundant CallManagers are essential for providing high availability.
- **Enterprise Edge:** The Enterprise Edge, with its modules (Remote Access and VPN and WAN modules, for example) can extend IP telephony from the Enterprise Campus to remote locations, the PSTN, and the Internet.

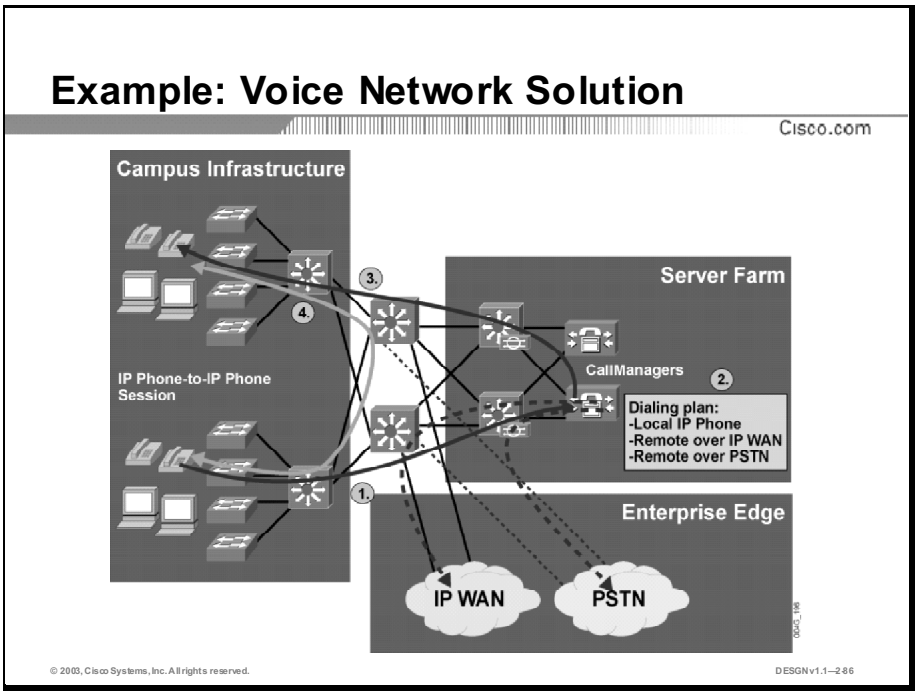

IP telephony requires modifications to the enterprise network infrastructure in terms of performance, capacity, and availability. It is an end-to-end solution, with clients (IP Phones) located in the Building Access submodule and the CallManager located in the Server Farm module.

The figure shows the voice network solution within the Enterprise Composite Network Model. It depicts how to modularize and evaluate the solution on a module-by-module basis. The figure shows how a call is initiated on an IP Phone, how the call setup goes through the CallManager, and how the end-to-end session between two IP Phones is established. CallManager is involved only in the call setup.

Three scenarios of IP call routing are introduced:

- Calls destined to remote locations traverse the Enterprise Edge through the WAN or Remote Access and VPN module.
- Calls destined to PSTNs are routed over the Enterprise Edge through the Remote Access and VPN module.
- Calls between IP Phones traverse the Server Farm module and Building Access, Building Distribution, and Campus Backbone submodules. Call setup uses all of these modules; speech transport employs only the Building Access, Building Distribution, and, in some cases, the Campus Backbone submodules.

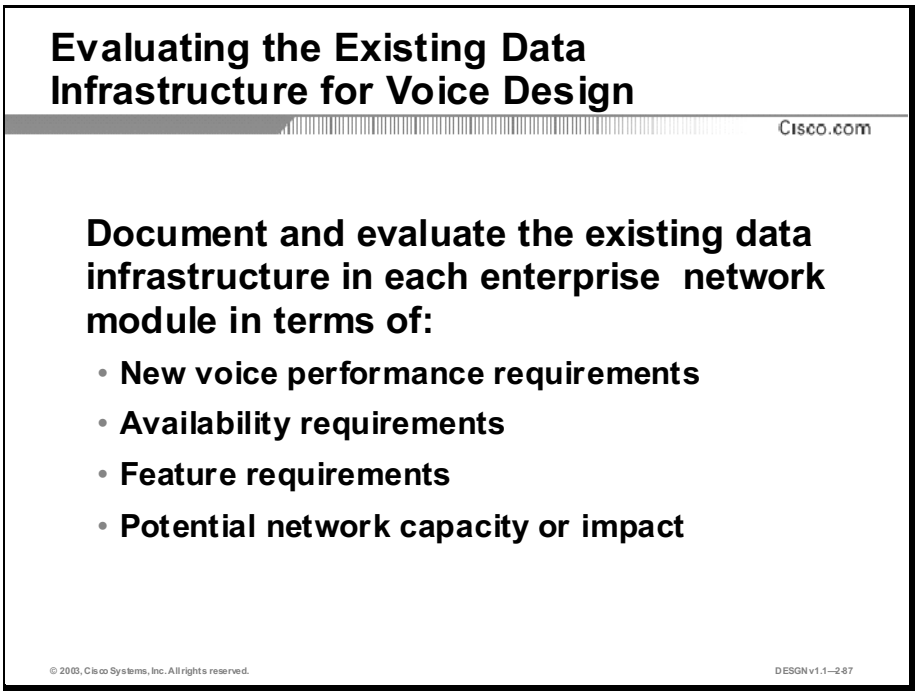

When designing IP telephony, you need to evaluate the existing data infrastructure in each enterprise network module to determine upgrade requirements for the IP telephony solution. Follow these guidelines when designing the data infrastructure to support voice:

- Provide infrastructure for additional bandwidth, consistent performance, or higher availability, if required, for the converging environment. Links and devices should have sufficient capacity for voice traffic. Links with high peak or busy-hour use may require an upgrade. Target devices for additional inspection and potential upgrades are those with high CPU use, high backplane use, high memory use, queuing drops, or buffer misses.
- Review the redundancy capabilities in all network modules to ensure that they can meet availability goals with the current network design (or new design) recommended for IP telephony.
- Evaluate device characteristics, including the chassis, module, and software inventory. This evaluation will prove useful in determining IP telephony feature capabilities in the existing environment.
- Evaluate overall network capacity and the impact of IP telephony on a module-by-module basis. This activity ensures that the network meets capacity requirements and that there will be no adverse impact on the existing network and application requirements.

Content Networking Network Solution

A Content Delivery Network (CDN) is a collection of comprehensive architectures and technologies that optimize web site performance and other content delivery. A CDN adds a layer of intelligence between the fundamental network functions and applications, as requested by the users. The goal of a CDN is to ensure that, transparent to the user, the network serves content with optimal resource usage. This topic describes content networking as a network solution.

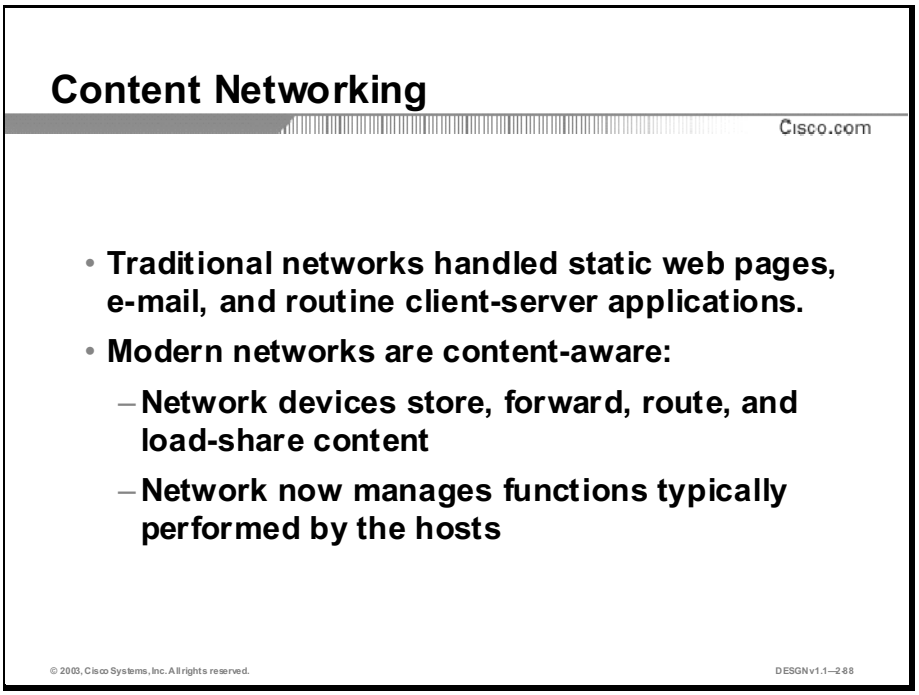

Traditional networks handled static web pages, e-mail, and client-server traffic. Modern enterprise networks need to handle more sophisticated types of network applications that include both voice and video, such as videoconferencing, corporate training, audio and video broadcasts, and so on.

The large amount of data and its variety require that the modern network be aware of the content carried across it to handle that content optimally. It is no longer enough to add more bandwidth as needs grow. Networks have had to become smarter.

With content networking, the network itself manages several functions that the hosts (servers and clients) typically performed. Dedicated network devices store, forward, route, and do load-sharing of the content.

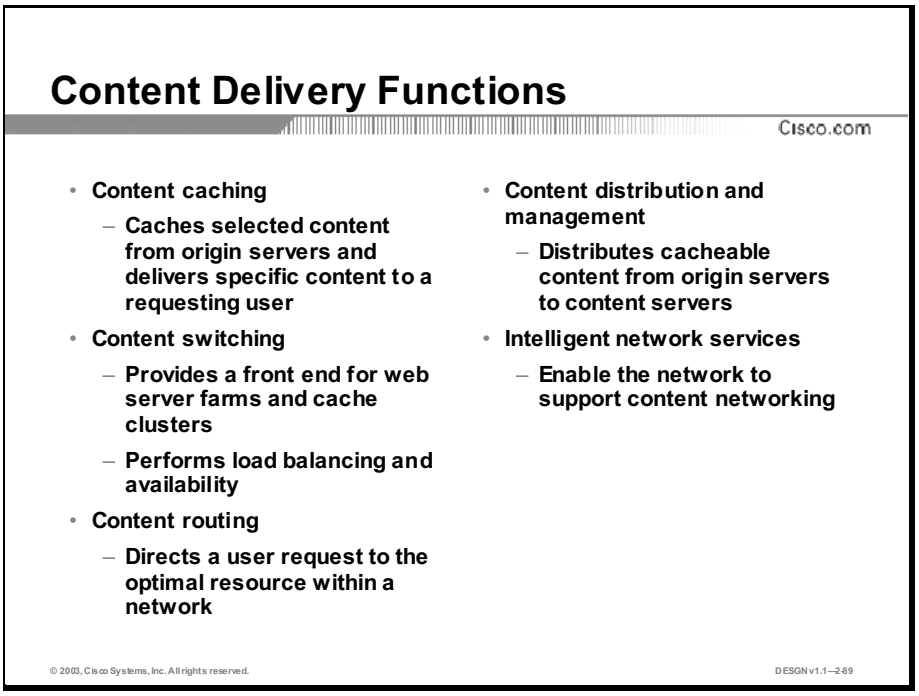

Content networking offers accelerated content delivery, hosting, and other content-based services. It addresses the need to distribute and receive high-bandwidth, media-rich content across the Internet or an intranet without performance losses or content delivery delays. Content networks typically have three delivery functions:

- **Caching:** A cache selects content from origin servers, caches the content, and delivers specific content to the requesting user. Today, content engines handle static and streaming media content. Soon, dynamic database content and applications will be cached.
- **Content switching:** Content switching provides a robust front end for web server farms and cache clusters, performing important functions such as load balancing of user requests across web server farms, policy-based web traffic direction based on full visibility of URLs, and so on.
- **Content routing:** Content routing directs a user request to the optimal resource within a global network based on user-defined policies, such as rules for specific content, availability of content, network health, current loads for web servers or caches, and various other network conditions.

A network administrator must manage content delivery in the enterprise network to ensure content freshness and the movement of the content to proper places in the network. The management tool is also needed to configure and monitor content networking devices.

Content networking is considered a network solution, and, as such, it requires proper intelligent network services to support it. In addition to security and QoS, content networking may require the intelligent network service IP multicasting for efficient delivery of the content to multiple destinations simultaneously.

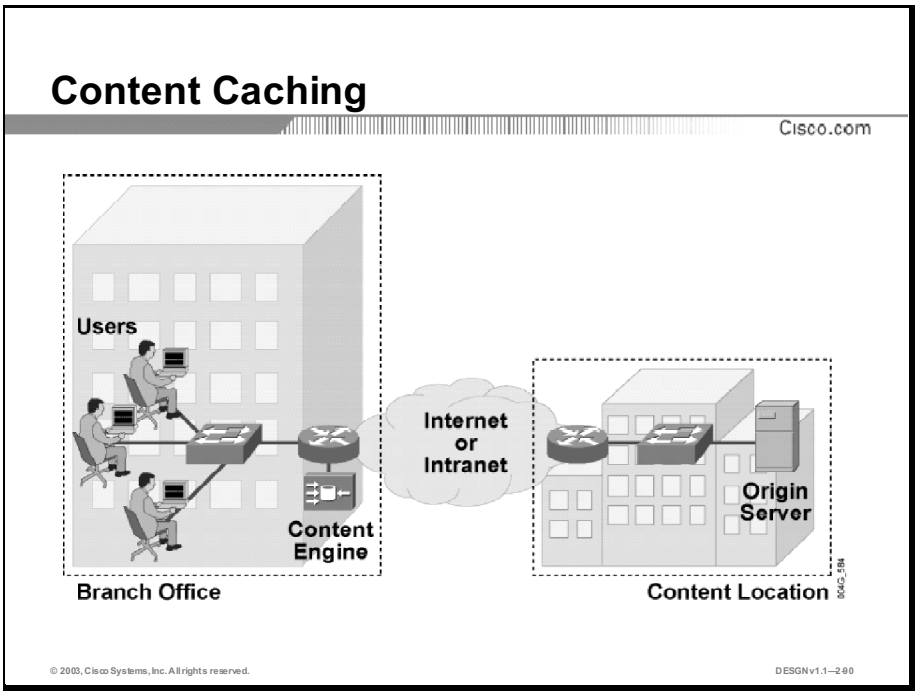

Content caches, such as the Cisco Content Engine, accelerate content delivery for end users by transparently caching frequently accessed content and then locally fulfilling content requests rather than traversing the Internet, the intranet, or both, to a distant server using the Web Cache Communication Protocol (WCCP). This solution helps protect the enterprise network from uncontrollable bottlenecks and accelerate content delivery, enabling enterprise employees to be more productive. Caches must store the most needed and most current data.

By caching streaming media and web content, caches minimize redundant network traffic that traverses WAN links. As a result, WAN bandwidth costs either decrease or grow less quickly. This bandwidth optimization increases network capacity for additional users and traffic and for new services, such as voice. Typical bandwidth savings range from 25 to 60 percent.

Deployment of Content Caches

You can deploy caches hierarchically at enterprise branch offices (within an Enterprise Edge) and at the central site (in the Internet Connectivity module) to provide optimal network response times. Caching at the central site reduces Internet access bandwidth consumption. Caching at the branch office reduces bandwidth consumption and improves response time for Internet and intranet connectivity.

Note: Placement of network caches is very sensitive. You must identify the content to be cached and the major traffic streams and directions of the traffic.

Because of content caching, employees can retrieve content more quickly and improve their productivity. Caching at the branch office may free up bandwidth on intersite links and make room for voice traffic. The benefits are accelerated content delivery, WAN bandwidth savings, and protection against uncontrollable bottlenecks, all resulting in higher productivity.

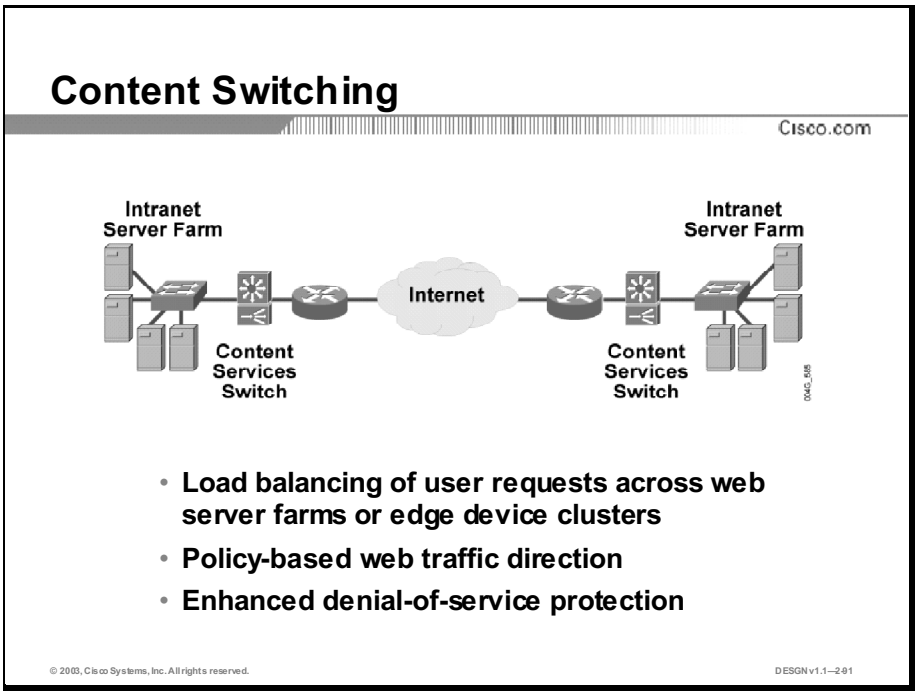

Content switching is a new generation of networking specifically designed to address the unique requirements of web traffic (thus, it is sometimes referred to as web switching). Content switching intelligently performs load balancing of traffic across multiple servers or cache devices based on content availability and the load on the server or cache device.

Content switches have sophisticated load-balancing capabilities and content-acceleration intelligence. The result is a consistently positive experience for website users.

Deployment of Content Switches

You can install content switches in front of the servers or cache devices in any Server Farm, E-Commerce, or Internet Connectivity module.

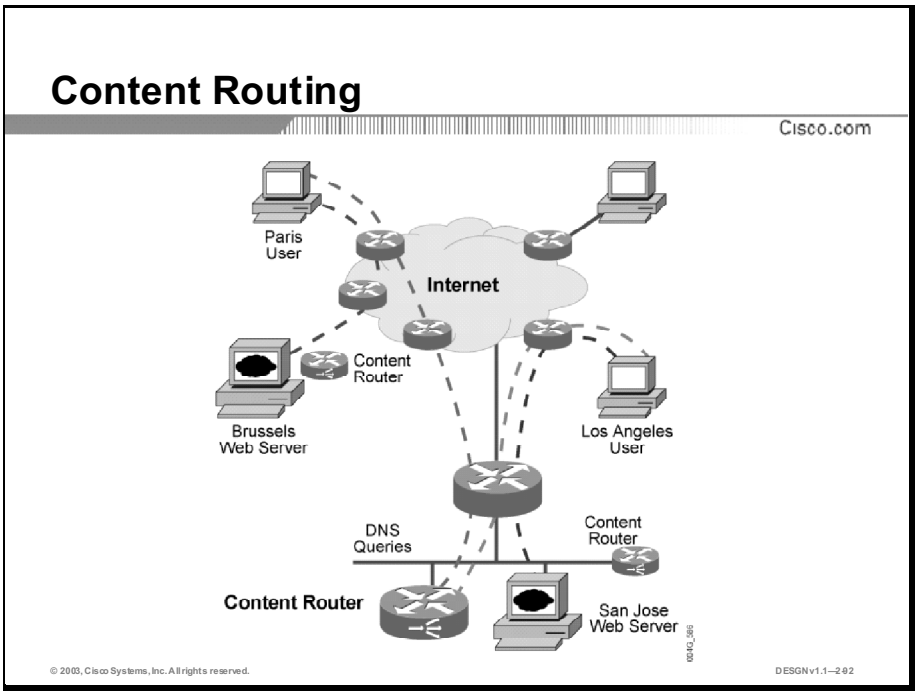

Content routing redirects an end-user request to the best server, based on a set of metrics such as delay, topology, and server load, and a set of policies such as location of content. This network feature enables the accelerated delivery of web content and streaming media. Content routing uses two methods: DNS server lookup and HTTP redirects.

As the number of users accessing content on the network grows, it becomes increasingly difficult to provide a high level of availability and rapid response from a single location. The solution to this problem is content routing among multiple locations.

Note:	Your understanding of the applications used in the networks is extremely important. Proper content router and DNS server configuration requires an excellent understanding of the behavior of the applications.

Content routing ensures the fast delivery of content, regardless of location, and provides high availability and improved server response.

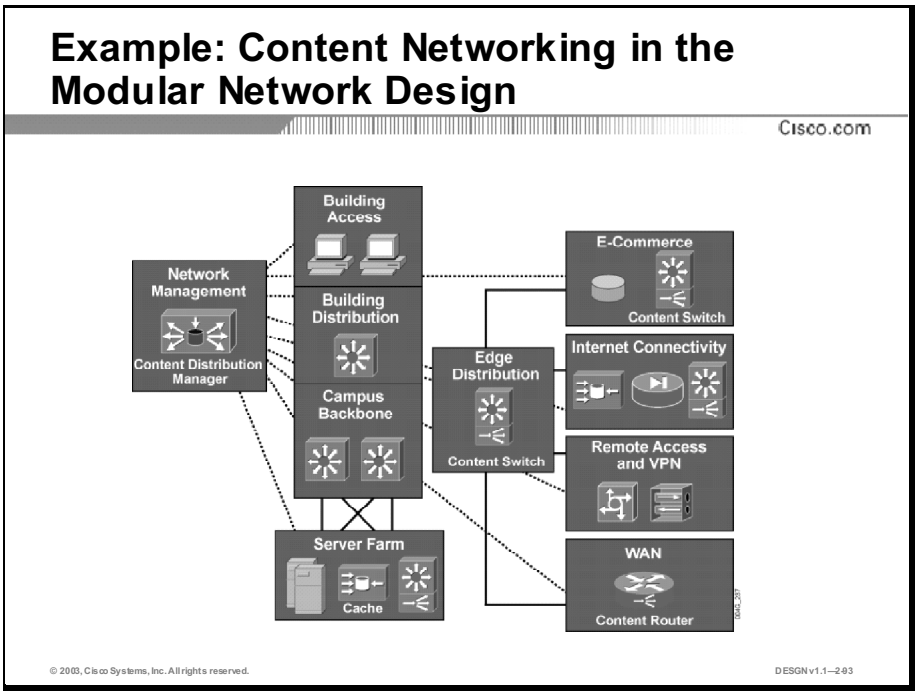

The figure shows how and where enterprises can implement content networking in the Enterprise Composite Network Model and how the network uses most of the content networking technologies.

You can deploy content switches inside, or at the edge of, all network modules that contain servers, to perform load sharing of requests and forward them to the least used server. The Server Farm module itself contains the cache device. The device servers cache in a more specific way, so the reverse-proxy cache relieves the server from serving the external requests by downloading the frequently accessed pages into the cache.

In the Network Management module, the Cisco Content Distribution Manager performs all management functions needed to control content distribution.

The content router is deployed in the WAN module to ensure that user requests are routed to the nearest servers.

Applications Influence Choice of Content Networking Technology

When considering future applications, you must identify the types of applications. For example, individual training often uses video on demand (VoD), which is normally done in a unicast manner. Conversely, corporate announcements are of a broadcast nature, so you can deliver them with multicast technology. Thus, the selection of the content networking technology and respective devices depends heavily on the type of the application.

Summary

This topic summarizes the key points discussed in this lesson.

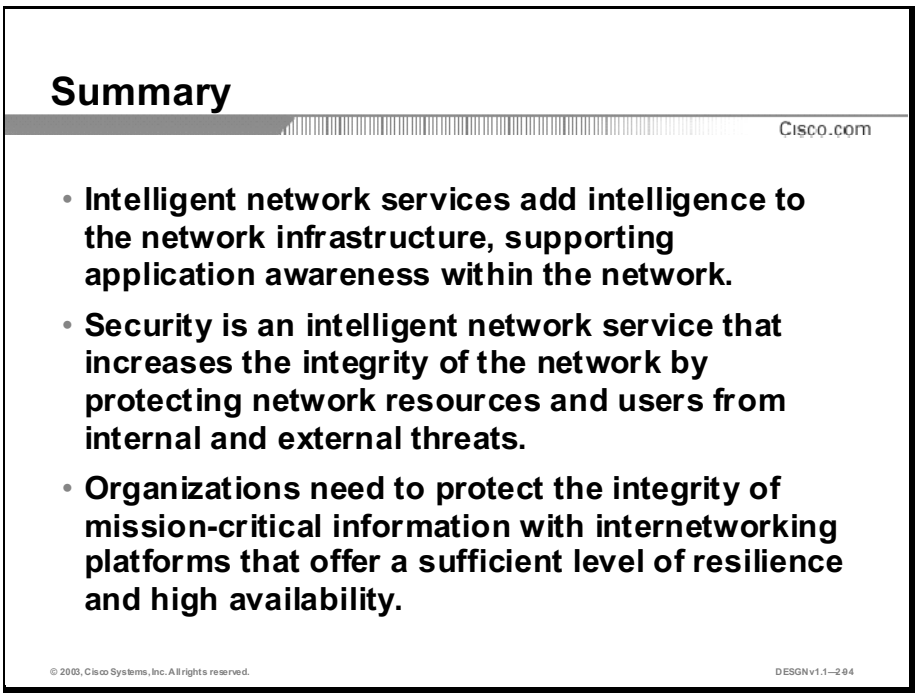

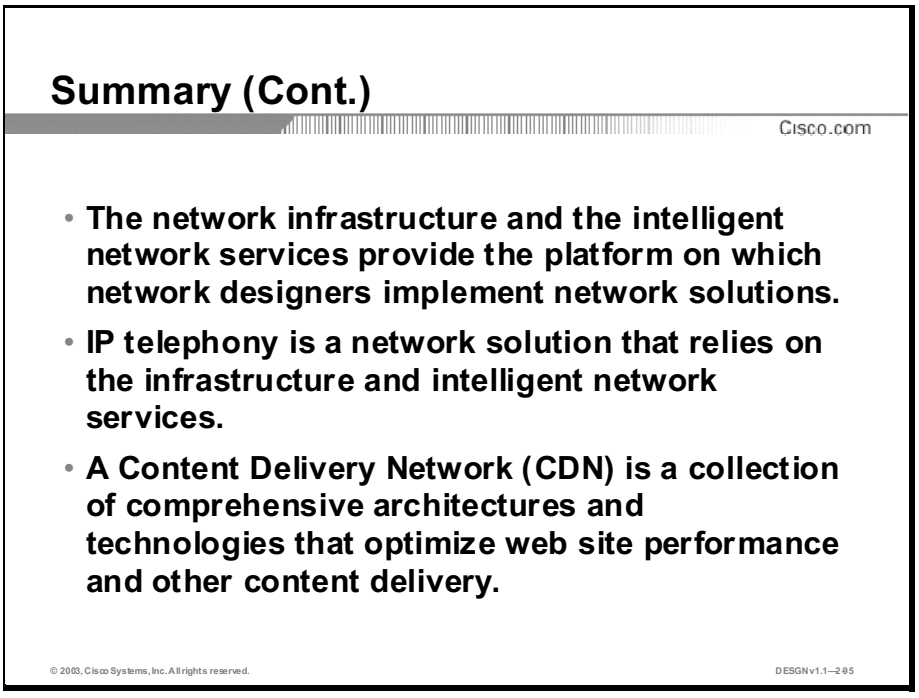

References

For additional information, refer to these resources:

- *SAFE White Paper: A Security Blueprint for Enterprise Networks,* http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm
- *Cisco Architecture for Voice, Video and Integrated Data White Paper,* http://www.cisco.com/warp/public/cc/so/neso/vvda/iptl/avvid_wp.htm
- *Cisco Enterprise Solutions: Cisco Content Delivery Networks,* http://www.cisco.com/en/US/netsol/ns110/ns49/net_solution_home.html

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Case Study 2: Designing a Network Hierarchy

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three choices are intelligent network services? (Choose three.)

A) IP telephony

B) IP forwarding

C) security

D) EIGRP

E) QoS

F) high availability

Q2) Which module of the Enterprise Composite Network Model is responsible for authentication, authorization, and accounting (AAA) of users, and stores One Time Passwords (OTP)?

A) Remote Access/VPN module

B) WAN module

C) Internet Connectivity module

D) Network Management module

E) Server Farm module

Q3) Which WAN technology is commonly used to provide redundancy in an Enterprise Edge environment?

A) ISDN

B) Frame Relay

C) leased lines

D) Point-to-Point Protocol

Q4) Which topology is best suited for connectivity in the Building Distribution submodule?

A) full mesh

B) hub and spoke

C) partial mesh

D) combination of full mesh and partial mesh

Q5) Select three benefits of using a modular approach for a network solution. (Choose three.)

A) facilitates changes

B) maintains integrity through the entire network

C) identifies interfaces between modules

D) encourages replication of design elements

E) confines complexity to a small subset of the network

Q6) Which three functions does the Building Access module of the enterprise network provide for an IP telephony network solution? (Choose three.)

A) call routing

B) connections to the PSTN and IP WAN

C) IP Phone connectivity

D) packet classification

E) power to the IP Phone

F) call setup

Q7) Match the lettered content delivery functions with the appropriate numbered descriptions.

A) content caching

B) content routing

C) content switching

_____ 1. redirects an end-user request to the best server based on a set of metrics

_____ 2. performs intelligent load balancing of traffic across multiple servers or cache devices

_____ 3. accelerates content delivery by transparently storing frequently accessed content

Quiz Answer Key

Q1) C, E, F

Relates to: Introducing Intelligent Network Services

Q2) D

Relates to: Security Intelligent Network Service

Q3) A

Relates to: High Availability Intelligent Network Service

Q4) C

Relates to: High Availability Intelligent Network Service

Q5) A, D, E

Relates to: Introducing Network Solutions

Q6) C, D, E

Relates to: IP Telephony Network Solution

Q7) 1=B, 2=C, 3=A

Relates to: Content Networking Network Solution

Case Study 2: Designing a Network Hierarchy

Complete this case study to practice what you learned in this module.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this module, you defined the Enterprise Composite Network Model. Upon completing this case study, you will be able to apply the Enterprise Composite Network Model to the requirements of DJMP Industries and meet these objectives:

- Provide an Enterprise Composite Network Model of the headquarters site
- Provide an Enterprise Composite Network Model of the regional sites
- Provide an Enterprise Composite Network Model for a remote office

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix, "Case Study Solutions"

Exercise Procedure

Complete these steps:

Step 1 Refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario before commencing the exercise. Focus on the network hierarchy requirements. Allow a maximum of 10 minutes for reading.

Step 2 Discuss the scenario and options for each functional area of the Enterprise Composite Network Model with your group. Allow 10 minutes for a discussion.

Step 3 Diagram the functional areas of the Enterprise Composite Network Model for the company headquarters network.

Step 4 Diagram the functional areas of the Enterprise Composite Network Model for the regional site networks.

Step 5 Diagram the functional areas of the Enterprise Composite Network Model for the remote office networks.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class.

Module 3

Designing Basic Campus-Switched Networks

Overview

The availability of multigigabit campus switches enables organizations to build high-performance, highly reliable networks. When you follow a systematic network design approach, performance, reliability, and manageability are achievable. This module describes a hierarchical campus network design approach called “multilayer design.” Multilayer design is modular, so organizations can increase their network capacity as needed. A multilayer campus design is based on a known campus topology, which aids troubleshooting.

This module first introduces general campus switching design considerations. You will also learn the basic difference between data link layer and multilayer switching and receive recommendations about where to use each. In addition, the module describes the switching modularity and scalability options appropriate for situations ranging from building-sized networks to large campus networks.

Module Objectives

Upon completing this module, you will be able to explain how to position switches in a campus network design.

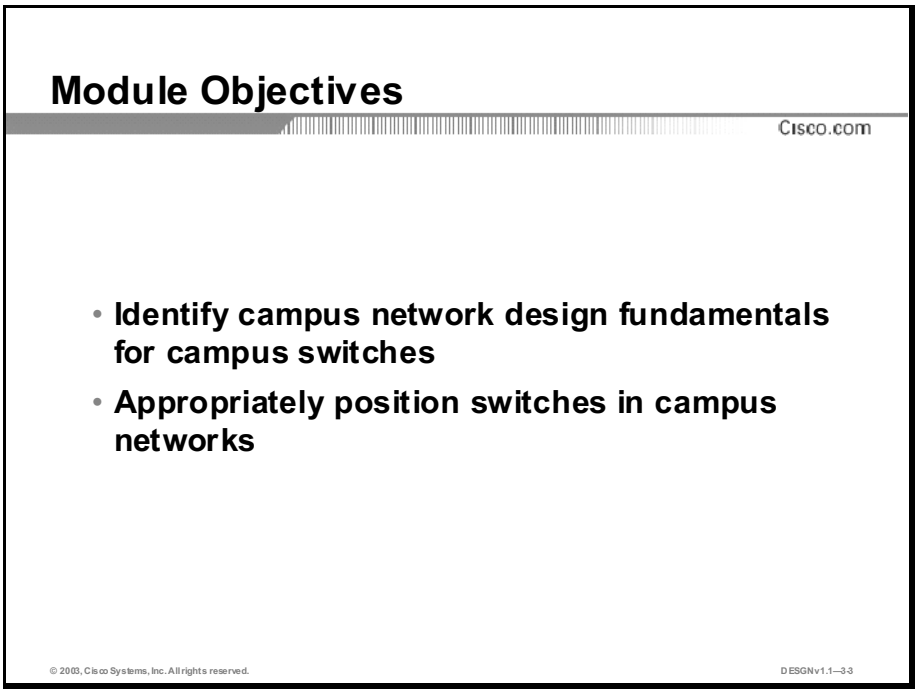

Module Outline

The outline lists the components of this module.

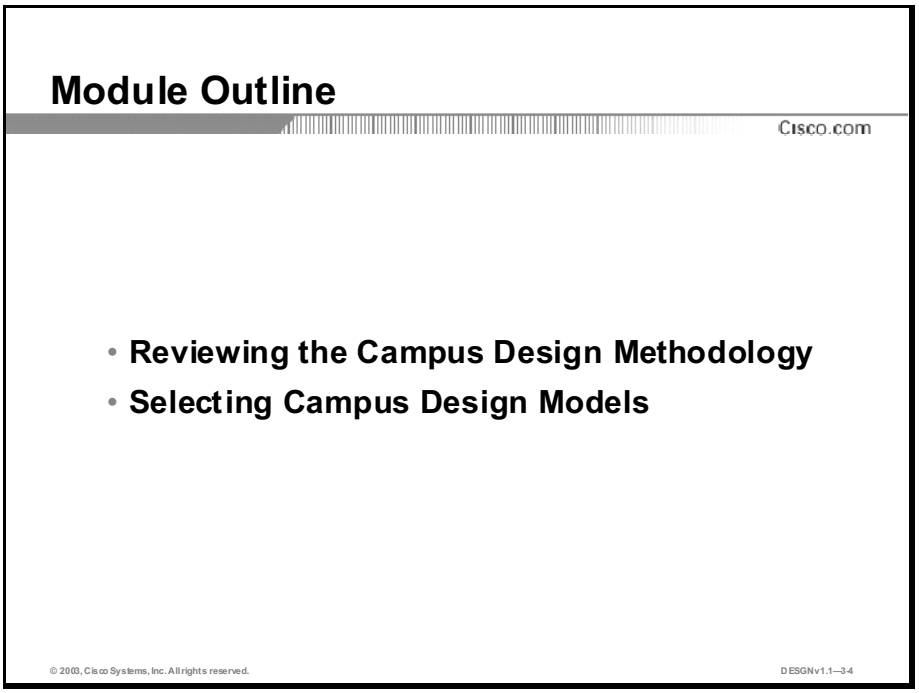

Reviewing the Campus Design Methodology

Overview

The multilayer approach to campus network design combines data link layer switching with multilayer switching to achieve robust, highly available campus networks. This lesson introduces the transmission media used in switched campus networks. The lesson further describes general campus switching considerations that you will analyze and consider in your campus designs. Several decision tables provide guidelines about where to use each technology.

Relevance

This lesson provides an overview of the general technologies used in campus switching design. You will need to be aware of the network topology options and network configuration parameters of the network devices so that you can effectively use them in your designs.

Objectives

Upon completing this lesson, you will be able to identify campus network design fundamentals for campus switches. This includes being able to meet these objectives:

- Identify the components of an enterprise campus design
- Identify the geography of the enterprise campus and its effects on the network design
- Determine the application requirements and their corresponding traffic requirements in enterprises
- Identify the correct use of copper and fiber cabling in a campus network design
- List the benefits of switched over shared technology
- Identify when to use data link layer and multilayer switching in an enterprise design

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, functions, and general LAN switching theory

Outline

The outline lists the topics included in this lesson.

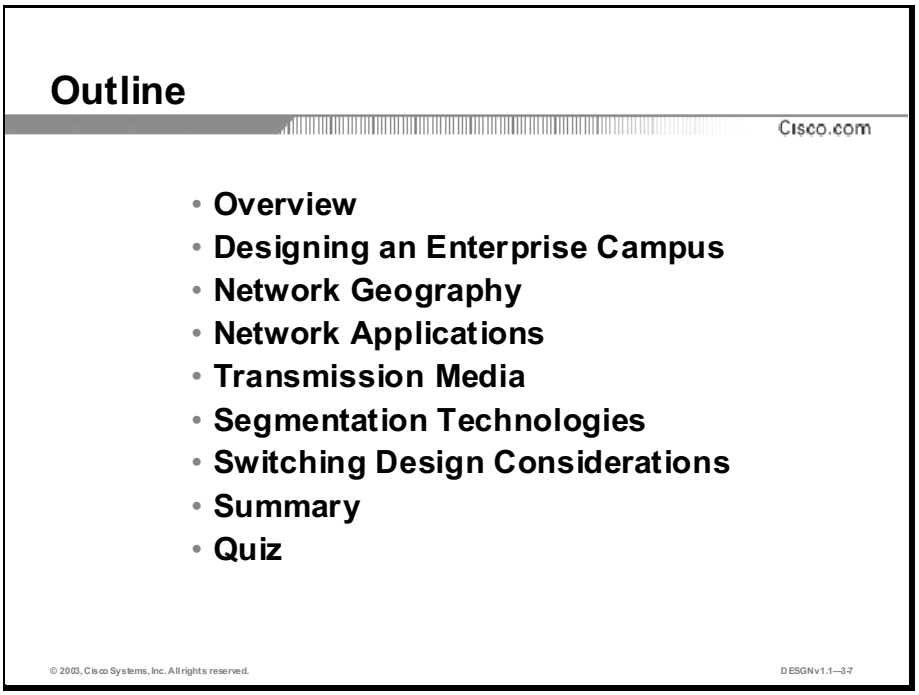

Designing an Enterprise Campus

Designing an enterprise campus network requires a broad view of the overall network. You must be familiar with both enterprise campus design methodologies and with Enterprise Campus modules of the Enterprise Composite Network Model. This topic describes an enterprise campus design methodology.

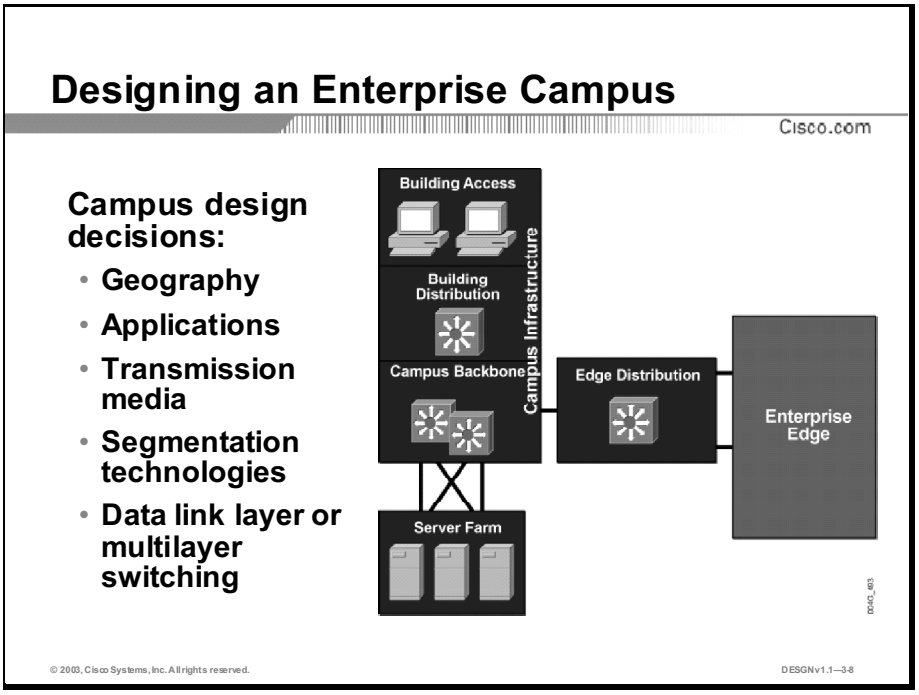

You should consider five factors when designing the campus network.

- **Geography:** The distribution of network nodes (hosts, network devices) and the distances between them affect the campus solution significantly.
- **Applications:** The application requirements, in terms of bandwidth and delay, place stringent requirements on a campus network solution.
- **Transmission media:** Cabling is one of the biggest long-term investments in network deployment. Therefore, the selection of the transmission media depends on the required bandwidth and distances as well as on the emerging technologies that might be deployed over the same infrastructure in the future. You will thoroughly evaluate the cost of the medium (including installation costs) and the available budget, in addition to the technical characteristics such as signal attenuation and electromagnetic interference. Two major cabling options exist: copper-based media and optical fiber.
- **Segmentation technologies:** The traditional approach, where devices share the available bandwidth, is being replaced with dedicated bandwidth LAN switching.
- **Data link layer or multilayer switching:** The network devices and their features support the network functions and contribute to overall network delay. The network designers usually consider data link layer switching (based on MAC address) versus multilayer switching (based on network layer address, transport layer, and application awareness).

Network Geography

The location of enterprise campus nodes and the distances between them influence the network topology. This topic describes how network geography affects network design.

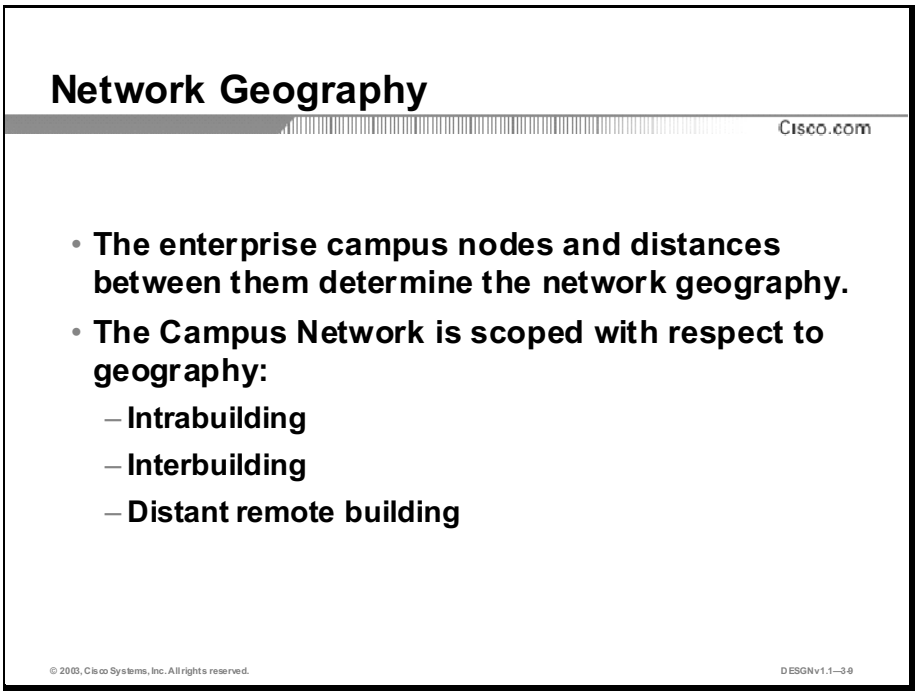

When designing an enterprise campus network, you must determine the location of nodes and the distances between nodes.

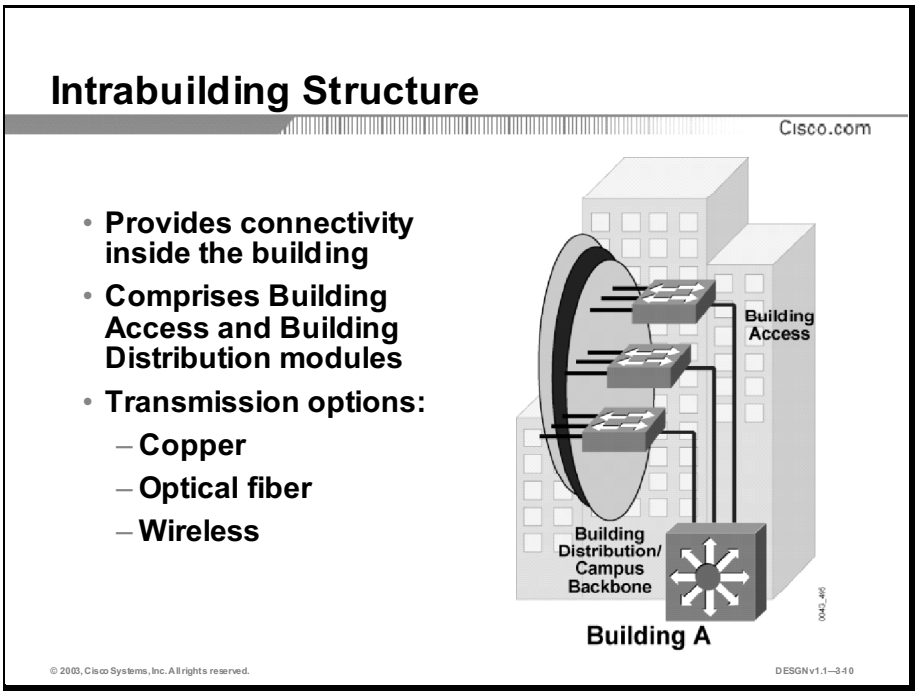

An intrabuilding campus network structure ensures connectivity for the end nodes, which are all located in the same building, and provides them with access to the network resources.

User workstations are normally attached via twisted-pair cables to the floor-wiring closet.

Wiring closets are normally connected to the Building Distribution switch over optical fiber. This offers better transmission performances and is less sensitive to environmental disturbances.

Wireless local area networking offers another intrabuilding solution. Wireless LANs enable users to establish and maintain a wireless network connection throughout or between buildings, without the limitations of wires or cables.

Interbuilding Structure

Cisco.com

- **Connectivity between buildings**
- **Distances between buildings within a few kilometers**
- **Building Distribution or Campus Backbone modules or both**
- **Transmission option optical fiber**

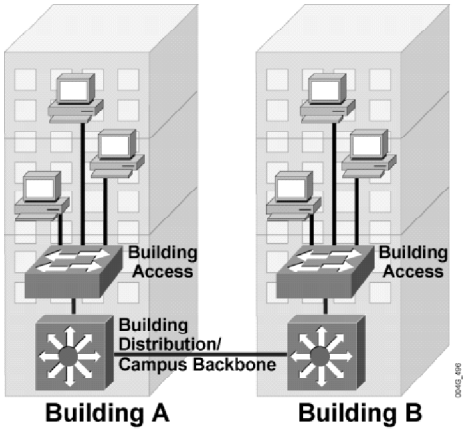

DESGN v1.1—3-11

An interbuilding network structure provides connectivity between the central switches of the individual campus buildings. These buildings are normally in close proximity and are typically only a few hundred meters or a few kilometers apart.

Because the nodes in all the campus buildings normally share common devices, such as servers, the demand for high-speed connectivity between the buildings is high. To provide high throughput without excessive interference from environmental conditions, the appropriate physical media must be deployed. This requirement narrows the recommended choices to optical fiber.

Distant Remote Building Structure

Cisco.com

Metropolitan-based network connectivity options:

- **Using company-owned fiber**
- **Through enterprise WAN**
- **Through service provider metro Ethernet**

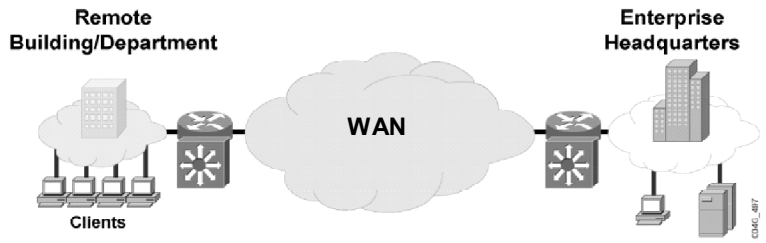

DESGN v1.1—3-12

To implement connectivity on distances that exceed a few kilometers (usually within a metropolitan area), the physical media is the most important factor that the network designer should consider.

Some companies may own their media, such as fiber, microwave, or copper lines. The speed and the cost of the network infrastructure depend heavily on these considerations.

When the bandwidth requirements are higher than the physical connectivity options can support, you need to identify the organization's critical applications and select the equipment that supports intelligent network services such as quality of service (QoS) and filtering capabilities.

If the organization does not own physical transmission media to certain remote locations, use WAN connectivity options from public service providers to connect the enterprise campus locations, or from a service provider metro Ethernet solution.

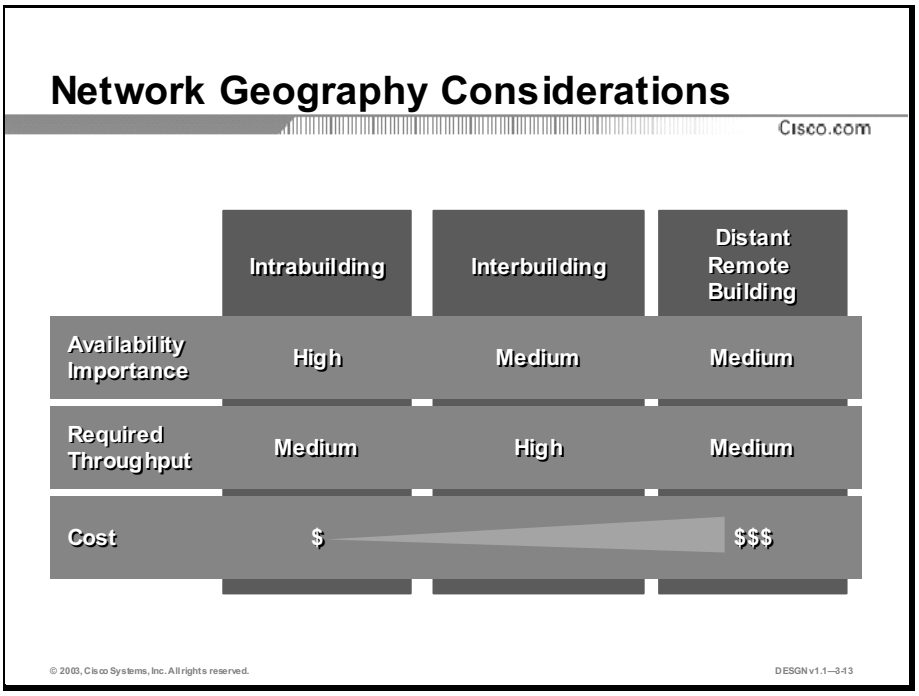

The figure compares the availability importance, required throughput, and expected cost for each geographical scope.

Typically, the importance of availability is very high within a building and decreases with distance into the Campus Backbone submodule. The Campus Backbone submodule should already provide high availability through fault-tolerance hardware and redundant links between buildings.

The throughput requirements increase close to the core of the network and close to the sites where the servers reside.

A balance between the desired bandwidth and available budget is required to keep the cost reasonable. The additional hardware and software cost to implement QoS features may prove more desirable than provisioning the sufficient bandwidth to meet peek requirements.

Network Applications

Application characterization provides information on the network usage and response times for each application. These parameters influence the selection of the transmission medium to provide the desired bandwidth. This topic describes how network application requirements influence network design.

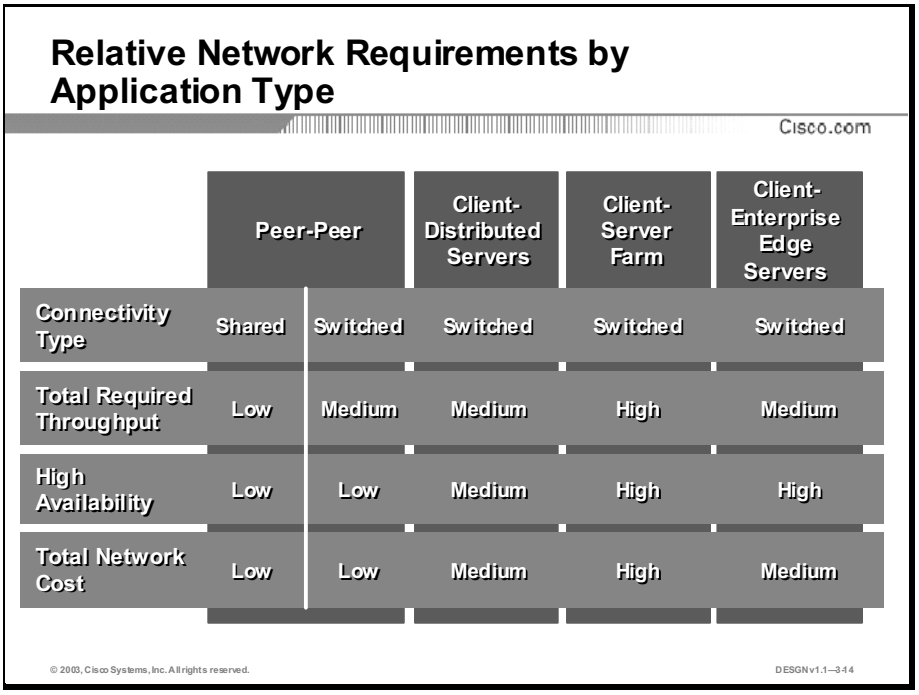

The figure compares the types of applications and their relation to important network parameters.

Connectivity Type

The exploding use of LAN switching at the data link layer has resulted in increased performance of throughput and responsiveness with greater utilization to satisfy the requirements of new organizational applications.

Note:	The shared media for peer-peer communication is suitable when the number of client workstations is very low, for example, in small home offices.

Total Required Throughput

The required throughput varies from application to application. The application that exchanges data between users in a workgroup usually does not require a high-capacity network infrastructure. However, organizational-level applications require a high-capacity link to servers, usually located in a server farm.

Applications located on servers in the Enterprise Edge are normally not highly bandwidth consuming (compared to the applications in the server farm).

High Availability

High availability is a function of the application as well as the whole network between a client workstation and a server located in the network. Although network availability is mainly determined by the network design, the mean time between failures of individual components is a factor. Adding redundancy in the Building Distribution and the Campus Backbone modules is highly desirable.

Total Network Cost

Depending on the application and resultant network infrastructure, the cost varies from low in a peer-peer environment to high in a network built with redundancy in the Campus Backbone and Building Distribution submodules and in the Server Farm module.

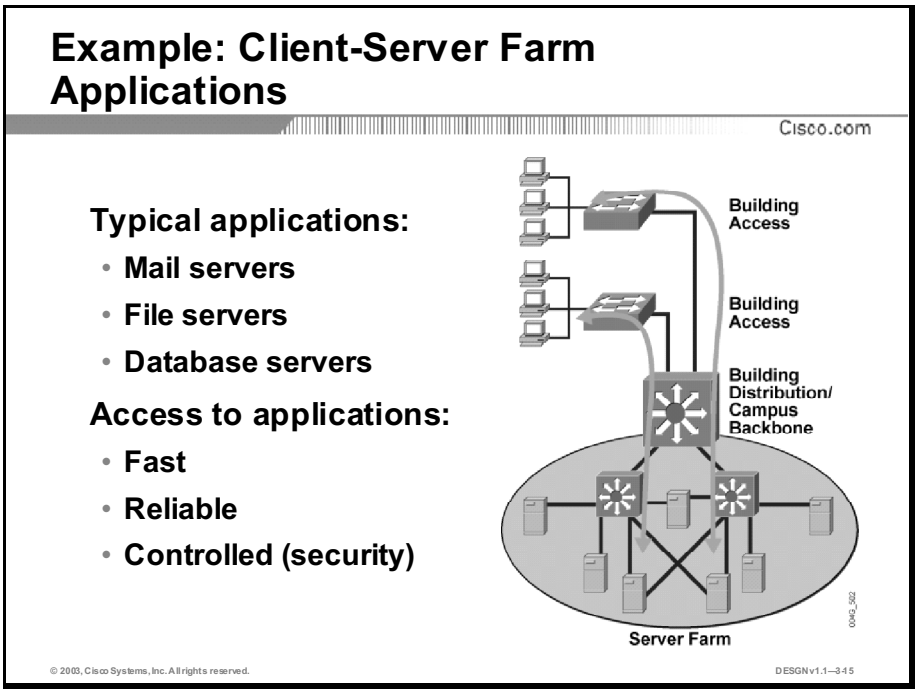

In a large organization, the organizational application traffic may cross more than one wiring closet or LAN to access applications in a server farm. Client-server farm applications apply the 20/80 rule where only 20 percent of the traffic remains on the local LAN segment, and 80 percent leaves the segment to reach centralized servers, the Internet, and so on. Client-server farm applications include:

- Organizational mail servers
- Common file servers
- Common database servers for organizational applications

Large organizations require users to have fast, reliable, and controlled access to the critical applications. To fulfill these demands and keep administrative costs low, the solution is to place the servers in a common server farm. The use of server farms requires a network infrastructure that is highly resilient, redundant, and that provides adequate throughput. Typically, high-end LAN switches with the fastest LAN technologies such as Gigabit Ethernet are deployed.

Example: Client-Distributed Server Applications

Cisco.com

- Servers are located close to clients.
- Servers and clients are in the same LAN.
- Request to servers from nonlocal LANs is rare.

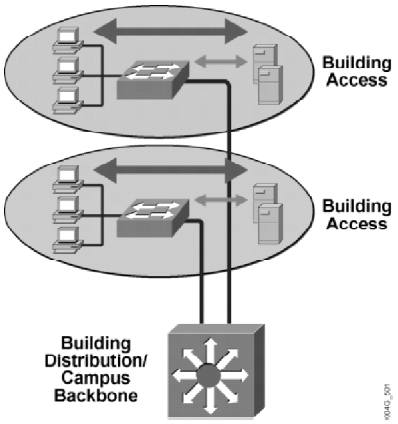

Historically, clients and servers were attached to a network device in a single LAN segment. The 80/20 workgroup rule for client-server applications indicated that 80 percent of the traffic is local and 20 percent leaves the LAN segment.

With increased traffic volumes on the network, an organization may decide to split the network into several isolated segments, with distributed servers for each application. Department administrators manage and control the servers. Most department traffic occurs in the same segment, but some data exchange may traverse the campus backbone. For traffic passing to another segment, the overall bandwidth requirement may not be crucial. For example, Internet access must go through a common segment, which demands less performance than the traffic to the local segment servers.

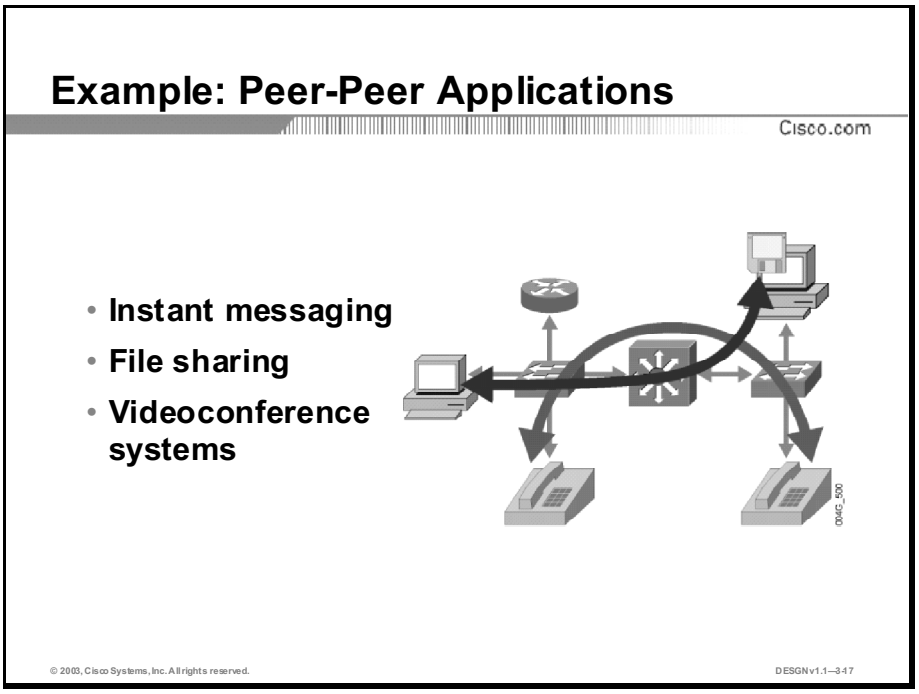

Peer-peer applications include applications in which the majority of network traffic passes from one network edge device to another through the organization's network. Typical peer-peer applications include:

- **Instant messaging:** Two peers establish communication between two end systems. When the connection is established, the conversation is direct.
- **File sharing:** Some operating systems or applications require direct access to data on other workstations.
- **Videoconference systems:** This application's network requirements are very high due to the bandwidth consumption and QoS requirements.

Example: Client-Enterprise Edge Applications

Cisco.com

Typical applications:

- **Internet applications**
  - **Mail servers**
  - **Web servers**
  - **Public Internet servers**
- **E-commerce applications**

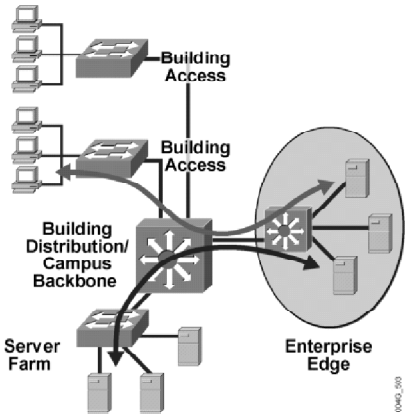

Client-enterprise edge applications use servers on the enterprise edge to exchange data between the organization and its public servers, usually using Internet technology. Examples of these applications include external mail servers and public web servers.

The most important communication issues between the enterprise campus network and the enterprise edge are security and high availability. An application installed on the enterprise edge may be crucial to organizational process flow; therefore, outage may result in increasing process cost.

The organizations that support their partnerships through e-commerce applications also place their e-commerce servers into the enterprise edge. Communication with the servers located on the enterprise campus network is vital because of two-way data replication. As a result, high redundancy and resiliency of the network are important requirements for these applications.

Transmission Media

Transmission media provide physical connectivity between network devices. The most common physical transmission media used in modern networks are twisted-pair cables (copper), wireless (satellite, microwave, 802.11b), and optical cables (fiber). A network designer must be aware of physical media characteristics because they influence the maximum distance and the maximum transmission speed as the main factors in cabling. This topic discusses how the choice of transmission media affects network design.

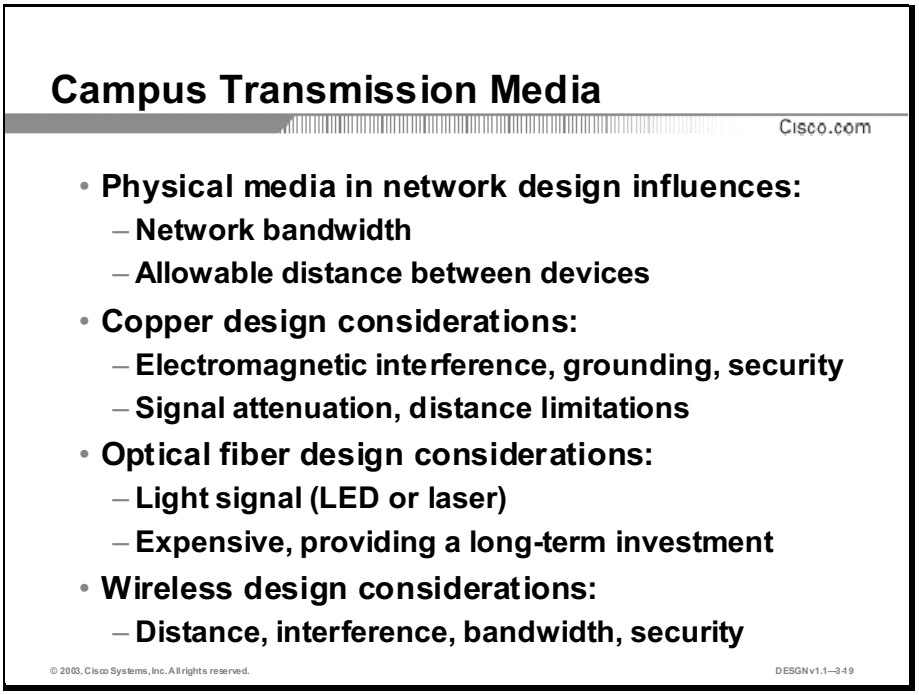

Deciding which type of cable to use is an important consideration when deploying or upgrading a network. Cabling infrastructure represents a long-term investment, which usually has an expected life of 10 years or more.

Copper

The characteristics of twisted-pair cable depend on the quality of material used. Twisted-pair cabling is widely used to interconnect workstations, servers, or other devices from their network interface card (NIC) to the wall outlet and beyond into the Building Access or Building Distribution switches. Category 5 or greater is recommended for speeds of 100 Mbps or higher. Because of the possibility of signal attenuation in the wires, the maximum cable length is usually limited to 100 meters.

A frequent consideration in the cabling design is electromagnetic interference. Because of high susceptibility to interference, twisted-pair cabling is not suitable for environments exposed to electromagnetic influences. Similarly, twisted-pair is not appropriate for environments that may be affected by the interference introduced by the cabling itself.

Note: Electromagnetic interference is also associated with some security issues, particularly eavesdropping if a hacker has access to the cabling infrastructure.

Distances longer than 100 meters may require Long-Reach Ethernet (LRE). Optical cable is also used, especially when immunity to electromagnetic interference is required.

Optical Fiber

The two main types of optical cable are multimode and single-mode optical cable. Multimode fiber is optical fiber that carries multiple light waves or modes concurrently, each at a slightly different reflection angle within the optical fiber core. Typically, LEDs are used with multimode fiber. The typical diameter of a multimode fiber is 50 or 62.5 micrometers.

Single-mode (monomode) fiber is optical fiber that carries a single wave (laser) of light. Typically, lasers are used with single-mode fiber. The typical diameter of a single-mode fiber core is from 2 to 10 micrometers.

Both multimode and single-mode cables have lower loss of signal on the cable than the twisted-pair cable. Therefore, optical cables enable longer distances. Optical fiber cable has very precise production and installation requirements; therefore it has a higher cost than twisted-pair cable.

Optical fiber requires a very precise technique to couple the cables together. Even a small deviation from the ideal position of optical connectors can result in either a loss of signal or a large number of frame losses. In environments where the cable does not consist of a single fiber from point to point, the loss of signal occurs very easily. Careful attention is imperative during optical fiber installation because of the high sensitivity of the traffic to coupling misalignment and bend radius.

Wireless

Wireless LAN technology can either replace a traditional wired network or extend its reach and capabilities. In-building wireless LAN equipment consists of PC client adapters and access points, which perform functions similar to wired networking hubs. To add functionality and range, access points can be incorporated to act as the center of a star topology and function as a bridge to an Ethernet network. The current standard, IEEE 802.11b, discusses speeds of 1 to 11 Mbps; the 802.11a standard goes beyond these limits, enabling faster speeds over wireless.

The inherent nature of wireless is that it does not require wires or lines to accommodate the data, voice, and video pipeline. As such, the system will carry information across geographical areas that are prohibitive in terms of distance, cost, access, or time. It also sidesteps the numerous issues of incumbent local exchange carrier (ILEC) colocation.

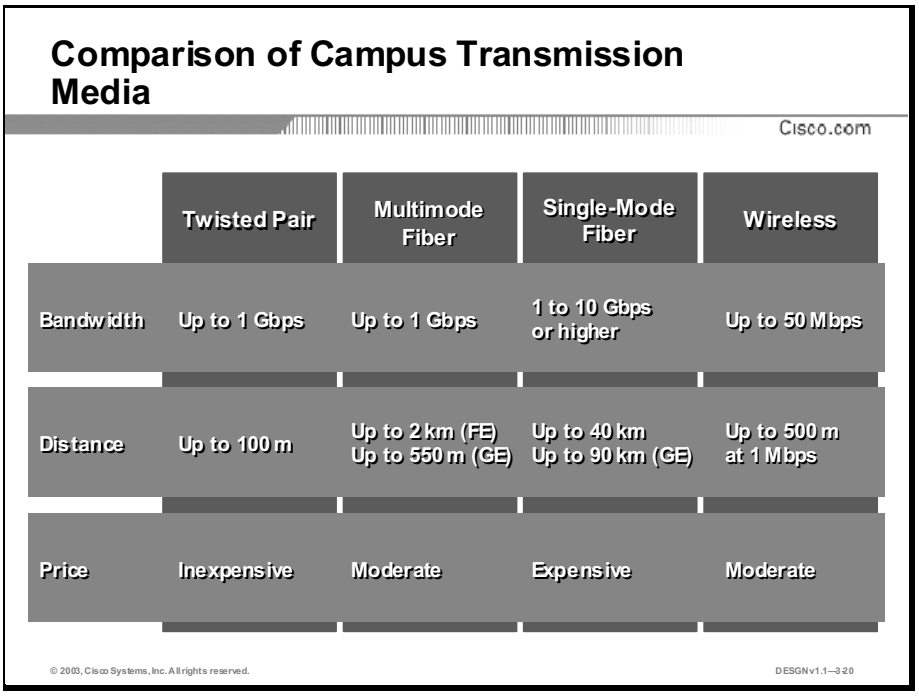

The figure presents the critical parameters that influence network transmission medium selection.

You can make an initial cabling decision based on these primary considerations:

- **Bandwidth:** The required bandwidth in a particular segment of the network or the connection speed between the nodes inside or outside the building
- **Distance:** The distances from network devices (workstations, servers, printers, and IP Phones, and so on) to network nodes and between the network nodes
- **Price:** While the cost of the medium is a clear factor, you must consider the costs of installation

Additional cabling decision criteria include:

- **Deployment area:** The wiring is required for wiring closet only (where users access the network) or for internode or even interbuilding connections
- **Electromagnetic interference:** The electromagnetic interference requirements may influence the selection of the media

Note: The table lists Ethernet as a technology that is compared by bandwidth and range. In addition to Ethernet, Fast Ethernet, and Gigabit Ethernet, you may consider Long-Reach Ethernet (LRE). LRE technology runs on copper cable but allows longer distances than traditional Ethernet. It is used as a distribution technology in broadband building access deployments otherwise constrained by existing cable installations.

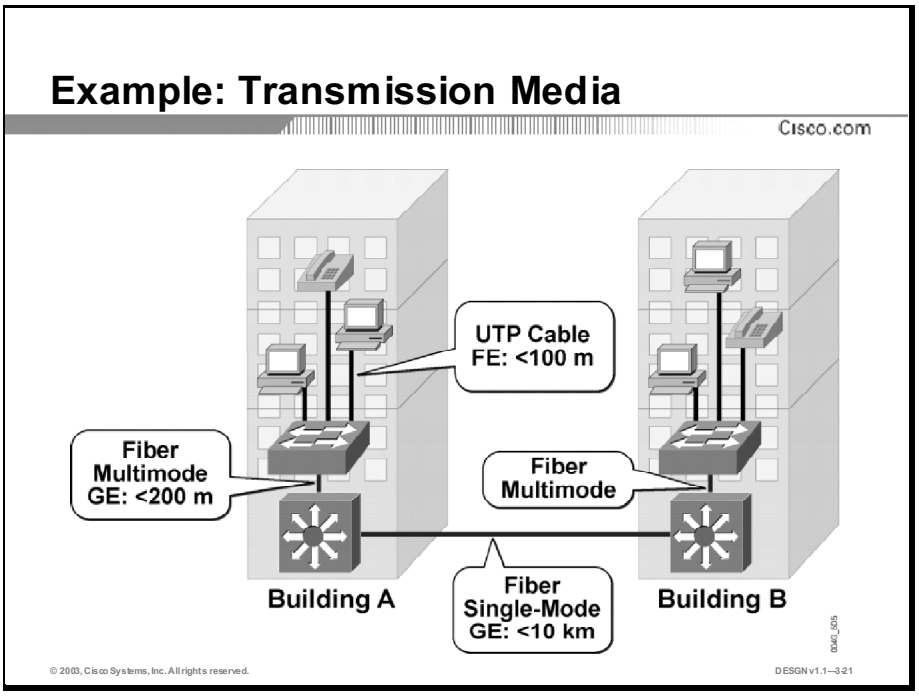

The figure illustrates a typical campus network structure. End devices, such as workstations, IP Phones, and printers, are no more than 100 meters away from the LAN switch. Twisted-pair wiring can easily handle the required distance and speed, while offering a reasonable price/performance ratio.

Optical fiber cables handle the higher speeds and distances required among switch devices. Inside the building, multimode optical cable is usually satisfactory. For interbuilding communication, organizations use multimode or single-mode optical cable depending on distance. If the distances are short (up to 500 meters), multimode fiber is a more cost-effective solution for speeds up to 1 Gbps.

If an organization's requirements are for longer distances or for future higher speeds (for example, 10 Gbps), they need to install single-mode fiber.

Segmentation Technologies

The increasing power of desktop processors and the requirements of client-server and multimedia applications drive the need for greater bandwidth, typically provided by switched technology. This topic describes segmentation technologies in network design.

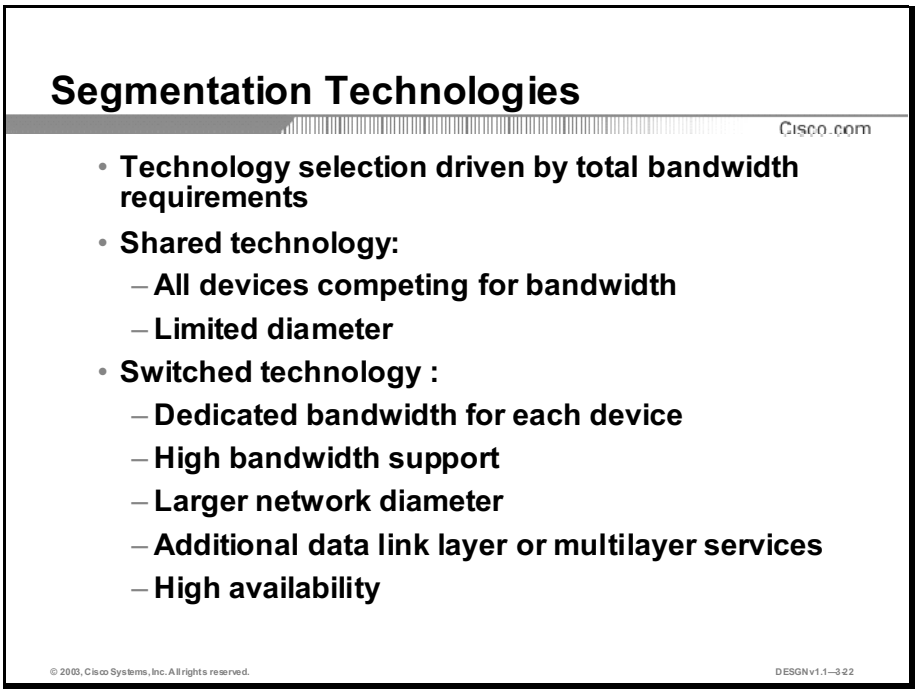

Shared Technology

Shared technology is based on all devices sharing a segment's bandwidth. A device must gain media access before forwarding its data frame to all other participants on the segment. All other devices must wait until completely receiving the data frame before attempting to gain access for their transmission. The major drawbacks of shared technology are that all network devices must compete for bandwidth and limited diameter of the network due to electrical and optical characteristics that limit the length of the segment.

Switched Technology

Switched technology provides dedicated network bandwidth for each device on the network.

Additional benefits of switched technology include:

- **High bandwidth support:** Switches can transfer multiple simultaneous frame flows.
- **Larger network diameter:** There is a single device on a physical segment, minimizing the need for collision detection.
- **Additional services:** LAN switches support frame and packet switching at the data link layer, and modern switches perform several functions at Layer 3 and at higher OSI layers.
- **High availability:** Switches can be interconnected with multiple links without creating loops in the network.

These benefits have essentially eliminated shared technologies, and the majority of new networks use only switched technologies. Shared technologies are only present in some parts of existing networks and in smaller home offices.

Note: A Layer 3-capable device separates network segments from each other. In a traditional network, the Layer 3 (network layer) device was a router. In a modern network, the preference is to deploy a multilayer switch, which performs at Layer 3 and above.

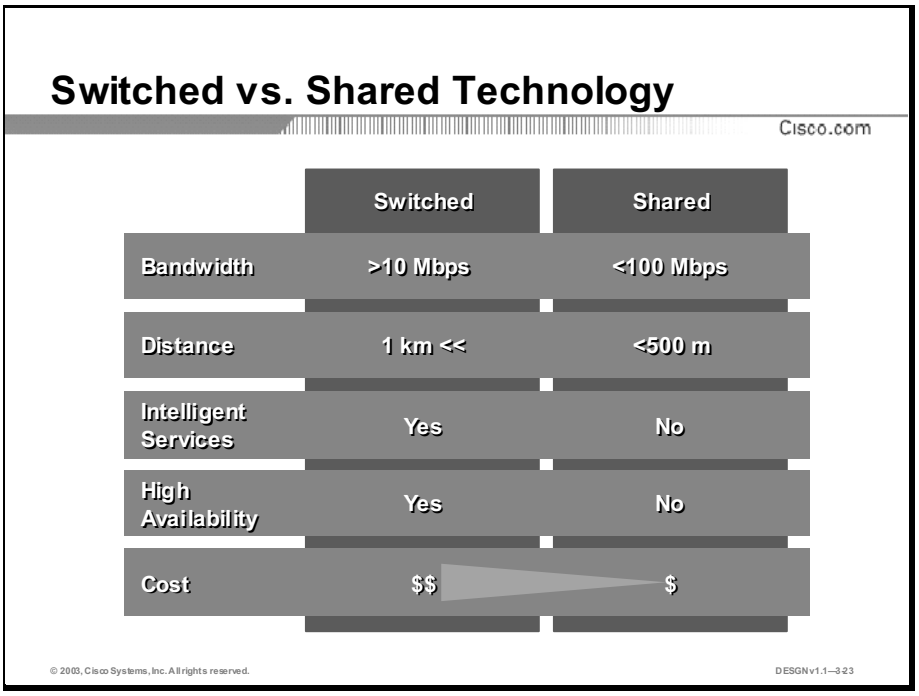

The figure presents some of the primary differences and benefits of switched technology as compared to shared technology.

Bandwidth

Bandwidth in shared technology is limited to the Fast Ethernet speed on a network segment. Switched technology supports speeds from 10 Mbps to more than 1Gbps, and enables multiple ports to forward frames through the switch simultaneously.

Distance

The distance or diameter of a segment is almost unlimited in a switched environment. The modern LAN switch stores all or part of a frame before forwarding. In a shared environment, every station on the segment must compete for resources and be able to detect if two or more network stations are transmitting at the same time. The Ethernet standard for shared technology defines how long the sending device must possess the bus before actually sending the data. Because of this time limitation, the range of the segment is defined and, in the best scenario, never reaches more than 500 meters.

Intelligent Services

Switched networks are required to support intelligent network services, such as QoS, security, and management. Shared technology does not have the capability to support these new features of the network, which become important with the increasing number of organizational client-server and multimedia applications.

High Availability

Many of the organizational processes that run on the network infrastructure are critical for the organization's success. As a consequence, high availability has become increasingly important. Shared networks do not offer the required capability, while switched environments do.

Cost

With all the benefits that LAN switches offer, the cost per port is expected to be higher on switches than on hubs. With wide deployment and availability, the price per port on LAN switches is almost the same as it is for hubs or repeaters.

Switching Design Considerations

To select the right type of switch for each network module requires that you fully understand the network topology and user needs. This topic describes the design considerations to help you choose between data link layer and multilayer switching.

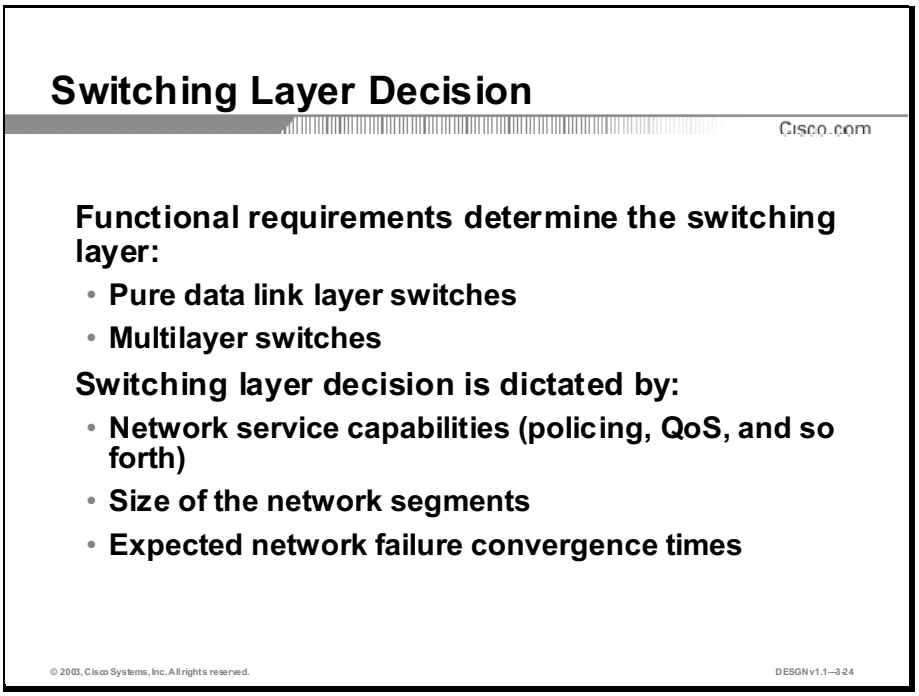

The difference between data link layer and multilayer switching is the type of information (inside the frame) used to determine the correct output interface. Data link layer switching forwards frames based on data link layer information (MAC address), while multilayer switching forwards frames based on network layer information (IP address).

To determine which LAN switch and features to deploy in an organizational network, answer these questions:

- **Network service capabilities (QoS, and so forth):** What network services does the organization require?
- **Size of the network segments:** Based on traffic characteristics, how will the network be segmented?
- **Convergence times:** What level of high availability is required and what is the maximum time for possible network outages?

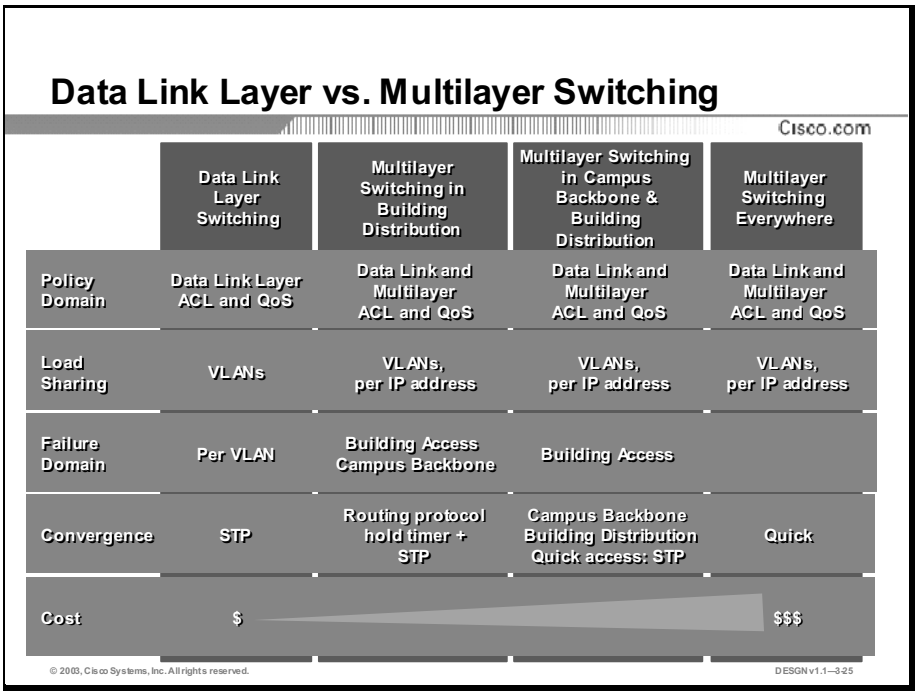

The figure compares data link layer and multilayer switching features in campus networks.

Policy Domain

A network policy is a formal set of statements that define how network resources are allocated among devices, for example, the time of day and client authorization priorities. You can apply policies to individual users, groups, or entire departments, as well as to selected hosts or applications. Network managers store policies in a policy repository or on a device. The device then applies the configured policies to network resources.

The policy domain is the scope of the network that is affected by a certain policy. The size of a policy domain depends on the type of the switching and on the mechanisms for policy implementation. With data link layer switching the policy domain overlaps with the boundaries of the switching domain. Multilayer switching offers more flexibility. In data link layer switching, you can only apply access control lists (ACL) and various QoS mechanisms to switched ports and MAC addresses. In multilayer switching, you can extend the ACL and QoS mechanisms to IP addresses applications (for example, TCP/User Datagram Protocol [UDP] ports).

Load Sharing

You can use multiple links for redundancy, traffic load sharing, or both. Data link layer switches only offer load sharing by distributing VLANs across different uplink ports. Multilayer switches, however, can perform load sharing between ports based on IP destinations.

Failure Domain

A failure domain defines the scope of the network affected by outages or misconfigurations usually associated by a data-link switch domain bounded by a VLAN or the network layer. Other failure domains are collision (bound by a switch or hub) and broadcast (bound by a LAN or VLAN).

A failure domain is:

- Bounded by multilayer switching
- Bounded by the VLAN when data link layer switching is deployed in an entire campus

Convergence

Loop prevention mechanisms in the data link layer topology cause the Spanning Tree Protocol (STP) to take between 30 and 50 seconds to converge. You should implement STP enhancements such as Portfast, UplinkFast, BackboneFast, and so on. To eliminate STP convergence issues in the Campus Backbone submodule, all the links connecting backbone switches should be routed links, not VLAN trunks. This approach also constrains the data-link broadcast and failure domains.

When multilayer switching is deployed everywhere, convergence takes place within seconds, depending on the routing protocol implemented. All the devices detect their connected link failure immediately and act upon it promptly, sending respective routing updates.

In a mixed data link layer/multilayer environment, you will consider both the multilayer factors (routing protocol timers such as holdtime and neighbor loss detection) and the STP convergence.

Using multilayer switching in a structured design will reduce the scope of spanning-tree domains. It is common to use a routing protocol, such as Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF), to handle load balancing, redundancy, and recovery in the backbone.

Cost

The cost of deploying multilayer switching in comparison to data link layer switching increases with the capabilities of its intelligent network services and device performance.

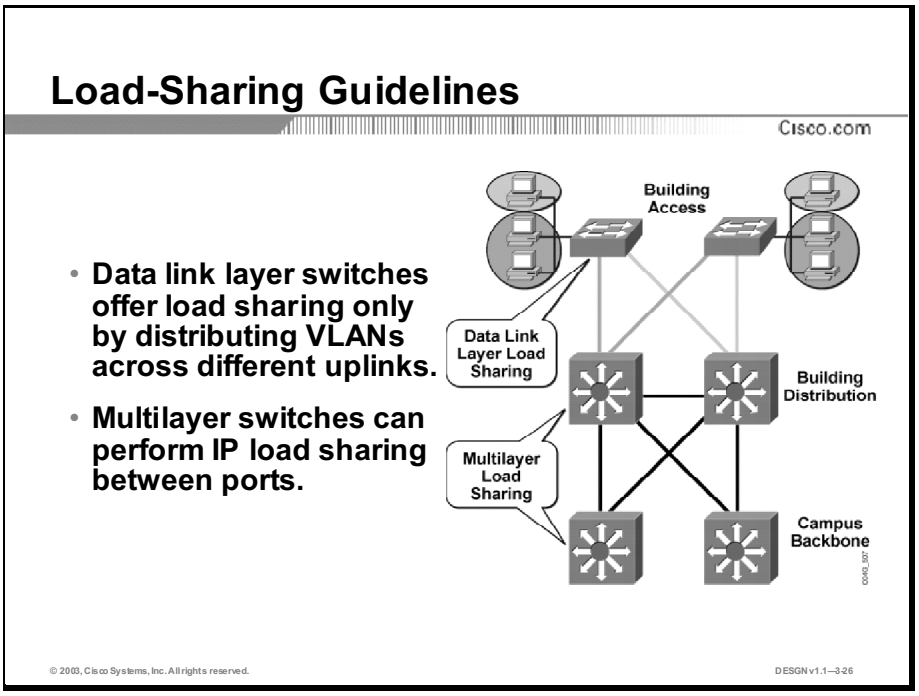

Data Link Layer Load-Sharing Guidelines

Data link layer switches are only aware of MAC addresses; they cannot perform any intelligent load sharing. In an environment with multiple VLANs per access switch and more than one connection to the uplink switch, the solution is to put all uplink connections into trunks. Each trunk carries all the VLANs. However, without an additional configuration, STP will disable all nonprimary uplink ports. This configuration may result in bandwidth shortage because the traffic for all VLANs will pass the same link. To overcome the problem, modify the STP parameters to carry some VLANs across one uplink and the rest of the VLANs across the other uplink. For example, configure one uplink to carry the VLANs with odd numbers and the other uplink to carry the VLANs with even numbers.

Multilayer Load-Sharing Guidelines

Switches that support multilayer capability can perform load sharing based on IP addresses. Most multilayer devices with load-sharing capability can balance the load per packet or per destination-source IP pair. The advantage of multilayer IP load sharing is that links are used more proportionately than with data link layer load sharing, which is based only on VLANs.

Multilayer load sharing is appropriate when the traffic load varies between VLANs. Multilayer switching supports dynamic adaptation to link utilization using routing protocol design. Multilayer switches support data link layer load sharing, and can apply per-VLAN load sharing when connected to other data link layer switches.

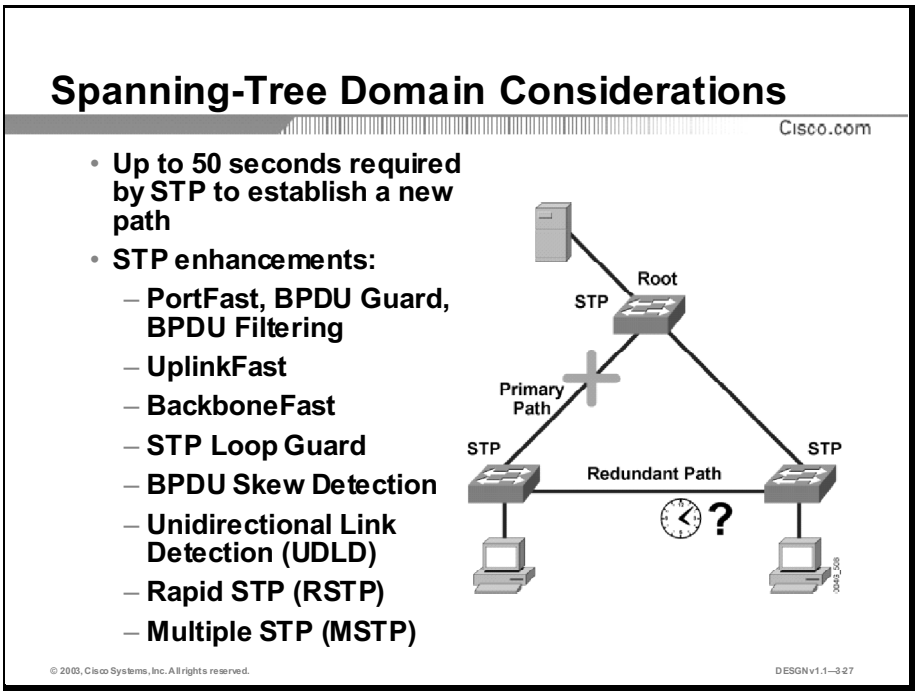

STP selects a root switch (root bridge according to IEEE 802.1D standard terminology) and redundant paths. After the switch comes online, it takes up to 50 seconds before the root switch and redundant links are detected. During this period, the port goes through listening and learning states and then to either the forwarding or the blocking state. No ordinary traffic can travel through the network at this time. Each state (listening, learning) takes approximately 20 seconds with default STP values.

STP Enhancements

Cisco Systems and standardization efforts have introduced new enhancements to STP:

- **PortFast:** Used for ports where end-user stations, servers, or both are directly connected. There is no delay in passing traffic because the switch puts the port to the forward state (skipping the listening and learning states) immediately. Associated with the PortFast feature are two additional measures that prevent potential STP loops:
  - **Bridge Protocol Data Unit (BPDU) Guard:** PortFast transitions the port into STP forwarding mode immediately upon linkup. Because the port still participates in STP, potential STP loops exist if a device attached to that port also runs STP. The BPDU Guard feature enforces the STP domain borders and keeps the active topology predictable. If a BPDU is received on the port, the port is transitioned into Errdisable state and an error message is reported.
  - **BPDU Filtering:** This feature blocks PortFast-enabled nontrunk ports from transmitting BPDUs. Essentially, spanning tree will not run on these ports.
- **UplinkFast:** If the link to the root switch goes down and the link is directly connected to the switch, UplinkFast enables the switch to put a redundant path (port) into the active state within a second.
- **BackboneFast:** If a link on the way to the root switch fails, but is not directly connected to the switch, BackboneFast reduces the convergence time from 50 seconds to 20 to 30 seconds. Enable BackboneFast on all the switches in the STP domain.

The following features prevent error conditions from causing unpredictable STP topology changes that could lead to STP loops:

- **STP Loop Guard:** When one blocking port in a physically redundant topology stops receiving BPDUs, the STP moves the port to forwarding state, creating a loop. The STP Loop Guard feature makes an additional check. If BPDUs are no longer received on a blocking port and the feature is enabled, that port is moved into the STP loop-inconsistent blocking state instead of moving to the listening / learning / forwarding state. This feature avoids loops in the network due to unidirectional failures or other software failures.
- **BPDU Skew Detection:** This feature allows the switch to keep track of late-arriving BPDUs (by default BPDUs are sent every 2 seconds) and notify the administrator by means of syslog messages. For every port on which a BPDU has ever arrived late (skewed), skew detection will generate a report. Report messages are rate-limited (one message every 60 seconds) to protect the CPU.
- **UniDirectional Link Detection (UDLD):** If the STP process running on the switch with a blocking port stops receiving BPDUs from its upstream (designated) switch on this port, STP will eventually age out the STP information for this port and move it to the forwarding state, creating an STP loop. The UDLD is a protocol that works with the L1 mechanisms to determine the physical status of a link. If the port does not see its own device or port ID in the incoming UDLD packets for a specific duration of time, the link is considered unidirectional from the L2 perspective. As soon as UDLD detects the unidirectional link, the respective port is disabled and the error message is generated.

The following standard enhancements allow STP to compare favorably with (or even to exceed) the convergence of routing protocols:

- **Rapid Spanning Tree Protocol (RSTP)** (defined in IEEE 802.1w**)**: RSTP provides rapid convergence of the spanning tree by assigning port roles and by determining the active topology. RSTP builds upon the IEEE 802.1D STP to select the switch with the highest switch priority as the root switch and then assigns the port roles (root, designated, alternate, backup, and disabled) to individual ports. These roles assist with rapid STP convergence, which can be extremely fast (within a second).
- **Multiple Spanning Tree Protocol (MSTP)** (defined in IEEE 802.1S): MSTP uses RSTP for rapid convergence. It enables several (topologically identical) VLANs to be grouped into a single spanning-tree instance, with each instance having a spanning-tree topology independent of other spanning-tree instances. This architecture provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of spanning-tree instances required to support a large number of VLANs.

Summary

This topic summarizes the key points discussed in this lesson.

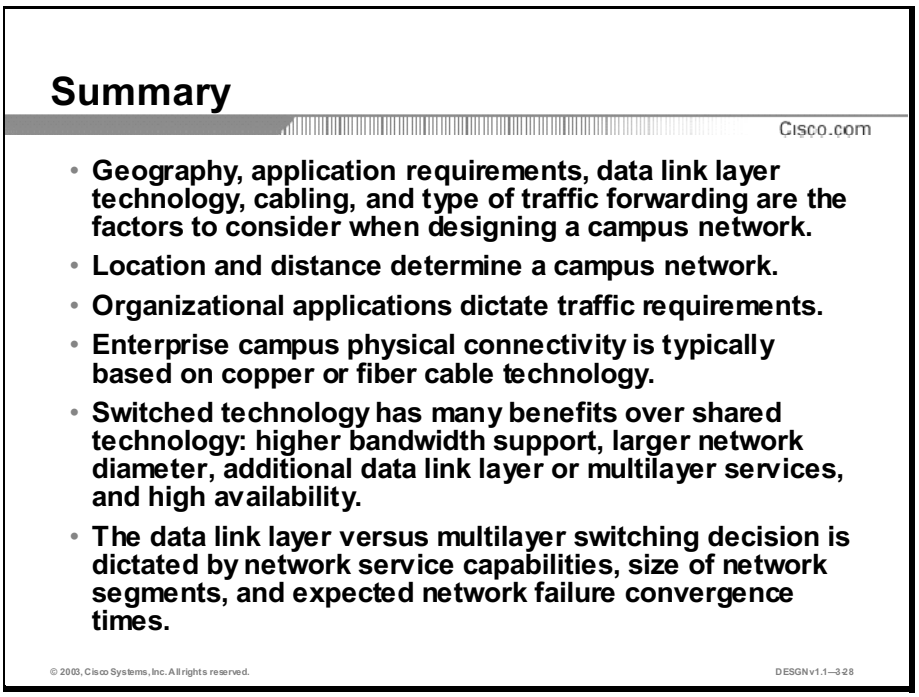

References

For additional information, refer to these resources:

- *Introduction to Gigabit Ethernet,* http://www.cisco.com/warp/public/cc/techno/media/lan/gig/tech/gigbt_tc.htm
- *Gigabit Networking Gigabit Ethernet Solutions,* http://www.cisco.com/warp/partner/synchronicd/cc/techno/lnty/etty/ggetty/tech/gesol_wp.htm
- *Gigabit Campus Network Design—Principles and Architecture,* http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which four factors affect the campus network design? (Choose four.)

A) application requirements

B) PC operating system used in an organization

C) transmission media characteristics

D) distribution of network nodes

E) remote-site connectivity

F) network operating system used in an organization

Q2) Organization XYZ owns several buildings at the organizational park. They want to connect the buildings into one organizational network. What geographical structure will they choose for a design?

A) intrabuilding

B) interbuilding

C) inter remote-building

D) intracampus

Q3) Two workstations located in separated VLANs must exchange data directly. The workstations are physically connected in different buildings. What type of communication will they use?

A) Enterprise Edge application

B) client-server

C) peer-peer

D) workgroup-workgroup

Q4) Which type of cable is the optimal solution for connecting a Building Access switch to a Building Distribution switch when the distance is 120 m?

A) power cable

B) UTP

C) multimode optical cable

D) single-mode optical cable

Q5) Why is LAN switching technology better than shared LAN technology?

A) Shared LANs do not consume all the available bandwidth.

B) Switched LANs allow more than one port to communicate simultaneously.

C) Switched LANs forward a unicast frame to all ports simultaneously.

D) Switched LANs offer no benefits compared to a hub.

Q6) Some users in a department use an organizational application that generates an increased number of broadcast frames, which results in up to 10-Mbps bandwidth utilization. Which solution is the best campus design choice?

A) Provide 100-Mbps connections or higher to all users in a domain.

B) Limit the number of broadcast frames in a domain for all department users.

C) Optimize the application.

D) Put the application users into a separate broadcast domain.

Quiz Answer Key

Q1) A, C, D, E

Relates to: Designing an Enterprise Campus

Q2) B

Relates to: Network Geography

Q3) C

Relates to: Network Applications

Q4) C

Relates to: Transmission Media

Q5) B

Relates to: Segmentation Technologies

Q6) D

Relates to: Switching Design Considerations

Selecting Campus Design Models

Overview

An enterprise campus network is constructed with network campus building blocks, where each block has specific goals to achieve. As a designer, you need to understand the organizational traffic to ensure that you select the appropriate features and building components.

The Enterprise Campus functional area of the Enterprise Composite Network Model is composed of the Campus Infrastructure, Server Farm, and Edge Distribution modules. Users access the network through Building Access submodules, which are interconnected via Building Distribution submodules. These interconnected submodules scale from individual buildings to a campus through the Campus Backbone submodule.

This lesson discusses advanced network traffic considerations and building design using data link layer and multilayer switching in the Building Access, Building Distribution, and Campus Backbone submodules. The lesson also introduces Server Farm module design, for which dual homing and redundancy are required. The lesson continues with design guidelines for connectivity to the Edge Distribution module.

Relevance

The design models presented will help you apply specific design requirements to effective, proven design models.

Objectives

Upon completing this lesson, you will be able to position switches appropriately in campus networks. This includes being able to meet these objectives:

- Describe basic design considerations for the campus network
- Describe basic design considerations for the Building Access and Building Distribution submodules
- Describe basic design considerations for the Campus Backbone submodule
- Describe basic design considerations for the Server Farm module

- Describe basic design considerations for the Enterprise Edge functional area

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, functions, and general switching theory

Outline

The outline lists the topics included in this lesson.

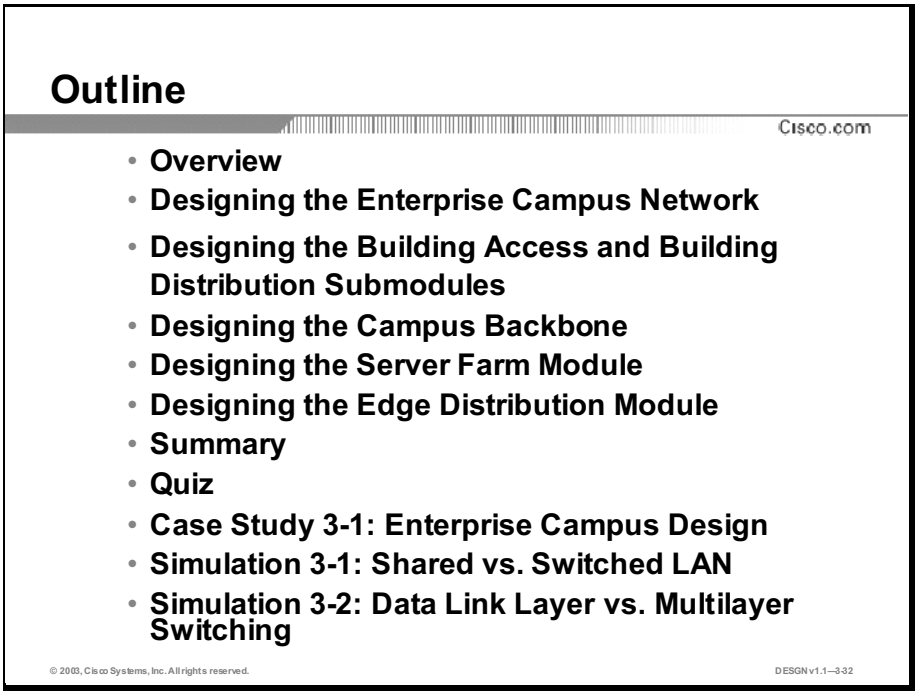

Designing the Enterprise Campus Network

The Enterprise Composite Network Model design approach results in a campus network that is scalable, manageable, and highly available. Designing an enterprise campus network involves balancing performance and cost, while providing scalability and high availability. This topic describes considerations for enterprise campus network design.

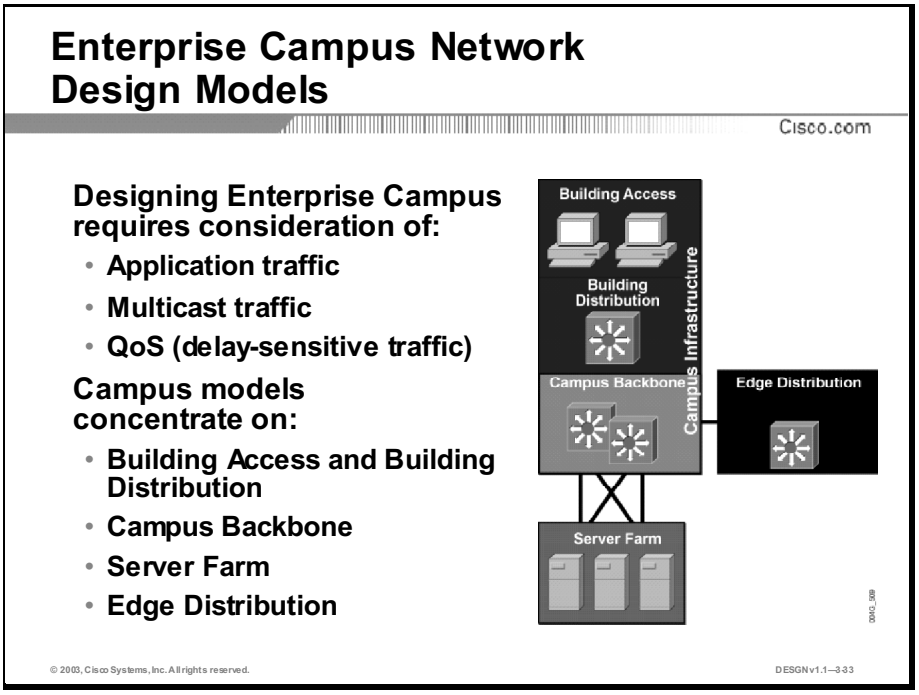

Designing an enterprise campus network requires that you consider:

- **Application traffic:** Identify the organizational traffic flows. This includes the type of the traffic, its bandwidth requirements, and traffic patterns.
- **Multicast traffic:** Identify the correct features that constrain multicast streams to the relevant ports. Multicast traffic, if not considered, may use a great amount of bandwidth.
- **QoS (delay-sensitive traffic):** Identify and incorporate the appropriate QoS mechanisms to manage requirements for delay and delay variations.

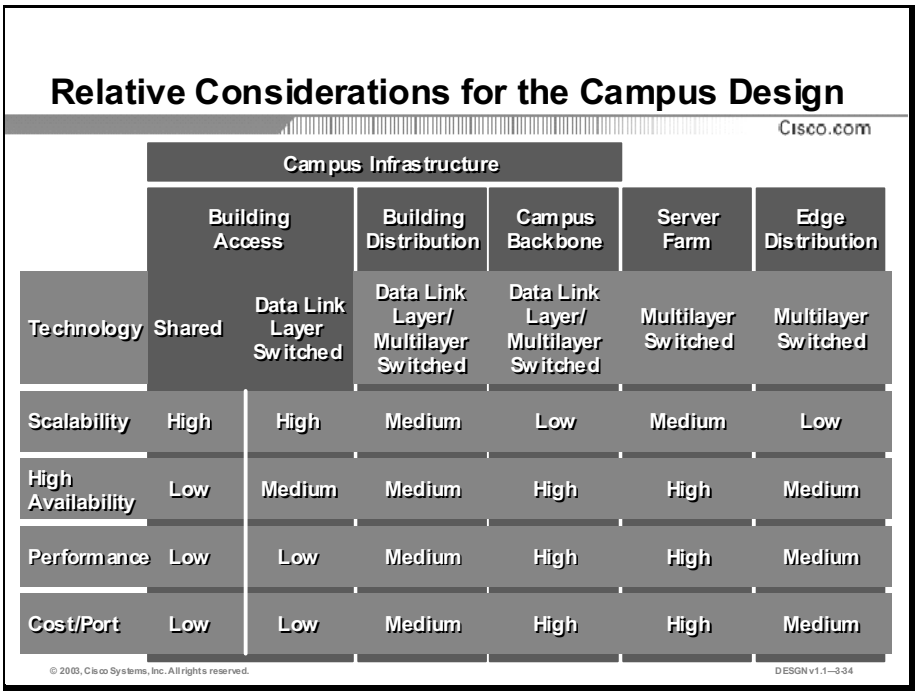

Each individual module of the enterprise campus has different requirements. The figure illustrates that modules closer to the user require different degrees of scalability. Therefore, you must consider options for future expansion of the campus network. Adding new workstations to a network should not result in either high additional costs or performance degradations.

End users may not require high performance and high availability, but these features are crucial in the Campus Backbone and Server Farm modules.

With increased performance and availability, the price per port increases. The Campus Backbone and Server Farm modules require a guarantee of higher throughput so that they can handle all traffic flows without introducing additional delays or drops to the network traffic.

The Edge Distribution module does not require the same performance as the Campus Backbone. However, it may require features such as security and policy enforcement that increase the overall cost.

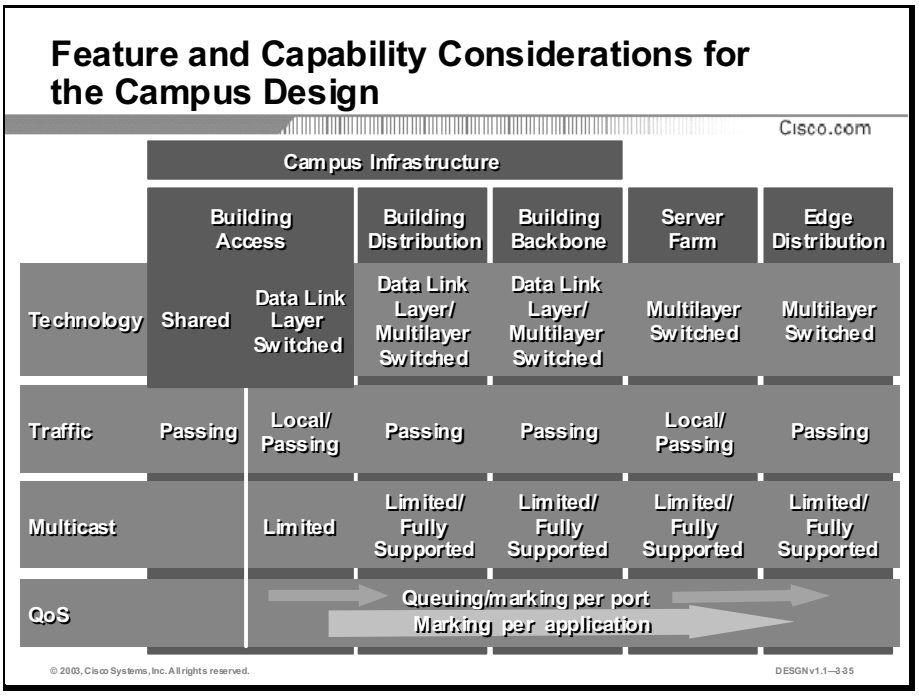

You can build an enterprise campus network on either a shared or switched foundation. In the Building Access submodule, you can connect workstations with low demand via shared technology, although shared technology is most suitable for some small (home) offices of a few devices without any special bandwidth requirements. Where higher speeds are required, LAN switching is the only viable option.

Consideration of the organizational applications and traffic is required to ensure the selection of the appropriate equipment for the individual modules. Traffic patterns, multicast traffic, and QoS are the most important network issues when designing the campus network.

Data link layer switches usually support multicast and QoS features but with limited capability. A multilayer switch with enhanced features is required if the data link layer switches do not offer satisfactory capabilities for the specific module.

A data link layer multicast aware switch that works closely with the multilayer device (router) can distinguish between the hosts that belong to the multicast stream and those that do not. Thus, the data link layer switch can forward the multicast stream only to selected hosts.

Data link layer QoS is usually limited to port marking capability and queuing only on uplink trunk ports, especially on low-end switches. The low-end switches are usually incapable of marking or queuing based on Layer 3 or higher parameters of packets. On recent platforms, Cisco added support for data link, network and transport layer class of service (CoS) and type-of-service packet marking and policing.

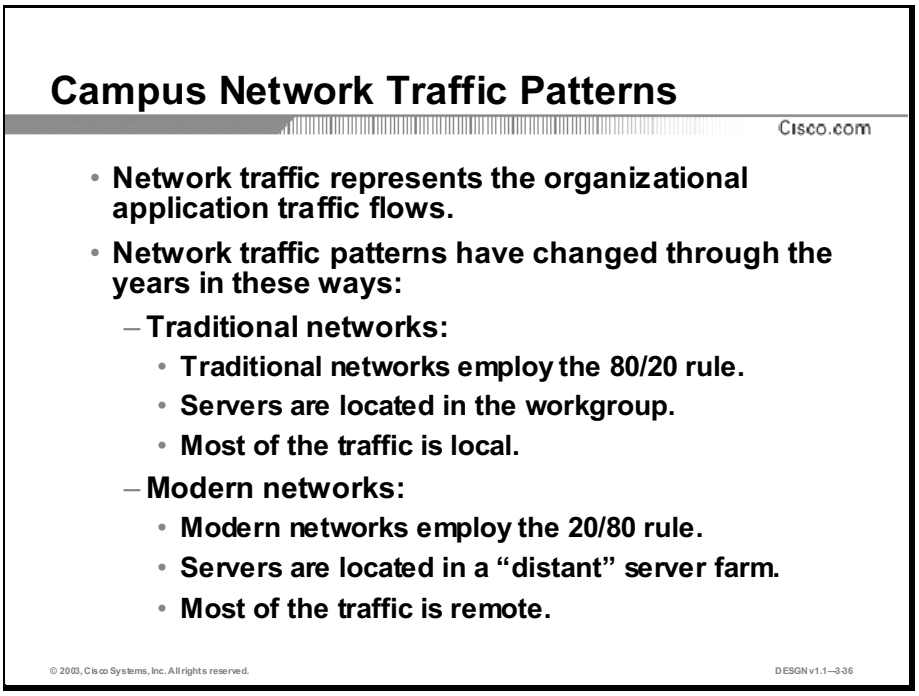

80/20 Rule in the Campus

When designing a switched campus, network designers ensure that each switched segment corresponds to a workgroup. By placing the workgroup server in the same segment as its clients, most of the traffic can be contained. Known as the 80/20 rule, this design principle refers to the goal of keeping at least 80 percent of the traffic within the local segment.

The campus-wide VLAN model relates to the 80/20 rule. If 80 percent of the traffic is within a workgroup or VLAN, the packets are locally switched at the data link layer. The 80/20 rule underlies traditional network design models.

20/80 Rule in the Campus

Many new and existing applications now use centralized data storage and retrieval. The traffic pattern is moving toward the 20/80 rule, where only 20 percent of traffic is local to the workgroup LAN and 80 percent of the traffic leaves the workgroup.

In a traditional network design, only a small amount of traffic passes through the multilayer devices (routers). Modern enterprise networks deploy servers located either in server farms or at the enterprise edge. With an increasing amount of traffic from clients to distant centralized servers, performance requirements in the Building Distribution and Campus Backbone submodules are higher.

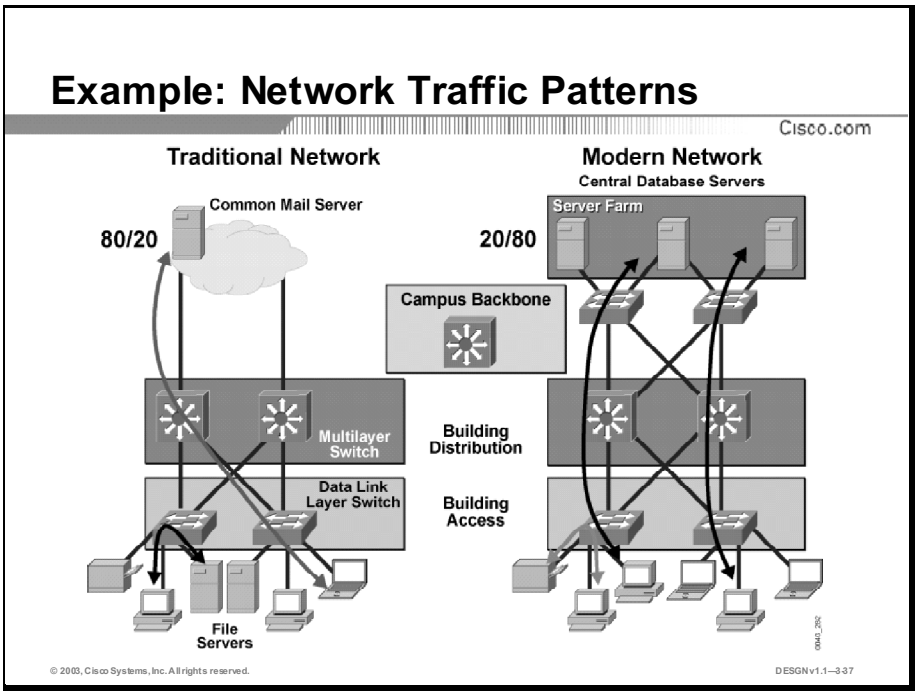

Example: 80/20 Campus Design

Company A, shown on the left side of the figure, has several independent departments. Each department has its own VLAN where the servers and printers are located. File transfers from other department servers or workstations are only occasionally necessary. Shared traffic must pass the Building Distribution submodule, represented by the L3 switch. The only common resource that the departments use is the mail server, located in the Campus Backbone.

Example: 20/80 Campus Design

Company B, shown on the right side of the figure, also has several departments, although they use common resources. They use file servers from their own department as well as services from a common data storage such as an Oracle database. This configuration requires a higher performance multilayer switch in the Building Distribution submodule. The Building Access data link layer switch associates and limits users to their VLANs. The servers on the other side of the network are organized into groups connected to data link layer switches. Building Distribution and Campus Backbone switches in the middle enable fast, reliable, and redundant communication between the groups on both sides of the network.

The figure shows that the majority of the communication takes place between servers and users, with only a small amount of traffic being switched inside the group.

Multicast Traffic in the Campus Network

Cisco.com

- IP multicast delivers a traffic stream to multiple destinations.
- Considerations on LAN switches:
  - Data link layer switches "flood" a multicast frame to every port.
  - Static entries specify which ports should receive the multicast traffic.

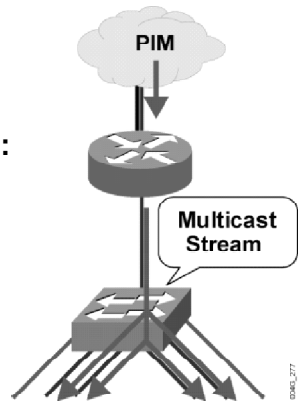

IP multicast is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to potentially thousands of corporate recipients. Videoconferencing; corporate communications; distance learning; and distribution of software, stock quotes, and news are the applications that take advantage of the multicast traffic stream. IP multicast delivers source traffic to multiple receivers.

IP multicast is based on the concept of a multicast group. Any group of receivers can express an interest in receiving a particular data stream. The group does not need to have any physical or departmental boundaries, so the hosts can be located anywhere on the corporate network. Hosts that are interested in receiving data flowing to a particular group must join the group by using the Internet Group Management Protocol (IGMP).

The figure illustrates a typical situation with IP multicast. Multicast-enabled routers ensure the proper delivery of traffic by using one of the multicast routing protocols such as Protocol Independent Multicast (PIM). The router to the switch port forwards the incoming multicast stream.

Note: Support for broadcast and multicast suppression is available on several switched platforms. The threshold can be set to any value between 0 and 100 percent (or as a number of packets when packet-based suppression is turned on). When the threshold is exceeded on the port, the switch suppresses further activity on the port for the remainder of a 1-second period.

Cisco switches support several methods to deal efficiently with multicast in an L2 switching environment. The most common are:

- **Cisco Group Management Protocol (CGMP):** Based on the communication between the multicast router and the switch, the multicast receiver registration (using the IGMP) is accepted by the router and communicated via CGMP to the switch from the router; the switch adjusts its forwarding table accordingly. CGMP is a Cisco proprietary solution implemented on all Cisco LAN switches.

- **IGMP snooping:** The switch intercepts multicast receiver registrations and adjusts the forwarding table accordingly. IGMP snooping requires that the switch is L3 aware because IGMP is a network layer protocol. Typically, IGMP packet recognition is hardware-assisted.

Note: Additional methods that address the problem of multicast frames in a switched environment include the GARP Multicast Registration Protocol (GMRP) and the Router-Port Group Management Protocol (RGMP). GMRP, used between the switch and the host, is not yet widely available. RGMP is a Cisco solution for router-only multicast interconnects in a switched environment.

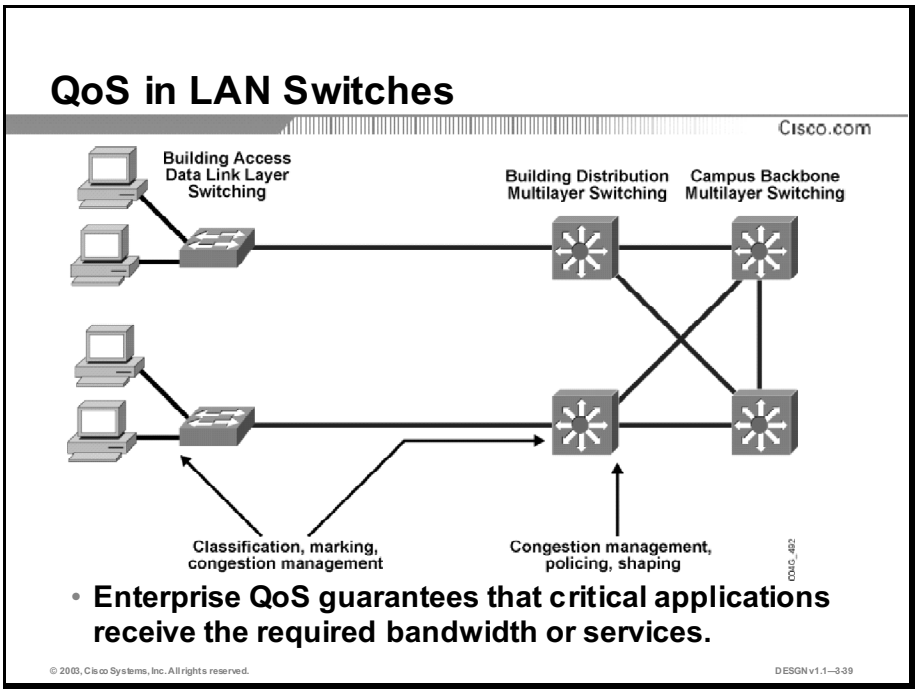

A campus network transports many types of applications and data, including high-quality video and delay-sensitive data, such as real-time voice. Bandwidth-intensive applications stretch network capabilities and resources but may also enhance many business processes. Networks must provide secure, predictable, measurable, and sometimes guaranteed services. Achieving the required QoS by managing delay, delay variation (jitter), bandwidth, and packet loss parameters on a network can be the key to a successful end-to-end business solution. QoS is the set of techniques used to manage network resources.

Bandwidth demand on the network may exceed the available bandwidth. To guarantee bandwidth to business-critical applications, the network needs to provide QoS. Stringent requirements for low delay and jitter further drive the need for QoS in LAN switches.

Most networks or individual network elements are oversubscribed. QoS is most often required on uplinks either from the Building Access submodule to the Building Distribution submodule or from the Building Distribution submodule to the Campus Backbone submodule. The sum of bandwidth of all ports on a switch where end devices are connected is usually greater than that of the uplink port. When the ports are fully used, congestion on the uplink port is unavoidable.

A network operator configures bandwidth management with QoS mechanisms on the Building Access, Building Distribution, or Campus Backbone switches, depending on traffic flow and oversubscription of the uplinks.

QoS Categorization

QoS, as implemented on LAN switches, can be categorized into these four areas:

- **Classification and marking:** Packet classification features allow you to partition traffic into multiple priority levels or classes of service. These features inspect the information in the frame and network header and determine the frame's priority. Marking is the term given to the process of changing the priority level (CoS) setting of a frame.

- **Scheduling:** Scheduling is the process that determines the order in which queues are serviced. CoS is used on data link layer switches to assist in the queuing process. Based on network, transport, and higher layer information, multilayer switches can also provide QoS scheduling. Layer 3 IP QoS queue selection uses the IP differential services code point (DSCP) or IP precedence field of the IP packet.
- **Congestion management:** Often, a network interface is congested (even at high speeds, transient congestion is observed), and queuing techniques are necessary to ensure that the critical applications receive the forwarding treatment necessary. For example, real-time applications such as VoIP and stock trading may need to be forwarded with the lowest latency and jitter.
- **Policing and shaping:** Policing is the process of reducing a stream of data to a predetermined rate or level. Unlike traffic shaping, where the frames can be stored in small buffers for a short period, policing drops or lowers the priority of the frame that is out of the profile.

When configuring QoS features, select the specific network traffic, prioritize it according to its relative importance, and use congestion management techniques to provide preferential treatment.

Building Access switches provide QoS classification and marking based on the input port because they are data link layer switches with no network or transport layer awareness. You can define traffic from a particular host as high-priority traffic on the uplink port. The scheduling mechanism on the output port of a Building Access switch ensures that traffic from such ports is served first. The proper marking of input traffic ensures the expected service behavior when traffic passes from Building Distribution to Campus Backbone switches.

Building Distribution and Campus Backbone switches are typically multilayer aware and can provide QoS selectively, on a port basis and according to higher-layer parameters, such as IP addresses, port numbers, or QoS bits in the IP packet. These switches differentiate the traffic based on the application and make QoS classification more selective. The policing for certain traffic is usually implemented on the Building Distribution switches.

Example: QoS for Voice Traffic Across the Switch

Cisco.com

- **Outbound port can be oversubscribed.**
- **Voice requires low jitter, low loss.**
- **Solution: Separate queues for voice and data.**

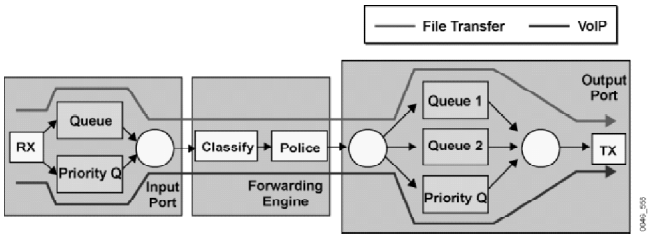

DESGN v1.1—3-40

QoS for Voice over IP (VoIP) allows packet loss and delay within a certain tolerance that does not affect voice quality. The straightforward solution is to provide sufficient bandwidth at all points in the network. A less costly alternative is to apply a QoS mechanism at the oversubscribed points in the network.

A reasonable design goal for end-to-end network delay for VoIP is 150 milliseconds. At this level, speakers do not notice the delay. To achieve guaranteed low delay for voice at campus speeds, you can provide a separate outbound queue for real-time traffic. Bursty data traffic, however, such as file transfer, is placed in a different queue.

QoS maps well to the multilayer campus design. Packet classification is a service applied as close to the originating source as possible, which is the ingress point to the network. A characteristic port number recognizes VoIP traffic flows. An IP CoS value indicating “low delay voice” classifies VoIP packets. Wherever the VoIP packets encounter congestion in the network, the local switch or router will apply the appropriate congestion management based on the type-of-service value.

Designing the Building Access and Building Distribution Submodules

In a conventional campus-wide VLAN design, you apply data link layer switching to the Building Access submodule, while the Building Distribution switches support multilayer capabilities. With multilayer capability, you can separate Building Access networks into failure and broadcast domains. In small networks, you can merge Building Access and Building Distribution submodules in a single switch. This topic discusses designing the Building Access and Building Distribution submodules.

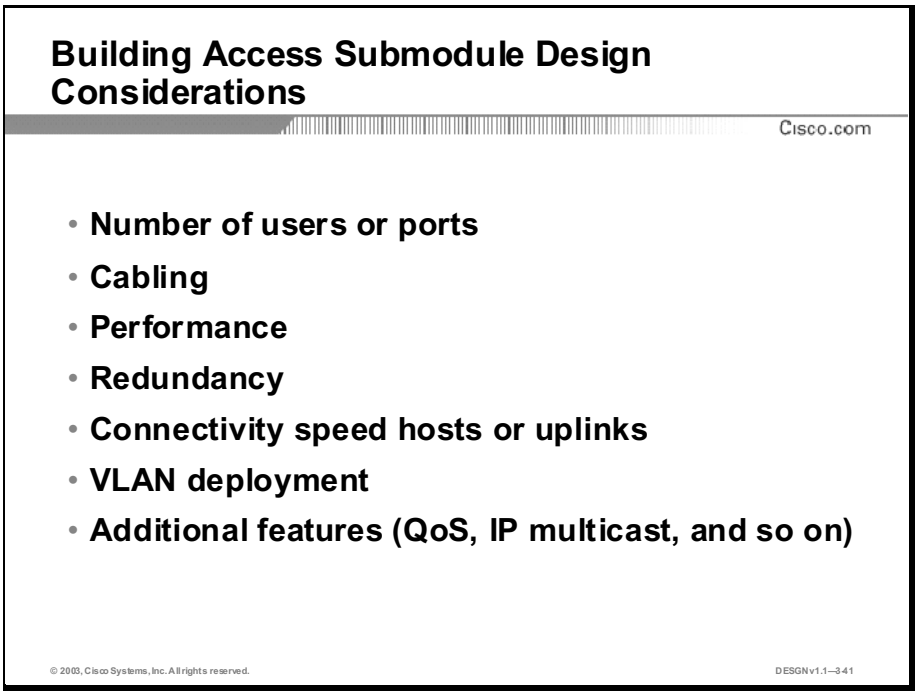

The Building Access submodule aggregates the workstations into a data link layer switch. When multiple workgroups share the same data link layer switch, implement VLANs to provide separate broadcast and failure (STP) domains.

The policies implemented on the Building Access switch are based on data link layer information. These policies focus on and include these features:

- Port security
- Access speeds
- Traffic classification priorities defined on uplink ports

When implementing the Building Access submodule of the Campus Infrastructure functional area, you must answer these questions:

- How many user or host ports are currently required in the wiring closet, and how many will be needed in the future? Should the switches support fixed or modular configuration?
- What cabling is currently available in the wiring closet, and what cabling options exist for uplink connectivity?
- What data link layer performance is needed for the node?
- What level of redundancy is needed?
- What is the requested link capacity to the Building Distribution switches?
- How will you deploy the VLANs and the STP? Will there be a single VLAN or several VLANs per Building Access switch? Will the VLANs reside on one switch or spread across multiple switches? Campus-wide VLANs are not desirable due to the large diameter of each VLAN's fault domain.
- Are additional features, such as port security, multicast traffic management, and QoS (traffic classification based on ports) needed?

The Building Access submodule should maintain the simplicity of traditional LAN switching, with the support of basic network intelligent services and business applications.

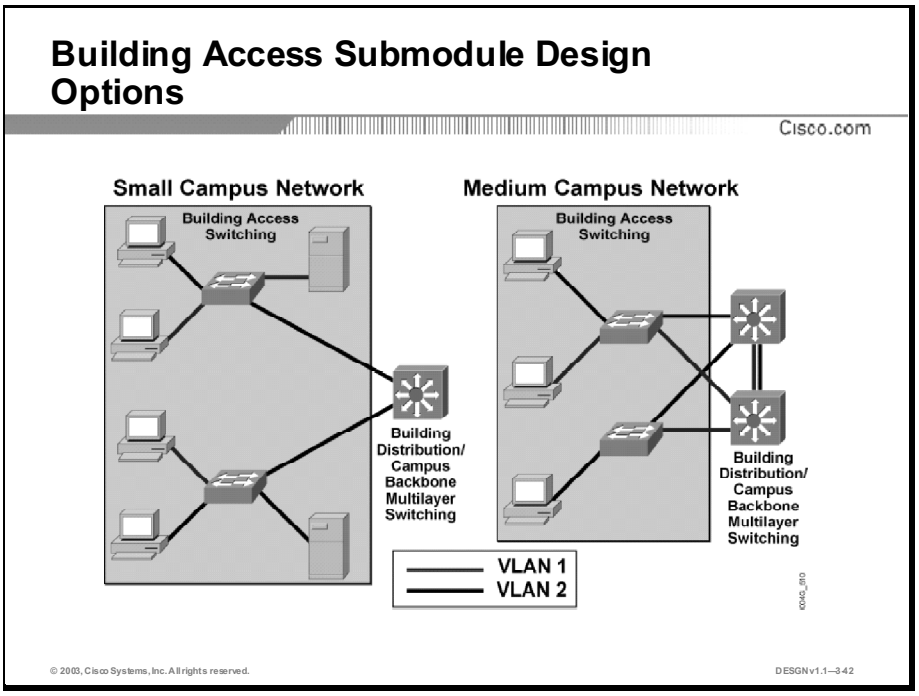

Small Campus Network Design Option

In small campus networks, network servers and workstations are connected to the same wiring closet. Switches in small campus networks do not usually require high-end performance. Therefore, you can combine Building Access and Building Distribution submodules. Low-end multilayer switches could provide the routing services closer to the end user when there are multiple VLANs.

Medium Campus Network Design Option

Connected by uplinks to the Building Distribution multilayer switches, medium-size campus networks are built on data link layer Building Access switches because of a medium-sized campus network's performance requirements. A medium-sized campus network forms a clear modular structure (Building Access and Building Distribution submodules). If redundancy is required, you can attach an additional multilayer switch to the aggregation point of the network with full link redundancy.

Network Uplink Redundancy

You can implement redundant paths for failover or for load balancing, if redundancy from the Building Access submodule to the Building Distribution submodule exists. Data link layer switches may support features that can accelerate STP timers and provide faster convergence and switchover of traffic to the redundant link. You should implement STP acceleration by using either the BackboneFast or UplinkFast feature or by implementing RSTP as specified in the IEEE 802.1w standard.

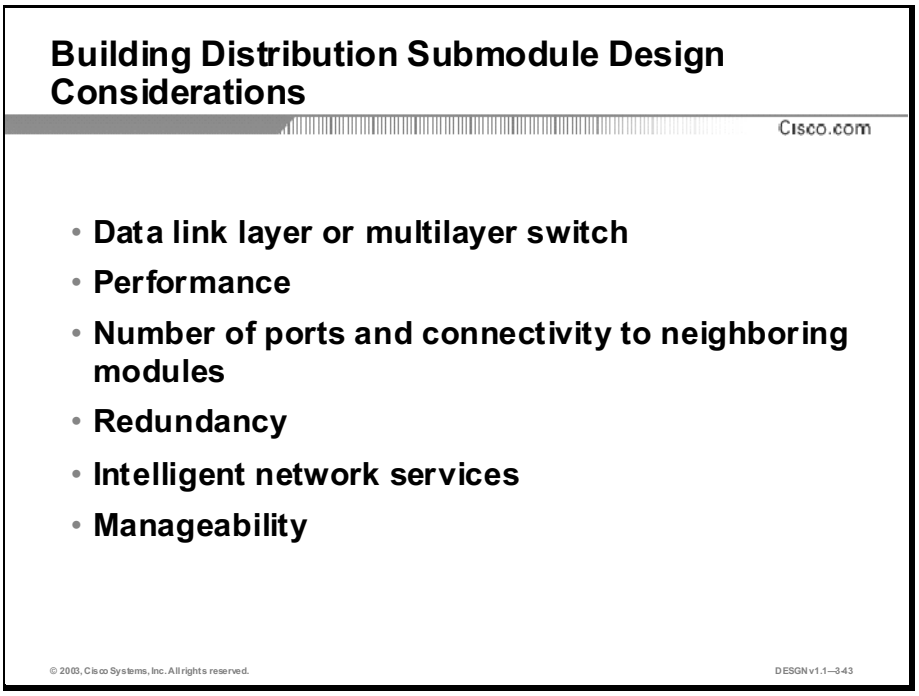

The Building Distribution submodule aggregates the Building Access switches and uses a combination of data link layer and multilayer switching to segment workgroups and isolate segments from failures and broadcast storms. The Building Distribution submodule implements a number of policies based on access lists and QoS settings. The Building Distribution submodule can protect the Campus Backbone from Building Access submodule impacts if you implement all the policies. To determine whether to deploy a data link layer or a multilayer switch in the Building Distribution submodule, answer these questions:

- How many users will the Building Distribution switch handle?
- What type and level of redundancy are needed?
- As intelligent network services are introduced, will the network continue to deliver high performance for all its applications, such as video on demand, IP multicast, or IP telephony?

You must pay special attention to the switching performance, support of intelligent services, manageability, and scalability.

- **Performance:** Building Distribution switches should provide wire-speed performance on all ports. This feature is important because of Building Access aggregation on one side and high-speed connectivity of the Campus Backbone submodule.
- **Intelligent network services:** Building Distribution switches should support fast switching and further incorporate intelligent network services such as high availability, QoS, security, and policy enforcement.
- **Manageability and scalability:** When expanding or reconfiguring Building Distribution devices, the process must be easy and efficient. These devices must support the required management features.

Note: Multilayer switches are generally preferred in the Building Distribution submodule because a Building Distribution switch usually must support intelligent network services such as QoS and traffic filtering.

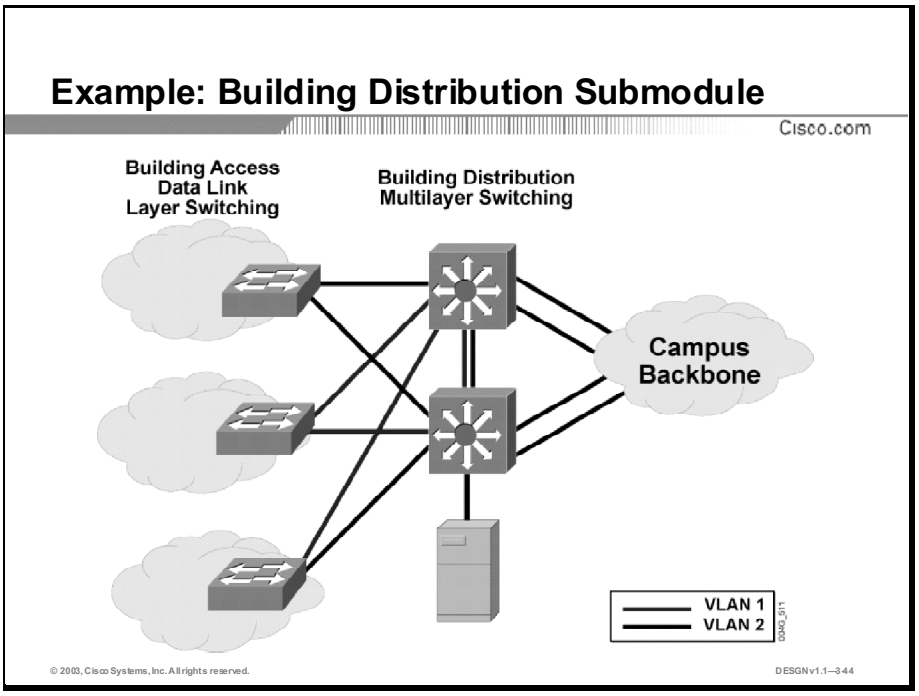

In a simple campus network, as illustrated in the figure, each Building Access submodule has two equal-cost paths to the Building Distribution switches, which have two equal-cost paths to the backbone that ensure fast failure recovery and possible load sharing. Because of built-in redundancy, the network designer addressed STP on the Building Distribution switches, because of the possibility for Layer 2 loops.

The figure illustrates the Building Access submodule connected to both Building Distribution switches, which are also directly interconnected. If the same VLAN is spread across all links, the loop exists so the STP must run on Building Access and Building Distribution switches. An additional concern is STP recovery time if the link to the Building Access submodule fails. STP features such as UplinkFast, BackboneFast, or RSTP can reduce the time to switch from one active link to another. If the connectivity to the Campus Backbone is based on network layer topology, the STP on those ports is not necessary.

Redundancy Options

You must decide where to implement redundancy and which mechanisms to use: Layer 2 with STP or Layer 3 with routing protocol redundant path configuration. If you do not implement advanced STP features such as UplinkFast, BackboneFast, or RSTP, Layer 2 redundancy and STP configuration may take up to 50 seconds. If you enable the advanced STP features on the switch, the switchover time could be from 1 second (in case of RSTP deployment) to 30 seconds. The routing protocols usually switch over in a few seconds. (EIGRP in redundant configuration is usually faster than OSPF because of the default 5-second shortest path first [SPF] delay recalculation.)

Designing the Campus Backbone

Low price per port and high port density may govern wiring-closet environments, but high performance and high availability drive the design for the Campus Backbone submodule. This topic describes design considerations for the Campus Backbone submodule.

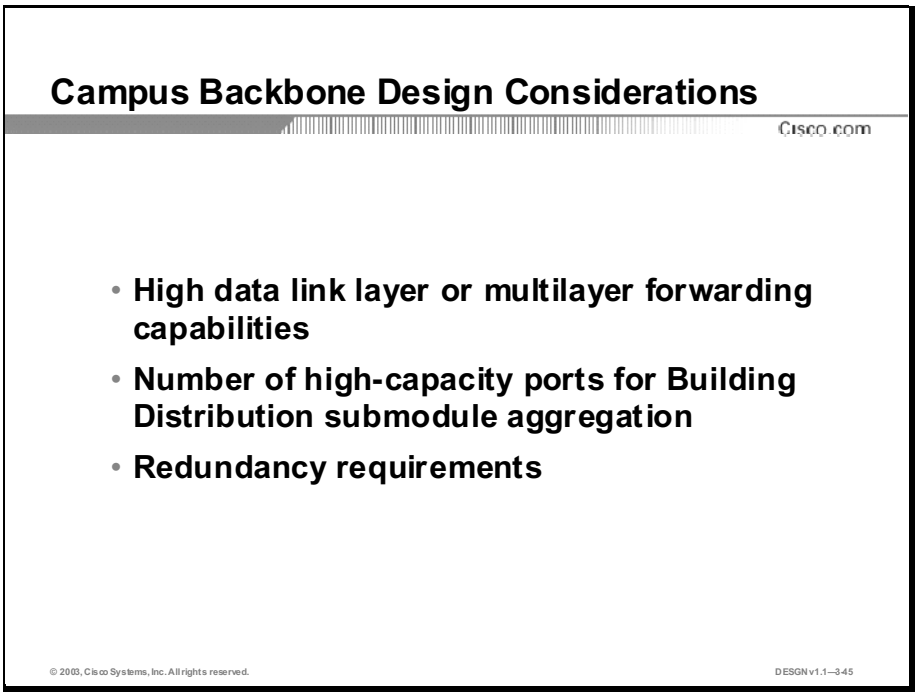

Deploy a dedicated Campus Backbone submodule to connect three or more buildings in an enterprise campus. Backbone switches reduce the number of connections between the Building Distribution switches and simplify the integration of the Server Farm and Edge Distribution modules. Campus Backbone switches focus primarily on wire-speed forwarding on all interfaces. Campus Backbone switches are differentiated by the level of performance achieved per port, rather than by high densities.

When implementing the Campus Backbone, you will select the switching mechanism, either data link layer, or multilayer. The issues to consider include these:

- The performance needed in the Campus Backbone network
- The number of high-capacity ports for Building Distribution submodule aggregation and connection to the Server Farm or Edge Distribution submodules
- The need for high availability: Deploy at least two separate switches, ideally located in different buildings

Example: Single VLAN Data Link Layer Campus Backbone Design

Cisco.com

- **Single VLAN in the backbone**
- **Limitation: Broadcast and multicast frames**

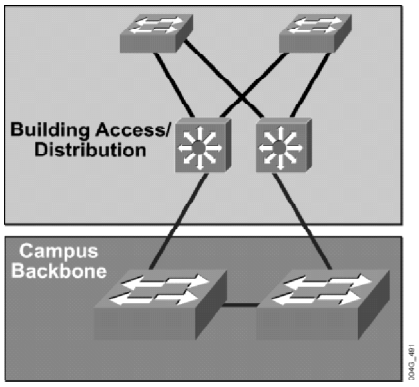

DESGN v1.1—3-46

The data link layer-based Campus Backbone example consists of a single data link layer switch that represents a single VLAN with a star topology towards Building Distribution switches. A single IP subnet is used in the Campus Backbone submodule, and each multilayer Building Distribution switch routes traffic across the Campus Backbone subnet. In this case, no loops exist, STP does not put any links in blocking mode, and STP convergence does not affect the backbone.

The figure illustrates a data link layer Campus Backbone with two switches and a single VLAN configured across both.

To prevent STP loops in the example, you must define the distribution switch links to the Campus Backbone as routed interfaces, not as VLAN trunks. The solution can lead to problems resulting from numerous Layer 3 peerings between the devices attached to the data link layer Campus Backbone, especially if the number of routers is high.

Note: One drawback of a data link layer-switched Campus Backbone is the lack of mechanisms to handle broadcast and multicast frames efficiently. The entire backbone is a single broadcast domain. Broadcast traffic increases CPU utilization on network devices and consumes available bandwidth in the Campus Backbone network, although the broadcast/multicast suppression feature can prevent the flood of such packets.

Example: Split Data Link Layer Campus Backbone Design

Cisco.com

- **Single VLAN per Campus Backbone switch**
- **Limitation: Broadcast and multicast frames**

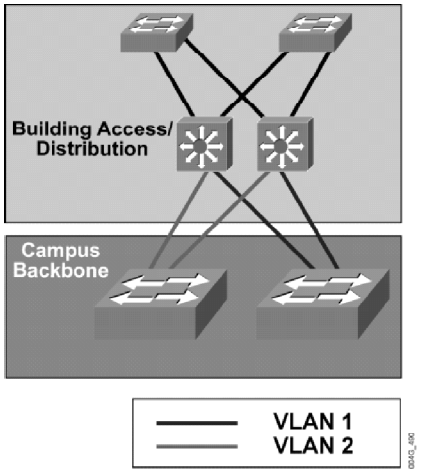

As an alternative, you can implement two VLAN domains, each on one switch but without a connection between the switches. The figure illustrates this solution, which is known as a split backbone. The advantage of this design is that two equal-cost paths across the backbone support fast convergence and possible load sharing.

Although this design increases availability, it suffers from data link layer problems, including inefficient handling of broadcast and multicast frames. In the figure, the broadcast domain is limited to a single switch which equals a single VLAN.

Example: Multilayer Switched Campus Backbone Design

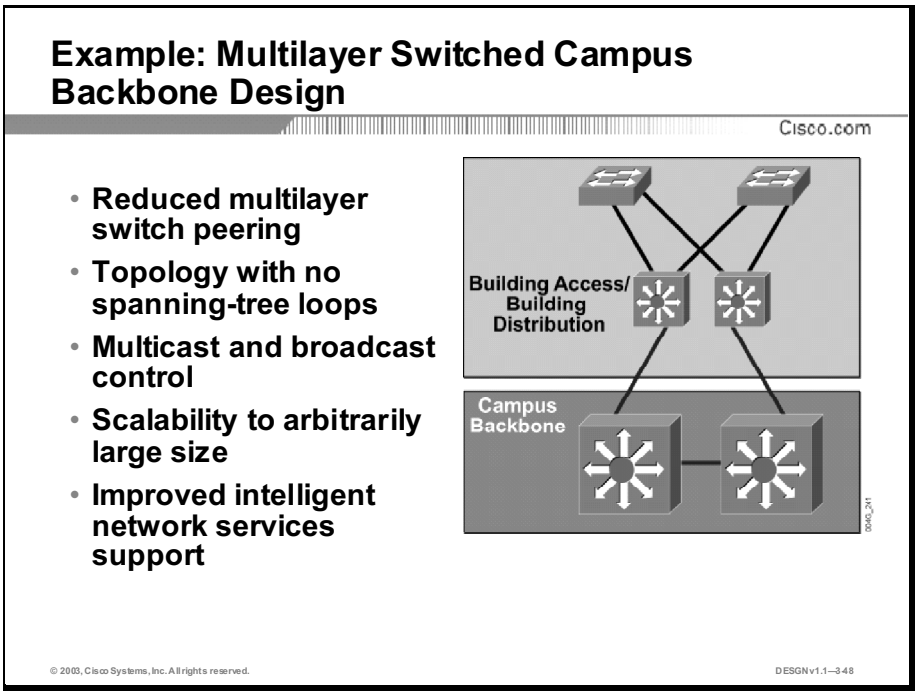

For large enterprise networks, the most flexible and scalable Campus Backbone submodule consists of multilayer switches, as illustrated in the figure.

Multilayer-switched Campus Backbones provide several improvements over the data link layer Campus Backbone:

- **Reduced multilayer switch peering (routing adjacencies):** Each multilayer Building Distribution switch connects only to one multilayer Campus Backbone switch. This implementation simplifies any-to-any connectivity between Building Distribution and Campus Backbone switches.
- **Topology with no spanning-tree loops:** There is no STP activity in the Campus Backbone or on the Building Distribution links to the Campus Backbone because all the links are routed links. Arbitrary topologies are supported because of the routing protocol used in the Campus Backbone.
- **Multicast and broadcast control**
- **Scalable to an arbitrarily large size**
- **Improved intelligent network services support:** Multilayer Campus Backbone switches provide better support for intelligent network services.

Multilayer switches are more sophisticated devices for high-speed packet routing. The switches support routing in the hardware although the hardware may not support all features. If a selected feature such as security or QoS is not supported in hardware, the switch must perform the function in software, which may dramatically reduce the data transfer rate.

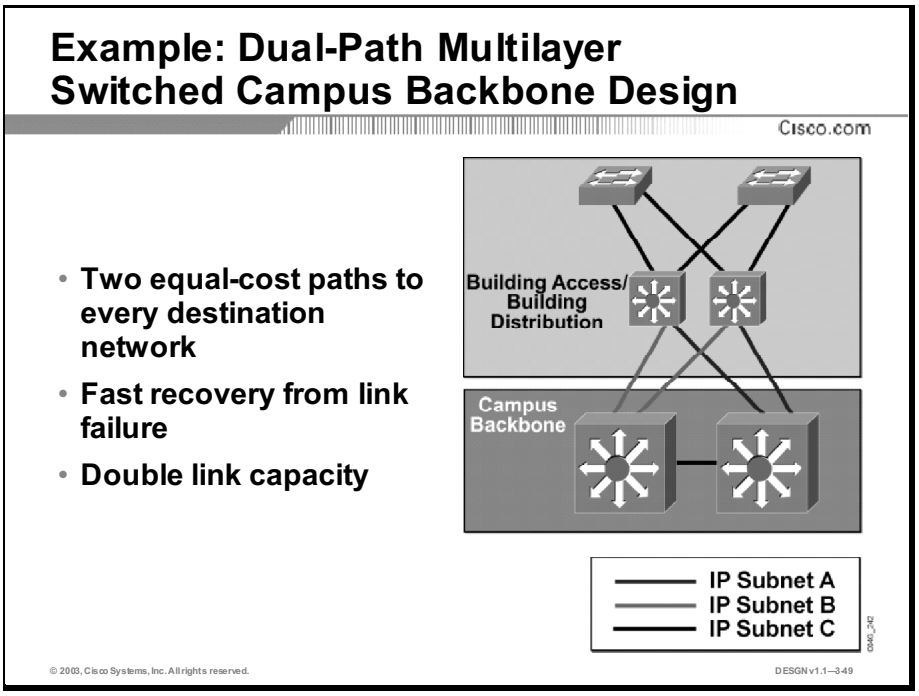

You usually deploy dual links to the Campus Backbone from each Building Distribution switch to provide redundancy and load sharing in the multilayer-switched Campus Backbone, as illustrated in the figure. This design maintains two equal-cost paths to every destination network. Thus, recovery from any link failure is fast and load sharing is possible, resulting in higher throughput in the Campus Backbone submodule.

The Campus Backbone switches should deliver high-performance, multilayer switching solutions for an enterprise campus, and should address requirements for:

- Gigabit density
- Data and voice integration
- LAN/WAN/MAN convergence
- Scalability
- High availability
- Intelligent multilayer switching in Campus Backbone, Building Distribution, and Server Farm environments

Note: You can implement the Campus Backbone with both data link layer and multilayer switching based on the need for auxiliary VLANs and private VLANs. The auxiliary VLAN feature allows you to place IP Phones into their own VLAN without any end-user intervention. The private VLAN feature places all servers in the same private VLAN, simplifying the Server Farm module design when servers have no need to communicate between themselves.

Note: Consider integrating the Network Management module with the Campus Backbone. Place the Network Management module in its own subnet and with traffic routed across the network.

Designing the Server Farm Module

Centralized servers are grouped into one or more Server Farm modules, which may be physically located in an enterprise campus or in a separate data center. This topic describes design considerations for a Server Farm module.

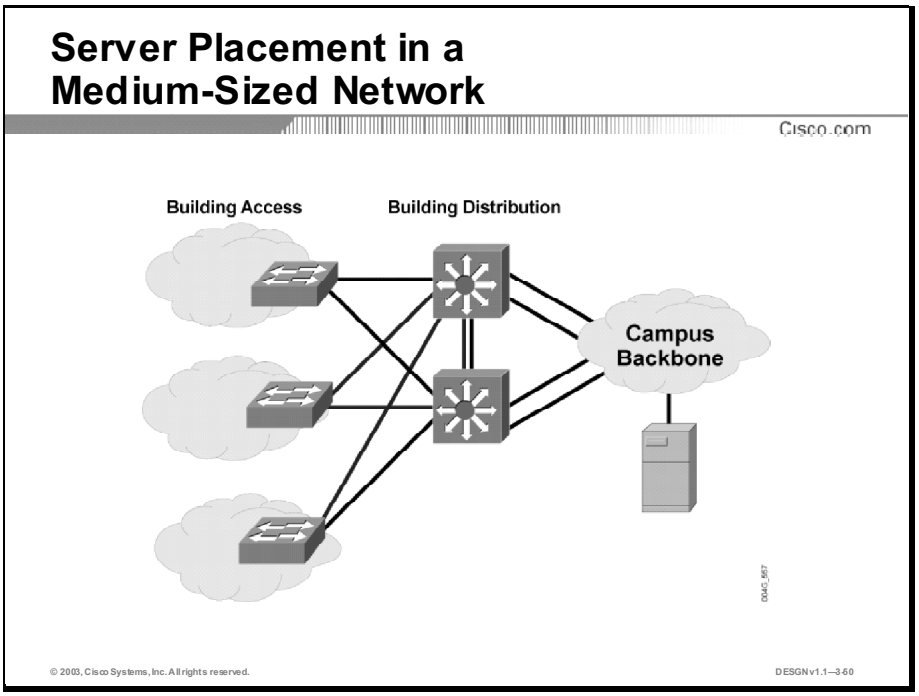

With centralized servers directly attached to the Campus Backbone, all client-server traffic crosses one hop from a subnet in the Building Access submodule to a subnet in the Campus Backbone submodule.

In general, the Campus Backbone provides fast transport of the traffic without any limitations. You can connect servers in medium-sized networks directly to Campus Backbone switches, making servers only one hop away from the users. However, the need to control traffic and access in the backbone then arises. Implement QoS and access control list (ACL) policy-based control to access the Server Farm in the Building Distribution or Edge Distribution modules.

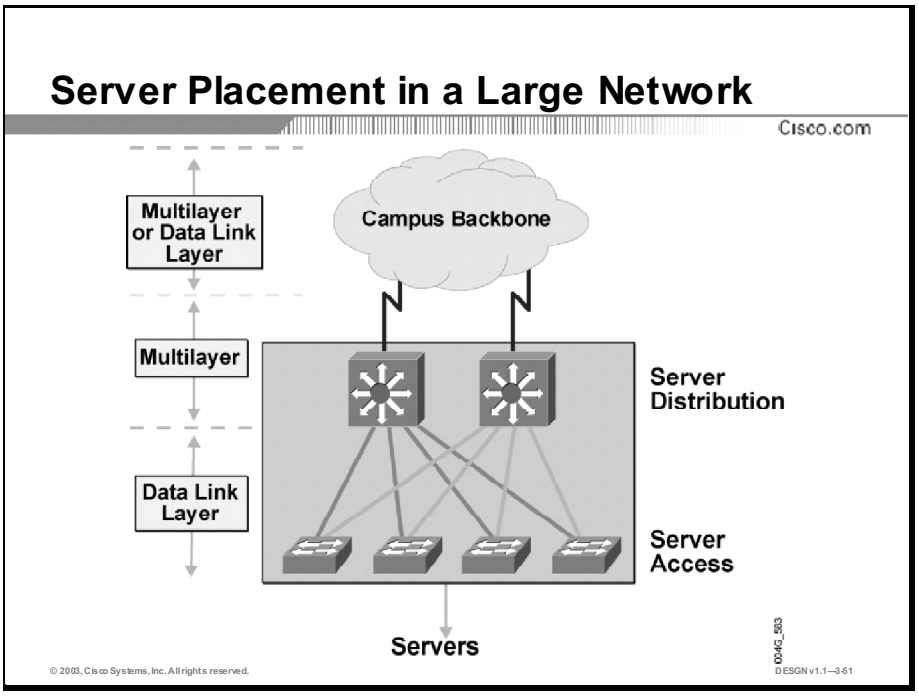

Larger enterprises place common servers in a Server Farm module and connect them to the Campus Backbone using multilayer Server Distribution switches. Because of high traffic load, the servers are usually Fast Ethernet-attached, Fast EtherChannel-attached, or Gigabit Ethernet-attached. Access lists at the multilayer Server Distribution switches in the Server Farm module control access to these servers. Fast failover is provided by redundant Server Distribution switches and solutions such as the Hot Standby Router Protocol (HSRP). In addition, the Server Farm module distribution switches keep all server-to-server traffic off the Campus Backbone.

Rather than being installed on only one server, modern applications are distributed among several servers. This approach improves application availability and responsiveness. Therefore, placing servers in a common group (Server Farm module) and using intelligent multilayer switches provides scalability, availability, responsiveness, throughput, and security for applications and services, as required.

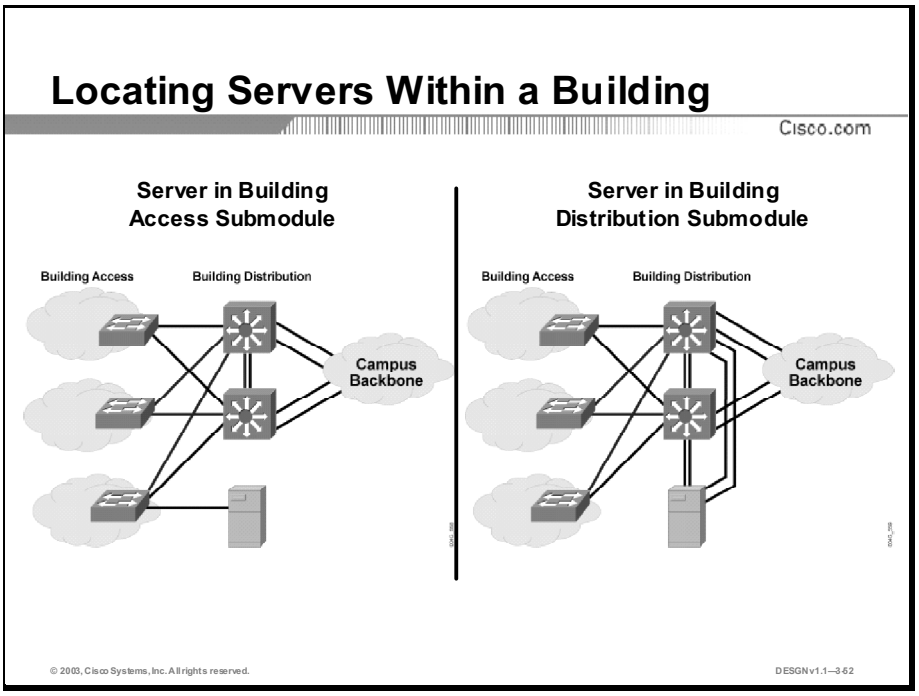

Local Server in a Building Access Submodule

If a server is local to a workgroup using a single VLAN and most of the server traffic is to or from the workgroup, you can connect the server directly to a switch port that participates in that VLAN. This scenario follows the conventional 80/20 workgroup rule for campus traffic distribution. Optionally, you could hide these servers from the enterprise by implementing an access list at the Building Distribution switch.

Server in a Building Distribution Submodule

In mid-size networks, you can attach servers to Building Distribution switches. You can define servers as building-level servers that communicate with clients in different VLANs while within the same physical building. You can create a direct L2-switched path between a server and clients in a VLAN in one of two ways:

- With multiple NICs making a direct attachment to each VLAN
- With a trunk connection or a separate VLAN on the Building Distribution switch for the common servers

If required, you can selectively hide servers from the rest of the enterprise by using an access list on the Building Distribution switch.

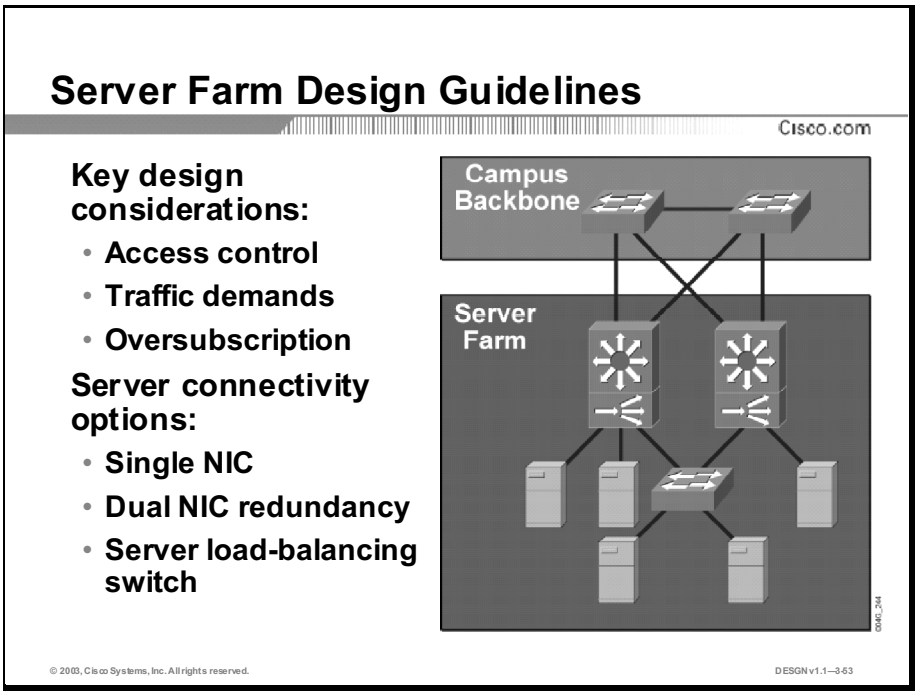

Using a modular design approach, you implement the Server Farm as a high-capacity building block attached to the Campus Backbone submodule. One of the main concerns regarding the Server Farm is that it receives the majority of the traffic from the whole campus. The uplink ports on switches are frequently oversubscribed, which can result in random frame drops. To reduce random frame drops during peak traffic, you can apply QoS mechanisms to the server links.

Note:	Switch oversubscription occurs when some switches allow more ports (bandwidth) in the chassis than the hardware of the switch can transfer through its internal structure.

It is extremely important to design the Server Farm switches with less oversubscription than the switches that reside in the Building Access or Building Distribution submodules. If the campus consists of a few Building Distribution submodules connected to the Campus Backbone with Fast Ethernet, then attach the Server Farm module to the Campus Backbone with Gigabit Ethernet or multiple Fast Ethernet links.

You must also evaluate the server's capabilities. Although server manufacturers support a variety of NIC connection rates, such as Gigabit Ethernet, the underlying device drivers or server operating system may not be able to transmit at line capacity. As such, you can raise oversubscription ratios, thereby reducing the overall cost of the Server Farm.

Server Connectivity Options

You can connect a server with a single or dual Fast Ethernet connections. If the server is dual-attached, one interface may be active while the other is in hot standby. Installing multiple single-port or multiport NICs in the server extends dual homing past the Server Farm module switches to the server itself.

Within the Server Farm, you use multiple VLANs to create multiple policy domains, as required. If one particular server or application has a unique access policy, you can create a unique VLAN and subnet for that server or application. If a group of servers that the application requires has a common access policy, you can place the whole group in a common VLAN and subnet and apply ACLs on the interfaces of the multilayer switches.

Note: There are several other solutions that improve the responsiveness of the servers and distribute the load evenly on them. The content switching solution is implemented by content switches that provide a robust front end for server farms and perform functions such as load balancing of user requests across server farms to achieve optimal performance, scalability, and content availability.

How Applications Affect Switch Performance

You must consider two main thresholds when designing the Server Farm module. Packets are generated at an average frequency with an average size based on the traffic patterns and number of enterprise application users.

Interactive applications such as conferencing generate high packet rates, but with small packet sizes. In terms of bandwidth for these applications, the packets-per-second (pps) limit of the switches may be more critical than the throughput.

Applications involving large movements of data, such as file repositories, transmit a high percentage of full-length packets. Uplink bandwidth and oversubscription ratios become key factors in the overall design. Actual switching capacities and bandwidths vary, based on the mix of applications.

Designing the Edge Distribution Module

The key design consideration for the Edge Distribution module is security. This topic discusses design considerations for the Edge Distribution module.

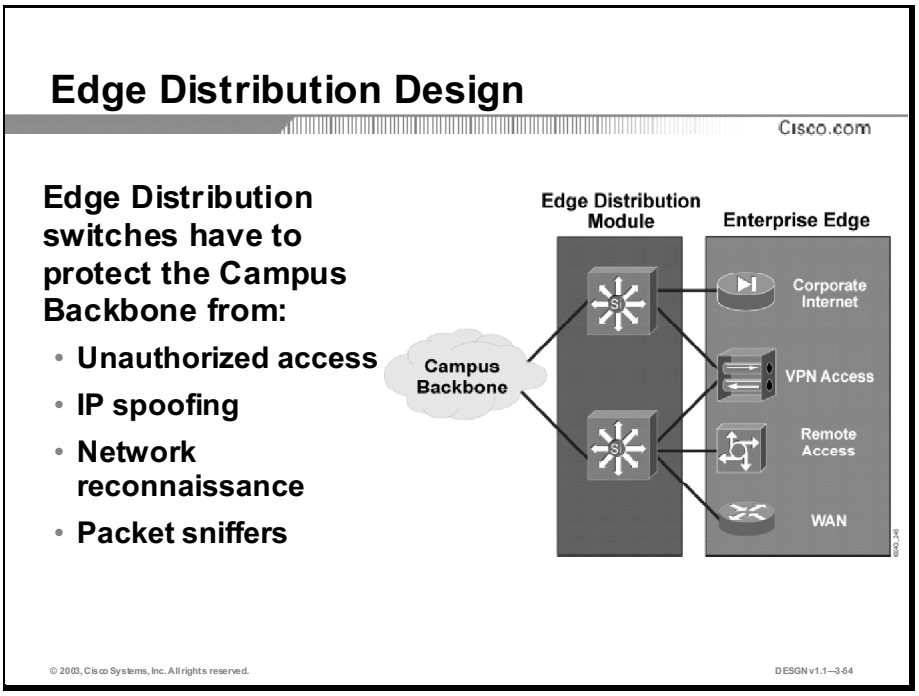

The Edge Distribution module filters and routes traffic into the Campus Backbone. Multilayer switches are the key devices that aggregate edge connectivity and provide advanced services. The speed of switching is not as important as security in the Edge Distribution module. The Edge Distribution module isolates and controls access to servers located in the Enterprise Edge modules. These servers are closer to the external users and, therefore, introduce a higher risk to the internal campus. To protect the core from threats, the switches in the Edge Distribution module must protect the campus against these dangers:

- **Unauthorized access:** The Edge Distribution module must verify each user and the user's rights before traffic passes to the Campus Backbone. Filtering mechanisms must provide control over specific edge subnets and their ability to reach areas within the campus.
- **IP spoofing:** IP spoofing is a hacker technique for impersonating another user's identity by using their IP address. Denial-of-service attacks use the IP spoofing technique to generate requests to the servers with the stolen IP address as a source. The server does not respond to the original source but does respond to the stolen IP address. Denial-of-service attacks are difficult to detect and defend against. A significant amount of this type of traffic makes the attacked server unavailable, thereby interrupting business.
- **Network reconnaissance:** Sending packets into the network and collecting responses from the network devices implements network reconnaissance (discovery). These responses provide basic information about the internal network topology. Network intruders use this approach to learn about network devices and the services they run. Therefore, filtering potential traffic that results from the network reconnaissance mechanisms before it enters the enterprise network may be crucial.

- **Packet sniffers:** Packet sniffers, devices that monitor and capture the traffic in the network, represent another threat. Packets belonging to the same broadcast domain are vulnerable to capture by packet sniffers, especially if the packets are broadcast or multicast. Multilayer switches can prevent a packet sniffer attack.

The Edge Distribution module provides the last line of defense for all external traffic destined to the Campus Infrastructure module. In terms of overall functionality, the campus Edge Distribution module is similar to the Building Distribution submodule. Both modules use access control to filter traffic, although the Edge Distribution module can rely on the other Enterprise Edge functional area modules to provide additional security. Both modules use multilayer switching to achieve high performance, but the Edge Distribution module can offer additional security functions because the performance requirements are not as high.

Alternative Designs for Edge Distribution Module

As you can with the Server Farm module and Building Distribution submodule, you can combine the Edge Distribution module with the Campus Backbone submodule, if performance requirements are not stringent. Campus Backbone switches do not require the Network Intrusion Detection System (NIDS), but you can place the Edge Distribution module there using intrusion detection line cards in the multilayer switches.

NIDS reduces the need for external appliances at the points where the critical Enterprise Edge modules connect to the campus.

Note: Performance or security concerns may dictate that you implement dedicated intrusion detection in the Enterprise Edge modules, rather than in the Edge Distribution module.

Summary

This topic summarizes the key points discussed in this lesson.

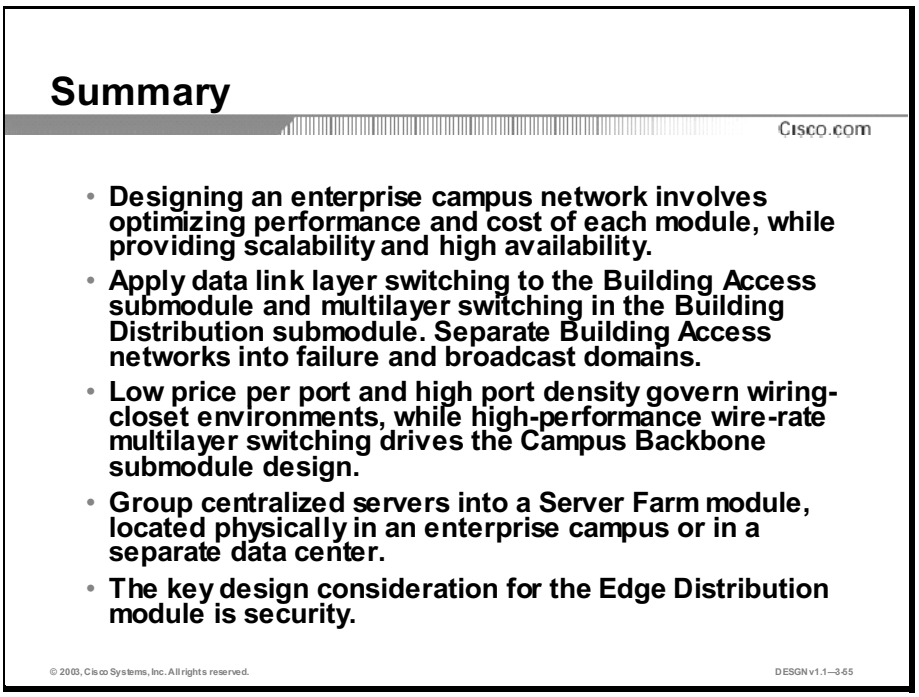

References

For additional information, refer to these resources:

- *A Security Blueprint for Enterprise Networks,* http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm
- *Gigabit Campus Network Design—Principles and Architecture,* http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm
- *LAN Design Guide for the Midmarket,* http://www.cisco.com/warp/public/cc/pd/si/casi/ca3500xl/prodlit/lan_dg.htm

Next Steps

For the associated case study and exercises, refer to the following sections that follow the Quiz:

- Case Study 3: Enterprise Campus Design
- Simulation 3-1: Shared vs. Switched LAN
- Simulation 3-2: Data Link Layer vs. Multilayer Switching

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) An organization placed servers to which all corporate users have access in a separate VLAN. The users are divided into organizational VLANs. However, some of the common workgroup devices are still located within these VLANs. What is the expected traffic flow?

A) Most of the traffic will be local.

B) Most of the traffic will have to be multilayer switched.

C) All traffic will require multilayer switching.

D) There is no need for multilayer switching.

Q2) A company has deployed Cisco IP/TV, a product that utilizes IP multicast to deliver video and audio streams. The routers are configured for IP multicast. Taking into account that the majority of the LAN switches are low-end switches, which protocol should you enable on the LAN switches to reduce flooding?

A) GMRP

B) VTP

C) CGMP

D) STP

E) PIM

Q3) A connection (trunk) from the Building Distribution to the Campus Backbone switch is 100 Mbps, and the average utilization is 80 percent or more. Employees on the network use web-based, business-critical applications. How would you minimize packet loss or delay on the multilayer Building Distribution to avoid impacting the business applications?

A) Implement more VLANs on the Building Access switch so that the business users are assigned to a separate VLAN.

B) Rewrite the priority bits on the Building Distribution switch.

C) Implement QoS with classification and policing on the Building Distribution switch.

D) Classify the users on the Building Access switch with different priority bits.

Q4) A corporate network is spread over three floors. There is one data link layer switch on each floor with more than one VLAN. One connection from each floor goes to the basement, where all WAN connections are terminated. Currently, servers are installed in each VLAN but the company will soon move the servers to the basement. Traffic between VLANs is essential. What network design should the company use?

A) Connect the multilayer switch on each floor to a data link layer switch in the basement.

B) Connect the data link layer switch on each floor to a multilayer switch in the basement.

C) Connect the data link layer switch on each floor to a data link layer switch in the basement and the VLAN router.

D) Connect the data link layer switch on each floor to a multilayer switch on the same floor for inter-VLAN communication.

Q5) An organization requires its Campus Backbone network to have high resilience. The servers that are directly attached to the Campus Backbone submodule produce a large amount of multicast traffic. What switch design is appropriate for this network?

A) multiple data link layer switches with full-mesh topology

B) a data link layer switch with redundant connectivity to other network modules

C) multiple multilayer switches with full-mesh topology and servers in a separate VLAN

D) a non-modular multilayer switch and servers attached to it in a separate VLAN

Q6) The departments of a corporation are spread across several buildings, while they use common servers. Network policy and security are important. Where should the corporation place the servers and how should they be attached to the network?

A) The company should create a Server Farm module with its own switches connected to the Campus Backbone submodule.

B) Each building should have some servers connected to the Building Distribution switches.

C) The servers should directly connect to the data link layer switches in the Campus Backbone submodule of the network.

D) The servers need to be close to users, so they should be attached to the Building Access switch.

Q7) A large corporation has a campus network composed of the Building Access, Building Distribution, Campus Backbone, and Server Farm modules. It needs to implement a WAN connection to remote locations and to support Internet access. How should the company implement the WAN and Internet connections to the campus network?

A) using Building Access switches

B) using Building Distribution switches

C) using Edge Distribution switches with security

D) using multilayer Campus Backbone switches

Quiz Answer Key

Q1) B

Relates to: Designing the Enterprise Campus Network

Q2) C

Relates to: Designing the Enterprise Campus Network

Q3) C

Relates to: Designing the Enterprise Campus Network

Q4) B

Relates to: Designing the Enterprise Campus Network

Q5) C

Relates to: Designing the Campus Backbone

Q6) A

Relates to: Designing the Server Farm Module

Q7) C

Relates to: Designing the Edge Distribution Module

Case Study 3: Enterprise Campus Design

Complete this case study to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the campus design guidelines and considered data link layer versus multilayer switching solutions. Upon completing this case study, you will be able to meet this objective:

- Propose the redundant campus design for a given network scenario

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Complete these steps:

Step 1 Refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario before commencing the exercise. Focus on the routing protocol issues. Allow a maximum of 10 minutes for reading.

Step 2 Discuss the scenario and options for campus design with your group. Allow 10 minutes for the discussion.

Step 3 Propose the optimal campus design that addresses the scenario requirements (switched solution, redundancy, servers in a separate segment, and so on).

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class and you have justified any major deviations from the case study solution.

Simulation 3-1: Shared vs. Switched LAN

Complete this exercise to practice what you learned in this lesson.

This exercise is a paper-only version of the simulation that was actually performed by the simulation tool, and it includes the results that the simulation provided.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the campus design methods and options. Upon completing this simulation, you will be able to meet this objective:

- Explain the benefits of using a switched LAN solution compared to a shared LAN solution

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Read the scenario and try to answer the questions that appear in the text. Discuss possible answers and explain your considerations in the classroom.

Shared vs. Switched LAN Scenario

The customer (DJMP Industries) plans to restructure its flat campus network, which consists of workstations and servers located in the central building and building A. The company is considering Ethernet switching technology as a replacement for the 10BASE-T Ethernet hubs. You have been asked to determine the effect that the introduction of the switches may have on the load of the links and to estimate the responsiveness and utilization of the network with respect to the existing applications.

To provide some proof of the future network efficiency, you will model FTP and HTTP performance on the network, using shared and then switched Ethernet platforms.

Client Accessing Server in Unloaded Shared Ethernet

The customer has provided you with information about their existing network and the number of users. You started the initial evaluation of the network behavior by simulating the load on the LAN links posed by a single client accessing the web server (as illustrated in the figure).

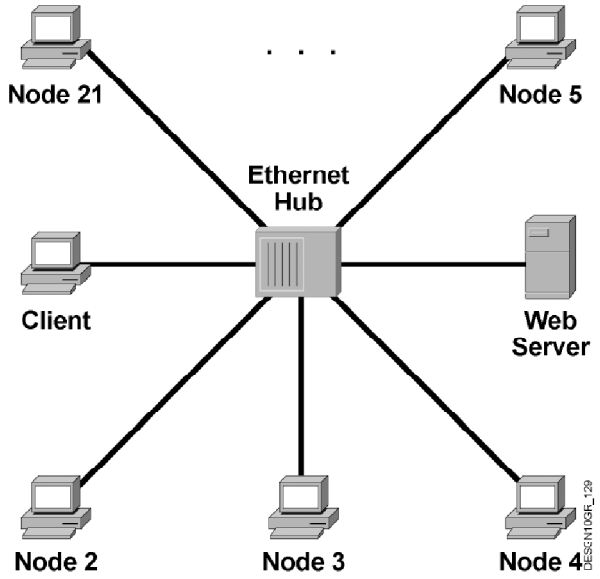

You performed the simulation (10-minute intervals), and the resulting graphs are illustrated in these figures. You should have observed the effect of traffic growth on these statistics and compared the results among different scenarios.

The relevant statistics of interest for this case are the link (Ethernet) utilization and the HTTP response times.

This graph describes the network load that resulted from the HTTP session between the client and the server. The HTTP traffic exchanged between the client and the server does not represent a significant load in the network, which is clearly indicated by the low Ethernet utilization number.

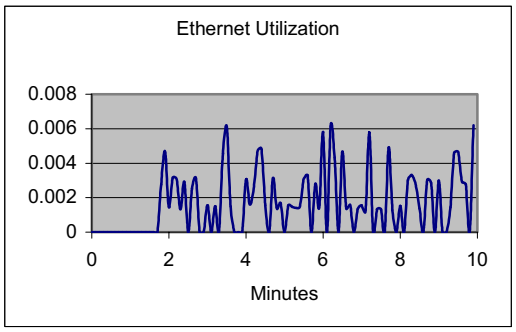

The next baseline test incorporates the measuring of the HTTP response times. On average, the *HTTP page response* times are within the range of 0.01 and 0.015 second, whereas the *HTTP object response* times vary by approximately 0.004 to 0.01 second (every HTTP page consists of several objects).

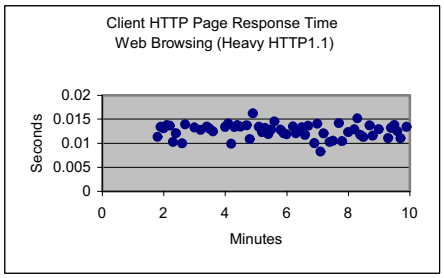

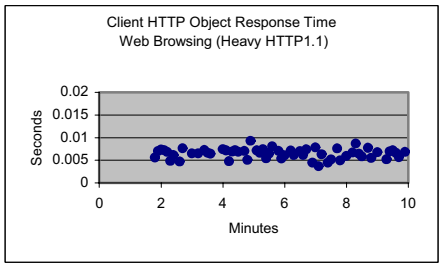

You may also use another point of view, the probability of HTTP response times. These graphs show the probability that the HTTP response time is equal to a particular value.

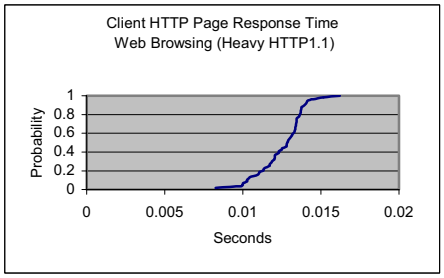

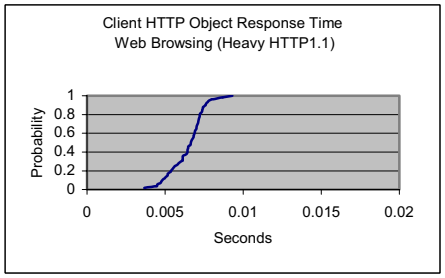

Q1) What can you observe from the graphs?

Client Accessing Server in Loaded Shared Ethernet

Your task now is to create a scenario in which the background traffic is simulated to provide a more realistic picture about the ongoing traffic in the network. The client will continue to access the web server while all the other clients concurrently initiate FTP sessions to a FTP server. A separate FTP server is introduced (as illustrated in the figure) to eliminate the effect of the server utilization. The HTTP session is thus tested in the heavily loaded shared Ethernet network.

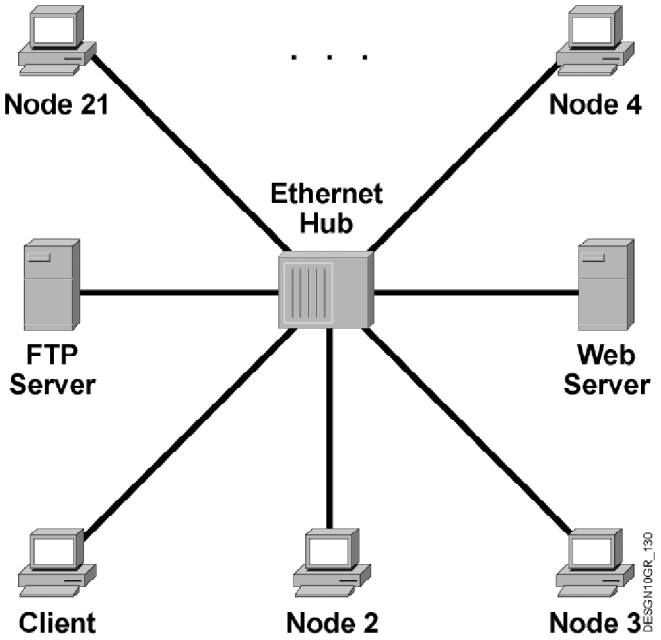

You performed the simulation, and these graphs were produced. Compare them with the results from the previous simulation. This graph describes the increased utilization of the network as a result of the concurrent FTP and HTTP conversations.

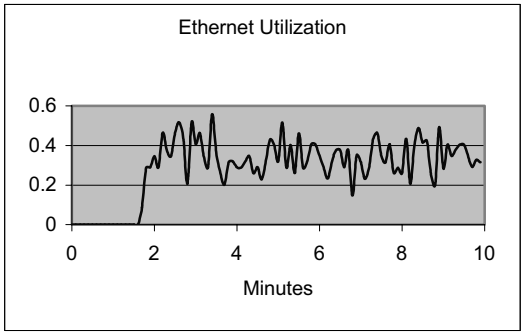

The next step is to observe the HTTP response times again. When examining this graph you can see that, in general, the results match those obtained in the unloaded network. There are some deviations, presumably because of the retransmissions that definitely lower the probability of immediate response. The delayed responses seem to be evenly distributed throughout the observed interval.

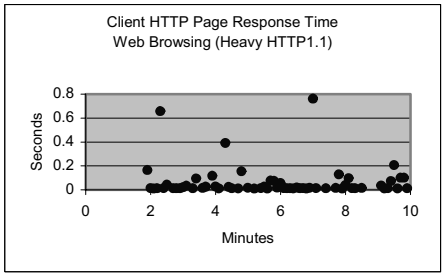

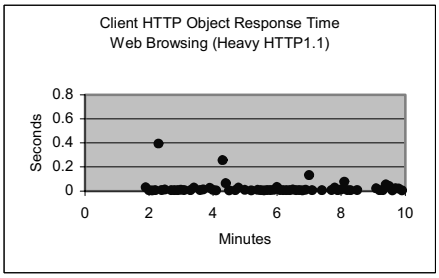

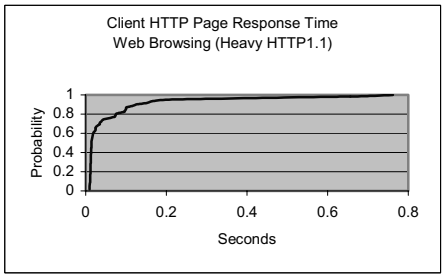

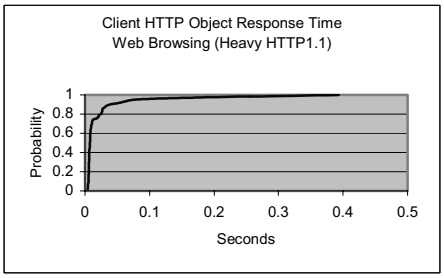

Q2) What can you determine from the results? What is the reason for the delayed HTTP responses?

_____
_____
_____
_____
_____

Introducing Switched Ethernet

In the third simulation scenario, the shared Ethernet is replaced with switched Ethernet, implemented with a single LAN switch (see the figure). The traffic pattern remains the same as in the previous scenario; the client is accessing a web server while all other clients are accessing an FTP server.

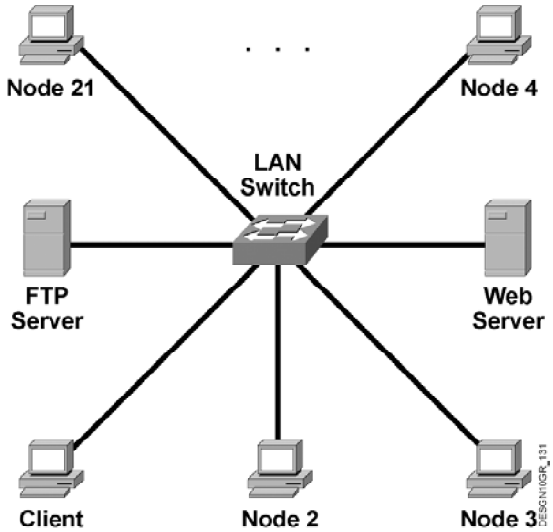

By examining the HTTP response time carefully, it seems that the background FTP traffic does not affect the web communication significantly. Everything is back to normal, the HTTP response times are constantly low, and there is no sign of individual deviations that could compromise the overall statistics.

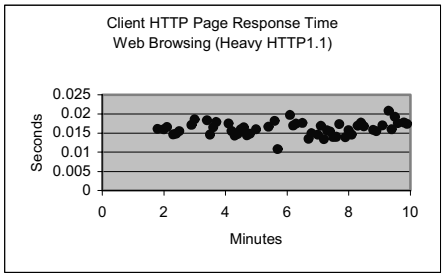

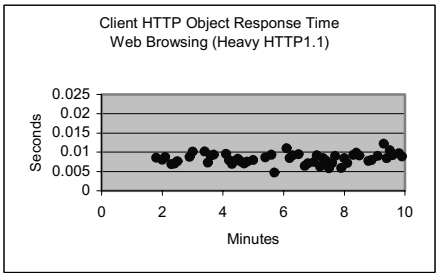

This graph illustrates the probability of receiving a prompt HTTP response. The possibility is almost as high as when a stand-alone HTTP session was simulated (with no background traffic). This leads you to the conclusion that switching technology may be the obvious solution.

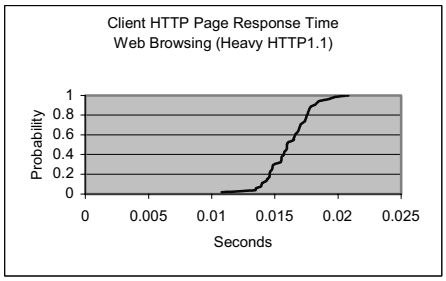

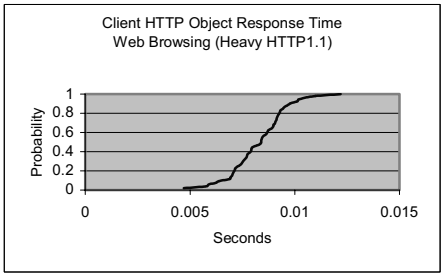

Q3) You came to a conclusion that the introduction of the data link layer switch represents a significant improvement for a given case. How is that determined from the graphs?

Simulation 3-2: Data Link Layer vs. Multilayer Switching

Complete this exercise to practice what you learned in this lesson.

This exercise is a paper-only version of the simulation that was actually performed by the simulation tool, and it includes the results that the simulation provided.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the campus design methods and options. Upon completing this simulation, you will be able to meet this objective:

- Explain the differences between using data link layer and multilayer switching solutions in the Enterprise Campus

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Read the scenario and try to answer the questions that appear in the text. Discuss possible answers and explain your considerations in the classroom.

Data Link Layer vs. Multilayer Switching Scenario

This simulation will demonstrate the impact of data link layer versus multilayer switching on the load in various parts of the structured campus network.

After successful deployment of the switching technology, the company is considering further improvements to its campus network design. They have already finished some baseline wiring work in the central building and in building A, facing some frequent data link layer and multilayer design issues.

You decided to model the company's network to match the existing situation, using this architecture:

- Each building contains Building Distribution switches to which the Building Access (wiring closet or data center concentrator) switches are connected.
- The Building Distribution devices are connected via two central Campus Backbone switches.
- The whole campus is fully redundant.

To provide comparable results, you need a reference traffic flow; therefore, you decided to focus solely on the communication between the two workstations WS_A and WS_B located in different floors of building A, with the server in the central building.

Initial Traffic

In the simulation, workstations A and B communicate with the server with various loads, as illustrated in this graph:

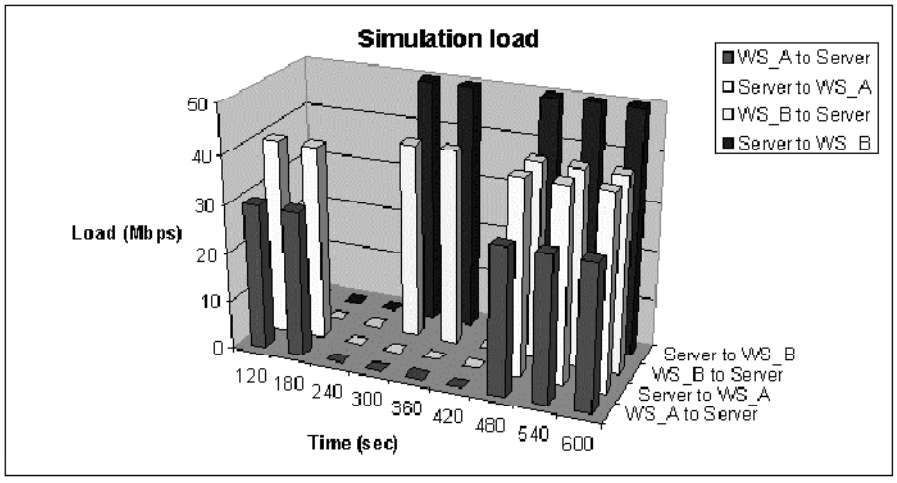

Data Link Layer-Only Design

You started the simulation by turning on the data link layer functionality on all switches in the campus network (see the figure). Soon you realized that, even in the highly redundant data link layer network, the number of possible paths reduces to only one, as determined by the Spanning Tree Protocol (STP). STP computes loop-free networks, and any redundant links belonging to the same LAN or VLAN are placed in the blocking state and cannot be used.

Loaded Network

This graph below shows the result of simulating 10 minutes of traffic, originated by both workstations toward the server, and vice versa. The average-loaded links (30 percent) appear in yellow (dashed) and the heavily loaded links (60 percent) in red (dotted). The resulting yellow (dashed) and red (dotted) arrows indicate that the load is not balanced; specifically, all traffic goes over a single path: DS_B – CS_A – DC_B.

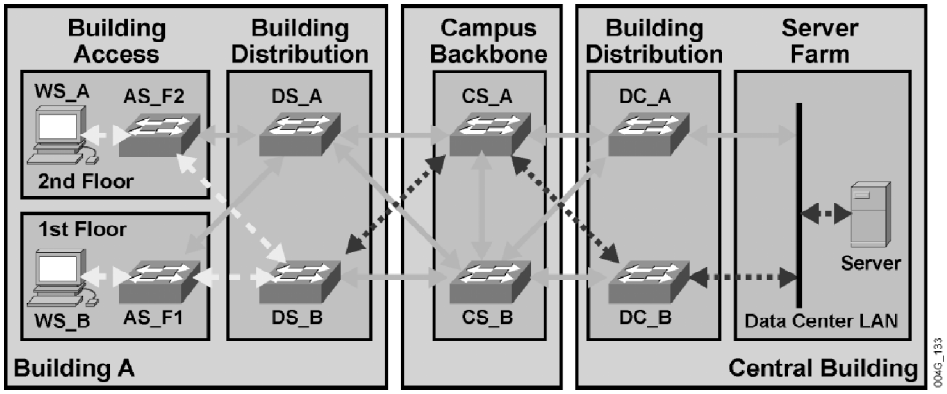

Link Failure

The use of redundant links that terminate at separated devices achieves network. This is especially true for the observed case, where you expect that the link or node failure will not impact the network (at least not for a longer period) and will not result in an imbalance of load.

To prove this, you studied the effect of the link and node failure on the network performance by tearing down the DS_B – CS_A link and afterward disabling the DC_B node. The resulting graph, illustrated in the figure, indicates that the traffic is simply redirected over the alternative path.

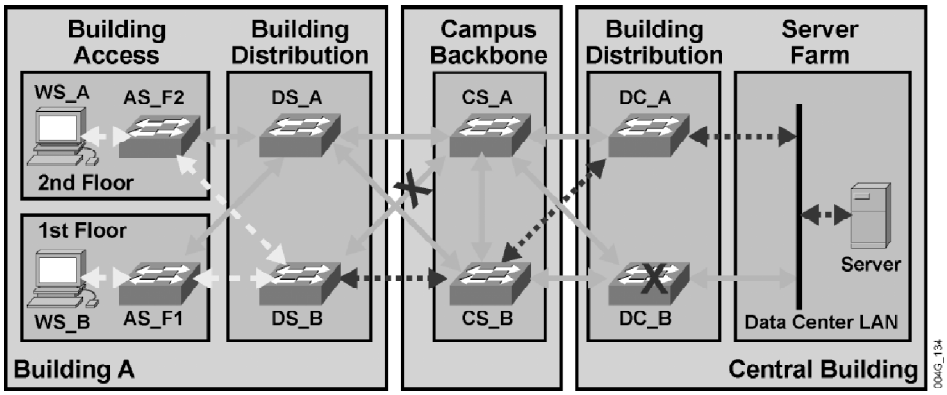

Q4) Will the traffic immediately start using the original path once the link or node has fully recovered?

Multilayer Switching in the Building Distribution Submodule

Next, you decided to replace Building Distribution data link layer switches with multilayer switches, eliminating the STP path selection restrictions. It was expected that this would improve the efficiency of the Building Distribution to Campus Backbone link usage.

The figure presents the result of the initiated simulation. The traffic is perfectly balanced, from the ingress multilayer switch all the way to the destination. The sharing is proportional on pairs of source-destination Building Distribution switches, so all the Building Distribution switches are equally loaded (see the arrows representing the load, dashed for average load and dotted for heavy load). The only remaining suboptimal paths are in the Building Access submodule.

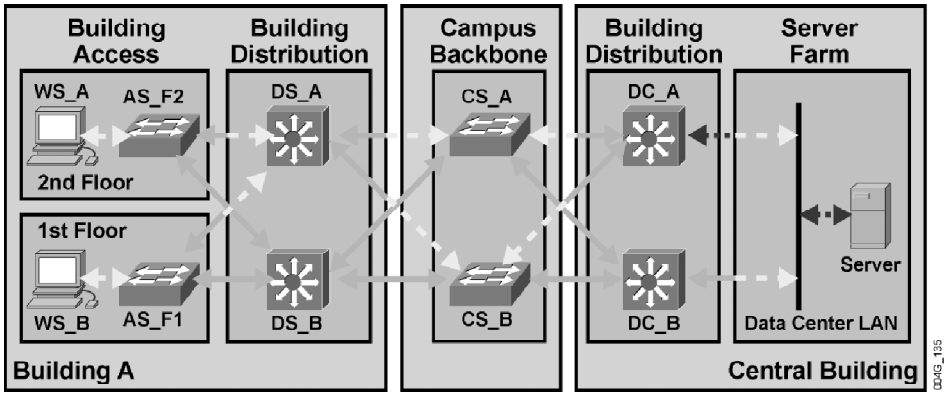

Q5) By examining the resulting graph, you may notice that there is no load sharing in the Building Access submodule of building A. Is this because of the default routing on the workstations using Building Distribution switch DS_A for the primary exit point or because of the attached data link layer switch placing the secondary port in the blocking mode?

Traffic Flow

This graph presents the path (thick lines) taken by the packet originated by the workstation WS_A and destined for the server in the central building. It is obvious that the network resources are more fairly used.

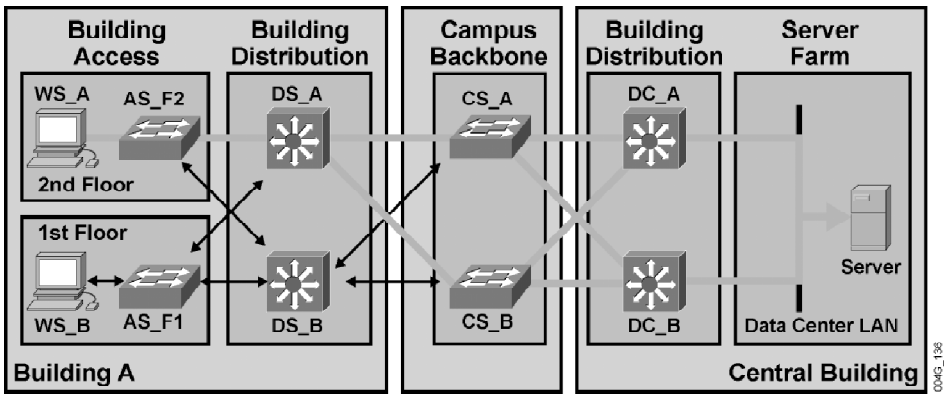

The next graph presents the path (thick lines) used by the packets in the opposite direction, from the server toward the workstation WS_A. The server uses a default routing to send the packets out of the local LAN, thus not utilizing the redundant path at all.

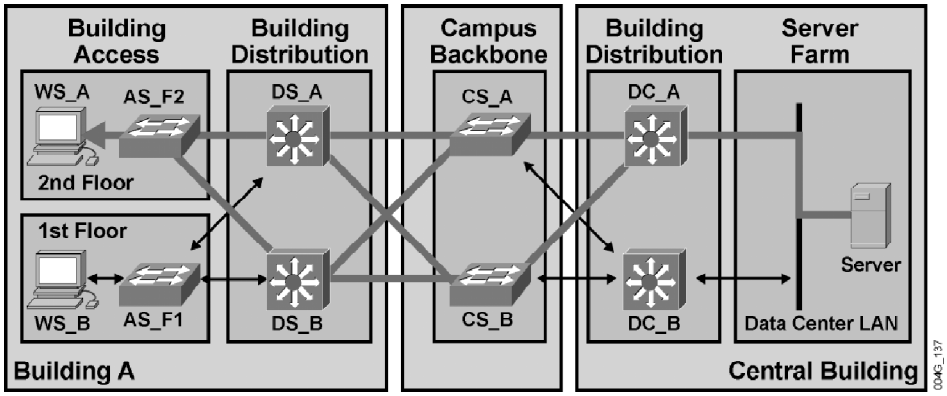

Failure Resilience

The network is now tested against severe failure events such as link loss situations, and the link CS_A – DC_A fails. As expected, the network does not change its behavior under link failure. The load balancing from the ingress multilayer switch to the destination is still perfect, although on a reduced topology. The load distribution ratio on DS_A – CS_A versus DS_A – CS_B is 1:2, because the load is shared between Building Distribution next hops (see the figure, thick lines).

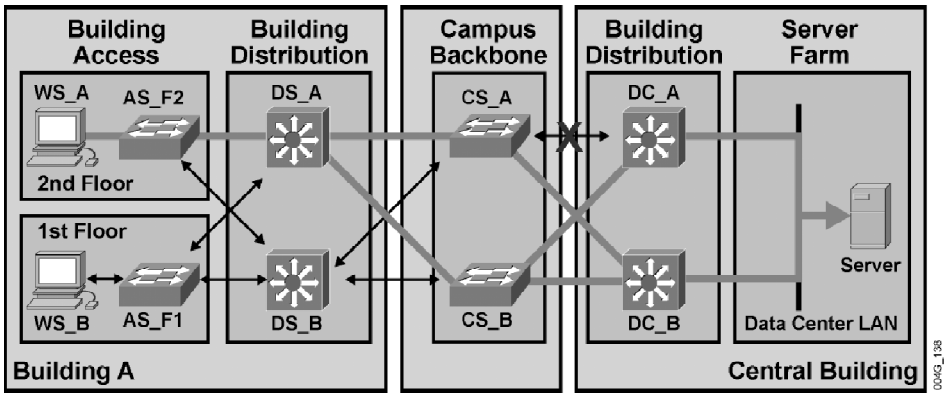

This graph illustrates the return path (thick lines).

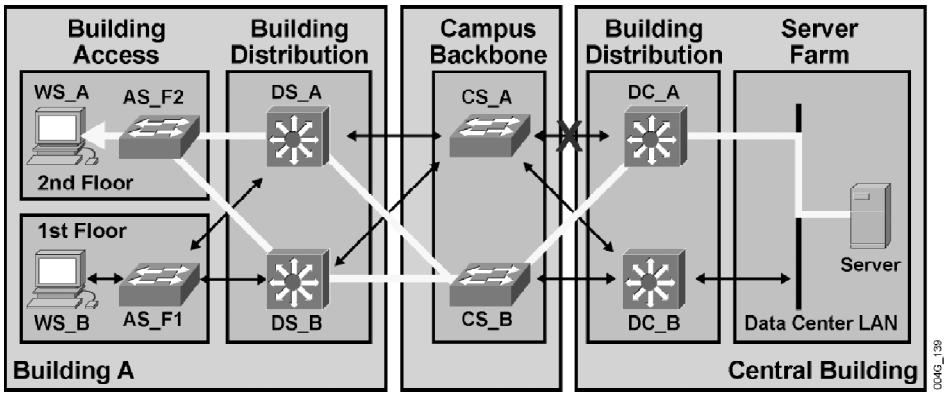

Q6) Why is the return path completely bypassing the CS_A switch?

_____

_____

_____

_____

Multilayer Switching in the Campus Backbone and Building Distribution Submodules

Campus Backbone and Building Distribution switches are now multilayer switches. The simulated load is perfectly shared (see the arrows, dashed for average load, dotted for heavy load) from the Building Distribution submodule across the Campus Backbone submodule on a hop-by-hop basis.

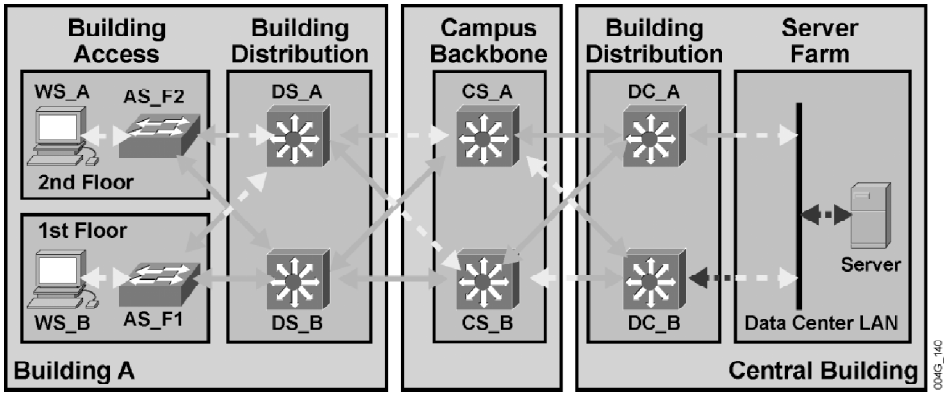

Load Sharing Under Failure

The link CS_B to DC_B failed. This graph illustrates the resulting path (thick lines) taken by the WS_A traffic to the server. The way the load sharing is done is comparable to the previous case with the Building Distribution multilayer switches and data-link switching in the Campus Backbone submodule.

Note: You can only see the actual impact of multilayer switches in the Campus Backbone if you take into account the convergence after the failure.

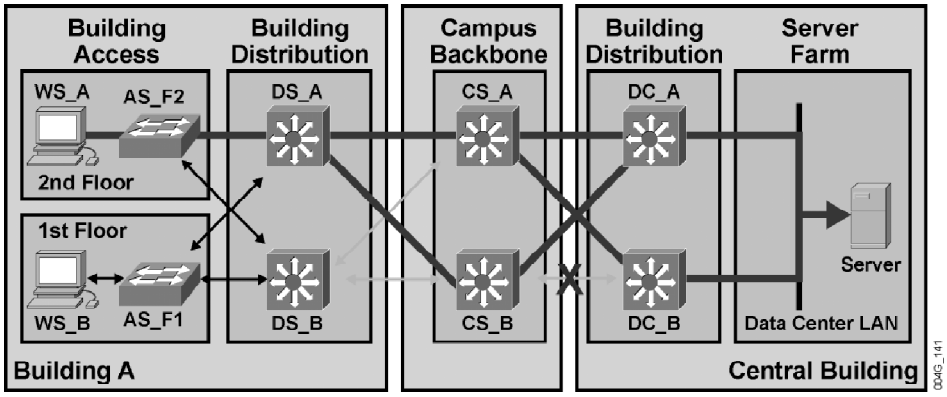

Q7) What is the load distribution ratio on DS_A – CS_A versus DS_A – CS_B link? Explain.

_____

_____

_____

_____

Multilayer Building Access Switch

There is no load sharing in the Building Access LAN or VLAN if the Building Access switch is a data link layer switch and if all the workstations use the same default gateway (Building Distribution switch). To achieve load sharing in the Building Access submodule, you must configure the workstations to use different next hops (DS_A and DS_B in our scenario) for their default routes.

Load Sharing

In this scenario, the AS_F1 Building Access switch has been upgraded to a multilayer switch to achieve more optimal load sharing in the Building Access submodule.

The result of the simulation is illustrated in the figure (thick lines); load sharing from AS_F1 toward DS_A and DS_B is perfect.

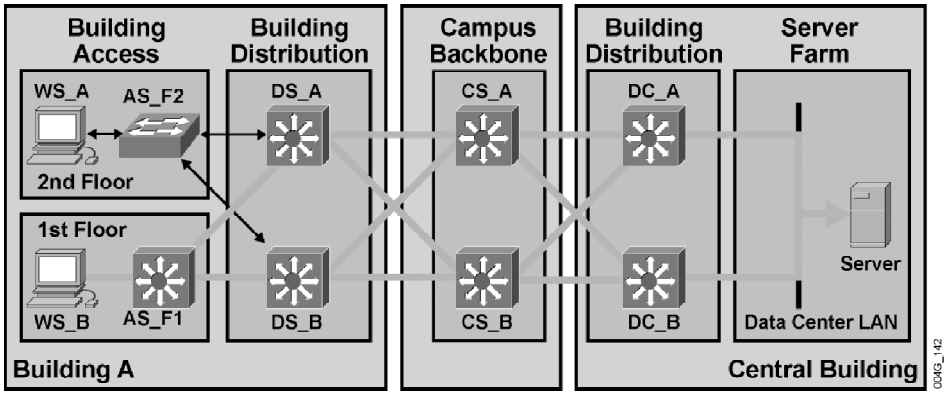

Q8) The workstation WS_B is not running any routing protocol; rather, it depends on the default routing. What is a proper next-hop address?

_____
_____
_____
_____

IP Routing Process on Server

In the last scenario, Open Shortest Path First (OSPF) is configured on the server. The server starts to participate in the campus routing and can rely on the OSPF data to load-share its traffic toward the workstations.

Load Sharing

The result of the server to the WS_A path simulation is presented in the figure (thick lines); the load distribution is achieved from the Building Access submodule to the destination.

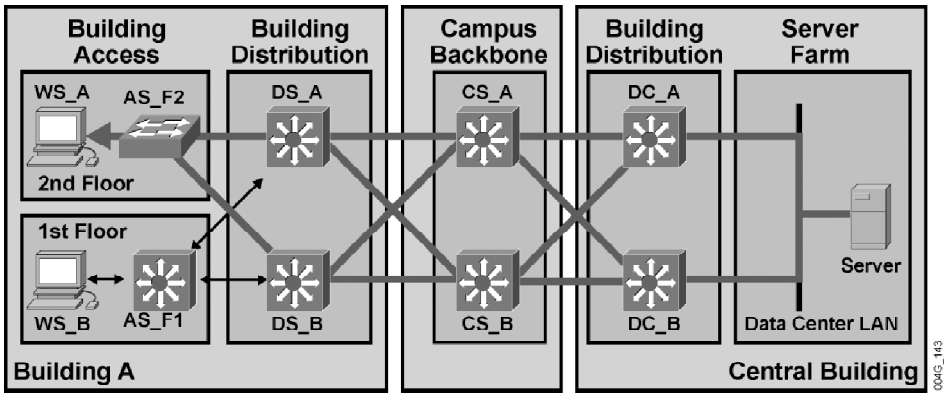

Q9) Running a routing protocol is one way to make the server forward packets to both Building Distribution switches. Can you think of any other option?

_____

_____

_____

_____

Module 4

Designing an Enterprise WAN

Overview

The Enterprise Edge functional area of the Enterprise Composite Network Model provides WAN access to the outside world through the WAN module. To design the Enterprise Edge, you will select WAN technologies and WAN transport media and consider ownership, reliability, and backup issues. In addition, WAN remote access choices include cable and digital subscriber line (DSL) technologies used with virtual private networks.

Note: The module “Evaluating Security Solutions for the Network” discusses the security aspect of WAN interconnections.

Module Objectives

Upon completing this module, you will be able to describe the methodology for designing an enterprise WAN.

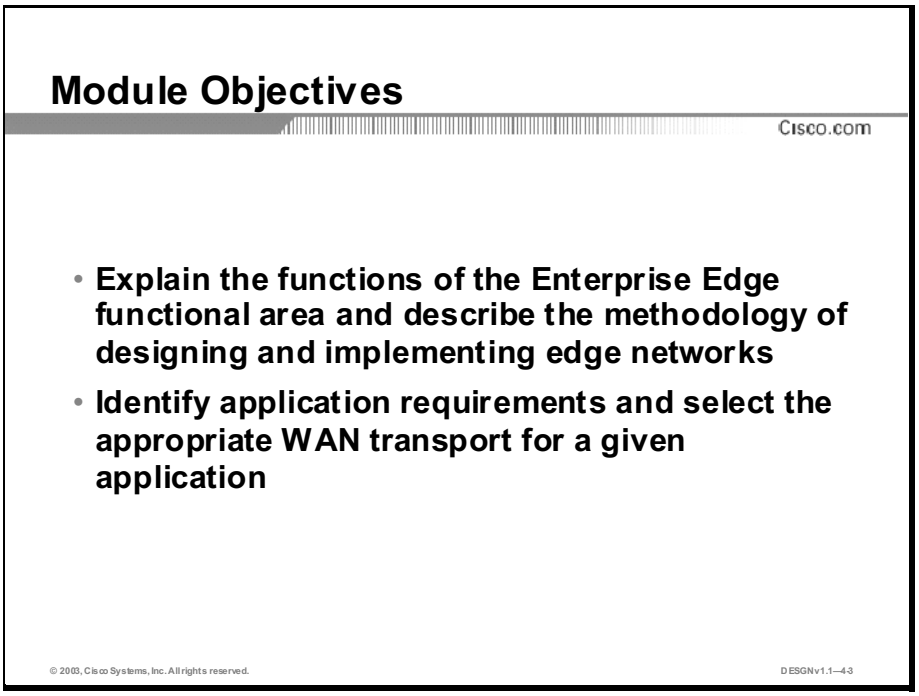

Module Outline

The outline lists the components of this module.

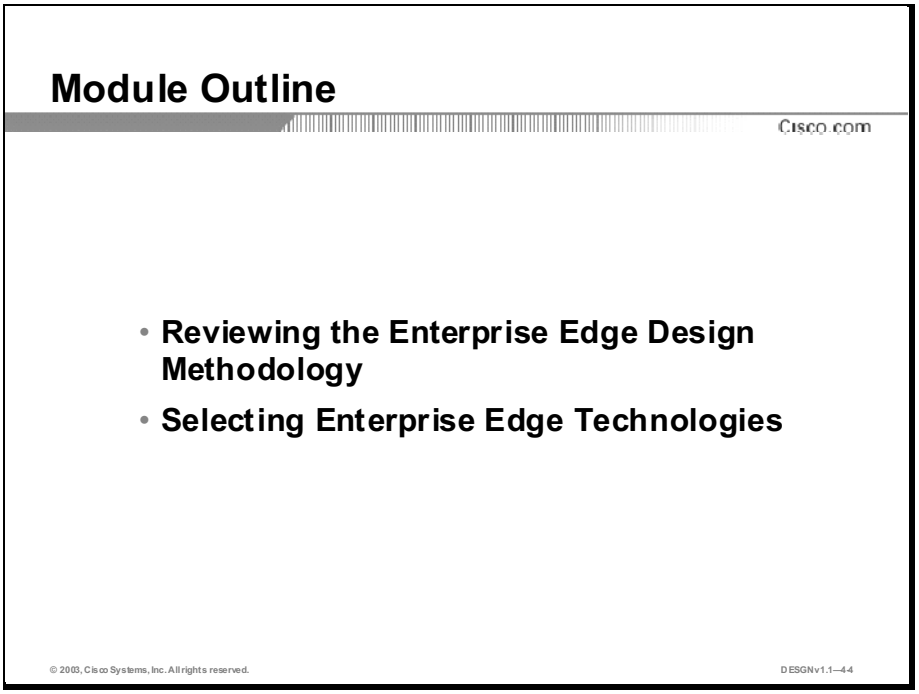

Reviewing the Enterprise Edge Design Methodology

Overview

The Enterprise Edge includes the E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules. This lesson focuses on the design of the WAN module. This lesson explains the role of a WAN and the ownership of the medium, such as leased media versus public networks. Technologies such as ISDN, Frame Relay, and ATM form the basis of the WAN. The lesson addresses the steps in a WAN design that help you address a range of requirements in order to achieve a reliable and efficient WAN design. The lesson concludes with a general discussion on how the design influences the process of selecting hardware components and software features.

Relevance

Designing the Enterprise Edge is a key network design function. Using a structured methodology will help you design WAN solutions that meet specified needs.

Objectives

Upon completing this lesson, you will be able to explain the functions of the Enterprise Edge network and describe the methodology of designing and implementing edge networks. This includes being able to meet these objectives:

- Describe the functions of the WAN and the technologies commonly used to build a WAN
- Describe the methodology used to identify Enterprise Edge requirements
- Describe the methodology used to select Enterprise Edge hardware components and software features

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in Cisco IOS software, particularly with regard to WANs

Outline

The outline lists the topics included in this lesson.

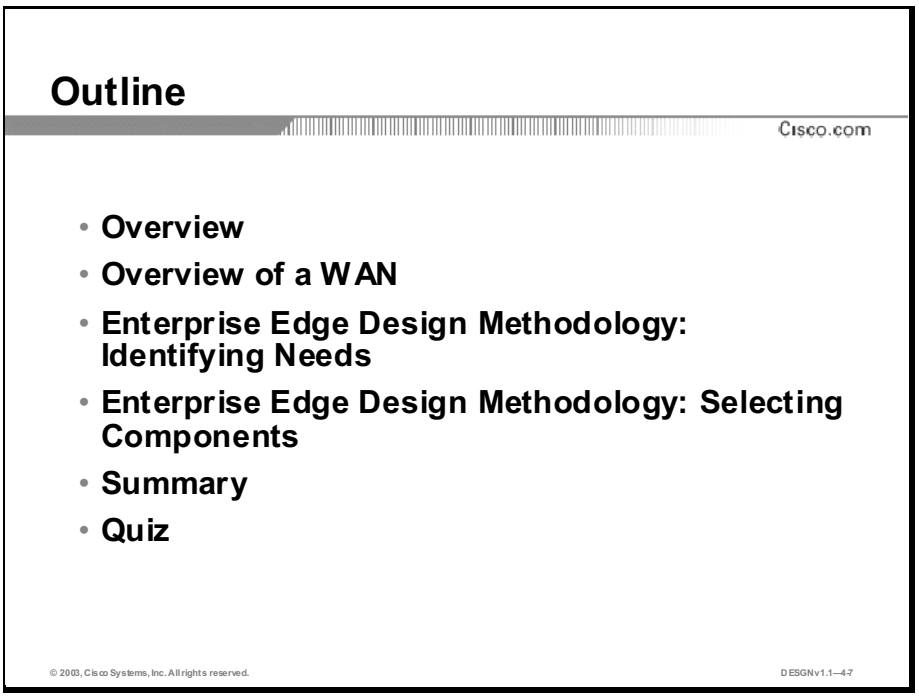

Overview of a WAN

A WAN is a data communications network that covers a relatively broad geographic area and most often uses the transmission facilities provided by service providers (carriers), such as telephone companies. This topic describes the primary objectives of WAN design.

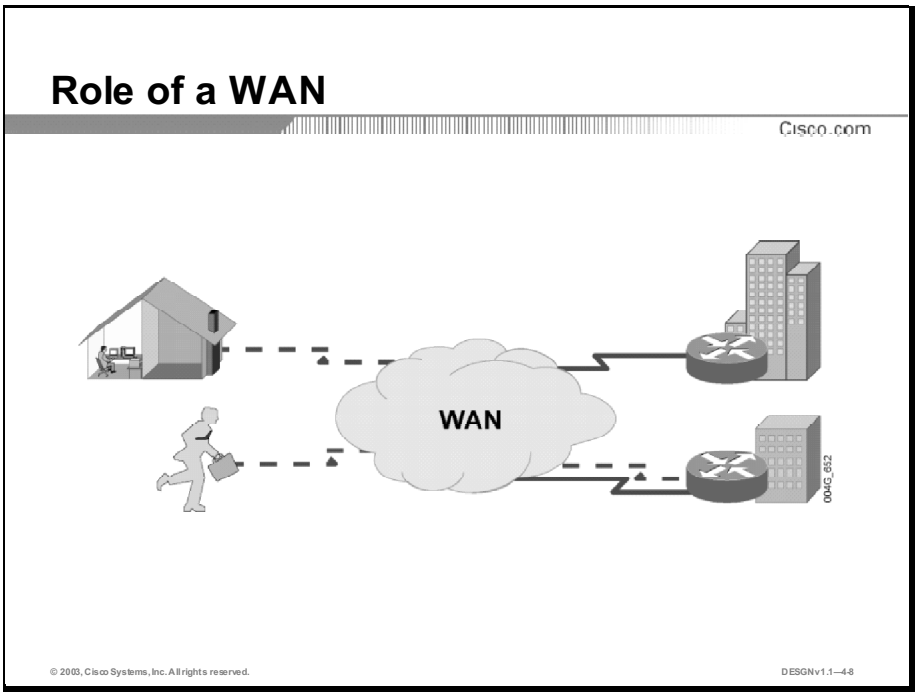

Switches that connect the WAN links are devices that relay information through the WAN and enable the services provided by the WAN. A network provider often charges users fees, called tariffs, for the services provided by the WAN. Therefore, WAN communication is often known as a service.

Designing a network can be a challenging task. The first step is to understand the networking requirements, driven by two primary goals:

- **Application availability:** Networks carry application information between computers. If the applications are not available to network users, the network is failing to achieve its design objectives.
- **Cost of investment and usage:** WAN designs are always subject to budget limitations. Selecting the right type of WAN technology is critical in providing reliable services for end-user applications in a cost-effective and efficient manner.

These are the objectives of an effective WAN design:

- A well-designed WAN must reflect an organization's goals, characteristics, and policies.
- The selected technology should be sufficient for current and (to some extent) future application requirements.
- The associated costs of investment and usage should stay within the budget limits.

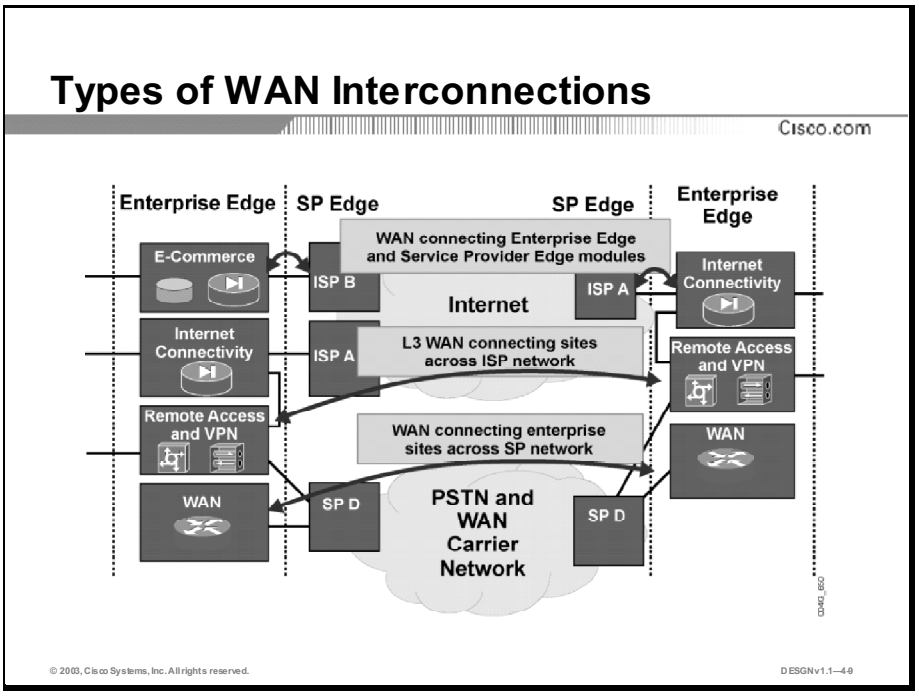

The figure illustrates the position of a WAN that connects the Enterprise Edge modules with the outside world, represented by the service provider network. Typically, the intent is to provide these connections:

- Connectivity between the Enterprise Edge and the Service Provider Edge modules
- L3 connectivity between enterprise sites across the Internet service provider (ISP) network
- Connectivity between enterprise sites across the service provider and Public Switched Telephone Network (PSTN) carrier network

Note: When seeking a WAN network solution, the service provider's offerings will often limit the choices available to you. Review the service provider's offerings before beginning your WAN design.

Enterprise Edge Design Methodology: Identifying Needs

When planning and designing the Enterprise Edge using the planning, design, implementation, operation, and optimization (PDIOO) methodology, you will analyze organization requirements, characterize the existing network, and, based on this information, design the topology and network solutions. This topic describes the WAN design methodology for identifying network needs.

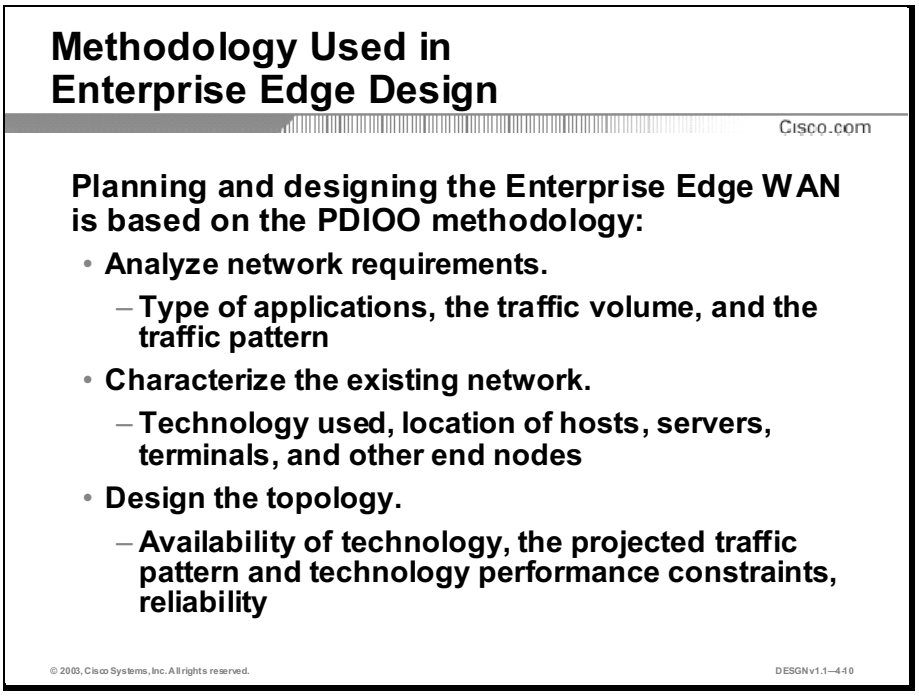

Step 1: Analyze Network Requirements

The initial step in the design methodology is to analyze the requirements of the network and its users. Network user needs and applications change constantly in response to changing business conditions and to changing technology. For example, as more voice- and video-based network applications become available, more network bandwidth intensity is needed.

Step 2: Characterize the Existing Network

The second step is to analyze the existing networking infrastructure and its capacity to migrate to a more appropriate design. Using the physical description of the network, you should evaluate the possibility of extending the network to support new sites and new features. For example, the integration of data and voice often requires considerable changes in the network. In this case, a detailed evaluation of current equipment, services, and topologies is important.

Step 3: Design the Topology and Network Solutions

The final step in the design methodology is to develop the overall network topology and its appropriate services. You will consider the projected traffic pattern, technology performance constraints, and network reliability. The design document should describe a set of discrete functions that the Enterprise Edge modules perform. The document should also describe the expected level of service provided by each selected technology, based on the services that a service provider offers.

Note: WAN connections are mainly characterized by the cost of renting transmission media (wire or fiber) from a service provider to connect two or more sites together. Because organizations often lease the WAN infrastructure from a service provider, WAN designs must optimize the cost of bandwidth and bandwidth efficiency.

A network design should be adaptable to include future technologies and should not include any design elements that limit the adoption of new technologies as they become available. You need to balance this consideration with the issue of cost-effectiveness throughout a network design and implementation. For example, many new internetworks are rapidly adopting Voice over IP (VoIP). Network designs should support VoIP without requiring a substantial upgrade by provisioning hardware and software that have future-proofed options for expansion and upgradability.

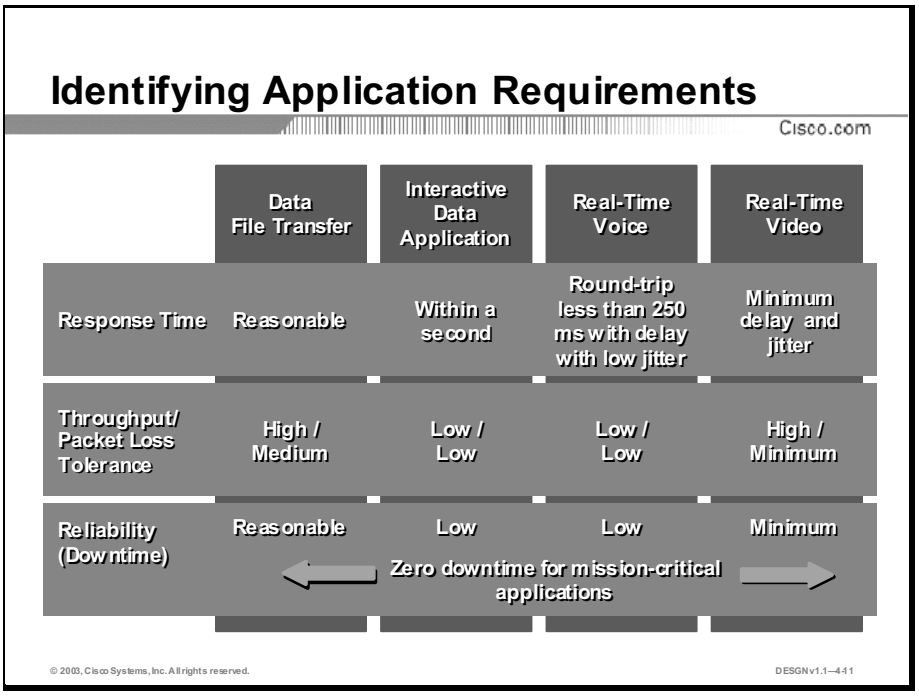

Most users want application availability in their networks. The chief components of application availability are response time, throughput, and reliability.

Response Time

Response time is the time between a user request and a response from the host system. Users accept response times up to some limit, at which point user satisfaction declines. Applications in which a fast response time is considered critical include interactive online services, such as point-of-sale machines.

Note: Voice and video applications use the terms “delay” and “jitter,” respectively, to express the responsiveness of the line and the variation of the delays.

Throughput

In data transmission, throughput is the amount of data moved successfully from one place to another in a given time period. Applications that put high-volume traffic onto the network have a high impact on throughput. In general, throughput-intensive applications involve file-transfer activities. For the most part, throughput-intensive applications have low response-time requirements, so you can schedule them when response-time sensitive traffic is low (for example, after normal work hours).

Reliability

Although reliability is always important, some applications have requirements that exceed typical needs. Organizations that require nearly 100 percent uptime for critical applications are financial services, securities exchanges, and emergency, police, and military operations. These organizations require a high level of hardware and topological redundancy. Determining the cost of any downtime is essential to identify the relative importance of the reliability of the network.

Determing the Maximum Offered Traffic

Cisco.com

- **WAN resources have finite capacity.**
- **End users require minimum response times.**
- **Network managers require maximum link utilization.**

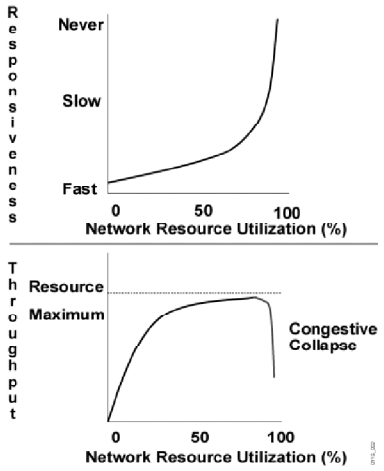

Response time is a problem only for users. It does not matter directly to the network manager if the query results are returned 120 ms sooner rather than later. Response time is a measure of usability for users. Users perceive the data processing experience in terms of how quickly the screen updates. They view data processing in terms of response time, not link utilization.

The figure illustrates the response time and link utilization. The response time increases with the offered traffic until it becomes unacceptable to the end user. Similarly, the link utilization increases with the offered traffic until the link becomes saturated. The goal of the designer is to determine the maximum offered traffic that is acceptable to both the end user and the network manager. Planning a WAN capacity increase should begin early, usually at about 50 percent link utilization. Additional bandwidth purchases should start at 60 percent utilization. 75 percent link utilization means a WAN capacity increase is already urgently needed.

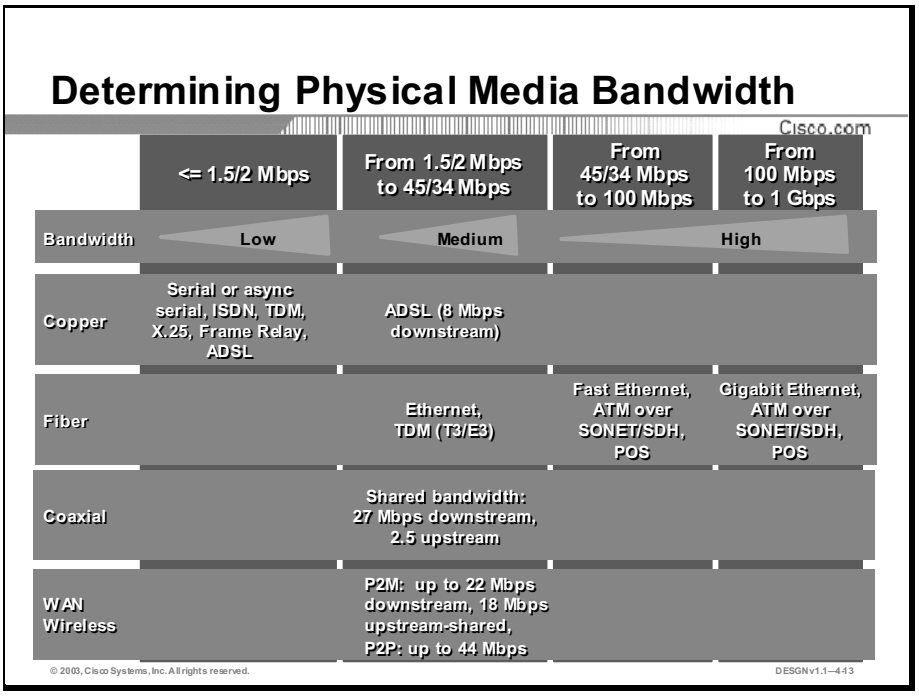

A major WAN design challenge is to provide sufficient bandwidth. The figure shows the range of bandwidths that each supports. Bandwidth is directly proportional to the amount of data transmitted or received per unit of time. In a qualitative sense, bandwidth is proportional to the complexity of the data for a given level of system performance. For example, it takes more bandwidth to download a photograph in one second than it takes to download a page of text in one second. Large sound files, computer programs, and animated videos require even more bandwidth for acceptable system performance.

Bandwidth is inexpensive in the LAN (equipment purchase and ongoing fees) and connectivity is limited only by hardware and implementation costs. In the WAN, bandwidth has been the overriding cost. Therefore, delay-sensitive traffic such as voice has remained separate from data. New applications and the economics of supporting them are forcing these conventions to change.

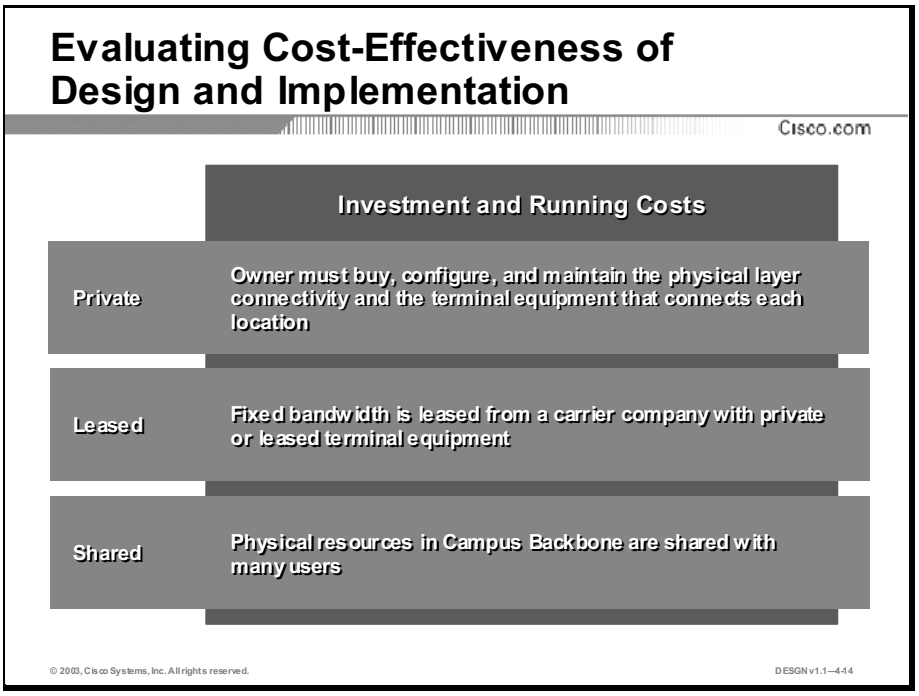

These are the usual fixed costs of the WAN environment:

- Equipment purchases, such as modems, channel service unit and data service units (CSU/DSUs), and router interfaces
- Circuit provisioning
- Network-management tools and platforms

Recurring costs include the service provider's monthly circuit fees and the support and maintenance of the WAN, including any network management center personnel.

Private, Leased, and Shared Ownership

From the ownership perspective, lines are divided into three broad categories:

- **Private:** A private WAN uses private transmission systems to connect distant LANs. The owner of a private WAN must buy, configure, and maintain the physical layer connectivity (copper, fiber, wireless, coaxial) and the terminal equipment required to connect locations. This makes private WANs expensive to build, labor-intensive to maintain, and difficult to reconfigure for constantly changing business needs. The advantages of using a private WAN include higher levels of security and transmission quality.

Note:	Transmission quality is not necessarily improved, nor reliability higher, when the WAN medium and devices are privately owned.

- **Leased:** A leased WAN uses dedicated bandwidth leased from a carrier company with either private or leased terminal equipment. However, the company pays for the allocated bandwidth, whether or not it is used, and operating costs tend to be high.

- **Shared:** A shared WAN shares the physical resources with many users. Carriers offer a variety of circuit or packet-switching transport networks, such as ATM or Frame Relay, for user traffic. Linking LANs and private WANs into a shared network involves a compromise between cost, performance, and security.

Note: Circuits often span regional or national boundaries so that several service providers may handle a connection in the toll network. In cases where multiple service providers handle a connection, the subscriber-owned devices and the leased or shared devices determine the path.

Enterprise Edge Design Methodology: Selecting Components

Once you identify Enterprise Edge requirements, you are ready to select the individual WAN components. This topic discusses selecting Enterprise Edge components.

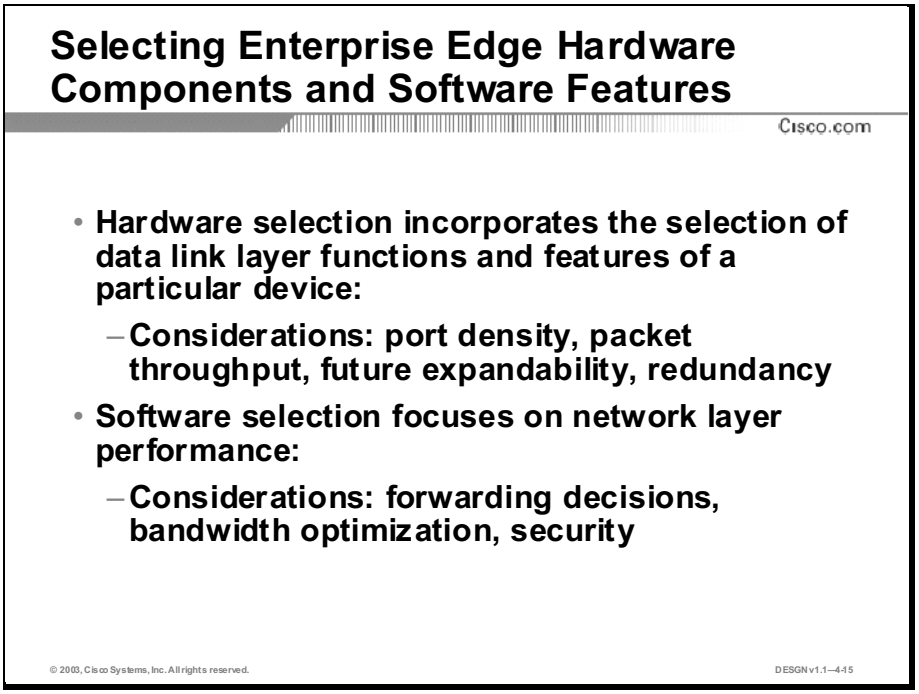

Hardware Selection

When selecting hardware, use vendor documentation to evaluate the WAN hardware components. You will consider the function and features of the particular devices, including their port densities, packet throughput, expandability capabilities, and readiness to provide redundant connections.

Software Selection

The next step is to select the appropriate Cisco IOS software features. Features such as network layer forwarding mechanisms support traffic forwarding, throughput and bandwidth optimization, and security (access lists).

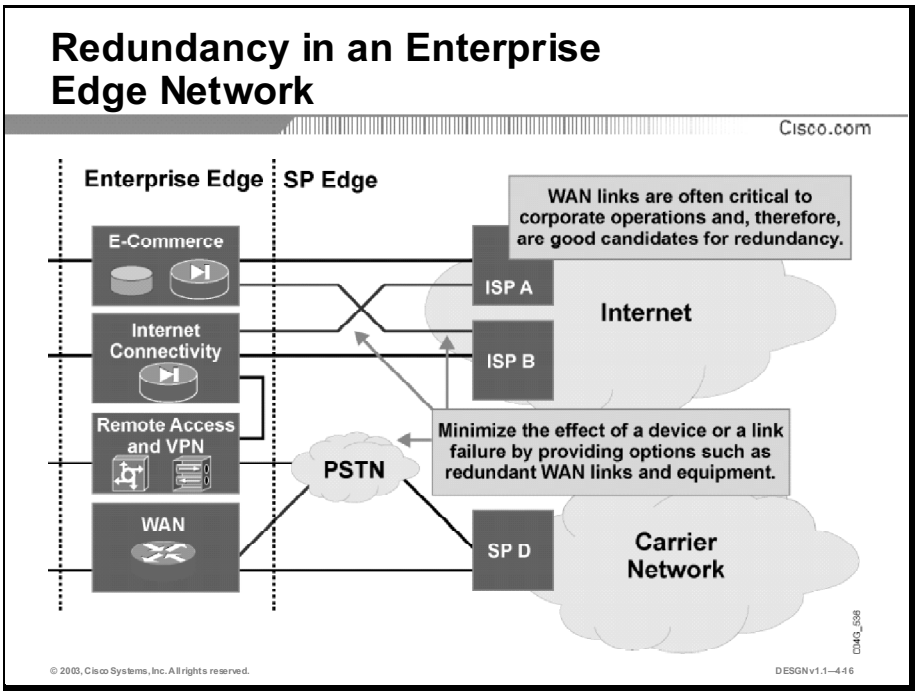

WAN links connect geographically dispersed sites. WAN links are relatively unreliable, and often much slower than the LANs they connect. The combination of uncertain reliability, lack of speed, and high importance makes the WAN link a good candidate for redundancy.

Because WAN links interconnect remote sites, they are critical components of the network and redundant media is often used. You should provision backup links so they become active when a primary link fails or becomes congested. Backup links often use different technologies; for example, leased lines are used with backup ISDN circuits.

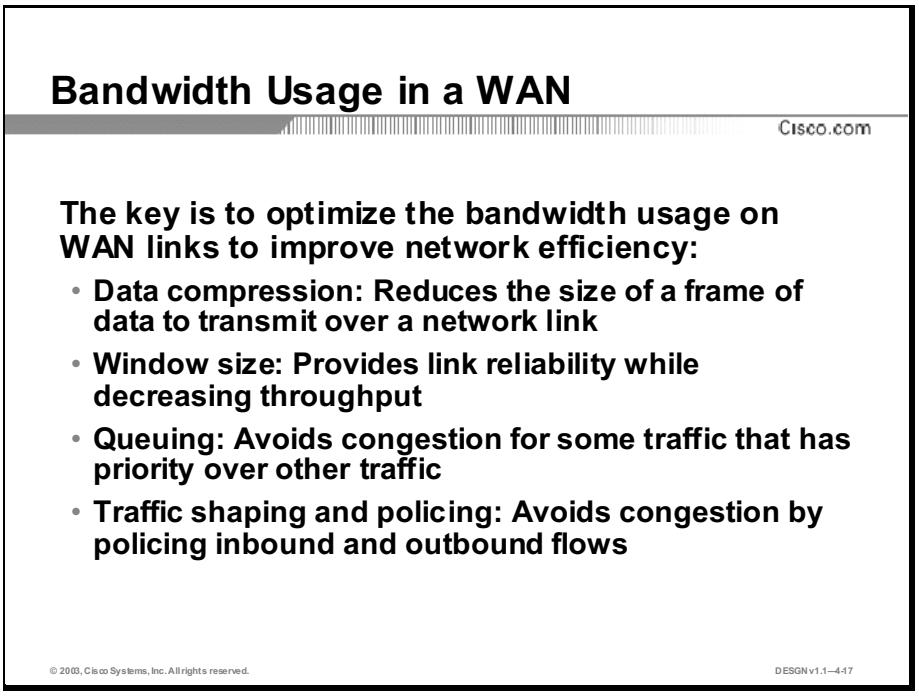

Transmitting data over a WAN is expensive. Therefore, you will specify data compression, queuing, limiting access rate, or traffic shaping to optimize bandwidth usage and improve overall efficiency.

Data Compression

Compression reduces data size to save transmission time. It enhances the efficient use of the available WAN bandwidth, which is often limited and congested.

Hardware-assisted data compression achieves the same goal as software-based data compression but accelerates compression rates by offloading the task from the main CPU to specialized compression circuits. In other words, compression is implemented in the compression hardware installed in a system slot.

IOS software supports these data software compression products:

- FRF.9 Frame Relay payload compression
- Link Access Procedure, Balanced (LAPB) payload compression using Lempel-Ziv Stac (LZS), commonly referred to as just “STAC”
- High-Level Data Link Control (HDLC) using LZS
- X.25 payload compression of encapsulated traffic
- PPP using LZS, Predictor
- Van Jacobson header compression for TCP/IP
- *Microsoft Point-to-Point Compression (MPCC)

Window Size

The window size specifies the maximum number of frames that are transmitted without receiving an acknowledgment. Acknowledgment procedures are particularly important in a protocol layer that provides reliability, such as hop-by-hop acknowledgment in a reliable link protocol or end-to-end acknowledgment in a transport protocol.

The current window is defined as the amount of data that may be sent without acknowledgement, which is always less than or equal to the window size. This form of data acknowledgement provides a means in which the network is “self-clocked” so data steadily flows between the two endpoints of the connection. For example, if the TCP window size is set to 8192, the sender must stop after sending 8192 bytes if no acknowledgment comes from the receiver. This may be unacceptable for long WAN links with significant delays. In these cases, you can adjust the window size to a higher value. Frequent retransmissions are a risk, however, because of links with high error rates which reduce the throughput dramatically.

Note: Adjustable windows and equipment that can adapt to line conditions are strongly recommended.

Queuing, Traffic Shaping, and Policing the Access Rate

Cisco Systems has developed quality of service (QoS) techniques to mitigate temporary congestion and provide preferential treatment for critical applications. QoS mechanisms, such as queuing, policing (limiting) the access rate, and traffic shaping enable network operators to deploy and operate large-scale networks that handle both bandwidth-hungry applications such as multimedia, and web traffic and mission-critical applications such as host-based applications efficiently.

Note: QoS does not create bandwidth. QoS optimizes the use of existing resources.

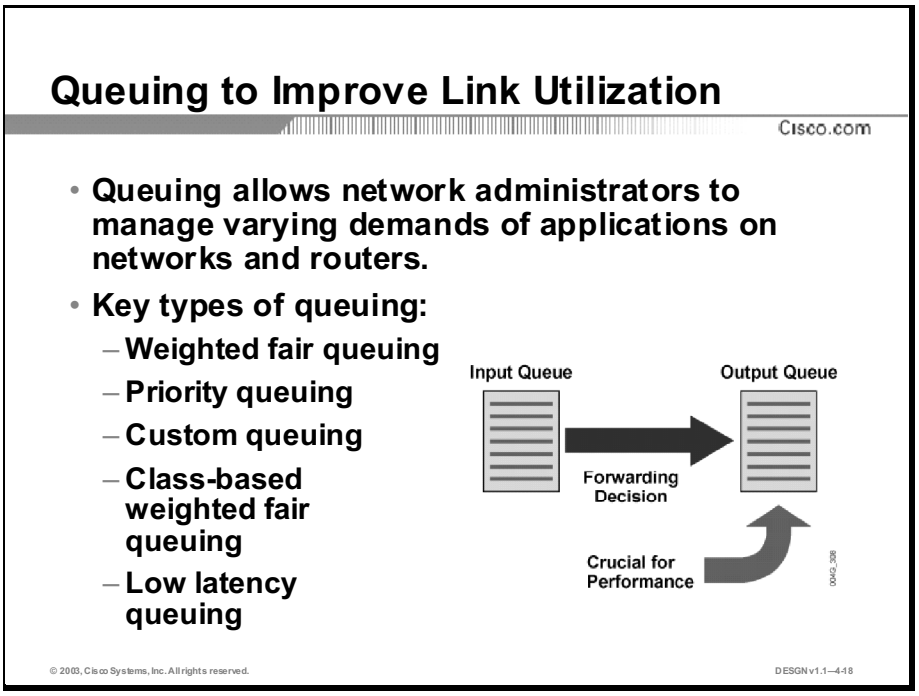

When positioning the role of queuing in networks, the primary issue is the duration of congestion. If WAN links are constantly congested, an organization either requires greater bandwidth or should use compression. Queuing is only required on congested WAN links.

Note:	Queuing occurs at the outbound interface and is appropriate for cases where WAN links are congested from time to time.

There are two types of queues:

- **Hardware queue:** Uses the first-in, first-out (FIFO) strategy, which is necessary for the interface drivers to transmit packets one by one. The hardware queue is sometimes referred to as the transmit queue or TxQ.
- **Software queue:** Schedules packets into the hardware queue based on the QoS requirements, custom queuing (CQ), priority queuing (PQ), weighted fair queuing (WFQ).

Weighted Fair Queuing

WFQ handles problems inherent in queuing schemes using a FIFO system. WFQ ensures that different traffic flows are sorted into separate streams, or conversation sessions, and alternately dispatched. WFQ is the default in the Cisco IOS software for links at or below 2.048 Mbps. Faster links use a hardware FIFO default.

Priority Queuing

Priority queuing is useful for time-sensitive, mission-critical protocols. It establishes four interface output queues that each serve a different priority level.

Custom Queuing

Custom queuing establishes up to 16 interface output queues. When the appropriate number of frames is transmitted from a queue, the transmission window size is reached and the next queue is checked. CQ is a fairer solution for mission-critical applications than PQ because it guarantees some level of service to all traffic.

Class-Based Weighted Fair Queuing

Class-based weighted fair queuing (CBWFQ) extends the standard weighted fair queuing (WFQ) functionality to provide support for user-defined traffic classes. For CBWFQ, you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class.

Once a class has been defined according to its match criteria, you can assign it characteristics. To characterize a class, you assign it bandwidth, weight, and maximum packet limit. The bandwidth assigned to a class is the guaranteed bandwidth delivered to the class during congestion.

To characterize a class, you also specify the queue limit for that class, which is the maximum number of packets allowed to accumulate in the queue for the class. Packets belonging to a class are subject to the bandwidth and queue limits that characterize the class.

Low Latency Queuing

Low latency queuing (LLQ) brings strict priority queuing to CBWFQ. Strict priority queuing allows delay-sensitive data such as voice to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.

Without LLQ, CBWFQ provides weighted fair queuing based on defined classes with no strict priority queue available for real-time traffic. CBWFQ allows you to define traffic classes and then assign characteristics to that class. For example, you can designate the minimum bandwidth delivered to the class during congestion.

For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth assigned to the class during configuration. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced fairly based on weight; no class of packets may be granted strict priority. This scheme poses problems for voice traffic that is largely intolerant of delay, and especially for voice traffic intolerant of variation in delay.

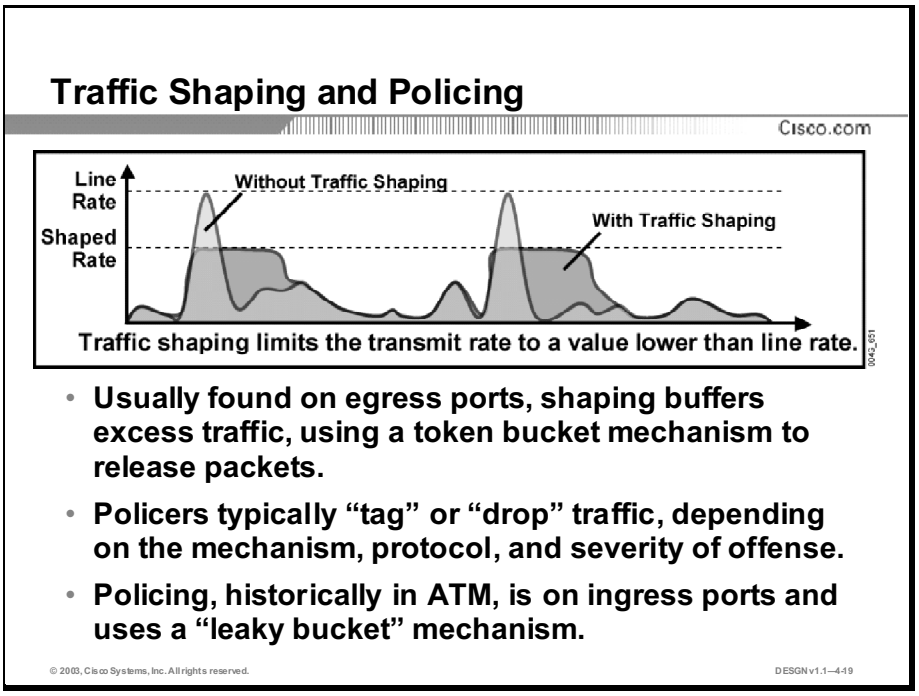

Traffic shaping and traffic policing, also referred to as committed access rate (CAR), are similar mechanisms that both inspect traffic and then take an action based on the characteristics of that traffic, usually that the traffic is over or under a given rate, or based on some bits in the headers, such as the differentiated services code point (DSCP) or IP precedence.

Policing either discards the packet or modifies some aspect of it, such as its IP precedence, when the policing agent determines that the packet meets given criteria. By comparison, traffic shaping adjusts the transmission rate of packets that match certain criteria. Traffic shaping holds packets in a buffer and releases them based on a preconfigured rate. Traffic shaping is only available on traffic leaving an interface.

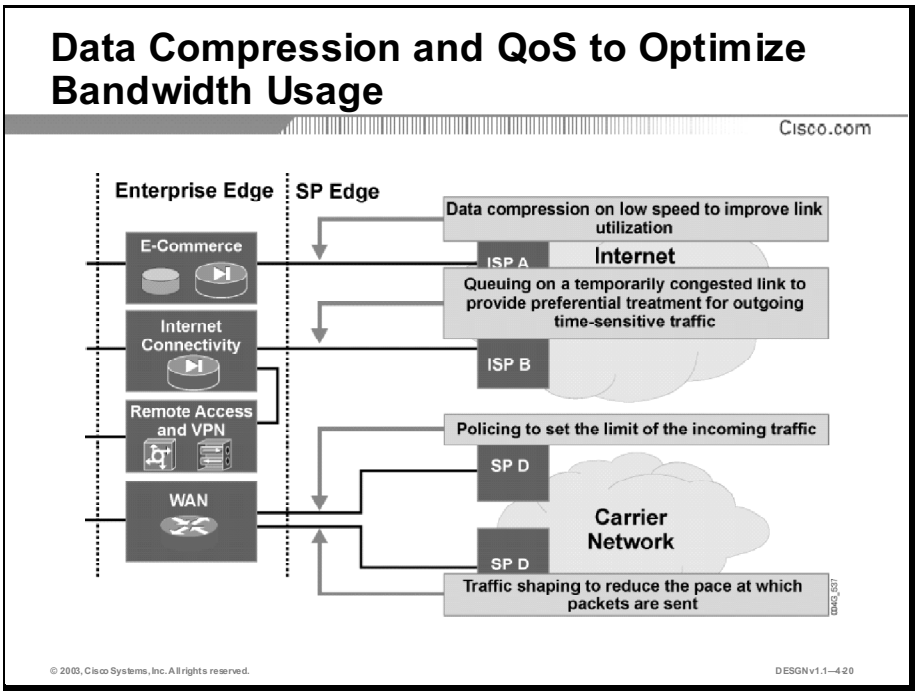

Compression allows higher throughput because it reduces the packet size, increasing the amount of data that can traverse a transmission resource in a given time period. Payload compression is performed on data link layer frames and, thereby, compresses the entire network layer packet.

An enterprise's policy management scheme could deem the traffic generated by a particular resource, such as voice, to be "first-class" traffic so that it receives a top priority marking. Other traffic, such as data, could drop to a lower priority class.

Topologies that have higher-speed links feeding into lower-speed links (such as from a central site to a branch office) often experience bottlenecks at the remote end. Traffic shaping helps eliminate the bottleneck situation by throttling back traffic volume at the source. The most common use of traffic shaping in the enterprise is to smooth the flow of traffic across a single link toward a service provider transport network in order to ensure compliance with the traffic contract. This avoids service provider policing at the receiving end. Shaping reduces the bursty nature of the transmitted data, and is most useful when the contract rate is less than the line rate. You can also use traffic shaping to respond to signaled congestion from the transport network when the traffic rates exceed the contract guarantee.

Summary

This topic summarizes the key points discussed in this lesson.

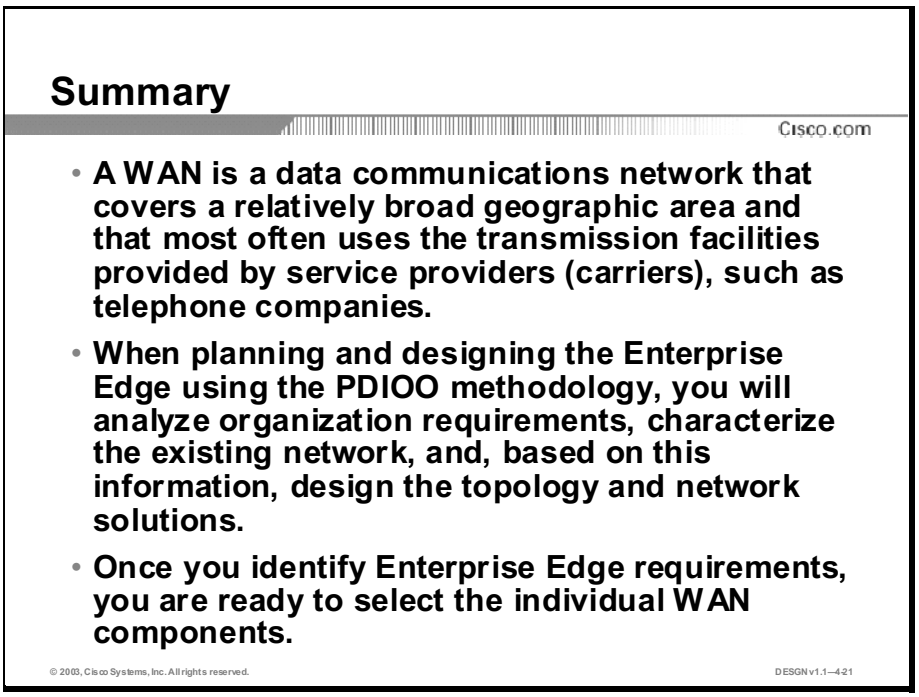

References

For additional information, refer to these resources:

- http://www.cisco.com/univercd
- http://www.whatis.com/

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match each element of technology to its transmission medium.

A) ADSL

B) VDSL

C) cable

_____ 1. copper twisted pair

_____ 2. coaxial cable

Q2) Which three technologies are suitable for connections in excess of 100 Mbps? (Choose three.)

A) Fast Ethernet

B) ISDN

C) Wireless

D) Gigabit Ethernet

E) packet over SONET/SDH (POS)

F) ATM over SONET/SDH

G) ADSL

Q3) Which two tasks should you complete to improve network efficiency? (Choose two.)

A) Apply compression on an entire-packet, a header-only, or a payload-only basis.

B) Apply compression to enable higher throughput by reducing the packet size.

C) Implement backup links with the same technology in use for the primary connection.

D) Apply queuing on the outbound interfaces in cases of constant WAN link congestion.

Quiz Answer Key

Q1) 1=A, B
2=C

Relates to: Overview of a WAN

Q2) D, E, F

Relates to: Enterprise Edge Design Methodology: Identifying Needs

Q3) A, B

Relates to: Enterprise Edge Design Methodology: Selecting Components

Selecting Enterprise Edge Technologies

Overview

Many WAN technologies exist today, and new technologies are constantly emerging. In general, the most appropriate WAN selection results in high efficiency and leads to user satisfaction. You must be aware of all possible WAN design choices while considering enterprise requirements. This lesson describes the characteristics of the most commonly deployed WAN technologies, including Virtual Private Networks (VPNs), which are built on top of public networks.

Relevance

An effective network design for the Enterprise Edge requires knowledge of the technologies available to implement the edge network.

Objectives

Upon completing this lesson, you will be able to identify application requirements and select the appropriate WAN transport for a given application. This includes being able to meet these objectives:

- Design the classic WAN module to meet specified needs
- Design remote-access networks to meet specified needs
- Select service provider offerings to create a network between dispersed enterprise sites
- Design VPNs using public networks
- Design WAN backup strategies to ensure network reliability
- Design WAN backup strategies using the Internet

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including the basic functions implemented in IOS software, particularly with regard to WAN networks

Outline

The outline lists the topics included in this lesson.

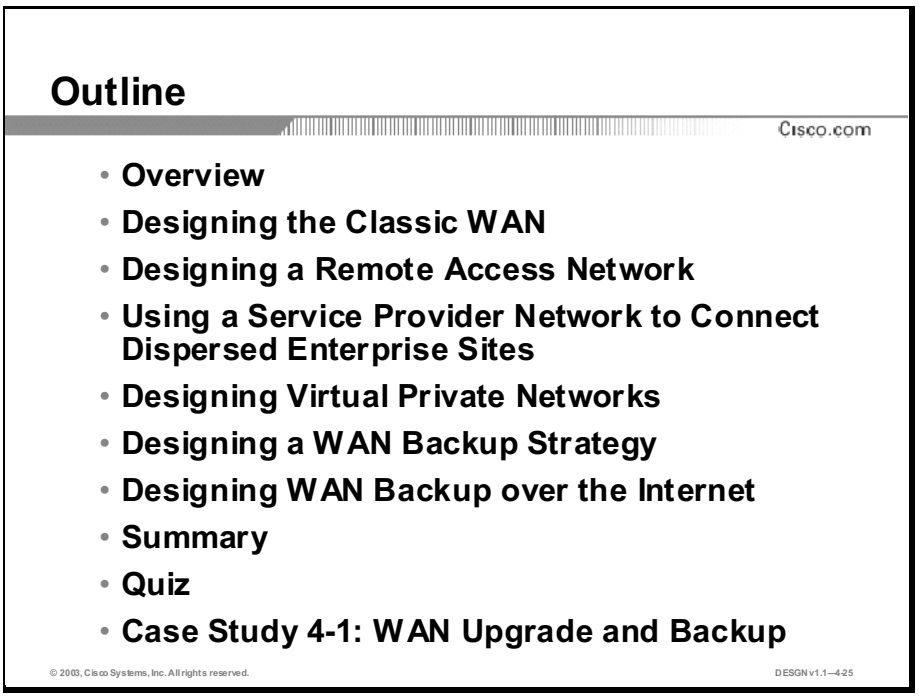

Designing the Classic WAN

Each WAN design is based on application requirements, the geography, and the available service provider offerings. This topic discusses designing the classic WAN.

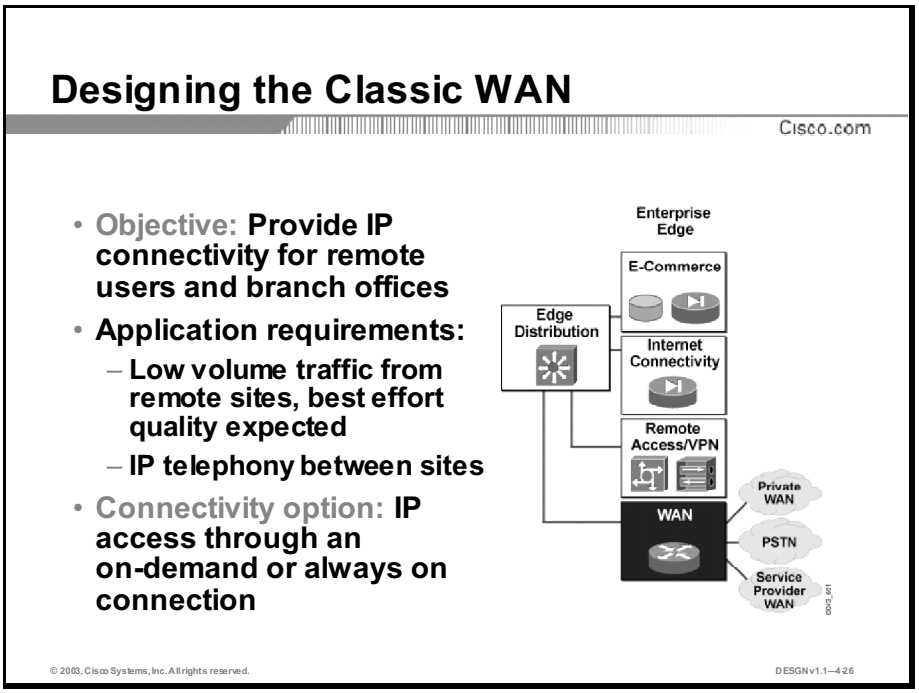

The primary requirements for the classic WAN are:

- Moderate-volume IP traffic from remote users and branch offices
- Best-effort quality expected

Different WAN technologies can functionally coexist. Service providers typically support more than one technology.

Note:	In general, the Internet service providers (ISPs) do not own WAN resources and lease them from WAN service providers (carriers).

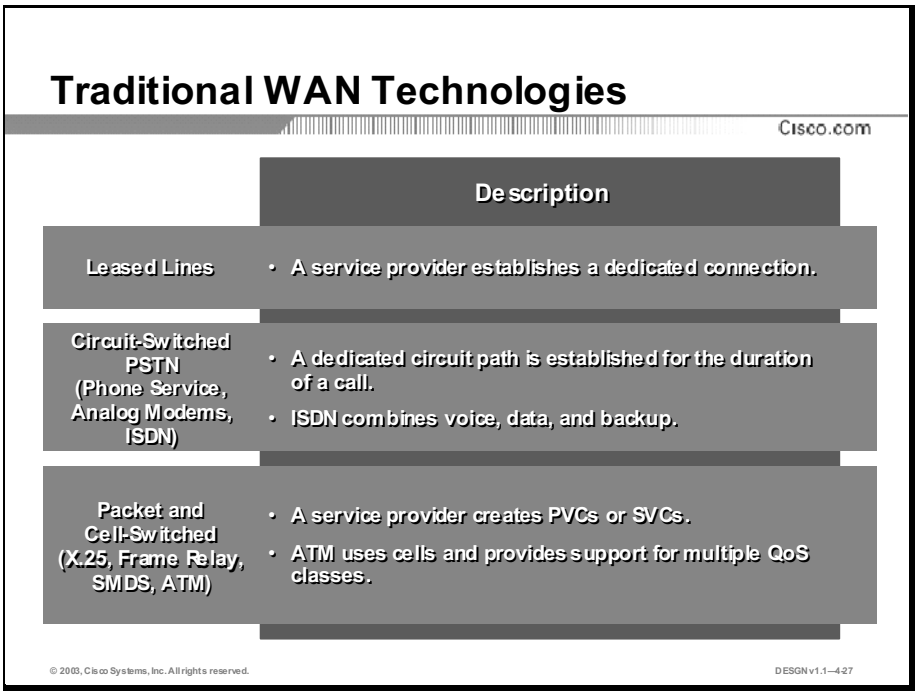

One of the main issues in WAN connections is the selection of the appropriate physical WAN technology. Options can include:

- **Leased lines:** Point-to-point connections that are reserved for transmissions rather than used only when transmission is required. The carrier establishes the connection to dedicate a physical wire or to delegate a channel using frequency modulation (FM) or time-division multiplexing (TDM). Usually, leased-line connections utilize synchronous transmission.
- **Circuit-switched networks:** A type of network that, for the duration of the connection, obtains and dedicates a physical path to a single connection between two endpoints in the network. Ordinary voice telephone service over the Public Switched Telephone Network (PSTN) is circuit-switched. The telephone company reserves a specific physical path to the number being called for the duration of the call. During that time, no one else can use the physical lines involved. Circuit-switched examples are asynchronous serial and ISDN.
- **Packet and cell-switched networks:** A carrier creates permanent virtual circuits (PVCs) or switched virtual circuits (SVCs) that deliver packets among different sites. Users share common carrier resources and can use different paths through the WAN. This allows the carrier to use its infrastructure more efficiently than with leased point-to-point links. Examples of packet-switching networks are X.25, Frame Relay, and Switched Multimegabit Data Service (SMDS).

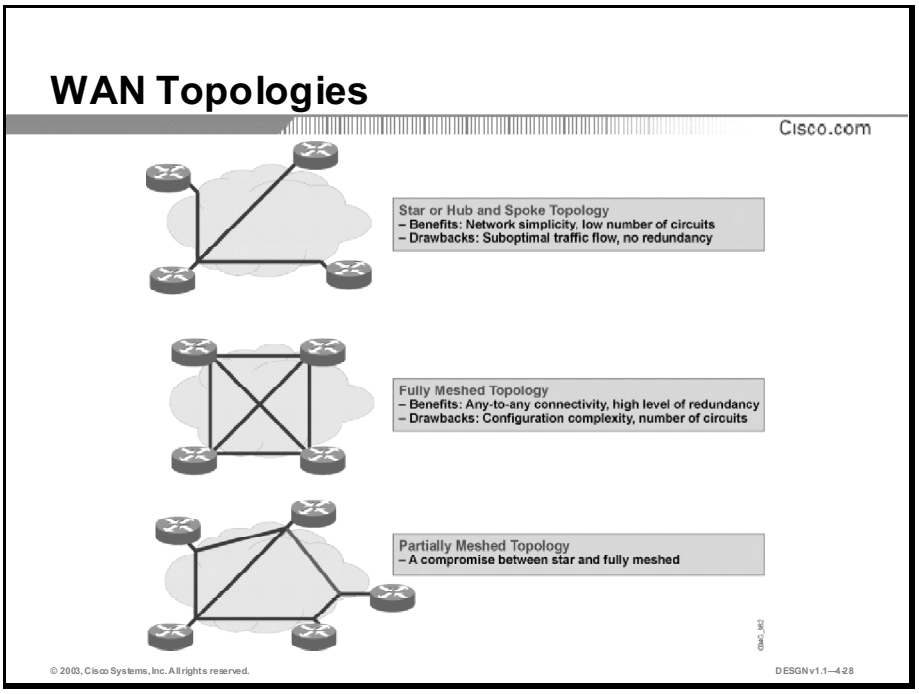

The three basic design approaches for packet-switched networks include star, full-mesh, and partial-mesh topologies.

Star Topology

A star topology features a single hub (central router) that provides access from remote networks into a core router. All communication between networks goes through the core router. The advantages of a star approach are simplified management and minimized tariff costs. However, the disadvantages are significant:

- The central router (hub) represents a single point of failure.
- The central router limits overall performance for access to centralized resources because it is a single pipe that handles all traffic intended either for the centralized resources or for the other regional routers.
- The topology is not scalable.

Fully Meshed Topology

In a fully meshed topology, each routing node on the periphery of a given packet-switching network has a direct path to every other node on the cloud. The key rationale for creating a fully meshed environment is to provide a high level of redundancy. A fully meshed topology supports is not viable in large packet-switched networks. These are the key issues with a fully meshed topology:

- The large number of virtual circuits required (one for every connection between routers)
- Problems associated with the requirement for large numbers of packet and broadcast replications
- The configuration complexity for routers without multicast support in nonbroadcast environments

Partially Meshed Topology

A partially meshed topology reduces the number of routers within a region that have direct connections to all other nodes in the region. All nodes are not connected to all other nodes. There are many forms of partially meshed topologies. In general, partially meshed approaches provide the best balance for regional topologies based on the number of virtual circuits, redundancy, and performance.

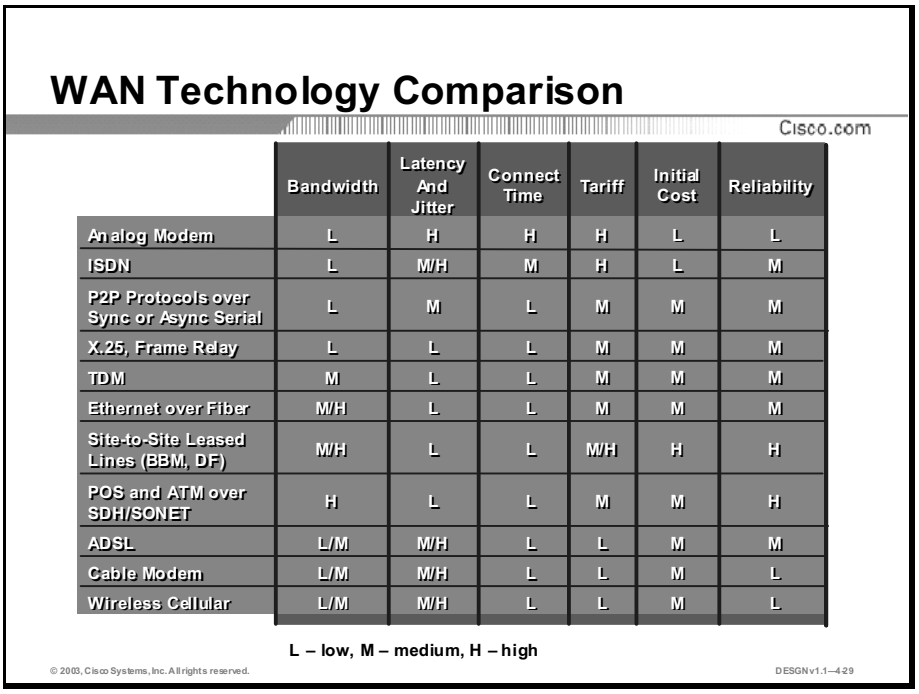

The figure reviews WAN technologies based on the main factors that influence technology selection. The figure provides baseline information to help you compare the performance and features that different technologies offer. Often, the service provider's offerings limit your technology decisions.

Designing a Remote-Access Network

When designing remote-access networks, the type of connection drives the technology selection, such as whether to choose a data link or a network layer connection. By analyzing the application requirements and service provider offerings, you can determine the most suitable of a wide range of remote-access technologies. This topic describes a methodology for designing a remote-access network.

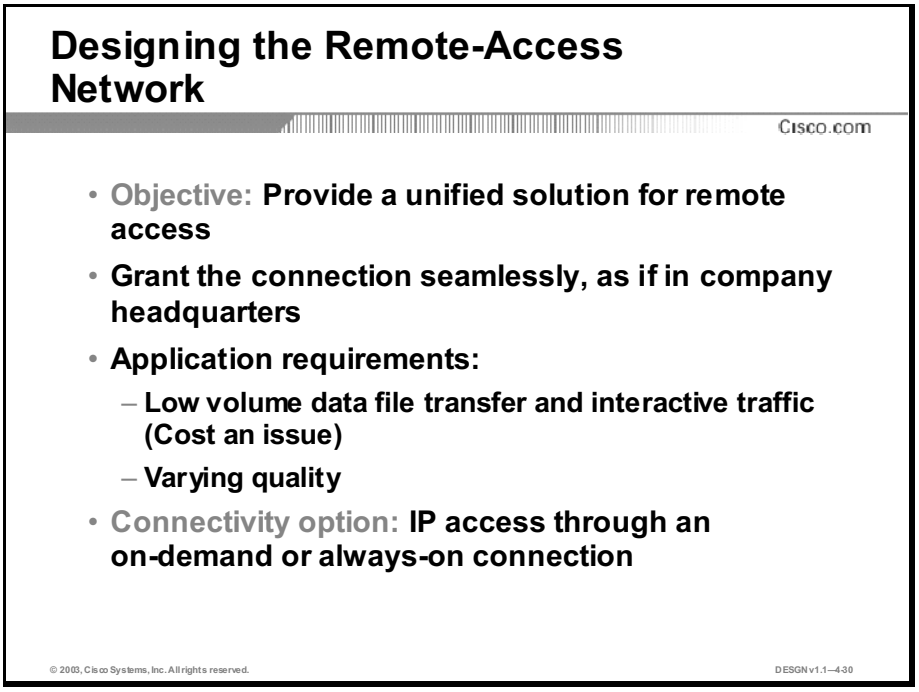

Here is a summary of remote-access requirements:

- Data link layer WAN technology from remote sites to the Enterprise Edge network (consider investment and running costs)
- Low-volume data file transfer and interactive traffic, without any specifics regarding the quality

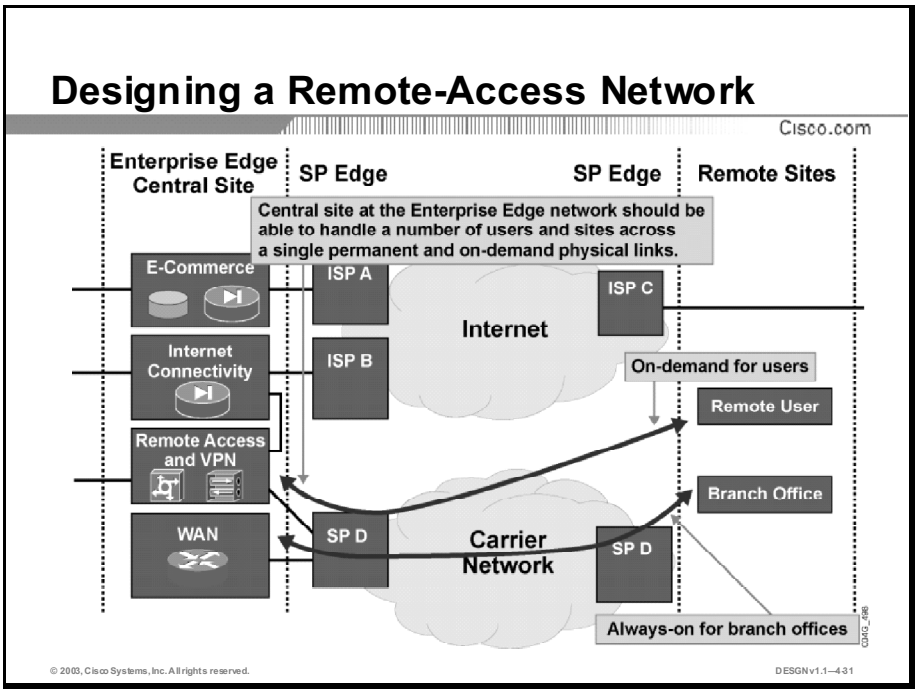

Remote access to the enterprise network is typically provided over permanent or dial-up connections. These are the initial design options:

- On-demand connections for remote users
- Permanent connections for remote branch offices

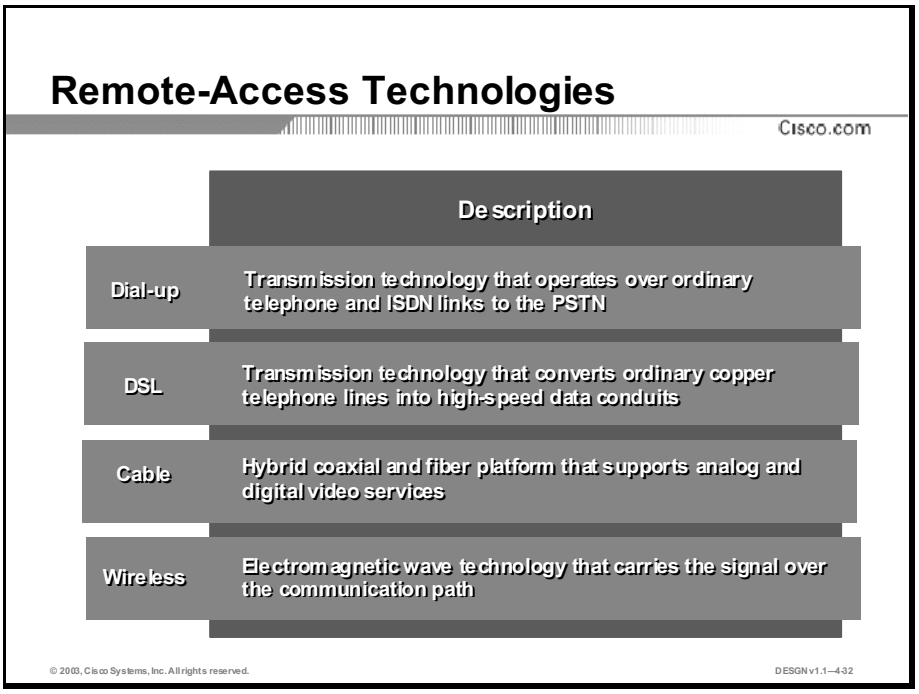

Remote-access technologies include:

- **Dial-up:** A technology that allows remote users to use the PSTN to access remote sites. Users can access the PSTN using ordinary analog connections or ISDN.
- **Digital subscriber line (DSL):** A technology that delivers high bandwidth over traditional telephone copper lines. The term xDSL covers a number of similar yet competing forms of DSL, including these:
  - **ISDN digital subscriber line (IDSL):** Similar to ISDN. The primary difference is that IDSL is always on, can reach speeds of 144 kbps, and is very capable, through compression, of reaching speeds of 512 kbps.
  - **High-data-rate DSL (HDSL):** Mature T1 technology that provides symmetric communications up to 1.54 Mbps. Data travels over two pairs of wires instead of one and does not support PSTN.
  - **HDSL-2:** Full-rate-only symmetric service that exists over a single twisted-pair wire. HDSL-2 was conceived specifically to provide spectral compatibility with asymmetric digital subscriber line (ADSL). This coexistence with ADSL is crucial, and in this regard the technology is superior to SDSL.
  - **Symmetric High Bit Rate Digital Subscriber Line (g.shdsl):** Combines the best of SDSL and HDSL-2. The standard defines multirates, such as SDSL, but provides the spectral compatibility of HDSL-2.
  - **Very-high-data-rate DSL (VDSL):** Extremely fast asymmetric DSL technology that provides data and PSTN service on a single twisted pair of wires. VDSL is reserved for users in close proximity to a central office.
- **Cable:** A technology for data transport that uses a hybrid of coaxial cable and fiber-optic media over cable distribution systems. This is a good option for environments where cable television is widely deployed.
- **Wireless:** A term used to describe telecommunications in which electromagnetic waves carry the signal. Common examples of wireless equipment include cellular phones and

pagers, global positioning system (GPS), cordless computer peripherals, satellite television, and wireless LANs. Wireless implementations include:

— **Broadband fixed wireless:** Designed to connect two or more networks, typically located n different buildings, at high data rates for data-intensive, line-of-sight applications.

— **Mobile wireless:** This includes cellular applications and others.

— **Wireless LAN:** Developed to meet demand for LAN connections over the air. It is often used in intrabuilding connections.

Note: An alternative to WAN connections is a service provider's IP network that spans remote sites of an enterprise network. Full cooperation at the IP layer between the Enterprise Edge and service provider (SP) network is required for this type of connection. DSL and cable are technologies frequently used for ISP access. This type of network service provides no guarantee of the quality of sessions and is considered a "best effort."

Remote-Access Networking Requirements

Cisco.com

	Bandwidth	Latency And Jitter	Connect Time	Tariff	Initial Cost	Reliability	Availability
Remote User (On-Demand Connection)	L	M	M	H	L	M	H
Central Site (On-Demand Connection)	L/M	M	M	H	L	M	H
Remote Branch Office (Always-On Connection)	L	M	L	L	L	M	H
Central Site (Always-on Connection)	L/M	M	L	L	L	M	H

L – low, M – medium, H – high

You will begin by evaluating the networking parameters to help select a suitable technology for a given network. You can consider always-on and on-demand connections for both central and remote sites, as the expected performance is slightly different.

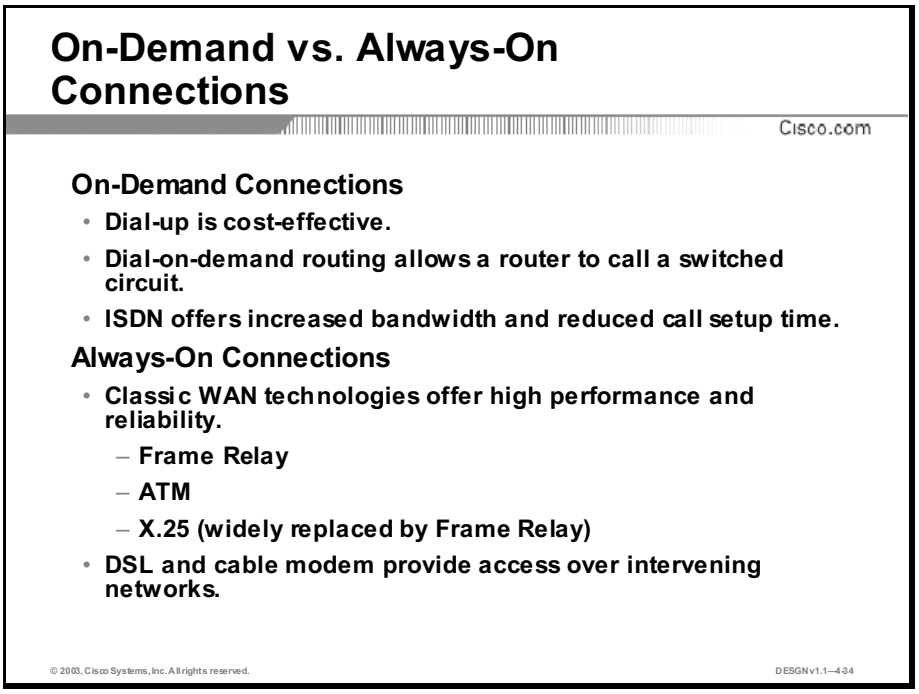

Analog Modem

Dial-up service offers a cost-effective solution for WAN connectivity. This is especially true when you combine dial-up service with applications such as dial-on-demand routing (DDR). DDR is a technique where a router initiates a call on a switched circuit when it needs to send data.

ISDN

Before ISDN was available, plain old telephone service (POTS) provided data connectivity over the PSTN using analog modems. Connectivity over ISDN offers increased bandwidth, reduced call setup time, reduced latency, and lower signal to noise ratios. Therefore, ISDN presents an effective solution for many remote user applications.

X.25, Frame Relay

Frame Relay is often described as a streamlined version of X.25, which offers fewer robust capabilities such as windowing and retransmission of lost data. Frame Relay operates over WAN facilities that offer more reliable connection services and a higher degree of reliability than the media for which the X.25 was designed. Frame Relay is a data link layer protocol suite, whereas X.25 provides services at the data link and network layers. Therefore, Frame Relay offers higher throughput and greater transmission efficiency than X.25 and is suitable for current WAN applications such as LAN interconnection.

Note: X.25 is a legacy technology that is still in use in various environments. It is being replaced by faster and more bandwidth efficient technologies such as Frame Relay.

TDM

TDM reserves point-to-point connection bandwidth indefinitely for transmissions, rather than only using bandwidth as required. The carrier establishes the connection by dedicating a channel with the use of TDM. By contrast, packet-switched networks traditionally offer the SP more flexibility and use network bandwidth more efficiently than TDM networks because the network resources are shared dynamically. Subscribers are charged on the basis of their guaranteed use of the network.

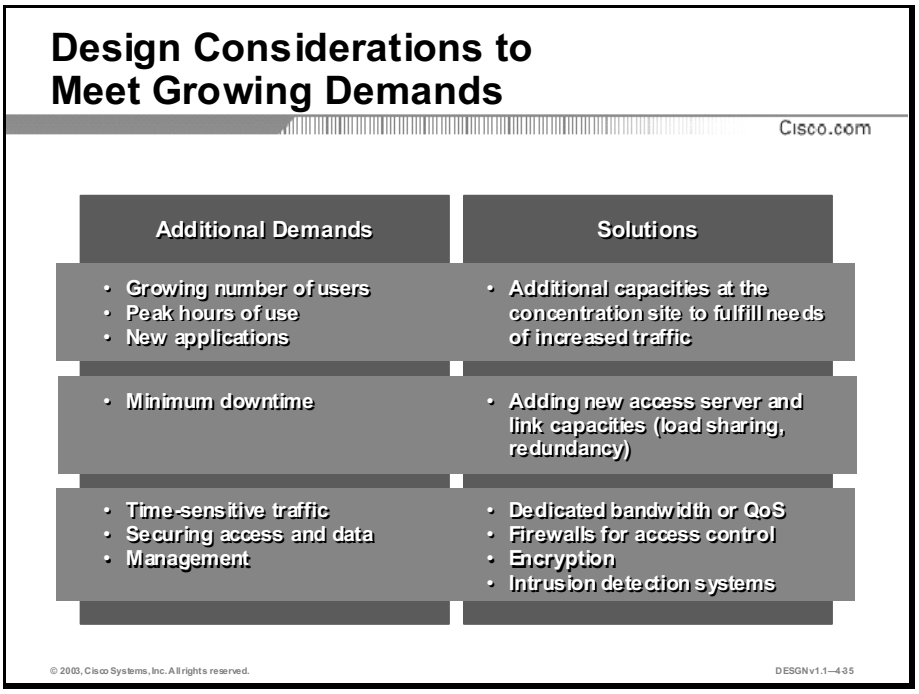

The remote-access design is an integral part of the total network solution and must scale to meet growing demand. It is up to the designer to analyze and estimate further requests for improved performance and for additional services. These demands (left column in the figure) may have several solutions (right column in the figure). The final goal is an infrastructure with efficient, reliable, and secure connections.

Remote-Access Technology Comparison

Cisco.com

	Bandwidth	Latency And Jitter	Connect Time	Tariff	Initial Cost	Reliability	Availability
Analog Modem	L	H	H	H	L	L	H
ISDN	L	M/H	M	H	L	M	H
P2P Protocols over Sync or Async Serial	L	M	L	M	M	M	H
X.25, Frame Relay	L	L	L	M	M	M	H
TDM	M	L	L	M	M	M	H
Ethernet over Fiber	M/H	L	L	M	M	M	M
Site-to-Site Leased Lines (BBM, DF)	M/H	L	L	M/H	H	H	L
POS and ATM over SDH/SONET	H	L	L	M	M	H	M
ADSL	L/M	M/H	L	L	M	M	M
Cable Modem	L/M	M/H	L	L	M	L	M
Wireless Cellular	L/M	M/H	L	L	M	L	L

L – low, M – medium, H – high

The figure compares the characteristics of candidate technologies for remote access. Based on the requirements, PPP over either analog modem or Integrated Services Digital Network (ISDN) is the best choice for on-demand connections, with X.25 or Frame Relay over TDM or synchronous serial leased line the best choices for permanent connections.

Note: X.25 or Frame Relay over TDM or synchronous serial leased line and PPP over ISDN provide the ability to connect multiple remote sites over a single physical connection at the central site. This reduces the number of point-to-point physical connections required to link sites together.

Note: When faced with multiple choices, you must perform a broad analysis, comparing the benefits and drawbacks of each technology. In the figure, a comparison is made between analog modem and ISDN for on-demand connections and between X.25 or Frame Relay for permanent connection.

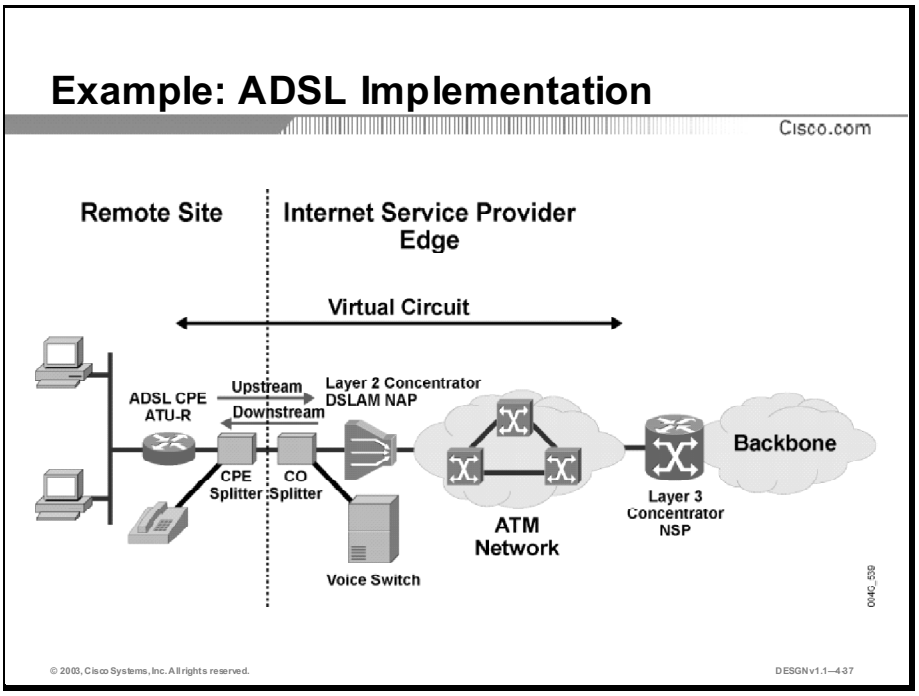

The figure illustrates a typical asymmetric DSL (ADSL) service architecture. The network consists of customer premises equipment (CPE), the Network Access Provider (NAP), and the network service provider (NSP):

- Customer premises equipment (CPE) refers to an end-user workstation, such as a PC, together with an ADSL modem or an ADSL Transmission Unit-Remote (ATU-R)
- The NAP provides ADSL line termination by using DSL access multiplexers (DSLAMs)
- The DSLAM forwards traffic to the local access concentrator, the NSP, which is used for L3 termination

An ADSL circuit connects an ADSL modem on each end of a twisted-pair telephone line. This creates three information channels:

- Medium-speed downstream channel
- Low-speed upstream channel
- Basic telephone service channel

Filters (splitters) split off the basic telephone service channel from the digital modem. This guarantees uninterrupted basic telephone service, even if ADSL fails.

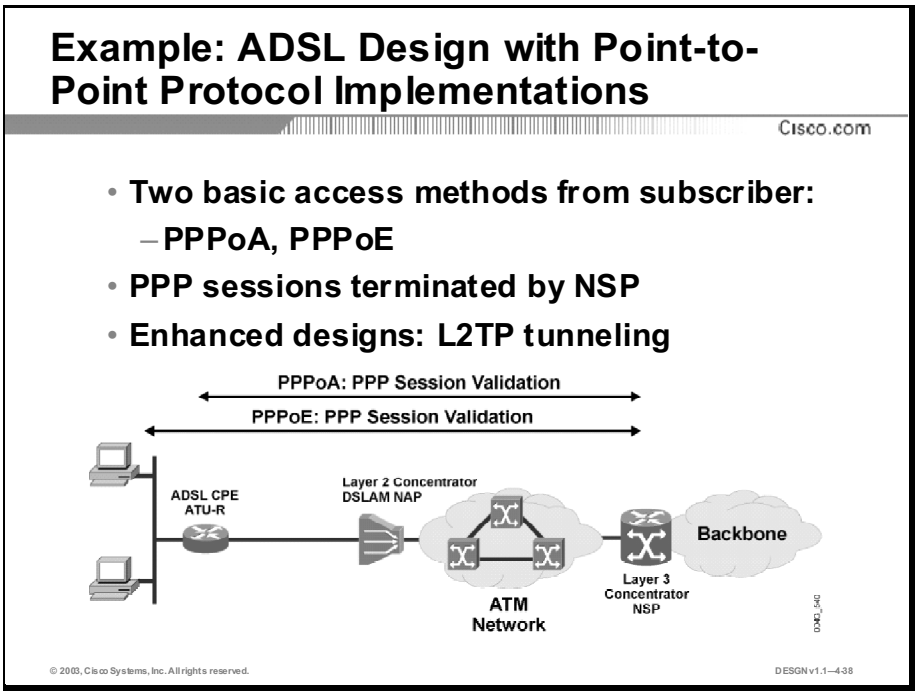

The figure illustrates a typical DSL network. It includes enterprise workstations and PCs on a LAN, CPE (DSL routers), a DSL access concentrator on an ATM transport network, an NSP concentrator, and both packet and ATM core networks.

Two very popular PPP implementations exist in ADSL designs, PPP over ATM (PPPoA), and PPP over Ethernet (PPPoE).

PPPoA Implementation

In the PPPoA architecture, the CPE acts as an Ethernet-to-WAN router and the PPP session is established between the CPE and the network layer access concentrator (the NSP). A PPPoA implementation involves configuring the CPE with PPP authentication information (login and password). The main advantage of PPPoA over pure bridging implementations is that it provides per-session authentication, authorization, and accounting (AAA).

The advantages of PPPoA include:

- The NAP has an existing provisioning system, which connects PVCs from one end of the ATM network to the other.
- ATM end-to-end is the basis of the infrastructure. Therefore, it may be easier to put a subscriber into a specific traffic class.
- Any protocol can ride transparently on top of the ATM virtual circuit. Therefore, addressing of these protocols is fully transparent to the network.
- The CPE requires little, if any, configuration.

PPPoE Implementation

In the PPPoE architecture, the CPE acts as an Ethernet-to-WAN bridge and the PPP session is established between the end user's PC or PPPoE router and the network layer access concentrator (the NSP).

The client initiates a PPP session by encapsulating PPP frames into a MAC frame and then bridging the frame (over ATM/DSL) to the gateway router (NSP). From this point, the PPP sessions can be established, authenticated, and addressed. The client receives its IP address using PPP negotiation from the termination point (NSP).

Long Range Ethernet

LRE is an innovative technology that uses Ethernet over existing, unconditioned, telephone-grade wire (copper twisted-pair). The technology allows Ethernet LAN transmissions to coexist with POTS, ISDN, or advanced PBX signaling services over the same pair of ordinary copper wires. LRE technology uses the newest coding and digital modulation techniques developed for DSL as well as Ethernet, the most popular LAN protocol.

Note: LRE is used in a LAN implementation and reuses existing wiring. It is included here because it is derived from DSL technologies.

An LRE system provides a point-to-point transmission that can deliver a symmetrical, full duplex, raw data rate of 11.25 Mbps over distances of up to 1 mile (1.6 km). Products utilizing LRE technology are simple to install and interface easily with any existing Ethernet solution.

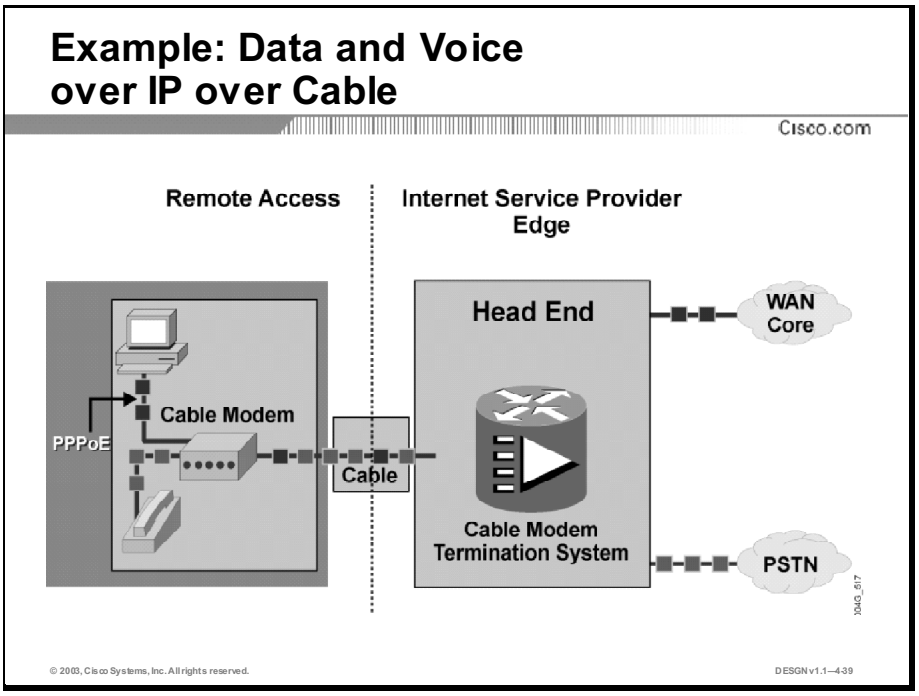

The Universal Broadband Router (uBR), also referred to as the cable modem termination system (CMTS), provides high-speed data connectivity and is deployed at the cable company's headend. The uBR forwards data upstream to connect with either the PSTN or the Internet. The cable modem (CM), also referred to as the cable access router, at the remote location supports voice, modem, and fax calls over the TCP/IP cable network.

In general, cable operators install CMs at the customer premises to support small businesses, branch offices, and corporate telecommuters.

The uBR is designed to be installed at a cable operator's headend facility or distribution hub and to function as the CMTS for subscriber-end devices.

The Data-over-Cable Service Interface Specifications (DOCSIS) protocol describes data-over-cable procedures that the equipment must support.

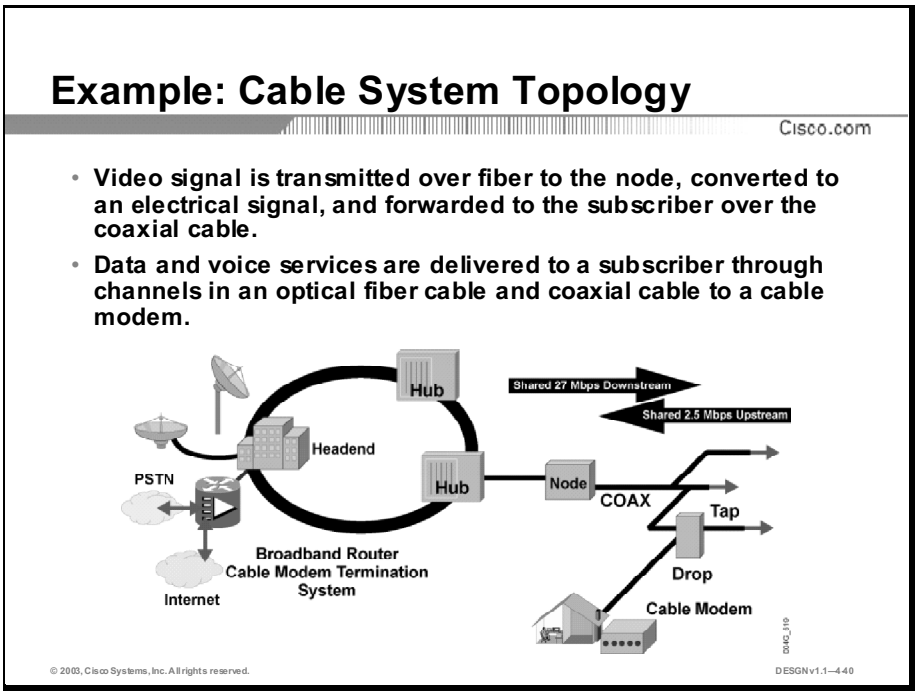

The figure illustrates a hybrid fiber-coaxial (HFC) topology, which uses a high-speed fiber backbone and coaxial cables to connect end users. The service provider backbone supports analog and digital video directly. IP packets running over the infrastructure are implemented into the data and packet telephony portion of the network.

Upstream and Downstream Data Flow

A data service is delivered to a subscriber through a coaxial cable or optical fiber cable to a cable modem installed externally or internally to a subscriber's computer or television set. One channel is used for upstream signals from the cable modem to the CMTS, and another channel is used for downstream signals from the CMTS to the cable modem.

When a CMTS receives signals from a cable modem, it converts these signals into IP packets, which are then sent to an IP router for transmission across the Internet. When a CMTS sends signals to a cable modem, it modulates the downstream signals for transmission across the cable, or across the optical fiber and cable, to the cable modem. All cable modems can receive from and send signals to the CMTS but not to other cable modems on the line.

The actual bandwidth for Internet service over a cable TV line is shared 27 Mbps on the download path to the subscriber with about 2.5 Mbps of shared bandwidth for interactive responses in the other direction. Under DOCSIS 1.1, upstream transmissions employ time division multiple access (TDMA) sharing using either the Quadrature Phase-Shift Keying (QPSK) or the quadruple amplitude modulation (QAM), while downstream transmissions employ TDM.

Cable Television Transmission

Before conversion to their respective channel assignments in the downstream frequency domain, signals from broadcasters and satellite services are descrambled.

The hub distributes the video signals to the optical node. The optical node converts the optical signal to an electrical signal and amplifies and forwards it downstream over coaxial cable for distribution to the cable operator's customers. Now the deployment of fiber technology into the network results in what is currently referred to as an HFC network.

Benefits of Hybrid Fiber-Coaxial Networks

The use of fiber instead of coaxial permits the elimination of numerous amplifiers in the cascade to support the customer serving area. Thus, the operating company is able to improve signal quality and eliminate numerous components susceptible to failure, while significantly reducing operating and maintenance costs.

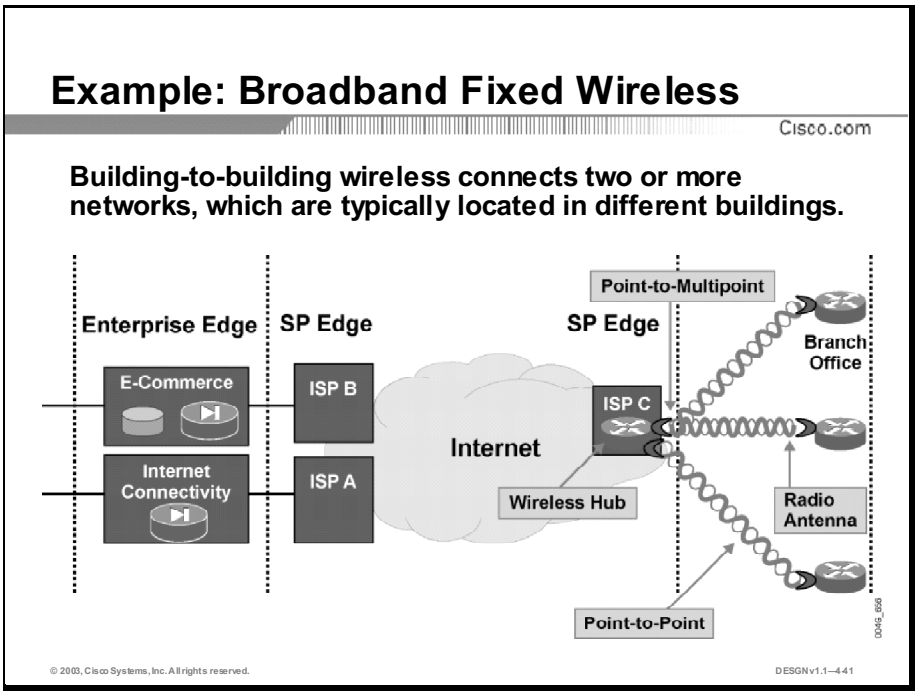

Building-to-building wireless connects two or more networks that are typically located in different buildings. A series of wireless bridges or routers can connect discrete distant sites into a single LAN, interconnecting hard-to-wire sites, noncontiguous floors, satellite offices, school or corporate campus settings, temporary networks, and warehouses.

The figure illustrates a situation where an ISP offers connectivity through access points using wireless technology. The routers are configured for point-to-point or point-to-multipoint applications and allow multiple line-of-sight sites to share a single, high-speed connection.

The system consists of a hub (or headend or base station), communicating with one or many customers through the use of a radio transmission system. The headend is an outdoor unit, or transformer, connected to a wireless modem card inside a Cisco uBR. The other transformer at the remote premises connects to a wireless network module in a router.

Some organizations require a data rate that is higher than the service provider can supply within the traffic capacity of the multipoint system. However, the service provider can satisfy the requirement by installing point-to-point links from the same hub as the point-to-multipoint system. Thus, the hub can support point-to-multipoint and point-to-point systems. In both cases, integrating the wireless card directly into the router integrates the IOS features and network management.

The headends or hubs utilize a Cisco uBR with a multipoint line card that delivers up to 22 Mbps downstream and 18 Mbps upstream bandwidth per 6-MHz channel pair. For deployment on the hub end, you need an outdoor, point-to-multipoint transceiver that converts intermediate frequency (IF) to radio frequency (RF) for transmitting or receiving voice or data with less chance of signal loss. At the user's site, an antenna sends and receives communications from other external locations. A cable connection to the wiring-closet links a Cisco 2600 or 3600 series router equipped with a multipoint subscriber network module, that also delivers up to 22 Mbps downstream and 18 Mbps upstream bandwidth per 6-MHz channel pair.

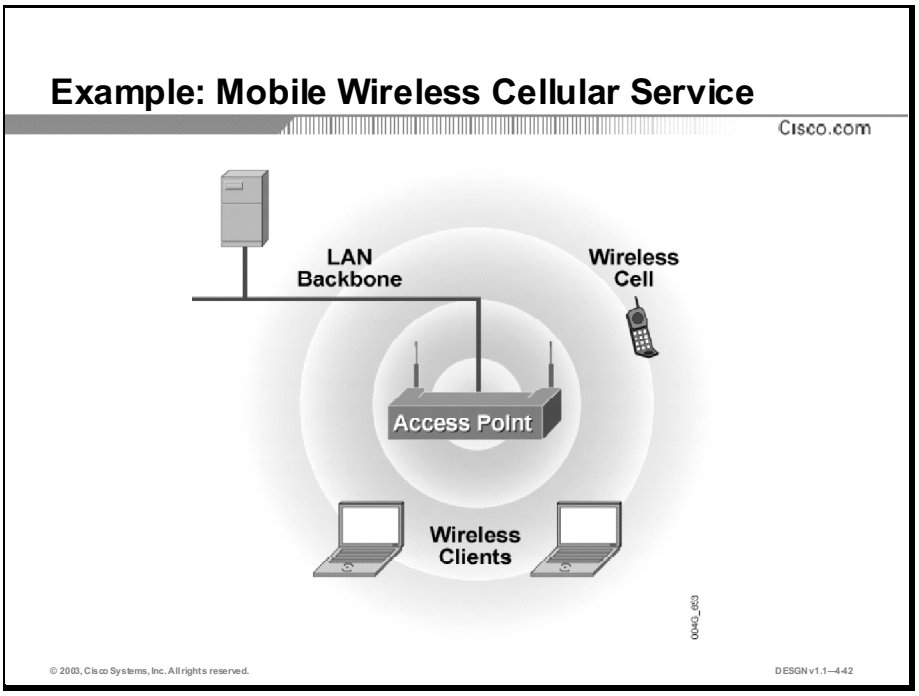

Real usage of wireless technologies coincided with the introduction of digital services on wireless. Second and third generation mobile phones offer connectivity and higher speeds:

- **Global system for mobile communication (GSM):** GSM is a digital mobile radio standard that uses TDMA technology in three different bands: 900, 1800, and 1900 MHz. The transfer data rate is 9.6 kbps. A unique benefit of GSM is its international roaming ability.
- **General packet radio service (GPRS):** GPRS extends the capability of GSM speed and supports intermittent and bursty data transfer. Speeds offered to the client are in the range of ISDN speeds (64 kbps to 128 kbps).
- **Universal Mobile Telecommunications Service (UMTS):** Also called "third-generation (3G)" broadband, UMTS provides packet-based transmission of text, digitized voice, video, and multimedia at data rates up to 2 Mbps. UMTS offers a consistent set of services to mobile computer and phone users, no matter where they are located in the world.

Benefits of Using Wireless Cellular Service

This list summarizes the main benefits of using wireless cellular service:

- **Completes the access technology portfolio:** Wireless enables a fully comprehensive access technology portfolio to work with existing dial, cable, DSL, and LRE technologies.
- **Goes where cable and fiber cannot:** Wireless will carry information across geographical areas that are prohibitive in terms of distance, cost, access, or time. It also sidesteps the issues of incumbent local exchange carrier (ILEC) colocation.

Using a Service Provider Network to Connect Dispersed Enterprise Sites

One option for connecting an enterprise's geographically dispersed sites is to establish WAN communications over a service provider network. This topic describes how to connect dispersed sites over a service provider network.

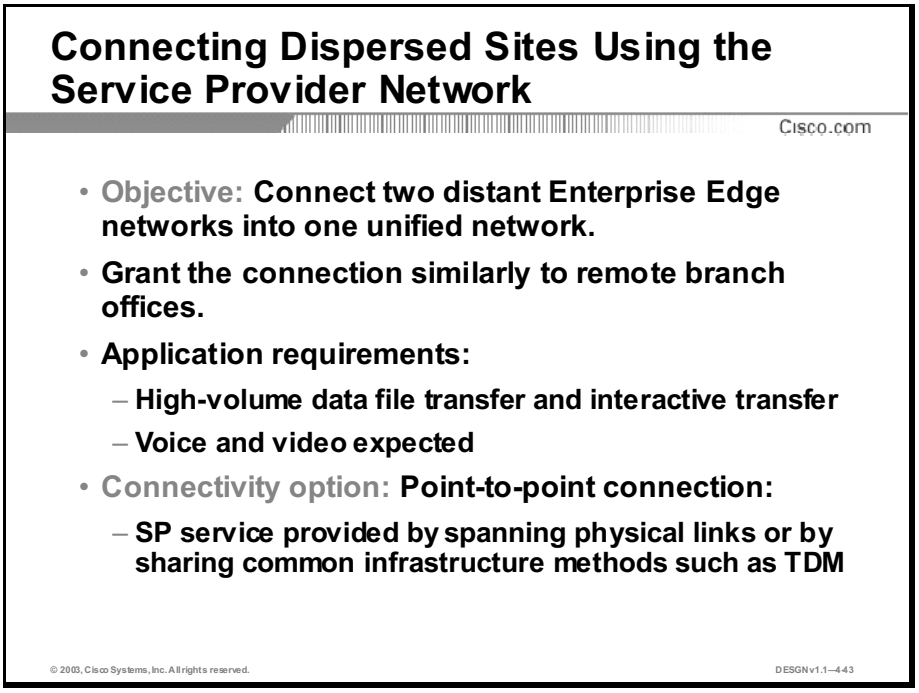

An enterprise needs to connect highly demanding distant sites through the use of L2 technology. The connection requirements are summarized in these two points:

- L2 WAN technology between distant Enterprise Edge networks
- High-volume data file transfer and interactive traffic, combined with time-sensitive voice and video applications with a good quality of transmission

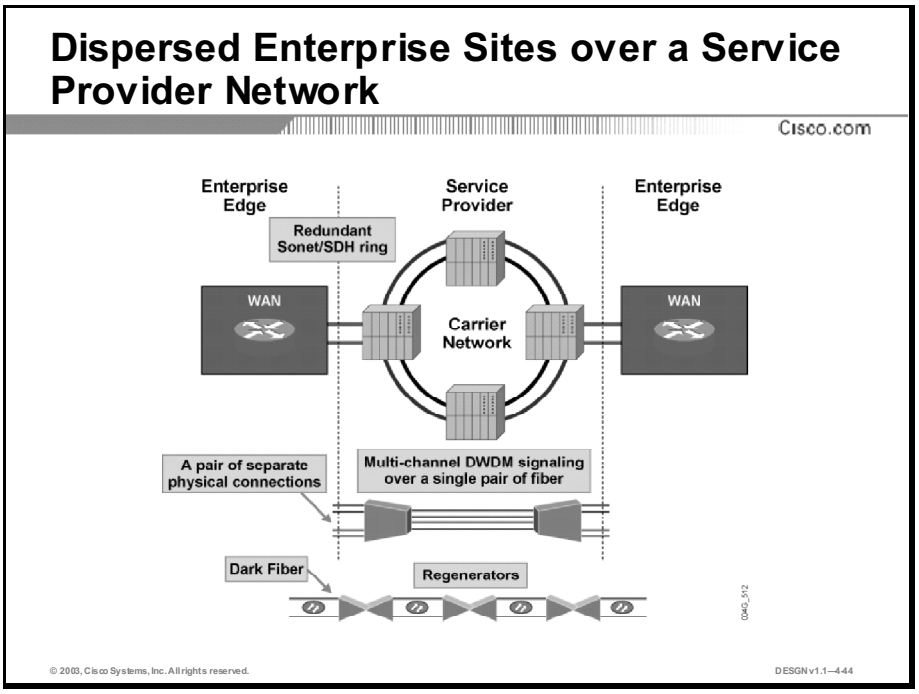

A point-to-point link provides a preestablished WAN communication path from the enterprise network, through a carrier network (such as a telephone network), to the enterprise's remote network. For a point-to-point line, the carrier allocates a physical medium to the subscriber and provides end-to-end connectivity. This is done by either spanning physical links or by sharing common infrastructure using, for example, frequency-division multiplexing (FDM) or TDM. The bandwidth requirements and the distance between the two connected points are usually the basis for the circuits' costs.

The figure illustrates a SONET/SDH ring, dense wavelength division multiplexing (DWDM), and dark fiber as widely deployed architectures in an SP environment. In terms of reliability, SONET/SDH networks offer advanced features, such as automatic backup and repair mechanisms to cope with system faults. Failure of a link or a network element does not lead to failure of the entire network.

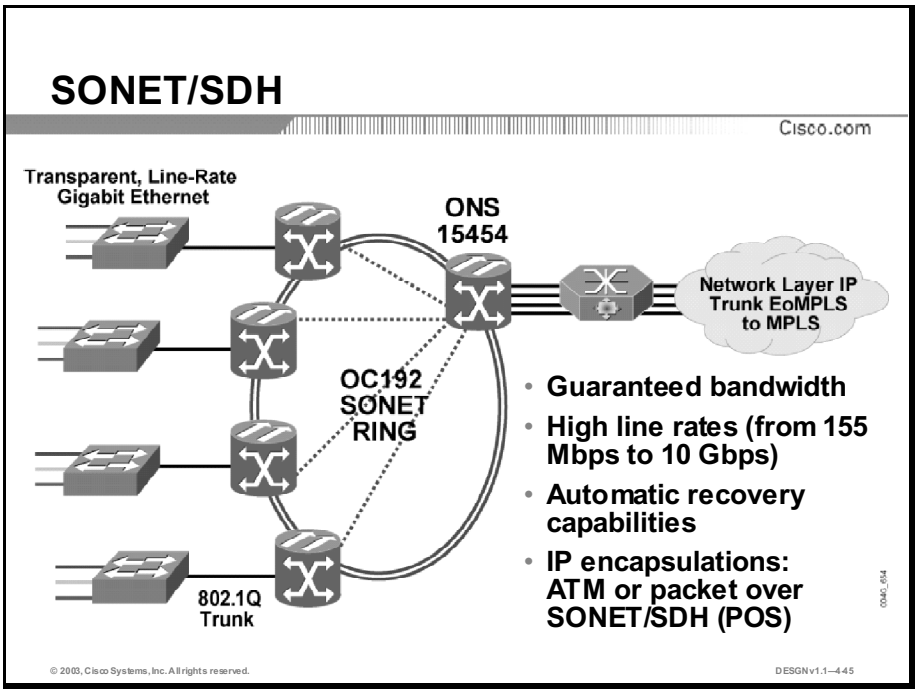

Circuit-based services architecture is the basis for SONET and Synchronous Digital Hierarchy (SDH). This technology uses TDM and delivers high-value services over an optical infrastructure. SONET/SDH provisions high-speed point-to-point connections that guarantee bandwidth, regardless of actual usage (for example, common bit rates are 155 Mbps and 622 Mbps, with a maximum of 10 Gbps).

SONET/SDH rings offer proactive performance monitoring and automatic recovery ("self-healing") via an automatic protection switching (APS) mechanism.

SONET/SDH rings support two IP encapsulations for user interfaces: ATM or packet over SONET/SDH (POS), which sends native IP packets directly over SONET/SDH frames.

Note: SONET and SDH have important differences in terminology. SONET is an American National Standards Institute (ANSI) specification. SDH is the SONET-equivalent specification proposed by the International Telecommunications Union (ITU). European carriers use SDH widely; Asian and Pacific Rim carriers use SONET more frequently.

DWDM and Dark Fiber

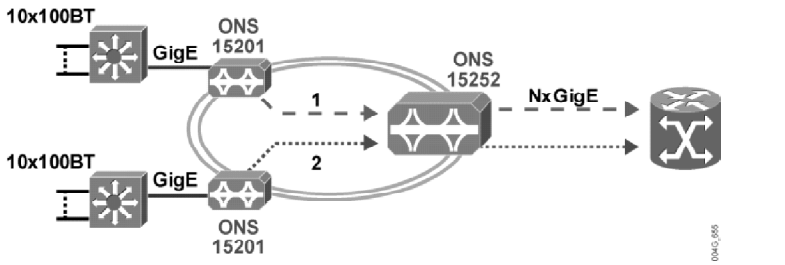

DWDM:

- **Improved signaling mechanisms to optimize bandwidth usage**
- **Used inside the SONET/SDH ring**

Dark Fiber:

- **Edge devices directly connected to regenerators or DWDM concentrators**
- **Edge devices can use any Layer 2 encapsulation**

Dense wavelength division multiplexing (DWDM) improves the utilization for optical fiber. Multichannel signaling on a single strand of fiber increases its available bandwidth. DWDM is a crucial component of optical networks. It maximizes the use of installed fiber cable and allows service providers to efficiently provision new services over the existing infrastructure. Flexible add-and-drop modules allow service providers to drop and insert individual channels along a route. An open architecture system allows a variety of devices, including SONET terminals, ATM switches, and IP routers, to be connected.

A dark fiber connection allows framing options other than SONET/SDH. The edge devices connect directly over the site-to-site dark fiber using other encapsulations, such as Gigabit Ethernet. To transmit data over significantly long distances, regenerators are inserted into the link to maintain signal integrity and provide appropriate jitter control.

Dispersed Site Technology Comparison

Cisco.com

	Bandwidth	Latency And Jitter	Connect Time	Tariff	Initial Cost	Reliability	Availability
Analog Modem	L	H	H	H	L	L	H
ISDN	L	M/H	M	H	L	M	H
Point-to-Point Protocols over Sync or Async Serial	L	M	L	M	M	M	H
X.25, Frame Relay	L	L	L	M	M	M	H
TDM	M	L	L	M	M	M	H
Ethernet over Fiber	M/H	L	L	M	M	M	M
Site-to-Site Leased Lines (BBM, DF)	M/H	L	L	M/H	H	H	L
POS and ATM over SDH/SONET	H	L	L	M	M	H	M
ADSL	L/M	M/H	L	L	M	M	M
Cable Modem	L/M	M/H	L	L	M	L	M
Wireless Cellular	L/M	M/H	L	L	M	L	L

L – low, M – medium, H – high

DESGN v1.1—4-47

The figure compares the characteristics of SONET and DWDM as candidate technologies for dispersed sites over a service provider network.

Designing Virtual Private Networks

A VPN is defined as connectivity deployed on a shared infrastructure with the same policies and performance as a private network. The infrastructure used can be the Internet, an IP infrastructure, or any WAN infrastructure such as a Frame Relay network or an ATM WAN. This topic describes how to design a VPN.

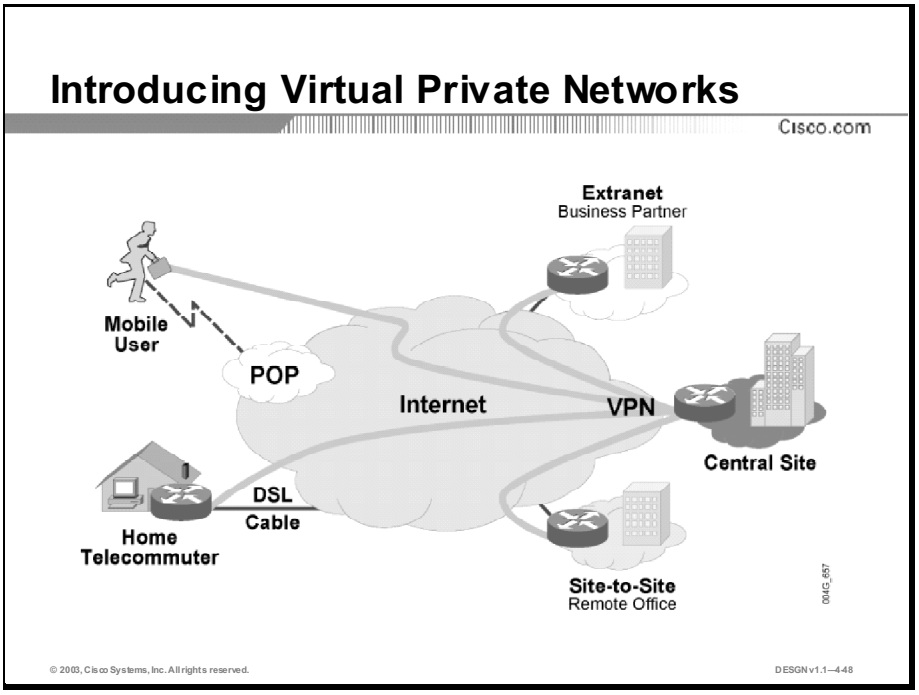

The three types of VPNs are grouped according to their applications:

- **Access VPN:** Access VPNs provide access to a corporate intranet over a shared infrastructure with the same policies as a private network. Remote-access connectivity is through dial-up, ISDN, DSL, wireless, and cable technologies. Access VPNs enable businesses to outsource their dial or other broadband remote-access connections without compromising their security policy. They include two architectural options: client-initiated connections or connections initiated by network access server (NAS). With client-initiated access VPNs, users establish an encrypted IP tunnel from their PCs across an SP's shared network to their corporate network. Another architecture for access VPNs defines the tunnels initiated from the NAS, where remote users dial into the local SP points of presence (POP), and the SP initiates a secure, encrypted tunnel to the corporate network.
- **Intranet VPN:** Intranet VPNs link remote offices. The intranet VPN services are typically based on dedicated access that extends the basic remote-access VPN to other corporate offices across the Internet or across the SP's IP backbone. These are the main benefits of intranet VPNs:
  - Reduced WAN infrastructure needs
  - Lower ongoing leased-line or Frame Relay charges
  - Operational savings

Note: With VPNs across the Internet, there are no performance guarantees.

- **Extranet VPN:** An organization uses either the Internet or an SP network to connect to its business partners. The security policy becomes very important at this point, as the organization does not want a hacker to spoof any orders from a business partner.

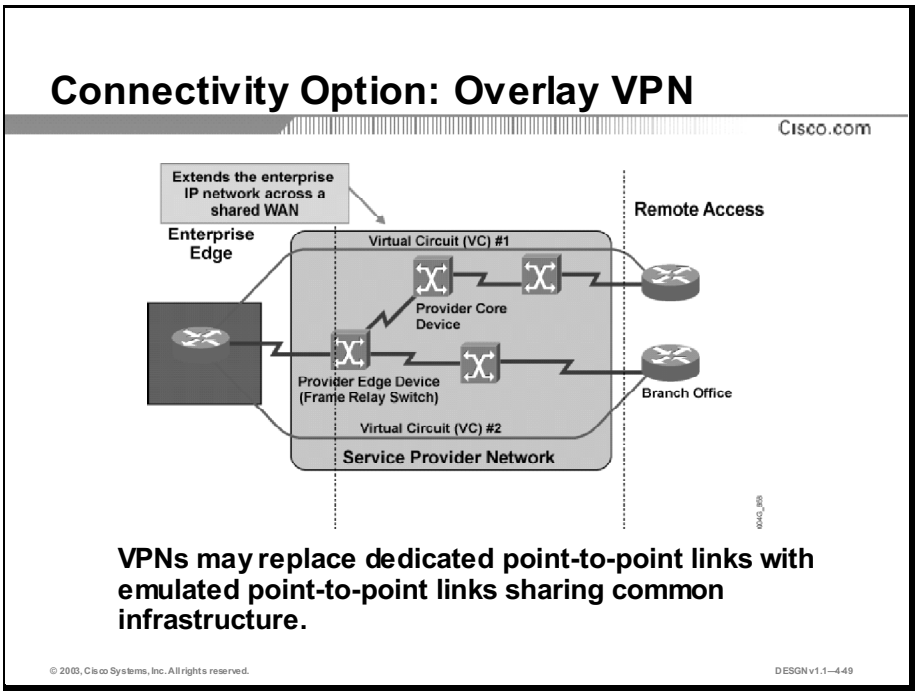

Overlay VPNs are implemented with traditional L1 and L2 technologies (ISDN, SONET/SDH, Frame Relay, ATM) and IP-based L3 solutions (generic routing encapsulation [GRE] and IPSec).

The overlay VPNs are more difficult to operate because of higher maintenance costs:

- Every individual virtual circuit needs to be provisioned.
- From the L3 perspective, the provider network is invisible. The customer routers are linked with emulated point-to-point links. The routing protocol runs directly between routers that establish routing adjacencies and exchange routing information.
- The provider is not aware of customer routing and has no information about customer routes. The only responsibility of the provider is the point-to-point data transport between customer sites.
- Optimum routing between customer sites requires a full mesh of virtual circuits between sites.
- You must provision bandwidth on a site-to-site basis.

An overlay VPN uses virtual circuits to establish the end-to-end connection over a shared SP infrastructure. A common infrastructure emulates point-to-point links and results in statistical sharing of provider infrastructure, and replaces the dedicated links. Statistical sharing of the infrastructure enables the SP to offer the connectivity for a lower price. This results in lower operational costs.

Example: Overlay VPN Statistical Sharing

The figure illustrates statistical sharing. The router in the Enterprise Edge module has one physical connection to the SP with two virtual circuits provisioned. Virtual circuit 1 (VC #1) provides connectivity to the routers on the right (Enterprise Edge in the branch office). Virtual circuit 2 (VC #2) provides the connectivity to the other branch office router.

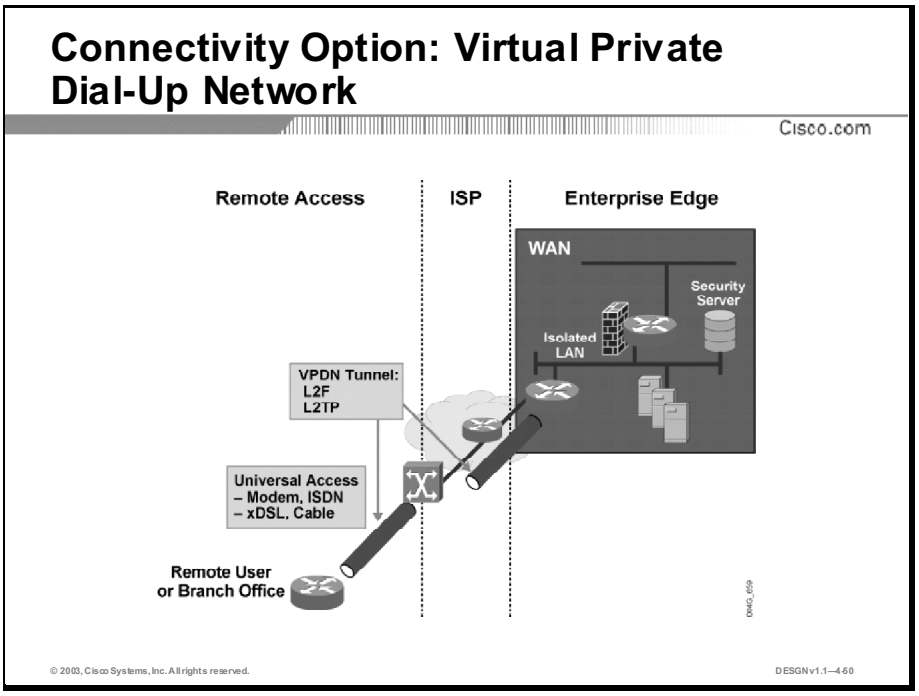

A Virtual Private Dial-Up Network (VPDN) enables an enterprise to configure secure networks that rely upon an ISP. The ISP terminates dial-up connections and forwards traffic through dynamically established tunnels. The dial-up SP agrees to forward the company's traffic from the ISP's POP to a company-run home gateway. Network configuration and security remains within the client's control. The dial-up SP supplies a virtual tunnel between the company's sites using Layer 2 Forwarding (L2F) Protocol or Layer 2 Tunneling Protocol (L2TP) tunnels.

The figure illustrates some of the features of remote-access VPNs based on VPDN. The features run over any access that is available, and ubiquity is important. This means they should work with a modem, ISDN, xDSL, or cable. The features provide potential operations and infrastructure cost savings because the company can outsource its dial-in remote access, thereby avoiding the remote-access server business.

Access VPN connectivity involves the configuration of VPDN tunnels. Voluntary tunnels are those initiated by the client PC. Compulsory tunnels require SP participation and awareness, leaving the client with no choice in influencing tunnel selection.

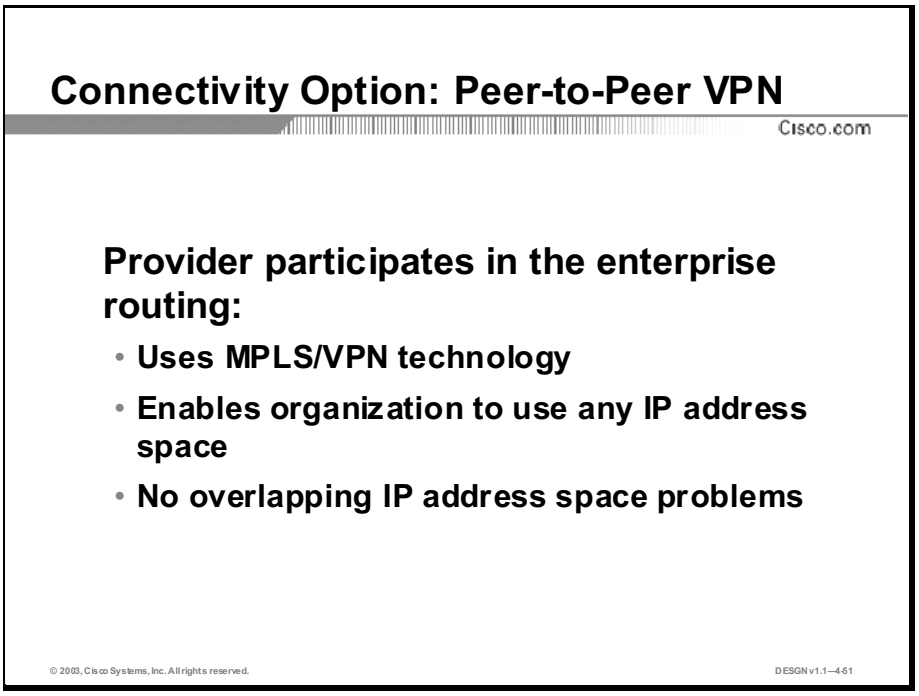

In peer-to-peer VPN, the provider participates actively in enterprise routing.

Traditional peer-to-peer VPNs are implemented with packet filters on shared provider edge routers or with dedicated per-customer provider edge routers. Along with high maintenance costs for the packet filter approach or equipment costs for the dedicated per-customer provider edge-router approach, both methods require the enterprise to accept the provider-assigned address space or to use public IP addresses in the private enterprise network. Modern Multiprotocol Label Switching (MPLS) VPNs provide all the benefits of peer-to-peer VPNs and alleviate most of the peer-to-peer VPN drawbacks, such as the need for common customer address.

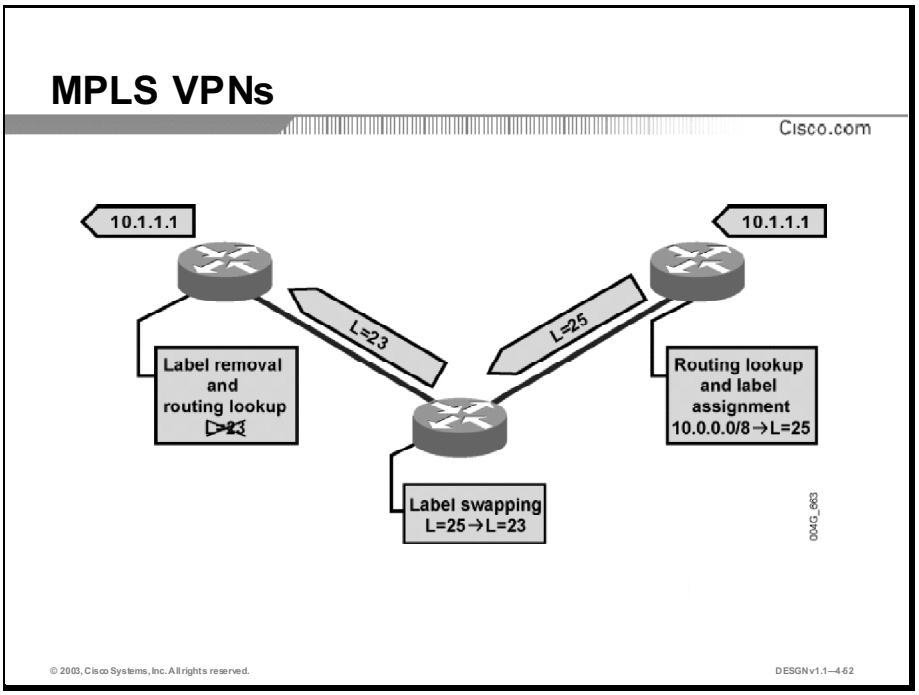

MPLS is a switching mechanism that uses labels (numbers) to forward packets. Labels usually correspond to Layer 3 destination addresses, making MPLS equal to destination-based routing. Labels can correspond to other parameters, such as a quality of service (QoS) value, source address, or a Layer-2 circuit identifier. Label switching occurs regardless of the Layer 3 protocol.

With MPLS VPNs, networks are learned with an Interior Gateway Protocol (IGP) routing protocol such as OSPF, External Border Gateway Protocol (EBGP), RIP version 2 [RIPv2] or with static addresses configured by an administrator, or with BGP from other internal routers.MPLS VPNs use an additional label to specify the VPN and the corresponding VPN destination network. This allows for overlapping addresses between VPNs.

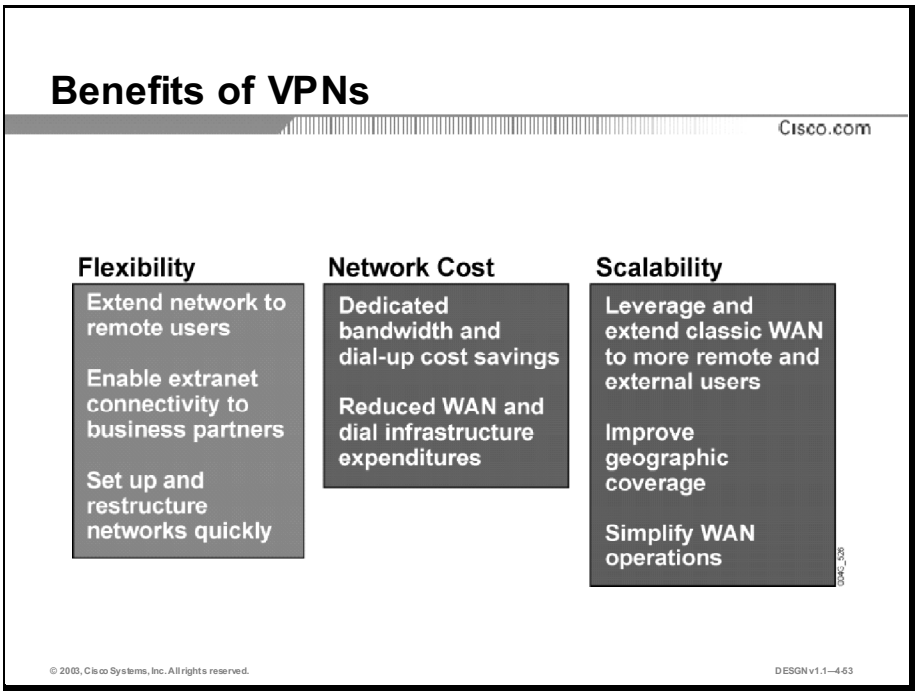

The benefits of using VPNs include flexibility, scalability, and lowered cost of communication. VPNs offer flexibility because you can quickly set up site-to-site and remote-access connections over an existing infrastructure. You can provision a variety of security policies in a VPN, enabling flexible interconnection of different security domains.

VPNs offer scalability over large areas because Internet transport is nearly universally available. This reduces the number of physical connections and simplifies the underlying structure of a customer's WAN.

Lower cost is a main reason for migrating from traditional connectivity options to a VPN connection. Customers may reuse existing links and take advantage of the statistical packet multiplexing features.

Designing a WAN Backup Strategy

Each Enterprise Edge solution requires a WAN backup to provide high availability between sites. Branch offices should experience minimum downtime, in case of primary link failure. You can establish backup connections using either dial-up or permanent connections. This topic describes designing a WAN backup strategy.

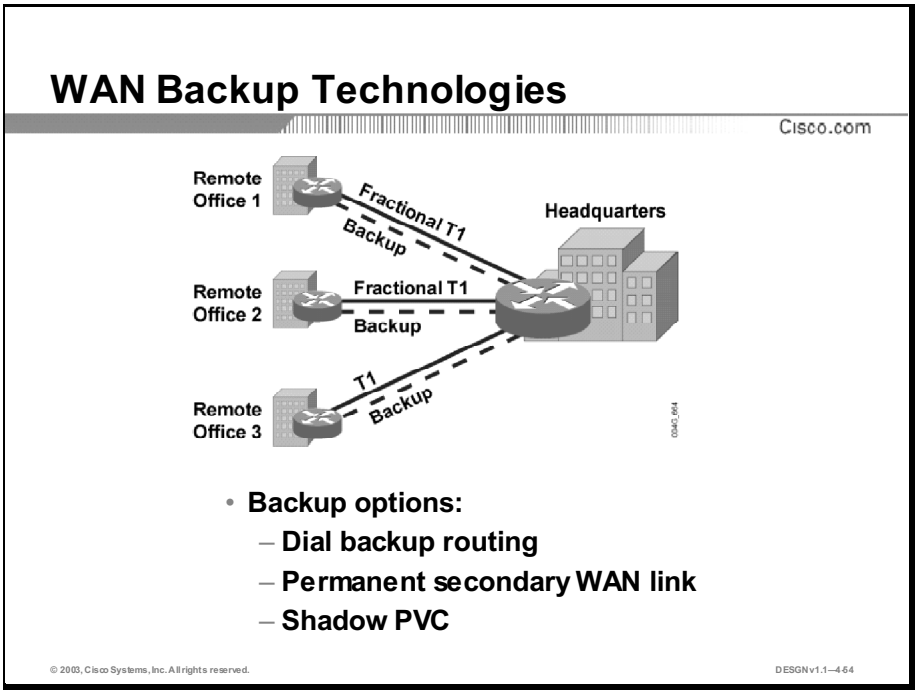

The primary WAN backup options are:

- **Dial backup routing:** Dial backup routing uses dial-up services for backup. The switched circuit provides the backup service for another type of circuit, such as point-to-point or Frame Relay. The router initiates the dial backup line when a failure is detected on the primary circuit. The dial backup line provides WAN connectivity until the primary circuit is restored, and then terminates.
- **Permanent secondary WAN link:** The deployment of an additional permanent WAN link between each remote office and the central office makes the network more fault tolerant. This offers two advantages:
  - **Backup link:** If a primary link connecting any remote office and the central office fails, the backup link is used. Routers can automatically compensate failed WAN links through floating static routing and routing protocols, such as the Internet Gateway Routing Protocol (IGRP), Open Shortest Path First (OSPF), and Intermediate System-to-Intermediate System (IS-IS). If one link fails, the routing software recalculates the routing algorithm and sends all traffic through another link. This allows applications to proceed in the event of a WAN link failure and thus improves application availability.
  - **Increased bandwidth:** Decreases response times; occurs if the routers support load balancing between two parallel links of equal cost. In this case, load balancing is performed automatically via routing protocol.

- **Shadow PVC:** The SP provides a user with a secondary PVC without any additional charge as long as the load does not exceed a given speed while the primary PVC is available,. The SP charges the client for two PVCs if the traffic limit on the shadow PVC is exceeded.

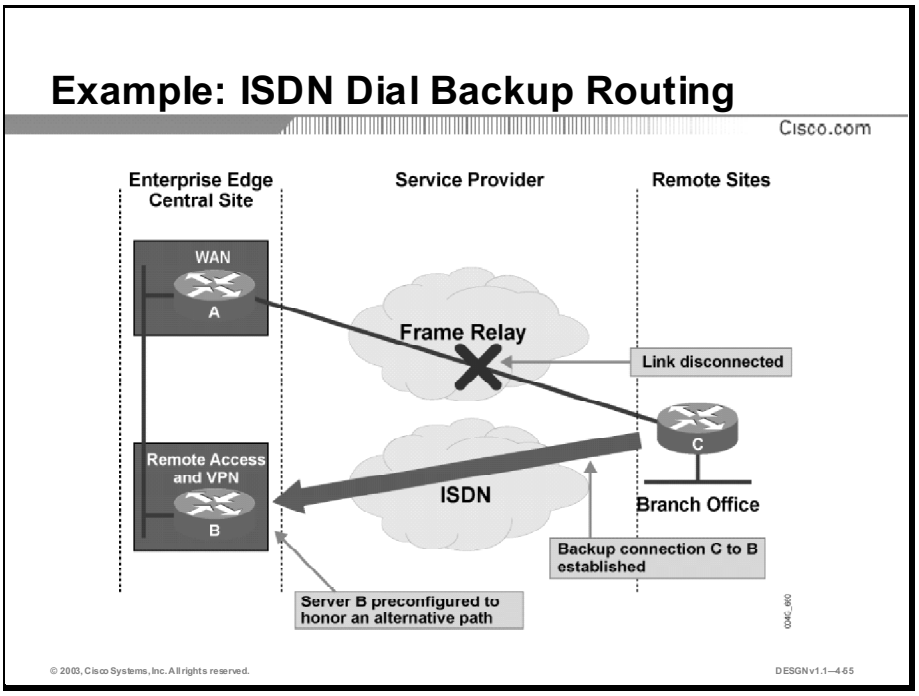

Using point-to-point subinterfaces on server A, with ISDN as a backup, is the preferred solution when providing dial backup in the Frame Relay environment. Because of the number of remote offices (a maximum of 20), a single ISDN PRI is allocated on server B to support the dial backup application.

A typical scenario for providing backup for remote locations because of a link failure includes these steps:

Step 1 The link between routers D and A fails. Sometimes interfaces remain in the "up" state even if the link fails, and the only way to detect that something went wrong is by neighbor loss detection.

Step 2 The backup interface feature on router D detects a data-link connection identifier (DLCI) loss.

Step 3 The backup interface selects an alternate ISDN dial backup interface to establish a connection to server B.

Step 4 A routing protocol recalculates the paths toward the remote sites. Upon completion, the network is converged.

Step 5 When reestablishing the primary connection, the ISDN connection becomes obsolete and is torn down.

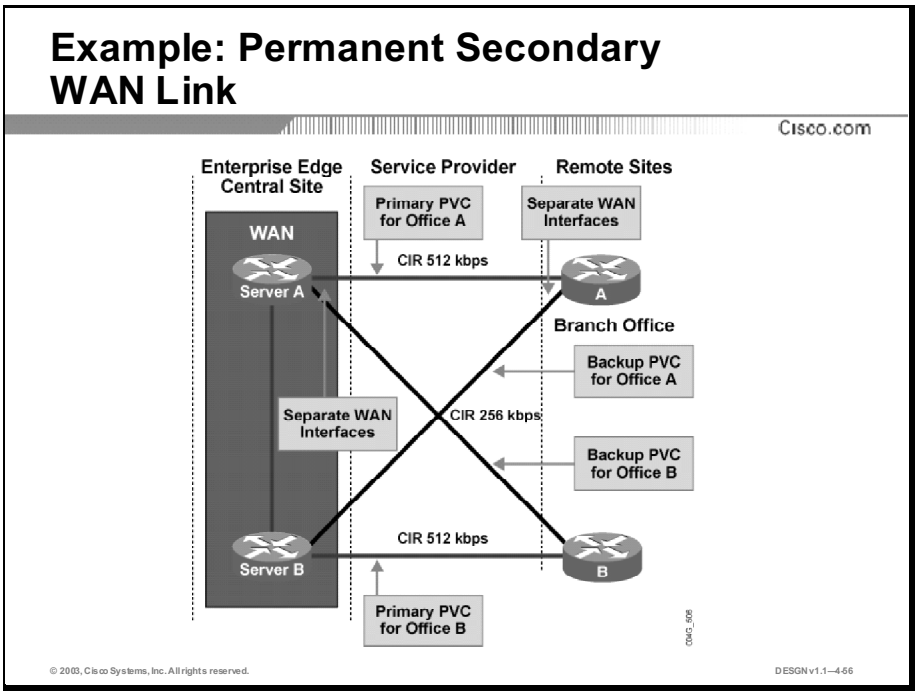

In the figure, the connections between the Central Site Enterprise Edge and remote sites use permanent primary and secondary WAN links for redundancy. To increase the utilization of the backup link, a routing protocol such as the Enhanced Interior Gateway Routing Protocol (EIGRP) is used to support load balancing over unequal paths on either a per-packet or a per-destination basis.

Load-Balancing Guidelines

If the WAN connections are relatively slow (for example, less than 56 kbps), use per-packet load balancing. If WAN connections are faster than 56 kbps, fast switching on the routers is more appropriate. Load balancing occurs on a per-destination basis when fast switching is enabled.

Cost is the primary disadvantage of duplicating WAN links to each remote office. In large star networks with more remote sites, 10 or 20 new virtual circuits might be needed, as well as new equipment such as new WAN router interfaces.

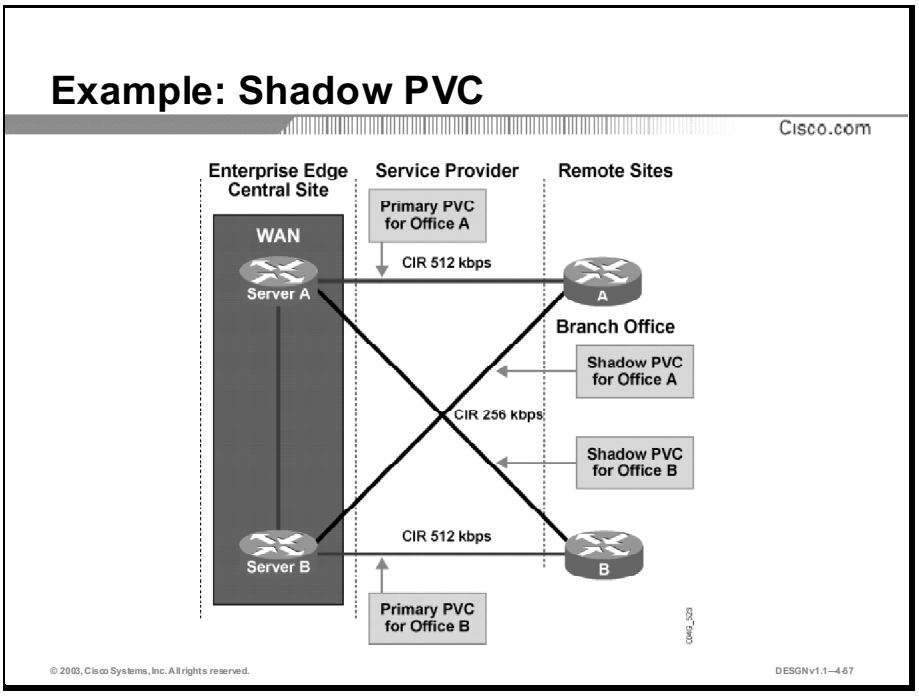

The figure illustrates redundant connections between remote sites and the Enterprise Edge using the shadow PVCs from the SP. The SP charges for two PVCs if the traffic limits are exceeded on the shadow PVC while the primary PVC is available. Consequently, it is very important that the routers not send any unnecessary data over the shadow PVC.

Designing WAN Backup over the Internet

You can use the Internet as an alternate option for a failed WAN connection. This type of connection is considered “best effort” and guarantees no bandwidth. This topic describes designing a WAN backup over the Internet.

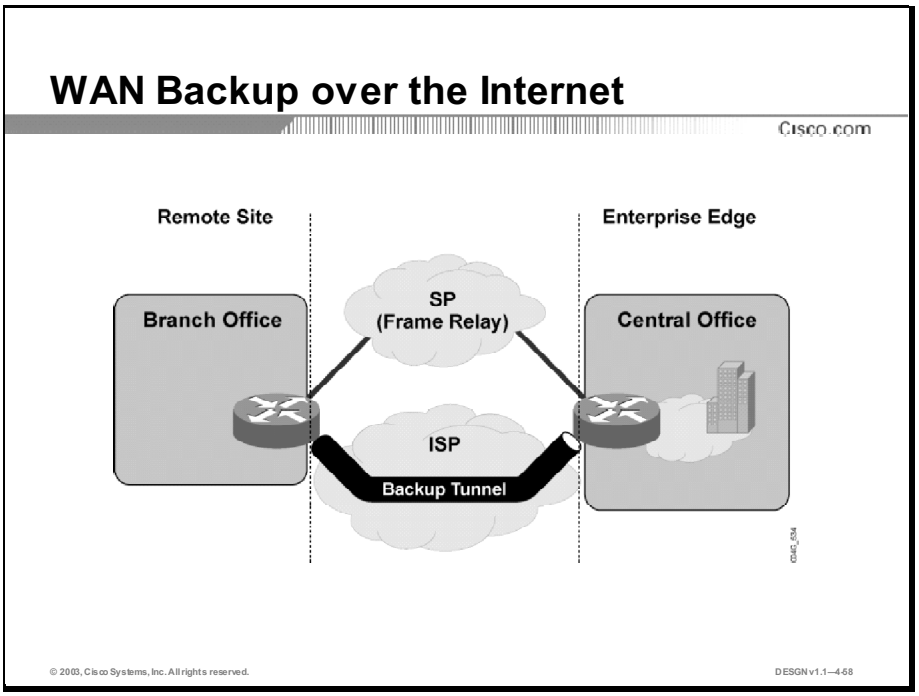

When relying on the Internet to provide a backup for branch offices, the enterprise must cooperate fully with the ISP and announce its networks to gain connectivity. The backup network thus becomes aware of the data as it is sent unencrypted.

The figure illustrates two noncontiguous networks connected over a point-to-point logical link implemented over an IP network using a Layer 3 tunnel. Such a tunnel is configured between a source (ingress) router and a destination (egress) router and is visible on each router as an interface.

The packets to be forwarded across the tunnel are already formatted with an encapsulation of the data with the standard protocol-defined packet header. The packets are further encapsulated with a new GRE or IPSec header, and placed into the tunnel with a destination address of the tunnel endpoint (the new next-hop). When the packet reaches the tunnel endpoint, the GRE or IPSec header is stripped away and the packet continues to be forwarded to the destination with the original IP packet header.

Layer 3 Tunneling with GRE and IPSec

Layer 3 tunneling uses a network layer protocol to transport over another Layer 3 network. Usually, Layer 3 tunneling is used either to connect two noncontiguous parts of a non-IP network over an IP network or to connect two IP networks over a backbone IP network. This can possibly hide the IP addressing details of the two networks from the backbone IP network.

These are the two methods of connecting noncontiguous private networks over a public IP network:

- **GRE**: A standardized Layer 3 carrier encapsulation, designed for generic tunneling of protocols. In IOS software, GRE tunnels IP over IP, which is useful when building a small-scale IP VPN network that does not require substantial security.
- **IPSec**: Provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts as the network layer, in tunneling or transport mode, and protects and authenticates IP packets between participating IPSec devices.

The table compares the characteristics of the GRE method versus the IPSec method:

GRE to IPSec Comparison

GRE	IPSec
GRE enables simple and flexible deployment of basic IP VPNs. Provisioning of tunnels is not very scalable in a full-mesh network because you must define every point-to-point association separately. Packet payload is not protected against sniffing and unauthorized changes, and there is no sender authentication. Using GRE tunnels as a mechanism for backup links has several drawbacks, including administrative overhead, scaling to large numbers of tunnels, and processing overhead of GRE encapsulation.	Data Confidentiality: An IPSec sender can encrypt packets before transmitting them across a network. Data Integrity: An IPSec receiver can authenticate packets sent by an IPSec sender to ensure that the data has not been altered during transmission. Data Origin Authentication: An IPSec receiver can authenticate the source of the IPSec packets sent. This service is dependent upon the data integrity service. Anti-Replay: An IPSec receiver can detect and reject replay. Easy deployment, with no change to the intermediate systems (ISP backbones). No change to existing applications (transparent). Uses the Internet Key Exchange (IKE) for automated key management. Has interoperability with public-key infrastructure (PKI). IPSec can be combined with GRE.

Summary

This topic summarizes the key points discussed in this lesson.

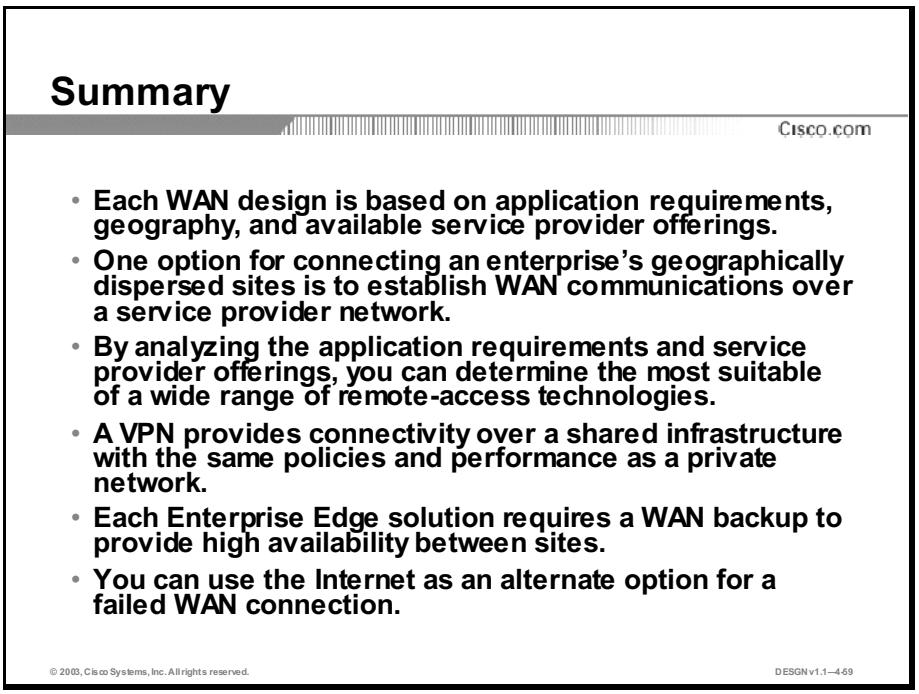

References

For additional information, refer to these resources:

- *Virtual Private Networks,* http://www.cisco.com/warp/public/779/largeent/learn/technologies/VPNs.html
- *Virtual Private Networks,* http://www.cisco.com/warp/public/779/servpro/services/vpn/

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Case Study 4: WAN Upgrade and Backup

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which statement is true when selecting a WAN technology?

A) X.25 is a legacy technology that is still in use in various environments, but is being replaced by faster and more efficient technologies such as Frame Relay.

B) In time-division multiplexing (TDM) networks, the network resources are shared dynamically and subscribers are charged on the basis of their use of the network.

C) X.25, Frame Relay, TDM, and ISDN concentrate multiple remote sites over a single physical connection at the central site.

D) The quality of data transport is significantly better for Frame Relay than for ATM.

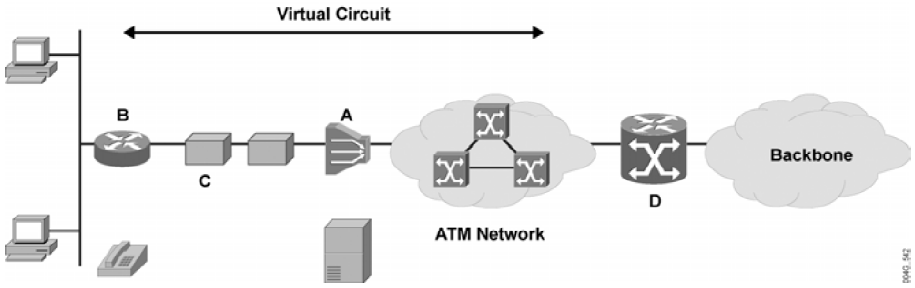

Q2) Identify the key ADSL devices in the picture.

_____ 1. L3 concentrator

_____ 2. L2 concentrator—DSLAM

_____ 3. splitter

_____ 4. ADSL CPE

Q3) Which two statements are true when describing the operation of cable networks? (Choose two.)

A) The Universal Broadband Router (uBR) at the headend of the network, also referred to as the cable modem termination system (CMTS), enables the coaxial users to connect with either the PSTN or the Internet.

B) The cable modem at the customer location operates in DOCSIS bridge mode.

C) In cable networks, bandwidth is evenly shared between upstream and downstream transmissions.

D) Provisioning to support return traffic is mainly required for video services.

Q4) Which two encapsulations are supported in a SONET/SDH network? (Choose two.)

A) IP (POS)

B) IP over DWDM

C) IP over ATM

D) PPP

Q5) _____is the type of the connection over SONET/SDH.

A) point-to-point

B) broadcast

C) point-to-multipoint

D) add-and-drop

Q6) Which three VPN models does the ISP infrastructure support? (Choose three.)

A) Layer 2 overlay model

B) Layer 3 overlay model

C) VPDN model

D) peer-to-peer model

Q7) Match the lettered backup solutions with the numbered functional descriptions of events.

A) dial-in

B) dial-out

C) dial-on-demand routing

D) dial backup

_____ 1. The router is configured to initiate the call when certain criteria are met.

_____ 2. The router is configured to detect a failure on a primary circuit and enable the secondary path.

_____ 3. Only remote sites call the central location.

_____ 4. Remote sites can dial to the central location.

Q8) Which method is the most appropriate when considering an Internet connection as a secure backup solution for an Enterprise WAN link?

A) IPSec tunnels

B) GRE tunnels

C) IP routing without encryption

D) IP routing with encryption

Quiz Answer Key

Q1) A

Relates to: Designing the Classic WAN

Q2) 1=D, 2=A, 3=C, 4=B

Relates to: Designing a Remote-Access Network

Q3) A, B

Relates to: Designing a Remote- Access Network

Q4) A, C

Relates to: Using a Service Provider Network to Connect Dispersed Enterprise Sites

Q5) A

Relates to: Using a Service Provider Network to Connect Dispersed Enterprise Sites

Q6) B, C, D

Relates to: Designing Virtual Private Networks

Q7) 1=C, 2=D, 3=A, 4=B

Relates to: Designing a WAN Backup Strategy

Q8) A

Relates to: Designing WAN Backup over the Internet

Case Study 4: WAN Upgrade and Backup

Complete this case study to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module, "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the WAN transport options and backup strategies. Upon completing this case study, you will be able to meet these objectives:

- Select the most optimal WAN transport
- Select the appropriate WAN backup strategy

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Complete these steps:

Step 1 Completely refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario before the exercise. Focus on the WAN issues and on the new international offices integration. Please also check the New Applications Simulation at the end of the module "Applying a Methodology to Network Design" for the results of WAN link simulation under the new expected load. Allow a maximum of 10 minutes for reading. Take into account that the current WAN bandwidths are 64 kbps and that all the links are leased lines. The existing equipment supports upgrades to higher speeds (no changes to router interfaces and cables are needed; there are all synchronous serial interfaces up to 2 Mbps). The routers at regional offices are already equipped with ISDN BRI interfaces. The central router at the headquarters has one WAN slot still free.

Step 2 Discuss the scenario and options for WAN upgrade and WAN backup with your group. Allow 10 minutes for a discussion.

Step 3 Propose the optimal WAN upgrade and WAN backup scenarios. Please address the international offices as well.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class and you have justified any major deviations from the case study solution.

Module 5

Designing IP Addressing for the Network

Overview

An efficient IP addressing solution is key for addressing and routing on an IP network. The module begins with an overview of IP addressing and general considerations when planning a network addressing scheme. It continues with a discussion of the specific considerations for IP version 4 (IPv4) and IP version 6 (IPv6). The discussion then continues with specific IPv4 and IPv6 considerations including migration strategies.

Module Objectives

Upon completing this module, you will be able to create an IP address plan for a network.

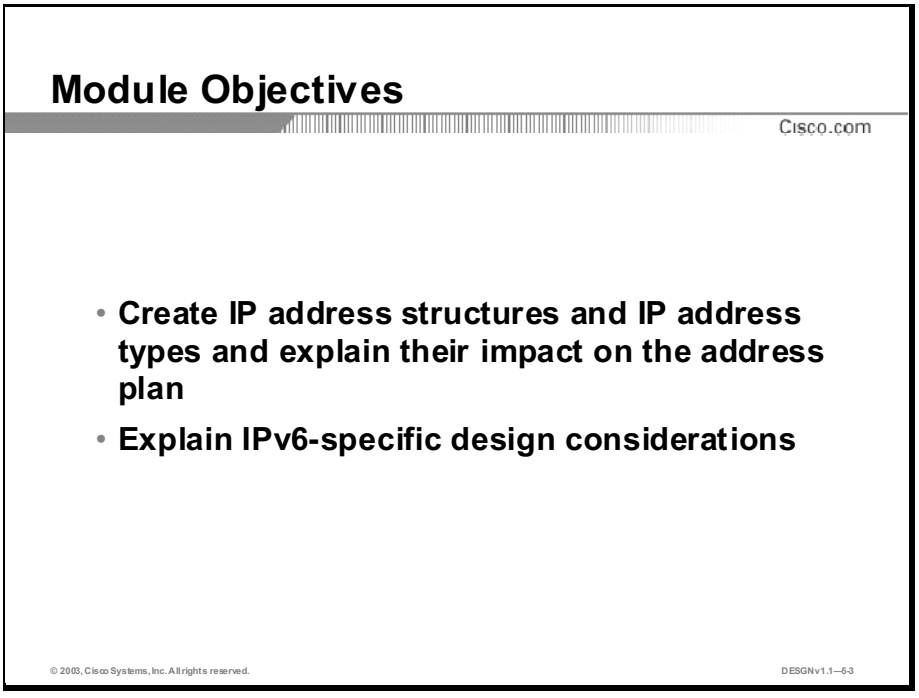

Module Outline

The outline lists the components of this module.

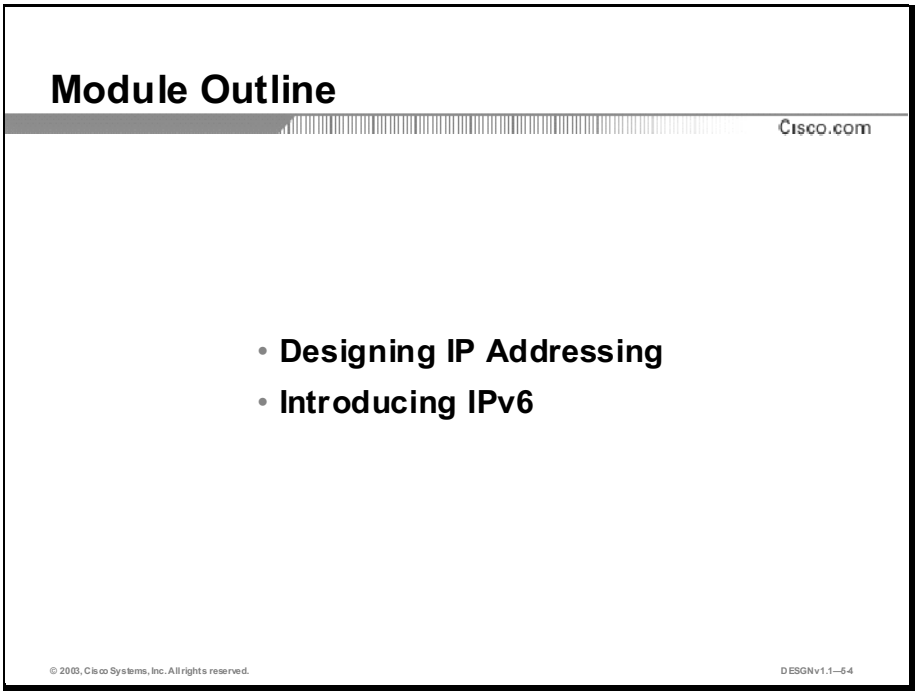

Designing IP Addressing

Overview

This lesson explains IP address structures, various IP address types, and their impact on the address plan. It discusses considerations about the routing protocol choice, describes various IP address assignment strategies, and explains name resolution.

Relevance

An effective and efficient IP addressing scheme is a critical component of the overall enterprise network design.

Objectives

Upon completing this lesson, you will be able to create IP address structures and IP address types and explain their impact on the address plan. This includes being able to meet these objectives:

- Describe the IPv4 address structure and classes
- Determine the network size based on the IP addressing plan
- Explain when to use private and public IP addresses
- Explain the impact of hierarchical addressing on routing protocol choice
- Describe and implement the appropriate method to determine host IP addressing
- Describe name resolution and its implementation in an enterprise network

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions

Outline

The outline lists the topics included in this lesson.

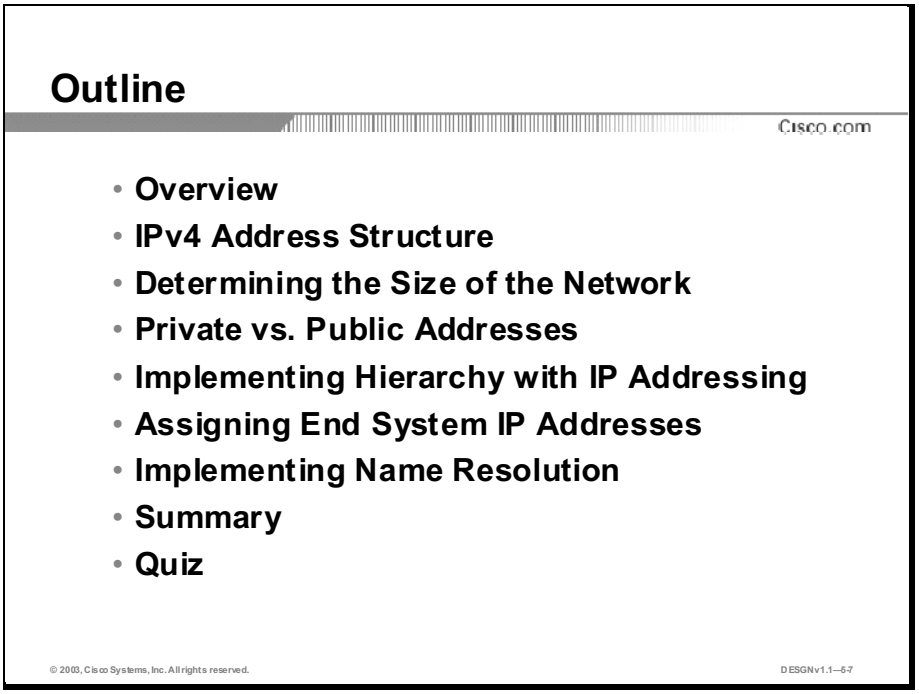

IPv4 Address Structure

IP addresses, IPv4 address structure, address classes, subnetting and masking are all key components of the IP addressing scheme. This topic describes the IPv4 address structure.

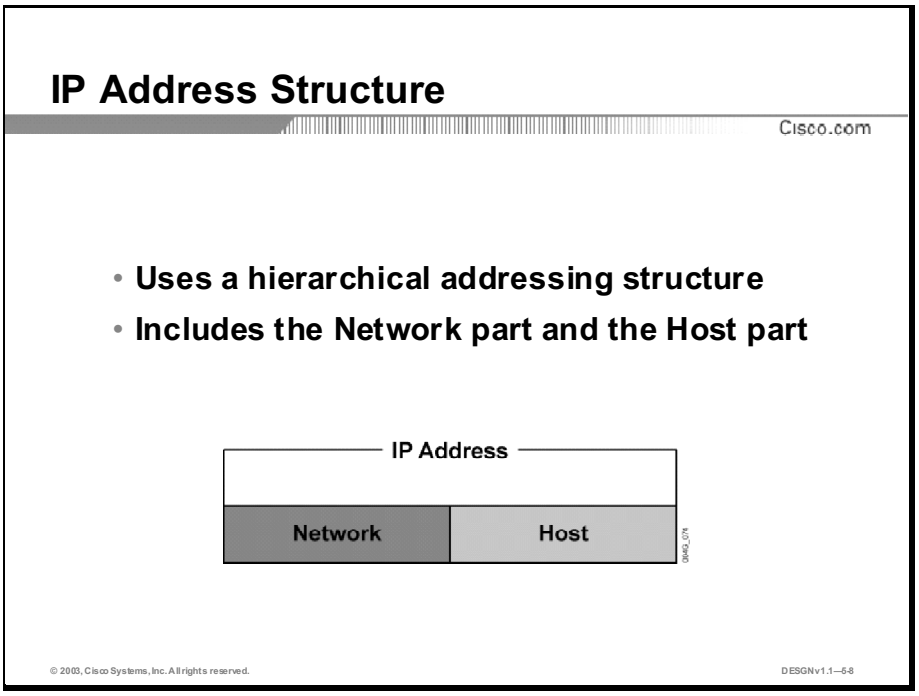

An IP address is either hierarchical or flat:

- **Hierarchical address:** Consists of related parts, which denote a hierarchy and provide more flexibility when locating the destination address. These are examples of a hierarchical address:
  - **A telephone system:** To make a long-distance call to another country, a user dials the country prefix followed by an area code, and finally the telephone number. The hierarchy is 1) the country code, 2) the area code, 3) the first three digits of the telephone number, and 4) the last four digits of the telephone number, which represent the individual telephone.
  - **A postal address:** This describes a person's location by country, state (where applicable), city, street address, and person's name. A postal code or ZIP code flattens the hierarchy. It defines state or province and city.
- **Flat address:** This type of address consists of only one part, and only the complete address has a meaning. Examples of a flat address are:
  - **MAC-layer addresses:** Each device on a LAN has a MAC-layer address that is used to communicate at the data link layer. The hierarchical manufacturer information is not operationally interesting.
  - **Various personal IDs:** A person has a social security number and library identification number.

IP addresses are hierarchical. IP addresses define networks and devices (nodes, hosts) associated with each network. The IP address consists of two parts:

- The network portion, which identifies a specific network. Routers use the network portion to decide where to send a packet (datagram).
- The host portion, which identifies the specific device in a network.

Each device is known in an internetwork by its unique IP address.

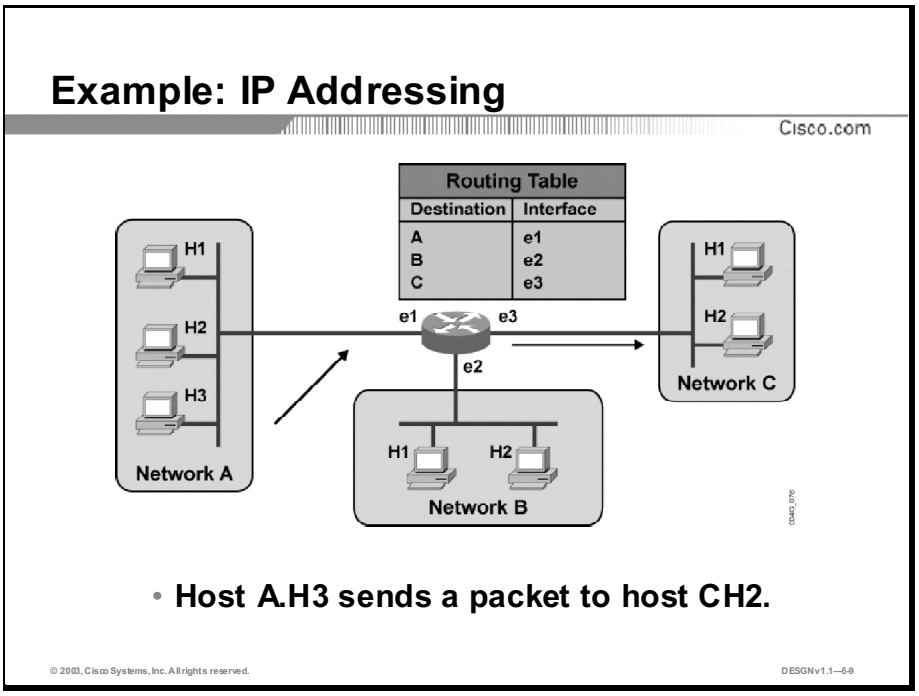

A router uses an IP address to route a packet. When a router receives a packet, the routing decision is made based upon the destination address, that is, its network part. The router looks up the destination address in its routing table and forwards the packet to the next hop or interface specified in its routing table.

After a packet leaves the router, the next-hop router forwards the packet to its final destination. If the router does not have a destination network in its routing table, the router forwards the packet to a predetermined default gateway, if configured; otherwise, the router discards the packet and informs the sending host that the network is unreachable. The routing learns known networks using these methods:

- Dynamic routing protocols (for example, Routing Information Protocol [RIP], Open Shortest Path First [OSPF], Intermediate System-to-Intermediate System [IS-IS], Enhanced Interior Gateway Protocol [EIGRP], Border Gateway Protocol [BGP])
- Static routes, which a network administrator manually enters
- Networks that are directly connected to a router interface

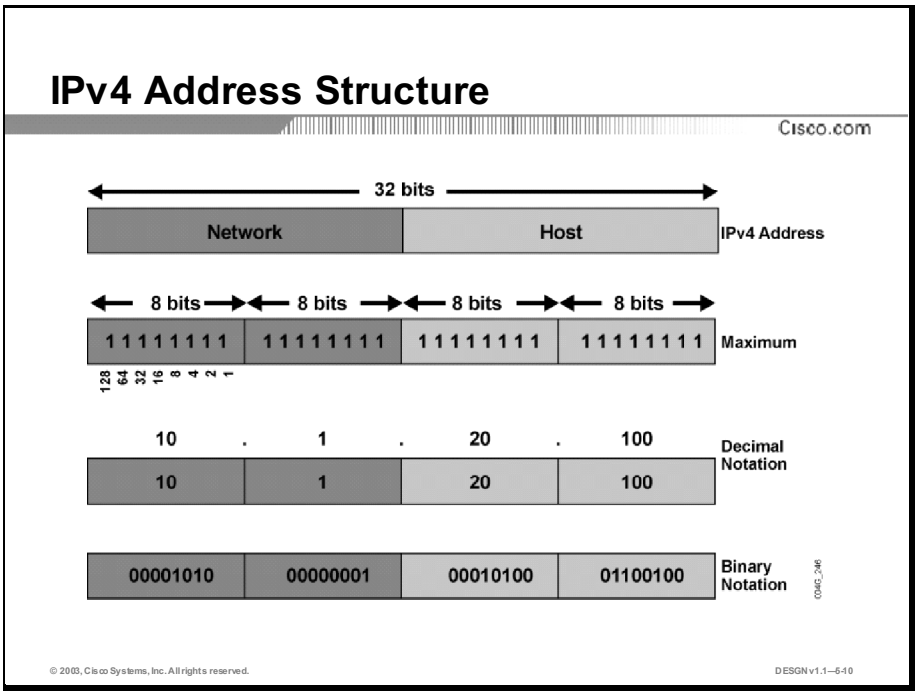

These are some characteristics of an IPv4 address:

- Hierarchically structured and consists of network and host parts. Such an address identifies a specific network and a specific host on that particular network.
- 32 bits long and divided into 4 octets (groups of 8 bits). Each bit in an octet has a binary weight (128, 64, 32, 16, 8, 4, 2, 1), with a minimum value of 0 (all bits set to 0) and maximum value of 255 (all bits set to 1).
- Usually written in dotted decimal notation. Each octet is divided by a point and written as a decimal value. Binary notation shows octets written with individual bits set to either 1 or 0. Hexadecimal, base 16, notation is sometimes used.

Example: IPv4 Addresses

Here are two examples of IPv4 addresses:

- 170.1.20.100, written in binary notation as 10101010.00000001.00010100.01100100
- 192.168.1.1, written in binary notation as 11000000.10101000.00000001.000000001

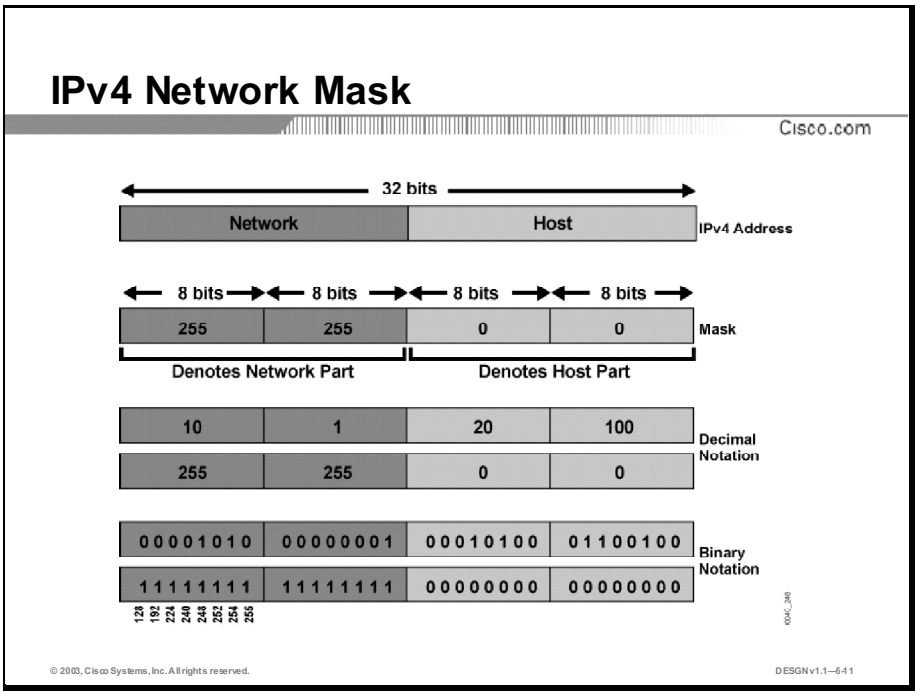

The network mask provides a distinction between the network part and the host part of the IP address. A network mask is 32 bits long, similar to the IPv4 address, is divided into four octets. As with the IPv4 address, each bit in an octet has a binary weight (128, 64, 32, 16, 8, 4, 2, 1). The mask consists of a specified number of recursive ones, with the remaining bits set to zero. The mask is used to interpret the address. The location where ones stop and zeros begin indicates the boundary between the network and host parts of the address. Write a network mask either in decimal notation (255.255.0.0) or in prefix notation (/16), which indicates the number of consecutive ones.

Calculating the Number of Host Addresses

To determine the number of available host addresses, use a 2^n-2, where n is the number of bits set to zero (that is, the mask). Subtract two because you cannot represent a host with a host part of either all zeros or all ones. You can determine the first IPv4 address available for hosts in a particular network, for the host part, by setting the last bit (reading from right to left) to one and all others to zero. You can determine the last IPv4 address available by setting the last bit in the host part to zero and all other bits to one.

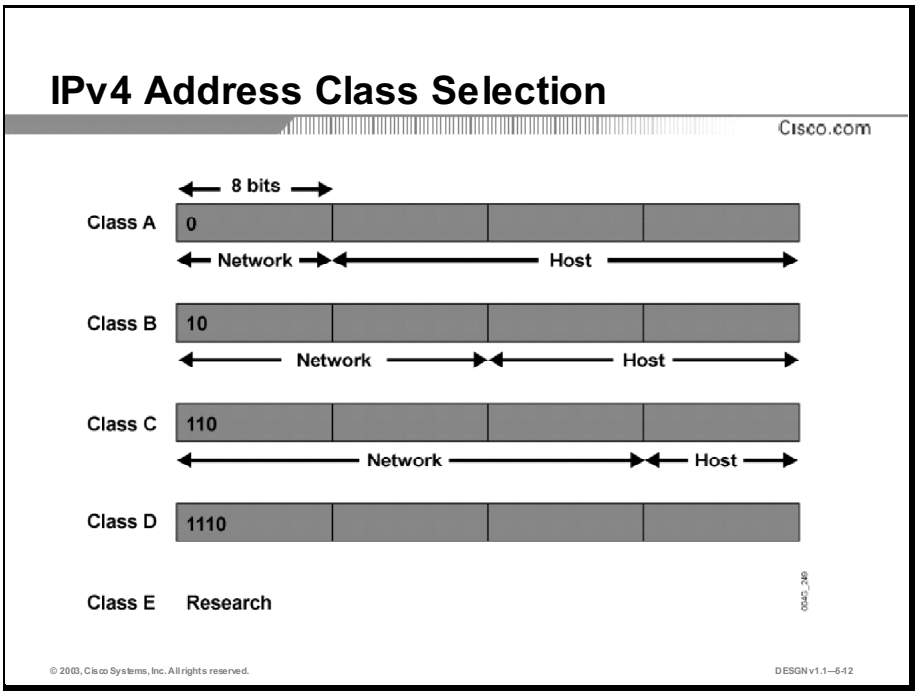

IPv4 address space is divided into address classes. The first octet rule and network mask determine an address class.

- **Class A:** The first octet starts with binary 0, resulting in a range of 0 to 127 for the first octet (0 and 127 are reserved). The network mask is set at 255.0.0.0 or /8. An individual Class A network has 16,777,214 available host addresses, if no subnetting is done.
- **Class B:** The first octet starts with binary 10, resulting in a range of 128 to 191 for the first octet. The network mask is set at 255.255.0.0 or /16. An individual Class B network has 65534 available host addresses, if no subnetting is done.
- **Class C:** The first octet starts with binary 110, resulting in a range of 192 to 233 for the first octet. The network mask is set at 255.255.255.0 or /24. An individual Class C network has 254 available host addresses, if no subnetting is done.
- **Class D:** The first octet starts with binary 1110, resulting in a range of 224 to 239 for the first octet. Class D is reserved for multicast addresses and you cannot use it to address hosts. Multicast addresses are, for example, used with the OSPF routing protocol (224.0.0.5, 224.0.0.6) and with the EIGRP routing protocol (224.0.0.9).
- **Class E:** The first octet starts with 1111, resulting in a range of 240 to 255 for the first octet. Class E is reserved for research.

The public number authority and Internet service providers (ISPs) assign the A, B, and C address classes to the individual applicants.

Example: Identifying the Address Class

To determine the appropriate address class for a certain address, inspect the first octet:

- The IPv4 address 193.18.9.45 is a Class C address because the 193 falls into the 192 to 223 range.
- The IPv4 address 172.31.1.2 is a Class B address because the 172 falls into the 128 to 191 range.

- The network consists of one location and requires 180 IP addresses. The required IPv4 address class is Class C, which offers 254 host addresses.
- The network consists of three locations and requires 490 IP addresses (300 for the first location, 100 for the second location, and 90 for the third location). The network size requires at least two Class C addresses, which are partitioned into smaller networks or subnets.

Defining Subnets

Subnetting uses a network mask (subnet mask) that is longer than the default mask for a certain address class. When subnetting, some bits from the host part are used for the network part. Subnetting provides network administrators with extra flexibility, more efficiently uses network addresses, and contains broadcast traffic, because a broadcasted traffic does not cross a subnet's boundaries (a router interface).

You can divide a given network address into many subnetworks. For example, 172.16.1.0, 172.16.2.0, 172.16.3.0, and 172.16.4.0 are all subnets of the Class B network 171.16.0.0.

Example: Subnet Masks

These are two examples of subnet masks:

- The IPv4 address is 172.1.20.5. The mask is 255.255.255.192 or /26.
  - Binary representation is 10101100.00000001**.**00010100.00000101.
  - Network is 172.1.20.0, with 26 bits used to denote the network part.
  - Host part is .5, with 6 bits used to denote the host part.
  - Number of addresses available to hosts is 2^6-2=62.
  - First host address is 172.1.20.1 and the last one is 172.1.20.62.
- The IPv4 address is 10.200.200.25. The mask is 255.255.128.0 or /17. In this case, an extra 9 bits are used to divide the network (originally network mask /8) further:
  - Binary representation is 00001010.11001000.11001000.00011001.
  - Network is 10.200.128.0, with 17 bits used to denote the network part.
  - Host part is 72.25, with 15 bits used to denote the host part, but when the boundary is not on the octet, the decimal notation for the host is nonsensical.
  - Number of addresses available to hosts is 2^15-2=32766.
  - First host address is 10.200.128.1, and the last one is 10.200.255.254.

Determining the Size of the Network

The first step in an IP addressing plan design is to determine the size of the network to establish how many IP addresses are needed. This topic describes how to determine the size of a network for an IP addressing plan.

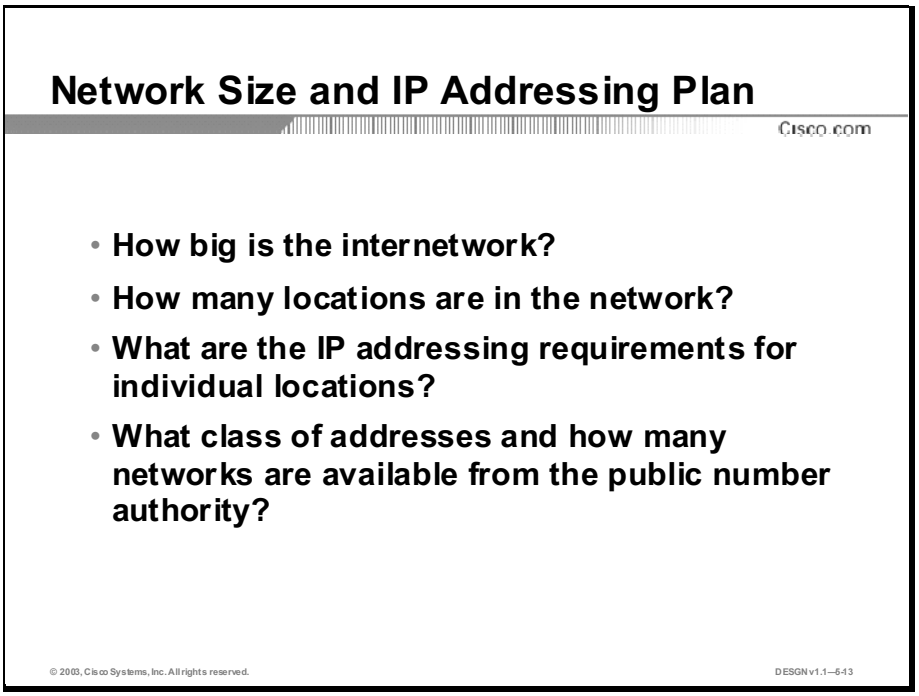

To determine the network size, answer these questions:

- **How big is the internetwork?** Determine the number of end systems, router interfaces, switches, firewall interfaces, and so forth.
- **How many locations are in the network?** Determine the number of locations and also identify their type.
- **What are the IP addressing requirements for individual locations?** Collect information about which systems will use Dynamic Host Configuration Protocol (DHCP) and which will use static addresses. In addition, obtain information on which systems can use private addresses instead of public addresses.
- **What class of addresses and how many networks are available from the public number authority, directly or indirectly?** Make decisions based on the collected information about the network size to apply for the required number of addresses. The Internet Assigned Numbers Authority (IANA) assigns addresses to regional authorities and ISPs for redistribution.

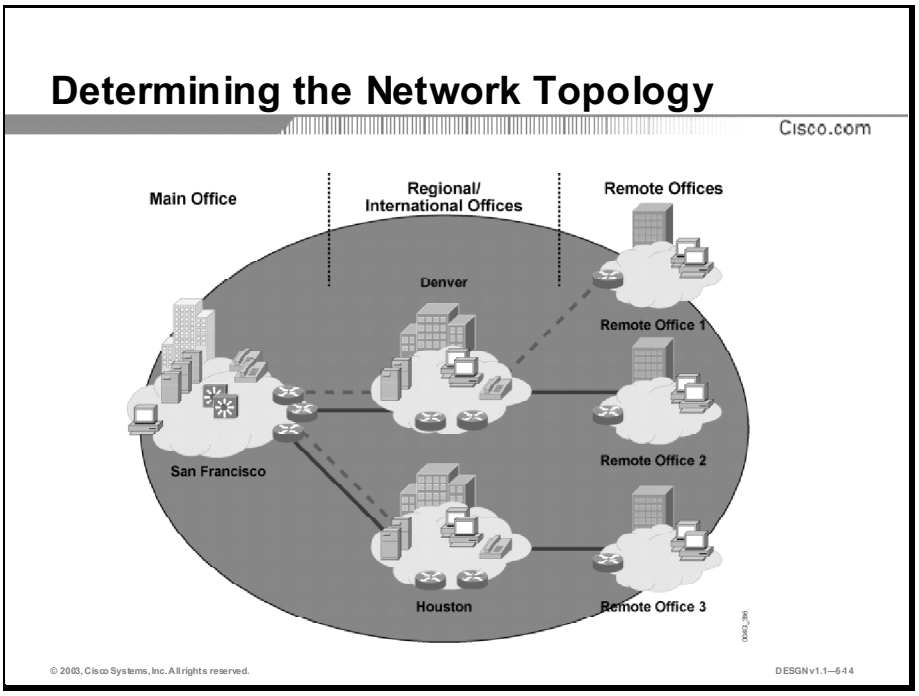

To gather the correct information about network size and its relation to the IP addressing plan, you should have a general picture of the network topology. With the general network topology information, you can determine the number of locations, location types, and their correlations. This information is the foundation for gathering the data about the network size and individual location requirements related to the IP addressing plan. The information for the figure is shown in the table.

Network Locations

Location	Type	Comments
San Francisco	Main office	The central location where the majority of users are located
Denver	Regional office	Connects to the San Francisco main office
Houston	Regional office	Connects to the San Francisco main office
Remote office 1	Remote office	Connects to the Denver regional office
Remote office 2	Remote office	Connects to the Denver regional office
Remote office 3	Remote office	Connects to the Houston regional office

Network Locations Worksheet

Use the table to record your own network information.

Location	Type	Comments

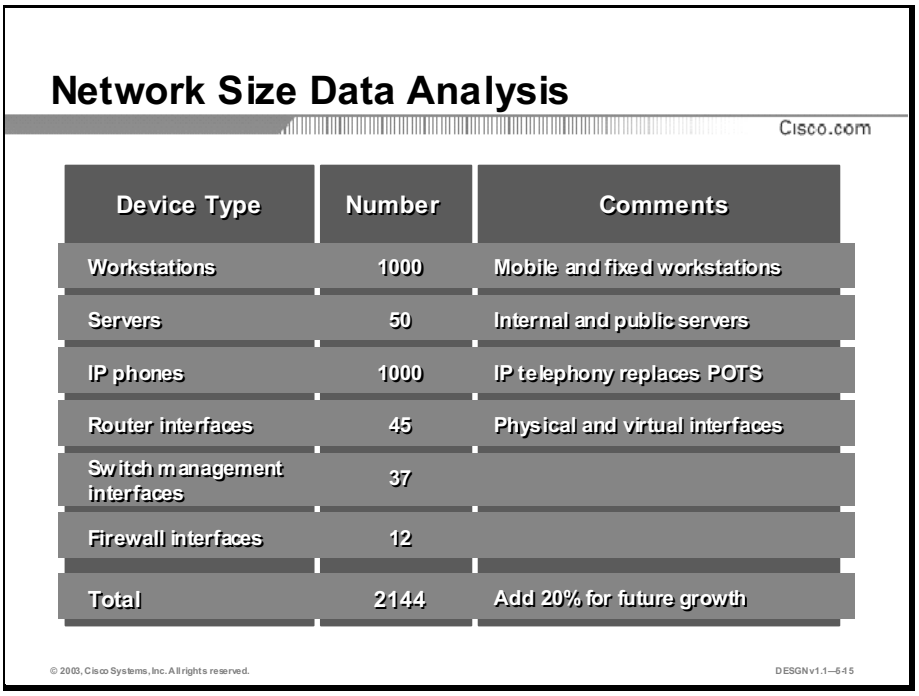

The network size, in terms of the IP addressing plan, relates to the number of devices and interfaces that need an IP address. To establish the overall network size, you must determine the number of workstations, servers, IP Phones, router interfaces, switch management interfaces, firewall interfaces, and so on. The summary provides the minimum overall number of IP addresses that are required to address the network. Because all networks tend to grow, keep a reserve of up to 20 percent for potential network expansion. This information can be presented in a table, as shown in the figure.

Overall Network Size Worksheet

Use the table to record network size information.

Device Type	Number	Comments
SUM		

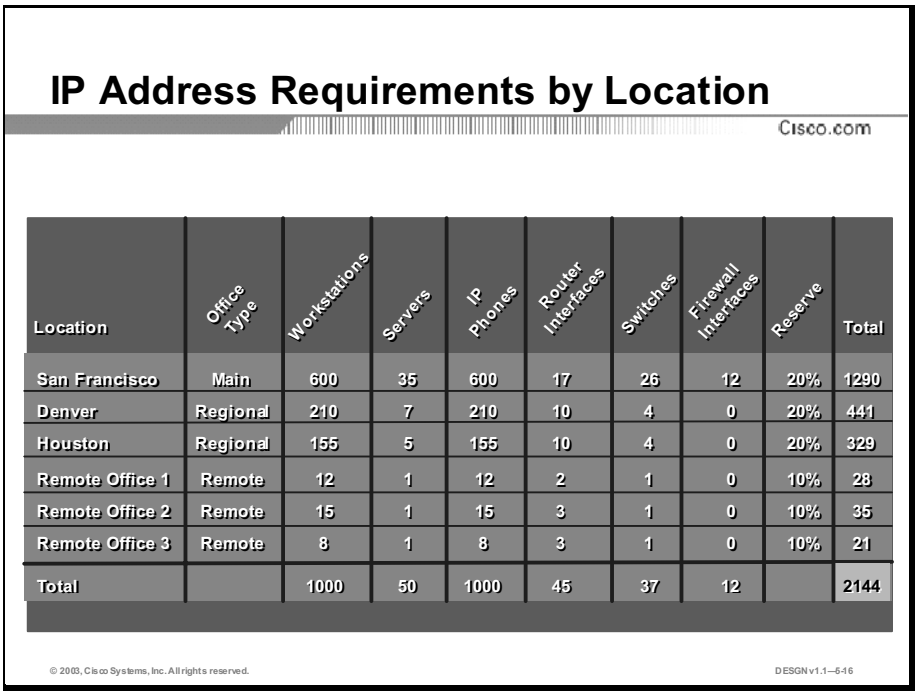

The information about the size of each individual location is closely related to the overall network size. The information about the size of individual locations is vital to determine the IP address range size. You need to be able to assign the appropriate IP addresses to all network devices. Collect the information when you count the number of workstations, servers, IP Phones, router interfaces, switches, firewall interfaces, and so on. Present this information in a table, as shown in the figure.

To allow for seamless network growth, make a reserve. A commonly suggested reserve is 20 percent for the main and regional offices and 10 percent for the remote offices. Carefully discuss the future network growth issue to ensure that you have a precise estimate of the required resources.

Network Locations Size Worksheet

Use the table to record the size of each network location.

Location	Office Type	Work-stations	Servers	IP Phones	Router Inter-faces	Switches	Firewall Inter-faces	Reserve	SUM
SUM									

Private vs. Public Addresses

The available number of public IPv4 addresses is too low, so the public number authority and ISPs assign only a subset of Class C addresses. In many cases, the number of public IPv4 addresses is inadequate to address the whole network. The solution to the problem is private IPv4 addresses. This topic discusses IPv4 private and public addresses and when to use each.

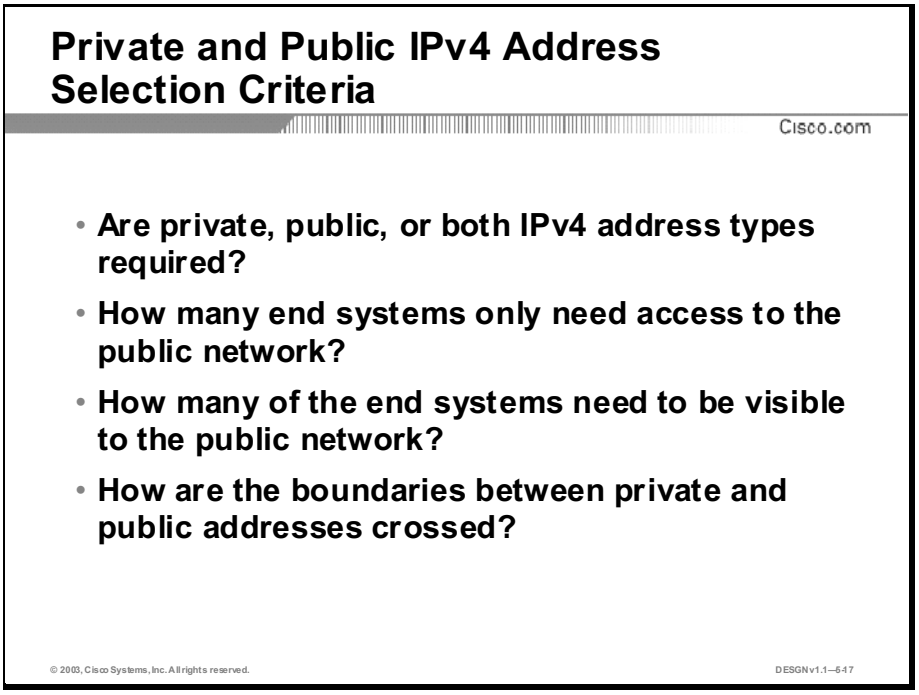

Initially, you must answer the questions listed in the figure.

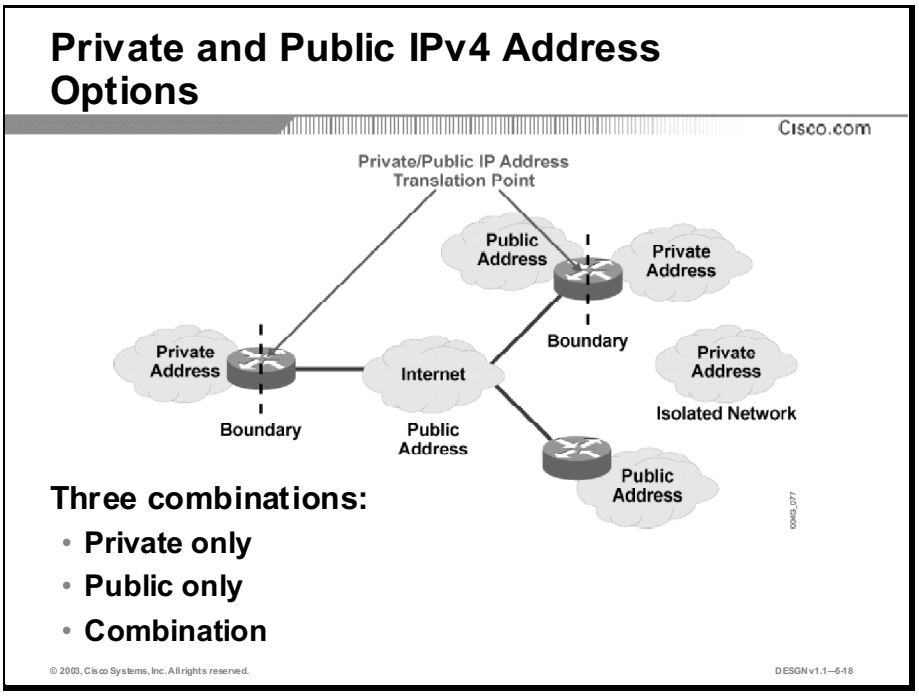

An IPv4 address space is divided into public and private sections. Public IPv4 addresses are used for external communication, and private IP addresses are used for internal communication.

- The public IPv4 address space is used in an internetwork when multiple networks are interconnected to share resources and exchange data.
- The private IPv4 address space is used in a network that requires no connection to the Internet. Such addresses are prohibited for use in public networks. You can reuse a private IPv4 address in different networks.
- A network can use either private IPv4 addresses only, public IPv4 addresses only, or both.

Note: You should save public (legal) addresses for the Internet (if, for example outside Domain Name System [DNS] servers, web servers, or FTP servers need to be added), and use private addresses in the internal network.

These IPv4 addresses are reserved for private use and you cannot use them in public networks:

- 10/8
- 172.16/12 –172.16.0.0 to 172.31.255.255
- 192.168/16

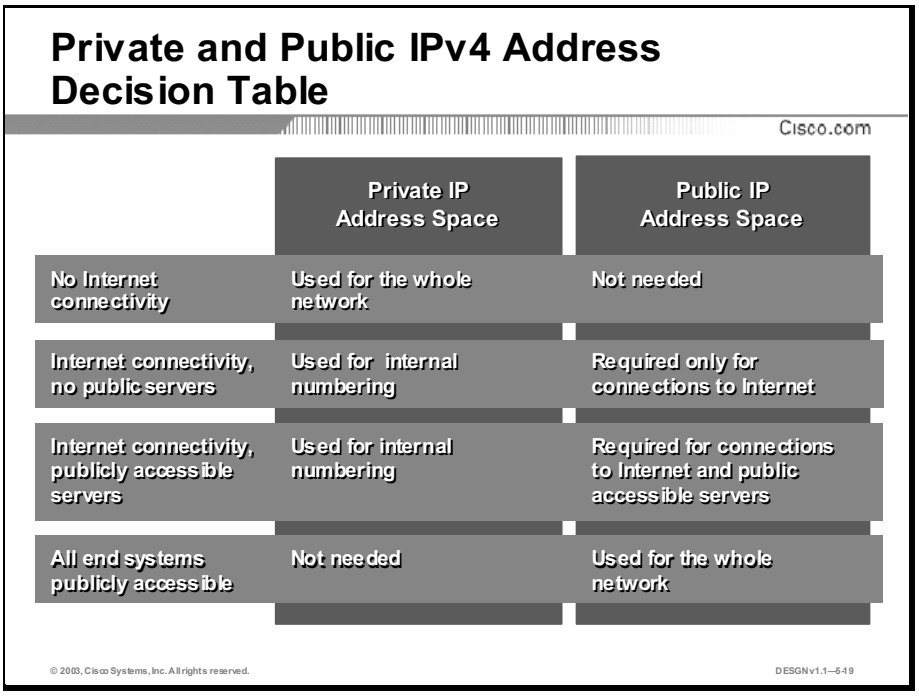

Design Consideration: Are Private, Public, or Both IPv4 Address Types Required?

Private or public IPv4 addresses are required in specific situations. The decision when to use private, public, or both address types depends on the Internet connection presence, the number of public visible servers, and the size of the network.

- **No Internet connectivity:** The network is isolated and there is no need to acquire public IPv4 addresses. You can address the whole network with private IPv4 addresses. Access to the public network is not required.
- **Internet connectivity, no publicly accessible servers:** The network is connected to the Internet and requires public IPv4 addresses. Very limited public IPv4 addresses and translation mechanisms are required to allow access to the Internet. The private IPv4 addresses are used to address the internal network.
- **Internet connectivity, publicly accessible servers:** Acquire public IPv4 addresses to address the connection to the Internet and all publicly accessible servers. The number of public addresses corresponds to the number of Internet connections and publicly accessible servers. Use private IPv4 addresses to address the internal network.
- **All publicly accessible end systems:** Only public IPv4 addresses are used to address the whole network.

Design Consideration: How Many End Systems Need Access to the Public Network Only?

This is the number of end systems that need a limited set of external services (for example, e-mail, FTP, web browsing) and do not need unrestricted external access.

Design Consideration: How Many End Systems Need to Be Visible to the Public Network?

This is the number of Internet connections and various servers that need to be visible to the public (public servers and servers used for e-commerce such as web servers, database servers, and application servers). It defines the number of required public IPv4 addresses. These end systems require IPv4 addresses that are globally unambiguous.

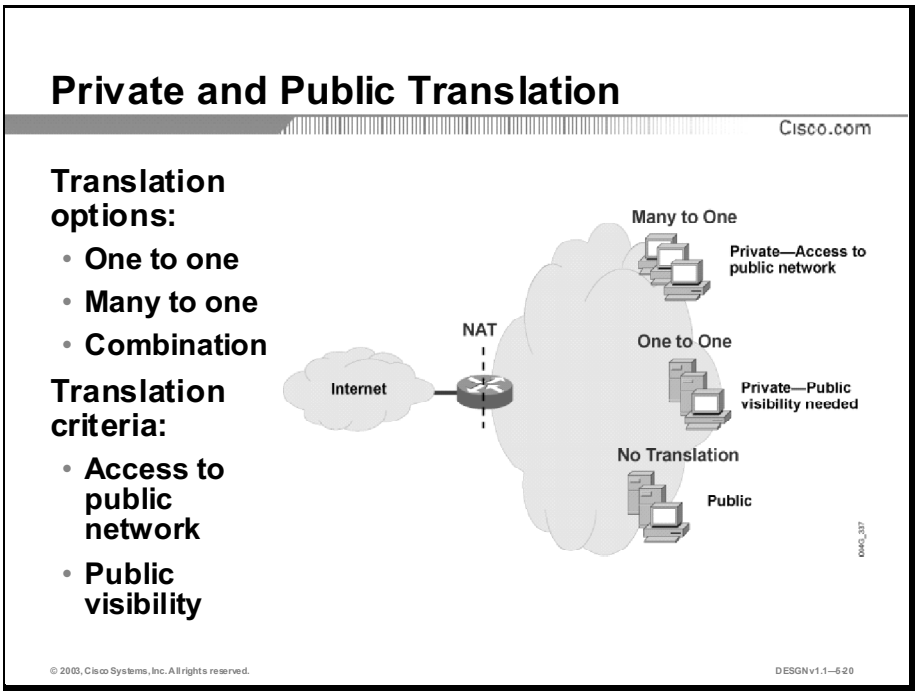

Design Consideration: How Are the Boundaries Between the Private and Public Addresses Crossed?

A network can use both public and private addresses. A router provides the interface between the private and public sections of the network. When private addresses are used and you need to connect the network to the Internet, you will use a translation mechanism such as Network Address Translation (NAT).

The criteria for NAT deployment are accessibility to public network and public visibility. Use NAT to translate:

- **One private address to one public address:** Used when servers from the internal network with private IPv4 addresses must be visible from the public network. Define the translation statically to translate from the public IPv4 address to the server private IPv4 address.
- **Many private addresses to one public address:** Used for end systems that need access to the public network and do not need to be visible to the outside world.
- **Combination:** You can combine both techniques throughout the network.

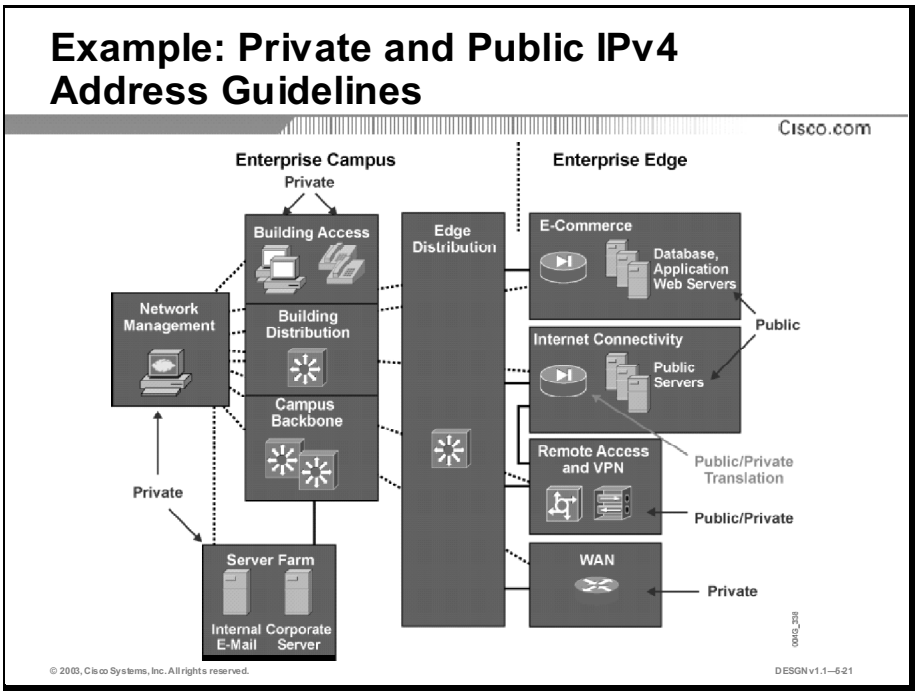

The typical enterprise network uses both private and public IPv4 addresses. Private IPv4 addresses are used throughout the enterprise network, except in:

- The Internet Connectivity module, where public IPv4 addresses are used for Internet connections and public accessible servers
- The E-Commerce module, where public IPv4 addresses are used for the database, application, and web servers
- The Remote Access and VPN module, where public IPv4 addresses are used for selected connections

Implementing Hierarchy with IP Addressing

The IP addressing hierarchy has a major impact on the routing protocol choice and vice versa. Features such as fixed length subnet masking, variable length subnet masking, and classful and classless routing protocols influence the IP addressing plan and the choice of routing protocol. This topic describes how to implement hierarchy with IP addressing.

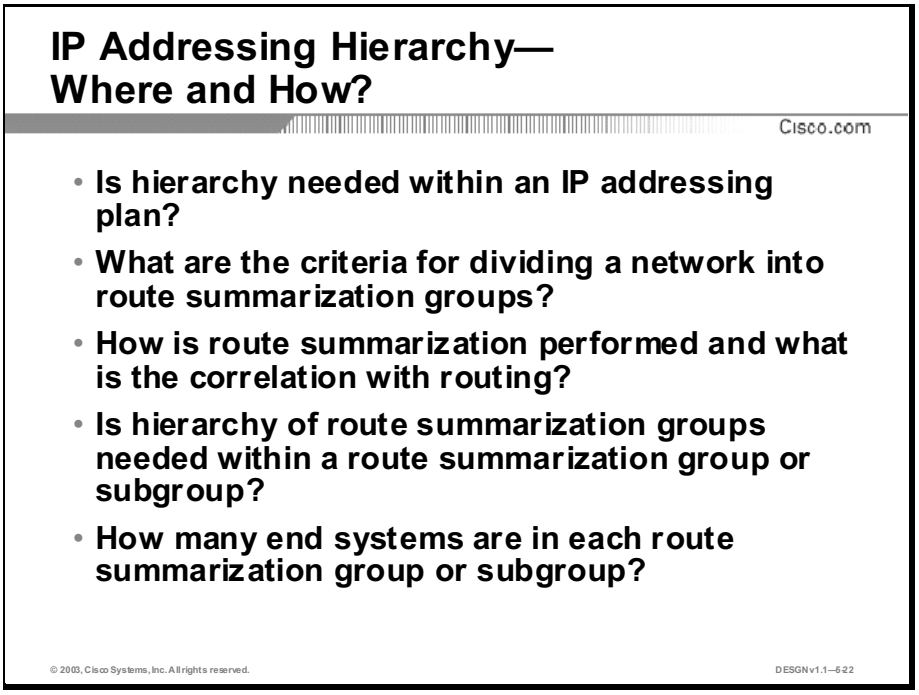

The decision on how to implement the IP addressing hierarchy is usually an administrative decision that is based on the questions listed in the figure.

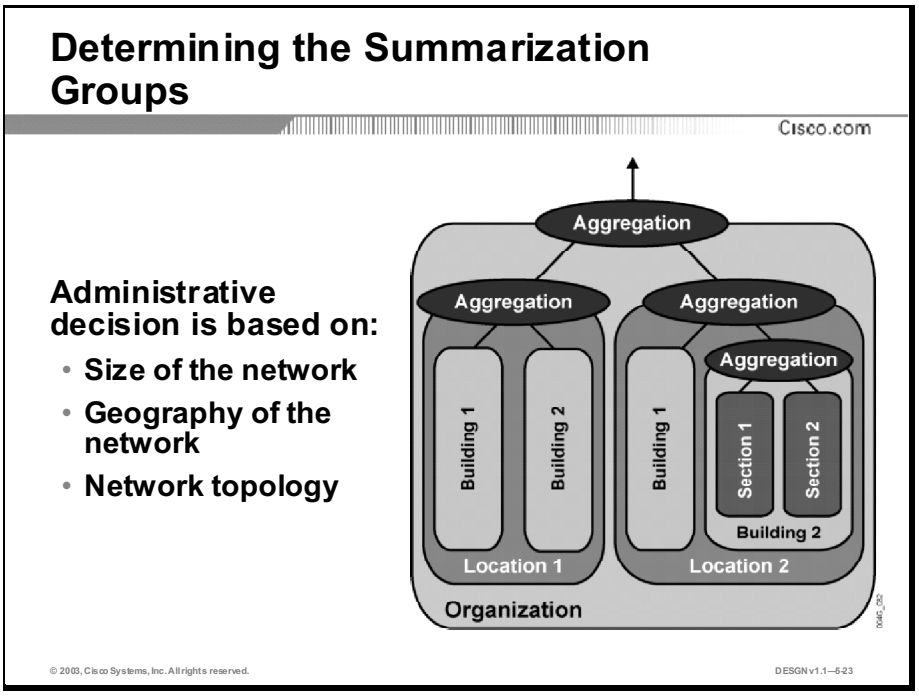

Design Consideration: Is Hierarchy Needed Within an IP Addressing Plan?

You will implement the IP addressing hierarchy based on the network size, geography, and topology. In large networks, hierarchy within the IP addressing plan is mandatory to have a stable network. When deciding whether to implement the hierarchy or not, consider these concerns:

- **Influence of IP addressing on routing:** An IP addressing plan influences the overall routing in the network. Before assigning IP addresses to devices and allocating blocks of IP addresses to different parts of the network, consider the criteria for an appropriate and effective IP addressing scheme. Routing stability, service availability, network scalability, and modularity are some of the key characteristics of every network that IP address allocation and deployment directly affect.
- **Modular design and scalable solutions:** Whether building a new network or adding a new service on an existing infrastructure, a modular design delivers a long-term, scalable solution. IP addressing modularity allows aggregation of routing information on a hierarchical basis.
- **Route aggregation:** To reduce routing overhead and improve the stability and scalability of routing, use route aggregation. To implement route aggregation, you must divide a network into contiguous IP address areas and have a solid understanding of IP address assignment effects on route aggregation (summarization) and hierarchical routing.

Design Consideration: What Are the Criteria for Dividing a Network into Route Summarization Groups?

The size of the network refers to the number of prefixes to advertise. To reduce the routing overhead in a large network, you need to implement a multilevel hierarchy. The depth of the hierarchy depends on the network size and the size of the upper-level summarization group. These are the levels of the hierarchy:

- **First level:** Represented by the locations within the network. Each location is typically a summarization group.
- **Second level:** Done within the first-level summarization group. You can divide a large location into smaller summarization groups represented by the buildings or cities within a certain location. All first-level summarization groups do not require second-level of hierarchy.
- **Third level:** May be performed within the second-level summarization group to minimize further potential routing overhead and instability. The floors within individual buildings may represent the third-level summarization group.

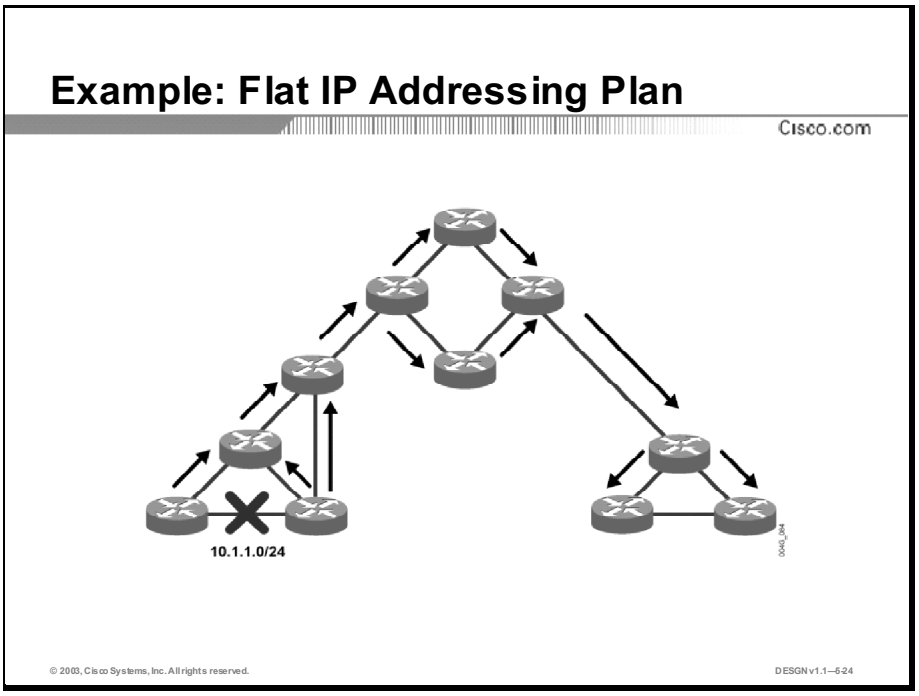

In general, poorly designed IP addressing results in IP network addresses that are randomly assigned on an as-needed basis. In most networks, the IP networks are likely dispersed through the organization. A poor design provides no option to divide the network into contiguous address areas to implement route summarization.

For example, suppose a certain link in part of the network is flapping (changing its state from Up to Down and vice versa) ten times per minute. Because dynamic routing is used, the routers that detect the change send routing updates to their neighbors, the neighbors to their neighbors, and so on. Thus, the routing update is propagated throughout the whole network, even if there is no need for a distance router to have knowledge of that link.

These are some of the results of poorly designed IP addressing:

- **Excess routing traffic consumes bandwidth:** Because of changes, routers need to send routing updates constantly and, thus, the routing traffic consumes more bandwidth.
- **Constant routing table recalculation:** Routing updates require routing table recalculation, which may affect the router's performance and ability to forward traffic.

Note: Well-designed IP addressing enables efficient aggregation of routing advertisements, narrowing the scope of flapping propagation.

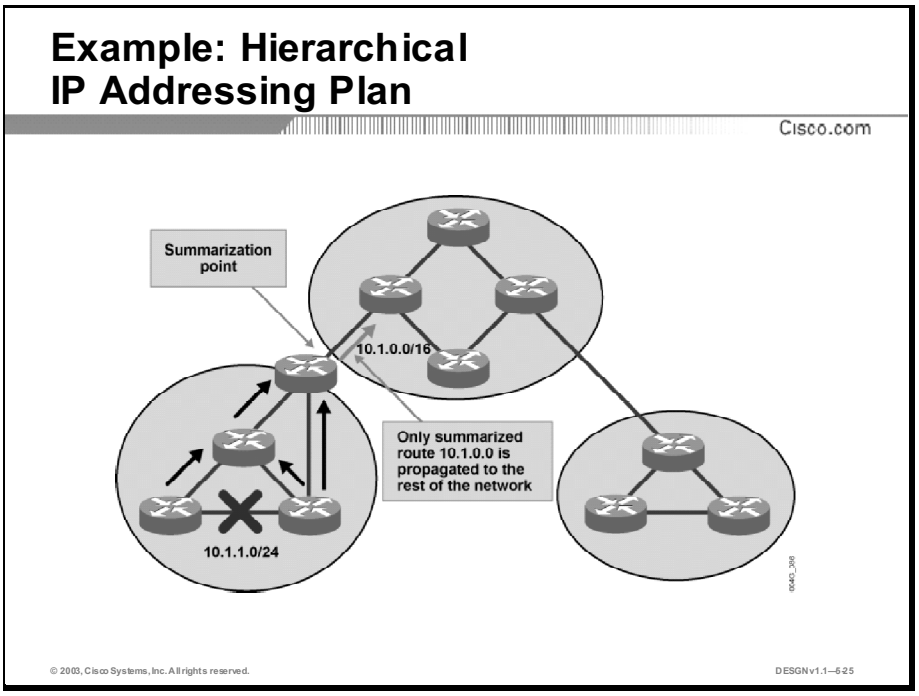

Route aggregation on border routers between contiguously addressed areas improves control over routing table growth. Implement route summarization (aggregation) on the area borders. In case of a link failure, routing updates are not propagated to the rest of the network but stay in the area. This reduces routing overhead from bandwidth consumption and relieves routers from unneeded routing table recalculation.

Note:	Efficient aggregation of routing advertisements narrows the scope of routing update propagation and decreases significantly the cumulative frequency of routing updates.

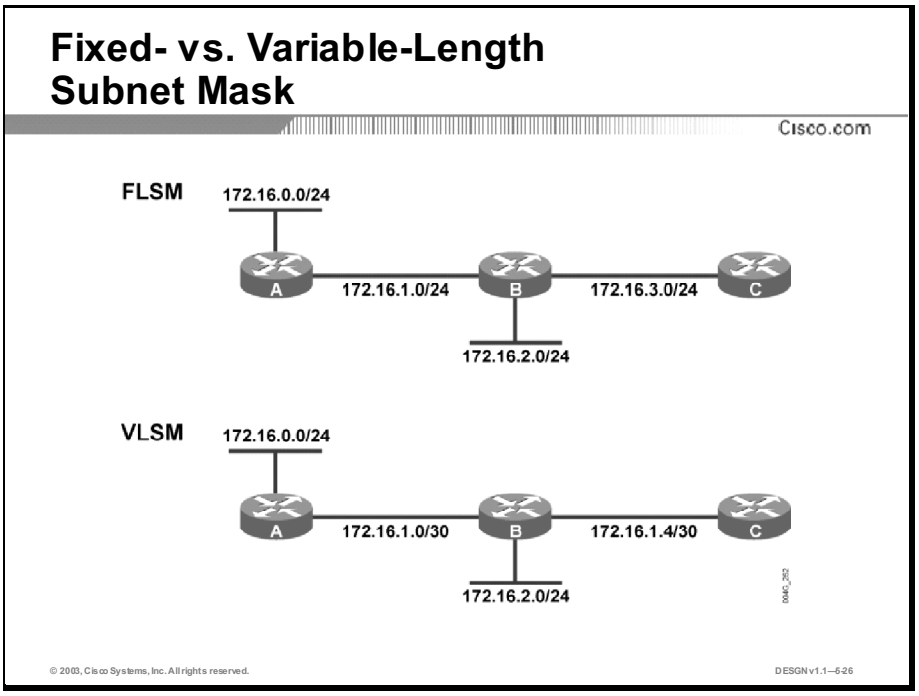

When performing subnetting, you can implement the same or different subnet masks throughout the network. This in turn influences the choice of protocol. If all subnet masks for an IP network must be the same size, use fixed-length subnet masking (FLSM). If all subnet masks can be different sizes, use variable-length subnet masking (VLSM). In modern networks, use VLSM to conserve IP addresses.

Fixed-Length Subnet Masking

FLSM is required by some routing protocols and has these issues:

- Requires that all subnets of a major network have the same subnet mask
- Results in less efficient address space allocation

Example: Subnetting with FLSM

Network 172.16.0.0/16 is subnetted using FLSM. Each subnet is given a /24 mask. The network is composed of multiple LANs connected by point-to-point WAN links. Because FLSM is used:

- All subnets have the same subnet mask
- The point-to-point links have a /24 subnet mask with 254 addresses available, though only two addresses are needed

Variable-Length Subnet Masking

VLSM:

- Makes it possible to subnet with different subnet masks
- Results in more efficient address space allocation
- Allows greater capability to perform route summarization (allows more hierarchical levels within an addressing plan)

Example: Subnetting with VLSM

Network 172.16.0.0/16 is subnetted using VLSM. The network is composed of multiple LANs connected by point-to-point WAN links. The point-to-point links receive subnet mask /30 with only two available addresses. The network requires less address space, leaving valuable addresses free for other applications.

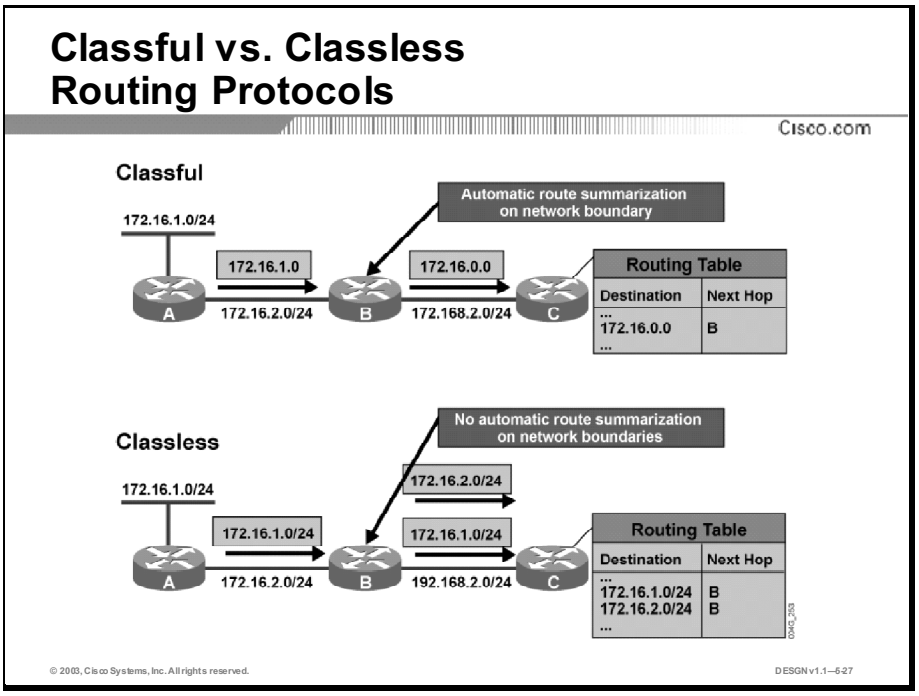

Decisions about IPv4 addressing plans and routing protocols are interdependent.

Classful Routing

When you use classful routing protocols, these rules apply:

- Subnet masks are not included in the routing updates.
- FLSM is required when subnetting is performed.
- When a route update is received and the network is from:
  - The same major network as configured on the interface, the configured subnet mask is applied where the update is received. Therefore, do subnetting with FLSM.
  - A different major network as configured on the interface, the major network mask is applied where the update is received. Therefore, make sure subnetted networks are contiguous.
- All subnets of the same major network (Class A, B, or C) must use the same subnet mask and be contiguous.
- Automatic route summarization is performed across network boundaries.

Typically, in modern networks classful routing is no longer used.

Note: Examples of classful routing protocols are Routing Information Protocol Version 1 (RIPv1) and Interior Gateway Routing Protocol (IGRP).

Classless Routing

When you use classless routing protocols, these rules apply:

- Subnet masks are included in the routing updates
- VLSM is supported
- Route summarization can be manually configured

All currently deployed modern networks should use classless routing.

Note: Examples of classless routing protocols are RIPv2, EIGRP, OSPF, and BGP.

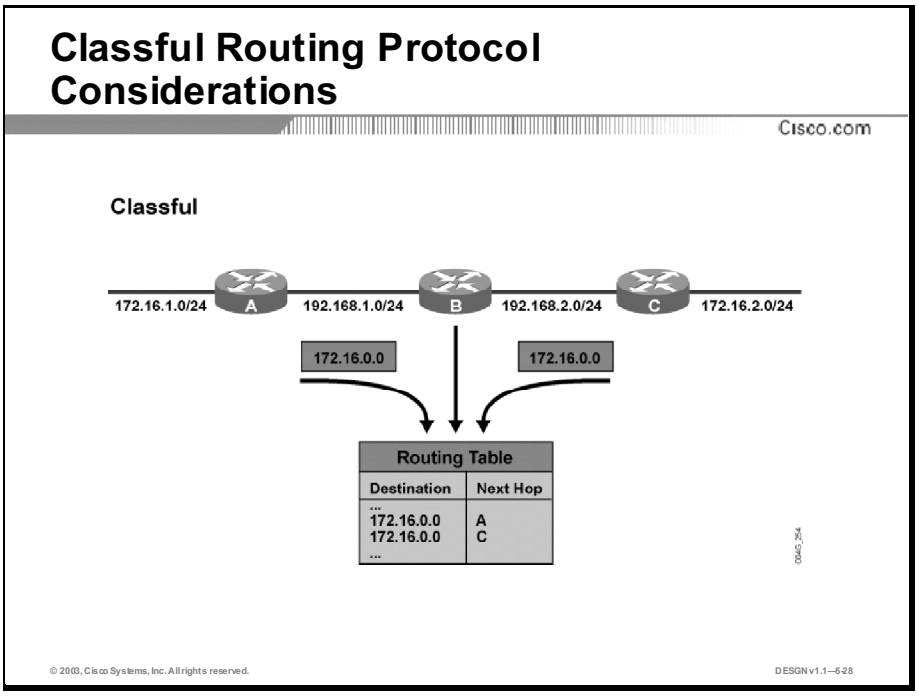

With classful routing, the routing updates do not carry the subnet mask. The figure illustrates that when a route is exchanged across a network boundary, information is not included for subnets 172.161.0/24 and 172.16.2.0/24. The network 172.16.0.0 is added from each interface. Because router B now has two entries for the major network 172.16.0.0, routing problems occur.

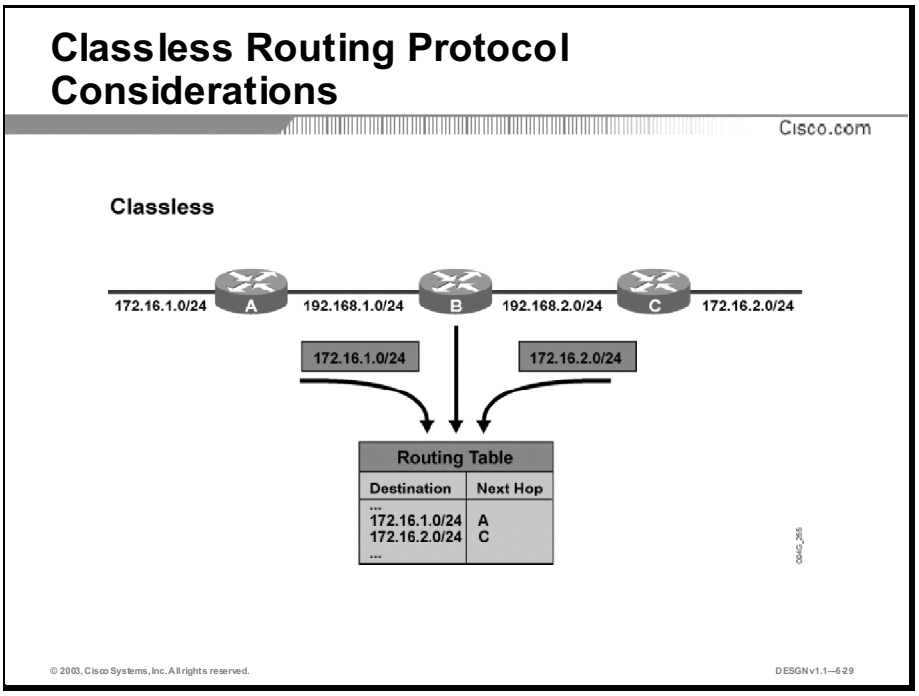

With classless routing, the routing updates carry the subnet mask. The figure illustrates how router B learns both networks 172.16.1.0/24 and 172.16.2.0/24, one from each interface, and the routing is performed correctly.

Assigning End System IP Addresses

The IP address assignment strategy you select can create a large administrative overhead. This topic describes IP address assignment methods and their influence on administrative overhead.

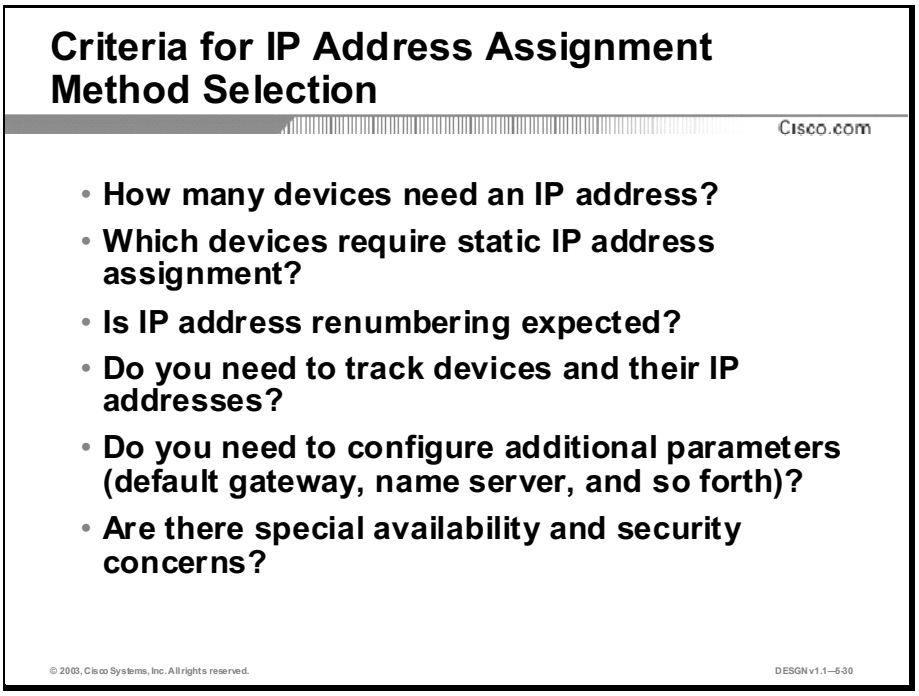

To determine the appropriate IP address assignment method, answer the questions listed in the figure.

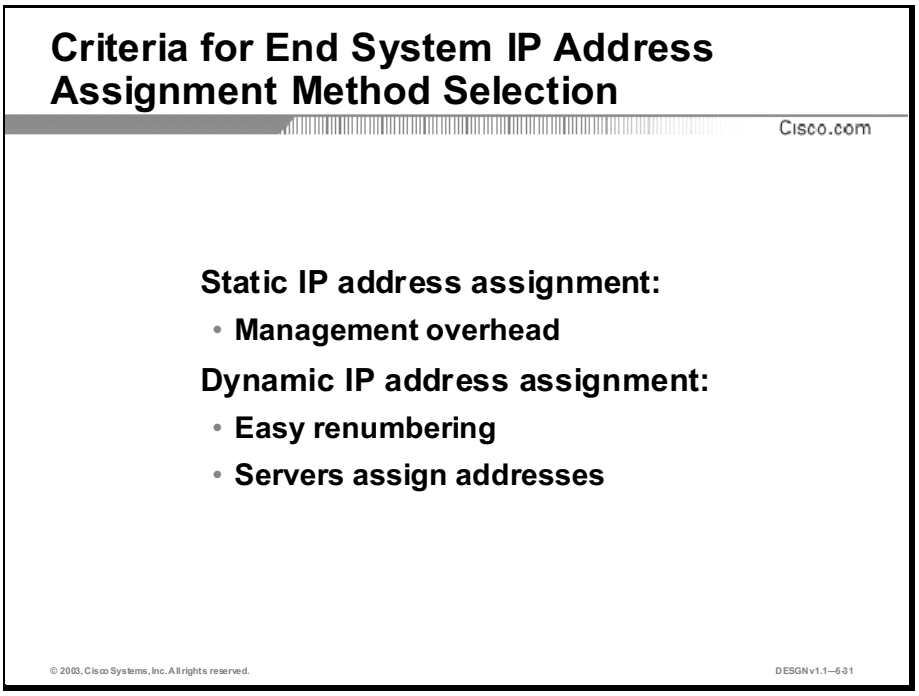

Address assignment includes assignment of an IP address along with the assignment of a default gateway, servers that resolve names to IP addresses, time servers, and so on. These are the two basic IP address assignment strategies:

- **Static:** Defines an IP address that you statically assign to a system. The network administrator configures the IP address, default gateway, and name servers manually using entries in a special file or files on the end system, using either a graphical or text interface. Static address assignment presents an extra burden for the administrator who must configure the address on every end system on the network. This is especially true in large-scale networks.
- **Dynamic:** Assigns addresses dynamically to the end systems. Dynamic address assignment relieves the administrator of manually assigning an address to every device on the network. The administrator must set up a server to assign the address. On that server, the administrator defines the address pools and additional parameters to send to the host (default gateway, name servers, time servers, and so on). On the host, the administrator must enable the host to acquire the address dynamically. When IP address reconfiguration is needed, the administrator reconfigures data on the server, which then renumbers the hosts for which it is responsible. Address assignment is accomplished by address assignment servers, such as Reverse Address Resolution Protocol (RARP) server, Bootstrap Protocol (BOOTP) server, and Dynamic Host Configuration Protocol (DHCP) server.

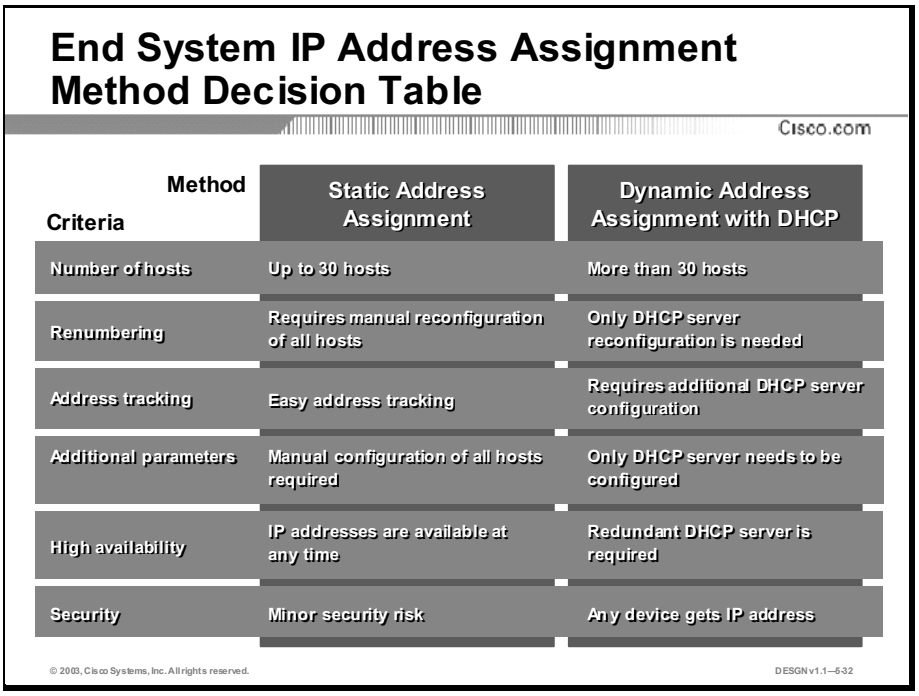

To select a static or dynamic end system IP address assignment method, or a combination of the two, consider these criteria:

- **The number of end systems:** If there are more than 30 end systems, use the dynamic method. Use the static method for 30 or fewer hosts.
- **Renumbering:** If renumbering is very likely to happen and there are many end systems, the dynamic address assignment method is the best choice.
- **Address tracking:** When the network policy requires address tracking, use the static address assignment method. However, address tracking is also possible with dynamic address assignment when you configure the DHCP server.
- **Additional parameters:** A DHCP server provides easy additional parameter configuration. You need only enter the parameters on the server, which sends those parameters and the address to the clients.
- **High availability:** Statically assigned IP addresses are available at any time. Dynamically assigned IP addresses are available only from a functioning server. To ensure reliability, implement a redundant DHCP server.
- **Security:** With the dynamic IP address assignment method, anyone who connects to the network can acquire a valid IP address. This imposes some security risk.

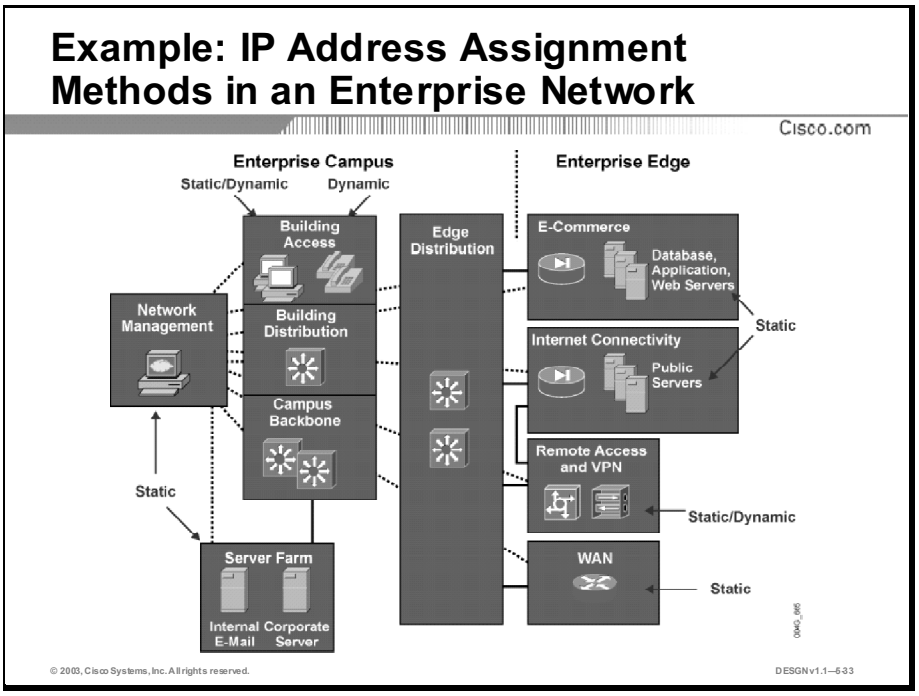

The typical enterprise network uses both static and dynamic address assignment methods. The static IP address assignment method is typically used in the Network Management and Server Farm modules of the Enterprise Campus functional area and in all modules of the Enterprise Edge functional area (the E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules). Static addresses are needed for systems such as servers or network devices, where the IP address needs to be known at all times for general access or management.

The dynamic IP address assignment method is used to assign IP addresses to workstations and IP Phones.

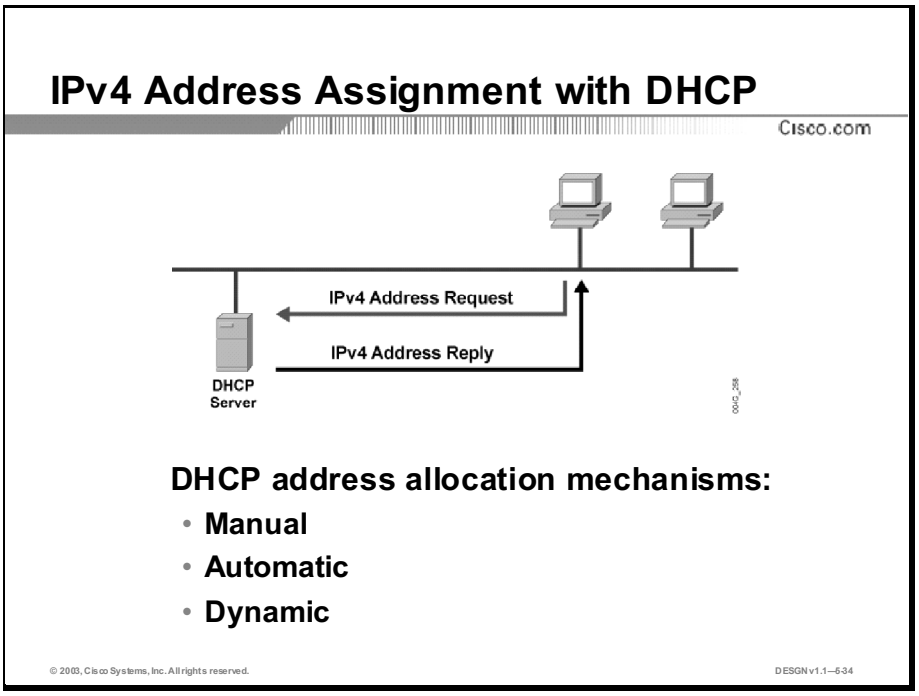

Dynamic address assignment servers rely on special software and may provide a name for address resolution. When using dynamic address assignment, you may use a DHCP server. The DHCP server allows you to send additional parameters (default gateway, DNS server, and so on), in addition to the IPv4 address and the subnet mask.

When a host is started, its address-acquire process sends an IPv4 address request by sending its physical hardware address to the network. The DHCP server intercepts the request and responds with the host IPv4 address, subnet mask, and additional IPv4 parameters.

DHCP Address Allocation Mechanisms

DHCP supports three possible address allocation mechanisms:

- **Manual:** The network administrator assigns the IPv4 address to a specific MAC address. DHCP is used to dispatch the assigned address to the host.
- **Automatic:** The IPv4 address is permanently assigned to a host.
- **Dynamic:** The IPv4 address is assigned to a host for a limited time or until the host explicitly releases the address. This mechanism supports automatic address reuse when the host to which it has been assigned no longer needs the address.

Implementing Name Resolution

Names are used to identify different hosts and resources on the network and to provide user-friendly interaction with computers. A name is much easier to remember than an IP address. Name resolution maps a name to an IP address. This topic discusses implementing name resolution.

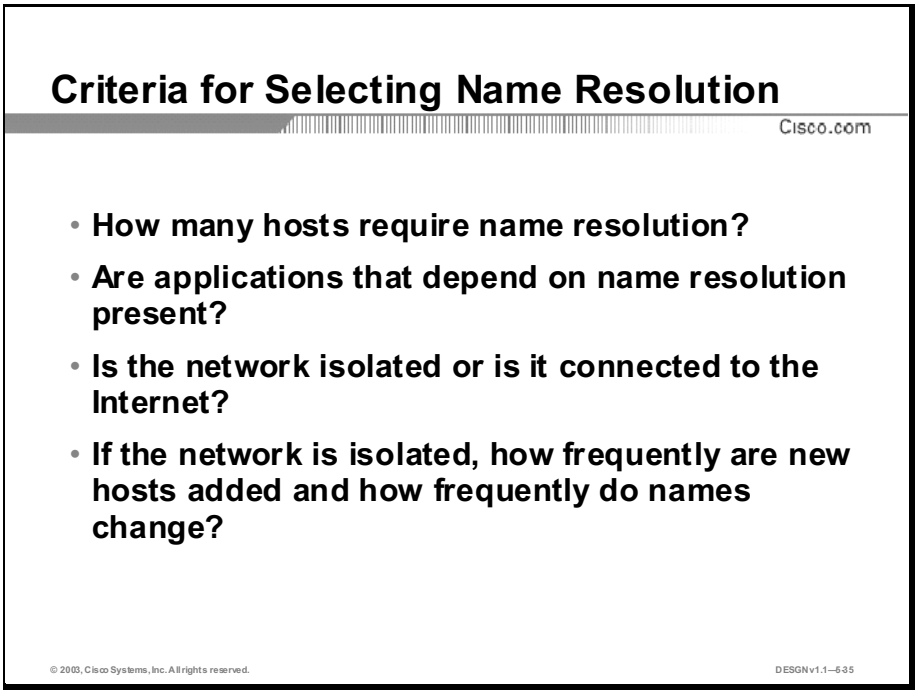

Network hosts identify themselves to each other by using naming schemes. Each computer on the network can have an assigned name to provide easier communication. Because the IP network layer protocol uses IP addresses to transport datagrams, a name used to identify a host must be mapped or "resolved" into an IP address. This is called name resolution. To select the appropriate name resolution method, answer the questions listed in the figure.

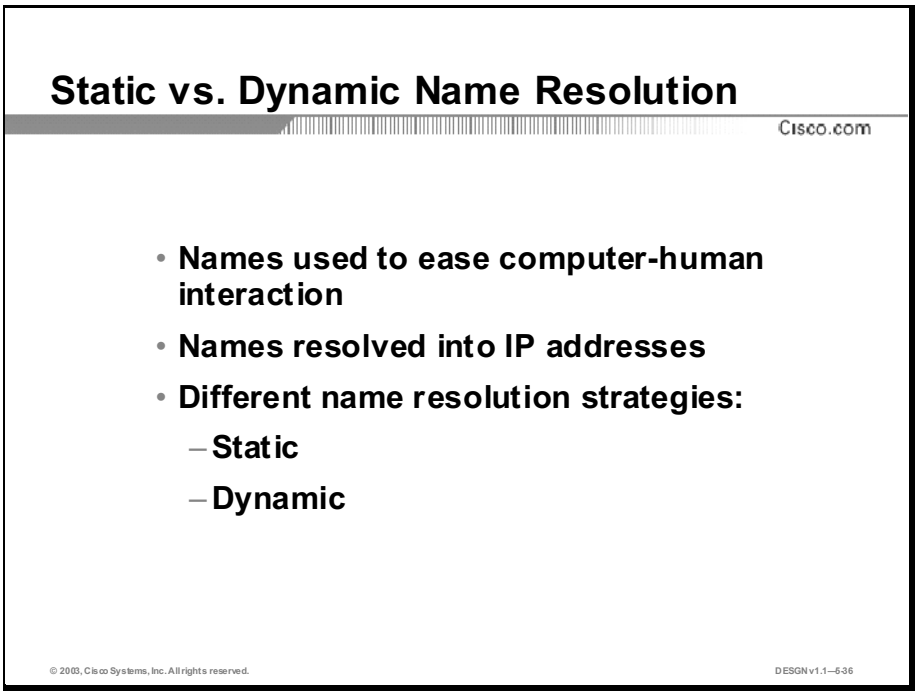

The process of resolving a host name to an IP address is either static or dynamic. These are the two name resolution methods:

- **Static:** With static name-to-IP-address resolution, the administrative overhead and configuration are similar to a static address assignment strategy. The network administrator manually defines name-to-IP-address resolutions in a special file by entering the name and IP address pairs, using either a graphical or text interface. Manual entries create additional work for the administrator, who must enter them on every host and are prone to errors and omissions.
- **Dynamic:** The dynamic name-to-IP-address resolution is very similar to the dynamic address assignment strategy. The administrator needs only to enter the name-to-IP-address resolutions on a special server and is relieved of repeating the task on every host. The server then performs the job of name-to-IP-address resolution. The dynamic name-to-IP-address resolution method facilitates renumbering and renaming.

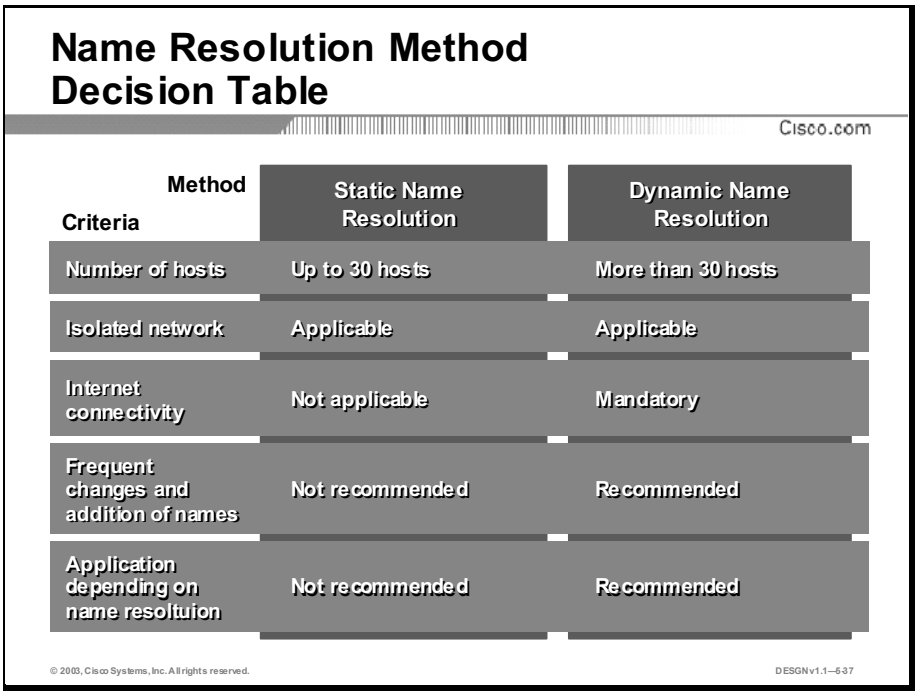

The selection of either static or dynamic end system IP address assignment method depends on:

- **The number of hosts:** If there are more than 30 end systems, use dynamic name resolution. Use static name resolution for 30 or fewer hosts.
- **Isolated network:** If the network is isolated (no connections to the Internet) and the number of hosts is small, use static name resolution. Alternatively, you can use the dynamic method.
- **Internet connectivity:** When Internet connectivity is present, the static name resolution is not an option and DNS name resolution is mandatory.
- **Frequent changes and addition of names:** When dealing with frequent changes and additions in an isolated network, the recommended name resolution is DNS.
- **Applications depending on name resolution:** If applications that depend on name resolution are used, the recommended name resolution is DNS.

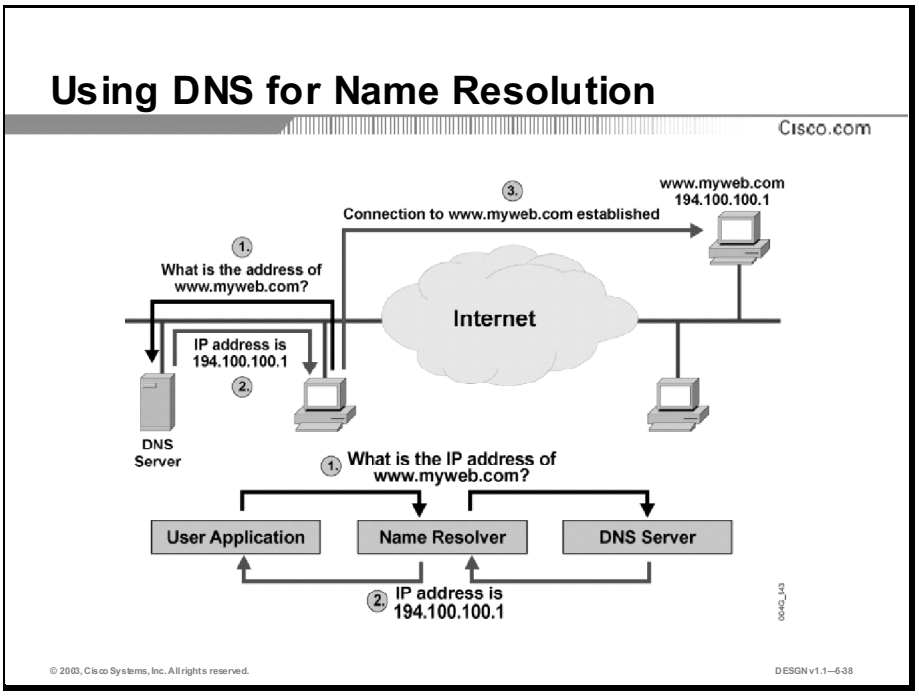

When numerous hosts or names must be resolved to IP addresses, statically defined resolutions in the HOSTS file are unwieldy to maintain. Use a DNS server for name resolution when you have a large number of hosts or fully qualified domain names (FQDNs) to resolve to IP addresses.

A DNS server is special software that usually resides on a dedicated host. DNS servers are organized in a hierarchical structure. A DNS server can also query other DNS servers to retrieve partial resolutions for a certain name; for example, one DNS server could resolve "myweb.com" and another could resolve "www."

These are example steps to resolve an IP address using a DNS server:

Step 1 A user wants to browse "www.myweb.com**."** Because the host does not know the IP address of that site, it queries the DNS server.

Step 2 A server responds with the appropriate IP address for "www.myweb.com**."**

Step 3 The host establishes a connection to the appropriate IP address (site).

Note: A FQDN is a complete domain name for a specific host on the Internet with the information needed to convert it into a specific IP address. The FQDN consists of a host name and a domain name. For example, "www.myweb.com" is the FQDN on the Internet for "myweb" web server. The host is "www," the domain is "myweb," and the top-level domain name is "com."

DNS Name Resolution

To resolve symbolic names to actual network addresses, hosts need to know the DNS server address. Queries on the name server are performed through resolver or name resolver programs, which are usually part of the host operating system. An application sends a query to a name resolver, which resolves the request with either the local database (HOSTS file) or the DNS server.

To enable DNS name resolution, the network administrator must set up the DNS server, enter information about host names and corresponding IP addresses, and configure the hosts to use the DNS server for name resolution.

Example: Communication Between Hosts

Host A is aware of DNS server B's presence. When host A wants to communicate with host X:

Step 1 Host A sends a request to DNS server B.

Step 2 DNS server B looks for the IP address in its database and replies back to host A with the host X IP address.

Step 3 Host A uses this information to establish communication with the host X.

Summary

This topic summarizes the key points discussed in this lesson.

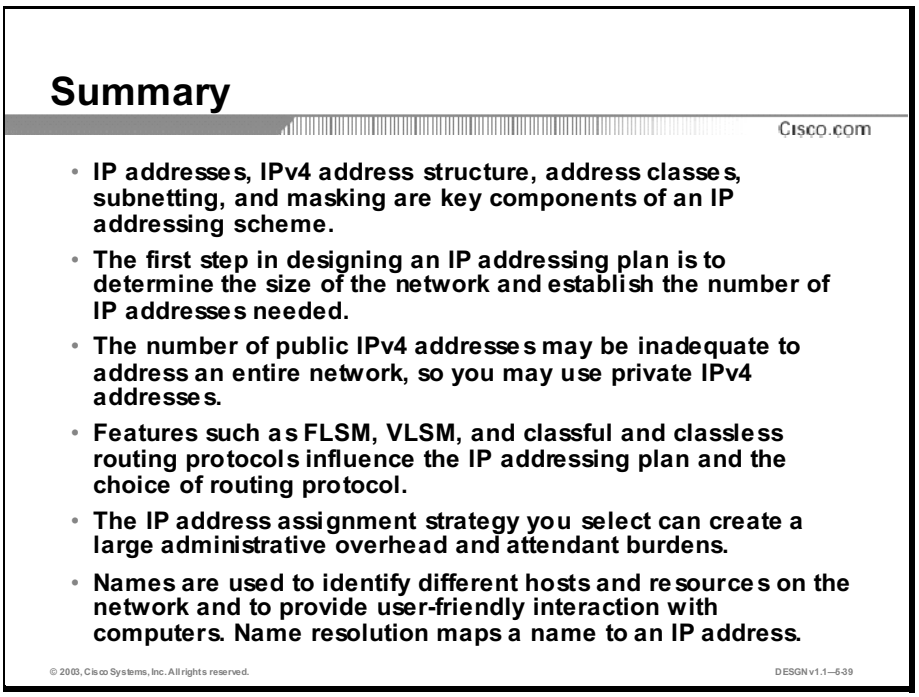

References

For additional information, refer to these resources:

- Comer, Douglas E., and D. L. Stevens. *Internetworking with TCP/IP*. Englewood Cliffs, New Jersey: Prentice-Hall; 1991.
- *Designing Large-Scale IP Internetworks,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm.
- *Subnetting an IP Address Space,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd20a.htm.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) How many hosts can you address with the IPv4 network and mask 168.222.240.240/28?

A) 14

B) 16

C) 1024

D) 65534

Q2) Which information must you collect to determine the overall network size?

A) number of users

B) number of departments

C) number of workstations, servers, IP Phones, router interfaces, switch management interfaces, and firewall interfaces

D) distances between central location and remote locations

Q3) Which two IPv4 addresses can you use in public networks? (Choose two.)

A) 192.167.20.1/24

B) 192.168.1.200/28

C) 172.30.100.33/24

D) 172.32.1.1/16

Q4) What is the appropriate method to reduce routing overhead?

A) random IP address assignment

B) choosing routers with more computing power

C) route summarization

Q5) _____-length subnet masking is allowed with classless routing protocols.

A) Variable

B) Fixed

C) Suffix

D) Prefix

Q6) What are the two disadvantages of using the Dynamic Host Configuration Protocol (DHCP) for address assignment? (Choose two.)

A) a need for a special server

B) additional traffic on the network

C) subnet mask must be configured on every device

D) additional parameters cannot be configured by DHCP address assignment

Q7) What name resolution method reduces administrative overhead?

A) dynamic name resolution

B) static name resolution

C) transparent name resolution

D) ISP name resolution

Quiz Answer Key

Q1) A

Relates to: IPv4 Address Structure

Q2) C

Relates to: Determining the Size of the Network

Q3) A, D

Relates to: Private vs. Public Addresses

Q4) C

Relates to: Implementing Hierarchy with IP Addressing

Q5) A

Relates to: Implementing Hierarchy with IP Addressing

Q6) A, B

Relates to: Assigning End System IP Addresses

Q7) A

Relates to: Implementing Name Resolution

Introducing IPv6

Overview

IPv6 was introduced to address deficiencies in IPv4. IPv6 offers a new address structure and address types, new name resolution strategies, and affects routing protocols. Because the migration from IPv4 to IPv6 does not happen automatically, you must plan a transition strategy. This lesson describes the IPv6 features and transition strategies.

Relevance

This will help you identify IPv6 addressing-related issues and coexistence issues.

Objectives

Upon completing this lesson, you will be able to explain IPv6-specific design considerations. This includes being able to meet these objectives:

- Describe the IPv6 address structure
- List different IPv6 address types
- Identify IPv6 routing protocol considerations
- Describe IPv6 address assignment strategies
- Describe the name resolution process in IPv6
- Describe strategies for transitioning from IPv4 to IPv6

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions

Outline

The outline lists the topics included in this lesson.

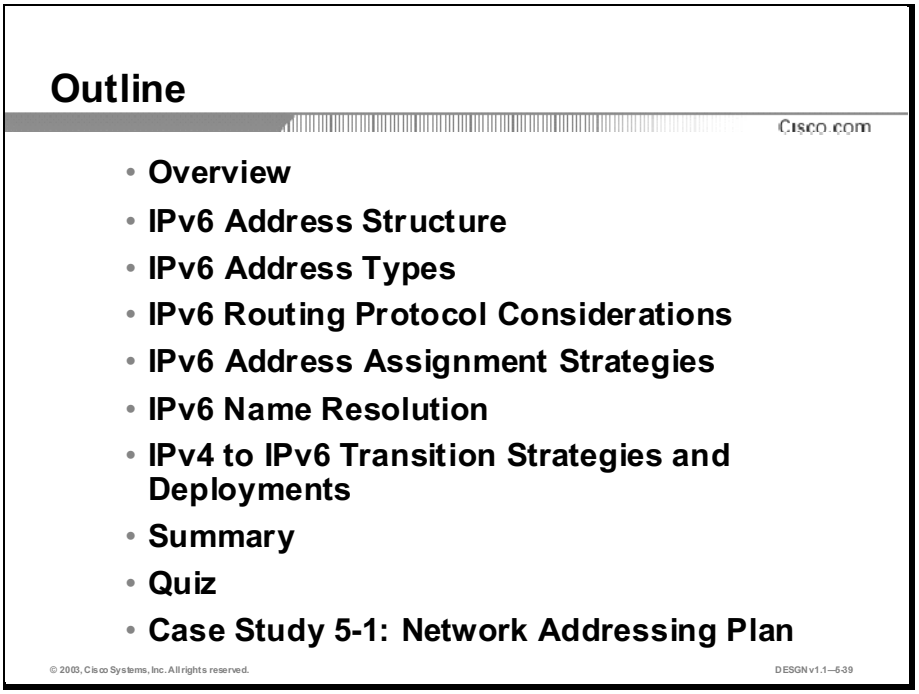

IPv6 Address Structure

IPv6 was designed as a successor to IPv4 to overcome IPv4 limitations. This topic discusses the IPv6 address structure and features.

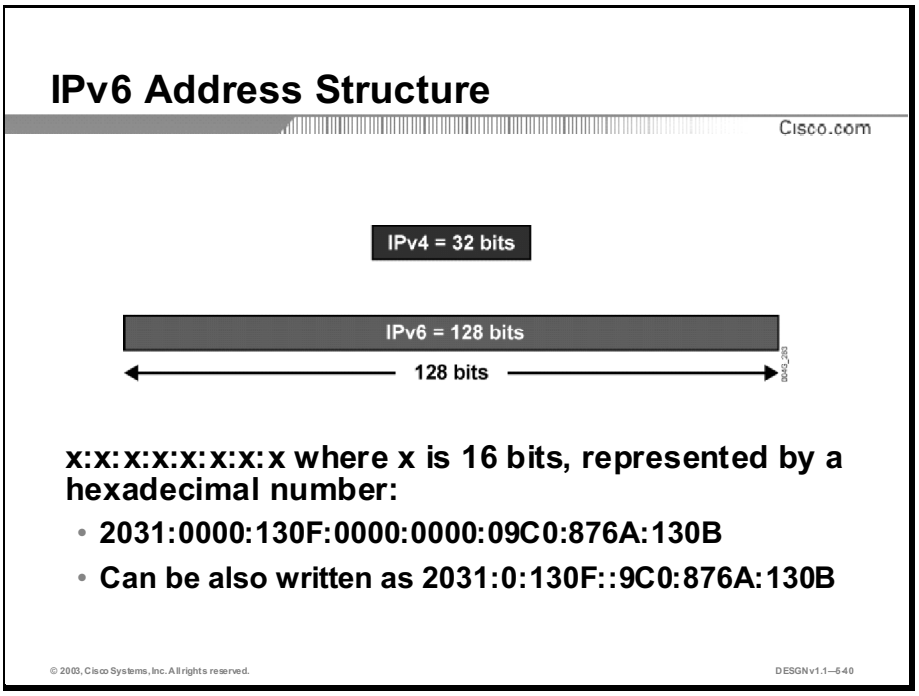

An IPv6 address is 128 bits long, a much larger address space than the address space in IPv4. It can provide approximately 3.40*10^38 addresses.

IPv6 addresses are represented as a series of 16-bit fields presented as a hexadecimal number and separated by colons (:), in the format: x:x:x:x:x:x:x:x. To shorten the writing of IPv6 addresses, a few techniques are available. Here are some of them:

- The leading zeros in a field are optional.
- IPv6 addresses often contain successive hexadecimal fields of zeros. To shorten IPv6 addresses, you can use two colons (::) to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). You can only shorten an IPv6 address one time. Otherwise you would not know how many zeros to add.

A single interface can have multiple IPv6 addresses of different types. IPv6 addresses contain a scope field that categorizes the types of applications that are suitable for the address.

Example: IPv6 Address

IPv6 address 2031:**0000**:130F:**0000:0000:0**9C0:876A:130B can be written as

RIGHT: 2031:**0**:130F::9C0:876A:130B

Benefits of IPv6

These are the main benefits of IPv6:

- **Larger address space:** IPv6 increases the IP address size from 32 bits to 128 bits. This increase supports more addressing hierarchy levels, a much greater number of addressable nodes, and simpler address autoconfiguration.
- **Global unique IP addresses:** Every node can have a unique global IP address, which eliminates the need for NAT.
- **Site multihoming:** IPv6 allows hosts to have multiple IPv6 addresses and networks to have multiple IPv6 prefixes. This facilitates connection to multiple ISPs without breaking the global routing table.
- **Header format efficiency:** A fixed header size makes processing more efficient.
- **Improved privacy and security:** IPv6 introduces optional security headers.
- **Flow labeling capability:** A new capability enables packet labeling to belong to particular traffic “flows” so the sender can request special handling, such as nondefault quality of service (QoS) or “real-time” service.
- Increased mobility and multicast capabilities

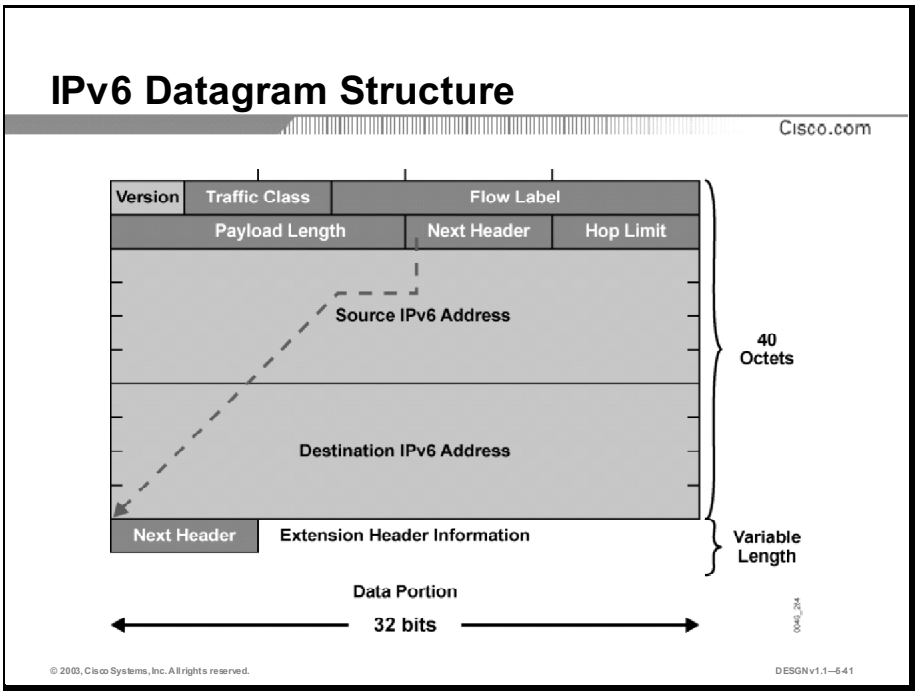

To provide more efficient processing, IPv6 introduces a new packet header structure with a fixed header size. The IPv6 header has these eight fields:

- **Version:** A 4-bit field that indicates the IP version, in this case, IPv6.
- **Traffic Class:** An 8-bit field that tags packets with a traffic class that is used in differentiated services.
- **Flow Label:** A 20-bit field that a source uses to label sequences of packets for which the source requests special handling by the IPv6 routers (for example, nondefault QoS or "real-time" service). If a host or router does not support the functions of the Flow Label field, the field is set to 0 for the packets that the device originated, left unchanged for the packets that are forwarded, and ignored for the received packets.

Note: The Flow Label field is still experimental and subject to change as the requirements for flow support in the Internet clarifies.

- **Payload Length:** A 16-bit field similar to the Total Length field in the IPv4 packet header. The Payload Length field indicates the total length of the data portion of the packet.
- **Next Header:** An 8-bit field similar to the Protocol field in the IPv4 packet header. The value of the Next Header field determines the type of information, following the basic IPv6 header (for example, TCP/UDP packet).
- **Hop Limit:** This 8-bit field specifies the maximum number of hops an IP packet can traverse and is similar to the Time to Live (TTL) field in the IPv4 packet header.
- **Source Address:** This 128-bit (16 octets) field contains the source address of the packet.
- **Destination Address:** This 128-bit (16 octets) field contains the destination address.

The header has a total length of 40 octets (320 bits). The six header fields other than the addresses are aligned to 64 bits, which enables more efficient processing. Extension headers, if any, follow the header fields, followed by the data portion of the packet.

IPv6 Address Types

Similar to IPv4, in IPv6 a single source can address datagrams to either one or many destinations at the same time. IPv6 supports three address types: link-local, site-local, and global aggregatable. This topic explains the IPv6 address types.

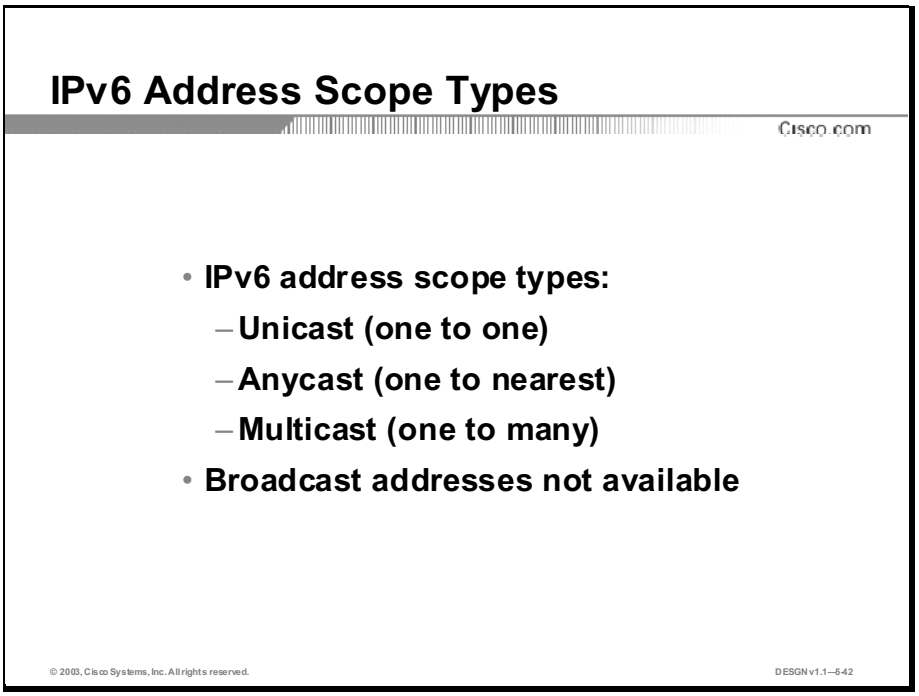

Here are some IPv6 address scope types:

- **Unicast (one-to-one):** The process is the same as IPv4. A single source sends data to a single destination. A packet sent to a unicast IPv6 address is delivered to the interface identified by that address. Different IPv6 unicast addresses are supported:
  - Link-local address
  - Site-local address
  - Global aggregatable address
  - IPv4-compatible IPv6 address
- **Anycast (one-to-nearest):** An identifier for a set of interfaces that typically belong to different nodes. A packet sent to an anycast address is delivered to the closest interface, as defined by the routing protocols in use, identified by the anycast address. In other words, receivers that share the same characteristics are assigned the same anycast address. A sender interested in contacting a receiver with those characteristics sends packets to the anycast address, and the routers deliver the packet to the receiver that is nearest to the sender.

  You can use anycast addresses for a service location. For example, an anycast address could be assigned to a set of replicated FTP servers. A user in China who wanted to retrieve a file would be directed to the Chinese server and a user in Europe would be directed to the European server. Anycast addresses are allocated from the unicast address space. You cannot use anycast addresses as the source address of an IPv6 packet. You must explicitly configure nodes to which the anycast address is assigned to recognize the anycast address.

- **Multicast (one-to-many):** The same as IPv4, an address for a set of interfaces (in a given scope) that typically belong to different nodes. A packet sent to a multicast address is delivered to all the interfaces that the address identifies.

Note: IPv6 has no concept of broadcast addresses. Multicast addresses are used instead.

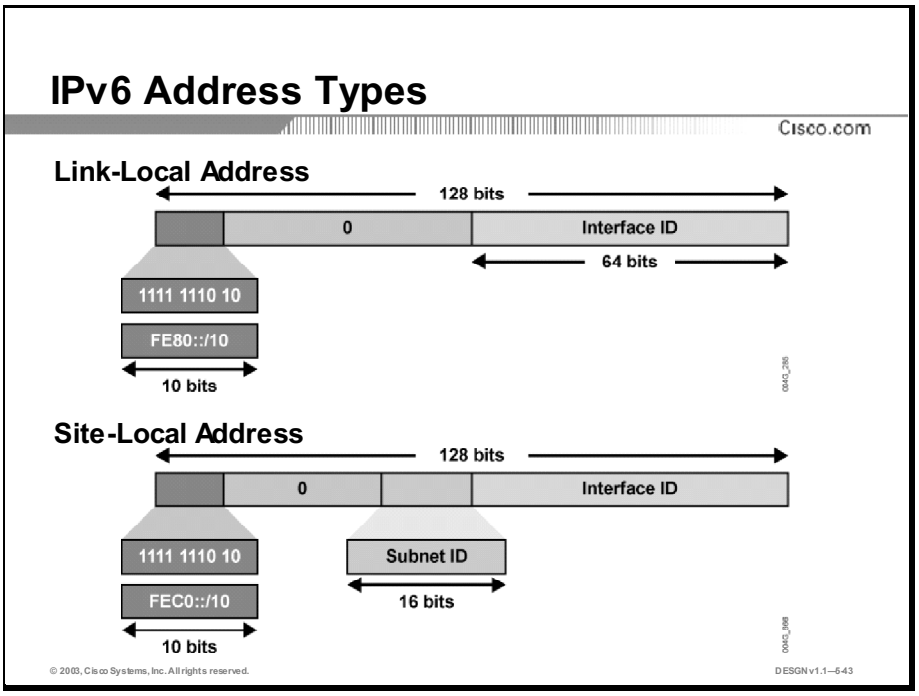

Link-Local Address

A link-local address is useful in the context of the local link network; its scope on the link is limited. A link-local address is an IPv6 unicast address that you can automatically configure on any interface by using the link-local prefix FE80::/10 (1111 111010) and the interface identifier. Link-local addresses are used in the neighbor discovery protocol and the stateless autoconfiguration process. Use link-local addresses to connect devices on the same local network without using either site-local or globally unique addresses. Many routing protocols use link-local addresses. An IPv6 router must not forward packets that have either link-local source or destination addresses to other links.

Site-Local Address

Use site-local addresses to number a complete site without a global prefix. Site-local addresses are IPv6 unicast addresses that use the prefix FEC0::/10 (1111 111011) and concatenate the subnet identifier (the 16-bit field) with the interface identifier. Think of site-local addresses as private addresses, and use them to restrict communication to a limited domain. For example, use them to assign numbers to a device such as a printer that will never communicate with the IPv6 Internet.

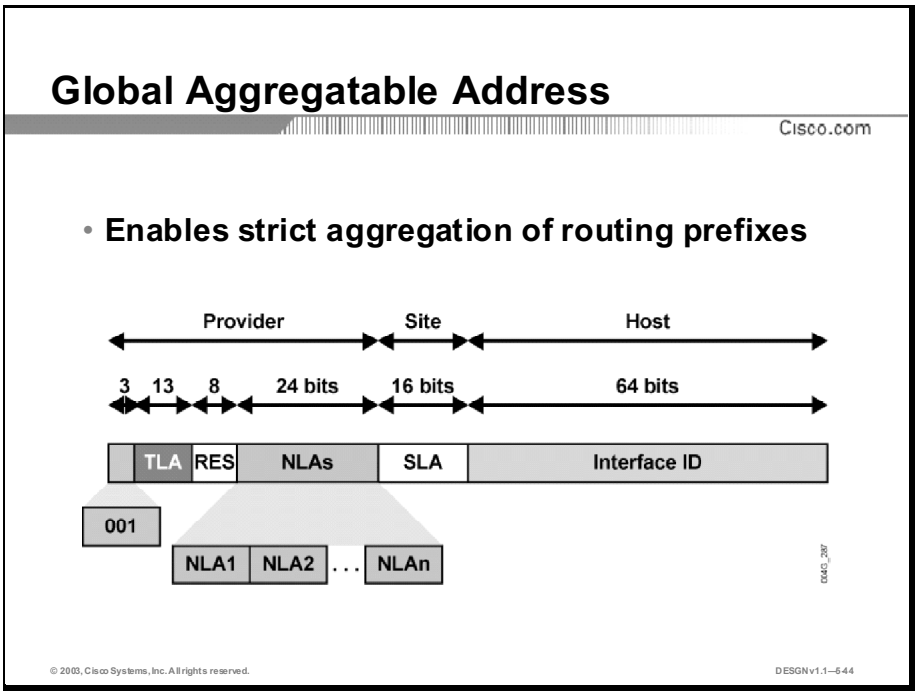

An aggregatable global address is an IPv6 address from the aggregatable global unicast prefix. Aggregatable global unicast addresses enable strict aggregation of routing prefixes that limits the number of routing table entries in the global routing table. Aggregatable global addresses used on links are aggregated upward through organizations, then to intermediate-level ISPs, and eventually to top-level ISPs. This is the structure:

- A fixed prefix of 2000::/3 (001) indicates an aggregatable global IPv6 address.
- At the top of the hierarchy, several international registries assign blocks of addresses to Top-Level Aggregators (TLA) of 13 bits. TLAs are essentially the public transit points (exchanges) where long-haul providers establish peer connections.
- A reserved field of 8 bits for the growth of the TLA and Next-Level Aggregator (NLA) fields. The field must always be equal to zero.
- TLAs allocate blocks of addresses to NLAs of 24 bits, which represent large providers and global corporate networks.
- When an NLA is a provider, it further allocates its addresses to its subscribers using a 16-bit Site-Level Aggregator (SLA). If the NLA is a global corporate network, you can use the SLA to identify individual sites or subnetworks.
- Individual organizations use a 16-bit field for the host, with the interface ID used to identify interfaces on a link. The 16-bit field must be unique to the link.

IPv6 Routing Protocol Considerations

The routing protocols available in IPv6 include Interior Gateway Protocols (IGPs) and exterior gateway protocols (EGPs). This topic discusses IPv6 routing protocol considerations.

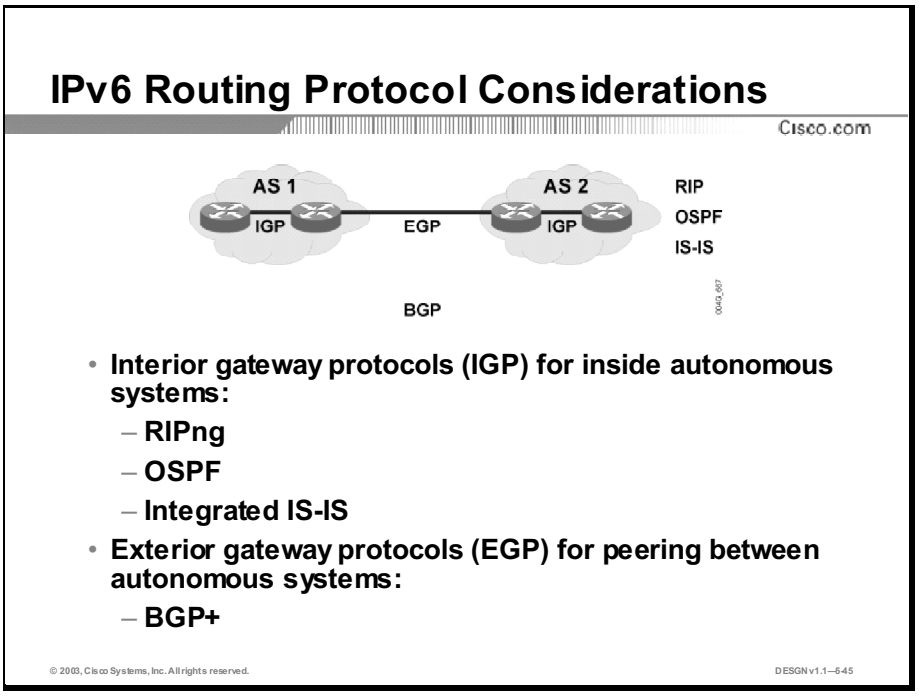

As does IPv4 classless interdomain routing (CIDR), IPv6 uses the same "longest-prefix match" routing. Recent routing protocol versions handle longer IPv6 addresses and different header structures. Currently, these updated routing protocols are available:

- **RIPng (RFC 2080):** Referred to as the RIP new generation. RIPng is a distance-vector protocol with a limit of 15 hops, which uses split-horizon and poison reverse to prevent routing loops. IPv6 update features include:
  - Based on IPv4 RIPv2 and similar to RIPv2
  - IPv6 prefix, next-hop IPv6 address
  - Multicast group FF02::9, the all-rip-routers multicast group, as the destination address for RIP updates
  - IPv6 for transport
- **OSPFv3 (RFC 2740):** The protocol implementation for IPv6 includes:
  - Similar to IPv4
  - Carries IPv6 addresses
  - Uses link-local addresses as source
  - IPv6 for transport

- **Integrated Intermediate System-to-Intermediate System (IS-IS):** Large address support facilitates the IPv6 address family:
  - Same as IPv4 with some extensions added:
    - Two new types, lengths, values (TLVs):
      - IPv6 reachability
      - IPv6 interface address
    - New protocol identifier
- **BGP4+ (Multiprotocol Extensions to BGP, RFC 2283), RFC 2545:**
  - Multiprotocol extensions for BGP4 enable other protocols besides IPv4
  - New identifier for the IPv6 address family
  - IPv6 specific extensions

IPv6 Address Assignment Strategies

As with IPv4, IPv6 presents two major address assignment strategies: static and dynamic (also called autoconfiguration). This topic discusses IPv6 address assignment strategies.

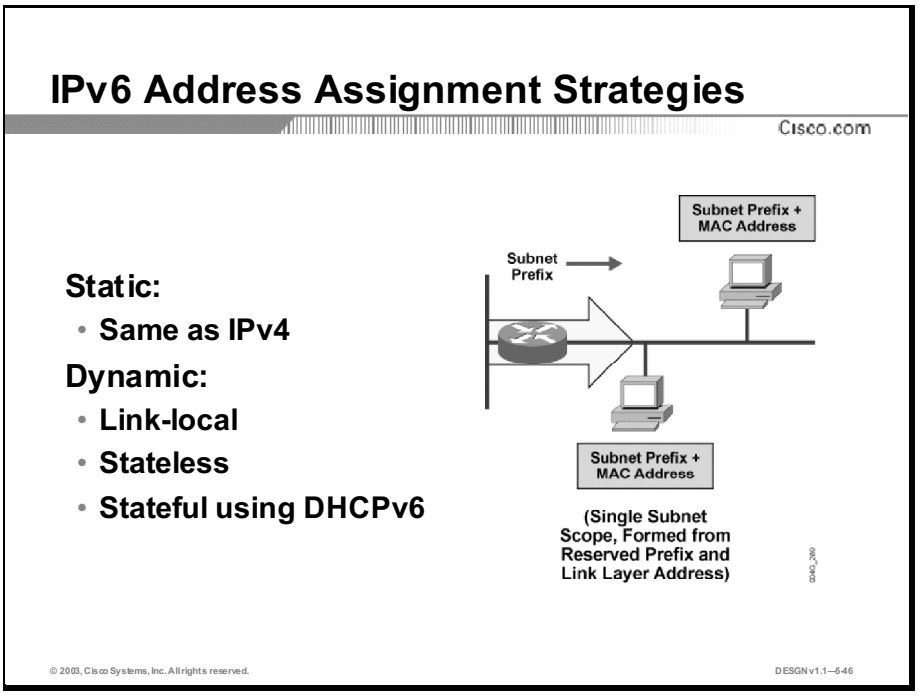

Static address assignment in IPv6 is the same as in IPv4. The administrator must manually configure the IPv6 address on every device in the network. IPv6 dynamic address assignment strategies include:

- **Link-local address:** The host configures its own link-local address autonomously, using identifiers for an interface and the link-local prefix FE80::0.
- **Stateless:** A router on the link advertises site-local and global prefixes and willingness to function as a default router for the link, either periodically or upon the host's request. Hosts can automatically generate site-local and global IPv6 addresses without the need for either manual configuration or the help of a server, such as a DHCP server, using these router messages.
- **Stateful using DHCPv6:** DHCPv6 is an updated version of DHCP for IPv4, which supports new addressing. It enables more control than stateless autoconfiguration and supports renumbering without routers. You can use DHCPv6 for automatic domain name registration using a dynamic DNS server. DHCPv6 uses multicast addresses.

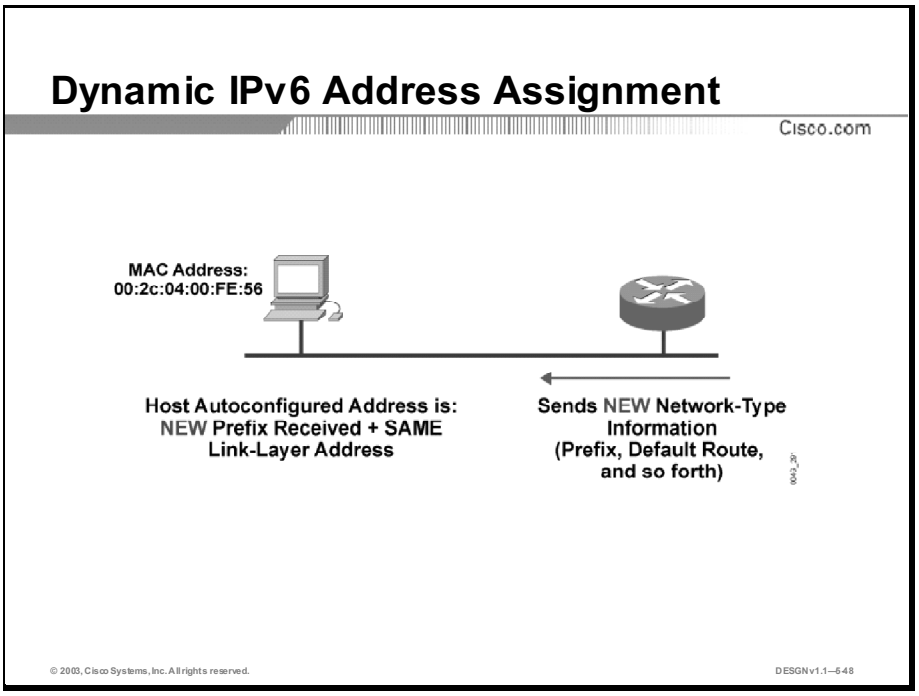

IPv6 eliminates the need to configure addresses by enabling automatic address configuration without any servers, such as DHCP servers. The process that the IPv6 autoconfiguration feature uses is as follows:

Step 1 The new prefix is added to the router advertisement messages sent on the link, so they contain both the old and new prefixes.

Step 2 The nodes use both the addresses created from the new prefix and from the existing address created from the old prefix. The old prefix has lifetime parameters configured to determine the transition period after which the old prefix is removed from the router advertisement messages.

Step 3 When the old prefix is removed from the router advertisements, only addresses that contain the new prefix are used on that link.

IPv6 Name Resolution

IPv6 and IPv4 static and dynamic name resolutions are very similar. This topic describes IPv6 name resolution.

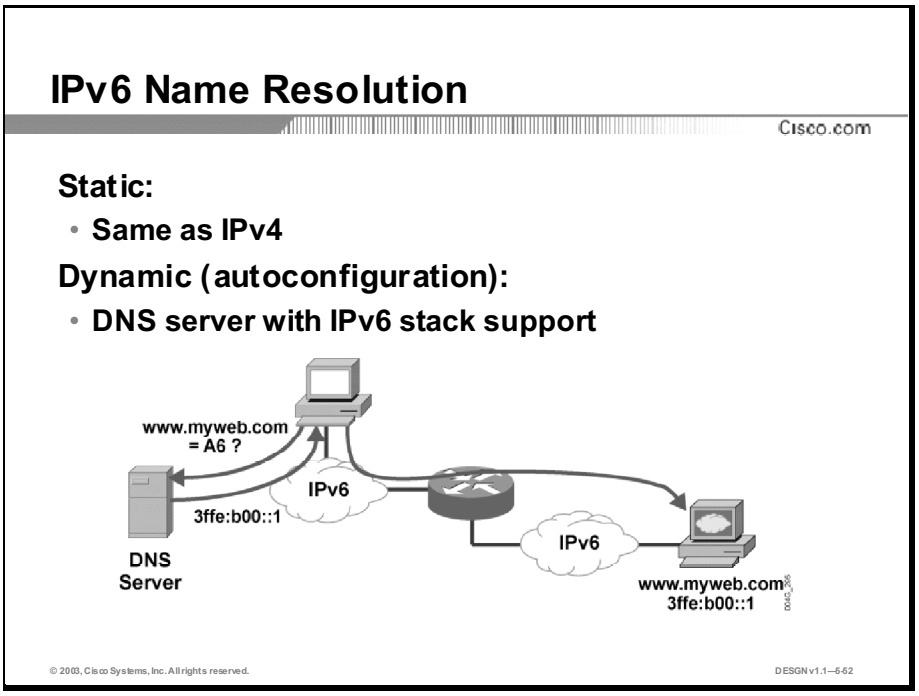

There are two name resolutions available with IPv6:

- **Static name resolution:** Relies on manual entries in the host's local configuration files.
- **Dynamic name resolution:** Uses a DNS server, which has built-in support for IPv6, usually with IPv4. An IPv6-aware application requests the IPv6 address of the destination host name (www.myweb.com) from the DNS with a request for an A6 record (an address record for the IPv6 host, a new DNS feature). The name resolver, usually part of the operating system, queries for the address. The network administrator must set up the appropriate DNS server with IPv6 support and must connect the DNS server to the IPv6 network with a valid IPv6 address. On the host side, the administrator must either enter valid DNS addresses manually or use DHCPv6.

IPv4 and IPv6 Aware Applications and Name Resolution

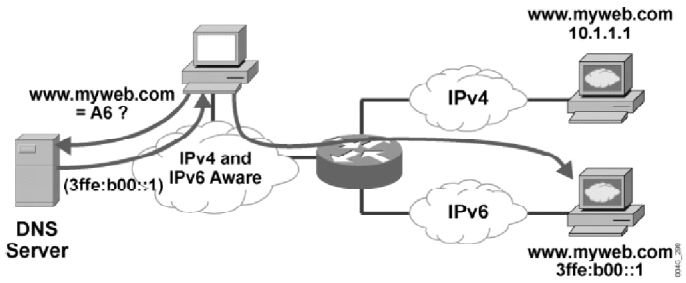

- **In a dual-stack case, an application is IPv4 and IPv6 enabled.**
- **The application decides which stack to use and asks DNS for the address.**

A dual-stack host is aware of both IPv4 and IPv6 protocol stacks, with a new application program interface (API) that supports both IPv4 and IPv6 addresses and DNS requests. A converted application can use both IPv4 and IPv6.

An IPv6- and IPv4-enabled application chooses which stack to use (the typical default is IPv6) and asks the DNS server the destination host name address (www.myweb.com). After receiving the response from the DNS server, the application then requests that the source host connect to the destination host using IPv6.

IPv4 to IPv6 Transition Strategies and Deployments

IPv4 to IPv6 migration requires careful planning. This topic discusses strategies for transitioning from IPv4 to IPv6.

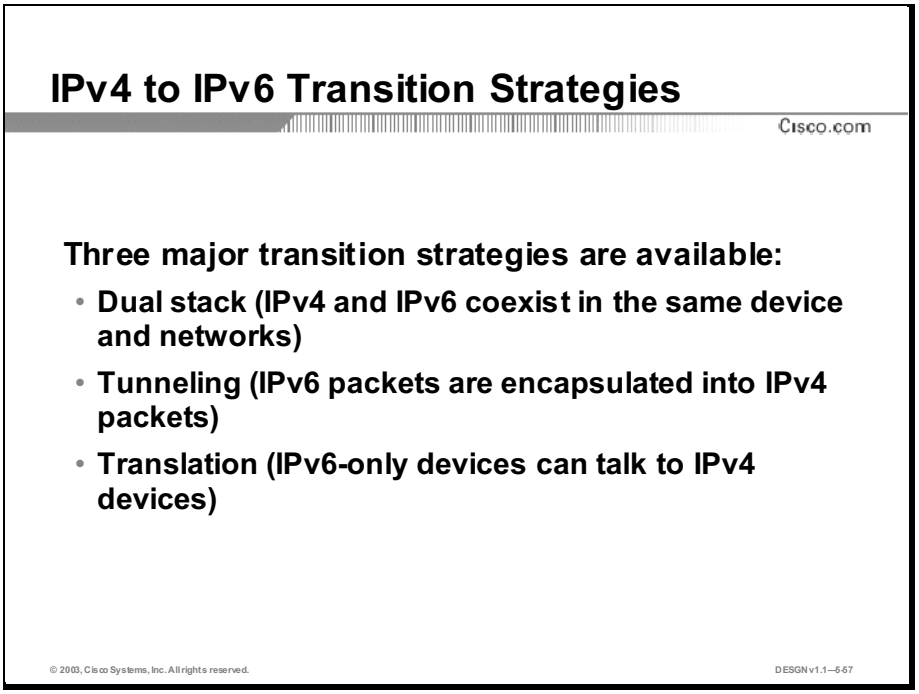

Transitioning from IPv4 to IPv6 can take several years due to the high cost of upgrading the equipment. During the transition, both IPv4 and IPv6 must coexist. To enable smooth and end-user-transparent communication between the IPv4 and IPv6 parts of a network, different solutions are available to the network administrator. These are three primary mechanisms that help with the transition from IPv4 to IPv6:

- **Dual stack:** Both the IPv4 and the IPv6 stacks run on a system. This system can communicate with both IPv6 and IPv4 devices.
- **Tunneling:** IPv6 packets are encapsulated to traverse IPv4 networks and vice versa.
- **Translation:** This mechanism translates one protocol to the other to facilitate communication between the two networks.

In addition, Cisco Systems designed the IPv6 Provider Edge router over Multiprotocol Label Switching (MPLS) feature (6PE), which supports smooth integration of IPv6 into MPLS networks. Because MPLS router switch packets are based on labels rather than address look-ups, customers with an MPLS backbone can scale IPv6 traffic easily and do not need to make costly hardware upgrades.

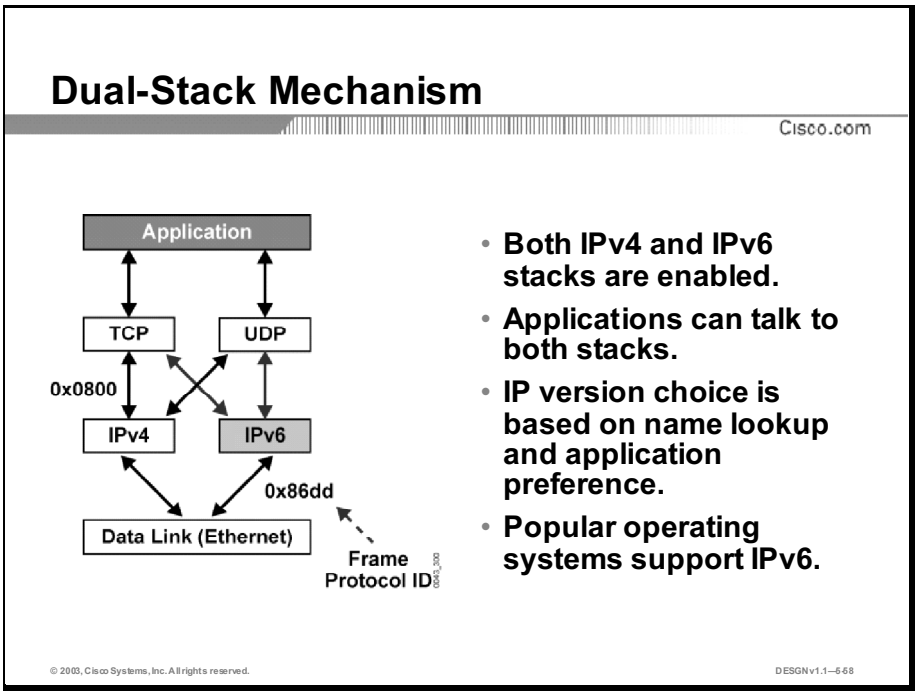

The dual-stack node supports both IPv4 and IPv6 stacks. Applications can communicate with both IPv4 and IPv6 stacks, and the IP version choice is based on name lookup and application preference. This is the most appropriate for the campus and access networks during the transition period and is the preferred technique for transition to IPv6. Operating systems use a dual-stack approach to support the maximum number of applications.

Among the operating systems that support IPv6 stack are FreeBSD, Linux, Sun Solaris, and Windows 2000/XP.

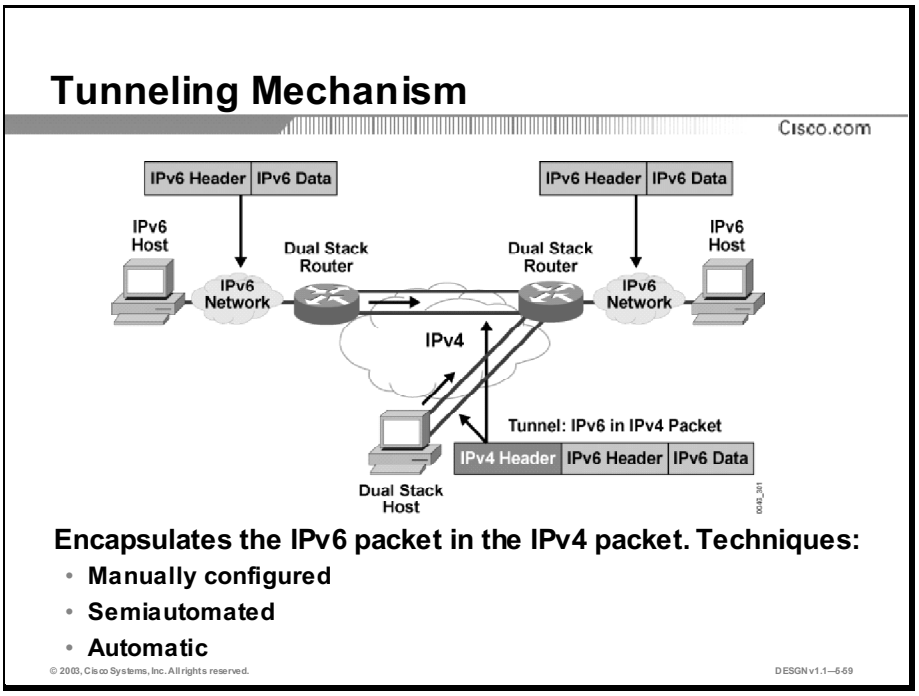

Tunneling encapsulates packets of one type into packets of another type. In the case of a transition to IPv6, tunneling encapsulates IPv6 packets into IPv4 packets. By using overlay tunnels, isolated IPv6 networks can communicate without upgrading the IPv4 infrastructure between them. Both routers and hosts can use tunneling with one of these techniques:

- **Manually configured:** Tunnel source and tunnel destination are manually configured with IPv4 and IPv6 addresses. You can configure tunnels between border routers or between a border router and a host.
- **Semiautomated:** A tunnel broker uses a web-based service to semiautomatically create a tunnel. A tunnel broker is a server on the IPv4 network that receives tunnel requests from dual-stack clients, configures the tunnel on the tunnel server or router, and associates the tunnel from the client to a tunnel server or router. A simpler model combines tunnel broker and server on one device.
- **Automatic:** These are automatic mechanisms to achieve tunneling:
  - **IPv4 compatible:** The tunnel is constructed using an IPv4-compatible address, an IPv6 address that consists of zeros, and an embedded IPv4 address in the last 32 bits. Because it does not scale, this mechanism is only appropriate for testing.
  - **6to4:** Each 6to4 site has a /48 prefix, which is the concatenation of 2002 and the IPv4 address of the edge router. 2002::/16 is a specially assigned address range for 6to4. When the edge router receives an IPv6 packet, the router extracts the IPv4 address embedded in the IPv6 destination address and encapsulates the IPv6 packet in an IPv4 packet with the extracted IPv4 address of the destination edge router. The destination edge router decapsulates the IPv6 packet from the received IPv4 packet and forwards the IPv6 packet to its final destination. To reach native IPv6 Internet, you need a 6to4 relay router that offers traffic forwarding to the IPv6 Internet.
  - **6over4:** You can use a router connected to a native IPv6 and with a 6over4-enabled interface to forward IPv6 traffic between 6over4 hosts and native IPv6. IPv6 multicast addresses are mapped into the IPv4 multicast addresses. The IPv4 becomes a “virtual Ethernet” for IPv6. To achieve that, you need an IPv4 multicast-enabled network.

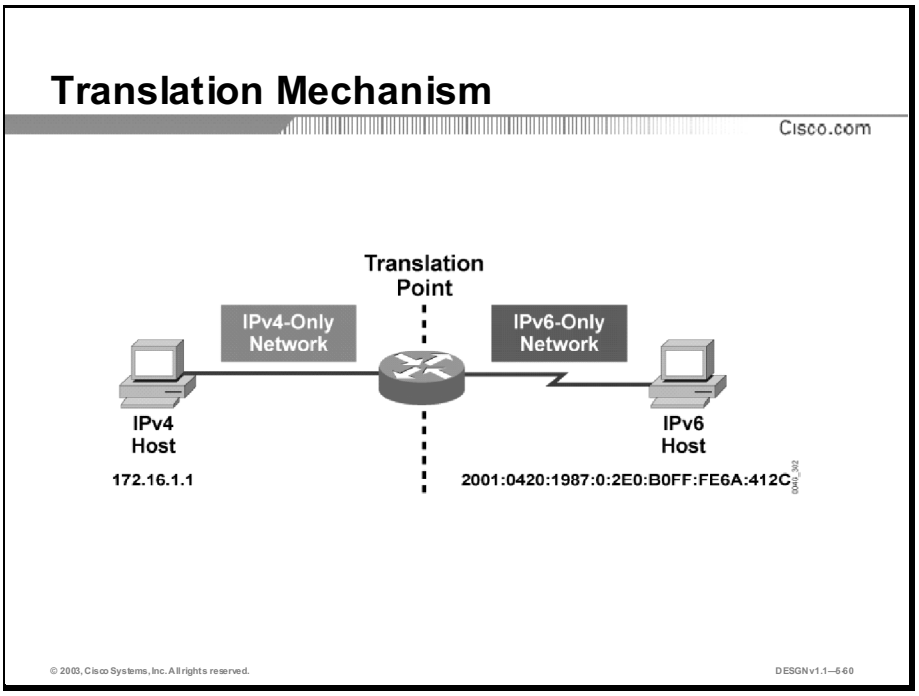

For legacy equipment that will not be upgraded to IPv6 and for some deployment scenarios, techniques that can connect IPv4-only nodes to IPv6-only nodes are available. Translation is basically an extension of NAT techniques.

An IPv6 node behind a translation device has full connectivity to other IPv6 nodes and offers NAT functionality to communicate with IPv4 devices. Possible solutions include:

- **Application Level Gateways (ALG):** Uses a dual-stack approach and enables a host in one domain to send data to another host in the other domain. This method requires that all application servers on a gateway run IPv6.
- **API:** You can install a specific module in a host's TCP/IP stack, for every host on the network. The module intercepts IP traffic through an API and converts it for the IPv6 counterpart.
- **Translation techniques:** Techniques are available to translate IPv4 addresses to IPv6 addresses and vice versa. You can translate at the transport layer or on the network layer as with current NAT devices. The two main network translation solutions are Network Address Translation-Protocol Translation (NAT-PT) and Dual Stack Transition Mechanism (DSTM).

Summary

This topic summarizes the key points discussed in this lesson.

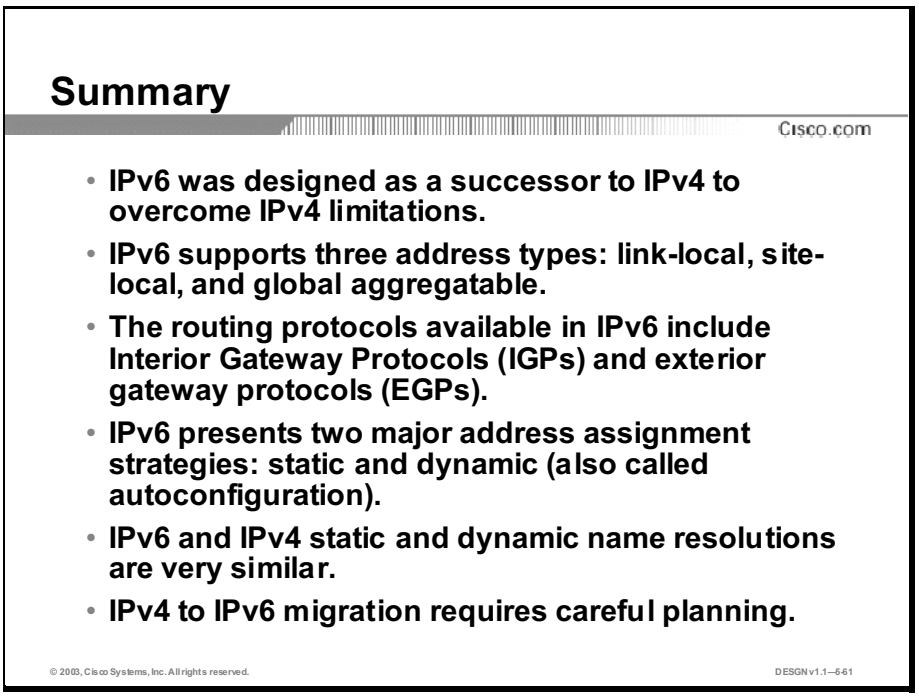

References

For additional information, refer to these resources:

- Comer, D. E., and D. L. Stevens. *Internetworking with TCP/IP*. Englewood Cliffs, New Jersey: Prentice-Hall; 1991.
- Bradner, S. O., and A. Mankin. *IPng Internet Protocol Next Generation*. Reading, Massachusetts: Addison-Wesley; 1995.
- *Designing Large-Scale IP Internetworks,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm
- *Subnetting an IP Address Space,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd20a.htm
- *Cisco IP Version 6 Solutions,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/ipv6_sol/index.htm
- *Cisco IOS IPv6,* http://www.cisco.com/warp/public/732/Tech/ipv6

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Case Study 5: Network Addressing Plan

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) An IPv6 address is _____ times larger than an IPv4 address.

A) two

B) four

C) six

D) eight

Q2) When compared to the IPv4 header, the IPv6 header provides _____ processing.

A) less efficient

B) more efficient

C) the same efficiency

D) inefficient

Q3) To shorten the IPv6 address 10a4:0000:aa1f:0000:0000:09c0:84a4:010b, it can be also written as _____.

A) 10a4:0:aa1f::9c0:84a4:010b

B) 10a4:0:aalf::09c:84a4:010b

C) 10a4::aa1f::9c0:84a4:010b

D) 10a4:0:aa1f::0::9c0:84a4

Q4) One-to-nearest communication means the use of _____ IPv6 addresses.

A) unicast

B) multicast

C) allcast

D) anycast

Q5) _____ IPv6 addresses are considered private addresses.

A) link-local

B) site-local

C) global aggregatable

D) universal

Q6) How many bits do global aggregatable addresses have for the site subnetting?

A) 8

B) 16

C) 32

D) 64

Q7) Which two IGP IPv6 routing protocols are currently supported on Cisco platforms? (Choose two.)

A) EGP

B) IGRP

C) IS-ISv6

D) RIPng

Q8) Which three address assignment strategies are available in IPv6? (Choose three.)

A) address assignment with DHCP

B) address assignment with DNS

C) static

D) stateless autoconfiguration

Q9) Which feature is an IPv6-related DNS features?

A) hierarchical addressing

B) A6 record: address record for the IPv6 host

C) distributed address resolution

D) dual-stack

Q10) Which three approaches offer valid IPv4 to IPv6 transition strategies? (Choose three.)

A) Application Level Gateways (ALG)

B) routing with RIPng

C) overlay tunnels

D) dual-stack approach

E) stateful address assignment

Quiz Answer Key

Q1) B

Relates to: IPv6 Address Structure

Q2) C

Relates to: IPv6 Address Structure

Q3) A

Relates to: IPv6 Address Structure

Q4) D

Relates to: IPv6 Address Types

Q5) A, B

Relates to: IPv6 Address Types

Q6) B

Relates to: IPv6 Address Types

Q7) C, D

Relates to: IPv6 Routing Protocol Considerations

Q8) A, C, D

Relates to: IPv6 Address Assignment Strategies

Q9) B

Relates to: IPv6 Name Resolution

Q10) A, C, D

Relates to: IPv4 to IPv6 Transition Strategies and Deployments

Case Study 5: Network Addressing Plan

Complete this case study to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the address design guidelines for IPv4 networks. Upon completing this case study, you will be able to meet these objectives:

- Propose the optimal IP addressing plan for a given network scenario
- Select the IP address assignment mechanisms that suit the customer needs

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Complete these steps:

Step 1 Refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario completely before the exercise. Allow a maximum of 10 minutes for reading.

Step 2 Discuss the scenario and options for IP addressing plan with your group. Allow 10 minutes for a discussion.

Step 3 Propose the optimal IP addressing plan for a given network scenario. Take into account that you will also have to design a new routing protocol. You will also need to ensure some WAN backup. The future campus will be completely restructured and more granular as well.

Step 4 Propose possible methods for IP address assignment.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class and you have justified any major deviations from the case study solution.

Module 6

Selecting Routing Protocols for a Network

Overview

Numerous factors are involved in the selection of a routing protocol, ranging from business requirements and technical features to the hierarchical structure of the newly designed network. This module presents general routing protocol features to evaluate during the selection process. You will weigh benefits and drawbacks of each routing protocol in your network design. The module concludes with a discussion of routing protocol deployment scenarios, covering convergence and redundancy, multiple routing protocols, redistribution, aggregation, and so forth.

Module Objectives

Upon completing this module, you will be able to select the appropriate routing protocol for a given network.

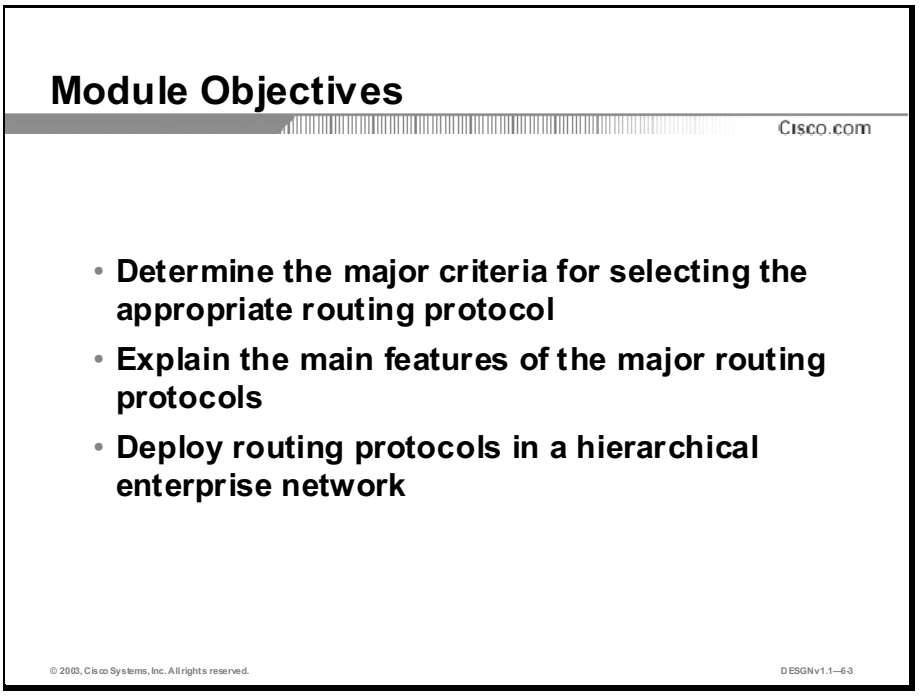

Module Outline

The outline lists the components of this module.

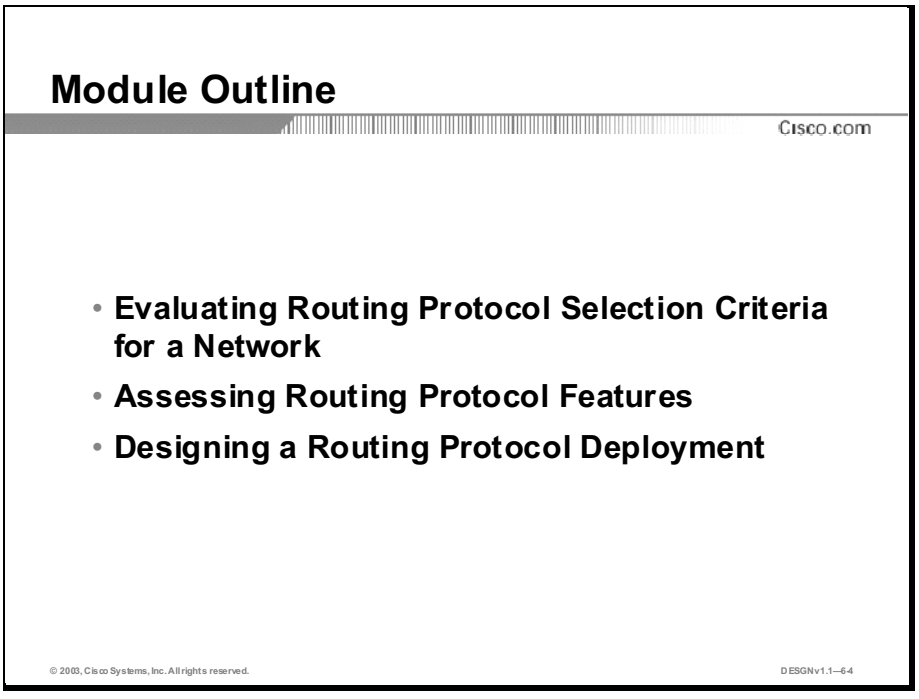

Evaluating Routing Protocol Selection Criteria for a Network

Overview

This lesson discusses the factors that dictate routing protocol use in a network. The lesson begins with the basic concepts of routing protocols and their features, in terms of metrics, scalability, convergence, resource utilization, and interoperability. This lesson evaluates each feature based on the business and technical requirements.

Relevance

You must consider the features of each routing protocol so you can select the best protocol for each environment.

Objectives

Upon completing this lesson, you will be able to determine the major criteria for selecting the appropriate routing protocol. This includes being able to meet these objectives:

- Select static or dynamic routing for a given network
- Describe the impact of each type of routing protocol on network resources
- Determine the need for interior and exterior routing protocols
- Evaluate routing protocol metrics
- Determine convergence requirements for the network
- Describe the differences between hierarchical and flat routing protocols

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of routing concepts

Outline

The outline lists the topics included in this lesson.

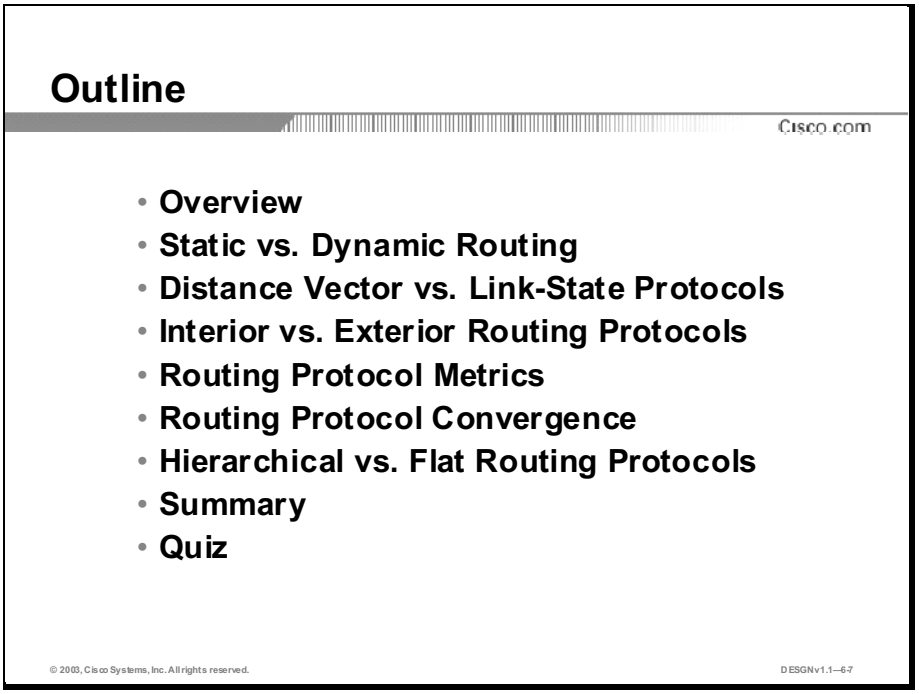

Static vs. Dynamic Routing

Static routing refers to manually configured routes while dynamic routing uses a routing protocol to make routing decisions. This topic describes static and dynamic routing and when each should be used.

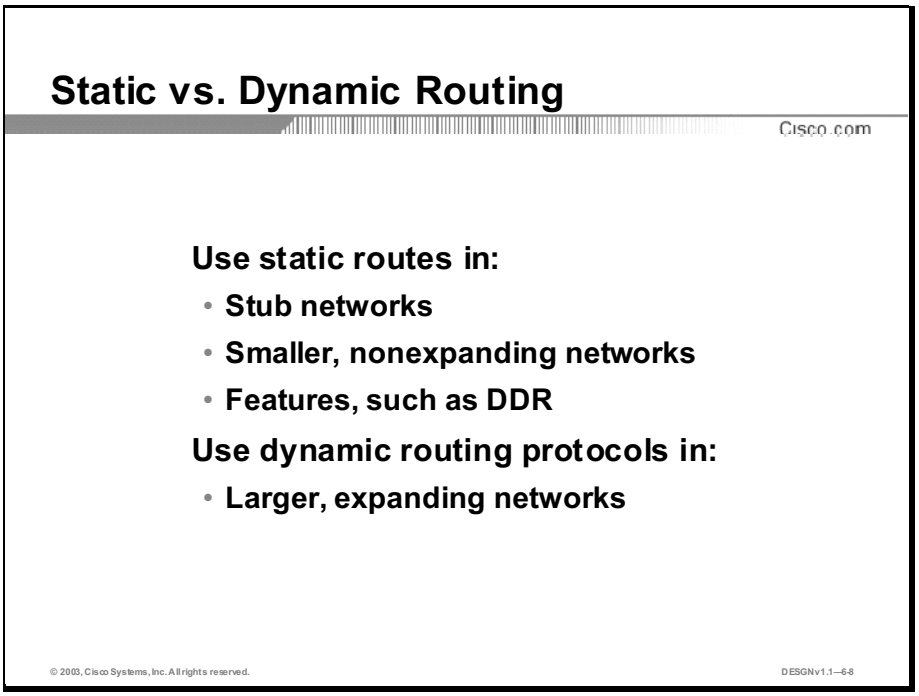

Static Routing Deployment

Static routing refers to the use of manually configured or injected static routes for traffic forwarding purposes. The authentication, authorization, and accounting (AAA) process injects the static routes. Static routing has these primary uses:

- Routing to and from stub networks. A stub network only carries traffic for local hosts. Typically, a stub network has only one entry or exit point. Even if it has paths to more than one other network, it does not carry traffic for other networks.
- Smaller networks that are not expected to grow significantly.
- Special features such as dial-on-demand routing (DDR).
- Specify routes toward dialing peers in dial-in environments.

Configuring and maintaining static routes is time consuming. It requires complete knowledge of the whole network for proper implementation.

Dynamic Routing Deployment

Dynamic routing refers to the use of a dynamic routing protocol to distribute routing information across the network and to select the best routes toward the destination. Dynamic routing protocols have two major advantages over static routes:

- Easy configuration and much less work for an administrator, even in small networks
- Dynamic adaptation to changes in the network

The use of dynamic routing protocols is favored in almost all network scenarios. Exceptions are, for example, DDR, a stub network, or a dial-in scenario.

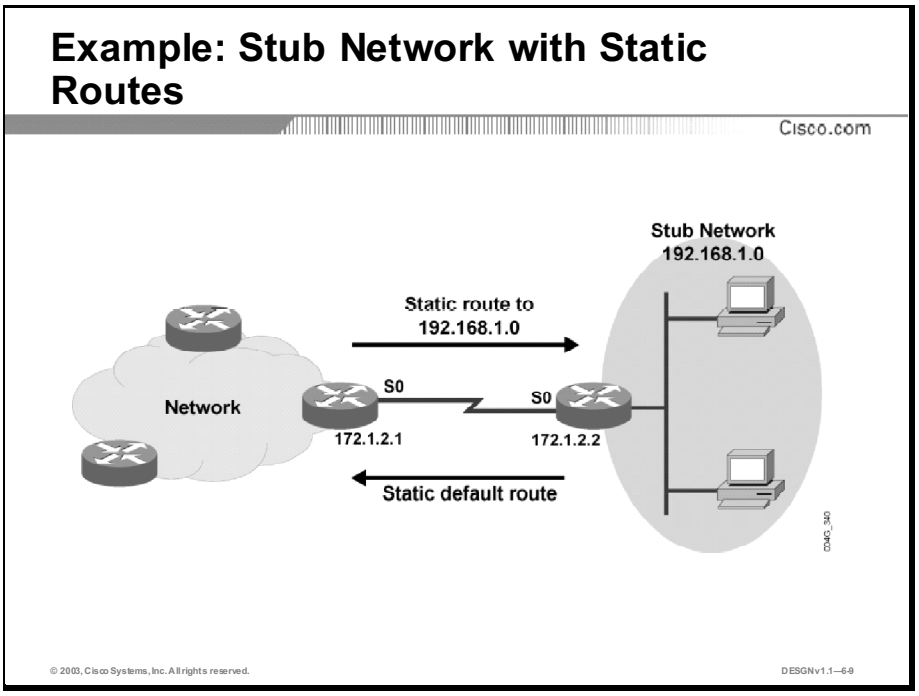

The stub network is a scenario in which the use of static routes is preferred over a dynamic routing protocol. The figure shows a stub network with a single entry or exit point over the S0 interface. On the stub network router, the static default route is configured so that all traffic toward destinations outside the stub network is forwarded via the S0 link. On the other side of the serial point-to-point connection, a static route toward the stub network is installed and then redistributed into the routing protocol, so that reachability information for the stub network is accessible to the rest of the network.

By using static and default static routes in a stub network, no control traffic from the dynamic routing protocol is present on the link or in the stub network. In addition, the processor and memory requirements for both routers are lower, and, in the stub network, a low-end router will suffice. Enterprises deploying a hub-and-spoke design of the WAN module in the Enterprise Edge functional area often deploy static and default static routes in a stub network.

On-Demand Routing (ODR) is a Cisco proprietary alternative to static routing. ODR uses the Cisco Discovery Protocol (CDP) to carry network information between spoke (stub) routers and the hub. ODR provides IP routing information with minimal overhead compared to a dynamic routing protocol. ODR requires less manual configuration than static routes. ODR is applicable in a hub-and-spoke topology only.

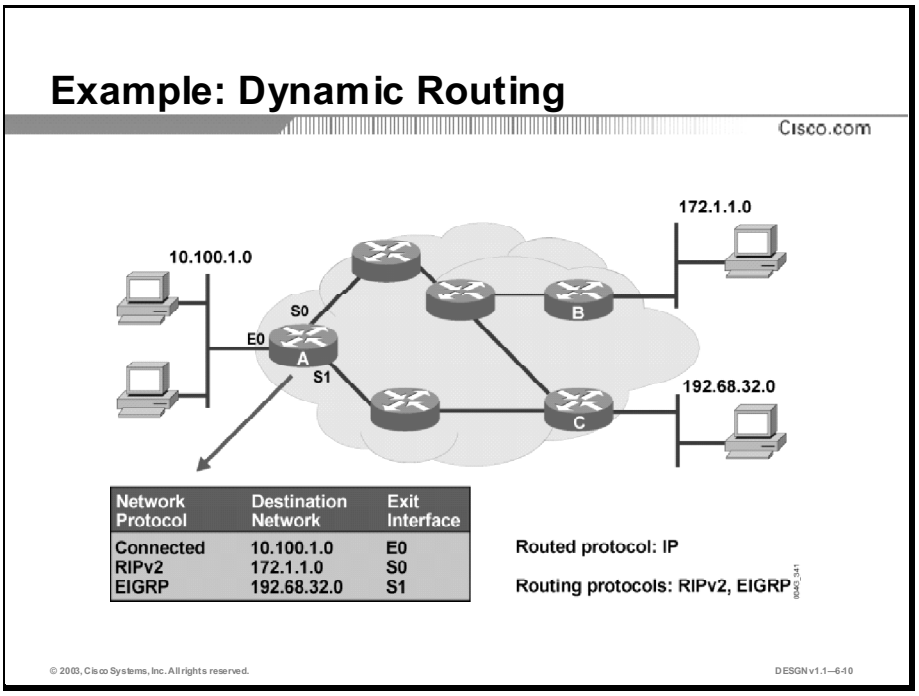

Dynamic routing protocols need to accomplish these tasks:

- Find sources from which to receive routing information (usually neighboring routers)
- Select best paths toward all reachable destinations based on received information
- Maintain this routing information
- Have a means to verify routing information (periodic updates or refreshes)

The figure shows the routers in a network running two dynamic routing protocols (some routers run Enhanced Interior Gateway Routing Protocol [EIGRP] and some run Routing Information Protocol Version 2 [RIPv2]), and IP is the routed protocol. IP reachability information is propagated throughout the network via RIPv2 and EIGRP.

Note: Running two or more routing protocols in the same network is sometimes necessary. As a consequence, information from one routing protocol is usually redistributed to another one.

Router A has all the required information to reach existing destinations. It selects the best paths for all three specified networks and inserts them into its IP routing table. The network is thus converged. The exchange of routing information is complete, and the network is in a stable state.

The path toward the 10.100.1.0 network is chosen as the only possible port, the connected E0 interface. The path toward network 172.1.1.0 is chosen via the S0 interface, following the information received from the RIPv2 routing process. The path toward the 193.64.32.0 network is chosen via the S1 interface, according to the information received from the EIGRP process.

Distance Vector vs. Link-State Protocols

The two basic categories of routing protocols are distance vector protocols and link-state protocols. This topic describes the differences between distance vector and link-state protocols.

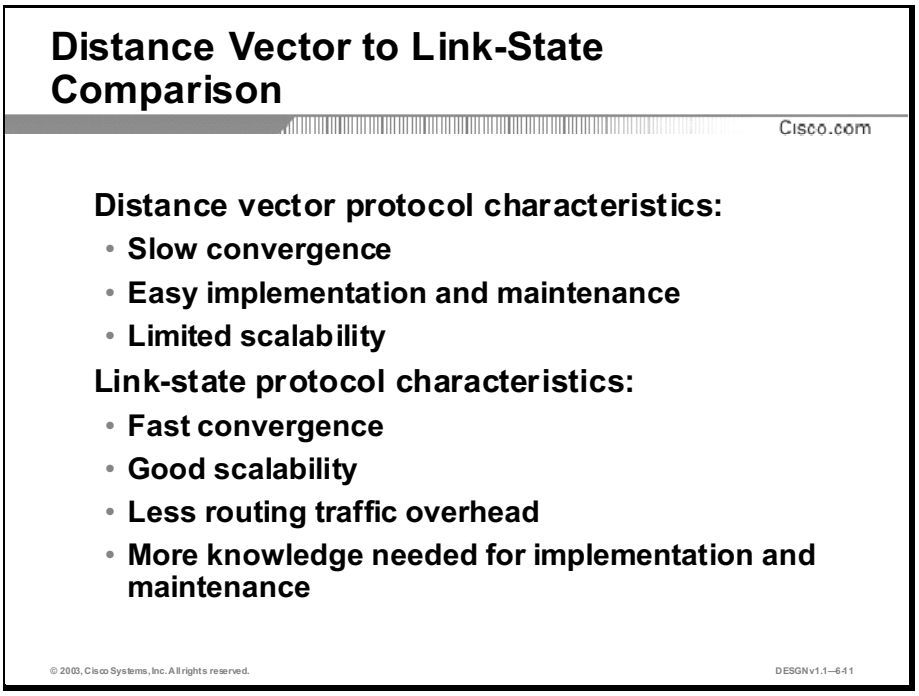

There are two types of interior gateway protocols:

- **Distance vector protocols:** "Routing by rumor" is the process where routing table maintenance decisions are made based on hearsay. Each router relies on its neighbor routers to maintain correct routing information. Each router passes only the results of local decisions to its neighbors. Examples of a distance vector protocol are Routing Information Protocol (RIP), RIPv2 (next-generation RIP), and Interior Gateway Routing Protocol (IGRP).
- **Link-state protocols:** Each router floods information about itself (its link-states) either to all other routers in the network or to a part of the network (area). Each router makes its own routing decision based on all received information using the common shortest path first (SPF) or Dijkstra's algorithm that calculates the shortest path to any destination. IP link-state protocols are Open Shortest Path First (OSPF) and Integrated Intermediate System-to-Intermediate System (IS-IS).

EIGRP has characteristics of both distance vector and link-state protocols. EIGRP supplements IGRP distance vector behavior with some link-state characteristics and some proprietary features. EIGRP is a very fast converging and scalable routing protocol.

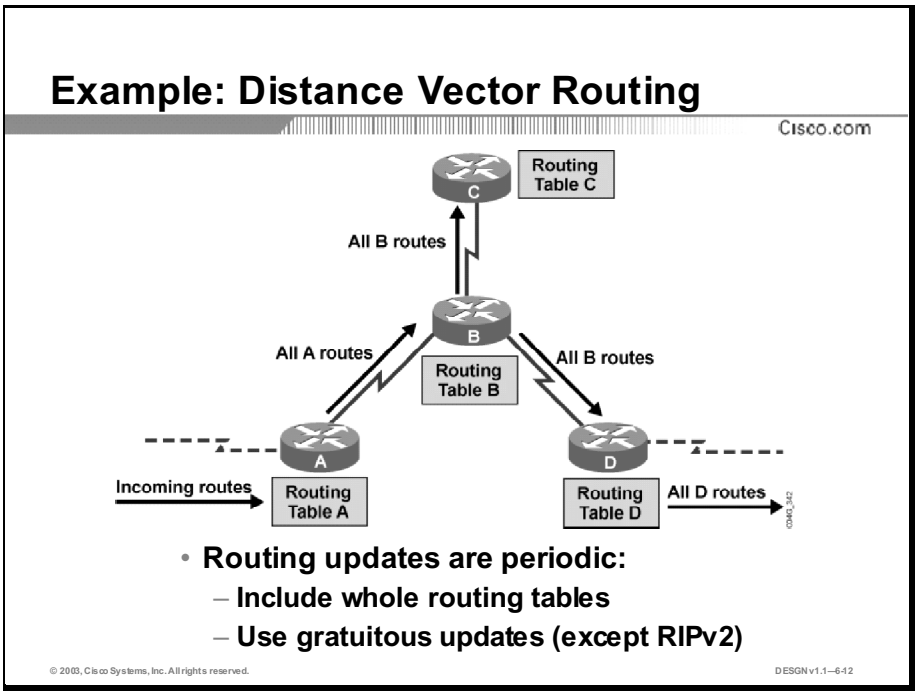

Distance vector routing is called “routing by rumor,” because routing table maintenance decisions are made locally based on hearsay from immediate neighbors. Distance vector protocols periodically send complete routing tables to all connected neighbors. Because triggered updates are not used (RIPv2 is the exception) and due to the speed of loop detection timers, convergence is slow.

In large networks, the routing table can become enormous, which causes significant traffic on the links.

An example of a distance vector protocol is RIPv2, which is a standardized protocol developed from the RIPv1 protocol. These are the characteristics of RIPv2:

- The metric for path selection is hop count.
- The maximum allowable hop count is 15.
- Routing updates are broadcast every 30 seconds by default.
- RIPv2 permits variable-length subnet masking on the network.

If a distance vector protocol seems adequate for network implementation and the network includes equipment from multiple vendors, RIPv2 might be a good solution.

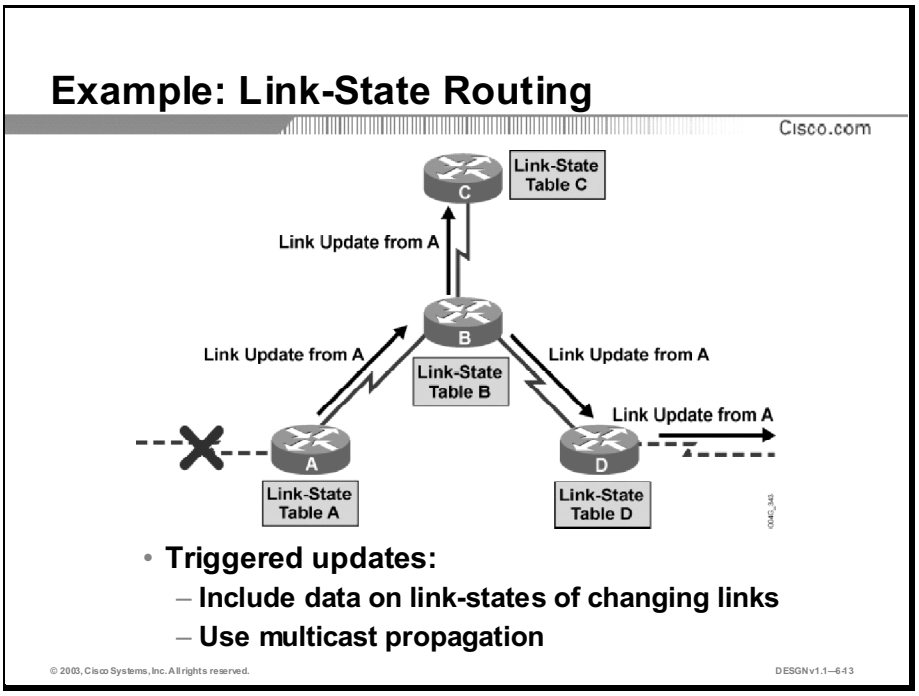

The link-state protocols OSPF and Integrated IS-IS use implementations of the hello protocol to establish neighbor relationships. From information shared through the neighbor relationship, each router understands a complete topology of the network. The SPF algorithm creates the shortest path tree for all reachable destinations and thus selects the best routes.

With link-state protocols, the information on connected links (subnets on those links) on all routers is flooded throughout the network, or to a specific area of the network. Therefore, all routers in the network have detailed knowledge of the whole network, unlike routers using distance vector routing protocols where routers receive knowledge only of the best routes from neighboring devices.

After the initial exchange of all link-states, and upon reaching the full state of operation, almost no periodic updates are sent through the network. (However, in OSPF, synchronizing periodic updates still take place every 30 minutes for each specific route, but not at the same time for all routes; this reduces the routing traffic volume.) Updates are triggered only when a change in a link-state occurs (for example, a link goes down or the bandwidth changes).

Most control packets used in link-state operations are sent via multicast. This can cause problems when deploying link-state protocols in nonbroadcast multiaccess (NBMA) networks, such as some Frame Relay and ATM topologies.

Choosing Between Distance Vector and Link-State Protocols

Guidelines are available to help ascertain which type of routing protocol to deploy.

Choose distance vector protocols when:

- It is a simple, flat network, which does not require a special hierarchical design.
- The administrators do not have enough knowledge to operate and troubleshoot link-state protocols.
- Specific types of networks, such as hub-and-spoke networks, are being implemented.
- Worst-case convergence times in a network are not a concern.

Choose link-state protocols when:

- The network design is hierarchical, usually large networks.
- The administrators have a good knowledge of the implemented link-state protocol.
- Fast convergence of the network is crucial.

Interior vs. Exterior Routing Protocols

Because the Internet is based on the autonomous system concept, two types of routing protocols are required: interior and exterior routing protocols. This topic describes the differences between interior and exterior protocols.

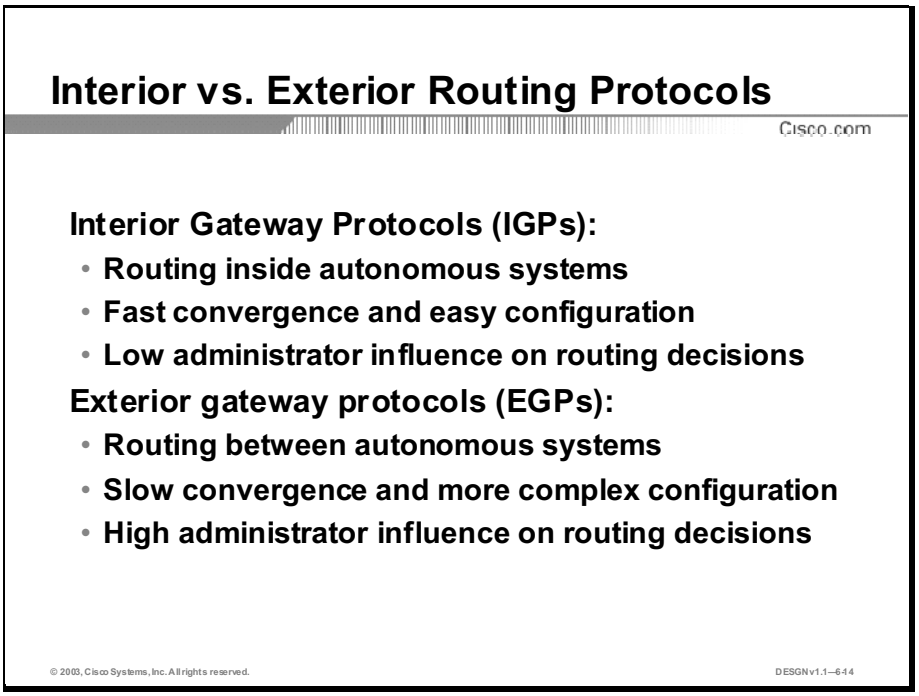

An autonomous system, otherwise known as a domain, is a collection of routers under a common administration, such as a company's internal network or an Internet service provider's network.

Therefore, protocols are required for:

- Intra-autonomous system (inside autonomous system) routing, Interior Gateway Protocols (IGPs).
- Inter-autonomous system (between autonomous system) routing, exterior gateway protocols (EGPs). The only widely used EGP protocol is Border Gateway Protocol (BGP).

Different types of protocols are needed because:

- Inter-autonomous system connections require more options to manually select routing characteristics. EGPs should be able to implement various policies.
- The speed of convergence and finding the shortest path to the destination are of crucial importance for intra-autonomous system routing protocols.

Therefore, the routing metrics of EGP protocols include more parameters, so the administrator can influence routing path selection. Alternatively, IGPs tend to use less complicated metrics to ease and speed up the decisions on best routing paths.

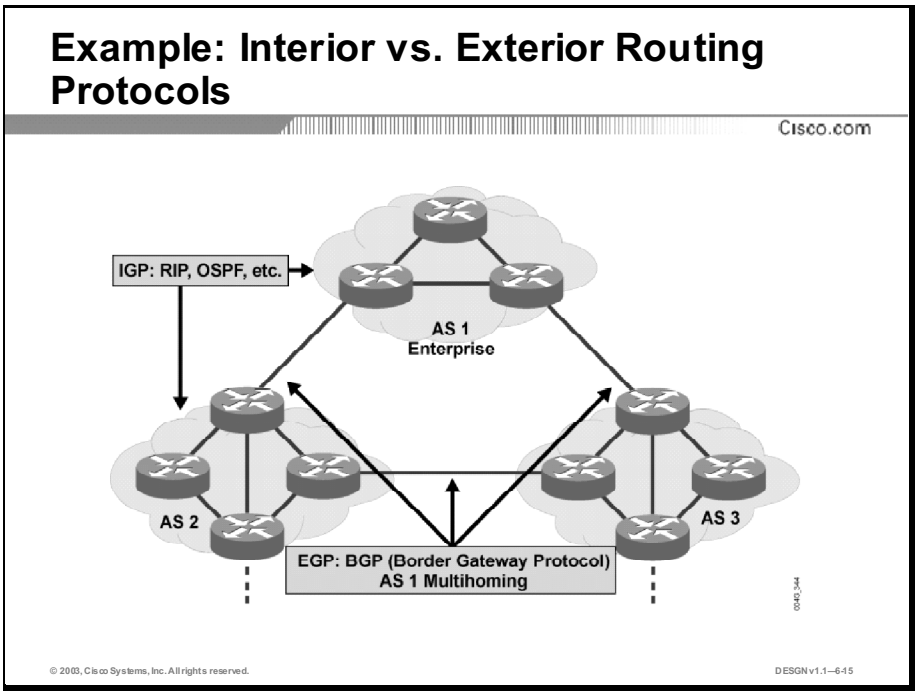

The figure shows three autonomous systems (domains). They are interconnected with inter-domain links and BGP. IGPs (RIP, RIPv2, OSPF, IGRP, IS-IS, and EIGRP) are implemented for intra-autonomous system (intra-domain) routing.

If autonomous systems need to communicate with each other, they require some form of inter-domain routing among the networks:

- In simple cases, static routes are used.
- Typically, an EGP is used.

BGP is the dominant EGP in use. BGP is especially useful when an autonomous system is connected to the Internet via multiple ISPs. This implementation of redundant connectivity for inter-autonomous system connections is called multihoming. An administrator can apply specific policies when using BGP (for example, traffic exit points, return traffic path, and levels of quality of service [QoS]) to comply with the contract requirements from specific Internet service providers (ISPs).

Routing Protocol Metrics

Routing protocols select the best paths to forward user data traffic across networks. They adapt dynamically to changes in the networks and try to maintain the best possible forwarding paths in all directions. To choose the most appropriate path through the network, routing protocols must be able to evaluate all the available paths. The value used in best path determination is called the routing protocol metric. This topic describes routing metrics and how they impact routing decisions.

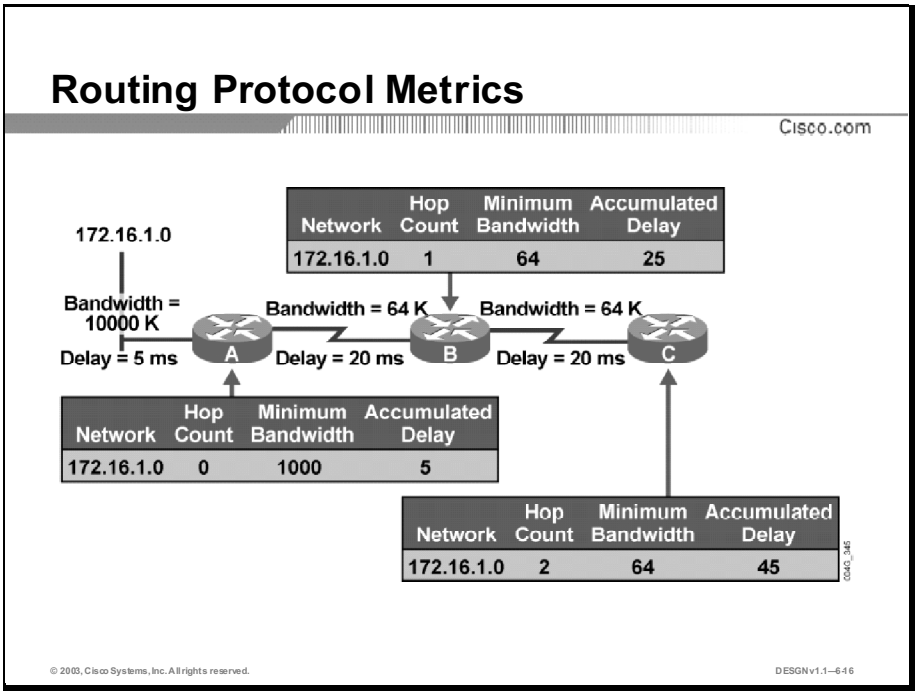

Different routing protocols use different parameters for routing metric calculation. The most popular are hop count (how many hops or routers away is the destination network), bandwidth (using the highest bandwith path), and delay (using the lowest latency path).

How are the parameters for routing metric calculation propagated?

The figure shows the left side network 172.16.1.0 connected to router A. The parameters for route metric calculation are forwarded in routing protocol updates using this procedure:

Step 1 Router A, which is the originator for the route 172.16.1.0, sends out the initial values to router B.

Step 2 Router B takes into account the parameters of its link toward router A, adjusts the parameters (bandwidth, delay, hop count) appropriately, calculates its metric toward the destination network, and sends the routing update toward router C.

Step 3 Router C again adjusts the parameters and calculates its metric toward the destination network 172.16.1.0 from those parameters.

In this case, the EIGRP method of route metric parameters was used, and the minimum bandwidth and the cumulative delay influenced the best path selection, the path with the highest minimum bandwidth and lowest delay is preferred.

In the case of RIP, only hop count is used to determine the best path (the path with the smallest hop count is preferred). In the case of link-state protocols, cumulative cost or metric is used (the lowest cost or metric path is selected). In general, the cost or metric reflects the bandwidth.

Note: In Cisco routers, you can manually configure the bandwidth and delay routing metric computation, which does not necessarily reflect the real speed of the link.

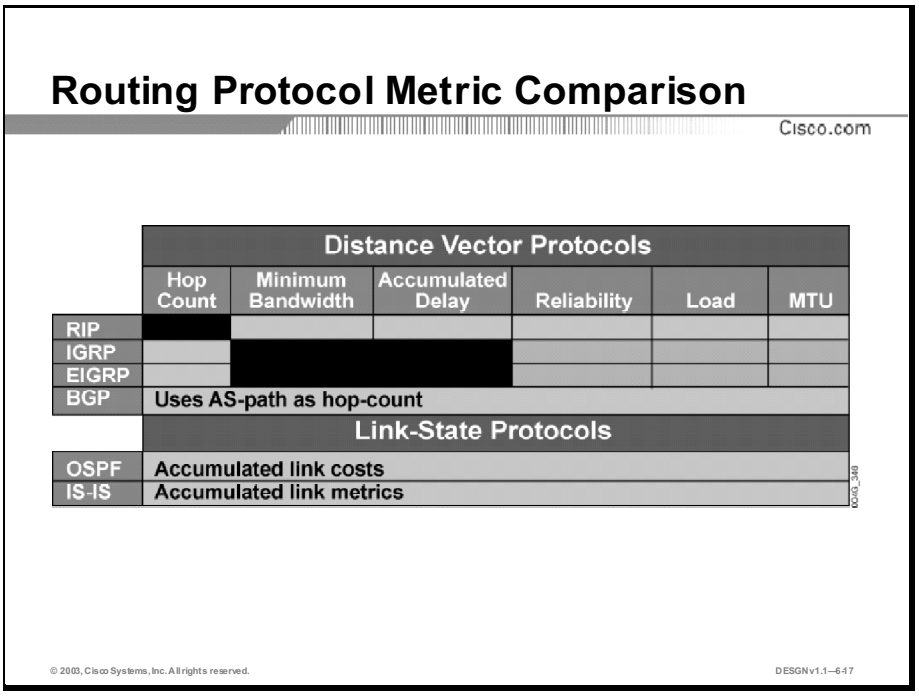

Different routing protocols calculate their routing metrics from different parameters and with different formulas. Some use very simple metrics (for example, RIP) and some quite complex metrics (IGRP, EIGRP).

Routing protocols that use only hop count for routing decisions (RIP, RIPv2) are not suitable for networks that use significantly different transmission speeds, because bandwidth is not considered. For networks using diverse media, routing protocols must consider bandwidth and possibly the delay of the links.

IGRP and EIGRP use almost the same method of metric calculation with bandwidth, delay, reliability, loading, and maximum transmission unit (MTU) as the metric criteria in a special formula. By default, only minimum bandwidth and accumulated delay of the path toward the destination network are considered. Use of the other parameters is not recommended because they can affect convergence and cause routing loops, if misconfigured.

OSPF and IS-IS use the cost or metric for path calculation, respectively. The cost in OSPF and metric in IS-IS normally reflect the bandwidth of the link. As a result, the highest accumulated bandwidth (lowest cost or metric) is used to select the best path.

BGP, as the only representative of the EGP family, uses the autonomous system path attribute as its metric. The length of the attribute (the number of autonomous systems to traverse for reaching a destination) is usually one of the factors that influence the path selection and can be compared to hop count. BGP incorporates additional path attributes that you can configure to influence routing decisions.

Routing Protocol Convergence

Convergence occurs whenever a network's topology changes and all routers in that network must learn the new topology. This topic discusses routing protocol convergence.

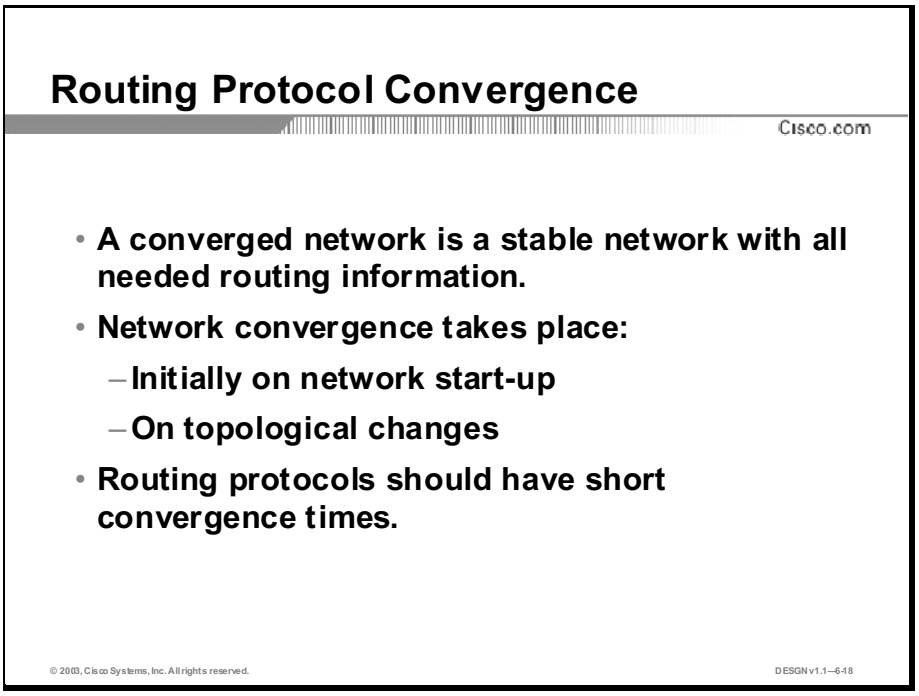

Convergence is both collaborative and independent. The routers share information with each other but must independently calculate the impacts of the topology change on their own routes. Because they develop an agreement with the new topology independently, they are said to converge on this consensus.

Convergence properties include the speed of propagation of routing information and calculation of optimal paths. The quicker the convergence, the better the routing protocol.

Network convergence or propagation of routing information occurs whenever a new routing protocol is started in the network. A network is not completely operable until the network has converged, so routing protocols require short convergence times.

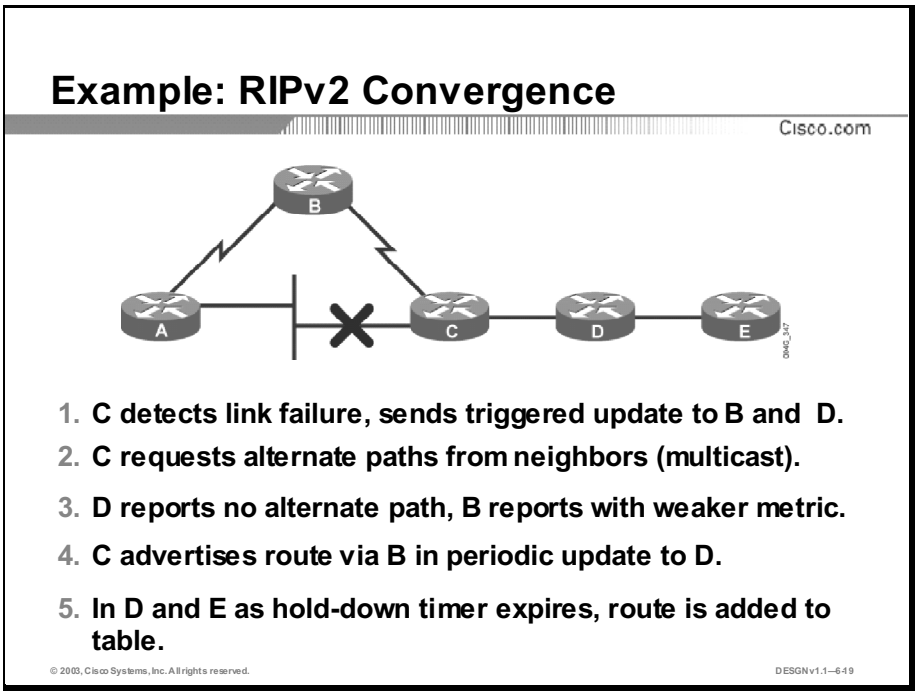

RIPv2 is a distance vector protocol that propagates its routing information periodically. Distance vector protocols use the principle of holddown to prevent routing loops. Putting a route in holddown after the route has failed means that if a routing update arrives with the same or worse metric, the new route will not be installed. The route in holddown is still used to forward traffic for the whole hold-down period, even though the destination may no longer be reachable.

In the figure, the Ethernet link between routers A and C has failed. The network runs RIPv2 as a routing protocol. The RIPv2 convergence steps are:

Step 1 Router C detects link failure and sends a triggered update to routers D and B.

This triggered update route is poisoned (hop count is too great) to B and D and removed from C's routing table.

Step 2 Router C sends a request to neighbors for alternate path.

A broadcast request is used for RIPv1 and a multicast request is used for RIPv2.

Step 3 Router D reports no alternate path; B reports route with weaker metric.

The route via router B is immediately placed in the routing table.

Step 4 Router C advertises the route via router B in a periodic update to D.

There is no change to router D's table because the route is in holddown.

Step 5 In router D, as hold-down timer expires, the route is added to the table and propagated to router E. The new route is propagated in a periodic update.

Convergence time at router E consists, therefore, of the hold-down time plus one or two update intervals.

Routing Protocol Convergence Comparison

Cisco.com

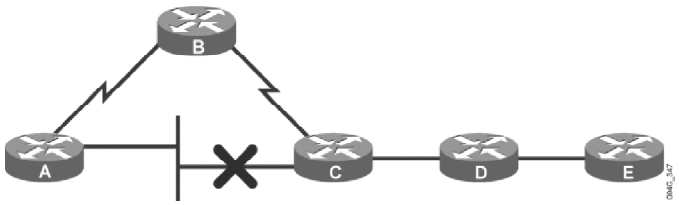

Protocol	Convergence Time to Router E
RIP	Holddown + 1 or 2 update intervals
IGRP	Holddown + 1 or 2 update intervals
EIGRP	Matter of seconds
OSPF	Matter of seconds

In a specific network, different routing protocols need different amounts of time to converge. Pure distance vector protocols are slower than link-state protocols, because they use periodic updates and the hold-down mechanism. Thus, when network convergence is of crucial importance in the design requirements, use the fast converging protocols.

Link-state protocols converge much faster because they instantly propagate routing updates. Therefore, whenever a change in a link-state occurs, a link-state update floods through the entire network, which results in fast convergence. There is no need to wait for the next periodic update, as you do with distance vector protocols.

Note: The default hold-down time is 180 sec for RIP and 280 sec for IGRP. You manually adjust these values.

EIGRP is a special case because it incorporates the distance vector principle of metric propagation (only best information is sent to the neighbors), but it has no periodic updates and does not implement the principle of holddown. EIGRP stores available backup routes in its topology table. When a certain lost destination has a backup route, the switchover to the best backup route is almost immediate and involves no action from other devices in the network. With proper EIGRP deployment, you can achieve very fast convergence.

Note: When all edge routers of a network converge, then the complete network has also converged.

Hierarchical vs. Flat Routing Protocols

Flat routing protocols propagate all routing information throughout the network while hierarchical routing protocols divide large networks into smaller areas. This topic provides a comparison of hierarchical and flat routing protocols.

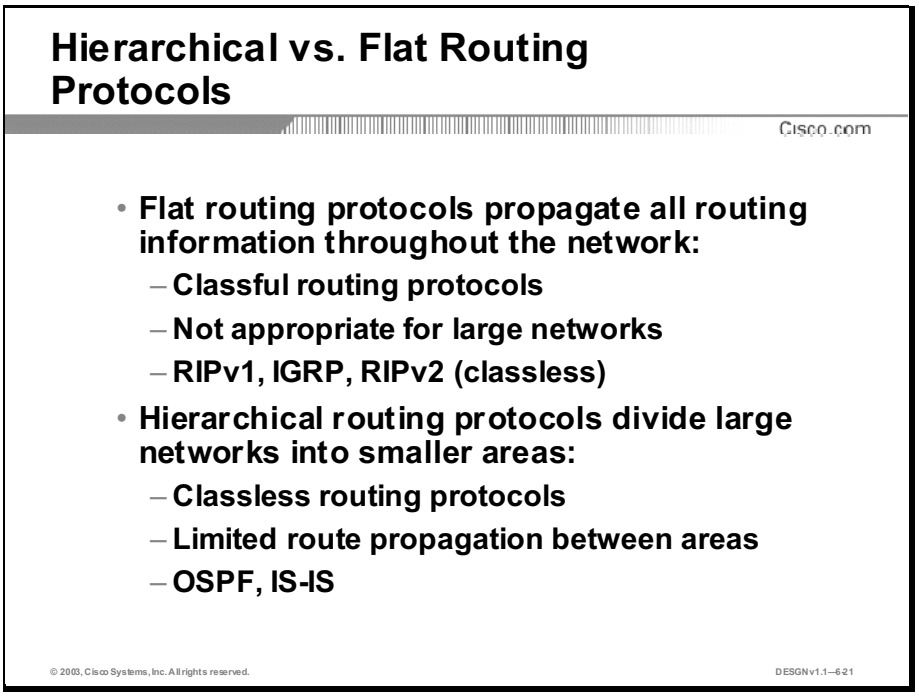

Flat routing protocols (classful distance vector protocols) cannot limit route propagation in a major network environment. Classful routing means that the protocol performs automatic summarization of network information only on major class network boundaries (Class A, B, or C). These protocols require fixed-length subnets. Flat protocols do not scale well because they produce significant volumes of routing information, which can easily affect a large network.

To support large networks, you can implement manual route summarization (creation of a nonmajor network summary address for a block of subnets) and area-based design.

Link-state protocols are hierarchical. Large networks are divided into multiple areas using link-state protocols such as OSPF and IS-IS. With route summarization, smaller routing updates are propagated among areas, resulting in higher scalability and a better fit for large networks. Link-state protocols are classless and support variable-length subnet masking (VLSM). You can implement the summarization on an arbitrary boundary within an IP address.

EIGRP is also a flat routing protocol, although it supports VLSM. Because EIGRP supports manual summarization, you can emulate a hierarchical network design using EIGRP. A hierarchical design is not necessary in EIGRP, but it is recommended for large networks.

Example: Flat and Hierarchical Networks

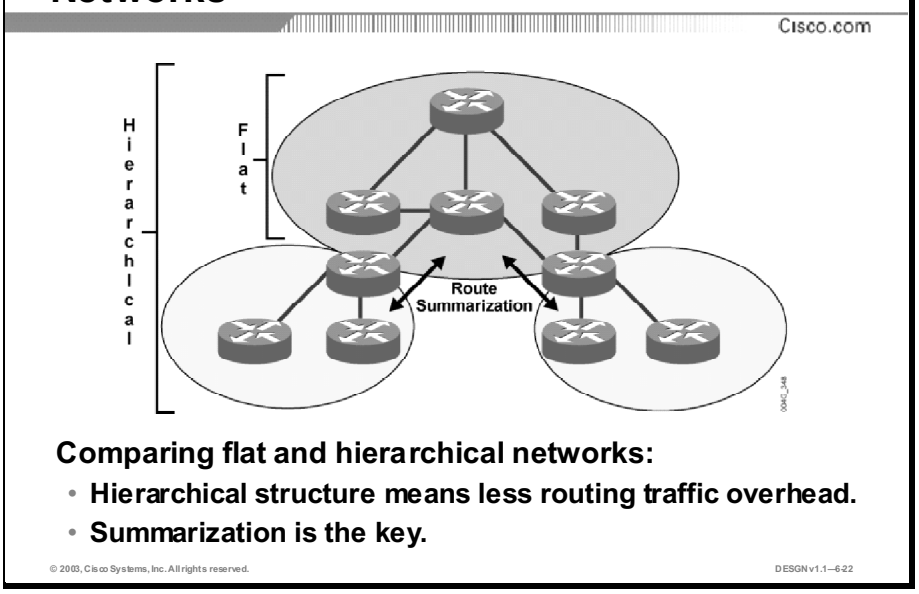

When using flat routing protocols, no summarization of route information is done within a major network. The automatic summarization of subnet prefixes takes place only on boundaries among major networks.

Running a large flat network is not scalable because routing traffic will consume too much of the network resources, which the routed traffic (application data, user traffic) should use.

You should implement a hierarchy in the network using manual summarization. This effectively decreases the amount of routing traffic among the different parts of the network.

With the help of summarization, you can isolate the impact of instabilities in a part of the network, greatly improving convergence.

Link-state protocols include hierarchies in their structures and implement the concept of backbone and nonbackbone areas. You can manually perform summarization, which is required in most cases.

Summary

This topic summarizes the key points discussed in this lesson.

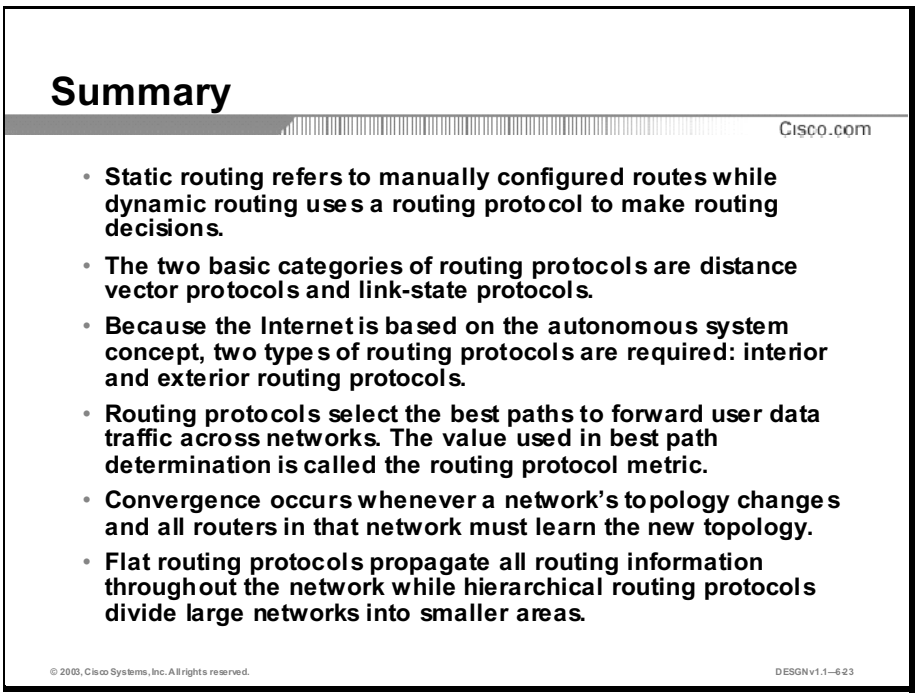

References

For additional information, refer to these resources:

- *Building Scalable Cisco Internetworks* (BSCI) course
- *Interconnecting Cisco Network Devices* (ICND) course

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match each type of networks with the type of routing most appropriate for it.

A) small, nonexpanding networks

B) dial-in networks

C) large, expanding networks

D) small and medium, quick-expanding networks

E) stub networks

_____ 1. dynamic routing

_____ 2. static routing

Q2) Which four features are implemented differently in distance-vector protocols compared to link-state protocols? (Choose four.)

A) routing updates

B) IP routing tables

C) routing information distribution

D) routing of data traffic

E) verification of routing information sources

F) routing information storage

Q3) What are the two differences between Interior Gateway Protocols (IGP) and Exterior Gateway Protocols (EGP)? (Choose two.)

A) IGPs must be able to converge quickly, but the speed of convergence is not of crucial importance for EGPs.

B) IGPs are better at finding the fastest paths across the network.

C) IGPs can be more easily substituted with static routing.

D) IGPs are used for interconnecting autonomous systems, while EGPs are used for intra-autonomous system connections.

Q4) Match the routing protocols with the metric parameters they use for metric calculation. (You may have to match more than one metric with a specific protocol.)

A) RIP

B) EIGRP

C) OSPF

D) IGRP

E) BGP

_____ 1. autonomous system path

_____ 2. minimum bandwidth

_____ 3. cumulative delay

_____ 4. Hop count

_____ 5. cumulative cost

Q5) What effect does the speed of convergence have on consistency of routing information throughout the network?

A) The speed of convergence influences the frequency of routing updates.

B) If the convergence is fast, the routing information is less consistent.

C) Consistent routing information has nothing to do with convergence speed.

D) If the convergence is fast, the routing information is more consistent.

Q6) Select the major advantage of a hierarchically structured network over a flat network.

A) Fewer IP addresses are used in flat networks.

B) Route updates are quicker in flat networks.

C) Hierarchical structure means less routing traffic overhead.

D) Hierarchical approach has no advantage over a flat approach.

Quiz Answer Key

Q1) 1=C, D
2=A, B, E

Relates to: Static vs. Dynamic Routing

Q2) A, C, E, F

Relates to: Distance Vector vs. Link-State Protocols

Q3) A, B

Relates to: Interior vs. Exterior Routing Protocols

Q4) 1=E
2=B, D
3=B, D
4=A
5=C

Relates to: Routing Protocol Metrics

Q5) D

Relates to: Routing Protocol Convergence

Q6) C

Relates to: Hierarchical vs. Flat Routing Protocols

Assessing Routing Protocol Features

Overview

Each routing protocol is suited to different environments. This lesson helps you select interior routing protocols, including On-Demand Routing (ODR), RIP, EIGRP, OSPF, and Integrated IS-IS. BGP, an exterior routing protocol, is briefly mentioned.

Relevance

Based on the suitability of each routing protocol within a certain design scenario, you will be able to select the most appropriate protocol.

Objectives

Upon completing this lesson, you will be able to explain the main features of the major routing protocols. This includes being able to meet these objectives:

- Describe the scenarios for using ODR as a routing protocol
- Explain the appropriateness of RIP for routing in a network
- Assess the EIGRP as a potential routing protocol in the Campus Backbone of a network
- Assess the OSPF as a potential routing protocol in the Campus Backbone of a network
- Describe the features of Integrated IS-IS that drive its use in a network
- Determine the features and need for BGP in a network
- Use decision routing tables to select an interior gateway routing protocol based on network requirements

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic knowledge of routing protocol concepts and existing routing protocols
- “Evaluating Routing Protocol Selection Criteria for the Network” lesson

Outline

The outline lists the topics included in this lesson.

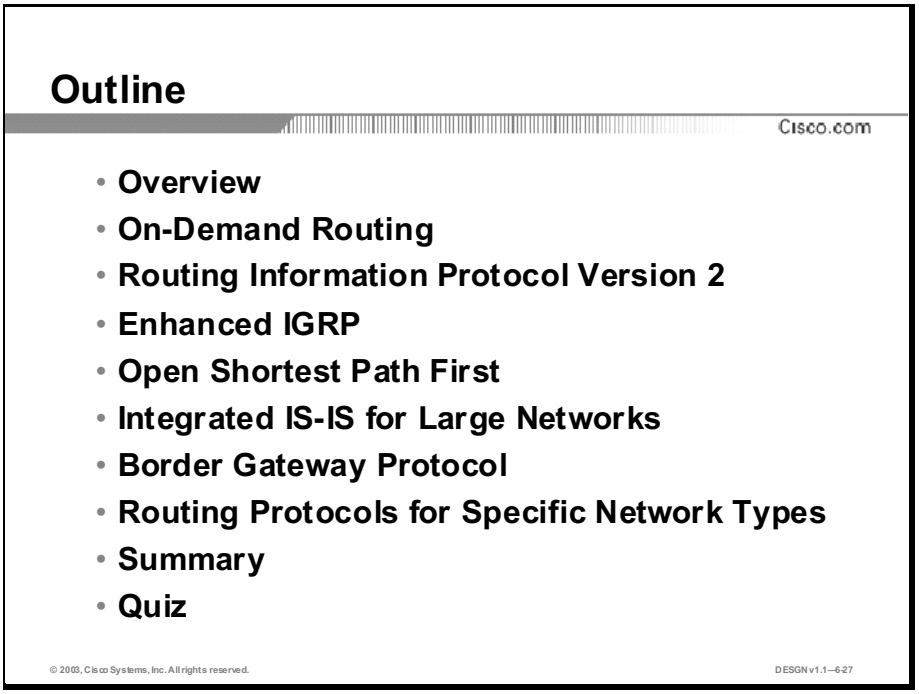

On-Demand Routing

ODR is a Cisco proprietary feature that provides IP routing with minimum overhead for stub networks. It avoids the overhead of a general, dynamic routing protocol without incurring the configuration and management overhead of static routing. This topic describes ODR.

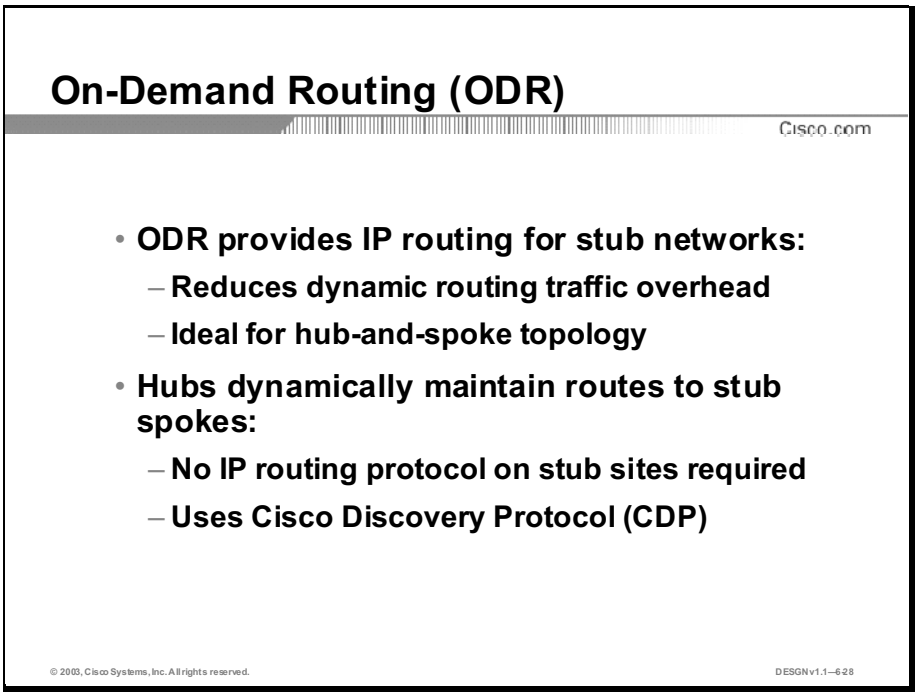

A stub router resembles a spoke router in a hub-and-spoke network topology, where the only router to which the spoke is adjacent is the hub router. The stub routers have a common WAN connection to the hub router, and a small number of LAN segments (stub networks) are directly connected to the stub router.

Stub networks may consist of only end systems and the stub router, and do not require the stub router to learn any dynamic IP routing information.

ODR supports easy installation of IP stub networks, where the hubs dynamically maintain routes to the stub networks. During installation, you do not need to configure an IP routing protocol on the stubs.

ODR uses the Cisco Discovery Protocol (CDP) to carry minimal routing information between the hub and the stub routers.

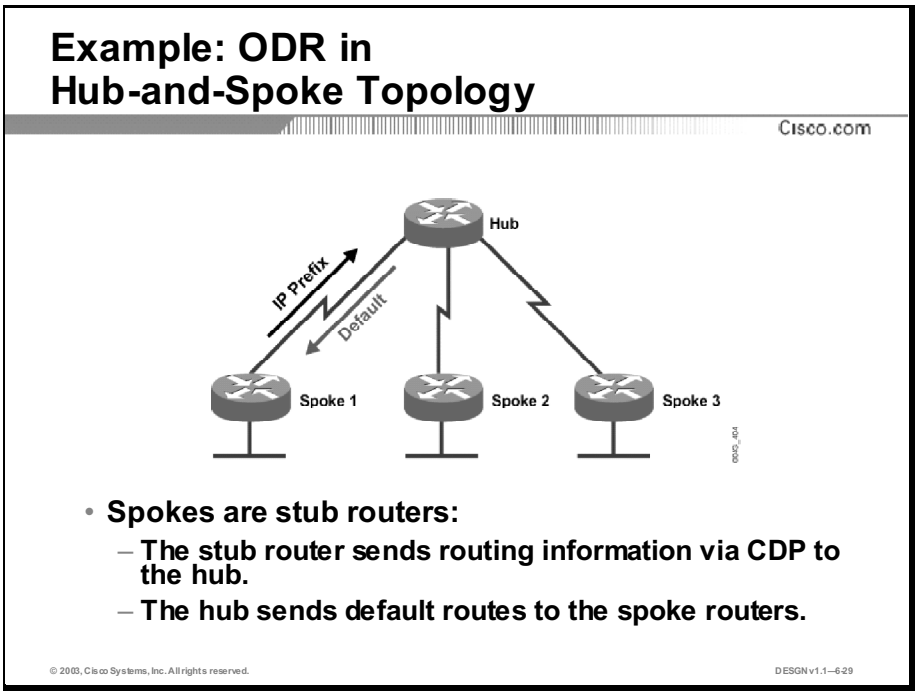

The hub router provides default route information to the stub routers, eliminating the need to configure a default route on each stub router. Spokes send IP prefixes to the hub.

On stub routers that support ODR, the stub router advertises IP prefixes that correspond to the IP networks configured on all directly connected interfaces. If an interface has multiple logical IP networks configured, only the primary IP network is advertised through ODR. Because ODR advertises IP prefixes, ODR is able to carry VLSM information.

ODR Operation

Once ODR is enabled on a hub router, the hub router begins to install stub network routes in the IP forwarding table. You can configure the hub router to redistribute these routes into any configured dynamic IP routing protocols.

On the stub router, you do not need to configure an IP routing protocol. From the standpoint of ODR, a router is automatically considered to be a stub when no IP routing protocols are configured.

Note:	ODR is a Cisco proprietary feature and, therefore, can only be implemented on connections between Cisco devices.

Routing Information Protocol Version 2

RIPv2 was introduced to address the need for a simple distance vector protocol that supported VLSM, manual summarization, and authentication of routing data. This topic describes RIPv2.

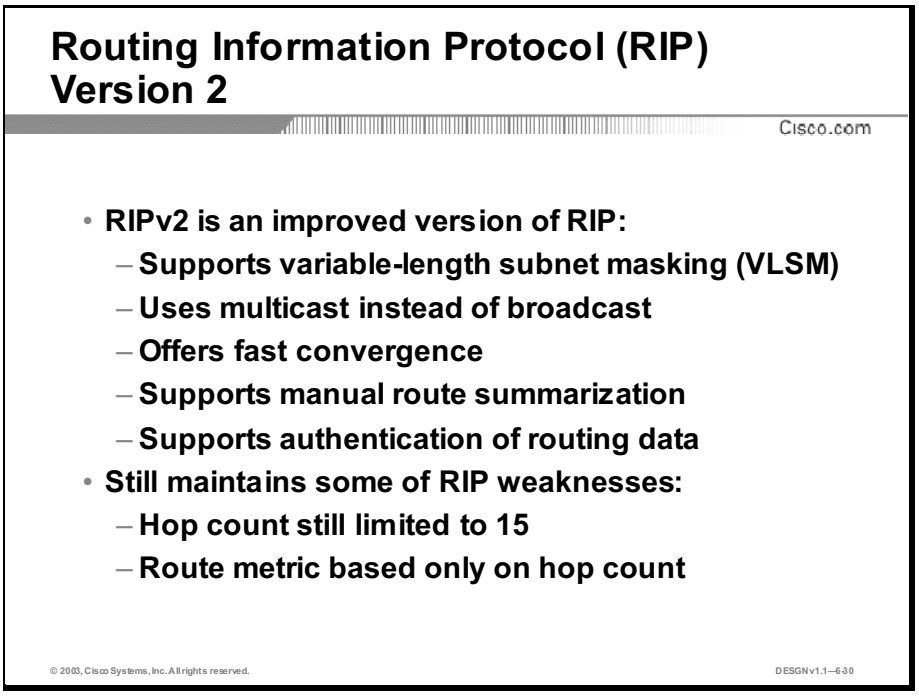

Advantages of RIPv2 over RIPv1

RIPv2 supports VLSM because it carries subnet prefixes in routing updates. It uses multicast for routing traffic propagation to reduce routing traffic overhead in multiaccess networks.

In RIPv2, manual summarization of routes is available, enabling a degree of hierarchy to be introduced in RIPv2 networks.

RIPv2 introduced routing packet authentication to lessen the possibility of attacks on the RIPv2 routing process. You can use authentication in plain text mode, where passwords are sent in a nonencrypted form (allowing replay attacks) or in Message Digest 5 (MD5) mode, where passwords are never sent but are used in the one-way encryption process. The MD5 encryption process creates a “fingerprint” of a packet, which the receiving router then verifies.

Inherited Weaknesses

RIPv2 still maintains two major weaknesses from RIPv1 that limit its applicability:

- The only parameter used in best path selection is the hop count.
- The limitation of the maximum hop count remains at 15.

Caution:	Be careful when combining RIPv2 routers and RIPv1 compatible hosts. Because it cannot apply the supplied subnet mask, a RIPv1 host may misinterpret route information.

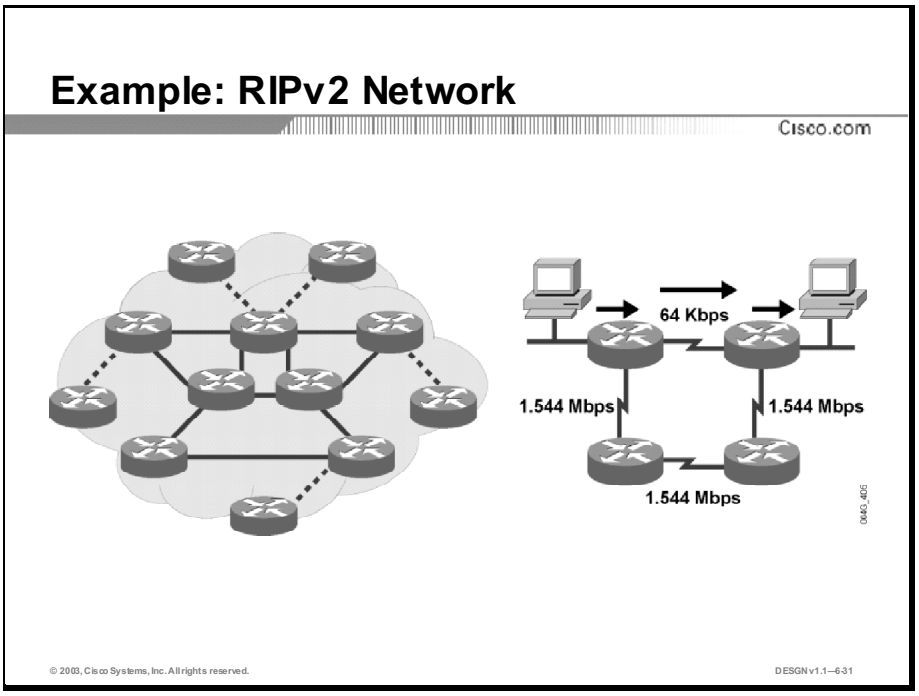

A RIPv2 network can have a network diameter of a maximum of 15 hops (hop is the transition of a packet from one router to the next). If correctly implemented, RIPv2 is suitable only in small to medium networks.

Note: The network diameter is the distance in hops between the most distant edge routers in the network.

A more serious problem is the use of hop count as the only metric parameter. In networks with similar media characteristics (bandwidth, delay of links), hop count does not pose a problem. In networks with diverse physical media speeds, RIPv2 is more likely to choose suboptimal paths for routed data traffic. The figure illustrates how RIPv2 will choose the path for data traffic from one PC to another across a slower 64-kbps serial link, instead of across the three much wider T1 (1.544-Mbps) links.

Enhanced IGRP

EIGRP is a Cisco proprietary protocol for routing IPv4, IPX, and AppleTalk traffic. This topic describes the basic characteristics of EIGRP.

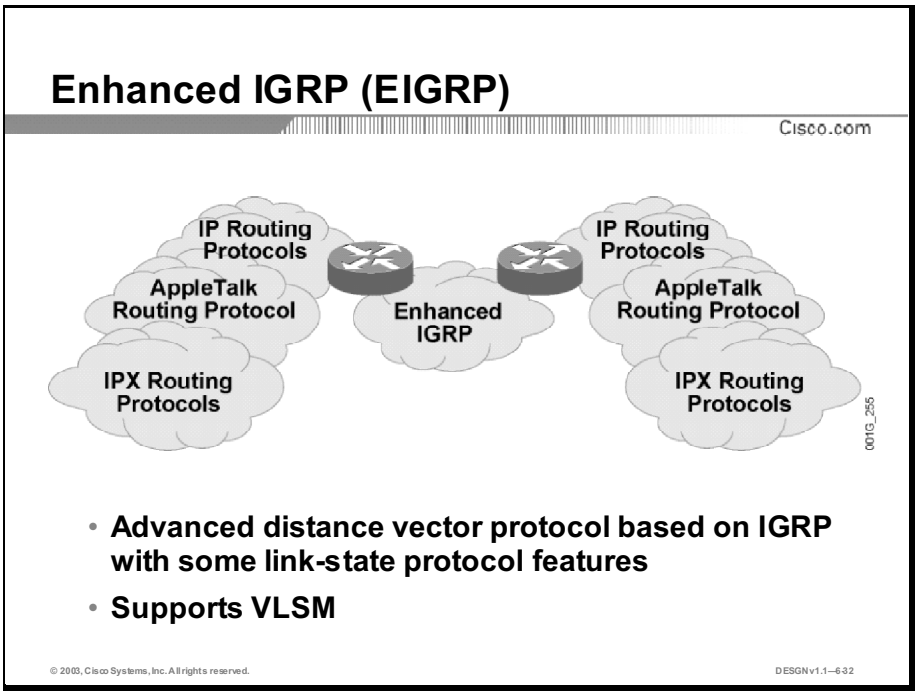

Enhanced IGRP was developed from the IGRP, which is a pure distance vector protocol. EIGRP is a hybrid routing protocol, which is a distance vector protocol with additional link-state protocol features.

These are some additional EIGRP:

- Triggered updates (EIGRP has no periodic updates)
- Use of a topology table to maintain all the routes received from neighbors (not only the best ones)
- Establishment of adjacencies with neighboring routers using the hello protocol

Other advantages of EIGRP are its support for VLSM and manual route summarization. These allow EIGRP to create hierarchically structured large networks.

Routes are propagated in EIGRP in a distance vector manner, from neighbor to neighbor, and only the best routes are sent onward. A router running EIGRP does not have a complete view of a network because it sees only the routes received from its neighbors. In a pure link-state operation (OSPF, IS-IS), all routers in the same area have identical information and, therefore, have a complete view of the area and its link states.

EIGRP can use minimum bandwidth, cumulative delay of the path, worst reliability between source and destination, worst loading on a link between source and destination, and the smallest MTU as parameters in the metric calculation, but, by default, only minimum bandwidth and cumulative delay of the path are used.

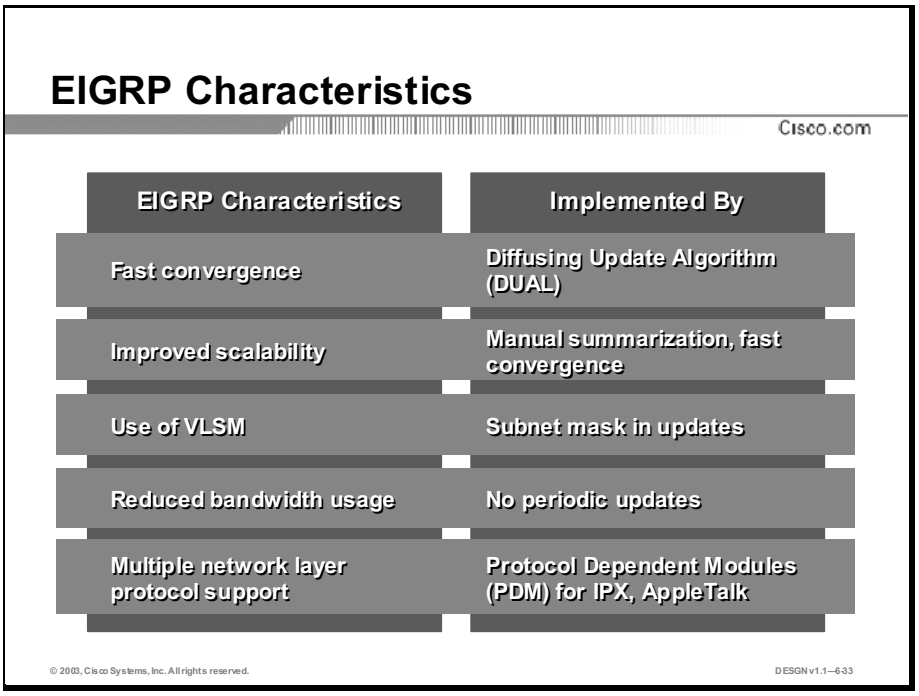

An advantage of EIGRP is its fast convergence Diffusing Update Algorithm (DUAL) route calculation mechanism. DUAL allows the insertion of backup routes (also known as feasible successors) into the EIGRP topology table, which are used in case of primary route failure. Because it is a local procedure, the switchover to the backup route is immediate and does not involve action in any other routers.

EIGRP, by default, summarizes routes on the classful network boundaries. You can turn off the autosummarization, and incorporate manual summarization. Manual summarization of subnet routes improves scalability and network performance because the routing protocol uses fewer resources.

Because EIGRP does not use periodic routing table updates, it uses less bandwidth, especially in large networks, where the number of routes becomes very large. On the other hand, EIGRP uses the hello protocol to establish and maintain adjacencies with its neighbors. If many neighbors are reachable over the same physical link, as is the case in NBMA networks, the hello protocol might create a significant routing traffic overhead. Therefore, you must design the network appropriately to use all of the EIGRP advantages.

EIGRP supports multiple network layer protocols through Protocol Dependent Modules, which include support for IPv4, IPX, and AppleTalk.

Note: EIGRP is a Cisco proprietary protocol and can only pass protocol information with licensed devices.

Open Shortest Path First

OSPF is a standardized protocol for routing IPv4. It was developed in 1988 by the Internet Engineering Task Force (IETF) to replace RIP in larger, more diverse media networks. Version 2 is described in RFC 2328 (http://www.ietf.org/rfc/rfc2328.txt?number=2328). This topic describes the basic characteristics of OSPF.

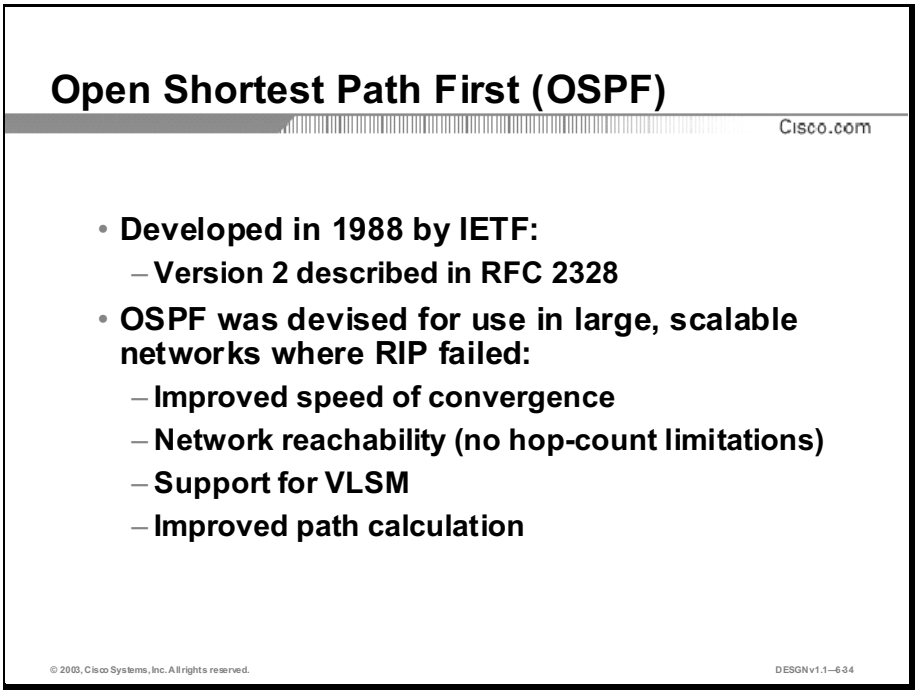

OSPF was developed for large, scalable networks where RIP failed to satisfy requirements because of its inherent limitations. OSPF is superior to RIP in all aspects. It has much faster convergence; supports VLSM, manual summarization, and hierarchical structure; better calculates the metric for best path selection; and has no hop-count limitations. At its inception, OSPF supported the largest networks.

In 1998, minor changes in OSPF version 2 addressed some of the problems of version 1, while maintaining full backward compatibility.

Note: Although OSPF was developed for large networks, its implementation requires proper design and planning, which is especially important for networks with 50 or more routers.

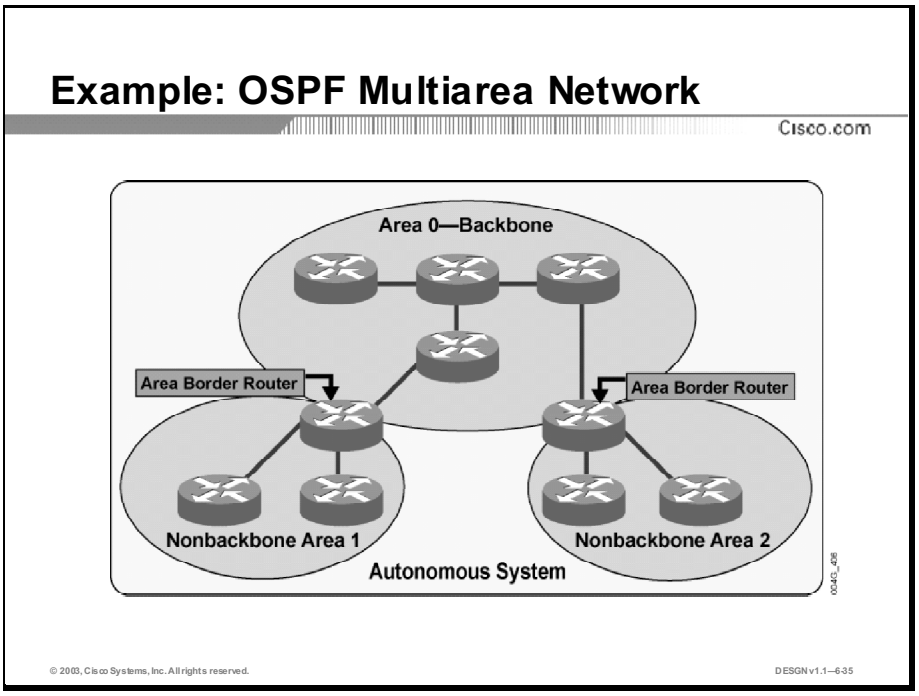

The concept of multiple separate areas inside one domain (or autonomous system) was implemented in OSPF to reduce the amount of routing traffic and to make networks more scalable. In OSPF, there must always be one backbone area, usually called Area 0, to which all other nonbackbone areas are attached. All nonbackbone areas must be directly attached to the backbone Area 0.

A router is a member of an OSPF area when at least one of its interfaces operates in that area. Routers that reside on boundaries between the backbone and a nonbackbone area are called Area Border Routers (ABRs) and have at least one interface in each area. The boundary between the areas is created in the ABR itself.

If external routes are propagated into the OSPF autonomous system, the router that redistributes those routes is called the Autonomous System Boundary Router (ASBR).

Careful design and correct mapping of areas to the network topology are important because you can perform manual summarization of routes only on ABRs and ASBRs.

When traffic is sent from one nonbackbone area to another, it crosses the backbone area. For example, in the figure, the Area 1 ABR must forward Area 1 traffic into the backbone. The Area 2 ABR receives the traffic and forwards it to the appropriate destination inside Area 2.

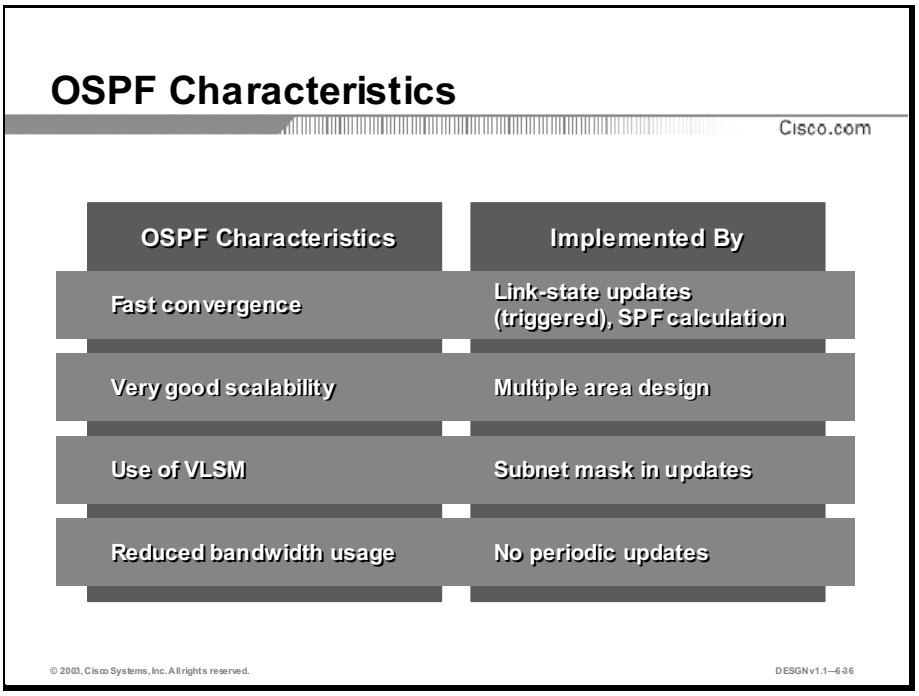

OSPF is a link-state protocol, which is ideal for the Campus Backbone submodule of a network. OSPF achieves fast convergence times by using triggered link-state updates, which include one or more link-state advertisements (LSAs). LSAs describe the state of links on specific routers and are propagated unchanged over one area. Therefore, all routers in the same area have identical topology tables, and each router has a complete view of all links and devices in the area. When LSAs cross into another area, the ABRs usually change them, depending on the type of LSA.

When the OSPF topology table is fully populated, the router applies the shortest path first (SPF) algorithm to calculate the shortest paths to the destination networks. Triggered updates and metric calculation based on the cost of a specific link assure a quick selection of the shortest path toward the destination.

Note: OSPF link cost is a value inversely proportional to the link's bandwidth.

A multiple area structure implemented in OSPF guarantees good scalability. However, strict area implementation rules require proper design so you can implement different scalability features such as manual summarization of routes on ABRs and ASBRs, stub areas, and not-so-stubby areas (NSSA). The stub and NSSA feature for nonbackbone areas decreases the amount of LSA propagation from the backbone Area 0 into nonbackbone areas. This allows low-end routers to run in the peripheral areas of the network. Fewer LSAs mean smaller OSPF topology tables, less OSPF memory usage, and lower CPU usage in stub area routers.

OSPF supports the use of VLSM and achieves better use of IP address space. Manual summarization limits the volume of link-state update propagation. OSPF, unlike Integrated IS-IS, works well over dial-up connections by suppressing the hello protocol. This mode of OSPF operation is called the OSPF Demand Circuit.

Fast convergence and good scalability make OSPF an excellent choice for the Campus Backbone routing protocol.

Integrated IS-IS for Large Networks

Integrated IS-IS is a link-state protocol that supports both IPv4 and OSI Connectionless Network Protocol (CLNP). Integrated IS-IS is a good choice for implementation in large networks. This topic discusses the characteristics of Integrated IS-IS.

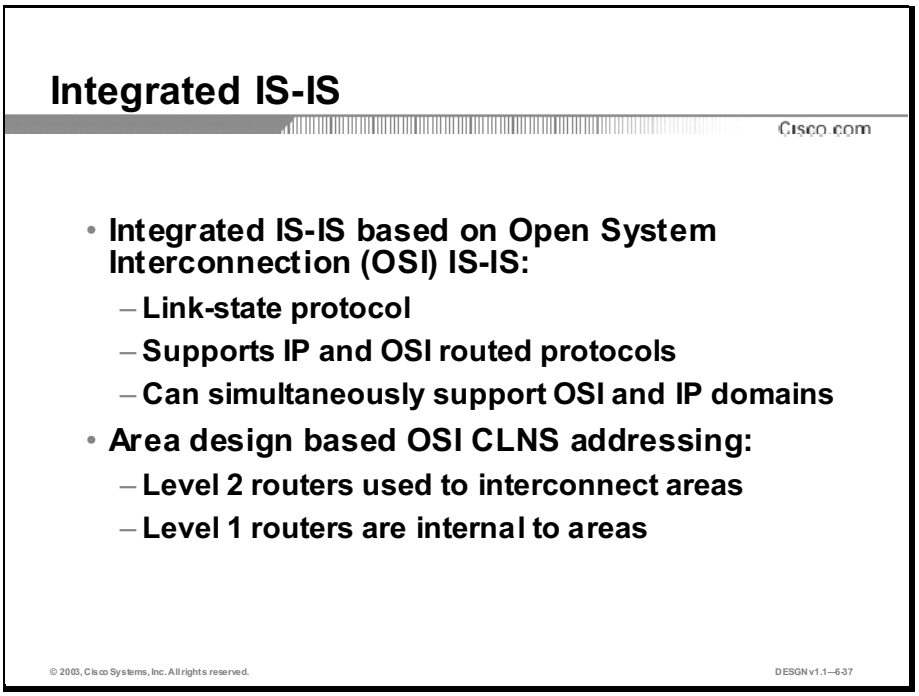

IS-IS is the primary OSI routing protocol for the Connectionless Network Services (CLNS), which used CLNP. CLNP is the OSI equivalent of IP. Because of worldwide recognition of IP as the main network layer protocol, IS-IS was adapted to the IP environment. The IPv4-capable version of IS-IS is called Integrated IS-IS, which simultaneously supports both CLNP and IP protocols.

IS-IS design is based on a Level 2 backbone to which multiple Level 1 areas are attached. The Level 2 backbone consists of Level 2 and Level 1/Level 2 routers. Level 1/Level 2 routers serve as intermediaries between the Level 2 backbone and Level 1 routers inside areas. Unlike with OSPF, with Integrated IS-IS you can easily expand the backbone area.

When designing IS-IS areas, you must assign OSI addresses to areas. This is proving to be a major disadvantage when implementing Integrated IS-IS, because OSI knowledge is not widespread in the networking community.

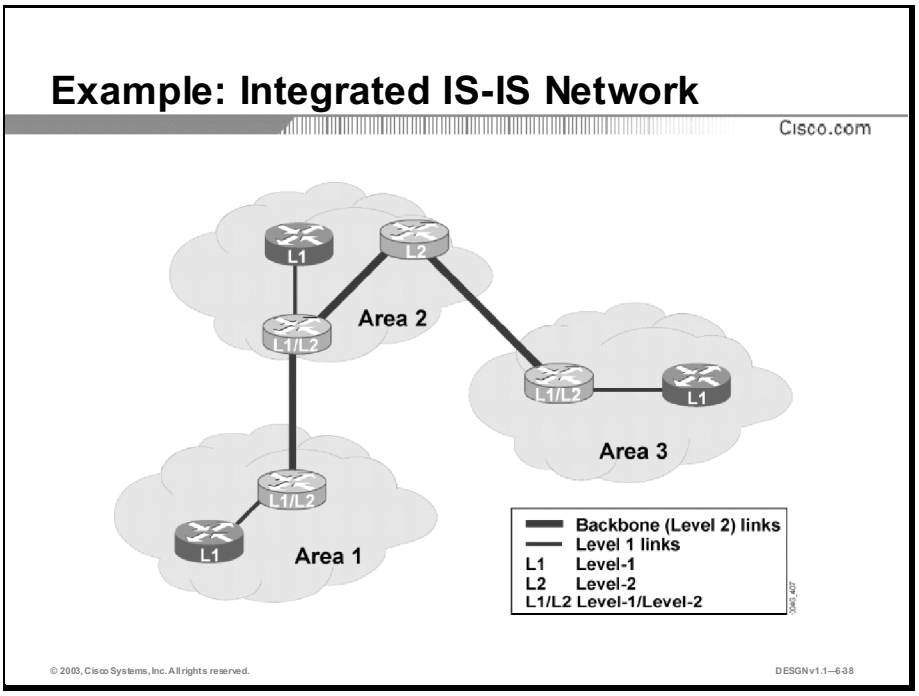

Use OSI addresses to accomplish area addressing in Integrated IS-IS. To design Integrated IS-IS, you must fully understand the principles behind OSI addressing.

In Integrated IS-IS, the backbone Level 2 and Level 1/Level 2 routers are not a part of a special backbone area, as is the case in OSPF. Level 2 routers belong to a specific Level 1 area and only form adjacencies with other Level 2 routers. An IS-IS Level 2 backbone resembles a chain of Level 2 (and Level 1/Level 2) routers, winding its way through Level 1 areas.

Level 1/Level 2 routers are used to establish adjacencies with both the Level 2 backbone routers and the Level 1 internal routers. Level 1/Level 2 routers pass internal area information into the Level 2 backbone, similarly to OSPF ABRs. Consider backbone link redundancy because the Level 2 backbone must remain contiguous at all times.

In the figure, the backbone routers are gray and nonbackbone or Level 1 internal routers are blue. Level 1/Level 2 routers form Level 1 adjacencies with Level 1 routers and Level 2 or backbone adjacencies with Level 2 and Level 1/Level 2 routers. When a backbone router has no Level 1 neighbors, you can only configure it as a Level 2 router.

Changing the type of specific Level 1 routers into Level 1/Level 2 or even Level 2 can easily expand the Integrated IS-IS backbone. In OSPF, you must renumber whole areas to expand the backbone.

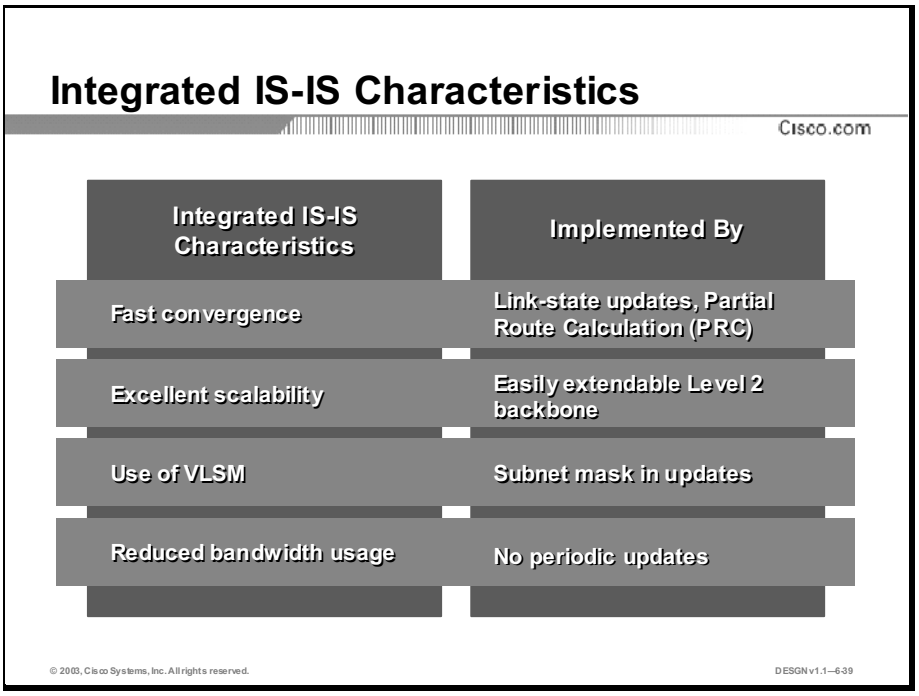

Integrated IS-IS is a proven protocol for very large networks. Live Integrated IS-IS networks with thousands of routers exist, because of the excellent scalability and convergence characteristics of Integrated IS-IS.

Like all link-state protocols, Integrated IS-IS supports VLSM.

Similar to OSPF, Integrated IS-IS owes its fast convergence characteristics to its link-state operation. The Partial Route Calculation (PRC) guarantees fast convergence and less CPU usage. Although Integrated IS-IS uses the same algorithm for best path calculation as OSPF, the full SPF calculation is initially done only on network start-up. When end IP subnets information changes, only PRC for the subnet in question is run on routers. This saves router resources and supports faster calculation. For each change in OSPF, the router must run a full SPF calculation.

Routing protocol scalability is an important characteristic for large networks. Integrated IS-IS delivers excellent scalability because of its hierarchical area-based network topology. Integrated IS-IS networks are more scalable and flexible than OSPF networks because the backbone area design is not as strict as OSPF and you can easily extend the backbone.

The disadvantage of Integrated IS-IS is its close association with the OSI world. Because few network administrators have an adequate knowledge of OSI addressing and operation, implementation of Integrated IS-IS may be difficult. Integrated IS-IS offers inherent support only for LAN and point-to-point environments, while NBMA point-to-multipoint environment support is not included. In NBMA environments, you must establish point-to-point links (subinterfaces) for correct Integrated IS-IS operation.

Border Gateway Protocol

BGP is representative of EGPs. It is primarily used to interconnect autonomous systems. Because the original exterior gateway protocol EGP is obsolete, BGP is the only EGP in current use. This topic describes the features of BGP.

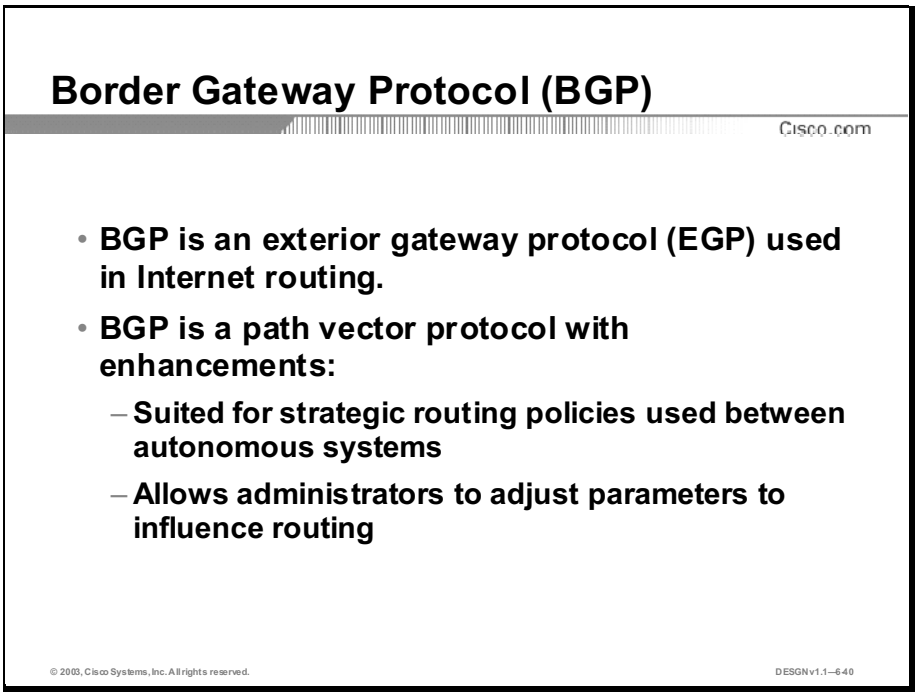

RFC 1771 (http://www.ietf.org/rfc/rfc1771.txt) defines a BGPv4 autonomous system, as "a set of routers under a single technical administration, using an IGP and common metrics to route packets within the autonomous system, and using an EGP to route packets to other autonomous systems."

In its core, BGP is a path vector protocol, which uses autonomous system path metrics as a basis for routing decisions. BGP has a number of additional metric parameters, called path attributes, which allow administrators to influence routing decisions in BGP. Inter-autonomous system routing involves a lot of strategic routing policy decisions for ISPs to comply with peering and other types of agreements.

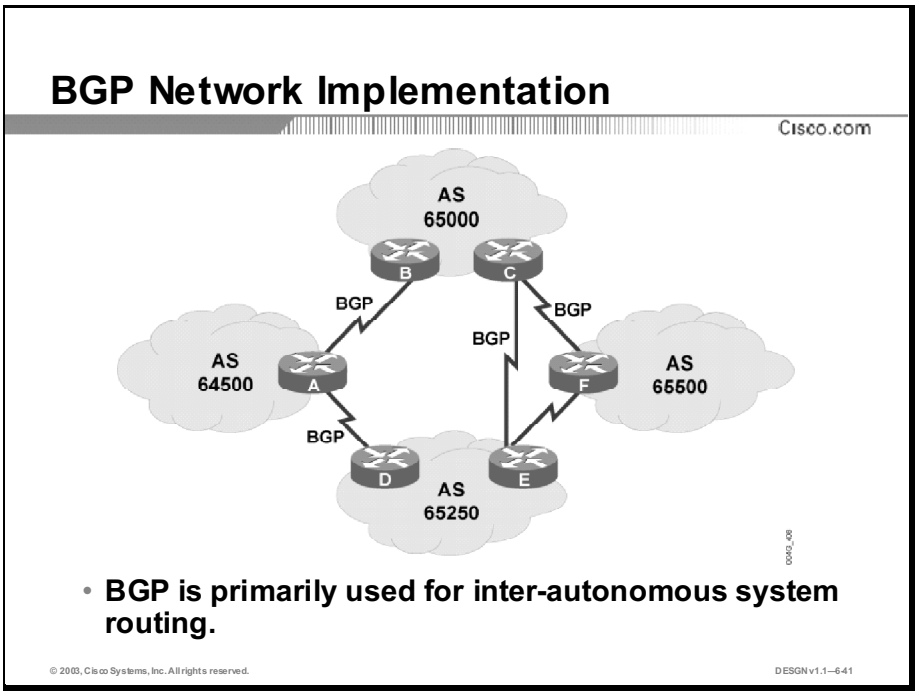

In the figure, BGP is used to interconnect multiple autonomous systems. Because of multiple connections between autonomous systems and the need for path manipulation, static routing is excluded. The redundant connectivity of one autonomous system to multiple ISP autonomous systems is called multihoming. The figure shows the autonomous system 65000 redundantly connected to three neighboring autonomous systems (65500, 65250, and 64500).

Note:	Private autonomous system numbers in the range of 64512 to 65535 are used as an example only. ISPs use public autonomous system numbers. Private autonomous system numbers are used only for non-ISPs or in special cases.

Use BGP for inter-autonomous system routing in these situations:

- If an autonomous system has multiple connections to other autonomous systems
- If an autonomous system is a transit autonomous system, meaning that it allows packets from other autonomous systems to transit themselves to reach another autonomous system (normal mode of operation for ISPs)
- If the traffic flowing to or from the autonomous system must be manipulated

The use of static routes is recommended for inter-autonomous system routing if none of these requirements exist.

Note:	BGP implementation requires considerable knowledge. Improper implementations can cause great damage, especially when complete BGP Internet tables are exchanged between neighbors (more than 100,000 routes).

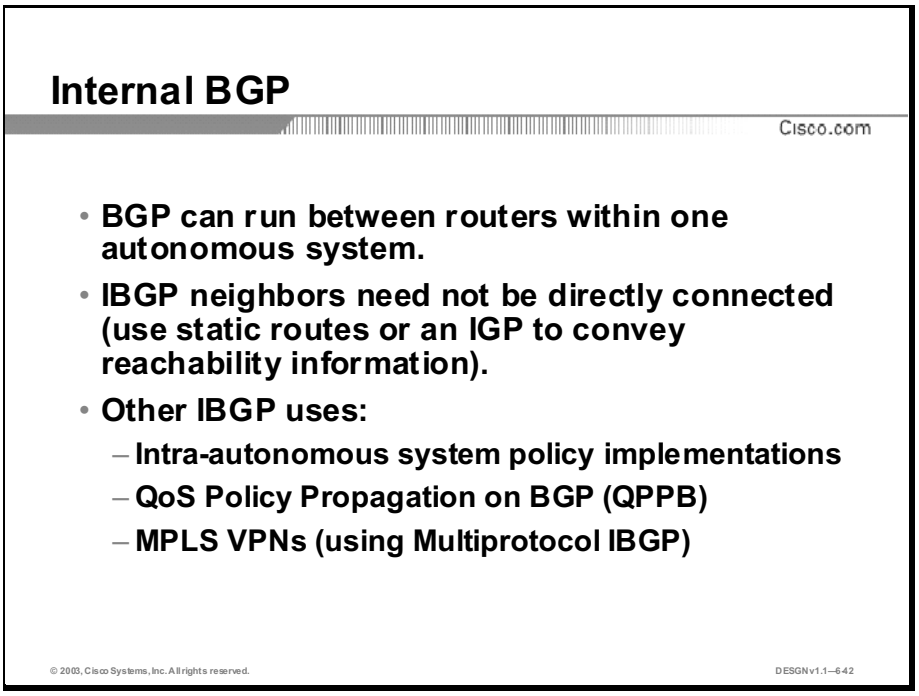

BGP that runs between routers of different autonomous systems is called External BGP (EBGP). BGP that runs between routers of the same autonomous system is called Internal BGP (IBGP).

IBGP runs on all routers or on specific routers inside the autonomous system. IBGP neighbors need not be directly connected, as long as they know how to reach each other. Neighbor reachability information may be acquired either by using the IGP running in the autonomous system or by using configured static routes.

IBGP is usually not the only protocol running in the autonomous system; it is used merely to avoid redistribution of a whole Internet table into an IGP. The EBGP information is in transit autonomous systems, where, for routing purposes, all Internet routes must be known on internal routers, transferred into IBGP, and returned into EBGP. In this way, the redistribution into IGP (which could throttle IGP routing because of the amount of routing data) is successfully avoided.

The primary use for IBGP is to carry EBGP (inter-autonomous system) routes through the autonomous system because the EBGP tables are too large for an IGP to handle. Even if an EBGP has a small table, you should prevent the loss of external routes triggering extensive computations in the autonomous system IGP.

These are other useful implementations of IBGP:

- Applying policy-based routing in internal autonomous systems with the help of BGP path attributes.
- QoS Policy Propagation on BGP (QPPB), which uses IBGP to spread common QoS parameters (for example, type of service) from one router to other routers in the network and results in a synchronized QoS policy.
- In Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs) the multiprotocol version of BGP is used to carry MPLS VPN information. In that way, the successful implementation of a VPN with MPLS is possible.

Routing Protocols for Specific Network Types

The selection of a routing protocol is based on the design goals and the physical topology of the network. This topic provides guidelines for choosing the most appropriate routing protocol for different networks.

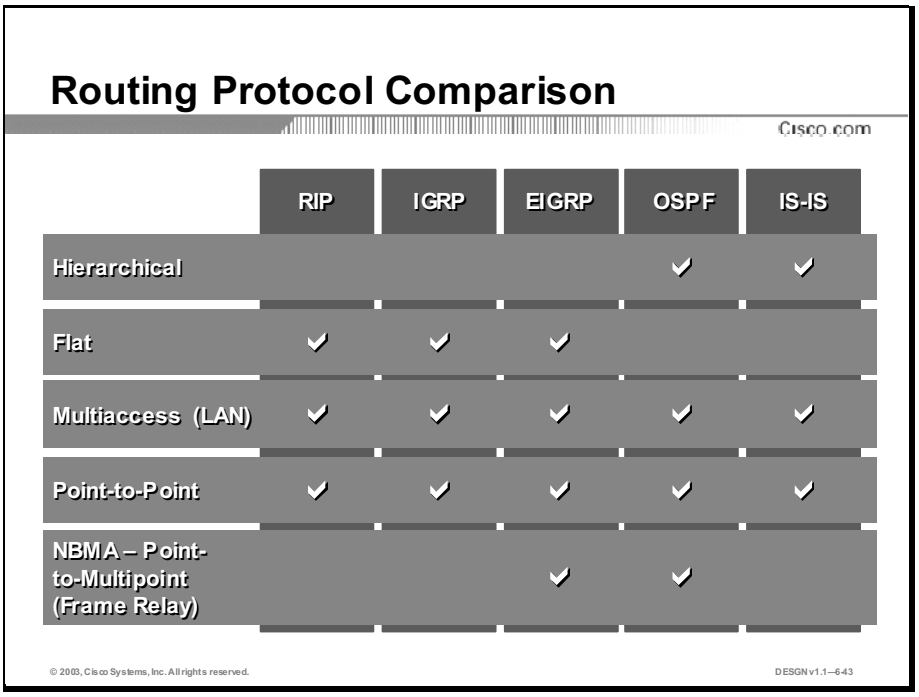

The figure compares the routing protocols to help you select the best protocol for a given situation.

When to Choose RIP or RIPv2

RIP is the oldest routing protocol and is simple in its operation. It is a classful distance vector protocol, its metric is based only on hop count, and it does not support VLSM and manual route summarization. RIP is not very common in modern networks and mainly encountered with hosts running it for router discovery. RIPv2 is an enhanced version of RIP that supports VLSM and flash updates. RIPv2 is implemented mainly in small networks, especially small hub-and-spoke networks using point-to-point links. RIPv2 is used in dial-up networks because it is able to freeze its routing table and wait for the dial-up link to connect to and start the exchange of routing information (sometimes called the snapshot routing feature). It is seldom used in LAN environments, because it has no notion of its neighbors and cannot detect a failure in neighboring routers to provide for fast convergence.

In Nonbroadcast Multiaccess (NBMA) environments, the main issue of RIPv2 is associated with the split-horizon rule, which prevents the propagation of routing updates to all connected routers reachable through the same physical interface but over different virtual circuits. Use of RIP and RIPv2 in NBMA networks is not appropriate. The workaround for the NBMA split-horizon problem is to use logical point-to-point subinterfaces, changing the NBMA network into a logical collection of point-to-point links.

When to Choose IGRP

IGRP is the original Cisco routing protocol. It is a classful distance vector protocol with a more complex metric calculation than RIP. It takes into account minimum bandwidth and accumulated delay. IGRP is suitable for small to medium networks, but has problems with the split-horizon feature in NBMA networks. Another problem of IGRP is its slow convergence because of its pure distance vector operation. IGRP has been replaced in most networks by the Enhanced IGRP (EIGRP). Only use IGRP in existing networks that have not yet migrated to EIGRP, OSPF, IS-IS, or RIPv2.

When to Choose EIGRP

EIGRP, which is based on IGRP, is a hybrid protocol that incorporates the best aspects of distance vector and link-state features (topology table, no periodic route propagation, and triggered updates). It is well suited to almost all environments, including LAN, point-to-point, and NBMA. In NBMA, you can disable the split-horizon functionality for EIGRP.

EIGRP is not suitable for dial-up environments because it must maintain the neighbor relationship and it uses periodic hello packets, effectively maintaining the dial-up connections all the time.

Note: EIGRP is a Cisco proprietary protocol licensed to limited vendors.

When to Choose OSPF

OSPF is a standards-based link-state protocol, based on the SPF or Dijkstra's algorithm for best path calculation. Initially, it was designed for networks that consist of point-to-point links, but later it was successfully adapted for operation in LAN and NBMA environments. OSPF is tuned for dial-up operation by suppressing the hello protocol over OSPF dial-up lines (sometimes called Demand Circuit operation). Because of the hierarchical design requirement, there are design considerations when using OSPF in larger networks. One backbone area is required, and you must attach all nonbackbone areas directly to that backbone area. Expansion of the backbone area can cause design issues, because the backbone area must remain contiguous.

When to Choose Integrated IS-IS

Integrated IS-IS is a standards-based link-state protocol similar in operation to OSPF. It uses the SPF algorithm for best path calculation. An IS-IS network consists of two areas: a backbone (Level 2 router) and connected nonbackbone (Level 1 router). In contrast to OSPF, you can easily expand the IS-IS backbone to accommodate new Level 1 areas. Integrated IS-IS is a proven protocol for very large networks. Integrated IS-IS has no adaptation for NBMA point-to-multipoint networks, which is one design consideration prior to implementing. Integrated IS-IS is not suited for dial-up networks because, unlike OSPF, it includes no hello protocol suppression capability. The deployment of Integrated IS-IS in networks requires more knowledge than other IGPs. Integrated IS-IS is based on the Open System Interconnection (OSI) IS-IS protocol, and the numbering of IS-IS areas is done in an OSI-based environment, not in IP.

Example: Routing Protocol Decision Table

Cisco.com

Parameters \ Options	RIPv2	IGRP	EIGRP	OSPF	IS-IS	Required Network Parameters
Size of Network (Small-Medium-Large-Very Large)	Medium ✗	Medium ✗	Large ✓	Large ✓	Very Large ✓	Large
Speed of Convergence (Very High-High-Low)	Medium ✗	Low ✗	Very High ✓	High ✓	High ✓	High
Use of VLSM (Yes-No)	Yes ✓	No ✗	Yes ✓	Yes ✓	Yes ✓	Yes
Mixed Vendor Devices (Yes-No)	Yes ✓	No ✗	No ✗	Yes ✓	Yes ✓	Yes
Network Support Staff Knowledge (Good-Poor)	Good ✓	Good ✓	Good ✓	Good ✓	Poor ✗	Good
	✗	✗	✗	✓	✗	

The figure displays an example of a decision table for selecting a routing protocol, based on multiple criteria. In this example, several routing protocols are considered as possible options (OSPF, IS-IS, IGRP, EIGRP, RIPv2), and their compliance to five different required parameters is tested.

The protocol you choose should have these properties:

- It should support a large network (up to 100 or more routers). RIPv2 and IGRP protocols do not meet this requirement.
- It must have a high speed of convergence. This precludes RIPv2 and IGRP.
- The use of VLSM is required. IGRP does not support VLSM.
- It should support Cisco Systems and other vendors' equipment. EIGRP and IGRP are Cisco proprietary protocols. Mixed vendor environments generally do not support them.
- Network staff should have a good knowledge of the chosen protocol, enabling them to troubleshoot the network. Most of the network administrators in this example have only a basic knowledge of IS-IS, because it is not a widely used protocol.

OSPF is the most suitable option because it satisfies all the given requirements. IS-IS and EIGRP are the next closest matches, failing on only one item, followed by RIPv2 and IGRP.

Protocol Selection Job Aid

You can use the table as a job aid during your assignments. Two additional rows are included so you can specify more parameters (for example, types of existing physical topologies), which might be of importance in your network.

Parameters (Options)	RIP v2	IGRP	EIGRP	OSPF	IS-IS	Required Network Parameters
Size of Network (Small-Medium-Large-Very Large)	Medium	Medium	Large	Large	Very Large	
Speed of Convergence (Very High-High-Low)	Medium	Low	Very High	High	High	
Use of VLSM (Yes-No)	Yes	No	Yes	Yes	Yes	
Mixed Vendor Devices (Yes-No)	Yes	No	No	Yes	Yes	
Network Support Staff Knowledge (Good-Poor)						

Summary

This topic summarizes the key points discussed in this lesson.

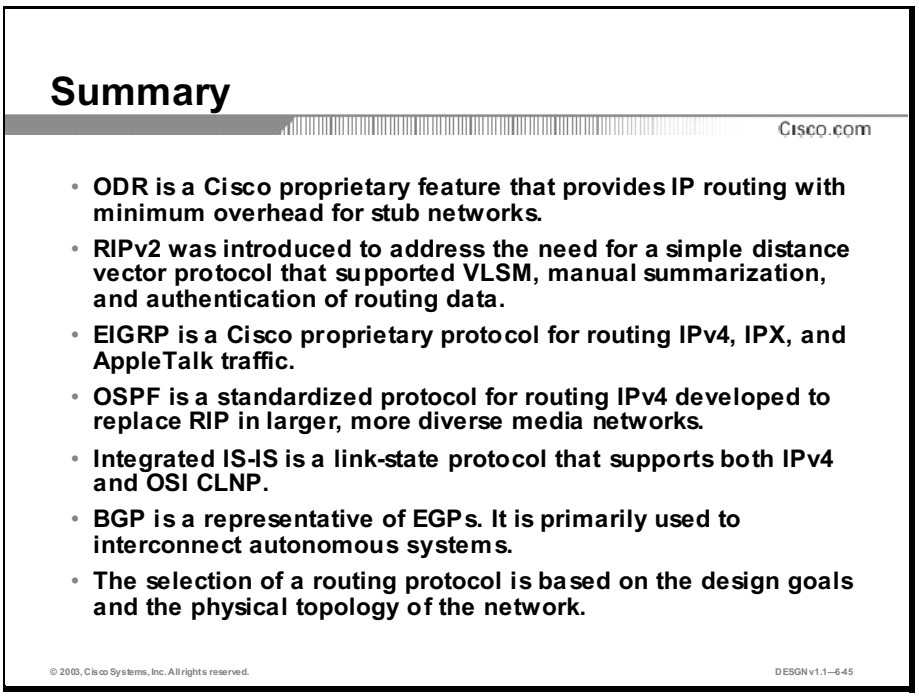

References

For additional information, refer to these resources:

- *Interconnecting Cisco Network Devices* (ICND) course
- *Building Scalable Cisco Internetworks* (BSCI) course

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which network topology is suitable for ODR implementation?

A) fully meshed NBMA network

B) Ethernet LAN

C) hub and spoke

D) Token Ring LAN

Q2) Select two statements that best describe RIPv2. (Choose two.)

A) RIPv2 allows authentication of packets.

B) RIPv2 is a link-state routing protocol.

C) RIPv2 does not carry subnet masks in its updates.

D) RIPv2 supports manual summarization.

E) RIPv2 does not send its routing table periodically via multicast.

Q3) Select two EIGRP features that make it an appropriate choice for a routing protocol. (Choose two.)

A) very fast convergence using DUAL

B) multivendor support

C) variable length subnet masking (VLSM) support

D) very slow convergence

E) area-based design

Q4) Select two statements that provide the most correct descriptions of OSPF. (Choose two.)

A) OSPF uses the SPF or Dijkstra's algorithm to calculate best paths.

B) Summarization of routes is required on all OSPF routers.

C) Area-based design is necessary for larger OSPF networks.

D) OSPF is an advanced distance vector protocol.

E) OSPF is used only on point-to-point links.

Q5) Name two advantages of IS-IS over OSPF for use in large networks. (Choose two.)

A) enhanced scalability (easier extending of the backbone)

B) support of OSI addressing

C) use of Level 1 routers

D) use of Partial Route Calculation (PRC)

E) widespread knowledge of IS-IS among administrators

Q6) Name two main reasons why BGP is used for inter-autonomous system connections instead of IGPs. (Choose two.)

A) BGP can handle much larger amounts of routing data.

B) IGPs have slower convergence times.

C) BGP allows administrators to influence traffic flow in detail.

D) BGP is better at calculation of best paths.

Q7) A large organization has decided to connect all their regional offices with their branch offices. Each regional office has a minimum of two, and a maximum of five branch offices to connect. The branch offices use low-end routers that are directly connected to the regional office router via Frame Relay permanent virtual circuit (PVC) links, effectively creating a hub-and-spoke topology (star network). No physical connections exist between the branch office routers. The protocol that runs between the regional and branch offices does not need to be the same as in the rest of the network where OSPF is run. Select the two better options for establishing IP connectivity. (Choose two.)

A) Deploy EIGRP.

B) Deploy IS-IS.

C) Deploy RIPv2 with default route for branch office connectivity to the rest of the network.

D) Use static routing with a default static route from branch to regional offices, and use static routes on regional routers toward branch networks.

Quiz Answer Key

Q1) C

Relates to: On-Demand Routing

Q2) A, D

Relates to: Routing Information Protocol Version 2

Q3) A, C

Relates to: Enhanced IGRP

Q4) A, C

Relates to: Open Shortest Path First

Q5) A, D

Relates to: Integrated IS-IS for Large Networks

Q6) A, C

Relates to: Border Gateway Protocol

Q7) C, D

Relates to: Routing Protocols for Specific Network Types

Designing a Routing Protocol Deployment

Overview

The selection and implementation of a routing protocol is based on specific needs and topologies. This lesson describes scenarios where routing protocols are deployed based on different needs. The network characteristics and customer requirements, as well as the multiprotocol multivendor environment, are considered. The lesson discusses route redistribution, summarization, and filtering, and provides examples of integrating internal networks with external domains.

Relevance

This lesson will help you apply routing protocol concepts to specific scenarios that you are likely to encounter.

Objectives

Upon completing this lesson, you will be able to deploy routing protocols in a hierarchical enterprise network. This includes being able to meet these objectives:

- Select the routing protocol for a specific hierarchical level of a network
- Design the route redistribution points in a network
- Identify opportunities for route filtering in a network
- Identify network locations for route summarization
- Explain the integration of interior routing protocols with external domains using BGP

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Solid understanding of routing protocol concepts and features of main routing protocols

Outline

The outline lists the topics included in this lesson.

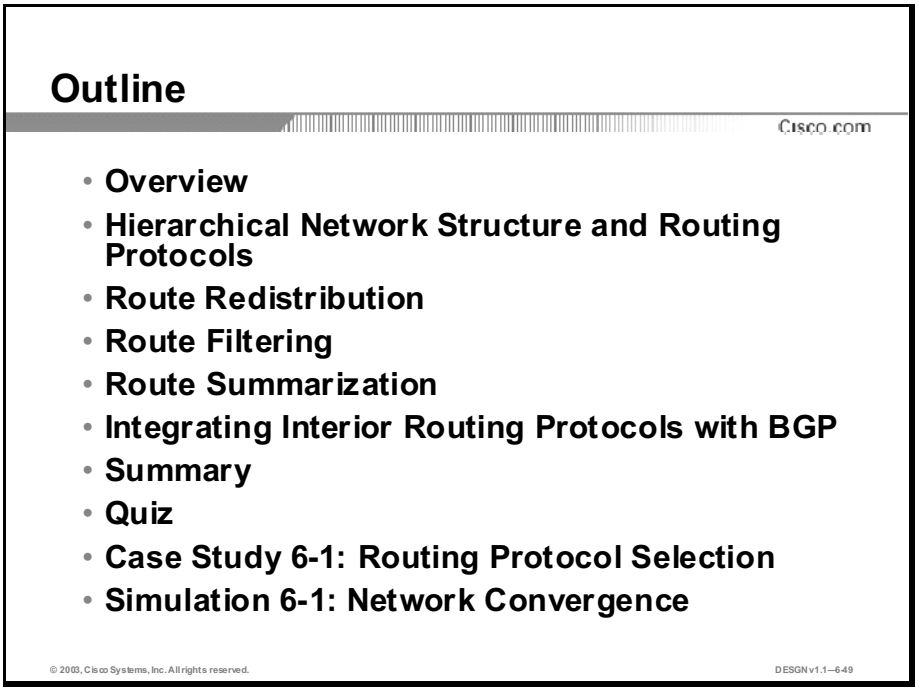

Hierarchical Network Structure and Routing Protocols

The choice of routing protocols is based on the network design goals. In large networks, you will likely select multiple protocols for different levels of the Enterprise Composite Network Model. This topic describes how routing protocols apply to the components of the enterprise network.

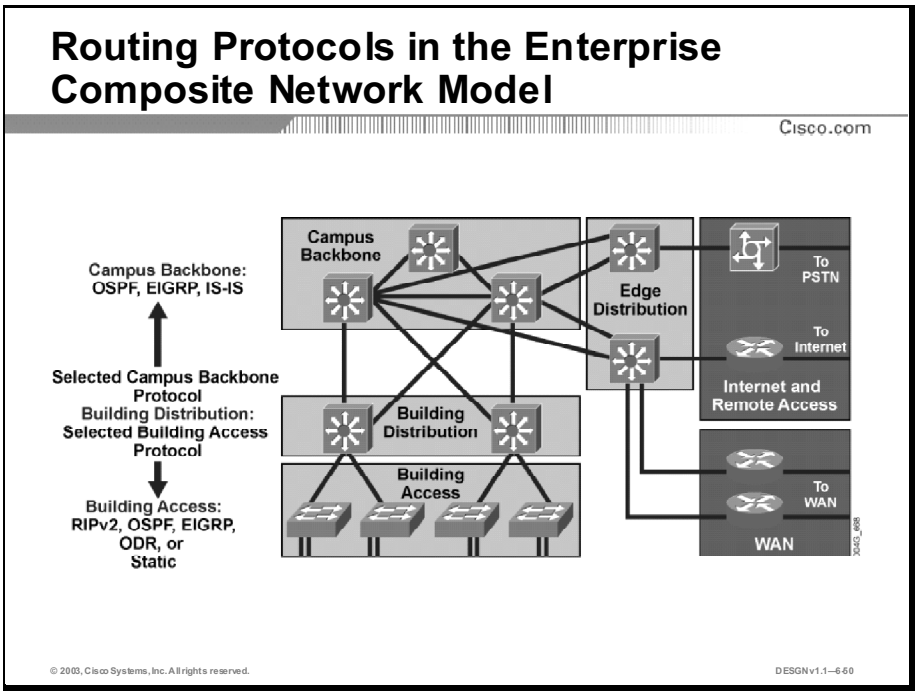

Routing in the Campus Backbone

The Campus Backbone provides high-speed data transmission between Building Distribution devices. The Campus Backbone is critical for connectivity and incorporates a high level of redundancy by using redundant links and load sharing between equal-cost paths. It must provide immediate response in the event of a link failure and must adapt very quickly to change.

The Campus Backbone must converge and adapt to changes quickly to provide a seamless transport service. EIGRP, OSPF, and IS-IS all adapt to changes quickly and have short convergence times. Make the decision to use EIGRP, OSPF, or IS-IS based on the underlying physical topology, IP addressing, equipment used, and possible issues related to the routing protocol in a particular situation.

These are the disadvantages of each routing protocol in the Campus Backbone:

- OSPF imposes a strict hierarchical design. OSPF areas must map to the IP addressing plan, which may be impossible.
- EIGRP restricts vendor selection because it is a Cisco proprietary protocol. To overcome this restriction, you can use multiple routing protocols with redistribution.
- IS-IS requires detailed knowledge for proper configuration.

- RIP is not recommended as a Campus Backbone routing protocol. Its convergence and response to change is slow and may result in disrupted connectivity, thus denying the basic Campus Backbone function. Because the RIP metric is based on hop count, it is unsuitable when diverse media is used and the number of hops is high.
- Using static routing in the Campus Backbone is not an option, because static routing requires administrative intervention for changes and link failures.

Routing in the Building Distribution Submodule

The Building Distribution submodule is the intermediate point between the Campus Backbone and Building Access submodules. The physical media, IP addressing, and the choice of routing protocols used in the Campus Backbone and Building Access submodules affect the routing protocol choice in the Building Distribution submodule. Routing protocols used in the Building Distribution submodule include EIGRP, OSPF, IS-IS, RIP, and ODR.

For example, if EIGRP is the Campus Backbone routing protocol and RIP is the Building Access routing protocol, then use both routing protocols on the Building Distribution devices and apply redistribution with filtering to provide connectivity.

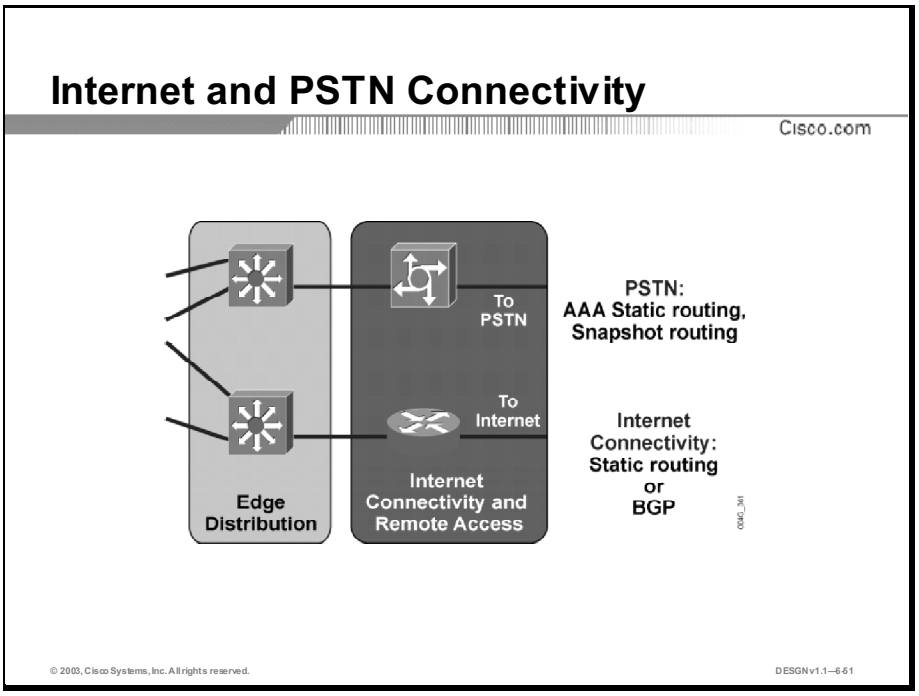

Routing in the Edge Distribution Module

The Edge Distribution module provides access to network resources for local and remote users. The underlying physical topology, IP addressing, and the deployed equipment drive the choice of routing protocol. If the Edge Distribution equipment is less powerful than the Building Distribution and Campus Backbone equipment, consider the available processing power and memory. The routing protocols in the Edge Distribution module are RIPv2, OSPF, EIGRP, ODR, and static routing. Routing protocols running in the Edge Distribution module are referred to as edge routing protocols.

These are the advantages and disadvantages of each routing protocol in the Edge Distribution module:

- IS-IS is not suitable for the Edge Distribution module because it demands a significant amount of knowledge to configure and is unsuitable for dial-up networks.
- RIPv2 is simple in its operation and suitable for small hub-and-spoke networks with point-to-point links as well as for dial-up networks. RIPv2 does not have a high demand for memory and processing power.
- EIGRP offers the administrator more influence on routing and is suitable for NBMA environments, where there is a split-horizon issue (for example, Frame Relay or ATM multipoint interfaces). When equipment from multiple vendors is part of the overall design, the use of EIGRP is restricted. EIGRP is unsuitable for dial-up environments.
- The limitations of using OSPF as an Edge Distribution module routing protocol are connected to its high memory and processing power requirements and strict hierarchical design. Use OSPF in environments such as LAN, NBMA, and dial-up. OSPF also demands significant configuration expertise.

Note: The high memory and processing power requirements can be surpassed with the use of summarization and careful area planning.

Remote Access and Internet Connectivity

The Remote Access and VPN module is used to provide connectivity to corporate networks for remote users via dial-up connections. In a dial-up environment, you can use AAA for static routing, when dynamic routing is not needed. If dynamic routing is required, use a snapshot-capable protocol such as RIP to implement dynamic routing and to minimize the costs related to the routing protocol.

Snapshot routing eliminates the constraints of static routes by deploying a hub-and-spoke topology in a DDR environment. Dynamic routing protocols will update the routing tables during active windows, and snapshot will prevent them from aging for a configurable period of time (Quiet Period).

For the Internet Connectivity module, use either static routes or BGP. Base the decision on what to use on whether multiple exit points exist and on redundancy requirements. Static routes present less overhead than BGP routing and are used when only one exit point exists. However, use BGP with multiple exit points and when multihoming is desired.

Route Redistribution

Advanced routing features such as redistribution, filtering, and summarization allow multiple routing protocols to coexist and provide greater scalability. This topic explains route redistribution and redistribution points.

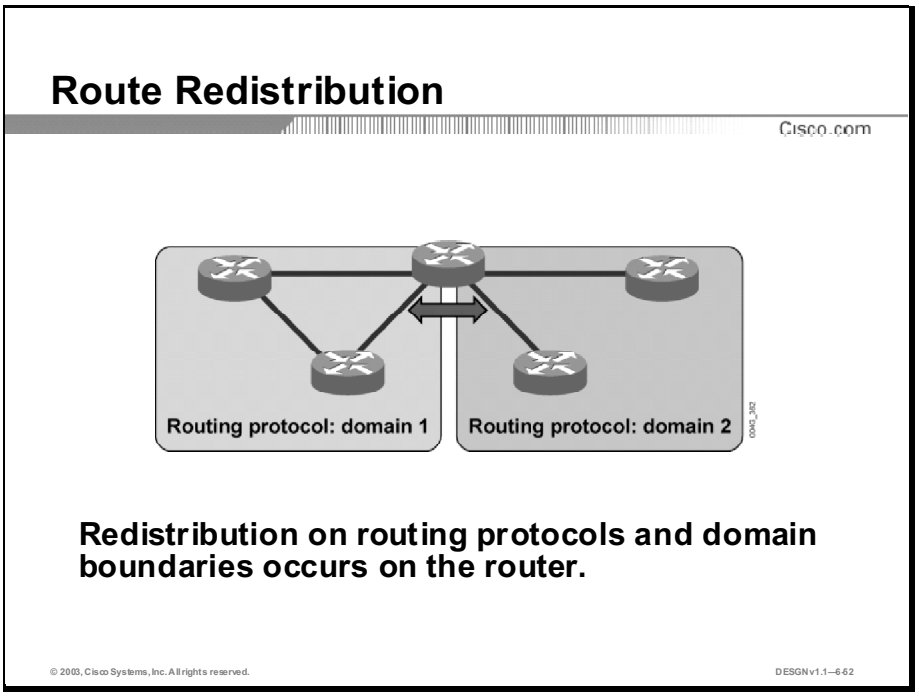

Redistribution between different routing protocols refers to passing routing knowledge from one protocol to another.

There are several scenarios when multiple protocols are used in a single network:

- When migrating from an older IGP to a new IGP, multiple redistribution points may exist until the new protocol has displaced the old one entirely (in the whole network).
- Different departments might not want to upgrade their routers, or they might not implement a sufficiently strict filtering policy. In these cases, redistribution between those protocols is needed.
- In a mixed vendor environment, you can use a Cisco-specific protocol in the Cisco portion of the network, and then use redistribution with common protocol to communicate with non-Cisco devices.

Redistribution occurs on routing protocols and domain boundaries on a router with interfaces that participate in multiple routing protocols and routing domains.

Redistribution of routes is needed in these situations:

- Multiple routing protocols are used in the network (for example, RIPv2, EIGRP, OSPF)
- Multiple routing domains are used in the network (for example, two EIGRP routing processes)

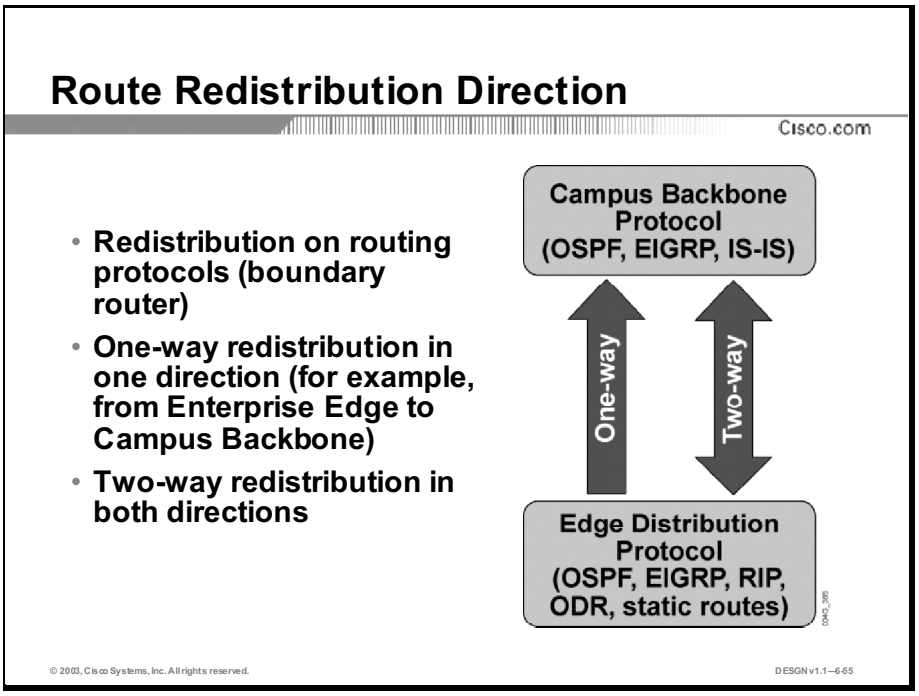

Redistribution is usually applied between the Campus Backbone and Edge Distribution protocols. Redistribution is possible in two ways:

- **One-way route redistribution:** Routing information is redistributed only from one routing protocol or domain to another but not vice versa. When this occurs, static or default routes are required in the opposite direction to provide connectivity.
- **Two-way redistribution:** Routing information is redistributed from one routing protocol or domain to another and vice versa. Static or default routes are not needed because all routing information is passed between two entities.

Route Redistribution Planning

When deciding where and how to use route redistribution, determine these three things:

- Routing protocols and domains to be used in the network
- Routing protocols and domain boundaries (boundary routers)
- Directions of route redistribution (one-way, two-way redistribution)

If you do not carefully design route redistribution, you can introduce suboptimal routing and routing loops into the network. This occurs when routes are redistributed back into a network that has redundant paths between dissimilar routing protocols or domains. The solution to this problem is achieved through route filtering.

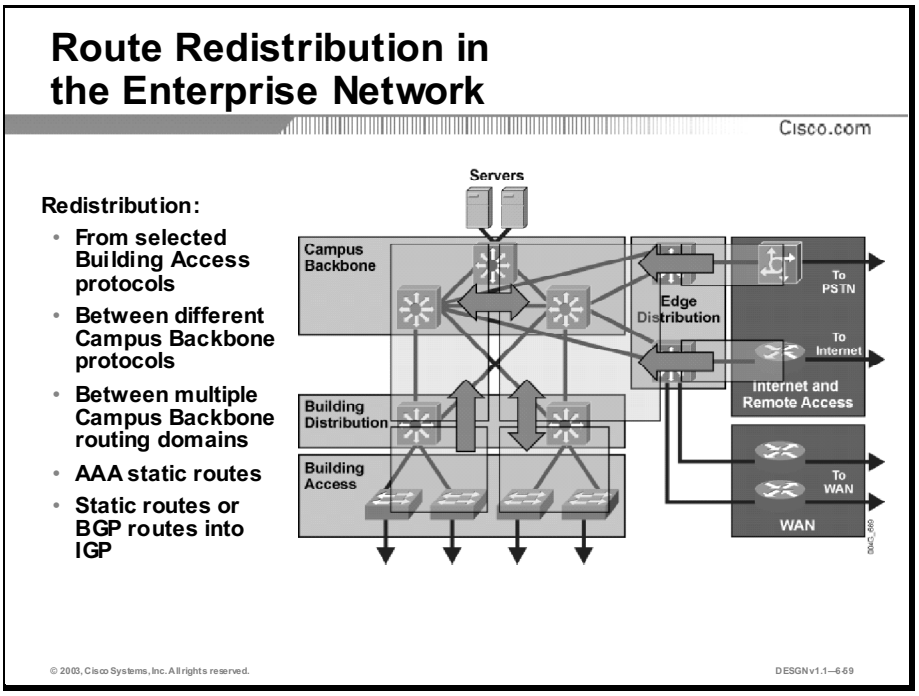

Redistribution occurs on routing protocols and domain boundaries. Therefore, redistribution is usually needed in the Building Distribution submodule when multiple routing protocols and domains collide.

The figure shows redistribution points in an enterprise network. The enterprise Campus Backbone uses OSPF as a routing protocol and RIP and EIGRP to provide dynamic routing toward remote sites. Some remote sites need connectivity only to the Server Farm module so you can perform one-way redistribution to inject routes from remote sites into the enterprise Campus Backbone. You can configure RIP to propagate only the default route down to the Building Access switch. The Building Access switch advertises its own LAN to the Building Distribution switch, greatly reducing the RIP update volume.

The other demand is connectivity to certain remote sites. Therefore, a two-way redistribution is required to provide connectivity.

The redistribution in the Campus Backbone is necessary when multiple routing protocols or multiple routing domains are used. For example, two OSPF processes run in the Campus Backbone. Redistribution between those processes is necessary to provide connectivity.

Part of the Campus Backbone could run EIGRP and part could run OSPF. In this instance, a two-way redistribution is required to provide connectivity.

Remote Access and Internet Connectivity Route Redistribution

Redistribution is necessary in the Remote Access and VPN or Internet Connectivity modules.

For the Remote Access and VPN module with AAA static routing, redistribution takes AAA static routes and injects them into the Campus Backbone routing protocol. In the opposite direction, default routing provides connectivity for the remote users.

For the Internet Connectivity module with only one exit point, that exit point is the default route for the Internet traffic and is propagated through the Campus Backbone routing protocol. When multiple exit points exist toward multiple ISPs, use BGP to provide Internet connectivity and redistribution.

Route Filtering

Route filtering prevents advertisement of certain routes through the routing domain. This topic describes route filtering.

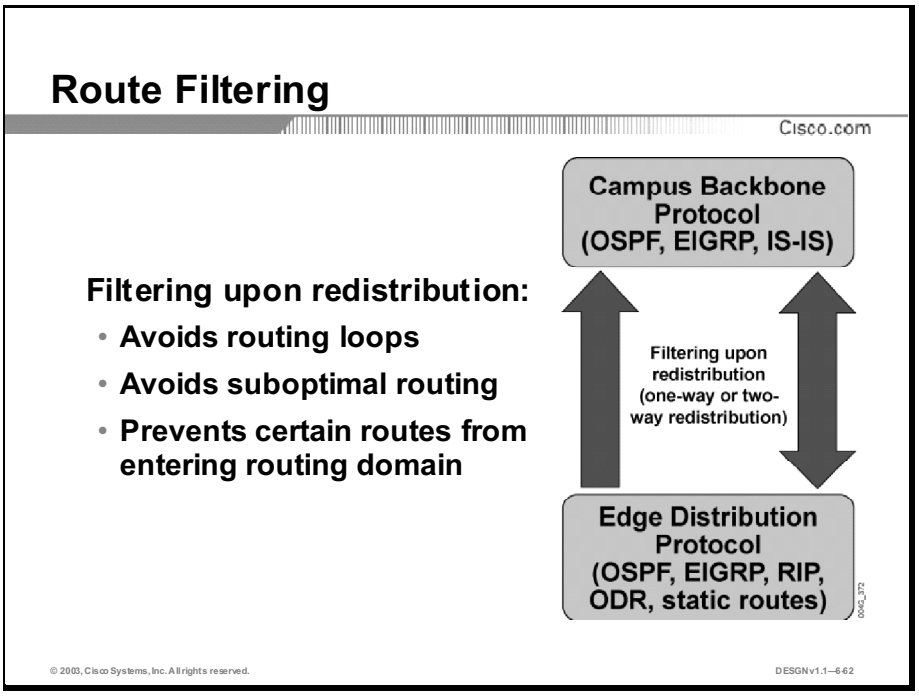

Filtering can occur either on the routing domain boundary where redistribution occurs or in the routing domain to isolate some parts of the network from other parts of the network.

Filtering is used in combination with route redistribution to prevent suboptimal routing and routing loops that might occur upon redistribution when routes are redistributed on multiple redistribution points.

Route filtering also prevents remote sites from receiving routes from a local IP address space.

Route Filtering and Internet Connectivity

In large networks, you can use BGP to provide the external connectivity, especially when the network is multihomed.

When BGP routes are exchanged with multiple ISPs, take care to prevent the network from becoming a transit network between the ISPs. Use route filtering to prevent advertisement of private addresses and addresses out of the official address scope.

Route Filtering in the Enterprise Network

Cisco.com

Route filtering occurs upon redistribution:

- **From WAN, Internet Connectivity, and Remote Access and VPN modules**
- **Between different routing protocols in the Campus Backbone**
- **Between multiple routing domains in the Campus Backbone (same protocol)**

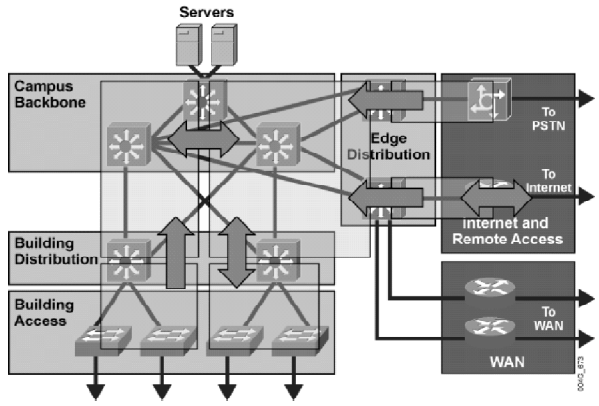

DESGN v1.1—6-66

In an enterprise network, route filtering usually occurs at redistribution points. Routes are redistributed from the Building Access submodule to the Building Distribution submodule and from the Building Distribution submodule to the Building Access submodule. Routes learned via redistribution from the Building Access submodule are filtered before redistributing them back into the Building Access submodule. Route filtering toward the Building Access submodule might block all routes except the default route, which is injected to the Building Access submodule. This filtering relieves the Building Access submodule switches from receiving too many updates and storing a large number of routes.

When redistributing between an IGP and BGP, use redistribution into the BGP domain with filtering to prevent announcement of invalid and private IP addresses. Because routers inside an enterprise network do not need data about every Internet route, apply route filtering during redistribution from the BGP domain into the IGP routing protocol.

Route Summarization

A hierarchy in the network reduces routing traffic and unnecessary route recomputation. To implement a hierarchy, you can divide a network into areas that enable route summarization. This topic describes route summarization.

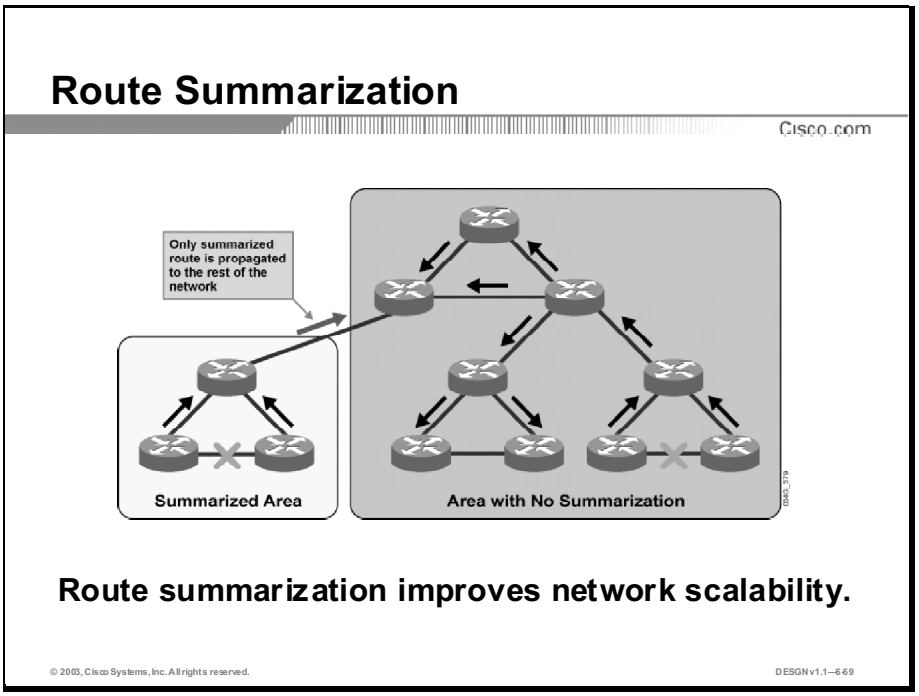

Routing traffic consumes considerable network resources, so a large flat network is not scalable. When a change occurs, it is propagated throughout the network. This requires processing time for route recomputation and for the bandwidth to propagate routing updates.

With summarization capabilities in place, a route recomputation occurring in one network area does not influence routing in other areas. Instabilities are isolated and convergence is improved, which reduces the amount of routing traffic, the size of the routing tables, and the required memory and processing power for routing.

You can implement summarization manually or automatically with routing protocols that provide such options.

Note: The underlying IP addressing plan must be modular to support route summarization.

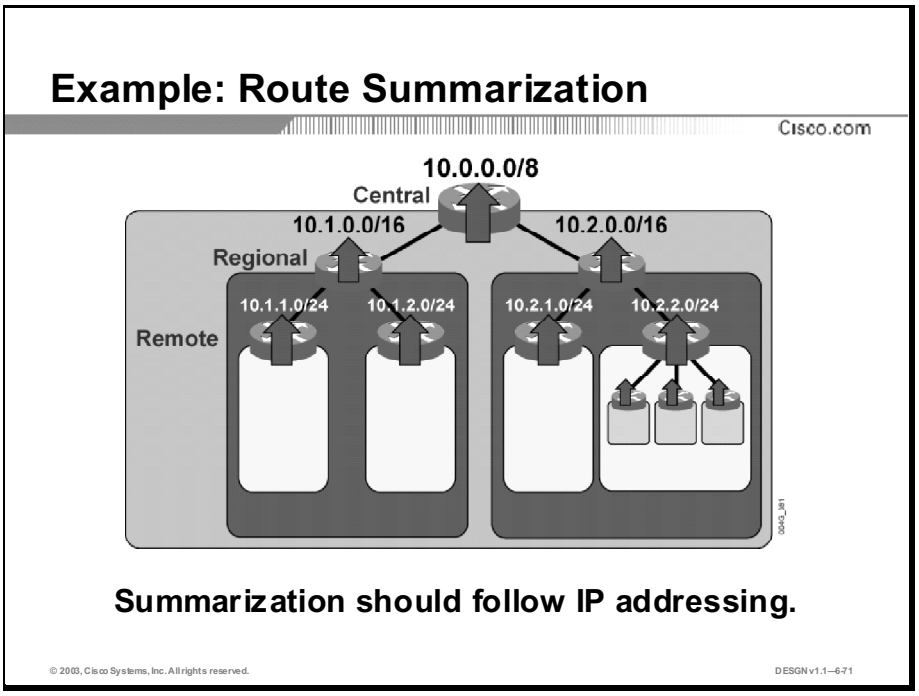

You can implement summarization in multiple levels, modeling the hierarchical (recursive) tree-like model. This allows multilayer partitioning of a network into variable size areas in a very scalable and flexible manner.

For example, the figure shows a network that consists of a central office, two regional offices, and some remote offices connected to the regional offices. The addressing plan uses 10.0.0.0/8 to address the network. The central office and regional offices are assigned /16 address blocks (that is, 10.0.0.0/16, 10.1.0.0/16, and 10.2.0.0/16). Each address block is further divided into smaller address blocks according to the needs (that is, 10.2.0.0/16 is divided into 10.2.1.0/24, 10.2.2.0/24, and 10.2.3.0/24 for LAN networks and 10.2.4.0/30, 10.2.4.4/30, and 10.2.4.8/30 for WAN connections). The summarization should follow the addressing plan and summarize networks from an individual location into a larger address block (that is, 10.2.1.0/24, 10.2.2.0/24, and 10.2.3.0/24; and 10.2.4.0/30, 10.2.4.4/30, and 10.3.4.8/30 are summarized to 10.3.0.0/16).

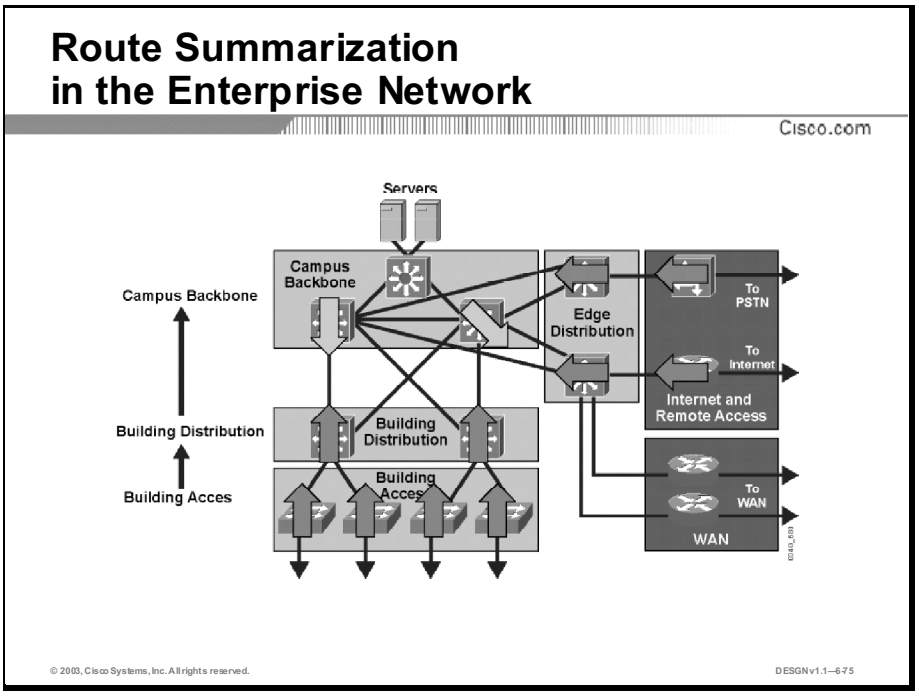

You can configure summarization in a large network in multiple levels:

- Implement the first level of summarization at WAN connectivity and remote-access points. Summarize remote networks into major networks, and only those major networks are then advertised to the Building Distribution submodule. Take special care not to advertise overlapping summary routes, which would result in disrupted connectivity.
- Implement the second level of summarization at the Building Distribution submodule in the same manner as the first level of summarization. All Building Distribution switches perform summarization of access networks on all interfaces toward the Campus Backbone.

You can implement summarization on Campus Backbone switches toward the Building Distribution submodule if the addressing in the Campus Backbone supports summarization of Campus Backbone networks. When all Building Distribution switches are connected to the Campus Backbone in a redundant way (primary or secondary link), perform summarization from the Campus Backbone toward the Building Distribution submodule carefully. The routing protocol must allow the summarization of each interface, enabling you to use different subnet masks for the summaries on the primary and backup links. When Building Distribution switches receive two summaries, the more specific route is installed in the routing table. Summaries on the primary links must use a longer subnet mask.

Note: Summarization relies on a solid network addressing design.

Integrating Interior Routing Protocols with BGP

You can implement BGP to integrate interior routing protocols. This topic explains the interaction between BGP and IGP routing protocols.

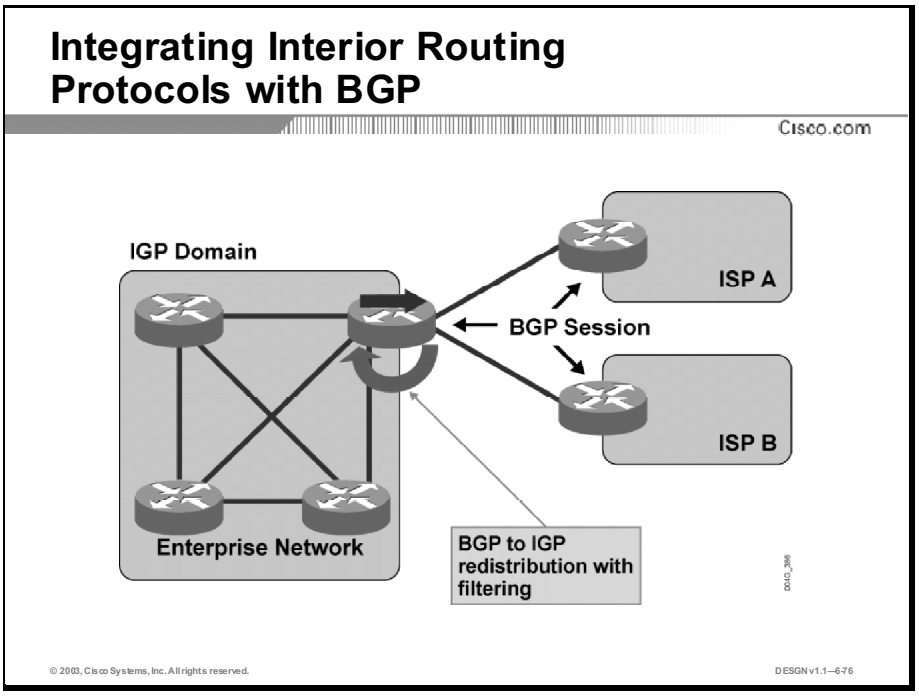

BGP on the border announces only the major network to the external domains (the prefix that was assigned to the enterprise network), without any details on subnets. You can redistribute internal networks into BGP, summarized into one major subnet (which covers the assigned public address space) and advertised to the external domains. In addition, you must implement filtering toward external domains for all private addresses and addresses that are not within the assigned address space.

In the opposite direction, implement conditional advertisement of a default candidate when one of the major networks is received via BGP updates from an external domain. On the border routers, the BGP tools select the best path to external networks. One of the major networks received via BGP is redistributed into IGP and marked as the default candidate, thus achieving effective conditional default advertising. To select between the primary and backup links when two exit points exist, assign a higher cost to the redistributed route on the secondary router.

The redistribution of all BGP routes into IGP is not recommended because non-BGP participating routers do not require full Internet routing and most IGP protocols are unable to process large amounts of advertised routes. It is not recommended to integrate interior routing protocols with BGP.

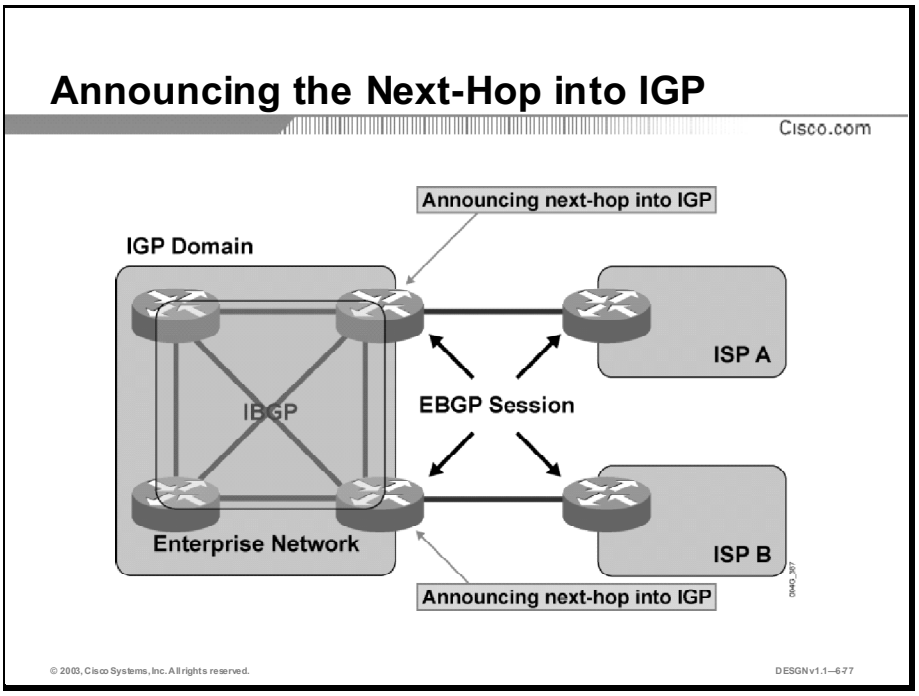

When using Internal BGP (IBGP) in the enterprise network, make sure that the next hop that EBGP advertises is announced in the network.

The IGP propagates only IP subnets used to provide BGP next-hop address reachability. IGP propagates enterprise subnets, including Building Access networks. All routes are redistributed into BGP at the domain borders (on Campus Backbone devices).

BGP is often used to offload the IGP in networks with large Campus Backbone and small Building Distribution sites. BGP runs in the Campus Backbone and the Campus Backbone IGP only propagates IP subnets used in the Campus Backbone itself to provide BGP next-hop address reachability.

Other IGP domains propagate their subnets, including Building Access networks. All routes are redistributed into BGP at the domain borders on Campus Backbone devices.

Summary

This topic summarizes the key points discussed in this lesson.

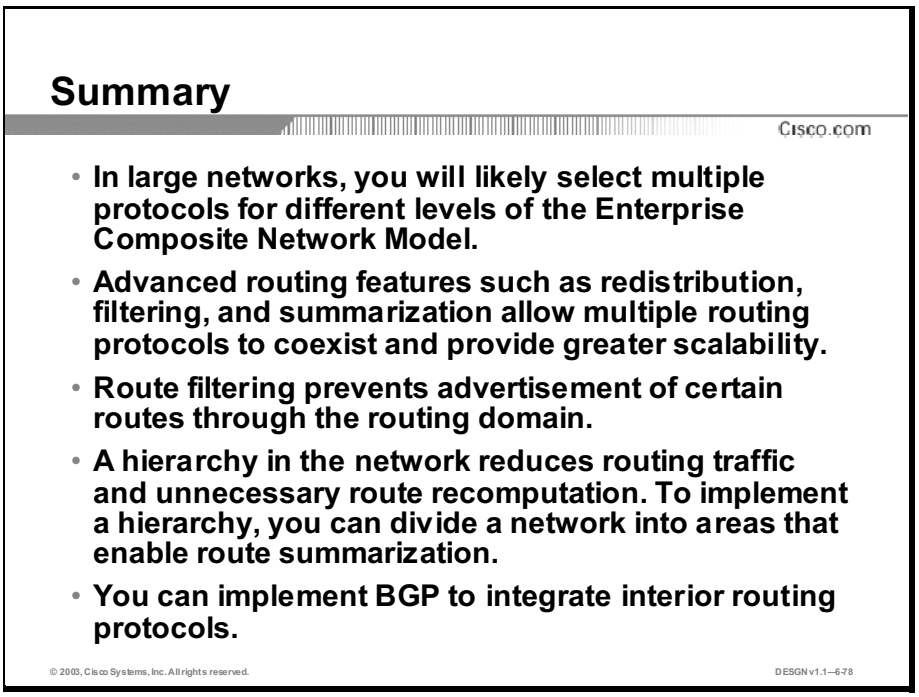

References

For additional information, refer to these resources:

- *Building Scalable Cisco Internetworks* (BSCI)
- *Interconnecting Cisco Network Devices* (ICND)
- *Designing Large-Scale IP Internetworks*, http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2003.htm

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Case Study 6: Routing Protocol Selection
- Simulation 6: Network Convergence

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match the routing protocols with the hierarchical level on which they are most likely to be implemented.

A) OSPF

B) IS-IS

C) static routing

D) EIGRP

E) RIPv2

F) BGP

G) ODR

_____ 1. core

_____ 2. WAN Access

_____ 3. Internet and remote access

Q2) Select two recommended redistribution points in a network. (Choose two.)

A) Network Management module

B) Building Distrbution submodule

C) Remote Access and Internet Connectivity modules

D) Campus Backbone submodule

E) Internet Connectivity module

Q3) Why is route filtering used?

A) to prevent routing loops and suboptimal routing upon redistribution

B) to prevent route leaking

C) to reduce the size of a routing table

D) to facilitate route summarization

Q4) Select the explanation that best describes why route summarization is used.

A) to impose a hierarchical IP addressing plan

B) to enable the use of more than one routing protocol

C) to reduce the routing overhead, such as by using smaller routing tables, and to improve the stability of routing

D) to enable coexistence of BGP with IGPs

Q5) How do you select between the primary and backup links when two exit points exist?

A) Assign a higher cost to the redistributed route on the primary router.

B) Assign a higher cost to the redistributed route on the secondary router.

C) Assign a static route on the primary router.

D) Assign a static route on the secondary router.

Quiz Answer Key

Q1) 1=A,B,D
2=A,D,E,G
3=C,F

Relates to: Hierarchical Network Structure and Routing Protocols

Q2) B, C

Relates to: Route Redistribution

Q3) A

Relates to: Route Filtering

Q4) C

Relates to: Route Summarization

Q5) B

Relates to: Integrating Interior Routing Protocols with BGP

Case Study 6: Routing Protocol Selection

Complete this case study to practice what you learned in this lesson.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the “Course Introduction” module
- DJMP Industries Case Study Scenario, presented at the end of the module ”Applying a Methodology to Network Design”
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the options for designing routing protocol deployment in the enterprise network. Upon completing this case study, you will be able to meet this objective:

- Select the suitable routing protocol for the network.

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the “Course Introduction” module
- Case study solutions, presented in the appendix “Case Study Solutions”

Exercise Procedure

Complete these steps:

Step 1 Refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario before commencing the exercise. Focus on the routing protocol issues. Allow a maximum of 10 minutes for reading.

Step 2 Discuss the scenario and options for a new routing protocol selection for the DJMP Industries network. Allow 10 minutes for the discussion.

Step 3 Propose the most suitable routing protocol and explain major deployment issues.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class and you have justified any major deviations from the case study solution.

Simulation 6: Network Convergence

Complete this exercise to practice what you learned in this lesson.

This exercise is a paper-only version of the simulation that was actually performed by the simulation tool and it includes the results that the simulation provided.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying Design Principles in Network Deployment"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the options for a routing protocol in the enterprise network design. One of the important issues in the network is its convergence after the changes. Upon completing this simulation, you will be able to meet this objective:

- Explain the importance of fast convergence in the network

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Read the simulation scenario and try to answer the questions that appear in the text. Discuss possible answers and explain your considerations in the classroom.

Network Convergence Scenario

This simulation addresses the network convergence issues in the campus network. The focus is on Layer 2 versus Layer 3 convergence details.

Although the time the network needs to recover depends basically on the combination of the routing protocol and the structure of the network, designers adhere in general to these two rules:

- The multilayer networks improve the convergence times around various failures.
- Pure distance vector protocols are slower than link-state protocols.

You decided to prove the above statements in the company's campus network by testing the network adaptability against some typical network failures.

Initial Traffic

For testing purposes, a campus design was proposed as a solution to the Campus Design Case Study. The company configured the application to support 20 Mbps of bidirectional traffic with the rate of 5000 pps between the workstation WS_A in building A and the server in the central building. This will serve as a model of reference for the traffic flow in the simulation tests.

Bridged Network Convergence

The first simulation tests the convergence in a pure data link layer network. The results will be used later for comparison with other models.

Initially, the reference (WS_A to Server) traffic was sent over the path as determined by the STP: AS_F2 – DS_B – CS_A – DC_B (as illustrated in the figure).

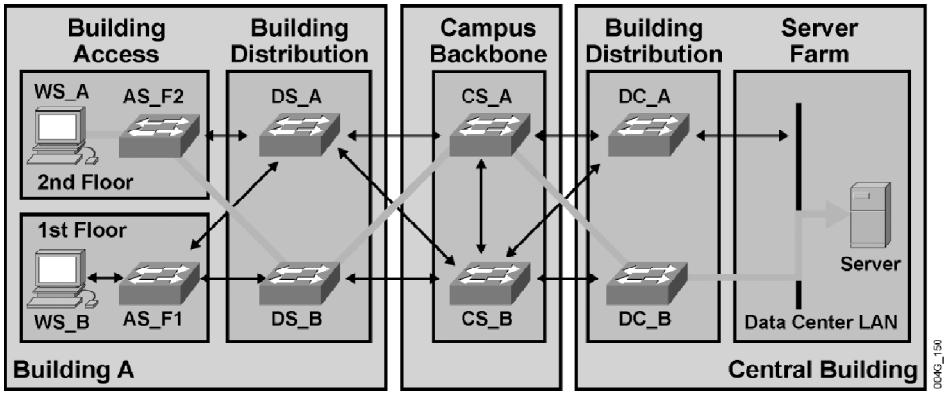

Failure Scenario

You start studying the effect of the link and node failure on the network performance by first tearing down and then restoring the CS_A – DC_B link. The event takes place between the 200th and the 300th second, followed by the 100 second disabling of the CS_A node triggered at the 400th second.

Test—Convergence Around Link Failure

At 200 seconds, the CS_A – DC_B link fails.

The simulation tool generates a route report (in the log file) that keeps track of the traffic on its way from the WS_A to the server. The last packet that reached the destination is seen at the 199th second. The subsequent packets were dropped. The figure illustrates the path (thick line) taken by the last successful packet. The packet took the best STP route.

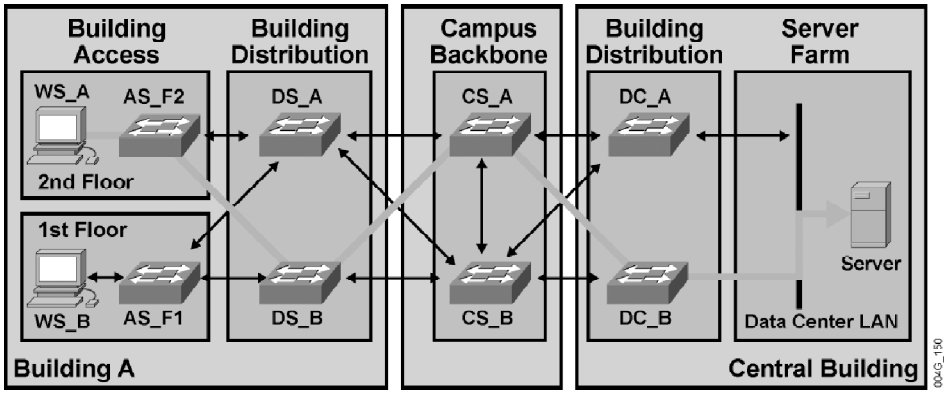

One second later the switch starts recalculating the STP tree, first entering the listening mode. As the CS_A switch is not yet able to provide an alternative path, the subsequent packets are dropped on the troubled port. (See the thick line in the figure.)

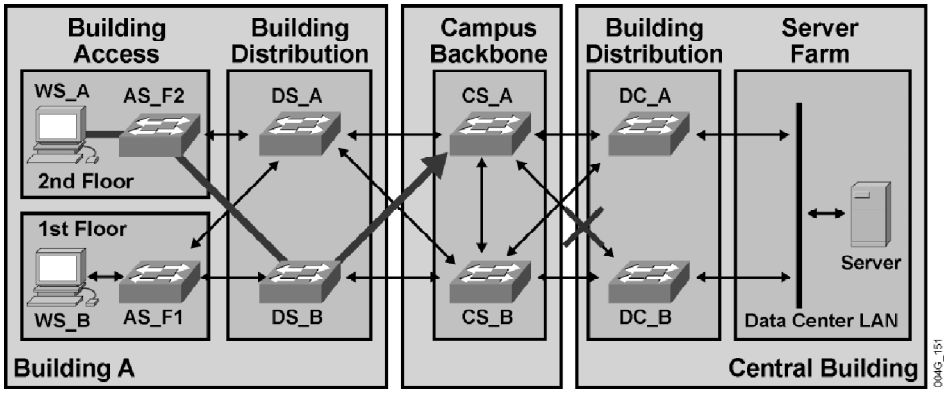

The first packet after the link failure is seen at its destination at the 248th second, after the STP has found an alternative path. The figure illustrates how the packet flows across an alternate route.

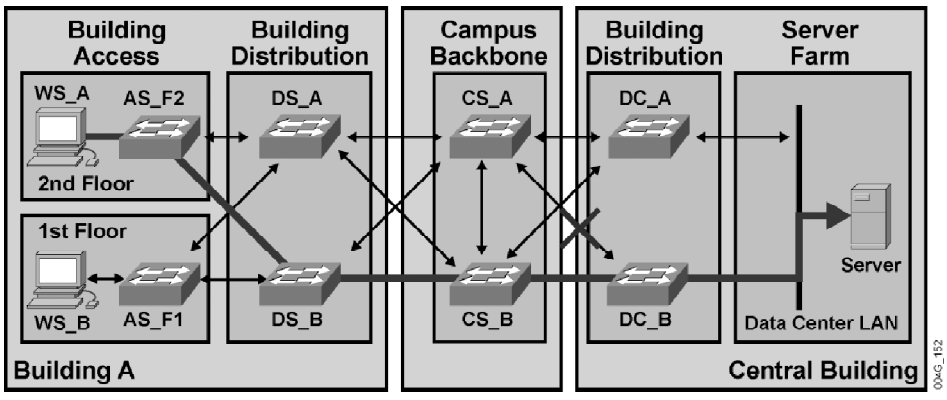

The packet flow is interrupted again after the CS_A – DC_B link is reestablished. The last packet is seen at the 299th second, and the first packet after link reestablishment is seen only at the 315th second. After the link reestablishment, the packets yet again flow over the previous path selected by the STP.

Conclusions

Because of the STP the bridged network is disrupted after a link failure for almost 50 seconds and there is another (approximately 15-second) disruption after link reestablishment.

Q1) The outage caused by the loss of the primary link can be explained as the consequence of the STP recalculation. What is the reason for the second delay after the primary path has physically recovered?

_____
_____
_____
_____

Test—Convergence Around Node Failure

At 400 seconds, the CS_A node fails. The last packet is seen at the 399th second and it takes the best STP route as illustrated in the figure (thick lines).

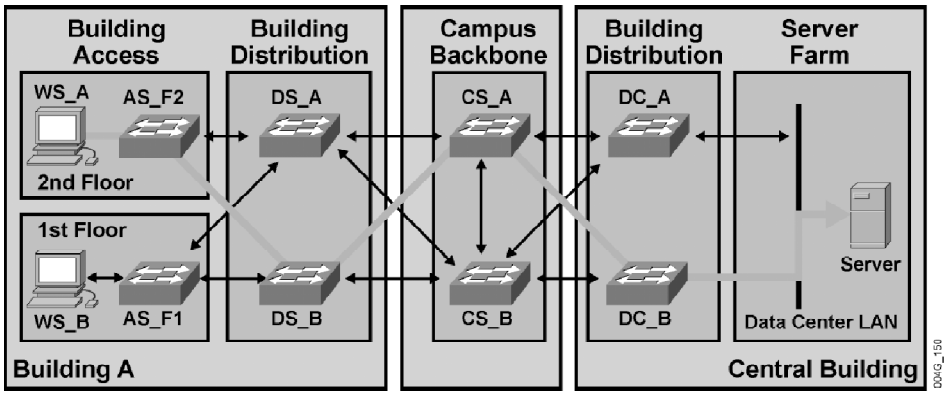

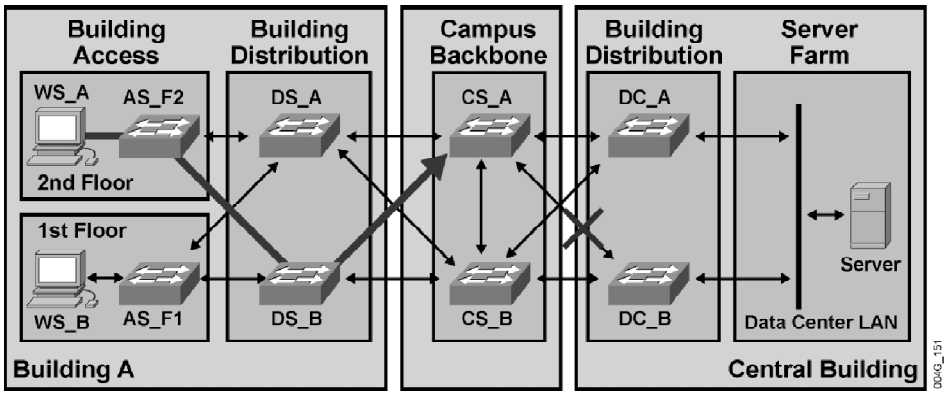

The first packet after link failure is seen at its destination at the 448th second. The convergence time is again almost 50 seconds. The packet flows across an alternate route as illustrated in the figure (thick lines).

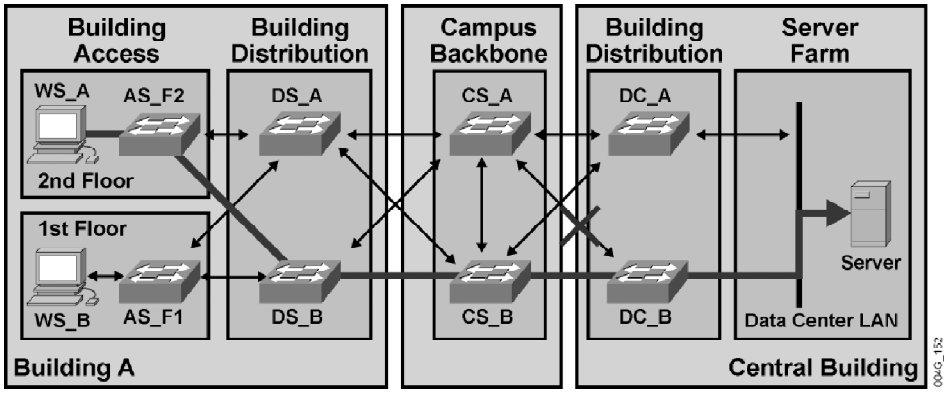

After the node CS_A recovers (at the 500th second), the packet flow is interrupted again. The last packet is seen at the 499th second and the first packet after the node recovery is seen only at the 515th second. After the node recovery, the packets yet again flow over the previous path selected by the STP.

Conclusions

Because of the STP, the bridged network is disrupted after a node failure for almost 50 seconds, and there is another 15-second disruption after node recovery.

Q2) Why do the link and node incidents impose the same time for network recovery?

_____
_____
_____
_____

Convergence in Mixed Data Link Layer/Multilayer Network

The Building Distribution switches in the campus have been upgraded to multilayer switches running OSPF.

Note: To avoid load sharing and thus simplifying the failure analysis, the OSPF cost on one path across the network has been lowered.

The traffic between the WS_A and the server flows over one path only: AS_F2 – DS_A – CS_A – DC_A—as illustrated with thick lines in the figure.

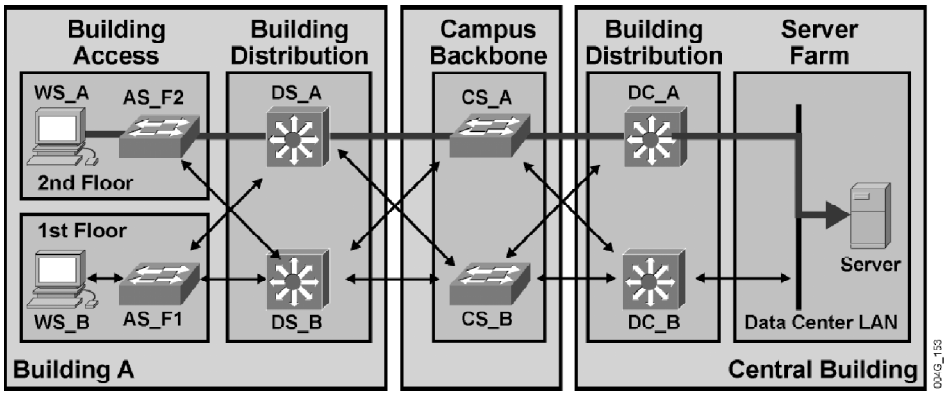

Failure Scenario

The failure scenario incorporates these subsequent events: the 100-second flap of the DC_A – CS_A link triggered at the 200th second, followed by the complete outage of the DC_A node operation between the 400th and 500th second.

Test—Convergence Around Link Failure

At 200 seconds, the CS_A – DC_A link fails. The link loss event immediately triggers the DC_A router to generate and distribute a new OSPF link-state packet, notifying other routers (multilayer switches) about a recent change. The convergence time is 6 seconds, approximately 1 second for the change in the link status to be propagated from DC_A to other routers and 5 seconds delay before the SPF is run in all routers.

With the failure of the CS_A – DC_A link, the first successful packet transmission is seen in the log at the 205th second. There are now two alternate equal-cost paths across the network and the packets are load shared across them (as illustrated by the thick lines in the figure).

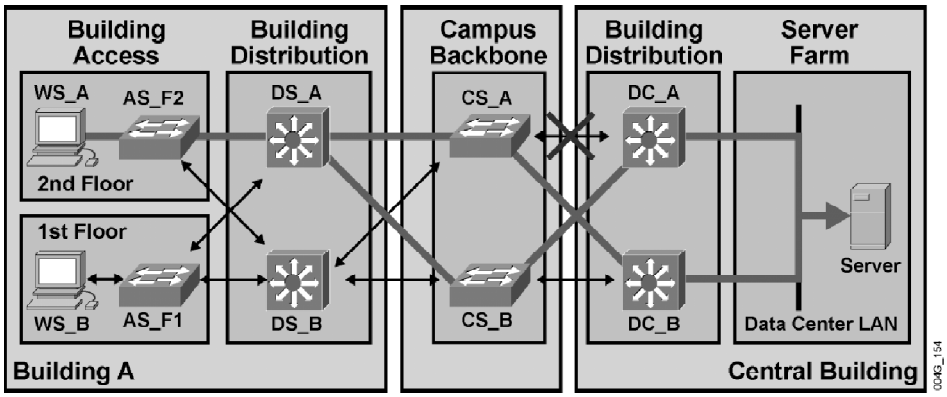

After the CS_A – DC_A link is reestablished at 300 seconds, the packet flow is not interrupted, as opposed to what occurs in the data link layer solution. However, the convergence to the recovered path takes approximately 30 seconds, the time that it takes for OSPF adjacency to be established across the recovered link, the adjacency database to be exchanged, the changes propagated across the network, and the SPF algorithm run on all routers. The packet traversing the network at the 330th second is again following the primary path (as configured initially by tuning the OSPF cost parameter).

Conclusions

Using OSPF, the network is disrupted for approximately 6 seconds following a link failure, the time that it takes the SPF algorithm to be run on all routers. Following the link recovery, there is no disruption in traffic flow, but the traffic returns to the primary path approximately 30 seconds after link recovery.

Test—Convergence Around Node Failure

At 400 seconds, the DC_A node fails. The last successful packet is seen at the 399th second, and the next packet is already dropped. (See the thick lines in the figure.)

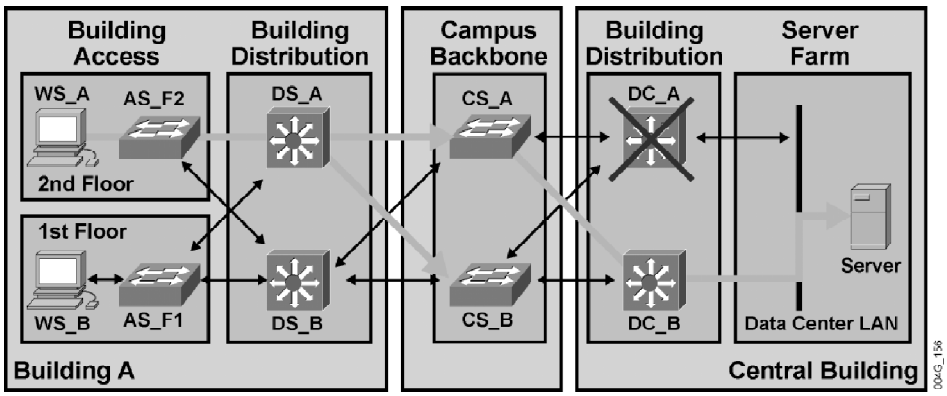

None of the OSPF routers are directly connected to the DC_A and are thus unable to detect the loss of a link and propagate a new update. However, the change is learned through a lost adjacency (missing hello packets).

The OSPF convergence time is between 35 and 45 seconds. This is the time that it takes the OSPF, through the OSPF dead timer, to discover that the neighbor is down (between 30 and 40 seconds with default timers on LAN interfaces), as well as the time that it takes for the change to be propagated across the network (less than 1 second) and the time for the SPF run (up to 5 seconds delay).

After 35 seconds, the router DS_A appears to be propagating the traffic across an alternate route (DS_A – CS_A – DC_B) and, at the same time, is sending some traffic to CS_B. The latter action is a result of the OSPF Hold time that has not yet expired. The DS_A is under impression that the failed DC_A can still be reached through the CS_B switch and, therefore, used as an alternate path. (See the thick lines in the figure.)

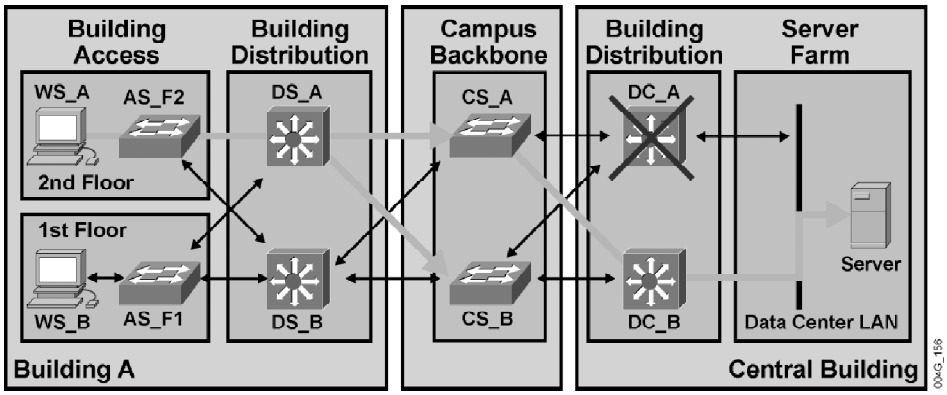

After 45 seconds the convergence is complete. The packet flows only across an alternate route, as illustrated by the thick lines in the figure.

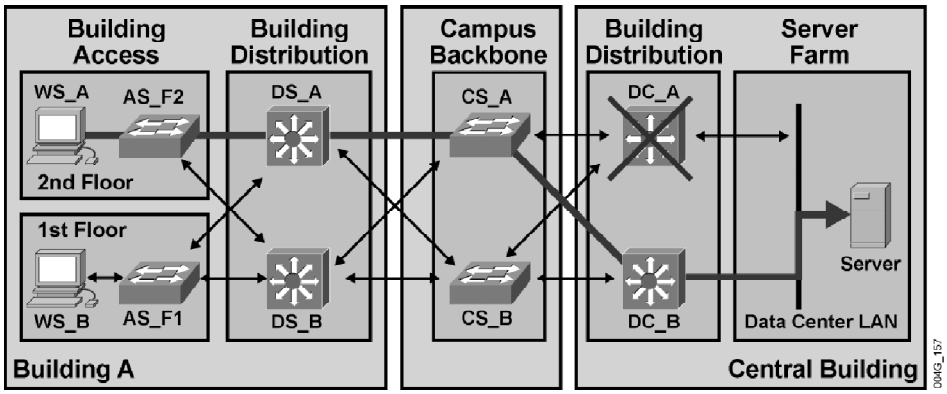

After the node DC_A recovers at 500 seconds, there is no interruption in packet flow. However, the path recovery takes longer. The packet traversing the network at the 519th second is still following the backup path. The path recovery takes approximately 20 seconds, with no interruptions in the packet flow.

Conclusions

Using OSPF in a mixed Layer 2/Layer 3 network, the network is disrupted for approximately 35 to 45 seconds after a multilayer switch failure. (The exact time depends on OSPF timer settings.) After node recovery, there is no disruption in the traffic flow, and the traffic is rerouted to the primary path only after the OSPF reestablishment is complete, approximately 20 seconds.

Q3) If you choose RIP instead OSPF, the convergence will change significantly under a link or node failure. Why does this occur, and what special procedures do RIP routers undergo?

Network Convergence in a Multilayer Network

Campus Backbone and Building Distribution devices are multilayer switches. The OSPF is run across the network. To avoid load sharing, the default costs are altered, simplifying the failure analysis. The traffic between the WS_A and the server flows over the following path: AS_F2 – DS_A – CS_A – DC_A. (See the thick lines in the figure.)

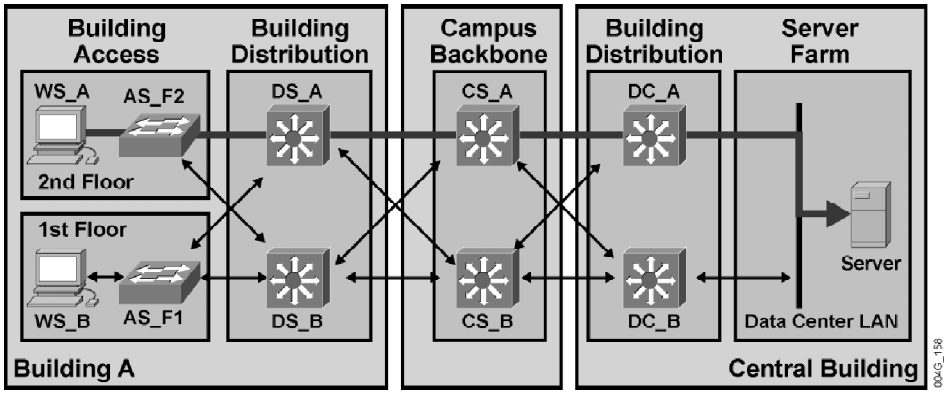

Failure Scenario

The failure scenario used in the previous simulation remains in place: the 100-second flap of the DC_A – CS_A link triggered at the 200th second, followed by the DC_A node outage between the 400th and 500th second.

Test—Convergence Times Around Link Failure

At 200 seconds, the CS_A – DC_A link fails. The packet flow is not interrupted, because both of the involved routers (CS_A and DC_A) propagate the loss of an attached link immediately through the use of a new OSPF link-state advertisement. The CS_A reroutes the traffic across the DC_B node. The first packet after the link failure takes the paths, as seen in the figure (solid lines).

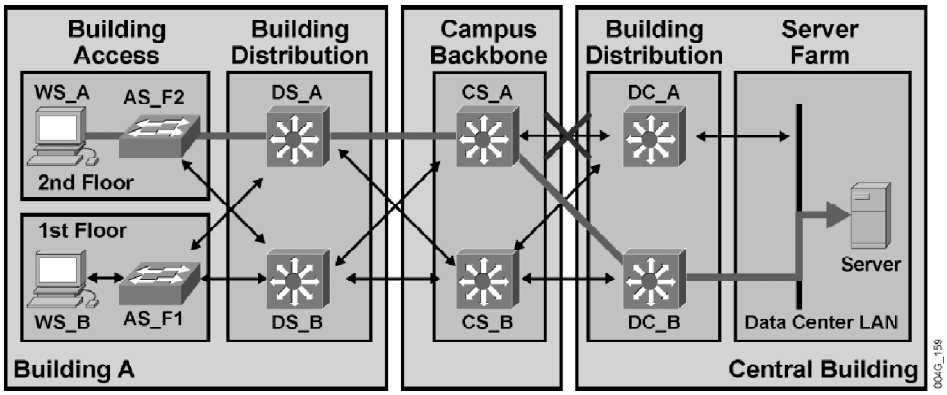

The other routers need some time to receive an update. In addition, because of OSPF behavior, the SPF recalculation is delayed for 5 seconds. It therefore takes approximately 5 to 6 seconds after the link failure for the other routers to recalculate the possible paths. As a result of the recalculation, there are now two alternate equal-cost paths across the network and the packets are load shared across them (as illustrated by the thick lines in the figure).

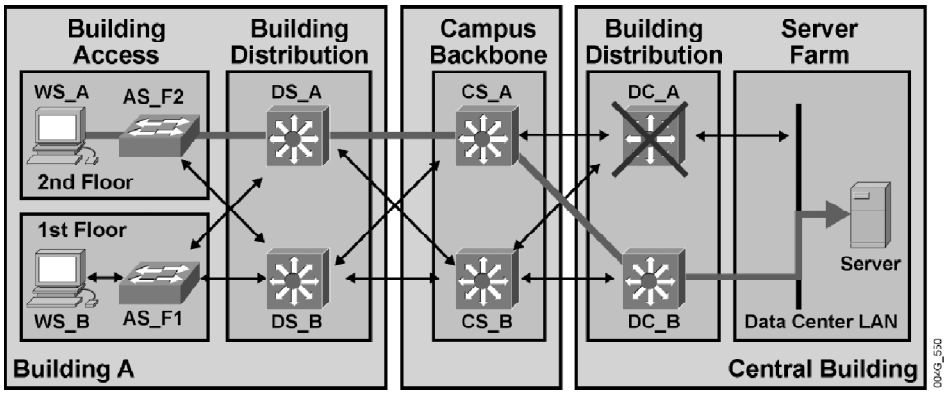

After the link CS_A – DC_A is reestablished at 300 seconds, the packets continue to flow over the backup paths until the OSPF on the primary link is fully reestablished.

Q4) From the simulation log, it appears that the switchover to the primary path takes too long, approximately 50 seconds. Why?

_____
_____
_____
_____

Conclusions

Using OSPF in the entire campus network, there is no disruption following a link failure. The router that is attached to the lost link reroutes traffic immediately over the alternative link, whereas the other routers need some time to be notified about a change and complete the SPF recalculation. Upon recovery of the link, there is no disruption in traffic flow, but the traffic is only rerouted on the primary path approximately 50 seconds after link recovery.

Test—Convergence Times Around Node Failure

At 400 seconds, the DC_A node fails. The OSPF convergence time is again immediate, with no packet loss.

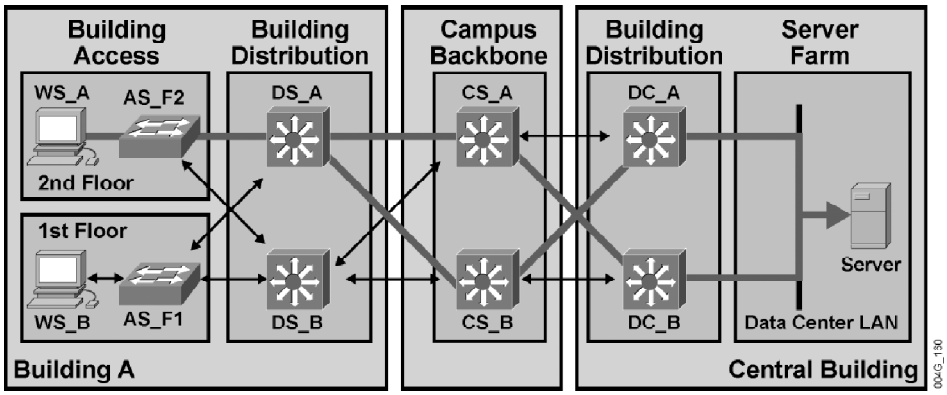

After the node DC_A recovers at 500 seconds, the path recovery again takes longer than expected. The packet traversing the network at the 550th second is finally switched back to the primary path.

Conclusions

There is no disruption in the traffic flow after node failure and recovery, but the traffic is rerouted to the primary path only after the OSPF establishment is complete, approximately 50 seconds.

Module 7

Evaluating Security Solutions for the Network

Overview

Network security is an essential network service that spans the entire network. The Enterprise Composite Network Model modularity helps you focus on a security problem within a particular network module and integrate a solution into a global network design. A modular approach simplifies the design and ensures that a security breach in one network module will remain isolated, not affecting the entire network.

This module introduces the Cisco SAFE blueprint, a security architecture that employs a modular approach to designing network security. This module evaluates security from the physical perspective up to the individual application security. Logging and monitoring are described as integral parts of any security solution and as preventive mechanisms to detect attacks before serious consequences.

Module Objectives

Upon completing this module, you will be able to select the appropriate security solutions for a given network.

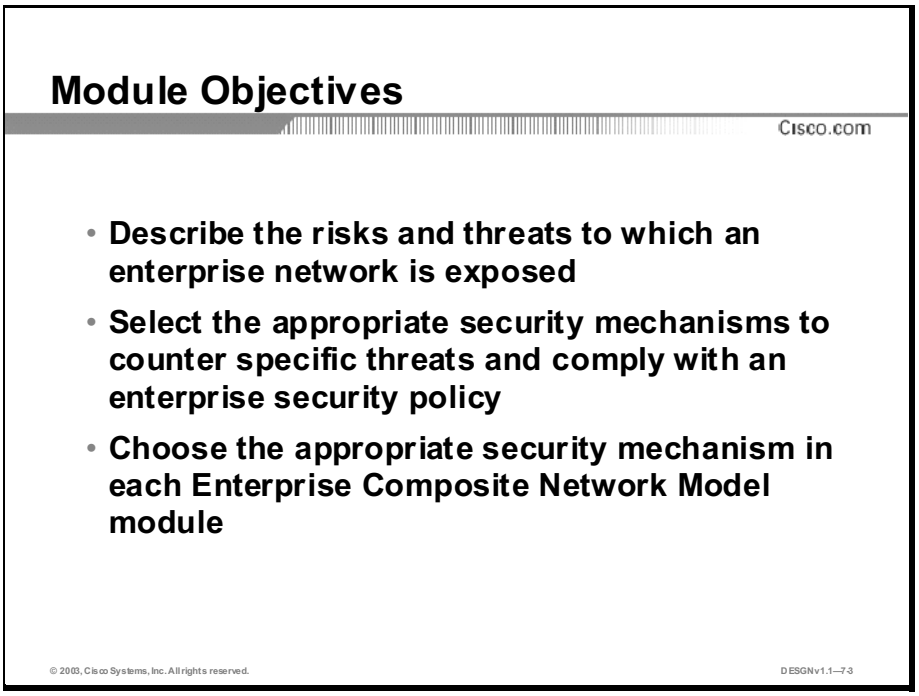

Module Outline

The outline lists the components of this module.

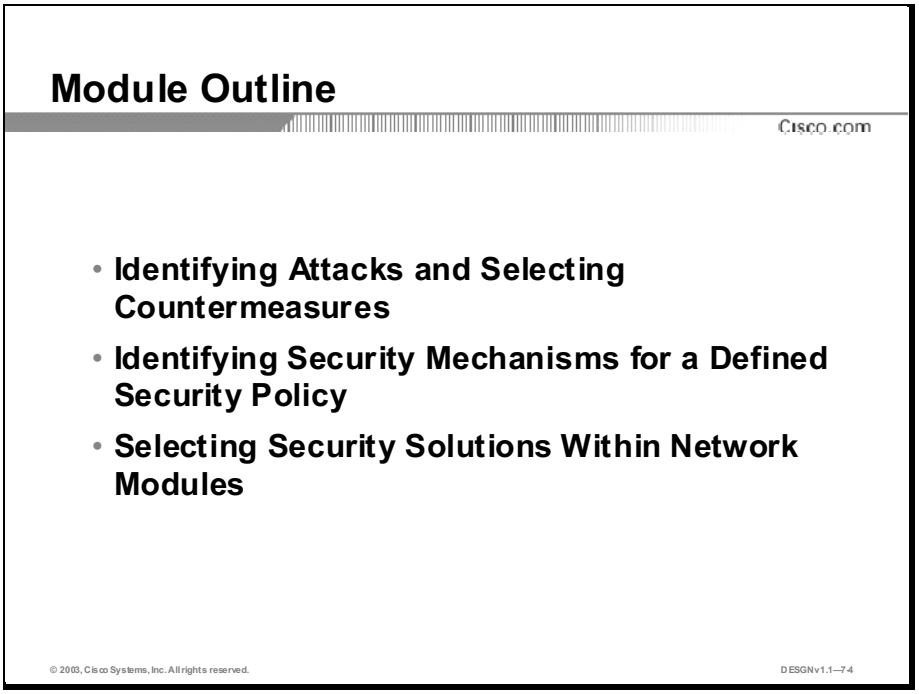

Identifying Attacks and Selecting Countermeasures

Overview

The scope of a network security solution is determined by organizational requirements and potential threats evaluated for each network component. The lesson focuses on network components as potential targets, and explains the risks associated with threats directed at them.

Relevance

Before you can design security for a network, you need to understand the threats posed.

Objectives

Upon completing this lesson, you will be able to describe the risks and threats to which an enterprise network is exposed. This includes being able to meet these objectives:

- Describe the main categories of security threats
- Identify the attacks directed at network devices such as routers and switches and describe possible protection mechanisms
- Describe attacks directed at an entire network and describe possible protection mechanisms
- Identify the security weak points of hosts and their applications

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in Cisco IOS software

Outline

The outline lists the topics included in this lesson.

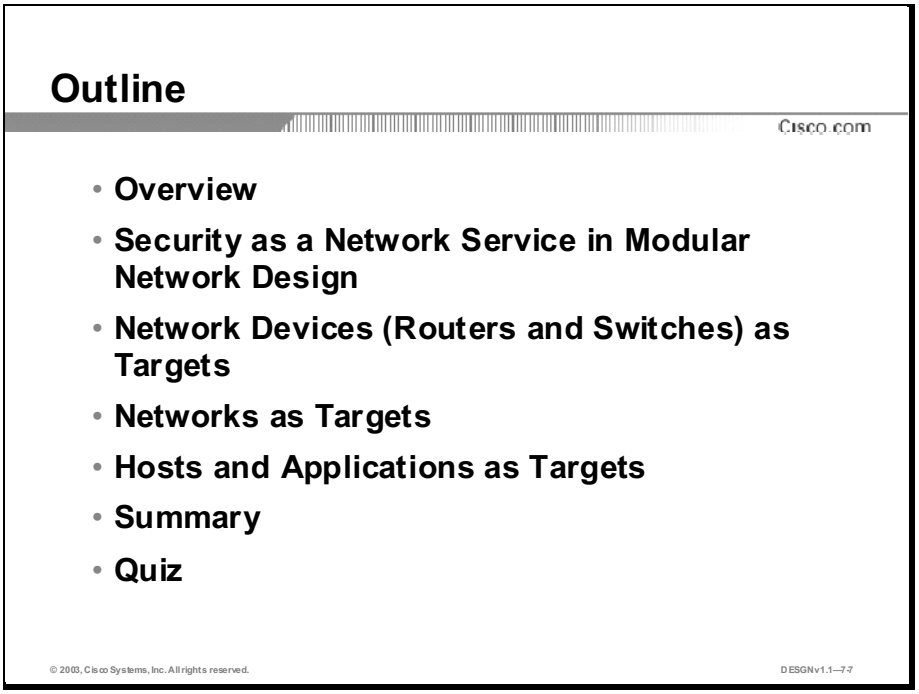

Security as a Network Service in Modular Network Design

As networks become increasingly interconnected and data flows more freely, security services become critical. In the commercial world, connectivity is no longer an option, and the possible risks do not outweigh the benefits. Therefore, security services must provide adequate protection to conduct business in a relatively open environment. This topic describes basic network security assumptions.

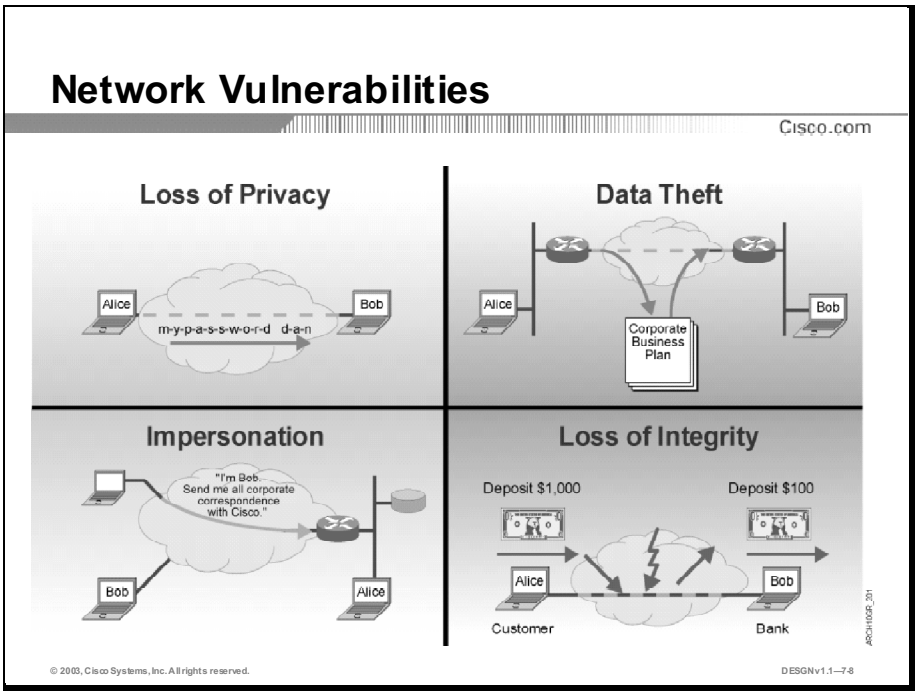

To provide adequate protection of network resources, the procedures and technologies deployed need to adequately provide these things:

- Confidentiality of data, which should ensure that only authorized subjects can view sensitive information
- Integrity of data, which should ensure that only authorized subjects can change sensitive information and guarantee the authenticity of data
- System and data availability, which should ensure uninterrupted access to important computing resources

When designing for security, you must be aware of the attacks that could compromise security and their associated risks.

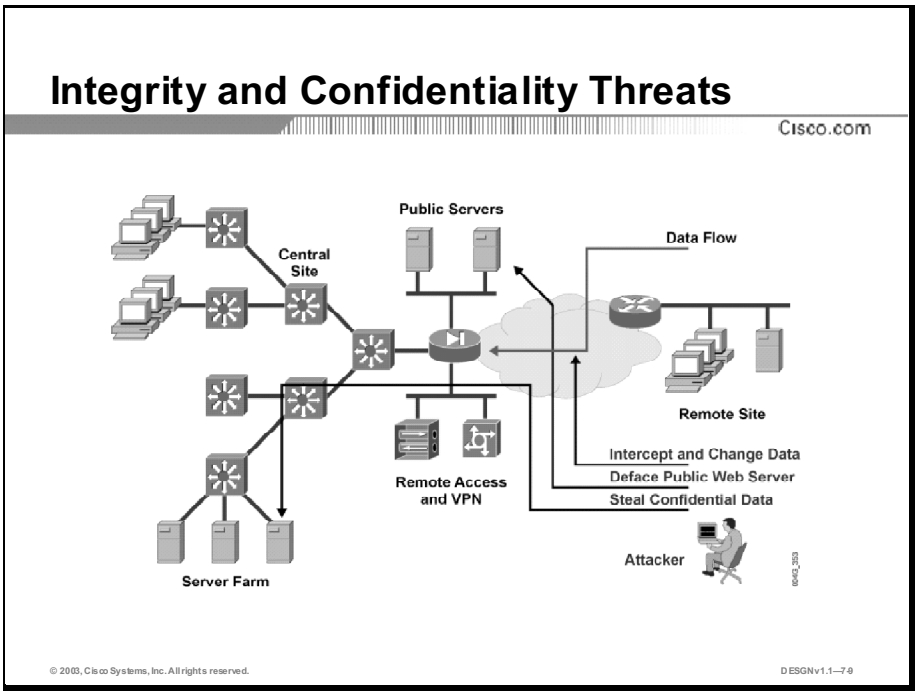

The major security threats are integrity violations and confidentiality breaches.

Integrity Violations

Integrity violations occur when the attacker attempts to change sensitive data without proper authorization. For example, the attacker obtains permission to write to sensitive data and changes or deletes it. The owner may not detect the change until it is too late, perhaps when the change has already resulted in tangible loss. Many businesses treat integrity violations as the most serious threat to their business because of the difficulty in detecting changes and the possible cascading consequences of late detection.

Confidentiality Breaches

Confidentiality breaches occur when an attacker attempts to read sensitive data. These attacks are extremely difficult to detect because the attacker can copy sensitive data without the owner's knowledge and without leaving a trace.

The risks of both integrity violations and confidentiality breaches are usually managed by enforcing access control in various ways, for example:

- Limiting access to network resources using network access control, such as physical separation of networks, restrictive firewalls, and VLANs
- Limiting access to files and objects using operating system-based access controls, such as UNIX host security and Windows NT domain security
- Limiting user access to data by application-level controls, such as different user profiles for different roles
- Using cryptography to protect data outside the application, such as encryption to provide confidentiality and secure fingerprints or digital signatures to provide data authenticity and integrity

Example: Integrity and Confidentiality Threats

The figure illustrates potential confidentiality and integrity threats to network resources that an outside attacker might exploit. If there is not adequate protection in place, an attacker might:

- Access an internal server and copy confidential data (a confidentiality breach)
- Deface the corporate web page (an integrity breach)
- Intercept data sent over the Internet between a branch office and the central site, and change or read data in transit (a confidentiality or integrity breach)

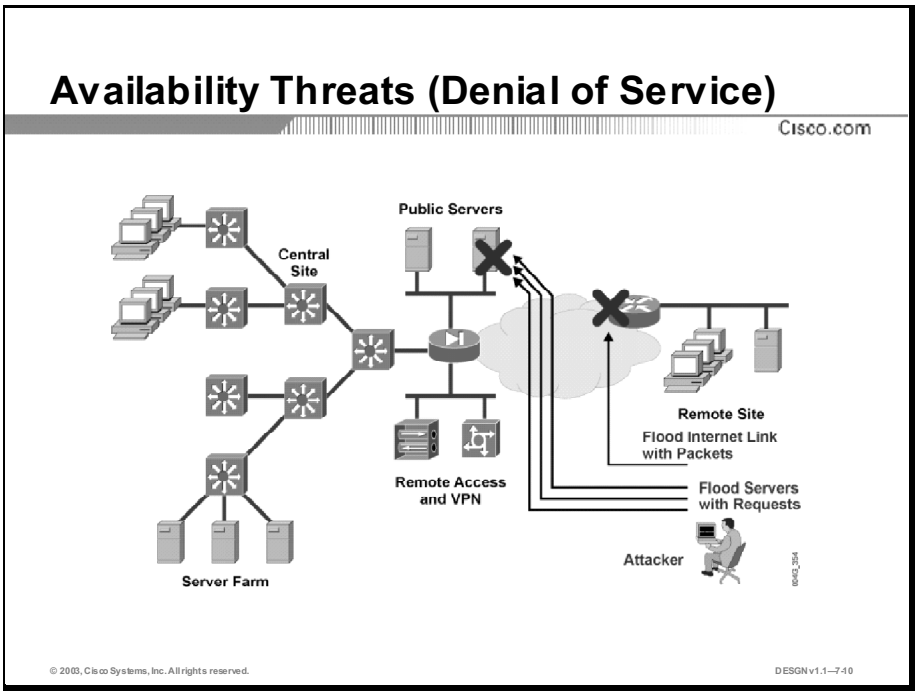

Denial-of-service (DoS) attacks attempt to compromise the availability of a network, host, or application. They are considered a major risk as they can easily interrupt a business process and cause significant loss. Denial-of-service attacks are relatively simple to conduct, even by an unskilled attacker.

Denial-of-service attacks are usually the consequence of one of these failures:

- A host's or application's failure to handle an unexpected condition, such as maliciously formatted input data, an unexpected interaction of system components, or simple resource exhaustion.
- A network's, host's, or application's inability to handle an enormous quantity of data, which crashes the system or brings it to a halt. The difficulty of defending against such an attack lies in the difficulty of distinguishing legitimate data from attacker data.

The figure depicts potential availability threats to network resources that an attacker might exploit. Without adequate protection in place, an attacker might:

- Flood a public server with an enormous number of connection requests, rendering the server unresponsive to legitimate users.
- Flood a bottleneck network connection with random traffic, in an attempt to consume as much bandwidth as possible. This may deny service to legitimate users of that connection.

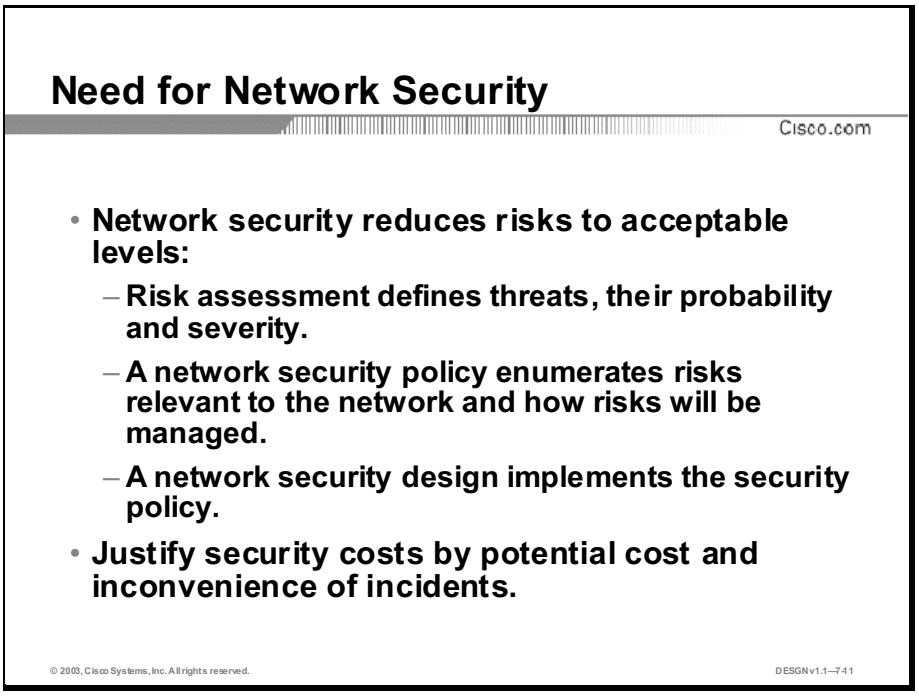

Network security employs risk management to reduce risk to acceptable levels. An acceptable level of risk is defined by the organization based on these factors:

- Value of the organization's data
- Expectation of loss in the event of compromise
- Severity and probability of risks

The weighing of these factors is called risk assessment, which is a continuous procedure of discerning these things:

- What assets to protect
- The value of the assets
- The cost of expected loss that would result from a compromise
- The severity and probability of attacks directed against the assets

Risk Assessment and Management

A risk assessment provides input to a network security policy, which documents the level of risk and suggests the methods of managing risk to an acceptable level.

The network security policy describes risk management measures as they relate to potential threats. It does not usually consider security implementation details but provides a more general security philosophy, which directs the implementation of security mechanisms.

Risk management and the consequent building of the security policy is a continuous process, as the severity and probability of risks change daily. A good example is the use of cryptography to provide confidentiality through encryption. A company's encryption algorithm and the length of the encryption key may require a change, if a relatively inexpensive and exceptionally fast code-cracking computer becomes available. The organization will need to choose a stronger algorithm to protect against the new threat.

For most scenarios, you can evaluate potential damage to some degree. Take care that the cost of security does not exceed the cost of potential security incidents. In the commercial world, it is common practice to build systems with just enough security to bring potential losses down to the desired level. Alternatively, organizations with higher security requirements may want to implement stronger measures than are deemed necessary to mitigate potential unforeseen risks.

Example: Valuing Assets

A security designer will evaluate how serious a particular risk is, which depends on the damage a successful attack could cause. However, it is often difficult to associate a value with an asset, as in these cases:

- A medical database of a large hospital system, where there are disastrous consequences if confidentiality is breached
- A corporate public web page, which if defaced (an integrity violation), can become a public relations nightmare, although it may not result in any serious breach of confidentiality

Network Devices (Routers and Switches) as Targets

Each device on the network, such as a router or switch, is a potential security target. This topic describes possible attacks against network devices and associated protection mechanisms.

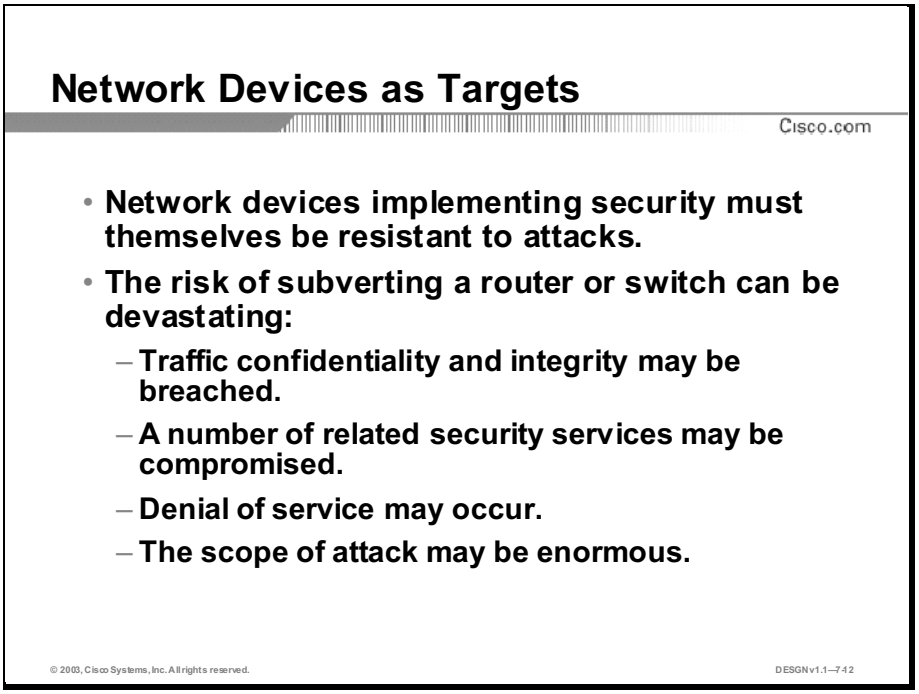

The classic guideline in developing trusted systems specifies that a trusted system must enforce an organization's security policy and the system itself must be secure against attacks. Apply the same principles to network security. In a secured network, network devices provide security services to network users, and the devices themselves must be resistant to attacks.

If an attacker can subvert a network device, that attacker can then launch many potential attacks from the compromised device, such as these:

- Interception of data flowing through the network, its analysis, and alteration (compromising confidentiality and integrity)
- Attacking related security services, which rely on trust among network devices (for example, injecting malicious routing information or subverting authentication protocols that are used by the compromised device)
- Denial of service by making the device unavailable or by changing its settings to deny connectivity

Example: Security Breaches to Network Devices

Depending on the location and the importance of a device, the results of such a security breach can be devastating. Consider these examples of the results of an attack on a network device:

- A core enterprise switch is compromised, resulting in an attacker's ability to analyze and alter all mission-critical data switched through it.
- An enterprise remote-access server is compromised, allowing an attacker to gather user passwords and use them for further intrusions and, perhaps, to dial out and attack other networks.
- A firewall is compromised, allowing an attacker to reconfigure it and open arbitrary connections into the protected network.
- A border router is compromised, allowing an attacker to intercept the company's Internet traffic and perhaps deny connectivity.
- A core ISP router is compromised, allowing an attacker to influence routing in the whole ISP network, intercept traffic of countless customers, and impersonate secure servers.

Example: Attack on an ISP Core Router

An attacker breaks into an Internet service provider (ISP) core router and configures the router to intercept all Domain Name System (DNS) traffic and send it to another system. The attacker can now change the data in any DNS reply and redirect any client to any server on the Internet. This illustrates the importance of protecting Internet core devices to prevent such large-scale attacks.

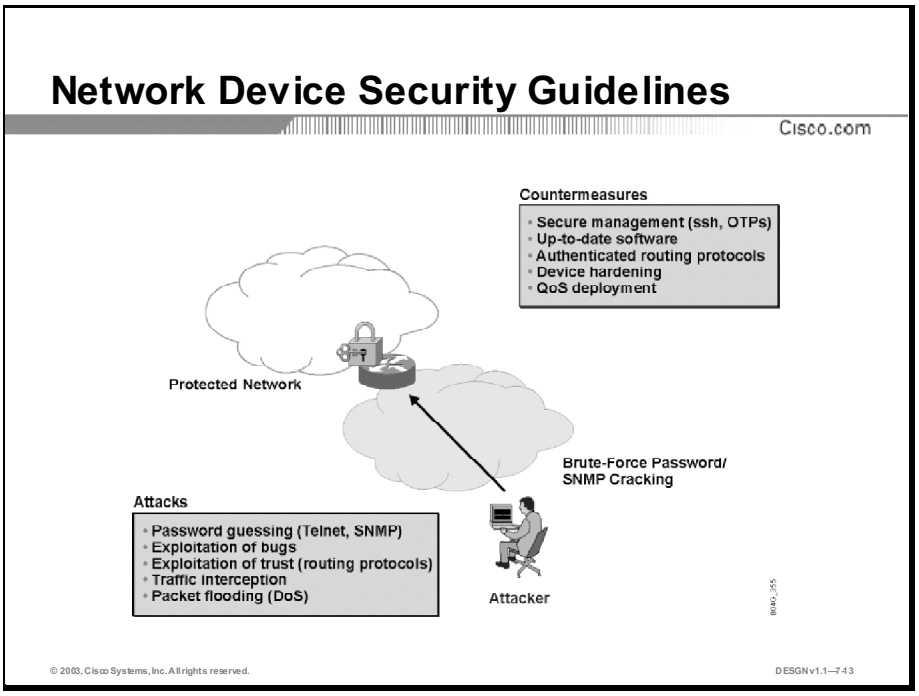

Securing network infrastructure involves several technical and nontechnical steps. These are the most important:

- Teaching the network administrators to use clearly defined, secure management procedures to avoid accidental human error.
- Treating each network device as a high-value host and hardening (strengthening) it against possible intrusions. This involves common practices such as running only the minimal necessary services and establishing trust only with authentic partners using authenticated routing protocols.
- Providing secure device management channels using strong authentication, session encryption, and change control (for example, using One Time Passwords [OTPs], Secure Shell Protocol [SSH], configuration command authorization, and administrator auditing).
- Patching the device software so it remains up to date with regard to security recommendations and known security issues.

Networks as Targets

Reconnaissance and denial-of-service attacks are directed at the network as a whole. This topic describes possible attacks against entire networks and associated protection mechanisms.

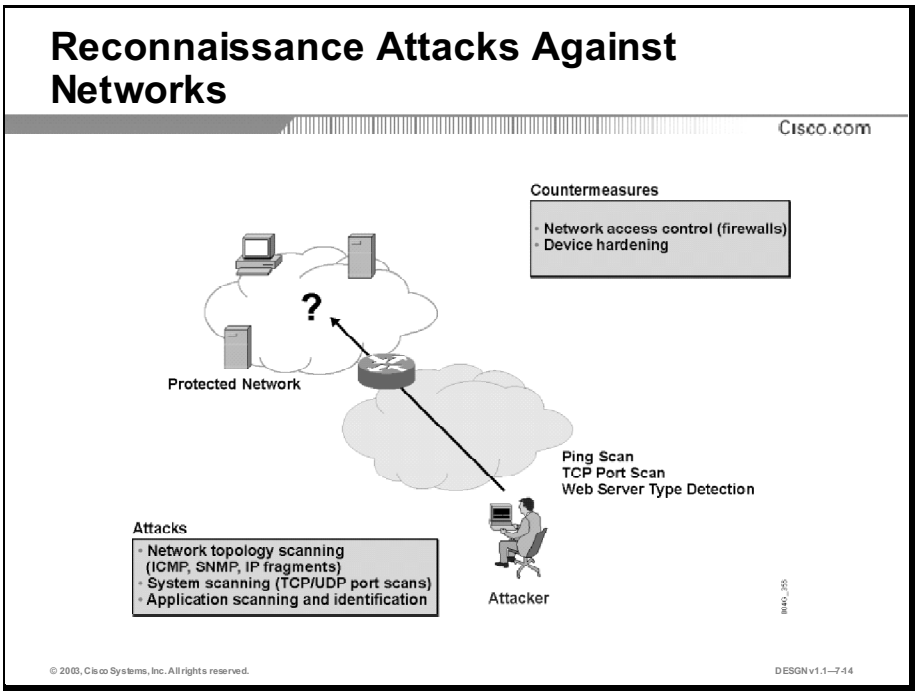

In a reconnaissance attack, an attacker searches the entire network for possible targets. Reconnaissance is usually the prelude to a more focused attack against a particular target. These are well-known reconnaissance methods used against networks to determine what subnets, hosts, and services are available as potential targets:

- Gathering information about a network from registries such as the DNS or the Internet registrars' databases.
- Discovering access possibilities using network mapping tools (for example, traceroute), "war dialing" (attempts to discover and connect to dial-up access points), and "war driving" (attempts to discover and connect to misconfigured wireless access points).
- Attempts of network mapping using tools that range from primitive (for example, the ping program or Simple Network Management Protocol [SNMP] queries) to very sophisticated (sending seemingly legitimate packets to map a network). Attackers use tools either to discover reachability of hosts and subnets (host scans), services (port scans), or to look for specific applications.

For an attacker, a reconnaissance attack provides a list of potential targets and weaknesses already discovered in the reconnaissance phase. The attacker can then prioritize the targets, determine how difficult an attack would be, and choose further actions.

You can stop most reconnaissance attacks by using perimeter defenses, such as firewalls that perform network access control to limit connectivity. In addition, devices should only give out the information necessary to support business needs.

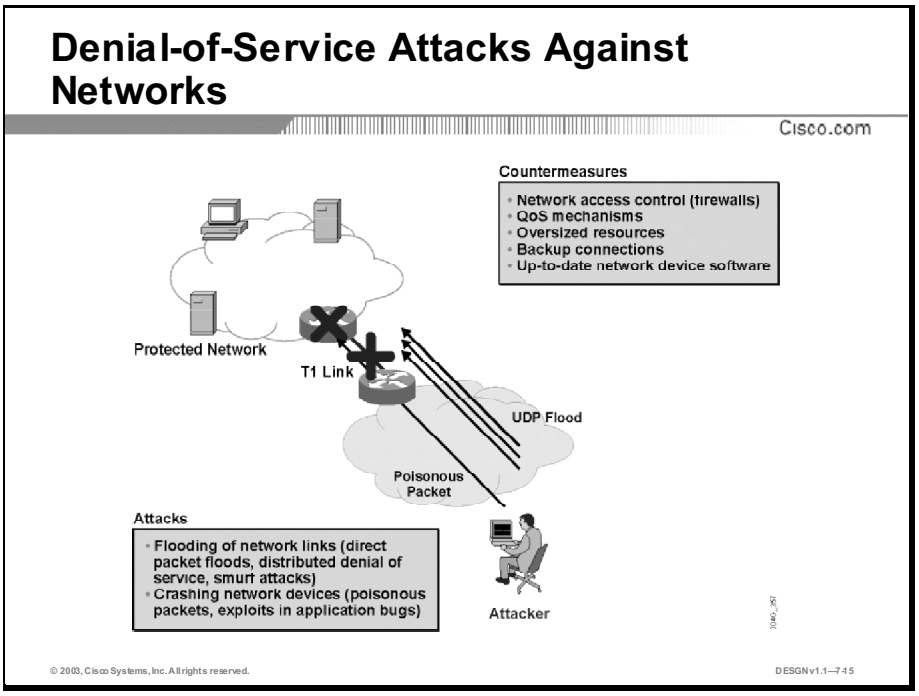

Denial-of-service attacks, directed at a network, can compromise connectivity to or from that network, effectively disconnecting it from its neighbors. The two methods of causing a denial-of-service attack, sending malformed data and sending a large quantity of data, apply to networks as a whole, as in these examples:

- An attacker sends a poisonous packet (an improperly formatted packet or a packet that the receiving device improperly processes) to a device, which causes it to crash or halt upon receipt. This may disrupt all communications over the device.
- An attacker sends a continuous stream of packets that overwhelms the available bandwidth of some network links. In most cases, it is impossible to differentiate between an attacker and legitimate traffic, and it is impossible to trace an attack quickly back to its source. In general, success correlates to bandwidth resources and whoever has more bandwidth, prevails. If an attacker compromises many systems in the Internet core, that attacker can take advantage of virtually unlimited bandwidth from many hosts to unleash packet storms at targets. This has already happened on the Internet and is called a distributed denial-of-service (DDoS) attack.

Most denial-of-service attacks are extremely difficult to trace and very difficult to defend against, especially those correlated to bandwidth resources. They can cause significant downtime and are considered among the most serious risks to networked businesses.

Example: Web Server Software Vulnerability

A critical vulnerabilty may be discovered in some of the software packages (such as web server software) that many businesses use on their servers. When vulnerabilities are known, many attackers scan networks to locate vulnerable hosts. Scans can quickly discover hundreds, and perhaps thousands, of vulnerable systems in an organization. A restrictive firewall at the network boundary can manage this risk by exposing only a few public servers.

Denial-of-Service Protection Guidelines

Denial-of-service attacks, which use “poisonous data” to disable network devices, are mitigated by limiting access to devices and using up-to-date software with security patches. In contrast, denial-of-service attacks, which take advantage of resource exhaustion, are difficult to trace and very expensive or even impossible to mitigate. The difficulty of separating the attacker’s data from the legitimate users’ data inside a data flood continues, as the attacker’s goal is to make this distinction difficult or impossible. Defending against these attacks requires a mixture of these resources:

- Heuristic-based defense mechanisms, which try to assess what data is malicious and discard it before it overwhelms a service. A good example of this is the TCP Intercept feature of Cisco IOS software and Cisco PIX Firewall systems.
- Having plenty of backup options, such as redundant servers and redundant connections, not all of which are vulnerable to the same type of denial-of-service attacks.

Hosts and Applications as Targets

The ultimate target of an attacker is often a host or an application that processes sensitive data that the attacker wants to obtain. This topic describes possible attacks against hosts and applications and associated protection mechanisms.

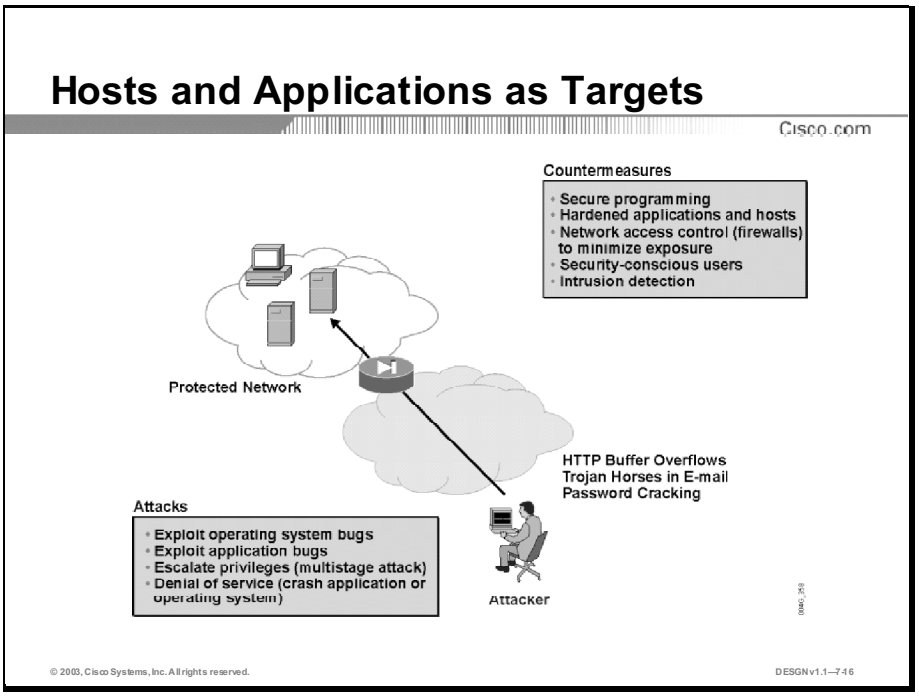

When an attacker can communicate directly with the host or the application hosting sensitive data, you must protect the host and the application against this potential threat. Attackers seek to obtain permissions to read or write to sensitive data, thereby compromising confidentiality and integrity.

Attacking Applications Directly

An attacker can attack applications directly. An attacker will find a flaw in the application and bypass its access controls to obtain read or write access to sensitive data. The complexity of current applications makes such flaws very common. In addition, secure development is too costly or not feasible for many businesses.

Privilege Escalation

In another scenario, an attacker gains access to sensitive data through a chain of compromises of other system components. For example, an attacker first obtains basic user-level access to the system on which the sensitive data resides. Then, by exploiting a flaw in any local application, the attacker attains system administration privileges (known as “privilege escalation”). Using those privileges, the attacker can read or write to most objects on the system, including sensitive data of the target application.

Example: Application Protection

Good examples of managing risks to application subversion can be found in military computer networks. Data in military systems is often labeled with their level of sensitivity (such as "secret," "top secret," and so on), and each application executes with a certain level of access ("secret," "top secret," and so on). If an application is subverted by an attacker, the operating system prevents the application from accessing data, which is more sensitive than its level of access. This enforcement is simple to implement and does not allow users to change security labels. It works well in military practice, because it adheres well to the organizational structure. However, in commercial environments, such systems break down as people and data are less hierarchically organized, and access control management is more complex and more difficult to implement securely.

Host and Application Protection Guidelines

To manage host and application risks, use multiple protection methods together to form a multilayered security system. A multilayered security system does not rely on a single security mechanism to perform a function, because that mechanism might be compromised. Instead, different security mechanisms provide a similar protection function and back up each other. This applies universally to any security system and is considered good practice in a security design.

Here are some host and application protection methods:

- Network access control methods allow access only to minimal services and only to select users (for example, firewalls).
- Strong host security policies protect the operating system and its services from compromise. Such a system will resist an attacker when alternate paths to sensitive data are attempted.
- Cryptography provides confidentiality, integrity, and authenticity guarantees for data. Cryptographic methods often protect data when it is outside the application's control. For example, the data on the disk drives is encrypted and only the application can read it. More commonly, all data between the client and the server is encrypted, providing confidentiality over an insecure network.
- Application access controls, coupled with secure programming, are the most significant cornerstones of application security. Secure development is typically expensive and slow. With the creation of safer high-level languages and developers' awareness of security in programming issues, many current systems can provide high levels of application security. Most stock software is vulnerable to simple attacks to defeat its security.

Summary

This topic summarizes the key points discussed in this lesson.

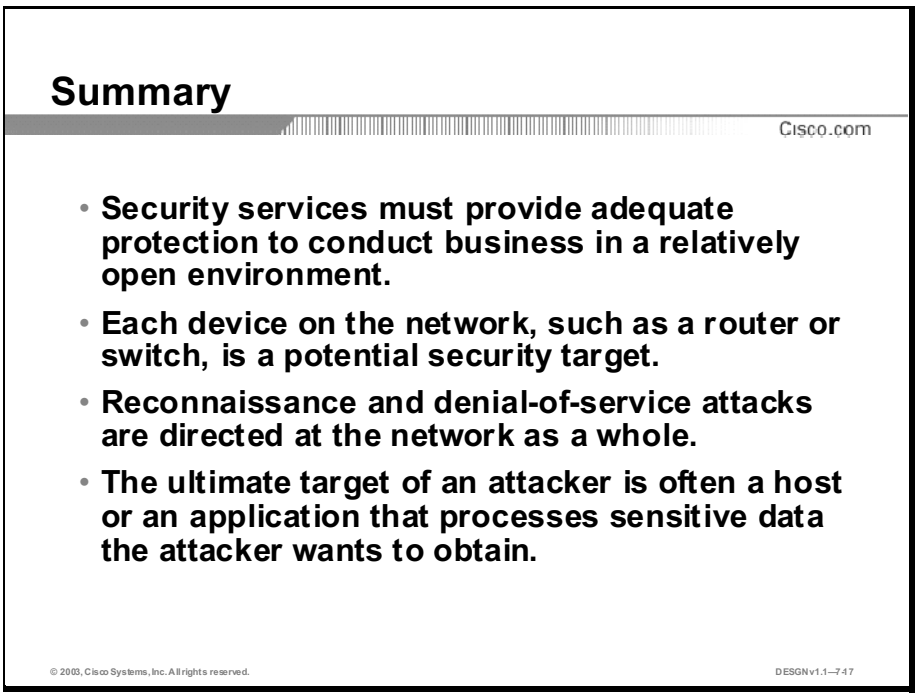

References

For additional information, refer to this resource:

- The Cisco *SAFE: A Security Blueprint for Enterprise Networks* white paper, http://www.cisco.com/go/safe.html.

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which situation best describes an integrity violation attack?

A) The attacker reads sensitive data.

B) The attacker compromises availability of a network.

C) The attacker changes sensitive data.

D) The attacker compromises availability of a host.

Q2) Which two steps can you take as countermeasures to mitigate the compromise of network devices? (Choose two.)

A) Enable only necessary services on devices.

B) Enable many routing protocols, so that at least one of them is secure.

C) Enable strong authentication to verify administrator identity.

D) Use devices from many different vendors to reduce the probability of a security issue in their software.

E) Update device software even though no security issues have been discovered.

Q3) Which two steps can you take as countermeasures to mitigate threats to entire networks?

A) Use minimal protection to ensure unobstructed data flow.

B) Rely on the ISP to protect the network from all denial-of-service attacks.

C) Give out as much information as possible to confuse the attacker.

D) Allow only minimal connectivity from external networks.

Q4) Which two steps can you take as countermeasures to mitigate the compromise of hosts and applications? (Choose two.)

A) Protect the host operating system and develop a secure application code.

B) Use well-tested, standardized software.

C) Use network access control to minimize exposure of hosts.

D) Offload all application protection to the firewall to increase application performance.

E) Use low-level programming languages to decrease the number of flaws.

Quiz Answer Key

Q1) C

Relates to: Security as a Network Service in Modular Network Design

Q2) A, C

Relates to: Network Devices (Routers and Switches) as Targets

Q3) D

Relates to: Networks as Targets

Q4) A, C

Relates to: Hosts and Applications as Targets

Identifying Security Mechanisms for a Defined Security Policy

Overview

A security policy is a critical component of network security. You will define the security policy during the initial gathering of requirements. This lesson describes possible security mechanisms and their impact on network design, operation, and performance.

Relevance

By knowing about the security mechanisms available to you, you will be able to develop and implement an effective security design.

Objectives

Upon completing this lesson, you will be able to select the appropriate security mechanisms to counter specific threats and comply with an enterprise security policy. This includes being able to meet these objectives:

- List the major components of an enterprise security policy
- Explain the mechanisms for physical security
- Describe network authentication and authorization mechanisms
- List techniques that ensure transmission confidentiality
- Describe the mechanisms for ensuring data integrity
- Identify security management and reporting mechanisms
- Describe security mechanisms that can be configured with AutoSecure

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking concepts, terms, and functions, including basic functions implemented in IOS software

Outline

The outline lists the topics included in this lesson.

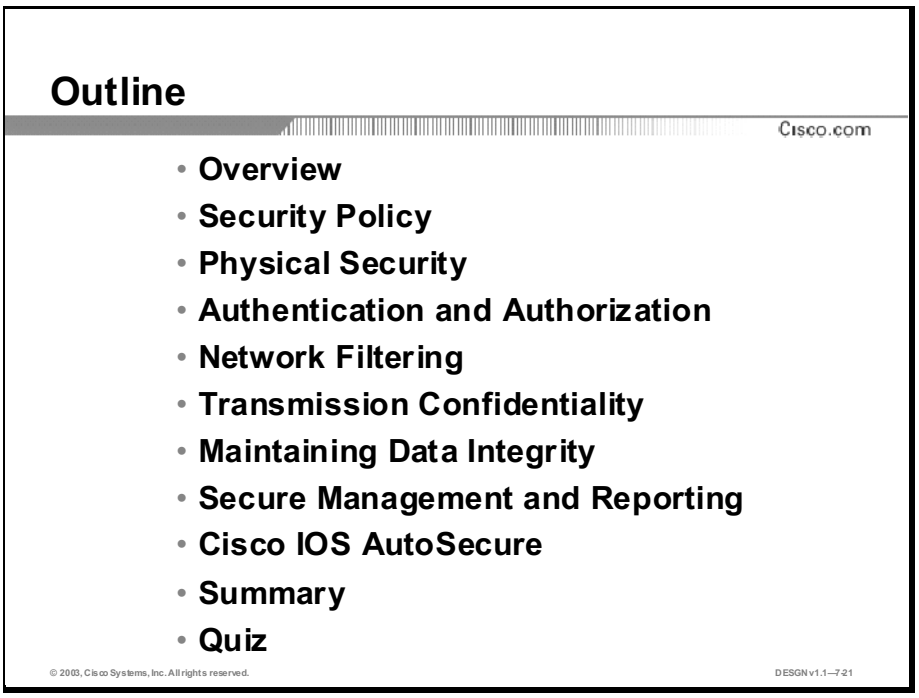

Security Policy

Network security employs risk management to lower risks to an acceptable level. A risk assessment identifies the risks relevant to a network. Risk assessment results in the development of a network security policy, which documents the level of risk and suggests the methods of managing the risk to an acceptable level. This topic describes the principles for building a security policy and explains how a network security policy reflects the risks associated with network computing.

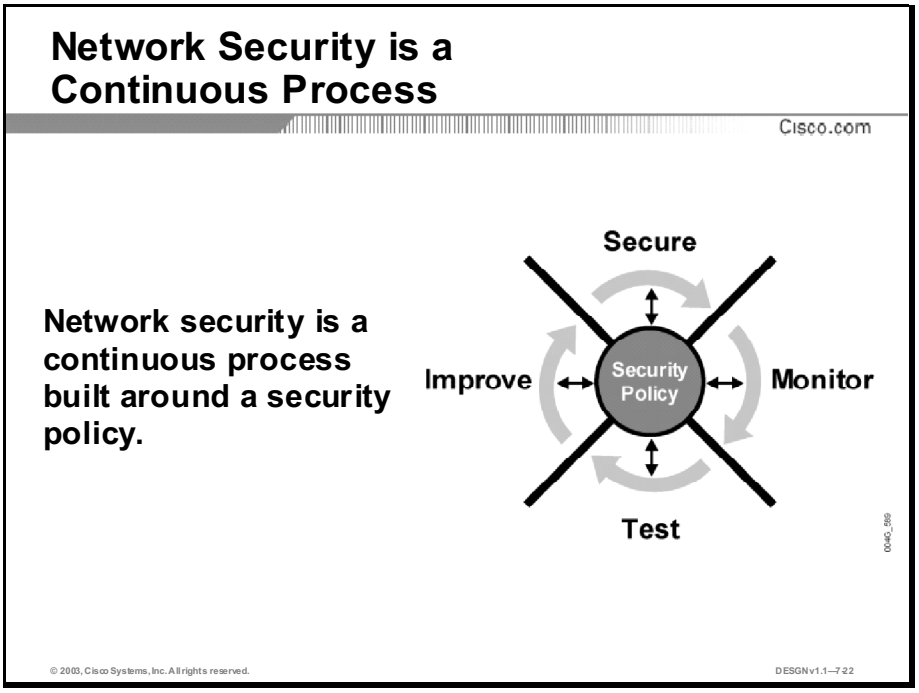

"Security Wheel"

Risks change over time and so should the security policy. The complete process of developing a security process is called the "security wheel." It comprises the major tasks of securing networks, the reevaluation of what is already secured, and the identification of new risks:

Step 1 Perform the initial risk assessment.

Step 2 Develop the initial security policy.

Step 3 Implement network security.

Step 4 Monitor security.

Step 5 Test security.

Step 6 Reassess the risks.

The security policy often changes to accommodate current risks.

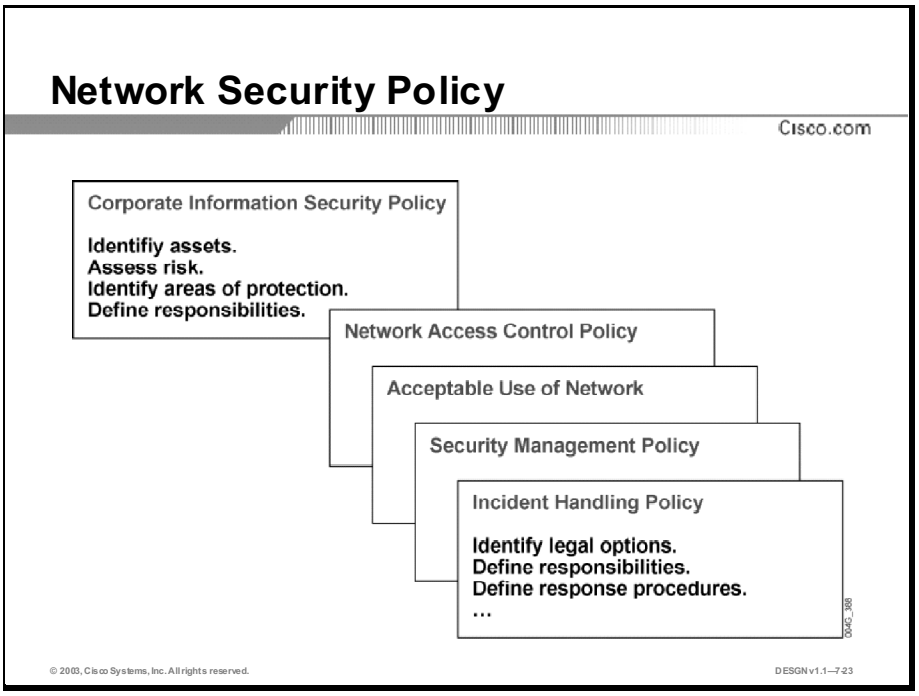

The figure illustrates the organization of a security policy and how to divide it into smaller parts that are applicable to the network segments. A general document describes the overall risk management policy, identifies the corporation's assets, and identifies where to apply protection. In addition, the document defines how responsibility for risk management is distributed throughout the enterprise.

Other documents might address more specific areas of risk management, such as these:

- **Network access control policy:** May document how data is categorized (for example, confidential, internal, top secret) and the general principles of access control implemented in the network.
- **Acceptable use of network document:** Usually written in easy-to-understand language and distributed among end users. This document informs users about their roles and responsibilities in risk management.
- **Security management policy:** May define how to perform secure computer infrastructure management.
- **Incident handling policy:** May document the procedures used to ensure reliable and acceptable emergency situation handling.

Numerous other areas may be covered in separate documents, depending on the organization's requirements.

Example: Internet Security Policy

As a part of an overall security policy, an organization wants to develop specific guidelines for Internet access. The organization's overall IT security policy states that access should be limited only to necessary services, that no user should install or run unauthorized applications, and that all transactions crossing security perimeters should be audited. As a result, the organization deploys these Internet access restrictions:

- The firewalls allow access only using HTTP and e-mail.
- Internet access to web sites should be related to business only (no personal use).
- The protected network does not allow downloading executable code from external sources.
- The network devices must log all connections to the Internet.

These restrictions are enforced using network devices, recorded in an "acceptable Internet use" document, and distributed to all employees.

Physical Security

Physical security is critical to a successful network security implementation and can significantly influence the strength of the total security design. This topic discusses physical security and its incorporation into an overall security policy.

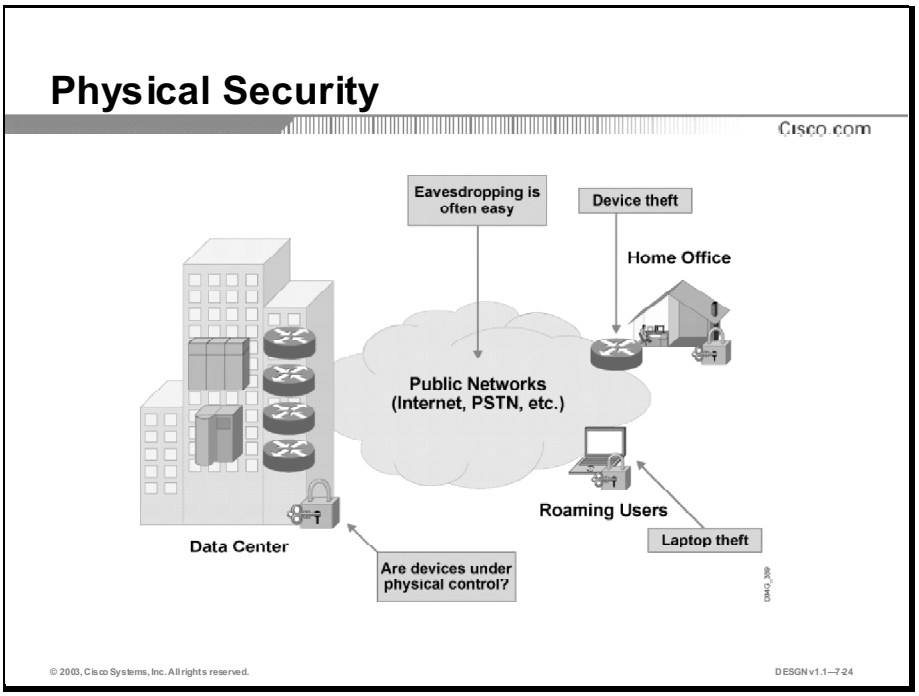

Physical security restricts physical access to a device or the communications media. A good security policy must anticipate possible physical attacks and assess their relevance in terms of possible loss, probability, and simplicity of attack.

Consider these potential physical threats:

- A network device does not always enforce all of its security settings when an attacker directly accesses the hardware through console access, memory probing, or installation of unreliable software.
- Access to the physical communication medium can allow an attacker to impersonate trusted systems and view, intercept, and change data flowing in a network.
- An attacker might use physically destructive attacks against devices and networks such as physical force, attacks over the power network, or electromagnetic surveillance and attacks.

Example: Physical Breaches of Network Security

The figure illustrates possible physical breaches of network security. An attacker might:

- Break into the computing center and obtain physical access to a firewall, compromise its physical connections to bypass it, or alter the security settings of routers and switches.
- Physically access the copper media of the corporate WAN and read or change sensitive data that is not encrypted.
- Steal a device such as a home office router or laptop computer and use it to access the corporate network.

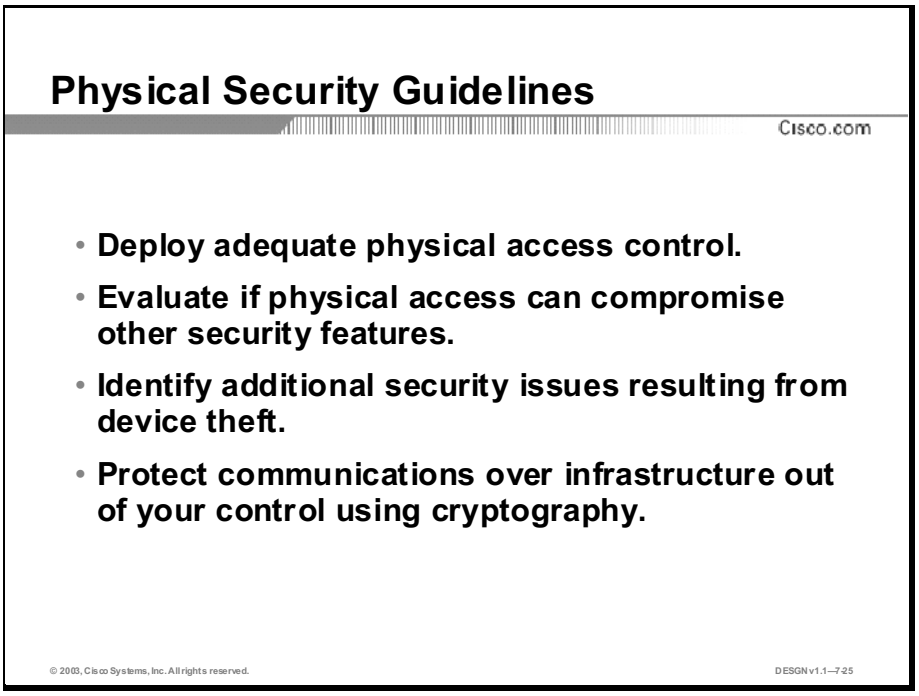

The traditional method to manage the risk of physical compromise is to deploy physical access controls using techniques such as locks or alarms. It is important to identify how a physical security breach might interact with network security mechanisms. For example, if an attacker obtains physical access to a switch port in a corporate building and acquires unrestricted access to the corporate network, there could be a significant risk. If the security policy assumes, in error, that only legitimate users can obtain such access, the attacker will be able to connect to the network without authentication, bypassing network access control.

A security designer must identify the additional consequences resulting from device theft in a secure network. For example, if a laptop computer is stolen from a roaming user, does it contain cryptographic keys that enable the attacker to connect to the enterprise network while impersonating a legitimate user? Moreover, does the network administrator have a scalable means to revoke credentials, which could be obtained by the attacker through physical theft?

Sometimes a significant portion of the network infrastructure is outside the physical control of the enterprise, so physical controls are not enforced at the media access level. For example, many enterprises rely on the fact that the physical infrastructure of the service provider's Frame Relay network is well protected, even though access to its wire conduits is easily available. To protect communications over unsafe networks, cryptography provides confidentiality and integrity protection, which are fully under the control of the enterprise.

Example: Protecting Communication over Public Infrastructure

Enterprises that simultaneously transmit sensitive and nonsensitive data over a public circuit such as an ATM or Frame Relay link can use IP Security (IPSec) protection. The sensitive traffic is protected when it is routed over the untrusted WAN, while other traffic is sent in the clear.

Another example might be a government intelligence agency, which is concerned about the theft of laptops that may have extremely sensitive data. To manage this risk, they deploy robust file encryption software that only decrypts sensitive files on special request. Sensitive information is hidden from a potential thief, who could otherwise read raw data from the laptop's disk.

Authentication and Authorization

Modern network security revolves around access control: the ability to enforce a policy that states which subjects can access which network resources. Access control aims to provide the desired confidentiality and integrity of sensitive data. Authorization mechanisms limit a subject's access to resources based on subject identity. This topic describes network authentication and authorization.

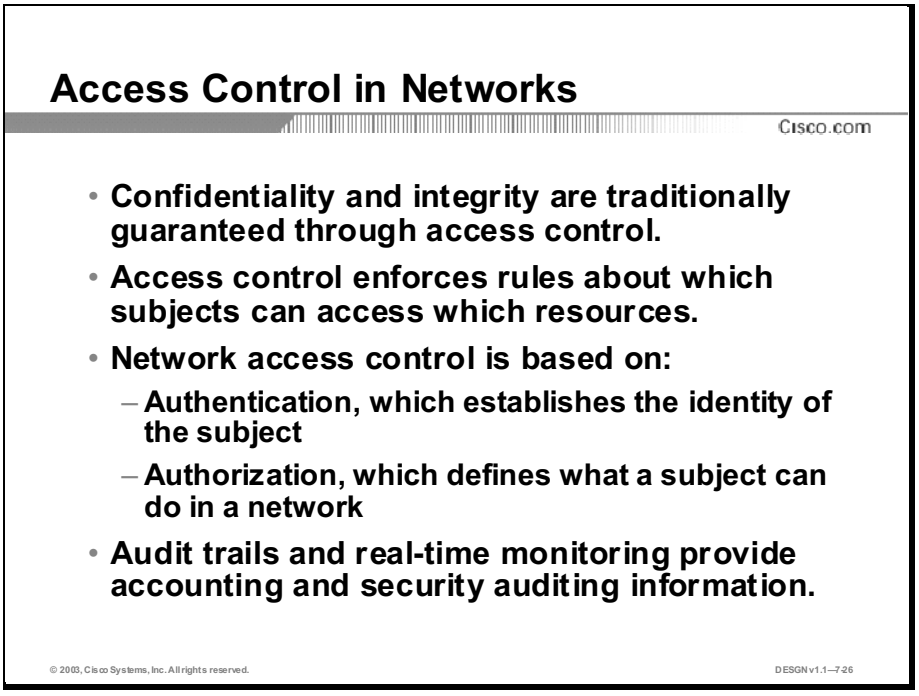

Access control mechanisms are usually classified in one of these ways:

- As authentication, which is used to establish subject identity
- As authorization, which is used to limit subject access to a network. This can define the granularity of access such as read or write.

An audit trail is needed to provide evidence of a subject's actions, and real-time monitoring provides security services such as intrusion detection.

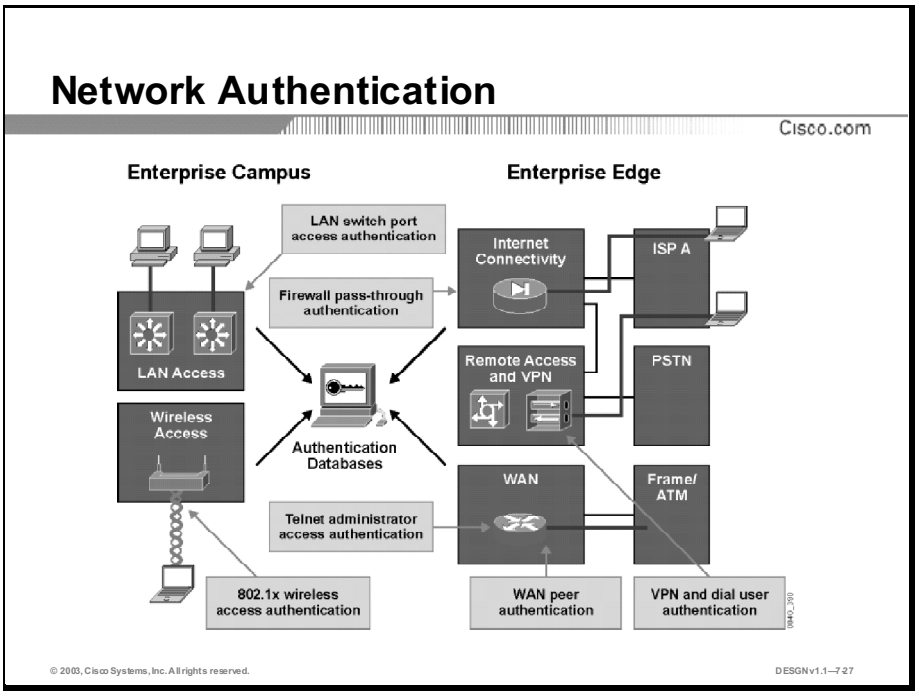

Authentication is used to establish the identity of a subject in a network. A subject can be a computer user, computer system, network device, or application. It is important to separate the concept of identification, where a subject presents its identity, and authentication, where a subject proves its identity. For example, to log on to a resource, a user is identified by a username and authenticated by a secret password.

Authentication, the proving of identity, is traditionally based on one of three proofs:

- **Something the subject knows:** Involves knowledge of a unique secret, which is usually shared by the authenticating parties. To a user, this secret appears as a classic password (PIN number) or a private cryptographic key.
- **Something the subject has:** Usually involves physical possession of an item, which is unique to the subject. Examples include password token cards, Smartcards, and hardware keys.
- **Something the subject is:** Involves verification of a unique physical characteristic of the subject, such as a fingerprint, retina pattern, voice, or face.

To achieve high assurance in authentication, many trusted systems require “two-factor authentication,” where a subject provides at least two types of proof of identity. An example is an access control system, which requires a Smartcard and a password. With two-factor authentication, a compromise of one factor does not lead to a compromise of the system. A password may become known, but it is useless without the Smartcard. Conversely, if the Smartcard is stolen, the thief cannot use it without the password.

Example: Using Network Authentication

These are examples of authentication:

- **Dial-up access points:** Any subject can establish a dial connection to the network and authentication is needed to distinguish between trusted and untrusted subjects. Well-known authentication protocols include Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and Extensible Authentication Protocol (EAP). PAP sends the secret password over the link in clear text, so use it only over trusted links or with One Time Passwords (OTPs), CHAP uses a challenge-response mechanism, which defeats password snooping on untrusted links, EAP is the next-generation authentication protocol, which supports many authentication methods carried within one authentication framework.
- **WAN and Virtual Private Network (VPN) infrastructure:** Network devices can authenticate each other on WAN or VPN links, mitigating the risk of infrastructure compromise or misconfiguration. WAN peer authentication involves PPP mechanisms and routing protocol authentication. In a VPN, authentication is embedded in the VPN security protocols, most often IPSec and Internet Key Exchange (IKE).
- **LAN access:** A network device (switch) authenticates the user before allowing access to the switched network. A standardized LAN authentication protocol is IEEE 802.1x.
- **Wireless access:** Only authenticated users can establish an association with a wireless access point (IEEE 802.1x).
- **Firewall authentication:** Users must prove their identity when entering a sensitive network protected by a firewall.

A network will authenticate management and administrative access to its infrastructure. Telnet, Secure Shell Protocol (SSH), and SNMP are examples of management protocols that offer authentication capabilities from simple cleartext passwords (Telnet, SNMPv2) to digital signatures and secure fingerprints (SSH, SNMPv3).

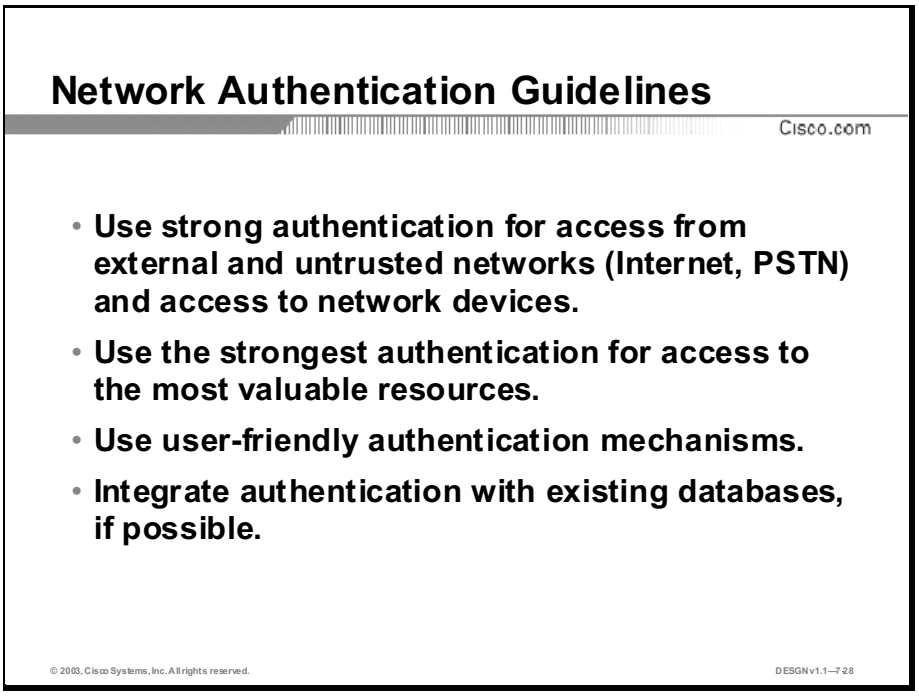

To authenticate access to very sensitive resources, stronger authentication is required. Examples include direct access to the corporate network from the Internet or terminal access to network devices. In these instances, an enterprise might use two-factor authentication for its users.

Authentication should also be user friendly, so that users do not willingly compromise it. An example of an unfriendly method is a system that enforces a user password with extremely restrictive rules about password randomness, and the users do not remember their passwords. The users then write down their passwords, which may be stolen. Alternatively, you can implement an OTP generator (token card), which displays the current password to the user, does not rely on the user remembering anything, and generates random passwords. Although such a system may be simple, it can require a significant investment in its technology and operations. To lower authentication costs, some enterprises standardize a small number of authentication databases and methods and reuse them for several systems.

Example: Network Authentication Over a VPN

An organization needs to deploy remote-access services to its network from the Internet. It has implemented remote-access VPN technology and requires proper user authentication before entering the protected network. The organization has had negative experiences using normal passwords and wants to deploy a very secure, yet simple system. OTP generators for remote users may be the ideal solution, because they are secure and simple to use.

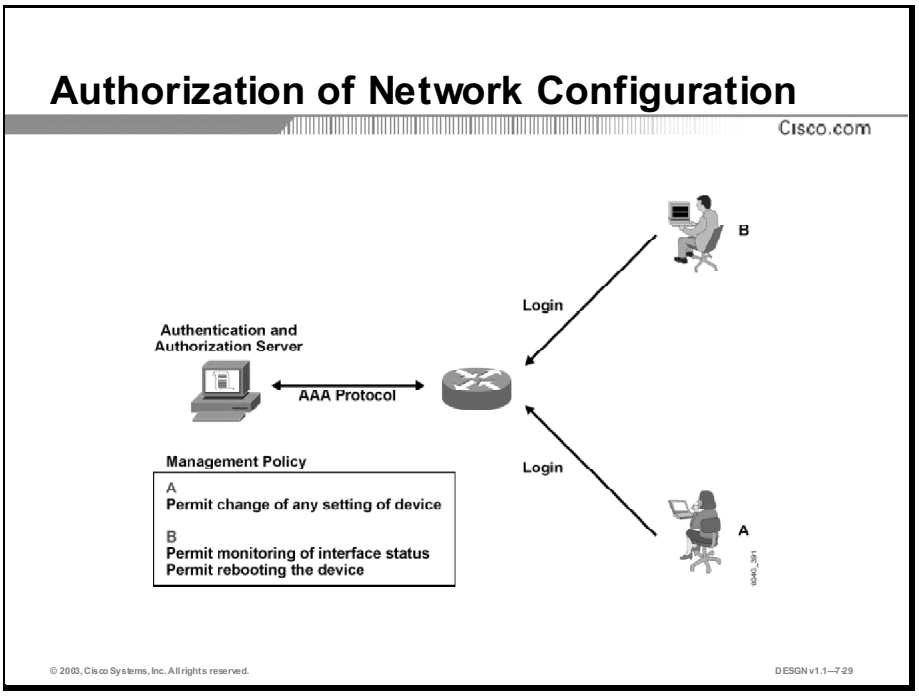

Most networks implement authorization using general access control lists, which define subjects and their access rights for each resource. An example is a router packet filtering access control list (ACL) that specifies which clients can connect to a sensitive server in a network. Another example is when network administrators are restricted to perform only a subset of configuration commands on a router.

In terms of network security, authorization mechanisms on a network include two categories:

- **Change to network infrastructure authorization:** Network devices are protected against unauthorized change of their settings by their operating-system authorization controls.
- **Access to network resources authorization:** The network controls a subject's network settings (such as the IP address) and limits its connectivity.

Example: Device Configuration Authorization

The figure illustrates an example of device configuration authorization. When multiple administrators manage a network, it is good practice to limit their rights to include only the permissions necessary to perform their tasks. In security, this is often called the "least privilege" concept. Some very trusted administrators may need full access to a device, while other administrators may only require monitoring rights and rights to resolve common problems.

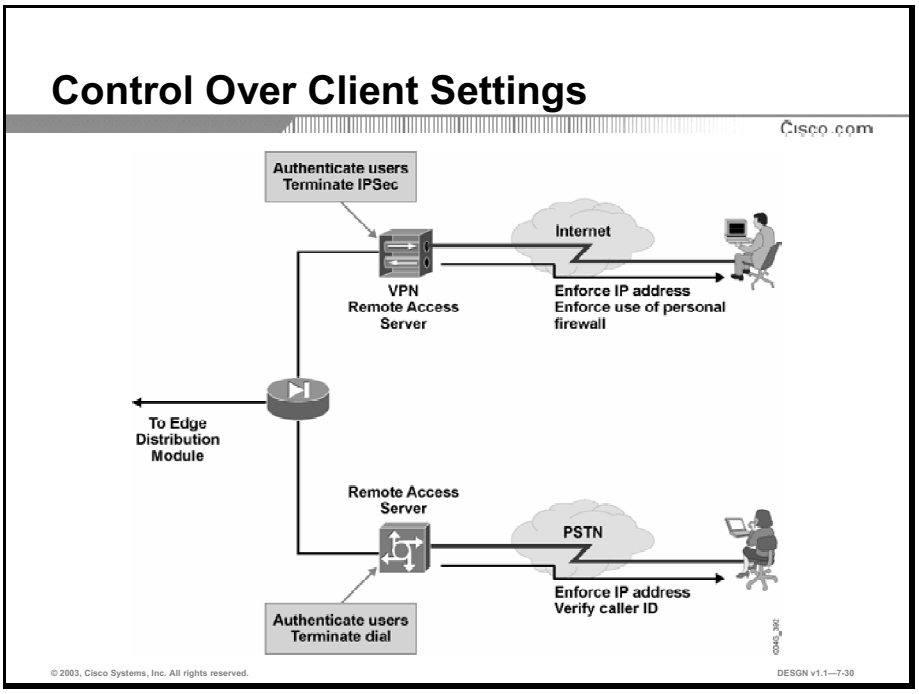

The figure illustrates how a network enforces client settings to simplify access control and manage the threat of identity spoofing. The access server enforces a particular IP address to the remote user and verifies the user's Caller ID address. In the remote-access VPN setup, the VPN concentrator configures the remote client with a particular IP address and enforces use of a personal firewall on the client system. Other access control mechanisms, such as firewalls in the enterprise network, use the assigned IP addresses to control the subject's access rights.

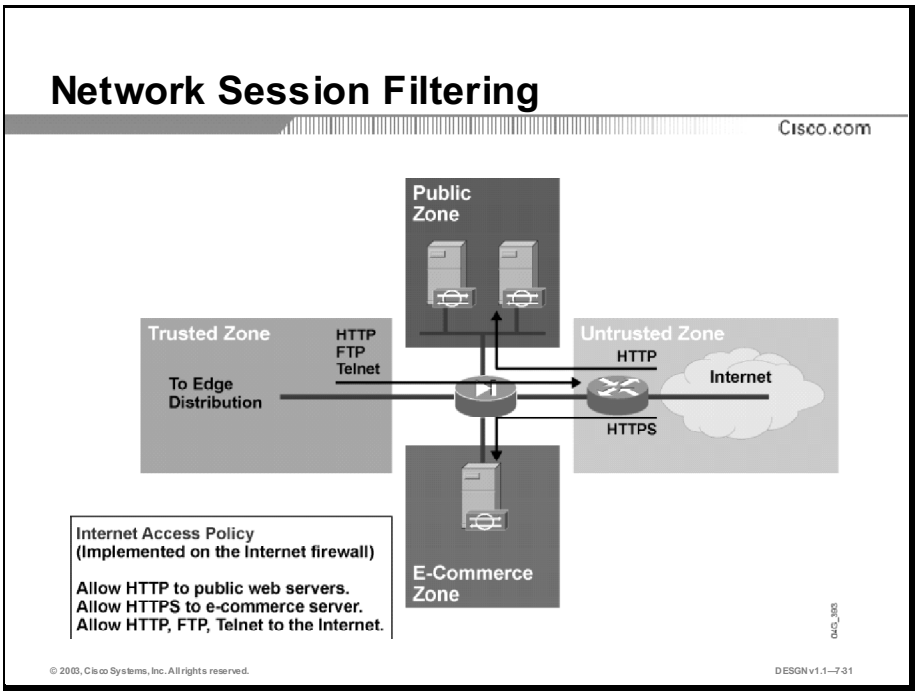

The figure illustrates the use of a network firewall to control access, which is a common use of network authorization. An enterprise network is usually divided into separate security perimeters or zones such as the untrusted Internet, trusted enterprise campus, perimeters of public and semipublic servers, and so on. Because all traffic must pass through the network firewall, it enforces the access and authorization policy in the network effectively by specifying which connections are permitted or denied between security perimeters.

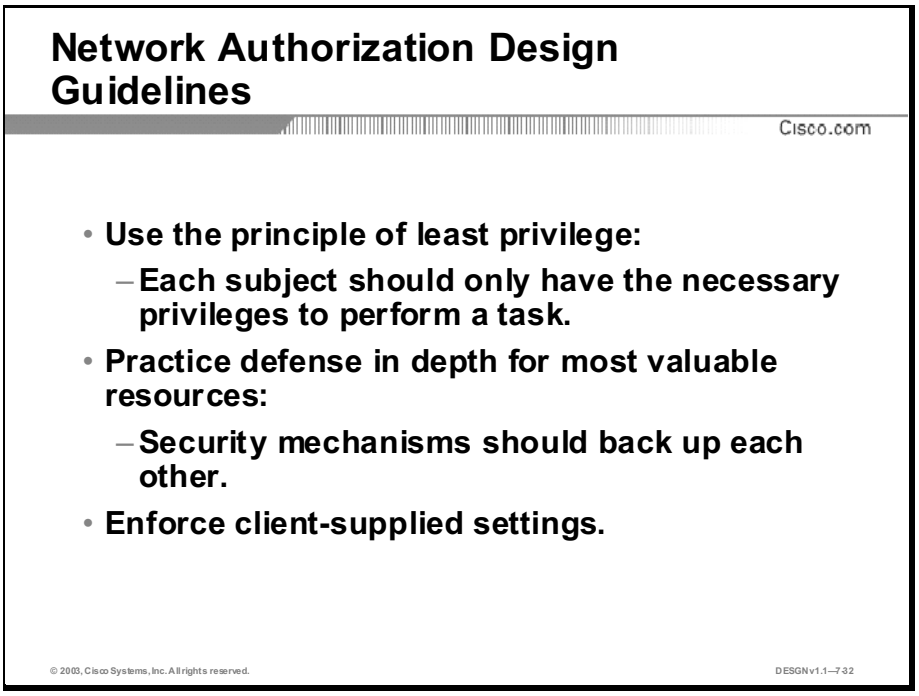

Use these guidelines in the design and implementation of access authorization:

- **Principle of least privilege:** This principle is based on the practice where each subject is given only the minimal necessary rights to perform a certain task. Therefore, an authorization mechanism should give a subject only the minimum access permissions. For example, if a user needs to access a particular web server, the firewall should allow that user to access only the specified web server. In reality, enterprises often introduce lenient rules that allow subjects more access than required. This may result in deliberate or accidental breaches of confidentiality and integrity.
- **Principle of defense in depth:** This principle suggests that security mechanisms should be fault tolerant. A security mechanism should have a backup or supplemental security mechanism. An example is to use a dedicated firewall to limit access to resources at a granular level and then to replicate with a failover firewall. You can supplement the dedicated firewall with perimeter routers that use access control lists to filter packets.
- **Enforce client-supplied settings:** This principle suggests that you need to control all networking parameters centrally. For example, the client need not configure a proper IP address when dialing in remotely. The access server should push those settings to the client.

Example: Security Policy for Data Through the Internet

An organization needs to provide some data through the Internet for its partners. The security policy requires proper authorization to access that data. An Internet application, running on a general-purpose server, needs an Internet connection. Implement these network authorization mechanisms to ensure protection:

- The firewall system filters all traffic to the server, except the minimal necessary access (concept of least privilege).
- The firewall does not allow any connection originating from the server to contain the attacker in the event that a server break-in occurs.
- The access router protects the server with additional filtering that supplements the firewall (defense in depth).

Beyond network authorization, other application-specific protection mechanisms serve as additional protection for sensitive data on the server.

Transmission Confidentiality

Transmission confidentiality protects data while it is transported over insecure networks. This topic discusses transmission confidentiality methods.

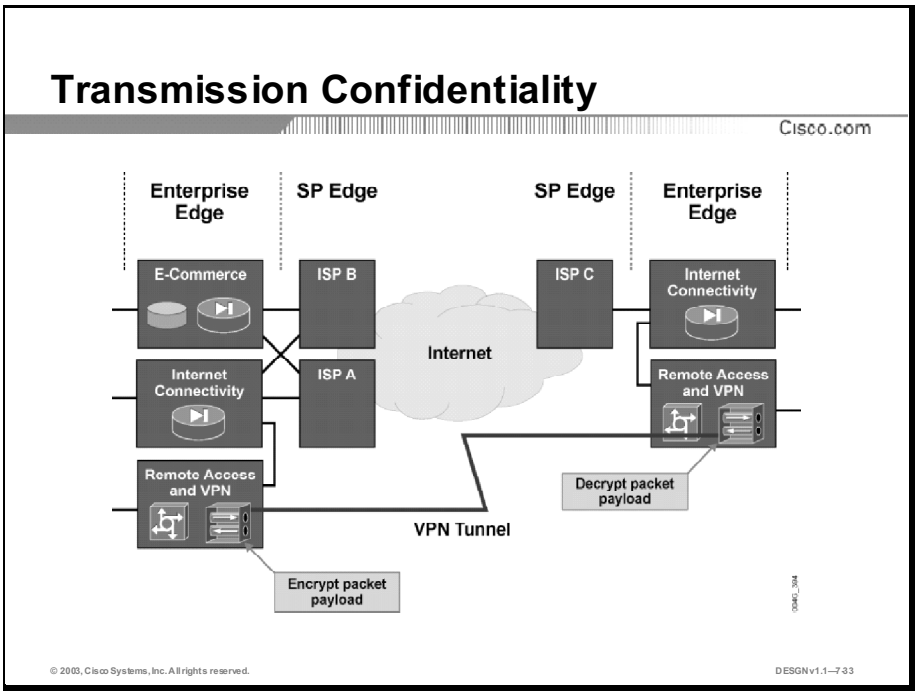

When connecting trusted and untrusted networks, for example, when connecting a corporate network with the Internet, data may be transmitted among trusted subjects over untrusted networks. Untrusted networks do not support classic access control mechanisms, because a corporation does not have control over the users and network resources in an untrusted network. Therefore, you need to protect data in transmission to ensure that no one in the untrusted network can view or change the transmitted data.

Encryption

Cryptography provides confidentiality through encryption. Encryption disguises a message to hide its original content. With encryption, plaintext (the readable message) is converted to ciphertext (the unreadable, disguised message), and decryption reverses the process. The purpose of encryption is to guarantee confidentiality. Only authorized entities can encrypt and decrypt data. With most modern algorithms, successful encryption and decryption require knowledge of the appropriate cryptographic keys. An example of data encryption is the use of encryption algorithms to hide the IP packet payloads using IPSec security protocols.

Example: Transmission Confidentiality

The figure shows two sites connected over an untrusted network: the Internet. To provide data confidentiality, a VPN technology, which supports encryption, creates a secured point-to-point association between the sites over the Internet. All packets leaving one site are encrypted, forwarded onto the untrusted network, and decrypted by a device on the remote site. Anyone eavesdropping on the untrusted network should not be able to decrypt the packet payloads and thereby read sensitive data.

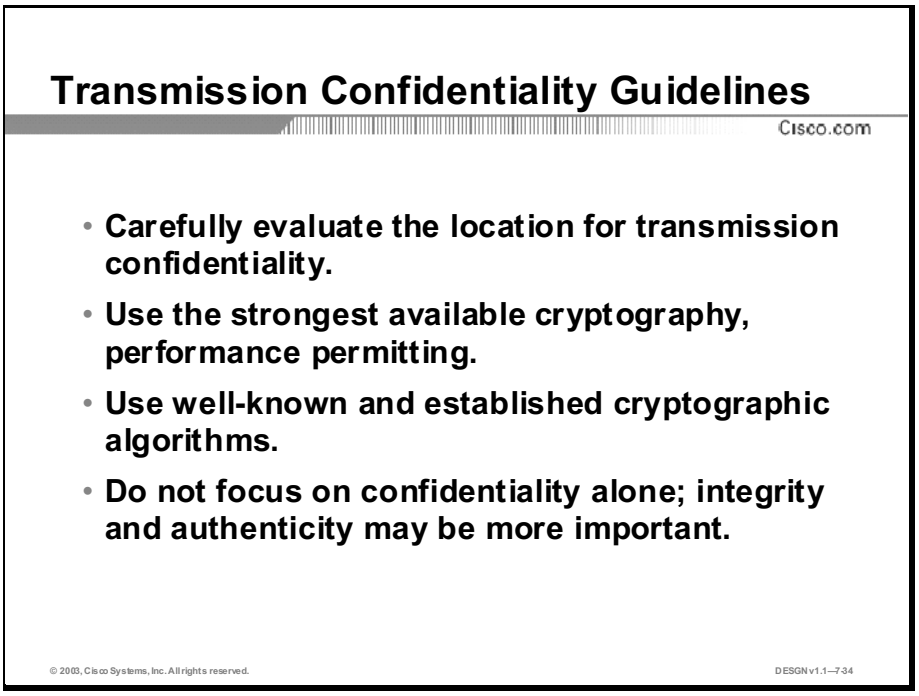

Specific cryptography guidelines to consider when designing and implementing a solution for transmission confidentiality include:

- Cryptography can become a performance bottleneck. You should perform a careful analysis to determine where to protect data. In general, if confidential or sensitive data travels over a network where an attacker could easily intercept communications (such as a network outside physical control or a network where device compromises are likely), then communications need to be protected as defined in the security policy.
- Companies can now export modern cryptography. Therefore, use the strongest available cryptography to provide sufficient protection. Take care because some cryptographic algorithms allow you to specify extremely long key lengths, which may not provide worthwhile confidentiality improvements.
- Only use well-known cryptographic algorithms, because only well-known algorithms are tested, analyzed, and considered trustworthy.
- Do not forget that encryption only provides confidentiality, and most organizations view data integrity and authenticity as equally important elements of security. If possible, use both confidentiality- and integrity-guaranteeing cryptographic algorithms.

Example: Providing Confidentiality Over the Internet

To lower communication costs, a health insurance company decides to connect some of its branch offices to its headquarters over the Internet. Because it needs to protect patient record confidentiality and attackers on the Internet are able to intercept communications, the company implements a VPN using the strongest possible encryption algorithms to guarantee data confidentiality. In the event of interception, it is unlikely that the attacker can decrypt the messages that are protected with as modern cryptographic algorithms, such as Triple Data Encryption Standard (3DES), Advanced Encryption Standard (AES), or RC4.

Maintaining Data Integrity

Cryptography provides integrity mechanisms, which can protect data in transmission over untrusted networks. Cryptographic protocols, such as secure fingerprints and digital signatures, can detect any violation of sensitive data integrity. This topic discusses selecting transmission integrity methods for security design.

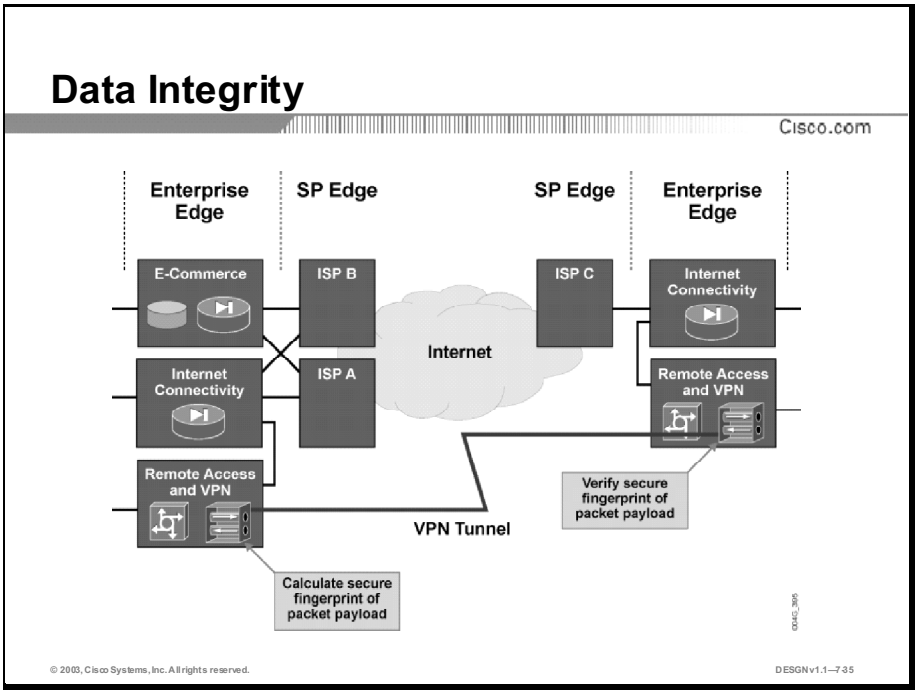

Secure fingerprints attach a cryptographically strong checksum to data, which is generated and verified using a secret key, known only to authorized subjects. By checking this checksum when needed, an authorized subject can verify the integrity of data. For example, a method of secure fingerprints, known as a Hash-based Message Authentication Code (HMAC) is implemented in the IPSec standards to provide packet integrity and authenticity in IP networks. The HMAC method is very fast and suitable for real-time traffic integrity and authentication.

Digital signing uses a similar cryptography method and attaches a digital signature to sensitive data. A unique signature generation key, known to one signer only, generates the signature. Other parties use the signer's signature verification key to verify the signature. The cryptography behind digital signing guarantees data authenticity and accuracy because the originator signed it. In the financial world, digital signatures also provide nonrepudiation of transactions, where a subject can prove to a third party that a transaction has occurred. Digital signature protocols are based on public-key cryptography and are not used for bulk protection because of their performance limitations.

Example: VPN Tunneling for Data Integrity

The figure illustrates a connection between two network sites over the Internet. To provide data integrity, a VPN technology supporting secure fingerprinting creates a secured point-to-point association over the Internet. All packets leaving one site are imprinted with a secure digital fingerprint (similar to a very strong checksum), which uniquely identifies the data at the sender side. The packets are forwarded onto the untrusted network, and a device on the remote site verifies the secure fingerprint to ensure that no one has tampered with the packet. Anyone eavesdropping on the untrusted network should not be able to change the packet payloads and thereby change sensitive data without being detected.

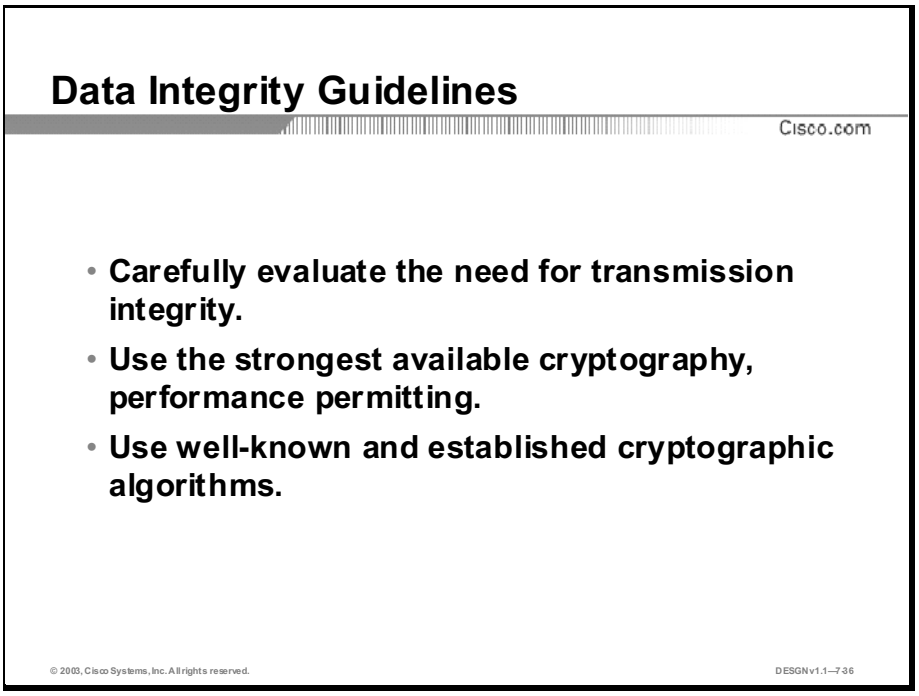

Integrity cryptographic mechanisms can apply the same guidelines and warnings as confidentiality mechanisms. Guidelines include:

- Evaluate the need for data integrity carefully, and enforce it only where justified by potential threats.
- Use the strongest available mechanisms for data integrity, taking the performance effects into account.
- Use only established and well-known cryptographic algorithms.

Example: Implementation of Digital Signatures

An organization needs to exchange stock market data over the Internet. Confidentiality is not its main concern. The main risk lies in an attacker changing data in transit and presenting false stock market data to the organization. The organization's preferred data exchange application is e-mail. Therefore, it decides to implement digital signatures of all e-mail messages when exchanging data over the Internet among partners.

Secure Management and Reporting

When network and security management authenticate administrator access to devices and provide authorization of configuration changes, it can apply many of the security mechanisms of authentication and authorization. You should consider how to perform secure management and reporting over the existing network. This topic discusses secure management and reporting.

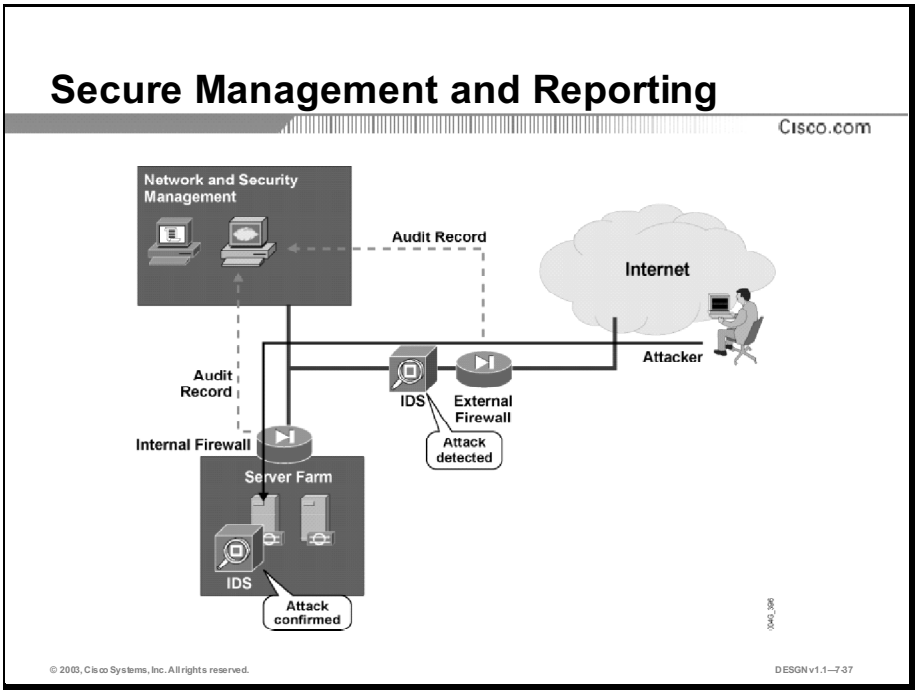

A major task of security management is to create and analyze an audit trail, which tracks user and system activity in a network. An audit trail consists of:

- Logging data, where systems and applications audit exceptional events for analysis
- Accounting data, where systems and applications audit user actions

In a network, the resulting audit trail records are examined for various purposes:

- To detect and respond to unauthorized actions, which endanger sensitive data (security monitoring)
- To identify new risks
- To collect usage statistics for the purpose of contingency planning and billing

Detecting and responding to unauthorized actions are primary security-related tasks. You should compare information about unauthorized actions to the risk profile defined in the security policy. Such monitoring is called intrusion detection, and is performed in real time or by using offline analysis.

Intrusion detection systems can range from basic (displaying attack statistics) to very complex (correlating events from multiple sources over a long period of time). Modern intrusion detection systems are:

- **Host intrusion protection systems (HIPS):** Installed on network servers; provide detection and protection against attacks within the host operating system. HIPS perform these tasks:
  - Proactively stop attacks before damage occurs
  - Proactively stop new and unknown attacks
  - Reduce the administrative burden associated with signatures and reactive alerts
- **Network intrusion detection systems (NIDS):** Installed on the network, where they capture and analyze network traffic for attacks. In Cisco Systems' case, a NIDS can be a sensor "appliance" that is installed passively on the network so it does not impede or introduce any delay or overhead on the network traffic itself. It can also be a feature of Cisco IOS, a PIX firewall, or even a card in routers and switches.

Example: Secure Management Protocols and Monitoring

The figure illustrates some secure management protocols and monitoring systems. Device and policy management systems are used to configure security devices according to appropriate policies.

An attacker on the Internet attempts to connect to a server in the enterprise network and copy some confidential data. The Internet firewall permits the connection, and the syslog protocol sends a log of the connection to the central logging server. The IDS behind the Internet firewall detects an attempted breach of security and reports it to its management station. In addition, the external IDS might respond to the event, and perhaps stop the attacker immediately. The intranet firewall permits the connection and sends a syslog audit record to the log repository. The IDS running on the target host confirms the security breach and responds by disallowing further access to the attacker. The security manager receives several alerts, correlates them, assesses the damage, and prepares a structured response to the attack.

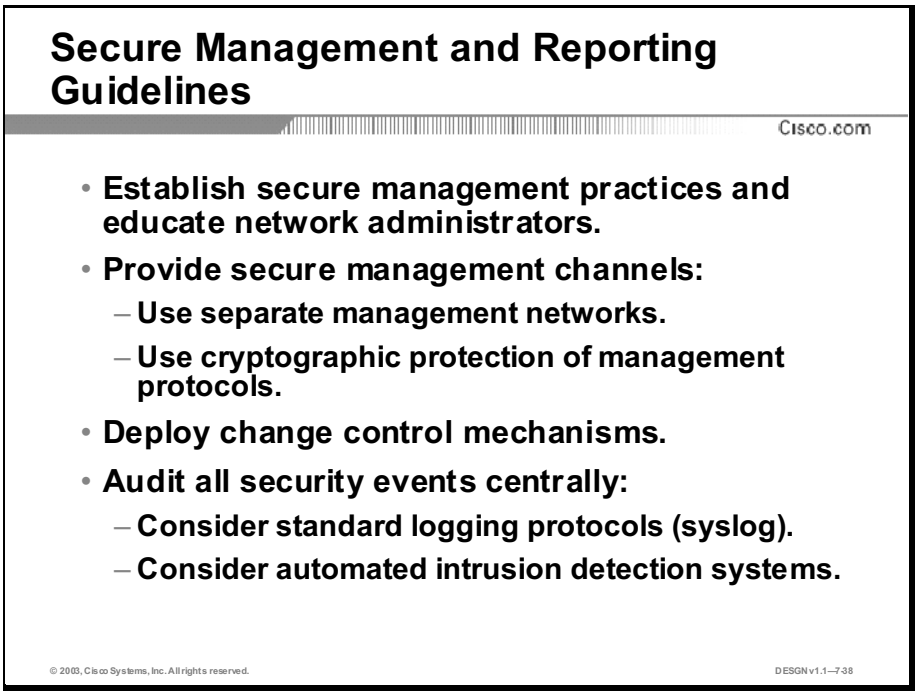

Follow these guidelines to design and implement secure management and reporting:

- Management personnel must be aware of secure management techniques and operations. They must know how to implement change control and monitoring, to prevent accidental damage.
- Network management and security management channels must be secure. This is usually accomplished by using a separate management network (for example, a separate VLAN) or cryptography to protect management protocols (for example, SSH for terminal access or IPSec protocols to protect SNMP traffic).
- Implement change control must be at the configuration and software level. Organizations must audit and version all device configurations and device software according to corporate standards.
- Centrally store and analyze the audit results, including logs and traps. In addition to classic logging using the syslog protocol, IDSs can provide automatic correlation and in-depth visibility into complex security events. This saves administrators a considerable amount of time.
- Deploy IDSs where the best security visibility is required, which is usually on the most important network segments (high-value servers, external connections) and hosts. Host and network intrusion detection should complement each other to provide the best probability of detection and allow the security administrator to respond to an incident quickly and reliably.

Example: Secure Management and Configuration Control

An organization needs to manage a remote router at its branch office. The security policy requires that the management channel is secured against eavesdropping and that strong authentication is required for managers. A local branch office administrator needs limited access to the router to check its interface status. The network solution is to design a secure management practice of only using encrypted terminal sessions (SSH) to the router, using OTPs to authenticate administrators, and configuration control to give different levels of access to different administrators.

Cisco IOS AutoSecure

Cisco IOS AutoSecure is a command-line interface (CLI)-based feature that provides fast router lockdown, easily disabling non-essential network services and enforcing secure access and forwarding. This topic describes AutoSecure.

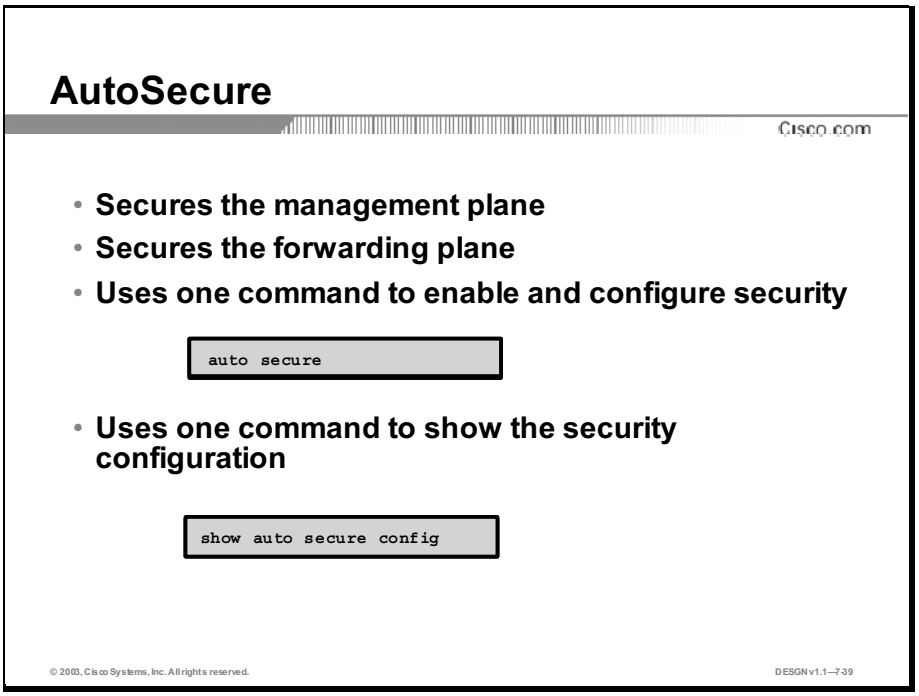

AutoSecure allows administrators to quickly secure the network without a thorough knowledge of Cisco IOS security features. AutoSecure creates a new CLI that automates the configuration of security features and disables certain features enabled by default that could be exploited for security holes. AutoSecure secures both the management plane and the forwarding plane.

Note: Visit Cisco's web site at http://www.cisco.com for current features and Cisco IOS version support.

AutoSecure secures the management plane by turning off certain global and interface services that could be exploited for security attacks and turning on global services that help mitigate the threat of attacks. Secure access and secure logging are also configured for the router. Administrators can secure the management plane of a router by using AutoSecure to accomplish these tasks:

- **Disable global services:** The following global services are disabled on the router without prompting the user:
  - Finger
  - Packet assembler/disassembler (PAD)
  - Small servers
  - Bootstrap Protocol (BOOTP) server
  - HTTP server

- Identification service
- Cisco Discovery Protocol (CDP)
- Network Time Protocol (NTP)
- Source routing

- **Disable per interface services:** The following per interface services are disabled on the router without prompting the user:
  - Internet Control Message Protocol (ICMP) redirects
  - ICMP unreachables
  - ICMP mask reply messages
  - Proxy Address Resoultion Protocol (ARP)
  - Directed broadcast
  - Maintenance Operations Protocol (MOP) service
- **Enable global services:** The following global services are enabled on the router without prompting the user:
  - The **service password-encryption** command
  - The **service tcp-keepalives-in** and **service tcp-keepalives-out** commands
- **Secure access to the router:** The following options to secure access to the router are available:
  - If a text banner does not exist, the user is prompted to add a banner.
  - The login and password are configured on the console, AUX, vty, and tty lines. The **transport input** and **transport output** commands are also configured on all of these lines. The **exec-timeout** command is configured on the console and AUX as 10.
  - When the image on the device is a crypto image, AutoSecure enables Secure Shell Protocol (SSH) and Secure Copy (SCP) for access and file transfer to and from the router.
  - If the AutoSecure user specifies that their device does not use Simple Network Management Protocol (SNMP), SNMP will be disabled if the community string is “public” or “private.” If AutoSecure is being run in interactive mode, the user is asked whether to disable SNMP regardless of the values of the community strings, which act like passwords to regulate access to the agent on the router.
  - If authentication, authorization, and accounting (AAA) is not configured, AutoSecure will prompt users to configure a local user name and password on the router.
- **Log for security:** The following logging options allow you to identify and respond to security incidents:
  - Sequence numbers and time stamps for all debug and log messages.
  - The **logging console critical** command sends system logging (syslog) messages to all available TTY lines and limits messages based on severity.
  - The **logging buffered** command copies logging messages to an internal buffer and limits messages logged to the buffer based on severity.

- The **logging trap debugging** command allows all commands with a severity higher than debugging to be sent to the logging server.

To minimize the risk of attacks on the router forwarding plane, AutoSecure provides the following functions:

- **Cisco Express Forwarding (CEF)**: Enables CEF or distributed CEF (dCEF) on the router whenever possible.
- **Access lists:** Configures the following named access lists used to prevent antispoofing source addresses:
  - autosec_iana_reserved_block: Address blocks reserved by IANA.
  - autosec_private_block: Private addresses that should not be used over the Internet.
  - autosec_complete_bogon: Combination of the autosec_iana_reserved_block, autosec_private_block, and additional addresses that are not permitted as source addresses.
- **TCP intercept:** If available, TCP intercept can be configured for connection timeout.
- **Unicast Reverse Path Forwarding (uRPF):** If available, strict uRPF can be configured to help mitigate problems that are caused by the introduction of forged (spoofed) IP source addresses.
- **Context-based access control (CBAC):** If the router is being used as a firewall, it can be configured for CBAC on public interfaces that are facing the Internet.

Summary

This topic summarizes the key points discussed in this lesson.

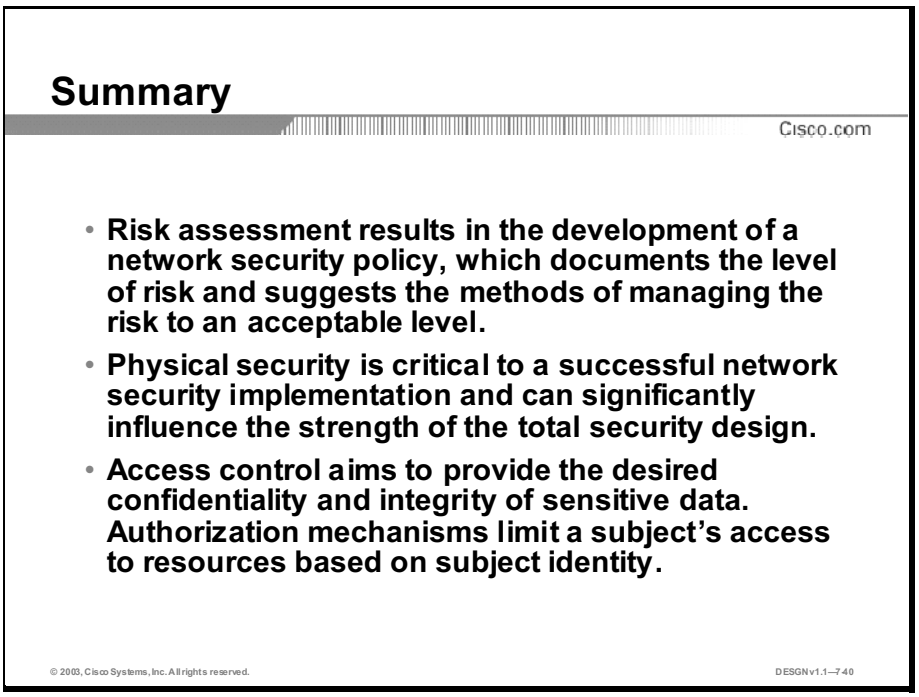

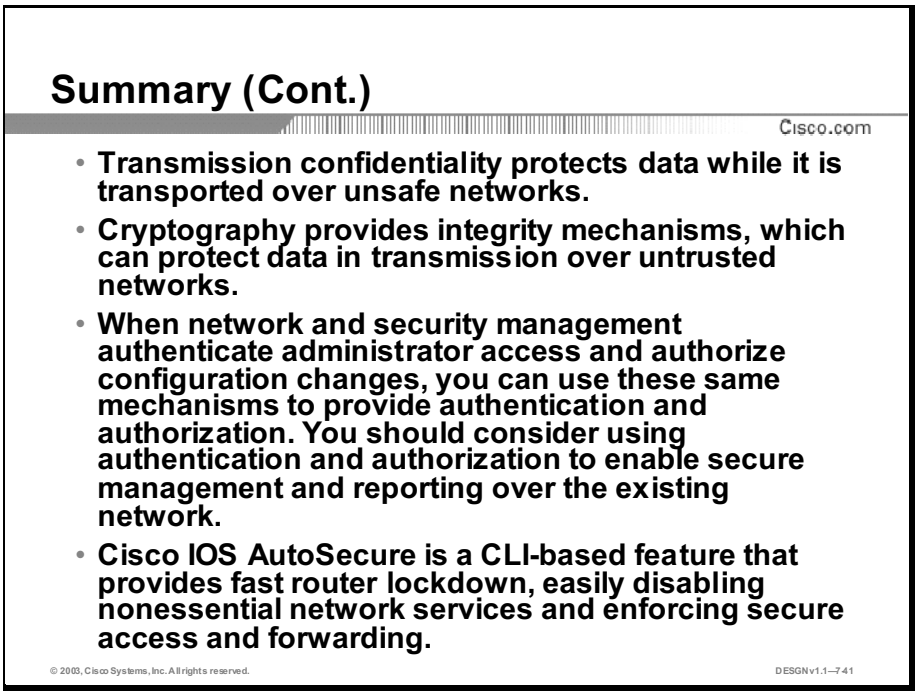

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers follow in the Quiz Answer Key.

Q1) Which three are components of a typical security policy? (Choose three.)

A) an “acceptable use” policy

B) a compilation of all local laws relating to computer intrusions

C) in-depth descriptions of individual network device configurations

D) incident handling guidelines

E) an access control policy

Q2) In which case does a network security mechanism mitigate the risk associated with a physical attack?

A) when a router containing clear text administrator passwords is stolen

B) when an attacker obtains physical access to the leased line, where only encrypted traffic is routed

C) when an attacker obtains physical access to a firewall

D) when the network management station is stolen

Q3) Place the authentication mechanisms in order of lowest to highest authentication strength.

_____ 1. plain passwords

_____ 2. a door, which needs a key, a password, and a fingerprint scan to unlock

_____ 3. token cards, which generate passwords and require a memorized PIN code to unlock

Q4) Which two attacks can you prevent using network filtering, such as a firewall? (Choose two.)

A) An attacker, who has a legitimate account on a UNIX server, uses locally available tools to obtain administrator privileges.

B) An attacker attempts to connect a sensitive nonpublic server of an organization to the Internet.

C) An attacker steals a bank ATM machine to obtain its cryptographic keys.

D) An attacker maps a company’s network using network management tools.

Q5) Which cryptographic mechanism provides transmission confidentiality for data?

A) encryption

B) digital signatures

C) key exchange

D) authentication

E) secure fingerprints

Q6) Which two cryptographic mechanisms provide transmission integrity for data? (Choose two.)

A) encryption

B) digital signatures

C) key exchange

D) authentication

E) secure fingerprints

Q7) Which two protection properties do secure management and auditing support? (Choose two.)

A) protection against device compromise over management protocols

B) protection against device theft

C) secure technology that eliminates human error

D) data about events, which can be used to detect intrusions

E) fault tolerance

Quiz Answer Key

Q1) A, D, E

Relates to: Security Policy

Q2) B

Relates to: Physical Security

Q3) Correct order: 1, 3, 2

Relates to: Authentication and Authorization

Q4) B, D

Relates to: Authentication and Authorization

Q5) A

Relates to: Transmission Confidentiality

Q6) B, E

Relates to: Maintaining Data Integrity

Q7) A, D

Relates to: Secure Management and Reporting

Selecting Security Solutions Within Network Modules

Overview

The Cisco SAFE Blueprint provides information and best practices on how to design and implement secure networks. The SAFE Blueprint identifies security threats in each network module and assesses possible security solutions. The emphasis is on external threats: network perimeter security (for example, the E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules), as well as on the Server Farm and Network Management modules of the Enterprise Composite Network Model. This lesson identifies and positions security solutions within the context of the Enterprise Composite Network Model.

Relevance

You can apply the SAFE Blueprint to each module you design for an enterprise network.

Objectives

Upon completing this lesson, you will be able to choose the appropriate security mechanisms in each Enterprise Composite Network Model module. This includes being able to meet these objectives:

- Describe the features and benefits of Cisco SAFE Blueprint
- Design SAFE security solutions for the Internet Connectivity module
- Design SAFE security solutions for the E-Commerce module
- Design SAFE security solutions for the Remote Access and VPN module
- Design SAFE security solutions for the WAN module
- Design SAFE security solutions for the Network Management module
- Design SAFE security solutions for the Server Farm module

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with networking and security concepts, terms, and functions, including basic functions implemented in IOS software

Outline

The outline lists the topics included in this lesson.

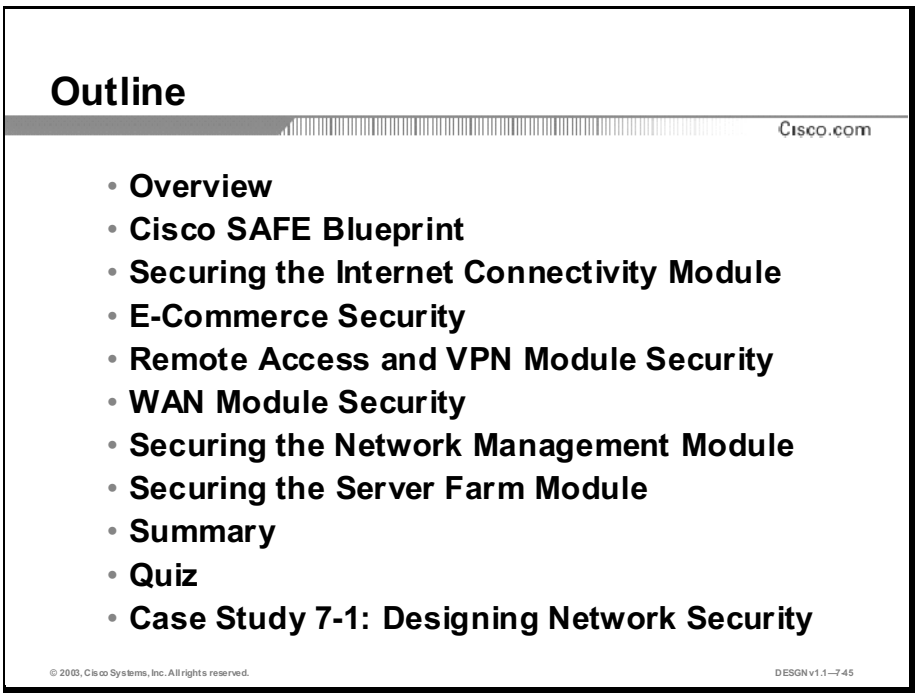

Cisco SAFE Blueprint

Cisco developed the SAFE Blueprint to provide guidelines for security implementation within the network infrastructure and beyond. The SAFE Blueprint provides information on best practices on designing and implementing secure networks. This topic describes the Cisco SAFE Blueprint.

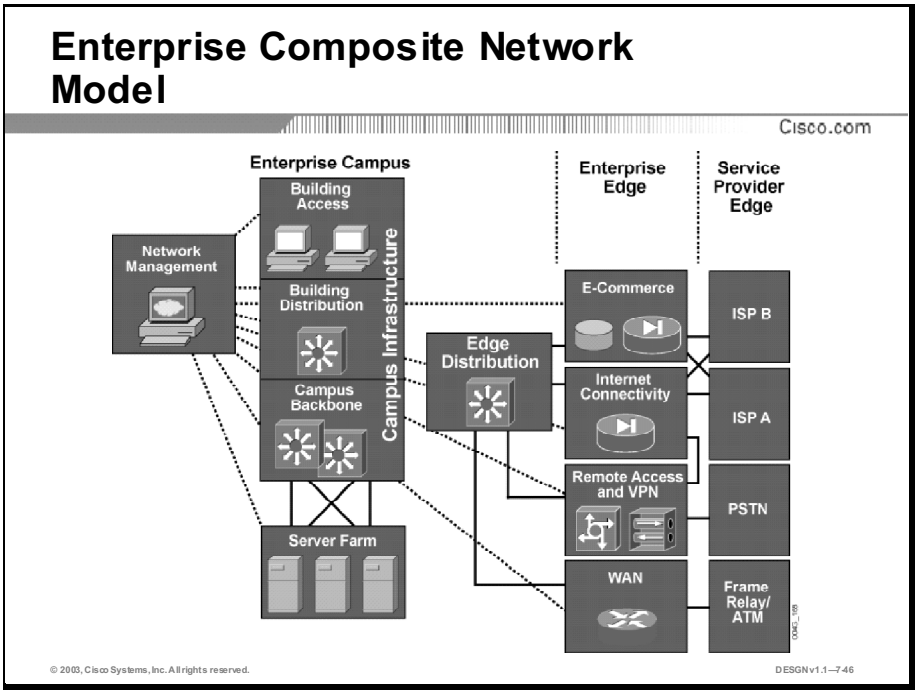

The SAFE security design follows the Enterprise Composite Network Model, in which each network module is analyzed to determine the threats, assess the severity of risks, and suggest best-practices design and implementation guidelines. SAFE often exceeds pure network security, as it relies on the concept of defense in depth, a multilayered security strategy where the failure of one security system is unlikely to lead to the compromise of network resources.

The figure shows how the SAFE Blueprint provides a modular view of the network infrastructure. By using a modular approach, a security designer reviews the security policy and then focuses on individual modules of the enterprise network to address the relevant risks. However, it is important to understand that SAFE does not apply to every network in the same way. Some networks may already have protection mechanisms in place and, therefore, cannot support SAFE designs, which generally assume the network is being designed securely from the beginning. Cisco can provide suggestions on how to improve existing security methods and provide defense in depth.

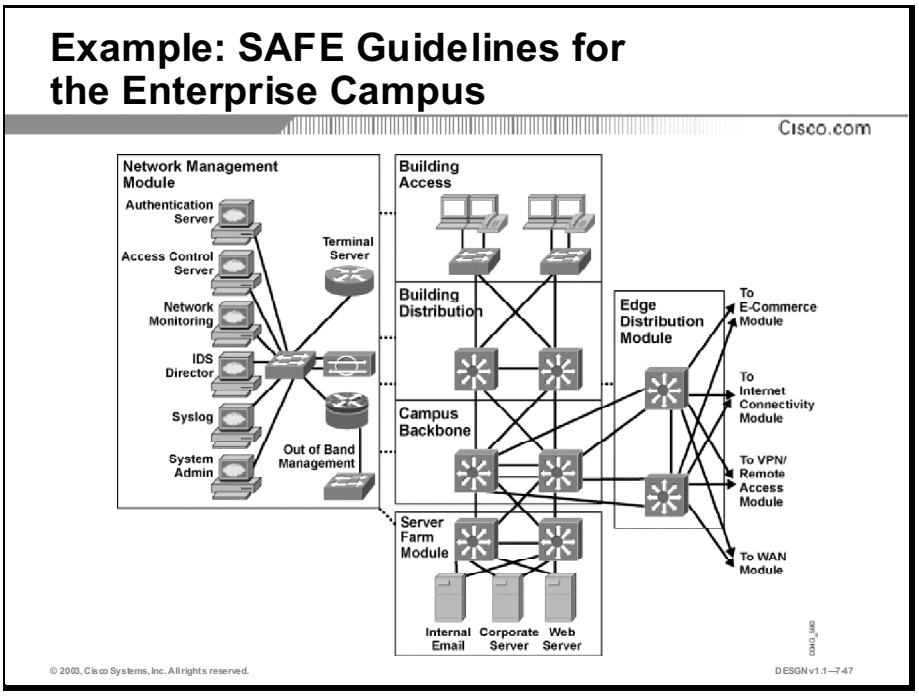

The figure illustrates how to apply SAFE guidelines to the Campus Infrastructure functional area. SAFE provides general design suggestions and selected configuration guidelines, which are recommended to increase security and comply with the risk management strategies suggested by the security policy.

Securing the Internet Connectivity Module

The Internet Connectivity module, part of the Enterprise Edge functional area, connects, directly or indirectly, all other network modules to the Internet, establishing potential contact with the most dangerous external network, where attackers' access to the network is easiest and attackers are difficult to trace. This topic discusses security mechanisms for the Internet Connectivity module.

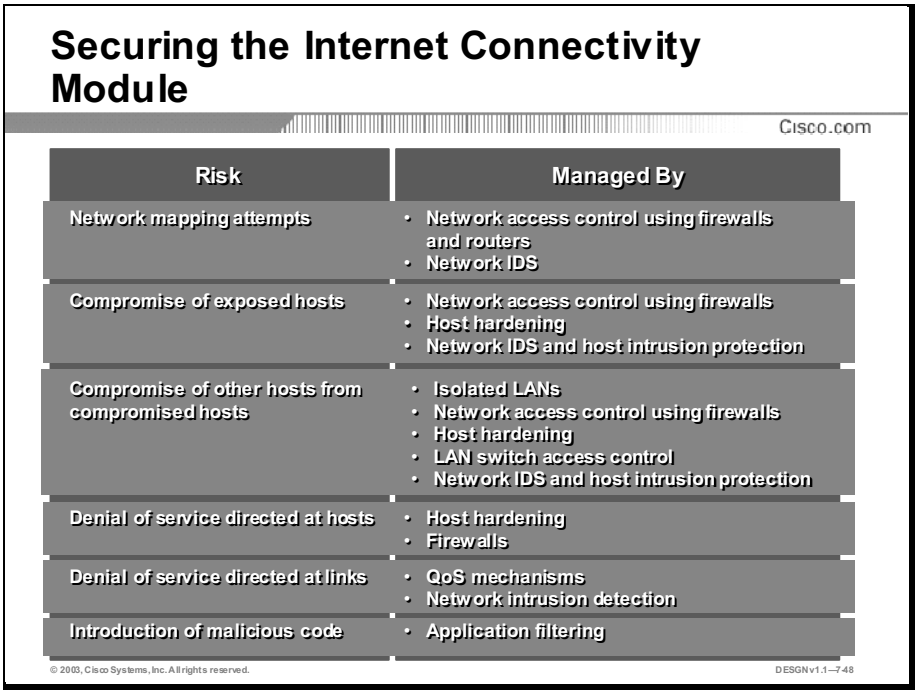

Public servers, such as the web and mail relay server of an enterprise network, are traditionally connected in the Internet Connectivity module, exposing them to the Internet. Historically, the majority of IP networks were strengthened at this point, which was reasonable because of the involved risks. However, this practice often introduced very relaxed security in other parts of the network. The SAFE Blueprint includes all parts of the network in a systematic security design, with each part secured according to the organization's policy.

Common risks associated with the Internet Connectivity module include:

- Reconnaissance threats from the Internet, where an attacker attempts to probe the network and its hosts in a network mapping effort, to discover reachable networks, hosts, and services running on exposed hosts. An example of such mapping is an attempt to ping all systems in the IP address range of the enterprise.
- The compromise of exposed hosts and their applications, which can lead directly to confidentiality breaches and integrity violations for data processed on exposed servers such as web servers and mail relay servers. An example is an attacker breaking into the mail relay server and being able to view or change all mail messages passing between the enterprise and the Internet.

- The compromise of other hosts from compromised hosts in the module, where an attacker can first compromise a host in the Internet Connectivity module and from that host compromise a host on another network module, such as the Enterprise Campus. An example is an attacker that breaks into a public web server and then breaks into the internal database server.
- Denial-of-service attacks directed at exposed hosts in this module. An example is an attacker sending a storm of connection requests (also known as a SYN flood) to the public web server, disabling its services.
- Denial-of-service attacks directed at network links such as the Internet connection of the enterprise network. An attacker could, for example, send a multigigabit stream of Internet Control Message Protocol (ICMP) messages towards the enterprise network, congesting its link to the Internet.
- Introduction of malicious code including viruses, Trojan horses, and malicious mobile code in Internet browsers over supported services such as e-mail and Internet access. A well-known example is an e-mail message with an executable attachment, containing a Trojan horse program, which sends the attacker the local user's passwords.

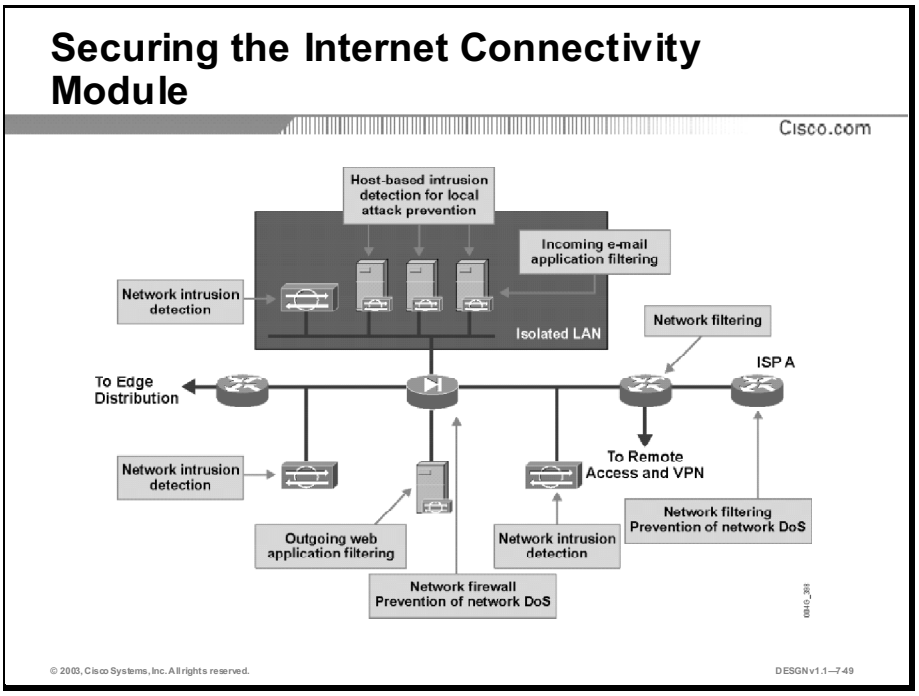

To manage the risks associated with the Internet Connectivity module, the Cisco SAFE Blueprint defines well-known and accepted solutions. Here are some solutions:

- To manage the risk of network mapping attempts, the routers and first-line firewall devices block all incoming connections, except those necessary, to the exposed hosts in the module. This prevents network mapping applications from reaching other hosts and networks. You can use access control lists on routers and firewalls to build filtering rules.
- To manage the risk of compromising an exposed host, the usual solution is to configure network firewalls to permit only the minimal required connections to exposed servers, and then secure the exposed server applications. Network and host intrusion detection systems then monitor individual hosts and subnets to detect signs of attack or malicious network activity, and identify potential successful breaches.
- To manage the risk of compromising other hosts from already compromised hosts, a common solution is to build an isolated LAN within the module. Isolated LANs connect to a leg of a firewall and host one or more servers. The purpose of an isolated LAN is to contain an attacker who has compromised a host, so all access from the potentially compromised host is filtered again by the firewall. This allows enforcement of an extremely strict connection policy, which denies all connections from public servers by default and prevents connectivity to hosts outside the isolated LAN. If multiple hosts are located in the same isolated LAN, LAN switch-based security access control mechanisms such as private VLANs can also effectively restrict communications among such hosts. Hardening of other hosts presents another defense-in-depth technique, and intrusion detection systems detect such attempts to further break-ins from compromised hosts.
- Denial-of-service attacks directed at hosts are managed by running reliable and patched operating systems and applications (host hardening), and using all available host denial-of-service protection tools on a firewall, such as connection rate limiting.
- Devices connected to potentially critical resources manage denial-of-service attacks directed at networks. The Internet router can limit the rate of specific traffic types over links to make flooding more difficult.

- Application code transferred between the Enterprise Campus and the Internet is often filtered using special application filtering servers in the Internet Connectivity module. Such servers examine the content of the application protocol, scanning, and possibly eliminating dangerous content, including viruses and mobile code.

The figure illustrates risk management techniques and their positioning within the Internet Connectivity module.

E-Commerce Security

The E-Commerce module, part of the Enterprise Edge functional area, hosts application servers, which serve e-commerce applications to Internet users. Connectivity to the Internet is available through the Internet Connectivity module, which already provides security mechanisms to manage some risks such as network denial of service. This topic describes security mechanisms for the E-Commerce module.

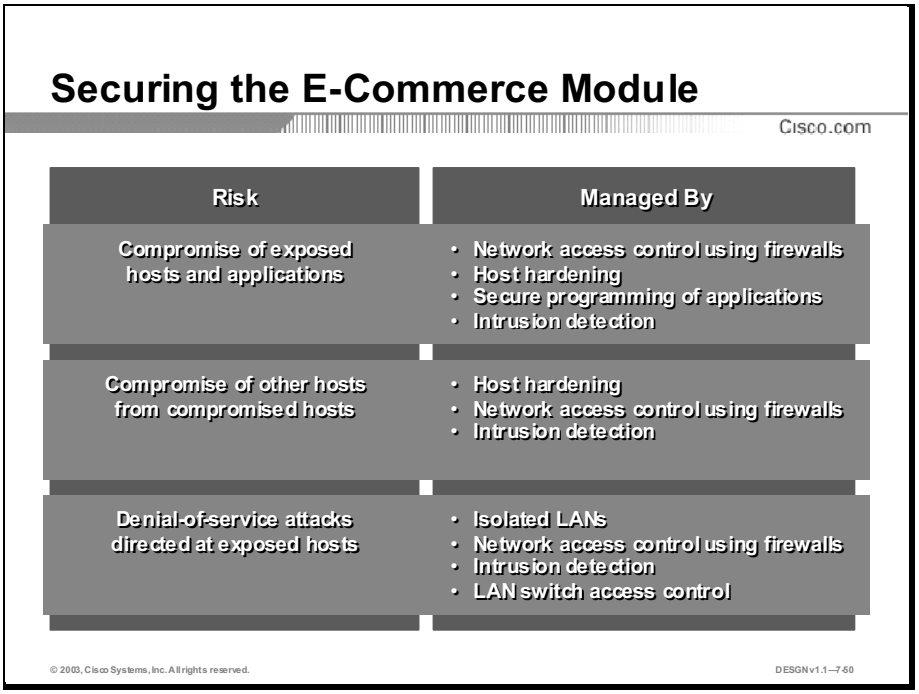

Here are some specific risks in the E-Commerce module:

- Compromise of exposed e-commerce hosts and applications, which is a risk also present in the Internet Connectivity module. However, the E-Commerce module usually hosts high-profile servers processing highly confidential and sensitive data, therefore a compromise results in a more substantial loss.
- Compromise of other hosts from compromised e-commerce servers in this module, including compromise of other e-commerce servers.
- E-commerce servers are often targets of denial-of-service attacks, directed at their operating systems or applications.

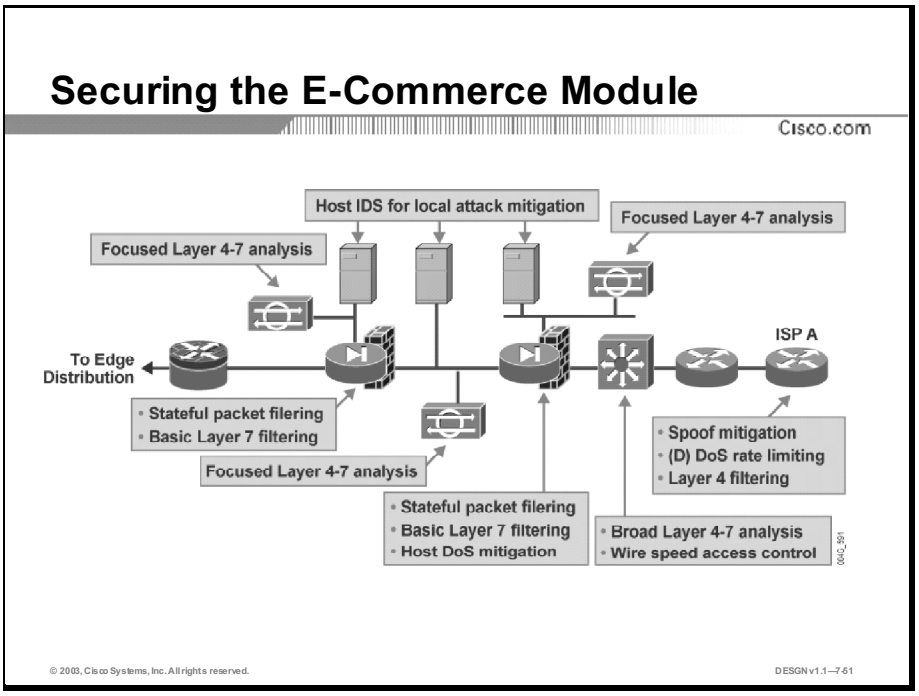

To manage the risks in the E-Commerce module, Cisco recommends these techniques:

- Host and application protection is managed through very tight network filtering on firewalls and through good host hardening and secure applications.
- To manage the risk of compromising other hosts from compromised e-commerce servers, similar techniques are used as in the Internet Connectivity module. Often, e-commerce applications are multitiered and run on multiple servers. For example, an e-commerce application front-end web server accepts encrypted sessions from Internet clients, processes the requests, and queries a database server, which contains sensitive data. Separate the multitiered server systems in their own isolated LANs to ensure there is a firewall system between them to protect more secure servers in the event of front-end compromise. Firewalls restrict connections from exposed e-commerce servers, so that compromise of any other host is less likely. To separate hosts on the same segment, use LAN switch access control mechanisms such as private VLANs. Network and host intrusion detection systems monitor individual hosts and subnets to detect signs of attacks and confirm potential successful breaches.
- For host-directed denial-of-service attacks, the same risk management mechanisms apply as in the Internet Connectivity module. Specifically, run reliable and patched operating systems and applications (host hardening), deploy host intrusion detection software, and use all available host denial-of-service protection tools on a firewall.

Example: Securing an E-Commerce Server Farm

An enterprise sets up an E-Commerce server farm, consisting of an application web server and a back-end database server, which replicates data from internal production servers. A good firewall design places both servers in separate isolated LANs, with firewalls in-between, and allows only the necessary connections of the E-Commerce application. Moreover, the application itself can be developed using high-level languages and security awareness and code checked by an outside verification agency.

Intrusion detection on both hosts, as well as network intrusion detection in most segments, detects and responds to attack attempts. You could simplify the design with a single multilegged firewall instead of multiple firewall devices separating all of the segments. Multiple firewall devices are often used to distribute access rules to multiple devices, making them simpler to understand, and to lessen the impact of firewall misconfiguration or potential bugs in firewall code, which impact only one of many devices.

The figure illustrates the outlined risk management techniques and their positioning within the E-Commerce module. Observe the multiple isolated LANs, which separate multitiered E-Commerce application servers in their own segments, separated by firewalls.

Remote Access and VPN Module Security

The Remote Access and VPN module, part of the Enterprise Edge functional area, hosts dial-up access servers, remote access VPN concentrators, and site-to-site VPN gateways. Connectivity to the Internet is available through the Internet Connectivity module, which already provides security mechanisms to manage some risks such as network denial of service. There is also local connectivity to external dial networks, such as the PSTN. This topic describes security mechanisms for the Remote Access and VPN module.

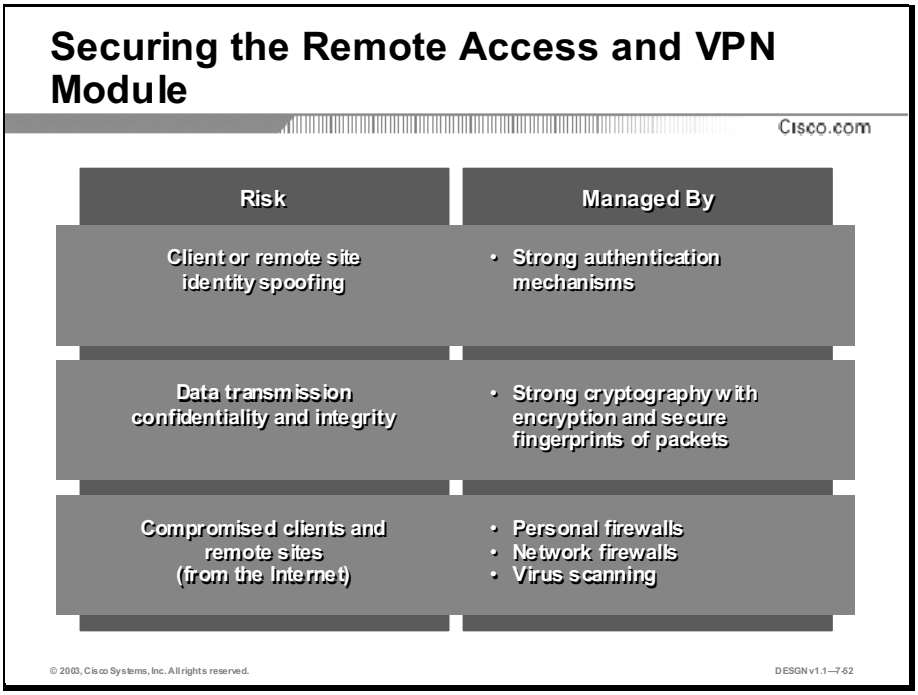

Here are some specific risks in the Remote Access and VPN module:

- Identity spoofing of remote clients or sites, where an attacker can impersonate a legitimate client and log in to the remote-access VPN connection. This can, for example, be possible if an attacker steals a legitimate user's credentials (such as a dial-up username and password pair) or guesses the authentication keys of a VPN connection.
- An attacker can obtain access to and change sensitive data by obtaining access to the network media, as VPN and dial-up connections are both transported over the Internet or PSTN. For example, an attacker might break into an ISP or a telco switch and eavesdrop on traffic.
- The compromise of a client or remote site, where an attacker successfully attacks the protected network over the VPN or dial-up connection through a legitimate client's system or a branch office. An example is a VPN client who has been compromised by a Trojan horse application. Such an application could turn the client system into a relay, so that when the client is connected to the enterprise network via an Internet remote-access VPN, the attacker can connect to the client over the Internet and then from the client to the protected enterprise network over the VPN.

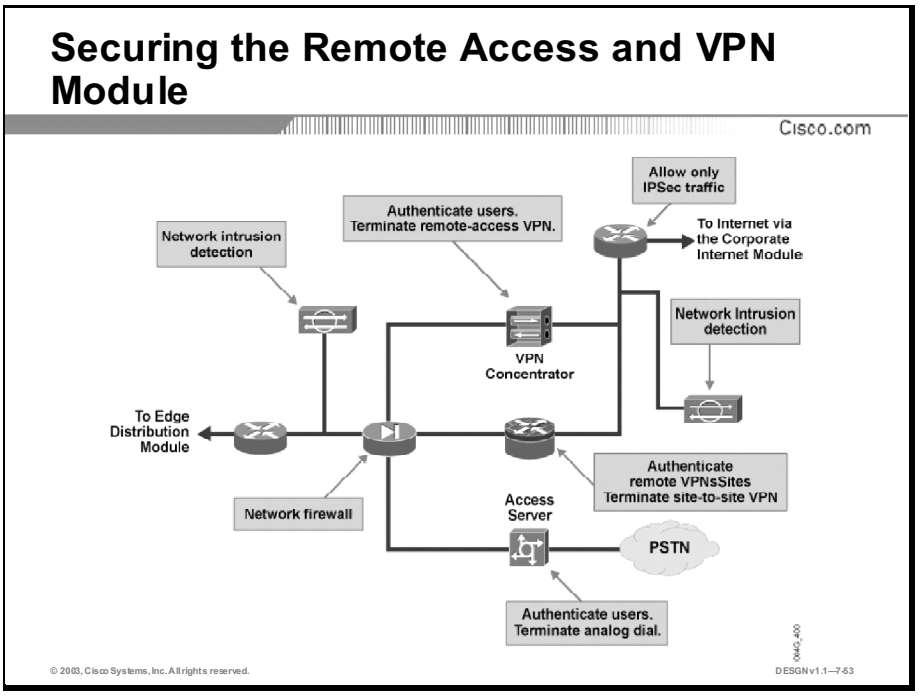

To manage risks associated with the Remote Access and VPN module, these techniques are available:

- To manage the risk of identity spoofing, an organization deploys very strong authentication to authenticate access from an external network because the external network does not provide any trusted authentication itself. Examples are token-card two-factor authentication for remote clients and public-key based certificate authentication for VPN sites.
- Modern cryptography and encryption algorithms provide confidentiality, while per-packet secure fingerprints provide data integrity and authenticity. The IPSec standard provides confidentiality, integrity, and authenticity in an IP network. IPSec provides a secure path between remote users and the VPN concentrator and between remote sites and the VPN site-to-site gateway. IPSec also protects classic dial-up users working over the PSTN.
- Limit access to the protected network using network firewalls, and enforce security standards on remote clients and sites to reduce the risk of compromised clients and remote sites. For example, a VPN concentrator can verify that a client is running a personal firewall and prevent communication with the Internet while connected to the VPN.

Example: Remote Access and VPN Security

An organization needs to migrate from a costly analog dial-up access system to an Internet remote-access and VPN system. When using the dial-up system, authentication is accomplished by using plain passwords (something that a user knows) and a callback to a predefined telephone number (something that a user "is"). When migrating to the Internet, the organization requires equivalent authentication strength and deploys OTP token cards. To guarantee confidentiality and integrity, the organization chooses strong encryption and integrity guarantees of the IPSec protocol suite.

WAN Module Security

The WAN module, part of the Enterprise Edge functional area, provides WAN connectivity among different parts of the enterprise network, usually connecting campuses with branch offices. Security is important whenever data is transferred between locations. This topic describes security mechanisms for the WAN module.

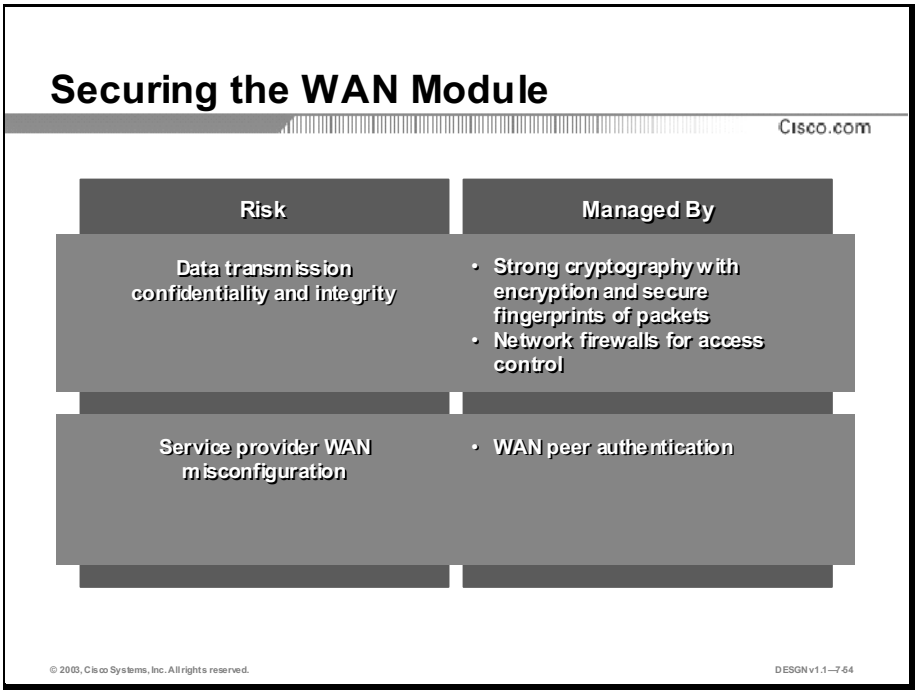

Here are some specific risks in the WAN module:

- Data transmission confidentiality and integrity, where an attacker obtaining physical access to the network media or to a service provider WAN switch can intercept WAN connections. An attacker might eavesdrop on any traffic or change data in transit.
- Accidental or deliberate misconfiguration of the WAN network, interconnecting different enterprises together. Some WAN protocols may establish automatic peering, and unwanted connectivity may become possible.

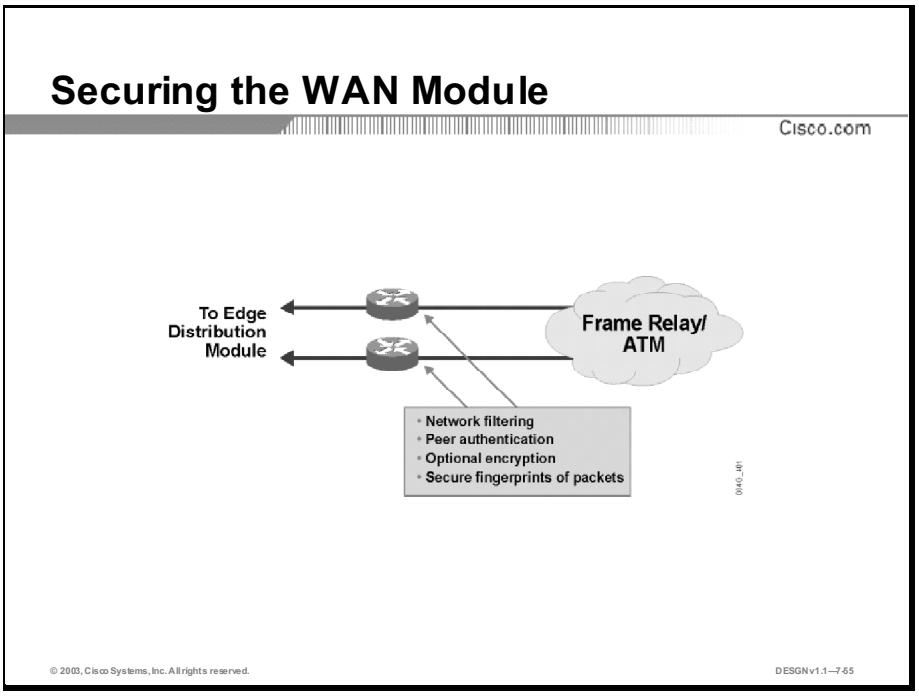

To manage the risks in the WAN module, apply these guidelines:

- Use the same technologies as with Internet VPNs to protect WAN communications. Encryption algorithms provide confidentiality, and per-packet secure fingerprints provide data integrity and authenticity. For example, IPSec protocols are often deployed on enterprise WAN networks to protect sensitive data.
- To prevent accidental WAN interconnection of different enterprises, WAN devices may require peer and routing protocol authentication over WAN links. Therefore, the WAN devices may not accept an unknown device, even though they both use the same WAN protocol.

Example: Security in the WAN

A service provider leased-line network failed several times in the last year and mixed customer circuits, accidentally interconnecting different enterprises. With address autoconfiguration (Serial Line Address Resolution Protocol [SLARP]) on some serial interfaces, and most enterprises running Enhanced Interior Gateway Routing Protocol (EIGRP) over Frame Relay, the networks have interconnected and traffic has been accidentally interchanged. As a result, most enterprises deployed PPP CHAP authentication on WAN links, managing the risk of rogue peers.

The figure illustrates how to deploy risk management techniques in the WAN module. Depending on the level of trust in the service provider's network, you can deploy encryption and secure fingerprinting in addition to basic filtering on WAN interfaces.

Securing the Network Management Module

The Network Management module provides network and security management to the entire enterprise network, hosting extremely sensitive data about network and security device configuration. It is usually connected to many devices directly over a separate management network, sometimes providing a potential management path around security mechanisms (such as firewalls), which is only used for management. This topic describes security mechanisms for the Network Management module.

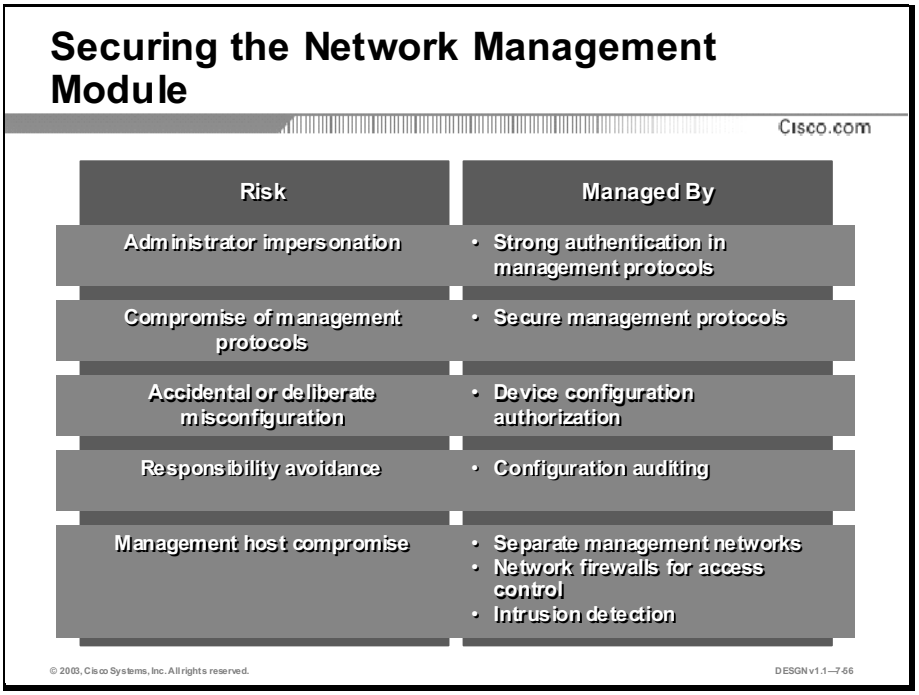

Here are some specific risks in the Network Management module:

- Impersonation of administrators, where an attacker might steal the usernames or passwords of an administrator, log on to network devices, and change their configuration
- Compromise of management protocols, where the attacker might send false management messages or listen to management protocols to obtain sensitive information
- Accidental or deliberate misconfiguration of network or security devices by inexperienced administrators
- Avoidance of responsibility among administrators, where an administrator might deny the actions that led to a security incident
- Compromise of management hosts, where very sensitive configuration, security, and audit data may be hosted

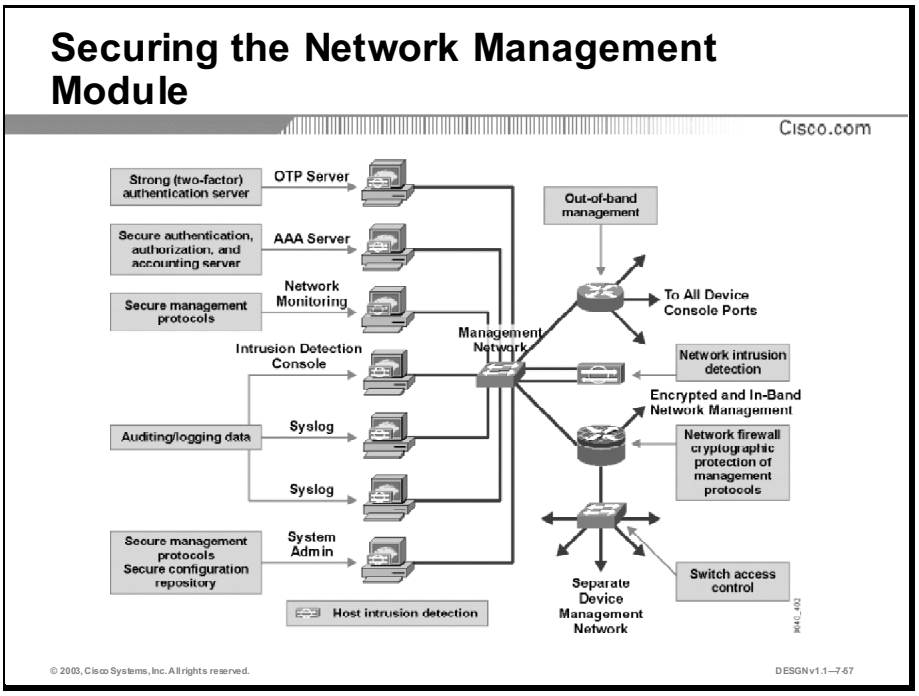

To manage the risks in the Network Management module, Cisco recommends these techniques:

- To manage the risk of administrator impersonation, provide the administrator with strong authentication mechanisms. A good example is a two-factor OTP system based on token cards.
- To manage the risk of a management protocols compromise, use protocols with cryptographic protection. Examples include SSH for terminal access (instead of Telnet), strongly authenticated SNMP instead of classic SNMP with basic authentication, and so on.
- To manage misconfiguration risks, enforce authorization on configuration mechanisms using well-known centralized change control and authorization servers that permit only specific changes by specific administrators.
- Prevent avoidance of responsibility by creating a good management audit trail, which logs all management events to a centralized and secured audit record repository. Good organizational practices must be in place to respond to such events.
- To manage the risk of management host compromise, separate the management network from the rest of the modules and implement good host security on the management stations.

Example: Securing the Network Management Module

An organization has experienced several incidents where untrusted users on the campus network intercepted management traffic. To manage the risk, the organization decided to move all management traffic to a separate VLAN (separate management network) to isolate it from user traffic.

The figure shows how different risk management mechanisms fit into the different elements of the Network Management module. A separate out-of-band management network and a network for network management system (NMS) servers exist, separated from the rest of the campus with a firewall router. A terminal server provides console access to network devices in an emergency and an additional switch connects the NMS server network to the out-of-band management network.

Securing the Server Farm Module

The Server Farm module hosts servers inside the main campus network and branch offices. Servers may contain the enterprise's most sensitive information and are available to a large number of users. Therefore, network performance is usually a critically important issue, which sometimes limits the choice of protection mechanisms. This topic describes security mechanisms for the Server Farm module.

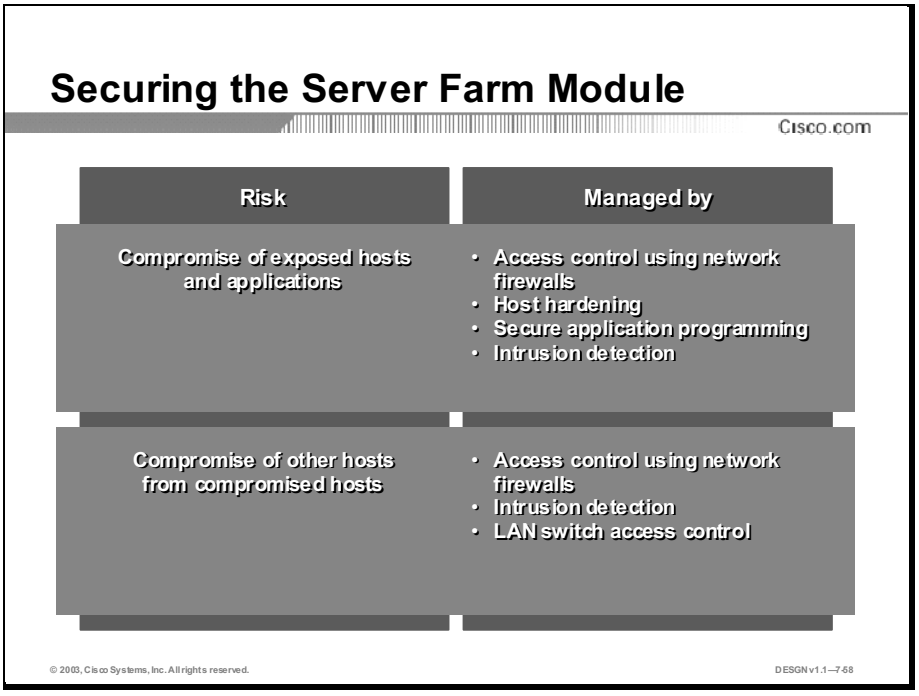

The specific risks in the Server Farm module are very similar to those in the E-Commerce module, except that sensitivity of data on internal servers is usually critical. Here are some specific risks in the Server Farm module:

- Direct compromise of exposed applications and unauthorized access to data
- Compromise of other hosts from compromised servers in this module

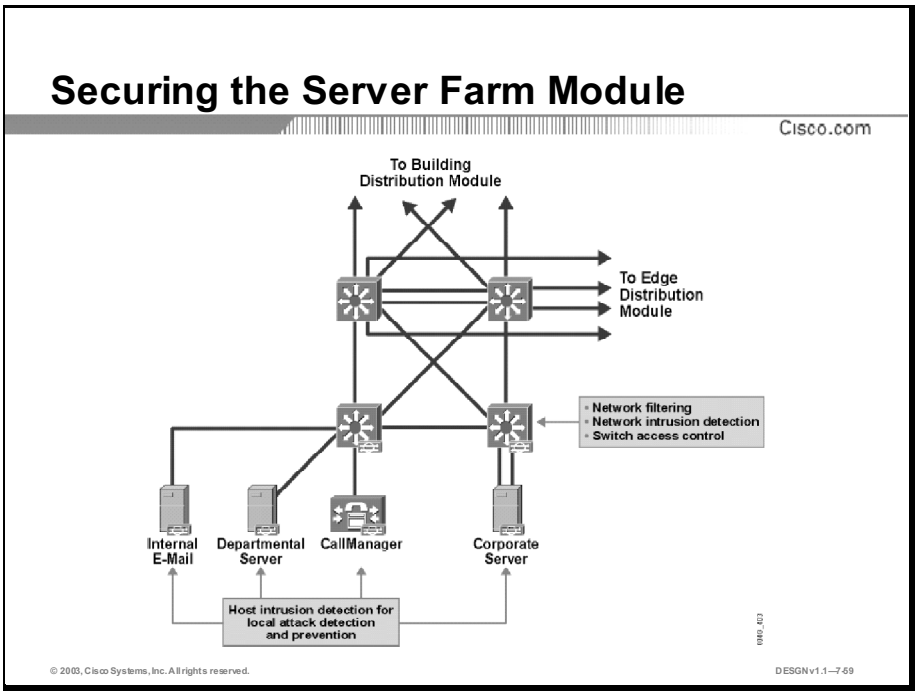

Risk management in the Server Farm module involves the same mechanisms as in the E-Commerce module, except that denial-of-service attacks in the internal network are less common. These guidelines apply:

- To manage the host and application protection, use good host hardening and secure applications. If performance permits, implement firewalls to allow only minimal connectivity to these servers.
- To manage the risk of compromise to other hosts from compromised servers, network filtering limits connectivity from the server. Separate the hosts on the same segment using LAN switch access control mechanisms, use network and host intrusion detection systems to monitor individual hosts, and use subnets to detect signs of attacks and confirm potential successful breaches.

Example: Firewall Security in the Server Farm Module

An organization became concerned when their operating system manufacturer published a number of security vulnerabilities. The server management personnel were unable to patch the system quickly enough, and internal server compromises occurred. To hide the majority of the unnecessary and vulnerable server services from the campus, the enterprise deployed a firewall quickly to protect their most valuable servers and limited user access to most of its services, thus making the servers less exposed to future vulnerabilities.

Summary

This topic summarizes the key points discussed in this lesson.

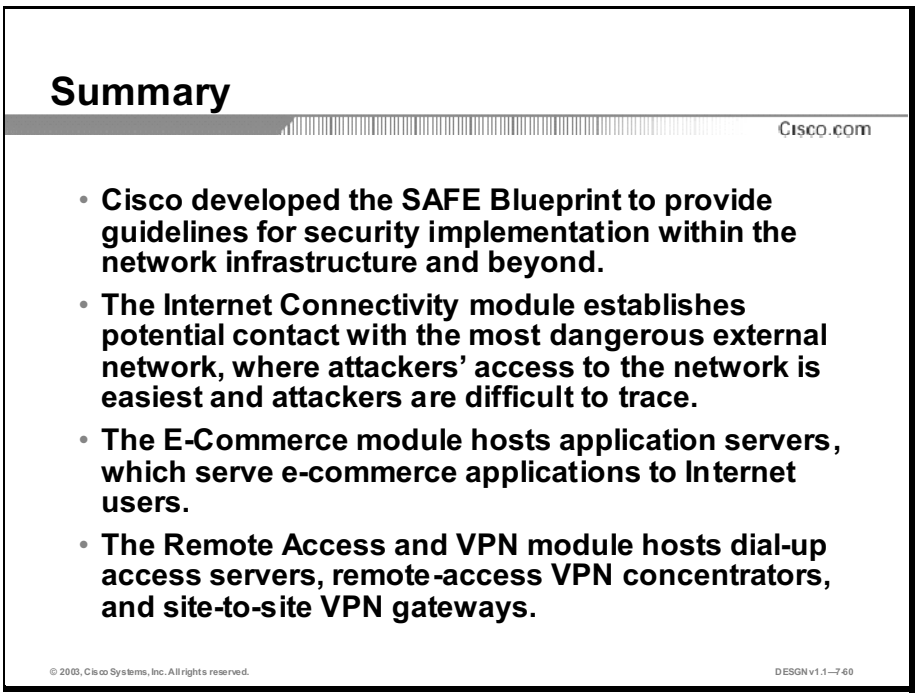

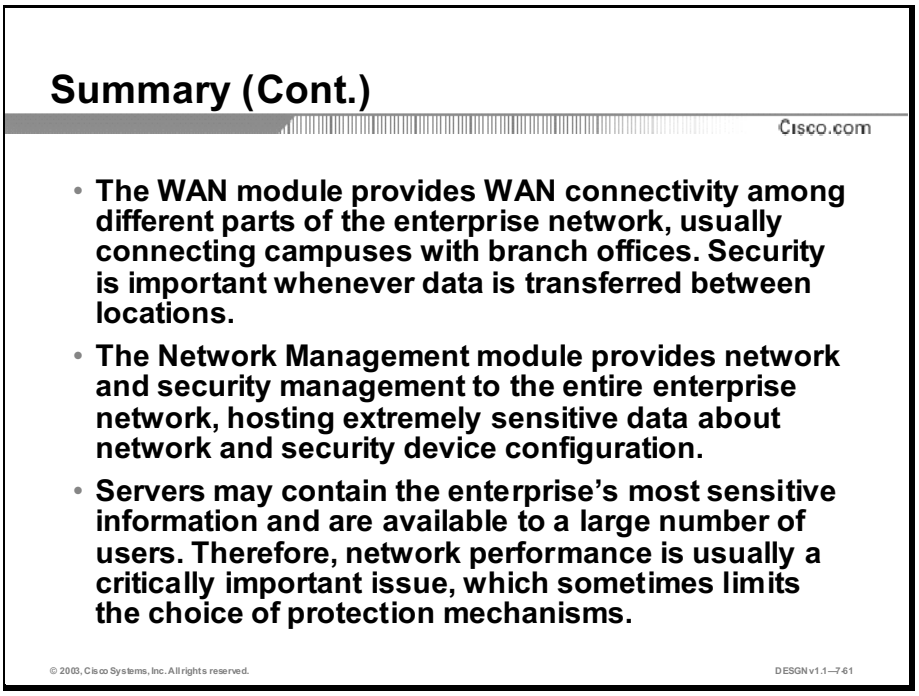

References

For additional information, refer to this resource:

- The Cisco *SAFE: A Security Blueprint For Enterprise Networks* white paper, http://www.cisco.com/go/safe.html

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Case Study 7: Designing Network Security

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which are the two main benefits the SAFE Blueprint offers to a network security designer? (Choose two.)

A) It provides recommended configurations for every security device in a network,

B) It offers guidelines on how to manage risks in several common design scenarios.

C) It applies to every network and can always be implemented in its entirety.

D) It allows network security to be implemented without regard to host and application security.

E) It offers suggestions on how to implement defense-in-depth.

Q2) Which two attacks can generally be prevented using network filtering access control lists in a router or firewall? (Choose two.)

A) network mapping attacks

B) sending viruses in e-mail messages

C) compromise of a host application, which should never be exposed to the attacker

D) downloading malicious Internet code inside HTTP sessions

E) flooding a link with seemingly legitimate traffic

Q3) Why is securing e-commerce servers generally more important than securing generic public web servers?

A) E-commerce servers are much more difficult to secure properly.

B) E-commerce servers generally process and store more sensitive data.

C) E-commerce servers are always more susceptible to attacks because of their complexity.

D) E-commerce servers have a much higher performance, and only a few security techniques can be used.

Q4) What is a method used to mitigate identity spoofing in the Remote Access and VPN module?

A) host hardening

B) QoS mechanisms

C) encryption algorithms

D) strong authentication

Q5) When should encryption be deployed on WAN links?

A) when there is a threat of integrity violations on the service provider WAN

B) when there is a threat of confidentiality breaches on the service provider WAN

C) when there is a need to authenticate peers in the WAN

D) when there is a need to perform network filtering in the WAN

Q6) How is the risk of administrator responsibility avoidance managed?

A) with management auditing and good organizational practices

B) with strong authentication of administrators only

C) with host intrusion detection on the management stations

D) with firewalls, performing access control to the management network

E) with secure management protocols

Q7) Which statement is true when securing the Server Farm module?

A) Firewalls can never be used to protect the Server Farm module.

B) Often, much higher performance is needed, which limits the choices of protection mechanisms.

C) Intrusion detection is not recommended as it lowers network performance.

D) Servers in the Server Farm module are never multitiered and do not need isolated LAN networks.

Quiz Answer Key

Q1) B, E

Relates to: Cisco SAFE Blueprint

Q2) A, C

Relates to: Securing the Internet Connectivity Module

Q3) B

Relates to: E-Commerce Security

Q4) D

Relates to: Remote Access and VPN Module Security

Q5) B

Relates to: WAN Module Security

Q6) A

Relates to: Securing the Network Management Module

Q7) B

Relates to: Securing the Server Farm Module

Case Study 7: Designing Network Security

Complete this case study to practice what you learned in this module.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this module, you discussed the Cisco SAFE Blueprint. Upon completing this case study, you will be able to apply the Cisco SAFE Blueprint to your earlier design based on the Enterprise Composite Network Model and the requirements of DJMP Industries, and you will be able to meet these objectives:

- Select appropriate security mechanisms to counter specific threats
- Position appropriate security mechanisms in each design module

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, present in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Complete these steps:

Step 1 Refresh your knowledge of the network topology and requirements as presented in the DJMP Industries Case Study Scenario completely before the exercise. Allow a maximum of 10 minutes for reading.

Step 2 Discuss the scenario, anticipated threats and the security options for each component of the composite model with your group. Allow 10 minutes for the discussion.

Step 3 Propose the optimal security design that addresses the scenario requirements.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class.

Module 8

Designing Networks for Voice Transport

Overview

Many enterprises are integrating their voice and data networks into a single, converged network. Converged voice networks can run the same applications as a telephony network but in a more cost-effective and scalable manner. To design a network, you need to consider both voice and data traffic.

This module reviews traditional voice architectures and features. It explains the reasons for migrating from a traditional architecture to integrated architectures. This module introduces three technologies for voice transmission across data networks: Voice over IP (VoIP), Voice over Frame Relay (VoFR), and Voice over ATM (VoATM). You will learn about voice traffic engineering on both the Public Switched Telephone Network (PSTN) and the VoIP network. The module concludes with useful guidelines for designing IP telephony networks.

Module Objectives

Upon completing this module, you will be able to design an enterprise network for voice transport.

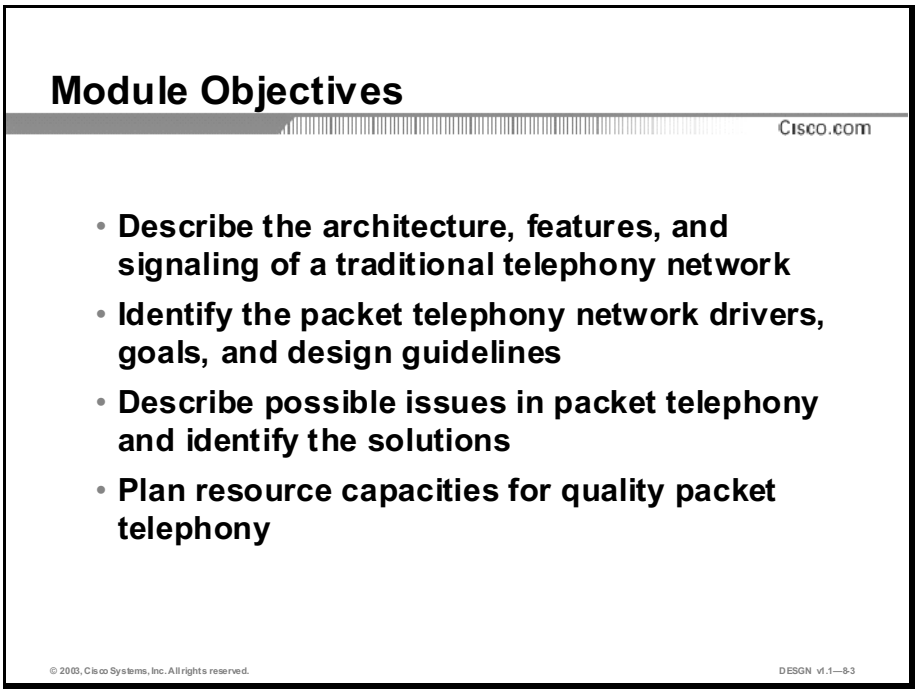

Module Outline

The outline lists the components of this module.

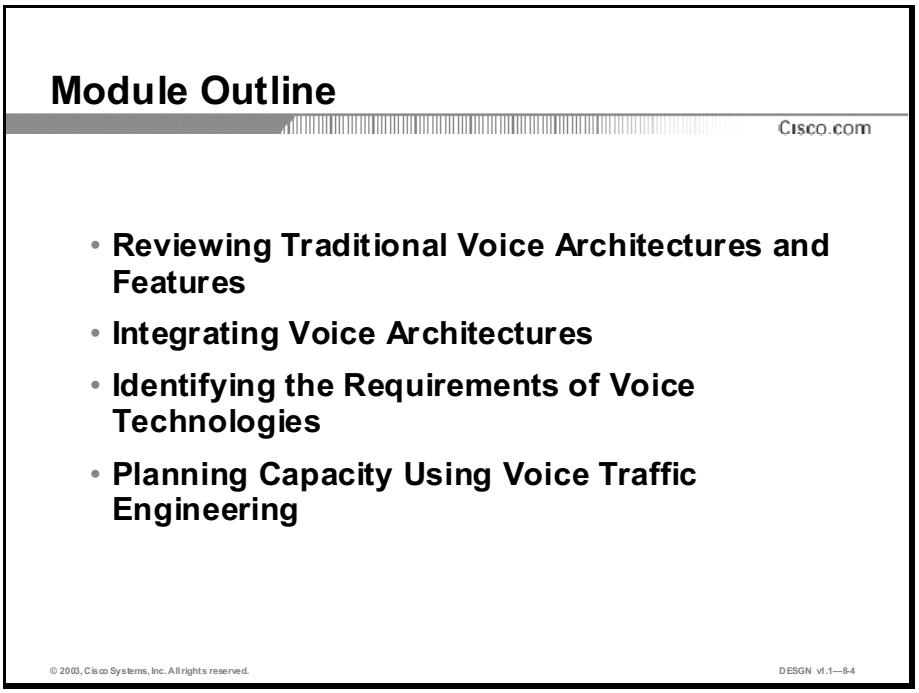

Reviewing Traditional Voice Architectures and Features

Overview

Traditional voice architectures and equipment are pervasive throughout organizations. Any new telephony solutions must integrate into existing environments and provide similar functionality.

This lesson introduces the traditional telephony infrastructure and explains its major components. The lesson describes how voice is routed and transmitted across a digital network and introduces basic circuit-switching concepts. The lesson concludes with a description of the services offered in traditional PSTNs.

Relevance

This lesson will help you identify traditional voice architectures and features so you can consider them in a voice network design.

Objectives

Upon completing this lesson, you will be able to describe the architectures, features, and signaling in a traditional telephony network. This includes being able to meet these objectives:

- Explain the difference between analog and digital signaling and describe their impact on voice transmission
- Distinguish between PBXs in the enterprise and voice switches in the PSTN network
- Describe the concept of local loops and trunks
- List and describe analog and digital telephony signaling types
- Explain how the PSTN numbering plan works and how PSTN switches route voice calls
- Describe the role of PBX and PSTN services and how they are implemented

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with traditional telephony concepts, terms, and functions

Outline

The outline lists the topics included in this lesson.

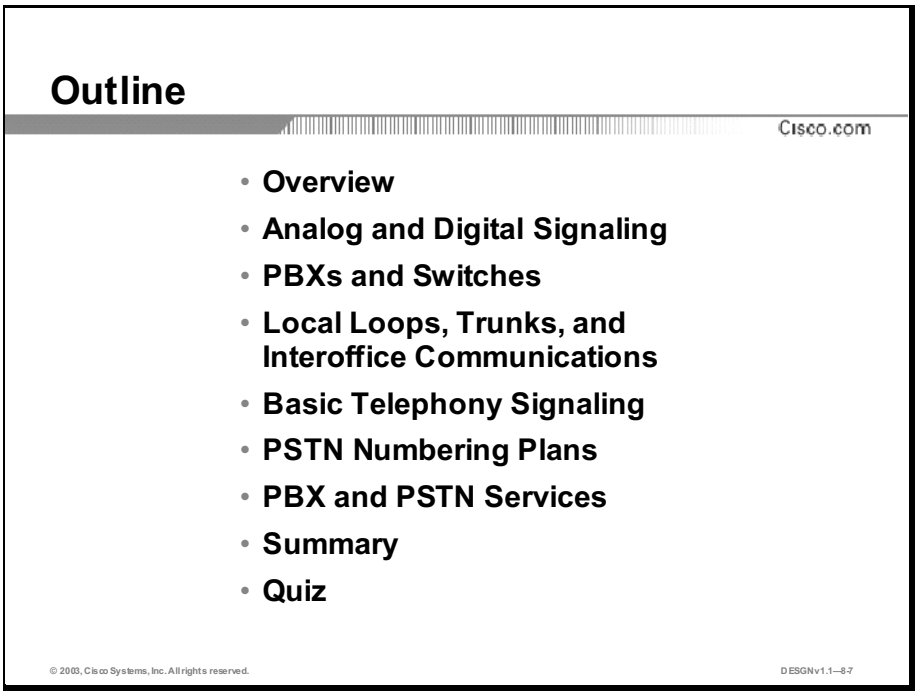

Analog and Digital Signaling

For clear voice connections, a phone system converts analog speech to a digital format and sends it over a digital network. At the other end of the connection, the phone system converts the digital signal back to an analog format: normal sound waves that the ear can pick up. This topic explains the difference between analog and digital signaling and describes their impact on voice transmission.

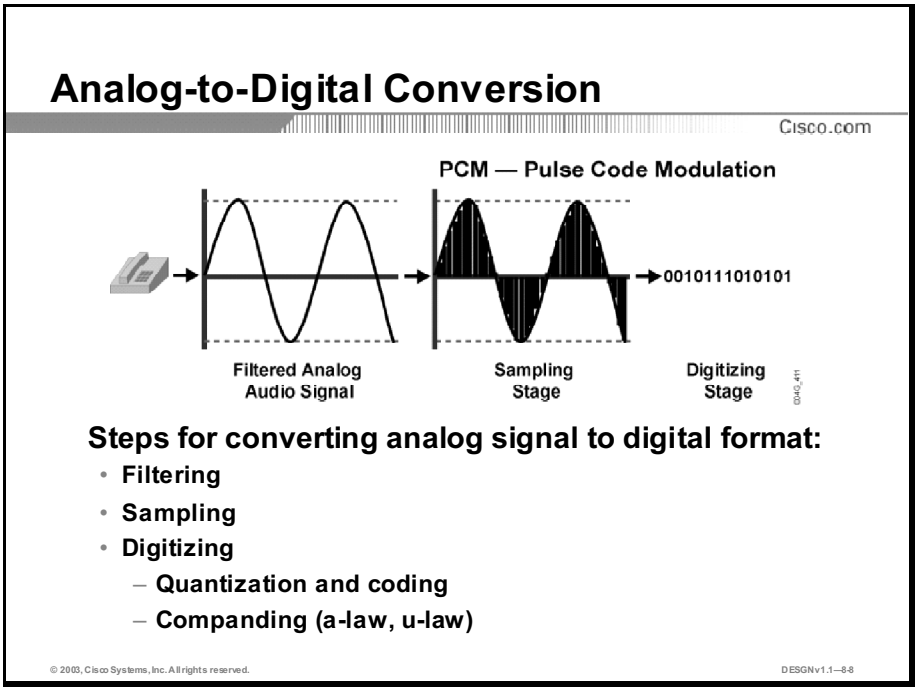

There are several steps to converting an analog signal into digital format:

- **Filtering:** Filters out the speech frequency component of the signal. Most of the energy of spoken language ranges from 300 Hz to (approximately) 3400 Hz. Therefore, early digital telephony researchers established a 3100-Hz bandwidth for standard speech. A coder-decoder (codec) puts the analog waveforms through a voice frequency filter to filter out anything greater than 4000 Hz.
- **Sampling:** Samples the filtered input signal at a constant sampling frequency using a process called pulse amplitude modulation (PAM). This step uses the original analog signal to modulate the amplitude of a pulse train that has a constant amplitude and frequency. The sampling rate is twice the highest frequency. This value is 4000 Hz, which is 8000 times a second or every 125 microseconds.
- **Digitizing:** Digitizes the samples to transmit over a telephony network using pulse code modulation (PCM). The only difference between PAM and PCM is that PCM takes the process one step further by encoding each analog sample using binary code words. PCM has an analog-to-digital converter on the source side and a digital-to-analog converter on the destination side.

The digitizing process is further divided into two subprocesses:

- **Quantization and coding:** A digitizing process that converts each analog sample value into a discrete value that is assigned a unique digital code word. As the input signal samples enter the quantization phase, they are assigned to a quantization interval. All quantization intervals are equally spaced throughout the dynamic range of the input analog signal. Each quantization interval is assigned a discrete value in the form of a binary code word. The standard word size used is 8 bits, enabling 256 possible quantizing intervals. The eight-bit code words allow for a bit rate of 64 kilobits per second (kbps). The rate is calculated by multiplying the sampling rate (twice the input frequency) by the size of the code word (2 x 4 kHz x 8 bits = 64 kbps).
- **Companding:** A process of first compressing an analog signal at the source, and then expanding this signal back to its original size when it reaches its destination. Combining the two terms, "compressing" and "expanding," into one word resulted in the term "companding." During the companding process, input analog signal samples are compressed into logarithmic segments and then each segment is quantized and coded using uniform quantization. The result is a more accurate value for smaller amplitudes and a uniform signal-to-noise quantization ratio across the input range.

  The a-law logarithmic companding standard is used in Europe, and u-law is used in North America and Japan. For communication between a u-law and an a-law country, the u-law country must change its signaling to accommodate the a-law country.

PBXs and Switches

Private branch exchanges (PBXs) and public telephone switches share many similarities, but they also have differences. This topic distinguishes between PBXs in the enterprise and voice switches in the PSTN network.

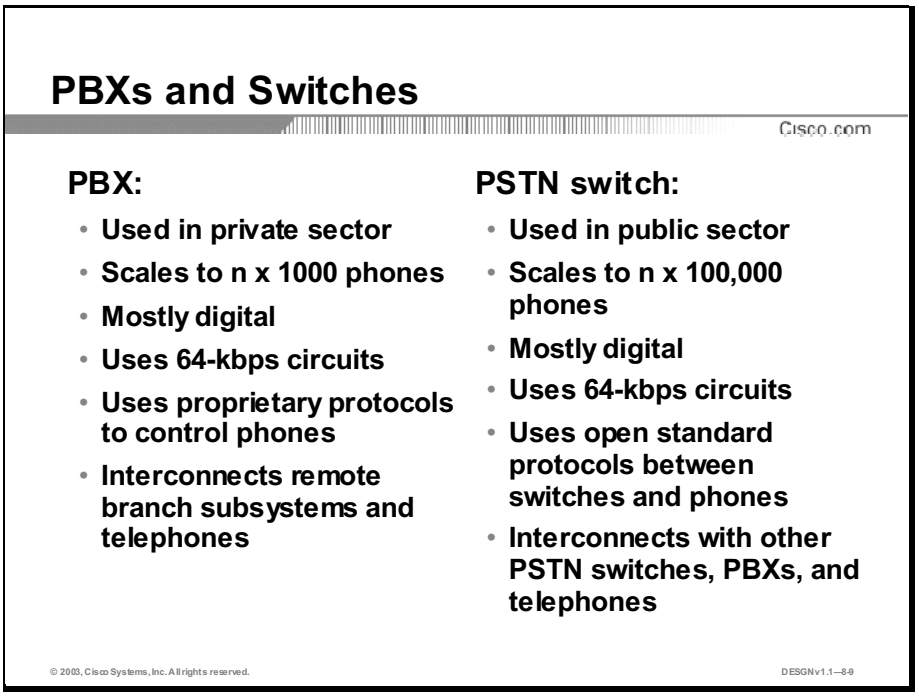

A PBX is a business telephone system that provides business features such as call hold, call transfer, call forward, follow-me, call park, conference calling, music on hold (MOH), call history, and voice mail. Most of these features are not available in public systems.

While the PBX and PSTN switch systems use the same type of circuits, the scale is very different. The PSTN switch can support hundreds of thousands of telephones while a PBX can support several thousand only.

The primary task of the PSTN switch is to provide residential telephony. The PBX, however, supports user telephones within a company. PBX vendors often create proprietary protocols to enable their PBXs to intercommunicate and carry additional features transparently through their voice network. A PBX typically supports its own vendor-specific phones. This encourages enterprise networks to consolidate to one brand of PBX to take full advantage of the offered features.

Note: Many vendors are implementing standards-based signaling protocols that enable interoperability between different PBXs. The two standards are Q Signaling (QSIG) and Digital Private Network Signaling System (DPNSS).

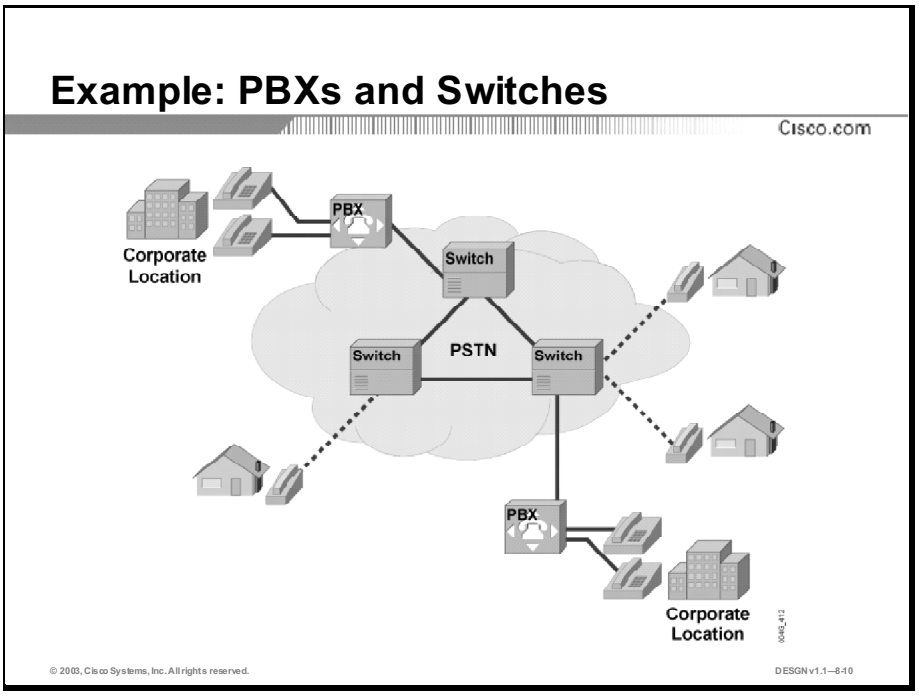

PBXs are typically located at corporate locations. PSTN switches build the PSTN network and are located in central offices (COs). PSTN switches connect residential and business users, and PBXs are used mainly for business purposes.

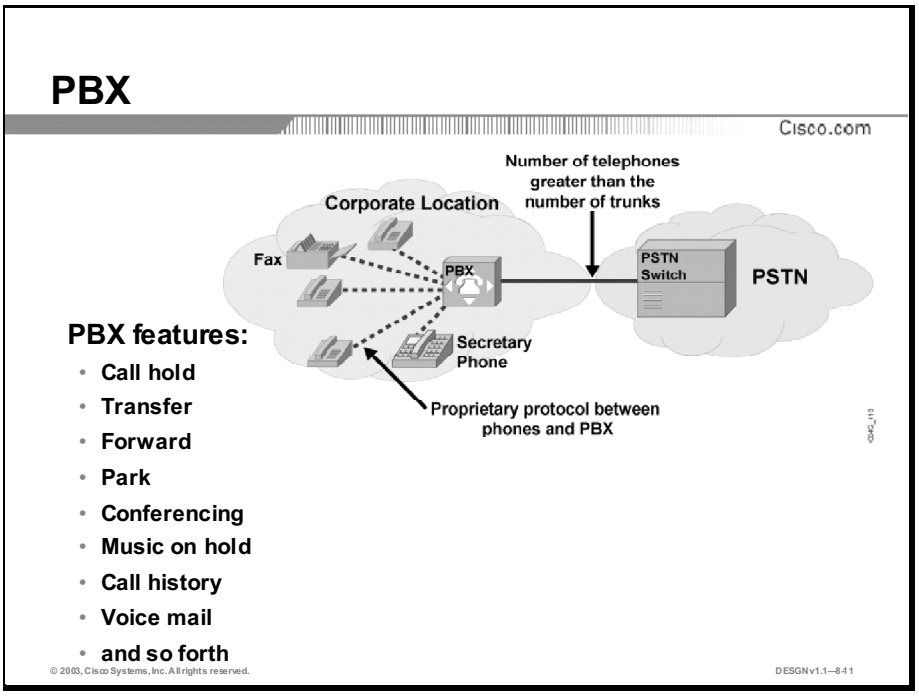

A PBX is a small version of a telephone switch, and is used for business purposes. A PBX provides many call features, such as call hold, transfer, forward, park, conferencing, MOH, call history, and voice mail, that business customers require. This switch often connects to the PSTN through a T1 or E1 digital circuit or circuits. The PBX supports end-to-end digital transmission, uses PCM switching technology, and supports both analog and digital proprietary telephones.

Advantages of and Disadvantages of PBXs

A local PBX provides several advantages:

- Local calls between telephones within the PBX or group of PBXs are free of charge.
- Most users of the PBX telephone system are not calling externally through T1 or E1 circuits at the same time; therefore, cost savings on PSTN trunks are realized because fewer outside lines than telephones are required.

When adding a new user, changing a voice feature, or moving a user to a different location, a customer does not need to contact the PSTN carrier. The local administrator can reconfigure the PBX. However, the PBX adds another level of complexity: the enterprise customer must configure and maintain the PBX.

All features, such as call forward, call hold, transfer, and voice mail, are available in the PBX.

Note: Enterprises install PBXs because the number of telephones is usually greater than the number of simultaneous calls to the PSTN network. Only a small percentage of telephones are active at one time. Companies with a PBX only need the number of external lines (to the PSTN) to be equal to the maximum desired number of simultaneous calls.

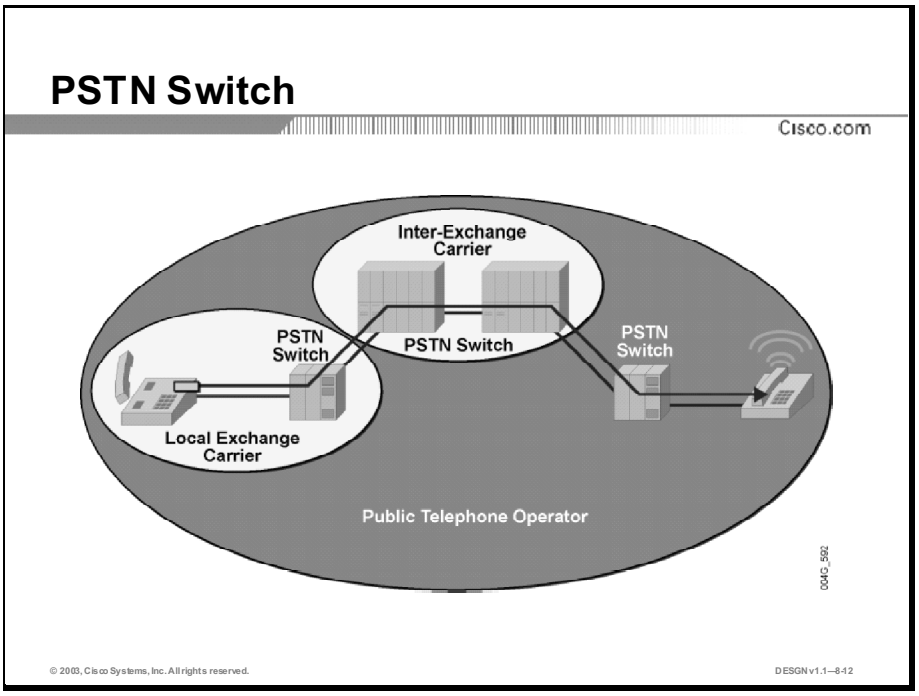

The PSTN resembles a single large network with telephone lines connected. In reality, the PSTN is composed of circuits, switches, signaling devices, and telephones. Many different companies may own and operate different systems within the PSTN.

The primary role of the PSTN switch is to connect the calling and called parties. If the two parties are physically connected to the same PSTN switch, the call remains local. Otherwise, the call is forwarded to the corresponding destination switch, which owns the called party. The PSTN switch connects business PBXs and public and private telephones with other PSTN switches. Large PSTN switches are located at COs, which provide circuits throughout the telephony network.

PSTN switches are deployed in hierarchies, to provide resiliency and redundancy to the PSTN network and avoid a single point of failure.

The PSTN signaling supports features such as caller ID and Direct Inward Dialing (DID).

Local Loops, Trunks, and Interoffice Communications

The telephone infrastructure includes loops, trunks, and interoffice trunks that interconnect PSTN switches, PBXs, and telephones. This topic describes the content of local loops and trunks.

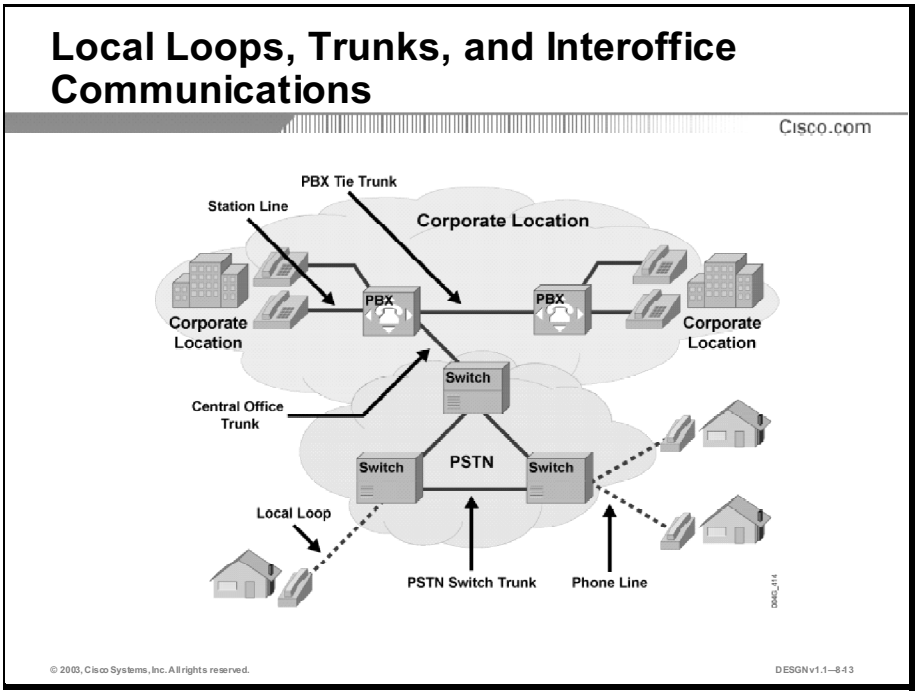

The telephone infrastructure starts with a simple pair of copper wires running to a home or business. The physical cabling is known as a local loop. The local loop physically connects the home telephone to the CO PSTN switch. Similarly, the connection between an enterprise PBX and its telephones is called the station line.

Trunk Communications

A trunk provides the communication path between two CO switches or telephony systems. Enterprises connect their PBXs to the PSTN over trunks. The telephone service provider is responsible for running trunks between its CO and the enterprise PBX.

Trunk Types

These trunk types are available:

- **Tie trunks:** Connect enterprise PBXs together without connecting to the CO (PSTN) switch
- **CO trunks:** Interconnect CO switches and enterprise PBXs
- **Interoffice trunks:** Interconnect CO switches

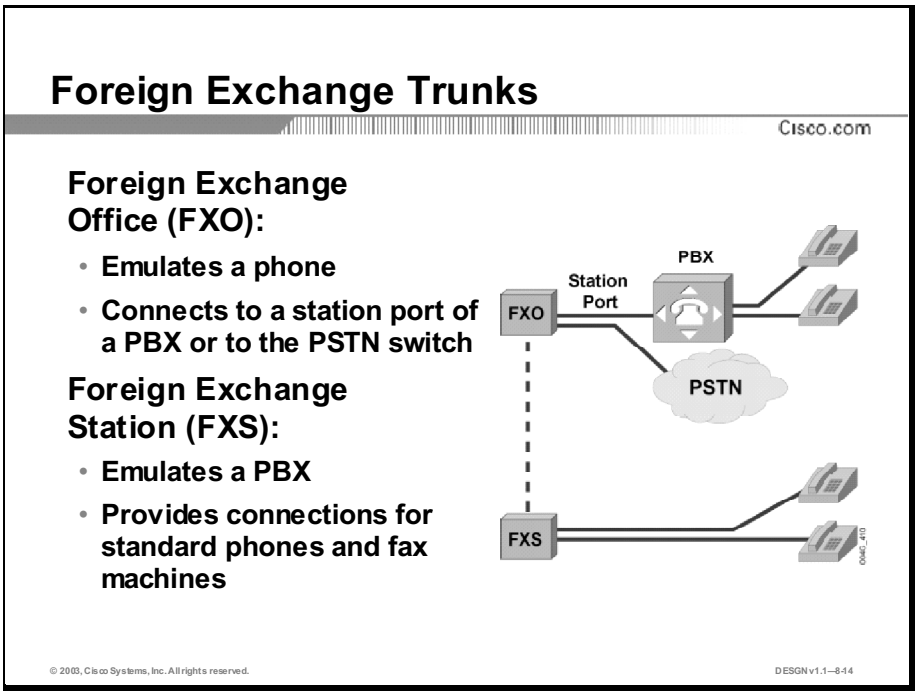

Foreign exchange trunks are analog or digital interfaces used to interconnect a PBX to telephones and other PBXs or the PSTN. These are the two types of foreign exchange trunks:

- **Foreign Exchange Office (FXO):** This interface creates an analog connection to a PSTN CO or to a station interface on another PBX. The FXO interface sits on the PBX end of the connection. It plugs directly into the line side of the PSTN or another PBX so that the PSTN or PBX thinks the FXO interface is a telephone. The FXO interface provides either pulse or dual tone multifrequency (DTMF) digits for outbound dialing.
- **Foreign Exchange Station (FXS):** This is an analog interface that typically terminates at a telephone, fax machine, or a similar device. The FXS interface functions like a switch and must supply line power, ring voltage and dial tone to the end device.

Basic Telephony Signaling

In a telephony system, a signaling mechanism is required to establish and disconnect telephone communications. This topic lists and describes analog and digital telephony signaling types.

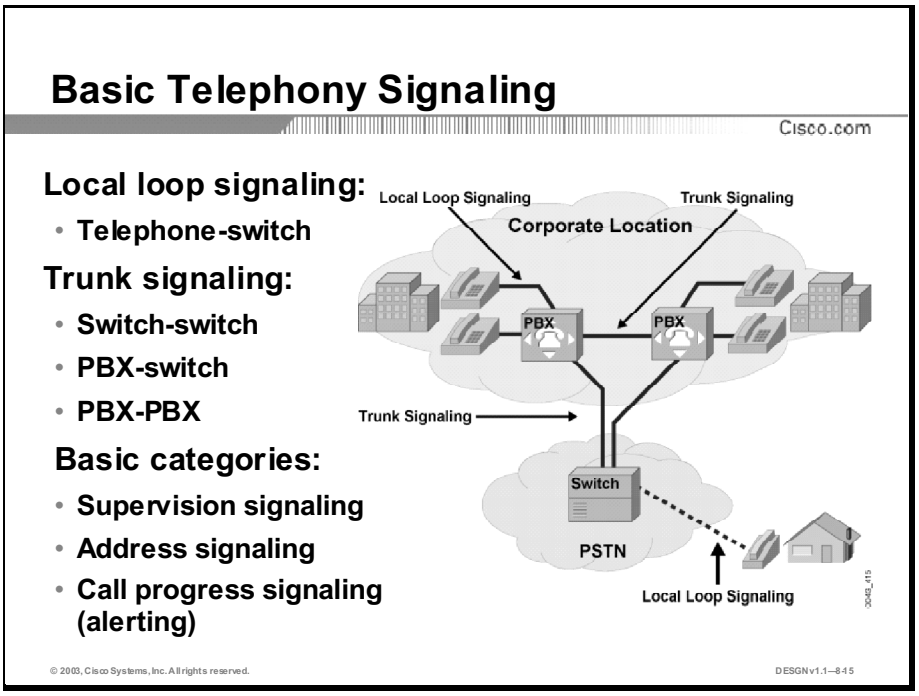

When a caller places a cross-country telephone call, many forms of signaling are used:

- Between the telephone and PBX
- Between the PBX and PSTN switch
- Between the PSTN switches
- Between two PBXs

In the broadest sense, there are two broad types of signaling:

- **Local loop signaling:** Between a PSTN switch and subscriber
- **Trunk signaling:** Between PSTN switches, between a PSTN switch and PBX, or between PBX switches

Three Basic Categories of Signaling

Simple examples of signaling include ringing of the telephone, a dial tone, and a ringback tone. There are three basic categories of signals commonly used in telephone networks:

- **Supervision signaling:** This is typically characterized as On-hook, Off-hook and Ringing. Supervisory signaling alerts the CO switch to the state of the telephone on each local loop.
- **Address signaling:** Involves the passing of pulse or dual tone multifrequency (DTMF) digits to a PBX or PSTN switch.
- **Informational signals:** This type of signaling is characterized by dial tone, busy tones, re-order tone, receiver off-hook, and no such number, such as those used with call progress indicators.

For a telephone call to take place, all three types of signaling occur.

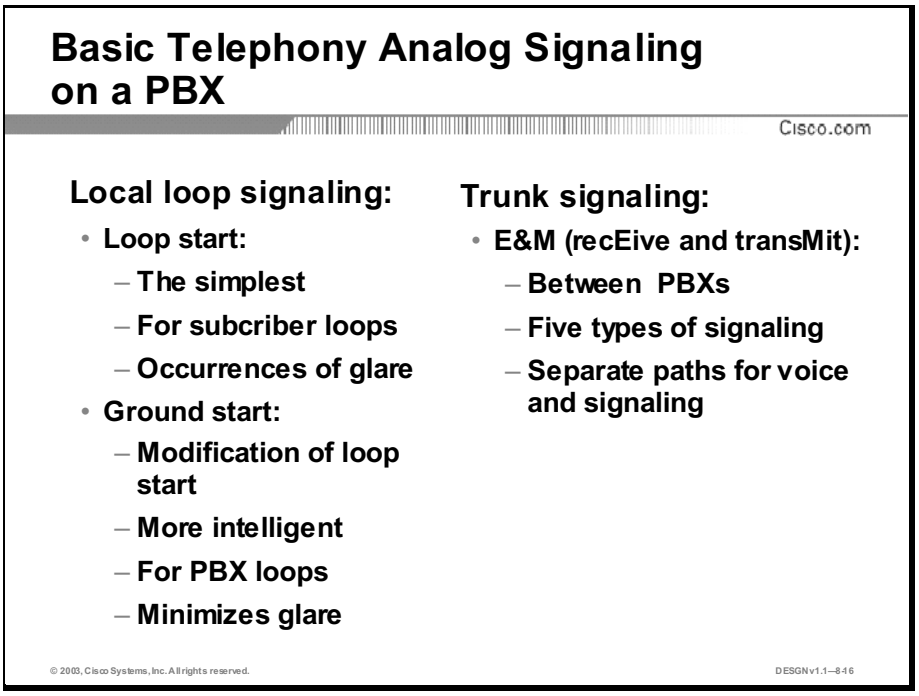

The most common methods of analog signaling are:

- **Loop start:** The simplest and least intelligent signaling protocol and the most common form of subscriber loop signaling. The creation of the electrical loop initiates a call, and opening of the loop terminates the call. PBXs generally do not support loop start signaling because it is subject to glare. Glare occurs when two endpoints try to seize the line at the same time, which results in two people being connected unknowingly. Because business callers use telephones regularly and the possibility for glare is high, loop start signaling is acceptable only for residential use.
- **Ground start:** Provides positive recognition of connects and disconnects. It uses current detection mechanisms at each end of the trunk. A ground start signal enables PBXs to agree which end is to seize the trunk before it is actually seized. This form of signaling minimizes the effect of glare. A ground start trunk is preferred when there is a high volume of calls; therefore, PBXs typically use this type of signaling.
- **E&M (recEive and transMit, sometimes also known as Ear and Mouth):** A common trunk signaling technique used between PBXs. In E&M, voice is transmitted over either two- or four-wire circuits, with five types of E&M signaling (Type I, II, III, IV, and V). E&M uses separate paths (or leads) for voice and signaling. The M (mouth) lead sends the signal, and the E (ear) lead receives the signal.

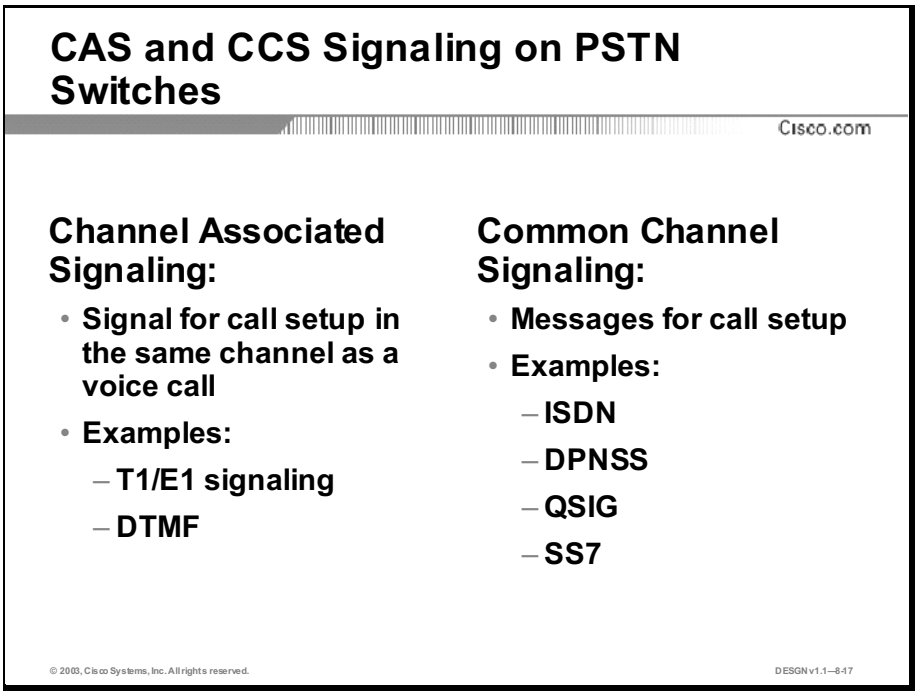

On PSTN switches, analog signaling is usually provided through current flow in closed electrical circuits, and digital signaling is provided through channel associated signaling (CAS) or common channel signaling (CCS).

Channel Associated Signaling

Different CAS varieties operate over various digital facilities. CAS associates defined bits in the T1/E1 bandwidth with the channels. Modern telecommunications networks require more efficient means of signaling, so they are moving toward CCS systems.

Here are examples of CAS signaling:

- T1 facilities using R1 signaling. This type of signaling is used in North America.
- E1 signaling
- Dual tone multifrequency (DTMF) signals used within the call path.

Common Channel Signaling

CCS uses a common link to carry signaling information messages for a number of trunks. It differs from CAS signaling because it uses a separate channel for call setup. This form of signaling may have faster connect times and the potential for a number of practically unlimited services.

Here are examples of CCS signaling:

- ISDN BRI
- ISDN PRI
- DPNSS
- QSIG
- Signaling System 7 (SS7)

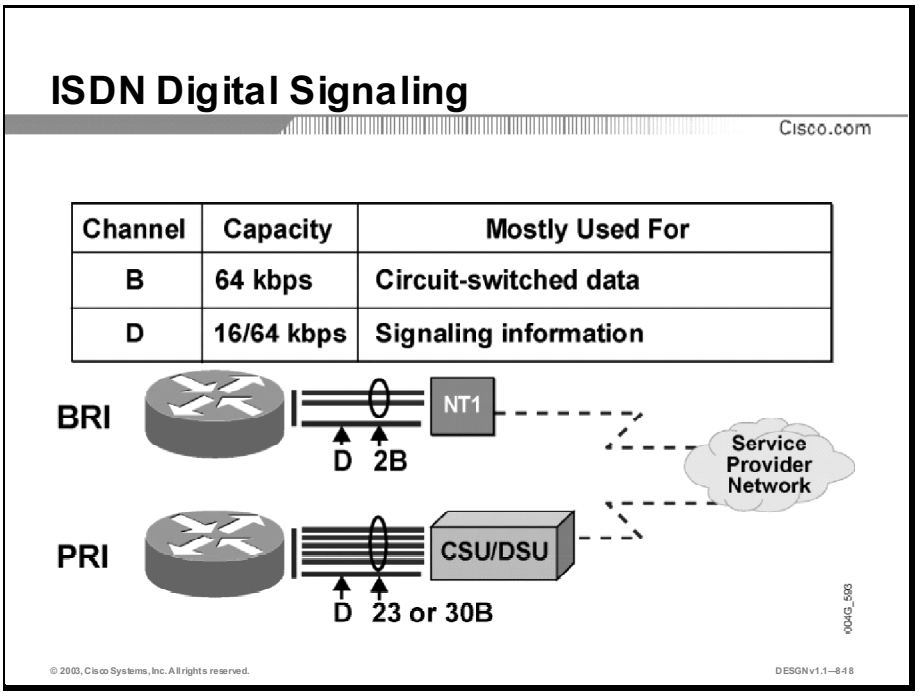

ISDN supports digital telephony and data transport services. ISDN involves the digitization of the telephone network. Digitization permits users to transmit voice, data, text, graphics, music, video, and other source material over the same facility.

ISDN enables PBXs to connect over the PSTN and create voice Virtual Private Networks (VPNs). This connection is accomplished by delivering PBX signaling over the network to distant PBXs.

ISDN Access Methods

ISDN supports two access methods:

- **ISDN BRI:** Offers two B-channels and one D-channel (2B+D). The BRI B-channel service operates at 64 kbps and carries user data and voice. The BRI D-channel service operates at 16 kbps and carries both control and signaling information. BRI is typically used for residential and small office, home office (SOHO) applications.
- **ISDN PRI:** Uses T1 or E1 circuits. The PRI service offers 23 B-channels and one D-channel (23B+D) in North America and 30 B-channels and one D-channel (30B+D) in Europe. The PRI B-channel service operates at 64 kbps and carries user data and voice. The PRI D-channel service also operates at 64 kbps and carries both control and signaling information. PRI is typically used for enterprise business applications.

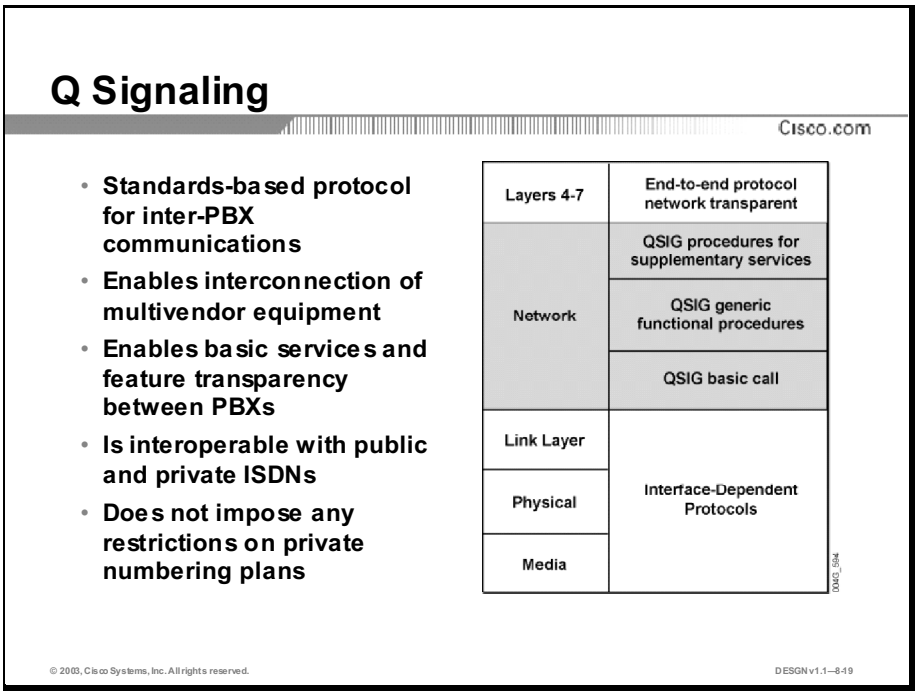

Q Signaling (QSIG) is a peer-to-peer signaling system used in corporate voice networking to provide standardized inter-PBX communications. It has a very important mechanism that provides a standard method for the transparent transportation of PBX features across a network.

Here are some QSIG features:

- Standards-based protocol, which enables interconnection of multivendor equipment
- Inter-PBX basic services, generic feature transparency, and supplementary services
- Interoperability with public and private ISDNs
- Operability in any network configuration and compatibility with many PBX-type interfaces
- No restrictions on private numbering plans

SS7 (C7) Signaling

Cisco.com

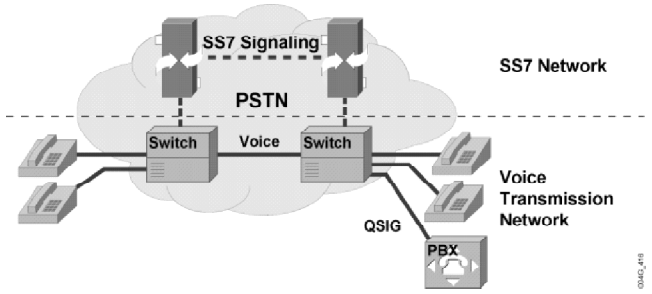

- **Used between PSTN switches**
- **Signaling implemented on a separate data network**
- **Trunk channels used solely for voice transmission**
- **Replaces per-trunk in-band signaling**

DESGN v1.1—8-20

SS7 is an international signaling standard within the PSTN. SS7 defines the architecture, network elements, interfaces, protocols, and management procedures for a network that transports control information between PSTN switches. SS7 is used between PSTN switches and replaces per-trunk in-band signaling.

Telephone service providers implement SS7 on a separate data network within the PSTN. The SS7 network provides call setup and teardown, network management, fault resolution, and traffic management services. The SS7 network is used solely for network control. Out-of-band signaling via SS7 provides numerous benefits for internetworking design, including reduced call setup time, bearer capability, and other progress indicators.

With SS7, all trunk channels are used for voice and user data while the associated signaling is carried separately over the SS7 network.

PSTN Numbering Plans

For any telephone network to function, each telephone must have a unique address. Numbering plans are unique to each country, based on the E.164 standard. This topic explains how the PSTN numbering plan works and how the PSTN routes voice calls.

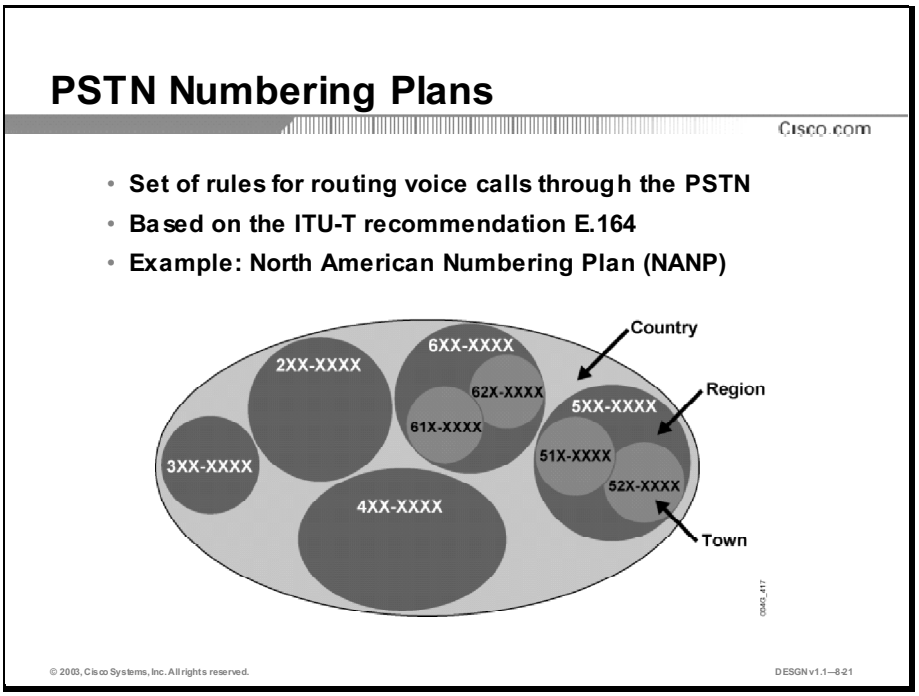

Voice addressing relies on a combination of international and national standards, local telephone company practices, and internal customer-specific codes. The International Telecommunication Union Telecommunication Standardization Sector (ITU-T) recommendation E.164 defines the international numbering plan. Each country's national numbering plan must conform to the E.164 recommendation and work in conjunction with the international numbering plan. Providers of a public switched telephone service must ensure that their numbering plan aligns with the E.164 recommendation and that the networks of each of their customers conform.

Call Routing and the Numbering Plan

Call routing is closely related to the numbering plan and signaling. Basic routing establishes a call from the source telephone to the destination telephone. However, most routing is more sophisticated and enables subscribers to select services or divert calls from one subscriber to another. Routing occurs based on a set of tables or rules within each switch. As each call arrives, the path to the desired destination and the type of feature services available are derived from these tables or rules.

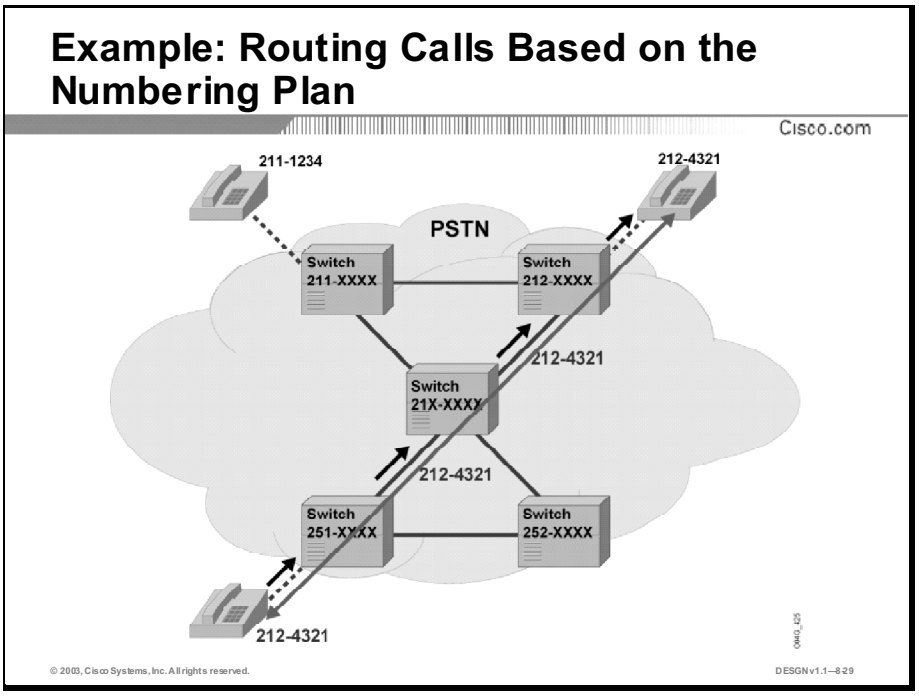

The North American Numbering Plan (NANP) is an example of the national PSTN numbering plan. It conforms to the ITU-T recommendation E.164. NANP numbers are 10 digits in length and occur in the following format: NXX-NXX-XXXX, where N is any digit from 2 to 9 and X is any digit from 0 to 9. The first three digits are called the area code. The second three digits are called the CO code or prefix. The final four digits are called the line number. The NANP is also referred to as 1+10. When a 1 is the first number dialed, a 10-digit number will follow. This enables the end-office switch to determine whether it should expect a 7- or 10-digit telephone number.

The figure illustrates how a numbering plan routes telephone calls. A PSTN switch forwards the signal as soon as it receives enough digits to make a routing decision and sends the call to the next switch. The last switch in the series receives all the digits and rings the destination telephone.

Note: SS7 uses out-of-band signaling to determine that there is a path to the destination and that the end station can accept the call, and then allocates the trunks.

PBX and PSTN Services

PBXs and the modern PSTN offer many different services, each with a desirable set of features and functionality. Service providers offer competitive services to differentiate themselves and generate additional revenue from these offerings. This topic describes the role of PBX and PSTN services and how they are implemented.

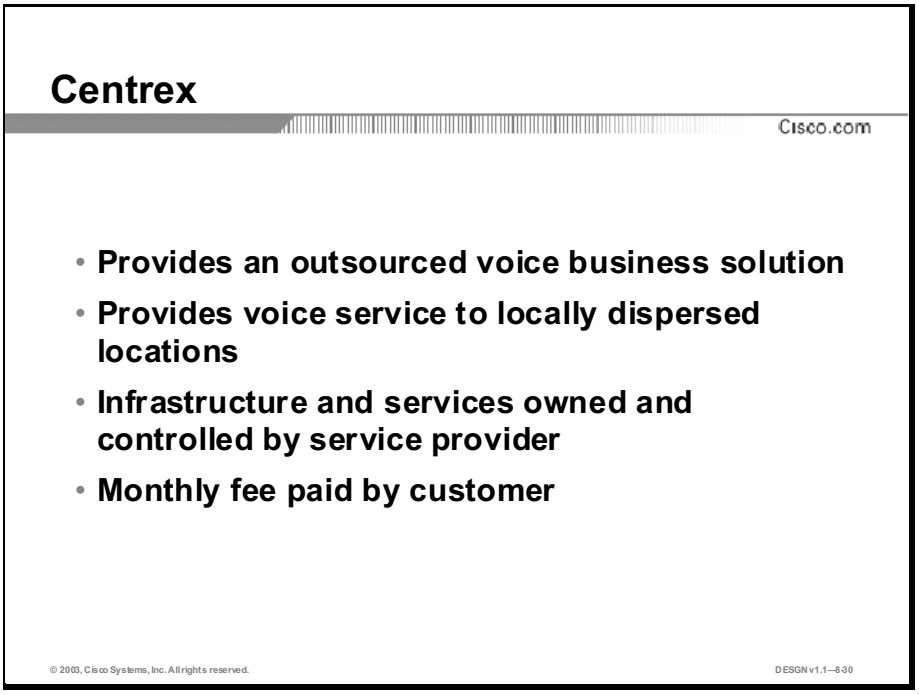

Centrex is a set of specialized business solutions primarily for voice service where the service provider owns and operates equipment providing call control and service logic functions on the service provider's premises. Centrex frees the enterprise from the costs and responsibilities of major equipment ownership.

Centrex enables the PSTN to offer features in a Closed User Group (CUG) where users can place a telephone call to all telephones within the group using only four to five digits. Although the telephones in the group may be at distant locations, they appear as a single location within the system.

Benefits of Centrex

The benefits of Centrex technology are:

- **Lower capital investment cost:** Because the service provider owns the infrastructure, the start-up costs for a Centrex service are much lower than for purchasing a traditional PBX.
- **Scalability:** Organizations can buy the exact number of lines that are needed and easily add or remove lines. In contrast, when an organization removes a station from a PBX, there are no cost savings because the equipment is owned. In addition, when an organization adds stations to a PBX, there are costs for the line cards and periodic costs for new common equipment, such as shelves to accommodate the line cards.
- **Simplicity:** The service provider is responsible for installing and configuring the service. With a PBX, the enterprise assumes this responsibility.
- **Operations and maintenance:** The service provider is responsible for the day-to-day Centrex operations and maintenance. This responsibility includes adding new lines and changing faulty components. With a PBX, the enterprise is responsible for these functions in addition to keeping an inventory of spare parts.
- **Upgrades:** Centrex service providers continue to upgrade the service. These upgrades include major evolutionary upgrades, such as analog to IP, and everyday upgrades, such as installing a new switch that offers more features.
- **Reliability:** The Centrex service provider monitors the network continuously and provides staff for immediate response to alarms and equipment failures.
- **Standardized telephones:** Centrex station equipment uses standardized protocols and conforms to an open interface. This standardization allows multiple equipment suppliers to manufacture Centrex telephones, so the enterprise can purchase any brand of telephone.
- **Space savings:** The Centrex infrastructure is located at the service provider premises. In contrast, PBX and key system solutions are located at the enterprise premises. The enterprise must provide floor space for the equipment and ensure that the storage rooms meet certain environmental requirements, such as air conditioning, humidity, and fire protection.

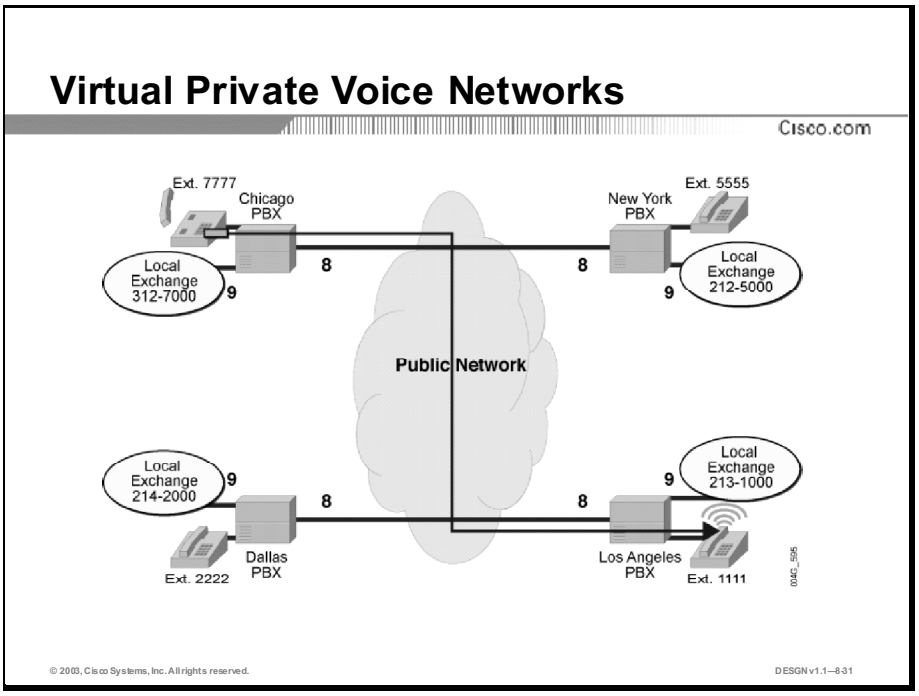

Virtual private voice networks interconnect corporate voice traffic among multiple locations over the PSTN. Virtual private voice networks are alternatives to tie lines among locations. Service providers offer competitively priced virtual private voice network services by maximizing the private use of the public infrastructure.

Benefits of Virtual Private Voice Networks

Because the same PSTN switch does not typically serve multiple locations and because tie trunks between locations are fairly expensive, virtual private voice networks are economical when multiple distant PBXs need to communicate.

Deployment of a voice-capable network eases the process of adding new and multiple sites to an existing virtual private voice network. Adding a new location and provisioning the appropriate translation and dialing plans is much easier with a virtual private voice network than with traditional tie trunks, where end-to-end connections are required between the new location and each existing location.

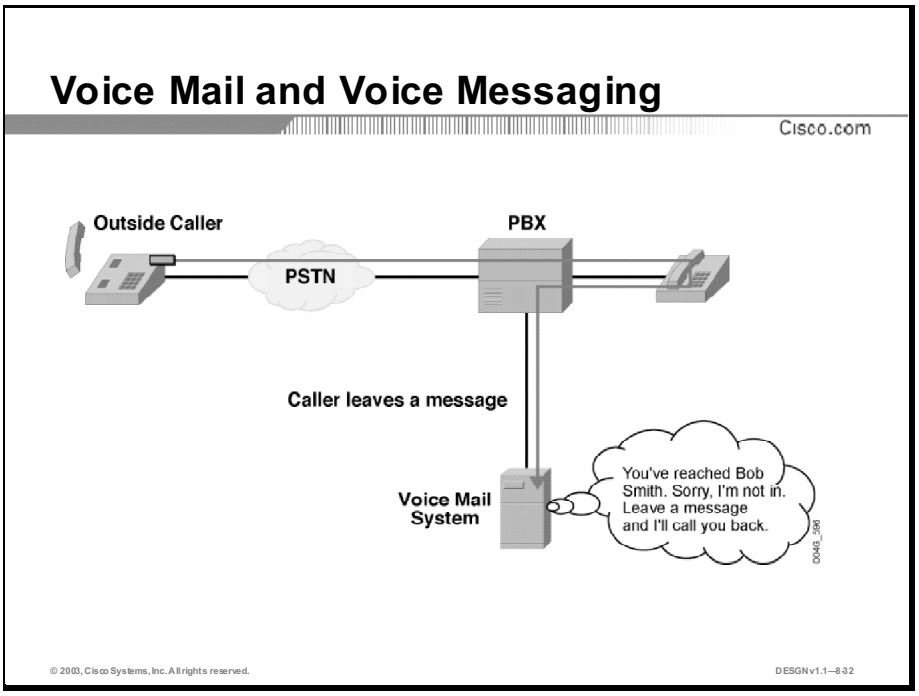

Voice messaging allows enterprises to divert their incoming PSTN calls to a voice mailbox when they are unable to answer their telephones.

Voice messaging allows residential and business subscribers access to their wireline or wireless mailbox for message retrieval via Message Waiting Indicator (MWI) services indicated with an LED, message display, special dial tone, or announcement.

Benefits of Voice Mail

The benefits of voice mail systems are:

- Improved communication because they allow people to communicate verbally in nonreal time
- Elimination of time zone and business hour issues
- Reduced labor for operators
- Fewer callbacks
- 24-hour availability

Example: Caller and Voice Mail Interaction

The steps explain how a caller might interact with a voice mail system:

Step 1 A caller places a telephone call.

Step 2 The recipient is unavailable. A prerecorded personal greeting answers: “Thank you for calling. Sorry, I am currently unavailable. Please leave me a message, and I will return your call.”

Step 3 The caller leaves a message.

Step 4 The recipient’s telephone displays an MWI lamp.

Step 5 After listening to the message, the recipient can store, delete, or forward it to another destination.

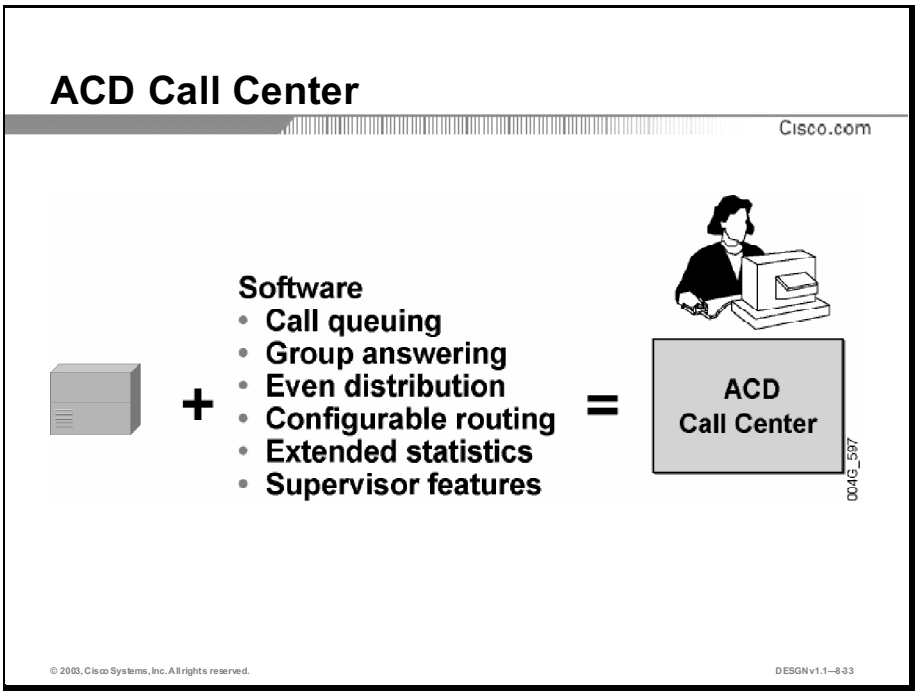

A call center is a place of doing business by telephone, combined with a centralized database that uses an automatic call distribution (ACD) system. Call centers require live agents to accept and handle calls.

An ACD greets callers with a customized announcement and then queues the calls until an agent is available to answer it. While queued, customers can hear music or customized announcements. The ACD system generally offers inbound call routing, with calls routed to a group and then routed to a specific agent within that group.

ACD Call Handling

The steps explain how a call center ACD handles an incoming call:

Step 1 The ACD accepts an incoming call.

Step 2 The ACD plays a recorded message that welcomes the caller.

Step 3 If an agent is available, the ACD assigns the call. If an agent is unavailable, the ACD queues the call.

ACD Features

An ACD typically supports these features:

- The ACD queues the calls until an available agent is assigned.
- Agents are divided into smaller skill groups, and each agent within a group can answer the call for that group.
- Call routing is configurable based on business needs.
- Statistics are available that provide information regarding the level of the customer service and the productivity of agents.

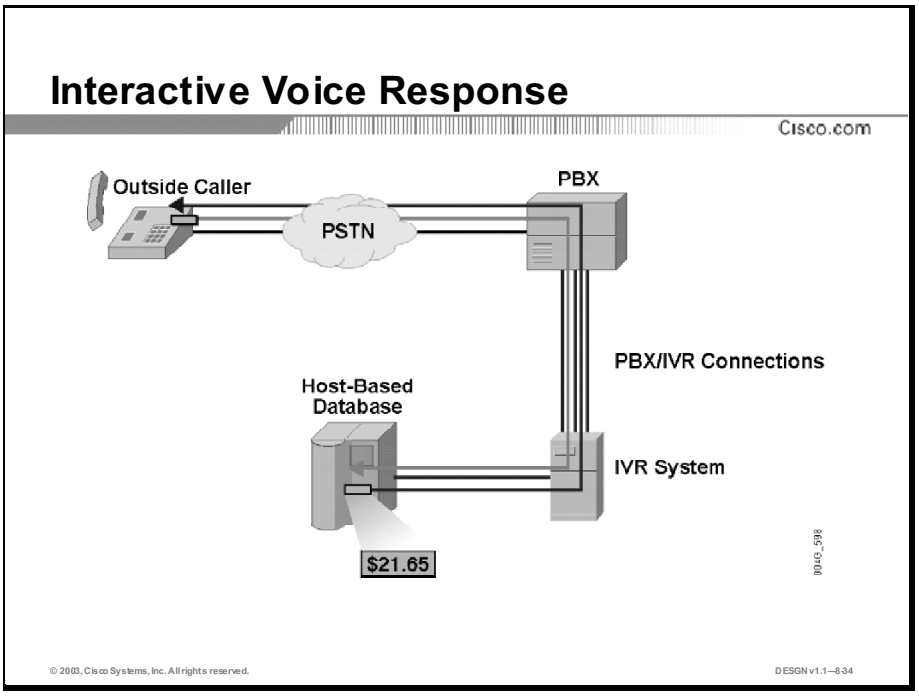

Interactive voice response (IVR) systems allow callers to exchange information over the telephone without an intermediary. The caller and the IVR system interact using a combination of spoken messages and DTMF tones. The IVR system plays a voice message that prompts a caller to enter information through the touch-tone telephone handset. Users enter numbers on a touch-tone phone, make menu selections, or answer simple directed questions. Each response takes the customer to another question. This sequence repeats until the caller receives the required information or completes the task.

Example: Banking with IVR

A customer contacts their bank to determine their checking account balance outside of normal bank hours. The customer dials a customer service telephone number and is prompted for their account number and access personal identification number (PIN). Then the IVR presents an array of options. The customer uses the touch-tone dial pad on their telephone or spoken words to make a selection and the IVR returns the information requested.

Summary

This topic summarizes the key points discussed in this lesson.

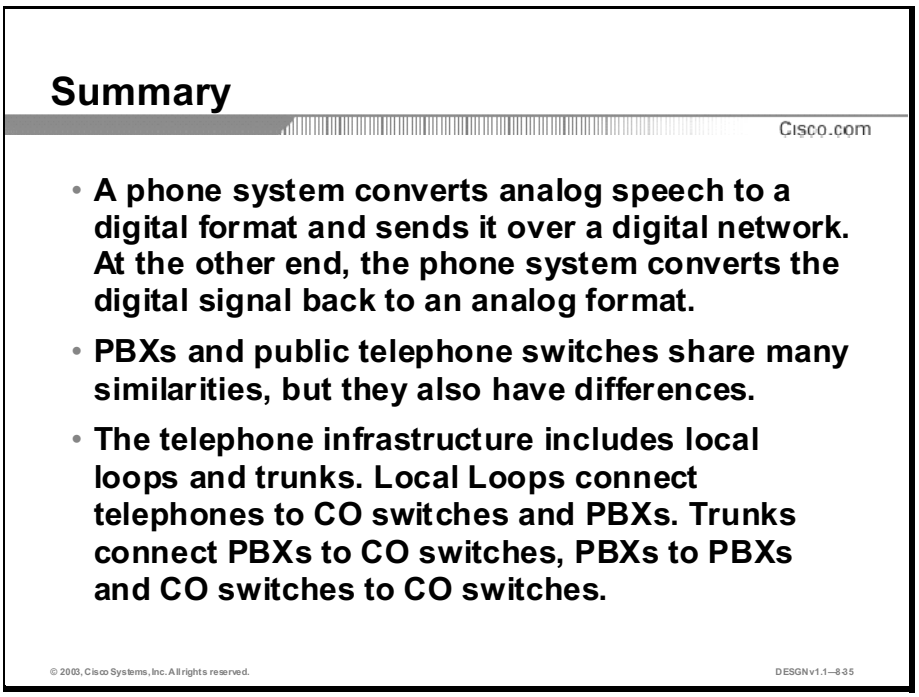

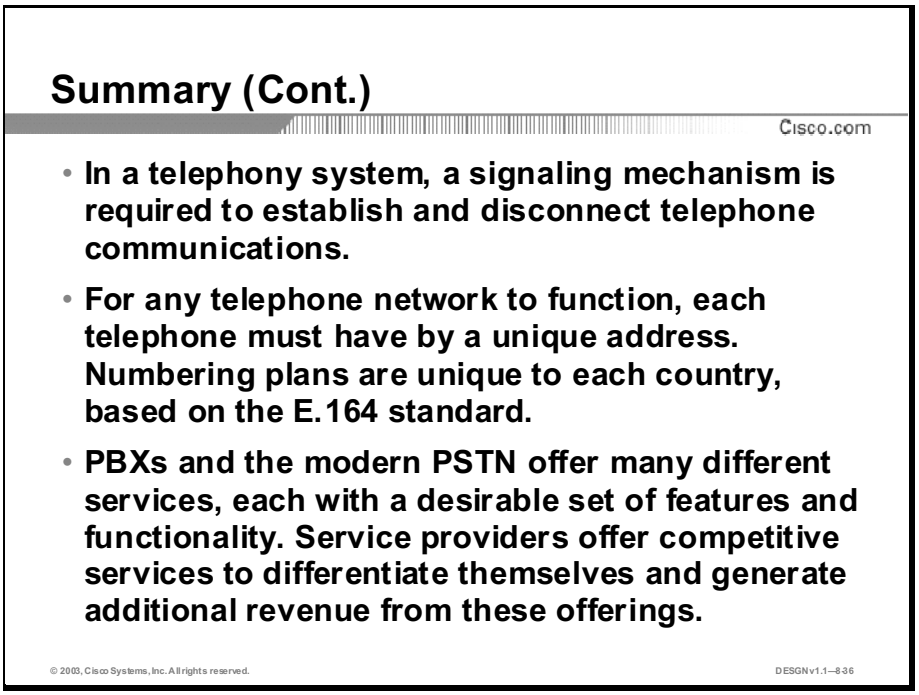

References

For additional information, refer to these resources:

- Davidson, J, and J. Peters. *Voice over IP Fundamentals*., Indianapolis, Indiana: Cisco Press; 2000.
- *Telephony Signaling (PSTN, PBX) Index Page,* http://www.cisco.com/warp/public/788/signalling/signalling.shtml
- *Voice Network Signaling and Control,* http://www.cisco.com/warp/public/788/signalling/net_signal_control.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which three processes are parts of analog to digital conversion of voice? (Choose three.)

A) companding

B) amplification

C) quantization

D) sampling

Q2) How often is the analog waveform sampled?

A) 300 times per second

B) 3400 times per second

C) 4000 times per second

D) 8000 times per second

Q3) Digital voice communication systems use a type of digital coding called _____.

A) PCM

B) companding

C) PAM

D) quantization

Q4) The process of compressing an analog signal and expanding it back to its original size is called _____.

A) quantization

B) companding

C) sampling

D) filtering

Q5) Which two statements best describe the characteristics of the PBX switch? (Choose two.)

A) is used in the private sector

B) scales up to 100,000 telephones

C) uses 46-kbps circuits

D) uses proprietary protocols between PBXs and telephones

E) is used in the PSTN

F) uses open standard protocols for its telephones

Q6) The connection between an enterprise PBX and its telephones is called the _____.

A) CO trunk

B) telephone line

C) station line

D) interswitch communication link

Q7) _____ trunks interconnect central office (CO) switches and enterprise PBXs.

A) CO

B) PSTN switch

C) PBX tie-lines

D) local loops

Q8) Within the PSTN, what signaling method is used today to control PSTN switches?

A) loop start

B) QSIG

C) ground start

D) R2 signaling

E) SS7 signaling

F) CAS signaling

Q9) Which standards-based protocol is preferred for interPBX (trunk) communications?

A) loop start

B) QSIG

C) ground start

D) E&M signaling

E) SS7 signaling

F) R2 signaling

Q10) When a telephone call is placed from location A to location B, what enables the PSTN switch to send the signal to the correct destination?

A) Each PSTN switch has a complete dialing directory that specifies the location of each telephone.

B) Using a PSTN numbering plan, the first switch routes the call, which tells the other switches how to route the call.

C) Each PSTN switch queries all its neighbors for the path to the destination telephone.

D) The first switch sends a query to a switch control point, which determines the availability of a circuit path to the destination.

Q11) _____ is a specialized business solution where the equipment providing the call control and service logic functions is owned and operated by the service provider.

A) ISDN

B) PBX

C) Centrex

D) IVR

Q12) With _____, no agent intervention is required.

A) IVR

B) call centers

C) virtual private voice networks

D) Centrex

Q13) In virtual private voice networks, the service provider offers _____.

A) telephony equipment

B) trunks

C) call routing

D) maintenance and support for enterprise PBXs

Quiz Answer Key

Q1) A,C, D

Relates to: Analog and Digital Signaling

Q2) D

Relates to: Analog and Digital Signaling

Q3) A

Relates to: Analog and Digital Signaling

Q4) B

Relates to: Analog and Digital Signaling

Q5) A, D

Relates to: PBXs and Switches

Q6) C

Relates to: Local Loops, Trunks, and Interoffice Communications

Q7) A

Relates to: Local Loops, Trunks, and Interoffice Communications

Q8) E

Relates to: Basic Telephony Signaling

Q9) B

Relates to: Basic Telephony Signaling

Q10) D

Relates to: PSTN Numbering Plans

Q11) C

Relates to: PBX and PSTN Services

Q12) A

Relates to: PBX and PSTN Services

Q13) C

Relates to: PBX and PSTN Services

Integrating Voice Architectures

Overview

Packet telephony introduces a new set of terms and standards. Each technology has a specific role to play in the network.

This lesson discusses the main drivers of the new packet telephony network. The lesson introduces the components required to successfully deploy voice on an existing data network. It provides an overview of standard voice components that are common to all packet networks and of IP telephony components found in enterprise packet networks. The lesson concludes with the introduction of VoFR and VoATM.

Relevance

By understanding integrated voice architecture concepts, components, mechanisms, and issues, you will be better prepared to design an integrated voice-data network.

Objectives

Upon completing this lesson, you will be able to identify packet telephony network drivers, goals, and design guidelines. This includes being able to meet these objectives:

- Describe why networks are migrating from separate voice and data networks to integrated networks, and list the main benefits
- Explain how speech is converted into packets and transferred over packet networks using the H.323 standard
- List additional components required for IP telephony in enterprise environments
- Describe the role of dial peers on an IP telephony network
- Describe VoIP control and transport protocols
- Describe the benefits of using VoFR when interconnecting PBXs
- Identify VoATM implementations

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with traditional telephony concepts and terms in addition to basic networking

Outline

The outline lists the topics included in this lesson.

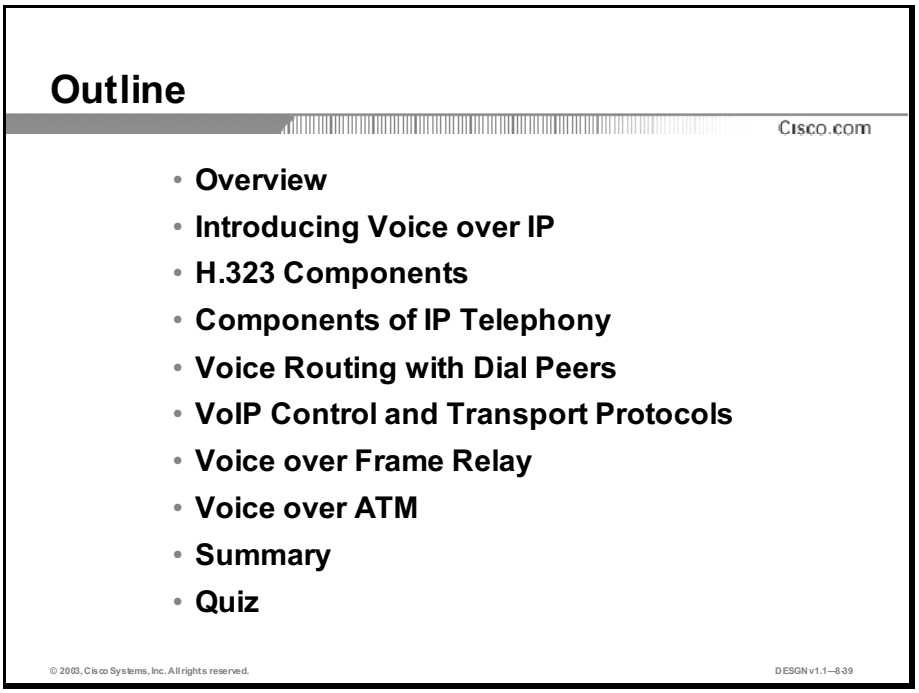

Introducing Voice over IP

Although the PSTN effectively carries voice signals, many business events are driving the need for converged voice and data networks. This topic describes why networks are migrating from separate voice and data networks to integrated networks, and lists the main benefits.

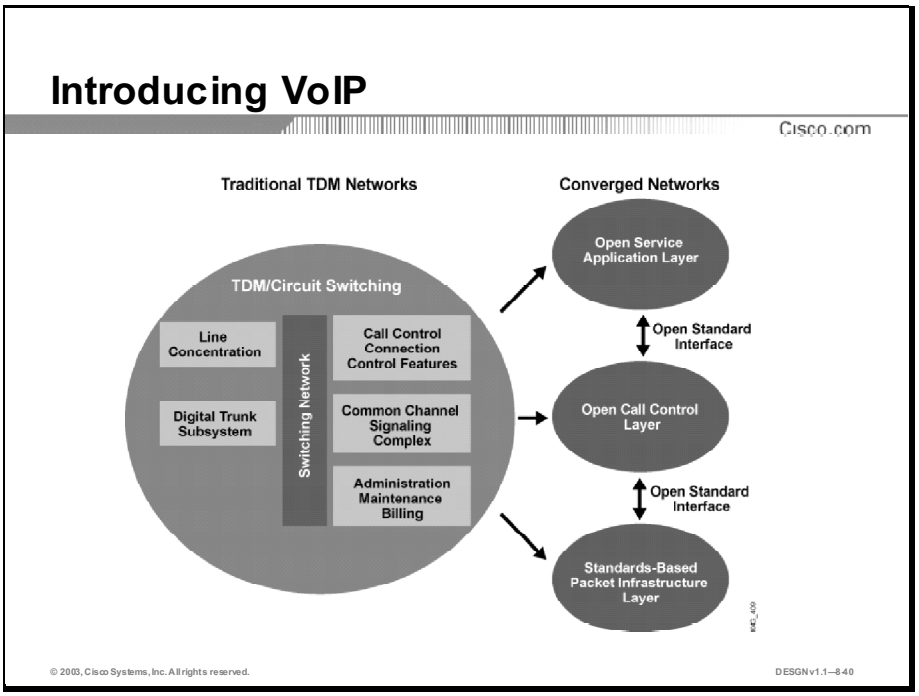

These events are driving the trend toward network convergence:

- Data has overtaken voice as the primary traffic on many voice networks.
- Companies want to reduce WAN costs by migrating to integrated networks that can carry any type of data.
- The PSTN cannot create and deploy features quickly enough.
- Data, voice, and video cannot converge on the current PSTN structure.
- The architecture built for voice is not flexible enough to carry data well.

Integrating data, voice, and video in a network changes the infrastructure and enables vendors to introduce new features. The converged network model enables you to distribute call routing, control, and applications functions that are based on industry standards. Enterprises can mix and match equipment from multiple vendors and geographically deploy these systems wherever they are needed.

On an IP network, enterprises can locate voice-call servers and application servers virtually anywhere. Over time, as with data application servers, the rationale for enterprises to maintain voice servers will diminish. As voice moves to IP networks using the public Internet for inter-enterprise traffic and private intranets for intra-enterprise traffic, service providers may host voice-call and application servers.

Converged Layers

The converged voice and data model contains three independent layers with open standard interfaces between them.

- **Packet infrastructure layer:** Replaces the circuit-switching infrastructure. The transport of upper-layer applications is based on IP because of its broad applicability. The underlying technology can be leased circuits, ATM virtual circuits, wireless technology, cable modems, SONET, Ethernet, and so on. Routers replace traditional circuit-switching and switch IP packets carrying voice.

  To minimize latency and jitter, the packet infrastructure layer requires QoS, unless bandwidth is substantially overprovisioned. The same packet infrastructure layer transports voice-control packets and the voice packets themselves.

- **Call control layer:** This layer directs voice calls to the appropriate destinations. It maps telephone numbers or user names into IP source or destination addresses, which the packet infrastructure layer understands. The main call control protocols in an IP network are H.323, Simple Gateway Control Protocol (SGCP), Media Gateway Control Protocol (MGCP), and session initiation protocol (SIP). The functions of these protocols are analogous to those of systems running SS7 in the circuit-switched world.

- **Open service application layer:** Application servers at this layer include voice mail, directory services, call distribution, accounting and billing, service provisioning, network management, and so on. Some applications keep the voice infrastructure up and running, while others provide value-added calling features.

  This layer enables developers to quickly and efficiently introduce new features. When building a network that has open interfaces from the packet layer to the call control layer, and from the call control layer to the application layer, telephony vendors no longer need to develop applications. They can write standard application programming interfaces (APIs) and allow other vendors to build the applications. The only limitations are that the applications need to run on IP and they need to conform to APIs.

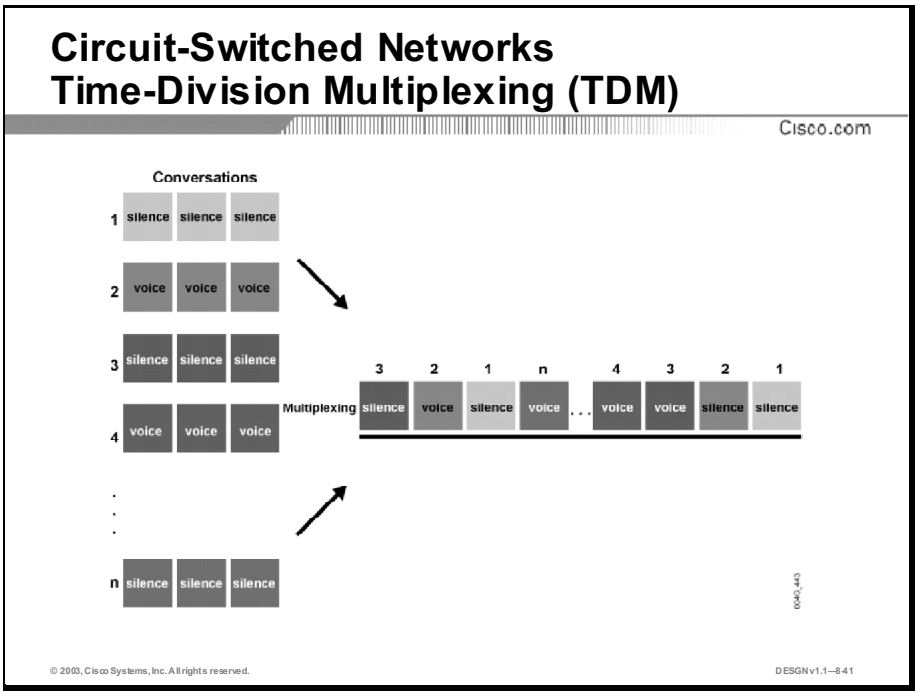

Time-division multiplexing (TDM) is a digital transmission technique for carrying multiple signals simultaneously over a single trunk line by interleaving octets of each signal into different time slots.

TDM works by converting all signals to a digital format. An analog signal is converted to a 64-kbps digital channel. Then, 24 or 30 such 64-kbps channels are multiplexed together in the telephone switch to construct the T1 or E1 trunk. In effect, information from up to 24 (or 30) different sources is placed on a single trunk, like train cars on a track.

Though TDM cannot allocate bandwidth on demand as packet switching does, its fixed bandwidth allocation ensures that a channel is never blocked because of competition for bandwidth resources and that performance does not degrade because of network congestion. Because TDM ensures time synchronization between sender and receiver, network designers often use it for delay-sensitive applications such as voice and video.

A circuit-switched call uses a dedicated 64 kbps for the entire duration of the call.

With time slot allocation, the number of simultaneous calls cannot exceed the number of TDM slots in the trunk. One call always allocates one TDM slot regardless of whether silence or speech is transmitted. Time slot allocation ensures that connections always have access to the trunk, which results in a very low delay and a low trunk efficiency. The low trunk efficiency of circuit-switched networks, is a major driver for converged packet-switched networks, where bandwidth is consumed only when the traffic flows.

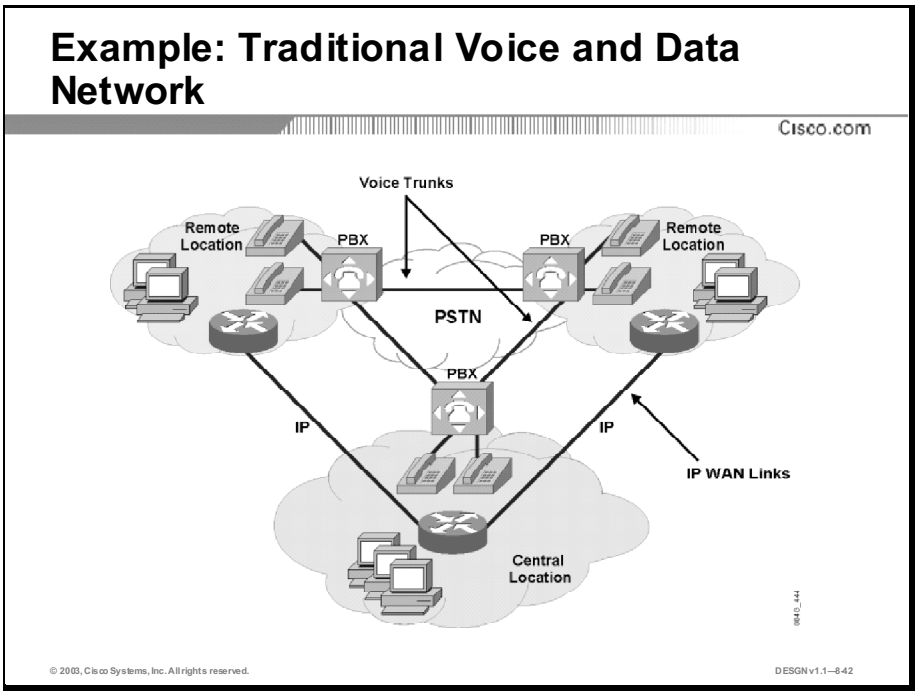

The figure illustrates a traditional voice network and an IP WAN carrying data.

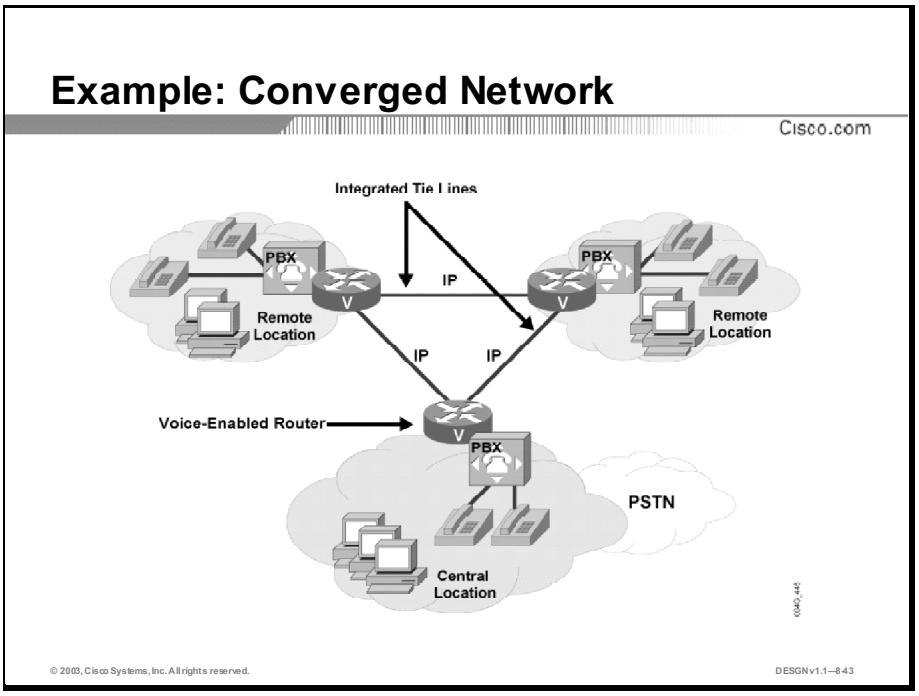

One way to create an integrated network is to replace the PBX trunk lines with IP connections. Voice traffic is converted to IP packets and directed over IP data networks through voice-enabled routers.

Another implementation solution is direct IP telephony. In this case, the IP telephones themselves convert voice into IP packets. The Cisco CallManager server is a dedicated network server running specialized call routing and PBX software. IP telephony uses no telephone cabling. Instead, all signals run over standard Ethernet network cabling.

IP telephony designs are both very cost-effective because of the reduced number of trunk lines and provide higher link efficiency. In addition, voice and data networks use the same WAN infrastructure, making it much easier to manage one single network than two separate networks with fewer administrators, a simpler management infrastructure, and lower administrator training costs.

H.323 Components

The H.323 standard is a foundation for audio, video, and data communications across IP-based networks, including the Internet. By complying with the H.323 standard, multimedia products and applications from multiple vendors can interoperate allowing users to communicate without concern for compatibility. This topic explains how speech is converted into packets and transferred over packet networks using the H.323 standard.

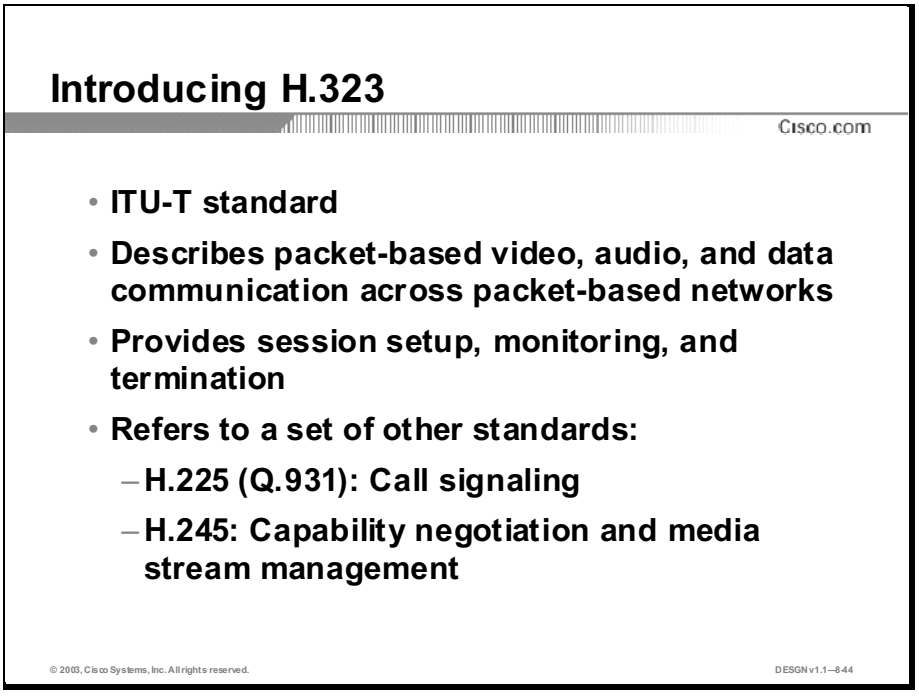

The H.323 standard is broad in scope and includes standalone devices (IP telephones, voice gateways), embedded personal computer technology (such as PCs with Microsoft NetMeeting), and point-to-point and multipoint conferences. H.323 addresses call control, multimedia management, and bandwidth management.

H.323 Call Signaling

Communications under H.323 are a mix of audio, video, data, and control signals. To establish a voice call, H.225 and H.245 signaling are required:

- **H.225 call signaling channel:** Uses the Q.931 protocol to establish a connection between two H.323 devices.
- **H.245 control channel:** A reliable channel that carries the control messages that govern the operation of the H.323 device. The control messages include capabilities exchange that negotiates audio, video and codec capabilities between endpoints; opening and closing of logical channels that carry the media stream; mode requests that request a change in mode or capability of the media stream; master/slave determination which determines which endpoint is master and which is slave, this is used to resolve conflicts during the call; and timer and counter values which establish values for timers and counters and agreement of those values by the endpoints.

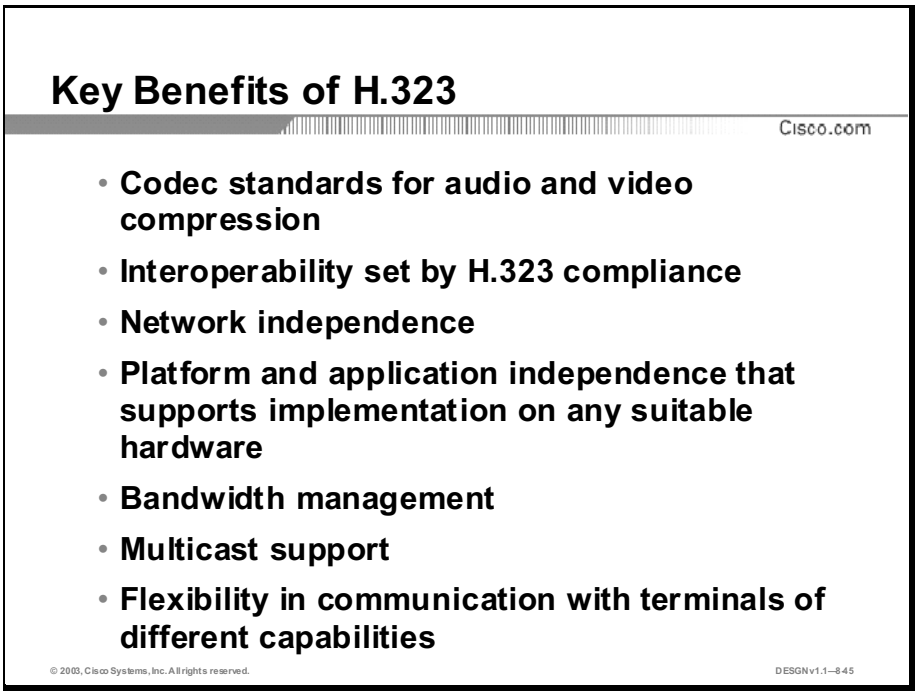

These are the key benefits of using the H.323 standard in voice, video, and data communications:

- **Codec standards:** H.323 standardizes compression and decompression of audio and video data streams, ensuring that equipment from different vendors can interoperate.
- **Interoperability:** Allows users to have a conference without worrying about compatibility at the receiving point. In addition to ensuring that the receiver can decompress information, H.323 establishes methods for receiving clients to negotiate capabilities to use with senders. The standard establishes common call setup and control protocols.
- **Network independence:** H.323 runs on top of a common network infrastructure.
- **Platform and application independence:** H.323 is not tied to any hardware or operating system. H.323-compliant platforms are available on many different devices including personal computers, IP-enabled telephone handsets, and voice-enabled gateways.
- **Bandwidth management:** Video and audio traffic is bandwidth-intensive. H.323 addresses this issue by providing bandwidth management. Network managers can limit the number of simultaneous H.323 connections within their network or the amount of bandwidth available to H.323 applications to ensure that critical traffic is not disrupted.
- **Multicast support:** H.323 supports multicast transport in multipoint conferences. Multicast transmission sends a single packet to a subset of destinations on the network without replication. By contrast, unicast transmission sends multiple point-to-point transmissions, while broadcast transmission sends to all destinations. In unicast or broadcast transmission, the network is used inefficiently because packets are replicated throughout the network. Multicast transmission uses bandwidth more efficiently because all the stations in the multicast group receive from a single data stream.
- **Flexibility:** An H.323 conference can include endpoints with different capabilities. For example, a terminal with audio-only capabilities can participate in a conference with terminals that have video and data capabilities. An H.323 multimedia terminal can share the data portion of a videoconference with a data-only terminal, while sharing voice, video, and data with other H.323 terminals.

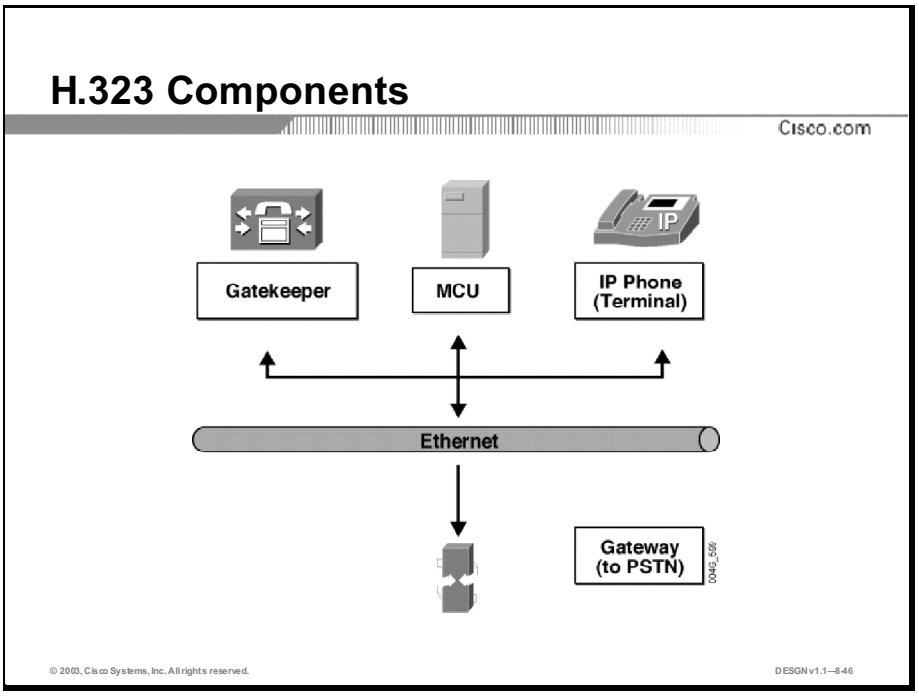

H.323 defines four major components for a network-based communications system: terminals, gateways, gatekeepers, and multipoint control units (MCUs).

Terminals

Terminals are endpoints that provide real-time, two-way voice (and optionally, video and data) communications with other endpoints such as H.323 terminals, gateways, or multipoint control units. An H.323 terminal must be able to transmit and receive G.711 64-kbps PCM-encoded voice and may support other formats, such as G.726 and G.723.1. Examples of H.323 terminals are PCs with NetMeeting software or IP telephones.

Gateways

An H.323 gateway is an optional element. Gateways provide many services such as translation between H.323 endpoints and other non-H.323 devices. This translation allows H.323 endpoints and non-H.323 endpoints to communicate. In addition, the gateway translates between audio, video, and data formats; converts call setup signals and procedures; and converts communication control signals and procedures.

Gateways are not required between two terminal connections because endpoints may directly communicate with each other. Terminals communicate with H.323 gateways using the H.245 and Q.931 protocols.

Gatekeepers

The gatekeeper is also an optional H.323 endpoint. An H.323 gatekeeper provides call control and services to H.323 endpoints. The scope of endpoints over which a gatekeeper exercises its authority is called a zone. H.323 defines a one-to-one relationship between a zone and a gatekeeper. When a gatekeeper is included, it must perform address translation, admission control, bandwidth control, and zone management. The gatekeeper may also perform call control signaling, call authorization, bandwidth management, and call management.

The gatekeeper can make decisions about balancing multiple gateways by integrating their addressing into the Domain Name System (DNS) or via Cisco IOS configuration options. For instance, if a call is routed through a gatekeeper, that gatekeeper forwards the call to the corresponding gateway based on some routing logic. In many ways, the H.323 gatekeeper functions as a virtual voice switch.

Multipoint Control Units

The MCU is an H.323 endpoint. It incorporates three functional components:

- Multipoint controller (MC): The multipoint controller supports conferences of three or more endpoints. The multipoint controller is not modeled as a standalone unit and may be located with an endpoint, terminal or gateway, gatekeeper, or MCU.
- Multipoint processor (MP): The multipoint processor can receive multiple streams of multimedia input, switch and mix the streams, and then retransmit the result to the conference members. Similar to the MC, an MP resides in an MCU.
- Multipoint control unit (MCU): The multipoint control unit is an endpoint that supports multipoint conferences by incorporating one MC and zero or more MPs.

Example: H.323 Components and Their Interactions

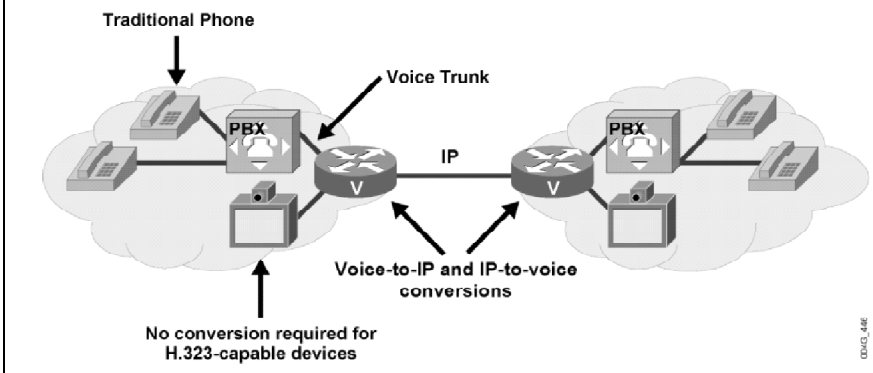

When voice is transported across the IP network using traditional telephones, a voice gateway is required on both sides. The gateway is a voice-enabled router that performs voice-to-IP and IP-to-voice conversions in a special piece of hardware called a digital signal processor (DSP). After the gateway converts voice into IP packets, it transmits the packets across the IP network. The receiving router performs the same function in the reverse order. It converts IP packets back to voice signals and forwards them through the PBX to the destination telephone.

When a network uses H.323-capable devices over IP, a voice gateway is required to provide the conversion capabilities between the IP network and the PSTN.

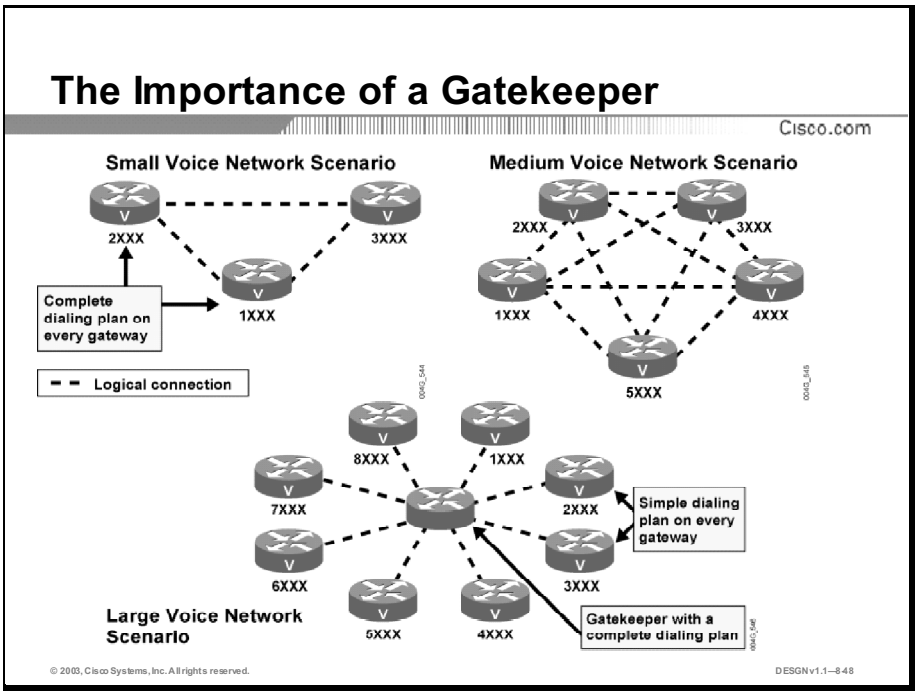

The figure illustrates different voice design options and emphasizes the importance of a gatekeeper, especially in large voice network designs. Voice network design depends primarily on the number of voice gateways and, consequently, the number of logical connections between them.

The maximum number of logical connections between voice gateways and, thus, the complexity of the network is represented by the formula (N * (N-1))/2, where N is the number of voice gateways in the system. For example, the maximum number of logical connections between three voice gateways is 3, between 5 voice gateways is 10, and between 8 voice gateways is 28. The complexity grows very fast, and adding one voice gateway to the existing network means reconfiguring all other voice gateways. Therefore, network maintenance can become very difficult.

The solution is to use a gatekeeper to store the dial plan for a zone. Gateways then only need to register with the gatekeeper, and the gatekeeper will provide all call control services to them.

Components of IP Telephony

IP telephony refers to communication services, and voice, facsimile, and voice-messaging applications that are transported via the IP network rather than the PSTN. This topic lists additional components required for IP telephony in enterprise environments.

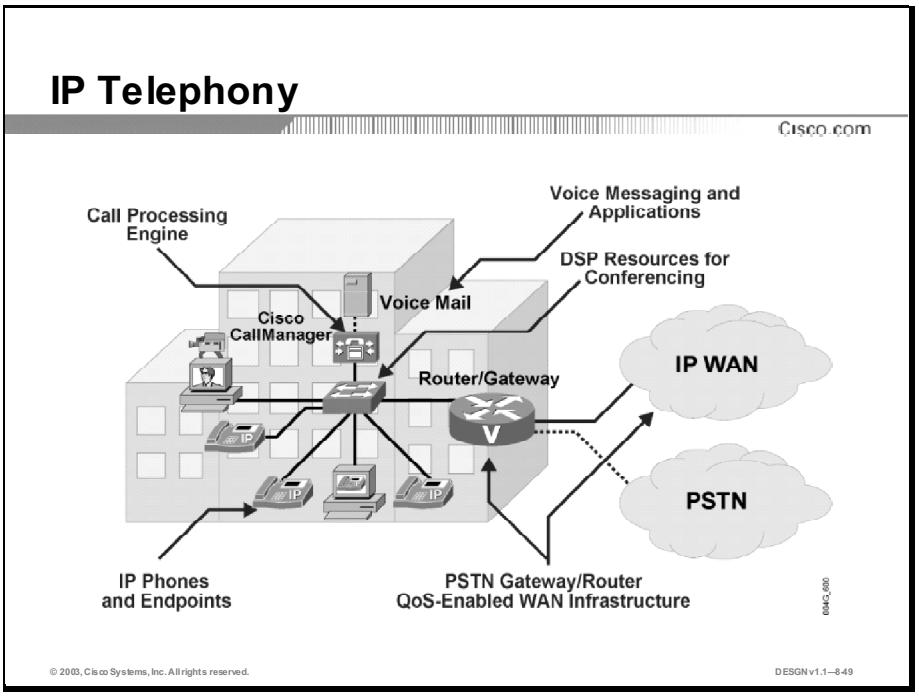

IP telephony architecture includes these distinct components:

- **Infrastructure:** The infrastructure is based on data link layer and network layer switches and voice-enabled routers that interconnect endpoints with the IP and PSTN network. Endpoints attach to the network using switched 10/100 Ethernet switch ports that sense the presence of IP devices, such as IP phones and wireless access points that require inline power. Voice-enabled routers convert voice between circuit-switched (PSTN) and IP networks.
- **Call processing:** Cisco CallManager is the software-based, call-processing component of the Cisco enterprise IP telephony solution. CallManager provides a scalable, distributable, and highly available enterprise IP telephony call-processing solution. CallManager performs in a similar manner as the PBX in a traditional telephone network.

Note: For additional Cisco CallManager features, please refer to: http://www.cisco.com/warp/public/cc/pd/nemnsw/callmn/prodlit/index.shtml.

- **Applications:** Applications use the existing IP telephony infrastructure and add features to the system. Unified messaging, voice mail, IVR, Contact Center, and Automated Attendant are among the applications that are available with IP telephony. The open service application layer allows third-party companies to develop software that interoperates with CallManager.

- **Client devices:** Client devices are IP telephones and software applications that allow communications across the IP network. Cisco CallManager centrally manages IP telephones through Ethernet connections in Building Access switches.

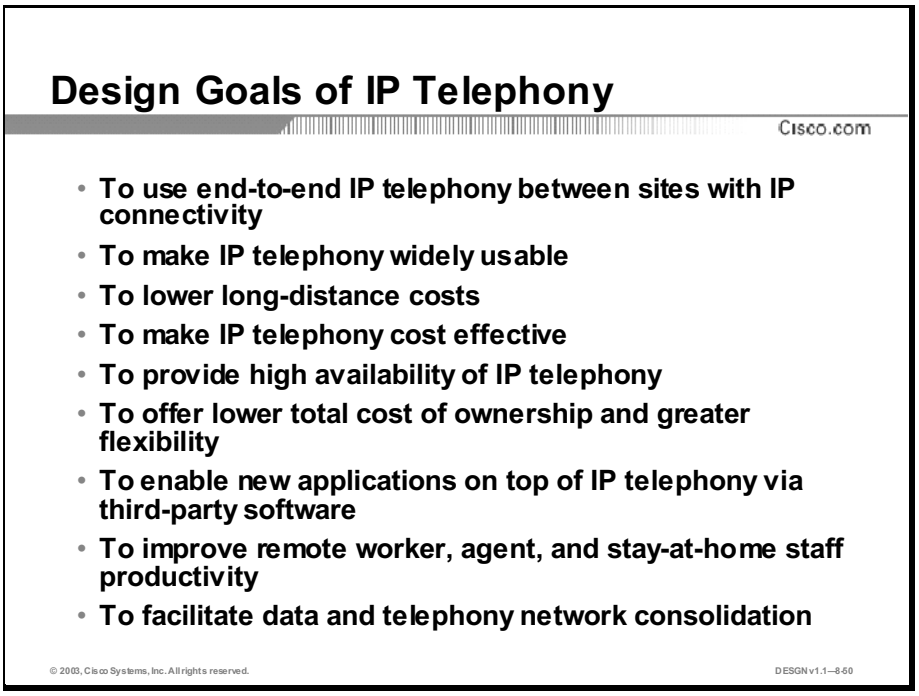

These are the overall design goals of an IP telephony network:

- **To use end-to-end IP telephony between sites where IP connectivity is already established:** You could simply deploy IP telephony as an overlaid service running on the existing infrastructure.
- **To make IP telephony widely usable:** For use to increase, voice quality needs to be on the same level as in traditional telephony.
- **To lower long-distance costs:** Enterprises can accomplish this goal by using the public Internet or private IP networks for routing telephone calls.
- **To make IP telephony cost effective:** This goal depends on more efficient use of existing WAN capacity and the cost of upgrading the existing IP network infrastructure to support IP telephony.
- **To provide high availability of IP telephony:** To meet this goal, make network components redundant, and provide back-up power to all network infrastructure components including routers, switches, and IP phones.
- To offer lower total cost of ownership and greater flexibility than traditional telephony.
- To enable new applications on top of IP telephony via third-party software.
- To improve remote worker, agent, and stay-at-home staff productivity.
- **To facilitate data and telephony network consolidation:** Thus to contribute to operational and equipment savings.

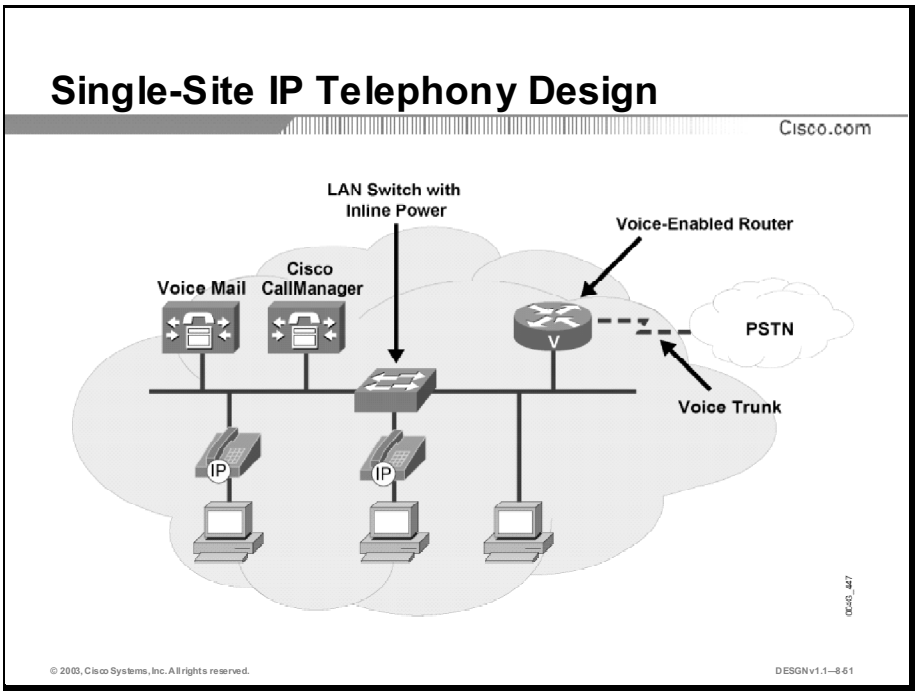

Single-site IP telephony design consists of a CallManager, IP telephones, switches with inline power, applications such as voice mail, and a voice-enabled router at the same physical location. A LAN switch powers IP telephones through the Ethernet interface. For users to make off-site calls, gateway trunks are connected to the PSTN.

Single-site deployment allows each site to be completely self-contained. Users place all calls to the outside world and remote locations across the PSTN. There is no dependency for service if an IP WAN failure occurs or there is insufficient bandwidth. There is no loss of call-processing service or functionality; the only requirements are a PSTN carrier and route diversity within the PSTN network.

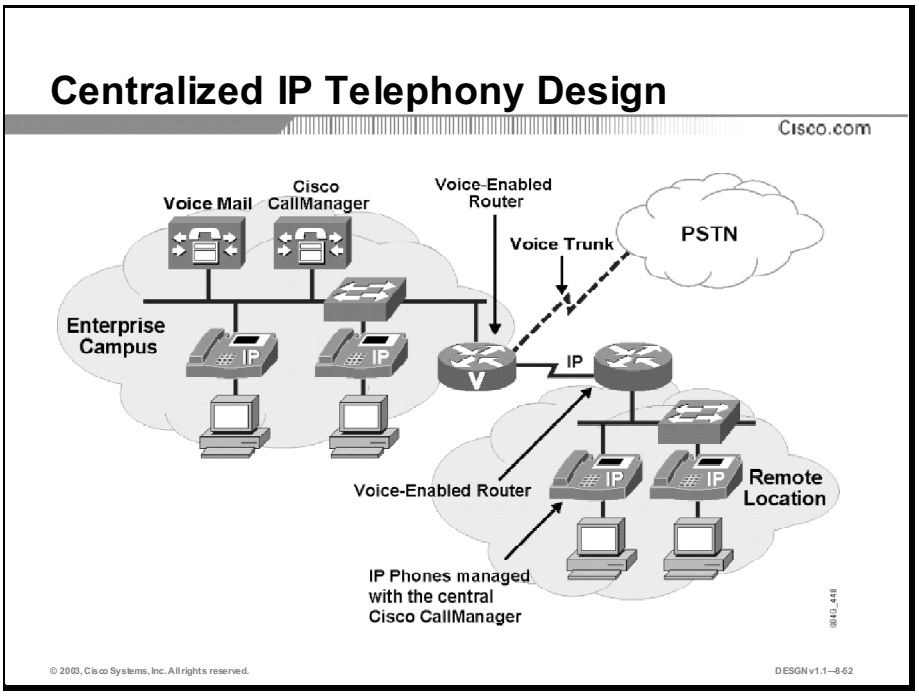

Remote IP telephones rely on the centralized CallManager to handle their call processing. Applications such as voice mail and IVR are also centralized, reducing the overall cost of ownership and centralized administration and maintenance.

The remote location requires IP connectivity with the enterprise campus network. IP telephones, powered by a local LAN switch, convert voice into IP packets and send them to the local LAN. The local router forwards the packets to the appropriate destination based on its routing table. In the event of a WAN failure, the voice-enabled router at the remote site can provide backup call processing functionality with Survivable Remote Site Telephony (SRST) services. SRST allows organizations to extend high availability IP telephony to their small branch offices by providing backup call processing functionality on voice-enabled routers.

Note: The router in the enterprise campus network is voice-capable to enable voice communication with the outside world through the PSTN.

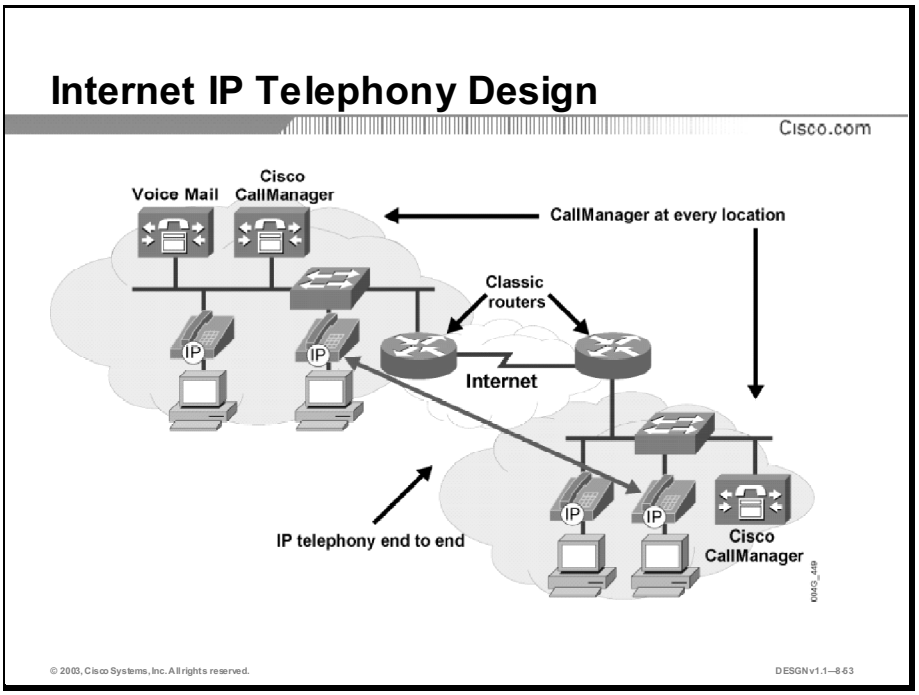

The Internet service provider (ISP) supplies Internet connectivity to enterprise locations. The ISP is not aware of IP telephony packets on its infrastructure. If an enterprise requires high-quality voice communication over the Internet, the service provider must implement QoS mechanisms. Enterprises and ISPs usually sign a service level agreement (SLA) that guarantees a bandwidth and latency levels suitable for voice transport.

Internet IP telephony requires every location to have a call-processing engine such as Cisco CallManager and IP telephones. Cisco routers that connect locations to the Internet see only data packets (some of which contain voice), so no additional hardware is generally required.

Voice Routing with Dial Peers

When an interface on the voice gateways carries voice data, it is referred to as a voice port. A voice port is a physical port that comes with a voice module, which makes a router voice-enabled. Dial peers are logical peers associated with physical voice ports. This topic describes the role of dial peers on an IP telephony network.

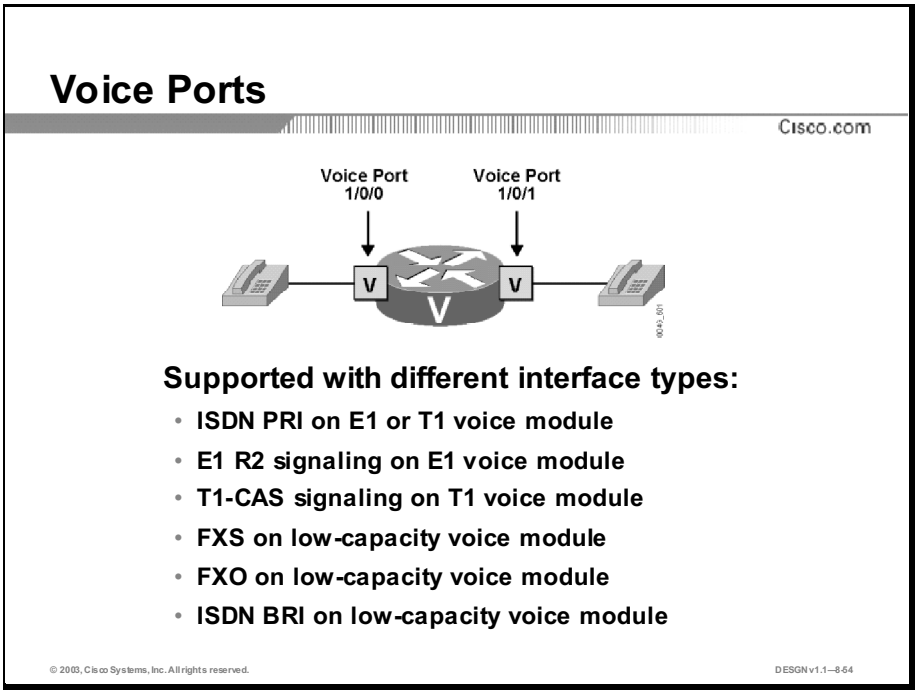

A voice module enables a voice gateway to communicate with traditional circuit-switched voice devices and networks and converts voice into IP packets and vice versa. The DSPs on the voice module code and compress voice. The voice processing of Cisco voice gateways supports the interface types and signaling listed in the figure.

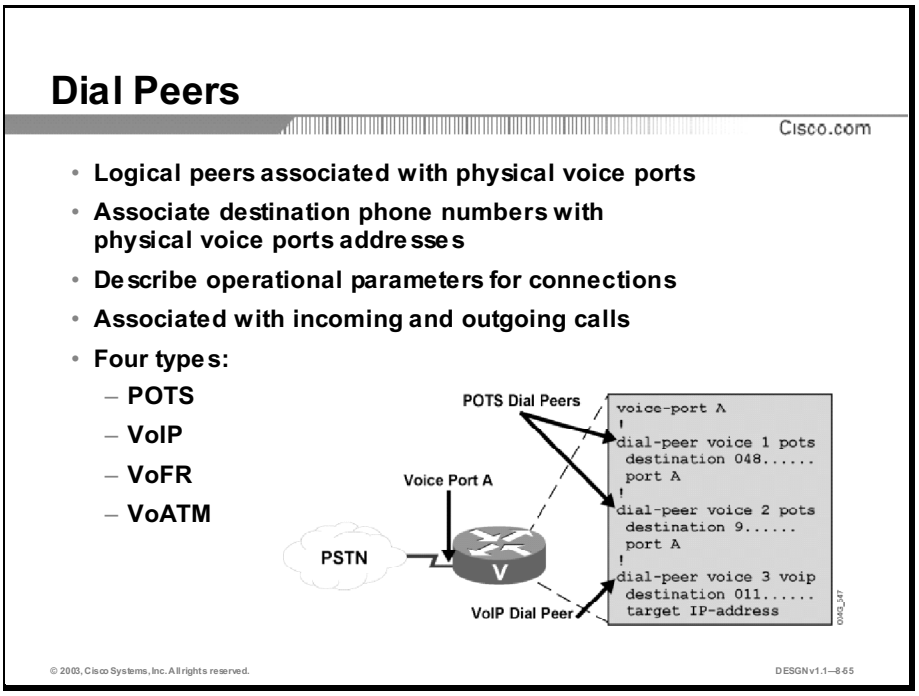

When using a voice gateway to enable communication between H.323 devices and traditional telephony, you need to configure the applicable PSTN and VoIP dial peers on the voice gateway to enable voice routing between them.

Dial peers are logical peers associated with physical voice ports. Based on the configuration of dial peers, the voice gateway establishes a connection. Four types of dial peers exist:

- **POTS:** Defines the characteristics of a traditional telephony network connection. The plain old telephone service (POTS) dial peer maps a telephone dial string to a specific voice port on the voice gateway. The voice port connects the voice gateway to the local PSTN, PBX, or telephone.
- **VoIP:** Defines how to direct VoIP calls that originate locally on the router to their destination in the IP cloud. The VoIP dial peer contains the address of the remote voice gateway or other VoIP device where the call is terminated. These are the ways to define destination IP addresses:
  - Statically configure the IP address of the gateway
  - Define the name of the gateway, which you can resolve through DNS name resolution
  - Use the registration, admission, and status (RAS) protocol. When using RAS, the gateway queries the H.323 gatekeeper to determine the destination target
- **VoFR:** Mapped to the Frame Relay data-link connection identifier (DLCI) of the interface from which the call exits the router. The destination telephone number is also mapped with the peer.
- **VoATM:** Mapped to the ATM virtual circuit of the interface from which the call exits the router. The destination telephone number is also mapped to the peer.

Relationship Between Dial Peers and Call Legs

Cisco.com

- **Call legs are router-centric (from the router perspective):**
  - **Inbound call leg originates outside the voice gateway.**
  - **Outbound call leg originates from the voice gateway.**
- **Four call legs conferenced together create an end-to-end call through the router.**

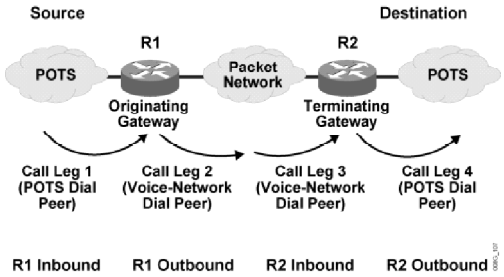

A voice call over a packet network is segmented into discrete call legs that are associated with dial peers. A call leg is a logical connection between two voice gateways or between a voice gateway and an IP telephony device. Dial peers are used to apply attributes to call legs and to identify call origin and destination. Attributes applied to a call include QoS, codec, voice activity detection (VAD), and fax rate.

There are two types of call legs: inbound and outbound. An inbound call leg originates outside the voice gateway, while an outbound call leg originates from the voice gateway.

- For inbound call legs, a dial peer might be associated to the calling number or the port designation.
- Outbound call legs always have an associated dial peer, and the destination pattern is used to identify the outbound dial peer. The call is associated with the outbound dial peer at setup time.

Unless a voice gateway is otherwise configured, when a call arrives, the gateway presents a dial tone and collects digits until it can identify the destination dial peer. Then, the gateway forwards the call through the next call leg to the destination.

An end-to-end voice call consists of four call legs: two from the originating router (R1) or gateway perspective and two from the terminating router (R2) or gateway perspective. An inbound call leg originates when an incoming call comes into the router or gateway. An outbound call leg originates when a call is placed from the router or gateway.

VoIP Control and Transport Protocols

Voice communication over IP relies on control signals and voice that is coded and encapsulated into IP packets. This topic describes the coding mechanisms used in packet telephony networks.

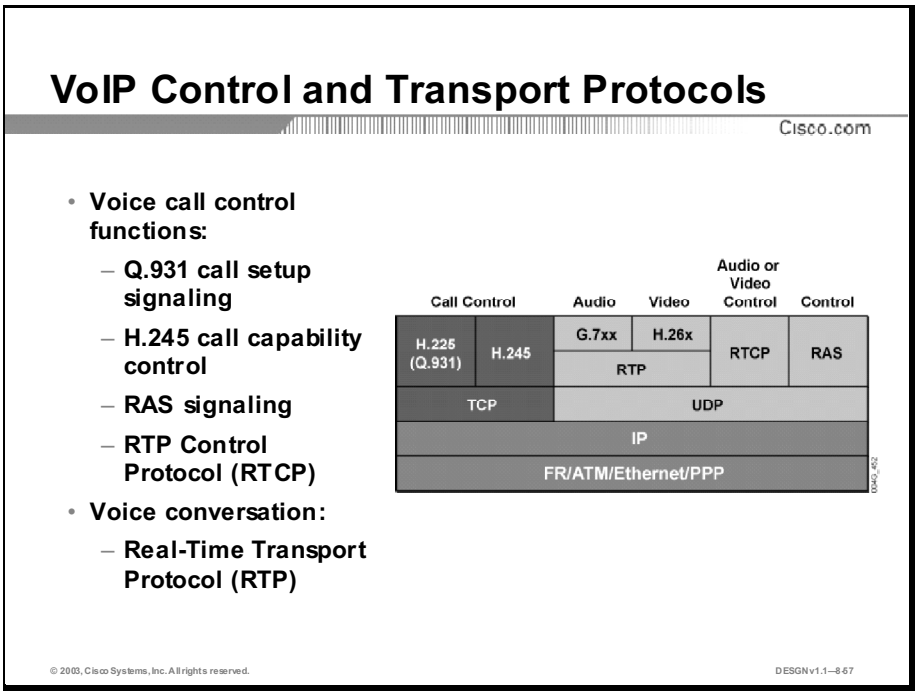

Control signals and data require reliable TCP/IP transport because users must receive the signals in the order in which they were sent, without any loss. However, voice loses its value with time. If a voice packet is delayed, it may lose its relevance to the recipient. Thus, voice conversation uses the more efficient, unreliable User Datagram Protocol (UDP)/IP transport.

Call Control Functions

Call control functions include signaling for call setup, capability exchange, signaling of commands and indications, and messages to open and describe the content of logical channels. Overall system control is provided by three separate signaling functions:

- **H.225 call signaling channel:** Uses Q.931 to establish a connection between two terminals.
- **H.245 control channel:** A reliable channel that carries control messages governing voice operation, including capabilities exchange, opening and closing of logical channels, preference requests, flow control messages, and general commands and indications. Capabilities exchange is one of the fundamental capabilities in the ITU recommendation.
- **RAS signaling:** Performs registration, admission, bandwidth changes, status, and disengage procedures between gateways and gatekeepers. The RAS protocol runs on UDP/IP. RAS is used only if an H.323 gatekeeper is present.
- **RTP Control Protocol (RTCP):** Provides a mechanism for hosts involved in a Real-Time Transport Protocol (RTP) session to exchange information about monitoring and controlling the session. RTCP monitors quality for such elements as packet counts, packet loss, and interarrival jitter.

Voice Conversation with RTP

RTP, which runs on top of UDP/IP, provides voice conversation between two IP endpoints. Because of the time-sensitive nature of voice transport, UDP/IP is the logical choice to carry voice.

With voice conversation, more information is needed on a packet-by-packet basis than UDP offers. Therefore, RTP carries packet sequence and time-stamping information. It uses sequence information to determine whether the packets are arriving in order, and it uses the time-stamping information to determine the interarrival packet time (jitter). These two bits of information are essential for high-quality VoIP conversations.

Using RTP is important for real-time traffic; however, a few drawbacks exist. The IP/UDP/RTP packet headers are 20, 8, and 12 bytes, respectively. These packet header sizes add up to a 40-byte header, which is twice as big as the payload (compressed voice) when using the G.729 codec with a 2 predictor, 20-ms payload. This large header adds considerable overhead to the voice traffic and reduces voice bandwidth efficiency.

Note: You can compress large IP/UDP/RTP headers by using RTP header compression (compressed Real-Time Transport Protocol [cRTP]) to 2 bytes without UDP checksums or 4 bytes with UDP checksums.

Voice over Frame Relay

Corporate data networks commonly rely on Frame Relay for its flexible bandwidth, widespread accessibility, support of a diverse traffic mix, and technological maturity. Enterprises can use the same infrastructure for voice transport with VoFR. This topic describes the benefits of using VoFR when interconnecting PBXs.

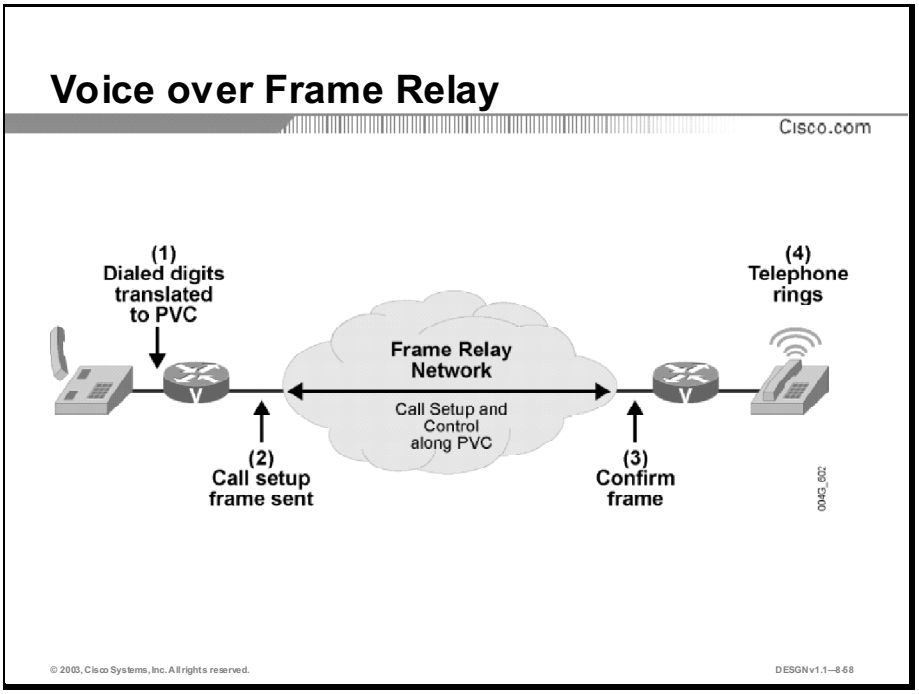

Designers often confuse VoFR with VoIP that runs over Frame Relay. No IP is involved in any stage of VoFR as it carries voice from source to destination. In VoFR implementations, a voice-enabled router encodes the voice received from the PBX directly into Frame Relay frames and forwards those frames onto the Frame Relay network.

Frame Relay is not a LAN solution and, therefore, you cannot implement it from end to end like VoIP. VoFR replaces traditional telephony trunks between PBXs and keeps the existing telephone equipment in place. VoFR is a logical progression for corporations already running data over Frame Relay and starts the process of converging voice and data networks.

Recommendations for VoFR

Deploy Frame Relay at low speeds (from about 64 kbps to 512 kbps). The key component in Frame Relay is the committed information rate (CIR) purchased for the link.

To ensure proper voice quality, the traffic stream must adhere strictly to the traffic shaping configuration. Traffic shaping shapes the total permanent virtual circuit (PVC) traffic to conform to the CIR, committed burst (Bc), and excess burst (Be). The router should mark streams as discard eligible (DE) during shaping to ensure that data traffic and not voice is dropped if the traffic exceeds the traffic contract.

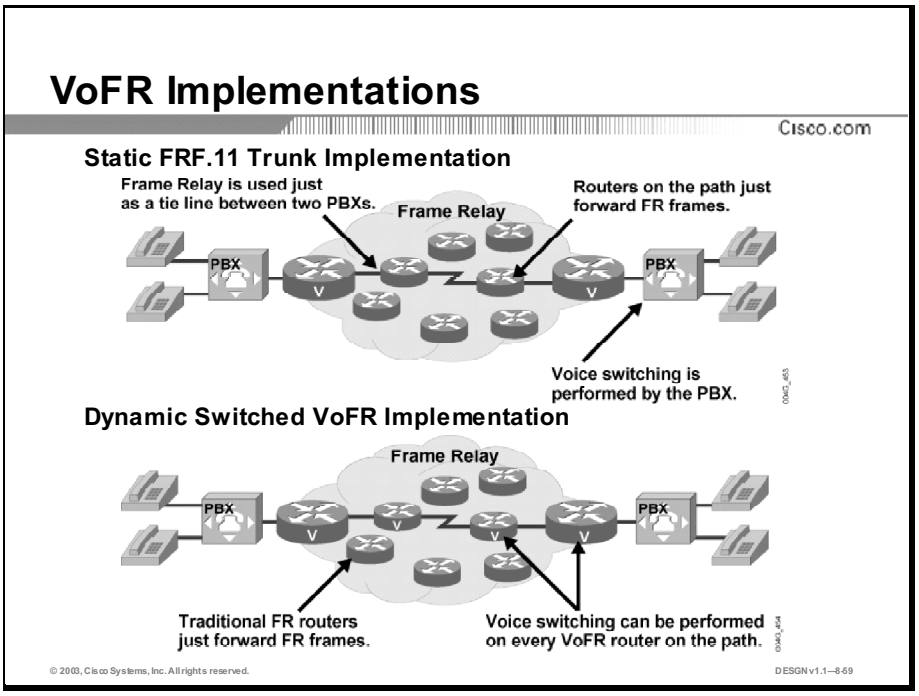

VoFR supports two main voice implementations:

- **Static FRF.11 trunks:** Used with permanent switched trunks to create fixed point-to-point connections, which are typically used to connect two PBXs. In this case, the VoFR system transports the voice connection channels but does not provide telephone call switching based on dial plan information. This functionality is sometimes referred to as “tie-line emulation.” In this scenario, PBXs perform all telephone call switching.
- **Dynamic switched VoFR calls:** The VoFR system includes dial plan information that is used to process and route calls based on the telephone numbers that callers dial. The dial plan information is contained within dial peer entries.

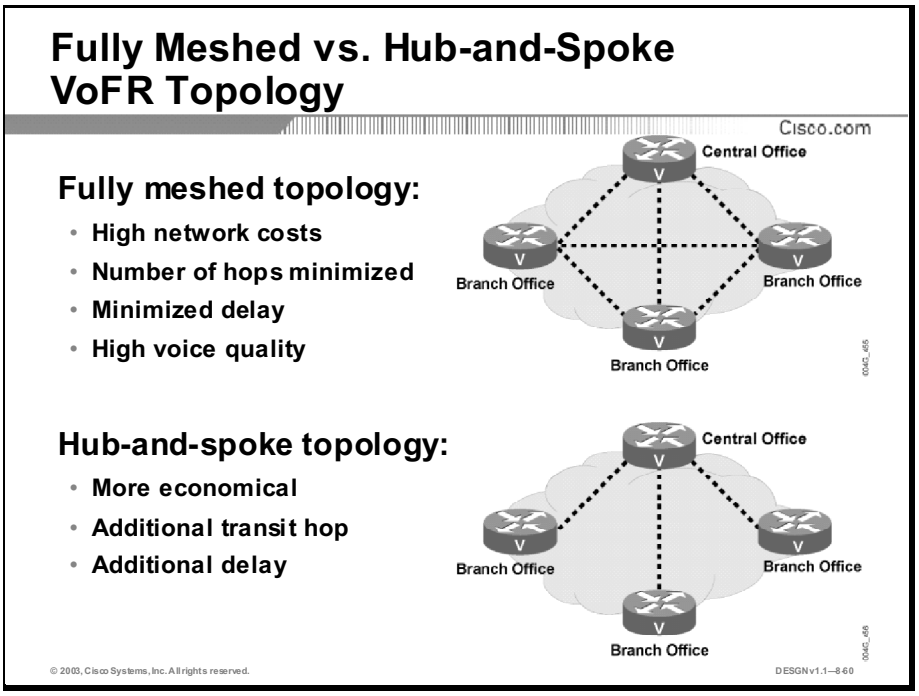

The two extremes for designing a VoFR network are:

- **Fully virtual meshed topology:** This topology is feasible in a private Frame Relay network because the enterprise can theoretically define a DLCI for each telephone number and establish any-to-any connectivity without paying surcharges. Implementing this type of interconnectivity in the public Frame Relay network is similar to building a leased-line meshed network, which can be prohibitively expensive. This solution minimizes the number of network transit hops and maximizes the ability to establish different qualities of service. A fully virtual meshed topology network minimizes delay and improves voice quality but produces the highest network cost.
- **Hub-and-spoke topology:** This topology is the optimal network topology for VoFR. It assumes that the majority of voice and data traffic is between a branch office and the central headquarters. It is more economical to designate single PVCs (DLCIs) between individual branches and the CO than to construct a network that maps PVCs (DLCIs) from branch to branch. Most Frame Relay providers charge based on the number of PVCs used. To reduce costs, you can configure both data and voice segments to use the same PVC, thereby reducing the number of PVCs required. In this design, the central site switch reroutes voice calls. This design can create a transit hop when voice needs to go from one remote office to another remote office, adding some extra delay.

  Traffic shaping is important in a hub-and-spoke topology. In most hub-and-spoke topologies, the hub typically has higher speed connection to the Frame Relay network than the spokes. As a result, traffic from the hub to a spoke can easily overwhelm the connection. The combined connection speeds of the spoke sites can easily overwhelm the hub. Therefore, you should deploy traffic shaping from both the hub to the spokes and from the spokes to the hub to maintain call quality.

Voice over ATM

VoATM enables you to transport voice traffic, such as telephone calls and faxes, over an ATM network. This topic identifies VoATM implementations.

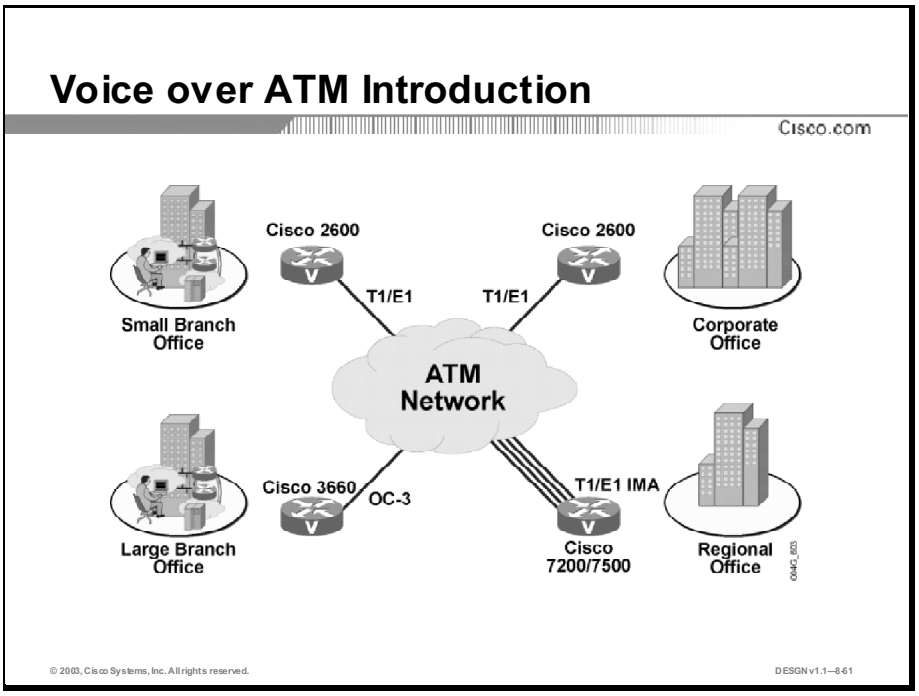

ATM is a multiservice, high-speed, scalable technology. It is a dominant switching fabric in carrier backbones and supports services with different transfer characteristics. ATM transports voice, data, graphics, and video simultaneously at very high speeds. It transmits information in fixed-length (53-byte) cells, based on application demand and priority. Five bytes are used for the ATM cell header, leaving 48 bytes for voice or data.

Characteristics of ATM

These are the basic characteristics of ATM:

- Uses small, fixed-sized cells (53 bytes)
- Is connection-oriented
- Supports multiple service types
- Applies to LAN and WAN traffic
- PSTN circuits emulated by ATM virtual circuits
- Minimizes delay and delay variation (jitter)

ATM Classes of Service

The ATM Forum and the ITU have specified different classes of service to represent different possible traffic types for VoATM:

- **Constant bit rate (CBR) and variable bit rate (VBR) classes:** Designed primarily for voice communications, these classes have provisions for transporting real-time traffic and are suitable for guaranteeing a certain level of service. CBR, in particular, allows the amount of bandwidth, end-to-end delay, and delay variation to be specified during the call setup.
- **Unspecified bit rate (UBR) and available bit rate (ABR) classes:** Designed principally for bursty traffic, these classes are more suitable for data applications. UBR, in particular, makes no guarantees about the delivery of the data traffic.

ATM networks typically have low overall jitter, and VBR real-time service from the network ensures that voice packets are transported efficiently within the backbone to minimize delay and jitter.

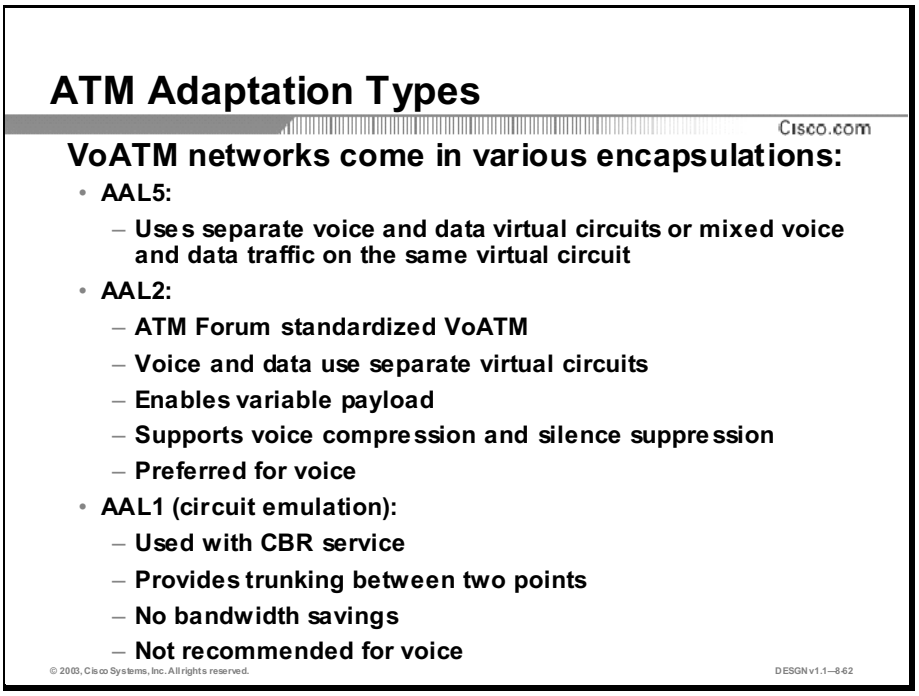

The method of transporting information through an ATM network depends on the nature of the traffic. Different ATM adaptation types have been developed for different traffic types, each with its benefits and drawbacks.

ATM adaptation layer 1 (AAL1) is the most common adaptation layer that network designers use with CBR services, while ATM adaptation layer 2 and layer 5 (AAL2 and AAL5) are widely used for voice over packet transports.

- **AAL5 adaptation type:** A solution that uses separate voice and data virtual circuits or mixed voice and data traffic on the same virtual circuit. This method of VoATM is most often used for "toll bypass" enterprise networks.
- **AAL2 adaptation type:** The ATM Forum standardized VoATM with the use of AAL2-based cells. These solutions are standards-based and are deployed for PSTN access where voice and data use separate virtual circuits, because data traffic does not use AAL2 cells. AAL2 enables a variable payload within cells. This functionality improves bandwidth efficiency over structured or unstructured circuit emulation using AAL1. In addition, AAL2 supports voice compression and silence suppression. It supports multiple voice channels with varying bandwidths on a single ATM connection.
- **AAL1 adaptation type:** This mode is also called circuit emulation service. It uses ATM to provide trunking between two points and typically uses AAL1 cells with CBR service. CBR, the highest-quality class of the ATM service, provides circuit emulation service (CES), which transmits a continuous bit-stream of information. CES allocates a constant amount of bandwidth to a connection for the duration of a transmission. Although it guarantees high-quality voice, CES monopolizes bandwidth that could be used for other applications. In the interest of reducing delay, CES might send the fixed-size ATM cells half empty rather than waiting 6 ms for 47 bytes of voice to fill the cell. This can waste more than 20 bytes of bandwidth per ATM cell. Using AAL1 for VoATM increases the overhead of voice transmissions and wastes bandwidth.

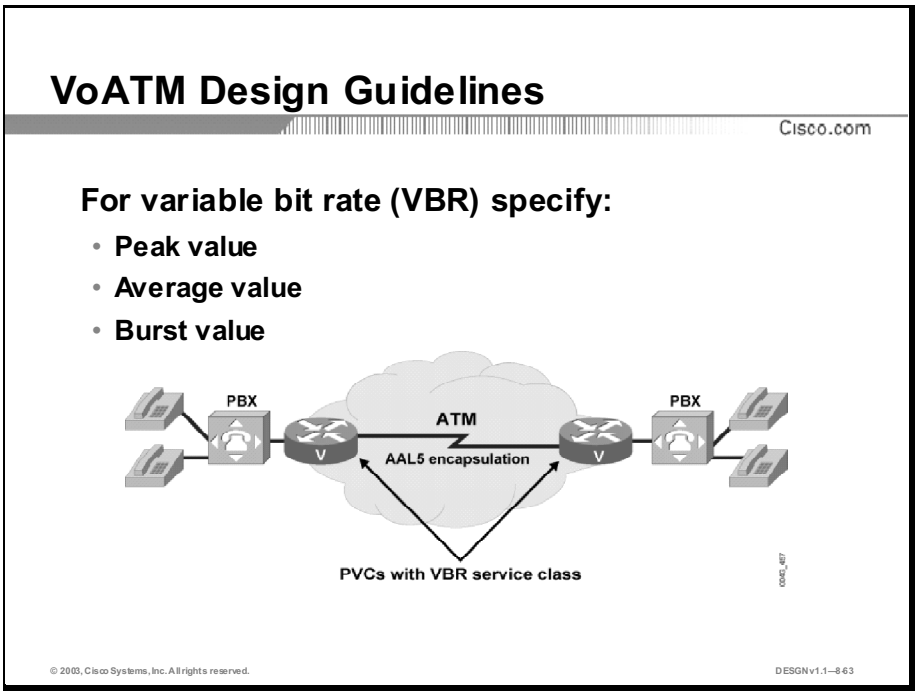

As voice traffic is sent over ATM, the signal is encapsulated using a special AAL5 encapsulation for multiplexed voice. You must configure the ATM PVC to support real-time voice traffic and must assign AAL5 voice encapsulation to the PVC. You must also configure the PVC to support VBR for real-time networks for traffic shaping between voice and data PVCs.

Traffic Shaping for VoATM

Traffic shaping ensures that the carrier does not discard the incoming calls. To configure voice and data traffic shaping, you must configure the peak, average, and burst options for voice traffic. The burst value is mandatory if the PVC carries bursty traffic. Based on these values, the PVC can effectively handle the bandwidth for the number of voice calls.

Role of the Voice Gateway in VoATM

To send VoATM, a voice gateway must code voice into cells for transport and decode those cells at the destination. A voice gateway must also handle the telephone signaling that the voice source and destination use to receive the called party's number and deliver call progress signals. A voice gateway must understand the signaling or addressing needed within the ATM network cloud to reach the destination voice gateways. This capability is important when the gateway is translating between a traditional voice network and an ATM network.

Summary

This topic summarizes the key points discussed in this lesson.

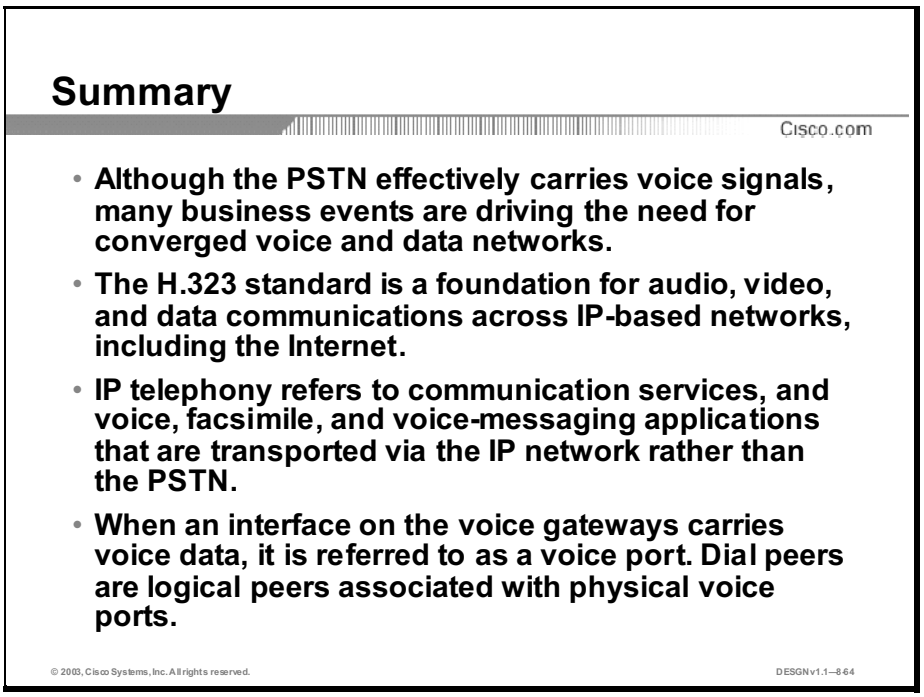

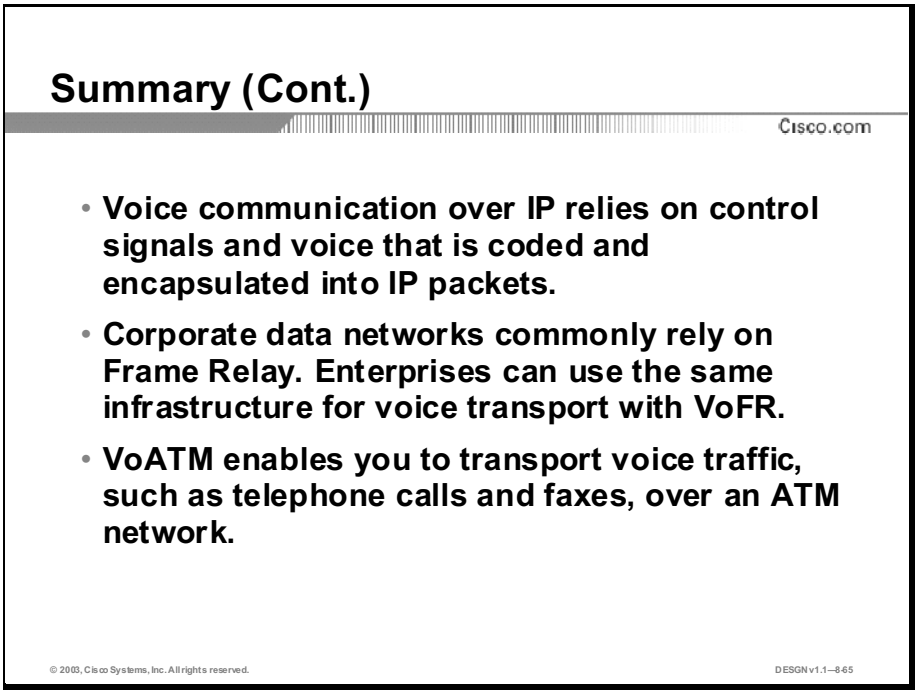

References

For additional information, refer to these resources:

- Davidson, J., and J. Peters. *Voice over IP Fundamentals*. Indianapolis, Indiana: Cisco Press; 2000.
- *Voice—Understanding How Inbound and Outbound Dial Peers Are Matched on Cisco IOS Platforms,* http://www.cisco.com/warp/public/788/voip/in_dial_peer_match.html
- *Architecture for Voice, Video and Integrated Data,* http://www.cisco.com/warp/public/cc/so/neso/vvda/iptl/avvid_wp.htm
- *VoIP—Understanding Codecs: Complexity, Support, MOS, and Negotiation,* http://www.cisco.com/warp/public/788/voip/codec_complexity.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Select the two benefits of using packet-switched networks. (Choose two.)

A) guaranteed access

B) open standards among the three layers

C) merged application, call control, and packet infrastructure layers

D) toll voice quality

E) fewer circuits

F) better tie line efficiency

G) low delay

Q2) Match the H.323 components with their descriptions.

A) terminal

B) gateway

C) gatekeeper

D) MCU

_____ 1. translates between H.323 endpoints and other non-H.323 devices, allowing them to communicate

_____ 2. enables three or more terminals and gateways to participate in a multipoint H.323 conference

_____ 3. a client endpoint on the LAN that helps provide real-time, two-way H.323 communications

_____ 4. provides call control services to registered H.323 endpoints

Q3) What are three benefits that H.323 offers? (Choose three.)

A) uses a proxy server to reduce call set up complexity

B) offers flexible communications with terminals of different capabilities

C) relies on a centralized call agent for signaling and control because endpoints are relatively unintelligent

D) uses codec standards for audio and video transmission

E) sets interoperability through H.323 compliance

Q4) Which component is the software-based, call-processing component of the Cisco enterprise IP telephony solution?

A) infrastructure

B) Cisco CallManager

C) call processing

D) applications

Q5) A voice-enabled router decides where to establish a connection based on which two parameters (components)? (Choose two.)

A) dial peer

B) call leg

C) destination number

D) calling number

E) voice port

Q6) Match each voice communication function with its description.

A) H.225 call signaling channel

B) H.245 control channel

C) RTP

_____ 1. carries packet sequence and time stamping information

_____ 2. carries messages governing voice operation, including capabilities exchange, opening and closing of logical channels, and preference requests

_____ 3. carries messages for call setup

Q7) Select two benefits of using hub-and-spoke topology versus a fully meshed topology in Voice over FR (VoFR). (Choose two.)

A) minimized delay

B) high voice quality

C) reduced network costs

D) reduced number of maximum possible hops

E) reduced number of PVCs

Q8) What is the preferred ATM class of service (COS) for Voice over ATM (VoATM) transport?

A) constant bit rate (CBR)

B) variable bit rate (VBR)

C) unspecified bit rate (UBR)

D) available bit rate (ABR)

Quiz Answer Key

Q1) B, E

Relates to: Introducing Voice over IP

Q2) 1=B, 2=D, 3=A, 4=C

Relates to: H.323 Components

Q3) B, D, E

Relates to: H.323 Components

Q4) B

Relates to: Components of IP Telephony

Q5) A, C

Relates to: Voice Routing with Dial Peers

Q6) 1=C, 2=B, 3=A

Relates to: VoIP Control and Transport Protocols

Q7) C, E

Relates to: Voice over Frame Relay

Q8) B

Relates to: Voice over ATM

Identifying the Requirements of Voice Technologies

Overview

To create a proper integrated network design, you need to know about the many considerations that affect voice traffic. This lesson focuses on the types of delay, bandwidth restrictions, and quality of service (QoS) issues that affect voice communications. It discusses the importance of coding, which affects bandwidth consumption and voice quality, and introduces mechanisms that impact voice quality.

Relevance

To design a network to transport voice, you need to understand and consider delay, packet loss, jitter, and the impact QoS can have on maintaining voice quality.

Objectives

Upon completing this lesson, you will be able to describe possible issues in packet telephony and identify the solutions. This includes being able to meet these objectives:

- Identify where network delays, jitter, and packet loss are likely to occur and describe how to minimize them
- Explain the problem of echo and the solution to the problem
- Estimate the possible number of simultaneous calls using a certain coding type on a WAN link
- Specify the QoS mechanisms to implement to effective service for voice traffic

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with modern telephony concepts and terms in addition to basic QoS mechanisms

Outline

The outline lists the topics included in this lesson.

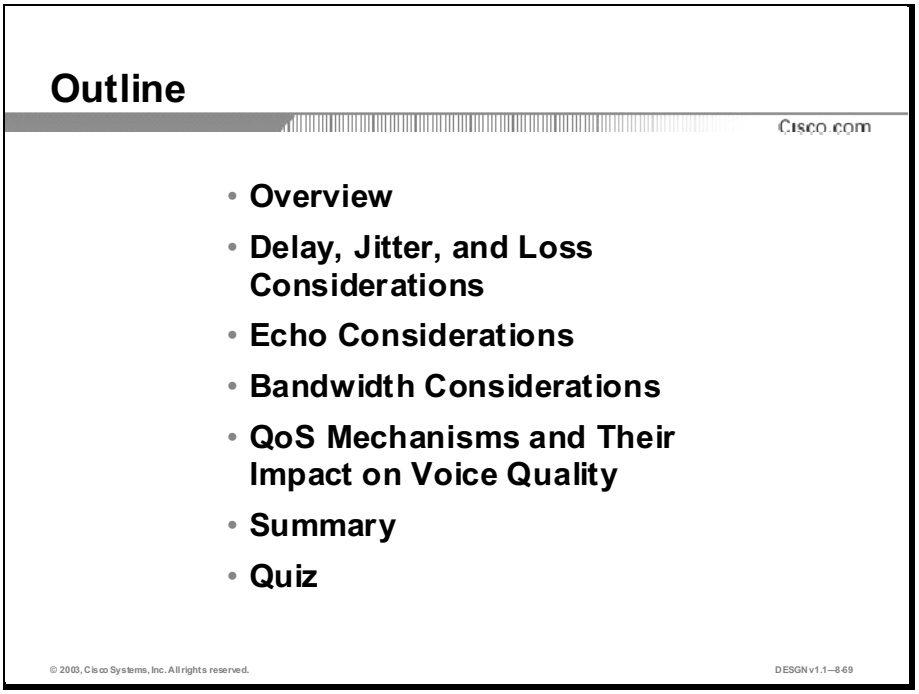

Delay, Jitter, and Loss Considerations

Voice quality in an IP network is directly affected by delay, jitter, and packet loss. Solutions are available to address all three issues. This topic identifies where network delays, jitter, and packet loss are likely to occur and describes how to minimize them.

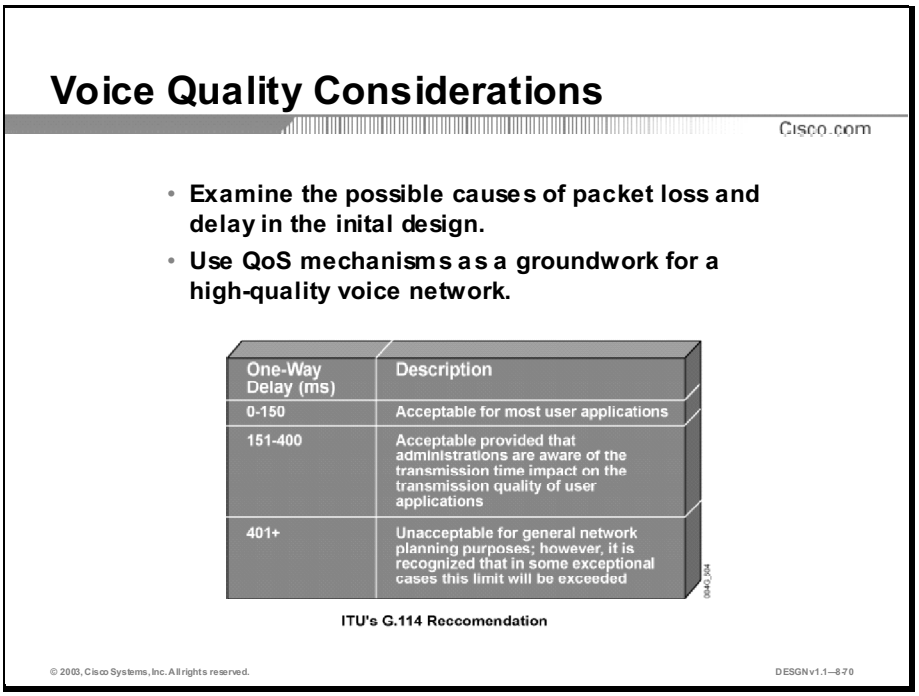

When designing networks that transport voice over packet, frame, or cell infrastructures, you must consider the delay components in the network. Correctly accounting for all potential delays ensures that overall network performance will be acceptable.

The two major types of delay are fixed delay and variable delay.

One-Way Network Delay Recommendations

The generally accepted limit for good-quality voice connection delay is 150 ms one-way. As delays rise, the communication between two people suffers. (For example, they speak at the same time, or both wait for the other to speak.) This condition is called talker overlap.

The International Telecommunication Union (ITU) describes network delay for voice applications in recommendation G.114. This recommendation, shown in the figure, defines three bands of one-way delay.

Voice packets will be delayed if the network is congested because of poor network design, traffic congestion, or insufficient bandwidth.

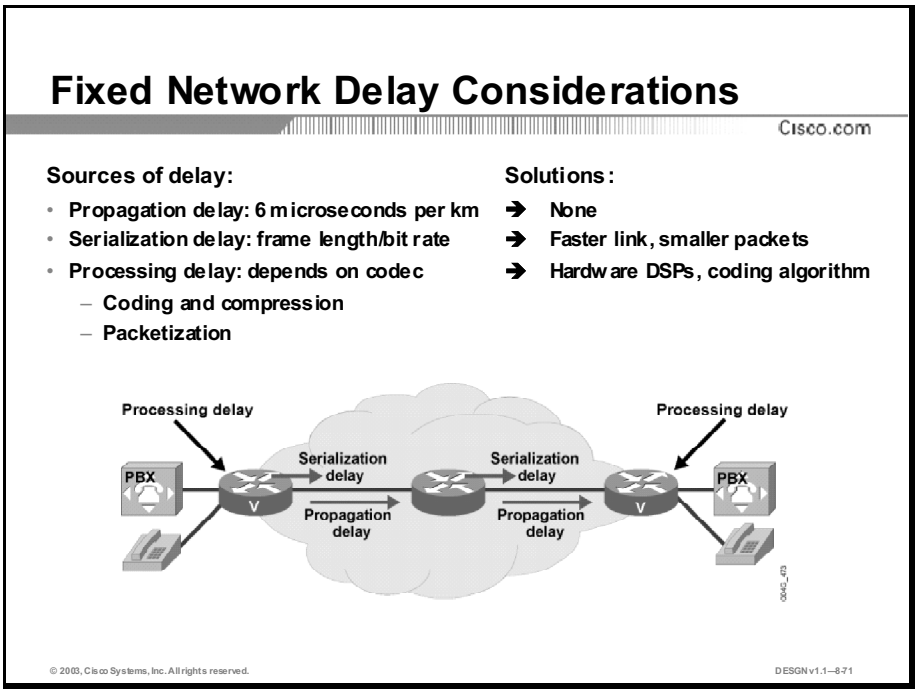

Fixed network delay includes three components: propagation delay, serialization delay, and processing delay.

Propagation Delay

Propagation delay is the delay of signals between the sending and receiving endpoints. You can ignore propagation delay, which is limited by the speed of light, for most designs because it is relatively small compared to other types of delay.

Note:	Propagation delay has a noticeable impact to the overall delay only on satellite links.

Serialization Delay

Serialization delay is the result of placing bits on the circuit. The higher the circuit speed, the less time it takes to place the bits on the circuit; the higher the speed, the lower the serialization delay.

Serialization delay is a constant function of link speed and packet size. Calculate serialization delay by dividing the packet length by the bit rate. This formula shows how large serialization delay can occur as a result of slow links or large packets. Serialization delay is always predictable; for example, using a 64-kbps link and an 80-byte frame, the delay is exactly 10 ms.

Note:	Serialization delay is a factor only on slow speed links up to 1 Mbps.

Processing Delay

Processing delays include:

- **Coding, compression, decompression, and decoding delays:** These functions are performed in either hardware or software, based on the algorithm used. By using specialized hardware, such as digital signal processors (DSPs), you can dramatically improve the quality and reduce the delay associated with different voice compression schemes.
- **Packetization delay:** The process of holding the digital voice samples to place into the payload until enough samples are collected to fill the packet or cell payload. To reduce excessive packetization delay associated with some compression schemes, the voice gateway can send partial packets.

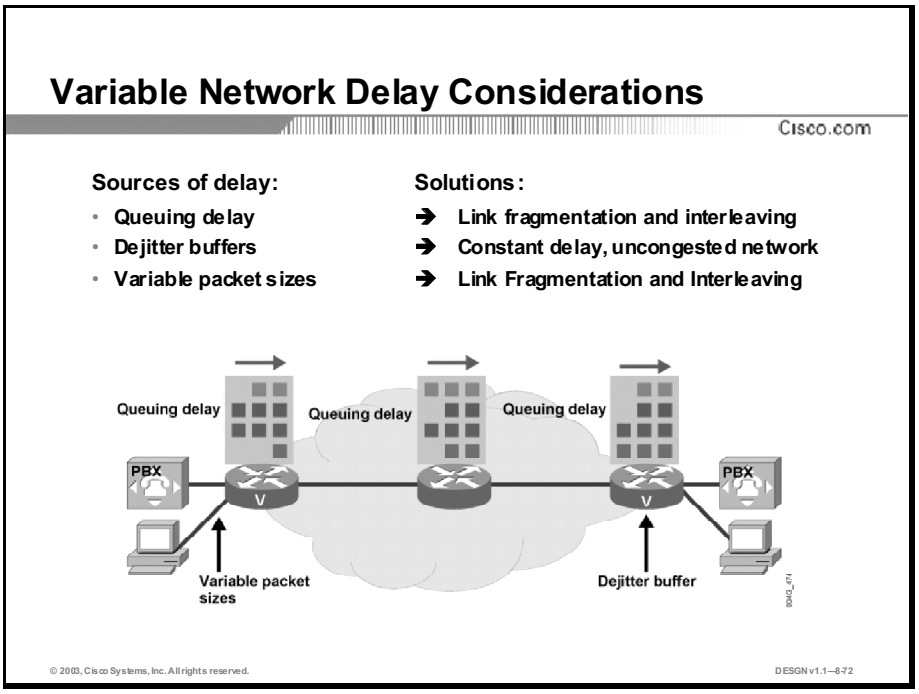

Variable network delay is more unpredictable and difficult to calculate than fixed network delay. Three factors that contribute to a variable network delay are queuing delay, dejitter buffers, and variable packet sizes.

Queuing Delay

Congested output queues on network interfaces are the most common sources of variable delay. Queuing delay occurs when a voice packet waits for others to be serviced first on the outgoing interface. This waiting time is statistically based on the arrival of traffic; the more inputs, the more likely that packets will contend for the trunk. Queueing delay is also based on the size of the packet currently being serviced.

For example, a 1500-byte data packet is queued before the voice packet. The voice packet must wait until the whole data packet is transmitted, which produces a delay in the voice path. If the link is slow (for example, 64 kbps or 128 kbps), the delay may be more than 200 ms, which is unacceptable for voice.

Link fragmentation and interleaving (LFI) is a solution for queuing delay situations. With LFI, the voice gateway fragments large packets into smaller frames and interleaves them with small voice packets. Thus, a voice packet does not experience a delay until the entire data packet is sent. LFI reduces delay and ensures a more predictable voice delay.

Dejitter Buffers

Because network congestion can occur at any point within a network, packets can fill interface queues instantaneously. This situation leads to a difference in delay times between packets in a single voice stream. This variable delay is called jitter.

Dejitter buffers work at the receiving end to smooth delay variability and allow time for decoding and decompression. On the first talk spurt, the dejitter buffer provides smooth voice playback. Setting the dejitter buffer too low causes overflows and loss of data, while setting it too high causes excessive delay.

In effect, a dejitter buffer reduces or eliminates delay variation by converting it to fixed delay. However, a dejitter buffer always adds delay to the total budget, depending on the variance of the delay.

Dejitter buffers work most efficiently when packets arrive in a fairly uniform delay. You can use QoS congestion avoidance mechanisms to manage delay and avoid network congestion. If there is no delay variance, you can disable dejitter buffers and reduce the constant delay.

Note: Delay is always added to the total delay budget when you use dejitter buffers. Keep the dejitter buffer as small as possible to reduce total delay to a minimum.

Variable Packet Sizes

Variable delay is influenced by the size of the packet currently being serviced; larger packets take longer to transmit than smaller packets. Therefore, a queue that combines large and small packets experiences varying lengths of delay.

LFI fragments packets into frames of equal size to solve the delay problem caused by variable packet size. Configure LFI fragmentation on a link to provide a fixed 10-ms delay. Set the fragment size so that voice packets do not become fragmented.

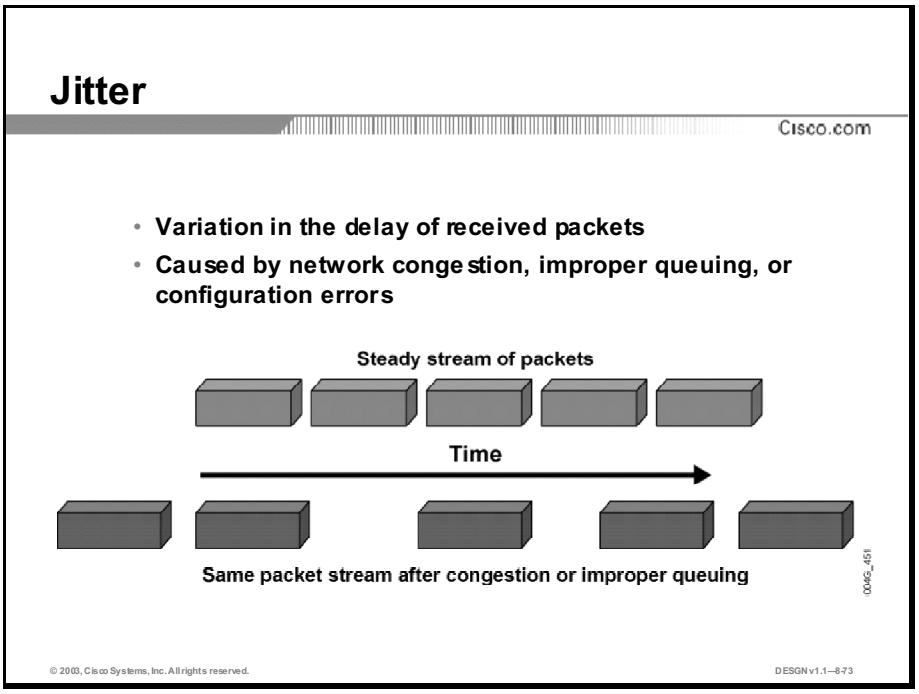

Jitter is defined as a variation in the delay of received packets. The originating voice gateway sends the packets in a continuous stream, spaced evenly apart. Because of network congestion, improper queuing, or configuration errors, the delay between each packet can vary instead of remaining constant.

When a voice gateway receives an audio stream for voice over IP (VoIP), it must compensate for the jitter that it encounters with the playout delay buffer. The playout delay buffer must buffer the packets and then play them out in a steady stream to the DSPs, which then convert the voice back to an analog audio stream. The playout delay buffer is also referred to as the dejitter buffer.

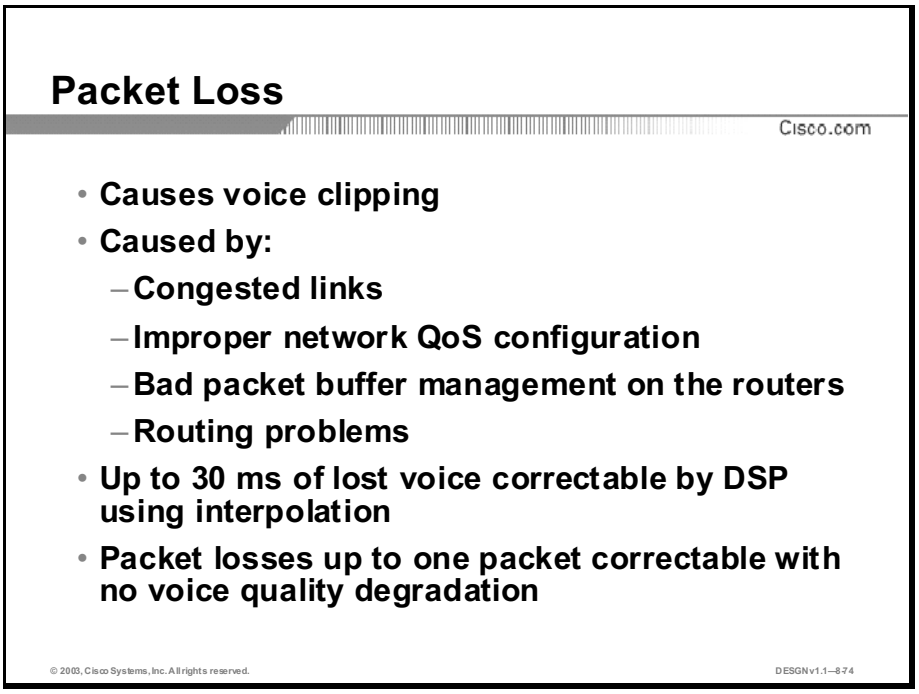

Packet loss causes voice clipping and skips. Loss may occur because of congested links, improper network QoS configuration, poor packet buffer management on the routers and switches, routing problems, and so on.

The industry standard codec algorithms used in the Cisco DSP can correct for up to 30 ms of lost voice. Cisco VoIP technology uses 20-ms samples of voice payload per VoIP packet. For the codec correction algorithms to be effective, only a single packet can be lost during any given time. For packet losses as small as one packet, the DSP interpolates the conversation with what it thinks the audio should be, and the packet loss is not audible.

Note: Losses also occur if packets are received out of range of the dejitter buffer, and the out-of-range packets are discarded.

Echo Considerations

In a voice telephone call, an echo occurs when callers hear their own words repeated. An echo is the audible leak of the caller's voice into the receive (return) path. This topic explains the problem of echo and the solution to the problem.

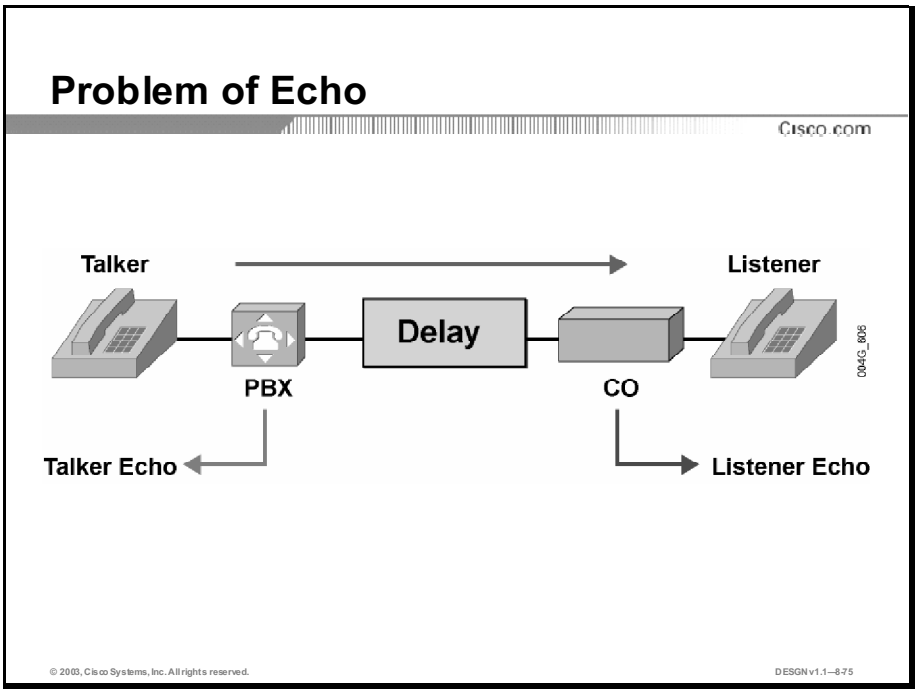

Echo is a function of delay and magnitude. The echo problem worsens as the delay and the loudness of the echo grow. When timed properly, echo is reassuring to the speaker. If the echo exceeds approximately 25 ms, it can be distracting and cause breaks in the conversation.

These elements in a VoIP network can affect echo:

- **Hybrid transformers:** Hybrid transformers are often prime culprits for signal leakage between analog transmit and receive paths. Echo is normally caused by a mismatch in impedance from the four-wire network switch conversion to the two-wire local loop. This type of echo is typically the result of an impedance mismatch in a PBX.
- **Telephones:**
  - The analog telephone itself presents a load to the PBX. This load should match the output impedance of the source device (FXS port). Some inexpensive telephones are not matched to the output impedance of the FXS port and are sources of echo. Headsets are particularly notorious for poor echo performance.
  - When digital telephones are used, the point of digital-to-analog conversion occurs inside the telephone. Extending the digital transmission segments closer to the actual telephone decreases the potential for echo.

Note: The belief that adding voice gateways to a voice network creates echo is a common misconception. Digital segments of the network do not cause leaks; technically, voice gateways cannot be the source of echo.

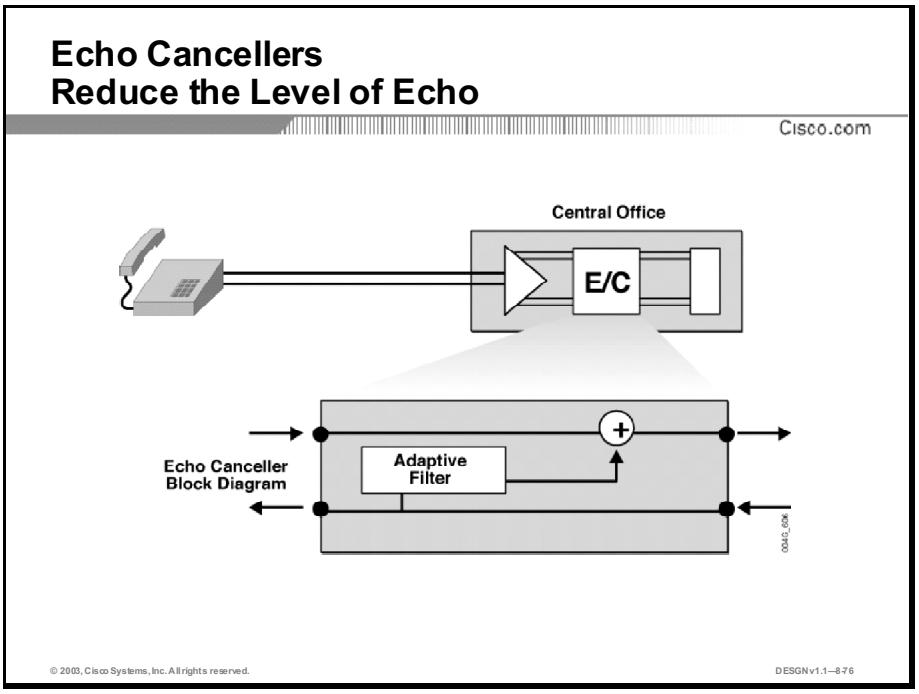

To improve the quality of telephone conversation, you may need to place an echo canceller in the network. An echo canceller in a voice gateway reduces the level of echo that has leaked from the receive path into the transmit path. Echo cancellers are built into low bit-rate codecs and are operated on each DSP. Echo cancellers are limited, by design, by the total amount of time that they wait for the reflected speech to be received, which is known as an echo trail. The echo cancellation time (echo trail) is normally between 16 and 32 ms.

Example: How Echo Cancellers Work

Assume that user A is talking to user B. When the speech of user A hits an impedance mismatch or other echo-causing environment, it bounces back to user A. User A can hear the delay several milliseconds after speaking.

To remove the echo from the line, user A's router needs to keep an inverse image of the speech of user A for a certain amount of time. This image is called inverse speech. The echo canceller in the router listens for sound coming from user B and subtracts the inverse speech of user A to remove any echo.

Echo Cancellation Guidelines

The International Telecommunication Union Telecommunication Standardization Sector (ITU-T) defines an "irritation zone" of echo loudness and delay. You do not need to suppress a short echo of around 15 ms, but higher echo delays require strong echo suppression. Therefore, echo cancellation is required in all networks that produce one-way time delays greater than 16 ms.

You also need to configure the appropriate echo cancellation time. If the echo cancellation time is set too low, callers still hear echo during the phone call. If the configured echo cancellation timer is set too high, it takes longer for the echo canceller to converge and eliminate the echo.

Attenuating the signal below the noise floor can also eliminate echo.

Bandwidth Considerations

A primary issue when network designers are designing voice on IP network is bandwidth availability. This topic shows you how to estimate the number of simultaneous calls using a certain coding type on a WAN link.

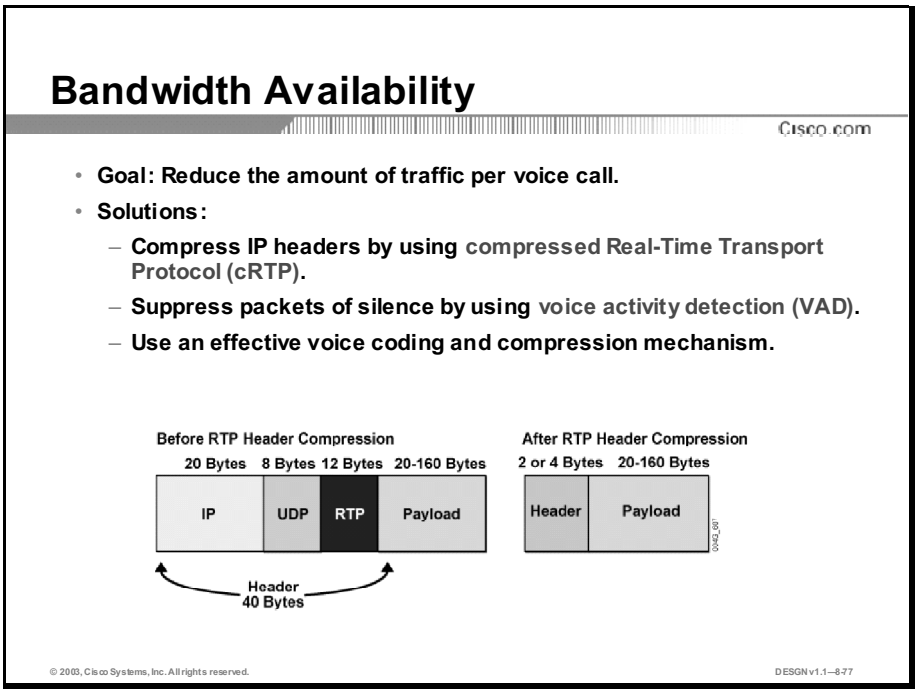

The amount of bandwidth per call increases or decreases greatly, depending on which codec is used and the number of voice samples required per packet. However, the best coding mechanism does not necessarily result in the best voice quality. For example, the better the compression, the worse the voice quality. You must decide which is more important, better voice quality or more efficient bandwidth consumption.

To reduce the amount of traffic per voice call and thus use available bandwidth more efficiently, you can use an effective voice coding and compression mechanism, use compressed Real-Time Transport Protocol (cRTP), or suppress silence using voice activity detection (VAD).

Reducing Voice Traffic with cRTP

All voice packets encapsulated into IP consist of two components: voice samples and IP/UDP/RTP headers. Although the DSP compresses voice samples, which vary in size based on the codec used, the headers are a constant 40 bytes. When compared to the 20 bytes of voice samples in a default G.729 call, the headers make up a considerable amount of overhead. cRTP compresses the headers to 2 or 4 bytes, which offers significant bandwidth savings. cRTP is sometimes referred to as RTP header compression.

Enabling compression on a low-bandwidth serial link can greatly reduce the network overhead and conserve WAN bandwidth if there is a significant volume of RTP traffic. In general, enable cRTP on slow links up to 2 Mbps. However, cRTP is not recommended for higher-speed links because of its high CPU requirements.

Note: Because cRTP compresses VoIP calls on a link-by-link basis, configure cRTP on all links on the path.

Reducing Voice Traffic with VAD

On average, about 35 percent of all calls produce silence. In traditional voice networks, all voice calls use a fixed bandwidth of 64-kbps links, regardless of how much of the conversation is speech and how much is silence. With VoIP networks, all conversation and silence is packetized. VAD suppresses packets of silence. Instead of sending VoIP packets of silence, VoIP gateways interleave data traffic with VoIP conversations to use network bandwidth more effectively.

Voice Coding and Compression

Cisco.com

- The quality of transmitted speech is a subjective listener response.
- MOS is a common benchmark to define sound quality.
- MOS scales from 1 (bad) to 5 (excellent).

	ITU Standard	Data Rate	MOS Score
PCM	G.711	64 kbps	4.1
ADPCM	G.726/G.727	16/24/32/40 kbps	3.85 or less
LDCELP	G.728	16 kbps	3.61
CS-ACELP	G.729	8 kbps	3.92
ACELP/MP-MLQ	G.723.1	6.3/5.3 kbps	3.9/3.65

Advances in technology have greatly improved the quality of compressed voice, resulting in a spectrum of coding and compression algorithms.

- **PCM:** The toll-quality voice expected from the PSTN. PCM runs at 64 kbps, provides no compression, and therefore provides no opportunity for bandwidth savings.
- **Adaptive differential pulse code modulation (ADPCM):** Provides three different levels of compression. The quality change is virtually imperceptible when compared to 64-kbps PCM. Some fidelity is lost as compression increases. Depending on the traffic mix, cost savings generally run at 25 percent for 32-kbps ADPCM, 30 percent for 24-kbps ADPCM, and 35 percent for 16-kbps ADPCM.
- **Low-delay code excited linear prediction (LDCELP) compression:** This algorithm models the human voice. Depending on the traffic mix, cost savings may be up to 35 percent for 16-kbps LDCELP.
- **Conjugate Structure Algebraic Code Excited Linear Prediction (CS-ACELP):** Provides eight times the bandwidth savings over PCM. CS-ACELP is a more recently developed algorithm, which is modeled after the human voice and delivers quality comparable to LDCELP and 32-kbps ADPCM. Cost savings are approximately 40 percent for 8-kbps CS-ACELP.
- **Code excited linear prediction compression (CELP):** Provides huge bandwidth savings over PCM. Cost savings may be up to 50 percent for 5.3-kbps CELP.

Voice Coding Standards (Codecs)

The ITU defines a series of standards for voice coding and compression:

- **G.711:** Uses the 64-kbps PCM voice coding technique. G.711-encoded voice is already in the correct format for digital voice delivery in the public telephone network or through PBXs.
- **G.726:** Uses ADPCM coding at 40, 32, 24, and 16 kbps. ADPCM voice may also be interchanged between packet voice networks and public telephone or PBX networks, providing the latter have ADPCM capability.
- **G.728:** Uses LDCELP voice compression, requiring only 16 kbps of bandwidth. CELP voice coding must be transcoded to a PCM-based coding before delivery to the PSTN for delivery to or through the telephone networks.
- **G.729:** Uses CS-ACELP compression, which enables voice to be coded into 8-kbps streams. There are three forms of this standard, and all provide speech quality similar to that of 32-kbps ADPCM.
- **G.723.1:** Uses a dual rate coder for compressing speech at very low bit rates. This standard has two bit rates associated with it: 5.3 kbps using ACELP and 6.3 kbps using Multipulse Multilevel Quantization (MP-MLQ).

Each codec provides a certain quality of speech. The quality of transmitted speech is a subjective listener response. A common measure used to determine the sound quality that specific codecs produce is the mean opinion score (MOS).

With MOS, a wide range of listeners judge the quality of a voice sample (produced by a particular codec) on a scale of 1 (bad) to 5 (excellent). The scores are averaged to provide the MOS for that sample. The figure shows the relationship between codecs and MOS scores. It is evident that MOS decreases with the increased compression. The data rates in the figure do not include any packetization overheads.

A newer, more objective measurement is available that is quickly overtaking MOS scores as the industry quality measurement of choice for coding algorithms. Perceptual Speech Quality Measurement (PSQM), specified in ITU standard P.861, provides a rating on a scale of 0 to 6.5, where 0 is best and 6.5 is worst. PSQM works by comparing the transmitted speech to the original input to produce a score.

PSQM is implemented in test equipment and monitoring systems that provide a PSQM score for a test voice call over a particular packet network. Some PSQM test equipment converts the 0-to-6.5 scale to a 0-to-5 scale to correlate to MOS.

Codec Complexity, DSPs, and Voice Calls

Codec is a technology for compressing and decompressing data implemented in DSPs. Some codec compression techniques require more processing power than others. Codec complexity is divided into medium and high complexity. The difference between a medium- and a high-complexity codec is the amount of processor power to process the codec algorithm and the number of voice channels that a single DSP can support. You can run medium-complexity codecs in high-complexity mode, but fewer channels (usually half of them) are available per DSP.

Assume that this is the relationship between coding mechanism, a particular DSP, and voice channels:

- Medium-complexity codecs allow the DSPs to process up to four voice calls per DSP.
- High-complexity codecs allow the DSPs to process up to two voice calls per DSP.

Codecs are divided into two complexity groups:

- Medium-complexity codecs are G.711, G.726, G.729a, and G.729ab.
- High-complexity codecs are G.728, G.723, G.729, and G.729b.

The relationship between codec algorithms and the number of voice channels per DSP is:

- G.711, G.726, G.729a, and G.729ab codecs allow four calls per DSP.
- G.728, G.723, G.729, and G.729b codecs allow two calls per DSP.

Example: DSP Resource Calculation

An enterprise implements a VoIP network between two locations that will use two voice gateways with 30 DSP resources. Because of the resulting high bandwidth savings, the company used the G.729a codec.

Because the G.729a is a medium-complexity codec, which allows up to four voice calls per DSP, the system can establish up to 120 voice channels between the two locations.

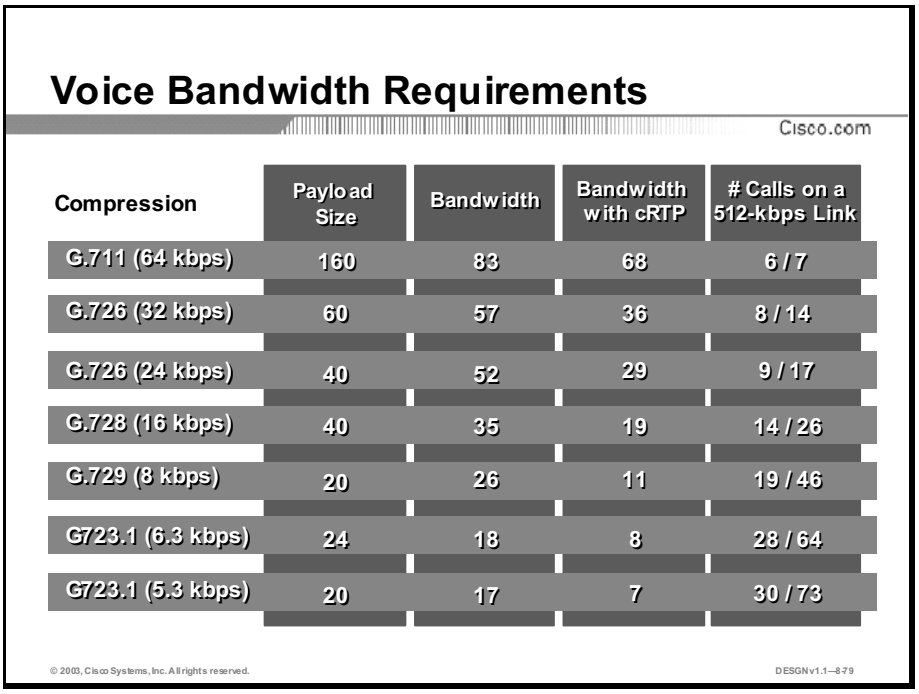

It might seem logical from a bandwidth consumption standpoint to convert all calls to low bit-rate codecs, saving bandwidth to decrease infrastructure costs. You must consider both expected voice quality and bandwidth consumption to choose the optimum codec. You should also consider the disadvantages of strong voice compression. One of the main disadvantages is signal distortion after multiple encodings. For example, when a G.729 voice signal is encoded three times, the MOS score drops from 3.92 (very good) to 2.68 (unacceptable). Another drawback is codec-induced delay with low bit-rate codecs.

Note: G.729 is the recommended voice codec for most WAN networks that do not support multiple encodings because of its relatively low bandwidth requirements and high MOS.

Make these assumptions in your bandwidth calculation:

- IP/UDP/RTP headers are 40 bytes.
- RTP header compression can reduce the IP/UDP/RTP headers to 2 or 4 bytes.
- A Layer 2 header adds 6 bytes.

The figure presents the codecs, their payload size, and the required bandwidth with and without RTP header compression. The last column indicates the number of normal compressed calls on a 512-kbps link.

These calculations are used to create the figure:

- Voice packet size = (Layer 2 header) + (IP/UDP/RTP header) + (voice payload)
- Voice packets per second (pps) = codec bit rate/voice payload size
- Bandwidth = voice packet size * pps

Example: Bandwidth Calculation

The bandwidth calculation for a G.729 call (8-kbps codec bit rate) with cRTP and default 20 bytes of voice payload is:

- Voice packet size (bytes) = (Layer 2 header of 6 bytes) + (compressed IP/UDP/RTP header of 2 bytes) + (voice payload of 20 bytes) = 28 bytes
- Voice packet size (bits) = (28 bytes) * 8 bits per byte = 224 bits
- Voice packets per second (pps) = (8-kbps codec bit rate) / (160 bits) = 50 pps
- Bandwidth per call = voice packet size (224 bits) * 50 pps **= 11.2 kbps**
- **Result:** The G.729 call with cRTP requires 11.2 kbps of bandwidth

QoS Mechanisms and Their Impact on Voice Quality

QoS mechanisms are important for networks that carry delay-sensitive traffic, such as voice. At the same time, you must address the needs of less time-dependent applications, such as file transfer. This topic specifies the QoS mechanisms to implement to provide effective service for voice traffic.

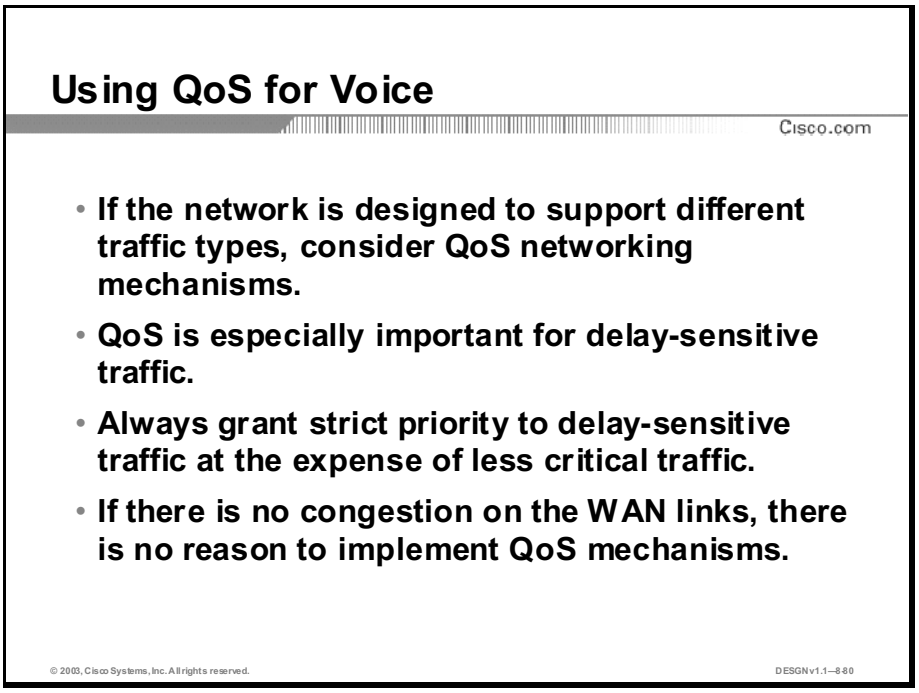

If a network supports a variety of traffic types on a single data path between routers, you must consider different QoS techniques to ensure that they treat each data type fairly.

Use these guidelines when determining whether or not you require QoS:

- Traffic prioritization is important for delay-sensitive, interactive, transaction-based applications.
- Prioritization is most effective on WAN links, where the combination of bursty traffic and relatively lower data rates can cause temporary congestion.
- Depending on the average packet size, prioritization is most effective when applied to links at T1/E1 bandwidth speeds or lower.
- If users of network applications notice poor response time, consider congestion management features.
- If there is no congestion on the WAN link, you do not need to implement traffic prioritization.

Designing Voice QoS

These steps summarize the aspects to consider when determining whether to implement QoS in the network:

Step 1 Determine if the WAN is congested, that is, whether users of an application perceive performance degradation.

Step 2 Determine the network goals and objectives based on the mix of traffic in the network. Consider which of these goals are important:

- To establish fair distribution of bandwidth allocation across all traffic types
- To grant strict priority to voice traffic at the expense of less critical traffic
- To customize bandwidth allocation so that network resources are shared among all applications, each having specific bandwidth requirements
- To analyze the types of traffic and determine how to distinguish them
- To review the available QoS mechanisms and determine which approach best addresses the requirements and goals

Step 3 Configure the voice gateways for the QoS strategy chosen, and observe the results.

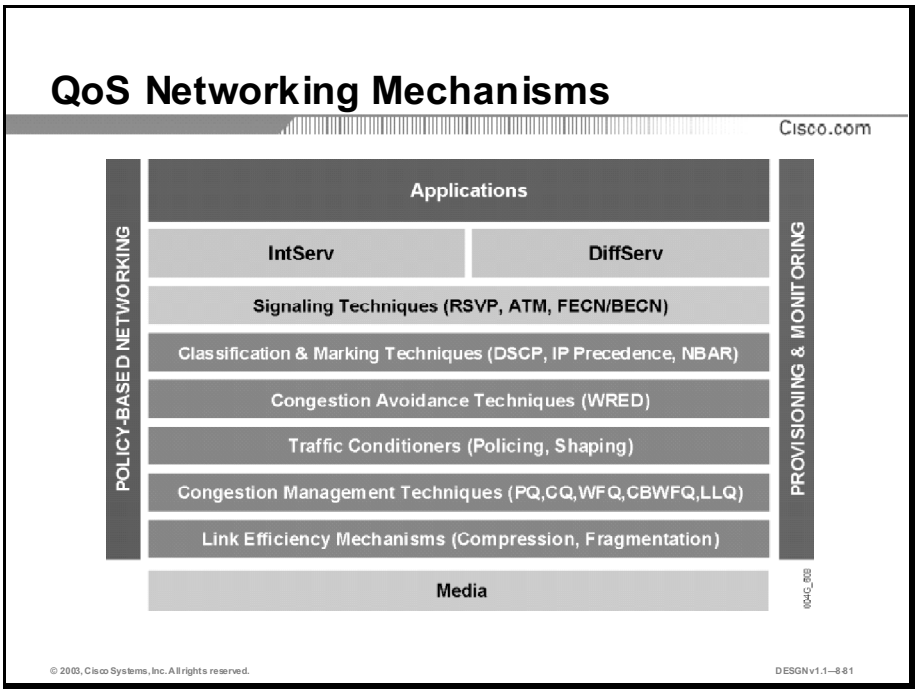

If voice traffic is sharing network resources with data traffic, consider IP QoS mechanisms to provide an effective level of service for the voice traffic. These are the categories of QoS mechanisms:

- **Classification and marking:** Classification is the process of identifying the class or group to which a packet belongs. Network devices use various match criteria to place traffic into a certain number of classes. Classification is accomplished with class maps, access control lists, or route maps and sets the IP precedence bits to the matched packets. Matches are based on the following criteria:
  - **dial-peer voice voip** global configuration command
  - Protocol such as stateful protocols, or a Layer 4 protocol
  - Input port
  - IP precedence or differentiated services code point (DSCP)
  - Ethernet 802.1p class of service (CoS)
- **Congestion avoidance:** Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network and internetwork bottlenecks before they become a problem. Congestion avoidance techniques provide preferential treatment for premium (priority) class traffic under congestion situations while concurrently maximizing network throughput and capacity utilization, and minimizing packet loss and delay. Weighted random early detection (WRED) and distributed random early detection (DWRED) are the Cisco IOS QoS congestion avoidance features.

- **Traffic conditioners:** Policers and shapers usually identify traffic descriptor violations in an identical manner. They usually differ, however, in the way they respond to violations, for example:
  - A policer typically drops traffic. For example, the CAR rate-limiting policer will either drop the packet or rewrite its IP precedence, resetting the type of service bits in the packet header.
  - A shaper typically delays excess traffic using a buffer, or queueing mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected. For example, Generic Traffic Shaping uses a weighted fair queue to delay packets in order to shape the flow, and Frame Relay traffic shaping uses either a priority queue, a custom queue, or a first in, first out (FIFO) queue for the same, depending on how you configure it.
- **Congestion management:** To manage congestion, network devices use a queuing algorithm to segregate traffic and determine a method to prioritize it on an output link. Configure QoS to provide sufficient bandwidth and priority forwarding for delay-sensitive traffic. Examples of congestion management techniques are FIFO, weighted fair queuing (**WFQ)**, priority queuing (PQ), custom queuing (CQ), class-based weighted fair queuing CBWFQ), and low latency queueing (LLQ).
- **Link efficiency:** Link efficiency mechanisms reduce delay on slower-speed links by breaking up large datagrams and interleaving low-delay traffic packets with the smaller packets resulting from the fragmented datagram.

You will select classification, congestion management, and congestion avoidance mechanisms to meet network service requirements over constrained network resources.

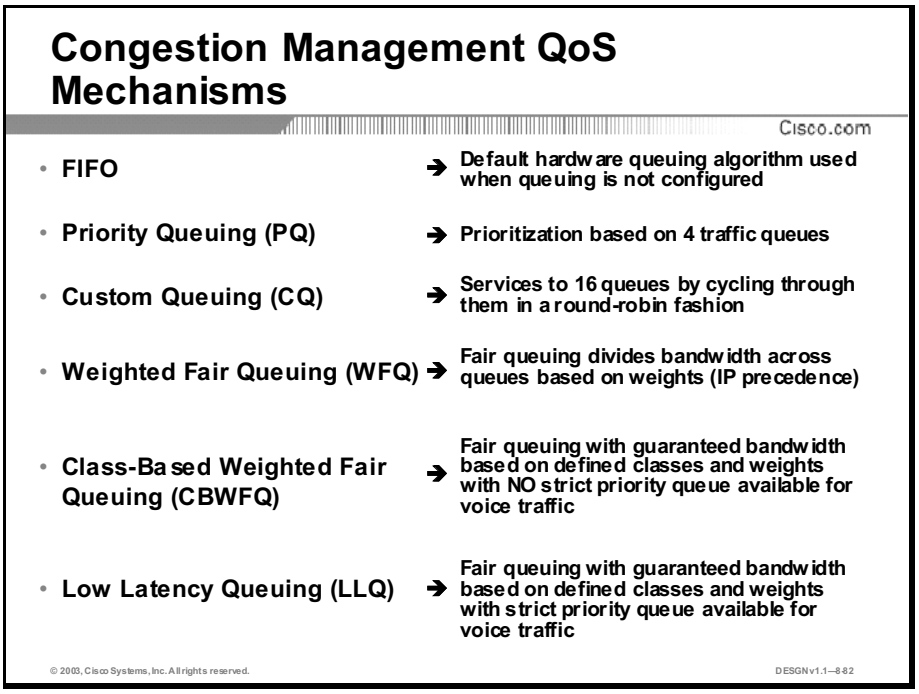

Congestion management mechanisms that address the requirements of end-to-end QoS and service differentiation for voice packet delivery are:

- **FIFO:** Provides hardware-based queueing. If you do not configure other queueing, FIFO is used on links with bandwidth greater than 2.048 Mbps.
- **PQ:** Defines how voice traffic is prioritized in the network. You can configure four traffic priorities. Define a series of filters based on packet characteristics, such as source IP address and port, to place voice traffic in the highest queue and other traffic in the lower three queues. The voice gateway services the queue with the highest priority and then the lower queues in sequence. This queuing method will continue to service the strict priority queue as long as there are packets in queue. This can lead to lower priority queues being starved.
- **CQ:** Allocates bandwidth proportionally for each class of traffic. Specify the number of bytes or packets drawn from the queue. CQ services queues by cycling through them in a round-robin fashion, sending the portion of allocated bandwidth for each queue before moving on to the next queue. If one queue is empty, the router sends packets from the next queue that has packets ready to send.
- **WFQ:** Offers dynamic fair queuing that divides bandwidth across queues of traffic based on weights. WFQ recognizes IP precedence. It classifies traffic into conversations and determines the amount of bandwidth that each conversation is allowed relative to other conversations. WFQ is a flow-based algorithm that simultaneously schedules voice traffic to the front of a queue to reduce response time and fairly shares the remaining bandwidth among high-bandwidth flows. WFQ is used on links with bandwidth less than 2.048 Mbps.
- **CBWFQ:** Provides WFQ based on defined classes with no strict priority queue available for real-time traffic. The voice gateway services all packets fairly based on weight, not based on strict priority. This scheme poses problems for voice traffic that is largely intolerant of delay, especially variation in delay. For voice traffic, variations in delay introduce transmission irregularities, which manifest as jitter in the heard conversation.

- **LLQ:** Brings strict priority queuing to CBWFQ. Strict priority queuing dequeues and sends delay-sensitive data such as voice to be dequeued and sent first (before packets in other queues are dequeued), which gives delay-sensitive data preferential treatment over other traffic. LLQ is a preferred queuing mechanism for designing voice on IP networks.

Queuing Guidelines for Voice

Use the latest version of IOS software to get the most advanced queuing features. If possible, use LLQ and classify voice in a priority queue. Set the bandwidth of the voice class to the aggregate voice bandwidth on the link, allowing for a little overhead.

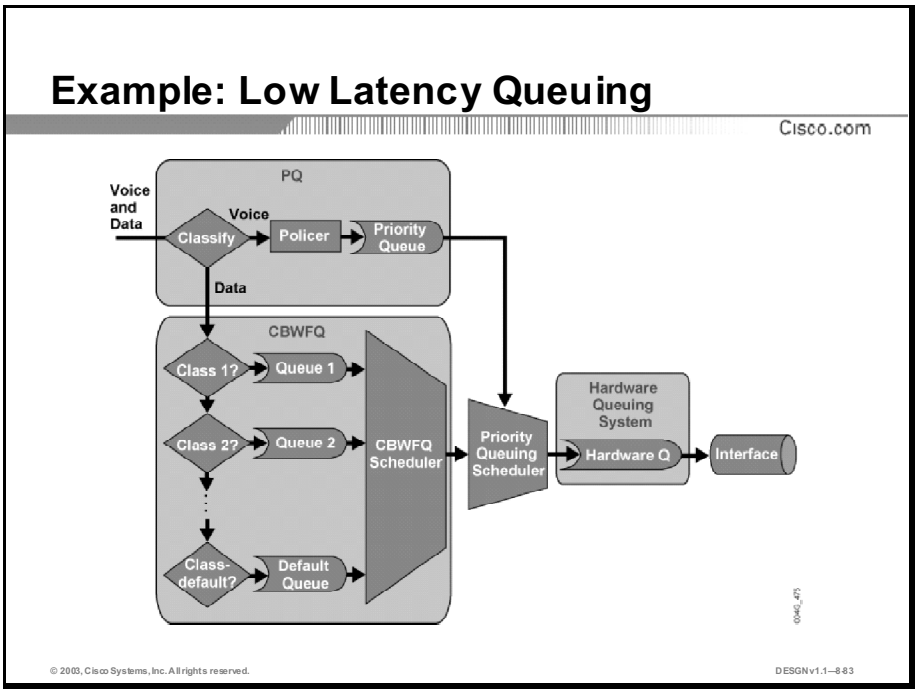

The figure shows how LLQ combines CBWFQ and PQ. Strict priority queuing sends delay-sensitive data first.

The advantage of LLQ is that its policing mechanism guarantees bandwidth for voice and gives it a priority. LLQ reduces jitter in voice conversations. The rest of the traffic is classified using CBWFQ.

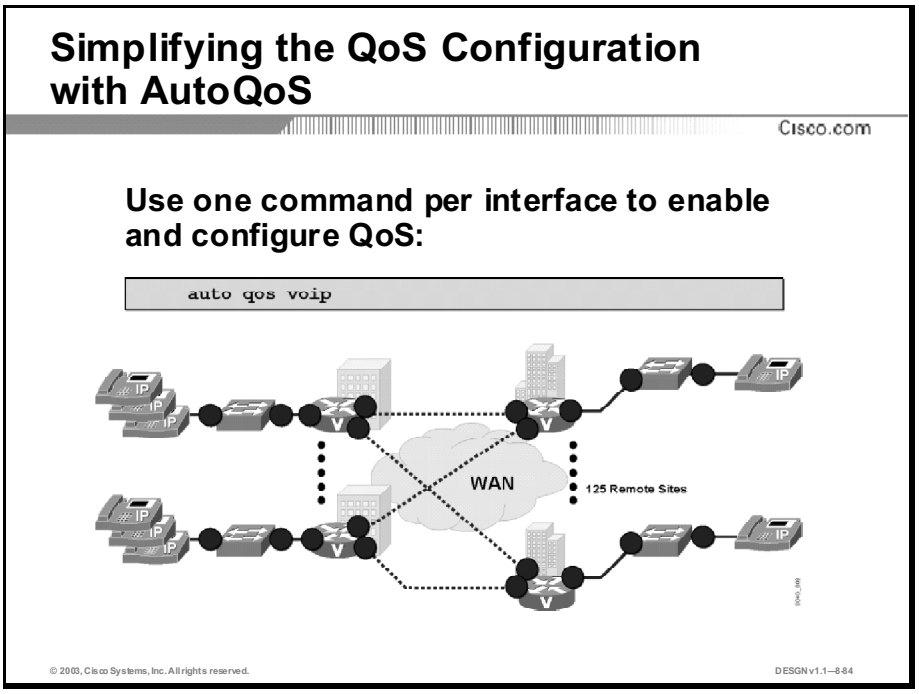

Cisco AutoQoS helps you deploy QoS on a converged voice and data network quickly and efficiently. It automatically configures the following, freeing a network administrator from configuring each device separately:

- Traffic classes
- Traffic policies

Therefore, when a network administrator configures AutoQoS at the interface or PVC, the traffic receives the required QoS treatment automatically. In-depth knowledge of the underlying technologies, service policies, link efficiency mechanisms, and Cisco QoS best practice recommendations for voice requirements is not required to configure AutoQoS.

Cisco AutoQoS is beneficial for these scenarios:

- Small-to-medium size businesses that need to deploy IP telephony quickly, but lack the experience and staffing to plan and deploy IP QoS services
- Large enterprises that need to deploy IP telephony on a large scale, while reducing the costs, complexity, and timeframe for deployment and ensuring that the appropriate QoS for voice applications is set in a consistent fashion
- International enterprises or service providers requiring QoS for VoIP where little expertise exists in different regions of the world and where provisioning QoS remotely and across different time zones is difficult
- Service providers requiring a template-driven approach to delivering managed services and QoS for voice traffic to large numbers of customer premise devices

Summary

This topic summarizes the key points discussed in this lesson.

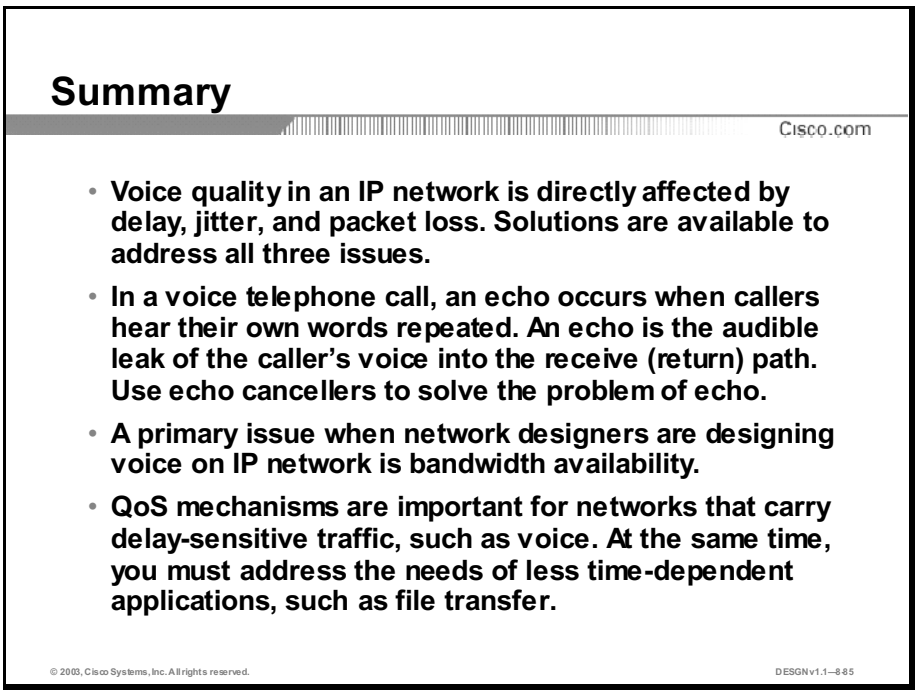

References

For additional information, refer to these resources:

- Davidson, J., and J. Peters. *Voice over IP Fundamentals*. Indianapolis, Indiana: Cisco Press; 2000.
- *Understanding Delay in Packet Voice Networks,* http://www.cisco.com/warp/public/788/voip/delay-details.html
- *VoIP—Understanding Codecs: Complexity, Support, MOS, and Negotiation,* http://www.cisco.com/warp/public/788/voip/codec_complexity.html
- *QoS Concepts*

  http://www.cisco.com/warp/customer/788/pkt-voice-general/bwidth_consume.html
- *QoS Features for Voice,* http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/qos_c/qcprt7/qcdvoice.htm
- *Configuring Quality of Service for Voice,* http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fvvfax_c/vvfqos.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) The solution for queuing delay where a voice packet is waiting for a large data packet to be serviced first is called _____.

A) dejitter buffer

B) DSP coding

C) header compression

D) LFI

E) payload compression

Q2) Dejitter buffers are used on the _____ side of the network to smooth delay variability.

A) originating

B) transit

C) receiving

D) originating and receiving

Q3) The uncompressed IP/UDP/RTP header is 40 bytes while the compressed IP/UDP/RTP header is _____ bytes.

A) 2 to 3

B) 2 to 5

C) 2 to 4

D) 1 to 5

Q4) Which two network features create voice echo on the network? (Choose two.)

A) delay

B) diameter

C) protocol

D) magnitude

E) topology

Q5) Arrange these codecs in descending order by bandwidth.

_____ 1. G.729

_____ 2. G.711

_____ 3. G.728

Q6) Which method uses a measurement-based approach to determining speech quality?

A) MOS

B) PSQM

C) SMQS

D) MOSQ

Q7) Match each ITU standard with its coding method.

A) PCM

B) ADPCM

C) CS-ACELP

D) LD-CELP

E) ACELP/MP-MLQ

_____ 1. G.711

_____ 2. G.723.1

_____ 3. G.729

_____ 4. G.728

_____ 5. G.726

Q8) Which queuing mechanism is recommended for most VoIP designs?

A) WFQ

B) CQ

C) LLQ

D) CBWFQ

E) IP RTP priority

Q9) Which two queuing mechanisms service the strict priority queue at the expense of lower queues leading to the potential of queue starvation? (Choose two.)

A) CQ

B) LLQ

C) FIFO

D) WFQ

E) CB-WFQ

F) PQ

Q10) Match each process with the category of QoS mechanism to which it belongs.

A) RED, WRED

B) class maps, access control lists

C) link fragmentation

D) queuing

E) traffic shaping, traffic policing

_____ 1. classification and marking

_____ 2. traffic conditioners

_____ 3. congestion avoidance

_____ 4. congestion management

_____ 5. link efficiency

Q11) Which queuing mechanism is based on WFQ and provides no strict priority queue for real-time traffic?

A) LLQ

B) PQ

C) CQ

D) CB-WFQ

Quiz Answer Key

Q1) D

Relates to: Delay, Jitter, and Loss Considerations

Q2) C

Relates to: Delay, Jitter, and Loss Considerations

Q3) C

Relates to: Delay, Jitter, and Loss Considerations

Q4) A, E

Relates to: Echo Considerations

Q5) Correct order: 2, 3, 1

Relates to: Bandwidth Considerations

Q6) B

Relates to: Bandwidth Considerations

Q7) A=1, B=5 , C=3 , D=4 , E=2

Relates to: Bandwidth Considerations

Q8) C

Relates to: QoS Mechanisms and Their Impact on Voice Quality

Q9) B, F

Relates to: QoS Mechanisms and Their Impact on Voice Quality

Q10) A=3, B=1, C=5, D=4, E=2

Relates to: QoS Mechanisms and Their Impact on Voice Quality

Q11) D

Relates to: QoS Mechanisms and Their Impact on Voice Quality

Planning Capacity Using Voice Traffic Engineering

Overview

Effective capacity-planning design minimizes degraded voice service in integrated networks. Capacity planning considers all network resources, from trunks and DSPs to WAN and the campus infrastructure. When sufficient resources are not available, you will need to reroute calls to the next available route. The lesson focuses on capacity planning based on grade of service, voice quality, and costs.

Relevance

Traffic engineering is a critical task to ensure that a network supports high quality voice.

Objectives

Upon completing this lesson, you will be able to plan resource capacities for quality packet telephony. This includes being able to meet these objectives:

- Describe the difference between on-net and off-net calling and how to use on- and off-net calling to minimize voice communication costs
- Calculate the probability of a voice call being blocked on an IP telephony network
- Provision trunks to provide sufficient voice capacity at the lowest possible cost
- Explain how to plan WAN link capacity for voice and how to prevent blockage or degradation of service
- Explain how to plan campus network capacity for voice and how to prevent blockage or degradation of service

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Familiarity with traditional and modern telephony concepts and terms, voice coding, and QoS mechanisms

Outline

The outline lists the topics included in this lesson.

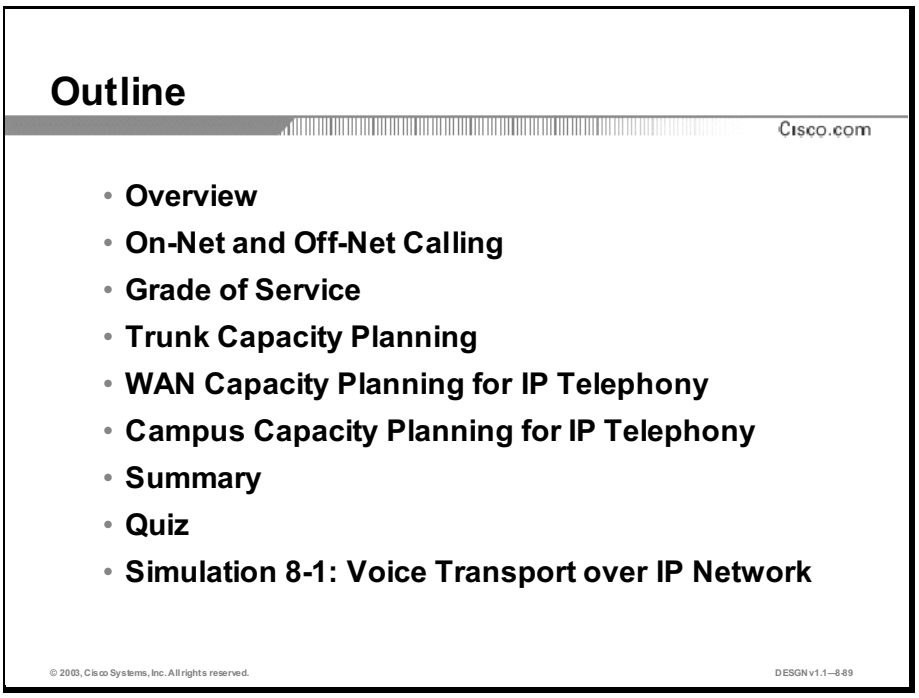

On-Net and Off-Net Calling

Many companies oversubscribe their private tie lines, assuming that not all callers use telephones and data connections over those lines at the same time. When private tie lines become congested and users cannot initiate calls, companies can use features called on-net calling and off-net calling. This topic introduces on-net and off-net calling and describes how on-net calls reduce telecommunications costs.

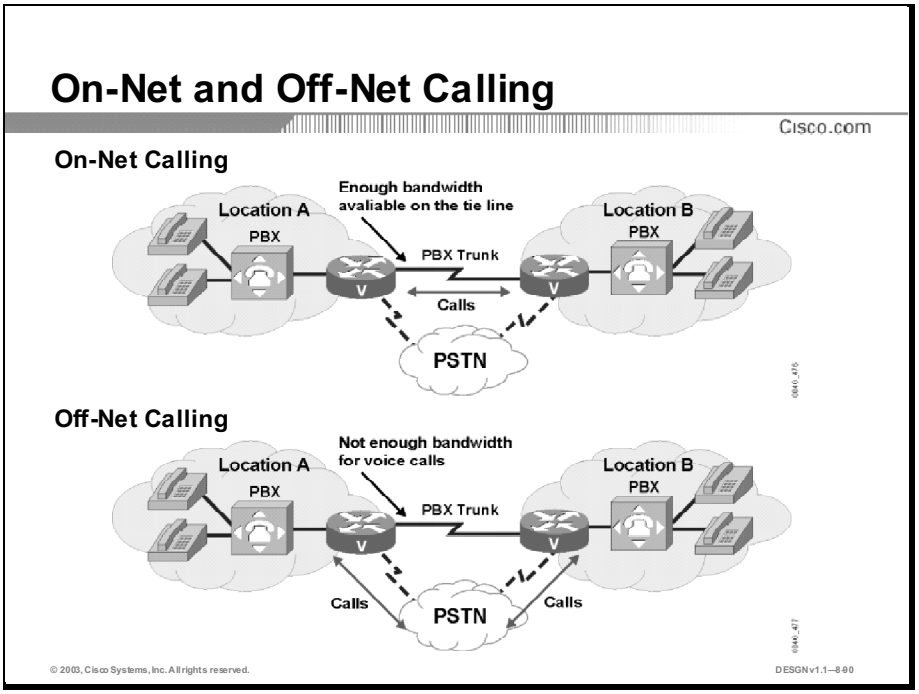

On-net and off-net calling are defined as follows:

- **On-net calling:** On-net calling transmits voice calls over private tie lines. With on-net calling, the call originates and terminates on a private network. Companies should use on-net calling whenever possible because on-net calling uses the existing private infrastructure.
- **Off-net calling:** A call is made to a destination located in the same city. The call is sent from the local voice-enabled router to the PSTN for call termination.

Two combinations of on-net and off-net calling are:

- **On-net to off-net calling:** If a private tie line is congested, voice gateways can select an alternate path, usually the PSTN. This is also known as automatic route selection. Another type of on-net to off-net call is when a call originates in the private network and is carried over private tie lines close to the destination of the call and then connected to the PSTN. This type of calling is sometimes referred to as Tail End Hop-Off (TEHO) or least cost routing.
- **Off-net to on-net calling:** A call originates at a site connected via the PTSN and terminates in the private network. This type of calling is sometimes referred to as Head End Hop-On (HEHO).

Note: CallManager and voice gateways allow on-net and off-net automatic route selection and manual route selection. Manual route selection enables callers to dial an access number, usually “9,” to access an outbound (PSTN) trunk or “8” to reach other locations on the corporate network. With this feature, a caller selects the path manually.

Example: On-Net to Off-Net Calling

In on-net calling, the user places a call from location A to location B across the data network using a private tie line.

In off-net calling, either the user dials an access code or the voice gateway automatically selects the path to gain access to the PSTN. When the call reaches the PSTN, the user hears a second dial tone. At this point, the user dials the destination number. If on-net resources are unavailable to complete the call, the voice gateway responds in these ways:

- If the far-end voice gateway has no available circuit to the PBX to which it is connected, the voice gateway sends the busy-back signal and the near-end voice gateway reroutes the call off-net.
- If the near-end voice gateway uses QoS to detect that there is not enough bandwidth to complete the call, it reroutes the call off-net.

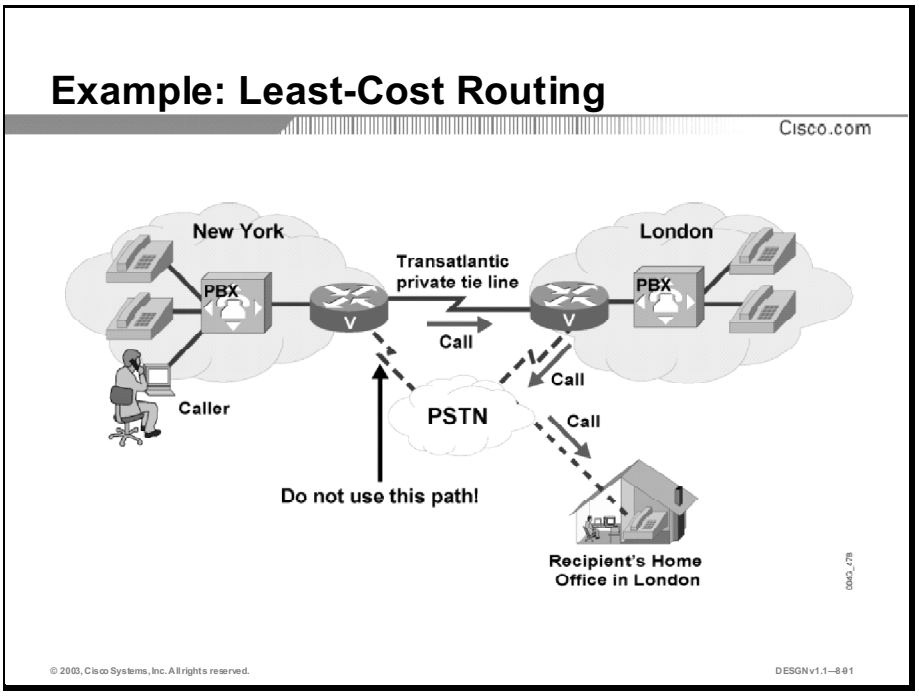

The figure shows an example of least-cost routing using private tie lines on a long-distance call and the PSTN locally to the destination number.

A caller in a branch office in New York places a call to a recipient in a home office in London. On the voice gateway in New York, the voice is coded and compressed into the VoIP call and transported over the private tie line to the London office, where the signal is decoded and decompressed and transported over the PSTN to the home office.

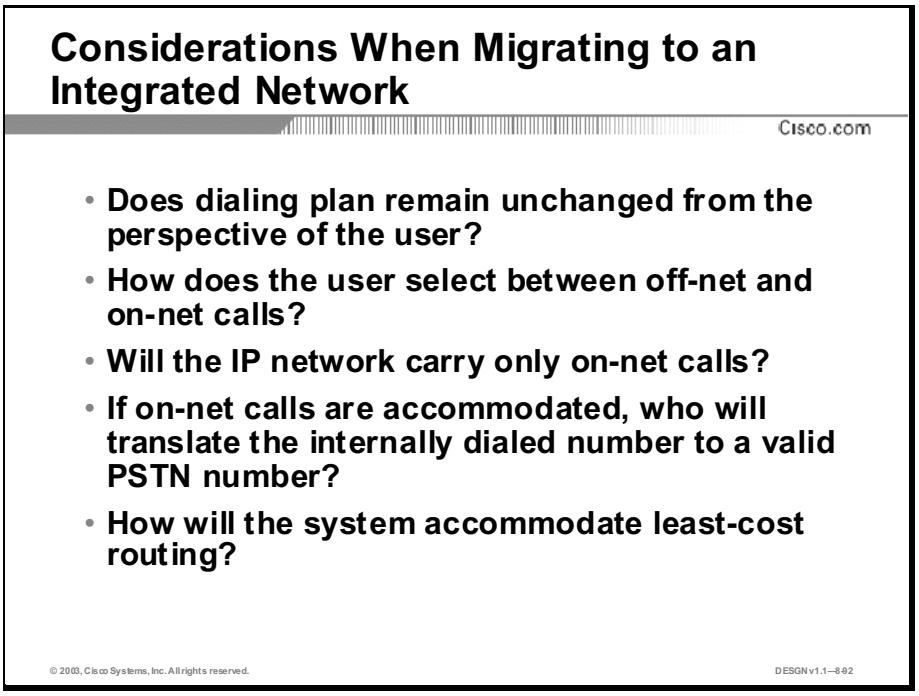

When migrating from separate data and voice networks to an integrated network, enterprises must consider several factors regarding the dialing plan to ensure a smooth migration and a manageable network.

Consider these factors:

- Does the dialing plan remain unchanged from the perspective of the user, or will you introduce a new access code to allow users to choose the IP network or the PSTN?
- Will the IP network carry only on-net calls, or will it also provide gateways to the PSTN?
- If the IP network will provide only on-net calls, then the enterprise needs to implement only its private dialing plan.
- If the IP network will accommodate off-net calls, the voice gateway must translate an internally dialed number to a publicly accessible PSTN number, including adding and deleting country codes and regional area or city codes.
- If least-cost routing is required, what device will provide it? Can the voice gateway support least-cost routing?

Grade of Service

Traffic engineering is a science of selecting the right number of lines and the proper types of service to accommodate users. Several measurements are available to help voice traffic engineers determine how to provision trunks. This topic shows you how to calculate the probability of a voice call being blocked on an IP telephony network.

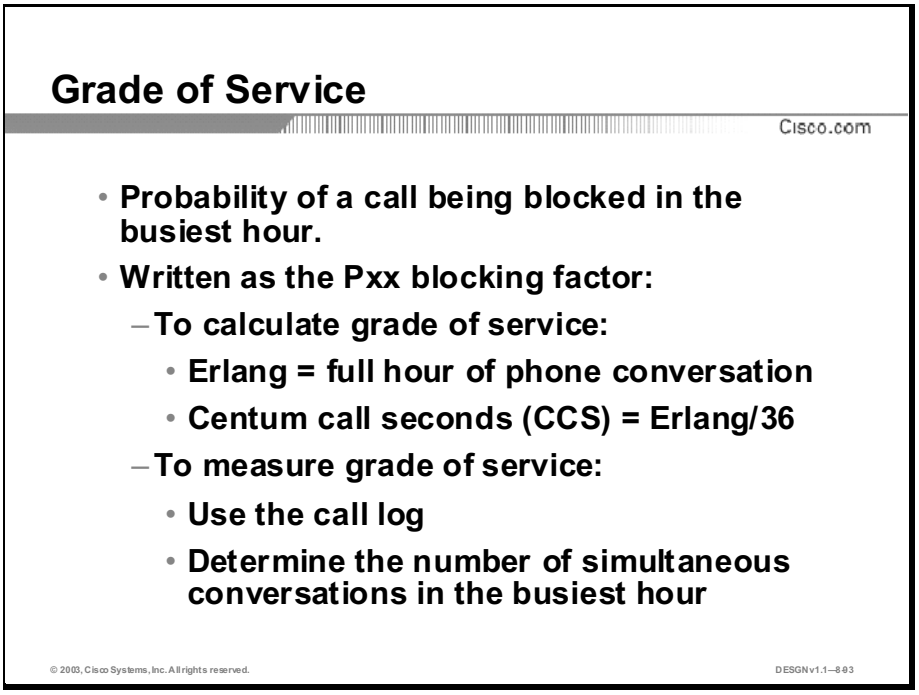

Grade of service is defined as the probability that a voice gateway will block calls while attempting to seize circuits. It is written as a decimal fraction, called the Pxx blocking factor, or blockage, where xx is the percentage of calls that are blocked for a traffic system. For example, traffic facilities requiring a P01 grade of service define a 1 percent probability of callers being blocked.

Measuring Grade of Service

These grade-of-service measurements exist:

- **Erlang:** One Erlang equals one full hour, or 3600 seconds, of telephone conversation. If a trunk carries 12.35 Erlangs during an hour, an average of a little more than 12 lines (connections) are busy.
- **Centum call second (CCS):** A centum (one hundred) call second represents 1/36th of an Erlang. To calculate a CCS, multiply the number of calls per hour by their average duration in seconds, and divide the result by 100. A system port that can handle a continuous one-hour call has a traffic rating of 36 CCS. Station traffic varies greatly among users, but the typical range is about 6 to 12 CCS per port. If you cannot obtain exact statistical data, assume that the average typical trunk traffic is 30 CCS per port.
- **Simultaneous conversations:** To determine traffic capacity, you can maintain a call logger and plot the number of simultaneous conversations on the network to determine the probability that exactly *x* simultaneous calls will occur. Provision voice systems to allow the maximum number of simultaneous conversations expected at the busiest time of day.

Example: Erlang Calculation

One hour of conversation, or one Erlang, may include 10 six-minute calls or 15 four-minute calls. Receiving 100 calls with an average length of 6 minutes is equivalent to 10 Erlangs, or 360 CCS. The formula to calculate Erlangs is:

- Number of calls x call duration in seconds / 1 Erlang in seconds = Number of Erlangs required
- 100 x (6 x 60) /3600 = Erlangs required
- 100 x 360 /3600 = Erlangs required
- 36000/3600 = 10 Erlangs required

The formula to calculate CCS is:

- Number of calls x call duration in seconds / 100 = centum seconds
- 100 x (6 x 60) / 100 = centum seconds
- 100 x 3600 /100 = centum seconds
- 36000/100 = 360 centum seconds

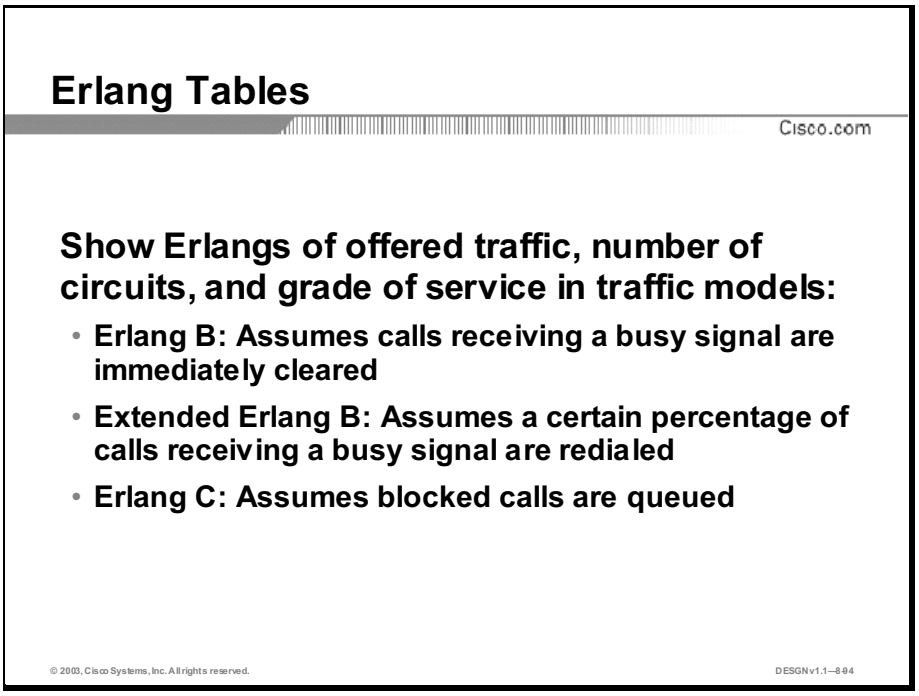

Erlang tables show the amount of traffic potential for specified numbers of circuits, including the probability of receiving a busy signal. The calculation results are usually stated in CCS. Erlang tables combine offered traffic, number of circuits, and grade of service in the following traffic models:

- **Erlang B:** This is the most common traffic model used to calculate the number of lines required if you know the traffic figure (in Erlangs) during the busiest hour. The model assumes that all blocked calls are immediately cleared.
- **Extended Erlang B:** This model is similar to Erlang B, but it considers the additional traffic load caused when blocked callers immediately trying to call again. You can specify the retry percentage.
- **Erlang C:** This model assumes that all blocked calls stay in the system until the voice gateway can handle them. You can apply the Erlang C model to a call center design where calls enter a queue if an agent is not available.

Note: You can find Erlang tables at http://www.erlang.com/.

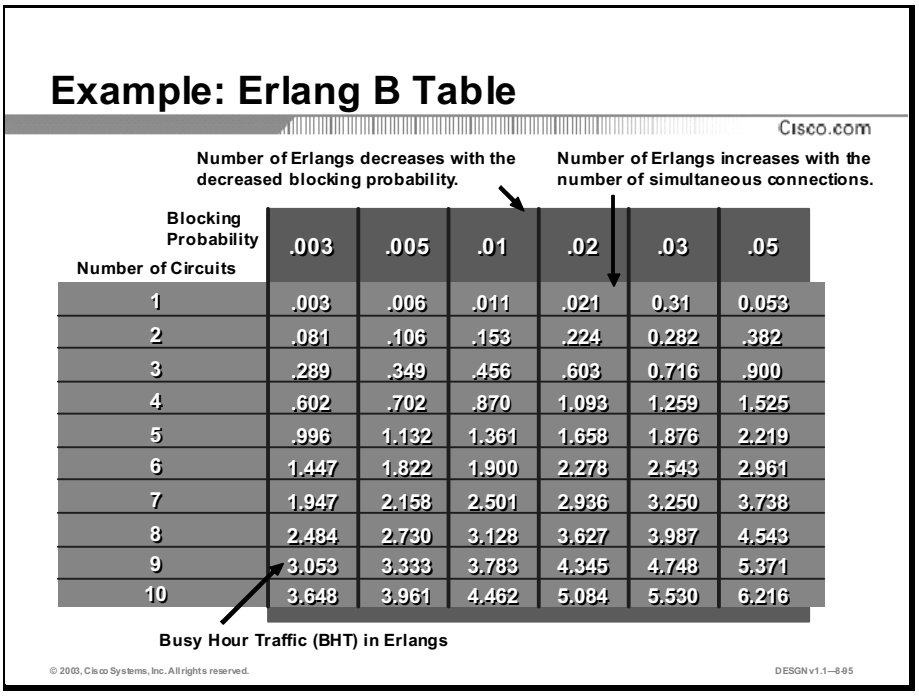

The figure shows the amount of traffic in Erlangs, number of circuits (simultaneous connections), and the grade of service. The Erlang B table presents:

- For 10 simultaneous connections with a grade of service P01 (1 percent block probability), 4.462 Erlangs of traffic is offered. 4.462 Erlangs equals approximately 160 CCS (4.462 * 36). Assuming that there are 20 users in the company, every user can talk for about 8 minutes every hour.
- 2.961 Erlangs are handled by six circuits at grade of service P05. 2.961 Erlangs equals approximately 107 CCS (2.961 * 36). Assuming that there are 10 users in the company, every user can talk for about 10.7 minutes every hour.

Trunk Capacity Planning

Trunk capacity planning, or provisioning, establishes the number of circuits (voice channels) from the PBX to the integrated network. After an organization establishes the number of circuits, the next step is to use that number to determine the required network bandwidth. This topic explains how to provision trunks to provide sufficient voice capacity at the lowest possible cost.

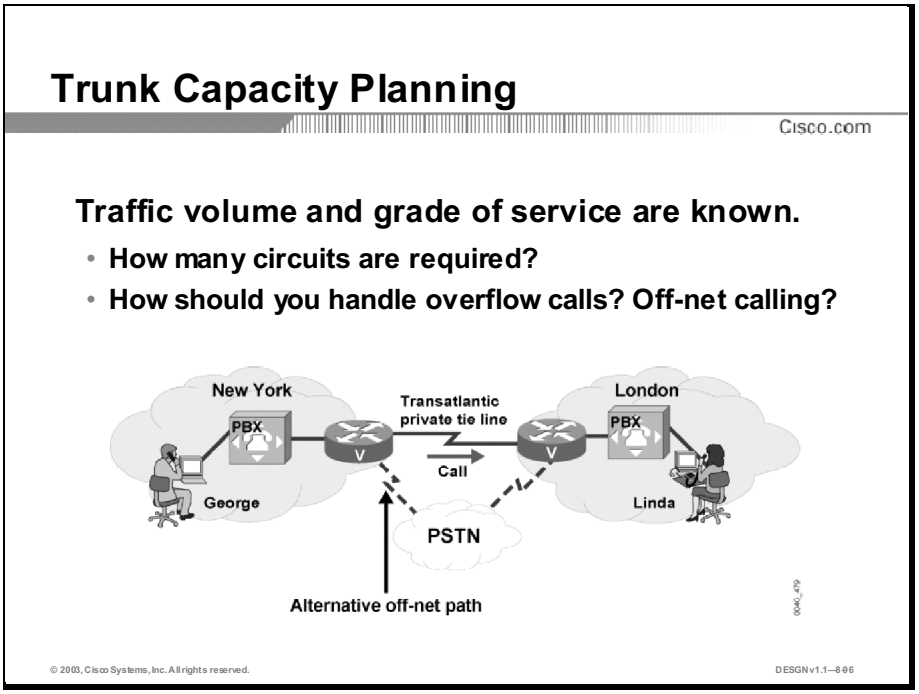

To design a successful voice operation, you need to estimate the number of circuits needed for a known amount of traffic. Organizations can use the traffic volume and flow, the selected grade of service or blocking factor, and other network-specific objectives to determine the correct number of circuits. Some organizations move circuits from the current network to the integrated network. Other organizations take this opportunity to conduct a traffic study and update their traffic engineering information.

It is important to design the network so that a voice gateway reroutes voice calls that oversubscribe the allocated network bandwidth via an alternate path such as the PSTN. A QoS category called Call Admission Control (CAC) serves as a base for handling oversubscribed calls off-net. For example, if the WAN access link from a branch office is provisioned to carry no more than five simultaneous calls, CAC prevents the sixth call from entering the network and impairing voice quality. This activity is also known as protecting voice from voice.

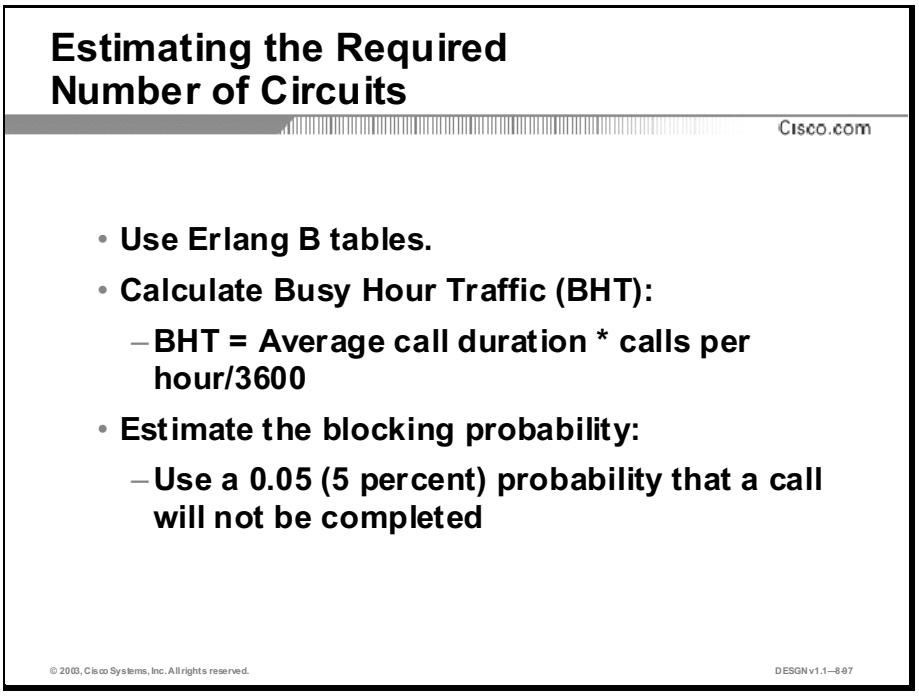

Using the Erlang B traffic model to estimate the number of circuits requires two inputs:

- **Busy Hour Traffic:** The BHT value represents the quantity of traffic, expressed in Erlangs. You will provide a BHT estimate, which represents the number of hours of traffic that is transported across a trunk group in its busiest hour. For example, if you know from your call logger that 350 calls are made on a trunk group in the busiest hour and the average call duration is 180 seconds, then the BHT value will be:
  - BHT = Average call duration * calls per hour/3600
  - BHT = 180 * 350/3600
  - BHT = 17.5 Erlangs
- **Blocking probability:** The blocking probability value describes the calls that are not completed because insufficient lines are available. A value of 0.01 means that the voice gateway will block 1 percent of calls.

  Most telecommunications organizations use a blocking factor of 5 percent. Degrading the service by 1 percentage point results in a large trunk savings, because circuit numbers (and cost) decrease exponentially with the grade of service.

Example: Estimating the Number of Circuits Required

After you establish the BHT and blocking probability, you can estimate of the number of circuits required using the Erlang B traffic model:

- BHT = 17.986 Erlangs
- Blocking = 0.01

According to the Erlang table, the number of required circuits is 27.

Note: You can find the Erlang table (calculator) at http://www.erlang.com/calculator/erlb/.

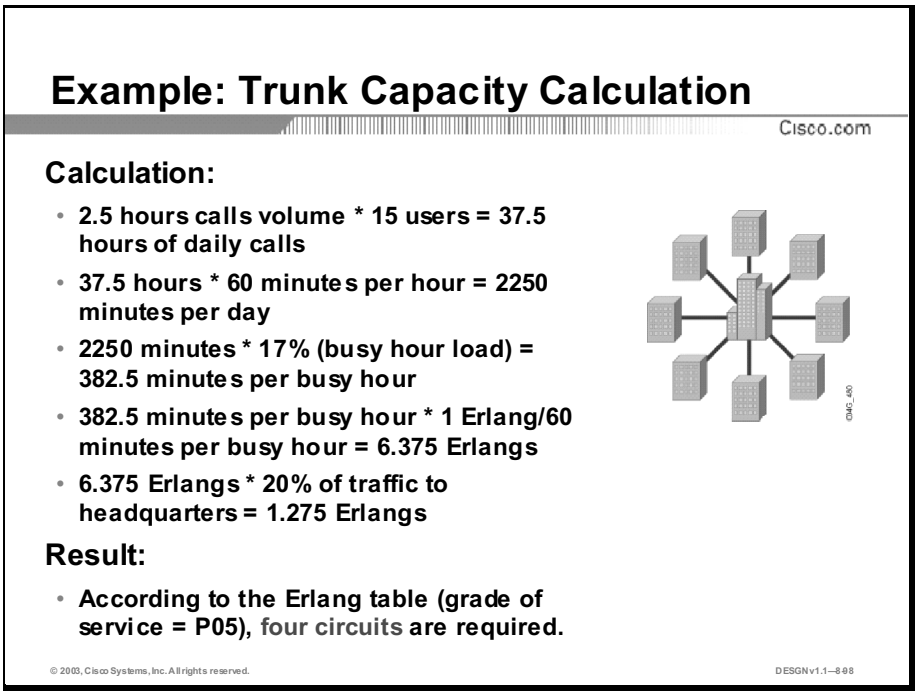

The figure shows an example of a trunk capacity calculation among branch offices and a main office. These assumptions guided the redesign of the network:

- Approximately 15 people occupy each branch office.
- The bidirectional voice and fax call volume totals about two and a half hours per person, per day, per branch office.
- Approximately 20 percent of the total call volume is between headquarters and the branch offices.
- The network design is a star topology that connects each branch office to the headquarters.
- A busy-hour loading factor of 17 percent is appropriate.
- One 64-kbps circuit supports one call.

The voice and fax traffic calculations are:

- 2.5 hours call volume per user per day * 15 users = 37.5 hours daily call volume per office
- 37.5 hours * 60 minutes per hour = 2250 minutes per day
- 2250 minutes * 17 percent (busy-hour load) = 382.5 minutes per busy hour
- 382.5 minutes per busy hour * 1 Erlang/60 minutes per busy hour = 6.375 Erlangs
- 6.375 Erlangs * 20 percent of traffic to headquarters = proposed volume of 1.275 Erlangs

To determine the appropriate number of trunks required to transport the traffic and given the desired grade of service, consult traffic engineering tables. This organization chose a P05 grade of service. Using the calculated Erlangs and the Erlang B table, the organization determined that four circuits are required for communication between the branch office and the main office.

Note: For information on Erlangs and to use online Erlang calculators, please refer to http://www.erlang.com.

Example: Off-Net Calls Cost Calculation

Cisco.com

- **Given a voice gateway controlling 10 circuits (connections) between New York and London, how many minutes per month will the system overflow to off-net calling?**
  - **Each call uses one circuit (connection)**
  - **Acceptable grade of service = P03**
  - **21 business days per month, 2 peak hours per day**
- **What are the monthly costs of the off-net calling between New York and London if one minute costs US$0.10?**

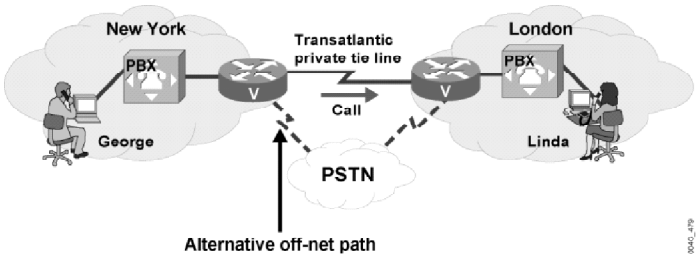

The figure shows an example of an off-net cost calculation between two locations: New York and London. The company uses the PSTN path when the transatlantic tie line cannot accept additional on-net calls.

Assume that all calls between New York and London use 64 kbps of bandwidth, which corresponds to one circuit, and the grade of service of .03 is acceptable. How many minutes of calls will use off-net calling due to a block of service on the transatlantic tie line? The transatlantic tie line can carry a maximum of 10 calls simultaneously. In the calculation, assume that the one-minute call between New York and London costs $0.10.

The calculation is:

- According to the Erlang B table, at P03 and 10 circuits, 5.53 Erlangs are available.
- At P03, 3 percent of the 5.53 calls will overflow and the system will send them off-net.
- Therefore, in the peak hour, .03 * 5.53 Erlangs * 60 minutes = 10 overflow minutes.
- If there are two peak hours per day and 21 business days per month, then 21 days * 2 peak hours per day * 10 overflow minutes = 420 minutes per month.
- 420 overflow minutes per month * $0.10 per overflow minute = $42.

The calculation shows that 420 minutes per month of off-net calling between New York and London is used, and costs $42.

The company should compare the off-net calling cost to the cost of additional circuits between New York and London.

WAN Capacity Planning for IP Telephony

WAN capacity planning for voice and data transport depends on a number of parameters, including the number of simultaneous voice calls, sampling rate, codec, link type, header compression techniques, and use of VAD. This topic explains how to plan WAN link capacity for voice and how to prevent blockage or degradation of service.

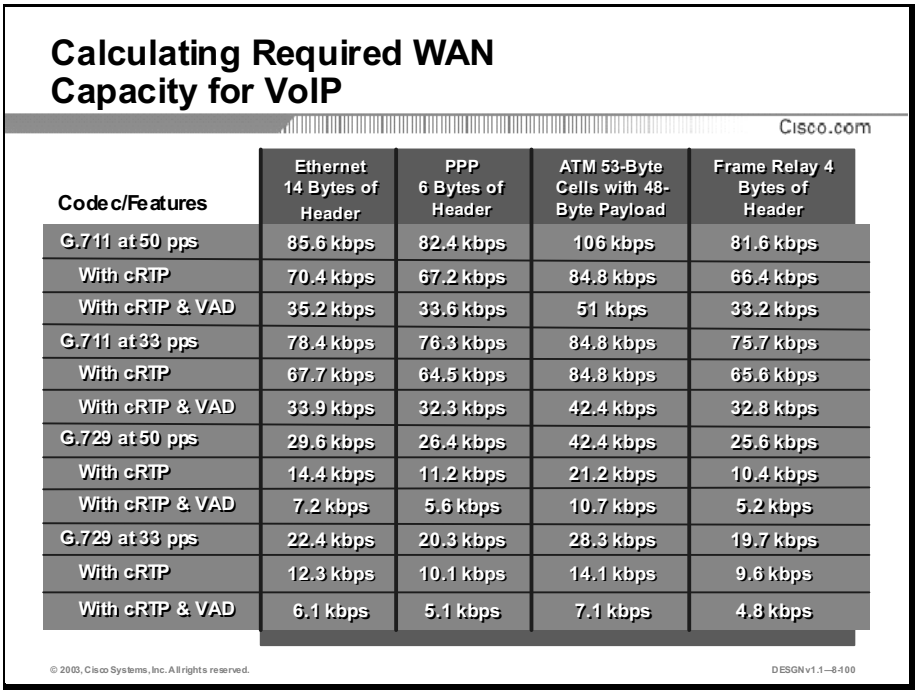

Codec/Features	Ethernet 14 Bytes of Header	PPP 6 Bytes of Header	ATM 53-Byte Cells with 48-Byte Payload	Frame Relay 4 Bytes of Header
G.711 at 50 pps	85.6 kbps	82.4 kbps	106 kbps	81.6 kbps
With cRTP	70.4 kbps	67.2 kbps	84.8 kbps	66.4 kbps
With cRTP & VAD	35.2 kbps	33.6 kbps	51 kbps	33.2 kbps
G.711 at 33 pps	78.4 kbps	76.3 kbps	84.8 kbps	75.7 kbps
With cRTP	67.7 kbps	64.5 kbps	84.8 kbps	65.6 kbps
With cRTP & VAD	33.9 kbps	32.3 kbps	42.4 kbps	32.8 kbps
G.729 at 50 pps	29.6 kbps	26.4 kbps	42.4 kbps	25.6 kbps
With cRTP	14.4 kbps	11.2 kbps	21.2 kbps	10.4 kbps
With cRTP & VAD	7.2 kbps	5.6 kbps	10.7 kbps	5.2 kbps
G.729 at 33 pps	22.4 kbps	20.3 kbps	28.3 kbps	19.7 kbps
With cRTP	12.3 kbps	10.1 kbps	14.1 kbps	9.6 kbps
With cRTP & VAD	6.1 kbps	5.1 kbps	7.1 kbps	4.8 kbps

The most important issue in the WAN capacity-planning process is the number of simultaneous calls that are permitted across a WAN link. You can limit the number of calls across a WAN link by using grade of service and CAC.

The next step is to understand the amount of bandwidth required for one voice call. Bandwidth requirements for a voice call can vary, depending upon the situation. It is important to understand the components of a VoIP packet and the variables that affect overall use.

Calculating Per-Call Bandwidth Requirements

Every voice packet includes payload, RTP header, UDP header, IP header, and link header. All media types include RTP, UDP, and IP headers at 40 bytes per packet; only the link header size differs. The sampling rate does not significantly impact bandwidth for the payload. However, when overhead is added, there are significant increases with 50 packets per second versus 33 packets per second.

Note: In any conversation, two streams are required: one in each direction.

Remember that WAN links are generally full duplex. Therefore, you should allocate an equal amount of bandwidth in each direction for one voice call. Be careful when estimating bandwidth by using VAD because the values in the table do not represent the actual required bandwidth in any one direction at any particular time.

Improving Voice Bandwidth Allocations

Network planners can improve bandwidth allocations by using RTP header compression and VAD. RTP header compression reduces the size of the IP/UDP/RTP header from 40 bytes to 2 to 4 bytes.

VAD reduces bandwidth requirements based on the theory that in any given voice call only one party is talking at a time, and that periods of nonactivity are not transmitted across the link. VAD can save up to 50 percent of overall bandwidth. You must be careful because one voice stream may use 100 percent of the available bandwidth at any given time.

Example: WAN Capacity Calculation to a Remote Site

An organization has a remote office with 20 permanent employees. The remote office has only three analog lines, so users can only place calls when an outside line is available. The network planner determines that users cannot place calls approximately 10 times a day, which is unacceptable. Therefore, the network planner deploys VoIP using the existing Frame Relay link between the main office and the remote office, providing bandwidth for four simultaneous voice calls.

The planner tests compression techniques and decides to use G.729 encoding with cRTP over Frame Relay at 50 packets per second (pps). Using the available information, the planner determines that each call will use approximately 10.4 kbps per stream. However, the planner is uncomfortable with only 41.6 kbps in each direction and allocates 64 kbps across Frame Relay for voice calls to guarantee voice quality.

Because the remote office currently has a 64 kbps CIR over Frame Relay for data, the planner doubles the CIR to 128 kbps and configures the appropriate QoS, traffic shaping, and FRF.12 Frame Relay fragmentation to provide acceptable voice quality.

Example: WAN Capacity Between Two Locations

Based on the Erlang table, 10 circuits are required between two locations to satisfy user demands. The designer implements a PPP link between the two locations to transport VoIP. The designer codes voice with the G.729 codec using 50 samples per second, and compresses the header with cRTP.

The per-call bandwidth table shows that 11.2 kbps of bandwidth is required for one voice call and 112 kbps of bandwidth is required between the two locations in each direction to carry 10 simultaneous voice calls.

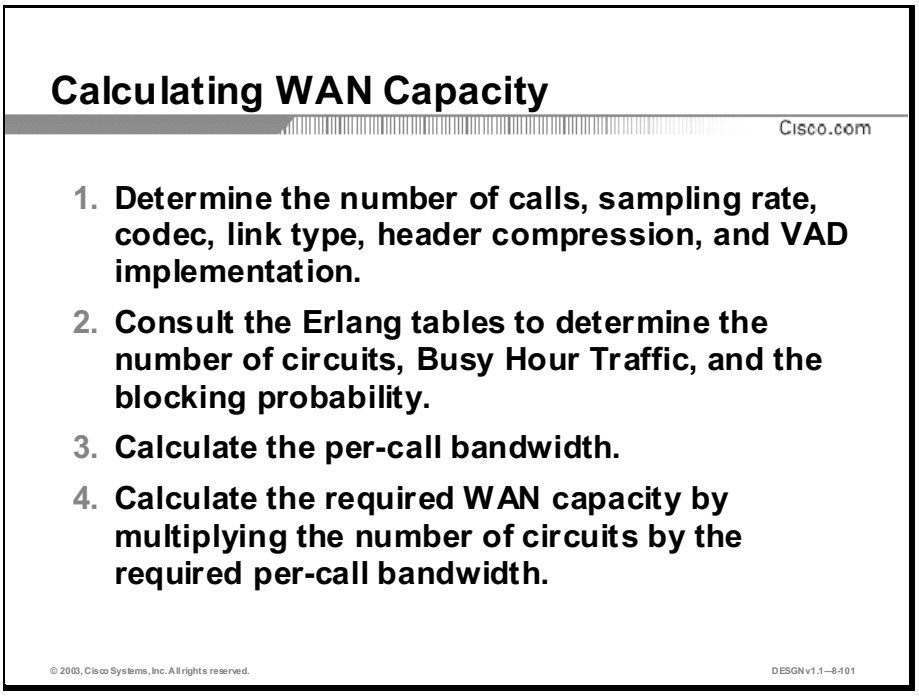

By determining the per-call bandwidth requirements and the number of simultaneous calls, you can calculate the amount of bandwidth with a certain grade of service. You can use the WAN capacity calculation to estimate the bandwidth that must be provided through an IP-based network to transport a given BHT level satisfactorily. The calculation is based on the Erlang B traffic model.

To determine the required WAN capacity, follow the steps listed in the figure.

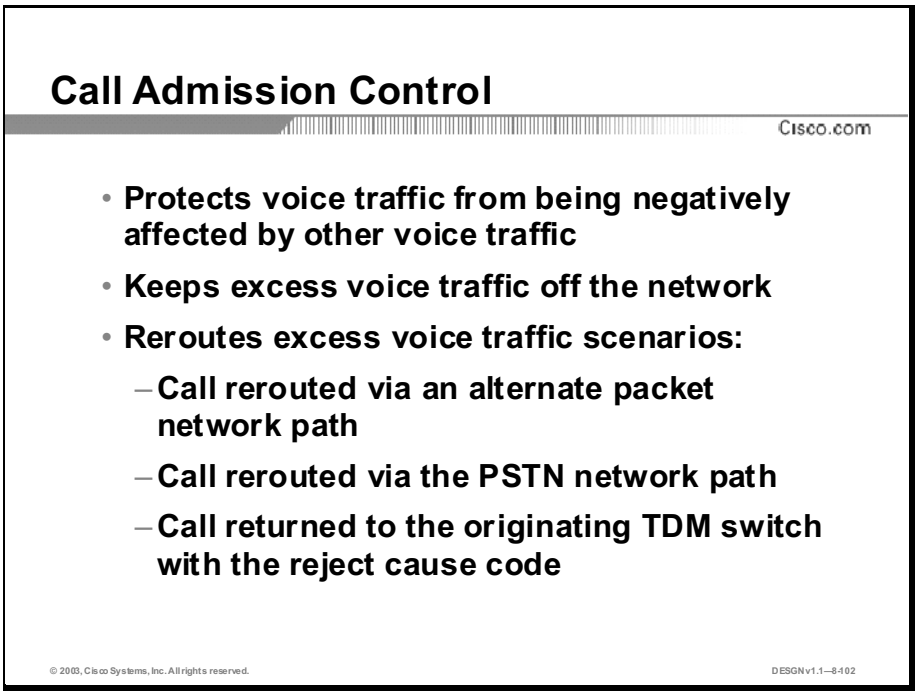

Call Admission Control (CAC) mechanisms extend the QoS capabilities to protect voice traffic from being negatively affected by other voice traffic and to keep excess voice traffic off the network.

If the WAN access link between two PBXs has the bandwidth to carry only two VoIP calls, admitting a third call will impair the voice quality of all three calls. The queuing mechanisms that provide policing, cause this problem, not CAC. If a voice gateway receives packets that exceed the configured or allowable rate, it tail-drops these packets from the queue. The queuing mechanism cannot distinguish which IP packet belongs to which voice call. The voice gateway drops packets that exceed the given arrival rate within a certain period of time. Thus, all three calls experience packet loss, which end users perceive as clipped speech.

Note: This problem is easier to solve for the Layer 2 voice transport mechanisms (VoFR and VoATM) but is particularly challenging for VoIP applications.

Call Rerouting Alternatives

When you implement CAC, the outgoing gateway detects that insufficient network resources are available to process a call. The gateway rejects the call, and the originating gateway must find another way to handle the call. In the absence of any specific configuration, the outgoing voice gateway sends a reorder tone to the calling party. The PSTN switch or PBX may then announce, “All circuits are busy. Please try your call again later.”

You can configure the outgoing voice gateway to accomplish the following tasks:

- Reroute the call via an alternate packet network path, if such a path exists
- Reroute the call via the PSTN network path
- Return the call to the originating TDM switch with the reject cause code

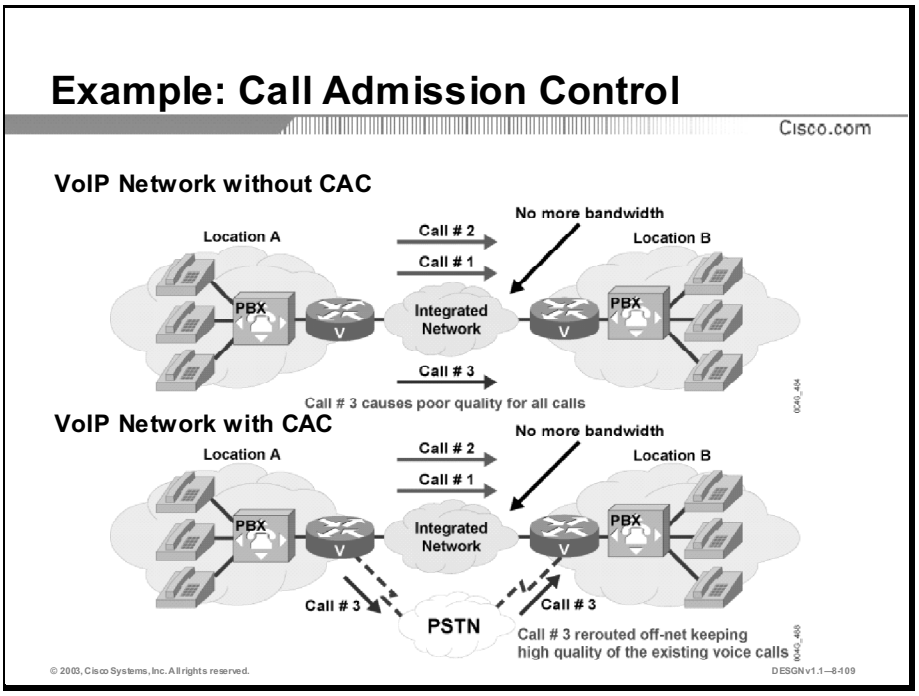

The first example illustrates a VoIP network without CAC. Suppose that the WAN access link between two PBXs has the bandwidth to carry only two VoIP calls. Admitting the third call impairs the voice quality of all three calls.

The second example illustrates a VoIP network with CAC. Suppose that the outgoing gateway detects that insufficient network resources are available to allow a call to proceed. The gateway automatically reroutes the call off-net, maintaining the voice quality of the two existing calls.

Campus Capacity Planning for IP Telephony

Capacity planning is a critical process for enterprise IP telephony migration and overall success. To plan capacity for a campus network that supports voice, you will configure Cisco CallManager requirements, network capacity and performance, and trunking capacity. This topic explains how to plan campus network capacity for voice and how to prevent blockage or degradation of service.

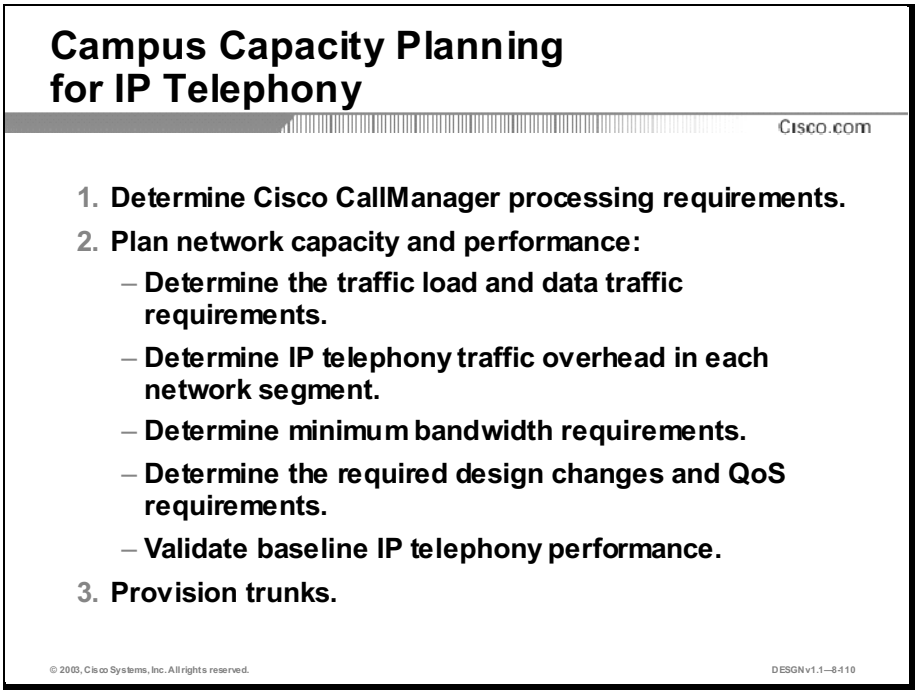

To perform capacity planning for the campus, follow these steps:

Step 1 **CallManager processing requirements:** Help to ensure that the CallManager servers have sufficient resources for normal call processing, voice conferencing, and other IP telephony services.

CallManager servers are typically deployed in clusters. Each cluster has a dedicated publisher for database replication. Individual IP telephony devices and Cisco CallManager services are supported based on the assigned weight for each device and service. The assigned weights impact server resource utilization and processor performance and directly relate to the number of devices and services each server in a cluster can support. Depending on the version of CallManager server software, each server in a cluster can control from 10,000 to 30,000 IP Phones, or from 2,500 to 7,500 IP Phones per active CallManager server in the cluster.

Step 2 **Plan network capacity and performance:** Ensure that the network can support the additional IP telephony traffic and that the traffic will meet delay and jitter requirements. This is the recommended process for network capacity and performance planning:

- Determine the current traffic load and data traffic requirements for a combined IP telephony and data architecture. Baselining network use helps to determine the current traffic load and data traffic requirements for a combined IP telephony and data architecture. A relatively homogeneous environment may require a

sample baseline calculation only from representative links rather than data from the entire network. Use peak use and average use statistics to describe link use.

- Determine IP telephony traffic overhead in required sections of the network based on busy-hour estimates, gateway capacities, and CallManager capacities. Differentiate buildings or sites where more or less voice traffic is required. Use an Erlang calculator to determine busy-hour call requirements. Multiply the busy-hour call requirements by the voice encoding method to determine busy-hour bandwidth requirements.
- Determine minimum bandwidth requirements adding the busy-hour data traffic and busy-hour voice traffic. Determine the growth requirements over time and perform link trending to determine overall bandwidth requirements after you implement both data and voice. Baseline and determine trends for link use over time.
- Determine the required design changes and QoS requirements, based on IP telephony design recommendations, voice bandwidth requirements, and data requirements. Configure QoS for both LAN and WAN environments to validate the network design with a performance baseline.
- Validate baseline IP telephony performance with a performance baseline before the voice network implementation. Investigate potential network issues, including queuing delay, CPU, link and buffer utilization, and error rates.

Step 3 **Provision trunks:** Determine the number of circuits required for PBX interconnectivity, voice mail connectivity, PSTN connectivity, and site-to-site connectivity. In addition, define the existing blocking factor for potential capacity issues. The recommended blocking factor (grade of service) is 1 percent.

Example: Campus Capacity Planning for IP Telephony

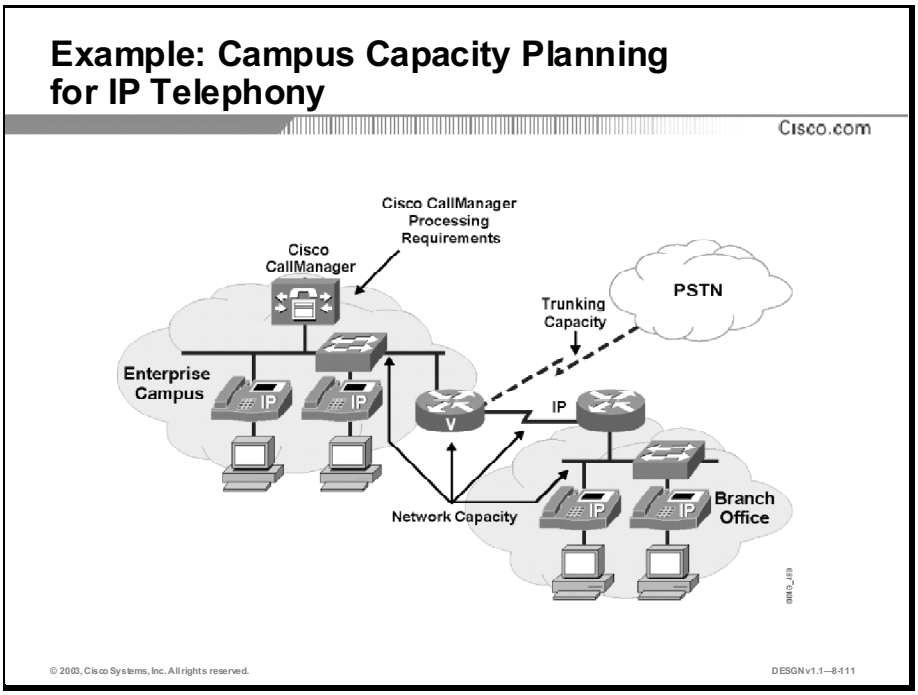

The example shows the three capacity-planning processes and their locations in the enterprise IP telephony network.

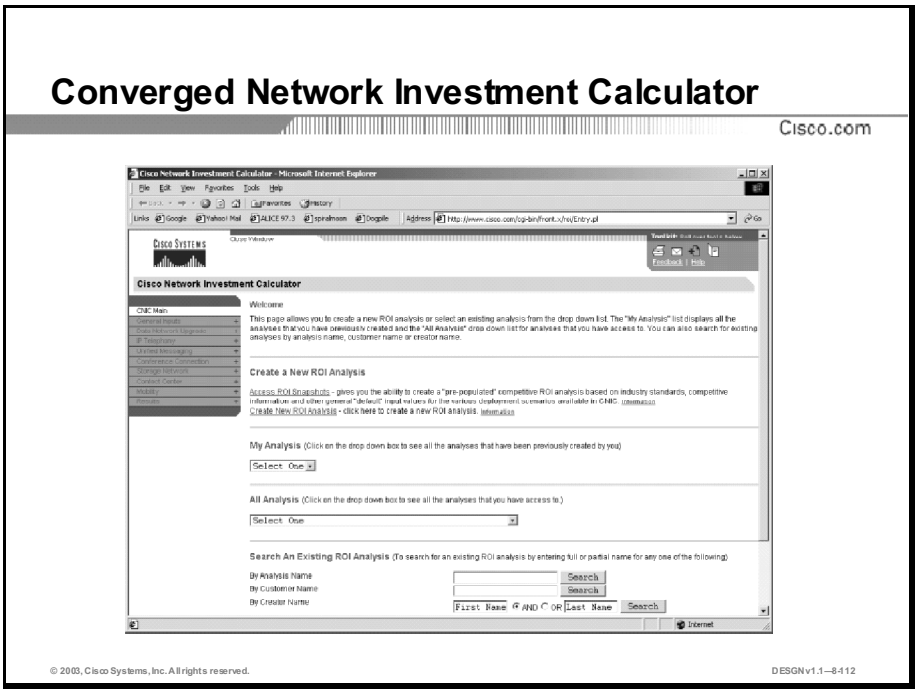

The Cisco Converged Network Investment Calculator (CNIC) helps Cisco partners create customer return on investment (ROI) data for converged networks. Using four common deployment scenarios, you can use the CNIC to gather and input information about the network and planned deployment, and interpret and present the results. CNIC considers the potential for overly optimistic ROI calculations by letting users enter in their own assumptions, from employee productivity gains to wiring drop cost estimates.

Using models for depreciation and other accounting principles, CNIC calculates detailed costs of IP telephony and unified messaging. CNIC is a web-based ROI calculator application that includes training and access to a helpline and numerous white papers, case studies, and presentations. CNIC is free and available to VoIP-specialized partners, who can access the application via a secure site once they have completed training. Partners complete a two-week data-gathering period using an extensive questionnaire that addresses sales and programmatic concerns.

Note: CNIC is available for authorized partners at http://www.cisco.com/cgi-bin/front.x/roi/Entry.pl.

CNIC calculates both ROI and payback. ROI is the average of the net benefits divided by the initial cost of the project, times 100. For instance, if a project cost $50 and returned $100 in the first year, the ROI in the first year would be 100/50, or 200%. Unfortunately, technology rarely covers its costs in the first year, so a more accurate calculation uses a 3-year horizon.

Payback period is the point where total benefits equal total costs. To calculate payback period if the initial year cost is less than the first year's benefit, simply divide the initial cost by the first year's benefit.

Summary

This topic summarizes the key points discussed in this lesson.

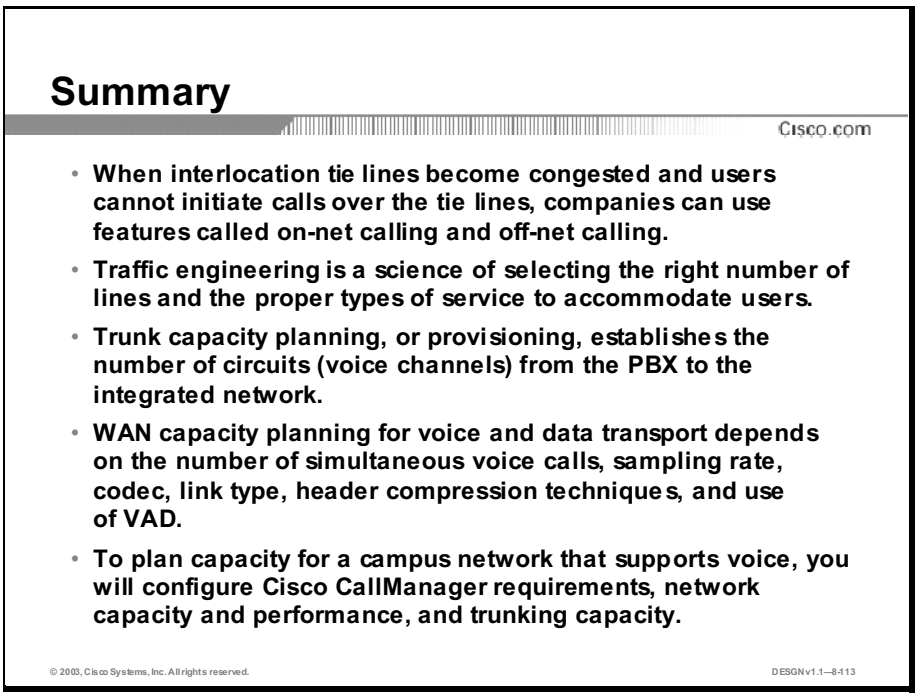

References

For additional information, refer to these resources:

- *Traffic Analysis for Voice over IP,* http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/voipsol/ta_isd.htm
- Westbay Engineers Limited home page http://www.erlang.com
- *VoIP - Understanding Codecs: Complexity, Support, MOS, and Negotiation,* http://www.cisco.com/warp/public/788/voip/codec_complexity.html
- *Cisco IP Telephony Network Design Guide,* http://cisco.com/en/US/netsol/ns110/ns163/ns165/ns268/networking_solutions_design_guidances_list.html

Next Steps

For the associated case study and exercises, refer to the following section that follows the Quiz:

- Simulation 8: Voice Transport over IP Network

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which feature allows a long-distance call to stay on-net as long as possible and go off-net locally, close to the called number?

A) off-net routing

B) least-cost routing

C) on-net routing

D) intelligent routing

E) destination routing

Q2) To estimate the number of circuits using the Erlang B traffic model, the following two inputs are required: _____ and _____. (Choose two.)

A) grade of service

B) the number of users

C) the amount of traffic

D) the average duration of a call

E) the number of calls per day

Q3) 10 Erlangs is equal to _____ centum call seconds.

A) 1/36

B) 1/100

C) 1

D) 36

E) 100

F) 360

Q4) The 1 percent probability that a call will be blocked while attempting to seize circuits is presented as _____.

A) P01

B) P001

C) P1%

D) 1%

E) P1.01

Q5) What two parameters are mandatory for the trunk capacity calculation? (Choose two.)

A) number of users per location

B) acceptable grade of service

C) average call duration

D) number of calls per hour

E) number of branch offices

F) traffic volume in the busiest hour

G) number of lines within the trunk

Q6) When calculating per-call bandwidth requirements, how is bandwidth calculated on a WAN link?

A) As a single bandwidth stream from the calling party to the called party.

B) As three streams, two equal amounts of bandwidth in each direction and a single bandwidth stream for call control.

C) As an equal amount of bandwidth in each direction.

D) As four streams of bandwidth, two in each direction carrying voice and call control traffic.

Q7) What is the formula used to calculate the Busy Hour Traffic?

A) BHT = Average call duration * calls per hour/3600

B) BHT = Total call duration * calls per hour/3600

C) BHT = Average call duration * blocking probability * calls per hour/3600

D) BHT = blocking probability * calls per hour/3600

Q8) Select four parameters that are required for WAN capacity planning to support IP telephony. (Choose four.)

A) payload compression techniques

B) packet rate

C) codec

D) DSP type

E) link type

F) QoS mechanism used

G) number of simultaneous voice calls

Q9) What is the purpose of the Call Admission Control (CAC) mechanism?

A) CAC routes voice calls off the network.

B) CAC protects data traffic from being negatively affected by voice traffic, and keeps excess voice traffic off the network.

C) CAC protects voice traffic from being negatively affected by other voice traffic, and keeps excess voice traffic off the network.

D) CAC protects voice traffic from being negatively affected by data traffic, and keeps excess voice traffic off the network.

Q10) Which three tasks should you complete to plan network capacity and performance to support voice in an enterprise campus network? (Choose three.)

A) Determine minimum bandwidth requirements.

B) Determine Cisco CallManager processing requirements.

C) Provision trunks.

D) Validate baseline IP telephony performance.

E) Determine the traffic load and data traffic requirements.

Quiz Answer Key

Q1) B

Relates to: On-Net and Off-Net Calling

Q2) A, C

Relates to: Grade of Service

Q3) F

Relates to: Grade of Service

Q4) A

Relates to: Grade of Service

Q5) B, F

Relates to: Trunk Capacity Planning

Q6) C

Relates to: WAN Capacity Planning for IP Telephony

Q7) A

Relates to: Trunk Capacity Planning

Q8) B, C, E, G

Relates to: WAN Capacity Planning for IP Telephony

Q9) C

Relates to: WAN Capacity Planning for IP Telephony

Q10) A, D, E

Relates to: Campus Capacity Planning for IP Telephony

Simulation 8: Voice Transport over IP Network

Complete this exercise to practice what you learned in this lesson.

This exercise is a paper-only version of a simulation that was actually performed with the simulation tool. The exercise includes the results of that simulation.

Required Resources

These are the resources required to complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- DJMP Industries Case Study Scenario, presented at the end of the module "Applying a Methodology to Network Design"
- A workgroup consisting of two students
- Blank sheets of paper and a pencil

Exercise Objective

In this lesson, you identified the capacity-planning issues in the data network and the impact of voice transport on the rest of the traffic. Upon completing this simulation, you will be able to meet this objective:

- Explain the effect of transporting voice traffic across data networks

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, present in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"

Exercise Procedure

Read the following scenario and try to answer the questions that appear in the text. Discuss possible answers and explain your considerations in the classroom.

Voice Transport over IP Network Scenario

This simulation focuses on the transport of voice over the IP network: VoIP.

DJMP Industries is considering the setup of new international remote offices located in Singapore and London. An Internet-based Virtual Private Network (VPN) will link both international offices with the San Jose headquarters.

Warned about the sensitive nature of voice traffic, the company established and tested a pilot VoIP network in San Jose. It was satisfied with the quality of the voice session but still has doubts about the deployment of voice in a real production network. The major concern is how data traffic will interfere with voice over the Internet.

To help answer the questions that the company still has, you have been asked to analyze the reference model, which simulates VoIP. Your analysis will provide the foundation for a future decision by DJMP.

The figure describes the architecture of the network.

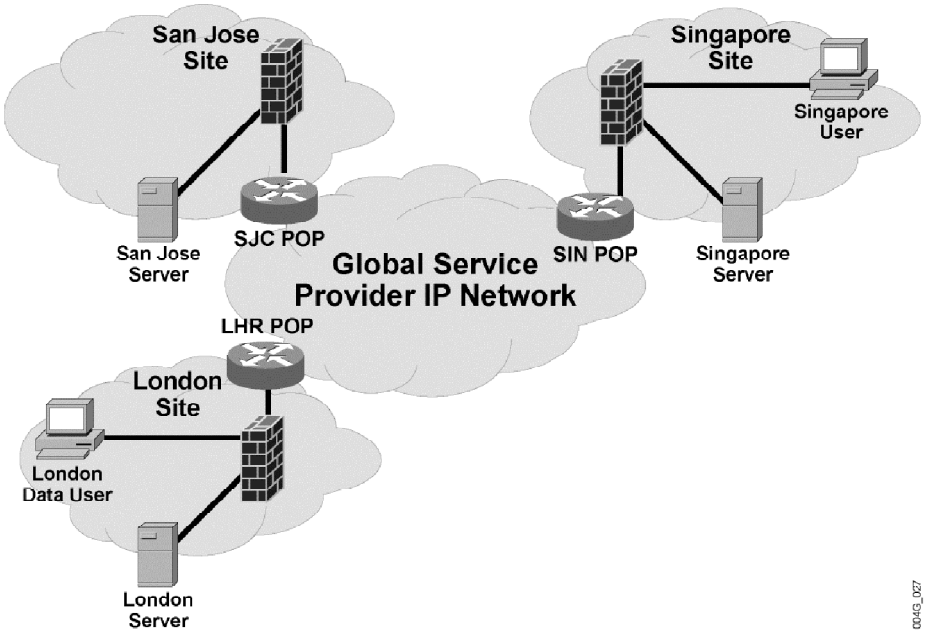

Initially, you decide to simulate the responsiveness and use of the WAN network without any voice traffic by initiating an application that represents a reference traffic flow. Clients in London and Singapore are accessing servers in all three locations. The tested application is FTP.

Testing the Data Load

The following graphs describe the transfer rates resulting from the FTP sessions initiated by the clients. The amount of sent and received traffic is measured in bytes per second.

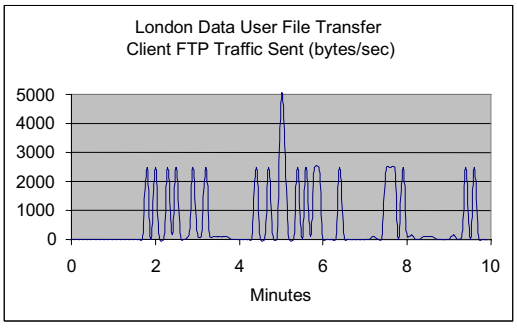

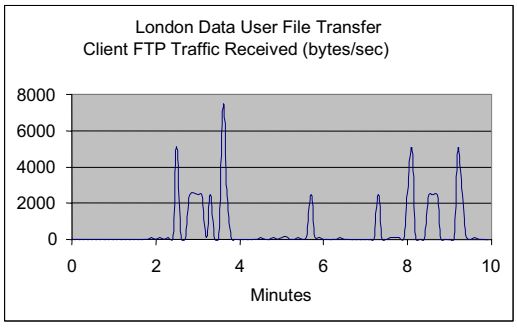

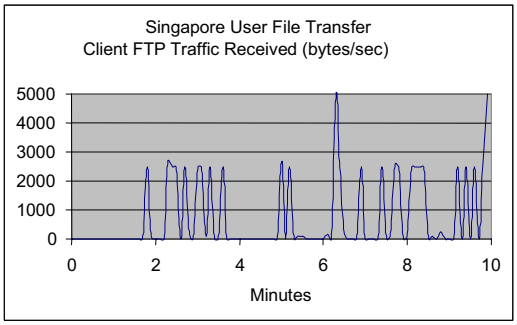

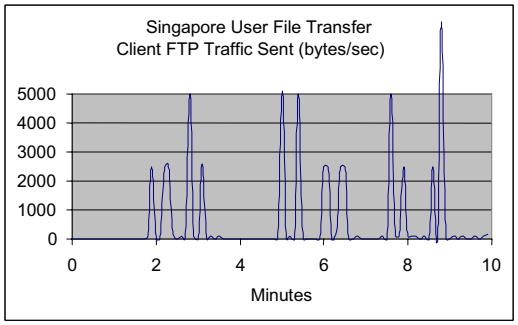

The next two graphs present the probability density of the FTP download response times over a certain period of time. As indicated by the peak values in the graph, the response times are variably distributed within the 7 seconds.

You conclude that most of the responses are received within a 10th of a second and usually between the 3rd and the 6th second.

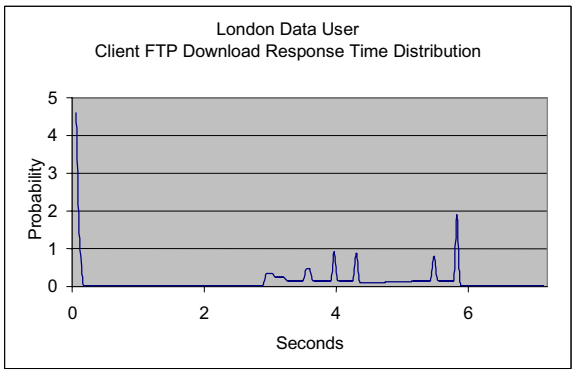

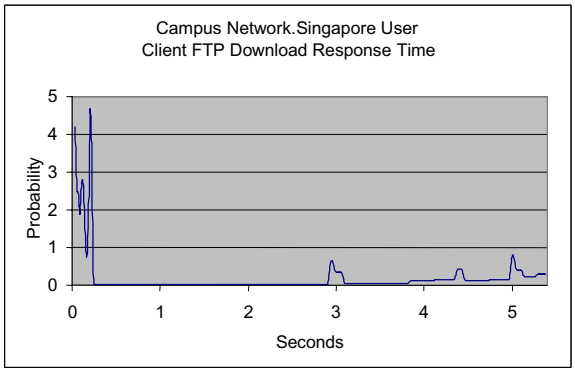

The following graphs show the use of the links as a result of heavy FTP traffic. The load that is placed on the links does not represent a significant burden for the links, but the effect that the data traffic may have on future voice sessions worries you.

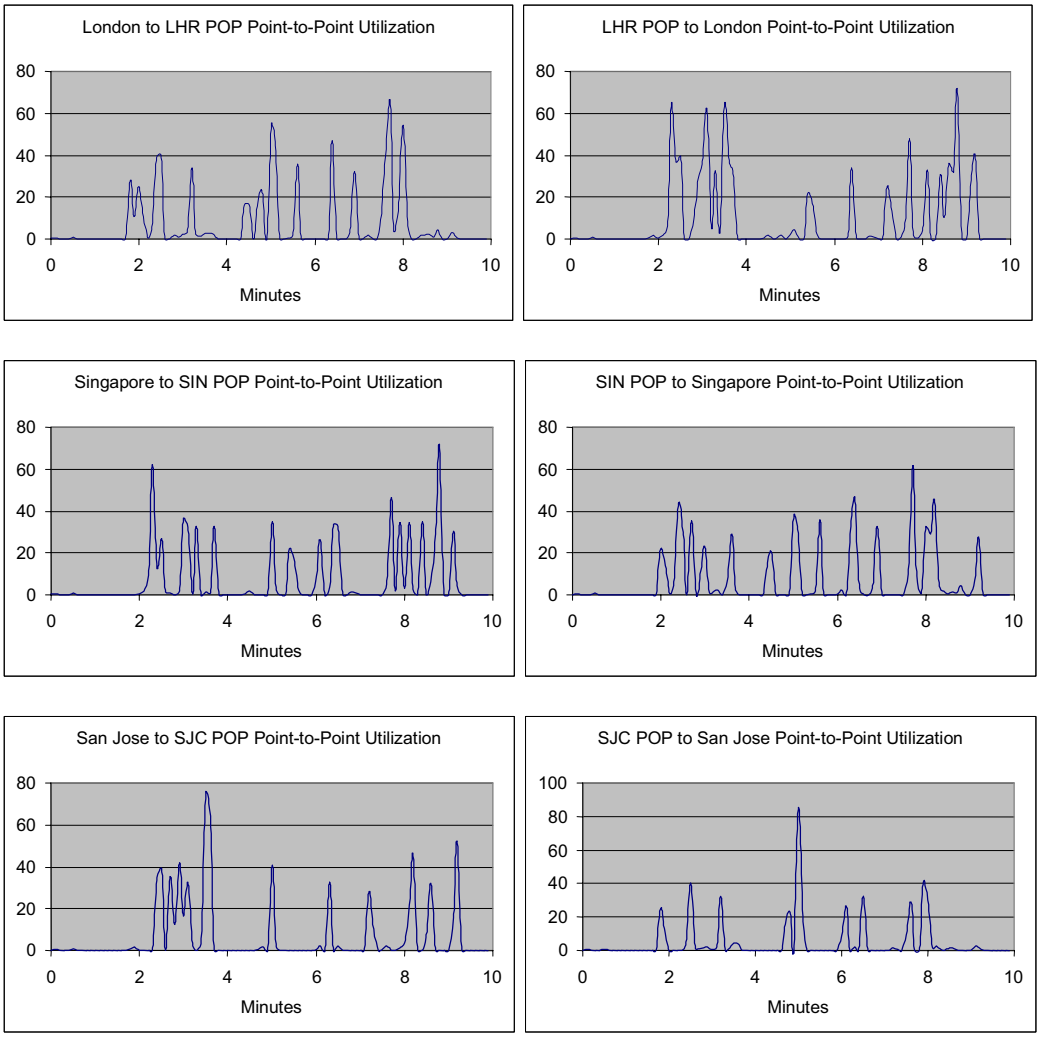

Voice over IP Pilot

The company was very satisfied with the results of its pilot VoIP network configured in San Jose—especially with the quality of sessions established over the weekend in the unloaded network.

Therefore, you decide to obtain comparable results by simulating the VoIP application and analyzing the results in a more realistic scenario—this time between distant locations.

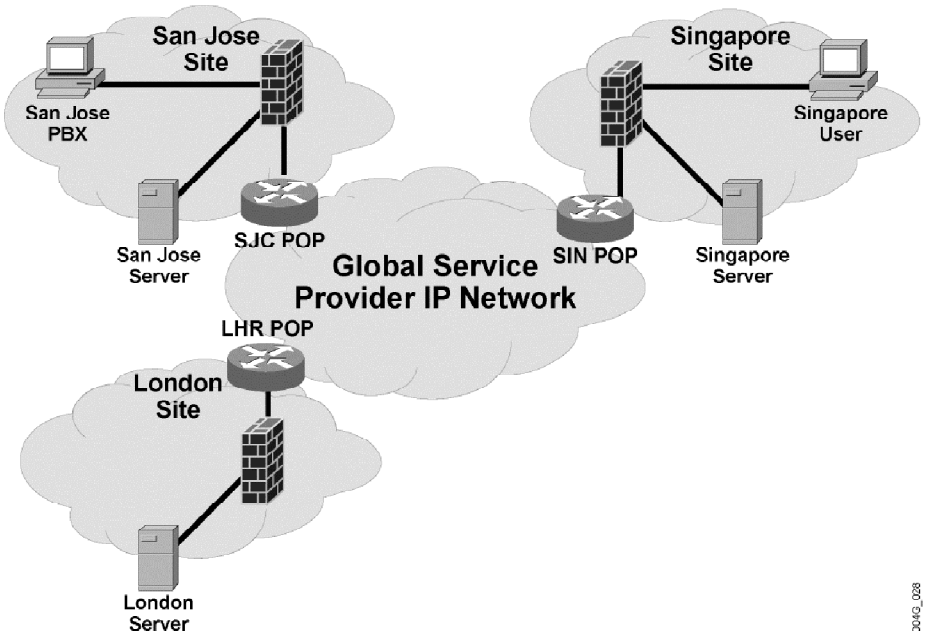

You decide to simulate the VoIP behavior by initiating an application that represents a reference voice traffic flow. The tested application is a VoIP session between the user in Singapore and a PBX switch in San Jose.

Voice Load on the Network

The following graph shows the quality that the company may expect in an unloaded network. The end-to-end delay of voice packets, measured in seconds, appears to be constant, slightly below 20 ms.

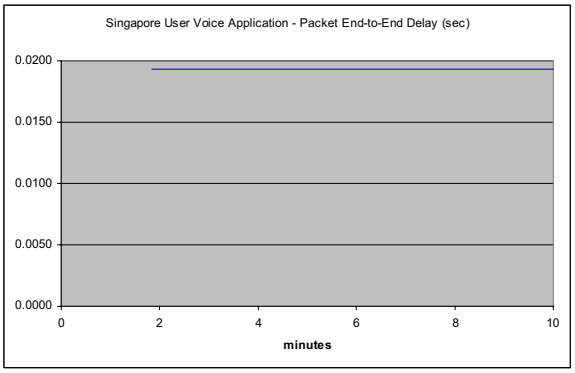

The use of the links that resulted from the initial VoIP traffic is presented in the following graphs. The voice traffic shows approximately 40 percent utilization at a constant level.

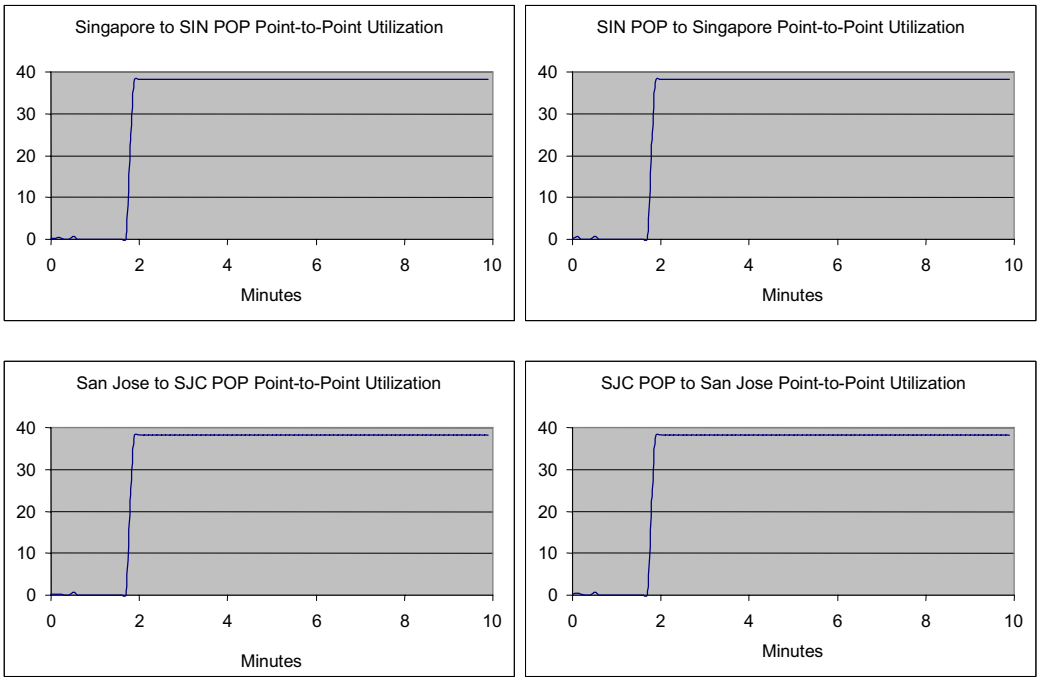

Conclusions

Assume that the voice transport over an unloaded network results in excellent VoIP quality, with no jitter, and that delay is well within acceptable limits. Further assume that the load on the links is not too heavy (40 percent use, but constant and predictable). The maximum load that is acceptable in a voice session is 150 ms.

Voice over IP in Production

When the company puts the VoIP pilot into production, the voice quality degrades dramatically during business hours. You decide to simulate the situation by placing a considerable load on the links and collecting the statistics.

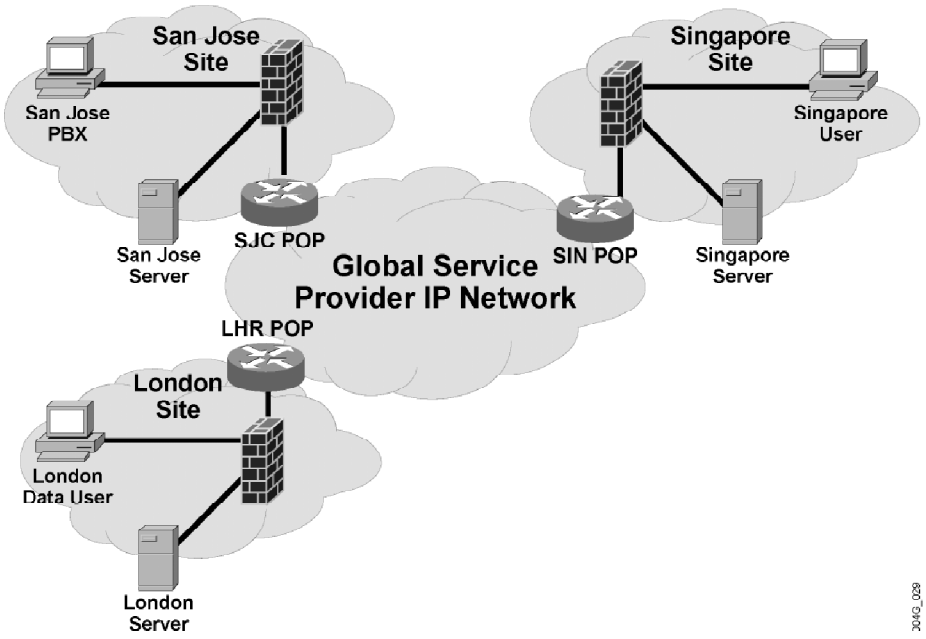

The reference traffic flow in this case incorporates a combination of data and voice traffic. Data users in London and Singapore communicate with servers at all three locations. A voice user in Singapore places a voice call to San Jose.

Testing the Loaded Network

You perform the simulation and produce the following graphs. You compare the data response times with the results from the previous simulation on the unloaded network.

The following graphs describe the use of the links. The load on the links has increased substantially because of the concurrent use of data and voice applications.

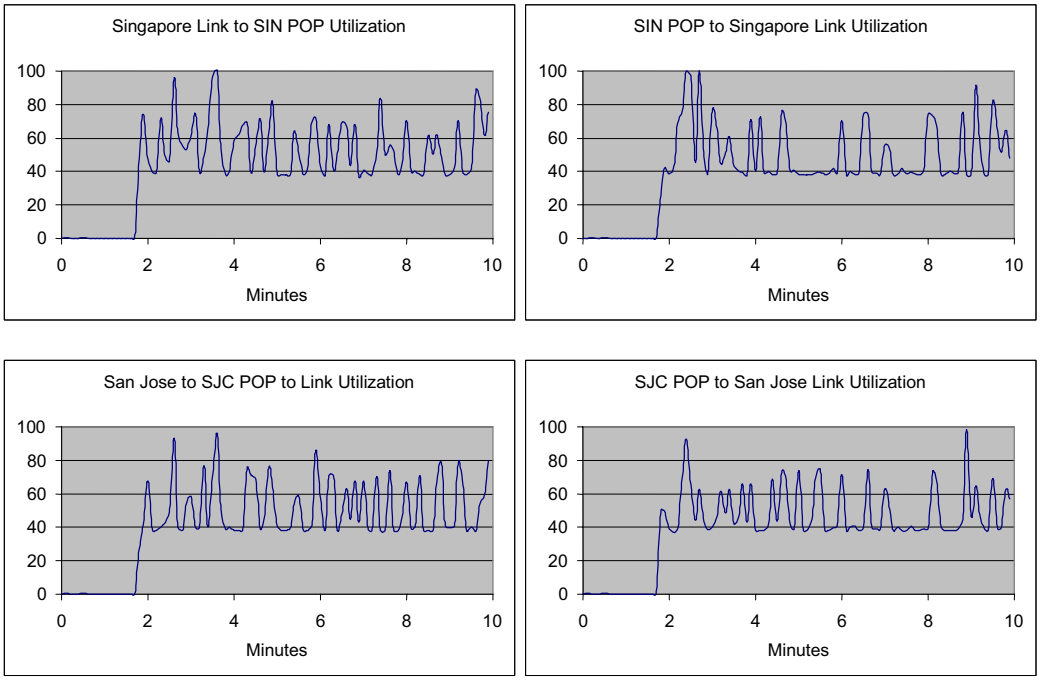

Data Response Times

The following graph compares the distribution of FTP download response times under different conditions. The solid line (blue) in the graph describes the responsiveness of the data network alone, and the dotted line (red) represents the distribution of response times in the combined data and voice network.

The download response time graph shows better response time distribution on the data-only network, where it goes up to 6 seconds. In the data and voice network, the response time distribution goes up to 10 seconds.

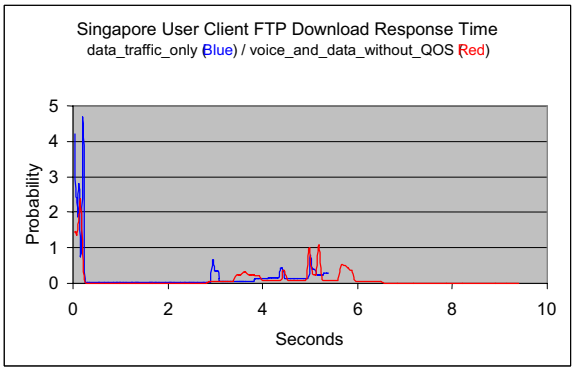

Voice Response Times

The data traffic significantly impacts the quality of service (QoS) provided for the voice session. The end-to-end delays of the voice packets now vary greatly (up to 1.4 seconds), resulting in considerable dissatisfaction with the voice service.

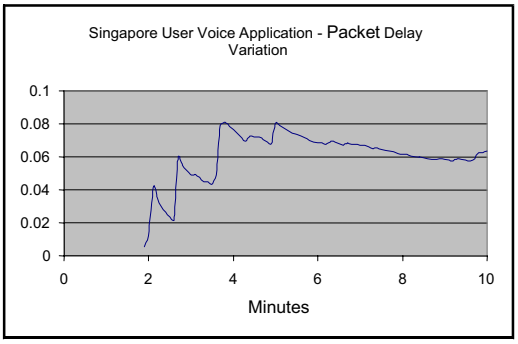

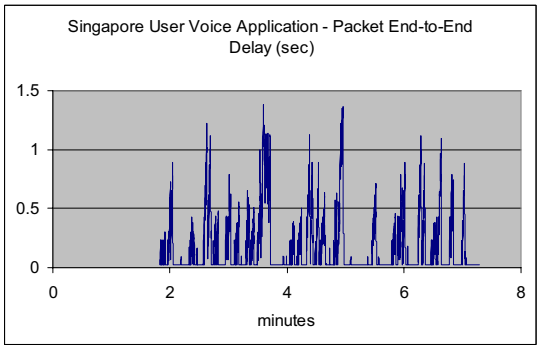

Conclusions

- The voice load placed on international links increases the data application response times. When designing voice and data networks, consider the effects of an increased load on your existing applications.
- The data traffic significantly impacts voice quality and makes the session unusable because of long variable delays. When deploying a converged voice and data network, you must deploy QoS mechanisms to ensure smooth propagation of voice packets.

Q1) Why does the jitter disturb the voice session?

_____

_____

_____

_____

Voice and Data with QoS

In order to improve the quality of the voice, you decide to deploy QoS features. You configure weighted fair queuing (WFQ) on the routers combined with Resource Reservation Protocol (RSVP) on routers and voice stations to support end-to-end bandwidth reservation and preferential forwarding of voice traffic.

The data flow remains as in the previous case. Data users in London and Singapore communicate with servers at all three locations. A voice user in Singapore places a voice call to San Jose.

Data Response Times

The following graphs compare the download response times with QoS (solid red line in the graph) or without QoS (dotted blue line in the graph). It is obvious that QoS significantly reduces the distribution of download response times.

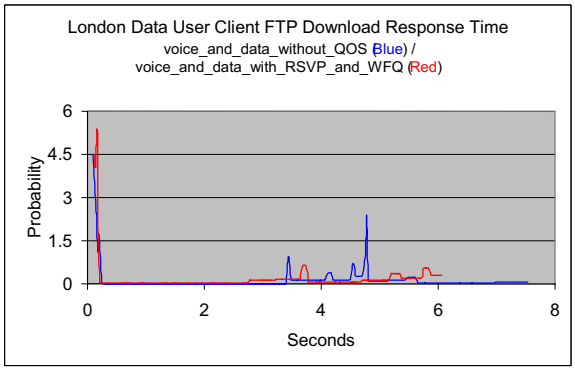

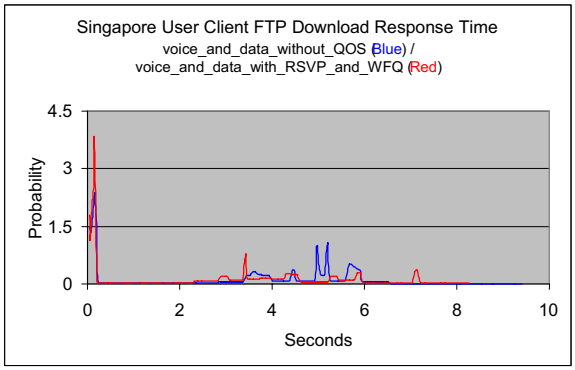

Voice Response Times

There is still some delay variation as seen in the following two graphs. This is mainly attributable to individual voice packets waiting in a queue behind a data packet.

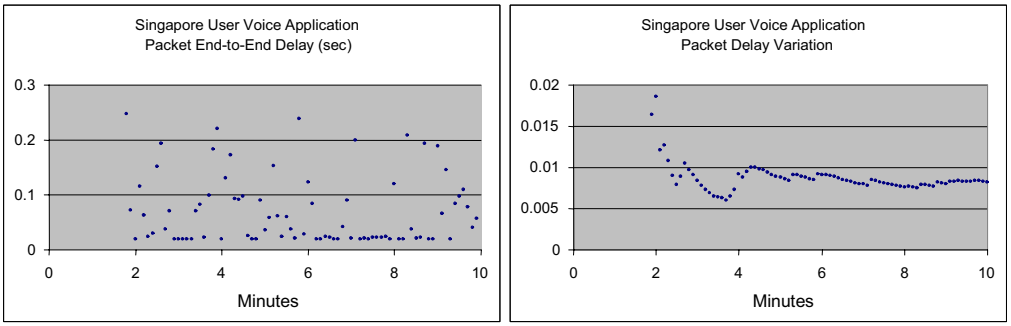

The following graph compares the end-to-end delays without QoS (solid blue line) and with QoS (red pixels). Deploying QoS makes voice perform almost as if there is no data traffic.

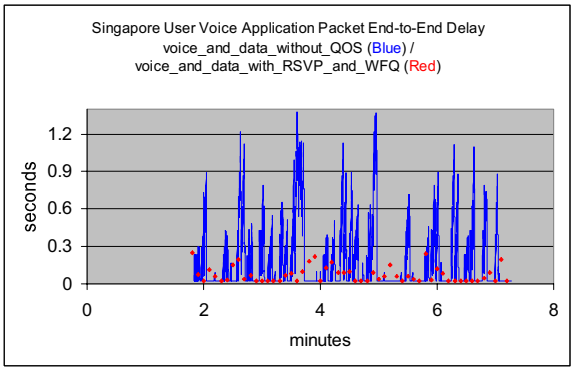

Conclusions

- QoS, when properly deployed, can significantly increase the voice quality in mixed voice and data networks. When designing voice and data networks over low-speed links or on congested networks, always plan on deploying QoS.
- QoS mechanisms improve the quality of voice sessions.

Q2) How can QoS mechanisms improve data propagation in congested networks?

_____

_____

_____

_____

Module 9

Applying Basic Network Management Design Concepts

Overview

Proper management is critical to an efficiently run network. Network management enables you to detect faults, monitor performance, track configuration changes in the network, and provide security and accounting management for both individual and group network resource usage. Controlled changes, such as configuration modifications, software updates, and cabling changes, as well as unexpected behaviors or failures, have an effect on a network. To maintain control, tools are required for tracking and monitoring all activity that may affect network performance.

This module introduces the protocols for managing networks. The functional areas of network management are also explained in detail. The module concludes with an explanation of the service levels that help guarantee a specified level of service.

Module Objectives

Upon completing this module, you will be able to describe basic network management design concepts.

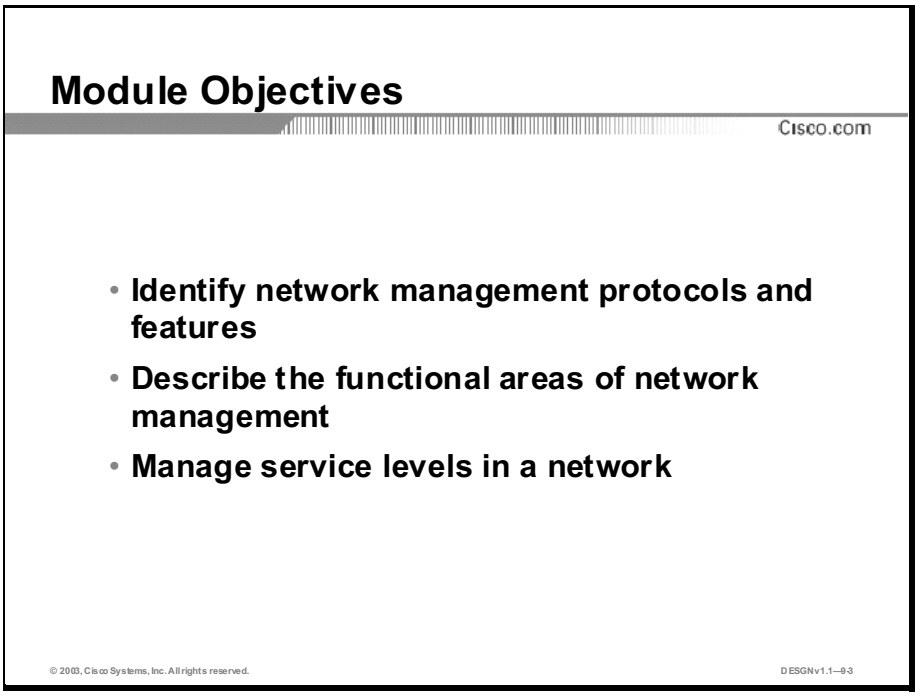

Module Outline

The outline lists the components of this module.

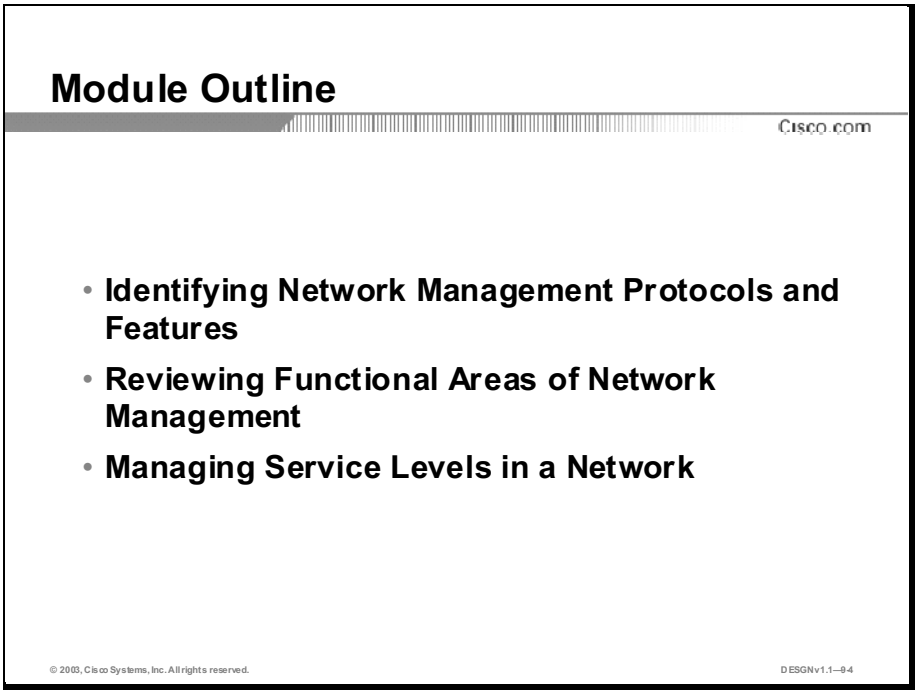

Identifying Network Management Protocols and Features

Overview

Network administrators need tools for monitoring the functionality of network devices, connections, and services. The Simple Network Management Protocol (SNMP) has become the standard for use in network management solutions, along with Remote Monitoring (RMON) and MIBs. Each managed device in the network has several variables that quantify the state of the device. By reading the values of these variables, you can monitor managed devices, and by writing values into these variables, you can control the managed devices.

This lesson introduces management protocols, describing the differences between SNMP versions 1, 2, and 3. The lesson also describes the role of MIBs in SNMP and RMON. It introduces Cisco Discovery Protocol (CDP), explaining its benefits and limitations. The lesson describes methods for gathering network statistics, messages, and alerts using network flow (NetFlow) and syslog.

Relevance

This lesson describes basic terms and functionality used in network management that you will need to design a network management solution.

Objectives

Upon completing this lesson, you will be able to identify network management protocols and features. This includes being able to meet these objectives:

- Explain the difference between SNMP polling and SNMP traps
- Describe a MIB and the information it stores
- Describe the benefits of using RMON over SNMP
- Explain when to use NetFlow
- Describe the primary role of CDP and list the devices that support it
- Explain when and why to use syslog

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of networking concepts

Outline

The outline lists the topics included in this lesson.

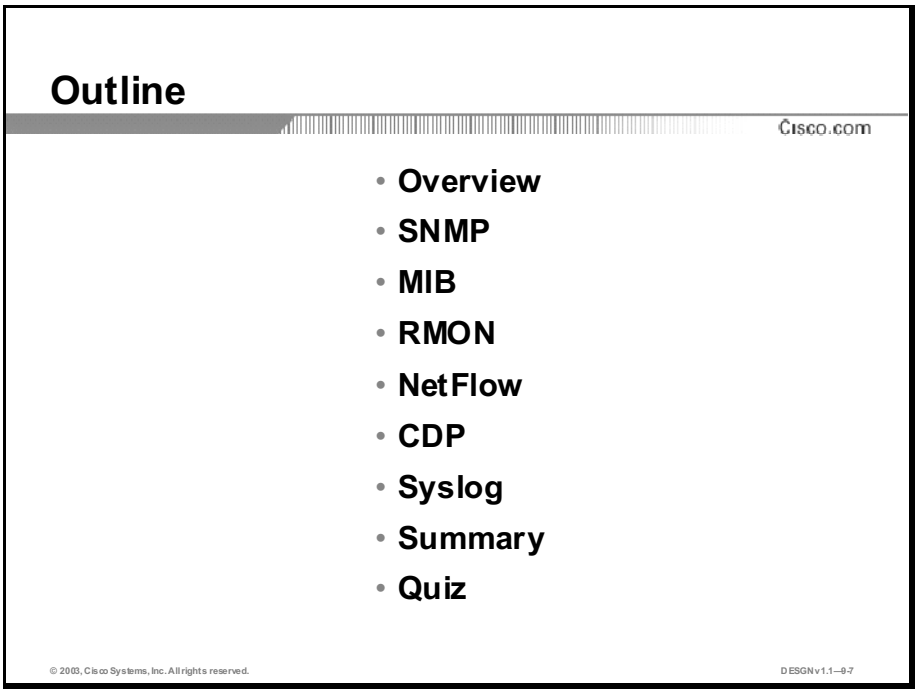

SNMP

SNMP has become the standard for network management. It is a simple solution that requires little code to implement, thus enabling vendors to easily build SNMP agents for their products. Therefore, SNMP is often the foundation of a network management architecture. This topic describes SNMP.

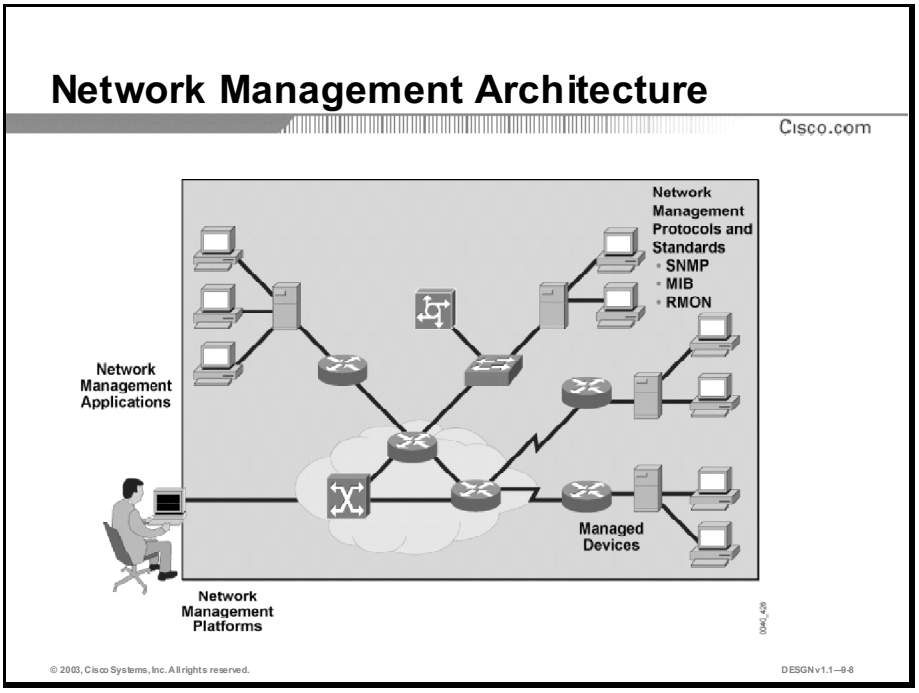

A network management architecture consists of these elements**:**

- **Network management systems:** A network management system (NMS) executes applications that monitor and control managed devices. NMSs provide the bulk of the processing and memory resources required for network management.
- **Network management protocols and standards that facilitate the exchange of management information between the NMS and managed devices:** The key network management protocols and standards are:
  - SNMP is, as its name implies, a simple network management protocol. An SNMP agent stores data specific to the managed device in a MIB.
  - A MIB is a detailed definition of the information on a network device accessible through a network management protocol like SNMP.
  - RMON is an extension to the standard MIB. RMON provides remote monitoring capability through the collection of network traffic data on remote links. The RMON agent resides on a managed device and collects specific groups of statistics, which an NMS can retrieve and use for long-term trend analysis.
- **Managed devices:** Those devices that are monitored and controlled by the NMS.
- **Management agents:** These typically reside on managed devices and include SNMP agents and RMON agents.
- **Management information:** This is commonly stored in MIBs.

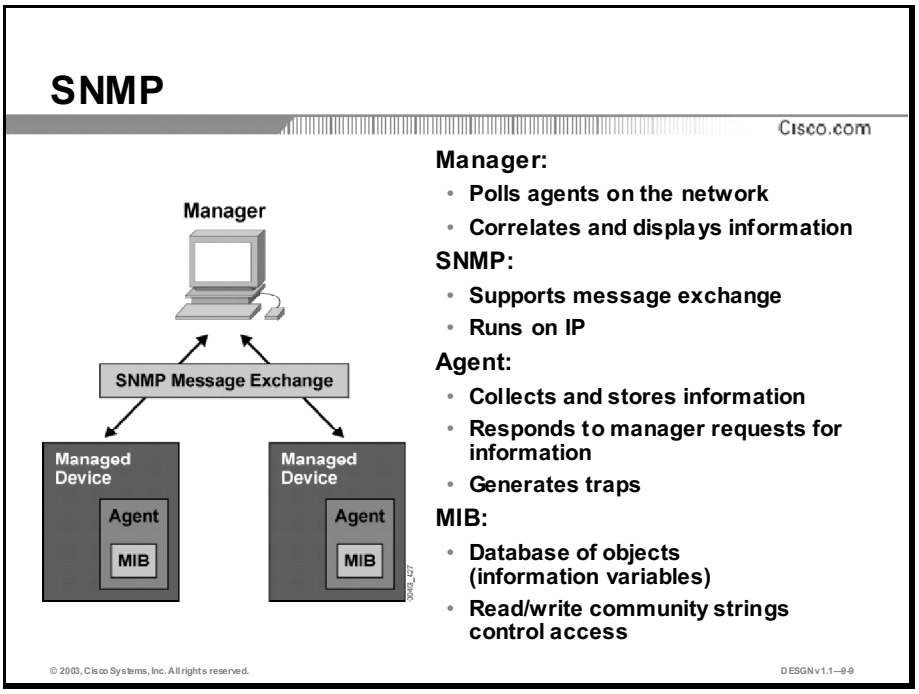

SNMP defines how management information is exchanged between network management applications and management agents. A network management application periodically polls the SNMP agents residing on managed devices and collects the data. A network management application can display the information in a GUI on the network manager.

SNMP uses the User Datagram Protocol (UDP) transport mechanism of IP to retrieve and send management information, such as MIB variables.

SNMP management agents that reside on managed devices collect and store information about the device and its operation, respond to manager's requests, and generate traps to inform the manager of certain events.

The management agent collects data and stores it locally in the MIB. Community strings control access to the MIB. To view or set MIB variables, the user must specify the appropriate community string for read or write access.

The initial version of the SNMP standard (SNMPv1) is defined in Request for Comments (RFC) 1157.

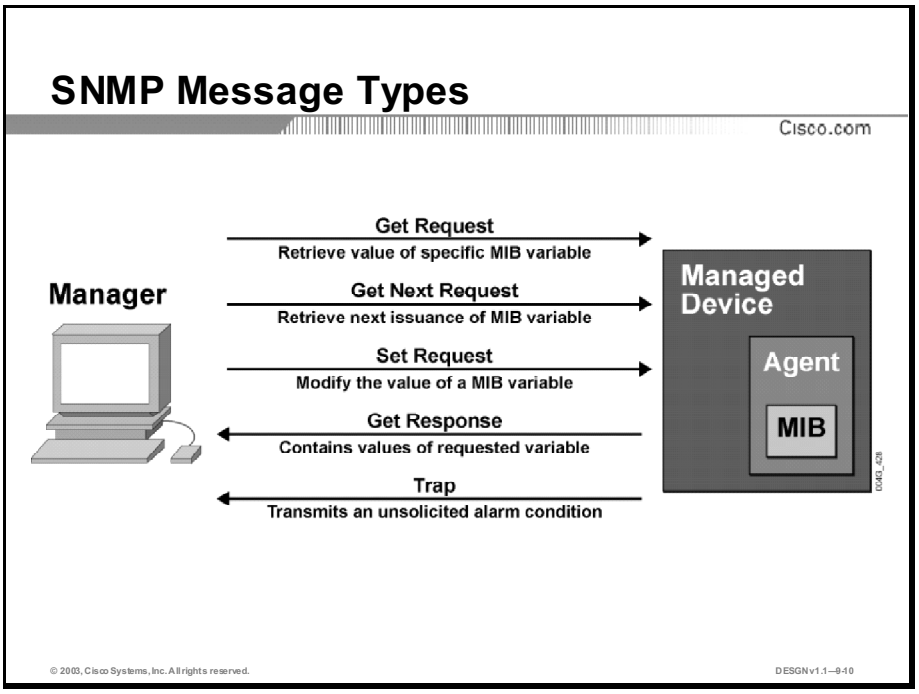

These are the basic SNMP messages that the network manager uses to transfer data from agents that reside on managed devices:

- **Get Request:** Used to request the value of a specific MIB variable from the agent.
- **GetNext Request:** Used after the initial get request to retrieve the next object instance from a table or a list.
- **Set Request:** Used to set a MIB variable on an agent.
- **Get Response:** Used by an agent to respond to a get request or get next request from a manager.
- **Trap:** Used by an agent to transmit an unsolicited alarm to the manager. An agent sends a trap message when a certain condition occurs, such as a change in the state of a device, device or component failure, or an agent initialization or restart.

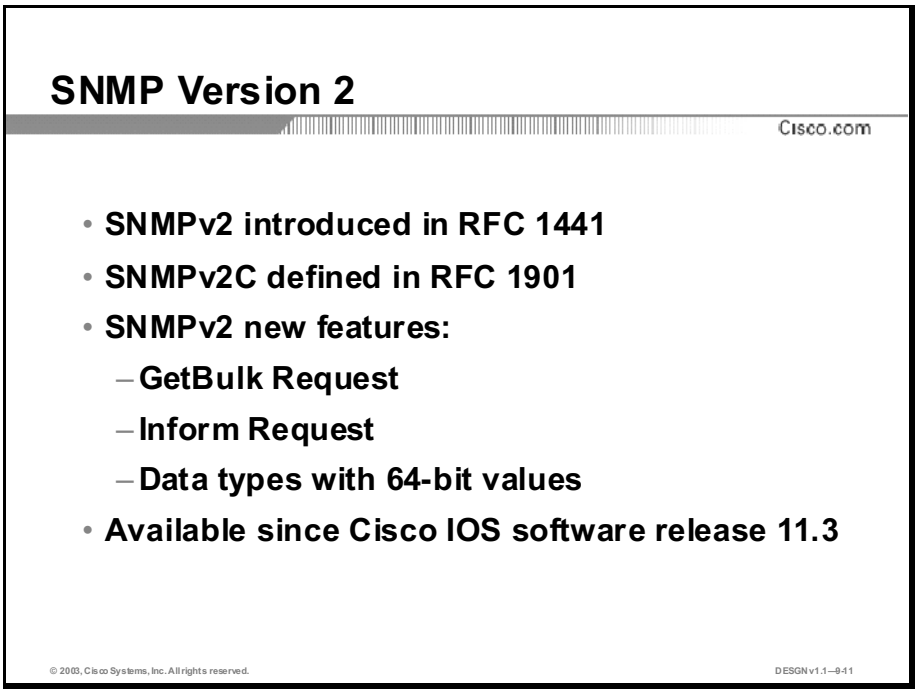

SNMP version 2 (SNMPv2) provides improved performance, security, confidentiality, and manager-to-manager communications.

SNMPv2 was introduced with RFC 1441, but members of the Internet Engineering Task Force (IETF) subcommittee could not agree on the security and administrative sections of the SNMPv2 specification. There were several attempts to achieve acceptance of SNMPv2 through the release of experimental modified versions.

Community-based SNMPv2 (SNMPv2C), defined in RFC 1901, is the most common implementation. SNMPv2C deploys the same administrative framework as defined in SNMPv1, which uses read and write community strings for administrative access.

SNMPv2 introduces two new message types:

- **GetBulk request:** Reduces repetitive requests and replies, improving the performance when retrieving large amounts of data (for example, tables).
- **Inform request:** Inform request messages alert an SNMP manager of specific conditions. Unlike SNMP trap messages, which are unconfirmed, the NMS acknowledges an Inform request by sending an Inform response message back to the requesting device.

SNMPv2 adds new data types with 64-bit counters, because 32-bit counters were quickly outmoded by fast network interfaces.

On Cisco routers, SNMPv2 is implemented in Cisco IOS software release 11.3 and later.

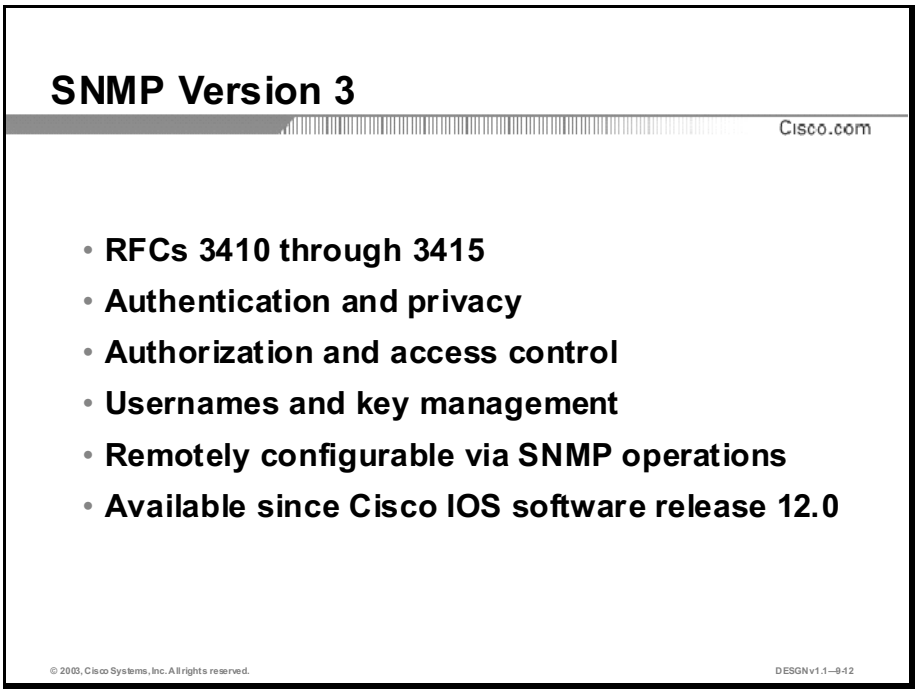

SNMP version 3 (SNMPv3) is the latest SNMP version to become a full standard. SNMPv3, described in RFCs 3410 through 3415, adds methods to ensure the secure transmission of critical data between managed devices.

SNMPv3 introduces three levels of security:

- **noAuthNoPriv:** No authentication is required and no privacy (encryption) is provided.
- **authNoPriv:** Authentication is based on Hash-based Message Authentication Code with Message Digest 5 (HMAC-MD5) or Hash-based Message Authentication Code with Secure Hash Algorithm (HMAC-SHA). No encryption is provided.
- **authPriv:** In addition to authentication, Cipher Block Chaining-Data Encryption Standard (CBC-DES) encryption is used as the privacy protocol.

Security levels implemented for each security model determine which SNMP objects a user can access for reading, writing, or creating, and the list of notifications its users can receive.

On Cisco routers, SNMPv3 is implemented in IOS software release 12.0 and later.

MIB

A MIB stores management information for the local agent on the managed device. This topic describes MIBs.

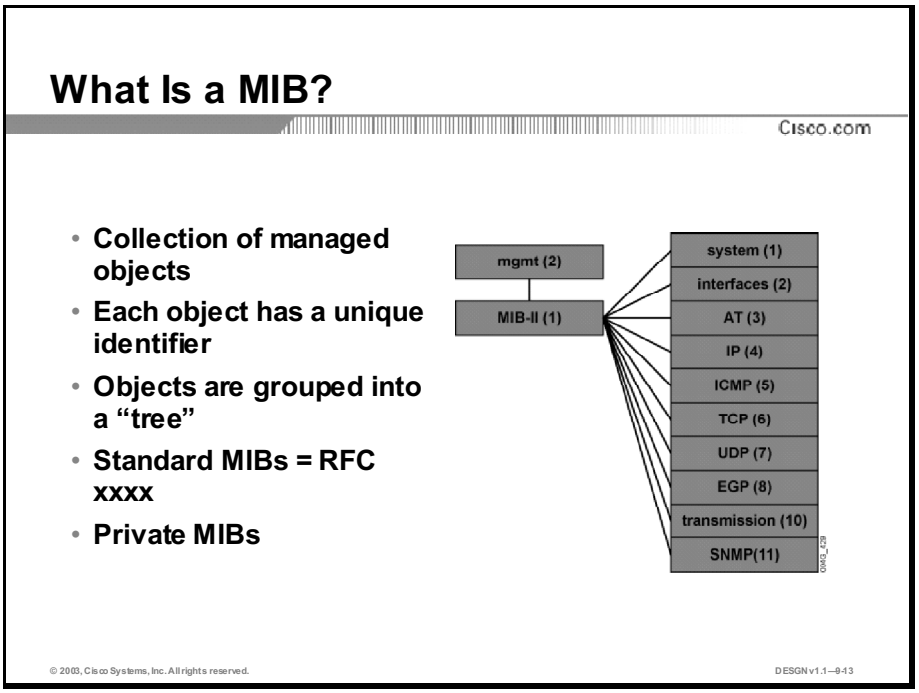

Each object in a MIB has a unique identifier which network management applications use to identify and retrieve a specific object's value. The MIB structure is a tree-like structure. Similar objects are grouped under the same branch of the MIB tree. For example, different interface counters are grouped under the interfaces branch of the MIB tree.

Standard MIBs are defined in different RFCs. RFC 1213 defines the TCP/IP MIB, RFC 1231 defines the Token Ring MIB, RFC 1243 defines the AppleTalk MIB, and so on.

In addition to standard MIBs, there are private or vendor-specific MIB definitions. Vendors can obtain their own branch for the definition of their private MIB subtree and create custom managed objects under that branch.

Example: A Router MIB

A Cisco router has a number of standard managed objects defined in the standard section of the MIB tree such as these:

- Interfaces
- Buffers
- Memory
- Standard protocols

The router has private managed objects introduced by Cisco in the private section of the MIB tree such as these:

- Small, medium, large, and huge buffers
- Primary and secondary memory
- Proprietary protocols

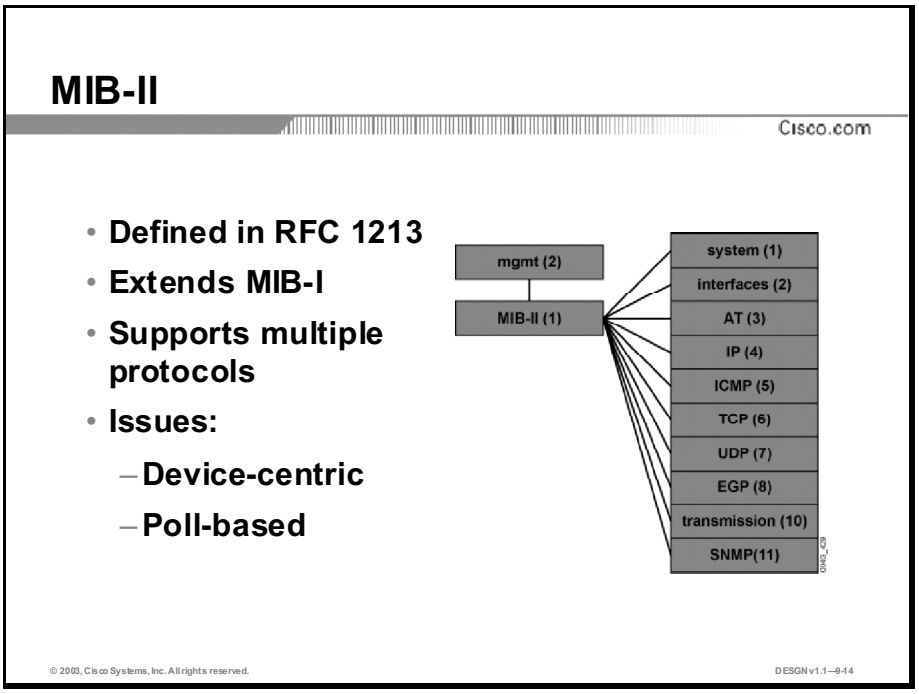

MIB-II, defined by RFC 1213, is an extension of MIB-I. MIB-II supports new protocols and provides more detailed and structured information. It remains compatible with the previous version, which is why MIB-II retains the same object identifier as MIB-I (1.3.6.1.2.1).

The location of MIB-II objects is under the iso.org.dod.internet.mgmt (1.3.6.1.2) subtree, where the top-level MIB objects are defined:

- system (1)
- interfaces (2)
- AT (3)
- IP (4)
- ICMP (5)
- TCP (6)
- UDP (7)
- EGP (8)
- transmission (10)
- SNMP (11)

Although the MIB-II definition is an improvement over MIB-I, unresolved issues exist. For example, MIB-II is still a device-centric solution and is poll-based.

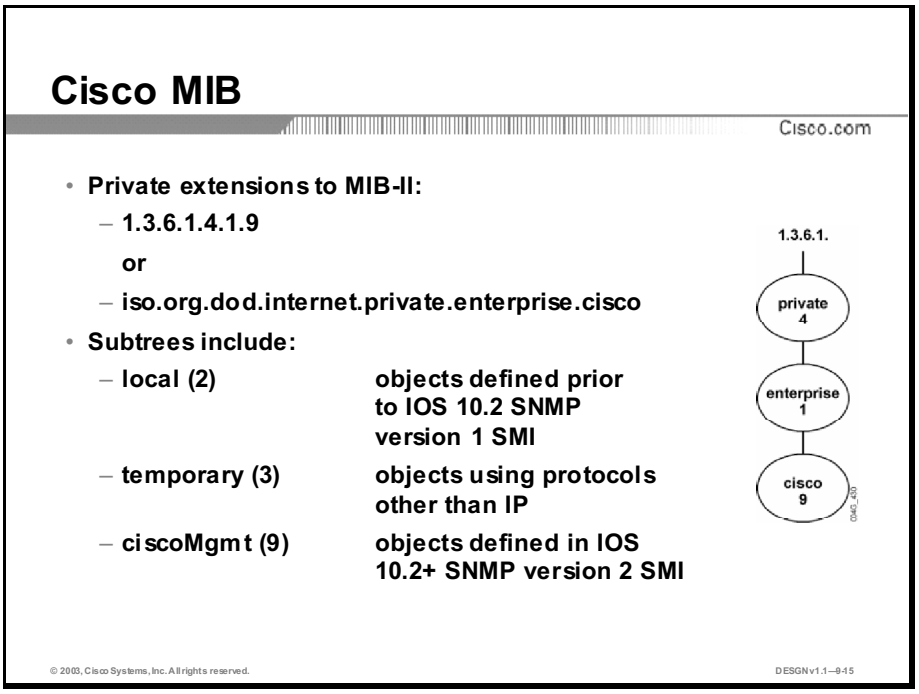

To use the private definitions of managed objects, the administrator must import the private definitions into the NMS. This process is useful for operators because the resulting outputs are descriptive variable names. The Cisco private MIB tree contains three subtrees:

- local (2)
- temporary (3)
- ciscoMgmt (9)

The local (2) subtree contains MIB objects defined before IOS software release 10.2, which implemented the SNMPv1 Structure of Management Information (SMI). SMI defines the structure of data residing within MIB-managed objects. Beginning with IOS software release 10.2, however, Cisco MIBs are defined according to SNMPv2 SMI.

MIBs defined with SNMPv2 SMI are placed in the ciscoMgmt subtree. Cisco is phasing out MIBs currently defined in the local subtree, replacing them with new objects defined in the ciscoMgmt subtree.

Cisco maintains its private MIB definitions under the Cisco MIB subtree (1.3.6.1.4.1.9). You can obtain Cisco MIB definitions that Cisco devices support at http://www.cisco.com/public/mibs.

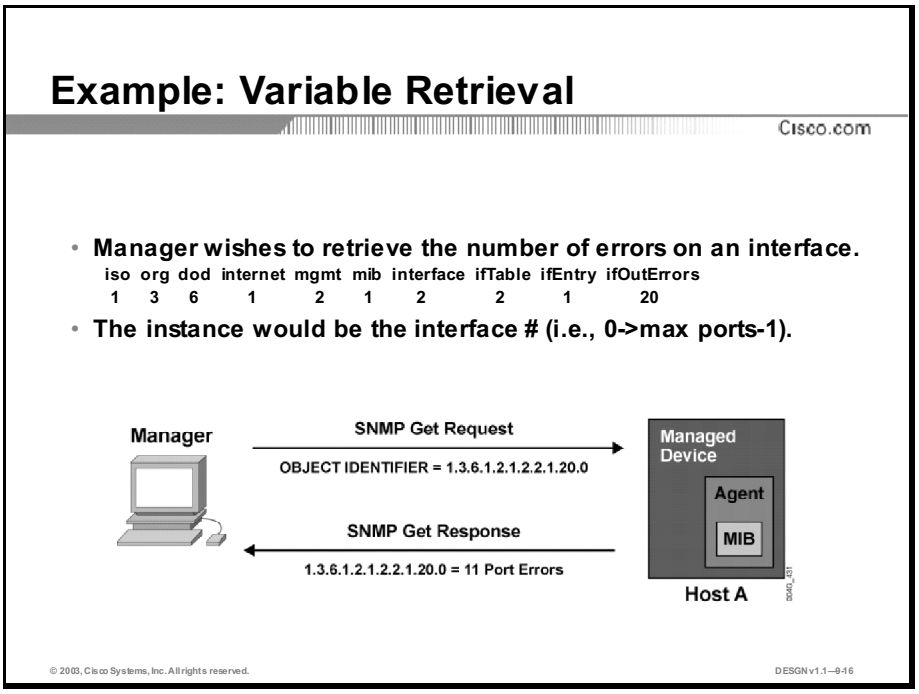

The figure shows an SNMP MIB variable retrieval in action.

The network manager wishes to retrieve the number of errors on the first interface. Starting with interface number 0, the valid range for interface numbers is 0 through "maximum ports minus one". Therefore, the manager creates the SNMP Get Request message with reference to the MIB variable 1.3.6.1.2.1.2.2.1.20.0, representing outgoing errors on interface 0.

The agent creates the SNMP Get Response message as a response to the manager's request. The referenced variable is included in the response. In the example, the agent returned 11, indicating that there were 11 outgoing errors on that interface.

Monitoring networks using SNMP requires that the NMS poll each managed device on a periodic basis to determine its status. Frequent polling of many devices or MIB variables on a device across a network to a central NMS may result in performance issues. Performance issues include congestion on slower links or at the NMS connection, or an overwhelming of NMS resources to adequately process all of the collected data. Therefore, some important polling guidelines include these:

- Restrict polling to only those MIB variables necessary for analysis.
- Increase polling intervals (reduce number of polls per period) over low-bandwidth links.
- For larger networks, consider the deployment of management domains, or a distributed model for deploying NMSs. Management domains permit polling to be more local to the managed devices, thus reducing overall management traffic across the network and the potential of one failed device or link from cutting off management visibility to the remaining network. Aggregated management data may still be centralized. This model is particularly appropriate for networks that already have separate administrative domains, or where large campuses or portions of the network are separated by slower WAN links.
- Analyze and use the data collected. Do not collect data if it is not analyzed.
- Leverage non-polling mechanisms such as SNMP traps, RMON, and syslog.

RMON

RMON is a MIB that supports proactive network management of remote networks. This topic describes RMON.

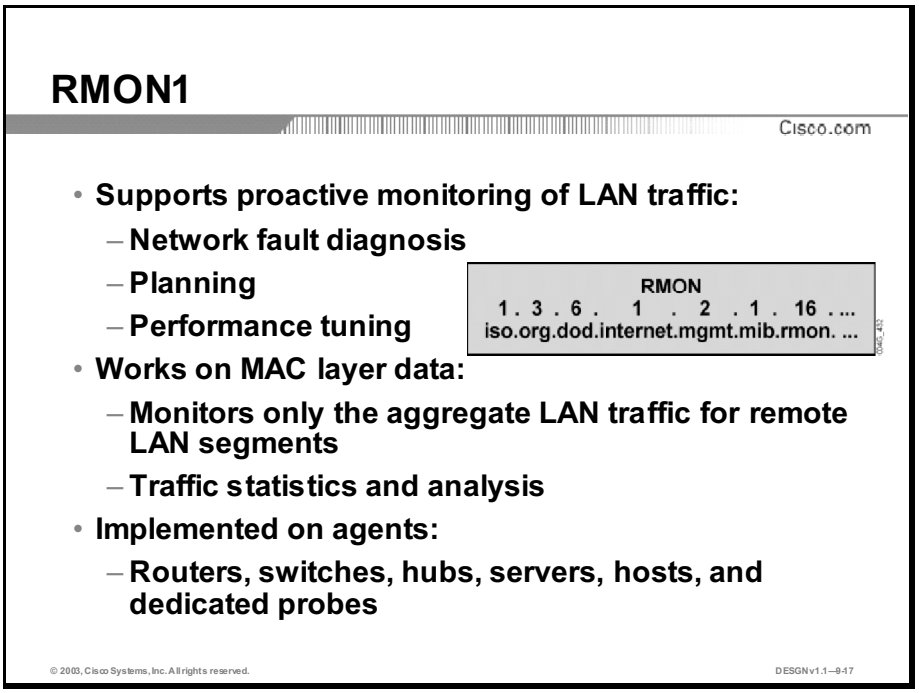

Using RMON, the managed device itself collects and stores MIB data.

RMON can set performance or error thresholds and only report if the threshold is breached, which helps to reduce management traffic. RMON provides effective network fault diagnosis, performance tuning, and planning for network upgrades.

RMON1 works on MAC-layer data and provides the aggregate LAN traffic, statistics, and analysis for remote LAN segments.

RMON agents must look at every frame on the network. Therefore, they can cause performance problems on a managed device with insufficient processing power and memory.

RMON agents can reside in routers, switches, and dedicated RMON probes. Because packet processing may become resource intensive and the amount of data that the RMON agent collects can be very large, network managers often deploy dedicated RMON probes instead of enabling RMON agents on routers and switches.

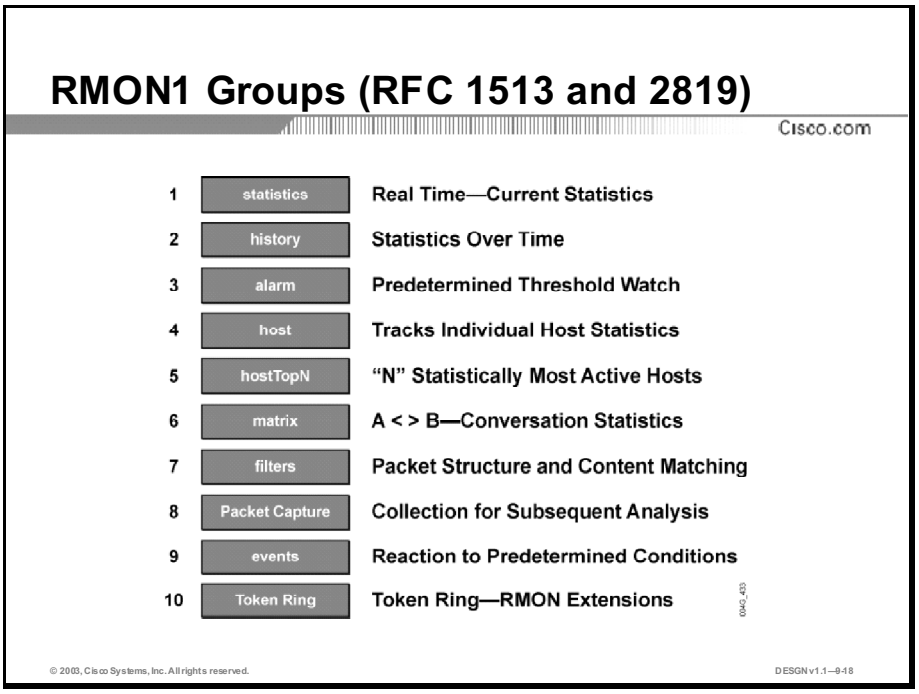

The RMON1 agents gather nine groups of statistics (10 with Token Ring). The agents then forward this information to a manager upon request, commonly through SNMP. These are RMON1 groups:

- **Statistics:** Contains statistics such as packets sent, bytes sent, broadcast packets, multicast packets, CRC errors, runts, giants, fragments, jabbers, collisions, and so on for each monitored interface on the device
- **History:** Used to store periodic statistical samples for later retrieval
- **Alarm:** Used to set specific thresholds for managed objects and to trigger an event upon crossing the threshold (requires an events group)
- **Host:** Contains statistics associated with each host discovered on the network
- **Host top N:** Contains statistics for hosts that top a list for a specific observed variable
- **Matrix:** Contains statistics for conversations between sets of two addresses (packets or bytes exchanged between two hosts)
- **Filters:** Contains rules for data packet filters, which generate events or are stored locally in a Packet Capture group
- **Packet Capture:** Contains data packets that matched rules set in the filters group
- **Events:** Controls the generation and notification of events from this device
- **Token Ring:** Contains Token Ring extensions:
  - **Ring station:** Provides detailed statistics on individual stations
  - **Ring station order:** An ordered list of stations currently on the ring
  - **Ring station configuration:** Configuration and insert/removal data on each station
  - **Source routing:** Statistics on source routing, such as hop counts

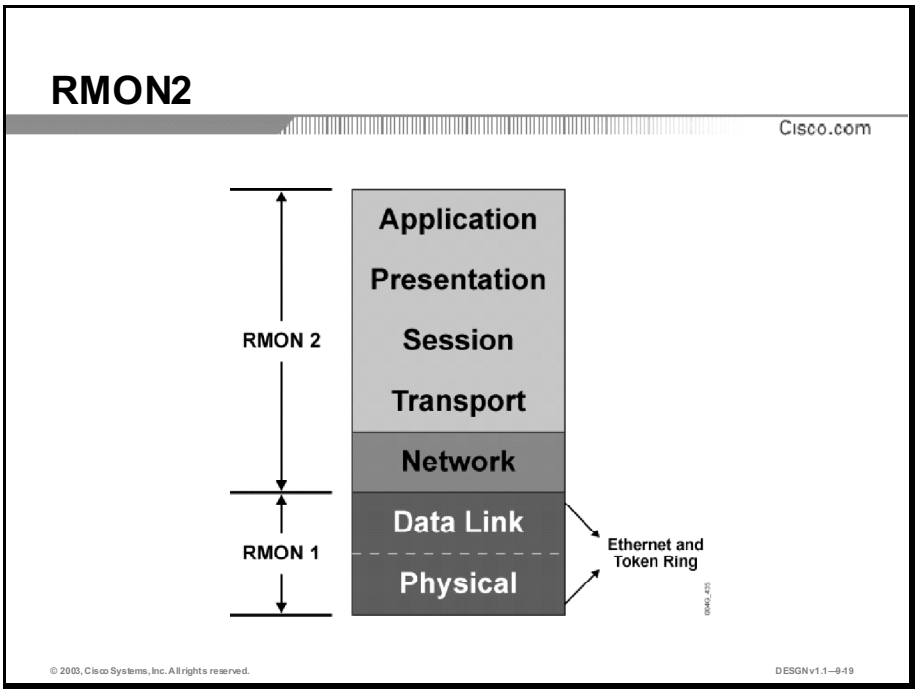

RMON1 only provides visibility into the data link and physical layers. Performance and fault analysis of the higher layers still require other capture and decode tools. Because of the limitations of RMON1, RMON2 was developed to extend functionality to upper-layer protocols. RMON2 provides full network visibility from the network layer to the application layer.

RMON2 is not a replacement for RMON1, but an extension to it. RMON2 extends RMON1 by adding nine more groups that provide visibility to the upper layers.

The visibility of upper-layer protocols enables the network manager to monitor any upper-layer protocol traffic for any device or subnet, in addition to the MAC-layer traffic. RMON2 also provides end-to-end views of network conversations per protocol.

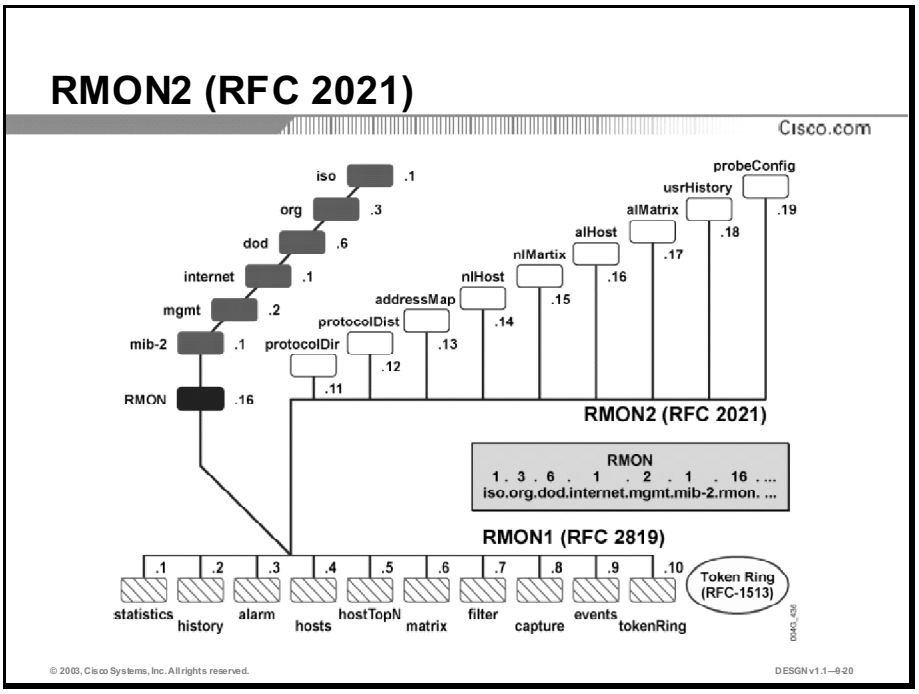

RMON2 supports the following new groups:

- **Protocol directory:** Provides the list of protocols that the device supports
- **Protocol distribution:** Contains the traffic statistics for each supported protocol
- **Address mapping:** Contains network layer to MAC layer address mappings
- **Network layer host:** Contains statistics for the network layer traffic to or from each host
- **Network layer matrix:** Contains network layer traffic statistics for conversations between pairs of hosts
- **Application layer host:** Contains statistics for the application layer traffic to or from each host
- **Application layer matrix:** Contains application layer traffic statistics for conversations between pairs of hosts
- **User history collection:** Contains periodic samples of user-specified variables
- **Probe configuration:** Provides a standard way to remotely configure probe parameters such as trap destination and out-of-band management

RMON2 collects statistics beyond the MAC layer of a specific segment. Network managers can view conversations at the network and application layers, including traffic generated by a specific host or even a specific application on that host.

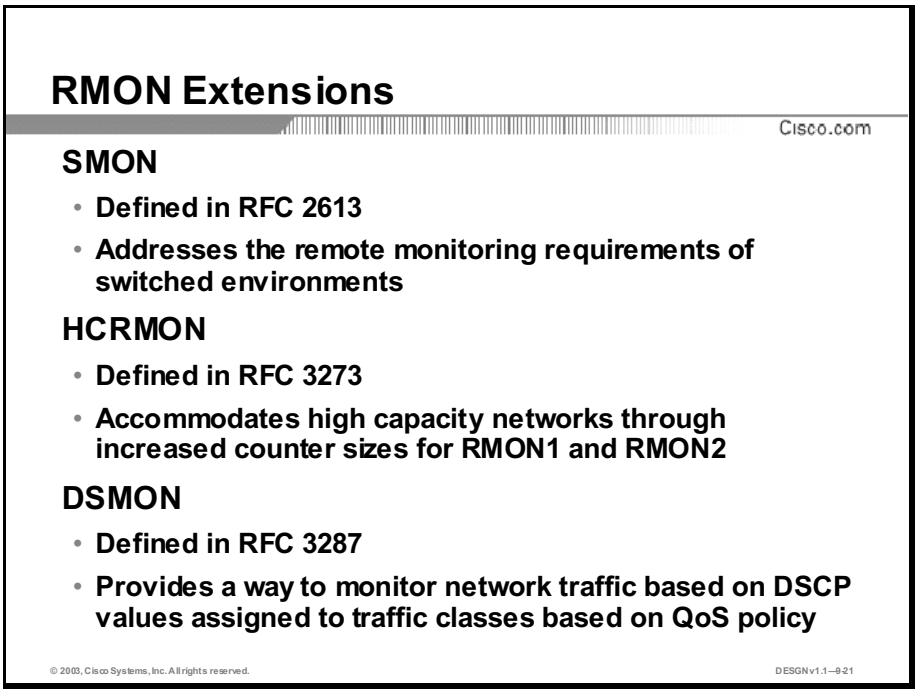

Switched networks present a unique challenge to RMON since monitoring for each agent is limited to a specific link, and switches typically concentrate many Ethernet links. To overcome this challenge, Cisco Catalyst switches embed RMON agent functionality at the port level and use application-specific integrated circuits (ASICs) specifically designed to process packets for RMON data analysis, thus offloading the resource requirements from the main CPU. Due to processing constraints, this embedded functionality is typically limited to four groups of RMON1 (Statistics, History, Alarms, and Events). This subset functionality is sometimes referred to as mini-RMON or RMON-Lite.

To monitor additional RMON1 and RMON2 groups, a dedicated RMON probe is often required, either as an external appliance or as an integrated module. The Catalyst 6500 switch, for example, supports a dedicated Network Analysis Module (NAM) that integrates directly into the switch fabric to provide high-performance RMON1/2 analysis for the entire switch.

Whether deployed as an external dedicated appliance or an integrated module, a Switched Port Analyzer (SPAN) on the switch needs to copy or mirror switch traffic on specific ports or virtual LANs (VLANs) to the RMON probe port for analysis. There are often restrictions on the number of ports or VLANs that the RMON probe port can monitor simultaneously based on the capacity of the probe, its port connection, or SPAN functionality.

When designing a network management solution using RMON, consider the network monitoring requirements, the type of traffic to analyze, the amount and type of RMON data to collect, and the performance characteristics of the RMON agent. This information will help you determine which network ports require RMON monitoring capability, which ports to enable for RMON and the level of analysis required, and where to position dedicated probes for advanced, high-performance monitoring of critical links.

Network managers need to carefully decide which agent functionality and groups to enable to reduce unnecessary RMON processing and device resources. At the same time, proper deployment and tuning of RMON thresholds (Alarms and Events) on critical links can greatly reduce the management traffic normally associated with polling mechanisms such as SNMP while providing a comprehensive alert notification solution if network utilization or error

conditions exceed preset limits. A common practice is to leverage embedded RMON agents for basic network monitoring, fault, and trending analysis, and to drill down using dedicated RMON devices for advanced performance, packet capture, and troubleshooting analysis of Layers 2 through 7.

Although many vendors have extended the functionality of RMON to support switched networks, it was not originally designed to do so. Switch Monitoring (SMON), defined in RFC 2613, is an extension to the RMON protocol. SMON was specifically developed to address the remote monitoring requirements of switched environments. It includes group objects to monitor a group of ports or all ports on a switch, and to monitor VLANs and port copy functions.

Other extensions to the RMON protocol include High Capacity RMON (HCRMON) and Differentiated Services Monitoring (DSMON). HCRMON, defined in RFC 3273, accommodates high capacity networks through increased counter sizes for RMON1 and RMON2. HCRMON is a necessity in today's high-speed switched networks where counters must support connection speeds of Gigabit Ethernet and higher.

DSMON, defined in RFC 3287, provides a standardized way to monitor network traffic based on differential services code point (DSCP) values, which are assigned to various classes of traffic based on a quality of service (QoS) policy. You can use the data gathered using DSMON to verify and tune application throughput for the various services associated with specific DSCP values.

NetFlow

NetFlow-collected data serves as the base for a set of applications including network traffic accounting, usage-based network billing, network planning, and network monitoring. NetFlow also provides the measurement base for QoS applications. NetFlow captures the traffic classification or precedence associated with each network flow, enabling differentiated charging based on QoS. This topic describes NetFlow.

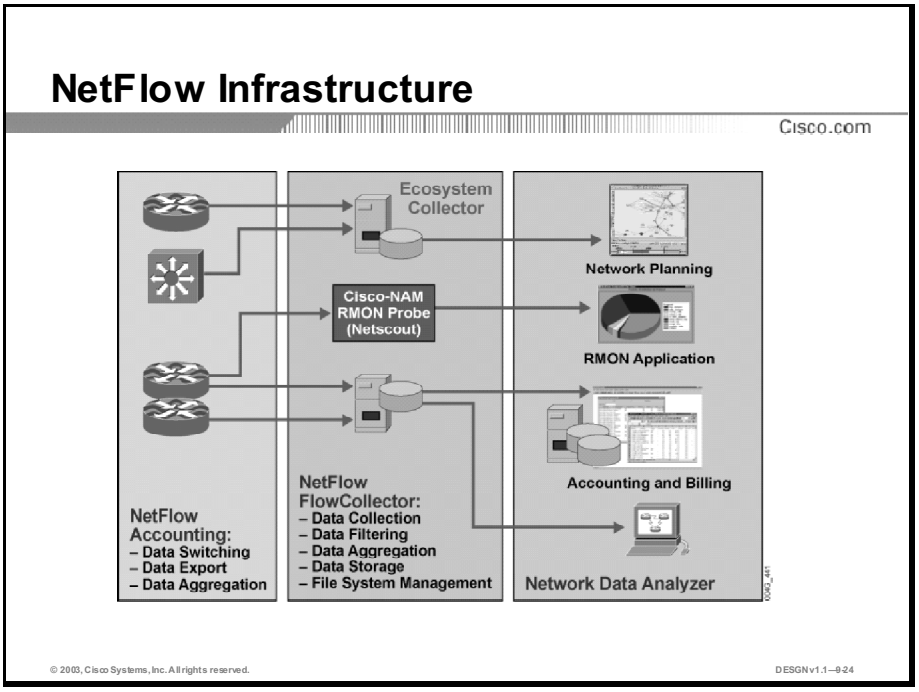

A network flow is defined as a unidirectional sequence of packets between source and destination endpoints. Flow endpoints are identified by IP address and transport layer application port numbers. NetFlow also identifies the flows by IP protocol type, class of service (CoS), and the input interface identifier.

Non-NetFlow enabled switching handles incoming packets independently, with separate serial tasks for switching, security, services, and traffic measurements applied to each packet. NetFlow-enabled switching applies security (access control list [ACL]) processing only to the first packet of a flow. Information from the first packet is used to build an entry in the NetFlow cache. Subsequent packets in the flow are handled via a single streamlined task that handles switching, services, and data collection concurrently.

You can configure NetFlow on individual interfaces to collect the following types of information:

- Source and destination interface numbers
- Source and destination IP addresses
- TCP/UDP source port and destination ports
- Number of bytes and packets in the flow
- Source and destination autonomous system numbers
- IP CoS

NetFlow services capitalize on the “flow” nature of traffic in the network to accomplish these tasks:

- Provide detailed data collection with minimal impact on router performance
- Process ACLs efficiently for packet filtering and security services

NetFlow enables these key applications:

- **Accounting and billing:** NetFlow data provides fine-grained metering for highly flexible and detailed resource utilization accounting. Service providers may use this information to migrate away from single-fee, flat-rate billing to more flexible charging mechanisms based on time of day, bandwidth usage, application usage, QoS, and so on. Enterprises may use the information for departmental cost recovery or cost allocation for resource utilization.
- **Network planning and analysis:** NetFlow data provides key information for sophisticated network architecture tools to optimize strategic planning (for example, whom to peer with, backbone upgrade planning, routing policy planning) and tactical network engineering decisions (such as adding additional resources to routers, upgrading link capacity). This has the benefit of minimizing the total cost of network operations while maximizing network performance, capacity, and reliability.
- **Network monitoring:** NetFlow data provides extensive, almost real-time, network monitoring. You can use flow-based analysis techniques to visualize traffic patterns associated with individual routers and switches on a network-wide basis providing aggregate traffic or application-based views. This analysis provides network managers with proactive problem detection, efficient troubleshooting, and rapid problem resolution.
- **Application monitoring and profiling:** NetFlow data enables network managers to gain a detailed, time-based view of application usage over the network. Content and service providers may use this information to plan and allocate network and application resources, such as web server sizing and location, to meet customer demands.
- **User monitoring and profiling:** NetFlow data identifies customer and user network utilization and resource application. Network managers use this information to plan efficiently; allocate access, backbone, and application resources; and detect and resolve potential security and policy violations.
- **NetFlow data warehousing and mining:** You can warehouse NetFlow data for later retrieval and analysis. NetFlow data enables service providers to create a wider range of offered services. For example, the service provider can easily determine the traffic characteristics and provide new services to the users, such as Voice over IP (VoIP), which requires a QoS adjustment.Deploy NetFlow on the edge or distribution router interfaces for service providers, or on WAN access router interfaces for enterprises.

Cisco recommends that you carefully plan a NetFlow deployment, with NetFlow services activated on strategically located routers. You can deploy NetFlow incrementally (interface by interface) and strategically (on select routers), rather than deploying NetFlow on every router on the network.

You can export the NetFlow data to network management applications that further process the information, resulting in display tables and graphs for accounting and billing, or as an aid for network planning.

NetFlow, compared to SNMP with RMON MIB, offers greater detail, time stamping, customized data collection according to interface, and greater scalability. The performance impact of NetFlow is much lower than that of RMON, and external probes are not required.

CDP

CDP is a Cisco proprietary protocol that enables you to discover Cisco devices on the network. This topic describes CDP.

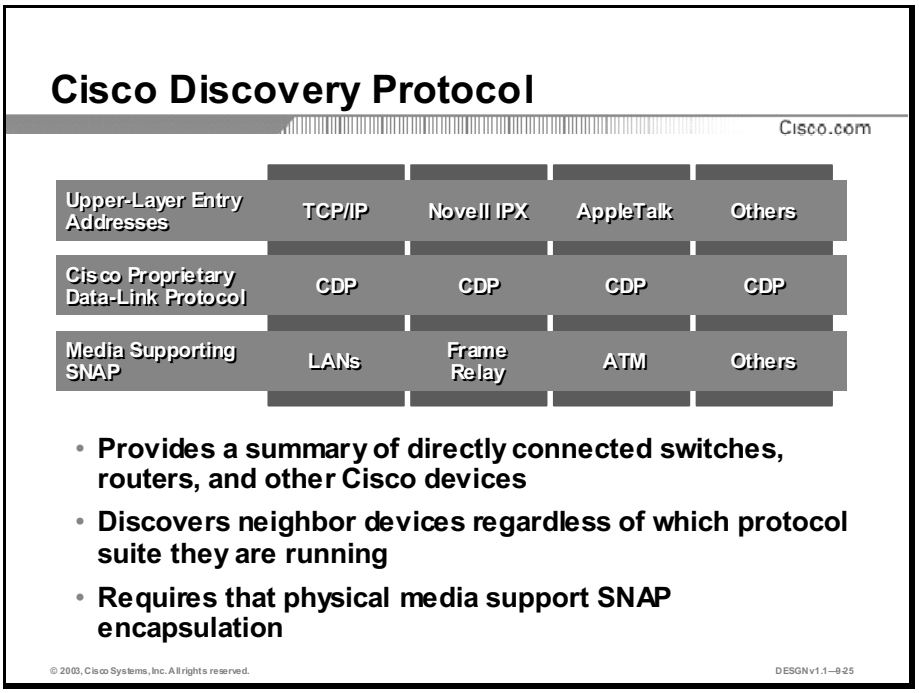

CDP is a media- and protocol-independent protocol that is enabled by default on each supported interface of a Cisco device, such as routers, access servers, and switches. The physical media must support Subnetwork Access Protocol (SNAP) encapsulation.

CDP runs over the data link layer, enabling two systems that support different network layer protocols to communicate.

Here is some of the information that Cisco devices exchange in the CDP packet:

- **Device ID:** The name of the neighbor device and either the MAC address or the serial number of this device
- **Local interface:** The local interface connected to the discovered neighbor
- **Hold time:** The remaining amount of time (in seconds) the current device will hold the CDP advertisement from a sending router before discarding it
- **Capability:** The type of device that is discovered (R—Router, T—Trans Bridge, B—Source Route Bridge, S—Switch, H—Host, I—IGMP, r—Repeater)
- **Platform:** The product number of the device
- **Port ID:** The port number on the discovered neighbor
- **Address:** This identifies all the network layer protocol addresses that have been configured on the interface (or on the box, in the case of protocols that are configured globally, for example IP, IPX, DECnet)

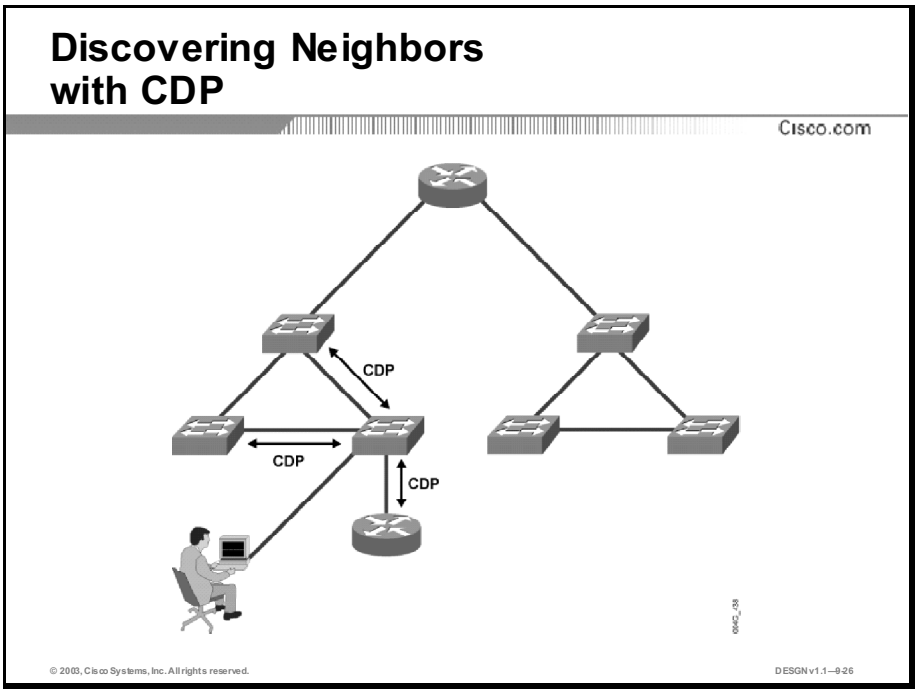

CDP is a “hello-based” protocol, and all Cisco devices running CDP will periodically advertise their attributes to their neighbors using a multicast address. CDP packets advertise a time-to-live value in seconds, which indicates the length of time to retain the packet before discarding it. Cisco devices send CDP packets with a time-to-live value that is nonzero after an interface is enabled. A time-to-live value of zero is sent immediately before an interface is idled down. Sending a CDP packet with a time-to-live value of zero allows a network device to quickly discover a lost neighbor.

All Cisco devices receive CDP packets and cache the information in the packet. The cached information is then available to an NMS using SNMP.

Note: Cisco devices never forward a CDP packet.

If any information changes from the last received packet, the device caches the new information and discards the previous information even if its time-to-live value has not yet expired.

Note: Do not run CDP in these places:

1. Do not run CDP on links you do not want discovered, such as Internet connections.
2. Do not enable CDP on links that do not go to Cisco devices.

Syslog

The system message and error reporting service (syslog) is an essential component of any network operating system. The system message service reports system state information to a network manager. This topic describes syslog.

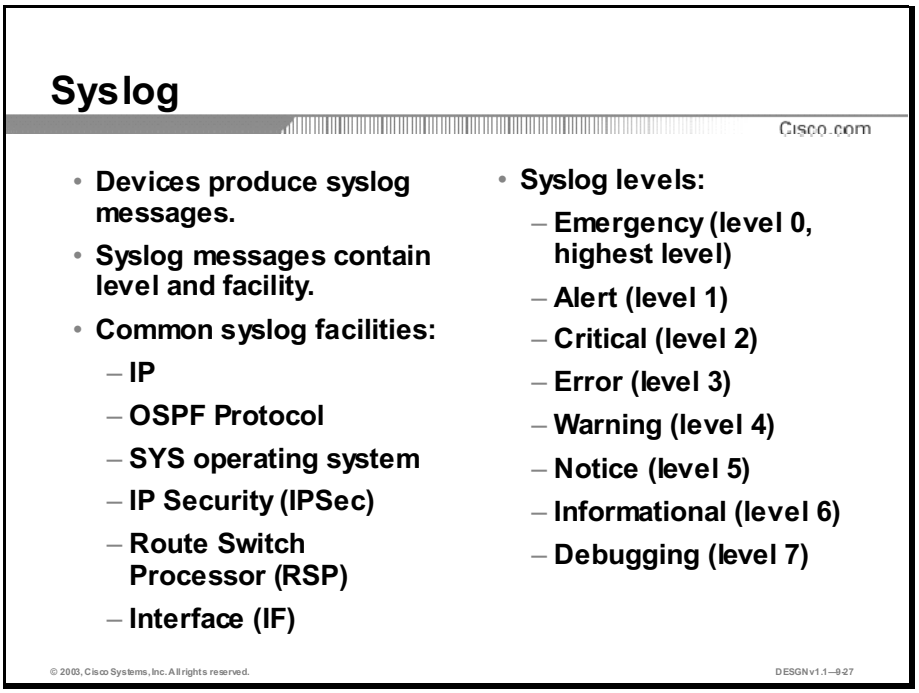

Cisco devices produce syslog messages as a result of network events. Every syslog message contains a timestamp, level, and facility. Many networking devices support syslog including routers, switches, application servers, firewalls, and other network appliances.

Syslog defines the levels listed in the figure.

Syslog facilities are service identifiers that are used to identify and categorize system state data for error and event message reporting. IOS software has more than 500 different facilities. The most common syslog facilities are listed in the figure.

Other facilities include CDP, Spanning Tree Protocol (STP), MCAST (multicast), SEC (IP security), TCP, Border Gateway Protocol (BGP), RADIUS, Telnet, and facilities related to QoS services.

Note: More syslog information is located at http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123sup/123sems/index.htm

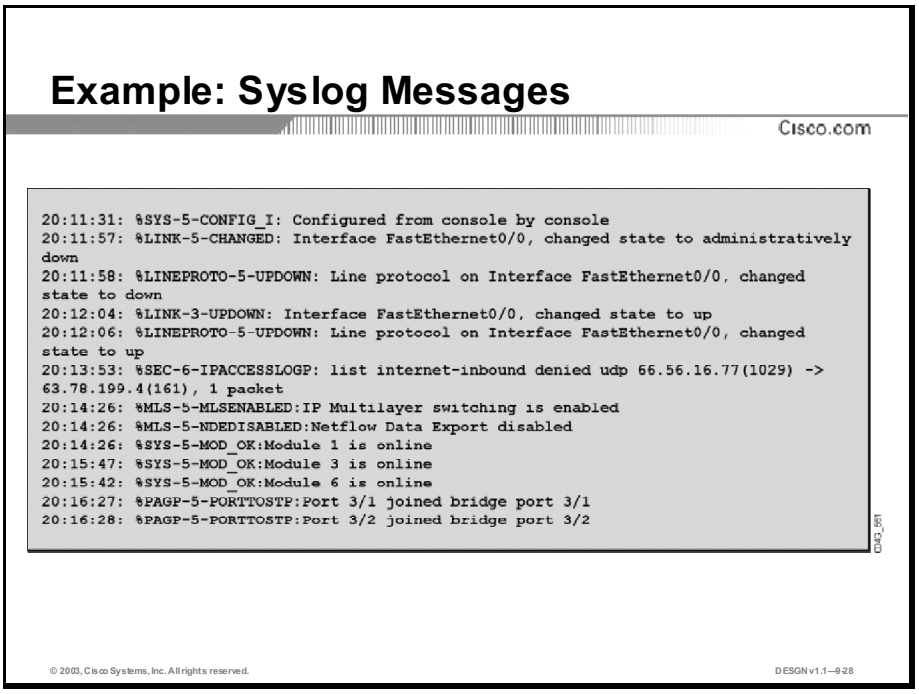

The example shows samples of syslog messages that IOS software produces. The most common messages are link up and down messages and messages that a device produces when it exits from Configuration Mode.

If ACL logging is configured, the device will generate syslog messages when packets match the parameter condition. ACL logging is very useful to detect packets that are denied access based on the security policy set by an ACL.

Syslog messages are displayed in this format:

- mm/dd/yy:hh/mm/ss:FACILITY-LEVEL-mnemonic:description

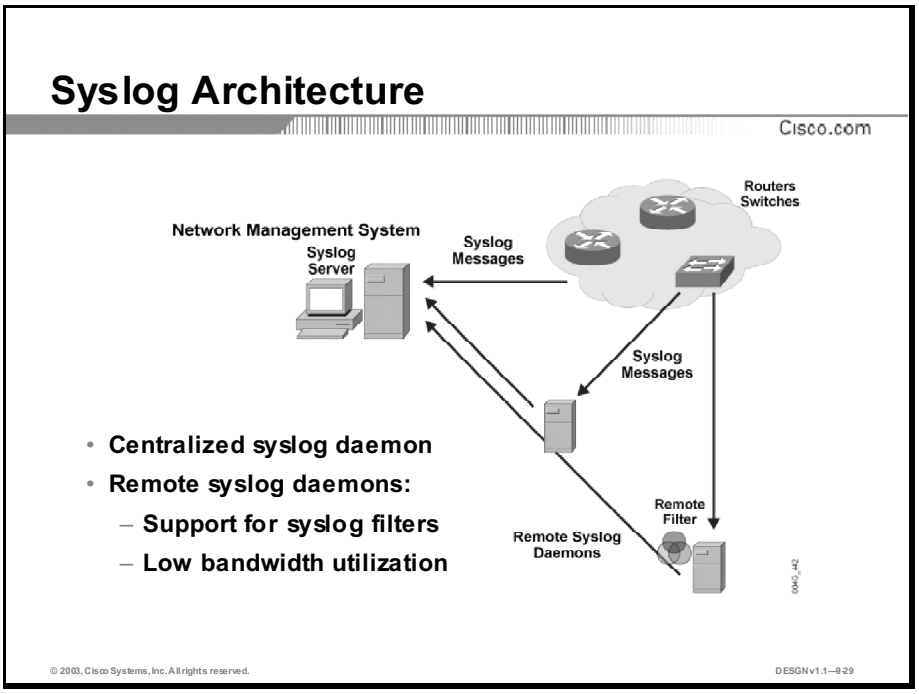

Syslog messages are sent by default to console sessions. To send syslog messages to the NMS, you must configure the device to send syslog messages to the address of the NMS running the syslog server function.

You can configure network devices to send syslog messages directly to the NMS or a remote network host on which a distributed syslog server is installed. A Syslog Analyzer conserves bandwidth on WAN links, because the remote analyzer usually applies different filters and sends only the predefined subset of all syslog messages it receives. The analyzer filters and periodically forwards messages to the central NMS.

Note: Syslog Analyzer is a CiscoWorks Resource Manager Essentials (RME) application that supports a distributed syslog server architecture for localized collection, filtering, aggregation, and forwarding of syslog data to a central syslog server for further processing and analysis. Syslog Analyzer also supports reporting functions to automatically parse the log data into predefined or custom formats for ease of use and readability.

Upon receiving a syslog message, the NMS applies filters to remove unwanted syslog messages. You can apply action filters to perform actions based on the received syslog message, such as paging or e-mailing the network manager.

Syslog data can consume large amounts of network bandwidth and may require very large storage capacity based on the number of devices sending syslog messages, the syslog facility and severity levels set for each, and any error conditions that may trigger excessive log messages. Therefore, it is important to enable logging only for network facilities of particular interest and to set the appropriate severity level to provide sufficient but not excessive detail. If the collected data is not analyzed, do not collect it in the first place. Selectively filter and aggregate syslog data that the distributed or centralized syslog servers receive based on the requirements.

Summary

This topic summarizes the key points discussed in this lesson.

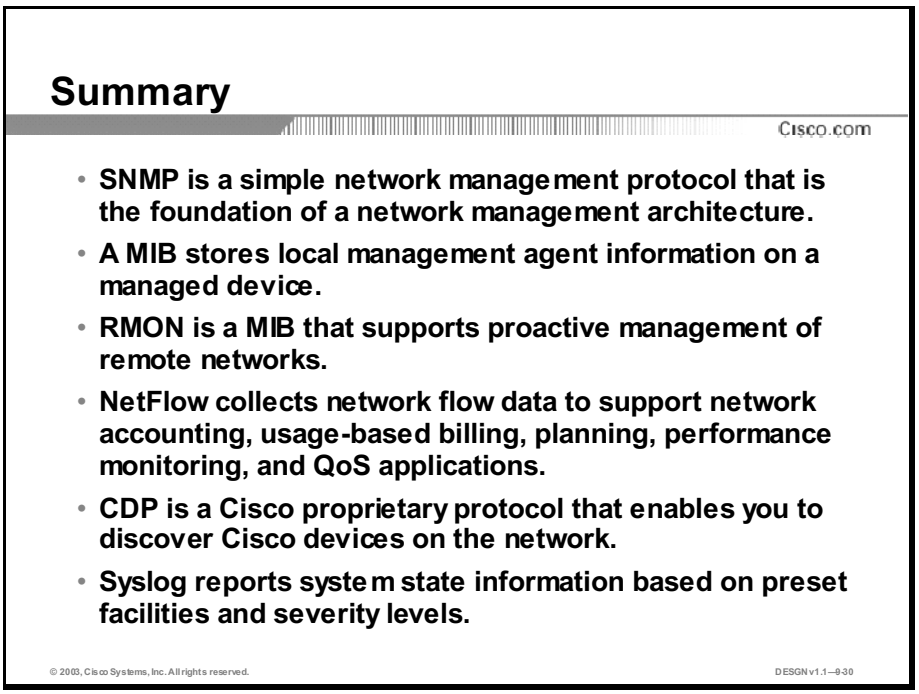

References

For additional information, refer to these resources:

- Stallings, W. *SNMP, SNMPv2 and CMIP*. Reading, Massachusetts: Addison-Wesley, 1996.
- Leinwand, F. and K. Fang. *Network Management*. Reading, Massachusetts: Addison-Wesley, 1995.
- *Cisco Management Information Base (MIB) User Quick Reference,* http://www.cisco.com/univercd/cc/td/doc/product/software/ios112/mbook/index.htm

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Which Simple Network Management Protocol (SNMP) message does the agent send without the explicit request from the network management system (NMS)?

A) Get Request

B) Set Request

C) Trap

D) Get Response

Q2) Which four standard managed objects does a Cisco router have defined in the standard section of the MIB tree such as? (Choose four.)

A) traps

B) interfaces

C) buffers

D) applications

E) memory

F) standard protocols

Q3) Which group is not an RMON1 group?

A) packet capture

B) token ring

C) protocol directory

D) history

Q4) What is the main reason for the deployment of NetFlow?

A) to improve security of the network traffic flows

B) to consolidate multiple flows of data traffic over a single interface

C) to provide detailed data collection with minimal performance impact

D) to improve congestion management on a device interface

Q5) On which OSI layer does CDP run?

A) application layer

B) session layer

C) network layer

D) data link layer

Q6) What is the highest (most critical) syslog priority level?

A) zero

B) one

C) six

D) seven

Quiz Answer Key

Q1) C

Relates to: SNMP

Q2) B, C, E, F

Relates to: MIB

Q3) C

Relates to: RMON

Q4) C

Relates to: NetFlow

Q5) D

Relates to: CDP

Q6) A

Relates to: Syslog

Reviewing Functional Areas of Network Management

Overview

The International Organization for Standardization (ISO) network management model defines five functional areas of network management: Fault management, Configuration management, Accounting management, Performance management, and Security management (FCAPS).

This lesson provides a high-level overview of each functional area of network management and gives practical recommendations to increase the overall effectiveness of current management tools and practices. It also provides design guidelines for future implementation of network management tools and technologies.

Relevance

You should use the FCAPS framework to ensure that your network management design provides a comprehensive solution.

Objectives

Upon completing this lesson, you will be able to describe the functional areas of network management. This includes being able to meet these objectives:

- Describe the FCAPS functional model
- Describe the goals of Fault management and how to achieve them
- List available configuration management tools
- Describe how to collect accounting data from routers and switches
- Describe solutions for collecting, storing, and presenting data from network devices
- Describe methods for controlling access on routers and switches

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of networking concepts and network management protocols

Outline

The outline lists the topics included in this lesson.

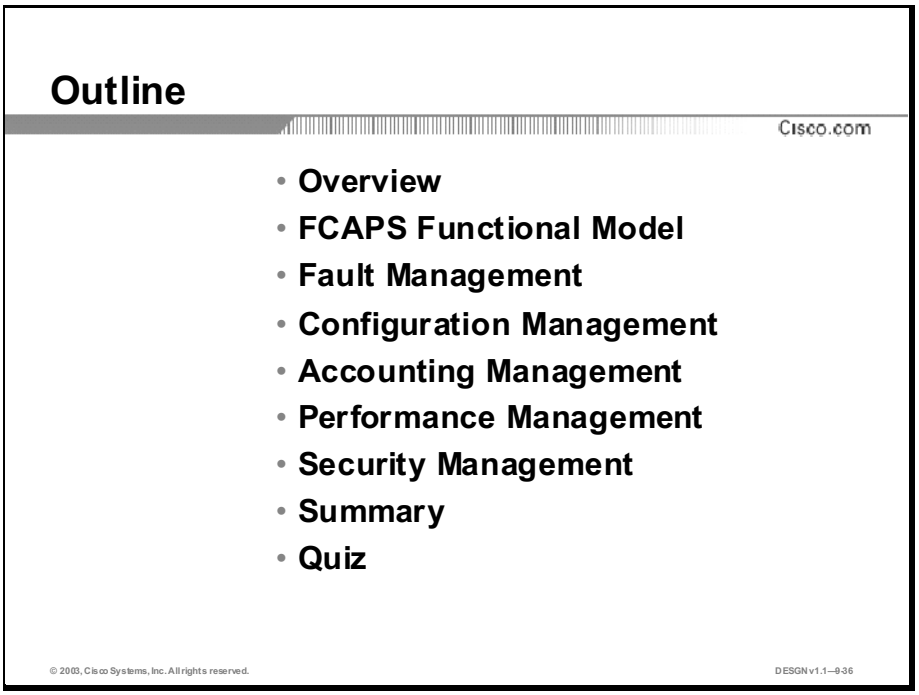

FCAPS Functional Model

ISO proposes a model for comprehensive network management called FCAPS. Each functional area influences the effectiveness of management tools and management practices. This topic describes the FCAPS functional model.

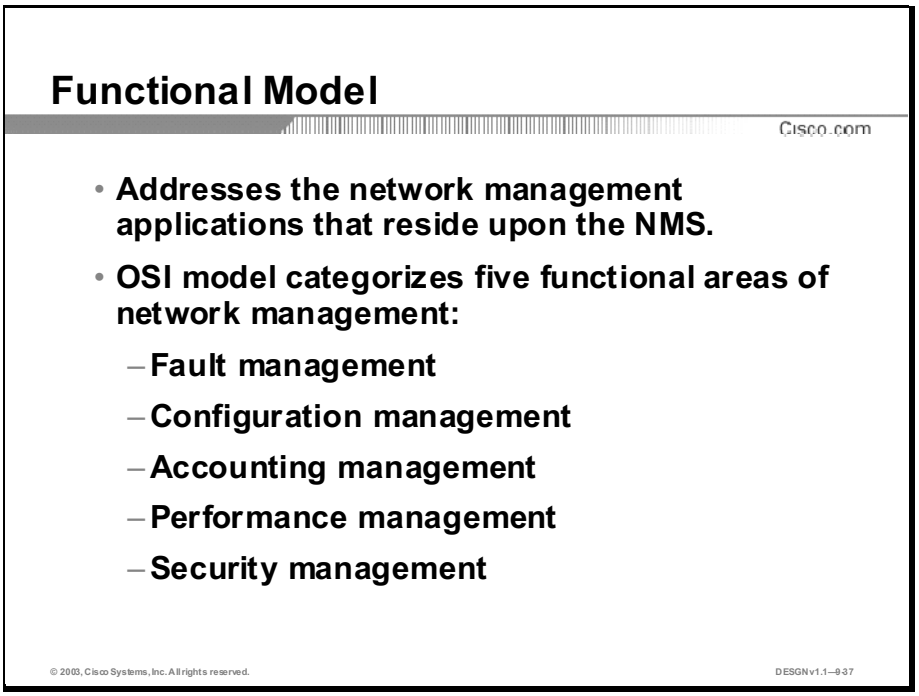

The network management model defines five functional areas of network management:

- **Fault management:** Detects, isolates, notifies, and corrects faults encountered in the network
- **Configuration management:** Manages the configuration of network devices, such as file management, inventory management, and software management
- **Accounting management:** Provides information on network resource usage
- **Performance management:** Monitors and measures various aspects of performance so that overall performance can be maintained at an acceptable level
- **Security management:** Provides access to network devices and corporate resources to authorized individuals

Fault Management

Fault management is designed to handle error conditions that cause users to lose the full functionality of a network resource. This topic discusses Fault management.

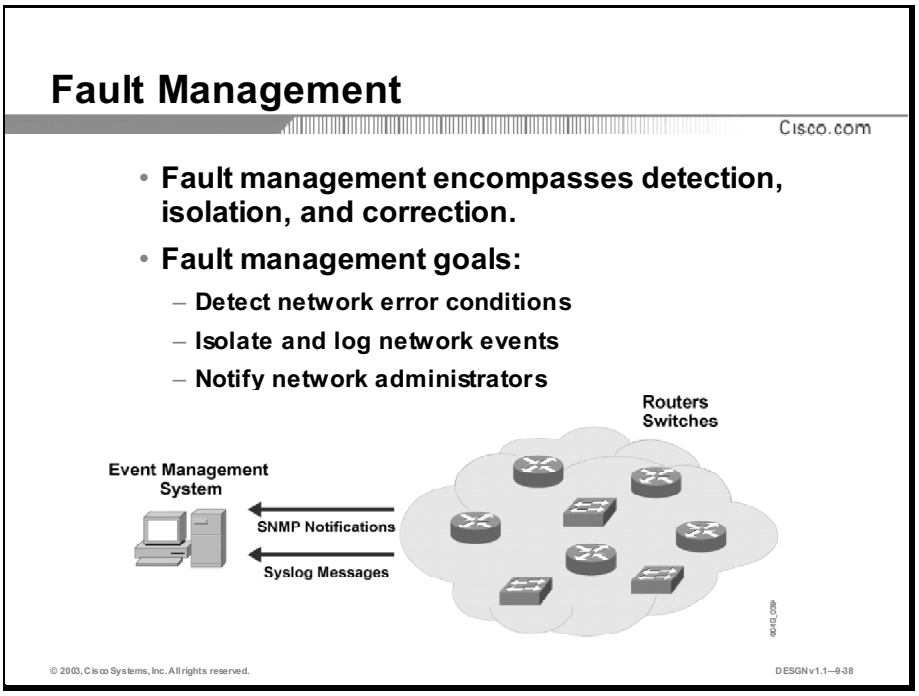

Fault management includes detection, isolation, and correction of abnormal network operation. It provides the means to receive and present faults, determine the cause of a network fault, isolate the fault, and perform a corrective action to keep the network running effectively.

Fault management is performed in five basic steps:

Step 1 **Fault determination:** Detects faults.

Step 2 **Diagnosis:** Determines the cause of the fault and an action to resolve it.

Step 3 **Bypass and recovery:** Attempts to bypass the fault until the fault is permanently fixed.

Step 4 **Resolution:** Eliminates the fault.

Step 5 **Fault tracking and control:** Tracks the fault through final resolution.

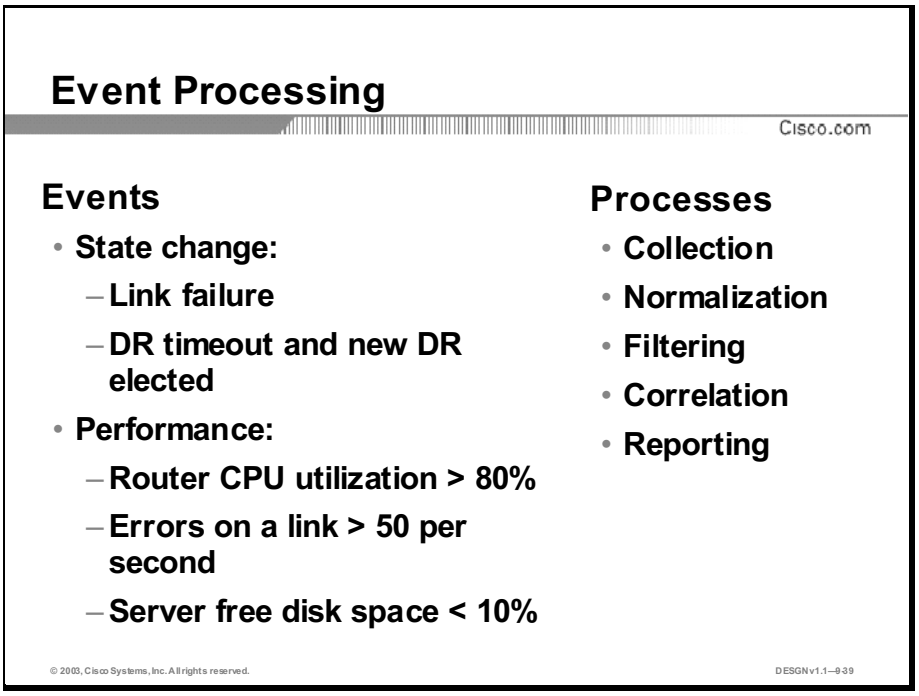

The fault management architecture consists of event generators and event collectors. Event generators, the entities that produce events, are typically network devices such as routers, switches, and hosts. The events generated are typically SNMP notifications (traps, informs, or RMON alarms and notifications), syslog messages, or results of SNMP polls. The event generators are configured with the address of the event collector, enabling them to forward events to the appropriate location.

Event collectors collect the events that event producers send and use predefined sets of rules to determine which events require which types of action.

Events

Events can be divided into two general categories:

- **State change events:** Triggered when a network device changes state. For example, you can configure network devices to send state change events when a link or a neighbor goes down, or when a current designated router times out and a new designated router is elected.
- **Performance events:** Generated when a network device detects a possible performance issue.

Processes

An event processor, typically part of the NMS, collects and processes events produced by event generators. Syslog messages are sent with a facility code (referring to a particular hardware device, a protocol, or a portion of the system software) and a priority. Event producers send messages only with priorities higher than those defined in a device, so that only messages with critical information are sent to the collector. You can store other messages not crucial for network activity on a device, avoiding collector and network utilization. The event collector can filter syslog messages containing certain facility or priority codes.

SNMP notifications are similar to syslog messages because the event generator sends them to the event collector. Because SNMP notifications do not contain the priority of the event, the event collector must maintain a table of SNMP notifications and their priorities.

Event correlation correlates multiple events from the same source or multiple sources, identifies duplicate or related events, and intelligently interprets the true source of the fault condition. Event correlation reduces the number of individual alarms, many of which can be misleading or simply “noise,” and permits a focused, singular action. A simple example is a central switch or router that loses power. Without event correlation, it is possible that an NMS would poll all network devices and incorrectly determine that all devices connected through that device have also failed.

Event reporting uses various methods to inform network engineers and operators of critical events in the network, such as displaying popup notifications, sending e-mails, or paging network staff.

Configuration Management

Configuration management is a collection of processes and tools that promote network consistency, track network changes, and provide up-to-date network documentation and visibility. By building and maintaining configuration management best practices, the network manager can expect benefits such as improved network availability and lower costs. This topic discusses configuration management.

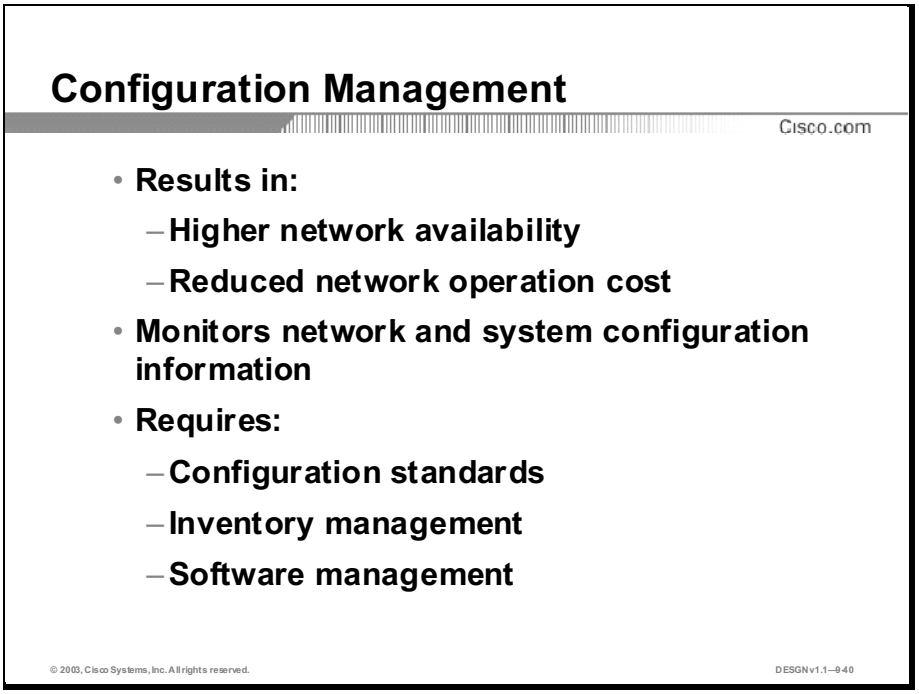

Configuration management includes functions to accomplish these tasks:

- Set the parameters that control routine operation and changes in the network (change control policy)
- Define naming and addressing standards for network devices
- Assign contact names and numbers for each network device
- Collect current configuration and status information for network resources on a periodic basis or on demand
- Receive notifications regarding network changes (what, where, when, and how)
- Manage (deploy and update) the configuration of network devices
- Generate reports on network configurations, hardware and software inventory, and changes

The goal of configuration management is to monitor network and system configuration information, enabling the network manager to track and manage various versions of hardware and software elements. These are some of the goals of configuration management:

- Lower support costs by decreasing reactive support issues
- Lower network costs through identification of unused network components
- Improve network availability through the decrease in reactive support costs and improve time to resolve problems

A lack of configuration management can result in issues such as these:

- Inability to determine the impact of network changes or outages on the end user
- Increased reactive support issues
- Increased time to resolve problems
- Higher network costs because of unused network components

Configuration Standards

The greater the number of devices in a network, the more critical it is to identify accurately the location of each device. This location information should provide a detailed description that is meaningful to those tasked with dispatching resources when a network problem occurs. To expedite a resolution of a network problem, it is critical to have the contact information of the person or department responsible for the devices.

You should plan and implement naming conventions for network devices and interfaces as part of the configuration standard. The naming convention for devices can use geographical location, building name, floor, and so on. The naming convention for interfaces can include information such as the segment to which a port is connected and the name of the connecting hub. On serial interfaces, the naming convention may include the actual bandwidth, the local data-link connection identifier (DLCI) number (for Frame Relay), the destination, the circuit ID, or information provided by the carrier.

You should develop a set of rules governing the process for implementing network changes to prevent unnecessary or excessive down time. Example rules may specify required authorization and approval processes, documentation procedures, back-out or contingency plans, and time periods when changes are permitted based on the criticality of an update or impact to users and business functions. Many companies, for example, have well-defined maintenance windows that dictate which changes are permitted and when, typically set at weekly or monthly intervals during periods of low usage or impact. Scheduled maintenance periods also permit users to plan around planned service disruptions.

Configuration File Management

When adding new configuration commands to existing network devices, network managers must verify the commands for integrity before completing the actual implementation. Managers must check configuration command parameters to avoid mismatches and incompatibility issues, and should review configurations with expert engineers on a regular basis.

You must track changes to configuration files to determine the impact of specific changes. You should maintain backup copies of configuration files to permit rollback to a previous configuration if changes have negative results or a device needs to return to a previous configuration version.

Inventory Management

An inventory database provides detailed configuration information about network devices. Common information includes models of hardware, installed modules, software images, microcode levels, and so on. The up-to-date listing of network devices collected during the discovery process creates a master list of inventory information collected using SNMP or program scripting. The discovery function of most network management platforms provides a dynamic listing of devices in the network.

Software Management

To successfully upgrade software images on network devices, you need a detailed analysis of the requirements such as memory, boot ROM, and microcode level. The requirements are normally documented and available on vendor web sites. To upgrade a network device running IOS software, download a correct image from Cisco.com, back up the current image, ensure that all hardware requirements are met, and then load the new image into the device.

Many organizations have a limited window of opportunity to complete device maintenance. In a large network environment with limited resources, you may need to schedule and automate software upgrades after business hours by using either a scripting language or a specific application such as CiscoWorks Resource Manager Essentials (RME).

You should track changes to software, such as IOS software images and microcode versions, to assist in the analysis phase when subsequent software maintenance is required. Using a modification history report minimizes the risk of loading incompatible images or microcode into network devices.

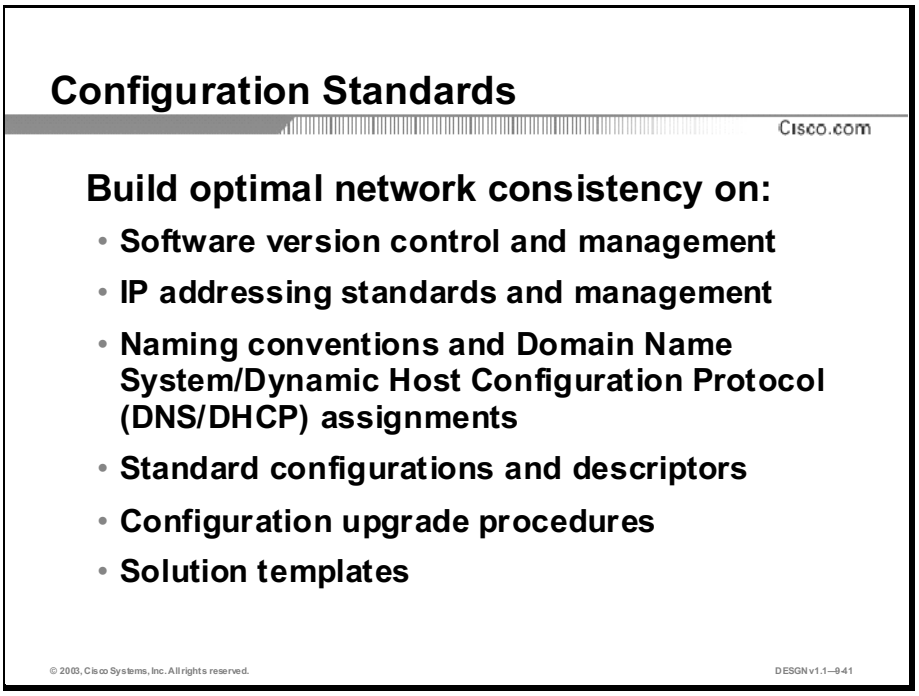

Creating standards for network consistency reduces network complexity, the amount of unplanned downtime, and exposure to events that negatively impact the network.

Software Version Control and Management

Software version control is the practice of deploying consistent software versions on similar network devices. Version control improves the opportunity for validation and testing on selected software versions and limits the amount of software defects and interoperability issues found in the network. Similar network devices should have the same software versions, to reduce the risk of unexpected behavior with user interfaces, command or management output, upgrades, and features. Overall, software version control improves network availability and helps to lower reactive support costs.

IP Addressing Standards and Management

IP address management is the process of allocating, recycling, and documenting IP addresses and subnets in a network. IP addressing standards define subnet size, subnet assignment, network device assignments, and dynamic address assignments within a subnet range. Recommended IP address management standards reduce the opportunity for overlapping or duplicate subnets, duplicate IP address device assignments, wasted IP address space, and unnecessary complexity.

Naming Conventions and DNS/DHCP Assignments

Consistent, structured use of naming conventions and DNS for devices enables network management in these ways:

- Creates a consistent location for network naming and address information related to a device
- Reduces ambiguity and the opportunity for duplicate device names and IP addresses
- Creates identification, such as location, device type, and purpose, that is simple for a device
- Improves inventory management by providing a simpler method to identify network devices

Standard Configuration and Descriptors

Create standard configurations for each device classification, such as a router, LAN switch, or voice gateway that defines the global, media, and protocol configuration necessary to maintain network consistency. Global configuration commands apply to all like devices and include parameters such as service, IP, and TACACS+/RADIUS commands, vty settings, banners, and configuration of SNMP, syslog, and Network Time Protocol (NTP).

Descriptors are interface commands used to describe an interface. Develop descriptors by creating a standard format that applies to each interface. The descriptor includes the purpose and location of the interface, other devices or locations connected to the interface, and circuit identifiers.

Configuration Upgrade Procedures

Upgrade procedures help to ensure that software and hardware upgrades, large and small, occur smoothly with minimal downtime. Upgrade procedures include vendor verification and vendor installation references, such as release notes, upgrade methodologies or steps, configuration guidelines, and testing requirements.

Create upgrade procedures in conjunction with new software deployment or standard releases. The procedures should define all steps for the upgrade, reference vendor documentation related to updating the device, and provide testing procedures for validating the device after the upgrade. Once upgrade procedures are defined and validated, reference the procedures in all change documentation appropriate to the particular upgrade.

Solution Templates

Use solution templates to define standard modular network solutions. A network module may be a wiring closet, a remote office, or an access concentrator. In each case, define, test, and document the solution to ensure that you can implement similar deployments in exactly the same way. This ensures that future changes occur at a much lower risk level to the organization because the behavior of the solution is well–defined and understood.

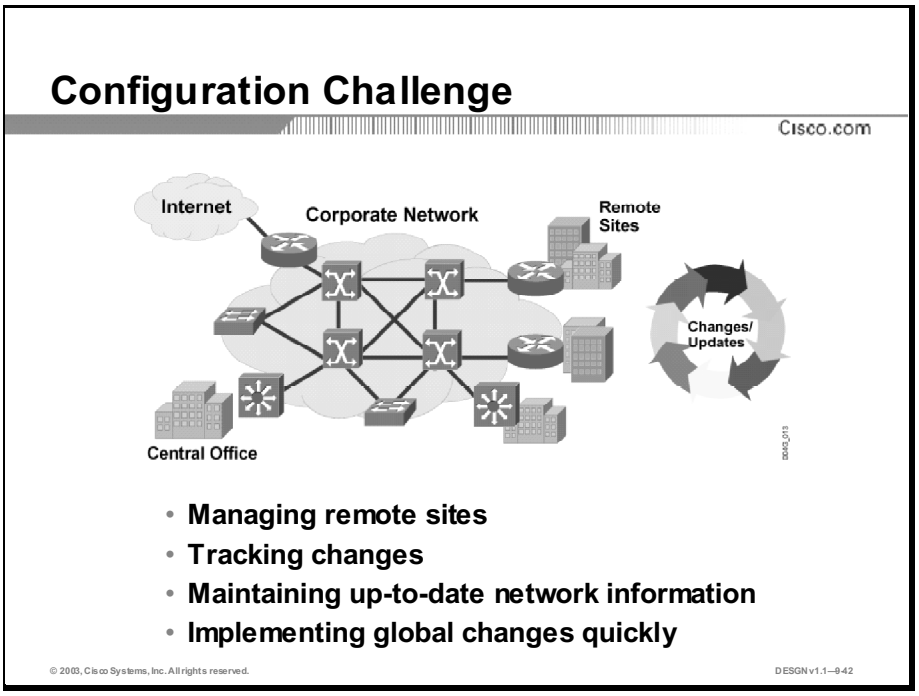

Network administrators require powerful, easy-to-use tools that provide information on software and hardware profiles, allow them to manage and update device configurations, and track changes to the network. These are some of the challenges that network managers may face:

- **Managing remote sites:** In a large, distributed network, network administrators are often responsible for devices in remote locations. Therefore, to troubleshoot problems efficiently, network managers require easy access to the devices for which they are responsible without having to be in the same physical location. At the same time, you need to control access to these devices to protect the security of the network and applications.
- **Tracking changes and maintaining up-to-date network information:** Many people often share responsibility for the network and its configuration, making it difficult to maintain accurate network documentation. To locate the source of problems quickly or to plan upgrades and expansion, you need up-to-date information on device inventory, software versions running, and specific device configurations maintained in one central place.
- **Implementing global changes quickly:** To upgrade software or configuration files on a large network with many devices, you need automated tools that reduce deployment time and errors.

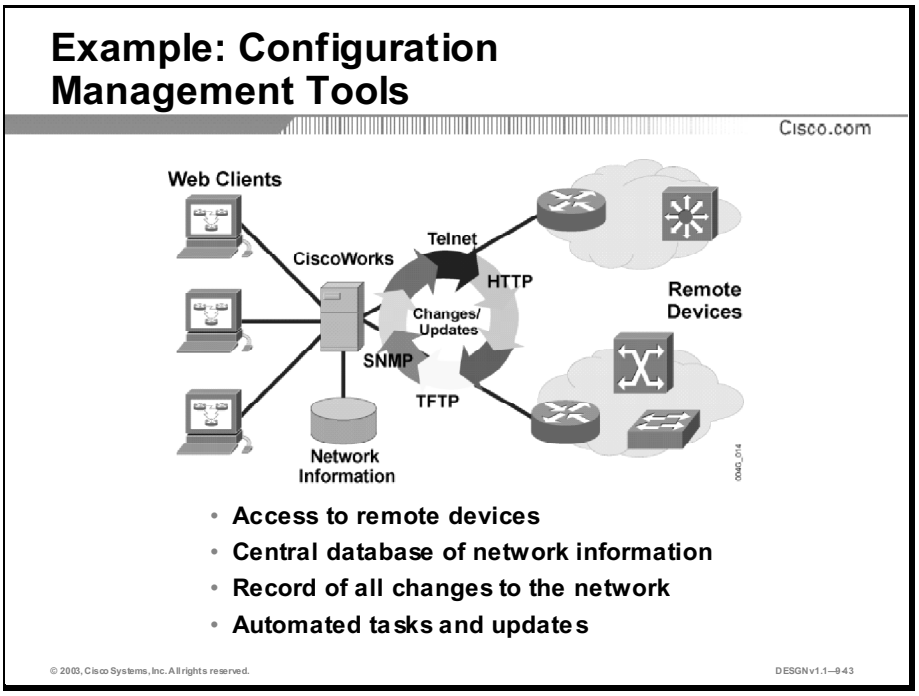

To maintain and troubleshoot a large network efficiently, network managers need access to current network information and all troubleshooting tools in one integrated place on the desktop. Network managers also need the ability to identify network changes and perform updates to multiple devices quickly.

The Cisco management tool called CiscoWorks RME, shown in the figure, collects information from the network devices via different methods. RME manages information about device configurations, inventory, and software versions.

CiscoWorks RME

CiscoWorks is a suite of web-enabled management solutions based on a common foundation of services for managing Cisco devices. The CiscoWorks RME application includes a configuration tool that supports industry-standard management protocols, including SNMP, SSH, Telnet, RCP and TFTP. All management data is centrally stored, allowing the network manager to access information on remote network devices through any web browser.

RME maintains a database of network information and provides reports used for configuration and troubleshooting. Once devices are added to the inventory, the network administrator can schedule tasks to periodically retrieve and update device information such as hardware, software, and configuration files to ensure that the most current network information is stored. In addition, RME automatically records any changes to network devices, making it easy to identify when changes are made and by whom.

Cisco Structured Wireless-Aware Network

The Cisco Structured Wireless-Aware Network (SWAN) is a comprehensive Cisco framework for deploying, operating, and managing Cisco Aironet access points using the Cisco infrastructure. The Cisco Structured Wireless-Aware Network extends to the wireless LAN the same level of security, scalability, and reliability provided in the wired LAN.

Cisco Structured Wireless-Aware Network infrastructure includes enhancements integrated in various Cisco Aironet access points and other networking devices. Other components of the solution include CiscoWorks Wireless LAN Solution Engine (WLSE) for management and monitoring, Cisco IOS Software, Cisco Secure Access Control Server for centralized authentication, and Cisco and Cisco-compatible client adapters for radio frequency (RF) monitoring and measurement.

Accounting Management

The goal of accounting management is to measure and regulate network utilization, minimizing network problems and maximizing network access for all users. This topic describes accounting management.

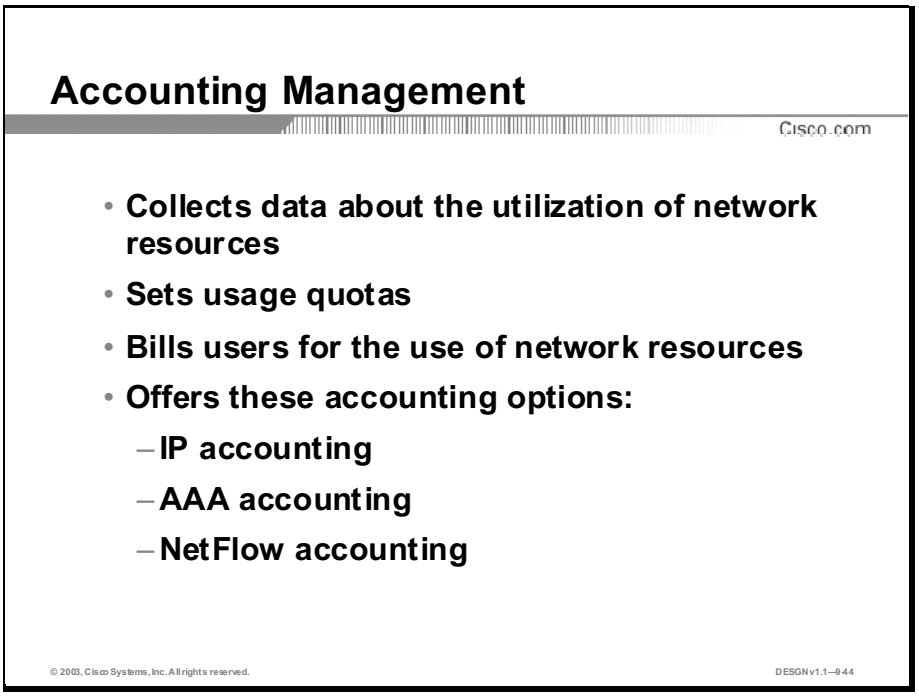

The first step toward appropriate accounting management is to measure utilization of all important network resources. By analyzing the results, you will understand current usage patterns so you can set usage quotas. Often, adjustments are necessary to achieve optimal access for all users.

Different accounting tools are available, depending on the accounting requirements:

- **IP accounting:** Used in Cisco IOS software to collect the information about the number of packets and bytes transferred between any pair of IP endpoints
- **Authentication, authorization, and accounting (AAA):** Used to log events such as connection times and utilization statistics for remote access (dialup, VPN) users
- **NetFlow accounting:** Similar to IP accounting except that it provides more detail by monitoring flows instead of pairs of IP addresses

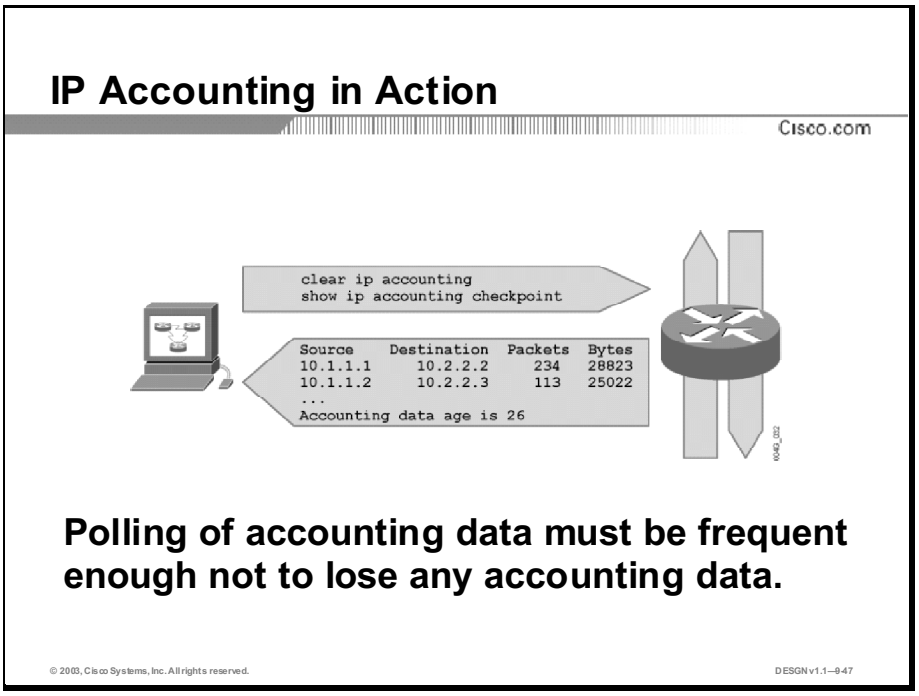

Use IP accounting to measure the amount of traffic that different IP hosts or subnets generate. A router will hold accounting information in its memory that identifies the source and destination IP address of each packet, as well as the number of packets and bytes exchanged between IP endpoints.

A server must poll to transfer the router accounting information to an accounting server. If the polling is not frequent, the router loses the accounting data.

IP accounting includes features such as these:

- Accounting of packets to and from devices based on MAC addresses
- Accounting of packets based on IP precedence values
- Accounting of access control list (ACL) violations

The figure illustrates the polling process that the accounting server initiates periodically. The server freezes the accounting database on the router and then downloads the accounting information.

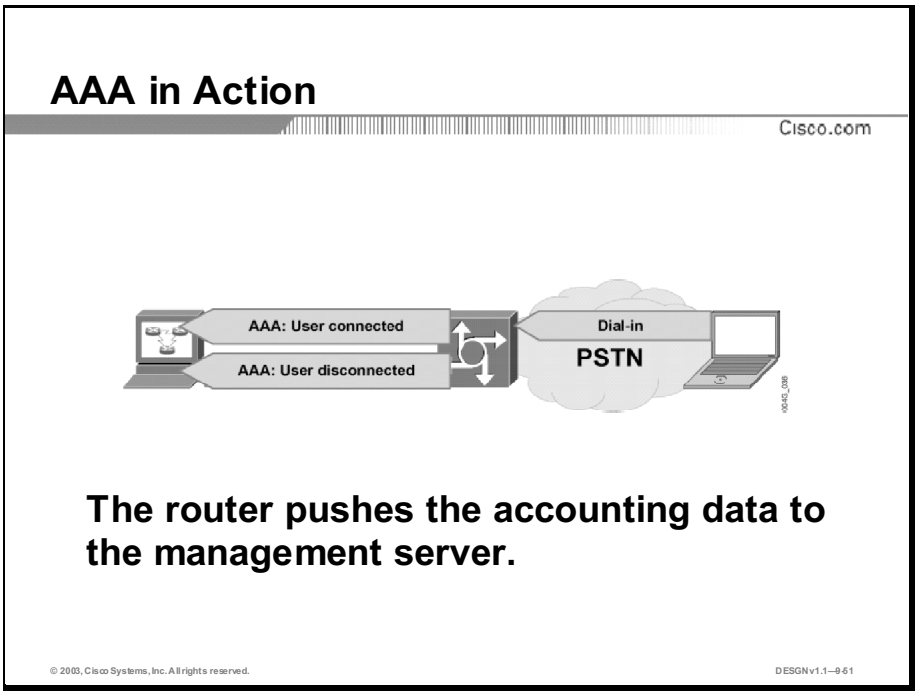

AAA network security services provide the framework through which a network administrator can establish access control on network entry points or network access servers. Authentication identifies a user; Authorization determines what that user can do. Accounting monitors the duration and other details of network usage for billing purposes.

AAA information is stored in an external database or remote AAA server such as a RADIUS or TACACS+ server, or locally on the access server or router. RADIUS and TACACS+ servers assign specific privileges to users by associating attribute value (AV) pairs, which define the access rights for a user.

The figure shows how AAA authenticates and maintains accounting records for a dial-up PPP user. A user dials a telephone number that corresponds to a port on a network access server (NAS) at the edge of the network. The server checks the user's ID and password in the AAA server's local database or external authentication server such as Windows NT/Active Directory to determine whether to permit or deny access to the network (authentication). If the user is permitted, the AAA server typically sends a configuration or AV pair to the NAS, which determines the type of access and service granted (authorization). The accounting function of AAA logs the user's connect and disconnect times, the duration of the session, number of packets transmitted, and other details that are useful for usage and billing purposes.

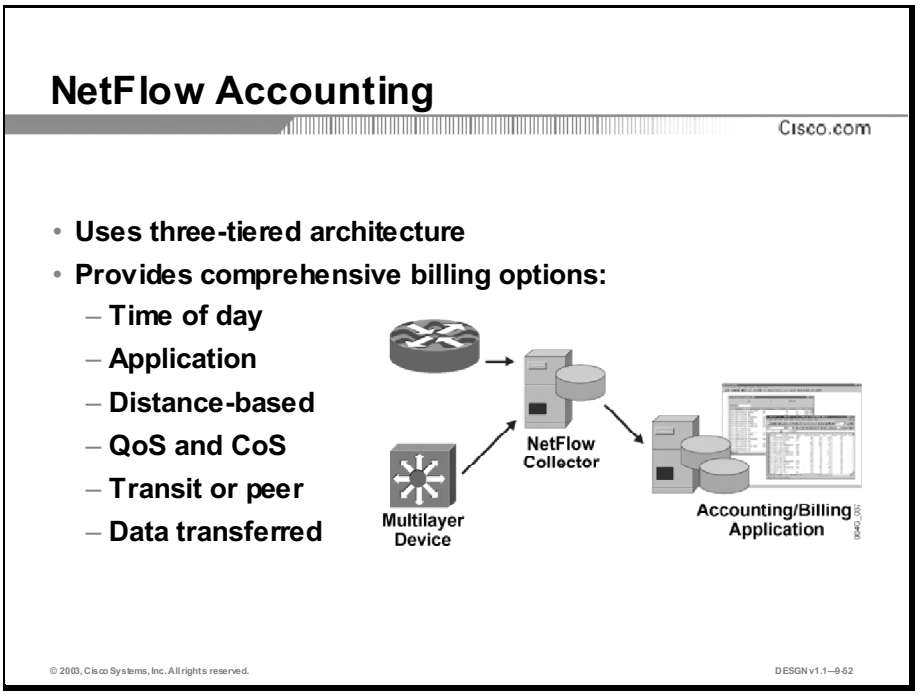

NetFlow services provide network administrators with access to IP flow information from their data networks. You can use exported NetFlow data for a variety of purposes, including network management and planning, enterprise accounting and departmental cost charges, Internet service provider (ISP) billing, data warehousing, and data mining.

The figure shows the three-tiered architecture of NetFlow service:

- **NetFlow-capable Layer 3 (L3) device:** Captures NetFlow accounting statistics for unicast ingress traffic on networking devices and exports the data to a collection device.
- **NetFlow collector:** Provides scalable and economical data collection from multiple NetFlow-enabled devices. The NetFlow collector provides these functionalities:
  - Consumes flows from multiple NetFlow-enabled devices
  - Reduces data volume through selective filtering and aggregation
  - Stores flow information in flat files on a disk for post-processing by NetFlow data consumers, third-party billing applications, and traffic analysis tools
- **Accounting and billing application:** Enables the network administrator to retrieve, display, and generate bills based on NetFlow data collected from NetFlow collector files.

NetFlow accounting and billing options provide detailed metering for highly flexible and detailed resource usage accounting. Service providers may use this information to migrate away from single-fee, flat-rate billing to more flexible charging mechanisms based on time of day, bandwidth usage, application usage, QoS, and so on. Enterprise customers may use the information for departmental cost recovery or cost allocation for resource usage.

Performance Management

Capacity planning is the process of identifying the network resources required to ensure that business-critical applications operate correctly. Performance management is the practice of managing network service response time, consistency, and quality for individual and overall services. This topic describes performance management.

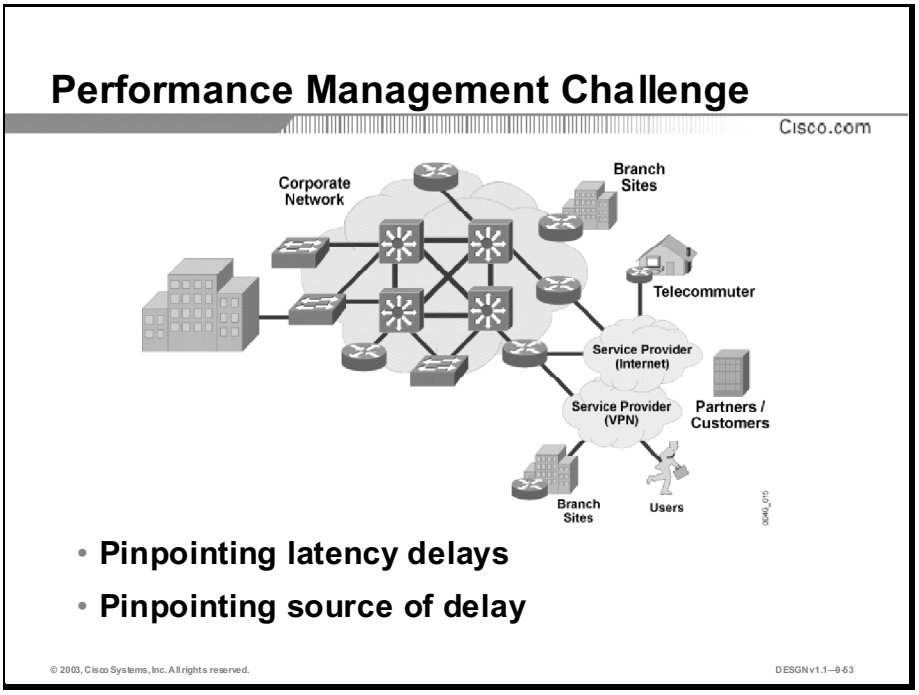

The goal of performance management is to guarantee internetwork performance at an acceptable level. Performance problems are usually related to capacity.

Some references to capacity planning and performance management mention the “data plane” and the “control plane.” The data plane refers to capacity and performance issues when data traverses the network. The control plane refers to resources required to maintain proper functionality of the data plane. Control plane functionality includes service overhead such as routing, spanning tree, interface keepalives, and SNMP device management. These control plane requirements use CPU, memory, buffering, queuing, and bandwidth just like the traffic that traverses the network. Many of the control plane requirements are essential to the overall functionality of the system. The network may fail if the control plane requirements do not have the resources they need.

The ability to measure network response time, determine device availability, resolve connectivity issues, analyze response time patterns, and provide critical reports (both real-time and historical) are increasingly important network management tasks.

Network managers require tools to isolate performance problems, locate bottlenecks, diagnose latency, and perform trend analysis in multiprotocol networks. Efficient diagnostic capabilities can lead to higher network availability by allowing network managers to resolve performance bottlenecks quickly.

Some challenges facing network managers include these:

- **Pinpointing network response time and availability problems:** Network managers need to be proactive, not reactive, identifying network delays and where delays reside within a network path. To diagnose problems quickly, managers need performance measurements for an entire path and for each hop within the path.
- **Pinpointing the source of the problem:** Different applications, hosts, network links, or networking devices within an organization may be responsible for poor performance of network applications and services. Identifying the source of a problem may be different in each case.

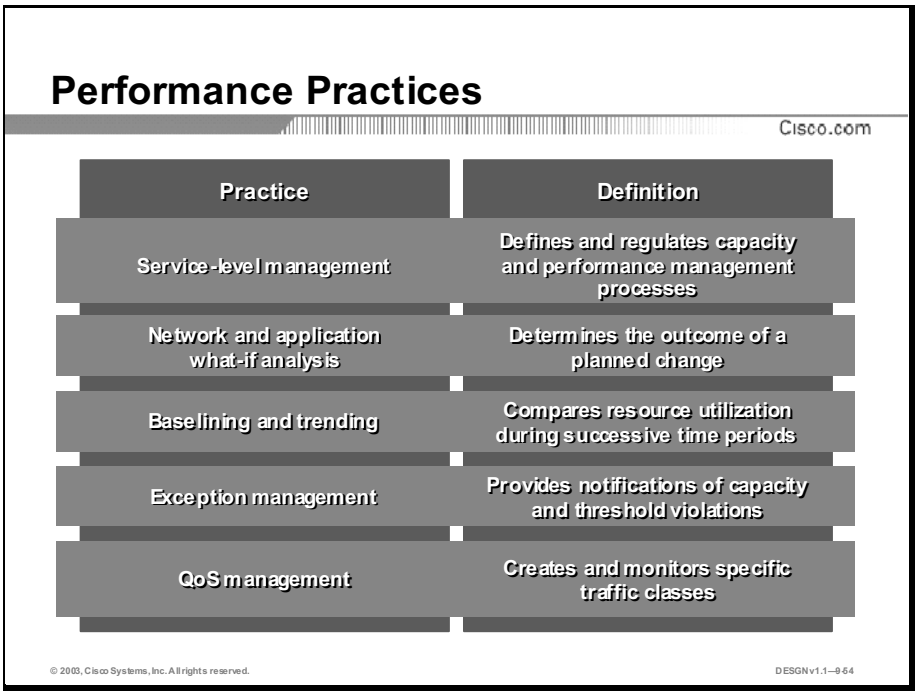

Here are some best practices for performance management:

- **Service-level management:** Service-level management (SLM) defines and regulates capacity and performance management processes. With the SLM methodology, network users and network service providers establish a service level agreement (SLA) that defines capacity and performance management. The SLA includes reports and recommendations to maintain service quality.
- **Network and application what-if analysis:** A network and application what-if analysis determines the outcome of a planned change. Without what-if analysis, organizations risk overall network availability and overall success. You can use automated tools to simulate a change on the network.
- **Baselining and trending:** Baselining and trending enable network administrators to plan and complete network upgrades before a capacity problem causes network downtime or performance problems. This process compares resource utilization during successive time periods or over time. Therefore, baselining and trending allow planners to view resource utilization for any specified time period. An issue with baselining and trending in large networks is the overwhelming amount of information. Many network management solutions provide information and graphs on capacity resource variables. Unfortunately, managers often use this data to provide reactive support rather than proactive support for potential problems, which is the purpose of baselining and trending.
- **Exception management:** Exception management identifies and resolves capacity and performance issues. The NMS notifies the network manager of capacity and performance threshold violations, enabling immediate investigation. For example, a network manager might receive an alarm about high CPU utilization on a router. The network administrator can log into the router to identify and resolve the problem.

- **QoS management:** QoS management involves creating and monitoring specific traffic classes. Network managers typically create traffic classes based on performance SLAs for business-critical applications and delay-sensitive applications such as voice. QoS management supports traffic class configuration and QoS mechanisms on network devices to maximize the performance of critical applications. QoS management also involves active monitoring of application performance to validate the QoS configuration and ensure that required service levels are met.

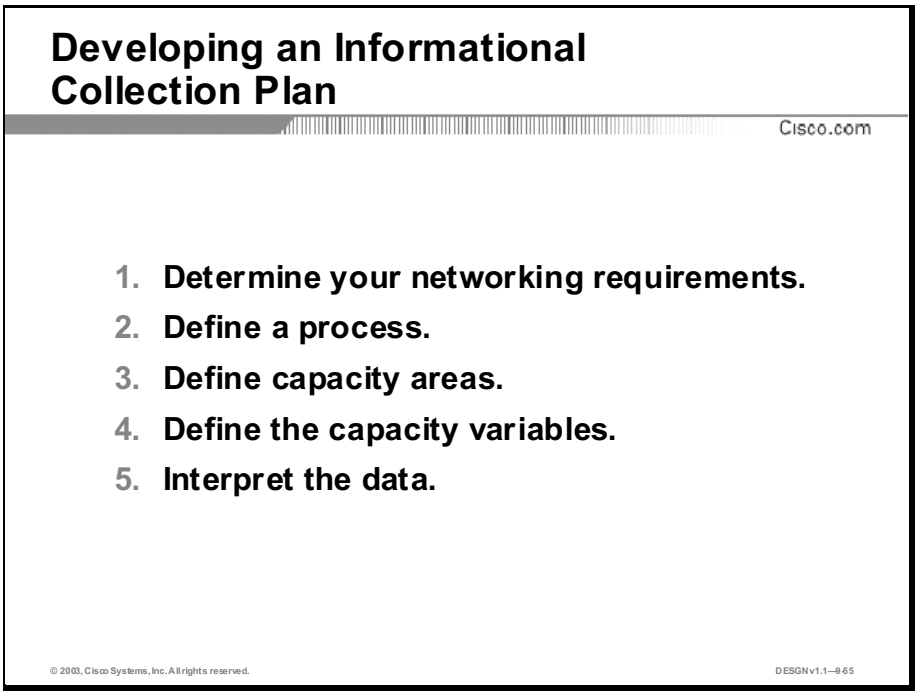

To collect capacity information, you need to follow the steps outlined in the table.

Step	Task	Description
1	Determine your networking requirements.	Identify the information required and its purpose. Determine what resources and tools are available, what gaps exist, and the tasks required to provide correct performance monitoring.
2	Define a process.	Define a process to ensure that network administrators use performance measurement tools successfully and consistently. Define how network administrators should react when threshold violations occur, and processes to follow for baselining, trending, and upgrading the network.
3	Define capacity areas.	Capacity areas define parts of the network that share a common capacity planning strategy; for example, the corporate backbone, remote offices, and critical WAN sites. Defining different areas is helpful for these reasons: ■ Different network areas may have different thresholds. For example, LAN bandwidth is less expensive than WAN bandwidth, so utilization thresholds could be lower. ■ Different areas may require monitoring different MIB variables. For example, forward explicit congestion notification (FECN) and backward explicit congestion notification (BECN) counters are critical in understanding Frame Relay capacity problems.
4	Define the capacity variables.	Define capacity variables based on the type of devices, media, and technology used within the network. Parameters such as CPU, memory, and link utilization are valuable. Queue depths, Frame Relay congestion notification, backplane utilization, buffer utilization, NetFlow statistics, broadcast volume, and RMON data may also be important.
5	Interpret the data.	Interpret the collected data to provide a high-quality service. For example, many organizations do not fully understand peak and average utilization levels.

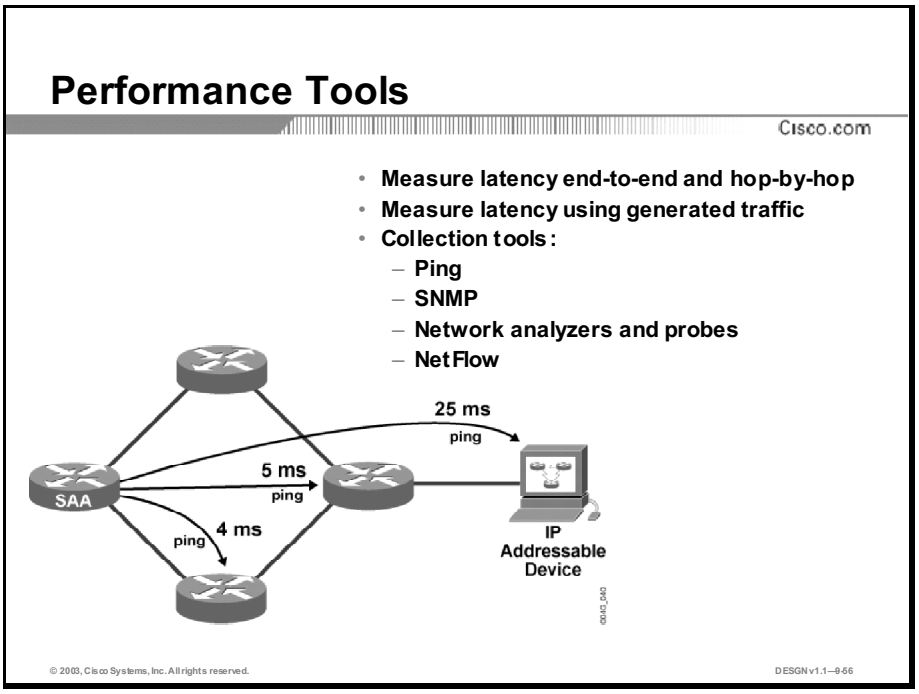

To find the location where network performance is lacking, you need to measure response time and availability end-to-end and hop-by-hop (network device-to-device) to isolate the trouble spots in the network.

Performance tools, such as Cisco's Service Assurance Agents (SAAs), embedded in Cisco routers or Catalyst Layer 3 switches, and Internetwork Performance Monitor (IPM), help network managers trace and identify performance degradation in a network. No dedicated hardware probe is needed to measure and monitor network performance statistics. It is important to measure the performance of business application traffic directly, that is, in the path that the data actually travels. Measuring the delay of voice data and other upper-layer protocols, such as TCP, UDP, DNS, and DHCP, can provide important information for optimizing a network.

Here are some performance collection tools:

- **Ping:** Ping is a utility that uses the Internet Control Message Protocol (ICMP) to test for basic connectivity to an IP address. Use ping to quickly determine a device's reachability and response time from the station to the target IP device.
- **SNMP:** SNMP managers periodically poll devices for utilization and error information that you can use to determine device-specific issues impacting overall network performance. For example, using SNMP to determine that the bandwidth utilization on a router's WAN interface is over 95 percent may help identify an existing or potential bottleneck.
- **Network Analyzers or Probes:** Network analyzers or probes monitor and troubleshoot a network segment's performance. For example, an RMON2 probe can analyze the existing network traffic and report on the connected segment's utilization, including major users and conversations broken down by upper-layer protocols. Probes can capture packets and analyze packet header information for an in-depth analysis of a network segment's activity. Dedicated probes are well suited to baselining or trending a network link, protocol, and application utilization, detecting error rates beyond a set threshold, and characterizing and identifying the major users and conversations.

- **NetFlow:** By analyzing NetFlow data, you can identify the cause of congestion, determine the class of service for each user and application, and identify the source and destination network for the traffic. NetFlow provides extremely detailed and accurate traffic measurements and high-level aggregated traffic collection.

Security Management

Security management aids administrators in creating a secure network environment. Key tasks include partitioning network resources into authorized and unauthorized areas, mapping groups of users to those areas, and monitoring, policing, and logging user access to resources in those areas. This topic describes security management.

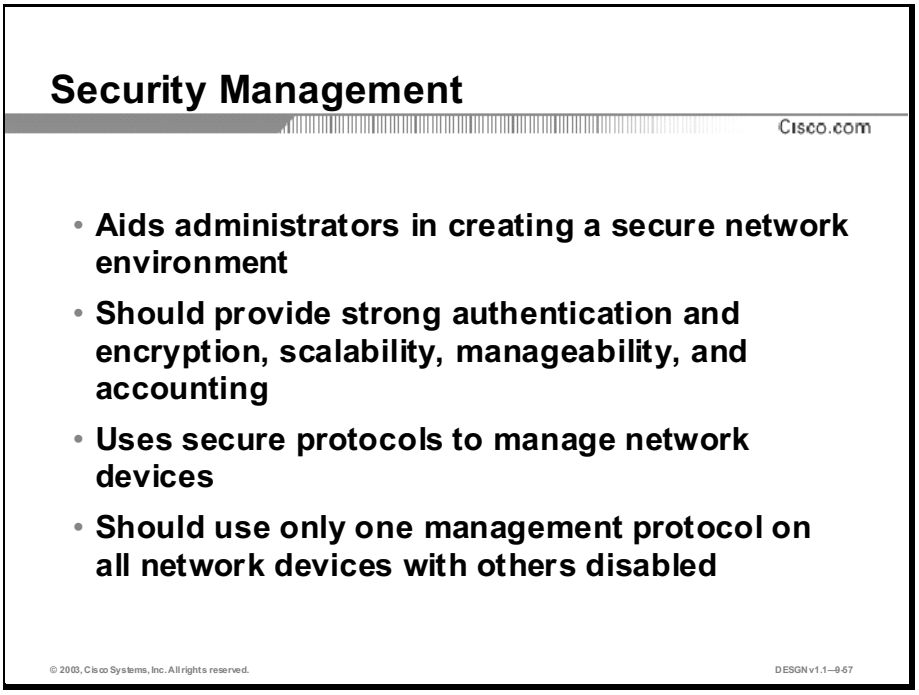

The purpose of security management is to support the application of security policies, including: the creation, deletion, and control of security services and mechanisms; the distribution of security-relevant information; and the reporting of security-relevant events. Security management controls access to network resources, and prevents network sabotage (intentional or unintentional) and unauthorized access to sensitive information.

Protocols that you can use for remote network device management include these:

- **Telnet:** Typically used to change network device configurations. Authentication is in clear text and, therefore, not secure. As an alternative, use Secure Shell (SSH) or encryption using IP Security (IPSec) to increase security. Offload authentication to a centralized authentication server that uses a One Time Password (OTP) approach to increase the level of security.
- **SNMP:** Typically used to manage and monitor network devices. Different SNMP versions offer different security mechanisms and levels of security.
- **HTTP:** Web-based technology used to monitor network devices. Use an authentication server such as AAA with an optional OTP system or IPSec for confidentiality. If available, use Secure HTTP (HTTPS) with Transport Layer Security (TLS) for secured web server access.
- **Remote shell protocol (RSH):** Used to modify configurations but not recommended because of its lack of security.

- **SSH:** A protocol that provides strong authentication and encryption of the management session but may not be supported on all devices.

To ensure enterprise network security, first secure the network devices. The default authentication is based on password-only authentication, which is shared among multiple administrators so you cannot identify who configured (or misconfigured) a device.

The solution is to use a username and password pair to identify individual administrators and a centralized user database to make security management more scalable. Cisco network devices support centralized security management using AAA.

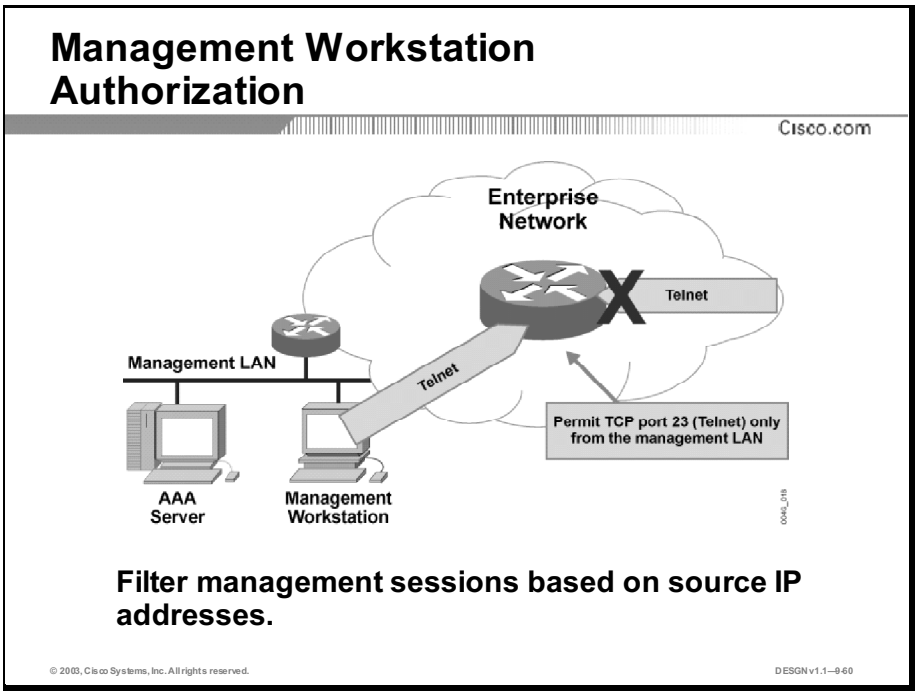

The figure illustrates the “first security level” where management sessions are authorized based on source IP addresses. The network device applies an access list to incoming packets on interfaces or to the vty lines. Subsequent authentication defends against IP spoofing. You can use IP filters to allow management access only from a designated management station on a secure management LAN.

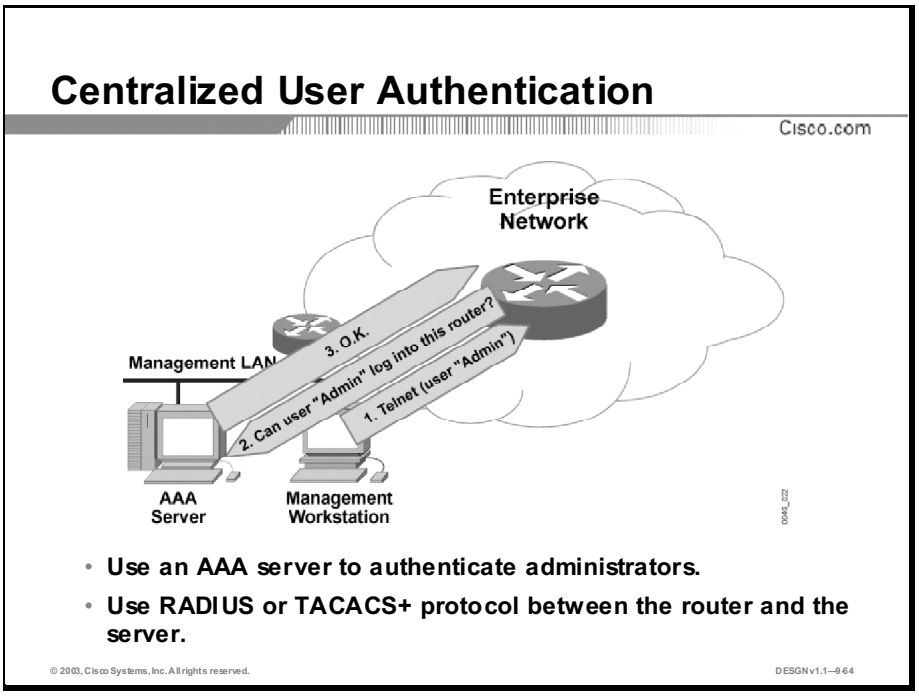

The figure illustrates the steps in the authentication process when an AAA server stores the user database. When a user starts the management session (Telnet in the example), the router asks the AAA server (by using the TACACS+ or RADIUS protocol) if this user has permission to manage this router. The AAA server allows the router to grant access if the user's credentials (username and password) match an entry in the user database.

Use an additional server that verifies OTPs to strengthen authentication.

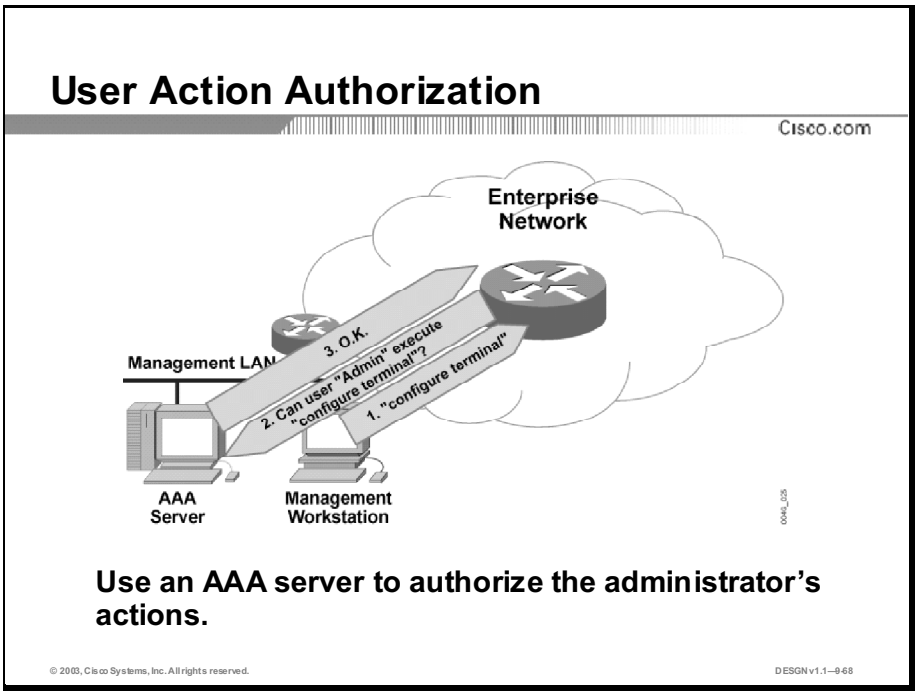

The figure illustrates the next level of security where the AAA server authorizes each user action. This allows a network administrator to control the actions that users can perform on network devices.

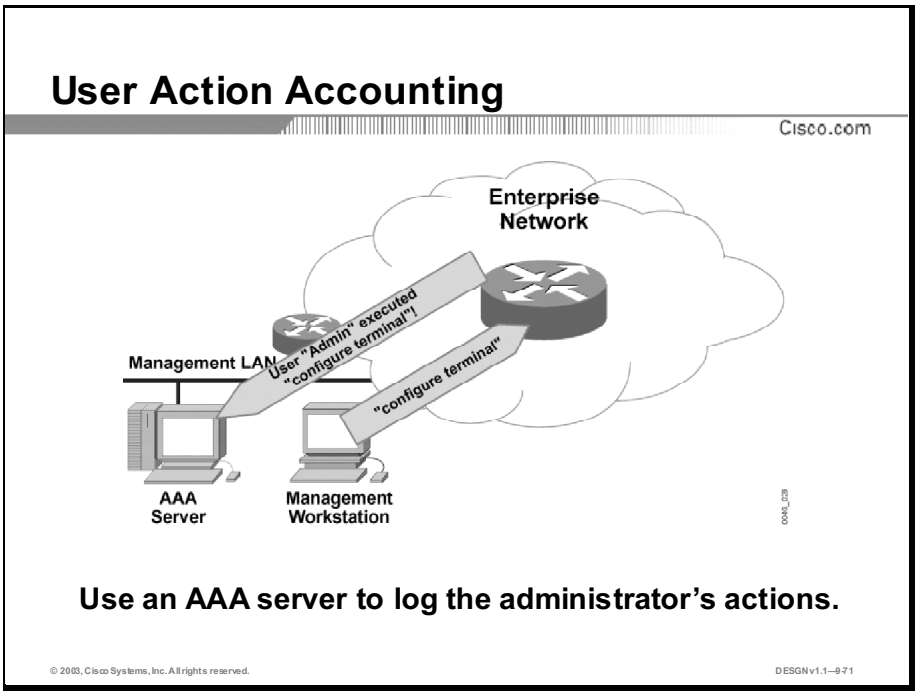

The last step is the accounting of all actions. This step is important for tracking changes to the network. Authorization identifies which commands users can use, while accounting logs which commands users have misused.

Summary

This topic summarizes the key points discussed in this lesson.

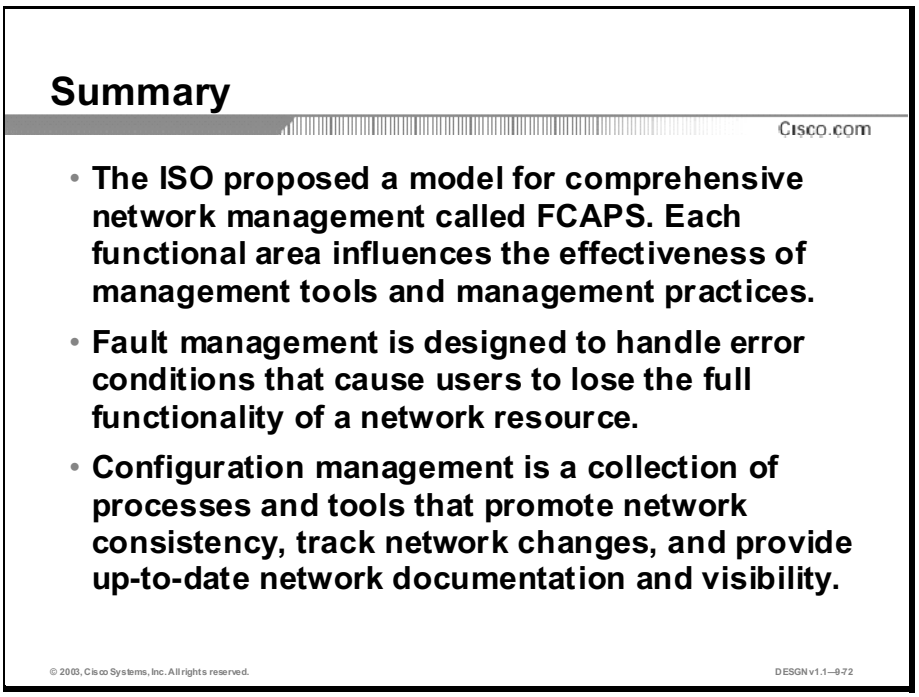

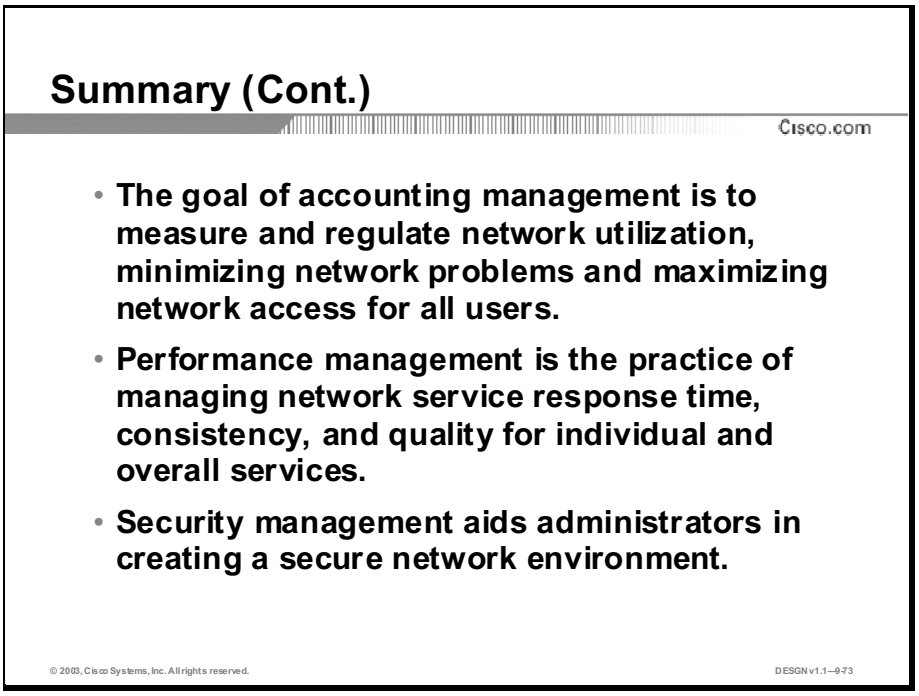

References

For additional information, refer to this resource:

- Technical Assistance Center (TAC), http://www.cisco.com/en/US/support/index.html

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Match each FCAPS functional area to the function it performs:

A) fault management

B) configuration management

C) accounting management

D) performance management

E) security management

_____ 1. monitors and measures various aspects of performance so that overall performance can be maintained at an acceptable level

_____ 2. provides information on network resource usage

_____ 3. detects, isolates, notifies, and corrects faults encountered in the network

_____ 4. provides access to network devices and corporate resources to authorized individuals

_____ 5. manages the configuration of network devices, such as file management, inventory management, and software management

Q2) Place the five basic steps of fault management in the order in which you would perform them.

_____ 1. Diagnosis

_____ 2. Resolution

_____ 3. Fault determination

_____ 4. Bypass and recovery

_____ 5. Fault tracking and control

Q3) Without configuration management, which four issues are you likely to experience? (Choose four.)

A) increased number of errors on the network

B) increased time to resolve problems

C) higher network costs because of inefficiently used network components

D) lower network performance

E) inability to determine the impact of network changes or outages on the end user

F) increased reactive support issues

Q4) The first step toward appropriate accounting management is to measure the _____ of all important network resources.

A) configuration

B) error rate

C) utilization

D) performance

Q5) The goal of performance management is to guarantee internetwork performance at an acceptable level. Performance problems are usually related to _____.

A) capacity

B) device errors

C) poorly designed databases

D) configuration errors

Q6) To ensure enterprise network security, you should first secure _____.

A) host devices

B) servers

C) gateways

D) network devices

Quiz Answer Key

Q1) A=3, B=5, C=2, D=1, E=4

Relates to: FCAPS Functional Model

Q2) 1=2, 2=4, 3=1, 4=3, 5=5

Relates to: Fault Management

Q3) B, C, E, F

Relates to: Configuration Management

Q4) C

Relates to: Accounting Management

Q5) A

Relates to: Performance Management

Q6) D

Relates to: Security Management

Managing Service Levels in a Network

Overview

Information technology providers are under increasing pressure to offer SLAs to their customers. Whether the service provider is the enterprise IT department or a third-party vendor such as an ISP, management needs contractual assurance that the provider will meet its business objectives. For the end user, the SLA provides assurance that critical network applications and services will be available when needed. Service levels can be divided into various components. Service level management combines SLAs with common management solutions, leading to comprehensive SLA measuring and reporting.

This lesson defines an SLA, describes its importance, defines its target customers, and discusses the typical requirements found in the SLA. The lesson concludes with an introduction to monitoring network response time and availability as a way to monitor SLAs.

Relevance

If you are involved in designing SLM solutions, this lesson helps identify the key concerns you will consider.

Objectives

Upon completing this lesson, you will be able to describe the management of service levels in a network. This includes being able to meet these objectives:

- Describe the need for SLAs
- Describe the requirements for using SLAs in network monitoring
- List the issues related to end-to-end SLM
- Identify the measuring methods SAA provides
- Identify the applications available for monitoring SLMs using network response and availability applications

Learner Skills and Knowledge

To benefit fully from this lesson, you must have these prerequisite skills and knowledge:

- Basic understanding of network management concepts and protocols

Outline

The outline lists the topics included in this lesson.

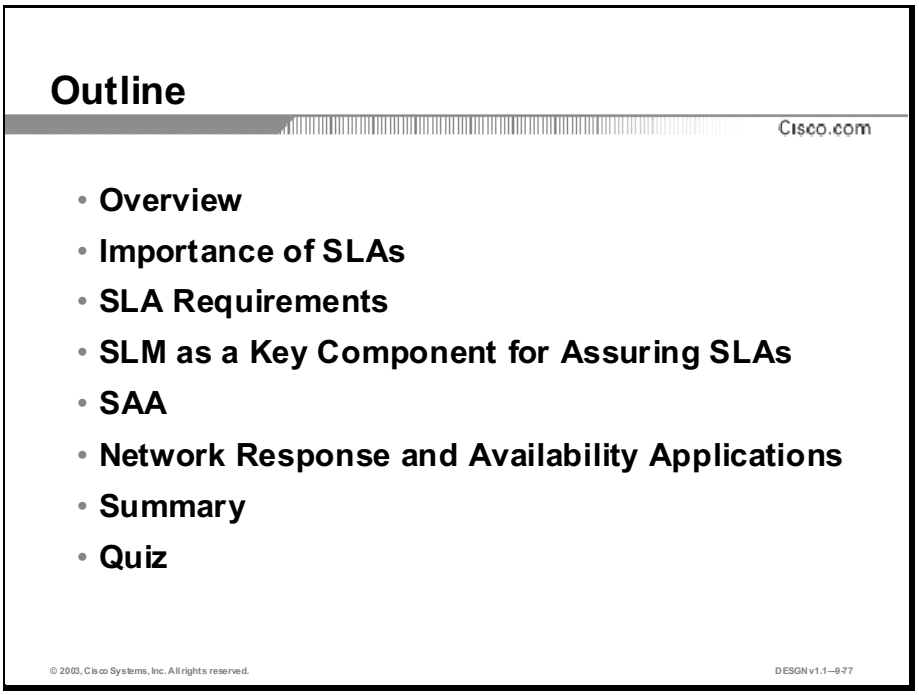

Importance of SLAs

An SLA is a key component of a service level contract (SLC). The SLC specifies connectivity and performance levels that a provider supplies for its customer. The service provider could be within the enterprise, for example, an IT department providing services to internal network users, or an external company such as an ISP providing wide-area or hosted application services. This topic introduces SLAs.

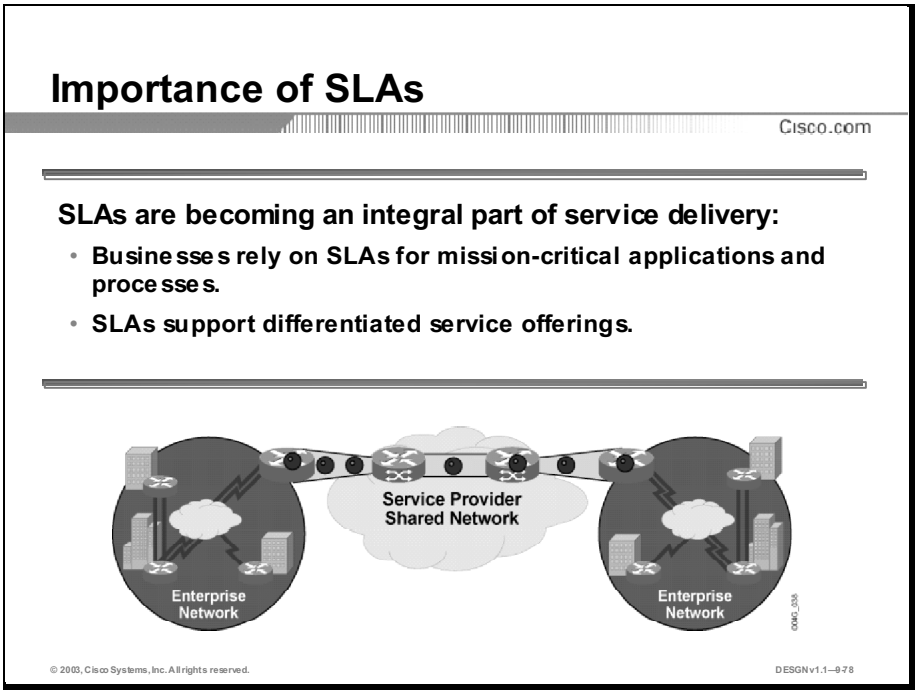

An SLC typically includes multiple SLAs. A violation of any particular SLA could create a violation of the overall SLC. The SLM solution must provide a way to manage all agreements that constitute a contract with the service provider. The SLM solution should enable the user to monitor multiple SLCs individually, examine SLA details, and monitor SLA conformance for a given SLC.

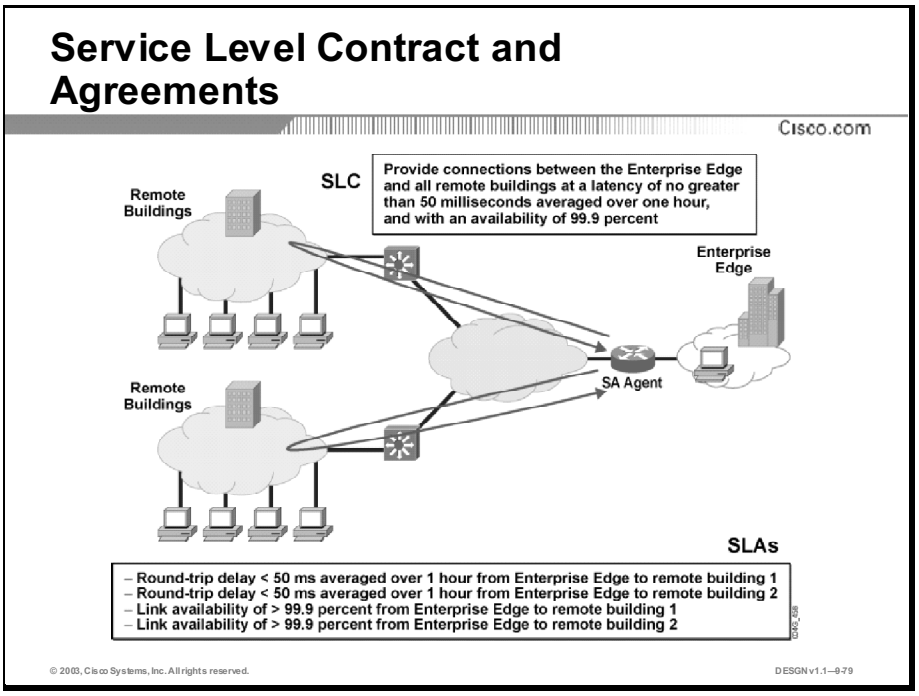

SLAs measure specific service performance between device pairs. Device pairs can include routers, servers, workstations, and other equipment. For example, an SLC for connectivity from remote building sites (numbered 1 and 2) to the Enterprise Edge may read, “A connection of 100 Mbps at a latency of no greater than 50 milliseconds averaged over one hour, and with an availability of 99.9 percent, is to be provided.” These are some supporting SLAs:

- Round-trip delay of no more than 50 ms averaged over one hour from remote building 1 to the Enterprise Edge
- Round-trip delay of no more than 50 ms averaged over one hour from remote building 2 to the Enterprise Edge
- Link availability of no less than 99.9 percent from remote building 1 to the Enterprise Edge
- Link availability of no less than 99.9 percent from remote building 2 to the Enterprise Edge

SLA Requirements

Organizations or departments need to monitor SLC compliance. Enterprises need to implement procedures to confirm that service levels are being met and not rely solely on reports generated by external service providers for confirmation. This topic describes SLA requirements.

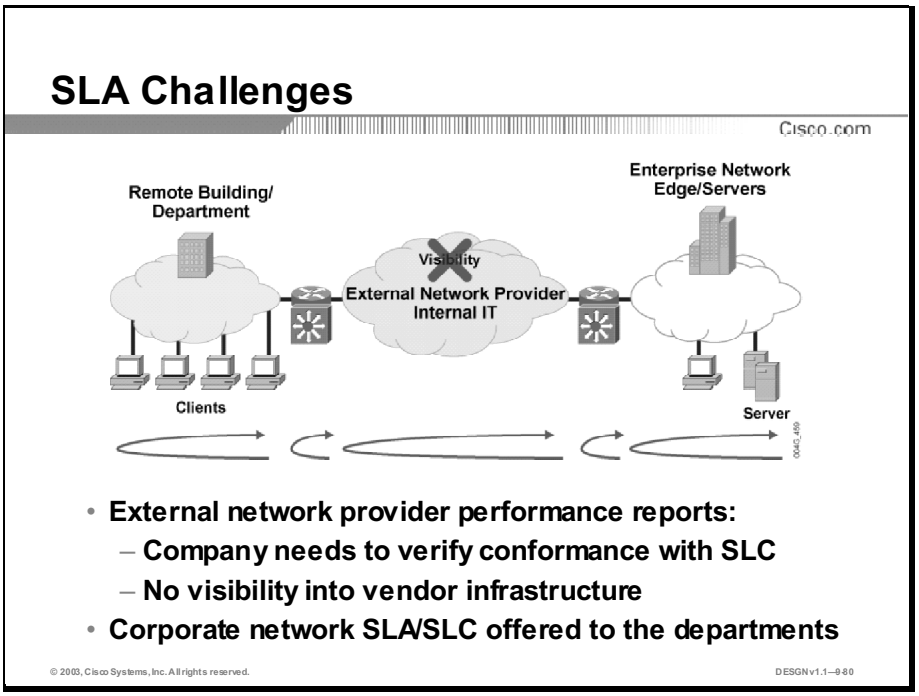

Monitoring a third-party vendor's provided services presents a unique challenge because individual components and data are transparent. To determine conformance you must collect service-level measurements from customer-owned and controlled devices. The collected data must accurately represent and report on SLA conformance. The network manager must be able to view the service from an end-user perspective and locate the site where the problem is occurring.

In a corporate network, the department receiving the service will sign an SLC with the corporate IT department, which controls and manages the enterprise network. The IT department must have proactive traffic measurement tools that determine when the terms of the SLC are being violated, enabling the IT department to respond with corrective action.

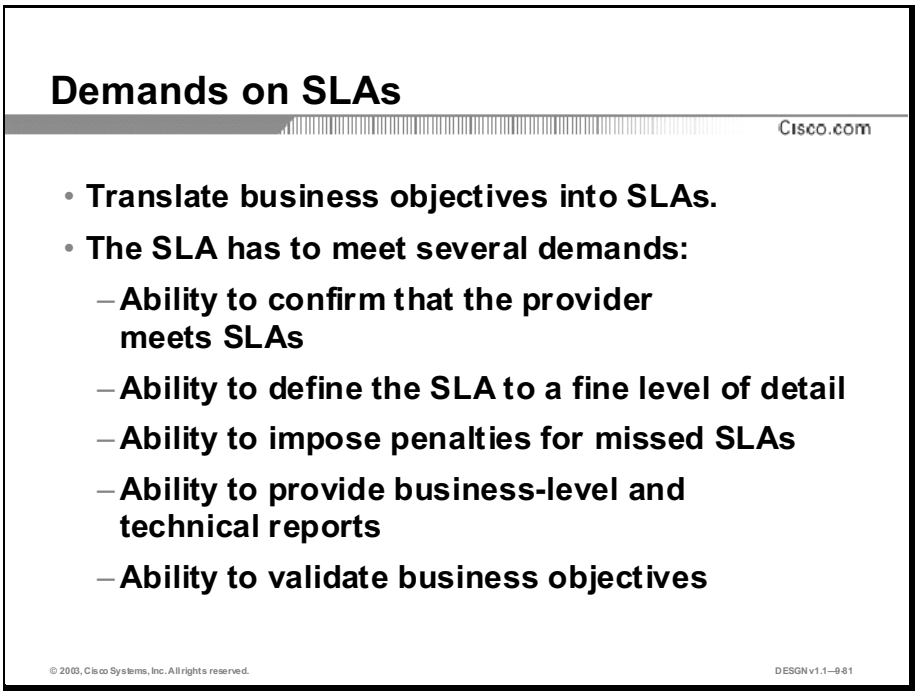

A key success factor in developing effective SLAs is to translate business objectives accurately into SLAs, with tangible service metrics that you can measure, report, and validate. Long, complex, and unrealistic agreements are difficult to manage.

The enterprise customer normally requires:

- **The ability to confirm that the provider meets SLAs, based on network connectivity and network application responses.** For instance, IP telephones may require low latency and low delay with other IP telephones in the corporate network. Therefore, an enterprise would require an SLA with particular latency and delay parameters to suit the IP telephones' communication demands.
- **The ability to identify SLAs to a fine level of detail.** In the event of an SLA violation, department managers need to know where and why the SLA was broken. For example, if an IP telephone call has poor voice quality, the department manager needs to know if the cause is related to the network or to the IP telephone system.
- **The ability to impose financial penalties for missed SLAs.** If an enterprise cannot conduct its business because of unavailable services, penalties are required. If the correct SLA measurements are applied, you can track the location of the faulty equipment and the time of the event.
- **The ability to provide business-level reports and detailed technical reports that a network manager can easily enhance and extend.** Reports identify the time of an SLA violation and show trends in network performance.
- **Validation that key business objectives are being met when deploying technologies.** When networks are upgraded, the SLA plays an important role in monitoring the performance of newly introduced technology.

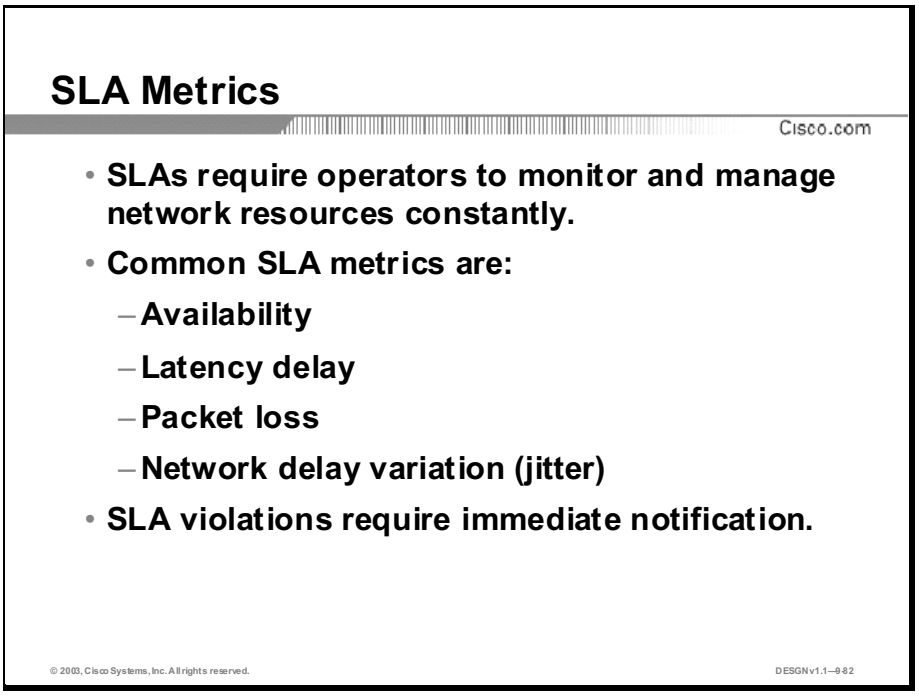

SLA customers need service management solutions that provide a method to validate that service levels are being met. Within an organization, departments using network services may require an SLA with the IT department to support their business-critical traffic. Each service may place different demands on the network. These varying network demands require a variety of measurements such as these:

- **Availability:** Constant device polling identifies device availability. Business-critical applications require constant availability of network devices and services. If the business-critical application cannot communicate, you must locate the point of failure. An obvious cause of broken communication could be services or devices that are unavailable and not responding.
- **Network delay:** Polling records the amount of time a device takes to send and receive a packet. High delay numbers can indicate congestion on an end device, an intermediate device such as a router or switch, or a link along the connection path.
- **Packet loss:** Packet loss can indicate a bottleneck on the network. With hop-by-hop polling, you can isolate the source of packet loss.
- **Network delay variation (jitter):** Jitter on the network is particularly critical with voice and video, which require a constant data stream on the network. Jitter causes voice applications to become garbled, making communication impossible.

Delivering the requested SLA means that the network operator must complete these tasks:

- Actively monitor the network to verify SLA implementation and make necessary changes to the network. The changes may only require minor adjustments to the network configuration, such as changing the policy-based parameters. However, network design changes may be required if the network runs low on resources. Proper measurement tools enable the network operator to detect, in advance, when and where the network will congest.
- Maintain proactive notification systems. Because of unpredictable traffic flows in the network, monitoring and periodic report generation of SLA performance is insufficient. Proactive notification informs network operators when network performance is not within SLA parameters.

SLA metrics vary widely, and the technologies for measuring those metrics vary, too. In managing service levels end-to-end, it is necessary to measure and collect data at every level of the protocol stack—from the Level 1 and Level 2 LAN/WAN layers, to the Level 3 and Level 4 network and network services layers, to the Level 5, Level 6, and Level 7 client-to-server layers.

SLM as a Key Component for Assuring SLAs

Guaranteeing a specific SLA requires an SLM agreement where a given service has several service components. This topic describes service-level management.

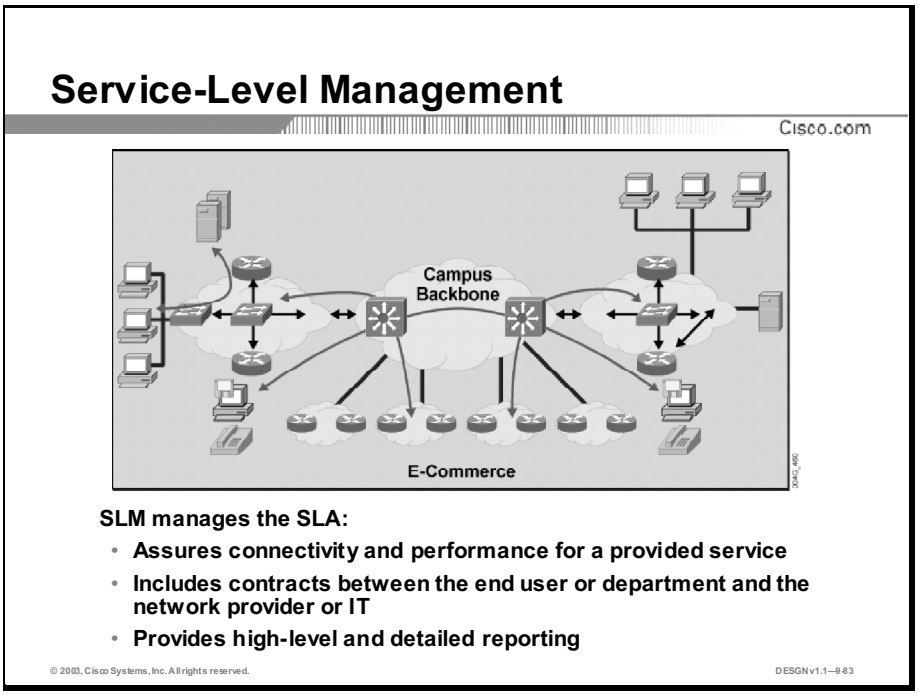

While service providers may offer different methods to manage service levels, only one definition and capability is meaningful to customers: end-to-end management of all aspects that relate to the connectivity, performance, and availability of the service or application.

An SLA contract ensures that the network provider will satisfy end-user requests. For example, corporation A may require many SLAs from their network provider, each for a different service:

- Low latency and jitter for IP telephony in their complete network
- High availability for users of the business application servers from inside and outside the corporate boundary
- High availability of business application services; for example, HTTP

End-to-End SLM Requirements

An end-to-end SLM must:

- Leverage component management products from multiple vendors
- Work with clients and equipment the customer does not own or control
- Adapt to new SLM metrics as new technologies are deployed
- Scale by orders of magnitude
- Collect the correct network and application SLM metrics at the appropriate times

High-Level and Detailed Reporting

An SLA application should measure data against the requested metrics, print basic SLA reports, such as whether or not an SLA has been met, and provide network administrators with high-level management reports and detailed reports.

When an SLA includes network management, enterprises translate business objectives to service provider agreements. Department managers must be aware of their business needs so that they can negotiate an SLA with the network provider. At the same time, the reports from the provider must be detailed enough so the department managers understand how the SLAs are met.

The SLA management system must provide the level of technical information that the day-to-day operations staff needs to help them resolve problems when SLAs are not met.

High-level reports include information such as these:

- SLCs that are not in conformance
- Percentage of SLAs that are out of conformance for a given contract
- Business objects that are being impacted
- Escalation procedures to resolve nonconformance of SLCs
- Performance of service levels over selected time periods

Detailed reports provide information such as:

- Device pairs related to an SLA
- Exceptions that occur and the frequency
- Round-trip latency and latency variations
- Mean time to repair and mean time to fail
- Link and device availability
- Throughput

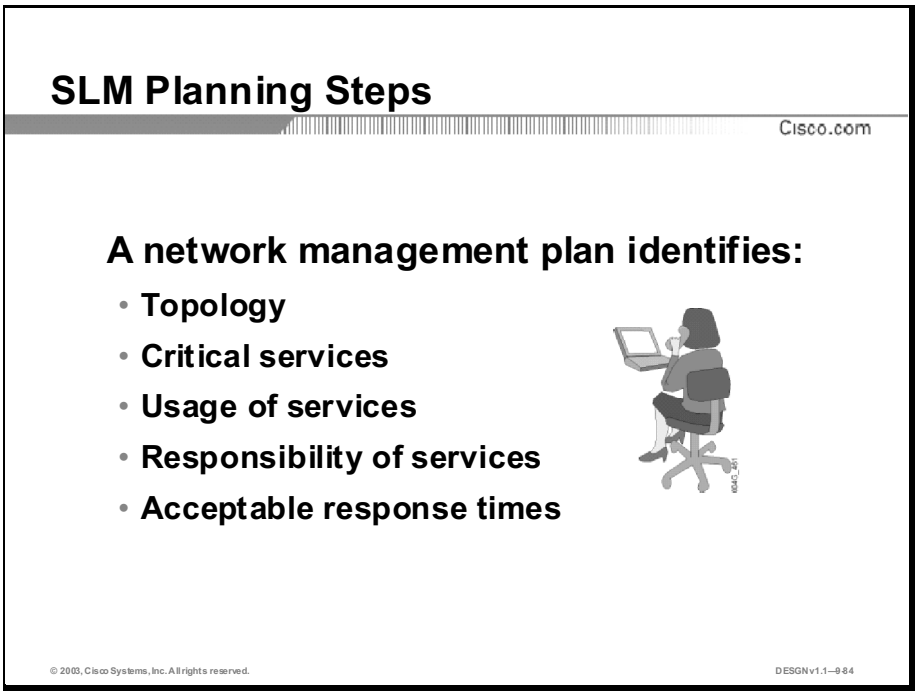

Defining a network management plan helps you understand business services and the network segments that these services traverse. The network management plan should include these things:

- A description of the network topology, its components, and the services the network can handle.
- A list of business-critical services to assist the administrator in choosing the operations that proactive monitoring agents are to emulate and measure.
- A description of the end users for the services and where they are located. Understanding the topology and the users of critical networked applications will help administrators choose the source routers and target devices.
- An explanation of who is responsible for the network upon which the services run. Businesses rely on service providers to link people and information sources in different locations. Outsourcing makes it difficult to troubleshoot poor response time and network connectivity, increasing the need for proactive performance tools.
- A description of the minimum acceptable response times for the protocols that these services use. If the minimum acceptable response time is unknown, use proactive performance tools to establish baselines on a monthly or weekly basis. Network designers often define acceptable performance, which does not always match end-user expectations.

SAA

The Service Assurance Agent (SAA), previously known as Response Time Reporter (RTR), is an IOS software feature on some IOS platforms. SAA allows enterprises to monitor network performance between a Cisco device and a remote device, which can be another Cisco device or an IP host. This topic describes the SAA.

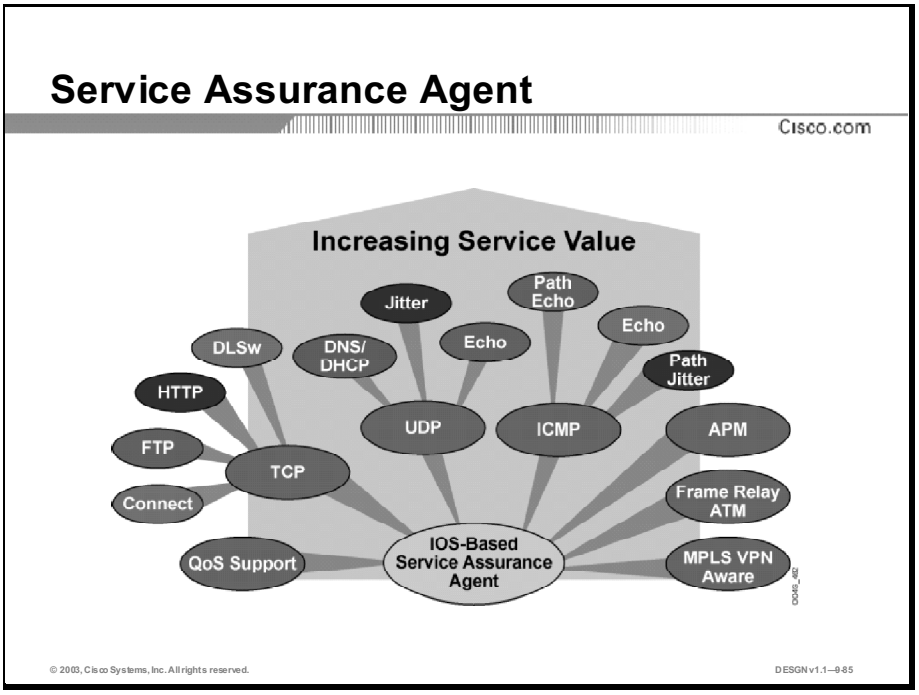

The SAA in Cisco IOS software allows users to perform troubleshooting, problem analysis, and notification based on the SAA statistics.

Here are some key capabilities of SAA:

- Accessibility using the IOS command-line interface (CLI) and SNMP
- Ability to define rising and falling thresholds to monitor SLAs, measure packet loss of SAA-generated packets, and notify the network management system
- Ability to measure ICMP, TCP, or UDP response time on a specific path
- Ability to measure response time between end points for a specific QoS, by using IP CoS bits
- Ability to measure voice application traffic response by using UDP jitter operations
- Ability to measure HTTP service performance including DNS lookup, TCP connect, and HTTP transaction time
- Ability to schedule an operation in the future and store historical data collected by SAA services

From the IP layer to the application layer, SAA can accomplish these tasks:

- Monitor file transfer performance using FTP
- Use the HTTP operation to monitor business-critical web sites to measure response time and availability of those sites
- Monitor the UDP jitter operation and store information such as jitter, round-trip delay, one-way latency, and packet loss
- Provide hop-by-hop measurements, for which the ICMP Path Echo operation is essential to isolate a troubled link quickly
- Generate traffic used for proactive measurements if QoS is implemented.

SAA provides a way to configure an IOS device to perform tests in the network to end systems or other IOS devices. The results of these tests are used to validate the SLA.

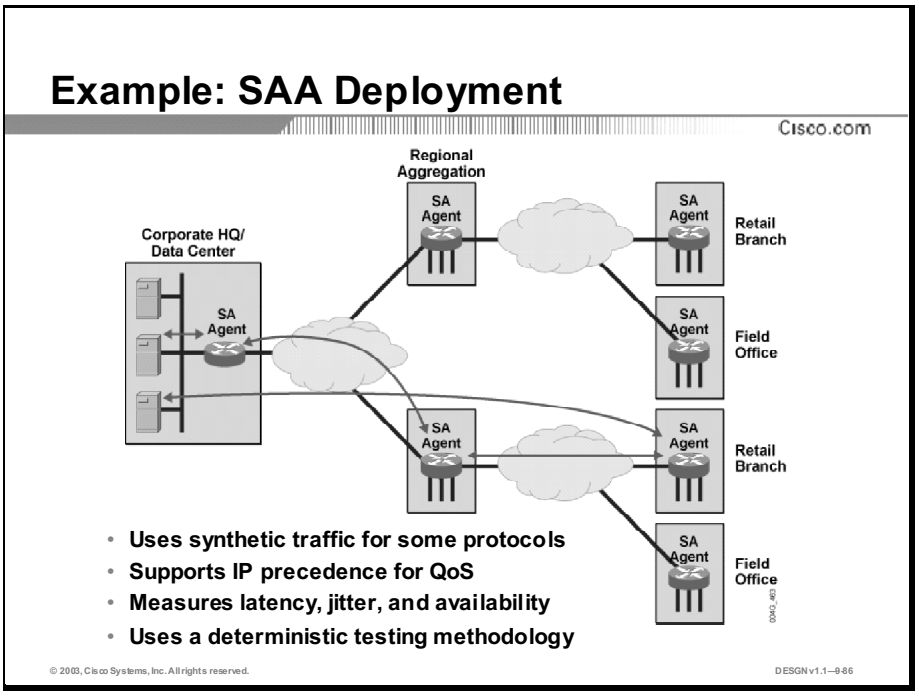

Organization ABC's network is organized into the corporate headquarters, regional aggregation, and retail branches networks. The company requires network device availability, as well as immediate notification if the end-to-end communication breaks or is not within the defined values. Additional requirements include availability of HTTP services in the campus Server Farm module and low jitter for IP telephony running in the campus.

Network managers can initiate network device availability testing from almost any device with the IP **ping** command. The network management station, which generates ping packets, can only test connectivity from that point of the network. The benefit of initiating a ping on network devices is that network managers can trigger it from any point in the network to any destination. SAA provides effective notification in the event a device is unavailable.

For measurement purposes, SAA emulates the role of an HTTP client, the source of web requests. The company placed another SAA close to the servers so it can monitor server availability. Reports from SAA-capable devices can identify where HTTP traffic is slow.

Voice traffic is assigned a higher priority to minimize delay and jitter. SAA measures the jitter for IP telephony by acting as both the client (source) and server (destination). Emulating the source of client requests, the source SAA sends test packets from the client side of the network. The destination SAA emulates the server function and accepts the packets, provides measurements, and responds back to the source SAA. The source SAA assesses the values in the reply packets and compares these values to the defined SLA. If the defined values are exceeded, the notification mechanism provides information to the management station, and the administrator can easily locate the source of problems.

Network Response and Availability Applications

Cisco and other vendors have extended the SAA capabilities to provide more measurements and more precise reporting. This topic describes network response and availability applications.

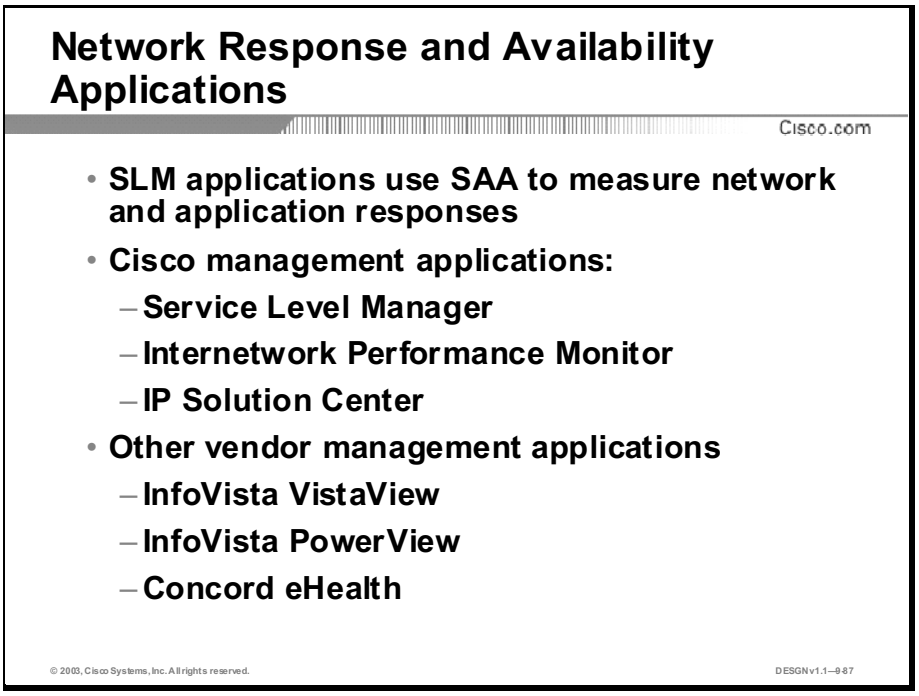

Here is a partial list of available management applications:

- Within Cisco, Service Level Manager and Internetwork Performance Monitor (IPM) are CiscoWorks applications that provide performance monitoring and SLM.
- IP Solutions Center (ISC) is a comprehensive solution for managed security and VPN services.
- Third-party management applications support SAA, including those provided by Concord, InfoVista, and Agilent. These applications receive measurement data from SAA, analyze the results, and present the performance statistics through various reports.

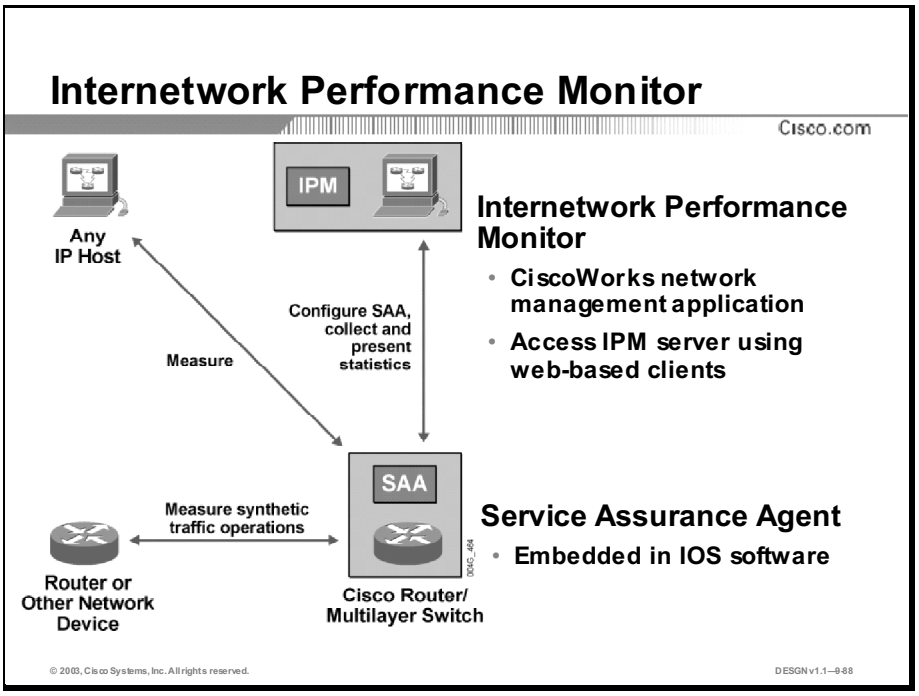

IPM monitors the performance of multiprotocol networks. IPM measures the latency and availability in IP networks on a hop-by-hop (router-to-router) basis and from the router to IP end station. It also measures latency between routers and mainframes in Systems Network Architecture (SNA) networks.

IPM can perform these tasks:

- Troubleshoots problems by checking the performance between devices
- Sends SNMP traps and SNA alerts when a user-configured threshold is exceeded, a connection is lost and reestablished, or a timeout occurs
- Analyzes potential problems before they occur by accumulating statistics, which are used to model and predict future network topologies
- Monitors latency, jitter, availability, and errors between two network endpoints

The IPM/SAA monitoring solution is composed of three parts: the IPM server, the IPM client application, and the SAA feature of the IOS software. The IPM network management application includes the server and the client.

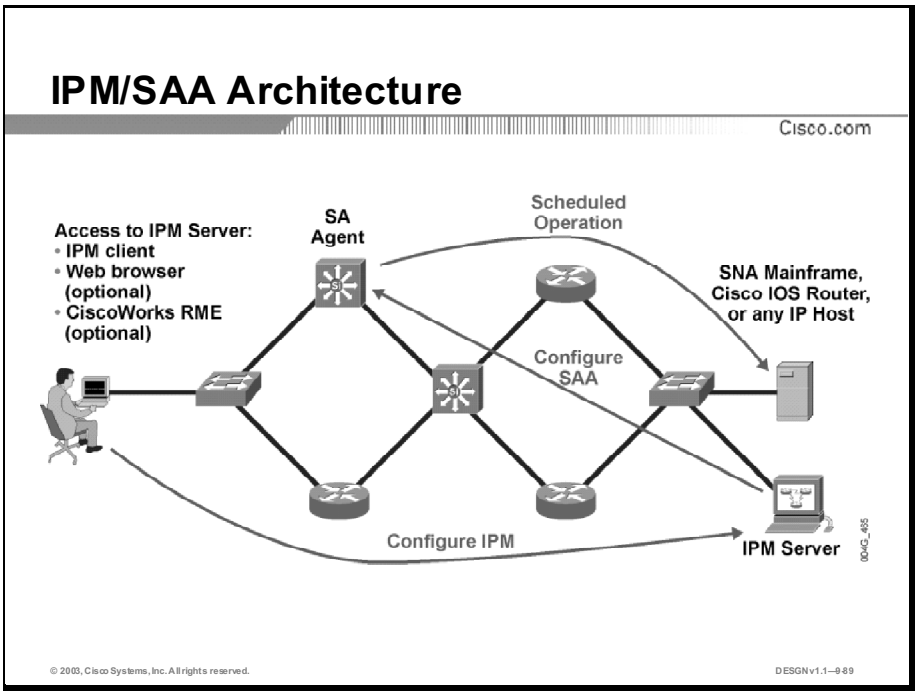

IPM is a client-server application. Using the IPM client, the user defines components such as the source routers where the measurements are to begin and identifies possible target devices and typical synthetic traffic operations in the network. The network manager then defines the measurement methods (ICMP, TCP, UDP, and so on). The IPM server's database stores the definitions so any other IPM client application can access them.

When the components are defined, the IPM server will configure the SAA feature in the source router to perform the operation at a specified interval. Each SAA takes measurements at specified intervals between the source router and the target device, using the specified traffic operation. The IPM server then extracts data from each source router and stores the data in the IPM database. The IPM client provides a real-time feature that allows an operator to display the data collected immediately. IPM, together with IOS SAA, provides reports on operation latency, device availability, and packet jitter between a source router and a target device for a specified operation.

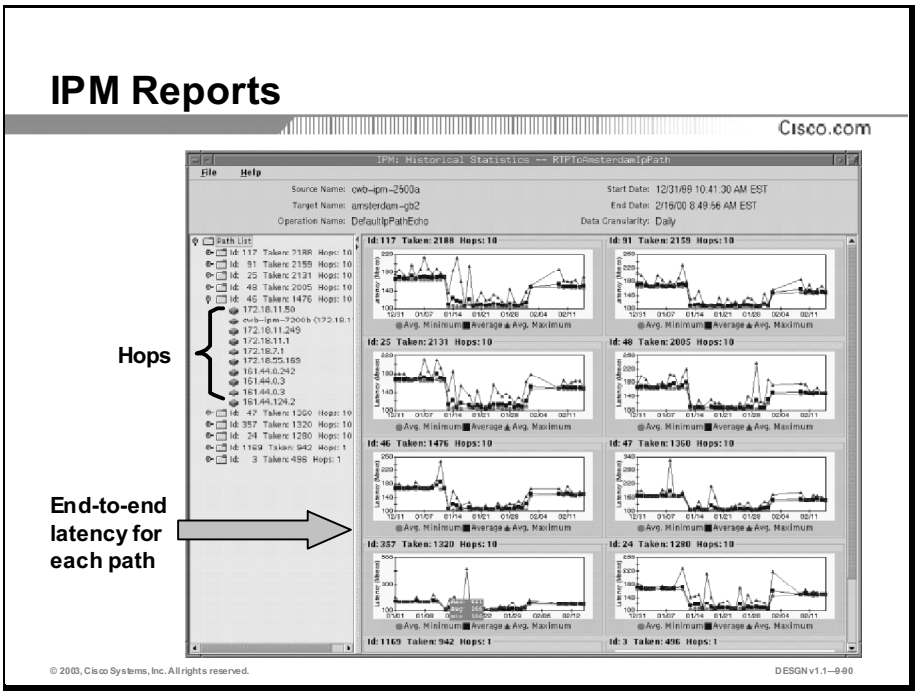

IPM provides comprehensive reports. The figure shows an example report. The report presents the “ICMP Path Echo” historical statistics, containing a list of the paths that the collector found between the source router and the target device. For each path between the source and the target, the IPM user (network operator) can display the hops (router hostnames or IP addresses), which are displayed on the left selection bar. The right side of the figure displays a separate graph for each unique path, which illustrates the end-to-end delay from the source router to the target device using that path.

IPM uses SNMP to configure SAA and read its data. Because SNMP is unreliable and is not secure, do not deploy IPM beyond the boundaries of a controlled network.

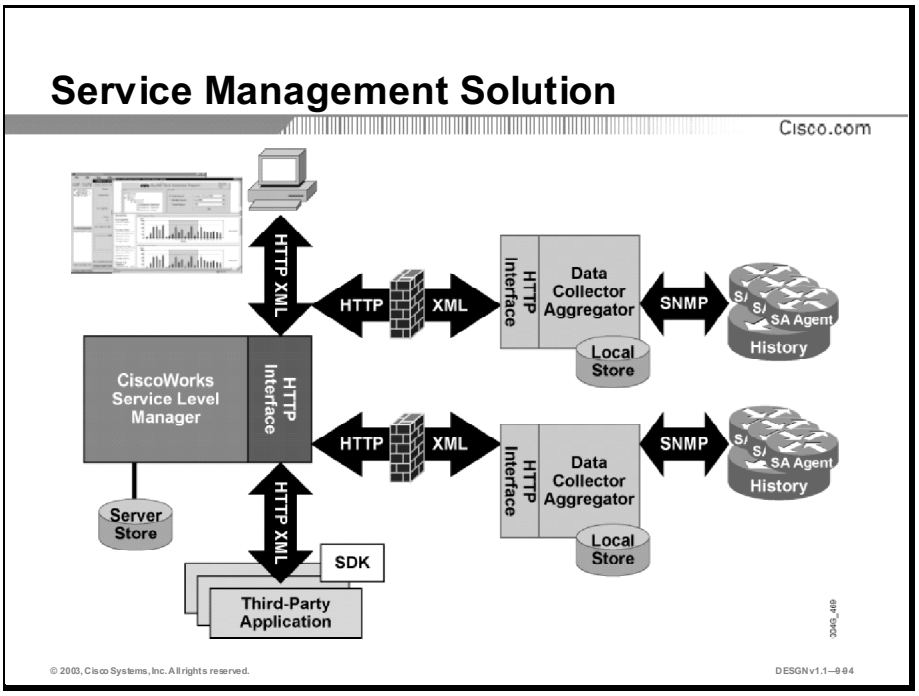

Service Level Manager manages service levels between enterprises and internal or external service providers. Working with the embedded IOS SAA, Service Level Manager defines and monitors SLAs, specifying traffic type, endpoints, and thresholds against key parameters such as latency, packet loss, and jitter. Service Level Manager sets real-time traps, allowing prompt detection and correction of possible service degradation.

Service Level Manager communication interfaces are industry-standard, including SNMP and extensible markup language (XML)/HTTP, allowing for data sharing and integration with third-party vendor applications. The architecture uses distributed data collection agents to poll locally via SNMP for service metrics recorded on SAAs.

The data collector is responsible for configuring the SAAs with the IOS commands necessary to perform the synthetic traffic operations as defined in the SLAs. Communication between the Service Level Manager server and the data collector is via HTTP and XML, as opposed to SNMP, which operates through a firewall.

Summary

This topic summarizes the key points discussed in this lesson.

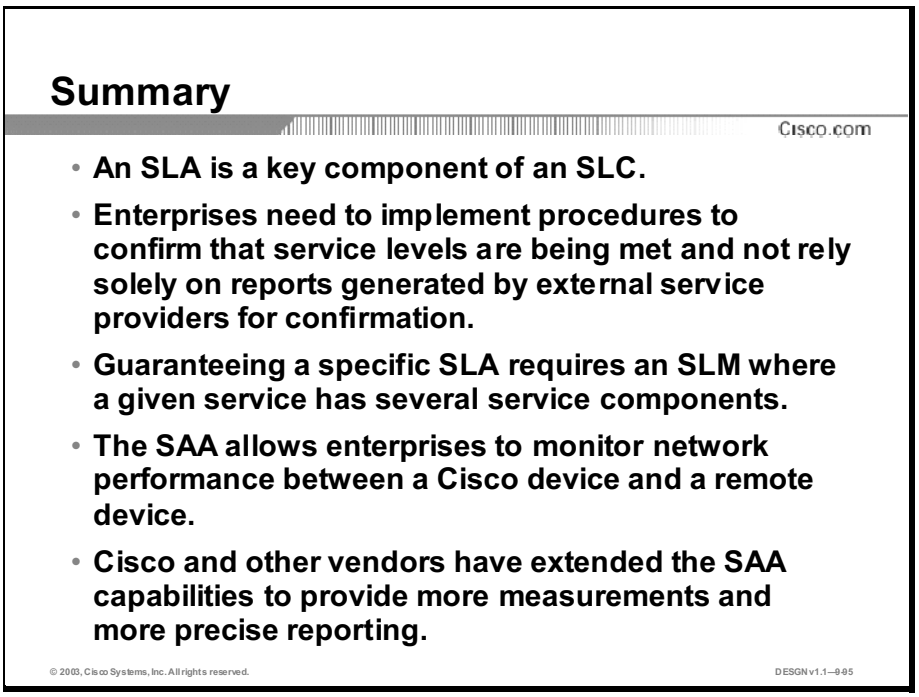

References

For additional information, refer to this resource:

- Service Assurance Agent (SAA), http://www.cisco.com/warp/public/732/Tech/nmp/saa/

Quiz

Use the practice items here to review what you learned in this lesson. The correct answers are found in the Quiz Answer Key.

Q1) Choose the best procedure for a department to enable business-critical services on the corporate network.

A) The department manager communicates the need to the network operator.

B) Department employees determine if the network meets their requirements before using the application.

C) Implement a service level agreement (SLA) between the department manager and network operator.

D) The department manager orders a publicly available SLA from the network operator.

Q2) Company A runs a business-critical, customer-service application over the corporate network. An external company, company B, controls this network. Company A manages the user workstations, servers, and the application. To meet business-critical requirements, the application demands a constant connection with the server. If employees experience greater than five minutes of downtime a day, the responsible department, or organization, is held accountable. Why is an SLA between company A and company B recommended in this situation?

A) An SLA enables company B to control services that company A is running.

B) If a violation of the SLA occurs, the reports from active monitoring identify the fault and penalties can be assessed.

C) If a violation of the SLA occurs, company B is alerted and can assess penalties against company A.

D) SLA reports are generated to show company A that company B performs proactive monitoring.

Q3) What are three requirements for an end-to-end SLM? (Choose three.)

A) The SLM must leverage management products from Cisco.

B) The SLM must incorporate clients and equipment the customer does not own or control

C) The SLM must adapt to new SLM metrics as new technologies are deployed.

D) The SLM must scale enough to meet the current needs.

E) The provider must collect the correct network and application SLM metrics at the appropriate times.

Q4) Corporate remote locations connected via an ISP are regularly sending files using FTP. An agreement with the ISP states that for FTP, 100 kbps of bandwidth will be guaranteed. Which SAA operations could be used for testing the SLA remote site latency?

A) FTP and ICMP

B) UDP and jitter

C) UDP and HTTP

D) IP telephony

Q5) Which Cisco application would provide the most comprehensive set of management reports on the SLAs?

A) Service Level Manager

B) IPSC

C) SAA report manager

D) IPM

Quiz Answer Key

Q1) C

Relates to: Importance of SLAs

Q2) B

Relates to: SLA Requirements

Q3) B, C, E

Relates to: SLM as a Key Component for Assuring SLAs

Q4) A

Relates to: SAA

Q5) A

Relates to: Network Response and Availability Applications

Final Case Study: MCMB Corporation Network Redesign

Complete this case study to practice what you learned in this course.

Required Resources

These are the resources required to complete this exercise:

- MCMB Corporation Network Redesign—Scenario
- Case Study Guidelines, presented in the "Course Introduction" module
- Case Study Solution, instructor to provide a sample Case Study Solution
- A student workgroup (two to three students)
- Blank sheets of paper and a pencil

Exercise Objective

In this exercise, you will identify the design methodology implementation details. Upon completing this case study, you will be able to meet these objectives:

- Propose a redundant campus switching design
- Select an appropriate WAN backup strategy
- Propose a suitable IP addressing plan
- Select a routing protocol for the network
- Propose a design for the Extranet

Job Aids

These job aids are available to help you complete this exercise:

- Case Study Guidelines, presented in the "Course Introduction" module
- Case study solutions, presented in the appendix "Case Study Solutions"
- Forms found within the course

Exercise Procedure

Complete these steps:

Step 1 Read the following scenario (MCMB Corporation Network Redesign—Scenario) completely before the exercise. Allow 10 to 15 minutes for reading.

Step 2 Discuss the case study scenario with your group. Allow 10 minutes for the discussion.

Step 3 Propose a campus redesign that solves current problems reported by the users. Consider redundancy when redesigning the LAN. If possible, keep the existing equipment in place.

Step 4 Propose a WAN backup design.

Step 5 Propose a redesign of the IP addressing. When designing the IP addressing, consider IP address summarization that should simplify the routing plan.

Step 6 Propose a new routing protocol to meet needs if the existing one does not fulfill all the requirements. If a new routing protocol is not required, outline any modifications that will be required for the existing routing protocol.

Step 7 Propose a design for the Extranet, which should enable secure access to the external servers for MCMB's customers and partners regardless of the connectivity option they choose. Consider alternative solutions and propose the one that will be appropriate for all customers with a reasonable investment in network equipment and monthly costs.

Exercise Verification

You have completed this exercise when your case study solution has been presented to the class.

MCMB Corporation Network Redesign—Scenario

This case study will involve analyzing the network infrastructure of MCMB Corporation, a fictitious manufacturer of tricycle tires. The company has provided you with a short description of their current situation and their plans. It is your task, as a network designer, to identify the customer requirements that will allow you to provide the most effective solution.

Company Facts

MCMB Corporation, a leading manufacturer of tricycle tires, is an international company with headquarters in Lyon, France, and offices around the world. The demand for the company's product is constantly increasing; therefore, the company faces the need to tighten the integration of its customers and partners into its information infrastructure.

This company employs approximately 8000 people across 75 sites. These sites are located globally and vary in size from a single part-time person working from a small office or home office (SOHO) to an office with 100 employees. The company's headquarters, located in Lyon, consists of three buildings containing approximately 1500 employees, 500 in each building. The headquarters also houses four major departments that are dispersed across all three buildings: development-manufacturing, technical support, marketing, and sales.

The company has one international office in each of the countries where they are present, and many smaller remote offices throughout each of these countries. All larger remote offices (more than five employees) within each of the countries are permanently connected to the international office, while smaller offices (less than five employees) are connected on-demand using ISDN. Each international office provides connectivity to the headquarters for all sites within that country.

Most of the time, the employees use universal software that covers all the business work flows within the company and runs on a mainframe. The e-mail server runs on a separate platform.

Current Situation

MCMB Corporation's WAN network is a typical hierarchical, three-tier aggregation network using serial lines, low bandwidth (64 kbps to 128 kbps) between larger remote offices and the international office, and increasingly higher bandwidth (up to 2 Mbps) between international offices and the headquarters. All remote locations with more than five employees are Layer 2 (L2)-switched while smaller locations use Ethernet hubs.

The current network does not provide any redundancy and backup strategy. Therefore, in the event of link or network equipment failures, the employees cannot perform their daily work, which results in lower productivity.

Internet connectivity with a maximum throughput of 1 Mbps is provided via a central firewall located next to the mainframe. The data traffic to the Internet is predominantly HTTP and e-mail.

All internal servers (e-commerce, e-mail, and so on) are located at the headquarters in one of the three buildings. Most of the time (approximately 75 percent), employees are accessing these servers, and the rest of the time they are accessing the Internet. The entire campus network at the headquarters is switched using Layer 2 switches with 10/100 Mbps ports for connecting users and GigaEthernet ports for the servers. The access and distribution switches are located in

every building. The distribution switches are connected to the core switch that builds the campus backbone. Currently, there is no redundancy implemented in the LAN.

The existing campus switching is shown in the figure.

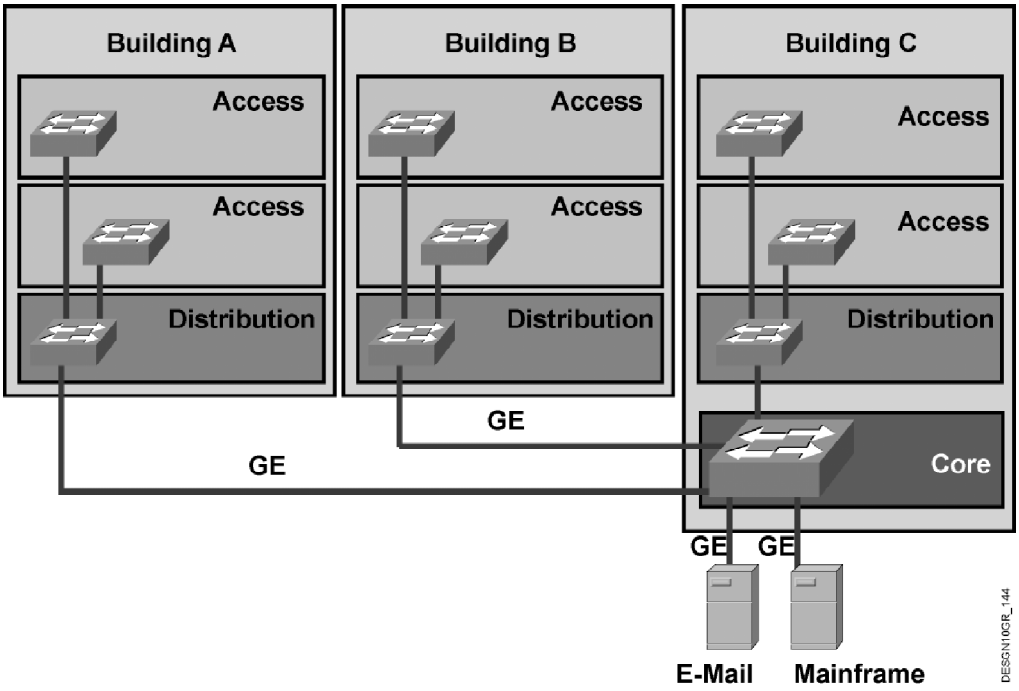

The MCMB Corporation Campus Switching Configuration

Users repeatedly report slow response times, although monitoring the campus backbone switch does not show any serious congestion or high CPU utilization. Furthermore, the average WAN utilization between the headquarters and the international offices is approximately 30 percent and rarely exceeds 80 percent. The WAN utilization between international offices and remote offices is even lower and does not seem to be the issue. The administrators are highly concerned about the reported slow response times, and believe this problem needs to be resolved immediately.

The network is based on IP, using the Enhanced Interior Gateway Routing Protocol (EIGRP). One C-class network from the range 192.168.x.x is allocated to each LAN and WAN link. In locations where more than 250 IP addresses are required, an additional C-class network is added. No route summarization is implemented on any router. Connections to small remote offices with ISDN lines are implemented with static routes that are redistributed into the routing protocol at the international offices.

MCMB's network with three international offices is illustrated in the figure.

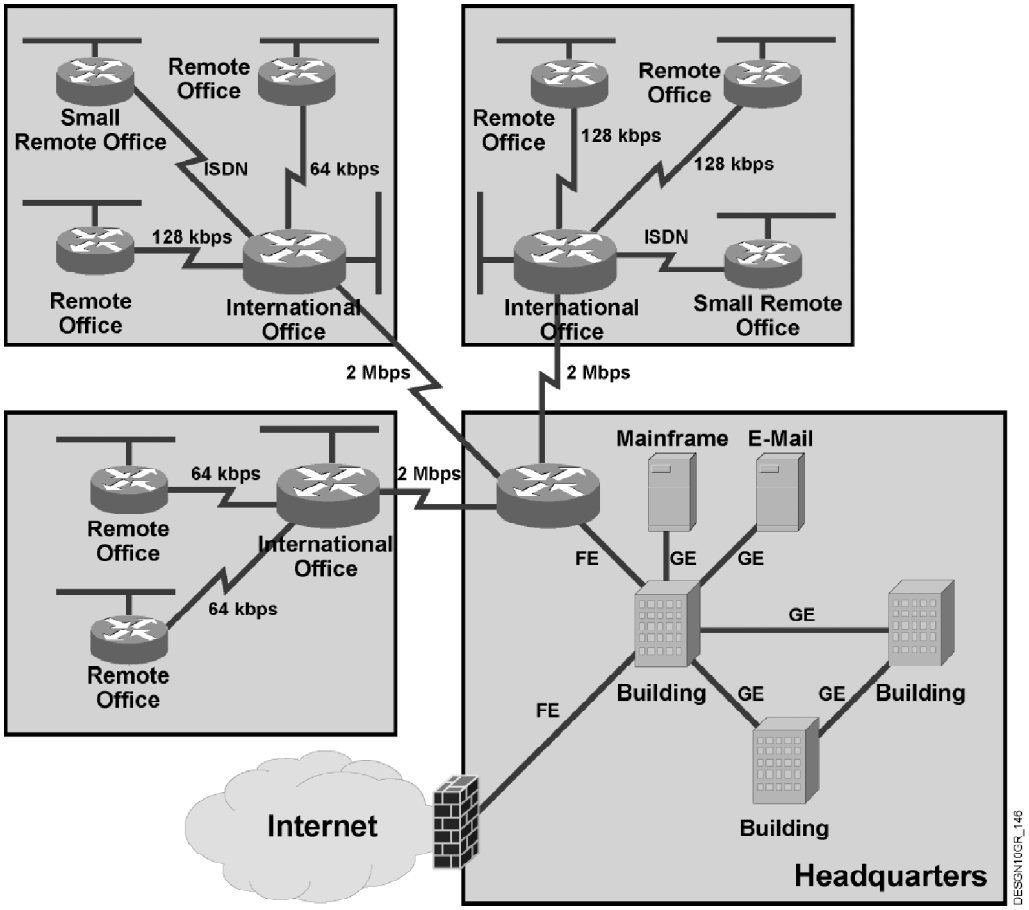

The MCMB Corporation Network

Plans and Requirements

The company urgently needs to resolve the slow response time issue. They are open to any solution that would solve this problem. An added value would be to include redundancy to the system. At the same time, they want to keep the existing equipment and thus minimize the investment costs. The redundancy should be implemented in the headquarters LAN, and on all WAN links. They are aware of potential business risks that may be caused by periodic network outages during implementation of the solution.

MCMB would also like to improve their IP addressing to provide a scalable and manageable solution that would allow simple expansion to new markets by simply adding new remote locations to the system. They would also like to separate the four departments in the campus and possibly deploy security policies among them in the future.

Additionally, the company wants to offer Extranet functionality to their partners and resellers, which would simplify their business. They are considering different connectivity options to the Internet, such as leased-line, ISDN, and asymmetric digital subscriber line (ADSL). MCMB wants to provide secure access to their external servers, which should be separated from the internal servers. The internal mainframe should upload and download the data from the external server in a secure manner. This solution must be universal, regardless of the connectivity option the partners will choose. They expect that about 50 partners and customers will use the Extranet at the beginning. The application requires approximately 64 kbps of bandwidth on the WAN for acceptable response times.

Exercise: Propose Your Network Redesign

The worksheets on the following pages are intended to help you develop your network redesign for this customer. Use additional sheets as necessary.

Campus Redesign

Propose a campus redesign that solves current problems reported by the users. Consider redundancy when redesigning the LAN.

Write a one-paragraph overview describing your redesign and articulate why you selected that solution.

_____
_____
_____
_____
_____

Use the following figure to create your diagram showing the proposed campus switching topology.

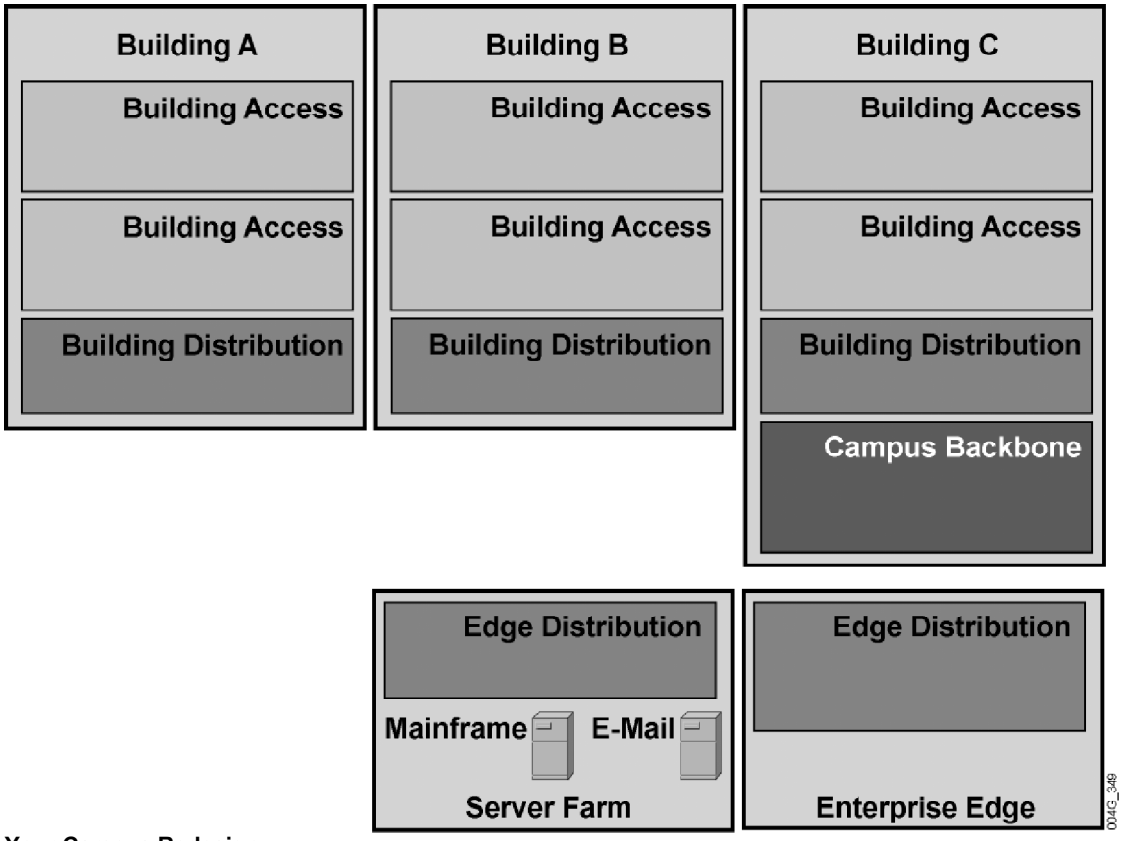

Your Campus Redesign

WAN Backup Design

Propose a WAN backup design to improve network reliability.

Write a one-paragraph overview describing your design and articulate why you selected that solution.

_____
_____
_____
_____
_____

Use the following figure to create your diagram showing the proposed WAN backup topology.

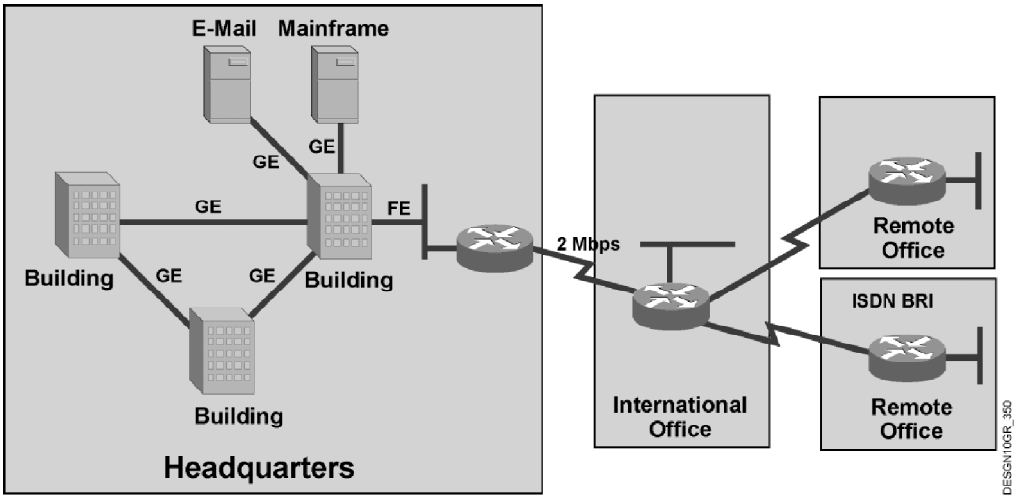

Your WAN Backup Design

IP Addressing Redesign

Propose a redesign for the IP addressing. When designing the IP addressing, consider IP address summarization that would simplify the routing plan.

Write a one-paragraph overview describing your redesign and articulate why you selected that solution.

_____

_____

_____

Use the following table to create your IP addressing redesign. Use additional rows for missing items.

IP Addressing Redesign

Location	IP Address Space
Headquarters	
Headquarters—development-manufacturing	
Headquarters—technical support	
Headquarters—marketing	
Headquarters—sales	
Headquarters—server farm	
Headquarters—edge distribution	
International network	
International Office LAN network	
Remote Office 1 LAN network	
Remote Office 2 LAN network	
International Office to Headquarters WAN connection	
International Office to Remote Office 1 WAN connection	
International Office to Remote Office 2 WAN connection	
International Office to Headquarters backup	
International Office to Remote Office 1 backup	
International Office to Remote Office 2 backup	

Use the following figure to create your diagram showing the proposed IP addressing in the campus and WAN, including backup. This diagram should correlate to the previous IP addressing table. Indicate IP address summarizations as well.

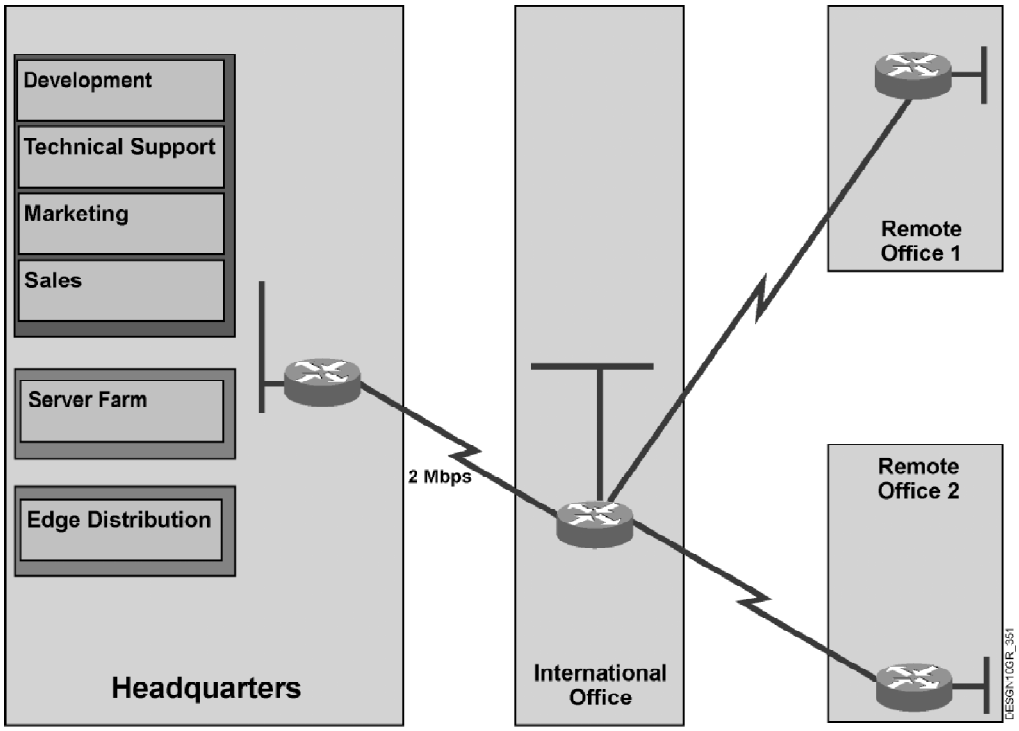

Your IP Addressing and Summarization Redesign

Routing Campus Redesign

Propose a new routing protocol to meet needs if the existing one does not fulfill all the requirements. If a new routing protocol is not required, outline any modifications that will be required for the existing routing protocol.

Write a one-paragraph overview describing your redesign and articulate why you selected that solution.

_____
_____
_____
_____
_____
_____

Use the following figure to create your diagram showing the proposed routing protocol modifications.

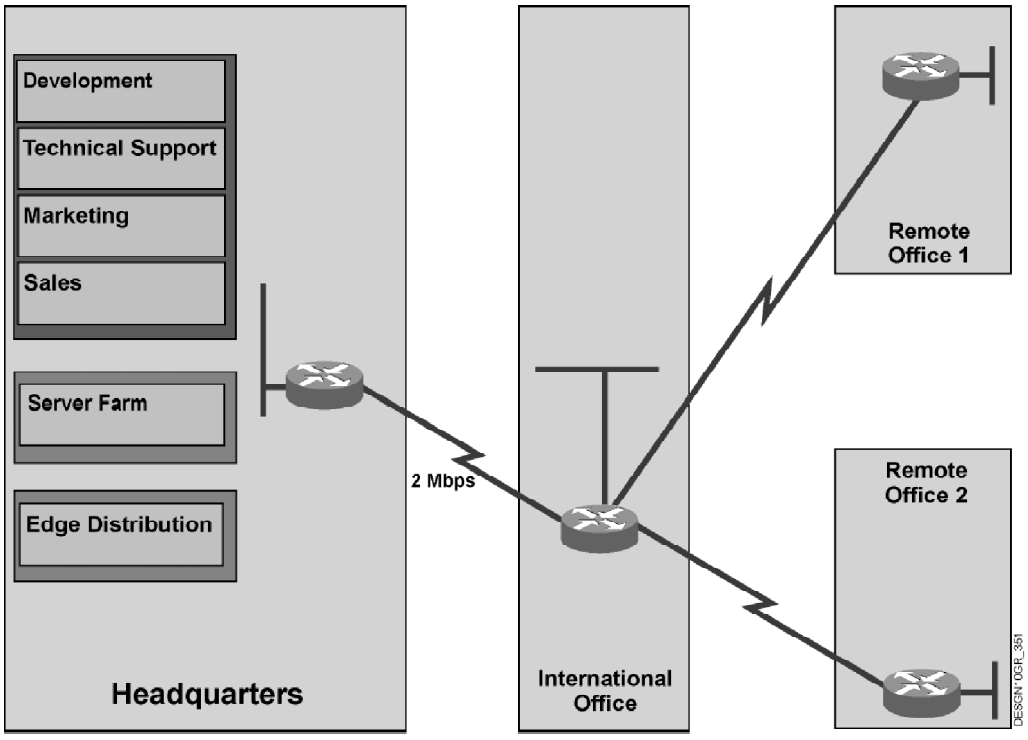

Your IP Routing Redesign

Extranet Design

Propose a design for the Extranet, which should enable secure access to the external servers for MCMB's customers and partners regardless of the connectivity option they choose. Consider alternative solutions and propose one that will be appropriate for all customers with a reasonable investment in network equipment and monthly costs.

Write a one-paragraph overview describing your design and articulate why you selected that solution.

_____
_____
_____
_____
_____

In the space provided below, create a diagram showing the proposed solution for Extranet.

DESGN

Course Glossary

The Course Glossary for *Designing for Cisco Internetwork Solutions* (DESGN) v1.1 highlights and defines key terms and acronyms used throughout this course. Many of these terms are also described in the Cisco Internetworking Terms and Acronyms resource, available via http://www.cisco.com/univercd/cc/td/doc/cisintwk/ita/.

Acronym	Definition
1000BASE-T	As the leading provider of Gigabit Ethernet and switched internetworking solutions, Cisco Systems is committed to the development of high-performance Ethernet technology and products that provide gigabit-per-second transmission rates to address both service provider and enterprise customer requirements.
100BASE-T	100-Mbps baseband Fast Ethernet specification using UTP wiring. Like the 10Base-T technology on which it is based, 100Base-T sends link pulses over the network segment when no traffic is present. However, these link pulses contain more information than those used in 10Base-T. Based on the IEEE 802.3 standard.
10BASE-T	10-Mbps baseband Ethernet specification using two pairs of twisted-pair cabling (Categories 3, 4, or 5): one pair for transmitting data and the other for receiving data. 10Base-T, which is part of the IEEE 802.3 specification, has a distance limit of approximately 328 feet (100 meters) per segment.
AAA	authentication, authorization, and accounting (pronounced “triple A”).
AAL1, AAL2, AAL3/4,AAL5	ATM adaptation Layer 1, 2, 3/4, and 5. The four AALs currently in the ITU-T BISDN recommendations and the ATM Forum specifications. AAL1 is used for connection-oriented, delay-sensitive services requiring constant bit rates, such as uncompressed video and other isochronous traffic. AAL2 is used for connection-oriented services that support a variable bit rate, such as some isochronous video and voice traffic. AAL3/4 is used for connection-oriented data services. It is closely aligned with SMDS. AAL5 is used to support connection-oriented VBR services primarily to transfer classical IP over ATM and LANE traffic. This least complex of the AAL recommendations uses simple and efficient AAL (SEAL), offering lower bandwidth costs and simpler processing requirements but also providing reduced bandwidth and error-recovery capacities.
ABR	1. available bit rate. Service category defined by the ATM Forum for ATM networks. ABR is used for connections that do not require timing relationships between source and destination. ABR provides no guarantees in terms of cell loss or delay, providing only best-effort service. Traffic sources adjust their transmission rate in response to information that they receive describing the status of the network and its capability to successfully deliver data. 2. Area Border Router. Router located on the border of one or more OSPF areas that connects those areas to the backbone network. ABRs are considered members of both the OSPF backbone and the attached areas. They therefore maintain routing tables describing both the backbone topology and the topology of the other areas.
access layer	One of the layers in the three-layer hierarchical model. The access layer provides users with access to the internetwork.
access server	Communications processor that connects asynchronous devices to a LAN or WAN through network and terminal emulation software. Performs both synchronous and asynchronous routing of supported protocols. Sometimes called a NAS.
ACD	1. automatic call distributor. Programmable device at a telephone call center that routes incoming telephone calls to agents (persons) within that call center. After the system determines the agent for a call, the call is sent to the ACD associated with that agent. 2. automatic call distribution. Device or service that automatically reroutes calls to customers in geographically distributed locations served by the same CO.
ACELP	algebraic code excited linear prediction. The process by which analog voice samples are encoded into high-quality digital signals.
ACL	access control list.
address	Data structure or logical convention used to identify a unique entity, such as a particular process or a network device.

Acronym	Definition
address mask	A bit combination used to describe which part of an address refers to the network or the subnet and which part refers to the host. Sometimes referred to simply as a mask.
administrative distance	Rating of the trustworthiness of a routing information source. Administrative distance often is expressed as a numerical value between 0 and 255. The higher the value, the lower the trustworthiness rating.
ADPCM	adaptive differential pulse code modulation. The process by which analog audio samples are encoded into compressed digital signals.
ADSL	asymmetric digital subscriber line. One of four DSL technologies. ADSL is designed to deliver more bandwidth downstream (from the central office to the customer site) than upstream. Downstream rates range from 1.5 to 9 Mbps, whereas upstream bandwidth ranges from 16 to 640 kbps. ADSL transmissions work at distances up to 18,000 feet (5488 meters) over a single copper twisted pair.
a-law	ITU-T companding standard used in the conversion between analog and digital signals in PCM systems. A-law is used primarily in European telephone networks and is similar to the North American u-law standard.
AM	amplitude modulation. A modulation technique whereby information is conveyed through the amplitude of the carrier signal.
amplitude	The maximum value of an analog waveform or a digital waveform. Also, the magnitude or strength of a varying waveform. Typically represented as a curve along the x-axis of a graph.
analog signal	The representation of information with a continuously variable physical quantity, such as voltage. Because of this constant changing of the wave shape with regard to its passing a given point in time or space, an analog signal might have a virtually indefinite number of states or values. This type of signal contrasts with a digital signal, which is expressed as a square wave and therefore has a very limited number of discrete states.
anti-replay	Security service where the receiver can reject old or duplicate packets in order to protect itself against replay attacks. IPSec provides this optional service by use of a sequence number combined with the use of data authentication.
API	application programming interface. The means by which an application program talks to system software and utilities. Standardized communications APIs allow application programs to be developed independently of the underlying method of communication. More specifically, APIs are a set of standard software calls, and data formats that computer application programs use to initiate contact with other devices (for example, network services, mainframe communications programs, or other program-to-program communications). Typically, APIs make it easier for software developers to create the links that an application needs to communicate with the operating system or with the network.
ARP	Address Resolution Protocol. Internet standard protocol used to map an IP address to an Ethernet MAC address. Defined in RFC 826.
ASBR	Autonomous System Boundary Router. ABR located between an OSPF autonomous system and a non-OSPF network. ASBRs run both OSPF and another routing protocol, such as RIP. ASBRs must reside in a nonstub OSPF area.
ASIC	application-specific integrated circuit.
ATM	Asynchronous Transfer Mode. The international standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media from T1/E1 through SONET/SDH rates to OC-192.
autonomous system	A collection of networks under a common administration sharing a common routing strategy. Autonomous systems are subdivided by areas. An Internet visible autonomous system must be assigned a unique 16-bit number by the IANA.

Acronym	Definition
AVVID	Architecture for Voice, Video and Integrated Data.
BackboneFast	A feature on a switch that reduces the link Spanning Tree Protocol convergence time from 50 seconds to from 20 to 30 seconds.
Backplane	The physical connection between an interface processor or card and the data buses and the power distribution buses inside a chassis.
Backup	A way of providing high availability by using redundant links. A backup connection can be established either via dial-up or by using permanent connections.
Baseband	Characteristic of a network technology where only one carrier frequency is used. Ethernet is an example of a baseband network. Also called narrowband.
Bc	committed burst. Negotiated tariff metric in Frame Relay internetworks. The maximum amount of data (in bits) that a Frame Relay internetwork is committed to accept and transmit at the CIR.
Be	excess burst. Negotiated tariff metric in Frame Relay internetworks. The number of bits that a Frame Relay internetwork attempts to transmit after Bc is accommodated. Be data, in general, is delivered with a lower probability than Bc data because Be data can be marked as DE by the network.
BECN	backward explicit congestion notification. Bit set by a Frame Relay network in frames traveling in the opposite direction of frames encountering a congested path. DTE receiving frames with the BECN bit set can request that higher-level protocols take flow control action as appropriate.
BER	bit error rate. Ratio of received bits that contain errors.
BGP	Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. Originally defined by RFC 1105, the current protocol definition is contained in RFC 1771.
BPDU	bridge protocol data unit. Spanning Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network.
Bps	bits per second.
Bps	Bytes per second
BRI	Basic Rate Interface. ISDN interface composed of two B channels and one D channel for circuit-switched communication of voice, video, and data.
Broadband	Term describing facilities or services that operate at the DS3 rate and above.
broadcast	Data packets that are sent to all nodes on a network.
broadcast address	A special address reserved for sending a message to all stations. Generally, a data link layer broadcast address is a MAC destination address of all 1s. An IPv4 broadcast address has all 1s in the host portion of the IP address.
broadcast domain	Set of all devices that receive broadcast frames originating from any device within the set. Data link layer broadcast domains typically are bounded by routers because routers do not forward broadcast frames.
broadcast storm	An undesirable network event in which many broadcasts are sent simultaneously across all network segments. A broadcast storm uses substantial network bandwidth and, typically, causes network timeouts.
BSCI	*Building Scalable Cisco Internetworks.*
BSCN	*Building Scalable Cisco Networks.*
Building Access submodule	A submodule within the Enterprise Composite Network Model. Contains end-user workstations, IP Phones, and Layer 2 access switches for connecting devices to the Building Distribution component.

Acronym	Definition
Building Block module	A module within the Enterprise Composite Network Model that comprises the Building Access and Building Distribution submodules.
Building Distribution submodule	A submodule within the Enterprise Composite Network Model. Provides aggregation of access networks using Layer 3 switching. Distribution performs routing, QoS, and access control.
CAC	Call Admission Control. CAC mechanisms extend the capabilities of the QoS tool suite to protect voice traffic from being negatively affected by other voice traffic and to keep excess voice traffic off the network.
CAG	Course Administration Guide.
Call leg	Discrete segment of a voice call connection. A call leg is a logical connection between the router and either a telephony endpoint over a bearer channel, or another endpoint using an H.323 protocol.
Campus	Enterprise Composite Network subdivision known as the “Campus.”
Campus Backbone submodule	A submodule within the Campus Infrastructure module of the Enterprise Composite Network Model that connects distribution modules.
campus core	One of the layers in the three-layer hierarchical model. The core layer connects multiple distribution layer devices and provides fast access.
Campus Infrastructure module	A module within the Enterprise Composite Network Model that comprises the Building Access, Building Distribution, and Campus Backbone submodules.
CAR	committed access rate. The CAR limits the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.
CAS	channel associated signaling. The transmission of signaling information within the voice channel. CAS signaling often is implemented as robbed-bit signaling because user bandwidth is being “robbed” by the network for other purposes.
Category 5	One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 5 cabling can transmit data at speeds of up to 100 Mbps.
CatOS	Catalyst software.
CATV	cable television. A communication system where multiple channels of programming material are transmitted to homes using broadband coaxial cable. Formerly called Community Antenna Television.
CBR	constant bit rate. Service category defined by the ATM Forum for ATM networks. CBR is used for connections that depend on precise clocking to ensure undistorted delivery.
CBWFQ	class-based weighted fair queuing extends the standard WFQ functionality to provide support for user-defined traffic classes.
CCDA®	Cisco Certified Design Associate.
CCDP®	Cisco Certified Design Professional.
CCIE®	Cisco Certified Internetwork Expert.
CCIP™	Cisco Certified Internetwork Professional.
CCNA®	Cisco Certified Network Associate.
CCNP®	Cisco Certified Network Professional.
CCS	common channel signaling. Signaling system used in telephone networks that separates signaling information from user data. A specified channel is exclusively designated to carry signaling information for all other channels in the system.
CCSI	Cisco Certified Systems Instructor.

Acronym	Definition
CDB	call data block. This block consists of several CDEs, related to a certain point in call (PIC).
CDE	call data element. A field that includes basic information within a billing record. Examples of CDEs are the calling number, called number, and so on.
CDN	Content Delivery Network.
CDP	Cisco Discovery Protocol.
CDR	Call Detail Record. 1. A record written to a database for use in postprocessing activities. CDRs comprise CDEs. CDR files consist of several CDBs. These activities include many functions but primarily are billing and network analysis. CallManager writes CDR records to a relational database as calls are made in a manner consistent with the configuration of each individual CallManager. Used in the original telephony networks, and now extended to mobile wireless network calls, the CDR contains billing information for charging purposes. In a GPRS network, the charging gateway sends the billing information within a CDR to the network service provider for that subscriber. 2. VNS record of voice or data SVCs, which includes calling and called numbers, local and remote node names, data and time stamp, elapsed time, and Call Failure Class fields.
CD-ROM	compact disk read-only memory.
CEF	Cisco Express Forwarding.
CELP	code excited linear prediction. Compression algorithm used in low bit-rate voice encoding. Used in ITU-T Recommendations G.728, G.729, and G.723.1.
Centrex	LEC service that provides local switching applications similar to those provided by an on-site PBX. With Centrex, there is no on-site switching; all customer connections go back to the CO.
CEO	chief executive officer.
CES	circuit emulation service. Enables users to multiplex or to concentrate multiple circuit emulation streams for voice and video with packet data on a single high-speed ATM link without a separate ATM access multiplexer.
CGMP	Cisco Group Management Protocol.
CHAP	Challenge Handshake Authentication Protocol. Security feature supported on lines using PPP encapsulation that prevents unauthorized access. CHAP does not itself prevent unauthorized access but merely identifies the remote end. The router or access server then determines whether that user is allowed access. It is defined in IETF RFC 1994.
CIP	Channel Interface Processor.
CIR	committed information rate. The rate at which a Frame Relay network agrees to transfer information under normal conditions, averaged over a minimum increment of time. CIR, measured in bits per second, is one of the key negotiated tariff metrics.
Cisco CallManager	Cisco CallManager is the software-based call-processing component of the Cisco enterprise IP telephony solution, and is a product enabled by Cisco AVVID. CallManager software extends enterprise telephony features and capabilities to packet telephony network devices such as IP Phones, media-processing devices, VoIP gateways, and multimedia applications.
Cisco.com	http://www.cisco.com.

Acronym	Definition
Cisco IOS	Cisco Internetwork Operating System. Cisco software that provides common functionality, scalability, and security for all Cisco products. IOS software allows centralized, integrated, and automated installation and management of internetworks while ensuring support for a wide variety of protocols, media, services, and platforms.
CiscoWorks	CiscoWorks is a network management solution that provides a powerful set of monitoring and configuration tools to simplify the administration of small to medium business networks and workgroups containing Cisco internetworking products (switches, routers, hubs, and access servers).
classful routing protocols	Routing protocols that perform automatic summarization of network information on major IPv4 address class network boundaries only (Class A, B, or C).
classless routing protocols	Routing protocols that propagate subnet mask information for each routing update and do not perform automatic route summarization on major class network boundaries (Class A, B, or C).
CLI	command-line interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments. Cisco IOS, UNIX operating system and DOS provide CLIs.
Client/server model	Common way to describe network services and the model user processes (programs) of those services. Examples include the name server/name resolver paradigm of the DNS and file server/file-client relationships, such as NFS and diskless hosts.
CLNP	Connectionless Network Protocol. The OSI network layer protocol that does not require a circuit to be established before data is transmitted.
CLP	1. Cisco Learning Partner. 2. cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are ensured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions to free up resources to handle ensured traffic.
CM	cable modem. Device used to connect a PC to a local cable TV line and receive data at much higher rates than ordinary telephone modems or ISDN. A cable modem can be added to or integrated with a set-top box, thereby enabling Internet access via a television set. In most cases, cable modems are furnished as part of the cable access service, and are therefore not purchased directly and installed by the subscriber.
CMTS	cable modem termination system, such as a router or a bridge, typically located at the cable headend. Any DOCSIS-compliant headend cable router, such as the Cisco uBR7246.
CN	Content Networking. An essential ingredient for optimization of content delivery, proactively distributing cacheable content from origin servers to content servers at the edges of the network, and keeping content consistent.
CO	central office. The local telephone company office to which all local loops in a given area connect and in which circuit switching of subscriber lines occurs.
codec	coder-decoder. Integrated circuit device (DSP) that transforms analog audio signals into a digital bit stream (code) and digital signals back into analog signals (decode). It may support multiple coding algorithms of various complexity and compression.
collision domain	In Ethernet, the network area within which frames collisions are propagated. Repeaters and hubs propagate collisions; LAN switches, bridges, and routers do not.
companding	Contraction derived from the opposite processes of compression and expansion. Part of the PCM process whereby analog signal amplitude values are rounded logically to discrete scale-step values on a nonlinear scale. The process is reversed at the receiving terminal using the same nonlinear scale. It provides resolution equal to a 12 to 13-bit sample using an 8-bit sample space.

Acronym	Definition
connectionless	Term used to describe data transfer without the existence of a virtual circuit.
connection-oriented	Term used to describe data transfer that requires the establishment of a virtual circuit.
Content Cache	A cache that accelerates content delivery for end users by transparently caching frequently accessed content and then locally fulfilling content requests rather than traversing the Internet/intranet to a distant server.
Content Distribution Manager	Cisco device that performs all the management functions needed to control content distribution.
convergence	1. Speed and ability of a group of internetworking devices running a specific routing protocol to agree on the topology of an internetwork after a change in that topology. 2. Speed and ability of a group of interconnected switches to rebuild a spanning tree following a link failure.
core layer	One of the layers in the three-layer hierarchical model. The core layer connects multiple distribution layer devices and provides fast access.
Core submodule	Part of the Campus Infrastructure module in the Enterprise Composite Network Model. It is also called the Campus Backbone submodule.
CoS	class of service. An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages. In SNA subarea routing, CoS definitions are used by subarea nodes to determine the optimal route to establish a given session. A CoS definition comprises a virtual route number and a transmission priority field. IP type of service, DiffServ and IEEE 802.1p are CoS implementations.
CPE	customer premises equipment. Terminating equipment, such as terminals, telephones, and modems, supplied by the telephone company, installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment residing on the customer site.
CPU	central processing unit.
crossbar	Multigigabit switching fabric implementation in high-end Cisco switches.
cRTP	Compressed Real-Time Transport Protocol. Protocol that compresses voice headers from 40 bytes to 2 or 4 bytes, which offers significant bandwidth savings. cRTP is sometimes referred to as RTP header compression.
CS-ACELP	conjugate structure algebraic code excited linear prediction. CELP voice compression algorithm providing 8 kbps, or 8:1 compression, standardized in ITU-T Recommendation G.729.
CSU	channel service unit. Digital interface device that connects end-user equipment to the local digital telephone loop. It terminates the service provider circuit. Often referred to together with DSU, *as* CSU/DSU.
CTI	computer telephony integration. The name given to the merger of traditional telecommunications (PBX) equipment with computers and computer applications. The use of caller ID to retrieve customer information automatically from a database is an example of a CTI application.
dark fiber	An optical fiber infrastructure composed of cabling and regenerators directly connected to edge devices and privately managed. It is the optical equivalent of a leased line.
data link layer	Layer 2 of the OSI reference model. Provides reliable transit of data across a physical link. The data link layer is concerned with physical addressing, network topology, line discipline, error notification, ordered delivery of frames, and flow control. The IEEE has divided this layer into two sublayers: the MAC sublayer and the LLC sublayer. The data link layer is sometimes simply called the link layer.

Acronym	Definition
DCE	data circuit-terminating equipment (ITU-T expansion). Devices and connections of a communications network that comprise the network end of the user-to-network interface. The DCE provides a physical connection to the network, forwards traffic, and provides a clocking signal used to synchronize data transmission between DCE and DTE devices. DSUs, modems, and interface cards are examples of DCE.
DCN	*Designing Cisco Networks.*
DDR	dial-on-demand routing. Technique whereby a router can automatically initiate and close a circuit-switched session as transmitting stations demand. The router spoofs keepalives so that end stations treat the session as active. DDR permits routing over ISDN or telephone lines using an external ISDN terminal adapter or modem.
DE	discard eligible. Frame relay header bit, that when set, allows traffic to be dropped preferentially. It is used when the network is congested, to ensure the delivery of unmarked traffic. The Frame Relay network sets it when a traffic stream violates its traffic contract. It may also be set by Frame Relay clients to identify less critical traffic.
dejitter buffer	Buffer used at the receiving end to smooth delay variability. Dejitter buffers delay the first talk spurt to provide smooth playback of voice traffic.
denial of service	An incident in which a user or organization is deprived of the services of a resource that it would normally expect to have. Typically, the loss of service is the inability of a particular network service, such as e-mail, to be available or the temporary loss of all network connectivity and services.
designated bridge	Bridge that incurs the lowest path cost when forwarding a frame from a segment to the root bridge.
designated router	OSPF router that generates LSAs for a multiaccess network and has other special responsibilities in running OSPF. Each multiaccess OSPF network that has at least two attached routers has a designated router that is elected by the OSPF hello protocol. The designated router enables a reduction in the number of adjacencies required on a multiaccess network, which in turn reduces the amount of routing protocol traffic and the size of the topological database.
DHCP	Dynamic Host Configuration Protocol. Protocol that provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them. It is defined in IETF RFC 2131.
dial backup	Feature that provides protection against WAN downtime by allowing the network administrator to configure a backup serial line through a circuit-switched connection.
dial peer	Addressable call endpoint. In Voice over IP, there are two kinds of dial peers: POTS and VoIP.
dial-up	Communications circuit that is established by a switched-circuit connection using the telephone company network.
digital signature	Value computed with a cryptographic algorithm and appended to a data object in such a way that any recipient of the data can use the signature to verify the data's origin and integrity.
distance vector routing protocols	Distance vector routing algorithms call for each router to send its entire routing table in each update, but only to its neighbors. Distance vector routing algorithms can be prone to routing loops, but are computationally simpler than link-state routing algorithms. Also called Bellman-Ford routing algorithm.
distribution layer	One of the layers in the three-layer hierarchical model.
DLCI	data-link connection identifier. Value that specifies a VC from a Frame Relay network to an interface on a client device. DLCIs are generally locally significant to an interface. The same DLCI value on a different interface may be a completely different connection. The two end devices of a Frame Relay VC need not use the same DLCI to connect to the network. The optional global addressing extension to the FRF Local Management Interface (LMI) specification makes DLCIs globally significant.

Acronym	Definition
DLSw	data-link switching. Interoperability standard, described in RFC 1434, that provides a method for forwarding SNA and Network Basic Input/Output System (NetBIOS) traffic over TCP/IP networks using data link layer switching and encapsulation. DLSw uses the service switching point (SSP) instead of source-route bridging (SRB), eliminating the major limitations of SRB, including hop-count limits, broadcast and unnecessary traffic, timeouts, lack of flow control, and lack of prioritization schemes.
DNS	Domain Name System. System used on the Internet for translating names of network nodes into addresses. It is defined in IETF RFC 1035, and has been incrementally updated by several later RFCs.
DOCSIS	Data-over-Cable Service Interface Specifications. Technical specifications for equipment at both subscriber locations and the headends of cable operators. Adoption of DOCSIS will accelerate the deployment of data-over-cable services and will ensure interoperability of equipment throughout the infrastructures of system operators.
DSL	digital subscriber line. Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. There are four types of DSL: ADSL, HDSL, SDSL, and VDSL. All are provisioned via modem pairs, with one modem located at a central office and the other at the customer site. Because most DSL technologies do not use the whole bandwidth of the twisted pair, there is usually room remaining for a voice channel.
DSLAM	digital subscriber line access multiplexer. A device that connects many digital subscriber lines to a network by multiplexing the DSL traffic onto one or more network trunk lines.
DSP	digital signal processor. A microprocessor that provides audio signal processing services, and segments the digital voice signal into frames.
DSU	data service unit. Device used in digital transmission that, with a CSU, adapts the clocked physical interface on a DTE device to a framed transmission facility, such as T1 or E1. The DSU also is responsible for such functions as signal timing. Often referred to together with CSU, as CSU/DSU.
DTE	data terminal equipment. Device at the user end of a user-network interface that serves as a data source, destination, or both. DTE connects to a data network through a DCE device (for example, a modem) and typically uses clocking signals generated by the DCE. DTE includes such devices as computers, protocol translators, and multiplexers.
DTMF	dual tone multifrequency. Tones generated when a button is pressed on a telephone.
DWDM	dense wavelength division multiplexing. Optical transmission of multiple signals over closely spaced wavelengths in the 1550-nanometer (nm) region. (Wavelength spacings are usually 100 GHz or 200 GHz, which corresponds to 0.8 nm or 1.6 nm.)
dynamic address resolution	Use of an address resolution protocol to determine and store address information on demand.
dynamic routing	Routing that adjusts automatically to network topology or traffic changes. Also called adaptive routing.
E&M	recEive and transMit (or ear and mouth or earth and magnet).Trunk signaling arrangement generally used for two-way switch-to-switch or switch-to-network connections. The Cisco analog E&M interface is an RJ-48 connector that allows connections to PBX trunk lines (tie-lines). E&M also is available on E1 and T1 digital interfaces.
E1	Wide-area digital transmission scheme used internationally except in North America and Japan that is clocked at a rate of 2.048 Mbps. Available data rates are 1.920 Mbps when the line carries voice signaling, 1.984 Mbps when just framed and, rarely, 2.048 over an unframed line. E1 lines can be leased for private use from common carriers.

Acronym	Definition
EAP	Extensible Authentication Protocol. Framework that supports multiple, optional authentication mechanisms for PPP, including cleartext passwords, challenge-response, and arbitrary dialog sequences.
echo	Telephony-audible and unwanted leak-through of one's own voice into the receive (return) path.
echo cancellation	Method for removing unwanted signals from the main received voice telephony signal.
e-commerce	electronic commerce.
E-Commerce module	A module within the Enterprise Composite Network Model. The E-Commerce module enables enterprises to successfully deploy e-commerce applications.
Edge Distribution module	A module within the Enterprise Composite Network Model that aggregates the connectivity from the various elements at the Enterprise Edge and routes the traffic into the Campus Backbone submodule.
EDI	electronic data interchange. Electronic communication of operational data, such as orders and invoices, between organizations.
EGP	exterior gateway protocol. Internet protocol for exchanging routing information between autonomous systems. Documented in RFC 904. Not to be confused with the general term “exterior gateway protocol.” EGP is an obsolete protocol that was replaced by BGP.
EIA	Electronic Industries Alliance. Group that specifies electrical transmission standards. The EIA and the TIA have developed numerous well-known communications standards, including EIA/TIA-232 and EIA/TIA-449.
EIGRP	Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency, and combines the advantages of link-state protocols with those of distance vector protocols.
e-mail	electronic mail. Widely used network application in which text messages are transmitted electronically between end users over various types of networks using various network protocols.
encryption	Application of a specific algorithm to data so as to alter the content of the data, making it incomprehensible to those who are not authorized to see the information.
Enterprise Campus	A functional area within the Enterprise Composite Network Model. Comprises the modules required to build a highly robust campus network in terms of reliability, availability, scalability, and flexibility. This area contains all the network elements for independent operation within one campus location.
Enterprise Composite Network Model	A modular model that organizes the Enterprise network into units with clear physical, logical, and functional boundaries.
Enterprise Edge	A functional area within the Enterprise Composite Network Model. Aggregates the connectivity to the various external elements at the edge of each enterprise campus.
enterprise network	Large and diverse network connecting most major points in a company or other organization.
Erlang	A standard measurement of voice traffic. One Erlang equals one full hour, or 3600 seconds, of telephone conversation. Erlang tables combine offered traffic, number of circuits, and grade of service, blocking probability.
Ethernet	Baseband LAN developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use carrier sense multiple access collision detect (CSMA/CD) and run over a variety of cable types at 10 Mbps. Ethernet is the basis of the IEEE 802.3 series of standards.

Acronym	Definition
failover	Backup operational mode in which the functions of a system component (such as a processor, server, network, or database, for example) are assumed by secondary system components when the primary component becomes unavailable through either failure or scheduled downtime.
Fast EtherChannel	Bundling multiple Fast Ethernet links that appear as one logical interface.
Fast Ethernet	Any of a number of 100-Mbps Ethernet specifications. Fast Ethernet offers a speed increase 10 times that of the 10BASE-T Ethernet specification while preserving such qualities as frame format, MAC mechanisms, and MTU. Such similarities allow the use of existing 10BASE-T applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification.
FCAPS	Five areas of network management: Fault management Configuration management Accounting management Performance management Security management
FDDI	Fiber Distributed Data Interface. LAN standard, defined by ANSI X3T9.5, specifying a 100-Mbps token-passing network using fiber-optic cable, with transmission distances of up to 2 km. FDDI uses a dual-ring architecture to provide redundancy.
FECN	forward explicit congestion notification. Bit set by a Frame Relay network to inform DTE receiving the frame that congestion was experienced in the path from source to destination. DTE receiving frames with the FECN bit set can request that higher-level protocols take flow-control action as appropriate.
FIB	Forwarding Information Base.
fiber optics	A method for the transmission of information (audio, video, data). Light is modulated and transmitted over high-purity, hair-thin fibers of glass. The bandwidth capacity of fiber-optic cable is much greater than that of conventional cable or copper wire.
FIFO	first-in, first-out. A simple queue service discipline, which serves the queue in the order of arrival. It is the default on high-speed interfaces (above 2.048 Mbps) in Cisco routers.
firewall	Routers, access servers, or dedicated devices at the edge of a network designated as an interface between connected networks. A firewall uses access lists and other methods to ensure the security of the protected network.
flash update	Routing update sent asynchronously in response to a change in the network topology.
flat addressing	Scheme of addressing that does not use a logical hierarchy to determine location. For example, MAC addresses are flat, so bridging protocols must flood packets throughout the network to deliver the packet to the appropriate location.
flooding	Traffic-passing technique used by switches and bridges in which traffic received on an interface is sent out all the interfaces of that device except the interface on which the information was received originally.
flow	Stream of data traveling between two endpoints across a network (for example, from one LAN station to another). Multiple flows can be transmitted on a single circuit.
FM	frequency modulation. Modulation technique in which different frequencies represent different data values within a signal path.
FQDN	fully qualified domain name. The full name of a system, rather than just its host name. For example, aldebaran is a host name, and aldebaran.interop.com is an FQDN.

Acronym	Definition
fragmentation	Process of breaking a packet into smaller units when transmitting over a network medium that cannot support the original size of the packet.
frame	Logical grouping of information sent as a data link layer unit over a transmission medium. Often refers to the header and the trailer, used for synchronization and error control, that surround the user data contained in the unit. The terms cell, datagram, message, packet, and segment also are used to describe logical information groupings at various layers of the OSI reference model and in various technology circles.
Frame Relay	A standard switched data link layer protocol that provides multiple virtual circuits to a single physical interface using an HDLC derived encapsulation between connected devices. Frame Relay is more bit efficient than X.25, the protocol for which it generally is considered a replacement.
FRF	Frame Relay Forum, an industry standards consortium.
FRF.11	Frame Relay Forum implementation agreement for Voice over Frame Relay. This specification defines multiplexed data, voice, fax, DTMF digit-relay, and CAS/robbed-bit signaling frame formats but does not include call setup, routing, or administration facilities.
FTP	File Transfer Protocol. Application protocol, part of the TCP/IP protocol stack, used for transferring files between network nodes. FTP is defined in RFC 959.
full mesh	Term describing a network topology in which devices are directly connected to every other device with either a physical circuit or a virtual circuit. A full mesh provides a great deal of redundancy. It can become prohibitively expensive to implement as the number of devices increases. This topology usually is reserved for network backbones.
FXO	Foreign Exchange Office. An interface that connects to the PSTN central office and is the interface offered on a standard telephone. Cisco's FXO interface is an RJ-11 connector that allows an analog connection at the PSTN central office or to a station interface on a PBX.
FXS	Foreign Exchange Station. An interface that connects directly to a standard telephone and supplies ring, voltage, and dial tone. Cisco's FXS interface is an RJ-11 connector that allows connections to basic telephone service equipment, keysets, and PBXs.
G.711	A 64-kbps PCM voice-coding technique, G.711-encoded voice is the expected format for digital voice delivery in the PSTN or through PBXs. Described in the ITU-T standard in its G series recommendations.
G.723	A compression technique that can be used for compressing speech signal components at a very low bit rate as part of the H.324 family of standards. This codec has two bit rates associated with it: 5.3 and 6.3 kbps. The higher bit rate is based on ML-MLQ technology and provides a somewhat higher quality of sound. The lower bit rate is based on CELP and provides system designers with additional flexibility. Described in the ITU-T standard in its G series recommendations.
G.729	A form of CELP compression where voice is coded into 8-kbps streams. There are four variations of this standard (G.729, G.729a, G.729b, and G.729ab) that differ mainly in computational complexity; all provide speech quality similar to 32-kbps ADPCM. G.729b and G.729ab also offer silence suppression using a VAD mechanism. Described in the ITU-T standard in its G series recommendations.
GARP	1. Group Address Registration Protocol, used in IEEE 802.1D/Q. 2. Generic Attribute Registration Protocol.

Acronym	Definition
gatekeeper	1. The component of an H.323 conferencing system that performs call address RAS bandwidth management. 2. Telecommunications: An H.323 entity on a LAN that provides address translation and control access to the LAN for H.323 terminals and gateways. The gatekeeper can provide other services to the H.323 terminals and gateways, such as bandwidth management and location of gateways. A gatekeeper maintains a registry of devices in the multimedia network. The devices register with the gatekeeper at start-up and request admission to a call from the gatekeeper.
gateway	In the IP community, an older term referring to a routing device. Today, the term router is used to describe nodes that perform this function, and “gateway” refers to a special-purpose device that performs an application-layer conversion of information from one protocol stack to another.
Gb	gigabit. Approximately 1,000,000,000 bits.
Gbps	gigabits per second.
Gigabit EtherChannel	A bundling of multiple Gigabit Ethernet links, which appear as one logical interface.
Gigabit Ethernet	Standard for a high-speed Ethernet, approved by the IEEE 802.3z standards committee in 1996.
GMRP	GARP Multicast Registration Protocol.
GPRS	general packet radio service. A service defined and standardized by the European Telecommunication Standards Institute (ETSI). GPRS is an IP packet-based data service for global system for mobile communication (GSM) networks.
grade of service	The probability that a call will be blocked while attempting to seize a circuit. It is written as a decimal fraction or Pxx, the blocking factor or blockage, where xx is the percentage of calls that are blocked for a traffic system. A common planning value in telecommunications traffic engineering is P05.
GRE	generic routing encapsulation. Tunneling protocol developed by Cisco that can encapsulate a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment. Defined in IETF RFC 2784.
ground-start signaling	A method of signaling used primarily on CO trunk lines to PBXs. A ground is placed on one side of the two-wire line to indicate that it is in use so that the other side of the two-wire interface does not attempt to use the line.
GUI	graphical user interface. A user environment that uses pictorial as well as textual representations of the input and the output of applications and the hierarchical or other data structure in which information is stored. Such conventions as buttons, icons, and windows are typical, and many actions are performed using a pointing device (such as a mouse). Microsoft Windows and the Apple Macintosh are prominent examples of platforms using a GUI.
H.225.0	An ITU standard that governs session establishment and packetization. H.225.0 actually describes several different protocols: RAS, use of Q.931, and use of RTP. It is specified for use by H.323
H.245	An ITU standard that governs endpoint control. It is specified for use by H.323
H.323	An ITU standard that allows dissimilar communication devices to communicate with each other by using a standardized communication protocol. H.323 specifies a common set of codecs, call setup and negotiating procedures, and basic data transport methods.

Acronym	Definition
HDSL	high-data-rate digital subscriber line. One of four DSL technologies. HDSL delivers 1.544 Mbps of bandwidth each way over two copper twisted pairs. Because HDSL provides T1 speed, telephone companies have been using HDSL to provision local access to T1 services whenever possible. The operating range of HDSL is limited to 12,000 feet (3658.5 meters), so signal repeaters are installed to extend the service. HDSL requires two twisted pairs, so it is deployed primarily for PBX network connections, digital loop carrier systems, interexchange POPs, Internet servers, and private data networks.
headend	Endpoint of a broadband network. All stations transmit toward the headend. The headend transmits toward the destination stations.
header	Control information placed before data when encapsulating that data for network transmission.
hello packet	Multicast packet that is used by routers for neighbor discovery and recovery. Hello packets also indicate that a client is still operating and network-ready.
hertz	Measure of frequency. Abbreviated Hz. Synonymous with cycles per second.
HIDS	Host Intrusion Detection System. Host-based security service that monitors and analyzes system events for the purpose of finding (and providing real-time or near real-time warning of) attempts to access system resources in an unauthorized manner.
hierarchical addressing	Scheme of addressing that uses a logical hierarchy to determine location. For example, IP addresses consist of network numbers, subnet numbers, and host numbers, which IP routing algorithms use to route the packet to the appropriate network.
hierarchical routing	A simplification of the complex problem of routing on large networks by reducing the size of the networks. This reduction is accomplished by breaking a network into a hierarchy of networks, where each level is responsible for its own routing.
high availability	An intelligent network service that, when carefully implemented, ensures adequate connectivity for mission-critical applications through fault tolerance, device redundancy, redundant physical connections, and route redundancy.
HMAC	Hash-based Message Authentication Code. A mechanism for message authentication using cryptographic hash functions. HMAC can be used with any iterative cryptographic hash function, for example, Message Digest 5 (MD5) or Secure Hash Algorithm 1 (SHA-1), in combination with a secret shared key. The cryptographic strength of HMAC depends on the properties of the underlying hash function.
holddown	State into which a route is placed so that routers neither advertise the route nor accept advertisements about the route for a specific length of time (the holddown period). Holddown is used to flush bad information about a route from all routers in the network. A route typically is placed in holddown when a link in that route fails.
HSRP	Hot Standby Router Protocol. Protocol that provides high network availability and transparent network topology changes. HSRP creates a hot standby router group with a lead router that services all packets sent to the hot standby address. The lead router is monitored by other routers in the group, and if it fails, one of these standby routers inherits the lead position and the hot standby group address. Described in IETF RFC 2281.
HTTP	Hypertext Transfer Protocol. The protocol used by web browsers and web servers to transfer files, such as text and graphics files.

Acronym	Definition
hub	1. Generally, a term used to describe a device that serves as the center of a star-topology network. 2. Hardware or software device that contains multiple independent but connected modules of network and internetwork equipment. Hubs can be active (where they repeat signals sent through them) or passive (where they do not repeat, but merely split, signals sent through them). 3. In Ethernet and IEEE 802.3, an Ethernet multiport repeater, sometimes called a concentrator.
IBGP	Internal Border Gateway Protocol. A BGP protocol used within an autonomous system.
ICMP	Internet Control Message Protocol. Network layer Internet protocol that reports errors and provides other information relevant to IP packet processing. Documented in RFC 792.
ICMP flood	Denial-of-service attack that sends a host more ICMP echo request ("ping") packets than the protocol implementation can handle.
ICND	*Interconnecting Cisco Network Devices.*
IDS	Intrusion Detection System. Security service that monitors and analyzes system events for the purpose of finding (and providing real-time or near real-time warning of) attempts to access system resources in an unauthorized manner.
IEEE	Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.
IEEE 802.1	IEEE specification that describes an algorithm that prevents bridging loops by creating a spanning tree. The original spanning tree algorithm was invented by Digital Equipment Corporation. The Digital algorithm and the IEEE 802.1 algorithm are not exactly the same, nor are they compatible.
IETF	Internet Engineering Task Force. Task force consisting of more than 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of Internet Society (ISOC).
IGMP	Internet Group Management Protocol. A protocol used by IP hosts to report their multicast group membership requests to an adjacent multicast router.
IGP	Interior Gateway Protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include IGRP, OSPF, and RIP.
IGRP	Interior Gateway Routing Protocol. IGP developed by Cisco to address issues associated with routing in large, heterogeneous networks.
IKE	Internet Key Exchange. A shared security policy that authenticates keys for services (such as IPSec). Before any IPSec traffic can be passed, each router/firewall/host must verify the identity of its peer. This verification can be done by manually entering preshared keys into both hosts or by a Certificate Authority (CA) service.
ILSG	Cisco Internet Learning Solutions Group.
ILT	instructor-led training.
IMA	inverse multiplexing over ATM. Standard protocol defined by the ATM Forum in 1997.
in-band signaling	Transmission within a frequency range normally used for information transmission.
Integrated IS-IS	Routing protocol based on the OSI routing protocol IS-IS but with support for IP and other protocols. Integrated IS-IS implementations send only one set of routing updates, making them more efficient than two separate implementations. Formerly called Dual IS-IS.

Acronym	Definition
intelligent network services	Services allow for application awareness within the network. Intelligent network services essentially add intelligence to the network infrastructure beyond just moving a datagram between two points. Examples of intelligent network services are network management, security, high availability, QoS, and IP multicasting.
interarea routing	Term used to describe routing between two or more logical areas.
Internet	The global internetwork, connecting tens of thousands of networks worldwide and having a "culture" that focuses on research and standardization based on real-life use. Many leading-edge network technologies come from the Internet community. The Internet evolved in part from Advanced Research Projects Agency Network (ARPANET). At one time, called the Defense Advanced Research Projects Agency (DARPA) Internet. Not to be confused with the general term internet.
Internet Connectivity module	A module within the Enterprise Composite Network Model. This module provides internal enterprise users with connectivity to external Internet services.
internetwork	Collection of networks interconnected by routers and other devices that functions (generally) as a single network. Sometimes called an internet, which is not to be confused with the Internet.
intra-area routing	Term used to describe routing within a logical area.
intranet	A closed, enterprise-wide network that includes LANs and WANs.
IP	Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless internetwork service. IP provides features for addressing, type of service specification, fragmentation and reassembly, and security. Defined in RFC 791.
IP address	32-bit address assigned to hosts using TCP/IP. An IP address is written as four octets separated by periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, and the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. Since 1993, classless interdomain routing (CIDR) has provided an alternate way of representing IP addresses. CIDR uses a prefix to separate the address into a network portion and a host portion, eliminating subnetworks. Also called an Internet address.
IP datagram	Fundamental unit of information passed across the Internet. Contains source and destination addresses along with data and a number of fields that define such things as the length of the datagram, the header checksum, and flags to indicate whether the datagram can be (or was) fragmented.
IPM	Internetwork Performance Monitor.
IP multicast	Internet Protocol multicast. Routing technique that allows IP traffic to be propagated from one source to a number of destinations or from many sources to many destinations. Rather than sending one packet to each destination, one packet is sent to a multicast group identified by a single IP destination group address.
IPng	IP next generation, the first name for the IPv6.
IP Phone	Device that allows communications across the IP network. IP telephones from Cisco are centrally managed by the CallManager and may be powered through their Ethernet connections.
IP precedence	Part of the IP header was implemented to provide prioritization. It is used in the router networks to make a more informed decision about routing the generated IP packets.
IPSec	Internet Protocol Security. A framework of standards that provides data confidentiality, data integrity, and data authentication between participating peers. IPSec provides these security services at the IP layer. IPSec uses IKE to handle the negotiation of protocols and algorithms based on local policy and to generate the encryption and authentication keys to be used by IPSec. IPSec can protect data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host.

Acronym	Definition
IP spoofing	A network attack that occurs when an attacker outside your network pretends to be a trusted user either by using an IP address that is within the range of IP addresses for your network or by using an authorized external IP address that you trust and to which you want to provide access to specified resources on your network. Should an attacker get access to your IPSec security parameters, that attacker can masquerade as the remote user authorized to connect to the corporate network.
IP telephony	Internet Protocol telephony. The transmission of voice calls over data networks that use the IP. IP telephony is the result of the transformation of the circuit-switched telephone network to a packet-based network that deploys voice compression algorithms and flexible and sophisticated transmission techniques, and to deliver services using only a fraction of the aggregate bandwidth of traditional digital telephony.
IPv4	Internet Protocol version 4.
IPv6	Internet Protocol version 6. Replacement for the current version of IP (version 4). IPv6 includes support for flow ID in the packet header, which can be used to identify flows. Formerly called IPng (IP next generation).
IP videoconferencing	Internet Protocol videoconferencing.
IPX	Internetwork Packet Exchange. NetWare network layer (Layer 3) protocol used for transferring data from servers to workstations. IPX is similar in purpose to IP.
ISDN	Integrated Services Digital Network. Communication protocol offered by telephone companies that permits telephone networks to carry data, voice, and other source traffic.
IS-IS	Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol based on DECnet Phase V routing, whereby ISs (routers) exchange routing information based on a single metric to determine network topology.
ISL	Inter-Switch Link. Cisco proprietary protocol that maintains VLAN information as traffic flows between switches and routers.
ISO	International Organization for Standardization. International organization that is responsible for a wide range of standards, including those relevant to networking. ISO developed the OSI reference model, a popular networking reference model.
ISP	Internet service provider. Company that provides Internet access to other companies and individuals.
ITU-T	International Telecommunication Union Telecommunication Standardization Sector. International body that develops worldwide standards for telecommunications technologies. The ITU-T carries out the functions of the former Consultative Committee for International Telegraph and Telephone (CCITT).
IVR	interactive voice response. Term used to describe systems that provide information in the form of recorded messages over telephone lines in response to user input in the form of spoken words or, more commonly, DTMF signaling. Examples include banks that allow you to check your balance from any telephone, and automated stock quote systems.
jitter	1. The interpacket delay variance, that is, the difference between interpacket arrival and departure. Jitter is an important QoS metric for voice and video applications. 2. Analog communication line distortion caused by the variation of a signal from its reference timing positions. Jitter can cause data loss, particularly at high speeds.
kbps	kilobits per second.
L2	Layer 2.
L2 switching (L2-switched)	Layer 2 switching. Switching based on Layer 2 (data link layer) information.

Acronym	Definition
L2TP	Layer 2 Tunneling Protocol. An IETF standards track protocol defined in RFC 2661 that provides tunneling of PPP. Based upon the best features of Layer 2 Forwarding (L2F) Protocol and PPTP, L2TP provides an industry-wide interoperable method of implementing VPDN. Communications transactions between the LAC and the LNS that support tunneling of a single PPP connection. There is a one-to-one relationship among the PPP connection, L2TP session, and L2TP call.
L3	Layer 3.
L3 switching (L3-switched)	Layer 3 switching. Emerging switching technology that integrates routing with switching to yield very high routing throughput rates in the millions-of-packets-per-second range. The movement to Layer 3 switching is designed to address the disadvantages of the current generation of Layer 2 switches, which functionally are equivalent to multiport bridges. These disadvantages include being subject to broadcast storms, spanning tree loops, and address limitations.
LAC	L2TP access concentrator. A node that acts as one side of an L2TP tunnel endpoint and is a peer to the L2TP network server (LNS). The LAC sits between an LNS and a remote system and forwards packets to and from each. Packets sent from the LAC to the LNS require tunneling with the L2TP protocol as defined in this document. The connection from the LAC to the remote system is either local or a PPP link.
LAN	local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.
LANE	LAN Emulation. Technology that allows an ATM network to function as a LAN backbone. The ATM network must provide multicast and broadcast support, address mapping (MAC-to-ATM), SVC management, and a usable packet format. LANE also defines Ethernet and Token Ring emulated LANs (ELANs).
latency	1. Delay between the time a device transmits a packet and when that packet is received at the destination. 2. Delay between the time that a device requests access to a network and the time that it is granted permission to transmit. 3. Delay between the time that a device receives a frame and the time that the frame is forwarded out the destination port.
LDCELP	low-delay code excited linear prediction compression. CELP voice compression algorithm providing 16 kbps, or 4:1 compression. Standardized in ITU-T Recommendation G.728.
leased line	Transmission line reserved by a communications carrier for the private use of a customer. A leased line is a type of dedicated line.
LEC	local exchange carrier. A telephone company that provides customer access to the worldwide public switched network through one of its central offices.
LFI	link fragmentation and interleaving. A solution for queuing delay situations. With LFI, large packets are fragmented into smaller frames and interleaved with small voice packets. Similar in effect to FRF.12, Frame Relay Fragmentation, available with Frame Relay.
LFIB	label forwarding information base. A data structure and way of managing forwarding in which destinations and incoming labels are associated with outgoing interfaces and labels.

Acronym	Definition
link-state routing protocols	Routing algorithm in which each router floods information regarding the cost of reaching each of its neighbors (link-state) to all nodes in the internetwork. Link-state algorithms create a consistent view of the network and therefore are not prone to routing loops. They achieve this at the cost of relatively greater computational difficulty, and more widespread traffic (compared with distance vector routing algorithms).
LLC	logical link control. The higher of the two data link layer sublayers defined by the IEEE. The LLC sublayer handles error control, flow control, framing, and MAC sublayer addressing. The most prevalent LLC protocol is IEEE 802.2, which includes both connectionless and connection-oriented variants.
LLQ	low latency queueing. Feature that brings strict priority queuing to CBWFQ. Strict priority queuing allows delay-sensitive data, such as voice, to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.
LNS	L2TP network server. A node that acts as one side of an L2TP tunnel endpoint and is a peer to the L2TP access concentrator (LAC). The LNS is the logical termination point of a PPP session that is being tunneled from the remote system by the LAC. Analogous to the Layer 2 Forwarding (L2F) home gateway (HGW).
load balancing, load sharing	In routing, the capability of a router to distribute traffic over all its network ports that are the same distance from the destination address. Good load-balancing algorithms use both line speed and reliability information. Load balancing increases the use of network segments, thus increasing effective network bandwidth.
local loop	Line from the premises of a telephone subscriber to the telephone company CO.
loop-start signaling	A method of signaling where a DC closure is applied to a phone line (loop), and the start of DC current flow indicates a change from on-hook to off-hook.
LSA	link-state advertisement. Broadcast packet used by OSPF that contains information about neighbors and path costs. LSAs are used by the receiving routers to maintain their routing tables. The equivalent in IS-IS is called an LSP.
MAC	Media Access Control. Lower of the two sublayers of the data link layer defined by the IEEE. The MAC sublayer handles access to shared media, such as whether token passing or contention will be used.
MAC address	Standardized data link layer address that is required for every port or device that connects to a LAN. Other devices in the network use these addresses to locate specific ports in the network and to create and update routing tables and data structures. MAC addresses are 6 bytes long and are controlled by the IEEE. Also known as a hardware address, MAC layer address, and physical address.
MAN	metropolitan-area network. Network that spans a metropolitan area. Generally, a MAN spans a larger geographic area than a LAN but a smaller geographic area than a WAN.
Management module	A module within the Enterprise Composite Network Model. This module performs intrusion detection, system logging, and TACACS+/RADIUS and OTP authentication, as well as network monitoring and general configuration management functions.
Mb	megabit. Approximately 1,000,000 bits.
Mbps	megabits per second. A bit rate expressed in millions of binary bits per second.
MCU	multipoint control unit.
MGCP	Media Gateway Control Protocol. A merging of Internet Protocol Device Control (IPDC) and Simple Gateway Control Protocol (SGCP). MGCP is defined in IETF RFC 2705.

Acronym	Definition
MIB	Management Information Base. Database of network management information that is used and maintained by a network management protocol, such as SNMP or Common Management Information Protocol (CMIP). The value of a MIB object can be changed or retrieved using SNMP or CMIP commands, usually through a GUI network management system. MIB objects are organized in a tree structure that includes public (standard) and private (proprietary) branches.
MIS	management information system.
MISTP	Multi-Instance Spanning Tree Protocol.
MLS	Multilayer Switching. Hardware-based switching based on Layer 3 and above.
modulation	Process by which the characteristics of electrical signals are transformed to represent information. Types of modulation include AM, FM, and PAM.
MPEG	Motion Picture Experts Group. Standard for compressing video. MPEG1 is a bit-stream standard for compressed video and audio optimized to fit into a bandwidth of 1.5 Mbps. MPEG2 is intended for higher-quality video on demand applications and runs at data rates between 4 and 9 Mbps. MPEG4 is a low-bit-rate compression algorithm intended for 64-kbps connections.
MPLS	Multiprotocol Label Switching. Switching method that forwards IP traffic using a label. This label instructs the routers and the switches in the network where to forward the packets based on preestablished IP routing information.
MRTG	Multi Router Traffic Grapher.
MTU	maximum transmission unit. Maximum packet size, in bytes, that a particular interface can handle.
multicast	Single packets sent to a multicast group, a specific subset of network devices, are copied within the network only as required to conserve bandwidth. The multicast group is specified in the destination address field.
multilayer switching	Hardware switching based on Layer 3 and above.
multimode fiber	A fiber-optic medium in which light travels in multiple modes.
MxU	Multi-Unit.
NANP	North American Numbering Plan.
NAS	1. network access server. Cisco platform (or collection of platforms) that interfaces between the packet world (for example, the Internet) and the circuit world (for example, the PSTN). 2. network attached storage.
NAT	Network Address Translation. Mechanism for reducing the need for globally unique IP addresses. NAT allows an organization with addresses that are not globally unique to connect to the Internet by translating those addresses into globally routable address space. Also known as Network Address Translator.
NBAR	network-based application recognition. IOS software commands used for network traffic analysis.
NBMA	nonbroadcast multiaccess. Term describing a multiaccess network that either does not support broadcasting (such as Frame Relay) or in which broadcasting is not feasible (for example, an SMDS broadcast group or an extended Ethernet that is too large).
NetFlow	network flow. A feature of some routers that allows them to categorize incoming packets into flows. Because packets in a flow often can be treated in the same way, this classification can be used to bypass some of the work of the router and accelerate its switching operation.

Acronym	Definition
NFS	Network File System. As commonly used, a distributed file system protocol suite developed by Sun Microsystems that allows remote file access across a network. In actuality, NFS is simply one protocol in the suite. NFS protocols include NFS, Remote Procedure Call (RPC), eXternal Data Representation (XDR), and others. These protocols are part of a larger architecture that Sun refers to as Open Network Computing (ONC).
NIC	network interface card. Board that provides network communication capabilities to and from a computer system. Also called an adapter.
NIDS	Network Intrusion Detection System. Security service that monitors and analyzes system events for the purpose of finding (and providing real-time or near real-time warning of) attempts to access system resources in an unauthorized manner.
NMS	network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer, such as an engineering workstation. NMSs communicate with agents to help keep track of network statistics and resources.
nonrepudiation service	Security service that provides protection against false denial of involvement in a communication.
nonstub area	OSPF area that carries a default route, static routes, intra-area routes, interarea routes, and external routes. Nonstub areas are the only OSPF areas that can have virtual links configured across them, and are the only areas that can contain an ASBR.
NSAP	network service access point. Network addresses, as specified by ISO. An NSAP is the point at which OSI network service is made available to a transport layer (Layer 4) entity.
NTP	Network Time Protocol. Protocol built on top of TCP that ensures accurate local timekeeping with reference to radio and atomic clocks located on the Internet. This protocol is capable of synchronizing distributed clocks within milliseconds over long time periods. Defined in IETF RFC 1305.
ODR	On-Demand Routing.
OPNET	Simulation tools used in the DESGN course.
OSI	Open System Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.
OSI protocol stack	Set of related communications protocols that operate together and, as a group, address communication at some or all of the seven layers of the OSI reference model. Not every protocol stack covers each layer of the model, and often a single protocol in the stack addresses a number of layers at once. TCP/IP is a typical protocol stack.
OSI reference model	Open System Interconnection reference model. Network architectural model developed by ISO and ITU-T. The model consists of seven layers, each of which specifies particular network functions, such as addressing, flow control, error control, encapsulation, and reliable message transfer. The lowest layer (the physical layer) is closest to the media technology. The lower two layers are typically implemented in hardware and firmware whereas the upper five layers are implemented in software. The highest layer (the application layer) is closest to the user. The OSI reference model is used as a framework for teaching and understanding network functionality.
OSPF	Open Shortest Path First. Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing. OSPF was derived from an early version of the IS-IS protocol. Defined in IETF RFC 2328.
OTP	One Time Password.

Acronym	Definition
out-of-band signaling	Transmission using frequencies or channels outside the frequencies or channels normally used for information transfer. Out-of-band signaling often is used for error reporting in situations in which in-band signaling can be affected by whatever problems the network might be experiencing.
p2mp	point-to-multipoint. Communication between a series of wireless receivers and transmitters to a central location. Cisco p2mp typically is set up in three segments to enable frequency reuse.
p2p	point-to-point. Wireless communication between one receiver and one transmitter. p2p provides higher bandwidth than p2mp because it requires less overhead to manage the data paths and because there is only one receiver per transmitter.
PAM	pulse amplitude modulation. Modulation scheme in which the modulating wave is caused to modulate the amplitude of a pulse stream.
PAP	Password Authentication Protocol. Authentication protocol that allows PPP peers to authenticate one another. The remote router attempting to connect to the local router is required to send an authentication request. Unlike CHAP, PAP passes the password and the host name or user name in the clear (unencrypted). PAP does not itself prevent unauthorized access but merely identifies the remote end. The router or access server then determines whether that user is allowed access. PAP is supported only on PPP lines.
partial mesh	Network in which devices are connected in less than a full-mesh topology. In a well-connected partial mesh all nodes are connected to at least two other nodes. This does not provide the level of redundancy of a full-mesh topology, but is less expensive to implement. Partial-mesh topologies generally are used to connect the peripheral networks that distribute traffic to a fully meshed backbone.
password sniffing	Passive wiretapping, usually on a local-area network, to gain knowledge of passwords.
PAT	port address translation. Translation method that allows the router to forward packets from several sessions or flows between a private internetwork and the Internet. Using a private IP network in the enterprise in conjunction with the PAT feature where Internet connectivity is required reduces the requirement for registered IP addresses.
PBX	private branch exchange. Digital or analog telephone switch located on the subscriber premises and used to connect private and public telephone networks.
PCM	pulse code modulation. Technique of encoding analog audio into a 64-Kb data stream by sampling with 8-bit values at a rate of 8000 times per second.
PDIOO	planning, design, implementation, operation, and optimization. A network life-cycle methodology that reflects the evolution of network activities over time.
PDN	public data network. Network operated either by a government (as in Europe) or by a private concern to provide computer communications to the public, usually for a fee. PDNs enable small organizations to create a WAN without the equipment costs of long-distance circuits.
pilot network	A design tested on a part of an existing live network.
PIM	Protocol Independent Multicast. Multicast routing architecture that allows the addition of IP multicast routing on existing IP networks. PIM is unicast routing protocol independent and can be operated in two modes: dense and sparse.
POP	point of presence. A physical location within a service provider network where users dial in.
PortFast	Feature used on switched ports where only end-user stations are directly connected. There is no delay in passing traffic, because the switch immediately puts the port to the forward state. It reduces the number and duration of SPT convergence events.
POS	packet over SONET/SDH.
POTS	plain old telephone service.

Acronym	Definition
PPP	Point-to-Point Protocol. Successor to Serial Line Internet Protocol (SLIP) that provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. Whereas SLIP was designed to work with IP, PPP was designed to work with several network layer protocols, such as IP, IPX, and AppleTalk Remote Access (ARA). PPP also has built-in security mechanisms, such as CHAP and PAP. PPP relies on two protocols: link control protocol (LCP) and Network Control Protocol (NCP). Defined in IETF RFC 1661.
pps	packet per second.
PPTP	Point-to-Point Tunneling Protocol. RFC 2637 describes the PPTP protocol.
PQ	priority queuing. Queuing that prioritizes traffic in the network. Four traffic priorities can be configured. A series of filters based on packet characteristics (source IP address and port) is defined to cause the router to place critical traffic in the highest queue and other traffic in the lower three queues. The queue with the highest priority is serviced first until empty; the lower queues are then serviced in sequence.
PQ-CBWFQ	priority queuing-class-based weighted fair queuing. Feature that brings strict priority queuing to CBWFQ. Strict priority queuing allows delay-sensitive data, such as voice, to be dequeued and sent first (before packets in any other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic. Also called low latency queueing (LLQ).
PQ-WFQ	priority queuing-weighted fair queuing.
PRI	Primary Rate Interface. ISDN interface to primary rate access. Primary rate access consists of a single 64-kbps D channel plus 23 (T1) or 30 (E1) B channels for voice or data.
prototype network	A separate (nonlive) network used to test network design and feature selections.
PSTN	Public Switched Telephone Network. General term referring to the variety of telephone networks and services in place worldwide. It includes POTS and ISDN services offered by telecommunications companies and agencies.
PVC	permanent virtual circuit. Virtual circuit that is permanently established. PVCs save equipment and operations costs associated with circuit establishment and teardown. In ATM terminology, called a permanent virtual connection.
PVST	Per VLAN Spanning Tree. Support for IEEE 802.1q trunks to map multiple spanning trees to a single spanning tree.
Q.931	ITU-T specification for signaling to establish, maintain, and clear ISDN network connections.
QLLC	Qualified Logical Link Control. Data link layer protocol defined by IBM that allows SNA data to be transported across X.25 networks.
QoS	quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.
QPPB	QoS Policy Propagation on BGP. Feature that classifies packets by IP precedence based on BGP community lists, BGP autonomous system paths, and access lists. After a packet is classified, other quality of service features such as committed access rate (CAR) and weighted random early detection (WRED) can specify and enforce policies to fit a business model.
QSIG	Q Signaling. Signaling standard. Common channel signaling protocol based on ISDN Q.931 standards and used by many digital PBXs.
queue	Generally, a time-ordered list of elements waiting to be processed.
queuing delay	Amount of time that data must wait before it can be transmitted onto a shared physical circuit.

Acronym	Definition
RADIUS	Remote Authentication Dial-In User Service. Database for authenticating modem and ISDN connections and for tracking connection time.
RARP	Reverse Address Resolution Protocol. Protocol in the TCP/IP stack that provides a method for finding IP addresses based on MAC addresses. Defined in IETF RFC 903.
RAS	1. registration, admission, and status protocol. Protocol that is used in H.323 between endpoints and the gatekeeper to perform management functions. The RAS signaling function performs registration, admissions, bandwidth changes, status, and disengage procedures between a VoIP gateway and the gatekeeper. 2. Remote Access Server
RDP	Router Discovery Protocol.
RED	random early detection.
Remote Access and VPN module	A module within the Enterprise Composite Network Model. This module terminates VPN traffic from remote users' remote sites and forwarded from the Internet Connectivity module.
RFC	Request for Comments. Document series published by the IETF used as the primary means for communicating information about the Internet. Some RFCs are designated as Internet standards. Many RFCs document protocol specifications, such as Telnet and FTP. Others suggest best practices. Some are humorous or historical. RFCs are available online from numerous sources.
RFI	request for information.
RFP	request for proposal.
RGMP	Router-Port Group Management Protocol.
RIO	Reusable Information Object.
RIP	Routing Information Protocol. IGP supplied with UNIX Berkeley Standard Distribution (BSD) systems. The most common IGP in the Internet. RIP uses hop count as a routing metric. Originally defined in RFC1058, RIP Version 2 is defined in IETF RFC 2453.
RIPv1	Routing Information Protocol Version 1.
RIPv2	Routing Information Protocol Version 2.
RLO	Reusable Learning Object.
RMON	Remote Monitoring. MIB agent specification described in RFC 1271 that defines functions for the remote monitoring of networked devices. The RMON specification provides numerous monitoring, problem detection, and reporting capabilities.
root bridge	The root, or start, of the spanning tree in a switched network. It exchanges topology information with designated bridges in a spanning-tree implementation to notify all other bridges in the network when topology changes are required. This exchange prevents loops and provides a measure of defense against link failure.
routing protocols	Protocols that accomplish routing through the implementation of a specific routing algorithm. Examples of routing protocols include IGRP, OSPF, and RIP.
RSP	Route Switch Processor. Processor module in the Cisco 7500 Series routers that integrates the functions of the route processor (RP) and the switch processor (SP).
RSRB	remote source-route bridging. SRB over WAN links.

Acronym	Definition
RSVP	Resource Reservation Protocol. Protocol that supports the reservation of resources across an IP network. Applications running on IP end systems can use RSVP to indicate to other nodes the nature (bandwidth, jitter, maximum burst, and so on) of the packet streams that they want to receive. RSVP depends on IPv6. Also known as Resource Reservation Setup Protocol.
RTCP	RTP Control Protocol. Protocol that monitors the QoS of an IPv6 RTP connection and conveys information about the ongoing session.
RTP	Real-Time Transport Protocol. Protocol commonly used with IP networks. RTP is designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulations, over multicast or unicast network services. RTP provides such services as payload type identification, sequence numbering, time stamping, and delivery monitoring to real-time applications.
SAA	Service Assurance Agent.
SAFE	Security Architecture for Enterprise.
SDH	Synchronous Digital Hierarchy is a standard technology for synchronous data transmission on optical media. It is the international equivalent of Synchronous Optical Network.
SDSL	single-line digital subscriber line. One of four DSL technologies. SDSL delivers 1.544 Mbps both downstream and upstream over a single copper twisted pair. The use of a single twisted pair limits the operating range of SDSL to 10,000 feet (3048.8 meters).
Server Farm module	A module within the Enterprise Composite Network Model. It contains internal e-mail and corporate servers providing application, file, print, e-mail, and DNS services to internal users.
service level	Various levels and quality of services defined for each service type. For example, the service type called quality of sound might have service levels defined for telephone, broadcast, and digital CD.
Service Provider Edge	A functional area described within the Enterprise Composite Network Model. The modules in this area are not implemented by the enterprise itself but are necessary to enable communication with other networks, and most often use different WAN technologies and ISPs.
SG	Student Guide.
SIP	session initiation protocol. Protocol developed by the IETF MMUSIC Working Group as an alternative to H.323. SIP features were originally defined in IETF RFC 2543, published in March 1999. SIP equips platforms to signal the setup of voice and multimedia calls over IP networks.
SLA	service level agreement.
SLARP	Serial Line Address Resolution Protocol.
SLB	server load balancing.
SLM	service-level management.
SMDS	Switched Multimegabit Data Service. High-speed, packet-switched, datagram-based WAN networking technology offered by the telephone companies.
SMI	Structure of Management Information. Document (RFC 1155) specifying rules used to define managed objects in the MIB.
SMTP	Simple Mail Transfer Protocol. Internet protocol providing e-mail services.
SNA	Systems Network Architecture. Large, complex, feature-rich network architecture developed in the 1970s by IBM. Similar in some respects to the OSI reference model but with a number of differences. SNA essentially is composed of seven layers.

Acronym	Definition
SNAP	Subnetwork Access Protocol. Internet protocol that operates between a network entity in the subnetwork and a network entity in the end system. SNAP specifies a standard method of encapsulating IP datagrams and ARP messages on IEEE networks. The SNAP entity in the end system makes use of the services of the subnetwork and performs three key functions: data transfer, connection management, and QoS selection.
SNMP	Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices and to manage configurations, statistics collection, performance, and security.
SONET	Synchronous Optical Network. A standard format for transporting a wide range of digital telecommunications services over optical fiber. SONET is characterized by standard line rates, optical interfaces, and signal formats.
SP	service provider.
SPF	shortest path first algorithm, or Dijkstra's algorithm. Routing algorithm that iterates on the length of the path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms.
SPI	security parameter index. A number that, together with a destination IP address and security protocol, uniquely identifies a particular security association. When using IKE to establish the security associations, the SPI for each security association is a pseudorandomly derived number. Without IKE, the SPI is manually specified for each security association.
spoofing	The act of a packet illegally claiming to be from an address from which it was not actually sent. Spoofing is designed to foil network security mechanisms, such as filters and access lists.
SS7	Signaling System 7. Standard CCS system used throughout the PSTN with POTS and ISDN. Developed by Bellcore.
SSH	Secure Shell Protocol.
static route	Route that is explicitly configured and entered into the routing table. Static routes take precedence over routes chosen by dynamic routing protocols.
statistical multiplexing	A technique for sharing line capacity in which packets or cells from multiple flows are interleaved on shared media with access conflicts resolved by weighted.
Storage Networking	Technology offering universal access to storage solutions and products in a standards-based architecture. Storage Networking combines intelligent Fiber Channel, Ethernet, and optical networking offerings to build scalable data center storage networks and extend storage networks through IP.
STP	Spanning Tree Protocol. Bridge protocol that uses the spanning-tree algorithm, enabling a learning bridge to dynamically work around loops in a network topology by creating a spanning tree. Bridges exchange BPDU messages with other bridges to detect loops and then remove the loops by shutting down selected bridge interfaces. Refers to both the IEEE 802.1 Spanning Tree Protocol standard and the earlier Digital Equipment Corporation Spanning Tree Protocol upon which it is based. The IEEE version supports bridge domains and allows the bridge to construct a loop-free topology across an extended LAN. The IEEE version generally is preferred over the Digital version. Sometimes abbreviated as STP.
stub area	OSPF area that carries a default route, intra-area routes, and interarea routes but does not carry external routes. Virtual links cannot be configured across a stub area, and they cannot contain an ASBR.
STUN	serial tunnel. Router feature allowing two Synchronous Data Link Control (SDLC) or High-Level Data Link Control (HDLC) compliant devices to connect to one another through an arbitrary multiprotocol topology (using Cisco routers) rather than through a direct serial link.

Acronym	Definition
subnet	In IP networks, a network sharing a particular subnet address. Also called a subnetwork, they are networks explicitly segmented by a network administrator to provide a multilevel, hierarchical routing structure while shielding the subnetwork from the addressing complexity of attached networks.
supernet	Aggregation of IP network addresses advertised as a single classless network address. For example, given four Class C IP networks—192.0.8.0, 192.0.9.0, 192.0.10.0, and 192.0.11.0—each having the intrinsic network mask of 255.255.255.0, one can advertise the network address 192.0.8.0 with a subnet mask of 255.255.252.0.
SVC	switched virtual circuit. Virtual circuit that is dynamically established on demand and is torn down when transmission is complete. SVCs are used in situations where data transmission is sporadic.
switch	1. Network device that filters, forwards, and floods frames based on the destination address of each frame. The switch operates at the data link layer of the OSI model. 2. General term applied to an electronic or mechanical device that allows a connection to be established as necessary and terminated when there is no longer a session to support. 3. In telephony, a general term for any device, such as a PBX, that connects individual phones to phone lines. See also PBX and PSTN.
switched LAN	LAN implemented with LAN switches.
switching	Process of taking an incoming frame from one interface and delivering it through another interface. Routers use Layer 3 switching to route a packet, and Layer 2 switches use Layer 2 switching to forward frames. See also Layer 2 switching and Layer 3 switching.
SYN flood	Denial-of-service attack that sends more TCP SYN packets than the protocol implementation can handle.
syslog	Dedicated server that logs system messages.
T1	Digital WAN carrier facility. T1 transmits DS-1-formatted data at 1.536 Mbps using a line rate of 1.544Mbps. 8 Kbps are used for framing. Coding options are alternate mark inversion (AMI) or binary 8-zero substitution (B8ZS).
TACACS+	Terminal Access Controller Access Control System plus. Authentication protocol, developed by Cisco that provides remote access authentication and related services, such as event logging. User passwords are administered in a central database rather than in individual routers, providing an easily scalable network security solution.
Tbps	terabits per second.
TCP	Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack. Defined in IETF RFC 793.
TDM	time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on common serial media based on preassigned time slots. Bandwidth is allocated to each channel regardless of whether there is data to transmit. Compare to statistical multiplexing.
Telnet	Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection, enabling users to log in to remote systems and use resources as if they were connected to a local system. Telnet is defined in RFC 854.
TFTP	Trivial File Transfer Protocol. Simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, user name and password). TFTP is defined in IETF RFC 1350.

Acronym	Definition
tie line	Connects enterprise PBXs together without requiring PSTN switch intervention. Tie lines allow implementation of private telephone networks over private and leased lines.
touch-tone	Term referring to the presence of push buttons that produce tones corresponding to numbers. Use the term as an adjective, not a noun; for example, touch-tone telephone, touch-tone telephone buttons, and so on.
traffic policing	Process used to measure the actual traffic flow across a given connection and compare it to the configured traffic flow for that connection. Traffic in excess of the agreed flow rate can be tagged (where some bit is set to 1) for discard en route if congestion develops or can be dropped immediately. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also known as admission control, permit processing, and rate enforcement. It is most often implemented on ingress ports to protect a transport network from greedy traffic flows.
traffic shaping	Use of queues to limit surges that can congest a network. Data is buffered and then sent into the network in regulated amounts to ensure that the traffic fits within the promised traffic envelope for the particular connection. Traffic shaping is used in ATM, Frame Relay, and other types of networks. Also known as metering, shaping, and smoothing. It is most often implemented on egress ports to ensure compliance with agreed connection traffic rates to avoid traffic policing.
Trojan horse	Computer program that appears to have a useful function but also has a hidden and potentially malicious function that evades security mechanisms, sometimes by exploiting legitimate authorizations of a system entity that invokes the program.
trunk	1. Physical and logical connection between two switches across which network traffic travels. A backbone is composed of a number of trunks. 2. In telephony, a phone line between two COs or between a CO and a PBX.
TTL	Time to Live.
tunneling	A technique that encapsulates frames of one protocol as data within frames of some other protocol. Typically both protocols are described at the same layer within the OSI/ISO reference model.
twisted pair	Twisted pair describes copper media in which the wires are twisted around each other in a spiral to reduce crosstalk or electromagnetic induction between the pairs of wires.
type of service	Type of service as defined in IETF RFC 1349, uses an octet in the IP header to provide CoS. As a service it has been obsolete in IP by differentiated services as defined in IETF RFC 2474, which uses the same bits in the IP header.
UBR	unspecified bit rate. Service category defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network, but there are no guarantees in terms of cell loss rate and delay.
UDP	User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols. UDP is defined in RFC 768.
µ-law	Companding technique commonly used in North America. mu-law is standardized in ITU-T G.711. Compare to a-law.
UMTS	Universal Mobile Telecommunications Service. A 3G mobile wireless telecommunications system whose standards are being developed by the 3rd Generation Partnership Project (3GPP).
unicast	Message sent to a single network destination.
UNIX	Operating system developed in 1969 at Bell Laboratories. UNIX has gone through several iterations since its inception. These iterations include UNIX 4.3 BSD, developed at the University of California at Berkeley, and UNIX System V, Release 4.0, developed by AT&T.

Acronym	Definition
UplinkFast	A spanning-tree maintenance mechanism that enables a switch to put a redundant path (port) into active state within a second.
UTP	unshielded twisted-pair. Four-pair wire medium used in a variety of networks. UTP does not require the fixed spacing between connections that is necessary with coaxial-type connections. Six types of UTP cabling are commonly used.
VACL	VLAN access control list.
VAD	voice activity detection. Technology that detects silence in the voice path and transmits only periods of audio activity across the connection. When VAD is enabled, the sound quality is slightly degraded but the connection consumes much less bandwidth.
VBR	variable bit rate. Service category defined by the ATM Forum for ATM networks. VBR is subdivided into a real-time class and nonreal-time class. VBR (RT) is used for connections in which there is a fixed timing relationship between samples. VBR (NRT) is used for connections in which there is no fixed timing relationship between samples but still need a guaranteed minimum bit rate.
VDSL	very-high-data-rate digital subscriber line. One of four DSL technologies. VDSL delivers 13 to 52 Mbps downstream and 1.5 to 2.3 Mbps upstream over a single twisted copper pair. The operating range of VDSL is limited to 1000 to 4500 feet (304.8 to 1372 meters).
VLAN	virtual local-area network. Group of devices that are configured (using management software) so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. Because VLANs are based on logical instead of physical connections, they are extremely flexible.
VLSM	variable-length subnet masking. Capability to specify a different subnet mask for the same network number on different subnets. VLSM can help optimize available address space.
VoATM	Voice over Asynchronous Transfer Mode. Technology that enables a router to carry voice traffic (for example, telephone calls and faxes) over an ATM network. When sending voice traffic over ATM, the router encapsulates voice traffic using a special AAL5 adaptation for multiplexed voice.
VoD	video on demand. System using video compression to supply video programs to viewers when requested via ISDN or cable.
VoFR	Voice over Frame Relay. Technology that enables a router to carry voice traffic (for example, telephone calls and faxes) over a Frame Relay network. When sending voice traffic over Frame Relay, the router segments voice traffic and frames it for transit using FRF.12 encapsulation.
voice mail	Voice messaging, an optional service for customers. Voice mail provides customers with the ability to divert their incoming PSTN calls to a voice mailbox when they are unable to answer their telephones.
VoIP	Voice over Internet Protocol. The capability to carry voice over an IP-based internet with POTS-like functionality, reliability, and voice quality. The voice packets are transported using RTP over UDP. Skinny Client Control Protocol (SCCP), H.323, and SIP provide session (call) control.
VPDN	virtual private dial-up network. A network that extends remote access to a private network using a shared infrastructure. VPDNs use Layer 2 tunnel technologies (L2F, L2TP, and PPTP) to extend Layer 2 and higher parts of the network connection from a remote user across an ISP network to a private network. VPDNs are a cost-effective method of establishing a long-distance, point-to-point connection between remote dial users and a private network. Also known as virtual private dial network.
VPN	Virtual Private Network. A network that enables IP traffic to travel securely over a public or shared TCP/IP network by encrypting all traffic between VPN access points to the public network. A VPN uses tunneling to transport the encrypted information at the IP level.

Acronym	Definition
VTP	VLAN Trunk Protocol.
vty	virtual terminal line.
WAN	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay, SMDS, and X.25 are examples of WANs.
WAN module	A module within the Enterprise Composite Network model. The WAN module uses different WAN technologies for carrying routed traffic between geographically separated sites.
WDM	wavelength division multiplexing. Optical transmission of signals in which multiple optical wavelengths share the same transmission fiber. The spectrum occupied by each channel must be adequately separated from the others.
web	World Wide Web (also called WWW). A client-server system based on HTML and HTTP.
WFQ	weighted fair queuing. Congestion management queue service algorithm that identifies conversations (in the form of traffic streams), separates packets that belong to each conversation, and ensures that capacity is shared fairly between these individual conversations. WFQ is an automatic way of stabilizing network behavior during congestion and results in increased performance and reduced retransmission. It is the default weighted discipline on interfaces that operate at 2,048 Mbps or less in Cisco routers.
wildcard mask	A 32-bit value used in conjunction with an IP address to determine which bits in an IP address should be ignored when comparing that address with another IP address. A wildcard mask is specified when an access list is configured.
wireline	Term that refers to standard telephone and data communications systems that use copper or fiber-optic cables in contrast to wireless, cellular, and satellite services. Also called “landline” or “land-based.”
wiring closet	Specially designed room used for wiring a data or voice network. Wiring closets serve as a central junction point for the wiring and the wiring equipment that is used for interconnecting devices.
workgroup	Collection of workstations and servers on a LAN that communicate and exchange data with one another while performing application activities.
WRED	weighted random early detection. Queuing method that ensures that high-precedence traffic has lower loss rates than other traffic during times of congestion.
X.25	ITU standard for defining how connections between DTE and DCE are maintained for remote terminal access and computer communications in PDNs.
xDSL	Group term used to refer to ADSL, HDSL, SDSL, and VDSL. All are digital technologies using the existing copper infrastructure provided by the telephone companies. xDSL is a high-speed alternative to ISDN.
XML	extensible markup language. A standard maintained by the World Wide Web Consortium (W3C). It defines a syntax that lets you use a markup language to specify information structures. Information structures define the type of information; for example, subscriber name or address, not how the information looks (bold, italic, and so on). External processes can manipulate these information structures and publish them in a variety of formats. This text markup language is designed to enable the use of Standardized Generalized Markup Language (SGML) on the World Wide Web. XML allows you to define your own customized extensions.
zone	Collection of all terminals, gateways, and multipoint control units (MCUs) managed by a single gatekeeper. A zone includes at least one terminal and can include gateways or MCUs. A zone can be independent of LAN topology and can consist of devices on multiple LAN segments connected using routers or other devices.

DESGN

Case Study Solutions

The lesson case study solutions are contained here. This is a presentation of all course case study items, with the exception of any e-learning drag and drop exercises which cannot be replicated in a printed form.

Applying a Methodology to Network Design

Case Study 1: Network Upgrade

Based on the scenario, the following tables include the proposed solutions. According to the Case Study Guidelines, there may be some minor variations in your solutions.

Case Study Solutions

Each of the steps that require solutions is listed below.

Step 1 The missing information that is needed in the design that the company did not present includes the items listed in this table. Before starting with the detailed design you should determine this information.

Missing Items

Missing Items	Comments
Type of the current WAN	There is no information on whether the WAN links to regional offices are leased lines, or some type of public network such as Frame Relay.
Current bandwidth of the links to the regional offices and to the Internet	The existing bandwidths are needed for evaluating the effect of the new applications and for planning new capacities.
WAN backup	There is no information on whether WAN links are backed up or not. The information is needed for designing resilience in the network.
Current IP addressing scheme	Although the information says that a flat addressing scheme is implemented, this is not sufficient. At a minimum the address range / address class is required.
Technical constraints (availability of certain WAN services at the regional offices, availability of the Internet at the planned new remote offices)	Although it is a designers' job to determine the availability of the WAN and Internet connectivity options, it is highly recommended to discuss the current situation with the customer and to extract this information.
Budget available for new solutions	The customer has not provided any information on the available budget. The available budget definitely affects the proposed solution and at least some hints on the budget are extremely helpful for a designer.
Responsible people	The customer has not supplied any contacts. You should determine at least the technical and business contacts for a project.
Business constraints, such as policies and goals, and criticality of applications	The customer has not mentioned any business constraints. The current vendor of the network equipment is unknown and no preferences are given. It is recommended that you determine the business constraints during the initial design phases.

Step 2 The table lists the major design areas that you have to address in your design project according to the given scenario). For each of the areas, your task is to evaluate possible solutions and propose the most optimal one based on the customer requirements. You will document your solution and propose the implementation plan.

Major Design Tasks

Major design areas	Comments
Redesign IP addressing	The flat addressing scheme and RIP as routing protocols are certainly not the features of scalable growing networks. New hierarchical addressing is required.
Redesign Campus LAN	The current campus LAN is shared and interconnects two buildings. Because there is also no redundancy, the designer needs to entirely redesign the campus, including the placement of servers.
Upgrade WAN links	The upgrade of the WAN links is essential because, according to the company, the current bandwidth seems insufficient. The introduction of new applications will result in a higher load because the existing applications will remain.
Design new routing protocol	The company is aware of the drawbacks of RIP. You should replace RIP with a routing protocol that is more scalable and that better fits into the planned hierarchical addressing scheme.
Integrate international offices into the company network	The two international offices that will open soon will likely use an Internet and VPN implementation. Additionally, the designer needs to propose a voice solution for these two offices.

Step 3 Depending on the major design areas you will address in your design project, you could use your simulation tool for evaluation of the solutions listed in this table. The simulation tool will allow you to simplify the prototype or pilot for this project.

Possible Simulation Scenarios

Possible simulation	Comments
Effect of new applications on the existing links	Based on the applications mix and information on their users, you could simulate the load on the WAN links. Before the simulation you must determine the current bandwidths of the links.
Comparison of shared versus switched LAN	As the customer wants some proof of all the benefits that the switched LAN will bring, you could simulate both scenarios and compare the results.
Switched campus LAN solution	You plan to completely redesign the campus LAN and introduce a switching solution that includes redundancy. You can use a simulation tool to check the effects of redundancy and link usage in the campus solution.

Possible simulation	Comments
Routing convergence	Due to your plan to replace RIP with a hierarchical protocol (possibly OSPF), you could examine the convergence of the protocol in the new campus. You could compare various campus scenarios including multilayer switching everywhere or only in some parts. You can also examine the load sharing options.
Voice over IP	For the international offices you can simulate the effect of adding voice on top of the data traffic.

If your answers significantly differ from the sample solutions, please justify them.

Simulation 1: New Applications

The simulation of the WAN link utilization was completed for this exercise. You were asked to answer a few questions. The following answers may differ from your answers in minor details. If there are significant differences, please explain your reasoning in your answers.

Q1) Observe the directions in which the load was higher. What can you determine from the results?

The utilization of the WAN links exceeded the predefined threshold in direction of headquarters towards regional offices. In the opposite direction, the traffic was light. These results indicate that the regional offices access the central servers and that the responses represent the majority of the traffic. A certain amount of the traffic also results from accessing the Internet via the central site.

Q2) What can you determine from the results? Compare the planned number of users and applications for each of the regional offices. In which direction are the links saturated?

The results show that with an increased number of applications and their respective users, the amount of traffic increases and leads to saturation on the San Jose – Houston link in the outbound direction.

Q3) When you compare the results from the initial traffic simulation with the results from the simulations of the new applications, you observe that the traffic from Denver to Headquarters is now significant. Why?

The users in Denver are mainly engineers. From their profile, it is evident that FTP is the major application they use. The heavy load in the outbound direction from Denver to the Headquarters is probably the result of uploading files to the central location.

Q4) What can you determine from the graphs?

The graphs show that, for approximately 25 percent of all web pages, the response time exceeded 10 seconds.

Q5) What can you observe from the graphs?

The link utilization in the outbound direction was jittering at approximately the 60 percent mark. The response time for over 90 percent of all web pages (HTTP requests) was below 10 seconds. For around 75 percent of all HTTP requests, the response time was below 5 seconds.

Structuring and Modularizing the Network

Case Study 2: Designing a Network Hierarchy

Based on the scenario, the following diagrams include the proposed solutions. According to the Case Study Guidelines, there may be some minor variations in your solutions.

Company Headquarters Network

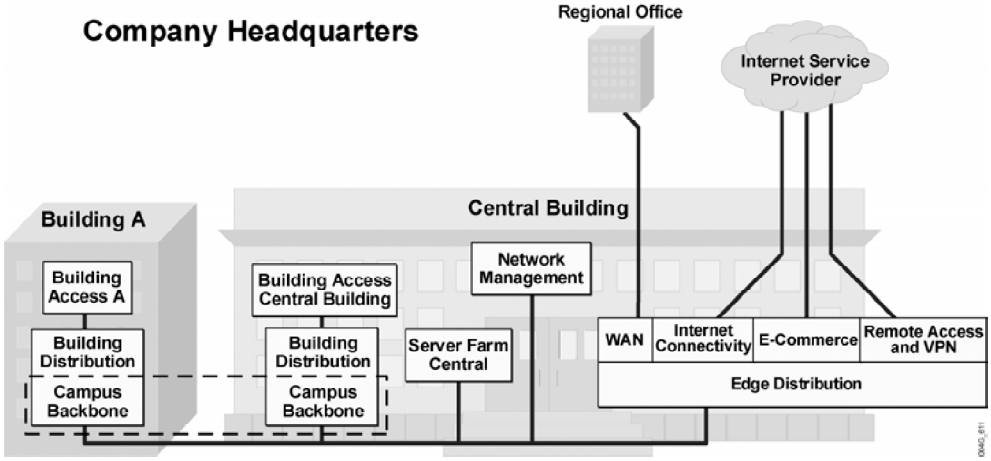

Regional Site Network

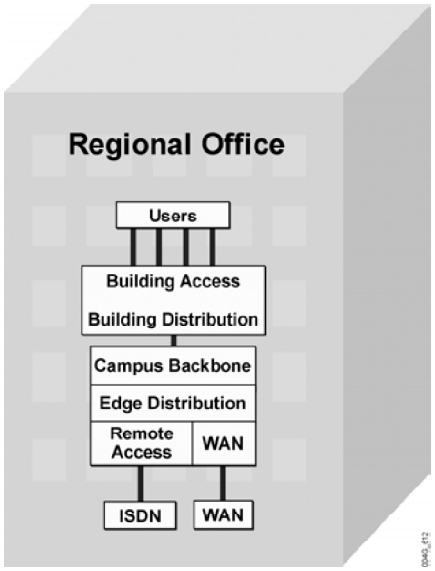

Remote Office Network

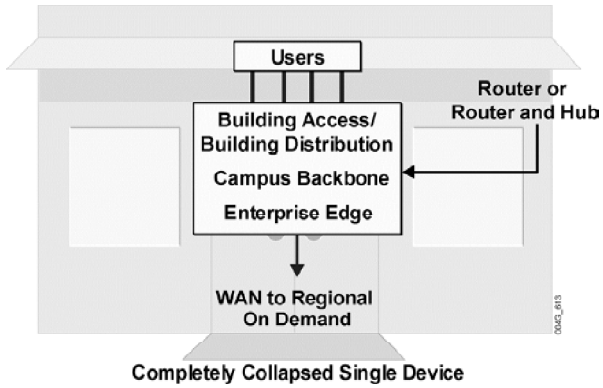

Designing Basic Campus Switched Networks

Case Study 3: Enterprise Campus Design

The DJMP Industries existing campus consists of two buildings located in San Jose. The customer has already decided to improve the performance of the campus LAN by introducing LAN switching. A simulation tool was used to prove the benefits of a switched versus a shared solution.

Due to requirements for increased reliability and performance, you could propose a redundant switched design with servers placed in a separate LAN (Server Farm). This proposed high-level solution is described in the figure.

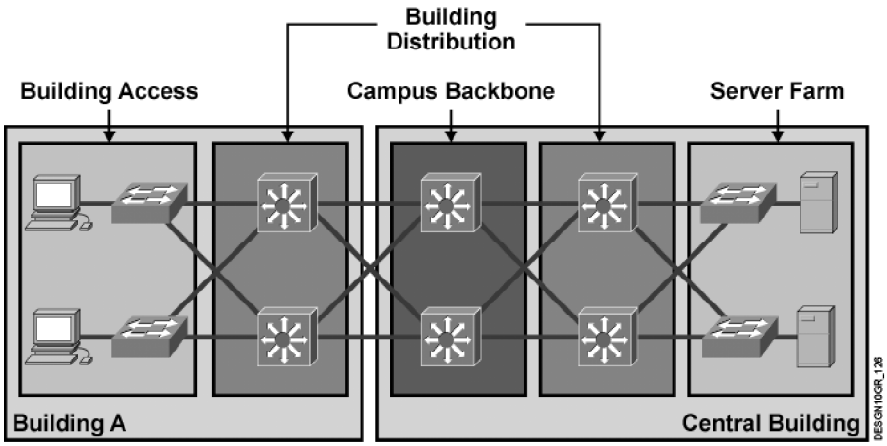

Proposed campus design for DJMP Industries

The proposed Campus Backbone and Building Distribution submodules consist of multilayer switches. All these switches are redundantly interconnected to ensure:

- High availability
- Load sharing capabilities at Layer 3

The Building Access submodule in both buildings consists of data link layer switches redundantly connected to upstream Building Distribution switches.

All the servers are placed in the Server Farm module and are attached to Building Access switches.

Note:	The design currently does not address the high availability of workstations (clients) and servers. With respect to initial requirements, this functionality is not needed at the moment. However, if there is a future requirement for the high availability of the servers, the company can easily implement dual-attached transceivers attached to two Building Access switches.

To increase high availability at the Building Access submodule, you could configure the Building Distribution switches configured with Cisco Hot Standby Router Protocol (HSRP).

To achieve load sharing between the Building Access switches and the Building Distribution switches, you must configure the VLANs on the Building Access switches and tune the

Spanning Tree Protocol (STP) per VLAN to utilize both upstream links to Building Distribution switches.

Simulation 3-1: Shared versus Switched LAN

Q1) What can you observe from the graphs?

The HTTP response times are very low and consistent due to the fact that the network is not loaded.

Q2) What can you determine from the results? What is the reason for the delayed HTTP responses?

The overall Ethernet utilization substantially increased compared to the unloaded network. The HTTP response averages are within the expected values. Occasionally, there are significant delays that decrease the probability of prompt responses. However, the majority of the response times are below 150 ms. Delays may result from collisions on the significantly loaded Ethernet and subsequent retransmissions of the frames carrying the HTTP requests.

Q3) You came to a conclusion that the introduction of the data link layer switch represents a significant improvement for a given case. How is that determined from the graphs?

The graphs are comparable to the graphs that resulted from an unloaded network. There are no significant delays in the HTTP responses and the probability of prompt responses, below 20 ms, is very high.

Simulation 3-2: Data link layer vs. Multilayer Switching

Q1) Will the traffic immediately start using the original path once the link or node has fully recovered?

No. The STP needs some time to recalculate the graph of best paths.

Q2) Examining the resulting graph, you may notice that there is no load sharing in the Building Access submodule of building A. Is this due to the default routing on the workstations using Building Distribution switch DS_A for the primary exit point, or because of the attached data link layer switch placing the secondary port in the blocking mode?

The reason for not using a secondary multilayer switch DS_B is because of the way in which default routing is done. There is no need for the STP when the switch is attached to the routed ports.

Q3) Why is the return path completely bypassing the CS_A switch?

The DC_A multilayer switch, which is used as a default gateway for the attached LAN, knows that the CS_A switch is down.

Q4) What is the load distribution ratio on DS_A – CS_A versus DS_A – CS_B link? Explain.

The ratio is 1:1. The load is evenly shared among symmetric equal-cost paths all the way to destination.

Q5) The workstation WS_B is not running any routing protocol; rather, it depends on the default routing. What is a proper next-hop address?

The interface network address of the attached AS_F1 multilayer switch.

Q6) Running a routing protocol is one way to make the server forward packets to both Building Distribution switches. Can you think of any other option?

The other option is to use two equal-cost default routes, pointing to the interface network addresses of both Building Distribution switches. Note that this will result in a serious disruption to the packet forwarding process if any interfaces or nodes fail.

Designing WAN Networks

Case Study 4: WAN Upgrade and Backup

The current situation with the DJMP Industries WAN network dictates an upgrade of the most loaded links. According to simulations, the expected load of newly introduced applications will result in the total congestion of the San Jose to Houston link, in both directions, and a heavy load on the other links.

During the initial redesign phase, the designers determined that the current WAN bandwidths on links to all regional offices were 64 Kbits. The links are leased lines. The simulation of new applications showed that the immediate upgrade was needed on the San Jose – Houston link and that 128 Kbps would be sufficient. Due to expected heavy load on other links as well, it is recommended to upgrade the entire WAN to 128 Kbits.

Taking into account that existing routers support higher speeds, you could consider the leased lines versus Frame Relay WAN transport as well. The flexibility of Frame Relay would ensure a gradual increase of bandwidth, in terms of Committed Information Rate (CIR) in the future as demand grows.

Note: You should perform a cost analysis before deciding whether to recommend a leased line upgrade or a Frame Relay transport. The existing equipment supports both leased lines and Frame Relay, so cost and flexibility are the decisive factors.

With respect to WAN backup, the ISDN is the most suitable option. The existing WAN bandwidths are low, and the 64 Kbits of the ISDN B channel is ideal in the current situation. The installed equipment on regional offices already supports ISDN for remote offices connectivity. One of the available ISDN B channels or additional Basic Rate Interface (BRI) would be sufficient. The backup will always be triggered at regional offices.

The central router at the Headquarters will be upgraded with an additional slot with multiple BRI interfaces, even though one BRI interface would be statistically sufficient for the three regional offices.

For the international offices in London and Singapore that will soon open, a VPN implementation is proposed. The existing DJMP Industries is already connected to the Internet at San Jose, and the new offices will also be connected to the Internet. Due to an insecure medium such as the Internet, the VPN implementation must ensure the security. An IPsec implementation is suggested.

In case the international offices require redundancy, the Generic Route Encapsulation (GRE) tunnels are an option with IPsec on the tunnels. The latter option requires an additional address block to be reserved, and routing to be turned on in the GRE tunnels.

The new WAN with the backup option, and the integration of the international offices are illustrated in the figure.

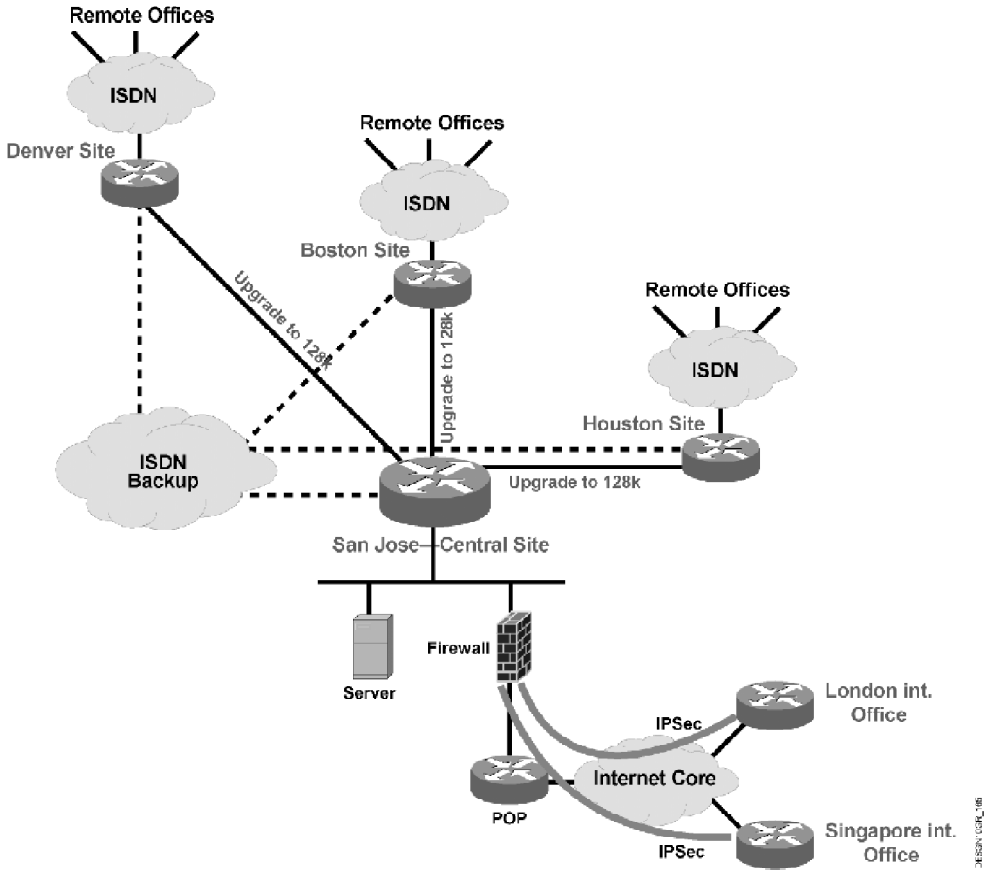

Proposed WAN network of the DJMP Industries

Designing IP Addressing for the Network

Case Study 5: Network Addressing Plan

The new IP addressing plan of the DJMP Industries network assumes the use of 10.0.0.0/8 private IP addresses. The address space is hierarchically designed and allows advanced routing design, including route summarization. Each location will be given a /16 address block from the 10.0.0.0/8 address pool. Each address block is further subnetted to address certain network segments and offers enough addresses for possible future growth. The proposed IP address space assignment by locations is listed in Table 1: IP Address Allocation.

IP Address Allocation

Location	IP Address Space
Headquarters in San Jose	10.1.0.0/16
Denver Regional Office	10.2.0.0/16
Boston Regional Office	10.3.0.0/16
Houston Regional Office	10.4.0.0/16
London International Office	10.5.0.0/16
Singapore International Office	10.6.0.0/16

The assigned IP address space on each location is divided into two blocks with /17 network masks. The first block is intended for LAN network(s). The second block is intended for various WAN connections, primary links as well as backup connections. Both IP address blocks provide enough IP address space for future growth. The point-to-point connections will have /30 network masks for optimal IP address assignment. The LAN networks will have /24 network masks to leave enough IP addresses for possible new users.

San Jose Headquarters

The proposed IP addressing allocation for the San Jose headquarters (10.1.0.0/16) is divided into 10.1.0.0/17 address block for LAN networks and 10.1.128.0/17 address block for WAN connections. There is currently one LAN network and three WAN connections to regional offices in Denver, Boston, and Houston. There will also be a backup connection for each WAN connection.

Note: For the international offices, a VPN implementation is planned and address space is also reserved for the tunnels, in case GRE tunnels are used. The details of VPN implementation are explained in the WAN Upgrade and Backup section, presented as a case study in the module “Designing IP Addressing for the Network” of the DESGN course.

San Jose Headquarters IP Address Allocation

Connection Type	Network
San Jose Headquarters LAN network	10.1.0.0/24
Headquarters to Denver WAN connection	10.1.128.0/30
Headquarters to Boston WAN connection	10.1.128.4/30
Headquarters to Houston WAN connection	10.1.128.8/30
Headquarters to London VPN tunnel (optional)	10.1.128.12/30
Headquarters to Singapore VPN tunnel (optional)	10.1.128.16/30
Headquarters to Denver backup connection	10.1.129.0/30
Headquarters to Boston backup connection	10.1.129.4/30
Headquarters to Houston backup connection	10.1.129.8/30

The proposed IP addressing plan for LAN networks allows the introduction of VLAN networks. Table 3: San Jose Headquarters VLAN lists the possible IP addressing scheme for VLANs at the Headquarters, including the planned Server Farm LAN.

San Jose Headquarters VLAN

VLAN or LAN Name	Network
Server Farm (LAN)	10.1.0.0/24
E-Commerce	10.1.2.0/24
Administration	10.1.3.0/24
The Next VLAN	10.1.4.0/24

Regional/International Offices

The proposed IP addressing allocation for the Denver, Boston and Houston Regional Office areas, including the remote offices, and the London and Singapore International Offices, is described in Table 4: Regional/International Offices IP Address Allocation. As the proposed IP address assignment is hierarchical, it enables route summarization towards San Jose headquarters, resulting in smaller routing tables.

Regional/International Offices IP Address Allocation

Connection Type	Network
Denver Regional Office LAN network	10.2.0.0/24
Denver to Remote Office 1 connection	10.2.128.0/30
Denver to Remote Office 2 connection	10.2.128.4/30
Denver to Remote Office 3 connection	10.2.128.8/30
Remote Office 1 LAN network	10.2.129.0/24
Remote Office 2 LAN network	10.2.130.0/24
Remote Office 3 LAN network	10.2.131.0/24
Boston Regional Office LAN network	10.3.0.0/24
Boston to Remote Office 1 connection	10.3.128.0/30
Boston to Remote Office 2 connection	10.3.128.4/30
Remote Office 1 LAN network	10.3.129.0/24
Remote Office 2 LAN network	10.3.130.0/24
Houston Regional Office LAN network	10.4.0.0/24
Houston to Remote Office 1 connection	10.4.128.0/30
Houston to Remote Office 2 connection	10.4.128.4/30
Houston to Remote Office 3 connection	10.4.128.8/30
Remote Office 1 LAN network	10.4.129.0/24
Remote Office 2 LAN network	10.4.130.0/24
Remote Office 3 LAN network	10.4.131.0/24
London International Office LAN network	10.5.0.0/24
Singapore International Office LAN network	10.6.0.0/24

IP Address Assignment Methods

The IP address assignment methods suitable for DJMP Industries network are either DHCP or manual configuration. The details are as follows:

- DHCP for the San Jose headquarters location. There are already 200 users with future growth expected, so the manual configuration is not an option.
- DHCP for the regional offices in Denver, Boston and Houston. The current number of users is 35 to 50.
- DHCP for the international offices in London and Singapore. The current number of users is 10 with future growth expected.
- Manual configuration for the remote offices with up to 5 users.

Selecting Routing Protocols for a Network

Case Study 6: Routing Protocol Selection

Based on the scenario, the following tables and text include the proposed solutions for routing deployment in the DJMP Industries network. According to the Case Study Guidelines, there may be some minor variations in your solutions.

Routing Protocol Selection

The new IP routing protocol choice for the DJMP Industries network is OSPF. Reasons for choosing the OSPF as the routing protocol are:

- OSPF is a standardized protocol for routing IP traffic and can therefore be used with multivendor equipment.
- OSPF offers fast convergence. The current routing protocol (RIPv1) is not suitable for growing networks that have to be scalable. The major reason to replace RIP is its slow convergence.
- OSPF supports VLSM. The IP address space is more economically used.
- OSPF offers multiple area design and manual summarization, which reduce the routing table size and overhead.
- OSPF supports dial-up connections with the OSPF Demand Circuit (DC) feature. Dial-up connections are used as backup links in the DJMP Industries network to provide a redundant connection from regional offices.

OSPF Area Design

The underlying IP address plan hierarchy maps to the OSPF routing protocol hierarchical structure. The OSPF Area 0, the backbone area, resides in the San Jose headquarters and includes LAN networks. There are currently three non-backbone areas, which include WAN connections to the regional offices and LAN networks from those locations. All non-backbone areas are OSPF Not-So-Stubby Area (NSSA) areas to allow the route redistribution from non-OSPF routing domains.

OSPF Areas

Location	OSPF Area	Networks in the area
Headquarters in San Jose	Area 0 (backbone area)	10.1.0.0/17
Denver Regional Office	Area 1	10.1.129.0/30, 10.1.128.0/30, 10.2.0.0/24
Boston Regional Office	Area 2	10.1.129.4/30, 10.1.128.4/30, 10.3.0.0/24
Houston Regional Office	Area 3	10.1.129.8/30, 10.1.128.8/30, 10.4.0.0/24

The redistribution is used to inject information about the remote offices LAN networks. The remote offices do not need to run the OSPF routing protocol. Instead, a default route is used on the remote offices to reach the other networks. A static route to reach the remote office LAN network is used on the regional office side. You can manually configure the route, or have the AAA server install it in case PPP encapsulation with authentication is used. The figure

illustrates the area assignments in the network. The backup links, which are dashed, are installed in the same area as the primary links to the regional offices.

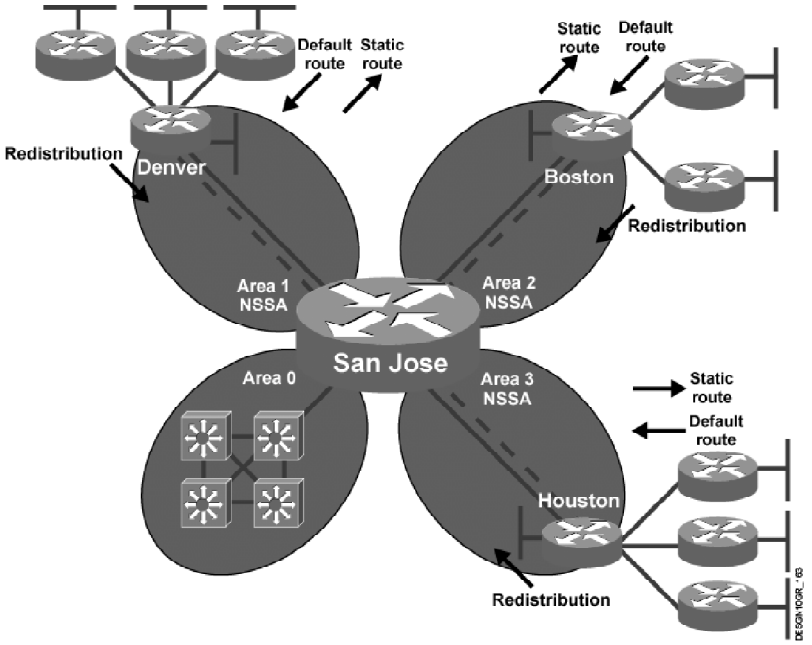

DJMP Industries OSPF Area Design

Route Summarization

The underlying IP addressing plan allows efficient route summarization. The manual summarization of routes is done on area borders. Networks assigned to a certain location are summarized into one summary route, which is then announced to the backbone area. The summarization is as follows:

- 10.2.0.0/16 from Area 1
- 10.3.0.0/16 from Area 2
- 10.4.0.0/16 from Area 3

The external routes (remote office LAN networks) injected into the OSPF area at the regional office routers are also summarized.

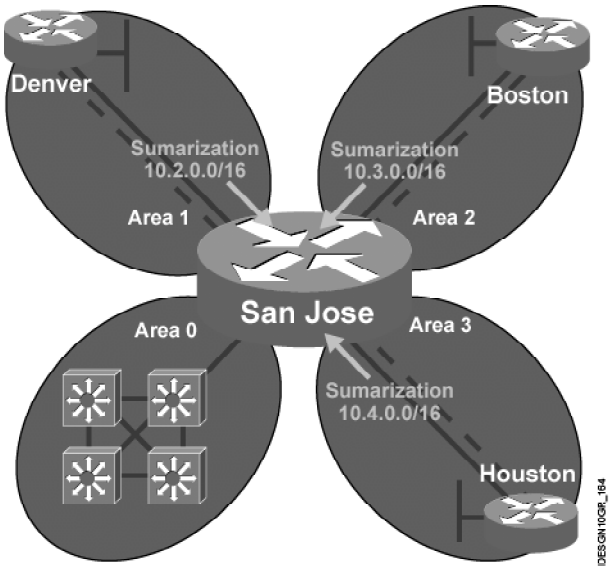

OSPF Route Summarization

The implemented route summarization narrows the scope of routing updates propagation when failure of a link internal to the area occurs.

Simulation 6: Network Convergence

Q1) The outage due to the loss of the primary link can be explained as the consequence of the STP recalculation. What is the reason for the second delay after the primary path has physically recovered?

The STP protocol needs some time to recalculate a graph of best paths based on the new circumstances.

Q2) Why do the link and node incidents impose the same time for network recovery?

In this situation, the STP protocol reacts equally on both disruptions.

Q3) If you chose RIP instead OSPF, the convergence would change significantly under a link or node failure. Why does this occur, and what special procedures do RIP routers undergo?

The network convergence depends heavily on the reacting speed of the routing protocol. Even with flash updates implemented, the distance-vector protocols undergo procedures such as invalid timeout and hold-down timeout, which additionally extend the convergence time.

Q4) From the simulation log it appears that the switchover to the primary path takes too long, approximately 50 seconds. Why?

The OSPF routers are waiting for DR/BDR election.

Evaluating Security Solutions for the Network

Case Study 7: Designing Network Security

Based on the scenario, the following diagrams include the proposed solutions. According to the Case Study Guidelines, there may be some minor variations in your solutions.

Company Headquarters Network

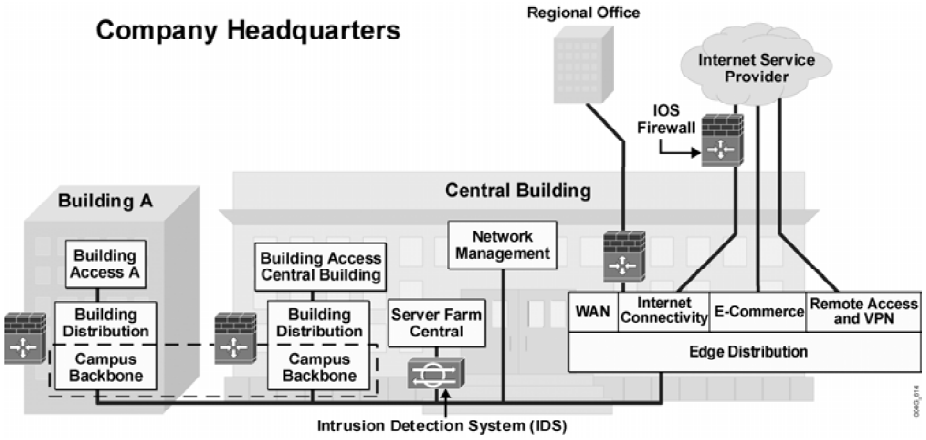

Regional Site Network

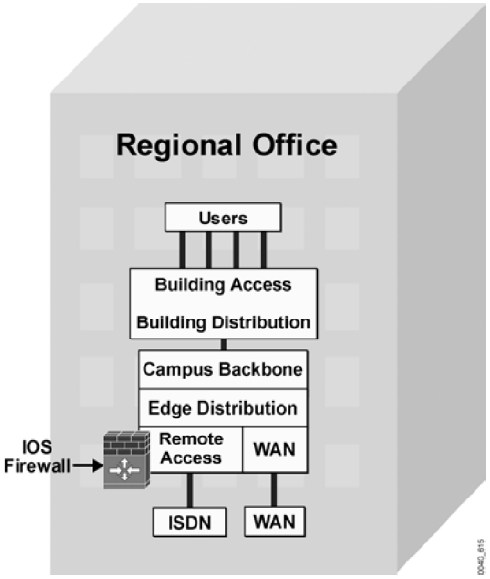

Remote Office Network

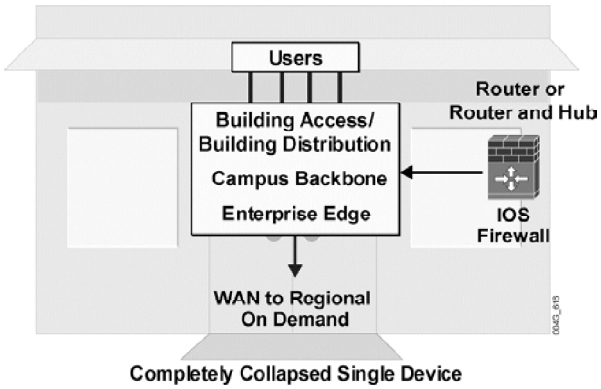

Designing Networks for Voice Transport

Simulation 8: Voice Transport over IP Network

Q1) Why does the jitter disturb the voice session?

It introduces a variation in delay, which voice is very sensitive to.

Q2) How can QoS mechanisms improve the data propagation in congested networks?

QoS mechanisms can even make data traffic more predictable in congested networks.

Applying Basic Network Management Design Concepts

This module does not include a case study exercise.

Final Case Study

MCMB Corporation Network Redesign

The possible reason for the low response times, reported by the users, is the number of broadcasts generated by the stations in the network. These broadcasts are flooded throughout the data link layer switched network since the data link layer switches do not stop broadcasts.

The administrators did not detect any congestion on the backbone switch because of its high bandwidth. The same broadcast cause serious problems on 100 Mbps and 10 Mbps user ports.

The only solution for this issue is to divide the campus into smaller networks and route between them. The Layer 3 functionality, provided by routers or multilayer switches, limits the broadcasts within the new network, and thus relieves end-user stations from processing broadcast requests.

In this design, multilayer switching is proposed in the Building Distribution submodule while the Building Access and the Campus Backbone submodules remain with data-link switching switching. In the Campus Backbone submodule, another data link layer switch is recommended to provide redundancy.

The Building Access submodule, using data link layer switches, supports VLANs and has redundant trunk uplinks to the Building Distribution multilayer switches. All routing between department VLANs and between buildings is performed there. The Campus Backbone submodule does not perform any routing and is dedicated to fast packet switching.

Multilayer switching is also proposed for the Server Farm and Edge Distribution modules. The Server Farm module includes all the internal servers while the Edge Distribution module connects the campus network with the remote locations and the Internet. You should consider the number of Gigabit Ethernet ports between the Campus Backbone submodule and the Server Farm module since most of the traffic flows from the mainframe towards the users and can be a possible bottleneck. In our design, two Gigabit Ethernet uplinks are provided between the Server Farm module and the Campus Backbone submodule.

For equal utilization of all links between switches, you must tune the Spanning Tree protocol (STP) on a per VLAN basis.

The proposed campus redesign is shown in the figure.

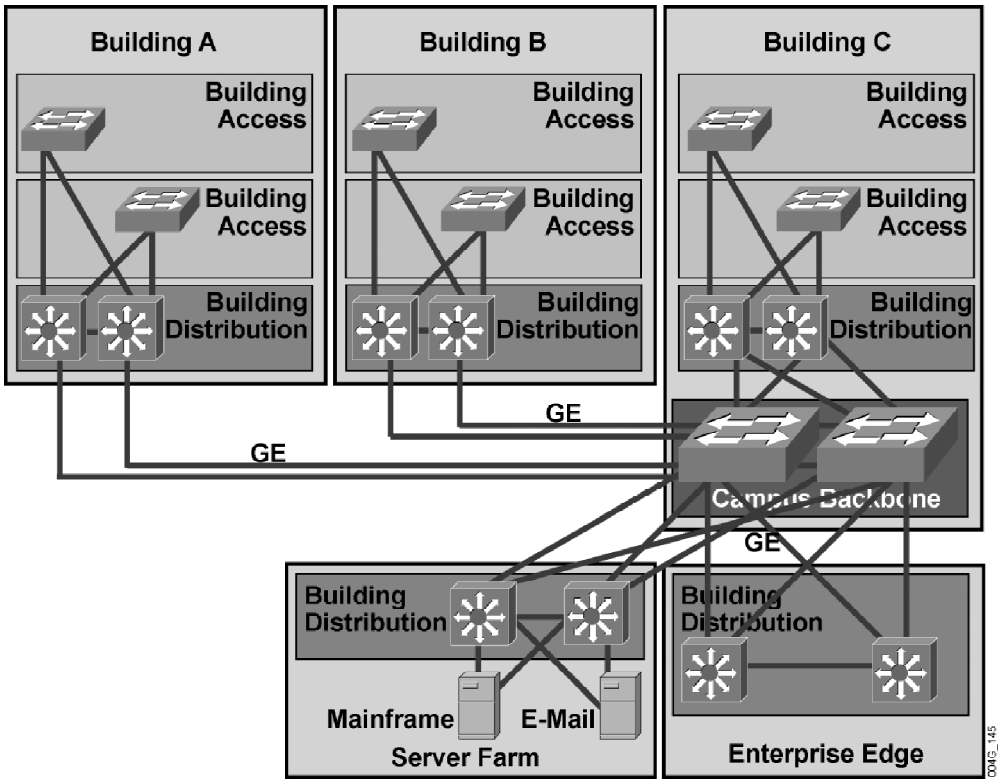

Proposed Campus Redesign

Proposed WAN Backup Design

The proposed WAN backup design assumes that all WAN links should have ISDN backup. Since the headquarters and the international office are interconnected with 2 Mbps links, the ISDN backup link should be of the same size. Therefore at least one ISDN PRI should be allocated in the headquarters and an international office.

A designer should consider the number of ISDN PRI links in the headquarters. If the company requires full WAN redundancy, the number of PRIs should equal to the number of WAN links from the headquarters. If full WAN redundancy is not required, the number of ISDN PRIs can be adequately smaller. Since the distance between international offices and the headquarters are large and the ISDN calls are international, the designer also needs to consider if the backup link needs to have the same bandwidth as the primary WAN link. Usually, the dial backup link has a slightly lower bandwidth, especially on expensive international calls. You can achieve variable bandwidth and the resulting dial-up connection savings by using the bandwidth on demand feature.

The proposed ISDN backup on international-to-remote office connections is based on ISDN BRI with one or two B-channels and the correspodning bandwidth of 64 or 128 kbps. Again, consider either a full or partial backup scenario. In the full backup scenario, the number of available ISDN B-channels in the international office should be equal to or larger than the number of connected remote offices.

Implement all backup links with floating static routes. This feature enables a dynamic use of backup when the destination is not reachable through the primary links. In the absence of dynamic routes, the static route with the higher administrative distance that points over the dial-up link is used. The floating static scenario enables smooth and efficient routing regardless of the WAN connection technique used.

The default route to the dial-up link with the administrative distance 200 is proposed for all upstream connections and remote locations. Of course, the default route needs to be advertised also through the routing protocol to the downstream routers and in typical situations, only the WAN link is utilized.

The proposed backup scenario is shown in the figure.

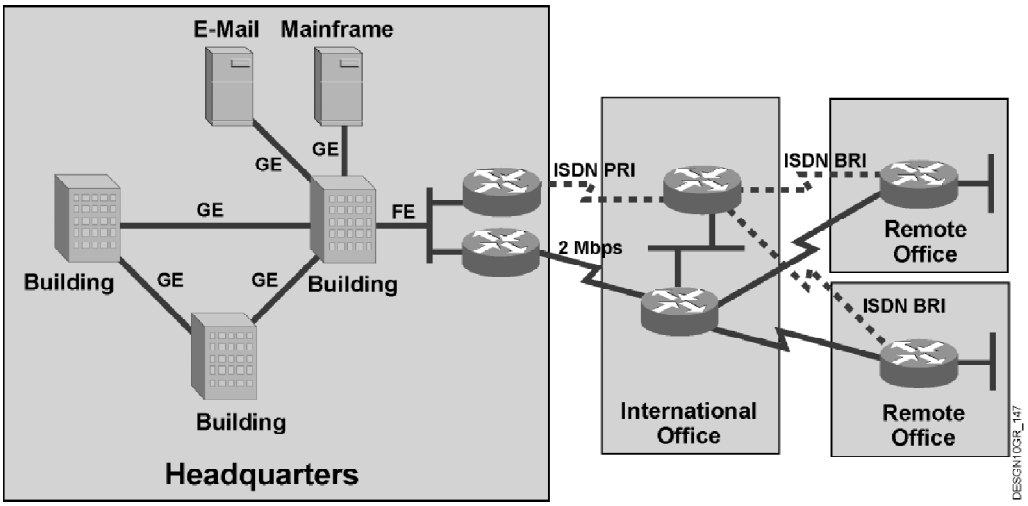

Proposed WAN backup design

Proposed IP Addressing Plan

The proposed IP addressing plan assumes the use of 10.0.0.0/8 private IP addresses. The address space is hierarchically designed, which allows advanced routing design including route summarization.

In the headquarters the /9 address block is assigned. This block is further divided into /12 blocks for each department and /16 for each building-department combination. For each country (international office with multiple remote offices), a /16 address block is used. This block is then divided into /24 blocks for international office and remote office LANs. All WAN links from the international office use a /32 address from one of the international /24 blocks. Furthermore, the ISDN backup links use very similar IP addresses from another /24 address block.

The table lists the proposed IP address space assignment by typical locations.

IP Address Allocation

Location	IP Address Space	
Headquarters	10.0.0.0/9	
Headquarters—development VLAN	10.0.0.0/12	
Headquarters—development VLAN—building 1	10.10.0/16	
Headquarters—development VLAN—building 2	10.2.0.0/16	
Headquarters—development VLAN—building 3	10.3.0.0/16	
Headquarters—technical support VLAN	10.16.0.0/12	
Headquarters—technical support VLAN—building 1	10.16.0.0/016	
Headquarters—technical support VLAN—building 2	10.17.0.0/16	
Headquarters—technical support VLAN—building 2	10.18.0/16	
Headquarters—marketing VLAN	10.32.0.0/12	
Headquarters—marketing VLAN—building 1	10.32.0.0/16	
Headquarters—marketing VLAN—building 2	10.33.0.0/16	
Headquarters—marketing VLAN—building 3	10.34.0.0/16	
Headquarters—sales VLAN	10.48.0.0/12	
Headquarters—sales VLAN—building 1	10.48.0.0/16	
Headquarters—sales VLAN—building 2	10.49.0.0/16	
Headquarters—sales VLAN—building 3	10.50.0.0/16	
Headquarters—server farm	10.64.0.0/12	
Headquarters—edge distribution	10.80.0.0/12	
International network	10.X.0.0/16	x > 127
International Office LAN network	10.X.0.0/24	x > 127
Remote Office 1 LAN network	10.X.Y1.0/24	x > 127, Y1 > 0
Remote Office 2 LAN network	10.X.Y2.0/24	x > 127, Y2 > 0 <> Y1
International Office to Headquarters WAN connection	10.X.Z.0/30	x > 127, Z <> Y
International Office to Remote Office 1 WAN connection	10.X.Z.4/30	x > 127, Z <> Y
International Office to Remote Office 2 WAN connection	10.X.Z.8/30	x > 127, Z <> Y
International Office to Headquarters ISDN backup	10.X.Z1.0/30	x > 127, Z1 <> Y <> Z
International Office to Remote Office 1 ISDN backup	10.X.Z1.4/30	x > 127, Z1 <> Y <> Z
International Office to Remote Office 2 ISDN backup	10.X.Z1.8/30	x > 127, Z1 <> Y <> Z

Note: When implementing VLANs for each department within the Campus Backbone and between buildings, the deployment of different security policies becomes very simple. All packets between different VLANs need to go through the Building Distribution switch where the security policies are implemented.

The proposed IP address allocation is shown in the figure.

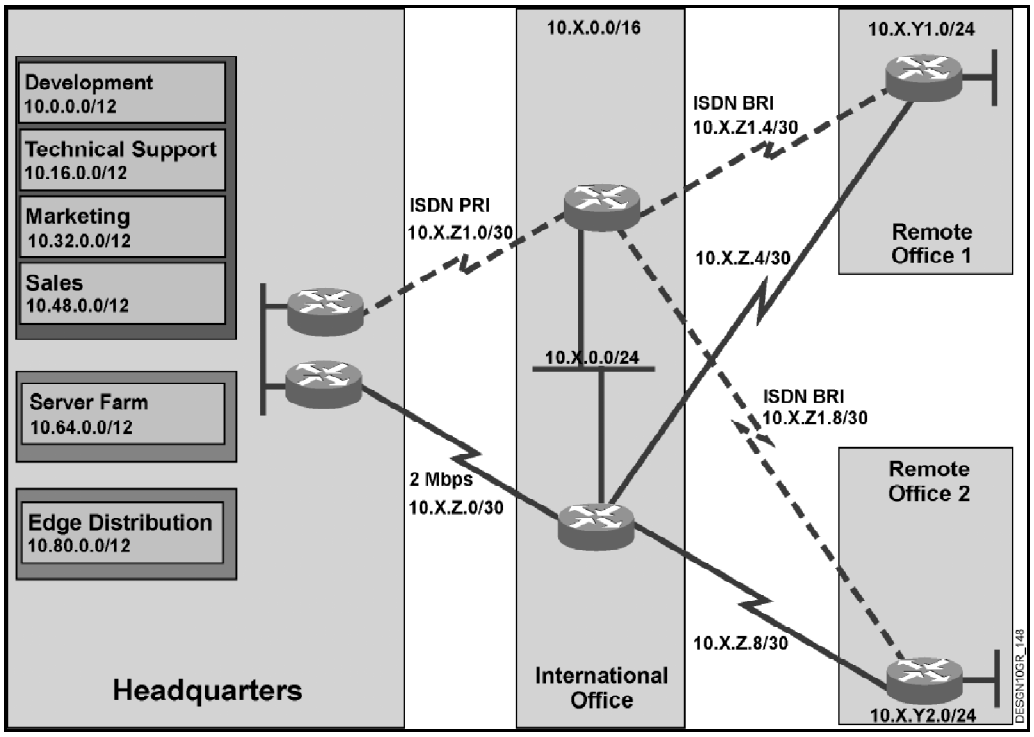

Proposed IP Address Allocation

The proposed IP address assignment is hierarchical and enables route summarization. The goal of route summarization is to have small and easily manageable routes in the routing table and to relieve small routers in the international offices from calculating large number of routes that can be a fairly CPU intensive task.

International routers summarize the network 10.X.0.0/16 towards the headquarters as well as downstream to the remote offices. The headquarters router summarizes the network 10.0.0.0/8 towards all international offices. However, if the default route is required (to access the Internet when there is no proxy in the firewall), the 0.0.0.0/0 network needs to be advertised through the routing protocol as well.

Note:	The following rule applies: the floating static routes between remote locations that point to the dial-up interfaces should not be more specific than the received summarized addresses. Otherwise, most of the traffic will flow through the dial-up links.

The proposed route summarization is shown in the figure.

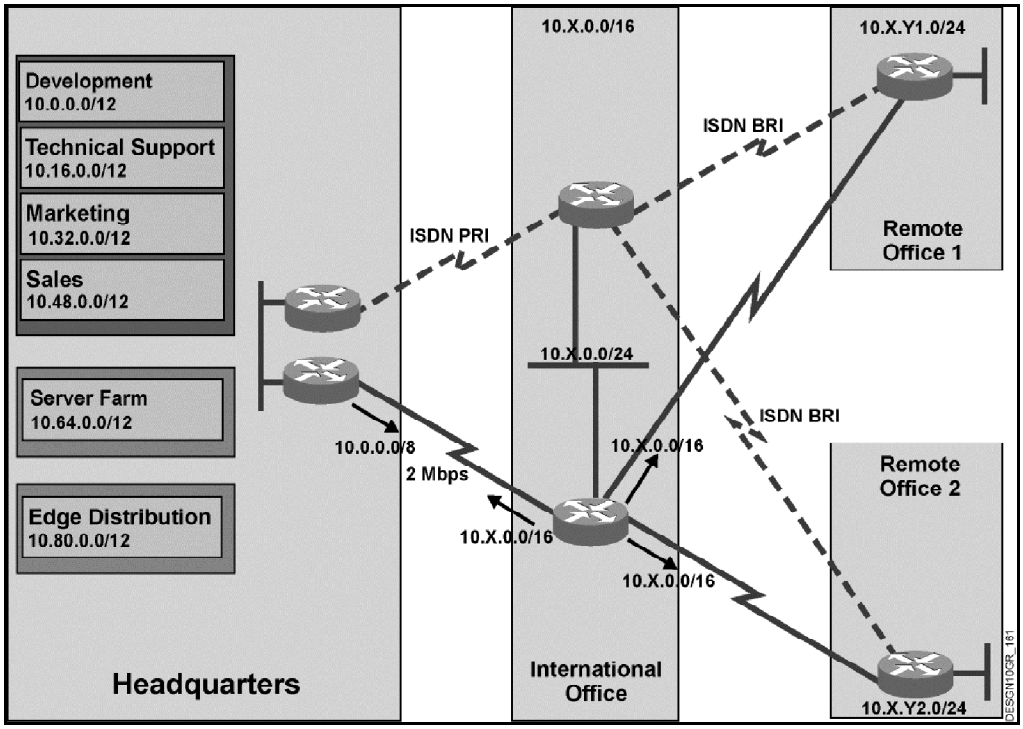

Proposed IP Address Summarization

Proposed Routing Protocol

The existing routing protocol, Enhanced Interior Gateway Routing Protocol (EIGRP), supports fast convergence and does not show any major disadvantages against other routing protocols. However, EIGRP will be kept in this design and summarization will be implemented, which reduces routing tables and simplifies the network.

Proposed Extranet Design

The Extranet design is based on the assumption that partners and customers may choose any connectivity option they want; the only requirement is that the connection must be secure due to the nature of the transferred data. The natural choice in this case is secure VPN with IP Security (IPSec) and 3DES that provides data confidentiality and integrity. In addition, every user accessing the application through the VPN tunnel needs to be authenticated. In this design, the user authentication is provided by using token-cards with one-time passwords.

The VPN tunnels are terminated on the PIX firewall in the E-commerce module where the Extranet servers are also located. The E-commerce module is connected to the campus network through the Edge Distribution module.

Due to the large number of connectivity options (ISDN, ADSL, leased line, and so on), the designers decided to outsource the access to a service provider and provide the access to the E-commerce module through the Internet. This solution requires the PIX firewall with the adequate performances and the Internet connection with enough bandwidth. According to the Erlang B table (http://www.erlang.com/calculator/erlb/) and the minimum blocking probability P01, the 2 Mbps link should be sufficient.

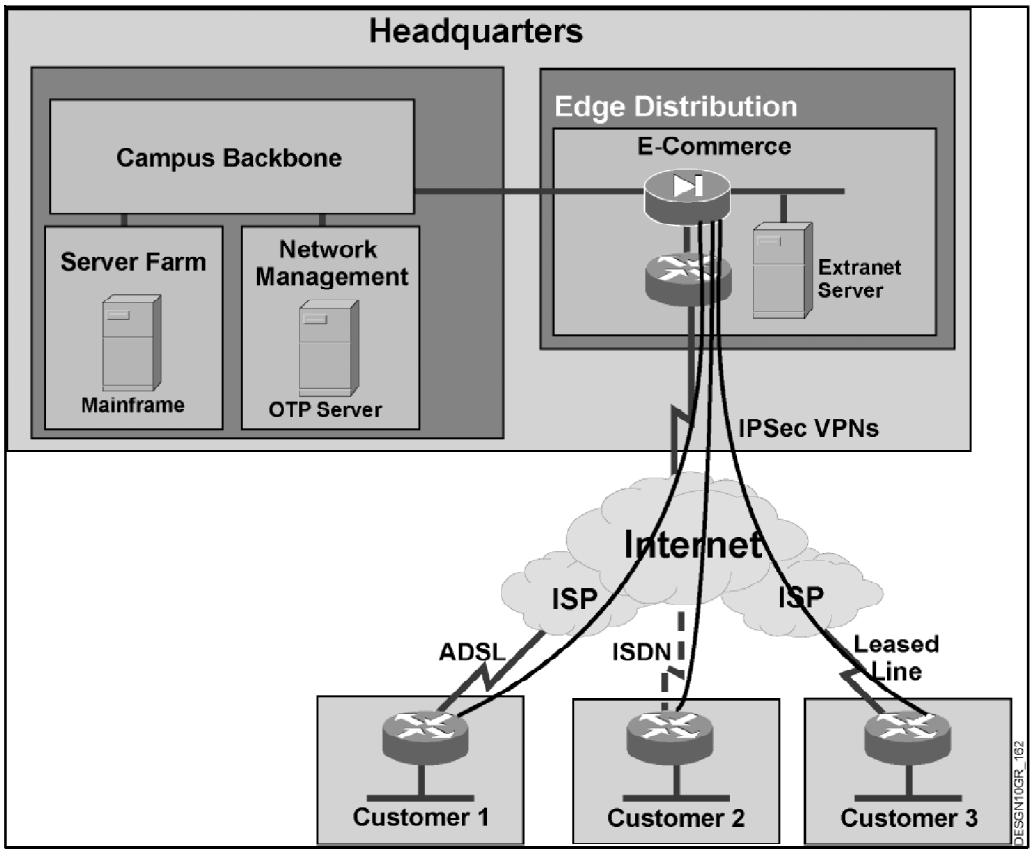

Proposed Extranet Design

Job Aids

The job aids described in this course are contained here. You can copy the job aids as needed to complete your network design tasks.

Planned Applications Worksheet

Use the table as a template to identify and evaluate planned applications in the course case study and for future design efforts. Remember to include your applications in the Application Type column.

Application Type	Application	Criticality	Comments
E-mail			
Groupware			
Voice networking			
Web browsing			
Video on demand			
Database			
Customer support applications			

Planned Intelligent Network Services Worksheet

Use the table as a template to identify and evaluate planned intelligent network services in the course case study and for future design efforts. Remember to include your services in the Service column.

Service	Comments
Security	
QoS	
Network management	
High availability	
IP multicast	

Organizational Goals Worksheet

Use the table as a template to identify organizational goals in the course case study and for future design efforts.

Goal	Data	Comments

Organizational Constraints Worksheet

Use the table as a template to help determine organizational constraints in the course case study and for future design efforts. Remember to include your constraints in the Constraint column.

Constraint	Data	Comments
Budget		
Personnel		
Policy		
Scheduling		

Technical Goals Worksheet

Use the table as a template to identify and evaluate technical goals in the course case study and for future design efforts. Remember to include your goals in the Technical Goals column.

Technical Goals	Importance	Comments
Performance		
Availability		
Manageability		
Security		
Adaptability		
Scalability		

Technical Constraints Worksheet

Use the table as a template to identify and evaluate technical constraints in the course case study and for future design efforts.

Technical Constraints	Gathered Data	Comments
Existing equipment		
Bandwidth availability		
Application compatibility		

Decision Table Worksheet

The table is provided as a template. Fill in the parameters on each row, the available options for each column, and make your own decision table.

Parameter						Required Network Parameters

Design Scope Worksheet

Use the table as a template to identify the scope of the network design in the course case study and for future design efforts.

Scope of Design	Comments

Network Locations Worksheet

Use the table to record your own network information.

Location	Type	Comments

Overall Network Size Worksheet

Use the table to record network size information.

Device Type	Number	Comments
SUM		

Network Locations Size Worksheet

Use the table to record the size of each network location.

Location	Office Type	Work-stations	Servers	IP Phones	Router Interfaces	Switches	Firewall Interfaces	Reserve	SUM
SUM									

Protocol Selection Job Aid

You can use the table as a job aid during your assignments. Two additional rows are included so you can specify more parameters (for example, types of existing physical topologies), which might be of importance in your network.

Parameters (Options)	RIP v2	IGRP	EIGRP	OSPF	IS-IS	Required Network Parameters
Size of Network (Small-Medium-Large-Very Large)	Medium	Medium	Large	Large	Very Large	
Speed of Convergence (Very High-High-Low)	Medium	Low	Very High	High	High	
Use of VLSM (Yes-No)	Yes	No	Yes	Yes	Yes	
Mixed Vendor Devices (Yes-No)	Yes	No	No	Yes	Yes	
Network Support Staff Knowledge (Good-Poor)						

