ARCH

Designing Cisco Network Service Architectures (ARCH) v1.1

Student Guide

Version 1.1

Course Introduction

Overview

Given enterprise business and technical requirements and constraints, you will learn how to perform the conceptual and intermediate design of a network infrastructure that supports desired network solutions over intelligent network services, to achieve effective performance, scalability, and availability.

You will learn the fundamental aspects of campus and edge network design, network management, high availability, security, quality of service (QoS), and IP multicast.

In addition, you will be able to design solutions for the network that are strategic to small, medium, and large enterprises, including virtual private networking, wireless, IP telephony, content networking, and storage networking.

Outline

The Course Introduction includes these topics:

- Overview
- Course Objectives
- This Course and the Design Process
- Cisco's Certification Track
- Learner Skills and Knowledge
- Learner Responsibilities
- General Administration
- Course Roadmap
- Icons and Symbols
- Learner Introductions
- Case Study and Simulations
- Course Evaluations

Course Objectives

This topic lists the course objectives.

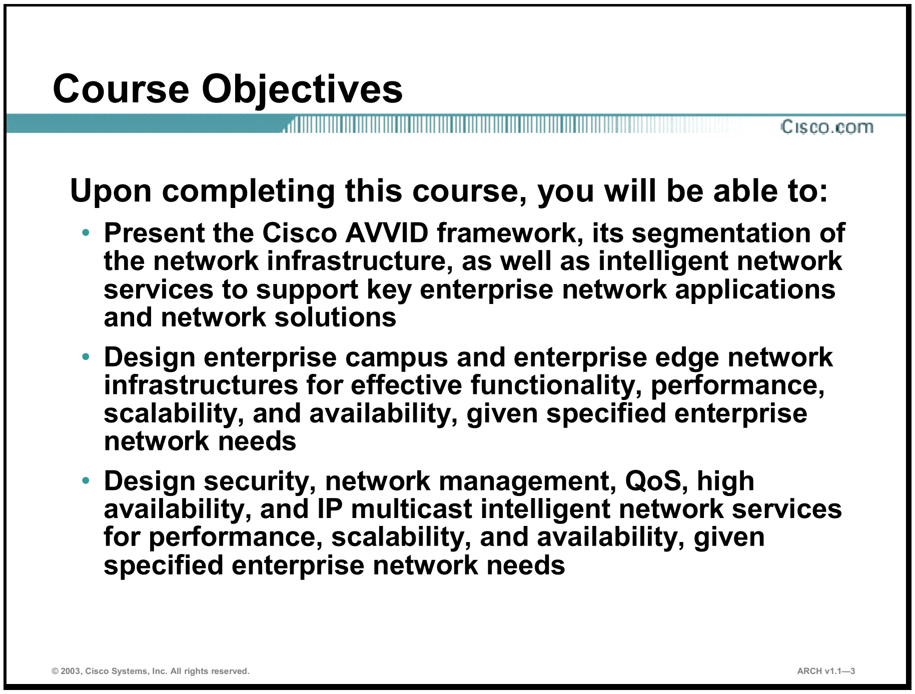

Course Objectives

Upon completing this course, you will be able to:

- Present the Cisco AVVID framework, its segmentation of the network infrastructure, as well as intelligent network services to support key enterprise network applications and network solutions
- Design enterprise campus and enterprise edge network infrastructures for effective functionality, performance, scalability, and availability, given specified enterprise network needs
- Design security, network management, QoS, high availability, and IP multicast intelligent network services for performance, scalability, and availability, given specified enterprise network needs

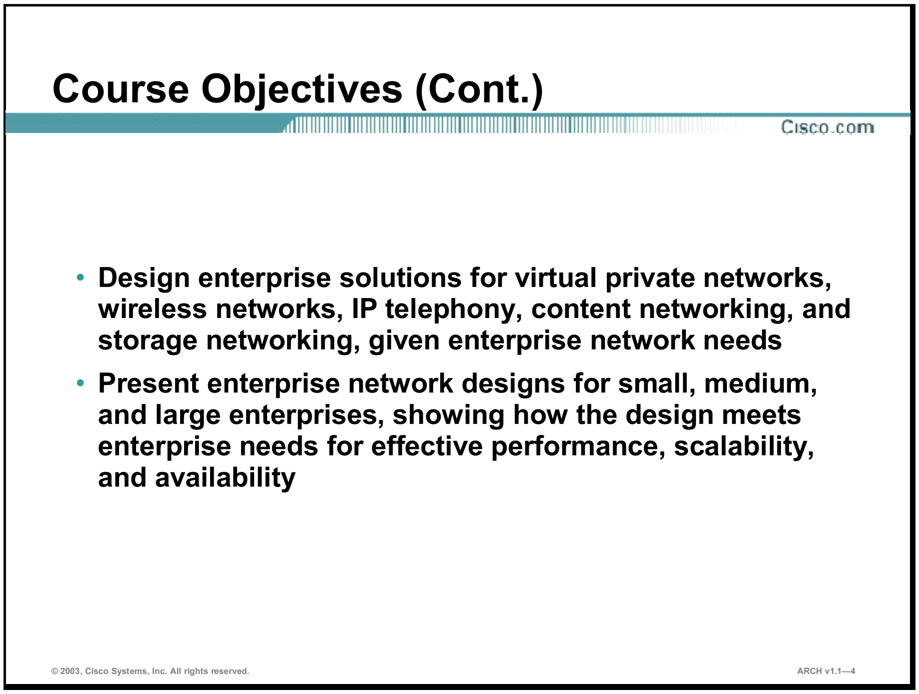

Course Objectives (Cont.)

- Design enterprise solutions for virtual private networks, wireless networks, IP telephony, content networking, and storage networking, given enterprise network needs
- Present enterprise network designs for small, medium, and large enterprises, showing how the design meets enterprise needs for effective performance, scalability, and availability

Upon completing this course, you will be able to:

- Present the Cisco Architecture for Voice, Video and Integrated Data (AVVID) framework, its segmentation of the network infrastructure, as well as intelligent network services to support key enterprise network applications and network solutions
- Design enterprise campus and enterprise edge network infrastructures for effective functionality, performance, scalability, and availability, given specified enterprise network needs
- Design security, network management, QoS, high availability, and IP multicast intelligent network services for performance, scalability, and availability, given specified enterprise network needs
- Design enterprise solutions for virtual private networks, wireless networks, IP telephony, content networking, and storage networking, given enterprise network needs
- Present enterprise network designs for small, medium, and large enterprises, showing how the design meets enterprise needs for effective performance, scalability, and availability

This Course and the Design Process

This topic describes the role of this course as it relates to developing a network design.

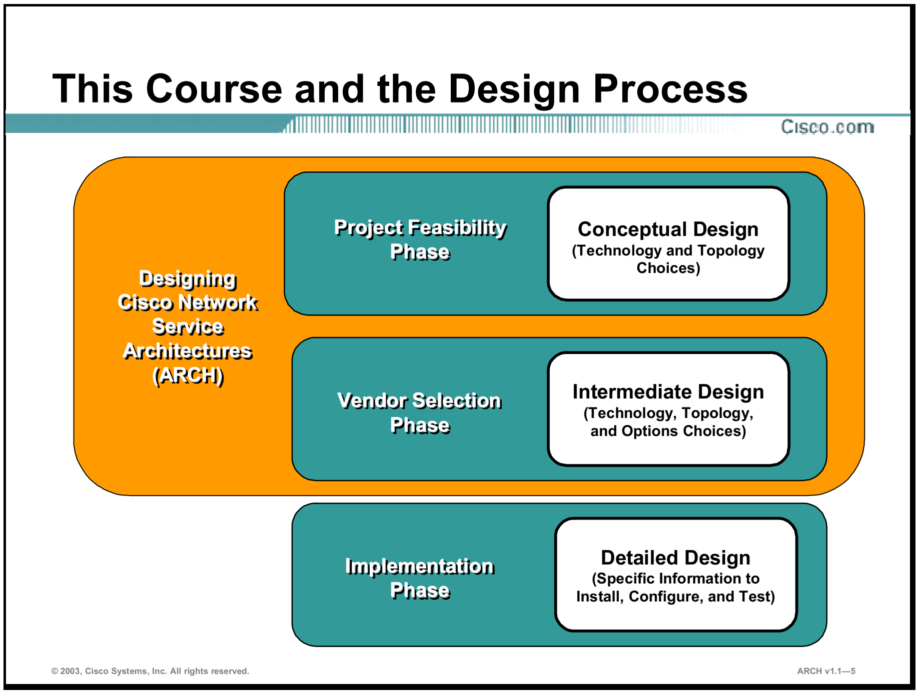

This Course and the Design Process

Designing Cisco Network Service Architectures (ARCH)

Project Feasibility Phase

Conceptual Design
(Technology and Topology Choices)

Vendor Selection Phase

Intermediate Design
(Technology, Topology, and Options Choices)

Implementation Phase

Detailed Design
(Specific Information to Install, Configure, and Test)

ARCH v1.1—5

A conceptual design is an outline description of a network solution, or a selection of network solutions, that an enterprise would plan to eventually engineer, implement, and test. The conceptual design is generally undertaken to establish the most suitable network architecture (a combination of the technology and topology) from a number of options.

An intermediate design is a description of a network solution, or a choice of network solutions, that an enterprise is proposing during the vendor selection phase of a project. The intermediate design is often developed through discussions with enterprise users, information technology personnel, and vendors. It may include costs.

A detailed design is a description of a network solution, which is to be implemented directly. It represents the final stage of diagrams and documentation before physical installation and configuration of the network solution. The detailed design describes everything necessary for an enterprise to order, stage, install, and physically configure network equipment.

This course, *Designing Cisco Network Service Architectures*, focuses on the conceptual and intermediate design phases of a project.

Cisco’s Certification Track

This topic lists the certification requirements of this course.

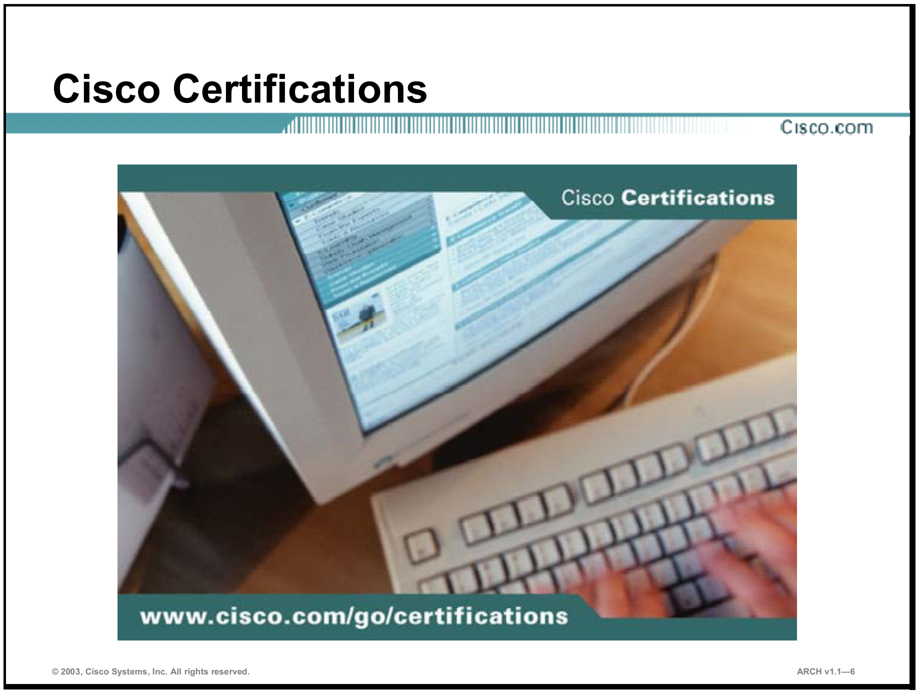

Cisco Certifications

Cisco.com

Cisco **Certifications**

www.cisco.com/go/certifications

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6

This education offering is a Cisco certified professional-level course. This course (*Designing Cisco Network Service Architectures*) is the recommended method of preparation for the Cisco CCDP® exam. The CCDP certification indicates a professional mastery of network design.

This course enables learners, applying solid Cisco network solution models and best design practices, to provide viable, stable enterprise internetworking solutions. The course presents concepts and examples necessary to design enterprise campus and edge networks. Advanced network infrastructure technologies such as Virtual Private Networks (VPNs) and wireless communications are also covered.

The course covers issues and considerations for fundamental intelligent network services including security, network management, QoS, high availability, and bandwidth use optimization through IP multicasting, as well as design models for network solutions such as voice networking and content and storage networking.

The CCDP exam is the final step necessary to achieve the status of Cisco CCDP, following the Cisco CCNA® and CCDA® exams. It affirms possession of some of the skills needed to achieve the status of Cisco CCIE®.

Learner Skills and Knowledge

This topic lists the course prerequisites.

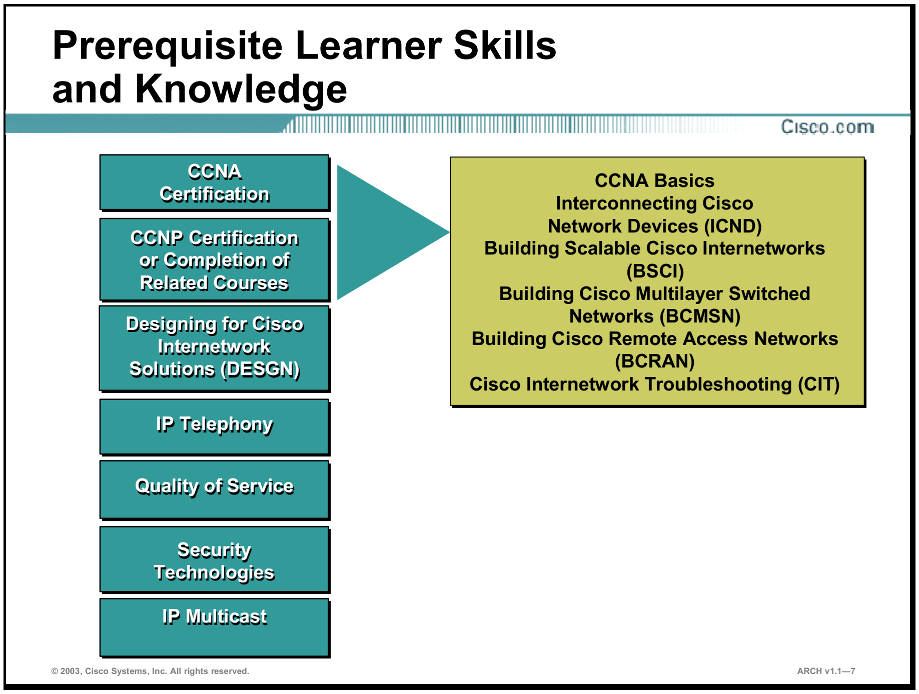

Prerequisite Learner Skills and Knowledge

Cisco.com

CCNA Certification

CCNP Certification or Completion of Related Courses

Designing for Cisco Internetwork Solutions (DESGN)

IP Telephony

Quality of Service

Security Technologies

IP Multicast

CCNA Basics
Interconnecting Cisco Network Devices (ICND)
Building Scalable Cisco Internetworks (BSCI)
Building Cisco Multilayer Switched Networks (BCMSN)
Building Cisco Remote Access Networks (BCRAN)
Cisco Internetwork Troubleshooting (CIT)

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—7

Before taking *Designing Cisco Network Service Architectures*, learners should be familiar with internetworking technologies, Cisco products, and Cisco IOS features. Specifically, before attending this course learners should be able to:

- Design the necessary services to extend IP addresses using variable-length subnet masking (VLSM) and route summarization
- Implement appropriate networking routing protocols, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP) on an existing internetwork
- Redistribute routes between different routing protocols
- Select the required Cisco products and services that enable connectivity and traffic transport for a multilayer campus network
- Select the necessary services at each layer of the network to enable all users to obtain membership in multicast groups in a working enterprise network
- Control network traffic by implementing the necessary admission policy at each layer of the network topology, given a working enterprise network
- Identify the appropriate hardware and software solutions for a given set of WAN technology requirements, including permanent or dial-up access between a central campus, branch offices, and telecommuters
- Select Cisco equipment to establish appropriate WAN connections, given a set of WAN topologies and specifications
- Enable protocols and technologies that allow traffic flow between multiple sites, while minimizing the amount of overhead traffic on each connection

- Implement QoS capabilities to ensure that mission-critical applications receive the required bandwidth within a given WAN topology
- Implement Cisco Voice over IP and IP telephony solutions

Learner Responsibilities

This topic discusses the responsibilities of the learners.

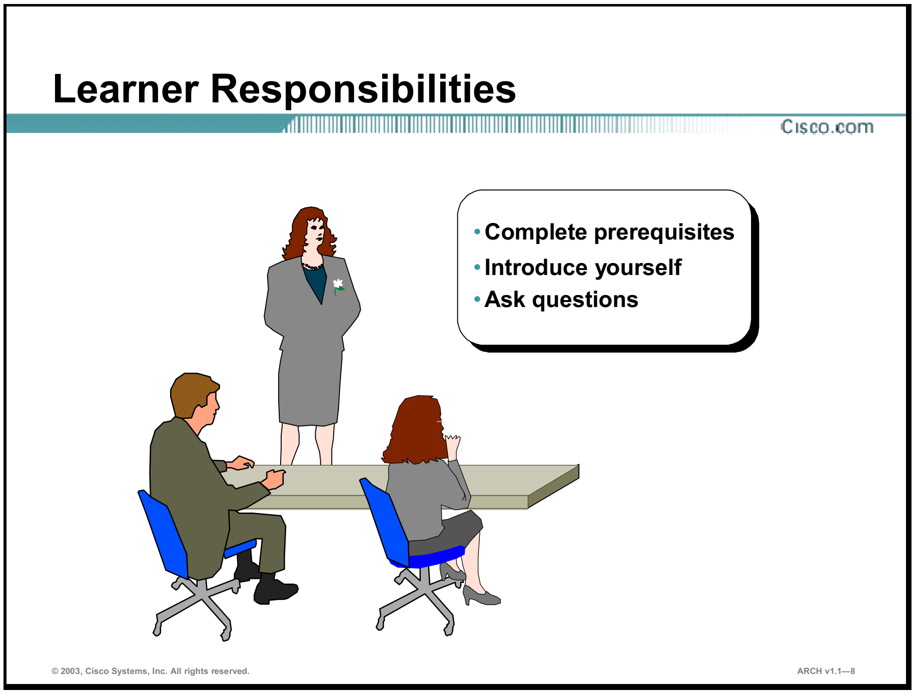

Learner Responsibilities

- Complete prerequisites
- Introduce yourself
- Ask questions

ARCH v1.1—8

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.

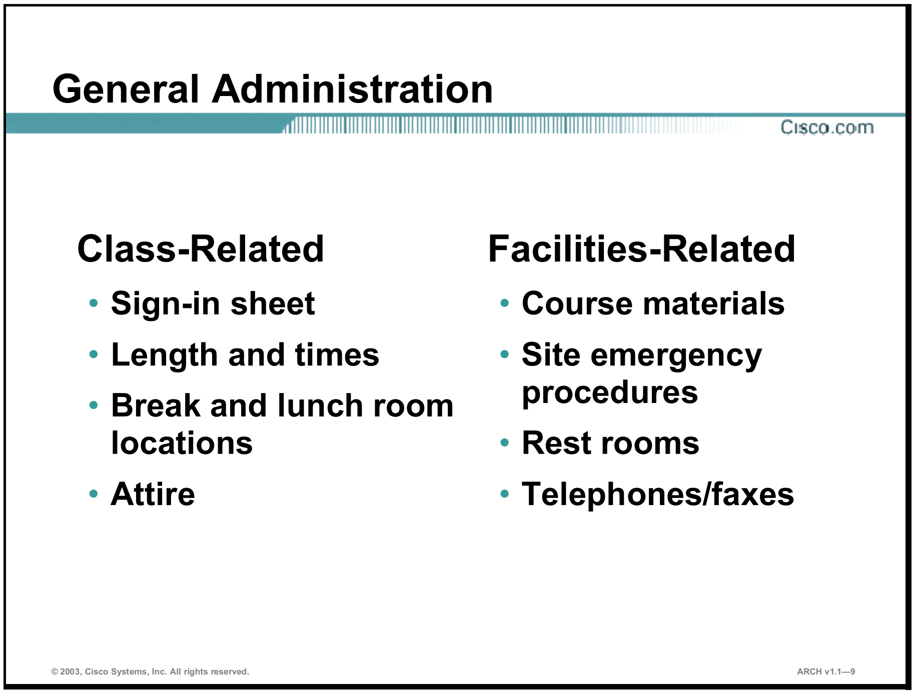

General Administration

Cisco.com

Class-Related

- Sign-in sheet
- Length and times
- Break and lunch room locations
- Attire

Facilities-Related

- Course materials
- Site emergency procedures
- Rest rooms
- Telephones/faxes

ARCH v1.1—9

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 Roadmap

This topic covers the suggested flow of the course materials.

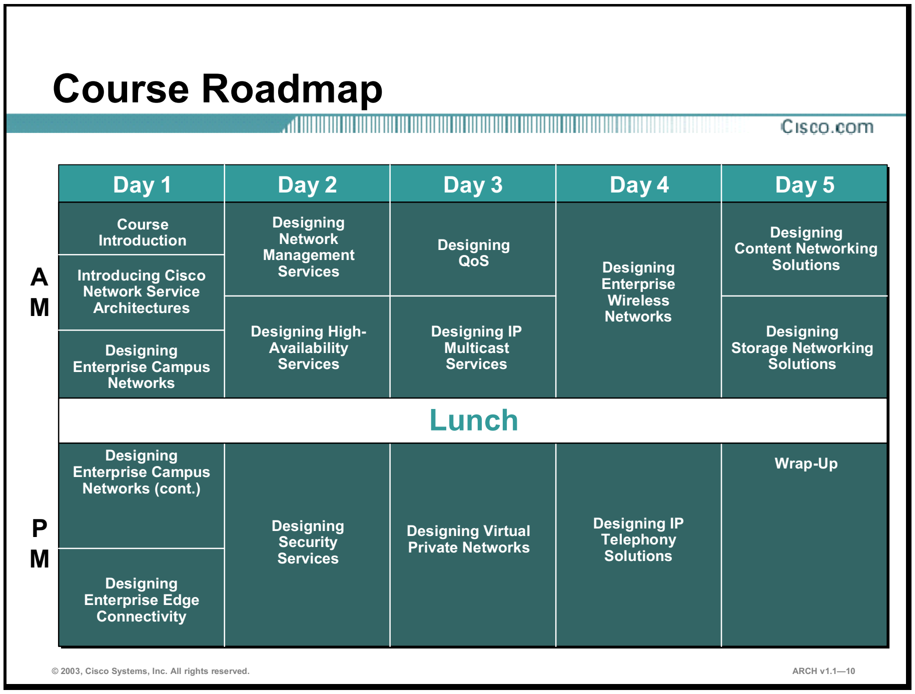

Course Roadmap

	Day 1	Day 2	Day 3	Day 4	Day 5
AM	Course Introduction Introducing Cisco Network Service Architectures Designing Enterprise Campus Networks	Designing Network Management Services Designing High-Availability Services	Designing QoS Designing IP Multicast Services	Designing Enterprise Wireless Networks	Designing Content Networking Solutions Designing Storage Networking Solutions
	Lunch				
PM	Designing Enterprise Campus Networks (cont.) Designing Enterprise Edge Connectivity	Designing Security Services	Designing Virtual Private Networks	Designing IP Telephony Solutions	Wrap-Up

ARCH v1.1—10

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 lesson assessments and exercises. The exact timing of the subject materials and labs depends on the pace of your specific class.

Icons and Symbols

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

Cisco Icons and Symbols

Cisco.com

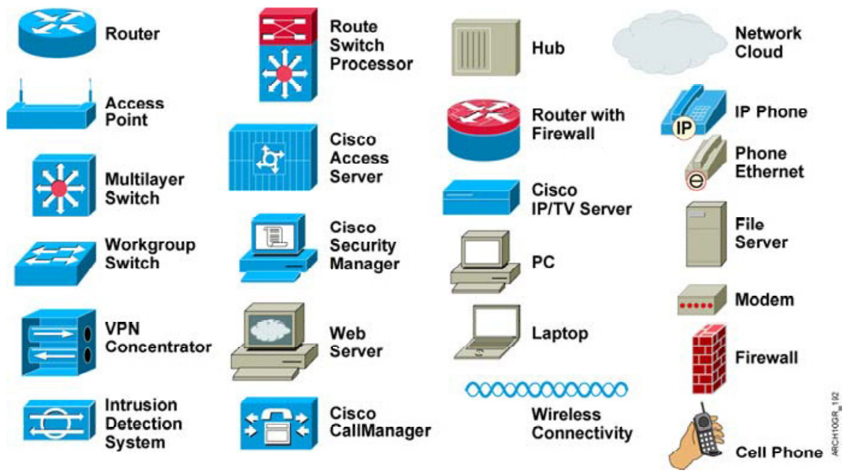

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11

Learner Introductions

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

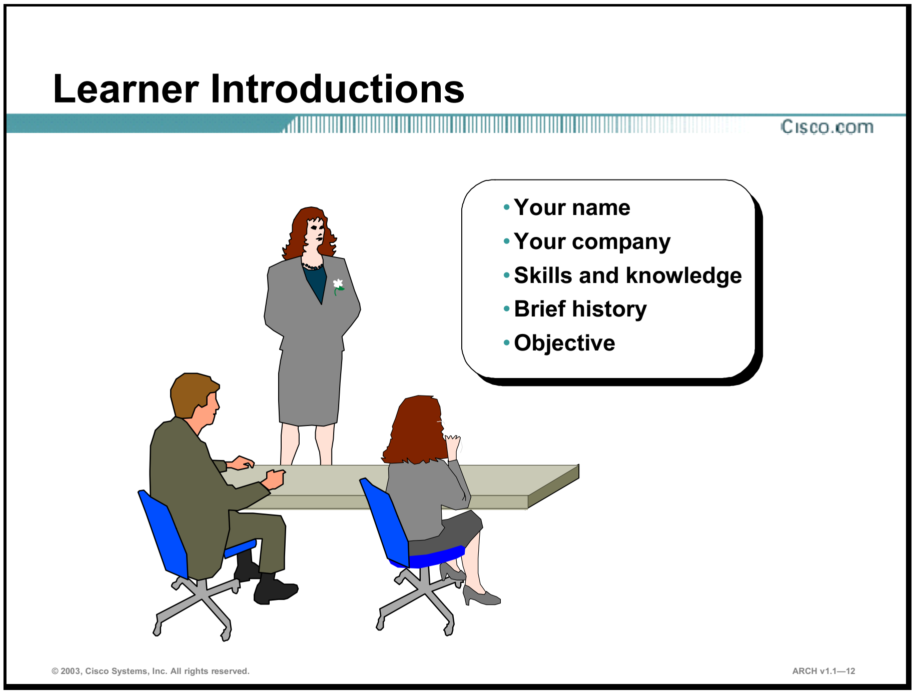

Learner Introductions

Cisco.com

- Your name
- Your company
- Skills and knowledge
- Brief history
- Objective

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12

Prepare to share this 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.

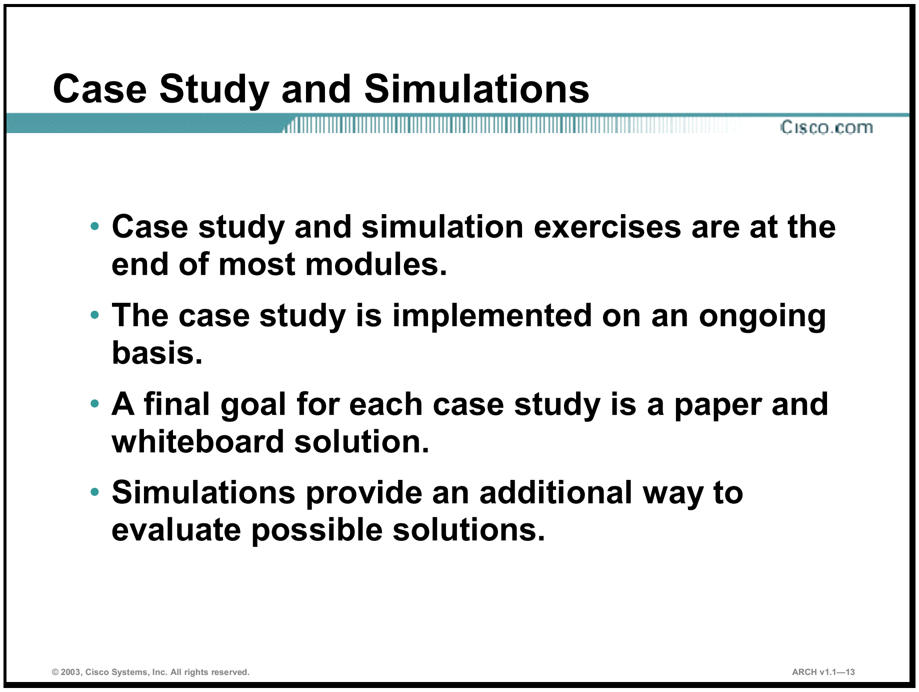

Case Study and Simulations

Cisco.com

- Case study and simulation exercises are at the end of most modules.
- The case study is implemented on an ongoing basis.
- A final goal for each case study is a paper and whiteboard solution.
- Simulations provide an additional way to evaluate possible solutions.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13

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 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. Your instructor will demonstrate the simulations.

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, such as technical factors and 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 a way to compare and contrast your solution with other possibilities.

Case Study Guidelines

Follow these guidelines as you complete the case study exercises:

1. 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.
2. You may use any and all documentation, books, white papers, and so on.
3. 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.
4. Use any design strategies that you feel are appropriate.
5. Use any internetworking technologies that you feel are appropriate.
6. A final goal for each case study is a paper and whiteboard solution. You do not need to provide the specific product names.

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.

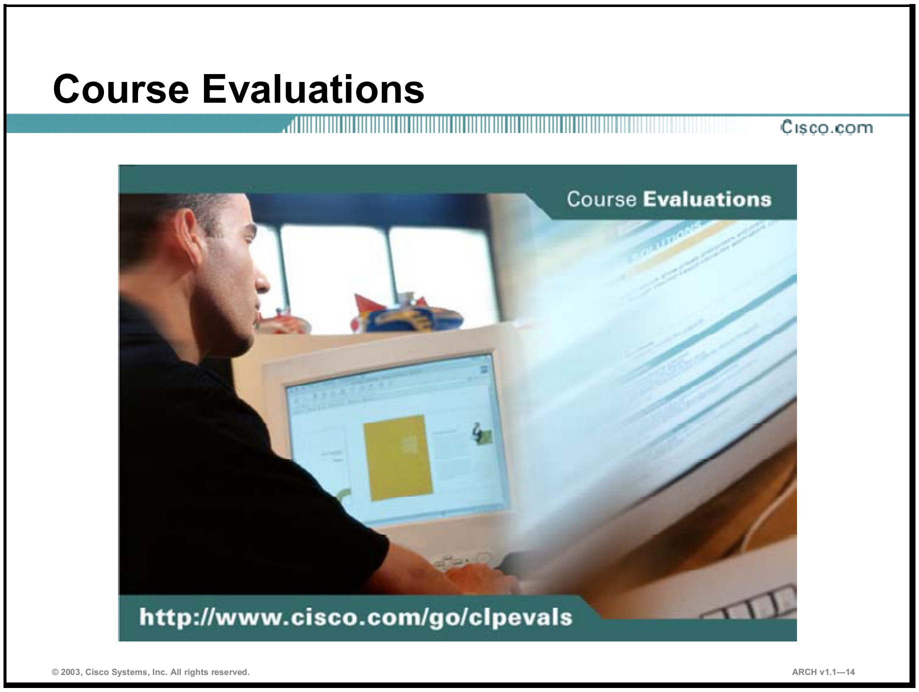

Course Evaluations

Cisco.com

Course Evaluations

http://www.cisco.com/go/clpevals

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—14

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

Introducing Cisco Network Service Architectures

Overview

Large enterprises increasingly seek an enterprise-wide infrastructure to serve as a solid foundation for emerging applications such as IP telephony, content networking, and storage networking. The Cisco Architecture for Voice, Video and Integrated Data (AVVID) framework, with its open communications interface, is the basis of Cisco's enterprise network architecture. The framework is designed to support the operation of concurrent solutions operating over a single infrastructure designed, tested, and fully documented with scalability, performance, and availability that meets end-to-end enterprise requirements.

Module Objectives

Upon completing this module, you will be able to present the Cisco Architecture for Voice, Video and Integrated Data (AVVID) framework, its segmentation of the network infrastructure, and intelligent network services to support key enterprise network applications and network solutions.

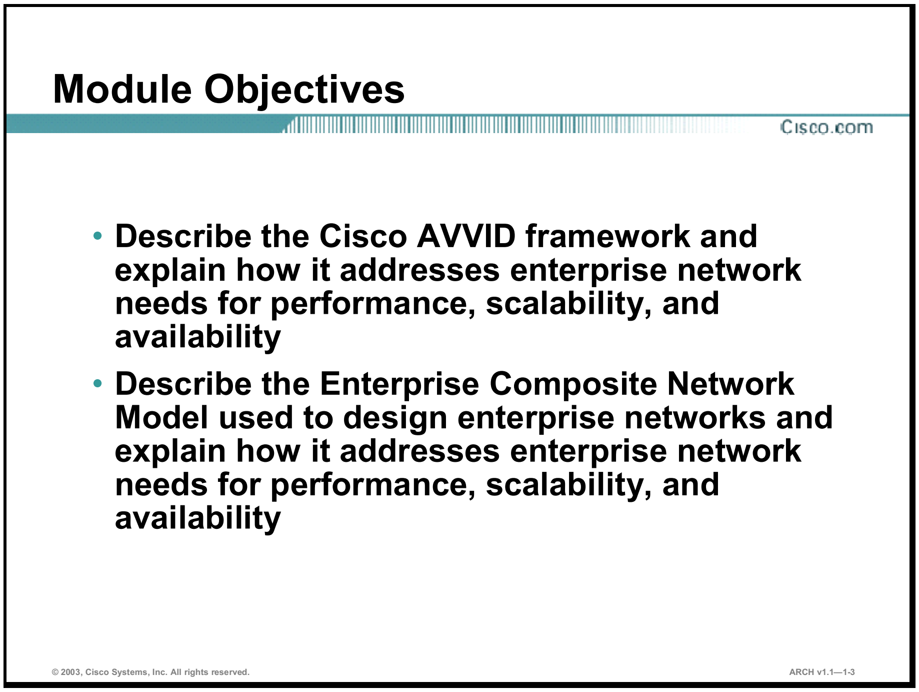

Module Objectives

- Describe the Cisco AVVID framework and explain how it addresses enterprise network needs for performance, scalability, and availability
- Describe the Enterprise Composite Network Model used to design enterprise networks and explain how it addresses enterprise network needs for performance, scalability, and availability

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—1-3

Module Outline

The outline lists the components of this module.

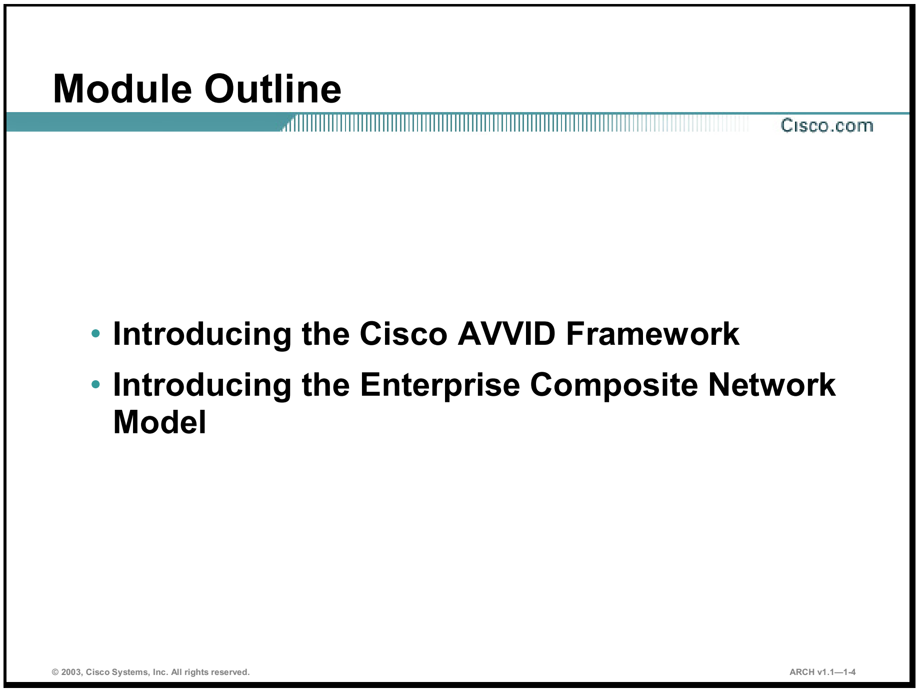

Module Outline

- Introducing the Cisco AVVID Framework
- Introducing the Enterprise Composite Network Model

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—1-4

Introducing the Cisco AVVID Framework

Overview

The Cisco AVVID framework provides an enterprise with a foundation that combines IP connectivity with performance and availability. Layering application solutions, such as voice, video, or content delivery networks, requires changes to an existing infrastructure. The Cisco AVVID framework provides effective design principles and practices to plan those changes. Each enterprise network is different because it is built to accommodate different topologies, media, and features that the specific enterprises may deploy.

Relevance

Network managers who design and build networks to support converged solutions combining data, voice, and video must consider the components that allow networks to operate properly. The Cisco AVVID framework provides an infrastructure on which to offer intelligent services to support network solutions and business applications.

Objectives

Upon completing this lesson, you will be able to describe the Cisco AVVID framework and explain how it addresses enterprise network needs for performance, scalability, and availability. This includes being able to meet these objectives:

- Describe the major components of the Cisco AVVID framework and explain why an architecture is important for enterprise networks
- Describe performance concerns when deploying an enterprise network
- Describe scalability concerns when deploying an enterprise network
- Describe availability concerns when deploying an enterprise network
- Describe the network infrastructure component of the Cisco AVVID framework, and explain how it satisfies enterprise requirements for performance, scalability, and availability

- Describe the intelligent network services of the Cisco AVVID framework, and explain how they support enterprise needs for performance, scalability, and availability
- Describe the Cisco AVVID network solutions offered to address the needs of enterprise applications

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

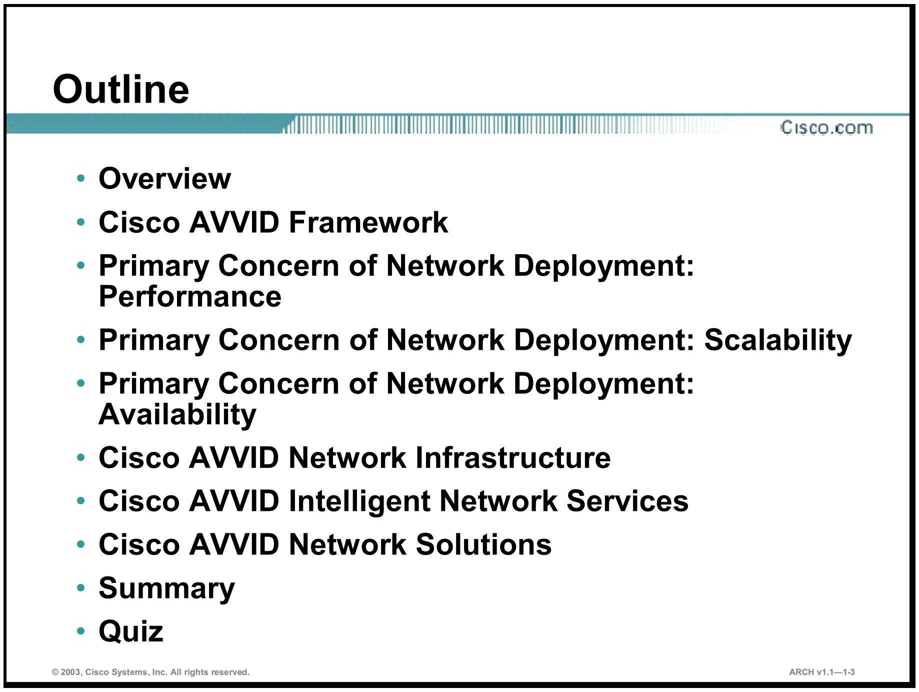

Outline

- **Overview**
- **Cisco AVVID Framework**
- **Primary Concern of Network Deployment: Performance**
- **Primary Concern of Network Deployment: Scalability**
- **Primary Concern of Network Deployment: Availability**
- **Cisco AVVID Network Infrastructure**
- **Cisco AVVID Intelligent Network Services**
- **Cisco AVVID Network Solutions**
- **Summary**
- **Quiz**

ARCH v1.1—1-3

Cisco AVVID Framework

Cisco AVVID provides the framework for today's Internet business solutions. Cisco AVVID is an enterprise-wide, standards-based network architecture that provides a roadmap for combining business and technology strategies into a cohesive model. This topic explains why an architecture is important for enterprise networks and describes the major components of the Cisco AVVID framework.

Cisco AVVID Framework

Cisco.com

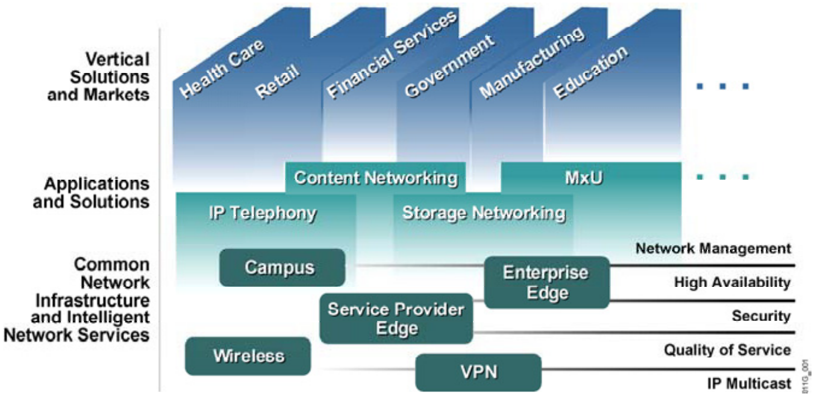

Vertical Solutions and Markets

Health Care

Retail

Financial Services

Government

Manufacturing

Education

. . .

Applications and Solutions

Content Networking

MxU

. . .

IP Telephony

Storage Networking

Common Network Infrastructure and Intelligent Network Services

Campus

Enterprise Edge

Service Provider Edge

Wireless

VPN

Network Management

High Availability

Security

Quality of Service

IP Multicast

8110_001

ARCH v1.1—1-4

A network architecture is a roadmap and guide for ongoing network planning, design, and implementation. It provides a coherent framework that unifies disparate solutions onto a single foundation.

The Cisco AVVID framework supports these key components:

- **Common network infrastructure:** Includes the hardware and software used to send, receive, and manage datagrams that are transmitted between end-user devices throughout the enterprise. It includes the transmission media and devices that control transmission paths, including private and public transport media. Examples of these devices are routers, LAN switches, WAN switches, PBXs, and so on.
- **Intelligent network services:** Allow the end user to operate in a controlled, secure environment in which differentiated services are provided. Intelligent network services essentially add intelligence to the network infrastructure beyond just moving a datagram between two points. The intelligent network services allow for application awareness. Intelligent network services include network management, high availability, security, quality of service (QoS), and IP multicast.

- **Network solutions:** Include the hardware and software that use the network infrastructure and intelligent network services to their advantage. Network solutions allow enterprises to make business decisions about the business itself as well as about networks and the technologies and applications that run on them. Network-based applications enable an enterprise organization to interact more effectively with customers, suppliers, partners, and employees. Customer service, commerce, supplier, and internal applications run over the network infrastructure enabled by intelligent network services. Some examples of network solutions are IP telephony, content networking, and storage networking, among others.

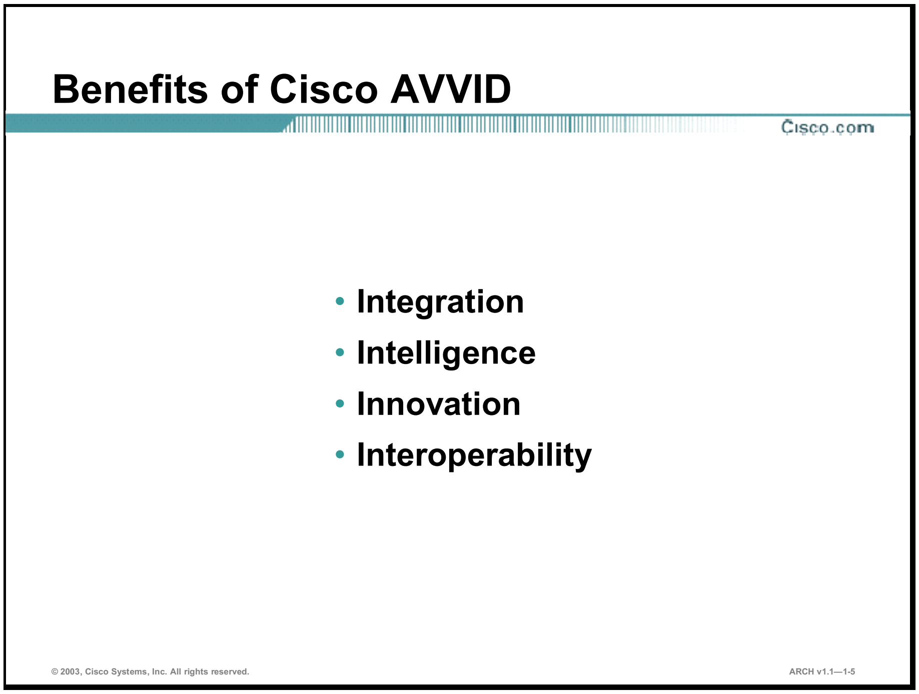

Benefits of Cisco AVVID

Cisco.com

- Integration
- Intelligence
- Innovation
- Interoperability

ARCH v1.1—1-5

Cisco AVVID offers these benefits:

- **Integration:** By leveraging the Cisco AVVID framework and applying the network intelligence inherent in IP, organizations can enable comprehensive tools to improve productivity.
- **Intelligence:** Traffic prioritization and intelligent networking services maximize network efficiency for optimized application performance.
- **Innovation:** Customers have the ability to adapt quickly in a competitive and changing business environment.
- **Interoperability:** Standards-based hardware and software interfaces allow open integration, providing organizations with choice and flexibility.

Combining the network infrastructure and services with new applications, Cisco AVVID accelerates the integration of technology strategy with business activities. Cisco AVVID is an enabler of Internet business solutions for enterprises via the network infrastructure.

Primary Concern of Network Deployment: Performance

While specific devices or applications may promise performance, effective performance is achieved only by considering and optimizing each component. Only a cohesive, integrated, and optimized network can ensure the best network performance. This topic describes the performance concerns when deploying an enterprise network.

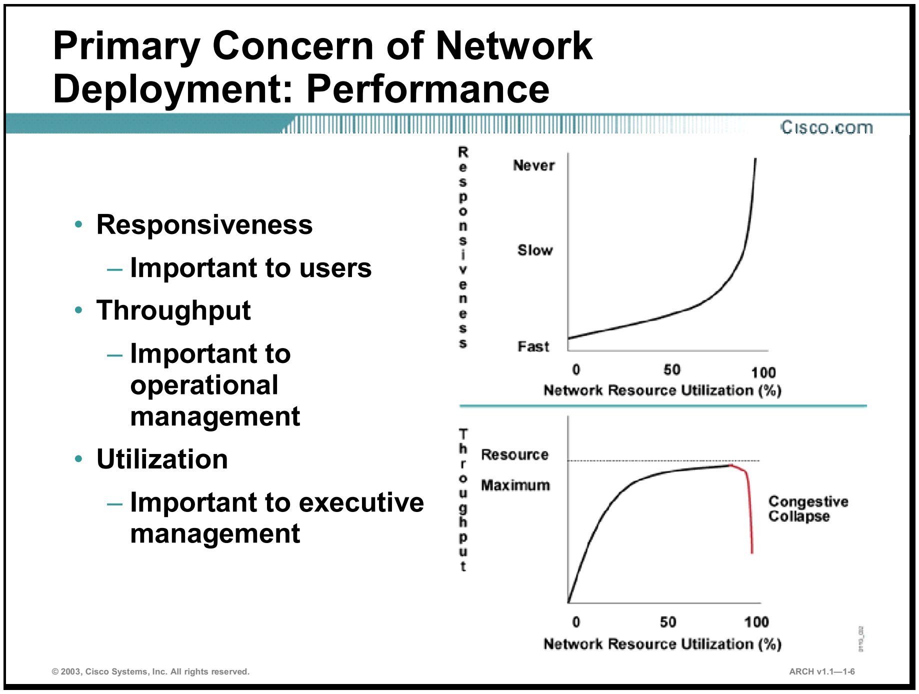

Primary Concern of Network Deployment: Performance

Cisco.com

- Responsiveness
  - Important to users
- Throughput
  - Important to operational management
- Utilization
  - Important to executive management

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-6

Performance might be the least understood term in networking. Typically, performance is defined as throughput and packets per second (pps). These are easy numbers to gauge and report, but these values relate to a single switch or router and make no sense when measuring an entire network. For example, one can state that a network should perform at 10,000 pps, but testing over the network might yield only 5000 pps. What happened? In fact, the network might consist of 1.536-Mbps (T1) WAN links with traffic shaping enabled, which in turn limits the packet rate through the entire network. In addition, forwarding traffic at that rate might impact the processor loads, limiting the overall throughput performance and placing the router at risk of not having enough resources, either to converge following a failure in the network or to enable additional features. The point is that there is no one metric for determining performance.

Instead, gauge network performance by these three metrics:

- **Responsiveness:** Indicates how a user or consumer perceives the performance of their applications. It is affected by link speeds, congestion, and features, and includes device and protocol responses. This is the most important metric in the network: if an application does not respond in an acceptable time, it does not matter how fast the network claims to be. This metric changes based on how an application responds to changes in the network. For example, many applications use TCP, which slows the transmission rate into the network if too much congestion or loss is present in the network.
- **Throughput:** Specifies the rate of information arriving at, and possibly passing through, a particular point in a network system. Throughput is closely related to utilization. As utilization increases, throughput approaches the theoretical maximum until driven to congestive collapse. Typically, throughput is measured in pps, kbps, Mbps, and Gbps.
- **Utilization:** Measures the use of a particular resource over time. The measure is usually expressed as a percentage, where the usage of a resource is compared with its maximum operational capacity. Through utilization measures, you can identify congestion (or potential congestion) throughout the network. You can also identify underutilized resources.

  Utilization is the principle measure to determine how full the network pipes (links) are. Analyzing CPU, interface, queuing, and other system-related capacity measurements allows you to determine the extent to which network system resources are being consumed. High utilization is not necessarily bad. Low utilization may indicate traffic flows in unexpected places. As lines become overutilized, the effects can become significant. Overutilization on a link occurs when there is consistently more traffic, which needs to pass through more of an interface than it can handle. Ultimately, there will be excessive queuing delays and even packet loss. Sudden jumps in resource utilization can indicate a fault condition.

Primary Concern of Network Deployment: Scalability

A network must be able to scale from where it is today to where it might be in the future. For example, a network administrator might need to design the WAN to support only 50 branch offices. However, over a year's time, 50 more branches might require connectivity. The design, IP address management, features, and WAN link speeds must all be able to accommodate this need for added connectivity without massive redesign of the network. This topic describes scalability concerns when deploying an enterprise network.

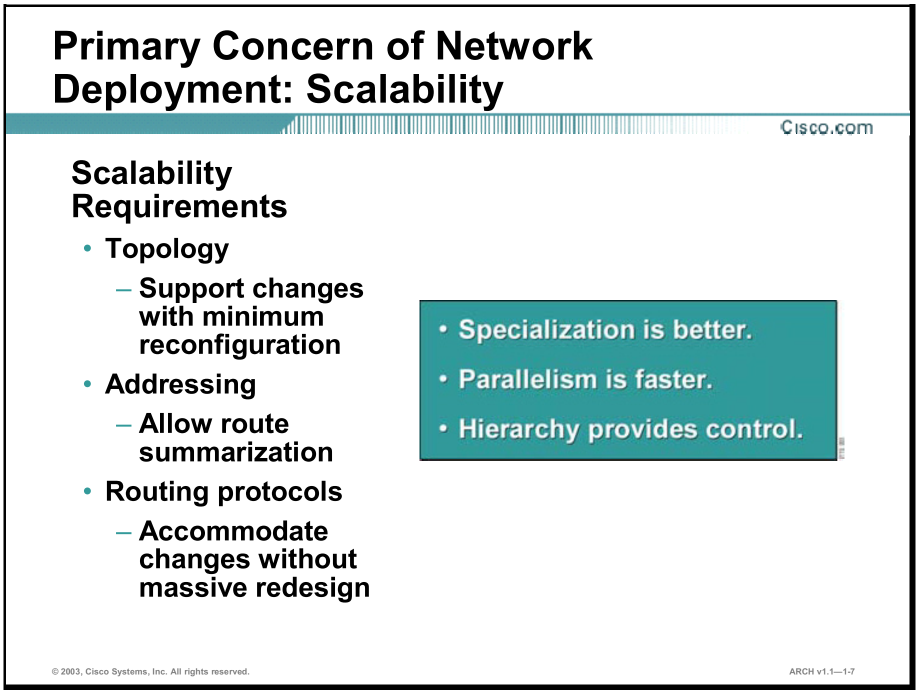

Primary Concern of Network Deployment: Scalability

Cisco.com

Scalability Requirements

- Topology
  - Support changes with minimum reconfiguration
- Addressing
  - Allow route summarization
- Routing protocols
  - Accommodate changes without massive redesign

- Specialization is better.
- Parallelism is faster.
- Hierarchy provides control.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—1-7

When designing an enterprise network, you should try to implement some basic principles throughout the network to improve scalability. Specialization of devices and card modules for specific functions makes it easy to upgrade each device as the network grows. Parallelism in the network design improves overall network performance. By implementing a hierarchy, you will achieve more control and manageability of the network.

Specific network scalability concerns include:

- **Topology:** Network topology must be such that additions or subtractions to the network do not cause major reconfigurations, create instability, affect deterministic performance, or adversely affect availability levels.
- **Addressing:** Distribution of IP addresses should facilitate route summarization. Additionally, it should be possible to create new subnets with a minimum impact on the addressing scheme and router load.
- **Routing protocols:** The routing protocol of choice must be able to accommodate additions, deletions, and changes without a massive redesign.

Primary Concern of Network Deployment: Availability

A major concern for network managers is how available the network is and how impervious it is to network changes. A network that takes ten seconds to converge is clearly superior to one that takes one minute to converge.

To the user, the network is down regardless of whether an application went down, a router died, or a piece of fiber was cut. For this reason, availability must be viewed from the user's perspective. This topic describes availability concerns when deploying an enterprise network.

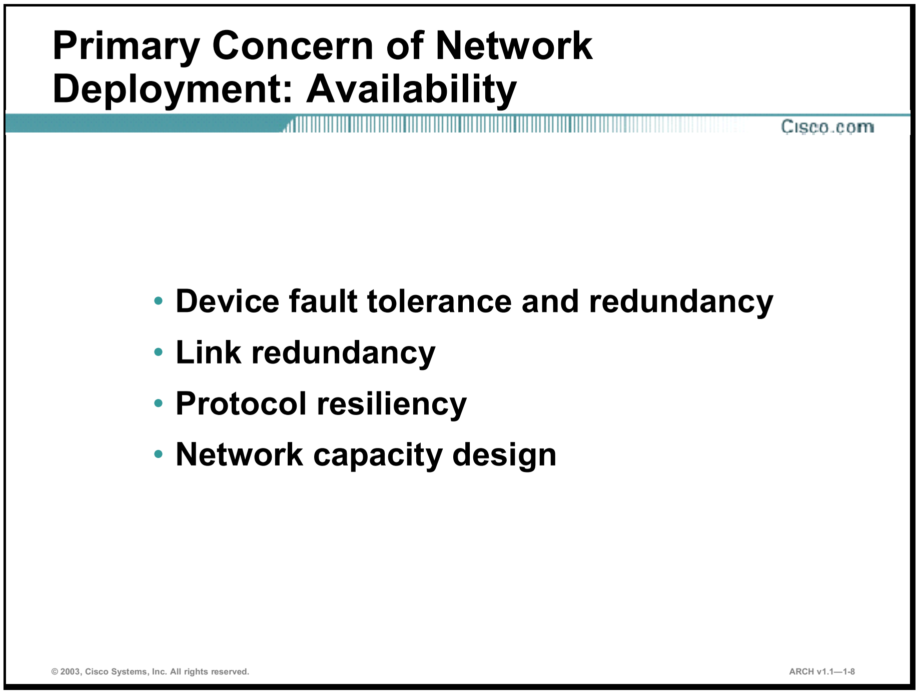

Primary Concern of Network Deployment: Availability

- Device fault tolerance and redundancy
- Link redundancy
- Protocol resiliency
- Network capacity design

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—1-8

Key availability issues to address include:

- **Device fault tolerance and redundancy:** This is often the first level of availability in the network. Fault-tolerant devices provide a high level of reliability. Cisco offers options for redundant supervisor engines and dual power supplies, which provide the first backstop against a network failure.
- **Link redundancy:** Link redundancy is critical in the network, and provides a high level of reliability in the event of a link failure. However, while some redundancy is good, more redundancy is not necessarily better.
- **Protocol resiliency:** Good design practices indicate how and when to use protocol redundancy, including load-sharing, convergence speed, and path redundancy handling.
- **Network capacity design:** Good design practices consider capacity planning. How much traffic can a connection handle in the worst-case scenario? Network designers must ascertain whether a link can handle twice the traffic when a redundant link fails.

Cisco AVVID Network Infrastructure

The Cisco AVVID framework consists of several building blocks that deliver solutions to accelerate the enterprise. The network infrastructure components include clients and servers, network platforms, and intelligent network services. This topic describes the network infrastructure component of the Cisco AVVID framework, and explains how it meets enterprise needs for performance, scalability, and availability.

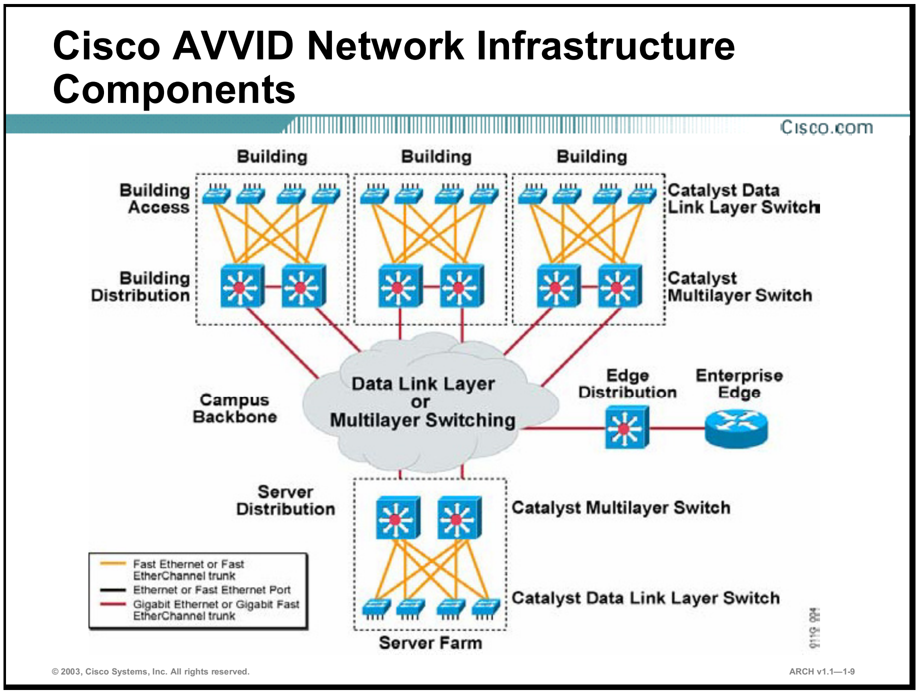

Cisco AVVID Network Infrastructure Components

Cisco.com
Building
Building
Building
Building Access
Building Distribution
Catalyst Data Link Layer Switch
Catalyst Multilayer Switch
Campus Backbone
Data Link Layer or Multilayer Switching
Edge Distribution
Enterprise Edge
Server Distribution
Catalyst Multilayer Switch
Catalyst Data Link Layer Switch
Server Farm
Fast Ethernet or Fast EtherChannel trunk
Ethernet or Fast Ethernet Port
Gigabit Ethernet or Gigabit Fast EtherChannel trunk

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-9

The Cisco AVVID network infrastructure consists of these hardware components:

- **Clients and application servers:** Network clients include workstations (both fixed and portable), IP phones, and wireless devices. Application servers provide services to clients, and may be located in a data center or other easily accessible network location.
- **Network platforms:** The network platforms comprise routers, gateways, switches, servers, firewalls, and other devices. This component of the architecture provides the basis for a complete networking solution.
- **Intelligent network services:** Intelligent network services include the platforms, network services, appliances, and management that allow business rules and policies to positively affect network performance.

The Cisco AVVID network infrastructure solution provides an enterprise foundation that combines IP connectivity with high performance and availability. Although layering application solutions such as voice, video, or content delivery networks require changes to the network infrastructure, this infrastructure provides a basis for good design principles and practices.

Each network is different because it is built to accommodate different topologies (mesh or hub-and-spoke), WAN technologies (such as Frame Relay, ATM, or PPP), and networks (LAN, WAN, or metropolitan-area network [MAN]) that enterprises deploy.

Network managers who design and build networks to support solutions such as voice and video must first consider the components that allow networks to operate properly. Thus, the network device often becomes the focus of design decisions. However, a single device, whether a switch, router, or other networking device, is only a component of the overall network. How the devices connect, what features and protocols are used, and how they are used form the foundation for the services that run on top of the network. If the foundation is unstable, layering solutions over the network can create problems.

By laying the foundation for basic connectivity and protocol deployment, the Cisco AVVID network infrastructure solution addresses the three primary concerns of network deployment: Performance, scalability, and availability.

Cisco AVVID Intelligent Network Services

Cisco AVVID network infrastructure supports the key intelligent network services, which comprise numerous networking technologies and topologies, with a corresponding large number of possible designs and architectures. The net result is a blueprint that blends equipment, features, and management tools that match business criteria. This topic describes the intelligent network services of the Cisco AVVID framework, and explains how they meet enterprise needs for performance, scalability, and availability.

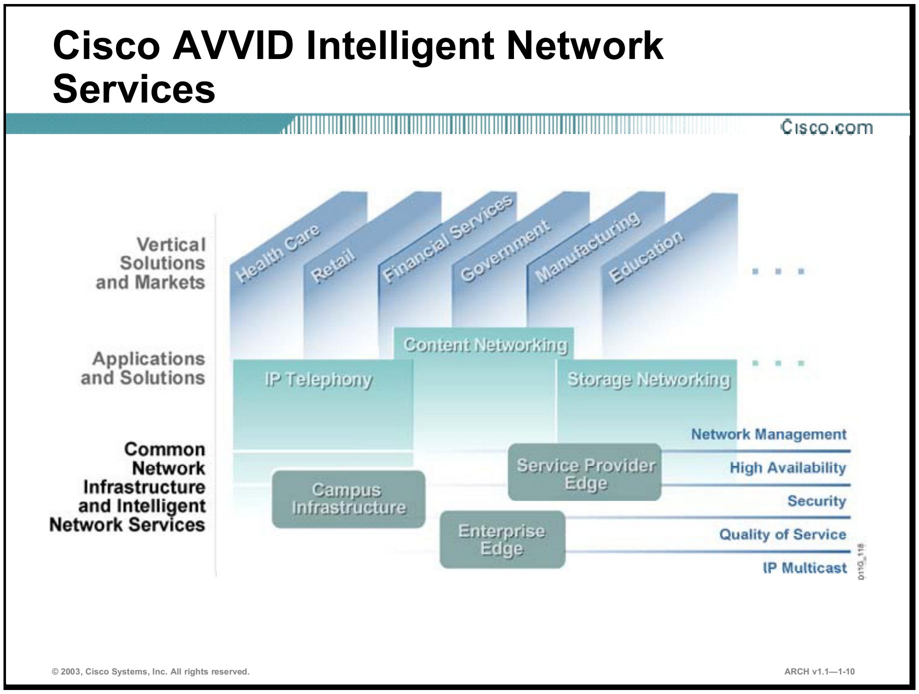

Cisco AVVID Intelligent Network Services

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-10

Cisco deploys these intelligent network services to keep the network at peak performance:

- **Network management:** Provides a number of related network management tools built on a Common Management Foundation (CMF). Tools include the LAN Management Solution for advanced management of Catalyst multilayer switches; the Routed WAN Management Solution for monitoring traffic management and providing access control to administer the routed infrastructure of multiservice networks; the Service Management Solution for managing and monitoring service level agreements; and the VPN/Security Management Solution for optimizing VPN performance and security administration.
- **High availability:** Refines design and tools to ensure end-to-end availability for services, clients, and sessions. Tools include reliable, fault-tolerant network devices to automatically identify and overcome failures, and resilient network technologies, such as Hot Standby Router Protocol (HSRP), to bring resilience to the critical junction between hosts and backbone links.
- **Security:** Ensures the security of the network through authentication, encryption, and failover. Security features include application-based filtering (context-based access control), intrusion detection in the network and at hosts, defense against network attacks, per-user authentication and authorization, and real-time alerts.

- **QoS:** Manages the delay, delay variation (jitter), bandwidth, and packet loss parameters on a network to meet the diverse needs of voice, video, and data applications. QoS features provide 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, and Resource Reservation Protocol (RSVP) signaling for admission control and reservation of resources.
- **IP multicast:** Provides bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of end-system clients. Multicasting enables distribution of videoconferencing, corporate communications, distance learning, distribution of software, and other applications. Multicast packets are replicated in the network by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols resulting in an efficient delivery of data to multiple receivers.

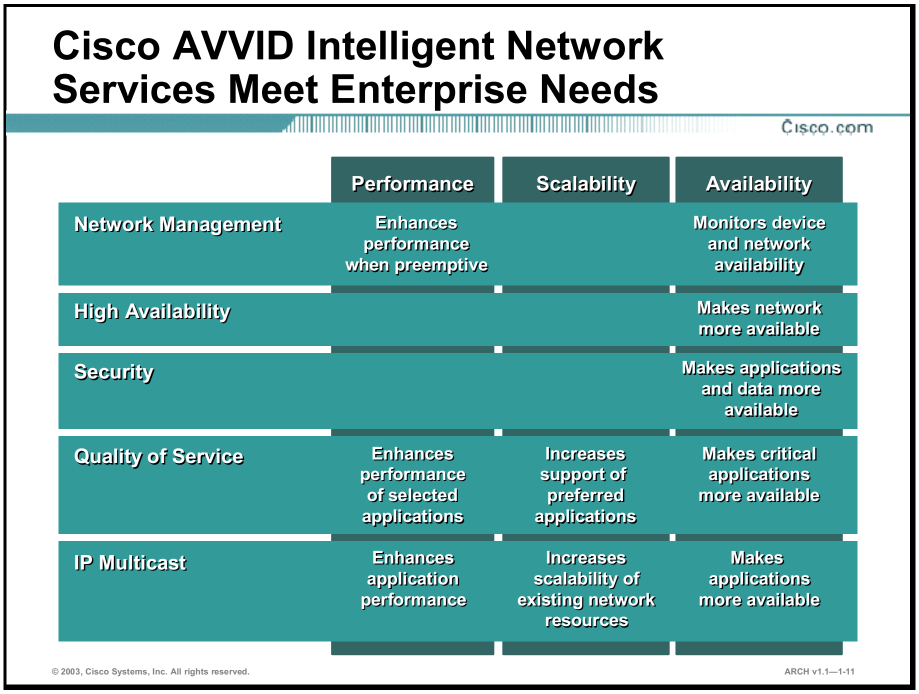

Cisco AVVID Intelligent Network Services Meet Enterprise Needs

Cisco.com

	Performance	Scalability	Availability
Network Management	Enhances performance when preemptive		Monitors device and network availability
High Availability			Makes network more available
Security			Makes applications and data more available
Quality of Service	Enhances performance of selected applications	Increases support of preferred applications	Makes critical applications more available
IP Multicast	Enhances application performance	Increases scalability of existing network resources	Makes applications more available

ARCH v1.1—1-11

The figure describes how the Cisco AVVID intelligent network services meet enterprise network needs for performance, scalability, and availability.

Cisco AVVID Network Solutions

Enterprises can make a competitive investment in their future by deploying specific solutions. The Cisco AVVID framework provides a foundation for applications and solutions. Cisco provides some solutions while third-party companies provide solutions through the Cisco AVVID Partner Program. This topic describes the Cisco AVVID network solutions that Cisco offers to address enterprise application needs.

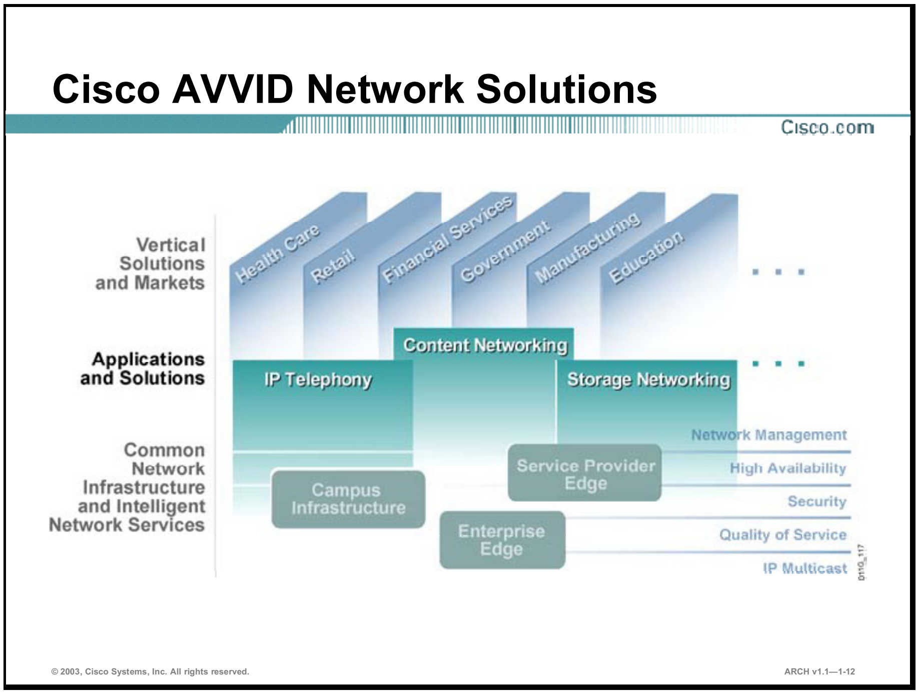

Cisco AVVID Network Solutions

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-12

Cisco provides these network infrastructure and application solutions, which are discussed in this course:

- **VPN (part of the Enterprise Edge):** VPNs use advanced encryption and tunneling to permit organizations to establish secure, end-to-end, private network connections over third-party networks, such as the Internet or extranets.
- **Wireless (part of the Campus Infrastructure):** Wireless and IP technology creates anytime, anywhere connections to the Internet and enterprise networks. In a campus environment or distant mobile location, wireless technology allows users to be constantly connected as they move between wireless cells, unconstrained by direct physical connections.
- **IP telephony:** The convergence of voice, video, and data on a single IP network is changing the way enterprises communicate. You can transport voice, video, and data on a single network infrastructure, lowering total network costs and optimizing enterprise communications.
- **Content networking:** Content networking provides an architecture that optimizes website performance and content delivery by positioning content near consumers in anticipation of use.

- **Storage networking:** Driven by workforce collaboration, e-commerce, and e-learning, storage networking has emerged as an important networking application. Cisco storage networking solutions provide high-capacity, low-latency networking for disaster recovery, data replication, and storage consolidation.

Summary

This topic summarizes the key points discussed in this lesson.

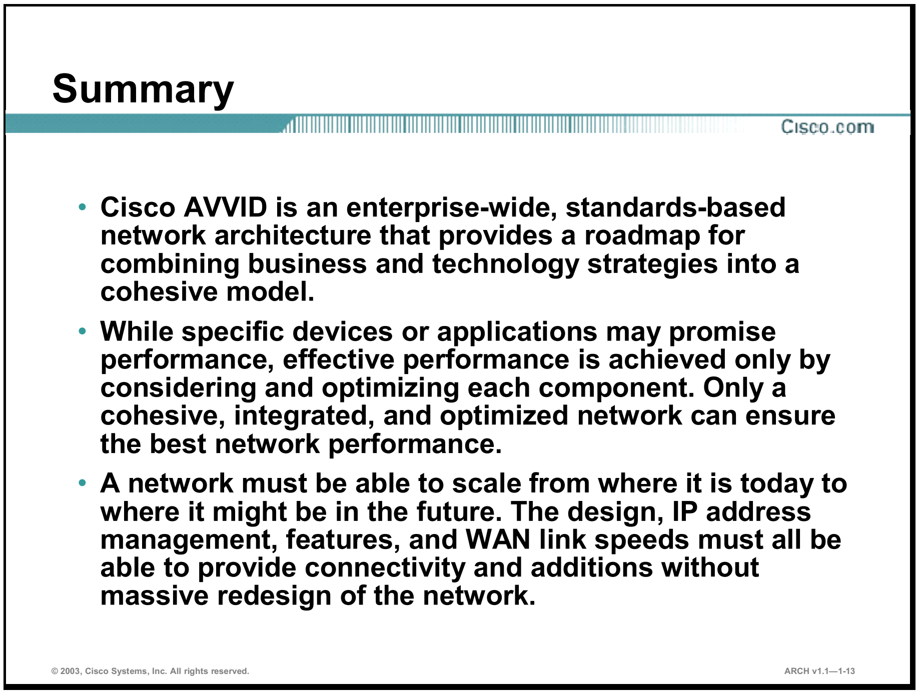

Summary

Cisco.com

- Cisco AVVID is an enterprise-wide, standards-based network architecture that provides a roadmap for combining business and technology strategies into a cohesive model.
- While specific devices or applications may promise performance, effective performance is achieved only by considering and optimizing each component. Only a cohesive, integrated, and optimized network can ensure the best network performance.
- A network must be able to scale from where it is today to where it might be in the future. The design, IP address management, features, and WAN link speeds must all be able to provide connectivity and additions without massive redesign of the network.

ARCH v1.1—1-13

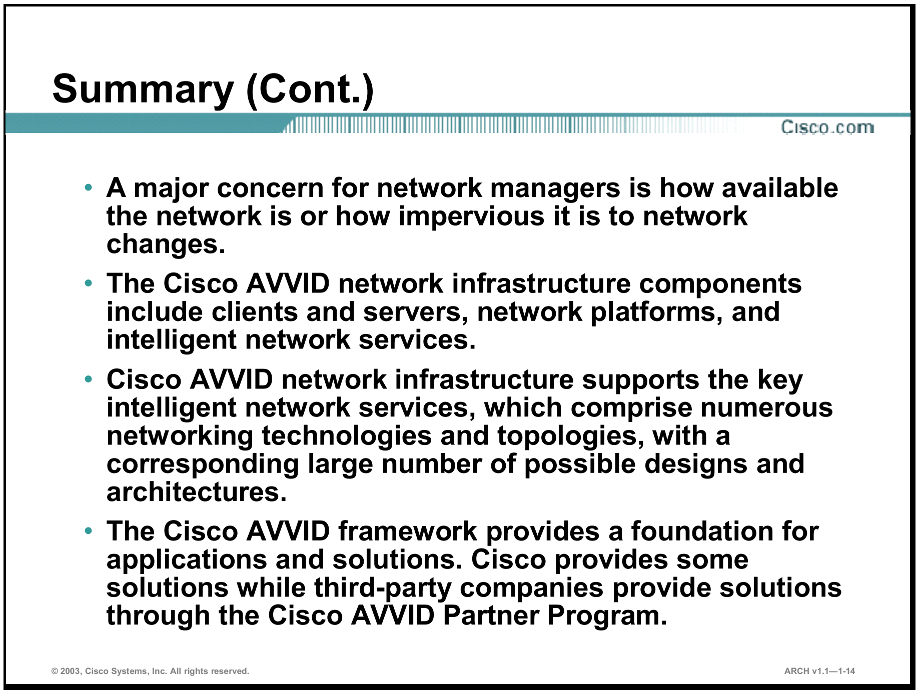

Summary (Cont.)

Cisco.com

- A major concern for network managers is how available the network is or how impervious it is to network changes.
- The Cisco AVVID network infrastructure components include clients and servers, network platforms, and intelligent network services.
- Cisco AVVID network infrastructure supports the key intelligent network services, which comprise numerous networking technologies and topologies, with a corresponding large number of possible designs and architectures.
- The Cisco AVVID framework provides a foundation for applications and solutions. Cisco provides some solutions while third-party companies provide solutions through the Cisco AVVID Partner Program.

ARCH v1.1—1-14

References

For additional information, refer to this resource:

- *Cisco AVVID: Enabling E-Business* at http://www.cisco.com/warp/public/779/largeent/avvid/cisco_avvid.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 functions does a network architecture facilitate? (Choose three.)

A) network design

B) network planning

C) network monitoring

D) network implementation

E) voice and data integration

Q2) Cisco AVVID meets the need for _____ by providing standards-based interfaces that allow open integration and that provide organizations with choice and flexibility.

A) innovation

B) integration

C) intelligence

D) interoperability

Q3) Which three metrics describe performance on the network? (Choose three.)

A) capacity

B) utilization

C) throughput

D) availability

E) responsiveness

Q4) Which statement best describes the requirement for a network topology to contribute to scalability?

A) The network topology must allow route summarization.

B) The network topology must support changes without reconfigurations.

C) The network topology must accommodate routing protocols without a massive redesign.

D) The network topology must allow creation of new subnets without affecting the addressing scheme.

Q5) What are three primary components of network availability? (Choose three.)

A) hierarchy

B) responsiveness

C) link redundancy

D) protocol resiliency

E) equipment fault tolerance

Q6) How does the Cisco AVVID network infrastructure solution address the primary concerns of network deployment (performance, scalability, availability)?

A) allows the network designer to support voice and video

B) supports technologies such as Frame Relay, ATM, and PPP

C) lays the foundation of basic connectivity and protocol deployment

D) provides box solutions that support any level of performance, scalability, and availability

Q7) Match each intelligent network service to its description.

_____ 1. security

_____ 2. IP multicast

_____ 3. high availability

_____ 4. quality of service

_____ 5. network management

A) ensures the integrity of the network through authentication, encryption, and failover

B) refines design to ensure resources are present end-to-end for services, clients, and sessions

C) manages the delay, delay variation, bandwidth, and packet loss parameters on a network to meet the diverse needs of critical applications

D) relies on the LAN Management Solution, Routed WAN Management Solution, Service Management Solution, and VPN/Security Management Solution

E) provides bandwidth-conserving technology that reduces traffic by simultaneously delivering a reduced number of information streams to thousands of client end stations

Q8) Match each Cisco AVVID network solution to its description.

_____ 1. VPN

_____ 2. wireless

_____ 3. IP telephony

_____ 4. storage networking

_____ 5. content networking

A) architecture that optimizes website performance and content delivery

B) allows users to be constantly connected as they move freely within different environments

C) transports voice, video, and data, lowering network costs and optimizing enterprise communications

D) provides high-capacity, low-latency networking for disaster recovery, data replication, and storage consolidation

E) uses advanced encryption and tunneling to permit organizations to establish secure, end-to-end, private network connections over third-party networks

Quiz Answer Key

Q1) A, B, D

Relates to: Cisco AVVID Framework

Q2) D

Relates to: Cisco AVVID Framework

Q3) B, C, E

Relates to: Primary Concern of Network Deployment: Performance

Q4) B

Relates to: Primary Concern of Network Deployment: Scalability

Q5) C, D, E

Relates to: Primary Concern of Network Deployment: Availability

Q6) C

Relates to: Cisco AVVID Network Infrastructure

Q7) 1-A, 2-E, 3-B, 4-C, 5-D

Relates to: Cisco AVVID Intelligent Network Services

Q8) 1-E, 2-B, 3-C, 4-D, 5-A

Relates to: Cisco AVVID Network Solutions

Introducing the Enterprise Composite Network Model

Overview

The Enterprise Composite Network Model provides a framework for designing the components of an enterprise network. The model relies on the principles of the Cisco AVVID, which provides a framework for the solutions presented in this course.

Relevance

Developing a common vocabulary and architecture is critical to designing modular enterprise networks that provide performance, scalability, and availability.

Objectives

Upon completing this lesson, you will be able to describe the Enterprise Composite Network Model used to design enterprise networks and explain how it addresses enterprise network needs for performance, scalability, and availability. This includes being able to meet these objectives:

- Describe the Enterprise Composite Network Model and explain how it addresses enterprise network modularity
- Describe the modules that comprise an enterprise campus network, and explain how the Enterprise Campus functional area meets the need for performance, scalability, and availability
- Describe the components and functionality at the Enterprise Edge, and explain how the Enterprise Edge functional area meets the need for performance, scalability, and availability
- Describe the components and functionality at the Service Provider Edge, and explain how the Service Provider Edge functional area meets the need for performance, scalability, and availability

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course or passing the Cisco CCDA® certification exam

Outline

The outline lists the topics included in this lesson.

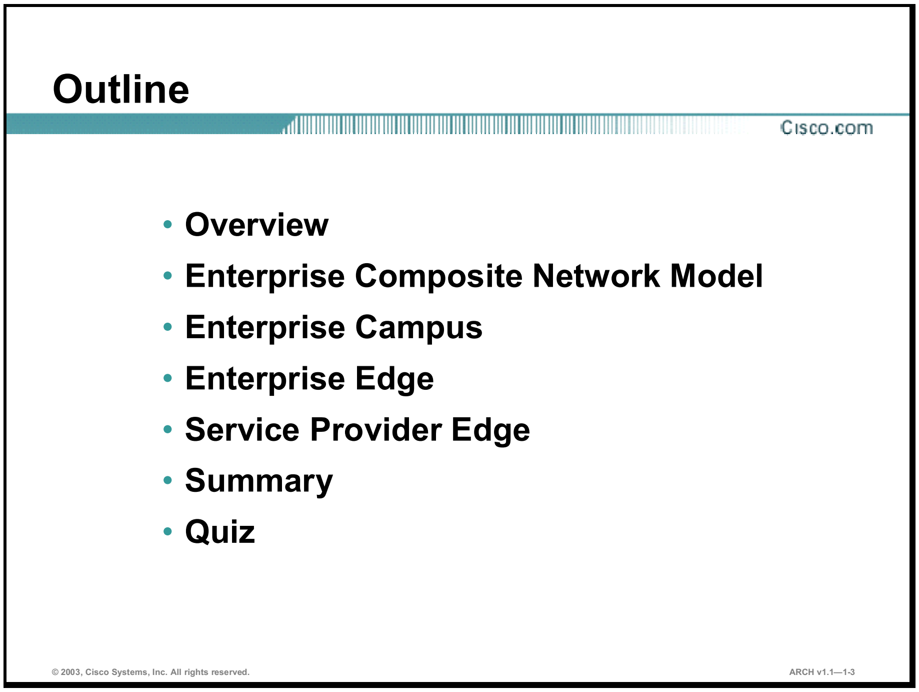

Outline

- Overview
- Enterprise Composite Network Model
- Enterprise Campus
- Enterprise Edge
- Service Provider Edge
- Summary
- Quiz

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-3

Enterprise Composite Network Model

The Enterprise Composite Network Model provides a modular framework for designing networks. The modularity within the model allows flexibility in network design and facilitates implementation and troubleshooting. This topic describes the Enterprise Composite Network Model and explains how it addresses enterprise network requirements for modularity.

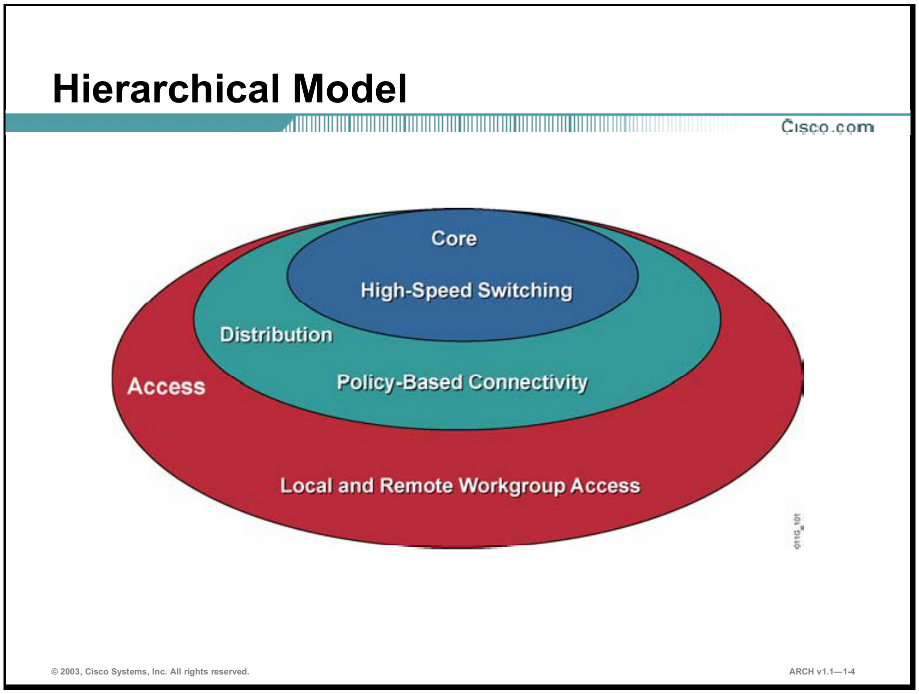

Hierarchical Model

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-4

Nearly a decade ago, Cisco introduced a hierarchical design model as a tool for network designers to approach network design from the physical, logical, and functional viewpoints. The hierarchical model divided networks into these layers:

- **Access layer:** The access layer is used to grant user access to network devices. At a network campus, the access layer incorporates shared, switched, or subnetted LAN devices with ports available to workstations and servers. In the WAN environment, the access layer can provide sites with access to the corporate network using a WAN technology.
- **Distribution layer:** The distribution layer aggregates the wiring closets and uses data link layer switching and multilayer switching to segment workgroups and isolate network problems, preventing them from impacting the core layer. Routing and packet manipulation occur in the distribution layer.
- **Core layer:** The core layer is a high-speed backbone and is designed to switch packets as fast as possible. Because the core is critical for connectivity, it must provide a high level of availability and must adapt to changes very quickly.

The hierarchical module was useful, but had weaknesses when implementing large, complex enterprise networks.

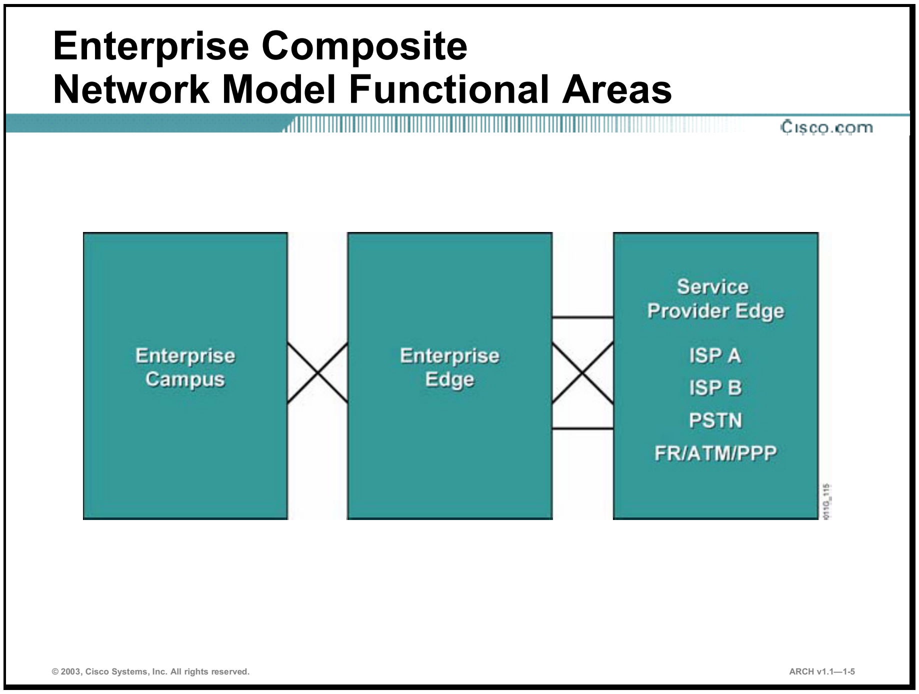

Enterprise Composite Network Model Functional Areas

Enterprise Campus

Enterprise Edge

Service Provider Edge

ISP A

ISP B

PSTN

FR/ATM/PPP

ARCH v1.1—1-5

The Enterprise Composite Network Model introduces additional modularity into the network structure. The entire network is divided into functional areas that contain the hierarchical model access, distribution, and core layers.

The Enterprise Composite Network Model contains three major functional areas:

- **Enterprise Campus:** Contains the modules required to build a hierarchical, highly robust campus network that offers performance, scalability, and availability. This functional area contains the network elements required for independent operation within a single campus. This functional area does not offer remote connections or Internet access.

  A campus is defined as one or more buildings, with multiple virtual and physical networks, connected across a high-performance, multilevel-switched backbone.

- **Enterprise Edge:** Aggregates connectivity from the various elements at the edge of the enterprise 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 of the network elements for efficient and secure communication between the Enterprise Campus and remote locations, remote users, and the Internet.

- **Service Provider Edge:** Provides functionality implemented by service providers. The Service Provider Edge functional area enables communications with other networks using different WAN technologies and Internet service providers (ISPs).

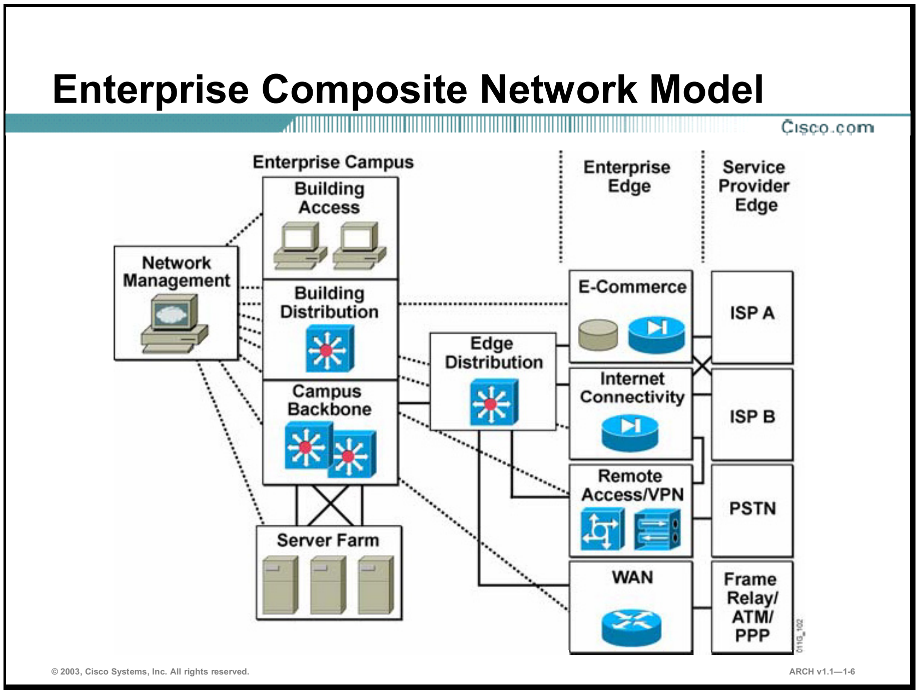

Enterprise Composite Network Model

ARCH v1.1—1-6

To scale the hierarchical model, Cisco introduced the Enterprise Composite Network Model that further divides the enterprise network into physical, logical, and functional boundaries. The Enterprise Composite Network Model contains functional areas, each of which has its own access, distribution, and core layers.

The Enterprise Composite Network Model meets the following criteria:

- Defines a deterministic network with clearly defined boundaries between modules. The model has clear demarcation points to aid the designer in knowing exactly where traffic is.
- Increases network scalability and eases the design task by making each module discrete.
- Provides scalability by allowing enterprises to add modules easily. As network complexity grows, designers can add new functional modules.
- Offers more integrity in network design, allowing the designer to add services and solutions without changing the underlying network design.

The Enterprise Composite Network Model provides a conceptual view of an enterprise network. The model describes the functions to include in a network design but does not specifically define the devices and connections that are required. As you design a network, you will need to select the devices, interfaces, wiring, and so on, that the specific design requires.

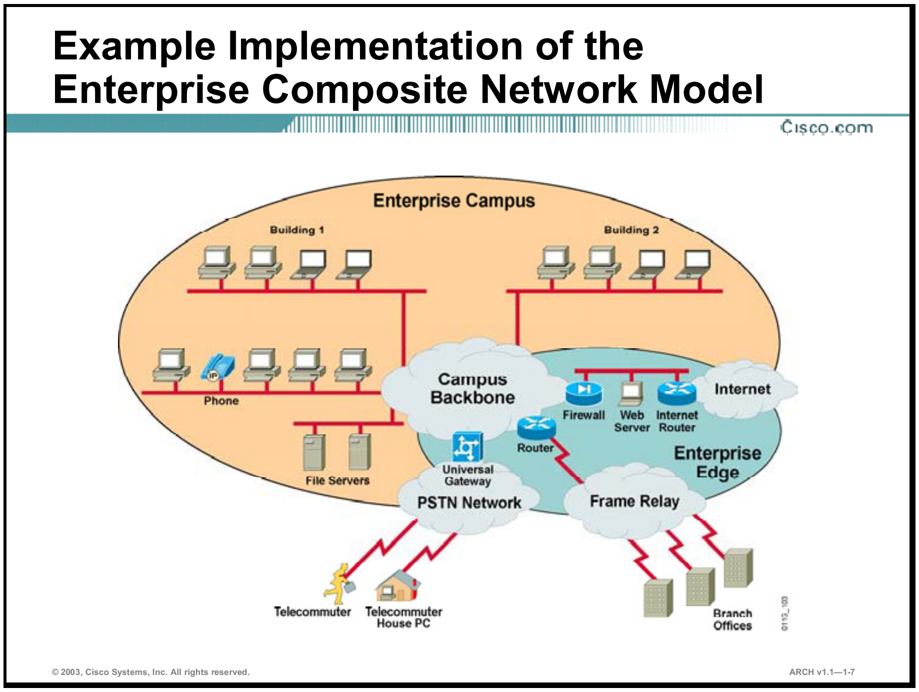

Example Implementation of the Enterprise Composite Network Model

ARCH v1.1—1-7

The figure shows an enterprise network with the major components of the Enterprise Composite Network Model. The network is divided into the Enterprise Campus and Enterprise Edge functional areas, connected by the Campus Backbone submodule.

Enterprise Campus

The Enterprise Campus functional area includes the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules. Each module has a specific function within the campus network. This topic describes the modules that comprise an enterprise campus network, and explains how the infrastructure meets the need for performance, scalability, and availability.

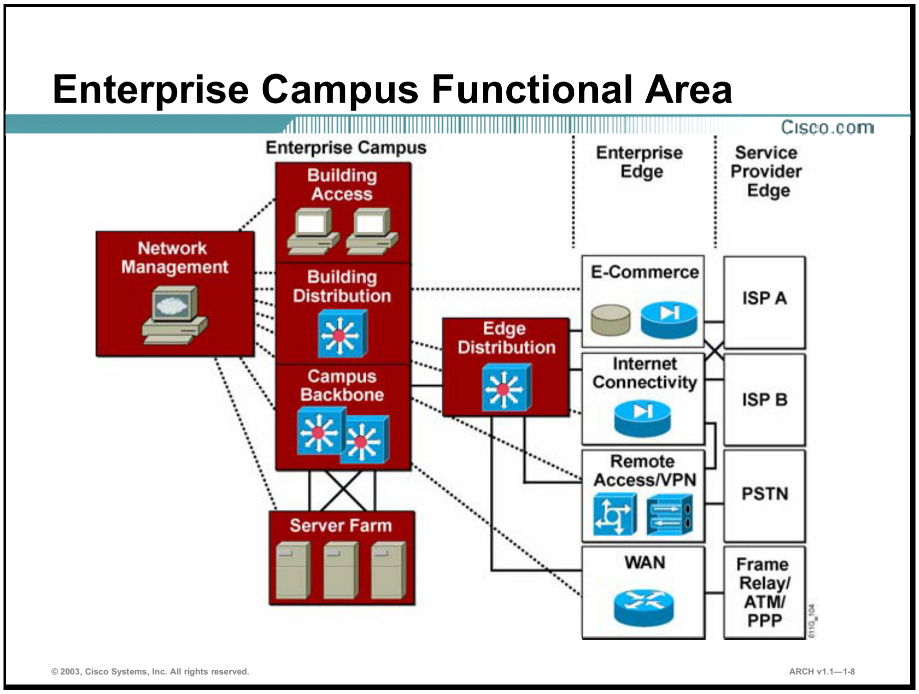

Enterprise Campus Functional Area

Enterprise Campus

Enterprise Edge

Service Provider Edge

Building Access

Network Management

Building Distribution

Campus Backbone

Edge Distribution

Server Farm

E-Commerce

Internet Connectivity

Remote Access/VPN

WAN

ISP A

ISP B

PSTN

Frame Relay/ ATM/ PPP

ARCH v1.1—1-8

The Enterprise Campus functional area includes these four major modules:

- Campus Infrastructure
- Network Management
- Server Farm
- Edge Distribution

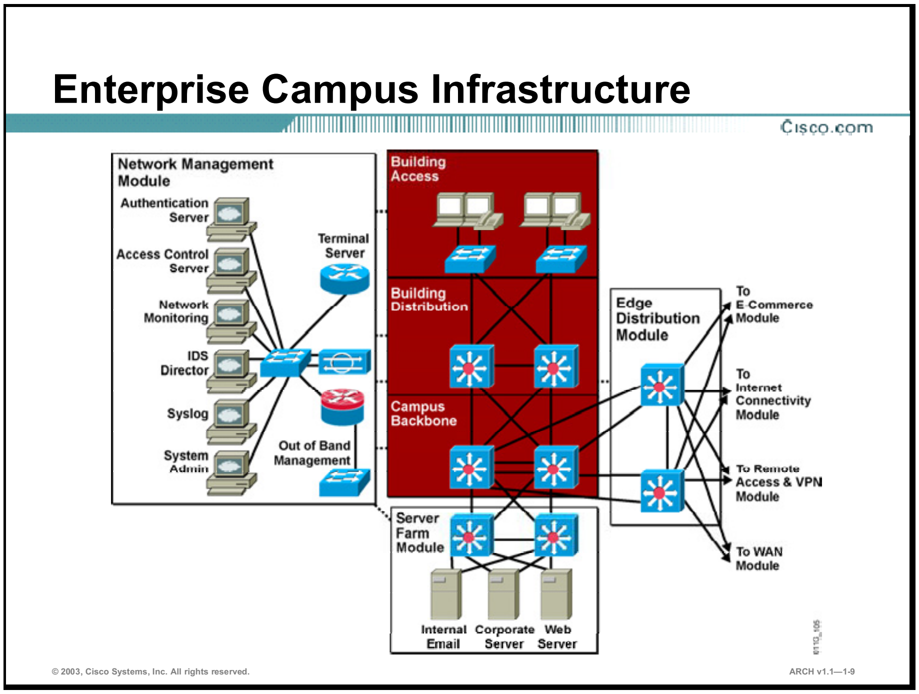

Enterprise Campus Infrastructure

Cisco.com

ARCH v1.1—1-9

The Campus Infrastructure module connects users within a campus with the Server Farm and Edge Distribution modules. This module is composed of one or more floors or buildings connected to the Campus Backbone submodule. Each building contains a Building Access and Building Distribution submodule.

The Campus Infrastructure module includes these submodules:

- **Building Access (also known as the access layer):** Contains end-user workstations, IP Phones, and data link layer access switches that connect devices to the Building Distribution submodule. Building Access performs important services such as broadcast suppression, protocol filtering, network access, and quality of service marking.
- **Building Distribution (also known as the distribution layer):** Provides aggregation of wiring closets, often using multilayer switching. The Building Distribution submodule performs routing, quality of service, and access control. Requests for data flow into the Building Distribution switches and to the campus backbone. The model provides fast failure recovery, since each Building Distribution switch maintains two equal-cost paths in the routing table to every destination network. When one connection to the campus backbone fails, all routes immediately switch over to the remaining path about one second after the link failure is detected.
- **Campus Backbone (also known as the core layer):** 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. This module uses data link layer switches or multilayer switches for high throughput functions with added routing, QoS, and security features.

In addition to the Campus Infrastructure module, the Enterprise Campus functional area includes these modules:

- **Network Management:** Performs intrusion detection, system logging, and authentication, as well as network monitoring and general configuration management functions. For management purposes, an out-of-band connection (a network on which no production traffic resides) to all network components is recommended. The Network Management module provides configuration management for nearly all devices in the network using Cisco routers and dedicated network management stations.
- **Server Farm:** 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, enabling full redundancy and load sharing. The Server Farm module switches are cross-connected with core-layer switches, enabling high reliability and availability of all servers.
- **Edge Distribution:** Aggregates the connectivity from the various elements at the Enterprise Edge functional area 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 edge distribution devices to perform additional security.

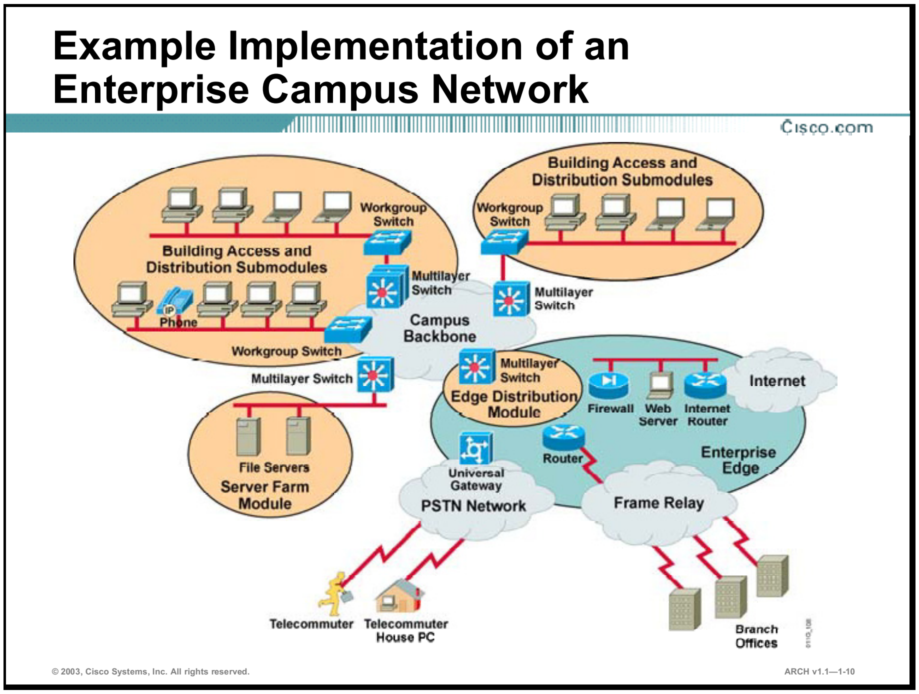

Example Implementation of an Enterprise Campus Network

ARCH v1.1—1-10

The figure shows how the enterprise campus network is divided into easily managed building blocks, including the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules. Notice that the Enterprise Edge functional area is reachable only through the Edge Distribution module.

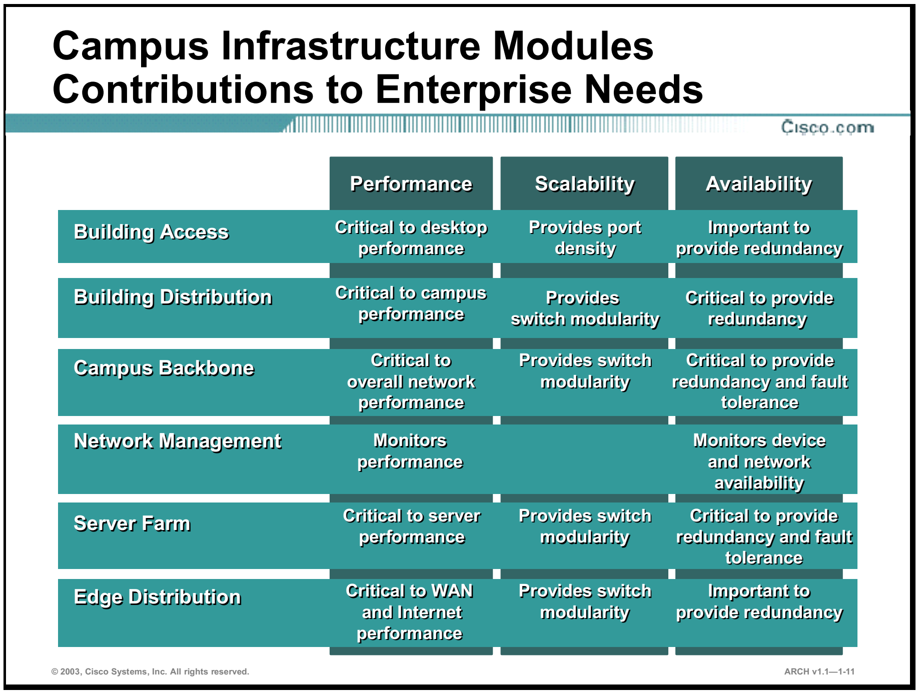

Campus Infrastructure Modules Contributions to Enterprise Needs

	Performance	Scalability	Availability
Building Access	Critical to desktop performance	Provides port density	Important to provide redundancy
Building Distribution	Critical to campus performance	Provides switch modularity	Critical to provide redundancy
Campus Backbone	Critical to overall network performance	Provides switch modularity	Critical to provide redundancy and fault tolerance
Network Management	Monitors performance		Monitors device and network availability
Server Farm	Critical to server performance	Provides switch modularity	Critical to provide redundancy and fault tolerance
Edge Distribution	Critical to WAN and Internet performance	Provides switch modularity	Important to provide redundancy

ARCH v1.1—1-11

The figure describes how the campus network meets the enterprise network needs for performance, scalability, and availability.

Enterprise Edge

The Enterprise Edge functional area is comprised of four modules: E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN. Each module connects to the Edge Distribution module, which connects the Enterprise Edge and Enterprise Campus functional areas. This topic describes the components and functionality at the Enterprise Edge functional area, and explains how the Enterprise Edge functional area meets the need for performance, scalability, and availability.

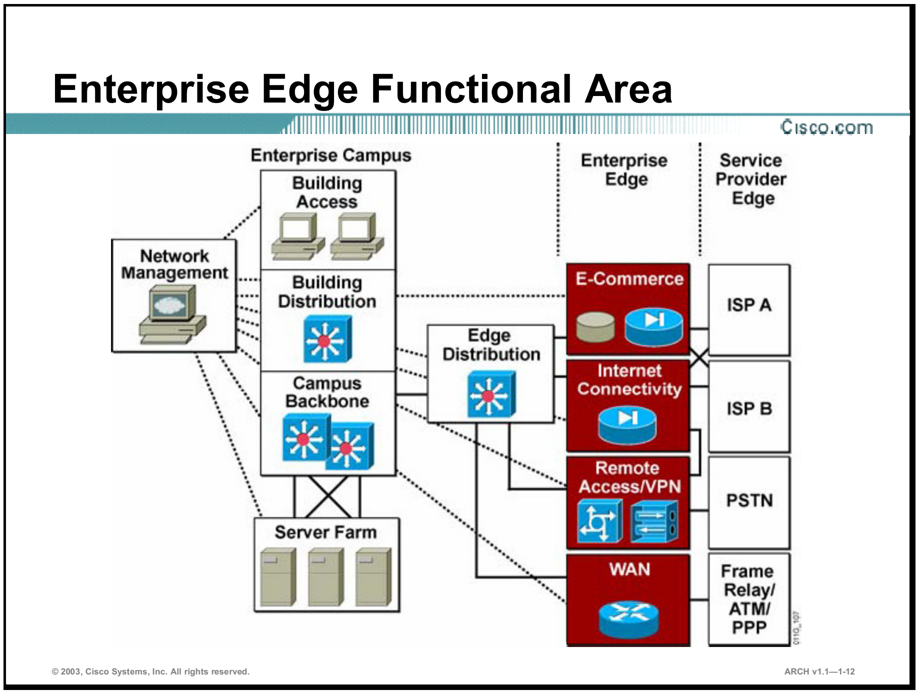

Enterprise Edge Functional Area

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-12

The Enterprise Edge module is comprised of four modules:

- E-Commerce
- Internet Connectivity
- Remote Access and VPN
- WAN

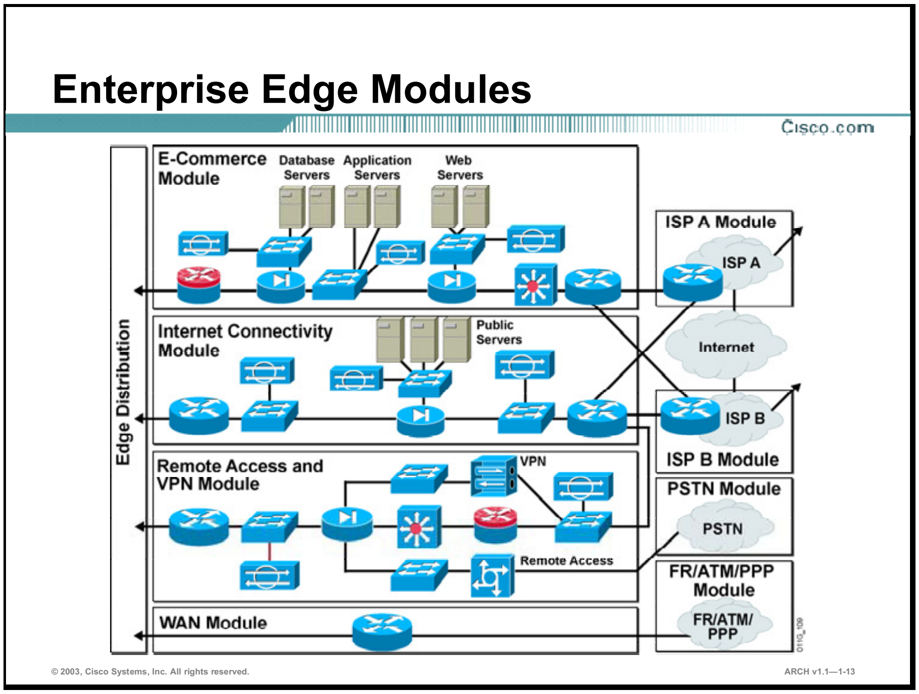

Enterprise Edge Modules

ARCH v1.1—1-13

The Enterprise Edge modules perform the following functions:

- **E-Commerce:** Enables enterprises to successfully deploy e-commerce applications and take advantage of the powerful competitive opportunities provided by the Internet. All e-commerce transactions pass through a series of intelligent services to provide performance, scalability, and availability within the overall e-commerce network design. To build a successful e-commerce solution, enterprises may deploy the following network devices:
  - **Web servers:** Act as the primary user interface for the navigation of e-commerce
  - **Application servers:** Support enterprise applications including online transaction processing systems and decision support applications
  - **Database servers:** Contain the critical information that is the heart of e-commerce business implementation
  - **Security servers:** Govern communication between the various levels of security in the system, often using firewalls and intrusion detection systems
- **Internet Connectivity:** Provides internal users with connectivity to Internet services. Internet users can access the information on publicly available servers. Additionally, this module accepts VPN traffic from remote users and remote sites and forwards it to the Remote Access and VPN module. The major components of the Internet Connectivity module are:
  - **E-mail servers:** Act as a relay between the Internet and the intranet mail servers
  - **DNS servers:** Serve as authoritative external DNS servers for the enterprise and relay internal requests to the Internet
  - **Public web servers:** Provide public information about the organization
  - **Security servers:** Govern communication between the various levels of security in the system, often using firewalls and intrusion detection systems
  - **Edge routers:** Provide basic filtering and Layer 3 connectivity to the Internet

- **Remote Access and VPN:** Terminates VPN traffic, forwarded by the Internet Connectivity module, from remote users and remote sites. It also initiates VPN connections to remote sites through 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. The 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 and intrusion detection systems:** Provide network-level protection of resources and stateful filtering of traffic; provide differentiated security for remote access users
  - **Data link layer switches:** Provide data link layer connectivity for devices
- **WAN:** Routes traffic between remote sites and the central site. The WAN module supports WAN physical technologies including leased lines, optical, cable, digital subscriber lines (DSLs), and wireless, as well as data link protocols such as Frame Relay, ATM, and PPP.

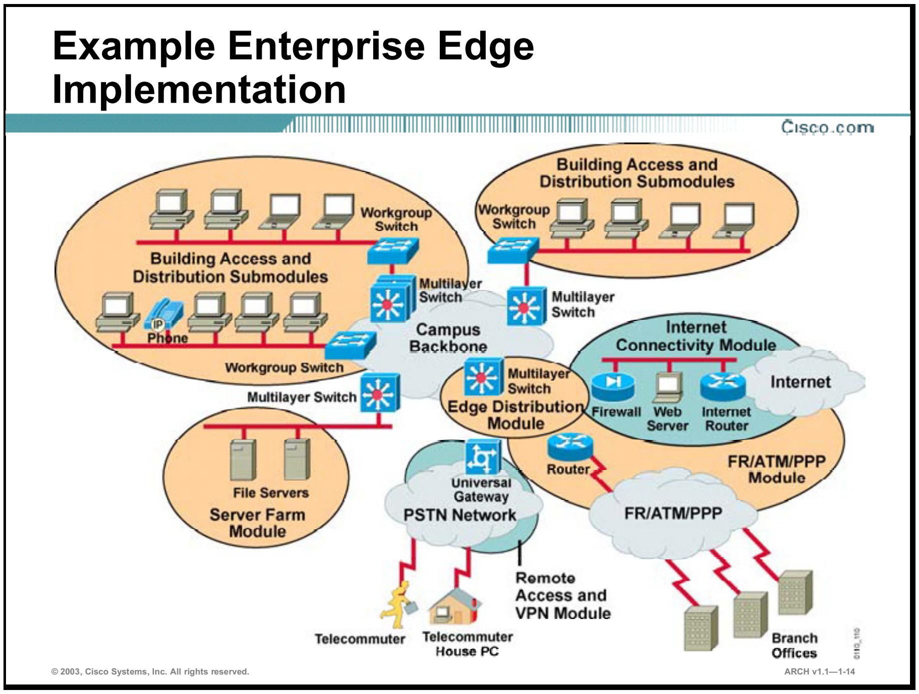

Example Enterprise Edge Implementation

ARCH v1.1—1-14

The figure shows how the Enterprise Edge functional area is divided into easily managed building blocks including E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules.

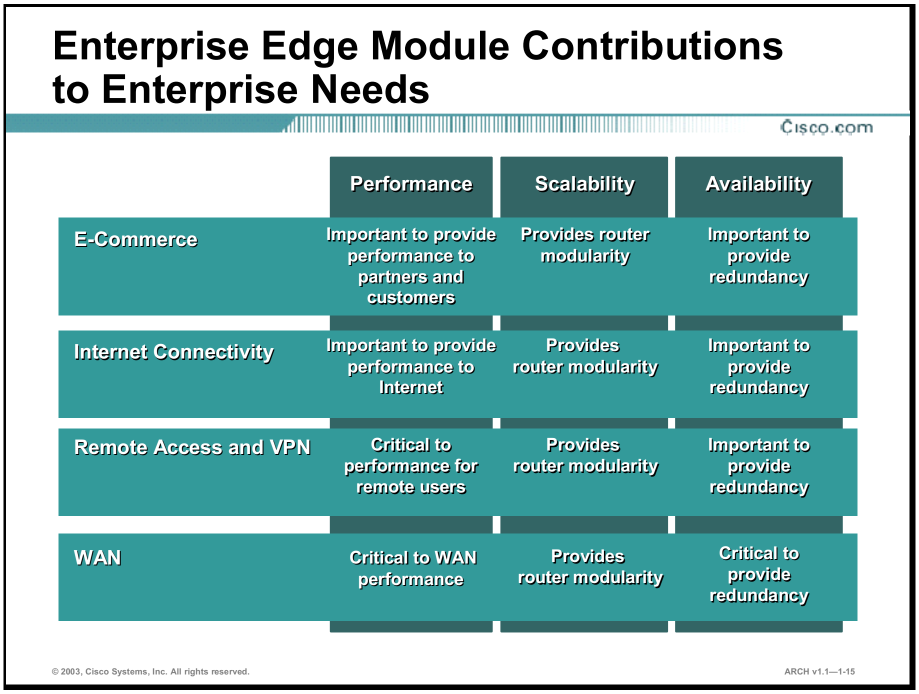

Enterprise Edge Module Contributions to Enterprise Needs

Cisco.com

	Performance	Scalability	Availability
E-Commerce	Important to provide performance to partners and customers	Provides router modularity	Important to provide redundancy
Internet Connectivity	Important to provide performance to Internet	Provides router modularity	Important to provide redundancy
Remote Access and VPN	Critical to performance for remote users	Provides router modularity	Important to provide redundancy
WAN	Critical to WAN performance	Provides router modularity	Critical to provide redundancy

ARCH v1.1—1-15

The figure describes how the Enterprise Edge functional area meets enterprise network needs for performance, scalability, and availability.

Service Provider Edge

The Service Provider Edge includes three modules: ISP, PSTN, and Frame Relay/ATM/PPP (FR/ATM/PPP). Each module has its own access, distribution, and core layers. This topic describes the components and functionality within the Service Provider Edge functional area, and explains how the Service Provider Edge functional area meets enterprise network requirements.

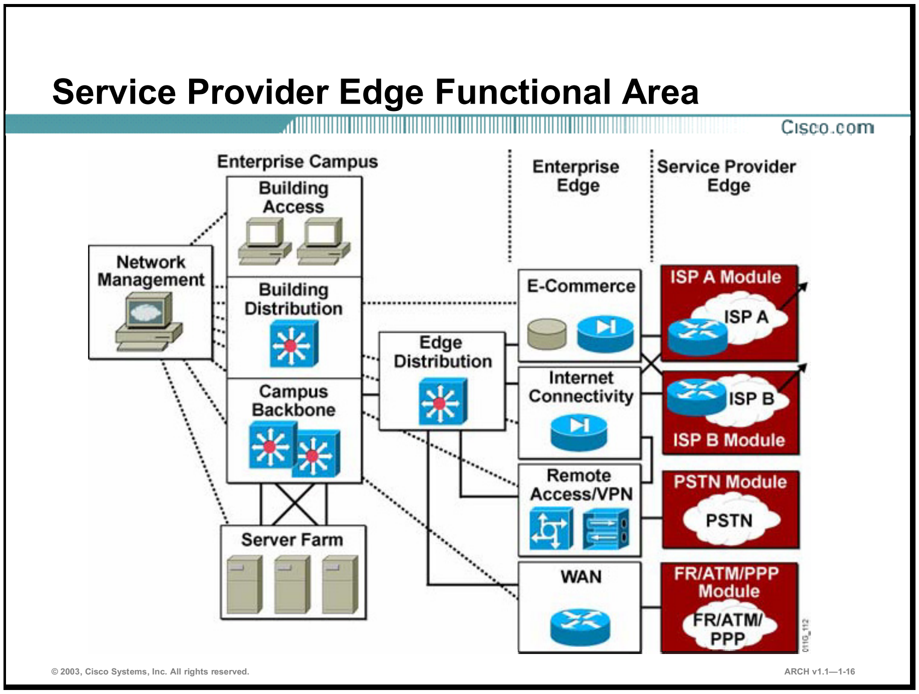

Service Provider Edge Functional Area

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-16

The functions provided by the Service Provider Edge modules are as follows:

- **ISP:** Enables enterprise connectivity to the Internet. This service is essential to enable Enterprise Edge services, such as E-commerce, Remote Access and VPN, and Internet Connectivity 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 WAN technologies.
- **PSTN:** Represents the dial-up infrastructure used to access the enterprise network using ISDN, analog, and wireless (cellular) technologies. Enterprises can also use the PSTN module to back up existing WAN links. Connections are established on demand and terminated when determined to be idle.
- **FR/ATM/PPP:** Includes all WAN technologies for permanent connectivity with remote locations. Frame Relay, ATM, and PPP are the most frequently used today. However, many technologies can fit into the same model.

The demarcation between the Enterprise Edge and the Service Provider Edge relate to ownership and control. The enterprise owns and controls the Enterprise Edge, while the service provider owns and controls the Service Provider Edge.

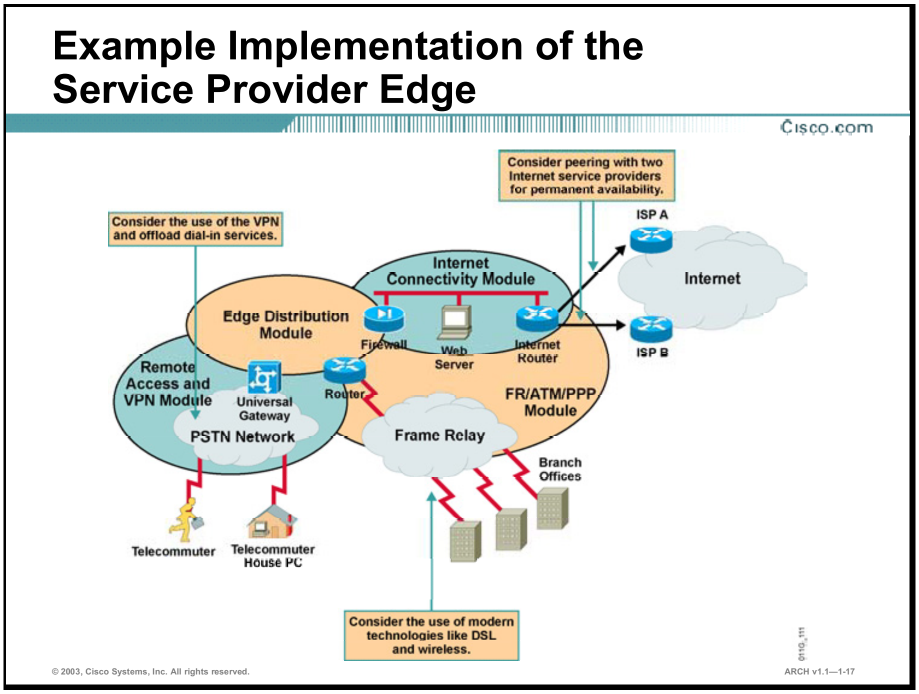

Example Implementation of the Service Provider Edge

Consider peering with two Internet service providers for permanent availability.

Consider the use of the VPN and offload dial-in services.

Consider the use of modern technologies like DSL and wireless.

ARCH v1.1—1-17

The figure shows how you can implement the Service Provider Edge functional area to meet enterprise networking needs. The E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules are independent.

Summary

This topic summarizes the key points discussed in this lesson.

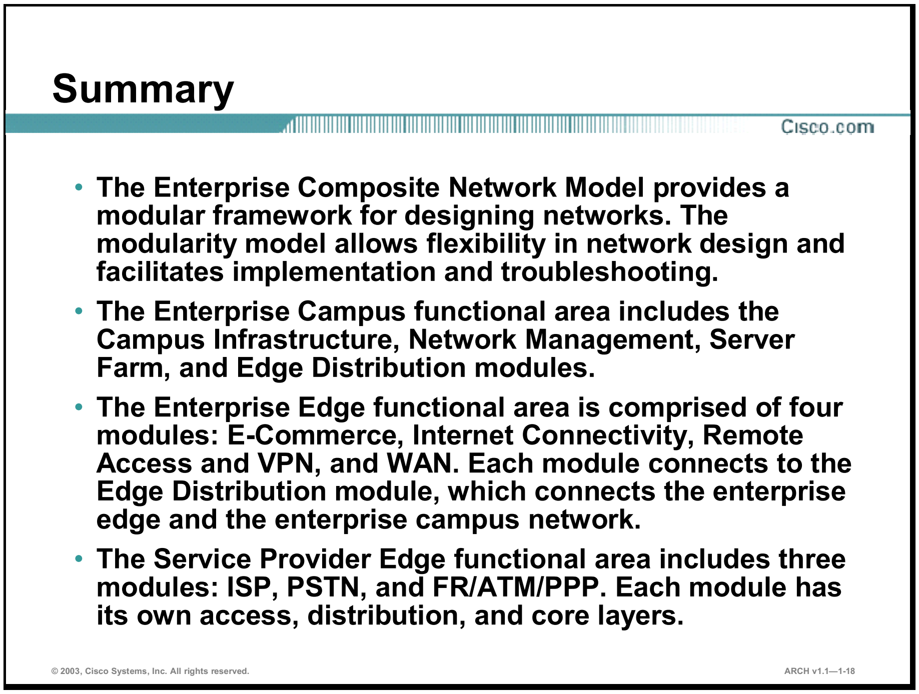

Summary

Cisco.com

- The Enterprise Composite Network Model provides a modular framework for designing networks. The modularity model allows flexibility in network design and facilitates implementation and troubleshooting.
- The Enterprise Campus functional area includes the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules.
- The Enterprise Edge functional area is comprised of four modules: E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN. Each module connects to the Edge Distribution module, which connects the enterprise edge and the enterprise campus network.
- The Service Provider Edge functional area includes three modules: ISP, PSTN, and FR/ATM/PPP. Each module has its own access, distribution, and core layers.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—1-18

References

For additional information, refer to this resource:

- *Cisco AVVID: Enabling E-Business* at http://www.cisco.com/warp/public/779/largeent/avvid/cisco_avvid.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 Enterprise Composite Network Model functional area with its description.

_____ 1. Enterprise Edge

_____ 2. Enterprise Campus

_____ 3. Service Provider Edge

A) aggregates connectivity from the various elements at the edge of the enterprise network

B) modules that enable communications with other networks using different WAN technologies and Internet service providers

C) contains the modules required to build a hierarchical, highly robust campus network that offers performance, scalability, and availability

Q2) Which two criteria does the Enterprise Composite Network Model meet that the Cisco hierarchical model does not meet? (Choose two.)

A) offers more network integrity in network design

B) supports the access and distribution layers in the Enterprise Edge

C) provides availability by allowing enterprises to add modules easily

D) increases network scalability and eases the design task by making each module discrete

Q3) Match each Enterprise Campus module to its definition.

_____ 1. Server Farm module

_____ 2. Edge Distribution module

_____ 3. Network Management module

_____ 4. Campus Infrastructure module

A) aggregates the connectivity from the various elements at the Enterprise Edge and routes the traffic to the campus core

B) contains internal e-mail and corporate servers providing application, file, print, e-mail, and DNS services to internal users

C) connects users within a campus, and includes the Building Access, Building Distribution, and Campus Backbone submodules

D) performs intrusion detection, system logging, and authentication, as well as network monitoring and general configuration management functions

Q4) Which three services does the Building Access submodule provide? (Choose three.)

A) routing

B) access control

C) network access

D) protocol filtering

E) broadcast suppression

F) aggregation of wiring closets

Q5) Match each Enterprise Edge submodule to its definition.

_____ 1. WAN

_____ 2. E-Commerce

_____ 3. Internet Connectivity

_____ 4. Remote Access and VPN

A) routes traffic between remote sites and the central site

B) terminates VPN traffic from remote users and remote sites

C) provides internal users with connectivity to Internet services

D) enables enterprises to successfully deploy e-commerce applications and take advantage of the powerful competitive opportunities provided by the Internet

Q6) Which Enterprise Edge module routes traffic between remote sites and the central site?

A) WAN

B) E-Commerce

C) Internet Connectivity

D) Remote Access and VPN

Q7) Match each Service Provider Edge module with its description.

_____ 1. PSTN

_____ 2. FR/ATM/PPP

_____ 3. Internet Service Provider

A) enables enterprise connectivity to the Internet

B) includes all WAN technologies for permanent connectivity with remote locations

C) represents the dial-up infrastructure used to access the enterprise network using ISDN, analog, and wireless technologies

Q8) Which three modules make up the Service Provider Edge functional area? (Choose three.)

A) PSTN

B) FR/ATM/PPP

C) Enterprise Edge

D) Campus Infrastructure

E) Internet Service Provider

Quiz Answer Key

Q1) 1-A, 2-C, 3-B

Relates to: Enterprise Composite Network Model

Q2) A, D

Relates to: Enterprise Composite Network Model

Q3) 1-B, 2-A, 3-D, 4-C

Relates to: Enterprise Campus

Q4) C, D, E

Relates to: Enterprise Campus

Q5) 1-A, 2-D, 3-C, 4-B

Relates to: Enterprise Edge

Q6) A

Relates to: Enterprise Edge

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

Relates to: Service Provider Edge

Q8) A, B, E

Relates to: Service Provider Edge

Module 2

Designing Enterprise Campus Networks

Overview

Enterprise sites, whether small or large, need a solid network infrastructure to support emerging solutions such as IP telephony, storage networking, broadband solutions, content networking, and the applications that surround them. The network foundation hosting these technologies for an emerging enterprise should be efficient, highly available, scalable, and manageable. The Cisco Architecture for Voice, Video and Integrated Data (AVVID) network infrastructure is designed to run a converged voice, video, and data network over IP with due consideration for quality of service, bandwidth, latency, and high performance demanded by network solutions.

This module provides design models for the Campus Backbone, Building Distribution, and Building Access submodules, and Server Farm and Edge Distribution modules of the Enterprise Composite Network Model.

Module Objectives

Upon completing this module, you will be able to design enterprise campus network infrastructures for effective functionality, performance, scalability, and availability, given specified enterprise network needs.

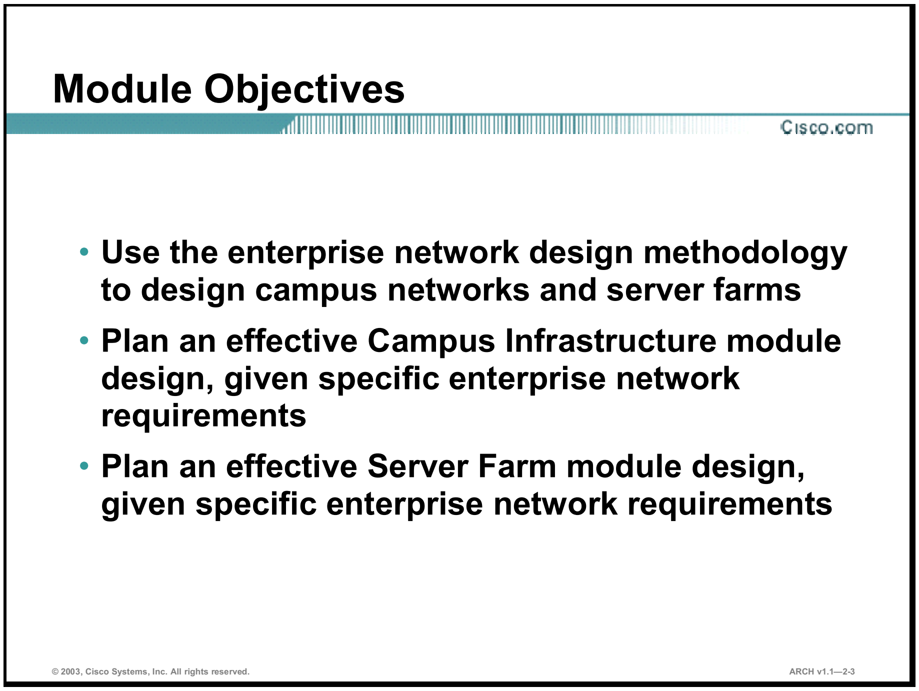

Module Objectives

Cisco.com

- Use the enterprise network design methodology to design campus networks and server farms
- Plan an effective Campus Infrastructure module design, given specific enterprise network requirements
- Plan an effective Server Farm module design, given specific enterprise network requirements

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—2-3

Module Outline

The outline lists the components of this module.

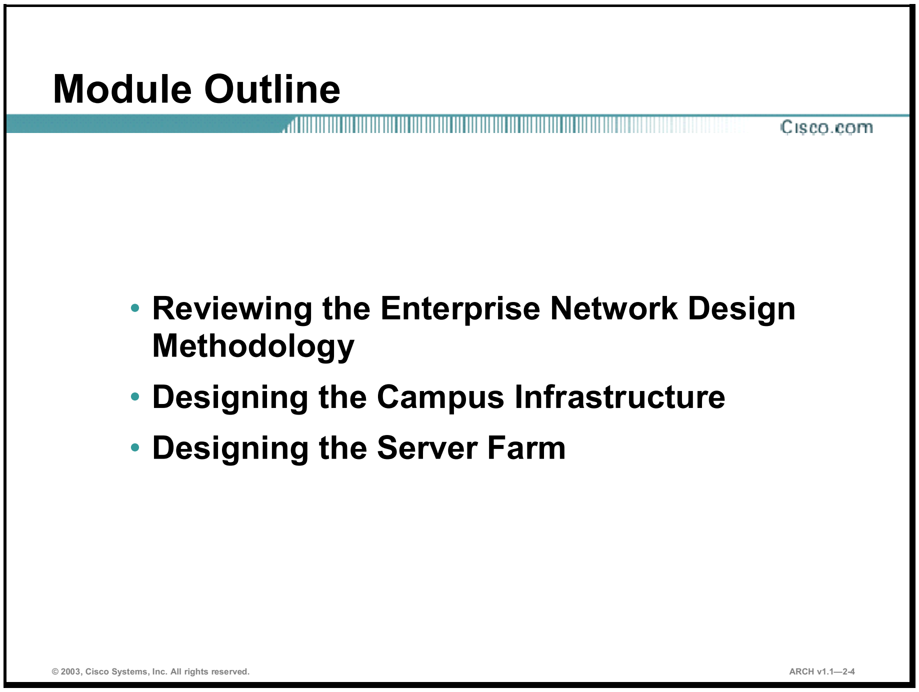

Module Outline

Cisco.com

- Reviewing the Enterprise Network Design Methodology
- Designing the Campus Infrastructure
- Designing the Server Farm

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—2-4

Reviewing the Enterprise Network Design Methodology

Overview

To facilitate effective network design, Cisco has developed a process that enables the network designer to assess requirements, design each module of the network, and determine the effectiveness of the design.

Relevance

The Enterprise Composite Network Model enables network designers to create a campus network made out of modular building blocks, which are scalable to meet evolving business needs. By deploying a step-by-step methodology, network designers can create an effective campus design that meets enterprise requirements for performance, scalability, and availability.

Objectives

Upon completing this lesson, you will be able to use the enterprise network design methodology to design campus networks and server farms. This includes being able to meet these objectives:

- Identify the modules and submodules that the network designers will design for the enterprise campus network
- Identify the performance, scalability, and availability design considerations for the Enterprise Campus functional area of the Enterprise Composite Network Model
- Describe a step-by-step methodology that network designers will use to design the Enterprise Campus functional area
- Analyze network traffic patterns typically found within the Enterprise Campus functional area

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

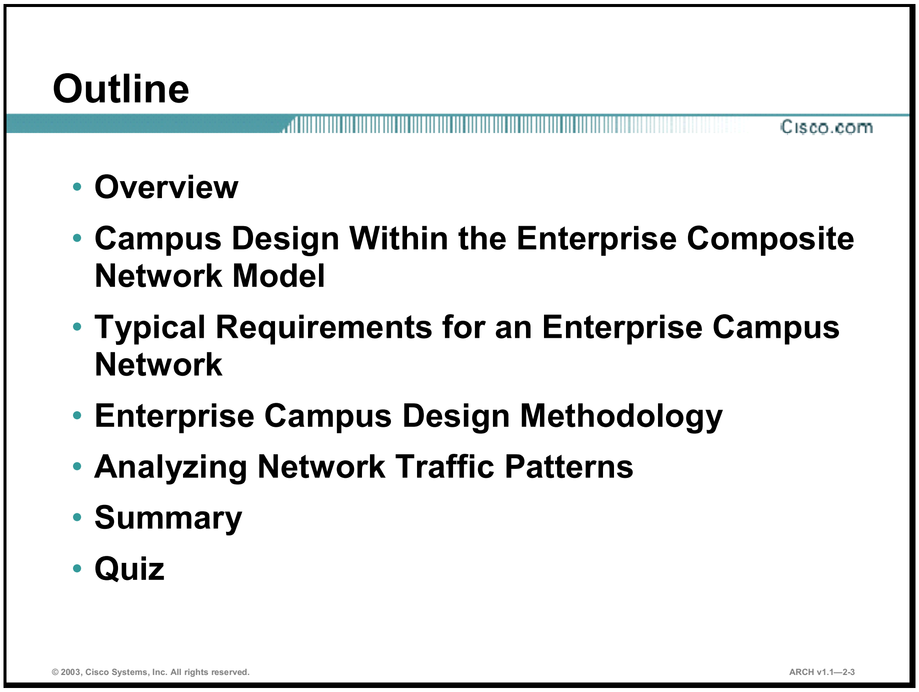

Outline

- Overview
- Campus Design Within the Enterprise Composite Network Model
- Typical Requirements for an Enterprise Campus Network
- Enterprise Campus Design Methodology
- Analyzing Network Traffic Patterns
- Summary
- Quiz

ARCH v1.1—2-3

Campus Design Within the Enterprise Composite Network Model

To design an enterprise campus network, you will design the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules. This topic identifies the modules and submodules that the network designer will design for the enterprise campus network.

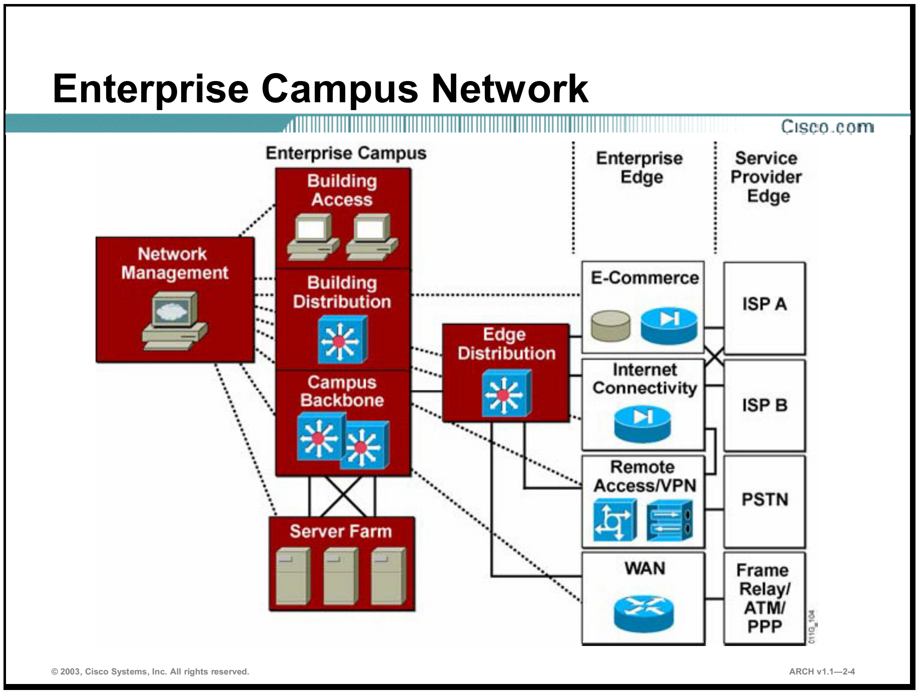

Enterprise Campus Network

ARCH v1.1—2-4

The Enterprise Campus functional area includes these four major modules:

- Campus Infrastructure
- Network Management
- Server Farm
- Edge Distribution

The Campus Infrastructure module connects users within a campus with the Server Farm and Edge Distribution modules. The Campus Infrastructure module is composed of one or more floors or buildings connected to the Campus Backbone submodule. Each building contains a Building Access and Building Distribution submodule.

The Campus Infrastructure module includes these submodules:

- **Building Access:** Contains end-user workstations, IP Phones, and data link layer access switches that connect devices to the Building Distribution submodule. Building Access performs important services such as broadcast suppression, protocol filtering, network access, and quality of service marking.

- **Building Distribution:** Provides aggregation of wiring closets, often using multilayer switching. Building Distribution performs routing, quality of service, and access control. Requests for data flow into the Building Distribution switches and then to the Campus Backbone. The model provides fast failure recovery, since each Building Distribution switch maintains two connections to the core, thus two equal cost paths in the routing table to every destination network. When one connection to the campus core fails, all routes immediately switch over to the remaining path about one second after the link failure is detected.
- **Campus Backbone:** 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. This submodule uses data link layer switches and multilayer switches for high throughput functions with added routing, quality of service, and security features.

Typical Requirements for an Enterprise Campus Network

An enterprise campus network must meet requirements for functionality, performance, scalability, availability, manageability, and cost-effectiveness. This topic identifies the design considerations for each layer of an enterprise campus network based on the Cisco Enterprise Composite Network Model.

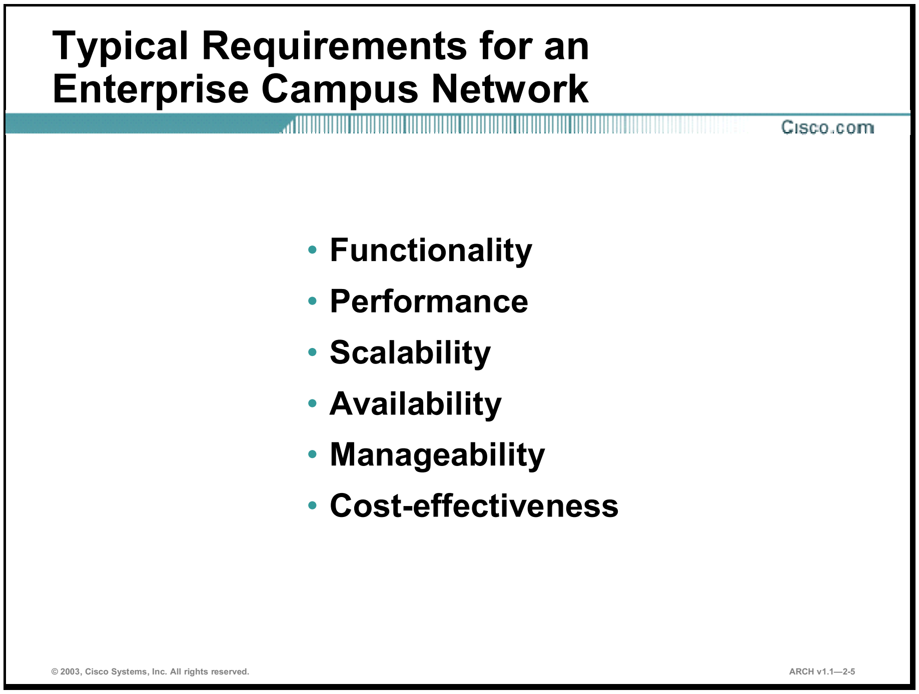

Typical Requirements for an Enterprise Campus Network

Cisco.com

- Functionality
- Performance
- Scalability
- Availability
- Manageability
- Cost-effectiveness

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-5

An enterprise campus network, as a whole, must meet these requirements:

- **Functionality:** The enterprise network must support the applications and data flows required, within the required time frames. Typical enterprise-wide applications include online transaction processing (OLTP) systems, decision support systems, e-mail, information sharing, as well as many others. Applications and data may require special peak-time processing, or they may require steady processing throughout a day.
- **Performance:** Performance includes three primary metrics: responsiveness, throughput (volume), and utilization. Each campus network will be measured in terms of how well it meets all three performance metrics.
- **Scalability:** Campus networks must provide scalability for future growth in the number of users and in the amount of data and applications that the network must support.
- **Availability:** Nearly 100 percent availability is required for most enterprise data networks. Networks providing converged services and solutions, and those providing support for critical applications, may be required to meet a standard of availability approaching 99.999 percent (“five nines”).
- **Manageability:** An enterprise campus network must be manageable across the entire infrastructure.

- **Cost-effectiveness:** Cost-effectiveness is a key concern for most enterprises, given limited budgets. The network designer's goal is to design the network for maximum effectiveness given affordability limitations. Design is often a matter of compromise. An important part of the design process is to discuss compromises and make cost-effective decisions.

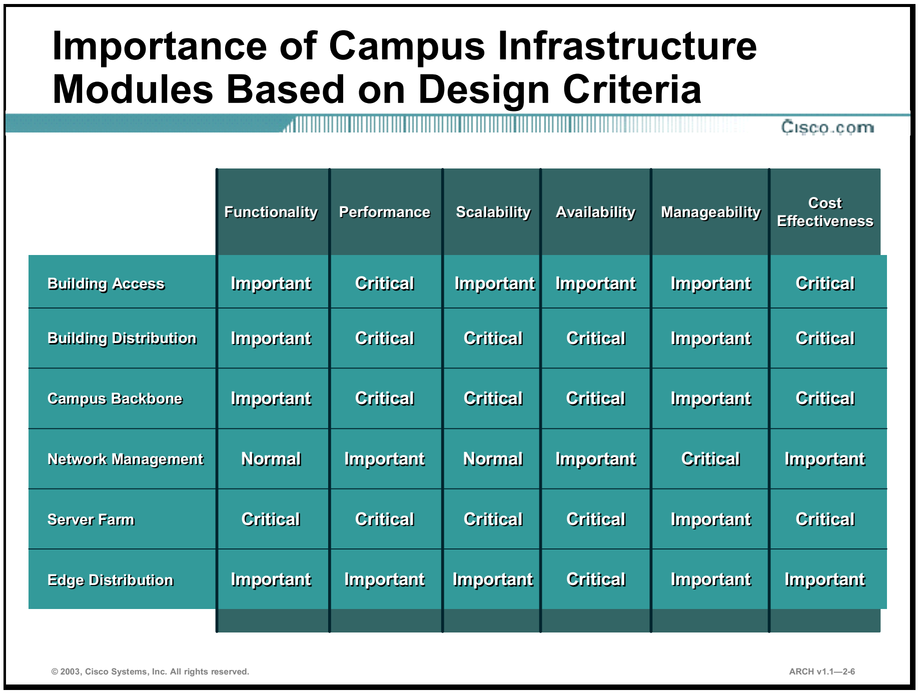

Importance of Campus Infrastructure Modules Based on Design Criteria

	Functionality	Performance	Scalability	Availability	Manageability	Cost Effectiveness
Building Access	Important	Critical	Important	Important	Important	Critical
Building Distribution	Important	Critical	Critical	Critical	Important	Critical
Campus Backbone	Important	Critical	Critical	Critical	Important	Critical
Network Management	Normal	Important	Normal	Important	Critical	Important
Server Farm	Critical	Critical	Critical	Critical	Important	Critical
Edge Distribution	Important	Important	Important	Critical	Important	Important

ARCH v1.1—2-6

The figure identifies the relative importance of needs within the various modules and submodules that comprise the Campus Infrastructure module. Each need is ranked in terms of its relative importance in the campus network, where Critical is highest in relative importance, followed by Important, and Normal.

Enterprise Campus Design Methodology

Cisco's network designers have developed a simple, seven-step process to design an enterprise campus network. This topic describes this step-by-step methodology that network designers can use to design the enterprise Campus Infrastructure and Server Farm modules.

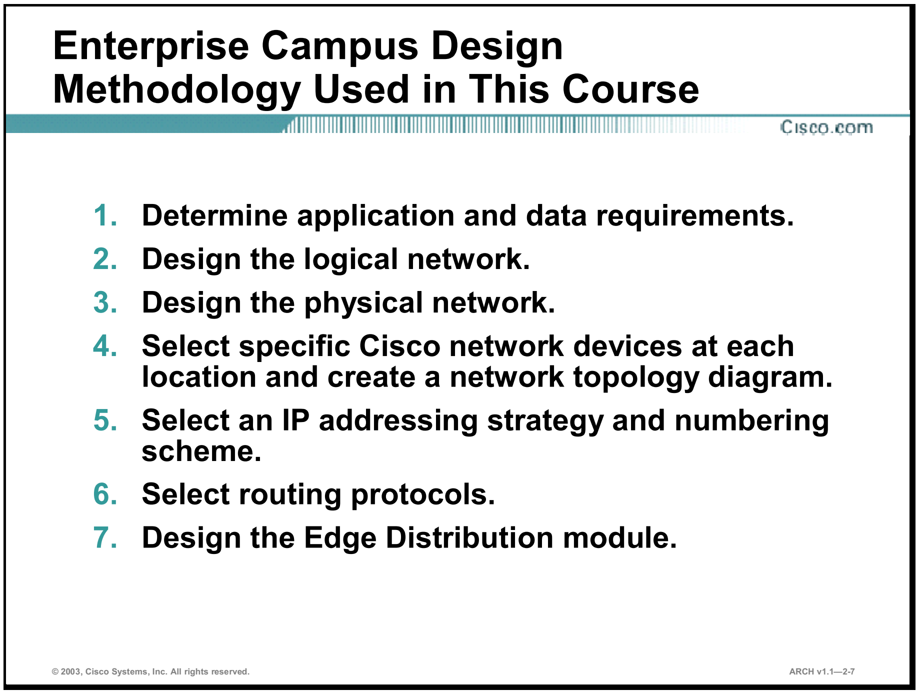

Enterprise Campus Design Methodology Used in This Course

Cisco.com

1. Determine application and data requirements.
2. Design the logical network.
3. Design the physical network.
4. Select specific Cisco network devices at each location and create a network topology diagram.
5. Select an IP addressing strategy and numbering scheme.
6. Select routing protocols.
7. Design the Edge Distribution module.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—2-7

To design an enterprise campus network for performance, scalability, and availability, you will complete a series of steps. The table describes the campus design methodology used throughout this course.

Activity	Description	Notes and Comments
1.	Determine application and data requirements for each campus location on the enterprise network.	Before beginning any network design, you must determine the enterprise application and data requirements.
2.	Design the logical network.	To design the logical network, you will identify the logical networks, usually created with VLANs or as separate networks.
3.	Design the physical network.	To design the physical network, you will identify these components: ■ Transmission media ■ Data link layer technology ■ Data link layer switching and multilayer switching strategy ■ STP[1] implementation ■ Method of connecting switches using trunks

Activity	Description	Notes and Comments
4.	Select specific Cisco network devices at each location and create a network topology diagram.	Based on the specific requirements at each location, you will: ■ Select specific Cisco network devices that meet specified criteria. ■ Select the hardware options that meet specified criteria. ■ Select the software options that meet specified criteria.
5.	Select an IP addressing strategy and numbering scheme.	You will determine if logical networks are single networks, subnetworks, or part of the larger network. You will determine the numbering scheme for each logical or physical network, and when to use route summarization.
6.	Select routing protocols.	You will select a routing protocol that meets the need for performance, scalability, and availability.
7.	Design the Edge Distribution module.	The Edge Distribution module provides connectivity between the core layer and the WAN modules. It generally consists of a multilayer switch.

[1] STP = Spanning Tree Protocol

Analyzing Network Traffic Patterns

Before designing the actual network, you should analyze the network traffic patterns for each application and location on the network. You will use the data to design the logical and physical network. This topic shows you how to analyze network traffic patterns typically found in enterprise campus networks.

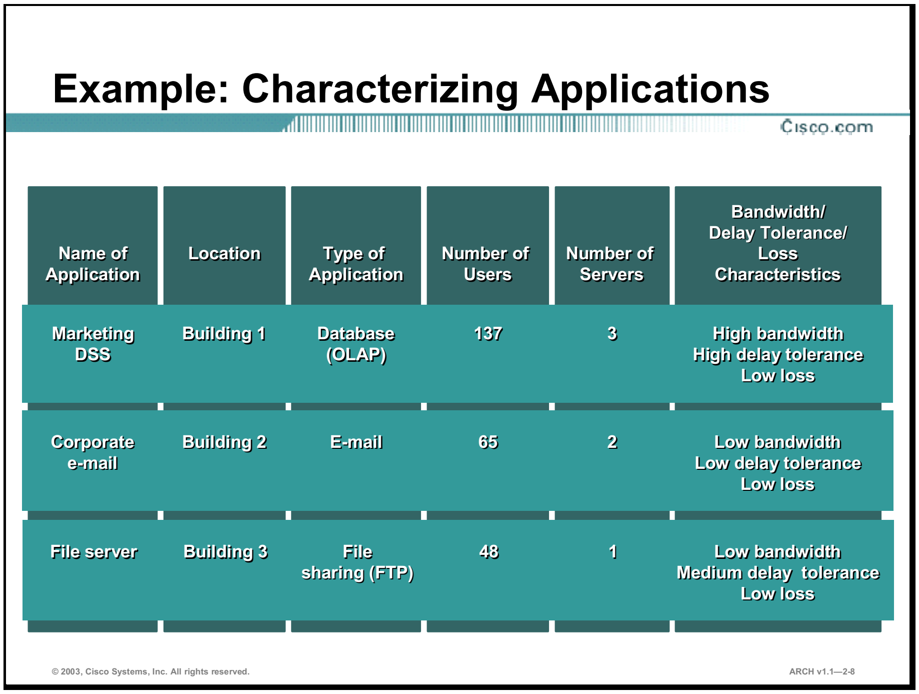

Example: Characterizing Applications

Name of Application	Location	Type of Application	Number of Users	Number of Servers	Bandwidth/ Delay Tolerance/ Loss Characteristics
Marketing DSS	Building 1	Database (OLAP)	137	3	High bandwidth High delay tolerance Low loss
Corporate e-mail	Building 2	E-mail	65	2	Low bandwidth Low delay tolerance Low loss
File server	Building 3	File sharing (FTP)	48	1	Low bandwidth Medium delay tolerance Low loss

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-8

You characterize the applications at each campus location on the network. The information you gather will help you determine the performance, scalability, and requirements for each location and network segment.

Use a table to characterize the applications at each network campus location, filling in the fields as indicated. The figure contains an example application table.

Name of Application	Building or Location	Type of Application	Number of Users	Number of Servers	Bandwidth/ Delay Tolerance/ Loss Characteristics

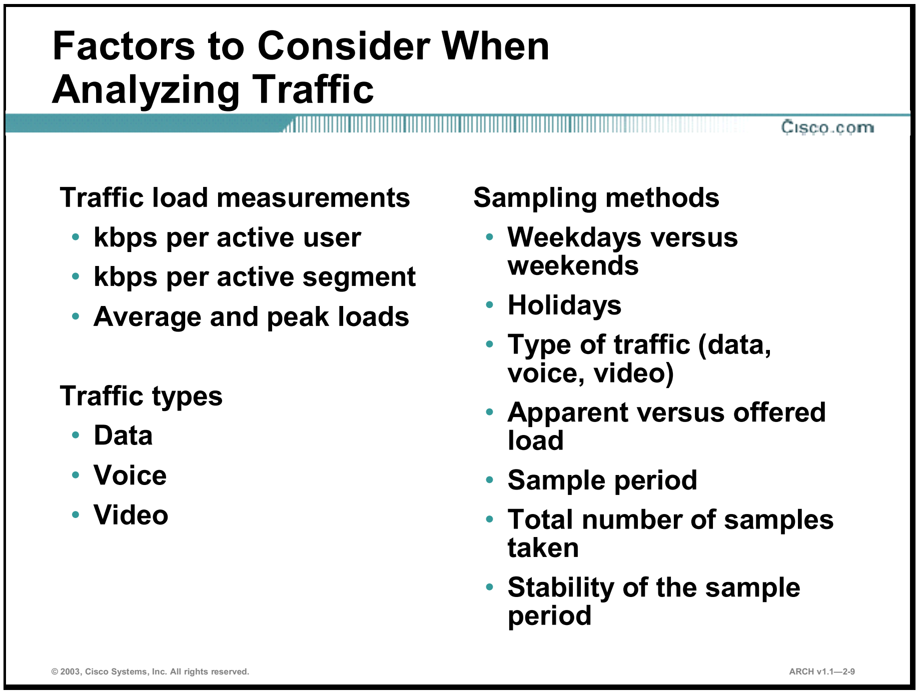

Factors to Consider When Analyzing Traffic

Cisco.com

Traffic load measurements

- kbps per active user
- kbps per active segment
- Average and peak loads

Traffic types

- Data
- Voice
- Video

Sampling methods

- Weekdays versus weekends
- Holidays
- Type of traffic (data, voice, video)
- Apparent versus offered load
- Sample period
- Total number of samples taken
- Stability of the sample period

ARCH v1.1—2-9

Network designers need a way to properly size network capacity, especially as networks grow. Traffic theory enables network designers to make assumptions about their networks based on past experience.

Traffic is defined as either the amount of data or the number of messages over a circuit during a given period of time. Traffic engineering addresses service issues by enabling you to define a grade of service or blocking factor. A properly engineered network has low blocking and high circuit utilization, which means that service is increased and your costs are reduced.

There are many different factors that you need to take into account when analyzing traffic. The most important factors are:

- **Traffic load measurement:** To measure the traffic load, you will gather statistics based on past experience in the enterprise. Specifically, you will determine the average number of kbps per user, and the average number of kbps per network segment.
- **Traffic types:** Traffic types may include data, voice, and video. The different types of data may include spreadsheets, word processing, or HTML documents, among others.
- **Sampling methods:** Probability theory suggests that to accurately assess network traffic, you need to consider at least 30 of the busiest hours of a network in the sampling period. Although this is a good starting point, other variables can skew the accuracy of this sample. You cannot take the top 30 out of 32 samples and expect that sampling to be an accurate picture of the network. To get the most accurate results, you need to take as many samples of the offered load as possible. Also, if you take samples throughout the year, your results can be skewed as the year-to-year traffic load increases or decreases.

Summary

This topic summarizes the key points discussed in this lesson.

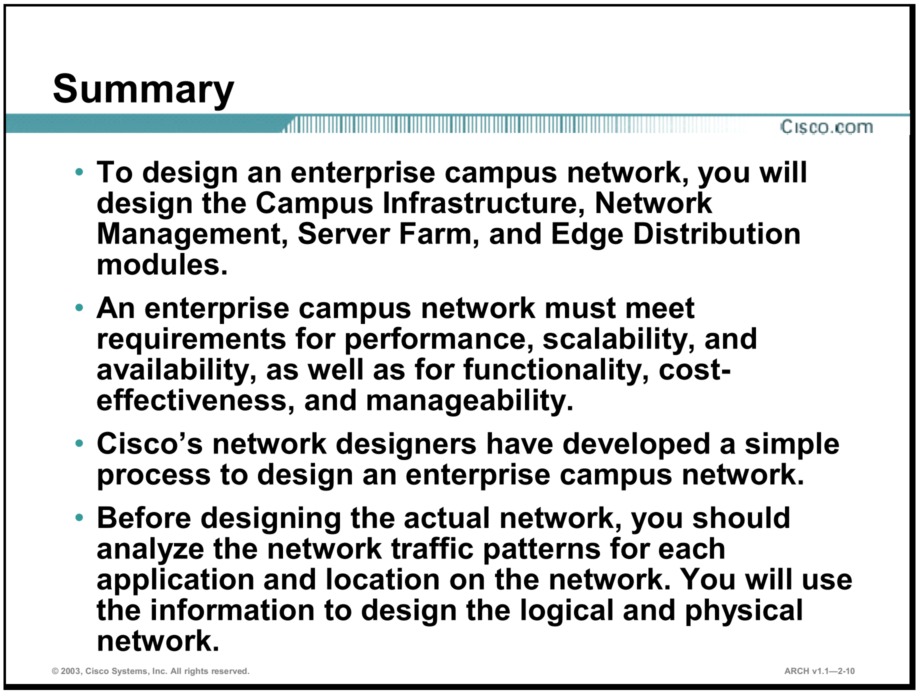

Summary

- To design an enterprise campus network, you will design the Campus Infrastructure, Network Management, Server Farm, and Edge Distribution modules.
- An enterprise campus network must meet requirements for performance, scalability, and availability, as well as for functionality, cost-effectiveness, and manageability.
- Cisco's network designers have developed a simple process to design an enterprise campus network.
- Before designing the actual network, you should analyze the network traffic patterns for each application and location on the network. You will use the information to design the logical and physical network.

ARCH v1.1—2-10

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 submodules of the Campus Infrastructure will you design? (Choose three.)

A) Building Access

B) Edge Distribution

C) Campus Backbone

D) Building Distribution

E) Network Management

Q2) To provide for future growth in the number of users and in the amount of data and applications that the network must support, campus networks must provide _____.

A) scalability

B) availability

C) performance

D) cost-effectiveness

Q3) What information do you need before you can design the logical network?

A) routing protocols

B) application and data requirements

C) Cisco network devices at each location

D) IP addressing strategy and numbering scheme

Q4) Which three types of information do you need to characterize applications on the network? (Choose three.)

A) name of users

B) type of application

C) number of subapplications

D) number of users and servers

E) bandwidth, delay, and loss characteristics

Quiz Answer Key

Q1) A, C, D

Relates to: Campus Design within the Enterprise Composite Network Model

Q2) A

Relates to: Typical Requirements for an Enterprise Campus Network

Q3) B

Relates to: Enterprise Campus Design Methodology

Q4) B, D, E

Relates to: Analyzing Network Traffic Patterns

Designing the Campus Infrastructure

Overview

The availability of hardware-accelerated multilayer switches that provide intelligent network services allows network designers to achieve data rates that were previously possible only on data link layer switches at the network layer and above. Cisco recommends use of data link layer switches at the network edge to avoid the complexity of extending multilayer switching in the wiring closet. Typically, Building Access devices will terminate on a unique subnet at the Building Distribution submodule.

This lesson describes how to design the Campus Infrastructure module for effective performance, scalability, and availability.

Relevance

Multilayer switching in the Campus Backbone and Building Distribution submodules offers speed and manageability advantages, while traditional data link layer workgroup switching at the Building Access submodule reduces complexity in the wiring closet. Design criteria for selecting data link layer switches and multilayer switches include functionality, performance, and cost requirements.

Objectives

Upon completing this lesson, you will be able to plan an effective Campus Infrastructure module design, given specific enterprise network requirements. This includes being able to meet these objectives:

- Select logical network segments and the segmentation method for the Campus Infrastructure module, given specific internetwork requirements
- Select transmission media, data link protocols, and spanning-tree strategy for the Building Access, Building Distribution, and Campus Backbone submodules of the Campus Infrastructure module, given specific internetwork requirements

- Select data link layer switching and multilayer switching solutions for the Building Access, Building Distribution, and Campus Backbone submodules of the Campus Infrastructure module, given specific internetwork requirements
- Select hardware, hardware options, and software options for a campus network infrastructure, given specific internetwork requirements
- Identify an IP addressing strategy for the campus network, given specific internetwork requirements
- Select routing protocols for the campus network that meet performance, scalability, and availability requirements
- Propose small, medium, and large campus network designs, given specific internetwork requirements

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

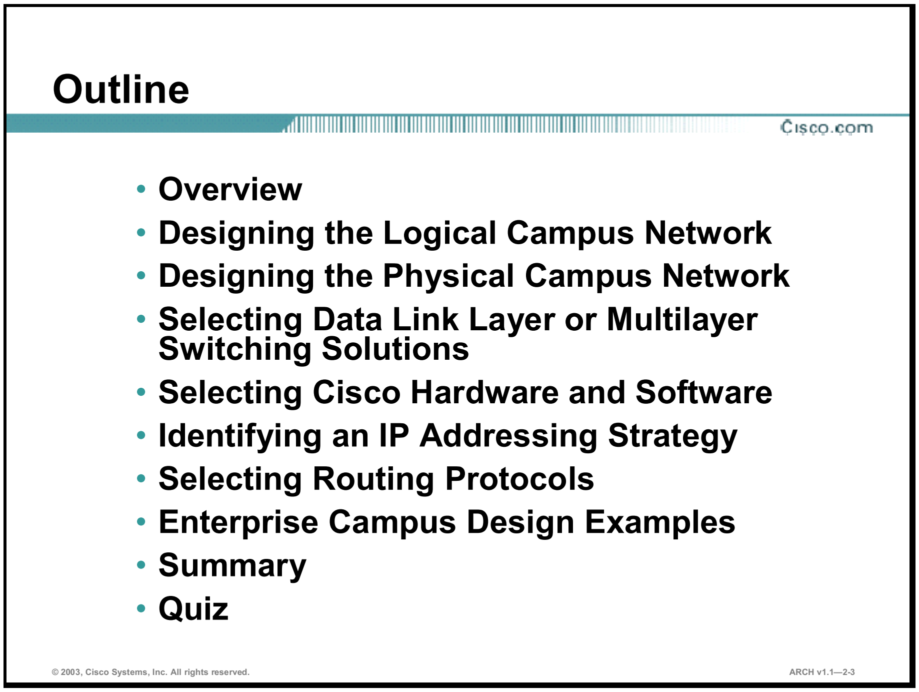

Outline

- **Overview**
- **Designing the Logical Campus Network**
- **Designing the Physical Campus Network**
- **Selecting Data Link Layer or Multilayer Switching Solutions**
- **Selecting Cisco Hardware and Software**
- **Identifying an IP Addressing Strategy**
- **Selecting Routing Protocols**
- **Enterprise Campus Design Examples**
- **Summary**
- **Quiz**

ARCH v1.1—2-3

Designing the Logical Campus Network

Once you have identified the application and data needs for an enterprise, you are ready to begin the logical network design by selecting logical network segments and choosing the methods you will implement to create them. In this topic, you will learn to select logical network segments and a segmentation method for a campus network infrastructure, given specific internetwork requirements.

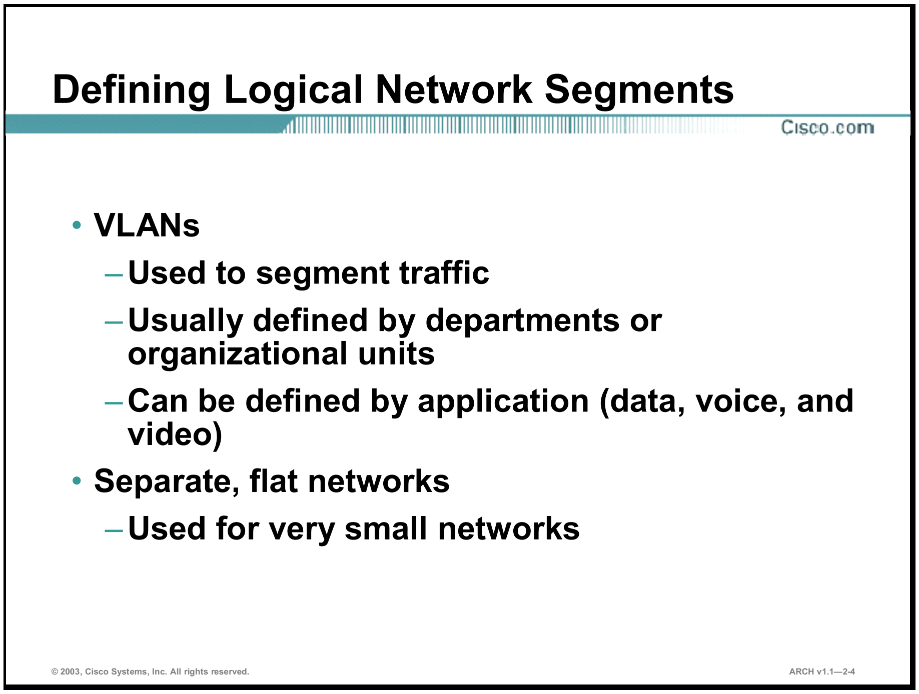

Defining Logical Network Segments

Cisco.com

- VLANs
  - Used to segment traffic
  - Usually defined by departments or organizational units
  - Can be defined by application (data, voice, and video)
- Separate, flat networks
  - Used for very small networks

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-4

The key to good network design is how you place clients in relation to resources. Ideally, client computers should be placed on the same logical network as the local resources they access most often, such as departmental servers, printers, and other local resources. You can define a logical connection in your network software so that users in one area of a building can be in the same logical network segment as related users and printers located at the opposite end of that building. This simple task minimizes the network-layer processing load on the Building Distribution and Campus Backbone submodules, which carry traffic between segments.

The two methods used to create logical network segments are:

- **VLANs:** A VLAN is an extended data link layer switched domain. If several VLANs coexist across a set of data link layer switches, each individual VLAN is an independent failure domain, broadcast domain, and spanning-tree domain.
- **Separate networks:** You can implement one or more bridged physical segments as a separate network.

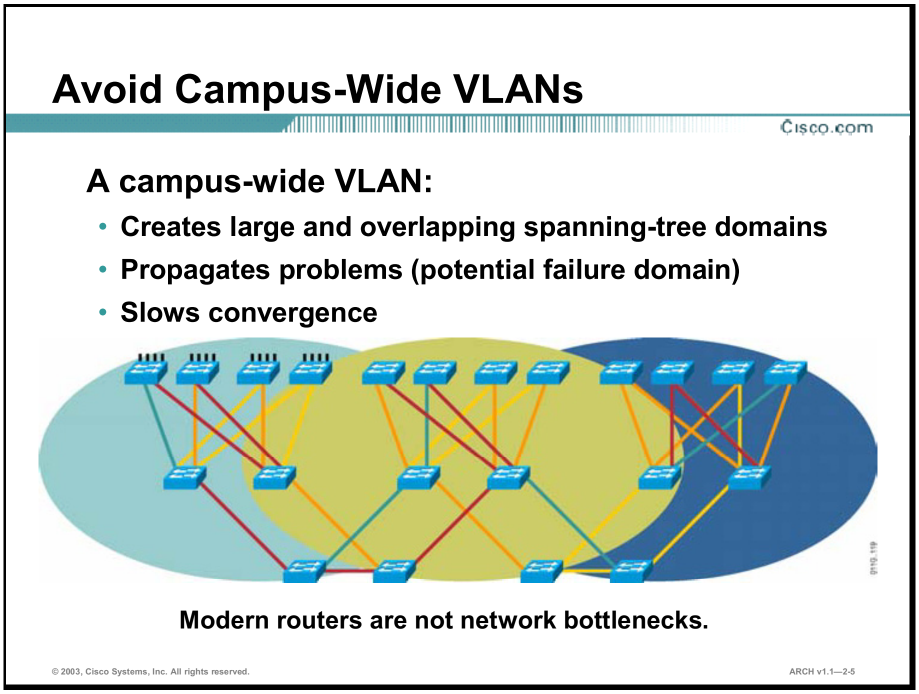

Avoid Campus-Wide VLANs

A campus-wide VLAN:

- **Creates large and overlapping spanning-tree domains**
- **Propagates problems (potential failure domain)**
- **Slows convergence**

Modern routers are not network bottlenecks.

Although you can use VLANs to segment the campus network logically, deploying pervasive VLANs throughout the campus introduces complexity and reduces the deterministic behavior of the network. Avoiding loops and restricting a set of unique VLANs to a single data link layer switch in one wiring closet minimizes the complexity, thus increasing manageability. Smaller spanning-tree domains can also converge faster in the event of failure. Weigh simplicity in the network design against the potential organizational advantages of logical networks that span site or campus, multiple buildings, or even Building Distribution submodules within a building. There is a constant balance between the logical (organizational) and the physical (segments and devices) in network design.

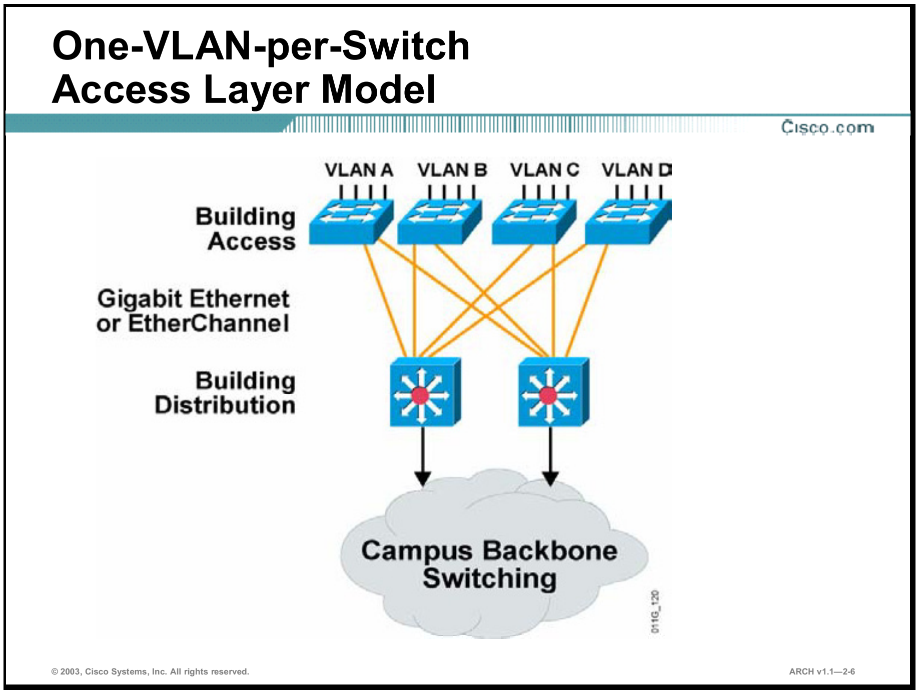

One-VLAN-per-Switch Access Layer Model

Cisco.com

ARCH v1.1—2-6

The Building Access submodule is the first point of entry into the network and essentially links end users to the remainder of the network. The figure shows a one-VLAN-per-switch Building Access submodule model.

In the figure, each uplink is part of only one VLAN, and therefore, no trunks are created between the wiring-closet switches and the Building Distribution submodule switches. Therefore, transporting traffic from one switch to another requires going through a multilayer device.

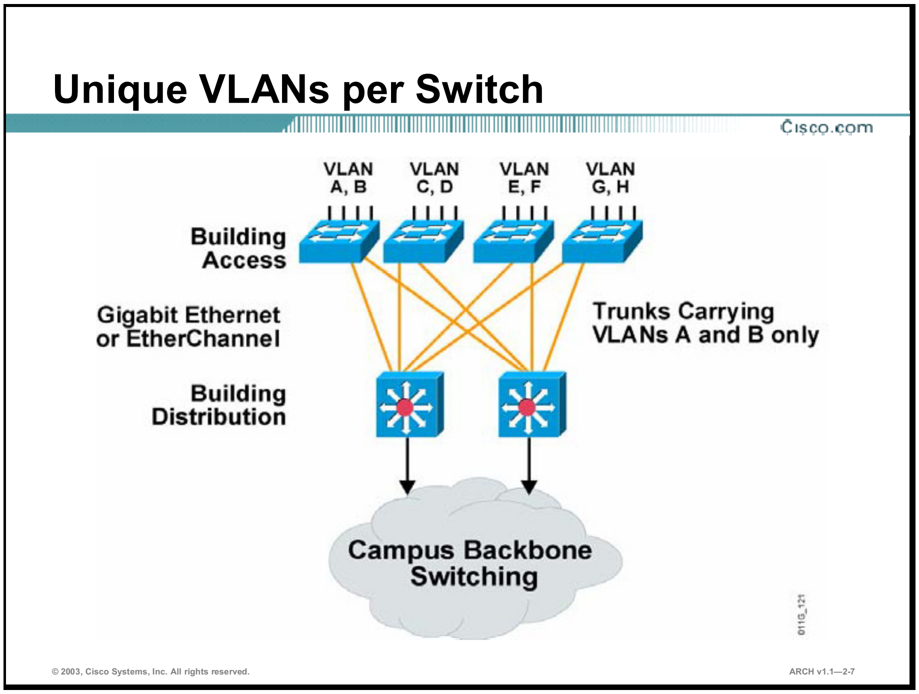

Unique VLANs per Switch

Cisco.com

VLAN A, B
VLAN C, D
VLAN E, F
VLAN G, H
Building Access
Gigabit Ethernet or EtherChannel
Trunks Carrying VLANs A and B only
Building Distribution
Campus Backbone Switching
011G_121

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-7

The figure represents a set of unique VLANs per switch. Because more than one VLAN exists per wiring-closet switch, trunks are defined and unnecessary VLANs are removed. This design still requires a multilayer device to transport traffic between VLANs.

Trunking is a way to carry traffic from several VLANs over a point-to-point link between the two devices. Two ways to implement Ethernet trunking are:

- Inter-Switch Link (ISL) (Cisco proprietary protocol)
- 802.1Q (IEEE standard)

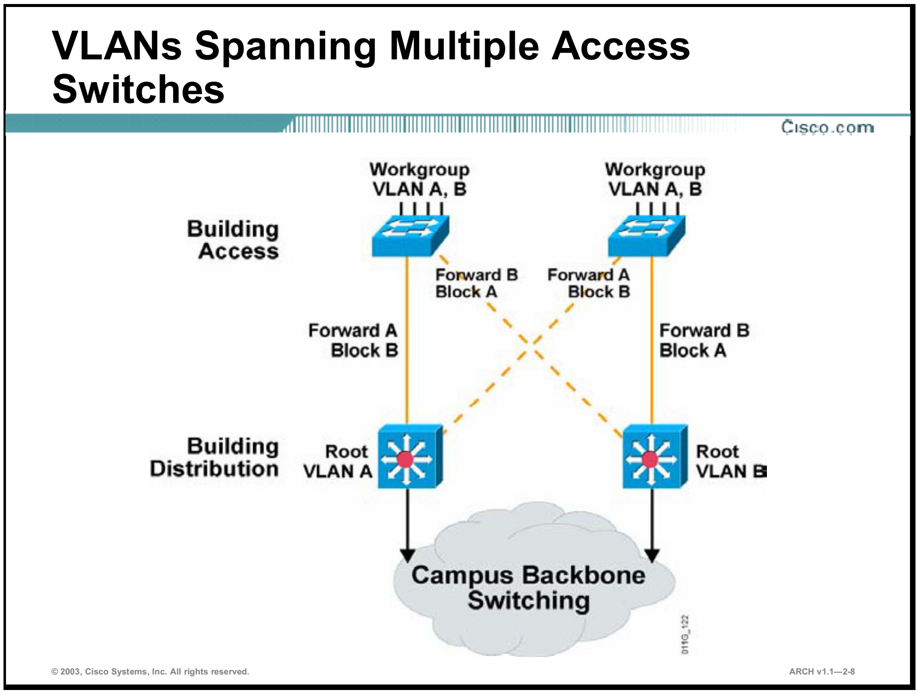

VLANs Spanning Multiple Access Switches

ARCH v1.1—2-8

The figure describes a more general concept, which can also support distributed workgroup servers attached to the Building Distribution submodule switches. These servers reside on VLAN A or B. If one or more VLANs span several Building Access submodule switches, install a trunk carrying those VLANs between the two Building Distribution submodule switches. Failing to do so may result in suboptimal traffic paths or introduce routing black holes in the network. A routing black hole occurs when a router advertises reachability to a network to its peers even though it cannot properly route the traffic to that network.

Designing the Physical Campus Network

The physical network design identifies the Layer 1 (physical) and Layer 2 (data link and spanning-tree) implementations for the enterprise network. This topic provides guidelines to help you select transmission media, data link protocols, and spanning-tree strategy for the Building Access, Building Distribution, and Campus Backbone submodules of the Campus Infrastructure module, given specific internetwork requirements.

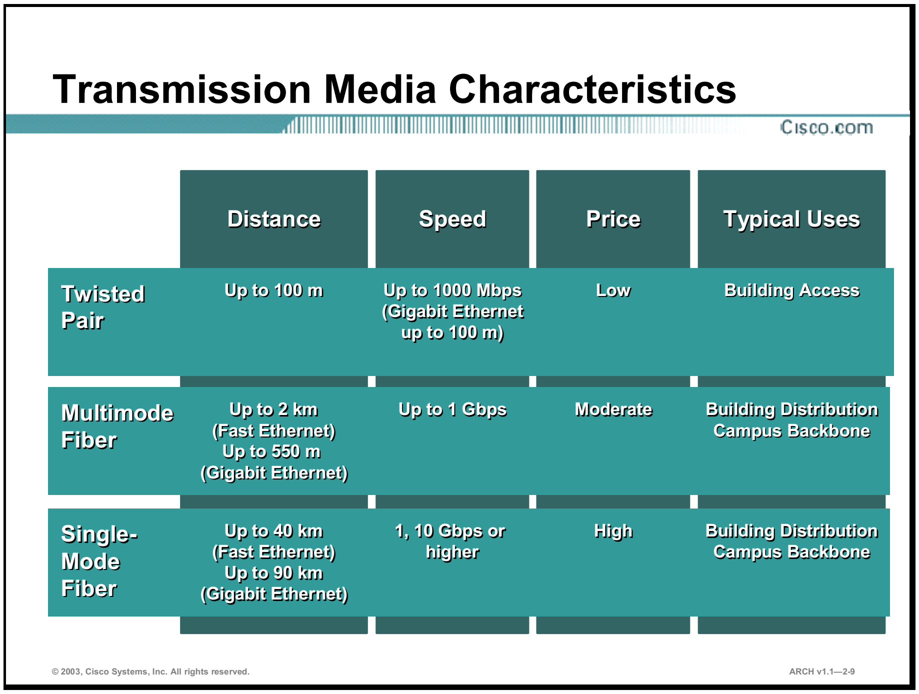

Transmission Media Characteristics

	Distance	Speed	Price	Typical Uses
Twisted Pair	Up to 100 m	Up to 1000 Mbps (Gigabit Ethernet up to 100 m)	Low	Building Access
Multimode Fiber	Up to 2 km (Fast Ethernet) Up to 550 m (Gigabit Ethernet)	Up to 1 Gbps	Moderate	Building Distribution Campus Backbone
Single-Mode Fiber	Up to 40 km (Fast Ethernet) Up to 90 km (Gigabit Ethernet)	1, 10 Gbps or higher	High	Building Distribution Campus Backbone

ARCH v1.1—2-9

The primary types of transmission media that are used in an enterprise include:

- **Twisted pair:** The two types of twisted pair cabling are unshielded twisted-pair (UTP) and shielded twisted-pair (STP). UTP is widely used to interconnect workstations, servers, and other devices from their network interface card (NIC) to a network device. STP is similar to UTP, but cables are wrapped in foil to protect them from external electromagnetic influences.
- **Multimode fiber:** Multimode fiber uses a LED as the light source. The low power output and modal dispersion limits the distance at which it can be distinguished reliably.
- **Single-mode fiber:** Single-mode optical fiber uses lasers as the light source and is designed for the transmission of a single wave or mode of light as a carrier. The single ray of light can be distinguished more reliably at longer distances compared to multimode fiber.

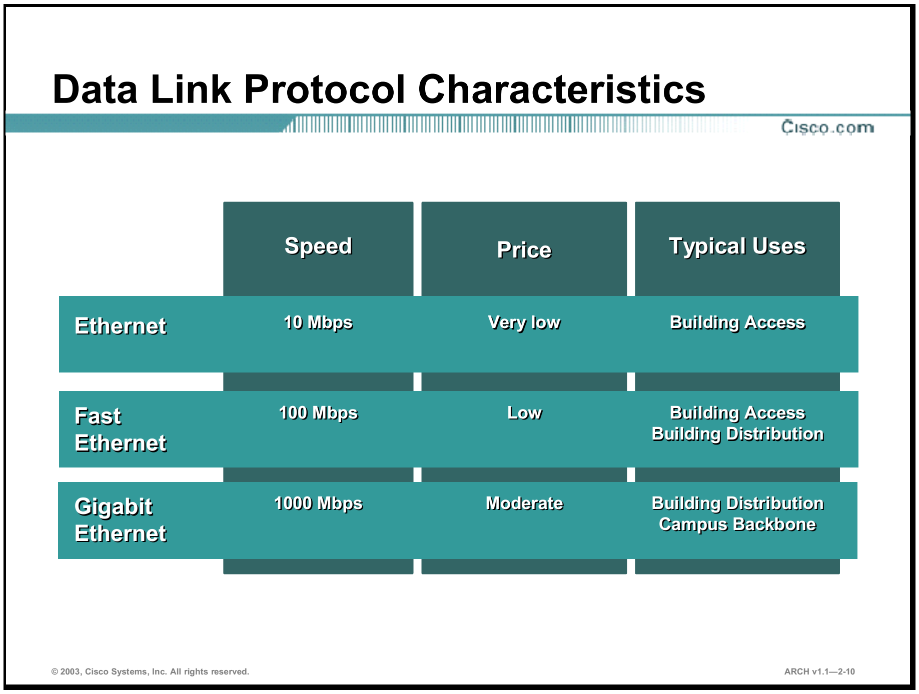

Data Link Protocol Characteristics

Cisco.com

	Speed	Price	Typical Uses
Ethernet	10 Mbps	Very low	Building Access
Fast Ethernet	100 Mbps	Low	Building Access Building Distribution
Gigabit Ethernet	1000 Mbps	Moderate	Building Distribution Campus Backbone

ARCH v1.1—2-10

The figure describes the speed, price, and typical uses for the key data link protocols deployed on campus networks. These three data link protocols are commonly used to build the campus network today:

- **10-Mbps Ethernet:** The slowest of the three technologies, it is considered a legacy media because the prices of the other media are low enough that it is just as economical to put in a 100-Mbps Fast Ethernet network as it is to put in a 10-Mbps Ethernet network.
- **100-Mbps Ethernet (Fast Ethernet):** Current media of choice because of its low cost and ability to service most user requirements and many server requirements. It is relatively inexpensive to implement.
- **1000-Mbps Ethernet (Gigabit Ethernet):** The gigabit network is a little more expensive to implement but gives a ten-fold increase in bandwidth, so is generally used between access and distribution and distribution and core.

EtherChannel technology provides incremental trunk speeds between Fast Ethernet and Gigabit Ethernet, or at speeds greater than Gigabit Ethernet. EtherChannel combines multiple Fast Ethernet links up to 800 Mbps or Gigabit Ethernet up to 8 Gbps. EtherChannel provides fault-tolerant, high-speed links between switches, routers, and servers. Without EtherChannel, connectivity options are limited to the specific line rates of the interface.

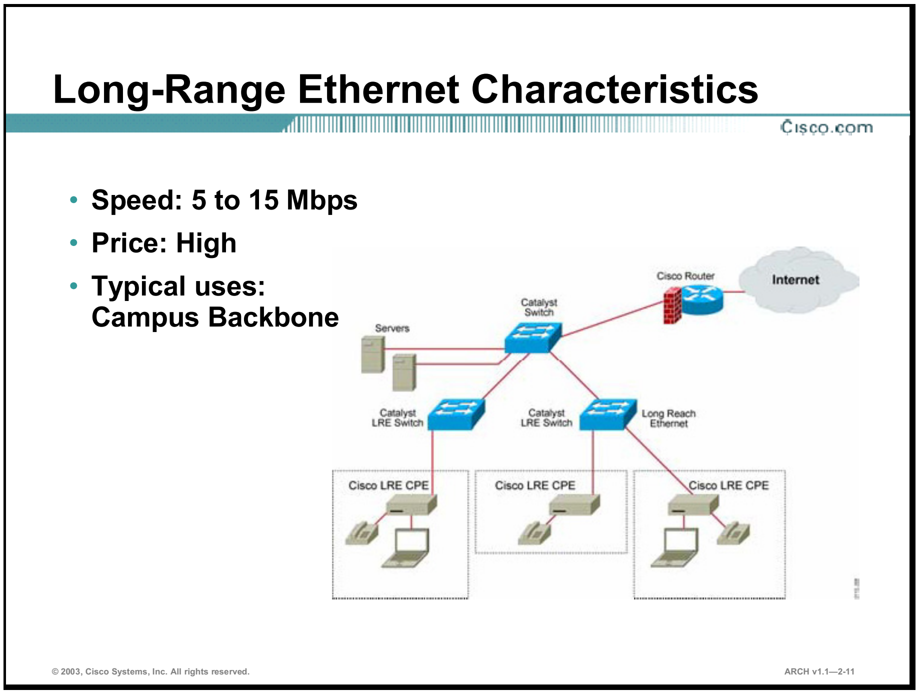

Long-Range Ethernet Characteristics

- Speed: 5 to 15 Mbps
- Price: High
- Typical uses: Campus Backbone

ARCH v1.1—2-11

For buildings with existing Category 1/2/3 wiring, Long-Range Ethernet (LRE) technology provides connectivity at speeds from 5 to 15 Mbps (full duplex) and distances up to 5000 feet. LRE technology delivers broadband service on the same lines as plain old telephone service (POTS), digital telephone, and ISDN traffic. LRE also supports modes compatible with asymmetric digital subscriber lines (ADSLs), allowing enterprises to implement LRE to buildings where broadband services currently exist.

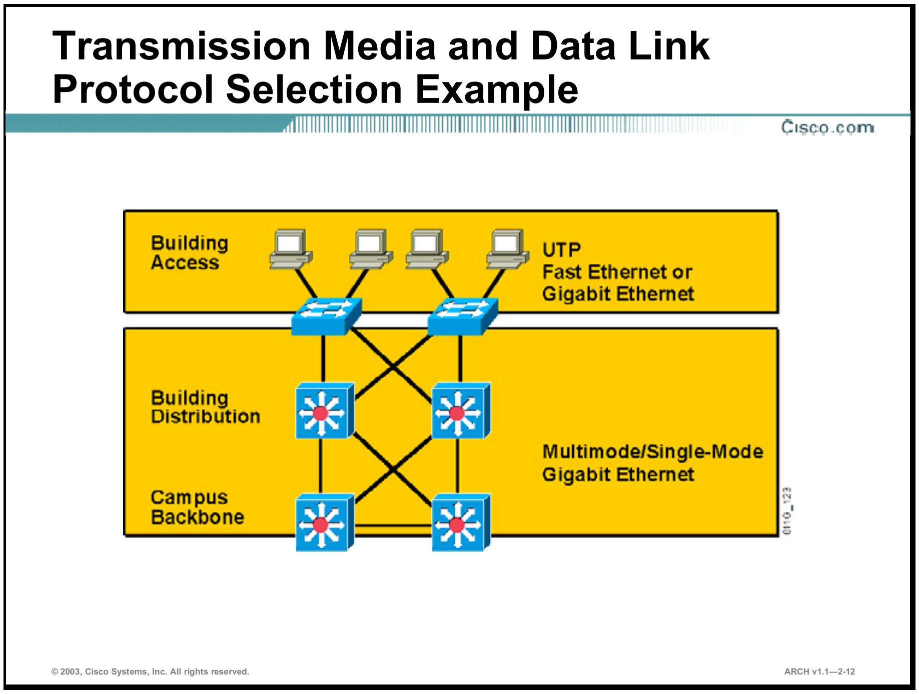

Transmission Media and Data Link Protocol Selection Example

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-12

The figure shows the transmission media for a typical campus network structure. Building Access submodule devices that are no more than 100 meters away from the LAN switch use UTP. The UTP wiring can easily handle the required distance and speed and is easy to install.

Multimode and single-mode optical cables are used to handle higher-speed requirements at the Campus Backbone and Building Distribution submodules. Multimode optical cable is usually satisfactory inside a building. For communications between buildings, multimode or single-mode cable is used, depending on distance limitations.

When selecting transmission media, scalability for the future is a key concern, as well as shifted and nonshifted dispersion fiber, the wavelength, and the diameter of the fiber.

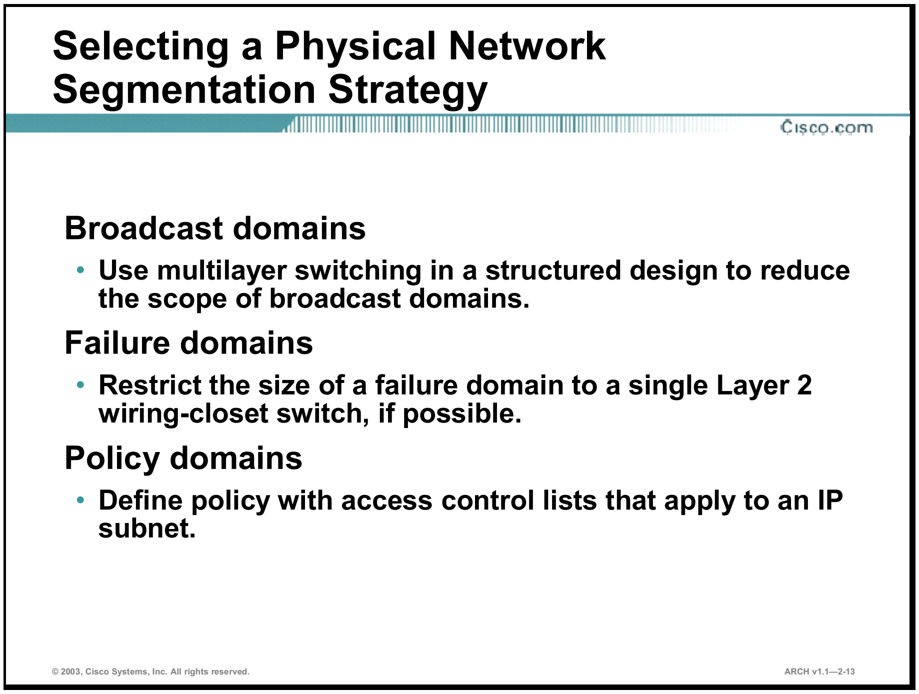

Selecting a Physical Network Segmentation Strategy

Broadcast domains

- Use multilayer switching in a structured design to reduce the scope of broadcast domains.

Failure domains

- Restrict the size of a failure domain to a single Layer 2 wiring-closet switch, if possible.

Policy domains

- Define policy with access control lists that apply to an IP subnet.

The next decision is to select a network segmentation strategy based on broadcast domains, failure domains, and policy domains. You will also determine how to implement STP to complement your segmentation strategy.

The effect of broadcast, failure, and policy domains on the campus network includes the following:

- **Broadcast domain:** MAC-layer broadcasts flood throughout the data link layer switched domain. Use multilayer switching in a structured design to reduce the scope of broadcast domains. In addition, intelligent, protocol-aware features of multilayer switches will further contain broadcasts such as Dynamic Host Configuration Protocol (DHCP) by converting them into directed unicasts. These protocol-aware features are a function of the Cisco IOS software, which is common to multilayer switches and routers.

- **Failure domain:** A group of data link layer switches connected together to extend a single network is called a Layer 2 switched domain. Consider the data link layer switched domain to be a failure domain, because a misconfigured or malfunctioning workstation can introduce errors that have an impact on or disable the entire domain. A jabbering network interface card (NIC) may flood the entire domain with broadcasts. A workstation with the wrong IP address can become a black hole for packets. Problems of this nature are difficult to localize.

  Restrict failure domains to one data link layer switch in one wiring closet, if possible, to reduce the scope of the failure domain. Restrict the deployment of VLANs and VLAN trunking. Ideally, each VLAN (IP subnet) is restricted to one wiring-closet switch. The gigabit uplinks from each wiring-closet switch connect directly to routed interfaces on multilayer switches in the Building Distribution submodule. One way to achieve load balancing is to configure two such VLANs in the wiring-closet switch, which is shown later.

- **Policy domain:** Policy is usually defined on the routers or multilayer switches in the campus network. A convenient way to define policy is with access control lists (ACLs) that apply to an IP subnet. Thus, a group of users or servers with similar quality of service (QoS) or security policies can be conveniently grouped together in the same IP subnet and the same VLAN. Other services, such as DHCP, are defined on an IP subnet basis.

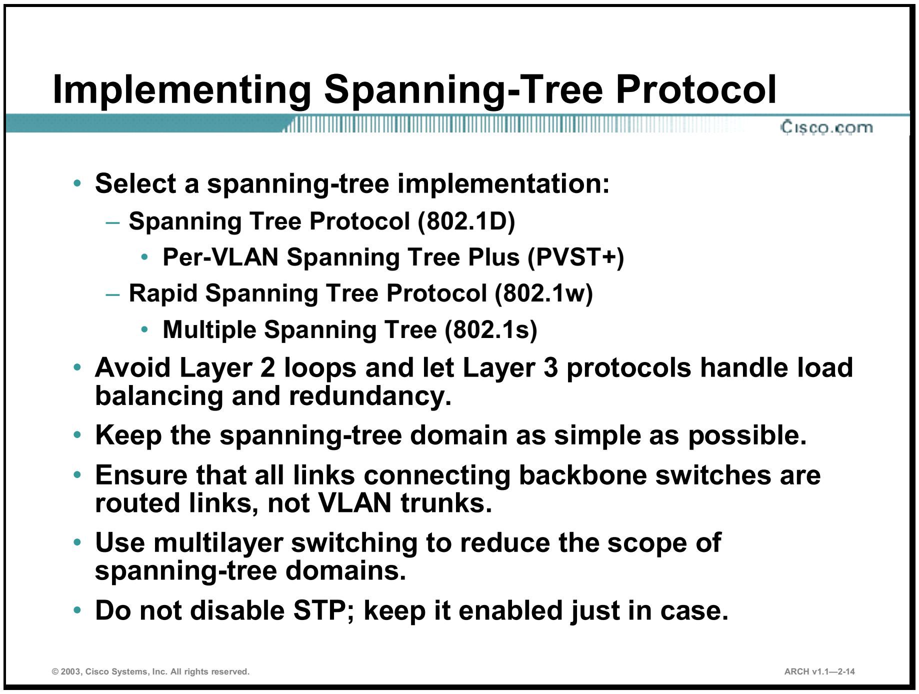

Implementing Spanning-Tree Protocol

- **Select a spanning-tree implementation:**
  - **Spanning Tree Protocol (802.1D)**
    - **Per-VLAN Spanning Tree Plus (PVST+)**
  - **Rapid Spanning Tree Protocol (802.1w)**
    - **Multiple Spanning Tree (802.1s)**
- **Avoid Layer 2 loops and let Layer 3 protocols handle load balancing and redundancy.**
- **Keep the spanning-tree domain as simple as possible.**
- **Ensure that all links connecting backbone switches are routed links, not VLAN trunks.**
- **Use multilayer switching to reduce the scope of spanning-tree domains.**
- **Do not disable STP; keep it enabled just in case.**

ARCH v1.1—2-14

Data link layer switches run STP to block loops in the Layer 2 topology. If loops are included in the data link layer design, redundant links are put in blocking mode and do not forward traffic. It is better to avoid Layer 2 loops by design and have the Layer 3 protocols handle load balancing and redundancy, so that all links are used for traffic.

Keep the spanning-tree domain as simple as possible and avoid loops. With loops in the Layer 2 topology, Spanning Tree Protocol takes 30 to 50 seconds to converge. Therefore, avoiding loops is especially important in the mission-critical parts of the network.

To prevent STP convergence events in the Campus Backbone submodule, ensure that all links connecting backbone switches are routed links, not VLAN trunks. This will also constrain the broadcast and failure domains.

Use multilayer switching in a structured design to reduce the scope of spanning-tree domains. Let a Layer 3 routing protocol, such as Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF), handle load balancing, redundancy, and recovery in the backbone.

Selecting Data Link Layer or Multilayer Switching Solutions

The Enterprise Composite Network Model combines data link layer switching with multilayer switching to achieve robust, highly available campus networks. In this topic, you will learn how to select data link layer switching and multilayer switching solutions for the Building Access, Building Distribution, and Campus Backbone submodules of the Campus Infrastructure module, given specific internetwork requirements.

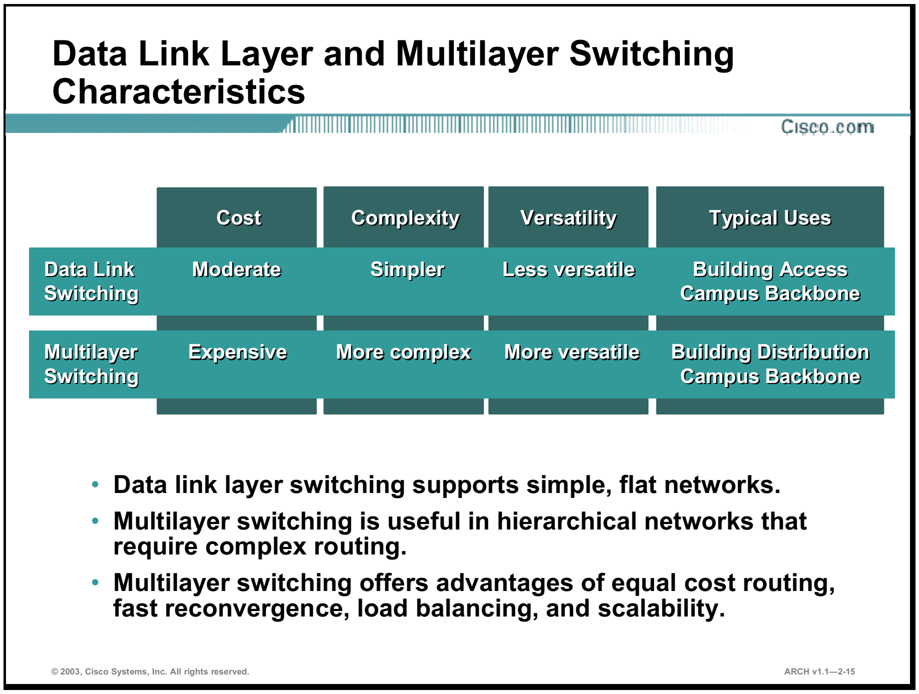

Data Link Layer and Multilayer Switching Characteristics

Cisco.com

	Cost	Complexity	Versatility	Typical Uses
Data Link Switching	Moderate	Simpler	Less versatile	Building Access Campus Backbone
Multilayer Switching	Expensive	More complex	More versatile	Building Distribution Campus Backbone

- **Data link layer switching supports simple, flat networks.**
- **Multilayer switching is useful in hierarchical networks that require complex routing.**
- **Multilayer switching offers advantages of equal cost routing, fast reconvergence, load balancing, and scalability.**

ARCH v1.1—2-15

The development of data link layer switching in hardware led to network designs that emphasized data link layer switching. These designs are often characterized as “flat” because they are most often based on the campus-wide VLAN model, in which a set of VLANs spans the entire network. This type of architecture favored the departmental segmentation approach in which, for example, all marketing or engineering users needed to use their own broadcast domain to avoid crossing slow routers. Because these departments could exist anywhere within the network, VLANs had to span the entire network.

Multilayer switching provides the identical advantages as routing, with the added performance boost from packet forwarding handled by specialized hardware. Adding multilayer switching in the Building Distribution and Campus Backbone submodules of the Campus Infrastructure module segments the campus into smaller, more manageable pieces, as defined in several different ways. This approach also eliminates the need for campus-wide VLANs, allowing for the design and implementation of a far more scalable architecture.

Today’s multilayer switches offer advances in semiconductor technology, enabling Cisco to offer more features. Therefore, you can implement multilayer functionality and multilayer control in the campus backbone at a cost-effective price point.

The figure summarizes the selection criteria for data link layer switching and multilayer switching for a campus network.

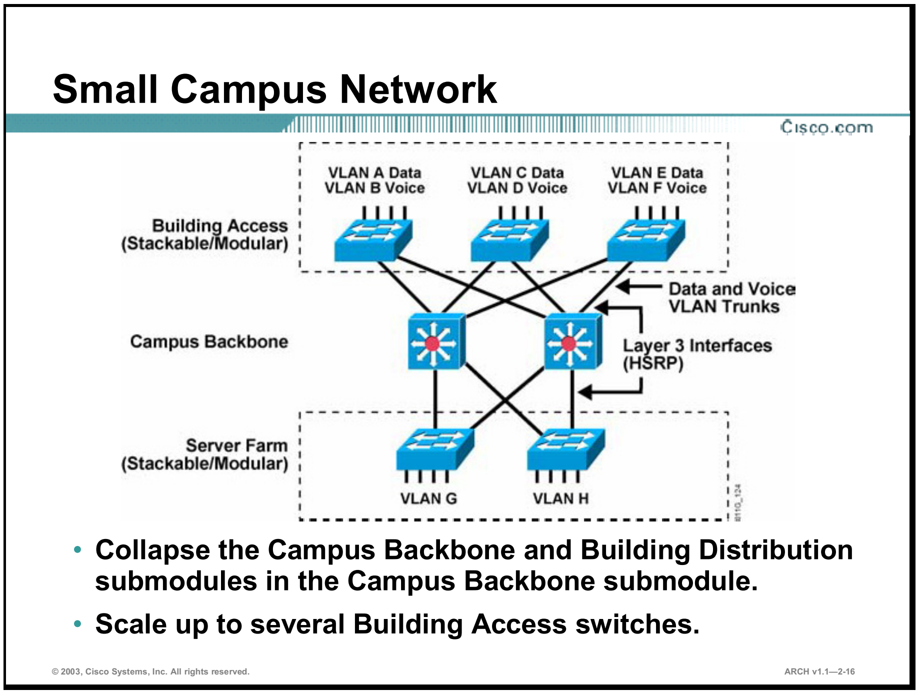

Small Campus Network

Cisco.com

VLAN A Data
VLAN B Voice
VLAN C Data
VLAN D Voice
VLAN E Data
VLAN F Voice
Building Access (Stackable/Modular)
Data and Voice VLAN Trunks
Campus Backbone
Layer 3 Interfaces (HSRP)
Server Farm (Stackable/Modular)
VLAN G
VLAN H
8110_124

- Collapse the Campus Backbone and Building Distribution submodules in the Campus Backbone submodule.
- Scale up to several Building Access switches.

ARCH v1.1—2-16

The small campus network design is appropriate for a building-sized network with up to several thousand networked devices. You can collapse the Campus Backbone and Building Distribution submodules into one layer for a small campus network. The Campus Backbone provides aggregation for Building Access switches. Cost-effectiveness in this model comes with a tradeoff between scalability and investment protection. The lack of distinct Campus Backbone and Building Distribution submodules and limited port density in the Campus Backbone restricts scaling in this model.

The building design shown in the figure comprises a single redundant building block. The two multilayer switches form a collapsed Campus Backbone. Data link layer switches are deployed in the wiring closets for desktop connectivity. Each data link layer switch has redundant gigabit uplinks to the backbone switches.

In the building design shown in the figure, servers are attached to data link layer switches or directly to the multilayer backbone switches, depending on performance and density requirements.

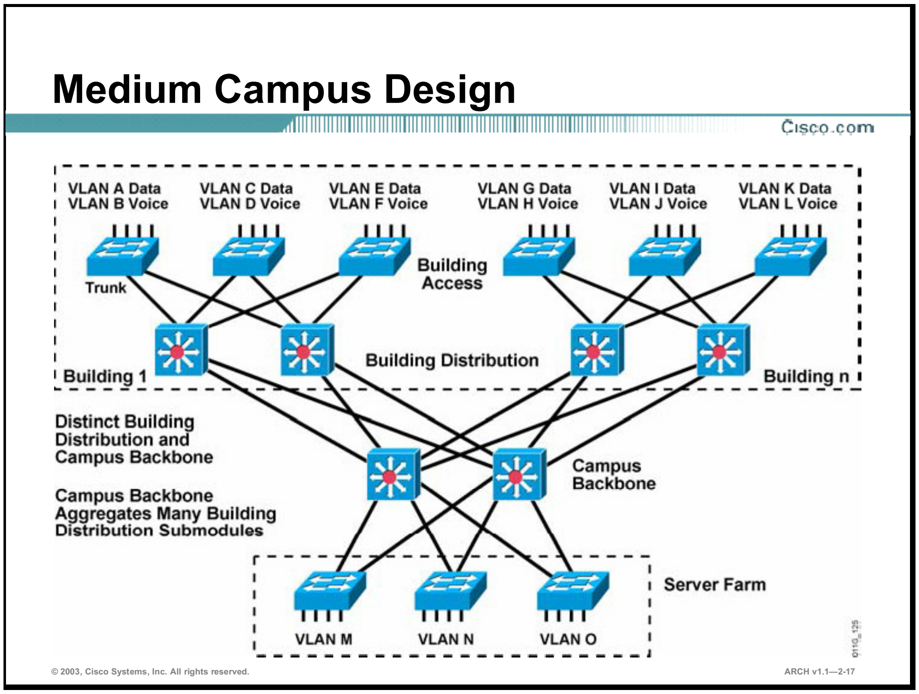

Medium Campus Design

Cisco.com

ARCH v1.1—2-17

A medium campus design with higher availability and higher capacity is shown in the figure. The most flexible and scalable campus backbone consists of multilayer switches, as shown in the figure. The backbone switches are connected by routed Gigabit Ethernet or Gigabit EtherChannel links. Multilayer-switched backbones offer these advantages:

- Reduced router peering
- Flexible topology with no spanning-tree loops
- Multicast and broadcast control in the backbone
- Scalability to arbitrarily large size

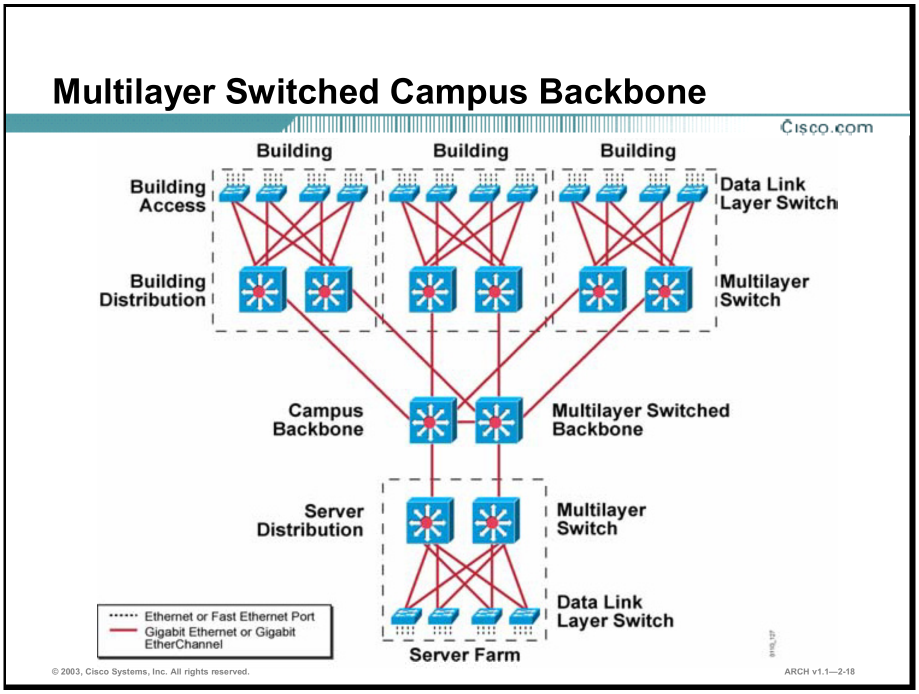

Multilayer Switched Campus Backbone

Cisco.com

Building
Building
Building
Building Access
Building Distribution
Data Link Layer Switch
Multilayer Switch
Campus Backbone
Multilayer Switched Backbone
Server Distribution
Multilayer Switch
Data Link Layer Switch
Server Farm

······ Ethernet or Fast Ethernet Port
—— Gigabit Ethernet or Gigabit EtherChannel

ARCH v1.1—2-18

The most flexible and scalable campus backbone consists of multilayer switches, as shown in the figure. The backbone switches are connected by routed Gigabit Ethernet or Gigabit EtherChannel links. Multilayer-switched backbones have several advantages:

- Reduced router peering
- Flexible topology with no spanning-tree loops
- Multicast and broadcast control in the backbone
- Scalability to an arbitrarily large size

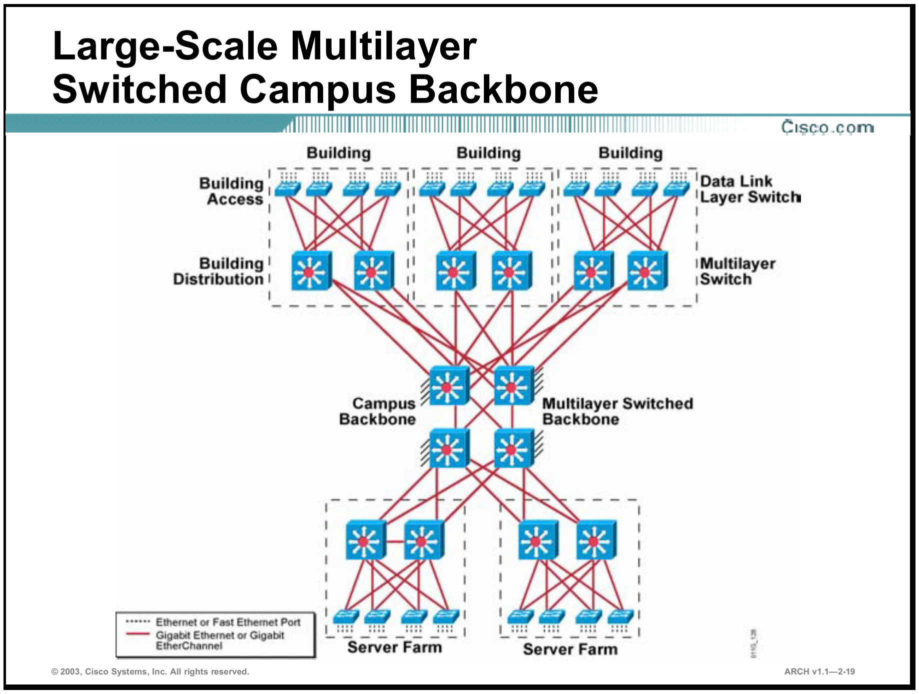

Large-Scale Multilayer Switched Campus Backbone

Cisco.com

ARCH v1.1—2-19

The figure shows the multilayer-switched campus backbone on a large scale. The multilayer-switched backbone has the advantage that arbitrary topologies are supported because a sophisticated routing protocol such as EIGRP or OSPF is used pervasively.

In the figure, the backbone consists of four multilayer switches with Gigabit Ethernet or Gigabit EtherChannel links. All links in the backbone are routed links, so there are no spanning-tree loops. The figure suggests the actual scale by showing several gigabit links connected to the backbone switches. Note that a full mesh of connectivity between backbone switches is possible but not required. Consider traffic patterns when allocating link bandwidth in the backbone.

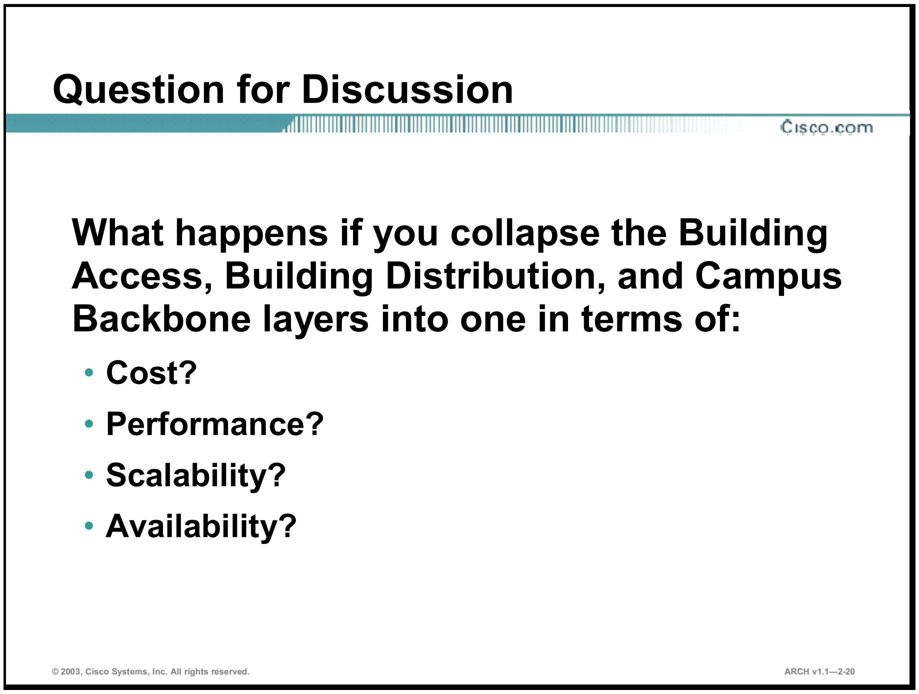

Question for Discussion

What happens if you collapse the Building Access, Building Distribution, and Campus Backbone layers into one in terms of:

- **Cost?**
- **Performance?**
- **Scalability?**
- **Availability?**

ARCH v1.1—2-20

Consider the questions listed in the figure.

Selecting Cisco Hardware and Software

The next step in the campus network design methodology requires that you select the Cisco hardware and software for each location on the network that you identified as part of your design. This topic helps you select hardware, hardware options, and software options for a campus network infrastructure, given specific internetwork requirements.

Selecting Hardware and Software

Cisco.com

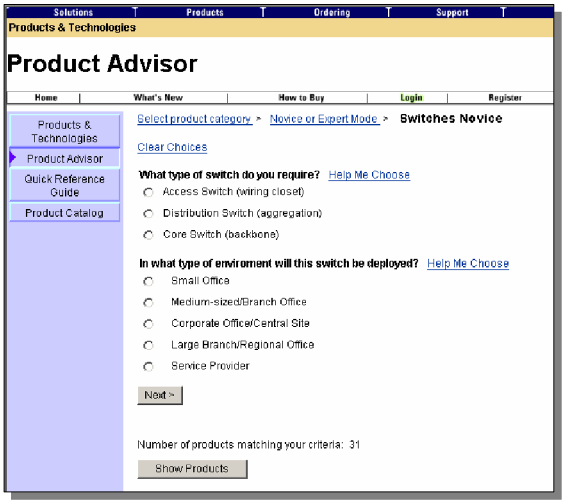

Product Advisor

Select product category > Novice or Expert Mode > **Switches Novice**

Clear Choices

What type of switch do you require? Help Me Choose

○ Access Switch (wiring closet)

○ Distribution Switch (aggregation)

○ Core Switch (backbone)

In what type of enviroment will this switch be deployed? Help Me Choose

○ Small Office

○ Medium-sized/Branch Office

○ Corporate Office/Central Site

○ Large Branch/Regional Office

○ Service Provider

Next >

Number of products matching your criteria: 31

Show Products

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-21

Cisco offers the Cisco Product Advisor to help you select the right switch solution for the enterprise campus network. The tool operates in two modes: novice and expert. To access the Product Advisor, go to http://www.cisco.com/en/US/products/prod_tools_index.html and click the Cisco Product Advisor link. Then click a device category. The Product Advisor will ask you questions to help select routers for particular needs. It does not include all products and features, but provides helpful information to help you select appropriate Cisco products.

The table summarizes the questions to answer to help select the right Cisco switches for each location in the design.

Step	Description	Notes and Comments
1.	Determine the size of the campus location or building for the switch.	Select one: ■ Small campus location/building ■ Medium campus location/building ■ Large campus location/building
2.	Determine the speed required for most ports (excluding uplinks).	Select one: ■ 10 MB ■ 10/100 MB ■ 10/100/1000 MB (copper ports) ■ 1000 MB only (GBIC ports)
3.	Determine the number of ports required.	Select any range from 1–24 up to 337 or more.
4.	Select the media types for uplinks.	Select one or more: ■ Fiber ■ Copper
5.	Determine the speed required for uplinks.	Select one or more: ■ 100 MB ■ 1000 MB ■ 10 GB
6.	Determine the number of uplinks required.	Select one: ■ 1 ■ 2 ■ more than 2
7.	Determine if you require redundant power.	Select Yes or No.
8.	Determine if the switch requires redundant switching engines.	Select Yes or No.
9.	Determine if the switch requires IP routing or multilayer switching.	Select Yes or No.
10.	Determine if the switch requires intrusion detection, server load balancing, or network analysis.	Select Yes or No.
11.	Determine if the switch requires inline power for IP phones or wireless access points.	Select Yes or No.
12.	Determine if you need IOS or Catalyst software for each data link layer switch and multilayer switch you selected.	
13.	Select the IOS version for each data link layer switch and multilayer switch you selected.	
14.	Select the intelligent network services and provision each feature for each data link layer switch and multilayer switch you selected.	

Identifying an IP Addressing Strategy

An effective IP addressing scheme is critical to the success and scalability of the network. This topic helps you identify an IP addressing strategy for the campus network, given specific internetwork requirements.

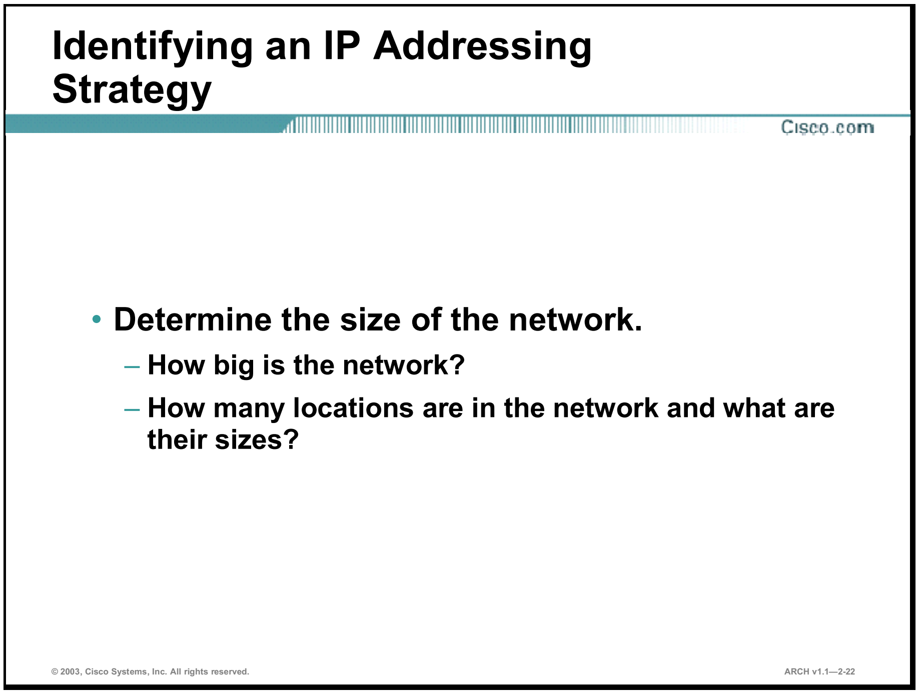

Identifying an IP Addressing Strategy

Cisco.com

- **Determine the size of the network.**
  - **How big is the network?**
  - **How many locations are in the network and what are their sizes?**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-22

The first step in an IP addressing plan design is to determine the size of the network in order to establish how many IP addresses are needed. To gather the required information, answer these questions:

- **How big is the network?** Determine the number of workstations, servers, IP Phones, router interfaces, switch management interfaces, firewall interfaces, and so on. The summary defines the minimum overall number of IP addresses required for the network. Because all networks tend to grow, allow a reserve of about 20 to 40 percent for potential network and application expansion.
- **How many locations are in the network and what are their sizes?** The information about the size and number of the individual locations is closely related to the overall network size. These values should correlate to the segmentation strategy chosen for the campus networks. These values impact the subnetwork addressing scheme deployed to accommodate all locations and the number of IP addresses required in each location.

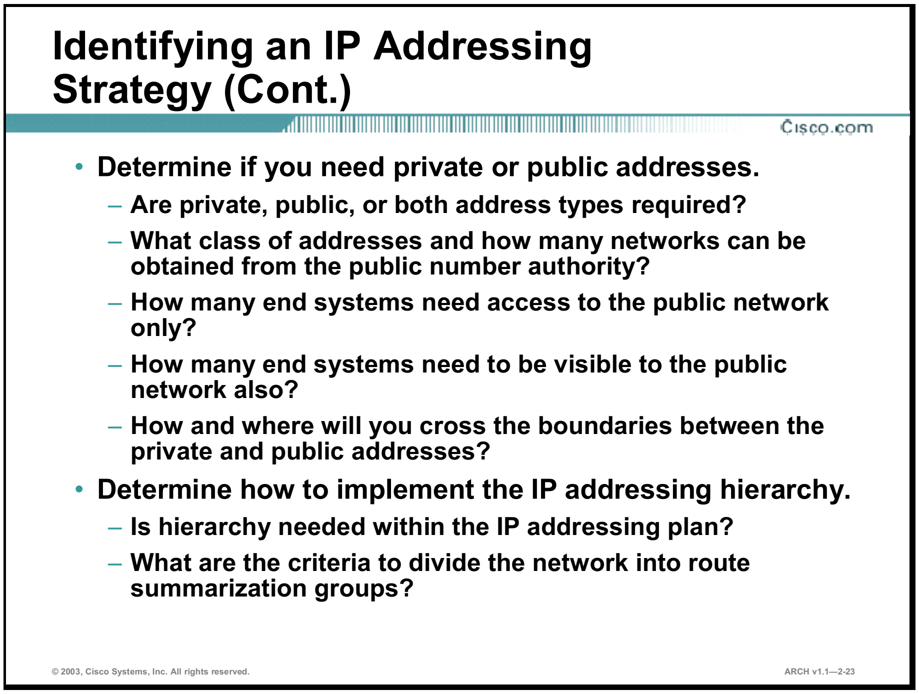

Identifying an IP Addressing Strategy (Cont.)

- **Determine if you need private or public addresses.**
  - Are private, public, or both address types required?
  - What class of addresses and how many networks can be obtained from the public number authority?
  - How many end systems need access to the public network only?
  - How many end systems need to be visible to the public network also?
  - How and where will you cross the boundaries between the private and public addresses?
- **Determine how to implement the IP addressing hierarchy.**
  - Is hierarchy needed within the IP addressing plan?
  - What are the criteria to divide the network into route summarization groups?

ARCH v1.1—2-23

Next, determine if you need private or public addresses based on these questions:

- **Are private, public, or both address types required?** The decision about when to use private, public, or both address types depends on the Internet connection presence and the number of publicly visible servers. Four situations are possible:
  - **No Internet connectivity:** The network is isolated and there is no need to acquire public addresses.
  - **Internet connectivity, no public-accessible servers:** The network is connected to the Internet and thus at least one public IPv4 address is required. Use one public IPv4 address and a translation mechanism such as port address translation (PAT) to allow access to the Internet from a single IP address. Private addresses are used to address the internal network.

Note Some applications do not support translation to a single IP address using PAT. Therefore, sufficient IPv4 addresses will be required to support one-to-one Network Address Translation (NAT) for each user concurrently accessing these applications through the Internet.

  - **Internet connectivity, public-accessible servers:** The public addresses are required to connect all public-accessible servers to the Internet. The number of public addresses corresponds to the number of Internet connections and public-accessible servers.
  - **All end systems should be public accessible:** Only public IPv4 addresses are required and used to address the whole network.

- **What class of addresses and how many networks can be obtained from the authority assigning public numbers, usually the ISP?** The required IPv4 address classes for the planned network are based on information about the network size, the number of locations, and the size of the individual locations.
- **How many end systems need access to the public network only (not publicly visible)?** 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.
- **How many end systems also 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) and defines the number of required public IPv4 addresses.
- **How and where will you cross the boundaries between the private and public IPv4 addresses?** When private addresses are used for addressing in a network, and this network needs to be connected to the Internet, a translation mechanism such as NAT must be used to translate from private to public addresses and vice versa.

The decision on how to implement the IP addressing hierarchy is an administrative decision that is based on these questions:

- **Is hierarchy needed within the IP addressing plan?** You will decide how to implement the IP addressing hierarchy based on the network size and the geography and topology of the network. In large networks, a hierarchical IP addressing plan is required to promote a stable network. Also, routing protocols such as OSPF rely on a hierarchical addressing plan.
- **What are the criteria to divide the network into route summarization groups?** The network is usually divided into route summarization groups based on the network size and topology. To reduce the routing overhead in a large network, a multilevel hierarchy may be required. The depth of hierarchy levels depends on the network size, topology, number of network layers, and the size of the upper-level summarization group.

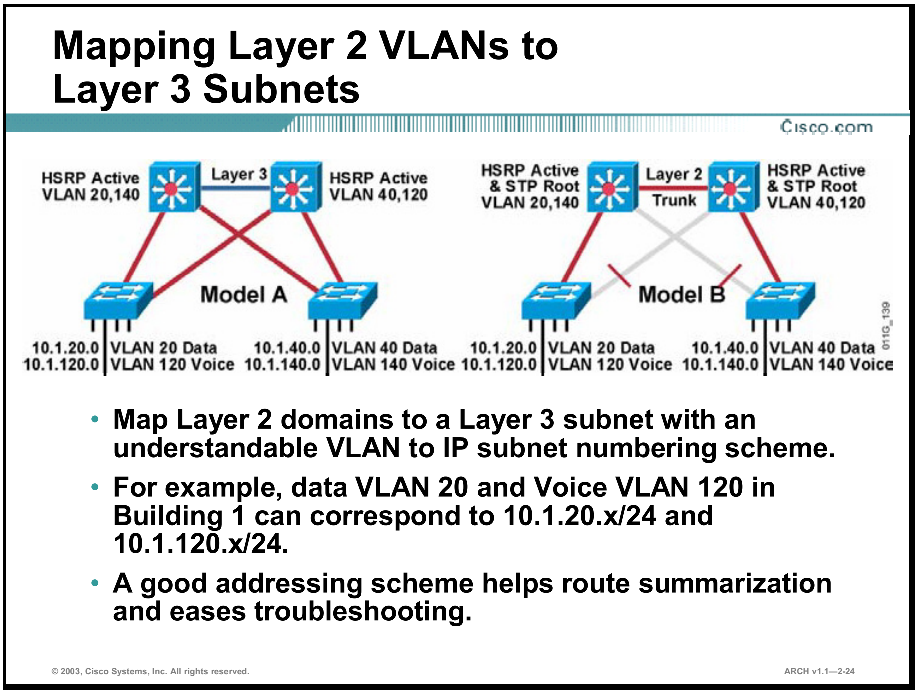

Mapping Layer 2 VLANs to Layer 3 Subnets

- **Map Layer 2 domains to a Layer 3 subnet with an understandable VLAN to IP subnet numbering scheme.**
- **For example, data VLAN 20 and Voice VLAN 120 in Building 1 can correspond to 10.1.20.x/24 and 10.1.120.x/24.**
- **A good addressing scheme helps route summarization and eases troubleshooting.**

ARCH v1.1—2-24

In a structured design, one IP subnet maps to a single VLAN, which is carried to the wiring-closet switch. A good IP addressing scheme can take advantage of multilayer switching features to exchange summarized routing information, rather than learning the path to every host in the whole network. Summarization is important to routing protocol scalability for OSPF and EIGRP.

In the figure, Model A shows multilayer switching only between the Building Distribution switches. Model B shows data link layer switching, without forwarding on both uplinks for separate VLANs.

Selecting Routing Protocols

The decision about which routing protocols to implement is based on the design goals and the physical topology of the network and the configuration of links for remote sites. This topic helps you select routing protocols for the campus network that meet performance, scalability, and availability requirements.

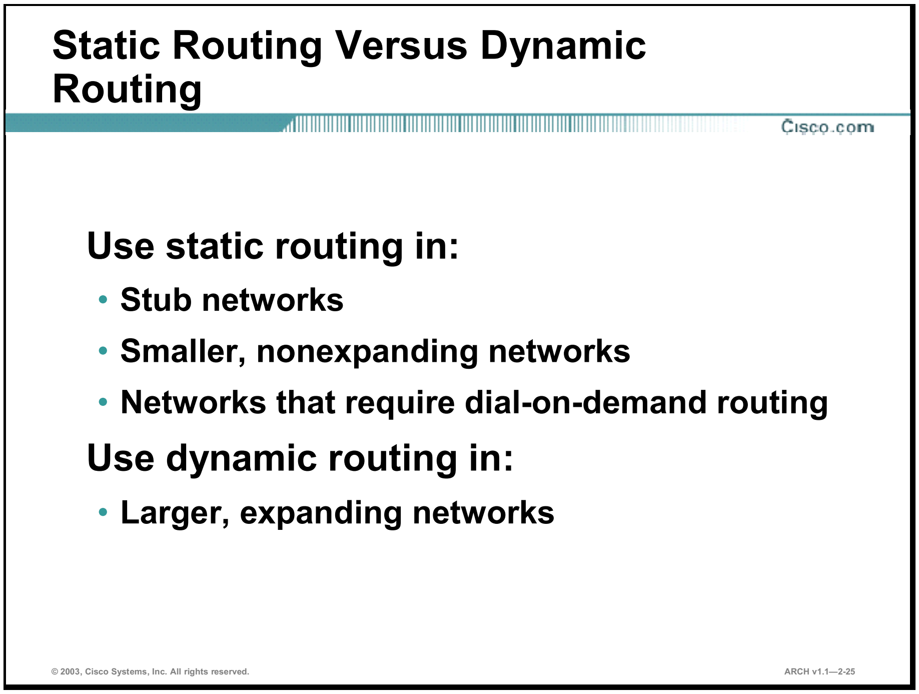

Static Routing Versus Dynamic Routing

Use static routing in:

- **Stub networks**
- **Smaller, nonexpanding networks**
- **Networks that require dial-on-demand routing**

Use dynamic routing in:

- **Larger, expanding networks**

ARCH v1.1—2-25

Static routing is primarily used for:

- Routing to and from stub networks. A stub network only carries traffic for local hosts, and typically has only one entry/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.
- Supporting special features such as dial-on-demand routing (DDR) and On-Demand Routing (ODR).
- Specifying routes toward dialing peers in dial-in environments.

Configuration and maintenance of static routes is time consuming. It requires that you have complete knowledge of the whole network to implement it properly.

Dynamic routing protocols have two major advantages over static routing protocols:

- 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, except for DDR, ODR, a stub network, or a dial-in scenario.

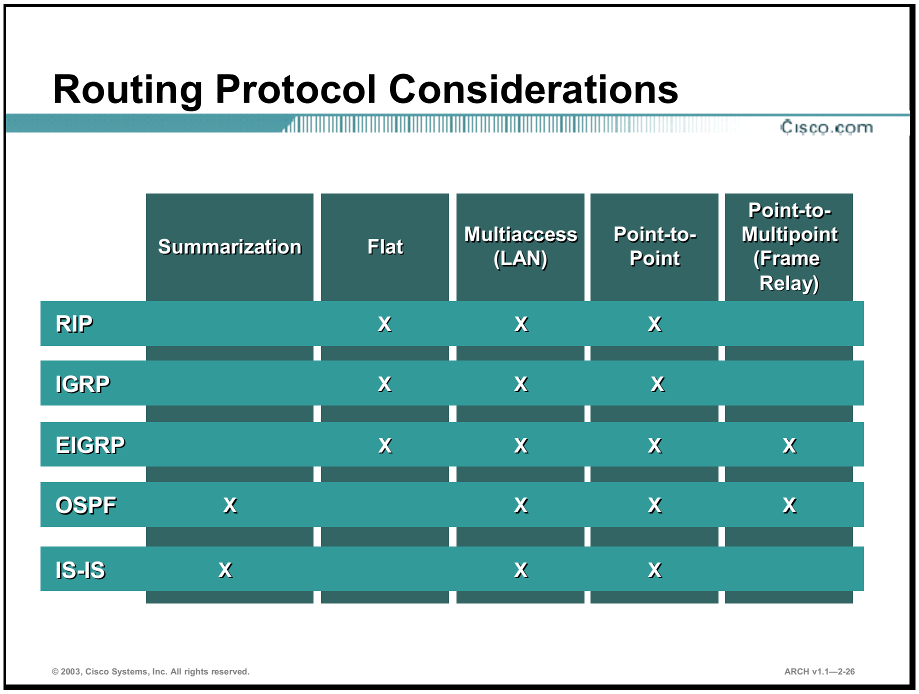

Routing Protocol Considerations

	Summarization	Flat	Multiaccess (LAN)	Point-to-Point	Point-to-Multipoint (Frame Relay)
RIP		X	X	X	
IGRP		X	X	X	
EIGRP		X	X	X	X
OSPF	X		X	X	X
IS-IS	X		X	X	

ARCH v1.1—2-26

When to Choose RIP or RIPv2

Routing Information Protocol (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 variable-length subnet masking (VLSM) and manual route summarization.

RIPv2 is an enhanced version of the original RIP protocol (now referred to as RIPv1) that supports VLSM. RIPv2 is implemented mainly in small networks, especially small hub-and-spoke networks using point-to-point links. RIPv2 with snapshot routing support 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 start exchange of routing information. It is seldom used in LAN environments because it is chatty and has a low hop count limit. 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 because of large bandwidth requirements.

When to Choose IGRP

Interior Gateway Routing Protocol (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 may be suitable for small to medium networks. Like RIP, it has problems with the split-horizon feature in NBMA networks. (You can disable split horizon, but distance vector protocols result in high-bandwidth usage if propagating entire tables.) Other problems of IGRP include its slow convergence due to its pure distance vector operation, and the potential for high bandwidth utilization when propagating entire routing tables to neighbors. IGRP is not typically recommended for new deployments.

When to Choose EIGRP

EIGRP, which was developed based on IGRP, is a very powerful routing protocol. EIGRP is a Cisco proprietary protocol licensed to limited vendors. It is an advanced distance vector protocol with some 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.

EIGRP is not suitable for dial-up environments because it must maintain the neighbor relationship and uses periodic hello packets, effectively maintaining the dial-up connections all the time.

EIGRP offers these features:

- VLSM support
- Advanced metrics
- Fast convergence
- Scalability
- Authentication
- Flexible summarization
- Configurable bandwidth usage
- Low bandwidth during normal conditions
- Load balancing across unequal cost paths
- Support for stub networks

When to Choose OSPF

OSPF is a standards-based link-state protocol, based on the shortest path first (SPF) or Dijkstra's algorithm for path calculation. Initially it was designed for networks that consisted of point-to-point links, but later it was successfully adapted for operation in LAN and NBMA environments. OSPF can be 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 all nonbackbone areas must be attached directly to that backbone area. Expansion of the backbone area can cause design issues, because the backbone area must remain contiguous.

OSPF offers these features:

- With OSPF, there is no limitation on the hop count. The intelligent use of VLSM is very useful in IP address allocation.
- OSPF uses IP multicast to send link-state updates. This ensures less processing on routers that are not listening to OSPF packets. Updates are only sent when routing changes occur rather than periodically. This ensures a better use of bandwidth.
- OSPF offers fast convergence because routing changes are propagated instantaneously and not periodically (a characteristic of distance vector routing protocols).
- OSPF allows for effective load balancing.
- OSPF allows for a logical definition of networks where routers can be divided into areas. This will limit the explosion of link-state updates over the whole network. This also provides a mechanism for aggregating routes and cutting down on the unnecessary propagation of subnet information.
- OSPF allows for routing authentication by using different methods of password authentication.
- OSPF allows for the transfer and tagging of external routes injected into an autonomous system. This keeps track of external routes injected by exterior protocols such as Border Gateway Protocol (BGP).

When to Choose Integrated IS-IS

Integrated Intermediate System-to-Intermediate System (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, the IS-IS backbone can easily be expanded 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 point to be considered prior to implementation. 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 for 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.

Selecting Areas or Networks

After selecting routing protocols, you will identify areas or networks for that routing protocol. The key considerations for routing areas or networks are based on:

- Number of multilayer devices and the type of CPU power and media available
- Type of topology, either full mesh or partial mesh
- Multilayer device memory

Enterprise Campus Design Examples

Small, medium, and large campus networks use the Campus Infrastructure model applied to specific situations. This topic contains example network designs to help you propose small, medium, and large campus data link layer switched and multilayer-switched network designs, given specific internetwork requirements.

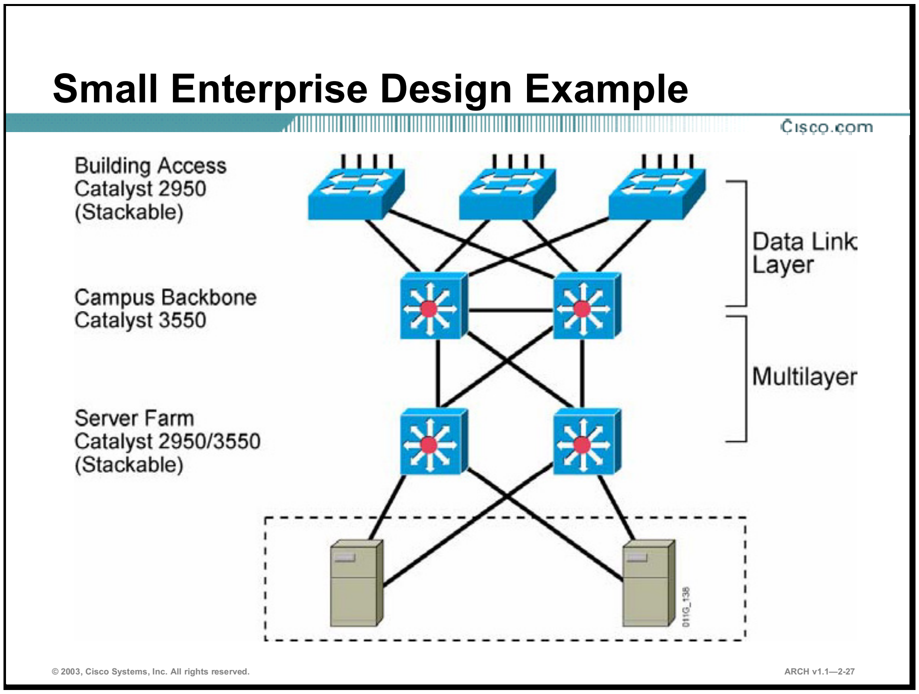

Small Enterprise Design Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-27

Company Background

This small company has about 200 users in one building with two separate floors. Their primary network user applications include e-mail, FTP file sharing, HTTP access to an intranet server, and dedicated Internet access for all employees.

Campus Design

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Notes and Comments
What is the logical network design?	Single corporate-wide VLAN	There is no requirement for voice so the entire network of 200 users will be in a single VLAN.
What physical network media will be used?	Twisted pair throughout the entire network	Twisted pair was selected because of its low cost.
What data link layer protocol will be used?	Fast Ethernet in the Campus Backbone Fast Ethernet in the Building Access to the desktop	Fast Ethernet is sufficient today. If more growth is realized in the future, the company could easily upgrade to Gigabit Ethernet.
What spanning-tree deployment will be used?	Spanning tree will be used and a Campus Backbone switch will be the root	For simplicity, the Catalyst 3550 was selected as the STP root because the Catalyst 3550 has the highest processing capability.
What is the data link layer/multilayer switching strategy for the network?	Network layer routing only needed on the Internet edge Data link layer switching in the network	Network layer routing is only needed for Internet access. Data link layer switching is used throughout the rest of the network.
Which Cisco products will be used?	Catalyst 3550 in the Campus Backbone submodule Catalyst 2950 in the Building Access submodule	The redundant Catalyst 3550 in the Campus Backbone is optional, and is provided for redundancy.
What IP addressing scheme will be used?	Approximately 254 IP addresses from ISP or regional number authority	
Which routing protocols will be used?	No routing protocol required, only a default route to the Internet	With the exception of a default route to the Internet, no routing protocol is required for this network design.

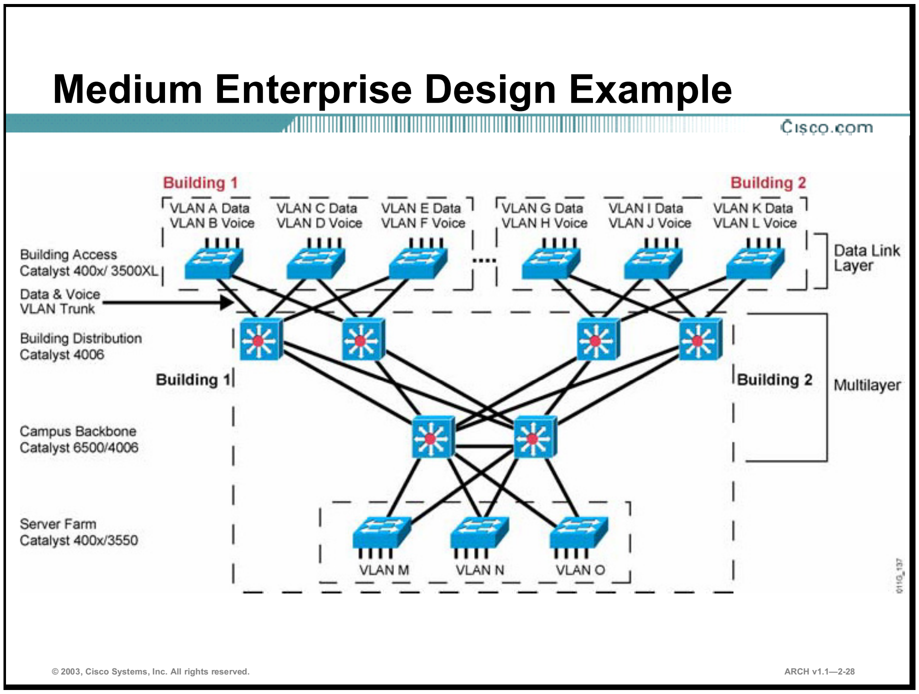

Medium Enterprise Design Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-28

Company Background

The medium-sized campus shown in the figure supports about a thousand users. The company's primary network applications are database applications. They are situated in two buildings that connect to each other via a 100-foot walkway. They have expressed an intent to add voice in the future and would like to plan for it in the infrastructure now.

Campus Design

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Notes and Comments
What is the logical network design?	Segmentation by type of traffic (voice and data) One VLAN for voice and one for data on each Building Access submodule switch To get both VLANs back to the Building Distribution submodule, the company chose ISL trunking from the access switch to the distribution switch The data center will have three server farms in unique VLANs that connect directly to the Campus Backbone submodule	The network requires at least 18 VLANs. The IT staff wants voice on one VLAN and data on a separate VLAN, which requires two VLANs in each access switch.
What physical network media will be used?	UTP is implemented from the workstations to the Building Access submodule Multimode fiber from the Building Access submodule to the Building Distribution submodule Multimode fiber from the Building Distribution submodule to the Campus Backbone submodule	The multimode fiber between the Building Access and Building Distribution submodules, and in the Campus Backbone submodule, provides an upgrade path to Gigabit Ethernet or faster technologies.

Design Question	Decision	Notes and Comments
What data link layer protocol will be used?	Gigabit Ethernet in the Campus Backbone and Building Distribution submodules Fast Ethernet in the Building Access submodule to the desktop	
What spanning-tree deployment will be used?	Spanning-tree root is the distribution device	
What is the data link layer/multilayer switching strategy for the network?	Multilayer switching in the Campus Backbone submodule Data link layer switching in the Building Distribution and Building Access submodules	
Which Cisco products will be used?	Catalyst 6500 and 400x in the Campus Backbone submodule Catalyst 4006 in the Building Distribution submodule Catalyst 400x and 3500XL PWR in the Building Access submodule Catalyst 400x and 3500XL in the Server Farm	Need inline power modules for voice equipment.
What IP addressing scheme will be used?	Private addressing Each VLAN is its own subnet Requires 18 subnets	Because of the number of IP addresses needed and no Internet access requirement, private addressing will be used in the network. The base network number for the voice VLANs will be different from the base network number for the data VLANs.
Which routing protocols will be used?	EIGRP	OSPF or even RIPv2 would also work because the network contains only two multilayer-switched devices. The company only has to maintain 18 to 24 different subnets.

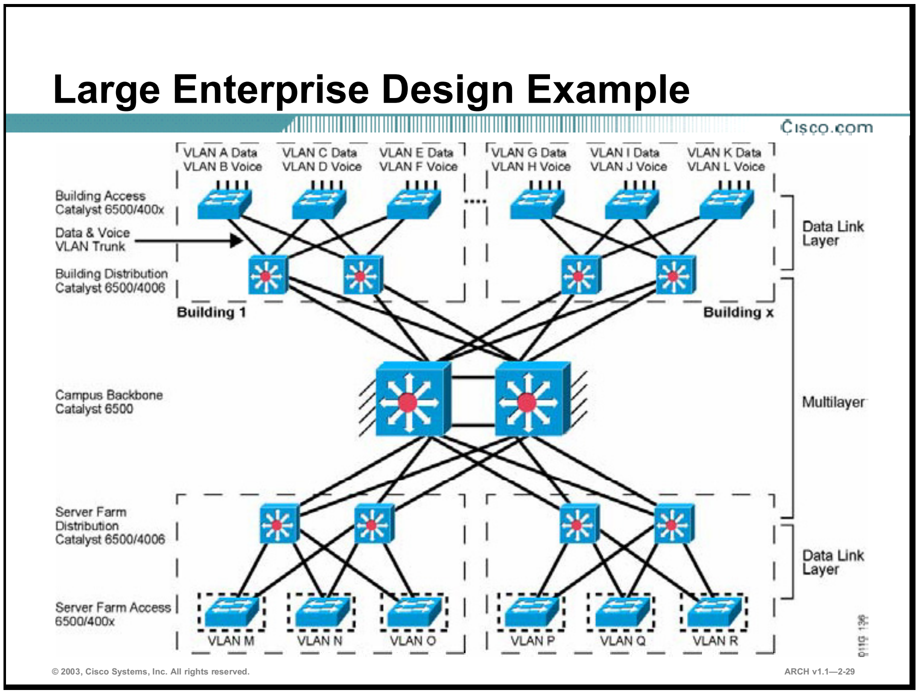

Large Enterprise Design Example

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-29

Company Background

The large enterprise network in the example supports 4000 users in four buildings. They are intending to implement voice in the future as part of an incremental deployment so they want to ensure that the infrastructure they plan today will support the known future requirement.

The information systems department has decided that each of the Building Distribution submodule devices will perform multilayer switching to limit broadcast and failure domains within each Building Distribution and Building Access switch.

The company has the budget to put in the required multiple strands of fiber between the buildings for redundancy purposes. Single-mode fiber is already installed in the risers of the buildings, enough to support the bandwidth needs.

Campus Design

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Notes and Comments
What is the logical network design?	Segmentation by type of traffic (voice and data) One VLAN for voice and one for data on each access layer switch	Each department will have its own VLAN. There will be separate subnetworks and VLANs designed for voice.
What physical network media will be used?	UTP to the desktop from the Building Access submodule Fiber to the Building Distribution and Campus Backbone submodules	Fiber runs between the buildings to create the meshed core, and will provide redundant connections between the Building Distribution and Campus Backbone submodules.
What data link layer protocol will be used?	Gigabit EtherChannel in the Building Distribution and Campus Backbone submodules Fast Ethernet in the Building Access submodule to the desktop	
What spanning-tree deployment will be used?	Spanning-tree root will be the distribution switch in each building	Eight separate spanning-tree domains, two per building.
What is the data link layer/multilayer switching strategy for the network?	Multilayer switching in the Building Distribution and Campus Backbone submodules Data link layer switching in the Building Access submodule	
Which Cisco products will be used?	Catalyst 6500 in the Campus Backbone submodule Catalyst 6500 or 4006 in the Building Distribution submodule Catalyst 400x in the Building Access submodule	
What IP addressing scheme will be used?	Class B addresses Private addressing two networks: one for data and one for voice	There will be multiple subnets used within each address range. The purpose of having two networks is to segregate the voice from the data completely. The company needs one Class C address from the service provider to support public access to the Internet.
Which routing protocols will be used?	OSPF Core will be area 0 and each building will be a different area	Given the size of the network, either EIGRP or OSPF is acceptable. Since the company wants standards-based routing protocols, they chose OSPF, with each building being its own area and the core being area 0.

Summary

This topic summarizes the key points discussed in this lesson.

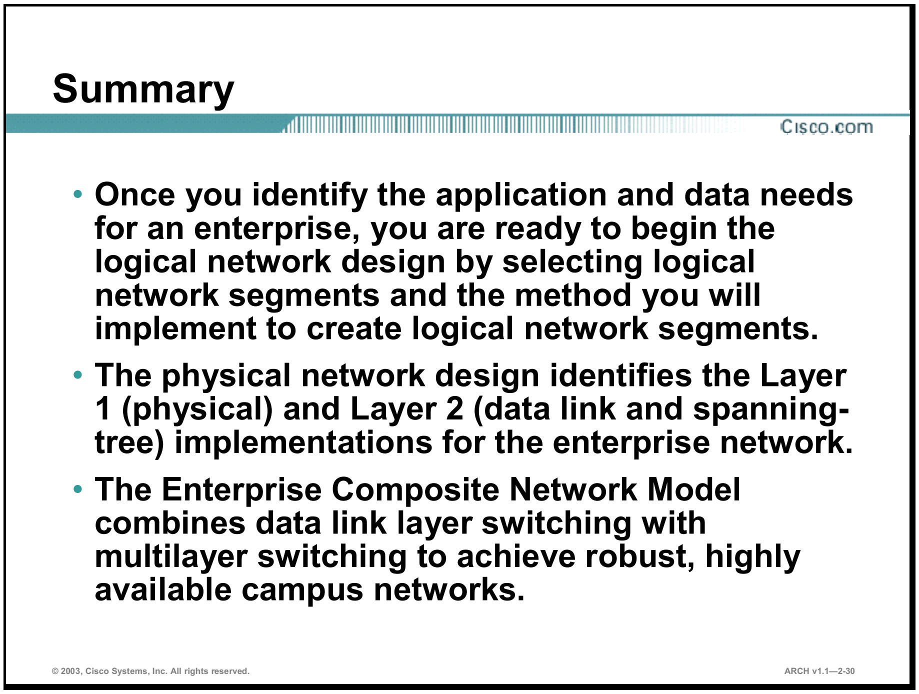

Summary

- Once you identify the application and data needs for an enterprise, you are ready to begin the logical network design by selecting logical network segments and the method you will implement to create logical network segments.
- The physical network design identifies the Layer 1 (physical) and Layer 2 (data link and spanning-tree) implementations for the enterprise network.
- The Enterprise Composite Network Model combines data link layer switching with multilayer switching to achieve robust, highly available campus networks.

ARCH v1.1—2-30

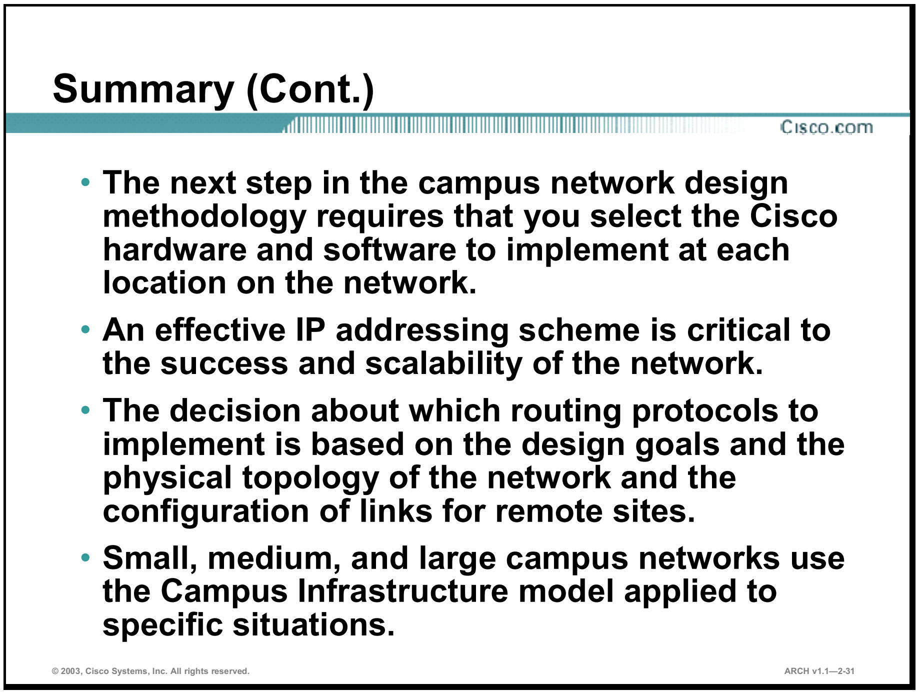

Summary (Cont.)

- The next step in the campus network design methodology requires that you select the Cisco hardware and software to implement at each location on the network.
- An effective IP addressing scheme is critical to the success and scalability of the network.
- The decision about which routing protocols to implement is based on the design goals and the physical topology of the network and the configuration of links for remote sites.
- Small, medium, and large campus networks use the Campus Infrastructure model applied to specific situations.

ARCH v1.1—2-31

References

For additional information, refer to these resources:

- *Gigabit Campus Network Design—Principles and Architecture* at http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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 about data link layer switches is correct?

A) Data link layer switches can route between VLANs.

B) Data link layer switches cannot route between VLANs.

C) Data link layer switches can route between VLANs on the same switch.

D) Data link layer switches can route between VLANs on different switches.

Q2) A campus-wide VLAN is not recommended because _____.

A) It can create routing holes.

B) It relies on recursive entries.

C) It adds complexity to the design.

D) It blocks the primary path to a destination.

Q3) How do you reduce the broadcast domain?

A) Use Layer 3 routing.

B) Use Layer 4 protocols.

C) Use data link layer switching.

D) You cannot reduce the broadcast domain.

Q4) The most flexible and scalable Campus Backbone submodule consists of _____ switches.

A) multilayer

B) aggregated

C) data link layer

D) nonaggregated

Q5) When using the Cisco Product Advisor to select products for a campus network, what should you keep in mind?

A) The tool does not cover all products and all features.

B) The tool is all encompassing, covering all products and features.

C) The tool can be reconfigured for your specific needs when entering the tool.

D) Any situation the tool does not cover does not need to be addressed in the design phase.

Q6) When planning an IP address scheme, which feature is critical to provide scalable routing?

A) scalability

B) flexibility

C) recoverability

D) summarization

Q7) Which three routing protocols does Cisco recommend for flat networks? (Choose three.)

A) IS-IS

B) IGRP

C) OSPF

D) RIPv1

E) EIGRP

Q8) Given a small campus network with about 200 users in one building, what type of VLAN implementation would you design?

A) Do not implement VLANs.

B) Use one VLAN across the company.

C) Implement VLANs for each department.

D) Implement VLANs for each type of traffic.

Q9) Given a large campus network with thousands of users, how would you implement Spanning Tree Protocol?

A) Do not use Spanning Tree Protocol.

B) Design a single spanning-tree domain.

C) Implement the spanning-tree root on each access device in the buildings.

D) Select the spanning-tree root for each distribution device in the buildings.

Quiz Answer Key

Q1) B

Relates to: Designing the Logical Campus Network

Q2) C

Relates to: Designing the Logical Campus Network

Q3) A

Relates to: Designing the Physical Campus Network

Q4) A

Relates to: Selecting Data Link Layer or Multilayer Switching Solutions

Q5) A

Relates to: Selecting Cisco Hardware and Software

Q6) D

Relates to: Identifying an IP Addressing Strategy

Q7) B, D, E

Relates to: Selecting Routing Protocols

Q8) B

Relates to: Enterprise Campus Design Examples

Q9) D

Relates to: Enterprise Campus Design Examples

Designing the Server Farm

Overview

A server farm (or data center) is the controlled environment that houses enterprise servers. The data center servers support the business applications accessed by users over the corporate intranet and can be centralized or distributed, thus offering high levels of performance, scalability, and availability.

Relevance

A large-scale server farm or data center must be able to support direct end systems and server connectivity, while offering performance, scalability, and availability.

Objectives

Upon completing this lesson, you will be able to plan an effective Server Farm module design, given specific enterprise network requirements. This includes being able to meet these objectives:

- Identify typical enterprise Server Farm module design requirements
- Select the most effective Server Farm module design, given specific enterprise network requirements
- Explain design options to enhance scalability in the Server Farm module
- List considerations for Server Farm module security and manageability

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

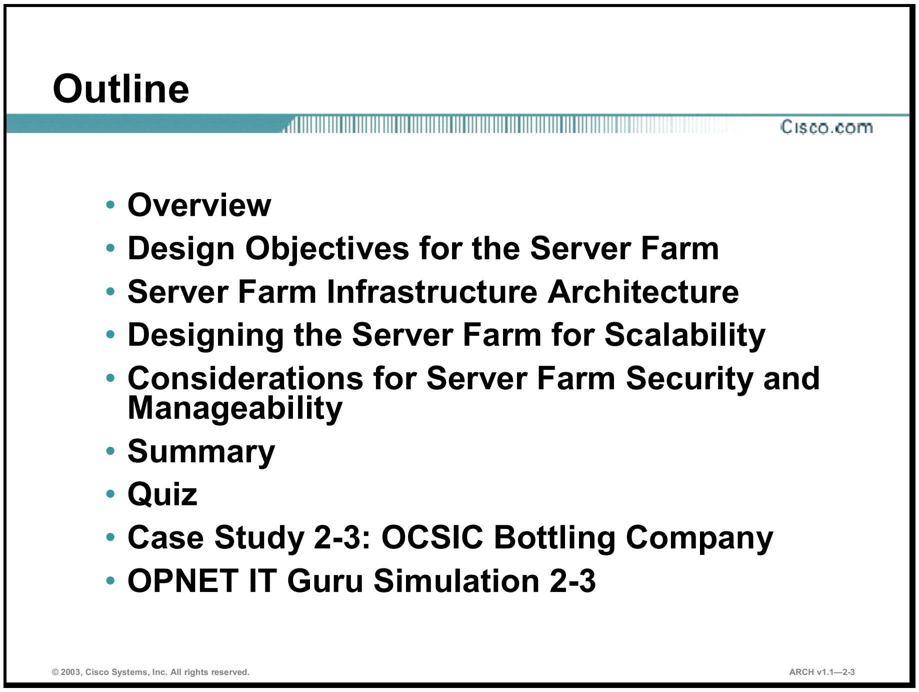

Outline

Cisco.com

- Overview
- Design Objectives for the Server Farm
- Server Farm Infrastructure Architecture
- Designing the Server Farm for Scalability
- Considerations for Server Farm Security and Manageability
- Summary
- Quiz
- Case Study 2-3: OCSIC Bottling Company
- OPNET IT Guru Simulation 2-3

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—2-3

Design Objectives for the Server Farm

The primary objectives in the design of an enterprise server farm (or data center) are performance, scalability, availability, security, and manageability. This topic identifies typical enterprise Server Farm module design requirements.

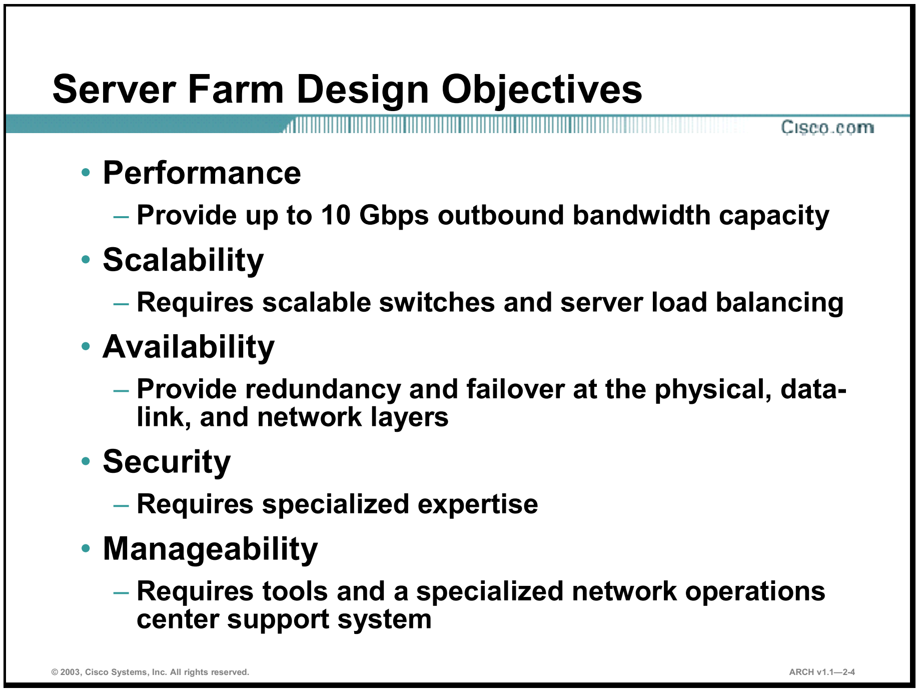

Server Farm Design Objectives

- Performance
  - Provide up to 10 Gbps outbound bandwidth capacity
- Scalability
  - Requires scalable switches and server load balancing
- Availability
  - Provide redundancy and failover at the physical, data-link, and network layers
- Security
  - Requires specialized expertise
- Manageability
  - Requires tools and a specialized network operations center support system

ARCH v1.1—2-4

A large-scale, shared server farm infrastructure must support direct end-system or server connectivity. The server farm must be able to support these properties:

- **Performance:** Up to 10 Gbps outbound bandwidth capacity is required from the server farm for most enterprises.
- **Scalability:** Scalability is a critical requirement in every server farm. Server load balancing is most often deployed. As the number of servers requiring higher-bandwidth connections increases, port densities can exceed the capacity of a single switch or server farm block. Applying a modular block design to server farm deployments permits flexible growth.
- **Availability:** Availability is generally ensured through the overall network design. Networks are designed to minimize the occurrence of service problems and the time to recover from problems, for example with backup recovery policies. You should design for high availability at each layer, with redundancy and failover provisions at the physical, data link, and network layers. The most effective solutions are those with consistent engineering considerations tightly integrated throughout the server farm.
- **Security:** Security is an integral part of the network design. A single vulnerability could compromise the enterprise. Specialized security expertise is often required. There are challenges in offering encryption, certification, directory services, network access, and other security capabilities that enable a fully secure network.

Note Security design is more fully explained in the Designing Security Service module because designing for security crosses functional areas, and even campuses or sites, in the Enterprise Composite Network Model.

- **Manageability:** Manageability means much more than knowing if a server or other network element is up or down. The ability to assess service levels on a “per user” basis is important to offering and maintaining required service levels. An operations center support system may track network configuration and application performance, and maintain and resolve errors and alarms. Good manageability tools and qualified personnel lowers operations costs and reduces wasted time, while resulting in overall higher satisfaction.

Note Designing to ensure manageability is more fully explained in the Designing Network Management Services module since this capability requires crossing functional areas, and even campuses or sites, within the Enterprise Composite Network Model.

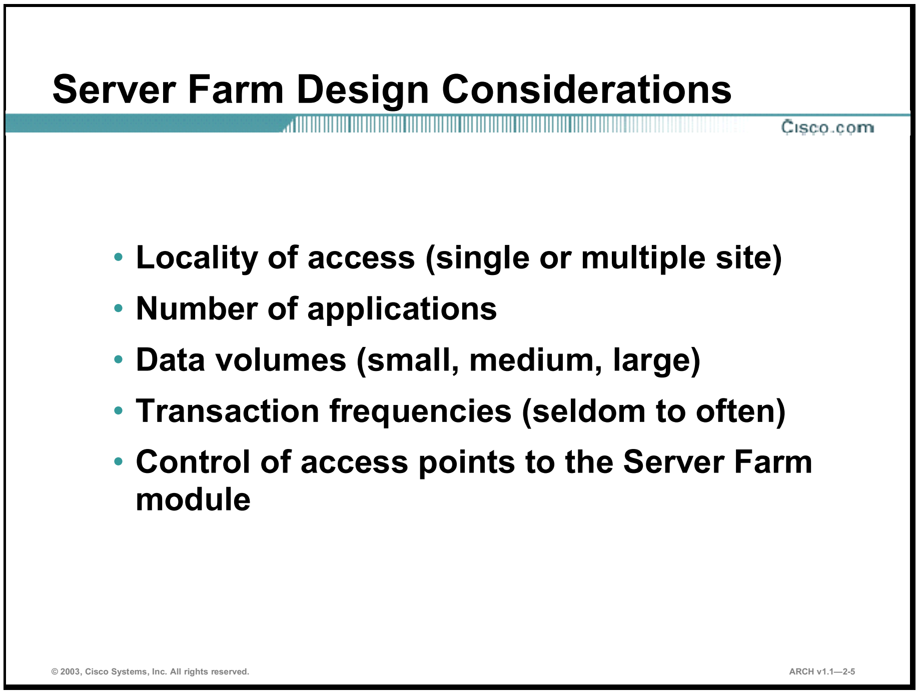

Server Farm Design Considerations

Cisco.com

- Locality of access (single or multiple site)
- Number of applications
- Data volumes (small, medium, large)
- Transaction frequencies (seldom to often)
- Control of access points to the Server Farm module

ARCH v1.1—2-5

The figure identifies additional considerations for the Server Farm module design.

Server Farm Infrastructure Architecture

The enterprise server farm infrastructure may contain an access layer and a distribution layer, similar to the Campus Infrastructure, and connected to the Campus Backbone submodule. This topic will help you select an effective Server Farm design model, given specific enterprise network requirements.

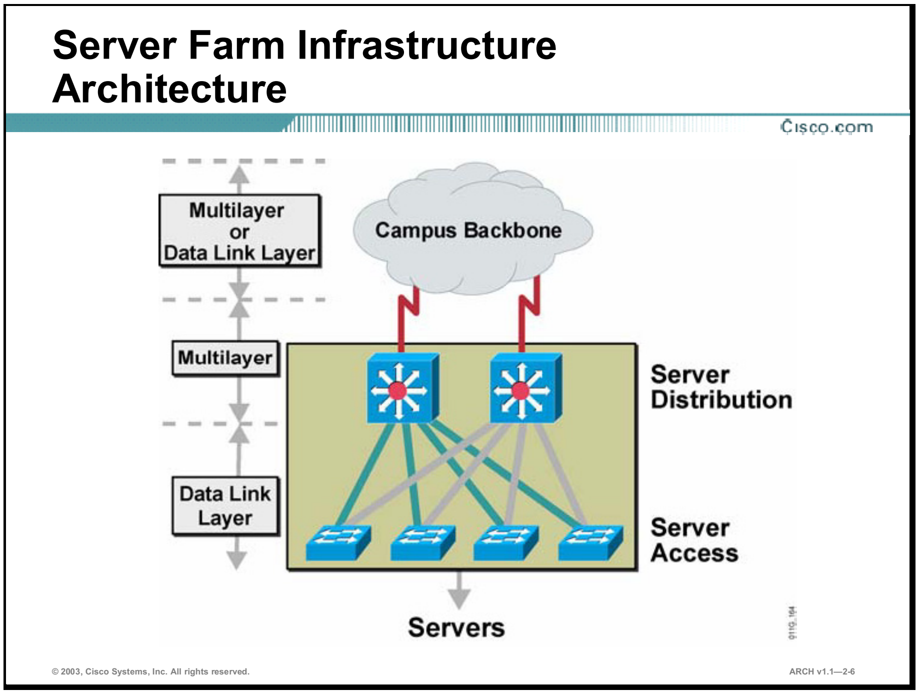

Server Farm Infrastructure Architecture

ARCH v1.1—2-6

You can logically divide the overall server farm infrastructure into three functional areas operating with different and very specific criteria:

- **Server Access layer:** The Server Access layer provides data link layer-based transport to directly connected servers. The data link layer provides flexibility and speed in provisioning. It also allows deployment of applications and systems, which may inherently expect data link layer-level connectivity.
- **Server Distribution layer:** The Server Distribution layer consists of devices in multiple sublayers, which provide transit for traffic from the data link layer into the network layer, and allow full deployment of the data link layer in the egress layer. The Server Distribution layer leverages multilayer switching scalability characteristics while benefiting from the flexibility of data link layer services.
- **Campus Backbone:** The Campus Backbone is shared between the Campus Infrastructure distribution devices, Enterprise Edge distribution devices, and the Server Farm distribution devices. It is composed of high-end switches providing network layer transport between the distribution and edge layers. Optionally, you can combine the Server Distribution and Campus Backbone layers physically and logically into a collapsed backbone to provide connectivity with the Server Access and Building Access layers.

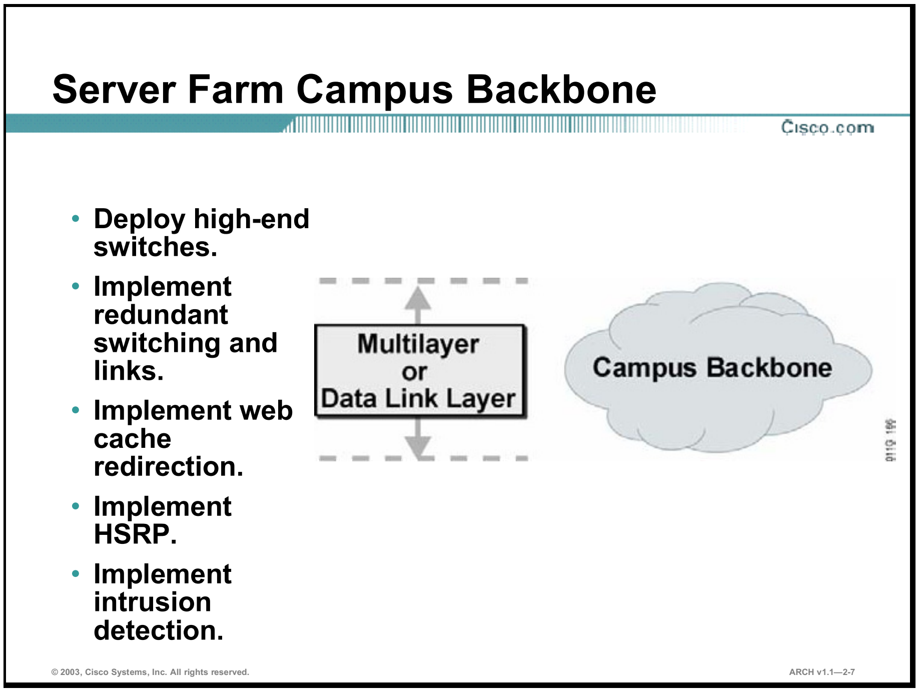

Server Farm Campus Backbone

- Deploy high-end switches.
- Implement redundant switching and links.
- Implement web cache redirection.
- Implement HSRP.
- Implement intrusion detection.

ARCH v1.1—2-7

At the Campus Backbone, consider these best practices to support the Server Farm module:

- Deploy high-end switches (such as the Catalyst 8500 or 6500 series).
- Implement highly redundant switching and links with no single points or paths of failure.
- Implement web cache redirection (using Cisco Content Networking solutions).
- Implement Hot Standby Router Protocol (HSRP) for failover protection.
- Implement intrusion detection with automatic notification of intrusion attempts in place.

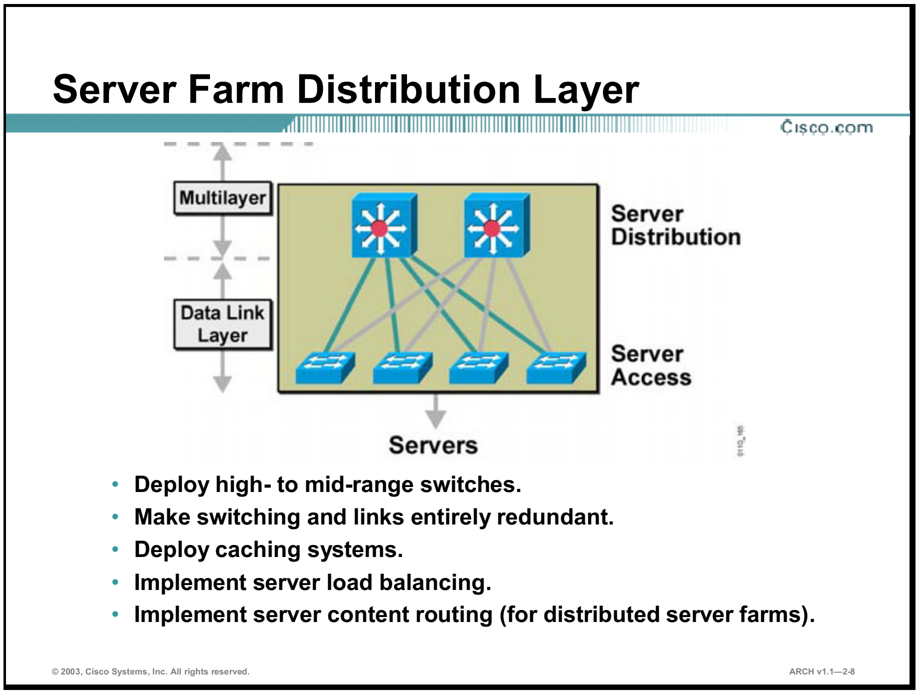

Server Farm Distribution Layer

- Deploy high- to mid-range switches.
- Make switching and links entirely redundant.
- Deploy caching systems.
- Implement server load balancing.
- Implement server content routing (for distributed server farms).

ARCH v1.1—2-8

At the Server Distribution layer, consider these best practices:

- Deploy high- to mid-range switches (such as Catalyst 6500 series).
- Implement redundant switching and links with no single points or paths of failure.
- Deploy caching systems where appropriate (using Cisco Content Networking solutions).
- Implement server load balancing (using Cisco Content Networking solutions).
- Implement server content routing (using Cisco Content Networking solutions).

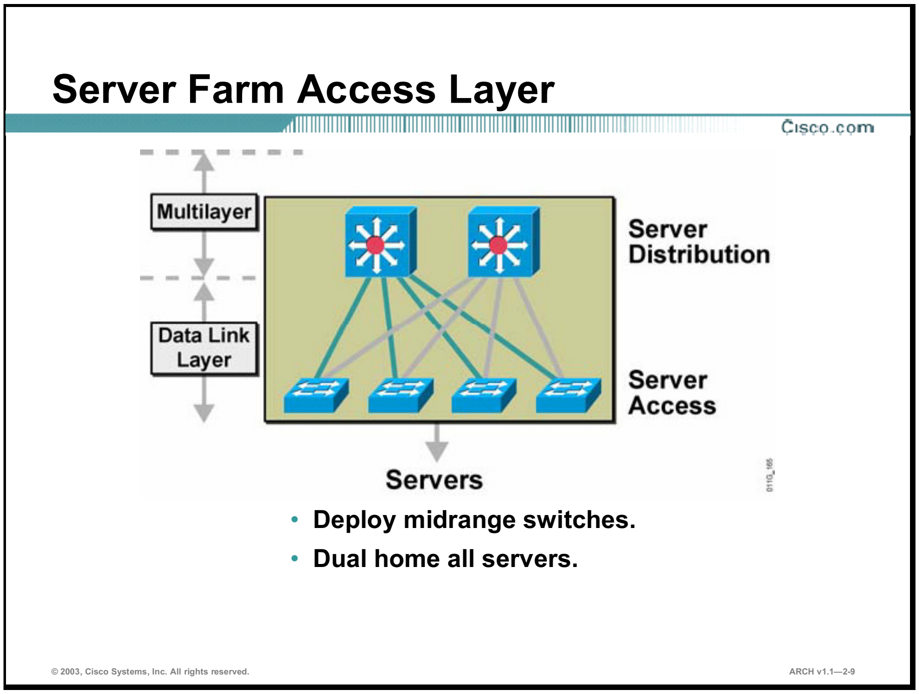

Server Farm Access Layer

Cisco.com

- Deploy midrange switches.
- Dual home all servers.

ARCH v1.1—2-9

At the Server Access layer, consider these best practices:

- Deploy mid-range switches (Catalyst 6500 or 4000 series).
- Dual home all servers.

Designing the Server Farm for Scalability

Scalability must be provided in every server farm. Scalability is provided in switches and routers, as well as with content networking solutions. This topic provides design options to enhance scalability in the server farm.

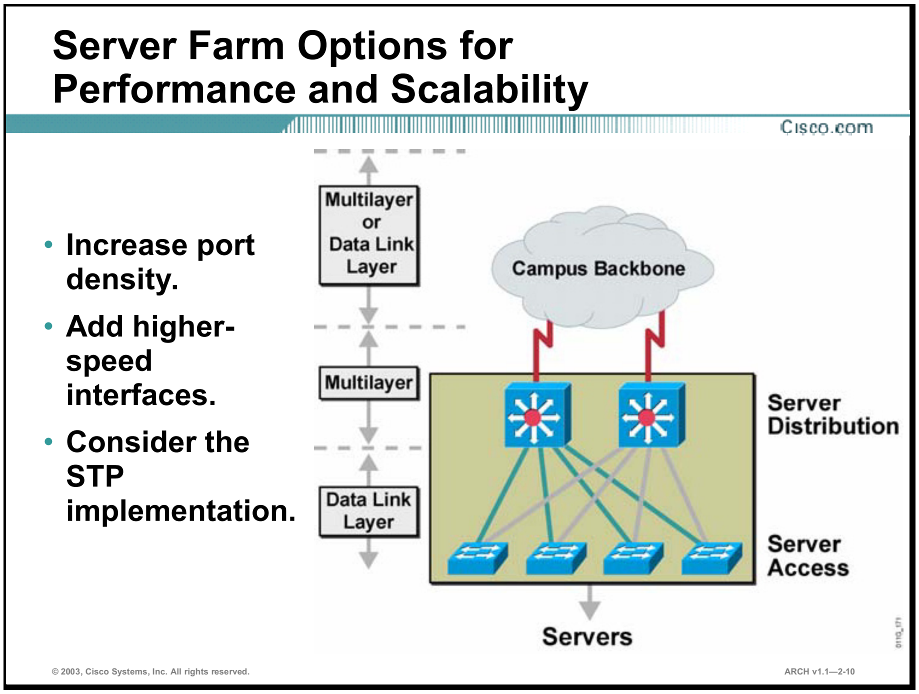

Server Farm Options for Performance and Scalability

Cisco.com

- Increase port density.
- Add higher-speed interfaces.
- Consider the STP implementation.

ARCH v1.1—2-10

The Server Farm architecture addresses scalability by providing flexible growth paths to deliver high bandwidth rates to the connected IP core. The methods used to grow the server farm capabilities include:

- **Increase port density:** You can increase raw port density for both end-user devices and infrastructure interconnecting links using a modular approach to add connectivity to the existing installation. You can consider the distribution layer and access layer switches as a module, which provides a predetermined number of data link layer ports. Expanding data link layer ports is only a matter of adding another module and updating the routing configuration to include new modules.
- **Add higher-speed interfaces:** Migrating to a higher-speed interface or EtherChannel technology is a way to deliver greater bandwidth capacity between devices.
- **Consider the spanning-tree implementation:** One of the main limiting factors in designing large Layer 2 implementations is the capacity of the system to handle and scale Spanning Tree Protocol. You can implement Multi-Instance Spanning Tree Protocol (802.1s) to reduce the total number of spanning-tree instances needed to support the infrastructure. This reduces the total number of spanning-tree instances needed to support the infrastructure.

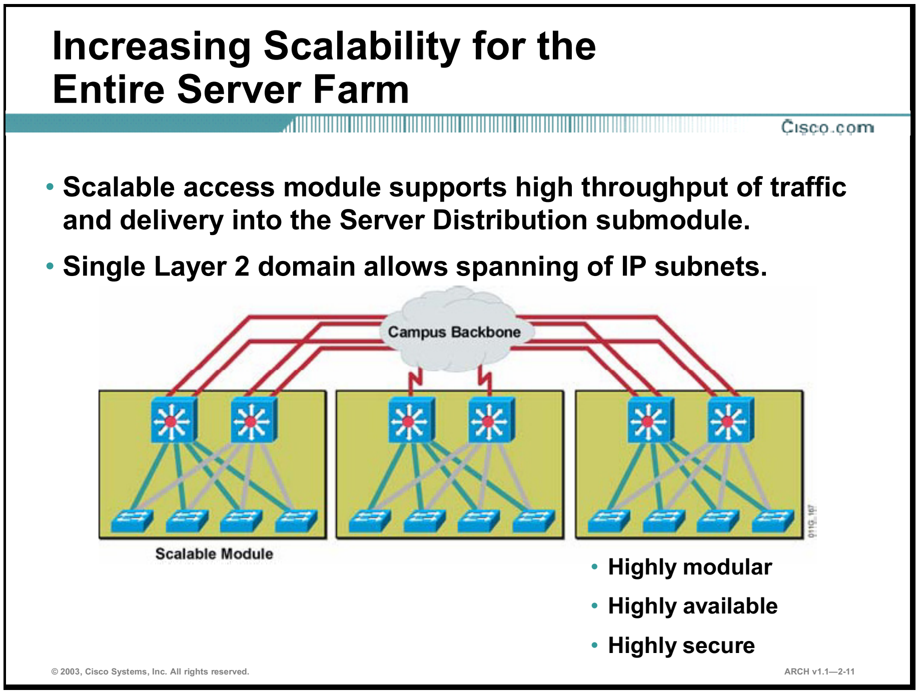

Increasing Scalability for the Entire Server Farm

Cisco.com

- Scalable access module supports high throughput of traffic and delivery into the Server Distribution submodule.
- Single Layer 2 domain allows spanning of IP subnets.

- Highly modular
- Highly available
- Highly secure

ARCH v1.1—2-11

You can increase scalability in the Server Farm module by implementing a modular design with devices that are easily upgradeable. Alternatively, you can implement a scalable Server Farm module, either at the same location or at another location.

Both the complications and available solutions increase greatly if the Server Farm module is part of a geographically dispersed set. A geographically dispersed Server Farm module allows content to be served closer to the requesting client. However, the multilocation Server Farm design must consider management of content, distribution of updates, synchronization of different sources, proper routing of requests, handling of downed servers, additional security, and so on.

Considerations for Server Farm Security and Manageability

The Server Farm design has important security and manageability considerations. This topic lists considerations for Server Farm module security and manageability.

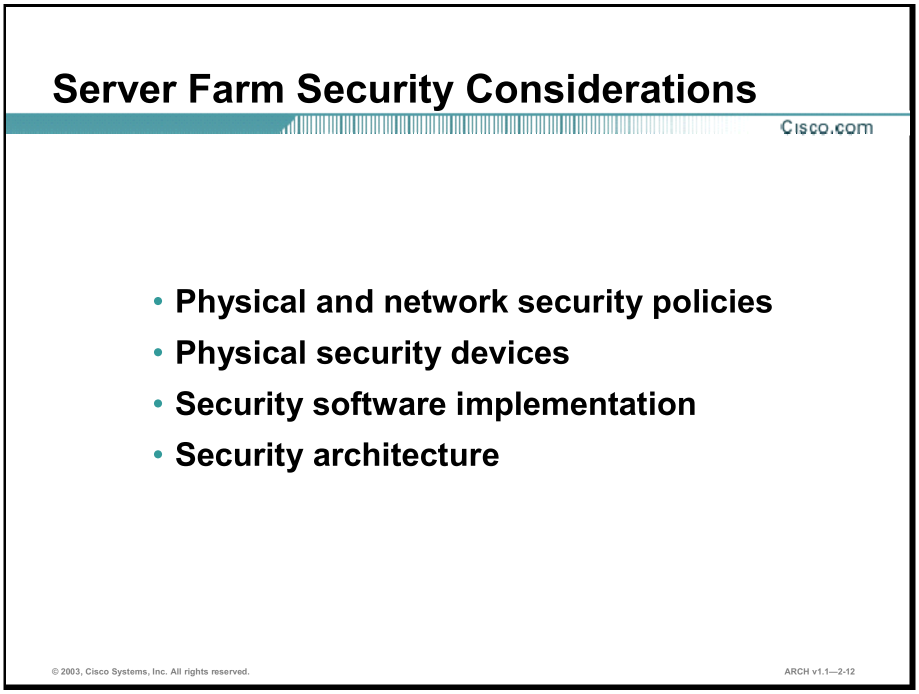

Server Farm Security Considerations

- Physical and network security policies
- Physical security devices
- Security software implementation
- Security architecture

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—2-12

To meet Server Farm module requirements for security, consider the items listed in the figure.

Note	Security is covered in detail in the Designing Security Services module in this course.

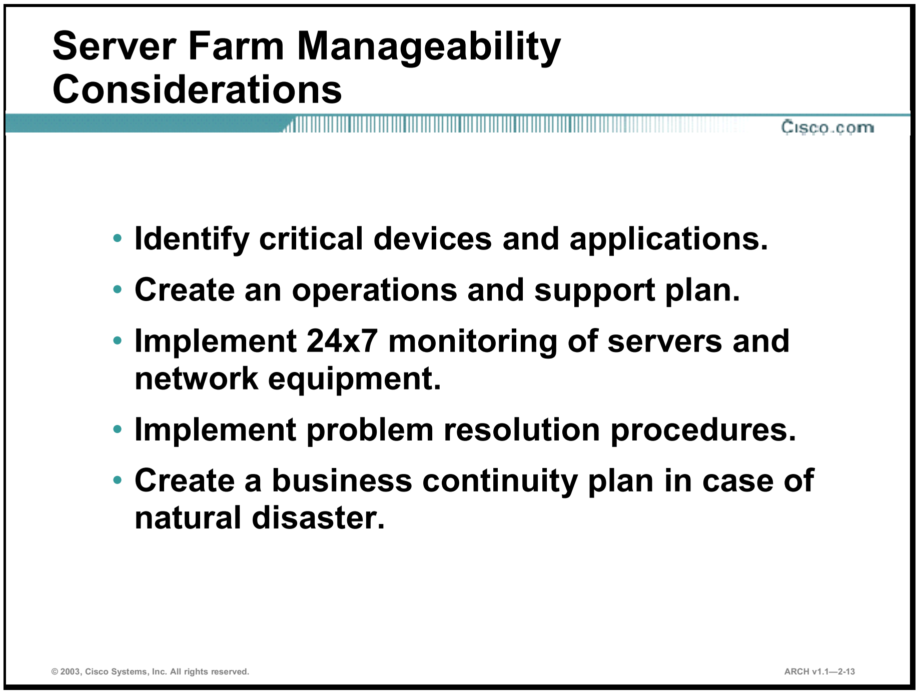

Server Farm Manageability Considerations

- Identify critical devices and applications.
- Create an operations and support plan.
- Implement 24x7 monitoring of servers and network equipment.
- Implement problem resolution procedures.
- Create a business continuity plan in case of natural disaster.

ARCH v1.1—2-13

Good manageability tools and qualified personnel supporting the infrastructure results in lower operations costs, since time is not wasted trying to resolve indications from conflicting management systems and higher user satisfaction. To meet Server Farm module manageability requirements, consider the items listed in the figure.

Note Manageability is covered in detail in the Designing Network Management Services module in this course.

Summary

This topic summarizes the key points discussed in this lesson.

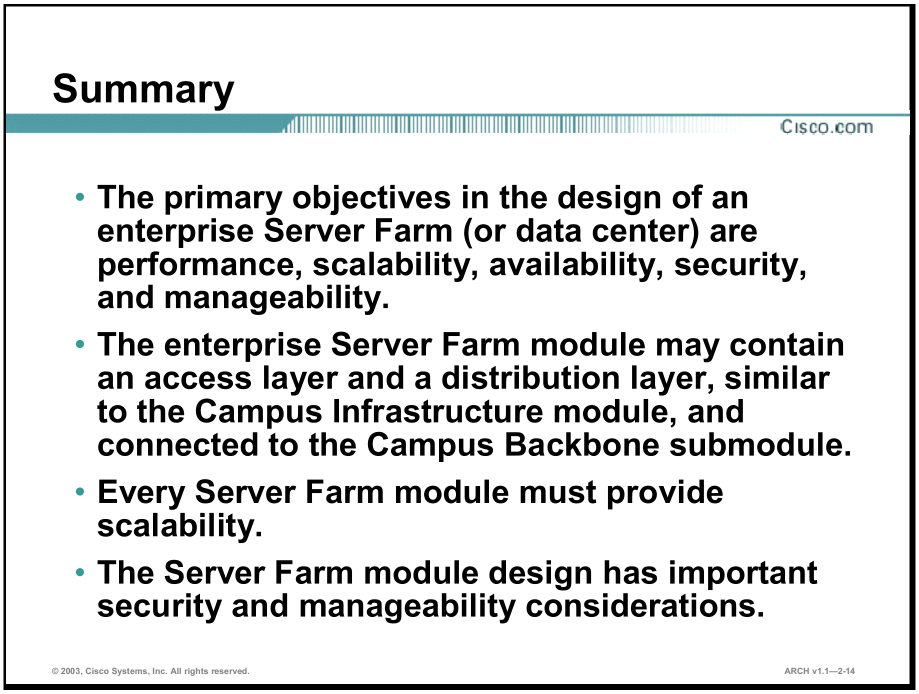

Summary

- The primary objectives in the design of an enterprise Server Farm (or data center) are performance, scalability, availability, security, and manageability.
- The enterprise Server Farm module may contain an access layer and a distribution layer, similar to the Campus Infrastructure module, and connected to the Campus Backbone submodule.
- Every Server Farm module must provide scalability.
- The Server Farm module design has important security and manageability considerations.

ARCH v1.1—2-14

References

For additional information, refer to these resources:

- *Gigabit Campus Network Design—Principles and Architecture* at http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.htm
- *Data Centers: Best Practices for Security and Performance* at http://www.cisco.com/warp/public/cc/so/neso/wnso/power/gdmdd_wp.pdf
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 2-3: OCSIC Bottling Company
- OPNET IT Guru Simulation 2-3

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) Server load balancing is a key feature that primarily contributes to _____.

A) security

B) scalability

C) functionality

D) manageability

Q2) Which three recommendations would you implement at the Server Distribution layer? (Choose three.)

A) Dual home all servers.

B) Implement server load balancing.

C) Deploy caching systems where appropriate.

D) Deploy mid-range switches (Catalyst 6500 or 4000 series).

E) Implement redundant switching and links with no single points or paths of failure.

Q3) Which four recommendations would you implement at the server farm core layer? (Choose four.)

A) Dual home all servers

B) Deploy caching systems, where appropriate

C) Deploy midrange switches (Catalyst 6500 or 4000 series)

D) Deploy high-end switches (such as the Catalyst 8500 or 6500 series)

E) Implement HSRP for failover protection

F) Implement web cache redirection (using Cisco Content Networking solutions)

G) Implement highly redundant switching and links with no single points or paths of failure

Q4) What is one method to increase scalability in the Server Farm module, while also adding complexity?

A) higher-speed interfaces

B) modules with more port density

C) geographically dispersed Server Farms

D) Multi-Instance Spanning Tree Protocol (MISTP)

Q5) Good manageability tools and _____ supporting the infrastructure results in lower operations costs, since time is not wasted trying to resolve indications from conflicting management systems and higher user satisfaction.

A) firewalls

B) qualified personnel

C) fault management systems

D) configuration control systems

Quiz Answer Key

Q1) B

Relates to: Design Objectives for the Server Farm

Q2) B, C, E

Relates to: Server Farm Infrastructure Architecture

Q3) D, E, F, G

Relates to: Server Farm Infrastructure Architecture

Q4) C

Relates to: Designing the Server Farm for Scalability

Q5) B

Relates to: Considerations for Server Farm Security and Manageability

Case Study 2-3: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

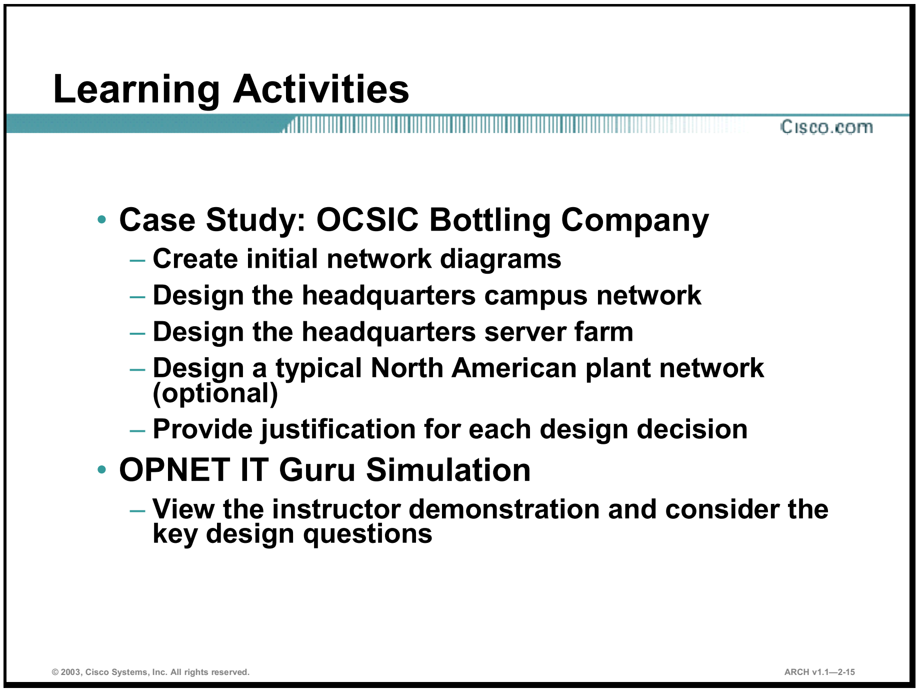

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Create initial network diagrams
  - Design the headquarters campus network
  - Design the headquarters server farm
  - Design a typical North American plant network (optional)
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—2-15

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

In this exercise, you will design a campus network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Create initial network diagrams
- Design the headquarters campus network
- Design the headquarters server farm
- Design a typical North American plant network (optional)

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

About the OCSIC Bottling Company

Company Background

OCSIC Bottling Company is a U.S.-based soft drink and beverage distributor. OCSIC's products include colas, flavored soft drinks, and bottled juices. Most of their products, particularly their soft drinks, are recognized throughout the world. Other products, such as juices, are currently available only in selected regions, although the company would like to distribute them worldwide.

OCSIC owns and operates 12 bottling, distribution, and sales facilities in North America. In South America, Europe, and the Asia Pacific regions, their products are sold, manufactured, and distributed by independently owned and operated distribution companies. These international distribution companies can sell any products they choose, and often do not carry the entire line of OCSIC products.

Company Business Goals

The company is completing a ten-year plan to move its business into the next century. They are concerned about losing their competitiveness if they do not innovate and become more responsive to their customers and their distributors. Through a series of planning sessions, the company has defined these goals:

- Develop and bring new products to market more quickly
- Provide faster order fulfillment to customers
- Increase product distribution outside the United States
- Improve communications between employees, customers, and partners
- Implement a supply chain management system that provides better integration and responsiveness to the plants and distributors
- Reduce operating costs and increase profitability

Headquarters Location

The facilities in North America consist of one main headquarters campus located in Memphis, Tennessee, and 12 wholly owned bottling, distribution, and sales facilities in the United States and Canada.

The table describes each building in the headquarters' campus. The campus has developed over a long period of time, and users continually move around the campus.

Building	Function	Estimated Number of Users	Building Characteristics
Building A	Executive Finance	1400	six-story building 210,000 square feet
Building B	Accounting Order Processing	600	two-story building 90,000 square feet
Building C	Marketing Distributor Relations	1,200	four-story building 180,000 square feet
Building D	Research and Development	800	three-story building 120,000 square feet
Building E	Research and Development	900	four-story building 120,000 square feet
Building F	Information Systems Data Center Help Desk	400	two-story building 80,000 square feet

Headquarters Network

The headquarters location has an aging Token Ring network, which supports traditional applications such as database, file and print sharing, and e-mail. Over the past ten years, the company has added IP capabilities to its existing network, but is finding the network increasingly difficult and expensive to maintain.

The company now wants to "move into the 21st century," replacing its campus network with a complete IP-based solution.

The following figure describes the headquarters campus and its network.

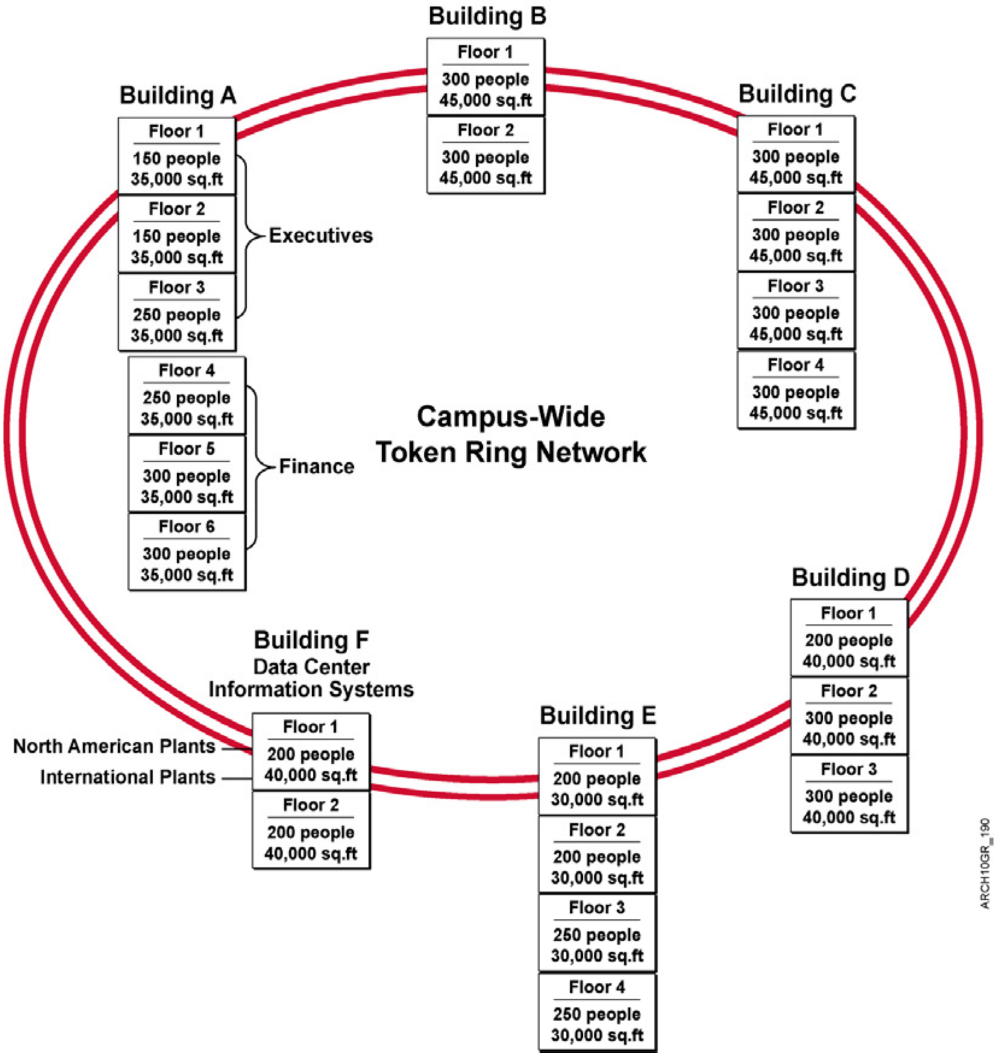

The cabling plant on the headquarters Token Ring network backbone and risers consists of copper using shielded twisted-pair wiring. Copper unshielded twisted-pair wiring goes to the desktops. Fiber is not currently used anywhere in the network.

Headquarters Applications

The IT department at OCSIC develops and maintains a variety of enterprise-wide applications. Some are available only at headquarters, while others are available to the North American plants and international manufacturing, distribution, and sales centers.

The table describes the primary applications that the company uses.

Application	Data Characteristics[1]	Primary Users	Notes and Description
SAP	Heavy	Accounting Finance Marketing Manufacturing Distribution Order Processing	SAP is used throughout the company to manage manufacturing, inventory, distribution, and order processing.
PeopleSoft	Moderate	Everyone	PeopleSoft applications are used for financial management and reporting throughout the company.
Custom Oracle database applications	Moderate	Everyone	The company has developed a number of custom Oracle database applications that are used throughout the company, primarily for reporting and decision support.
Electronic mail	Moderate	Everyone	E-mail is used as the primary means of electronic communication throughout the company. E-mail messages consist primarily of text, but some users send rich e-mail with graphics and video.
Intranet web site	Moderate	Everyone	The company maintains an intranet web site that provides up-to-date, corporate-wide information to the employees.
Extranet web site (planned)	Moderate	International distributors	The company wants to add an extranet web site that provides up-to-date information for the international distributors. Information to be contained on the extranet web site includes: ■ Marketing information, such as product data ■ Order status information ■ Sales data ■ Inventory data

[1] A light application generates hundreds of bytes of data per minute. A moderate application generates thousands of bytes of data per minute. A heavy application generates tens of thousands of bytes of data per minute. An extremely heavy application, such as near broadcast-quality video, generates hundreds of thousands of bytes per minute.
Application activity is spread across five to seven applications, and includes network activity only about ten percent of the time. The remainder is dedicated to server activity, workstation interaction, thought research, planning, and workgroup interaction.

The table describes the traffic volume on the network from each building to the data center in Building F.

Building From	**Building To**	**Volume (kbps)** **(to nearest 100 kbps)**
Building A	Building F	1900 kbps
Building B	Building F	1200 kbps
Building C	Building F	2300 kbps
Building D	Building F	500 kbps
Building E	Building F	500 kbps
Building F	Building F	300 kbps

North American Plants

Each district and local building facility is similar. They each include a manufacturing floor, warehouse facilities, distribution and logistics offices, and sales offices. The district office facilities have slightly more staff to handle the administrative and logistical functions of their district.

The table describes each district and local facility in North America.

Function	Location	Estimated Number of Users	Building Characteristics
Eastern District Office/Plant	Boston, MA	175	60,000-square-foot manufacturing, distribution, and sales office building
Midwestern District Office/Plant	Kansas City, MO	175	60,000-square-foot manufacturing, distribution, and sales office building
Southern District Office/Plant	Dallas, TX	175	60,000-square-foot manufacturing, distribution, and sales office building
Western District Office/Plant	Los Angeles, CA	175	60,000-square-foot manufacturing, distribution, and sales office building
New York Regional Office/Plant	New York, NY	150	50,000-square-foot manufacturing, distribution, and sales office building
Toronto Regional Office/Plant	Toronto, Canada	150	50,000-square-foot manufacturing, distribution, and sales office building
Chicago Regional Office/Plant	Chicago, IL	150	50,000-square-foot manufacturing, distribution, and sales office building
Omaha Regional Office/Plant	Omaha, NB	150	50,000-square-foot manufacturing, distribution, and sales office building
Orlando Regional Office/Plant	Orlando, FL	150	50,000-square-foot manufacturing, distribution, and sales office building
Denver Regional Office/Plant	Denver, CO	150	50,000-square-foot manufacturing, distribution, and sales office building
San Francisco Regional Office/Plant	San Francisco, CA	150	50,000-square-foot manufacturing, distribution, and sales office building
Seattle Regional Office/Plant	Seattle, WA	150	50,000-square-foot manufacturing, distribution, and sales office building

North American Plant Applications

The North American plants use the same applications as headquarters, although their usage varies. The data center is located at the Memphis headquarters, so each plant accesses the headquarters applications over the network on a regular basis.

The table describes the current and planned applications at the North American plants.

Application	Data Characteristics	Primary Users	Notes and Description
SAP	Moderate	Accounting Finance Marketing Manufacturing Distribution	SAP is used in the plants to manage manufacturing, inventory, and distribution.
PeopleSoft	Light	Everyone	PeopleSoft applications are used to obtain financial reports.
Custom Oracle database applications	Moderate	Everyone	The plants access custom Oracle database applications, primarily for reporting and decision support.
Electronic mail	Heavy	Everyone	E-mail is used as the primary means of electronic communication throughout the company. Most e-mail messages are sent within a plant, with lighter traffic to the headquarters.
Intranet web site	Heavy	Everyone	The company maintains an intranet web site that provides up-to-date, corporate-wide information to the employees in the plants, who would not be able to obtain information otherwise. Information contained on the intranet web site includes: ■ Forms used for human resources functions ■ Accounting functions such as purchase requests and supply ordering ■ Marketing information such as product data ■ IT application and help desk information

Networking Strategy and Goals

To better support the overall business goals and reduce costs, OCSIC is developing an integrated information systems project plan that includes six building blocks or components:

- Replacement of older, slower PCs with new, faster personal computer workstations while maintaining the current workload
- Implementation of advanced network solutions, such as IP telephony
- Implementation of a corporate intranet and extranet that better serves employees, customers, and partners
- Replacement of the existing campus and plant networks
- Upgrade to the WAN
- Streamline operations and lower total costs through business process reengineering

Task 1: Create Initial Network Diagrams

Complete these steps:

Step 1 On an overhead transparency, create a global network diagram for the company, to include the headquarters location, district offices, regional offices, and international plants. This network diagram should show the main sites in each country and the main WAN links. Label each location.

Step 2 On an overhead transparency, create a country-level network diagram for the company that identifies the locations in North America. This network diagram shows the main sites in each town and the main WAN links. Label each location.

Task 2: Design the Headquarters Campus Network

Complete these steps:

Step 1 On an overhead transparency, create a campus network diagram for the company headquarters site. If desired, create a logical map showing the extent of each VLAN for the headquarters site. Your network diagram should include each building and the Campus Backbone submodule. Label each location.

Step 2 Complete the table to design the details about your headquarters campus network.

Design Question	Decision	Justification
What is the logical network design?		
What type of VLAN trunking will be used?		
What physical network media will be used in the Campus Backbone submodule?		
What physical network media will be used in the Building Distribution submodule?		
What physical network media will be used in the Building Access submodule?		
Which data link layer protocol will be used at each location?		
What spanning-tree deployment and version will be used?		

Design Question	Decision	Justification
What is the data link layer/multilayer strategy for the Campus Backbone submodule?		
What is the data link layer/multilayer strategy for the Building Distribution submodule?		
What is the data link layer/multilayer strategy for the Building Access submodule?		
Which Cisco products and options will be used in the Campus Backbone submodule?		
Which Cisco products and options will be used in the Building Distribution submodule?		
Which Cisco products and options will be used in the Building Access submodule?		
What IP addressing scheme will be used? Is NAT/PAT required?		
Which routing protocols will be used in each area of the network?		
What type of switching will be deployed at the Edge Distribution module?		

Task 3: Design the Headquarters Server Farm

Complete these steps:

Step 1 Create a Server Farm network diagram for the OCSIC Bottling Company data center. This network diagram shows the physical layout and how the data center relates to the Campus Backbone module. Label each device and location.

Step 2 Complete the table to design the details about your Server Farm network.

Design Question	Decision	Justification
What is the logical network design?		
What type of VLAN trunking will be used?		
What physical network media will be used?		
What data-link layer protocol will be used?		
What spanning-tree deployment will be used?		
What is the data link layer/multilayer strategy for the Server Distribution layer?		
What is the data link layer/multilayer strategy for the Server Access layer?		
Which Cisco products and options will be used in the Server Distribution layer?		

Design Question	Decision	Justification
Which Cisco products and options will be used in the Server Access layer?		
What IP addressing scheme will be used? Is NAT/PAT required?		
Which routing protocols will be used?		

Task 4: Design a Typical North American Plant Network (Optional)

Complete these steps:

Step 1 Create a campus network diagram for a typical plant in North America. This network diagram shows the physical layout and the Campus Backbone module. Label each location or area of the building.

Step 2 Complete the table to design the details about a typical North American plant network.

Design Question	Decision	Justification
What is the logical network design?		
What type of VLAN trunking will be used?		
What physical network media will be used in the Campus Backbone submodule?		
What physical network media will be used in the Building Distribution submodule?		
What physical network media will be used in the Building Access submodule?		
What data-link layer protocol will be used?		
What spanning-tree deployment will be used?		

Design Question	Decision	Justification
What is the data link layer/multilayer strategy for the Campus Backbone submodule?		
What is the data link layer/multilayer strategy for the Building Distribution submodule?		
What is the data link layer/multilayer strategy for the Building Access submodule?		
Which Cisco products and options will be used in the Campus Backbone submodule?		
Which Cisco products and options will be used in the Building Distribution submodule?		
Which Cisco products and options will be used in the Building Access submodule?		
What IP addressing scheme will be used? Is NAT/PAT required?		
Which routing protocols will be used?		

Task 5: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- A completed network design for the OCSIC Bottling Company headquarters location that includes a logical network diagram (VLANs), physical network diagram, network devices, IP addressing strategy, and routing protocols
- A completed network design for the OCSIC Bottling Company server farm that includes a logical network diagram (VLANs), physical network diagram, network devices, IP addressing strategy, and routing protocols
- (Optional) A completed network design for one OCSIC Bottling Company North American plant that includes a logical network diagram (VLANs), physical network diagram, network devices, IP addressing strategy, and routing protocols

OPNET IT Guru Simulation 2-3

OPNET Technologies, Inc. provides intelligent network management software that enables enterprises to optimize the performance, cost, and availability of networks and applications. The OPNET IT Guru provides intelligent network management, empowering enterprise IT managers to diagnose application performance problems, validate changes to server and router configurations, and plan for growth and high availability.

Your instructor will demonstrate a series of OPNET IT Guru simulations based on the case study exercise used throughout this class.

In this simulation, you will see three different campus designs for the OCSIC Bottling Company, as follows:

- The first simulation demonstrates a minimal design.
- The second simulation demonstrates an optimal design.
- The third simulation demonstrates a high-end design.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your campus network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 3

Designing Enterprise Edge Connectivity

Overview

Enterprises commonly use WANs, on-demand connections, and the Internet to build an intranet (site-to-site) between corporate offices, connect with customers and business suppliers over the Internet, conduct electronic commerce, and provide remote-access capabilities to their partners and employees.

Module Objectives

Upon completing this module, you will be able to design enterprise edge network infrastructures for effective functionality, performance, scalability, and availability, given specified enterprise network needs.

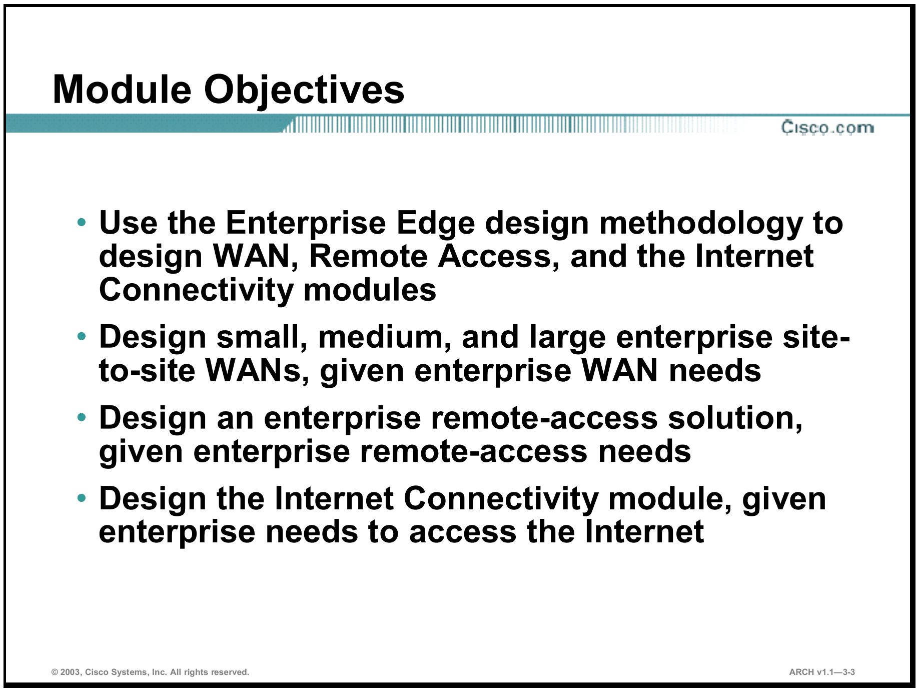

Module Objectives

- Use the Enterprise Edge design methodology to design WAN, Remote Access, and the Internet Connectivity modules
- Design small, medium, and large enterprise site-to-site WANs, given enterprise WAN needs
- Design an enterprise remote-access solution, given enterprise remote-access needs
- Design the Internet Connectivity module, given enterprise needs to access the Internet

ARCH v1.1—3-3

Module Outline

The outline lists the components of this module.

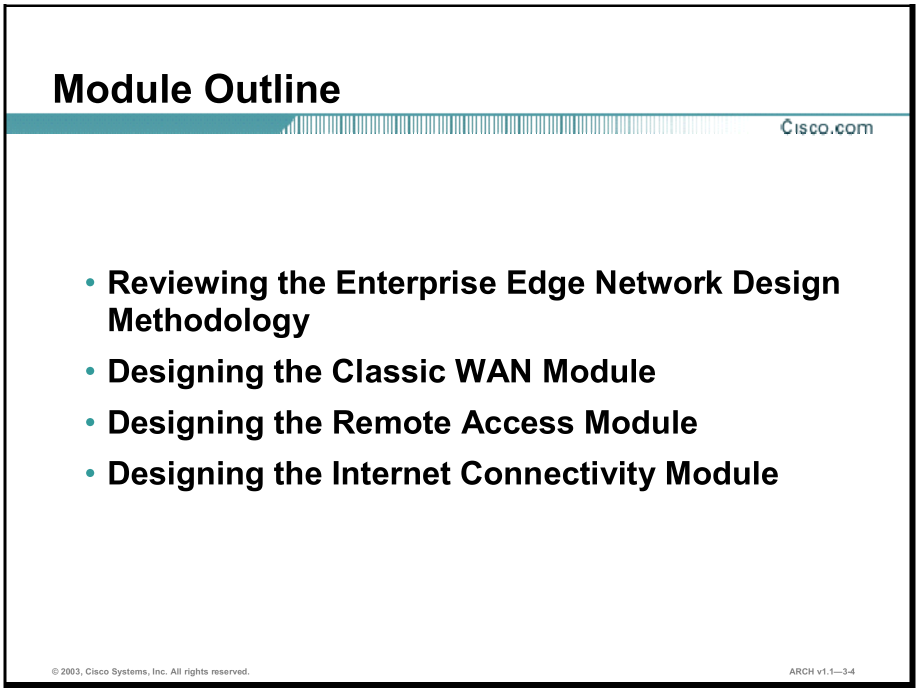

Module Outline

Cisco.com

- Reviewing the Enterprise Edge Network Design Methodology
- Designing the Classic WAN Module
- Designing the Remote Access Module
- Designing the Internet Connectivity Module

ARCH v1.1—3-4

Reviewing the Enterprise Edge Network Design Methodology

Overview

To facilitate effective network design, Cisco has developed a process that enables the network designer to assess requirements, design each module of the network, and determine the effectiveness of the design.

Relevance

The Enterprise Composite Network Model enables network designers to design the Enterprise Edge functional area from modular building blocks, which are scalable enough to meet evolving business needs. By deploying a step-by-step methodology, network designers can create an effective enterprise edge design that meets enterprise needs for performance, scalability, and availability.

Objectives

Upon completing this lesson, you will be able to use the Enterprise Edge design methodology to design WAN, Remote Access, and Internet Connectivity modules. This includes being able to meet these objectives:

- Identify the modules that the network designers will design for the Enterprise Edge functional area
- Identify the performance, scalability, and availability design considerations for the Enterprise Edge functional area of the Enterprise Composite Network Model
- Describe a step-by-step methodology that network designers will use to design the Enterprise Edge functional area
- Access and analyze network traffic patterns typically found in enterprise edge networks

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

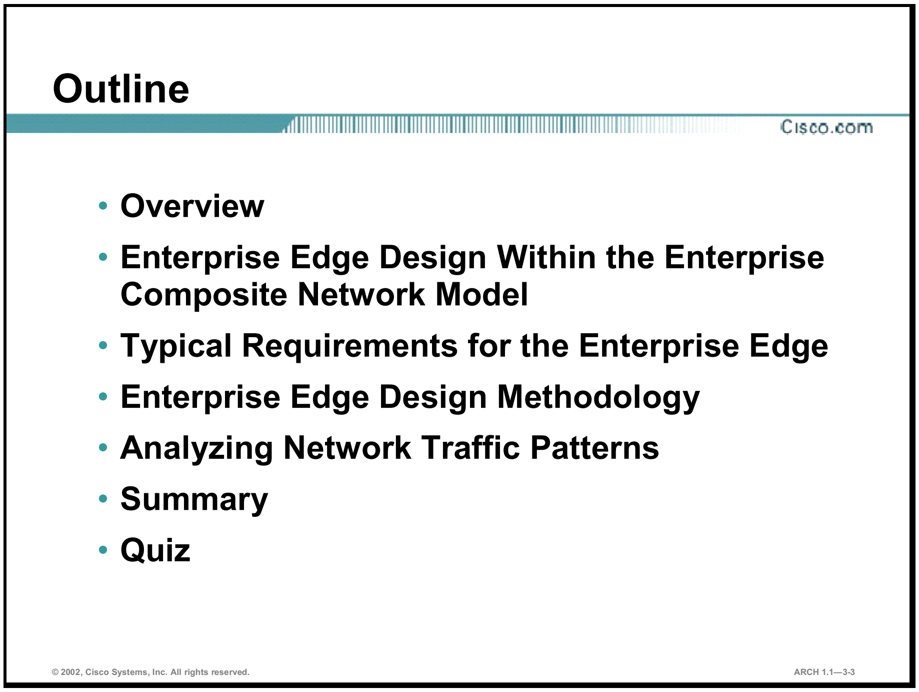

Outline

- Overview
- Enterprise Edge Design Within the Enterprise Composite Network Model
- Typical Requirements for the Enterprise Edge
- Enterprise Edge Design Methodology
- Analyzing Network Traffic Patterns
- Summary
- Quiz

ARCH 1.1—3-3

Enterprise Edge Design Within the Enterprise Composite Model

The Enterprise Edge functional area contains the E-Commerce, Internet Connectivity, VPN/Remote Access, and WAN modules. Each module has unique design requirements. This topic identifies the modules and submodules that the network designer will design for the Enterprise Edge.

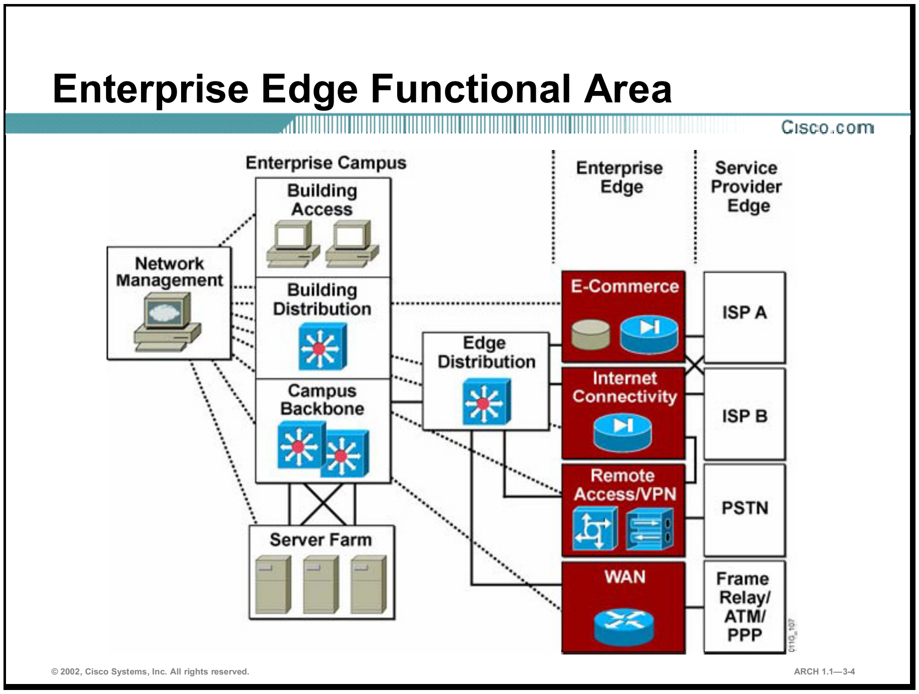

Enterprise Edge Functional Area

Cisco.com

© 2002, Cisco Systems, Inc. All rights reserved.

ARCH 1.1—3-4

The Enterprise Edge functional area is comprised of four modules:

- E-Commerce
- Internet Connectivity
- Remote Access and VPN
- WAN

Each module connects to the Edge Distribution module in the Enterprise Campus functional area, which connects the Enterprise Edge and the Enterprise Campus functional areas.

The Enterprise Edge modules perform these functions:

- **E-Commerce:** Enables enterprises to deploy e-commerce applications and take advantage of the Internet. All e-commerce transactions pass through a series of intelligent services to provide performance, scalability, and availability within the overall e-commerce network design.
- **Internet Connectivity:** Provides internal users with connectivity to Internet services. Internet users can access the information on publicly available servers. Additionally, this module accepts Virtual Private Network (VPN) traffic from remote users and remote sites and forwards it to the Remote Access and VPN module.

Note VPN design is discussed in the Designing Virtual Private Networks module in conjunction with security, a required concern when implementing VPNs over a public network.

- **Remote Access and VPN:** This 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. It also terminates VPN traffic forwarded by the Internet Connectivity module from remote users and remote sites, and initiates VPN connections to remote sites through the Internet Connectivity module.
- **WAN:** Routes traffic between remote sites and the central site using dedicated media or circuits. The WAN module supports any WAN technology, including leased lines, Frame Relay, ATM, optical, cable, and wireless. It may also use PSTN dial-on-demand for occasional access and availability.

Note This module focuses on designing the WAN, Remote Access, and Internet Connectivity modules.

Typical Requirements for the Enterprise Edge

The Enterprise Edge module requires functionality, performance, scalability, availability, manageability, and cost-effectiveness. This topic identifies the features and functionality to consider when designing the Enterprise Edge functional area.

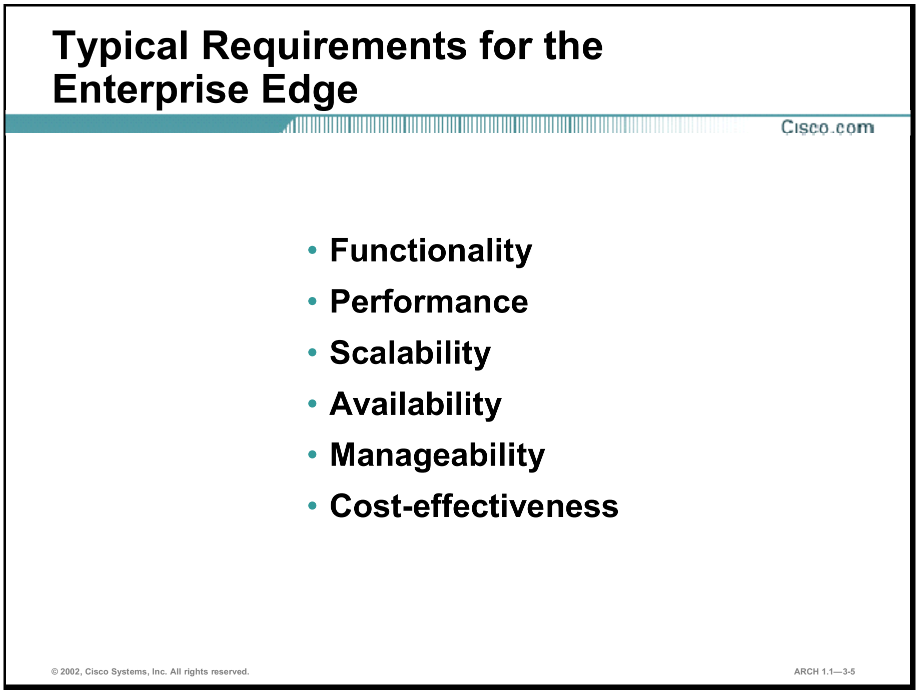

Typical Requirements for the Enterprise Edge

- Functionality
- Performance
- Scalability
- Availability
- Manageability
- Cost-effectiveness

ARCH 1.1—3-5

The Enterprise Edge must meet these requirements:

- **Functionality:** The enterprise network must support the applications and data flows required, within the required time frames. Typical enterprise-wide applications include online transaction processing (OLTP) systems, decision support systems (DSS), e-mail, information sharing, and many other functions. Applications and data may require special peak-time processing, or they may require steady processing throughout a day.
- **Performance:** Performance includes three primary metrics: responsiveness, throughput, and utilization. Each Enterprise Edge link and device will be measured in terms of how well it meets all three performance metrics.
- **Scalability:** The Enterprise Edge functional area must provide scalability for future growth in the number of users and in the amount of data and applications that the network may support.
- **Availability:** Users perceive that the network is down, regardless of where a failure may occur. A typical standard for most enterprise data networks is 99.9 percent availability.
- **Manageability:** The Enterprise Edge module must be manageable across the entire infrastructure.
- **Cost-effectiveness:** Cost-effectiveness is a key concern for most enterprises, given limited budgets. The network designer’s goal is to design the network for maximum efficiency given affordability limitations. Affordability for the Enterprise Edge functional area includes one-time costs for equipment, as well as ongoing tariffs or service charges.

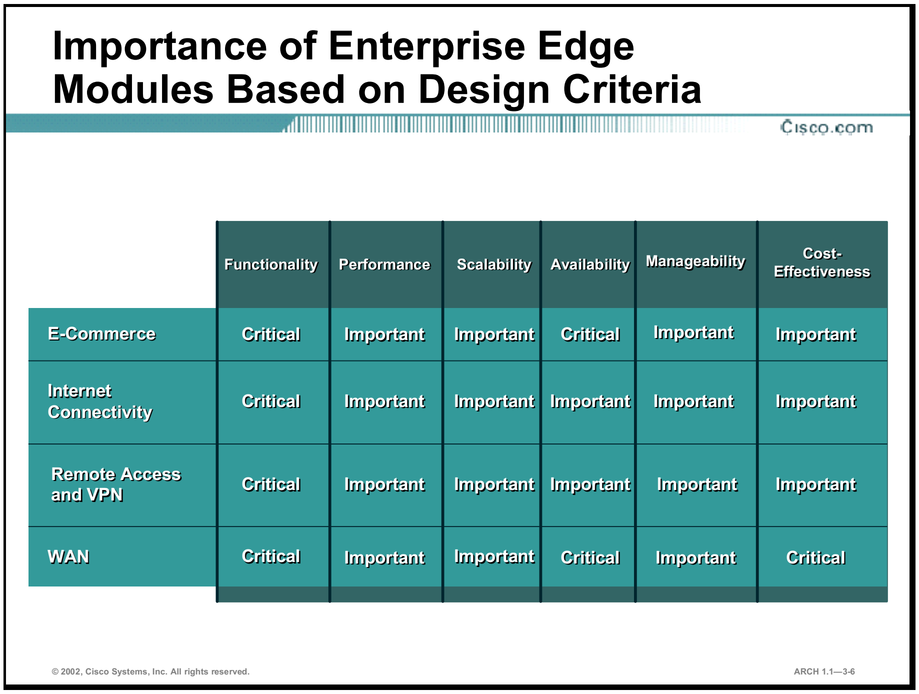

Importance of Enterprise Edge Modules Based on Design Criteria

	Functionality	Performance	Scalability	Availability	Manageability	Cost-Effectiveness
E-Commerce	Critical	Important	Important	Critical	Important	Important
Internet Connectivity	Critical	Important	Important	Important	Important	Important
Remote Access and VPN	Critical	Important	Important	Important	Important	Important
WAN	Critical	Important	Important	Critical	Important	Critical

The figure describes which Enterprise Edge modules meet enterprise needs for functionality, performance, scalability, availability, manageability, and cost-effectiveness, and the importance of each component in meeting that need. Each need is ranked in terms of its relative importance in the campus network, where Critical is highest in relative importance, followed by Important and Normal. For example, functionality is critical (that is, absolutely required) to the E-Commerce module, while scalability is important (desirable) to the E-Commerce module.

Enterprise Edge Design Methodology

Cisco has developed a methodology that you can use to design the Enterprise Edge functional area. The design is based on application characteristics, and involves selecting topology, an ISP, data link layer and physical layer technologies, features, specific Cisco devices, and routing protocols. This topic describes a step-by-step methodology that network designers will use to design the Enterprise Edge functional area.

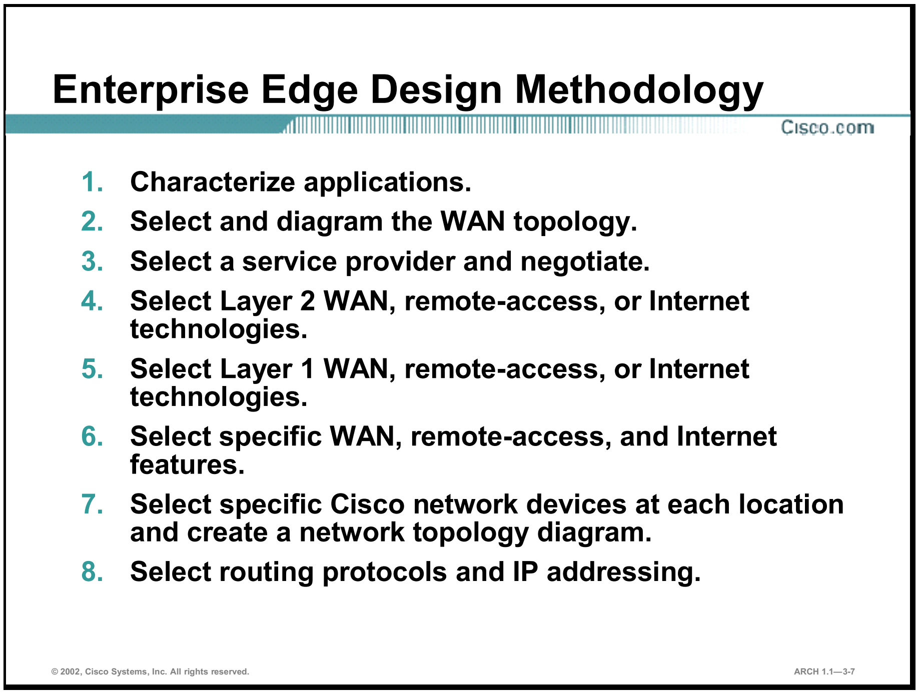

Enterprise Edge Design Methodology

1. **Characterize applications.**
2. **Select and diagram the WAN topology.**
3. **Select a service provider and negotiate.**
4. **Select Layer 2 WAN, remote-access, or Internet technologies.**
5. **Select Layer 1 WAN, remote-access, or Internet technologies.**
6. **Select specific WAN, remote-access, and Internet features.**
7. **Select specific Cisco network devices at each location and create a network topology diagram.**
8. **Select routing protocols and IP addressing.**

ARCH 1.1—3-7

To design the Enterprise Edge functional area, you will complete a series of steps. The table describes the Enterprise Edge design methodology used throughout this course.

Step	Description	Notes and Comments
1.	Characterize applications for the Enterprise Edge functional area.	The important application characteristics are bandwidth, delay, and loss requirements.
2.	Select and diagram the WAN topology.	Based on the geography and data-sharing requirements, you will design the WAN topology.
3.	Select a service provider and negotiate price and features.	Each service provider will offer different services, rates, and quality guarantees. Once you have selected a service provider, you can complete the remaining steps based on the features available to you.
4.	Select a data link layer WAN, remote-access, or Internet technology for each link on the enterprise network.	The data link layer technology selection is based on application requirements and the features a service provider has to offer.
5.	Select a physical layer WAN, remote-access, or Internet technology for each link on the enterprise network.	Based on the data link layer technology selection and the services offered by the service provider, you can select the physical layer technology.
6.	Select specific WAN, remote-access, and Internet features for each link on the enterprise network.	WAN features are based on application requirements and the features a service provider has to offer.

Step	Description	Notes and Comments
7.	Select specific Cisco network devices and hardware and software options at each location and create a network topology diagram.	Based on the specific requirements at each location, select specific Cisco network devices that meet specified criteria.
8.	Select routing protocols and IP addressing for the Enterprise Edge functional area.	Similar to the Enterprise Campus functional area, you will select routing protocols and an IP addressing strategy for the Enterprise Edge functional area.

Analyzing Network Traffic Patterns

The important characteristics of network applications at the Enterprise Edge functional area are bandwidth, delay, and loss. This topic helps you assess and analyze network traffic patterns typically found in the Enterprise Edge functional area.

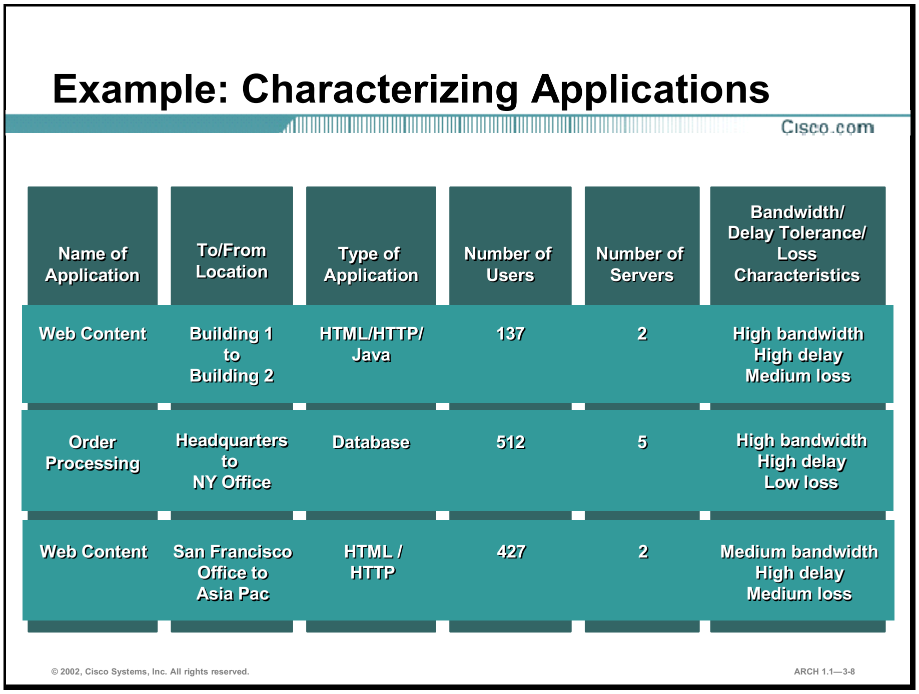

Example: Characterizing Applications

Cisco.com

Name of Application	To/From Location	Type of Application	Number of Users	Number of Servers	Bandwidth/ Delay Tolerance/ Loss Characteristics
Web Content	Building 1 to Building 2	HTML/HTTP/ Java	137	2	High bandwidth High delay Medium loss
Order Processing	Headquarters to NY Office	Database	512	5	High bandwidth High delay Low loss
Web Content	San Francisco Office to Asia Pac	HTML / HTTP	427	2	Medium bandwidth High delay Medium loss

ARCH 1.1—3-8

You will characterize the applications that are shared between any two or more sites on the network. The information you gather will help you determine the performance, scalability, and requirements for each WAN, remote-access, or Internet link.

Use the table to characterize the applications at each network campus location, filling in the fields as indicated. The figure contains an example application table.

Name of Application	To/From Location	Type of Application	Number of Users	Number of Servers	Bandwidth/ Delay/Loss Characteristics

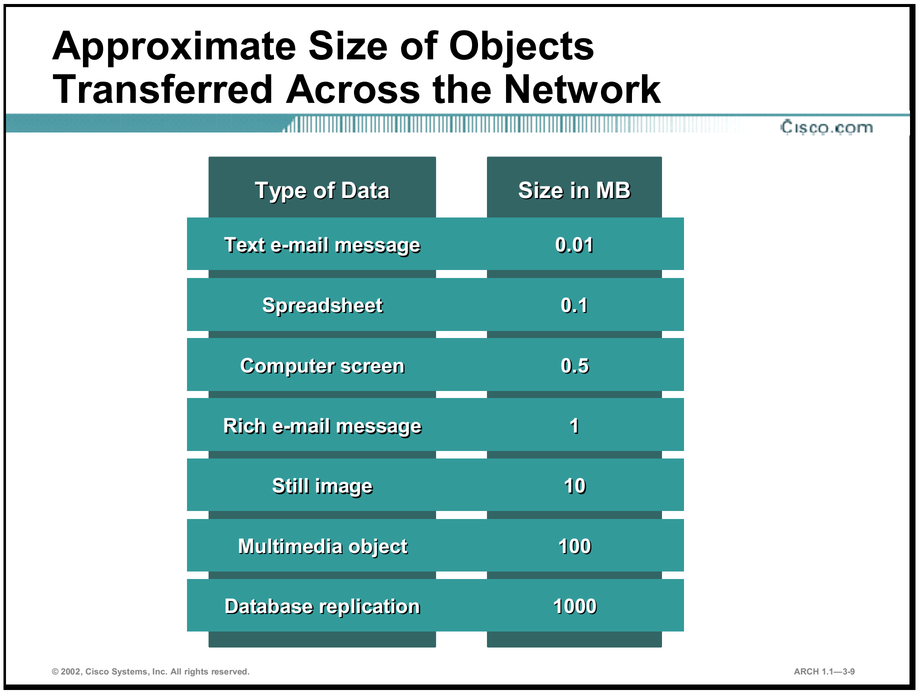

Approximate Size of Objects Transferred Across the Network

Type of Data	Size in MB
Text e-mail message	0.01
Spreadsheet	0.1
Computer screen	0.5
Rich e-mail message	1
Still image	10
Multimedia object	100
Database replication	1000

ARCH 1.1—3-9

Use the table to help characterize traffic load for applications that an enterprise wants to implement. The data in the table is approximate and does not take the place of a thorough analysis of the network in question.

Summary

This topic summarizes the key points discussed in this lesson.

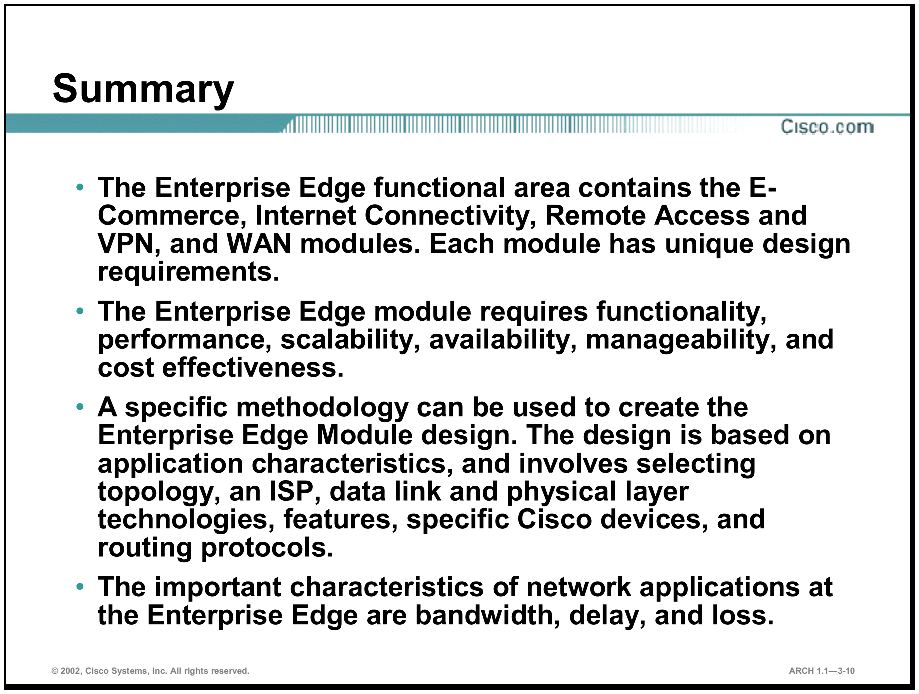

Summary

Cisco.com

- The Enterprise Edge functional area contains the E-Commerce, Internet Connectivity, Remote Access and VPN, and WAN modules. Each module has unique design requirements.
- The Enterprise Edge module requires functionality, performance, scalability, availability, manageability, and cost effectiveness.
- A specific methodology can be used to create the Enterprise Edge Module design. The design is based on application characteristics, and involves selecting topology, an ISP, data link and physical layer technologies, features, specific Cisco devices, and routing protocols.
- The important characteristics of network applications at the Enterprise Edge are bandwidth, delay, and loss.

ARCH 1.1—3-10

References

For additional information, refer to this resource:

- *Wide Area Network Design* at http://www.cisco.com/warp/public/779/largeent/design/wan_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) Which module of the Enterprise Edge model terminates VPN traffic?

A) WAN

B) E-Commerce

C) Internet Connectivity

D) Remote Access and VPN

Q2) Which requirement for the Enterprise Edge can you measure, in part, by utilization?

A) availability

B) functionality

C) performance

D) cost-effectiveness

Q3) Which two Enterprise Edge modules have availability as a critical requirement? (Choose two.)

A) WAN module

B) E-Commerce module

C) Internet Connectivity module

D) Remote Access and VPN module

Q4) What information do you need before you can begin to design the Enterprise Edge functional area?

A) IP addressing scheme

B) identified WAN topology

C) routing protocol selection

D) application characteristics

Q5) Which three types of information do you need to characterize applications on the network? (Choose three.)

A) name of users

B) age of application

C) type of application

D) number of users and servers

E) bandwidth requirement, delay tolerance, and loss characteristics

Quiz Answer Key

Q1) C

Relates to: Enterprise Edge Design within the Enterprise Composite Network Model

Q2) C

Relates to: Typical Requirements for the Enterprise Edge

Q3) A, B

Relates to: Typical Requirements for the Enterprise Edge

Q4) D

Relates to: Enterprise Edge Design Methodology

Q5) C, D, E

Relates to: Analyzing Network Traffic Patterns

Designing the Classic WAN Module

Overview

Wide-area networking provides communications to users across a broad geographic area. Wide-area networks (WANs) typically include routers and switches that link sites and remote offices around the world. Network designers can select WAN links of many sizes and those which utilize different technologies, depending on the requirements of each link.

Relevance

Cisco WAN solutions help network designers build scalable networks and deliver business-critical services using the communications infrastructure.

Objectives

Upon completing this lesson, you will be able to design small, medium, and large enterprise site-to-site WANs, given enterprise WAN needs. This includes being able to meet these objectives:

- Identify typical enterprise needs for site-to-site WANs
- Recommend a WAN topology, given criteria that affect the selection of WAN topologies
- Identify the services and service level agreements desired of WAN service providers
- Recommend data link layer technologies, given criteria that affect the selection of data link layer technologies
- Recommend physical layer protocols, given criteria that affect the selection of physical layer protocols
- Select WAN features that meet specified enterprise requirements
- Select edge routing solutions, based on specific enterprise requirements
- Propose routing protocols and IP addressing strategies for a site-to-site WAN, based on specific enterprise requirements
- Design small, medium, and large enterprise site-to-site WANs, given Enterprise Edge requirements

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

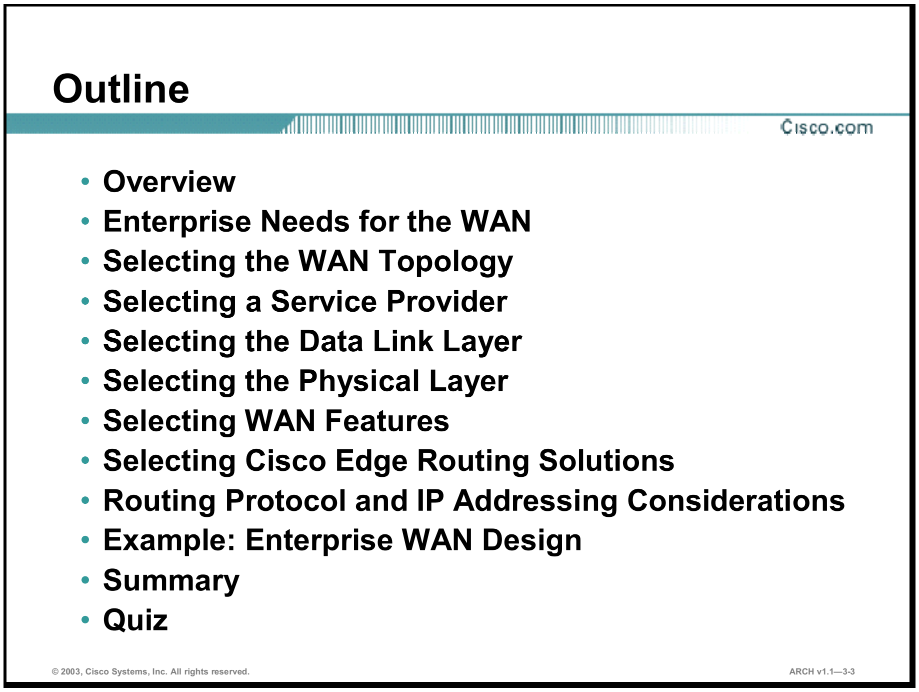

Outline

Cisco.com

- Overview
- Enterprise Needs for the WAN
- Selecting the WAN Topology
- Selecting a Service Provider
- Selecting the Data Link Layer
- Selecting the Physical Layer
- Selecting WAN Features
- Selecting Cisco Edge Routing Solutions
- Routing Protocol and IP Addressing Considerations
- Example: Enterprise WAN Design
- Summary
- Quiz

ARCH v1.1—3-3

Enterprise Needs for the WAN

An enterprise site-to-site WAN network must meet requirements for bandwidth, link quality, reliability, data link protocol characteristics, always-on or on-demand characteristics, and cost. This topic identifies typical enterprise needs for the site-to-site WAN.

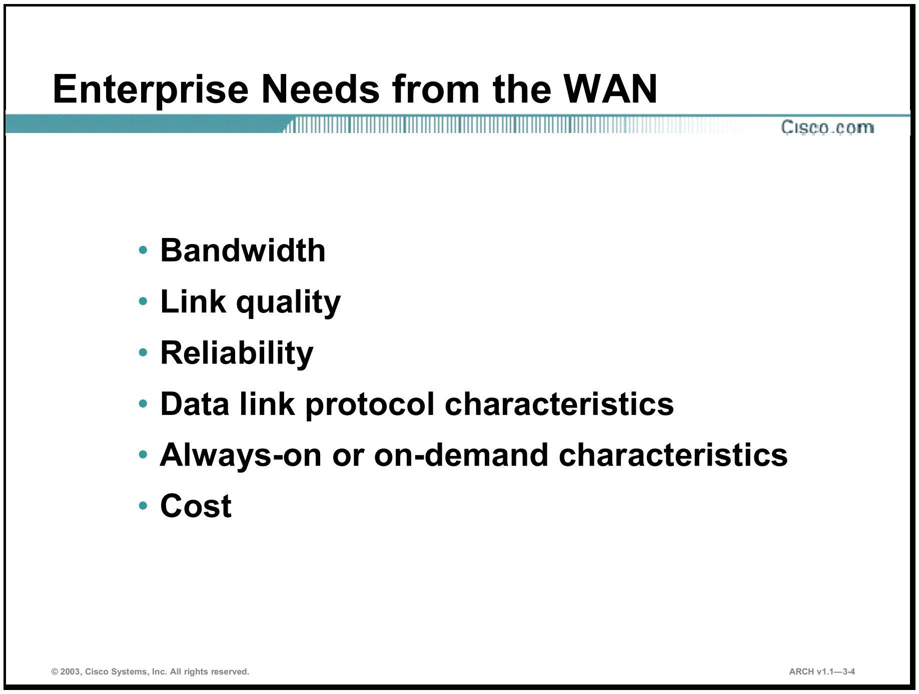

Enterprise Needs from the WAN

- Bandwidth
- Link quality
- Reliability
- Data link protocol characteristics
- Always-on or on-demand characteristics
- Cost

ARCH v1.1—3-4

Applications drive the design of the site-to-site WAN. To determine the site-to-site requirements, understanding the number of users at each location and which applications will be used will dictate the service and bandwidth requirements for each regional and branch site.

Traditionally, WAN communication has been characterized by relatively low throughput and high error rates. Because the WAN infrastructure is generally rented from a service provider, WAN network designs optimize the cost of bandwidth and pursue bandwidth efficiency.

Based on application needs, enterprises typically have these requirements for a site-to-site WAN solution:

- **Bandwidth:** Sufficient bandwidth is required to support applications.
- **Link quality:** High link quality is required to ensure end-to-end delivery of packets.
- **Reliability:** Reliability and availability are critical to ensure end-to-end delivery of packets.
- **Data link protocol characteristics:** Each data link protocol offers services that make the protocol ideal for certain applications.
- **Always-on or on-demand characteristics:** Some applications require that a WAN be available all the time (always on) or as needed (on demand).
- **Cost:** Cost-effectiveness is a concern for any network solution. The WAN solution needs to consider fixed and recurring costs.

Selecting the WAN Topology

The WAN topology includes the physical and logical WAN topology. The topology is closely related to the geographical structure of the enterprise. This topic helps you recommend a WAN topology, using criteria that affect the selection of WAN topologies.

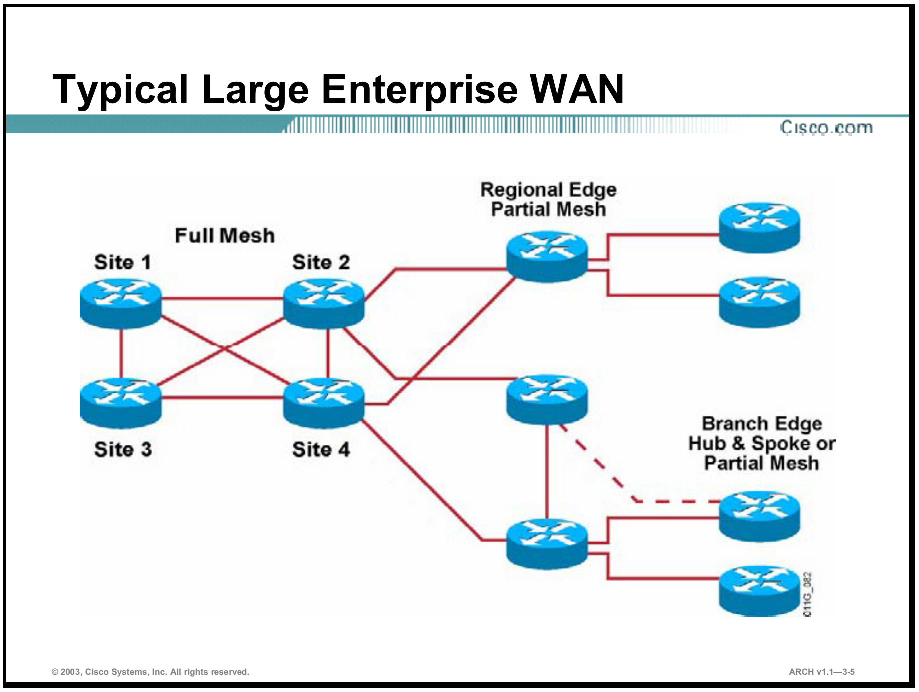

Typical Large Enterprise WAN

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-5

The site-to-site WAN may include a branch, regional, and core hierarchy. You will identify the reliability, availability, and service levels for each level of the hierarchy in the WAN during the design phase.

The actual design will often mirror the enterprise's organizational structure. The WAN requirements at each point in the network will be unique. Although you can group sites that have similar requirements, the actual bandwidth and service requirements will be different for each site on the network.

For example, given that the required bandwidth from a branch office to its regional office is 256 kbps, what amount of bandwidth is needed from the regional office that supports four branch offices to the central site? In the calculation, do not forget to include the bandwidth requirements for the regional office itself, and then add it to the total bandwidth (4 * 256 kbps) coming from the four branch offices.

To determine bandwidth requirements, you can work from the branch and remote-access devices into the central site. This allows you to determine the aggregation needs and to identify bandwidth limitations early in the process.

Starting with the branch office and working back to the Campus Backbone submodule makes it easier to see bottlenecks or potential aggregation issues.

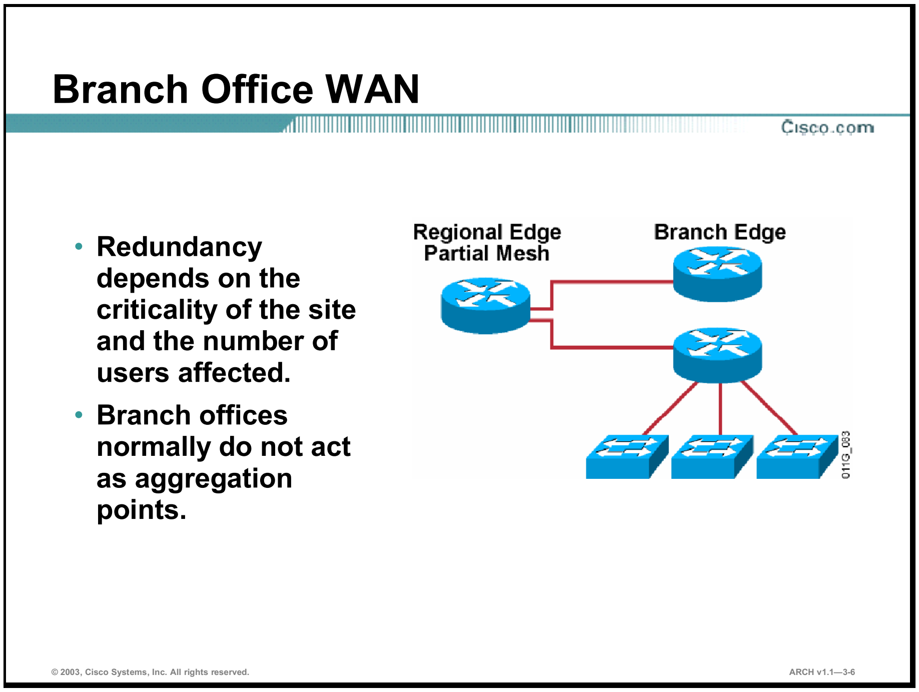

Branch Office WAN

- Redundancy depends on the criticality of the site and the number of users affected.
- Branch offices normally do not act as aggregation points.

ARCH v1.1—3-6

The branch office is typically the end of the network. When designing the WAN to reach the branch office, ask these questions:

- How many users are in the branch?
- What are the per-application bandwidth requirements?
- What is the total bandwidth needed for applications?
- What type of routing protocol is going to be used?
- What are the redundancy needs of the site? What is the effect on the business if the site is unreachable or if the site cannot reach the central servers?
- Is the site supporting on-demand connectivity to other sites or users?

After you determine the number of users and bandwidth, you can determine the total amount of bandwidth for the branch site using this formula:

- ((users * bandwidth) * 1.5) = amount of bandwidth for the site

If the branch office requires redundant links, the design will use either dual WAN links to two different regions or connect to another branch that connects to a regional site. The link between two branch offices is generally the minimum amount of bandwidth to support each branch. In that case, oversize the link between the branch and the regional site to support a fraction of the bandwidth (usually half) of the other branch site. A third method is to implement a dial-on-demand circuit through either ISDN or the PSTN.

The regional office typically sits between the Campus Backbone submodule and the branch offices. It may house the application servers that the branch sites use, or it may simply provide access to the Campus Backbone submodule. The regional office will have its own set of requirements for WAN connectivity.

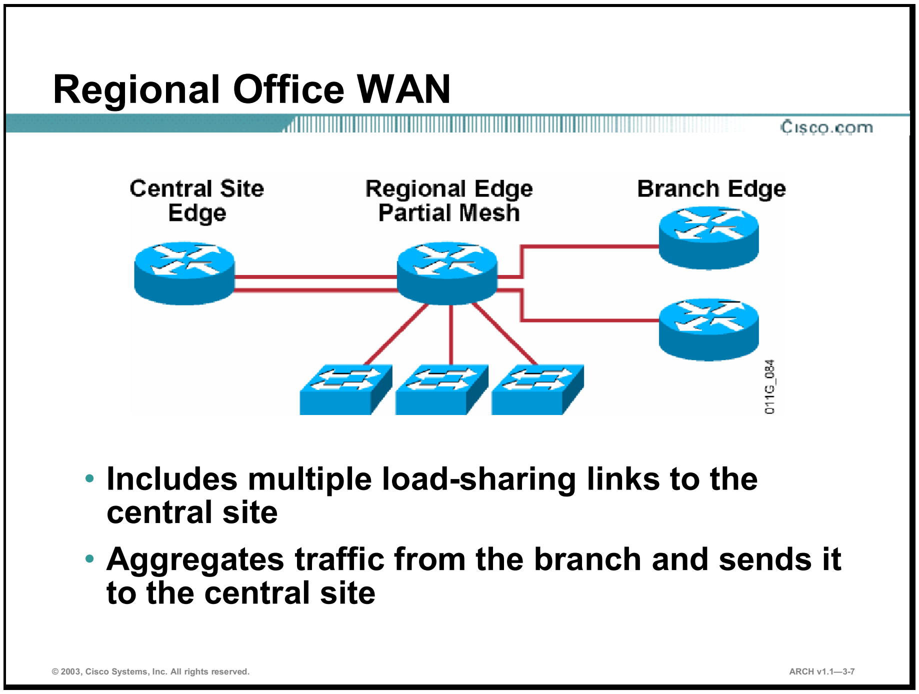

Regional Office WAN

Cisco.com

- Includes multiple load-sharing links to the central site
- Aggregates traffic from the branch and sends it to the central site

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—3-7

Before developing the WAN module to support a regional office WAN, ask these questions:

- How many users are in the regional office?
- What are the per-application bandwidth requirements?
- What is the total bandwidth needed for applications?
- What type of routing protocol is going to be used?
- What are the redundancy needs of the site? What is the effect on the business if the site is not reachable or the site cannot reach the central servers?
- Is the site supporting on-demand connectivity to other sites or users?
- Is the site a rally point for traffic from other sites to pass through?
- Does the regional site have servers or services that are shared with other offices, either branch or core? Does this change the amount of bandwidth that the branch offices need to the core?

It is common for the regional office to be an aggregation point for multiple branch offices. When aggregation is done it is imperative to ensure there is enough bandwidth from the regional office to the core to provide the expected level of service to all branch offices that connect to that regional site. The regional office typically has multiple load-sharing links between the regional office and the central site Enterprise Edge functional area.

The Campus Backbone can be a single router or a collection of routers used to terminate the WAN links from the regional layer. The central site Enterprise Edge functional area is typically a fully meshed environment with multiple load-sharing links, able to distribute all of their aggregated traffic from the regional office.

When load balancing, attempt to avoid asymmetrical routing.

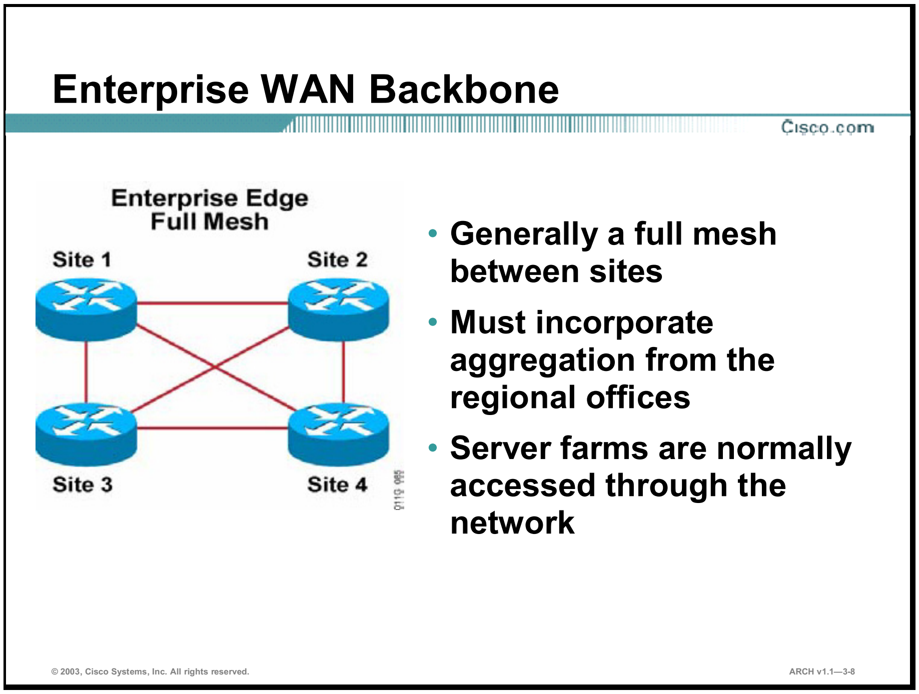

Enterprise WAN Backbone

- Generally a full mesh between sites
- Must incorporate aggregation from the regional offices
- Server farms are normally accessed through the network

ARCH v1.1—3-8

The Campus Backbone WAN is normally the center of the enterprise. The requirement for high-speed connectivity between routers is critical to ensure that all of the outlying regions and branches/remote sites maintain their service levels. When designing the enterprise WAN backbone, ask these questions:

- What are the per-application bandwidth requirements?
- What is the total bandwidth needed for applications?
- What type of routing protocol is going to be used?
- What are the redundancy needs of the site? What is the effect on the business if the site is not reachable?
- Is the site supporting on-demand connectivity to other sites or users?
- Is the site a rally point for traffic from other sites to pass through?

If the core is a single router or the Enterprise Edge routers of a large-scale enterprise, it must support the total aggregation of all of the traffic from the rest of the network.

A full mesh solution may have an impact on convergence time depending on the routing protocol implemented. The fiber reliability that is available today is eliminating the need for a full mesh design. Consider this question: What is the effect on convergence of adding additional links?

Selecting a Service Provider

Once the bandwidth, redundancy, and service level requirements are defined, it is time to talk to several providers to determine what transport is available to implement the design. Do not be surprised if you have to do some redesign based on the features, services, and costs that the provider offers. This topic describes criteria for selecting a WAN service provider.

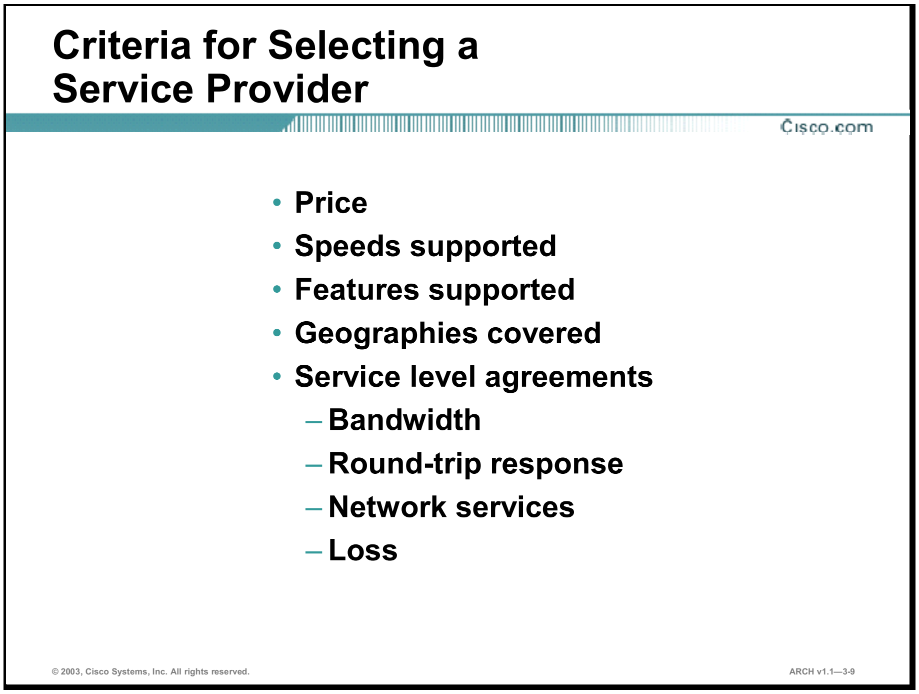

Criteria for Selecting a Service Provider

Cisco.com

- Price
- Speeds supported
- Features supported
- Geographies covered
- Service level agreements
  - Bandwidth
  - Round-trip response
  - Network services
  - Loss

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-9

Each service provider offers a range of prices, speeds, features, and geographical coverage, which will affect your selection. Once you select a service provider, you can redesign the site-to-site WAN based on the services available to you. The selection of a service provider depends on these criteria:

- **Price:** Price, including one-time and fixed costs, is one of the most important criteria when selecting a service provider.
- **Speeds supported:** Different service providers offer a different range of speeds for different technologies. Speed is often closely related to price. Distance may also affect price. Make sure the service provider you select offers the speeds you require.
- **Features supported:** Different service providers offer different WAN features and technologies. Features offered may affect price. Make sure the service provider you select offers the appropriate range of features possible.
- **Geographies covered:** The service provider must service the geographies you need to include in the WAN. Several different service providers may be needed to provide the full geographical coverage the network requires.
- **Service level agreements:** A service level agreement is a key component of a service level contract (SLC). The SLC specifies connectivity and performance agreements for an end-user service from a provider of service. A service provider may provide wide-area or hosted application services. The table describes an example service level agreement and how it is measured.

Network Area	Availability Target	Measurement Method	Average Network Response Time Target	Maximum Response Time Accepted	Response Time Measurement Method
WAN	99.9%	Impacted user minutes	Under 100 ms (round-trip ping)	150 ms	Round-trip ping response
Critical WAN and extranet	99.95%	Impacted user minutes	Under 100 ms (round-trip ping)	150 ms	Round-trip ping response

Selecting the Data Link Layer

For the data link layer, you will select technologies including PPP, Frame Relay, ATM, and X.25. This topic helps you select data link layer technologies, using criteria that affect the selection of data link layer technologies.

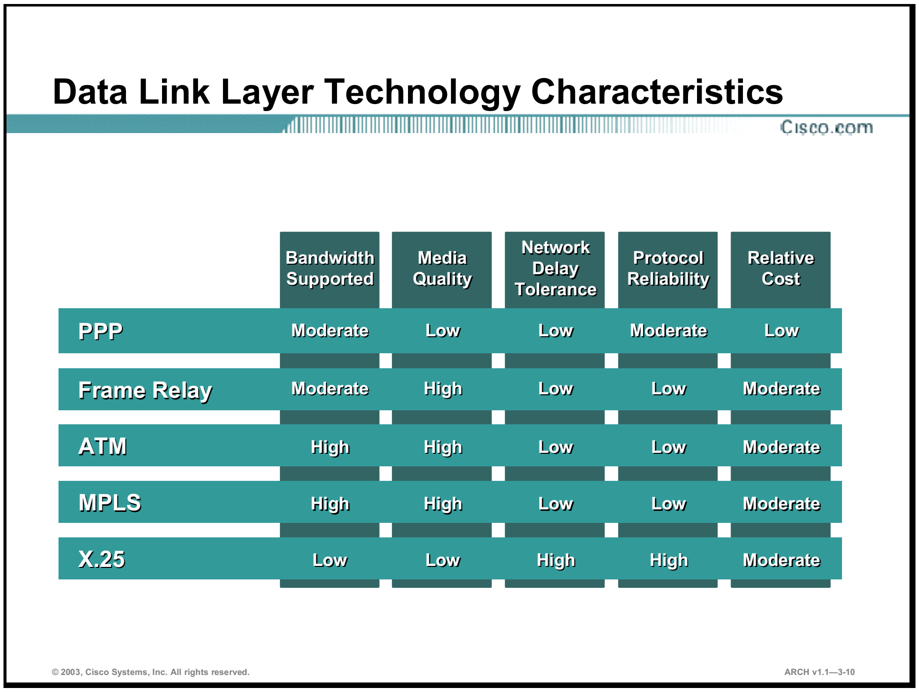

Data Link Layer Technology Characteristics

Cisco.com

	Bandwidth Supported	Media Quality	Network Delay Tolerance	Protocol Reliability	Relative Cost
PPP	Moderate	Low	Low	Moderate	Low
Frame Relay	Moderate	High	Low	Low	Moderate
ATM	High	High	Low	Low	Moderate
MPLS	High	High	Low	Low	Moderate
X.25	Low	Low	High	High	Moderate

ARCH v1.1—3-10

The figure describes the characteristics of data link layer technologies for the WAN. Your selection of a data link technology depends on services that your service provider offers and the features your network requires.

These data link layer technologies are available for site-to-site WANs:

- **PPP:** PPP provides router-to-router and host-to-network connections over synchronous and asynchronous circuits. PPP is typically used for the transmission of IP packets over serial lines and ISDN.
- **Frame Relay:** Frame Relay is a switched data link layer protocol that handles multiple virtual circuits using HDLC-derived encapsulation between connected devices. Frame Relay is more bandwidth efficient than X.25, the protocol for which it is generally considered a replacement. Frame Relay provides cost-effective, high-speed, low-latency virtual circuits between sites. Frame Relay runs over DS0, T1/E1, and serial links.
- **ATM:** ATM is 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, reducing transit delays. ATM is designed to take advantage of high-speed transmission media, such as T1/E1, T3/E3, and optical.
- **Multiprotocol Label Switching (MPLS):** MPLS is a switching mechanism that uses labels (numbers) to forward packets. Labels usually correspond to multilayer 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 data link layer circuit identifier. Label switching occurs regardless of the multilayer protocol.

- **X.25:** X.25 is an International Telecommunication Union Telecommunication Standardization Sector (ITU-T) standard for use in packet data networks. X.25 is an older protocol that is often being replaced by Frame Relay. X.25, which runs over DS0, T1/E1, and serial links.

Selecting the Physical Layer

For the physical layer, you will select physical layer technologies including leased-line, digital subscriber line (DSL), dial-up, ISDN, or optical. This topic helps you recommend physical layer technologies, using criteria that affect the selection of physical layer technologies.

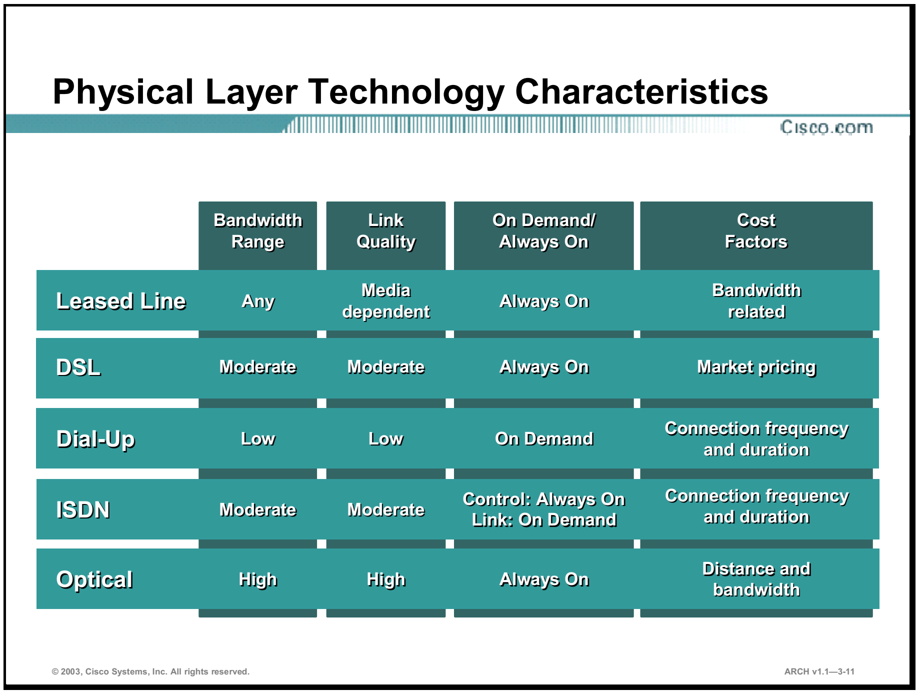

Physical Layer Technology Characteristics

	Bandwidth Range	Link Quality	On Demand/ Always On	Cost Factors
Leased Line	Any	Media dependent	Always On	Bandwidth related
DSL	Moderate	Moderate	Always On	Market pricing
Dial-Up	Low	Low	On Demand	Connection frequency and duration
ISDN	Moderate	Moderate	Control: Always On Link: On Demand	Connection frequency and duration
Optical	High	High	Always On	Distance and bandwidth

ARCH v1.1—3-11

The figure describes the characteristics of physical layer technologies. These physical layer technologies are available for site-to-site WANs:

- **Leased line (synchronous or asynchronous serial; or time-division multiplexing):** Leased lines can be used for PPP networks and hub-and-spoke topologies, or for backup for another type of link.
- **DSL:** DSL is becoming widely available as an always-on Internet connectivity option.
- **Dial-up:** Dial-up is a low-speed but cost-effective technology for intermittent, on-demand WAN requirements including access to corporate data networks.
- **ISDN:** ISDN can be used for cost-effective remote access to corporate networks. It provides support for voice and video as well as a backup for other links.
- **SONET/Synchronous Digital Hierarchy (SDH):** Establishes Optical Carrier (OC) levels from 51.8 Mbps (capable of transporting a T3 circuit) to 40 Gbps.

Selecting WAN Features

After you select the data link layer and physical layer technologies, you can select specific WAN features. Each data link layer technology has its own WAN features to select. This topic helps you select WAN features that meet specified enterprise requirements.

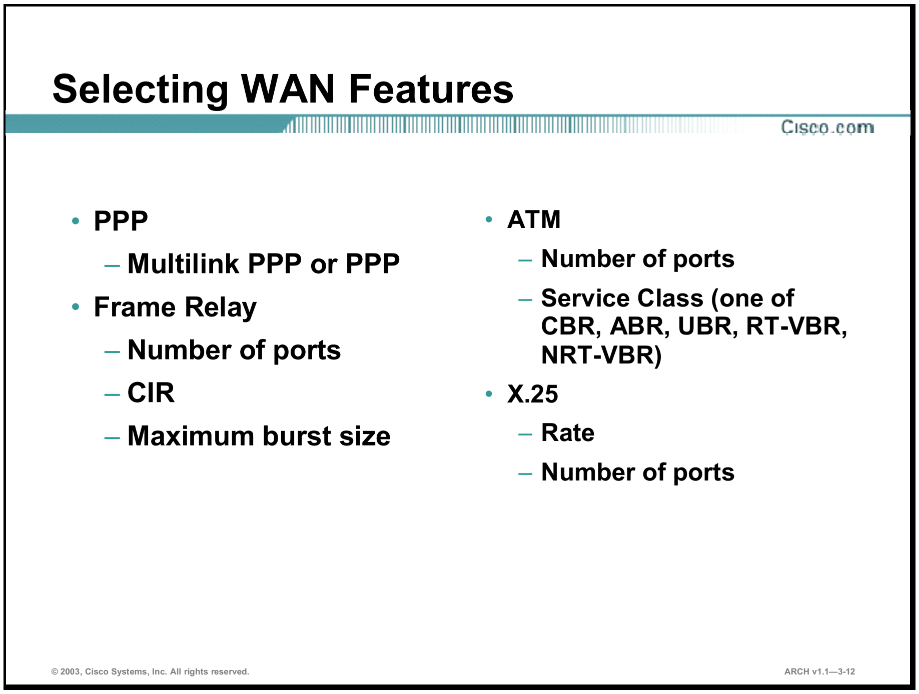

Selecting WAN Features

Cisco.com

- PPP
  - Multilink PPP or PPP
- Frame Relay
  - Number of ports
  - CIR
  - Maximum burst size
- ATM
  - Number of ports
  - Service Class (one of CBR, ABR, UBR, RT-VBR, NRT-VBR)
- X.25
  - Rate
  - Number of ports

ARCH v1.1—3-12

The WAN features available depend on the service provider you selected. For each data link layer technology, you can select these features:

- **PPP:** The rates available for PPP depend on the type of connection, synchronous or asynchronous. Multilink PPP (MP) allows devices to send data over multiple point-to-point data links to the same destination by implementing a virtual link. The MP connection has a maximum bandwidth equal to the sum of the bandwidths of the component links. MP can be configured for either multiplexed links, such as ISDN and Frame Relay, or for multiple asynchronous lines.
- **Frame Relay:** You can select a number of ports, committed information rate (CIR), and maximum burst size for Frame Relay. The number of ports depends on the number of connections required at any point in time as well as bandwidth requirements. The CIR is fixed.

- **ATM:** You can select the number of ports and one of these service classes:
  - **Constant bit rate (CBR):** This traffic category has a bandwidth guarantee. Use it for traffic least tolerant of delay or loss.
  - **Available bit rate (ABR):** This traffic type is scheduled at the same priority as nonreal time (NRT) variable bit rate (VBR). Use it for medium priority traffic.
  - **Unspecified bit rate (UBR):** This traffic category is "best effort." Use it only for least important traffic.
  - **Real-time variable bit rate (RT-VBR):** This traffic category has a higher priority than NRT-VBR and a lower priority than CBR. Use it for medium priority traffic.
  - **Nonreal time variable bit rate:** This traffic type has a higher priority than UBR, but lower than RT-VBR. Use it for medium priority traffic.
- **X.25:** You can select the number of ports and speed for X.25, although rates will be lower than those available for Frame Relay.

Selecting Cisco Edge Routing Solutions

The Cisco Product Advisor is a useful tool for selecting edge routing solutions. The tool is interactive, and provides a list of options from which to choose. This topic shows you how to select Cisco edge routing solutions, based on specific enterprise requirements.

Selecting Cisco Edge Routing Solutions

Cisco.com

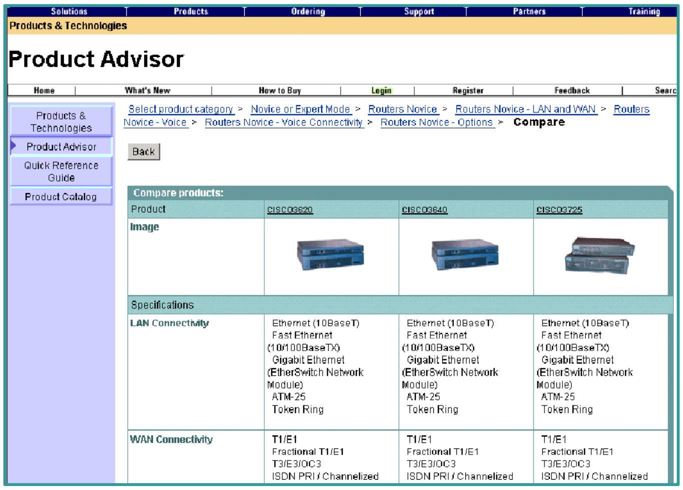

Select product category > Novice or Expert Mode > Routers Novice > Routers Novice - LAN and WAN > Routers Novice - Voice > Routers Novice - Voice Connectivity > Routers Novice - Options > **Compare**

Compare products:

Product	CISCO3620	CISCO3640	CISCO3725
Image			
Specifications			
LAN Connectivity	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX) Gigabit Ethernet (EtherSwitch Network Module) ATM-25 Token Ring	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX) Gigabit Ethernet (EtherSwitch Network Module) ATM-25 Token Ring	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX) Gigabit Ethernet (EtherSwitch Network Module) ATM-25 Token Ring
WAN Connectivity	T1/E1 Fractional T1/E1 T3/E3/OC3 ISDN PRI / Channelized	T1/E1 Fractional T1/E1 T3/E3/OC3 ISDN PRI / Channelized	T1/E1 Fractional T1/E1 T3/E3/OC3 ISDN PRI / Channelized

ARCH v1.1—3-13

Cisco offers the Product Advisor to help you select the right routing solution for the enterprise edge network. The tool operates in two modes: novice and expert. To access the Product Advisor, go to http://www.cisco.com/en/US/products/products_cisco_product_advisor_tool_launch.html and click the Cisco Product Advisor link. Then click a device category. The Product Advisor will ask you questions to help select routers for particular needs. It does not include all products and features, but provides helpful information to help you select appropriate Cisco products.

The table summarizes the questions to answer to help select the right Cisco router. The table describes only the options relevant to a site-to-site WAN, not all options that may be available.

Step	Description	Notes and Comments
1.	Determine the type of environment in which the router will be deployed.	Select one: ■ Small office/branch office ■ Corporate office/central site ■ Large branch/regional office
2.	Determine how the router will be used.	Select one or more: ■ To connect to the Internet ■ To connect offices together ■ To connect employees to the network remotely ■ To deploy IP telephony on the network ■ To provide security for the network
3.	Determine the type of configuration required.	Select one: ■ Fixed configuration (less expensive, not highly scalable) ■ Modular configuration (more expensive, highly scalable) ■ No preference
4.	Determine the type of LAN connectivity required.	Select one or more: ■ Ethernet (10BASE-T) ■ Fast Ethernet (10/100) ■ Gigabit Ethernet
5.	Determine the number of LAN ports required.	Select one: ■ 1 ■ 2 ■ more than 2
6.	Select the types of WAN connectivity required.	Select one or more: ■ T1/E1 ■ Fractional T1/E1 ■ ISDN PRI/Channelized T1/E1 ■ ISDN BRI ■ Synchronous Serial ■ Asynchronous Serial ■ T3/E3/OC3 ■ DSL ■ ATM ■ Frame Relay

Step	Description	Notes and Comments
7.	Determine the number of WAN ports required.	Select one: ■ 1 ■ 2–10 ■ 11–30 ■ more than 30
8.	Determine if voice will be deployed on the WAN, now or in the future.	Answer Yes or No.
9.	Determine if a redundant power supply is required.	Answer Yes or No.
10.	Determine if a rack-mountable router is required.	Answer Yes or No.
11.	Determine which software features are required now and in the future.	Select one or more: ■ VPN ■ Firewall ■ Internet/WAN access
12.	Select the Cisco IOS version for each router you selected.	

Routing Protocol and IP Addressing Considerations

The decision about which routing protocols to implement is based on the design goals, the physical topology of the network, and the configuration of links for remote sites. Routing protocol selection is closely related to IP addressing strategies. This topic helps you select routing protocols and IP addressing for a site-to-site WAN, based on specific enterprise requirements.

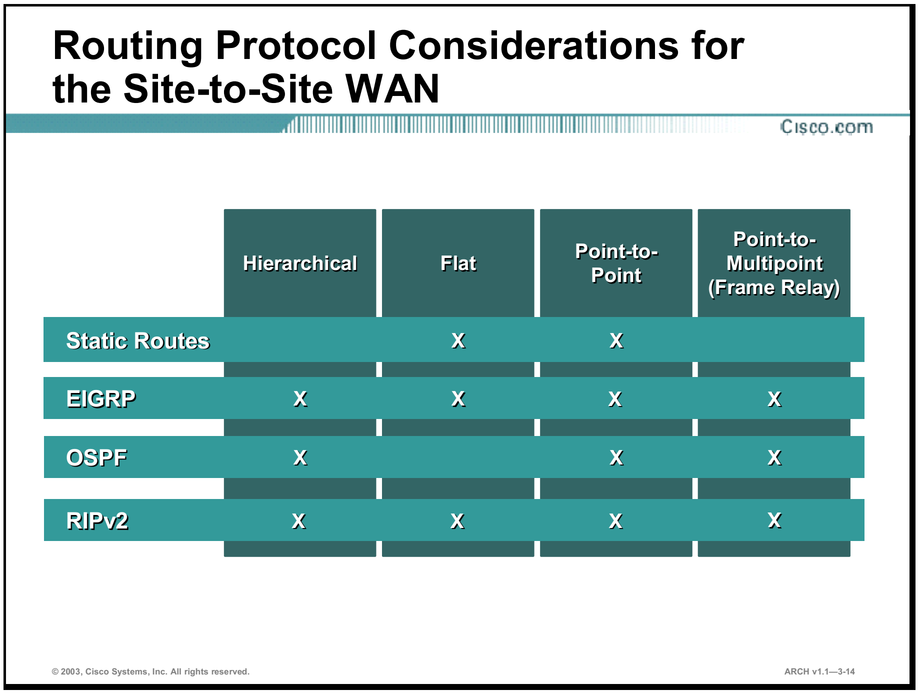

Routing Protocol Considerations for the Site-to-Site WAN

	Hierarchical	Flat	Point-to-Point	Point-to-Multipoint (Frame Relay)
Static Routes		X	X	
EIGRP	X	X	X	X
OSPF	X		X	X
RIPv2	X	X	X	X

ARCH v1.1—3-14

Edge routing protocols include static routes, Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), and Routing Information Protocol version 2 (RIPv2).

Select a site-to-site WAN routing protocol based on these considerations:

- Static routing is useful in smaller environments where there are few WAN connections.
- EIGRP is suitable for nonbroadcast multiaccess (NBMA) environments where there are split horizon issues, such as with Frame Relay or ATM multipoint interfaces. When equipment from multiple vendors is part of the overall design, the use of EIGRP is restricted.
- OSPF is useful in NBMA and dial-up environments. OSPF requires more knowledge for proper configuration.
- RIPv2 is useful for small- to medium-sized networks that do not exceed the 15-hop limit.

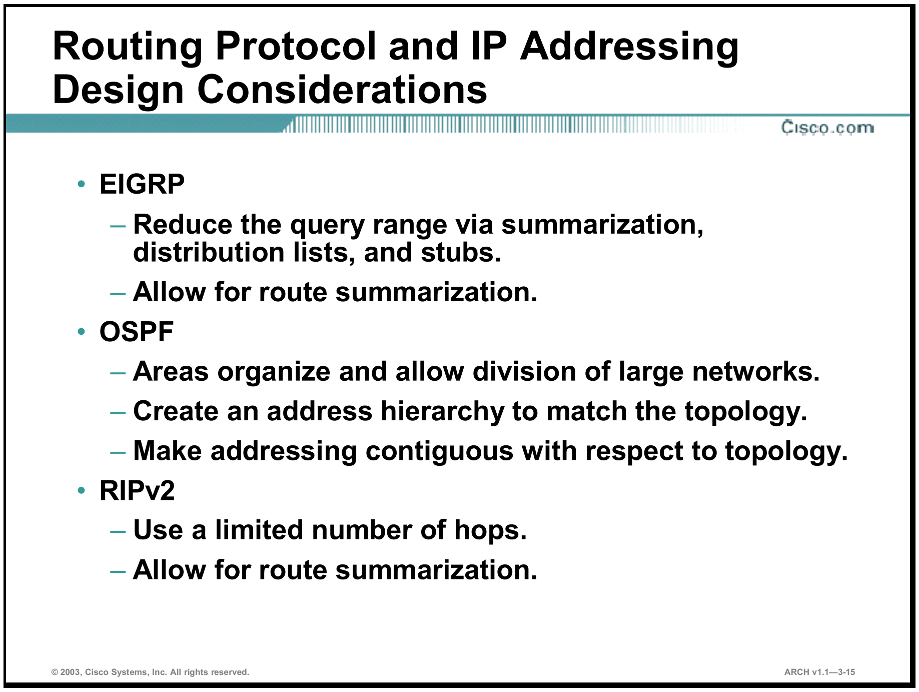

Routing Protocol and IP Addressing Design Considerations

- EIGRP
  - Reduce the query range via summarization, distribution lists, and stubs.
  - Allow for route summarization.
- OSPF
  - Areas organize and allow division of large networks.
  - Create an address hierarchy to match the topology.
  - Make addressing contiguous with respect to topology.
- RIPv2
  - Use a limited number of hops.
  - Allow for route summarization.

ARCH v1.1—3-15

The routing protocol and IP addressing design considerations are closely related. Design considerations for routing protocols and IP addressing are:

- **EIGRP:** Reduce the query change using summarization, distribution lists, and stubs. Use route summarization whenever possible.
- **OSPF:** An area is a logical collection of OSPF routers and links. A router within an area must maintain a topological database for the area to which it belongs. The router does not have any detailed information about networks outside of its area, thereby reducing the size of its database.
- **RIPv2:** RIPv2 is a good choice for hub-and-spoke environments. To design RIPv2, send the default route from the hub to the spokes. The spokes then advertise their connected interface via RIP. RIPv2 can be used when there are secondary addresses on the spokes that need to be advertised or if several vendors' routers are used.

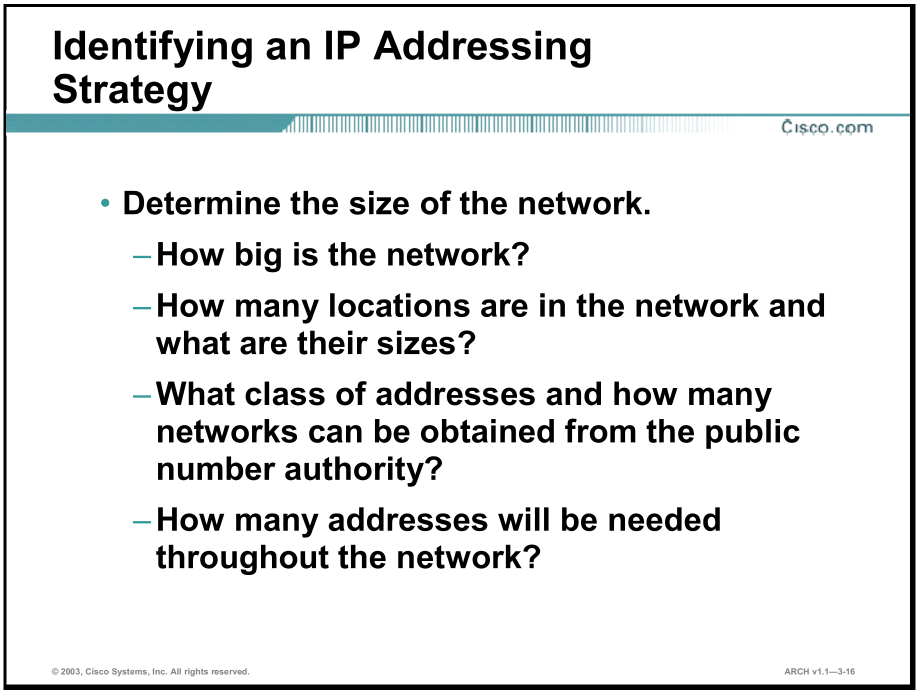

Identifying an IP Addressing Strategy

- **Determine the size of the network.**
  - **How big is the network?**
  - **How many locations are in the network and what are their sizes?**
  - **What class of addresses and how many networks can be obtained from the public number authority?**
  - **How many addresses will be needed throughout the network?**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-16

The questions to ask to help identify the IP addressing strategy for the WAN are the same as those you would ask for the enterprise campus network.

The first step in IP addressing plan design is to determine the size of the network in order to establish how many IP addresses are needed. To gather the required information, answer these questions:

- **How big is the network?** Determine the number of workstations, servers, IP Phones, router interfaces, switch management interfaces, firewall interfaces, and so on. The summary defines the minimum overall number of IP addresses required for the network. Because all networks tend to grow, allow a reserve of about 20 percent for potential network expansion.
- **How many locations are in the network and what are their sizes?** The information about the sizes of the individual locations is closely related to the overall network size.
- **What class of addresses and how many networks can be obtained from the public number authority?** The required IP address classes for the planned network are based on information about the network size, the number of locations, and the size of the individual locations.
- **How many addresses will be needed throughout the network?** Determine the number of addresses needed for workstations, servers, IP phones, network devices, and so on.

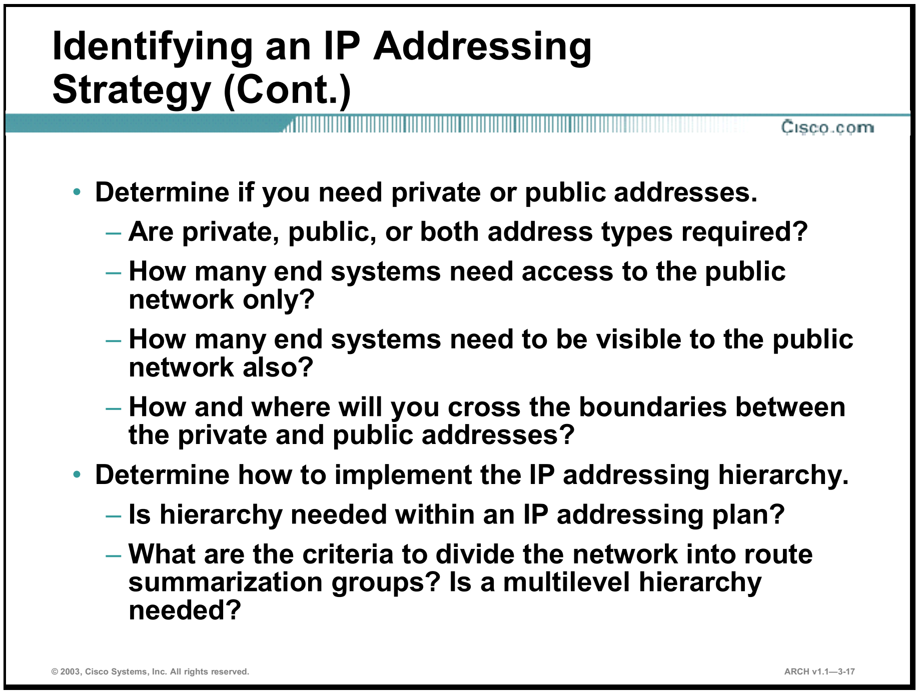

Identifying an IP Addressing Strategy (Cont.)

Cisco.com

- Determine if you need private or public addresses.
  - Are private, public, or both address types required?
  - How many end systems need access to the public network only?
  - How many end systems need to be visible to the public network also?
  - How and where will you cross the boundaries between the private and public addresses?
- Determine how to implement the IP addressing hierarchy.
  - Is hierarchy needed within an IP addressing plan?
  - What are the criteria to divide the network into route summarization groups? Is a multilevel hierarchy needed?

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—3-17

Next, determine if you need private or public addresses based on these questions:

- **Are private, public, or both address types required?** The decision when to use private, public, or both address types depends on the Internet presence and the number of publicly visible servers. Four situations are possible:
  - **No Internet connectivity:** The network is isolated and there is no need to acquire public addresses.
  - **Internet connectivity, no public accessible servers:** The network is connected to the Internet and thus public addresses are required. Use one public address and translation mechanism to allow access to the Internet. Private addresses are used to address the internal network.
  - **Internet connectivity, public accessible servers:** The public addresses are required to connect all public accessible servers to the Internet. The required number of public addresses varies, but in most instances a public address is required for the routers that connect to the Internet, and any publically accessible servers, plus a range of addresses needs to be available for the corporate users that are accessing the Internet.
  - **All end systems should be publicly accessible:** Only public addresses are required and used to address the whole network.
- **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.

- **How many end systems need to be visible to the public network also?** This is the number of Internet connections and 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), which defines the number of required public addresses. These end systems require addresses that are globally unambiguous.
- **How and where will you cross the boundaries between the private and public addresses?** When private addresses are used for addressing in a network, and this network needs to be connected to the Internet, a translation mechanism such as Network Address Translation (NAT) must be used to translate from private to public addresses and vice versa.

The decision on how to implement the IP addressing hierarchy is an administrative decision that is based on these questions:

- **Is hierarchy needed within an IP addressing plan?** You will decide how to implement the IP addressing hierarchy based on the network size and the geography and topology of the network. In large networks, the hierarchy within the IP addressing plan is required in order to have a stable network.
- **What are the criteria to divide the network into route summarization groups?** The network is usually divided into route summarization groups based on the network size and topology.

Example: Enterprise WAN Design

This example provides an opportunity to look at the design of the branch, regional, and Campus Backbone WAN components.

Company Background

A national insurance brokerage has four main data centers and each data center has four regional offices and each regional office has four branch offices. There are a total of 64 branch offices, 16 regional offices, and 4 core offices.

Their primary WAN applications are e-mail and database applications with low bandwidth requirements. The IT department has done a study and determined that each branch office needs 128 kbps worth of bandwidth, and each regional office needs 256 kbps of bandwidth.

The first part of the design project was to determine how the branch offices will be connected to the regional offices. The company identified these network needs:

- **Total bandwidth required for applications:** 128 kbps per branch.
- **Redundancy needs of the site:** The IT staff determined that 32 of the 64 offices did not require a redundant link, since the downtime at those offices would not drastically impact the company. The other 32 sites needed redundancy. The decision was to have those branch offices dual homed to two different regional offices.

Example Branch Office WAN

Cisco.com

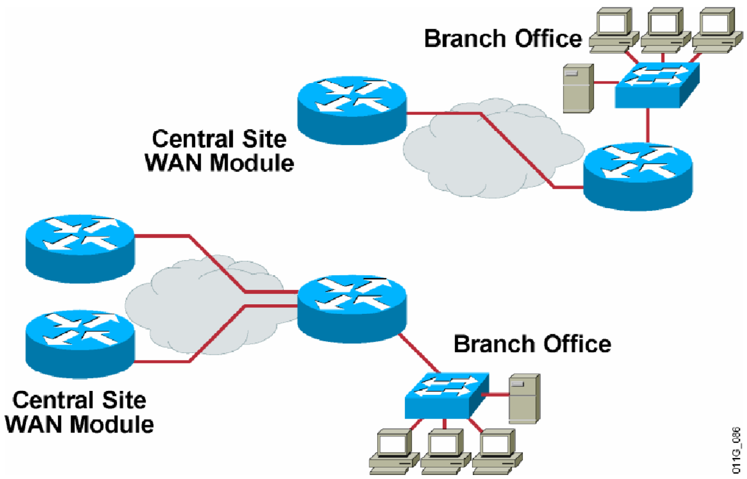

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-18

From the information gathered, the company decided to implement two different design scenarios. The requested bandwidth for the single run site will be 128 kbps. For the dual-homed branch offices, each link will be 128 kbps.

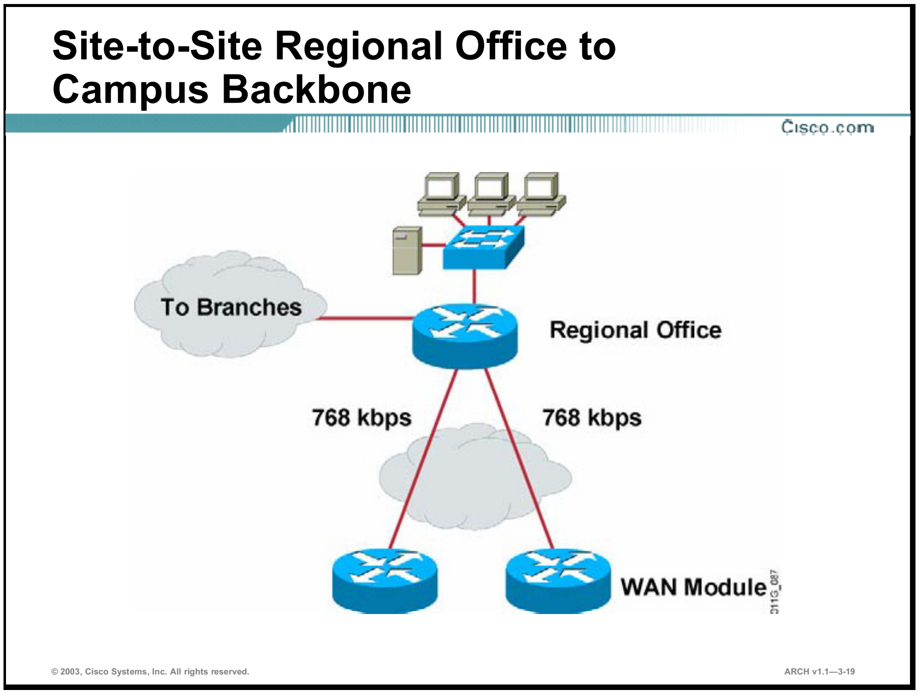

Site-to-Site Regional Office to Campus Backbone

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-19

The next goal was to design the regional office WAN layout. The IT department wanted the regions to be connected through a closed delta design, each region connected to two different core sites. To determine the amount of bandwidth needed between the regions and the Campus Backbone, they calculated the aggregate bandwidth from the branches (4 * 128 = 512) and then added it to the bandwidth requirements for the regional site (512 + 256 = 768). Since there were two load-sharing paths to the Campus Backbone submodule, they were sized using the current requirements of 768 kbps per link.

To complete the design decision, the company identified these network needs:

- **Total bandwidth needed for applications:** 256 kbps. There are four branch offices requiring a total of 512 kbps of throughput to the Campus Backbone submodule.
- **Redundancy needs of the site:** There is a requirement for load-sharing redundant links from the regional offices to the Campus Backbone submodule.

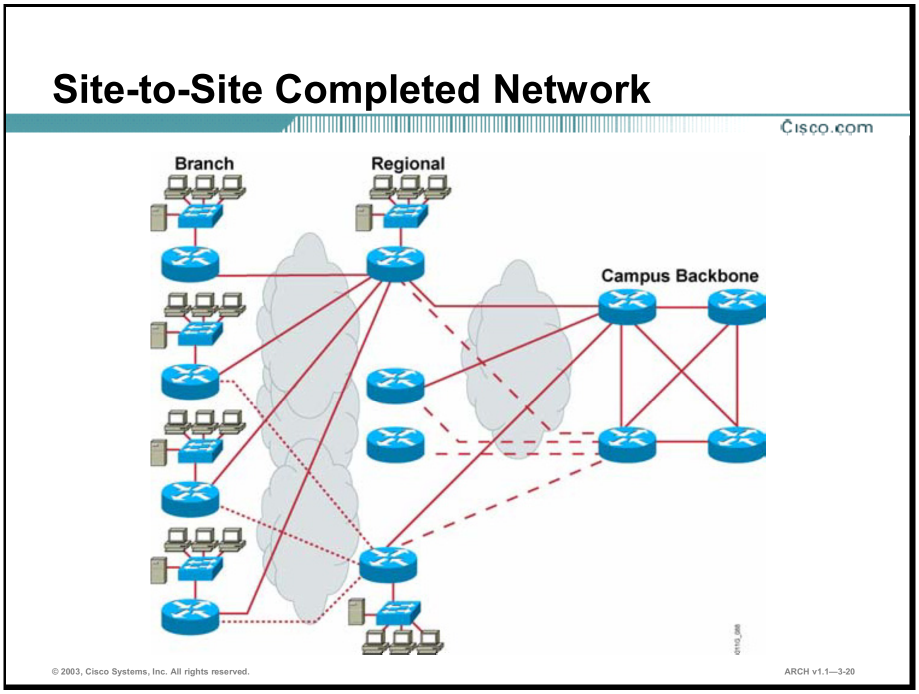

Site-to-Site Completed Network

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-20

The next step was to design the Campus Backbone connectivity. Between the core and the regional office, 768-kbps links were implemented. The IT team determined that each core could be self-sufficient, but the company wanted sufficient bandwidth to support database and server replication. According to their experts, two T1s of bandwidth were required. From each site, they required 768 * 4 = 3072 kbps of bandwidth. Therefore, the two T1s = 3072 + 3000 = 6 Mbps worth of bandwidth between the Campus Backbone sites. The service provider was willing to provide a T3 for the price of 4 T1s. Each core router had three T3s going to the other core devices.

Site-to-Site WAN Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Notes and Comments
What topology will be used for the WAN?	Mixture of partial mesh and full mesh in the core	Given the size and requirements of the network several topologies are used.
What service provider will be selected?	National carrier	A national carrier is required to provide geographical coverage.
What data link layer protocol will be used?	Frame Relay where available and PPP leased lines where Frame Relay is not available	
What physical network media will be used?	Copper or fiber	
Which Cisco products will be used?	Branch offices: Cisco 1720 Regional offices: Cisco 2620 Core: Cisco 3640	
Which routing protocols will be used?	OSPF hierarchical design	Given the number of sites and the way the design leads to an easy division of areas, OSPF was chosen.
What IP addressing scheme will be used?	Access to the Internet and NAT are required	A single Class C address provides Internet connectivity. A Class B is required used internally and NAT is used outside the corporate network.

Summary

This topic summarizes the key points discussed in this lesson.

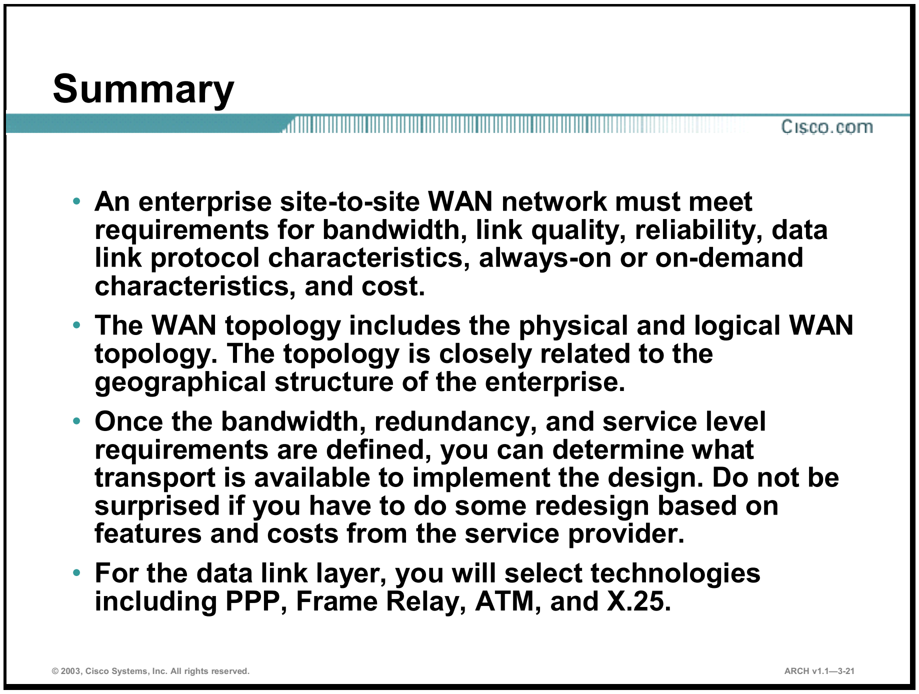

Summary

- An enterprise site-to-site WAN network must meet requirements for bandwidth, link quality, reliability, data link protocol characteristics, always-on or on-demand characteristics, and cost.
- The WAN topology includes the physical and logical WAN topology. The topology is closely related to the geographical structure of the enterprise.
- Once the bandwidth, redundancy, and service level requirements are defined, you can determine what transport is available to implement the design. Do not be surprised if you have to do some redesign based on features and costs from the service provider.
- For the data link layer, you will select technologies including PPP, Frame Relay, ATM, and X.25.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—3-21

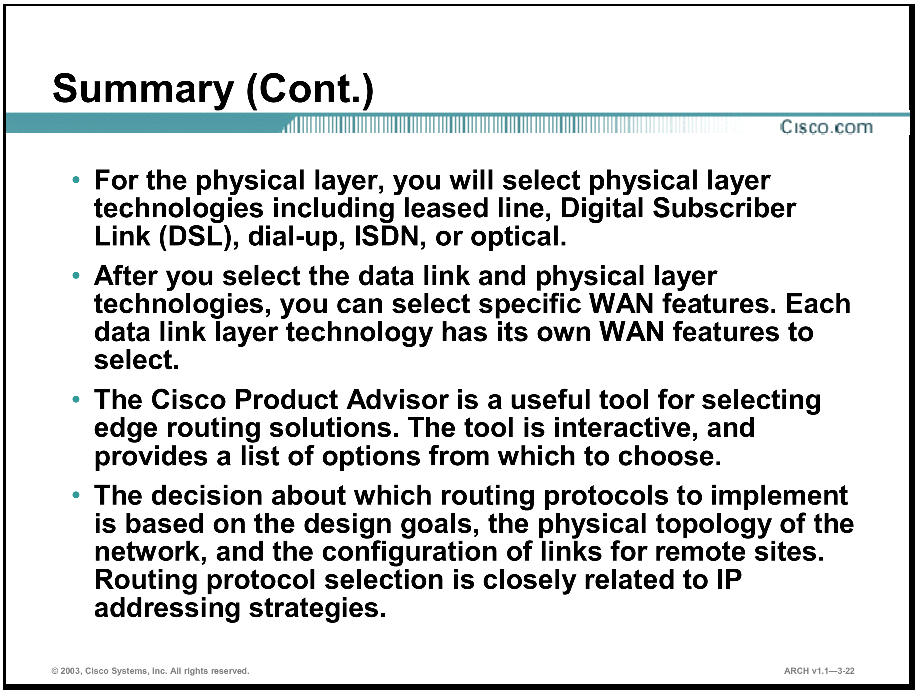

Summary (Cont.)

- For the physical layer, you will select physical layer technologies including leased line, Digital Subscriber Link (DSL), dial-up, ISDN, or optical.
- After you select the data link and physical layer technologies, you can select specific WAN features. Each data link layer technology has its own WAN features to select.
- The Cisco Product Advisor is a useful tool for selecting edge routing solutions. The tool is interactive, and provides a list of options from which to choose.
- The decision about which routing protocols to implement is based on the design goals, the physical topology of the network, and the configuration of links for remote sites. Routing protocol selection is closely related to IP addressing strategies.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—3-22

References

For additional information, refer to these resources:

- *Wide Area Network Design* at http://www.cisco.com/warp/public/779/largeent/design/wan_index.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

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 two requirements help ensure end-to-end delivery of packets? (Choose two.)

A) reliability

B) bandwidth

C) on-demand

D) link quality

E) protocol characteristics

Q2) What type of topology does a Campus Backbone WAN generally support?

A) full mesh

B) partial mesh

C) point-to-point

D) switched backbone

Q3) Which formula helps define the amount of bandwidth required between sites in a site-to-site WAN?

A) ([users * bandwidth] * 1) = amount of bandwidth for the site

B) ([users * bandwidth] * 2) = amount of bandwidth for the site

C) ([users * bandwidth] * 1.5) = amount of bandwidth for the site

D) ([peak times * bandwidth] * 1.5) = amount of bandwidth for the site

Q4) Which service provider selection criterion is closely related to speed?

A) price

B) features

C) geographies

D) service-level agreements

Q5) What reason might you have for choosing ATM instead of Frame Relay for the data link layer?

A) cost

B) bandwidth

C) reliability

D) network delay

Q6) Which physical layer technology might be the best choice for a cost-effective remote-access solution that natively supports voice?

A) DSL

B) ISDN

C) dial-up

D) SONET

Q7) In which situation would dial-up be a viable physical layer choice?

A) remote access to video applications

B) intermittent telecommuter connections

C) always-on branch office to corporate connection

D) high bandwidth data center to corporate connection

Q8) For which WAN technology can you select maximum burst size?

A) PPP

B) X.25

C) ATM

D) Frame Relay

Q9) How does the Cisco Product Advisor help you to select a routing solution?

A) by listing all products available

B) by categorizing products by solution needs

C) by matching your criteria with specific products

D) by presenting product feature and configuration tables

Q10) For which routing protocol is route summarization recommended whenever possible?

A) BGP

B) OSPF

C) EIGRP

D) static routes

Q11) Which data link layer technology might be a suitable WAN alternative to Frame Relay?

A) PPP

B) ATM

C) X.25

D) EoPPP

Quiz Answer Key

Q1) A, D

Relates to: Enterprise Needs for the WAN

Q2) A

Relates to: Selecting the WAN Topology

Q3) C

Relates to: Selecting the WAN Topology

Q4) A

Relates to: Selecting a Service Provider

Q5) B

Relates to: Selecting the Data-Link Layer

Q6) B

Relates to: Selecting the Physical Layer

Q7) B

Relates to: Selecting the Physical Layer

Q8) D

Relates to: Selecting WAN Features

Q9) C

Relates to: Selecting Cisco Edge Routing Solutions

Q10) C

Relates to: Routing Protocol and IP Addressing Considerations

Q11) B

Relates to: Example: Enterprise WAN Design

Designing the Remote Access Module

Overview

Easy connectivity solutions and consistency are important for enterprises relying on remote access. Customers, employees, and partners should be able to connect as if they are at the company site. They also must count on the ability to log in and to remain connected at an expected level of performance.

The number of telecommuters and mobile users is growing every day. Their communications needs are expanding from simple data transmission to the need for voice, fax, and data transmission, including conferencing.

Note Remote-Access Virtual Private Networks (VPNs) are discussed in the “Designing Virtual Private Networks” module.

Relevance

The remote-access server is an integral part of the total network solution and must scale to meet growing demand.

Objectives

Upon completing this lesson, you will be able to design an enterprise remote-access solution, given enterprise remote-access needs. This includes being able to meet these objectives:

- Identify typical enterprise needs for remote access
- List guidelines to help select the type of remote access based on specified needs and the type of termination required for that solution
- Select physical layer protocols for the remote-access solution
- Select data link layer protocols for the remote-access solution

- Select Cisco access routing solutions for users and aggregation points, based on specific enterprise requirements
- Design small, medium, and large enterprise remote-access solutions, given enterprise remote-access needs

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

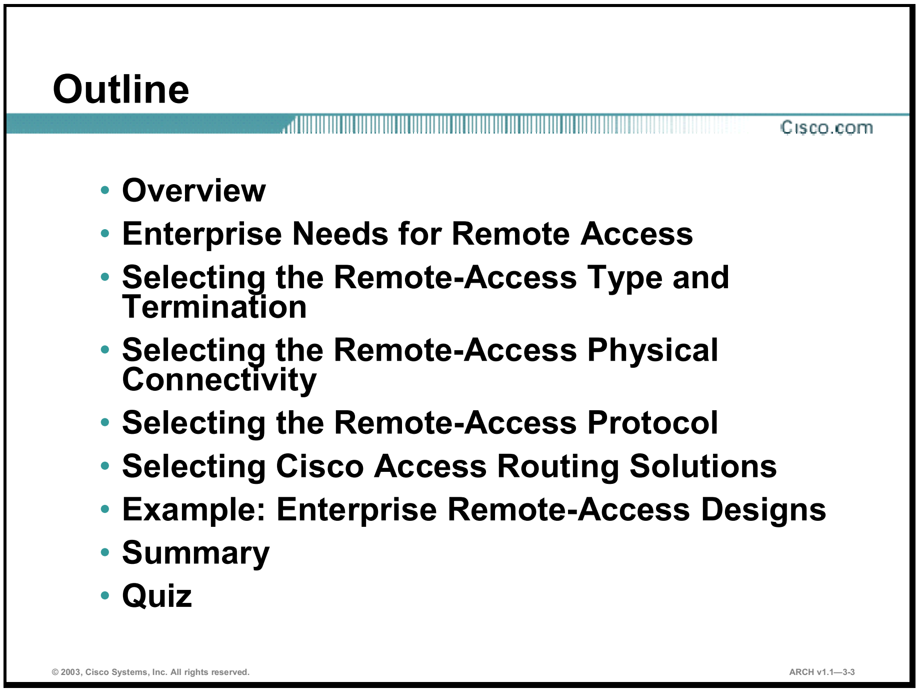

Outline

- **Overview**
- **Enterprise Needs for Remote Access**
- **Selecting the Remote-Access Type and Termination**
- **Selecting the Remote-Access Physical Connectivity**
- **Selecting the Remote-Access Protocol**
- **Selecting Cisco Access Routing Solutions**
- **Example: Enterprise Remote-Access Designs**
- **Summary**
- **Quiz**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-3

Enterprise Needs for Remote Access

Telecommuters, remote users, and branch offices all require remote access to a central site. The key concerns for an information technology (IT) department are functionality, performance, scalability, availability, manageability, and security. This topic identifies typical enterprise needs for remote access.

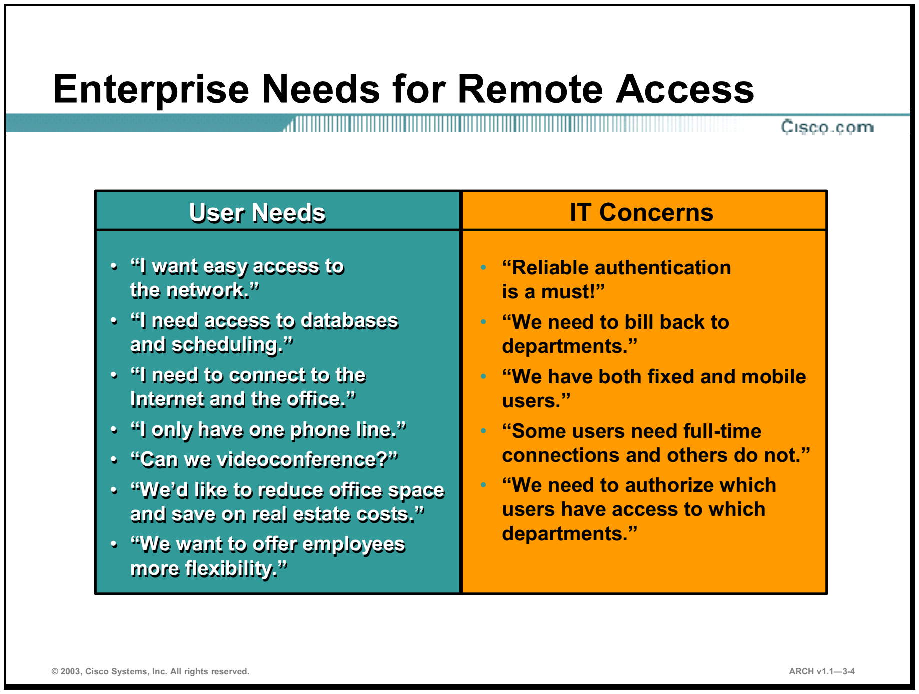

Enterprise Needs for Remote Access

User Needs	IT Concerns
• "I want easy access to the network." • "I need access to databases and scheduling." • "I need to connect to the Internet and the office." • "I only have one phone line." • "Can we videoconference?" • "We'd like to reduce office space and save on real estate costs." • "We want to offer employees more flexibility."	• "Reliable authentication is a must!" • "We need to bill back to departments." • "We have both fixed and mobile users." • "Some users need full-time connections and others do not." • "We need to authorize which users have access to which departments."

ARCH v1.1—3-4

Remote connections link single users (mobile users and/or telecommuters) and branch offices to a local campus or the Internet. Typically, a remote site is a small site that has few users and therefore needs a smaller WAN connection. The remote requirements of an internetwork, however, usually involve a large number of remote single users or sites, and therefore an internetwork usually needs a larger WAN connection.

Since there are so many remote single users or sites, the aggregate WAN bandwidth cost is proportionally greater in remote connections than in WAN connections. The WAN media rental charge from a service provider is the largest cost component of a remote network. Unlike WAN connections, smaller sites or single users seldom need to connect 24 hours a day.

Easy connections and consistency are crucial to companies relying on remote access. Customers, employees, and partners should be able to connect seamlessly, as if they were in one office. They must also count on the ability to log in and remain connected at an expected level of performance.

Security is a high priority for remote access. Security solutions are discussed in the Designing Security Services module.

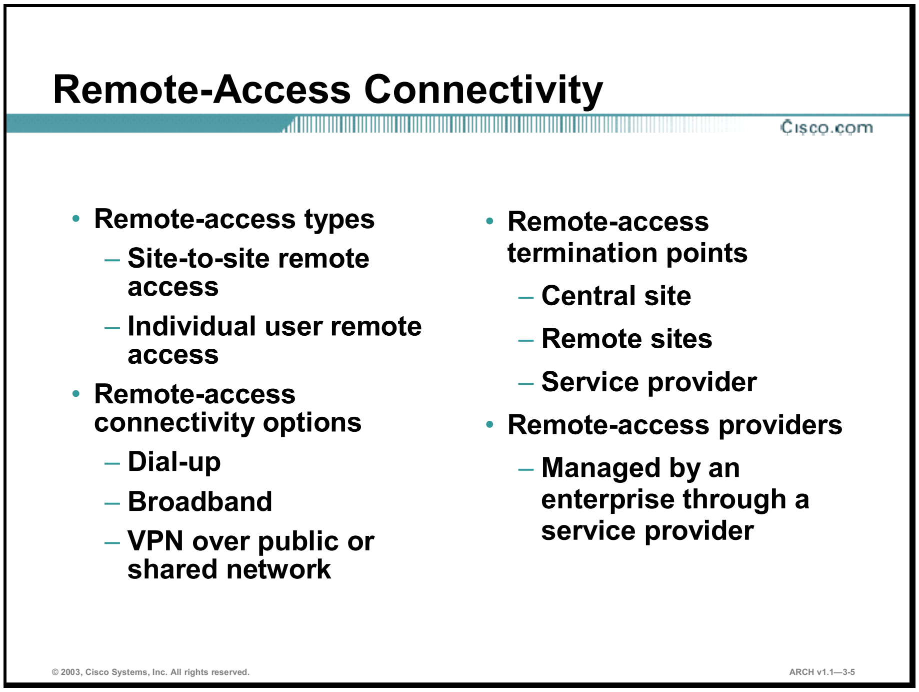

Remote-Access Connectivity

- Remote-access types
  - Site-to-site remote access
  - Individual user remote access
- Remote-access connectivity options
  - Dial-up
  - Broadband
  - VPN over public or shared network
- Remote-access termination points
  - Central site
  - Remote sites
  - Service provider
- Remote-access providers
  - Managed by an enterprise through a service provider

ARCH v1.1—3-5

When designing the remote-access connectivity there are four functional areas to address. The basic questions to ask are:

- **What type of remote access is needed?** Determine if there is a group of users in a remote location that need intermittent data exchanges with an enterprise site, or if there are individuals located in different places that need their own connectivity solution.
- **What types of remote-access connectivity is needed in the environment?** The solution may consist of a single-access methodology or a combination of different methodologies. The most common today is to run PPP through either an analog dial-up circuit, a digital trunk, or through the current WAN service provider to provide a Virtual Private Network through either their network or a public network.

  An enterprise may decide to build a remote-access VPN if they have analog and digital circuits that terminate with a service provider, and the expense of moving those circuits is prohibitive. An enterprise site could also terminate the physical layer locally, and terminate the VPN inside the enterprise network.
- **Where is the remote-access termination point going to be?** Most often, the remote-access termination is located at a central site, but it could be at a remote site or even within a service provider network. If the enterprise decides to host its own remote-access termination point, they must decide if it is less expensive to bring all the termination back to the central office or to distribute it between regional or remote branches to save on toll charges.
- **Who is going to provide the actual endpoint for termination of the remote-access device?** The options include having the equipment in-house with internal IT management, or outsourcing with a link to the outsourced party.

Once you answer these questions, you can select remote-access physical, data link, and network layer technologies based on functionality, reliability, and cost-effectiveness.

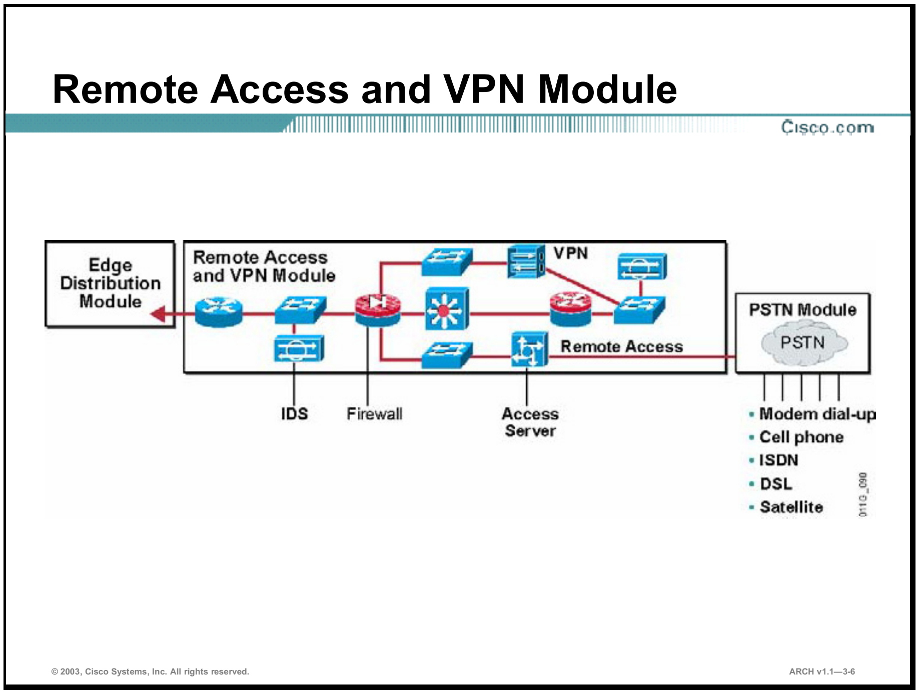

Remote Access and VPN Module

ARCH v1.1—3-6

The Remote Access and VPN module of the Enterprise Composite Network Model provides remote access to end users. The primary components within this model are the circuits, the access server (which provides authentication and authorization), firewalls, and, optionally, Intrusion Detection Systems (IDSs).

From the PSTN to the Remote Access and VPN module, you can deploy many effective data link layer technologies.

Note Refer to the “Designing Security Services” module to learn how to design security services for remote access and the entire enterprise.

Note Refer to the “Designing Virtual Private Networks” module to learn more about designing an enterprise VPN solution.

Selecting the Remote-Access Type and Termination

The two primary types of remote access are site-to-site and user-to-site. This topic provides guidelines to help you select the type of remote access based on specified needs and the type of termination required.

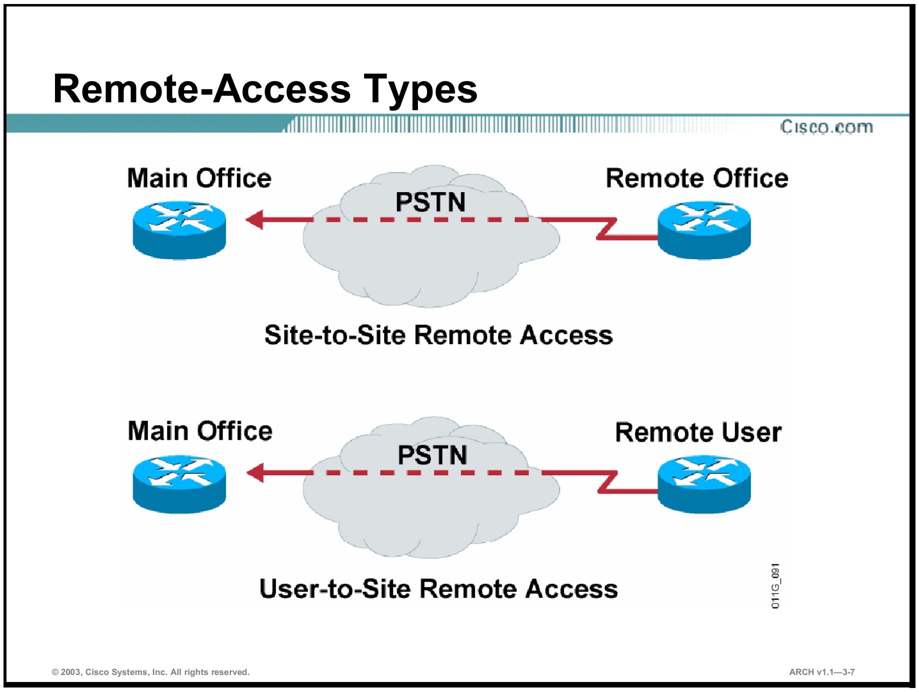

Remote-Access Types

Cisco.com

Main Office

PSTN

Remote Office

Site-to-Site Remote Access

Main Office

PSTN

Remote User

User-to-Site Remote Access

011G_091

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-7

The site-to-site remote-access model is appropriate when there are a group of users in the same vicinity that can share an on-demand connection to either their local branch office or central site. The criteria for selecting a site-to-site remote-access solution include:

- Sporadic need for enterprise network connectivity, not requiring an “always up” connection
- Multiple users at a facility sharing the on-demand access
- Prohibitive cost of putting in a dedicated always-on connection

The most common remote-access model involves a single user, needing connectivity to the corporate network, who is somewhere where there are no immediate network resources to connect to. They may dial in through a dial-up mechanism or they may connect through an always-on connection and access the network through a VPN.

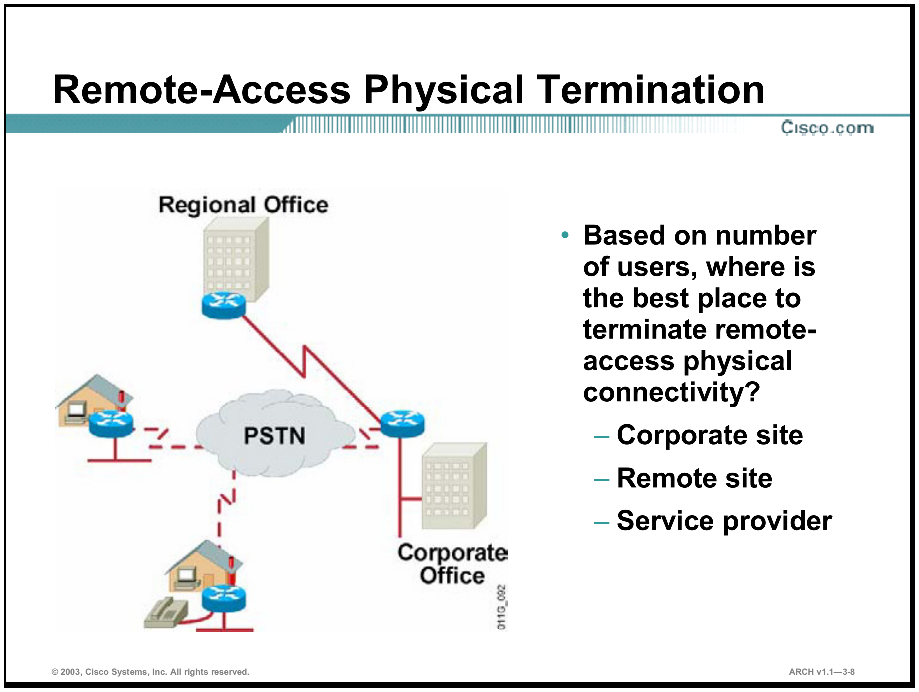

Remote-Access Physical Termination

- Based on number of users, where is the best place to terminate remote-access physical connectivity?
  - Corporate site
  - Remote site
  - Service provider

ARCH v1.1—3-8

The choice on where to terminate the physical remote access connectivity depends on who is doing the termination. If a service provider provides termination, remote users will dial into a local point of presence (POP) to limit toll charges.

If an enterprise is going to provide the remote-access physical termination at one of their sites, you need to answer these questions:

- What are the requirements on the termination ports? Do they have to support voice, data, and fax?
- What is the cost of bringing all the users into a central site, versus the cost of maintaining modem pools in several sites? Where will the connectivity be most reliable?
- How many users are going to simultaneously use the remote-access system?
- Are the users mobile or fixed?
- How many fixed users have access to always-on technology?
- Are sites, or individual users, being terminated?

Once these questions are answered, you can select the physical and data link protocols.

Selecting the Remote-Access Physical Connectivity

The physical connectivity technologies include remote dial-up access and broadband technologies. Dial-up access is provided through modems, cell phones, and ISDN. Broadband technologies include DSLs, cable modems, and satellites. This topic provides guidelines to help you recommend physical layer protocols for remote access.

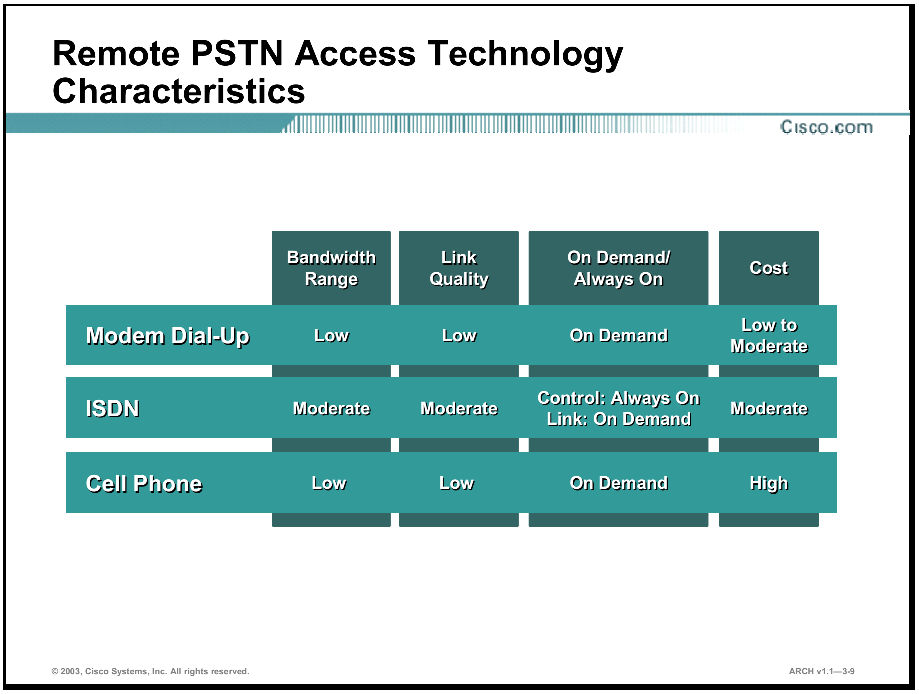

Remote PSTN Access Technology Characteristics

	Bandwidth Range	Link Quality	On Demand/ Always On	Cost
Modem Dial-Up	Low	Low	On Demand	Low to Moderate
ISDN	Moderate	Moderate	Control: Always On Link: On Demand	Moderate
Cell Phone	Low	Low	On Demand	High

Cisco.com
© 2003, Cisco Systems, Inc. All rights reserved.
ARCH v1.1—3-9

These physical layer technologies are available for remote dial-up access to an enterprise site:

- **Modem dial-up:** Analog modems using basic telephone service are asynchronous transmission-based. Basic telephone service is available everywhere, is easy to set up, dials anywhere on demand, and is very inexpensive.
- **ISDN:** ISDN offers digital dial-up connections and a short connection setup time. ISDN is a good solution for telecommuters. Instead of leasing a dedicated line for high-speed digital transmission, ISDN offers the option of dial-up connectivity, incurring charges only when the line is active.
- **Cell phone:** Cell phones use the public cell phone network to access the central site. The primary benefit of using cell phones is mobility, although the expense can be high with limited functionality.

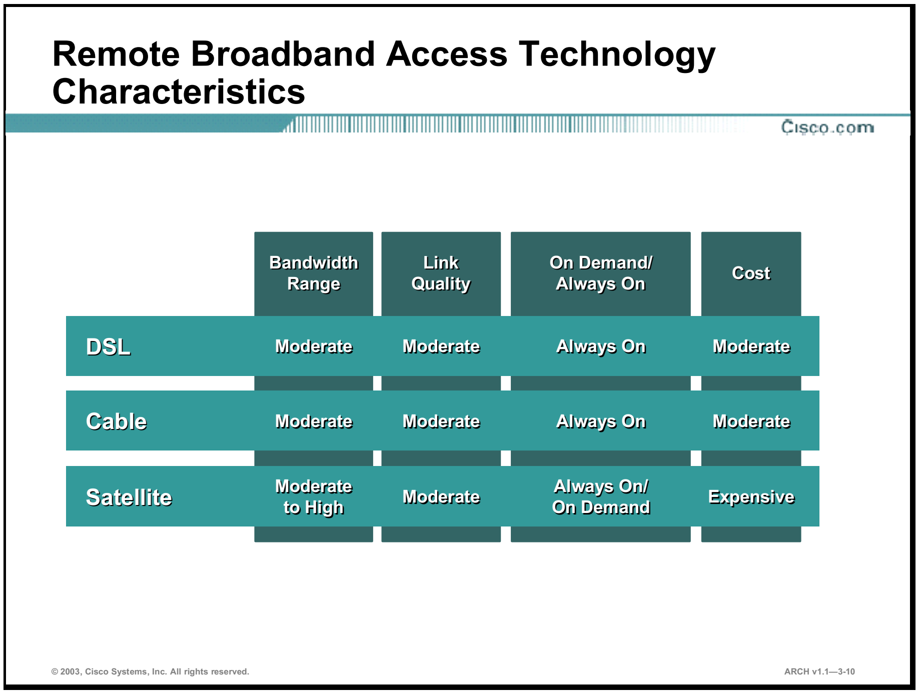

Remote Broadband Access Technology Characteristics

	Bandwidth Range	Link Quality	On Demand/ Always On	Cost
DSL	Moderate	Moderate	Always On	Moderate
Cable	Moderate	Moderate	Always On	Moderate
Satellite	Moderate to High	Moderate	Always On/ On Demand	Expensive

ARCH v1.1—3-10

These physical layer technologies offer broadband remote access:

- **DSL:** Enterprises are increasingly turning to DSL to expand the use of telecommuting, reduce costs, and provide Internet-based services. DSL offers always-on access, allowing users to work at remote offices as if they were on-site.
- **Cable:** Cable is increasingly available. Cable offers always-on access, allowing users to work at remote offices and at home as if they were on-site.
- **Satellite:** Wireless communications devices usually connect to a satellite. A transponder receives and transmits radio signals at a prescribed frequency range. After receiving the signal, a transponder will broadcast the signal at a different frequency. Satellites provide high quality, at a high cost, primarily to support mobility.

Selecting the Remote-Access Protocol

The most common method used to transport packets from user equipment to a termination point is a form of PPP. Each version of PPP has its targeted purpose, but each basically encapsulates the IP packet and delivers it at the other end. This topic provides guidelines to help you recommend data link layer protocols for remote access.

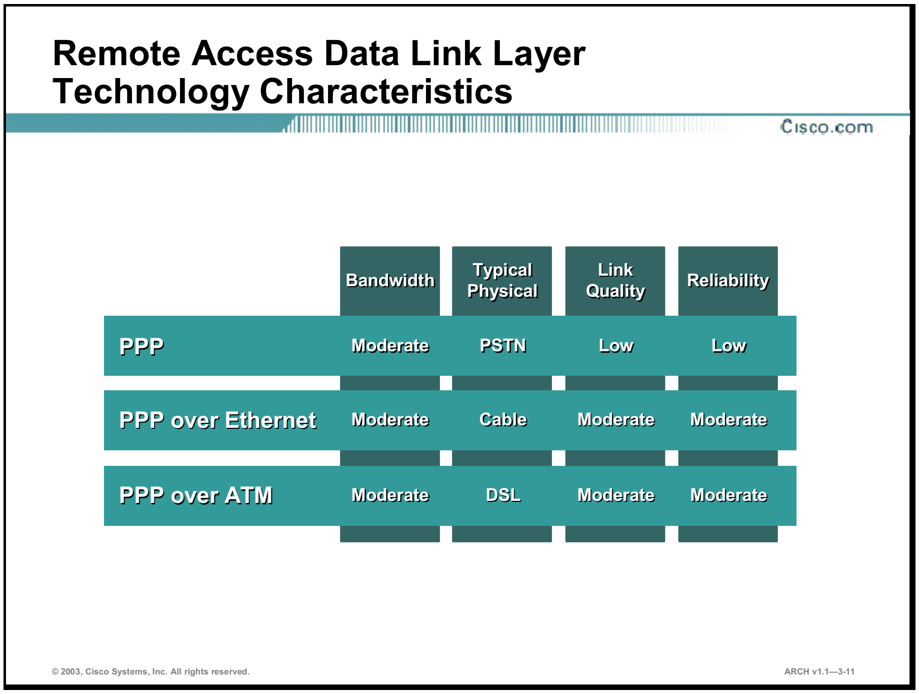

Remote Access Data Link Layer Technology Characteristics

Cisco.com

	Bandwidth	Typical Physical	Link Quality	Reliability
PPP	Moderate	PSTN	Low	Low
PPP over Ethernet	Moderate	Cable	Moderate	Moderate
PPP over ATM	Moderate	DSL	Moderate	Moderate

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-11

These data link protocols are typically deployed for remote-access networks:

- **PPP:** PPP in a remote-access environment defines methods of sending IP packets over circuit lines and is an inexpensive way of connecting PCs to a network. Refer to Requests for Comments (RFCs) 1331 and 1332 for more information about PPP.
- **PPP over Ethernet:** PPP over Ethernet (PPPoE) allows a PPP session to be initiated on a simple Ethernet-connected client. PPPoE can be used on existing customer premise equipment. PPPoE preserves the point-to-point session used by ISPs in the current dial-up model. It is the only protocol capable of running point-to-point over Ethernet without requiring an intermediate IP stack. PPPoE is most often used to connect a host to a cable modem.
- **PPP over ATM:** PPP over ATM (PPPoA) adaptation Layer 5 (AAL5) (specified in RFC 2364) uses AAL5 as the framed protocol, which supports both permanent virtual circuits (PVCs) and switched virtual circuits (SVCs). PPPoA was primarily implemented as part of an asymmetric digital subscriber line (ADSL). It relies on RFC1483, operating in either logical link control-Subnetwork Access Protocol (LLC-SNAP) or virtual circuit multiplexer mode. A customer premise equipment (CPE) device encapsulates the PPP session PPPoA for transport across the ADSL loop and the DSLAM.

Selecting Cisco Access Routing Solutions

For an enterprise that is providing their own remote-access termination, access routing solutions are required at both the remote location and at a central site. The requirements at each site will be different. This topic shows you how to use the Product Advisor to select Cisco access routing solutions for both the remote site and central site, based on specific enterprise requirements.

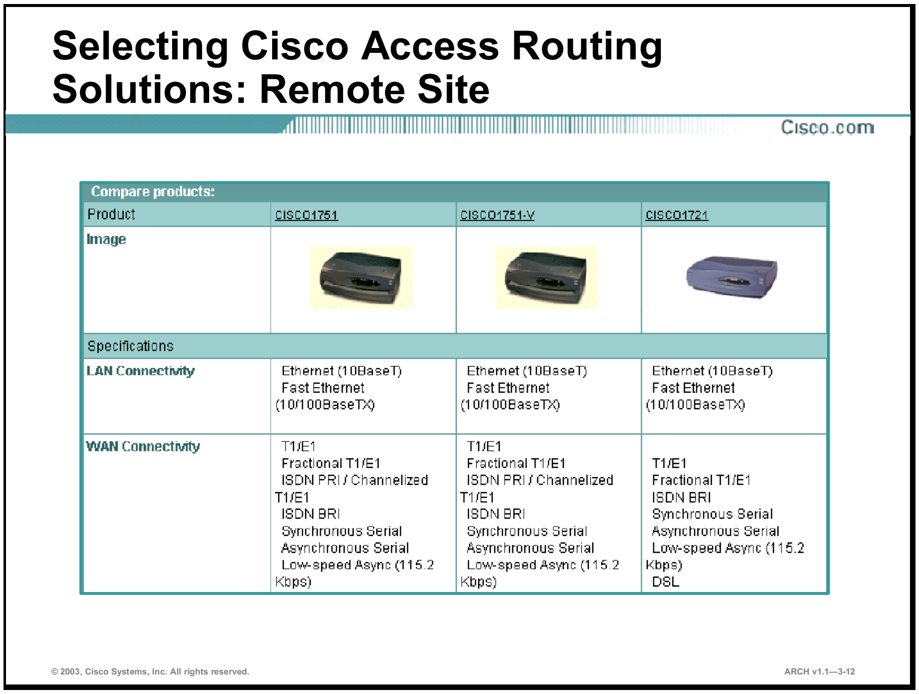

Selecting Cisco Access Routing Solutions: Remote Site

Cisco.com

Compare products:			
Product	CISCO1751	CISCO1751-V	CISCO1721
Image			
Specifications			
LAN Connectivity	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX)	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX)	Ethernet (10BaseT) Fast Ethernet (10/100BaseTX)
WAN Connectivity	T1/E1 Fractional T1/E1 ISDN PRI / Channelized T1/E1 ISDN BRI Synchronous Serial Asynchronous Serial Low-speed Async (115.2 Kbps)	T1/E1 Fractional T1/E1 ISDN PRI / Channelized T1/E1 ISDN BRI Synchronous Serial Asynchronous Serial Low-speed Async (115.2 Kbps)	T1/E1 Fractional T1/E1 ISDN BRI Synchronous Serial Asynchronous Serial Low-speed Async (115.2 Kbps) DSL

ARCH v1.1—3-12

Cisco offers the Product Advisor to help you select the right routing solution for the Enterprise Edge network. The tool operates in two modes: Novice and expert. To access the Product Advisor, go to http://www.cisco.com/en/US/products/products_cisco_product_advisor_tool_launch.html and click the Cisco Product Advisor link. Then click a device category. The Product Advisor will ask you questions to help select routers for particular needs. It does not include all products and features, but provides helpful information to help you select appropriate Cisco products.

The table summarizes the questions to answer to help select the right Cisco remote-access solution for the remote site.

Step	Description	Notes and Comments
1.	Determine the type of environment in which the router will be deployed.	Select one: ■ Telecommuter/home office ■ Small office ■ Medium-sized/branch office
2.	Determine how the router will be used.	Select: ■ To connect employees to the network remotely
3.	Determine the type of configuration required.	Select one: ■ Fixed configuration (less expensive, not highly scalable) ■ Modular configuration (more expensive, highly scalable) ■ No preference
4.	Determine the type of LAN connectivity required.	Select one or more: ■ Ethernet (10BASE-T) ■ Fast Ethernet (10/100) ■ Gigabit Ethernet
5.	Determine the number of LAN ports required.	Select one: ■ 1 ■ 2 ■ more than 2
6.	Select the types of WAN connectivity required.	Select one or more: ■ ISDN PRI[1]/Channelized T1/E1 ■ ISDN BRI[2] ■ Synchronous serial ■ Asynchronous serial ■ DSL ■ ATM ■ Frame Relay ■ and so on
7.	Determine the number of WAN ports required.	Select one: ■ 1 ■ 2 10 ■ 11 30 ■ more than 30

Step	Description	Notes and Comments
8.	Determine if voice will be deployed on the WAN, now or in the future.	Answer Yes or No.
9.	Determine if a redundant power supply is required.	Answer Yes or No.
10.	Determine if a rack-mountable router is required.	Answer Yes or No.
11.	Determine which software features are required now and in the future.	Select one or more: ■ VPN ■ Firewall ■ Internet/WAN access
12.	Select the IOS version for each router you selected.	

[1] PRI = Primary Rate Interface

[2] BRI = Basic Rate Interface

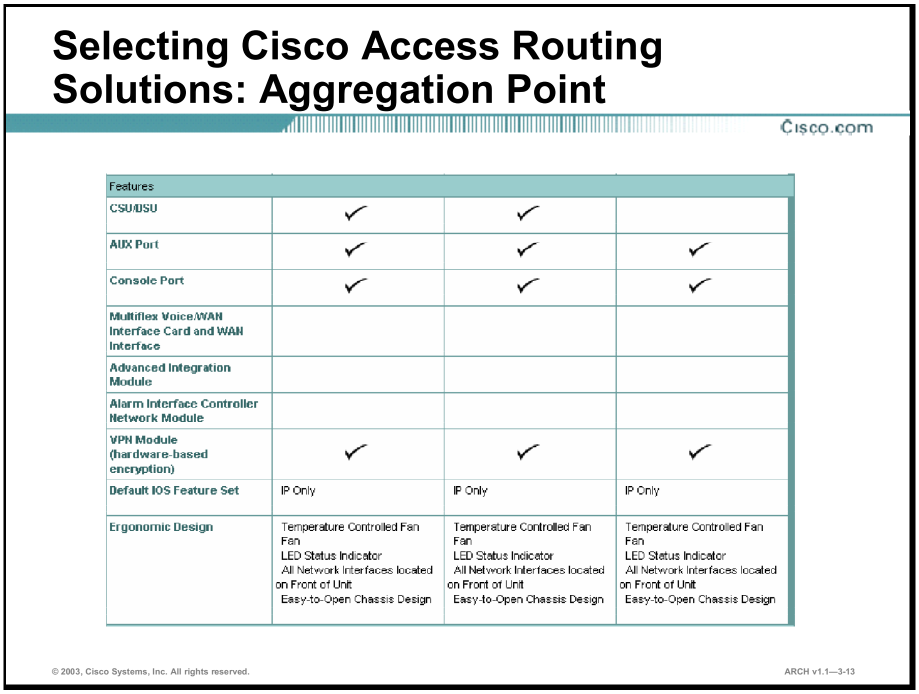

Selecting Cisco Access Routing Solutions: Aggregation Point

Cisco.com

Features			
CSU/DSU	✓	✓	
AUX Port	✓	✓	✓
Console Port	✓	✓	✓
Multiflex Voice/WAN Interface Card and WAN Interface			
Advanced Integration Module			
Alarm Interface Controller Network Module			
VPN Module (hardware-based encryption)	✓	✓	✓
Default IOS Feature Set	IP Only	IP Only	IP Only
Ergonomic Design	Temperature Controlled Fan Fan LED Status Indicator All Network Interfaces located on Front of Unit Easy-to-Open Chassis Design	Temperature Controlled Fan Fan LED Status Indicator All Network Interfaces located on Front of Unit Easy-to-Open Chassis Design	Temperature Controlled Fan Fan LED Status Indicator All Network Interfaces located on Front of Unit Easy-to-Open Chassis Design

ARCH v1.1—3-13

Use the Product Advisor to select remote-access servers for the site.

The table summarizes the first questions to answer to help select the right Cisco remote-access solution for each site.

Step	Description	Notes and Comments
1.	Determine the type of environment in which the router will be deployed.	Select: ■ Corporate office/site
2.	Determine how the router will be used.	Select: ■ To provide WAN aggregation services
3.	Complete Steps 3-12 for the remote site solution.	

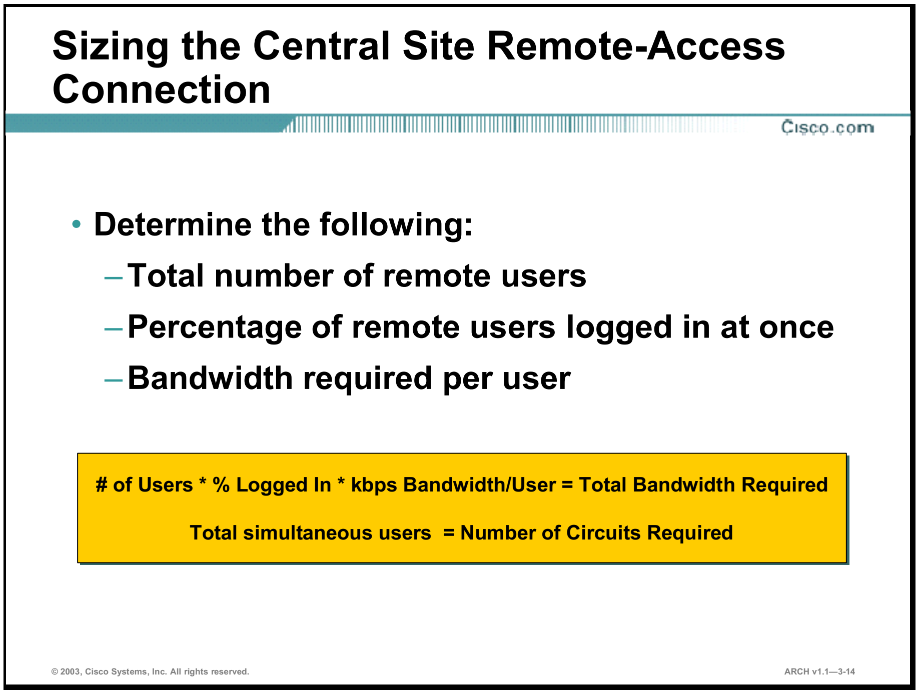

Sizing the Central Site Remote-Access Connection

- **Determine the following:**
  - **Total number of remote users**
  - **Percentage of remote users logged in at once**
  - **Bandwidth required per user**

# of Users * % Logged In * kbps Bandwidth/User = Total Bandwidth Required

Total simultaneous users = Number of Circuits Required

ARCH v1.1—3-14

Planning for peak usage is crucial. The mix and time of connections help determine the peak requirements. For example, 200 modem users calling between 1 p.m. and 3:30 p.m. would require 200 DS-0s (digital signal level 0) or nine PRI circuits.

To determine the peak bandwidth, use the following formula:

- Total number of remote users * % of users logged in at one time (expressed as 0.*nn*) * bandwidth required per user (expressed as kbps) = total bandwidth required

Based on the number of simultaneous users, you can make the assumption that the number of circuits equals the number of simultaneous users. Most telephony circuits are 64 kbps each and, unless there is a methodology for multiplexing multiple steams over the same circuit, the end point will expect a 64-kbps circuit.

Example: Enterprise Remote-Access Designs

A remote-access network design includes data link and physical network technologies, access routing at both the remote site and the central site, and security services. This topic helps you design small, medium, and large enterprise remote-access solutions, given enterprise remote-access needs.

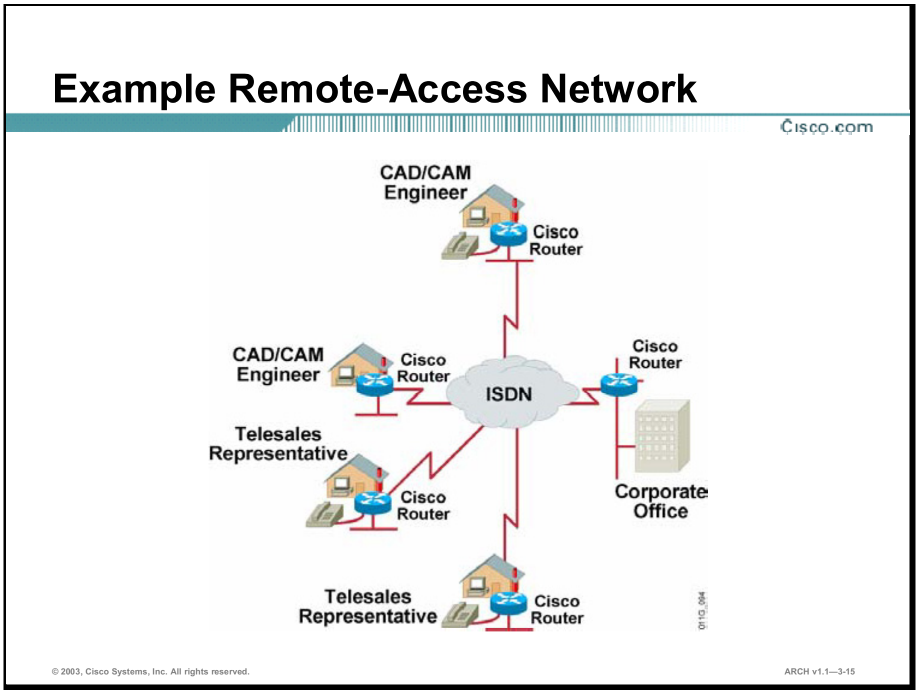

Example Remote-Access Network

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-15

Company Background

The executive management of a tool and die company has identified a number of factors that point to the need for some form of dial-in access for cost reduction and to improve employee morale and productivity.

The company's pool of 25 engineers spends over 90 percent of its time sitting at a workstation working on computer-aided design and manufacturing program-based projects. Some engineers are on part-time contracts, yet the company is paying for office space to put them in front of a PC all day. Some of the full-timers put in long hours and aren't happy about the time spent away from home. Others would work more if they could do it from home.

Additionally, the company employs five telesales representatives and two customer service representatives who, using their phones and PCs, almost never leave their cubes. They are questioning why they must deal with their ugly commutes when they could do it all from home. In addition, company management is wondering why it pays for their office space.

Remote-Access Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the WAN?	Hub and spoke	ISDN is a hub-and-spoke technology, in which each user authenticates with the hub to use the services in the network.
What service provider will be selected?	Local phone company	The local phone company is capable of providing the services and geographic coverage required.
What data link layer protocol will be used?	PPP	PPP is easy to configure and maintain.
What physical network media will be used?	ISDN (always on)	ISDN provides the bandwidth requirements and the always-on feature needed for continuous remote access.
Which Cisco products will be used?	Engineers: Cisco 802[1] Telesales representatives: Cisco 804* Central site: Cisco 3640 with T1 controller module for PRI*	The 800 series routers provide ISDN remote access. The Cisco 802 and 804 include integrated NT-1, and the Cisco 804 provides phone ports for the telesales representatives. The 3600 series provides sufficient features and modularity to support current needs and future expansion.
Which routing protocols will be used?	OSPF (used on backbone network)	OSPF routing supports the remote routing needs of this application.
What IP addressing scheme will be used?	DHCP is used to dynamically assign addresses	Automatic IP addressing simplifies administration of the remote sites.

[1] The example platforms are accurate as of the date this course was published.

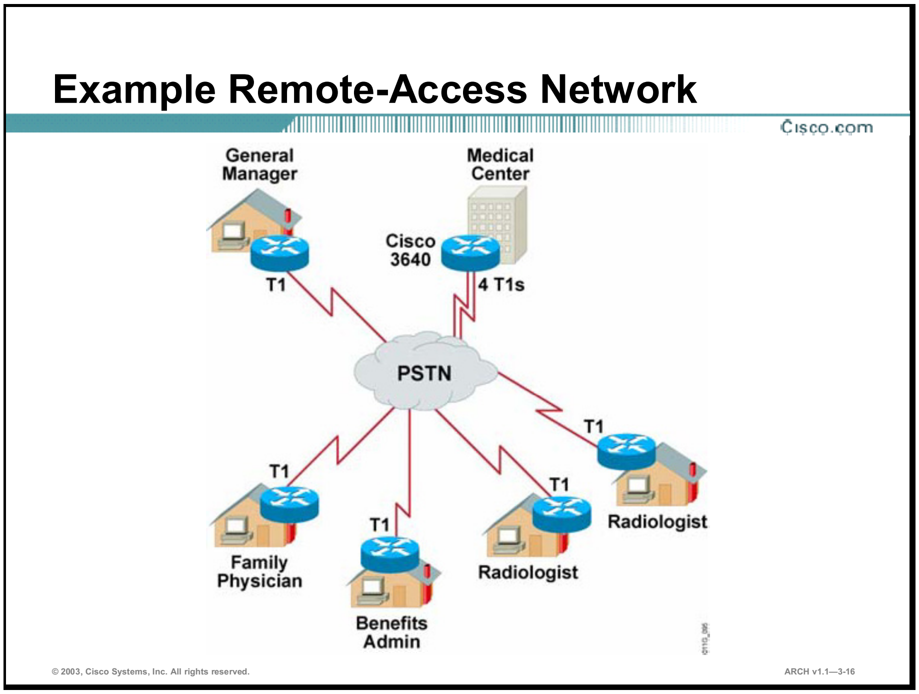

Example Remote-Access Network

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-16

Company Background

The Welleville Medical Center employs a number of highly skilled, highly paid professionals who spend much of their time driving to and from the medical center simply to access medical records, images, and files. Welleville is feeling the financial pressures of the health care industry and must find a way to reduce the diagnosis time of its patients and increase the productivity of its professionals.

Remote-Access Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the WAN?	Star topology	A star topology is sufficient for this application, due to the intermittent access needs.
What service provider will be selected?	Local ISDN provider	The local provider can provide the head end for the ISDN termination and is willing to extend the network to the corporate office over a PRI.
What data link layer protocol will be used?	PPP	PPP is the protocol that the ISDN provider offers.
What physical network media will be used?	ISDN Multi-BRI (MBRI)	ISDN provides for up to a T1 worth of bandwidth.
Which Cisco products will be used?	Medical center: Cisco 1700 series[1] Central site: Cisco 3600	1720 series routers were provided to the medical center's insurance benefits administrator and three of the family physicians for after-hours access to the center's patient records. The Cisco 3600 supports routing of IP, IPX, and AppleTalk protocols, and features EIGRP for integration with the medical center's existing network backbone.
Which routing protocols will be used?	EIGRP	EIGRP was selected for integration with existing network backbone.
What IP addressing scheme will be used?	DHCP is used to dynamically assign addresses	Automatic IP addressing simplifies administration for the remote sites.

[1] The example platforms are accurate as of the date this course was published.

Summary

This topic summarizes the key points discussed in this lesson.

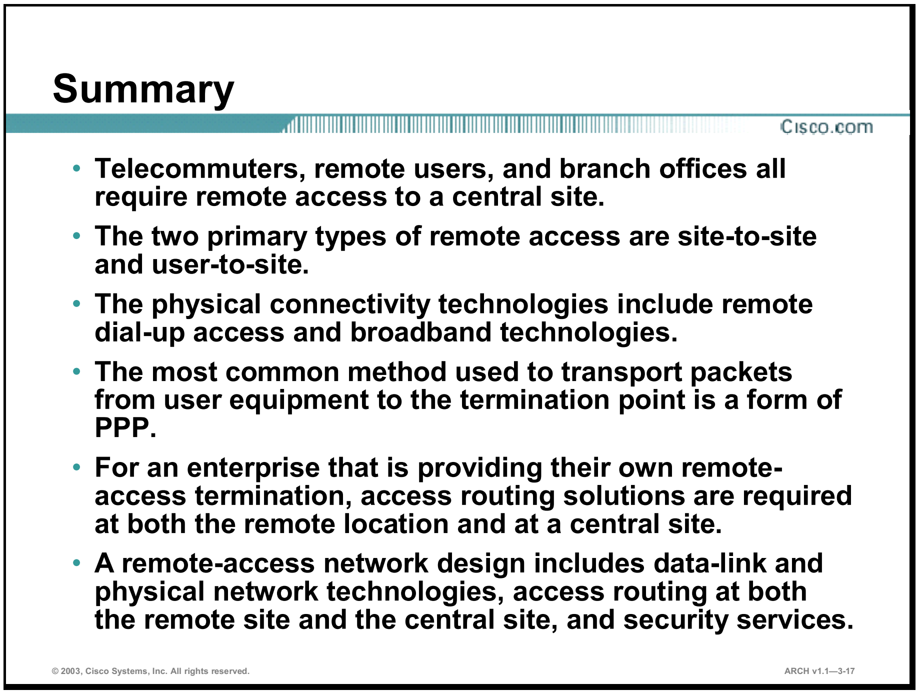

Summary

- Telecommuters, remote users, and branch offices all require remote access to a central site.
- The two primary types of remote access are site-to-site and user-to-site.
- The physical connectivity technologies include remote dial-up access and broadband technologies.
- The most common method used to transport packets from user equipment to the termination point is a form of PPP.
- For an enterprise that is providing their own remote-access termination, access routing solutions are required at both the remote location and at a central site.
- A remote-access network design includes data-link and physical network technologies, access routing at both the remote site and the central site, and security services.

ARCH v1.1—3-17

References

For additional information, refer to these resources:

- *Remote-Access Networking* at http://www.cisco.com/warp/public/779/largeent/learn/topologies/remote_access.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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) Why do the remote requirements of an enterprise network cause the aggregate WAN charge to be exaggerated?

A) because each remote site must connect to the others

B) because all remote sites require the same size WAN connection

C) because they usually involve a large number of remote users or sites

D) because each remote user requires a separate connection, regardless of location

Q2) When there are a group of users in the same vicinity that can share an on-demand connection to their local branch office or central site, the _____ is appropriate.

A) virtual private network

B) point of presence model

C) site-to-site remote-access model

D) service provider termination remote-access model

Q3) What is the primary benefit of using the cell phone network for remote access?

A) mobility

B) reduced cost

C) high bandwidth

D) high link quality

Q4) Which cost component of a remote-access solution does the use of dial-up PPP eliminate?

A) switch

B) modem

C) leased lines

D) access router

Q5) For which Product Advisor question is Gigabit Ethernet a valid response?

A) software features

B) configuration type

C) LAN connectivity

D) WAN connectivity

Q6) Which three items are used to calculate the total bandwidth required for remote access? (Choose three.)

A) number of users

B) bandwidth required per site

C) bandwidth required per user

D) % of users logged in at one time

E) number of concurrent connections

Q7) Which technology is easy to configure and maintain for remote access?

A) PPP

B) SLIP

C) PPPoE

D) Frame Relay

Q8) What are two reasons an enterprise might consider providing remote access for employees? (Choose two.)

A) reduced office costs

B) reduced staffing costs

C) reduced equipment costs

D) improved data reliability

E) improved employee morale

Quiz Answer Key

Q1) C

Relates to: Enterprise Needs for Remote Access

Q2) C

Relates to: Selecting the Remote Access Type and Termination

Q3) A

Relates to: Selecting the Remote Access Physical Connectivity

Q4) C

Relates to: Selecting the Remote Access Protocol

Q5) C

Relates to: Selecting Cisco Access Routing Solutions

Q6) A, C, D

Relates to: Selecting Cisco Access Routing Solutions

Q7) A

Relates to: Example: Enterprise Remote Access Designs

Q8) A, E

Relates to: Example: Enterprise Remote-Access Designs

Designing the Internet Connectivity Module

Overview

Not long ago, enterprises owned and operated numerous networks to deliver multiple services to their customers and employees. Voice communications required the telephone network. Video broadcasting utilized a broadband cable network to broadcast video onto the network. To transport computer application data, enterprises built data networks.

Relevance

Enterprises can now share information throughout the network that previously existed in isolation and was accessed through one medium. Customers, partners, and remote employees can access the information when connected across the Internet.

Objectives

Upon completing this lesson, you will be able to design the Internet Connectivity module, given enterprise needs to access the Internet. This includes being able to meet these objectives:

- Identify typical enterprise needs for the Internet
- Explain when to use NAT for Internet connectivity
- Present Cisco ISP connectivity solutions and design guidelines to support availability and load balancing
- Design small, medium, and large enterprise Internet access solutions, given enterprise Internet needs

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

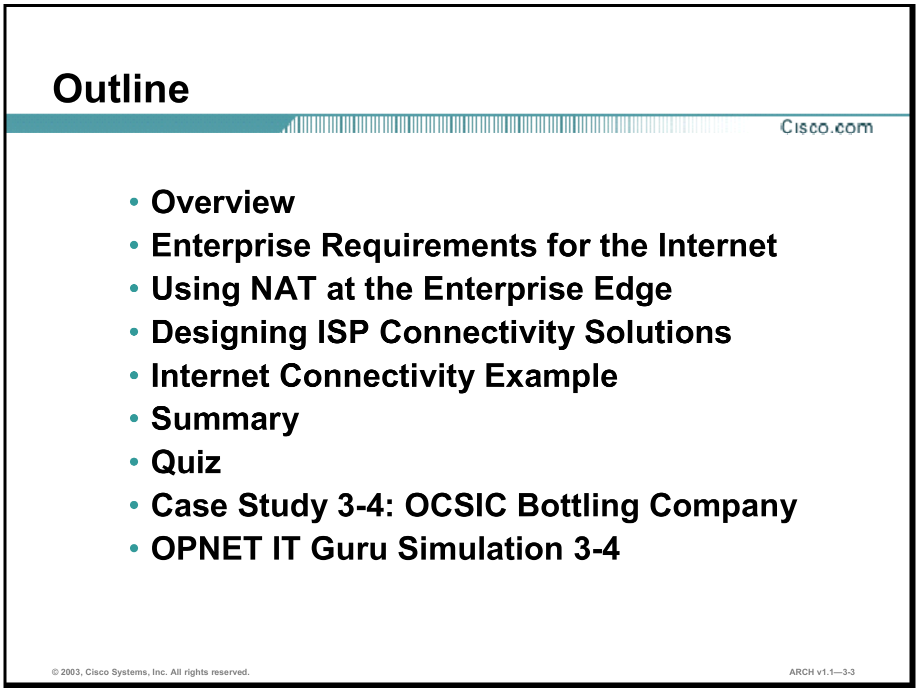

Outline

- Overview
- Enterprise Requirements for the Internet
- Using NAT at the Enterprise Edge
- Designing ISP Connectivity Solutions
- Internet Connectivity Example
- Summary
- Quiz
- Case Study 3-4: OCSIC Bottling Company
- OPNET IT Guru Simulation 3-4

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-3

Enterprise Requirements for the Internet

Enterprises frequently require access from the internal network to the Internet, and may provide access to public servers for outside users over the Internet. This topic identifies typical enterprise needs for the Internet.

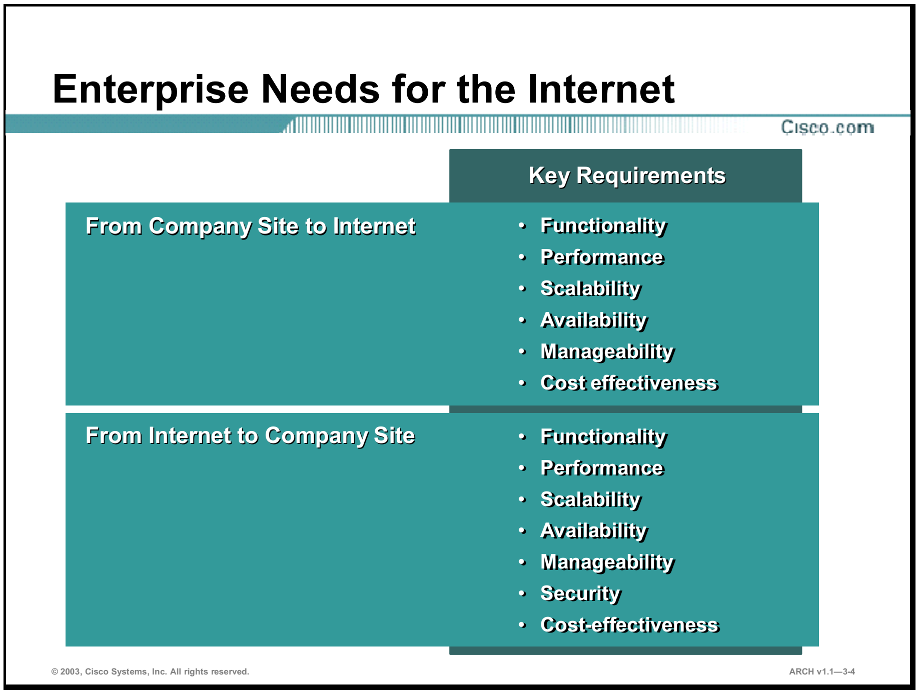

Enterprise Needs for the Internet

Cisco.com

	Key Requirements
From Company Site to Internet	• Functionality • Performance • Scalability • Availability • Manageability • Cost effectiveness
From Internet to Company Site	• Functionality • Performance • Scalability • Availability • Manageability • Security • Cost-effectiveness

ARCH v1.1—3-4

Enterprises must design to support connectivity between company-owned sites and the Internet, and from the Internet to the company-owned sites. The requirements in both cases are similar. The Internet-to-company site situation requires an increased level of security to ensure that unauthorized users do not gain access to the corporate network.

The first step in developing a network design for the Internet Connectivity module is to determine connectivity requirements. Then you must ask these questions:

- Is a single Internet service provider (ISP) required or are two connections to the same or different ISPs needed?
- If multiple ISPs are used, how will load balancing be done?
- Which routing protocol will advertise the Internet internally, and advertise publicly available subnets externally?
- Is NAT or port address translation (PAT) required at a router or transition device between the public and corporate network?
- What security measures are required to protect the corporate network?

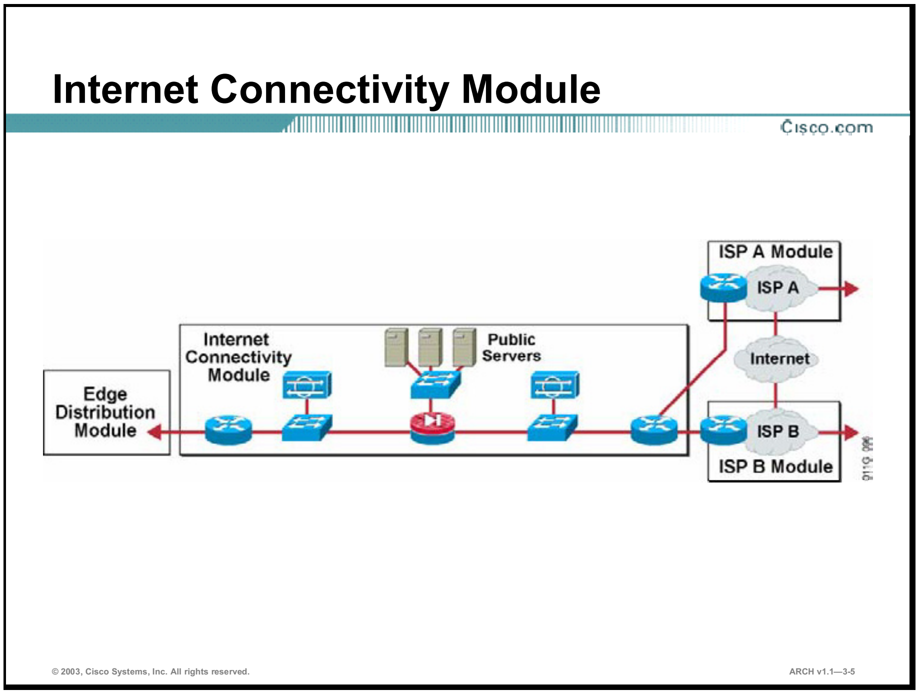

Internet Connectivity Module

Cisco.com

ARCH v1.1—3-5

The Internet Connectivity module of the Enterprise Composite Network Model provides services to enterprise users who want to access the Internet and to outside users who want access to company public services, typically web servers or application servers.

The primary components at the central site are the Internet access device (a router, access server, or VPN equipment), a publicly available network with publicly accessible servers, and, optionally, firewalls and IDSs.

Note Refer to the "Designing Security Services" module to learn how to design security services for Internet access and the entire enterprise.

Using NAT at the Enterprise Edge

NAT is designed to simplify and conserve IP addresses, as it enables private IP internetworks that use nonregistered IP addresses to connect to the Internet. This topic describes when to use NAT for Internet connectivity.

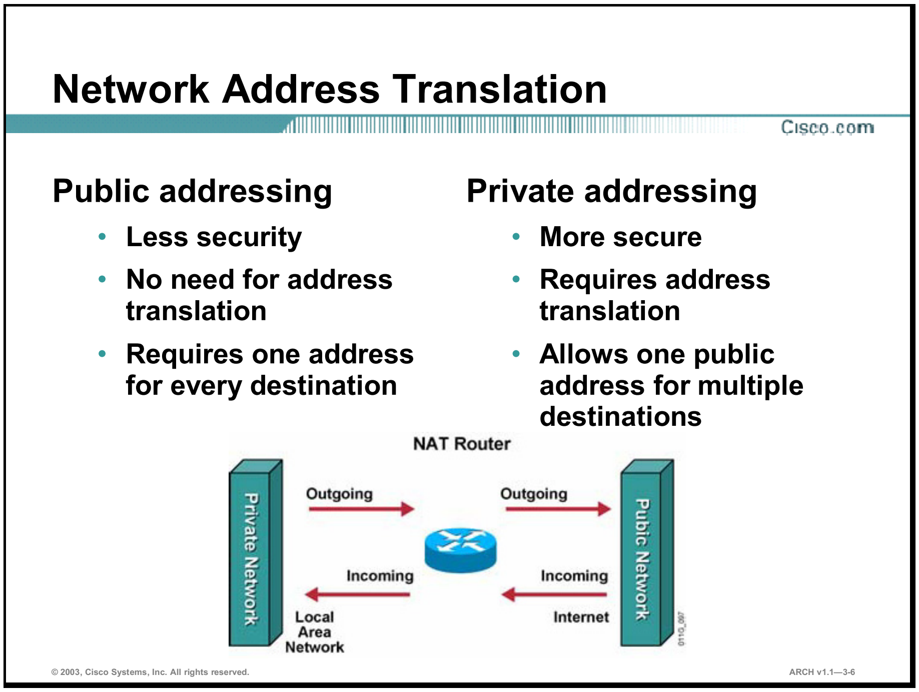

Network Address Translation

Cisco.com

Public addressing

- **Less security**
- **No need for address translation**
- **Requires one address for every destination**

Private addressing

- **More secure**
- **Requires address translation**
- **Allows one public address for multiple destinations**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-6

NAT operates on a router, usually connecting two networks together, and translates private addresses in the internal network into legal public addresses before forwarding packets onto another network. As part of NAT functionality, you can configure NAT to advertise only one address for the entire network to the outside world. The figure lists the major features of public and private addresses.

NAT provides additional security, effectively hiding the entire internal network from the world behind that address. NAT takes these forms:

- **Static NAT**: Maps an unregistered IP address to a registered IP address on a one-to-one basis. Static NAT is particularly useful when a device needs to be accessible from outside the network.
- **Dynamic NAT**: Maps an unregistered IP address to a registered IP address from a group of registered IP addresses.
- **Overloading**: A form of dynamic NAT that maps multiple unregistered IP addresses to a single registered IP address by using different ports. Known also as PAT, single-address NAT, or port-level multiplexed NAT.
- **Overlapping**: When the IP addresses used on your internal network are registered IP addresses in use on another network, the router must maintain a lookup table so that it can intercept and replace these registered addresses with unique IP addresses. The NAT router must translate the internal addresses into registered unique addresses. It must also translate the external registered addresses to addresses that are unique to the private network through static NAT or with dynamic NAT with DNS.

Designing ISP Connectivity Solutions

There are two methods to connect to an ISP: a single run or multi-homed method. The simplest method is to provision a single connection to the ISP. If redundancy and load sharing are a requirement, a multi-homed system is required. This topic presents Cisco ISP connectivity solutions and design guidelines to support availability and load balancing.

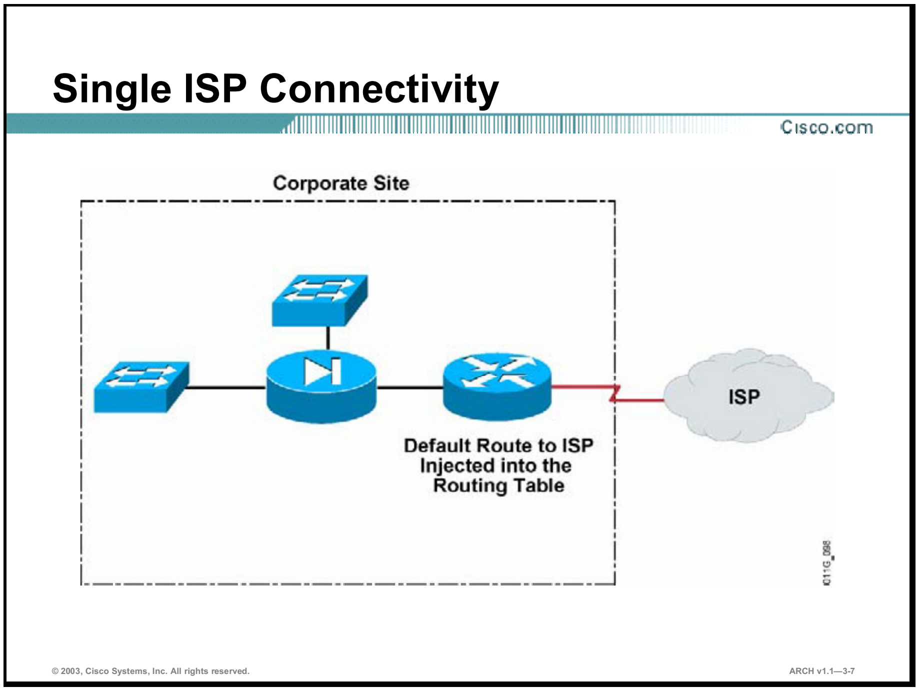

Single ISP Connectivity

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-7

The ISP connectivity that has no redundancy and is the easiest to configure is the single run, or single-homed system, which connects the ISP to the corporate site.

When implementing a single-homed system, the routing decision is to use default routes pointing to the network that connects the site to the ISP. The default route is then advertised throughout the corporate site, so that any packets with an unknown destination are forwarded to the ISP. The IP addressing is accomplished with real addressing, if it is available, or through NAT software. If NAT is used, the publicly available servers must have addresses that are statically mapped either to a public address or to a PAT table. The ISP will use static routes that point to the enterprise site and then advertise those routes within their network and to those with whom they have peering arrangements.

The questions to ask when implementing a single-homed connection are:

- What are the consequences if the Internet connection is lost?
- Can the enterprise afford the consequences of an outage?
- Will public addressing or private addressing be used in the network?
- If private addressing is used inside, how many public addresses are needed to support the hosts that need static addressing? How many addresses are needed in the address pool for the users?
- When selecting the ISP, what services and support do they provide?

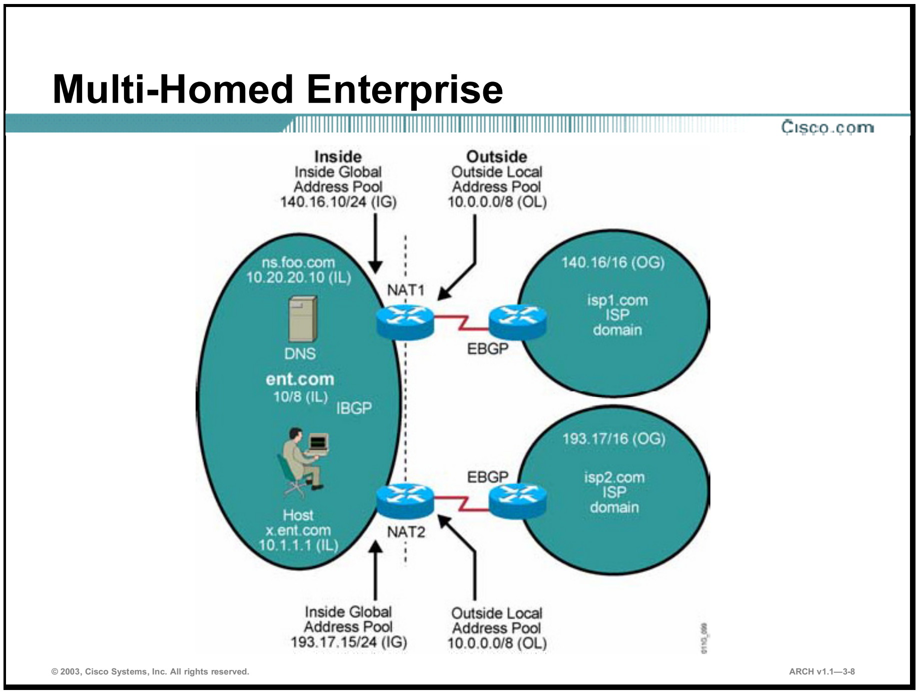

Multi-Homed Enterprise

ARCH v1.1—3-8

ISP multi-homing solutions improve availability and load balancing for WANs that use the Internet. Multiple connections, known as multi-homing, reduce the chance of a potentially catastrophic shutdown if one of the connections should fail.

In addition to maintaining a reliable connection, multi-homing allows a company to perform load balancing by lowering the number of computers connecting to the Internet through any single connection. Distributing the load through multiple connections optimizes performance and can significantly decrease wait times.

Multi-homed networks are often connected to several different ISPs. Each ISP assigns an IP address (or range of IP addresses) to the company. Routers use Border Gateway Protocol (BGP) to route between networks using different protocols. In a multi-homed network, the router utilizes Internal BGP (IBGP) on the stub domain side and External BGP (EBGP) to communicate with other routers.

Multi-homing really makes a difference if one connection to an ISP fails. As soon as the router assigned to connect to that ISP determines that the connection is down, it will reroute all data through one of the other routers.

There are two common methods used when implementing a multi-homed ISP solution. One method is to perform auto-route injection into the network to increase the availability and backup of the Internet connection. The second method is to perform non-direct BGP peering to enable load balancing.

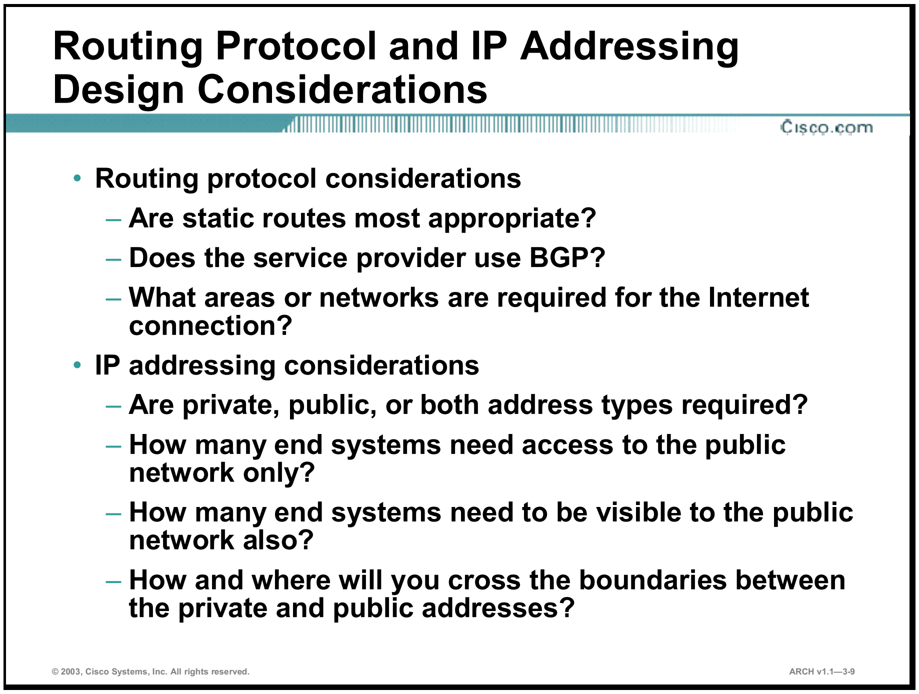

Routing Protocol and IP Addressing Design Considerations

- Routing protocol considerations
  - Are static routes most appropriate?
  - Does the service provider use BGP?
  - What areas or networks are required for the Internet connection?
- IP addressing considerations
  - Are private, public, or both address types required?
  - How many end systems need access to the public network only?
  - How many end systems need to be visible to the public network also?
  - How and where will you cross the boundaries between the private and public addresses?

ARCH v1.1—3-9

For Internet connectivity, use either static routes or IBGP based on these criteria:

- Use static routes when you want lower overhead and when only one exit point exists.
- Use IBGP with multiple exit points and when multi-homing is required.

IP addressing considerations relate to the need for private or public addressing. IP addressing considerations are similar to those for the WAN and Remote Access modules.

Internet Connectivity Example

An Internet Connectivity example provides guidelines to help you design your own solution. This topic provides an example of an enterprise Internet access solution, given enterprise Internet needs.

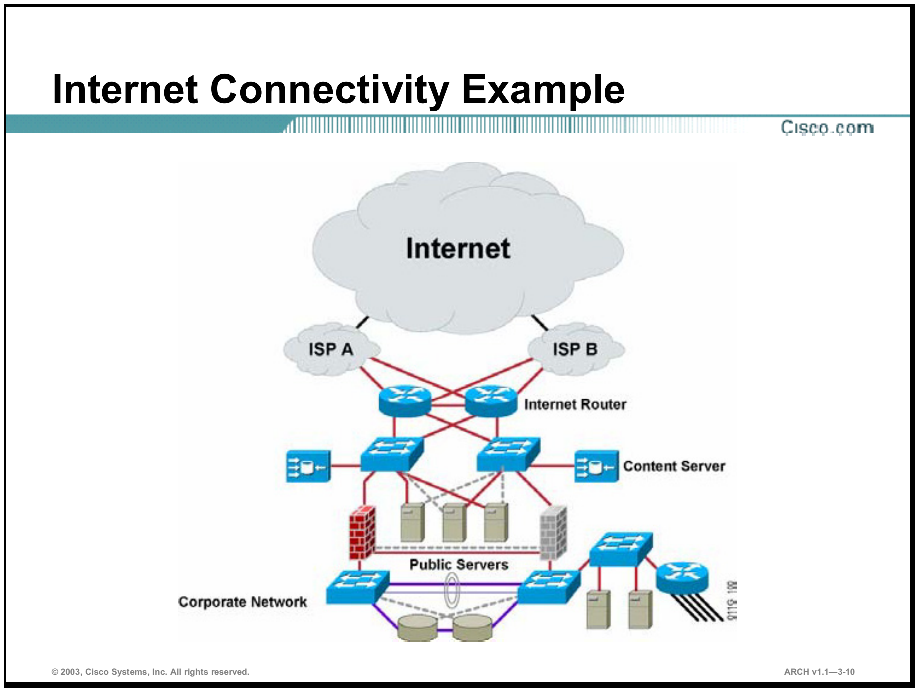

Internet Connectivity Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—3-10

Company Description

Jessie and partners have outgrown their current single ISP connection and want to move to a multi-homed network connection using two different ISPs to provide them with reliable, redundant service. They currently have 3400 employees and, at any given time, 400 people need simultaneous access to the Internet. The company recently developed a new process for ordering their products online and expects a large demand for their online ordering and online technical support services.

Internet Connectivity Solution

The IT staff has decided to use ISP A and ISP B as their service providers and each provider is willing to create BGP connections into their network. The selected ISPs provide solutions and services, POPs, and high reliability and support. Both ISPs are willing to support BGP peer advertising of both public addresses that have been assigned to Jessie and partners.

Jessie and partners have received two individual class C addresses from the InterNIC that they want both of the ISPs to advertise.

The company expects a large response from the outside to the publicly accessible or corporate network, but it will be difficult to judge without baselining the time it takes to use the products.

Basic security is currently provided by firewalls that are used to keep outsiders out of the corporate network.

Services provided will be FTP for drivers and support documentation and HTTP access for papers and documentation. The internal users are expected to have the same kinds of traffic, that is, FTP and HTTP.

NAT will be used in two ways: overlapping for most employees and several key systems will have static addresses. The public addresses will be used in the isolated LAN and on the routers that connect to the ISPs. There will be no public addressing behind the firewalls.

The company will inject BGP routes into the routing table and redistribute them to the rest of the world.

Summary

This topic summarizes the key points discussed in this lesson.

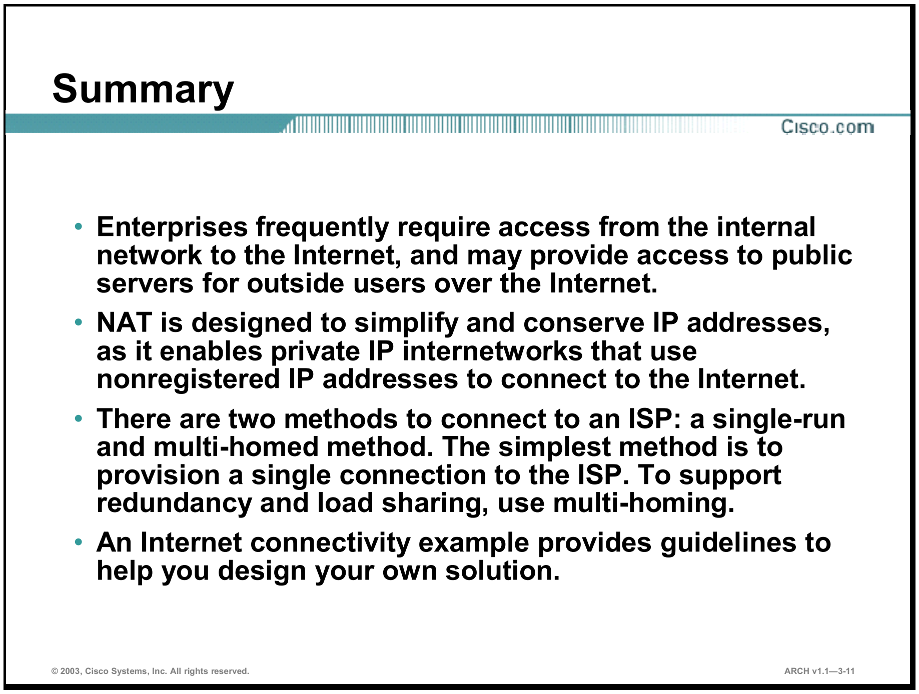

Summary

Cisco.com

- Enterprises frequently require access from the internal network to the Internet, and may provide access to public servers for outside users over the Internet.
- NAT is designed to simplify and conserve IP addresses, as it enables private IP internetworks that use nonregistered IP addresses to connect to the Internet.
- There are two methods to connect to an ISP: a single-run and multi-homed method. The simplest method is to provision a single connection to the ISP. To support redundancy and load sharing, use multi-homing.
- An Internet connectivity example provides guidelines to help you design your own solution.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—3-11

References

For additional information, refer to these resources:

- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 3-4: OCSIC Bottling Company
- OPNET IT Guru Simulation 3-4

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 requirement is of more concern for Internet-to-company connectivity than for company-to-Internet connectivity?

A) security

B) scalability

C) performance

D) cost-effectiveness

Q2) Which two types of NAT might you deploy in the Internet Connectivity module? (Choose two.)

A) physical

B) dynamic

C) perpetual

D) virtualistic

E) overloading

Q3) Which two advantages does multi-homing provide? (Choose two.)

A) reliability

B) reduced costs

C) load balancing

D) simplified routing

E) reduced equipment needs

Q4) Which Internet Connectivity module implementation choice supports load balancing?

A) firewalls

B) multi-homing

C) mirrored servers

D) redundant access routers

Quiz Answer Key

Q1) A

Relates to: Enterprise Requirements for the Internet

Q2) B, E

Relates to: Using NAT at the Enterprise Edge

Q3) A, C

Relates to: Designing ISP Connectivity Solutions

Q4) B

Relates to: Internet Connectivity Example

Case Study 3-4: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

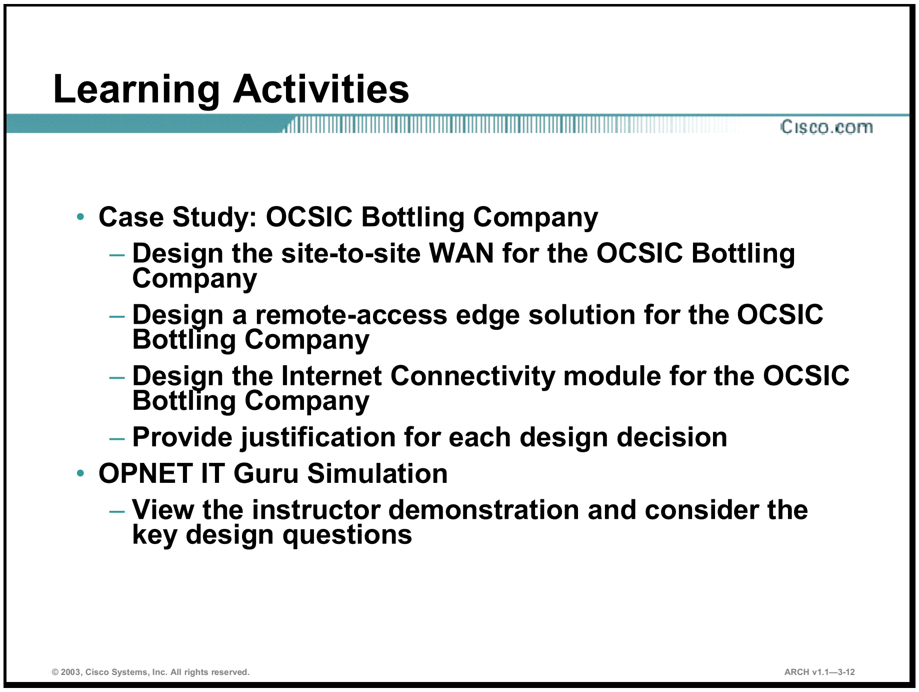

Learning Activities

- Case Study: OCSIC Bottling Company
  - Design the site-to-site WAN for the OCSIC Bottling Company
  - Design a remote-access edge solution for the OCSIC Bottling Company
  - Design the Internet Connectivity module for the OCSIC Bottling Company
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

ARCH v1.1—3-12

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

In this exercise, you will design the enterprise edge network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design the site-to-site WAN for the OCSIC Bottling Company
- Design a remote-access edge solution for the OCSIC Bottling Company
- Design the Internet Connectivity module for the OCSIC Bottling Company

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

About the OCSIC Bottling Company

Following is additional information you will need about the OCSIC Bottling Company WAN to complete the case study exercise.

North American Plant-Headquarters Wide Area Network

The networks at each plant were developed at different times and, therefore, are quite different. The support and maintenance expenses have grown out of control and are placing an increasing burden on IT. The company wants to replace the existing networks with a complete IP-based solution.

The following figure describes the current network between the North American district and regional plants, and the headquarters office.

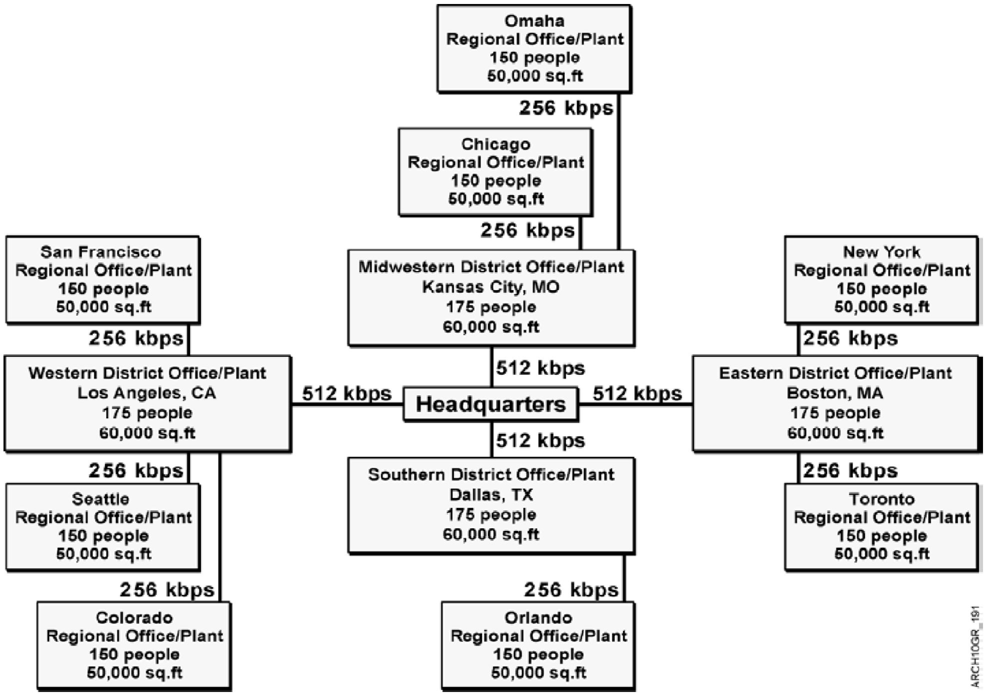

The plants currently use a Frame Relay over Fractional T1 WAN to the headquarters office. While relatively inexpensive, the company finds these links too slow.

The table describes the traffic volume on the network from each North American plant office to the data center at headquarters.

Building From	Building To	Volume (kbps) (to nearest 50 kbps)
District Office/Plant	Building F (server farm)	600 kbps
Regional Office/Plant	District Office/Plant	450 kbps

Remote-Access Requirements

The salespeople work out of their local plant, but are often on the road visiting with existing and potential customers. From the road, they need access from their PCs to the service access point (SAP) and custom Oracle applications. Each office has approximately 20 to 30 users that require remote access to the network. Security is key for those users as much of the data is financial in nature.

It is assumed that 300 people are on the road at one time. If 40 percent are active at one time, 120 ports will be required.

All remote access is through the headquarters office.

Internet Connectivity Requirements

The company wants all employees at headquarters to have access to the Internet. The projected traffic is 512 kbps.

International Manufacturing, Distribution, and Sales Plants

The international manufacturing, distribution, and sales plants are similar to their North American equivalents. That is, each location supports 150 to 175 users, in a 50–60,000 square foot facility. However, each plant is independently owned and operated.

The South American manufacturing, distribution, and sales plants are located in these cities:

- Sao Paolo, Brazil
- Santiago, Chile
- Caracas, Venezuela
- San Jose, Costa Rica
- Mexico City, Mexico

The European manufacturing, distribution, and sales plants are located in these cities:

- Hanover, Germany
- London, England
- Paris, France
- Rome, Italy
- Dublin, Ireland
- Madrid, Spain
- Prague, Czech Republic

The Asia-Pacific manufacturing, distribution, and sales plants are located in these cities:

- Singapore
- Tokyo, Japan
- Hong Kong
- Taiwan, China
- Sydney, Australia

International Plant Networks and Applications

Since the international plants are independently owned and operated, they have their own networks, applications, and IT staff. The offices do not communicate among themselves. To meet its own corporate goals, OCSIC wants to enhance communications and information sharing with its distributors, which require a higher-speed, inexpensive connection between the headquarters and each plant.

Today, distributors use high-speed dial-up connections to headquarters, as their needs are intermittent. However, many have very high toll calls to support this connectivity. OCSIC believes that the Internet will provide an ideal mechanism to share applications on an as-needed basis with these companies.

Security will be a key consideration as the company implements a networking solution between its headquarters and its international partners.

The international plants need access to these applications at headquarters:

- SAP
- E-mail
- Extranet web site (planned)

The table describes the anticipated traffic volume on the network from international location to the data center at headquarters.

Building From	Building To	Estimate Volume (kbps)
Each Office/Plant (today)	Building F (server farm)	90 kbps
Each Office/Plant (with extranet)	Building F (server farm)	170 kbps

Task 1: Design the Wide Area Network

Complete these steps:

Step 1 Refer to the global network diagram for the company that includes the headquarters location, district offices, regional offices, and international plants.

Step 2 Refer to the country-level network diagram for the company that identifies the locations in North America and the WAN links between the locations. Make a copy of the diagram for this exercise.

Step 3 Complete the table to design the details about the WAN. Assume that the service provider selected offers all of the popular wide-area networking services available on the market today, and their service level agreement is acceptable.

Design Question	Decision	Justification
What topology will be used for the WAN?		
What service provider will be selected?		
What data-link layer protocol will be used?		
What physical network media will be used?		
What additional services would you select for each WAN link? ■ If you selected Frame Relay, choose the number of ports, committed information rate (CIR), committed burst (Bc), excess burst (Be), transmission convergence (TC), and maximum burst size. ■ If you selected ATM, choose the service class, either CBR, ABR, UBR, RT-VBR, or NRT-VBR. ■ If you selected PPP, the services depend on the Layer 1 technology you selected.		

Design Question	Decision	Justification
Which Cisco products will be used?		
Which routing protocols will be used?		
What IP addressing scheme will be used?		

Step 4 Update the WAN network diagram to indicate the products and services you selected at each location.

Task 2: Design the Remote-Access Network

Complete these steps:

Step 1 Complete the table to design the details about the remote-access network. Assume that the service provider selected offers all of the popular remote-access services available on the market today, and their service level agreement is acceptable.

Design Questions	Decision	Justification
What topology will be used for the WAN?		
What service provider will be selected?		
What data-link layer protocol will be used?		
What physical network media will be used? How many trunks are required?		
Which Cisco products will be used?		
Which routing protocols will be used?		
What IP addressing scheme will be used?		

Step 2 Update the WAN network diagram to indicate the products and services you selected at each location.

Step 3 Is authentication required for remote users? Why or why not?

Step 4 Are access control lists required for remote users? Why or why not?

Step 5 Are firewalls or intrusion detection systems required? Why or why not?

Task 3: Design the Internet Connectivity Module

Complete these steps:

Step 1 Complete the table to design the details about the Internet Connectivity module network. Assume that the service provider selected offers all of the popular WAN and Internet services available on the market today, and their service level agreement is acceptable.

Design Questions	Decision	Justification
What topology will be used for the WAN?		
What service provider will be selected?		
What data-link layer protocol will be used?		
What physical network media will be used?		
Which Cisco products will be used?		
Which routing protocols will be used?		
What IP addressing scheme will be used?		

Step 2 Update the WAN network diagram to indicate the products and services you selected at each location.

Task 4: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- A completed intermediate design for the OCSIC Bottling Company site-to-site WAN that includes a WAN topology, routing protocol, data link layer technologies, physical layer technologies, and network devices
- A completed intermediate design for one OCSIC Bottling Company remote-access WAN that includes a remote-access WAN topology, routing protocol, data link layer technologies, physical layer technologies, and network devices
- A completed intermediate design for the OCSIC Bottling Company Internet Connectivity module that includes a Corporate Internet topology, routing protocol, data link layer technologies, physical layer technologies, and network devices

OPNET IT Guru Simulation 3-4

This simulation demonstrates different WAN, remote-access, and Internet connectivity options for the OCSIC Bottling Company. Specifically:

- The first simulation demonstrates three different Frame Relay WAN scenarios.
- The second simulation demonstrates a remote-access scenario.
- The third simulation demonstrates three different Internet connectivity scenarios.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your edge network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet

Module 4

Designing Network Management Services

Overview

Network management is a fundamental service in any enterprise network. Network management solutions require special consideration to network design. Appropriate network management will help network engineers achieve efficient performance, scalability, and availability in their network. This module discusses the importance, requirements, and considerations for implementing network management in the overall enterprise design. The module considers network management design models with respect to the functional areas of network management.

Module Objectives

Upon completing this module, you will be able to design network management intelligent network services for performance, scalability, and availability, given specified enterprise network needs.

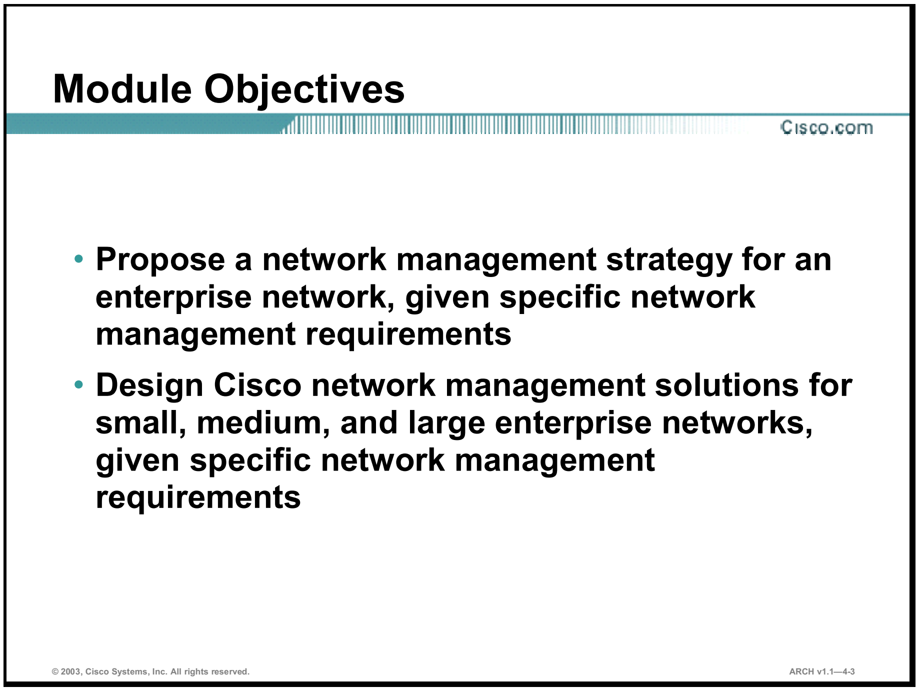

Module Objectives

- Propose a network management strategy for an enterprise network, given specific network management requirements
- Design Cisco network management solutions for small, medium, and large enterprise networks, given specific network management requirements

ARCH v1.1—4-3

Module Outline

The outline lists the components of this module.

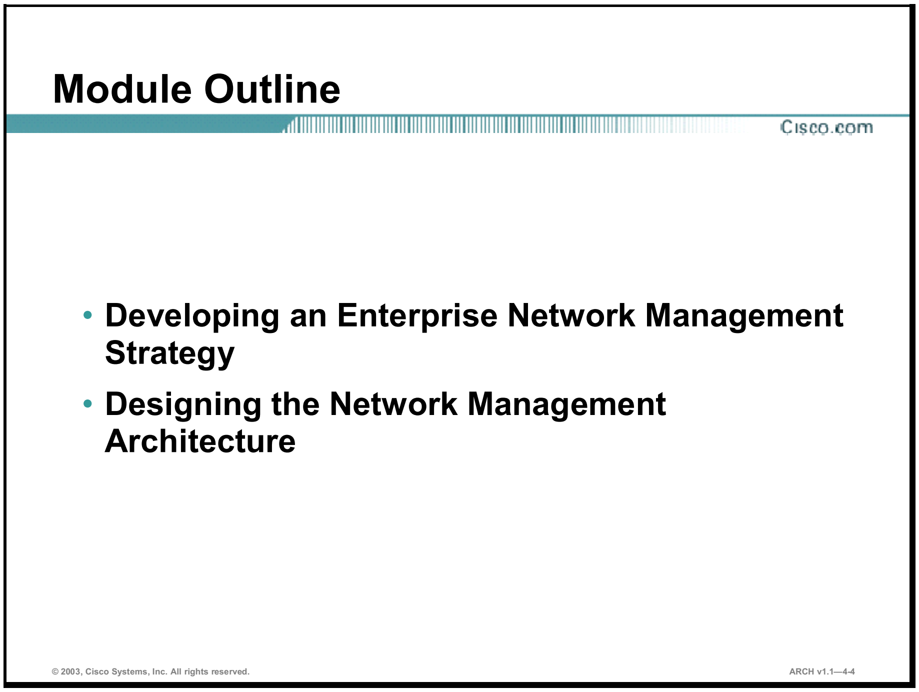

Module Outline

- Developing an Enterprise Network Management Strategy
- Designing the Network Management Architecture

ARCH v1.1—4-4

Developing an Enterprise Network Management Strategy

Overview

Network management includes a broad range of policies, procedures, and purpose-built hardware and software used to manage computer networks. Network management affects the performance, reliability, and security of the entire network.

Relevance

An effective enterprise network management strategy is critical to guarantee performance, reliability, and security for an enterprise network. Ineffective network management will adversely affect performance on the network at some point.

Objectives

Upon completing this lesson, you will be able to propose a network management strategy for an enterprise network, given specific network management requirements. This includes being able to meet these objectives:

- Identify enterprise goals for network management solutions
- List recommendations for developing effective network management policies and procedures
- Identify the infrastructure components that the Network Management module provides
- Describe the Cisco network management strategy, given specific network management requirements
- Describe the CiscoWorks network management solution, given specific network management requirements
- Select CiscoWorks LAN management tools, given specific network management requirements
- Select CiscoWorks WAN management tools, given specific network management requirements

Learner Skills and Knowledge

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

- Basic understanding of network management and the protocols used within the network management arena

Outline

The outline lists the topics included in this lesson.

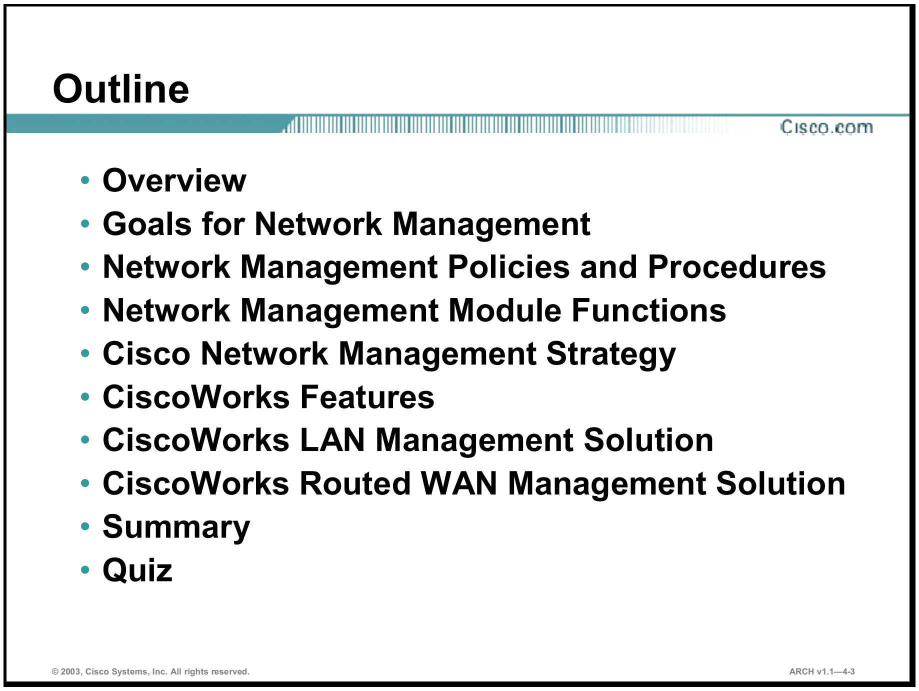

Outline

Cisco.com

- Overview
- Goals for Network Management
- Network Management Policies and Procedures
- Network Management Module Functions
- Cisco Network Management Strategy
- CiscoWorks Features
- CiscoWorks LAN Management Solution
- CiscoWorks Routed WAN Management Solution
- Summary
- Quiz

ARCH v1.1—4-3

Goals for Network Management

Network management is a growing field that is under constant change. Initially, network management included only fault notification. Today, it has grown to incorporate multiple stages including fault, configuration, accounting, performance, and security management. This topic identifies enterprise goals for network management solutions.

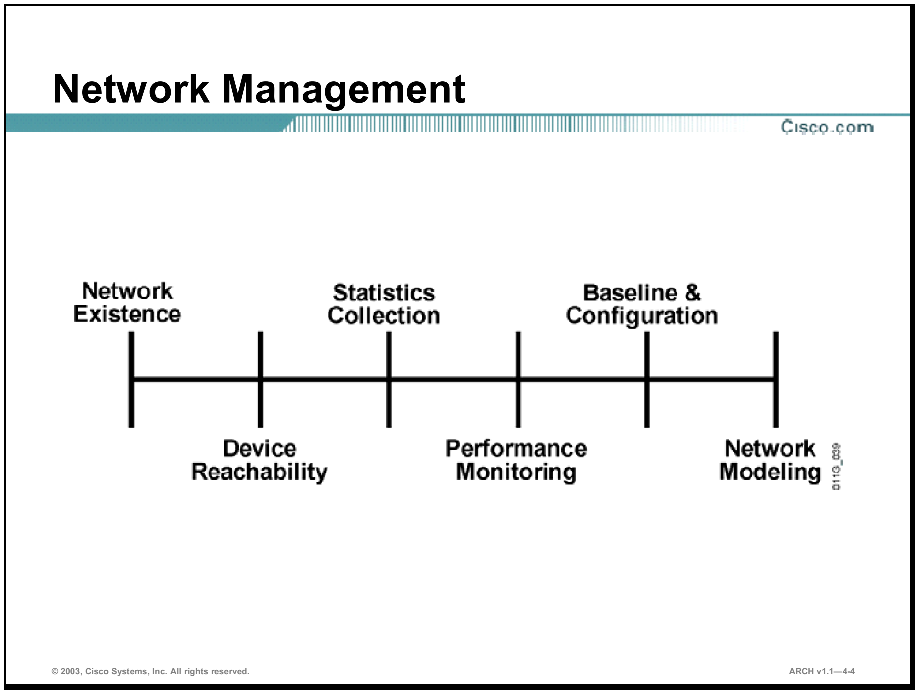

Network Management

Cisco.com

Network Existence

Device Reachability

Statistics Collection

Performance Monitoring

Baseline & Configuration

Network Modeling

ARCH v1.1—4-4

Network management evolves in these ways:

- **Network existence:** There is no systematic way to determine if a fault occurs. When something goes wrong in the network, the people in charge of the network are told via word of mouth that there is a problem.
- **Device reachability:** Basic device reachability is established. This allows the people in charge of the network to see when a device fails or becomes unreachable but does not give any reason for the outage.
- **Statistics collection:** The people managing the network begin gathering statistics on the behavior of a device. When a device goes down, historical data is available that may help solve the problem. This statistics gathering is extremely helpful for resolving recurring problems that are difficult to reproduce. Statistics collection also helps you become more proactive in your network management efforts.
- **Performance monitoring:** The statistics gathering became the input to measuring the performance of the devices being monitored. By monitoring the performance of the devices, the management staff could then tell when a device was about to have problems based on the device performance.
- **Baseline and configuration:** The performance data was used to create a baseline of how the network was performing on the average. Thresholds could then be set to determine when the network was performing outside the expected thresholds. You can use the baseline and configuration data to monitor the network and be proactive in your network management efforts.

- **Network modeling:** Take a snapshot of the network and then model the network so "what if" scenarios can be run to determine when and where a failure of the network could occur. Having this information beforehand allows the management staff to predict when a problem will happen and to proactively fix the potential problem area before it ever becomes a problem. Modeling also provides capacity planning to help determine which upgrades may be necessary to support a proposed architecture, new applications, and growth.

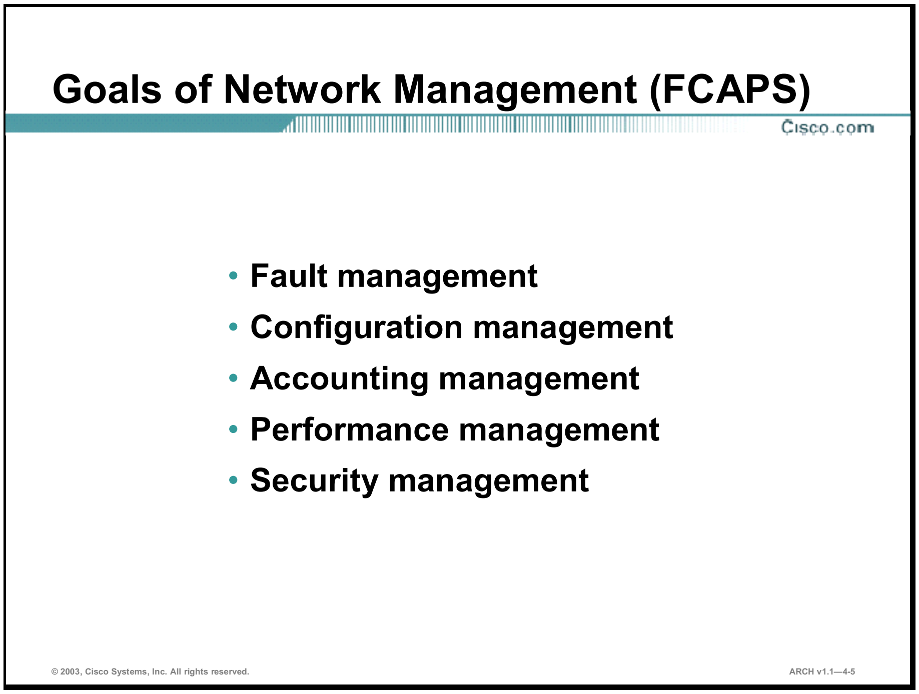

Goals of Network Management (FCAPS)

- Fault management
- Configuration management
- Accounting management
- Performance management
- Security management

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-5

As network management has become an integral part of the network, the International Organization for Standardization (ISO) developed a framework for network management.

The five functional areas defined by ISO for network management are:

- **Fault management:** The ability to detect, isolate, and notify when a fault occurs within the network.
- **Configuration management:** The ability to track and maintain device configurations within the network. This includes configuration file management, device inventory, and software management.
- **Accounting management:** The ability to track the usage of the devices and the network resources. Accounting management tracks the hardware and software inventory, and should provide change control (versioning).
- **Performance management:** The ability to gather performance information from the devices and then interpret that information to determine the performance levels of the links and devices within the network.
- **Security management:** The ability to restrict and log access to the network resources.

Network Management Policies and Procedures

As network management systems are implemented, it is vital to put policies and procedures in place. This topic lists recommendations for developing effective network management policies and procedures.

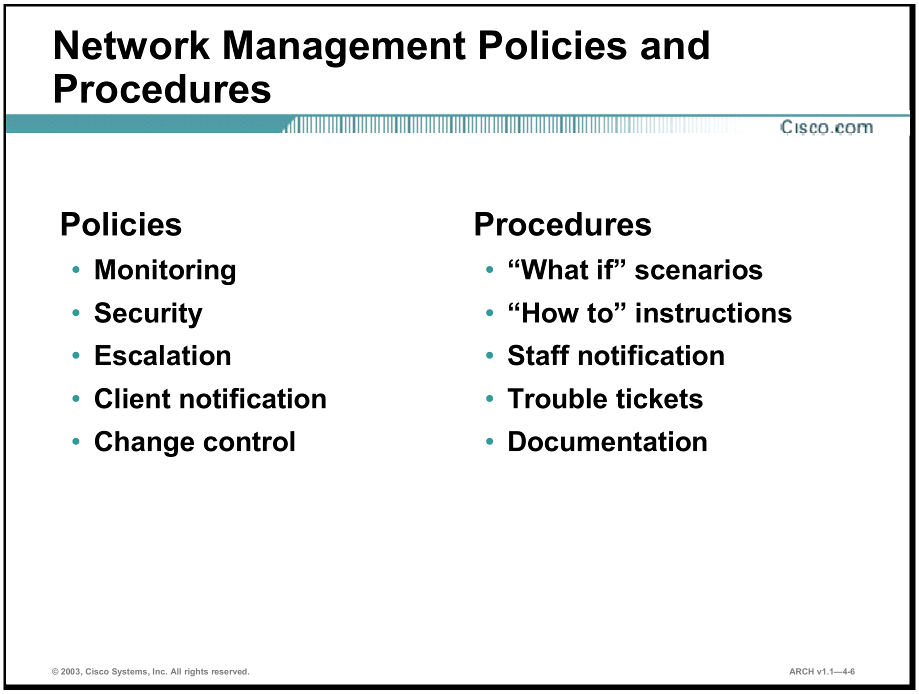

Network Management Policies and Procedures

Cisco.com

Policies

- Monitoring
- Security
- Escalation
- Client notification
- Change control

Procedures

- "What if" scenarios
- "How to" instructions
- Staff notification
- Trouble tickets
- Documentation

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-6

Policies are implemented to define the plan for the network management system while the procedures are in place to define how to react to an event or how to interpret collected performance data.

Policies are written to provide details about whom to manage and about what to manage. They outline which devices are monitored for reachability, which devices are monitored for performance measurements, the plan for escalation, and how and when users are notified of network issues.

Procedures are written to describe the steps to take when an event happens. These procedures should instruct the problem solver about the steps to take to solve the event or escalate the problem for resolution, and how to document the findings.

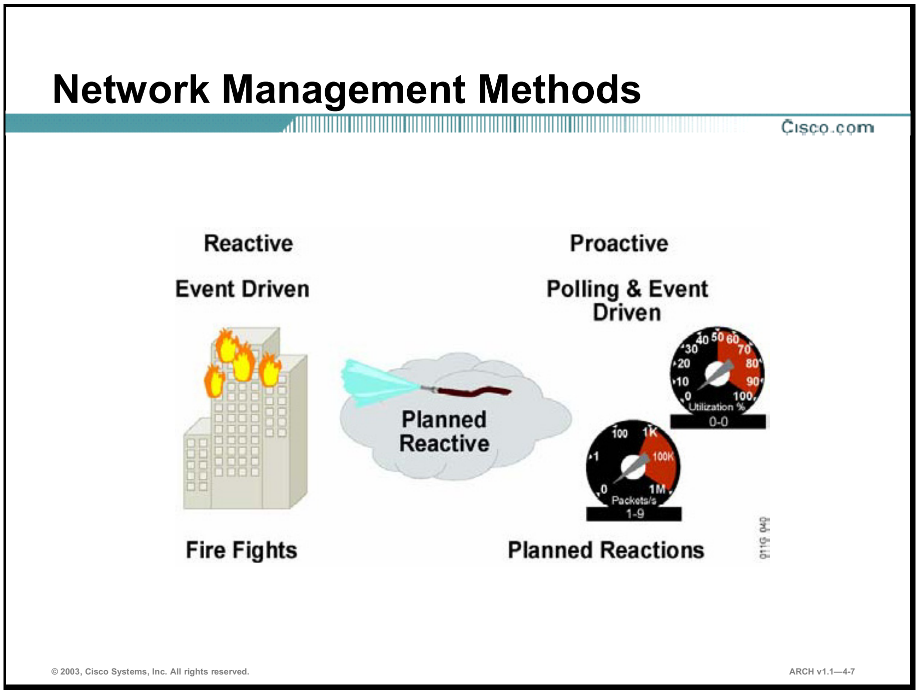

Network Management Methods

Cisco.com

Reactive

Event Driven

Fire Fights

Planned Reactive

Proactive

Polling & Event Driven

Planned Reactions

ARCH v1.1—4-7

Two network management styles are typically used in the network management arena: reactive and proactive.

The reactive management style is very common in the enterprise that has not thought out its network management strategy. Many enterprise IT departments are understaffed, or their network is so poorly planned that the staff is constantly on the defensive, having to react to network issues. Another form of reactive management occurs when management makes a decision to only monitor events sent from the device, or devices adjacent to the device that is having a problem. The problem with such an event-driven solution is that it keeps the staff in a reactive mode. The staff only learns of a problem after the problem has manifested itself as a network issue, when a device sends an event out to the network management station.

The proactive management style is the most preferred style by most management staff. It uses event-driven information and polling to determine the health of the network. The polling and collection of data from devices allow the management team to identify problems before they happen so they can take steps to reduce or remove issues before users see any problems. Though in the proactive style it is still necessary to put out fires, there should be fewer issues and better plans developed to resolve issues.

To develop a network management strategy for the enterprise, complete these tasks:

Step 1 Plan which devices will be managed and which will not.

Step 2 Determine what information to gather or receive from network devices. If the management system is going to be event driven, determine what kind of information to collect from the traps, which are events sent to the management station from a device. If the management style is polling driven, determine what information to collect from the devices to meet the management goals.

Step 3 Set realistic, measurable goals for network management. If the management style is event-driven, set goals that you can measure based on the events received. If the management style is based on polling or a combination of polling and event driven, set measurable goals to determine if the network management system is performing as expected.

Step 4 Identify the tools available to collect the required data.

Step 5 Identify the monitoring goals and thresholds. Then set the appropriate traps and alerts on the specific devices best positioned to report the activity.

Step 6 Create plans and procedures to handle "what if" scenarios, so that when a network problem is identified, there are some basic procedures in place to resolve the problem.

Network Management Module Functions

To implement network management goals, enterprises will implement individual infrastructure components that meet specific needs. This topic identifies the infrastructure components that the Network Management module provides.

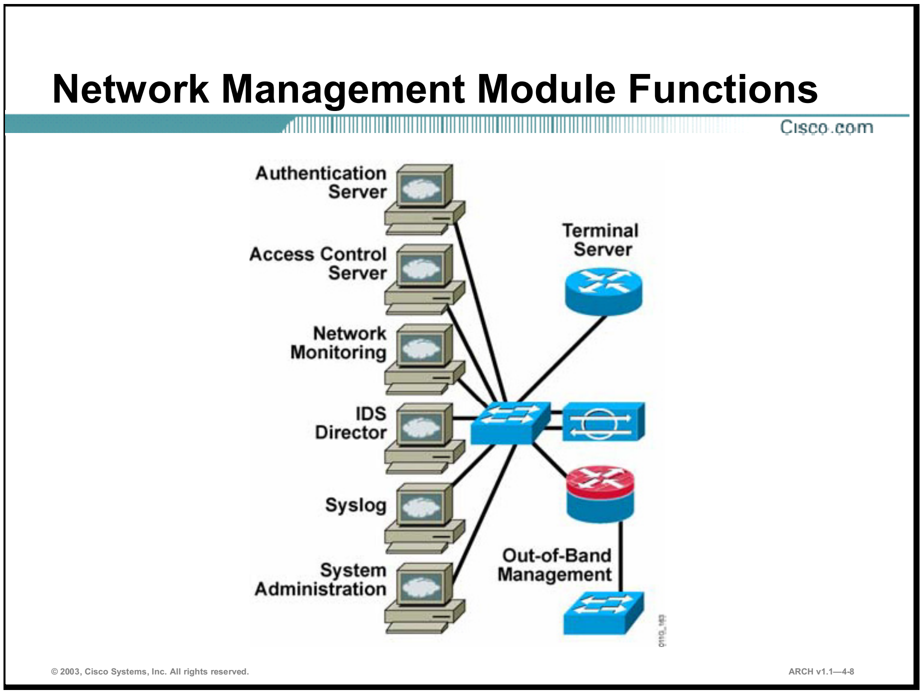

Network Management Module Functions

Cisco.com

ARCH v1.1—4-8

The Enterprise Composite Network Model includes the Network Management module. The Network Management module may contain one or more of these services:

- **Authentication server:** Provides authentication services for remote and local users on the network
- **Access control server:** Provides centralized command and control for all user authentication, authorization, and accounting
- **Network monitoring server:** Responsible for monitoring the devices in the network
- **IDS Director:** Provides a comprehensive, pervasive security solution for combating unauthorized intrusions, malicious Internet worms, and bandwidth and application attacks
- **Syslog:** Provides collection point for network events and traps
- **System administration server:** Management station used to configure network management and other network devices

These tools provide a network administrator with access to the devices in the Network Management module:

- **Out-of-band management:** Provides the ability to access devices through a path external to that taken by production network traffic
- **Terminal server:** Provides a way to perform out-of-band management to multiple devices connected serially to the terminal server, which is in turn connected to a modem

Cisco Network Management Strategy

Enterprise network managers are often faced with the task of developing a strategy for managing very large networks. Networks continue to grow in size, and the number of different management tools and products is often large as well, making the task more difficult. This topic describes the Cisco network management strategy, given specific network management requirements.

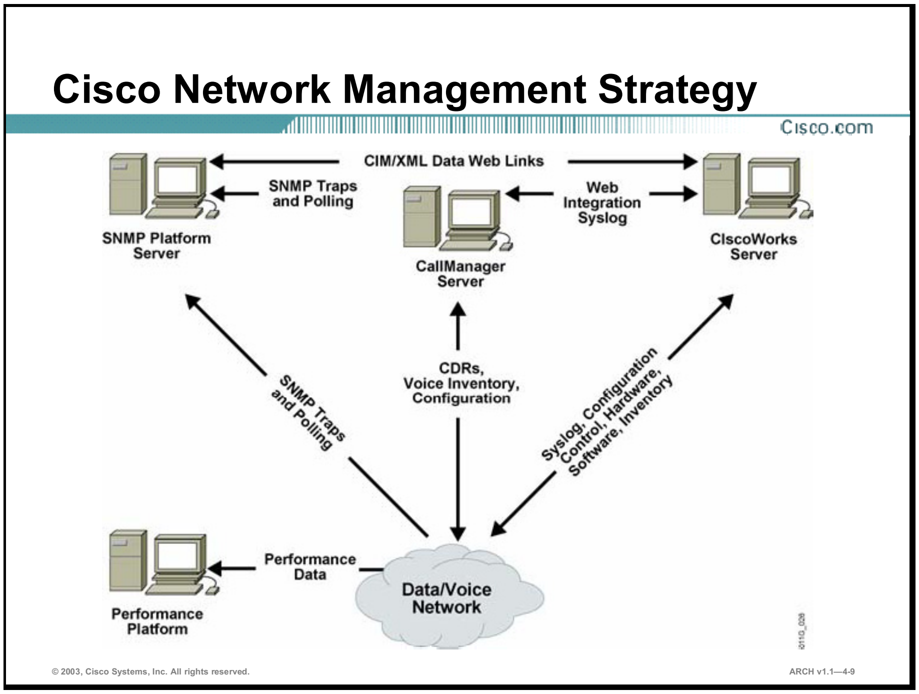

Cisco Network Management Strategy

ARCH v1.1—4-9

The Cisco network management strategy includes a web-based model that offers these characteristics:

- Simplification of tools, tasks, and processes
- Flexible but secure user access via a common web browser
- Web-level integration with network management system (NMS) platforms and general management products
- Capability to provide end-to-end solutions for managing routers, switches, and access servers
- Creation of a management intranet by integrating discovered device knowledge with Cisco.com and third-party application knowledge

CiscoWorks Features

CiscoWorks is a family of products based on Internet standards for managing Cisco enterprise networks and devices. This topic describes the CiscoWorks network management solution, given specific network management requirements.

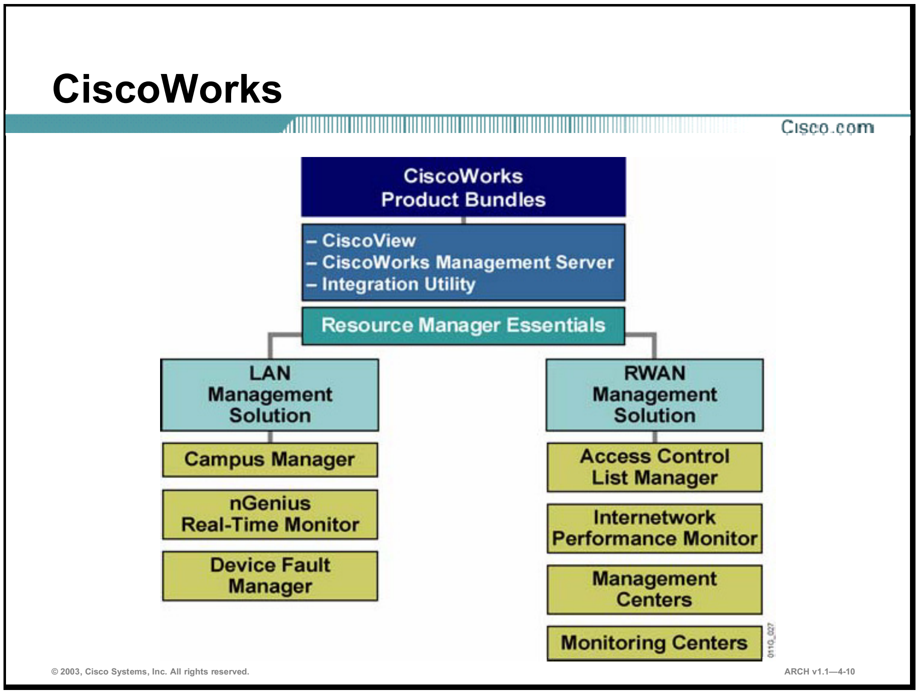

CiscoWorks

Cisco.com

CiscoWorks Product Bundles

- CiscoView
- CiscoWorks Management Server
- Integration Utility

Resource Manager Essentials

LAN Management Solution

- Campus Manager
- nGenius Real-Time Monitor
- Device Fault Manager

RWAN Management Solution

- Access Control List Manager
- Internetwork Performance Monitor
- Management Centers
- Monitoring Centers

0110_027

ARCH v1.1—4-10

The CiscoWorks product line offers a set of solutions designed to manage the enterprise network. The solutions include the LAN Management Solution (LMS) and the Routed WAN Management Solution (RWAN). These solutions focus on key areas in the network such as optimization of the WAN, administering switch-based LANs, and measuring service level agreements within all types of networks. The expanding CiscoWorks product line offers the flexibility to deploy end-to-end network management solutions.

Cisco Architecture for Voice, Video and Integrated Data (AVVID) solutions rely on a stable foundation of optimally functioning Catalyst multiprotocol, multilayer LAN switches. Proper network management is essential to maintain such an environment. However, administrators often find that it is difficult to adequately manage a converged network using manual processes.

CiscoWorks network management can provide the network administrator with a scalable tool, easily learned, and capable of automating most common network management tasks. When properly deployed, the CiscoWorks network management solution can provide considerable savings by reducing labor cost and increasing network availability.

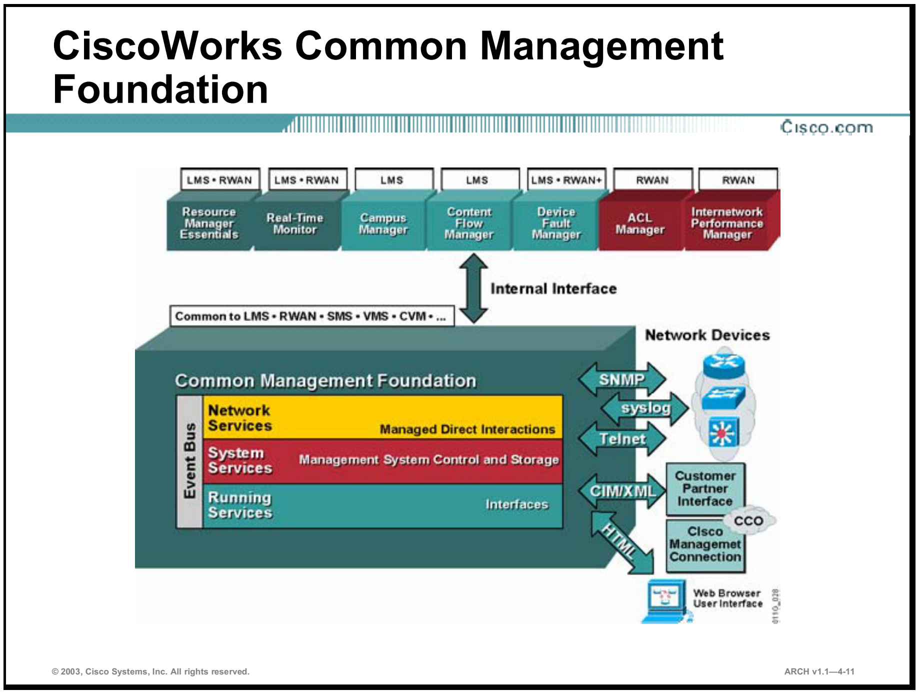

CiscoWorks Common Management Foundation

Cisco.com

LMS • RWAN | LMS • RWAN | LMS | LMS | LMS • RWAN+ | RWAN | RWAN

Resource Manager Essentials | Real-Time Monitor | Campus Manager | Content Flow Manager | Device Fault Manager | ACL Manager | Internetwork Performance Manager

Internal Interface

Common to LMS • RWAN • SMS • VMS • CVM • ...

Common Management Foundation

Event Bus

Network Services — Managed Direct Interactions

System Services — Management System Control and Storage

Running Services — Interfaces

SNMP

syslog

Telnet

CIM/XML

HTML

Network Devices

Customer Partner Interface

CCO

Cisco Managemet Connection

Web Browser User Interface

0110_028

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-11

The CiscoWorks architecture consists of a Common Management Foundation (CMF) with a web-based desktop as a single point of management. An additional component provides data collection services using Simple Network Management Protocol (SNMP), Cisco Discovery Protocol (CDP), and Integrated Local Management Interface (ILMI) tables. SNMP, Syslog, Telnet, Common Information Model/extensible markup language (CIM/XML), and HTML.

Discovery of the network occurs via a seed device, typically a Catalyst switch, through which the CMF discovers the network by reading its CDP or ILMI table and SNMP variables. This information permits the discovery of the seed's neighboring devices, which in turn are queried for CDP/ILMI and SNMP information. This process continues outward from each discovered device to build a network topology map. The CMF also provides granular security, process control, and device information retrieval via SNMP.

CiscoWorks LAN Management Solution

The CiscoWorks LMS provides a solid foundation of basic and advanced management applications that enable network operators to efficiently manage the LAN. This topic helps you select CiscoWorks LAN management tools, given specific network management requirements.

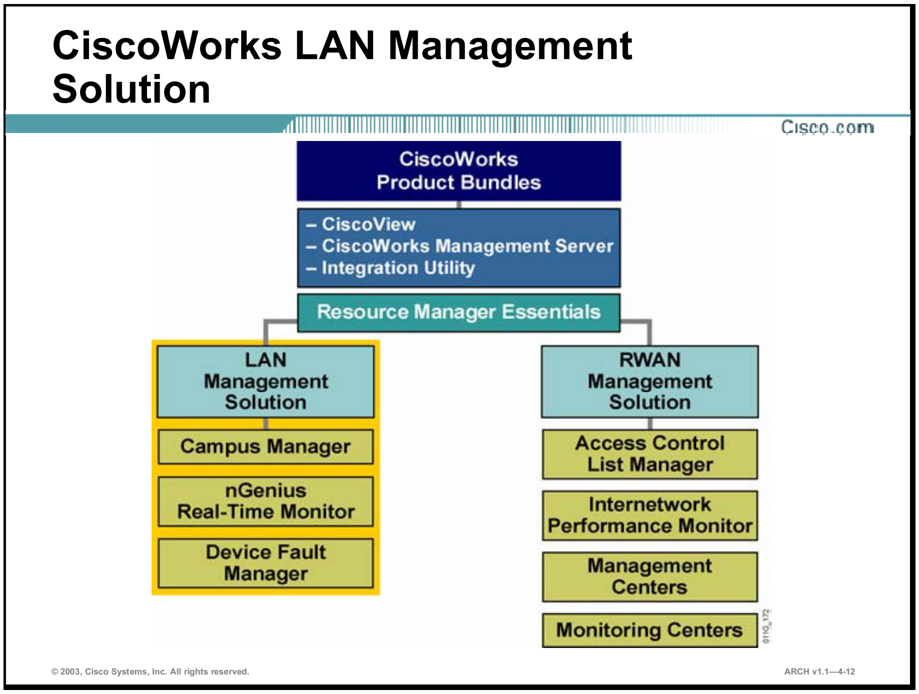

CiscoWorks LAN Management Solution

Cisco.com

CiscoWorks Product Bundles

– CiscoView
– CiscoWorks Management Server
– Integration Utility

Resource Manager Essentials

LAN Management Solution

Campus Manager

nGenius Real-Time Monitor

Device Fault Manager

RWAN Management Solution

Access Control List Manager

Internetwork Performance Monitor

Management Centers

Monitoring Centers

0110_172

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-12

All CiscoWorks product bundles include these common modules:

- **CiscoView:** CiscoView is a web-based device management application providing dynamic status, monitoring, and configuration for a managed Cisco device.
- **CiscoWorks Management Server:** CiscoWorks Management Server includes the database engine, online help, security, login, application launching (from the CiscoWorks desktop), job and process management, and a web server for client access. CiscoWorks Management Server also includes CMF services and Cisco Management Connection to integrate other web-based management applications and tools into the CiscoWorks desktop.
- **Integration Utility:** The Integration Utility offers integration with third-party network management platforms by adding CiscoView device-specific information to the platform and provides launch points to other Cisco applications.
- **Resource Manager Essentials (RME):** RME contains the configuration management tools necessary for Cisco devices. RME contains these six applications to aid the administrator’s efforts:
  - **Inventory Manager:** Inventory Manager provides current inventory of all Cisco devices (routers, switches, firewalls) in the network, including support for Cisco CallManager, VPN concentrators, and WAN switches. Hardware and software summary information, includes detailed reports for groups of devices, memory, flash, software version, interface, and stack modules.

- **Configuration Manager:** Configuration Manager maintains an active archive of configuration changes, and can modify stored configuration changes across multiple Cisco routers and switches. When Configuration Manager detects a configuration change (applied via command-line interface (CLI), Telnet, or via CiscoWorks), it automatically updates the archive data, and logs the change information to the Change Audit Service. Configuration Manager also provides powerful editing capabilities of the archived configuration data, including find, search, replace, copy, cut, paste, as well as compare and change detail. Modified files can be saved locally or downloaded to the target device.

  Cisco-provided templates simplify the configuration change process for SNMP community, Terminal Access Controller Access Control System plus (TACACS+), enable, syslog, SNMP trap destinations, Cisco Discovery Protocol (CDP), and Domain Name System (DNS).
- **Software Image Manager:** Software Image Manager simplifies the version management and routine deployment of software updates to Cisco routers and switches through wizard-assisted planning, scheduling, downloading, and monitoring of software updates. Links to Cisco.com compare the latest Cisco online software update information with the IOS and Catalyst software images deployed in the network. Software Image Manager also allows the administrator to validate the target switch or router's inventory data with the hardware requirements of a new image to help ensure a successful upgrade. When multiple devices are being updated, Software Image Manager synchronizes download tasks and allows the user to monitor job progress. Scheduled jobs are controlled through a sign-off process, enabling managers to authorize a technician's activities before initiating each upgrade task.
- **Change Audit:** Change Audit displays changes made to managed devices. Summary information includes: types of changes made, the person responsible for the change, time of change, and whether the changes were made from a Telnet or console CLI or from a CiscoWorks application. Reports detailing the nature of the changes, such as cards added or removed, memory changes, and configuration changes, are available as well.
- **Availability Manager:** Availability Manager displays the operational status of critical routers and switches. Drilling down on a particular device allows the administrator to view historical information with regard to a given device, including response time, availability, reloads, protocols, and interface status.
- **Syslog Analyzer:** Syslog Analyzer filters syslog messages logged by Cisco switches, routers, access servers, and IOS firewalls, and generates reports in an easily digested format. Its reports are based on user-defined filters that highlight specific errors or severity conditions, and help identify when specific events occurred (such as a link-down condition or a device reboot). Syslog Analyzer also allows syslog messages to be linked to customized information, such as web-based administrative tips or to launch a Common Gateway Interface (CGI) script to take corrective actions.

If you turn on RME features, such as Availability Manager, you need to evaluate the number of network management stations needed to manage the network. Each feature requires significant computing resources.

CiscoWorks LMS is composed of these three modules:

- **Campus Manager:** Campus Manager provides the administrator with tools to configure, manage, and understand the physical and logical aspects of a Catalyst-based local-area network. Campus Manager offers these applications:
  - **Topology Services:** This is the principal interface to large-scale topology maps, tabular summaries, reports, and configuration services of the data link layer network. A directory-like tree interface lists physical data link layer and the logical, Virtual Terminal Protocol (VTP), and ATM domain views, along with table summaries of the devices and the interface associated with these views. This tree structure acts as the launching point for topology maps, discrepancy reporting functions, and configuration services.
  - **User Tracking:** Designed to assist in locating end-station connections at the access switch, this application is used in troubleshooting or connectivity analysis. Through automated acquisition, a table of end-user stations and data link layer connection information is constructed. This table can be sorted and queried, allowing administrators to easily find users. Users can be identified by name, IP handset, or MAC and IP address, as well as the switch port and switch on which they are connected, and the VLAN and VTP assignment of the port can also be identified. Predefined reports, such as duplicate MAC addresses per switch port, or duplicate IP addresses, enable managers to locate mobile users or violations in port policies.
  - **Path Analysis:** An application for switched network management, this is a powerful tool for connectivity troubleshooting. Path Analysis uses topology services, user tracking,and real-time spanning-tree information to determine data link layer connectivity and multilayer connectivity between two endpoints in the network. The resulting trace appears in graphical views that illustrate the devices, path direction, and link types. A tabular format provides specific interface, IP address, VLAN, and link-type information.
  - **VLAN Port Assignment:** Campus Manager provides a graphical interface to create, modify, or delete VLANs and LAN emulation (LANE) elements, plus assign switch ports to VLANs. As you create or modify VLANs, port and user changes are updated and transmitted to the switches, eliminating the need to update and configure each participating switch individually. As you select VLANs, the table view shows the participating ports, port status, and switch information, and you can launch the topology map to highlight participating devices and links for VLAN connections.
  - **Discrepancy Reports:** Discrepancy reports are used to view the physical and logical discrepancies discovered on the network.
- **nGenius Real-Time Monitor:** nGenius Real-Time Monitor is a web-based tool that delivers multi-user access to real-time Remote Monitoring (RMON) and RMON2 information, and is used for monitoring, troubleshooting, and maintaining network availability. nGenius Real-Time Monitor can graphically analyze and report device-, link-, and port-level RMON-collected traffic data obtained from RMON-enabled Catalyst switches, internal Network Analysis Modules, and external LAN and WAN probes.
- **Device Fault Manager:** Device Fault Manager provides real-time fault analysis for managed Cisco devices. Device Fault Manager actively monitors a wide range of problems that Cisco has identified can occur within Cisco devices. Depending on the type of device, Device Fault Manager will actively monitor different conditions via Internet Control Message Protocol (ICMP) polling, SNMP Management Information Base (MIB) interrogation, and SNMP trap reception, and track only those conditions known to help cause higher-level problems in that particular device.

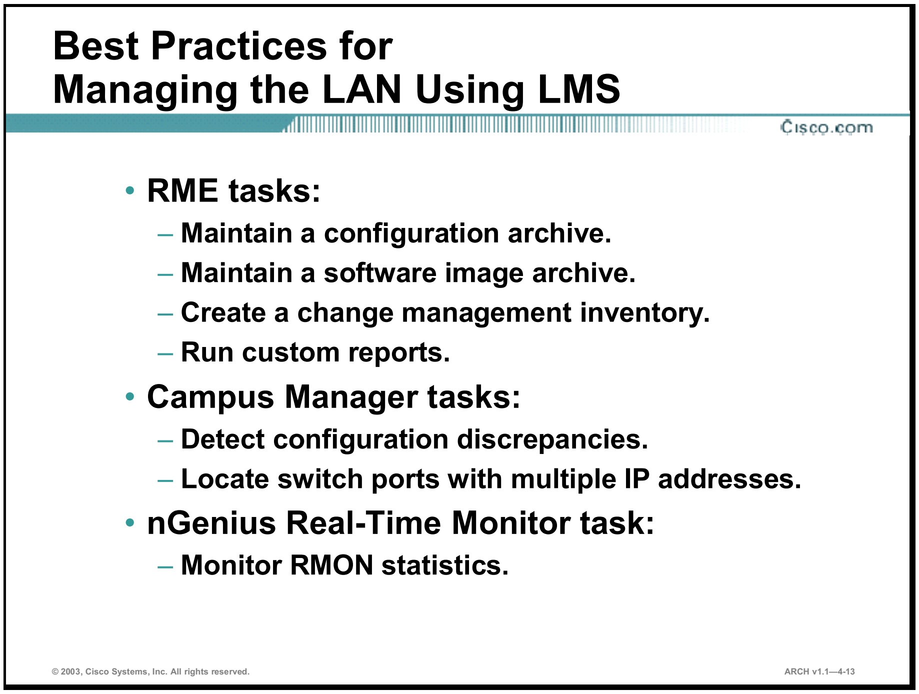

Best Practices for Managing the LAN Using LMS

Cisco.com

- RME tasks:
  - Maintain a configuration archive.
  - Maintain a software image archive.
  - Create a change management inventory.
  - Run custom reports.
- Campus Manager tasks:
  - Detect configuration discrepancies.
  - Locate switch ports with multiple IP addresses.
- nGenius Real-Time Monitor task:
  - Monitor RMON statistics.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—4-13

Network management tasks will differ depending on the needs and capabilities of the network and available support resources. However, you should plan to conduct these tasks with RME regardless of the network's individual characteristics:

- **Maintain a configuration archive:** Automatically conducted by RME on all devices in the RME inventory.
- **Maintain a software image archive:** RME can import the software images on all devices, and then automatically poll devices to determine if their image is backed up in the archive for disaster recovery purposes.
- **Create a change management inventory:** Automatically conducted on all inventoried devices.
- **Run custom reports:** RME Syslog Analyzer can run custom reports based on a device level, syslog message priority, or timeframe.

Use Campus Manager to perform these tasks:

- **Detect configuration discrepancies:** Use Campus Manager to automatically validate connectivity for proper duplex, speed, and trunking configuration.
- **Locate switch ports with multiple IP addresses:** Campus Manager's user tracking function will locate switch ports with multiple IP addresses.

Use nGenius Real-Time Monitor to perform this task:

- **Monitor RMON statistics:** Monitor RMON statistics using nGenius Real-Time Monitor to detect application errors and link usage.

CiscoWorks Routed WAN Management Solution

The CiscoWorks Routed WAN (RWAN) Management Solution extends the CiscoWorks product family by providing a collection of powerful management applications to configure, administer, and maintain a Cisco RWAN. This topic helps you select CiscoWorks WAN management tools, given specific network management requirements.

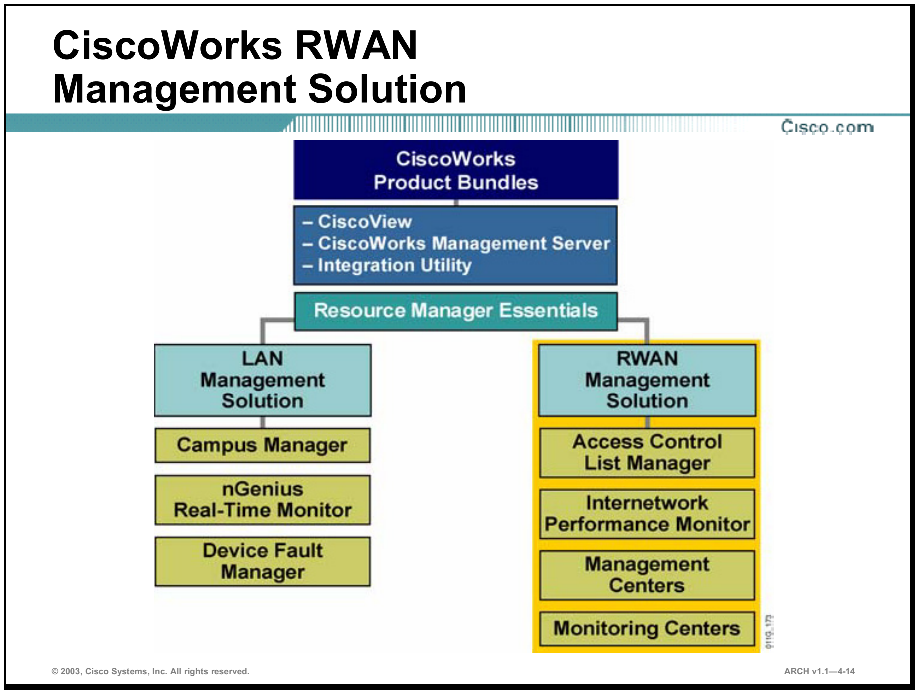

CiscoWorks RWAN Management Solution

Cisco.com

CiscoWorks Product Bundles

- CiscoView
- CiscoWorks Management Server
- Integration Utility

Resource Manager Essentials

LAN Management Solution

Campus Manager

nGenius Real-Time Monitor

Device Fault Manager

RWAN Management Solution

Access Control List Manager

Internetwork Performance Monitor

Management Centers

Monitoring Centers

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-14

The RWAN solution addresses the management needs of WANs by improving the accuracy, efficiency, and effectiveness of your network administrators and operations staff, while increasing the overall availability of your network through proactive planning, deployment, and troubleshooting tools.

Many management tasks that are essential in the LAN environment are also equally important in the WAN environment. Therefore, there is some overlap in the functionality between LMS and RWAN.

The CMF is common in all CiscoWorks solutions. RWAN shares the Resource Manager Essentials with LMS.

Additional RWAN modules include:

- **Internetwork Performance Monitor (IPM):** IPM measures network performance based on synthetic traffic generation technology generated by the Service Assurance Agent (SAA) contained in IOS. IPM gives the network manager the ability to obtain baseline performance data, useful in troubleshooting situations and to validate the network infrastructure for new multiservice applications. IPM can generate a network response time baseline for any of these network traffic types:
  - Internet Control Message Protocol (ICMP) echo
  - IP path echo
  - Systems Network Architecture (SNA) echo
  - User Datagram Protocol (UDP)
  - UDP jitter
  - Voice over IP
  - TCP connect
  - DNS
  - Dynamic Host Configuration Protocol (DHCP)
  - HTTP
  - Data-link switching (DLSw)
- **Access Control List Manager (ACL Manager):** ACL Manager manages the access lists of Cisco devices. The ACL Manager provides tools to set up and manage IP and IPX filtering and device access control. These tools include: access list editors, policy template managers, network and service class managers for scalability, access list navigation tools for troubleshooting, optimization of access lists, and automated distribution of access list updates.

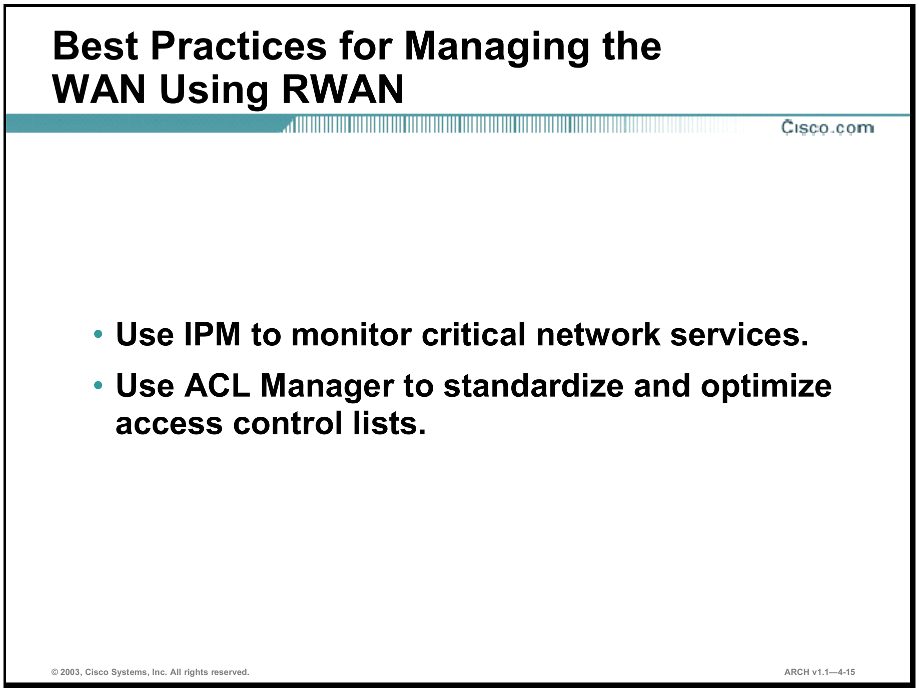

Best Practices for Managing the WAN Using RWAN

Cisco.com

- Use IPM to monitor critical network services.
- Use ACL Manager to standardize and optimize access control lists.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-15

You should plan to conduct these tasks, regardless of the network's individual characteristics:

- **Maintain a configuration archive:** Automatically conducted by RME on all devices in the RME inventory.
- **Maintain a software image archive:** RME can import the software images on all devices, and then automatically poll devices to determine if their image is backed up in the archive for disaster recovery purposes.
- **Change management:** Automatically conducted on all inventoried devices.
- **Syslog reporting and monitoring:** RME Syslog Analyzer can run custom reports based on a device level, syslog message priority, or timeframe.

You should perform these WAN-specific management tasks:

- **Use IPM to monitor critical network services:** Monitor critical network services such as DNS and DHCP, as well as response times to critical servers with IPM.
- **Use ACL Manager to standardize and optimize access control lists:** ACL Manager presents a user-friendly graphical user interface that allows you to concentrate on the security of your network without having to learn complex IOS syntax.

Summary

This topic summarizes the key points discussed in this lesson.

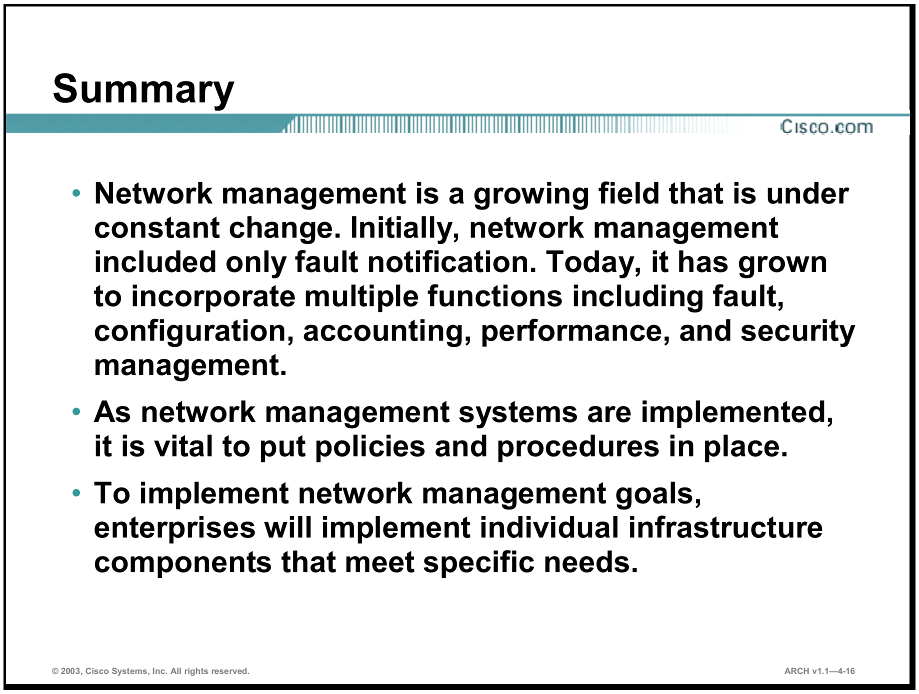

Summary

- Network management is a growing field that is under constant change. Initially, network management included only fault notification. Today, it has grown to incorporate multiple functions including fault, configuration, accounting, performance, and security management.
- As network management systems are implemented, it is vital to put policies and procedures in place.
- To implement network management goals, enterprises will implement individual infrastructure components that meet specific needs.

ARCH v1.1—4-16

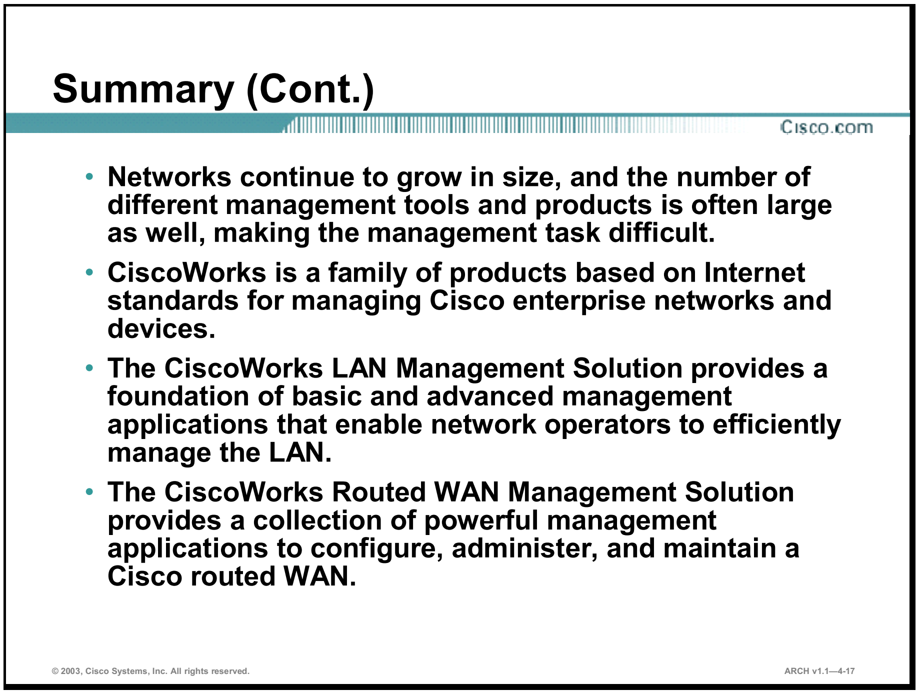

Summary (Cont.)

- Networks continue to grow in size, and the number of different management tools and products is often large as well, making the management task difficult.
- CiscoWorks is a family of products based on Internet standards for managing Cisco enterprise networks and devices.
- The CiscoWorks LAN Management Solution provides a foundation of basic and advanced management applications that enable network operators to efficiently manage the LAN.
- The CiscoWorks Routed WAN Management Solution provides a collection of powerful management applications to configure, administer, and maintain a Cisco routed WAN.

ARCH v1.1—4-17

References

For additional information, refer to these resources:

- *Network Management System: Best Practices White Paper* at http://www.cisco.com/warp/public/126/NMS_bestpractice.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

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 component of the FCAPS model enables you to detect, isolate, and notify when a fault occurs within the network?

A) fault management

B) security management

C) performance management

D) configuration management

Q2) You are assigned the task of moving your company's network management strategy from a reactive to a more proactive style of management. Which three tasks should you complete to reach your goal? (Choose three.)

A) rely on event-driven problem detection

B) identify monitoring goals and thresholds

C) follow standard guidelines for setting event thresholds

D) create plans and procedures to handle and resolve problems

E) implement a staff paging system to handle problems when they occur

F) determine what information to gather or receive from network devices

Q3) What does a network management policy do?

A) defines reachability statistics

B) defines bandwidth requirements

C) defines how to react to a network event

D) defines the plan for the network management system

Q4) After you identify the business and management requirements for an enterprise network, what is the first decision point when designing a network management system?

A) identify what to manage

B) identify measurable goals

C) select a management style

D) identify tools to manage the devices

Q5) What are two network management styles?

A) event and fault

B) reactive and proactive

C) fault and performance

D) polling and performance

Q6) Which component provides for managing devices using dedicated communication paths?

A) VPN management

B) WAN management

C) out-of-band management

D) infrastructure management

Q7) Which four features describe the Cisco network management strategy? (Choose four.)

A) simplification of tools, tasks, and processes

B) complex tools to manage the overall network

C) single-point to multipoint solutions for workstations

D) infinite reporting capabilities on the status of individual WAN links

E) complex functionality to work with proprietary management systems

F) web-level integration with NMS platforms and general management products

G) creation of a management intranet by integrating discovered device knowledge with Cisco.com and third-party application knowledge

Q8) On which three key areas do the CiscoWorks product bundles focus for enterprise network management? (Choose three.)

A) design of hierarchical networks

B) management of the routed WAN

C) software distribution and monitoring

D) administration of switch-based LANs

E) measurement of service level agreements

Q9) You just installed the CiscoWorks LAN Management Solution. Which four tasks can you perform with Resource Manager Essentials? (Choose four.)

A) monitor RMON statistics

B) manage access control lists

C) maintain a configuration archive

D) detect and correct discrepancies

E) maintain a software image archive

F) create a change management inventory

G) create and run custom reports tailored to your needs

Q10) When managing a LAN using the LAN Management Solution, which two tasks should the Campus Manager perform? (Choose two.)

A) run custom reports

B) maintain a configuration archive

C) detect configuration mismatches

D) create a change management inventory

E) locate switch ports with multiple IP addresses

Q11) Which CiscoWorks management tool could you use to obtain baseline performance data?

A) nGenius

B) Remote Monitor

C) LAN Management Station

D) Internetwork Performance Monitor

Q12) Given a need to perform active monitoring of the network devices, which RWAN solution would you select?

A) Availability Manager

B) Software Image Manager

C) Access Control List Manager

D) Internetwork Performance Monitor

Quiz Answer Key

Q1) A

Relates to: Goals for Network Management

Q2) B, D, F

Relates to: Goals for Network Management

Q3) D

Relates to: Network Management Policies and Procedures

Q4) A

Relates to: Network Management Policies and Procedures

Q5) B

Relates to: Network Management Policies and Procedures

Q6) C

Relates to: Network Management Module Functions

Q7) A, F, G

Relates to: Cisco Network Management Strategy

Q8) B, D, E

Relates to: CiscoWorks Features

Q9) C, E, F, G

Relates to: CiscoWorks LAN Management Solution

Q10) C, E

Relates to: CiscoWorks LAN Management Solution

Q11) D

Relates to: CiscoWorks Routed WAN Management Solution

Q12) A

Relates to: CiscoWorks Routed WAN Management Solution

Designing the Network Management Architecture

Overview

Enterprise network managers are often faced with the problem of managing very large networks. Networks continue to grow in size, and the number of different management tools and products is often large as well. To design the network management architecture, you will consider the infrastructure, data collection and management strategy, and server sizing.

Relevance

The network management architecture affects the network management system itself as well as the network as a whole. Selecting the right components and defining strategies are critical to the success and availability of the enterprise network.

Objectives

Upon completing this lesson, you will be able to design Cisco network management solutions for small, medium, and large enterprise networks, given specific network management requirements. This includes being able to meet these objectives:

- List design considerations for a network management system
- Provide strategies for designing network management deployments, given specific enterprise management needs
- Design Cisco network management solutions for small, medium, and large enterprise networks, given specific network management requirements

Learner Skills and Knowledge

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

- An understanding of the network infrastructure
- Knowledge of the requirements a management system puts on a network

Outline

The outline lists the topics included in this lesson.

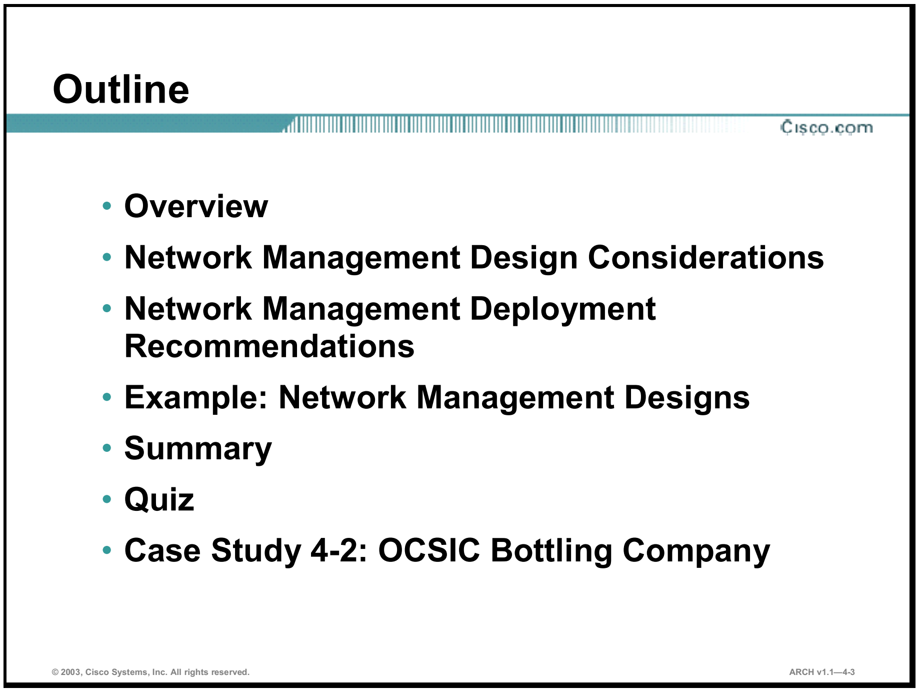

Outline

- Overview
- Network Management Design Considerations
- Network Management Deployment Recommendations
- Example: Network Management Designs
- Summary
- Quiz
- Case Study 4-2: OCSIC Bottling Company

ARCH v1.1—4-3

Network Management Design Considerations

Network management must be an integral part of the overall network design, since network management can adversely affect network performance. This topic identifies design considerations for a network management system.

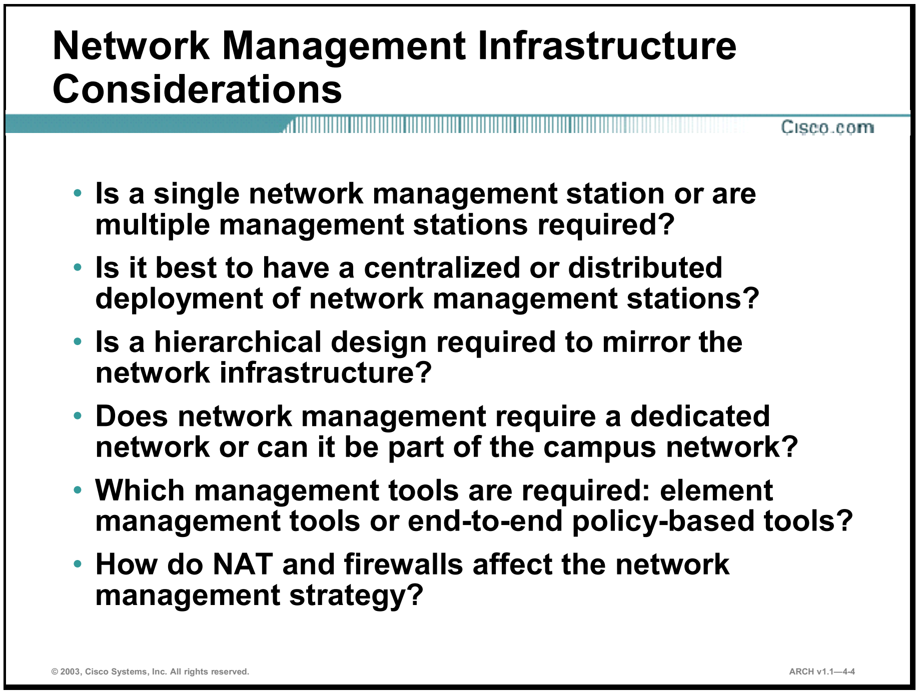

Network Management Infrastructure Considerations

Cisco.com

- Is a single network management station or are multiple management stations required?
- Is it best to have a centralized or distributed deployment of network management stations?
- Is a hierarchical design required to mirror the network infrastructure?
- Does network management require a dedicated network or can it be part of the campus network?
- Which management tools are required: element management tools or end-to-end policy-based tools?
- How do NAT and firewalls affect the network management strategy?

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-4

To develop a network management strategy for the network management infrastructure, you will answer the following questions:

- **Is a single network management station required or multiple management stations?** The number of network management stations required depends on the number of end-user devices to be managed.
- **Is it best to have a centralized or distributed deployment of network management stations?** The decision often depends on the organization’s ability to support the solution.
- **Is a hierarchical design required to mirror the network infrastructure?** Many enterprises implement a hierarchical network management system design that mirrors the network infrastructure from the campus, to the enterprise edge, to the branch offices. The decision depends on the size of the organization, and the organization’s ability to support the solution.
- **Does network management require a dedicated network or can it be part of the campus network?** In other words, will HP OpenView, or CiscoWorks, or some combination take care of all my needs?
- **Which management tools are required: element management tools or end-to-end policy-based tools?** The size of the enterprise usually dictates the types of tools required.
- **How do Network Address Translation (NAT) and firewalls affect the network management strategy?** Enterprises often use NAT and firewalls that block or break access via common management protocols. Solutions are to disable SNMP on security devices, use more advanced authentication options, restrict SNMP access from specific devices, and use more secure protocols.

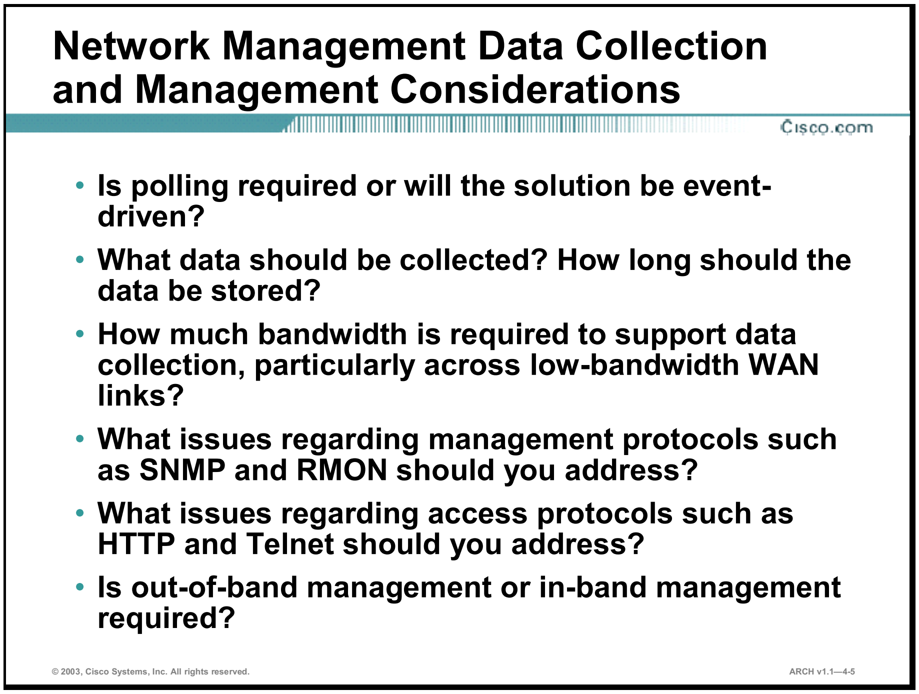

Network Management Data Collection and Management Considerations

- Is polling required or will the solution be event-driven?
- What data should be collected? How long should the data be stored?
- How much bandwidth is required to support data collection, particularly across low-bandwidth WAN links?
- What issues regarding management protocols such as SNMP and RMON should you address?
- What issues regarding access protocols such as HTTP and Telnet should you address?
- Is out-of-band management or in-band management required?

ARCH v1.1—4-5

To develop a network management strategy for data collection and management, you will answer these questions:

- **Is polling required or will the solution be event-driven?** To maintain a proactive management strategy, you will generally choose to implement polling. However, polling can place a large burden on the network infrastructure. Some enterprises implement a separate infrastructure to support network management and polling.
- **What data should be collected? How long should the data be stored?** Consider which data is critical to the management effort and whether the data will actually be used before storing it. Stored data should be purged at regular intervals.
- **How much bandwidth is required for polling, particularly across low-bandwidth WAN links?** To reduce bandwidth requirements, you can reduce polling intervals, implement RMON, and distribute network management systems so they are close to the devices they manage.
- **What issues regarding management protocols such as SNMP and RMON should you address?** SNMP and RMON pose inherent security risks. When managing devices that connect to less secure areas of the network, security must be considered.
- **What issues regarding access protocols such as HTTP and Telnet should you address?** Web-based management tools that use HTTP or Telnet access are not usually options due to security concerns.
- **Is out-of-band management or in-band management required?** Depending on the network management solution deployed and the location of support personnel, out-of-band and/or in-band management will be required.

A challenge is to ensure the availability of the devices being managed. The network management station will poll the essential data to determine the status of the device and to meet the goals for managing each type of device.

To provide availability, you need to build a reliable, redundant network and network management system. Then, if one link goes down, the whole network management system does not go down. You should have a plan for disaster recovery if the situation dictates that the system be manageable from another location in case the primary location goes down.

Another challenge is to implement a proactive system versus a totally reactive system. To build the proactive system, you need to have the staff, procedures, and policies in place and train users to use the network management tools.

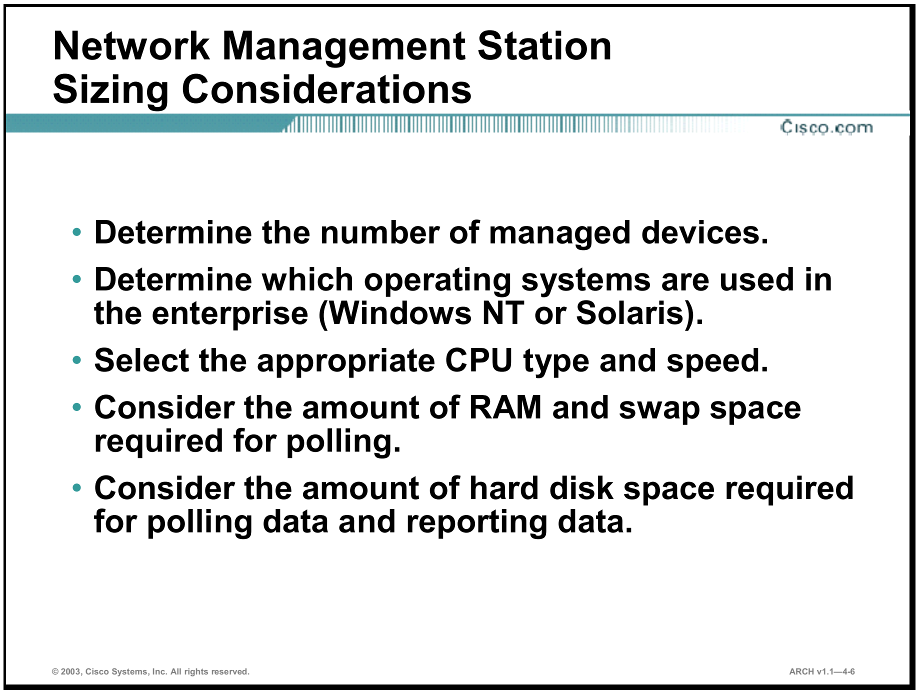

Network Management Station Sizing Considerations

- Determine the number of managed devices.
- Determine which operating systems are used in the enterprise (Windows NT or Solaris).
- Select the appropriate CPU type and speed.
- Consider the amount of RAM and swap space required for polling.
- Consider the amount of hard disk space required for polling data and reporting data.

ARCH v1.1—4-6

Correctly sizing network management stations is important as a server can easily become overloaded, adversely affecting performance. Each network management vendor publishes guidelines to help you select the right size platform for your needs. To size a network management station, consider the items listed in the figure.

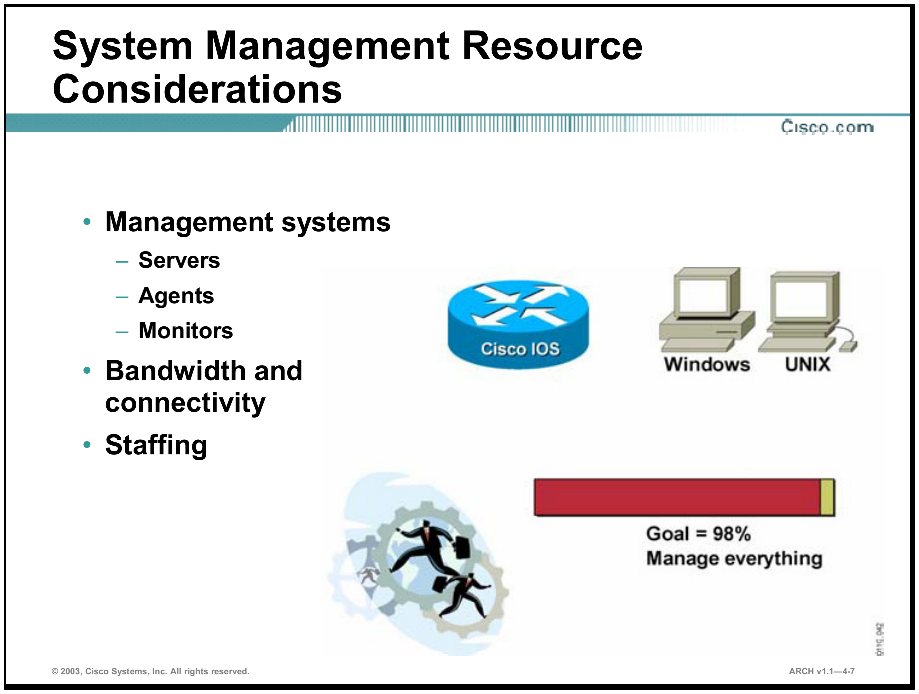

System Management Resource Considerations

- Management systems
  - Servers
  - Agents
  - Monitors
- Bandwidth and connectivity
- Staffing

Cisco IOS

Windows UNIX

Goal = 98%
Manage everything

ARCH v1.1—4-7

Resources not considered may adversely impact the overall effectiveness of the network management system.

To manage the network, you need to consider management system servers and agents as part of the network management plan. In today’s management environment, no single platform can do everything. Most platforms are based on the UNIX or Microsoft Windows operating systems. The devices being managed may utilize their own internal agents.

Depending on the data collection requirements, the network management system may use a substantial amount of the bandwidth over WAN links. Monitoring remote devices across a WAN requires that you plan for the bandwidth and connectivity requirements to the remote sites. The bandwidth requirements for management and user traffic may exceed the total amount of bandwidth available.

If a full-scale enterprise management system is implemented, an appropriate staff is required to monitor and maintain the systems. Staff needs training on the procedures, goals, and the use of the tools to properly manage the network.

Network Management Deployment Recommendations

For networks that require more than a single workstation or server, you may need to use multiple workstations for a single management domain (a single managed network) by distributing applications across multiple workstations. The result will be better performance and maximum scaling. This topic provides strategies to help you design network management deployments, given specific enterprise management needs.

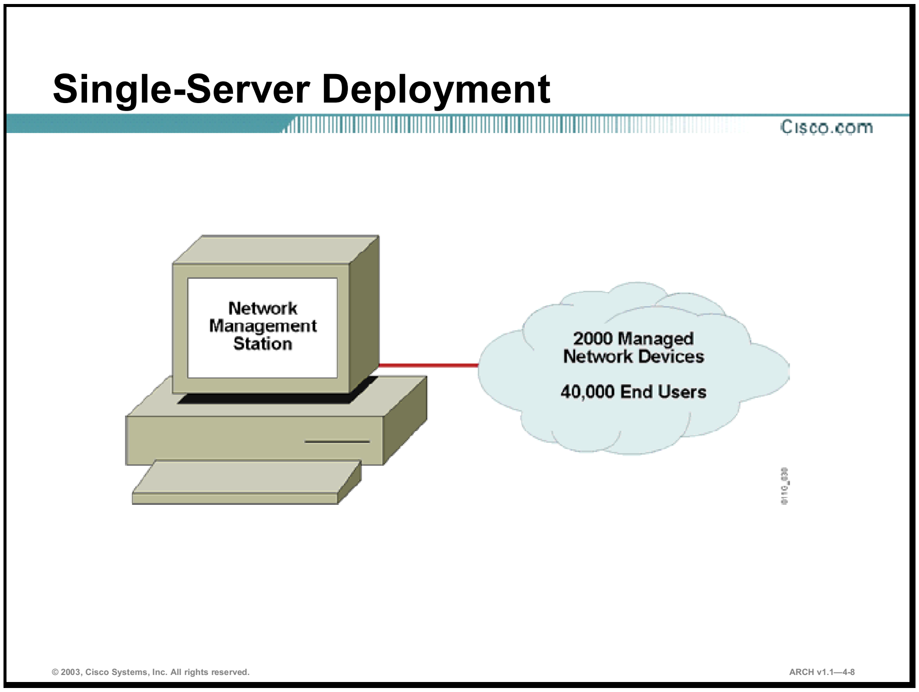

Single-Server Deployment

ARCH v1.1—4-8

For a single server or workstation as a network management platform, a LAN Management Solution (LMS) is recommended for networks with up to 2000 managed network devices or 40,000 user end stations. The addition of other applications depends on specific application scaling factors and the size of your network. Performance can be an issue, so it is important to monitor system resources, and use a multiple processor if appropriate, or additional memory and disk space as needed. Using a single server may not be practical if all bundled applications are on one server. For performance reasons, it may still be necessary to use more than one machine and distribute individual applications across several machines.

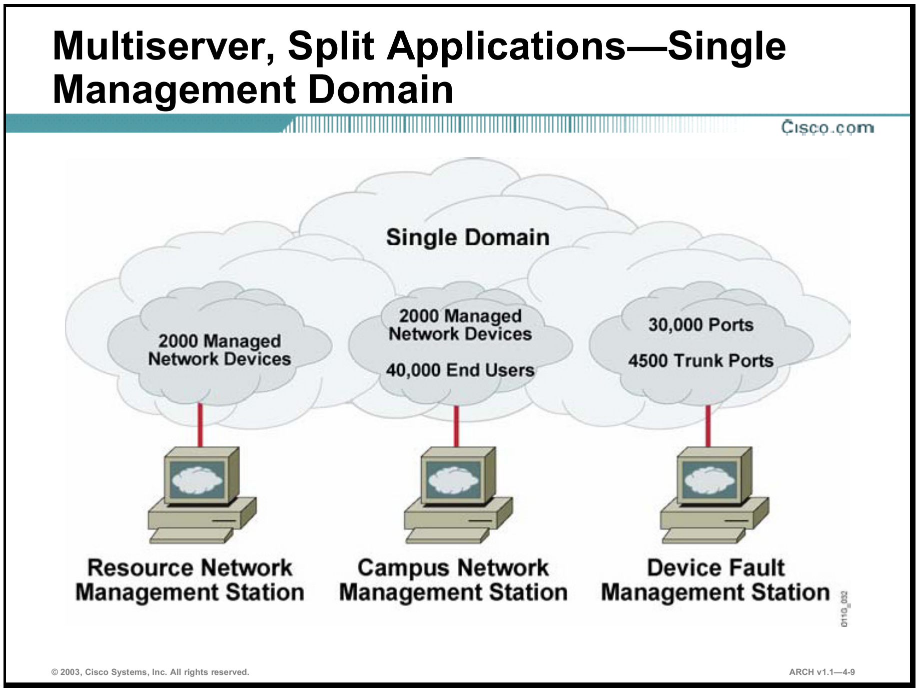

Multiserver, Split Applications—Single Management Domain

Single Domain
2000 Managed Network Devices
2000 Managed Network Devices
40,000 End Users
30,000 Ports
4500 Trunk Ports
Resource Network Management Station
Campus Network Management Station
Device Fault Management Station
0110_032

ARCH v1.1—4-9

If a single workstation cannot handle the load, due to lack of capacity, when multiple applications are required, one solution is to distribute the applications across several servers. Applications should be distributed based on the biggest resource users.

In the example shown in the figure, three management workstations are used to manage a network of 2000 network devices, with the heaviest applications from the LMS distributed across the three workstations.

Other bundled applications can be installed where it makes sense.

Larger networks will have to be split into multiple management domains, or multiple groups managed by individual management servers or groups of servers. When a network is split into multiple domains, you can make the division by administrative groups, geographically, or however fits your needs.

Multiple Management Domains

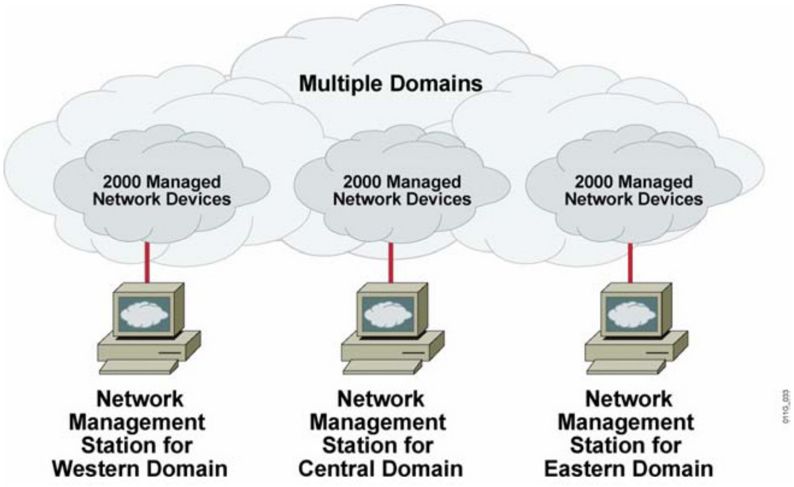

Multiple Domains

2000 Managed Network Devices

2000 Managed Network Devices

2000 Managed Network Devices

Network Management Station for Western Domain

Network Management Station for Central Domain

Network Management Station for Eastern Domain

ARCH v1.1—4-10

When the size of the network is larger than 2000 network devices, you should divide the network into multiple management domains, and multiple management servers (or groups of servers) to ensure that you do not exceed the resource requirements for each server. In some cases, it may be preferable to implement multiple management domains for administrative reasons, even if the numbers do not require the division.

Look for logical ways to segment the network based on the following:

- Virtual Terminal Protocol (VTP) domains
- IP address ranges
- LAN/WAN boundaries

Look for administrative logic with separate management teams, regions, or administrative groupings. You need to determine which management workstation is managing which device (and vice versa), and remember to leave room for future growth.

Consider the scenario shown in the figure, in which a network of 6000 network devices is broken into three groups with up to 2000 devices each (with no more than 40,000 end stations), with a separate LAN server for each segment of the network.

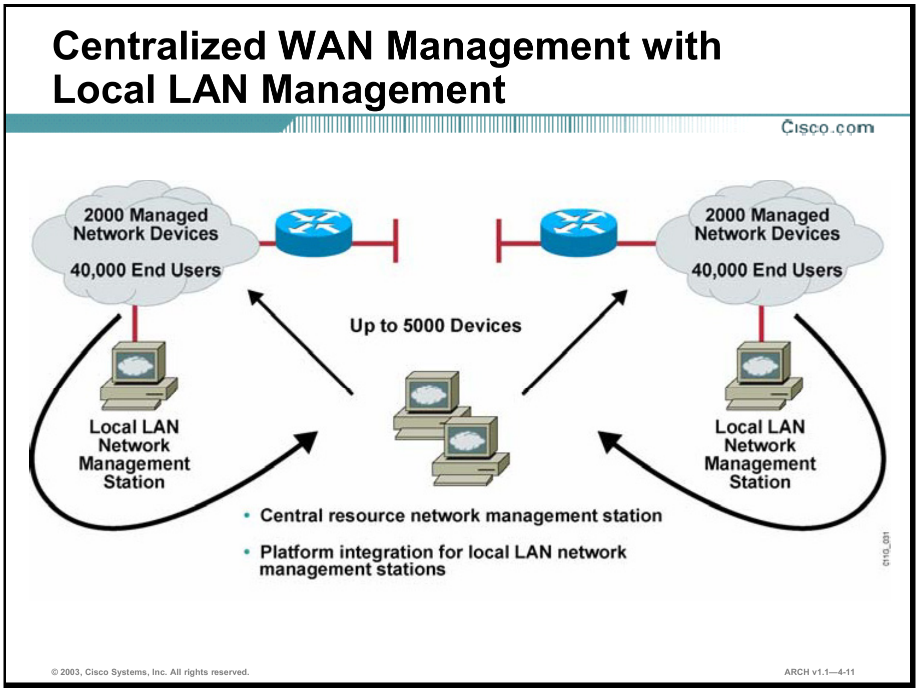

Centralized WAN Management with Local LAN Management

2000 Managed Network Devices
40,000 End Users
2000 Managed Network Devices
40,000 End Users
Up to 5000 Devices
Local LAN Network Management Station
Local LAN Network Management Station
011G_031

- Central resource network management station
- Platform integration for local LAN network management stations

ARCH v1.1—4-11

One option for networks of up to 5000 user devices is to install a single central resource management server, and combine that with multiple campus management servers.

The figure shows a centralized resource management design that provides a single reporting server for inventory, configurations, changes, software distribution, and bulk changes for up to 5000 network devices. For larger networks, one centralized resource server per network partition might be another option. The local LAN management server would have campus and other management applications, but resource management would not be installed on these machines. The resource management server would receive inventory data from each local LAN management machine, and push changes in device credentials back to each local LAN management machine.

For large deployments, it may be necessary to distribute network management applications across multiple servers, either for performance reasons or simply to accommodate larger numbers of network devices.

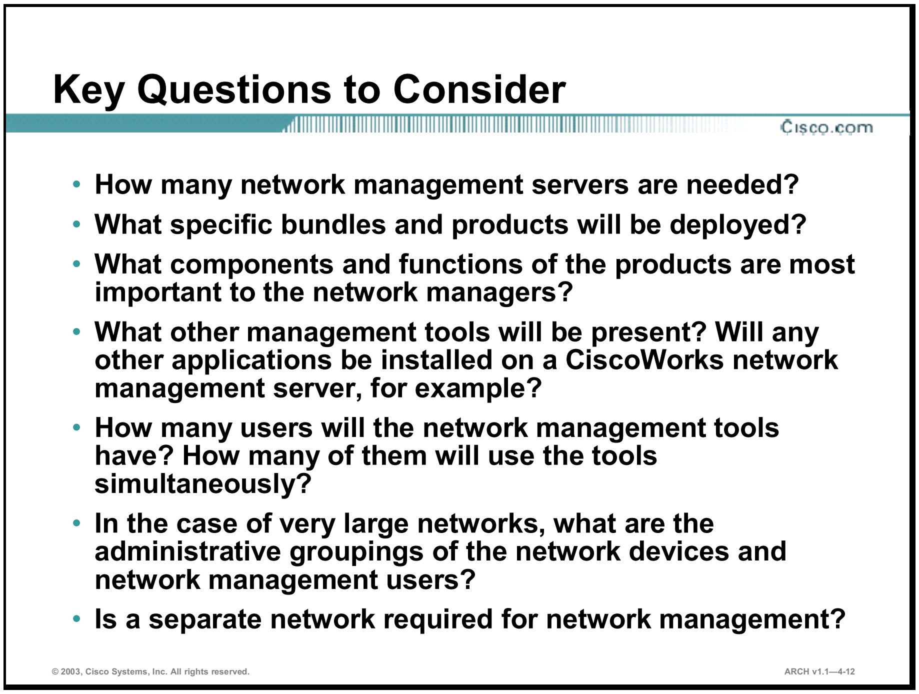

Key Questions to Consider

- How many network management servers are needed?
- What specific bundles and products will be deployed?
- What components and functions of the products are most important to the network managers?
- What other management tools will be present? Will any other applications be installed on a CiscoWorks network management server, for example?
- How many users will the network management tools have? How many of them will use the tools simultaneously?
- In the case of very large networks, what are the administrative groupings of the network devices and network management users?
- Is a separate network required for network management?

ARCH v1.1—4-12

The most common question, “What size workstation do I need for a network management station in order to manage X number of devices?” is difficult to answer. Consider the questions indicated in the figure as you define network management services to manage an enterprise network.

Example: Network Management Designs

When designing a small or large network management solution, you will consider the number of management stations, functionality requirements, resource utilization, and many other factors. In this topic, you will learn how to design Cisco network management solutions for small, medium, and large enterprise networks, given specific network management requirements.

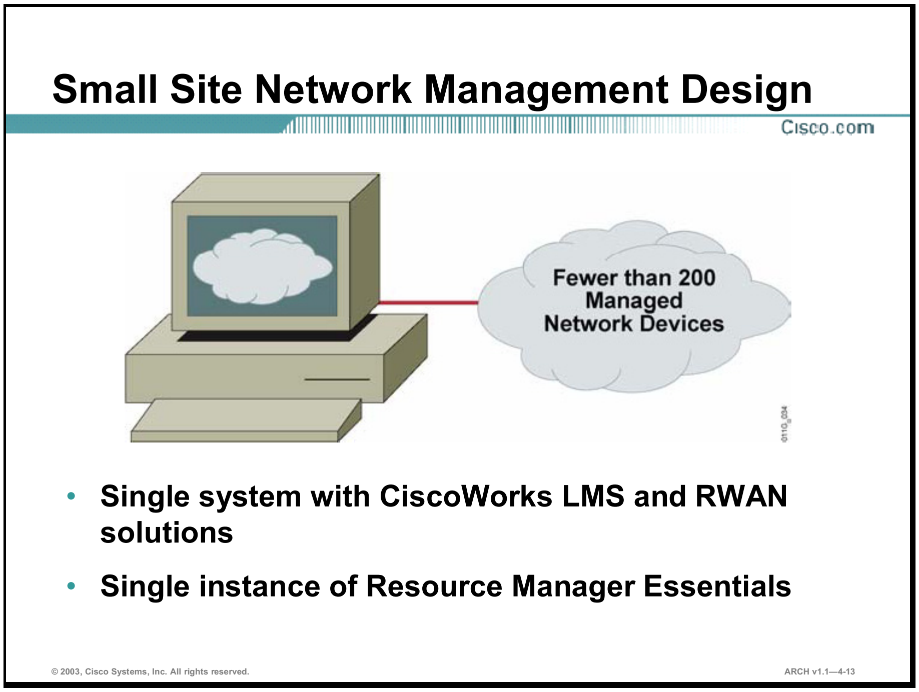

Small Site Network Management Design

- **Single system with CiscoWorks LMS and RWAN solutions**
- **Single instance of Resource Manager Essentials**

ARCH v1.1—4-13

For a small site with fewer than 200 network devices, a single CiscoWorks system with LMS and RWAN Management Solutions is likely sufficient. A single instance of RME can manage the entire network.

The table summarizes the design decisions that a small enterprise would make to meet their requirements.

Design Question	Decision	Justification
How many management domains does the enterprise require?	1 domain for the company	The domain will be managed centrally.
How many devices need to be managed?	Fewer than 200	
What are the key components and functions required?	LAN Management Solution Routed WAN Management Solution	For a small network, one network management server is sufficient.
How many servers are required?	One server with: ■ LAN Management Solution ■ RWAN Management Solution	
What administrative grouping of network devices will work for this enterprise?	All network devices within a single administrative grouping	Given the small size of the network, only one administrative grouping of network devices is required.
What administrative grouping of network management users will work for this enterprise?	All management users within a single administrative grouping	Given the small size of the network, only one administrative grouping of management users is required.

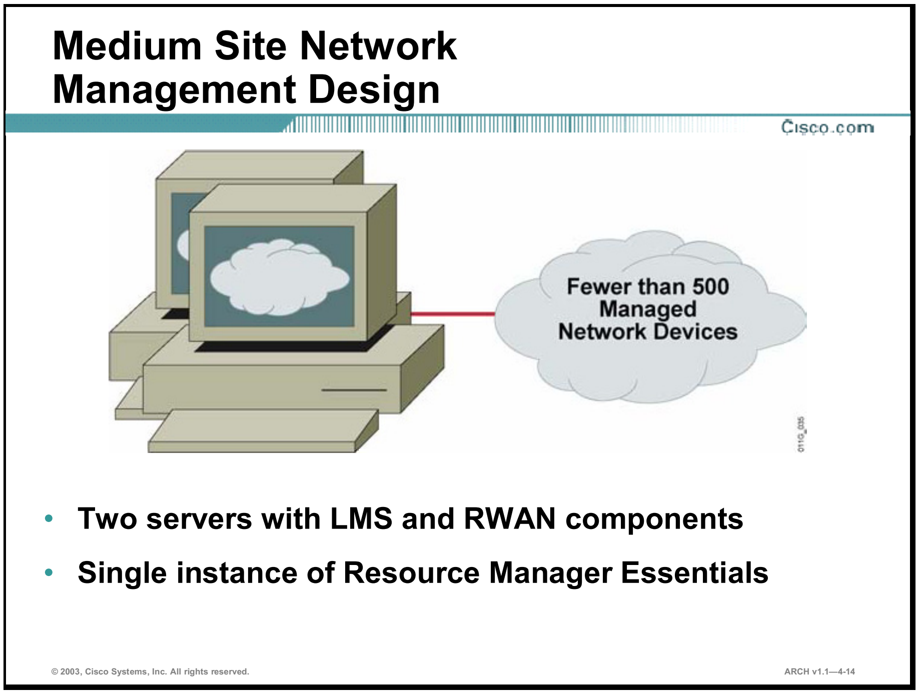

Medium Site Network Management Design

Cisco.com

Fewer than 500 Managed Network Devices

011G_035

- Two servers with LMS and RWAN components
- Single instance of Resource Manager Essentials

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-14

If a CiscoWorks RME Availability Manager application is to be used, Cisco recommends that there be no more than 500 devices monitored from a single RME server on an adequately equipped system. If a high-end system is used with multiple processors, each RME server can support up to 1000 network devices, depending on the full usage of the system running the RME software.

Note The functions provided by RME Availability Manager are often provided by third-party SNMP management platforms. Therefore, the 500-device limitation may not be an issue with common deployment scenarios, which include these other management servers.

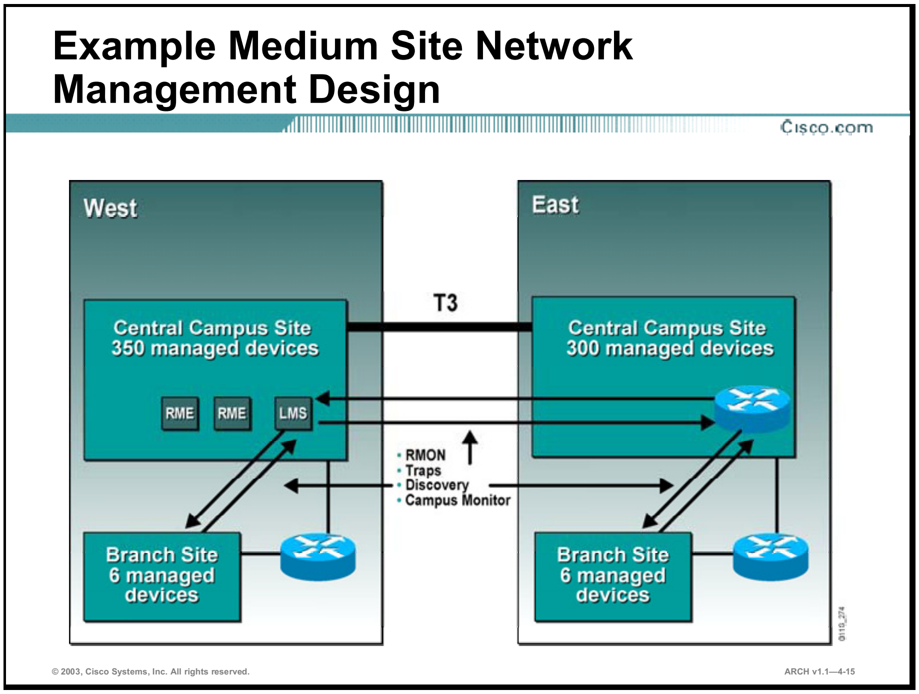

Example Medium Site Network Management Design

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-15

Company Background

In this scenario, there are two main sites: the west site in Los Angeles, California, and the east site in Pensacola, Florida. The west site supports 4500 users with 350 managed network devices. It is the hub for the west coast branch offices and connects to 10 branch offices with 6 managed devices at each office. The east site supports 7000 users with 300 managed network devices. The east site is the hub for 15 east coast branch offices with 6 managed devices per branch. There is a T3 between the east and west facilities. The company wants to use the Availability Manager to determine the status of their key devices. Each domain has up to 450 key devices to be managed. The network operations center is located in Los Angeles. All management will be done from the Los Angeles site.

Network Management Considerations

The company is planning to use Cisco RME for configuration and software management.

The IT staff is expecting to use a single domain while splitting management functions over multiple servers.

Availability Manager will be used to determine device operational status.

Network Management Design

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
How many management domains does the enterprise require?	One domain for the company	The domain will be managed from a central site.
How many devices need to be managed?	West: 410 managed devices East: 390 managed devices	
What are the key components and functions required?	LAN Management System including: ■ RME ■ Availability Manager ■ Campus Manager ■ Fault Manager RWAN Management Solution including: ■ Performance Monitor ■ ACL Manager	
How many servers are required?	Three total: ■ Two RME with Availability Manager ■ One for LMS/RWAN Management	The recommended number of devices associated with the Availability Manager is 500, so two servers are required.
What administrative grouping of network devices will work for this enterprise?	One administrative grouping	Given the requirements, there are two domains but one administrative group.
What administrative grouping of network management users will work for this enterprise?	One administrative group for all locations	

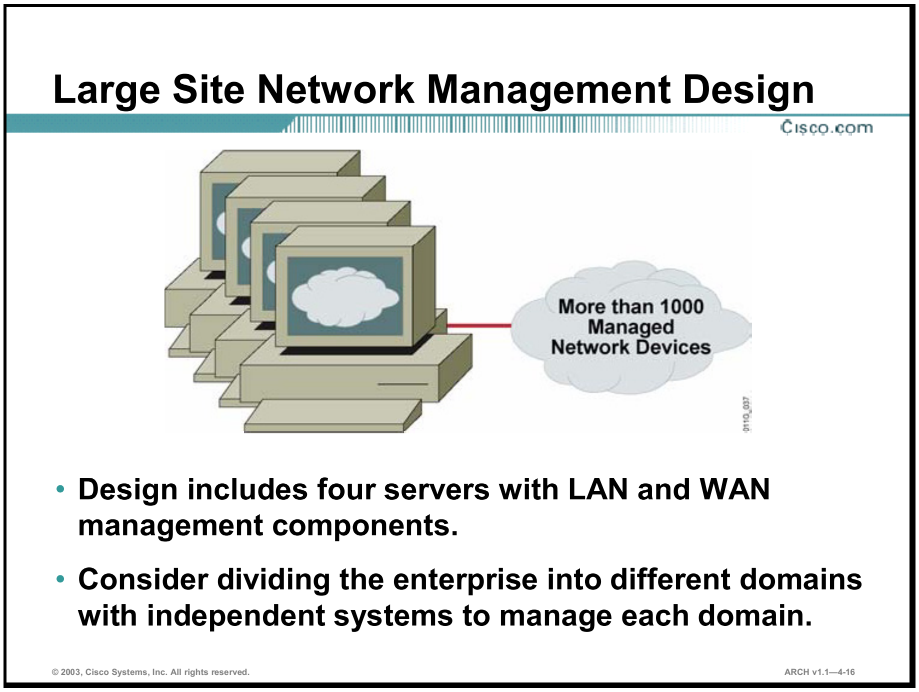

Large Site Network Management Design

Cisco.com

More than 1000 Managed Network Devices

- Design includes four servers with LAN and WAN management components.
- Consider dividing the enterprise into different domains with independent systems to manage each domain.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—4-16

When designing large and very large network management solutions, dividing the network into regions or domains, each with its own set of servers that manage individual domains, becomes a viable solution.

The scaling issues for a large site network management design include:

- Cisco RME
  - Total number of objects in inventory
  - Inventory updates of largest devices
  - Availability monitoring
  - Web GUI performance
  - Software update jobs for large numbers of devices at one time
  - Configuration-change jobs for large numbers of devices at one time
  - Syslog traffic level
- Cisco Campus Manager
  - Total number of devices or objects discovered and kept in database
  - Campus topology maps, which get very crowded and difficult to use for a large number of devices
  - User tracking limit on total number of end stations (number of rows in table)
  - User tracking discovery time
  - User tracking **ping** sweep (can disrupt network traffic)

- Cisco Device Fault Manager
  - Intelligent fault interpretation and device modeling resident in memory
  - Number of managed objects (chassis, modules, ports, and so on)
  - CPU and I/O utilization, which depend on polling intervals
- Other CiscoWorks Products
  - ACL Manager: The major scaling issue is the size of the access control lists (ACLs).
  - CiscoView: Server performance can be affected if too many simultaneous users are active. Five to ten is the recommended limit on the number of simultaneous users.
- Bundles
  - Bundles combine various products; they can be mixed together in a single deployment.
  - It is usually possible to run all products in the bundle on a single server for small to moderate size networks if sufficient system resources are available; for larger networks, it will be necessary to deploy multiple servers.
  - Plan around largest resource users.

Example Large Site Network Management Design

Cisco.com

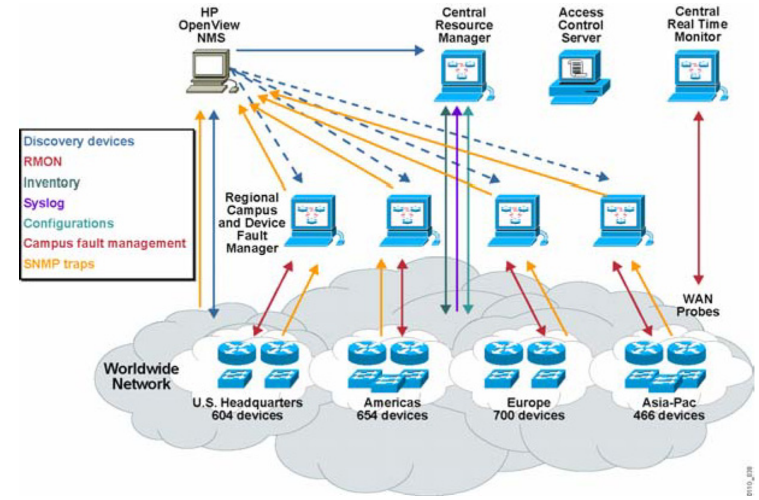

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—4-17

Company Background

In this scenario, the IT staff for AngelFish (a global fish reseller) wants to implement a centralized RME approach with each domain having their own LMS functionality. In this design, the planners have decided to split the global network into four manageable domains: the U.S. headquarters in Chicago, the Americas in Sao Paolo, Europe in Munich, and Asia-Pac in Seoul. Each domain will monitor and discover the devices within their domain.

Network Management Considerations

Each site manages approximately 600 routers. Chicago is designated the primary site where the RME system will be located. All four of the LMS systems will discover their regional asset information based on IP addresses. They will synchronize the device credential information with the central RME system. The Chicago site will need an RME and LMS server; all other domains would only need an LMS server. Any necessary polling is done by each of the four workstations, with polling intervals set appropriately.

Regional and local management of switched networks is performed with Cisco Campus Manager (under 1000 network devices) as well as RME. IP address ranges segment the regional networks. Different management groups do not share community strings, making configuration of CiscoWorks more difficult, but this policy is the preferred one for this customer.

Network Management Design

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
How many managed domains are there?	Four domains	The headquarters domain is assumed to be where the RME server will be located.
How many managed devices are there?	Headquarters: 950 Americas: 275 Europe: 867 Asia-Pac: 600	
What are the key components and functions required?	LAN Management System with: ■ RME ■ Campus Manager ■ Fault Manager RWAN Management Solution with: ■ Performance Monitor ■ ACL Manager	
How many servers are required?	One Server for RME in headquarters One HP OpenView Server One LMS/RWAN server at headquarters One LMS/RWAN server for Americas One LMS/RWAN server for Europe One LMS/RWAN server for Asia Pac	All domains will have at least one LMS/RWAN server. Corporate headquarters will utilize three servers: ■ One for HP OpenView ■ One for RME ■ One for LMS/RWAN
What administrative grouping of network devices will work for this enterprise?	Headquarters Americas Europe Asia-Pac	
What administrative grouping of network management users will work for this enterprise?	Headquarters Americas Europe Asia-Pac	

Summary

This topic summarizes the key points discussed in this lesson.

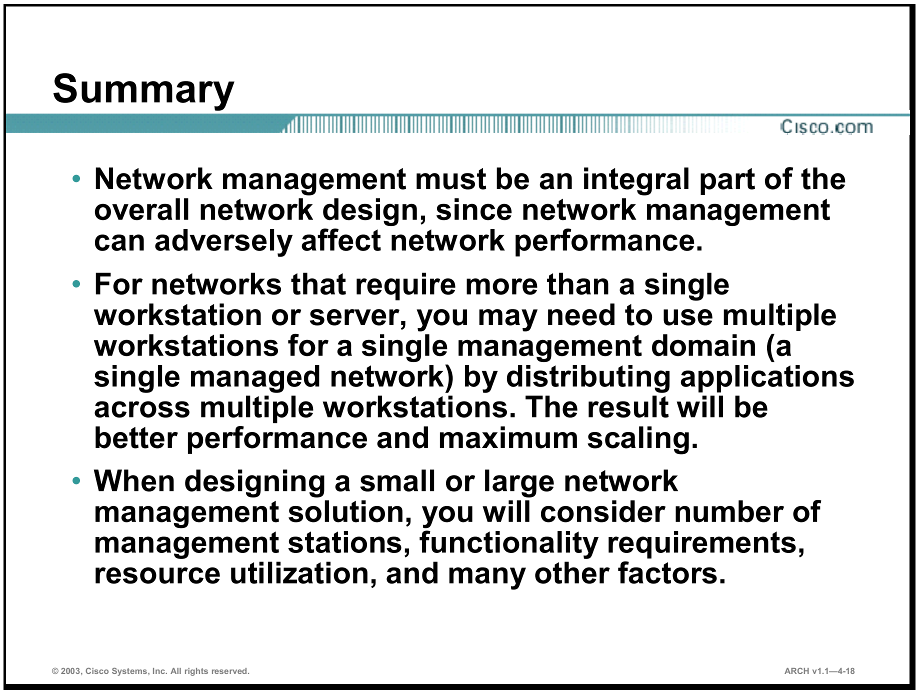

Summary

Cisco.com

- Network management must be an integral part of the overall network design, since network management can adversely affect network performance.
- For networks that require more than a single workstation or server, you may need to use multiple workstations for a single management domain (a single managed network) by distributing applications across multiple workstations. The result will be better performance and maximum scaling.
- When designing a small or large network management solution, you will consider number of management stations, functionality requirements, resource utilization, and many other factors.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—4-18

References

For additional information, refer to these resources:

- *CiscoWorks in Large-Scale Network Environments* available at http://www.cisco.com/warp/public/cc/pd/wr2k/prodlit/ckspp_wp.htm.
- *Network Management System: Best Practices White Paper* at http://www.cisco.com/warp/public/126/NMS_bestpractice.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study, refer to the following section:

- Case Study 4-2: OCSIC Bottling Company

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) When used for network management, SNMP, HTTP, and Telnet pose _____.

A) RMON errors

B) difficulty in use

C) bandwidth problems

D) inherent security risks

Q2) What feature most often impacts an enterprise's decision about whether to use a distributed network management strategy?

A) ease of use

B) amount of security required

C) ability to support the solution

D) availability of web-based tools

Q3) Which two items will determine the required number of management stations for a network? (Choose two.)

A) server performance

B) SNMP polling requirements

C) number of users in the enterprise

D) location of devices to be managed

E) number of managed network devices

Q4) In a large enterprise network with multiple administrative domains, what is the recommended solution for management domains?

A) Use one system to manage all domains on the network.

B) Plan the management so that each domain is managed separately.

C) Plan the management so that all domains are managed by a central system.

D) Use a large-scale system to manage multiple domains in one system. If system capacity is exceeded, add an additional system to handle the other domains.

Q5) Assume that you are designing a network management solution for a network with more than 5000 managed network devices. The design calls for several servers to handle the management needs. Which design solution would be a viable approach to dividing up the network?

A) syslog domains

B) regions or domains

C) switched LAN segments

D) global network operations centers

Quiz Answer Key

Q1) D

Relates to: Network Management Design Considerations

Q2) C

Relates to: Network Management Design Considerations

Q3) A, E

Relates to: Network Management Deployment Recommendations

Q4) B

Relates to: Example: Network Management Designs

Q5) B

Relates to: Example: Network Management Designs

Case Study 4-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

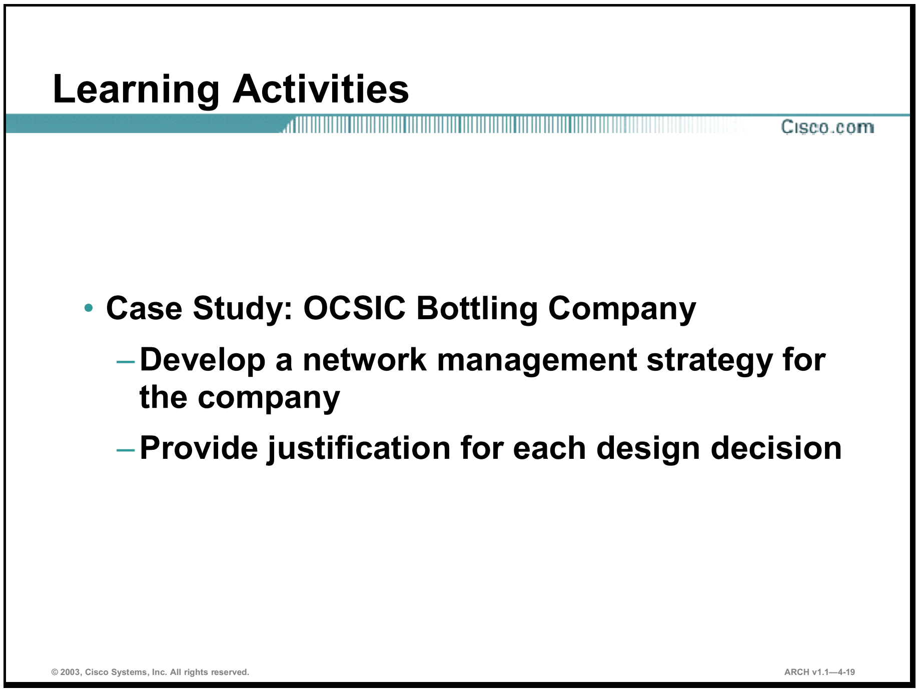

Learning Activities

- Case Study: OCSIC Bottling Company
  - Develop a network management strategy for the company
  - Provide justification for each design decision

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—4-19

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

In this exercise, you will design network management services that meet the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Develop a network management strategy for the company

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Develop a Network Management Strategy for the Company

Complete these steps:

Step 1 Complete the table to design the details about the Network Management solution for the OCSIC Bottling Company.

Design Question	Decision	Justification
How many management domains does the enterprise require?		
How many devices need to be managed?		
What are the key components and functions required?		
How many servers are required?		
What administrative grouping of network devices will work for this enterprise?		
What administrative grouping of network management users will work for this enterprise?		

Step 2 Describe policies and procedures to implement for the enterprise network management.

Step 3 What type of polling will you deploy on the network? How will polling affect network performance? How long would you recommend to maintain polling information?

Step 4 Update your campus network diagram to indicate the components of the Network Management module.

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have created a network management design that includes a completed campus diagram showing the location of each network management station, specification of the services and functionality that run on each server, and other servers required for the network, indicating their location and expected functionality.

Module 5

Designing High-Availability Services

Overview

As enterprises rely more and more heavily on their IP network for core business practices, a high degree of network availability becomes critical. System downtime translates into significant productivity and revenue losses.

Maximizing network uptime requires the use of operational best practices and redundant network designs in conjunction with high-availability technologies within network elements. Several high availability technologies are embedded in Cisco IOS software.

Module Objectives

Upon completing this module, you will be able to design high-availability intelligent network services for performance, scalability, and availability, given specified enterprise network needs.

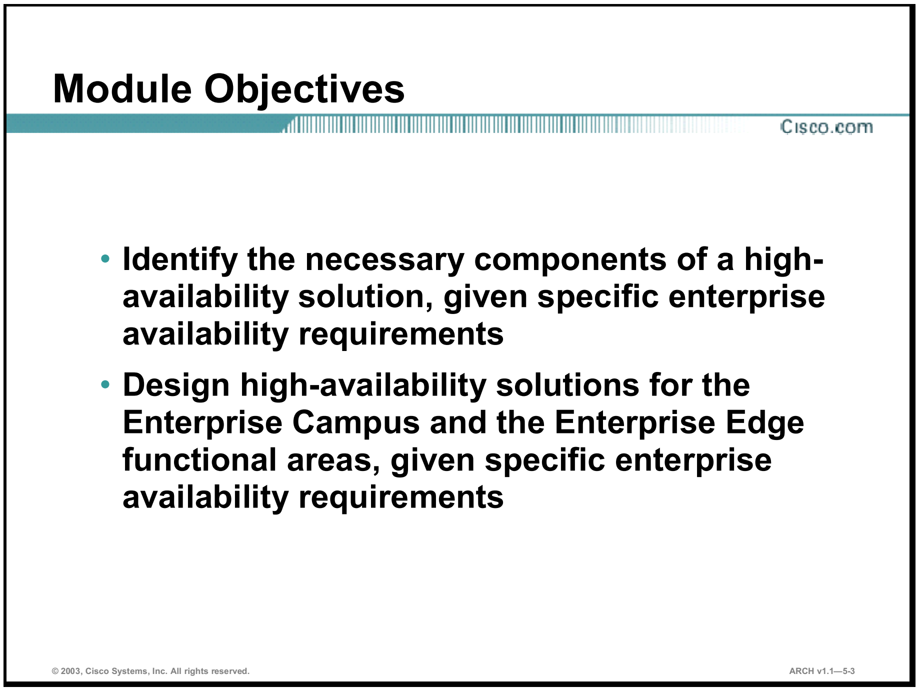

Module Objectives

- Identify the necessary components of a high-availability solution, given specific enterprise availability requirements
- Design high-availability solutions for the Enterprise Campus and the Enterprise Edge functional areas, given specific enterprise availability requirements

ARCH v1.1—5-3

Module Outline

The outline lists the components of this module.

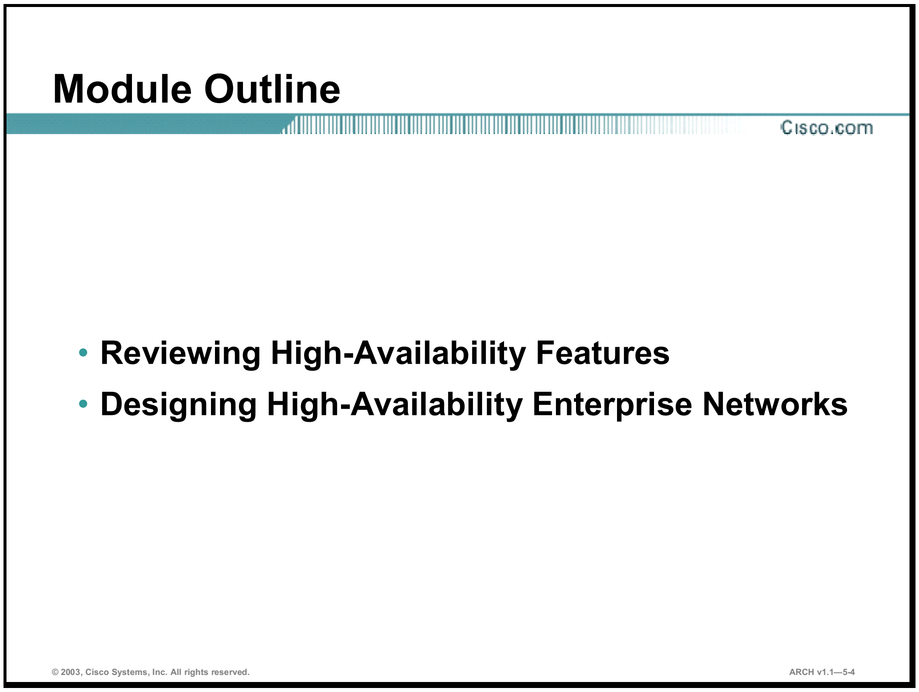

Module Outline

- Reviewing High-Availability Features
- Designing High-Availability Enterprise Networks

ARCH v1.1—5-4

Reviewing High-Availability Features

Overview

IOS high-availability technologies provide network redundancy and fault tolerance. Reliable network devices, redundant hardware components with automatic failover, and protocols like Hot Standby Router Protocol (HSRP) are used to maximize network uptime.

Relevance

Enterprises rely heavily on their network to run their business. Any downtime translates into lost revenue and productivity.

Objectives

Upon completing this lesson, you will be able to identify the necessary components of a high-availability solution, given specific enterprise availability requirements. This includes being able to meet these objectives:

- Identify requirements for network high availability
- Identify the necessary components of a high-availability solution, given specific high-availability requirements
- Determine when to implement fault tolerant network devices and redundant topologies, given specific high-availability requirements
- Determine when to use HSRP, Virtual Router Redundancy Protocol (VRRP), and options for Layer 3 redundancy, given specific high-availability requirements
- Determine when to use Spanning Tree Protocol and Layer 2 redundancy, given specific high-availability requirements

Learner Skills and Knowledge

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

- Designing Enterprise Campus Networks module
- Designing Enterprise Edge Connectivity module

Outline

The outline lists the topics included in this lesson.

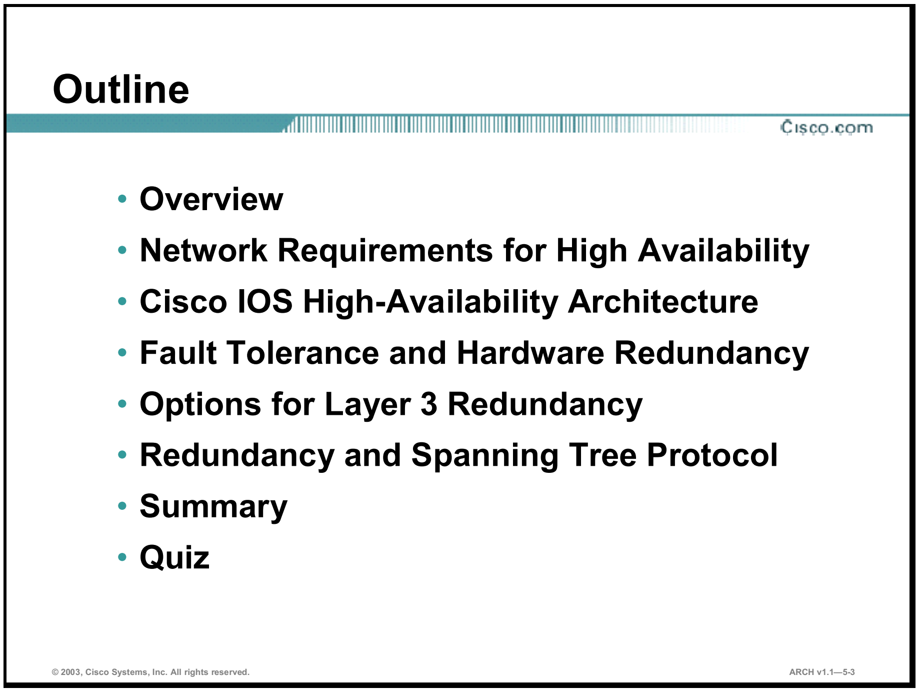

Outline

- Overview
- Network Requirements for High Availability
- Cisco IOS High-Availability Architecture
- Fault Tolerance and Hardware Redundancy
- Options for Layer 3 Redundancy
- Redundancy and Spanning Tree Protocol
- Summary
- Quiz

ARCH v1.1—5-3

Network Requirements for High Availability

An enterprise requires its network to be highly available to ensure that its mission-critical applications are available. Increased availability is a measurable quantity that translates into real cost savings. This topic describes the enterprise requirements for network high availability.

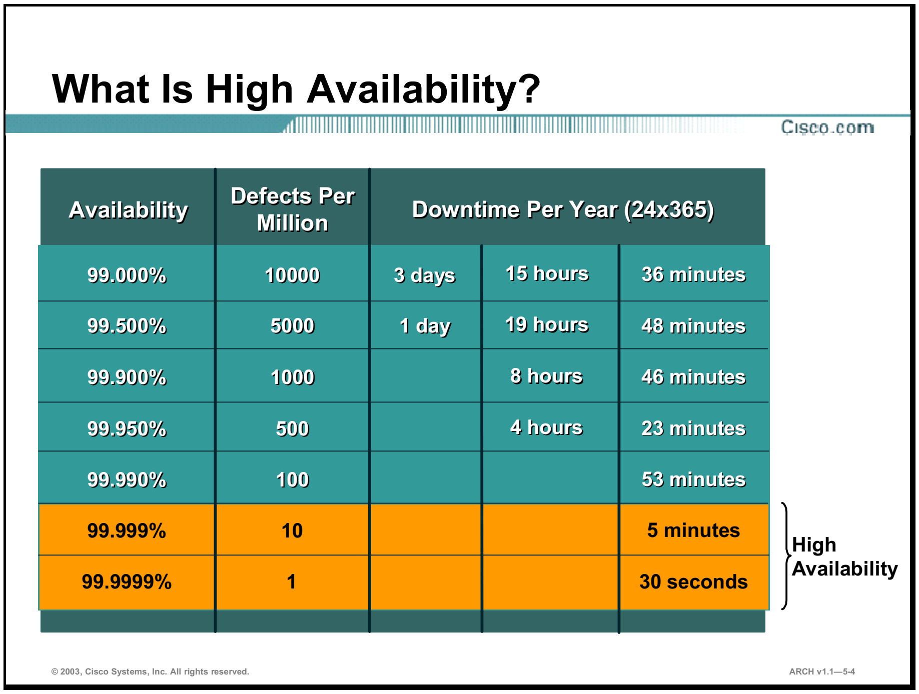

What Is High Availability?

Cisco.com

Availability	Defects Per Million	Downtime Per Year (24x365)		
99.000%	10000	3 days	15 hours	36 minutes
99.500%	5000	1 day	19 hours	48 minutes
99.900%	1000		8 hours	46 minutes
99.950%	500		4 hours	23 minutes
99.990%	100			53 minutes
99.999%	10			5 minutes
99.9999%	1			30 seconds

High Availability

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—5-4

Reliability implies that the system performs its specified task correctly. Availability means that the system is ready for immediate use.

For example, an airplane needs to be reliable. It does not need to be available to fly 24 hours a day. However, it must be reliable when the user determines it is ready to be available.

Today's networks need to be available 24 hours a day, 365 days a year. To meet that objective, 99.999 or 99.9999 percent availability is required.

Enterprises implement high availability to meet these requirements:

- **Ensure that mission-critical applications are available:** The purpose of an enterprise network is to enable applications. When those applications are not available, the enterprise ceases to function properly. Making the network highly available helps ensure that the enterprise's mission-critical applications are available.
- **Improve employee and customer satisfaction and loyalty:** Network downtime can cause frustration among both employees and customers attempting to access applications. Ensuring a highly available network helps to improve and maintain satisfaction and loyalty.
- **Reduce reactive information technology (IT) support costs, resulting in increased IT productivity:** Designing a network to incorporate high-availability technologies allows IT to minimize the time spent fighting fires and maximize the time providing proactive services.

- **Reduce financial loss:** An unavailable network, and therefore an unavailable application, can translate directly into lost revenue for an enterprise. Downtime can mean unbillable customer access time, lost sales, and contract penalties.
- **Minimize lost productivity:** When the network is down, employees cannot perform their functions efficiently. Lost productivity means increased cost to the enterprise.

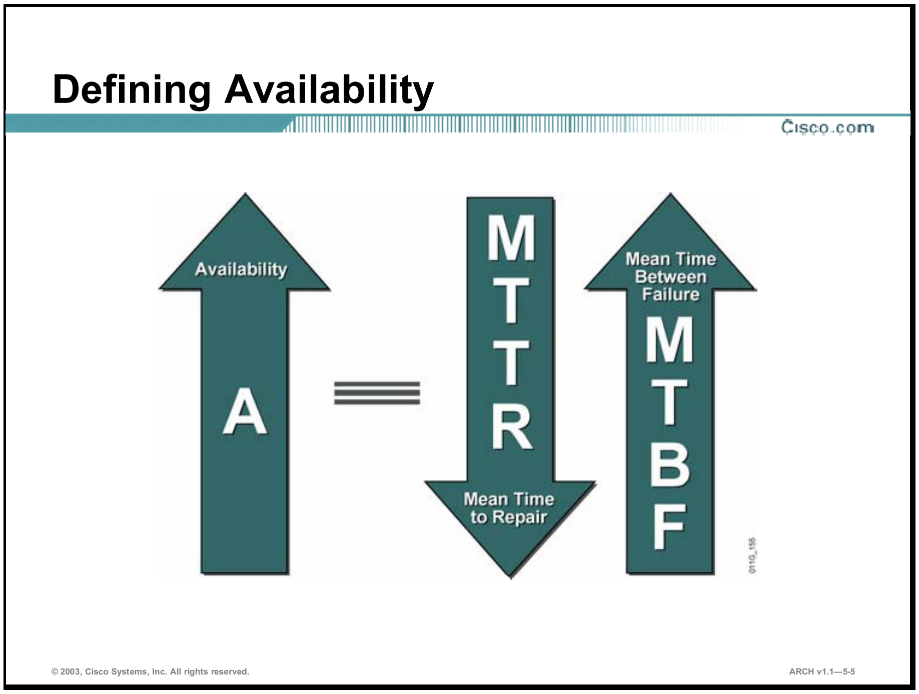

Defining Availability

Cisco.com

ARCH v1.1—5-5

Availability is a measurable quantity. The factors affecting availability are mean time to repair (MTTR), the time it takes to recover from a failure, and mean time between failure (MTBF), the time that passes between network outages or device failures.

Decreasing MTTR and increasing MTBF increase availability. Dividing MTBF by the sum of MTBF and MTTR results in a percentage indicating availability.

A common goal for availability is to achieve 99.999 percent (“five nines”). For example:

- Power supply MTBF = 40,000 hours
- Power supply MTTR = eight hours
- Availability = 40,000/(40,000 + 8) = 0.99980 or 99.98% availability

To calculate the availability of a complex system or device, multiply the availability of all of its parts. For example:

- Switch fabric availability = .99997
- Route processor availability = .99996
- System availability = .99997 * .99996 = 0.99992

As system complexity increases, availability decreases. If a failure of any one part causes a failure in the system as a whole, it is called serial availability.

Cisco IOS High-Availability Architecture

A highly available network requires reliable and fault-tolerant devices, resilient network technologies, optimized design, and implementation of best practices. This topic describes the components of a Cisco high-availability solution.

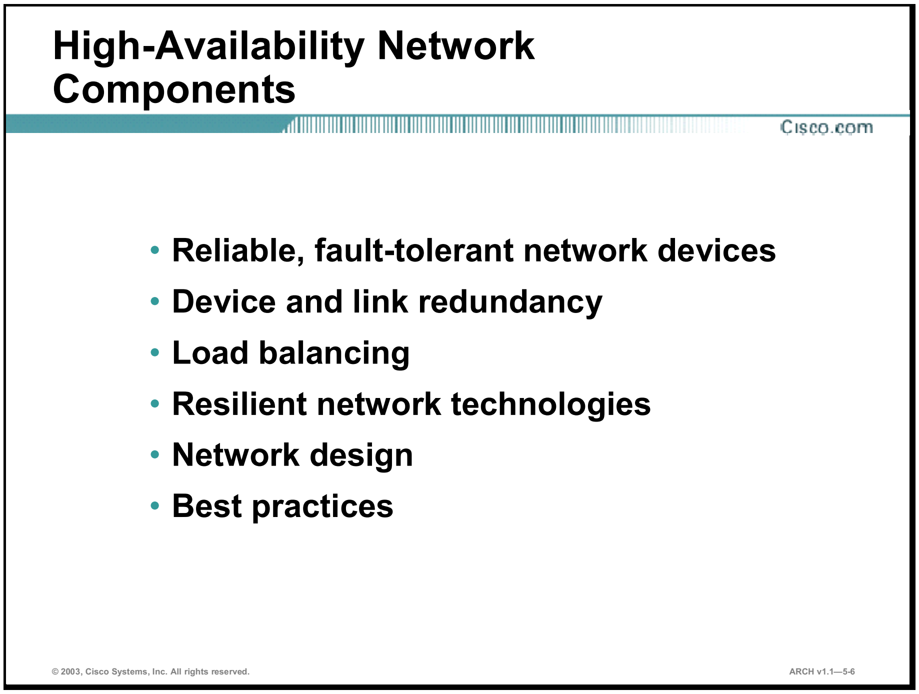

High-Availability Network Components

Cisco.com

- Reliable, fault-tolerant network devices
- Device and link redundancy
- Load balancing
- Resilient network technologies
- Network design
- Best practices

ARCH v1.1—5-6

To achieve high network availability, these network components are required:

- **Reliable, fault-tolerant network devices:** Hardware and software reliability to automatically identify and overcome failures
- **Device and link redundancy:** Entire devices, modules within devices, and links can be redundant
- **Load balancing:** Allows a device to take advantage of multiple best paths to a given destination
- **Resilient network technologies:** Intelligence that ensures fast recovery around any device or link failure
- **Network design:** Well-defined network topologies and configurations designed to ensure there is no single point of failure
- **Best practices:** Documented procedures for deploying and maintaining a robust e-commerce network infrastructure

High availability implies that a device or network is ready for use as close to 100 percent of the time as possible. Fault tolerance indicates the ability of a device or network to recover from the failure of a component or device. Achieving high availability relies on eliminating any single point of failure and on distributing intelligence throughout the architecture. You can increase availability by adding redundant components, including redundant network devices and connections to redundant Internet services. With the proper design, no single point of failure will impact the availability of the overall system.

Fault Tolerance and Hardware Redundancy

Enterprises can achieve high availability through the use of fault-tolerant devices, with redundancy provided within each device, and by the provision of multiple devices, with redundancy provided by the topology. This topic describes when to implement fault-tolerant network devices and redundant topologies.

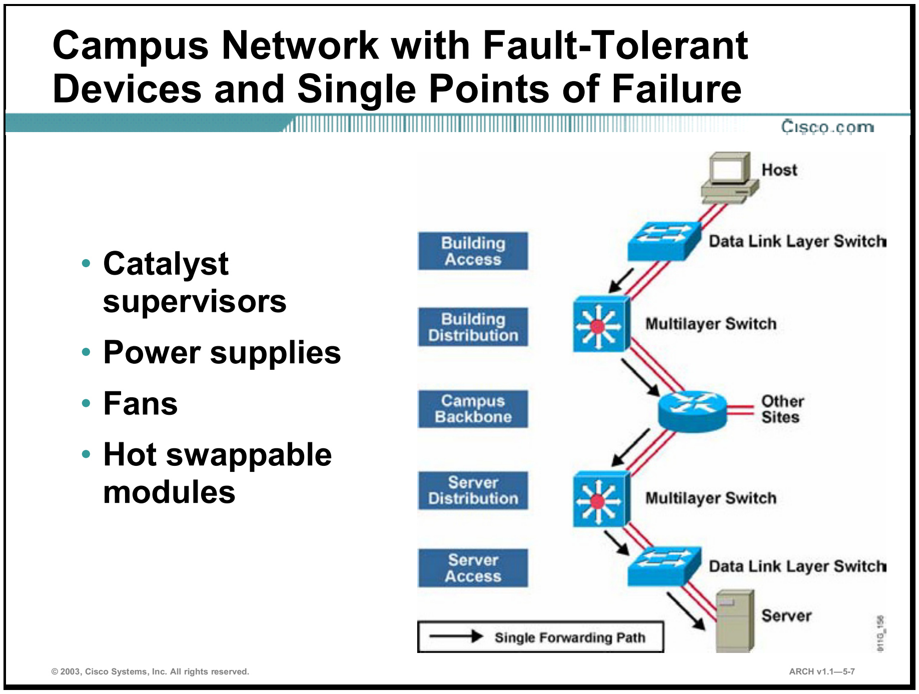

Campus Network with Fault-Tolerant Devices and Single Points of Failure

- Catalyst supervisors
- Power supplies
- Fans
- Hot swappable modules

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—5-7

One approach to building highly available networks is to use extremely fault-tolerant network devices throughout the network. To achieve high availability end-to-end, the fault tolerance of each device is optimized. This is achieved by providing redundant backup within the device for each of its key components. Fault tolerance offers these benefits:

- Minimizes time periods during which the system is nonresponsive to call-routing requests (for example, while the system is being reconfigured due to recovery)
- Eliminates all single points of failure that would cause the system to stop
- Provides disaster protection by allowing the major system components to be geographically separated

Achieving high network availability solely through device-level fault tolerance has drawbacks.

- Massive redundancy within each device adds significantly to its cost, while at the same time reducing physical capacity by consuming slots that could otherwise house network interfaces or provide useful network services.
- Redundant subsystems within devices are often maintained in a hot standby mode, in which they cannot contribute additional performance because they are only fully activated when the primary component fails.
- Focusing on device-level hardware reliability may result in overlooking a number of other failure mechanisms. Network elements are not standalone devices, but are components of a network system in which internal operations and system-level interactions are governed by configuration parameters and software.

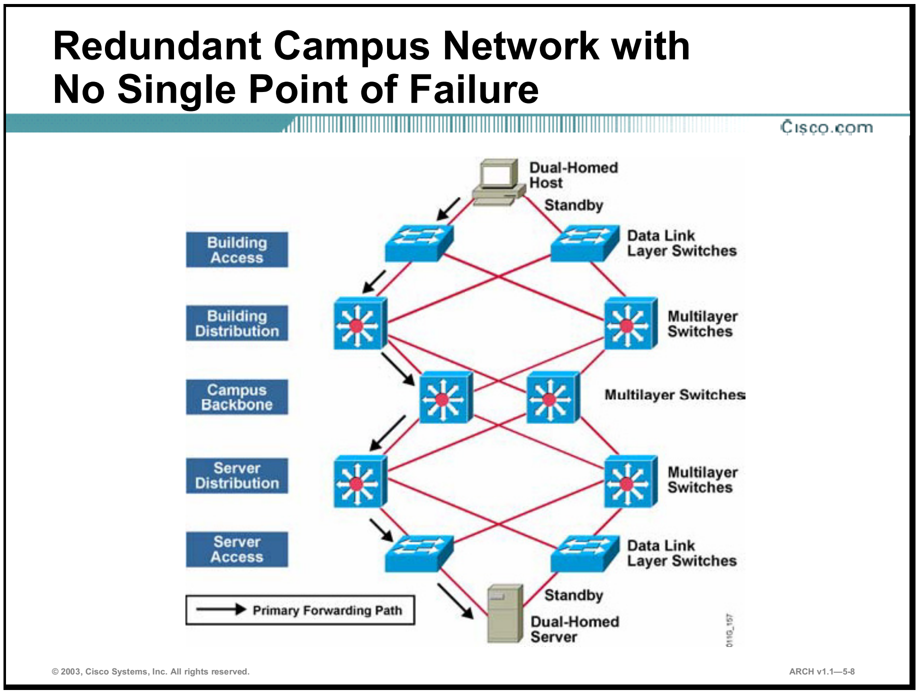

Redundant Campus Network with No Single Point of Failure

ARCH v1.1—5-8

A complementary way to build highly available networks is to provide reliability through redundancy in the network topology rather than primarily within the network devices themselves. In the campus network design shown in the figure, there is a backup for every link and every network device in the path between the client and server. This approach to network reliability offers these advantages:

- The network elements providing redundancy need not be located with the primary network elements. This reduces the probability that problems with the physical environment will interrupt service.
- Problems with software bugs and upgrades or configuration errors and changes can be dealt with separately in the primary and secondary forwarding paths without completely interrupting service. Therefore, network-level redundancy can also reduce the impact of non-hardware failure mechanisms.
- With the redundancy provided by the network, each network device no longer needs to be configured for optimal standalone fault tolerance. Device-level fault tolerance can be concentrated in the Campus Backbone and Building Distribution submodules of the network, where a hardware failure would affect a larger number of users. By partially relaxing the requirement for device-level fault tolerance, the cost per network device is reduced, to some degree offsetting the requirement for more devices.
- With carefully designed and implemented resiliency features, you can share the traffic load between the respective layers of the network topology (that is, Building Access and Building Distribution submodules) between the primary and secondary forwarding paths. Therefore, network-level redundancy can also provide increased aggregate performance and capacity.
- You can configure redundant networks to automatically failover from primary to secondary facilities without operator intervention. The duration of service interruption is equal to the time it takes for failover to occur. Failover times as low as a few seconds are possible.

Fast EtherChannel (FEC) is a trunking technology based on grouping together multiple full-duplex Fast Ethernets to provide fault-tolerant high-speed links between switches, routers, and servers. FEC uses a peer-to-peer control protocol that provides autoconfiguration and minimal convergence times for parallel links.

The drawbacks of link redundancy include:

- Increased media costs
- More difficult management and troubleshooting

As a data link layer feature, deterministic load distribution (DLD) adds reliability and predictable packet delivery with load balancing between multiple links.

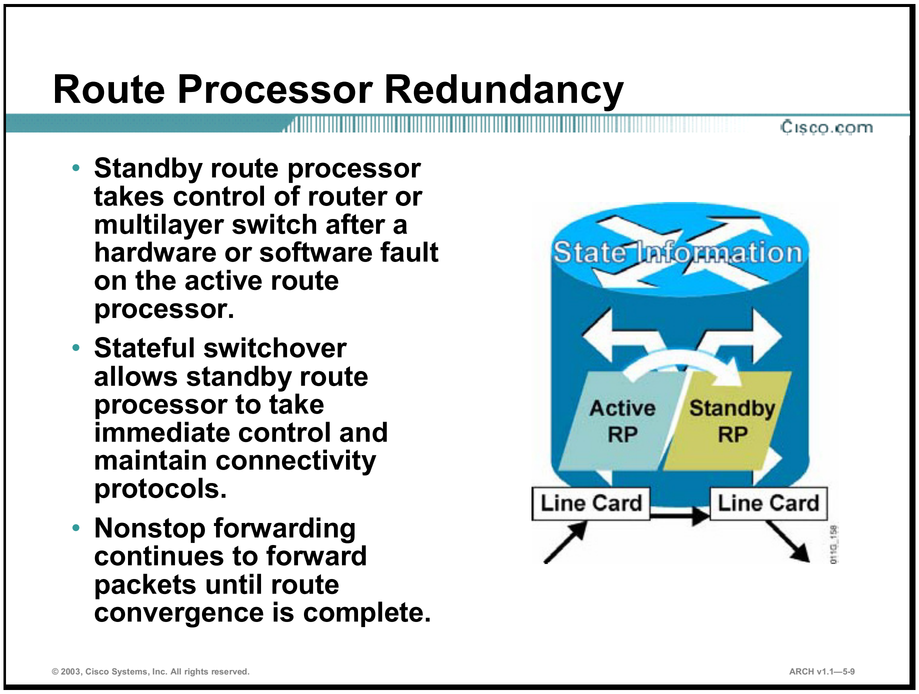

Route Processor Redundancy

Cisco.com

- Standby route processor takes control of router or multilayer switch after a hardware or software fault on the active route processor.
- Stateful switchover allows standby route processor to take immediate control and maintain connectivity protocols.
- Nonstop forwarding continues to forward packets until route convergence is complete.

ARCH v1.1—5-9

Route Processor Redundancy (RPR) provides a high system availability feature for some Cisco switches and routers. A system can reset and use a standby Route Switch Processor (RSP) in the event of a failure of the active RSP.

Using RPR, you can reduce unplanned downtime. RPR enables a quicker switchover between an active and standby RSP in the event of a fatal error on the active RSP. When you configure RPR, the standby RSP loads a Cisco IOS image upon bootup and initializes itself in standby mode. In the event of a fatal error on the active RSP, the system switches to the standby RSP, which reinitializes itself as the active RSP, reloads all of the line cards, and restarts the system.

RPR reduces the amount of unplanned downtime of a switch or router by enabling a faster startup time of a standby RSP.

RPR+ allows a failover to occur without reloading the line cards. The standby route processor takes over the router without affecting any other processes and subsystems. In addition, the RPR+ feature ensures that:

- The redundant processor is fully booted and the configuration is parsed
- The IOS running configuration is synchronized between active and standby route processors
- No link flaps occur during failover to the secondary router processor

The Cisco Catalyst 6500 offers software redundancy features that include Dual Router Mode (DRM) and Single Router Mode (SRM). These features provide redundancy between Multilayer Switch Feature Cards (MSFC) within the device.

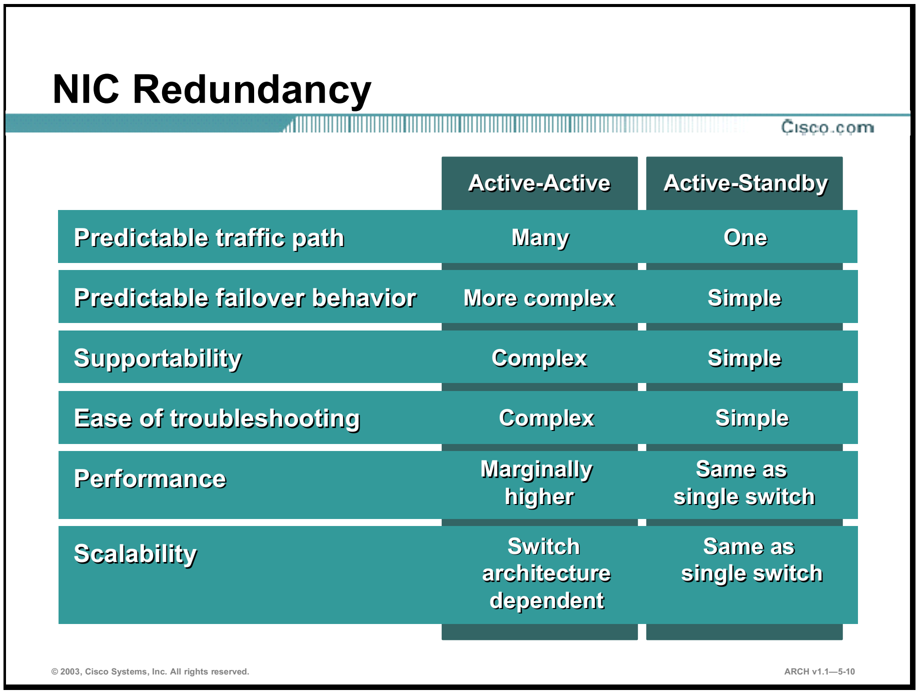

NIC Redundancy

	Active-Active	Active-Standby
Predictable traffic path	Many	One
Predictable failover behavior	More complex	Simple
Supportability	Complex	Simple
Ease of troubleshooting	Complex	Simple
Performance	Marginally higher	Same as single switch
Scalability	Switch architecture dependent	Same as single switch

ARCH v1.1—5-10

The option of dual-homing connected end systems is available. Most network interface cards (NICs) operate in an active-standby mode with a mechanism for MAC address portability between them. During a failure, the standby NIC becomes active on the new access switch.

Other end-system redundancy options include NICs operating in active-active mode, in which each host is available through multiple IP addresses. Either end-system redundancy mode requires more ports at the Building Access submodule.

Active-active redundancy implies that two redundant switches in a high-availability pair are concurrently load balancing traffic to server farms. Since both switches are active, you can support the same virtual IP address on each switch at the same time. This is known as shared VIP address. However, the use of active-active schemes supporting shared VIP configurations is not recommended.

Active-standby redundancy implies an active switch and a standby switch. The standby switch does not forward or load balance any traffic. The standby switch is only active in participating in the peering process that determines which switch is active and which is on standby. The peering process is controlled by the redundancy protocol used by the content switches.

Options for Layer 3 Redundancy

Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP) enable a set of routers to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. Protocols for router redundancy allow one router to automatically assume the function of the second router if the second router fails. This topic describes when to use HSRP, VRRP, and options for Layer 3 redundancy.

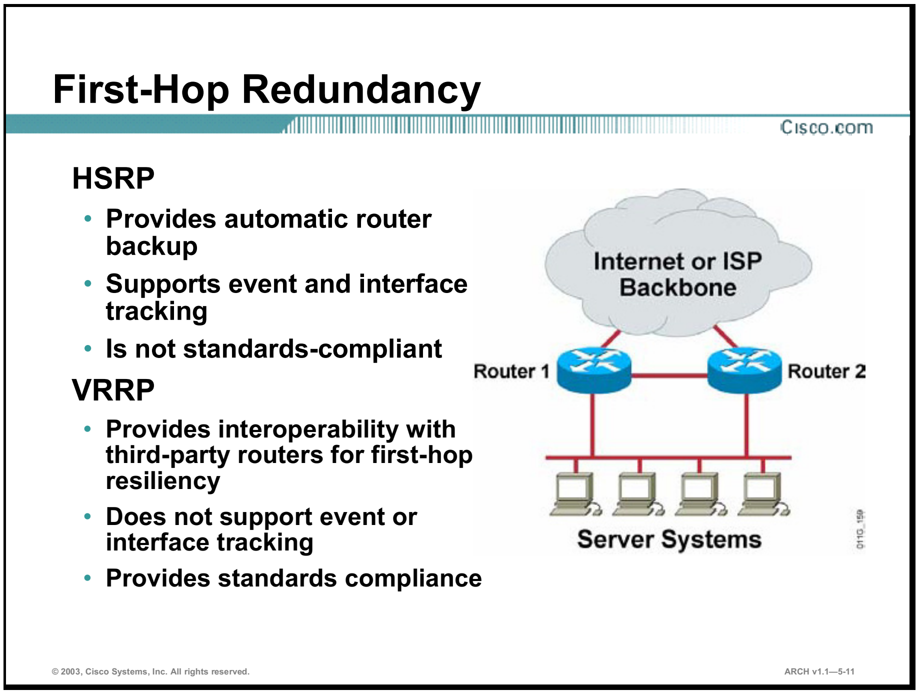

First-Hop Redundancy

Cisco.com

HSRP

- **Provides automatic router backup**
- **Supports event and interface tracking**
- **Is not standards-compliant**

VRRP

- **Provides interoperability with third-party routers for first-hop resiliency**
- **Does not support event or interface tracking**
- **Provides standards compliance**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—5-11

For IP, HSRP allows one router to automatically assume the function of the second router if the second router fails. HSRP is particularly useful when the users on one subnet require continuous access to resources in the network.

HSRP enables a set of routers to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. HSRP is particularly useful in environments where critical applications are running and fault-tolerant networks have been designed. By sharing an IP address and a hardware address, two or more routers acting as one virtual router are able to seamlessly assume routing responsibility in the case of a defined event or unexpected failure. This enables hosts on a LAN to continue to forward IP packets to a consistent IP and MAC address, enabling the changeover of devices doing the routing to be transparent to them and their sessions.

HSRP works by allowing an administrator to configure hot standby groups to share responsibility for an IP address. Each router can be given a priority to enable an administrator to weight the prioritization of routers for active router selection. Of the routers in each group, one will be selected as the active forwarder, and one will be selected as the standby router. These selections are done in accordance with the configured priorities of the router. The router with the highest priority wins and, in the case of a tie in priority, the greater value of their configured IP addresses will break the tie. Other routers in this group will monitor the active and standby routers' status to enable further fault tolerance. All HSRP routers participating in a standby group will watch for hello packets from the active and the standby routers. From the active router in the group, they will all learn the hello and dead timer as well as the standby IP address to be shared, if these parameters are not explicitly configured on each individual router. If the active router becomes unavailable due to scheduled maintenance, power failure, or other reasons, the standby can assume this functionality transparently within a few seconds. This will occur if the dead timer is reached, by missing three successive hello packets, and the standby router will promptly take over the virtual addresses, identity, and responsibility.

Multigroup HSRP (MHSRP) is an extension of HSRP that allows a single router interface to belong to more than one hot standby group. MHSRP requires the use of Cisco IOS software Release 10.3 or later and is supported only on routers that have special hardware that allows them to associate an Ethernet interface with multiple unicast MAC addresses, such as the Cisco 7000 series.

VRRP defines a standard mechanism that enables a pair of redundant (1 + 1) devices on the network to negotiate ownership of a virtual IP address. One device is elected to be active and the other to be standby. If the active fails, the backup takes over. An advantage of this scheme is that it achieves 1 + 1 redundancy without requiring any special intelligence. However, this scheme only works for n = 1 capacity and k = 1 redundancy; it will not scale above 1 + 1. VRRP is described in RFC 2338, at: http://www.ietf.org/rfc/rfc2338.txt.

Note Different Cisco routers and switches support different standby protocols. Verify the standby protocol support on a device before you select it.

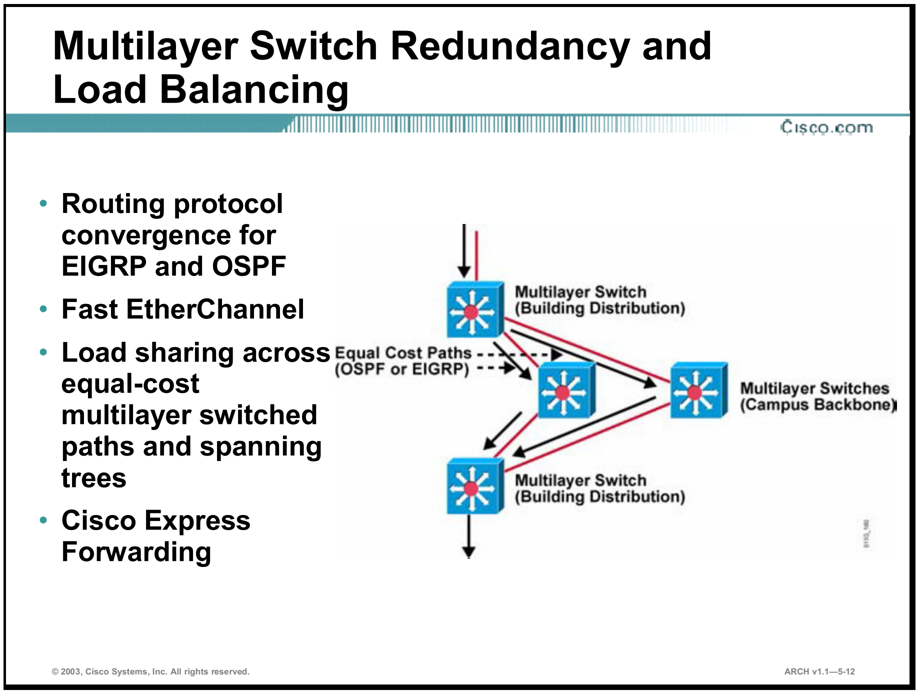

Multilayer Switch Redundancy and Load Balancing

Cisco.com

- Routing protocol convergence for EIGRP and OSPF
- Fast EtherChannel
- Load sharing across equal-cost multilayer switched paths and spanning trees
- Cisco Express Forwarding

Multilayer Switch (Building Distribution)
Equal Cost Paths (OSPF or EIGRP)
Multilayer Switches (Campus Backbone)
Multilayer Switch (Building Distribution)

ARCH v1.1—5-12

In addition to HSRP and VRRP, Cisco IOS software provides additional network redundancy features:

- Routing protocol convergence with Enhanced Interior Gateway Routing Protocol (EIGRP) and Open Shortest Path First (OSPF). EIGRP provides superior convergence properties and operating efficiency for Layer 3 load balancing and backup across redundant links and Cisco IOS devices to minimize congestion.
- Fast EtherChannel technology, which uses multiple Fast Ethernet links to scale bandwidth between switches, routers, and servers.
- Load sharing across equal-cost Layer 3 paths and spanning trees (for Layer 2-based networks).
- Cisco Express Forwarding (CEF) distributed switching architecture.

Redundancy and Spanning Tree Protocol

Cisco's spanning-tree implementation provides a separate spanning-tree domain for each VLAN, providing high availability while allowing traffic between the access and distribution layers of the network to be load balanced over redundant connections. This topic describes when to use Spanning Tree Protocol and Layer 2 redundancy.

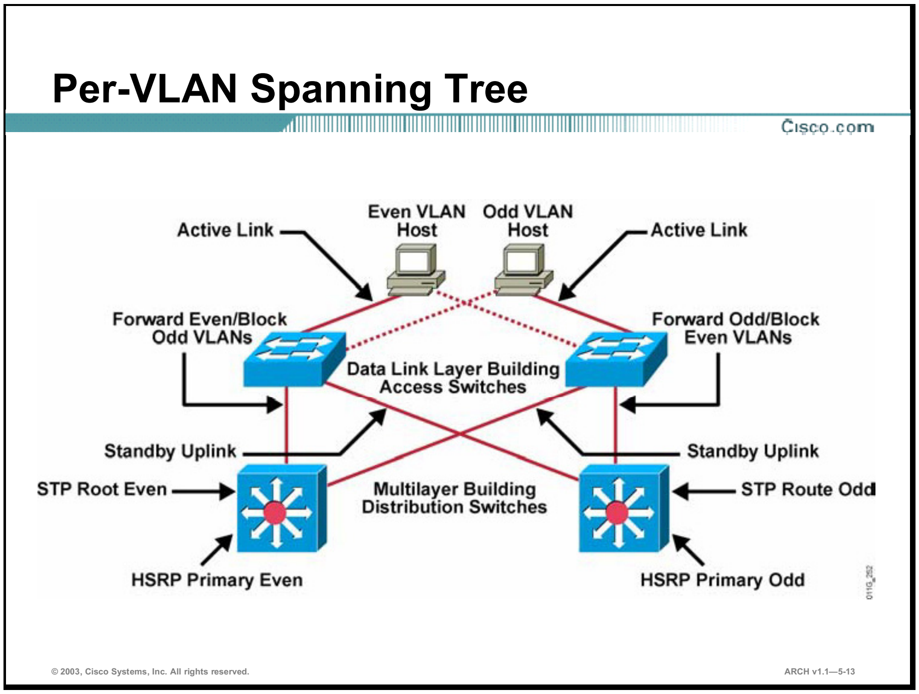

Per-VLAN Spanning Tree

© 2003, Cisco Systems, Inc. All rights reserved.   ARCH v1.1—5-13

The Spanning Tree Protocol was designed to prevent loops, but also provides advantages for redundancy. Cisco's spanning-tree implementation provides a separate spanning-tree domain for each VLAN, or Per VLAN Spanning Tree (PVST). PVST allows the bridge control traffic to be localized within each VLAN and supports configurations where the traffic between the access and distribution layers of the network can be load balanced over redundant connections. Cisco supports PVST over both Inter-Switch Link (ISL) and 802.1Q trunks. In the figure, the dotted lines represent alternate paths for VLAN traffic.

ISL and 802.1Q VLAN tagging also play an important role in load sharing across redundant links. All of the Layer 2 ISL between Building Access and Building Distribution switches are configured as trunks for all of the access VLANs. In the event of failure of the access switch or uplink, the most appropriate remaining uplink carries the traffic from all of the VLANs. With ISL or 802.1Q configured on the link Building Distribution switches, you can configure workgroup servers on either switch as part of either odd or even VLAN subnets.

Rapid Spanning Tree Protocol (RSTP, 802.1w) significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP selects one switch as the root of a connected spanning-tree active topology and assigns port roles to individual ports of the switch, depending on whether that port is part of the active topology.

RSTP provides rapid connectivity following the failure of a switch, switch port, or a LAN. A new root port and the designated port on the other side of the bridge transition to forwarding using an explicit handshake between them. RSTP allows switch port configuration so that the ports can transition to forwarding directly when the switch reinitializes.

RSTP as specified in 802.1w supersedes STP specified in 802.1D, but remains compatible with STP.

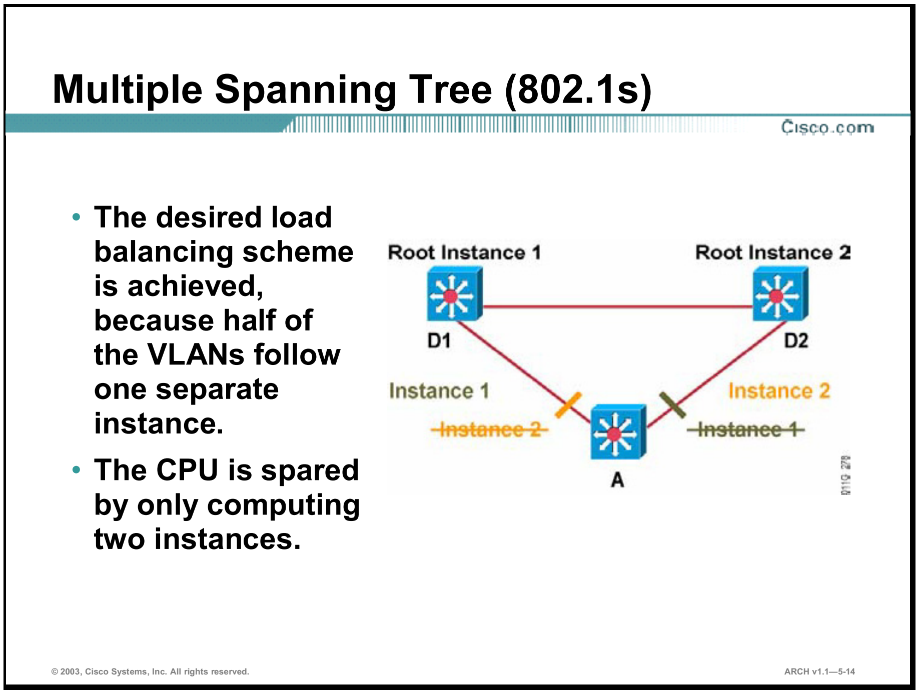

Multiple Spanning Tree (802.1s)

- **The desired load balancing scheme is achieved, because half of the VLANs follow one separate instance.**
- **The CPU is spared by only computing two instances.**

ARCH v1.1—5-14

Multiple Spanning Tree (MST, IEEE 802.1s) allows you to map several VLANs to a reduced number of spanning-tree instances, because most networks do not need more than a few logical topologies. In the topology described in the figure, there are only two different final logical topologies, so only two spanning-tree instances are really necessary. There is no need to run a thousand instances. If you map half of the 1000 VLANs to a different spanning-tree instance, as shown in the figure, the following is true:

- The desired load balancing scheme is realized, because half of the VLANs follow one separate instance.
- The CPU is spared by only computing two instances.

From a technical standpoint, MST is the best solution. From an end-user's perspective, the only drawbacks associated with migrating to MST are mainly due to the fact that MST is a new protocol, and these issues arise:

- The protocol is more complex than the usual spanning tree and requires additional training of the staff.
- Interaction with legacy bridges is sometimes challenging.

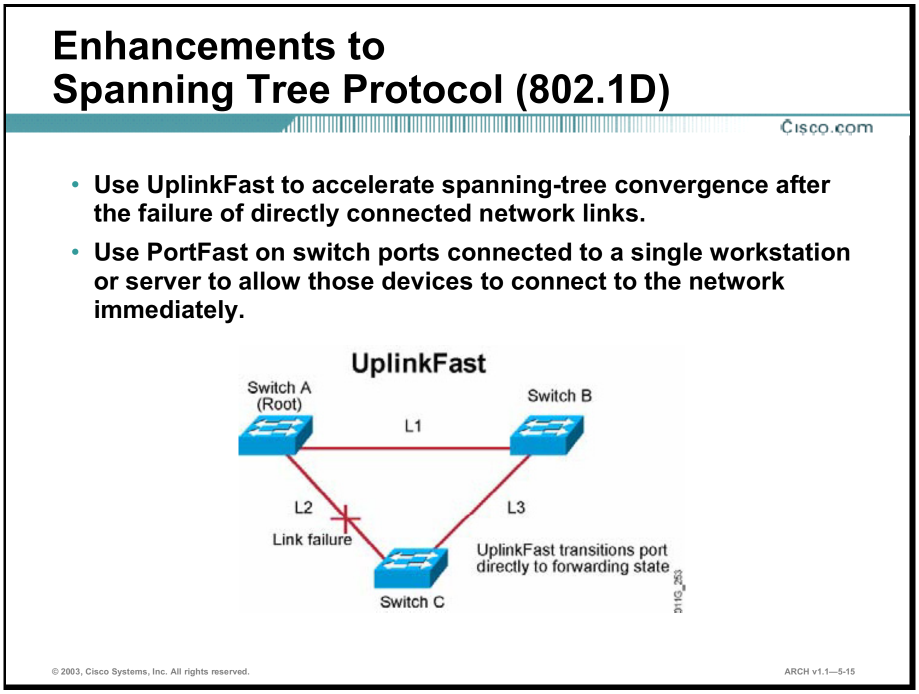

Enhancements to Spanning Tree Protocol (802.1D)

Cisco.com

- **Use UplinkFast to accelerate spanning-tree convergence after the failure of directly connected network links.**
- **Use PortFast on switch ports connected to a single workstation or server to allow those devices to connect to the network immediately.**

ARCH v1.1—5-15

The Spanning Tree Protocol (802.1D) was designed for robust, plug-and-play operation in bridged networks, or arbitrary connectivity (looping), and almost unlimited flatness.

To improve spanning-tree convergence, you can implement PortFast. PortFast is a feature that you can enable on Catalyst switch ports dedicated to connecting single servers or workstations. PortFast allows the switch port to begin forwarding as soon as the end system is connected, bypassing the listening and learning states and eliminating up to 30 seconds of delay before the end system can begin sending and receiving traffic. PortFast is used when an end system is initially connected to the network or when the primary link of a dual-homed end system or server is reactivated after a failover to the secondary link. Since only one station is connected to the segment, there is no risk of PortFast creating network loops.

In the event of a failure of a directly connected uplink that connects a Building Access switch to a Building Distribution switch, you can increase the speed of spanning-tree convergence by enabling the UplinkFast feature on the Building Access switch. With UplinkFast, each VLAN is configured with an uplink group of ports, including the root port that is the primary forwarding path to the designated root bridge of the VLAN, and one or more secondary ports that are blocked. When a direct uplink fails, UplinkFast unblocks the highest priority secondary link and begins forwarding traffic without going through the spanning-tree listening and learning states. Bypassing listening and learning reduces the failover time after uplink failure to approximately the bridge protocol data unit (BPDU) hello interval (1 to 5 seconds). With the default configuration of standard STP, convergence after uplink failure can take up to 30 seconds.

Summary

This topic summarizes the key points discussed in this lesson.

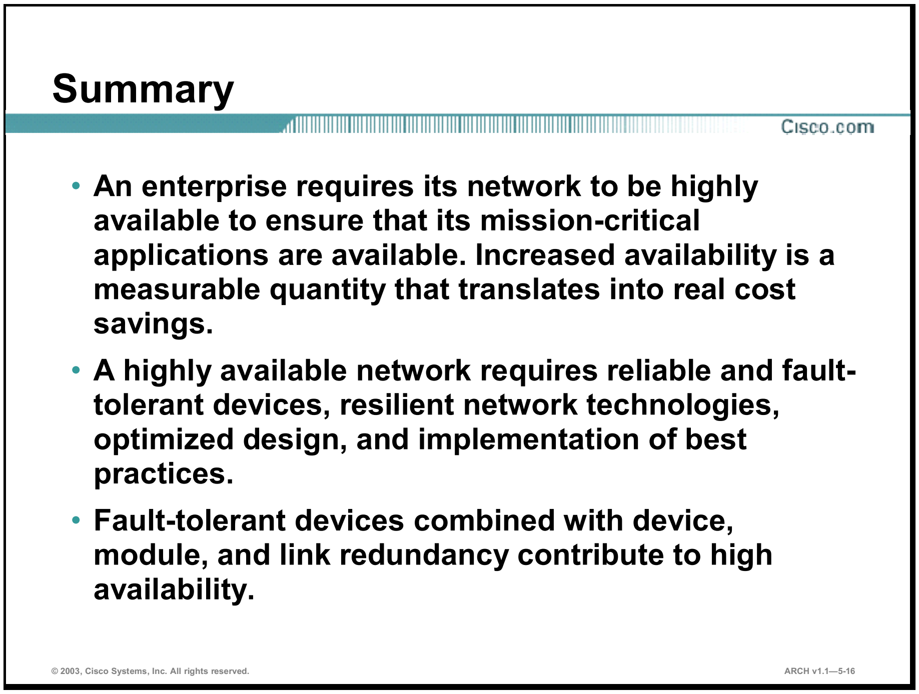

Summary

- An enterprise requires its network to be highly available to ensure that its mission-critical applications are available. Increased availability is a measurable quantity that translates into real cost savings.
- A highly available network requires reliable and fault-tolerant devices, resilient network technologies, optimized design, and implementation of best practices.
- Fault-tolerant devices combined with device, module, and link redundancy contribute to high availability.

ARCH v1.1—5-16

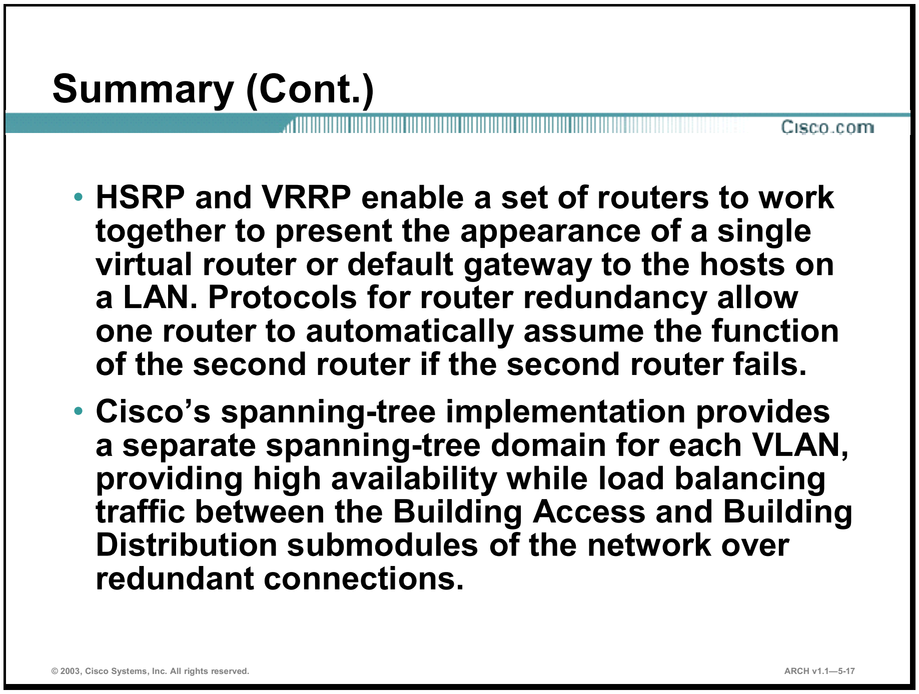

Summary (Cont.)

- HSRP and VRRP enable a set of routers to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. Protocols for router redundancy allow one router to automatically assume the function of the second router if the second router fails.
- Cisco's spanning-tree implementation provides a separate spanning-tree domain for each VLAN, providing high availability while load balancing traffic between the Building Access and Building Distribution submodules of the network over redundant connections.

ARCH v1.1—5-17

References

For additional information, refer to these resources:

- *High Availability Services* at http://www.cisco.com/warp/public/779/largeent/learn/technologies/availability.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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) What are two measurable costs of network downtime? (Choose two.)

A) revenue losses

B) productivity losses

C) reduced customer loyalty

D) minimized competitive edge

E) reduced employee satisfaction

Q2) Which two components are required for achieving high availability? (Choose two.)

A) providing redundant Internet services

B) eliminating any single point of failure

C) centralizing intelligence within the architecture

D) distributing intelligence throughout the architecture

E) providing a redundant backup for each network component

Q3) Which three drawbacks can apply when attempting to achieve reliability solely through extremely fault-tolerant devices? (Choose three.)

A) increased device cost

B) increased cabling cost

C) increased device efficiency

D) limited administrative control

E) components not contributing to performance

Q4) Which method would be a good choice to provide high availability while keeping hardware costs down?

A) HSRP

B) device fault tolerance

C) redundant network topology

D) redundant systems in hot standby mode

Q5) Which HSRP state specifies that a router is transferring packets?

A) active

B) standby

C) listening

D) speaking and listening

Q6) Which technology would you use if your users require continuous, uninterrupted Layer 3 network access?

A) RRI

B) xSTP

C) HSRP

D) RIPv2

Q7) Which Cisco IOS feature supports configurations where the traffic between the access and distribution layers of the network can be load balanced over redundant connections?

A) ISL

B) PVST

C) 802.1 Q

D) PortFast

Quiz Answer Key

Q1) A, B

Relates to: Network Requirements for High Availability

Q2) B, D

Relates to: Cisco IOS High Availability Architecture

Q3) A, D, E

Relates to: Fault Tolerance and Hardware Redundancy

Q4) C

Relates to: Fault Tolerance and Hardware Redundancy

Q5) A

Relates to: Options for Layer 3 Redundancy

Q6) C

Relates to: Options for Layer 3 Redundancy

Q7) B

Relates to: Redundancy and Spanning-Tree Protocol

Designing High-Availability Enterprise Networks

Overview

The Enterprise Campus and the Enterprise Edge require the availability of the network resources in environments to provide effective performance, scalability, and availability. The network designer must incorporate high-availability features into each location on the network.

Relevance

High-availability solutions must be implemented throughout the network to ensure maximum system availability.

Objectives

Upon completing this lesson, you will be able to design high-availability solutions for the Enterprise Campus and the Enterprise Edge functional areas, given specific enterprise availability requirements. This includes being able to meet these objectives:

- List design guidelines for each component of an enterprise network
- Identify best practices recommendations to ensure high availability of the network
- Describe the basic guidelines for designing the Campus Infrastructure functional area for high availability
- Describe the basic guidelines for designing the Enterprise Edge functional area for high availability
- Describe a high-availability strategy for an enterprise site

Learner Skills and Knowledge

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

- Designing Enterprise Campus Networks module
- Designing Enterprise Edge Connectivity module
- Reviewing Cisco High-Availability Features lesson

Outline

The outline lists the topics included in this lesson.

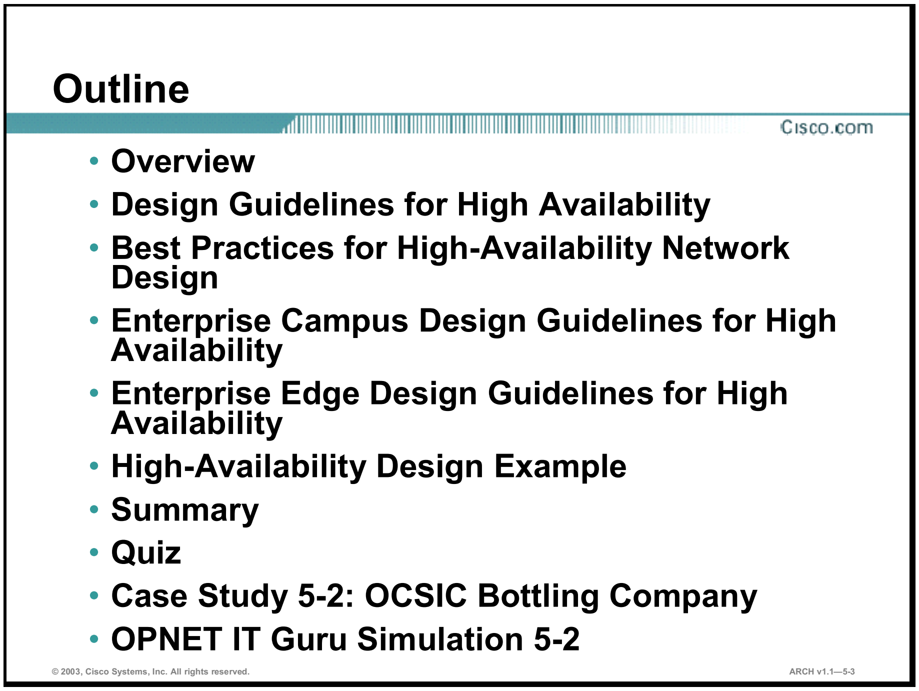

Outline

Cisco.com

- Overview
- Design Guidelines for High Availability
- Best Practices for High-Availability Network Design
- Enterprise Campus Design Guidelines for High Availability
- Enterprise Edge Design Guidelines for High Availability
- High-Availability Design Example
- Summary
- Quiz
- Case Study 5-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 5-2

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—5-3

Design Guidelines for High Availability

When designing a network for high availability, you will consider the reliability of each network hardware and software component, redundancy, protocol attributes, circuits and carriers, and environmental and power features that contribute to the overall availability of the network. This topic lists design guidelines for each component of an enterprise network.

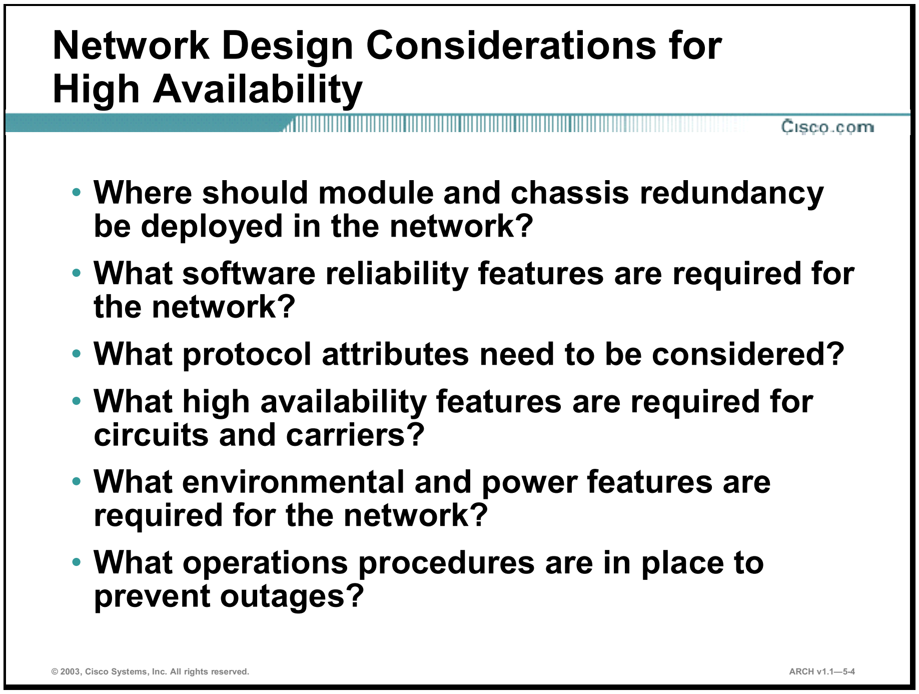

Network Design Considerations for High Availability

- Where should module and chassis redundancy be deployed in the network?
- What software reliability features are required for the network?
- What protocol attributes need to be considered?
- What high availability features are required for circuits and carriers?
- What environmental and power features are required for the network?
- What operations procedures are in place to prevent outages?

ARCH v1.1—5-4

To design high-availability services for an enterprise network, you will answer the questions listed in the figure. Each question is considered in this lesson.

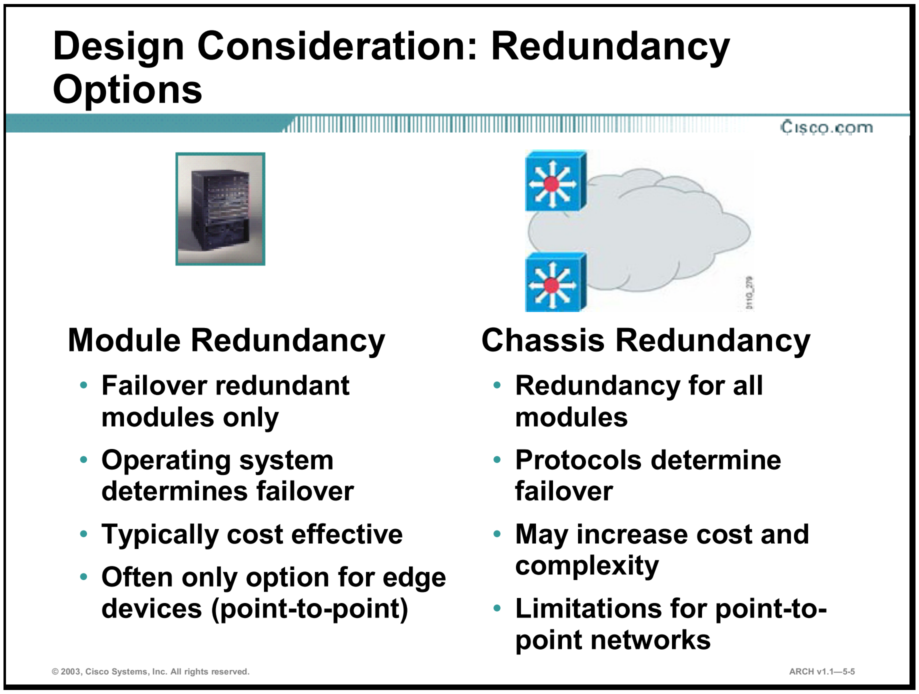

Design Consideration: Redundancy Options

Module Redundancy

- Failover redundant modules only
- Operating system determines failover
- Typically cost effective
- Often only option for edge devices (point-to-point)

Chassis Redundancy

- Redundancy for all modules
- Protocols determine failover
- May increase cost and complexity
- Limitations for point-to-point networks

ARCH v1.1—5-5

The options for device redundancy include both module and chassis redundancy. Both types of redundancy are usually most important at the Building Distribution and Campus Backbone submodules. The decision about which to use is based on the criticality of the resource and the cost of redundancy.

With module redundancy, only selected modules are selected for failover. In the event that the primary module fails, the device operating system determines the failover. Module redundancy is typically the most cost-effective redundancy option available, and is the only option (over chassis redundancy) for edge devices in point-to-point topologies.

With chassis redundancy, the entire chassis and all modules within it are redundant. In the event of a failure, the protocols running on the network, such as HSRP or VRRP, determine how the failover occurs. Chassis redundancy increases the cost and complexity of the network, which are factors to consider when selecting device redundancy. Chassis redundancy is also limited for point-to-point edge networks. To calculate the theoretical advantage gained with redundant modules or chassis, use the following formula:

- Availability = 1 – [(1 – availability$_1$) * (1 – availability$_2$)]

For example, if you implement a redundant switch fabric with 100 percent failure detection, calculate availability as follows:

- Availability = 1 – [(1 – .99997) * (1 – . 99997)]
- 1 – [(.00003) * (.00003) = 1 – [.0000000009] = 0.99999

Therefore, redundant switch fabrics increase the availability of the component to 99.9999 percent. This is known as parallel availability.

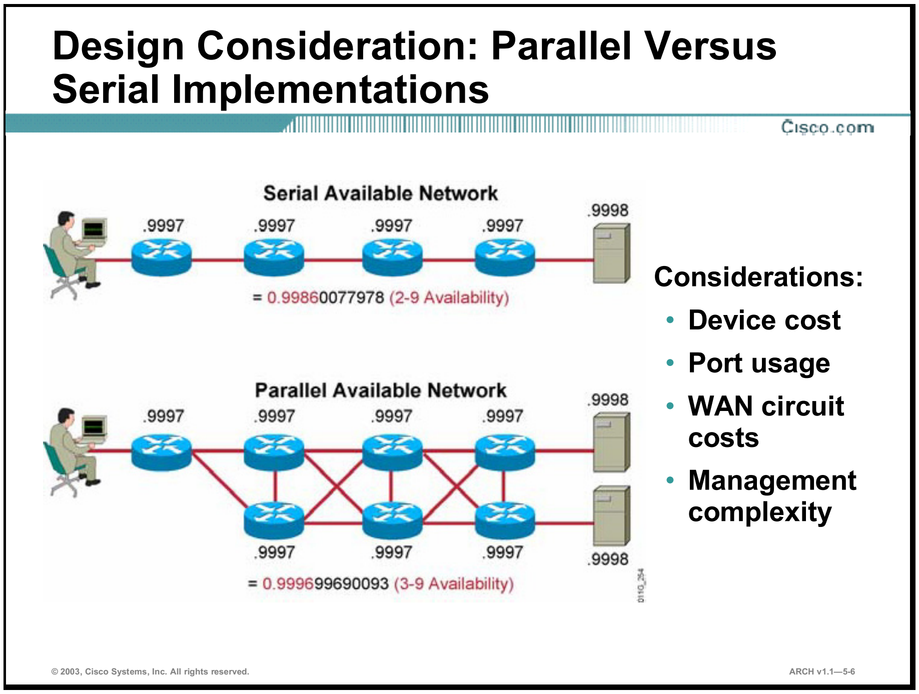

Design Consideration: Parallel Versus Serial Implementations

Cisco.com

Considerations:

- **Device cost**
- **Port usage**
- **WAN circuit costs**
- **Management complexity**

ARCH v1.1—5-6

Link redundancy, implemented through parallel or serial implementations, can increase availability significantly.

To calculate the theoretical advantage gained with redundant links, use the following formula:

- Availability = $[1 – (1 – \text{availability}_1)^2] * [1 – (1 – \text{availability}_2)^2] * [1 – (1 – \text{availability}_3)^2]$

In the example shown in the figure, a serial available network is available 99.86 percent of the time, while the parallel available network is available 99.97 percent of the time.

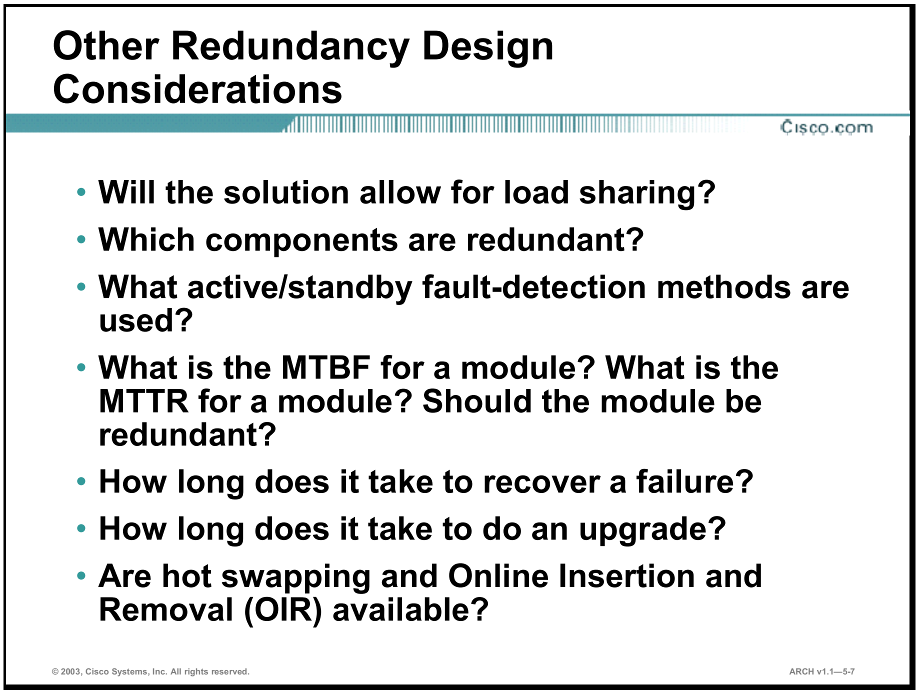

Other Redundancy Design Considerations

- Will the solution allow for load sharing?
- Which components are redundant?
- What active/standby fault-detection methods are used?
- What is the MTBF for a module? What is the MTTR for a module? Should the module be redundant?
- How long does it take to recover a failure?
- How long does it take to do an upgrade?
- Are hot swapping and Online Insertion and Removal (OIR) available?

ARCH v1.1—5-7

To fully determine the benefit of device, chassis, and link redundancy, you will ask the questions listed in the figure.

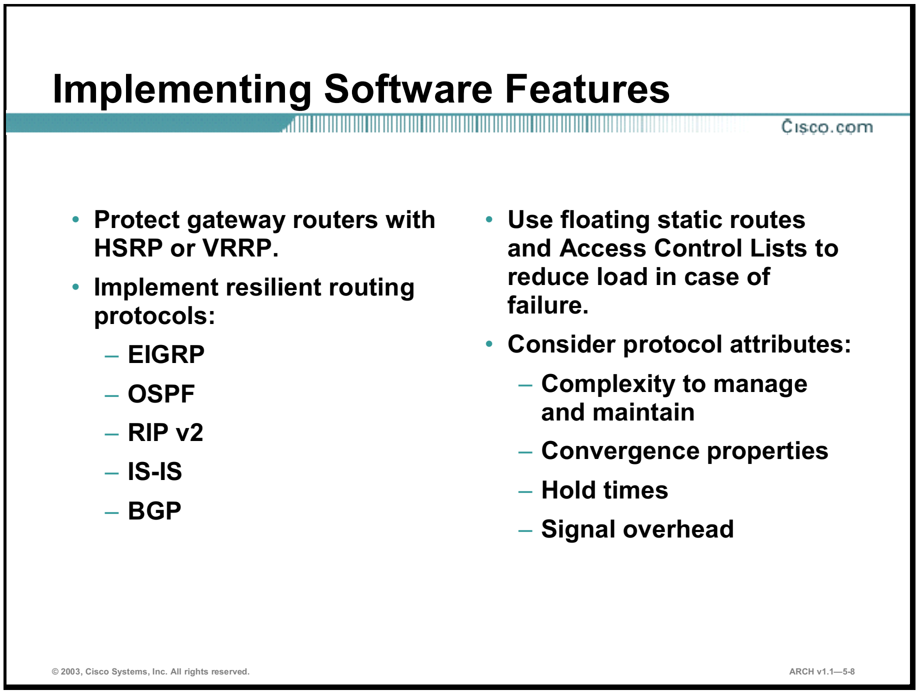

Implementing Software Features

- Protect gateway routers with HSRP or VRRP.
- Implement resilient routing protocols:
  - EIGRP
  - OSPF
  - RIP v2
  - IS-IS
  - BGP
- Use floating static routes and Access Control Lists to reduce load in case of failure.
- Consider protocol attributes:
  - Complexity to manage and maintain
  - Convergence properties
  - Hold times
  - Signal overhead

ARCH v1.1—5-8

Cisco recommends that you implement the software features listed in the figure.

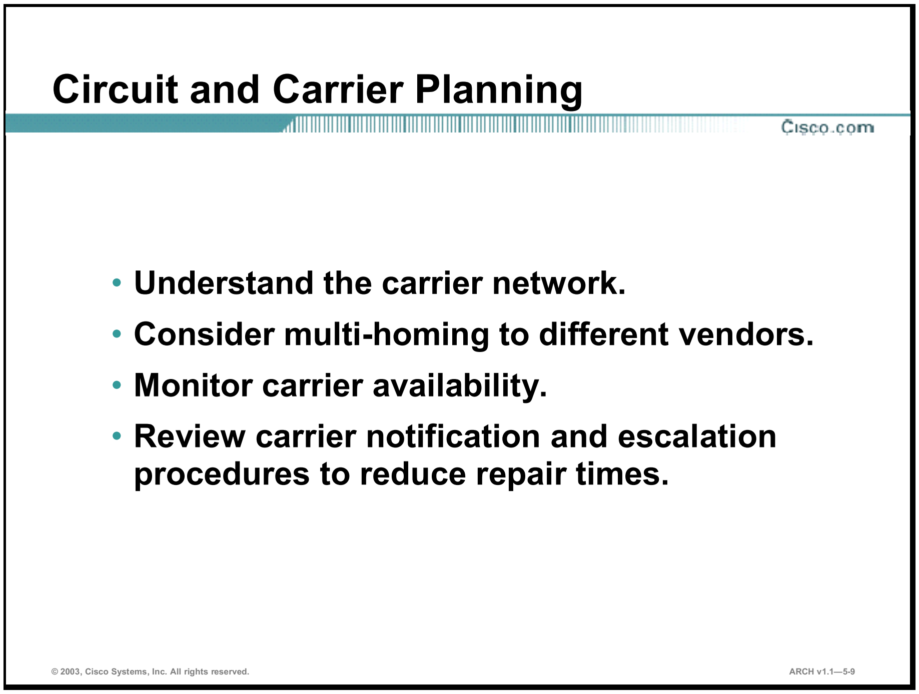

Circuit and Carrier Planning

- Understand the carrier network.
- Consider multi-homing to different vendors.
- Monitor carrier availability.
- Review carrier notification and escalation procedures to reduce repair times.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—5-9

Since the carrier network is an important component of the enterprise network and its availability, you should carefully consider these points about the carrier network in your design:

- **Understand the carrier network:** You should model and understand carrier availability, including the carrier diversity strategy and how that will affect the availability of your network design. Make sure you have a service level agreement that specifies availability and offers alternate routes in case of failure.
- **Consider multi-homing to different vendors:** Multi-homing to different vendors provides protection if one carrier goes down.
- **Monitor carrier availability:** Determine if the carrier offers enhanced services such as a guaranteed committed information rate (CIR) for Frame Relay or differentiated services. Use carrier service level agreements.
- **Review carrier notification and escalation procedures to reduce repair times:** Review the carrier's notification and escalation procedures to ensure that they can reduce down times.

Ensure that the carrier offers diversity. You want to ensure that dual paths to an ISP, for example, do not terminate at the same location (a single point of failure).

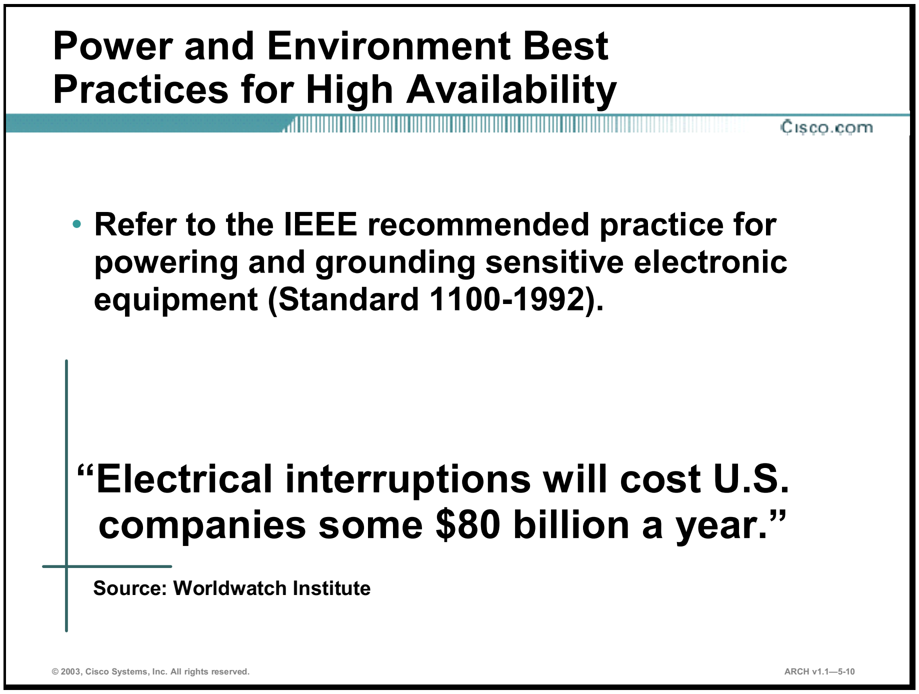

Power and Environment Best Practices for High Availability

- Refer to the IEEE recommended practice for powering and grounding sensitive electronic equipment (Standard 1100-1992).

> **"Electrical interruptions will cost U.S. companies some $80 billion a year."**

Source: Worldwatch Institute

ARCH v1.1—5-10

Power and environmental availability affect overall network availability. By implementing uninterruptible power supplies (UPS), you can increase availability. The table describes the effect of UPS and power array generators on overall availability.

	Raw AC Power	**5-Minute UPS**	**1-Hour UPS**	**UPS with Generator**	**Power Array with Generator**
Event outages	15 events	1 event	.15 event	.01 event	.001 event
Annual downtime	189 minutes	109 minutes	10 minutes	1 minute	.1 minute
Availability	99.96%	99.979%	99.998%	99.9998%	99.99999%

Source American Power Conversion, Tech Note #26.

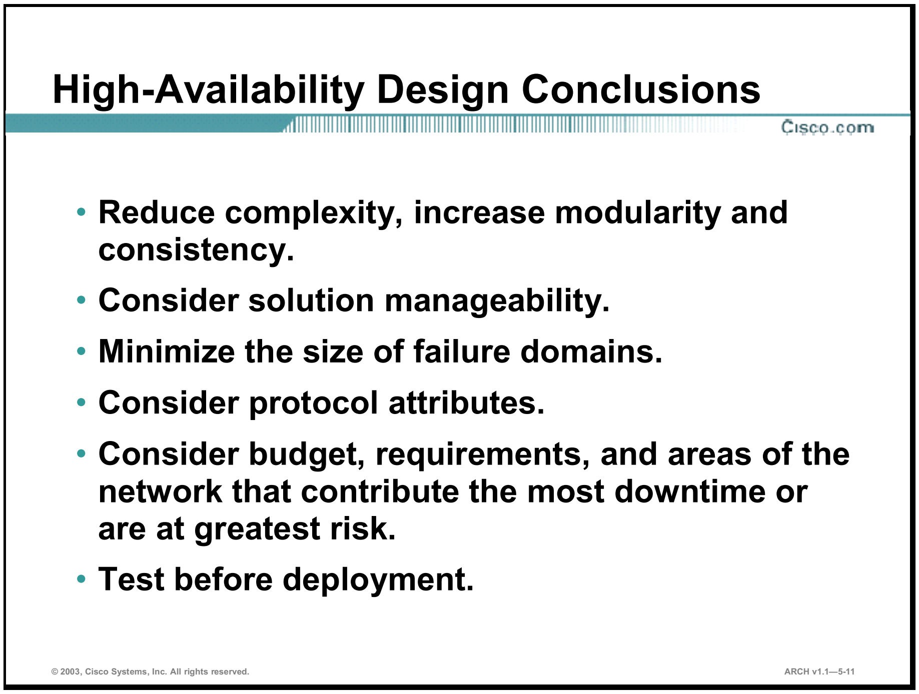

High-Availability Design Conclusions

- Reduce complexity, increase modularity and consistency.
- Consider solution manageability.
- Minimize the size of failure domains.
- Consider protocol attributes.
- Consider budget, requirements, and areas of the network that contribute the most downtime or are at greatest risk.
- Test before deployment.

ARCH v1.1—5-11

The figure lists overall design considerations as you approach your design for the network.

In terms of costs, consider:

- **One-time costs:** Calculate the cost of additional components or hardware, software upgrades, new software costs, and installation.
- **Recurring costs:** Consider the costs of additional WAN links.
- **Complexity costs:** Keep in mind that availability may be more difficult to manage and troubleshoot. More training may be required.

Best Practices for High-Availability Network Design

Cisco has developed a set of best practices recommendations to ensure high availability of the network. This topic identifies those best practices.

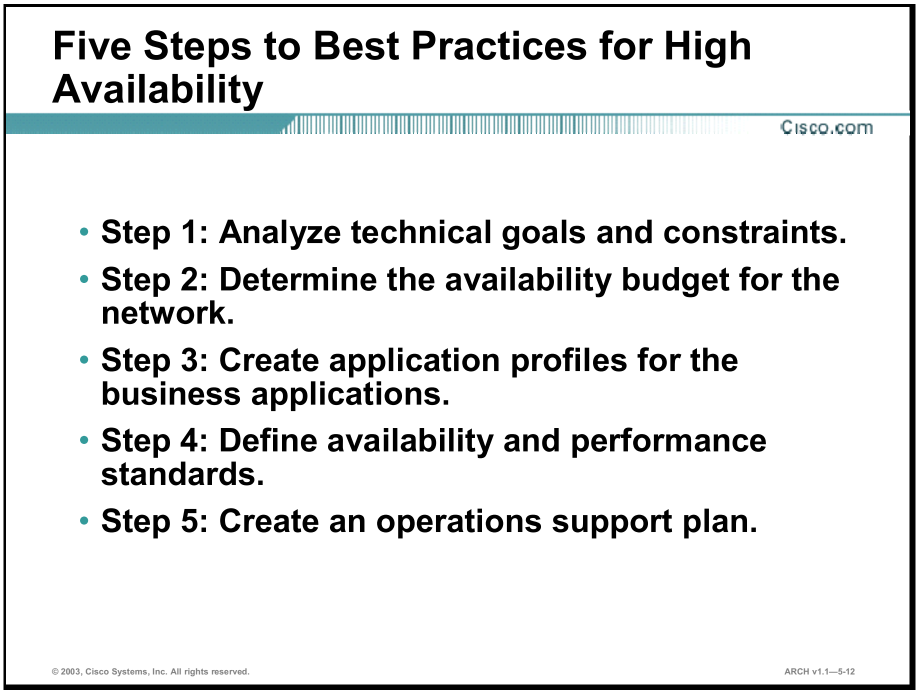

Five Steps to Best Practices for High Availability

Cisco.com

- Step 1: Analyze technical goals and constraints.
- Step 2: Determine the availability budget for the network.
- Step 3: Create application profiles for the business applications.
- Step 4: Define availability and performance standards.
- Step 5: Create an operations support plan.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—5-12

Complete these steps to implement a highly available network:

Step 1 **Analyze technical goals and constraints.** Technical goals include availability levels, throughput, jitter, delay, response time, scalability requirements, introductions of new features and applications, security, manageability, and cost. Investigate constraints, given the available resources. Prioritize goals and lower expectations that can still meet business requirements. Prioritize constraints in terms of the greatest risk or impact to the desired goal.

Step 2 **Determine the availability budget for the network.** Determine the expected theoretical availability of the network. Use this information to determine the availability of the system to help ensure the design will meet business requirements.

Step 3 **Create application profiles for business applications.** Application profiles help to align network service goals with application or business requirements by comparing application requirements, such as performance and availability, with realistic network service goals or current limitations.

Step 4 **Define availability and performance standards.** Availability and performance standards set the service expectations for the organization.

Step 5 **Create an operations support plan.** Define the reactive and proactive processes and procedures used to achieve the service level goal. Determine how the service process will be managed and measured. Each organization should know their role and responsibility for any given circumstance. The operations support plan should also include a plan for spare components.

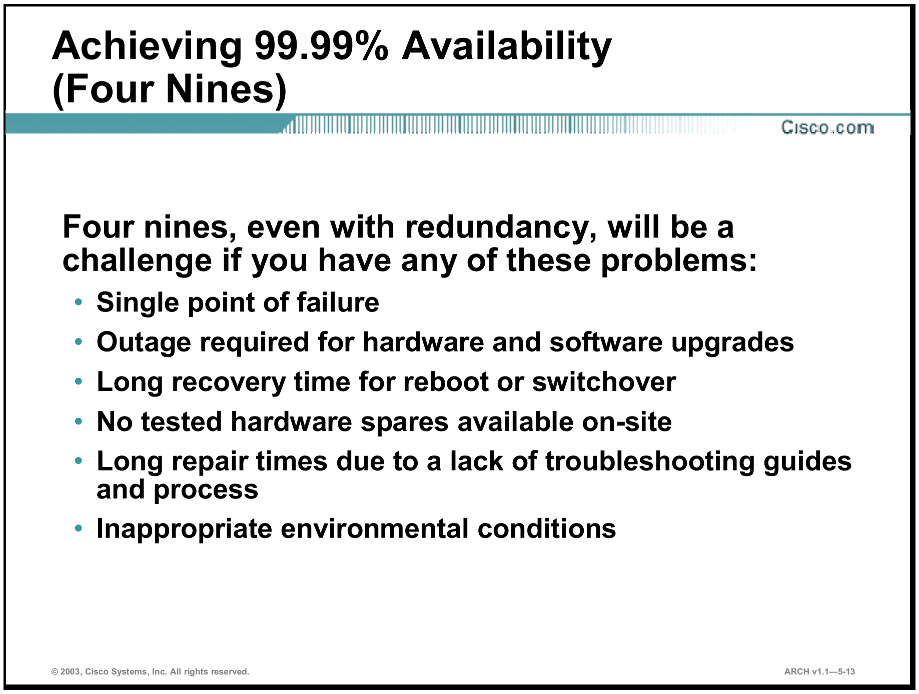

Achieving 99.99% Availability (Four Nines)

Four nines, even with redundancy, will be a challenge if you have any of these problems:

- Single point of failure
- Outage required for hardware and software upgrades
- Long recovery time for reboot or switchover
- No tested hardware spares available on-site
- Long repair times due to a lack of troubleshooting guides and process
- Inappropriate environmental conditions

ARCH v1.1—5-13

To achieve 99.99 percent availability (often referred to as “four nines”), you need to eliminate the problems listed in the figure.

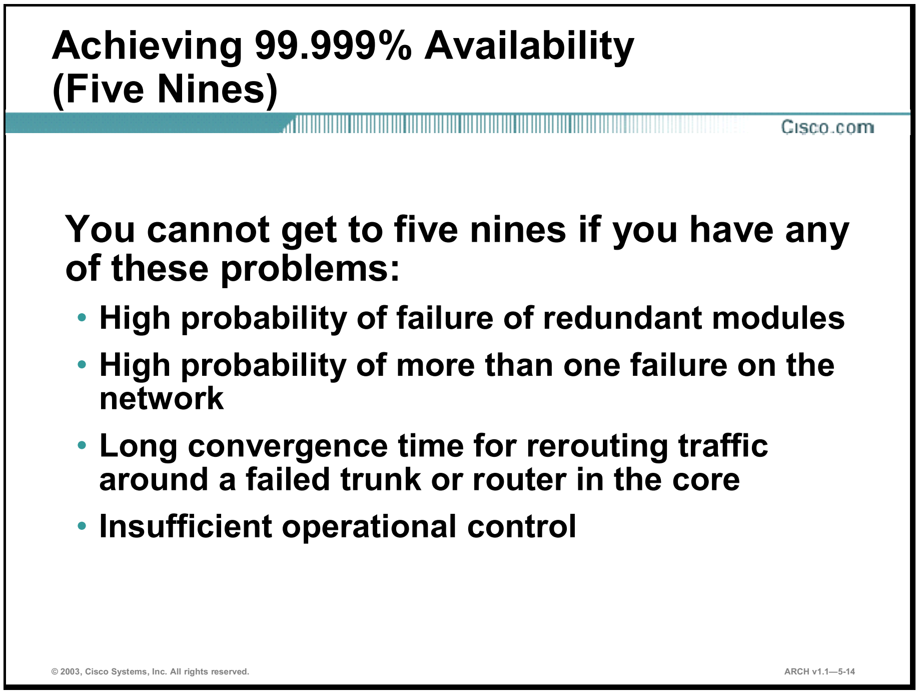

Achieving 99.999% Availability (Five Nines)

You cannot get to five nines if you have any of these problems:

- High probability of failure of redundant modules
- High probability of more than one failure on the network
- Long convergence time for rerouting traffic around a failed trunk or router in the core
- Insufficient operational control

ARCH v1.1—5-14

To achieve 99.999 percent availability (often referred to as “five nines”), you need to eliminate the problems listed in the figure.

Enterprise Campus Design Guidelines for High Availability

Each submodule of the Campus Infrastructure module should incorporate fault tolerance and redundancy features to provide an end-to-end highly available network. This topic describes the basic guidelines for designing a campus network for high availability.

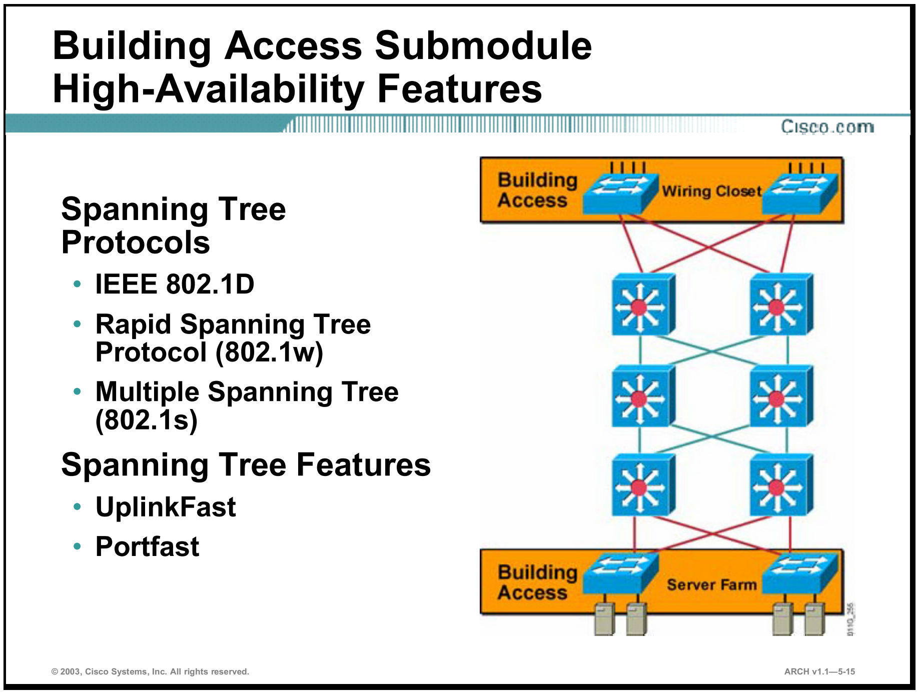

Building Access Submodule High-Availability Features

Spanning Tree Protocols

- IEEE 802.1D
- Rapid Spanning Tree Protocol (802.1w)
- Multiple Spanning Tree (802.1s)

Spanning Tree Features

- UplinkFast
- Portfast

ARCH v1.1—5-15

In the Building Access submodule, Cisco recommends that you implement STP along with the UplinkFast and PortFast enhancements.

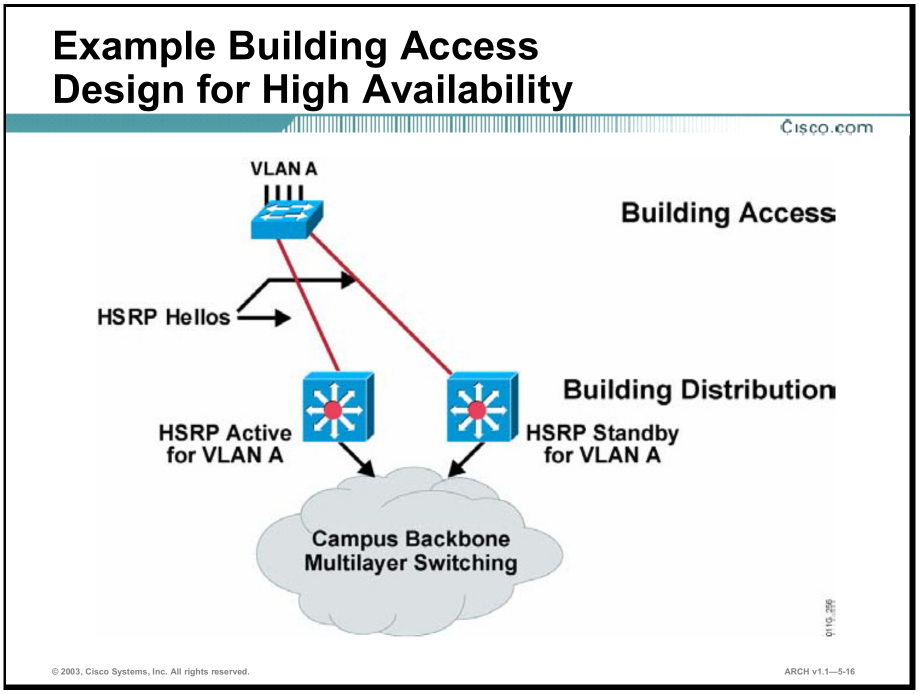

Example Building Access Design for High Availability

ARCH v1.1—5-16

You can implement HSRP in the Building Distribution submodule, with HSRP hellos to switches in the Building Access submodule.

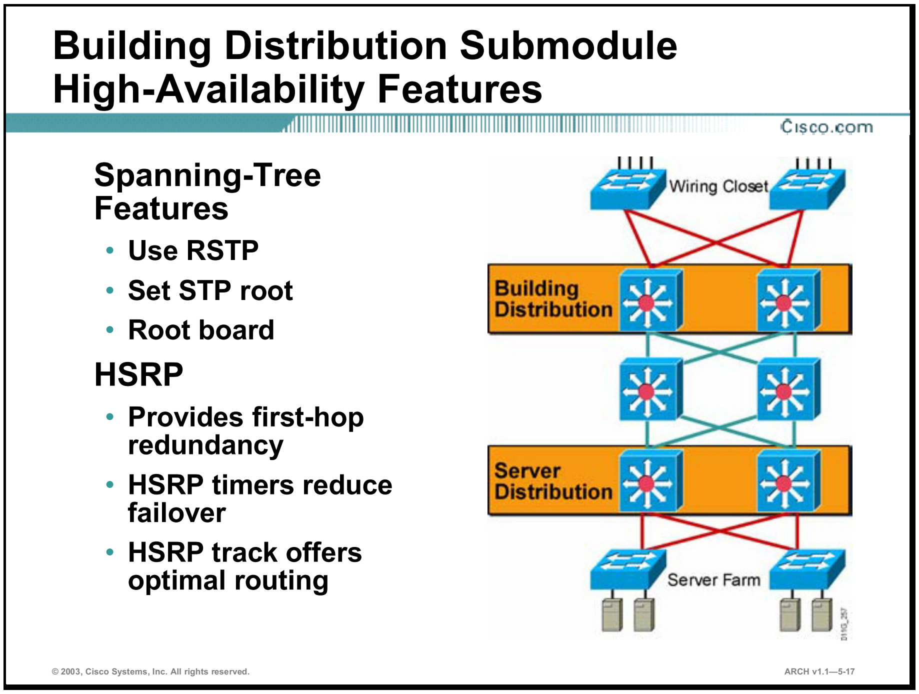

Building Distribution Submodule High-Availability Features

Spanning-Tree Features

- Use RSTP
- Set STP root
- Root board

HSRP

- Provides first-hop redundancy
- HSRP timers reduce failover
- HSRP track offers optimal routing

ARCH v1.1—5-17

At the Building Distribution submodule, Cisco recommends that you implement Spanning Tree Protocol as well as HSRP for first-hop redundancy.

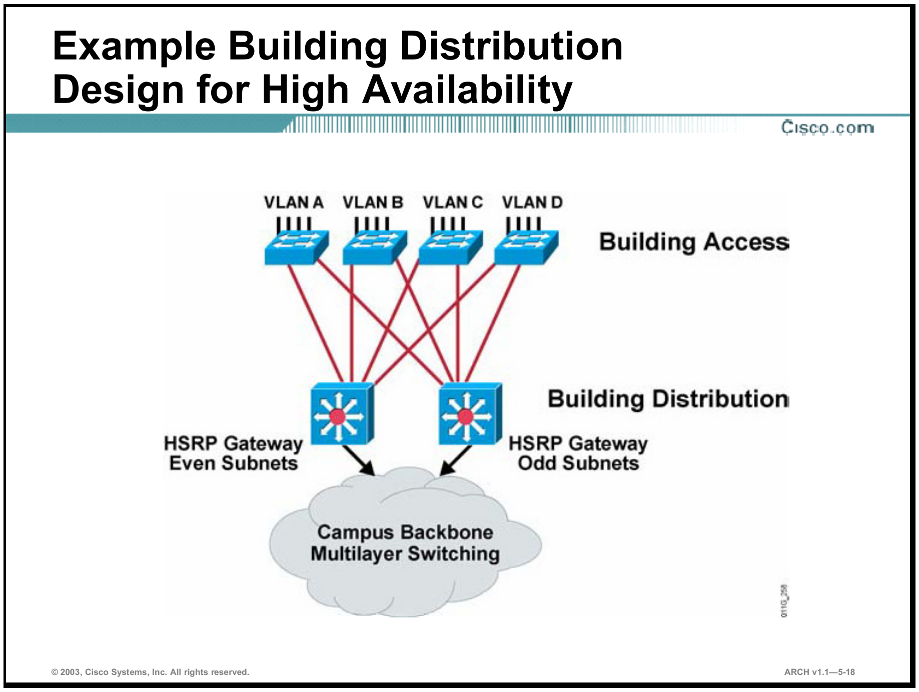

Example Building Distribution Design for High Availability

ARCH v1.1—5-18

The figure shows an example of a Building Distribution design that provides redundant links to each Building Access switch. HSRP provides failover redundancy in the Building Distribution submodule.

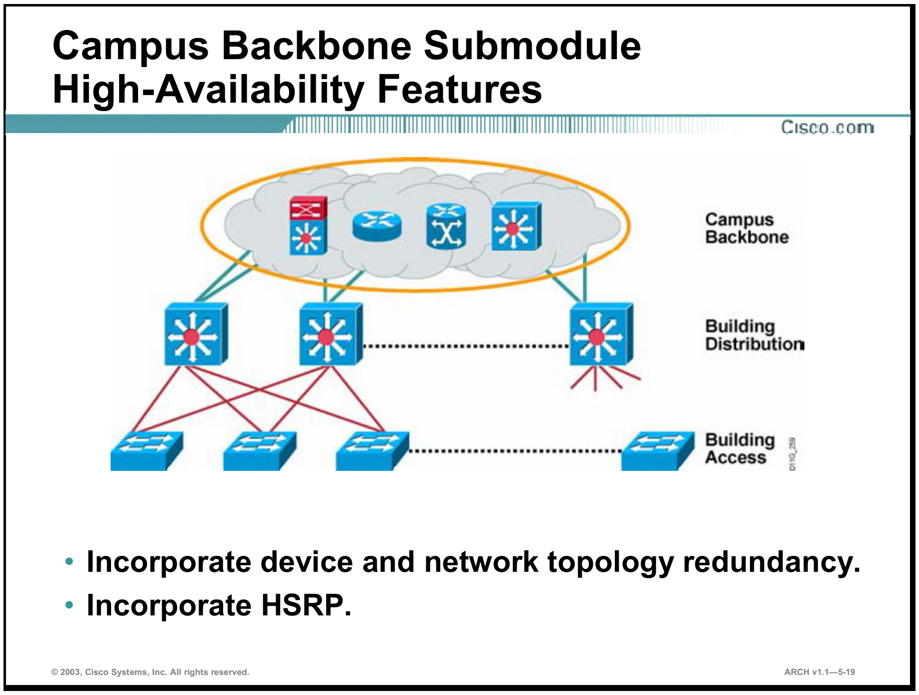

Campus Backbone Submodule High-Availability Features

- Incorporate device and network topology redundancy.
- Incorporate HSRP.

ARCH v1.1—5-19

The Campus Backbone submodule is a critical resource to the entire network. Cisco recommends that you incorporate device and network topology redundancy at the Campus Backbone, as well as HSRP for failover.

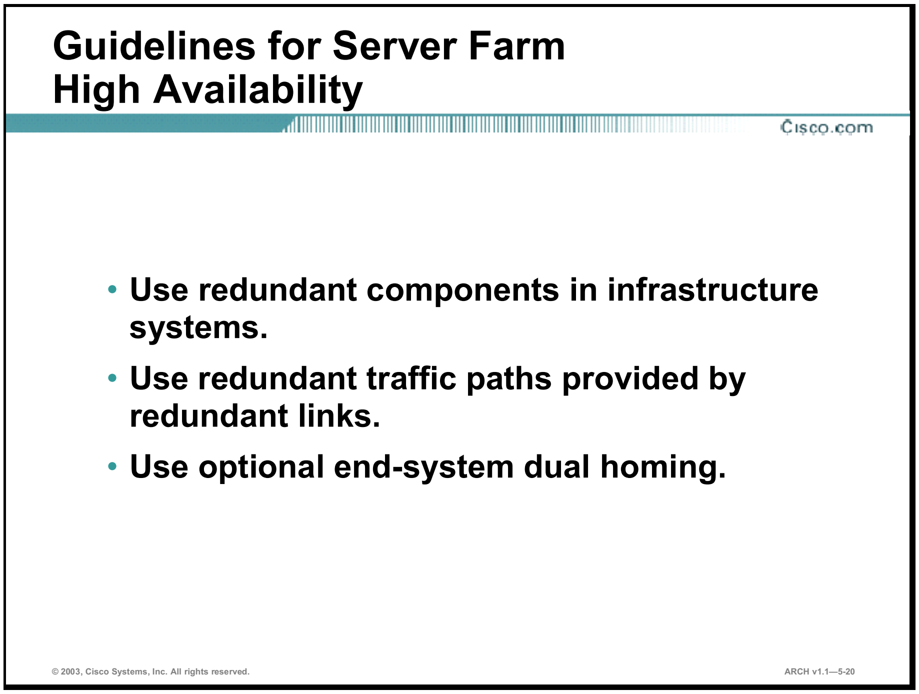

Guidelines for Server Farm High Availability

Cisco.com

- **Use redundant components in infrastructure systems.**
- **Use redundant traffic paths provided by redundant links.**
- **Use optional end-system dual homing.**

ARCH v1.1—5-20

The primary design objective for a server farm is to ensure high availability in the infrastructure architecture by implementing these features:

- Redundant components in infrastructure systems, where such a configuration is practical, cost effective, and considered optimal
- Redundant traffic paths provided by redundant links between infrastructure systems
- Optional end-system dual homing to provide a higher degree of availability

By leveraging the flexibility of data link layer connectivity in the Building Access switches, the option of dual-homing the connected end systems is available. Most NICs operate in an active-standby mode with a mechanism for MAC address portability between pairs. During a failure, the standby NIC becomes active on the new Building Access switch.

Another end-system redundancy option is for a NIC to operate in active-active mode, in which each host is available through multiple IP addresses. Either end-system redundancy mode requires more ports in the Building Access submodule.

Enterprise Edge Design Guidelines for High Availability

Each module of the Enterprise Edge functional area should incorporate high-availability features from the service provider edge to the enterprise campus network. This topic describes the basic guidelines for designing the Enterprise Edge for high availability.

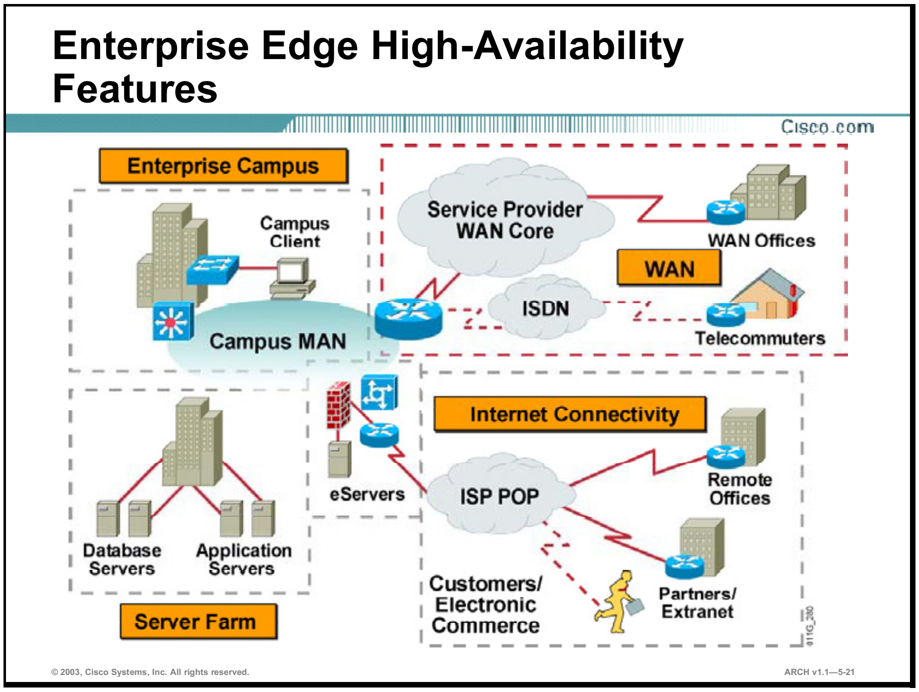

Enterprise Edge High-Availability Features

ARCH v1.1—5-21

Within the Enterprise Edge functional area, consider the following for high availability:

- **Service level agreement:** Ask your service provider to write into your service level agreement that your backup path terminates into separate equipment at the service provider, and that your lines are not trunked into the same paths as they traverse the network.
- **Link redundancy:** Use separate ports, preferably on separate routers, to each remote site. Having backup permanent virtual circuits (PVCs) through the same physical port accomplishes little or nothing, since a port is more likely to fail than any individual PVC.
- **Load balancing:** Load balancing occurs when a router has two (or more) equal cost paths to the same destination. EIGRP also allows unequal-cost load sharing. You can implement load sharing on a per-packet or per-destination basis. Load sharing provides redundancy, because it provides an alternate path if a router fails. OSPF will load share on equal-cost paths by default. EIGRP will load share on equal-cost paths by default, and can be configured to load share on unequal-cost paths. Unequal-cost load sharing is discouraged because it can create too many obscure timing problems and retransmissions.
- **Policy-based routing:** If you have unequal cost paths, and you do not want to use unequal-cost load sharing, you can use policy-based routing to send lower priority traffic down the slower path.
- **Routing protocol convergence:** The convergence time of the routing protocol chosen will affect overall availability of the Enterprise Edge. The main area to examine is the impact of the Layer 2 design on Layer 3 efficiency.

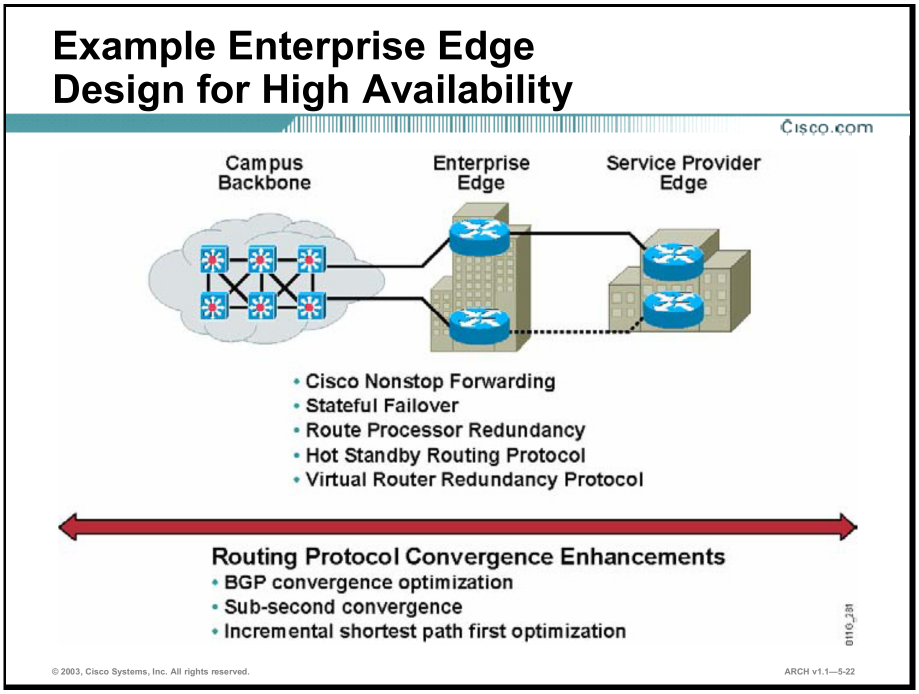

Example Enterprise Edge Design for High Availability

Cisco.com

- Cisco Nonstop Forwarding
- Stateful Failover
- Route Processor Redundancy
- Hot Standby Routing Protocol
- Virtual Router Redundancy Protocol

Routing Protocol Convergence Enhancements

- BGP convergence optimization
- Sub-second convergence
- Incremental shortest path first optimization

ARCH v1.1—5-22

The example implements these features at the Enterprise Edge functional area:

- **Cisco Nonstop Forwarding:** Enables continuous packet forwarding during route processor takeover and route convergence
- **Stateful Failover:** Allows a backup route processor to take immediate control from the active route processor while maintaining WAN connectivity protocols
- **Route Processor Redundancy:** Allows a standby route processor to load an IOS image configuration, parse the configuration, and reset and reload the line cards, thereby reducing reboot time
- **HSRP:** Enables two or more routers to work together in a group to emulate a single virtual router to the source hosts on the LAN
- **VRRP:** Enables a group of routers to form a single virtual router by sharing one virtual router IP address and one virtual MAC address

High-Availability Design Example

Providing high availability in the enterprise site can involve deploying highly fault-tolerant devices, incorporating redundant topologies, implementing Spanning Tree Protocol, and configuring HSRP. This topic describes a high availability strategy for an enterprise site.

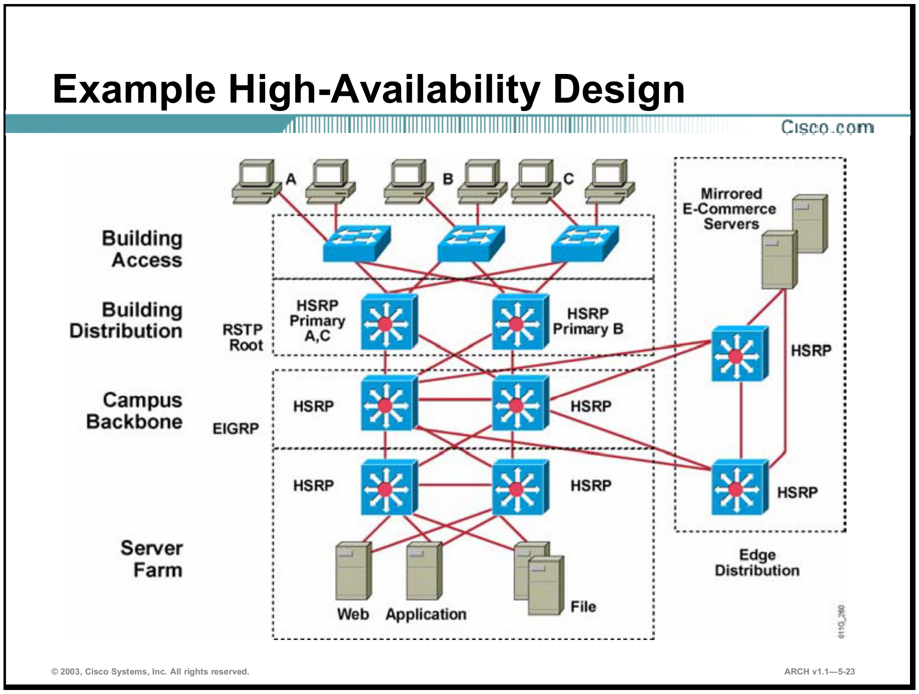

Example High-Availability Design

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—5-23

The figure shows an example enterprise site design that incorporates high-availability features.

- **Building Access submodule:** The Building Access switches all have uplinks terminating in a pair of redundant multilayer switches at the Building Distribution submodule, which act as an aggregation point. Only one pair of Building Distribution switches is needed per building. The number of wiring-closet switches was based on port density requirements. Each Building Access switch includes fault tolerance to reduce MTBF. Since the failure of an individual switch would have a smaller impact than a device failure in the Building Distribution and Campus Backbone submodules, device redundancy is not provided.
- **Building Distribution submodule:** First-hop redundancy and fast failure recovery is achieved with HSRP, which runs on the two switches in the distribution layer. HSRP provides end stations with a default gateway in the form of a virtual IP address that is shared by a minimum of two routers. HSRP routers discover each other via hello packets, which are sent through the Building Access switches with negligible latency.
- **Campus Backbone submodule:** In the Campus Backbone submodule layer, two multilayer switches are deployed, each configured for high fault tolerance. HSRP is implemented to allow for device redundancy. EIGRP is used to provide load balancing and fast convergence.
- **Server Farm module:** In the Server Farm module, two multilayer switches with HSRP configured provide redundancy. The file servers are mirrored for added protection.
- **Enterprise Edge module:** At the enterprise edge, fault-tolerant switches are deployed with link redundancy and HSRP to enable failover. Outward-facing e-commerce servers are mirrored to ensure availability.

Summary

This topic summarizes the key points discussed in this lesson.

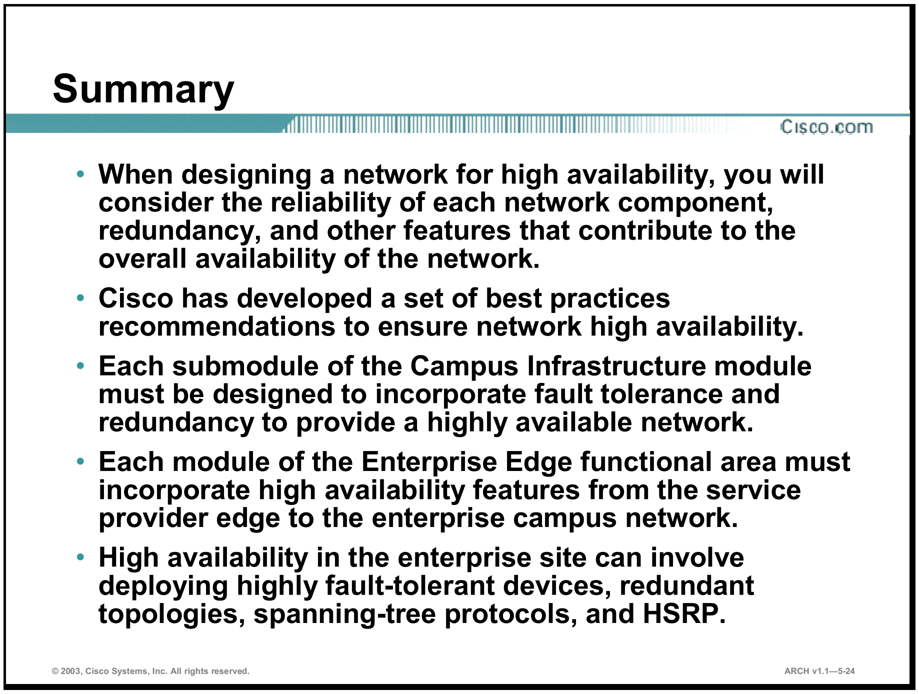

Summary

Cisco.com

- When designing a network for high availability, you will consider the reliability of each network component, redundancy, and other features that contribute to the overall availability of the network.
- Cisco has developed a set of best practices recommendations to ensure network high availability.
- Each submodule of the Campus Infrastructure module must be designed to incorporate fault tolerance and redundancy to provide a highly available network.
- Each module of the Enterprise Edge functional area must incorporate high availability features from the service provider edge to the enterprise campus network.
- High availability in the enterprise site can involve deploying highly fault-tolerant devices, redundant topologies, spanning-tree protocols, and HSRP.

ARCH v1.1—5-24

References

For additional information, refer to these resources:

- *High Availability Services* at http://www.cisco.com/warp/public/779/largeent/learn/technologies/availability.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 5-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 5-2

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 formula calculates the theoretical advantage gained by implementing redundant modules or chassis in the network?

A) Availability = MTBF/MTTR * $[1 – (1 – \text{availability}_1)^2]$

B) Availability = $1 – [(1 – \text{availability}_1) * (1 – \text{availability}_2)]$

C) Availability = $[1 – (1 – \text{availability}_1)^2] * [1 – (1 – \text{availability}_2)^2]$

D) Availability = $[1 – (1 – \text{availability}_1)^2] * [1 – (1 – \text{availability}_2)^2] * [1 – (1 – \text{availability}_3)^2]$

Q2) What is one task you should perform to determine if the high-availability design meets the business requirements of the organization?

A) determine the IP address strategy

B) determine the routing protocol to use

C) determine the amount of expected downtime

D) determine the availability budget of the network

Q3) What is the first step to complete when developing an operations plan for high availability?

A) Create an operations support plan.

B) Analyze technical goals and constraints.

C) Define availability and performance standards.

D) Determine the availability budget for the network.

Q4) To design a network for five nines (99.999% availability), which three problems must you eliminate from the network? (Choose three.)

A) moderate recovery times

B) network operations planning

C) high probability of double failures

D) high probability of redundancy failure

E) long convergence times for rerouting traffic around a failed trunk or router

Q5) Which two options will improve availability within the Server Farm module? (Choose two.)

A) IP multicast

B) OSPF routing protocol

C) quality of service features

D) redundant NICs for servers

E) redundant traffic paths provided by redundant links between infrastructure systems

Q6) Which availability component is affected by network-level redundancy?

A) device costs

B) administrative complexity

C) mean time between failures

D) average duration of service interruptions

Q7) Ask your service provider to write into your _____ that your backup path terminates into separate equipment at the service provider.

A) network diagram

B) service level agreement

C) committed information rate

D) routing protocol redundancy plan

Q8) In a typical enterprise network, the Building Access data link layer switches all have uplinks terminating in a pair of _____ multilayer switches at the Building Distribution submodule.

A) secure

B) scalable

C) redundant

D) manageable

Q9) In planning for a new Campus Backbone design, you determined the need for rapid convergence and load balancing to meet the business requirements of your organization. Which protocol will best satisfy the requirement?

A) BGP

B) HSRP

C) RSTP

D) EIGRP

Quiz Answer Key

Q1) B

Relates to: Design Guidelines for High Availability

Q2) D

Relates to: Design Guidelines for High Availability

Q3) B

Relates to: Best Practices for High-Availability Network Design

Q4) C, D, E

Relates to: Best Practices for High-Availability Network Design

Q5) D, E

Relates to: Enterprise Campus Design Guidelines for High Availability

Q6) D

Relates to: Enterprise Campus Design Guidelines for High Availability

Q7) B

Relates to: Enterprise Edge Design Guidelines for High Availability

Q8) C

Relates to: High-Availability Design Example

Q9) D

Relates to: High-Availability Design Example

Case Study 5-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

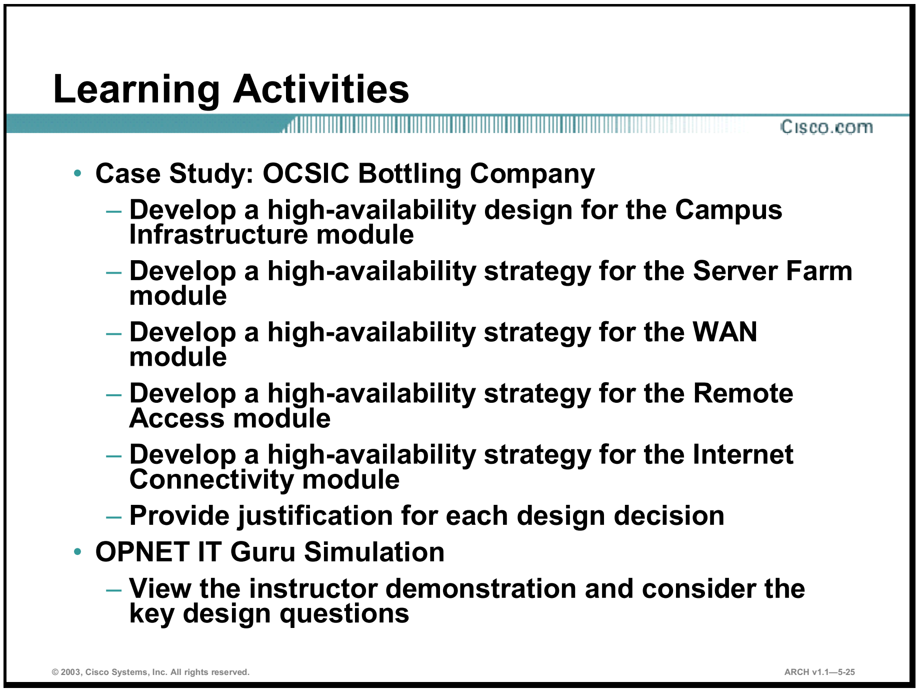

Learning Activities

- Case Study: OCSIC Bottling Company
  - Develop a high-availability design for the Campus Infrastructure module
  - Develop a high-availability strategy for the Server Farm module
  - Develop a high-availability strategy for the WAN module
  - Develop a high-availability strategy for the Remote Access module
  - Develop a high-availability strategy for the Internet Connectivity module
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

ARCH v1.1—5-25

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

OCSIC Bottling Company wants to revisit the design and ensure that the Campus Infrastructure, Server Farm, WAN, Remote Access, and Internet Connectivity modules are highly available.

In this exercise, you will design high-availability services that meet the needs of the OCSIC Bottling Company. After completing this exercise, you will be able to:

- Develop a high-availability design for the headquarters campus network
- Develop a high-availability strategy for the Server Farm module
- Develop a high-availability strategy for the WAN module
- Develop a high-availability strategy for the Remote Access module
- Develop a high-availability strategy for the Internet Connectivity module

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Note	You can complete all tasks within your group, or complete the task assigned to your group by the instructor.

Task 1: Develop a High-Availability Strategy for the Headquarters Campus Network

Complete these steps:

Step 1 Complete the table to design a high availability solution for the headquarters campus network.

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Task 2: Develop a High-Availability Strategy for the Server Farm Module

Complete these steps:

Step 1 Complete the table to design a high availability solution for the Server Farm module.

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Task 3: Develop a High-Availability Strategy for the WAN Module

Complete these steps:

Step 1 Complete the table to design a high availability solution for the WAN module.

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Task 4: Develop a High-Availability Strategy for the Remote Access Module

Complete these steps:

Step 1 Complete the table to design a high availability solution for the Remote Access module.

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Task 5: Develop a High-Availability Strategy for the Internet Connectivity Module

Complete these steps:

Step 1 Complete the table to design a high availability solution for the Internet Connectivity module.

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Step 2 Update your network diagrams to reflect your high-availability strategy.

Task 6: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have a high-availability design for the headquarters campus network with a completed network diagram.
- You have a high-availability design for the Server Farm module with a completed network diagram.
- You have a high--availability design for the WAN module with a completed network diagram.
- You have a high-availability design for the Remote Access module with a completed network diagram.
- You have a high-availability design for the Internet Connectivity module with a completed network diagram.

OPNET IT Guru Simulation 5-2

This simulation demonstrates the affect of failures on the network and application performance. Specifically:

- The first simulation demonstrates the effect of a device failure on application response times.
- The second and third simulations demonstrate the effect of link failures on application response times.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 6

Designing Security Services

Overview

As enterprises continually expand their mission-critical networks with new intranet, extranet, and e-commerce applications, network security is increasingly vital to prevent corruption and intrusion, and eliminate network security vulnerabilities. Without precautions, enterprises could experience major security breaches, resulting in serious damages or loss.

A key component of Cisco Architecture for Voice, Video and Integrated Data (AVVID), network security services improve the network's ability to support mission-critical Internet applications while providing authentication, authorization, and data integrity.

Module Objectives

Upon completing this module, you will be able to design security-intelligent network services for performance, scalability, and availability, given specified enterprise network needs.

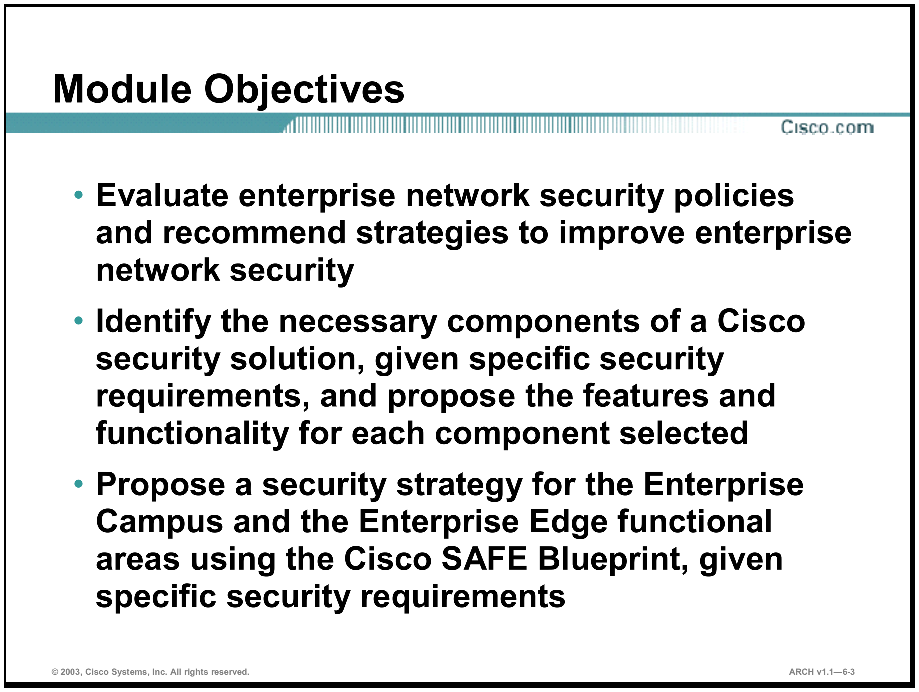

Module Objectives

- **Evaluate enterprise network security policies and recommend strategies to improve enterprise network security**
- **Identify the necessary components of a Cisco security solution, given specific security requirements, and propose the features and functionality for each component selected**
- **Propose a security strategy for the Enterprise Campus and the Enterprise Edge functional areas using the Cisco SAFE Blueprint, given specific security requirements**

Module Outline

The outline lists the components of this module.

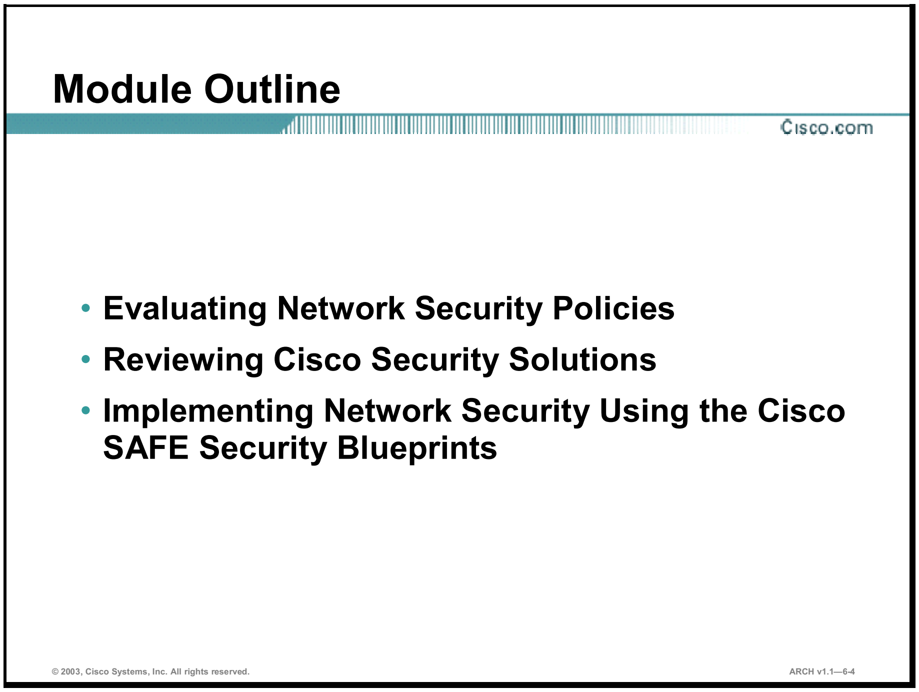

Module Outline

- Evaluating Network Security Policies
- Reviewing Cisco Security Solutions
- Implementing Network Security Using the Cisco SAFE Security Blueprints

ARCH v1.1—6-4

Evaluating Network Security Policies

Overview

Network security policies are critical to an overall security architecture for an enterprise. The three main phases of developing network security policies are establishing a security policy, implementing network security technologies, and auditing the network on a recurring basis. You will use the results of the audits to modify the security policy and the technology implementation as needed.

The enterprise security strategy includes establishing a security policy that defines the security goals of the enterprise, and implementing network security technologies in a comprehensive and layered approach so that the enterprise does not rely upon only one type of technology to solve all security issues.

Relevance

Maintaining a high level of network security requires a continuous effort. Evaluating network security on an ongoing basis is critical to maintaining the most effective security.

Objectives

Upon completing this lesson, you will be able to evaluate enterprise network security policies and recommend strategies to improve enterprise network security. This includes being able to meet these objectives:

- Describe the primary network vulnerabilities and their countermeasures
- Describe the purpose and components of a security policy
- Describe the process of maintaining network security
- Explain how to assess an existing network's risk from network attacks

Learner Skills and Knowledge

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

- Successful completion of *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

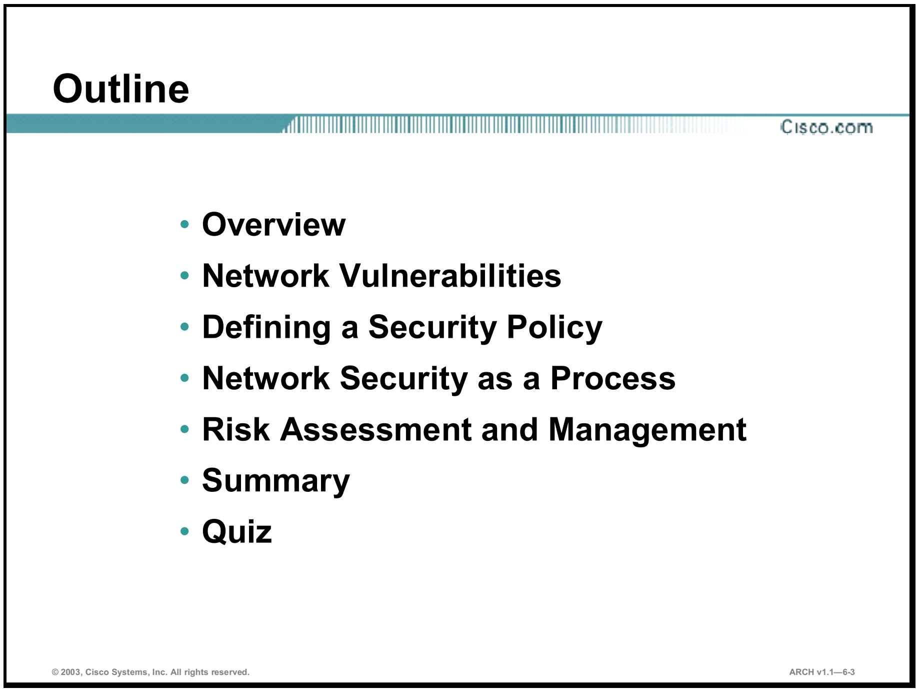

Outline

- Overview
- Network Vulnerabilities
- Defining a Security Policy
- Network Security as a Process
- Risk Assessment and Management
- Summary
- Quiz

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-3

Network Vulnerabilities

As time has passed, the sophistication of tools used to attack a network has increased while the technical knowledge needed to use those tools has decreased. Networks are vulnerable to a variety of threats, classified as loss of privacy, data theft, impersonation, and loss of integrity. This topic describes the primary network vulnerabilities and their countermeasures.

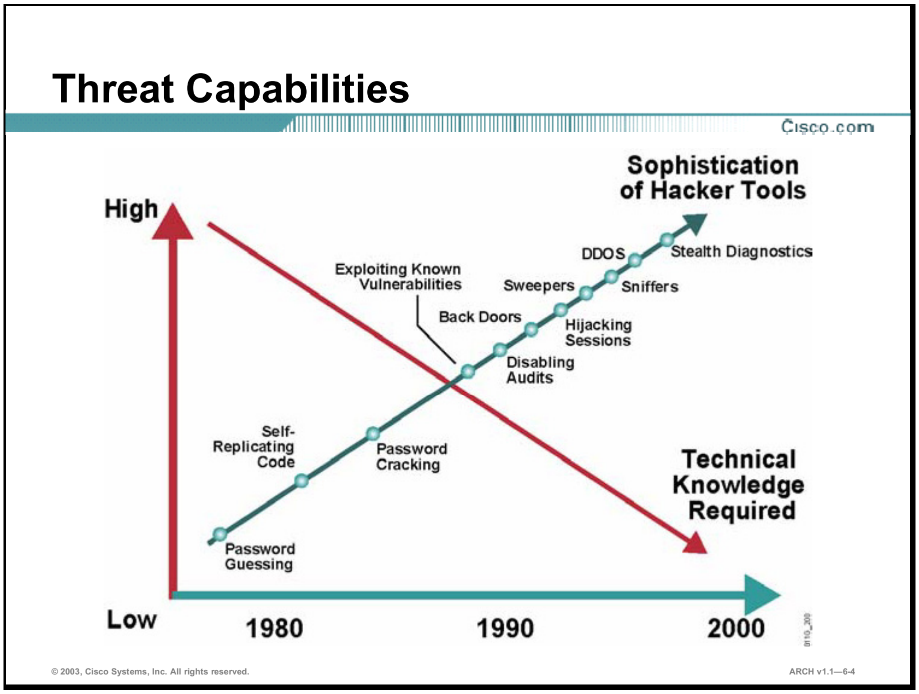

Threat Capabilities

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-4

Many people have the impression that Internet hacking incidents are highly complex, technical attacks that take a genius to create. The reality is that a few sophisticated people develop these highly complex, technical attacks, but they then use the Internet to share the information and the tools required to execute the attack. The open sharing of hacking information and tools allows individuals with minimal technical knowledge to duplicate an attack. Often, it is as easy as downloading the attack tool from the Internet and launching it against targets. A hacker need not know anything other than how to run the attack tool.

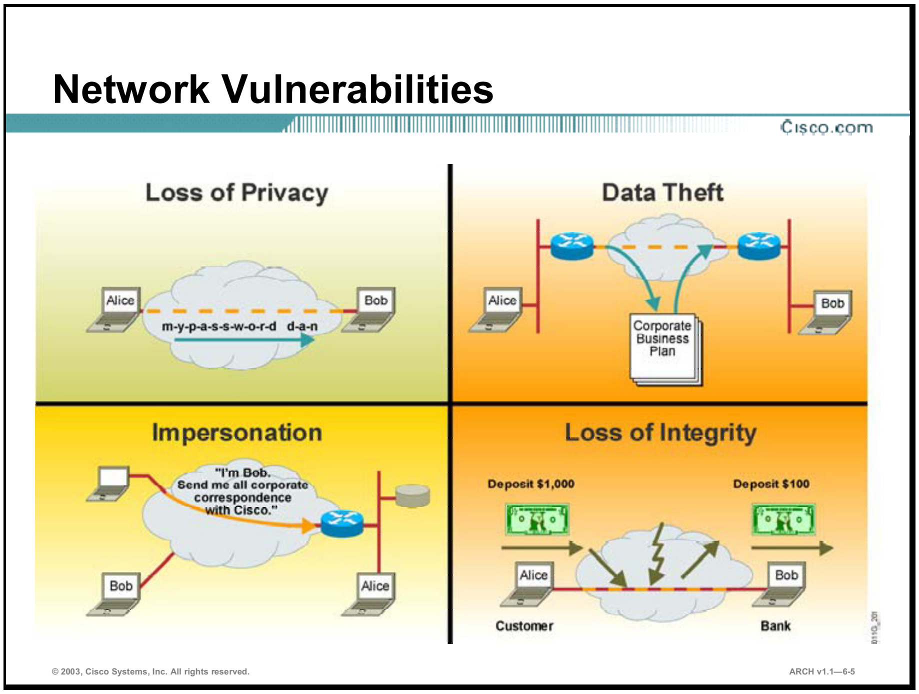

Network Vulnerabilities

Cisco.com

ARCH v1.1—6-5

Total data security assurance results from a comprehensive strategy that addresses each type of network vulnerability.

To counteract the problems of loss of privacy and data theft, where data is accessed or even removed, security protocols provide confidentiality for sensitive information as it travels across an untrusted or public network. Protocols that provide confidentiality typically employ encryption techniques that scramble data in a way that is undecipherable to unauthorized access attempts.

To counteract problems associated with impersonation, authentication protocols both validate and guarantee the identity of communicating parties. Authentication protocols are implemented in many ways, but most commonly take the form of digital signatures, digital certificates, or shared keys.

To counteract the problem of loss of integrity, where an external entity may not be able to see the data content but still alter it, security protocols validate the integrity of information traveling across an untrusted or public network. Such protocols are typically hashing algorithms that generate a value unique to the data content. Hashing algorithms do not prevent alteration of data, but rather allow communicating parties to detect when alteration occurs.

Effective data security assurance, from a protocol perspective, requires methods for ensuring data confidentiality, integrity, and authentication.

Defining a Security Policy

Network security efforts are based on a security policy. The policy should identify what is being protected, how users are identified and trusted, how the policy is to be enforced, the consequences of a violation, and the response to a violation. This topic describes the purpose and components of a security policy.

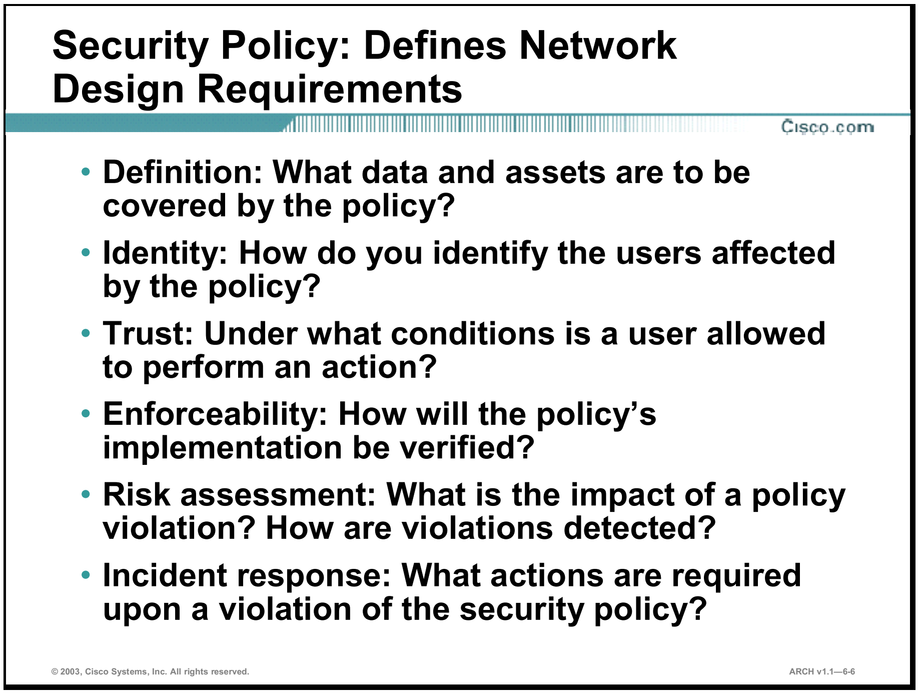

Security Policy: Defines Network Design Requirements

- **Definition: What data and assets are to be covered by the policy?**
- **Identity: How do you identify the users affected by the policy?**
- **Trust: Under what conditions is a user allowed to perform an action?**
- **Enforceability: How will the policy's implementation be verified?**
- **Risk assessment: What is the impact of a policy violation? How are violations detected?**
- **Incident response: What actions are required upon a violation of the security policy?**

ARCH v1.1—6-6

As stated in RFC 2196, a security policy "is a formal statement of the rules by which people who are given access to an organization's technology and information assets must abide." Security policies apply to all aspects of running an organization, including building, maintaining, and using the network. Maintaining a high level of network security requires a continuous cycle of efforts based on a security policy. A security policy should contain these elements:

- **Definition:** What data and assets are to be covered by the policy?
- **Identity:** How do you identify the users (including hosts and applications) affected by the policy?
- **Trust:** Under what conditions is a user trusted to perform an action?
- **Enforceability:** How will the policy's implementation be verified?
- **Risk assessment:** What is the impact of a policy violation? How are violations detected?
- **Incident response:** What actions are required upon a violation of the security policy?

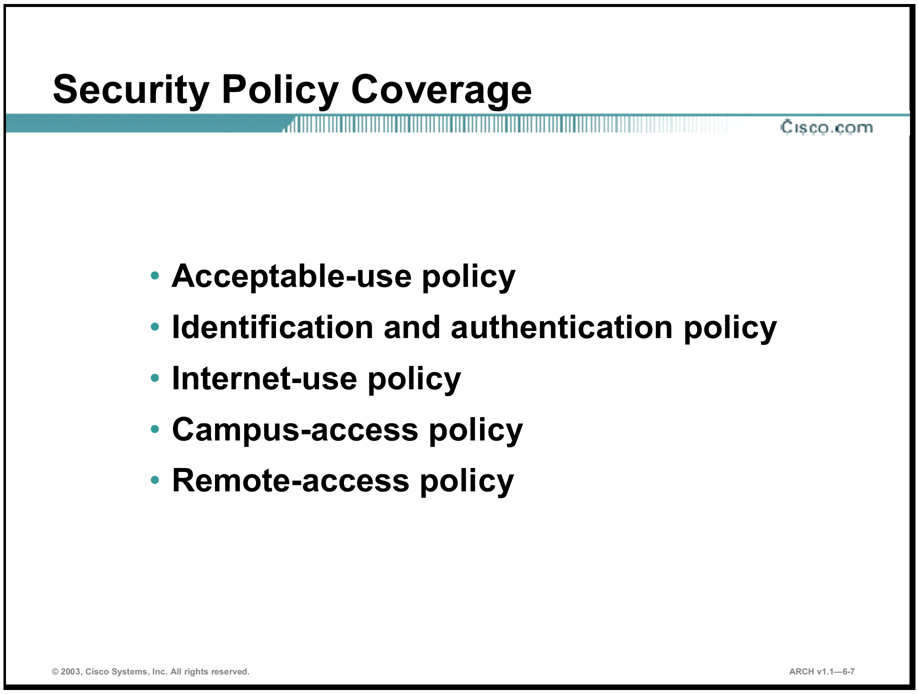

Security Policy Coverage

- Acceptable-use policy
- Identification and authentication policy
- Internet-use policy
- Campus-access policy
- Remote-access policy

ARCH v1.1—6-7

Network security policies typically define these situations, at a minimum:

- **Acceptable-use policy:** What constitutes acceptable and appropriate use of the network? What uses are not allowed? How does the policy differ for users, partners, and administrators?
- **Identification and authentication policy:** What standards and methodologies are used to identify and authenticate network users?
- **Internet-use policy:** What is the policy regarding the purposes for which users are allowed to access the Internet? Are any specific uses identified?
- **Campus-access policy:** Under what conditions are users allowed to access the campus network internally?
- **Remote-access policy:** What is the policy for users accessing the network from a remote location?

Network Security as a Process

Maintaining network security is based on a security policy. The ongoing steps include securing the network, monitoring network security, testing security, and improving security. This topic describes the process of maintaining network security.

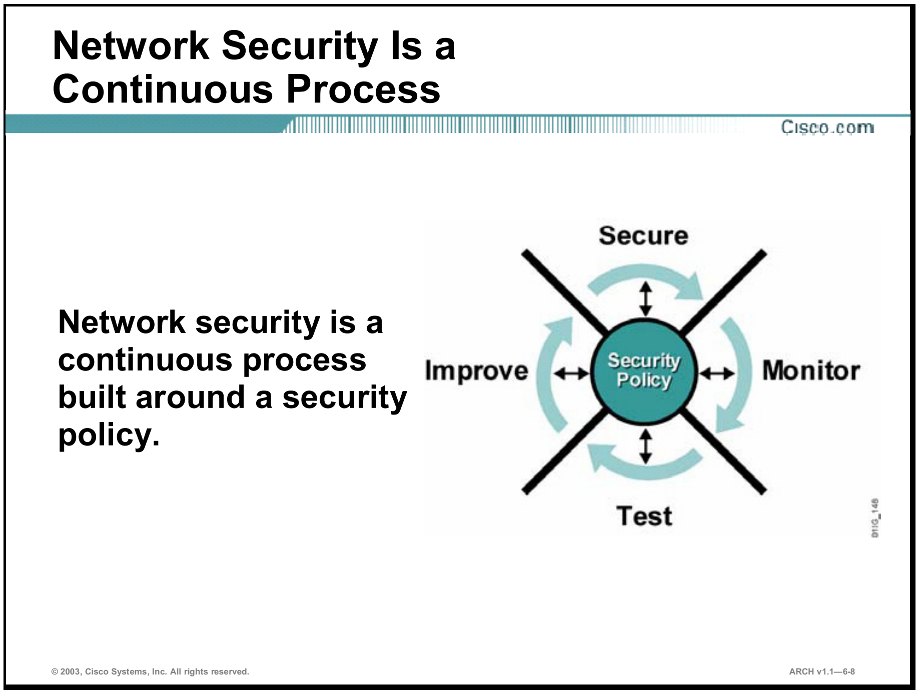

Network Security Is a Continuous Process

Cisco.com

Network security is a continuous process built around a security policy.

ARCH v1.1—6-8

After setting appropriate policies, a company or organization must methodically consider security as part of normal network operations. The policy could be as simple as configuring routers to reject unauthorized addresses or services, or as complex as installing firewalls, intrusion detection systems, centralized authentication servers, and encrypted Virtual Private Networks (VPNs).

After developing a security policy, secure your network using a variety of point products (firewalls, intrusion detection, and so on). Before you can secure your network, however, you need to consider your users, the assets that require protection, and the network's topology.

Securing the Network

The four solutions that you can implement to secure a network are as follows:

- **Authentication:** The recognition of each individual user, and mapping of their identity, location, and use policy, plus authorization of their network services.
- **Encryption:** A method for ensuring the confidentiality, integrity, and authenticity of data communications across a network. Cisco's solution combines several standards, including the Data Encryption Standard (DES) and Triple DES (3DES).
- **Firewalls:** A firewall is a set of related programs, located at a network gateway server, which protects the resources of a private network from users from other networks. (The term also implies the security policy that is used with the programs.)
- **Vulnerability patching:** Identifying and patching possible security holes that could compromise a network.

Monitoring Security

To ensure that a network remains secure, you should monitor the state of security preparation. Intrusion detection systems can monitor and respond to security events as they occur. Using security monitoring solutions, organizations can obtain visibility into both the network data stream and the security posture of the network.

Testing Security

Testing security is as important as monitoring. Testing the policy allows you to verify the effectiveness of the policy or identify weaknesses. Network vulnerability scanners can proactively identify areas of weakness.

Improving Security

Monitoring and testing provide the data necessary to improve network security. With the information gathered during monitoring and testing, you can improve the security implementation to better enforce the security policy and modify the policy to incorporate responses to new risks.

Risk Assessment and Management

A risk assessment identifies risks to your network, network resources, and data. The information gathered during a risk assessment aids in assessing the validity of a network security implementation and should be performed periodically. This topic describes how to assess an existing network's risk from network attacks.

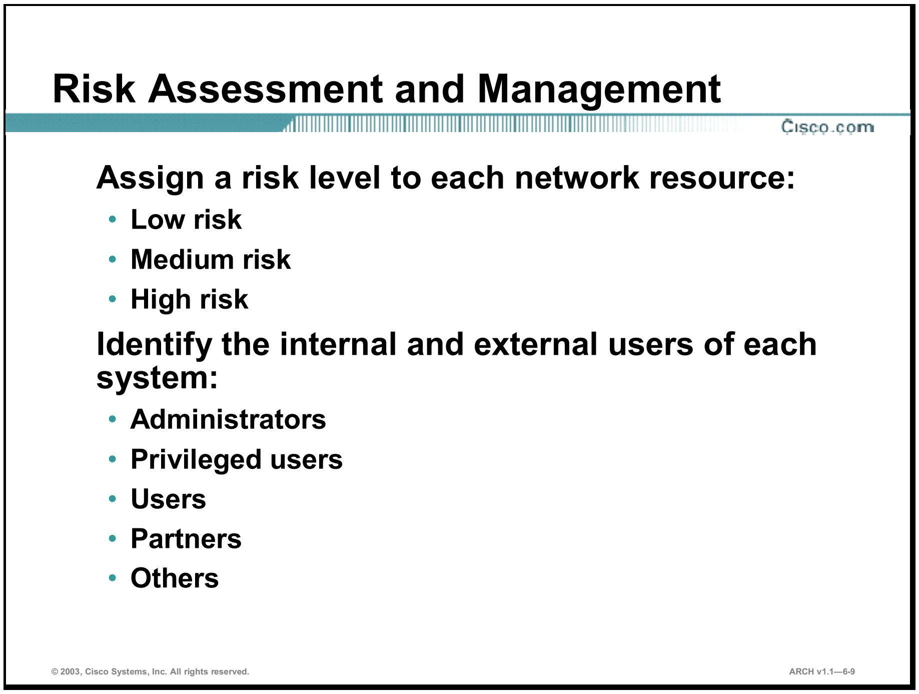

Risk Assessment and Management

Assign a risk level to each network resource:

- **Low risk**
- **Medium risk**
- **High risk**

Identify the internal and external users of each system:

- **Administrators**
- **Privileged users**
- **Users**
- **Partners**
- **Others**

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-9

A risk analysis should identify the risks to your network, network resources, and data. This doesn't mean you should identify every possible entry point to the network, nor every possible means of attack. The intent of a risk analysis is to identify portions of your network, assign a threat rating to each portion, and apply an appropriate level of security. This helps maintain a workable balance between security and required network access. The risk assessment should be carried out in conjunction with the established security policy.

Assign each network resource one of these three risk levels:

- **Low-risk systems or data:** If compromised (data viewed by unauthorized personnel, data corrupted, or data lost), these systems would not disrupt the business or cause legal or financial ramifications. The targeted system or data can be easily restored and does not permit further access of other systems.
- **Medium-risk systems or data:** If compromised (data viewed by unauthorized personnel, data corrupted, or data lost), these systems would cause a moderate disruption in the business, minor legal or financial ramifications, or provide further access to other systems. The targeted system or data requires a moderate effort to restore, or the restoration process is disruptive to the system.
- **High-risk systems or data:** If compromised (data viewed by unauthorized personnel, data corrupted, or data lost), these systems would cause an extreme disruption in the business, cause major legal or financial ramifications, or threaten the health and safety of a person. The targeted system or data requires significant effort to restore, or the restoration process is disruptive to the business or other systems.

Assign a risk level to each of the following: core network devices, distribution network devices, access network devices, network monitoring devices (Simple Network Management Protocol [SNMP] monitors and Remote Monitoring [RMON] probes), network security devices, e-mail systems, network file servers, network print servers, network application servers, data application servers, desktop computers, and other devices.

Network equipment such as switches, routers, Domain Name System (DNS) servers, and Dynamic Host Configuration Protocol (DHCP) servers can allow further access into the network, and are therefore either medium- or high-risk devices. It is also possible that corruption of this equipment could cause the network itself to collapse. Such a failure can be extremely disruptive to the business.

Once you assign a risk level, you will identify the types of users of that system. The five most common types of users are as follows, although you might consider including nonhuman users, such as applications:

- **Administrators:** Internal users responsible for network resources
- **Privileged:** Internal users with a need for greater access
- **Users:** Internal users with general access
- **Partners:** External users with a need to access some resources
- **Others:** External users or customers

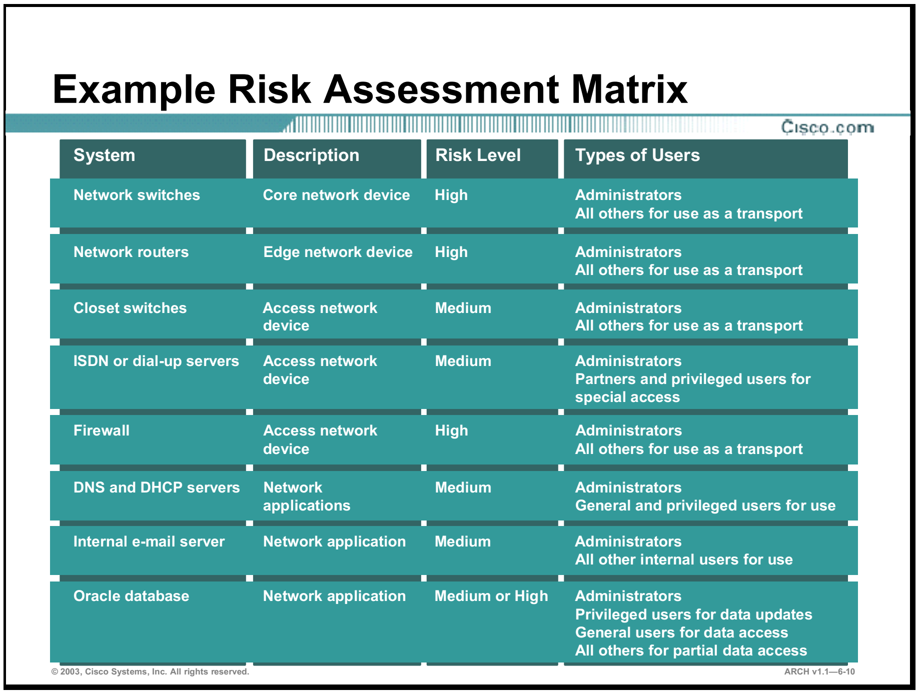

Example Risk Assessment Matrix

System	Description	Risk Level	Types of Users
Network switches	Core network device	High	Administrators All others for use as a transport
Network routers	Edge network device	High	Administrators All others for use as a transport
Closet switches	Access network device	Medium	Administrators All others for use as a transport
ISDN or dial-up servers	Access network device	Medium	Administrators Partners and privileged users for special access
Firewall	Access network device	High	Administrators All others for use as a transport
DNS and DHCP servers	Network applications	Medium	Administrators General and privileged users for use
Internal e-mail server	Network application	Medium	Administrators All other internal users for use
Oracle database	Network application	Medium or High	Administrators Privileged users for data updates General users for data access All others for partial data access

ARCH v1.1—6-10

The identification of the risk level and the type of access required of each network system forms the basis of the following security matrix. The security matrix provides a quick reference for each system and a starting point for further security measures, such as creating an appropriate strategy for restricting access to network resources. The figure shows an example of a security matrix.

Use the table to create your own security matrix.

System	Description	Risk Level	Types of Users

Summary

This topic summarizes the key points discussed in this lesson.

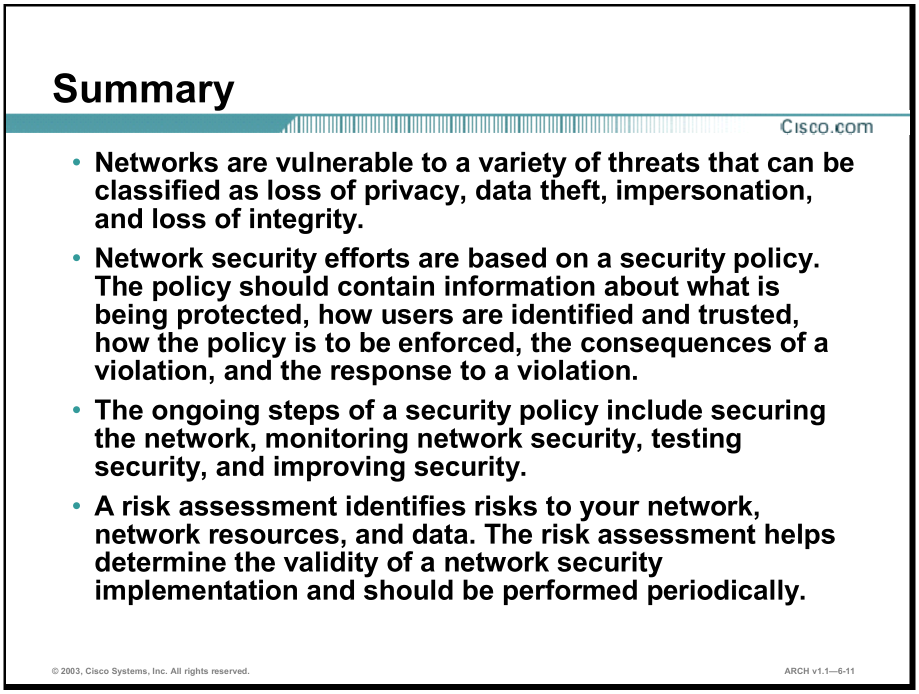

Summary

- Networks are vulnerable to a variety of threats that can be classified as loss of privacy, data theft, impersonation, and loss of integrity.
- Network security efforts are based on a security policy. The policy should contain information about what is being protected, how users are identified and trusted, how the policy is to be enforced, the consequences of a violation, and the response to a violation.
- The ongoing steps of a security policy include securing the network, monitoring network security, testing security, and improving security.
- A risk assessment identifies risks to your network, network resources, and data. The risk assessment helps determine the validity of a network security implementation and should be performed periodically.

ARCH v1.1—6-11

References

For additional information, refer to this resource:

- *Network Security Policy: Best Practices White Paper* at http://www.cisco.com/warp/public/126/secpol.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) To counteract problems associated with impersonation, you can validate and guarantee the identity of communicating parties by using _____.

A) firewalls

B) authentication protocols

C) security assurance policies

D) intrusion detection systems

Q2) Which component of a security policy specifies how violations are detected?

A) identity

B) acceptable use

C) risk assessment

D) incident response

Q3) Which implementation solution helps ensure data confidentiality and integrity?

A) firewalls

B) encryption

C) authentication

D) vulnerability patching

Q4) Which step in the security process ensures that the security posture of the network is being met?

A) testing

B) securing

C) monitoring

D) improving

Q5) What is the basis for assigning a risk level to a resource?

A) the purchase cost of that resource

B) the impact caused by the destruction of that resource

C) the availability of redundant options for that resource

D) the impact caused by a security violation on that resource

Quiz Answer Key

Q1) B

Relates to: Network Vulnerabilities

Q2) C

Relates to: Defining a Security Policy

Q3) B

Relates to: Network Security as a Process

Q4) C

Relates to: Network Security as a Process

Q5) D

Relates to: Risk Assessment and Management

Reviewing Cisco Security Solutions

Overview

Cisco offers an array of enterprise network security solutions to make the implementation and maintenance of good network security easier and more cost effective. These solutions include dedicated appliances, software, and security capabilities embedded into other Cisco network products.

Relevance

The enterprise security strategy does not include a single product or solution, but encompasses a range of solutions, strategies, and ongoing audits to provide optimal security.

Objectives

Upon completing this lesson, you will be able to identify the necessary components of a Cisco security solution, given specific security requirements, and propose the features and functionality for each component selected. This includes being able to meet these objectives:

- Identify the key components of a Cisco security solution, given specific enterprise security attacks
- Identify the primary security attacks on an enterprise network and propose solutions
- Propose features and functionality for firewall solutions, given specific enterprise security requirements
- Propose features and functionality for Intrusion Detection Systems, given specific enterprise security requirements
- Propose features and functionality for software security solutions, given specific enterprise security requirements
- Describe security options for the specific components of a network

Learner Skills and Knowledge

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

- *Designing for Cisco Internetwork Solutions* (DESGN) course
- Evaluating Network Security Policies lesson

Outline

The outline lists the topics included in this lesson.

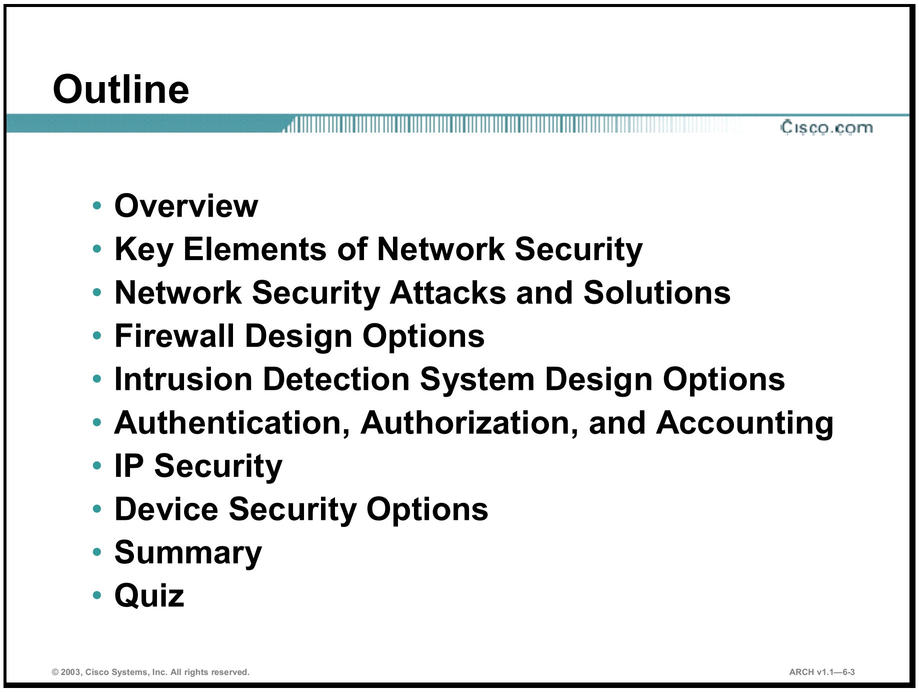

Outline

- Overview
- Key Elements of Network Security
- Network Security Attacks and Solutions
- Firewall Design Options
- Intrusion Detection System Design Options
- Authentication, Authorization, and Accounting
- IP Security
- Device Security Options
- Summary
- Quiz

ARCH v1.1—6-3

Key Elements of Network Security

An effective security solution includes secure connectivity, perimeter security, intrusion protection, identity, and security management. This topic describes the elements of Cisco security solutions.

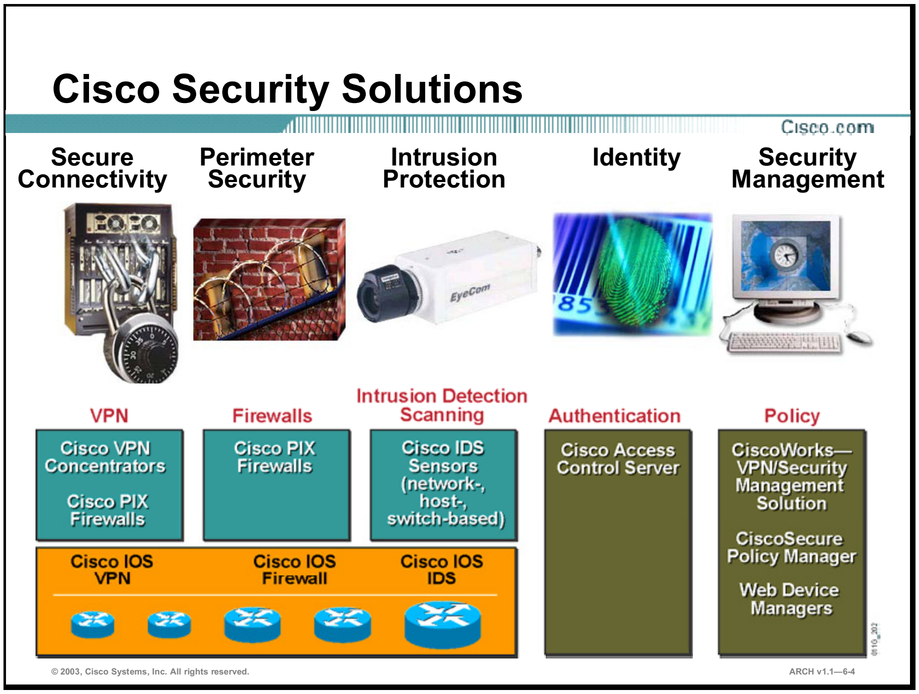

Cisco Security Solutions

Cisco.com

Secure Connectivity | Perimeter Security | Intrusion Protection | Identity | Security Management

VPN: Cisco VPN Concentrators; Cisco PIX Firewalls

Firewalls: Cisco PIX Firewalls

Intrusion Detection Scanning: Cisco IDS Sensors (network-, host-, switch-based)

Authentication: Cisco Access Control Server

Policy: CiscoWorks—VPN/Security Management Solution; CiscoSecure Policy Manager; Web Device Managers

Cisco IOS VPN | Cisco IOS Firewall | Cisco IOS IDS

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-4

The Cisco security solution is comprised of five key elements:

- **Secure connectivity:** When you must protect information from eavesdropping, the ability to provide authenticated, confidential communication on demand is crucial. Sometimes, data separation using tunneling technologies, such as generic routing encapsulation (GRE) or Layer 2 Tunneling Protocol (L2TP), provides effective data privacy. Often, however, additional privacy requirements call for the use of digital encryption technology and protocols such as IP Security (IPSec). This added protection is especially important when implementing VPNs.

- **Perimeter security:** This element provides the means to control access to critical network applications, data, and services so that only legitimate users and information can pass through the network. Routers and switches with Access Control Lists, stateful firewall implementations, and dedicated firewall appliances provide perimeter security control. Complementary tools, including virus scanners and content filters, also help control network perimeters.

- **Intrusion protection:** To ensure that a network remains secure, it is important to regularly test and monitor the state of security preparation. Network vulnerability scanners can proactively identify areas of weakness, and intrusion detection systems can monitor and respond to security events as they occur. Using security monitoring solutions, organizations can obtain unprecedented visibility into both the network data stream and the security posture of the network.

- **Identity:** Identity is the accurate and positive identification of network users, hosts, applications, services, and resources. Standard technologies that enable identification include authentication protocols such as Remote Authentication Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System plus (TACACS+), Kerberos, and one-time password tools. New technologies such as digital certificates, smart cards, and directory services are beginning to play increasingly important roles in identity solutions.
- **Security management:** As networks grow in size and complexity, the requirement for centralized policy management tools grows as well. Sophisticated tools that can analyze, interpret, configure, and monitor the state of security policy, with browser-based user interfaces, enhance the usability and effectiveness of network security solutions.

Network Security Attacks and Solutions

Attacks against network security come in many forms. Each has corresponding actions that you can take to prevent or mitigate the consequences of an attack. This topic describes the necessary components of a Cisco security solution, given specific enterprise security attacks.

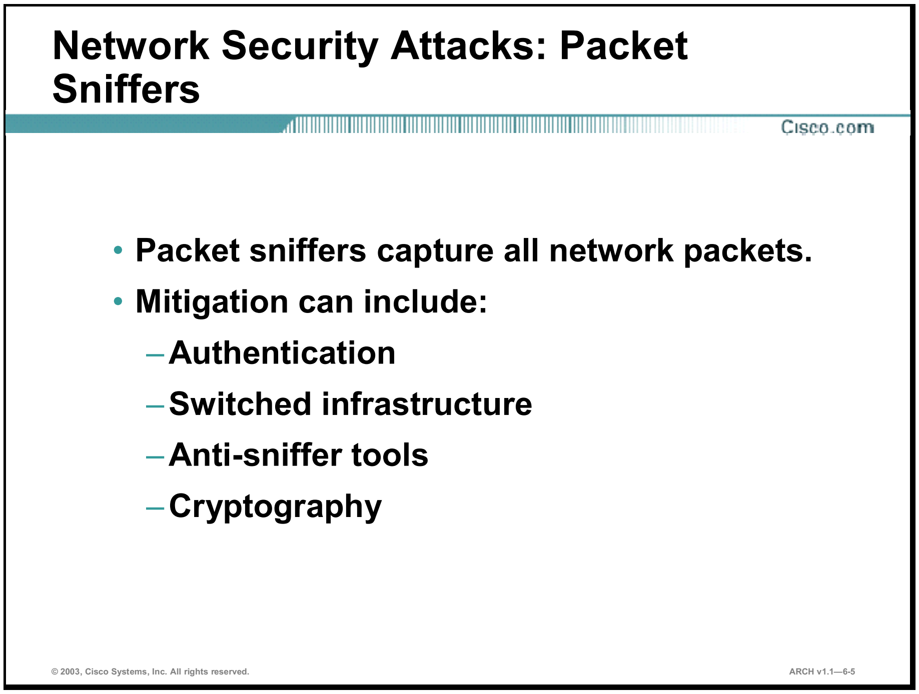

Network Security Attacks: Packet Sniffers

Cisco.com

- Packet sniffers capture all network packets.
- Mitigation can include:
  - Authentication
  - Switched infrastructure
  - Anti-sniffer tools
  - Cryptography

ARCH v1.1—6-5

A packet sniffer is a software application that uses a network adapter card in promiscuous mode (a mode in which the network adapter card sends all packets received on the physical network wire to an application for processing) to capture all network packets that are sent across a particular collision domain. Sniffers are used legitimately in networks today to aid in troubleshooting and traffic analysis. However, because several network applications send data in clear text, a packet sniffer can access meaningful and often sensitive information, such as usernames and passwords.

You can mitigate the threat of packet sniffers in several ways:

- **Authentication:** Using strong authentication is a first option for defense against packet sniffers.
- **Switched infrastructure:** Another method to counter the use of packet sniffers in your environment is to deploy a switched campus network infrastructure.
- **Anti-sniffer tools:** A third method used against sniffers is to employ software and hardware designed to detect the use of sniffers on a network. These tools wait for the sniffer owner to initiate an attack.
- **Cryptography:** An effective method for countering packet sniffers; does not prevent or detect packet sniffers, but rather renders them irrelevant.

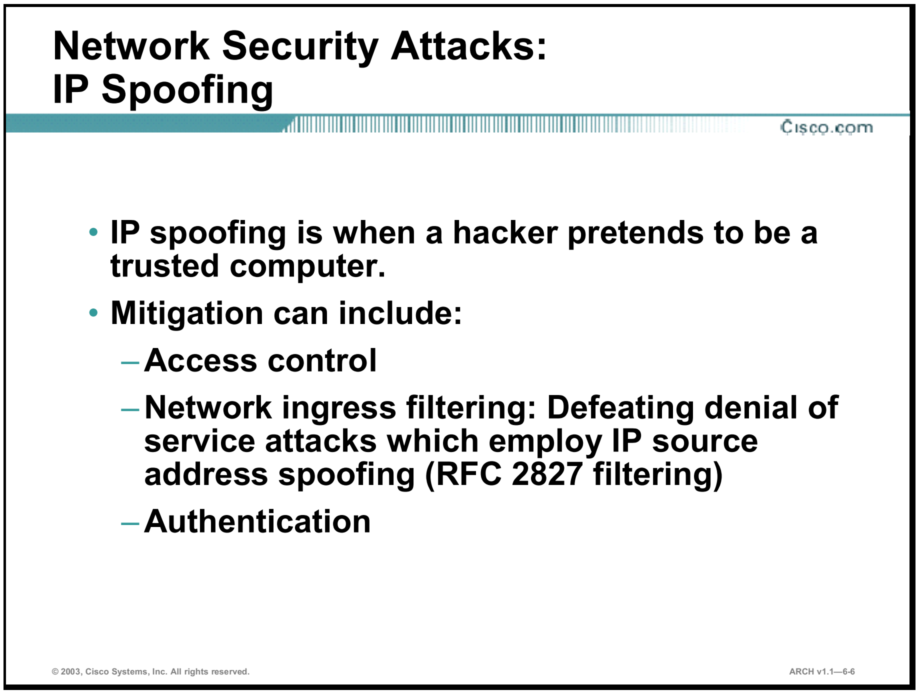

Network Security Attacks: IP Spoofing

Cisco.com

- IP spoofing is when a hacker pretends to be a trusted computer.
- Mitigation can include:
  - Access control
  - Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing (RFC 2827 filtering)
  - Authentication

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-6

An IP spoofing attack occurs when a hacker inside or outside a network pretends to be a trusted computer. A hacker can use an IP address that is within the range of trusted IP addresses for a network, or an authorized external IP address that is trusted and to which access is provided to specified resources on a network. IP spoofing attacks are often a launch point for other attacks. The classic example is to launch a denial-of-service attack using spoofed source addresses to hide the hacker's identity. The spoofed address need not be trusted.

You can reduce, but not eliminate, the threat of IP spoofing through these measures:

- **Access control:** The most common method for preventing IP spoofing is to properly configure access control. To reduce the effectiveness of IP spoofing, configure access control to deny any traffic from the external network that has a source address that should reside on the internal network.
- **Network ingress filtering: Defeating denial-of-service attacks that employ IP source address spoofing (RFC 2827 filtering):** You can also prevent a network's users from spoofing other networks by preventing any outbound traffic on your network that does not have a source address in your organization's own IP range. Your ISP can also implement this type of filtering, which is collectively referred to as RFC 2827 filtering. This filtering denies any traffic that does not have the source address that was expected on a particular interface.
- **Authentication:** IP spoofing can function correctly only when devices use IP address-based authentication. Therefore, if you use additional authentication methods, IP spoofing attacks are irrelevant.

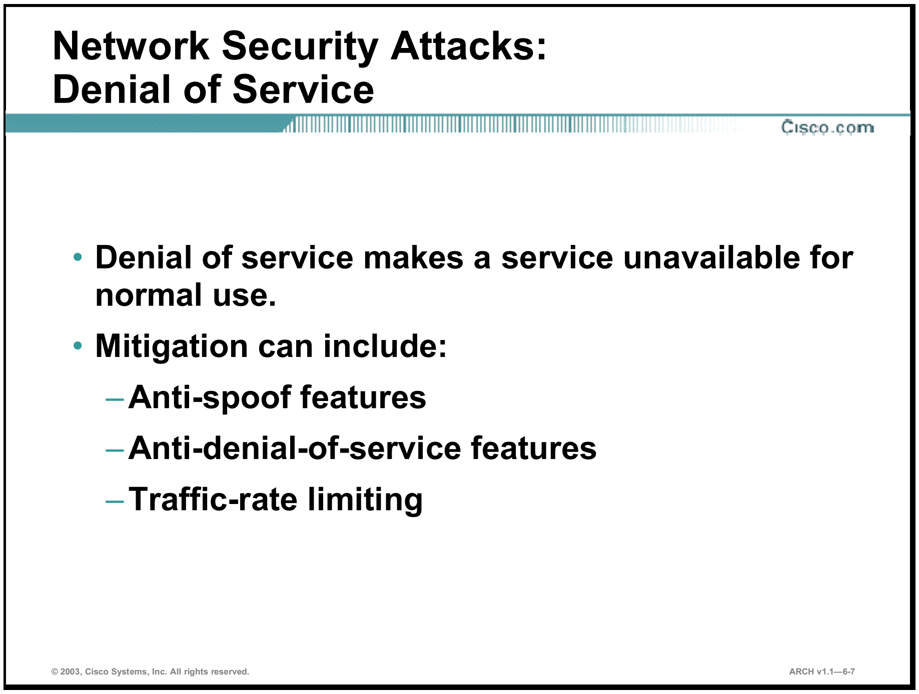

Network Security Attacks: Denial of Service

- Denial of service makes a service unavailable for normal use.
- Mitigation can include:
  - Anti-spoof features
  - Anti-denial-of-service features
  - Traffic-rate limiting

ARCH v1.1—6-7

Denial-of-service attacks are different from most other attacks because they are generally not targeted at gaining access to your network or the information on your network. These attacks focus on making a service unavailable for normal use, which is typically accomplished by exhausting some resource limitation on the network or within an operating system or application.

When involving specific network server applications, such as a web server or an FTP server, these attacks can focus on acquiring and keeping open all the available connections supported by that server, effectively locking out valid users of the server or service. Hackers can implement denial-of-service attacks using common Internet protocols, such as TCP and Internet Control Message Protocol (ICMP). Most denial-of-service attacks exploit a weakness in the overall architecture of the system being attacked rather than a software bug or security hole. Some attacks compromise the performance of your network by flooding the network with undesired, and often useless, network packets, and by providing false information about the status of network resources. This type of attack is often the most difficult to prevent as it requires coordination with your upstream network provider. If traffic meant to consume your available bandwidth is not stopped there, denying it at the point of entry into your network will do little good because your available bandwidth has already been consumed. When this type of attack is launched from many different systems at the same time, it is referred to as a distributed denial of service attack (DDoS).

You can reduce the threat of denial-of-service attacks by using these three methods:

- **Anti-spoof features:** Proper configuration of anti-spoof features on your routers and firewalls can reduce your risk. This includes RFC 2827 filtering at a minimum. If hackers cannot mask their identities, they might not attack.
- **Anti-denial-of-service features:** Proper configuration of anti-denial-of-service features on routers and firewalls can help limit the effectiveness of an attack. These features often involve limits on the amount of half-open connections that a system allows open at any given time.

- **Traffic-rate limiting:** An organization can implement traffic-rate limiting with its ISP. This type of filtering limits the amount of nonessential traffic that crosses network segments. A common example is to limit the amount of ICMP traffic allowed into a network that is used only for diagnostic purposes. ICMP-based denial-of-service attacks are common.

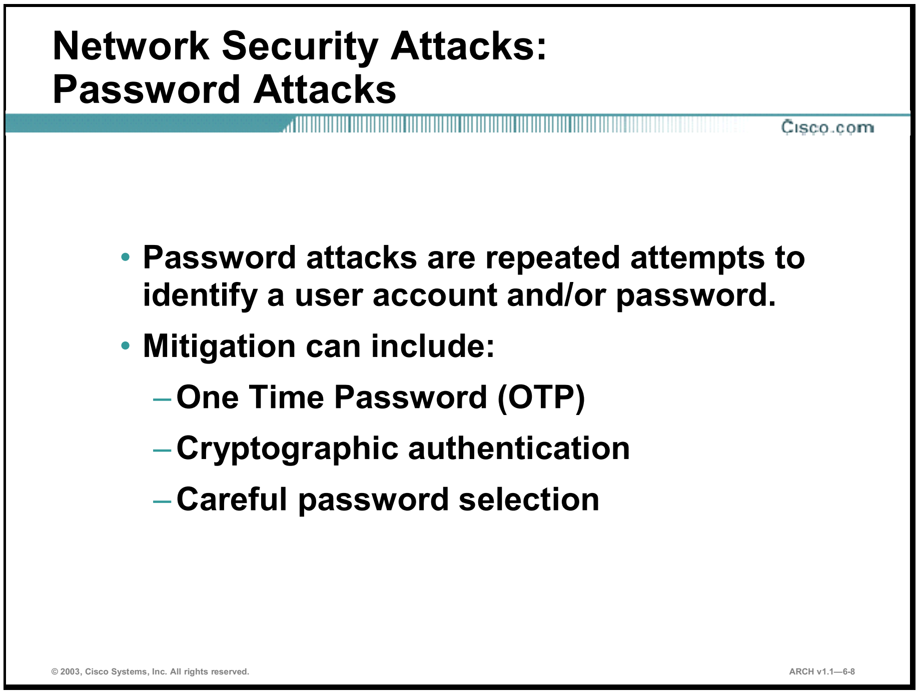

Network Security Attacks: Password Attacks

- Password attacks are repeated attempts to identify a user account and/or password.
- Mitigation can include:
  - One Time Password (OTP)
  - Cryptographic authentication
  - Careful password selection

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-8

Hackers can implement password attacks using several different methods, including brute-force attacks, Trojan horse programs, IP spoofing, and packet sniffers. Although packet sniffers and IP spoofing can yield user accounts and passwords, password attacks usually amount to repeated attempts to identify a user account and/or password, called brute-force attacks.

A brute-force attack is often performed using a program that runs across the network and attempts to log in to a shared resource, such as a server. When hackers successfully gain access to resources, they have the same rights as the compromised account users to gain access to those resources. If the compromised accounts have sufficient privileges, the hackers can create back doors for future access without concern for any status and password changes to the compromised user accounts.

Another problem is created when users use the same password on every system they utilize, including personal systems, corporate systems, and systems on the Internet. Because a password is only as secure as the host that contains it, if that host is compromised, hackers can try the same password on a whole range of hosts.

You can most easily eliminate password attacks by not relying on plain-text passwords in the first place. Using One Time Password (OTP) and/or cryptographic authentication can virtually eliminate the threat of password attacks. Unfortunately, not all applications, hosts, and devices support these authentication methods. When standard passwords are used, it is important to choose a password that is difficult to guess. Passwords should be at least eight characters long and contain uppercase letters, lowercase letters, numbers, and special characters (#%$, and so on). The best passwords are randomly generated but are very difficult to remember, often leading users to write down their passwords.

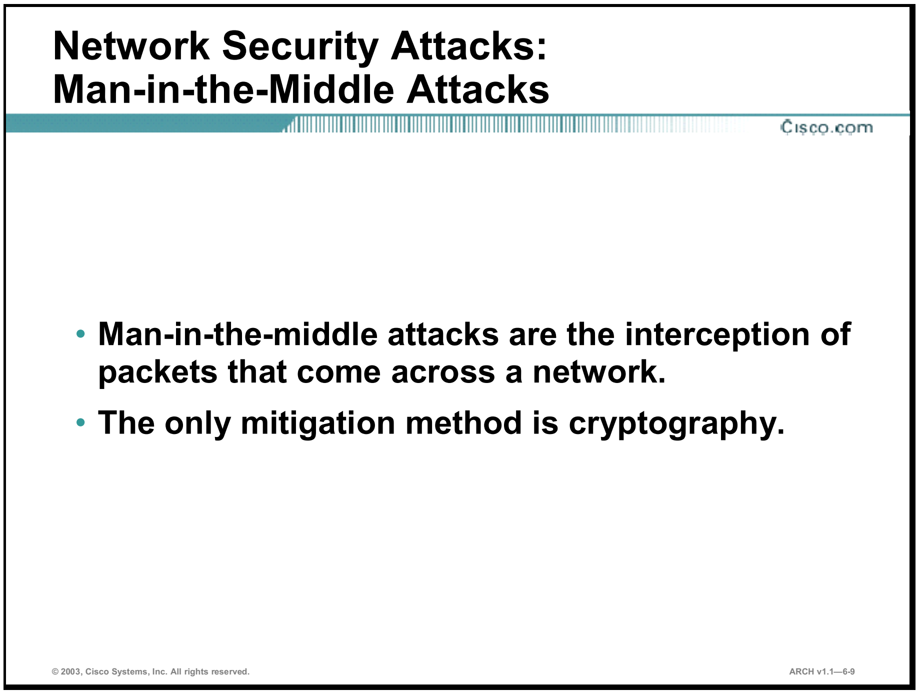

Network Security Attacks: Man-in-the-Middle Attacks

- Man-in-the-middle attacks are the interception of packets that come across a network.
- The only mitigation method is cryptography.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-9

A man-in-the-middle attack requires that the hacker have access to network packets that come across a network. An example of such a configuration could be someone who is working for an ISP who has access to all network packets transferred between his employer's network and any other network. Such attacks are often implemented using network packet sniffers and routing and transport protocols. The possible uses of such attacks are theft of information, hijacking of an ongoing session to gain access to private network resources, traffic analysis to derive information about a network and its users, denial of service, corruption of transmitted data, and introduction of new information into network sessions.

Man-in-the-middle attacks are effectively mitigated only through the use of cryptography. If someone hijacks data in the middle of a cryptographically private session, all the hacker will see is cipher text and not the original message. If a hacker can learn information about the cryptographic session (such as the session key), man-in-the-middle attacks are still possible.

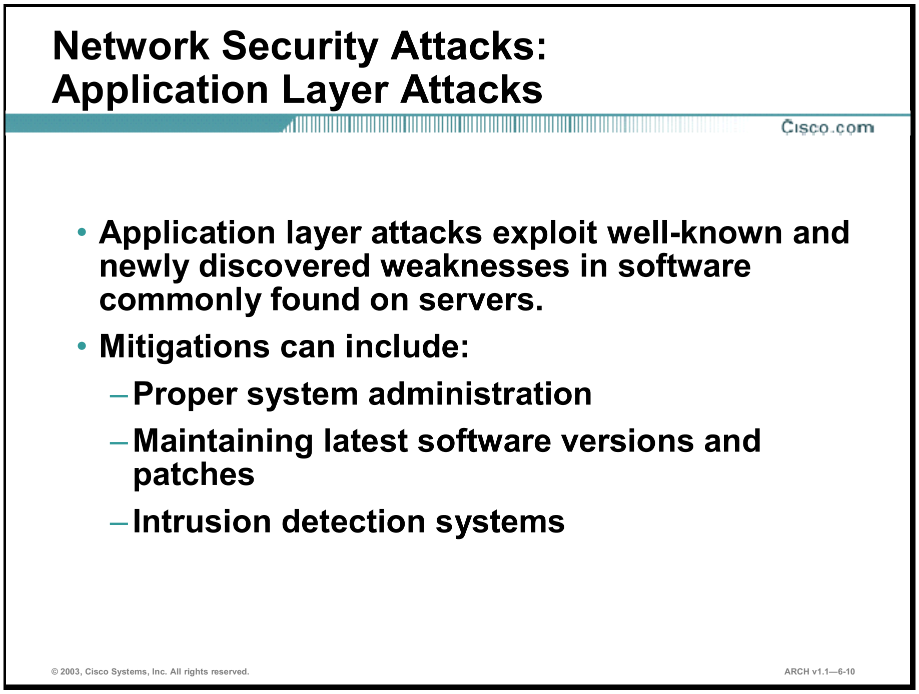

Network Security Attacks: Application Layer Attacks

Cisco.com

- Application layer attacks exploit well-known and newly discovered weaknesses in software commonly found on servers.
- Mitigations can include:
  - Proper system administration
  - Maintaining latest software versions and patches
  - Intrusion detection systems

ARCH v1.1—6-10

Application layer attacks exploit well-known or newly discovered weaknesses in software that are commonly found on servers, such as sendmail, HTTP, and FTP. By exploiting these weaknesses, hackers can gain access to a computer with the permissions of the account running the application, which is usually a privileged system-level account. These application-layer attacks are often widely publicized to encourage administrators to correct the problem with a patch. Unfortunately, many hackers also subscribe to these same mailing lists, which results in their learning about the vulnerabilities in the software.

The primary problem with application-layer attacks is that they often use ports that allow traffic through a firewall. For example, a hacker executing a known vulnerability against a web server often uses TCP port 80. Because the web server serves pages to users, a firewall needs to allow access on that port. From the perspective of a firewall, the hacker's input is merely standard port 80 traffic.

You can never completely eliminate application-layer attacks. Hackers continually discover and publicize new vulnerabilities. The best way to reduce your risk is to practice good system administration. A few measures you can take to further reduce your risks are:

- Read operating system and network log files and analyze them using log analysis applications.
- Subscribe to mailing lists that publicize vulnerabilities such, as Bugtraq (http://www.securityfocus.com/) and the CERT (http://www.cert.org/).
- Keep your operating systems and applications current with the latest patches.
- Use Intrusion Detection Systems (IDSs) to minimize application-layer attacks.

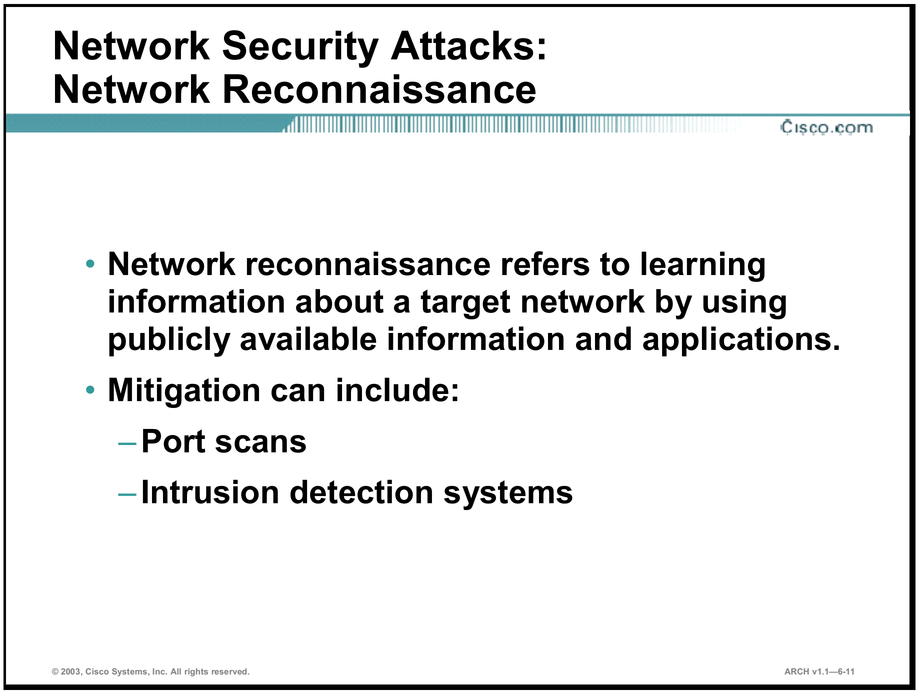

Network Security Attacks: Network Reconnaissance

- Network reconnaissance refers to learning information about a target network by using publicly available information and applications.
- Mitigation can include:
  - Port scans
  - Intrusion detection systems

ARCH v1.1—6-11

Network reconnaissance refers to the act of learning information about a target network by using publicly available information and applications. When hackers attempt to penetrate a particular network, they often need to learn as much information as possible about the network before launching attacks, often using DNS queries and **ping** sweeps. DNS queries can reveal information such as who owns a particular domain and what addresses are assigned to that domain. **ping** sweeps of the addresses revealed through DNS queries present a picture of the live hosts in a particular environment. The hackers can examine the characteristics of the applications that are running on the hosts. This can lead to specific information that is useful when the hacker attempts to compromise that service.

You cannot prevent network reconnaissance entirely. If ICMP echo and echo-reply is turned off on edge routers, for example, you can stop **ping** sweeps, but at the expense of network diagnostic data. However, you can run port scans without full **ping** sweeps; they simply take longer because they need to scan IP addresses that might not be live. IDS at the network and host levels can usually notify an administrator when a reconnaissance gathering attack is underway. This allows the administrator to better prepare for the coming attack or to notify the ISP who is hosting the system that is launching the reconnaissance probe.

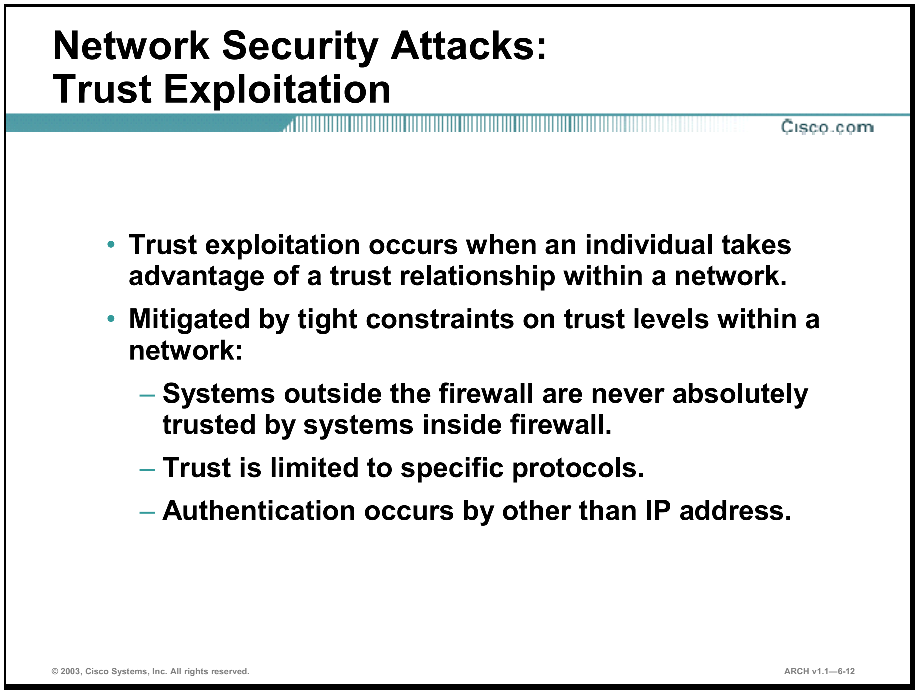

Network Security Attacks: Trust Exploitation

- Trust exploitation occurs when an individual takes advantage of a trust relationship within a network.
- Mitigated by tight constraints on trust levels within a network:
  - Systems outside the firewall are never absolutely trusted by systems inside firewall.
  - Trust is limited to specific protocols.
  - Authentication occurs by other than IP address.

ARCH v1.1—6-12

Trust exploitation refers to an attack where an individual takes advantage of a trust relationship within a network. The classic example is a perimeter network connection from a corporation, which often houses enterprise-wide servers. Because they all reside on the same segment, a compromise of one system can lead to the compromise of other systems, since they might trust other systems attached to the same network. Another example is a system on the outside of a firewall that has a trust relationship with a system on the inside of a firewall. When the outside system is compromised, it can leverage that trust relationship to attack the inside network.

You can mitigate trust exploitation-based attacks through tight constraints on trust levels within a network. Systems on the inside of a firewall should never absolutely trust systems on the outside of a firewall. Limit trust to specific protocols and authenticate with more than an IP address where possible.

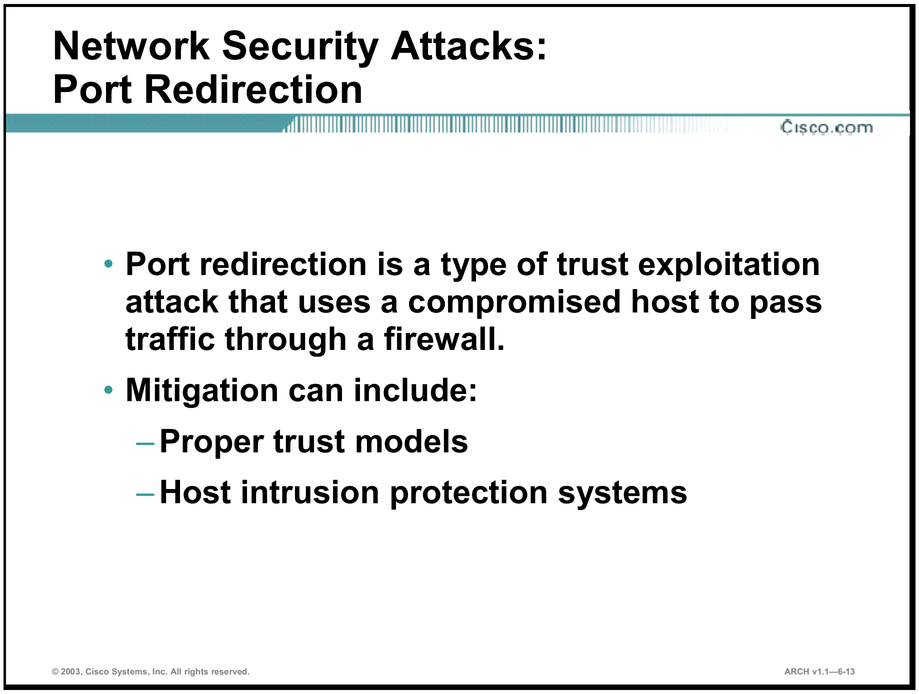

Network Security Attacks: Port Redirection

Cisco.com

- Port redirection is a type of trust exploitation attack that uses a compromised host to pass traffic through a firewall.
- Mitigation can include:
  - Proper trust models
  - Host intrusion protection systems

ARCH v1.1—6-13

Port redirection attacks are a type of trust exploitation attack that uses a compromised host to pass traffic through a firewall that would otherwise be dropped. Consider a firewall with three interfaces and a host on each interface. The host on the outside can reach the host on the public services segment (commonly referred to as a perimeter LAN), but not the host on the inside. The host on the public services segment can reach the host on both the outside and the inside. If hackers could compromise the public services segment host, they could install software to redirect traffic from the outside host directly to the inside host. Though neither communication violates the rules implemented in the firewall, the outside host has now achieved connectivity to the inside host through the port redirection process on the public services host.

You can mitigate most port redirection using proper trust models. Assuming a system under attack, a host intrusion protection system (HIPS) can help detect and prevent a hacker installing such utilities on a host.

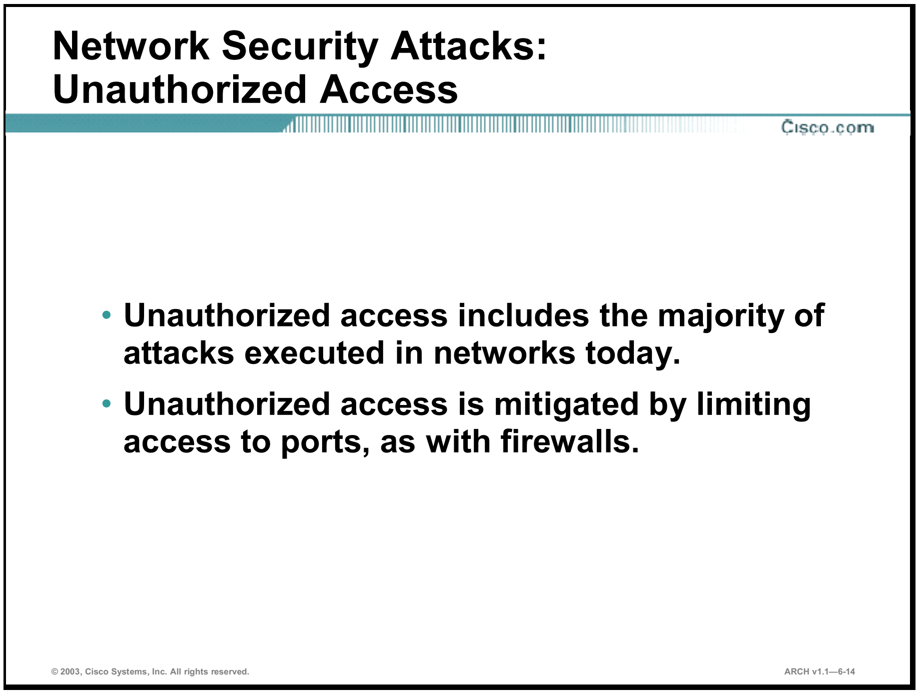

Network Security Attacks: Unauthorized Access

Cisco.com

- Unauthorized access includes the majority of attacks executed in networks today.
- Unauthorized access is mitigated by limiting access to ports, as with firewalls.

ARCH v1.1—6-14

While not a specific type of attack, unauthorized access attacks refer to the majority of attacks executed in networks today. For someone to brute-force a Telnet login, they must first get the Telnet prompt on a system. Upon connection to the Telnet port, a message might indicate: “authorization required to use this resource.” If the hacker continues to attempt access, the actions become “unauthorized.” A hacker can initiate these kinds of attacks both from the outside and inside of a network.

Mitigation techniques for unauthorized access attacks are very simple. They involve reducing or eliminating the ability of hackers to gain access to a system using an unauthorized protocol. An example would be preventing hackers from having access to the Telnet port on a server that needs to provide web services to the outside. If a hacker cannot reach that port, it is very difficult to attack it. The primary function of a firewall in a network is to prevent simple unauthorized access attacks.

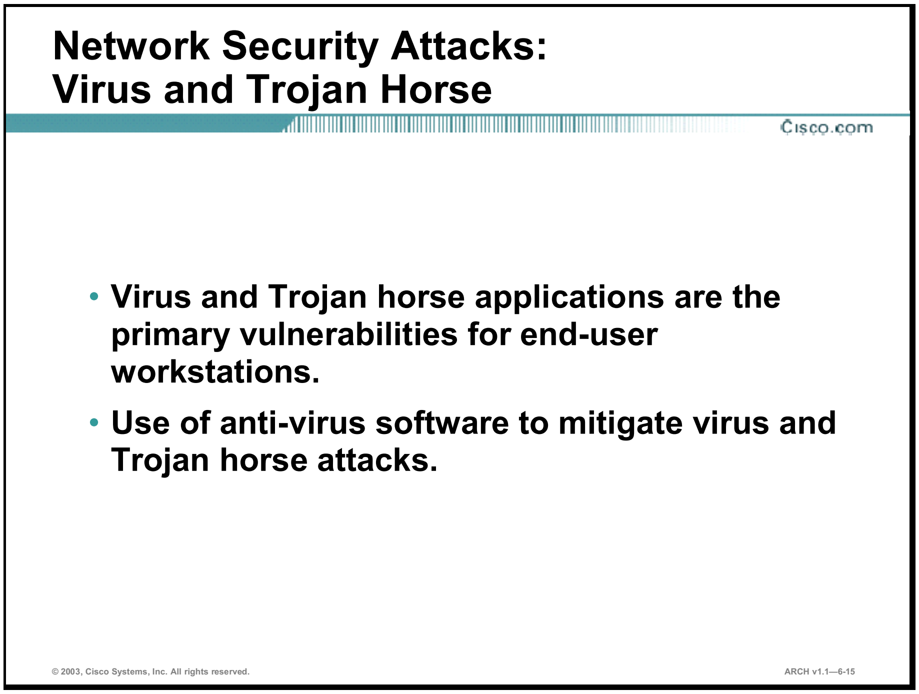

Network Security Attacks: Virus and Trojan Horse

- Virus and Trojan horse applications are the primary vulnerabilities for end-user workstations.
- Use of anti-virus software to mitigate virus and Trojan horse attacks.

ARCH v1.1—6-15

The primary vulnerabilities for end-user workstations are viruses and Trojan horse attacks. Viruses refer to malicious software that is attached to another program to execute a particular unwanted function on a user’s workstation. An example of a virus is a program that is attached to command.com (the primary interpreter for Windows systems), which deletes certain files and infects any other versions of command.com that it can find. A Trojan horse is different only in that the entire application was written to look like something else, when in fact it is an attack tool. An example of a Trojan horse is a software application that runs a simple game on the user’s workstation. While the user is occupied with the game, the Trojan horse mails a copy of itself to every user in the user’s address book. Then other users get the game and play it, thus spreading the Trojan horse.

You can contain Trojan horse applications through the effective use of anti-virus software at the user level and potentially at the network level. Anti-virus software can detect most viruses and many Trojan horse applications and prevent them from spreading in the network. Keeping current with the latest mitigation techniques can also lead to a more effective posture against these attacks. As new virus or Trojan applications are released, enterprises need to keep up to date with the latest anti-virus software and application versions.

Firewall Design Options

Firewalls provide perimeter security by preventing unauthorized access to the internal network. Identifying the type of traffic that is not allowed to pass the firewall and how such traffic will be prevented are two of the primary decisions to make about a firewall implementation. This topic describes features and functionality for firewall solutions.

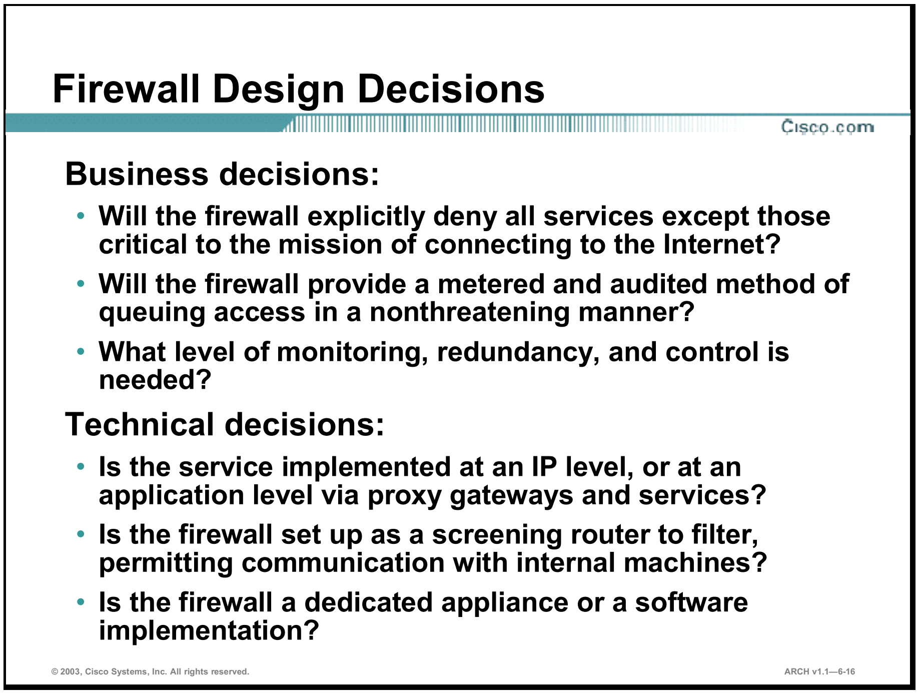

Firewall Design Decisions

Cisco.com

Business decisions:

- **Will the firewall explicitly deny all services except those critical to the mission of connecting to the Internet?**
- **Will the firewall provide a metered and audited method of queuing access in a nonthreatening manner?**
- **What level of monitoring, redundancy, and control is needed?**

Technical decisions:

- **Is the service implemented at an IP level, or at an application level via proxy gateways and services?**
- **Is the firewall set up as a screening router to filter, permitting communication with internal machines?**
- **Is the firewall a dedicated appliance or a software implementation?**

ARCH v1.1—6-16

The first and most important firewall design decision reflects the policy of how the enterprise wants to operate the system. Is the firewall in place explicitly to deny all services except those critical to the mission of connecting to the Internet, or is the firewall in place to provide a metered and audited method of queuing access in a nonthreatening manner?

Another question is: What level of monitoring, redundancy, and control do you want? You can form a checklist of what should be monitored, permitted, and denied. In other words, you start by figuring out your overall objectives as described in your security policy, and then combine a needs analysis with a risk assessment, and sort the requirements into a list that specifies what you plan to implement.

On the technical side, there are also decisions to make. A firewall is a static traffic routing service placed between the network service provider's router and the internal network. The traffic routing service may be implemented at an IP level via something like screening rules in a router, or at an application level via proxy gateways and services.

The decision to make is whether to place an exposed firewall on the outside network to run proxy services for Telnet, FTP, news, and so on, or whether to set up a screening router as a filter, permitting communication with one or more internal machines. There are pluses and minuses to both approaches, with the proxy machine providing a greater level of audit and security in return for increased cost in configuration and a decrease in the level of service that may be provided, since a proxy needs to be developed for each desired service.

Finally, you need to decide whether you will deploy a dedicated firewall hardware appliance or use an integrated software solution on a router. The integrated functionality is often attractive because you can implement it on existing equipment, or because the features can interoperate with the rest of the device to provide a better functional solution. Appliances are often used when the depth of functionality required is very advanced or when performance needs require using specialized hardware. Make your decisions based on the capacity and functionality of the appliance versus the integration advantage of the device.

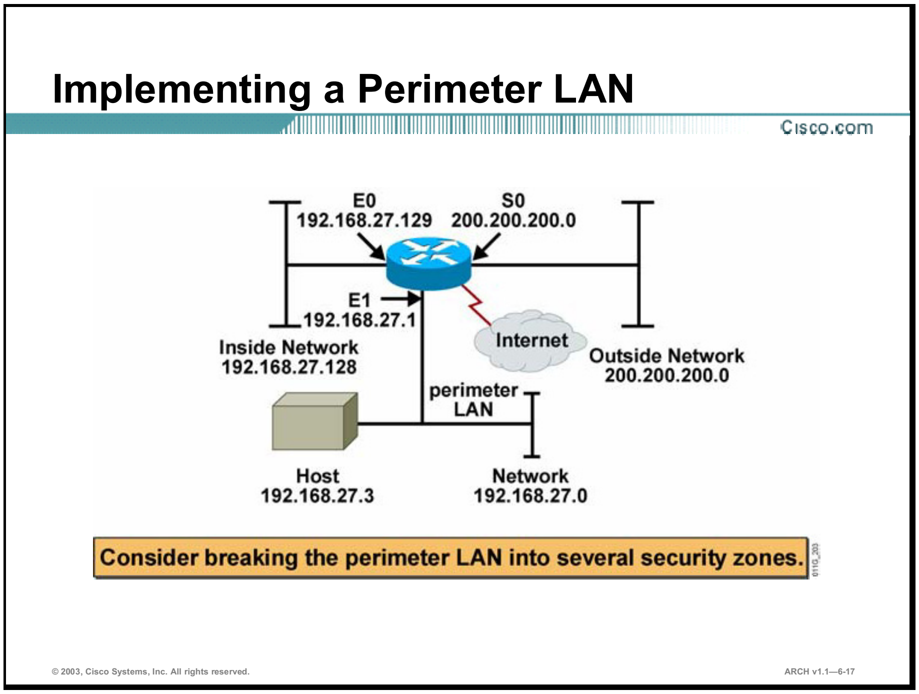

Implementing a Perimeter LAN

Cisco.com

Consider breaking the perimeter LAN into several security zones.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-17

Perimeter LAN is another term for a demilitarized zone. In the context of firewalls, this refers to a part of the network that is neither part of the internal network nor directly part of the Internet. Typically, this is the area between the Internet access router and the bastion host, though it can be between any two policy-enforcing components of the network architecture.

You can create a perimeter LAN by putting access control lists on your access router. This minimizes the exposure of hosts on the external LAN by allowing only recognized and managed services on those hosts to be accessible by hosts on the Internet. For example, a web server running Microsoft Windows NT might be vulnerable to a number of denial-of-service attacks against such services. These services are not required for the operation of a web server, so blocking TCP connections to ports 135, 137, 138, and 139 on that host will reduce the exposure to a denial-of-service attack. In fact, if you block everything but HTTP traffic to that host, an attacker will only have one service to attack.

A common approach that an attacker uses is to break into a host that is vulnerable to attack, and exploit trust relationships between the vulnerable host and more interesting targets.

If you are running a number of services that have different levels of security, you might want to consider breaking the perimeter LAN into several security zones. For example, the access router could feed two Ethernet segments, both protected by access control lists, and therefore in the perimeter LAN.

On one of the Ethernet segments, you might have hosts that provide Internet connectivity. These will likely relay mail, news, and host DNS. On the other Ethernet segment, you could have web servers and other hosts that provide services for the benefit of Internet users.

In many organizations, services for Internet users tend to be less carefully guarded and are more likely to be doing insecure things. (For example, in the case of a web server, unauthenticated and untrusted users might be running Common Gateway Interface [CGI] or other executable programs. This might be reasonable for the web server, but brings with it a certain set of risks that need to be managed. It is likely these services are too risky for an organization to run them on a bastion host, where a slipup can result in the complete failure of the security mechanisms.)

By splitting services not only by host, but by network, and limiting the level of trust between hosts on those networks, you can greatly reduce the likelihood of a break-in on one host being used to break into another.

You can also increase the scalability of your architecture by placing hosts on different networks. The fewer machines that there are to share the available bandwidth, the more bandwidth that each will get.

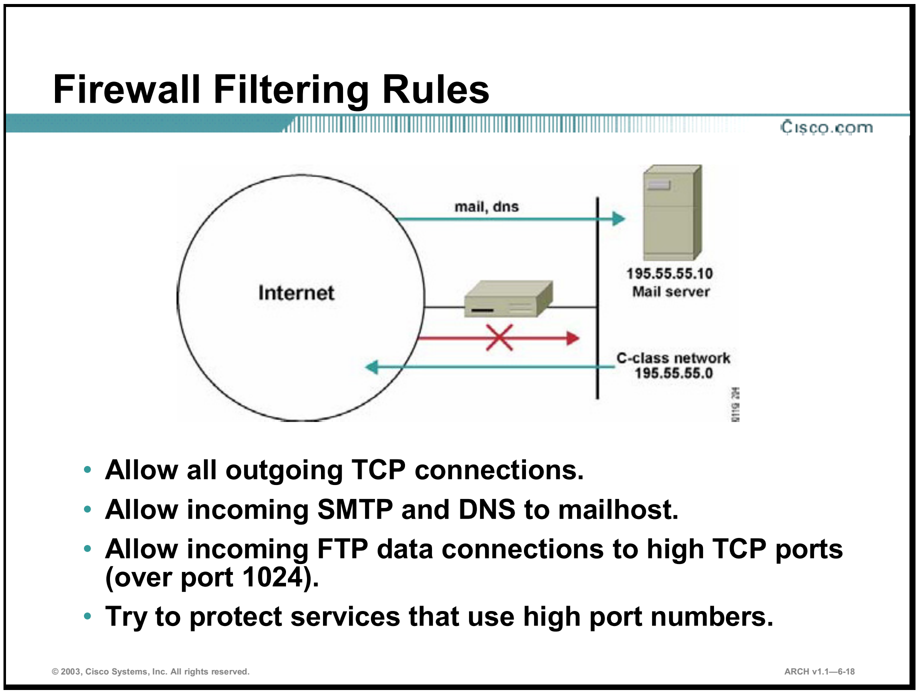

Firewall Filtering Rules

Cisco.com

- Allow all outgoing TCP connections.
- Allow incoming SMTP and DNS to mailhost.
- Allow incoming FTP data connections to high TCP ports (over port 1024).
- Try to protect services that use high port numbers.

ARCH v1.1—6-18

The figure shows one possible example for using the Cisco router as the filtering router for a specific policy. The company has Class C network address 195.55.55.0. The company network is connected to the Internet via an ISP. Company policy is to allow everybody access to Internet services, so all outgoing connections are accepted. All incoming connections go through mailhost. Mail and DNS are only incoming services. Only incoming packets from the Internet are checked.

The firewall in the figure provides these security services:

- Allows all outgoing TCP connections
- Allows incoming SMTP and DNS to mailhost
- Allows incoming FTP data connections to high TCP port (over port 1024)
- Tries to protect services that live on high port numbers

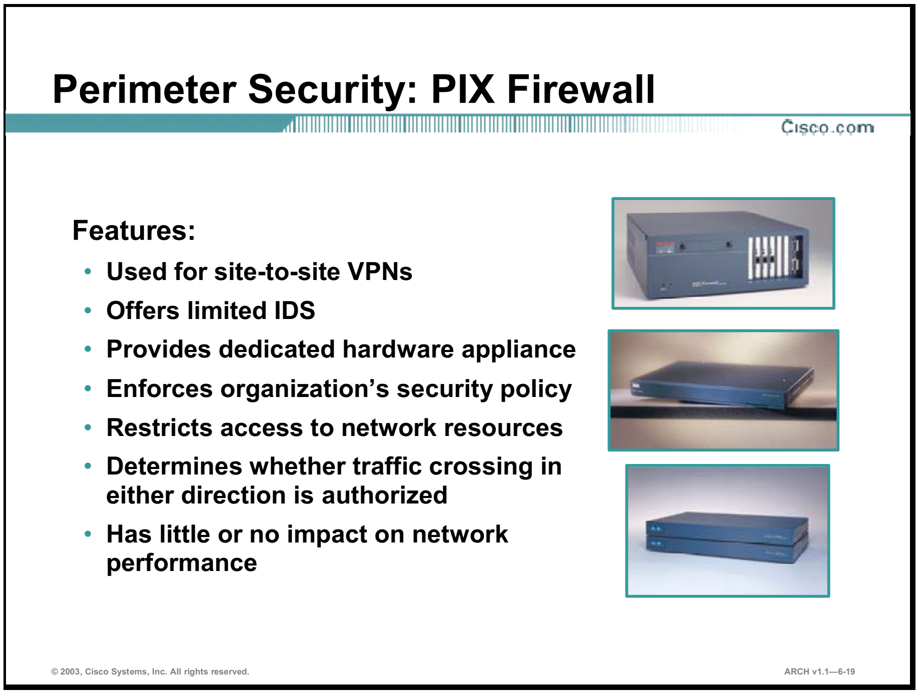

Perimeter Security: PIX Firewall

Features:

- **Used for site-to-site VPNs**
- **Offers limited IDS**
- **Provides dedicated hardware appliance**
- **Enforces organization’s security policy**
- **Restricts access to network resources**
- **Determines whether traffic crossing in either direction is authorized**
- **Has little or no impact on network performance**

ARCH v1.1—6-19

The Cisco PIX Firewall is a dedicated hardware appliance that implements firewall services. Built upon a proprietary operating system for security services, PIX OS, PIX Firewalls provide a range of security services, including:

- Network Address Translation (NAT)
- Port address translation (PAT)
- Content filtering (Java/ActiveX)
- URL filtering
- Authentication, authorization, and accounting (AAA) RADIUS/TACACS+ integration
- Support for X.509 public-key infrastructure (PKI) solutions
- DHCP client/server
- PPP over Ethernet (PPPoE) support

PIX Firewalls support VPN clients, including Cisco hardware and software VPN clients, and Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP) clients found within Microsoft Windows operating systems. A PIX Firewall does not add any noticeable latency on the network.

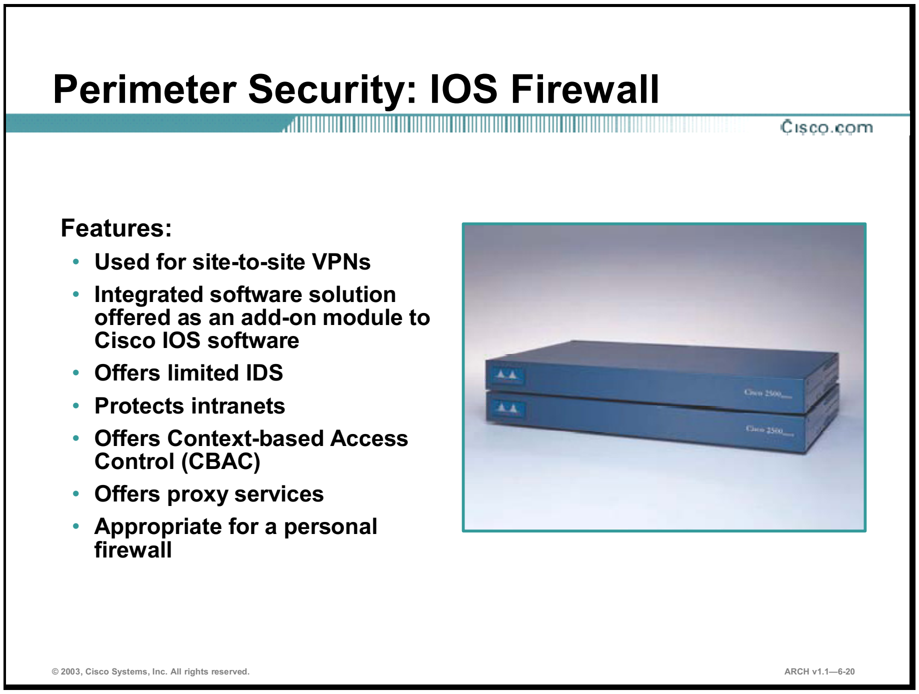

Perimeter Security: IOS Firewall

Features:

- **Used for site-to-site VPNs**
- **Integrated software solution offered as an add-on module to Cisco IOS software**
- **Offers limited IDS**
- **Protects intranets**
- **Offers Context-based Access Control (CBAC)**
- **Offers proxy services**
- **Appropriate for a personal firewall**

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-20

As an alternative to a dedicated hardware device, the Cisco IOS firewall is an add-on software module that allows a router to provide firewall services without additional hardware. It integrates firewall functionality and intrusion detection by providing stateful, application-based filtering; dynamic per-user authentication and authorization; defense against network attacks; Java blocking; and real-time alerts. In combination with IPSec and other Cisco IOS software technologies, the Cisco IOS firewall provides a complete VPN solution.

IOS firewalls support VPN clients, including Cisco hardware and software VPN clients, and PPTP and L2TP clients found within Microsoft Windows operating systems.

Intrusion Detection System Design Options

An Intrusion Detection System (IDS) detects and responds to attacks. Host intrusion protection systems (HIPS) protect individual hosts, while network IDSs (NIDSs) protect the overall network. This topic describes the features and functionality of IDSs.

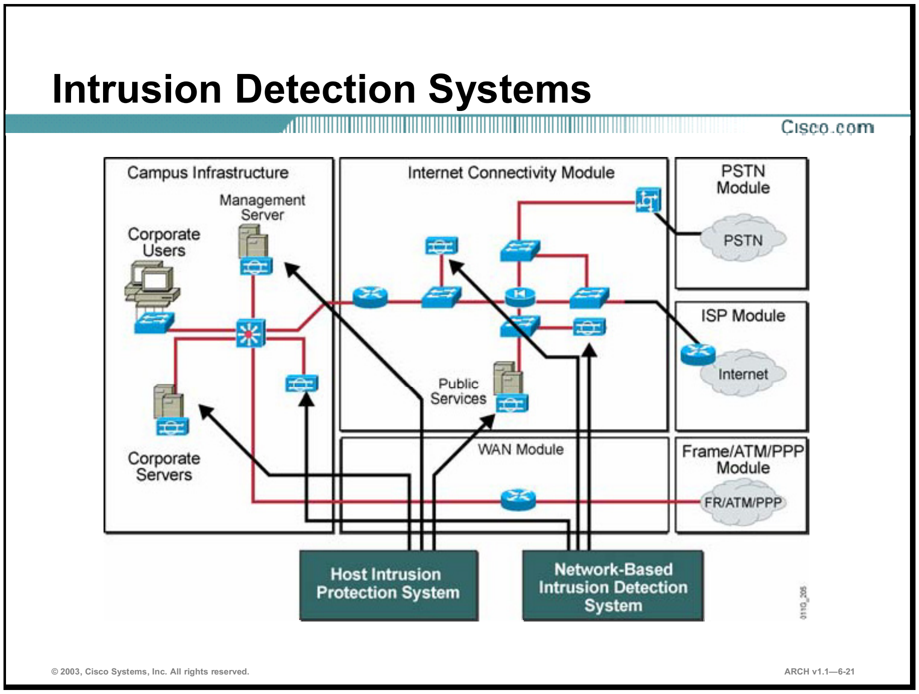

Intrusion Detection Systems

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-21

IDSs act like an alarm system in the physical world. There are two complementary IDS technologies:

- HIPS operate by inserting agents into the host to be protected. It is then concerned only with attacks generated against that one host.
- NIDS operate by watching all packets traversing a particular collision domain. When NIDS sees a packet or series of packets that match a known or suspect attack, it can flag an alarm and/or terminate the session.

IDSs operate by using attack signatures. Attack signatures are the profile for a particular attack or kind of attack. They specify certain conditions that must be met before traffic is deemed to be an attack. In the physical world, IDSs are most closely compared to an alarm system or security camera. When an IDS detects something that it considers an attack, it can either take corrective action itself or notify a management system for actions by the administrator.

Some systems are more or less equipped to respond to and prevent such an attack. Host intrusion detection can work by intercepting operating system and application calls on an individual host. It can also operate by after-the-fact analysis of local log files. The former approach allows better attack prevention; whereas, the latter approach dictates a more passive attack-response role. Because of the specificity of their role, HIPS are often better at preventing specific attacks than NIDS, which usually only issue an alert upon discovery of an attack. However, that specificity causes a loss of perspective to the overall network. This is where NIDS excels. Cisco recommends a combination of the two systems, HIPS on critical hosts and NIDS looking over the whole network, for a complete intrusion detection system.

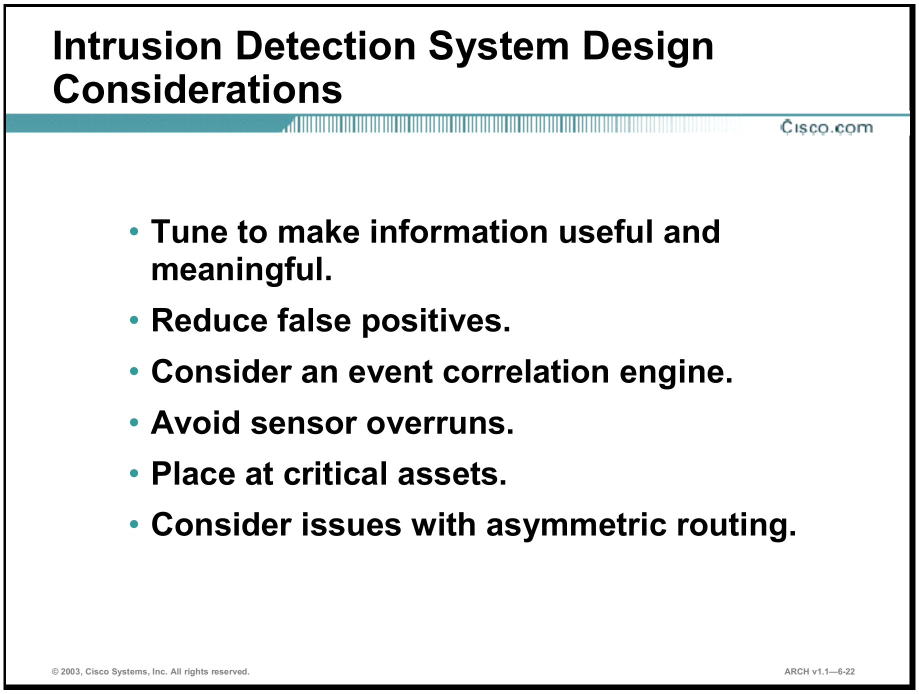

Intrusion Detection System Design Considerations

- Tune to make information useful and meaningful.
- Reduce false positives.
- Consider an event correlation engine.
- Avoid sensor overruns.
- Place at critical assets.
- Consider issues with asymmetric routing.

ARCH v1.1—6-22

When you deploy an IDS, you must tune its implementation to increase its effectiveness and remove false positives. False positives are alarms caused by legitimate traffic or activity. False negatives are attacks that the IDS fails to see. When you tune an IDS, you can configure it more specifically to its threat-mitigation role. You should configure HIPS to stop most valid threats at the host level, because it is well prepared to determine that certain activity is a threat. However, configuring a HIPS is often difficult.

Remember that the first step prior to implementing any threat-response option is to adequately tune NIDS to ensure that any perceived threat is legitimate. When deciding on mitigation roles for NIDS, you have two primary options:

- The first option, and potentially the most damaging if improperly deployed, is to shun traffic by using access control filters on routers and firewalls. When a NIDS detects an attack from a particular host over a particular protocol, it can block that host from coming into the network for a predetermined amount of time. To mitigate the risks of shunning, you should generally use it only on TCP traffic, which is much more difficult to successfully spoof than User Datagram Protocol (UDP). Use it only in cases where the threat is real and the chance that the attack is a false positive is very low.
- The second option for NIDS mitigation is the use of TCP resets. As the name implies, TCP resets operate only on TCP traffic and terminate an active attack by sending TCP reset messages to the attacking and attacked host. Because TCP traffic is more difficult to spoof, you should consider using TCP resets more often than shunning. Keep in mind that TCP resets in a switched environment are more challenging than when a standard hub is used, because all ports don’t see all traffic without the use of a Switched Port Analyzer (SPAN) or mirror port. Make sure this mirror port supports bidirectional traffic flows and can have SPAN port MAC learning disabled.

Both mitigation options require around-the-clock staffing to watch the IDS consoles. Because IT staff is often overworked, consider outsourcing your IDS management to a third party. Another option for reducing monitoring requirements is to deploy a third-party event correlation engine.

From a performance standpoint, NIDS observes packets on the wire. If packets are sent faster than the NIDS can process them, there is no degradation to the network because the NIDS does not sit directly in the flows of data. However, the NIDS will lose effectiveness and packets could be missed, causing both false negatives and false positives. Be sure to avoid exceeding the capabilities of IDSs, so that you can benefit from their services.

From a routing standpoint, IDS, like many state-aware engines, does not operate properly in an asymmetrically routed environment. Packets sent out from one set of routers and switches and returning through another will cause an IDS to see only half the traffic, causing false positives and false negatives.

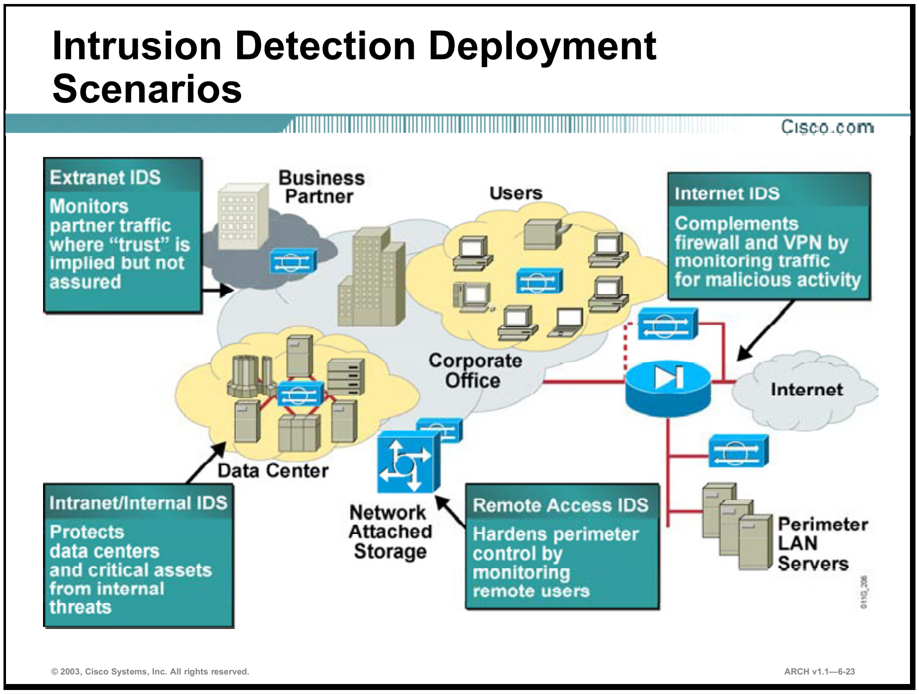

Intrusion Detection Deployment Scenarios

Cisco.com

Extranet IDS
Monitors partner traffic where “trust” is implied but not assured

Internet IDS
Complements firewall and VPN by monitoring traffic for malicious activity

Intranet/Internal IDS
Protects data centers and critical assets from internal threats

Remote Access IDS
Hardens perimeter control by monitoring remote users

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-23

Consider placing an IDS wherever there is a need to protect critical assets from the threat of intrusion. Network ingress points like the connections with the Internet and extranets are prime candidates, as are remote-access points. Also consider placing IDS internally at critical points to protect assets from internal threats.

Authentication, Authorization, and Accounting

Authentication, authorization, and accounting (AAA) is a term for a software mechanism to enhance network security by providing authentication services. This topic describes the software components of AAA.

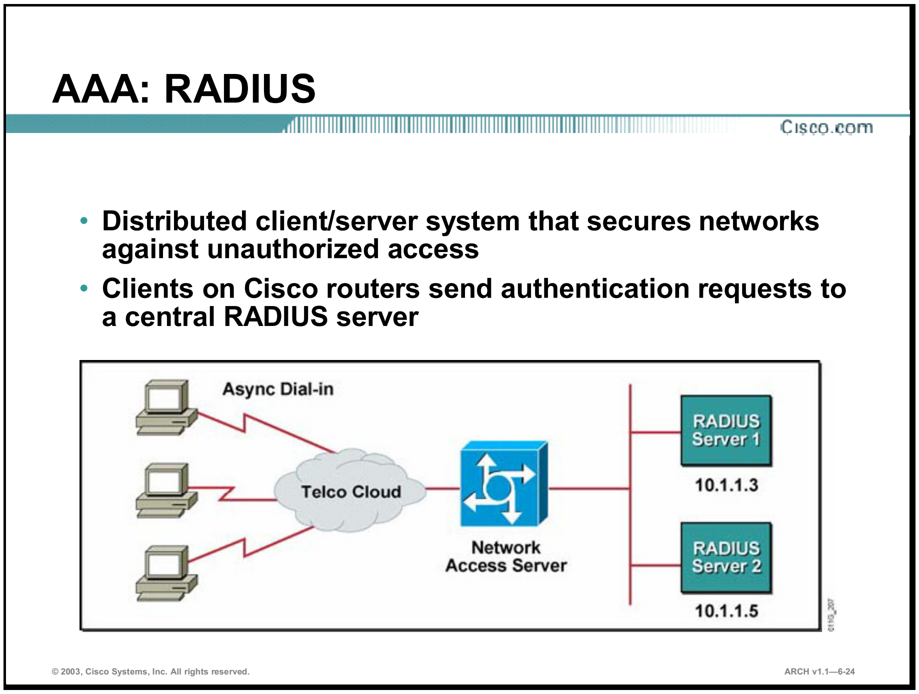

AAA: RADIUS

Cisco.com

- Distributed client/server system that secures networks against unauthorized access
- Clients on Cisco routers send authentication requests to a central RADIUS server

ARCH v1.1—6-24

The RADIUS protocol was developed by Livingston Enterprises, Inc., as an access server authentication and accounting protocol. The RADIUS authentication protocol is documented separately from the accounting protocol, but you can use the two together for a comprehensive solution.

A client/server model forms the basis for the RADIUS protocol. A network access server (NAS) such as a Cisco access server operates as a client of RADIUS. The client is responsible for passing user information to a designated RADIUS server and then acting on the response that is returned.

A RADIUS server (or daemon) can provide authentication and accounting services to one or more client NAS devices. RADIUS servers are responsible for receiving user connection requests, authenticating users, and then returning all configuration information necessary for the client to deliver service to the users. A RADIUS access server is generally a dedicated server connected to the network.

Communication between a NAS and a RADIUS server is based on the UDP. The authors of the RADIUS protocol selected UDP as the transport protocol for technical reasons. Generally, the RADIUS protocol is considered to be a connectionless service. The RADIUS-enabled devices, rather than the transmission protocol, handle issues related to server availability, retransmission, and timeouts.

Typically, a user login consists of a query (Access-Request) from the NAS to the RADIUS server and a corresponding response (Access-Accept or Access-Reject) from the server. The Access-Request packet contains the username, encrypted password, NAS IP address, and port. The format of the request also provides information on the type of session that the user wants to initiate.

Authentication is the most troublesome aspect of remote security because of the difficulty associated with positively identifying a user. To ensure the identity of a remote user, the RADIUS protocol supports several methods of authentication, including Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and token cards.

The RADIUS accounting functions allow data to be sent at the start and end of sessions, indicating the amount of resources (such as time, packets, bytes, and so on) used during the session.

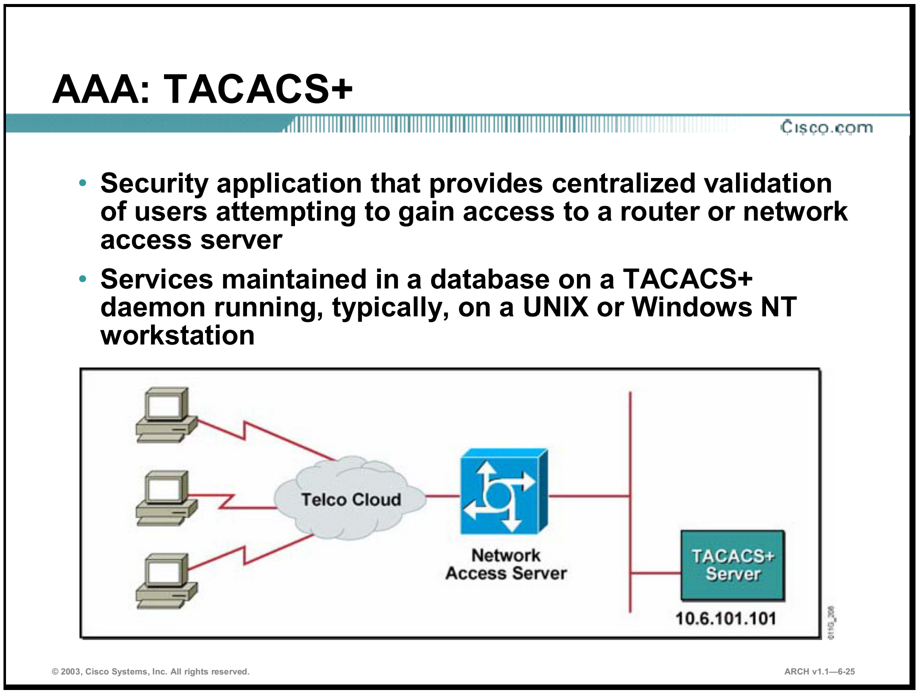

AAA: TACACS+

Cisco.com

- Security application that provides centralized validation of users attempting to gain access to a router or network access server
- Services maintained in a database on a TACACS+ daemon running, typically, on a UNIX or Windows NT workstation

ARCH v1.1—6-25

TACACS+ is a security application that provides centralized validation of users attempting to gain access to a router or NAS. TACACS+ services are maintained in a database on a TACACS+ daemon running, typically, on a UNIX or a Microsoft Windows NT server. You must have access to and must configure a TACACS+ server before the configured TACACS+ features on your NAS are available. TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide authentication, authorization, and accounting independently. You can tie each service into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

TACACS+ has three major components: the protocol support within the access servers and routers, the protocol specification, and the centralized security database. Similar to an internal security database, TACACS+ supports these features:

- **Authentication:** The TACACS+ protocol forwards many types of username and password information. This information is encrypted over the network with MD5, an encryption algorithm. TACACS+ can forward the password types for AppleTalk Remote Access (ARA), Serial Line Internet Protocol (SLIP), PAP, CHAP, and standard Telnet. This allows clients to use the same username and password for different protocols. TACACS+ is extensible to support new password types like KCHAP.
- **Authorization:** TACACS+ provides a mechanism to tell an access server which access list that a user connected to port 1 uses. The TACACS+ server and location of the username and password information identify the access list through which the user is filtered. The access list resides on the access server. The TACACS+ server responds to a username with an Accept and an access list number that causes that list to be applied.

- **Accounting:** TACACS+ provides accounting information to a database through TCP to ensure a more secure and complete accounting log. The accounting portion of the TACACS+ protocol contains the network address of the user, the username, the service attempted, protocol used, time and date, and the packet-filter module originating the log. For Telnet connections, it also contains source and destination port, action carried (communication accepted, rejected), log, and alert type. Formats are open and configurable. The billing information includes connect time, user ID, location connected from, start time, and stop time. It identifies the protocol that the user is using and may contain commands being run if the users are connected through EXEC and Telnet.

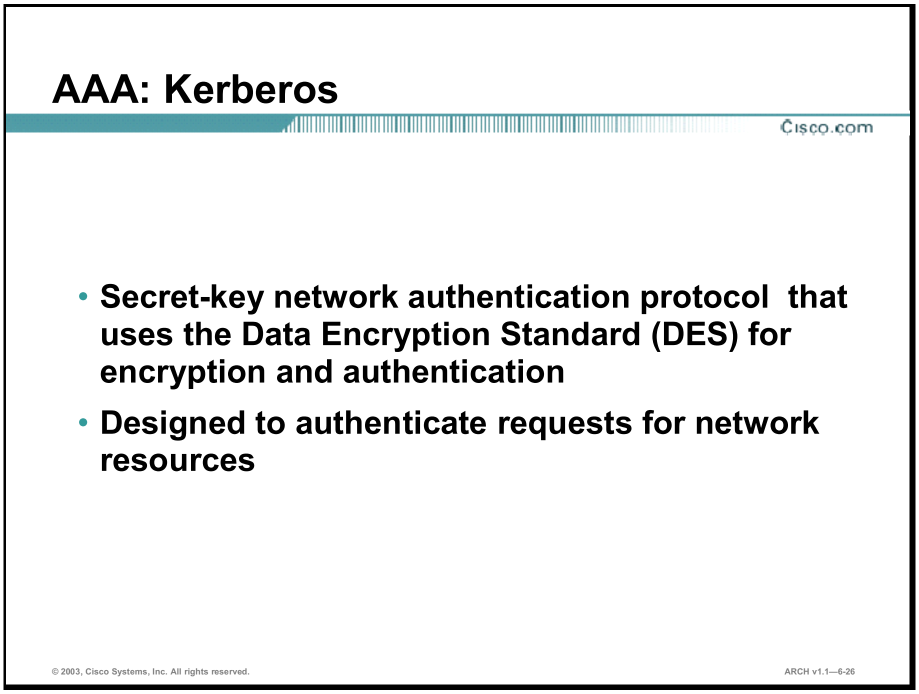

AAA: Kerberos

Cisco.com

- **Secret-key network authentication protocol that uses the Data Encryption Standard (DES) for encryption and authentication**
- **Designed to authenticate requests for network resources**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-26

Kerberos, developed at MIT, is a secret-key network authentication protocol that uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication. Kerberos was designed to authenticate requests for network resources. Kerberos, like other secret-key systems, is based on the concept of a trusted third party that performs secure verification of users and services. In the Kerberos protocol, this trusted third party is called the key distribution center (KDC). The primary use of Kerberos is to verify that users and the network services they use are really who and what they claim to be. To accomplish this, a trusted Kerberos server issues tickets to users. These tickets, which have a limited lifespan, are stored in a user's credential cache. You can use the tickets in place of the standard username and password authentication mechanism.

Kerberos is a trusted third-party authentication service. Each of its clients trusts Kerberos' judgment as to the identity of each of its other clients. Timestamps (large numbers representing the current date and time) have been added to the original model to aid in the detection of replay, which occurs when a message is stolen off the network and resent later.

Kerberos uses private key encryption. Each Kerberos principal is assigned a large number, its private key, known only to that principal and Kerberos. In the case of a user, the private key is the result of a one-way function applied to the user's password.

Because Kerberos knows these private keys, it can create messages that convince one client that another is really who it claims to be. Kerberos also generates temporary private keys, called session keys, which are given to two clients and no one else. A user or application can use a session key to encrypt messages between two parties.

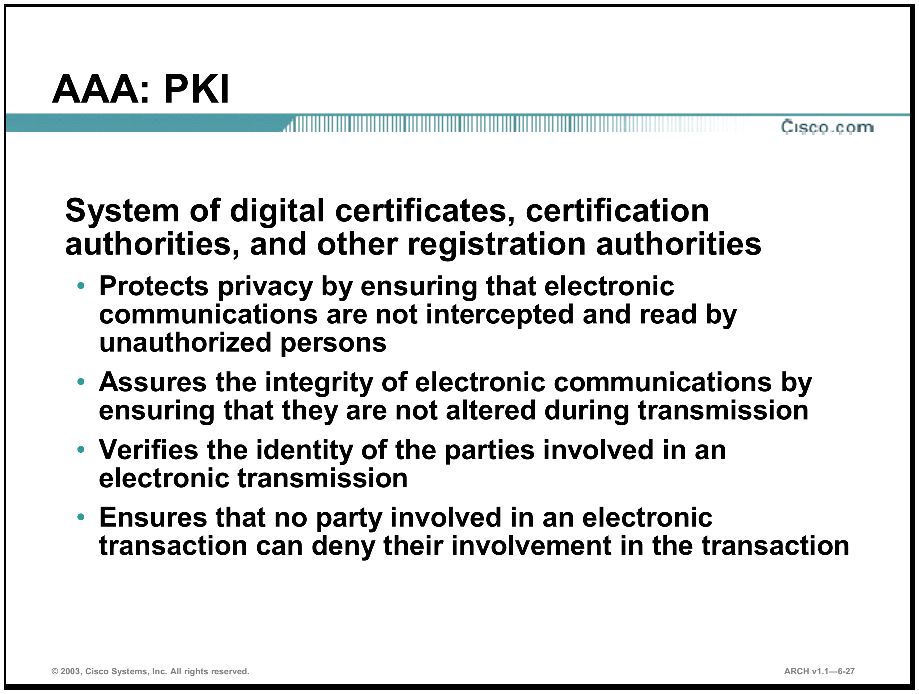

AAA: PKI

System of digital certificates, certification authorities, and other registration authorities

- Protects privacy by ensuring that electronic communications are not intercepted and read by unauthorized persons
- Assures the integrity of electronic communications by ensuring that they are not altered during transmission
- Verifies the identity of the parties involved in an electronic transmission
- Ensures that no party involved in an electronic transaction can deny their involvement in the transaction

ARCH v1.1—6-27

A PKI is a management system designed to administer asymmetrical cryptographic keys and public key certificates. It acts as a trusted component that guarantees the authenticity of the binding between a public key and security information, including identity, involved in securing a transaction with public-key cryptography.

A certificate is a cryptographically signed structure, called the digital certificate, which guarantees the association between at least one identifier and a public key. It is valid for a limited period of time (called the validity period), for a specific usage, and under certain conditions and limitations described in a certificate policy. The authority that issues this certificate is called the certification authority.

The initialization process consists of setting the necessary configuration for a PKI entity to communicate with other PKI entities. For example, the initialization of an end entity involves providing it with the public-key certificate of a trusted certification authority. The initialization of a certification authority involves the generation of its key pair.

During the registration process, an end entity makes itself known to a certification authority through a registration authority, before that certification authority issues a certificate. The end entity provides its name and other attributes to be included in its public key certificate, and the certification authority (or the registration authority, or both) verifies the correctness of the provided information.

The key pair generation for an end entity may either take place in its own environment or is done by the certification authority (or registration authority). If the key pair is not generated by the end entity itself, then the generated private key must be distributed to the end entity in a secure way (for example, through a secure key distribution protocol, or by using a physical token such as a smart card).

The certification process occurs at the certification authority. After verifying the correctness of the end entity's name and attributes, and that the end entity possesses the corresponding private key, the certification authority issues a certificate for the end entity's public key. That certificate is then returned to the end entity and/or posted in a repository where it is publicly available.

IP Security

IPSec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices. This topic describes the functionality of IPSec.

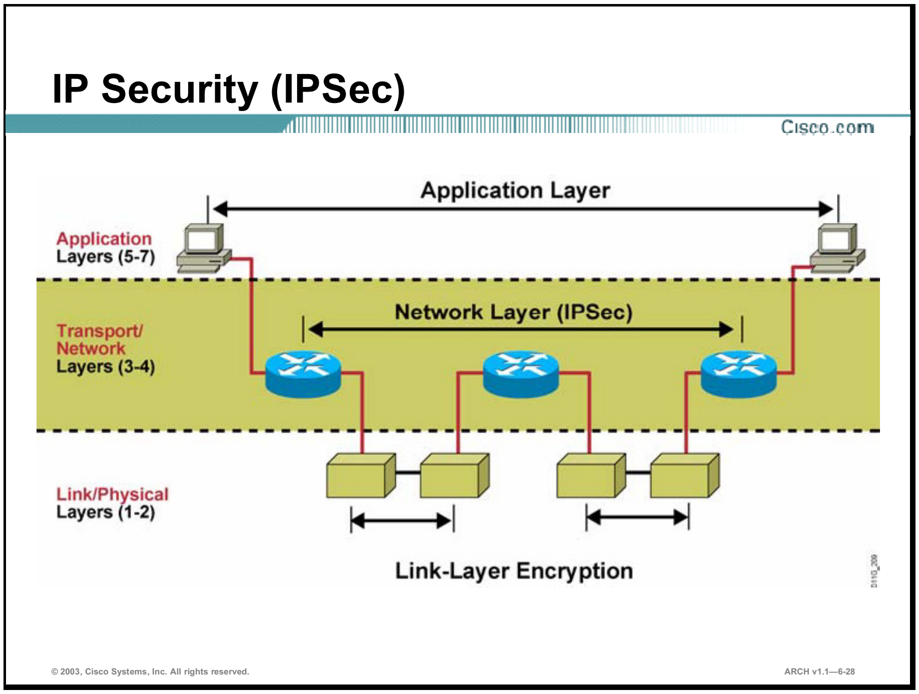

IP Security (IPSec)

Cisco.com

ARCH v1.1—6-28

IPSec is a framework of open standards developed by the Internet Engineering Task Force (IETF). IPSec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices (peers), such as routers.

IPSec provides these optional network security services, dictated by local security policy:

- **Data confidentiality**: The IPSec sender can encrypt packets before transmitting them across a network.
- **Data integrity**: The IPSec receiver can authenticate packets sent by the IPSec sender to ensure that the data has not been altered during transmission.
- **Data origin authentication**: The IPSec receiver can authenticate the source of the IPSec packets sent. This service is dependent upon the data integrity service.
- **Anti-replay**: The IPSec receiver can detect and reject replayed packets.

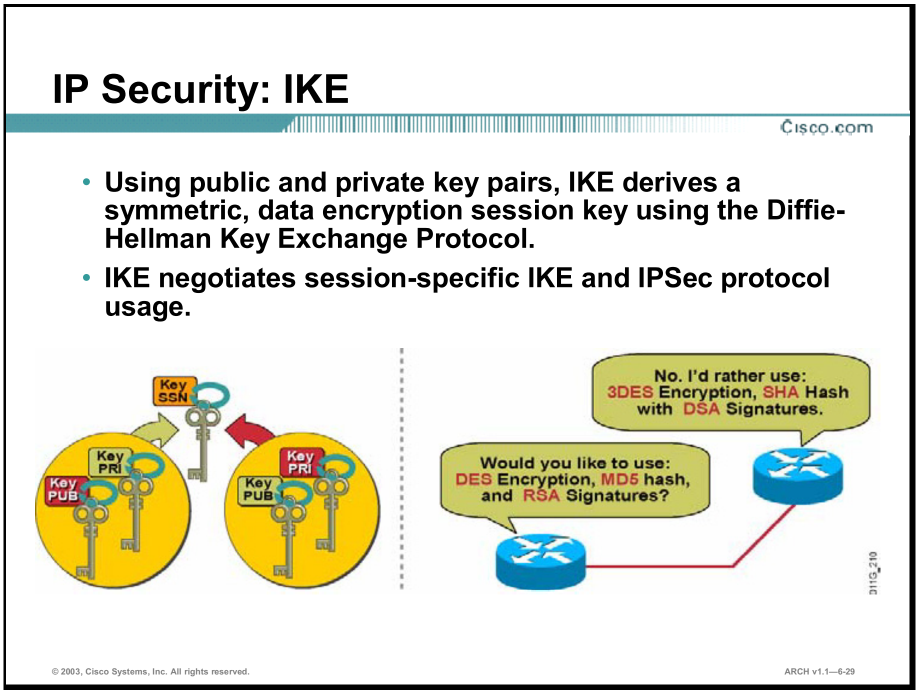

IP Security: IKE

- Using public and private key pairs, IKE derives a symmetric, data encryption session key using the Diffie-Hellman Key Exchange Protocol.
- IKE negotiates session-specific IKE and IPSec protocol usage.

Cisco.com
Key SSN
Key PRI
Key PUB
Key PRI
Key PUB
No. I'd rather use: 3DES Encryption, SHA Hash with DSA Signatures.
Would you like to use: DES Encryption, MD5 hash, and RSA Signatures?
D11G_210

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-29

IPSec offers a standard way of establishing authentication and encryption services between endpoints. This means not only standard algorithms and transforms, but also standard key negotiation and management mechanisms, via Internet Security Association and Key Management Protocol (ISAKMP) and Oakley, to promote interoperability between devices by allowing for the negotiation of services between these devices.

Negotiation refers to the establishment of policies or security associations (SAs) between devices. An SA is a policy rule that maps to a specific peer, with each rule identified by a unique security parameter index (SPI). A device may have many SAs stored in their security association database, created in dynamic random-access memory (DRAM) and indexed by SPI. As an IPSec datagram arrives, the device will use the enclosed SPI to reference the appropriate policy that needs to be applied to the datagram.

SAs are negotiated for both Internet Key Exchange (IKE) and IPSec, and it is IKE itself that facilitates this SA establishment.

IKE is a form of ISAKMP/Oakley specifically for IPSec. ISAKMP describes the phase of negotiation; Oakley defines the method to establish an authenticated key exchange. This method may take various modes of operation and is also used to derive keying material via algorithms such as Diffie-Hellman.

ISAKMP Phase 1 is used when two peers establish a secure, authenticated channel with which to communicate. Oakley main mode is generally used here. The result of main mode is the authenticated bidirectional IKE security association and its keying material. ISAKMP Phase 2 is required to establish SAs on behalf of other services such as IPSec, which needs key material or parameter negotiation, and uses Oakley quick mode. The result of quick mode is two to four (depending on whether Authentication Header [AH] or Encapsulating Security Payload [ESP] was used) unidirectional IPSec security associations and their keying material.

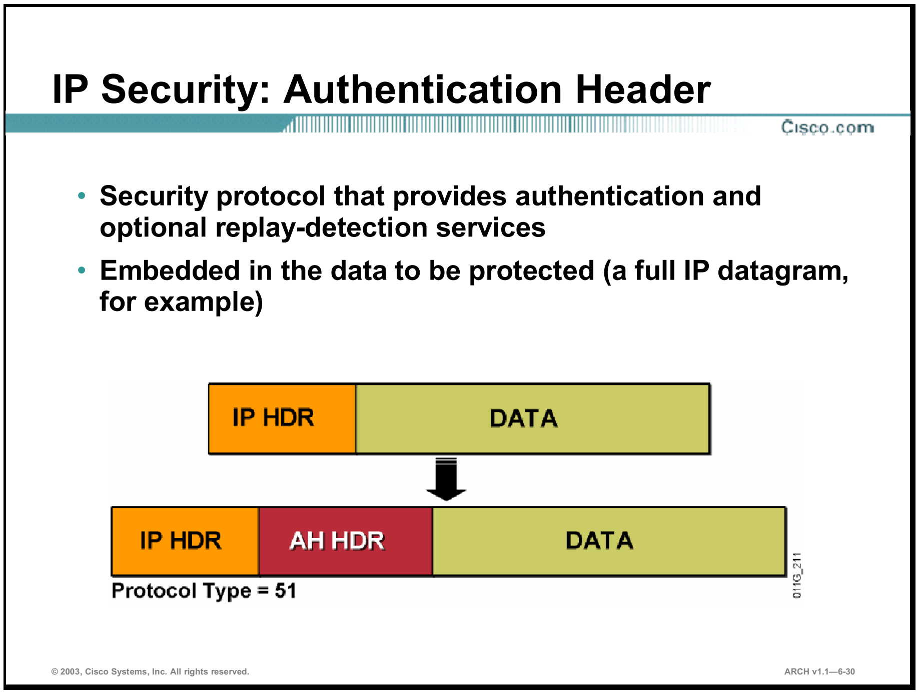

IP Security: Authentication Header

- Security protocol that provides authentication and optional replay-detection services
- Embedded in the data to be protected (a full IP datagram, for example)

ARCH v1.1—6-30

The AH is a mechanism for providing strong integrity and authentication for IP datagrams. It might also provide nonrepudiation, depending on which cryptographic algorithm is used and how keying is performed. When using AH, the data is not encrypted.

The AH may appear after any other headers that are examined at each hop, and before any other headers that are not examined at an intermediate hop. The IPv4 or IPv6 header immediately preceding the AH will contain the value 51 in its Next Header (or Protocol) field.

Using AH is resource-intensive in terms of bandwidth and the networking device.

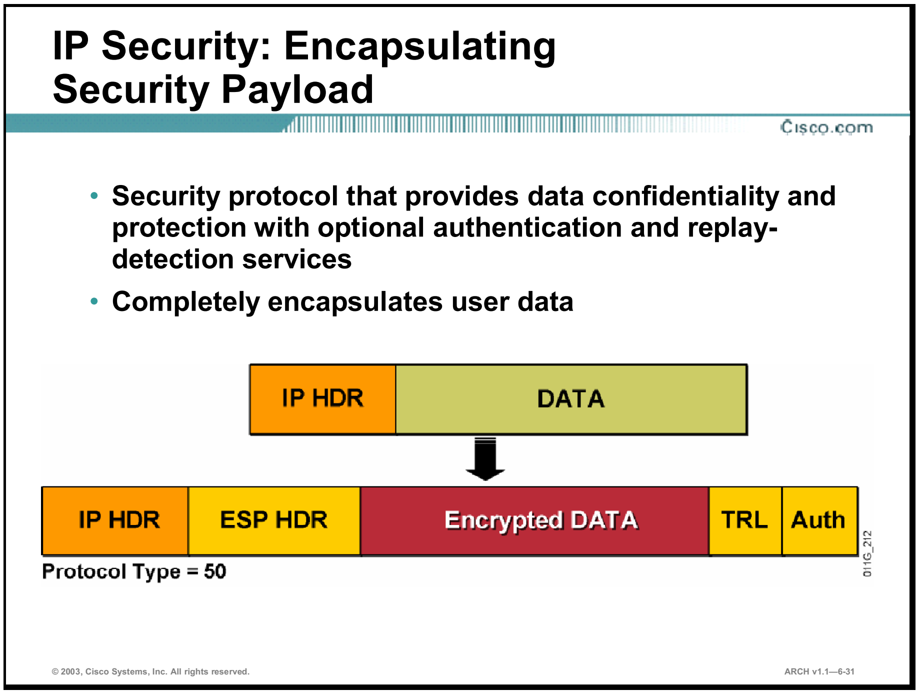

IP Security: Encapsulating Security Payload

Cisco.com

- Security protocol that provides data confidentiality and protection with optional authentication and replay-detection services
- Completely encapsulates user data

IP HDR
DATA
IP HDR
ESP HDR
Encrypted DATA
TRL
Auth
Protocol Type = 50
011G_212

ARCH v1.1—6-31

The ESP may appear anywhere after the IP header and before the final transport-layer protocol. The Internet Assigned Numbers Authority (IANA) has assigned Protocol Number 50 to ESP. The IP ESP seeks to provide confidentiality and integrity by encrypting data to be protected and placing the encrypted data in the data portion of the IP Encapsulating Security Payload. In tunnel-mode ESP, the original IP datagram is placed in the encrypted portion of the ESP and that entire ESP frame is placed within a datagram having unencrypted IP headers. The information in the unencrypted IP headers is used to route the secure datagram from origin to destination. In transport-mode ESP, the ESP header is inserted into the IP datagram immediately before the transport-layer protocol header, such as TCP, UDP, or ICMP. In this mode, bandwidth is conserved because there are no encrypted IP headers or IP options.

In tunnel-mode ESP, the original IP datagram is placed in the encrypted portion of the ESP and that entire ESP frame is placed within a datagram having unencrypted IP headers. The information in the unencrypted IP headers is used to route the secure datagram from origin to destination. An unencrypted IP routing header might be included between the IP header and the ESP.

Using ESP is resource-intensive in terms of bandwidth and the networking device.

Device Security Options

To secure a network, the individual components that make up the network must be secure. You can take actions to ensure security specific to routers, switches, hosts, applications, and the network as a whole. This topic describes security options for the specific components of a network.

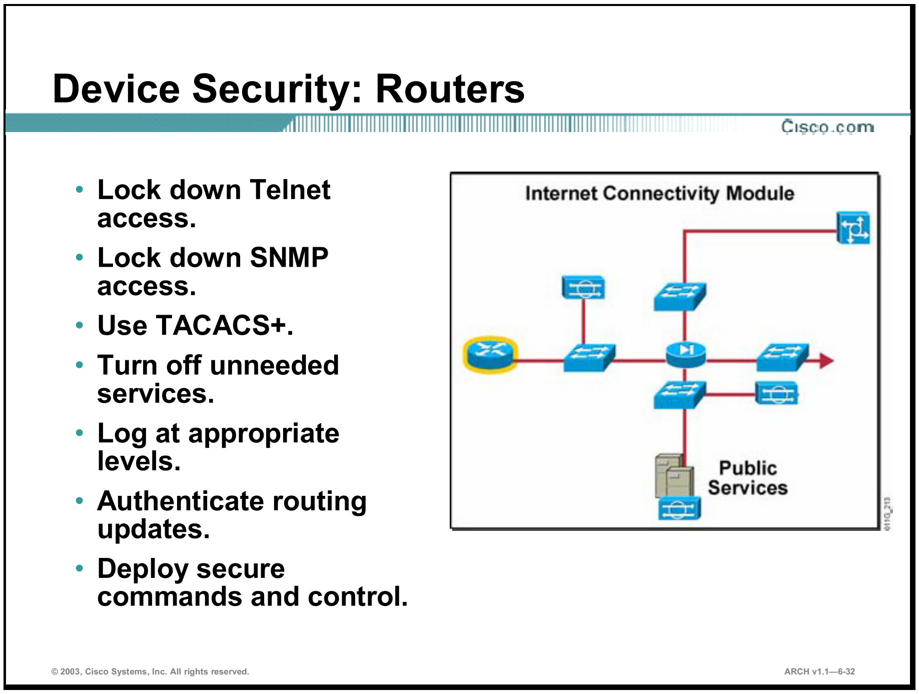

Device Security: Routers

Cisco.com

- Lock down Telnet access.
- Lock down SNMP access.
- Use TACACS+.
- Turn off unneeded services.
- Log at appropriate levels.
- Authenticate routing updates.
- Deploy secure commands and control.

ARCH v1.1—6-32

Routers control access from network to network. They advertise networks and filter who can use them, and they are a huge potential threat for a hacker. Router security is a critical element in any security deployment. By their nature, routers provide access and, therefore, you should secure them to reduce the likelihood that they will be directly compromised. There are several tasks that you can complete to secure a router:

- Lock down Telnet access.
- Lock down SNMP access.
- Control access through the use of TACACS+.
- Turn off unneeded services.
- Log at appropriate levels.
- Authenticate routing updates.
- Deploy secure commands and control.

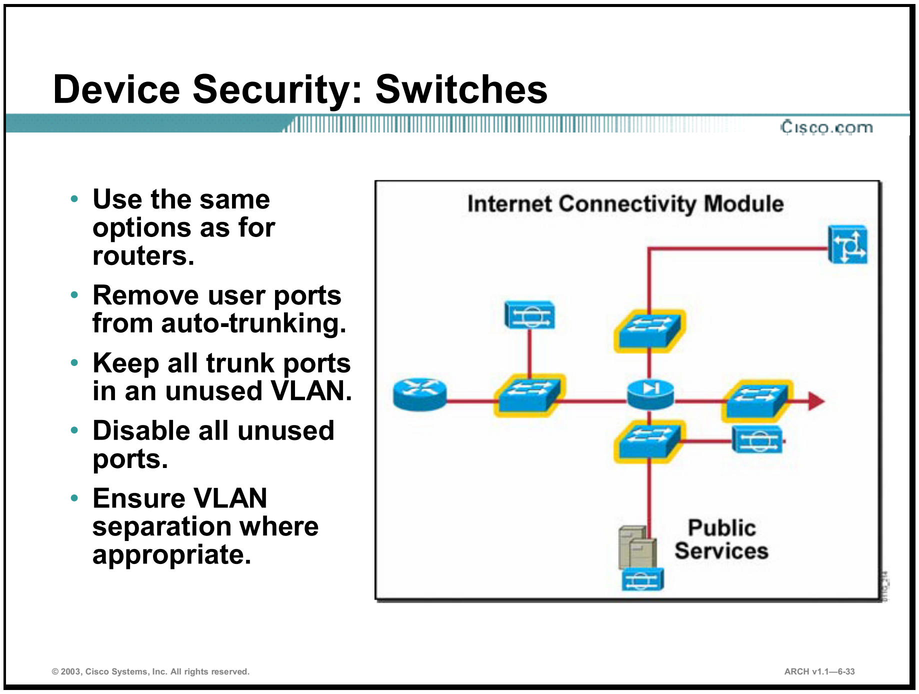

Device Security: Switches

Cisco.com

- Use the same options as for routers.
- Remove user ports from auto-trunking.
- Keep all trunk ports in an unused VLAN.
- Disable all unused ports.
- Ensure VLAN separation where appropriate.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-33

Like routers, data link layer switches and multilayer switches have their own set of security considerations. Most of the security techniques that apply to routers also apply to switches. In addition, consider taking these precautions:

- **Use the same options as for routers.** For example, lock down Telnet and SNMP access, use TACACS+, turn off unneeded services, log at appropriate levels, and deploy secure commands and control.
- **Remove user ports from auto-trunking.** For ports without any need to trunk, set any trunk settings to off, as opposed to auto. This setup prevents a host from becoming a trunk port and receiving all traffic that would normally reside on a trunk port.
- **Keep all trunk ports in an unused LAN.** If you are using older versions of software for your Ethernet switch, make sure that trunk ports use a virtual LAN (VLAN) number not used anywhere else in the switch. This setup prevents packets tagged with the same VLAN as the trunk port from reaching another VLAN without crossing a Layer 3 device.
- **Disable all unused ports on a switch.** This setup prevents hackers from plugging in to unused ports and communicating with the rest of the network.
- **Ensure VLAN separation where appropriate.** Avoid using VLANs as the sole method of securing access between two subnets. The capability for human error, combined with the understanding that VLANs and VLAN tagging protocols were not designed with security in mind, makes their use in sensitive environments inadvisable. When VLANs are needed in security deployments, be sure to pay close attention to the security configurations and guidelines. Within an existing VLAN, private VLANs provide some added security to specific network applications.

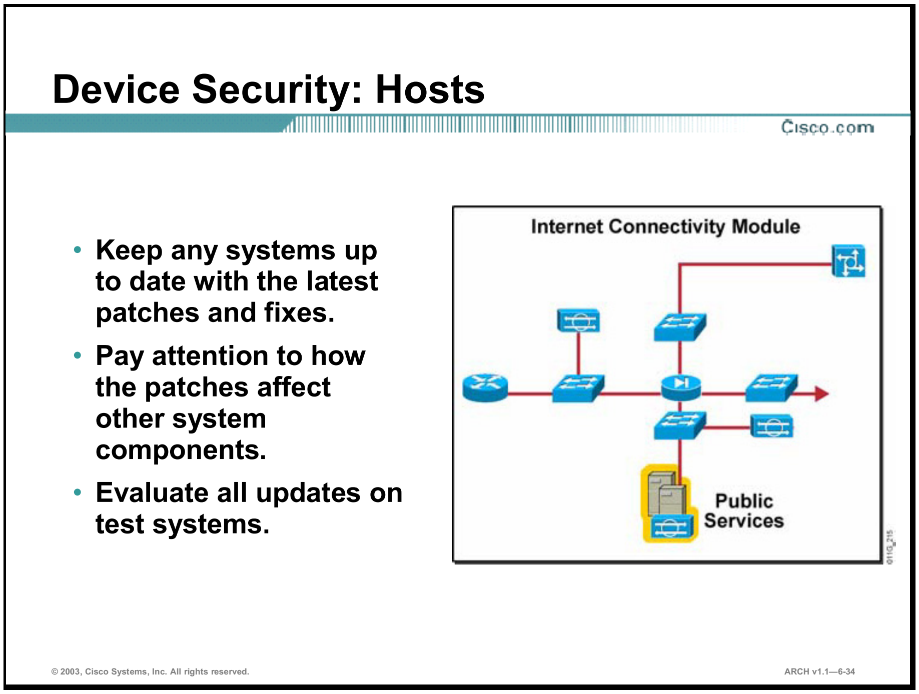

Device Security: Hosts

Cisco.com

- Keep any systems up to date with the latest patches and fixes.
- Pay attention to how the patches affect other system components.
- Evaluate all updates on test systems.

ARCH v1.1—6-34

As the most likely target during an attack, the host presents some of the most difficult challenges from a security perspective. There are numerous hardware platforms, operating systems, and applications, all of which have updates, patches, and fixes available at different times. Because hosts provide the application services to other hosts that request them, they are extremely visible within the network. Because of the visibility, hosts are the most frequently attacked devices in any network intrusion attempt.

To secure hosts, pay careful attention to each of the components within the systems. Keep any systems up to date with the latest patches and fixes. Pay attention to how these patches affect the operation of other system components. Evaluate all updates on test systems before you implement them in a production environment. Failure to do so might result in the patch itself causing a denial of service.

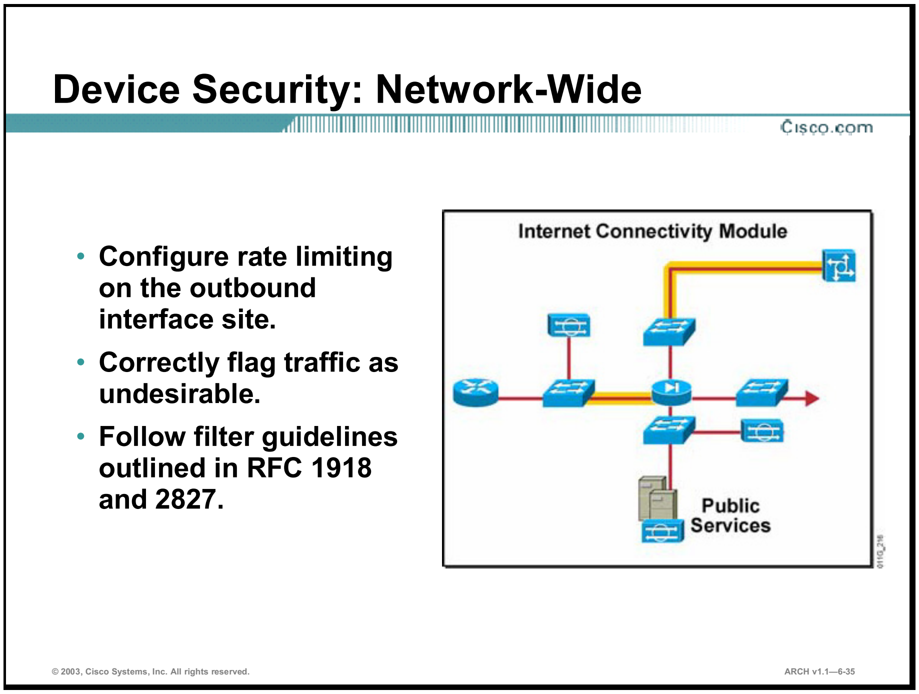

Device Security: Network-Wide

- Configure rate limiting on the outbound interface site.
- Correctly flag traffic as undesirable.
- Follow filter guidelines outlined in RFC 1918 and 2827.

ARCH v1.1—6-35

Network attacks are among the most difficult attacks to deal with because they typically take advantage of an intrinsic characteristic in the way your network operates. These attacks include Address Resolution Protocol (ARP) and MAC-based Layer 2 attacks, sniffers, and distributed denial-of-service attacks. You can mitigate some of the ARP and MAC-based Layer 2 attacks through best practices on switches and routers.

Distributed denial-of-service attacks work by causing tens or hundreds of machines to simultaneously send spurious data to an IP address. Only through cooperation with its Internet service provider (ISP) can an e-commerce company hope to thwart such an attack. An ISP can configure rate limiting on the outbound interface to the company’s site. This rate limiting can drop most undesired traffic when it exceeds a prespecified amount of the available bandwidth. The key is to correctly flag traffic as undesired.

Common forms of distributed denial-of-service attacks are ICMP floods, TCP SYN floods, or UDP floods. In an e-commerce environment, this type of traffic is fairly easy to categorize. Only when limiting a TCP SYN attack on port 80 (HTTP) does an administrator run the risk of locking out legitimate users during an attack. Even then, it is better to temporarily lock out new legitimate users and retain routing and management connections than to have the router overrun and lose all connectivity. One approach to limiting this sort of attack is to follow filtering guidelines for networks outlined in RFC 1918 and RFC 2827.

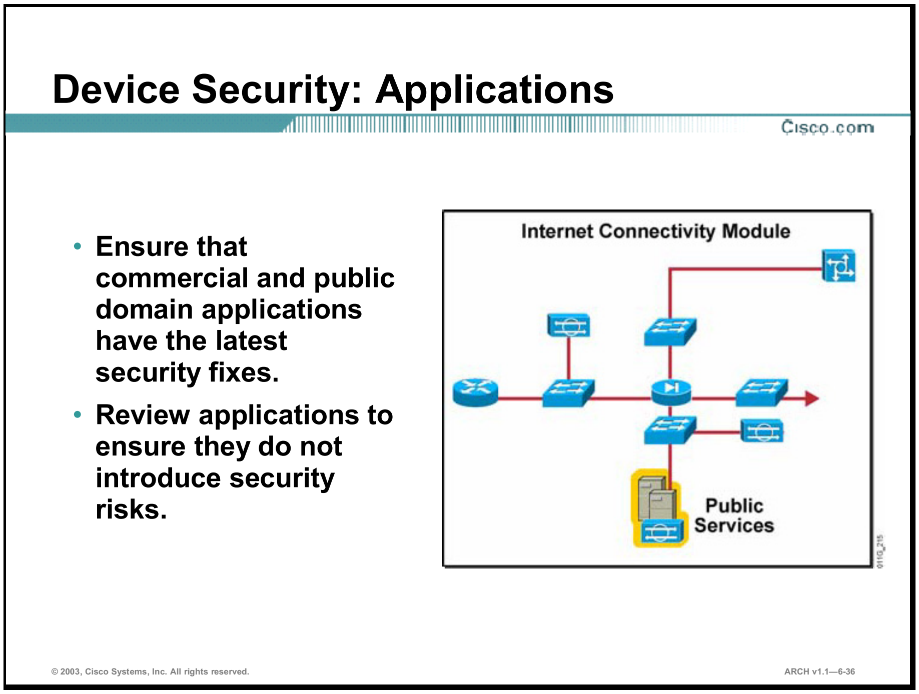

Device Security: Applications

Cisco.com

- Ensure that commercial and public domain applications have the latest security fixes.
- Review applications to ensure they do not introduce security risks.

Internet Connectivity Module
Public Services

ARCH v1.1—6-36

Applications are usually coded by human beings and, as such, are subject to numerous errors. These errors can be benign, such as an error that causes your document to print incorrectly, or malignant, such as an error that makes the credit card numbers on your database server available via anonymous FTP. Care needs to be taken to ensure that commercial and public domain applications are up to date with the latest security fixes. Public domain applications, as well as custom-developed applications, also require code review to ensure that the applications are not introducing any security risks caused by poor programming. This programming can include scenarios such as how an application makes calls to other applications or the operating system itself, the privilege level at which the application runs, the degree of trust that the application has for the surrounding systems, and finally, the method the application uses to transport data across the network.

Summary

This topic summarizes the key points discussed in this lesson.

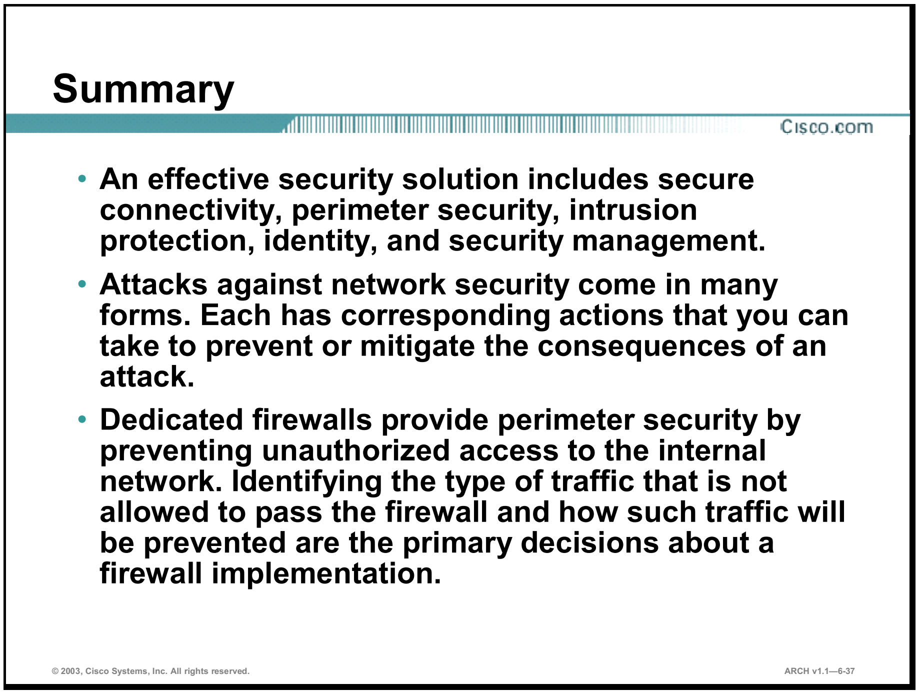

Summary

- An effective security solution includes secure connectivity, perimeter security, intrusion protection, identity, and security management.
- Attacks against network security come in many forms. Each has corresponding actions that you can take to prevent or mitigate the consequences of an attack.
- Dedicated firewalls provide perimeter security by preventing unauthorized access to the internal network. Identifying the type of traffic that is not allowed to pass the firewall and how such traffic will be prevented are the primary decisions about a firewall implementation.

ARCH v1.1—6-37

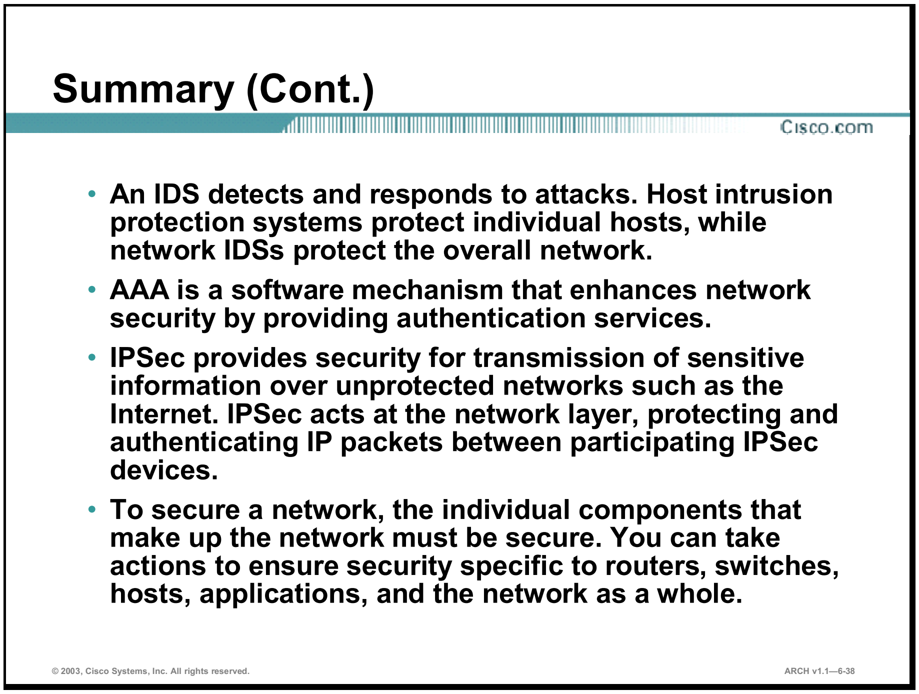

Summary (Cont.)

- An IDS detects and responds to attacks. Host intrusion protection systems protect individual hosts, while network IDSs protect the overall network.
- AAA is a software mechanism that enhances network security by providing authentication services.
- IPSec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec devices.
- To secure a network, the individual components that make up the network must be secure. You can take actions to ensure security specific to routers, switches, hosts, applications, and the network as a whole.

ARCH v1.1—6-38

References

For additional information, refer to these resources:

- *Network Security* at http://www.cisco.com/warp/public/779/largeent/issues/security/
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

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 technology enables secure connectivity?

A) RADIUS

B) digital signature

C) digital encryption

D) stateful firewalling

Q2) What is the only effective means of mitigating man-in-the-middle attacks?

A) cryptography

B) authentication

C) access control

D) switched infrastructure

Q3) What is the purpose of a perimeter LAN?

A) to allow internal access to some parts of the Internet

B) to allow external access to some parts of the internal network

C) to allow internal access to some external services and provide access to the Internet

D) to allow external access to some services without providing access to the internal network

Q4) What is the primary function of the perimeter firewall?

A) to prevent trust exploitation

B) to prevent denial of service attacks

C) to prevent unauthorized access attacks

D) to prevent virus and Internet worm access

Q5) How does a Cisco network IDS determine that traffic is an attack?

A) by comparing traffic to an attack signature

B) by inserting agents into the host to be protected

C) by intercepting operating system and application calls

D) by watching all packets traversing a particular collision domain

Q6) Which AAA service provides for separate and modular authentication, authorization, and accounting facilities?

A) PKI

B) Kerberos

C) RADIUS

D) TACACS+

Q7) Which authentication protocol is embedded into data?

A) AH

B) IKE

C) ESP

D) IPSec

Q8) What action can you take on a switch to prevent a host from becoming a trunk port?

A) Disable all unused ports.

B) Turn off unneeded services.

C) Turn off auto-trunking on all unnecessary ports.

D) Make sure that trunk ports use a unique VLAN number.

Quiz Answer Key

Q1) C

Relates to: Key Elements of Network Security

Q2) A

Relates to: Network Security Attacks and Solutions

Q3) D

Relates to: Firewall Design Options

Q4) C

Relates to: Firewall Design Options

Q5) A

Relates to: Intrusion Detection System Design Options

Q6) D

Relates to: Authentication, Authorization, and Accounting

Q7) A

Relates to: IP Security

Q8) C

Relates to: Device Security Options

Implementing Network Security Using the Cisco SAFE Security Blueprints

Overview

The Cisco Security Architecture for Enterprise (SAFE) Blueprints are frameworks for security and are based on Cisco Architecture for Voice, Video and Integrated Data (AVVID).

SAFE serves as a guide for network designers considering the security requirements of their network. It takes a defense-in-depth approach to network security design. This type of design focuses on the expected threats and their methods of mitigation based on the best practices that Cisco has developed. The SAFE strategy results in a layered approach to security where the failure of one security system is not likely to lead to the compromise of network resources.

Relevance

The guidelines in this lesson do not guarantee a secure environment, or that a designer will prevent all intrusions. However, designers will achieve reasonable security by establishing a good security policy, following the best practices outlined in this lesson, staying up to date on the latest developments in the hacker and security communities, and maintaining and monitoring all systems using sound system administration practices.

Objectives

Upon completing this lesson, you will be able to propose a security strategy for the Enterprise Campus and the Enterprise Edge functional areas using the SAFE Blueprint, given specific security requirements. This includes being able to meet these objectives:

- Describe the SAFE architecture and explain how it meets enterprise security needs
- Propose a security strategy for each component of the Enterprise Campus functional area in a small network, given specific security requirements
- Propose a security strategy for each component of the Enterprise Campus functional area in a medium network, given specific security requirements
- Propose a security strategy for each component of the Enterprise Campus functional area in a large network, given specific security requirements
- Propose a security strategy for the Enterprise Edge functional area, given specific security requirements

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course
- Evaluating Network Security Policies lesson
- Reviewing Cisco Security Solutions lesson

Outline

The outline lists the topics included in this lesson.

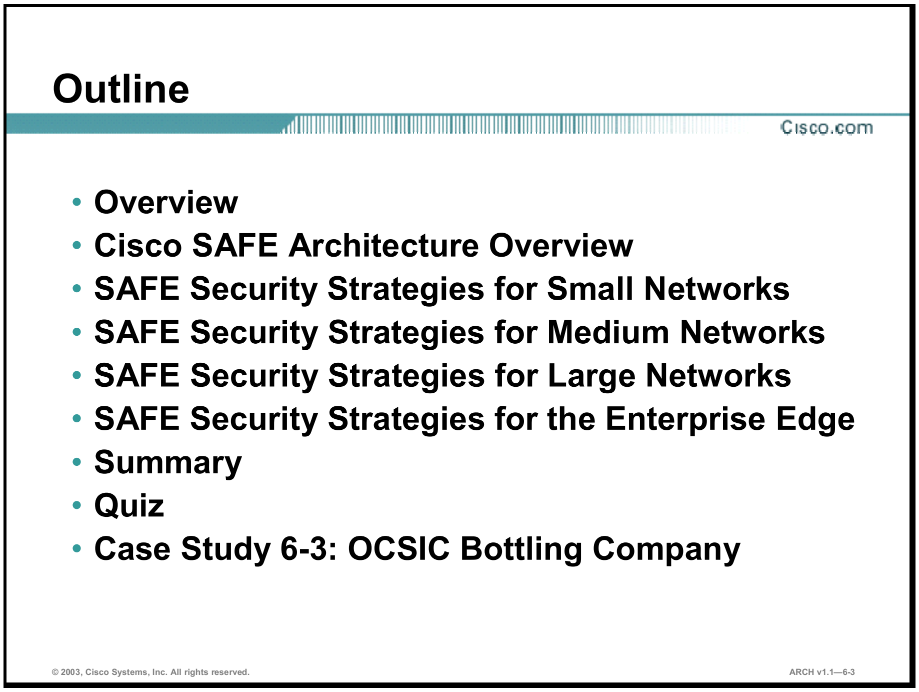

Outline

- **Overview**
- **Cisco SAFE Architecture Overview**
- **SAFE Security Strategies for Small Networks**
- **SAFE Security Strategies for Medium Networks**
- **SAFE Security Strategies for Large Networks**
- **SAFE Security Strategies for the Enterprise Edge**
- **Summary**
- **Quiz**
- **Case Study 6-3: OCSIC Bottling Company**

ARCH v1.1—6-3

Cisco SAFE Architecture Overview

The principal goal of the Cisco SAFE Blueprint for enterprise networks is to provide best practices information to interested parties on designing and implementing secure networks. SAFE serves as a guide to network designers considering the security requirements of their network. This topic describes the Cisco SAFE architecture and design objectives.

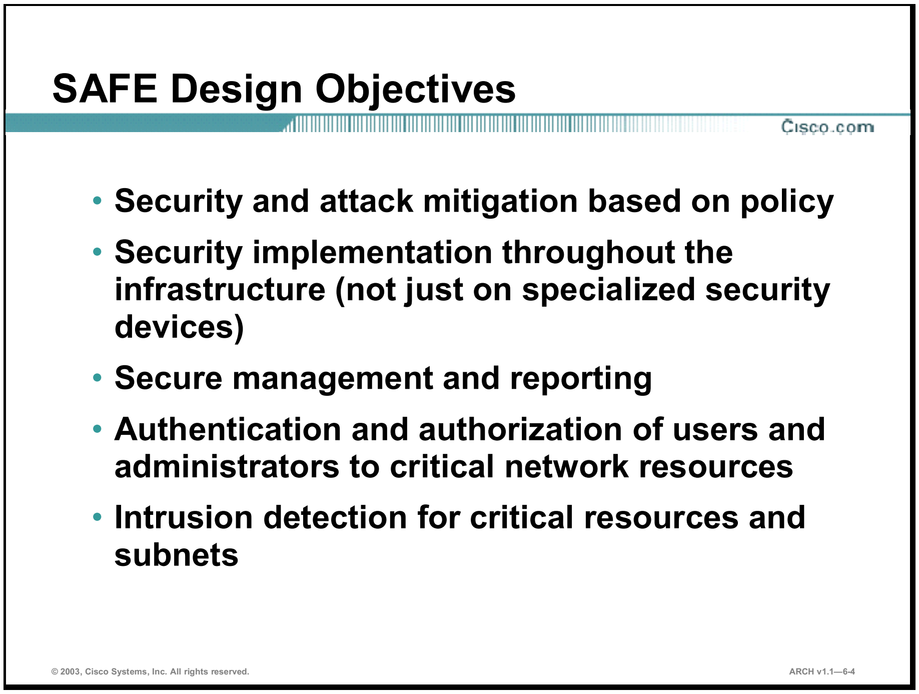

SAFE Design Objectives

- **Security and attack mitigation based on policy**
- **Security implementation throughout the infrastructure (not just on specialized security devices)**
- **Secure management and reporting**
- **Authentication and authorization of users and administrators to critical network resources**
- **Intrusion detection for critical resources and subnets**

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-4

SAFE emulates as closely as possible the functional requirements of today’s networks. Implementation decisions vary, depending on the network functionality required. However, these design objectives, listed in order of priority, guide the decision-making process:

- Security and attack mitigation based on policy
- Security implementation through the infrastructure (not just on specialized security devices)
- Secure management and reporting
- Authentication and authorization of users and administrators to critical network resources
- Intrusion detection for critical resources and subnets

First and foremost, SAFE is a security architecture. It must prevent most attacks from successfully affecting valuable network resources. The attacks that succeed in penetrating the first line of defense, or originate from inside the network, must be accurately detected and quickly contained to minimize their effect on the rest of the network. However, in being secure, the network must continue to provide critical services that users expect. Proper network security and good network functionality can be provided at the same time. The SAFE architecture is not a revolutionary way of designing networks, but merely a blueprint for making networks secure.

SAFE is resilient and scalable. Resilience in networks includes physical redundancy to protect against a device failure, whether through misconfiguration, physical failure, or network attack. The SAFE architecture for small, midsize, and remote networks was designed without resiliency, due to the cost-effectiveness and limited complexity of smaller designs.

At many points in the network design process, you need to choose between using integrated functionality in a network device and using a specialized functional appliance. The integrated functionality is often attractive because you can implement it on existing equipment, and because the features can interoperate with the rest of the device to provide a better functional solution. Appliances are often used when the depth of functionality required is very advanced or when performance needs require using specialized hardware. Make your decisions based on the capacity and functionality of the appliance rather than on the integration advantage of the device. For example, sometimes you can choose an integrated higher-capacity IOS router with IOS firewall software as opposed to a smaller IOS router with a separate firewall. When design requirements do not dictate a specific choice, you should choose to go with integrated functionality in order to reduce the overall cost of the solution.

Although most networks evolve with the growing IT requirements of an organization, the SAFE architecture uses a modular approach. A modular approach has two main advantages: First, it allows the architecture to address the security relationship between the various functional blocks of the network. Second, it permits designers to evaluate and implement security on a module-by-module basis, instead of attempting the complete architecture in a single phase. The security design of each module is described separately, but is validated as part of the complete design.

Although it is true that most networks cannot be easily dissected into clear-cut modules, this approach provides a guide for implementing different security functions throughout the network. The authors do not expect network engineers to design their networks identical to the SAFE implementation, but rather to use a combination of the modules described and integrate them into the existing network.

SAFE Security Strategies for Small Networks

The SAFE design for a small network includes only an Internet Connectivity module that provides access to the external network, and the Campus Infrastructure module containing the internal network. This topic describes the Cisco SAFE architecture and security strategies for a small network.

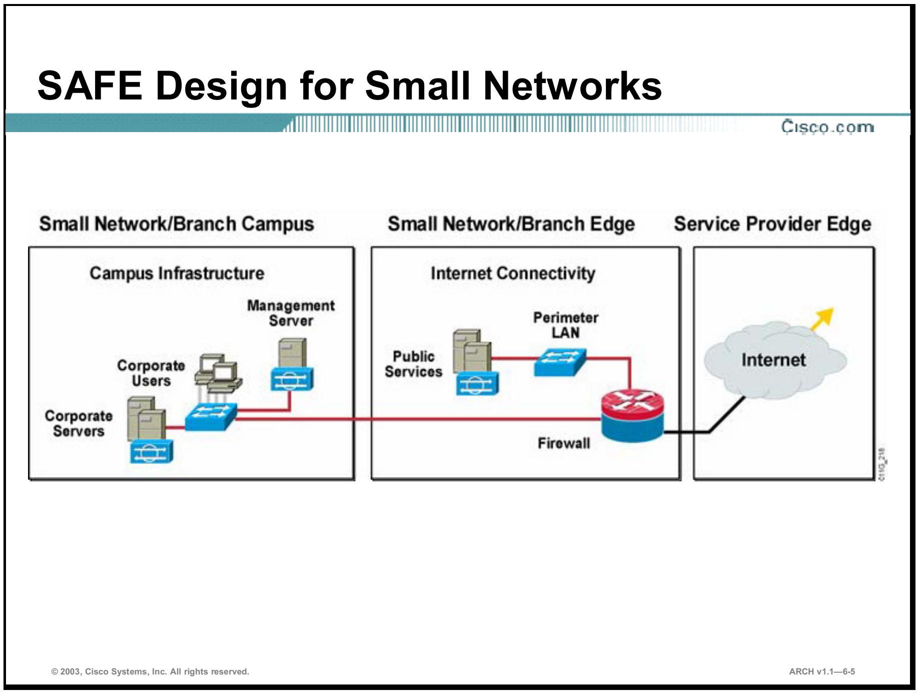

SAFE Design for Small Networks

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-5

The small network design has two modules: the Internet Connectivity module and the Campus Infrastructure module. The Internet Connectivity module has connections to the Internet and terminates VPN and public services (DNS, HTTP, FTP, Simple Mail Transfer Protocol [SMTP]) traffic. The Campus Infrastructure module contains the data link layer switching and all the users, as well as the management servers and intranet servers.

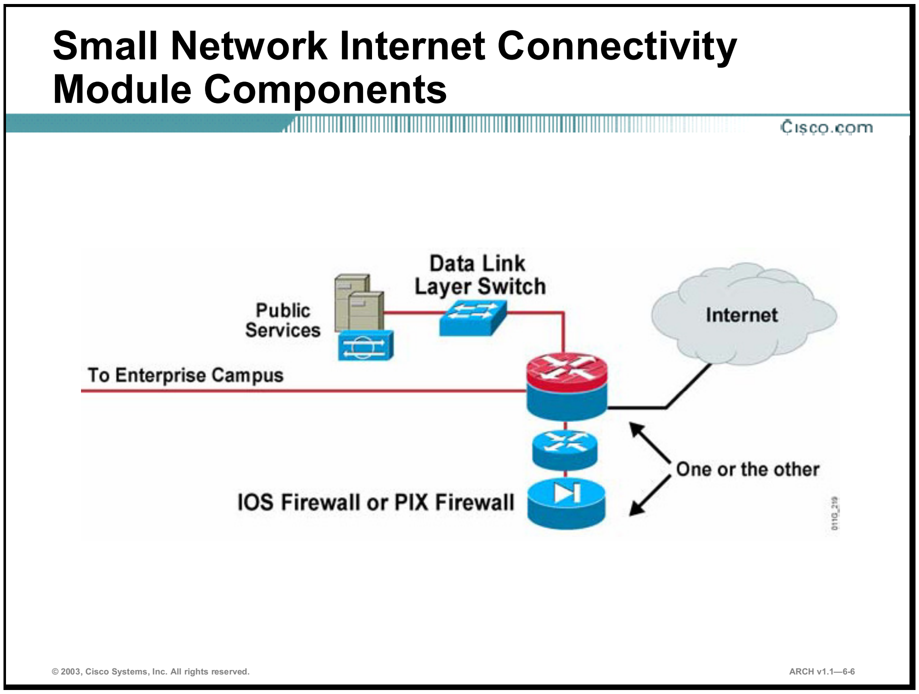

Small Network Internet Connectivity Module Components

Cisco.com

ARCH v1.1—6-6

The Internet Connectivity module provides internal users with connectivity to Internet services and Internet users access to information on public servers. This module is not designed to serve e-commerce type applications.

The primary devices included in this module are:

- **SMTP server:** Acts as a relay between the Internet and the intranet mail servers
- **DNS server:** Serves as authoritative external DNS server for the enterprise; relays internal DNS requests to the Internet
- **FTP/HTTP server:** Provides public information about the organization
- **Firewall or firewall router:** Provides network-level protection of resources, stateful filtering of traffic, and VPN termination for remote sites and users
- **Data link layer switch (with private VLAN support):** Ensures that data from managed devices can only cross directly to the IOS firewall
- **HIPS:** Provides host-level intrusion detection

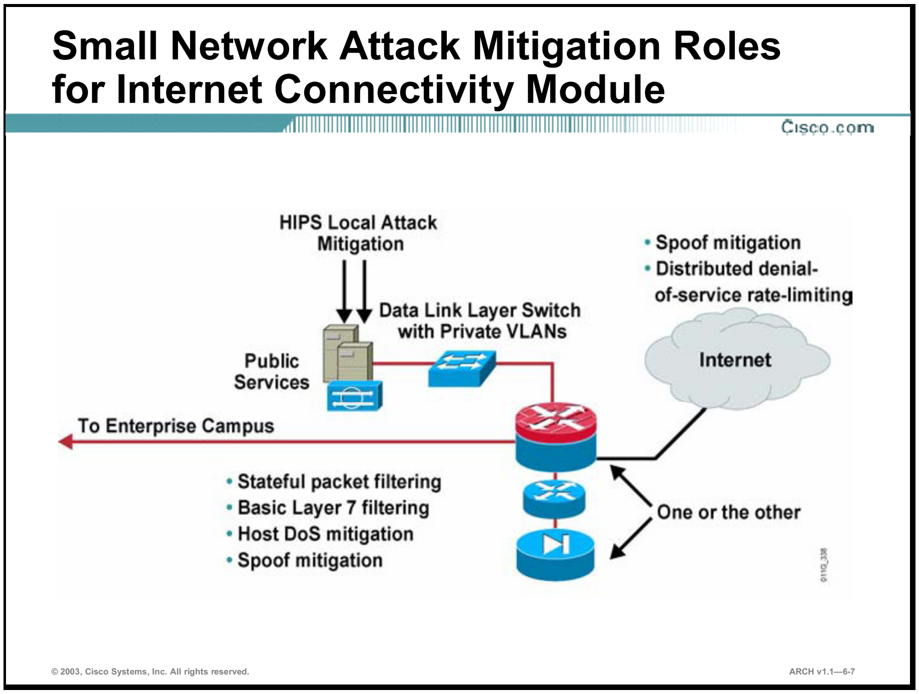

Small Network Attack Mitigation Roles for Internet Connectivity Module

Cisco.com

HIPS Local Attack Mitigation

Data Link Layer Switch with Private VLANs

Public Services

- Spoof mitigation
- Distributed denial-of-service rate-limiting

Internet

To Enterprise Campus

- Stateful packet filtering
- Basic Layer 7 filtering
- Host DoS mitigation
- Spoof mitigation

One or the other

0110_338

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-7

The Internet Connectivity module includes publicly addressable servers that are the most likely points of attack. The expected threats to this module and the security features used to mitigate their effect are:

- **Unauthorized access:** Mitigated through filtering at the firewall
- **Application layer attacks:** Mitigated through HIPS on the public servers
- **Virus and Trojan horse attacks:** Mitigated through virus scanning at the host level
- **Password attacks:** Limited services available to brute-force attack; operating systems and IDSs can detect the threat
- **Denial of service:** Committed access rate (CAR) at ISP edge and TCP setup controls at firewall to limit exposure
- **IP spoofing:** RFC 2827 and 1918 filtering at ISP edge and local firewall
- **Packet sniffers:** Switched infrastructure and HIPS to limit exposure
- **Network reconnaissance:** HIPS detects reconnaissance; protocols filtered to limit effectiveness
- **Trust exploitation:** Restrictive trust model and private VLANs to limit trust-based attacks
- **Port redirection:** Restrictive filtering and HIPS to limit attack

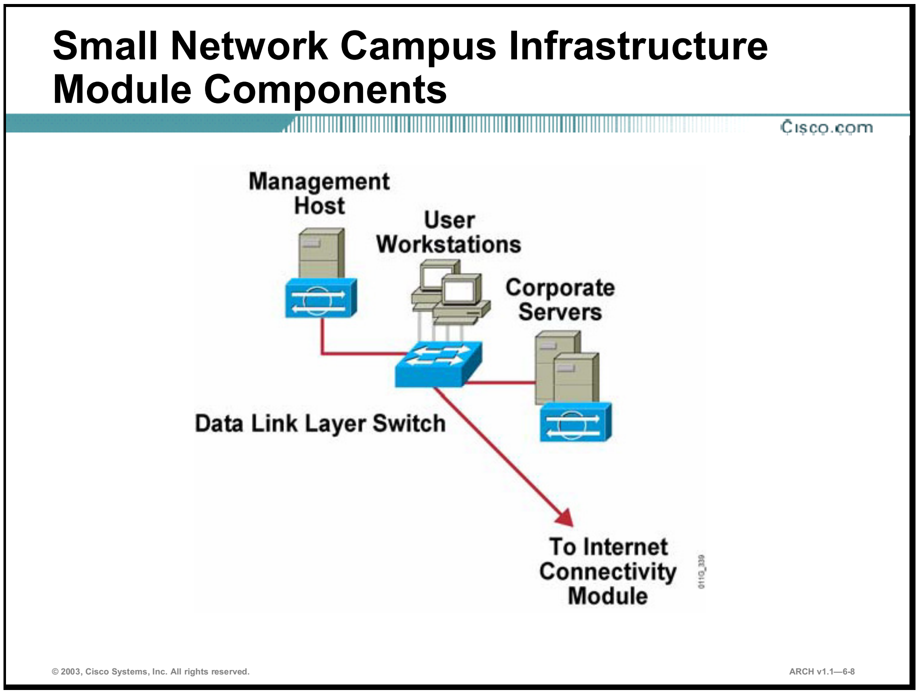

Small Network Campus Infrastructure Module Components

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—6-8

The Campus Infrastructure module contains end-user workstations, corporate intranet servers, management servers, and the associated data link layer infrastructure required to support the devices. Within the small network design, the data link layer functionality is combined into a single switch.

The figure shows these key devices for the Campus Infrastructure module:

- **Data link layer switching (with private VLAN support):** Provides data link services to user workstations
- **Corporate servers:** Provide e-mail (SMTP and POP3) services to internal users, as well as delivering file, print, and DNS services to workstations
- **User workstations:** Provide data services to authorized users on the network
- **Management host:** Provides HIPS, syslog, TACACS+/RADIUS, and general configuration management

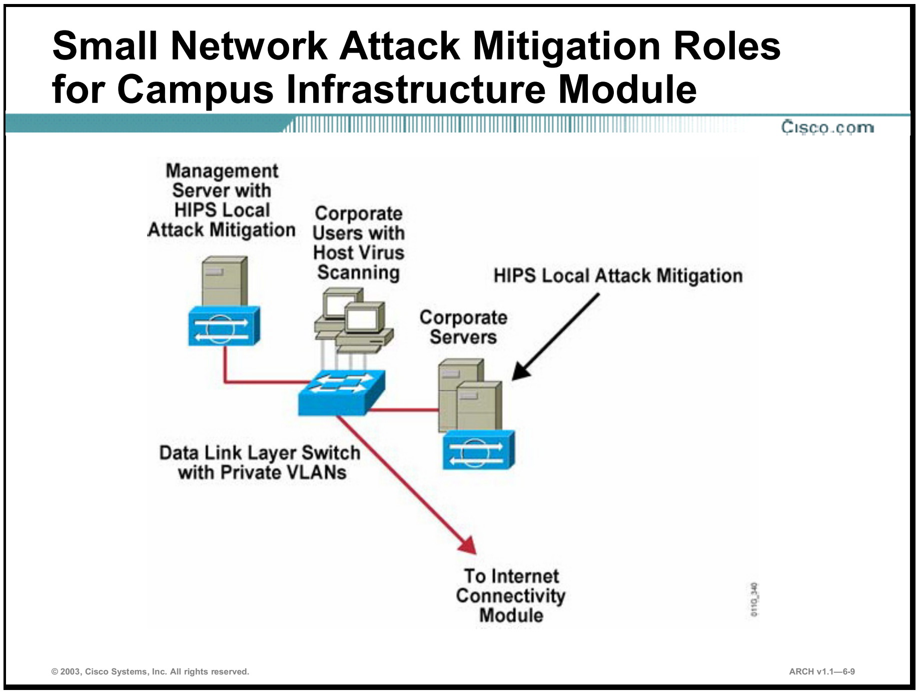

Small Network Attack Mitigation Roles for Campus Infrastructure Module

Management Server with HIPS Local Attack Mitigation
Corporate Users with Host Virus Scanning
HIPS Local Attack Mitigation
Corporate Servers
Data Link Layer Switch with Private VLANs
To Internet Connectivity Module

ARCH v1.1—6-9

The expected threats for the Campus Infrastructure module and the associated mitigating factors are:

- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.
- **Virus and Trojan horse applications:** Host-based virus scanning prevents most viruses and many Trojan horses.
- **Unauthorized access:** This type of access is mitigated through the use of HIPS and application access control.
- **Application layer attacks:** Operating systems, devices, and applications are kept up to date with the latest security fixes and are protected by HIPS.
- **Trust exploitation:** Private VLANs prevent hosts on the same subnet from communicating unless necessary.
- **Port redirection:** HIPS prevents port redirection agents from being installed.

SAFE Security Strategies for Medium Networks

The SAFE medium network design consists of the Internet Connectivity module, the Campus Infrastructure module, and the WAN module. This topic describes the SAFE architecture and security strategies for a medium-sized network.

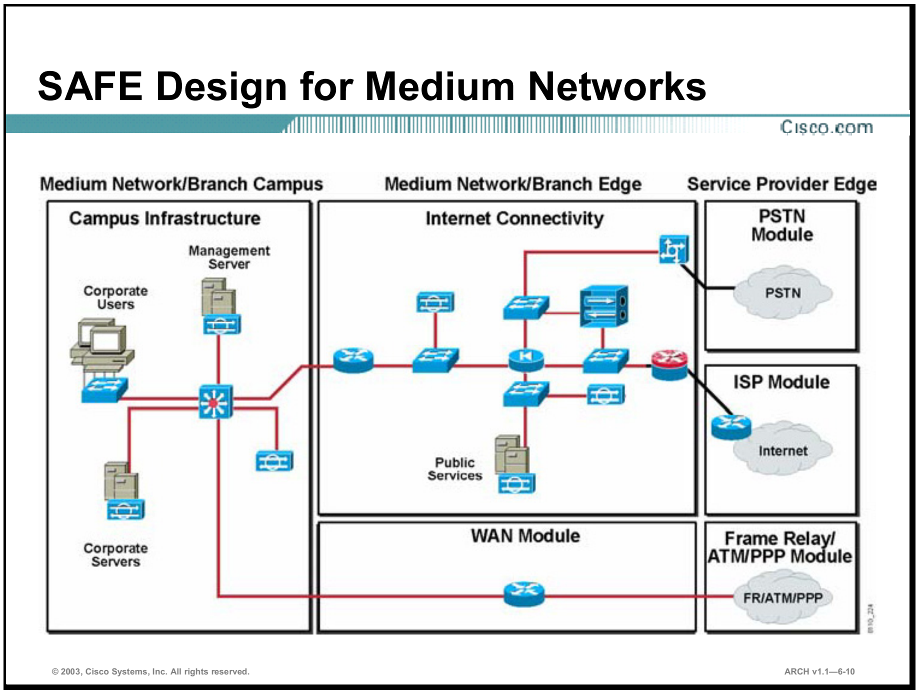

SAFE Design for Medium Networks

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-10

The SAFE medium network design consists of three modules: the Internet Connectivity module, the Campus Infrastructure module, and the WAN module. As in the small network design, the Internet Connectivity module has the connection to the Internet and terminates VPNs and public services (DNS, HTTP, FTP, and SMTP) traffic. Dial-in traffic also terminates at the Internet Connectivity module.

The Campus Infrastructure module contains the data link layer and multilayer switching infrastructure along with all the corporate users, management servers, and intranet servers.

From a WAN perspective, there are two options for remote sites connecting into the medium design. The first is a private WAN connection using the WAN module. The second is an IPSec VPN into the Internet Connectivity module.

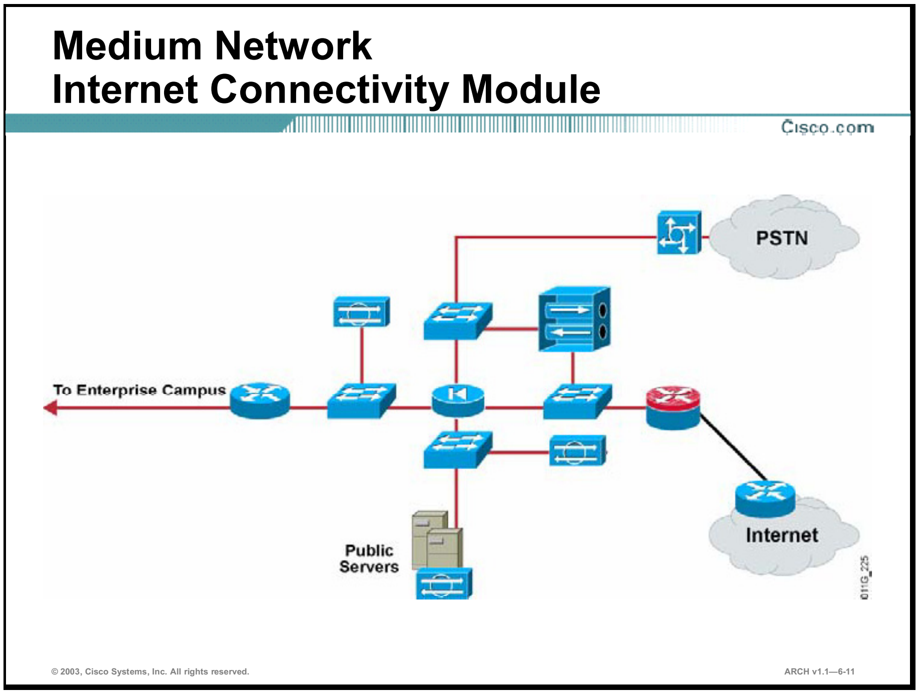

Medium Network Internet Connectivity Module

Cisco.com

ARCH v1.1—6-11

The goal of the Internet Connectivity module is to provide internal users with connectivity to Internet services and Internet users access to information on the public servers (HTTP, FTP, SMTP, and DNS). Additionally, this module terminates VPN traffic from remote users and remote sites as well as traffic from traditional dial-in users. The Internet Connectivity module is not designed to serve e-commerce type applications.

The Internet Connectivity module contains these devices:

- **Dial-in server:** Authenticates individual remote users and terminates analog connections
- **DNS server:** Serves as authoritative external DNS server for the medium network; relays internal DNS requests to the Internet
- **FTP/HTTP server:** Provides public information about the organization
- **Firewall:** Provides network-level protection of resources and stateful filtering of traffic; provides differentiated security for remote-access users; authenticates trusted remote sites; and provides connectivity using IPSec tunnels
- **Data link layer switches (with private VLAN support):** Provide data link layer connectivity for devices
- **NIDS appliance:** Provides Layer 4 to Layer 7 monitoring of key network segments in the module
- **SMTP server:** Acts as a relay between the Internet and the intranet mail servers; inspects content
- **VPN concentrator:** Authenticates individual remote users and terminates their IPSec tunnels
- **Edge router:** Provides basic filtering and Layer 3 connectivity to the Internet

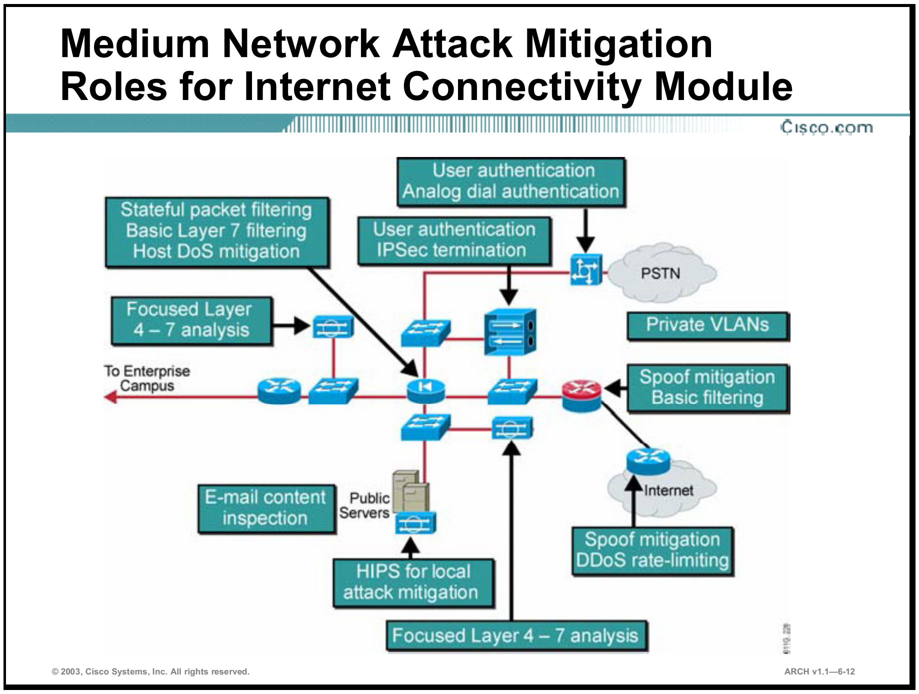

Medium Network Attack Mitigation Roles for Internet Connectivity Module

Cisco.com

User authentication
Analog dial authentication

Stateful packet filtering
Basic Layer 7 filtering
Host DoS mitigation

User authentication
IPSec termination

PSTN

Focused Layer 4 – 7 analysis

Private VLANs

To Enterprise Campus

Spoof mitigation
Basic filtering

E-mail content inspection

Public Servers

Internet

Spoof mitigation
DDoS rate-limiting

HIPS for local attack mitigation

Focused Layer 4 – 7 analysis

ARCH v1.1—6-12

The publicly addressable servers are likely points of attack within the Internet Connectivity module. The expected threats and the security features used to address them are:

- **Unauthorized access:** Mitigated through filtering at the ISP, edge router, and corporate firewall
- **Application layer attacks:** Mitigated through IDS at the host and network levels
- **Virus and Trojan horse attacks:** Mitigated through e-mail content filtering, HIPS, and host-based virus scanning
- **Password attacks:** Limited services available to brute-force attack; OS and IDS can detect the threat
- **Denial of service:** CAR at ISP edge and TCP setup controls at firewall
- **IP spoofing:** RFC 2827 and 1918 filtering at ISP edge and medium network edge router
- **Packet sniffers:** Switched infrastructure and HIPS to limit exposure
- **Network reconnaissance:** IDS detects reconnaissance, protocols filtered to limit effectiveness
- **Trust exploitation:** Restrictive trust model and private VLANs to limit trust-based attacks
- **Port redirection:** Restrictive filtering and HIPS to limit attacks

The remote-access and site-to-site VPN services are also points of attack. Expected threats in these areas are:

- **Network topology discovery:** ACLs on the ingress router limit access to the VPN concentrator and firewall (when used to terminate IPSec tunnels from remote sites) to IKE and ESP from the Internet.
- **Password attack:** OTPs mitigate brute-force password attacks.
- **Unauthorized access:** Firewall services after packet decryption prevent traffic on unauthorized ports.
- **Man-in-the-middle attacks:** These attacks are mitigated through encrypted remote traffic.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.

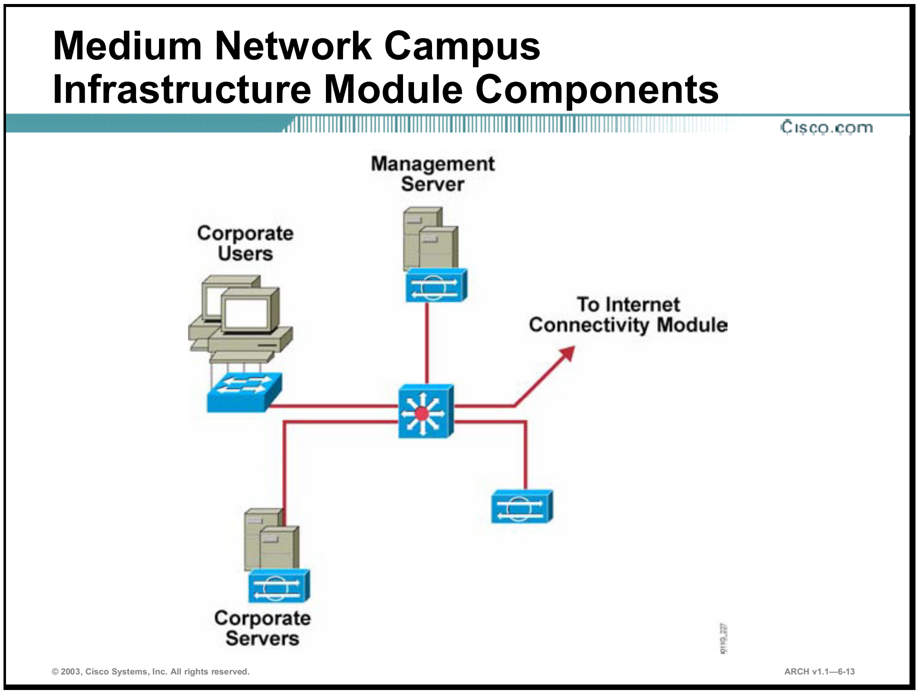

Medium Network Campus Infrastructure Module Components

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-13

The Campus Infrastructure module contains end-user workstations, corporate intranet servers, management servers, and the associated data link layer and multilayer-switched infrastructure required to support the devices. This configuration more accurately reflects the smaller size of medium networks, and reduces the overall cost of the design. As in the Internet Connectivity module, the redundancy normally found in an enterprise design is not reflected in the medium network design.

The Campus Infrastructure module includes these devices:

- **Multilayer switch:** Routes and switches production and management traffic within the campus module, provides distribution layer services to the building switches, and supports advanced services such as traffic filtering
- **Data link layer switches (with private VLAN support):** Provides data link layer services to user workstations
- **Corporate servers:** Provides e-mail (SMTP and POP3) services to internal users, as well as delivering file, print, and DNS services to workstations
- **User workstations:** Provides data services to authorized users on the network
- **SNMP management host:** Provides SNMP management for devices
- **NIDS host:** Provides alarm aggregation for all NIDS devices in the network
- **Syslog host**: Aggregates log information for firewall and NIDS hosts
- **Access control server:** Delivers authentication services to the network devices
- **OTP Server:** Authorizes one-time password information relayed from the access control server
- **System admin host:** Provides configuration, software, and content changes on devices
- **NIDS appliance:** Provides Layer 4 to Layer 7 monitoring of key network segments in the module

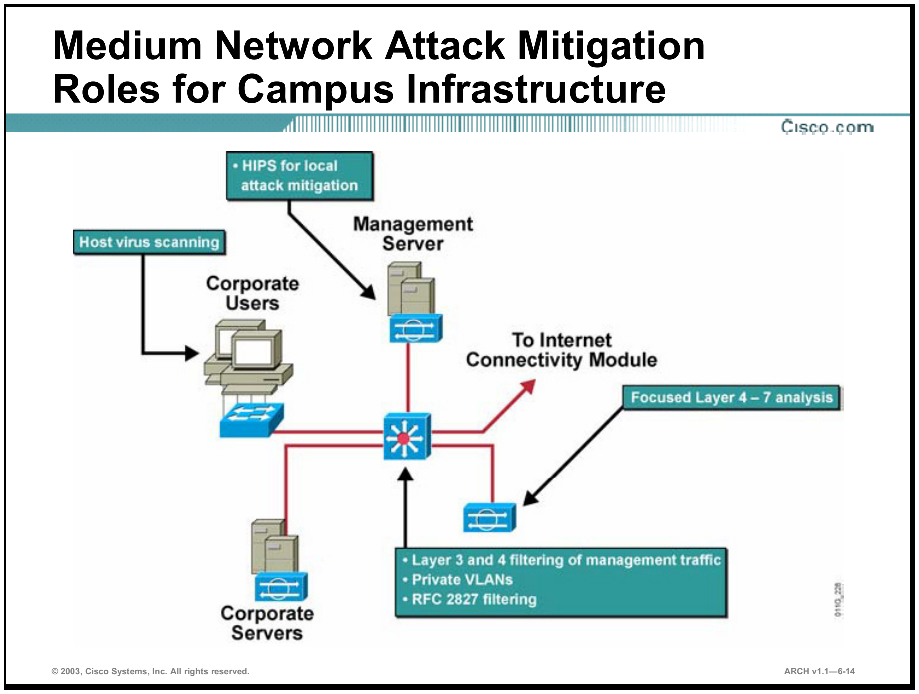

Medium Network Attack Mitigation Roles for Campus Infrastructure

Cisco.com

- HIPS for local attack mitigation

Host virus scanning

Management Server

Corporate Users

To Internet Connectivity Module

Focused Layer 4 – 7 analysis

- Layer 3 and 4 filtering of management traffic
- Private VLANs
- RFC 2827 filtering

Corporate Servers

ARCH v1.1—6-14

The expected threats and their solutions for the Campus Infrastructure module are:

- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.
- **Virus and Trojan horse applications:** Host-based virus scanning prevents most viruses and many Trojan horses.
- **Unauthorized access:** These types of attacks are mitigated through the use of host-based intrusion detection and application access control.
- **Password attacks:** The access control server allows for strong authentication for key applications.
- **Application layer attacks:** Operating systems, devices, and applications are kept up to date with the latest security fixes, and HIPS protects them.
- **IP spoofing:** RFC 2827 filtering prevents source-address spoofing.
- **Trust exploitation:** Trust arrangements are very explicit; private VLANs prevent hosts on the same subnet from communicating unless necessary.
- **Port redirection:** HIPS prevents port redirection agents from being installed.

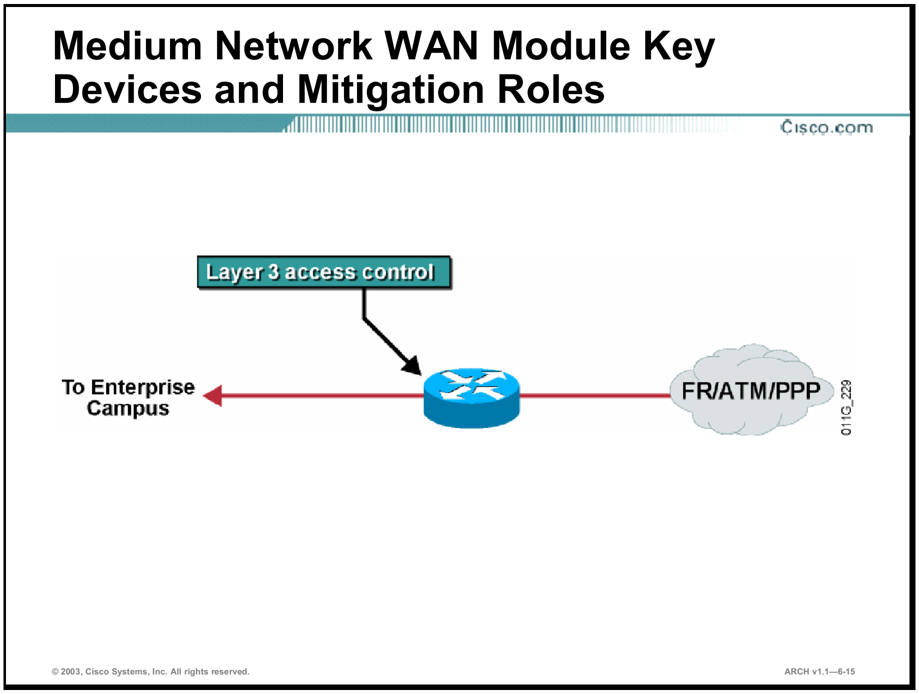

Medium Network WAN Module Key Devices and Mitigation Roles

Cisco.com

Layer 3 access control

To Enterprise Campus

FR/ATM/PPP

ARCH v1.1—6-15

The WAN module is included in the medium network only when connections to remote locations over a private network are required. This requirement may occur when an IPSec VPN cannot meet stringent quality of service (QoS) requirements, or when legacy WAN connections are in place without a compelling cost justification to migrate to IPSec.

The key device for this module is a Cisco router, which provides routing, access control, and QoS mechanisms to remote locations.

The threats mitigated by the IOS router include:

- **IP spoofing:** IP spoofing can be mitigated through Layer 3 filtering
- **Unauthorized access:** Simple access control on the router can limit the types of protocols, applications, networks, and devices to which branches have access

SAFE Security Strategies for Large Networks

The SAFE large network design consists of the entire Enterprise Campus functional area. This topic describes the SAFE architecture and security strategies for a large network.

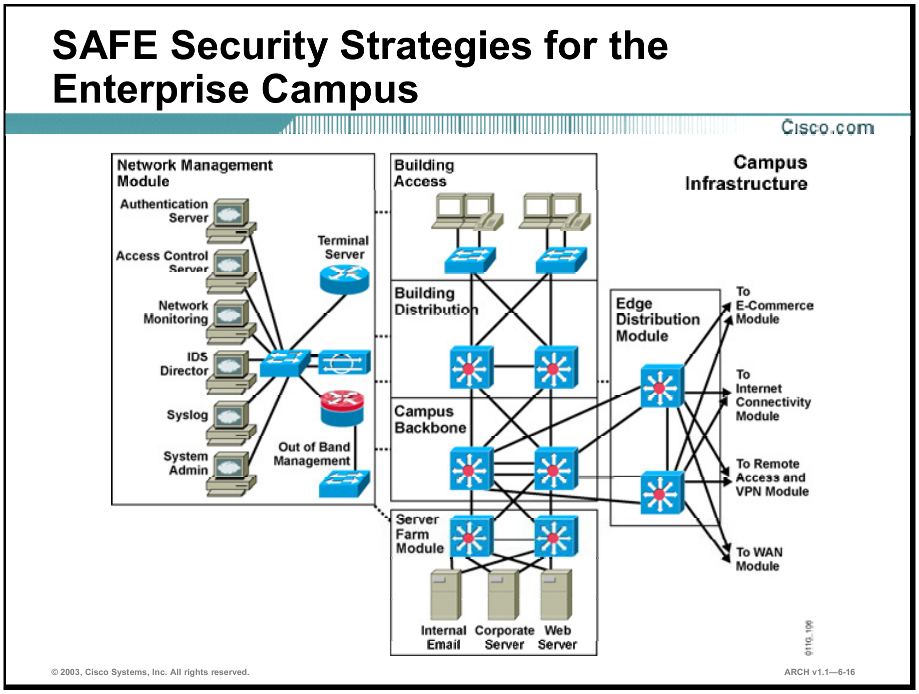

SAFE Security Strategies for the Enterprise Campus

ARCH v1.1—6-16

The figure illustrates the modules of the Enterprise Campus functional area. Security considerations for each module differ based on the function of the module.

The Campus Infrastructure module is composed of three submodules: Campus Backbone, Building Distribution, and Building Access.

As shown in the figure, multilayer switches are used to route and switch production network data from one module to another. By using a switched Campus Backbone, the effectiveness of packet sniffers is limited.

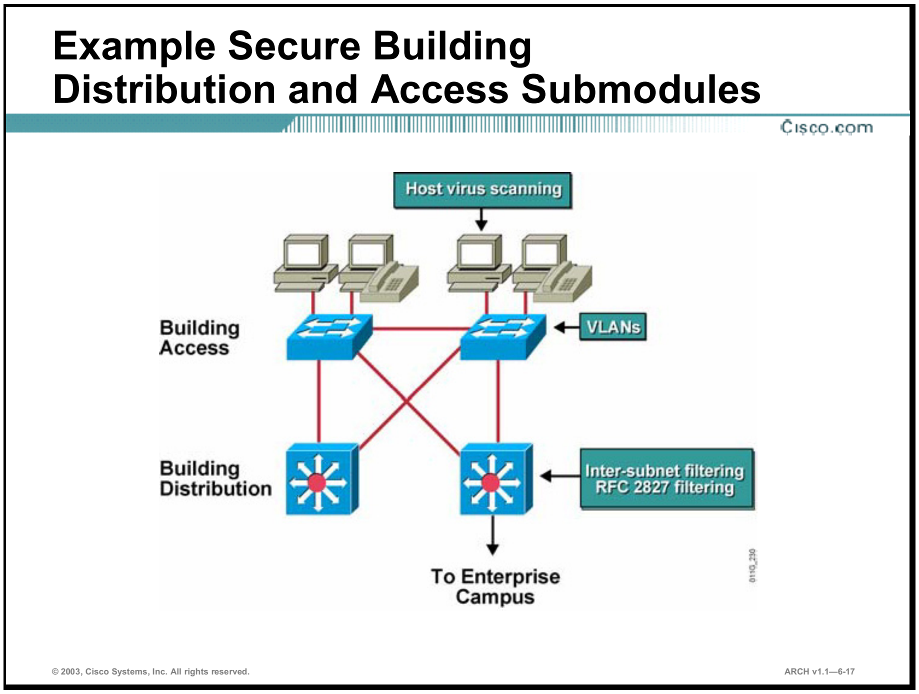

Example Secure Building Distribution and Access Submodules

ARCH v1.1—6-17

The goal of the Building Distribution submodule is to provide distribution layer services to the building switches. Services include routing, QoS, and access control. Requests for data flow into these switches and onto the core, and responses follow the identical path in reverse.

The security features implemented help mitigate these attacks:

- **Unauthorized access:** Attacks against server module resources are limited by Layer 3 filtering of specific subnets.
- **IP spoofing:** RFC 2827 filtering stops most spoofing attempts.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.

SAFE defines the Building Access submodule as the extensive network portion that contains end-user workstations, phones, and their associated Layer 2 access points. Its primary goal is to provide services to end users.

The security strategy implemented at the Building Access submodule addresses these types of attacks:

- **Packet sniffers:** A switched infrastructure and default VLAN services limit the effectiveness of sniffing.
- **Virus and Trojan horse applications:** Host-based virus scanning prevents most viruses and many Trojan horses.

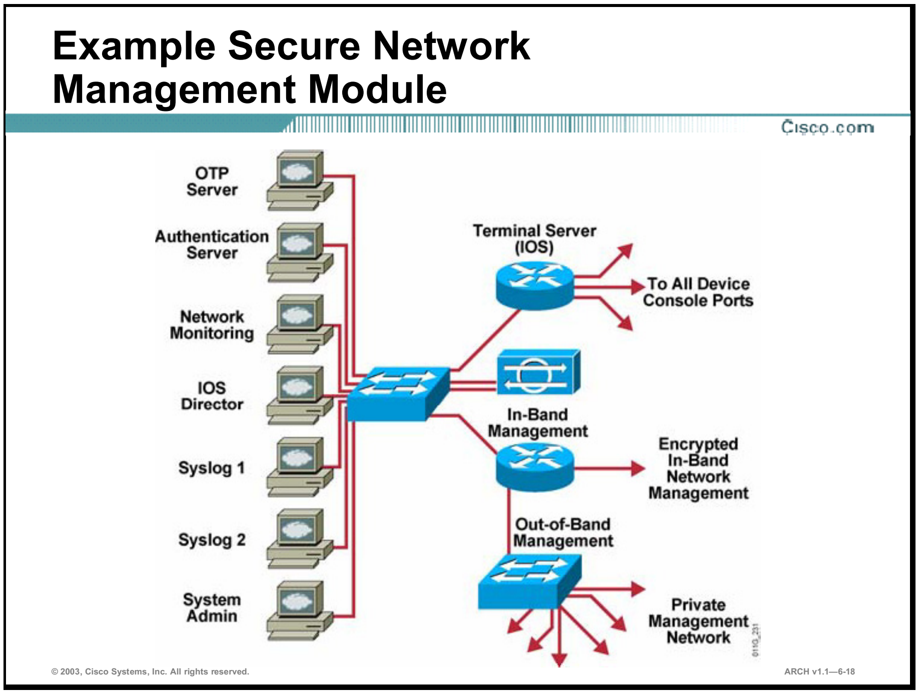

Example Secure Network Management Module

Cisco.com

ARCH v1.1—6-18

The primary goal of the Network Management module is to facilitate the secure management of all devices and hosts within the enterprise SAFE architecture. Logging and reporting information flows from the devices through to the management hosts, while content, configurations, and new software flow to the devices from the management hosts.

These primary devices are used in the Network Management module:

- **SNMP management host:** Provides SNMP management for devices
- **NIDS host:** Provides alarm aggregation for all NIDS devices in the network
- **Syslog hosts:** Aggregates log information for firewall and NIDS hosts
- **Access control server:** Delivers one-time, two-factor authentication services to the network devices
- **OTP server:** Authorizes One Time Password information relayed from the access control server
- **System administration host:** Provides configuration, software, and content changes
- **NIDS appliance:** Provides Layer 4 to Layer 7 monitoring of key network segments
- **IOS firewall:** Allows granular control for traffic flows between the management hosts and the managed devices
- **Data link layer switch (with private VLAN support):** Ensures data from managed devices can only cross directly to the IOS firewall

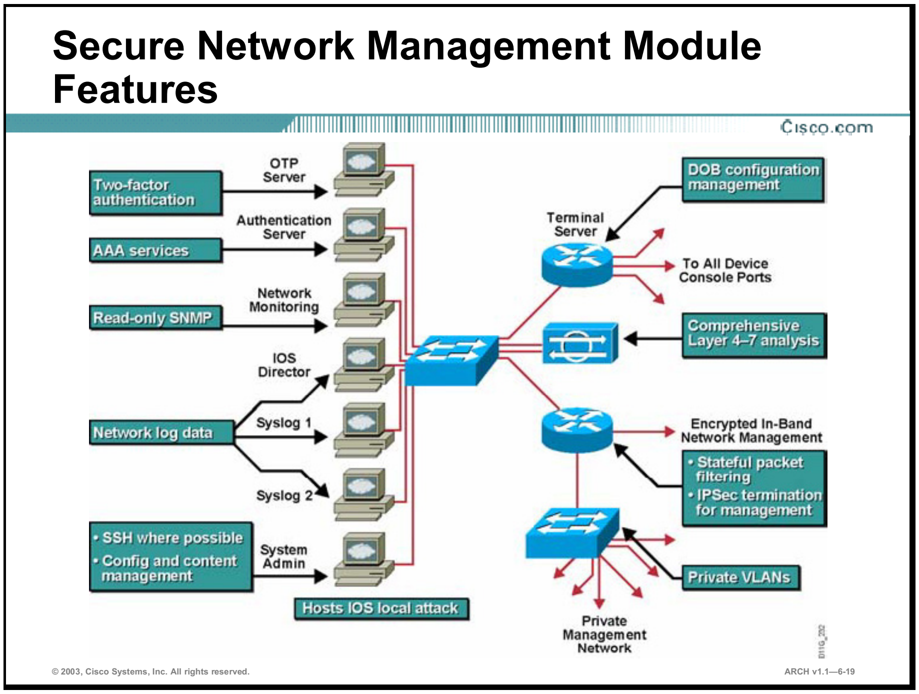

Secure Network Management Module Features

Cisco.com
Two-factor authentication
OTP Server
AAA services
Authentication Server
Read-only SNMP
Network Monitoring
IOS Director
Network log data
Syslog 1
Syslog 2
• SSH where possible
• Config and content management
System Admin
Hosts IOS local attack
DOB configuration management
Terminal Server
To All Device Console Ports
Comprehensive Layer 4–7 analysis
Encrypted In-Band Network Management
• Stateful packet filtering
• IPSec termination for management
Private VLANs
Private Management Network
B11G_252

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-19

The security features implemented in the Network Management module are shown in the figure, and help mitigate these attacks:

- **Unauthorized access:** Filtering at the IOS firewall stops most unauthorized traffic in both directions.
- **Man-in-the-middle attacks:** Management data is crossing a private network making man-in-the-middle attacks difficult.
- **Network reconnaissance:** Because all management traffic crosses this network, it does not cross the production network where it could be intercepted.
- **Password attacks:** The access control server allows for strong two-factor authentication at each device.
- **IP spoofing:** Spoofed traffic is stopped in both directions at the IOS firewall through Layer 3 filtering.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.
- **Trust exploitation:** Private VLANs prevent a compromised device from masquerading as a management host.

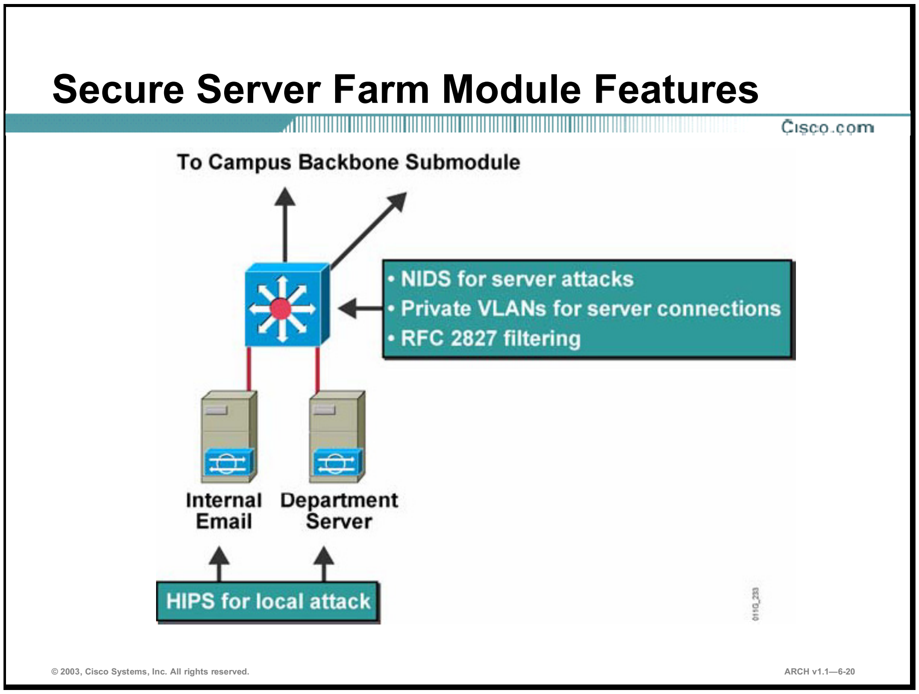

Secure Server Farm Module Features

Cisco.com

To Campus Backbone Submodule

- NIDS for server attacks
- Private VLANs for server connections
- RFC 2827 filtering

HIPS for local attack

ARCH v1.1—6-20

The Server Farm module's primary goal is to provide application services to end users and devices. On-board intrusion detection within the multilayer switches inspects traffic flows in the Server Farm module.

The security strategy shown in the figure addresses these attack threats:

- **Unauthorized access:** Mitigated through the use of HIPS and application access control.
- **Application layer attacks:** Operating systems, devices, and applications are kept up to date with the latest security fixes and protected by HIPS.
- **IP spoofing:** RFC 2827 filtering prevents source address spoofing.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.
- **Trust exploitation:** Trust arrangements are very explicit; private VLANs prevent hosts on the same subnet from communicating unless necessary.
- **Port redirection:** HIPS prevents port redirection agents from being installed.

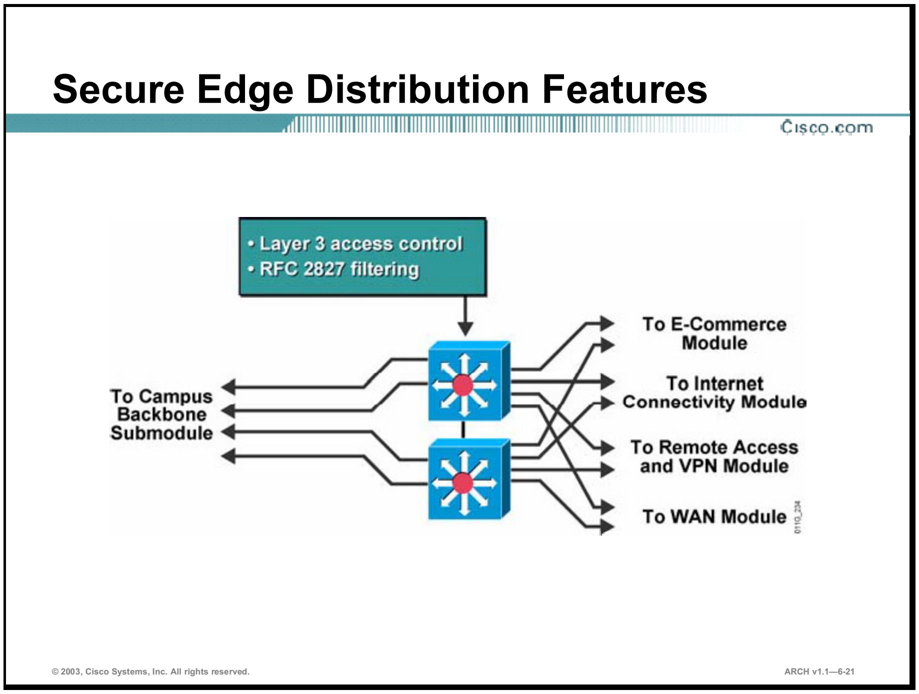

Secure Edge Distribution Features

ARCH v1.1—6-21

The Edge Distribution module aggregates the connectivity from the various elements at the edge. Traffic is filtered and routed from the edge modules and routed into the Campus Backbone submodule.

The key devices in this module are multilayer switches used to aggregate edge connectivity and provide advanced services.

The security strategy illustrated in the figure addresses these threats:

- **Unauthorized access:** Filtering provides granular control over specific edge subnets and their ability to reach areas within the campus.
- **IP spoofing:** RFC 2827 filtering limits locally initiated spoof attacks.
- **Network reconnaissance:** Filtering limits nonessential traffic from entering the campus, limiting a hacker’s ability to perform network reconnaissance.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.

SAFE Security Strategies for the Enterprise Edge

The SAFE architecture defines the Enterprise Edge functional area as containing the Internet Connectivity module, E-Commerce module, Remote Access and VPN module, and the WAN module. This topic describes a security strategy for the Enterprise Edge.

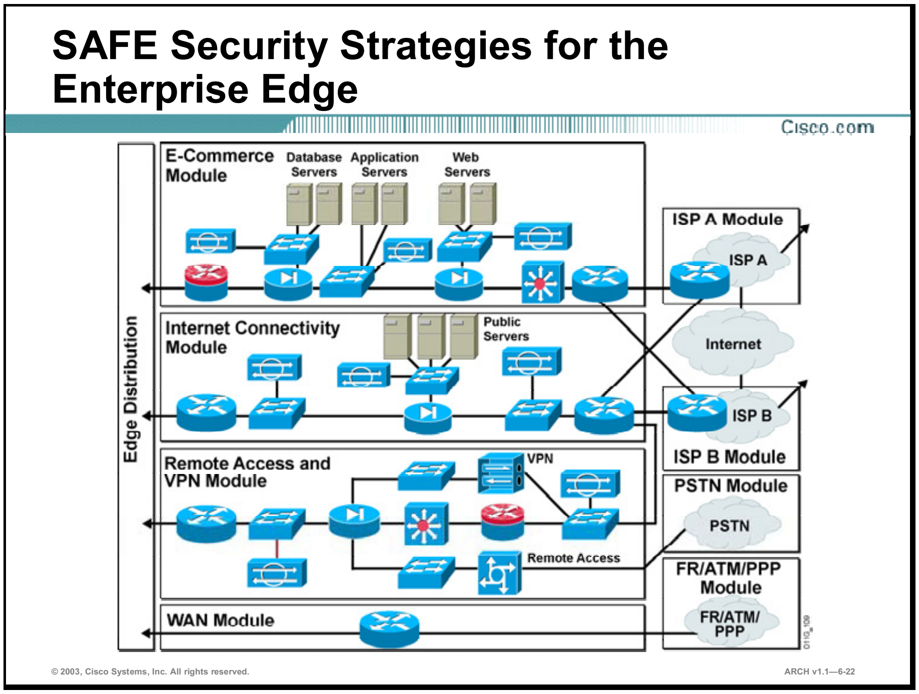

SAFE Security Strategies for the Enterprise Edge

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—6-22

As illustrated in the figure, these key devices are deployed in the E-Commerce module:

- **Web server:** Acts as the primary user interface for the navigation of the e-commerce store
- **Application server:** Platform for the various applications required by the web server
- **Database server:** Critical information that is the heart of the e-commerce business implementation
- **Firewall:** Governs communication between the various levels of security and trust in the system
- **NIDS appliance:** Provides monitoring of key network segments in the module
- **Multilayer switch with IDS module:** The scalable e-commerce input device with integrated security monitoring

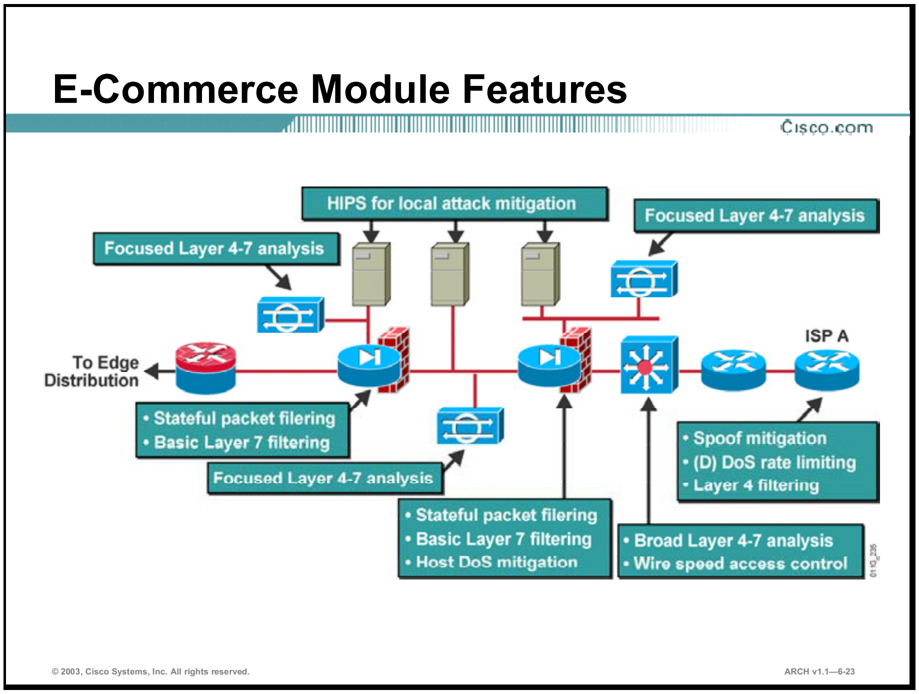

E-Commerce Module Features

Cisco.com
HIPS for local attack mitigation
Focused Layer 4-7 analysis
Focused Layer 4-7 analysis
ISP A
To Edge Distribution
• Stateful packet filering
• Basic Layer 7 filtering
Focused Layer 4-7 analysis
• Spoof mitigation
• (D) DoS rate limiting
• Layer 4 filtering
• Stateful packet filering
• Basic Layer 7 filtering
• Host DoS mitigation
• Broad Layer 4-7 analysis
• Wire speed access control

ARCH v1.1—6-23

The security features illustrated in the figure for the E-Commerce module mitigate these threats:

- **Unauthorized access:** Stateful firewalls and ACLs limit exposure to specific protocols.
- **Application layer attacks:** Attacks are mitigated through the use of IDS.
- **Denial of service:** ISP filtering and rate-limiting reduce denial-of-service potential.
- **IP spoofing:** RFC 2827 and 1918 prevent locally originated spoofed packets and limit remote spoof attempts.
- **Packet sniffers:** A switched infrastructure and HIPS limit the effectiveness of sniffing.
- **Network reconnaissance:** Ports are limited to only what is necessary. ICMP is restricted.
- **Trust exploitation:** Firewalls ensure communication flows only in the proper direction on the proper service.
- **Port redirection:** HIPS and firewall filtering limit exposure to these attacks.

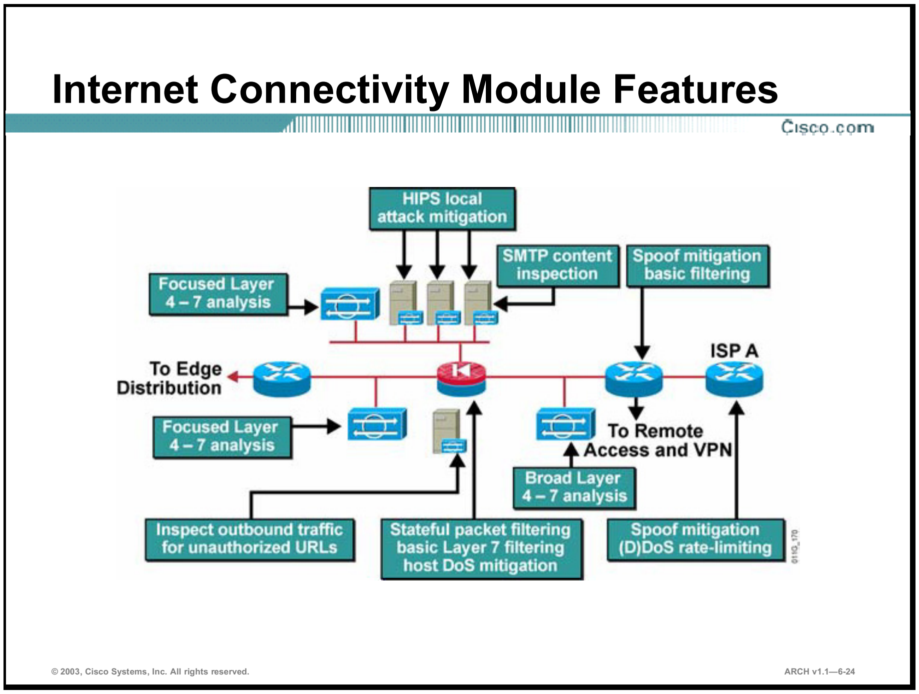

Internet Connectivity Module Features

Cisco.com
HIPS local attack mitigation
SMTP content inspection
Spoof mitigation basic filtering
Focused Layer 4 – 7 analysis
ISP A
To Edge Distribution
Focused Layer 4 – 7 analysis
To Remote Access and VPN
Broad Layer 4 – 7 analysis
Inspect outbound traffic for unauthorized URLs
Stateful packet filtering basic Layer 7 filtering host DoS mitigation
Spoof mitigation (D)DoS rate-limiting
0110_170

The security features illustrated in the figure for the Internet Connectivity module mitigate these threats:

- **Unauthorized access:** Mitigated through filtering at the ISP, edge router, and corporate firewall.
- **Application layer attacks:** Mitigated through IDS at the host and network levels.
- **Virus and Trojan horse:** Mitigated through e-mail content filtering and HIPS.
- **Password attacks:** Limited services available to brute force; operating systems and IDSs can detect the threat.
- **Denial of service:** Rate limiting at ISP edge and TCP setup controls at firewall.
- **IP spoofing:** RFC 2827 and 1918 filtering at ISP edge and enterprise edge router.
- **Packet sniffers:** Switched infrastructure and HIPS limits exposure.
- **Network reconnaissance:** IDS detects reconnaissance; protocols filtered to limit effectiveness.
- **Trust exploitation:** Restrictive trust model and private VLANs limit trust-based attacks.
- **Port redirection:** Restrictive filtering and HIPS limit attacks.

Remote Access and VPN Module Features

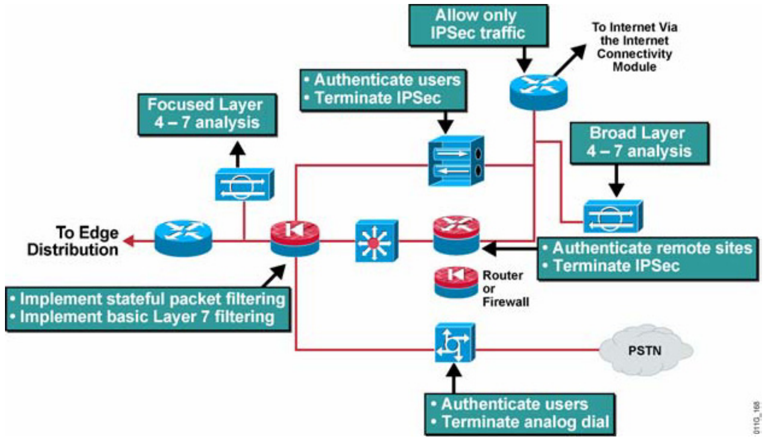

Allow only IPSec traffic

To Internet Via the Internet Connectivity Module

- Authenticate users
- Terminate IPSec

Focused Layer 4 – 7 analysis

Broad Layer 4 – 7 analysis

To Edge Distribution

- Authenticate remote sites
- Terminate IPSec

Router or Firewall

- Implement stateful packet filtering
- Implement basic Layer 7 filtering

PSTN

- Authenticate users
- Terminate analog dial

ARCH v1.1—6-25

The security features illustrated in the figure for the Remote Access and VPN module mitigate these threats:

- **Network topology discovery:** Only IKE and ESP are allowed into this segment from the Internet.
- **Password attack:** OTP authentication reduces the likelihood of a successful password attack.
- **Unauthorized access:** Firewall services after packet decryption prevent traffic on unauthorized ports.
- **Man-in-the-middle:** Mitigated through encrypted remote traffic.
- **Packet sniffers:** A switched infrastructure limits the effectiveness of sniffing.

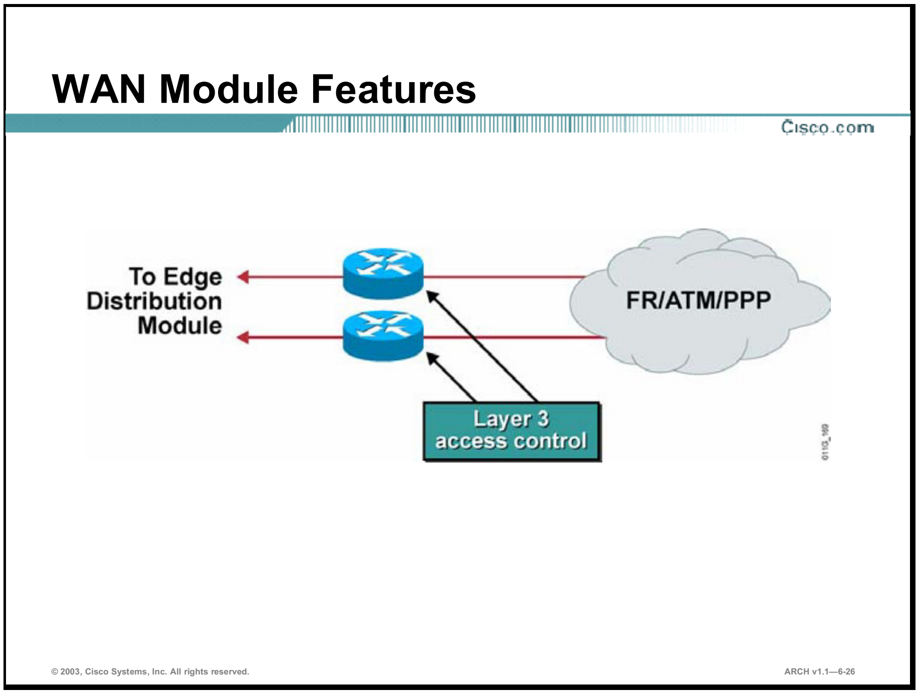

WAN Module Features

Cisco.com

ARCH v1.1—6-26

The WAN module is often not addressed in a security context. You can mitigate man-in-the-middle attacks initiated through an ISP with these features:

- **Access control:** Access control is required for Layer 3 and Layer 4 network functionality.
- **IPSec encryption:** IPSec encryption is needed if the level of trust for Layer 2 WAN technology is not high.

Summary

This topic summarizes the key points discussed in this lesson.

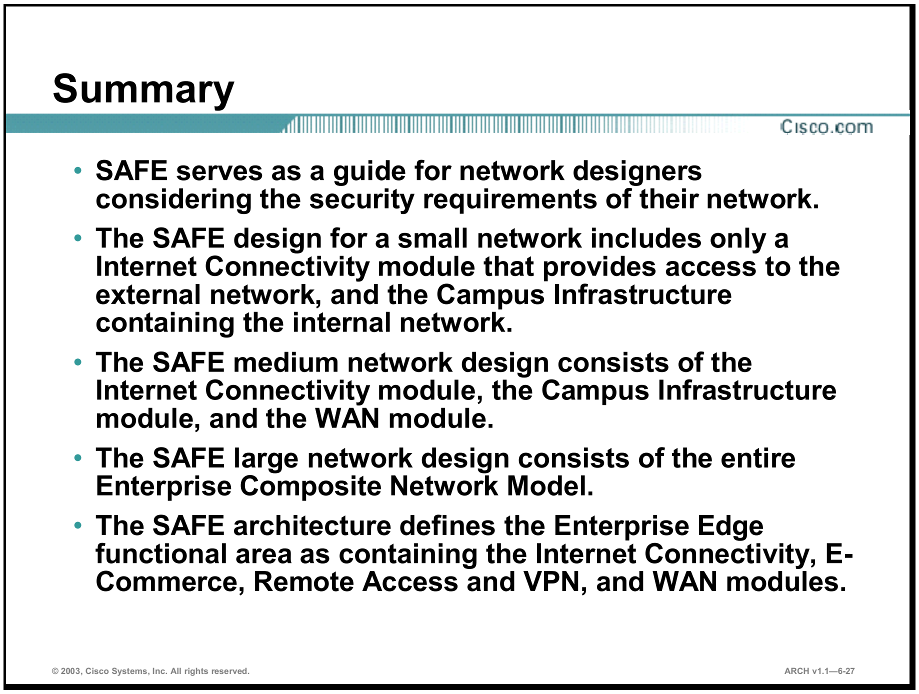

Summary

Cisco.com

- SAFE serves as a guide for network designers considering the security requirements of their network.
- The SAFE design for a small network includes only a Internet Connectivity module that provides access to the external network, and the Campus Infrastructure containing the internal network.
- The SAFE medium network design consists of the Internet Connectivity module, the Campus Infrastructure module, and the WAN module.
- The SAFE large network design consists of the entire Enterprise Composite Network Model.
- The SAFE architecture defines the Enterprise Edge functional area as containing the Internet Connectivity, E-Commerce, Remote Access and VPN, and WAN modules.

ARCH v1.1—6-27

References

For additional information, refer to these resources:

- *Cisco SAFE Blueprint* at http://www.cisco.com/warp/public/779/largeent/issues/security/safe.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study, refer to the following section:

- Case Study 6-3: OCSIC Bottling Company

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) What is one reason for using a modular approach to security design?

A) It reduces the overall cost of the solution.

B) It allows the network to continue to provide critical services.

C) It prevents most attacks from successfully affecting valuable network resources.

D) It addresses the security relationship between the functional blocks of the network.

Q2) Which security feature mitigates application layer attacks in the small network Internet Connectivity module?

A) private VLANs

B) a switched infrastructure

C) network intrusion detection system (NIDS) appliance

D) host intrusion protection system (HIPS) on the public servers

Q3) Which security feature mitigates man-in-the-middle attacks in the medium-size Internet Connectivity module?

A) HIPS

B) private VLANs

C) RFC 2827 filtering

D) remote traffic encryption

Q4) Which security feature mitigates IP spoofing attacks in the Building Distribution submodule?

A) private VLANs

B) RFC 2827 filtering

C) remote traffic encryption

D) host intrusion protection system

Q5) Which two threats to the Server Farm module can a HIPS mitigate? (Choose two.)

A) IP spoofing

B) packet sniffers

C) port redirection

D) trust exploitation

E) unauthorized access

Q6) What security feature can mitigate the threat of packet sniffers in the Edge Distribution module?

A) HIPS

B) NIDS

C) filtering

D) switched infrastructure

Q7) Which two security features can mitigate denial-of-service attacks at the Internet Connectivity module? (Choose two.)

A) ACLs

B) ISP filtering

C) rate limiting

D) intrusion protection

Quiz Answer Key

Q1) D

Relates to: Cisco SAFE Architecture Overview

Q2) D

Relates to: SAFE Security Strategies for Small Networks

Q3) D

Relates to: SAFE Security Strategies for Medium Networks

Q4) B

Relates to: SAFE Security Strategies for Large Networks

Q5) C, E

Relates to: SAFE Security Strategies for Large Networks

Q6) D

Relates to: SAFE Security Strategies for Large Networks

Q7) B, C

Relates to: SAFE Security Strategies for the Enterprise Edge

Case Study 6-3: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

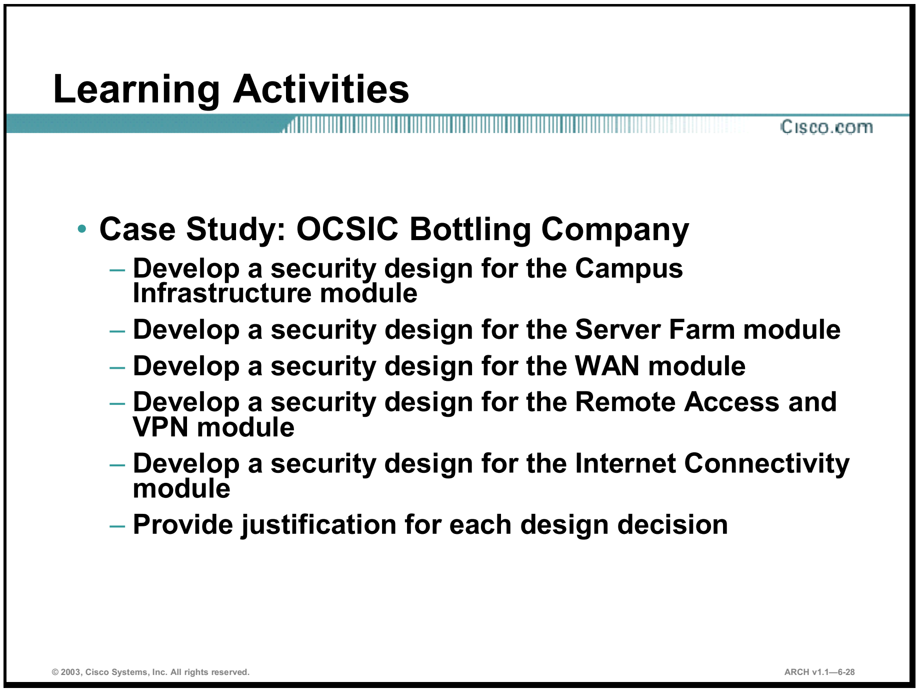

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Develop a security design for the Campus Infrastructure module
  - Develop a security design for the Server Farm module
  - Develop a security design for the WAN module
  - Develop a security design for the Remote Access and VPN module
  - Develop a security design for the Internet Connectivity module
  - Provide justification for each design decision

ARCH v1.1—6-28

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

OCSIC bottling company wants a basic security solution that will provide them with authentication for the remote-access users, and users trying to gain access to the network devices. They also want to implement a firewall with intrusion detection on their connection going out to the Internet. Management would like to do some type of filtering between the finance network and the rest of the corporate network. Finally, they want to be able to filter based on IP addresses.

The security requirements for the enterprise network are:

- Provide remote access to the corporate database to all authorized users.
- Allow R&D users to protect their data by limiting salespeople from accessing their network and by encrypting sensitive data as it is stored on the servers.
- Provide authenticated access to the company accounting system.
- Allow all internal users access to the Internet.

In this exercise, you will design security services that meet the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Develop a security design for the headquarters campus network
- Develop a security design for the Server Farm module
- Develop a security design for the WAN module
- Develop a security design for the Remote Access and VPN module
- Develop a security design for the Internet Connectivity module

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Note You can complete all tasks within your group, or complete the task assigned to your group by the instructor.

Task 1: Develop a Security Policy for the Network

Complete these steps:

Step 1 Propose a comprehensive security policy for the company network.

Task 2: Develop a Security Design for the Headquarters Campus Network

Complete these steps:

Step 1 Complete the table to design your security solution for the headquarters campus network.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Task 3: Develop a Security Design for the Server Farm Module

Complete these steps:

Step 1 Complete the table to design your security solution for the Server Farm module.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Task 4: Develop a Security Design for the WAN Module

Complete these steps:

Step 1 Complete the table to design your security solution for the WAN module.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Task 5: Develop a Security Design for the Remote Access Module

Complete these steps:

Step 1 Complete the table to design your security solution for the Remote Access module.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Task 6: Develop a Security Design for the Internet Connectivity Module

Complete these steps:

Step 1 Complete the table to design your security solution for the Internet Connectivity module.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Step 2 Update your campus network diagram to reflect your security design.

Task 7: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have a design for the headquarters campus network with a completed network diagram.
- You have a design for the Server Farm module with a completed network diagram.
- You have a design for the WAN module with a completed network diagram.
- You have a design for the Remote Access and VPN module with a completed network diagram.
- You have a design for the Internet Connectivity module with a completed network diagram.

Module 7

Designing QoS

Overview

An efficient, well-designed network forms the backbone of a successful enterprise. These networks transport many applications and data, including high-quality video and critical application data. Bandwidth-intensive applications may stretch network capabilities and resources. Networks must provide secure, predictable, measurable, and sometimes guaranteed levels of service.

Achieving the required quality of service (QoS) by managing delay, delay variation (jitter), bandwidth, and packet loss on a network is critical for the end-to-end network. QoS tools give network designers the techniques to manage the network resources.

Module Objectives

Upon completing this module, you will be able to design QoS intelligent network services for performance, scalability, and availability, given specified enterprise network needs.

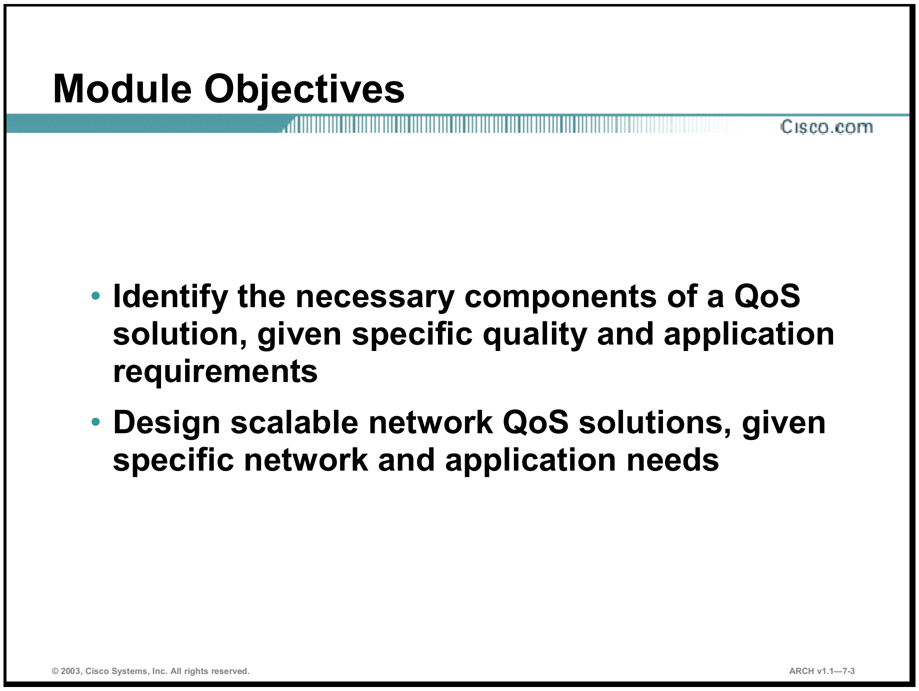

Module Objectives

Cisco.com

- Identify the necessary components of a QoS solution, given specific quality and application requirements
- Design scalable network QoS solutions, given specific network and application needs

ARCH v1.1—7-3

Module Outline

The outline lists the components of this module.

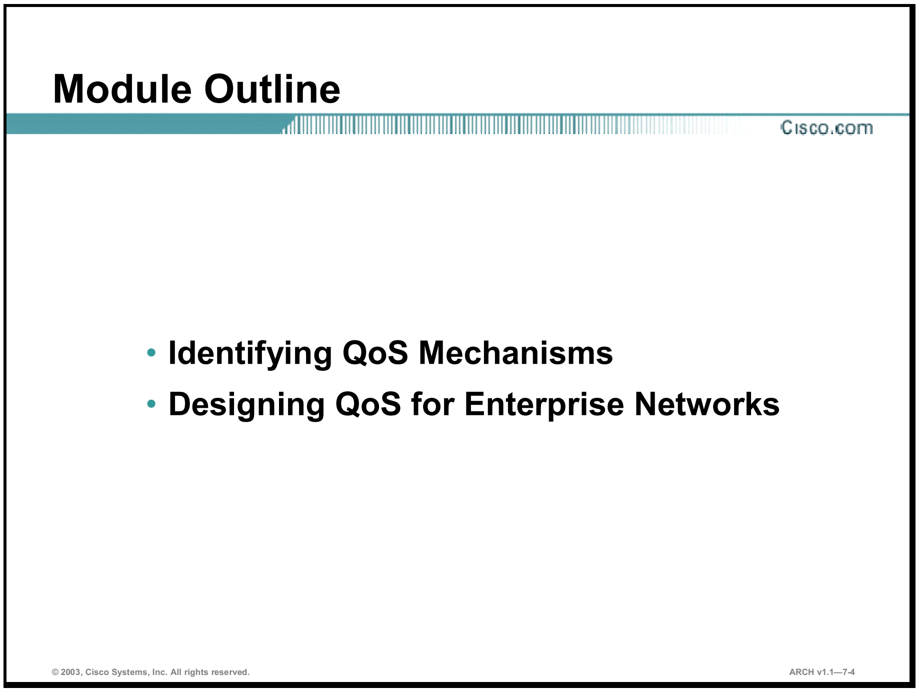

Module Outline

Cisco.com

- Identifying QoS Mechanisms
- Designing QoS for Enterprise Networks

ARCH v1.1—7-4

Identifying QoS Mechanisms

Overview

Cisco IOS software provides a range of QoS tools that address the needs of voice, video, and data applications. Cisco IOS QoS technology allows the network designer to implement complex networks that predictably control services to a variety of networked applications and traffic types.

Using the QoS toolset in Cisco IOS software, enterprises can design and implement networks that conform to either the Internet Engineering Task Force (IETF) Integrated Services (IntServ) model or the Differentiated Services (DiffServ) model. Cisco IOS QoS tools provide additional functionality such as network-based application recognition (NBAR) for classifying traffic on an application basis, a service assurance agent (SAA) for end-to-end QoS measurements, and Resource Reservation Protocol (RSVP) signaling for admission control and reservation of resources.

Relevance

Bandwidth, delay, jitter, and packet loss can be effectively controlled on an enterprise network. QoS features lead to efficient, predictable services for business-critical applications.

Objectives

Upon completing this lesson, you will be able to identify the necessary components of a QoS solution, given specific quality and application requirements. This includes being able to meet these objectives:

- Identify enterprise requirements for quality of service
- Describe the IntServ and DiffServ QoS architectures, and explain when to use each one
- Identify the classification and marking components of a Cisco QoS solution, given specific quality and application requirements
- Identify the congestion avoidance components of a Cisco QoS solution, given specific quality and application requirements
- Identify the congestion management components of a Cisco QoS solution, given specific quality and application requirements

- Identify the traffic conditioning components of a Cisco QoS solution, given specific quality and application requirements
- Identify the signaling components of a Cisco QoS solution, given specific quality and application requirements
- Identify the link efficiency components of a Cisco QoS solution, given specific quality and application requirements
- Summarize the key IOS software QoS features and explain when to use each one

Learner Skills and Knowledge

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

- CCNP® curriculum courses

Outline

The outline lists the topics included in this lesson.

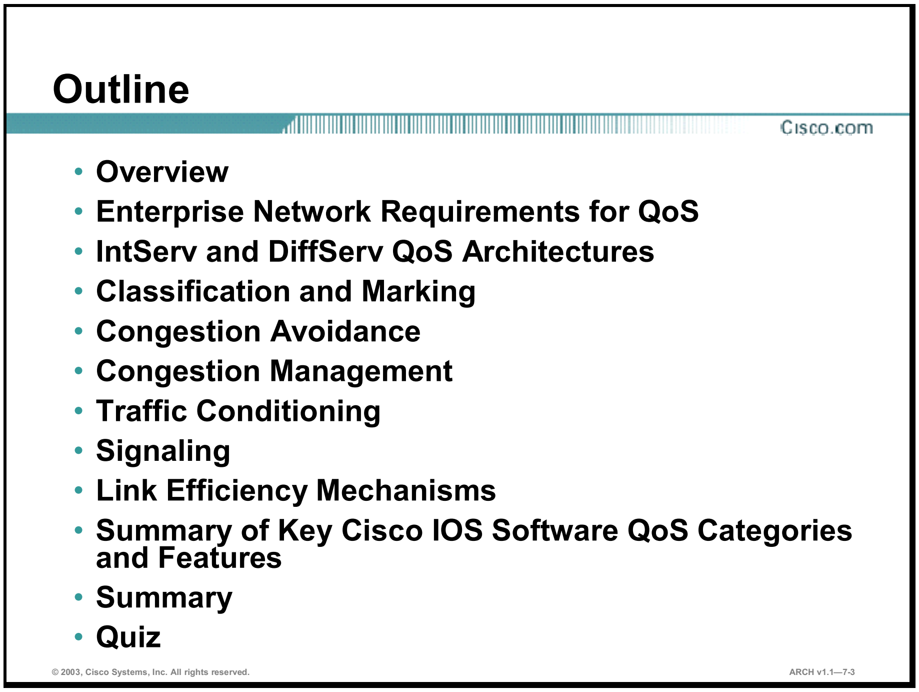

Outline

- Overview
- Enterprise Network Requirements for QoS
- IntServ and DiffServ QoS Architectures
- Classification and Marking
- Congestion Avoidance
- Congestion Management
- Traffic Conditioning
- Signaling
- Link Efficiency Mechanisms
- Summary of Key Cisco IOS Software QoS Categories and Features
- Summary
- Quiz

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—7-3

Enterprise Network Requirements for QoS

QoS tools are required to manage bandwidth and minimize loss, delay, and delay variation between enterprise sites and within a campus. Between sites, bandwidth availability is the most frequent concern. Within the campus infrastructure, buffer management issues dominate. This topic identifies enterprise requirements for quality of service.

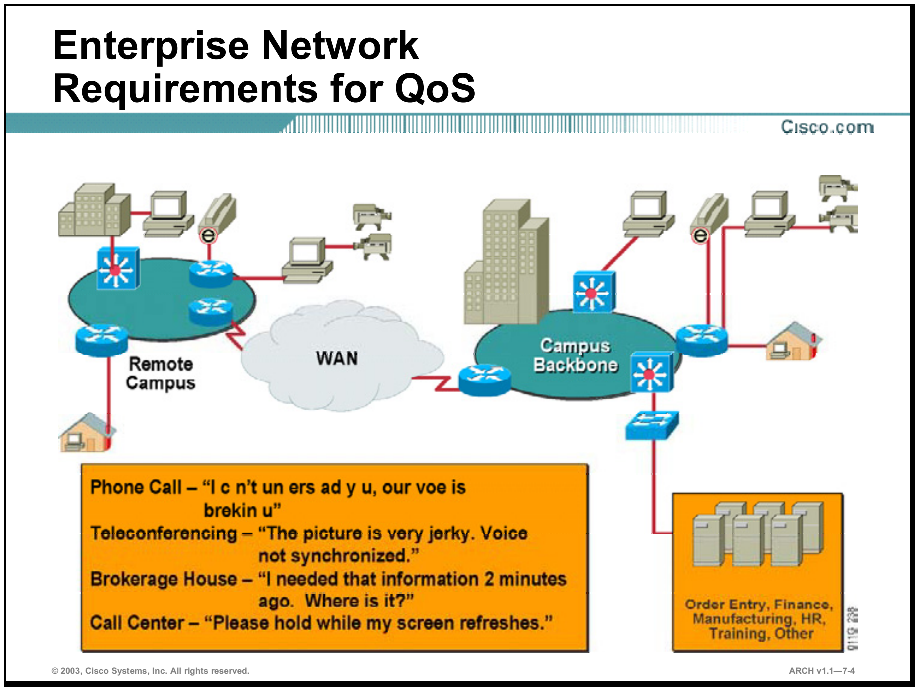

Enterprise Network Requirements for QoS

Phone Call – "I c n't un ers ad y u, our voe is brekin u"
Teleconferencing – "The picture is very jerky. Voice not synchronized."
Brokerage House – "I needed that information 2 minutes ago. Where is it?"
Call Center – "Please hold while my screen refreshes."

ARCH v1.1—7-4

QoS is the application of features and functionality needed to satisfy networking requirements for loss, delay, and delay variation (jitter) sensitive applications, and needed to guarantee the availability of bandwidth for critical application flows.

QoS provides control and predictable service for a variety of networked applications and traffic types in complex networks. Almost any network can take advantage of QoS to optimize efficiency. Effective QoS configurations provide these benefits:

- **Control over resources:** You have control over which resources (bandwidth, equipment, wide-area facilities, and so on) are being used. For example, you can limit the bandwidth consumed over a backbone link by FTP transfers or give priority to an important database access.
- **More effective use of network resources:** Using network analysis management and accounting tools, you will know what your network is being used for and you can configure the system to provide the most effective use of resources.
- **Tailored services:** The control and visibility provided by QoS enables carefully tailored grades of service differentiation to applications and customers within large enterprises and service provider networks.
- **Coexistence of mission-critical applications:** QoS technologies make certain that the WAN is used efficiently by mission-critical applications that are most important to the business, that bandwidth and minimum delays required by time-sensitive multimedia and voice applications are available, and that other applications using the link get fair service without interfering with mission-critical traffic.

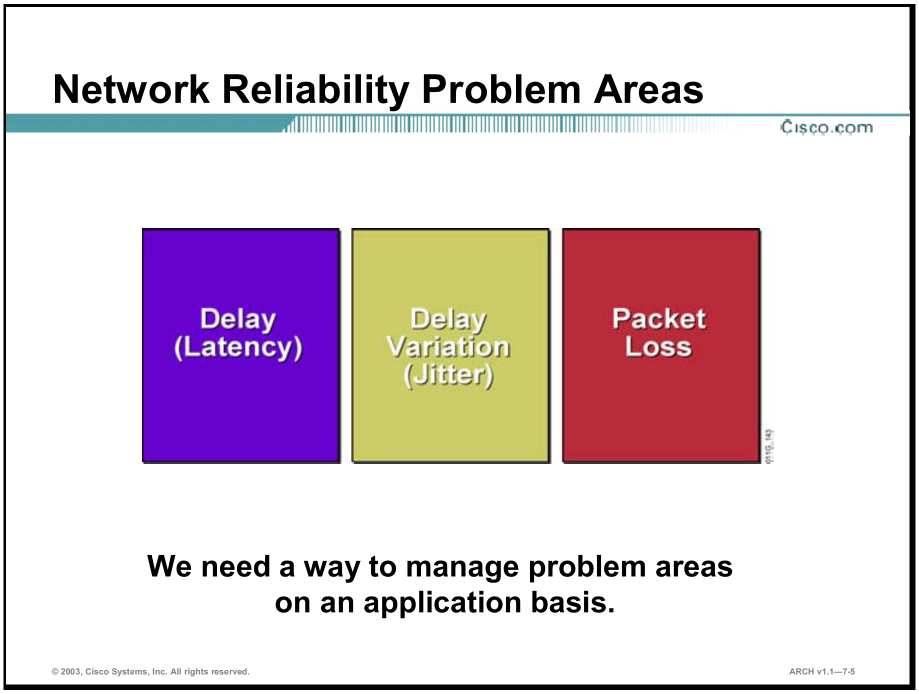

Network Reliability Problem Areas

Cisco.com

Delay (Latency)

Delay Variation (Jitter)

Packet Loss

We need a way to manage problem areas on an application basis.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-5

An enterprise network may experience any of these network reliability problems:

- **Delay:** Delay (or latency) is the amount of time it takes a packet to reach the receiving endpoint after being transmitted from the sending endpoint. This time period is referred to as the "end-to-end delay" and includes two components: fixed network delay and variable network delay. Fixed network delay includes encoding and decoding time (for voice and video), as well as the time required for the electrical and optical pulses to traverse the media en route to their destination. Variable network delay is generally caused by network conditions, such as congestion, that may affect the overall time required for transit.

  In converged data networks, there are three types of delay:

  — **Packetization delay:** The amount of time that it takes to packetize the content. With voice and video, this includes the time to sample and encode the analog signals.

  — **Serialization delay:** The amount of time that it takes to place the bits of the data packets onto the physical media.

  — **Propagation delay:** The amount of time it takes to transmit the bits of a packet across the physical media links.

- **Delay variation:** Delay variation (or jitter) is the difference in the end-to-end delay between packets. For example, if one packet required 100 ms to traverse the network from the source-endpoint to the destination-endpoint, and the following packet required 125 ms to make the same trip, then the delay variation is calculated as 25 ms.

  Each end station in a voice or video conversation has a jitter buffer. Jitter buffers are used to smooth out changes in arrival times of data packets containing voice. A jitter buffer is often dynamic and can adjust for approximately 30 ms changes in arrival times of packets. If you have instantaneous changes in arrival times of packets that are outside of the capabilities of a jitter buffer's ability to compensate, you will have one of these situations:

  — A jitter buffer underrun occurs when the arrival time of packets increase to the point where the jitter buffer is exhausted and contains no packets to be processed. The effect is unnatural silence in the case of voice, or a black screen in the case of video.

- A jitter buffer overrun occurs when packets containing voice or video arrive faster than the jitter buffer can dynamically resize itself to accommodate. When this happens, packets are dropped. When it is time to play voice or video samples, voice quality is degraded.

- **Loss:** Loss (or packet loss) is a comparative measure of packets faithfully transmitted and received to the total number that were transmitted. Loss is expressed as the percentage of packets that were dropped.

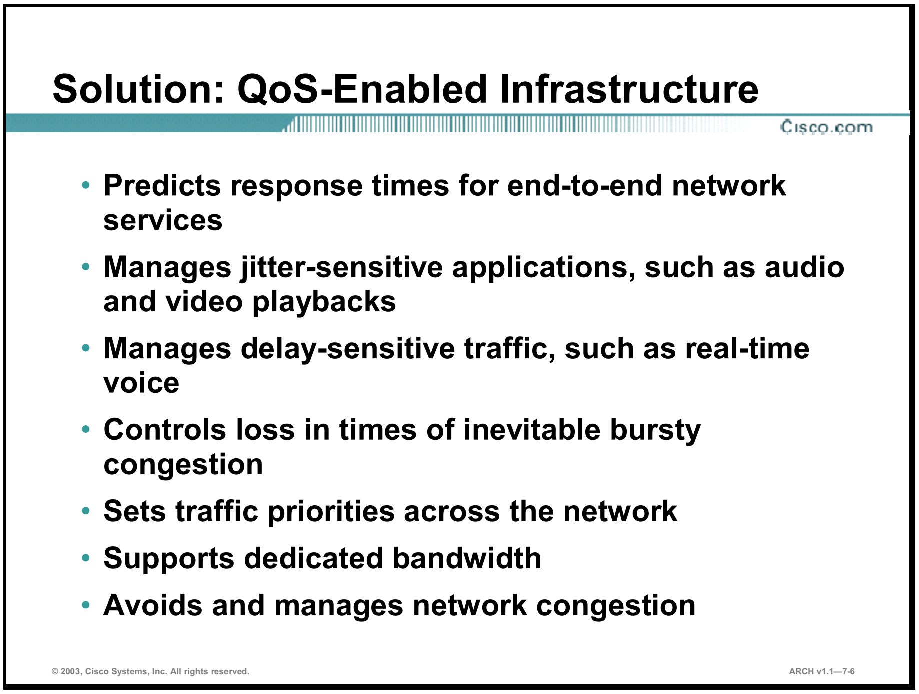

Solution: QoS-Enabled Infrastructure

- Predicts response times for end-to-end network services
- Manages jitter-sensitive applications, such as audio and video playbacks
- Manages delay-sensitive traffic, such as real-time voice
- Controls loss in times of inevitable bursty congestion
- Sets traffic priorities across the network
- Supports dedicated bandwidth
- Avoids and manages network congestion

Managing QoS becomes increasingly difficult in a converged network because many applications deliver individually unpredictable bursts of traffic. For example, usage patterns for web, e-mail, and file transfer applications are virtually impossible to predict, yet network managers need to be able to support mission-critical applications even during peak periods.

QoS technologies allow IT managers and network managers to:

- Predict response times for end-to-end network services
- Manage jitter-sensitive applications, such as audio and video playbacks
- Manage delay-sensitive traffic, such as real-time voice
- Control loss in times of inevitable bursty congestion
- Set traffic priorities across the network
- Support dedicated bandwidth
- Avoid and manage network congestion

IntServ and DiffServ QoS Architectures

The two QoS architectures used in IP networks when designing a QoS solution are the IntServ) and DiffServ models. In this topic, you will learn about the features of the IntServ and DiffServ models, and when to use each one.

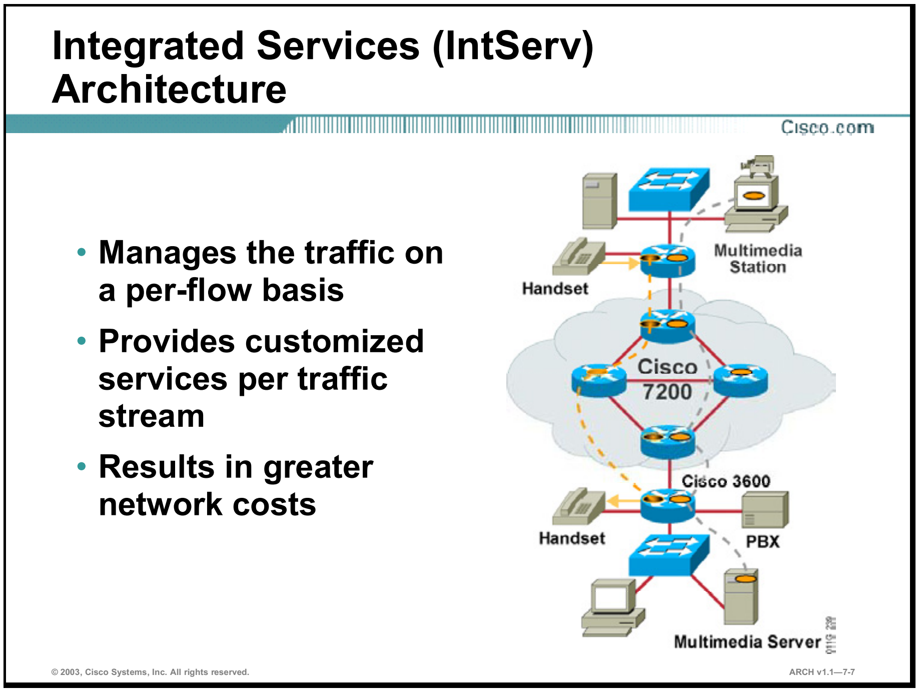

Integrated Services (IntServ) Architecture

Cisco.com

- Manages the traffic on a per-flow basis
- Provides customized services per traffic stream
- Results in greater network costs

ARCH v1.1—7-7

IntServ is a multiple service model that can accommodate multiple QoS requirements. The application requests a specific kind of service from the network before it sends data. Explicit signaling makes the request. The application informs the network of its traffic profile and requests a particular kind of service that can encompass its bandwidth and delay requirements. The application is expected to send data only after it gets a confirmation from the network. It is also expected to send data that lies within its described traffic profile.

The network performs admission control based on information from the application and available network resources. It also commits to meeting the QoS requirements of the application as long as the traffic remains within the profile specifications. The network fulfills its commitment by maintaining per-flow state and then performing packet classification, policing, and intelligent queuing based on that state.

Cisco IOS includes these features that provide controlled load service, which is a kind of integrated service:

- **RSVP:** Used by applications to signal their QoS requirements to the routers through the network.
- Intelligent queuing mechanisms: Used with RSVP to provide:
  - **Guaranteed rate service:** Allows applications to reserve bandwidth to meet their requirements. For example, a Voice over IP (VoIP) application can reserve 32 kbps end-to-end using this kind of service. IOS QoS uses weighted fair queuing (WFQ) with RSVP to provide this kind of service. Guaranteed rate service is implemented using a queue-service discipline.

- **Controlled load service:** Allows applications to have low delay and high throughput even during times of congestion. For example, adaptive real-time applications such as playback of a recorded conference can use this kind of service. IOS QoS uses RSVP with weighted random early detection (WRED) to provide this kind of service. Controlled load service is a queue-entry discipline that accelerates packet discard as congestion increases.

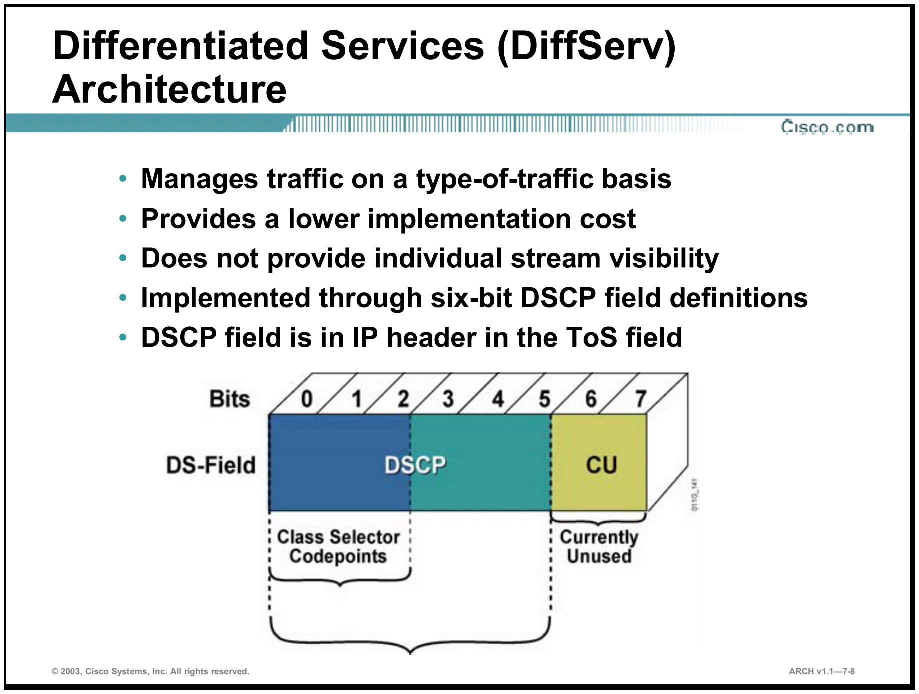

Differentiated Services (DiffServ) Architecture

- Manages traffic on a type-of-traffic basis
- Provides a lower implementation cost
- Does not provide individual stream visibility
- Implemented through six-bit DSCP field definitions
- DSCP field is in IP header in the ToS field

ARCH v1.1—7-8

Differentiated service is a multiple service model that can satisfy differing QoS requirements. However, unlike in the IntServ model, an application using differentiated service does not explicitly signal the router before sending data.

For differentiated service, the network tries to deliver a particular kind of service based on the QoS per-hop behavior (PHB) associated with the differential services code point (DSCP) within each packet. You can base the DSCP assignment on different criteria, for example, using the IP precedence bit settings in IP packets or source and destination addresses. The network uses the QoS specification to classify, mark, shape, and police traffic, and to perform intelligent queuing.

You can use the differentiated service model for critical applications and to provide end-to-end QoS to traffic aggregates. This service model performs a relatively coarse level of traffic classification, and no information about individual flows is required in the network. DiffServ consumes less network resources than IntServ.

Cisco IOS QoS includes these features that support the differentiated services model:

- **Committed access rate (CAR):** Performs packet classification through IP precedence and QoS group settings. CAR performs metering and policing of traffic, providing bandwidth management.
- **Intelligent queuing schemes:** Includes distributed WRED (DWRED) and distributed WFQ and their equivalent features on the Versatile Interface Processor (VIP). You can use these features with CAR to deliver differentiated services.

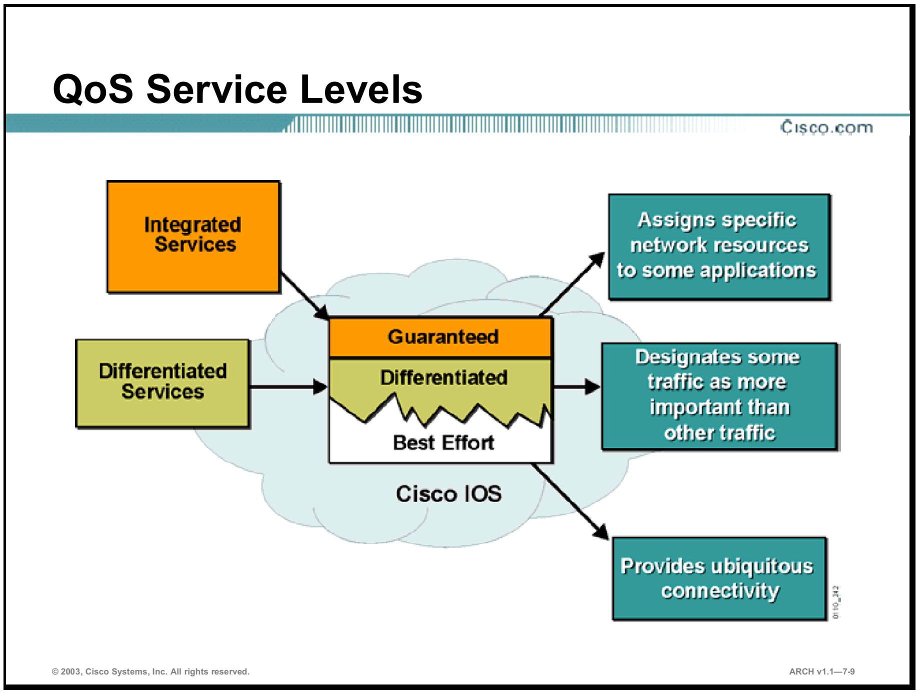

QoS Service Levels

Cisco.com

Integrated Services

Differentiated Services

Guaranteed

Differentiated

Best Effort

Cisco IOS

Assigns specific network resources to some applications

Designates some traffic as more important than other traffic

Provides ubiquitous connectivity

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-9

The services differ in their level of QoS strictness, which describes how tightly the services are bound by specific bandwidth, delay, jitter, and loss characteristics.

Three basic levels of end-to-end QoS can be provided across a heterogeneous network:

- **Best-effort service (also called lack of QoS):** Best-effort service is basic connectivity with no guarantees. This is often characterized by queues that have no differentiation between flows.
- **Differentiated service (also called soft QoS):** Some traffic classes (aggregates) are treated better than the rest (faster handling, more average bandwidth, and lower average loss rate). This is a statistical preference, not a hard-and-fast guarantee to any particular traffic flow. You can provide differentiated services by classifying traffic and using QoS tools such as priority queuing (PQ), custom queuing (CQ), WFQ, and WRED.
- **Guaranteed service (also called hard QoS):** This is an absolute reservation of network resources for specific traffic. Guaranteed services use QoS tools including RSVP and class-based weighted fair queuing (CBWFQ).

Selecting the type of service to deploy in the network depends on:

- **Application supported or problem being solved:** Each of the three types of service is appropriate for certain applications. This does not imply that an enterprise must migrate to differentiated and then to guaranteed service (although many probably eventually will). A differentiated service, or even a best-effort service, may be appropriate, depending on the application requirements.
- **Speed to upgrade the infrastructure:** There is a natural upgrade path from the technology needed to provide differentiated service to that needed to provide guaranteed service.
- **Cost:** The cost of implementing and deploying guaranteed service is likely to be more than that for differentiated service.

Classification and Marking

Classification tools mark packets with a value used for prioritization, shaping, or policing within the network. This marking establishes a trust boundary that must be enforced. This topic helps you select the classification and marking components of a Cisco QoS solution, given specific quality and application requirements.

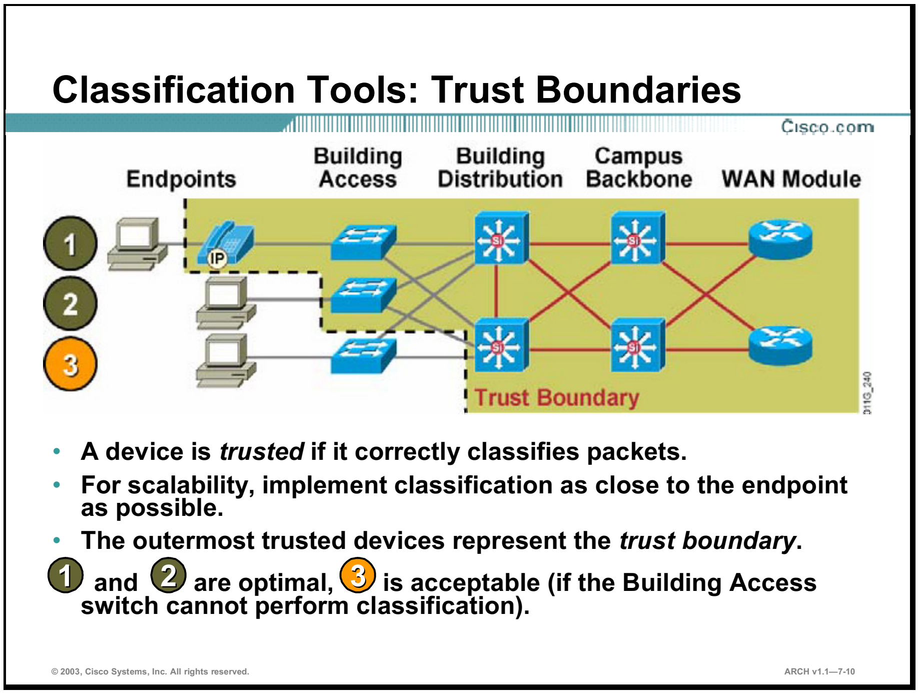

Classification Tools: Trust Boundaries

- A device is *trusted* if it correctly classifies packets.
- For scalability, implement classification as close to the endpoint as possible.
- The outermost trusted devices represent the *trust boundary*.

① and ② are optimal, ③ is acceptable (if the Building Access switch cannot perform classification).

ARCH v1.1—7-10

The first element of a QoS policy is to identify the traffic that is to be treated differently. Classification tools mark a packet or flow with a specific identifier. Classification at the trust boundary marks packets by examining any of the following:

- **Layer 2 parameters:** 802.1Q class of service (CoS) bits, MAC address, Multiprotocol Label Switching (MPLS) label
- **Switching:** MPLS experimental values
- **Layer 3 parameters:** IP precedence, DSCPs, source or destination address, protocol
- **Layer 4 parameters:** TCP or User Datagram Protocol (UDP) ports
- **Layer 7 parameters:** Application signatures

You can apply policy only after traffic is positively identified. You should identify and mark traffic as close to the source of the traffic as possible. The network edge where markings are accepted (or rejected) is referred to as the *trust boundary*.

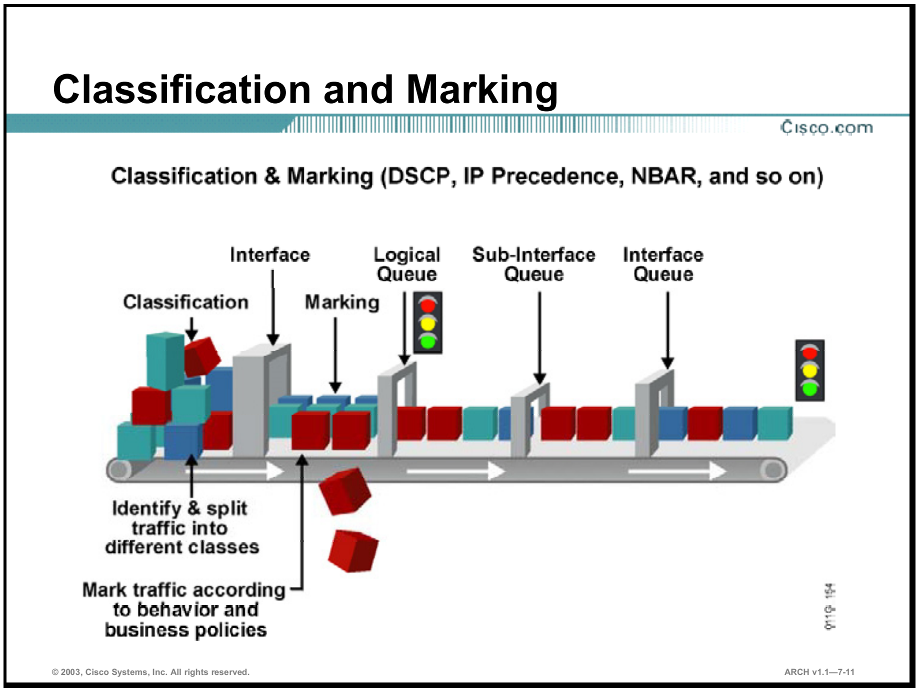

Classification and Marking

Cisco.com

Classification & Marking (DSCP, IP Precedence, NBAR, and so on)

Interface

Logical Queue

Sub-Interface Queue

Interface Queue

Classification

Marking

Identify & split traffic into different classes

Mark traffic according to behavior and business policies

9119-154

ARCH v1.1—7-11

Packet classification features provide the capability to partition network traffic into multiple priority levels or classes of service. For example, by using the three precedence bits in the CoS field of the IP packet header (two of the values are reserved for other purposes), you can categorize packets into a limited set of up to six traffic classes. After you classify packets, you can utilize other QoS features to assign the appropriate traffic handling policies, including congestion management, bandwidth allocation, and delay bounds for each traffic class.

You can also classify packets by external sources, that is, by a location or by a downstream network provider. You can either allow the network to accept the classification or override it and reclassify the packet according to a policy that you specify.

You can set policies such as classification based on physical port, source, or destination IP or MAC address, application port, IP protocol type, and other criteria that you can specify by using access lists or extended access lists.

Congestion Avoidance

Congestion avoidance techniques monitor network traffic loads in an effort to anticipate and reduce the impact of congestion at common network and internetwork bottlenecks before it becomes a significant problem. These techniques are designed to provide preferential treatment for premium (priority) traffic under congestion situations while concurrently maximizing network throughput and capacity utilization and minimizing packet loss and delay. WRED and DWRED are the IOS QoS congestion avoidance features. This topic helps you select the congestion avoidance components of a Cisco QoS solution, given specific quality and application requirements.

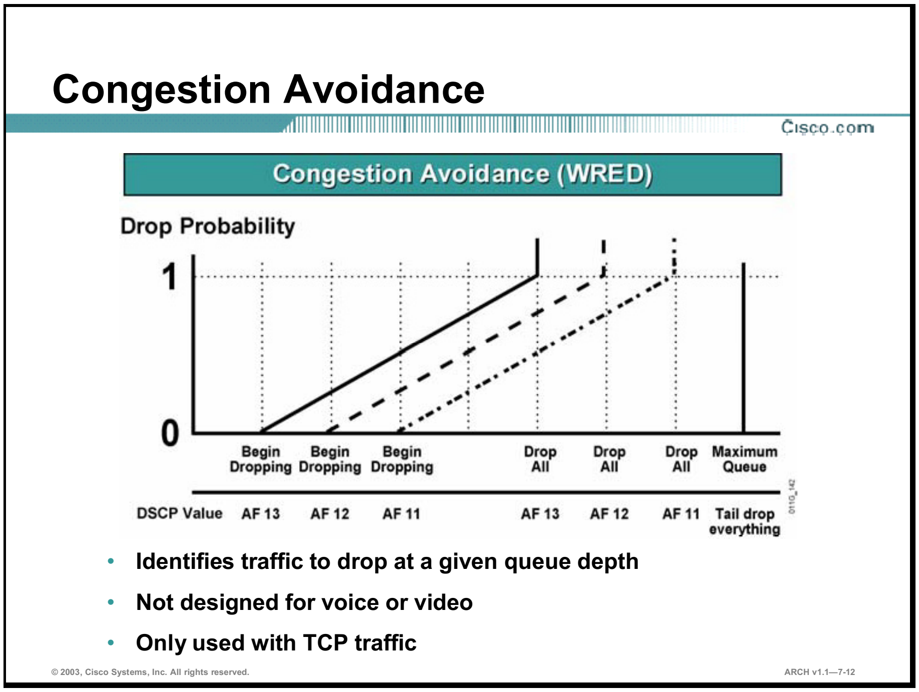

Congestion Avoidance

Cisco.com

Congestion Avoidance (WRED)

- **Identifies traffic to drop at a given queue depth**
- **Not designed for voice or video**
- **Only used with TCP traffic**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-12

Default router behavior allows interface queues to fill during periods of congestion using tail drop to resolve the problem of full queues, unless WRED is configured. When the queue has filled, a potentially large number of packets from numerous connections are discarded because of lack of buffer capacity. This behavior can result in waves of congestion followed by periods during which the transmission link is not fully used. WRED mitigates this situation proactively, and preferentially. Congestion avoidance for selected traffic is provided by monitoring buffer depth and performing a probabilistic (random) discard on packets from traffic streams configured for early discard, instead of waiting for buffers to fill and dropping all arriving packets.

WRED is a Cisco implementation of the random early detection (RED) class of congestion avoidance algorithms. When RED is used with TCP and the source detects the dropped packet, the source slows its transmission. WRED can also be configured to use the DSCP value when it calculates the drop probability of a packet, enabling WRED to be compliant with the DiffServ standard being developed by the IETF.

WRED combines the capabilities of the RED algorithm with IP precedence to provide preferential traffic handling for higher-priority packets. It can selectively discard lower-priority traffic when the interface begins to get congested and provide differentiated performance characteristics for different classes of service. WRED is also RSVP-aware. DWRED is the high-speed version of WRED. The DWRED algorithm was designed with service providers in mind; it allows an ISP to define minimum and maximum queue depth thresholds and drop capabilities for each class of service. It is implemented on VIP-capable routers.

The flow-based RED feature forces RED to afford greater fairness to all flows on an interface in regard to how packets are dropped.

To provide fairness to all flows, flow-based RED has these features:

- It ensures that flows that respond to RED packet drops by backing off packet transmission are protected from flows that do not respond to WRED packet drops.
- It prohibits a single flow from monopolizing the buffer resources at an interface.

The DiffServ-compliant WRED feature extends the functionality of WRED to enable support for DiffServ and Assured Forwarding (AF) per-hop behavior (PHB). This feature enables customers to implement AF PHB by coloring packets according to DSCP values and then assigning low, preferential drop probabilities to those packets.

Congestion Management

Congestion management features control congestion once it occurs. One way that network elements handle an overflow of arriving traffic is to use a queuing algorithm to sort the traffic into one or more logical queues, and then determine a method of prioritizing it onto an output link. Each queuing algorithm solves a specific network traffic problem and has a particular effect on network performance. This topic helps you select congestion management components of a Cisco QoS solution, given specific quality and application requirements.

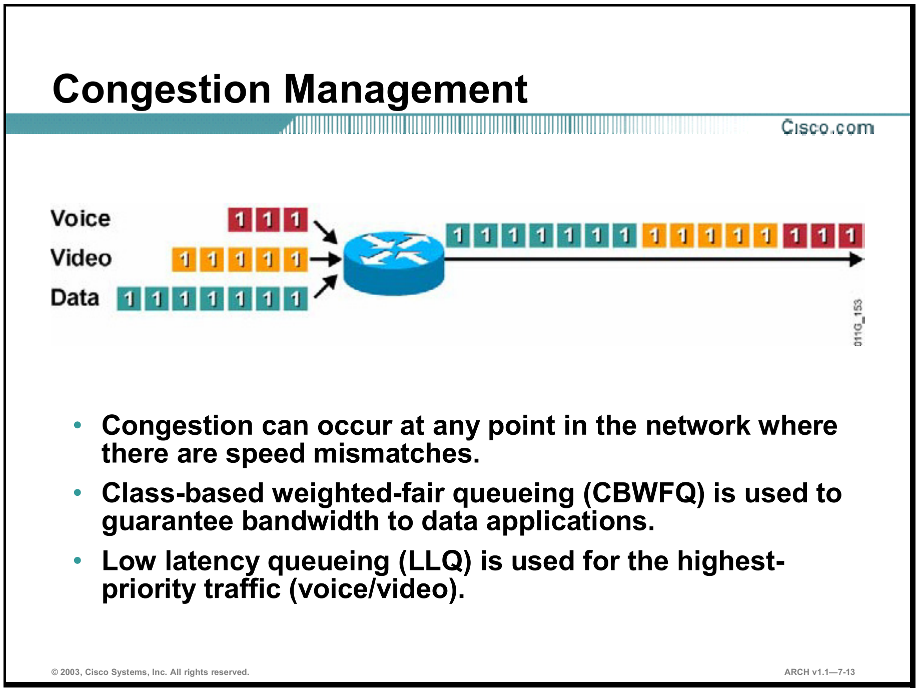

Congestion Management

Cisco.com

- **Congestion can occur at any point in the network where there are speed mismatches.**
- **Class-based weighted-fair queueing (CBWFQ) is used to guarantee bandwidth to data applications.**
- **Low latency queueing (LLQ) is used for the highest-priority traffic (voice/video).**

ARCH v1.1—7-13

The software congestion management features include the following:

- **First-in, first-out (FIFO):** Provides basic store and forward capability. FIFO is the default queuing algorithm on high-speed interfaces, requiring no configuration.
- **Weighted fair queuing (WFQ) (flow-based, class-based, and distributed):** Applies priority (or weights) to identified traffic to classify traffic into conversations and determine how much bandwidth each conversation is allowed relative to other conversations. WFQ classifies traffic into different flows based on such characteristics as source and destination address, protocol, and port and socket of the session. WFQ is the default queuing discipline on links at and below 2.048 Mbps.

  To provide large-scale support for applications and traffic classes requiring bandwidth allocations and delay bounds over the network infrastructure, IOS QoS includes a version of WFQ that runs only in distributed mode on VIPs. This version is called VIP-distributed WFQ (DWFQ). It provides increased flexibility in terms of traffic classification, weight assessment, and discard policy, and delivers Internet-scale performance on the Cisco 7500 series platforms.

- **Low latency queueing (LLQ), distributed LLQ, and LLQ for Frame Relay:** LLQ provides strict priority queuing on ATM virtual circuits (VCs) and serial interfaces. This feature allows you to configure the priority status for a class within CBWFQ, and is not limited to UDP port numbers, as is IP RTP Priority. LLQ and IP RTP Priority can be configured at the same time, but IP RTP Priority takes precedence.

  Additionally, the functionality of LLQ has been extended to allow you to specify the committed burst (Bc) size in LLQ and to change (or vary) the number of packets contained in the per-VC hold queue (on ATM adapters that support per-VC queuing).

  The distributed LLQ feature provides the ability to specify low latency behavior for a traffic class on a VIP-based Cisco 7500 series router. LLQ allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued. The distributed LLQ feature also introduces the ability to limit the depth of a device transmission ring.

  LLQ for Frame Relay provides strict PQ for voice traffic and WFQs for other classes of traffic. Before the release of this feature, LLQ was available at the interface and ATM VC levels. It is now available at the Frame Relay VC level when Frame Relay traffic shaping is configured. Strict PQ improves QoS by allowing delay-sensitive traffic such as voice to be pulled from the queue and sent before other classes of traffic. LLQ for Frame Relay allows you to define classes of traffic according to protocol, interface, or access lists. You can then assign characteristics to those classes, including priority, bandwidth, queue limit, and WRED.
- **PQ:** Designed to give strict priority to important traffic, PQ ensures that important traffic gets the fastest handling at each point where PQ is used. PQ can flexibly prioritize according to network protocol (such as IP, Internetwork Packet Exchange [IPX], or AppleTalk), incoming interface, packet size, source/destination address, and so on. A strict priority scheme can starve traffic flows with a lower priority. There is no minimum service guarantee.
- **Frame Relay permanent virtual circuit (PVC) Interface Priority Queuing (FR PIPQ):** Provides an interface-level PQ scheme in which prioritization is based on destination PVC rather than packet contents. For example, FR PIPQ allows you to configure PVC transporting voice traffic to have absolute priority over a PVC transporting signaling traffic, and a PVC transporting signaling traffic to have absolute priority over a PVC transporting data.

  FR PIPQ provides four levels of priority: high, medium, normal, and low. The Frame Relay packet is examined at the interface for the data-link connection identifier (DLCI) value. The packet is then sent to the correct priority queue based on the priority level configured for that DLCI.
- **CQ:** Reserves a percentage of the available bandwidth of an interface for each selected traffic type. If a particular type of traffic is not using the bandwidth reserved for it, other traffic types may use the remaining reserved bandwidth.
- **CBWFQ and distributed CBWFQ (DCBWFQ):** CBWFQ and DCBWFQ extend the standard WFQ functionality to provide support for user-defined traffic classes. They allow you to specify the exact amount of bandwidth to be allocated for a specific class of traffic. Taking into account available bandwidth on the interface, you can configure up to 64 classes and control distribution among them. DCBWFQ is intended for use on the VIP-based Cisco 7000 series routers with the Route Switch Processors (RSPs), and the Cisco 7500 series routers.

- **IP RTP Priority and Frame Relay IP RTP Priority:** The IP RTP Priority feature provides a strict priority queuing scheme that allows delay-sensitive data such as voice to be dequeued and sent before packets in other queues are dequeued. Use this feature on serial interfaces and Frame Relay PVCs in conjunction with either WFQ or CBWFQ on the same outgoing interface. In either case, traffic matching the range of UDP ports specified for the priority queue is guaranteed strict priority over other CBWFQ classes or WFQ flows; packets in the priority queue are always serviced first.

  The Frame Relay IP RTP Priority feature provides a strict priority queuing scheme on a Frame Relay PVC for delay-sensitive data such as voice. Voice traffic can be identified by its Real-Time Transport Protocol (RTP) port numbers and classified into a priority queue. The result of using this feature is that voice is serviced as strict priority in preference to other nonvoice traffic.

- **Modified Deficit Round Robin (MDRR):** MDRR is a traffic latency control function that allows the operators to guarantee traffic latency for differentiated flows by controlling the packet dequeuing process. Packet classification is based on IP precedence. There are two basic modes of operation that govern how packets are dequeued from the low latency queue in relation to other queues; they are:

  — **Alternate priority:** Queues are serviced by alternating between the low latency queue and the other queues in round robin.

  — **Strict priority:** The low-latency queue is continually serviced to keep it empty.

Consideration of the behavior of congested systems is not simple because traffic rates do not simply rise to a level, stay there a while, then subside. Periods of traffic congestion can be quite long, with losses that are heavily concentrated. A slight increase in the number of active connections can result in a large increase in the packet loss rate. This understanding of the behavior of congested networks suggests that, because the level of busy period traffic is not predictable, it is difficult to economically size networks to reduce congestion.

Traffic Conditioning

The IOS QoS software solutions include two traffic-shaping tools, Generic Traffic Shaping (GTS) and Frame Relay traffic shaping (FRTS), to manage traffic and congestion on the network. The IOS policing tool is CAR. This topic helps you select the traffic conditioning components of a Cisco QoS solution, given specific quality and application requirements.

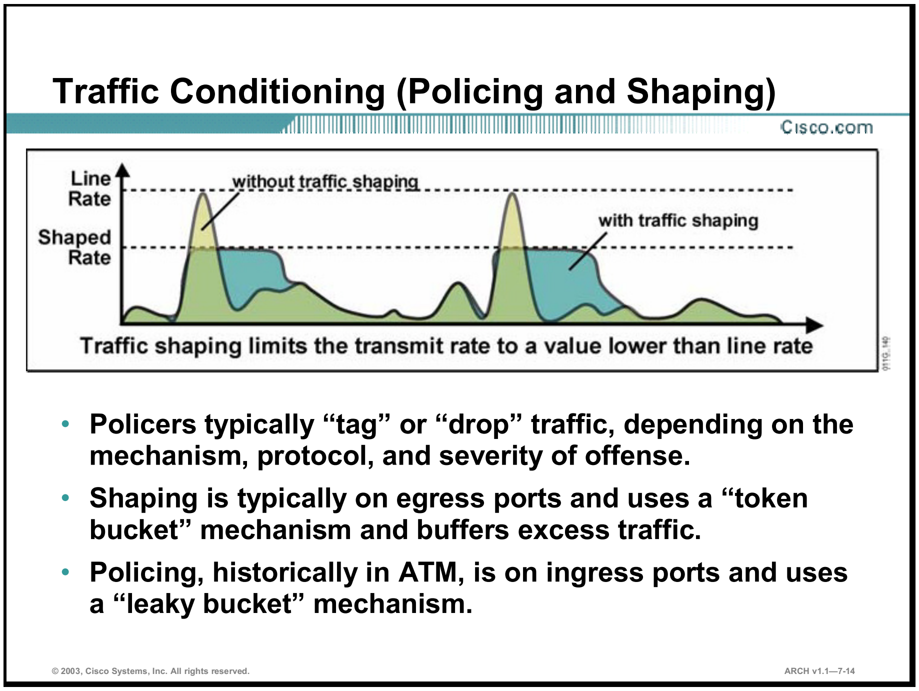

Traffic Conditioning (Policing and Shaping)

Traffic shaping limits the transmit rate to a value lower than line rate

- Policers typically "tag" or "drop" traffic, depending on the mechanism, protocol, and severity of offense.
- Shaping is typically on egress ports and uses a "token bucket" mechanism and buffers excess traffic.
- Policing, historically in ATM, is on ingress ports and uses a "leaky bucket" mechanism.

ARCH v1.1—7-14

CAR: Managing Access Bandwidth Policy and Performing Policing

QoS provides priority either by raising the priority of one flow or by limiting the priority of another. CAR is used to limit the bandwidth of a flow in order to favor another flow.

The IOS implementation of CAR transmits, drops, sets IP precedence bits, and continues (this refers to cascading CAR statements). This flexibility allows for a number of ways to act upon traffic. Conforming traffic can be transmitted, and exceeding traffic can be reclassified to a lower IP precedence setting and then sent to the next CAR statement for additional conditions.

The IOS CAR implementation also provides an excess burst capability. Excess burst allows additional tokens above the original (or normal) burst. When these tokens are used, the packet has the possibility of being dropped (even if the action is to transmit). A RED-like algorithm is used that says, "The more excess burst tokens you use, the higher probability that the next packet will be dropped." This allows the flow to be scaled back slowly as in WRED, while still maintaining the opportunity to send traffic in excess of the normal burst.

Traffic Shaping: Controlling Outbound Traffic Flow

QoS software contains these traffic shaping features to manage traffic and congestion on the network:

- **GTS:** Provides a mechanism to control the flow of outbound traffic on a particular interface. It reduces outbound traffic flow to avoid congestion by constraining specified traffic to a particular bit rate. Traffic adhering to a particular profile can be shaped to meet downstream requirements, eliminating bottlenecks in topologies with data-rate mismatches.
- **Class-based shaping:** Provides the means for configuring GTS on a class, rather than only on an access control list (ACL). You can enable class-based shaping on any interface that supports GTS. Using the class-based shaping feature, you can perform these tasks:
  - Configure GTS on a traffic class.
  - Specify average rate or peak rate traffic shaping.
  - Configure CBWFQ inside GTS.
- **Distributed traffic shaping (DTS):** Provides the means for managing the bandwidth of an interface to avoid congestion, to meet remote site requirements, and to conform to a service rate that is provided on that interface. DTS uses queues to buffer traffic surges that can congest a network.
- **Frame Relay traffic shaping (FRTS):** Provides parameters that are useful for managing network traffic congestion, such as:
  - Committed information rate (CIR)
  - Forward and backward explicit congestion notification (FECN/BECN)
  - The discard eligible (DE) bit

FRTS applies only to Frame Relay permanent PVCs and switched virtual circuits (SVCs).

Signaling

QoS signaling is a form of network communication that allows an end station or network node to communicate with, or signal, its neighbors to request special handling of certain traffic. This topic helps you select the signaling components of a Cisco QoS solution, given specific quality and application requirements.

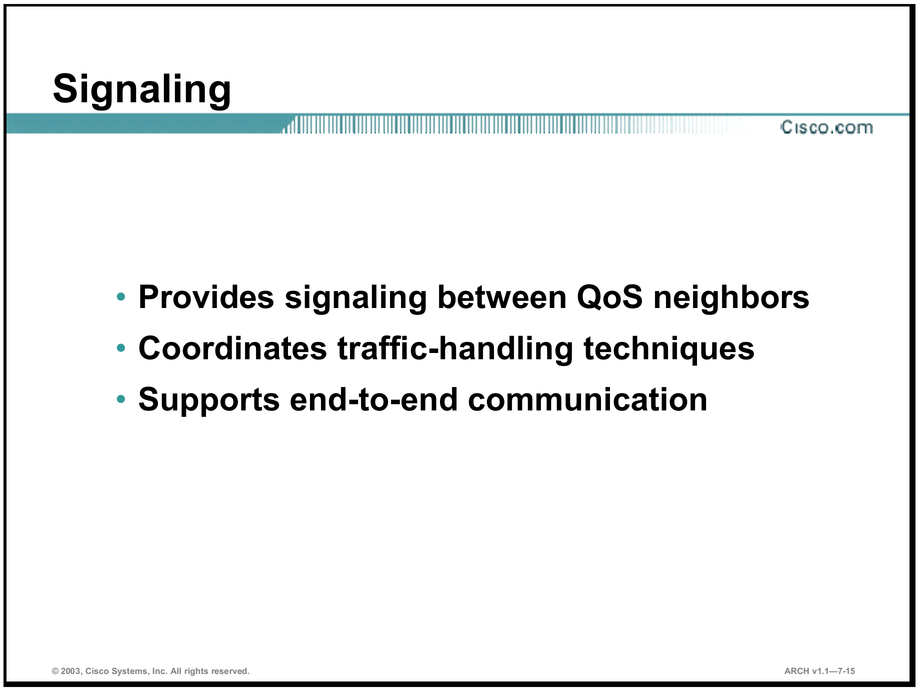

Signaling

Cisco.com

- Provides signaling between QoS neighbors
- Coordinates traffic-handling techniques
- Supports end-to-end communication

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-15

QoS signaling is useful for coordinating the traffic handling techniques provided by other QoS features. It plays a key role in configuring successful overall end-to-end QoS service across your network.

True end-to-end QoS requires that every element in the network path, including switches, routers, firewalls, hosts, and clients, deliver its part of QoS, and that all of these entities be coordinated with QoS signaling.

Many viable QoS signaling solutions provide QoS at some places in the infrastructure, but they often have limited scope across the network. To achieve end-to-end QoS, signaling must span the entire network.

IOS QoS software takes advantage of IP to meet the challenge of finding a robust QoS signaling solution that can operate over heterogeneous network infrastructures. It overlays data link layer technology-specific QoS signaling solutions with network layer IP QoS signaling methods of the RSVP and IP precedence features.

An IP network can achieve end-to-end QoS, for example, by using part of the IP packet header to request special handling of priority or time-sensitive traffic. Given the ubiquity of IP, QoS signaling that takes advantage of IP provides powerful end-to-end signaling. Both RSVP and IP precedence fit this category.

QoS signaling provides a way for an end station or network node to signal its neighbors to request special handling of certain traffic. QoS signaling is useful for coordinating the traffic handling techniques provided by other QoS features. It plays a key role in configuring successful overall end-to-end QoS service across your network.

To achieve the end-to-end benefits of IP precedence and RSVP signaling, QoS software offers ATM User-Network Interface (UNI) signaling and the Frame Relay Local Management Interface (LMI) to provide signaling into their respective backbone technologies.

To achieve centralized monitoring and control of RSVP signaling, IOS software offers Common Open Policy Service (COPS) with RSVP.

To enable admission control over IEEE 802-styled networks, IOS QoS software offers Subnetwork Bandwidth Manager (SBM).

To provide support for controlled load service using RSVP over an ATM core network, IOS QoS software offers the RSVP-ATM QoS Internetworking feature.

Link Efficiency Mechanisms

QoS software offers three link efficiency mechanisms that work in conjunction with queuing and traffic shaping to improve the predictability of the application services levels: Link fragmentation and interleaving (LFI), Compressed Real-Time Transfer Protocol (cRTP), and Distributed Compressed Real-Time Transfer Protocol (dcRTP). This topic helps you select the link efficiency components of a Cisco QoS solution, given specific quality and application requirements.

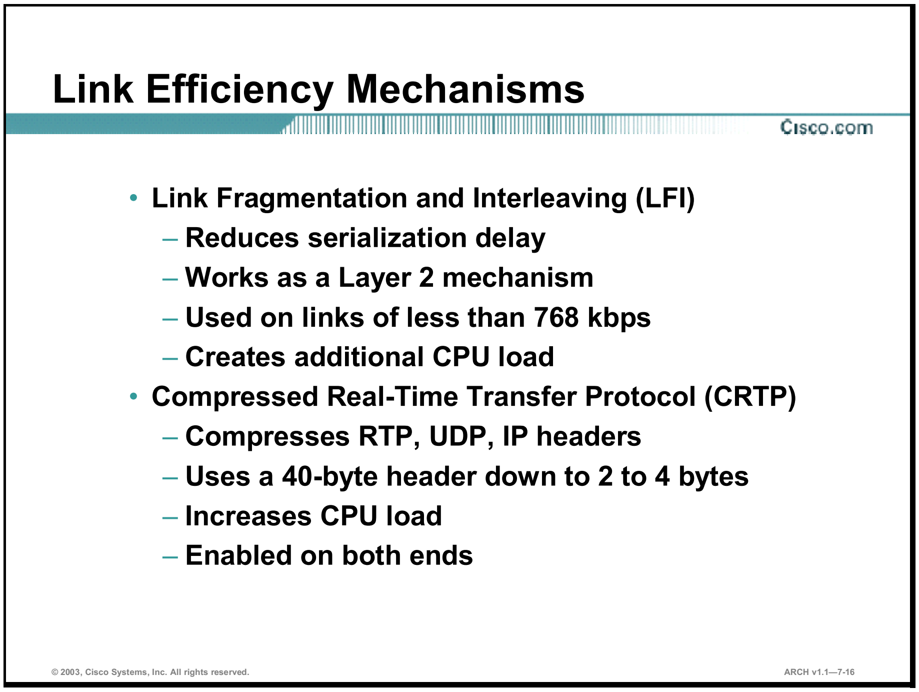

Link Efficiency Mechanisms

- Link Fragmentation and Interleaving (LFI)
  - Reduces serialization delay
  - Works as a Layer 2 mechanism
  - Used on links of less than 768 kbps
  - Creates additional CPU load
- Compressed Real-Time Transfer Protocol (CRTP)
  - Compresses RTP, UDP, IP headers
  - Uses a 40-byte header down to 2 to 4 bytes
  - Increases CPU load
  - Enabled on both ends

ARCH v1.1—7-16

Link Fragmentation and Interleaving

Interactive traffic, such as Telnet and VoIP, is susceptible to increased latency and jitter when the network processes large packets, such as LAN-to-LAN FTP transfers traversing a WAN link. This susceptibility increases as the traffic is queued on slower links. QoS LFI reduces delay and jitter on slower speed links by breaking up large datagrams and interleaving low-delay traffic packets with the resulting smaller packets. Alternatives to using LFI include:

- Reduced maximum transmission unit (MTU)
- Frame Relay Fragmentation (FRF.12)
- ATM

Using LFI with multilink PPP reduces 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.

Compressed Real-Time Protocol

RTP is a host-to-host protocol used for carrying newer multimedia application traffic, including packetized audio and video, over an IP network. RTP provides end-to-end network transport functions intended for applications sending real-time requirements, such as audio, video, or simulation data multicast or unicast network services.

To avoid the unnecessary consumption of available bandwidth, the RTP header compression feature, referred to as cRTP, is used on a link-by-link basis. The cRTP feature compresses the combined 40-byte IP/UDP/RTP packet headers to two to four bytes. It must be configured on both ends of each link on which it is desired. This compression reduces the packet size, improves the speed of packet transmission, and reduces packet latency.

Distributed Compressed Real-Time Protocol

The dcRTP feature is the implementation of cRTP on a Cisco 7500 series router with a VIP in distributed fast-switching and distributed Cisco Express Forwarding (dCEF) environments. It offers the same benefits and costs.

Summary of Key IOS Software QoS Categories and Features

Each QoS feature has an important role to play in the network. This topic summarizes the key IOS software QoS features so you can determine when to use each one.

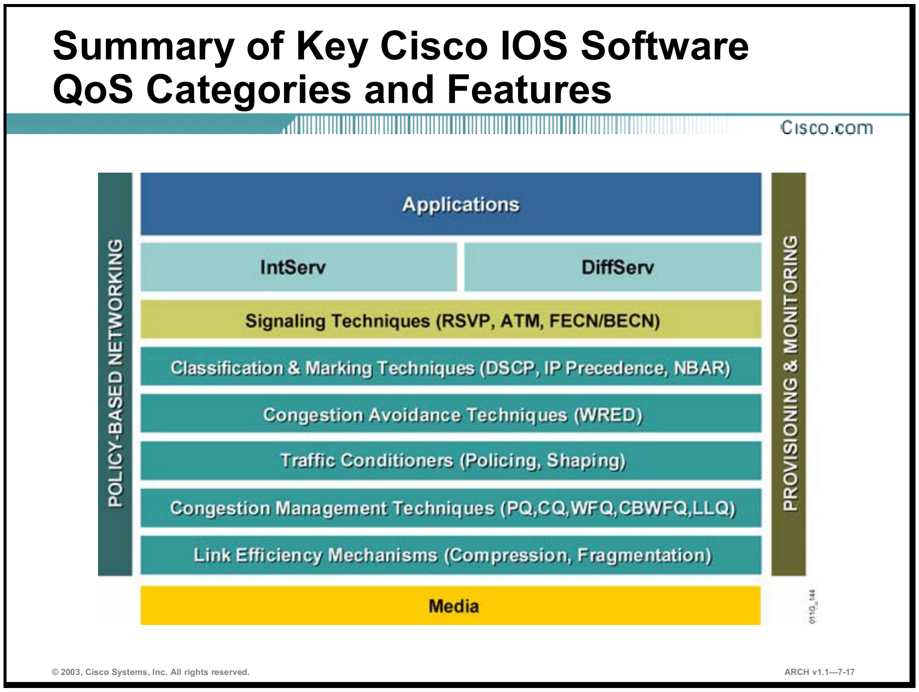

Summary of Key Cisco IOS Software QoS Categories and Features

Cisco.com

POLICY-BASED NETWORKING

Applications

IntServ

DiffServ

Signaling Techniques (RSVP, ATM, FECN/BECN)

Classification & Marking Techniques (DSCP, IP Precedence, NBAR)

Congestion Avoidance Techniques (WRED)

Traffic Conditioners (Policing, Shaping)

Congestion Management Techniques (PQ,CQ,WFQ,CBWFQ,LLQ)

Link Efficiency Mechanisms (Compression, Fragmentation)

Media

PROVISIONING & MONITORING

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-17

The figure summarizes the key IOS software QoS categories and features. The base layer contains the transport protocols.

The middle topic lists the tools that are used in deploying the QoS: classification and marking, congestion avoidance, traffic conditioners, congestion management, and link efficiency.

The top row contains the different applications that benefit from QoS.

Below the top row are IntServ and DiffServ. In the bars below that are RSVP and DSCP, the two marking tools of IntServ and DiffServ respectively. The IntServ architecture defines fine-grained (flow-based) methods of performing IP traffic admission control that uses RSVP. The DiffServ architecture defines methods of classifying IP traffic into coarse-grained service classes and defines forwarding treatment based on these classifications.

Summary

This topic summarizes the key points discussed in this lesson.

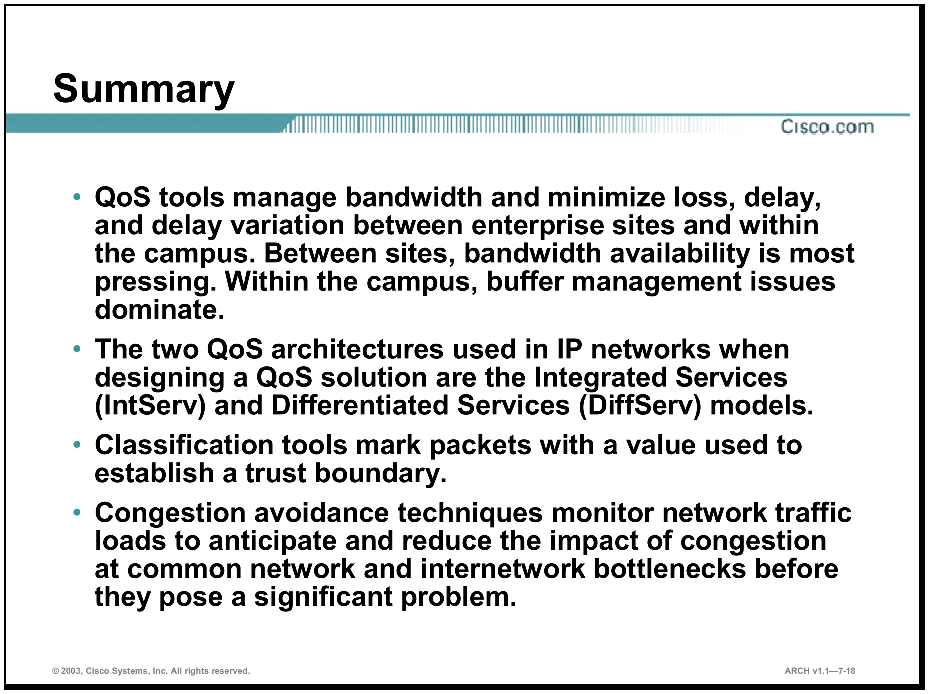

Summary

Cisco.com

- QoS tools manage bandwidth and minimize loss, delay, and delay variation between enterprise sites and within the campus. Between sites, bandwidth availability is most pressing. Within the campus, buffer management issues dominate.
- The two QoS architectures used in IP networks when designing a QoS solution are the Integrated Services (IntServ) and Differentiated Services (DiffServ) models.
- Classification tools mark packets with a value used to establish a trust boundary.
- Congestion avoidance techniques monitor network traffic loads to anticipate and reduce the impact of congestion at common network and internetwork bottlenecks before they pose a significant problem.

ARCH v1.1—7-18

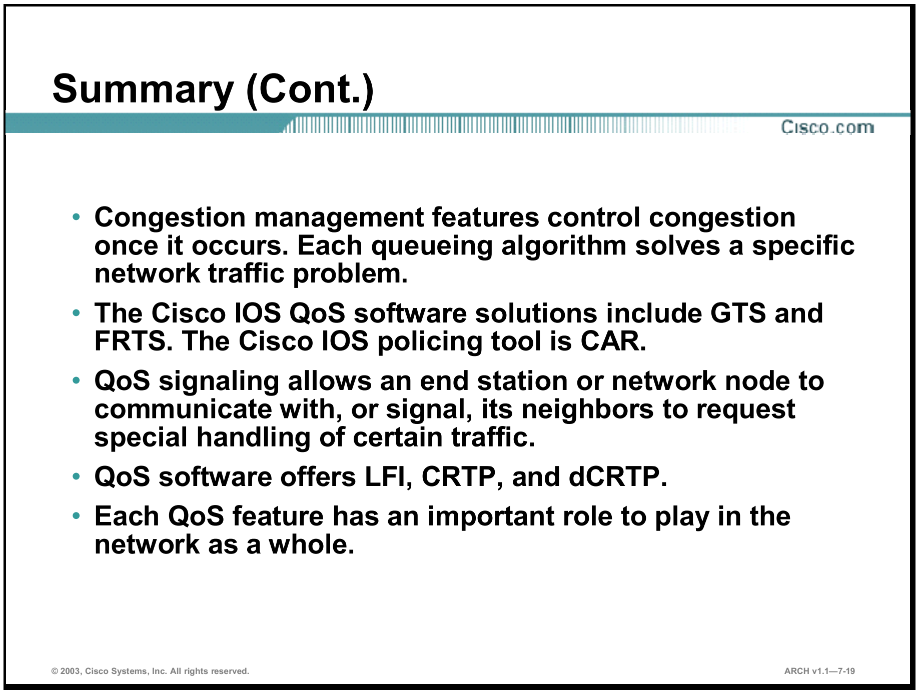

Summary (Cont.)

Cisco.com

- Congestion management features control congestion once it occurs. Each queueing algorithm solves a specific network traffic problem.
- The Cisco IOS QoS software solutions include GTS and FRTS. The Cisco IOS policing tool is CAR.
- QoS signaling allows an end station or network node to communicate with, or signal, its neighbors to request special handling of certain traffic.
- QoS software offers LFI, CRTP, and dCRTP.
- Each QoS feature has an important role to play in the network as a whole.

ARCH v1.1—7-19

References

For additional information, refer to these resources:

- *Cisco IOS Quality of Service* at http://www.cisco.com/warp/public/732/Tech/qos/
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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 quality issues does QoS address? (Choose three.)

A) gain

B) jitter

C) delay

D) stutter

E) bandwidth

F) packet loss

Q2) What is the difference between IntServ and DiffServ?

A) IntServ is flow-based while DiffServ is type-based.

B) IntServ is type-based while DiffServ is flow-based.

C) IntServ is for voice only and DiffServ is for data only.

D) IntServ is for data only and DiffServ is for voice only.

Q3) Which is a service level agreement?

A) first-effort service

B) last-effort service

C) best-effort service

D) differential service

Q4) What is the purpose of trust boundaries?

A) to determine where in the network packets are dropped

B) to determine where in the network we reclassify packets

C) to determine where in the network packet marking takes place

D) to determine where in the network we trust the packet contents

Q5) When using congestion avoidance mechanisms, when is all traffic discarded?

A) at any buffer overflow

B) when a minimum threshold is reached

C) when a maximum threshold is reached

D) at regular intervals during buffer overflows

Q6) Which three features are congestion management tools? (Choose three.)

A) LFI

B) LLQ

C) WFQ

D) CLLLQ

E) CBWFQ

F) CCBWFQ

Q7) What function does traffic shaping perform?

A) delays excess traffic based on inbound bandwidth restrictions

B) delays excess traffic based on outbound bandwidth restrictions

C) discards excess traffic immediately based on outbound traffic restrictions

D) discards excess traffic immediately based on destination bandwidth restrictions

Q8) What is a primary purpose of signaling for QoS?

A) useful for managing network traffic congestion

B) allows output buffers to fill during periods of congestion

C) useful for coordinating the traffic handling techniques provided by other QoS features

D) provides priority either by raising the priority of one flow or by limiting the priority of another

Q9) Which QoS mechanism reduces delay and jitter for smaller packets on slower-speed multilink PPP links by breaking up large datagrams and interleaving low-delay traffic packets with the resulting smaller packets?

A) LFI

B) cRTP

C) dcRTP

D) FRF.12

Q10) Which QoS model requires end-to-end configuration?

A) LLQ

B) IntServ

C) DiffServ

D) traffic conditioning

Quiz Answer Key

Q1) B, C, F

Relates to: Enterprise Network Requirements for QoS

Q2) A

Relates to: IntServ and DiffServ QoS Architectures

Q3) C

Relates to: IntServ and DiffServ QoS Architectures

Q4) C

Relates to: Classification and Marking

Q5) D

Relates to: Congestion Avoidance

Q6) B, C, E

Relates to: Congestion Management

Q7) B

Relates to: Traffic Conditioning

Q8) C

Relates to: Signaling

Q9) A

Relates to: Link Efficiency Mechanisms

Q10) B

Relates to: Summary of Key Cisco IOS Software QoS Categories and Features

Designing QoS for Enterprise Networks

Overview

Different applications have different requirements for bandwidth, delay, jitter, and loss. For example, voice applications have stringent delay requirements and can tolerate limited packet loss. Alternatively, a commercial transaction may be less sensitive to delay but very sensitive to packet drops. Cisco QoS features provide the ability to manage traffic intelligently across the infrastructure.

QoS includes the mechanisms that give network managers the ability to control the mix of bandwidth, delay, variances in delay (jitter), and packet loss in the network in order to deliver a network service such as voice over IP; define different service level agreements (SLAs) for divisions, applications, or organizations; or simply prioritize traffic across a WAN link.

Relevance

The end-to-end performance of a network is critical for the enterprise. The network designers must examine the entire network to determine the impact of any proposed service on the network, and the applications that share the network.

Objectives

Upon completing this lesson, you will be able to design scalable network QoS solutions, given specific network and application needs. This includes being able to meet these objectives:

- List design guidelines to consider when developing a QoS solution for enterprise networks
- Design scalable network QoS solutions for the enterprise network, given specific network and application needs
- Design QoS solutions, given specific network and application needs

Learner Skills and Knowledge

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

- CCNP curriculum courses
- "Identifying Cisco QoS Mechanisms" lesson

Outline

The outline lists the topics included in this lesson.

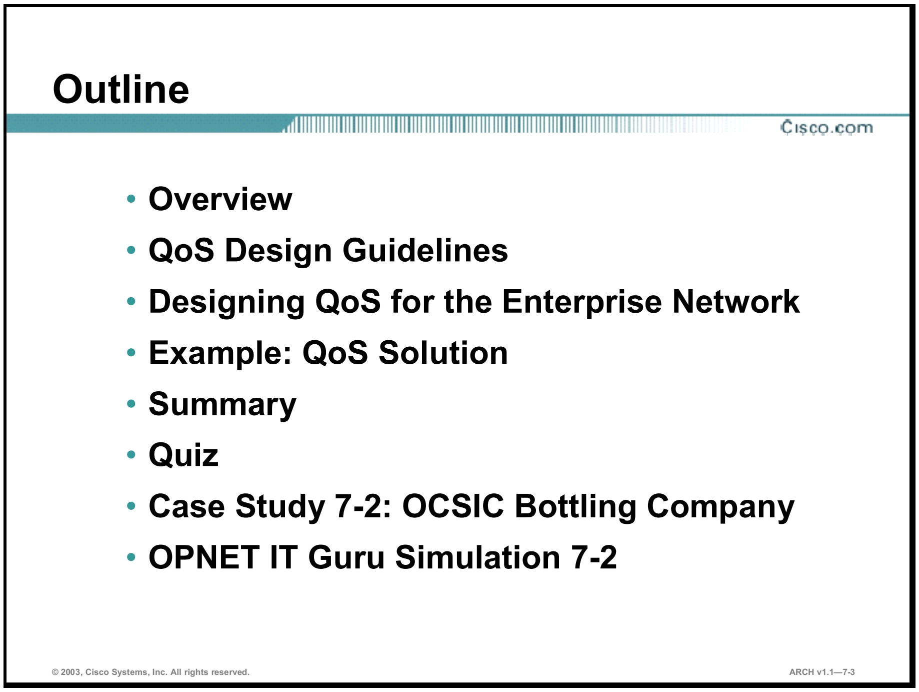

Outline

Cisco.com

- Overview
- QoS Design Guidelines
- Designing QoS for the Enterprise Network
- Example: QoS Solution
- Summary
- Quiz
- Case Study 7-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 7-2

ARCH v1.1—7-3

QoS Design Guidelines

You will configure QoS features throughout a network to provide for end-to-end QoS delivery. This topic lists design guidelines to consider when developing a QoS solution for enterprise networks.

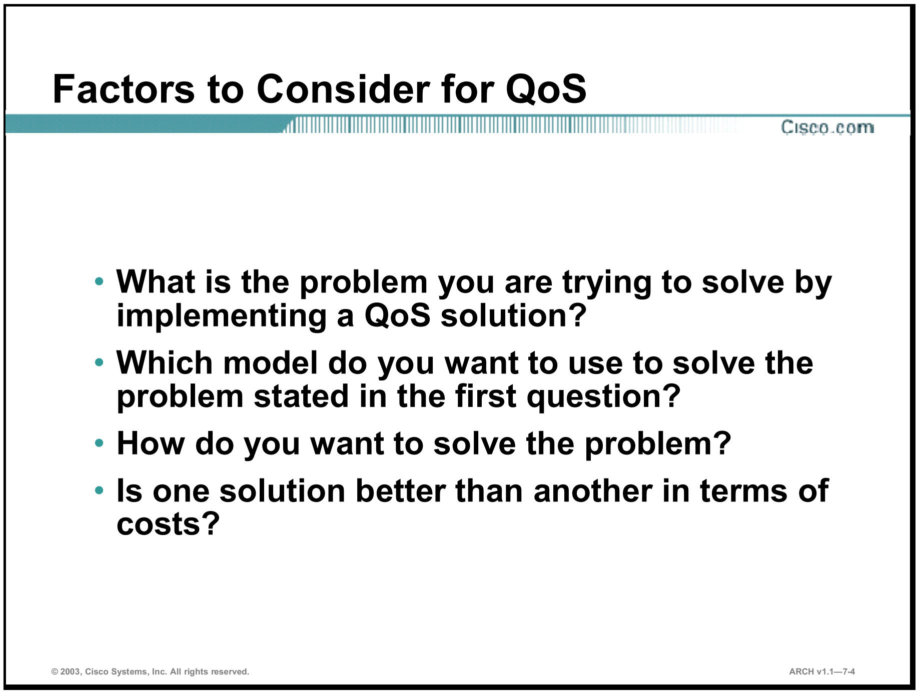

Factors to Consider for QoS

Cisco.com

- What is the problem you are trying to solve by implementing a QoS solution?
- Which model do you want to use to solve the problem stated in the first question?
- How do you want to solve the problem?
- Is one solution better than another in terms of costs?

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—7-4

The figure lists the important design questions to ask when implementing a QoS solution.

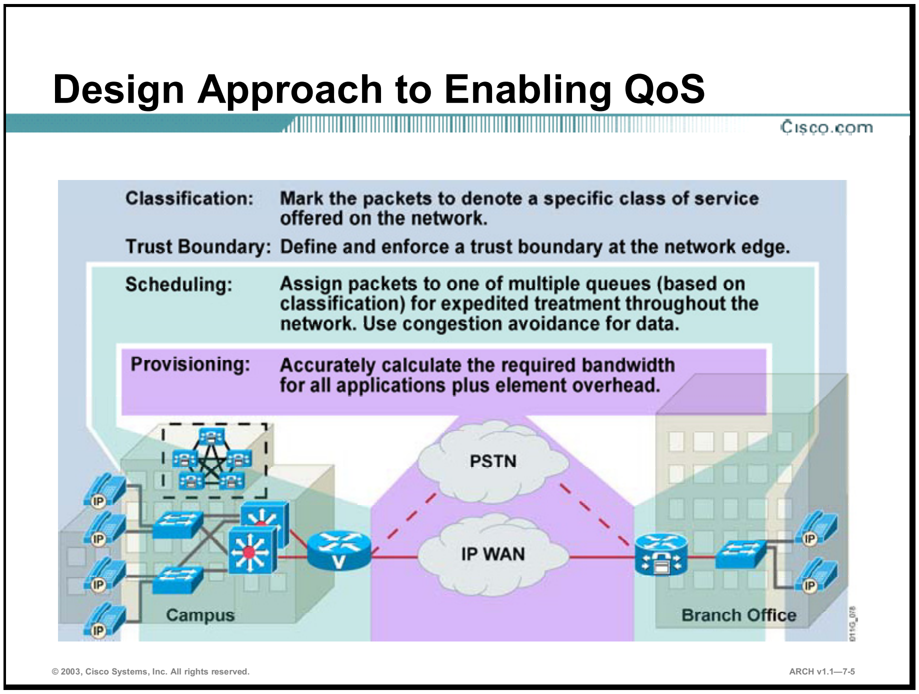

Design Approach to Enabling QoS

Cisco.com

Classification: Mark the packets to denote a specific class of service offered on the network.

Trust Boundary: Define and enforce a trust boundary at the network edge.

Scheduling: Assign packets to one of multiple queues (based on classification) for expedited treatment throughout the network. Use congestion avoidance for data.

Provisioning: Accurately calculate the required bandwidth for all applications plus element overhead.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-5

The QoS tools are a set of mechanisms to increase voice quality on data networks by decreasing dropped voice packets during times of network congestion and by minimizing both the fixed and variable delays encountered in a given voice connection.

Classification tools mark a packet or flow. Only classifications made, or confirmed, at a trust boundary are honored through the network.

The table lists packet priority classifications and the corresponding IP precedence and DSCP values.

Packet Priority Classifications Layer 2 Class of Service	IP Precedence	Differentiated Services Code Point
CoS 0	Routine (IP precedence 0)	0–7
CoS 1	Priority (IP precedence 1)	8–15
CoS 2	Immediate (IP precedence 2)	16–23
CoS 3	Flash (IP precedence 3)	24–31
CoS 4	Flash-override (IP precedence 4)	32–39
CoS 5	Critical (IP precedence 5)	40–47
CoS 6	Internet (IP precedence 6)	48–55
CoS 7	Network (IP precedence 7)	56–63

Scheduling tools refer to the set of tools that determine how a frame or packet exits a node. When packets enter a device faster than they can exit, a point of congestion will occur at some output interface. Multiple inputs being directed to a single output or by a higher-speed input bursting to a lower-speed output may cause congestion. Devices queue packet buffers to allow for scheduling higher-priority packets to exit sooner than lower priority ones. Most prioritized queuing mechanisms use multiple logical queues to separate the traffic. The queues are then serviced on a strict priority basis or in some kind of a round-robin scheme that achieves a fairer service for all traffic.

Queues have a finite capacity and act very much like a funnel for water. If water continually enters the funnel faster than it exits, eventually the funnel will begin overflowing from the top, and water is lost. When queues begin overflowing, packets drop as they arrive (tail-drop). Traffic congestion management algorithms begin to drop packets arriving from selected low-priority traffic flows before the queue fills. There are three mechanisms involved in queue management:

- Placement of packets into queues on arrival based on prior marking or other values in the packet header
- Servicing of queues as output capacity becomes available
- Determination of which arriving packets are sacrificed as a queue fills to preserve space for other traffic

When calculating the required amount of bandwidth for a converged WAN, remember that all the application traffic (voice, video, and data traffic), when added together, should equal no more than 75 percent of the provisioned bandwidth. The remaining 25 percent is reserved for short-demand peaks or outages, and administrative overhead, such as routing and signaling protocols.

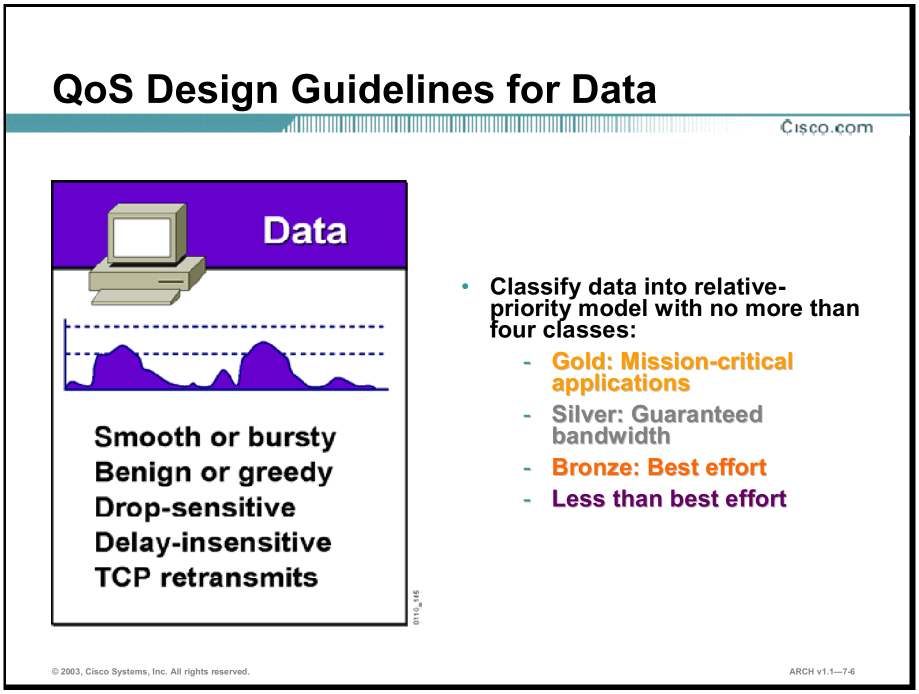

QoS Design Guidelines for Data

Cisco.com

Data

Smooth or bursty
Benign or greedy
Drop-sensitive
Delay-insensitive
TCP retransmits

- Classify data into relative-priority model with no more than four classes:
  - Gold: Mission-critical applications
  - Silver: Guaranteed bandwidth
  - Bronze: Best effort
  - Less than best effort

0110_145

ARCH v1.1—7-6

Different data applications have different traffic characteristics. Even different versions of the same application can have different traffic characteristics. You can classify data into a relative-priority model in one of four classes: Gold, silver, bronze, and less than best effort.

Traffic analysis and lab testing are required to determine bandwidth requirements for data applications.

Cisco recommends that you implement a relative-priority model of these four classes, which has been proven to work well in most enterprise environments.

For highly mission-critical data, Cisco recommends that you implement the following: DSCP AF21-23, IP precedence 2, and CoS 2.

For second-tier mission-critical data, Cisco recommends that you implement DSCP AF11-AF13, IP precedence 1, and CoS 1.

Application updates, fluctuation in the numbers of users, varying business environments, and the time of day, month, and year, all affect bandwidth requirements of data applications. Therefore, rather than attempting to determine exact kbps of bandwidth requirements for data applications, a simpler and proven approach is to assign relative priorities to data applications.

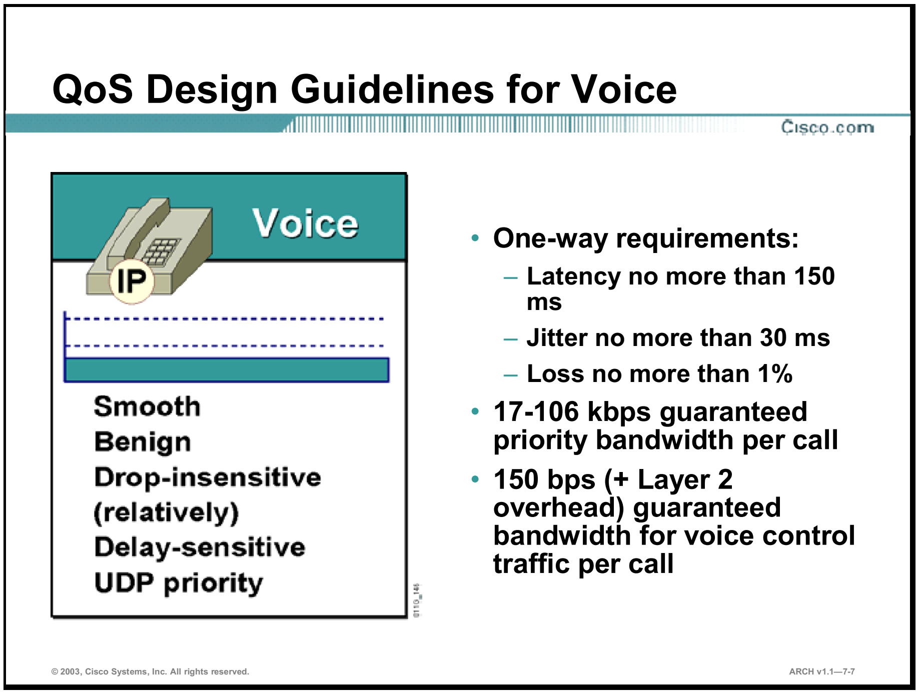

QoS Design Guidelines for Voice

Cisco.com

Voice

IP

Smooth
Benign
Drop-insensitive (relatively)
Delay-sensitive
UDP priority

- One-way requirements:
  - Latency no more than 150 ms
  - Jitter no more than 30 ms
  - Loss no more than 1%
- 17-106 kbps guaranteed priority bandwidth per call
- 150 bps (+ Layer 2 overhead) guaranteed bandwidth for voice control traffic per call

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-7

Bandwidth per call for voice depends on the codec, duration of the sample or number of predictors, and data link layer media. To calculate the bandwidth that voice streams consume, add the packet payload and all headers (in bits), and then multiply by the packet rate per second.

In centralized call processing designs, the IP Phones use a TCP control connection to communicate with a call processing server. If there is not enough bandwidth provisioned for these lightweight control connections, the end user might be adversely affected.

When addressing the QoS needs of voice traffic, keep in mind:

- One-way latency should be no more than 150 to 200 ms.
- Jitter should be no more than 30 ms.
- Loss should be no more than 1 percent.
- 17 to 106 kbps of guaranteed priority bandwidth is required per call (depending on the sampling rate, codec, and Layer 2 overhead).
- 150 bps (+ Layer 2 overhead) per phone of guaranteed bandwidth is required for voice control traffic.

In an enterprise Architecture for Voice, Video and Integrated Data (AVVID) network, you should classify packets that contain voice and video traffic into the appropriate queues as follows:

- **Voice:** The IETF draft recommends a DSCP PHB label of EF for VoIP traffic. To remain backwardly compatible with IP precedence, use an IP precedence value of five and a CoS marking of five. These markings can be used as selection criteria for entry into the priority queue, where it exists, or the queue with the highest service weight and lowest drop probability in a weighted round-robin (WRR)/WRED scheduling scheme.

- **Video:** Use a value that is different from that used for VoIP traffic. In places where policing is needed, you can protect voice from video or vice versa. To accomplish this separation, you should use a DSCP PHB label of AF41, an IP precedence value of 4, and a CoS marking of 4.

If you marked both voice and video as EF, IP precedence 5, and CoS 5, it would be more difficult to differentiate between the two types of traffic if you wanted to rate-limit (police) or otherwise control the amount of voice or video traffic through the network.

Typically, a voice gateway or voice application server will mark its RTP and control traffic with the appropriate DSCP and CoS markings. However, some end devices may not have the capability to correctly classify their own traffic. To provide control and security, do not trust the CoS and markings that an end device assigns.

Assign signaling traffic a DSCP PHB label of AF31, an IP precedence value of 3, and a CoS marking of 3.

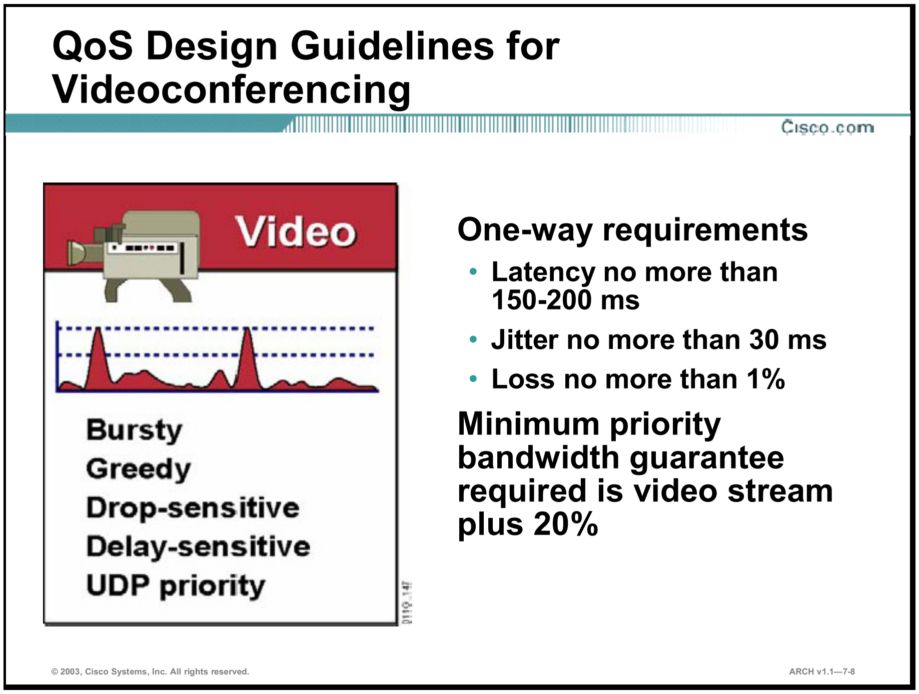

QoS Design Guidelines for Videoconferencing

Cisco.com

Video

0110-147

Bursty
Greedy
Drop-sensitive
Delay-sensitive
UDP priority

One-way requirements

- Latency no more than 150-200 ms
- Jitter no more than 30 ms
- Loss no more than 1%

Minimum priority bandwidth guarantee required is video stream plus 20%

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—7-8

There are two main types of video applications: streaming video (such as IP/TV, which may be either on-demand or multicast) and interactive video (such as videoconferencing). When addressing the QoS needs of videoconferencing traffic, keep these considerations in mind:

- One-way latency should be no more than 150 to 200 ms.
- Jitter should be no more than 30 ms.
- Loss should be no more than 1 percent.
- The minimum bandwidth guarantee is the size of the videoconferencing session plus 20 percent, meaning that a 384-kbps videoconferencing session requires 460 kbps of guaranteed priority bandwidth.

When addressing the QoS needs of streaming video traffic, keep these considerations in mind:

- Latency should be no more than four to five seconds, depending on the video application's buffering capabilities.
- There are no significant jitter requirements.
- Loss should be no more than 2 percent.
- Bandwidth requirements depend on the encoding and rate of video stream.
- Nonentertainment streaming video should be provisioned into the silver (guaranteed bandwidth) data-traffic class.

For video content distribution, keep the following considerations in mind:

- Streaming video content distribution is delay and delay-variation insensitive.
- Streaming video requires large file transfers (traffic patterns similar to FTP sessions).
- Try to restrict distribution to less-busy times of day.
- Provision video as less-than-best-effort data.

In enterprise networks, videoconferencing over IP has similar loss, delay, and delay-variation requirements to that of VoIP traffic. You must classify IP videoconferencing traffic so that network devices can recognize it and provide the appropriate treatment during periods of congestion. In enterprise networks, videoconferencing packets should be marked with a DSCP PHB label of AF41. For backward compatibility, an IP precedence of 4 should be used. Additionally, a Layer 2 802.1p CoS value of 4 should be used for IPVC traffic in 802.1Q environments.

Streaming-video applications, such as IPTV video on demand (VoD) programs, are relatively high-bandwidth applications with a high tolerance for loss, delay, and delay variation. As such, significant QoS tools are not required to meet the needs of these applications. However, in most enterprise environments, these types of applications are considered more important than regular background applications (such as e-mail and web browsing) and should be given preferential treatment. A Layer 2 classification of CoS 1 in 802.1Q/802.1p environments should be used for these applications. To remain backwardly compatible, an IP precedence classification of 1 should be used. Because streaming video is not drop- or delay-sensitive, you can use the high-drop precedence DSCP PHB. The DSCP label AF13 should be used as it has the highest drop precedence in the AF2x or silver data class.

Designing QoS for the Enterprise Network

The QoS implementation for a campus network differs at the Campus Backbone, Building Access, and Building Distribution submodules. This topic helps you select scalable network QoS solutions for the enterprise network, given specific network and application needs.

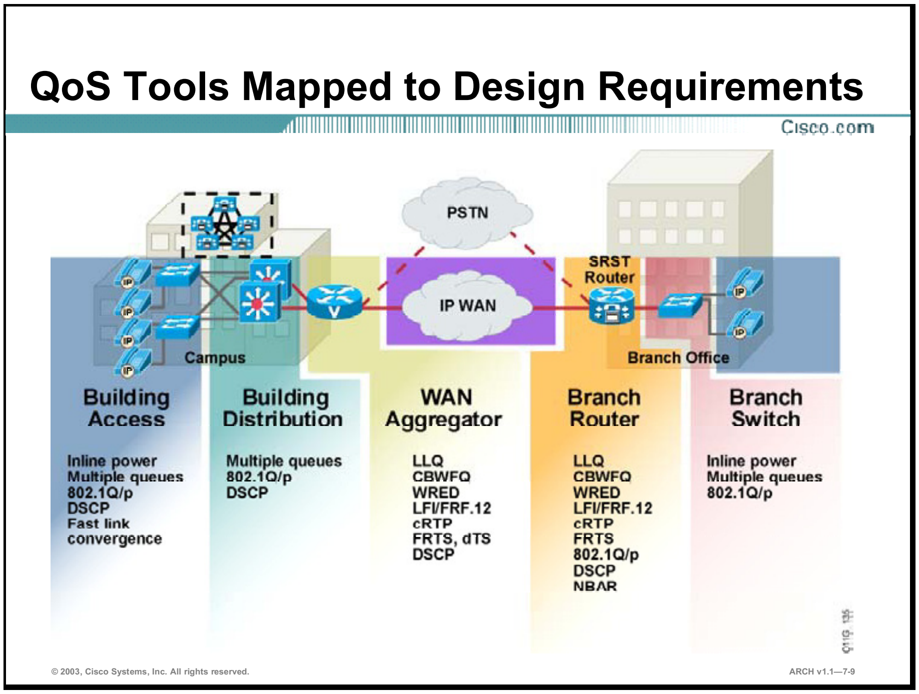

QoS Tools Mapped to Design Requirements

Cisco.com

Building Access
Inline power
Multiple queues
802.1Q/p
DSCP
Fast link convergence

Building Distribution
Multiple queues
802.1Q/p
DSCP

WAN Aggregator
LLQ
CBWFQ
WRED
LFI/FRF.12
cRTP
FRTS, dTS
DSCP

Branch Router
LLQ
CBWFQ
WRED
LFI/FRF.12
cRTP
FRTS
802.1Q/p
DSCP
NBAR

Branch Switch
Inline power
Multiple queues
802.1Q/p

ARCH v1.1—7-9

Not all QoS tools are appropriate in all areas of the network. The diagram points out what types of QoS to implement in each module and submodule of the Enterprise Composite Model. As pictured in the figure, you would select different QoS solutions based on where QoS is being implemented in the network.

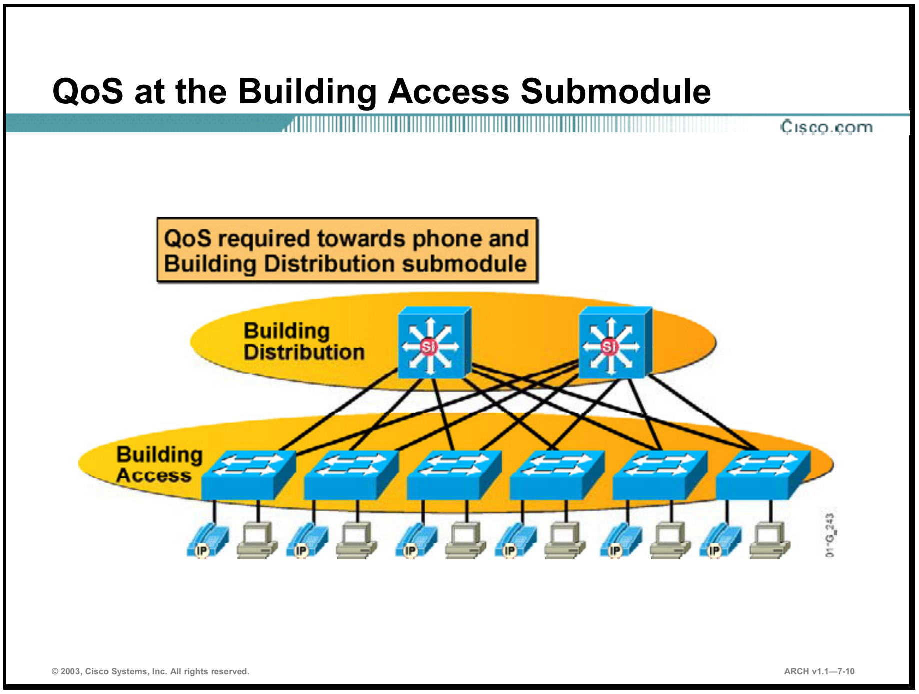

QoS at the Building Access Submodule

Cisco.com

QoS required towards phone and Building Distribution submodule

ARCH v1.1—7-10

The Building Access submodule is typically where the trust boundary is formed. That is where the precedence is set for the packets and then trusted throughout the rest of the network. A switch at the Building Access submodule must support these capabilities:

- Multiple VLANs on the access port to which an end user is attached
- Manipulation of the QoS or CoS values provided by an end device
- Extension of the trust boundary for the CoS or DSCP marking toward the end devices

There are times when the devices attached to the campus network do not classify their traffic with the appropriate Layer 2 and Layer 3 markings. When considering your choices for access layer devices, consider the switch's ability to classify and mark traffic at the edge of the network using ACLs and service policies. This allows you to offer QoS as a service throughout the network, and administer it at the edge of the network where CPU resources are plentiful, rather than at the Campus Backbone and Building Distribution submodules where QoS classification and marking could adversely affect network responsiveness and utilization.

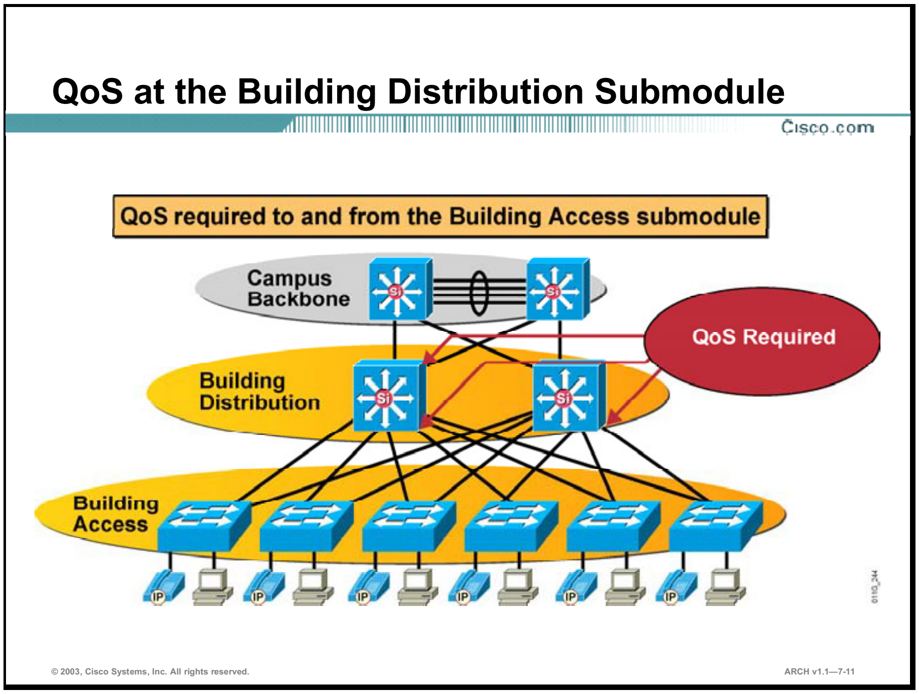

QoS at the Building Distribution Submodule

QoS required to and from the Building Access submodule

ARCH v1.1—7-11

QoS at the Building Distribution layer requires these changes to the implementation of the switches:

- Enable QoS.
- Change the default CoS to the DSCP table so that CoS and DSCP behavior can be maintained throughout the network.
- Configure service policies to classify traffic that does not contain a CoS to DSCP marking that you can trust. Finally, enable CoS or DSCP trust on the ports where trust is appropriate, DSCP for Layer 3 aware access, and CoS for Layer 2 only access.

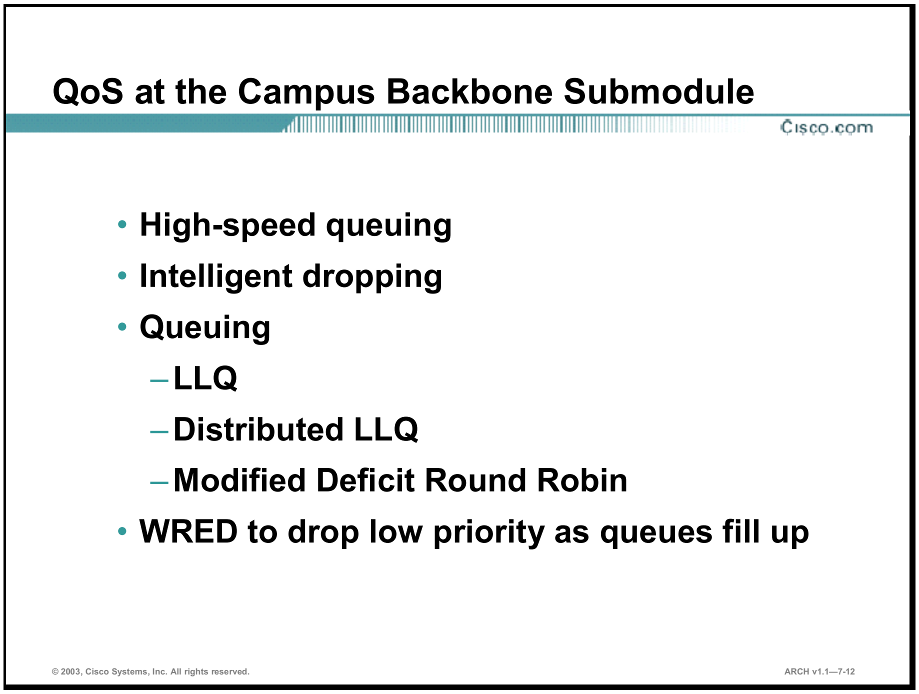

QoS at the Campus Backbone Submodule

Cisco.com

- High-speed queuing
- Intelligent dropping
- Queuing
  - LLQ
  - Distributed LLQ
  - Modified Deficit Round Robin
- WRED to drop low priority as queues fill up

ARCH v1.1—7-12

To maintain the QoS policy throughout the network, the core device can provide some QoS congestion management features. The traffic going to the Campus Backbone should already be classified and marked at the Building Access or Building Distribution submodules, so the Campus Backbone should be able to process the traffic quickly using a simple queuing mechanism. There should be no need to run QoS classification and marking tools within the Campus Backbone of the network.

If QoS is implemented in the Campus Backbone, keep it to a minimum to facilitate high-speed queuing with some form of intelligent dropping. The typical queue service mechanisms are LLQ and modified deficit round-robin scheduling.

Example: QoS Solution

Small, medium, and large enterprise QoS solutions encompass classification and marking, congestion avoidance, congestion management, and link-efficiency mechanisms. This topic provides a QoS design solution to help you design your own solutions.

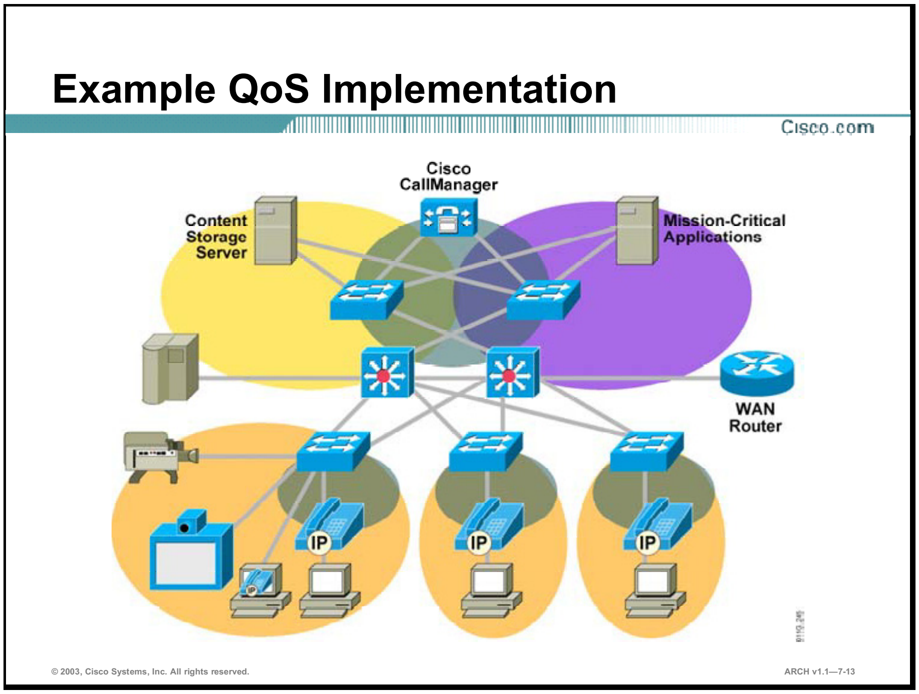

Example QoS Implementation

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—7-13

The Celestal Curtain Manufacturing Company headquarters has recently redesigned their network and has come up with plans to implement a QoS solution for their network, particularly on WAN links. Their new design includes voice and access to critical applications.

Voice applications have the top priority and the other critical applications take second priority. All other traffic can be processed on a best-effort service.

The company has a choice of three QoS models:

- **Best-effort model:** Works for the best-effort type of traffic, but not for the other two types of traffic.
- **IntServ model:** Works for all types of traffic but is harder to configure and maintain, as RSVP must be enabled throughout the network and traffic flow is actively managed.
- **DiffServ model:** Allows fairly easy configuration to maintain traffic priority by packet type.

The company selected the DiffServ model for its ease of implementation and traffic type-based prioritization.

The company decided to build the trust boundary at the access layer, eliminating the need for QoS classification and marking in the Campus Backbone or Building Distribution submodule.

The Catalyst 3524 switch is the base access device, and supports these QoS enablers:

- Two transmit queues
- CoS-based QoS (802.1p)
- Trust CoS
- Port-based CoS marking for unmarked traffic
- GigaStack: Only point-to-point stacks are recommended
- Inline power

On the access devices, the company enabled multiple VLAN identifiers per port. They use the IP Phones to classify their traffic. The Catalyst 3524 reclassifies the traffic from behind the end-user device to the low queue.

The access layer includes one Catalyst 4000 that supports these QoS enablers:

- Two transmit queues
- CoS (802.1p)-based QoS
- CoS values mapped to output queues in pairs
- Queues are serviced in a round-robin fashion
- Switch-wide CoS marking for unmarked traffic
- All ports are considered trusted
- Inline power via privacy enhanced mail (PEM) and power shelf

Once QoS is enabled, all CoS values are mapped to queue number 1 by default. CoS queue values must be mapped. The company implemented two Catalyst 6500 switches in the Campus Backbone and Building Distribution submodules that support these QoS enablers:

- Redundant supervisors
- Transmit and receive queues
- Priority queues and multiple-drop thresholds
- CoS, DSCP, or ACL-based QoS (policy maps)
- Trust DSCP or CoS
- Settings by port DSCP or CoS (marked or unmarked)
- Mapping from CoS to DSCP or DSCP to CoS
- Port can trust DSCP, IP Precedent, or CoS
- Recommended: trust-CoS, with access to receive priority queuing (Rx PQ)
- 10/100 network interface cards require an additional step of configuring ACL to trust traffic

Output scheduling consists of:

- Assigning traffic to queues based on CoS
- Configuring threshold levels
- Modifying buffer sizes (expert mode)
- Assigning weights for WRR (expert mode)

The Building Distribution submodule is the aggregation point in this network. Since the Catalyst 3524 switches cannot mark and classify multiple data types, classification and marking takes place in the distribution layer.

On the Building Distribution devices, the administrator creates ACLs to classify the mission-critical data into a higher priority. The company implemented a gold, silver, and bronze best-effort model. Voice is classified as gold, signalling is classified as silver, and critical data is classified as bronze. All other traffic is in the best-effort queue.

Summary

This topic summarizes the key points discussed in this lesson.

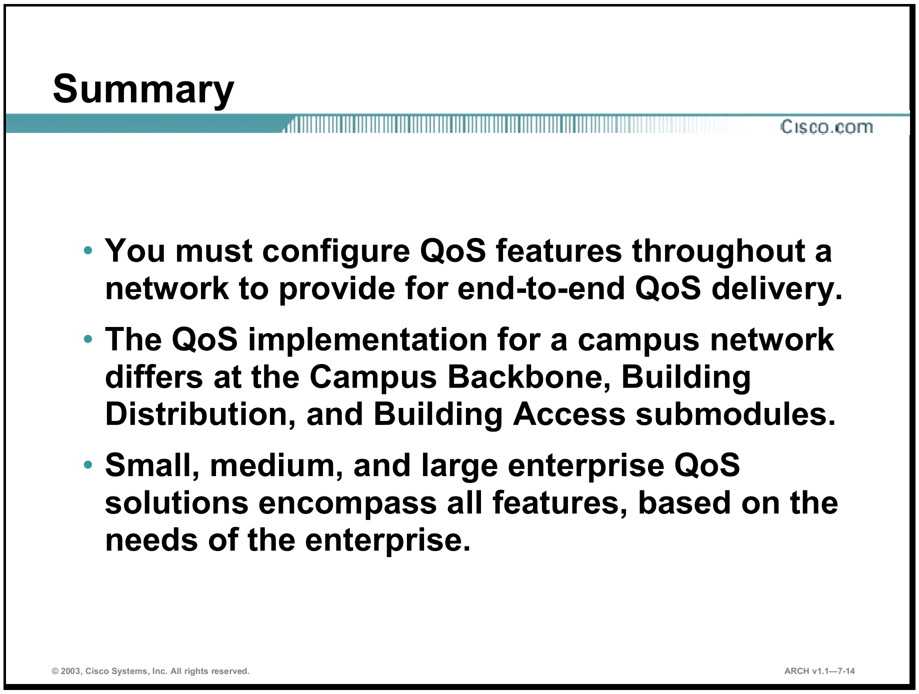

Summary

- You must configure QoS features throughout a network to provide for end-to-end QoS delivery.
- The QoS implementation for a campus network differs at the Campus Backbone, Building Distribution, and Building Access submodules.
- Small, medium, and large enterprise QoS solutions encompass all features, based on the needs of the enterprise.

ARCH v1.1—7-14

References

For additional information, refer to these resources:

- *Cisco IOS QoS Articles and Reports* at http://www.cisco.com/warp/public/732/Tech/qos/qos_articles.shtml
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 7-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 7-2

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) What are the three primary areas to consider in a QoS design?

A) classification, scheduling, provisioning

B) tool selection, provisioning, management

C) classification, management, QoS tool design

D) scheduling, tool implementation, equipment selection

Q2) Which QoS feature should access devices support?

A) inline power

B) support for voice signaling

C) support for multiple VLANs per port

D) support for BECN/FECN acknowledgements

Q3) What can happen to the network if QoS is not applied to the Building Distribution submodule?

A) While aggregating traffic, lower-priority traffic might be sent before high-priority traffic.

B) While aggregating traffic, the TX buffer will use a last-in, first-out scheduling mechanism.

C) Nothing, traffic will flow through the switch as it was prioritized at the all access layers.

D) The Building Distribution submodule will use the QoS configuration at the access switches to ensure QoS compliance.

Q4) What happens when traffic from untrusted sources is allowed in the Campus Backbone?

A) Nothing happens.

B) The non-trusted traffic is deleted.

C) It could degrade the anticipated QoS for other traffic.

D) The traffic remains in the queue until the device finishes processing all other traffic.

Q5) In which area of the Enterprise Campus should you set the trust boundary?

A) Building Access submodule

B) Campus Backbone submodule

C) Building Distribution submodule

D) between backbone devices

Q6) Why is QoS applied in the distribution switches on a campus network?

A) to modify service levels

B) to remove service levels

C) to maintain service levels

D) to create new service levels

Q7) When should you implement QoS in the Campus Backbone?

A) only for trusted traffic

B) only when the core is congested

C) when traffic is inserted into the network that has never been classified and marked by a trusted device

D) when the Campus Backbone devices have enough bandwidth to allow for classification and scheduling

Quiz Answer Key

Q1) A

Relates to: QoS Design Guidelines

Q2) C

Relates to: Designing QoS for the Enterprise Network

Q3) A

Relates to: Designing QoS for the Enterprise Network

Q4) C

Relates to: Designing QoS for the Enterprise Network

Q5) A

Relates to: Designing QoS for the Enterprise Network

Q6) C

Relates to: Designing QoS for the Enterprise Network

Q7) C

Relates to: Designing QoS for the Enterprise Network

Case Study 7-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

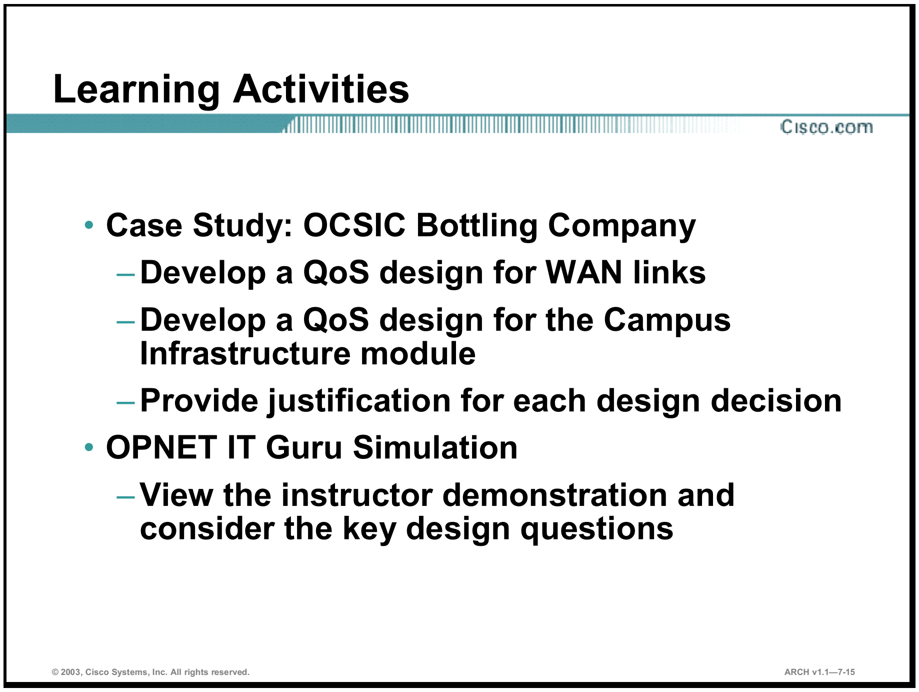

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Develop a QoS design for WAN links
  - Develop a QoS design for the Campus Infrastructure module
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—7-15

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

OCSIC Bottling Company wants to prioritize the traffic going over the WAN links to ensure that the SAP/Oracle traffic always gets first priority in the network. The PeopleSoft and e-mail applications get second priority. All traffic within the intranet receives third priority, and all other traffic not specified gets a low priority. OCSIC wants the WAN links to adjust their rate, based on Frame Relay BECN and FECN activity. There has been a decision to include a VoIP solution, which means the voice traffic will require the highest priority.

In this exercise, you will design QoS services that meet the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Develop a QoS design for WAN links
- Develop a QoS design for the campus network

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Develop a QoS Design for the Site-to-Site WAN

Complete these steps:

Step 1 Complete the table to design your QoS solution.

Design Questions	Decision	Justification
What classification and marking tools and settings are required?		
What congestion avoidance tools and settings are required?		
What congestion management tools and settings are required?		
What traffic conditioning tools and settings are required?		
What signaling tools and settings are required?		
What link efficiency tools and settings are required?		

Task 2: Develop a QoS Design for the Campus Network

Complete these steps:

Step 1 Complete the table to design your QoS solution.

Design Questions	Decision	Justification
What classification and marking tools and settings are required?		
What congestion avoidance tools and settings are required?		
What congestion management tools and settings are required?		
What traffic conditioning tools and settings are required?		
What signaling tools and settings are required?		
What link efficiency tools and settings are required?		

Task 3: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have developed a QoS design for the site-to-site WAN that identifies classification and marking, congestion avoidance, congestion management, traffic conditioning, signaling, and link-efficiency mechanisms.
- You have developed a QoS design for the headquarters' campus network that identifies classification and marking, congestion avoidance, congestion management, traffic conditioning, signaling, and link-efficiency mechanisms.

OPNET IT Guru Simulation 7-2

The QoS simulation demonstrates the affect of QoS on network and application performance.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 8

Designing IP Multicast Services

Overview

Multicasting provides bandwidth conservation that reduces traffic load by simultaneously delivering a single stream of information to multiple recipients. Multicasting enables distribution of video conferencing, corporate communications, distance learning, distribution of software, and other applications. Multicast packets are replicated in the network by routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, resulting in an efficient delivery of data to multiple receivers.

Module Objectives

Upon completing this module, you will be able to design IP multicast intelligent network services for performance, scalability, and availability, given specified enterprise network needs.

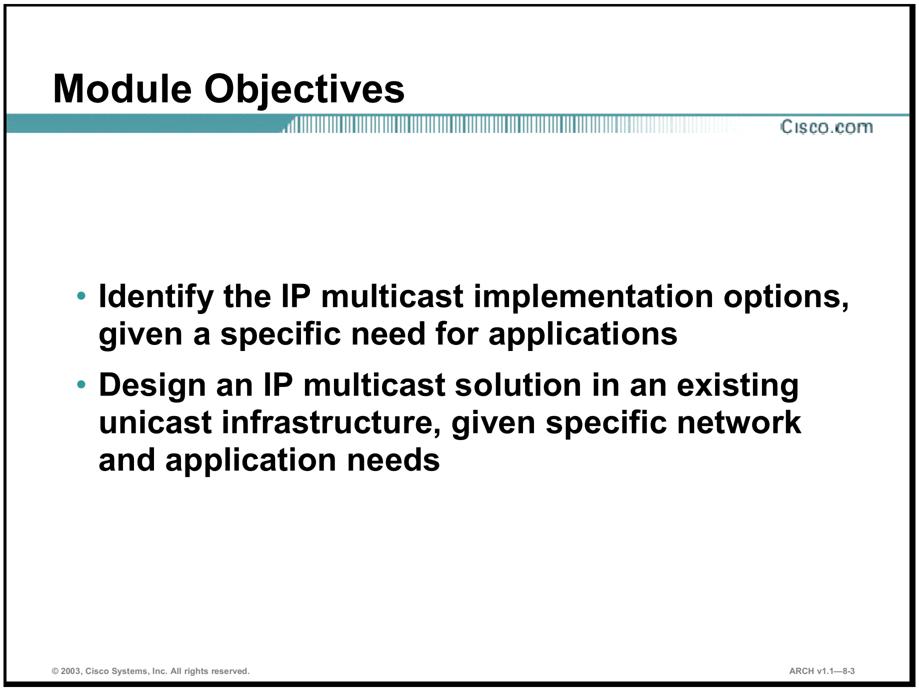

Module Objectives

- Identify the IP multicast implementation options, given a specific need for applications
- Design an IP multicast solution in an existing unicast infrastructure, given specific network and application needs

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—8-3

Module Outline

The outline lists the components of this module.

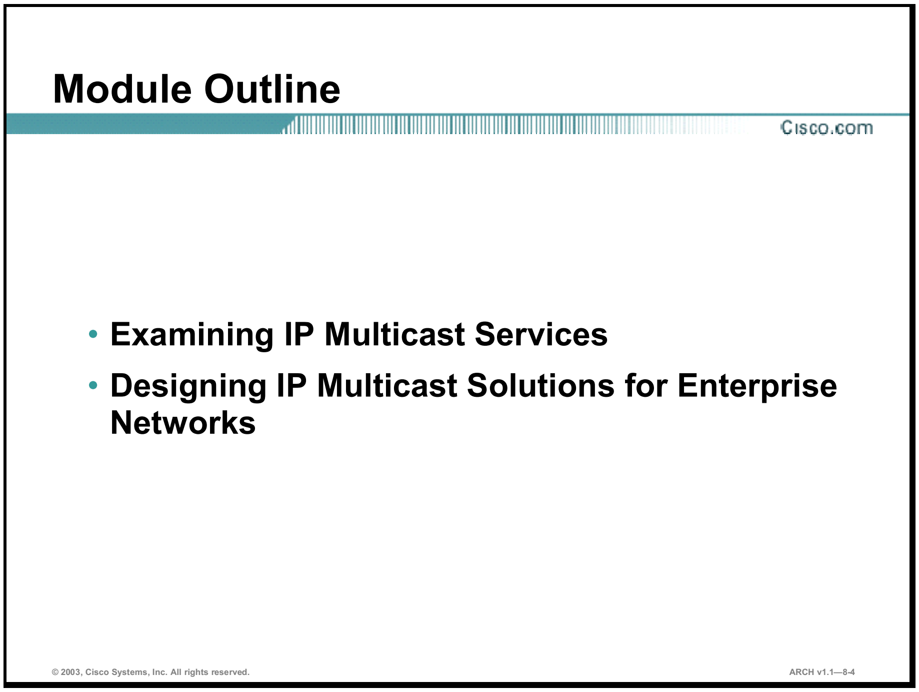

Module Outline

- Examining IP Multicast Services
- Designing IP Multicast Solutions for Enterprise Networks

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—8-4

Examining IP Multicast Services

Overview

Traditional IP communication allows a host to send packets to a single host (unicast transmission) or to all hosts (broadcast transmission). IP multicast provides a third possibility: allowing a host to send packets to a subset of all hosts as a group transmission. This lesson provides an overview of IP multicast, a review of intradomain multicast protocols, and an introduction to interdomain protocols.

Relevance

As one of the many capabilities of Cisco IOS software, the IP multicast technologies enable massively scalable, efficient distribution of data, voice, and video streams to hundreds, thousands, even millions of users. IOS multicast enables corporate communications, video conferencing, e-learning, Internet broadcast, hoot & holler, and streaming media applications.

Objectives

Upon completing this lesson, you will be able to identify the IP multicast implementation options, given a specific need for applications. This includes being able to meet these objectives:

- Identify enterprise requirements for IP multicast intelligent network services
- Explain the benefits of using IP multicast as an underlying transport mechanism in an enterprise network
- Explain how multicast forwarding works with Reverse Path Forwarding (RPF)
- Explain the purpose and use of IP multicast receiver group membership
- Describe the operation of PIM for IP multicast
- Explain the purpose and function of IP multicast control mechanisms

Learner Skills and Knowledge

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

- CCNP® certification courses

Outline

The outline lists the topics included in this lesson.

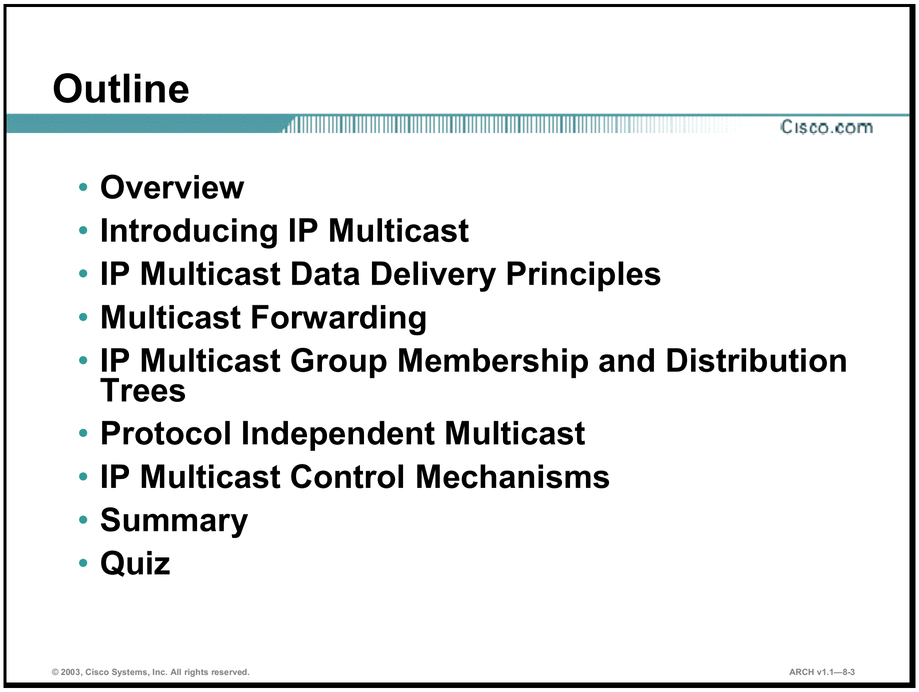

Outline

Cisco.com

- Overview
- Introducing IP Multicast
- IP Multicast Data Delivery Principles
- Multicast Forwarding
- IP Multicast Group Membership and Distribution Trees
- Protocol Independent Multicast
- IP Multicast Control Mechanisms
- Summary
- Quiz

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-3

Introducing IP Multicast

IP multicast, as an alternative to unicast and broadcast, sends packets to a subset of network hosts simultaneously. By requiring only a single copy of each packet to be sent on each interface, multicast helps reduce network traffic. This topic identifies the enterprise requirements for IP multicast intelligent network services.

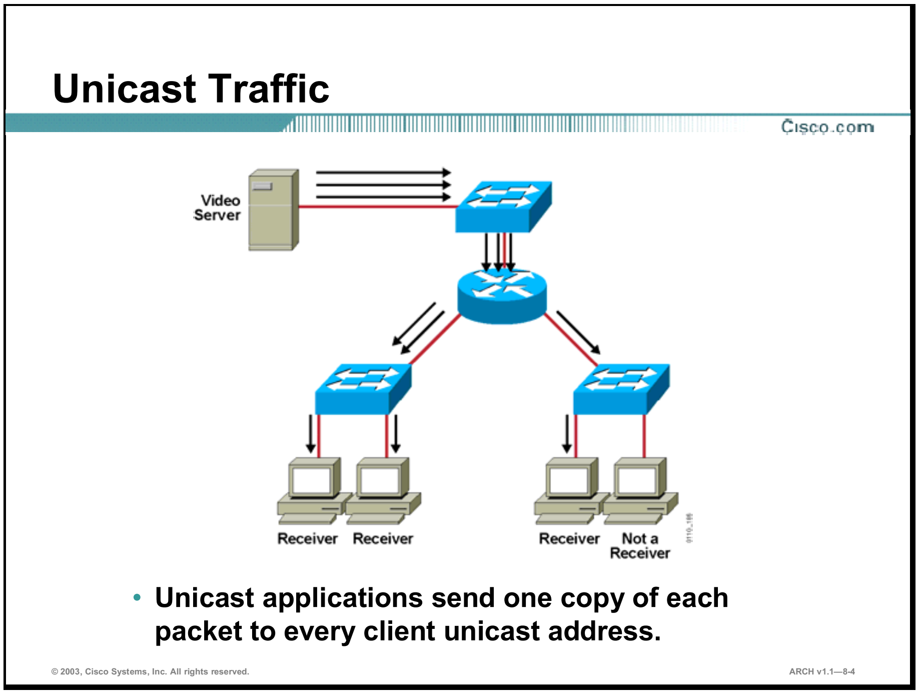

Unicast Traffic

- **Unicast applications send one copy of each packet to every client unicast address.**

ARCH v1.1—8-4

With a unicast design, an application sends one copy of each packet to every client unicast address. Unicast transmission requires a large amount of bandwidth, as the same information has to be carried multiple times, even on shared links. A large number of clients can impact the scalability of the network.

Two areas of unicast traffic that are of concern to network managers are:

- Number of user connections
- Amount of replicated unicast transmissions

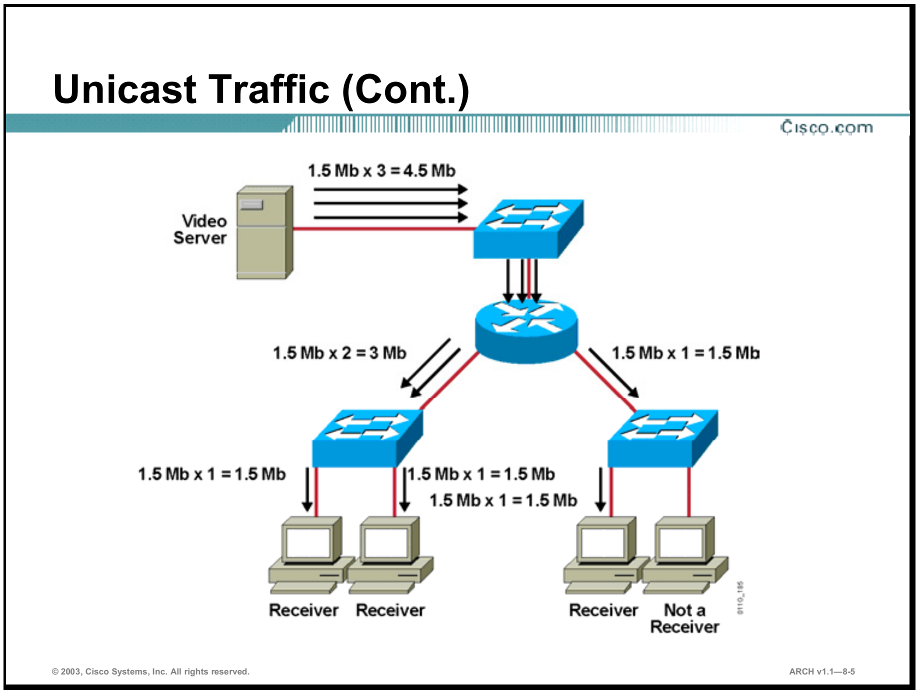

Unicast Traffic (Cont.)

ARCH v1.1—8-5

A server in a unicast environment must send a separate video stream for each client requesting access to the application.

The figure illustrates how the number of clients in a unicast environment can quickly consume bandwidth.

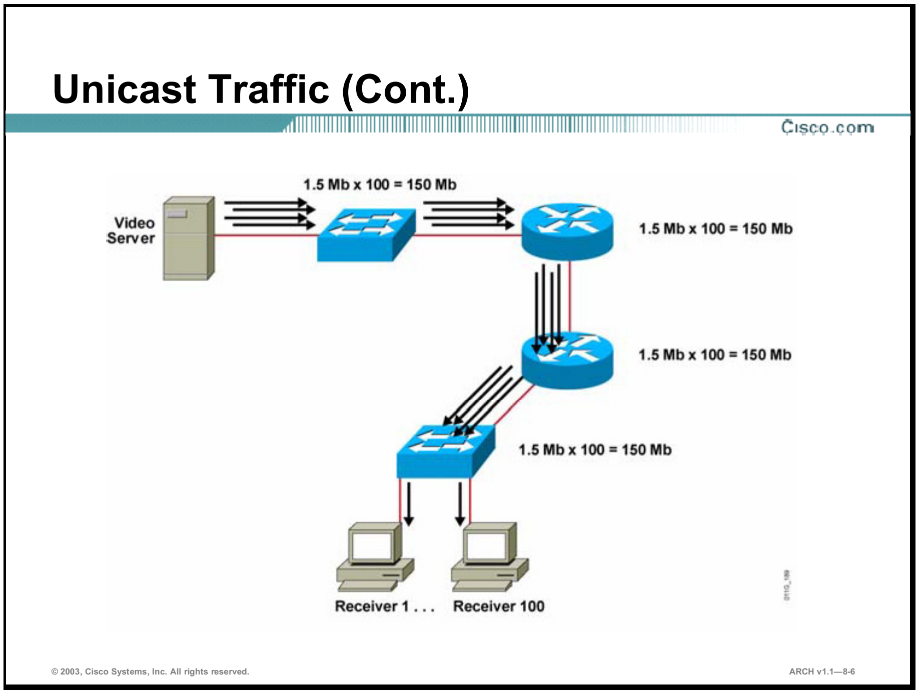

Unicast Traffic (Cont.)

ARCH v1.1—8-6

Similarly, the number of clients in a unicast environment can quickly consume network bandwidth, as shown in the figure.

Replicated unicast transmissions consume bandwidth within the network. In addition, you must consider the number of router and switch hops that occur in the path between server and client. Clients in a unicast environment can overload these intermediate devices by causing them to replicate the required number of packets.

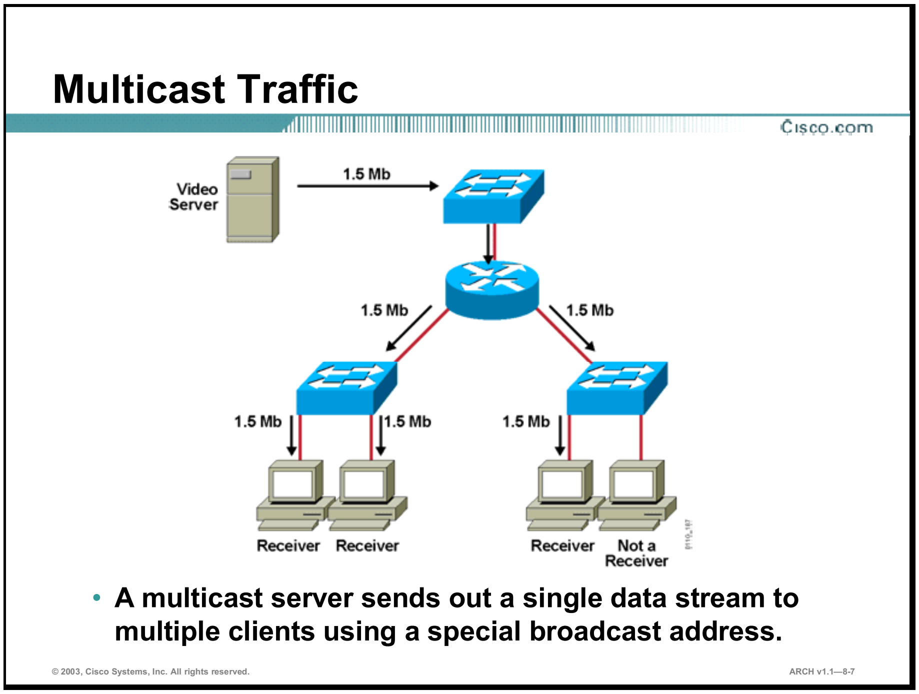

Multicast Traffic

- A multicast server sends out a single data stream to multiple clients using a special broadcast address.

ARCH v1.1—8-7

A multicast server sends out a single data stream to multiple clients using a special broadcast address. Client devices decide whether or not to listen to the multicast address.

The figure demonstrates how multicasting saves bandwidth and controls network traffic by forcing the network to replicate packets only when necessary.

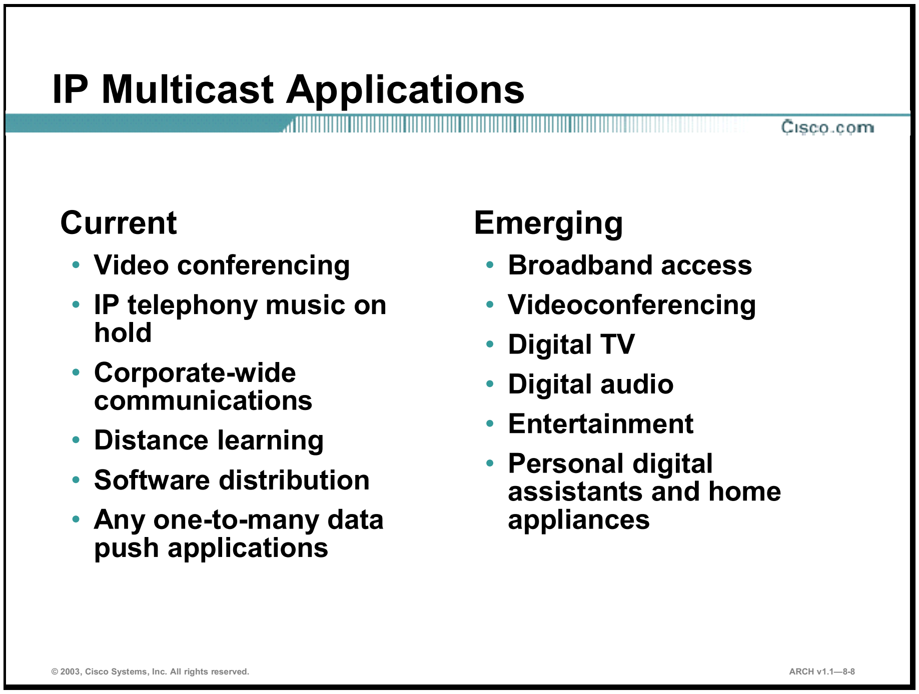

IP Multicast Applications

Current

- Video conferencing
- IP telephony music on hold
- Corporate-wide communications
- Distance learning
- Software distribution
- Any one-to-many data push applications

Emerging

- Broadband access
- Videoconferencing
- Digital TV
- Digital audio
- Entertainment
- Personal digital assistants and home appliances

ARCH v1.1—8-8

Any application that requires information to be delivered to multiple users concurrently can benefit from IP multicast. Examples include video or audio broadcast applications like video conferencing, IP telephony music on hold, corporate-wide communications, and distance learning, as well as software distribution.

IP Multicast Data Delivery Principles

IP multicast packets are replicated only at routers where paths diverge to reach the intended recipients. This topic explains the benefits of using IP multicast as an underlying transport mechanism in an enterprise network.

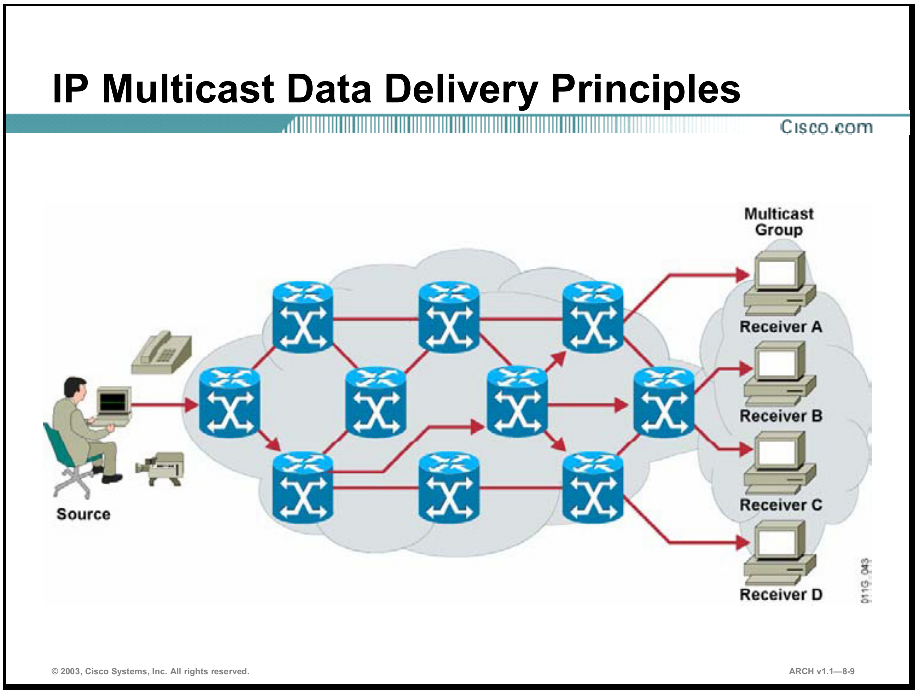

IP Multicast Data Delivery Principles

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-9

IP multicast delivers source traffic to multiple receivers while using a minimum of network bandwidth. Multicast packets are replicated in the network where paths diverge at routers enabled with PIM and other supporting multicast protocols, resulting in the most efficient delivery of data to multiple receivers.

Many alternatives to IP multicast require the source to send more than one copy of the data. Some, such as application-level multicast, require the source to send an individual copy to each receiver. Even low-bandwidth applications can benefit from using IP multicast when there are thousands of concurrent receivers. High-bandwidth applications, such as MPEG video, may require a large portion of the available network bandwidth for a single stream. In these applications, IP multicast is the only practical way to send to more than one receiver simultaneously. The figure shows how IP multicast is used to deliver data from one source to many interested recipients.

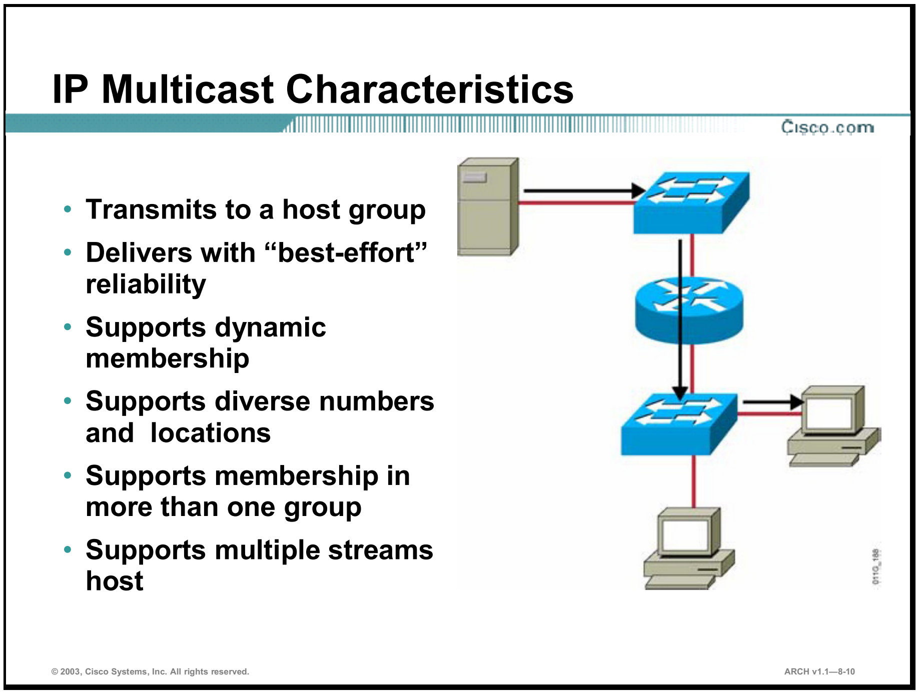

IP Multicast Characteristics

- Transmits to a host group
- Delivers with “best-effort” reliability
- Supports dynamic membership
- Supports diverse numbers and locations
- Supports membership in more than one group
- Supports multiple streams host

ARCH v1.1—8-10

IP multicast is the transmission of an IP data frame to a host group that is defined by a single IP address. IP multicasting is an extension to the standard IP network-level protocol and is described in RFC 1112, Host Extensions for IP Multicasting.

IP multicasting has these characteristics:

- Transmits IP datagrams to a host group identified by a single IP destination address. A host group is dynamic and can contain zero or more host devices at any given time.
- Delivers a multicast packet to all members of the destination host group with the same best-effort reliability as regular unicast IP datagrams.
- Supports dynamic membership of a host group.
- Supports all host groups regardless of the location or number of members.
- Supports the membership of a single host in one or more multicast groups.
- Upholds multiple data streams at the application level for a single group address.
- Supports a single group address for multiple applications on a host.

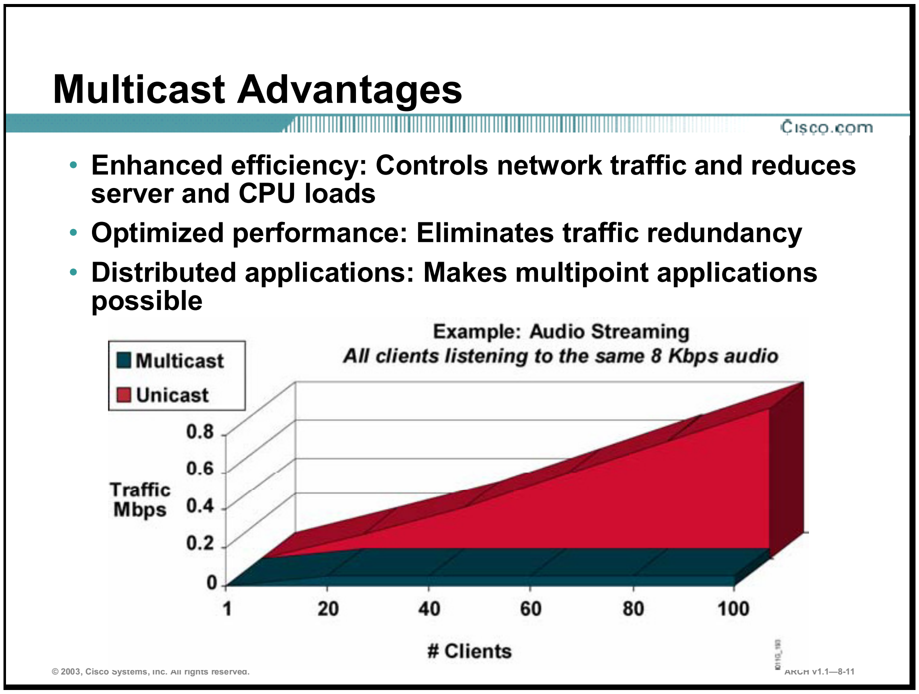

Multicast Advantages

Cisco.com

- Enhanced efficiency: Controls network traffic and reduces server and CPU loads
- Optimized performance: Eliminates traffic redundancy
- Distributed applications: Makes multipoint applications possible

ARCH v1.1—8-11

Multicast transmission affords many advantages over unicast transmission in a one-to-many or many-to-many environment, including:

- **Enhanced efficiency:** Available network bandwidth is utilized more efficiently since multiple streams of data are replaced with a single transmission.
- **Optimized performance:** Fewer copies of data require forwarding and processing.
- **Distributed applications:** In a unicast transmission, multipoint applications will not be possible as demand and usage grow because unicast transmission will not scale.

Traffic levels and clients increase at a 1:1 rate with unicast transmission. Traffic levels increase at a greatly reduced rate compared to clients when you use multicast transmission.

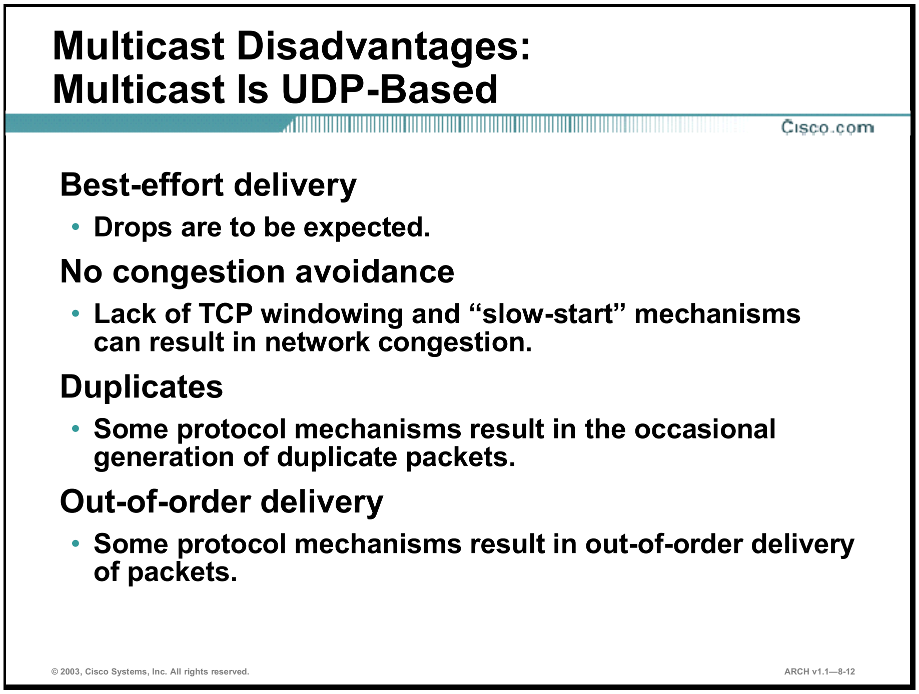

Multicast Disadvantages: Multicast Is UDP-Based

Best-effort delivery

- Drops are to be expected.

No congestion avoidance

- Lack of TCP windowing and “slow-start” mechanisms can result in network congestion.

Duplicates

- Some protocol mechanisms result in the occasional generation of duplicate packets.

Out-of-order delivery

- Some protocol mechanisms result in out-of-order delivery of packets.

ARCH v1.1—8-12

IP multicast intrinsically has a number of disadvantages. Most IP multicast applications are User Datagram Protocol (UDP)-based. This results in some undesirable side effects when compared to unicast TCP applications.

Best-effort delivery results in occasional packet drops. Many multicast applications that operate in real time (such as audio and video) may be impacted by these losses. It is not feasible to request retransmission of the lost data at the application layer for real-time applications. There also can be heavy drops on voice applications. Packet drops can result in jerky, missed speech patterns that can make content unintelligible when the drop rate gets high enough.

Moderate to heavy drops in video are sometimes fairly well tolerated by the human eye, appearing as unusual “artifacts” on the picture. However, some compression algorithms can be severally impacted by even low drop rates causing the picture to become jerky or freeze for several seconds while the decompression algorithm recovers.

Lack of congestion control can result in overall network degradation as the popularity of UDP-based multicast applications grows. The network can occasionally duplicate packets as multicast network topologies change. Multicast applications should tolerate occasional duplicate packets.

Multicast Forwarding

Multicast routing uses RPF to flood packets out from all interfaces, except incoming packets. This topic explains how multicast forwarding works with RPF.

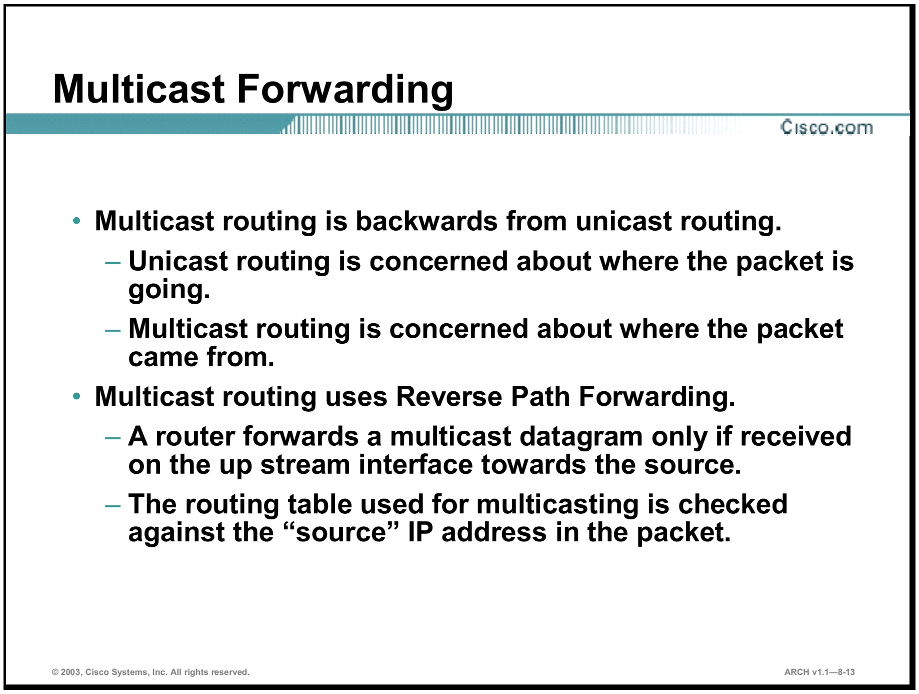

Multicast Forwarding

Cisco.com

- **Multicast routing is backwards from unicast routing.**
  - **Unicast routing is concerned about where the packet is going.**
  - **Multicast routing is concerned about where the packet came from.**
- **Multicast routing uses Reverse Path Forwarding.**
  - **A router forwards a multicast datagram only if received on the up stream interface towards the source.**
  - **The routing table used for multicasting is checked against the “source” IP address in the packet.**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-13

In multicast transmissions, routers must know packet origination, rather than destination, which is the opposite of how unicast works. In multicast transmissions:

- The origination IP address denotes the known source.
- The destination IP address denotes an unknown group of receivers.

With RPF, a broadcast floods packets out all interfaces except where packets are incoming from the source. The broadcast initially assumes that every host on the network is part of the multicast group. Pruning eliminates tree branches without multicast group members and cuts off transmission to LANs without interested receivers.

Selective forwarding requires its own integrated unicast routing protocol.

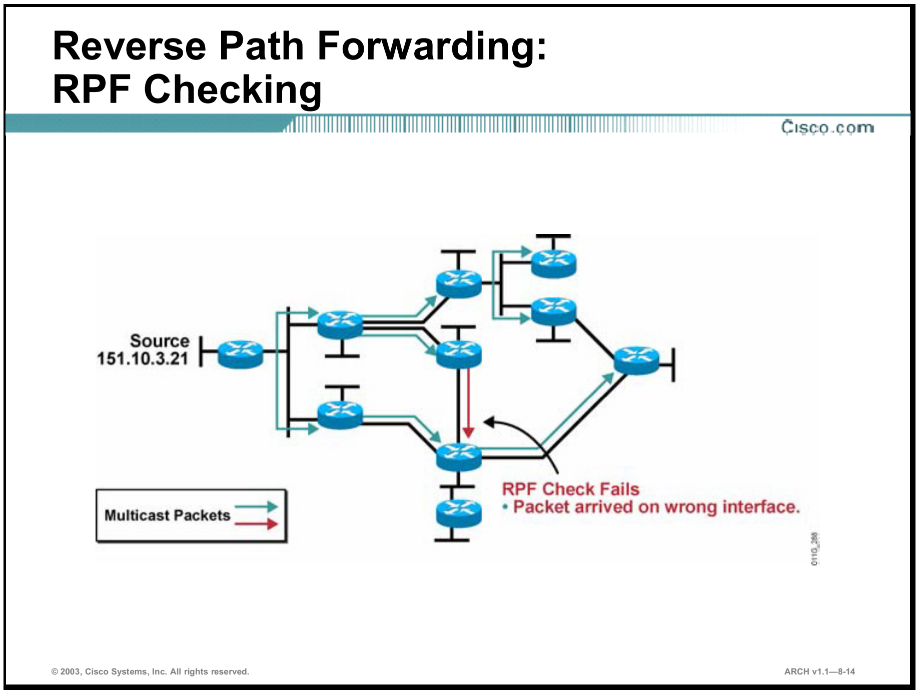

Reverse Path Forwarding: RPF Checking

ARCH v1.1—8-14

RPF checking with IP multicast forwarding operates as follows:

- The source floods the network with multicast data.
- Each router has a designated incoming interface (RPF interface) on which multicast data is received from a given source.
- Each router receives multicast data on one or more interfaces, but performs an RPF check to prevent duplicate forwarding.

In the example shown in the figure, assume that a router receives multicast data on two interfaces as follows:

- Interface 1 performs an RPF check on multicast data received on interface E0. The RPF check succeeds because data was received on the specified incoming interface from source 151.10.3.21. The data is forwarded through all outgoing interfaces on a multicast distribution tree.
- Interface 2 performs an RPF check on multicast data received on interface E1. RPF check fails because data was not received on specified incoming interface from source 151.10.3.21. Data is silently dropped.

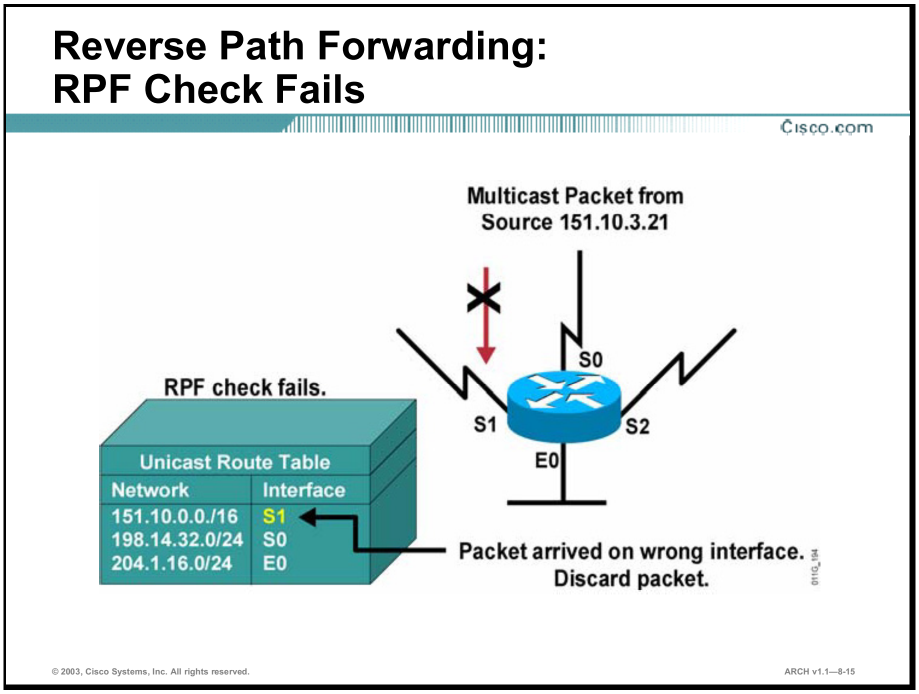

Reverse Path Forwarding: RPF Check Fails

Cisco.com

Multicast Packet from Source 151.10.3.21

RPF check fails.

Unicast Route Table	
Network	Interface
151.10.0.0./16	S1
198.14.32.0/24	S0
204.1.16.0/24	E0

Packet arrived on wrong interface.
Discard packet.

ARCH v1.1—8-15

The figure demonstrates IP multicast forwarding when an RPF check fails.

The router can only accept multicast data from source 151.10.3.21 on interface S1. Therefore, multicast data is silently dropped because it arrived on an interface not specified in the RPF check (S0).

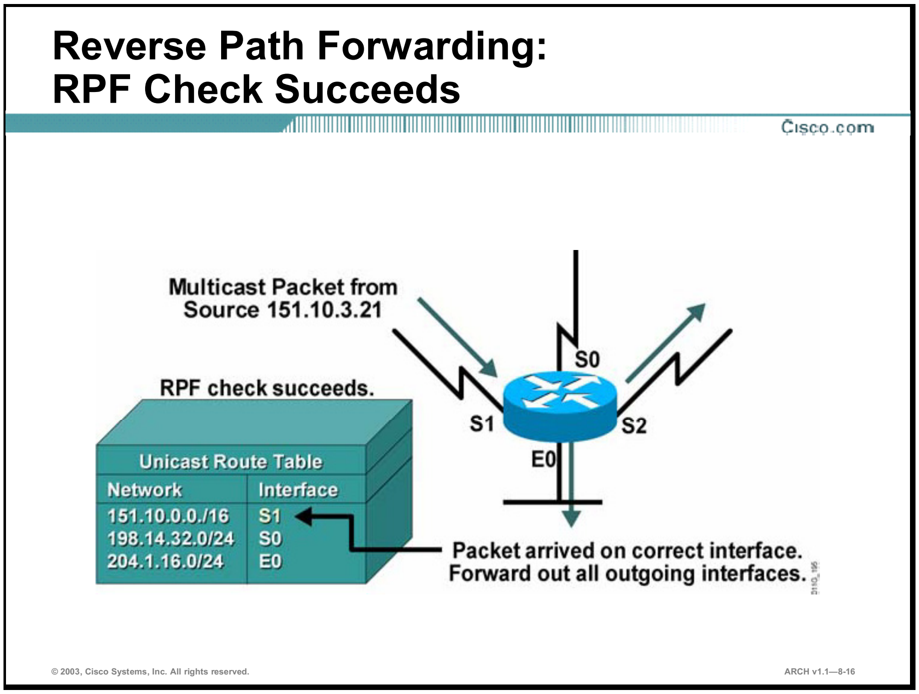

Reverse Path Forwarding: RPF Check Succeeds

Cisco.com

Multicast Packet from Source 151.10.3.21

RPF check succeeds.

Unicast Route Table	
Network	Interface
151.10.0.0./16	S1
198.14.32.0/24	S0
204.1.16.0/24	E0

Packet arrived on correct interface. Forward out all outgoing interfaces.

ARCH v1.1—8-16

The figure illustrates multicast forwarding when an RPF check succeeds.

The router can only accept multicast data from Source 151.10.3.21 on interface S1. Multicast data is forwarded out through all outgoing interfaces because the data arrived on the incoming interface specified in the RPF check (S1).

IP Multicast Group Membership and Distribution Trees

The destination of an IP multicast packet is a virtual group address. Members join a group and then begin receiving multicast packets addressed to that group. This topic explains the purpose and use of IP multicast receiver group membership.

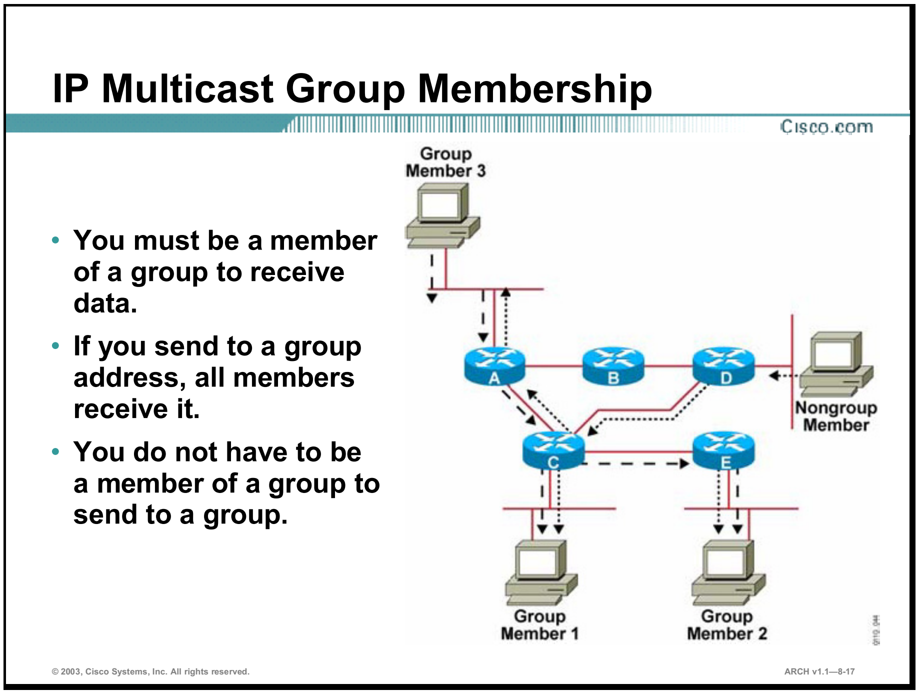

IP Multicast Group Membership

Cisco.com

- You must be a member of a group to receive data.
- If you send to a group address, all members receive it.
- You do not have to be a member of a group to send to a group.

ARCH v1.1—8-17

IP multicast relies on the concept of a virtual group address. In normal TCP/IP routing, a packet is routed from a source address to a destination address, traversing the IP network on a hop-by-hop basis. In IP multicast, the packet's destination address is not assigned to a single destination. Instead, receivers join a group, and when they join, packets addressed to that group begin flowing to them. All members of the group receive the packet. You must be a member of the group to receive the packet. Senders to the group do not need to be members of that group.

In the figure, packets that are sent to group addresses go to all group members, but no non-group members.

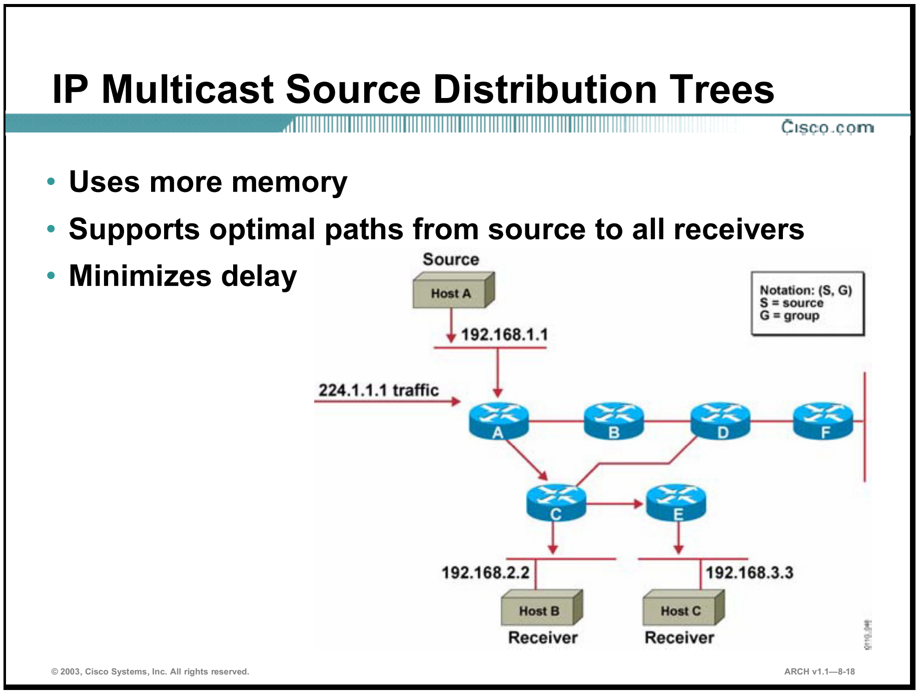

IP Multicast Source Distribution Trees

- Uses more memory
- Supports optimal paths from source to all receivers
- Minimizes delay

ARCH v1.1—8-18

Multicast-capable routers create distribution trees that control the path that IP multicast traffic takes through the network to deliver traffic to all receivers. The two basic types of multicast distribution trees are source trees and shared trees.

The simplest form of a multicast distribution tree is a source tree with its root at the source and branches forming a spanning tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT).

The figure shows an example of an SPT for group 224.1.1.1 rooted at the source, host A, and connecting two receivers, hosts B and C.

The special notation of (S,G), pronounced “S comma G,” enumerates an SPT where S is the IP address of the source and G is the multicast group address. Using this notation, the SPT for the example shown in the figure would be (192.168.1.1, 224.1.1.1).

The (S,G) notation implies that a separate SPT exists for each individual source sending to each group, which is correct. For example, if host B is also sending traffic to group 224.1.1.1 and hosts A and C are receivers, a separate (S,G) SPT would exist with a notation of (192.168.2.2, 224.1.1.1).

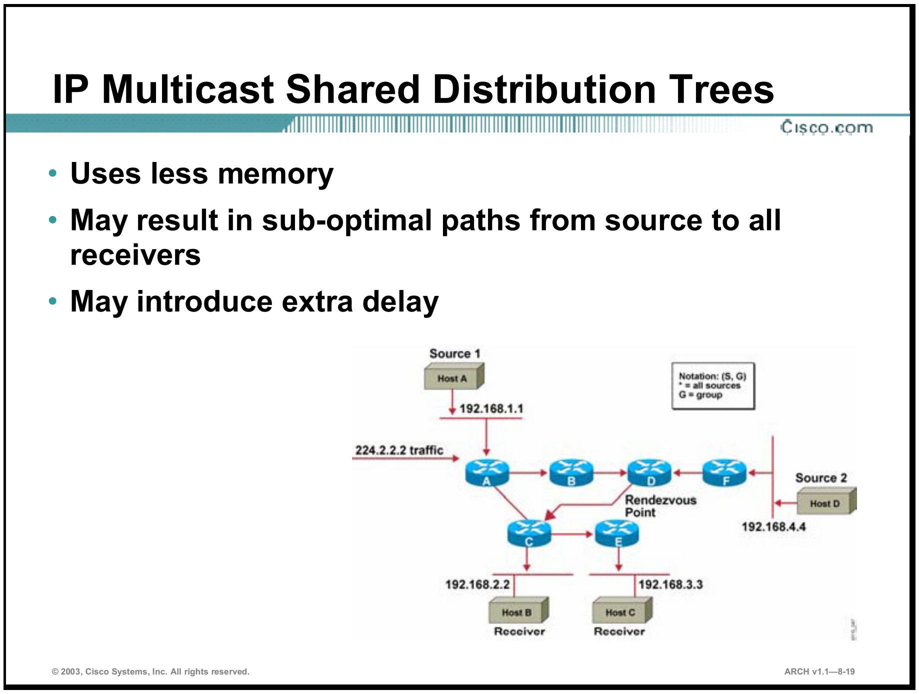

IP Multicast Shared Distribution Trees

Cisco.com

- Uses less memory
- May result in sub-optimal paths from source to all receivers
- May introduce extra delay

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-19

Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called a rendezvous point (RP). The figure shows a shared tree for the group 224.2.2.2 with the root located at router D. This shared tree is unidirectional. Source traffic is sent towards the RP on a source tree. The traffic is then forwarded down the shared tree from the RP to reach all of the receivers, unless the receiver is located between the source and the RP, in which case it will be serviced directly.

With multicast Layer 2 addresses, flowing the multicast traffic through an RP saves memory. Alternatively, the multicast traffic can flow through the RP and then through the SPT.

In the example, multicast traffic from the sources, hosts A and D, travels to the root (router D) and then down the shared tree to the two receivers, hosts B and C. Because all sources in the multicast group use a common shared tree, a wildcard notation written as (*,G), pronounced “star comma G,” represents the tree, followed by an (S,G) entry with a subset outgoing interface list.

Source Trees Versus Shared Trees

Both source trees and shared trees are loop-free. Messages are replicated only where the tree branches.

Members of multicast groups can join or leave at any time; therefore the distribution trees must be dynamically updated. When all the active receivers on a particular branch stop requesting the traffic for a particular multicast group, the routers prune that branch from the distribution tree and stop forwarding traffic down that branch. If one receiver on that branch becomes active and requests the multicast traffic, the router will dynamically modify the distribution tree and start forwarding traffic again.

Source trees have the advantage of creating the optimal path between the source and the receivers. This advantage guarantees the minimum amount of network latency for forwarding multicast traffic. However, this optimization comes at a cost: the routers must maintain path information for each source. In a network that has thousands of sources and thousands of groups, this overhead can quickly become a resource issue on the routers. Memory consumption from the size of the multicast routing table is a factor that network designers must take into consideration. The multicast routing table is required to maintain current values, called state, that determine multicast routing behavior.

Shared trees have the advantage of requiring the minimum amount of state in each router. The disadvantage of shared trees is that under certain circumstances the paths between the source and receivers might not be the optimal paths, which might introduce some latency in packet delivery. For example, in the figure, the shortest path between host A (source 1) and host B (a receiver) would be router A and router C. Because we are using router D as the root for a shared tree, the traffic must traverse routers A, B, D and then C.

These characteristics describe distribution trees:

- **Source or shortest path trees:** These trees use more router memory, but benefit from optimal paths from source to all receivers. They also minimize delay.
- **Shared trees:** These trees use less memory, but may result in sub-optimal paths from source to all receivers. They may introduce extra delay.

Protocol Independent Multicast

PIM uses unicast routing information to perform the multicast forwarding function. Although PIM is called a multicast routing protocol, it actually uses the unicast routing table to perform the RPF check function instead of building up a completely independent multicast routing table. This topic describes the operation of PIM for IP multicast.

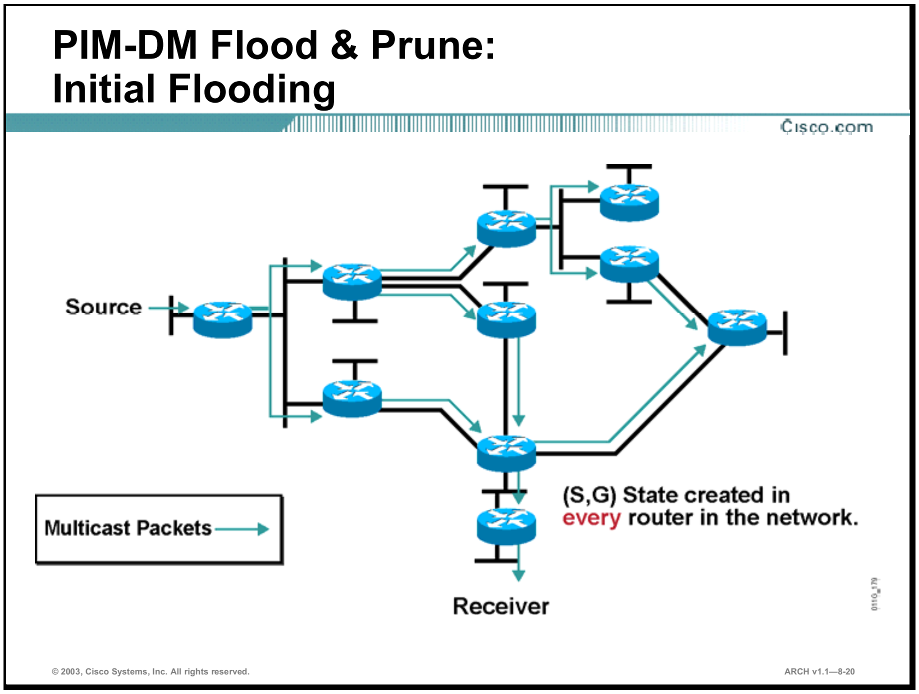

PIM-DM Flood & Prune:
Initial Flooding

Cisco.com

Source

Receiver

Multicast Packets

(S,G) State created in every router in the network.

0110_179

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-20

PIM dense mode (PIM-DM) uses a push model to flood multicast traffic to every corner of the network. This push model is a brute-force method for delivering data to the receivers. This method is efficient in deployments where there are active receivers on every subnet in the network.

PIM-DM initially floods multicast traffic throughout the network. Routers that have no downstream neighbors prune back the unwanted traffic. This process repeats every three minutes.

Routers accumulate state information by receiving data streams through the flood and prune mechanism. These data streams contain the source and group information so that downstream routers can build up their multicast forwarding table. PIM-DM supports only source trees, that is, (S,G) entries, and cannot be used to build a shared distribution tree.

The figure shows how PIM-DM initially floods the network, sending the (S,G) state to very router on the network.

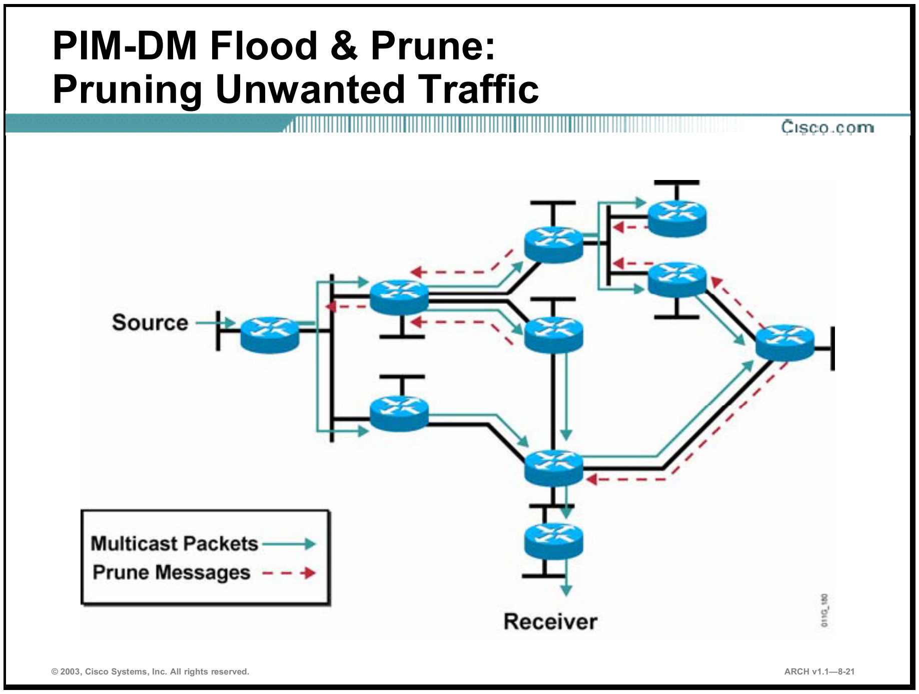

PIM-DM Flood & Prune: Pruning Unwanted Traffic

ARCH v1.1—8-21

After flooding, PIM-DM sends prune messages to each router on the network to determine the location of the receiver.

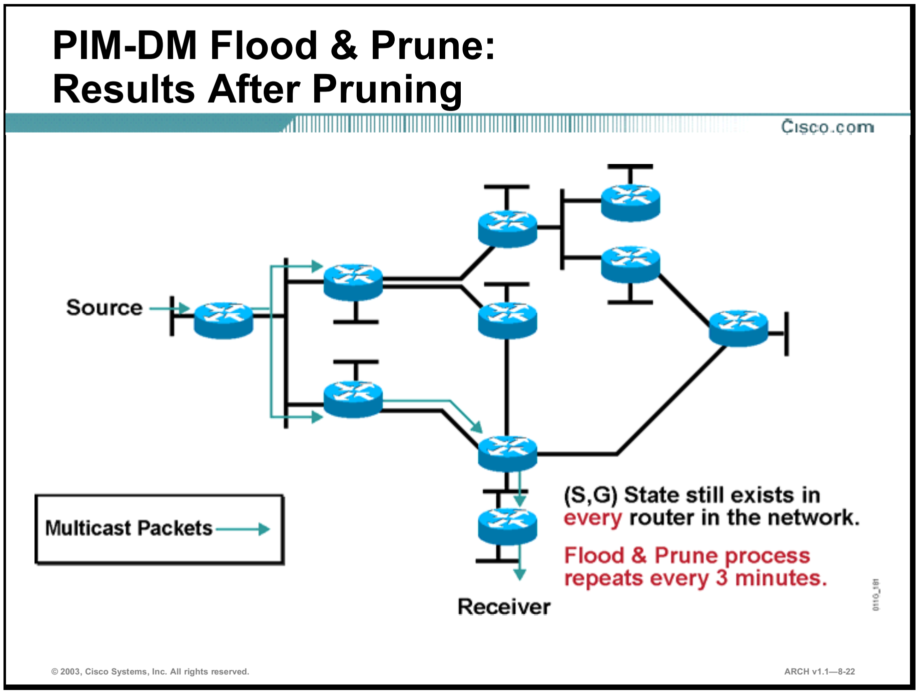

PIM-DM Flood & Prune: Results After Pruning

(S,G) State still exists in every router in the network.

Flood & Prune process repeats every 3 minutes.

ARCH v1.1—8-22

After pruning, the multicast packets are sent only to the receiver. The flood and prune process repeats every three minutes.

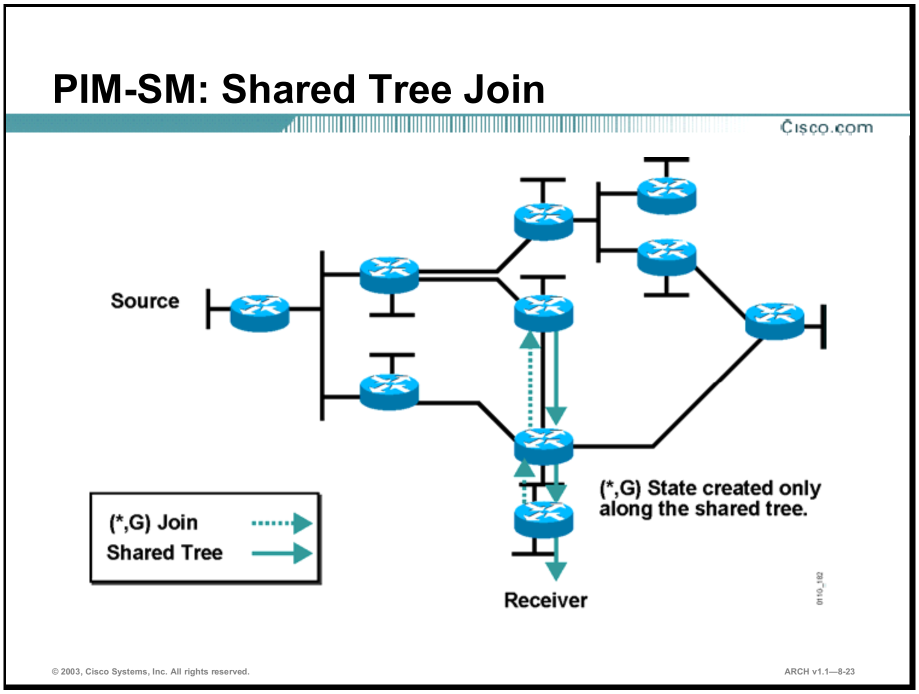

PIM-SM: Shared Tree Join

(*,G) State created only along the shared tree.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-23

PIM sparse mode (PIM-SM) uses a pull model to deliver multicast traffic. Only network segments with active receivers that have explicitly requested the data will receive the traffic.

PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees (at least, initially), it requires the use of an RP. The RP must be administratively configured in the network.

Sources register with the RP and then data is forwarded down the shared tree to the receivers. The edge routers learn about a particular source when they receive data packets on the shared tree from that source through the RP. The edge router then sends PIM (S,G) Join messages toward that source. Each router along the reverse path compares the unicast routing metric of the RP address to the metric of the source address. If the metric for the source address is better, it will forward a PIM (S,G) Join message towards the source. If the metric for the RP is the same or better, then the PIM (S,G) Join message will be sent in the same direction as the RP. In this case, the shared tree and the source tree would be considered congruent.

In the figure, an active receiver has joined multicast group G. The router knows the IP address of the RP for group G and it sends a (*,G) Join for this group towards the RP. This (*,G) Join travels hop-by-hop to the RP, building a branch of the shared tree that extends from the RP to the last-hop router directly connected to the receiver. At this point, group G traffic can flow down the shared tree to the receiver.

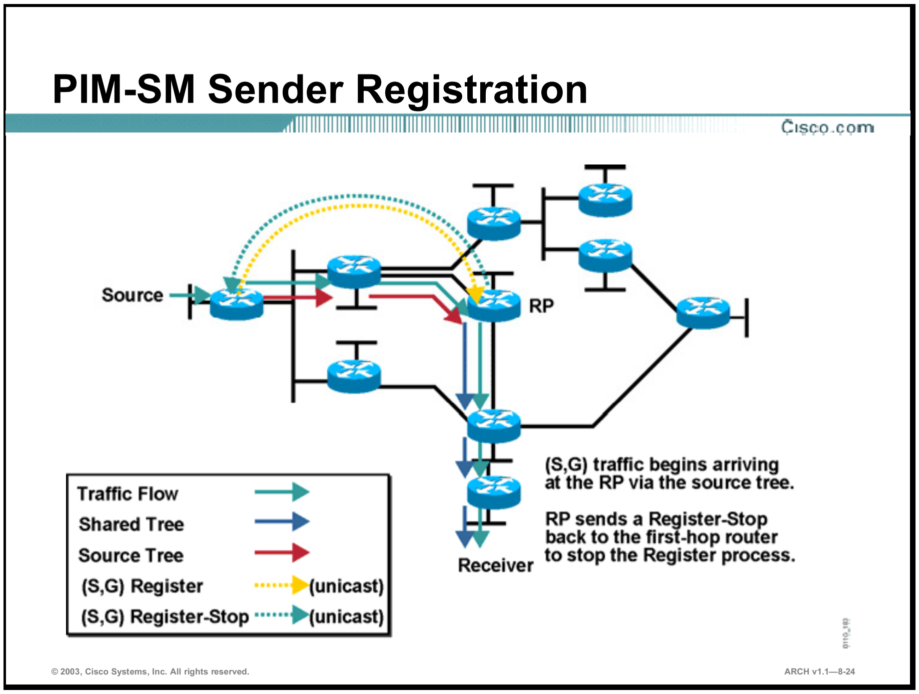

PIM-SM Sender Registration

(S,G) traffic begins arriving at the RP via the source tree.

RP sends a Register-Stop back to the first-hop router to stop the Register process.

ARCH v1.1—8-24

As soon as an active source for group G sends a packet to the router that is attached to this source, the router is responsible for “registering” the source with the RP and requesting the RP to build a tree back to that router. The source router encapsulates the multicast data from the source in a special PIM-SM message called the Register message and unicasts that data to the RP.

When the RP receives the Register message it does two things:

- The RP unencapsulates the multicast data packet inside of the Register message and forwards it down the shared tree.
- The RP sends an (S,G) Join back towards the source network S to create a branch of an (S,G) shortest path tree. This results in (S,G) state being created in all the routers along the SPT, including the RP.

As soon as the shortest path tree is built from the source router to the RP, multicast traffic begins to flow from source S to the RP. Once the RP begins receiving data (that is, down the shortest path tree) from source S, it sends a register stop to the first-hop router of the source to inform it that it can stop sending the unicast Register messages.

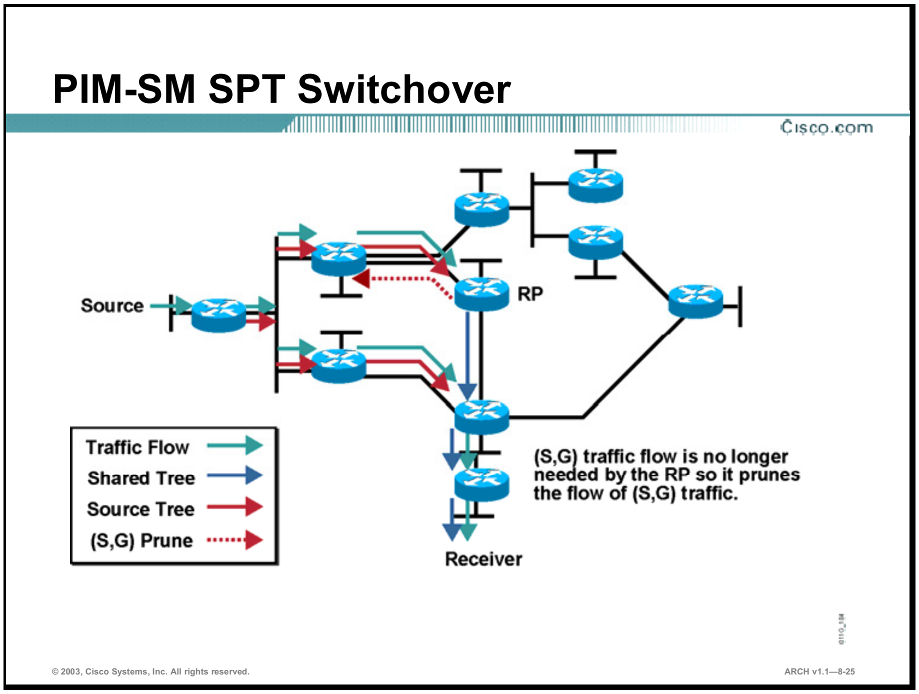

PIM-SM SPT Switchover

(S,G) traffic flow is no longer needed by the RP so it prunes the flow of (S,G) traffic.

ARCH v1.1—8-25

PIM-SM includes the capability for last-hop routers (that is, routers with directly connected members) to switch to the shortest path tree and bypass the RP if the traffic rate is above a set threshold, called the SPT-threshold.

The default value of the SPT-threshold in Cisco routers is zero. This means that the default behavior for PIM-SM leaf routers attached to active receivers is to immediately join the shortest path tree to the source as soon as the first packet arrives via the (*,G) shared tree.

In the figure, the last-hop router sends an (S,G) Join message toward the source to join the shortest path tree and bypass the RP.

The (S,G) Join messages travel hop-by-hop to the first-hop router (the router connected directly to the source), thereby creating another branch of the shortest path tree. This also creates (S,G) state in all the routers along this branch of the shortest path tree.

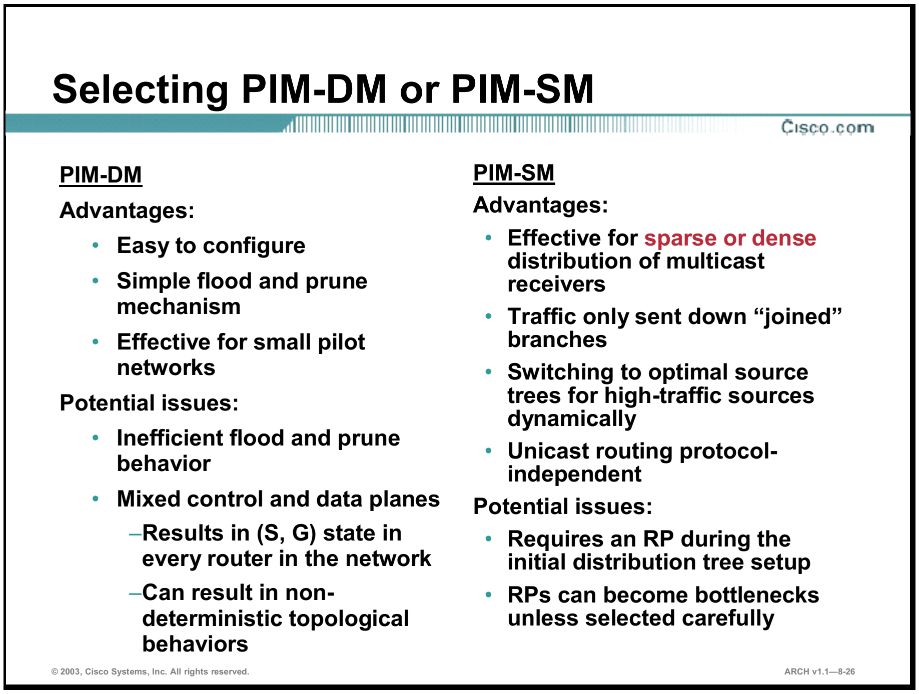

Selecting PIM-DM or PIM-SM

Cisco.com

PIM-DM

Advantages:

- Easy to configure
- Simple flood and prune mechanism
- Effective for small pilot networks

Potential issues:

- Inefficient flood and prune behavior
- Mixed control and data planes
  - Results in (S, G) state in every router in the network
  - Can result in non-deterministic topological behaviors

PIM-SM

Advantages:

- Effective for sparse or dense distribution of multicast receivers
- Traffic only sent down “joined” branches
- Switching to optimal source trees for high-traffic sources dynamically
- Unicast routing protocol-independent

Potential issues:

- Requires an RP during the initial distribution tree setup
- RPs can become bottlenecks unless selected carefully

ARCH v1.1—8-26

PIM-DM is appropriate for a large number of densely distributed receivers located in close proximity to the source. It offers these advantages:

- Minimal number of commands required for configuration (two)
- Simple mechanism for reaching all possible receivers and eliminating distribution to uninterested receivers
- Simple behavior is easier to understand and therefore easier to debug

When configuring a router, use PIM sparse-dense mode to ease the configuration. The router uses PIM-SM if it hears from a rendezvous point. Otherwise, it uses PIM-DM.

A potential issue with PIM-DM is the need to flood frequently because prunes expire after three minutes.

Use PIM-SM for sparse or dense distribution of multicast receivers (no necessity to flood). It offers these advantages:

- Traffic is sent only to registered receivers that have explicitly joined the multicast group.
- RP can be switched to the optimal shortest-path tree when high-traffic sources are forwarding to a sparsely distributed receiver group.

Potential issues with PIM-SM include:

- PIM-SM requires an RP during the initial setup of the distribution tree (it can switch to the shortest-path tree once RP is established and determined as optimal). RPs can become bottlenecks if not selected with great care.
- PIM-SM’s complex behavior is difficult to understand and therefore difficult to debug.

IP Multicast Control Mechanisms

By default, a data link layer switch will forward all multicast traffic to every port that belongs to the destination LAN. IP multicast control mechanisms limit multicast traffic to the ports that need to receive the data. This topic explains the purpose and function of IP multicast control mechanisms.

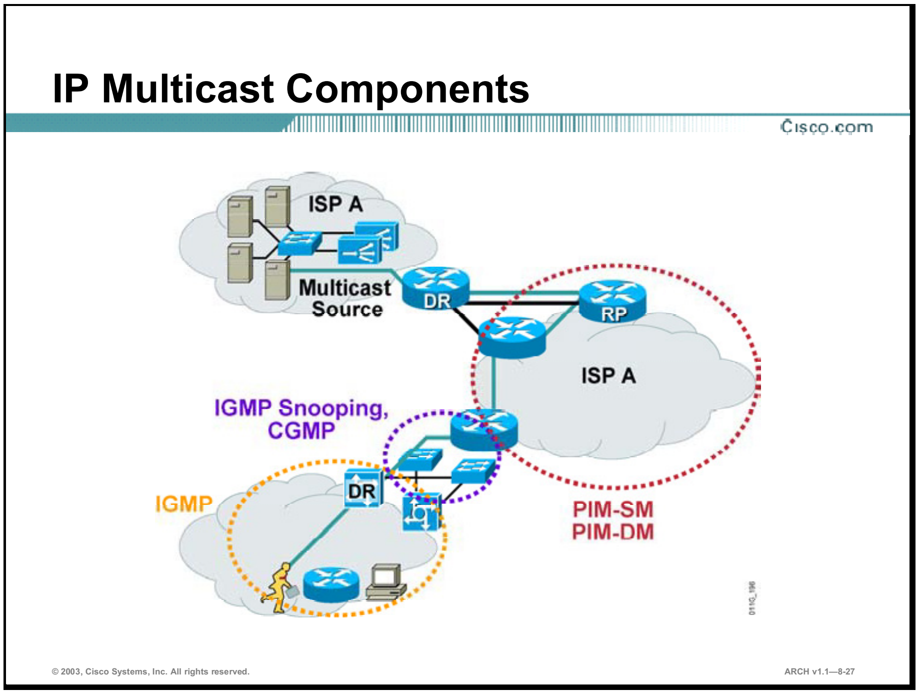

IP Multicast Components

ARCH v1.1—8-27

Intradomain multicast supports multicast applications within an enterprise campus. IOS multicast technologies leverage network resources for massively scalable content distribution applications. The table describes the IP multicast control mechanisms that IOS supports.

Feature	Description
CGMP[1]	■ Cisco-developed protocol that allows data link layer switches to leverage IGMP[2] information on Cisco routers to make data link layer forwarding decisions ■ Provides management of group membership on switched Ethernet LANs ■ Allows switches to forward multicast traffic to only those ports that are interested in the traffic ■ Used in low-end or older Catalyst series switches that do not support IGMP snooping ■ Fully interoperable with IGMP snooping
IGMP	■ Used by IP routers and their immediately connected IPv4 hosts to communicate multicast group membership states to neighboring multicast routers ■ Version 3 of IGMP adds support for "source filtering," the ability for a system to report interest in receiving packets only from specific source addresses, or from all but specific source addresses sent to a particular multicast address

Feature	Description
IGMP snooping	■ Requires the LAN switch to examine, or "snoop," some network layer information in the IGMP packet sent from the host to the router ■ Used in higher-end, hardware-enabled platforms
PIMv2[3]	■ Provides intradomain multicast forwarding for all underlying unicast routing protocols ■ Independent from any underlying unicast protocol such as OSPF or BGP[4] ■ Supports explicit join (sparse mode), flood-and-prune (dense mode), or hybrid sparse-dense modes ■ Sparse Mode: Relies upon an explicit joining method before attempting to send multicast data to receivers of a multicast group ■ Dense Mode: Actively attempts to send multicast data to all potential receivers (flooding) and relies upon their self-pruning (removal from group) to achieve desired distribution

[1] CGMP = Cisco Group Management Protocol

[2] IGMP = Internet Group Management Protocol

[3] PIMv2 = Protocol Independent Multicast v2

[4] BGP = Border Gateway Protocol

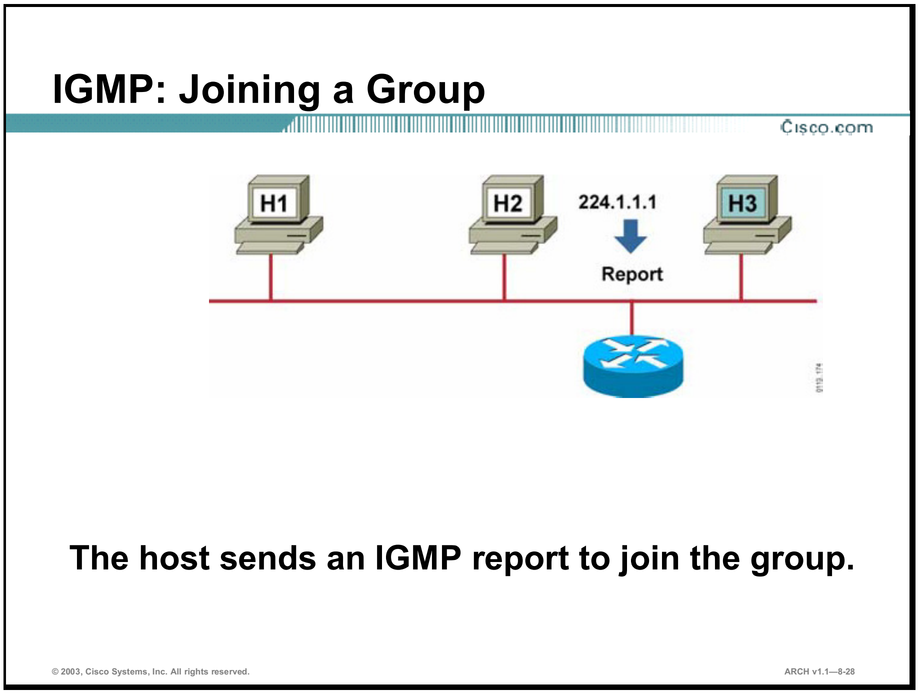

IGMP: Joining a Group

Cisco.com

The host sends an IGMP report to join the group.

ARCH v1.1—8-28

With Internet Group Management Protocol (IGMP), members joining a group do not have to wait for a query to join. They send in an unsolicited report indicating their interest. This reduces join latency for the end-system joining if no other members are present.

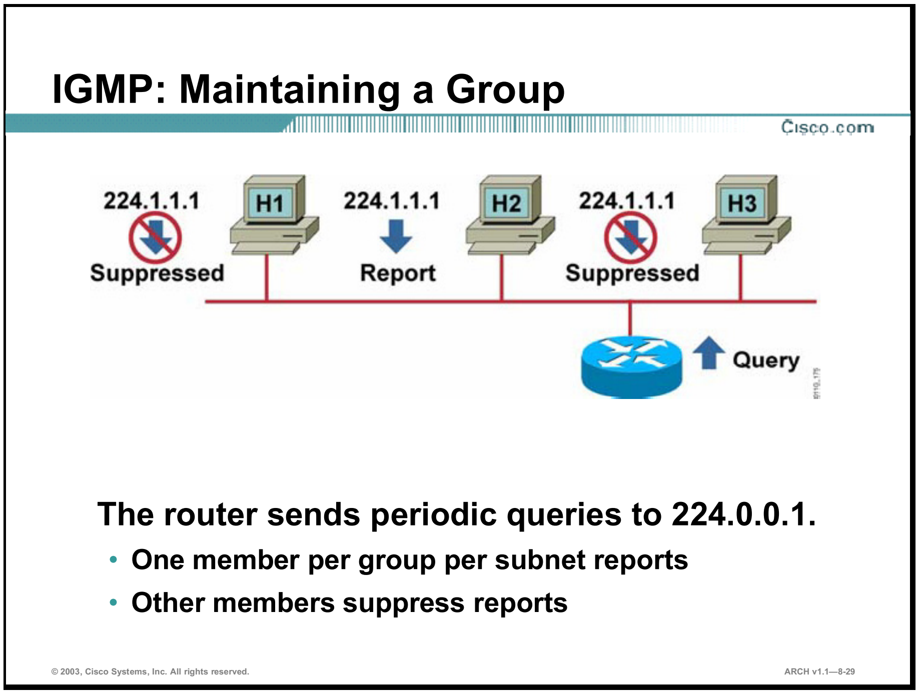

IGMP: Maintaining a Group

The router sends periodic queries to 224.0.0.1.

- **One member per group per subnet reports**
- **Other members suppress reports**

ARCH v1.1—8-29

The router multicasts periodic IGMPv1 membership queries to the “all hosts” (224.0.0.1) group address.

Only one member per group responds with a report to a query to save bandwidth on the subnet network and processing by the hosts. This process is called response suppression. The response suppression mechanism is accomplished as follows:

- When a host receives the query, it starts a countdown timer for each multicast group of which it is a member. The countdown timers are each initialized to a random count within a given time range. (In IGMPv1 this was a fixed range of ten seconds. Therefore, the countdown timers were randomly set to some value between zero and ten seconds.)
- When a countdown timer reaches zero, the host sends a membership report for the group associated with the countdown timer to notify the router that the group is still active.
- However, if a host receives a membership report before its associated countdown timer reaches zero, it cancels the countdown timer associated with the multicast group, thereby suppressing its own report.

In the figure, H2’s time expired first so it responded with its membership report. H1 and H3 cancelled their timers associated with the group, thereby suppressing their reports.

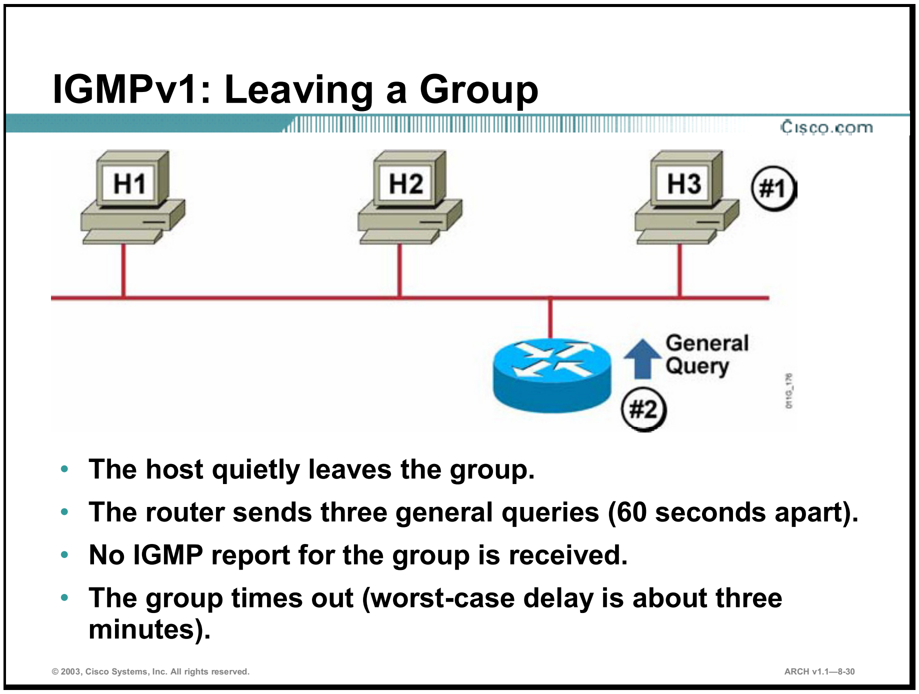

IGMPv1: Leaving a Group

- The host quietly leaves the group.
- The router sends three general queries (60 seconds apart).
- No IGMP report for the group is received.
- The group times out (worst-case delay is about three minutes).

ARCH v1.1—8-30

When hosts want to leave a multicast group, they can either ignore the periodic general queries sent by the multicast router (IGMP v1 host behavior), or they can send an IGMP leave (IGMP v2 host behavior). When the switch receives a leave message, it sends out a MAC-based general query on the port on which it received the leave message to determine if any devices connected to this port are interested in traffic for the specific multicast group.

If no IGMP report is received for any of the IP multicast groups that map to the MAC multicast group address, the port is removed from the multicast forwarding entry. If the port is not the last non-multicast-router port in the entry, the switch suppresses the IGMP leave (it is not sent to the router). If the port is the last non-multicast-router port in the entry, the IGMP leave is forwarded to the multicast router ports and the MAC group forwarding entry is removed.

When the router receives the IGMP leave, it sends several IGMP group-specific queries. If no Join messages are received in response to the queries, and there are no downstream routers connected through that interface, the router removes the interface from the IP multicast group entry in the multicast routing table.

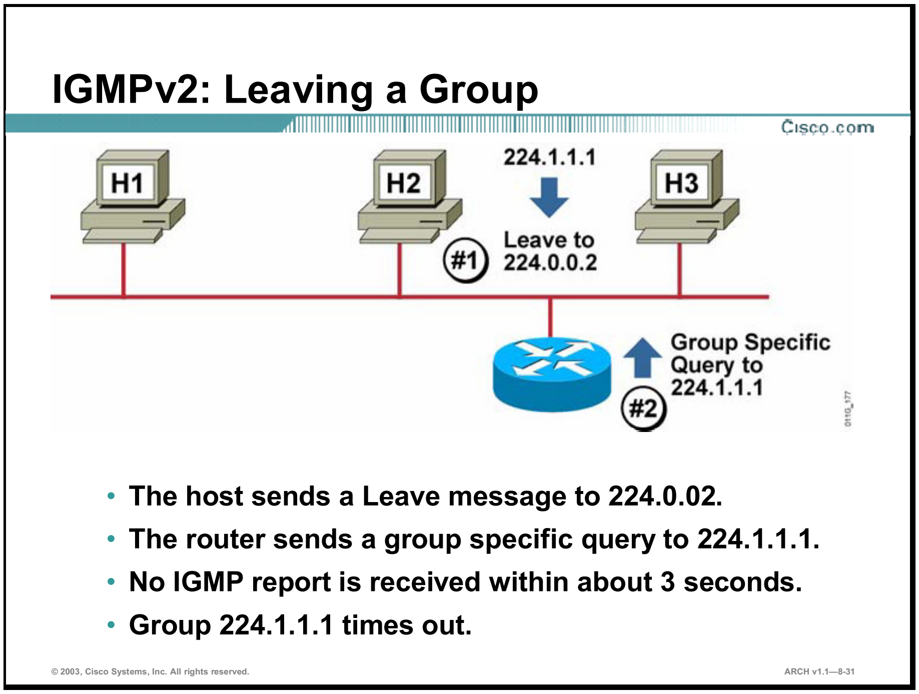

IGMPv2: Leaving a Group

- The host sends a Leave message to 224.0.02.
- The router sends a group specific query to 224.1.1.1.
- No IGMP report is received within about 3 seconds.
- Group 224.1.1.1 times out.

ARCH v1.1—8-31

The figure shows how a host leaves a multicast group with IGMP v2. The host can send an IGMP leave (IGMP v2 host behavior).

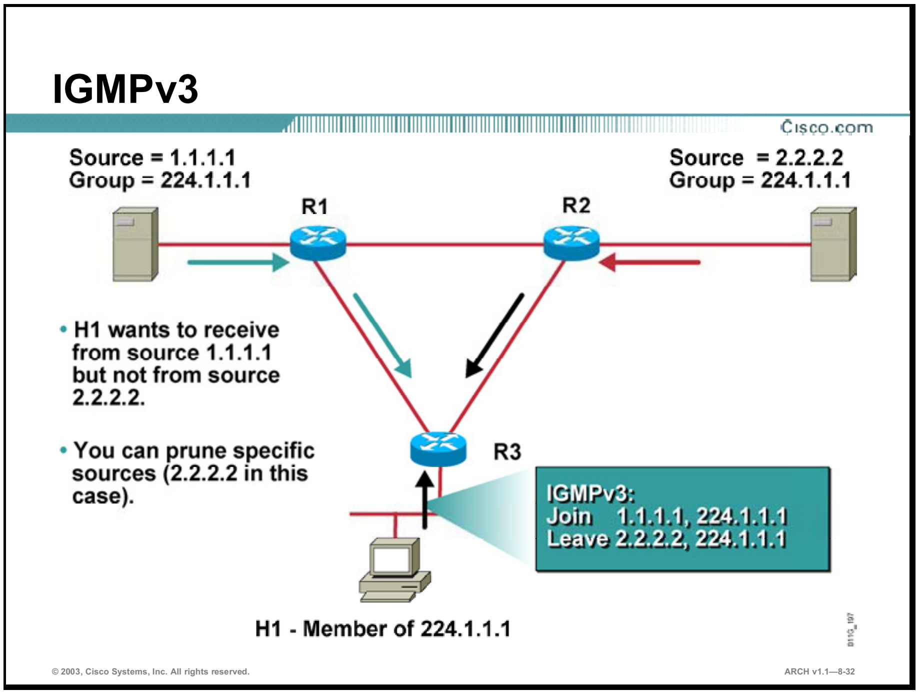

IGMPv3

- H1 wants to receive from source 1.1.1.1 but not from source 2.2.2.2.
- You can prune specific sources (2.2.2.2 in this case).

IGMPv3:
Join 1.1.1.1, 224.1.1.1
Leave 2.2.2.2, 224.1.1.1

ARCH v1.1—8-32

IGMPv3 is the third version of the IETF standards-track protocol in which hosts signal membership to last-hop routers of multicast groups. IGMPv3 introduces the ability for hosts to signal group membership with filtering capabilities with respect to sources. A host can either signal that it wants to receive traffic from all sources sending to a group except for some specific sources (called Exclude mode), or that it wants to receive traffic only from some specific sources sending to the group (called Include mode).

In the example, host H1 has joined group 224.1.1.1 but only wishes to receive traffic from Source 1.1.1.1. Using an unspecified IGMPv3 mechanism, the host can inform the designated router, R3, that it is only interested in multicast traffic from Source 1.1.1.1 for Group 224.1.1.1. Router R3 could then potentially prune this specific (S,G) traffic source.

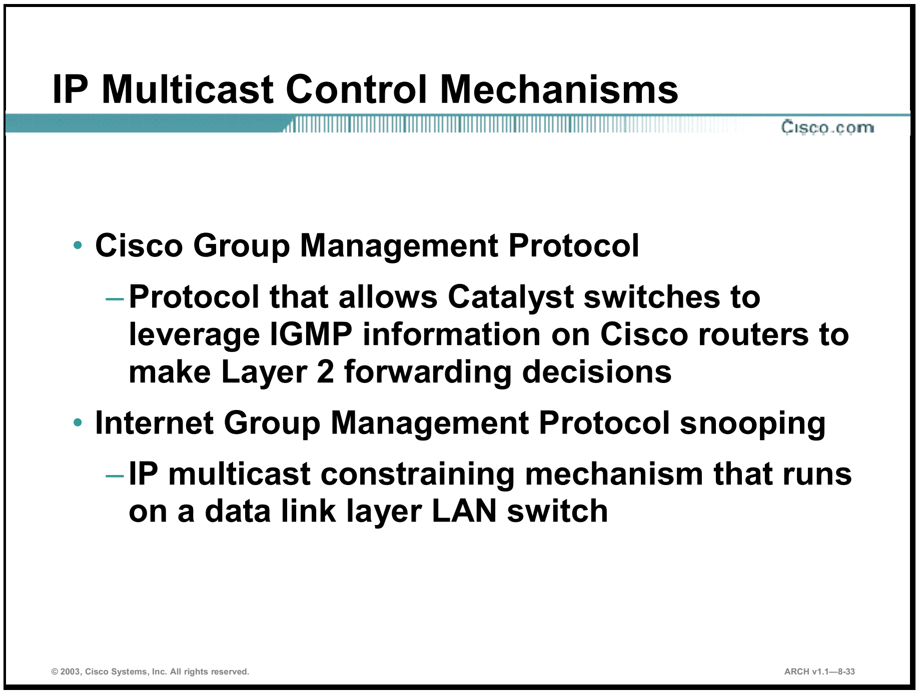

IP Multicast Control Mechanisms

- Cisco Group Management Protocol
  - Protocol that allows Catalyst switches to leverage IGMP information on Cisco routers to make Layer 2 forwarding decisions
- Internet Group Management Protocol snooping
  - IP multicast constraining mechanism that runs on a data link layer LAN switch

ARCH v1.1—8-33

The default behavior for a data link layer switch is to forward all multicast traffic to every port that belongs to the destination LAN on the switch. This behavior reduces the bandwidth efficiency of the switch because it does not limit traffic to only the ports that need to receive the data.

Two methods more efficiently handle IP multicast in a data link layer-switching environment:

- **Cisco Group Management Protocol (CGMP):** CGMP is a Cisco-developed protocol that allows Catalyst switches to leverage IGMP information on Cisco routers to make data link layer-forwarding decisions. You must configure CGMP on the multicast routers and the data link layer switches. With CGMP, IP multicast traffic is delivered only to those Catalyst switch ports that are attached to interested receivers. All other ports that have not explicitly requested the traffic will not receive it unless these ports are connected to a multicast router. Multicast router ports must receive every IP multicast data packet.
- **IGMP snooping:** IGMP snooping is an IP multicast constraining mechanism that runs on a data link layer LAN switch. IGMP snooping requires the LAN switch to examine, or “snoop,” some network layer information (IGMP Join/Leave messages) in the IGMP packets sent between the hosts and a router or multilayer switch. When the switch hears the IGMP host report from a host for a particular multicast group, the switch adds the port number of the host to the associated multicast table entry. When the switch hears the IGMP leave group message from a host, the switch removes the table entry of the host.

CGMP and IGMP snooping are used on subnets that include end users or receiver clients. Some older Cisco routers support only CGMP. However, IGMP snooping is preferred if the router supports it.

Summary

This topic summarizes the key points discussed in this lesson.

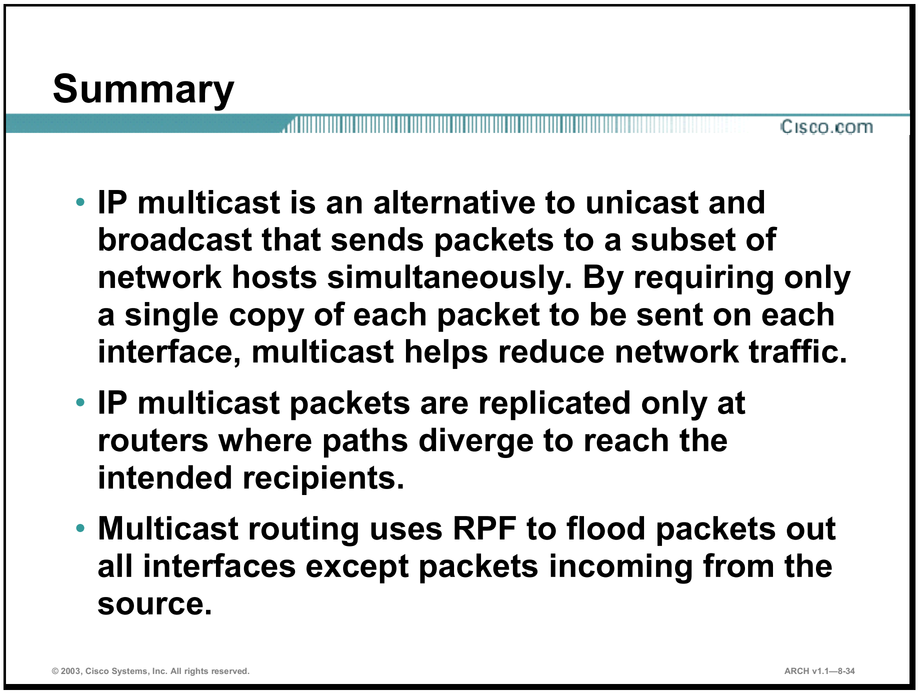

Summary

- IP multicast is an alternative to unicast and broadcast that sends packets to a subset of network hosts simultaneously. By requiring only a single copy of each packet to be sent on each interface, multicast helps reduce network traffic.
- IP multicast packets are replicated only at routers where paths diverge to reach the intended recipients.
- Multicast routing uses RPF to flood packets out all interfaces except packets incoming from the source.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—8-34

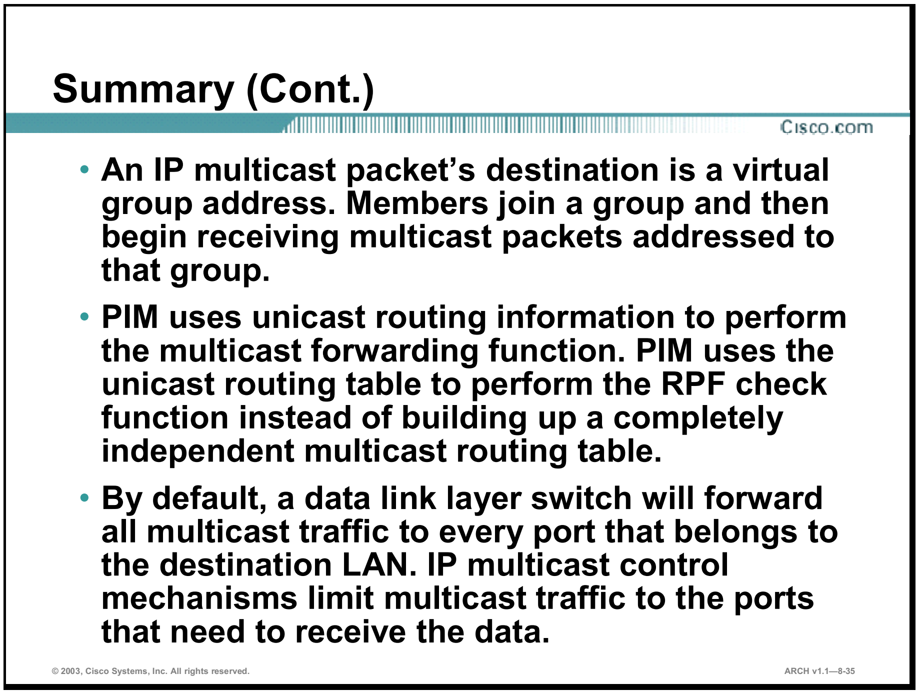

Summary (Cont.)

- An IP multicast packet's destination is a virtual group address. Members join a group and then begin receiving multicast packets addressed to that group.
- PIM uses unicast routing information to perform the multicast forwarding function. PIM uses the unicast routing table to perform the RPF check function instead of building up a completely independent multicast routing table.
- By default, a data link layer switch will forward all multicast traffic to every port that belongs to the destination LAN. IP multicast control mechanisms limit multicast traffic to the ports that need to receive the data.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—8-35

References

For additional information, refer to these resources:

- *Multicast Services* at http://www.cisco.com/warp/public/732/Tech/multicast/index.shtml
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

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) What differentiates IP multicast from other transmission modes?

A) IP multicast sends packets to a single host.

B) IP multicast sends packets to a subset of hosts.

C) IP multicast sends packets to all hosts sequentially.

D) IP multicast sends packets to all hosts simultaneously.

Q2) What is a benefit of using IP multicast to deliver source traffic to multiple receivers?

A) It guarantees packet delivery.

B) It reduces network bandwidth consumption.

C) It is highly efficient for sending single stream application traffic.

D) It replicates packets at all network devices to enable multiple client requests.

Q3) Where in the network are IP multicast packets replicated?

A) at the source

B) at the destination

C) at routers where the paths to the recipients diverge

D) at each router included in the path to each recipient

Q4) In order for a broadcast to flood packets out all interfaces, except those incoming from the source, multicast routing utilizes _____.

A) unicasting

B) multicast routers

C) Reverse Path Forwarding (RPF)

D) Open Shortest Path First (OSPF)

Q5) What is one potential drawback to source distribution trees compared to shared distribution trees?

A) increased latency

B) increased memory overhead

C) sub-optimal path calculations

D) increased bandwidth utilization

Q6) In what type of deployments would PIM-DM be efficient?

A) deployments that use shared distribution trees

B) deployments where switches are used pervasively

C) deployments where only a few receivers need IP multicast content

D) deployments where there are active receivers on every subnet in the network

Q7) You are designing the network to support company-wide video distribution of important announcements from the head of your company. You expect almost all local users across the entire network to tune in live. Users in other locales or those who work off hours will view the presentation using video on demand. What would be the best method of delivery for this announcement?

A) PIM dense mode

B) source tree mode

C) PIM sparse mode

D) reverse path forwarding

Q8) Which IOS feature provides IP multicast forwarding?

A) PIM

B) IGMP

C) CGMP

D) MBGP

Q9) What purpose is served by IGMP in IP multicast?

A) IGMP performs the RPF check.

B) IGMP registers hosts in a multicast group.

C) IGMP provides reliable multicast transport.

D) IGMP performs the multicast forwarding function.

Quiz Answer Key

Q1) B

Relates to: Introducing IP Multicast

Q2) B

Relates to: Introducing IP Multicast

Q3) C

Relates to: IP Multicast Data Delivery Principles

Q4) C

Relates to: Multicast Forwarding

Q5) B

Relates to: IP Multicast Group Membership and Distribution Trees

Q6) D

Relates to: Protocol Independent Multicast

Q7) C

Relates to: Protocol Independent Multicast

Q8) C

Relates to: IP Multicast Control Mechanisms

Q9) B

Relates to: IP Multicast Control Mechanisms

Designing IP Multicast Solutions for Enterprise Networks

Overview

Multicast deployments require three elements: the application, the network infrastructure, and client devices. Cisco IOS multicast technologies reside in the network infrastructure in Cisco routers and switches. Unlike first-generation video broadcast applications that require a separate stream for each viewer, IOS multicast is highly scalable. Multicast comprises a single content stream that is replicated by the network at branch points closest to viewers. This uses bandwidth much more efficiently, and also greatly decreases load on content servers, reaching more users at a lower cost per user.

Relevance

Most enterprises find it essential to deploy customer care, e-learning, e-commerce, and supply-chain management applications over their data networks. Companies are investing in these and other network-enabled applications to attain and keep a competitive edge in an increasingly fast-paced economy. IP multicast technologies enable the efficient deployment of these applications.

Objectives

Upon completing this lesson, you will be able to design an IP multicast solution in an existing unicast infrastructure, given specific network and application needs. This includes being able to meet these objectives:

- Describe the design considerations for IP multicast implementations in the enterprise
- Design IP multicast services for small enterprise campus networks
- Design IP multicast services for large enterprise campus networks
- Design IP multicast services that operate over a WAN

Learner Skills and Knowledge

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

- CCNP certification curriculum
- Examining IP Multicast Services lesson

Outline

The outline lists the topics included in this lesson.

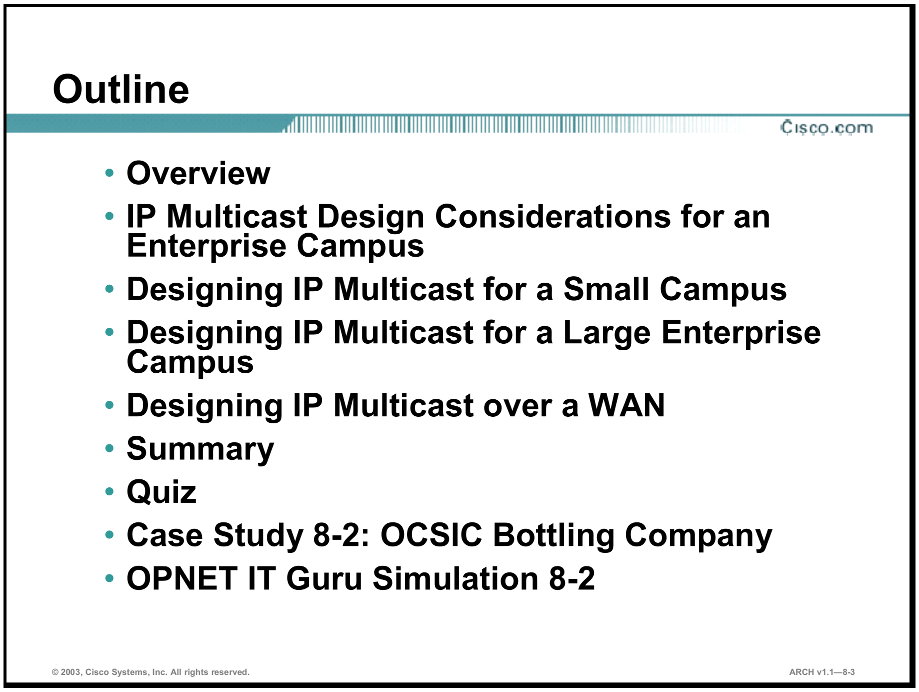

Outline

- Overview
- IP Multicast Design Considerations for an Enterprise Campus
- Designing IP Multicast for a Small Campus
- Designing IP Multicast for a Large Enterprise Campus
- Designing IP Multicast over a WAN
- Summary
- Quiz
- Case Study 8-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 8-2

ARCH v1.1—8-3

IP Multicast Design Considerations for an Enterprise Campus

When designing a network for IP multicast, you will want to consider the servers and hosts, IP multicast control mechanisms, PIM mode, and router provisioning. This topic describes the design considerations for IP multicast implementations in the enterprise.

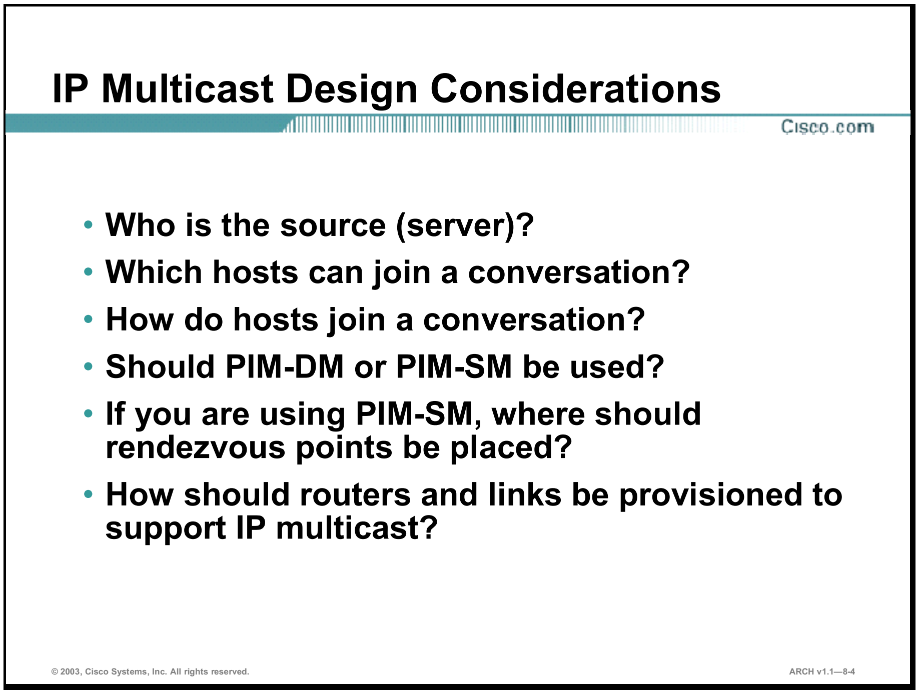

IP Multicast Design Considerations

- Who is the source (server)?
- Which hosts can join a conversation?
- How do hosts join a conversation?
- Should PIM-DM or PIM-SM be used?
- If you are using PIM-SM, where should rendezvous points be placed?
- How should routers and links be provisioned to support IP multicast?

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-4

The questions to ask when designing a network for IP multicast are:

- **Who is the source (server)?** For each IP multicast application, you must identify the source or server, such as a Cisco IP/TV server.
- **How do hosts join a conversation?** To avoid Layer 2 flooding of multicast traffic, implement either IGMP snooping or CGMP. IGMP snooping is best in devices that provide application-specific integrated circuit (ASIC) hardware support for this purpose. Without hardware support, snooping can seriously degrade device performance, and CGMP should be implemented instead. Verify the IP multicast support for each platform prior to design and implementation.
- **Should you use PIM dense mode or sparse mode?** These guidelines can help you select the correct PIM mode:
  - PIM dense mode (PIM-DM) is easy to configure and provides a simple flood-and-prune mechanism. However, it causes inefficient flood-and-prune behavior and does not support shared trees.
  - PIM sparse mode (PIM-SM) is very efficient, using an explicit-join model, and flowing traffic only where it is needed. Cisco recommends that you use PIM-SM whenever possible.
- **If you are using PIM-SM, where should rendezvous points be placed?** Place rendezvous points close to the source of the multicast traffic. Determine the routers that will be used as RPs.

- **How should routers be provisioned to support IP multicast?** Links that support IP multicast need sufficient bandwidth. Similarly, each router that supports IP multicast needs sufficient processing power. The IP multicast routing table stores an entry for each source and host. As the number of sources and hosts grows, the memory requirement for the router grows also.

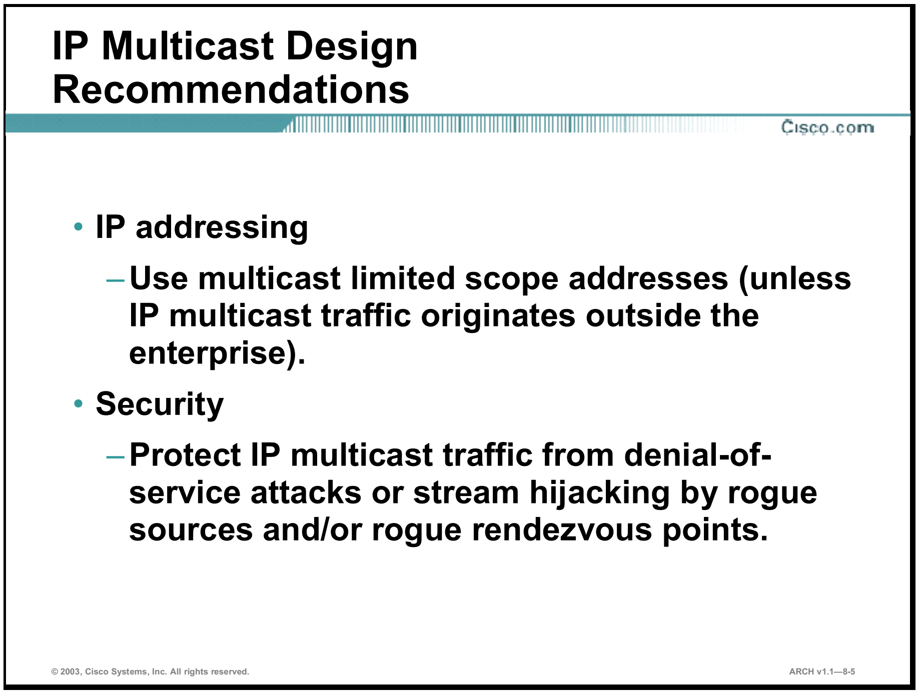

IP Multicast Design Recommendations

- IP addressing
  - Use multicast limited scope addresses (unless IP multicast traffic originates outside the enterprise).
- Security
  - Protect IP multicast traffic from denial-of-service attacks or stream hijacking by rogue sources and/or rogue rendezvous points.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-5

To implement IP multicast efficiently in enterprise networks, you must first ensure that all data link layer switches are able to constrain multicast flooding. Second, the network must support a sparse-mode multicast routing protocol. Clients and servers must have an IP protocol stack-supporting multicast as specified in the Internet RFC 1112 or 2236 standard, and enterprise applications must support IP multicast. All modern operating systems support IGMP and IP multicast. All IOS software-based platforms support multicast, including the Cisco Catalyst family of switches and routers.

While planning your multicast implementation, consider:

- **IP addressing:** Unless IP multicast traffic will be originating from or sent outside the enterprise, you should use the multicast limited scope addresses. You can aid network management by subdividing the available addresses into separate ranges for different purposes. For example, you can use a different range of addresses for each multicast application.
- **Security:** Security is an important topic for all areas of network design. Consider protecting your IP multicast traffic from denial-of-service attacks or stream hijacking by rogue sources and/or rogue RPs. You can configure your RPs to accept source registrations only from a defined access control list (ACL) or route map, eliminating rogue multicast sources, and can filter the source of RP announcements, eliminating rogue RPs.

Designing IP Multicast for a Small Campus

In a small campus design with a collapsed core and distribution layer, the backbone switches act as the rendezvous points for multicast forwarding. This topic describes the design of IP multicast services for small Enterprise Campus networks.

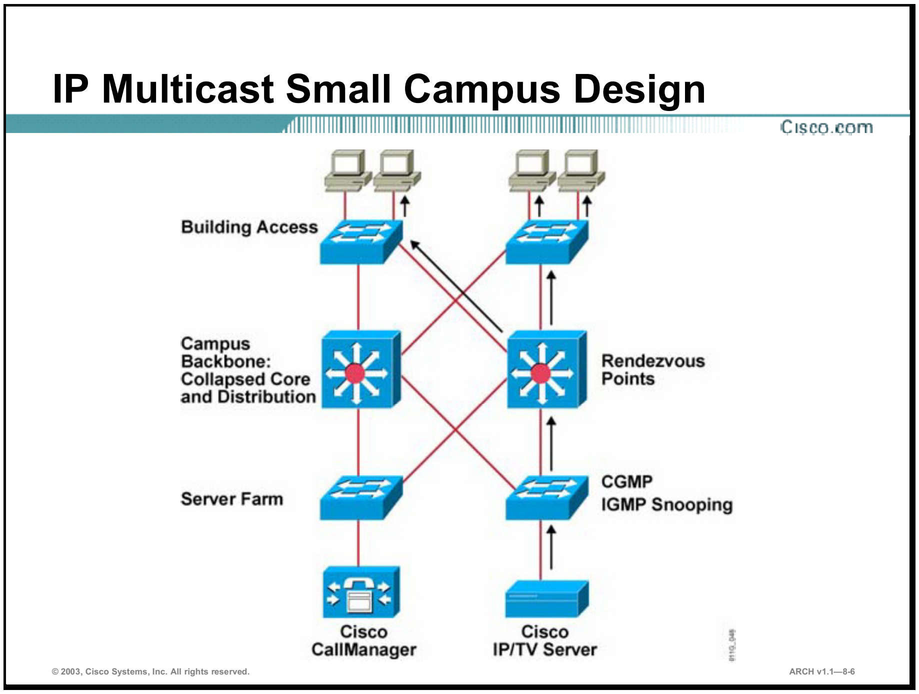

IP Multicast Small Campus Design

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—8-6

In the example, the small campus consists of two small buildings with no more than 200 users and end devices. The network uses a collapsed Campus Backbone and Building Distribution layer referred to as the Campus Backbone. The backbone switches have network layer interfaces to be used for Hot Standby Router Protocol (HSRP). The multicast sources are Cisco CallManager with music on hold and a Cisco IP/TV server attached through the Server Farm module.

As an IP multicast control mechanism**,** IGMP snooping is enabled on all switches, except for those switches that do not provide hardware support for IGMP snooping. CGMP will be implemented on those devices.

The two backbone switches act as multicast RPs.

The Campus Backbone would not normally do anything other than transmit packets. However, it must take over the function of the Building Distribution submodule when you implement a collapsed backbone.

Designing IP Multicast for a Large Enterprise Campus

IP multicast design for a large campus needs to balance between granular administrative control and simplicity. This topic describes the design of IP multicast services for large Enterprise Campus networks.

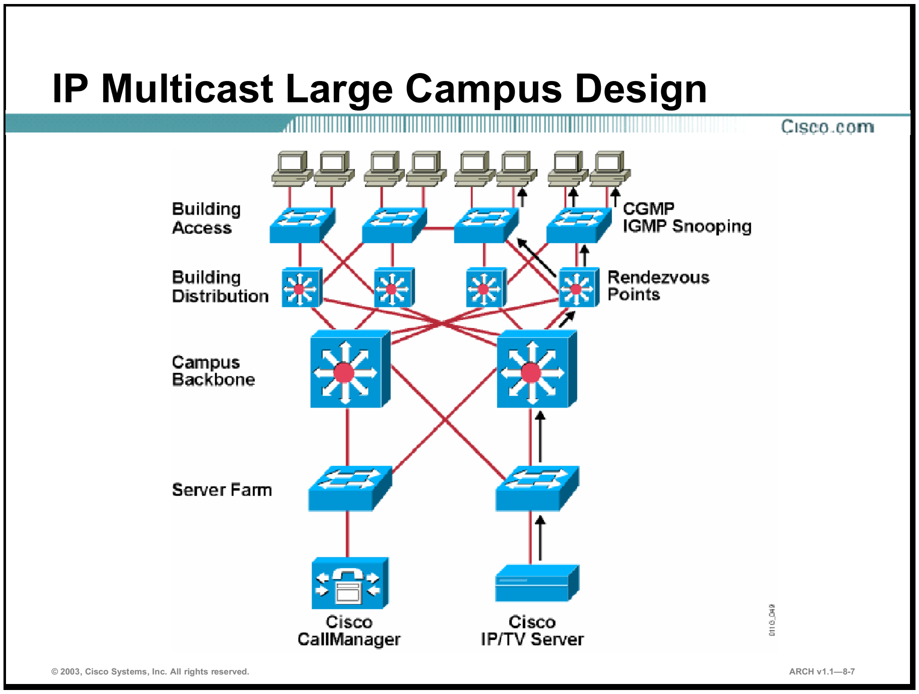

IP Multicast Large Campus Design

ARCH v1.1—8-7

Multicast design in a large enterprise network can be difficult to administer if the design is too granular. For example, you could place redundant RPs throughout the network, each having responsibility for a specific multicast address range. The optimal design is one that provides fast failover, proper RP placement, and a simplistic approach to traffic control. Although there are many possible combinations of multicast deployment in a large campus, and even more combinations for each type of multicast application, it is generally best to focus on keeping administration simple and the traffic reliable.

In this large campus network design, each access switch in the Building Access submodule and the Server Farm module is dual-connected to a pair of distribution routers running HSRP. For multicast, one of the two routers is the designated router and the other is the IGMP querier. The IP unicast routing protocol is configured such that the trunk from the access switch to the DR is always preferred, forcing unicast and multicast paths to be the same.

The Building Access submodule uses switches that support IGMP snooping and good port density to serve the end stations. The Building Distribution and Campus Backbone modules both use switches that have good multicast forwarding performance, can support large multicast routing tables, and are able to house multiple-gigabit links. The Server Farm module uses switches that support IGMP snooping. The Campus Backbone switches need to support a dense population of gigabit ports for connectivity to the Building Distribution submodule and other Campus Backbone switches.

The multicast applications in this design use an architecture with few sources to many receivers. The sources are located in the Server Farm module. The Campus Backbone switches are designated as the multicast RPs. (An alternate choice would be to have the Server Farm access switches serve this role.)

Designing IP Multicast over a WAN

When IP multicast traffic is to cross a WAN link, the primary consideration is that the WAN bandwidth not be overwhelmed by unnecessary traffic. This topic describes the design of IP multicast services over a WAN.

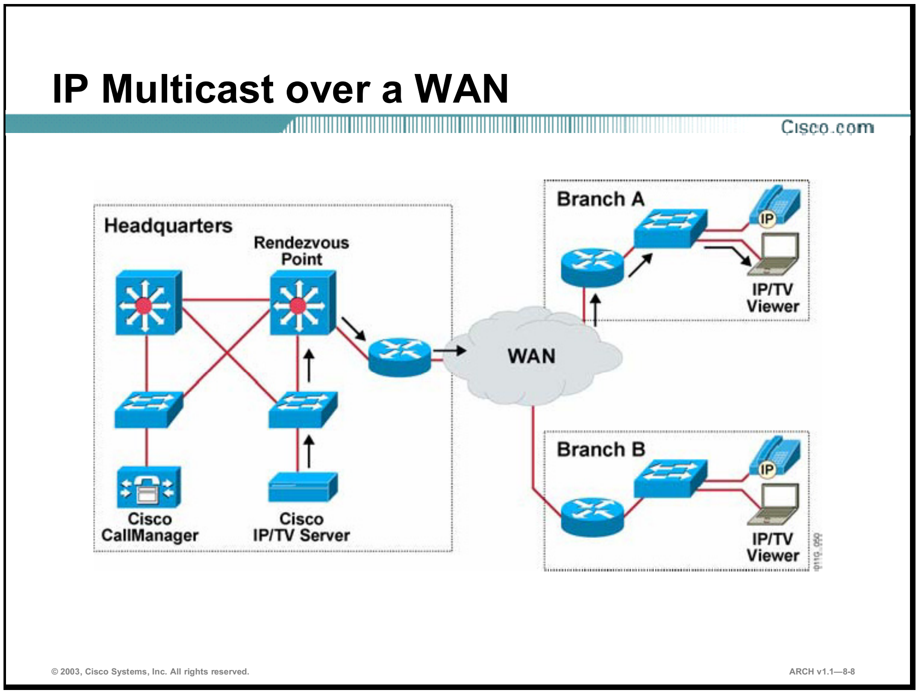

IP Multicast over a WAN

ARCH v1.1—8-8

The design decisions for the headquarters network in this example are the same as in any large campus design: place the RPs close to the multicast source and constrain Layer 2 flooding with IGMP snooping or CGMP.

The primary focus in this design is the filtering of multicast traffic at the WAN access router. The IP addressing scheme will facilitate the design goals.

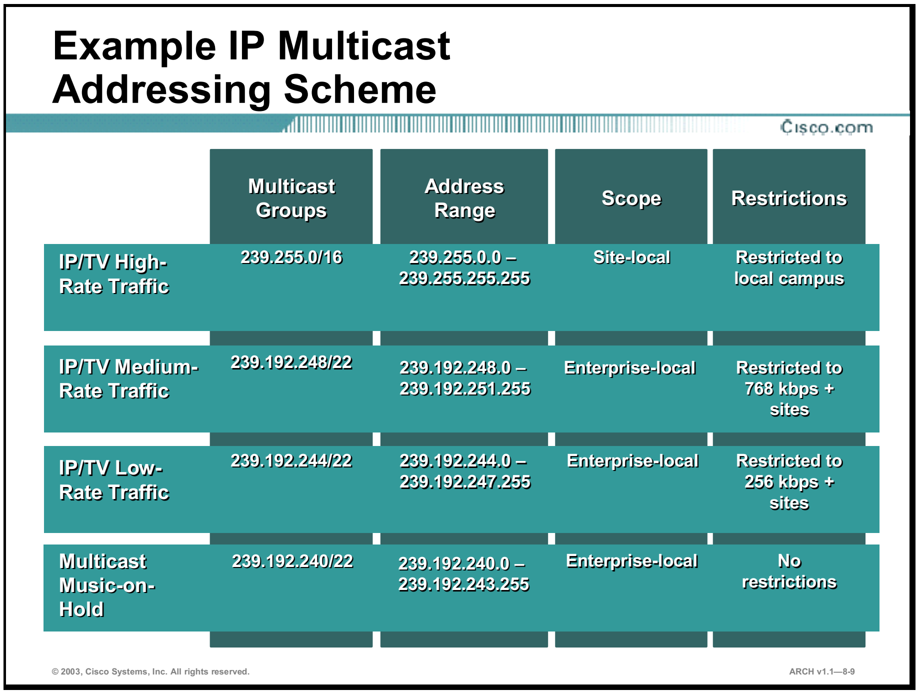

Example IP Multicast Addressing Scheme

Cisco.com

	Multicast Groups	Address Range	Scope	Restrictions
IP/TV High-Rate Traffic	239.255.0/16	239.255.0.0 – 239.255.255.255	Site-local	Restricted to local campus
IP/TV Medium-Rate Traffic	239.192.248/22	239.192.248.0 – 239.192.251.255	Enterprise-local	Restricted to 768 kbps + sites
IP/TV Low-Rate Traffic	239.192.244/22	239.192.244.0 – 239.192.247.255	Enterprise-local	Restricted to 256 kbps + sites
Multicast Music-on-Hold	239.192.240/22	239.192.240.0 – 239.192.243.255	Enterprise-local	No restrictions

ARCH v1.1—8-9

As shown in the table, you can assign multicast group address ranges based on bandwidth requirements for particular types of multicast traffic. With this scheme in place, you can configure the WAN access router with filters and multicast boundaries to ensure that multicast traffic is only forwarded to sites over links that support the necessary bandwidth.

Summary

This topic summarizes the key points discussed in this lesson.

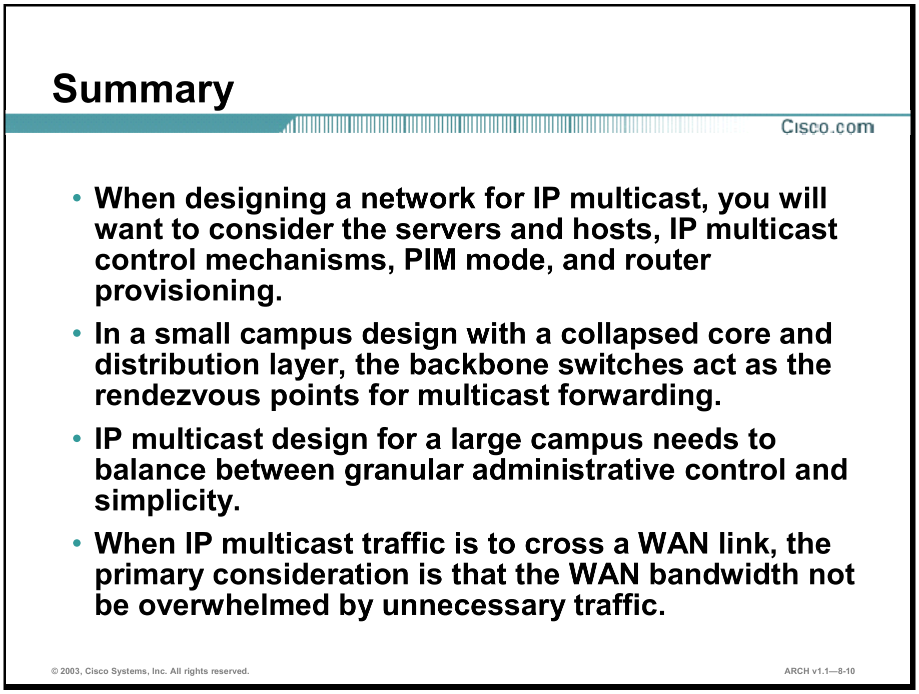

Summary

Cisco.com

- When designing a network for IP multicast, you will want to consider the servers and hosts, IP multicast control mechanisms, PIM mode, and router provisioning.
- In a small campus design with a collapsed core and distribution layer, the backbone switches act as the rendezvous points for multicast forwarding.
- IP multicast design for a large campus needs to balance between granular administrative control and simplicity.
- When IP multicast traffic is to cross a WAN link, the primary consideration is that the WAN bandwidth not be overwhelmed by unnecessary traffic.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—8-10

References

For additional information, refer to these resources:

- *Multicast Services* at http://www.cisco.com/warp/public/732/Tech/multicast/index.shtml
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 8-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 8-2

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) Under what circumstance should you implement CGMP instead of IGMP snooping?

A) when the switch is located in a very small network

B) when the switch provides hardware support for CGMP

C) when Layer 2 flooding of multicast traffic is not an issue

D) when the switch does not provide hardware support for IGMP snooping

Q2) What devices best serve as rendezvous points in a small campus network with a collapsed backbone layer?

A) edge switches

B) backbone switches

C) Server Farm switches

D) access layer switches

Q3) What is a possible drawback of a large IP multicast design that is extremely granular?

A) increased latency

B) administrative complexity

C) degraded network performance

D) increased bandwidth consumption

Q4) Which two basic considerations should you take into account when designing a multicast deployment? (Choose two.)

A) reliable traffic

B) simple administration

C) location of where to deploy RPF

D) IP addressing summarization points

E) switches that employ IGMP snooping

Q5) How can a good IP addressing scheme support an IP multicast WAN solution?

A) It can decrease administrative complexity.

B) It can enable filtering at the WAN access router.

C) It can reduce bandwidth in the Building Access submodule.

D) It can enable multicast boundaries at the Server Farm access layer.

Quiz Answer Key

Q1) D

Relates to: IP Multicast Design Considerations for an Enterprise Campus

Q2) B

Relates to: Designing IP Multicast for a Small Campus

Q3) B

Relates to: Designing IP Multicast for a Large Enterprise Campus

Q4) A, B

Relates to: Designing IP Multicast for a Large Enterprise Campus

Q5) B

Relates to: Designing IP Multicast over a WAN

Case Study 8-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

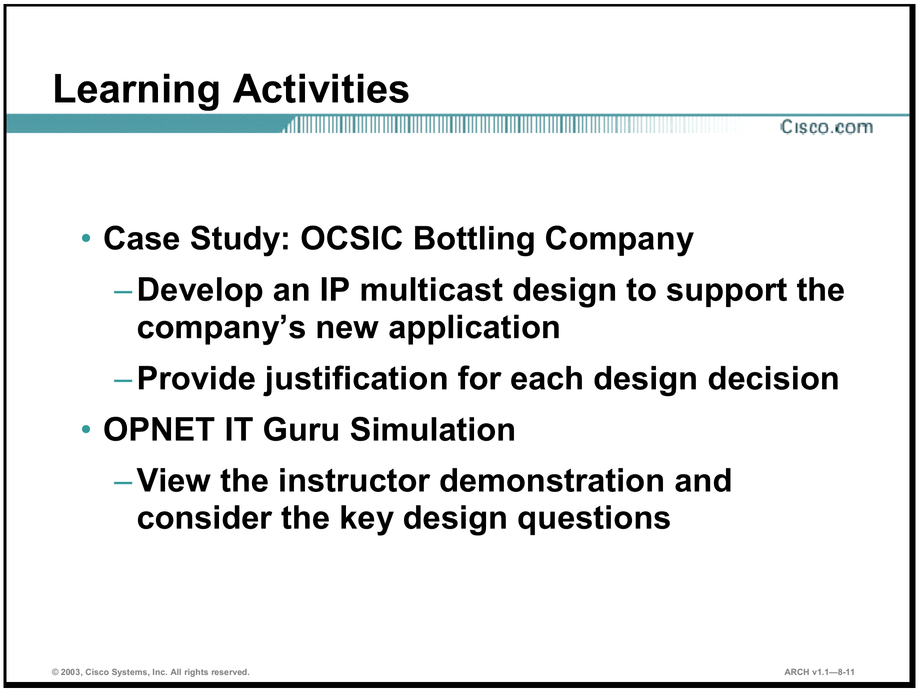

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Develop an IP multicast design to support the company’s new application
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

ARCH v1.1—8-11

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

The OCSIC Bottling Company headquarters is sending massive updates to the North American plants for new-product kickoffs. The kickoff events contain a live video feed. The updates are usually saved so people can view the video at their leisure.

In this exercise, you will design IP multicast services that meet the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Develop an IP multicast design to support the company’s new application

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Develop an IP Multicast Design

Complete these steps:

Step 1 Complete the table to design your IP multicast solution.

Design Questions	Decision	Justification
Where will you implement multicast applications on the network?		
What IP multicast control mechanism will you use?		
Will you use PIM-DM or PIM-SM? If you selected PIM-SM, determine the location of rendezvous points.		
What security is needed to support IP multicasting on the network?		
Which network devices require upgrade (memory, IOS version, new hardware) to support the IP multicast applications?		

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have developed an IP multicast strategy that identifies the method hosts use to join a conversation, the IP multicast control mechanism, PIM mode, rendezvous points (if required), security options, and router provisioning.

OPNET IT Guru Simulation 8-2

This simulation demonstrates the benefits of deploying IP multicast applications over unicast applications.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 9

Designing Virtual Private Networks

Overview

Virtual Private Networks (VPNs) are networks deployed on a public or private network infrastructure. VPNs are useful for telecommuters, mobile users, and remote offices as well as for customers, suppliers, and partners.

For enterprises, VPNs are an alternative WAN infrastructure, replacing or augmenting existing private networks that utilize dedicated WANs based on leased-line, Frame Relay, ATM, or other technologies. Increasingly, enterprises are turning to their service providers for VPNs and other complete service solutions tailored to their particular business.

Module Objectives

Upon completing this module, you will be able to design enterprise solutions for virtual private networks, given enterprise network needs.

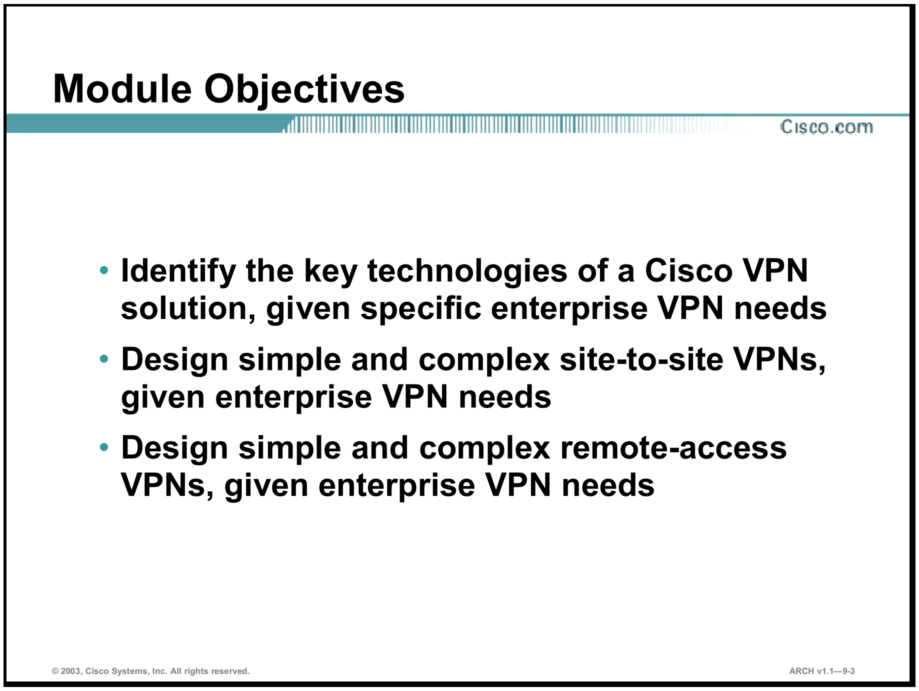

Module Objectives

Cisco.com

- Identify the key technologies of a Cisco VPN solution, given specific enterprise VPN needs
- Design simple and complex site-to-site VPNs, given enterprise VPN needs
- Design simple and complex remote-access VPNs, given enterprise VPN needs

ARCH v1.1—9-3

Module Outline

The outline lists the components of this module.

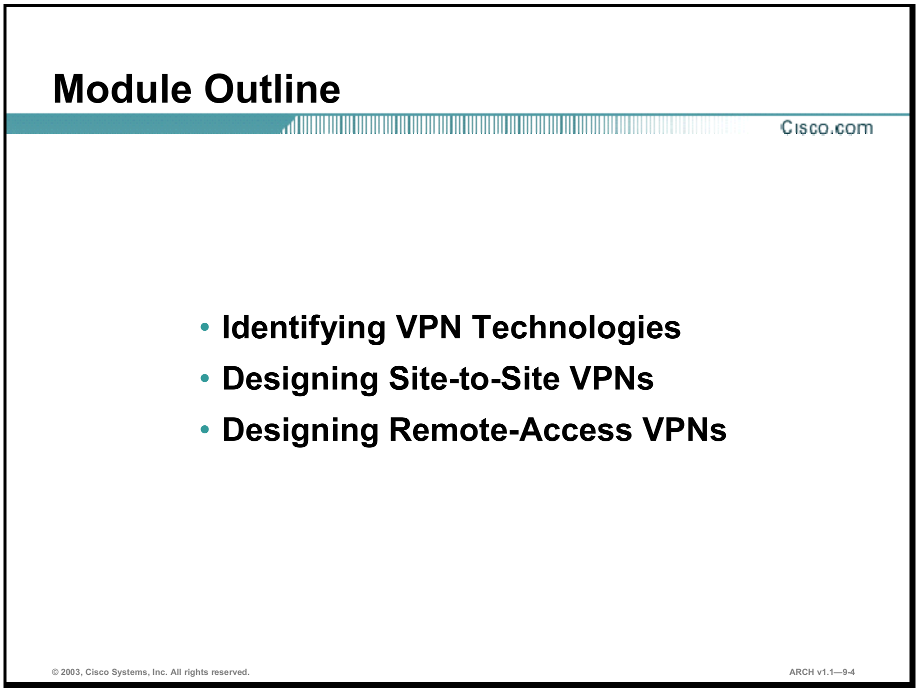

Module Outline

Cisco.com

- Identifying VPN Technologies
- Designing Site-to-Site VPNs
- Designing Remote-Access VPNs

ARCH v1.1—9-4

Identifying VPN Technologies

Overview

VPNs use advanced encryption and tunneling to permit organizations to establish secure, end-to-end, private network connections over third-party networks, such as the Internet or extranets. VPNs are included in Cisco Architecture for Voice, Video and Integrated Data (AVVID), the enterprise architecture that provides an intelligent network infrastructure for today’s Internet business solutions.

Relevance

VPNs do not inherently change private WAN requirements, such as support for multiple protocols, high reliability, and extensive scalability. They often meet these requirements cost-effectively and with great flexibility.

Objectives

Upon completing this lesson, you will be able to identify the key technologies of a Cisco VPN solution, given specific enterprise VPN needs. This includes being able to meet these objectives:

- Identify enterprise security requirements for site-to-site and remote-access VPNs
- Describe the tunneling technology that enables VPNs
- Describe the primary security technologies that enable VPNs
- Identify the necessary termination components of a VPN solution, given specific enterprise VPN requirements
- Identify the necessary VPN management components of a Cisco VPN solution, given specific enterprise VPN requirements

Learner Skills and Knowledge

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

- Designing Enterprise Edge Connectivity module
- Designing Security Solutions module

Outline

The outline lists the topics included in this lesson.

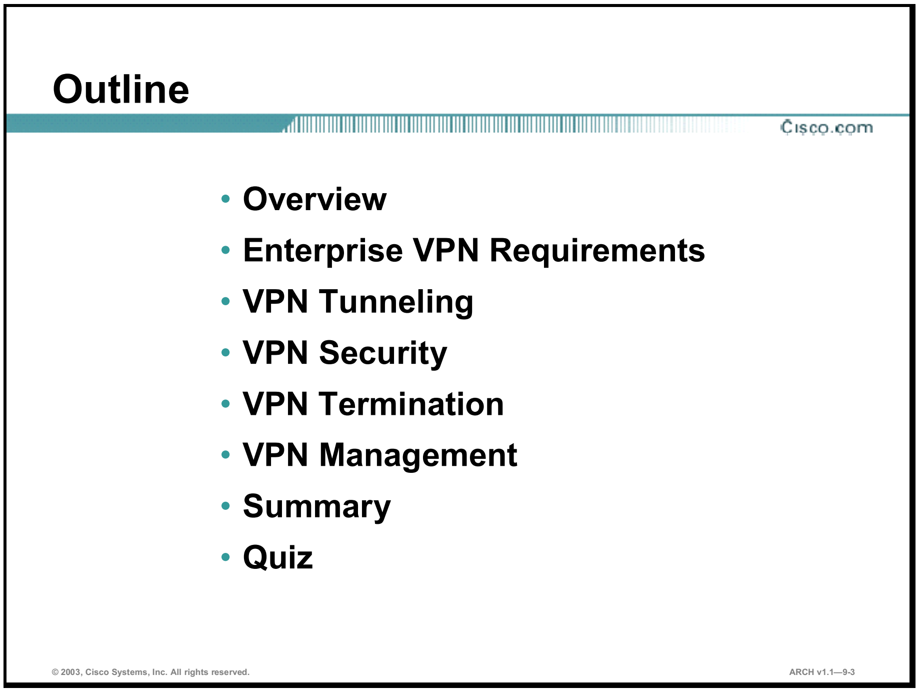

Outline

- Overview
- Enterprise VPN Requirements
- VPN Tunneling
- VPN Security
- VPN Termination
- VPN Management
- Summary
- Quiz

ARCH v1.1—9-3

Enterprise VPN Requirements

VPNs enable network connectivity for an organization, its business partners, and customers over a shared infrastructure, delivering the same policies as a private network. This topic identifies enterprise requirements for site-to-site and remote-access VPNs.

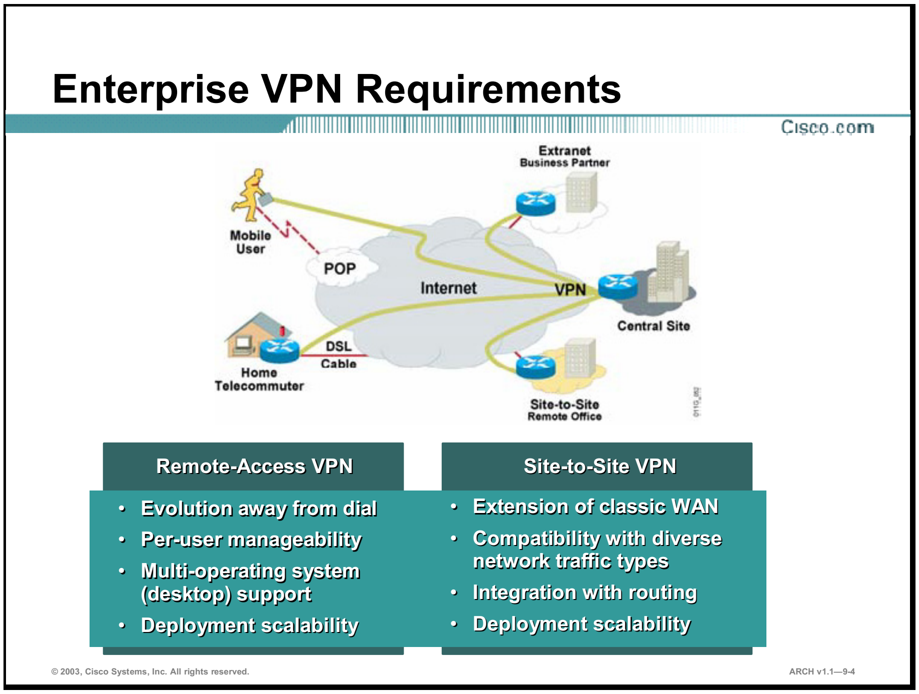

Enterprise VPN Requirements

Cisco.com

Remote-Access VPN

- Evolution away from dial
- Per-user manageability
- Multi-operating system (desktop) support
- Deployment scalability

Site-to-Site VPN

- Extension of classic WAN
- Compatibility with diverse network traffic types
- Integration with routing
- Deployment scalability

ARCH v1.1—9-4

For VPNs, the IP-based infrastructure can be a private enterprise network, an ISP network, or the Internet, providing organizations with flexibility in connectivity and cost. In order to deliver the same policies as a private network, organizations must consider implementing enhanced security, sophisticated traffic engineering, and policy-based management.

Because VPNs are typically new technologies to most enterprise networks, it is important to use a four-phase structured methodology for their implementation. The strategy phase should analyze the business requirements for implementing a VPN. The design phase sets the technical and functional requirements for the VPN. The implementation phase focuses on evolving the existing network and rolling out new network components. Finally, the administrative phase defines the metrics and tools to ensure that the new technology meets the technical and business goals of the project.

VPN applications in enterprise networks are divided into two main categories: Site-to-site and remote access.

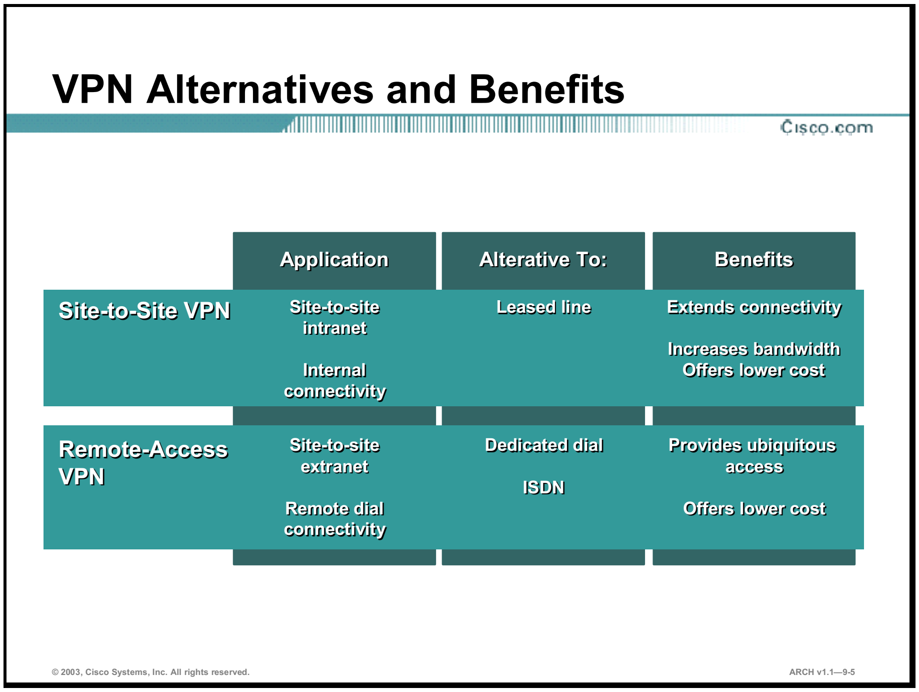

VPN Alternatives and Benefits

	Application	Alterative To:	Benefits
Site-to-Site VPN	Site-to-site intranet Internal connectivity	Leased line	Extends connectivity Increases bandwidth Offers lower cost
Remote-Access VPN	Site-to-site extranet Remote dial connectivity	Dedicated dial ISDN	Provides ubiquitous access Offers lower cost

ARCH v1.1—9-5

Site-to-site VPNs focus on connecting geographically dispersed offices without requiring dedicated circuits. Extranet VPNs, a type of site-to-site VPN, add interconnections between multiple organizations. Remote-access VPNs focus on remote users and partners who access the network on an as-needed basis.

For the large-scale enterprise VPN, particularly one that must support multiple broadband connections, purpose-built communication devices are strongly advised. The required attributes of a VPN solution are:

- Robust architecture
- Scalability
- Easy management
- Flexibility

VPN Tunneling

Virtual point-to-point connectivity is typically provided by a tunneling protocol. This topic describes the tunneling technology that enables VPNs.

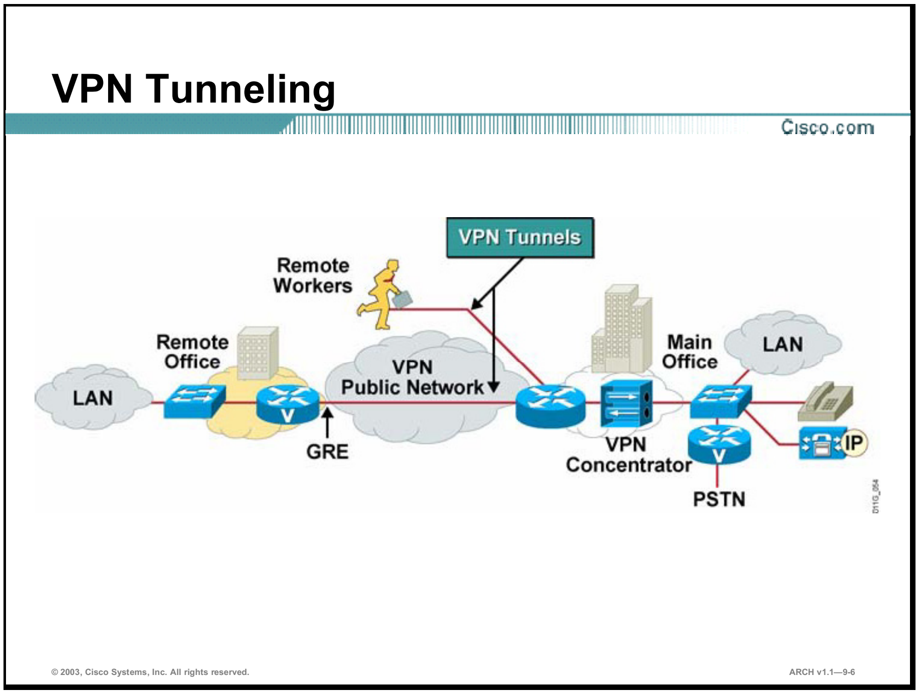

VPN Tunneling

ARCH v1.1—9-6

Tunneling is a technique where packets of one type of protocol are encapsulated by another protocol. Tunneling is often used to transport or route packets across networks of disparate protocols. In the context of VPN, tunneling is used to encapsulate private messages and apply encryption algorithms to the payload.

Several different tunneling protocols have evolved, and each is based on encapsulating a Layer 2 protocol (Layer 2 tunneling) or a Layer 3 protocol (Layer 3 tunneling).

You can implement Layer 3 tunneling in native form or nest it within other tunneling protocols, such as Layer 2 Tunneling Protocol (L2TP). For example, Microsoft Windows 2000 includes a native VPN desktop client that uses L2TP over IP Security (IPSec) as the transport protocol. This combination provides the routing advantages of L2TP with the security of IPSec.

In site-to-site VPNs, the principal tunneling is generic routing encapsulation (GRE). If only IP-unicast packets are being tunneled, simple encapsulation provided by IPSec is sufficient. GRE is used when multicast, broadcast, and non-IP packets need to be tunneled.

VPN Security

VPN security is provided by IPSec, user authentication, and encryption. This topic describes the primary security technologies that enable VPNs.

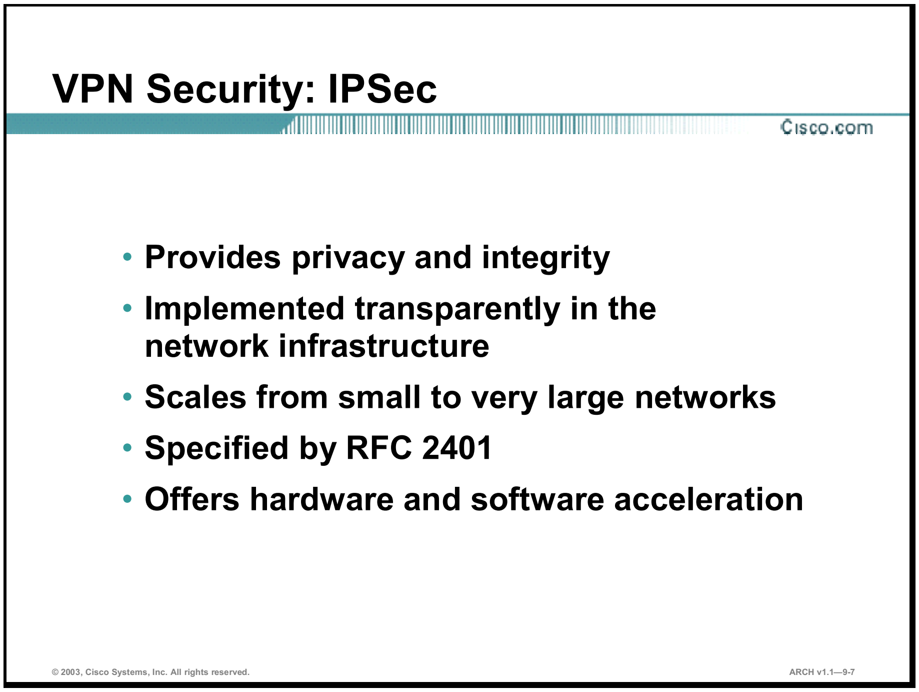

VPN Security: IPSec

- Provides privacy and integrity
- Implemented transparently in the network infrastructure
- Scales from small to very large networks
- Specified by RFC 2401
- Offers hardware and software acceleration

ARCH v1.1—9-7

A VPN solution based on an IP network must provide secure services. Without security measures, standard IP packets are susceptible to numerous security breaches. For example, individual packets or streams of packets may be intercepted in transit and modified by a third party without the knowledge of the user or application. Source and destination addresses may be changed. TCP/IP does not provide any methods of securing a data stream.

IPSec is a set of standards that specify various options for providing VPN data privacy. Packet authentication protects the information flow from being tampered with or even repeated, thereby minimizing disruption to the communication.

An IPSec networking architecture has been specifically designed to address these issues. The framework set forth by the IPSec working group, Internet Engineering Task Force (IETF), provides data integrity checks for tamper detection, source address verification, and data privacy for the packet data and data path.

Today, IPSec is widely accepted as a robust and secure method of securing IP traffic. A range of services are performed in IPSec-compliant implementations, including key management, data encryption, and data authentication.

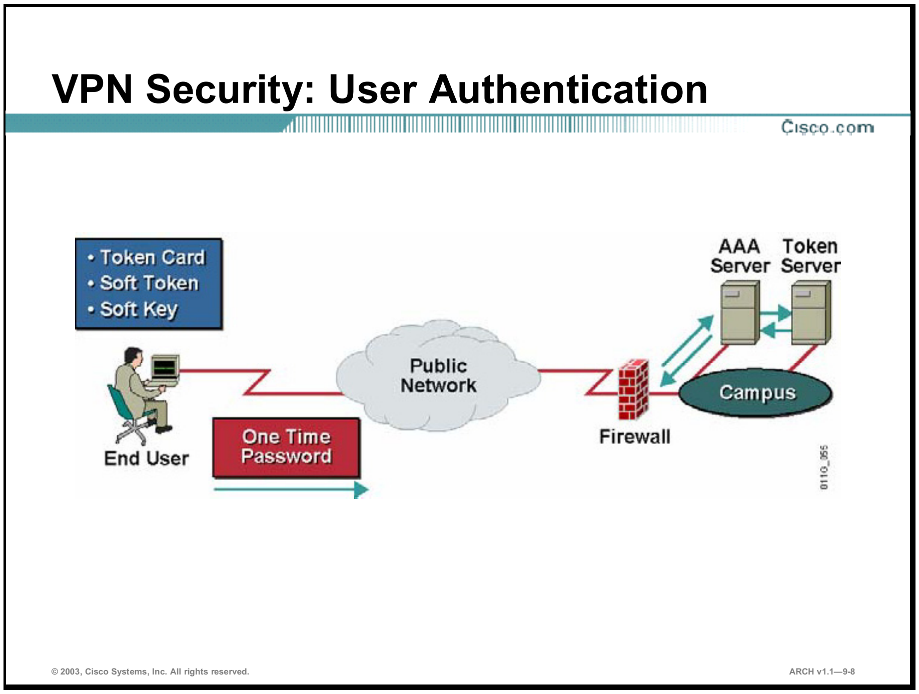

VPN Security: User Authentication

Cisco.com

- Token Card
- Soft Token
- Soft Key

End User

One Time Password

Public Network

Firewall

AAA Server

Token Server

Campus

0110_055

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-8

VPN implementations cannot afford to ignore security threats. One of the primary reasons VPNs are popular is the ubiquity and easy access afforded by IP services. This results in security threats, which must be guarded against. In many cases, enterprises that are migrating to VPN from a dial-up remote-access infrastructure want to use their existing authentication servers such as RADIUS Microsoft Windows NT domain authentication, and token-based security servers.

In cases where a selected VPN device does not support the authentication database of choice, you can use RADIUS proxy. Authentication requests are sent from the VPN device through the VPN gateway to the RADIUS server, which communicates with the actual authentication server. In some cases, both the RADIUS server executable and the authentication database may reside on a single server. Once the proxy service is defined, the authentication process is seamless to the remote-access user.

Public key infrastructure (PKI) can be used for authenticating remote-access users and generating keys used for Internet Key Exchange (IKE) negotiations. X.509 digital certificates are used to exchange key information for each remote-access user. The PKI provides a suite of security services for distribution, management, and revocation of digital certificates. A certificate authority (CA) digitally signs each certificate and validates the integrity of the certificate.

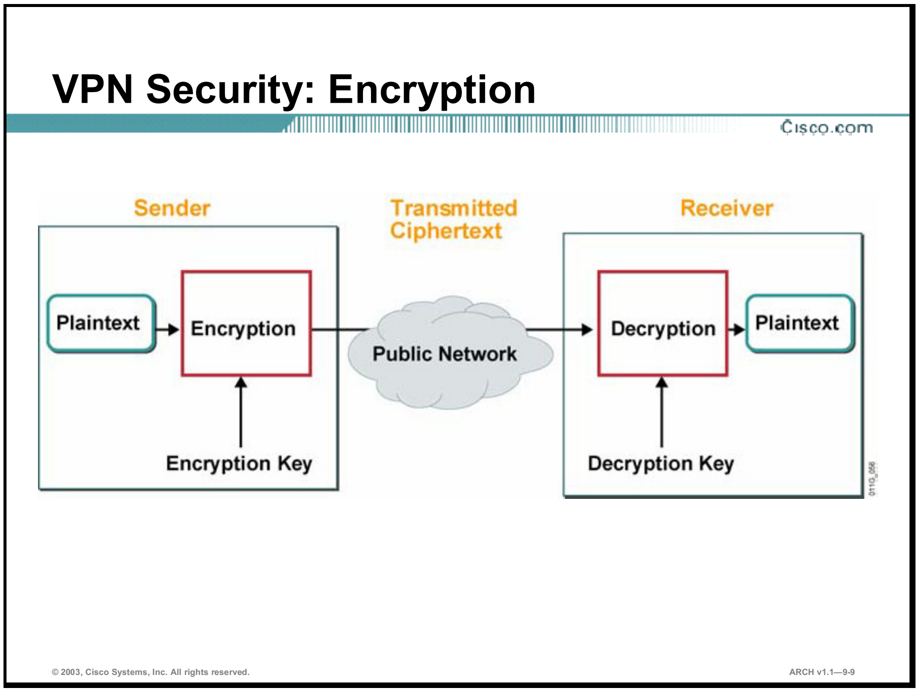

VPN Security: Encryption

Cisco.com

Sender

Plaintext

Encryption

Encryption Key

Transmitted Ciphertext

Public Network

Receiver

Decryption

Plaintext

Decryption Key

ARCH v1.1—9-9

Encryption is the process of converting data through an algorithm to a form that cannot be easily read. The resultant encrypted data stream cannot be read without decrypting the contents of the packets that transport the data.

Data encryption algorithms are generally configurable. Today, Triple Data Encryption Standard (3DES) is widely used. It is considered secure enough for enterprise deployments and sensitive data applications. The National Institute of Standards and Technology has endorsed 3DES since 1975. Data authentication protocols are also configurable, the most popular today being the Secure Hash Algorithm 1 (SHA-1) and Message Digest 5 (MD5). Key exchanges identify each party in the remote-access session. A popular, automated method of key exchange is called IKE, or Internet Key Exchange (formerly known as ISAKMP/Oakley). Advanced Encryption Standard (AES) is also gaining in popularity in addition to Data Encryption Standard (DES) and 3DES. Cisco is now using AES as well in VPN devices.

VPN Termination

A wide variety of VPN products are on the market. For high availability and scalability in enterprise applications, it is best to seek out a dedicated VPN concentrator platform, purpose-built to aggregate a large volume of simultaneous VPN sessions. This topic helps you select the necessary termination components of a VPN solution, given specific enterprise VPN requirements.

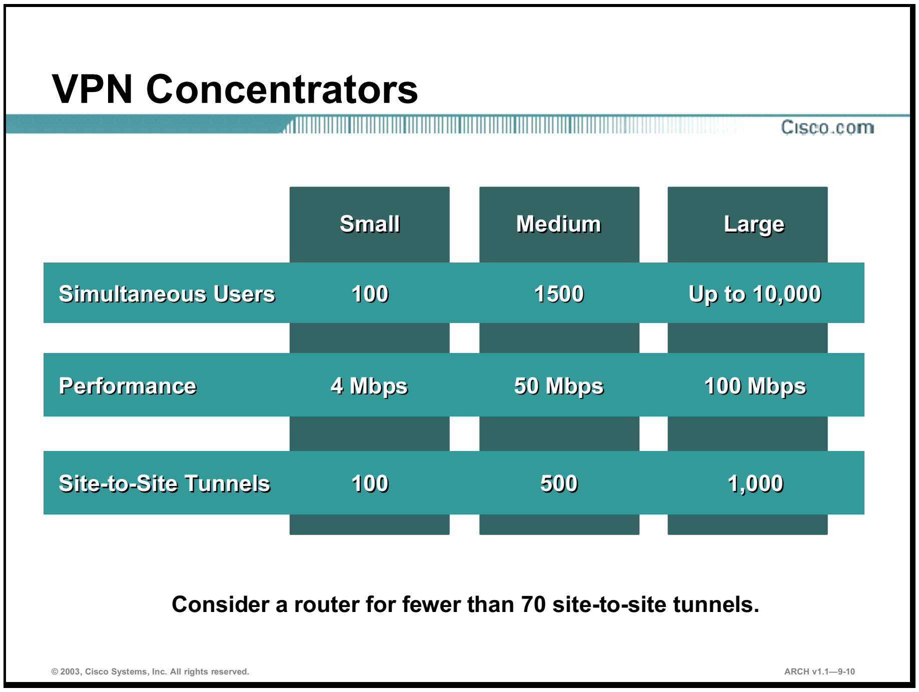

VPN Concentrators

Cisco.com

	Small	Medium	Large
Simultaneous Users	100	1500	Up to 10,000
Performance	4 Mbps	50 Mbps	100 Mbps
Site-to-Site Tunnels	100	500	1,000

Consider a router for fewer than 70 site-to-site tunnels.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-10

When evaluating the capabilities of the concentrator device, answer these questions:

- **Is it a purpose-built communications platform?** Many vendors have attempted to enter the market with a general-purpose PC to build their VPN devices. This approach suffers from the low mean time between failures (MTBF) of consumer-grade PC technology, introduction of failure-prone mechanical components such as disk drives, and lack of scalability to deal with performance requirements. Like routers and switches, the VPN device is part of the communications infrastructure and should be implemented using a purpose-built communications platform.
- **Is physical redundancy designed in?** A high-availability platform should provide fully redundant power, multiple fans, field-swappable components, and easy swapping of a failed unit. Look for a form factor that readily allows rack mounting and hot swapping in case of failure, with a size and weight that allows overnight shipment of replacement parts.
- **Does the product utilize a proven software operating system?** Even if the hardware uses a robust, fault-tolerant design, system uptime can still be compromised by a software architecture that uses proprietary, unproven operating system software. Ideally, the software platform will be based on an industry-standard OS with significant field experience in embedded system products.

- **Can the device be proactively managed for high availability?** The VPN design should include tools that monitor critical components and warn in advance of potential failure conditions. It should provide status monitoring for power supply, fan speed, and internal temperature. System failure from these elements tends to occur over time, so proactive monitoring and reporting can allow intervention before an actual system failure occurs. In addition, the system software should provide automatic recovery in case of a system crash. To ensure high availability, the unit should recover from software failure without external intervention.

To support fewer than 70 individual tunnels, a router would be acceptable. Between 70 and 100 individual tunnels, you could select a router or a VPN concentrator. The main design question to answer is: Do you want all of your WAN and VPN termination through one device, or do you want to split the functionality?

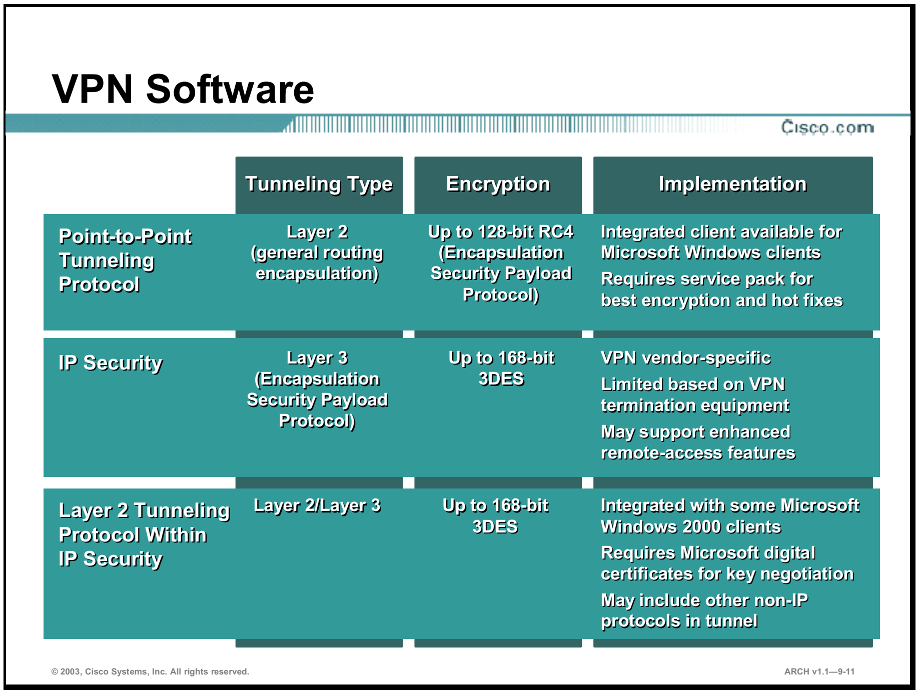

VPN Software

Cisco.com

	Tunneling Type	Encryption	Implementation
Point-to-Point Tunneling Protocol	Layer 2 (general routing encapsulation)	Up to 128-bit RC4 (Encapsulation Security Payload Protocol)	Integrated client available for Microsoft Windows clients Requires service pack for best encryption and hot fixes
IP Security	Layer 3 (Encapsulation Security Payload Protocol)	Up to 168-bit 3DES	VPN vendor-specific Limited based on VPN termination equipment May support enhanced remote-access features
Layer 2 Tunneling Protocol Within IP Security	Layer 2/Layer 3	Up to 168-bit 3DES	Integrated with some Microsoft Windows 2000 clients Requires Microsoft digital certificates for key negotiation May include other non-IP protocols in tunnel

ARCH v1.1—9-11

Remote-access client software runs on the remote desktop and provides the user interface and transparent underlying VPN security protocols for establishing the tunnel to the VPN termination.

The enterprise must weigh the pros and cons of vendor-specific VPN client software implementations versus the embedded clients in Windows 95, Windows 98, Windows NT, and Windows 2000. There are also third-party offerings for MacOS, Solaris, and Linux. Operating system-integrated desktop clients can be easier to deploy, especially if no service packs are required. However, consider carefully any security and compatibility trade-offs with existing remote-access infrastructure (token security, digital certificates, RADIUS).

The client software should be easy to install and simple to operate. Open software distribution licenses provide the network administrator with fixed costs during the buildout and upgrade phases. With this model, spikes in demand do not trigger uncapped incremental expenses.

When evaluating proprietary client software, the development philosophy behind the implementation needs to be understood. Is this a client application that can be modified or tuned by the end user? If so, support calls and misconfigurations may plague a large-scale deployment. Is the client well integrated with a dialer? Does the client support roaming dial services for traveling users? How is the client distributed to end users? Can it be preconfigured for mass deployment? Are there methods for revision control and automatic software updates? Does the client support a larger policy management strategy? How is split tunneling handled? Can users be grouped together and assigned common LAN privileges?

VPN Management

Robust VPN management is another critical success factor for any large remote-access deployment. Management features can be grouped into broad categories such as configuration, monitoring, and alert functionality. This topic helps you select the necessary VPN management components of a Cisco VPN solution, given specific enterprise VPN requirements.

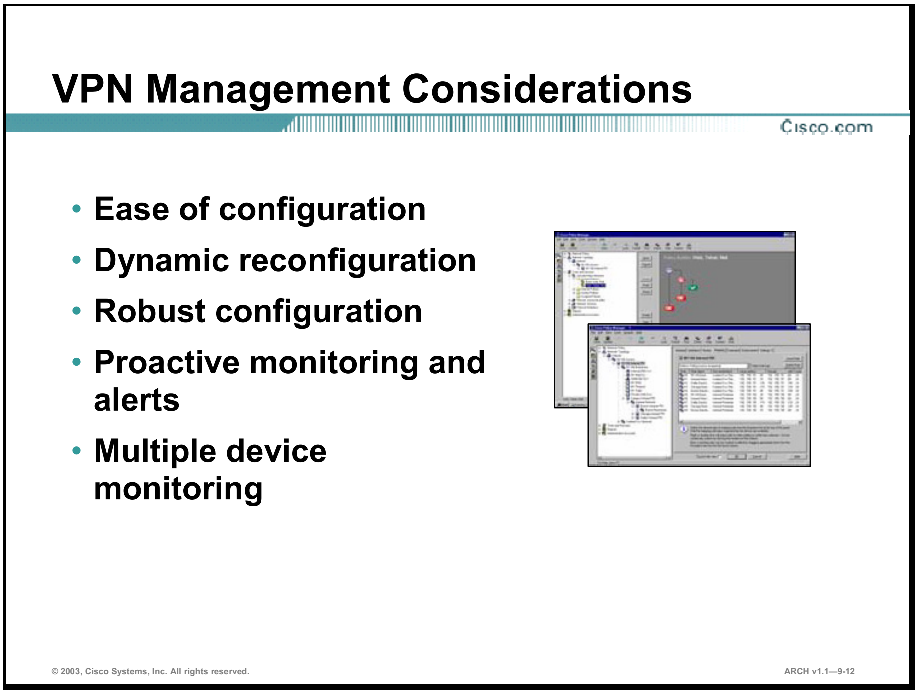

VPN Management Considerations

- Ease of configuration
- Dynamic reconfiguration
- Robust configuration
- Proactive monitoring and alerts
- Multiple device monitoring

ARCH v1.1—9-12

Consider the flexibility of management tools and the varied audiences that may interface with them. In many cases, having both a browser interface and command-line interface (CLI) is beneficial. Consider the following deployment issues in the management context:

- **Ease of configuration:** VPN is a sophisticated technology with many configuration options. A VPN solution should provide an intuitive interface and allow the systems administrator to quickly configure and manage new devices. The management functionality should also provide flexibility for extended configurations so that the VPN may be optimized for specific applications.
- **Dynamic reconfiguration:** All configuration changes should take effect without requiring a reboot of the device. Disruption of service with a fully loaded VPN device can potentially impact thousands of individual users.
- **Robust configuration:** Any time the device is reconfigured or new software is loaded, there is a real possibility of disrupted operation due to operator error (incorrect configuration or faulty download of a software image). It is critical for the unit to be able to check the validity of the configuration and downloaded software, and automatically restore operation to the last known configuration or software image in case of error.
- **Proactive monitoring and alerts:** To ensure high availability, the device must support a wide range of system monitoring (both hardware and software) to constantly monitor operational status. At a minimum, the solution should include tools that can allow rapid isolation, diagnosis, and reporting of faults to allow rapid repair and recovery. Ideally, the system will also incorporate intelligence to identify trends that can predict a potential failure, alerting the system manager to take action before a fault condition occurs.

- **Multiple device monitoring:** In a redundant topology, VPNs need a management tool to allow status viewing of multiple devices as a single service entity. This tool should allow top-level monitoring of overall service operation, as well as automatic highlighting and notification of abnormal operating conditions. Facilities that allow an operator to aggregate and analyze a high volume of data are essential. Archival, graphing, and trend analysis of management status are also critical.

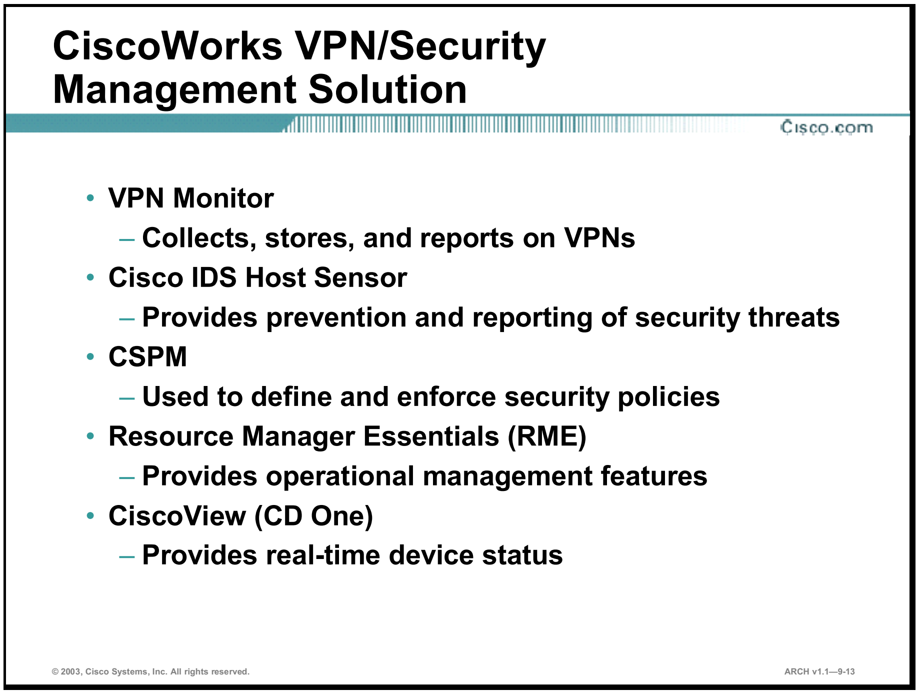

CiscoWorks VPN/Security Management Solution

Cisco.com

- VPN Monitor
  - Collects, stores, and reports on VPNs
- Cisco IDS Host Sensor
  - Provides prevention and reporting of security threats
- CSPM
  - Used to define and enforce security policies
- Resource Manager Essentials (RME)
  - Provides operational management features
- CiscoView (CD One)
  - Provides real-time device status

ARCH v1.1—9-13

CiscoWorks VPN/Security Management Solution offers web-based applications for configuring, monitoring, and troubleshooting enterprise VPNs, firewall security, and network and host-based intrusion detection systems.

The CiscoWorks VPN/Security Management Solution includes these modules:

- **VPN Monitor:** Collects, stores, and reports on IPsec-based site-to-site and remote-access VPNs. VPN Monitor supports the Cisco VPN concentrators and routers.
- **Cisco IDS Host Sensor:** Provides prevention and reporting of security threats to critical servers. Includes both the management console and the evaluation sensor agents. Agents provide protection to operating systems and protection to servers. Agents are purchased separately.
- **Cisco Secure Policy Manager (CSPM):** Used to define and enforce security policies on Cisco PIX firewalls, and to report and alert about intrusions when Cisco Network Intrusion Detection Systems (NIDSs) are deployed. CSPM is also used to define IDS and IOS security policies.
- **Resource Manager Essentials (RME):** Provides the operational management features required by enterprises: software distribution, change audit and authorization, device inventory, credentials management, and Syslog analysis for problem solving and notification of VPN and security operational problems.
- **CiscoView (CD One):** Provides administrators with browser access to real-time device status, and operational and configuration functions.

Summary

This topic summarizes the key points discussed in this lesson.

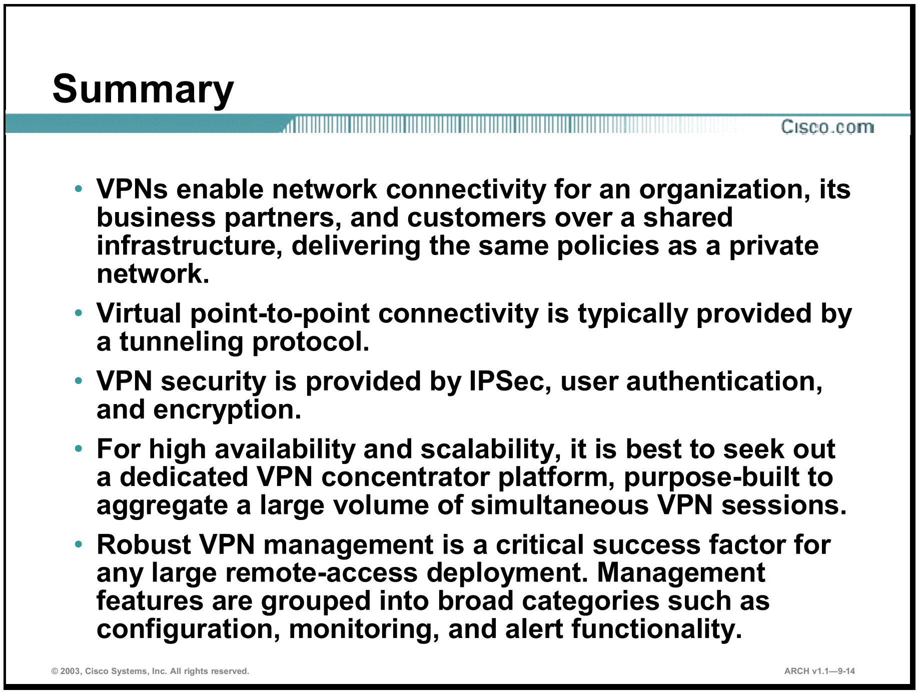

Summary

- VPNs enable network connectivity for an organization, its business partners, and customers over a shared infrastructure, delivering the same policies as a private network.
- Virtual point-to-point connectivity is typically provided by a tunneling protocol.
- VPN security is provided by IPSec, user authentication, and encryption.
- For high availability and scalability, it is best to seek out a dedicated VPN concentrator platform, purpose-built to aggregate a large volume of simultaneous VPN sessions.
- Robust VPN management is a critical success factor for any large remote-access deployment. Management features are grouped into broad categories such as configuration, monitoring, and alert functionality.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—9-14

References

For additional information, refer to these resources:

- *Virtual Private Network Design* at http://www.cisco.com/warp/public/779/largeent/design/vpn.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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 is a type of VPN?

A) Internet VPN

B) external VPN

C) remote-access VPN

D) site-to-Internet VPN

Q2) Which tunneling technique is used when one protocol serves as the transport for another protocol?

A) routed

B) secured

C) switched

D) encapsulated

Q3) Which three functions does IPSec provide? (Choose three.)

A) Layer 3 tunneling

B) user authentication

C) source address verification

D) data integrity checks for tamper detection

E) data privacy for the packet data and data path

Q4) What is the purpose of user authentication?

A) to safeguard the network from local users

B) to keep outsiders from connecting to the inside

C) to keep insiders from connecting to the outside

D) to safeguard the network from unauthorized access

Q5) What is a VPN concentrator?

A) router designed to terminate VPNs

B) switch designed to terminate VPNs

C) multilayer switch designed to terminate VPNs

D) purpose-built device designed to terminate VPNs

Q6) Which VPN client protocol can you implement if your users have Microsoft Windows 2000 workstations?

A) PPTP

B) L2TP

C) ATM

D) PPPOE

Q7) Which module in the CiscoWorks VPN/Security Management Solution collects, stores, and reports on IPSec-based site-to-site and remote-access VPNs?

A) CiscoView

B) VPN Monitor

C) Resource Manager Essentials

D) CiscoWorks Inventory Services

Quiz Answer Key

Q1) C

Relates to: Enterprise VPN Requirements

Q2) D

Relates to: VPN Tunneling

Q3) C, D, E

Relates to: VPN Security

Q4) D

Relates to: VPN Security

Q5) D

Relates to: VPN Termination

Q6) B

Relates to: VPN Termination

Q7) B

Relates to: VPN Management

Designing Site-to-Site VPNs

Overview

Site-to-Site VPNs are an alternative WAN infrastructure used to connect branch offices, home offices, or business partners' sites to all or portions of an enterprise's network. VPNs do not inherently change private WAN requirements, such as support for multiple protocols, high reliability, and extensive scalability, but instead meet these requirements more cost-effectively and with greater flexibility. Site-to-site VPNs utilize the most pervasive transport technologies available today, such as the public Internet or service provider IP networks, by employing tunneling and encryption for data privacy and quality of service (QoS) for transport reliability.

Relevance

Using Internet transport, site-to-site VPNs often reduce WAN costs and can be easily and quickly extended to new locations and business partners. VPNs enable secure use of cost-effective, high-speed links. VPNs encrypt and authenticate traffic traversing the WAN to deliver true network security in an insecure, networked world.

Objectives

Upon completing this lesson, you will be able to design simple and complex site-to-site VPNs, given enterprise VPN needs. This includes being able to meet these objectives:

- Identify typical requirements for an enterprise site-to-site VPN
- List the high-availability and resiliency design considerations for enterprise site-to-site VPNs
- List the routing protocol design considerations for enterprise site-to-site VPNs
- List the packet fragmentation design considerations for enterprise site-to-site VPNs
- Design simple and complex site-to-site VPN solutions, given enterprise VPN requirements

Learner Skills and Knowledge

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

- Designing Enterprise Edge Connectivity module
- Designing Security Solutions module
- Identifying VPN Technologies lesson

Outline

The outline lists the topics included in this lesson.

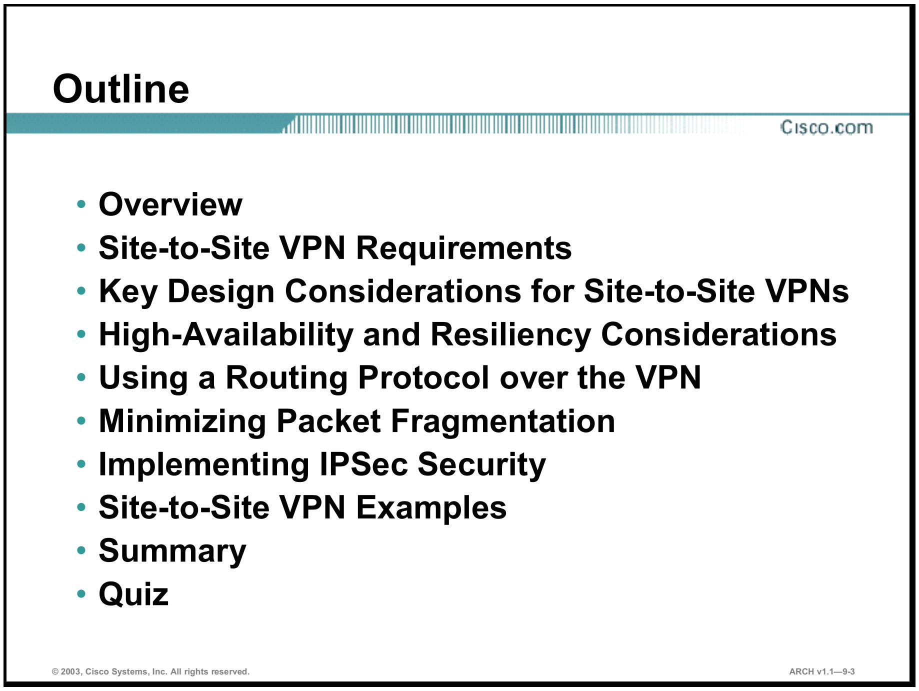

Outline

- **Overview**
- **Site-to-Site VPN Requirements**
- **Key Design Considerations for Site-to-Site VPNs**
- **High-Availability and Resiliency Considerations**
- **Using a Routing Protocol over the VPN**
- **Minimizing Packet Fragmentation**
- **Implementing IPSec Security**
- **Site-to-Site VPN Examples**
- **Summary**
- **Quiz**

ARCH v1.1—9-3

Site-to-Site VPN Requirements

Site-to-site VPNs extend the classic WAN by providing large-scale encryption between multiple fixed sites such as remote offices and central offices, over a shared private or public network, such as the Internet. This topic identifies typical requirements for site-to-site VPNs.

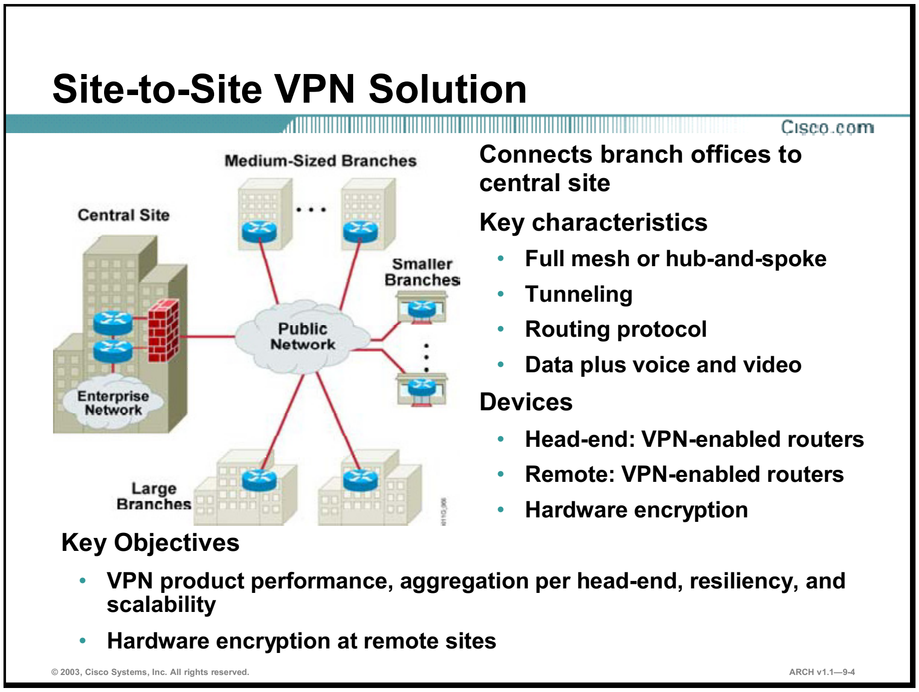

Site-to-Site VPN Solution

Cisco.com

Connects branch offices to central site

Key characteristics

- **Full mesh or hub-and-spoke**
- **Tunneling**
- **Routing protocol**
- **Data plus voice and video**

Devices

- **Head-end: VPN-enabled routers**
- **Remote: VPN-enabled routers**
- **Hardware encryption**

Key Objectives

- **VPN product performance, aggregation per head-end, resiliency, and scalability**
- **Hardware encryption at remote sites**

ARCH v1.1—9-4

Site-to-site VPNs are primarily deployed to connect office locations of an enterprise. They provide an alternative to the WAN infrastructure, while offering significant cost benefits. They enable new infrastructure applications such as extranet, and extend and enhance network connectivity.

Enterprise WAN requirements for traditional private WAN services, such as multiprotocol support, high-availability, scalability, and security, are also requirements for VPNs. VPNs can often meet these requirements more cost-effectively and with greater flexibility than private WAN services using leased lines or virtual circuit technologies such as Frame Relay and ATM.

The key components of site-to-site VPN design include:

- **Cisco head-end VPN routers:** Serve as VPN head-end termination devices at a central campus (head-end devices)
- **Cisco VPN access routers:** Serve as VPN branch-end termination devices at branch office locations (branch-end devices)
- **IPSec and GRE tunnels:** Interconnect the head-end and branch-end devices in the VPN
- **Internet services from ISPs:** Serve as the WAN interconnection medium

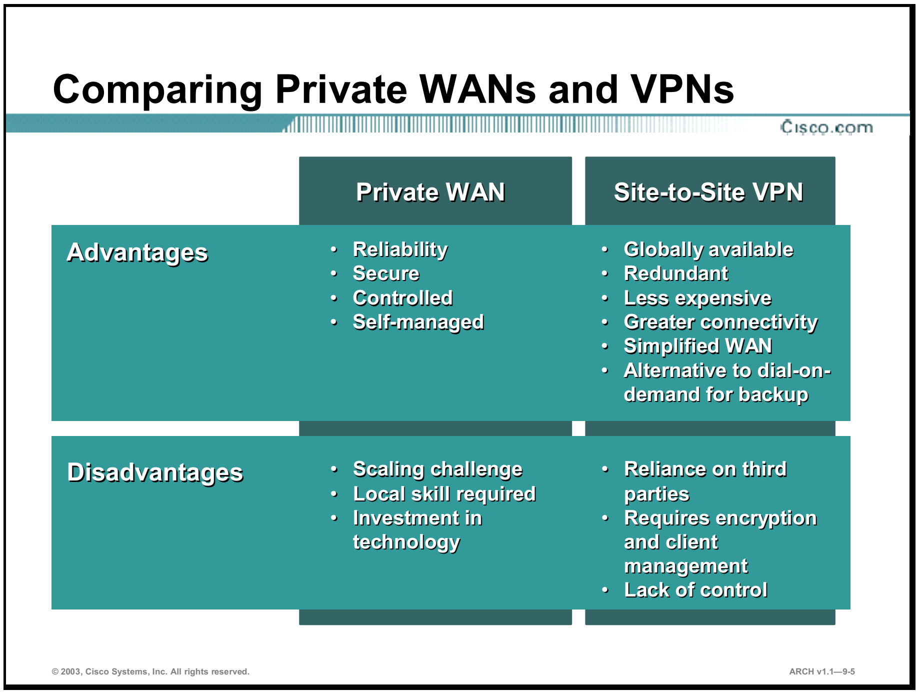

Comparing Private WANs and VPNs

	Private WAN	Site-to-Site VPN
Advantages	• Reliability • Secure • Controlled • Self-managed	• Globally available • Redundant • Less expensive • Greater connectivity • Simplified WAN • Alternative to dial-on-demand for backup
Disadvantages	• Scaling challenge • Local skill required • Investment in technology	• Reliance on third parties • Requires encryption and client management • Lack of control

ARCH v1.1—9-5

The first decision for an enterprise is whether to replace a traditional WAN with a VPN. The enterprise needs to compare the features and benefits of each technology.

The figure describes the advantages and disadvantages associated with traditional private WAN solutions and VPN solutions.

Key Design Considerations for Site-to-Site VPNs

When designing the site-to-site VPN, you need to design the topology, and incorporate resiliency and failover mechanisms. Cisco products support all these design criteria. This topic lists the design considerations for enterprise site-to-site VPNs.

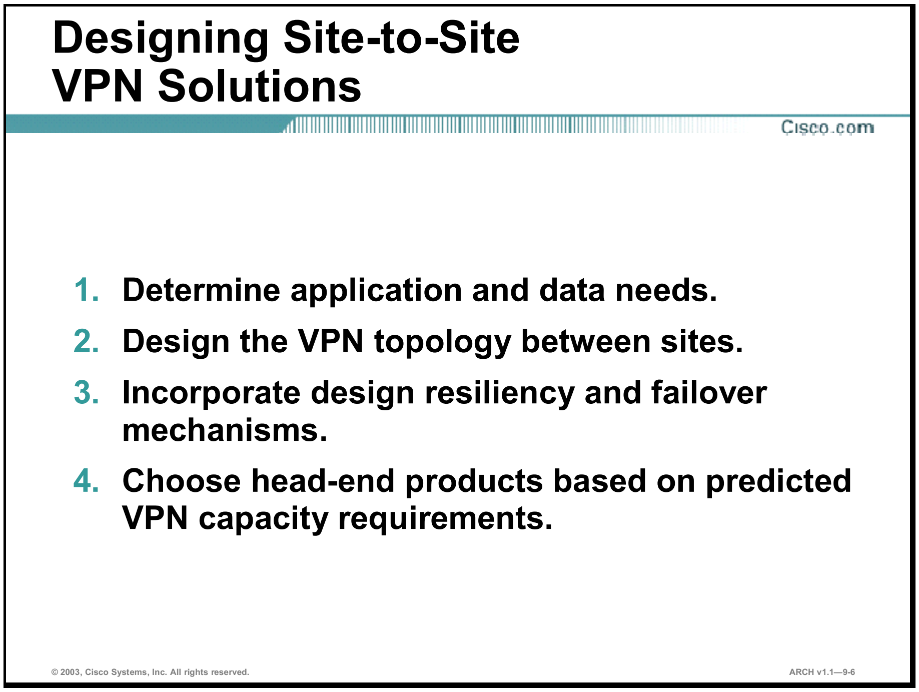

Designing Site-to-Site VPN Solutions

Cisco.com

1. Determine application and data needs.
2. Design the VPN topology between sites.
3. Incorporate design resiliency and failover mechanisms.
4. Choose head-end products based on predicted VPN capacity requirements.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—9-6

The four major steps to design a site-to-site VPN solution are:

Step 1 Determine the application and data needs for the VPN solution.

Step 2 Design the VPN topology between sites.
These types of site-to-site VPNs are deployed most often:

- Hub-and-spoke topology; Multiple hub-and-spoke topology
- Mesh topology
- Hierarchical network topology

Step 3 Incorporate design resiliency and failover mechanisms.

Step 4 Choose head-end products based on predicted VPN capacity requirements.

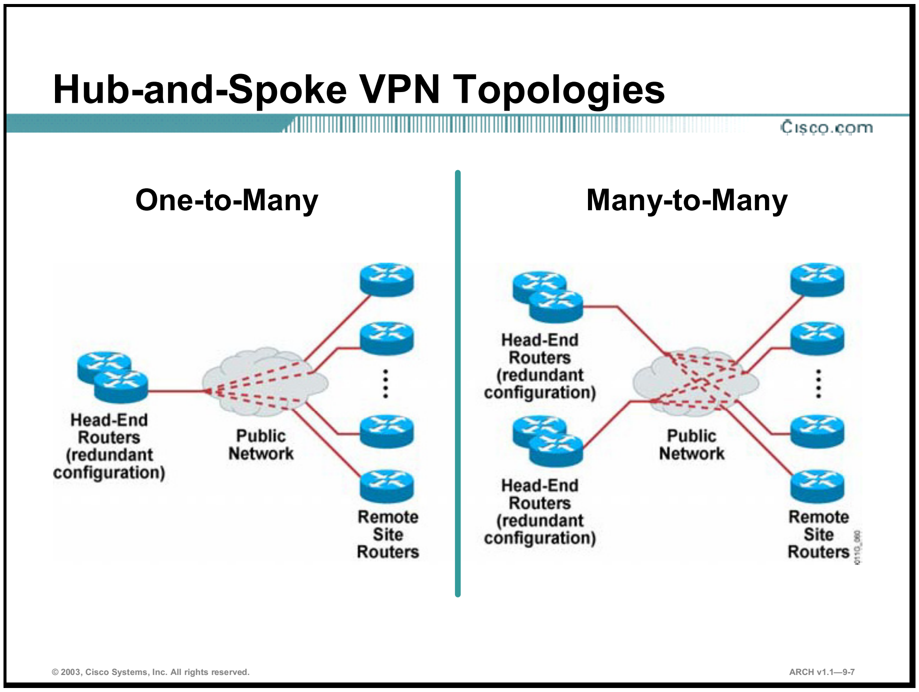

Hub-and-Spoke VPN Topologies

Cisco.com

One-to-Many

Many-to-Many

Head-End Routers (redundant configuration)
Public Network
Remote Site Routers
Head-End Routers (redundant configuration)
Head-End Routers (redundant configuration)
Public Network
Remote Site Routers

ARCH v1.1—9-7

The hub-and-spoke design is used when there is a single regional or headquarters location with a large number of remote offices, and the majority of all traffic is between the remote sites and the regional or headquarters location. The variations of the hub-and-spoke topology are:

- The remote sites communicate exclusively to the head-end location.
- The remote sites require communication with each other as well as to the head-end location.

The table lists the advantages and disadvantages of the hub-and-spoke topology for VPNs.

Advantages	Disadvantages
This design typically minimizes the device configuration and complexity of the solution by having a single IPSec connection or a single GRE tunnel from each remote location back to the regional or headquarters location.	This design does not scale well when there is a requirement for a high degree of traffic flow between remote sites. Traffic flow between remote sites can be accomplished by using GRE tunneling with IPSec, and routing traffic through the regional or headquarters location.
—	This design does not allow for any redundancy in the VPN network in the event of a failure of the single headquarters or regional office location.

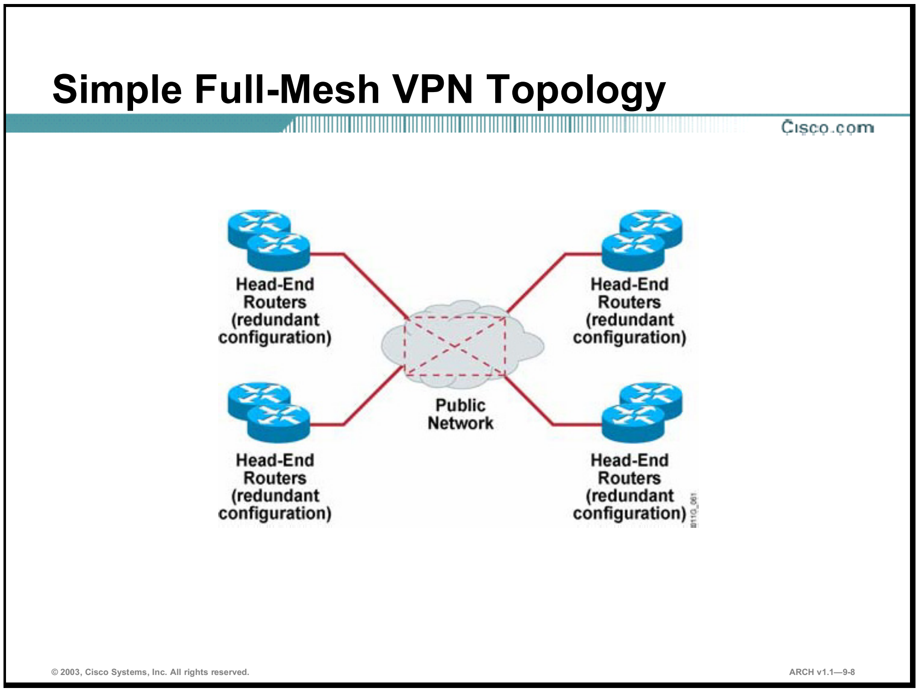

Simple Full-Mesh VPN Topology

Cisco.com

Head-End Routers (redundant configuration)
Head-End Routers (redundant configuration)
Public Network
Head-End Routers (redundant configuration)
Head-End Routers (redundant configuration)
B11G_061

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-8

Mesh VPN designs can either be fully meshed, providing any-to-any connectivity, or partially meshed, providing some-to-some connectivity, depending upon the customer requirements. The meshed topology is the appropriate design to use when there are a small number of total locations (regional, headquarters, or remote locations), with a large amount of traffic flowing between some (partial mesh) or all (full mesh) of the sites.

The table lists the advantages and disadvantages of the meshed topology for VPNs.

Advantages	Disadvantages
In a meshed design, the loss of a single location only affects traffic to or from that location. All other locations remain unaffected.	This design does not scale well when there are a large number of sites, due to the large number of straight IPSec connections and/or GRE tunnels with IPSec that have to be configured on each device.

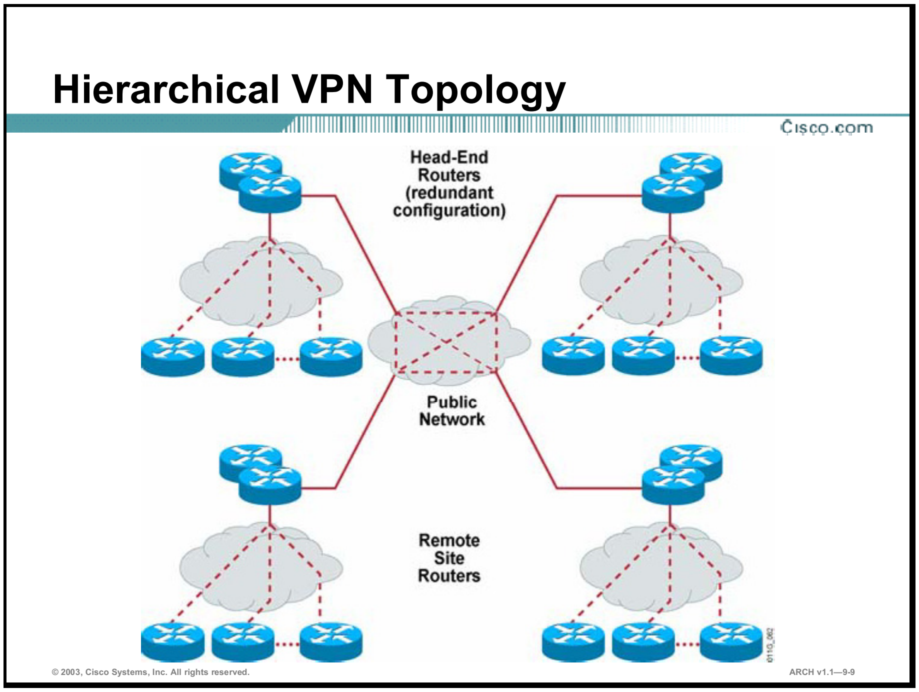

Hierarchical VPN Topology

Cisco.com

Head-End Routers (redundant configuration)

Public Network

Remote Site Routers

ARCH v1.1—9-9

A hierarchical topology design consists of a full or partial mesh core, with peripheral sites connecting into the core using a hub-and-spoke design.

A hierarchical topology is the appropriate design to use with larger networks that contain both a large number of remote offices, which have little traffic interaction between them, and a smaller number of headquarters or regional offices, with a large amount of traffic flowing between them. Hierarchical designs are the most scalable design for large networks, which require either partial or full connectivity between sites.

The criteria for the core components are similar to those for a meshed design. Likewise, the criteria for the peripheral components are similar to those for a hub-and-spoke design. The design differences depend upon whether a single set of routers will be used for both the head-end VPN termination device for the core component and the VPN termination device for the peripheral component; or if two sets of routers will be used.

This design is the most complex of the designs in terms of configuration, and may have a combination of GRE tunnels with IPSec running over them, or straight IPSec connections.

High-Availability and Resiliency Considerations

When remote user or branch office connectivity is critical to the successful operation of the enterprise, downtime for the VPN is not an option. Enterprises need a systematic approach to examine all the essential elements in delivering a high-availability site-to-site VPN. This topic lists the high-availability and resiliency design considerations for enterprise site-to-site VPNs.

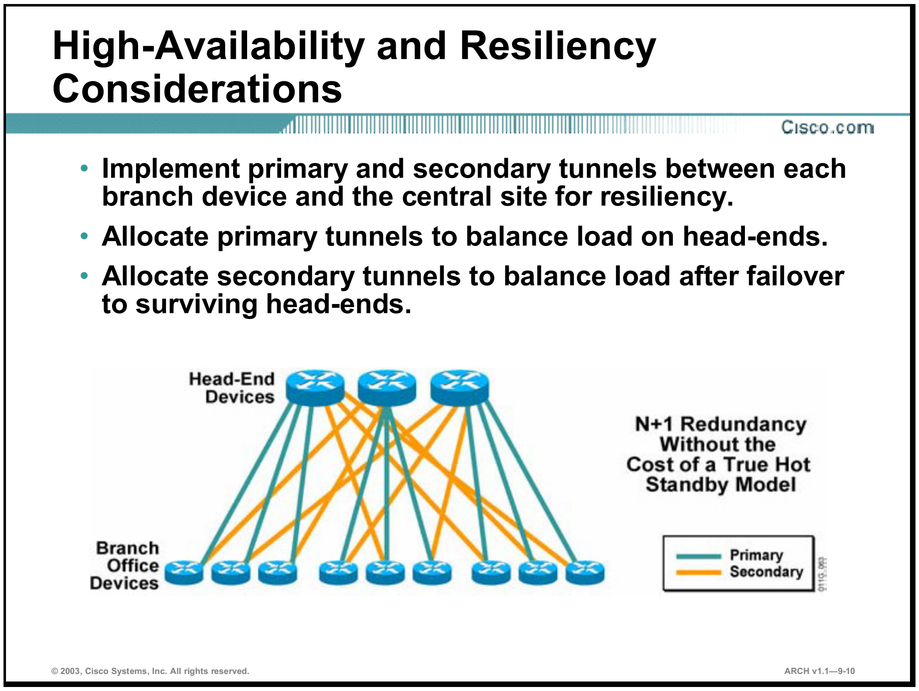

High-Availability and Resiliency Considerations

Cisco.com

- **Implement primary and secondary tunnels between each branch device and the central site for resiliency.**
- **Allocate primary tunnels to balance load on head-ends.**
- **Allocate secondary tunnels to balance load after failover to surviving head-ends.**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-10

A typical IPSec VPN may involve a number of site-to-site connections. IPSec protects information as it travels from one part of the private network to another part over the public network. For each unique connection across the public network, a unique IPSec connection is established between the two peer points at the boundary of the private and public networks. An IPSec connection consists of one IKE security association and at least two dependent IPSec security associations. Security associations, identified by a unique security parameter index, are stateful relationships between the two peer points. The state information includes common secret keys, security parameters, and peer identity. This state information is established during the main-mode negotiation for the IKE security association (SA) and the quick-mode negotiation for IPSec security associations. If there is a prolonged loss of IP connectivity between the two peers, you must set up a new set of relationships through stateless failover.

There are two steps that must occur for stateless failover before the process is successful:

- One of the peers must detect the loss of connectivity.
- Once detected, the peer must take action to reconnect with another peer to reach the part of the private network at the same site.

Loss of IP connectivity can be caused by local-link failure, full-loss connectivity by the service provider, or device failure. For a typical remote site, a dedicated or dial-on-demand path to the head-end site can protect against failure. Protection against a local-device failure at a remote site is not usually provided unless the importance of connectivity for the remote site warrants the cost. For a typical head-end site, you can achieve redundancy by implementing multiple provider connections and by deploying multiple head-end routers.

Using a Routing Protocol over the VPN

A site-to-site VPN solution will support static and dynamic routing protocols that are implemented elsewhere in the network. This topic lists the routing protocol design considerations for enterprise site-to-site VPNs.

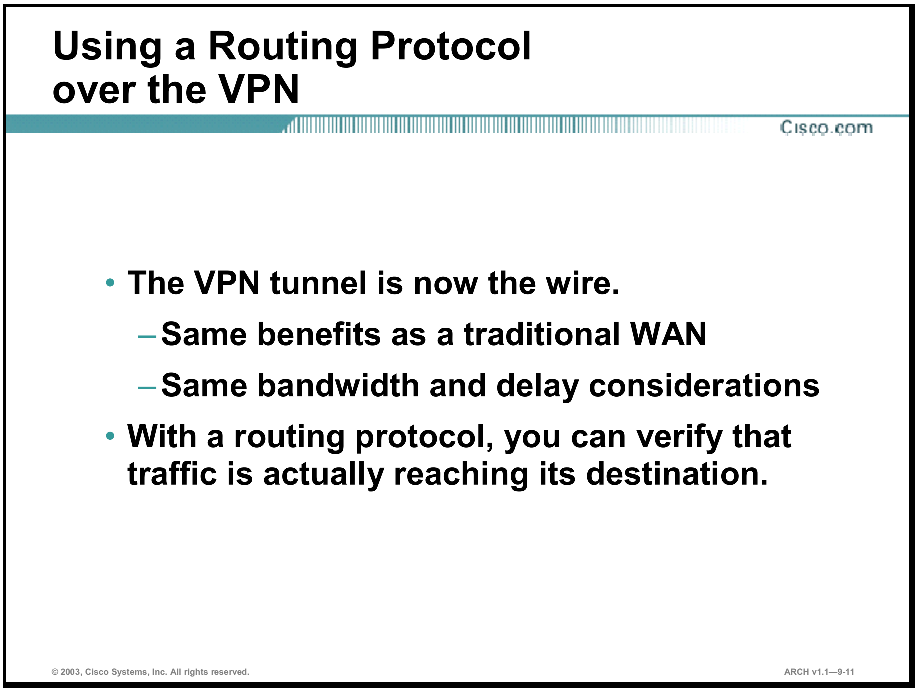

Using a Routing Protocol over the VPN

- The VPN tunnel is now the wire.
  - Same benefits as a traditional WAN
  - Same bandwidth and delay considerations
- With a routing protocol, you can verify that traffic is actually reaching its destination.

ARCH v1.1—9-11

Routing functionality enables a central-site VPN device to efficiently learn a remote network during initial installation, and dynamically update the connections over time. This scenario drastically reduces installation time and the management overhead of maintaining static routing tables.

Routing functionality in VPN devices dramatically improves the flexibility of these devices. In many cases, enterprises will choose to connect one or more remote sites through a remote access-based VPN deployment. You can implement routing protocols within GRE tunnels to support functionality such as any-to-any connectivity of remote sites within a hub-and-spoke design without creating a full mesh of IPSec security associations. This can be particularly useful in large implementations, which require all remote sites to communicate with each other. Redundancy, through the use of routing protocols within GRE tunnels, is an additional functionality that can be supported.

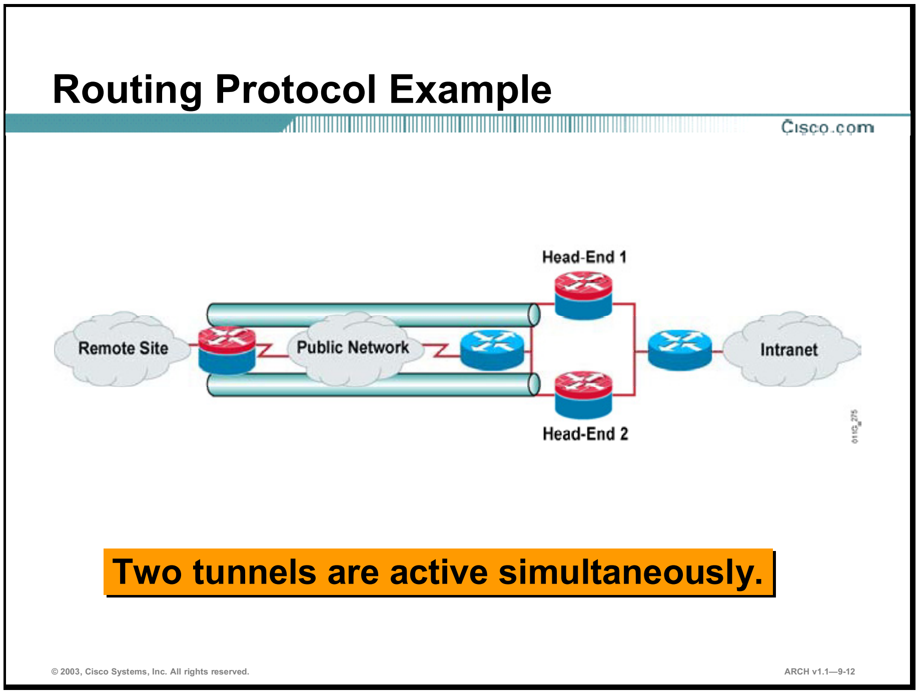

Routing Protocol Example

Cisco.com

Two tunnels are active simultaneously.

ARCH v1.1—9-12

Routing protocols use two paths, as they are sent over IPSec-protected GRE tunnels to track remote network reachability. A remote site using routing protocols for high availability will establish two IPSec-protected GRE tunnels, one to each head-end. Routing updates traverse both tunnels to the remote site, which will then forward the traffic to the head-end that has reachability to the destination network. From the perspective of the remote site, there are two paths to the head-end. Consider defining one of the tunnels as the primary tunnel to avoid asymmetric routing, by adjusting the routing protocol cost for one of the links. In case of tunnel failure, convergence will occur as soon as the routing protocol realizes the path is no longer available. After failure recovery, remote sites using routing protocols will optionally revert back to their primary preferred path.

Concentrator and firewall head-ends often support failover capabilities in an active or standby configuration. When the primary fails, the secondary unit assumes the IP and MAC address of the primary, and the tunnel reestablishment commences. Routers function in an active-active configuration. Both head-end devices will allow tunnel establishment. You could consider using IKE keepalives in the head-end for heterogeneous remote-site device support. There is no IETF standard for keepalives today, and thus this mechanism will work only with products from a single vendor. If a momentary loss of connectivity occurs at a remote site, it may establish a new tunnel with the secondary (but always active) head-end device. Because tunnel establishment does not affect the routing table unless routing protocols are running over the tunnel, the routing state in the head-end will not change. Flapping occurs when the remote site temporarily loses WAN connectivity. When the tunnel switches between the head-ends because of remote-site flapping, the next-hop router will not be able to determine which active head-end device has a valid path to the remote site. Failover of IPSec between devices today is stateless and requires new tunnel establishment on the secondary device.

In summary, when using VPN concentrators or firewalls at the head-end, use IKE keepalives for high availability. When using VPN routers at the head-end, use routing protocol resilience for high availability.

Minimizing Packet Fragmentation

IPSec and GRE headers increase the size of packets being transported over a VPN. If the size of a packet before encryption is near the maximum transmission unit (MTU) of the transmission media, the encrypted packet with the additional IPSec and GRE headers can exceed the MTU of the transmitting interface. This topic lists the packet fragmentation design considerations for enterprise site-to-site VPNs.

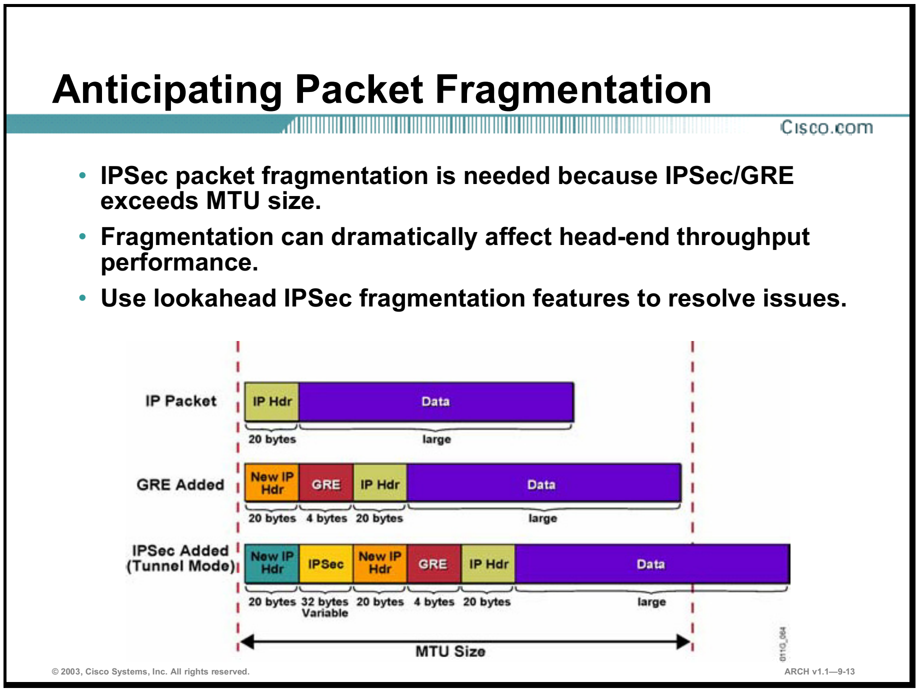

Anticipating Packet Fragmentation

Cisco.com

- **IPSec packet fragmentation is needed because IPSec/GRE exceeds MTU size.**
- **Fragmentation can dramatically affect head-end throughput performance.**
- **Use lookahead IPSec fragmentation features to resolve issues.**

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-13

The Layer 3 packet fragmentation requires these packets to be reassembled prior to the decryption process. In some enterprise networks, fragmentation results in less-than-optimal throughput performance.

To avoid the fragmentation problem, you can:

- Employ the MTU discovery path.
- Set the MTU to allow for packet expansion, such as 1400 bytes.

For some networks, it might not be possible to easily manage MTU size. For these situations, Cisco has implemented prefragmentation for IPSec VPNs.

IP MTU discovery can eliminate the possibility of fragmentation when it is supported by the end stations. This is a procedure that is run between two end stations with the participation of the network devices between them. For this process to work over an IPSec network with GRE, the GRE tunnel MTU should be set to a value low enough to ensure that the packet will make it through the encryption process without exceeding the MTU on the outbound interface, usually 1400 bytes.

Implementing IPSec Security

IPSec tunnels are not virtual interfaces. They can only carry unicast IP packets, and have no end-to-end interface management protocol. To add resiliency to IPSec tunnels, you can implement it in transport mode on top of a robust tunnel technology such as GRE. This topic lists the IPSec security design considerations for enterprise site-to-site VPNs.

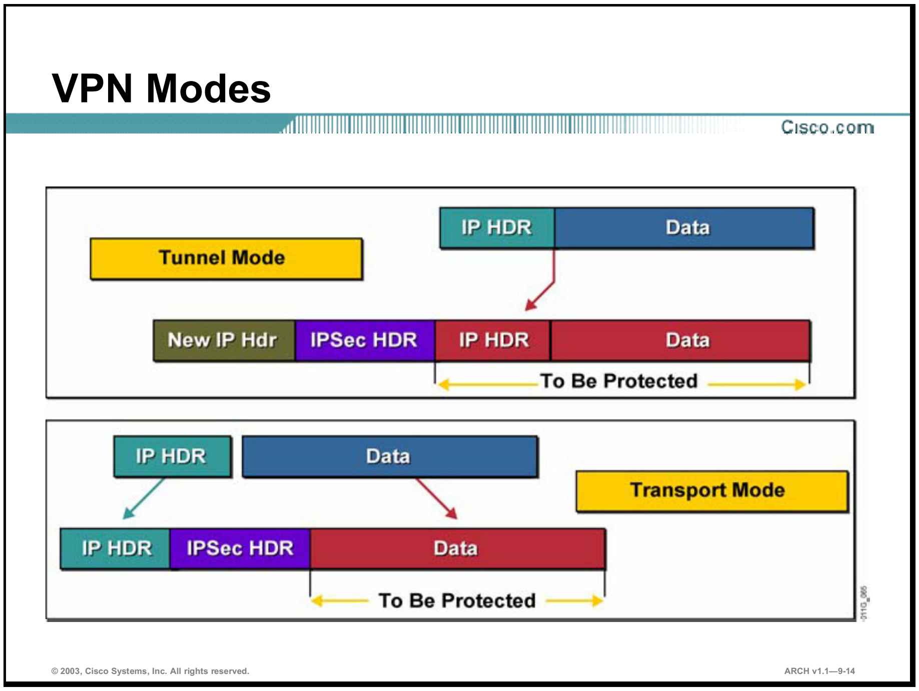

VPN Modes

Cisco.com

Tunnel Mode

IP HDR | Data

New IP Hdr | IPSec HDR | IP HDR | Data

To Be Protected

Transport Mode

IP HDR | Data

IP HDR | IPSec HDR | Data

To Be Protected

ARCH v1.1—9-14

A GRE tunnel is a virtual interface in a router and provides many of the features associated with physical interfaces. This is generally preferable except when the head-end router connects to thousands of remote sites.

GRE tunnels provide the ability to encapsulate multicast broadcast packets and non-IP protocols. Enabling this feature may enhance performance and scalability for site-to-site VPN services. Since GRE tunnels are unique interfaces, they can each be assigned their own cryptographic maps. When the source router needs to send a packet on the VPN destined for the other end of the GRE tunnel, it first makes a routing decision to send it out an interface and then does a search of the security parameter index (SPI) table to find the corresponding SA. With GRE tunnels, the router must make a routing decision across a multitude of GRE tunnels. Once the GRE tunnel is chosen, there are then a limited number of SAs from which to choose.

Since GRE provides the tunneling mechanism, IPSec can be configured in transport mode, eliminating a portion of the IPSec overhead that is present in IPSec tunnel mode. By using transport mode, the IPSec packets do not need to be tunneled in IPSec as well as GRE for traffic when the endpoints are both the source and destination of the traffic.

For VPN resilience, configure the remote site with two GRE tunnels: One to the primary head-end VPN router and the other to the backup VPN router. Both GRE tunnels are secured via IPSec. Since GRE can carry multicast and broadcast traffic, it is possible and often desirable to configure a routing protocol for these virtual links. Once a routing protocol is configured, it provides failover capabilities. The hello/keepalive packets sent by the routing protocol over the GRE tunnels provide a mechanism to detect loss of connectivity. In other words, if the primary GRE tunnel is lost, the remote site will detect this event by the loss of the routing protocol hello packets. Once virtual-link loss is detected, the routing protocol will choose the next best route: The backup GRE tunnel. VPN resilience is obtained by the automatic behavior of the routing protocol. Since the backup GRE tunnel is already up and secured, the failover time is determined by the hello packet mechanism and the convergence time of the routing protocol.

Using GRE tunnels with IPSec in transport mode provides a robust resilience mechanism for hub-and-spoke topologies. The three main issues when using this feature are:

- Manageability of the Cisco IOS configuration for the head-end router
- Overhead on the network and the router processor
- Scalability

Site-to-Site VPN Examples

You can implement site-to-site VPNs in both small and large enterprise environments. This topic provides examples of simple and complex site-to-site VPN solutions that meet specific enterprise VPN requirements.

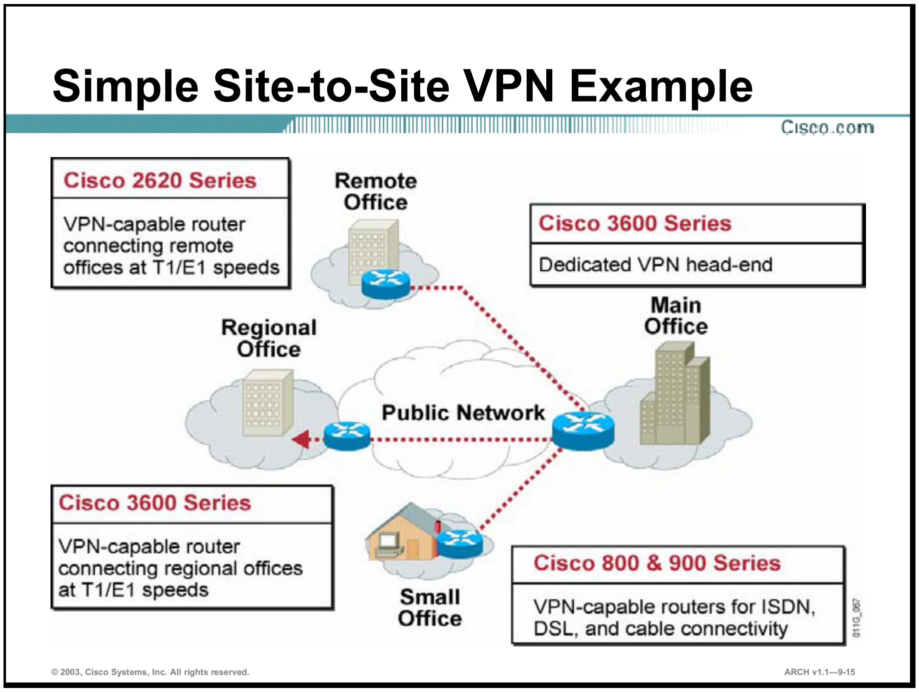

Simple Site-to-Site VPN Example

Cisco.com

Cisco 2620 Series

VPN-capable router connecting remote offices at T1/E1 speeds

Remote Office

Cisco 3600 Series

Dedicated VPN head-end

Main Office

Regional Office

Public Network

Cisco 3600 Series

VPN-capable router connecting regional offices at T1/E1 speeds

Small Office

Cisco 800 & 900 Series

VPN-capable routers for ISDN, DSL, and cable connectivity

ARCH v1.1—9-15

Company Background

The small publishing company shown in the figure has done a comparative analysis between a leased-line WAN solution and a VPN. The company is using a centralized data store that resides in the corporate network. It accesses the data store through a web-based interface. The company uses e-mail as a primary means of communication and the Microsoft Office package for office productivity. Employees use extensive file sharing with files that include large graphics.

The company's main office has 150 people who access the corporate network. In addition, these offices need access to the corporate network:

- Regional office: 75 users
- Remote office: 26 users
- Six small offices: 15 users each

The regional office and remote office have T1 access to the headquarters office. The small offices are using asymmetric digital subscriber line (ADSL) at 512 kbps to connect to the Internet. Each office has two physical connections to the Internet.

Site-to-Site VPN Solution

The table summarizes the design decisions that the enterprise made to meet its requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the site-to-site VPN?	Hub-and-spoke topology	All corporate traffic will pass through the corporate network since all of the shared file storage, database, and e-mail engines are at the main office.
What type of tunneling will be deployed?	GRE between each site and the main office Single tunnel to each location without VPN client software on each client computer	This solution maintains simplicity. Since the company has two interfaces to the Internet, one is designated for the tunnel and one for Internet access.
What type of security will be deployed?	Basic security on each tunnel using IPSec Eight tunnels connect from the hub to individual sites Firewall at each site	Basic security keeps any confidential traffic from going over the Internet in plain text. Each site has a firewall either in the same device as the VPN tunnels or separate from the VPN tunnels.
Is NAT required?	Yes	NAT is not needed over the VPN tunnels but will be used on the link going to the Internet.
What VPN hardware will be used?	Main office: Cisco 3600 series[1] Regional office: Cisco 3600 series Remote office: Cisco 2621 Small office: Cisco 800	Given the requirement for dedicated VPN hardware with hardware encryption, the company selected the Cisco 3600 series at the main office. Given 75 users at the regional office, the company selected the Cisco 3600 series. The remote office has 26 users. Since there is no need for dedicated VPN hardware, a Cisco 2621 was selected to support two Ethernet segments, and one- to two-port serial WIC[2]. The Cisco 2621 leaves growth for voice and the ability to insert a hardware encryption card. The small office requires VPN access over DSL. The two most common hardware solutions are standard routers, and DSL specialty devices. The company selected a DSL specialty device, the Cisco 800.
What type of high availability will be deployed?	No additional redundancy or failover, but a separate interface will be used for Internet access at each site	A separate interface for Internet traffic means that Internet traffic will not affect the performance of the entire network. The IT department agreed that a temporary outage would not stop the business from running.

[1] The example platforms are accurate as of the date this course was published.

[2] WIC = WAN interface card

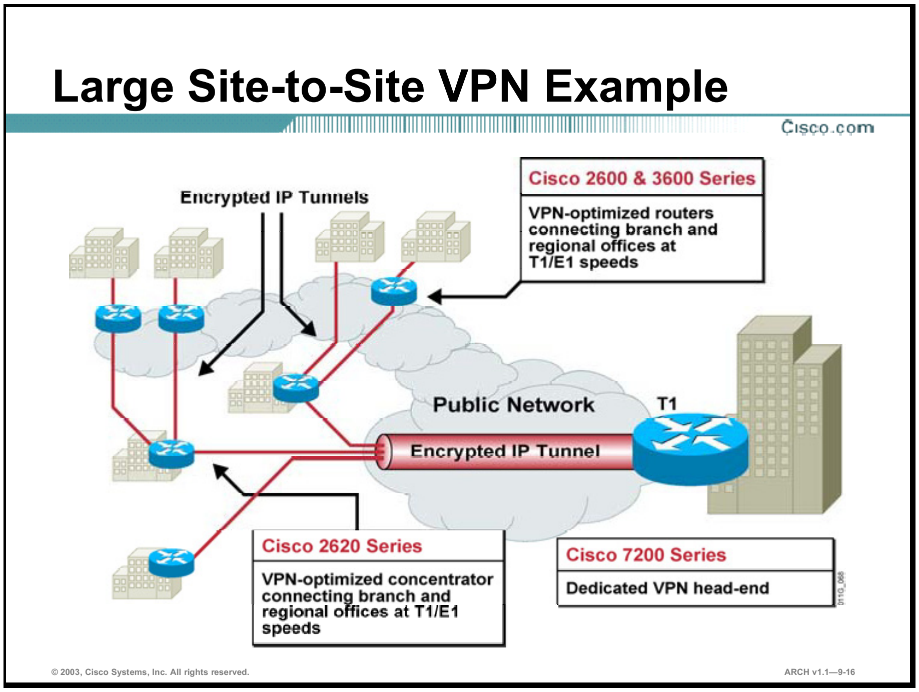

Large Site-to-Site VPN Example

Cisco.com

Cisco 2600 & 3600 Series

VPN-optimized routers connecting branch and regional offices at T1/E1 speeds

Cisco 2620 Series

VPN-optimized concentrator connecting branch and regional offices at T1/E1 speeds

Cisco 7200 Series

Dedicated VPN head-end

ARCH v1.1—9-16

Company Description

The large publishing company shown in the figure needs a VPN solution between its headquarters office and a number of large branch offices. VPN access is required from the remote sites to the regional sites.

The headquarters office has 1600 people, and contains master file servers and data servers. It has a file distribution system that pushes a copy of the production files to a server in each region. It relies heavily on e-mail with attachments for corporate correspondence. The primary e-mail server is located at the headquarters corporate network.

Each regional office has approximately 286 people, and has a support server farm where they get a copy of the latest data store and published files on a nightly basis. Each regional facility has multiple T1s for access to its service provider.

The remote site offices have approximately 45 people at each facility. They are logically part of their regional office and retrieve all of their data store information from their regional office. The remote site offices also submit their updated data to their regional office at the end of each business day.

Site-to-Site VPN Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the site-to-site VPN?	Meshed VPN solution	The company implemented a meshed VPN solution between headquarters and the regional offices, and between the regional offices and the remote sites. The final design is a hierarchical design. All Internet traffic goes through the headquarters office. Headquarters needs the ability to forward out to the Internet.
What type of tunneling will be deployed?	14 site-to-site tunnels using IPSec and GRE	A total of six tunnels is required between the headquarters and the regional sites. From the regional office to the remote offices, eight tunnels are required. Hardware encryption was required in all devices and dedicated VPN hardware at the headquarters and regional sites.
What type of security will be deployed?	Any data sent over the Internet will have the original IP header and data encrypted Basic security required through tunneling and firewalls	There is no outside access to the Internet except through the headquarters facility. Each site requires a firewall, either in the same device as the VPN tunnels or separate from the VPN tunnels.
Is NAT required?	Yes	Since the headquarters site handles Internet access, there is a need for NAT going out to the Internet, offered on a separate system from the VPN solution.
What VPN hardware will be used?	Headquarters: Cisco 7200[1] Regional offices: Cisco 2600 and 3600 series Remote offices: Cisco 2621	Since there is a requirement for dedicated VPN hardware with hardware encryption at the headquarters site, the Cisco 7200 series was selected to provide hardware encryption and cryptography. At the regional offices, based on the number of users, the Cisco 2600 or 3600 series was selected. Since there was no need for dedicated VPN hardware at the remote offices, the Cisco 2621 was selected, leaving room for growth.
What type of high availability will be deployed?	Redundant paths between headquarters and regional site, and paths between regional sites Failover from regional sites to the remote sites	

[1] The example platforms are accurate as of the date this course was published.

Summary

This topic summarizes the key points discussed in this lesson.

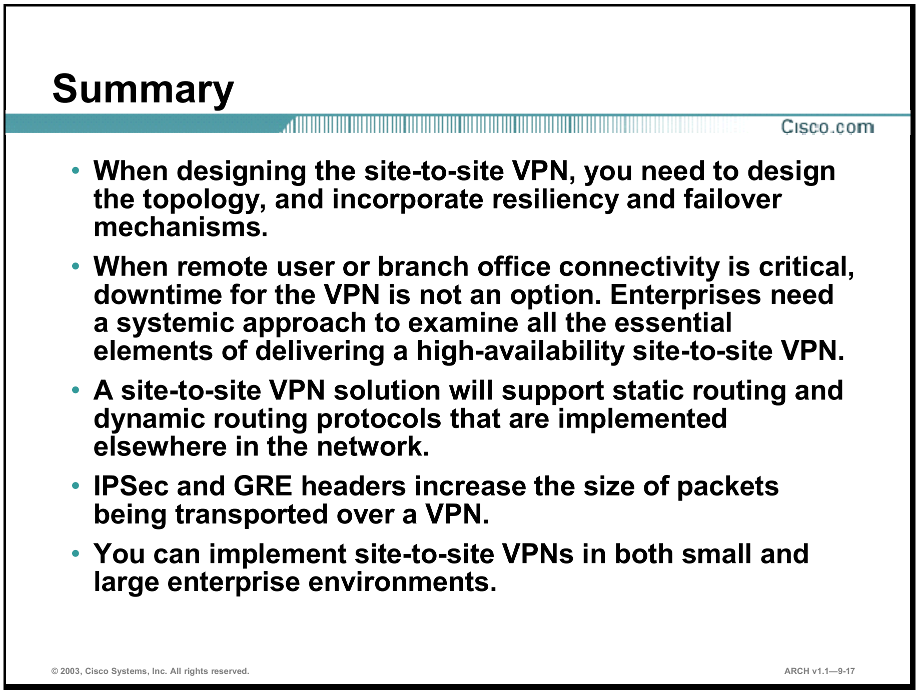

Summary

Cisco.com

- When designing the site-to-site VPN, you need to design the topology, and incorporate resiliency and failover mechanisms.
- When remote user or branch office connectivity is critical, downtime for the VPN is not an option. Enterprises need a systemic approach to examine all the essential elements of delivering a high-availability site-to-site VPN.
- A site-to-site VPN solution will support static routing and dynamic routing protocols that are implemented elsewhere in the network.
- IPSec and GRE headers increase the size of packets being transported over a VPN.
- You can implement site-to-site VPNs in both small and large enterprise environments.

ARCH v1.1—9-17

References

For additional information, refer to these resources:

- *A Primer for Implementing a Cisco Virtual Private Network* at http://www.cisco.com/warp/public/cc/so/neso/vpn/vpne/vpn21_rg.htm
- *SAFE VPN: IPSec Virtual Private Networks in Depth* at http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safev_wp.htm
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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 statements often describe VPNs? (Choose three.)

A) VPNs are more secure than traditional WANs.

B) VPNs are more scalable than traditional WANs.

C) VPNs are more flexible than traditional WANs.

D) VPNs are more expensive than traditional WANs.

E) VPNs are more cost effective than traditional WANs.

Q2) Which topology is the most scalable design for large networks, which require either partial or full connectivity between sites?

A) full mesh

B) hierarchical

C) partial mesh

D) hub-and-spoke

Q3) When planning for high availability, how many tunnels should be planned from each remote site?

A) two, because of the IPSec timeout value

B) only one, because IPSec can only handle one

C) only one because it is too expensive to get more thanone1 tunnel

D) two, one to a primary endpoint and one to a secondary endpoint

Q4) What function is improved by running a routing protocol on VPN devices?

A) cost

B) flexibility

C) throughput

D) performance

Q5) Why is fragmentation a problem?

A) It requires more bandwidth to the remote site.

B) It requires the use of bundling to reassemble the packets.

C) It requires additional processor time to manipulate the packets.

D) It requires the sender to create a header to go around the packet.

Q6) When using IPSec, which additional tunneling mechanism is recommended, if needed?

A) GRE

B) LLP

C) L2TP

D) 3DES

Q7) To offer basic VPN security when the traffic includes multicast, what functionality should you select?

A) fixed firewalls

B) asynchronous DSL

C) IPSec and NAT to the Internet

D) encrypted tunnels using IPSec and GRE

Q8) You are designing a large site-to-site VPN solution that will support multiple protocols. Some telecommuters will need remote access. Which head-end VPN device will meet these requirements?

A) Cisco NAS server

B) Cisco PIX firewall

C) Cisco VPN concentrator

D) Cisco VPN-enabled router

Quiz Answer Key

Q1) B, C, E

Relates to: Site-to-Site VPN Requirements

Q2) B

Relates to: Key Design Considerations for Site-to-Site VPNs

Q3) D

Relates to: High Availability and Resiliency Considerations

Q4) B

Relates to: Using a Routing Protocol over the VPN

Q5) C

Relates to: Minimizing Packet Fragmentation

Q6) A

Relates to: Implementing IPSec Security

Q7) D

Relates to: Site-to-Site VPN Examples

Q8) D

Relates to: Site-to-Site VPN Examples

Designing Remote-Access VPNs

Overview

Remote-access VPNs permit secure, encrypted connections between mobile or remote users and their corporate networks via a third-party network, such as a service provider. Deploying a remote-access VPN enables enterprises to reduce communications expenses by leveraging the local dial-up infrastructures of Internet service providers.

Relevance

To fully realize the benefits of high-performance remote-access VPNs, an organization needs to deploy a robust, highly available VPN solution to their mobile and remote authorized users.

Objectives

Upon completing this lesson, you will be able to design simple and complex remote-access VPNs, given enterprise VPN needs. This includes being able to meet these objectives:

- Identify typical requirements for an enterprise remote-access VPN
- List the key design considerations for enterprise remote-access VPNs
- Plan the capacity for remote-access VPNs, given a number of users
- List issues with Network Address Translation (NAT) for remote-access VPNs
- Design small and large remote-access VPN solutions that meet specific enterprise VPN requirements

Learner Skills and Knowledge

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

- Designing Enterprise Edge Connectivity module
- Designing Security Solutions module
- Identifying VPN Technologies lesson

Outline

The outline lists the topics included in this lesson.

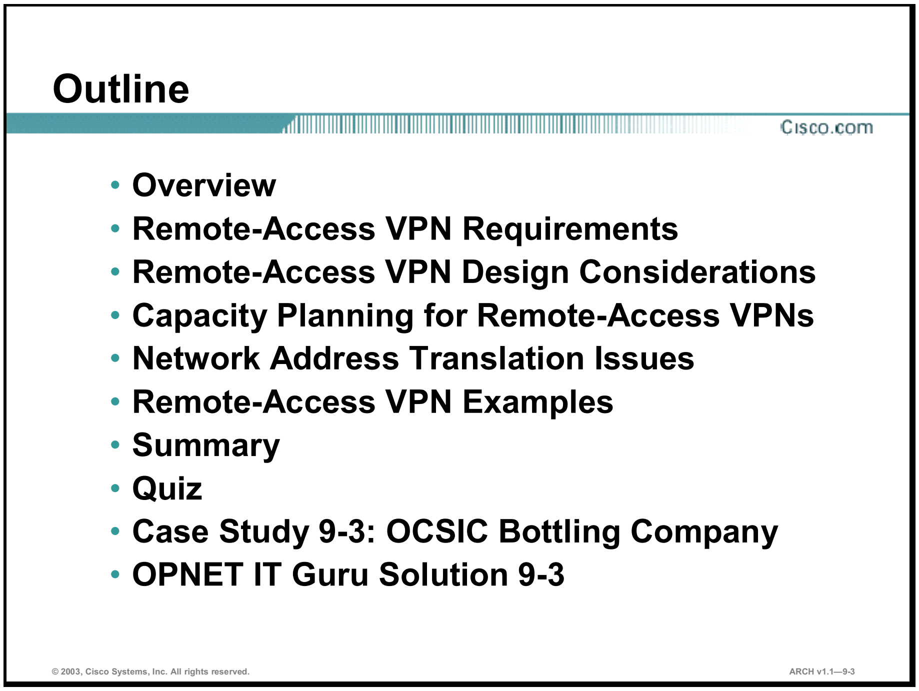

Outline

- Overview
- Remote-Access VPN Requirements
- Remote-Access VPN Design Considerations
- Capacity Planning for Remote-Access VPNs
- Network Address Translation Issues
- Remote-Access VPN Examples
- Summary
- Quiz
- Case Study 9-3: OCSIC Bottling Company
- OPNET IT Guru Solution 9-3

ARCH v1.1—9-3

Remote-Access VPN Requirements

Remote-access VPNs typically begin as a replacement technology for traditional remote-access servers. As high-speed Internet access and broadband connectivity continue to emerge as cost-effective choices for consumers and businesses, the VPN paradigm takes on greater strategic significance. This topic identifies typical requirements for a remote-access VPN.

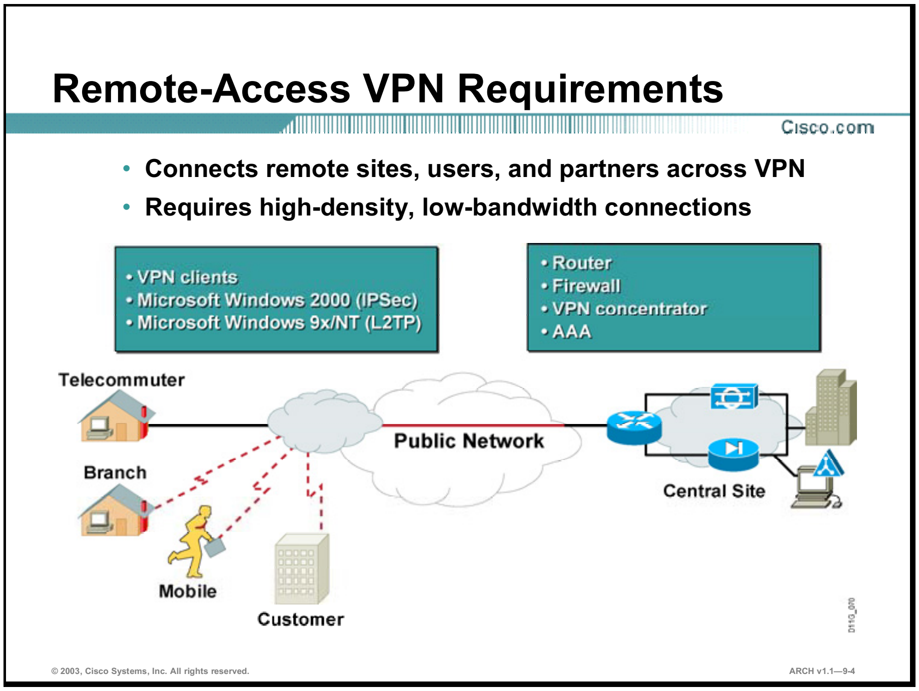

Remote-Access VPN Requirements

Cisco.com

- Connects remote sites, users, and partners across VPN
- Requires high-density, low-bandwidth connections

- VPN clients
- Microsoft Windows 2000 (IPSec)
- Microsoft Windows 9x/NT (L2TP)

- Router
- Firewall
- VPN concentrator
- AAA

ARCH v1.1—9-4

Remote-access VPNs encompass analog, dial-up, ISDN, digital subscriber line (DSL), mobile IP, and cable technologies to securely connect mobile users, telecommuters, or branch offices.

Remote-access VPNs require high-density, relatively low-bandwidth connections between remote users and a central site.

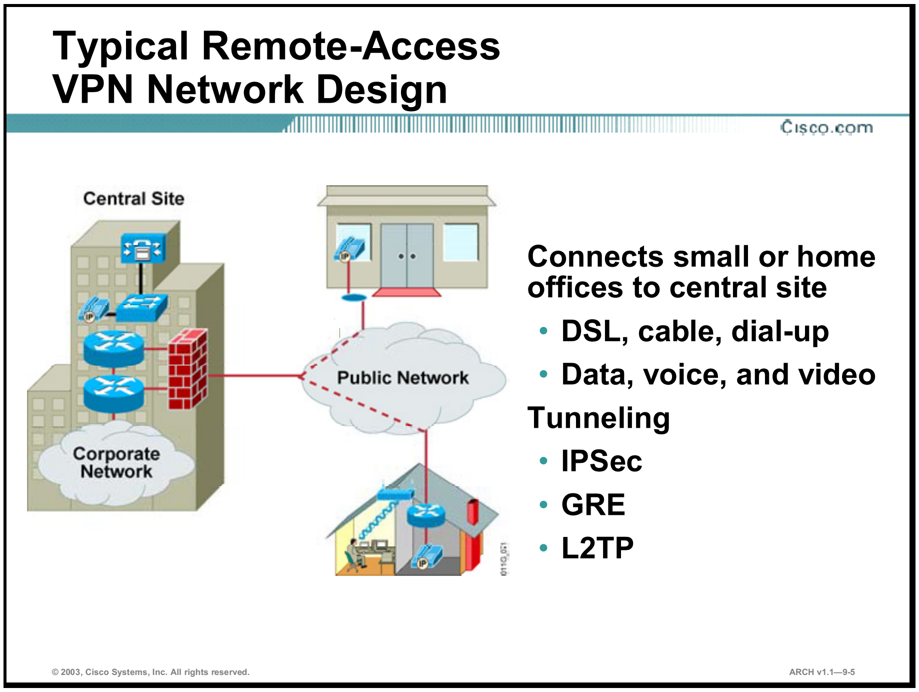

Typical Remote-Access VPN Network Design

Cisco.com

Connects small or home offices to central site

- DSL, cable, dial-up
- Data, voice, and video

Tunneling

- IPSec
- GRE
- L2TP

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-5

The principal advantages of a remote-access VPN solution include:

- **Reduced access charges:** VPNs use the Internet or a service provider’s shared IP network as the transport medium for private LAN access. Typically, enterprises deploy remote-access infrastructures in centralized topologies or in multiple locations for redundancy purposes. Toll-free access is the norm, but this becomes expensive as companies engage large numbers of remote workers. However, ubiquitous low-cost Internet access through ever-increasing point-of-presence (POP) coverage means that corporate users can connect to the enterprise using local numbers, thereby dramatically reducing or eliminating long-distance access charges.
- **Reduced overhead cost*:*** VPNs reduce the management overhead and capital investments required for remote access. In traditional remote-access infrastructures, modems are a core technology. In many cases, however, the modem is the weak link, frequently compromising the service availability and performance of the access infrastructure, due to the need for revision control, upgrades, and patches. This costly overhead consumes valuable IT resources. When this infrastructure is not maintained, remote-access users experience degraded service and service-call spikes.
- **Reduced service and support costs*:*** VPNs allow IT managers to take advantage of their service provider’s modem pool and remote-access server infrastructure for remote access from local POPs. Since top-tier service providers constantly monitor, maintain, and upgrade their dial-up infrastructure, you do not have to perform these functions. This drastically reduces primary service calls because the service provider can offer technical support at all times for any and all modem connectivity issues. This benefit is subtle but extremely important, because a large percentage of remote-access usage typically occurs after normal business hours.
- **More successful connections:** By dialing a local POP, end users typically achieve higher connection rates than by dialing directly into the enterprise over long-distance lines.

- **Improved scalability:** The cost and complexity of scaling and delivering remote-access services is drastically reduced with VPNs. With remote-access servers, scalability is limited to the number of integrated modems, typically 24 to 48 modems. To increase capacity, the IT manager must install additional systems and the T1/Primary Rate Interface (PRI) feeder (or individual measured business lines). Telco installation times vary greatly, and adding new remote-access systems is a considerable cost. With VPN devices, maximum capacities approach thousands of simultaneous sessions in a single chassis. The IT manager can simply add new logins and distribute client software as demand for the remote-access service increases, incrementally increasing the bandwidth of the connection to the ISP as actual use requires.
- **Improved availability:** A relatively unknown feature of modem-based, remote-access servers is that IT managers must estimate a suitable port ratio for the usage pattern of their particular enterprise, because, under normal circumstances, remote-access users dial in at different times during the day and stay online for different duration periods. Port ratios enable the enterprise to significantly reduce the cost of their remote-access infrastructure (cost savings are proportional to the port ratio) and conserve equipment charges.

Remote-Access VPN Design Considerations

To design a remote-access VPN, you will determine the primary applications and requirements for the system. This topic lists design considerations for remote-access VPNs.

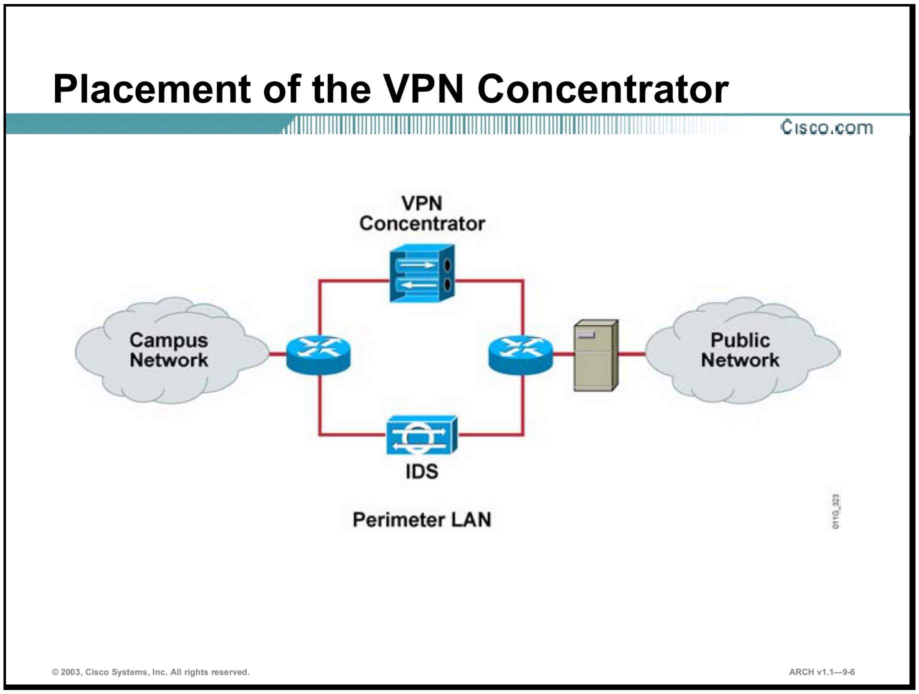

Placement of the VPN Concentrator

Cisco.com

ARCH v1.1—9-6

The VPN concentrator used to support a remote-access VPN solution resides within the Remote Access and VPN module. Generally, enterprises place an authentication, authorization, and accounting (AAA) server between the public network and the first router within the Remote Access and VPN module. The VPN concentrator then resides between routers so that all traffic goes through the VPN concentrator before being routed to the campus network.

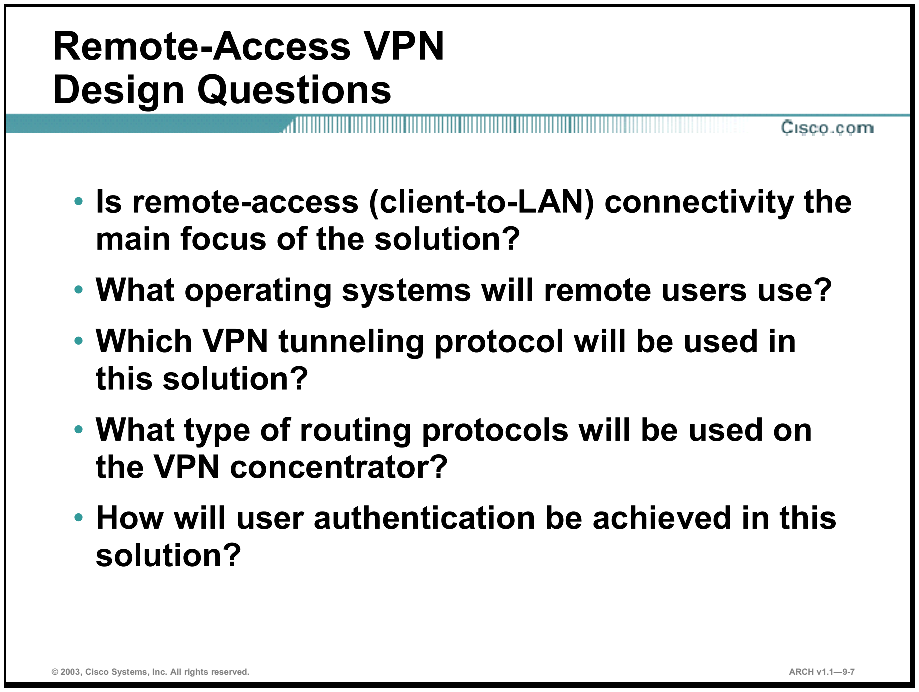

Remote-Access VPN Design Questions

- Is remote-access (client-to-LAN) connectivity the main focus of the solution?
- What operating systems will remote users use?
- Which VPN tunneling protocol will be used in this solution?
- What type of routing protocols will be used on the VPN concentrator?
- How will user authentication be achieved in this solution?

To design a remote-access VPN, ask the following questions:

- Is remote-access connectivity the main focus of the solution?
- What operating systems will remote users use? Determine if the VPN hardware and software client supports these operating systems.
- Which VPN tunneling protocol will be used in this solution?
- What type of routing protocols will be used on the VPN concentrator?
- How will user authentication be achieved in this solution?

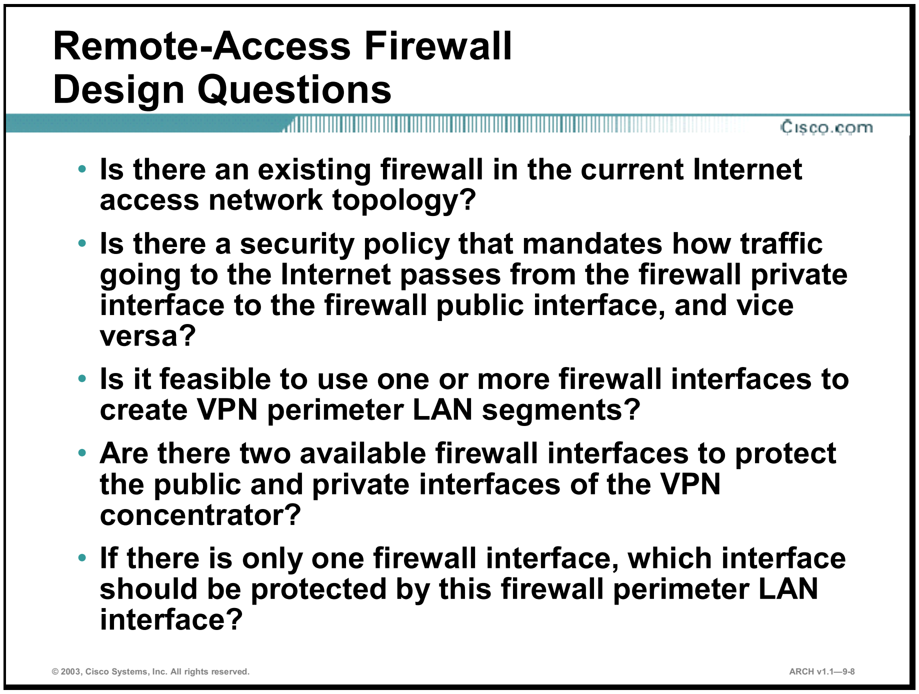

Remote-Access Firewall Design Questions

- **Is there an existing firewall in the current Internet access network topology?**
- **Is there a security policy that mandates how traffic going to the Internet passes from the firewall private interface to the firewall public interface, and vice versa?**
- **Is it feasible to use one or more firewall interfaces to create VPN perimeter LAN segments?**
- **Are there two available firewall interfaces to protect the public and private interfaces of the VPN concentrator?**
- **If there is only one firewall interface, which interface should be protected by this firewall perimeter LAN interface?**

ARCH v1.1—9-8

To design a firewall solution for the VPN, ask the following questions:

- Is there an existing firewall in the current Internet access network topology? If not, is there any plan to introduce a firewall into the new VPN design?
- Is there an internal corporate security policy in place that mandates that all traffic going to the Internet must pass from the private interface of the firewall to the public interface of the firewall, and vice versa, rather than traversing a perimeter LAN interface off the firewall?
- Is it feasible to use one or more available firewall interfaces to create VPN perimeter LAN segments?

At this point in the design formulation, we have determined that there exists at least one available firewall interface that can be used to protect one of the interfaces (public, private, or both) of the VPN concentrator.

- Are there two available firewall interfaces to protect both the public and private interfaces of the VPN concentrator?
- If there is only one available firewall interface, which firewall interface (public or private) should be protected by this firewall perimeter LAN interface?

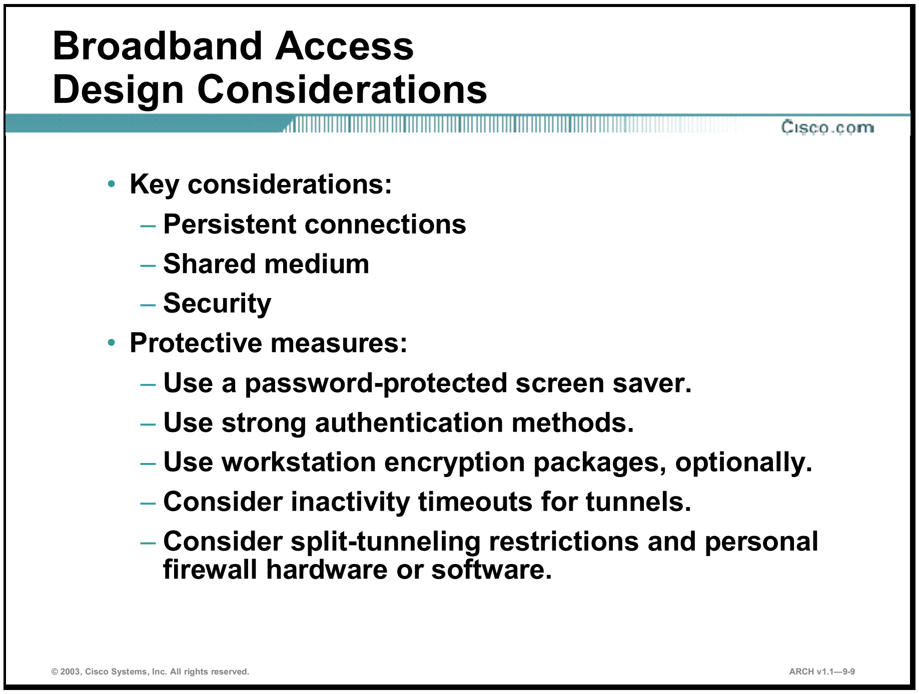

Broadband Access Design Considerations

- Key considerations:
  - Persistent connections
  - Shared medium
  - Security
- Protective measures:
  - Use a password-protected screen saver.
  - Use strong authentication methods.
  - Use workstation encryption packages, optionally.
  - Consider inactivity timeouts for tunnels.
  - Consider split-tunneling restrictions and personal firewall hardware or software.

Broadband can play a very important role in an enterprise VPN strategy. Today, cable and DSL service providers can offer corporate accounts, centralized billing, and turnkey VPN solutions. Always-on broadband connections are typically more cost effective than toll-free dial-up for high-usage users. In addition, user satisfaction is directly proportional to the speed of their remote-access connection. Broadband VPN does have security implications. Key areas to consider are:

- **Persistent connections:** The remote computer can be on all the time, either to the Internet, the corporate LAN via a tunnel, or both simultaneously.
- **Shared medium:** The enterprise LAN is being accessed over a shared medium. In some cases, the service provider may bridge traffic from one or more residential areas.
- **Security:** Abuse of the high-bandwidth connection can occur. For example, a disgruntled employee can easily download megabytes of confidential information in a very short period of time, which would be impractical at dial-up speeds. In some cases, the corporate PC will be attached to a private home LAN. Policies must ensure that sensitive corporate data is protected from family members, cohabitants, and visitors. You can use several protective measures, including:
  - A password-protected screen saver
  - Strong authentication methods, including token-based security and digital certificates
  - Workstation encryption packages
  - Inactivity timeouts for tunnels to protect the remote desktop from intrusion over the Internet or home LAN
  - Split-tunneling restrictions and personal firewall hardware or software

Capacity Planning for Remote-Access VPNs

You must select a VPN concentrator for a remote-access VPN based on current and future capacity projections. This topic will help you plan the capacity for remote-access VPNs, given the number of users.

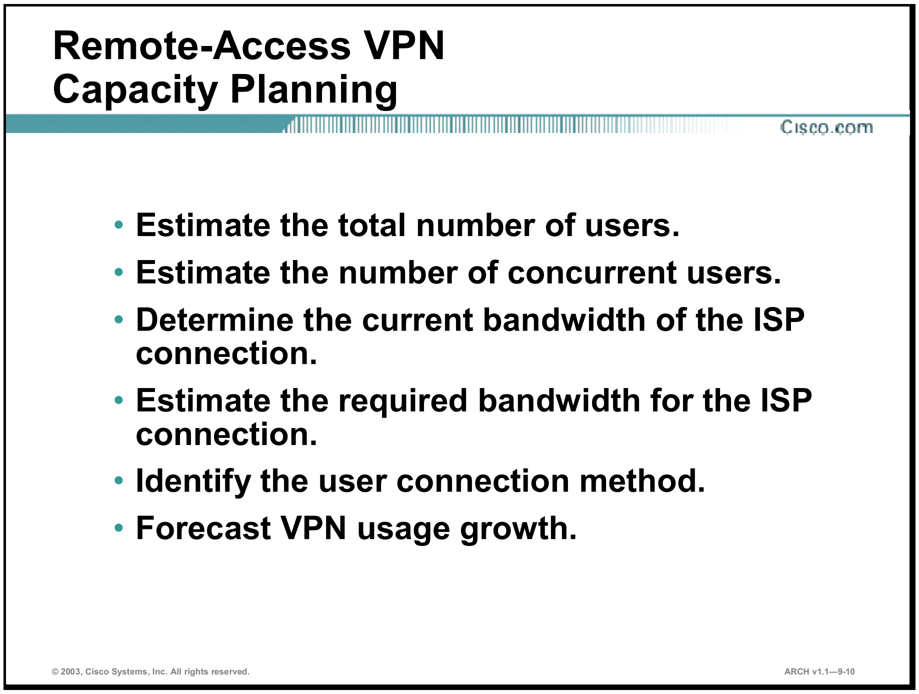

Remote-Access VPN Capacity Planning

- Estimate the total number of users.
- Estimate the number of concurrent users.
- Determine the current bandwidth of the ISP connection.
- Estimate the required bandwidth for the ISP connection.
- Identify the user connection method.
- Forecast VPN usage growth.

ARCH v1.1—9-10

Answer the following questions to help document the capacity planning design aspects of the remote-access VPN solution:

- What is an estimate for the total number of users who plan to take advantage of the remote-access VPN solution in the initial phase of deployment?
- What is an estimate for the number of concurrent users who plan to connect to the remote-access VPN solution in the initial phase of deployment?
- What is the current bandwidth of the links to the ISP or Internet?
- List the estimated number of users connecting to the remote-access VPN solution using each of the following methods.
  - Analog dial
  - Dedicated line
  - ISDN
  - Wireless
  - DSL
  - Frame Relay
  - Cable modem
  - Other (list)
- List any forecasted growth estimates for the remote-access VPN solution.

Based on this information, choose the appropriate VPN series concentrator model, asking the following question: Is the peak number of simultaneous users accessing the VPN concentrator expected to reach 100, or is the estimated amount of user traffic expected to exceed 4 Mbps in the near future? If the answer is yes, a VPN concentrator is required.

Network Address Translation Issues

Network Address Translation (NAT), along with IPSec, presents issues for the remote-access VPN. You can use NAT statically or dynamically for the remote-access VPN. This topic lists issues with NAT for remote-access VPNs.

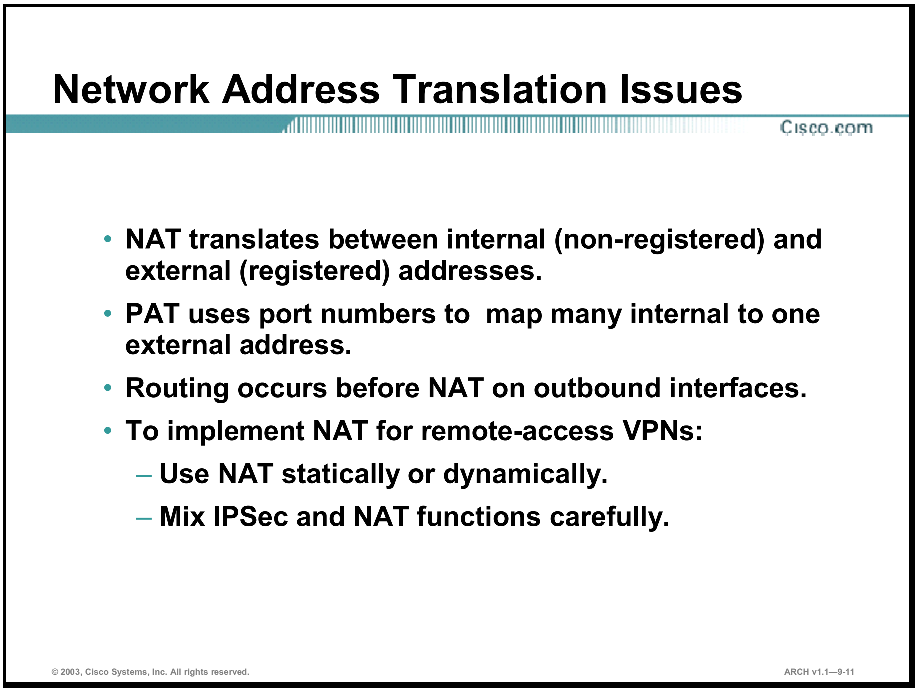

Network Address Translation Issues

Cisco.com

- NAT translates between internal (non-registered) and external (registered) addresses.
- PAT uses port numbers to map many internal to one external address.
- Routing occurs before NAT on outbound interfaces.
- To implement NAT for remote-access VPNs:
  - Use NAT statically or dynamically.
  - Mix IPSec and NAT functions carefully.

ARCH v1.1—9-11

NAT refers to the process of converting an IP address to a virtual IP address either to protect, or hide, the original IP address, or to convert a private (RFC 1918) IP address to an address that can be legally used and routed over the public Internet.

IPSec runs directly on top of IP and encrypts all information regarding the connection. Although this provides a very high degree of security, it poses a problem for devices performing NAT. NAT is the function used to disguise source addresses from network eavesdroppers and is a method used by companies to conserve IP addresses. To operate correctly, NAT needs access to some basic information in the transmitted packets (in particular, a unique identifier per session) that does not exist with IPSec packets. In short, since IPSec runs directly over IP, there is no port information and no unique unencrypted identifier for NAT to examine and map. Most devices that can perform NAT support two types of translations: One-to-one and many-to-one.

One-to-one translation does not require the use of ports and, therefore, does allow IPSec traffic to pass. Using this mechanism, an internal IP address block (for example, 10.2.4.0/24) is mapped to an external address block (for example, 125.100.2.0/24). A machine on the internal network (for example, 10.2.4.44) that would like to reach an external site such as Yahoo! can have all packets translated so that they are sent out with a source address of 125.100.2.44. This type of NAT, while not common in most enterprise deployments, will allow IPSec traffic to pass because one-to-one NAT mapping does not involve changing anything associated with ports, only modifying the source IP address on outbound packets and the destination address of inbound packets.

The most common NAT scenario is many-to-one NAT. With this setup, a block of internal addresses (for example, 10.2.4.0/24) are all mapped to appear as a single external IP address (for example, 125.100.2.254), allowing an administrator to hide internal IP addresses from the outside world while conserving IP address space. Routers and firewalls supporting NAT keep a translation table that contains a map of internal source addresses and ports, and they provide a new external port (visible address) that will replace the port address of the existing outgoing packet. This setup will allow return packets to come back to the source destination, mapped to the correct internal address based on the unique identifier, and then be routed to the appropriate device.

For the many-to-one type of NAT, most routers support mapping of only two main protocols, TCP and User Datagram Protocol (UDP). When TCP is used, source and destination ports are available for NAT so that the substitution can take place. This scenario is not possible with IPSec, which does not have any ports and encrypts all critical information, making parameters such as internal addresses invisible to the NAT process. This is why NAT typically fails when used with IPSec.

NAT Traversal (NAT-T) lets IPSec peers establish a connection through a NAT device. It does this by encapsulating IPSec traffic in UDP datagrams, using port 4500, thereby providing NAT devices with port information. NAT-T autodetects any NAT devices, and only encapsulates IPSec traffic when necessary.

Remote-access clients that support both NAT-T and IPSec/UDP methods first attempt NAT-T, and then IPSec/UDP (if enabled) if a NAT device is not autodetected, allowing IPSec traffic to pass through firewalls that disallow IPSec. The VPN concentrator implementation of NAT-T supports IPSec peers behind a single NAT or port address translation (PAT) device as:

- One Microsoft L2TP IPSec client
- One LAN-to-LAN connection

Either a LAN-to-LAN connection or multiple remote-access clients, but not a mixture of both, are allowed.

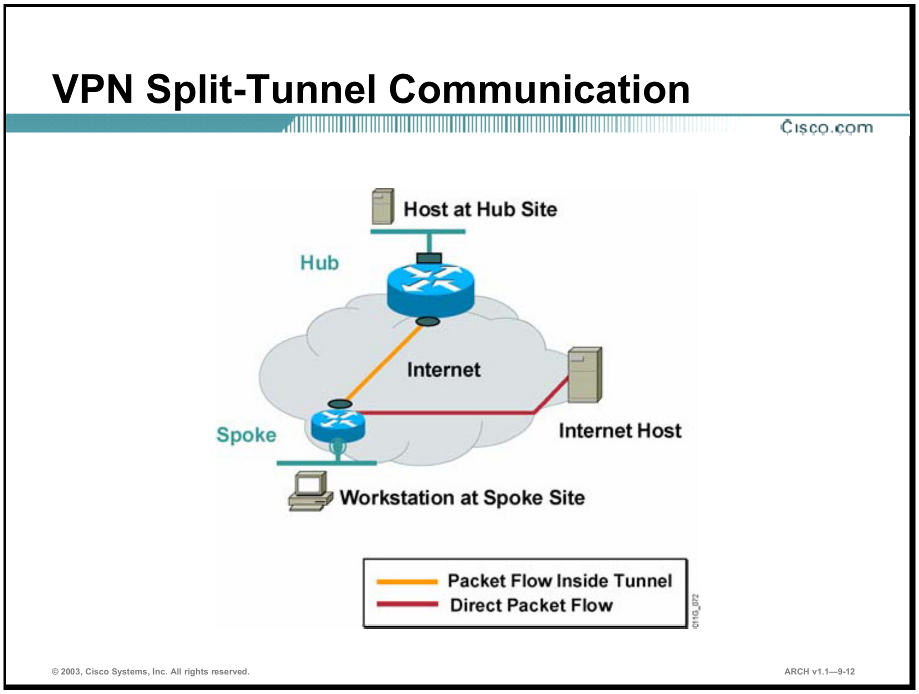

VPN Split-Tunnel Communication

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-12

Split tunneling is a configuration involving tunnel interfaces or protocols that encapsulate packets so data can flow either inside or outside of a particular tunnel.

This tunnel could be an IPSec tunnel, a GRE tunnel, or a combination of these or almost any other tunneling protocol. When split tunneling is not enabled, all the data flowing out a router's (or a PC of a client) egress interface will be encapsulated by the tunneling protocol and be sent directly to the peer router or concentrator of that particular tunnel. When split tunneling is enabled, that traffic may or may not be encapsulated in the tunneling protocol. This decision is normally made on a destination-by-destination basis by routing table entries, or via an access list entry.

When the split-tunneling feature is enabled, all the traffic leaving a particular site through a router will be forwarded directly to the next hop closest to the destination of that traffic. Only traffic destined to the remote site served by the tunnel will be encapsulated. This will decrease the load on a head-end or central site device when access is made to hosts not at that central site, such as the Internet. In the case of a home office network, this keeps traffic not specifically business-related out of the core of the network, and casual Internet surfing by family members does not appear as if it originated from the corporate network.

Because traffic is no longer guaranteed to be tunneled, access to the remote site may be made directly from outside networks without passing through any access controls placed at the central site. In the case of the Internet, this access will constitute a security problem because access to the remote site may not be authenticated. For this reason, the use of a stateful inspection firewall feature is recommended at each remote site with split tunneling enabled when the Internet is used for the VPN. There will then be two instances of firewalls that must be managed.

Management of the firewalls includes ensuring that the policies implemented on each router and firewall are synchronized. Failure to keep these policies synchronized causes a possible vulnerability. Managing the firewalls also may include intrusion detection. If an Intrusion Detection System (IDS) is deployed at the corporate location, you should replicate this function at each remote location that uses split tunneling.

When split tunneling is involved with a network that uses the Internet to create a VPN, it is highly recommended that a stateful inspection firewall function be configured at the spoke sites. Split tunneling allows an end station to directly access those hosts on the Internet. These end stations will not be protected by any firewall configured at the hub location. NAT is not a complete firewall solution. When running a routing protocol, a tunnel method other than IPSec tunnel mode must be used, since IPSec currently supports IP unicast traffic only.

When tunnels are present, the routing device should see two paths to a particular end point: one path through the tunnel and another path, which would appear longer from a routing standpoint, through the network unencapsulated. When data is encapsulated inside the tunnel, the destination may appear many hops closer than it actually is. This can cause problems if the state of the tunnel is not properly fed back into the routing table. This is often the case with static routes and GRE tunnels. When running a routing protocol, do not mix routes from the inside networks with the outside routes, because a routing loop or recursive routing could occur. A routing loop occurs if the route from the tunnel interface is many hops shorter than the route learned via the outside routing protocol. A packet could be sent into a router to be encapsulated and then sent out the same interface it entered the router on.

Remote-Access VPN Examples

You can implement remote-access VPNs in any network, from a small company to large enterprise environments. This topic provides examples of a small and a large remote-access VPN solution that meets specified VPN requirements.

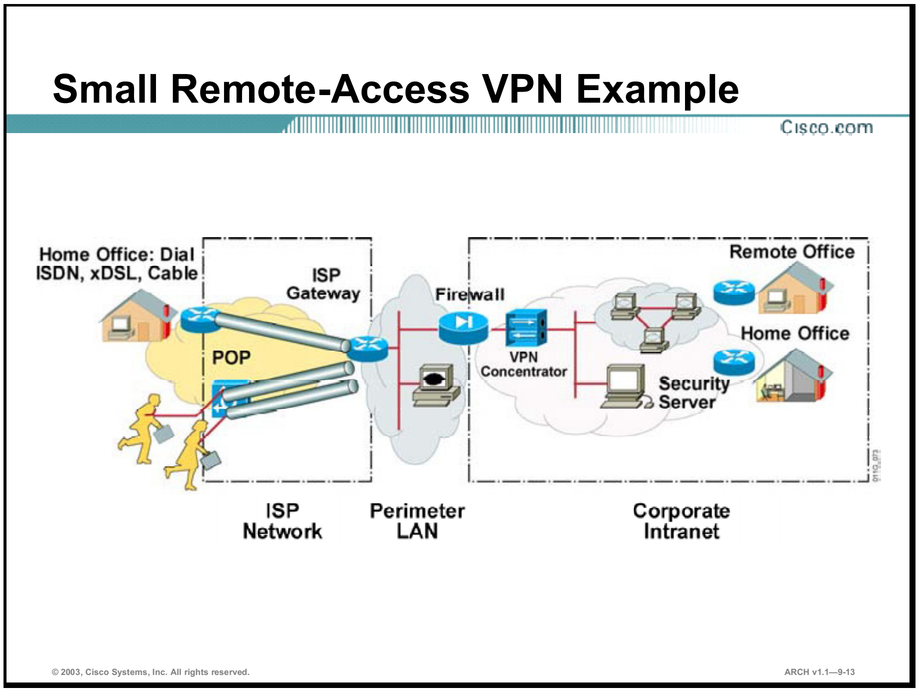

Small Remote-Access VPN Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—9-13

Company Description

A small training company has 19 office-based employees, 8 salespeople who work either in the office or from home, and 160 instructors and 55 course developers. Remote users need secure remote access to information on the corporate network, including e-mail and documents such as course materials, learner lists, and administration forms.

Internally, there is a single-switch network that connects to all of the office equipment. The company has a DSL router that connects to the Internet, and it has a public Class C address assigned to them.

The company thought about allowing Internet access to their corporate FTP server, but a new client needing courseware development is providing them with highly sensitive material on a new product to be released in a year. The arrangement to win the bid specified that the technical materials be safeguarded and, if transmitted over the Internet, they would need to be secured through VPN tunnels.

Remote-Access VPN Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the remote-access VPN?	Hub-and-spoke topology	The corporate LAN is the hub and each individual user becomes a spoke.
What type of tunneling will be deployed?	Encrypted tunnels using IPSec	
What type of security will be deployed?	RADIUS server Firewall	The RADIUS server authenticates remote users. The firewall augments a DSL router. It allows access to the Internet for the IP addresses specified in a filter list. Users who dial in will be given an IP address from a pool of addresses that are blocked from Internet access.
Is NAT required?	No	All addresses are public and the remote users are not allowed access to the Internet through the corporate network.
What VPN hardware will be used?	Dial-in access provided by an ISP Headquarters: Cisco 3005 VPN concentrator[1] VPN client for remote users	By allowing remote users to dial in to an ISP, the company avoids maintaining its own modem pool. As a restriction, none of the remote users can access the Internet while they are connected to the corporate LAN. At headquarters, there is no requirement for specialized hardware except that it has to connect to DSL and provide VPN termination with the ability to authenticate against a RADIUS server. Therefore, a Cisco 3005 VPN concentrator is used at the corporate site. Each remote system has a VPN client installed to allow for authentication against the RADIUS server over an encrypted IPSec connection.
What type of high availability will be deployed?	None, except services offered by the ISP	The only resiliency incorporated into the design is that the ISP offers the VPN service and has many numbers that users can call to access their global system. If the corporate network fails, there is no corporate access.

[1] The example platforms are accurate as of the date this course was published.

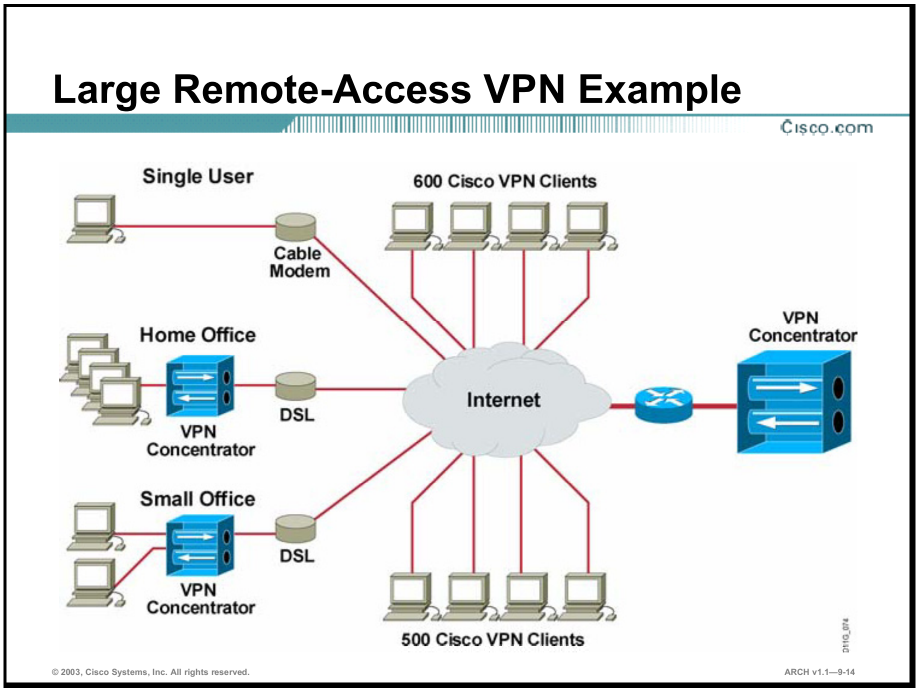

Large Remote-Access VPN Example

ARCH v1.1—9-14

Company Description

An old-fashioned cosmetics and cleaning material manufacturer is looking for a way to keep their sales force updated on the latest inventories, while utilizing the Internet as the basis of connectivity for their remote salespeople. They are enlisting the help of a nationwide service provider to provide access points for dial-in connectivity. Each remote user will be able to use the VPN to access the corporate network and then access the Internet from within the corporate backbone. The primary applications are e-mail and file sharing between users and corporate data devices.

The company has 1400 users in the local corporate campus and 1100 salespeople in the field. There are 40 small and home office combinations with fewer than 30 people at each site. The IT department wants to have only one VPN device at each of the fixed small and home office locations.

Remote-Access VPN Solution

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Notes and Comments
What topology will be used for the remote-access VPN?	Hub-and-spoke topology	The corporate LAN is the hub and each individual user becomes a spoke.
What type of tunneling will be deployed?	Encrypted tunnels with IPSec	
What type of security will be deployed?	RADIUS authentication Firewall with filters	Remote users will authenticate themselves as they connect to the corporate network. Each remote system has an installed VPN client to allow for authentication with the RADIUS server over the encrypted IPSec connection.
Is NAT required?	Yes	NAT is accomplished using the VPN concentrators.
What VPN hardware will be used?	■ Headquarters: Cisco 303x VPN concentrator ■ Remote fixed sites: Cisco 3002 VPN hardware client[1] ■ Remote sites: Cisco VPN client	The Cisco 3030 VPN concentrator at headquarters supports a RADIUS server for authentication. It uses hardware encryption and is upgradeable for more functionality. The 40 fixed sites utilize the Cisco 3002 VPN hardware client. The 3002 VPN hardware client offers the following features: ■ Includes DHCP Client and a DHCP Server for up to 253 stations behind the 3002 ■ Support for Network Address Port Translation for hiding stations behind 3002 ■ Optional 8-port 10/100 Mbps switch ■ Supports client mode as well as network extension mode for application flexibility ■ Works with any operating system ■ Eliminates the need to add or support VPN applications on the PC or workstation ■ Seamless operation with existing applications The Cisco VPN client was selected for the remote devices connecting to the corporate LAN because of its ease of use and availability on a wide range of operating systems, including Windows, MAC, Solaris, and Linux.
What type of high availability will be deployed?	Redundant link from the corporate office to the service provider for failover.	There is no other redundancy planned for the network besides the redundant link. The ISP offers the VPN service and has many phone numbers available for access to their global system. If the corporate network fails, a second corporate connection comes online.

[1] The example platforms are accurate as of the date this course was published.

Summary

This topic summarizes the key points discussed in this lesson.

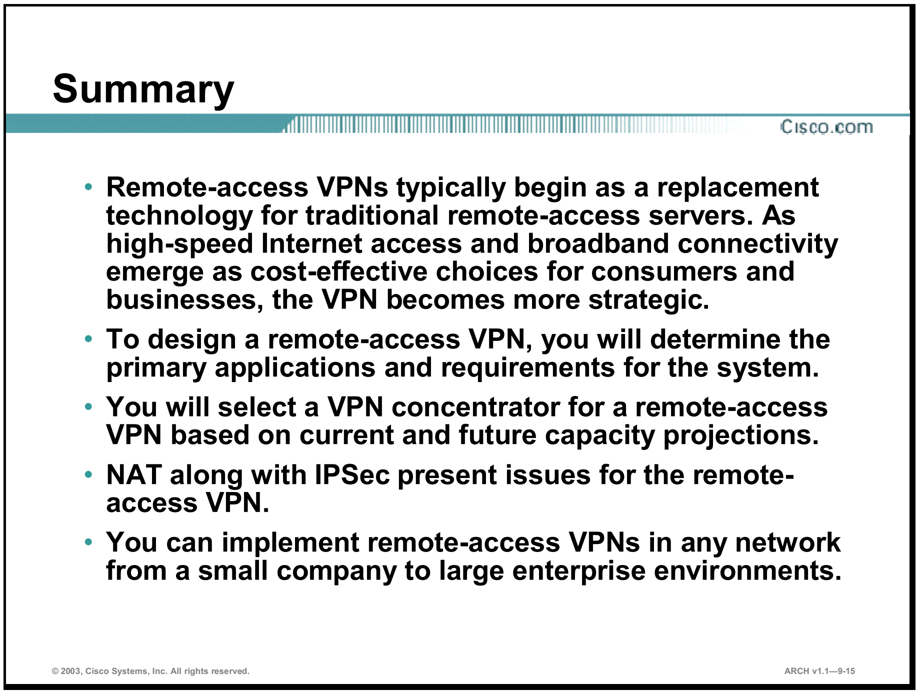

Summary

Cisco.com

- Remote-access VPNs typically begin as a replacement technology for traditional remote-access servers. As high-speed Internet access and broadband connectivity emerge as cost-effective choices for consumers and businesses, the VPN becomes more strategic.
- To design a remote-access VPN, you will determine the primary applications and requirements for the system.
- You will select a VPN concentrator for a remote-access VPN based on current and future capacity projections.
- NAT along with IPSec present issues for the remote-access VPN.
- You can implement remote-access VPNs in any network from a small company to large enterprise environments.

ARCH v1.1—9-15

References

For additional information, refer to these resources:

- *A Primer for Implementing a Cisco Virtual Private Network* at http://www.cisco.com/warp/public/cc/so/neso/vpn/vpne/vpn21_rg.htm
- *SAFE VPN: IPSec Virtual Private Networks in Depth* at http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safev_wp.htm
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 9-3: OCSIC Bottling Company
- OPNET IT Guru Solution 9-3

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) What are two reasons for implementing a remote-access VPN solution? (Choose two.)

A) to reduce staffing needs

B) to reduce cost of ownership

C) to increase cost of ownership

D) to replace a site-to-site WAN

E) to reduce corporate application traffic

Q2) Which three questions are relevant to designing remote VPN access? (Choose three.)

A) Is remote access the main focus?

B) What type of routing protocol will be used?

C) What type of NIC card is in the user device?

D) Which VPN tunneling protocol will be used?

Q3) When planning capacity for a remote-access VPN, which question should you ask?

A) What is the current bandwidth of the links to a remote office?

B) What is an estimate for the total number of users who plan to take advantage of the remote-access VPN solution in the initial phase of deployment?

C) What is an estimate for the number of intermittent users who do not plan to connect to the remote-access VPN solution in the initial phase of deployment?

D) Is the peak number of simultaneous users accessing the VPN concentrator expected to reach 50, or is the estimated amount of user traffic expected to exceed 2 Mbps in the near future?

Q4) To operate correctly, NAT needs access to _____ in the transmitted packets that does not exist with IPSec packets.

A) the data

B) the IP address

C) the source address

D) a unique identifier per session

Q5) Which function does NAT-traversal (NAT-T) provide?

A) encapsulates IPSec in TCP

B) encapsulates IPSec in UDP

C) reverses NAT port information

D) bypasses NAT devices to retain port information

Q6) Suppose that a company incorporates resiliency into its design using a service from a service provider. What would happen if the corporate network fails?

A) There is no corporate access.

B) The tunnel to the ISP is no longer secure.

C) There will be no authentication for users.

D) Users dial in through a second service provider.

Q7) You are designing a large remote-access VPN solution. There may be a future need to deploy some small site-to-site locations, as well. All traffic will be strictly based on an IP. Which back-end device is the best choice to meet the requirements of the solution?

A) Cisco PIX firewall

B) Cisco VPN concentrator

C) Cisco VPN-enabled router

D) Cisco router with VPN and firewall feature set

Quiz Answer Key

Q1) A, B

Relates to: Remote-Access VPN Requirements

Q2) A, B, D

Relates to: Remote-Access VPN Design Considerations

Q3) B

Relates to: Capacity Planning for Remote-Access VPNs

Q4) D

Relates to: Network Address Translation Issues

Q5) B

Relates to: Network Address Translation Issues

Q6) A

Relates to: Remote-Access VPN Examples

Q7) B

Relates to: Remote-Access VPN Examples

Case Study 9-3: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

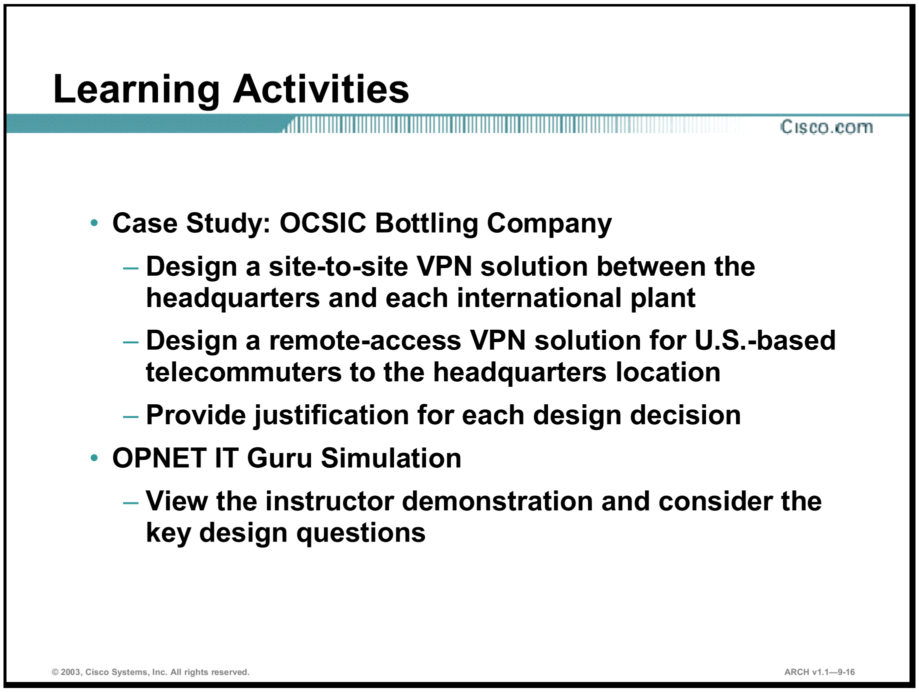

Learning Activities

- Case Study: OCSIC Bottling Company
  - Design a site-to-site VPN solution between the headquarters and each international plant
  - Design a remote-access VPN solution for U.S.-based telecommuters to the headquarters location
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—9-16

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

The OCSIC Bottling Company international offices require an always-on connection to the headquarters to share data and e-mail. However, the cost of a dedicated connection is too high. Each international plant requires a connection to their local ISP, and a VPN for secure communications over the Internet to the headquarters office.

In this exercise, you will design VPNs that meet the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design a site-to-site VPN solution between the headquarters and each international plant
- Design a remote-access VPN solution for U.S.-based telecommuters to the headquarters location

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Design a Site-to-Site VPN Solution Between the Headquarters and Each International Plant

Complete these steps:

Step 1 Complete the table to design your site-to-site VPN solution.

Design Questions	Decision	Justification
What topology will be used for the site-to-site VPN?		
What type of tunneling will be deployed?		
What type of security will be deployed?		
Is NAT required?		
What VPN hardware will be used?		
What type of high availability will be deployed?		

Step 2 Update your global network diagram to reflect your VPN strategy.

Task 2: Design a Remote-Access VPN Solution for U.S.-Based Telecommuters to the Headquarters Location

Complete these steps:

Step 1 Complete the table to design your remote-access VPN solution.

Design Questions	Decision	Justification
What topology will be used for the remote-access VPN?		
What type of tunneling will be deployed?		
What type of security will be deployed?		
Is NAT required?		
What VPN hardware will be used?		
What type of high availability will be deployed?		

Step 2 Update your global network diagram to reflect your VPN strategy.

Task 3: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have designed a site-to-site VPN solution between the headquarters' site and the international locations. Your VPN solution identifies the data link and physical layer protocols, head-end and remote-end termination solutions, VPN client software, VPN security and tunneling, and NAT. Your global network diagram shows the VPN solution.
- You have designed a remote-access VPN solution between telecommuters and the headquarters' site. Your VPN solution identifies the data link and physical layer protocols, head-end and remote end termination solutions, VPN client software, VPN security and tunneling, and NAT. Your global network diagram shows the VPN solution.

OPNET IT Guru Simulation 9-3

This simulation demonstrates the affect of a VPN on network and application performance.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 10

Designing Enterprise Wireless Networks

Overview

Simply put, a wireless local-area network (WLAN) does not require wires. It uses radio frequencies instead of a traditional media such as copper or fiber. It provides real-time access to the network. With a WLAN, the covered area is essentially comparable to that of a wired LAN. In large applications, by deploying a series of access points throughout a building or campus, enterprises can achieve more coverage than a wired network, with the benefits of high-speed data rates and the freedom of mobility from being able to access broadband data anywhere within the WLAN.

Module Objectives

Upon completing this module, you will be able to design enterprise solutions for wireless networks, given enterprise network needs.

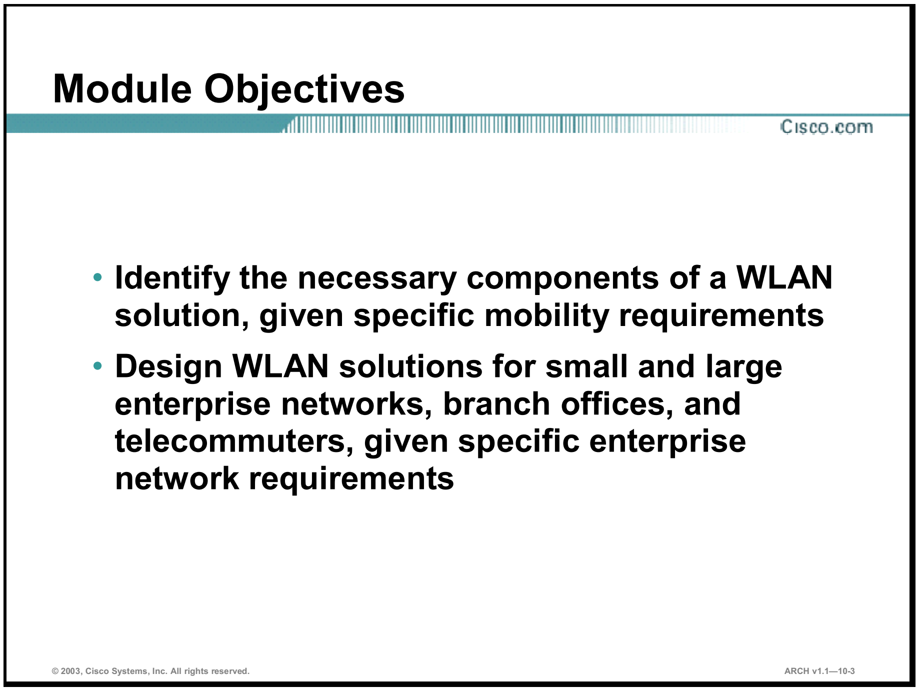

Module Objectives

- Identify the necessary components of a WLAN solution, given specific mobility requirements
- Design WLAN solutions for small and large enterprise networks, branch offices, and telecommuters, given specific enterprise network requirements

ARCH v1.1—10-3

Module Outline

The outline lists the components of this module.

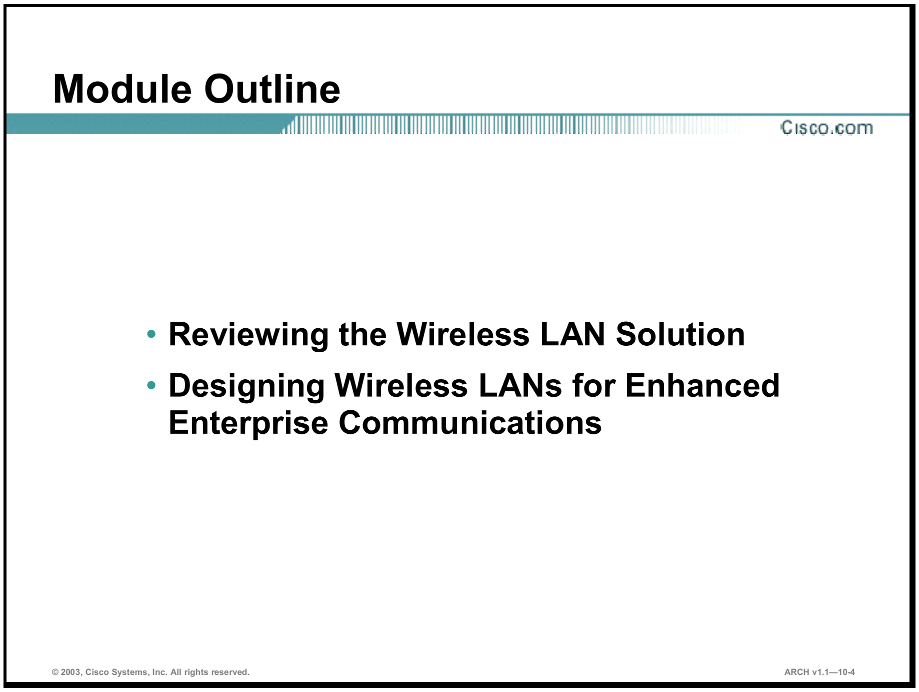

Module Outline

- Reviewing the Wireless LAN Solution
- Designing Wireless LANs for Enhanced Enterprise Communications

ARCH v1.1—10-4

Reviewing the Wireless LAN Solution

Overview

Wireless local-area networks (WLANs) enable network designers to establish and maintain a wireless network connection throughout or between buildings, without the limitations of wires or cables. Cisco provides a family of WLAN products that combine the mobility and flexibility users want from a wireless system with the throughput and security they get from a wired LAN.

Relevance

Mobility and ease of installation have made WLAN technology a key technology in markets such as health care, education, and retail. WLANs are also making inroads in general business environments.

Objectives

Upon completing this lesson, you will be able to identify the necessary components of a WLAN solution, given specific mobility requirements. This includes being able to meet these objectives:

- Describe how WLANs can meet enterprise requirements
- Compare the WLAN architecture to a typical wired LAN configuration
- Describe the 802.11 standards and their differences
- Identify the necessary components of a wireless solution, given specific mobility requirements

Learner Skills and Knowledge

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

- Basic knowledge of local-area networking
- Basic knowledge of 802.11 wireless standards

Outline

The outline lists the topics included in this lesson.

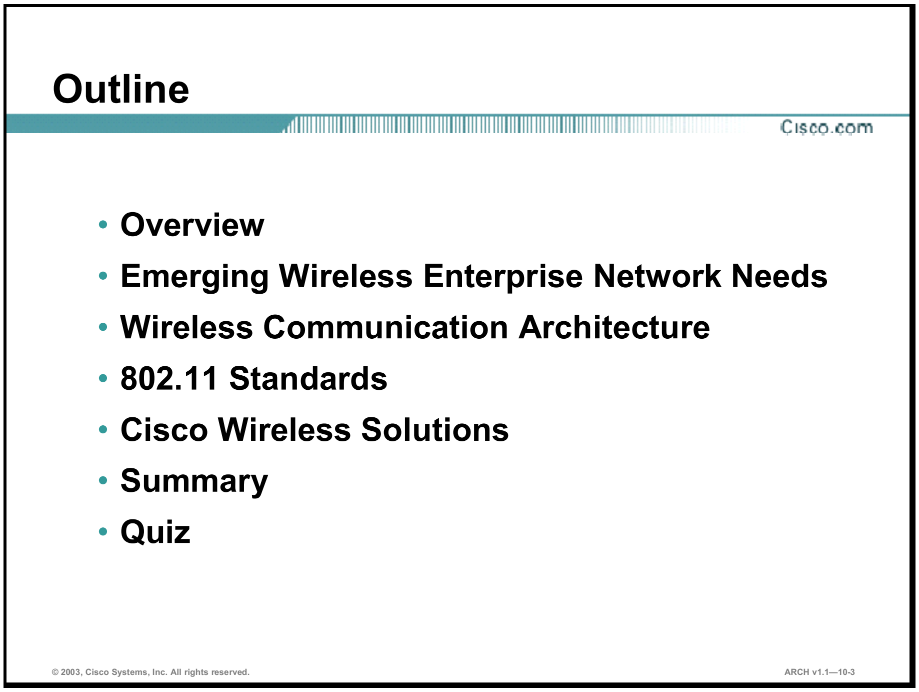

Outline

- Overview
- Emerging Wireless Enterprise Network Needs
- Wireless Communication Architecture
- 802.11 Standards
- Cisco Wireless Solutions
- Summary
- Quiz

ARCH v1.1—10-3

Emerging Wireless Enterprise Network Needs

A WLAN implementation adds a mobile component to the traditional office LAN and provides LAN coverage in areas where cabling is impractical or inefficient. This topic describes how WLANs can meet enterprise requirements.

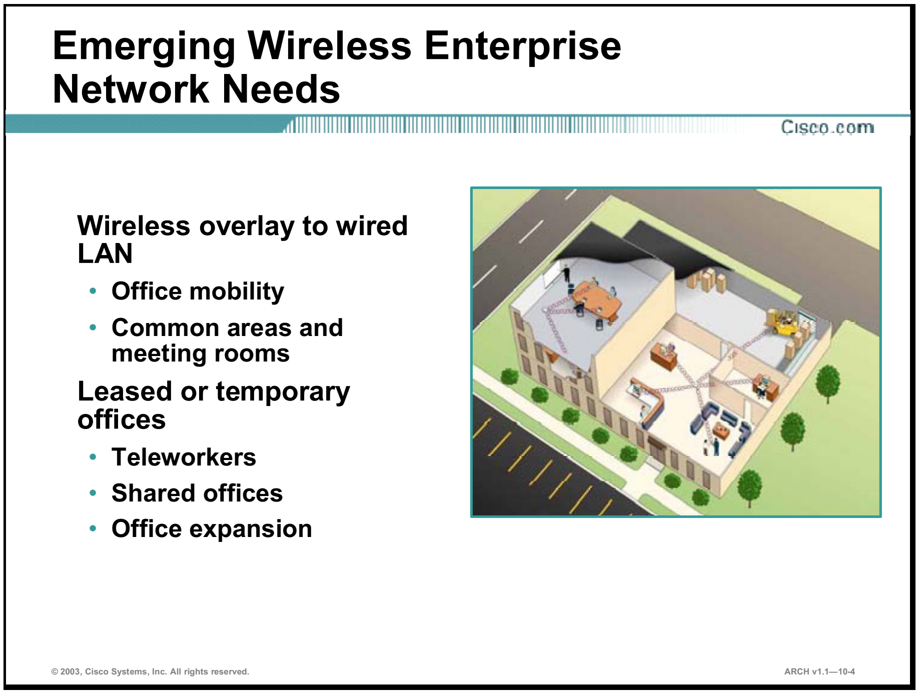

Emerging Wireless Enterprise Network Needs

Cisco.com

Wireless overlay to wired LAN

- **Office mobility**
- **Common areas and meeting rooms**

Leased or temporary offices

- **Teleworkers**
- **Shared offices**
- **Office expansion**

ARCH v1.1—10-4

In a high-performance, switched environment, wireless technology can deliver Ethernet-level speeds to open areas on the campus, or high-density areas like auditoriums and conference rooms that require network access for a large number of users. Typically, a WLAN does not replace the wired LAN. It is typically an adjunct or overlay to the wired LAN. WLAN technology enables deployment of LANs in offices where Category 5 cabling may not be cost effective or timely. The WLAN adds a nomadic or mobile dimension to the traditional office LAN. The WLAN also facilitates connectivity in meeting rooms, cafeterias, and other common areas where wired LANs are impractical.

Wireless Communication Architecture

A WLAN consists of an access point communicating over radio frequency to wireless clients. The data rate, power level, and antenna choice affect the size of the coverage area of a single wireless cell, which in turn affects how many access points are required in a specific implementation. This topic describes the wireless communication architecture.

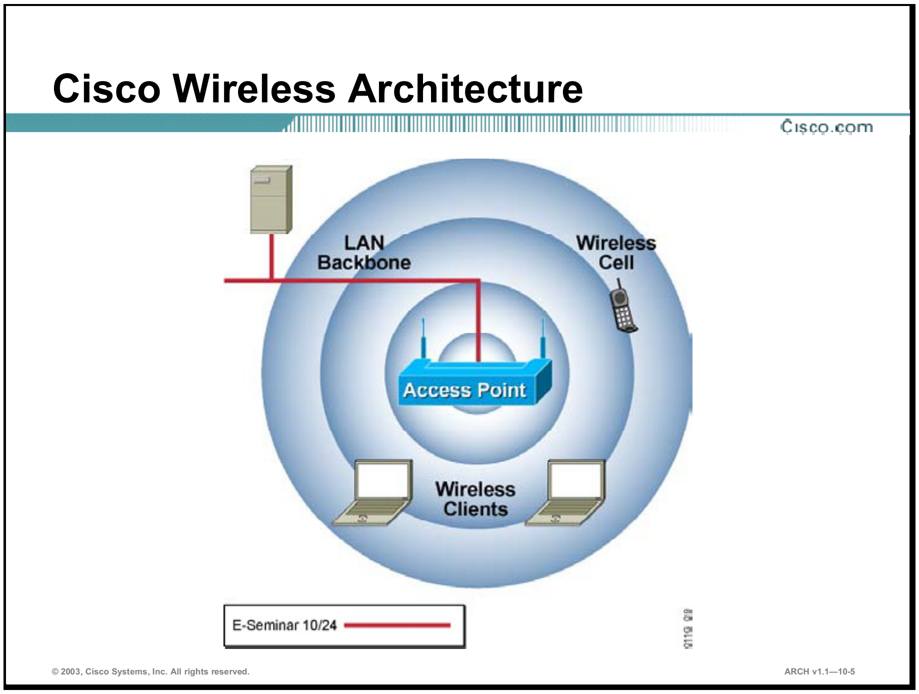

Cisco Wireless Architecture

ARCH v1.1—10-5

Only one station in a wireless cell, including the access point, can send data at any one time. The bandwidth is shared among all stations. If a station wishes to send, it listens and waits for an available slot. WLANs use carrier sense multiple access collision avoidance (CSMA/CA).

The protocols used in a wireless network cover the physical and data link layers. Therefore, a wireless LAN can transport a variety of LAN and network layer protocols such as IP, AppleTalk, NetBIOS Extended User Interface (NetBEUI), and so on.

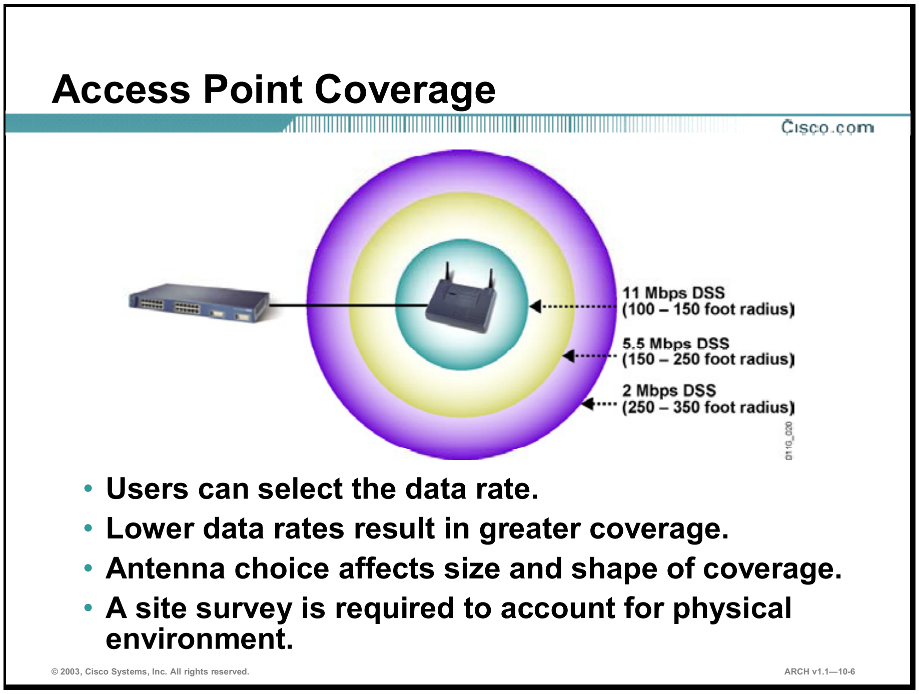

Access Point Coverage

- Users can select the data rate.
- Lower data rates result in greater coverage.
- Antenna choice affects size and shape of coverage.
- A site survey is required to account for physical environment.

ARCH v1.1—10-6

Data rates affect cell size. Lower data rates can extend further from the access point than can higher data rates. Hence the data rate (and power level) will affect cell coverage and the number of access points required.

The factors that affect the coverage are as follows:

- Selected data rate (1, 2, 5.5, 11 Mbps)
- Power level, as follows (maximum power setting will vary according to individual country regulations):
  - 100 mW (20 dBm)
  - 50 mW (17 dBm)
  - 30 mW (15 dBm)
  - 20 mW (13 dBm)
  - 5 mW (7 dBm)
  - 1 mW (0 dBm)
- Antenna choice (dipole, omni, wall mount)

For a given data rate, the WLAN designer can alter the power level or choose a different antenna to change the coverage area and/or coverage shape.

Access points have an aggregate throughput of about 6 Mbps. With this in mind, the maximum suggested number of active clients is between 10 and 30. The precise number of active clients depends on the data rates supported. That is, active clients with higher data rates necessitate fewer active clients for each access point.

Cisco wireless products support 11-Mbps communications with an indoor range (within a building) of 40 meters (130 feet) and outdoor range of 244 meters (800 feet). However, as the distance between the wireless station and the access point grows, it becomes necessary to lower transmission speed to maintain channel quality. To avoid loss of data for mobile users, Cisco wireless products automatically switch to lower speeds, enabling users to maintain a connection in open areas of up to 610 meters (2000 feet) from the access point.

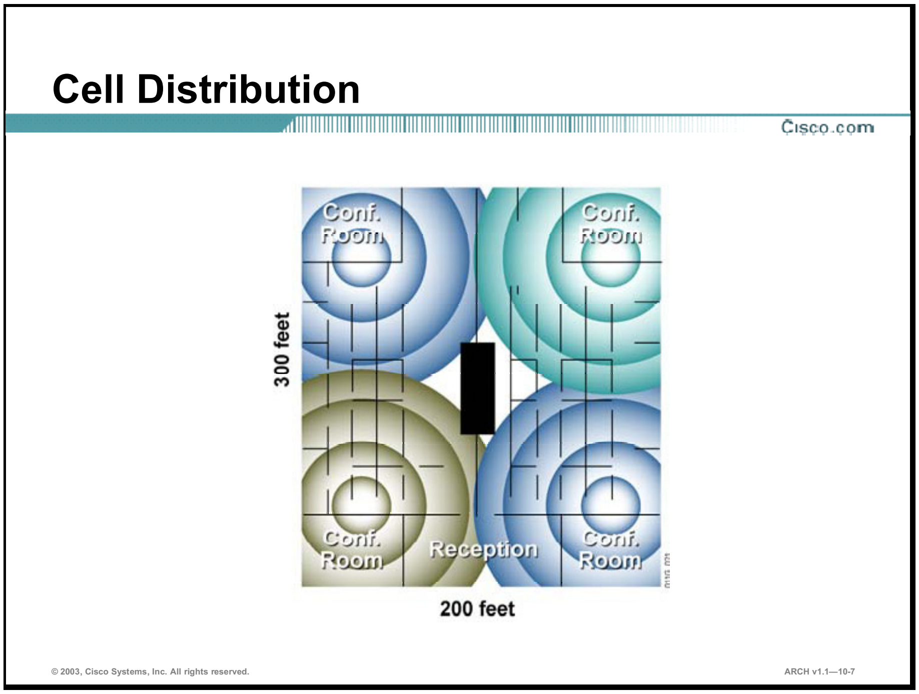

Cell Distribution

Cisco.com

ARCH v1.1—10-7

A large cell size may lead to too many clients sharing the available bandwidth. By reducing the access point power or antenna gain, you can reduce the cell size and share it with fewer clients. This will result in more access points for a given coverage area, but will provide better and more equitable performance for clients.

802.11 Standards

The IEEE 802.11 standard is a group of protocol specifications for WLANs. This topic describes the IEEE 802.11 standards and their differences.

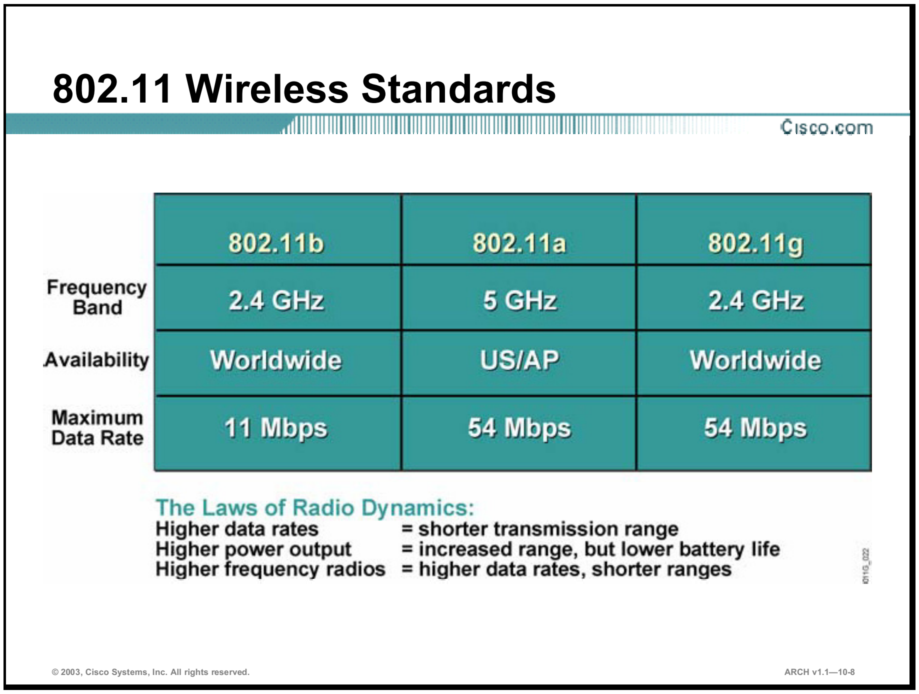

802.11 Wireless Standards

Cisco.com

	802.11b	802.11a	802.11g
Frequency Band	2.4 GHz	5 GHz	2.4 GHz
Availability	Worldwide	US/AP	Worldwide
Maximum Data Rate	11 Mbps	54 Mbps	54 Mbps

The Laws of Radio Dynamics:

Higher data rates	= shorter transmission range
Higher power output	= increased range, but lower battery life
Higher frequency radios	= higher data rates, shorter ranges

ARCH v1.1—10-8

The 802.11 standard is a group of specifications for WLANs created by the IEEE. The 802.11 standards define the communication protocols between wireless workstations and the network access points that bridge wireless and wired networks. The original 802.11 standard specified support for 1-Mbps and 2-Mbps peak transfer rates. Subsequent advances in the 802.11 standard have greatly increased wireless transfer rates. The High Rate amendment to the 802.11 standard, 802.11b, added 5.5-Mbps and 11-Mbps transmission speeds.

The next evolution in the standard is 802.11a, which supports transmission speeds of 36 Mbps, 48 Mbps, and 54 Mbps, and ultimately 108 Mbps, utilizing dual channels and operating on the 5-GHz Unlicensed National Information Infrastructure (U-NII) band. The 802.11a standard is not compatible with the 802.11b standard.

The 802.11g standard operates in the same unlicensed portion of the 2.4-GHz spectrum as 802.11b. Both the IEEE 802.11g and 802.11a standards provide a 54-Mbps data rate. IEEE 802.11g provides the benefit of backward compatibility with IEEE 802.11b equipment, preserving users' investment in their existing WLAN infrastructure. In addition, because it builds on 802.11b technology, 802.11g will take advantage of the years of 802.11b silicon integration and the resulting reduction in power consumption, form factor size, and cost. However, because 802.11g is limited to the same three channels as 802.11b, scalability may become a factor as WLAN user density increases.

Cisco Wireless Solutions

The Cisco wireless solution includes access points, client adapters, workgroup bridges, bridges, antennas, and accessories. This topic describes the components of a Cisco wireless solution.

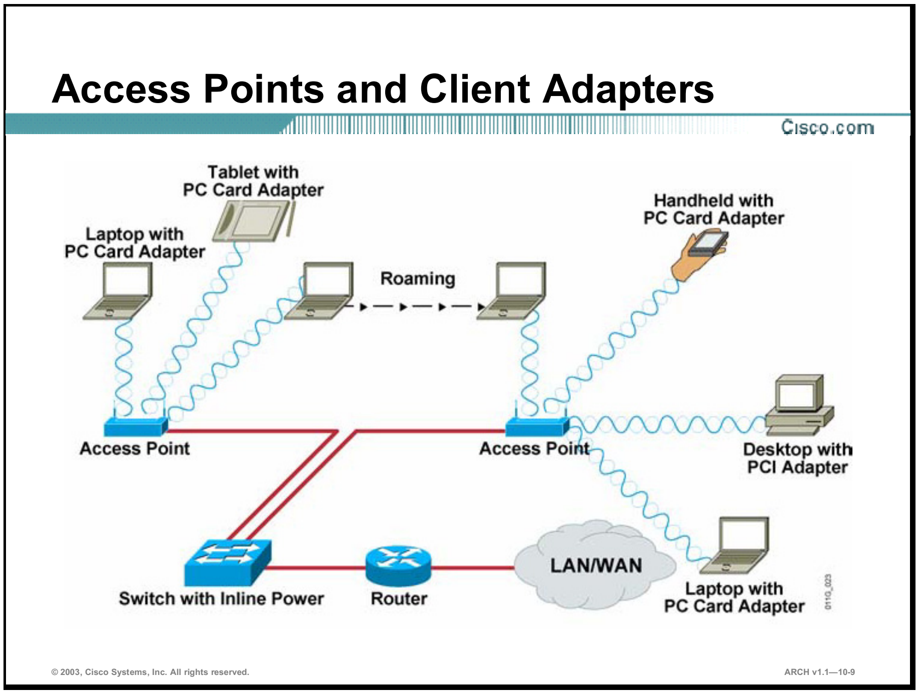

Access Points and Client Adapters

Cisco.com

Tablet with PC Card Adapter
Laptop with PC Card Adapter
Roaming
Handheld with PC Card Adapter
Access Point
Access Point
Desktop with PCI Adapter
Switch with Inline Power
Router
LAN/WAN
Laptop with PC Card Adapter

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—10-9

An access point is the center point in an all-wireless network, or serves as a connection point between a wired and wireless network. You can place multiple access points throughout a facility to give users with WLAN client adapters the ability to roam freely throughout an extended area while maintaining uninterrupted access to all network resources.

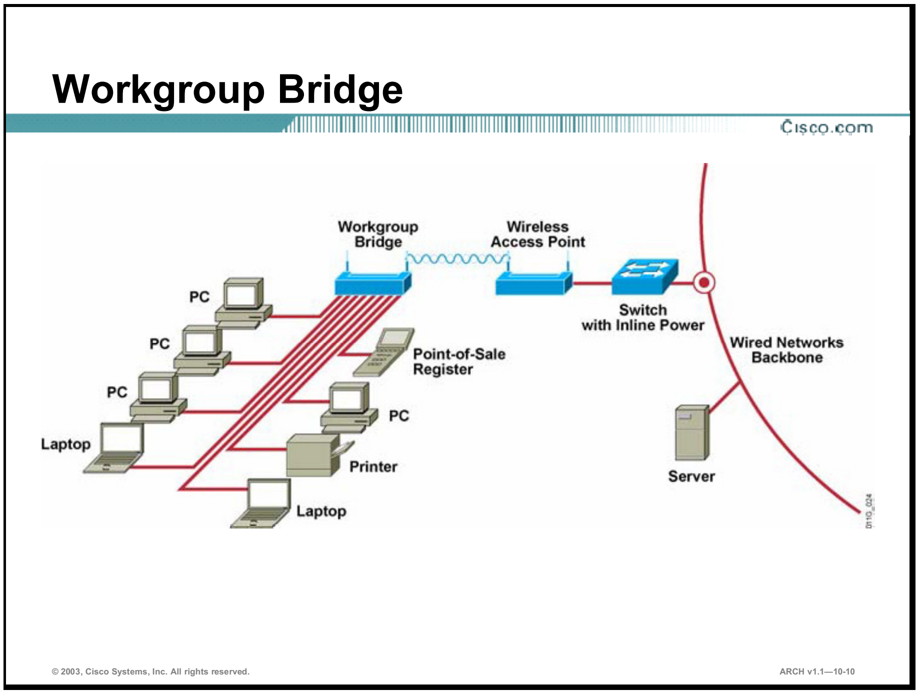

Workgroup Bridge

Cisco.com

ARCH v1.1—10-10

A workgroup bridge provides wireless connectivity to an Ethernet-enabled device. A workgroup bridge connects Ethernet-enabled laptops or other portable computers to a WLAN, providing the link from these devices to an access point or wireless bridge.

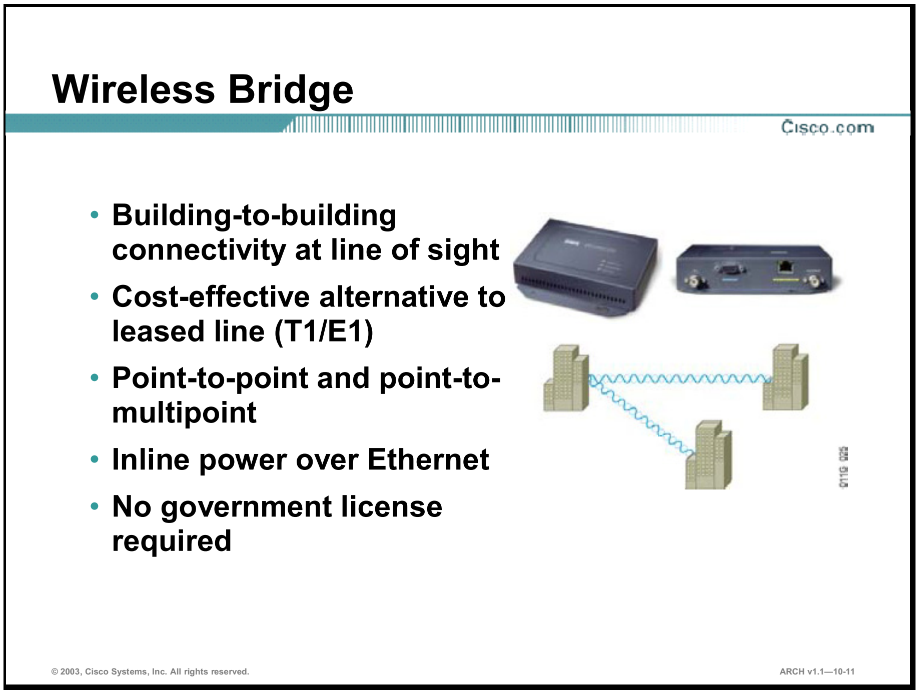

Wireless Bridge

Cisco.com

- Building-to-building connectivity at line of sight
- Cost-effective alternative to leased line (T1/E1)
- Point-to-point and point-to-multipoint
- Inline power over Ethernet
- No government license required

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—10-11

Bridging provides a means to connect two or more remote Ethernet segments over a wireless connection.

A wireless bridge is capable of communicating over greater distances than 802.11b access points and clients. Bridges do this by stretching the timing constraints 802.11 puts on the return times for packet acknowledgments. This alteration of the timing values violates 802.11 specifications.

With the appropriate antennas, a bridge in access-point mode can communicate with clients in the line of sight. In contrast, access points are compliant with 802.11 and are subject to timing characteristics that impose a maximum distance limitation of one mile for communication with clients at any speed. The actual distances that access points achieve varies up to one mile depending on the antennas used.

Note You should plan to work with an RF designer to determine appropriate antennas for a WLAN installation.

Summary

This topic summarizes the key points discussed in this lesson.

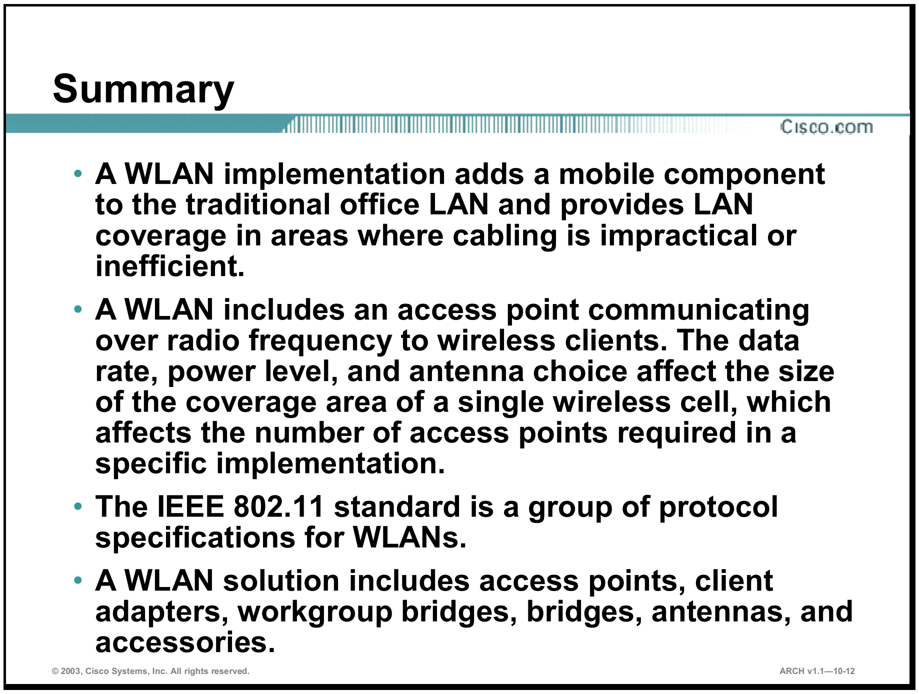

Summary

Cisco.com

- A WLAN implementation adds a mobile component to the traditional office LAN and provides LAN coverage in areas where cabling is impractical or inefficient.
- A WLAN includes an access point communicating over radio frequency to wireless clients. The data rate, power level, and antenna choice affect the size of the coverage area of a single wireless cell, which affects the number of access points required in a specific implementation.
- The IEEE 802.11 standard is a group of protocol specifications for WLANs.
- A WLAN solution includes access points, client adapters, workgroup bridges, bridges, antennas, and accessories.

ARCH v1.1—10-12

References

For additional information, refer to these resources:

- *Wireless Solutions* at http://www.cisco.com/warp/public/44/solutions/network/wireless.shtml
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

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) In a practical enterprise WLAN implementation, what relationship does the WLAN have to the existing wired LAN?

A) The WLAN would extend the wired LAN.

B) The WLAN would replace the wired LAN.

C) The WLAN would duplicate the wired LAN.

D) The WLAN would be completely separate from the wired LAN.

Q2) Why is the number of clients per access point on a WLAN limited?

A) because stations cannot avoid collisions

B) because all clients must share the available bandwidth

C) because each client must maintain a minimum distance from the other clients

D) because each access point can only communicate with a specific number of clients

Q3) Which two characteristics differentiate 802.11a from both 802.11b and 802.11g? (Choose two.)

A) availability

B) power output

C) coverage area

D) frequency band

E) minimum data rate

Q4) What is a significant drawback to the 802.11a standard?

A) It has a slow data rate.

B) It requires additional equipment.

C) It requires government licensing.

D) It is not compatible with the 802.11b.

Q5) What device serves as the connection between the wireless network and the wired network?

A) access point

B) client adapter

C) wireless bridge

D) workgroup bridge

Quiz Answer Key

Q1) A

Relates to: Emerging Wireless Enterprise Network Needs

Q2) B

Relates to: Wireless Communication Architecture

Q3) A, D

Relates to: 802.11 Standards

Q4) D

Relates to: 802.11 Standards

Q5) A

Relates to: Cisco Wireless Solutions

Designing Wireless LANs for Enhanced Enterprise Communications

Overview

WLAN is generally deployed in an enterprise campus or branch office for increased efficiency and flexibility. WLANs are emerging as an effective method to connect to an enterprise network. They are an access technology intended for LAN implementations. This lesson presents design recommendations for the WLAN infrastructure in enterprises.

Relevance

By addressing common deployment schemes, an enterprise WLAN will provide mobility within a building or site, convenience, and flexibility.

Objectives

Upon completing this lesson, you will be able to design WLAN solutions for small and large enterprise networks, branch offices, and telecommuters, given specific enterprise network requirements. This includes being able to meet these objectives:

- Identify the design considerations for an enterprise WLAN solution
- Describe the WLAN security extensions available and the differences between them
- Design WLAN solutions for small enterprise networks, given specific network requirements
- Design WLAN solutions for large enterprise networks, given specific network requirements
- Design WLAN solutions for remote access and telecommuters, given specific network requirements

Learner Skills and Knowledge

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

- Basic knowledge of local-area networking
- Basic knowledge of 802.11 wireless standards

Outline

The outline lists the topics included in this lesson.

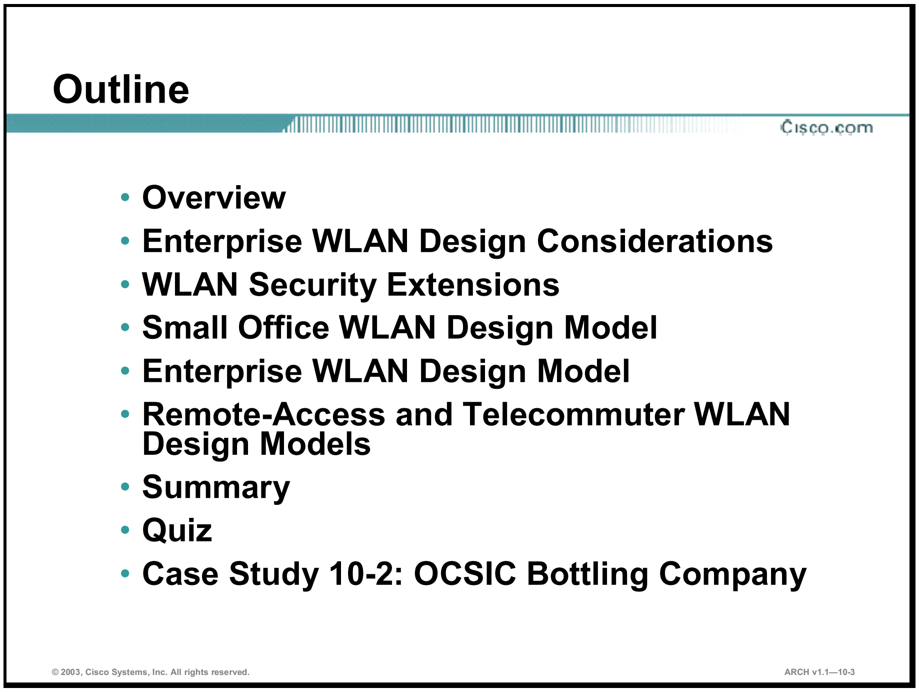

Outline

- Overview
- Enterprise WLAN Design Considerations
- WLAN Security Extensions
- Small Office WLAN Design Model
- Enterprise WLAN Design Model
- Remote-Access and Telecommuter WLAN Design Models
- Summary
- Quiz
- Case Study 10-2: OCSIC Bottling Company

ARCH v1.1—10-3

Enterprise WLAN Design Considerations

When designing an enterprise wireless network solution, you must consider the radio frequency (RF) design, the campus infrastructure, high availability, roaming, IP multicast, and quality of service (QoS). This topic identifies the design considerations for an enterprise WLAN solution.

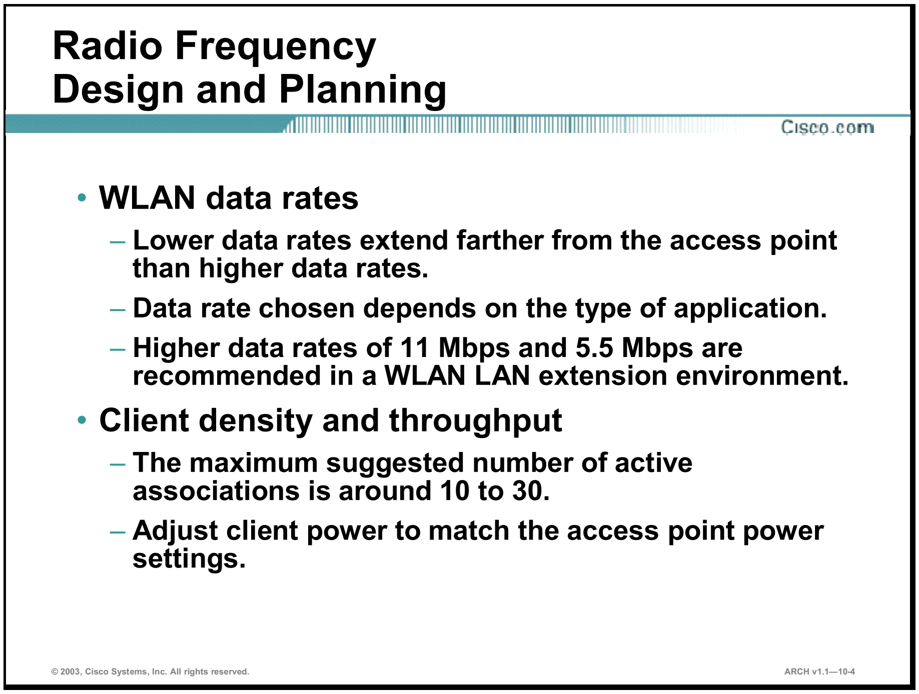

Radio Frequency Design and Planning

Cisco.com

- **WLAN data rates**
  - **Lower data rates extend farther from the access point than higher data rates.**
  - **Data rate chosen depends on the type of application.**
  - **Higher data rates of 11 Mbps and 5.5 Mbps are recommended in a WLAN LAN extension environment.**
- **Client density and throughput**
  - **The maximum suggested number of active associations is around 10 to 30.**
  - **Adjust client power to match the access point power settings.**

ARCH v1.1—10-4

WLAN Data Rates

Data rates affect cell size. Lower data rates (such as 1 Mbps) can extend farther from the access point than can higher data rates (such as 11 Mbps). Therefore, the data rate (and power level) affects cell coverage and consequently the number of access points required.

The required data rate has a direct impact upon the number of access points required for the design. While six access points with a data rate of 2 Mbps might adequately service an area, it might take twice as many access points to support a data rate of 5 Mbps, and more again to support data rates of 11 Mbps.

The data rate chosen is dependent on the type of application to be supported. In a WLAN LAN extension environment, the higher data rates of 11 Mbps and 5.5 Mbps are recommended. This gives maximum throughput and should minimize performance-related support issues. In a WLAN environment, the data rates selected are determined by the application requirements. Some clients might not support the higher data rates and might require the use of lower data rates.

It might seem logical to choose the default configuration of access points and clients, thereby allowing all data rates. However, there are three key reasons for limiting the data rate to the highest rate at which full coverage is obtained:

- Broadcast and multicast are sent at the slowest data rate (to ensure that all clients can see them), which reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.

- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced.
- If an 11-Mbps service has been specified and provisioned with access points to support this level of service, allowing clients to associate at lower rates will create a coverage area greater than planned, increasing the security exposure and potentially interfering with other WLANs.

Client Density and Throughput

Access points have an aggregate throughput of approximately 6 Mbps. Therefore, the maximum suggested number of active associations (active clients) is around 10 to 30. You can adjust this number depending on the particular application.

A large cell size can result in an overloading of available capacity with too many clients sharing network access via the same access point. By reducing the access point power or antenna gain, you can reduce the cell size and share it among fewer clients. More access points will be required for a given coverage area, but clients receive better performance.

Client power should be adjusted to match the access point power settings. Maintaining a high setting on the client does not result in higher performance, and it can cause interference in nearby cells.

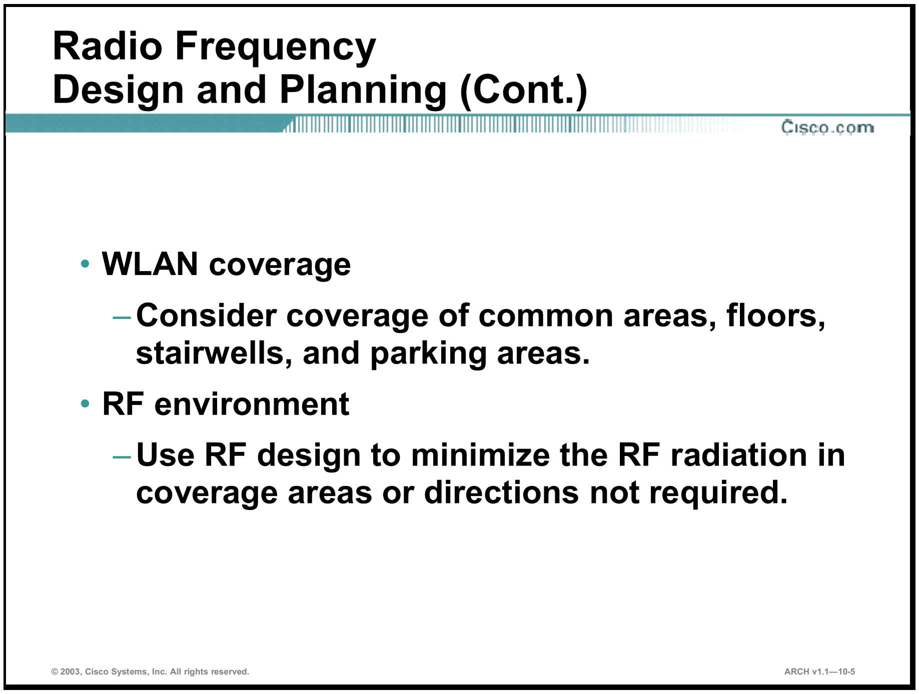

Radio Frequency
Design and Planning (Cont.)

Cisco.com

- WLAN coverage
  - Consider coverage of common areas, floors, stairwells, and parking areas.
- RF environment
  - Use RF design to minimize the RF radiation in coverage areas or directions not required.

ARCH v1.1—10-5

WLAN Coverage

Different enterprises have different requirements. Some need a WLAN to cover specific common areas. Others need WLANs to cover each floor of a building, to cover the entire building including stairwells and elevators, or to cover the entire campus including parking areas and roads.

Apart from impacting the number of access points required, the coverage requirements can introduce other issues, such as specialized antennas, outdoor enclosures, and lightning protection.

RF Environment

You can use RF design to minimize the RF radiation in coverage areas or directions not required. For example, if WLAN coverage is required only in the buildings, then you can minimize the amount of RF coverage outside the building through access-point placement and directional antennas.

The performance of the WLAN and its equipment depends upon its RF environment. Some examples of variables that can adversely affect RF performance are:

- 2.4-GHz cordless phones
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Transformers
- Heavy-duty electric motors
- Fire walls and fire doors
- Concrete
- Refrigerators

- Sulphur plasma lighting (Fusion 2.4-GHz lighting systems)
- Air conditioning ductwork
- Other radio equipment
- Microwave ovens
- Other WLAN equipment

You should perform a site survey to ensure that the required data rates are supported in all the required areas, despite the environmental variables.

The site survey should consider the three-dimensional space occupied by the WLAN. For example, a multistory building WLAN with different subnets per floor might require a different RF configuration than the same building with a single WLAN subnet per building. In the multiple-subnet instance, a client attempting to roam to a different access point on the same floor might acquire an access point from an adjacent floor. Switching access points in a multisubnet environment would change the roaming activity from a seamless data link layer roam to a network layer roam, which would in turn disrupt sessions and might require user intervention.

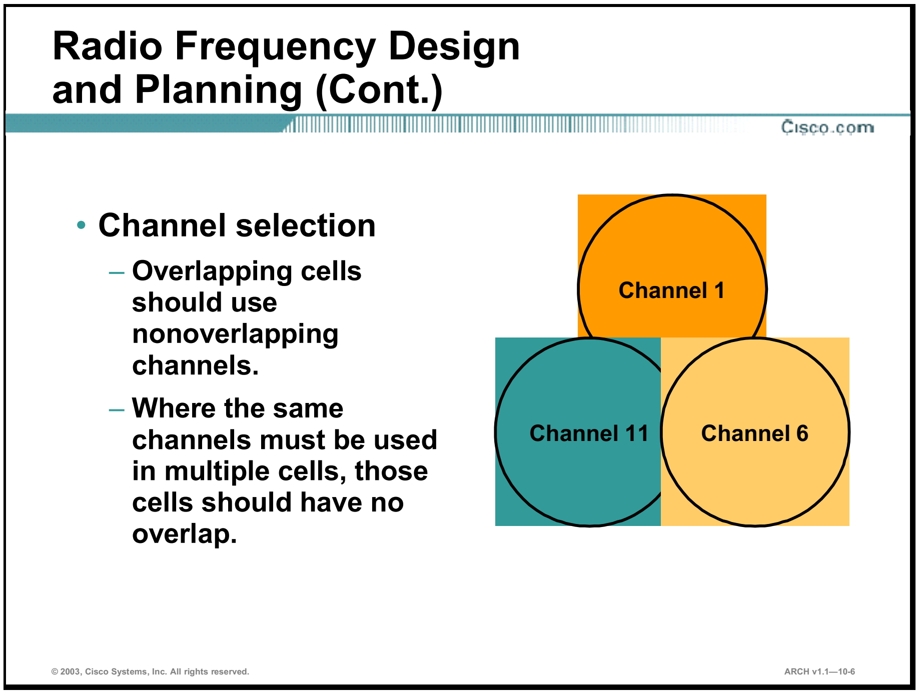

Radio Frequency Design and Planning (Cont.)

Cisco.com

- Channel selection
  - Overlapping cells should use nonoverlapping channels.
  - Where the same channels must be used in multiple cells, those cells should have no overlap.

Channel 1
Channel 11
Channel 6

ARCH v1.1—10-6

Channel Selection

Channel selection depends on the frequencies and channels permitted for the particular region. The North American and European Telecommunication Standards Institute (ETSI) channel sets allow the allocation of three nonoverlapping channels: 1, 6, and 11.

You should allocate the channels to the cells as follows:

- Overlapping cells should use nonoverlapping channels.
- Where the same channels are required in multiple cells, make sure those cells have no overlap.

In multistory buildings, check the cell overlap between floors according to the overlap guidelines. Some resurveying and relocating of access points might be required. Retest the site using the selected channels and check for interference.

You can configure an access point to automatically search for the best channel.

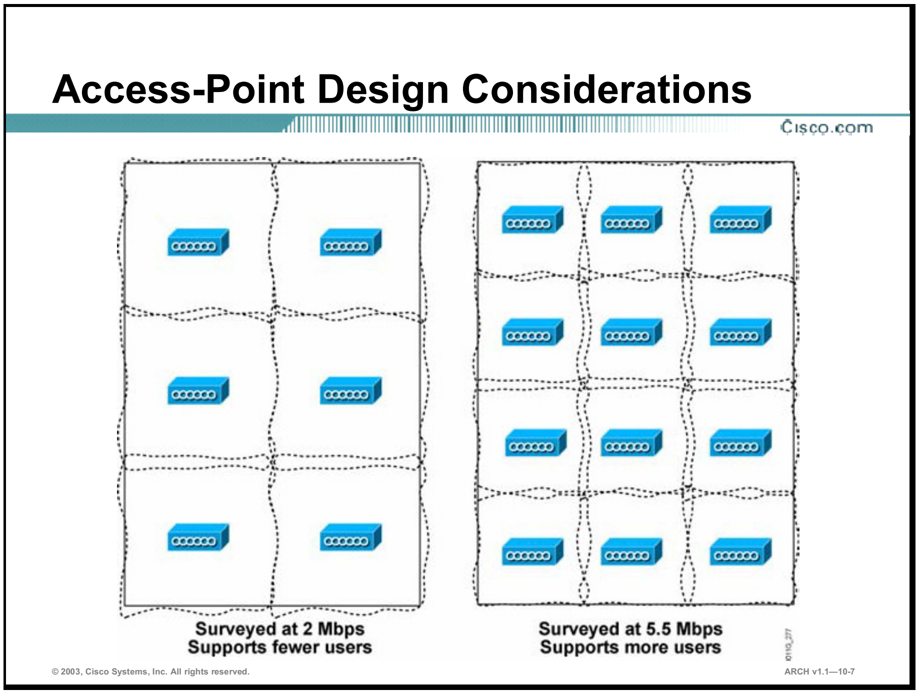

Access-Point Design Considerations

Cisco.com

Surveyed at 2 Mbps
Supports fewer users

Surveyed at 5.5 Mbps
Supports more users

ARCH v1.1—10-7

Access-Point Placement and Number

The required data rate has a direct impact on the number of access points needed in a design. The example shows this point. While six access points with a data rate of 2 Mbps might adequately service an area, it would take twice as many access points to support a data rate of 5 Mbps.

The data rate selected is dependent on the type of application to be supported. In a WLAN LAN extension environment, higher data rates of 11 Mbps and 5.5 Mbps are recommended to provide maximum throughput and to minimize performance-related support issues.

When choosing an access point and associated client cards, consider:

- Processor power
- Throughput requirements
- Inline power support
- Output power support

Choose the access point that best serves the application needs and then select the appropriate client cards and accessories.

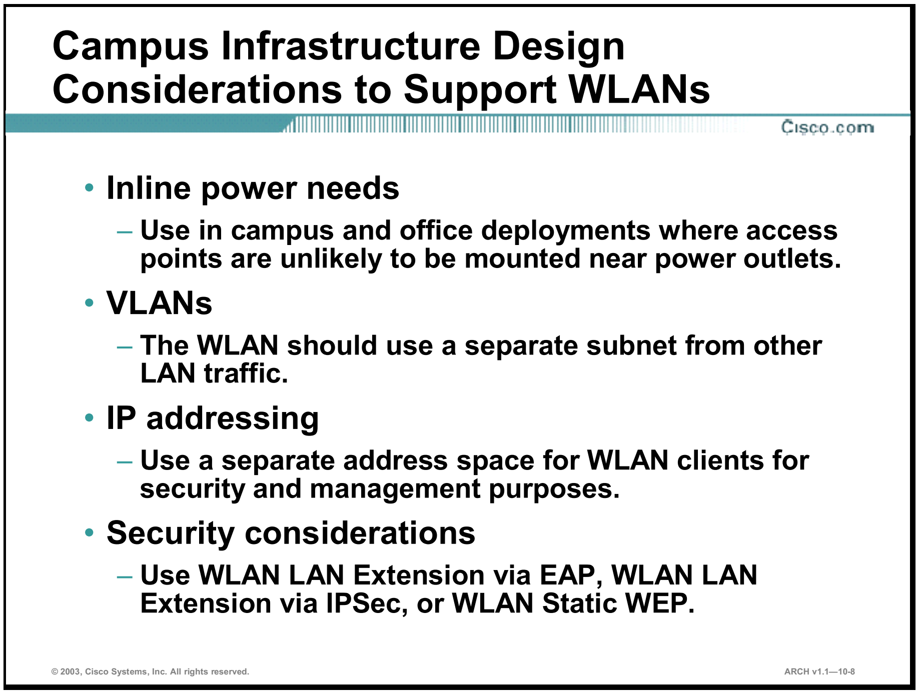

Campus Infrastructure Design Considerations to Support WLANs

- **Inline power needs**
  - **Use in campus and office deployments where access points are unlikely to be mounted near power outlets.**
- **VLANs**
  - **The WLAN should use a separate subnet from other LAN traffic.**
- **IP addressing**
  - **Use a separate address space for WLAN clients for security and management purposes.**
- **Security considerations**
  - **Use WLAN LAN Extension via EAP, WLAN LAN Extension via IPSec, or WLAN Static WEP.**

ARCH v1.1—10-8

Inline Power

Inline power is particularly useful in campus and office deployments where access points are unlikely to be mounted near power outlets. Inline power eliminates the requirement for site customization to provide power outlets in ceilings or walls to support access points.

Power options include the following:

- A switch with inline power
- An inline power patch panel
- A Cisco Aironet power injector (used when inline power is not available)

VLANs

Whenever possible, the WLAN should be a separate subnet from other LAN traffic, for these reasons:

- To optimize overall network performance. The WLAN media is shared and therefore any unnecessary broadcast or multicast traffic can impair network performance.
- To clearly identify the source of traffic for management and security issues.
- To increase the number of WLAN clients on a single VLAN and increase the possible Layer 2 roaming domain.

When a WLAN is an overlay network extension, it is not expected that WLAN VLANs mirror the wired VLANs. For example, you can implement separate VLANs per floor for the wired LAN, but only a single WLAN VLAN for the building.

Separate VLANs for WLANs are mandatory for solutions using IP Security (IPSec) Virtual Private Networks (VPNs) or static Wired Equivalent Privacy (WEP). Protocol and address filtering is applied to traffic on these VLANs, and this filtering interferes with the traffic of users. In addition the filters are required between the wireless network and the wired network to protect the wired clients from attack from the wireless network.

Some WLAN applications, especially those using static WEP, might require one VLAN to be extended across the entire campus to support application roaming.

IP Addressing

The IP addressing of the WLAN has no direct impact on its behavior, but the design should, if possible, use a separate address space for WLAN clients. Separating address spaces in this way can ease security and management. Filters and Intrusion Detection Systems (IDSs) are easier to configure and clients are easier to identify.

You should also apply RFC 2827 filtering, which mitigates the use of enterprise networks as launching points for security breaches of other networks, on the WLAN segments to prevent spoofing.

Security Considerations

There are three different security deployment options:

- **WLAN LAN extension via EAP:** Using the Extensible Authentication Protocol (EAP) to provide dynamic per-user per-session WEP to ensure privacy, in combination with 802.11x to provide access control.
- **WLAN LAN extension via IPSec:** Using IPSec to ensure privacy and using access-point filtering, router filtering, and the VPN concentrator to provide access control.
- **WLAN static WEP:** Using whatever privacy mechanism is available and using access point filtering, router filtering, and the hardened application servers to provide access control. This security option is not recommended for open network access (access must be limited to specific applications).

Guidelines for selecting security deployment options are included later in this lesson.

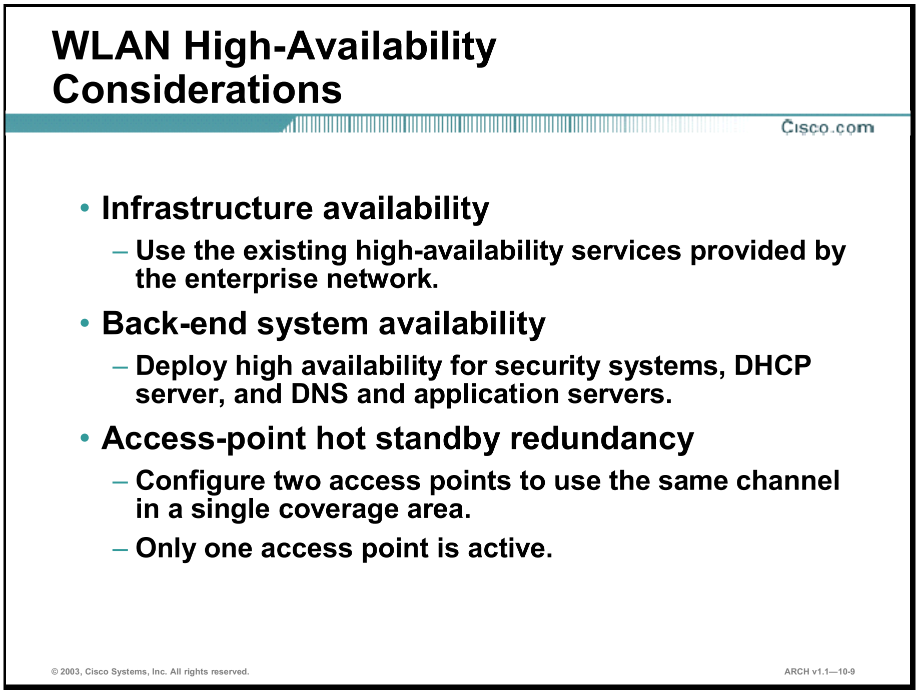

WLAN High-Availability Considerations

- Infrastructure availability
  - Use the existing high-availability services provided by the enterprise network.
- Back-end system availability
  - Deploy high availability for security systems, DHCP server, and DNS and application servers.
- Access-point hot standby redundancy
  - Configure two access points to use the same channel in a single coverage area.
  - Only one access point is active.

ARCH v1.1—10-9

Infrastructure Availability

The WLAN can use the existing high-availability services provided by the enterprise network, such as Hot Standby Router Protocol (HSRP) and Layer 2 and Layer 3 redundancies. Therefore, it is simply an issue of providing the required availability in the WLAN.

There are three ways to address the availability of the WLAN:

- **As an overlay network:** The hard-wired LAN is the users' backup if the WLAN fails.
- **As a mobile network:** Users can move to a location where connectivity is available if the WLAN fails.
- **As a mission-critical network:** Network redundancies ensure that the failure of one component does not impact WLAN users.

Back-End System Availability

High-availability designs are required for the back-end systems to ensure that individual component failures do not have widespread impact. The back-end systems vary depending upon the WLAN security deployment options:

- **WLAN LAN extension using EAP:** Back-end systems are the RADIUS asynchronous communications servers (ACSs) used to authenticate users.
- **WLAN LAN extension using IPSec:** Back-end systems are the VPN concentrators and associated servers.
- **WLAN static WEP:** Back-end systems are the application servers for the mobility application.

A mission-critical network would require the use of high-availability back-end systems such as Dynamic Host Configuration Protocol (DHCP), and Domain Name System (DNS) and application servers.

Access-Point Hot Standby Redundancy

The use of access-point hot standby is independent of the security model chosen. In the hot standby redundancy case, two access points are configured to use the same channel in a single coverage area. Only one of the access points is active. The standby access point passively monitors the network and the primary access point. If the primary access point fails, the secondary access point seamlessly takes over to provide cell coverage. A Simple Network Management Protocol (SNMP) trap is generated to alert the administrator that the primary access point has failed.

Hot standby is not the same as HSRP. Hot standby mode designates an access point as a backup for another access point. The standby access point is placed near the access point it monitors and is configured exactly the same as the monitored access point (except for its role in the radio network and IP address). The standby access point associates with the monitored access point as a client and queries the monitored access point regularly through both the Ethernet interface and the radio interface. If the monitored access point fails to respond, the standby access point comes online, signals the primary access point radio to become quiescent, and takes the monitored access point's place in the network.

As soon as the primary access point failure is detected, user intervention is required. The user must return the backup access point (which is now in root mode) to standby mode when the primary access point comes back online. Failure to reset the standby access point results in both the primary and standby access points to operate concurrently on the same channel.

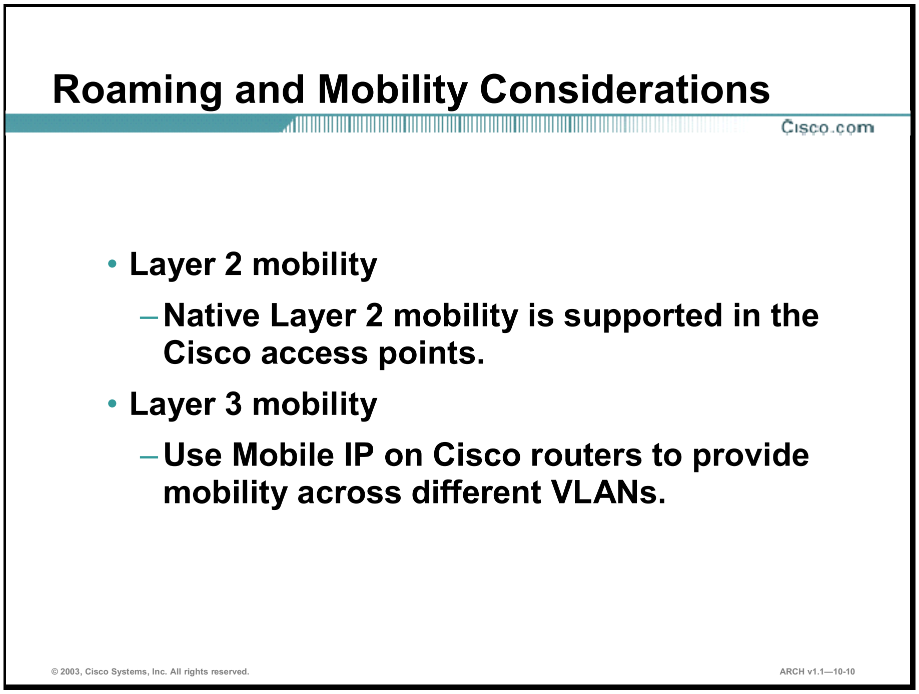

Roaming and Mobility Considerations

- Layer 2 mobility
  - Native Layer 2 mobility is supported in the Cisco access points.
- Layer 3 mobility
  - Use Mobile IP on Cisco routers to provide mobility across different VLANs.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—10-10

The native Layer 2 mobility of the Cisco access points can support devices that stay within a single subnet. Devices that need to move from subnet to subnet must acquire a new IP address and can lose packets that might have been buffered when roaming between access points on different subnets. Seamless roaming on WLAN requires that the access point involved be included within a single VLAN.

The Layer 3 mobility of the Cisco access points allows you to use Mobile IP on Cisco routers to provide mobility across different VLANs. When a WLAN is used as an overlay network, it is possible to span the same WLAN VLAN across a building floor or multiple floors of a building. This arrangement should be sufficient to meet most user mobility requirements.

Clients that require continuous connectivity, but only within a building, should be able to be accommodated by implementing a single VLAN within all but the largest buildings.

Applications that require continuous connectivity within a campus present a challenge in providing mobility across different VLANs. Cisco routers support Mobile IP, which is designed to provide this type of mobility. The issue is that there are almost no standard Mobile IP clients available, and those that are available do not support the operating systems of mobile clients that are likely to need Mobile IP support, such as scanners and 802.11 phones.

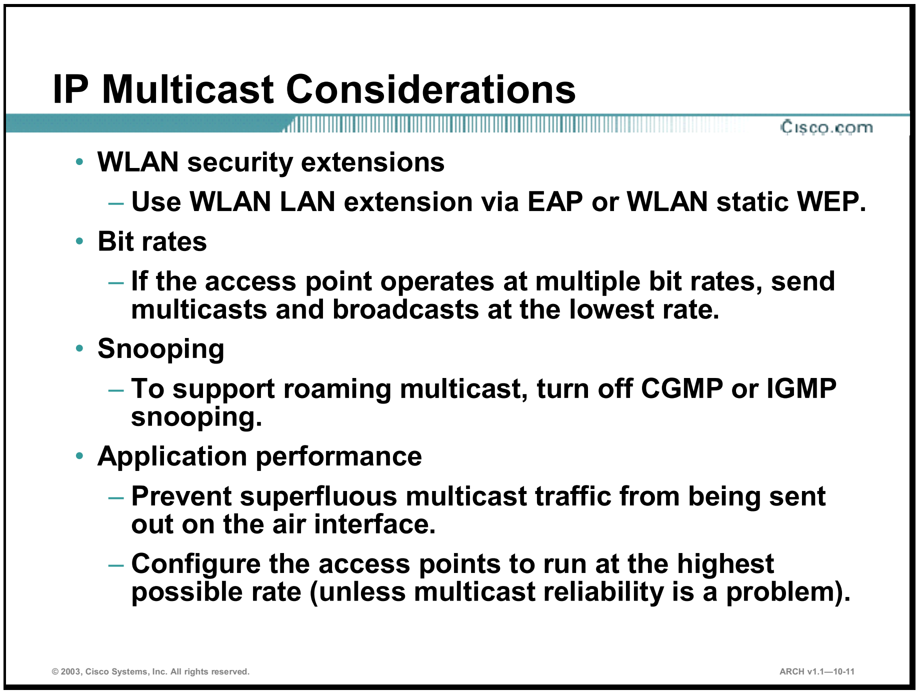

IP Multicast Considerations

- WLAN security extensions
  - Use WLAN LAN extension via EAP or WLAN static WEP.
- Bit rates
  - If the access point operates at multiple bit rates, send multicasts and broadcasts at the lowest rate.
- Snooping
  - To support roaming multicast, turn off CGMP or IGMP snooping.
- Application performance
  - Prevent superfluous multicast traffic from being sent out on the air interface.
  - Configure the access points to run at the highest possible rate (unless multicast reliability is a problem).

ARCH v1.1—10-11

These considerations apply to IP multicast in WLAN environments:

- **WLAN security extensions:** The WLAN LAN extension via EAP and WLAN static WEP solutions can support multicast traffic on the WLAN. The WLAN LAN extension via IPSec solution cannot support multicast traffic.
- **Bit rates:** The WLAN available bit rate must be shared by all clients of an access point. If the access point is configured to operate at multiple bit rates, multicasts and broadcasts are sent at the lowest rate to ensure that all clients receive them. This reduces the available throughput of the network because traffic must queue behind traffic that is being clocked out at a slower rate.
- **Snooping:** WLAN clients can roam from one access point to another seamlessly within the same subnet. If roaming multicast is to be supported, Cisco Group Management Protocol (CGMP) and/or Internet Group Management Protocol (IGMP) snooping must be turned off, because a multicast user roaming from one access point to another is roaming from one switch port to another. The new switch port might not have this stream setup and it has no reliable way of determining the required multicast stream. Therefore, to deliver multicast reliably to roaming clients, the multicast must be flooded.
- **Application performance:** Multicast and broadcast from the access point are sent without requiring link-layer acknowledgement. Every unicast packet is acknowledged and retransmitted if unacknowledged. The purpose of the acknowledgement is to overcome the inherent unreliable nature of wireless links. Broadcasts and multicasts are unacknowledged due to the difficulty in managing and scaling the acknowledgements. This means that a network that is seen as operating well for unicast applications can experience degraded performance in multicast applications.

  To ensure that multicast operates effectively and that it has minimal impact upon network performance, follow these strategies:

  - Prevent superfluous multicast traffic from being sent out on the air interface. The first step is to have the WLAN on its own subnet. The second step is to determine which multicasts must be permitted by filters, and then only allow these multicasts.

— To gain the highest performance for multicast traffic and for the access point, configure the access points to be running at the highest possible rate. This removes the requirement for multicast to clock out at a slower rate. This can impact the range of the access point and must be taken into account in the site survey.

If multicast reliability is a problem (seen as dropped packets), use a slower data rate (base rate) for multicast. This will give the multicast a better signal-to-noise ratio and can reduce the number of bad packets.

Test the multicast application for its suitability in the WLAN environment. Determine the application and user-performance effects when packet loss is higher than seen on wired networks.

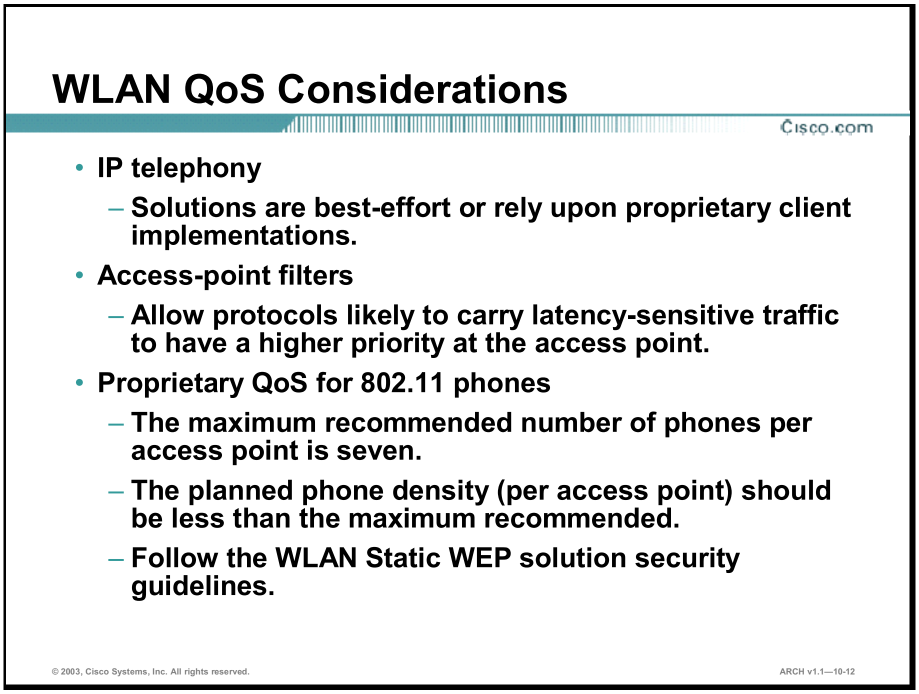

WLAN QoS Considerations

Cisco.com

- IP telephony
  - Solutions are best-effort or rely upon proprietary client implementations.
- Access-point filters
  - Allow protocols likely to carry latency-sensitive traffic to have a higher priority at the access point.
- Proprietary QoS for 802.11 phones
  - The maximum recommended number of phones per access point is seven.
  - The planned phone density (per access point) should be less than the maximum recommended.
  - Follow the WLAN Static WEP solution security guidelines.

ARCH v1.1—10-12

IP Telephony

The QoS available over a WLAN is a critical concern for certain latency-sensitive applications, particularly Voice over IP (VoIP). There is no existing standard QoS mechanism for 802.11. Solutions are best-effort or rely upon proprietary client implementations. In general, this is because 802.11 is a shared-medium protocol and the provision of QoS is more challenging over shared media, compared with switched media.

Access-Point Filters

While the access point (being a central point) is able to provide a level of queuing and prioritization for downstream traffic (access point to client), the clients must rely on other mechanisms to prioritize their upstream traffic.

Access-point filters provide strict priority queuing for downstream traffic. Filters are used to assign priority on EtherType, IP port, or protocol. Therefore, protocols likely to carry latency-sensitive traffic can have a higher priority at the access point.

Proprietary QoS for 802.11 Phones

Certain client devices such as some 802.11 phones have an upstream proprietary QoS mechanism.

If VoIP over 802.11 is required, you should consider the use of WLAN static WEP solutions that use downstream prioritization and the proprietary upstream prioritization.

The general recommendations are:

- The maximum recommended number of phones per access point is seven. This limitation is due to the number of packets that can be forwarded per second over an 802.11 link and minimizing transmission delays, rather than a bandwidth limitation of the link.
- There are no additional control mechanisms, so the planned phone density (per access point) should be less than the maximum recommended to reduce the probability of over-subscription. The impact on data throughput when carrying VoIP is unknown.
- VoIP installations should follow the WLAN static WEP solution security guidelines, because 802.11 phones currently only support static WEP. Consider implementing a different call policy on the wireless network to prevent phone fraud if a WEP key is compromised.

Any application that has stringent QoS requirements must be examined carefully to determine its suitability for 802.11 unlicensed wireless networking.

Wireless QoS implementations must allow the available network resources to be prioritized. Wireless interference in the unlicensed 802.11 frequency bands can deplete these network resources, making prioritization ineffective. This might be acceptable if it only means some dropped voice or video frames, but might be unacceptable if it means dropped frames in a real-time control system.

Apart from the QoS mechanisms in 802.11, be sure to consider the throughput and forwarding rate of the systems.

The maximum throughput of the Cisco 802.11b access point is approximately 6 Mbps under good RF conditions and at an 11-Mbps bit rate. Under poor RF or lower bit rates, the throughput will be less.

WLAN Security Extensions

You can choose from EAP, IPSec, and WEP as the security model for a WLAN implementation. The choice of security model has far-reaching design implications. Whenever possible, EAP should be implemented in a WLAN. This topic describes the available WLAN security extensions and the differences among them.

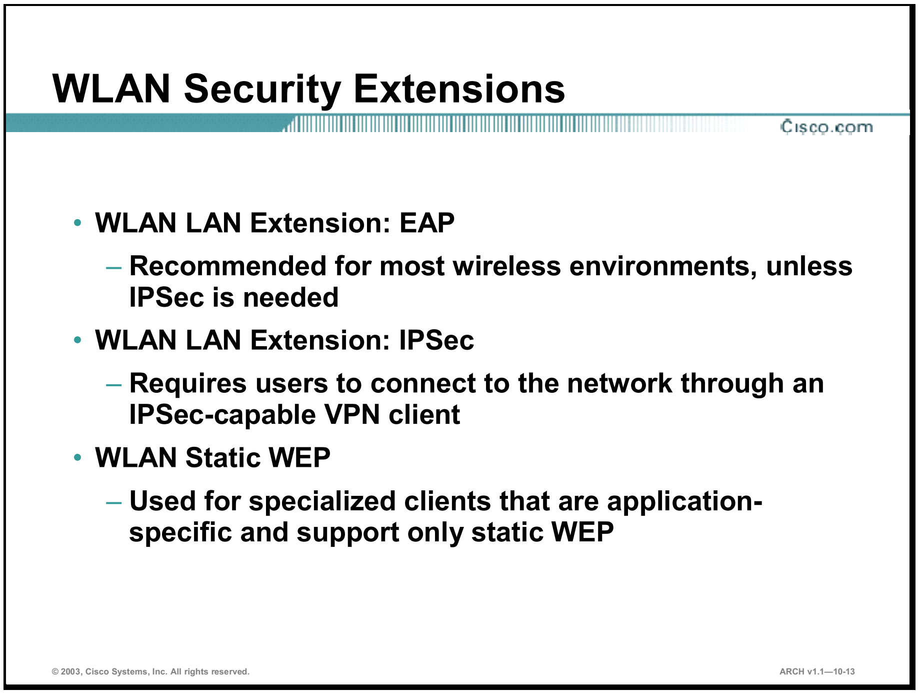

WLAN Security Extensions

Cisco.com

- WLAN LAN Extension: EAP
  - Recommended for most wireless environments, unless IPSec is needed
- WLAN LAN Extension: IPSec
  - Requires users to connect to the network through an IPSec-capable VPN client
- WLAN Static WEP
  - Used for specialized clients that are application-specific and support only static WEP

ARCH v1.1—10-13

The security model selected for a given WLAN implementation has a substantial impact on the overall WLAN design. The three security models to consider are:

- WLAN LAN extension: Extensible Authentication Protocol
- WLAN LAN extension: IP Security
- WLAN static WEP

WLAN LAN Extension: Extensible Authentication Protocol

EAP-Cisco provides these advantages:

- Requires no user intervention
- Provides per-user authentication
- Automatically provides a dynamic WEP key, thus overcoming key management issues associated with WEP
- Supports accounting
- Does not require any additional filtering or access control
- Is multiprotocol and can carry protocols other than IP over the WLAN
- Supports the same filtering requirements at the network access layer as those for wired implementations

While EAP is the recommended option, it may not be suitable in all cases for these reasons:

- EAP requires EAP-aware access points and WLAN clients, and a client might not be available for your operating system. In this case, there is no EAP solution for your preferred authentication type. EAP-Cisco is available only from Cisco and Apple (client only). EAP-Transport Layer Security (TLS) protoacol is supported by Microsoft on XP clients. You can use either the Microsoft client or the Cisco Aironet client utility. Cisco access points support all EAP solutions that conform to the 802.1x and EAP standard.
- You may require the security features offered by IPSec, such as Triple Data Encryption Standard (3DES) encryption, One Time Password (OTP) support, or per-user policies.
- WLAN clients, such as scanners or 802.11 phones, might not support EAP.
- Where seamless roaming within a Layer 2 domain is required, EAP clients can take longer to roam between access points, compared to those using static WEP, which can impact some applications, such as VoIP over 802.11.

WLAN LAN Extension: IP Security

As with the WLAN EAP solution, IPSec provides a WLAN LAN extension service. However, this solution requires users to connect to the network through an IPSec-capable VPN client, even within a campus environment.

Typical characteristics of a WLAN using IPSec VPNs are:

- It does not require the use of EAP, and allows any client adaptor to be used with a 3DES encryption.
- It allows the use of multifactor authentication systems, such as OTP systems.
- It requires the implementation of extensive filters on the network edge to limit network access to IPSec-related traffic destined to the VPN concentrator network.
- It requires user intervention. The users must launch the VPN client before they attach to the network.
- It does not support multicast applications.
- It requires local traffic to go through the VPN concentrator, causing traffic to cross the network multiple times, increasing traffic across the network and degrading performance.
- Clients such as scanners or 802.11 phones might not support IPSec.

WLAN Static WEP

WLAN static WEP addresses specialized clients that are application-specific and support only static WEP.

Within each enterprise, small application verticals exist that can benefit from WLAN applications (specialized applications that run on specialized clients designed for mobile use). Applications requiring this type of solution might also require uninterrupted seamless coverage. Examples of potential WLAN applications that may use static WEP are as follows:

- VoIP over 802.11
- Messaging applications
- Workflow applications
- Security applications
- Package-tracking applications

The following table compares the three security implementation models in detail.

	WLAN LAN Extension via EAP	WLAN LAN Extension via IPSec	WLAN Static WEP
Protocols	Multiprotocol	Unicast only	Multiprotocol
Network interface cards	Cisco and Apple	Any 802.11b-compliant card	Any 802.11b-compliant card
Connection to network	Integrated with Windows login; non-Windows users enter user name and password	User must launch a VPN client and log in	Transparent to user
Clients	Laptop PCs, high-end PDAs; a wide range of operating systems supported	Laptop PCs, high-end PDAs; a wide range of operating systems supported	Any 802.11 client
Authentication	User name and password or certificates	OTP or user name and password	Matching WEP key required
Privacy	Dynamic WEP with time-limited keys and TKIP[1] enhancements	3DES	Static WEP (with TKIP enhancements for Cisco clients); problematic key management
Impact on existing network architecture	Additional RADIUS server required	Additional infrastructure WLAN will be on a perimeter LAN and require VPN concentrators, authentication servers, DHCP servers	Option of additional firewall software or hardware at access layer
Filtering	None required	Extensive filtering required, limiting network access until VPN authentication has occurred	Extensive filtering required, limiting wireless access to only certain predetermined applications

	WLAN LAN Extension via EAP	WLAN LAN Extension via IPSec	WLAN Static WEP
Layer 2 roaming	Transparent Automatically reauthenticates without client intervention (may be slower than VPN or WEP)	Transparent May be easier to extend Layer 2 domain due to reduced broadcast and multicast traffic	Transparent
Layer 3 roaming	Requires IP address release or renewal, or Mobile IP solution	Requires IP address release or renewal, or Mobile IP solution	Requires IP address release or renewal, or Mobile IP solution
Management	Network is open to existing network management systems	Filtering must be adjusted to support management applications	May have application-specific management requirements; filtering must be adjusted to support the management applications
QoS	Best-effort QoS Proprietary client schemes exist but do not currently support EAP	Best-effort QoS Proprietary client schemes exist but do not currently support IPSec; IPSec tunnel prevents the use of NBAR[2] until after the VPN concentrator	Best-effort QoS unless proprietary client schemes are used, such as Symbol and Spectralink
Multicast	Supported	Not supported	Supported
Performance	WEP encryption performed in hardware on Cisco NICs for EAP-Cisco May be performed in software for other EAP solutions	3DES performed in software, an expected throughput hit of 20–30 percent	WEP encryption performed in hardware on Cisco NICs

[1] TKIP = Temporal Key Integrity Protocol

[2] NBAR = network-based application recognition

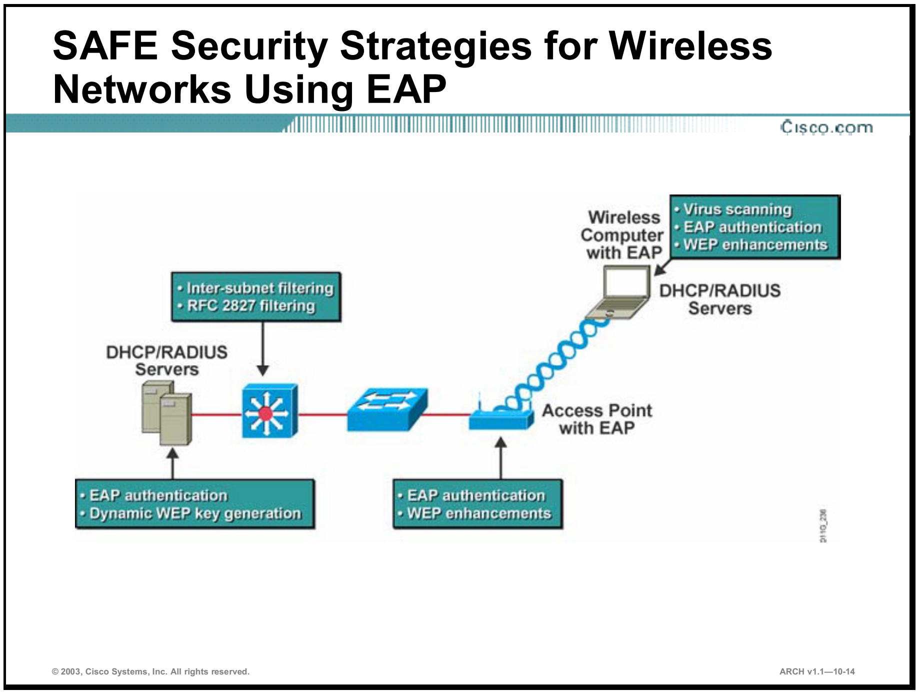

SAFE Security Strategies for Wireless Networks Using EAP

Cisco.com

Wireless Computer with EAP

- Virus scanning
- EAP authentication
- WEP enhancements

DHCP/RADIUS Servers

- Inter-subnet filtering
- RFC 2827 filtering

DHCP/RADIUS Servers

Access Point with EAP

- EAP authentication
- Dynamic WEP key generation

- EAP authentication
- WEP enhancements

2110_236

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—10-14

In most cases, WLAN access points are connected to existing data link layer access switches. RADIUS and DHCP servers are located in the server module of the corporate network. Security in the design is maintained by preventing network access in the event of a RADIUS service failure, since most of the mitigation against security risks relies on the RADIUS service. Overall, management of the security solution is hindered if DHCP services fail. The wireless clients and application processors use EAP to authenticate the WLAN client devices and end users against the RADIUS servers. Be sure to require (and check) that users choose strong passwords and set account lockouts after a small number of incorrect login attempts. This configuration can be made at the RADIUS server.

For scalability and manageability purposes, the WLAN client devices are configured to use the DHCP protocol for IP configuration. DHCP occurs after the device and end user are successfully authenticated via EAP Cisco Wireless (also called Cisco LEAP). After successful DHCP configuration, the wireless end user is allowed access to the corporate network. Filtering in place at the first Layer 3 switch prevents the wireless network from accessing portions of the wired network as dictated by an organization's security policy. In SAFE, for example, filtering was put in place to prevent wireless access to any department servers, voice networks, or other user networks. Network designers should give special consideration to the location of the RADIUS and DHCP servers used by EAP.

The design shown in the figure mitigates the following attacks:

- **Wireless packet sniffers:** Wireless packet sniffers can take advantage of any of the known WEP attacks to derive the encryption key. These threats are mitigated by WEP enhancements, and key rotation using LEAP.
- **Unauthenticated access:** Only authenticated users are able to access the wireless and wired network. Optional access control on the Layer 3 switch limits wired network access.
- **Man-in-the-middle:** The mutual authentication nature of EAP combined with the Message Integrity Check (MIC) prevents a hacker from inserting itself in the path of wireless communications.
- **IP spoofing:** Hackers cannot perform IP spoofing without first authenticating to the WLAN. Authenticating optional RFC 2827 filtering on the Layer 3 switch restricts any spoofing to the local subnet range.
- **Address Resolution Protocol (ARP) spoofing:** Hackers cannot perform ARP spoofing without first authenticating to the WLAN. Authenticating ARP spoofing attacks can be launched in the same manner as in a wired environment to intercept other user's data.
- **Network topology discovery:** Hackers cannot perform network discovery if they are unable to authenticate. When authenticated via EAP, standard topology discovery can occur in the same way that is possible in the wired network.

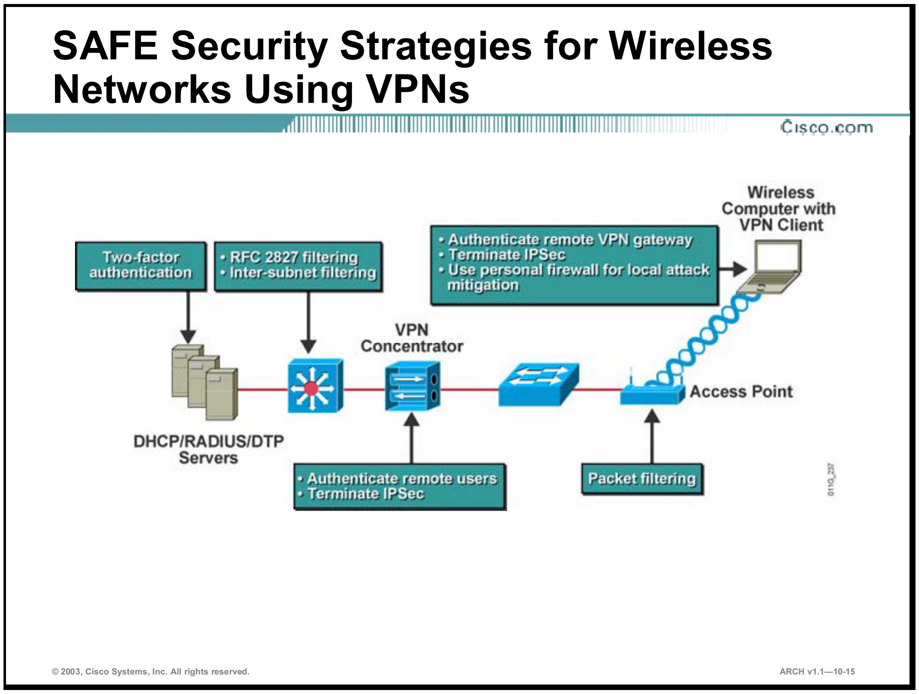

SAFE Security Strategies for Wireless Networks Using VPNs

Cisco.com

Two-factor authentication
• RFC 2827 filtering
• Inter-subnet filtering
• Authenticate remote VPN gateway
• Terminate IPSec
• Use personal firewall for local attack mitigation
Wireless Computer with VPN Client
VPN Concentrator
Access Point
DHCP/RADIUS/DTP Servers
• Authenticate remote users
• Terminate IPSec
Packet filtering
0110_237

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—10-15

WLAN access points connect to data link layer switches in the Campus Backbone on a dedicated VLAN, and forward traffic from the WLAN to the wired LAN using IPSec to protect the flows until they reach the wired network. It is important to point out that WEP is not enabled in this design. The wireless network itself is considered an untrusted network, suitable only as a transit network for IPSec traffic. In order to isolate this untrusted network, administrators should not mix the VLAN for the WLAN users with a wired network. This configuration would potentially allow hackers on the wireless network to attack users on the wired network. The WLAN clients associate with a wireless access point to establish connectivity to the campus network at the data link layer. The wireless clients then use DHCP and DNS services in the Server Farm module to establish connectivity to the campus at the network layer. When the wireless client is communicating with the campus network, but before the IPSec tunnel is established, the client traffic is not considered secure. All the noted WLAN security issues are still present until the wireless client can secure communications with an IPSec VPN. Therefore, two mitigation techniques are recommended:

- First, you should configure the access point with the ether type protocol and port filters based on a company's wireless usage policy. SAFE WLAN recommends restrictive filters that allow only the necessary protocols required for establishing a secure tunnel to a VPN gateway. These protocols include DHCP for initial client configuration, DNS for name resolution of the VPN gateways, and the VPN-specific protocols, Internet Key Exchange (IKE) on User Datagram Protocol (UDP) port 500, and Encapsulating Security Payload (ESP) (IP Protocol 50). The DNS traffic is optional, dependent on whether the VPN client needs to be configured with a DNS name for the VPN gateway or if only an IP address is suitable.

- Secondly, you should use personal firewall software on the wireless client to protect the client while it is connected to the untrusted WLAN network, without the protection of IPSec. In general terms, the VPN gateway delineates between the trusted wired network and the untrusted WLAN. The wireless client establishes a VPN connection to the VPN gateway to start secure communication to the corporate network. In the process of doing so, the VPN gateway provides device and user authentication via the IPSec VPN.

Even with this filtering, the DNS and DHCP servers are still open to direct attack on the application protocols themselves. Extra care should be taken to ensure that these systems are as secure as possible at the host level. This includes keeping them up to date with the latest OS and application patches and running a Host Intrusion Detection System (HIDS).

The VPN gateway can use digital certificates or preshared keys for wireless device authentication. The VPN gateway then takes advantage of OTPs to authenticate users. Without OTP, the VPN gateways are open to brute-force login attempts by hackers who have obtained the shared IPSec key used by the VPN gateway. The VPN gateway takes advantage of RADIUS services, which in turn contact the OTP server for user authentication. The VPN gateway uses DHCP for IP address configuration for the WLAN client to communicate through the VPN tunnel. Security in the design is maintained by preventing network access if a VPN gateway or RADIUS service fails. Both services are required in order for the client to reach the wired network with production traffic.

Network designers may still consider enabling static WEP keys on all devices in an effort to add an additional deterrent against hackers. Although enhancements to WEP, such as the MIC and WEP key hashing, provide effective risk mitigation to currently identified WEP vulnerabilities, the management overhead of dealing with static key changes makes this alternative less than ideal for large WLAN deployments. This management overhead could be mitigated by never changing the static WEP key, but this solution falls strongly into the "security through obscurity" category.

To further secure the DNS and DHCP services, network designers should consider using dedicated hosts for the VPN WLAN DHCP and DNS deployment. This mitigates against two potential threats that could affect wired resources:

- Denial-of-service attacks against the DHCP and DNS services that could affect wired users
- Network reconnaissance through the use of DNS queries or reverse-lookups

As an alternative to dedicated DNS servers, designers may consider hard-coding the IP address of the VPN gateway for the VPN clients. The drawback of this solution is that if the IP address of the VPN gateway changes, every client will need to update their gateway entry.

The design shown in the figure mitigates the following attacks:

- **Wireless packet sniffers:** These threats are mitigated by IPSec encryption of wireless client traffic.
- **Man-in-the-middle:** These threats are mitigated by IPSec encryption of wireless client traffic.
- **Unauthorized access:** The only known protocols for initial IP configuration, DHCP and VPN access protocols (DNS, IKE, and ESP), are allowed from the WLAN to the corporate network through filtering at the access point and Layer 3 switch. Optionally, you can enforce authorization policies on the VPN gateway for individual user groups.
- **IP spoofing:** Hackers can spoof traffic on the WLAN, but only valid, authenticated IPSec packets will ever reach the wired network.
- **ARP spoofing:** ARP spoofing attacks can be launched, but data is encrypted to the VPN gateway, so hackers will be unable to read the data.
- **Password attacks:** These threats are mitigated through good password policies, auditing, and OTP.

- **Network topology discovery:** Only IKE, ESP, DNS, and DHCP are allowed from this segment into the corporate network.

This threat is not mitigated:

- **MAC/IP spoofing from unauthenticated users:** ARP spoofing and IP spoofing are still effective on the WLAN subnet until the wireless client uses IPSec to secure the connection.

Small Office WLAN Design Model

In a small office, implement a WLAN to extend the network reach to areas where physical constraints, cost, or speed of deployment are issues. This topic describes the design of Cisco WLAN solutions for small enterprise networks.

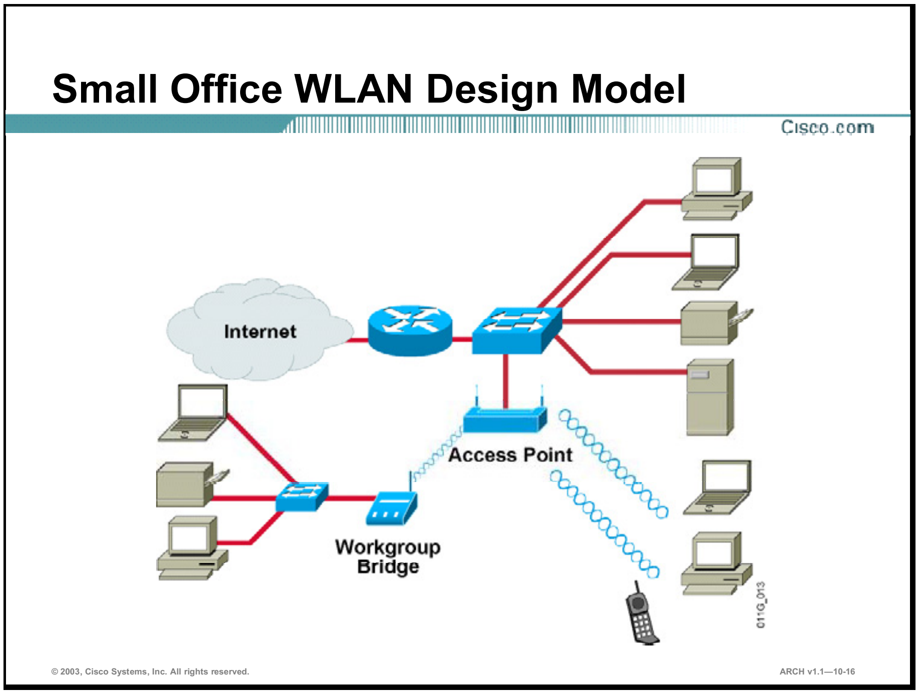

Small Office WLAN Design Model

Cisco.com

ARCH v1.1—10-16

In the small office, the WLAN may be used as an overlay network, or it may be used as a replacement network. Generally WLAN is not recommended as a replacement network, but you can use it quite effectively to extend network reach to areas where physical constraints, cost, or speed of deployment are issues.

The workgroup bridge is a useful aid on these types of extensions, as you can extend the WLAN without having to have WLAN cards in the devices. The workgroup bridge can support up to eight devices.

This figure shows the WLAN on the same subnet as other users. You may want to consider filters on the access points to limit the amount of broadcast and multicast traffic sent to the WLAN.

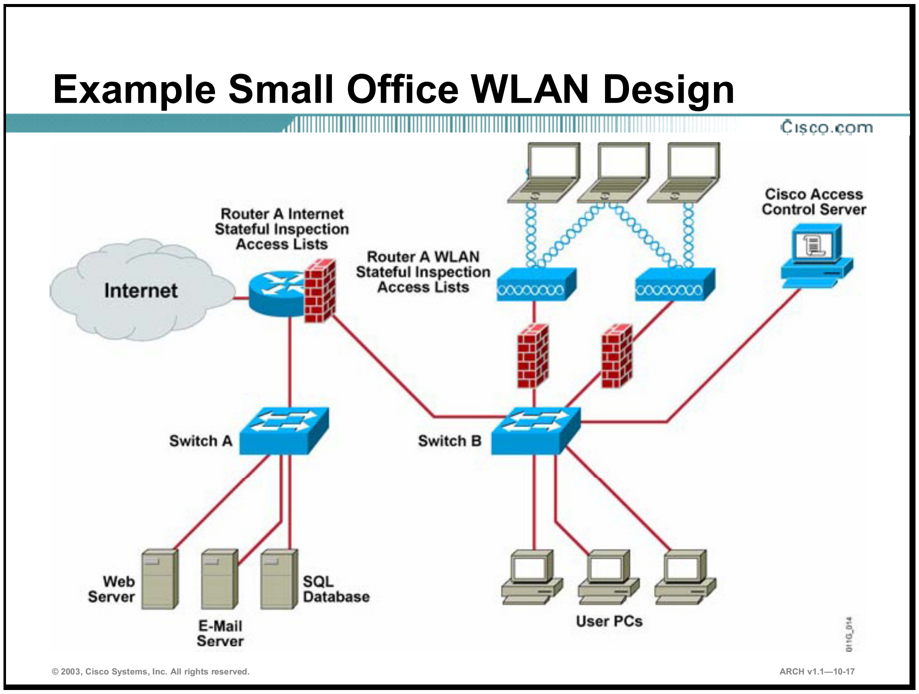

Example Small Office WLAN Design

Cisco.com
Router A Internet Stateful Inspection Access Lists
Internet
Router A WLAN Stateful Inspection Access Lists
Cisco Access Control Server
Switch A
Switch B
Web Server
E-Mail Server
SQL Database
User PCs

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—10-17

The environment illustrated in the figure assumes a WLAN that provides access to a server farm.

Key management via EAP systems provides the least intrusive form of security for the small office and offers the lowest cost of ownership, as it is compatible with all topologies and should require the least configuration and maintenance.

Enterprise WLAN Design Model

A WLAN is typically deployed in an enterprise network as an extension rather than a replacement. The goal is usually to obtain the benefits of WLAN with as little disruption as possible to the existing infrastructure. This topic describes the design of WLAN solutions for large enterprise networks.

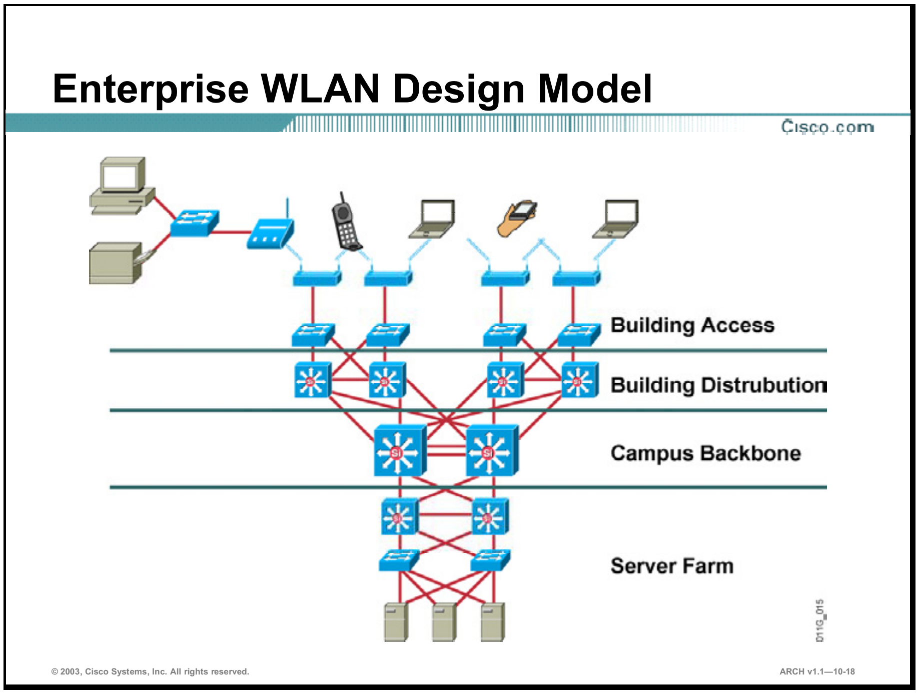

Enterprise WLAN Design Model

Cisco.com

ARCH v1.1—10-18

The goal for the introduction of WLAN into the enterprise environment is to obtain the benefits of WLAN with as little disruption as possible to the existing infrastructure. The idealized solution would be to simply attach the WLAN to the existing LAN infrastructure and go from there. The caveat in this is that the WLAN is a shared media where broadcast and multicast traffic are sent at the slowest connection speed, and are buffered for equipment in power save mode.

To create an environment where multicasts and broadcasts can be more easily controlled, it is recommended that the WLAN be a separate VLAN to the wired LAN.

You need to take into account the number of expected users on an access point and their traffic requirements in the site survey and RF design of the sites. The shared nature of WLAN impacts the location and density of WLAN equipment, but normally does not impact the architecture.

The biggest influence on the network architecture is the mechanism to secure the WLAN. The WLAN shown in the figure needs a key management solution, such as network EAP, EAP-TLS, or IPSec VPNs to make it secure enough for enterprise use.

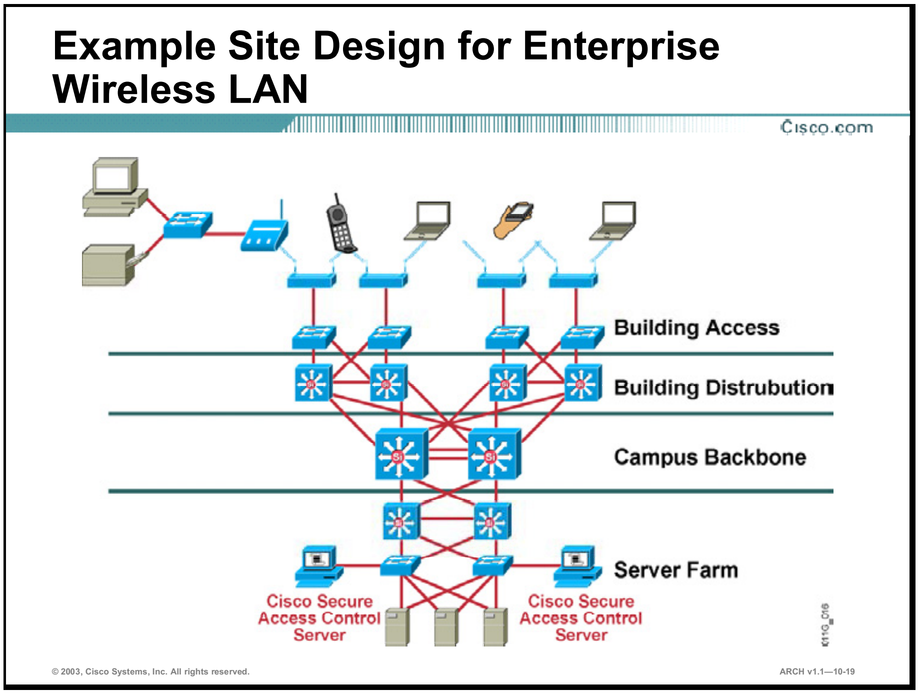

Example Site Design for Enterprise Wireless LAN

ARCH v1.1—10-19

Acme Corporation has decided to extend their network with wireless access in two conference rooms. The first decision to make is the choice of security implementation. To obtain maximum security, Acme chooses EAP. This requires the addition of Cisco ACS authentication servers to their server farm. For the wireless components themselves, two access points per conference room are installed, providing access for up to 50 simultaneous users.

The EAP solution provides privacy over the WLAN via the dynamic WEP key, and controls access to the network by the WEP key in combination with 802.1x and the RADIUS server authentication. This creates end-to-end network security.

Remote-Access and Telecommuter WLAN Design Models

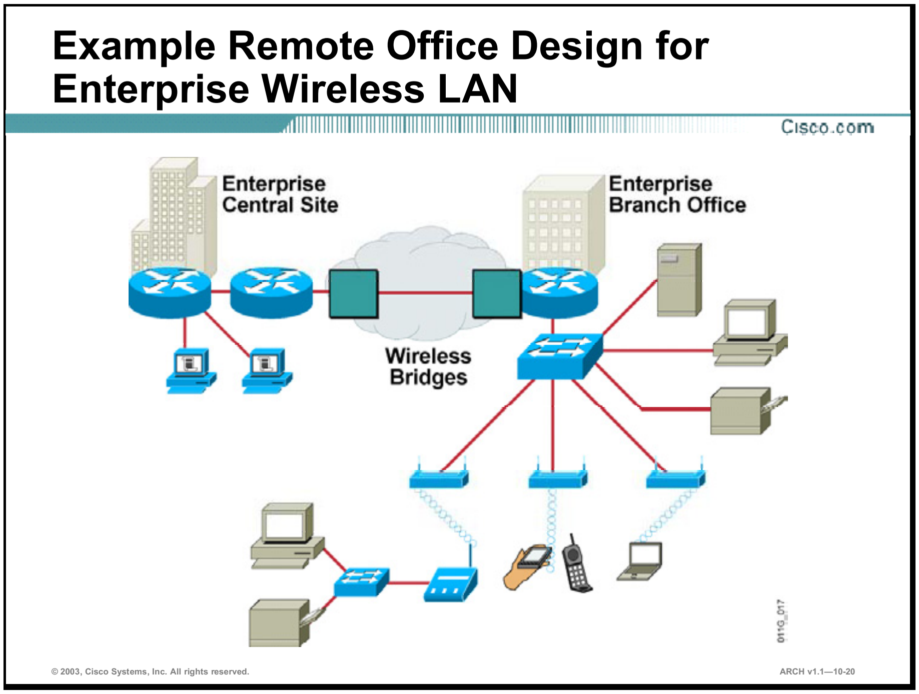

Example Remote Office Design for Enterprise Wireless LAN

ARCH v1.1—10-20

Corporate growth has required that a company acquire additional office space. No space is available in their headquarters building, but space is located within the same complex a half-mile away. Wireless bridges are used to connect the branch office LAN to the headquarters LAN. The same headquarters ACS servers are used to provide security to the wireless components of the remote office.

The access points are added as an overlay to the existing wired LAN, providing some degree of mobility within the branch office. Since this is a single network using access points, clients can roam seamlessly about the office (from access point to access point) without disruption.

You can use a wireless base station to provide access for a location with a small number of users, such as a small remote office or telecommuter. This topic describes the design of Cisco WLAN solutions for remote access and telecommuters.

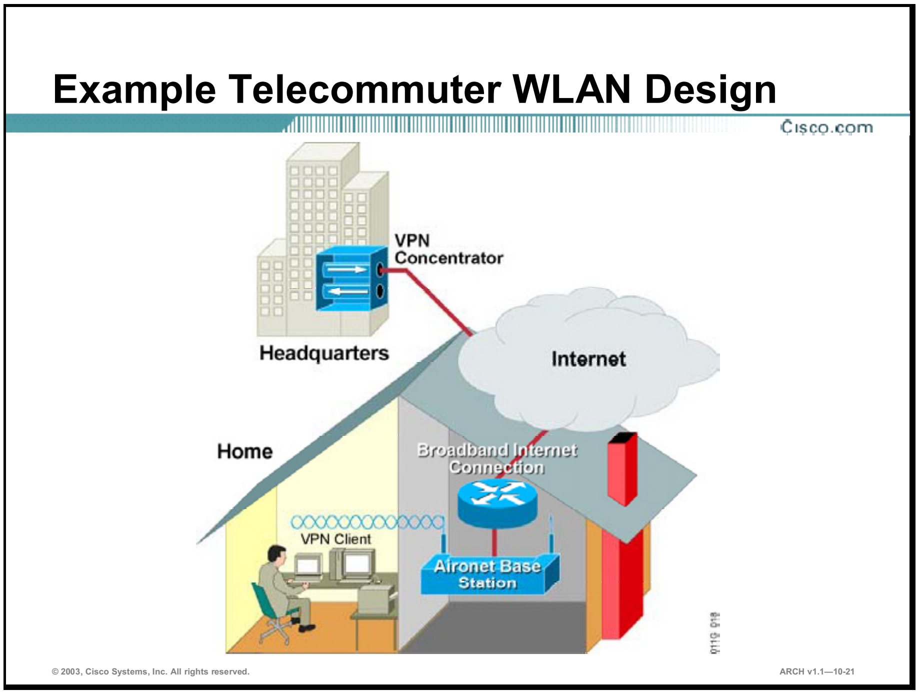

Example Telecommuter WLAN Design

ARCH v1.1—10-21

Several of Acme's remote sales offices are staffed with only a few employees, with individual salespeople frequently operating from home. To provide these employees with access while maintaining the mobility required, Aironet base stations are installed in the smallest offices and the telecommuter homes.

Strong client authentication must be required before access from the home network to the corporate network is permitted. An IPSec-based, client-initiated VPN authenticates the client and ensures that only authorized people can access the corporate network from the home network. As well as authenticating the individual access to the corporate network, VPNs provide data privacy via strong encryption algorithms such as 3DES.

This design uses a broadband connection to the Internet such as DSL or cable. The broadband router supplying this access includes an internal firewall to protect the home network from intrusion.

The base station supports static 128-bit WEP to prevent unauthorized access to the home network. Data privacy is provided, independent of WEP, by the VPN for traffic to and from the corporate site, and by Secure Socket Layer (SSL) when required for general web traffic.

Personal firewall software should be installed on client PCs to protect the PCs from attacks from the WLAN or the Internet. If an attacker gains access to a client PC, he can use that to launch other attacks on the corporate network at a later time.

Summary

This topic summarizes the key points discussed in this lesson.

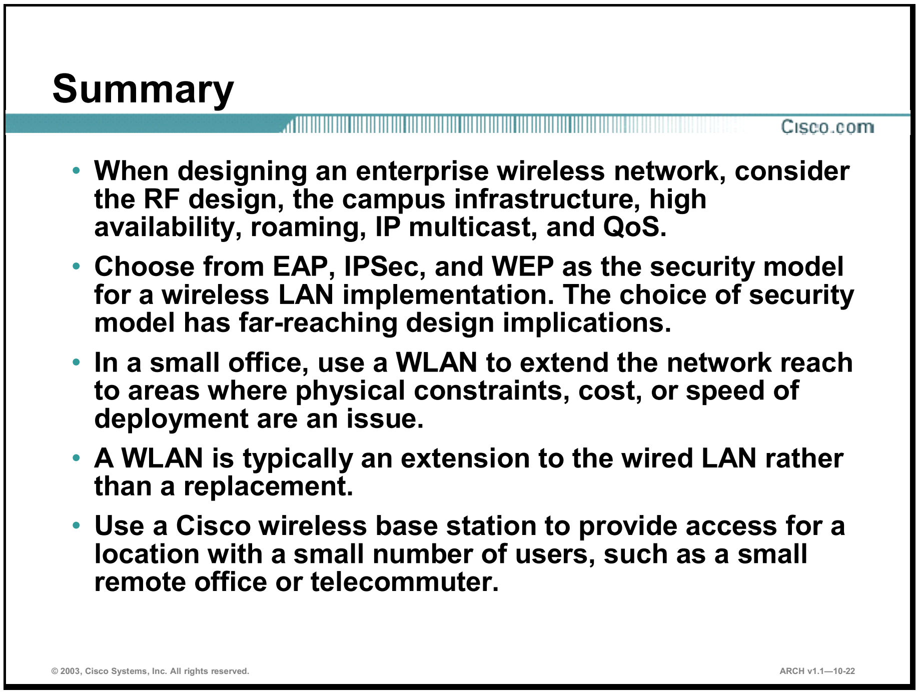

Summary

Cisco.com

- When designing an enterprise wireless network, consider the RF design, the campus infrastructure, high availability, roaming, IP multicast, and QoS.
- Choose from EAP, IPSec, and WEP as the security model for a wireless LAN implementation. The choice of security model has far-reaching design implications.
- In a small office, use a WLAN to extend the network reach to areas where physical constraints, cost, or speed of deployment are an issue.
- A WLAN is typically an extension to the wired LAN rather than a replacement.
- Use a Cisco wireless base station to provide access for a location with a small number of users, such as a small remote office or telecommuter.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—10-22

References

For additional information, refer to these resources:

- *Wireless Solutions* at http://www.cisco.com/warp/public/44/solutions/network/wireless.shtml
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study, refer to the following section:

- Case Study 10-2: OCSIC Bottling Company

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) What are two drawbacks to configuring an access point to operate at multiple data rates? (Choose two.)

A) Equipment costs increase.

B) Clients operating at the lower rate degrade overall throughput.

C) Clients operating at the higher rate degrade overall throughput.

D) Coverage area increases, potentially interfering with other WLANs.

E) Coverage area decreases, potentially leaving some clients unsupported.

Q2) What two advantages are gained by placing WLAN devices on a separate VLAN? (Choose two.)

A) increased security

B) decreased equipment costs

C) optimized network performance

D) increased Layer 2 roaming domain

E) increased Layer 3 roaming domain

Q3) If you plan to implement wireless IP telephony, which security implementation should be used?

A) EAP.

B) IPSec.

C) static WEP.

D) Any security implementation will work.

Q4) What are two reasons a small office might decide to implement a wireless solution? (Choose two.)

A) physical constraints

B) speed of deployment

C) ease of administration

D) security requirements

E) government regulations

Q5) If EAP security is implemented, what additional equipment is required?

A) RADIUS server

B) Cisco ACS server

C) workgroup bridge

D) access point server

Q6) For a wireless telecommuter solution, what security implementation should you select?

A) EAP

B) WEP

C) IPSec

D) VLANs

Quiz Answer Key

Q1) B, D

Relates to: Enterprise WLAN Design Considerations

Q2) C, D

Relates to: Enterprise WLAN Design Considerations

Q3) C

Relates to: WLAN Security Extensions

Q4) A, B

Relates to: Small Office WLAN Design Model

Q5) B

Relates to: Enterprise WLAN Design Model

Q6) C

Relates to: Remote Access and Telecommuter WLAN Design Models

Case Study 10-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

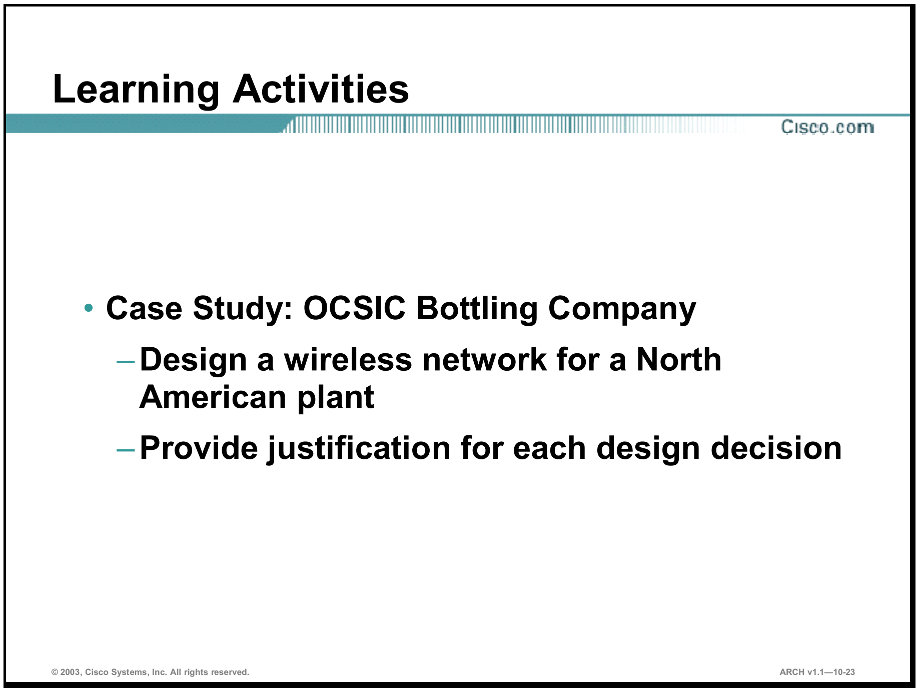

Learning Activities

- Case Study: OCSIC Bottling Company
  - Design a wireless network for a North American plant
  - Provide justification for each design decision

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—10-23

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

The OCSIC Bottling Company wants to attach wireless devices to key control systems in the plants to be able to monitor and reconfigure them remotely. They also want to implement a tracking system to track their mobile plant material handling equipment, and provide network access for inventory accounting and control.

In this exercise, you will design a wireless network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design a wireless network for a North American plant

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Design a Wireless Network for a North American Plant

Complete these steps:

Step 1 On an overhead transparency, create a campus network diagram indicating your wireless LAN design for one of the North American plants. Label each location.

Step 2 Complete the table to design the details about the wireless network.

Design Questions	Decision	Justification
How many access points are required within a typical OCSIC 60,000-square-foot district office or plant? Where should the access points be placed?		
How many active devices can each access point support?		
How are channels identified for the design?		
How will you meet the inline power requirements for the design?		
What is the high-availability (redundancy) strategy for the design?		

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have created a WLAN design that includes the following components:
  - Number and placement of access points
  - Number of devices each access point will support
  - Channel selection
  - Inline power requirements
  - High availability

Module 11

Designing IP Telephony Solutions

Overview

Built on the Cisco Architecture for Voice, Video and Integrated Data (AVVID) network infrastructure, a Cisco IP telephony solution delivers high-quality IP voice and fully integrated communications by allowing voice to be originated on and transmitted over a single network infrastructure along with data and video. Cisco IP telephony solutions provide feature functionality with straightforward configuration and maintenance requirements and interoperability with a wide variety of other applications.

Module Objectives

Upon completing this module, you will be able to design enterprise solutions for IP telephony, given enterprise network needs.

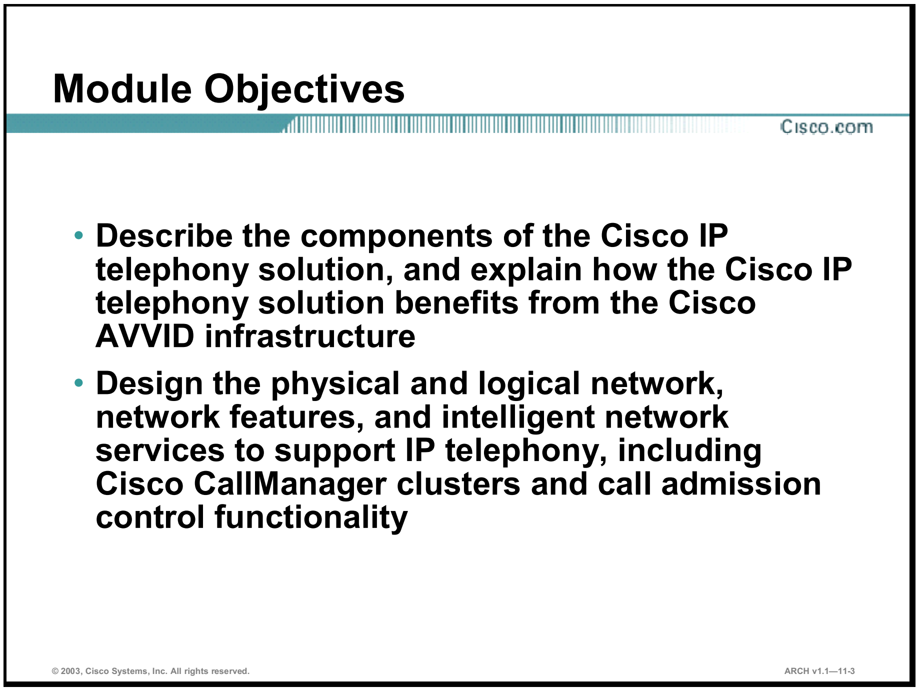

Module Objectives

- Describe the components of the Cisco IP telephony solution, and explain how the Cisco IP telephony solution benefits from the Cisco AVVID infrastructure
- Design the physical and logical network, network features, and intelligent network services to support IP telephony, including Cisco CallManager clusters and call admission control functionality

ARCH v1.1—11-3

Module Outline

The outline lists the components of this module.

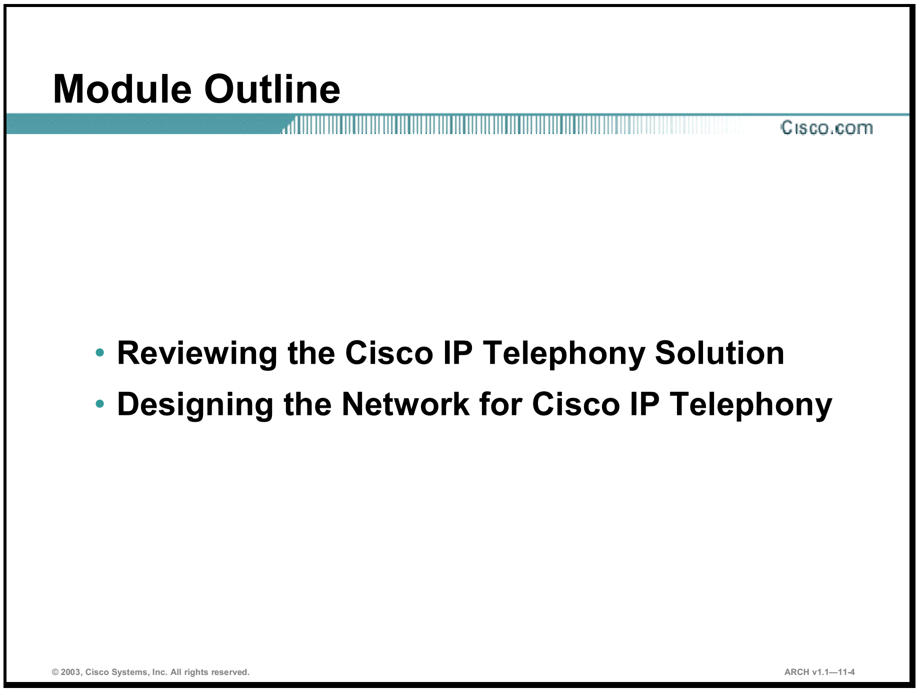

Module Outline

- Reviewing the Cisco IP Telephony Solution
- Designing the Network for Cisco IP Telephony

ARCH v1.1—11-4

Reviewing the Cisco IP Telephony Solution

Overview

The flexibility and functionality of the Cisco AVVID network infrastructure provides a framework that permits rapid deployment of IP telephony applications.

Relevance

The Cisco AVVID framework, combined with multicast services and the Cisco IP telephony solution, provides universal transport for data, voice, and video applications today.

Objectives

Upon completing this lesson, you will be able to describe the components of the Cisco IP telephony solution and to explain how the Cisco IP telephony solution benefits from the Cisco AVVID infrastructure. This includes being able to meet these objectives:

- List the components of the Cisco IP telephony solution and explain the role of each component
- Describe Cisco CallManager features and deployment models
- Identify the gateway components of the Cisco IP telephony solution and explain the role of each component
- Describe the types and role of the transcoder and conferencing components of an IP telephony solution
- List the telephony application components of the Cisco IP telephony solution and explain the role of each component

Learner Skills and Knowledge

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

- IP telephony basics
- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

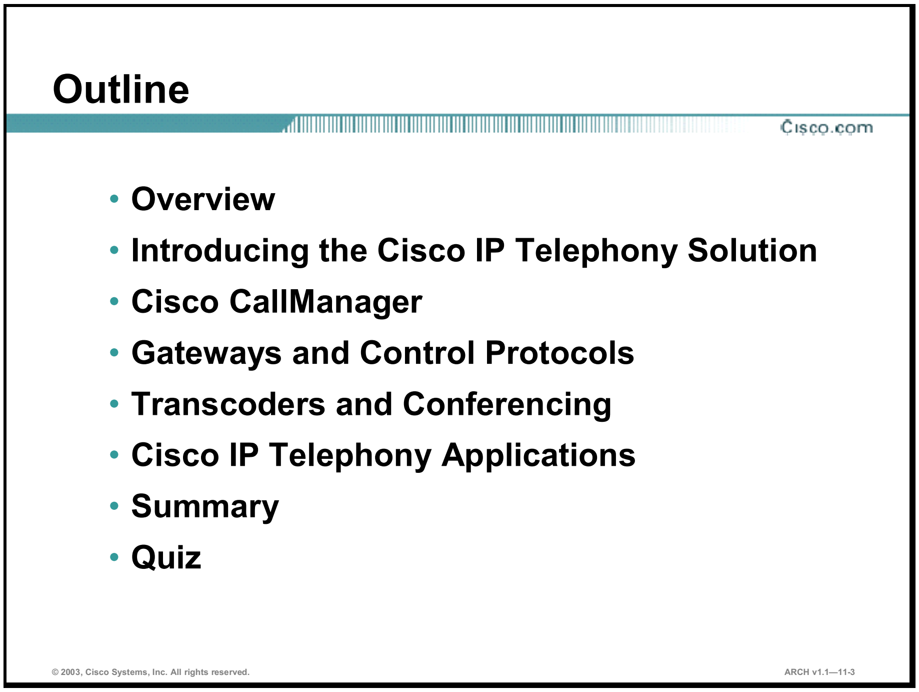

Outline

- Overview
- Introducing the Cisco IP Telephony Solution
- Cisco CallManager
- Gateways and Control Protocols
- Transcoders and Conferencing
- Cisco IP Telephony Applications
- Summary
- Quiz

ARCH v1.1—11-3

Introducing the Cisco IP Telephony Solution

The Cisco IP telephony solution includes infrastructure components of the traditional IP network, as well as devices dedicated to voice. This topic lists the components of the Cisco AVVID IP telephony solution and explains the role of each component in the overall IP telephony solution.

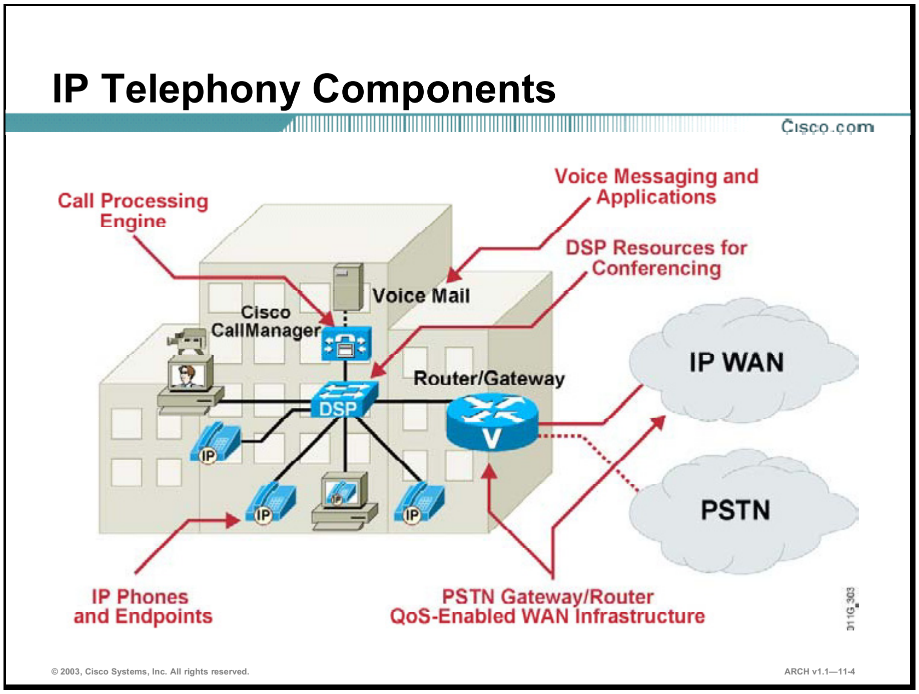

IP Telephony Components

Cisco.com

Call Processing Engine
Voice Messaging and Applications
DSP Resources for Conferencing
Cisco CallManager
Voice Mail
Router/Gateway
DSP
V
IP
IP WAN
PSTN
IP Phones and Endpoints
PSTN Gateway/Router QoS-Enabled WAN Infrastructure
311G_303

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-4

The Cisco IP telephony solution makes it possible to implement a phone system that is transparent to the users. The user sees a phone that offers the same service and quality offered in a PBX system, but that has unique data characteristics as well.

The telephony components found within the architecture include:

- **Call processing engine:** The call processing engine routes calls on the network. CallManager is the primary call processing engine for Cisco IP telephony networks. It provides call control and signaling services to client endpoints, which may include telephones or gateways.
- **IP Phones:** IP Phones convert analog voice to digital IP packets so the IP network can transport them. Other endpoints require gateways.
- **Gateways:** Gateways provide access to endpoints other than IP telephones, such as applications, WAN facilities, the Public Switched Telephone Network (PSTN), and other IP telephony and Voice over IP (VoIP) installations.
  - **Applications:** Voice applications include voice mail, automated attendant, interactive voice response (IVR), call distribution, and others.
  - **Voice-enabled routers:** Voice-enabled routers route traffic between the campus voice network and WAN facilities.
  - **PSTN gateways:** PSTN gateways enable enterprises to send calls to other enterprises and individuals over the PSTN.

— **VoIP gateways:** VoIP gateways, often software running within a network device, enable independent call processing agents to coordinate calls across the Internet without reliance on the PSTN.

- **Digital signal processor (DSP) resources for transcoding and conferencing:** A DSP changes the digitization (G.711 to G.729) or provides services such as conferencing (replication of a single inbound voice stream for transport to multiple endpoints).

Cisco CallManager

CallManager provides call processing for the Cisco AVVID IP telephony solution. This topic describes the CallManager features and introduces the deployment models.

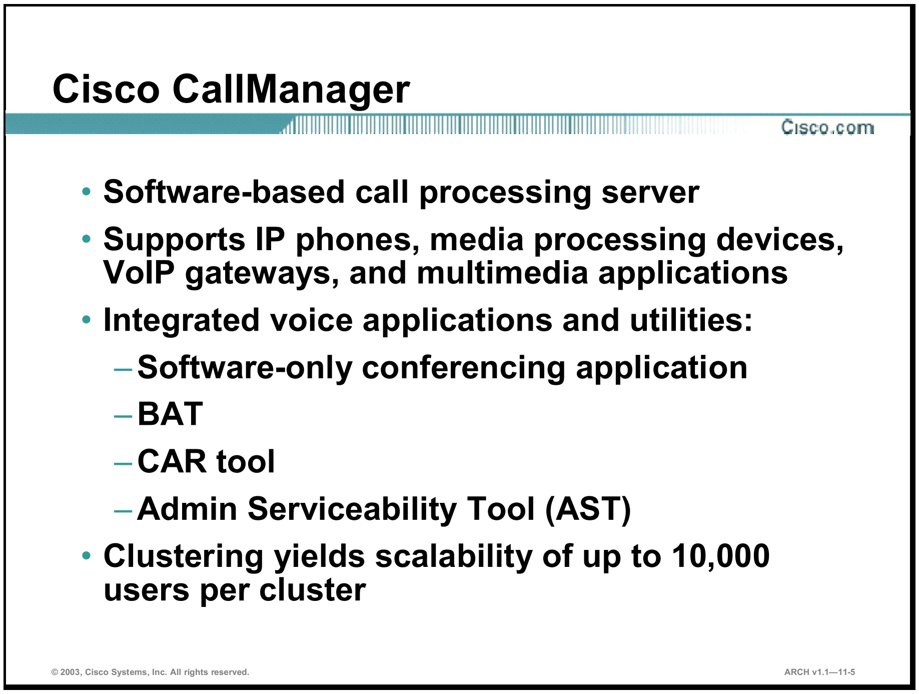

Cisco CallManager

- Software-based call processing server
- Supports IP phones, media processing devices, VoIP gateways, and multimedia applications
- Integrated voice applications and utilities:
  - Software-only conferencing application
  - BAT
  - CAR tool
  - Admin Serviceability Tool (AST)
- Clustering yields scalability of up to 10,000 users per cluster

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—11-5

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. Additional data, voice, and video services such as unified messaging, multimedia conferencing, collaborative contact centers, and interactive multimedia response systems interact with the IP telephony solution through Cisco CallManager's telephony application programming interfaces (APIs).

CallManager is installed on the Cisco Media Convergence Server and selected third-party servers. CallManager includes a suite of integrated voice applications and utilities; a software-only conferencing application; the Bulk Administration Tool (BAT); the CDR Analysis and Reporting (CAR) tool; and the Admin Serviceability Tool (AST).

Multiple CallManager servers are clustered and managed as a single entity. Clustering multiple call-processing servers on an IP network is a unique capability in the industry and highlights the leading architecture provided by Cisco AVVID. CallManager clustering yields scalability of up to 10,000 users per cluster. By interlinking multiple clusters, system capacity can be increased to as many as one million users in a 100-site system. Clustering aggregates the power of multiple, distributed CallManagers.

Call admission control ensures that voice quality of service (QoS) is maintained across constricted WAN links, and automatically diverts calls to alternative PSTN routes when WAN bandwidth is not available.

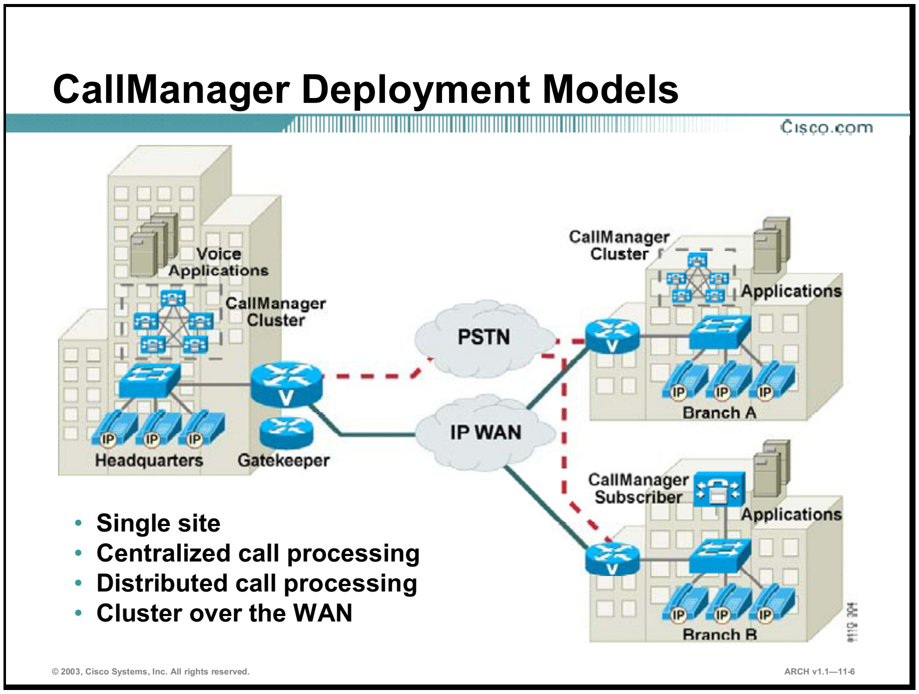

CallManager Deployment Models

- Single site
- Centralized call processing
- Distributed call processing
- Cluster over the WAN

ARCH v1.1—11-6

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.

There are several basic models for deploying the call processing capabilities of CallManager, depending on the size, geographical distribution, and functional requirements of your enterprise:

- Single-site call processing model
- Multisite WAN model with centralized call processing
- Multisite WAN model with distributed call processing
- Clustering over the IP WAN

The overall goals of an IP telephony network are to:

- Provide end-to-end IP telephony for network users
- Use the IP WAN as the primary voice path with the PSTN as the secondary voice path between sites
- Lower the total cost of ownership with greater flexibility
- Enable new applications

Note	The CallManager deployments are described in more detail in the Designing the Network for Cisco AVVID IP Telephony lesson.

Gateways and Control Protocols

Voice, video, and data gateway protocols are required for call routing and interoperability on a network that supports IP telephony. This topic lists voice, video, and data gateway components of the Cisco IP telephony solution and explains the role of each component in the overall IP telephony solution.

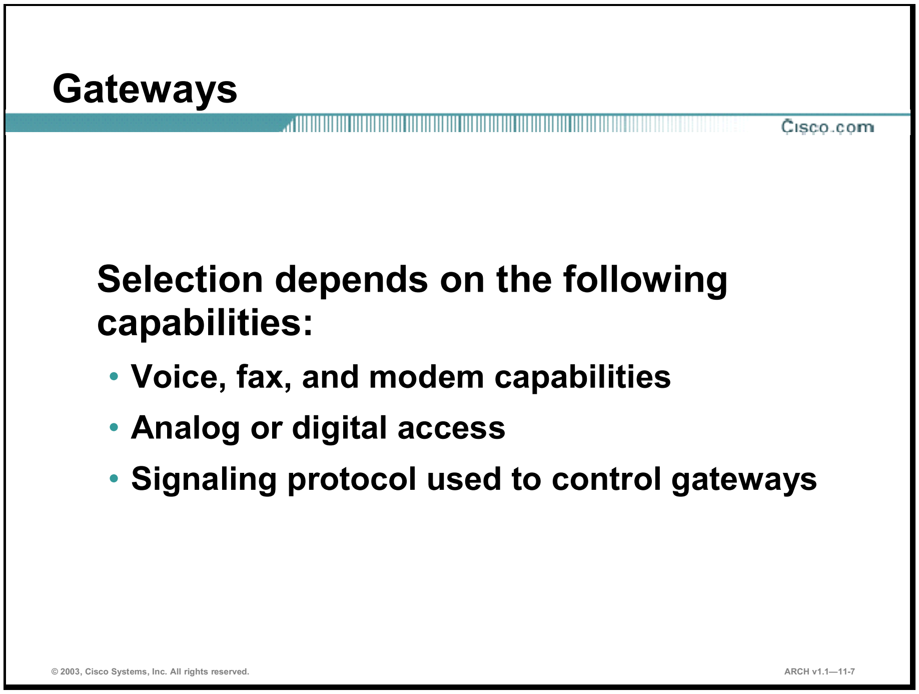

Gateways

Selection depends on the following capabilities:

- **Voice, fax, and modem capabilities**
- **Analog or digital access**
- **Signaling protocol used to control gateways**

ARCH v1.1—11-7

Gateway selection depends on the following requirements:

- Voice, fax, and modem capabilities
- Analog or digital access
- Signaling protocol used to control gateways

Gateway selection will depend on the type of platform already deployed. For example, a large campus with many Cisco Catalyst 6000 switches may opt to use the cards that fit within that chassis. A small site might use an existing Cisco IOS router with voice interface modules as an integrated solution. A non-IP voice mail system might also require gateways, which, in most cases, would be Media Gateway Control Protocol (MGCP) gateways.

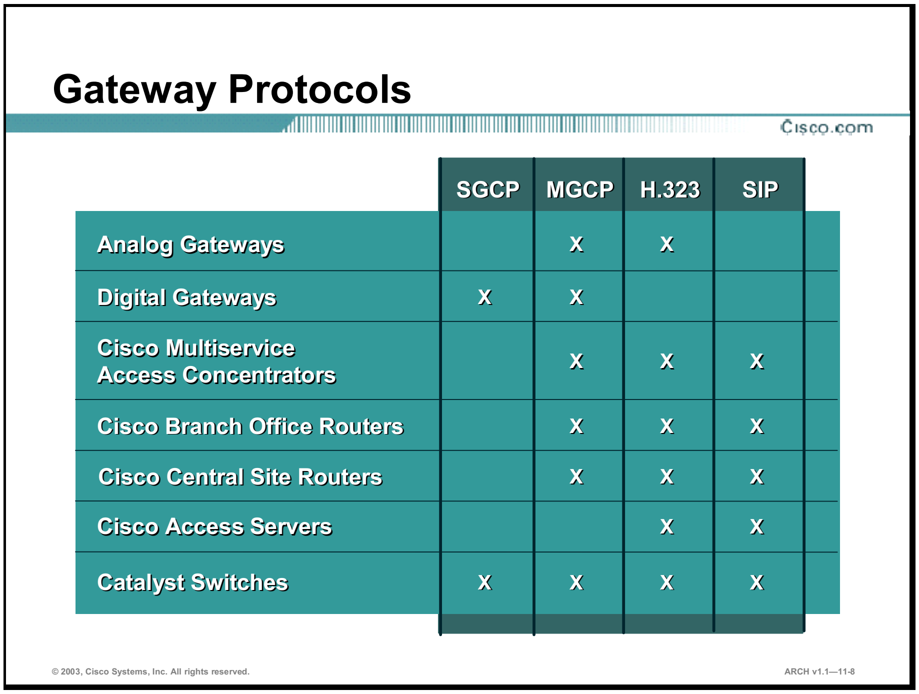

Gateway Protocols

Cisco.com

	SGCP	MGCP	H.323	SIP
Analog Gateways		X	X	
Digital Gateways	X	X		
Cisco Multiservice Access Concentrators		X	X	X
Cisco Branch Office Routers		X	X	X
Cisco Central Site Routers		X	X	X
Cisco Access Servers			X	X
Catalyst Switches	X	X	X	X

ARCH v1.1—11-8

Protocol selection depends on site-specific requirements and the installed base of equipment. CallManager supports the following gateway protocols:

- **Simple Gateway Control Protocol (SGCP):** Provides control between the CallManager and digital gateways.
- **MGCP:** Provides control between the CallManager and analog gateways.
- **H.323:** Provides control between Cisco IOS integrated router gateways and CallManager.
- **Session initiation protocol (SIP):** Internet Engineering Task Force (IETF) standard for multimedia conferencing over IP. SIP is an ASCII-based, application-layer control protocol (defined in RFC 2543) that can be used to establish, maintain, and terminate calls between two or more endpoints.

Most remote branch locations have Cisco 2600 or 3600 series routers installed that support the H.323 and MGCP protocols. For gateway configuration, an enterprise might prefer to implement MGCP over H.323 because MGCP offers simpler configuration and call survivability during a CallManager switchover from a primary to a secondary CallManager. Alternatively, H.323 might be preferred over MGCP because of the robustness of the interfaces supported.

Note SGCP functionality will be included in H.323 v.3.

Transcoders and Conferencing

CallManager provides access to a variety of media resources. A media resource is a software- or hardware-based entity that performs media processing functions on the voice data streams to which it is connected. Media processing functions include mixing multiple streams to create one output stream, passing the stream from one connection to another, or transcoding the data stream from one compression type to another. This topic describes the types and role of the transcoder and conferencing components of an IP telephony solution.

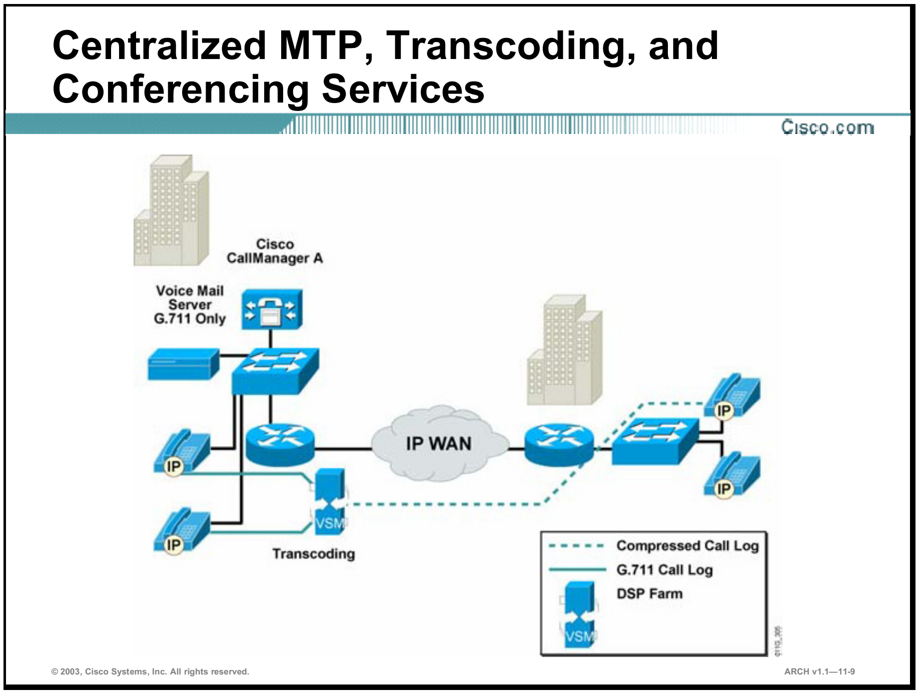

Centralized MTP, Transcoding, and Conferencing Services

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-9

The DSP resources provide hardware support for the CallManager IP telephony features. These features include hardware-enabled voice conferencing, hardware-based media termination point (MTP) support for supplementary services, and transcoding services.

Catalyst-enabled conferencing supports voice conferences in hardware. DSPs convert VoIP sessions into time-division multiplexing (TDM) streams for multiparty conference calls.

The Cisco MTP service can act either like the original software MTP resource or as a transcoding MTP resource. An MTP service can provide supplementary services such as hold, transfer, and conferencing when using gateways and clients that do not support H.323.

Transcoding compresses and decompresses voice streams to maximize use of WAN resources for VoIP traffic. A transcoder takes the output stream of one codec and converts it in real time (transcodes it) into an input stream for a different codec type. In other words, a transcoder converts a stream of one compression type into a stream of another compression type.

Transcoding is, in effect, an IP-to-IP voice gateway service. A transcoding node can convert a G.711 voice stream into a low bit-rate compressed voice stream, such as G.729a, which is critical for enabling applications such as IVR, voice messaging, and conference calls over low-speed IP WANs. In addition, a transcoder provides the capabilities of an MTP. You can use transcoding to enable supplementary services for H.323 endpoints when required.

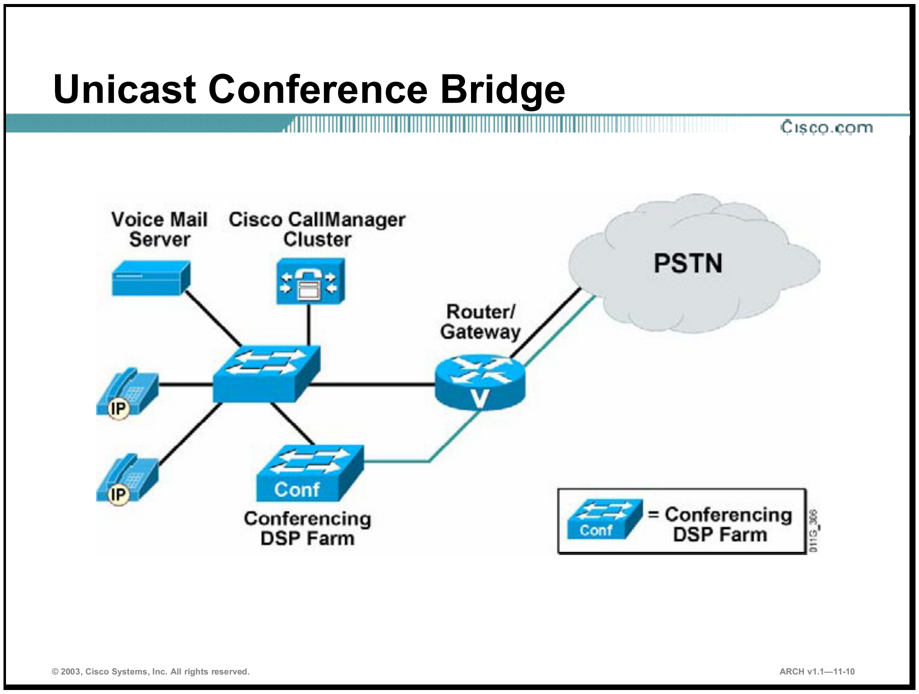

Unicast Conference Bridge

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-10

A unicast conference bridge is a device that accepts multiple connections for a given conference. It can accept any number of connections for a given conference, up to the maximum number of streams allowed for a single conference on that device. There is a one-to-one correspondence between media streams connected to a conference and participants connected to the conference. The conference bridge mixes the streams together and creates a unique output stream for each connected party. The output stream for a given party is usually the composite of the streams from all connected parties minus their own input stream. Some conference bridges mix only the three loudest talkers on the conference and distribute that composite stream to each participant (minus their own input stream if they are one of the talkers).

There are two types of conference bridges:

- **Software conference bridge:** A software unicast conference bridge is a standard conference mixer that is capable of mixing G.711 audio streams. Both a-law and mu-law streams may be connected to the same conference. The number of parties that can be supported on a given conference depends on the server where the conference bridge software is running and the configuration for that device.
- **Hardware conference bridge:** A hardware conference bridge has all the capabilities of a software conference bridge. In addition, some hardware conference bridges can support multiple low-bit-rate stream types such as G.729, global system for mobile communication (GSM), or G.723. This allows some hardware conference bridges to handle mixed-mode conferences. In a mixed-mode conference, the hardware conference bridge transcodes G.729, GSM, and G.723 streams into G.711 streams, mixes them, and then encodes the resulting stream into the appropriate stream type for transmission back to the user. Some hardware conference bridges support only G.711 conferences.

Cisco IP Telephony Applications

The primary Cisco IP telephony applications include Cisco Customer Response Solution, Cisco Conference Connection, Cisco IP Contact Center, and Cisco Unity. This topic lists the Cisco IP telephony application components of the Cisco IP telephony solution and explains the role of each component in the overall IP telephony solution.

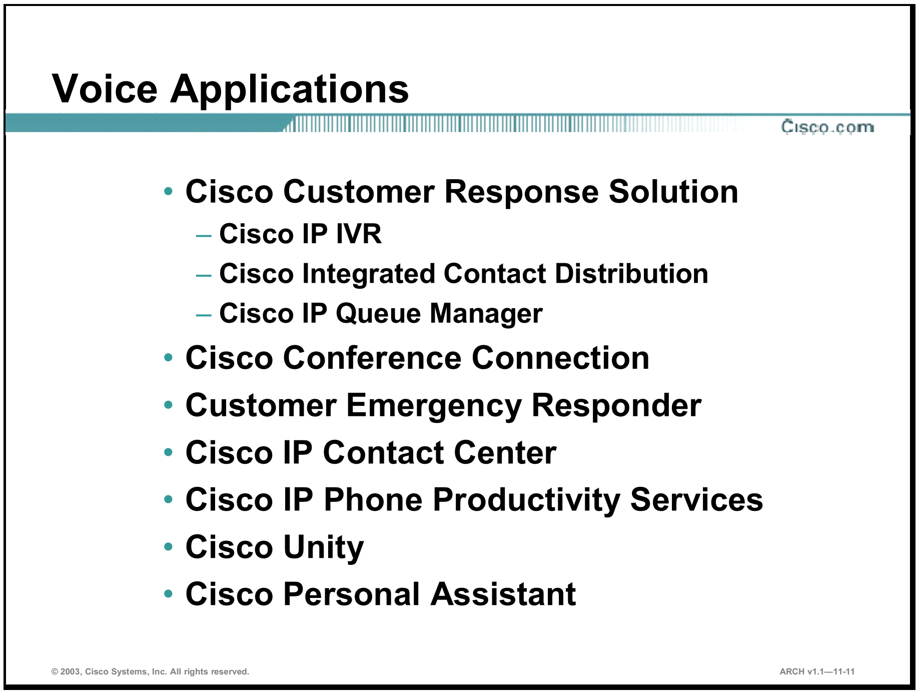

Voice Applications

- Cisco Customer Response Solution
  - Cisco IP IVR
  - Cisco Integrated Contact Distribution
  - Cisco IP Queue Manager
- Cisco Conference Connection
- Customer Emergency Responder
- Cisco IP Contact Center
- Cisco IP Phone Productivity Services
- Cisco Unity
- Cisco Personal Assistant

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—11-11

Cisco offers these voice applications:

- **Cisco Customer Response Solution (CRS):** Cisco CRS is an integrated platform that simplifies business integration, eases agent administration, increases agent flexibility, and provides efficiency gains in network hosting. This single-server integrated platform includes Cisco IP IVR, Cisco IP Integrated Contact Distribution (ICD), and Cisco IP Queue Manager with automatic call distribution (ACD) features, such as skills-based routing and priority queuing.
- **Cisco Conference Connection:** Cisco Conference Connection is a meet-me audio conference server that provides integrated operation with CallManager. Conferences are scheduled from an intuitive web-based conference scheduler. Conference participants call in to a central number, enter a meeting identification, and are then placed into the conference.
- **Cisco Emergency Responder:** Cisco Emergency Responder addresses the need to identify the location of 911 callers in an emergency, with no administration required when phones or people move from one location to another.
- **Cisco IP Contact Center:** The Cisco IP Contact Center (IPCC) delivers intelligent call routing, network-to-desktop computer telephony integration (CTI), and multimedia contact management to contact center agents over an IP network. It combines software automatic call distribution (ACD) functionality with IP telephony in a unified solution.

- **Cisco IP Phone Productivity Services:** The Cisco IP Phone services software developers kit (SDK) makes it easier for web developers to format and deliver content to IP Phones by providing web-server components for Lightweight Directory Access Protocol (LDAP) directory access, web proxy, and graphics conversion. It also contains several sample applications.
- **Cisco Unity:** Cisco Unity provides convergence-based communication services such as voice and unified messaging on a platform that offers the utmost in reliability, scalability, and performance. With Cisco Unity, users can listen to e-mail over the telephone, check voice messages from the Internet, and (when integrated with a supported third-party fax server) forward faxes to any local fax machine, increasing organizational productivity while improving customer service and responsiveness.
- **Cisco Personal Assistant:** Cisco Personal Assistant is a telephony application that operates with Cisco Unity and streamlines communications by helping users manage how and where they want to be reached.

Summary

This topic summarizes the key points discussed in this lesson.

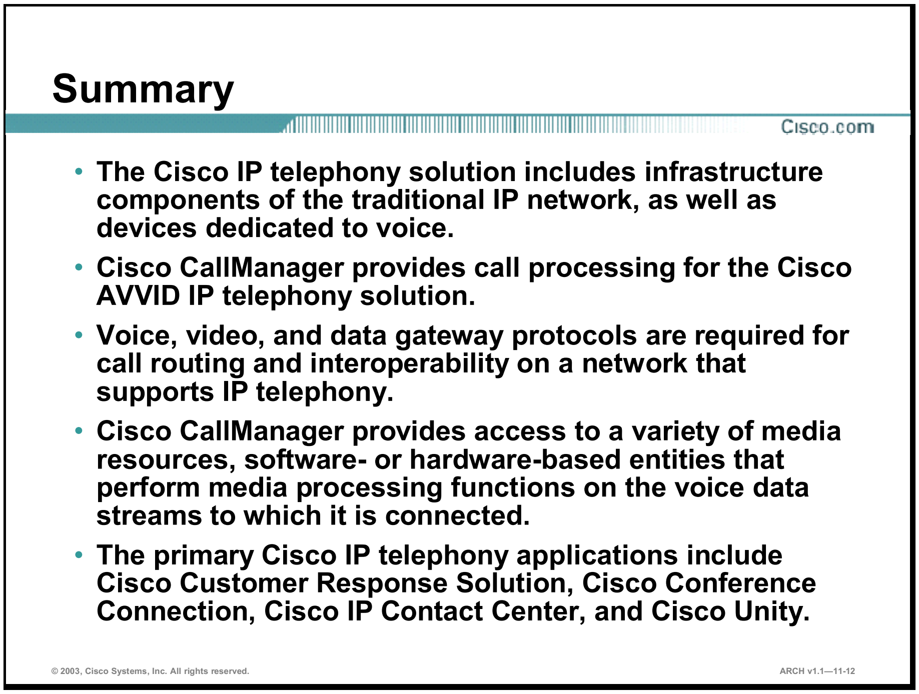

Summary

Cisco.com

- The Cisco IP telephony solution includes infrastructure components of the traditional IP network, as well as devices dedicated to voice.
- Cisco CallManager provides call processing for the Cisco AVVID IP telephony solution.
- Voice, video, and data gateway protocols are required for call routing and interoperability on a network that supports IP telephony.
- Cisco CallManager provides access to a variety of media resources, software- or hardware-based entities that perform media processing functions on the voice data streams to which it is connected.
- The primary Cisco IP telephony applications include Cisco Customer Response Solution, Cisco Conference Connection, Cisco IP Contact Center, and Cisco Unity.

ARCH v1.1—11-12

References

For additional information, refer to this resource:

- *Voice Solutions* at http://www.cisco.com/warp/public/44/solutions/network/voice.shtml

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 layer of the architecture must be in place before you add IP telephony to the network?

A) clients

B) applications

C) infrastructure

D) call processing

Q2) Which CallManager deployment model is self-contained, only requiring access to the PSTN for connectivity?

A) single-site

B) distributed

C) centralized

D) cluster over IP WAN

Q3) What is a major goal of implementing an IP telephony network?

A) to raise the cost of ownership

B) to lower the total cost of ownership

C) to raise the number of employees in a company

D) to lower the number of employees in a company

Q4) What is the purpose of a gateway?

A) to provide access only to the PSTN

B) to provide connection admission control

C) to provide compatibility for legacy telephony equipment

D) to provide PSTN access and limit access to the voice network

Q5) Which protocol is used to communicate between a gateway and a CallManager?

A) MGCP

B) SCGP

C) SCCP

D) MCGP

Q6) Which resource can provide both transcoding and conferencing from the same pool of resources?

A) Cisco CallManager

B) hardware transcoding

C) software conferencing

D) hardware conferencing

Q7) What application does Cisco Unity provide?

A) IP services

B) fax services

C) automated attendant

D) voice and unified messaging

Q8) Which application delivers intelligent call routing, network-to-desktop computer telephony integration (CTI), and multimedia contact management to contact center agents over an IP network?

A) Cisco Unity

B) Cisco IP Contact Center

C) Cisco Personal Assistant

D) Cisco IP Phone Productivity Services

Quiz Answer Key

Q1) C

Relates to: Introducing the Cisco IP Telephony Solution

Q2) A

Relates to: Introducing the Cisco IP Telephony Solution

Q3) B

Relates to: Cisco CallManager

Q4) C

Relates to: Gateways and Control Protocols

Q5) A

Relates to: Gateways and Control Protocols

Q6) D

Relates to: Transcoders and Conferencing

Q7) D

Relates to: Cisco IP Telephony Applications

Q8) B

Relates to: Cisco IP Telephony Applications

Designing the Network for Cisco IP Telephony

Overview

You can deploy Cisco IP telephony solutions in single-site, multisite, and clustering over IP WAN configurations. A multisite configuration deployment may support centralized or distributed call processing. Each deployment model has its own network design considerations.

Relevance

Each call processing modeling offers its own benefits for specific enterprise needs and situations. The challenge is to select the right model and design each component. Effective design of the underlying network infrastructure and Cisco AVVID components will provide an efficient, scalable, and available foundation for IP telephony networks.

Objectives

Upon completing this lesson, you will be able to design the physical and logical network, network features, and intelligent network services to support IP telephony, including Cisco CallManager clusters and call admission control functionality. This includes being able to meet these objectives:

- Design Cisco CallManager clusters, given specific telephony requirements
- Design standalone IP telephony solutions that communicate with other IP telephony systems over the PSTN
- Design a multisite IP telephony solution with centralized call processing, given two or more locations over an IP WAN
- Design a multisite IP telephony solution with distributed call processing, given two or more locations over an IP WAN
- Design a cluster over IP WAN telephony solution, given two or more locations
- Design the physical network and network features to support Cisco IP telephony solutions
- Propose intelligent network services required to support Cisco IP telephony solutions

Learner Skills and Knowledge

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

- IP telephony basics
- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course
- Reviewing the Cisco IP Telephony Solution lesson

Outline

The outline lists the topics included in this lesson.

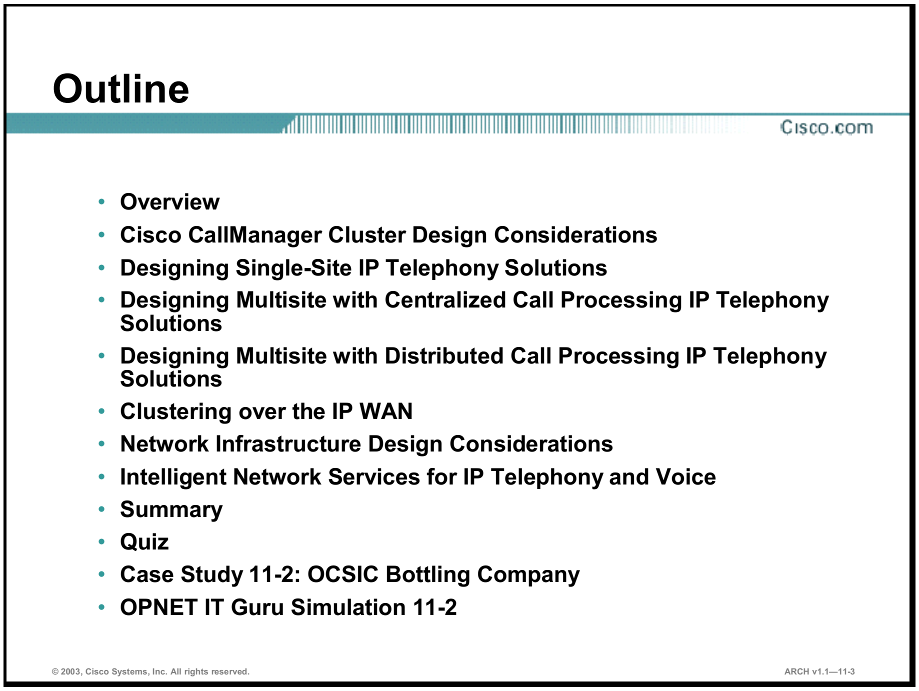

Outline

- Overview
- Cisco CallManager Cluster Design Considerations
- Designing Single-Site IP Telephony Solutions
- Designing Multisite with Centralized Call Processing IP Telephony Solutions
- Designing Multisite with Distributed Call Processing IP Telephony Solutions
- Clustering over the IP WAN
- Network Infrastructure Design Considerations
- Intelligent Network Services for IP Telephony and Voice
- Summary
- Quiz
- Case Study 11-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 11-2

ARCH v1.1—11-3

Cisco CallManager Cluster Design Considerations

A Cisco CallManager cluster may contain as many as 8 servers, of which a maximum of 6 provide CallManager call processing to support up to 10,000 phones. This topic helps you design CallManager clusters.

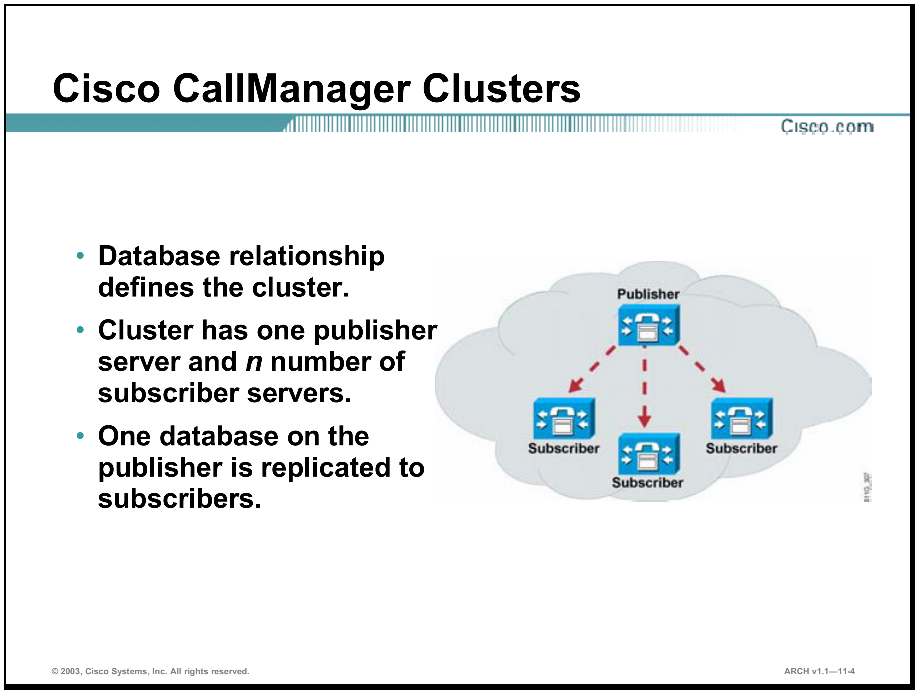

Cisco CallManager Clusters

Cisco.com

- Database relationship defines the cluster.
- Cluster has one publisher server and *n* number of subscriber servers.
- One database on the publisher is replicated to subscribers.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-4

A CallManager cluster consists of two or more CallManager servers that work together. A database defines the servers in the cluster. A cluster has one publisher (main) server and up to seven subscriber servers. The publisher maintains one database, which it replicates to the other servers in the cluster. Logically, a CallManager cluster is a single CallManager instance.

In addition to the CallManager, you can configure other servers, including:

- **Dedicated database publisher:** The database publisher synchronizes configuration changes and captures Call Detail Records (CDRs).
- **TFTP server:** The TFTP server stores configuration files, device loads (operating code), and ring types for downloading.
- **Music on hold (MoH) server:** MoH servers provide the music-on-hold functionality for a CallManager cluster.
- **Computer telephony integration (CTI) manager:** CTI manager servers are used to manage Telephony Application Programming Interface (TAPI), Java TAPI (JTAPI), or CTI devices.
- **Media streaming server:** A media streaming server, such as a conference bridge or MTP, is a separate server that registers with the CallManager cluster.

For large systems, Cisco recommends a dedicated database publisher and either a dedicated TFTP server or multiple load-balanced TFTP servers co-resident with CallManager. For smaller systems, you can combine the functions of database publisher and TFTP server.

These general guidelines apply to all clusters:

- A cluster may contain a mix of server platforms, but all CallManagers in the cluster must run the same software version.
- Within a cluster, you can enable a maximum of six servers with the CallManager service, and use other servers for more specialized functions such as TFTP, database publisher, music on hold, and so forth.
- All operating system and network services that are not required on a server should be disabled to maximize the server’s available resources.

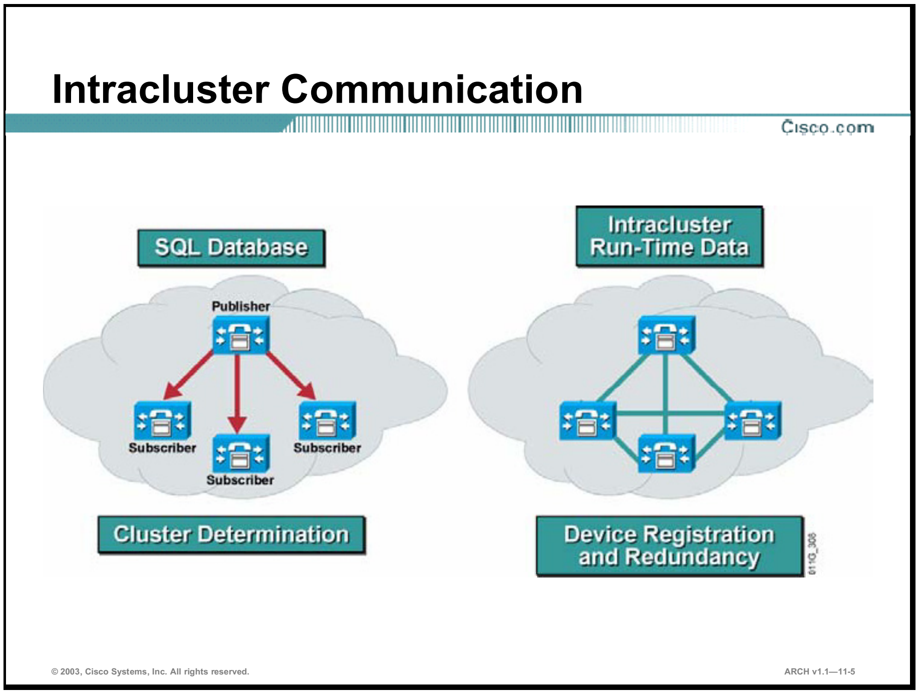

Intracluster Communication

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-5

There are two primary kinds of communication within a CallManager cluster. The first is a mechanism for distributing the database that contains all the device configuration information. The configuration database is stored on a publisher server and replicated to the subscriber members of the cluster. Changes made on the publisher are communicated to the subscriber databases, ensuring that the configuration is consistent across the members of the cluster, as well as facilitating spatial redundancy of the database.

The second intracluster communication is the propagation and replication of run-time data such as registration of devices, location bandwidth, and shared media resources. This information is shared across all members of a cluster running the CallManager Service, and it assures optimum routing of calls between members of the cluster and associated gateways.

LDAP directory information is also replicated between all servers in a cluster. The publisher replicates the LDAP directory to all other servers. You can integrate CallManager information into a corporate LDAP directory, such as Microsoft’s Active Directory or Netscape Directory. The replication is dependent on the integration method deployed.

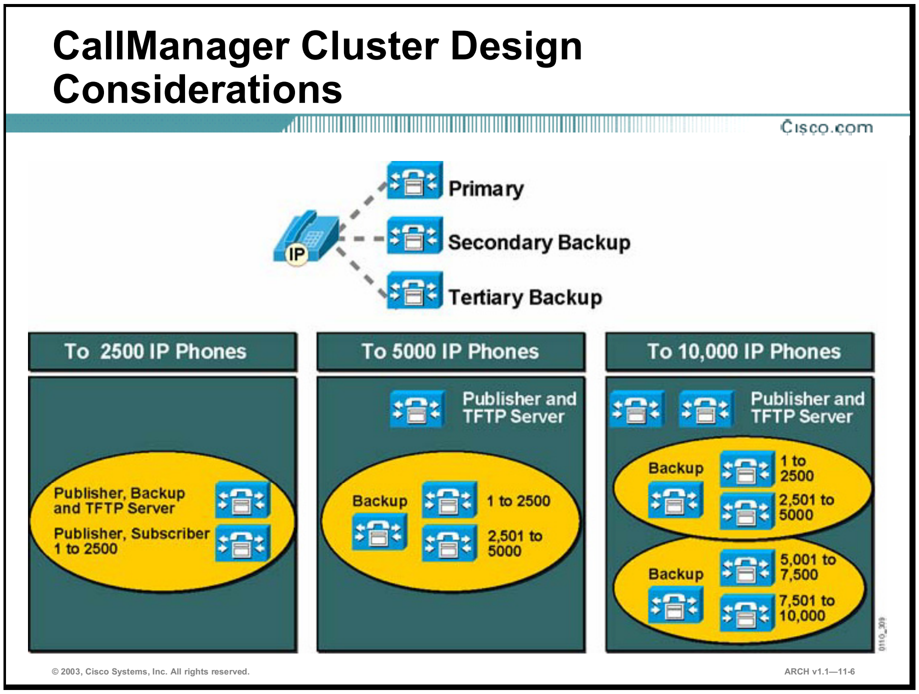

CallManager Cluster Design Considerations

ARCH v1.1—11-6

The figure shows the recommended cluster configuration for the small, medium, and large cluster.

The clustering options relate to the grouping of devices, usually phones and gateways. With the limits imposed on a single CallManager (device weights of 5000 per CallManager, not IP Phones) and good design practices, the minimum configuration consists of 2 CallManagers, which will support up to 2500 IP Phones. Cisco recommends 4 CallManagers to support 5000 IP Phones and up to 6 CallManagers to support up to 10,000 IP Phones.

In a small-scale environment (up to 2500 phones), Cisco recommends that you have a publisher and a subscriber. The publisher stores the master copy of the database while the subscriber is the device to which the phones register. In this scenario, the publisher is the backup server and the subscriber is the primary server. If the subscriber fails, the publisher becomes the primary CallManager for the cluster.

In the medium-sized cluster, there are 4 servers. The publisher acts as the TFTP server and is separate from the primary and backup servers. There would then be 2 primary systems that could register up to 5000 phones, and 1 server as the backup server to the 2 primary servers.

In the large-cluster design, there are up to 8 CallManagers in the cluster. The same basic scenario exists as for the 5000-phone service, but in this case there are 4 primary CallManager systems and 2 backup systems. The publisher is separate from the primary and backup systems, and holds the master database. The TFTP server acts independently of the publisher.

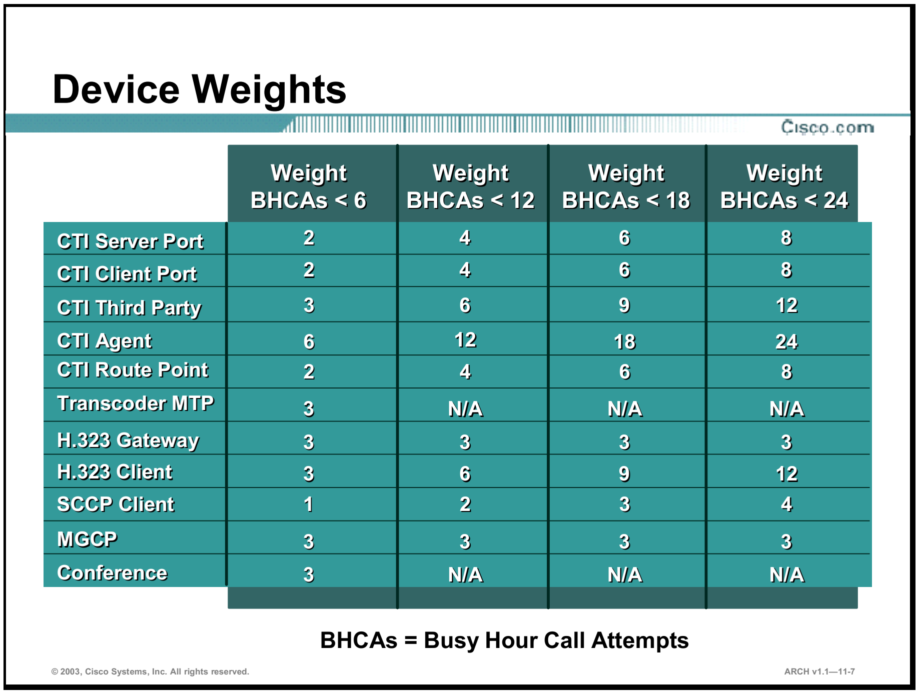

Device Weights

Cisco.com

	Weight BHCAs < 6	Weight BHCAs < 12	Weight BHCAs < 18	Weight BHCAs < 24
CTI Server Port	2	4	6	8
CTI Client Port	2	4	6	8
CTI Third Party	3	6	9	12
CTI Agent	6	12	18	24
CTI Route Point	2	4	6	8
Transcoder MTP	3	N/A	N/A	N/A
H.323 Gateway	3	3	3	3
H.323 Client	3	6	9	12
SCCP Client	1	2	3	4
MGCP	3	3	3	3
Conference	3	N/A	N/A	N/A

BHCAs = Busy Hour Call Attempts

ARCH v1.1—11-7

Each telephony device carries a different weight based on the amount of resources it requires from the server platform with which it is registered. The required resources include memory, processor, and I/O. Each device then consumes additional server resources during transactions, which are normally in the form of calls. For example, a device that makes only 6 calls per hour consumes fewer resources than a device making 12 calls per hour. As a common starting point, the base weight of a device is calculated with the assumption that it makes 6 or fewer calls per hour during its busiest hour, or 6 busy hour call attempts (BHCA).

The maximum number of IP Phones that can register with a single CallManager is 2500 on the largest server platforms, even if only IP Phones are registered. To calculate the number of IP Phones that can register with a CallManager in a specific deployment, subtract the weighted value of non-IP Phone resources from the maximum number of device units allowed for that platform. The maximum number of IP Phones may be lower than this calculated number, depending on the call volume per phone. In the case of the largest servers, the maximum number of device units is 5000.

The co-resident CTI manager allows a maximum of 800 CTI connections or associations per server, or a maximum of 3200 CTI connections or associations per cluster when they are equally shared between the 4 active CallManager servers. Associations are defined as devices that have been associated with a particular user in the CallManager user configuration.

CTI route points require some additional consideration. The base weight is 2, but the multiplier is based on the number of busy hour call completions (BHCC). To calculate the BHCC of a route point, we need to know how many calls we can expect to redirect to other ports through the route point. For example, in a typical IP IVR application, the IP IVR is expected to handle 10 simultaneous calls. The configuration for this requires a CTI route point and 10 CTI ports. If each IP IVR port expects 6 BHCC, then the route point can expect to redirect 6 calls per hour for each port, or a total of 60 calls per hour for the route point.

The multiplier for a CTI route point is calculated by taking the sum of the BHCC for all the ports associated with the CTI route point, dividing that sum by 6, and rounding up to the nearest whole number.

Designing Single-Site IP Telephony Solutions

The single-site IP telephony model offers the ability to use the PSTN for all off-net calls. In the LAN environment, there is sufficient bandwidth for voice traffic and the bandwidth allocations are less of a concern. In this topic, you will learn how to design standalone IP telephony solutions that communicate with other IP telephony systems over the PSTN.

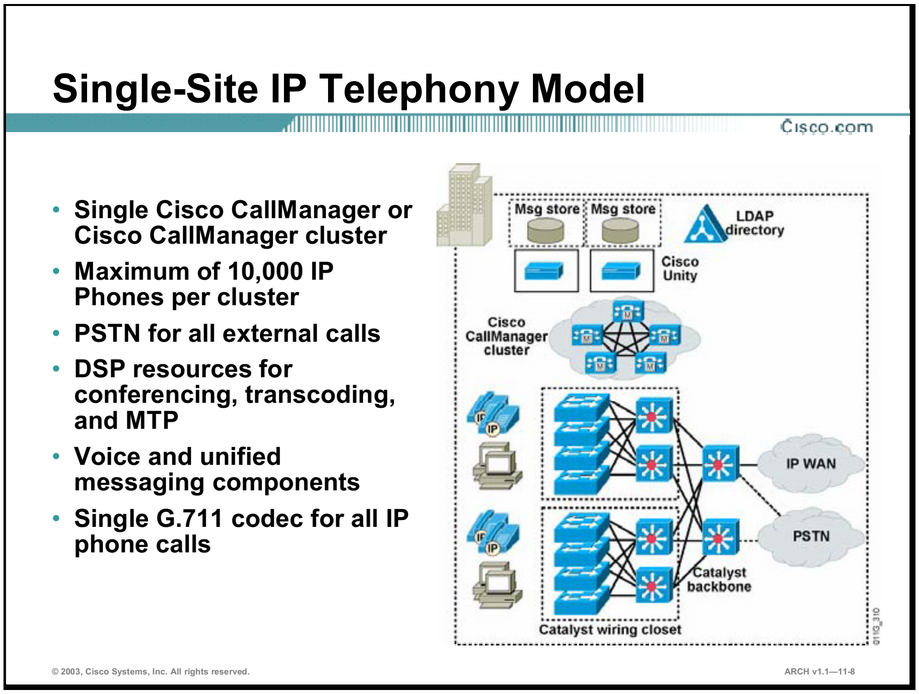

Single-Site IP Telephony Model

Cisco.com

- Single Cisco CallManager or Cisco CallManager cluster
- Maximum of 10,000 IP Phones per cluster
- PSTN for all external calls
- DSP resources for conferencing, transcoding, and MTP
- Voice and unified messaging components
- Single G.711 codec for all IP phone calls

ARCH v1.1—11-8

The single-site model for IP telephony consists of a call processing agent located at a single site, with no telephony services provided over an IP WAN. An enterprise would typically deploy the single-site model over a LAN or metropolitan-area network (MAN), which carries the VoIP traffic within served buildings or geographic areas. In this model, calls beyond the LAN or MAN use the PSTN.

The single-site model has the following design characteristics:

- Single CallManager or CallManager cluster
- Maximum of 10,000 IP Phones per cluster
- PSTN for all external calls
- DSP resources for conferencing, transcoding, and MTP
- Voice mail and unified messaging components
- G.711 codec used for all IP Phone calls (80 kbps of IP bandwidth per call, uncompressed)
- Capability to integrate with legacy PBX and voice mail systems

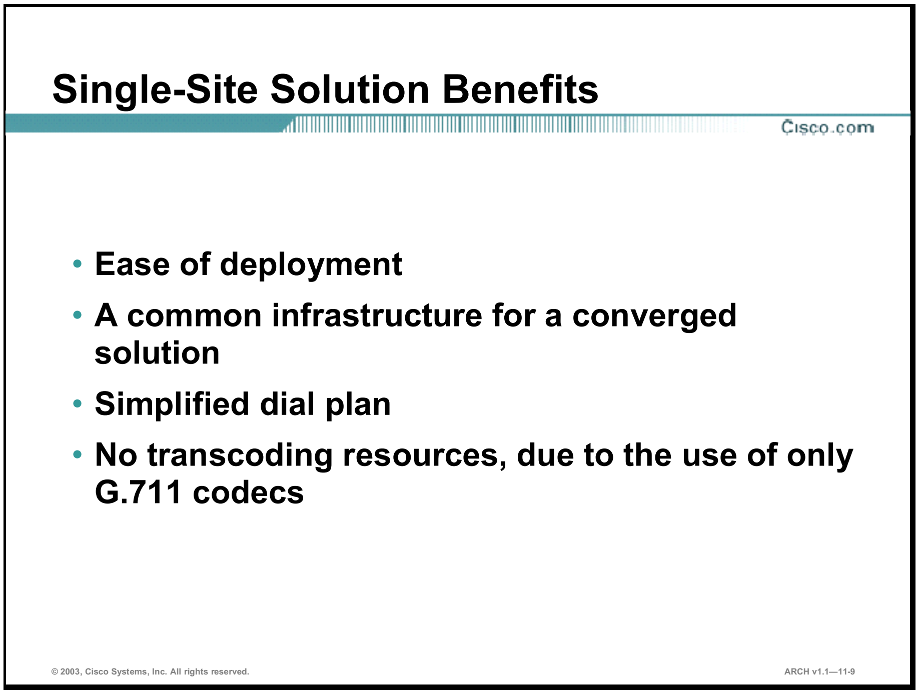

Single-Site Solution Benefits

- Ease of deployment
- A common infrastructure for a converged solution
- Simplified dial plan
- No transcoding resources, due to the use of only G.711 codecs

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-9

A single infrastructure for a converged network solution provides significant cost benefits and enables IP telephony to take advantage of the many IP-based applications in the enterprise. Single-site deployment also allows each site to be completely self-contained. There is no dependency for service in the event of an IP WAN failure or insufficient bandwidth, and there is no loss of call processing service or functionality.

In summary, the main benefits of the single-site model are:

- Ease of deployment
- A common infrastructure for a converged solution
- Simplified dial plan
- No transcoding resources required, due to the use of only G.711 codecs

The dial plan for the single-site model is usually the simplest of all the deployment models because of reliance on the PSTN for all off-net calls. However, there are some requirements for the dial plan for a single site, mainly to offer various classes of service, calling restrictions, 911 and E911 services, and security.

The CallManager dial-plan architecture can handle these general types of calls:

- All internal calls within the site
- External calls through a PSTN gateway

The complexity of your dial-plan configuration depends on the number of classes of service required by your specific enterprise policy. A class of service is a set of calling restrictions applied to a certain group of devices. Some examples are.

- Internal calls only
- Internal and local PSTN calls (no long-distance PSTN)
- Unrestricted calls (internal, local, and long-distance PSTN)

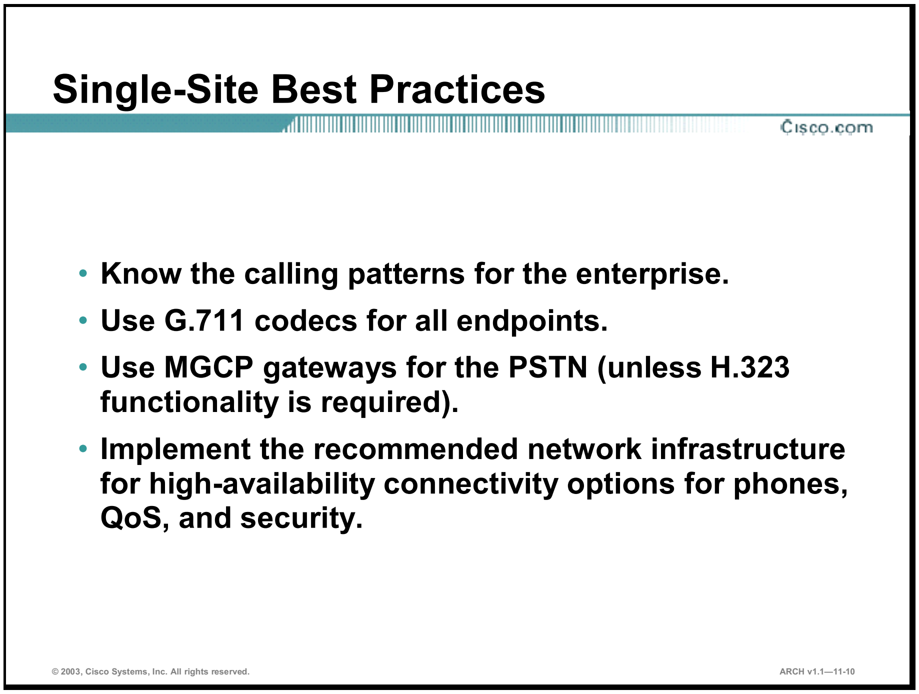

Single-Site Best Practices

Cisco.com

- Know the calling patterns for the enterprise.
- Use G.711 codecs for all endpoints.
- Use MGCP gateways for the PSTN (unless H.323 functionality is required).
- Implement the recommended network infrastructure for high-availability connectivity options for phones, QoS, and security.

ARCH v1.1—11-10

A single-site deployment is a subset of the distributed and centralized deployment models. The deployment should be over a stable infrastructure that allows easy migration to a voice and video network. Some guidelines specific to a single-site IP telephony deployment include:

- Know the calling patterns for the enterprise. Use the single-site model if most of the calls from the enterprise are within the same site or to PSTN users outside of the enterprise.
- Use G.711 codecs for all endpoints. This practice eliminates the consumption of DSP resources for transcoding, and those resources can be allocated to other functions such as conferencing and MTPs.
- Use MGCP gateways for the PSTN if the enterprise does not require H.323 functionality. This practice simplifies the dial-plan configuration. H.323 might be required to support specific functionality not offered with MGCP, such as support for Signaling System 7 (SS7) or Non-Facility Associated Signaling (NFAS).
- Implement the recommended network infrastructure for high-availability connectivity options for phones (inline power), QoS mechanisms, and security.

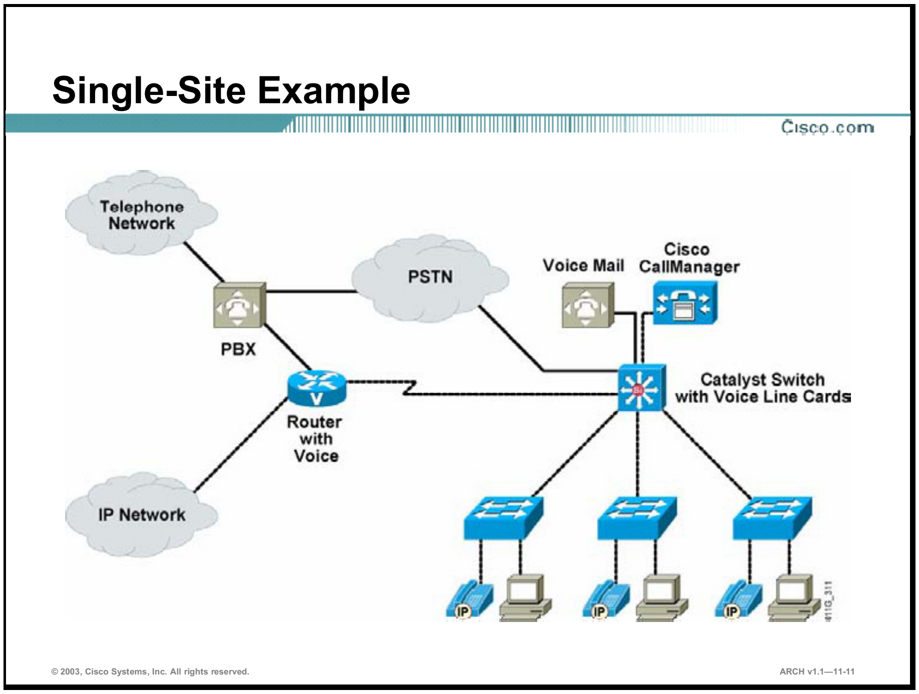

Single-Site Example

Cisco.com

ARCH v1.1—11-11

Company Background

NB is a global publishing company. The Singapore operation consists of approximately 6000 sales, printing, and writing staff members. The NB staff resides in a number of office buildings located in the vicinity of each other. In late 2000, NB added another building, the DBS building, to their campus. This additional building houses 750 employees. Rather than deploy a PBX in the new building, NB decided to deploy an IP Telephony solution. The deployment includes the network component.

IP Telephony Solution

The NB IP telephony solution is a single-site design. All IP telephony users are located in the DBS building, and are distributed across five floors. CallManagers, PSTN gateway, and voice mail are also physically located in the DBS building.

A MAN link connects the DBS building to the NBAP building less than 1 km away. This MAN link carries voice traffic across to NBAP, where a gateway connects into the worldwide NB PBX network. The diagram shows the DBS and NBAP buildings.

The DBS building has approximately 750 IP Phone 7960s. IP Phones connect to 10/100 ports on the Catalyst 4006, and receive inline power from the switch. Workstations connect to switch ports in the back of the IP Phone.

IP Phones and workstations are on separate VLANs and IP subnets.

Cisco CallManager and Call Admission Control

The NB CallManager deployment model is single-site with centralized call processing. No call admission control is required across the MAN. This is because the number of calls across the MAN is implicitly limited by the number of trunks connecting the gateway to the PBX.

One CallManager performs the database publishing function, and the other subscribes to the database. All IP Phones register with the subscriber as the primary CallManager, and use the publisher as the secondary CallManager.

Voice Mail Integration

CallManager connects to the company's voice mail system by means of two 24-port foreign exchange station (FXS) cards in the Catalyst 6509 switch. Only 30 of the available 48 ports are used. A 9600-bps simplified message desk interface (SMDI) link connects the primary CallManager to the voice mail device.

In addition to the voice mail system, there are other single points of failure:

- The SMDI link is not redundant. A failure of the primary CallManager will take the voice mail system out of service. Should this situation occur, the NB strategy is to manually move the SMDI cable to the backup CallManager. Alternatively, an SMDI splitter would allow both CallManagers to be connected at the same time, and allow for automatic failover.
- Currently both 24-port FXS cards reside in the same Catalyst 6509 chassis. A Catalyst 6509 failure will take the voice mail system out of service. As discussed earlier in this document, there is much to be gained in terms of resilience by adding a second Catalyst 6509.

Gateway Integration

The following three types of gateways exist:

- **H.323 gateway:** One Cisco 7200 connecting to legacy NB PBX network
- **PSTN gateway:** Three Catalyst 6509 E1 ports connecting to PSTN
- **Voice mail gateway:** Two Catalyst 6509 24-port FXS cards with 30 ports connecting to voice mail

Designing Multisite with Centralized Call Processing IP Telephony Solutions

In a centralized call processing system, CallManagers are centrally located at the hub or aggregation site, with no local call processing at the branch or remote office location. This topic shows you how to design a multisite IP telephony solution with centralized call processing, given two or more locations over an IP WAN.

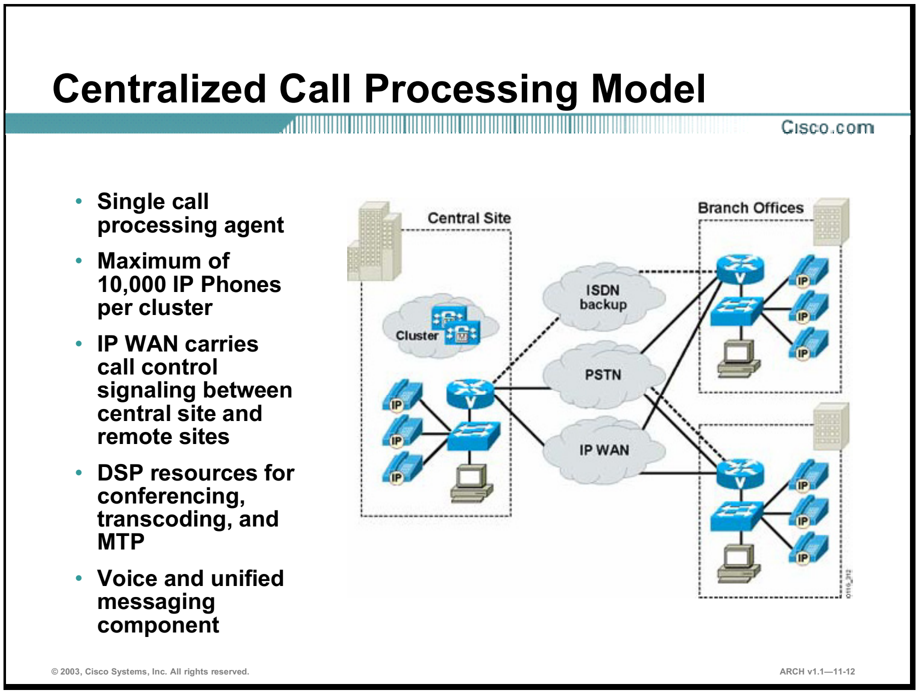

Centralized Call Processing Model

Cisco.com

- Single call processing agent
- Maximum of 10,000 IP Phones per cluster
- IP WAN carries call control signaling between central site and remote sites
- DSP resources for conferencing, transcoding, and MTP
- Voice and unified messaging component

ARCH v1.1—11-12

With centralized call processing, the CallManager cluster is located at the central site. Cisco CallManager supports a cluster of 10,000 IP Phones.

A primary advantage of the centralized IP telephony model is the ability to centralize call processing. This reduces the equipment required at the remote branch, and eliminates the administration of multiple PBXs or key systems. Dedicated plain old telephone service (POTS) lines or cellular phones can provide backup services.

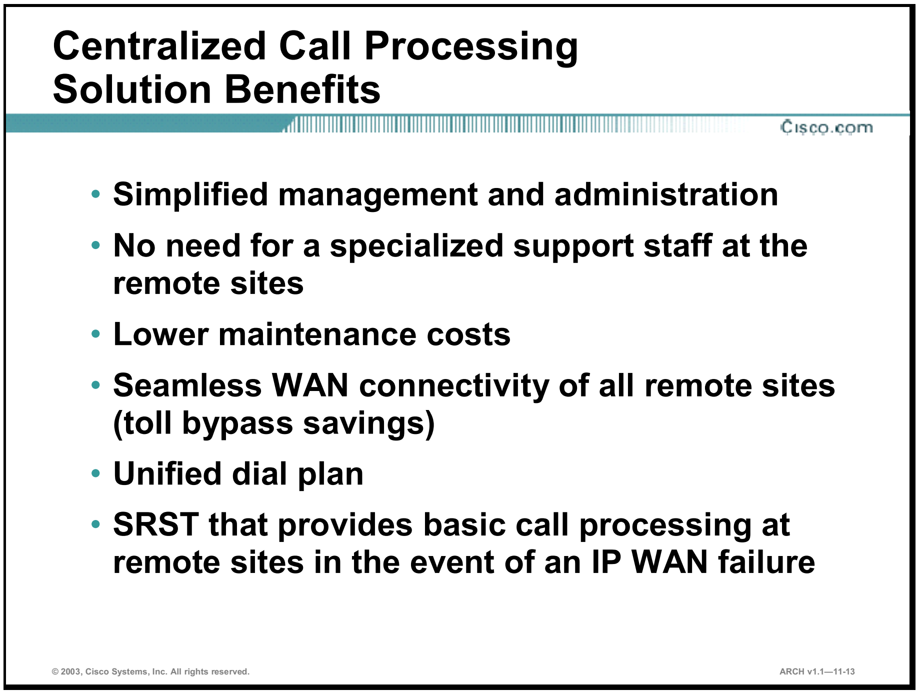

Centralized Call Processing Solution Benefits

- Simplified management and administration
- No need for a specialized support staff at the remote sites
- Lower maintenance costs
- Seamless WAN connectivity of all remote sites (toll bypass savings)
- Unified dial plan
- SRST that provides basic call processing at remote sites in the event of an IP WAN failure

The primary advantage of a distributed model is centralized call processing and applications. Centralized services reduce the equipment required at the remote sites and eliminate the administration and maintenance costs of multiple PBXs or key systems used in traditional telephony systems.

In summary, the multisite WAN model with centralized call processing provides the following benefits:

- Simplified management and administration
- No need for a specialized support staff at the remote sites
- Lower maintenance costs
- Seamless WAN connectivity of all remote sites (toll bypass savings)
- Unified dial plan
- Survivable Remote Site Telephony (SRST) that provides basic call processing at remote sites in the event of an IP WAN failure

Note In deployments where IP WAN bandwidth is either scarce or expensive with respect to PSTN charges, you can configure a remote site to place all external calls through the PSTN. In this scenario, the WAN link carries only regular data and call control signaling between the centralized CallManager cluster and the remote IP Phones and gateways. With the centralized call processing approach, there is no need for PBX equipment at the remote sites.

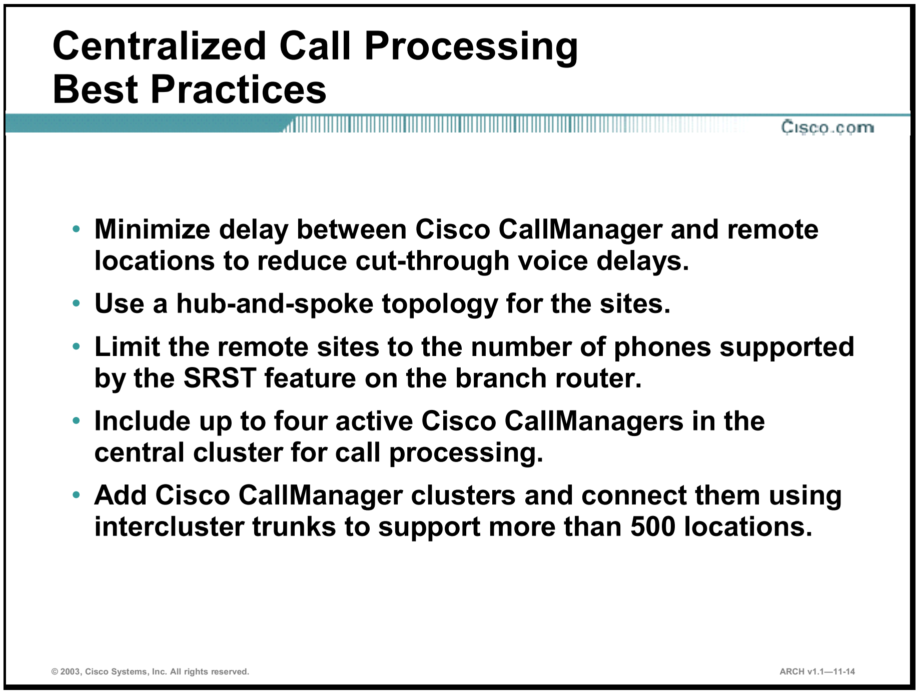

Centralized Call Processing Best Practices

- Minimize delay between Cisco CallManager and remote locations to reduce cut-through voice delays.
- Use a hub-and-spoke topology for the sites.
- Limit the remote sites to the number of phones supported by the SRST feature on the branch router.
- Include up to four active Cisco CallManagers in the central cluster for call processing.
- Add Cisco CallManager clusters and connect them using intercluster trunks to support more than 500 locations.

ARCH v1.1—11-14

Follow these guidelines and best practices when implementing the multisite WAN model with centralized call processing:

- Minimize delay between CallManager and remote locations to reduce voice cut-through delays (also known as clipping).
- Use a hub-and-spoke topology for the sites. The call admission control relies on the hub-and-spoke topology and records only the bandwidth into and out of each location.
- Limit the remote sites to the number of phones supported by the SRST feature on the branch router, or provide more branch routers.
- Configure up to four active CallManagers in the central cluster for call processing. This configuration can support a maximum of 10,000 IP Phones (or 20,000 device units) when CallManager runs on the largest supported server. Devices such as gateways, conferencing resources, voice mail, and other applications consume device units according to their relative device weights.
- Each CallManager cluster can support up to 500 locations configured with call admission control. If you need more remote sites, add CallManager clusters and connect them using intercluster trunks, as in the distributed call processing model.

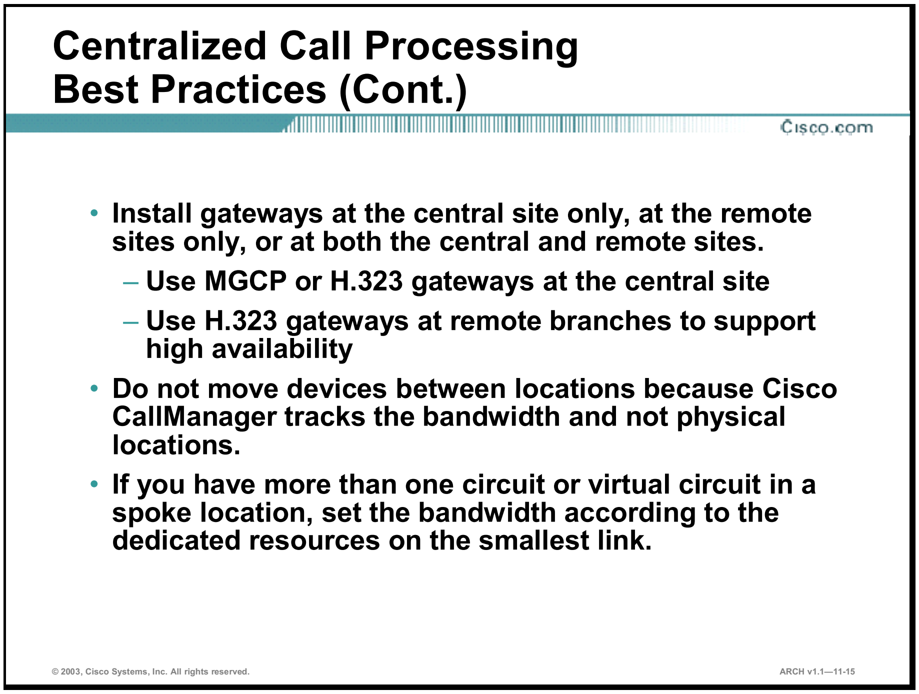

Centralized Call Processing Best Practices (Cont.)

- Install gateways at the central site only, at the remote sites only, or at both the central and remote sites.
  - Use MGCP or H.323 gateways at the central site
  - Use H.323 gateways at remote branches to support high availability
- Do not move devices between locations because Cisco CallManager tracks the bandwidth and not physical locations.
- If you have more than one circuit or virtual circuit in a spoke location, set the bandwidth according to the dedicated resources on the smallest link.

ARCH v1.1—11-15

When using the CallManager locations mechanism for call admission control, follow these recommendations:

- You can install gateways at the central site only, at the remote sites only, or at both the central and remote sites. Use the CallManager locations mechanism to provide call admission control for the gateways at the remote sites, but not at the central site. You do not need a Cisco IOS gatekeeper under these circumstances.
- Do not move devices between locations, because Cisco CallManager keeps track of the bandwidth only for the configured location of the device and not for the physical location.
- If you have more than one circuit or virtual circuit in a spoke location, set the bandwidth according to the dedicated resources allocated on the smallest link.

Designing Multisite with Distributed Call Processing IP Telephony Solutions

In the distributed call processing model, each CallManager cluster has its own set of resources and connects to the other clusters within the network via intercluster trunk links. This topic helps you design a multisite IP telephony solution with distributed call processing, given two or more locations over an IP WAN.

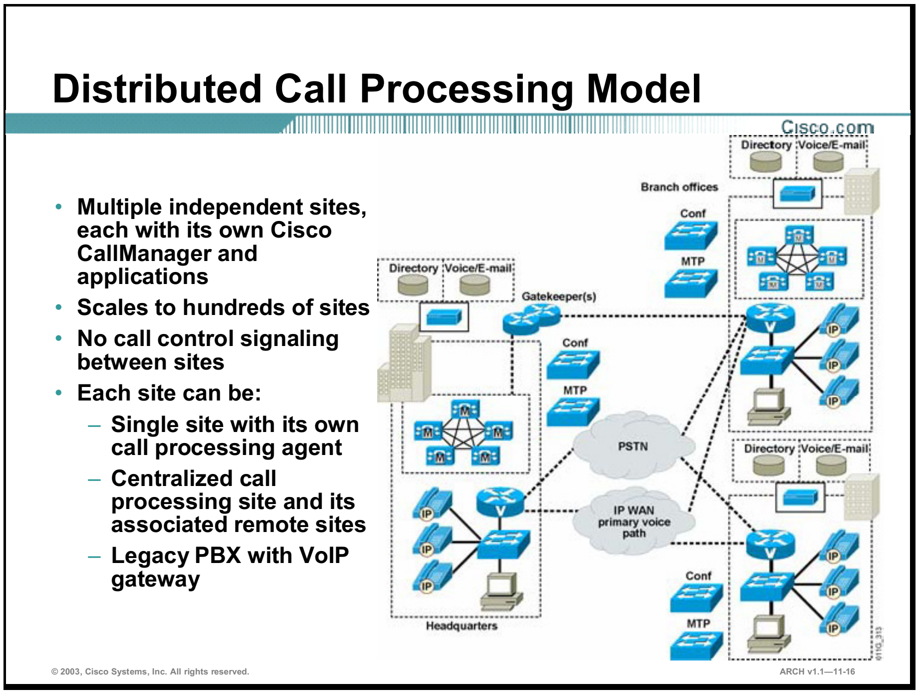

Distributed Call Processing Model

- Multiple independent sites, each with its own Cisco CallManager and applications
- Scales to hundreds of sites
- No call control signaling between sites
- Each site can be:
  - Single site with its own call processing agent
  - Centralized call processing site and its associated remote sites
  - Legacy PBX with VoIP gateway

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-16

CallManager supports up to 100 sites for distributed call processing. Voice calls between sites can use the IP WAN as the primary path, and the PSTN as the secondary path in the event the IP WAN is down or has insufficient resources to handle additional calls. Whether calls use the IP WAN or the PSTN can be transparent to both the calling party and the called party.

The primary advantage of the distributed call processing model is that, by using local call processing, it provides the same level of features and capabilities whether the IP WAN is available or not. Each site can have from one to eight CallManager servers in a cluster, based on the number of users. This is the predominant deployment model for sites with greater than 50 users, and each site can support up to 10,000 users. In addition, there is no loss of service if the IP WAN is down.

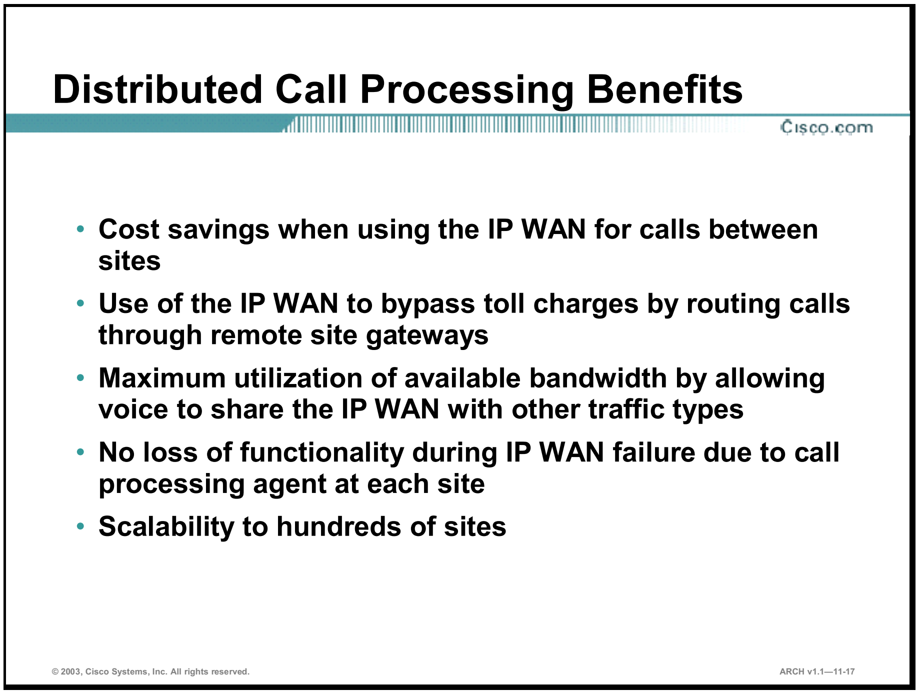

Distributed Call Processing Benefits

Cisco.com

- Cost savings when using the IP WAN for calls between sites
- Use of the IP WAN to bypass toll charges by routing calls through remote site gateways
- Maximum utilization of available bandwidth by allowing voice to share the IP WAN with other traffic types
- No loss of functionality during IP WAN failure due to call processing agent at each site
- Scalability to hundreds of sites

ARCH v1.1—11-17

The multisite WAN model with distributed call processing provides these benefits:

- Cost savings when using the IP WAN for calls between sites.
- Use of the IP WAN to bypass toll charges by routing calls through remote-site gateways, closer to the PSTN number dialed. This practice is known as tail-end-hop-off.
- Maximum utilization of available bandwidth by allowing voice to share the IP WAN with other types of traffic.
- No loss of functionality during IP WAN failure, because there is a call processing agent at each site.
- Scalability to hundreds of sites.

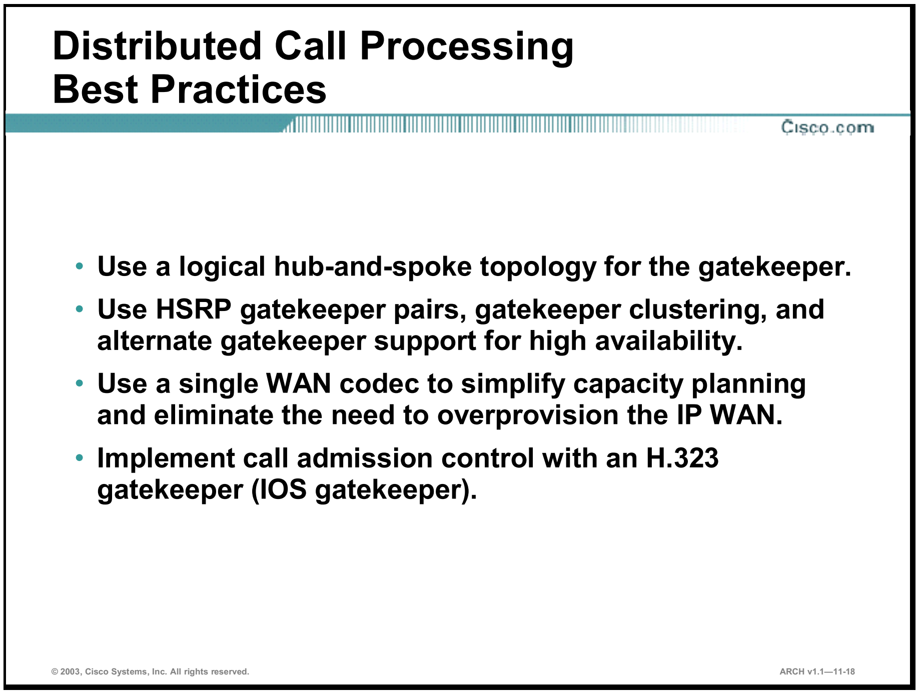

Distributed Call Processing Best Practices

- Use a logical hub-and-spoke topology for the gatekeeper.
- Use HSRP gatekeeper pairs, gatekeeper clustering, and alternate gatekeeper support for high availability.
- Use a single WAN codec to simplify capacity planning and eliminate the need to overprovision the IP WAN.
- Implement call admission control with an H.323 gatekeeper (IOS gatekeeper).

A multisite WAN with distributed call processing is a superset of the single-site deployment. The best practices should include those from the single-site deployment as well as those listed in the figure. The key difference in the single-site and the multisite distributed site is how to control the voice calls over the WAN. The gatekeeper is responsible for controlling voice calls in the distributed mode. The gatekeeper performs two main functions:

- Call admission control
- E.164 dial-plan resolution

In the distributed site, the number of sites is limited by the hub-and-spoke environment, while the gatekeeper can scale to hundreds of sites.

A gatekeeper is one of the key elements in the multisite WAN model with distributed call processing.

A gatekeeper is an H.323 device that provides call admission control and PSTN dial-plan resolution. The following best practices apply to the use of a gatekeeper:

- Use a logical hub-and-spoke topology for the gatekeeper. A gatekeeper can manage the bandwidth into and out of a site, or between zones within a site, but it is not aware of the topology.
- To provide high availability of the gatekeeper, use Hot Standby Router Protocol (HSRP) gatekeeper pairs, gatekeeper clustering, and alternate gatekeeper support. In addition, use multiple gatekeepers to provide redundancy within the network.
- Use a single WAN codec because the H.323 specification does not allow for Layer 2, IP, User Datagram Protocol (UDP), or Real-Time Transport Protocol (RTP) header overhead in the bandwidth request. (Header overhead is allowed only in the payload or encoded voice part of the packet.) Use one type of codec on the WAN to simplify capacity planning, by eliminating the need to overprovision the IP WAN to allow for the worst-case scenario.
- Gatekeeper networks can scale to hundreds of sites, and the design is limited only by the hub-and-spoke topology.

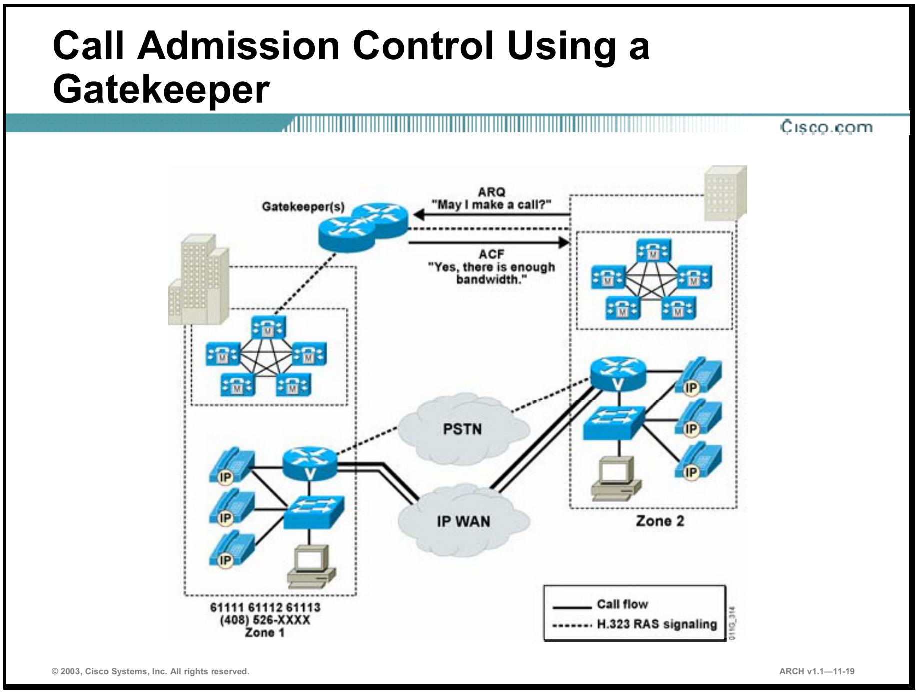

Call Admission Control Using a Gatekeeper

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-19

A distributed call processing system requires call admission control, just as a centralized call processing system does. However, the mechanism for implementing call admission control differs greatly in these two types of systems.

For distributed call processing systems, you can implement call admission control with an H.323 gatekeeper. In this design, the call processing agent registers with the IOS gatekeeper, and queries it each time the agent wants to place an IP WAN call. The IOS gatekeeper associates each call processing agent with a zone that has specific bandwidth limitations. Therefore, the IOS gatekeeper can limit the maximum amount of bandwidth consumed by IP WAN voice calls into or out of a zone.

The figure shows call admission control with a gatekeeper. When the call processing agent wants to place an IP WAN call, it first requests permission from the gatekeeper. If the gatekeeper grants permission, the call processing agent places the call across the IP WAN. If the gatekeeper denies the request, the call processing agent places the call across the IP WAN, and the call processing agent can try a secondary path, such as the PSTN, or it can simply fail the call.

This design essentially consists of a call accounting method for providing admission control, in which the gatekeeper keeps track of the bandwidth consumed by the IP WAN calls. When you set the maximum bandwidth for a zone, take into account the limitation that voice traffic should not consume more than 75 percent of the WAN link.

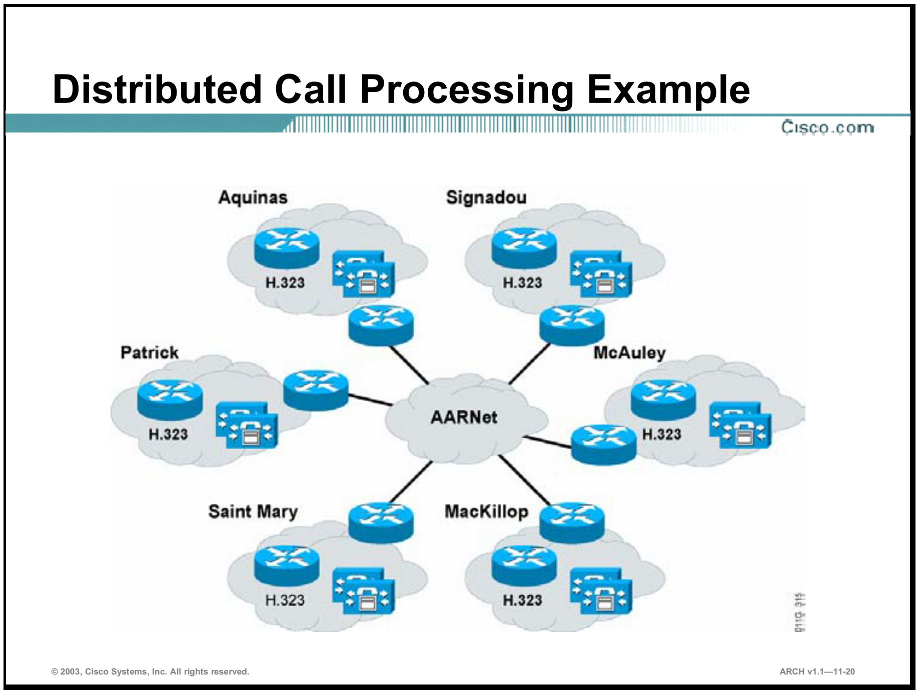

Distributed Call Processing Example

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-20

Company Background

The ACU is a public government-funded university established in 1991. The university has approximately 10,000 learners and approximately 1000 staff. There are six campuses spread across Ireland.

The current campus design does not comply with the Cisco-recommended QoS design guidelines for IP telephony. These concerns are in regards to QoS:

- The broadcast domain is very large. IP Phones may be affected by the excessive amount of broadcasts they process.
- Catalyst 1900 switches within each campus are not QoS-capable. If an IP Phone and PC are connected to the same switch port, voice packets may be dropped if the PC is receiving data at a high rate.

Significant improvements can be achieved by redesigning parts of the campus infrastructure. A hardware upgrade is not necessarily required.

IP Telephony Solution

The ACU recently deployed an IP telephony solution. The solution consists of a two-CallManager cluster and a Cisco 3640 gateway at each campus, along with a number of IP Phones. The six campuses are interconnected by a WAN provided by an international service provider.

The campus is now split into a voice VLAN and a data VLAN. Phones and PCs that connect to a Catalyst 1900 switch must now connect to different ports to achieve the VLAN separation. An additional uplink from each Catalyst 1900 switch to a Cisco 3500 Series XL switch is added. One of the two uplinks is a member of the voice VLAN; the other uplink is a member of the data VLAN. Using Inter-Switch Link (ISL) trunking as an alternative to having two uplinks is not recommended, as this will not provide the voice and data traffic with separate queues. The gigabit Ethernet links from the Catalyst 3500XL to the Catalyst 6000 must also be converted to 802.1Q trunks so that both voice and data VLAN can be carried across this core switch.

Ports on the Catalyst 3500XL that are in the data VLAN have a default class of service (CoS) of zero. Ports that are a member of the voice VLAN have a default CoS of 5. As a result, the voice traffic will be correctly prioritized once it arrives at a Catalyst 3500/Catalyst 6500 core.

In the rare case when IP Phones connect directly to a Catalyst 3500XL, a PC may be connected to the rear switch port on the IP Phone. In this case, the IP Phones connect to the switch by means of an 802.1Q trunk. This allows voice and data packets to travel on separate VLANs, and packets can be given the correct CoS at ingress. As the network evolves over time, and the Catalyst 1900s reach end-of-life, they should be replaced with Catalyst 3500XL switches or other QoS-capable switches. This topology then becomes the standard method of connecting IP Phones and PCs to the network.

Gateways

Each of the six ACU campuses has a Cisco 3640 router acting as an H.323 gateway. These gateways connect to the PSTN by means of ISDN. The number of PRIs and bearer channels (B Channels) varies depending on the size of the campus.

For direct outward dialing (DOD), these gateways are used only as secondary gateways. The service provider’s gateways are used as the primary gateways. For direct inward dialing (DID), the ACU gateways are always used.

Cisco CallManager

Each of the six campuses has a cluster consisting of two CallManager servers. One CallManager is the publisher and the other CallManager is the subscriber. The subscriber acts as the primary CallManager for all IP Phones.

Each cluster is configured with two regions: one for intracampus calls (G.711) and the other for intercampus calls (G.729).

Location-based Call Admission Control (CAC) is not appropriate for ACU, as all IP Phones served by each cluster are on a single campus. There are merits to a gatekeeper-based CAC for intercampus calls, but this is not currently implemented. There are, however, plans to do so in the near future.

Each CallManager is configured with 22 H.323 gateways. This is made up of intercluster trunks to the five other CallManager clusters, six service-provider PSTN gateways, and one ACU gateway at each campus.

International calls are gatekeeper routed and not sent through the local gateway. This is significant, as the service provider may deploy international gateways in the future. If a gateway was later deployed in the United States, a simple gatekeeper configuration change would allow universities to place calls to the United States at U.S. domestic rates.

Voice Mail

Prior to the migration to IP telephony, ACU had three Active Voice Repartee OS/2-based voice mail servers with Dialogic phone boards. The plan is to reuse these servers in the IP telephony environment. When implemented, each Repartee server will connect to a CallManager by means of SMDI and a Catalyst 6000 24-port FXS card. This provides voice mail for three of the six campuses, leaving three campuses without voice mail. It is not possible to properly share one Repartee server between users on two CallManager clusters, as there is no way of propagating the Message Waiting Indicator (MWI) across the intercluster H.323 trunk.

ACU is also considering purchasing three Cisco Unity servers for the remaining three campuses. These servers will be Skinny Station Protocol-based, so no gateways will be required.

Media Resources

Hardware DSPs are not currently deployed at ACU. Local conferencing uses the software-based conference bridge on the CallManager. Intercluster conferencing is not currently supported. Transcoding is currently not required. Only G.711 and G.729 codecs are used and supported on all deployed end devices.

Fax and Modem Support

Fax and modem traffic is not currently supported by the ACU IP telephony network. The university is planning to utilize the Catalyst 6000 24-port FXS card for fax/modem relay in the future.

Clustering over the IP WAN

You may deploy a single CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled. This topic provides a brief overview of how to design a cluster over WAN telephony solution, given two or more locations.

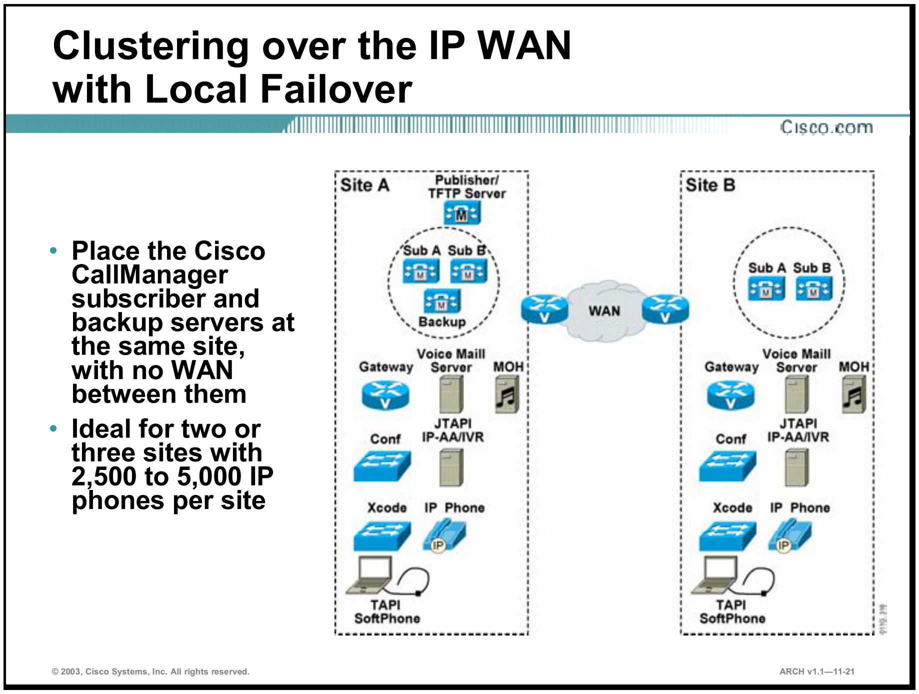

Clustering over the IP WAN with Local Failover

Cisco.com

- Place the Cisco CallManager subscriber and backup servers at the same site, with no WAN between them
- Ideal for two or three sites with 2,500 to 5,000 IP phones per site

ARCH v1.1—11-21

Clustering over the WAN can support two types of deployments: local failover and remote failover. In addition, clustering over the WAN supports a single point of administration for IP Phones for all sites within the cluster, feature transparency, shared line appearances, extension mobility within the cluster, and a unified dial plan. These features make this solution ideal as a disaster recovery plan for business continuance sites or as a single solution for small or medium sites.

Local failover requires that you place the CallManager subscriber and backup servers at the same site with no WAN between them. This deployment model is ideal for two or three sites with CallManager servers and a maximum of 5000 and 2500 IP Phones per site, respectively. This model allows for up to 10,000 IP Phones in the two-site configuration and 7500 IP Phones in the three-site configuration.

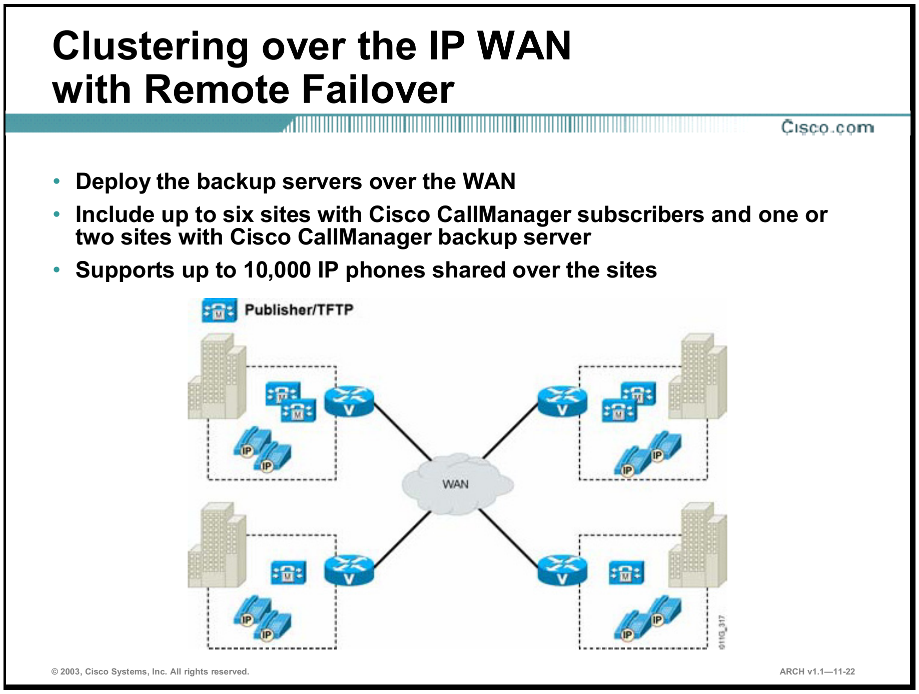

Clustering over the IP WAN with Remote Failover

Cisco.com

- Deploy the backup servers over the WAN
- Include up to six sites with Cisco CallManager subscribers and one or two sites with Cisco CallManager backup server
- Supports up to 10,000 IP phones shared over the sites

ARCH v1.1—11-22

Remote failover allows you to deploy the backup servers over the WAN. Using this employment model, you may have up to six sites with CallManager subscribers and one or two sites containing the CallManager backup server. This deployment allows for up to 10,000 IP Phones shared over the required number of sites.

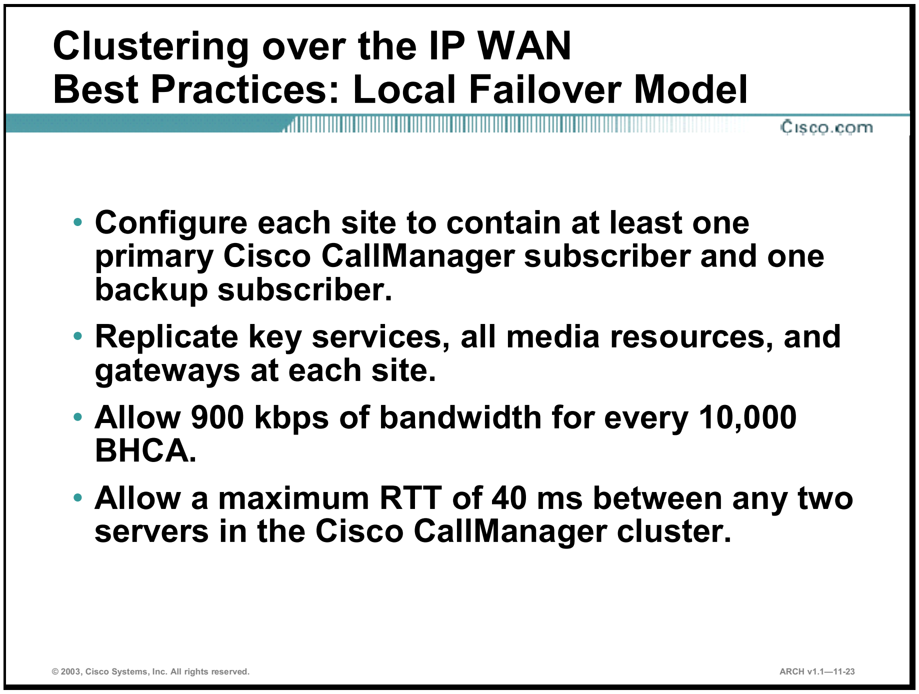

Clustering over the IP WAN
Best Practices: Local Failover Model

Cisco.com

- Configure each site to contain at least one primary Cisco CallManager subscriber and one backup subscriber.
- Replicate key services, all media resources, and gateways at each site.
- Allow 900 kbps of bandwidth for every 10,000 BHCA.
- Allow a maximum RTT of 40 ms between any two servers in the Cisco CallManager cluster.

ARCH v1.1—11-23

The local failover deployment model provides the most resilience for clustering over the WAN. Each of the sites in this model contains at least one primary CallManager subscriber and one backup subscriber. This configuration allows for either a two-site deployment with 5000 IP Phones per site or a three-site deployment with 2500 IP Phones per site.

In summary, observe these guidelines when implementing the local failover model:

- Configure each site to contain at least one primary and one backup CallManager subscriber.
- Cisco highly recommends that you replicate key services (TFTP, DNS, Dynamic Host Configuration Protocol [DHCP], LDAP, and IP Phone Services), all media resources (conference bridges and MoH), and gateways at each site to provide the highest level of resiliency. You could also extend this practice to include a voice mail system at each site. In the event of a WAN failure, only sites without access to the publisher database might lose a small amount of functionality.
- Every 10,000 BHCAs in the cluster requires 900 kbps of bandwidth for Intra-Cluster Communication Signaling (ICCS). This is a minimum bandwidth requirement, and bandwidth is allocated in multiples of 900 kbps.
- Allow a maximum round-trip time (RTT) of 40 ms between any two servers in the CallManager cluster. This time equates to a 20-ms maximum one-way delay, or a transmission distance of approximately 1860 miles (3000 km) under ideal conditions.

Cisco CallManager Provisioning

Provisioning of the CallManager cluster for the local failover model should follow the design guidelines for device weights. If calls are allowed across the WAN between the sites, you must configure CallManager locations in addition to the default location for the other sites, to provide call admission control between the sites. If the bandwidth is over-provisioned for the number of devices, it is still best to configure call admission control based on locations. Because call admission control based on locations does not provide automatic failover to the PSTN, Cisco recommends that you over-provision the WAN for intersite calls.

As the delay increases between the CallManager servers, the bandwidth information shared between the servers for call admission control will also be delayed, possibly allowing additional calls to be set up until the ICCS bandwidth message arrives. This situation is a small risk because, even at full capacity, only a few additional calls might be set up before the bandwidth information arrives. To provide for this situation, Cisco recommends that you over-provision the priority queue for voice traffic by a few extra calls.

To improve redundancy, Cisco recommends that you enable TFTP service on at least one of the CallManager servers at each location. You can run the TFTP service on either a publisher or a subscriber server, depending on the site. The TFTP server option must be correctly set on the DHCP servers for each site. If DHCP is not in use or the TFTP server is manually configured, you should configure the correct address for the site.

Other services, which may affect normal operation of CallManager during WAN outages, should also be replicated at all sites to ensure uninterrupted service. These services include DHCP servers, DNS servers, corporate directories, and IP Phone services. On each DHCP server, set the DNS server address correctly for each location.

IP Phones may have shared line appearances between the sites. The ICCS bandwidth provisioned between the sites allows for the additional ICCS traffic that shared line appearances generate. During a WAN outage, call control for each line appearance is segmented, but call control returns to a single CallManager server once the WAN is restored. During the WAN restoration period, there is extra traffic between the two sites. If this situation occurs during a period of high call volume, the shared lines might not operate as expected during that period. This situation should not last more than a few minutes, but if it is a concern, you can provision an extra 2 Mbps of bandwidth to minimize the effects.

Gateways

Normally, gateways should be provided at all sites for access to the PSTN. The device pools should be configured to register the gateways with the CallManager servers at the same site. Partitions and calling search spaces should also be configured to select the local gateways at the site as the first choice for PSTN access and the other site gateways as a second choice for overflow. Take special care to ensure emergency service access at each site.

You can centralize access to the PSTN gateways if access is not required during a WAN failure and if sufficient additional bandwidth is configured for the number of calls across the WAN. For E911 requirements, additional gateways may be needed at each site.

Voice Mail

You can deploy Cisco Unity at all sites and integrate it into the CallManager cluster. This configuration provides voice mail access even during a WAN failure and without using the PSTN. Because Cisco Unity requires a unique pilot number for voice mail, you have to configure a translation pattern at each location to translate the “virtual” messages number at each site to the correct pilot number. You should place this translation pattern in a partition that is in the calling search space for the devices at that location. If extension mobility is not being used, then users from one site who are visiting another site will have to dial the voice mail pilot number directly.

Music on Hold

MoH servers should be provisioned at each site, with sufficient capacity for the expected load. Through the use of media resource groups (MRGs) and media resource group lists (MRGLs), MoH is provided by the on-site resource and is available during a WAN failure.

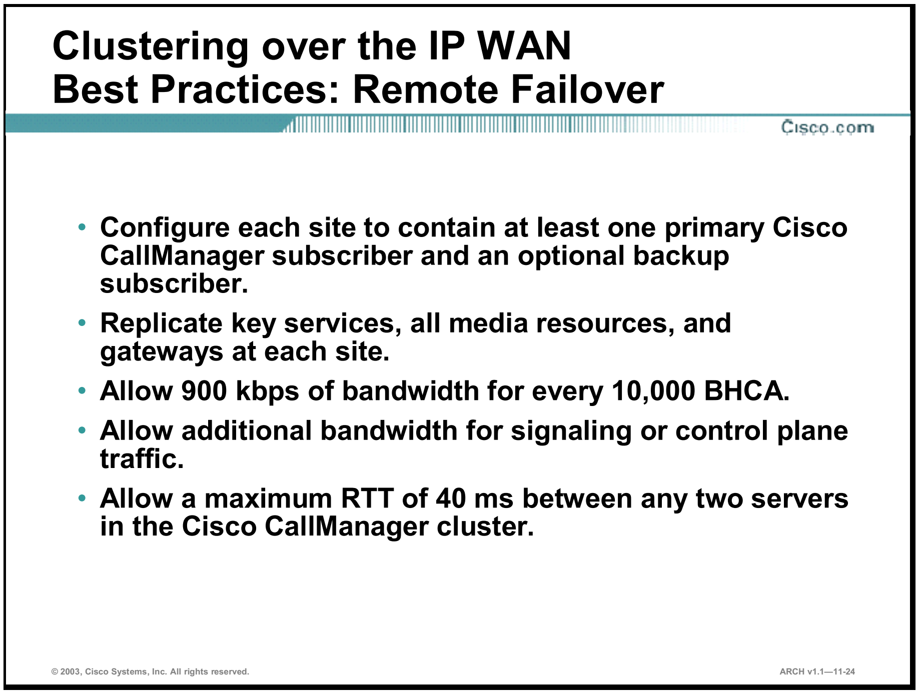

Clustering over the IP WAN Best Practices: Remote Failover

- Configure each site to contain at least one primary Cisco CallManager subscriber and an optional backup subscriber.
- Replicate key services, all media resources, and gateways at each site.
- Allow 900 kbps of bandwidth for every 10,000 BHCA.
- Allow additional bandwidth for signaling or control plane traffic.
- Allow a maximum RTT of 40 ms between any two servers in the Cisco CallManager cluster.

ARCH v1.1—11-24

The remote failover deployment model provides flexibility for the placement of backup servers. Each site contains at least one primary CallManager subscriber and may or may not have a backup subscriber.

Remote failover allows for a deployment of three to six sites with IP Phones and other devices normally registered, with a maximum of four servers. In summary, observe these guidelines when implementing the remote failover model:

- Configure each site to contain at least one primary CallManager subscriber and an optional backup subscriber if desired.
- Cisco highly recommends that you replicate key services (TFTP, DNS, DHCP, LDAP, and IP Phone services), all media resources (conference bridges and MoH), and gateways at each site with IP Phones to provide the highest level of resiliency. You could also extend this practice to include a voice mail system at each site. In the event of a WAN failure, only sites without access to the publisher database might lose a small amount of functionality.
- Every 10,000 BHCA in the cluster requires 900 kbps of bandwidth for ICCS. This is a minimum bandwidth requirement, and bandwidth is allocated in multiples of 900 kbps.
- Signaling or control plane traffic requires additional bandwidth when devices are registered across the WAN with a remote CallManager server within the same cluster.
- Allow a maximum RTT of 40 ms between any two servers in the CallManager cluster. This time equates to a 20-ms maximum one-way delay, or a transmission distance of approximately 1860 miles (3000 km) under ideal conditions.

Cisco CallManager Provisioning

If calls are allowed across a WAN between sites, then you must configure CallManager locations in addition to the default location for the other sites, to provide call admission control between the sites. If the bandwidth is over-provisioned for the number of devices, it is still best to configure call admission control based on locations. Because call admission control based on locations does not provide automatic failover to the PSTN, Cisco recommends that you over-provision the WAN for intersite calls.

As the delay increases between the CallManager servers, the bandwidth information shared between the servers for call admission control will also be delayed, possibly allowing additional calls to be set up until the ICCS bandwidth message arrives. This situation is a small risk because, even at full capacity, only a few additional calls might be set up before the bandwidth information arrives. To provide for this situation, Cisco recommends that you over-provision the priority queue for voice traffic by a few extra calls.

To improve redundancy, Cisco recommends that you enable the TFTP service on at least one of the CallManager servers at each location. You can run the TFTP service on either a publisher or a subscriber server, depending on the site. The TFTP server option must be correctly set on the DHCP servers for each site. If DHCP is not in use or the TFTP server is manually configured, you should configure the correct address for the site.

Other services, which may affect normal operation of CallManager during WAN outages, should also be replicated at all sites to ensure uninterrupted service. These services include DHCP servers, DNS servers, corporate directories, and IP Phone services. On each DHCP server, set the DNS server address correctly for each location.

IP Phones may have shared line appearances between the sites. The ICCS bandwidth provisioned between the sites allows for the additional ICCS traffic that shared line appearances generate. During a WAN outage, call control for each line appearance is segmented, but call control returns to a single CallManager server once the WAN is restored. During the WAN restoration period, there is extra traffic between the two sites. If this situation occurs during a period of high call volume, the shared lines might not operate as expected during that period. This situation should not last more than a few minutes, but if it is a concern, you can provision an extra 2 Mbps of bandwidth to minimize the effects.

Gateways

Normally, gateways should be provided at all user sites for access to the PSTN. The device pools may be configured to allow the gateways to register with a remote CallManager server as backup if the local CallManager server is unavailable. Partitions and calling search spaces should also be configured to select the local gateways at the site as the first choice for PSTN access and the other site gateways as a second choice for overflow. Take special care to ensure emergency service access at each site.

You can centralize access to the PSTN gateways if access is not required during a WAN failure and if sufficient additional bandwidth is configured for the number of calls across the WAN. For E911 requirements, additional gateways may be needed at each site.

Voice Mail

You can deploy Cisco Unity at all sites and integrate it into the CallManager cluster. This configuration provides voice mail access even during a WAN failure and without using the PSTN. Because Cisco Unity requires a unique pilot number, you have to configure a translation pattern at each location to translate the “virtual” messages number at each site to the correct pilot number. You should place this translation pattern in a partition that is in the calling search space for the devices at that location. If extension mobility is not being used, users from one site who are visiting another site will have to dial the voice mail pilot number directly.

Music on Hold

MoH servers should be provisioned at each site, with sufficient capacity for the expected load. Through the use of MRGs and MRGLs, MoH is provided by the on-site resource and is available during a WAN failure.

Network Infrastructure Design Considerations

IP telephony places strict requirements on the network infrastructure. The network must provide sufficient bandwidth and quick convergence after network failures or changes. This topic describes the physical network and network features required to support Cisco IP telephony solutions.

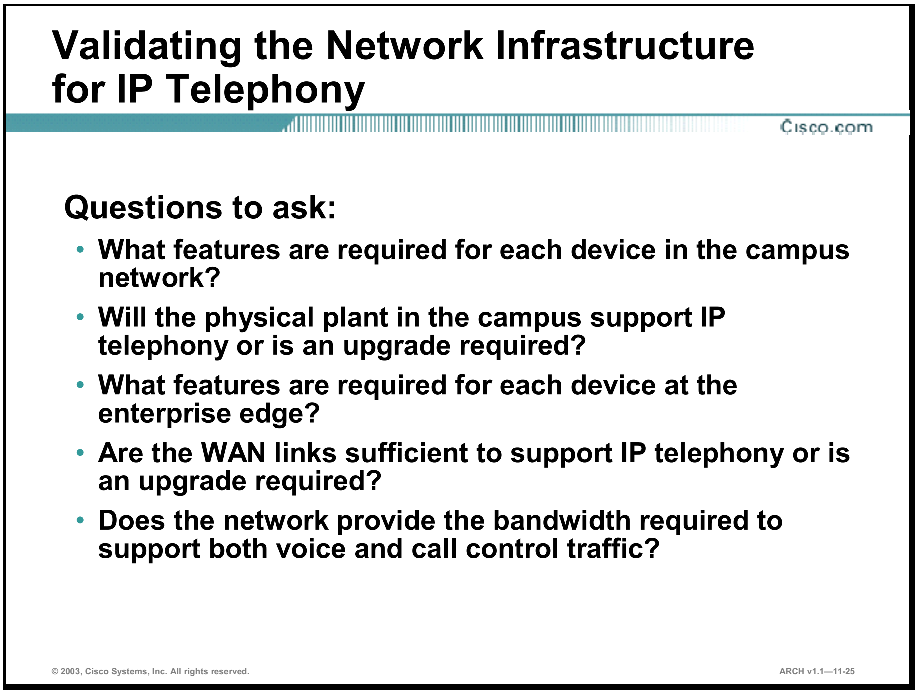

Validating the Network Infrastructure for IP Telephony

Cisco.com

Questions to ask:

- **What features are required for each device in the campus network?**
- **Will the physical plant in the campus support IP telephony or is an upgrade required?**
- **What features are required for each device at the enterprise edge?**
- **Are the WAN links sufficient to support IP telephony or is an upgrade required?**
- **Does the network provide the bandwidth required to support both voice and call control traffic?**

ARCH v1.1—11-25

Most IP telephony installations are built on an existing network infrastructure, but the infrastructure may require enhancements. In the voice solution, it is critical that you prepare to allow voice traffic to have priority over all other traffic in the network. To design the infrastructure to support voice, ask the following questions:

- **What features are required for each device in the campus network?** IP Phones require power. Most enterprises put IP telephony applications on a separate VLAN with prioritization. The infrastructure must support voice to ensure success.
- **Will the physical plant in the campus support IP telephony or is an upgrade required?** The wiring and cabling plant are critical for IP telephony. Category 5 cabling running at least 100 Mbps throughout is generally considered critical for IP telephony.
- **What features are required for each device at the enterprise edge?** At the edge, quality of service is critical. You need to know what features are implemented in the network. Without considering the edge, users may think that the voice equipment is not performing adequately when, in fact, the network is not performing adequately.
- **Are the WAN links sufficient to support IP telephony or is an upgrade required?** Evaluate each WAN link to determine if it will support voice.

- **Does the network provide the bandwidth required to support both voice and call control traffic?** Bandwidth must consider both voice traffic and call control traffic. Consider both in your traffic engineering efforts. In the campus network, bandwidth provisioning requires careful planning of the LAN infrastructure so that the available bandwidth is always considerably higher than the load and there is no steady-state congestion over the LAN links. This ensures that the network is responsive to the offered traffic.

Note You should plan to work with a voice specialist to complete a traffic engineering analysis for the network.

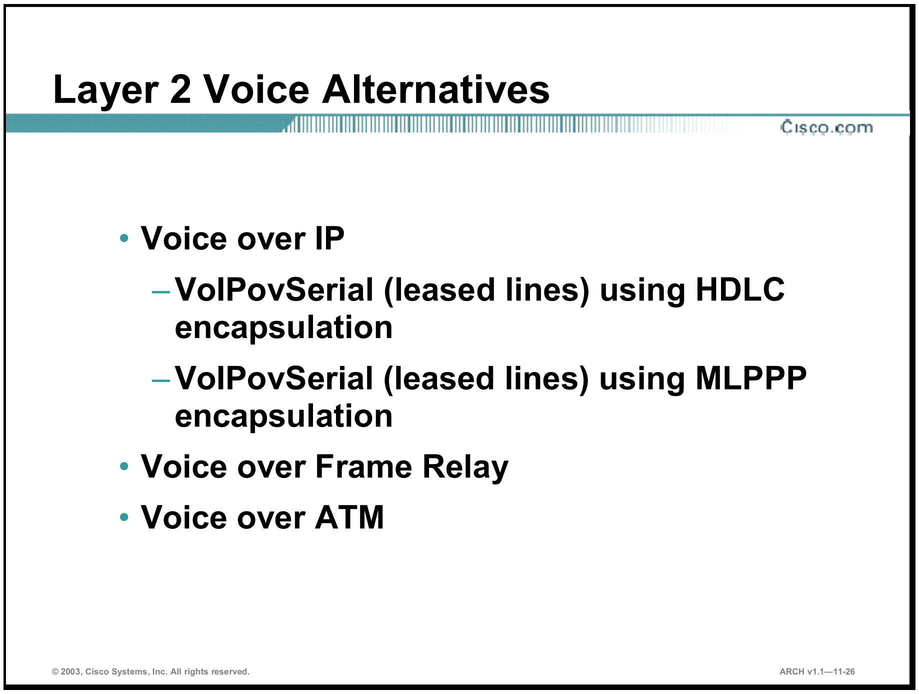

Layer 2 Voice Alternatives

- Voice over IP
  - VoIPovSerial (leased lines) using HDLC encapsulation
  - VoIPovSerial (leased lines) using MLPPP encapsulation
- Voice over Frame Relay
- Voice over ATM

ARCH v1.1—11-26

As an alternative to a CallManager VoIP-based voice network, you can use the data link layer as the transport. The most common use of a Layer 2 voice transport is to replace tie lines or create circuits or trunks from one voice-enabled network device to another voice-enabled network device, providing a point-to-point circuit.

Many enterprises have an existing Frame Relay or ATM network. Voice over Frame Relay or Voice over ATM is an option for those enterprises. For enterprises already using constant-bit-rate ATM, there is a natural match with voice capabilities. Their savings for packet voice coding is dependent on the networks support for subrate multiplexing.

For enterprises that have an existing Frame Relay network or the variable bit rate forms of ATM, transport of compressed voice can represent significant savings while providing additional flexibility of features.

Voice over IP over Leased Lines

VoIP over leased lines provides a method for enterprises that currently use leased lines (VoIPovSerial) with either PPP or High-Level Data Link Control (HDLC) encapsulation to integrate voice into their current data leased lines.

Voice over Frame Relay

Voice-enabled routers can integrate voice, LAN, synchronous data, video, and fax traffic for transport over a public or private Frame Relay network. Cisco optimizes network bandwidth over network links by multiplexing voice and data on the same circuit or physical interface. These features are offered with Voice over Frame Relay:

- Enables real-time, delay-sensitive voice traffic to be carried over slow Frame Relay links
- Allows replacement of dedicated 64-kbps TDM telephony circuits with more economical Frame Relay PVCs or SVCs
- Uses voice compression technology that conforms to International Telecommunication Union Telecommunication Standardization Sector (ITU-T) specifications
- Allows intelligent setup of proprietary-switched Voice over Frame Relay connections between two Voice over Frame Relay endpoints
- Supports standards-based FRF.11 and FRF.12 functionality

Voice over ATM

ATM is a switching method for transmitting information in fixed-length cells, based on application (voice, video, and data) demand and priority. ATM supports a number of service classes. Each ATM service class is given a guaranteed minimum bandwidth, which ensures deterministic behavior under load, and supports complex classes of service that voice and video can take advantage of to guarantee they receive a higher priority through the network. Voice over ATM offers these features:

- Uses small, fixed-sized cells (53 bytes)
- Is a connection-oriented protocol
- Supports multiple service types
- Supports LAN and WAN traffic
- Emulates PSTN circuits using ATM virtual circuits
- Minimizes delay and delay variation

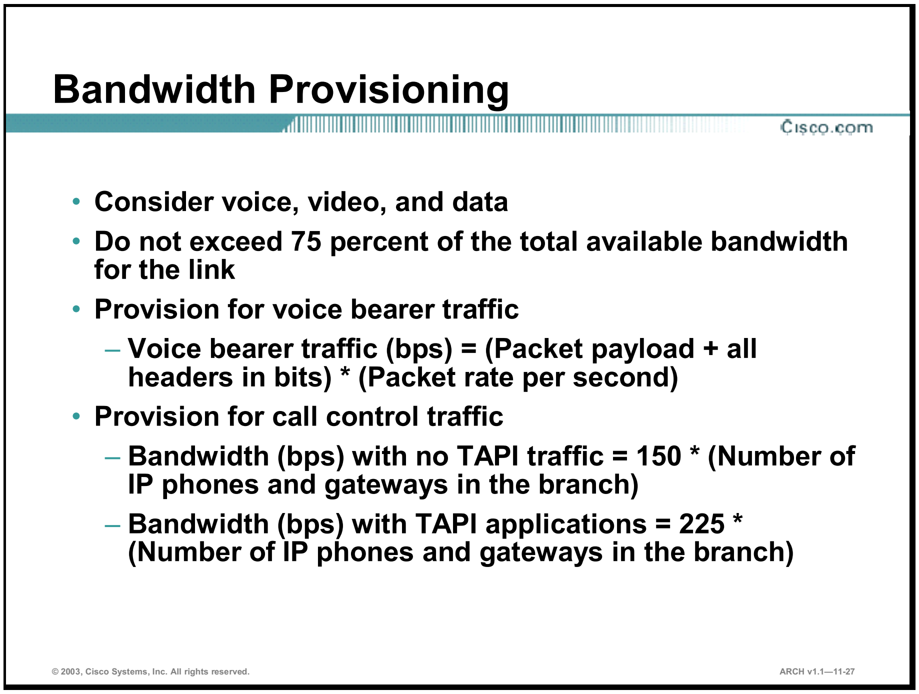

Bandwidth Provisioning

- Consider voice, video, and data
- Do not exceed 75 percent of the total available bandwidth for the link
- Provision for voice bearer traffic
  - Voice bearer traffic (bps) = (Packet payload + all headers in bits) * (Packet rate per second)
- Provision for call control traffic
  - Bandwidth (bps) with no TAPI traffic = 150 * (Number of IP phones and gateways in the branch)
  - Bandwidth (bps) with TAPI applications = 225 * (Number of IP phones and gateways in the branch)

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—11-27

Properly provisioning the network bandwidth is a major component of designing a successful IP telephony network. You can calculate the required bandwidth by adding the bandwidth requirements for each major application, including voice, video, and data. This sum then represents the minimum bandwidth requirement for any given link, and it should not exceed approximately 75 percent of the total available bandwidth for the link.

From a traffic standpoint, an IP telephony call consists of two parts:

- The voice carrier stream, which consists of RTP packets that contain the actual voice samples.
- The call control signaling, which consists of packets belonging to one of several protocols, according to the endpoints involved in the call, for example, H.323 or MGCP. Call control functions are, for instance, those used to set up, maintain, tear down, or redirect a call.

Bandwidth provisioning must include not only the voice stream traffic but also the call control traffic.

Provisioning for Voice Bearer Traffic

A VoIP packet consists of the payload, IP header, UDP header, RTP header, and Layer 2 link header. At a packetization period of 20 ms, voice packets have a 160-byte payload for G.711 or a 20-byte payload for G.729. However, G.729 does not support fax or modem traffic. The IP header is 20 bytes, the UDP header is 8 bytes, and the RTP header is 12 bytes. The link header varies in size according to the Layer 2 media used.

Use the following formula to calculate the bandwidth that voice streams consume:

- (Packet payload + all headers in bits) * Packet rate per second (50 packets per second when using a 20 ms packet period)

Provisioning for Call Control Traffic

When provisioning bandwidth for call control traffic using centralized call control from remote sites, consider:

- When a remote branch phone places a call, the control traffic traverses the IP WAN (even if the call is local to the branch) to reach the CallManager at the central site.
- Signaling Connection Control Part (SCCP) and TAPI are the most common signaling protocols on the IP WAN. Other deployment patterns might use H.323, MGCP, or SIP. All the control traffic is exchanged between a CallManager and the central site, and endpoints or gateways at the remote branches.

As a consequence, the area in which to provision bandwidth for control traffic lies between the branch routers and the WAN aggregation router at the central site. The control traffic that traverses the WAN includes two categories:

- **Maintenance traffic:** Includes keepalive messages periodically exchanged between the branch IP Phones and CallManager, regardless of phone activity
- **Call-related traffic:** Includes signaling messages exchanged between the branch IP Phones and/or gateways and the CallManager at the central site when a call needs to be set up, torn down, forwarded, and so on

To estimate the generated call control traffic, you must determine the average number of calls per hour made by each branch IP Phone.

You can determine the recommended bandwidth needed for call control traffic with this formula:

- Voice bearer traffic (bps) = (Packet payload + all headers in bits) * (Packet rate per second)

If a TAPI application is deployed at the remote branches, the recommended bandwidth is affected because the TAPI protocol requires more messages to be exchanged between CallManager and the endpoints.

The following formula takes into account the impact of a TAPI application:

- Bandwidth (bps) with TAPI applications = 225 * (Number of IP Phones and gateways in the branch)

If you consider the fact that the smallest bandwidth that you can assign to a queue on an IOS router is 8 kbps, you can summarize the values of minimum and recommended bandwidth for different branch office sizes.

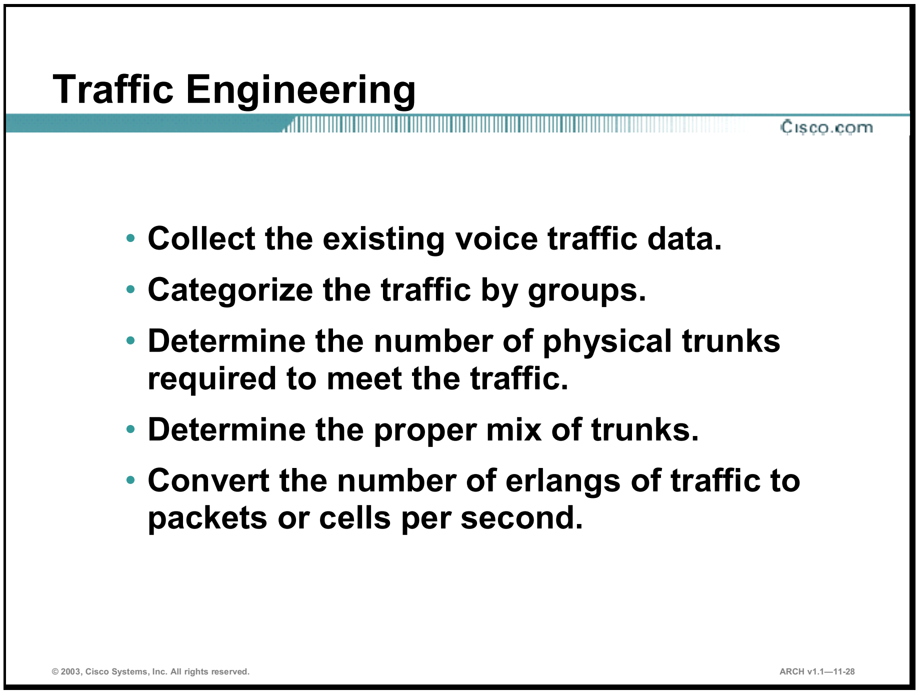

Traffic Engineering

- Collect the existing voice traffic data.
- Categorize the traffic by groups.
- Determine the number of physical trunks required to meet the traffic.
- Determine the proper mix of trunks.
- Convert the number of erlangs of traffic to packets or cells per second.

ARCH v1.1—11-28

Traffic engineering, as it applies to traditional voice networks, is determining the number of trunks necessary to carry a required amount of voice calls during a period of time. For designers of a voice network, the goal is to properly size the number of trunks and provision the appropriate amount of bandwidth necessary to carry the amount of traffic determined.

There are two different types of connections to consider: lines and trunks. Lines allow telephone sets to be connected to telephone switches, like PBXs and central office (CO) switches. Trunks connect switches together. An example of a trunk is a tie line interconnecting PBXs (ignore the use of “line” in the tie line statement; it’s actually a trunk).

Companies use switches to act as concentrators because the number of telephone sets required is usually greater than the number of simultaneous calls that need to be made. For example, a company may have 600 telephone sets connected to a PBX, but may only have 15 trunks connecting the PBX to the central office switch.

To perform traffic engineering on a voice network, follow the process described in the table.

Step	Description	Notes and Comments
1.	Collect the existing voice traffic.	From the carrier, gather the following information: ■ Peg counts for calls offered, calls abandoned, and all trunks busy ■ Grade of service rating for trunk groups ■ Total traffic carried per trunk group ■ Phone bills, to see the carrier's rates The terms used above will be covered in more detail in the next few topics. For best results, get two weeks' worth of traffic. The internal telecommunications department can provide CDRs for PBXs. This information typically records calls that are offered, but does not provide information on calls that were blocked because all trunks were busy.
2.	Categorize the traffic by groups.	In most large enterprises, it is more cost effective to apply traffic engineering to groups of trunks serving a common purpose. For example, you might separate inbound customer service calls into a separate trunk group distinctly different from general outgoing calls. Start by separating the traffic into inbound and outbound directions. Group outbound traffic into distances called (for example, local, local long distance, intrastate, interstate, and so on). It is important to break the traffic by distance because most tariffs are distance sensitive. Determine the purpose of the calls. For example, what were the calls for? Were they used for fax, modem, call center, 800 for customer service, 800 for voice mail, telecommuters, and so on?
3.	Determine the number of physical trunks required to meet the traffic needs.	If you know the amount of traffic generated and the GoS required, then you can calculate the number of trunks required to meet your needs. Use the following simple equation to calculate traffic flow: $A = C * T$ A is the traffic flow, C is the number of calls originating during a period of 1 hour, and T is the average holding time of a call. It is important to note that C is the number of calls originated, not carried. Typically, the information received from the carrier or from the company's internal CDRs are in terms of carried traffic and not offered traffic, as is usually provided by PBXs.
4.	Determine the proper mix of trunks.	The proper mix of trunks is more of an economic decision than a technical decision. Cost per minute is the most commonly used measurement for determining the price breakpoint of adding trunks. Care must be taken to ensure that all cost components are considered, such as accounting for additional transmission, equipment, administration, and maintenance costs. There are two rules to follow when optimizing the network for cost: ■ Use average usage figures instead of the busy hour, which would overstate the number of call minutes. ■ Use the least costly circuit until the incremental cost becomes more expensive than the next best route.

Step	Description	Notes and Comments
5.	Convert the number of erlangs of traffic to packets or cells per second.	The final step in traffic engineering is to equate erlangs of carried traffic to packets or cells per second. One way to do this is to convert one erlang to the appropriate data measurement, then apply modifiers. The following equations are theoretical numbers based on PCM voice and fully loaded packets. 1 PCM voice channel requires 64 kBps 1 erlang is 60 minutes of voice Therefore: 1 erlang = 64 kBps * 3600 seconds * 1 byte/8 bits = 28.8 MB of traffic in one hour

Note Use an erlang calculator and an experienced voice network designer to perform traffic engineering for a voice network. Refer to http://www.erlang.com/calculator/ for more information and an erlang calculator.

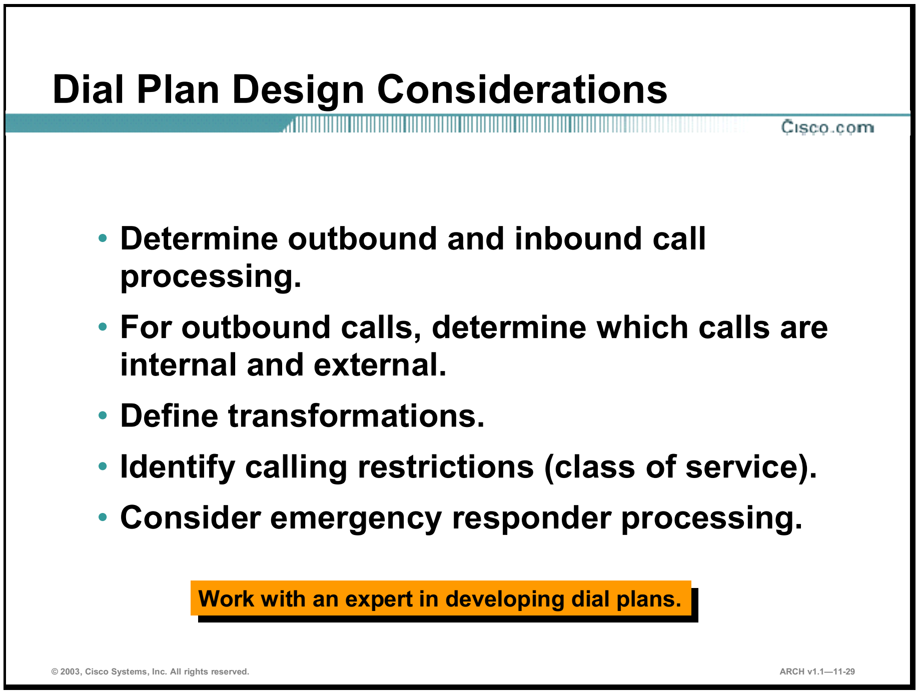

Dial Plan Design Considerations

Cisco.com

- Determine outbound and inbound call processing.
- For outbound calls, determine which calls are internal and external.
- Define transformations.
- Identify calling restrictions (class of service).
- Consider emergency responder processing.

Work with an expert in developing dial plans.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-29

A well-designed dial plan is a vital component of any IP telephony network, and all the other network elements rely on it. The dial plan is essentially the endpoint selection for IP voice calls.

CallManager distinguishes between internal and external calls. The design questions to ask are:

- How many calls are internal calls?
- How many calls are external calls?

CallManager provides digit transformation (translation), which is the ability to transform a called or calling number into another number. Digit translation can be used on internal as well as external calls, whether inbound or outbound.

CallManager provides the ability to restrict calls on each phone individually or on groups of phones in the same CallManager cluster. Users can be grouped into communities of interest on the same CallManager, yet share the same gateways and have overlapping dial plans. These capabilities help support multisite IP WAN deployments with centralized call processing and multi-tenant deployments.

Access to emergency services (911 and E911) is required by law in most areas. Because emergency services have requirements that can affect the overall design of your network, you should consider them early in the design phase. Gateways for 911 and E911 services must be highly available, and you can distribute them in many locations to meet this requirement. You can also install redundant gateways at each location to connect to the PSTN and provide routing of 911 calls.

Note Creating dial plans is difficult and involved. It requires an expert in dial planning to develop an effective solution. Consult with an expert before attempting to create an enterprise dial plan.

Intelligent Network Services for IP Telephony and Voice

Intelligent network services such as network management, high availability, security, and QoS extend to incorporate voice-specific solutions. This topic explains the role that intelligent network services play in a Cisco IP telephony solution.

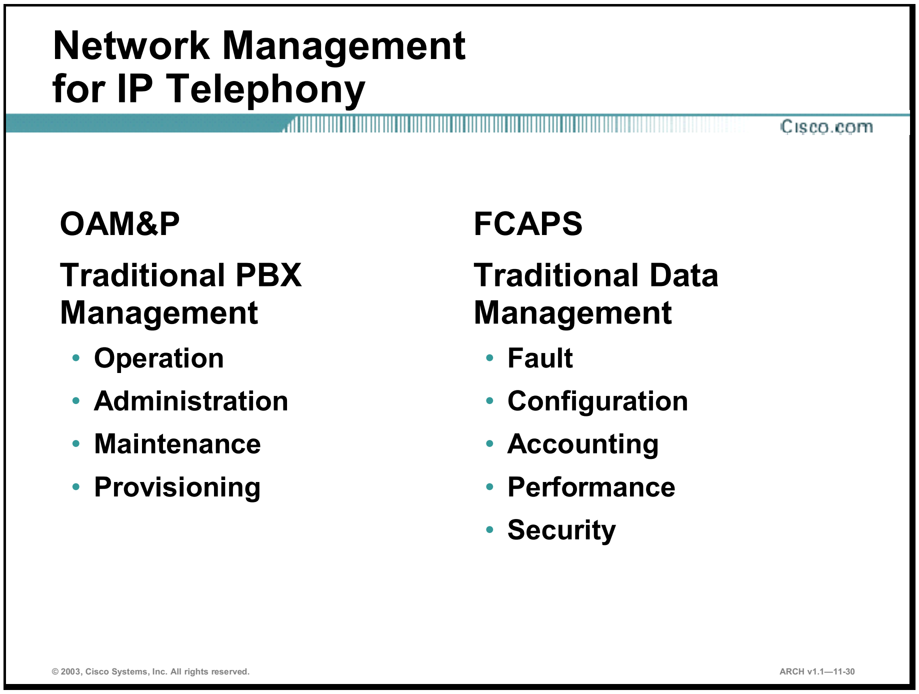

Network Management for IP Telephony

Cisco.com

OAM&P

Traditional PBX Management

- Operation
- Administration
- Maintenance
- Provisioning

FCAPS

Traditional Data Management

- Fault
- Configuration
- Accounting
- Performance
- Security

ARCH v1.1—11-30

In traditional voice networks, there is a distinct set of voice management concepts and processes. The convergence of voice and data has brought about a similar merge of data network and voice-only management.

In fact, this merging of management tasks and processes is one of the key benefits of using a converged network as opposed to a dedicated voice-only network. However, it is still necessary to comprehend the traditional voice-only management concepts in order to relate the features available in that technology to the converged network management techniques.

Network management of circuit-switched data networks is often summarized by the Fault, Configuration, Accounting, Performance, Security (FCAPS) model, defined in the functional model of the Open System Interconnection (OSI) management architecture and adopted by ITU-T as part of the Telecommunication Management Network (TMN) initiative. In a traditional voice environment, the approach to management is usually referred to as operations, administration, management, and provisioning (OAM&P).

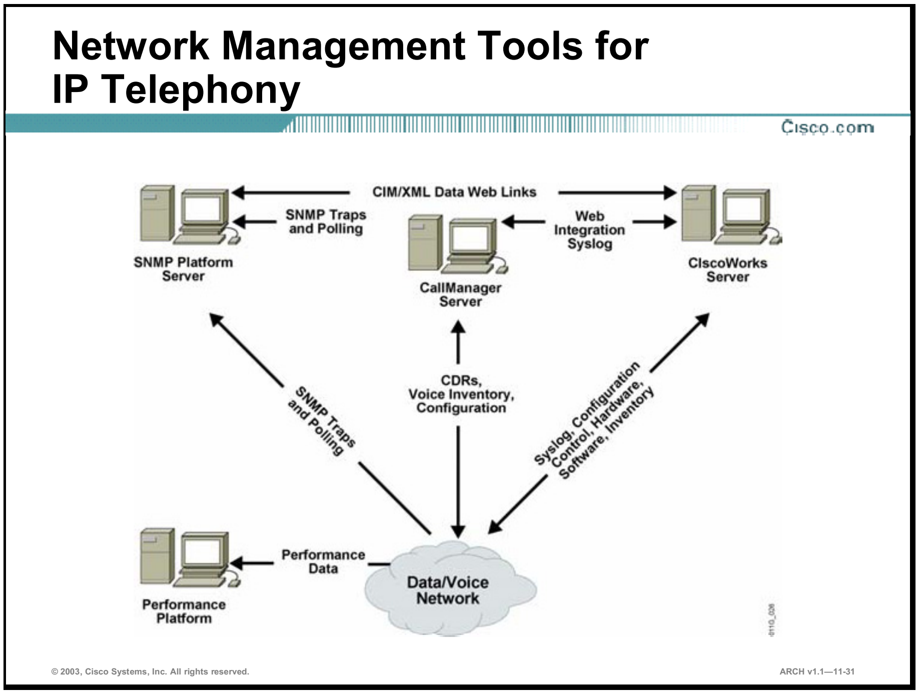

Network Management Tools for IP Telephony

ARCH v1.1—11-31

The CiscoWorks family of tools, CallManager, and various third-party tools provide management functions for IP telephony networks. This topic proposes network management services required to support Cisco IP telephony solutions.

In addition to managing the routers and switches of the Cisco AVVID infrastructure, the CiscoWorks tools communicate with devices providing voice services, such as CallManager, Cisco Conference Connection, and Cisco Emergency Responder, which themselves may provide other voice management services.

CiscoWorks VoIP Health Monitor (VHM) proactively monitors Cisco voice servers and polls for reachability, interface status, environmental conditions (power supply, fan, and temperature), and application status. In addition to monitoring the voice server, VHM verifies the availability of key voice services by performing synthetic transactions, wherein the VHM server emulates the behavior of a Cisco IP Phone and performs a specific transaction. The value of synthetic transactions is that, by actually accessing these critical voice services, it is possible to verify that the services are available.

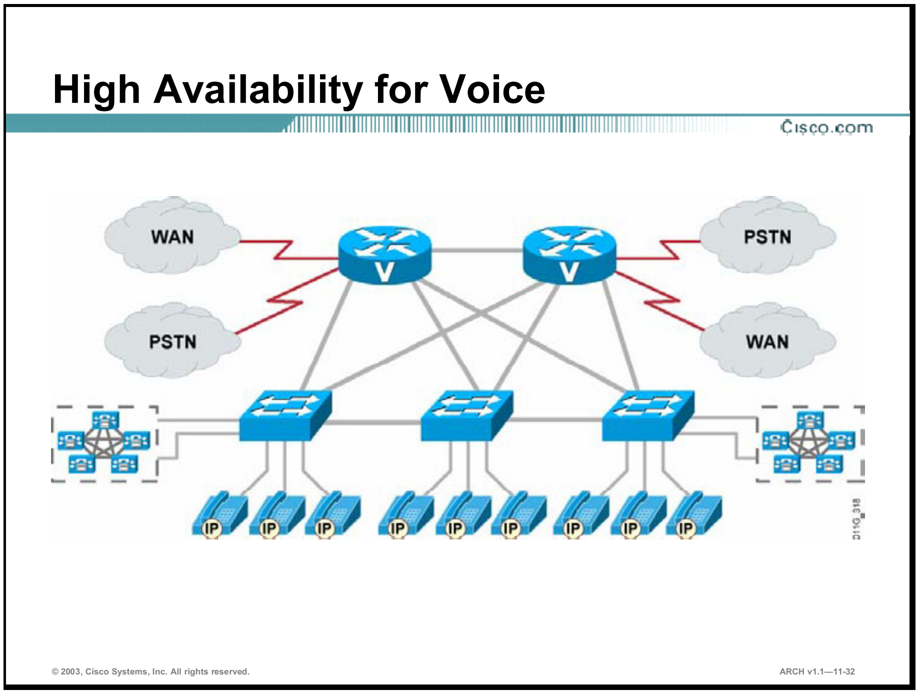

High Availability for Voice

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—11-32

Cisco AVVID IP telephony is based on a distributed model for high availability. CallManager clusters support CallManager redundancy. The gateways must support the ability to "re-home" to a secondary CallManager in the event that a primary CallManager fails, providing CallManager redundancy. This differs from call survivability in the event of a CallManager or network failure, where the call is routed to an alternate gateway, such as an MGCP gateway.

As with any capability within the network, you need to plan redundancy for critical components such as the CallManager and the associated gateway and infrastructure devices that support the voice network.

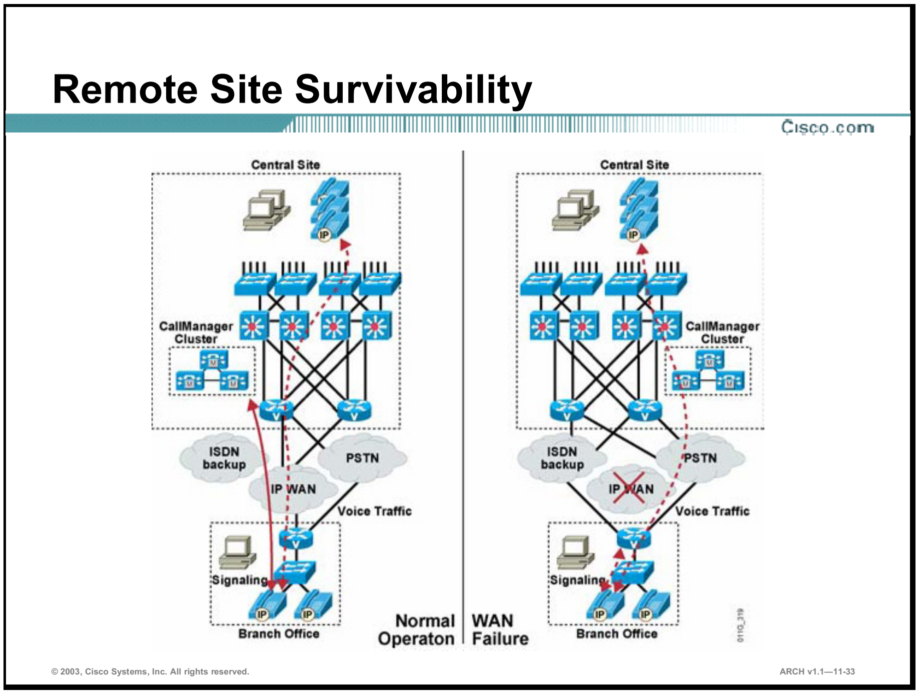

Remote Site Survivability

Cisco.com

ARCH v1.1—11-33

When deploying IP telephony across a WAN with the centralized call processing model, you should take additional steps to ensure that data and voice services at the remote sites are highly available. SRST provides high availability for voice services only, by providing a subset of the call processing capabilities within the remote office router and enhancing the IP Phones with the ability to "re-home" to the call processing functions in the local router if a WAN failure is detected.

If the WAN link to the branch office fails, or if some other event causes loss of connectivity to the CallManager cluster, the branch IP Phones reregister with the branch router. The branch router queries the IP Phones for their configuration and uses this information to build its own configuration automatically. The branch IP Phones can then make and receive calls either internally or through the PSTN. The phone displays the message "CM fallback mode," and some advanced CallManager features are unavailable and are grayed out on the phone display.

When WAN connectivity to the central site is reestablished, the branch IP Phones automatically reregister with the CallManager cluster and resume normal operation. The branch router deletes its information about the IP Phones and reverts to its standard routing or gateway configuration.

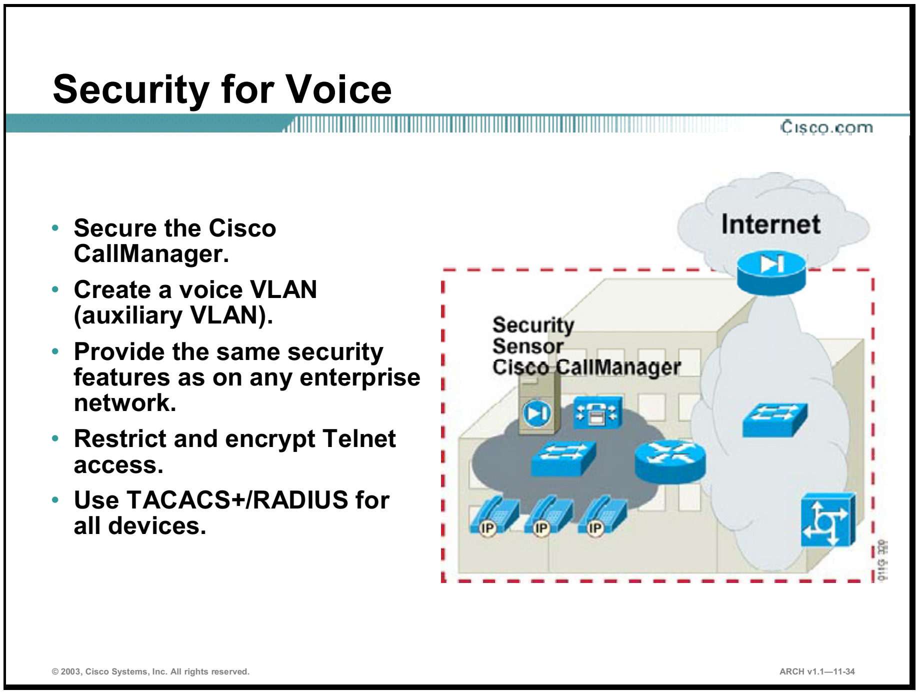

Security for Voice

Cisco.com

- Secure the Cisco CallManager.
- Create a voice VLAN (auxiliary VLAN).
- Provide the same security features as on any enterprise network.
- Restrict and encrypt Telnet access.
- Use TACACS+/RADIUS for all devices.

ARCH v1.1—11-34

The subject of securing voice communications, always a very sensitive topic for today’s communications architects, has received even more visibility recently as network convergence becomes the accepted design model. With the advent of IP telephony, which uses IP data network devices for voice communication, the potential exists for malicious attacks on call processing components and telephony applications.

To help safeguard against attacks, you should implement the same security precautions as in the rest of the enterprise network.

Securing the voice call processing platform and installed applications is perhaps the most vital step in securing Cisco AVVID networks.

Every enterprise should have a predefined security policy for all devices, applications, and users to follow. The strictness of the security policy depends upon the level of caution required. This document does not describe how to establish a security policy for your enterprise, but it does present a set of security guidelines and recommendations to help you configure the IP telephony network to conform to your enterprise security policy.

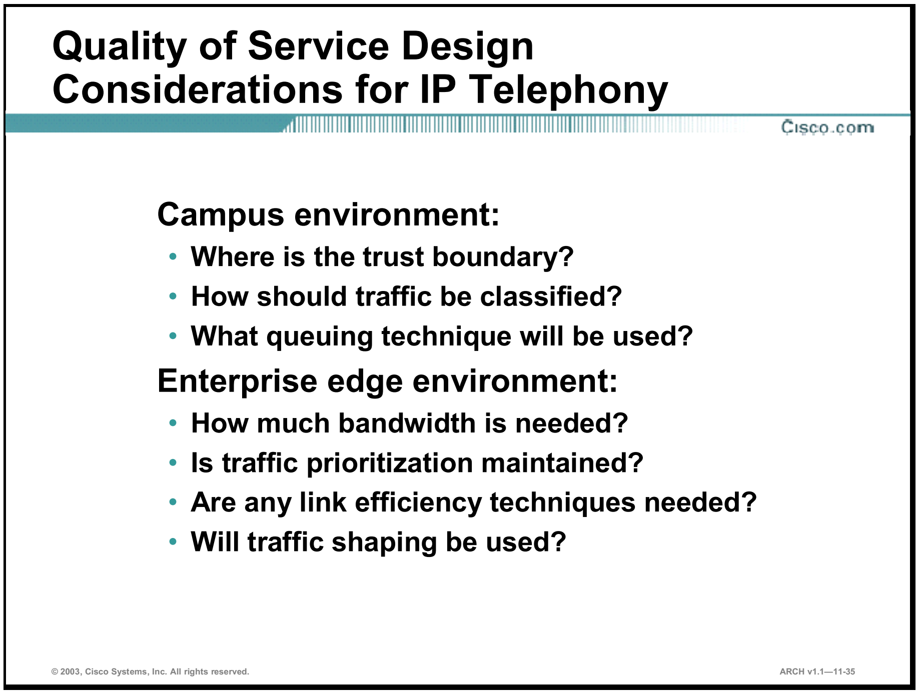

Quality of Service Design Considerations for IP Telephony

Campus environment:

- **Where is the trust boundary?**
- **How should traffic be classified?**
- **What queuing technique will be used?**

Enterprise edge environment:

- **How much bandwidth is needed?**
- **Is traffic prioritization maintained?**
- **Are any link efficiency techniques needed?**
- **Will traffic shaping be used?**

ARCH v1.1—11-35

Voice, as a class of IP network traffic, has strict requirements concerning delay and delay variation (also known as jitter). Compared to most data, it is relatively tolerant of loss. To meet the requirements for voice traffic, the Cisco AVVID IP telephony solution uses a wide range of IP QoS features such as classification, queuing, congestion detection, traffic shaping, and compression.

The overall goal of QoS in the network is to be able to manage which applications are less likely to be affected by loss, delay, and jitter. When a network becomes congested, some traffic will be delayed or even, at times, lost. The goal is to give critical applications a higher priority for service so that they are least likely to be delayed or dropped in times of congestion. In many converged networks voice is the most critical application. In others, voice may be used to opportunistically use bandwidth not required for data, and fall back to the PSTN in times of congestion.

When you configure campus QoS, you must consider:

- **Trust boundary:** Define and enforce a trust boundary at the network edge. Classification should take place at the network edge, typically in the wiring closet or within the endpoints themselves.
- **Traffic classification:** Network design practice emphasizes that you should classify or mark traffic as close to the edge of the network as possible. Traffic class is a criterion for queue selection in the various queuing schemes used at interfaces within the campus switches and WAN devices. When you connect an IP Phone using a single cable, the phone becomes the edge of the managed network. As the edge device, the IP Phone can and should classify and mark traffic flows so that network QoS mechanisms can correctly handle the packets.
- **Interface queuing**: To guarantee voice quality, you must design for QoS and enable it within the campus infrastructure. By enabling QoS on campus switches, you can configure voice traffic to use separate queues, virtually eliminating the possibility of dropped voice packets as an interface buffer fills, and minimizing delay.

Although network management tools may show that the campus network is not congested, QoS tools are still required to guarantee voice quality. Network management tools typically show only the average congestion over a relatively long sample time, minutes to tens of minutes. While useful, this average does not show the congestion peaks on a campus interface, which can occur in a period of seconds to tens of seconds.

Transmit interface buffers within a campus tend to congest in small, finite intervals as a result of the bursty nature of network data traffic. When severe congestion occurs, any packets destined for that transmit interface are delayed or even dropped, thereby affecting voice quality. The way to prevent dropped voice traffic is to configure multiple queues on campus switches. Most Cisco Ethernet switches support the enhanced queuing services using multiple queues that can guarantee voice quality in the campus.

WAN QoS techniques depend on the speeds of the links. At speeds above 768 kbps, voice priority queuing is used to reduce jitter and possible packet loss if a burst of traffic oversubscribes a buffer. This queuing requirement is similar to the one for the LAN infrastructure.

In addition, the WAN requires traffic shaping for two reasons:

- To remain within the contracted traffic agreement with the ATM or Frame Relay network to avoid being policed and incurring dropped packets.
- To maintain comparable traffic speeds between sites linked to the Frame Relay or ATM network by different line speeds. For example, the headquarters site might use DS-3 and the other sites use DS-1, which can result in buffer overruns within the network and, thus, in packet loss. Traffic shaping helps prevent buffer overruns and packet loss.

In choosing from among the many available prioritization schemes, the major factors to consider include the type of traffic involved and the type of media on the WAN. For multiservice traffic over an IP WAN, Cisco recommends low latency queuing (LLQ) for low-speed links. LLQ allows up to 64 traffic classes with the ability to specify, for example, strict priority queuing behavior for voice and interactive video, a minimum bandwidth for Systems Network Architecture (SNA) data and market data feeds, and weighted fair queuing (WFQ) for other traffic types.

Because wide-area bandwidth is often expensive, only low-speed circuits might be available or affordable when interconnecting remote sites. In these cases, it is important to achieve maximum bandwidth efficiencies by transmitting as many voice calls as possible over the low-speed link. Compression schemes, such as G.729, can compress a 64-kbps call into an 8-kbps payload. Cisco gateways and IP Phones support a range of codecs that can improve bandwidth efficiency on these low-speed links in exchange for slightly increased delay.

Note You should not use G.729 codecs with conferencing applications. With G.729 codecs, the quality of any conversation between multiple parties is unacceptable when more than one person talks at one time.

You can increase link efficiency further by using compressed Real-Time Transport Protocol (CRTP) on selected links. This protocol compresses a 40-byte IP, UDP, and RTP header to approximately two to four bytes across an appropriately configured link.

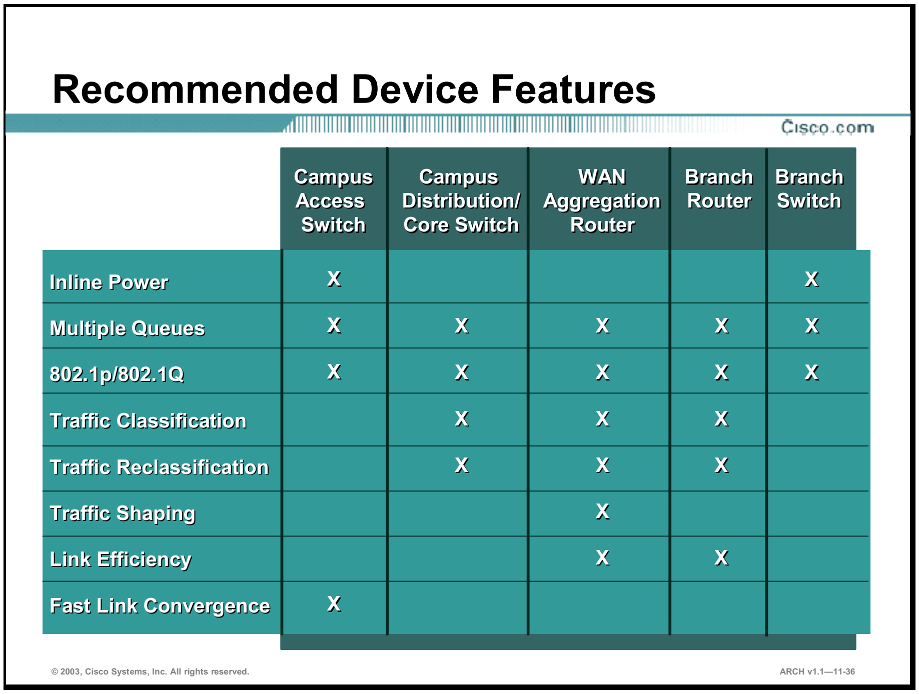

Recommended Device Features

	Campus Access Switch	Campus Distribution/ Core Switch	WAN Aggregation Router	Branch Router	Branch Switch
Inline Power	X				X
Multiple Queues	X	X	X	X	X
802.1p/802.1Q	X	X	X	X	X
Traffic Classification		X	X	X	
Traffic Reclassification		X	X	X	
Traffic Shaping			X		
Link Efficiency			X	X	
Fast Link Convergence	X				

ARCH v1.1—11-36

IP telephony places strict requirements on IP packet loss, packet delay, and delay variation (or jitter). Therefore, you need to enable most of the QoS mechanisms available on Cisco switches and routers throughout the network.

The figure indicates the requirements for each device that forms the network infrastructure. The features are key components of the infrastructure. If the infrastructure is not planned properly, adding voice to an improperly planned network will have a disastrous effect on the voice quality. In a small to maybe medium-sized network, the features listed in the figure are not as critical as they are for a large enterprise. You should implement each feature to provide a reliable quality voice and integrated data network.

Summary

This topic summarizes the key points discussed in this lesson.

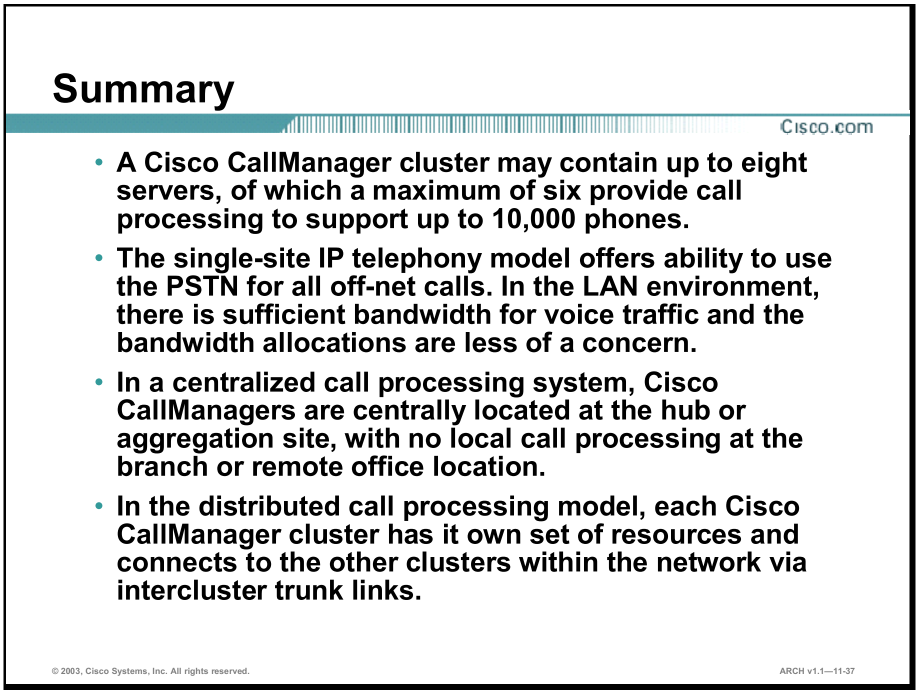

Summary

Cisco.com

- A Cisco CallManager cluster may contain up to eight servers, of which a maximum of six provide call processing to support up to 10,000 phones.
- The single-site IP telephony model offers ability to use the PSTN for all off-net calls. In the LAN environment, there is sufficient bandwidth for voice traffic and the bandwidth allocations are less of a concern.
- In a centralized call processing system, Cisco CallManagers are centrally located at the hub or aggregation site, with no local call processing at the branch or remote office location.
- In the distributed call processing model, each Cisco CallManager cluster has it own set of resources and connects to the other clusters within the network via intercluster trunk links.

ARCH v1.1—11-37

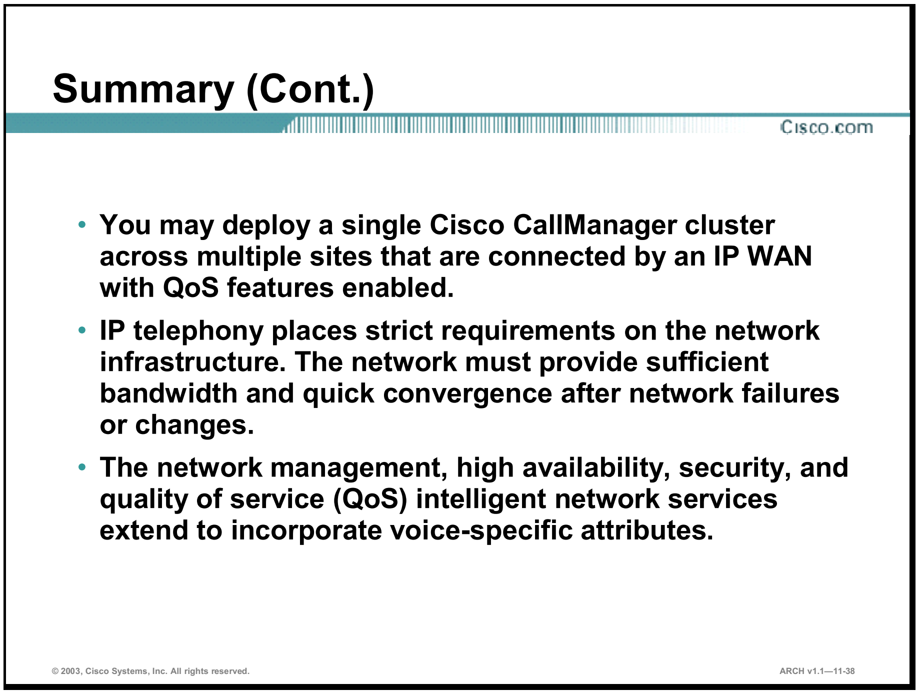

Summary (Cont.)

Cisco.com

- You may deploy a single Cisco CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled.
- IP telephony places strict requirements on the network infrastructure. The network must provide sufficient bandwidth and quick convergence after network failures or changes.
- The network management, high availability, security, and quality of service (QoS) intelligent network services extend to incorporate voice-specific attributes.

ARCH v1.1—11-38

References

For additional information, refer to this resource:

- *Cisco CallManager Clusters* at http://www.cisco.com/univercd/cc/td/doc/product/voice/ip_tele/network/dgclustr.htm

Next Steps

For the associated case study and simulation, refer to the following sections:

- Case Study 11-2: OCSIC Bottling Company
- OPNET IT Guru Simulation 11-2

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 CallManager systems does Cisco recommend to support up to 5000 devices?

A) 2

B) 4

C) 8

D) 10

Q2) When the busy hour call attempts rise, what happens to the CallManager device weights?

A) Nothing.

B) The device weights increase.

C) The device weights decrease in relationship to the BHCA.

D) The devices decrease on a factor of 2 per BHCA increased.

Q3) What is a benefit of single-site deployments?

A) ease of deployment

B) remote phone access

C) distributed processing

D) handling of 20,000 phones

Q4) In a single-site deployment cluster, how many phones can be in the cluster?

A) 2500

B) 5000

C) 7500

D) 10,000

Q5) How many phones does a centralized CallManager support?

A) 2500

B) 5000

C) 7500

D) 10,000

Q6) What is used to determine when a CallManager is supporting its maximum number of devices?

A) device load

B) device weights

C) device statements

D) device performance variable

Q7) Which type of CallManager call admission control is used in a centralized environment?

A) region-based

B) locations-based

C) gatekeeper-based

D) device pool-based

Q8) What is a benefit of a distributed CallManager deployment?

A) scalability

B) common dial plan

C) no load-sharing with PSTN

D) loss of functionality when WAN goes down

Q9) What device does a distributed CallManager deployment use to route calls outside its cluster?

A) phones

B) gateways

C) voice mail

D) voice applications

Q10) What is a benefit of the cluster over the WAN deployment model?

A) separate dial plans

B) SRST at the remote site

C) no security needed at remote site

D) extension mobility within the cluster

Q11) With remote failover for cluster over the WAN, how many CallManagers are needed in the remote site?

A) 1

B) 2

C) 3

D) 4

Q12) What are the two components of an IP telephony call? (Choose two.)

A) data stream

B) UDP header

C) TAPI applications

D) voice carrier stream

E) call control signaling

Q13) Which VoX implementation solution provides for deterministic behavior?

A) ATM

B) Frame Relay

C) VoIP over PPP

D) VoIP over HDLC

Q14) Which mechanism can you use to segregate voice and data traffic for security purposes?

A) VCR

B) VPN

C) VLAN

D) multicast

Q15) As with any service within the network, for components such as the CallManager and the associated devices that support the voice network, you need to plan _____.

A) resiliency

B) redundancy

C) on-site support

D) multicast capability

Quiz Answer Key

Q1) B

Relates to: Cisco CallManager Cluster Design Considerations

Q2) B

Relates to: Cisco CallManager Cluster Design Considerations

Q3) A

Relates to: Designing Single-Site IP Telephony Solutions

Q4) D

Relates to: Designing Single-Site IP Telephony Solutions

Q5) D

Relates to: Designing Multisite with Centralized Call Processing IP Telephony Solutions

Q6) B

Relates to: Designing Multisite with Centralized Call Processing IP Telephony Solutions

Q7) B

Relates to: Designing Multisite with Centralized Call Processing IP Telephony Solutions

Q8) A

Relates to: Designing Multisite with Distributed Call Processing IP Telephony Solutions

Q9) B

Relates to: Designing Multisite with Distributed Call Processing IP Telephony Solutions

Q10) D

Relates to: Clustering Over the IP WAN

Q11) A

Relates to: Clustering Over the IP WAN

Q12) D, E

Relates to: Network Infrastructure Design Considerations

Q13) A

Relates to: Network Infrastructure Design Considerations

Q14) C

Relates to: Intelligent Network Services for IP Telephony and Voice

Q15) B

Relates to: Intelligent Network Services for IP Telephony and Voice

Case Study 11-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

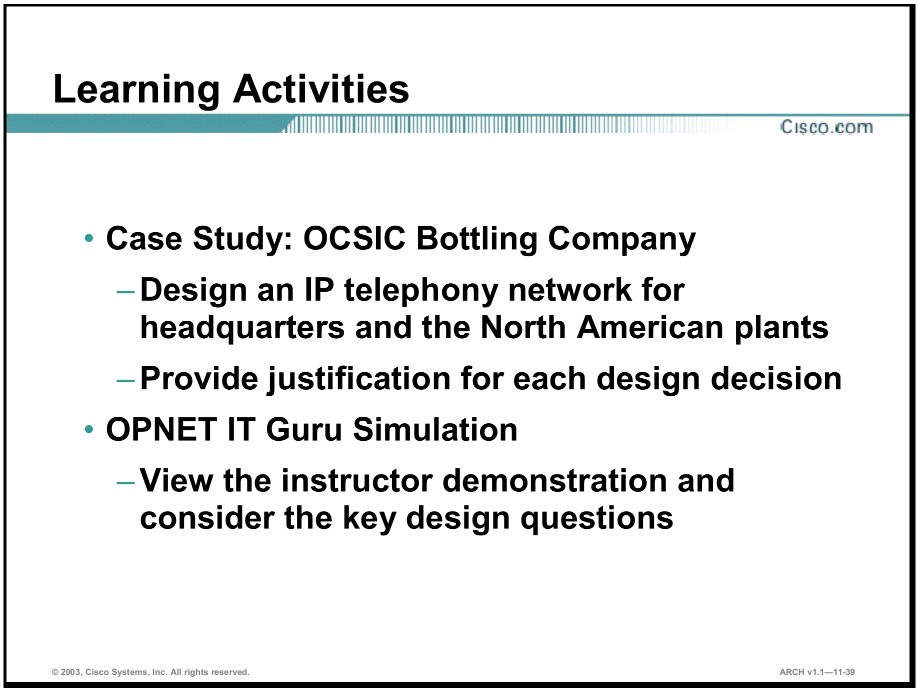

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Design an IP telephony network for headquarters and the North American plants
  - Provide justification for each design decision
- OPNET IT Guru Simulation
  - View the instructor demonstration and consider the key design questions

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—11-39

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

OCSIC Bottling Company wants to design a multisite telephony solution for the headquarters and North American plant locations.

In this exercise, you will design an IP telephony network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design an IP telephony network for headquarters and the North American plants

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Design an IP Telephony Network for OCSIC Bottling Company

Complete these steps:

Step 1 Complete the table to design the details about the IP telephony network.

Design Questions	Decision	Justification
Is centralized or distributed call processing most appropriate?		
How many CallManager servers will you deploy? In what locations? Which server will be the publisher? Which servers will be the subscribers? What other servers would you include in the cluster?		
What gateways will you deploy? Where will the gateways be located? What function will each gateway serve?		
What QoS strategy will you deploy to support the solution?		
What DSP resources are required for the solution?		
What transcoding resources are required for the solution?		
What are the network bandwidth and traffic engineering considerations?		

Step 2 Update your network diagrams to indicate the location of each CallManager server at headquarters and at the North American plants.

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have designed a voice network for the headquarters locations that includes CallManagers, dial plan, gateways, QoS, DSP, and transcoding resources.
- You have designed capabilities into the North American plant network to support IP telephony.

OPNET IT Guru Simulation 11-2

This simulation demonstrates the affect IP telephony, combined with QoS, has on network and application performance.

Review the simulation as presented by your instructor. When the simulation is complete, consider the following questions:

- How would you modify your network design based on the OPNET IT Guru simulation?
- Which option is most effective for the network in terms of performance, scalability, availability, and cost-effectiveness?

For more information about OPNET, visit www.opnet.com or send an e-mail to opnet_info@opnet.com.

Module 12

Designing Content Networking Solutions

Overview

Content networking is the replication and placement of copies of content closer to the user or groups of users through caching and Content Delivery Network (CDN) technologies. Caching is the process by which content is copied and delivered to the user upon demand. CDN refers to the positioning of content closer to the user in advance of the user’s request. These technologies support applications such as corporate communication and e-learning by providing services such as live broadcast, video and audio on demand (media on demand), and rebroadcast.

Module Objectives

Upon completing this module, you will be able to design enterprise solutions for content networking, given enterprise network needs.

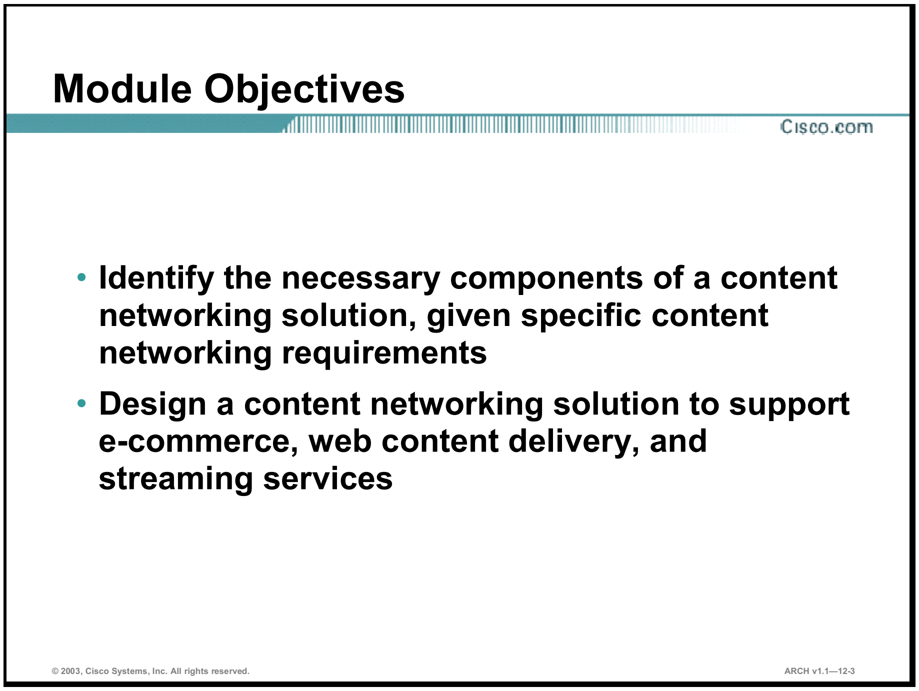

Module Objectives

- Identify the necessary components of a content networking solution, given specific content networking requirements
- Design a content networking solution to support e-commerce, web content delivery, and streaming services

ARCH v1.1—12-3

Module Outline

The outline lists the components of this module.

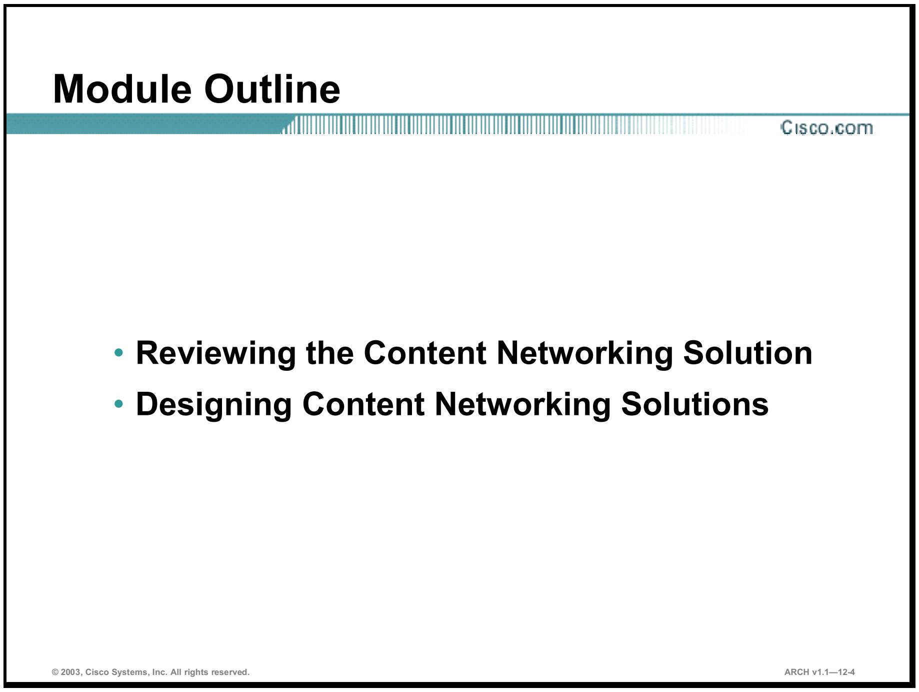

Module Outline

- Reviewing the Content Networking Solution
- Designing Content Networking Solutions

ARCH v1.1—12-4

Reviewing the Content Networking Solution

Overview

More and more applications require large volumes of content to be delivered over the network. Streaming video, e-learning, and graphic-intensive web sites can place a strain on an enterprise's network resources. A content networking architecture comprises several components that allow an enterprise to optimize web site performance and content delivery.

Importance

Content networking extends an IP service infrastructure to enable content services opportunities for web-driven enterprises. Content networking makes networks more efficient in delivering the content required by many applications.

Objectives

Upon completing this lesson, you will be able to identify the necessary components of a content networking solution, given specific content networking requirements. This includes being able to meet these objectives:

- Identify enterprise requirements for content networking
- Describe the purpose of each component of a content networking architecture
- Identify the content distribution and management components of a content networking solution, given specific content networking requirements
- Identify the content routing components of a content networking solution, given specific content networking requirements
- Identify the content switching components of a content networking solution, given specific content networking requirements
- Identify the content edge delivery components of a content networking solution, given specific content networking requirements
- Explain how to integrate intelligent networking services to support content networking solutions

Learner Skills and Knowledge

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

- Successful completion of the *Designing for Cisco Internetwork Solutions* (DESGN) course

Outline

The outline lists the topics included in this lesson.

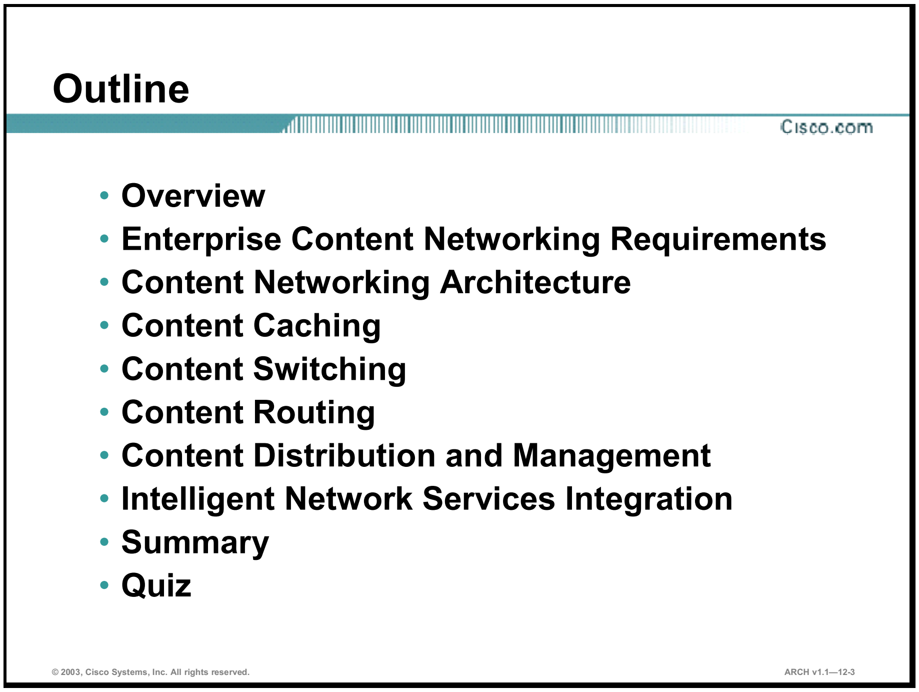

Outline

Cisco.com

- Overview
- Enterprise Content Networking Requirements
- Content Networking Architecture
- Content Caching
- Content Switching
- Content Routing
- Content Distribution and Management
- Intelligent Network Services Integration
- Summary
- Quiz

ARCH v1.1—12-3

Enterprise Content Networking Requirements

As the data (content) being accessed over the network increases in size and complexity, content networking becomes a requirement. A content networking solution supports bandwidth optimization, server scalability, response-time reduction, and the large-scale deployment of rich content such as video or audio. This topic identifies enterprise requirements for content networking.

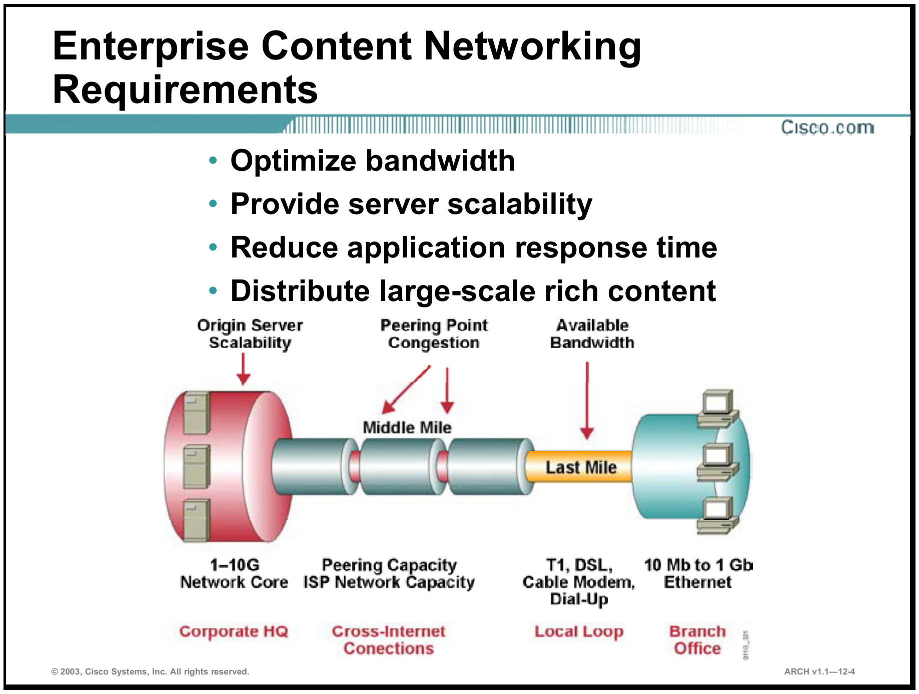

Enterprise Content Networking Requirements

Cisco.com

- Optimize bandwidth
- Provide server scalability
- Reduce application response time
- Distribute large-scale rich content

ARCH v1.1—12-4

Users are accessing web content more and more frequently. In addition, the content being delivered continues to increase in size and complexity. With the increase in activity and content size, enterprises can experience bandwidth bottlenecks, overloaded servers, and degraded response times. New applications such as broadcast video and media-on-demand can place an untenable strain on network resources.

Implementing a content networking solution can address these issues in several ways:

- Moving content closer to the users reduces bandwidth requirements over the WAN or Internet and improves response time.
- Offloading server content, load balancing, and redundancy allow enterprises to scale server availability without necessarily increasing the number of servers.
- Distributing content to multiple locations in a controlled manner can enable new applications without overstressing the available network resources.

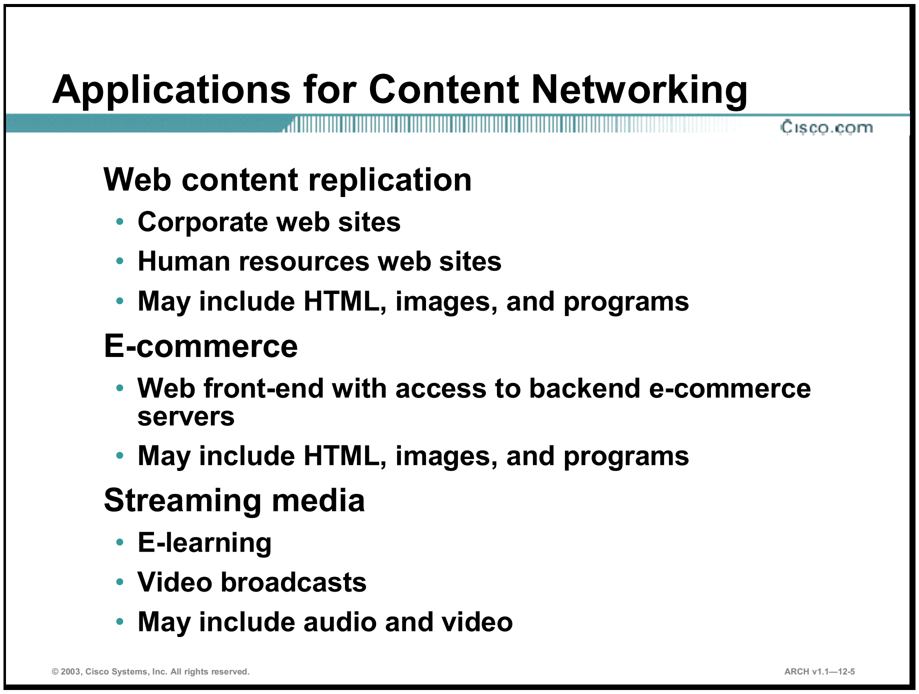

Applications for Content Networking

Cisco.com

Web content replication

- Corporate web sites
- Human resources web sites
- May include HTML, images, and programs

E-commerce

- Web front-end with access to backend e-commerce servers
- May include HTML, images, and programs

Streaming media

- E-learning
- Video broadcasts
- May include audio and video

ARCH v1.1—12-5

Content networking supports rich-media-intensive, distributed e-business applications for both live and on-demand streaming application support.

Content Networking Architecture

Cisco’s content networking solutions include content edge delivery (caching), content switching, content routing, and content delivery and management. This topic describes the purpose of each component of the content networking architecture.

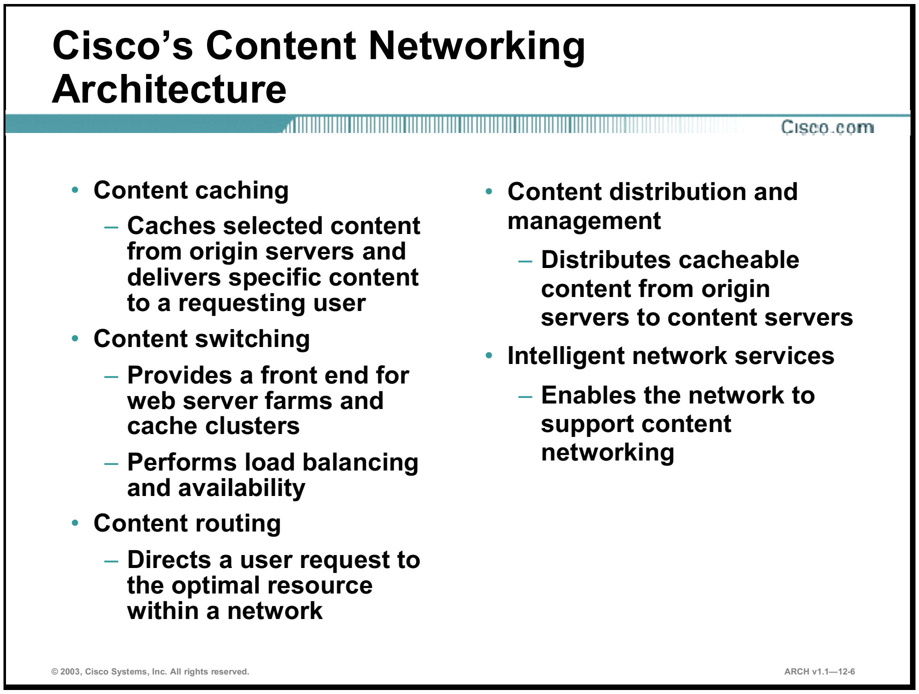

Cisco’s Content Networking Architecture

Cisco.com

- Content caching
  - Caches selected content from origin servers and delivers specific content to a requesting user
- Content switching
  - Provides a front end for web server farms and cache clusters
  - Performs load balancing and availability
- Content routing
  - Directs a user request to the optimal resource within a network
- Content distribution and management
  - Distributes cacheable content from origin servers to content servers
- Intelligent network services
  - Enables the network to support content networking

ARCH v1.1—12-6

The content networking architecture recommended by Cisco is made up of five components:

- **Content caching:** Caches selected content from origin servers and delivers specific content to a requesting user.
- **Content switching:** Provides a front end for web server farms and cache clusters, performing functions such as load balancing and availability.
- **Content routing:** Directs a user request to the optimal resource within a network based on user-defined policies, such as rules for specific content, availability of content, health or current loads for web servers or caches, and various other network conditions.
- **Content distribution and management:** Provides the mechanism for proactively distributing cacheable content from origin servers to content servers throughout the network.
- **Intelligent network services:** Enables content networking by the use of tightly integrated intelligent network services, such as high availability, security, quality of service (QoS), and IP multicast.

Content Caching

A Content Engine (CEs) caches copies of data from origin servers, allowing the content to be served locally. A CE can be deployed in transparent mode, proxy mode, or reverse proxy mode. This topic describes the content caching components of a Cisco content networking solution.

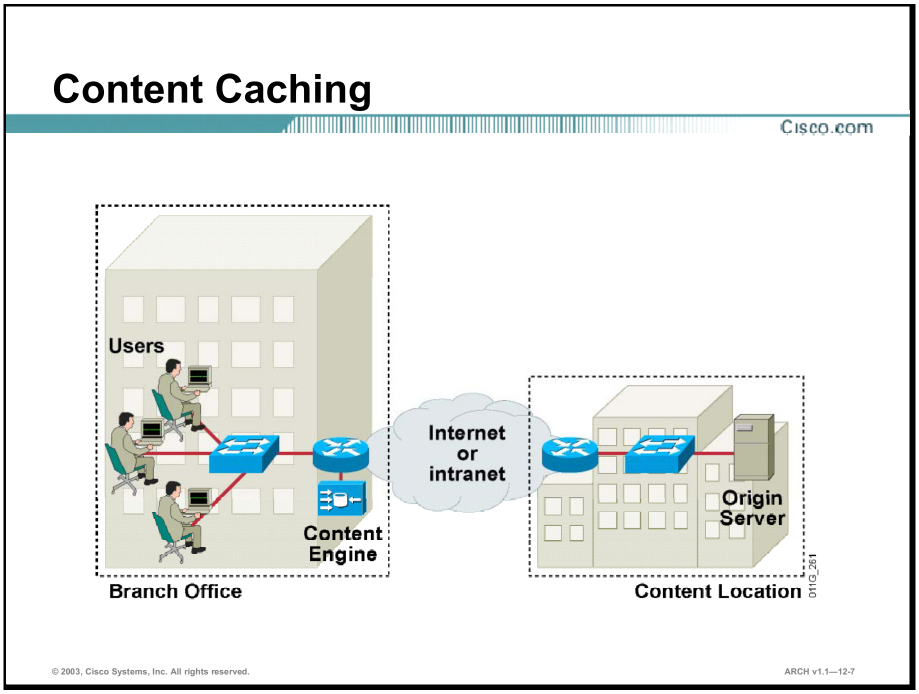

Content Caching

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-7

Caching is a demand-based replication and storage of cacheable content for the purpose of serving that same content to users making subsequent requests for it. Cacheable content is typically static application data associated with a file type and file extension, such as:

- Graphics files: .gif, .jpeg, .bmp
- Compressed files: .zip, .gz, .tar
- Document files: .txt, .pdf
- Multimedia files: .avi, .mov, .mpeg, .wav

CEs store or cache copies of content from origin servers. The CE can then process user requests rather than the origin server, reducing both network traffic and server processing. Web objects are typically cached, rather than entire web pages. Prime candidates for caching are graphics files, document files, compressed files, and multimedia files. Content engines can be proactively populated with specific content or simply allowed to cache requested content automatically.

The diagram shows a CE caching content in a location closer to end users than the origin server.

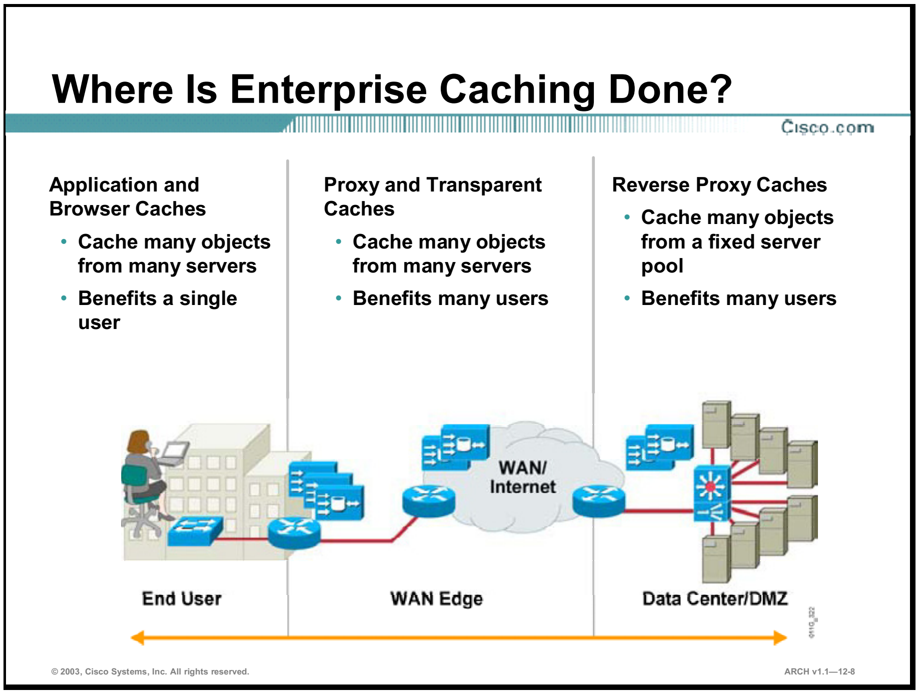

Where Is Enterprise Caching Done?

Application and Browser Caches

- **Cache many objects from many servers**
- **Benefits a single user**

Proxy and Transparent Caches

- **Cache many objects from many servers**
- **Benefits many users**

Reverse Proxy Caches

- **Cache many objects from a fixed server pool**
- **Benefits many users**

ARCH v1.1—12-8

Enterprise caching can be accomplished between a central site and any remote site. At the central site, caching is deployed at the data center or within a perimeter LAN. Typically, caching takes place at the remote site, close to users.

The three deployment models for content caching include: transparent, proxy, and reverse proxy. With each caching mechanism, you can specify how content is stored and when it is cleared from the content server.

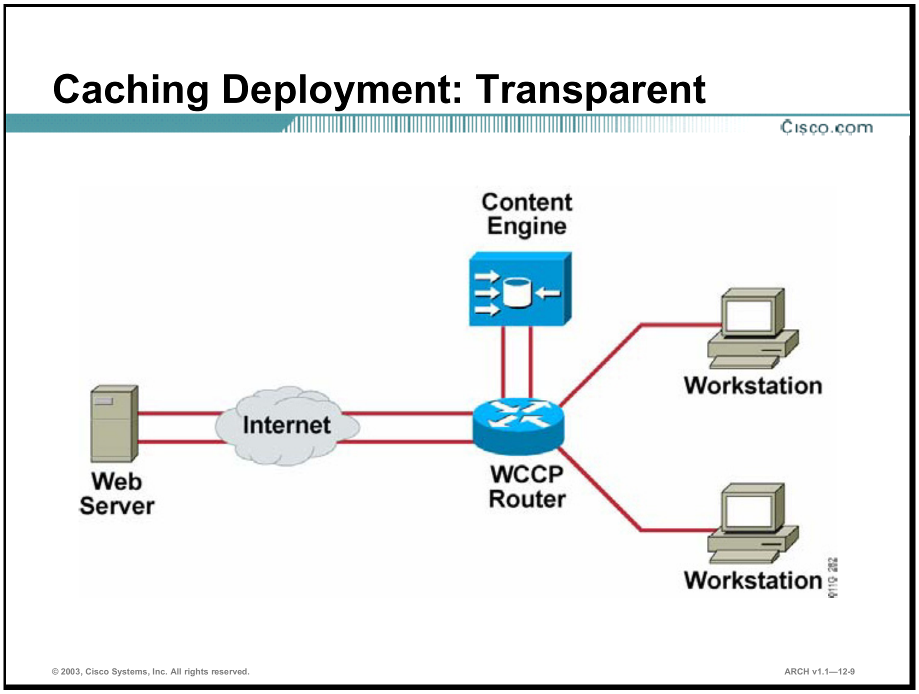

Caching Deployment: Transparent

Cisco.com

ARCH v1.1—12-9

With transparent caching, a user request for content (usually through a web browser) is redirected by a router or switch to a CE. On the first request for a specific piece of content, the CE initiates a request to the origin server to retrieve the content, then stores it locally and returns the content to the user making the request. On subsequent requests for that same piece of content, the content can be served locally, eliminating the need for a request to travel to the origin server. The CEs deployed in transparent mode are placed close to the users making the requests.

To deploy transparent caching, consider these best practices:

- Place the cache at the network edge between the user and the servers, close to the user population.
- Place the cache in the path of outbound network traffic at the network edge.
- Understand the network topology and traffic patterns.

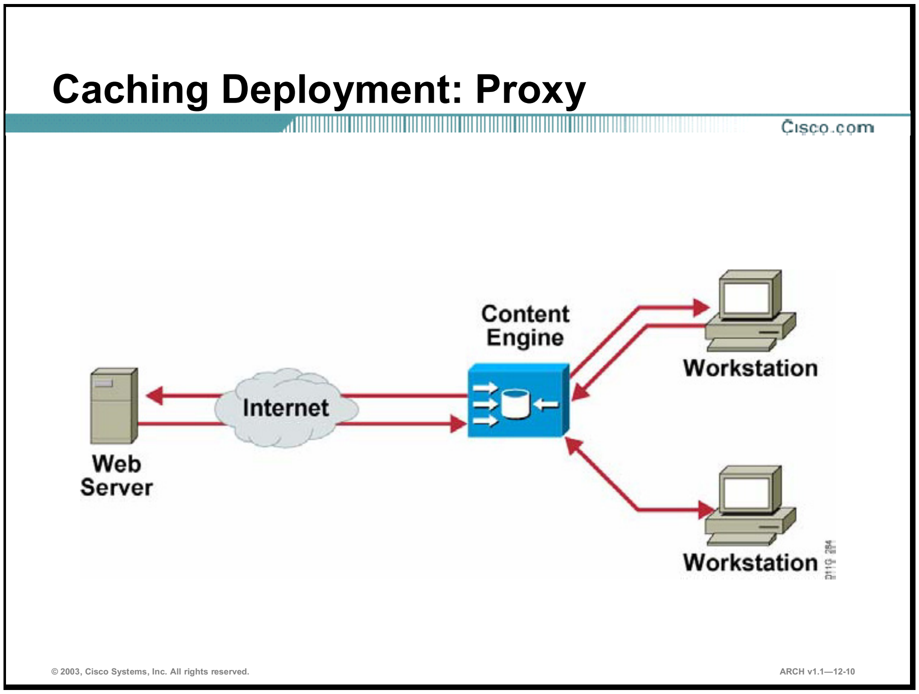

Caching Deployment: Proxy

ARCH v1.1—12-10

With proxy caching, the CE acts as a proxy for all content requests. The user's browser is typically modified to send requests directly to the CE rather than to the origin server. The CE then forwards the request to the origin server, if necessary, or serves the content locally. CEs in proxy mode are placed close to the users making the requests, but far enough along the traffic path to the servers to intercept all such traffic. Typically, these CEs are placed in the edge, either at a WAN or Internet access point.

In proxy mode, end-user web browsers need to be explicitly configured to the IP address or host name of the CE, and there is no need for additional hardware such as Layer 4 switches or Web Cache Communication Protocol (WCCP)-enabled routers to intercept user requests, as in transparent caching. Enterprises are normally interested in deploying transparent network caching, but some enterprises may have a legacy requirement for a proxy (nontransparent) cache.

To deploy proxy caching, consider these best practices:

- Position the cache at the Internet edge.
- Manually configure client browsers to send all requests to the proxy. The proxy acts on behalf of the client.

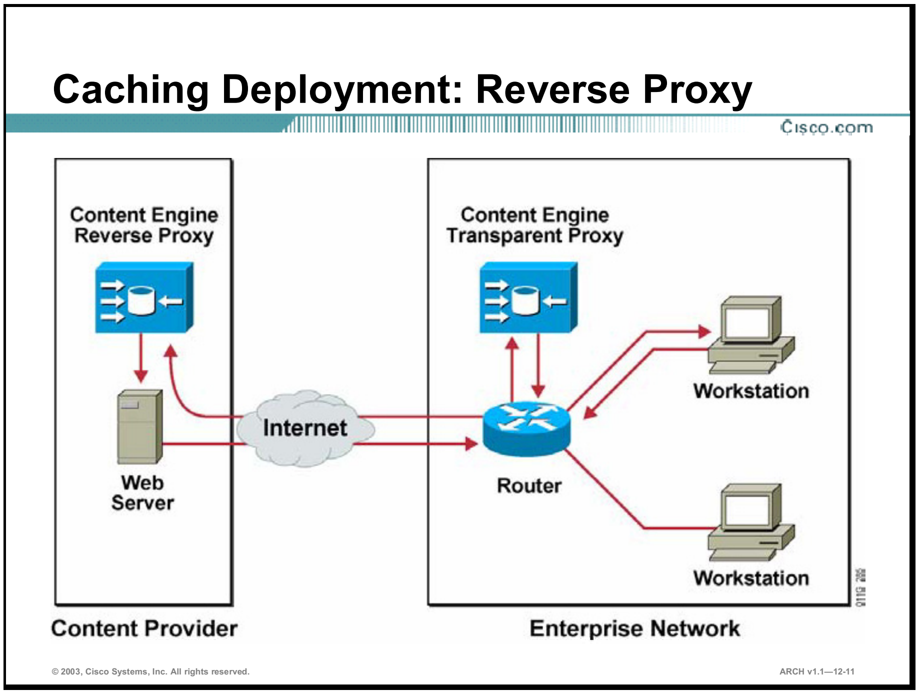

Caching Deployment: Reverse Proxy

ARCH v1.1—12-11

With reverse proxy caching, static content from a server is offloaded onto a CE. Requests destined for the server are directed to the CE, which serves the content locally, freeing the server from processing multiple requests for static content. CEs in reverse proxy mode are placed close to the servers being offloaded, typically in an internal server farm or an edge server farm providing outbound web or e-commerce services.

Use reverse proxy on CEs that are deployed in front of web server farms to increase the server farm capacity and improve web site performance.

To deploy reverse proxy caching, consider these best practices:

- Position the reverse proxy in front of web server farms to minimize the transaction processing for the back-end infrastructure.
- Deploy reverse proxies in data centers and perimeter LANs for Internet-facing web server farms.

Content Switching

Content switches provide load balancing and availability features for multiple content or application servers. This topic describes the content switching components of a Cisco content networking solution.

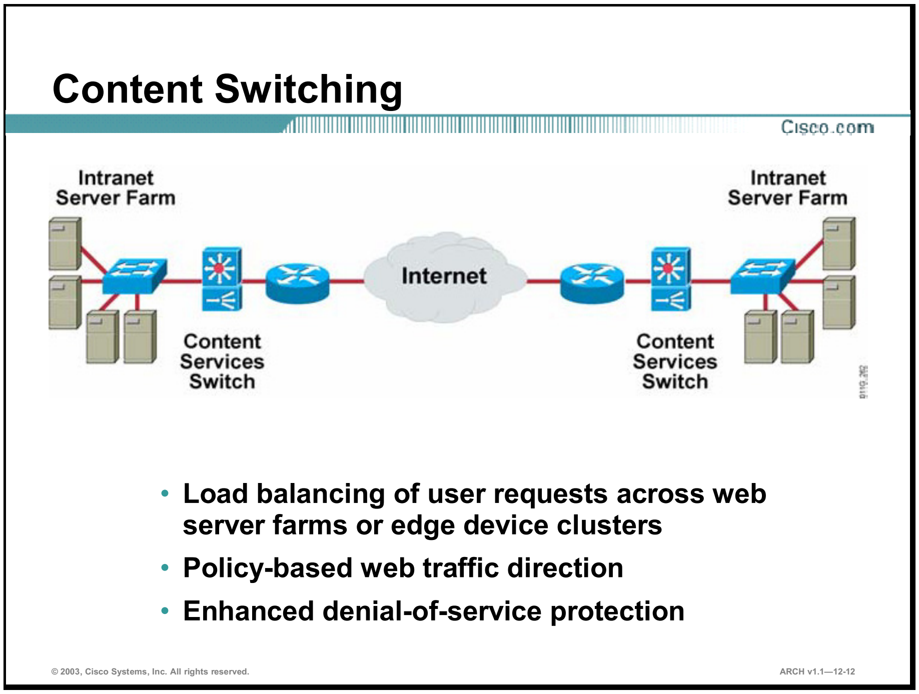

Content Switching

- Load balancing of user requests across web server farms or edge device clusters
- Policy-based web traffic direction
- Enhanced denial-of-service protection

ARCH v1.1—12-12

Content switching intelligently load balances traffic across servers or CEs based on the availability of the content and the load on the server. Content Services Switches (CSSs) provide these services:

- Local and global load balancing of user requests across server farms or CE clusters to improve performance, scalability, and content availability
- Policy-based web traffic direction based on full visibility of URLs, host tags, and cookies
- Enhanced denial-of-service protection, cache and firewall load balancing, and flash-crowd management (flash-crowd refers to unpredictable, event-driven traffic surges that swamp servers and disrupt site services)

With content switches, multiple web or application servers can be represented with a single IP address. The CSS load-balancing algorithms intelligently distribute content requests based on round-robin, weighted round-robin, least connections, weighted least connections, HTTP header content, or URL. The CSS performs server health checks to verify content and server availability.

Content switches are "smart" devices with sophisticated load-balancing capabilities and content-acceleration intelligence. Use them to add scalability to large networks.

Content Routing

Content routing redirects an end-user request to the best site based on a set of metrics such as delay, topology, or server load, and a set of policies such as location of content. This topic describes the content routing components of a content networking solution.

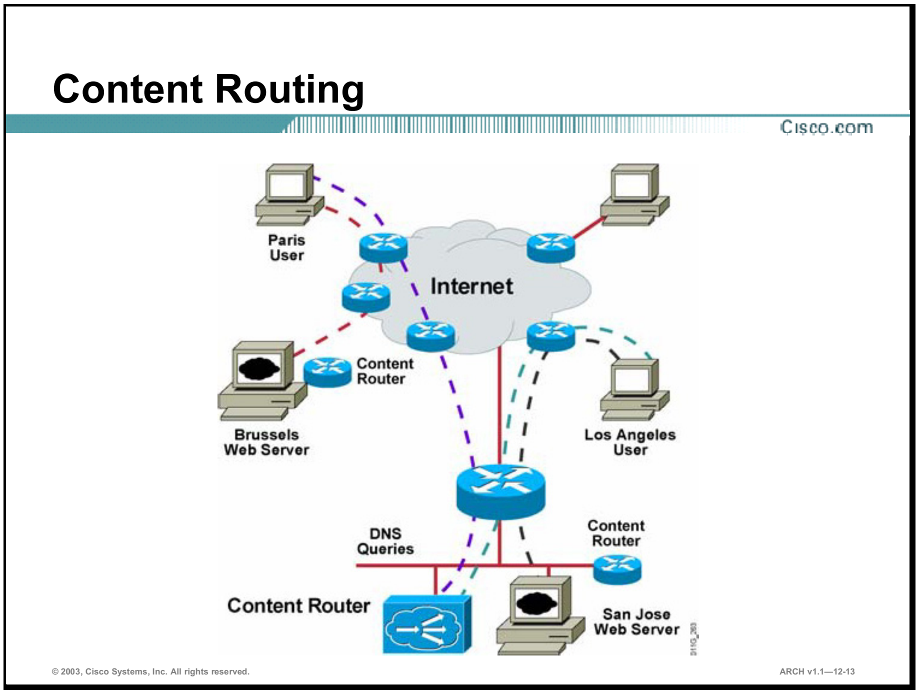

Content Routing

ARCH v1.1—12-13

Cisco routers and switches incorporate Web Cache Communication Protocol (WCCP) software to enable content routing capabilities. Additionally, Cisco offers content routers specifically designed to support large-scale mirrored web sites.

Content routing routes user requests to the replicated-content site (typically a mirror site) that can serve them most quickly and efficiently. The content routing software redirects a user request to the closest (best) replicated-content site, based on network delay, using a software process called boomerang. The content routing software load balances up to 500 sites for each domain it is configured to support.

A content router can be deployed on a network in two different ways: It can be set up in direct mode, or it can be set up in WCCP mode. Both deployments involve setting up a content routing agent at each content site within each domain you want the content router to support. Content routing agents are machines (such as CEs) that have been configured to interact with the content router.

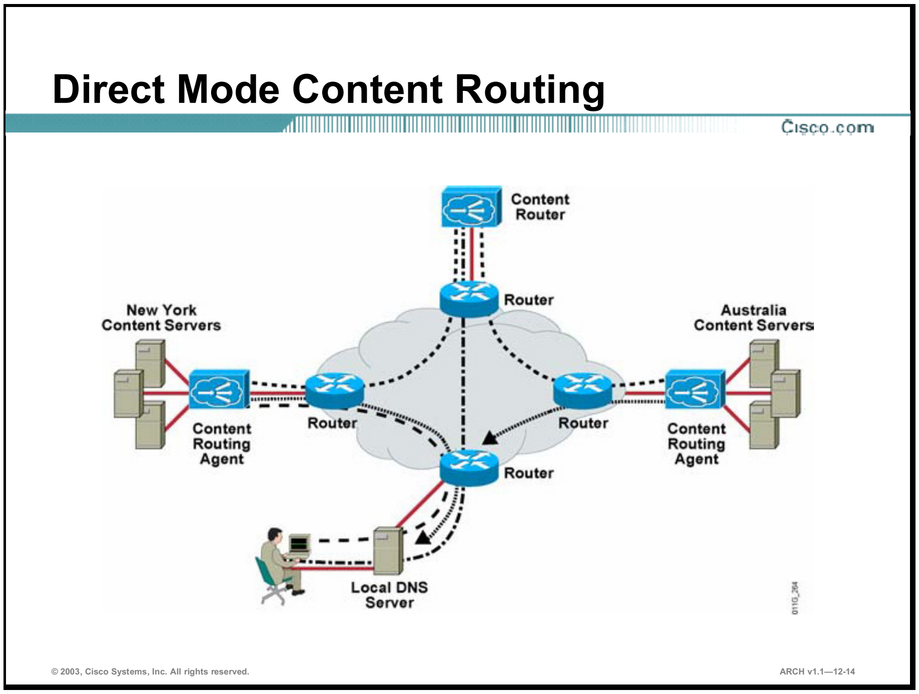

Direct Mode Content Routing

ARCH v1.1—12-14

In direct mode, the content router is configured as the authoritative DNS server for the fully qualified domain name being routed. For example, to route www.enterprise.com, the address record in the primary DNS server for www.foo.com is changed to a name server (NS) record pointing to the content router. All requests for the IP address of www.enterprise.com are handled by the content router and its content routing agents. When a request arrives, a specific number of agents respond at exactly the same time. This is called the DNS race because the agent that sent the first response received is the winner of the DNS race and is therefore the site to which the user will connect. The figure provides an overview of the direct mode content routing process, which follows these steps:

Step 1 A user requests a connection from a web browser or other service.

Step 2 A DNS request is sent to the content router.

Step 3 The content router forwards the request to content routing agents.

Step 4 Agents simultaneously send responses back to the local DNS server. The first response through the network contains the IP address of the best site.

Step 5 The user connects to the best site.

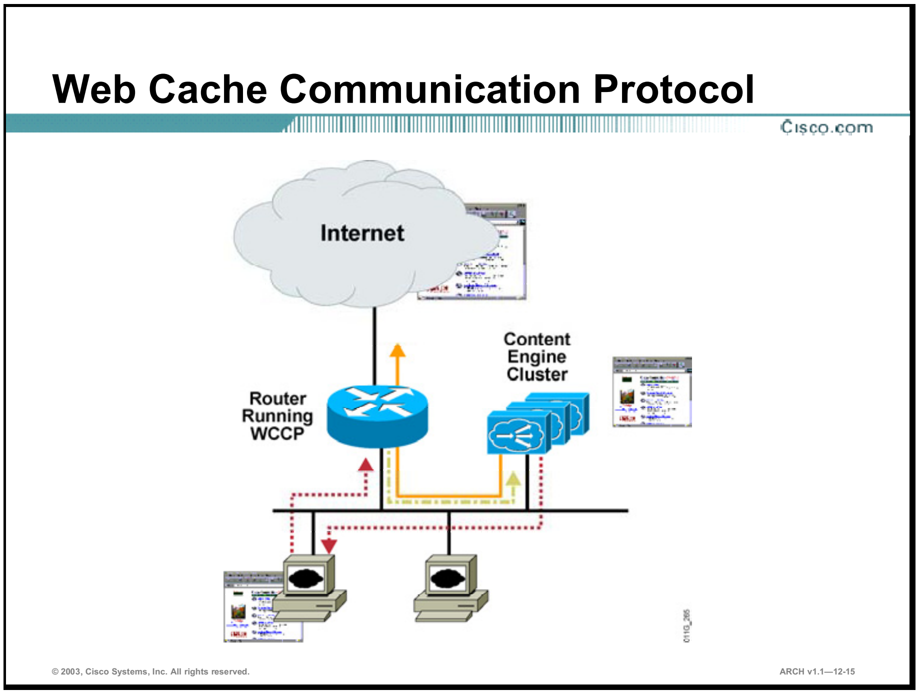

Web Cache Communication Protocol

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-15

WCCP provides the mechanism for content routers and other WCCP-enabled routers and switches to redirect user requests to the appropriate destination. WCCP includes load balancing, scaling, fault tolerance, and service-assurance (failsafe) mechanisms, and includes these steps:

A user requests a connection from a web browser or other service.

Step 1 The WCCP-enabled router analyzes the request, and based on TCP port number, determines if it should transparently redirect it to a CE. Access lists can be applied to control which requests are redirected.

Step 2 If a CE does not have the requested content, it sets up a separate TCP connection to the end server to retrieve the content. The content returns to, and is stored on, the CE.

Step 3 The CE sends the content to the client. Upon subsequent requests for the same content, the CE transparently fulfills the requests from its local storage.

Content Distribution and Management

Content distribution and management provides the mechanism for importing, replicating, and managing content throughout the enterprise network. This topic describes the Cisco content distribution and management components of a content networking solution.

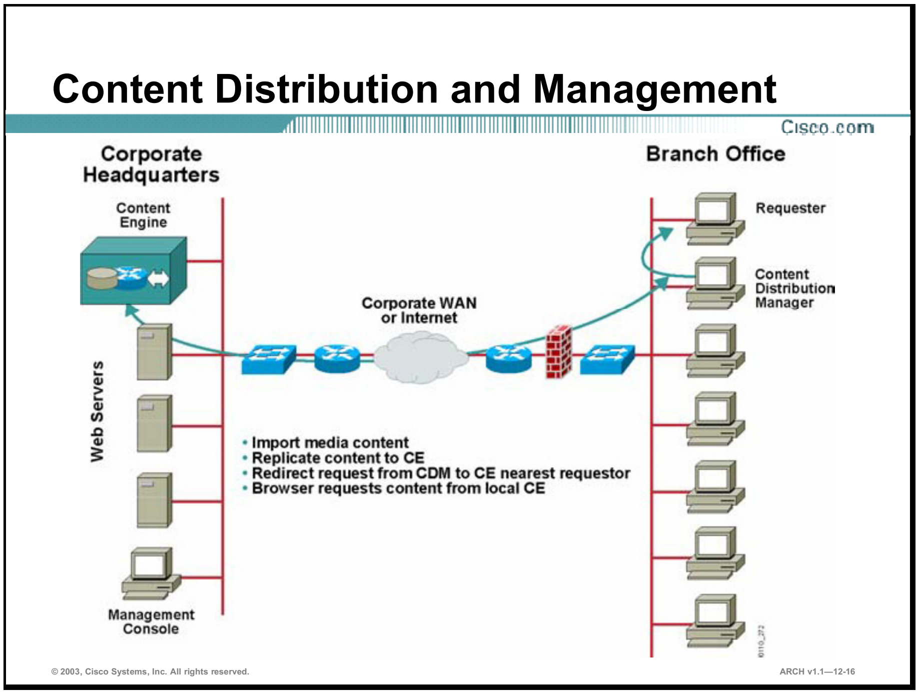

Content Distribution and Management

- Import media content
- Replicate content to CE
- Redirect request from CDM to CE nearest requestor
- Browser requests content from local CE

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—12-16

The Cisco CDM is a web-browser-based tool that provides the administrative function for content networking. With the CDM, you can configure and monitor CEs, import and preview media, and generate media URLs for access from web sites. You also set maximum bandwidth usage over the WAN from the CDM to the remote CEs, as well as maximum LAN bandwidth usage from the CEs to end-user desktops. You can schedule content distribution by time of day and day of week, allowing content replication to occur during off-peak hours.

The CDM usually resides at the corporate headquarters data center. Client requests for content are sent first to the CDM, which redirects the request to the most appropriate CE. The CDM can also capture usage and billing data.

Content management is used to automatically import, maintain copies, and configure content at the edge of the network. The CDM maintains a hierarchy of CEs, allowing them to replicate content to CEs lower in the hierarchy. Replication is performed without waiting for the entire content to be received at one location before distribution to the next location begins.

The Self-Organizing Distributed Architecture (SODA) is a Cisco technology allowing content networking devices to self-organize into a single, cooperating system. The CDM defines the network policies and then automatically stores the SODA network information, building routing tables for specific content. When a device is added to the network, it automatically configures itself in the network based on the network topology and content requirements.

Intelligent Network Services Integration

Content networking requires the support of intelligent networking services such as security, QoS, high availability, and IP multicast. This topic explains how to integrate intelligent networking services to support content networking solutions.

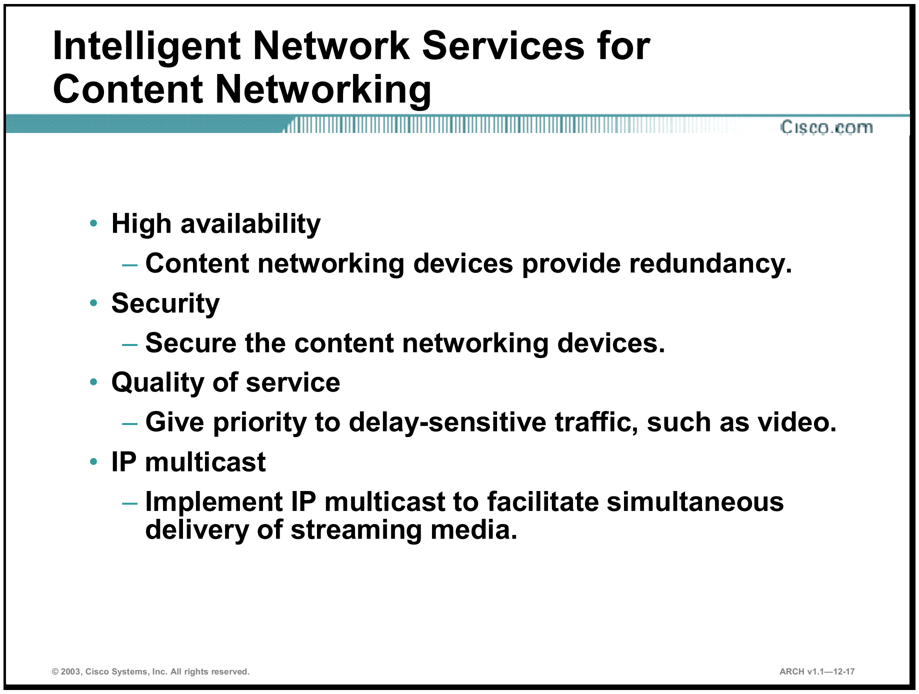

Intelligent Network Services for Content Networking

Cisco.com

- High availability
  – Content networking devices provide redundancy.
- Security
  – Secure the content networking devices.
- Quality of service
  – Give priority to delay-sensitive traffic, such as video.
- IP multicast
  – Implement IP multicast to facilitate simultaneous delivery of streaming media.

ARCH v1.1—12-17

When implementing a content networking solution, Cisco intelligent services should be integrated to provide these functions:

- **High availability:** Some of the content networking devices, like content switches and routers, are used to provide high availability for content. Redundant CEs, switches, and routers can further ensure availability for the most critical content.
- **Security:** The content being served can be highly sensitive, so in addition to the security features implemented throughout the network, you should take additional care to secure the content networking devices themselves.
- **QoS:** When content networking involves the delivery of delay-sensitive data such as audio or video, QoS features allow such traffic to be given priority over other data traffic.
- **IP multicast:** When content networking is used to facilitate simultaneous delivery of streaming media, IP multicast must be implemented throughout the network to support the broadcast requirements.

Summary

This topic summarizes the key points discussed in this lesson.

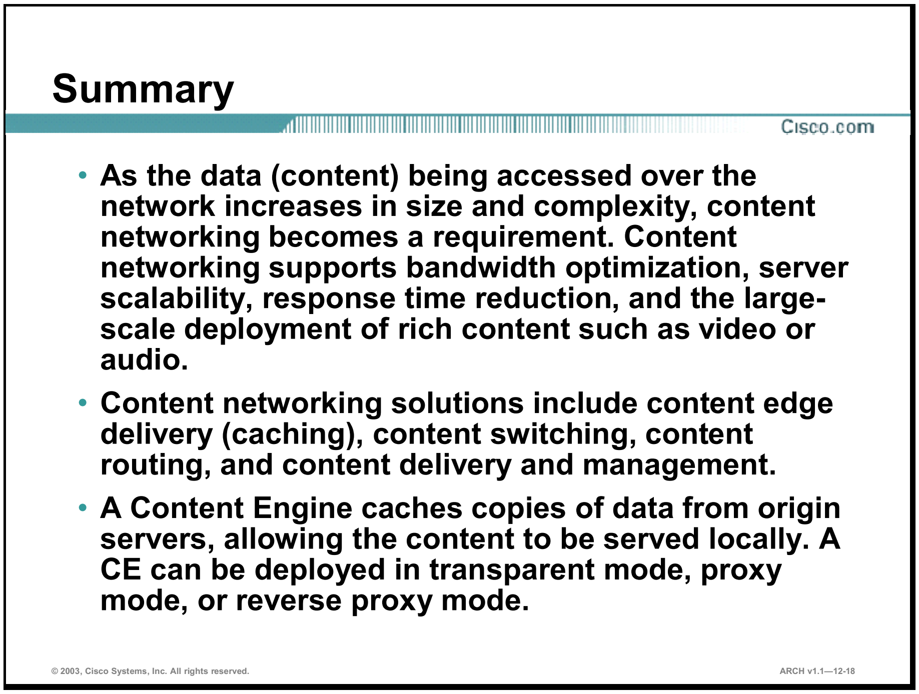

Summary

Cisco.com

- As the data (content) being accessed over the network increases in size and complexity, content networking becomes a requirement. Content networking supports bandwidth optimization, server scalability, response time reduction, and the large-scale deployment of rich content such as video or audio.
- Content networking solutions include content edge delivery (caching), content switching, content routing, and content delivery and management.
- A Content Engine caches copies of data from origin servers, allowing the content to be served locally. A CE can be deployed in transparent mode, proxy mode, or reverse proxy mode.

ARCH v1.1—12-18

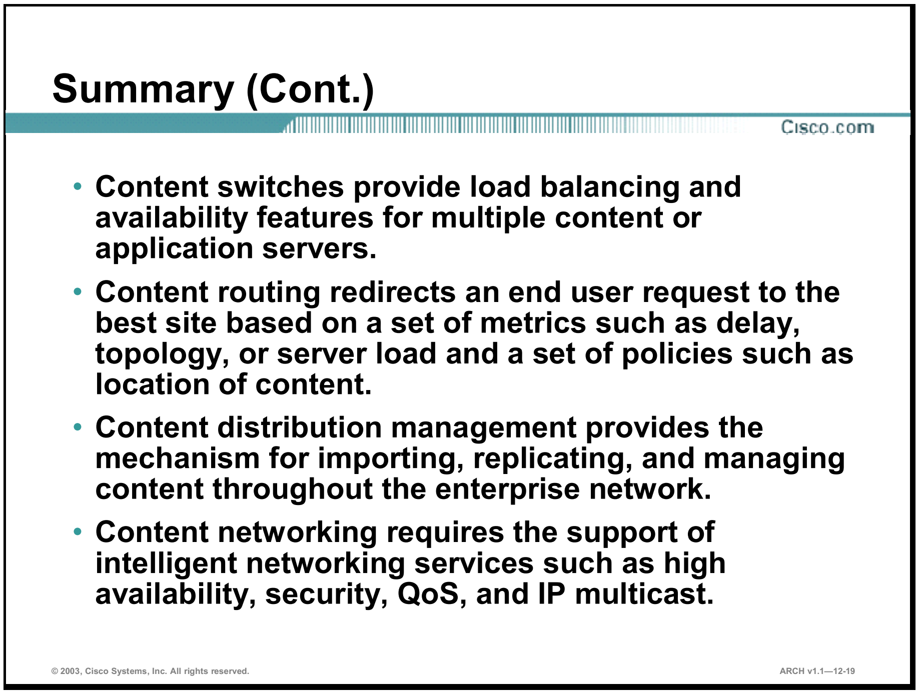

Summary (Cont.)

Cisco.com

- Content switches provide load balancing and availability features for multiple content or application servers.
- Content routing redirects an end user request to the best site based on a set of metrics such as delay, topology, or server load and a set of policies such as location of content.
- Content distribution management provides the mechanism for importing, replicating, and managing content throughout the enterprise network.
- Content networking requires the support of intelligent networking services such as high availability, security, QoS, and IP multicast.

ARCH v1.1—12-19

References

For additional information, refer to these resources:

- *Accelerating Web Applications with Cisco Enterprise Content Delivery Networks* at http://www.cisco.com/warp/public/cc/pd/cxsr/ces/prodlit/awace_wp.htm
- *Network Caching* at http://www.cisco.com/warp/public/cc/pd/cxsr/500/tech/cds_wp.htm
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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 two benefits can an enterprise network realize by moving content closer to users? (Choose two.)

A) server scalability

B) improved reliability

C) improved response time

D) reduced LAN bandwidth

E) reduced WAN bandwidth utilization

Q2) Which component of the Cisco content networking architecture directs a request to the best destination location to serve that request?

A) content routing

B) content caching

C) content switching

D) content distribution and management

Q3) Where in the network are Content Engines in reverse proxy mode typically deployed?

A) close to the server

B) close to the end users

C) in the Building Access submodule

D) in the Campus Backbone submodule

Q4) What are two benefits that can be achieved by implementing content services switches? (Choose two.)

A) reduced server load

B) increased throughput

C) reduced network traffic

D) improved server scalability

E) increased content availability

Q5) What is the role of a WCCP-enabled router in a content networking solution?

A) to serve flows

B) to terminate flows

C) to initiate flows to the origin server

D) to coordinate flow delivery decisions with the delivery device

Q6) What are the two ways in which you can deploy a content router? (Choose two.)

A) in direct mode

B) in WCCP mode

C) as a failsafe server

D) as a CE

E) close to the DNS server

Q7) What two functions of the CDM allow content to be replicated with minimal impact on a production network? (Choose two.)

A) scheduling

B) bandwidth limiting

C) content previewing

D) distribution hierarchy

E) content request redirection

Q8) Which intelligent service is required to support broadcast of streaming media?

A) QoS

B) security

C) IP multicast

D) high availability

Quiz Answer Key

Q1) C, E

Relates to: Enterprise Content Networking Requirements

Q2) A

Relates to: Content Networking Architecture

Q3) A

Relates to: Content Caching

Q4) D, E

Relates to: Content Switching

Q5) D

Relates to: Content Routing

Q6) A, B

Relates to: Content Routing

Q7) A, B

Relates to: Content Distribution and Management

Q8) C

Relates to: Intelligent Network Services Integration

Designing Content Networking Solutions

Overview

Cisco content networking solutions enable enterprises to build networks that can deliver e-commerce, web content delivery, and streaming media. Content networking solutions give enterprises control in allocating site resources for effective utilization.

Relevance

Enterprises can deploy end-to-end content networking solutions that enable Internet data centers and branch sites to provide content and delivery services.

Objectives

Upon completing this lesson, you will be able to design a content networking solution to support e-commerce, web content delivery, and streaming services. This includes being able to meet these objectives:

- List key design considerations for content networking solutions
- Select the necessary components of a content networking solution for web content delivery
- Select the necessary components of a content networking solution for e-commerce
- Select the necessary components of a content networking solution for streaming media

Learner Skills and Knowledge

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

- Reviewing the Content Networking Solution lesson

Outline

The outline lists the topics included in this lesson.

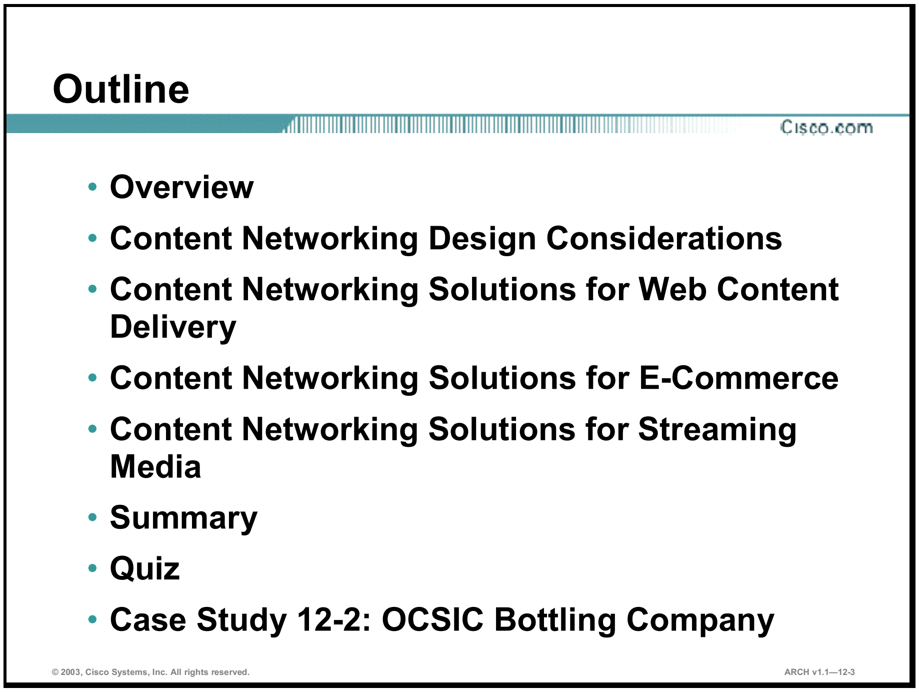

Outline

Cisco.com

- Overview
- Content Networking Design Considerations
- Content Networking Solutions for Web Content Delivery
- Content Networking Solutions for E-Commerce
- Content Networking Solutions for Streaming Media
- Summary
- Quiz
- Case Study 12-2: OCSIC Bottling Company

ARCH v1.1—12-3

Content Networking Design Considerations

When planning to implement a content networking solution, you should make sure that your design doesn't interfere with existing network topology or redundancy and that the solution is scaled sufficiently to avoid creating traffic bottlenecks. This topic describes the design considerations for content networking solutions.

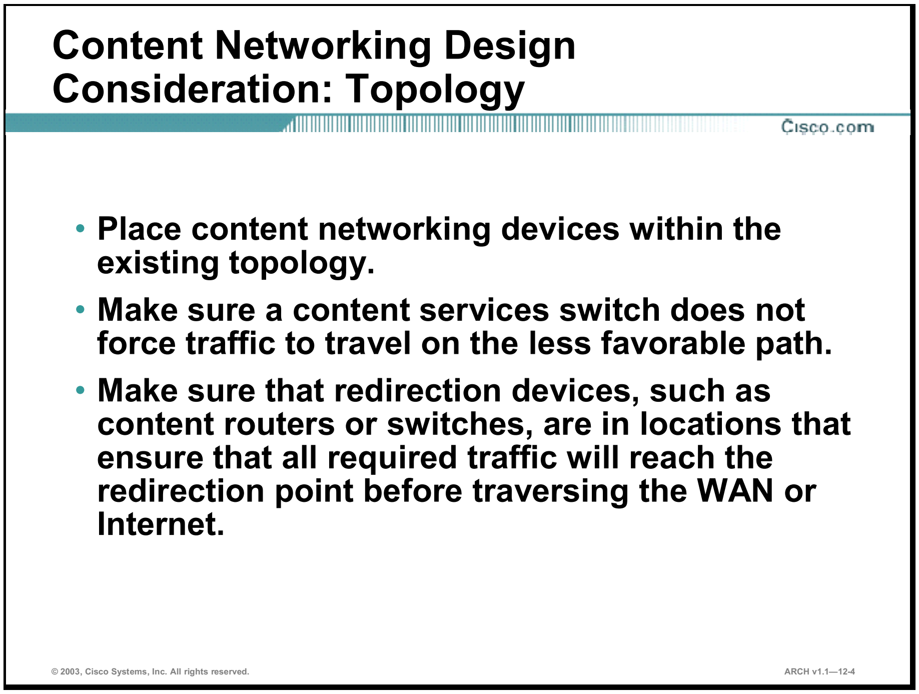

Content Networking Design Consideration: Topology

- Place content networking devices within the existing topology.
- Make sure a content services switch does not force traffic to travel on the less favorable path.
- Make sure that redirection devices, such as content routers or switches, are in locations that ensure that all required traffic will reach the redirection point before traversing the WAN or Internet.

ARCH v1.1—12-4

Place your content networking devices with consideration for the existing topology. Make sure that anywhere that your design redirects traffic for content purposes does not break the existing topology. For example, the existing network may have asymmetric traffic path possibilities. When you introduce a CSS, for example, make sure you do not force traffic to travel on the less favorable path. Make sure that you place redirection devices, such as content routers or switches, in locations that will ensure that all required traffic will reach the redirection point before traversing the WAN or Internet.

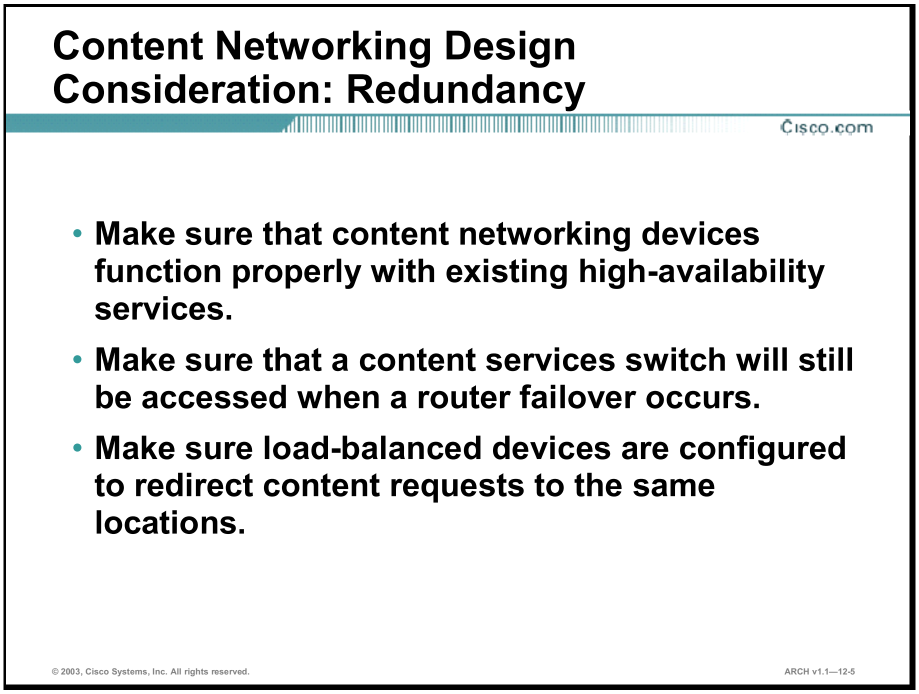

Content Networking Design Consideration: Redundancy

Cisco.com

- Make sure that content networking devices function properly with existing high-availability services.
- Make sure that a content services switch will still be accessed when a router failover occurs.
- Make sure load-balanced devices are configured to redirect content requests to the same locations.

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-5

Make sure that your content networking devices function properly with existing high-availability services. For example, make sure that a CSS will still be accessed when a router failover occurs, and that load-balanced devices are all configured to redirect content requests to the same locations.

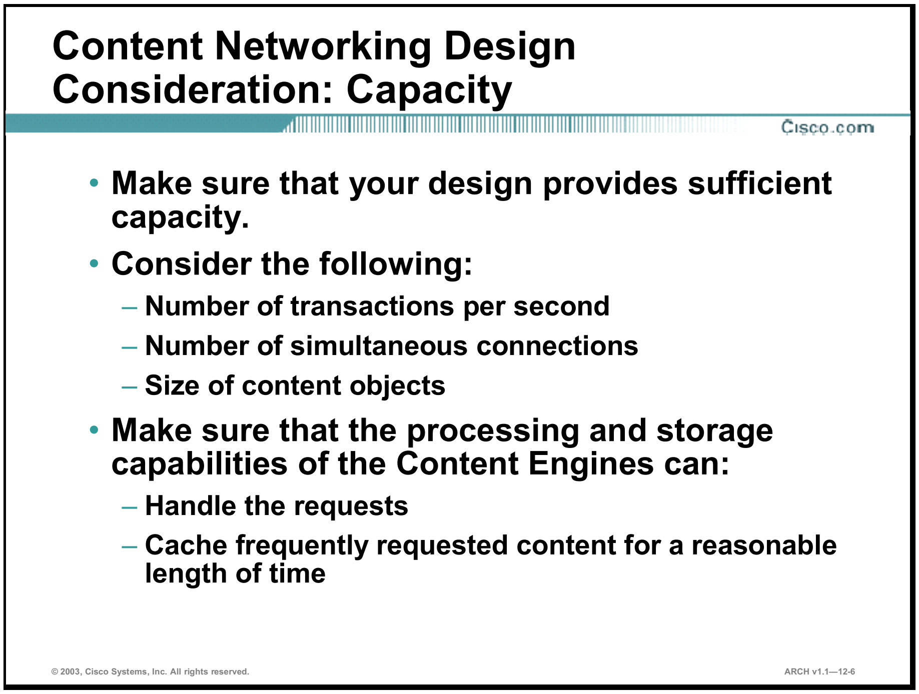

Content Networking Design Consideration: Capacity

- Make sure that your design provides sufficient capacity.
- Consider the following:
  - Number of transactions per second
  - Number of simultaneous connections
  - Size of content objects
- Make sure that the processing and storage capabilities of the Content Engines can:
  - Handle the requests
  - Cache frequently requested content for a reasonable length of time

ARCH v1.1—12-6

Make sure that your design provides sufficient capacity to avoid introducing undesirable effects. For example, when deciding how many CEs to install for a particular application, take into account the number of transactions per second, the number of simultaneous connections, and the size of the content objects that need to be supported. Make sure that the processing and storage capabilities of the CEs are sufficient to handle the requests and to cache frequently requested content for a reasonable length of time (at least 24 and preferably 72 hours).

There are three important parameters to consider when determining the sizing required for a particular cache installation:

- **Number of transactions per second required:** Identifies the amount of HTTP traffic
- **Number of simultaneous connections:** Defines the total number of HTTP flows that the cache will see at any single point in time
- **Size of objects:** Determines the amount of disk storage required, and the object hold time

To allow a cache to function effectively and reduce the amount of network bandwidth, the minimum cache storage time should be between 24 and 72 hours. The longer objects remain on disk, the larger the hit rate, but the larger the storage requirements. Sizing the disk requirements is a function of the number of transactions per second (TPS), average object size in bytes, and the expected cache hit rate in bytes. Use the following formula to determine the cache storage required for 24 hours:

- Average TPS * Average object size (bytes) * Number of seconds in 24 hours * (1 – expected cache hit rate [bytes]) / 1,000,000,000

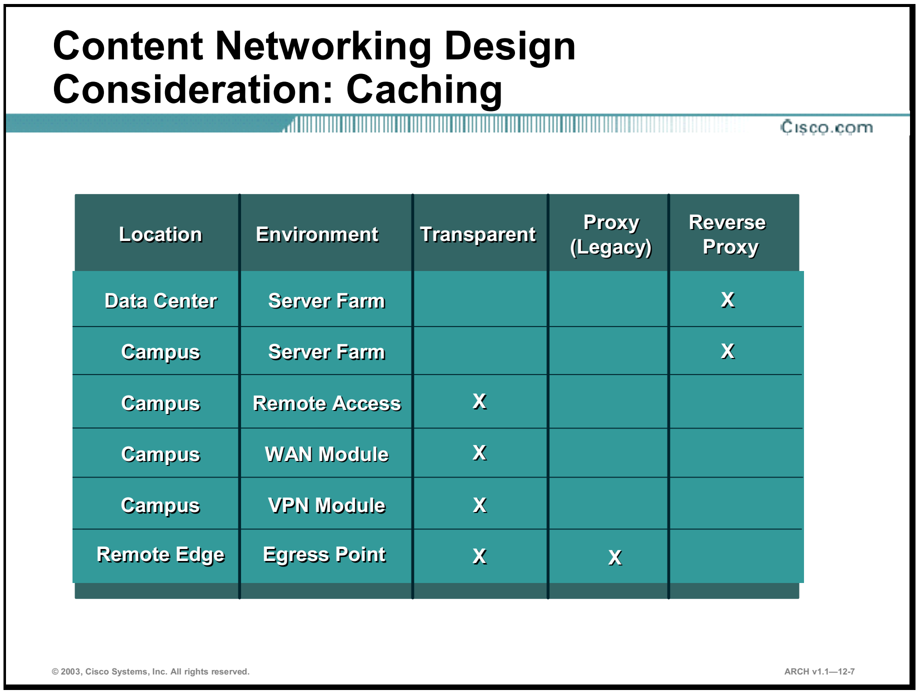

Content Networking Design Consideration: Caching

Location	Environment	Transparent	Proxy (Legacy)	Reverse Proxy
Data Center	Server Farm			X
Campus	Server Farm			X
Campus	Remote Access	X		
Campus	WAN Module	X		
Campus	VPN Module	X		
Remote Edge	Egress Point	X	X	

ARCH v1.1—12-7

The figure describes the recommended locations for deploying caching on the network.

Content Networking Solutions for Web Content Delivery

A content networking solution for web content delivery incorporates CEs to cache content closer to the users and offload internal web servers, content switches to provide load balancing, and WCCP-enabled routers to provide traffic redirection. This topic describes enterprise content networking requirements for web content delivery.

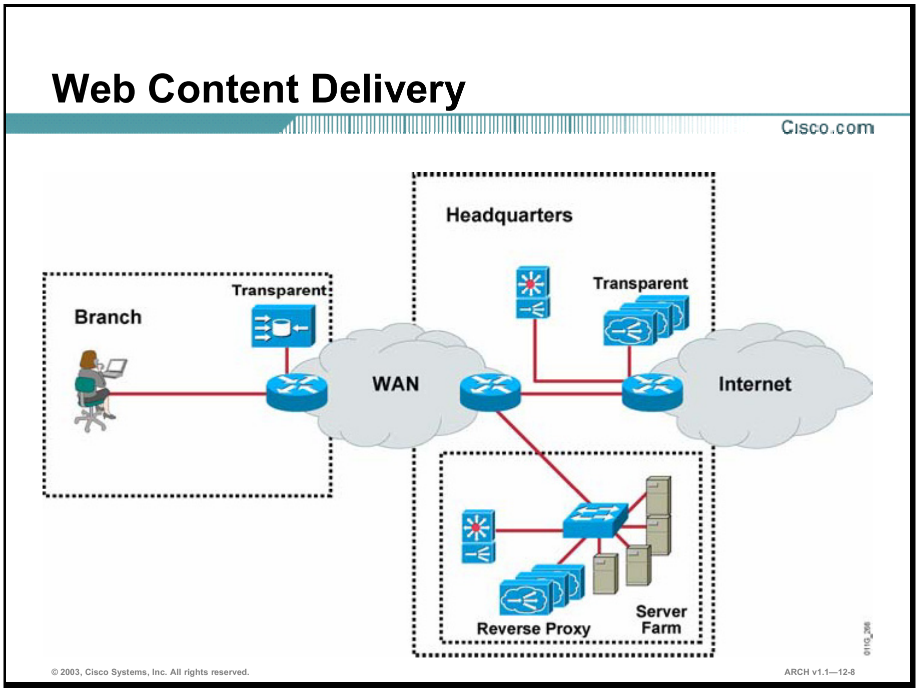

Web Content Delivery

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-8

A content networking implementation for web delivery should take into account access to both internal and external web servers from throughout the enterprise. The example implementation shown in the figure includes these elements:

- **Caching:** Content Engines in transparent mode are placed at the edge of the remote offices to cache content originating from the headquarters web servers or the over the Internet. CEs in transparent and/or proxy mode are placed at the Internet edge in the headquarters site, with another cluster of CEs in reverse proxy mode placed in the server farm to offload the internal web servers.
- **Content switching:** Because the caching requirements at the remote offices are minimal, no CSS is required. The clusters of CEs at the server farm and Internet edge each have an associated CSS to perform load balancing and manage redundancy.
- **Content routing:** WCCP is configured on selected Building Access or Building Distribution routers or switches at all sites to perform redirection of content requests to the appropriate CE or CSS.

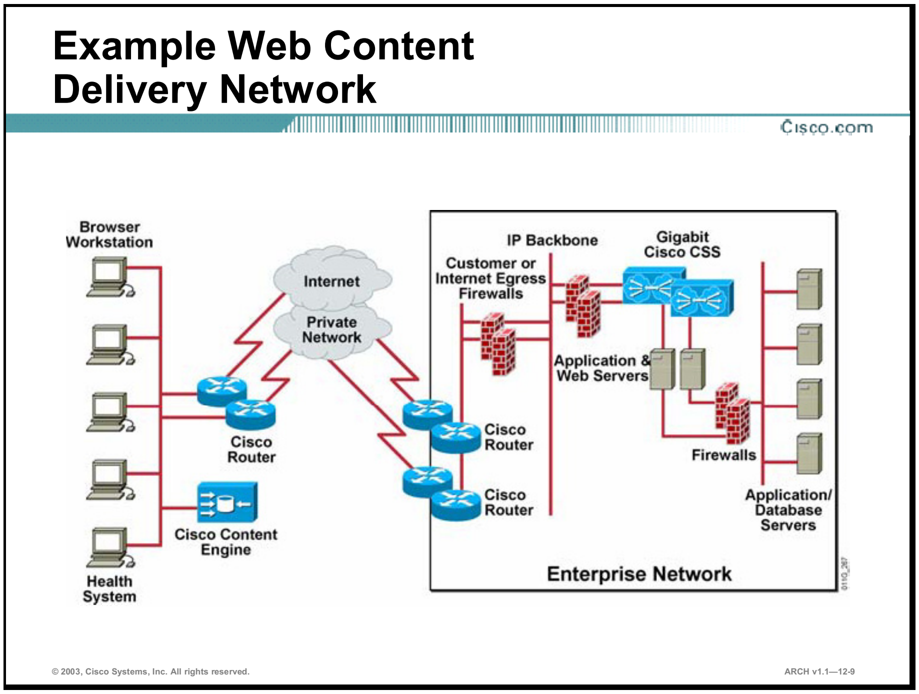

Example Web Content Delivery Network

ARCH v1.1—12-9

Company Background

This company is a health care provider that requires the network to be available 24 hours a day, 7 days a week, with no exceptions. The company supports complex health care applications including clinical information systems, financial information systems, and the Picture Archive and Communications System. The company is heavily regulated within the United States.

The company serves 6 of the top 10 healthcare systems, 22 public health institutions, and over 300 university teaching facilities in the United States. They also operate acute care institutions, clinics, and hospitals worldwide.

The company's goal is to accelerate content delivery for customers and accelerate the delivery of Internet and intranet content for employees located at their headquarters.

Content Networking Solution

The company deployed CEs at customer sites as part of its managed service offerings, which include a web-based order processing and registration application and a hospital scheduling application. Customers that outsource these applications deploy a CE at the customer premises. The CE supports specific Java-based health care applications. The CE at the customer site means that a customer downloads the applet once for all users to access.

At the company's headquarters, CEs were deployed on the campus network. The content is then delivered to the local-area network, which accelerates content delivery and saves WAN bandwidth.

The company network comprises redundant Cisco 3500 switches and redundant Cisco Catalyst 6500 switches. Everything within the environment is redundant, including the switch fabric back to the Cisco content services switches, firewalls, and routers, across a private WAN and out to a Cisco router at the customer site.

Content Networking Solutions for E-Commerce

A content networking solution for e-commerce will typically incorporate caching and content switching. A special consideration for e-commerce is session persistence, or stickiness. This topic describes enterprise content networking requirements for e-commerce.

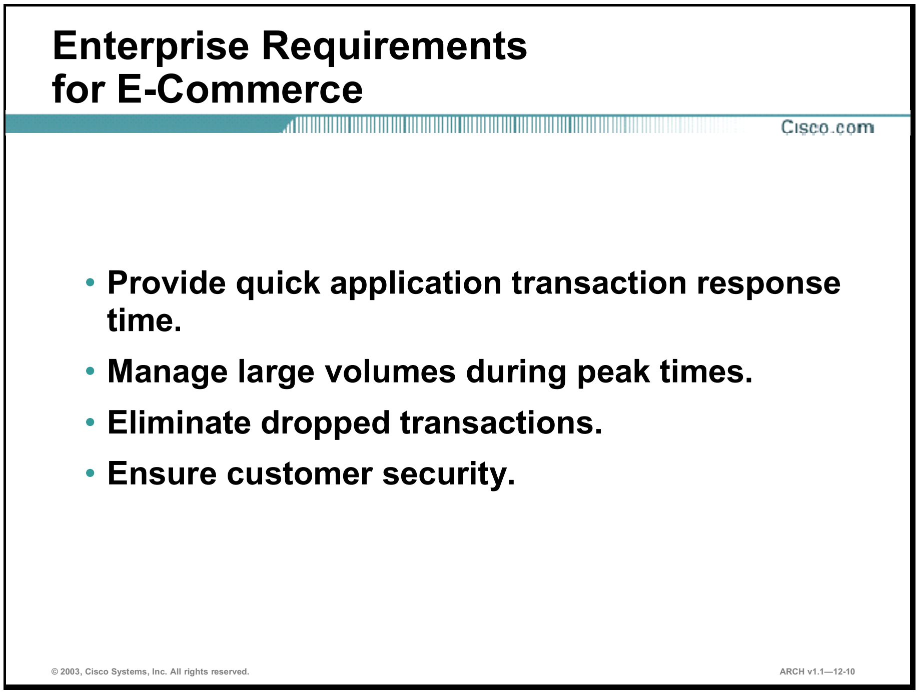

Enterprise Requirements for E-Commerce

Cisco.com

- Provide quick application transaction response time.
- Manage large volumes during peak times.
- Eliminate dropped transactions.
- Ensure customer security.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—12-10

Businesses using e-commerce need to deliver rapid transaction response time and manage peak-period volume levels, whether from seasonal increases in traffic or from unexpected surges in customer demand. Customers will return to e-commerce sites that offer consistently high levels of reliability, and they will avoid sites that deliver slow response times, difficult shopping experiences, or failed attempts to make purchases. A key component in ensuring that e-commerce sites remain open for business requires support for persistent, "sticky" network connections between customers and e-commerce servers, so shopping carts are not lost before purchase transactions are completed.

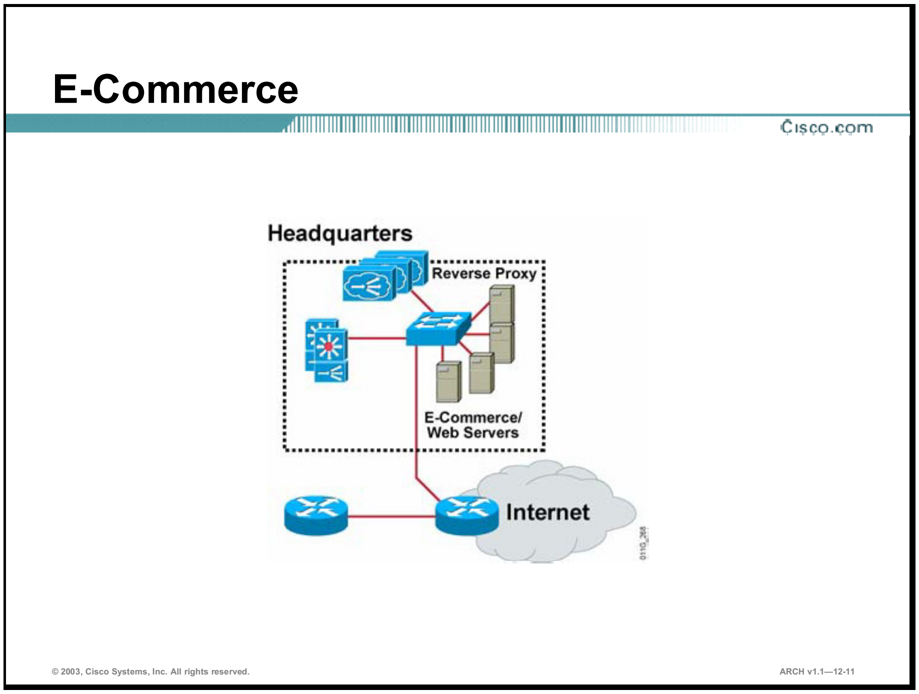

E-Commerce

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-11

The content networking solution for e-commerce shown in the figure includes these components:

- **Caching:** A cluster of CEs is placed in the perimeter LAN containing the outbound-facing web and e-commerce servers. These CEs are deployed in reverse proxy mode to offload content from the servers.
- **Content switching:** Redundant CSSs are used to load balance and direct traffic to the appropriate CE.

When users are connected to a server, the CSS ensures that they stay connected to a single server for the duration of their transaction using “cookies” embedded in the user’s request. The only reliable way to maintain sticky connections is to use cookies to identify individual customers. CSS switches can read the entire cookie to identify the user and route the request to the correct server.

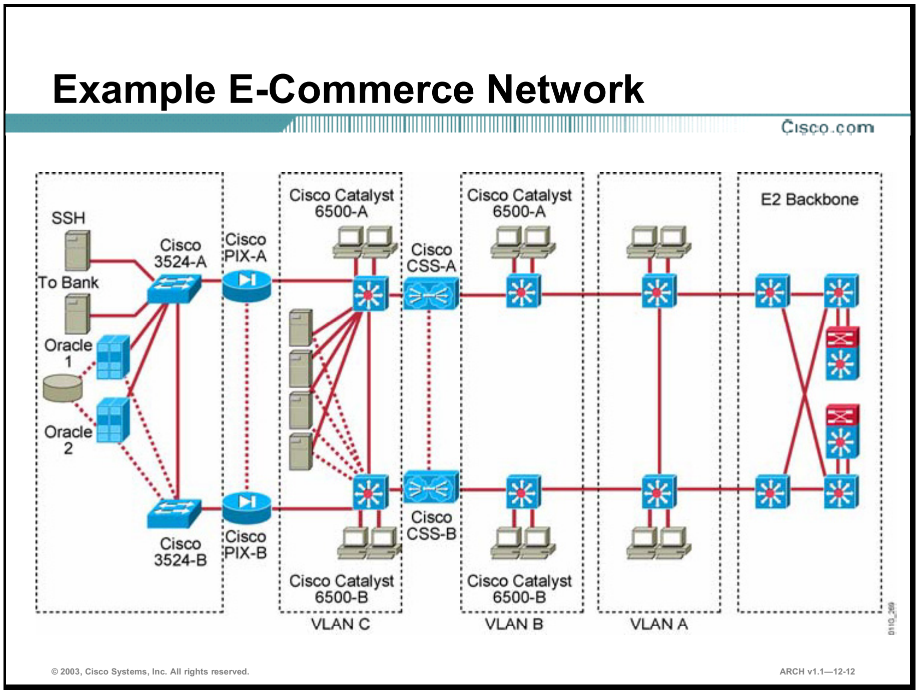

Example E-Commerce Network

ARCH v1.1—12-12

Company Background

In the United Kingdom, a one-day event was staged to raise money for charities around the world. The goal was to end poverty and social injustice. The group sponsoring the event wanted to build a reliable e-commerce web site that would enable them to effectively collect donations online during a TV broadcast. The web site had to provide 100 percent availability and the ability to handle extremely heavy traffic during the event. Requirements for the web were:

- Provide 100 percent availability with a response time of less than 7 seconds
- Process up to 200 credit card transactions per second
- Handle forecasted peaks in traffic (440 hits per second, 75 kb average downloaded page size)

Content Networking Solution

The group deployed Oracle database servers and web server software, along with real-time authentication software. The company implemented Cisco content services switches that could process up to 200 credit card transactions per second, and Cisco Catalyst switches.

The Cisco content services switches were used to front-end 19 web servers. The group used the intelligent keepalive capabilities of the Cisco CSS switches to provide 100 percent availability.

Content Networking Solutions for Streaming Media

A content networking solution for streaming media typically incorporates the entire range of content networking devices. The content distribution manager is the central management device for both media-on-demand and broadcast streaming media. This topic describes the design of content networking solutions to support streaming media.

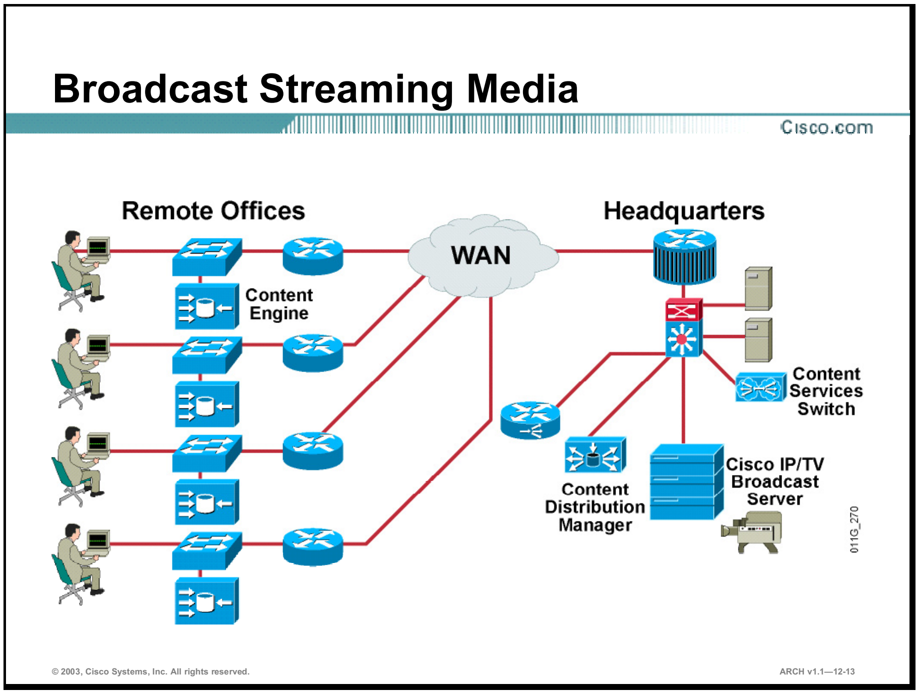

Broadcast Streaming Media

ARCH v1.1—12-13

Applications such as e-learning and corporate communications frequently involve broadcasting media streams that are to be accessed simultaneously by multiple users throughout the network. The content networking solution shown in the figure includes these components:

- **CDM:** The CDM enables you to configure bandwidth and distribution settings such that the streaming content will not interfere with other network traffic. It is also the central control point where the CEs that will carry the broadcast media are identified. The CDM is typically located in the server farm.
- **CE:** The CEs stream live or on-demand content to the desktop. A user logs on to a web page or application to access the high-bandwidth media over the LAN.
- **Event capture and delivery:** The Cisco IP/TV Broadcast Server and Control Server can capture and deliver live events.
- **IP multicast:** Multicast technology makes it possible for organizations to deliver live broadcasts to desktops and meeting rooms. IP multicast must be configured throughout the network to enable broadcast.

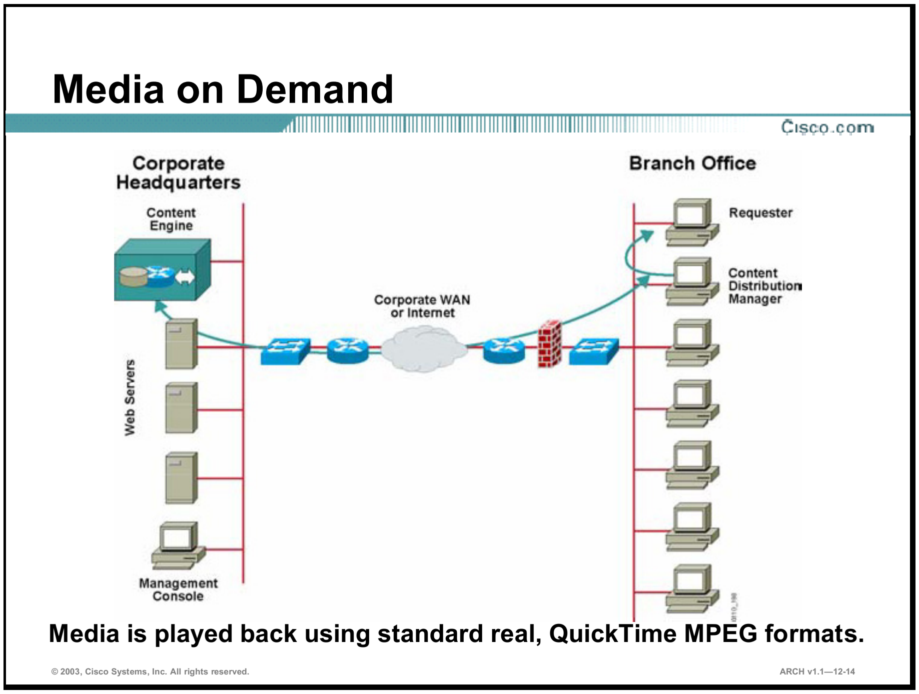

Media on Demand

Media is played back using standard real, QuickTime MPEG formats.

ARCH v1.1—12-14

A media-on-demand solution typically involves the distribution of large video and/or audio files. The content networking solution shown in the figure includes these components:

- **CDM:** Media (content) is imported into the CDM and then replicated to the CEs according to a schedule and bandwidth constraints set by the administrator.
- **CEs:** The CEs serve the content on receiving a request from a user.

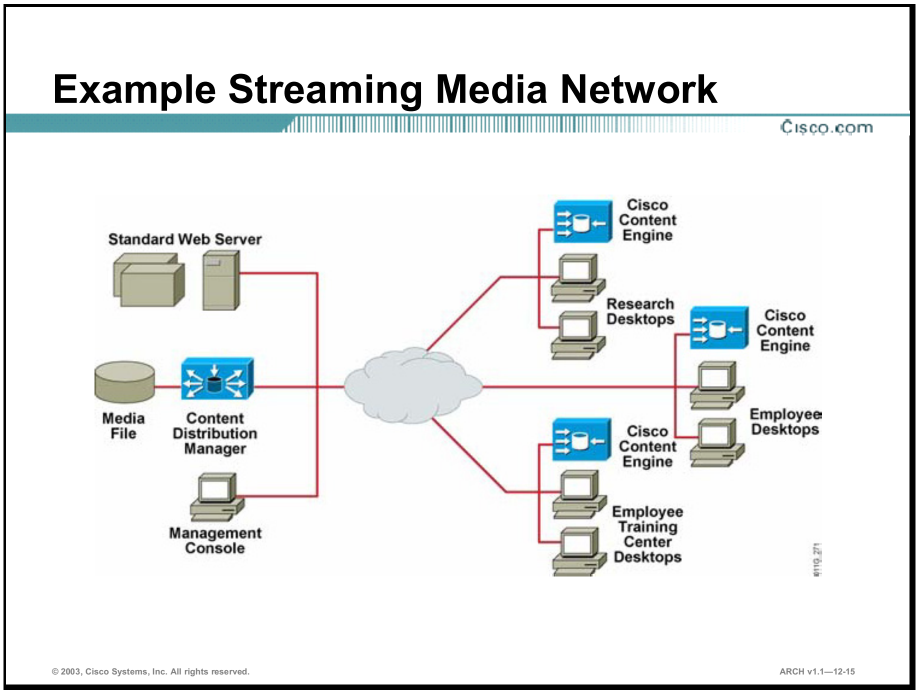

Example Streaming Media Network

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—12-15

Company Background

A major financial institution wanted to deliver timely corporate information, including corporate messages and company initiatives, through a sophisticated video content delivery system. They wanted a network-based system that would enable it to distribute high-quality video worldwide over IP versus videotape. Core requirements for a system included the ability to stage content, fully control content distribution worldwide, and limit the need for administration at remote sites.

Content Networking Solution

The company deployed an IP-based television service for employees worldwide. Components of the content delivery network included Cisco Content Distribution Manager, Cisco Content Engine, Cisco CSS 11000 Series Content Services Switch, and Cisco Content Router. The CEs are deployed in major offices throughout the world, while the Cisco Content Distribution Manager is located at the main office in Europe. The Cisco CSS 11000 Series Content Services Switches provide load balancing between the Content Engines and the CDM. The Content Engine Router is deployed for redundancy to provide redirection, speed, and performance. All video clips are stored on the company intranet and, when activated, are streamed via unicast.

Summary

This topic summarizes the key points discussed in this lesson.

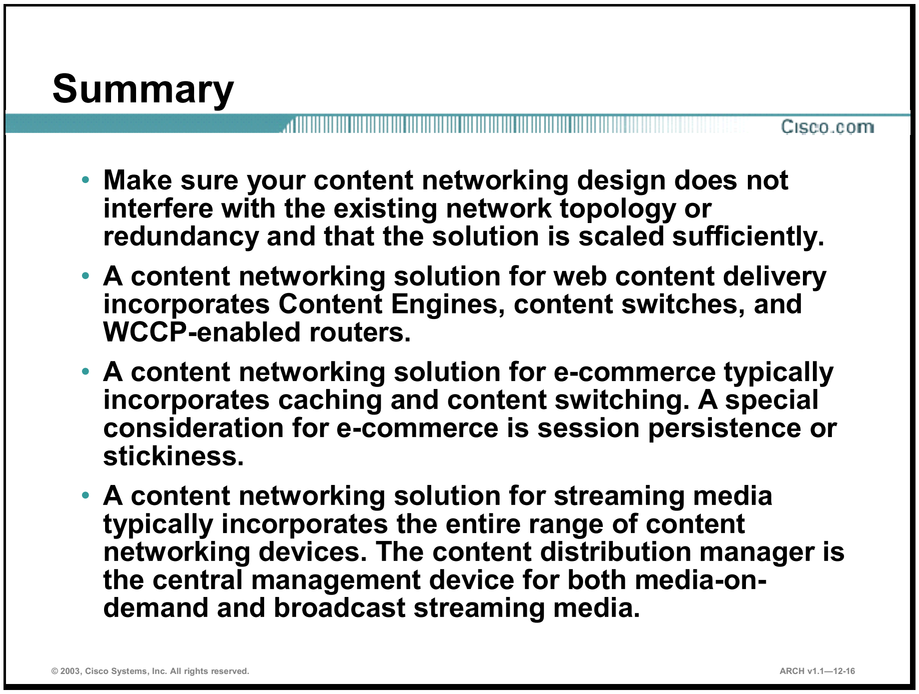

Summary

- Make sure your content networking design does not interfere with the existing network topology or redundancy and that the solution is scaled sufficiently.
- A content networking solution for web content delivery incorporates Content Engines, content switches, and WCCP-enabled routers.
- A content networking solution for e-commerce typically incorporates caching and content switching. A special consideration for e-commerce is session persistence or stickiness.
- A content networking solution for streaming media typically incorporates the entire range of content networking devices. The content distribution manager is the central management device for both media-on-demand and broadcast streaming media.

ARCH v1.1—12-16

References

For additional information, refer to these resources:

- *Accelerating Web Applications with Cisco Enterprise Content Delivery Networks* at http://www.cisco.com/warp/public/cc/pd/cxsr/ces/prodlit/awace_wp.htm
- *Technology Solutions: Content Networking* at http://www.cisco.com/offer/tdm_home/content_network/index.shtml
- *Network Caching* at http://www.cisco.com/warp/public/cc/pd/cxsr/500/tech/cds_wp.htm
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study, refer to the following section:

- Case Study 12-2: OCSIC Bottling Company

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 two factors influence the number of CEs required for a particular application? (Choose two.)

A) cacheable object size

B) number of web servers

C) existing bandwidth capabilities

D) number of transactions per second

E) distance between users and servers

Q2) Where in the network do the CEs configured in reverse proxy mode reside?

A) the core

B) the server farm

C) the WAN edge

D) the Internet edge

Q3) What feature provided by a content services switch is of particular importance in an e-commerce solution?

A) failover

B) stickiness

C) redundancy

D) load balancing

Q4) In an e-commerce content networking design, for which two reasons would you include redundant content service switches? (Choose two.)

A) to load balance

B) to read an entire cookie

C) to provide content caching

D) to direct traffic to the appropriate CE

E) to provide persistent, “sticky” network connections

Q5) What additional component is required for a content networking solution that will support one to many?

A) CEs

B) IP multicast

C) content routers

D) content switches

Quiz Answer Key

Q1) A, D

Relates to: Content Networking Design Considerations

Q2) B

Relates to: Content Networking Solutions for Web Content Delivery

Q3) B

Relates to: Content Networking Solutions for E-Commerce

Q4) A, D

Relates to: Content Networking Solutions for E-Commerce

Q5) B

Relates to: Content Networking Solutions for Streaming Media

Case Study 12-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

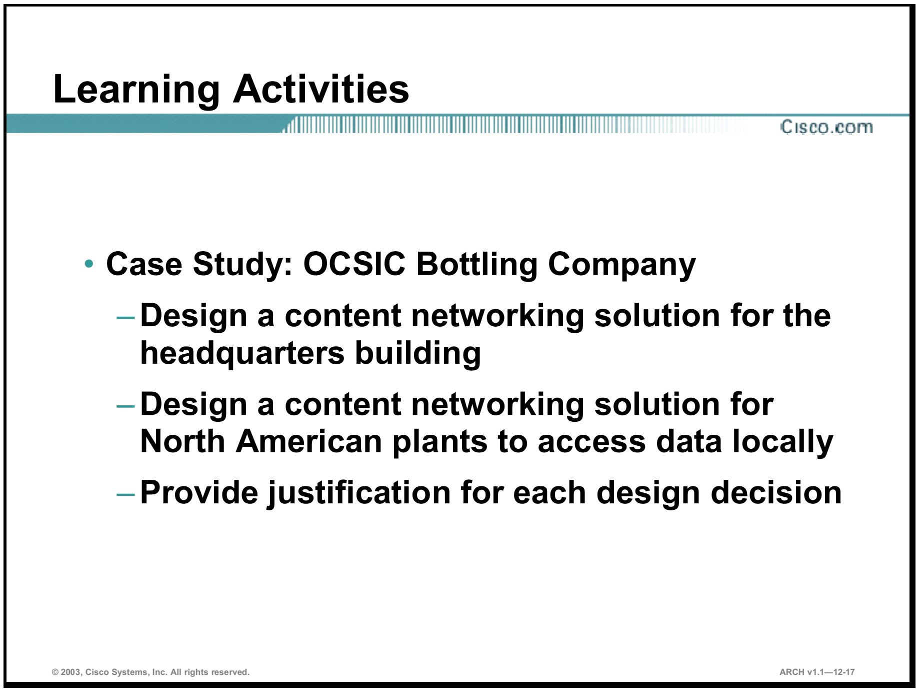

Learning Activities

- Case Study: OCSIC Bottling Company
  - Design a content networking solution for the headquarters building
  - Design a content networking solution for North American plants to access data locally
  - Provide justification for each design decision

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—12-17

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

OCSIC believes that content networking will help provide faster access to data and reduce the load on the network between buildings at the headquarters campus and between the headquarters' office and the North American plants.

In this exercise, you will design a content network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design a content networking solution for the headquarters building
- Design a content networking solution for North American plants to access data locally

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Design a Content Networking Solution for the Company Network

Complete these steps:

Step 1 Determine each location where users will need to access cached content. On an overhead transparency, create a network diagram for the headquarters and North American plants indicating the location of each content networking device including content switches, content routers, content managers, and content distribution managers.

Step 2 Complete the table to design the details about the content network.

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your content networking solution?		
What security strategy will you deploy to support your content networking solution?		
What QoS strategy will you deploy to support your content networking solution?		

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have created a content networking solution for the North American network. The solution specifies the locations of each content networking solution, content networking devices, high availability, security, and QoS.
- You have created a content networking solution for North American plants to access data locally. The solution specifies the locations of each storage solution, networking technologies, networking devices, high availability, security, and QoS.

Module 13

Designing Storage Networking Solutions

Overview

The Cisco storage networking solutions combine storage and networking technologies in an IP-enabled enterprise network. Storage networking allows interconnection, access, and sharing of stored data over IP, using a variety of media that includes Gigabit Ethernet, Fibre Channel, and optical networks. These technologies provide interconnection and access for network-attached storage (NAS) and storage area networking (SAN) environments.

Module Objectives

Upon completing this module, you will be able to design enterprise solutions for storage networking, given enterprise network needs.

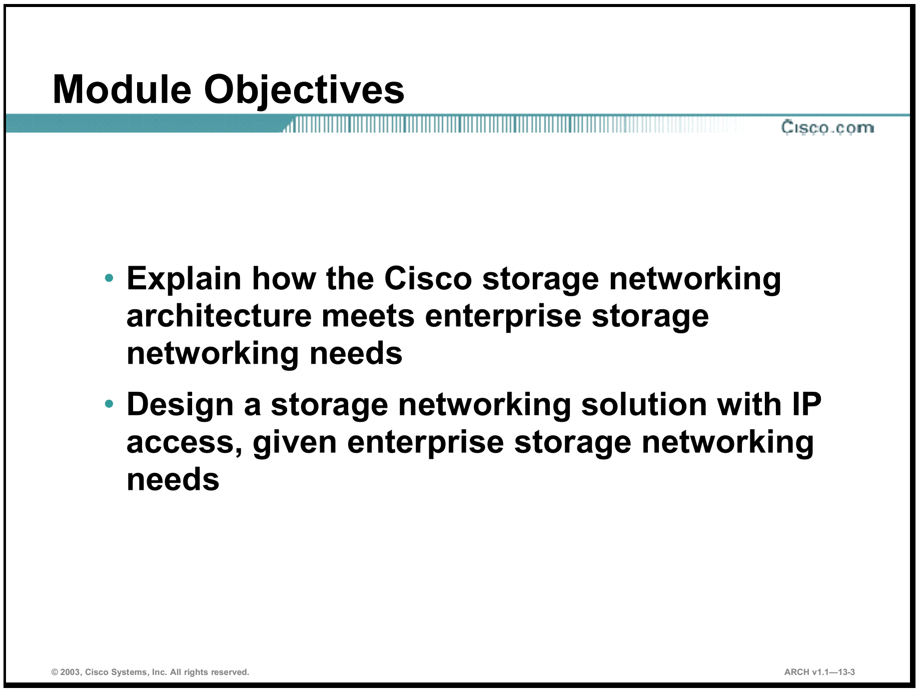

Module Objectives

- Explain how the Cisco storage networking architecture meets enterprise storage networking needs
- Design a storage networking solution with IP access, given enterprise storage networking needs

ARCH v1.1—13-3

Module Outline

The outline lists the components of this module.

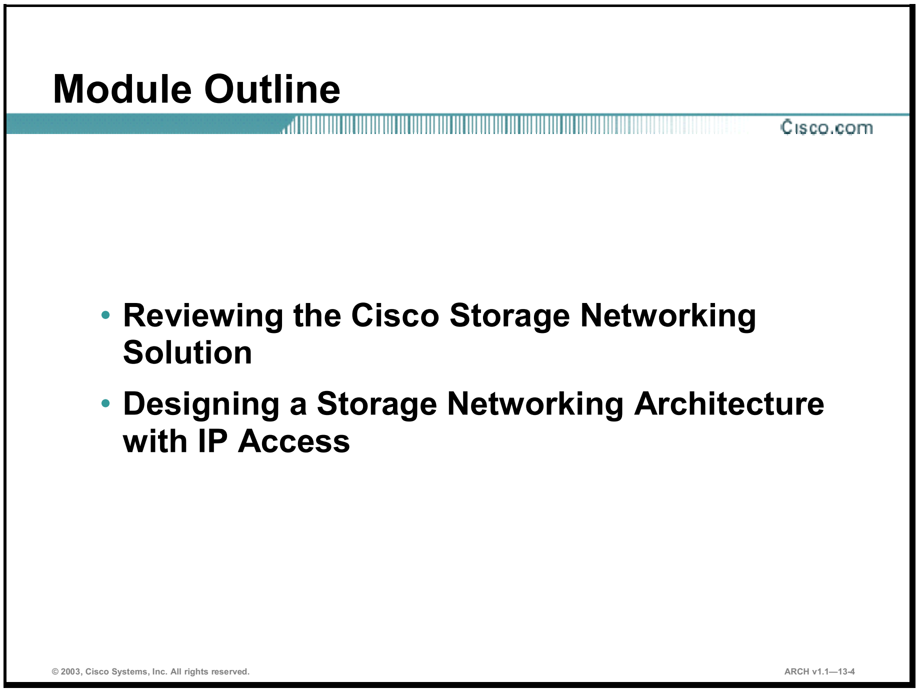

Module Outline

Cisco.com

- Reviewing the Cisco Storage Networking Solution
- Designing a Storage Networking Architecture with IP Access

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-4

Reviewing the Cisco Storage Networking Solution

Overview

Enterprise data storage is changing as Fibre Channel (FC) and IP networks converge toward an integrated storage networking infrastructure. This eliminates the limitations imposed by separate SAN islands. The convergence of storage and the network enables the Cisco Architecture for Voice, Video and Integrated Data (AVVID) storage networking solution to integrate the technology already deployed in IP networks with new standards, protocols, and products.

Relevance

The Cisco storage networking solution fully integrates with the Cisco AVVID network infrastructure and is based on industry standards.

Objectives

Upon completing this lesson, you will be able to explain how the Cisco storage networking architecture meets enterprise storage networking needs. This includes being able to meet these objectives:

- Identify enterprise requirements for storage networking
- Explain how the Cisco storage networking architecture meets enterprise storage networking needs
- Explain how the Cisco storage networking models utilize the underlying networking technology and intelligent network services
- Describe the important standards that make storage networking possible in an enterprise network
- Describe how to implement intelligent network services to support storage networking

Learner Skills and Knowledge

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

- Designing Enterprise Campus Networks module
- Designing Enterprise Edge Connectivity module

Outline

The outline lists the topics included in this lesson.

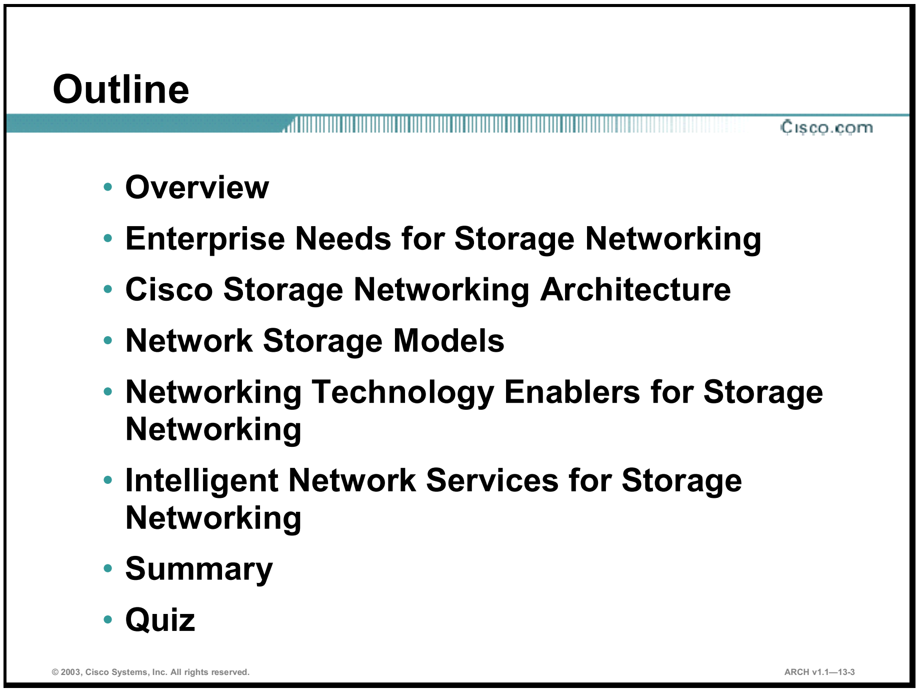

Outline

- Overview
- Enterprise Needs for Storage Networking
- Cisco Storage Networking Architecture
- Network Storage Models
- Networking Technology Enablers for Storage Networking
- Intelligent Network Services for Storage Networking
- Summary
- Quiz

ARCH v1.1—13-3

Enterprise Needs for Storage Networking

With storage networking, you can distribute storage for remote access through the network's IP infrastructure. Storage consolidation and business continuance are two common applications for storage networking. This topic identifies enterprise requirements for storage networking.

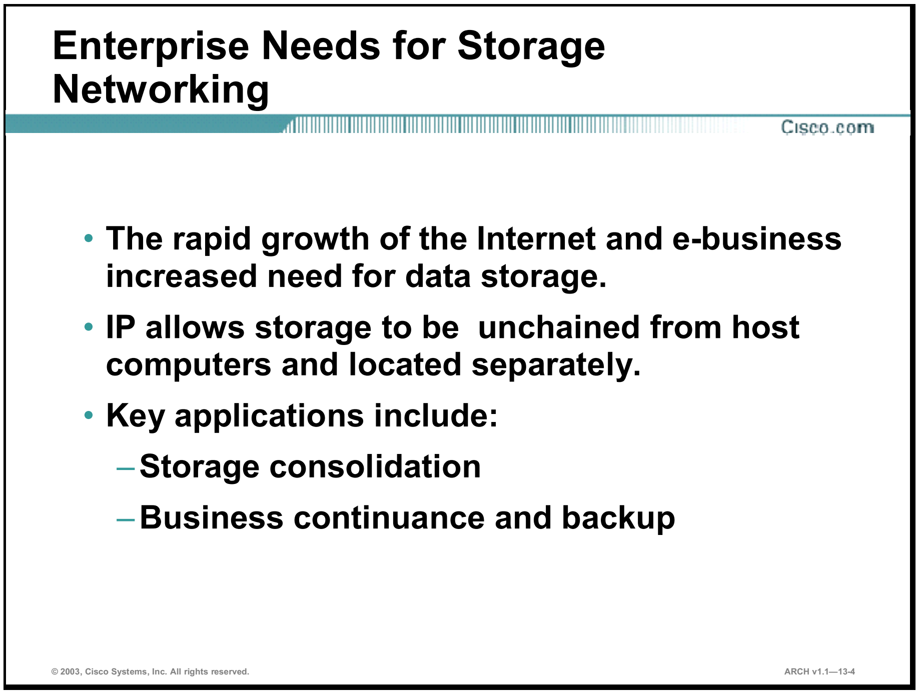

Enterprise Needs for Storage Networking

- The rapid growth of the Internet and e-business increased need for data storage.
- IP allows storage to be unchained from host computers and located separately.
- Key applications include:
  - Storage consolidation
  - Business continuance and backup

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-4

The rapid growth of the Internet and e-business has made data storage more important than ever. Business applications such as e-commerce, e-learning, supply chain management, customer care, and workforce optimization add to storage requirements. IP network technology is bringing some new capabilities to the storage world.

Instead of being tied to local host computers, storage device controllers are accessed directly by distant hosts using IP technology. This lowers costs and simplifies IT infrastructures, while allowing storage to be located remotely, subject to user and application storage access latency and data synchronization requirements.

Storage Consolidation

Enterprises that have already implemented SAN or NAS architectures want to leverage the existing infrastructure to consolidate storage. With the growth of digital information, the amount of data and servers has also increased. System administrators are faced with the challenging task of managing storage and making it scalable to accommodate future needs.

With storage directly attached to the server, scalability is difficult. The storage expansion capability is limited to the capacity of the server (for example, as measured by the number of I/O controllers and devices per controller configured in the server). The nature of the small computer system interface (SCSI) bus commonly used to connect commodity disks to a commodity server makes it difficult to allocate more disk storage without interrupting and rebooting the server, and thus affecting applications.

To accommodate growth, deployment of additional storage must be accomplished quickly and have minimum or no impact on the availability of applications or data. A pool of storage devices attached to the network creates rapid and simplified scalability. These storage devices provide server file access (NAS) or block access (SAN).

The addition of servers with directly attached storage resources to accommodate rapid growth results in a more difficult environment to manage and poor use of resources. A best-practice approach is to provide all servers that do not have local access to the SAN with IP access to the storage and allocate storage on demand. This storage consolidation provides centralization and simplification of storage management.

Business Continuance and Backup

Today's storage and networking architectures do not tolerate any interruption of operation, and enterprises are implementing stringent disaster recovery plans to guarantee the recovery of services in a timely and cost-effective manner. Enterprise-wide information is protected and archived according to its level of criticality, the length of time allowed to recover the information, and whether the potential loss of information is acceptable.

The highest level of information protection is achieved through mirror sites: a complete replicated storage, server, network, and application infrastructure in two or more locations. The information is synchronously replicated, in real time, between the mirror sites. The cost for this level of protection is significant.

A lower level of protection includes asynchronous mirroring: The information is replicated in an asynchronous manner at a selected frequency. In this asynchronous mode of replication, snapshots, or point-in-time copies of the data, are taken and transferred to a remote site. The point-in-time image is used for disaster recovery, and can potentially be used to perform tasks that do not require tight synchronization with the online applications (for example, data mining, reporting, backup).

Disaster recovery applications, such as mirroring and replication, protect against equipment or site failures to provide a highly available infrastructure. However, they do not protect against user errors or data corruption. Backup, with offsite storage and restore, provides protection against any kind of data loss for the retention period of the backup. Backup media are also used for long-term off-site archiving for protection of key information for future reference or audit purposes. With the growth of stored data the backup volume required has increased, yet the time frame allowed for backup is limited. Local backups impact the performance of the application servers, and management of backups has become increasingly more complex as the amount of information grows.

Backup has evolved from an architecture in which tape drives were directly attached to the servers. Network backup uses tape drives attached to a centralized backup server or to a SAN-based, networked tape library. When storage area networks first emerged, tape autoloader or tape libraries constituted a consolidated backup system. Snapshot technology makes it possible to back up systems online with minimum impact on the applications, by archiving only the changed data. Snapshot technology improves backup efficiency and reduces the impact of backup activities while increasing the restoration requirements.

Cisco Storage Networking Architecture

Storage networking is the hardware and software that enables you to consolidate, share, and access storage over a networked infrastructure. Cisco views the storage network as another component of the Cisco AVVID common infrastructure. This topic explains how the Cisco storage networking architecture meets enterprise storage networking needs.

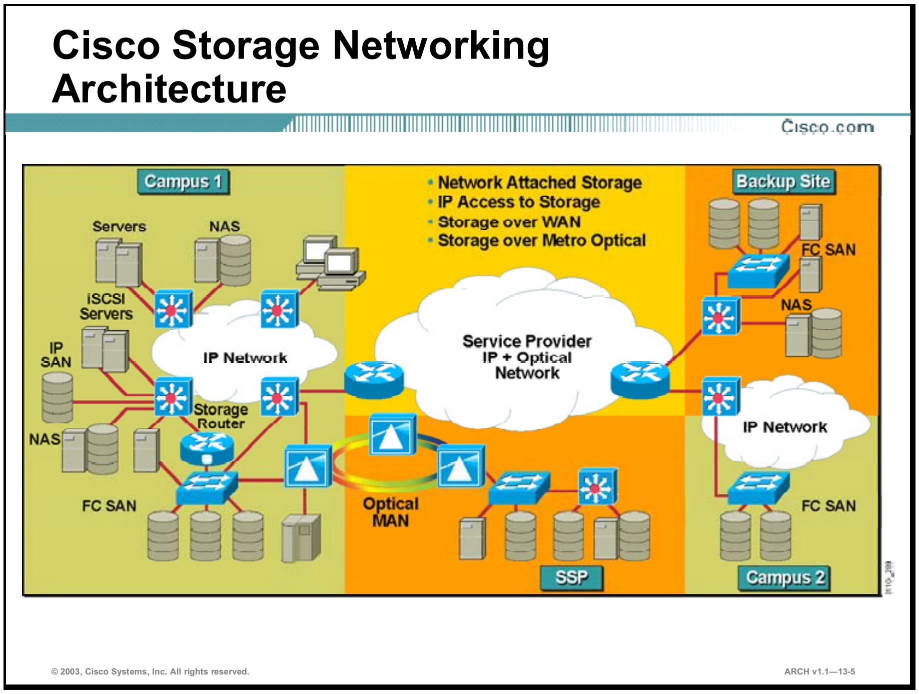

Cisco Storage Networking Architecture

Cisco.com

- Network Attached Storage
- IP Access to Storage
- Storage over WAN
- Storage over Metro Optical

ARCH v1.1—13-5

Key components of the Cisco storage networking solution are:

- Network-attached storage for file-oriented access to storage
- IP access to storage for block-oriented host-to-storage communication in storage area networks
- Storage over the WAN, for the interconnection of all storage environments
- Metro optical connectivity for the efficient, high-performance transport of storage traffic types, including Fibre Channel, Internet Small Computer System Interface (iSCSI), and Enterprise System Connection (ESCON), over data link protocols, which include 1/10 Gigabit Ethernet and other optical technologies

The Cisco AVVID common infrastructure supports both SAN and NAS models. They are complementary technologies that simultaneously use the Cisco AVVID common infrastructure:

- SAN provides block-oriented access to native disk storage. It is based on a shared or switched infrastructure, often Fibre Channel. You can extend SAN to an IP infrastructure. New protocols and products are emerging that allow the integration of SANs with the IP network. Historically, SANs have been well suited to high-volume, write-intensive, transaction-driven applications.
- NAS provides file-oriented access over an IP network. NAS is implemented using customized storage appliances that run Network File System (NFS) for UNIX environments and Common Internet File System (CIFS) for Microsoft Windows NT and Microsoft Windows environments. NAS is deployed for high-performance file sharing applications such as engineering collaboration, NT file systems, e-mail, and web content storage.

Most enterprises deploy a combination of NAS and SAN strategies to meet the wide range of application environments. Differentiation between these technologies will diminish as storage architectures converge to provide both file and block-based services.

Network Storage Models

The two network storage models are NAS and SANs. This topic describes the two storage networking models.

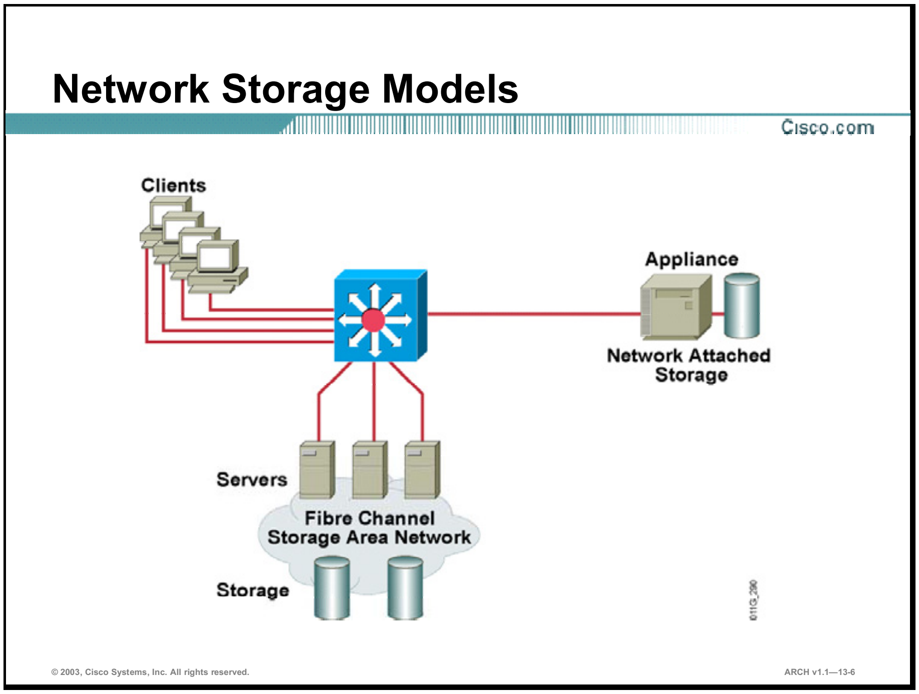

Network Storage Models

ARCH v1.1—13-6

A SAN describes a dedicated, high-performance network infrastructure deployed between servers and storage resources. The storage area infrastructure is a separate, dedicated network entity optimized for the efficient movement of a large amount of raw block data to the application servers, which coordinate client access to data. In effect, SAN is an extended link between server and storage. SANs enable the extension of the storage access protocol over longer distances.

SANs are typically built using the SCSI and Fibre Channel protocols (SCSI-FCP). Fibre Channel is well suited to this application because it can transfer large blocks of data (as is required with SCSI) while at the same time being able to transfer these blocks over longer distances (unlike SCSI media). Fibre Channel topologies, either loop or fabric, are built using specially designed devices that closely resemble the hubs, switches, and gateways used to build typical packet-based LANs and WANs. The SAN market has historically addressed high-end, enterprise-class storage applications where performance, redundancy, and availability are paramount in support of mission-critical, transaction-driven business systems.

LANs made it possible to connect multiple file servers with a common infrastructure for the purpose of file sharing. LANs accelerated the development of distributed multitier computing. The concept of distributed technology involves using an arrangement of inexpensive microcomputers and storage devices (disk, tape, and so forth) to reduce cost and move processing nearer the user. As computers proliferated, there were many incompatibilities, which complicated shared data access. The advent and widespread deployment of local-area networking encouraged workgroup clusters offering file sharing, interoperability, and cost savings. NAS consists of a specialized file server and storage, as shown in the figure. NAS servers run optimized file systems and are installed as preconfigured storage "appliances."

Because NAS systems are connected to the IP network, clients are able to transfer data to and from the storage devices associated with the NAS system.

In addition, NAS can directly process file system protocols such as NFS and CIFS. Client machines mount volumes on these disk resources, allowing virtual access to productivity applications and file data. While these devices are capable of hosting distributed applications, these applications typically are placed on application-specific server platforms that have no responsibilities for directly attached data storage.

Note SANs are typically used to access blocks of data within files for applications such as clustered database access. NAS is typically used to access files.

Networking Technology Enablers for Storage Networking

The core storage networking technology enablers are IP, Gigabit Ethernet, Fibre Channel, and optical networking, which provide access and interconnection for NAS and SAN environments. This topic describes important standards that make storage networking possible in an enterprise network.

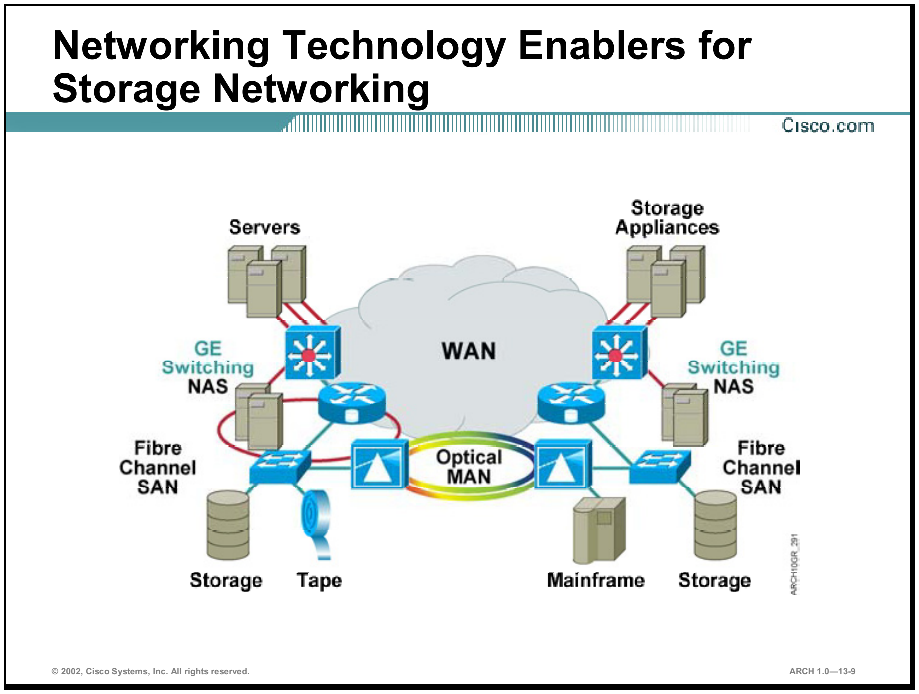

Networking Technology Enablers for Storage Networking

Cisco.com

© 2002, Cisco Systems, Inc. All rights reserved.

ARCH 1.0—13-9

The iSCSI protocol enables access to storage over TCP/IP, while Fibre Channel over IP (FCIP) links Fibre Channel SANs over IP.

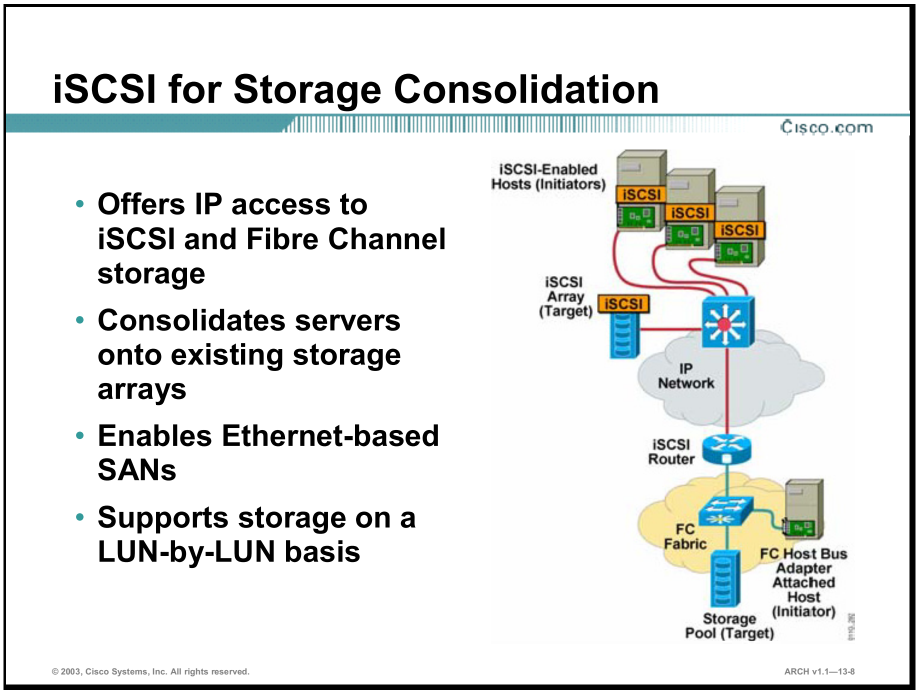

iSCSI for Storage Consolidation

Cisco.com

- Offers IP access to iSCSI and Fibre Channel storage
- Consolidates servers onto existing storage arrays
- Enables Ethernet-based SANs
- Supports storage on a LUN-by-LUN basis

ARCH v1.1—13-8

The iSCSI protocol encapsulates a SCSI request in an IP packet. iSCSI is a SCSI transport protocol for mapping block-oriented storage requests over TCP/IP networks. Making direct access to storage over IP possible, the iSCSI protocol allows IP-connected hosts to access iSCSI or Fibre Channel-connected storage.

iSCSI is fully compatible with existing IP infrastructures. With iSCSI, users can access storage across campus and wide-area networks, allowing data center storage to scale across the enterprise.

Relying on existing IP network infrastructures, local- and wide-area routers and switches transparently extend storage access across the WAN for applications such as remote disk copy and tape backup and restore. In the WAN environment, TCP/IP ensures data reliability, manages network congestion, and controls WAN retransmissions.

The figure shows how block accesses to storage are made through the interconnected IP and FC network. The storage router acts as a bridge between Fibre Channel and the IP network. One or more TCP sessions support the communication between the SCSI initiator and SCSI targets. The key technologies that enable the transfer of blocks of data over the IP network are:

- **iSCSI routers:** Enable connection of iSCSI hosts to Fibre Channel-connected storage. Along with the iSCSI router, iSCSI device drivers provide the interface between the operating system and the TCP/IP stack.
- **Optical media:** The bandwidth necessary for the timely transfer of I/O data is typically provided over optical media. Gigabit Ethernet is often used in this solution.

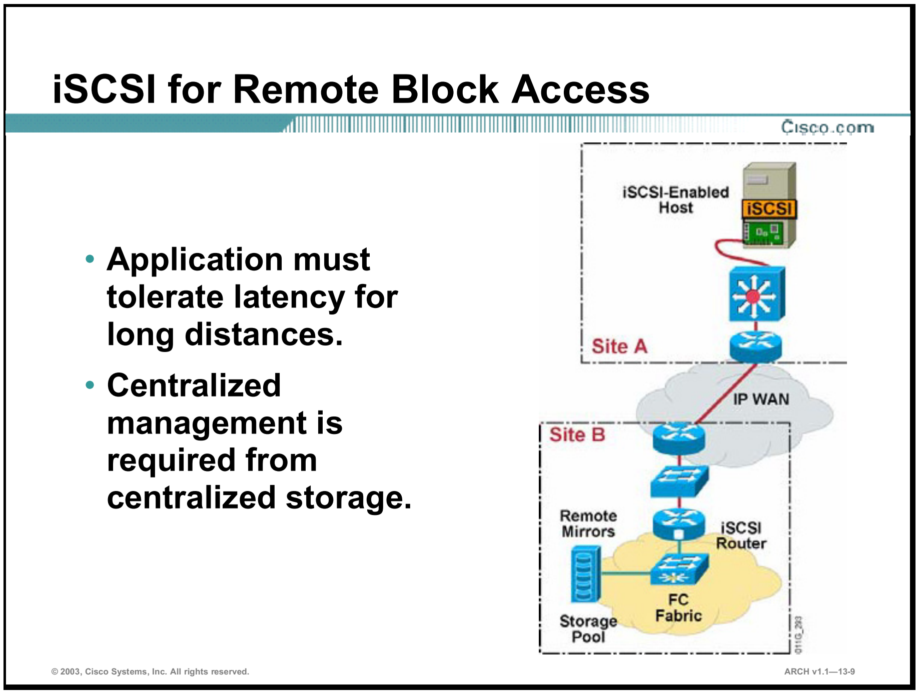

iSCSI for Remote Block Access

Cisco.com

- Application must tolerate latency for long distances.
- Centralized management is required from centralized storage.

ARCH v1.1—13-9

You can deploy iSCSI to provide:

- Remote storage over IP
- Remote backup of devices over IP to provide centralized storage management

An application must tolerate relatively high access latency to support a remote iSCSI solution over long distances. Metro Ethernet services offer a low-latency, high-volume transport alternative where available.

The figure shows how storage devices use iSCSI to provide remote access to servers through the high-speed network. Enterprises that already have a consolidated storage area network may want to preserve their investment in a centralized pool of storage and use an iSCSI router to access the SAN island. iSCSI technology with Gigabit Ethernet enables the connection of storage appliances or devices to the IP network.

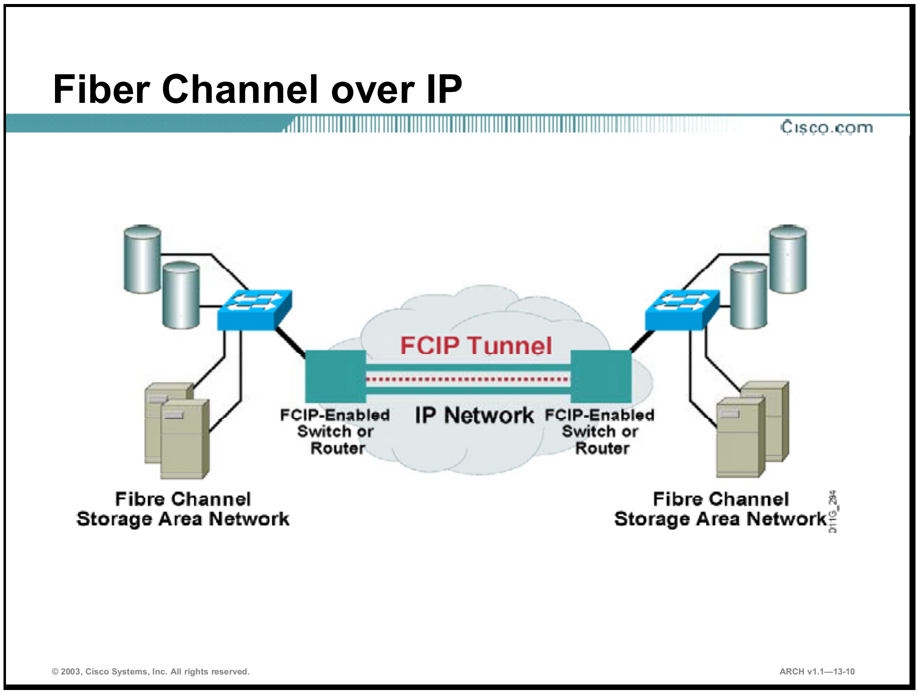

Fiber Channel over IP

Cisco.com

FCIP Tunnel
IP Network
FCIP-Enabled Switch or Router
FCIP-Enabled Switch or Router
Fibre Channel Storage Area Network
Fibre Channel Storage Area Network

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-10

An important technology for linking Fibre Channel SANs is FCIP. FCIP and iSCSI are complementary solutions for enabling company-wide access to storage. FCIP transparently interconnects Fibre Channel SAN islands over IP networks through FCIP tunnels, while iSCSI allows IP-connected hosts to access iSCSI or FC-connected storage.

iSCSI and FCIP are typically used for different purposes. With iSCSI, SCSI commands and data frames are encapsulated in IP to support disk access over an IP network. With FCIP, Fibre Channel frames are encapsulated in IP so that both SCSI and non-SCSI frames can be transported over an IP network.

FCIP overcomes many shortcomings of direct-attached storage by offering these features:

- Addressing for up to 16 million nodes (24 bits)
- Loop (shared) and fabric (switched) transport
- Speeds of 1000 or 2000 Mbps (1 or 2 Gbps)
- Distance of up to 10 km with extenders
- Support for multiple protocols

The combination of FCIP and iSCSI allows enterprises to:

- Interconnect SAN islands
- Provide applications including remote backup and replication, in addition to performing Fibre Channel I/O communication

Businesses expanding their storage infrastructures are faced with business continuance issues. Applications associated with disaster recovery and high availability can use FCIP as a solution for protecting their data. You can use FCIP to connect two geographically dispersed Fibre Channel storage arrays for the purpose of synchronous data storage. If the local storage array becomes unavailable, an application could utilize the FCIP link to access the data on the "hot backup" storage system at the remote site. It is also possible to implement remote tape backups to further protect customers' valuable information in the event of a disaster at the primary site.

FCIP differs from iSCSI as follows:

- iSCSI encapsulates SCSI commands and data in a TCP/IP packet. In this case, an IP-connected host running an iSCSI driver is accessing block-level data over an IP network.
- FCIP encapsulates Fibre Channel in IP packets. In this case, any Fibre Channel frame, SCSI/FCP or otherwise, is transported transparently in an IP packet. FC hosts and storage communicate on both sides of an FCIP link.

FCIP extends storage-area networks over the MAN and provides peer-to-peer connection between SANs. With FCIP, enterprises can use their existing MAN infrastructure. Remote data replication, or backup, is performed when use is low. Asynchronous data replication, or backup applications, require high bandwidth but are less sensitive to latency. A response time of tens of milliseconds is acceptable for backup or asynchronous replication applications, unlike interactive applications, which require the least possible storage access latency to meet overall delay bounds, including processing and multiple storage accesses.

Intelligent Network Services for Storage Networking

Standard IP services, including network management, high availability, security, and QoS, are used to support storage networking. This topic describes how to implement services to support storage networking.

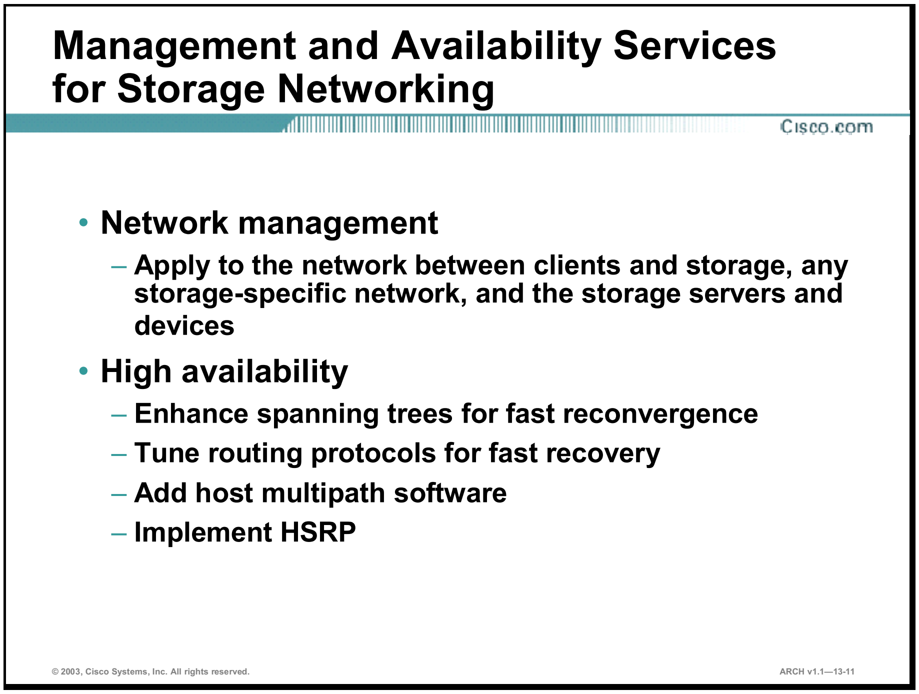

Management and Availability Services for Storage Networking

- Network management
  - Apply to the network between clients and storage, any storage-specific network, and the storage servers and devices
- High availability
  - Enhance spanning trees for fast reconvergence
  - Tune routing protocols for fast recovery
  - Add host multipath software
  - Implement HSRP

ARCH v1.1—13-11

Storage area networking relies on network services to deliver performance, scalability, and availability. With IP-based storage networking, the existing IP services are used for storage networking purposes. Services supported by storage networking include:

- **Network management:** Apply network management principles to the network between clients and storage, any storage-specific network, and the storage servers and devices.
- **High availability:** For storage networking, availability applies not only to the network between clients and storage, but also to the storage router and the storage devices themselves. Consider implementing these features:
  - **Spanning-tree enhancements:** Use 802.1w for fast spanning-tree reconvergence in Ethernet environments. Use UplinkFast and PortFast on Catalyst switches for fast recovery from a topology change.
  - **Routing protcols:** Tune Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) routing protocols for very fast convergence following a network link or router failure.
  - **Host multipath support:** Use multipath software above the host-based iSCSI layer to add path resiliency to an IP SAN.
  - **Hot Standby Router Protocol (HSRP):** Use HSRP to create redundancy default gateways for iSCI initiator hosts to ensure fast recovery if a gateway fails.

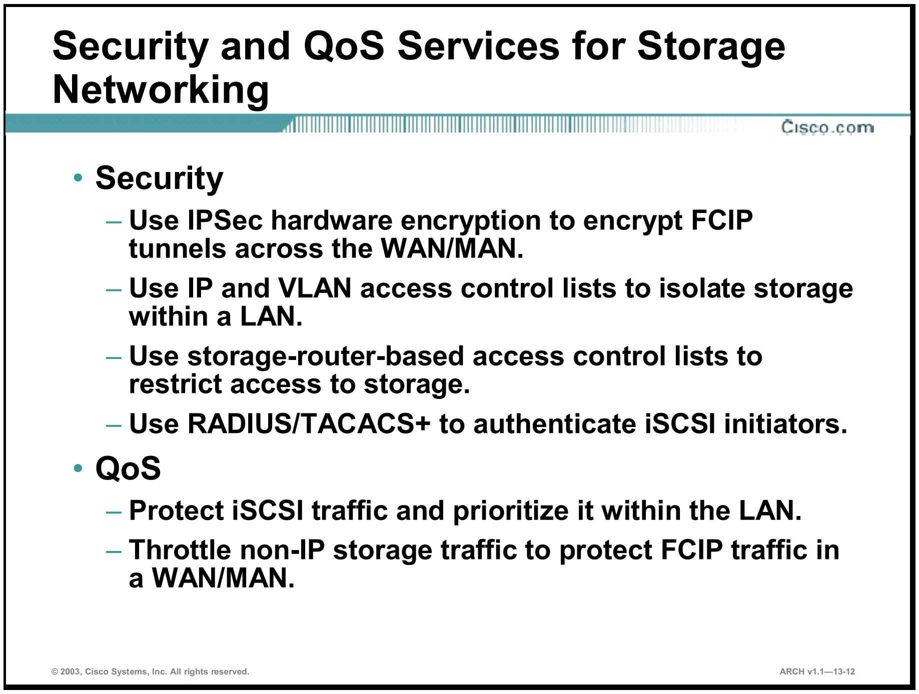

Security and QoS Services for Storage Networking

Cisco.com

- Security
  - Use IPSec hardware encryption to encrypt FCIP tunnels across the WAN/MAN.
  - Use IP and VLAN access control lists to isolate storage within a LAN.
  - Use storage-router-based access control lists to restrict access to storage.
  - Use RADIUS/TACACS+ to authenticate iSCSI initiators.
- QoS
  - Protect iSCSI traffic and prioritize it within the LAN.
  - Throttle non-IP storage traffic to protect FCIP traffic in a WAN/MAN.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—13-12

Two important considerations regarding storage networking are:

- **Security**: You can use IP Security (IPSec) hardware encryption to encrypt FCIP tunnels across the WAN/MAN. The iSCSI standard calls for IPSec support and requires hardware acceleration from clients. Use VLANs to isolate storage traffic within a LAN, or consider using private VLANs. You can use IP and VLAN access control lists to isolate storage within a LAN and storage-router-based access control lists to restrict access to storage. Implement RADIUS/TACACS+ to provide authentication for iSCSI initiators. You can use firewalls to prevent attacks due to static TCP ports.
- **QoS**: Use QoS to protect iSCSI traffic and prioritize it within the LAN for higher-priority queuing and switching. Use QoS to throttle non-IP storage traffic to protect FCIP traffic in a WAN/MAN.

The best QoS solution for storage networking is a separate network. Physical separation with adequate bandwidth provides an absolute QoS guarantee. A dedicated network may be expensive to provision and increase management complexity, but is more predictable.

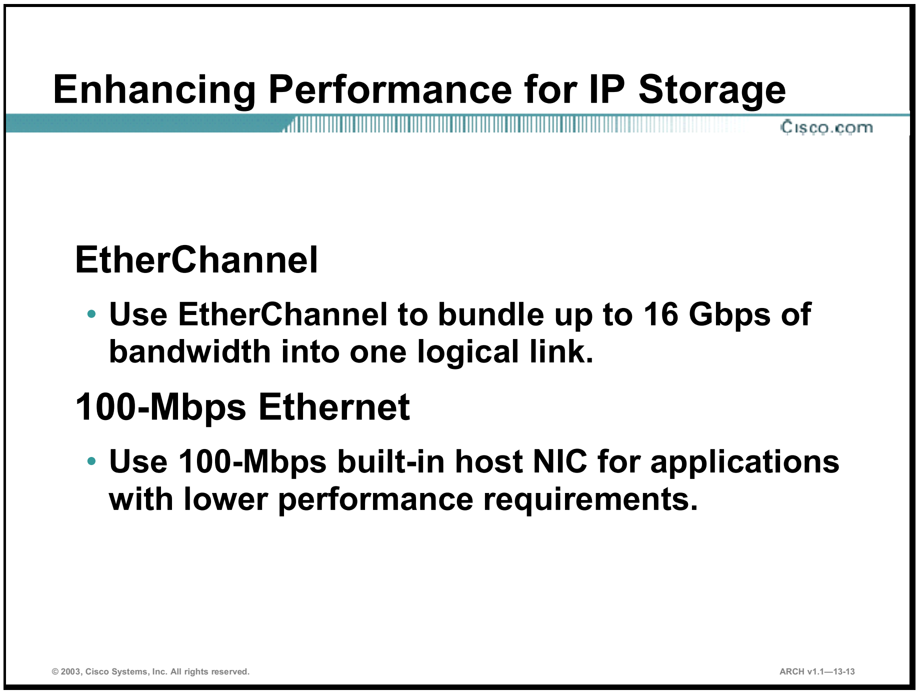

Enhancing Performance for IP Storage

Cisco.com

EtherChannel

- Use EtherChannel to bundle up to 16 Gbps of bandwidth into one logical link.

100-Mbps Ethernet

- Use 100-Mbps built-in host NIC for applications with lower performance requirements.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—13-13

Many services exist within IP to ensure performance in an IP SAN as well as to protect IP storage traffic from potential bottlenecks. You can implement these services to enhance performance:

- **EtherChannel:** Use EtherChannel to bundle up to 16 Gbps of bandwidth into one logical link within the LAN.
- **100-Mbps Ethernet:** Potentially use 100-Mbps built-in host network interface cards (NICs) for applications with lower performance requirements.

Summary

This topic summarizes the key points discussed in this lesson.

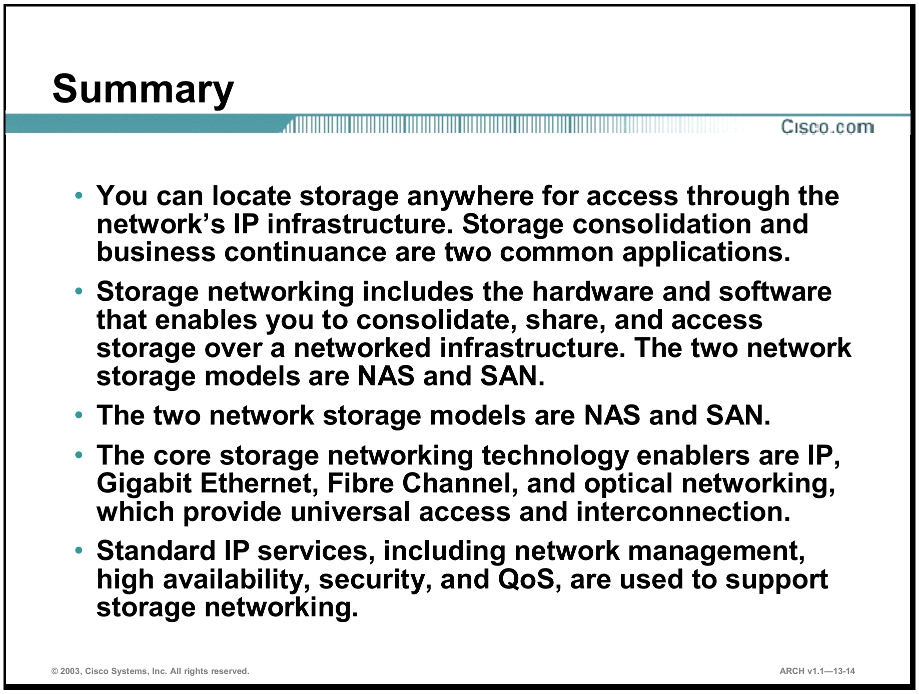

Summary

Cisco.com

- You can locate storage anywhere for access through the network's IP infrastructure. Storage consolidation and business continuance are two common applications.
- Storage networking includes the hardware and software that enables you to consolidate, share, and access storage over a networked infrastructure. The two network storage models are NAS and SAN.
- The two network storage models are NAS and SAN.
- The core storage networking technology enablers are IP, Gigabit Ethernet, Fibre Channel, and optical networking, which provide universal access and interconnection.
- Standard IP services, including network management, high availability, security, and QoS, are used to support storage networking.

ARCH v1.1—13-14

References

For additional information, refer to these resources:

- *Storage Networking* at http://www.cisco.com/warp/public/779/largeent/learn/technologies/storage/apps.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - Go to: http://www.cisco.com/.
  - In the Search box, enter "SRND" and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - Select the Networking Solutions Design Guide that meets your needs.

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) What are two key purposes for storage networking solutions? (Choose two.)

A) e-learning

B) high availability

C) video broadcasting

D) business continuance

E) storage consolidation

Q2) What are two advantages of storage networking over directly attached storage? (Choose two.)

A) improved scalability

B) simplified management

C) increased storage capacity

D) decreased bandwidth utilization

E) improved application response time

Q3) Which Cisco storage networking solution is well suited to high-performance, write-intensive database applications?

A) universal IP access

B) storage-area networking

C) network-attached storage

D) metro optical connectivity

Q4) Which component of the Cisco storage networking solution provides only file-oriented access to storage?

A) IP access

B) storage area networking

C) network-attached storage

D) metro optical connectivity

Q5) Match each storage networking technology with the benefit it provides for storage networking.

_____ 1. iSCSI

_____ 2. Gigabit Ethernet

_____ 3. Fibre Channel over IP

A) allows IP-connected hosts to access iSCSI or FC-connected storage

B) provides the bandwidth necessary for the transfer of I/O intensive data

C) transparently interconnects Fibre Channel SAN islands over IP networks

Q6) What protocol is typically used to build storage area networks?

A) IP

B) NFS

C) Ethernet

D) Fibre Channel

Q7) What information is encapsulated in an IP packet with FCIP?

A) data blocks

B) any FC frame

C) SCSI commands

D) SCSI data frames

Q8) What security feature does iSCSI require for encryption?

A) IPSec

B) firewalls

C) access control lists

D) RADIUS authentication

Quiz Answer Key

Q1) D, E

Relates to: Enterprise Needs for Storage Networking

Q2) A, B

Relates to: Enterprise Needs for Storage Networking

Q3) B

Relates to: Cisco Storage Networking Architecture

Q4) C

Relates to: Cisco Storage Networking Architecture

Q5) 1-A, 2-B, 3-C

Relates to: Cisco Storage Networking Architecture

Q6) D

Relates to: Network Storage Models

Q7) B

Relates to: Networking Technology Enablers for Storage Networking

Q8) A

Relates to: Intelligent Network Services for Storage Networking

Designing a Storage Networking Architecture with IP Access

Overview

Block storage access has been associated with Fibre Channel SANs, and file access with NAS. SANs and NAS are converging, as Fibre Channel and IP networks enable an integrated storage networking architecture. Storage networks are an infrastructure that enables file access and block access over interconnected Fibre Channel and IP networks. Cisco supports the development of new protocols that allow access to Fibre Channel SANs through enterprise IP networks.

Relevance

The limitations of locally connected storage solutions have led to the development of new network storage technologies, including serial SCSI and iSCSI to provide a scalable data storage infrastructure.

Objectives

Upon completing this lesson, you will be able to design a storage networking solution with IP access, given enterprise storage networking needs. This includes being able to meet these objectives:

- Describe a methodology to design a storage networking architecture
- Design an IP access storage networking solution using Cisco products and features, given enterprise storage networking needs
- Design a storage over WAN solution using Cisco products and features, given enterprise storage network needs
- Design a network-attached storage networking solution using Cisco components, given enterprise storage networking needs

Learner Skills and Knowledge

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

- Designing Enterprise Campus Networks module
- Designing Enterprise Edge Connectivity module

Outline

The outline lists the topics included in this lesson.

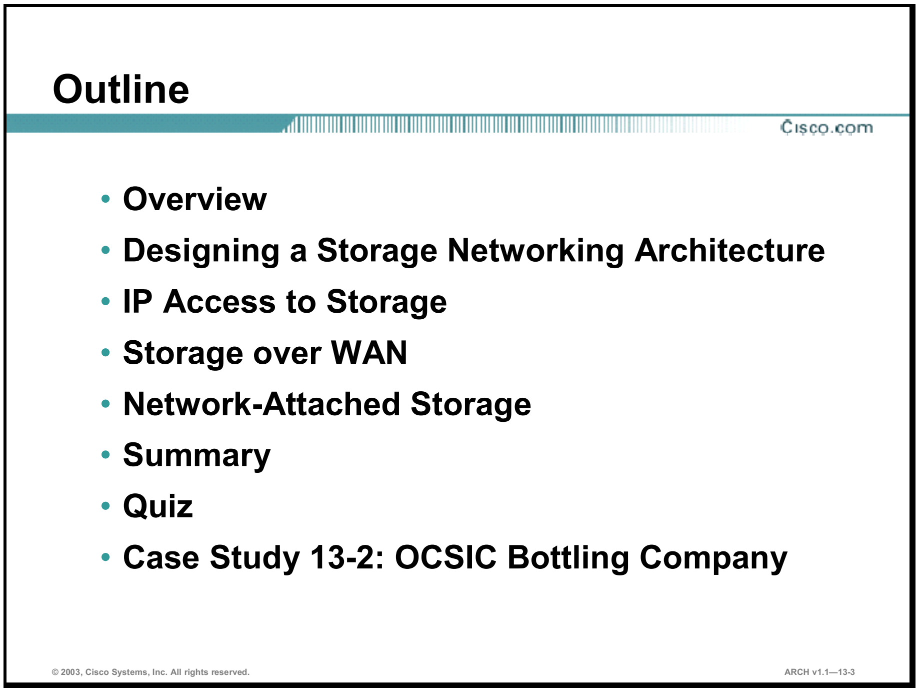

Outline

- Overview
- Designing a Storage Networking Architecture
- IP Access to Storage
- Storage over WAN
- Network-Attached Storage
- Summary
- Quiz
- Case Study 13-2: OCSIC Bottling Company

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—13-3

Designing a Storage Networking Architecture

The architecture of a storage networking solution depends on the applications accessing the storage. To determine the design requirements you must know the I/O profiles, throughput, and availability needs of the applications. This topic describes a methodology that a network designer may use to design a storage networking architecture.

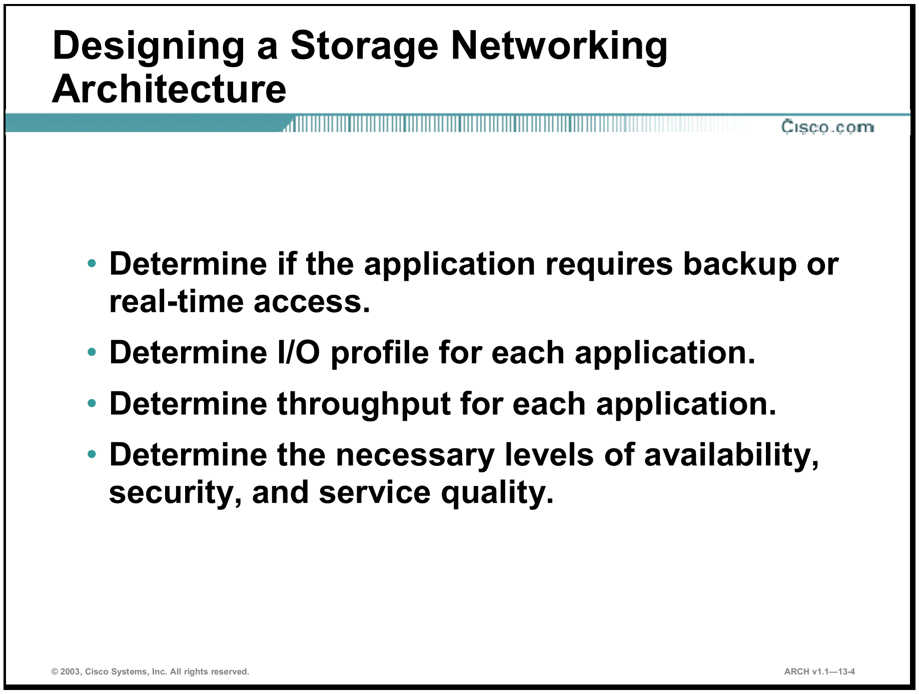

Designing a Storage Networking Architecture

Cisco.com

- Determine if the application requires backup or real-time access.
- Determine I/O profile for each application.
- Determine throughput for each application.
- Determine the necessary levels of availability, security, and service quality.

© 2003, Cisco Systems, Inc. All rights reserved. ARCH v1.1—13-4

To plan for a storage networking implementation, you must identify the needs of the application accessing the storage, considering these factors:

- **Access pattern:** The type of access required has a major impact on the storage networking architecture. A storage-area network is typically used for real-time application access, while network-attached storage is typically used for individual access and backup.
- **I/O profile:** How many bytes are being read or written per second? How much bandwidth will the data transfer related to storage access use? What's the most important aspect of the data I/O for the application: throughput or latency? The transaction rate of the application can be affected by the storage networking solution. Can the application tolerate latency for transfer between storage and the server? IP access to storage may introduce a latency factor that is unacceptable for some applications. Enterprises may require a Fibre Channel SAN in such cases. For latency-sensitive applications, the geographic placement of servers relative to storage devices can be an issue, as well as the bandwidth limitations between devices.
- **Throughput:** What are the sustained throughput requirements of the application? How will the requirements scale over time? The storage networking solution will need to address the infrastructure and device needs for the expected throughput.
- **Availability:** How critical is the application data to the enterprise? What is the effect on the enterprise of the application becoming unavailable? What high-availability measures should you consider for the storage network and devices? Mirrored servers and redundant network infrastructure may be required for some applications.

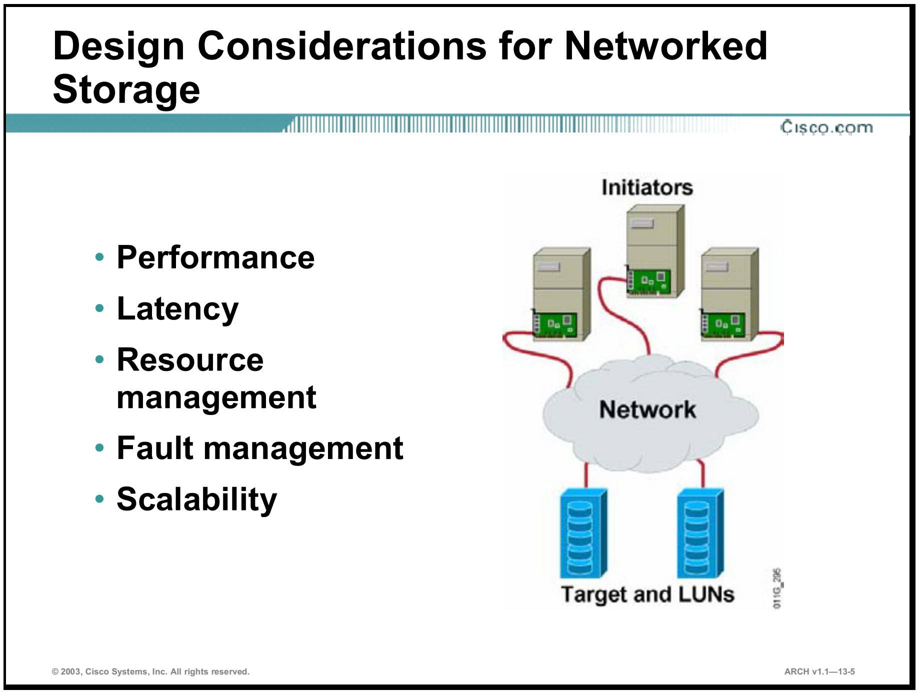

Design Considerations for Networked Storage

- Performance
- Latency
- Resource management
- Fault management
- Scalability

ARCH v1.1—13-5

The figure lists several considerations for storage networking:

- **Performance:** Ensures that the networked channel has adequate bandwidth and flow control; the performance requirements vary depending on whether the application is for real-time or archival requirements
- **Latency:** Ensures that the channel does not experience sufficient latency to compromise application integrity
- **Resource management:** Ensures that network resources are monitored
- **Fault management:** Ensures that the proper tools exist to detect, evaluate, and act on faults in the network
- **Scalability:** Ensures that the network can be expanded without jeopardizing network stability

IP Access to Storage

IP access to storage provides access to block-oriented storage over IP networks. This topic describes the design of an IP access storage networking solution.

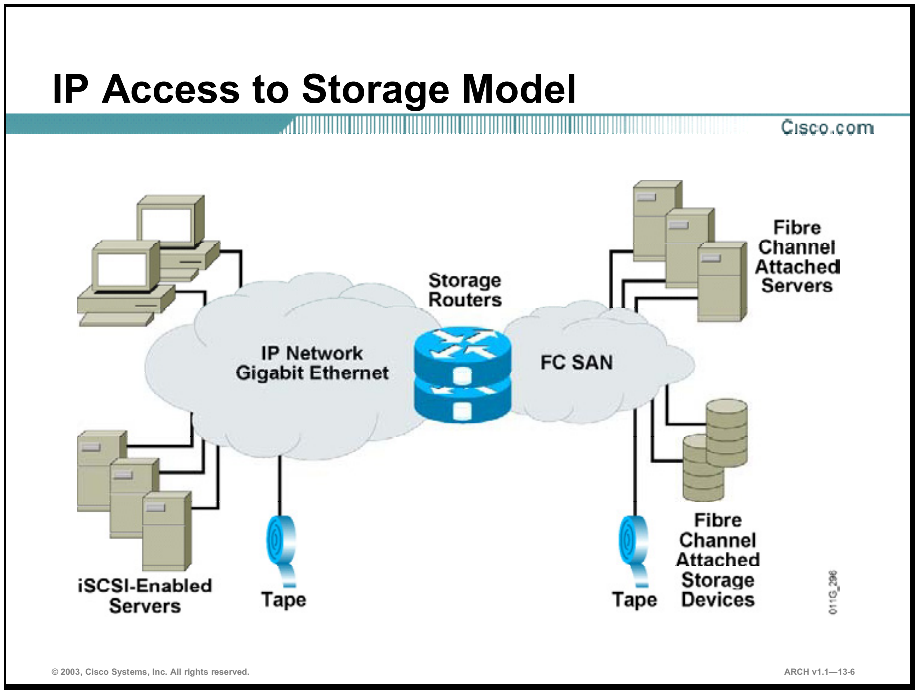

IP Access to Storage Model

Cisco.com

Storage Routers
IP Network Gigabit Ethernet
FC SAN
Fibre Channel Attached Servers
iSCSI-Enabled Servers
Tape
Tape
Fibre Channel Attached Storage Devices
011G_296

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-6

Cisco extends storage consolidation beyond the data center by enabling block-oriented storage I/O over an IP network. This allows servers that may not have otherwise been eligible to connect to shared storage resources. This could include Microsoft Windows server farms, remote servers, and small server clusters in remote offices.

iSCSI encapsulates the SCSI command set and data frame into TCP/IP, and thereby allows hosts to communicate with storage over a high-speed IP infrastructure, transparently to the application.

You can implement iSCSI to access Fibre Channel-connected storage or to access storage devices that are natively connected to the IP network. The figure shows both IP access to Fibre Channel-connected and native iSCSI-connected storage from iSCSI-enabled servers.

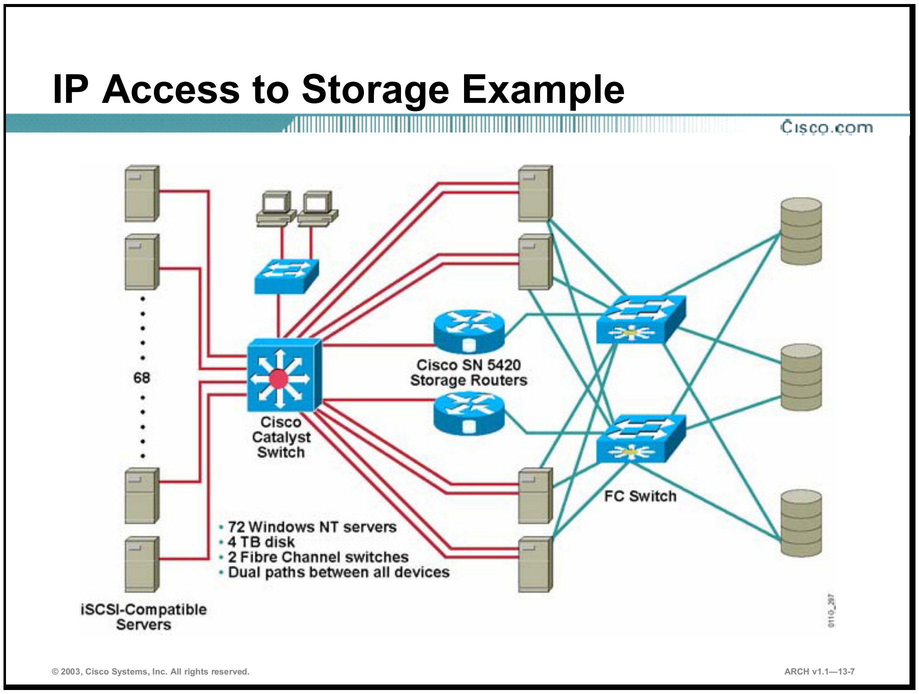

IP Access to Storage Example

Cisco.com

- 72 Windows NT servers
- 4 TB disk
- 2 Fibre Channel switches
- Dual paths between all devices

ARCH v1.1—13-7

Company Background

A major medical research facility's database was growing quickly, and users needed access. Hundreds of users in five separate buildings were networked into a powerful server farm made up of application servers that shared a very large, 4-terabyte Oracle database. The organization wanted to install a storage network to support the storage access requirements at the remote locations that would support their performance and cost needs, as well as grow with them in the future.

Storage Networking Solution

The researchers compared a Fibre Channel-only approach with a hybrid approach that coupled an iSCSI storage network with Fibre Channel for storage access. Both configuration alternatives provided redundant Gigabit Ethernet paths from every server to the facility's 4-terabyte database.

The Fibre Channel-only alternative required a mesh of Fibre Channel switches. This approach involved stacking a large number of switches and could require multiple hops through the switches to reach the right storage array.

Although the Fibre Channel approach met the basic access requirements, that fact was overshadowed by a high probability of congestion on the interswitch links. In addition, the architecture cost exceeded the medical research facility's budget allocation.

The Cisco storage router delivers redundant iSCSI paths to a pair of Fibre Channel switches. iSCSI takes advantage of the connection-oriented TCP protocol for reliable service. Ethernet was already part of the IT network. This meant trained personnel were on board, and simplified the storage networking installation shown in the figure.

Cost was an important factor in choosing iSCSI. Because the research facility already had TCP/IP and Gigabit Ethernet networks installed, the iSCSI solution fit their budget and met their storage networking needs.

The Cisco storage router met the budget constraints of the research facility and allowed the company to tune storage performance to meet the facility's requirements for volume, activity rates, and management of access conflicts. For the medical research facility, storage access was the primary need. Given that the solution would include high-performance servers, the researchers determined that it was important to directly connect those servers to the Fibre Channel switches.

The end result used the scalability and cost advantages of iSCSI via Gigabit Ethernet for storage networking and retained Fibre Channel for storage access. This hybrid approach permitted the medical research facility to reduce capital costs and meet their operational needs.

The applications interfaced with the generic SCSI layer in the Windows hosts, which see only SCSI. The storage router shielded the host from any Fibre Channel considerations. This transparency gave the medical research facility complete flexibility in designing their storage network.

Storage over WAN

Storage over WAN enables fast, secure, and highly available storage networks interconnecting over wide-area networks. This topic describes the design of an IP access storage networking solution over a WAN.

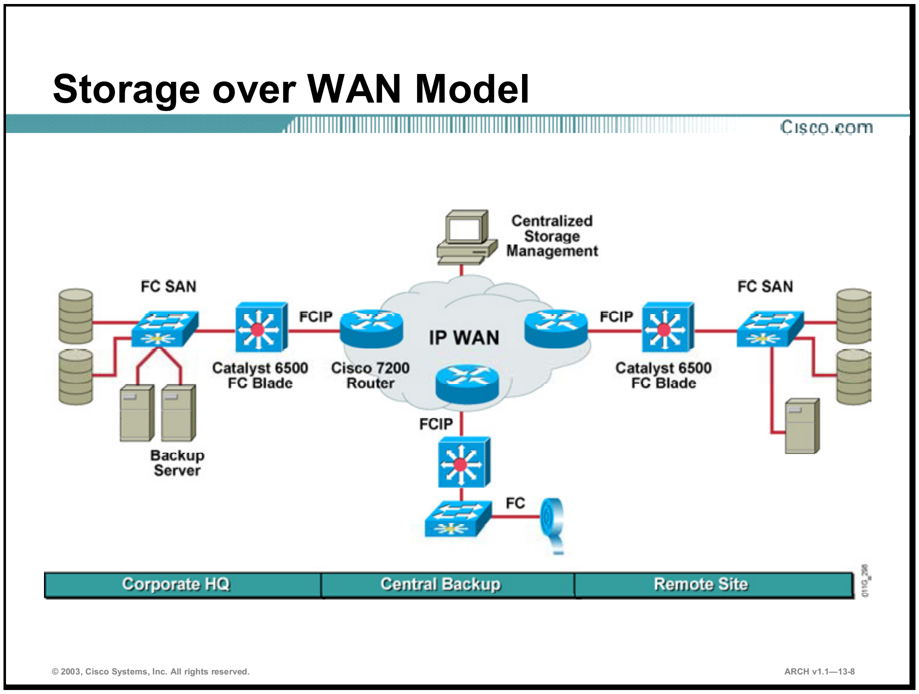

Storage over WAN Model

ARCH v1.1—13-8

To build a truly enterprise-wide storage resource, companies need the ability to interconnect and manage storage across WANs. The growing requirements to replicate, distribute, and manage data over relatively long distances results from several needs.

The most critical need is for backup and restore services to ensure business continuance and data protection at an offsite location. Other applications include centralized storage management, efficient data center migration, or the aggregation of multiple production databases into a single data warehouse for analysis and data mining.

For IP-based networked storage, either NAS or SAN (the traditional wide-area IP routing technologies) allow enterprises to manage storage.

For Fibre Channel- or ESCON-based MAN applications up to a 10-km distance, you must bridge the storage environment to the IP network infrastructure. For Fibre Channel, the FCIP protocol provides transparent encapsulation of the complete FC frame. This allows you to transport both host-to-storage traffic and storage-to-storage replication traffic transparently across the MAN.

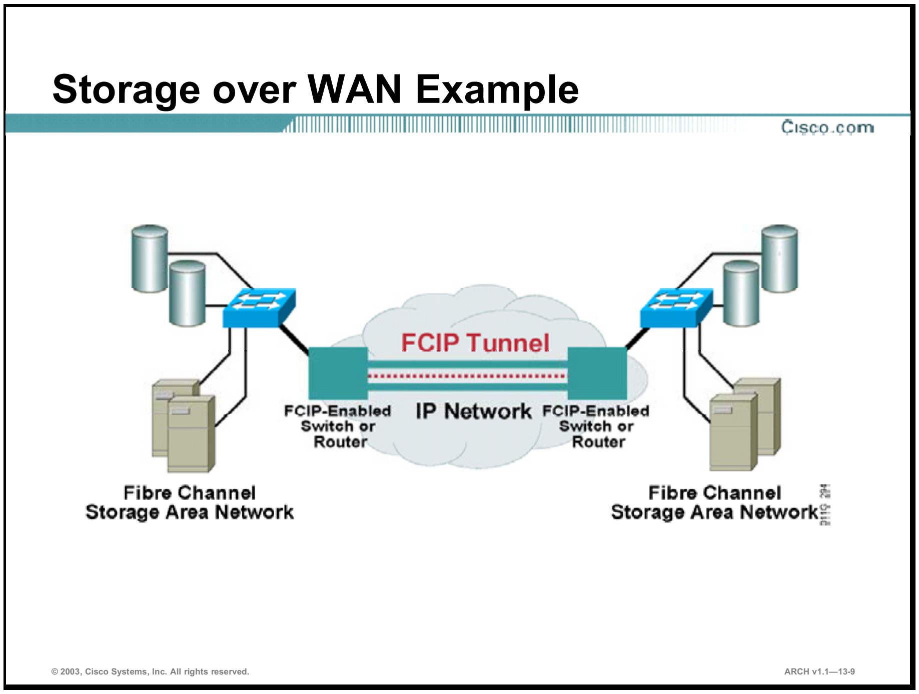

Storage over WAN Example

Cisco.com

FCIP Tunnel

FCIP-Enabled Switch or Router

IP Network

FCIP-Enabled Switch or Router

Fibre Channel Storage Area Network

Fibre Channel Storage Area Network

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-9

Company Background

A large financial institution decided that their critical data needed additional protection through expanded backup procedures. They already had multiple data centers in place, each with an FC SAN.

Storage Networking Solution

An FCIP solution allowed the financial enterprise to add storage to each data center, then use FCIP to perform asynchronous backup of data from one site to another. This allowed them to take advantage of their existing IP infrastructure over the MAN and to provide additional locations for all critical data.

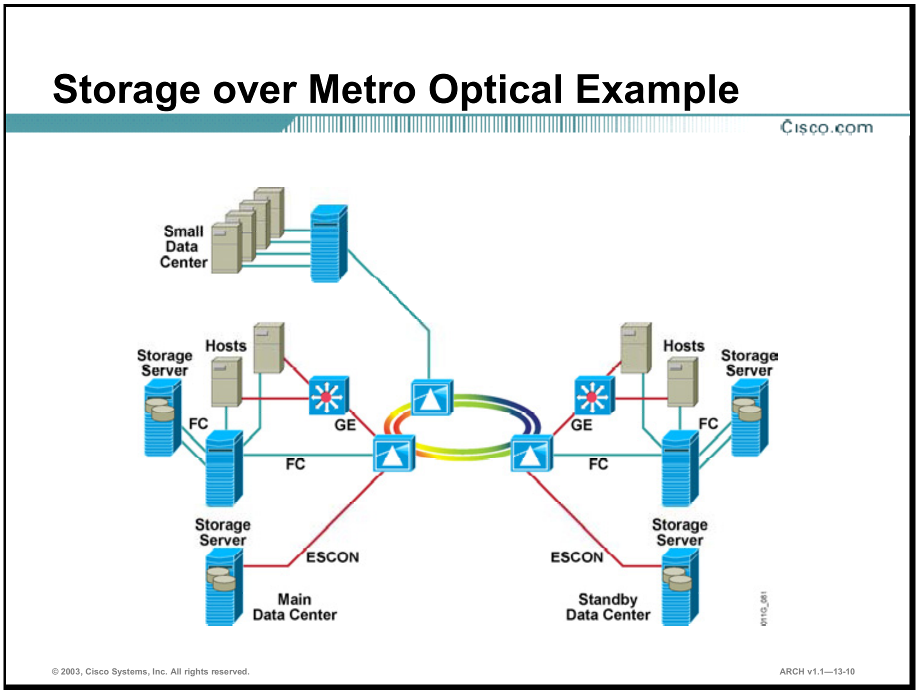

Storage over Metro Optical Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-10

Company Background

A financial management company is concerned about maintaining good customer relationships. They want to adopt strategies that keep application services up and running while protecting business-critical information from corruption and loss. An outage would be devastating to the business. Brokerage firms and other financial institutions can lose millions of dollars per hour when systems are down. Even retail sales organizations can lose hundreds of thousands of dollars per hour when customers cannot place orders.

Storage Networking Solution

The company is implementing a business continuance strategy for storage that addresses both data backup and disaster recovery. Backup and replication includes data archiving for protection against data loss and corruption, remote replication of data for distribution of content, application testing, disaster protection, and data-center migration. Real-time disaster recovery implemented with synchronous mirroring allows the company to safeguard data by guaranteeing that mission-critical data is securely and remotely mirrored to avoid any data loss in the event of a disaster, ensuring uninterrupted services to employees, customers, and partners.

The business continuance strategy and associated SAN technology, implemented with Fibre Channel and ESCON, requires a fault-tolerant, high-bandwidth, and low-latency network. For synchronous mirroring, the high bit rate of an optical network minimizes the time necessary to complete a data transfer. This is critical to avoid negative impact on application performance.

The figure shows a metro optical ring, implemented with the Cisco extended services platform, providing high-bandwidth transport for multiple Gigabit Ethernet LANs, FC SANs, and mainframe environments across a single fiber pair.

Network-Attached Storage

Network-attached storage provides high-performance access, data protection, and disaster recovery for file sharing over an IP network. This topic describes the design of a network-attached storage networking solution.

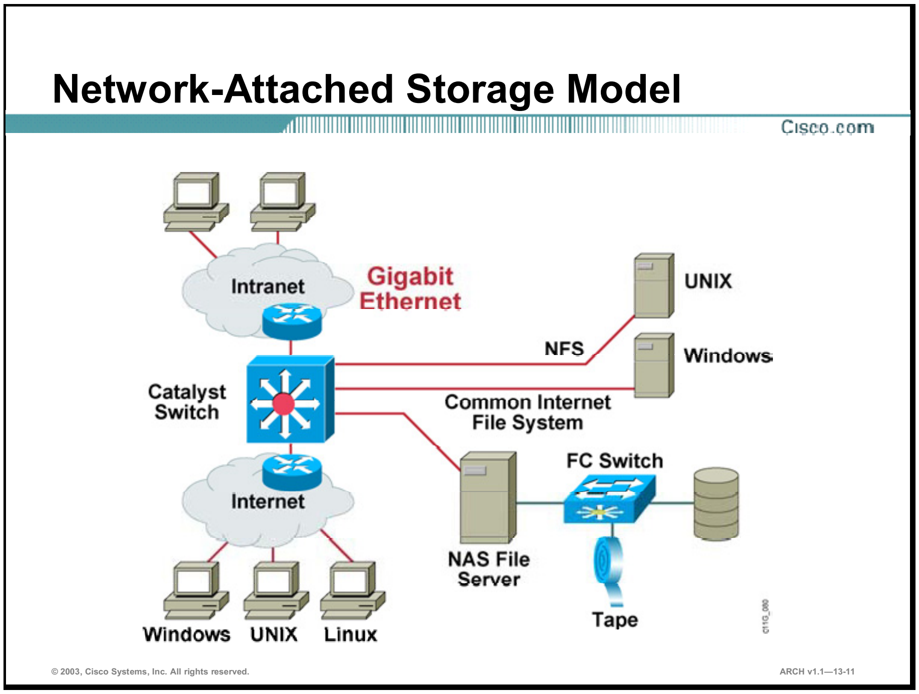

Network-Attached Storage Model

ARCH v1.1—13-11

NAS makes mainstream deployment of IP-based storage consolidation and file sharing possible. NAS is popular for many applications including collaborative development, engineering, e-mail, web serving, and general file serving. In particular, because NAS abstracts storage to the file-system level, it can manage the sharing of files effectively between multiple users and applications. The figure describes a generic NAS deployment across an IP, Gigabit Ethernet network.

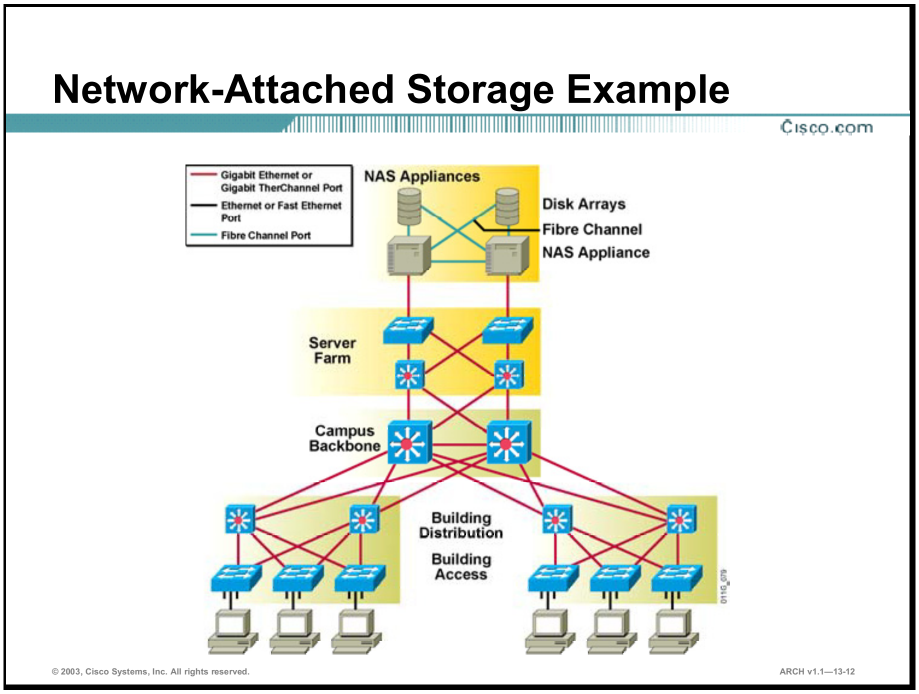

Network-Attached Storage Example

Cisco.com

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-12

Company Background

An expanding software development company found that as they added engineering sites, their current solution for centralized code storage was becoming cumbersome and inefficient. Each site was needing to maintain separate servers with dedicated storage, introducing the possibility of synchronization issues, or they had to access file servers over the WAN. Latency wasn't an issue, since entire files were downloaded to individual engineering workstations for modification, then uploaded back to the servers.

Storage Networking Applications

By deploying NAS devices in their network, the company leveraged their existing infrastructure and allowed each site to maintain their own application servers while keeping the data itself centralized, eliminating synchronization problems. NAS requires file locking.

Summary

This topic summarizes the key points discussed in this lesson.

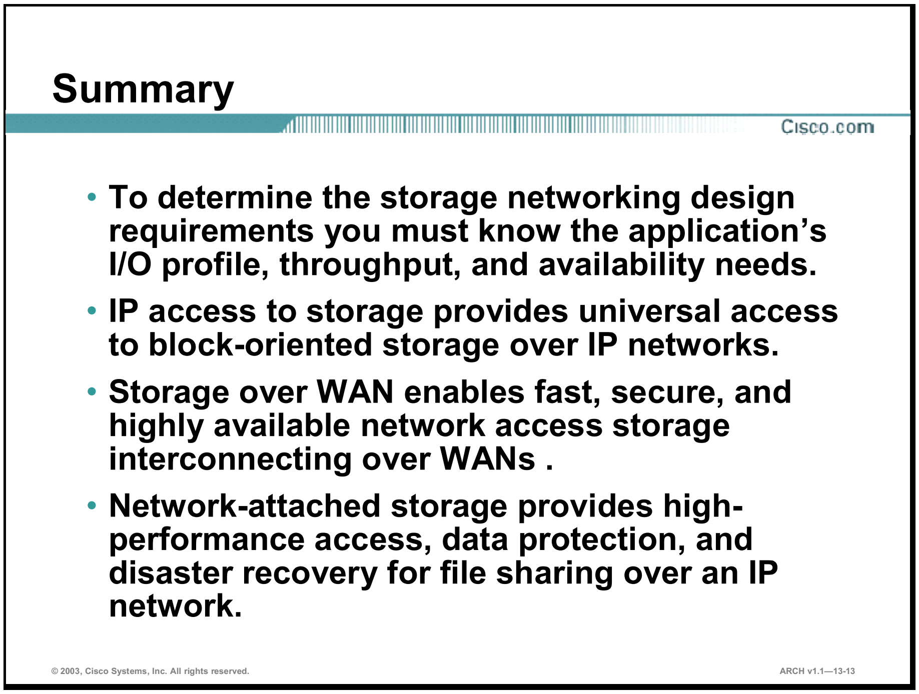

Summary

Cisco.com

- To determine the storage networking design requirements you must know the application’s I/O profile, throughput, and availability needs.
- IP access to storage provides universal access to block-oriented storage over IP networks.
- Storage over WAN enables fast, secure, and highly available network access storage interconnecting over WANs .
- Network-attached storage provides high-performance access, data protection, and disaster recovery for file sharing over an IP network.

ARCH v1.1—13-13

References

For additional information, refer to these resources:

- *Storage Networking* at http://www.cisco.com/warp/public/779/largeent/learn/technologies/storage/apps.html
- Solutions Reference Network Design (SRND) Networking Solutions Design Guides; to locate these documents:
  - — Go to: http://www.cisco.com/.
  - — In the Search box, enter “SRND” and click **Go**. A list of SRND Networking Solutions Design Guides appears.
  - — Select the Networking Solutions Design Guide that meets your needs.

Next Steps

For the associated case study, refer to the following section:

- Case Study 13-2: OCSIC Bottling Company

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) What application requirement can make geographic placement of storage devices an issue?

A) limited latency

B) high bandwidth

C) low throughput

D) high availability

Q2) What is one advantage of introducing iSCSI to a high-volume Fibre Channel-only storage solution?

A) improves data availability

B) reduces the application response time

C) reduces the bandwidth used to transport the storage data

D) reduces the number of hops to reach the correct storage device

Q3) What is one drawback of a high-volume Fibre Channel-only storage solution?

A) cost

B) data availability

C) application response time

D) bandwidth used to transport the storage data

Q4) How is synchronous mirroring supported by storage networking over an optical metro area network?

A) by providing security

B) by providing redundancy

C) by providing high bandwidth

D) by providing high availability

Q5) At what level does NAS provide access to storage?

A) at the file level

B) at the byte level

C) at the block level

D) at the application level

Quiz Answer Key

Q1) A

Relates to: Designing a Storage Networking Architecture

Q2) D

Relates to: IP Access to Storage

Q3) A

Relates to: IP Access to Storage

Q4) C

Relates to: Storage Over WAN

Q5) A

Relates to: Network-Attached Storage

Case Study 13-2: OCSIC Bottling Company

Complete this case study to practice the key design skills discussed in this module.

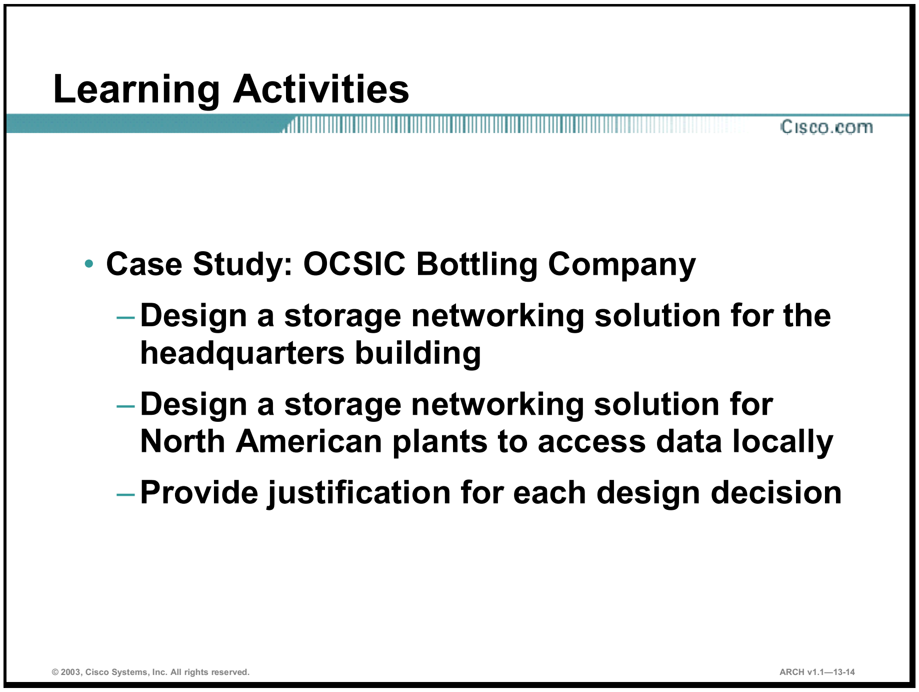

Learning Activities

Cisco.com

- Case Study: OCSIC Bottling Company
  - Design a storage networking solution for the headquarters building
  - Design a storage networking solution for North American plants to access data locally
  - Provide justification for each design decision

© 2003, Cisco Systems, Inc. All rights reserved.

ARCH v1.1—13-14

Required Resources

There are no resources required to complete this exercise.

Exercise Objective

The OCSIC Bottling Company believes that storage networking will help provide faster access to data and reduce the load on the network between buildings at the headquarters campus and between the headquarters' office and the North American plants.

In this exercise, you will design a storage network that meets the needs of the OCSIC Bottling Company.

After completing this exercise, you will be able to:

- Design a storage networking solution for the headquarters building
- Design a storage networking solution for North American plants to access data locally

Each component of the design will have multiple options. Consider each option carefully, given the case study constraints. As you identify each component of your design, provide justification for your decision. The justification explains the options you considered and why you chose the option you selected. Remember, there are no wrong answers to this exercise.

Task 1: Design a Storage Networking Solution

Complete these steps:

Step 1 Determine the location of each storage device on the network. On an overhead transparency, create a campus network diagram for the headquarters location and the North American plants, indicating the location of each storage networking solution.

Step 2 Complete the table to design the details about the storage network.

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your storage networking solution?		
What security strategy will you deploy to support your storage networking solution?		
What QoS strategy will you deploy to support your storage networking solution?		

Task 2: Present Your Design

Present your design to the class. Be prepared to justify each design decision.

Exercise Verification

You have completed this exercise when you attain these results:

- You have created a storage networking solution for the headquarters building. The solution specifies the locations of each storage solution, networking technologies, networking devices, high availability, security, and QoS.
- You have created a storage networking solution for North American plants to access data locally. The solution specifies the locations of each storage solution, networking technologies, networking devices, high availability, security, and QoS.

ARCH

Course Glossary

The Course Glossary for *Designing Cisco Network Service Architectures* (ARCH) v2.0 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.

Acronym or Term	Definition
1000BASE-T	1000 Mbps baseband Gigabit Ethernet specification using twisted pair based on the IEEE 802.3ab standard.
100BASE-T	100 Mbps baseband Fast Ethernet specification using UTP wiring. Like the 10BaseT technology on which it is based, 100BaseT sends link pulses over the network segment when no traffic is present. However, these link pulses contain more information than those used in 10BaseT. Based on the IEEE 802.3u 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. 10BaseT, 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 (said “triple A”). Network security services that provide the primary framework through which you set up access control on your router or access server. AAA protocol requirements for network access are defined in RFC 2989.
ABR	available bit rate. A service category defined by the ATM Forum for ATM networks. It relates traffic characteristics and QoS requirements to network behavior. 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 they receive describing the status of the network and its capability to successfully deliver data. Area Border Router. Router located on the border of one or more OSPF areas that connect 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 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 or RAS.
ACD	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. automatic call distribution. Device or service that automatically reroutes calls to customers in geographically distributed locations served by the same CO.
ACL	access control list.
address	Data structure or logical convention used to identify a unique entity, such as a particular process, network interface (IP) or a network device (DECnet).
address mask	A bit combination used to describe which part of an IP network address refers to the network or the subnet and which part refers to the host. Sometimes referred to simply as mask.
administrative distance	Rating of the trustworthiness of a routing information source. Administrative distance is expressed as a numerical value between 0 and 255 in Cisco IOS. The higher the value, the lower the trustworthiness rating. It is configurable.
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 may work at distances up to 18,000 feet (5,488 meters) over a single copper twisted pair. Available data rates are affected by distance, wire gauge, and the quality of the local cable plant.

Acronym or Term	Definition
analog-signal	The representation of information with a continuously variable physical quantity, such as frequency and amplitude. Because of the constant changing of the wave shape with regard to a given point in time or space, an analog signal has a virtually infinite number of states or values. This contrasts with a digital signal that 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 to protect itself against replay attacks. IPSec provides this optional service by use of a sequence number combined with the use of data authentication.
ASBR	Autonomous System Border Router. A router that connects multiple OSPF routing domain.
ASIC	application-specific integrated circuit. Chip that is built for a specific application.
ASN.1	Abstract Syntax Notation number One. A formal notation used to describe data transmitted by telecommunications protocols.
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 through OC-192c.
AVVID	Architecture for Voice, Video and Integrated Data.
BackboneFast	A feature on the switch that reduces the Spanning Tree Protocol convergence time from 50 seconds to 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. 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 in excess of 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.
BGP	Border Gateway Protocol. Interdomain routing protocol that replaces EGP. BGP exchanges reachability information with other BGP systems. It is defined by 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.

Acronym or Term	Definition
broadband	Describes facilities or services that operate at the DS1 rate and above. Describes media in which multiple frequencies are available to transmit information. This allows information to be multiplexed and sent on many different frequencies or channels within the band concurrently; contrast with baseband.
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 broadcast address is a MAC destination address of all ones. An IPv4 broadcast address is one in which the host portion of the address is all ones. There is no corresponding capability in IPv6.
broadcast domain	Set of all devices that receive broadcast frames originating from any device within the set. Routers typically bound data link broadcast domains because routers do not forward data link 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 time-outs.
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.
Building Distribution submodule	A submodule within the Enterprise Composite Network Model. Provides aggregation of access networks using Layer 3 switching. Performs routing, QoS, and access control.
CA	certification authority In Internet Key Exchange, the CA is a trusted agent responsible for certificate management.
CAC	Call Admission Control. CAC mechanisms protect existing traffic from being negatively affected by new traffic requests and keeps excess traffic off the network. Connection Admission Control. The ATM Forum Specification defines CAC as the set of actions taken by the network during the call set-up phase to determine whether a connection request can be accepted or should be rejected (or whether a request for re-allocation can be accommodated).
campus	One or more buildings with multiple virtual and physical networks, connected across a high-performance backbone.
Campus Backbone submodule	A module within the Enterprise Composite Network Model that connects distribution modules.
Campus Infrastructure module	A module within the Enterprise Composite Network Model that comprises Building Access and Building Distribution submodules.
CAR	committed access rate. In Cisco IOS, the CAR limits the input or output transmission rate on an interface or subinterface based on a flexible set of configured criteria.
Category 5	One of several grades of UTP cabling described in the EIA/TIA-568 and ISO 11801 standards. Category 5 cabling can transmit data at speeds up to 100 Mbps over limited distance.
CBR	Constant Bit Rate. A service category defined by the ATM Forum for ATM networks. It relates traffic characteristics and QoS requirements to network behavior. CBR is used for connections that depend on precise clocking to ensure undistorted delivery.
CBWFQ	class-based weighted fair queuing extends WFQ functionality to provide support for user-defined traffic classes.
CDN	Content Delivery Network.

Acronym or Term	Definition
CDP	Cisco Discovery Protocol. Media and protocol independent device-discovery protocol that runs on Cisco systems, including routers, access servers, and switches. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN or on the remote site of a WAN segment.
CDR	Call Detail Record. Used in telephony networks devices, including 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.
CEF	Cisco Express Forwarding. Scalable, distributed Layer 3 switching technology designed to enhance network performance within supported platforms.
CELP	code-excited linear prediction. A family of algorithms for compression of a digital audio stream that simulates speech using a combination of tone generators and filters. It does not allow exact reproduction of the input.
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.
CIDR	classless interdomain routing. A way to allocate and specify IPv4 addresses more flexibly than with the original system of address classes.
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 negotiated tariff metrics.
Cisco CallManager	Cisco CallManager is the software-based call-processing component of the Cisco enterprise IP telephony solution. Cisco 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 IOS	Cisco Internetwork Operating System. Cisco software that provides common functionality, scalability, and security for Cisco products. Cisco IOS 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 comprehensive web-based network management solution that provides monitoring and configuration tools to simplify the administration of 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 class network boundaries only (class A, B, or C).
classless routing protocols	Routing protocols that propagate subnet mask information with each routing update to enable route summarization anywhere in the IP address, not just on major class network boundaries (class A, B, or C).
CLI	command-line interface. A syntactic user interface that allows interaction with the application or operating system though commands and optional arguments entered from a keyboard. Cisco IOS, UNIX operating systems and DOS provide CLIs. Contrast with GUI.

Acronym or Term	Definition
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 not purchased directly or installed by the subscriber. Cisco CallManager is the software-based call-processing component of the Cisco enterprise IP telephony solution. Cisco 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.
CMF	Common Management Foundation. A set of shared application services for the CiscoWorks network management solution.
CO	central office. The local telephone company office in which all local loops in a given area terminate and where circuit switching of subscriber lines occurs.
codec	COder-DECoder. Integrated circuit device that transforms analog acoustic signals into a digital bit stream (coder) and digital signals back into analog signals (decoder). Software that uses a DSP software algorithm to compress/decompress digital speech or audio signals.
collision domain	A single CSMA/CD network in which there will be a collision if two devices attached to the system transmit at the same time. Ethernet uses CSMA/CD. Repeaters and hubs extend the collision domain; LAN switches, bridges, and routers do not.
connectionless	Term used to describe data transfer without the existence of a physical or virtual circuit.
connection-oriented	Term used to describe data transfer that requires the establishment of a physical or virtual circuit connecting the end points of the transfer.
content cache	A device 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	A device that performs all the management functions needed to control content distribution accessible through a browser interface.
content networking	A technology for optimization of web content delivery that proactively distributes cacheable content from origin servers to content servers at the edges of the network, and keeps content fresh.
convergence	The agreement of a group of interconnected internetworking devices running a specific routing protocol on the network topology of an internetwork after a change in that topology. Speed and ability of a group of interconnected switches to rebuild a spanning tree following a topology change.
CoS	class of service. An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages. In Ethernet networks, CoS is signaled using three bits in the frame header. Closely related to type of service in networks implemented using Cisco routers and switches. 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.
CPE	customer premises equipment. Terminating equipment, such as terminals, telephones, and modems installed at customer sites, and connected to the telephone company network. Can also refer to any telephone equipment supplied by the telephone company residing on the customer site.

Acronym or Term	Definition
CPU	central processing unit. Computing part of a computer or networking device.
crossbar	A type of high-performance switching fabric found in high-end Cisco switches.
cRTP	compressed Real-Time Transfer Protocol. A method to conserve bandwidth, which compresses IP headers from 40 bytes to 2 or 4 bytes, offering significant bandwidth savings. CRTP is sometimes referred to as RTP header compression. Configured on each link individually, it is often used within networks transporting delay-sensitive traffic over narrow links.
cryptography	The principles, means, and methods for making plain information unintelligible, and for restoring the processed information to intelligible form.
CSMA/CD	carrier sense multiple access/collision detect
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 installed optical fiber infrastructure through which no light is being transmitted, or installed fiber optic cable not carrying a signal.
data link layer	Layer 2 of the OSI reference model. This layer responds to service requests from the network layer and issues service requests to the physical layer. It 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 divided this layer into two sublayers: the MAC sublayer and the LLC sublayer. Sometimes simply called link layer.
DCE	data circuit-terminating equipment. The equipment that (a) performs functions, such as signal conversion and coding, at the network end of the line between the DTE and the line, and (b) may be a separate or an integral part of the DTE or of intermediate equipment. The interfacing equipment that may be required to couple the DTE into a transmission circuit or channel and from a transmission circuit or channel into the DTE.
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 adaptor 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.
DES	Data Encryption Standard. Standard cryptographic algorithm developed by the U.S. National Bureau of Standards.
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.

Acronym or Term	Definition
DHCP	Dynamic Host Configuration Protocol. Provides a mechanism for allocating IP addresses dynamically so that addresses can be reused when hosts no longer need them. DHCP is defined in 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	Dial peers are used in packet voice networks to identify call source and destination endpoints and to define the characteristics applied to each call leg in the call connection. Attributes that are defined in a dial peer and applied to the call leg include codec, QoS, VAD, and fax rate. In Cisco devices plain old telephone service and voice-network dial peers can be defined.
dial-up	Communications circuit that is established by a switched-circuit connection using the dial telephony circuits, usually over the PSTN.
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	Class of routing algorithms that use a relatively simple measure, such as the number of hops in a route, to define the best path toward a destination network. 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. A common algorithm used to build the routing table is the Bellman-Ford algorithm.
DLCI	data-link connection identifier. Value that identifies a VC in the physical link connecting client equipment and a Frame Relay network edge device. In the basic Frame Relay specification, DLCIs are locally significant. Connected devices might use different values to specify the same destination device, and the same DLCI also might be used on different physical links. The optional global addressing extension to the Frame Relay Forum (FRF) LMI specification makes DLCIs globally significant.
DNS	Domain Name System. DNS is used on the Internet for translating names of network nodes into addresses. DNS is Internet Standard 13. The protocol definitions are spread over several RFCs.
DOCSIS	Data-over-Cable Service Interface Specifications. Technical specifications for equipment at both subscriber locations and the head-ends of cable operators. Adoption of DOCSIS will ensure interoperability of equipment throughout the infrastructures of system operators.
denial of service	An incident in which a user or organization is deprived of the services of a resource they 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.
DSL	digital subscriber line. Public network technology that delivers high bandwidth over conventional copper wiring at limited distances. The types of DSL include ADSL, RADSL, 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 may be capacity 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 software-configurable CPU that processes analog to digital (and vice versa) data streams. Performs audio/video coding and transcoding, including compression.
DTMF	dual tone multifrequency. Tones generated when a button is pressed on a telephone.

Acronym or Term	Definition
DWDM	dense wavelength division multiplexing. Optical transmission of multiple signals in a single optical fiber over closely spaced wavelengths in the 1550 nm region. (Frequency 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	earth and magnet (more commonly “ear and mouth”) signaling arrangement generally used for switch-to-switch or switch-to-network trunks. E&M is available on analog and digital interfaces. The term originally comes from the term earth and magnet. Earth represents electrical ground and magnet represents the electromagnet used to generate tone.
E1	Channelized digital transmission scheme used internationally that carries data in up to 32 64 Kbps channels at an aggregate 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.
EAP	Extensible Authentication Protocol. Framework that supports multiple, optional authentication mechanisms for PPP, including clear text passwords, challenge-response, and arbitrary dialog sequences. EAP is defined in RFC 2284.
echo	Audible and unwanted leak-through of one's own voice into one's own receive (return) path. A signal from the transmission path is returning to one's ear through the receive 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 module 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.
EIGRP	Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency. A hybrid, it combines the advantages of link state protocols with those of distance vector protocols.
e-mail	electronic mail. Widely used application in which text messages are transmitted electronically between end users over various types of networks using various network protocols. Underlying network application protocols include SMTP and POP.
encryption	Application of a specific algorithm to data so as to alter the representation of the data making it incomprehensible to those who do not have access to the algorithm and key required to reverse the process.

Acronym or Term	Definition
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 performance, scalability, and availability. This area contains all the network elements for independent operation within one geographic location.
Enterprise Composite Network Model	A model of enterprise campus networks that logically and physically segregates the campus along functional boundaries.
Enterprise Edge	A functional area within the Enterprise Composite Network Model. Aggregates the connectivity from the various elements at the edge of each enterprise campus network.
enterprise network	The comprehensive network that connects an organization. It includes all LAN, campus, metropolitan, and WAN links, and equipment.
Ethernet	Baseband LAN specification invented by Xerox Corporation and developed jointly by Xerox, Intel, and Digital Equipment Corporation. Ethernet networks use CSMA/CD or switches and run over a variety of cable types at 10 Mbps. Current Ethernet implementations are defined in the IEEE 802.3 series of standards.
failover	A 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 down time.
failure domain	A group of Layer 2 switches connected together is called a Layer 2-switched domain. The Layer 2-switched domain can be considered as a failure domain because a misconfigured or malfunctioning workstation can introduce errors that will impact or disable the entire domain.
Fast EtherChannel	Bundled 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 10BaseT Ethernet specifications while preserving such qualities as frame format, MAC mechanisms, and MTU. These similarities allow the use of existing Ethernet applications and network management tools on Fast Ethernet networks. Based on an extension to the IEEE 802.3 specification.
FCAPS	Fault, Configuration, Accounting, Performance and Security. A model of network management that divides the required activities into fault management, configuration management, accounting management, performance management, and security management.
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.
fiber optics	A medium used for the transmission of information (audio, video, data). Light is modulated and transmitted over high-purity, hair-thin fibers of glass or plastic. The bandwidth capacity of fiber optic cable is much greater than that of conventional coaxial cable or copper wire.
firewall	A network appliance, or software running in a router or access server, which provides isolation between any connected public networks and a private network. A firewall router uses access lists and other methods to ensure the security of the protected network.
flat addressing	Scheme of network addressing that does not use a logical hierarchy to determine association. For example, MAC addresses are flat. Bridging protocols must flood packets throughout a flat network to deliver the packet to the appropriate location.
fragmentation	Process of breaking a packet into smaller units when transmitting over a network medium that cannot support the original size of the packet.

Acronym or Term	Definition
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. Transmission unit within time division multiplexed media. For example, a T1 frame is 193 bits, while an E1 frame is 256 bits.
Frame Relay	Industry-standard, switched data link layer protocol that handles multiple virtual circuits using HDLC derived encapsulation between connected devices. Frame Relay is more bit efficient and less robust than X.25, the protocol for which it generally is considered a replacement.
FTP	File Transfer Protocol. Application protocol, part of the TCP/IP protocol family, used for transferring files between network nodes. FTP is an Internet Standard 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 usually is reserved for network backbones because it can be very expensive to implement.
FXO	Foreign Exchange Office. An FXO interface is intended to connect to the PSTN central office. It is the interface offered on a standard telephone. Cisco's FXO interface allows an analog connection to the PSTN central office or to a station interface on a PBX.
FXS	Foreign Exchange Station. An FXS interface connects directly to a standard telephone and supplies ring, voltage, and dial tone. Cisco's FXS interface allows connections to basic telephone service equipment, key sets, and PBXs.
G.711	Defines the 64-Kbps PCM voice coding technique. One of the ITU-T G-series recommendations, G.711 encoded voice is the expected format for digital voice delivery in the PSTN or through PBXs. This coding technique allows reproduction of the input stream.
G.723	Defines a compression technique that can be used for compressing speech or audio 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. It is one of the ITU-T G-series recommendations.
G.729	Describes a CELP compression in which voice is coded into 8-kbps streams. There are several variations of this standard (G.729, G.729 Annex A, and G.729 Annex B) that differ mainly in computational complexity; all provide speech quality similar to 32-kbps ADPCM. It is an ITU-T G-series recommendation.
gatekeeper	The component of an H.323 conferencing system that performs call address RAS bandwidth management. Telecommunications: H.323 entity 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 locating 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	A device that performs application-layer conversion of information from one protocol stack to another. In the IP community, it originally referred to a routing device. Today, the term router is used to describe nodes that perform this function.
Gb	gigabit. Approximately 1,000,000,000 bits.
Gbps	gigabits per second.

Acronym or Term	Definition
GBps	gigabytes per second.
GFR	Guaranteed Frame Rate. A service category defined by the ATM Forum for ATM networks. It relates traffic characteristics and QoS requirements to network behavior. GFR is a frame-aware service that only applies to VCCs since frame delineation is not usually visible at the virtual path level.
Gigabit EtherChannel	Bundled multiple Gigabit Ethernet links, which appear as one logical interface.
Gigabit Ethernet	Standard for a high-speed Ethernet at 1 Gbps, approved by the IEEE 802.3z standards committee in 1996.
GPRS	general packet radio service. A service defined and standardized by the ETSI. GPRS is an IP packet-based data service for global system for mobile communication (GSM) networks.
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 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. GRE is defined in RFC 2784.
GUI	graphical user interface. A navigational user interface that allows interaction with the application or operating system through selection of menu items or icons using a pointing device such as a mouse. Apple MAC OS and Microsoft Windows provide GUIs. Contrast with CLI.
H.225.0	An ITU recommendation 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 recommendation that governs endpoint control. It is specified for use by H.323.
H.323	An ITU-T recommendation that allows dissimilar packet communication devices to communicate with each other by using a standardized communication protocol. H.323 defines a common set of codecs, call setup and negotiating procedures, and basic data transport methods.
headend	End point of a broadband cable network. All stations transmit toward the head-end; the head-end transmits toward the destination stations.
header	Control information placed before data when encapsulating that data for network transmission.
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 location.
hierarchical routing	The complex problem of routing on large networks can be simplified by reducing the size of the networks. This 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.

Acronym or Term	Definition
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.
HSRP	Hot Standby Router Protocol. 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. Other routers in the group monitor the lead router, and if it fails, one of the standby routers inherits the lead position and the Hot Standby group address. HSRP is documented in RFC 2281.
HTML	Hypertext Markup Language. Simple hypertext document formatting language that uses tags to indicate how a given part of a document should be interpreted by a viewing application, such as a web browser. The World Wide Web Consortium (W3C) maintains the HTML standard.
HTTP	Hypertext Transfer Protocol. The protocol used by web browsers and web servers to transfer files, such as text and graphic files. HTTP is defined in RFC2616.
IBGP	Internal Border Gateway Protocol. IBGP is a variant of the 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*. Cisco training course.
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. Digital Equipment Corporation invented the original algorithm. 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 over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.
IGMP	Internet Group Management Protocol. Used by IP hosts to report their multicast group membership requests to an adjacent multicast router. IGMP is defined in RFC 3376.
IGP	Interior Gateway Protocol. Internet protocol used to exchange routing information within an autonomous system. Examples of common Internet IGPs include EIGRP, OSPF, IS-IS, 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. IKE establishes a shared security policy and authenticates keys for services (such as IPSec) that require keys. Before any IPSec traffic can be passed, each router/firewall/host must verify the identity of its peer. This can be done by manually entering pre-shared keys into both hosts or by a CA service. IKE is defined in RFC2409.

Acronym or Term	Definition
ILMI	Integrated Local Management Interface. Specification developed by the ATM Forum for incorporating interface management capabilities into the ATM UNI.
in-band signaling	Transmission of control information within a content stream also 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 it more efficient than two separate implementations. Formerly called Dual IS-IS. Use of Integrated IS-IS for routing in a TCP/IP network is defined in RFC 1195.
intelligent network services	Services that enable application awareness within the network. Intelligent network services add intelligence to the network infrastructure beyond that required to just move a datagram between two points. Examples of intelligent network services include network management, security, high availability, QoS, and IP multicasting.
interarea routing	Term used to describe routing between two or more logical areas.
Internet	The largest 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 Edge functional area of the Enterprise Composite Network Model. This module provides internal enterprise users with connectivity to 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	Intranet is a closed, organization-wide network that includes LANs and WANs. It frequently uses open standards such as TCP/IP instead of proprietary protocols traditionally used for LANs and WANs.
Intrusion Detection	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.
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	An address assigned to hosts using TCP/IP. An IPv4 network address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated by periods (dotted decimal format), or 8 hexadecimal digits. Each address consists of a network portion, which includes the network number and an optional subnetwork number, and a host number. A subnet mask, or CIDR prefix, is used to extract the network portion from the IP address. The network number, extended using a single subnet mask for that network, can be used for routing in all routing protocols. More modern routing protocols can route on the network portion constructed using variable length subnet masks. The host number is used to address an individual host within the network or subnetwork. Also called an Internet address. In IPv6 the address is 128 bits that are displayed using eight 16-bit groups. Each group is presented using up to four hexadecimal digits and separated by colons.
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 and the header checksum. The IPv6 header is similar in purpose, though very different in structure.

Acronym or Term	Definition
IPM	Internetwork Performance Monitor. CiscoWorks tool used to isolate performance problems, locate bottlenecks, diagnose latency and jitter, and perform trend analysis of network response time.
IP multicast	Internet Protocol multicast. A packet routing technique that allows IP traffic to be propagated efficiently from one source to a number of destinations, or from many sources to many destinations. Rather than sending duplicate packets, one to each destination, only one packet is sent out each interface on which a multicast group identified by a single IP destination group address is registered. This can greatly reduce the required bandwidth.
IPng	IP next generation, the first name for the IPv6.
IP Phone	Device that enables termination of voice communications within the IP network. Cisco IP telephones are centrally managed by the Cisco CallManager. They may be powered inline through Ethernet connections, reducing the need to protect wall power outlets to maintain voice power.
IP precedence	Use of three bits from the term-of-service octet in the IP header to provide limited prioritization for IP packets in a routed network.
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.
IP spoofing	A network attack that occurs when an attacker outside your network pretends to be a trusted user 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 Internet Protocol (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 to deliver services using only a fraction of the aggregate bandwidth required by 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).
ISDN	Integrated Services Digital Network. Communication protocol offered by telephone companies that extended the digital network to the customer premises to carry data, digital voice, and other source traffic integrated as a single service.
IS-IS	Intermediate System-to-Intermediate System. OSI link-state hierarchical routing protocol whereby intermediate systems (routers) exchange routing information based on a single configurable metric to determine network topology.
ISL	Inter-Switch Link. Cisco-proprietary protocol that maintains VLAN information as traffic from multiple VLANs flows between switches and routers on a single physical link.
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.
ISOC	Internet Society.

Acronym or Term	Definition
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. A United Nations agency, 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 DTMF signaling. Examples include banks that allow you to check your balance from any telephone and automated stock quote systems.
jitter	The interpacket delay variance; that is, the difference between interpacket arrival and departure. Jitter is an important QoS metric for voice and video applications. 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.
jitter buffer	Dejitter buffers are used at the receiving end to smooth delay variability and allow time for decoding and decompression. They help on the first talk spurt to provide smooth playback of voice traffic.
Kbps	kilobits per second.
Kerberos	Standard for authenticating network users. Kerberos offers two key benefits: it functions in a multivendor network, and it does not transmit passwords over the network. Kerberos is described in RFC 1510.
L2 switching (L2-switched)	Switching based on Layer 2 (data link layer) information. The current generation of Layer 2 switches are functionally equivalent to bridges. The exposures in a large bridged network include broadcast storms, spanning-tree loops, and address limitations.
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 L2F and PPTP, L2TP provides an industry-wide interoperable method of implementing Virtual Private Dial-up Networks (VPDN). Communications transactions between the LAC and the LNS support tunneling of a single PPP connection. There is a one-to-one relationship among the PPP connection, L2TP session, and L2TP call.
L3 switching	Integrates routing with switching to yield very high routing throughput rates typical of L2 switches while offering Network Layer (L3) routing services and Data Link Layer (L2) termination.
LAN	local-area network. High-speed, low-error rate 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 at various rates is the most widely used LAN implementation technology.
latency	Delay between the time a device transmits a packet and when that packet is received at the destination. Delay between the time a device requests access to a network and the time that it is granted permission to transmit. Delay between the time a device receives a frame and the time the frame is forwarded out the destination port.
leased line	Transmission line reserved by a communications carrier for the private use of a customer. The enterprise perceives a leased line as a type of dedicated line.
LEC	local exchange carrier. A telephone company that provides customer access to the public switched telephone network through one of its central offices.

Acronym or Term	Definition
LFI	link fragmentation and interleaving is a solution for queuing delay situations. With LFI, large packets are fragmented into smaller frames and interleaved with small voice packets. It is similar in effect to FRF.12, Frame Relay Fragmentation, available with Frame Relay.
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 complexity 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.
LMS	LAN Management Solution.
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.
LRE	Long-Range Ethernet. Ethernet standard frames over single-pair wiring at distances of up to 5000 feet.
LSA	link-state advertisement. Broadcast packet used by OSPF that contains information about neighbors and path costs. The receiving routers use LSAs 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 an Ethernet-based 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.
Mb	megabit. Approximately 1,000,000 bits.
MB	megabyte. Depending on the context it can mean either 1,000,000 or 1,048,576 (2^20) bytes.
Mbps	megabits per second. A bit rate expressed in millions of binary bits per second.
MCU	Multipoint Control Unit. An H.323 endpoint that provides the capability for three or more terminals and gateways to participate in multipoint conferences.

Acronym or Term	Definition
MGCP	Media Gateway Control Protocol. A merging of the IPDC and SGCP protocols. MGCP is defined in RFC 2705.
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 defined in ASN.1.
MM fiber	Multimode Fiber. A less costly fiber-optic medium in which light travels in multiple modes.
MPEG	Motion Picture Experts Group. Standards 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 to provide acceptable quality over dedicated connections as narrow as 64 Kbps. MPEG-7 is a standard for description and search of audio and visual content. Work started on MPEG-21 "Multimedia Framework" mid-2000.
MPLS	Multiprotocol Label Switching. Switching method that forwards Network Layer traffic using a label. This label instructs the routers and the switches in a network where to forward the packets based on preestablished routing information determined as the packet entered the network. MPLS is defined in RFC 3031.
MTU	maximum transmission unit. Maximum packet size, in bytes, that a particular interface can transmit without fragmentation.
multicast	The transmission of packets from a single source to multiple destinations in a way which conserves network bandwidth by reducing the duplication of packets sent.
NANP	North American Numbering Plan. A specification for assigning telephone numbers in North America. It implements the E.164 international standard in this region.
NAS	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). 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. Cisco IOS software feature used for real-time network traffic analysis to support QoS requirements.
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).
Network Management module	A module within the Enterprise Composite Network Model. This module performs intrusion detection logging, system logging, and Terminal Access Controller Access Control System Plus (TACACS+)/RADIUS and One Time Passwords (OTP) authentication, as well as network monitoring and general configuration management functions.
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).

Acronym or Term	Definition
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.
non-repudiation 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, intraarea 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.
NTP	Network Time Protocol. Protocol built on top of TCP that ensures accurate local time keeping 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. The current version of NTP is defined in RFC 1305.
ODR	On-Demand Routing. A routing mechanism that uses Cisco Discovery Protocol (CDP) to propagate the IP prefix. ODR appropriate for hub-and-spoke topologies.
OPNET	The vendor of the simulation tool used in the ARCH 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 implemented in hardware and firmware whereas the upper five layers are implemented only in software. The highest layer (the application layer) is closest to the user application. The OSI reference model is used universally as a method 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. The current version of OSPF is defined in RFC 2328.
OTP	One Time Password. Type of authentication that permits a user to enter a password one time for all applications and systems.
out-of-band signaling	Transmission of control traffic using frequencies or channels other than 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. The SS7 signaling system of the PSTN is an example of out-o- band signaling.
p2mp	Point-to-Multipoint. Wireless communication between a series of receivers and transmitters to a central location. Cisco p2mp typically is set up in three segments to enable frequency re-use.

Acronym or Term	Definition
p2p	Point-to-Point. Wireless communication between one receiver and one location. p2p has a higher bandwidth than p2mp for reasons including it has less overhead to manage the data paths and there is only one receiver per transmitter.
packet sniffer	Device that monitors traffic on a network and reports on problems on the network.
partial mesh	Network topology in which devices are only directly connected to some other nodes in the network. A network topology in which every node is connected to at least two of the nodes is described as well connected. A partial mesh does not provide the level of link redundancy of a full mesh topology but is far less expensive to implement. Partial mesh topologies generally are used to connect peripheral networks that distribute traffic to a fully meshed backbone.
password sniffing	Passive traffic interception, usually on a local-area network, to gain knowledge of passwords.
PAT	port address translation. IP address translation method that allows a router to forward packets from several sessions or flows between a private internetwork and the Internet. PAT allows the router to forward packets between a private IP network and the Internet using a single public IP address to support multiple actual users.
PBX	private branch exchange. Digital or analog telephone switch located on the subscriber premises and used to connect private and public telephone networks.
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 part of an existing live network used to test designs, hardware compatibility, and new software.
PIM	Protocol Independent Multicast. Multicast routing protocol 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.
policy domain	A collection of networks under single management.
POP	point of presence. A physical location within service provider network where users dial- n. Post Office Protocol. Internet application protocol providing e-mail services. An Internet Standard, POP is defined by RFC 1939.
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. Technology that enables core routers to send native IP packets directly over SONET/SDH frames. Essentially an IP packet is placed into PPP (RFC 1661), then encapsulated in HDLC like framing (RFC 1662), and finally placed into a SONET/SDH payload (RFC 2615).
POTS	plain old telephone service. Basic service supplying standard single-line telephones, telephone lines, and access to the public switched network.
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, PAP, and EAP. PPP relies on two protocols: link control protocol (LCP) and Network Control Protocol (NCP). PPP is defined in IETF RFC 1661.
pps	packets per second.

Acronym or Term	Definition
PPTP	Point-to-Point Tunneling Protocol. RFC 2637 describes the PPTP protocol.
PQ	priority queuing. Queue management and service discipline that prioritizes traffic at a network interface. 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. It is possible for higher-priority traffic to starve lower-priority traffic by consuming all the bandwidth.
PQ-CBWFQ	priority queuing-class-based weighted fair queuing (PQ-CBWFQ). Feature that joins strict priority queuing and CBWFQ. Strict priority queuing allows delay-sensitive data, such as voice, to be dequeued from a single priority queue and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic. It is also called low latency Queueing (LLQ).
PQ-WFQ	priority queuing-weighted fair queuing. Also called IP RTP Priority. Queuing mechanism that provides a strict priority queuing scheme for delay-sensitive data such as voice.
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 (non-live) network used to test designs, new hardware, and software versions before deployment.
PSTN	Public Switched Telephone Network. General term referring to the variety of telephone networks and services in place worldwide. The PSTN includes POTS and ISDN services.
PVC	permanent virtual circuit. Virtual circuit that is permanently established. PVCs save bandwidth, and operations and processing costs associated with circuit provisioning and tear down in situations where certain virtual circuits must exist all the time. 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.
QoS	quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability.
queue	Generally, a list of elements waiting to be processed.
queuing delay	Amount of time that a data packet must wait in a queue before it can be transmitted onto a statistically multiplexed physical circuit.
RADIUS	Remote Authentication Dial-In User Service. Responsible for receiving user connection requests, authenticating the user, and then returning all configuration information necessary for the client to deliver service to the user. It includes a database for authenticating connections and for tracking connection time. RADIUS is defined in RFC 2865.
RAS	registration, admission, and status protocol. Protocol that is used between endpoints and the gatekeeper to perform management functions. RAS signaling function performs registration, admissions, bandwidth changes, status, and disengage procedures between the VoIP gateway and the gatekeeper. Remote Access Server. One of a family of devices that provide remote-access services to clients.
RED	random early detection. Congestion avoidance algorithm in which some percentage of packets are dropped when congestion is detected and before the queue in question overflows completely.

Acronym or Term	Definition
Remote Access and VPN module	A module within the Enterprise Edge functional area of the Enterprise Composite Network Model. This module terminates VPN traffic, forwarded by the Internet Connectivity module, from remote users and remote sites.
RFC	Request for Comments. Document series used as the primary means for communicating information about Internet protocols and related technical details. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications, such as Telnet and FTP, but some humorous or historical. RFCs are available online from numerous sources.
RIP	Routing Information Protocol. IGP supplied with UNIX BSD systems. The most common IGP in the Internet. RIP uses hop count as a routing metric. Version 2 of RIP, RIPv2 is defined in RFC 2453. RIPng is defined for use with IPv6.
RMON	Remote Monitoring. Remote Network Monitoring MIB specification described in RFC 2819 that defines objects for managing remote network monitoring 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 instance and notifies 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 EIGRP, 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).
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 they want to receive. 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 designed to provide end-to-end network transport functions for applications transmitting real-time data, such as audio, video, or simulation data, 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.
RWAN	Routed WAN Management Solution. CiscoWorks solution used to configure, administer, and maintain a Cisco routed wide-area network (WAN).
SAA	Service Assurance Agent. Network performance measurement agent in Cisco IOS software that provides a scalable, cost-effective solution for service level monitoring.
SAFE	Security Architecture for Enterprise Blueprint. Flexible, dynamic blueprints for security and VPN networks, built on the Cisco Architecture for Voice, Video and Integrated Data (AVVID), that enables businesses to securely take advantage of e-business economies and compete in the Internet economy.
SAN	storage area network.
SDH	Synchronous Digital Hierarchy is a standard technology for synchronous data transmission on optical media. It is the international equivalent of SONET.
SDSL	single-line digital subscriber line. One of four DSL technologies. SDSL delivers 1544 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).

Acronym or Term	Definition
Server Farm module	A module within the Enterprise Composite Network Model. It contains servers providing application, file, print, e-mail, and Domain Name System (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. It most often uses different WAN technologies provided by SPs.
SIP	session initiation protocol. Protocol developed by the IETF as an alternative to H.323. SIP features are compliant with 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.
SMI	Structure of Management Information. RFC 1155 specifying rules used to define managed objects in the MIB.
SMTP	Simple Mail Transfer Protocol. Internet application protocol providing e-mail services. SMTP is defined in RFC 2821.
SN	Storage Networking provides customers with universal access to storage solutions and products an open standards-based architecture. SN combines intelligent Fibre Channel, Ethernet, and optical networking offerings to build scalable data center storage networks and extend storage networks through IP and optical technologies. Major technology areas include NAS and SAN.
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 framing 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.
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 blocking 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.
SPF	shortest path first algorithm or Dijkstra's algorithm. Routing algorithm that iterates on length of path to determine a shortest-path spanning tree. Commonly used in link-state routing algorithms, it runs on every routing device in the network.
spoofing	The act of constructing a packet stream claiming to be from an address other than the actual source. Spoofing is designed to foil network security mechanisms, such as filters and access lists.
static route	Route that is explicitly configured and entered into the routing table. Static routes take precedence over routes chosen by dynamic routing protocols.

Acronym or Term	Definition
Storage Networking (SN)	Storage Networking provides customers with universal access to storage solutions and products in an open standards-based architecture. Storage Networking combines intelligent Fibre Channel, Ethernet, and optical networking offerings to build scalable data center storage networks and extend storage networks through IP and optical technologies. Major technology areas include NAS and SAN.
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 Equipment version. Sometimes abbreviated as STP.
stub area	OSPF area that carries a default route, intraarea 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.
subnet	In IP networks, a network sharing a particular subnet address. Subnetworks are networks arbitrarily segmented by a network administrator to provide a multilevel, hierarchical routing structure while shielding the subnetwork from the addressing complexity of attached networks.
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 to another interface for transmission. Routers use Layer 3 switching to route a packet, and traditional LAN switches use Layer 2 switching to forward frames. See also Layer 2 switching and Layer 3 switching.
syslog	Dedicated server that logs system messages.
T1	Digital WAN carrier facility. T1 transmits DS-1-formatted data at 1.544 Mbps through a TDM network, using AMI or B8ZS coding.
TACACS+	Terminal Access Controller Access Control System Plus. Authentication protocol extended 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 a scalable network security solution.
Tbps	terabits per second. 1,000,000,000,000 bits 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.
TDM	time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots.

Acronym or Term	Definition
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, username and password). TFTP is defined in RFC 1350.
tie-line	A dedicated circuit that connects enterprise PBXs together.
type of service	An indication of how an upper-layer protocol requires a lower-layer protocol to treat its messages. In many IP networks this is signaled using the 3 precedence bits in the type-of-service octet of the IP header. More recent implementations use DSCP in the DS field of the IP header.
touch-tone	Use as adjective, not noun; for example, touch-tone telephone, touch-tone telephone buttons, and so forth.
traffic policing	Process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the agreed upon flow can be discarded immediately or tagged (where some field is changed) and discarded en route if congestion develops. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also known as admission control, permit processing, and rate enforcement. Most often implemented on ingress ports to protect a transport network from greedy traffic flows, it is frequently implemented using a leaky bucket algorithm.
traffic shaping	Use of queues to smooth 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. Most often configured on egress ports to ensure compliance with agreed connection traffic rates to avoid traffic policing, it is frequently implemented using a token bucket algorithm.
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, between a CO and a PBX, or between two PBXes.
TTL	Time to Live.
tunneling	A dual encapsulation mechanism by which a protocol at some layer in the protocol stack is transported by another protocol operating at the same layer.
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. The ordinary copper wire that connects homes and many business computers to the PSTN uses a single pair for each analog telephone line.
UBR	unspecified bit rate. A service category defined by the ATM Forum for ATM networks. It relates traffic characteristics and QoS requirements to network behavior. 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.

Acronym or Term	Definition
unicast	Traffic from a single source 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 include UNIX 4.3 BSD (Berkeley Standard Distribution), developed at the University of California at Berkeley, and UNIX System V, Release 4.0, developed by AT&T.
UplinkFast	A spanning-tree maintenance mechanism that enables the 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. A VACL contains an ordered list of access control entries (ACEs).
VAD	voice activity detection. When enabled on a voice port or a dial peer, silence is not transmitted over the network, only audible speech. When VAD is enabled, the sound quality is slightly degraded but the connection uses much less bandwidth.
VBR	variable bit rate. A service category defined by the ATM Forum for ATM networks. It relates traffic characteristics and QoS requirements to network behavior. VBR is subdivided into a real time (RT) class and non-real time (NRT) 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 that still need a guaranteed QoS.
VLAN	virtual LAN. Group of devices on one or more LANs 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.
voice mail	Voice messaging is a service expected by most users of a telephone system. It provides the facility to divert their incoming calls to a voice mailbox when they are unable to answer their telephones.
VoIP	Voice over IP. The capability to carry voice over an IP-based Internet with POTS-like functionality, reliability, and voice quality. VoIP enables a router to carry voice traffic (for example, telephone calls) over an IP network. In VoIP, the DSP output is collected over 20 or 30 milliseconds and placed in UDP datagrams. These datagrams are transported using IP packets with RTP. Skinny Client Control Protocol (SCCP), H.323 and SIP provide session (call) control.
VPN	Virtual Private Network. Uses tunneling and encryption at the Network Layer to enable IP traffic to travel securely over a public TCP/IP network.
VTP	VLAN Trunking Protocol. VTP reduces administration in a switched network by distributing VLAN information through all switches in the VTP domain. VTP is a Cisco-proprietary protocol that is available on most of the Cisco Catalyst Family products.
vty	virtual type terminal. Commonly used as virtual terminal lines.
WAN	wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. ATM, Frame Relay, PPP, SMDS, and X.25 are examples of common data link layer protocols found in WANs.

Acronym or Term	Definition
WAN module	A module within the Enterprise Edge functional area of the Enterprise Composite Network Model. The WAN module includes all WAN technologies that provide circuits between geographically separated locations. FR, ATM, and PPP are frequently encountered data link technologies.
web	World Wide Web (also called WWW). A client/server system based on HTML and HTTP.
WFQ	weighted fair queuing. Queuing 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 on interfaces at and below 2.048 Mbps.
wildcard mask	A 32-bit quantity 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 configuring access lists.
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. They are sometimes called distribution facilities.
workgroup	Collection of workstations and servers on a LAN that are designed to communicate and exchange data with one another.
WRED	weighted random early detection. Queuing method that ensures that high-precedence traffic has lower loss rates than other traffic during times of congestion by dropping some percentage of packets when congestion is detected and before the queue in question overflows. The drop probability can be configured differently for each of multiple traffic classes.
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 emerging 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 create markup languages 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. Text markup language designed to enable the use of SGML on the World Wide Web. XML allows you to define your own customized markup language.
zone	In H.323, the 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 has only one gatekeeper. A zone can be independent of LAN topology and can be comprised of multiple LAN segments connected using routers or other devices.

B

Case Study Solutions

The lesson case study solutions are contained here. This is a presentation of all course case study items.

Module 1: Introducing Cisco Network Service Architectures

There is no case study for this module.

Module 2: Designing Enterprise Campus Networks

Case Study 2-3: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Create Initial Network Diagrams

Following are suggested solutions.

Step 1 On an overhead transparency, create a global network diagram for the company, to include the headquarters location, district offices, regional offices, and international plants. This network diagram should show the main sites in each country and the main WAN links. Label each location.

Refer to your network diagram.

Step 2 On an overhead transparency, create a country-level network diagram for the company that identifies the locations in North America. This network diagram shows the main sites in each town and the main WAN links. Label each location.

Refer to your network diagram.

Task 2: Design the Headquarters Campus Network

Following are suggested solutions.

Step 1 On an overhead transparency, create a campus network diagram for the company headquarters site. If desired, create a logical map showing the extent of each VLAN for the headquarters site. Your network diagram should include each building and the Campus Backbone submodule. Label each location.

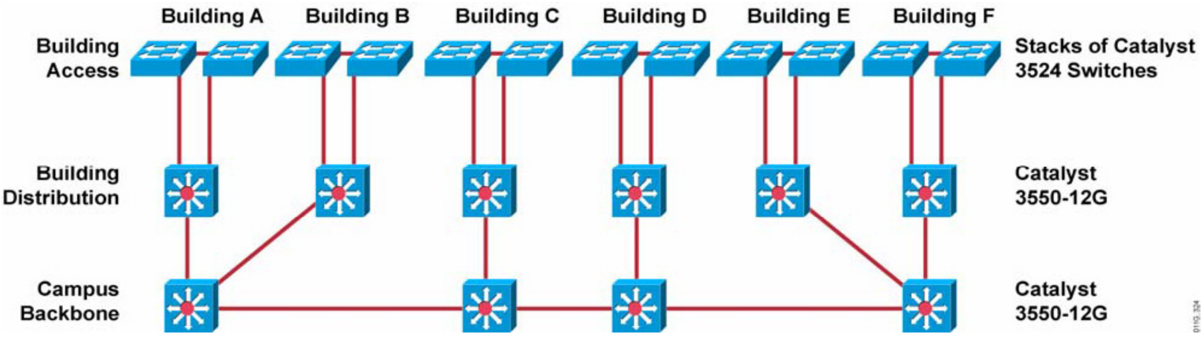

Step 3 Complete the table to design the details about your headquarters campus network.

The table summarizes one possible set of design decisions that meet the OCSIC Bottling Company's requirements.

Design Question	Decision	Justification
What is the logical network design?	Use logical Layer 2 segmentation by VLAN for each floor by department	The network is logically segmented based on departments.
What type of VLAN trunking will be used?	802.1Q	802.1Q is an industry standard that provides interoperability with other vendors' equipment
What physical network media will be used in the Campus Backbone submodule?	Multimode fiber in the Campus Backbone and building risers	The current cabling plant is outdated and will not support the Gigabit Ethernet requirements of the company. Multimode fiber provides the bandwidth required in the Campus Backbone submodule.
What physical network media will be used in the Building Distribution submodule?	Multimode fiber in the Building Distribution submodule	Multimode fiber provides the bandwidth required in the Building Distribution submodule.
What physical network media will be used in the Building Access submodule?	Category 5 cabling to the desktop	Category 5 cabling provides the bandwidth required in the Building Access submodule.
Which data link layer protocol will be used at each location?	Fast Ethernet from the switch to the desktops Gigabit Ethernet through the risers and between the floor switches	Fast Ethernet and Gigabit Ethernet provide the performance required.
What spanning-tree deployment and version will be used?	Spanning tree (RSTP/MST) will be used The associated Building Distribution switch will be the root	For simplicity, the Catalyst 3550 was selected as the STP root because it provides a logical break between the data link and network layers.
What is the data link layer/multilayer strategy for the Campus Backbone submodule?	Multilayer switched in the Campus Backbone submodule with 12 fiber pairs to each building	Multilayer switching in the Campus Backbone provides flexibility.
What is the data link layer/multilayer strategy for the Building Distribution submodule?	Multilayer switched in the Building Distribution submodule with 4 fiber pairs	Multilayer switching in the Building Distribution submodule provides flexibility.
What is the data link layer/multilayer strategy for the Building Access submodule?	Data link layer switched in the Building Access submodule with inline power	Data link layer switching provides performance and simplicity at the wiring closet.
Which Cisco products and options will be used in the Campus Backbone submodule?	Catalyst 3550-12G switches in the Campus Backbone submodule	The selected switches provide cost-effective solution, supporting effective performance, scalability, and availability.
Which Cisco products and options will be used in the Building Distribution submodule?	Catalyst 3550-12G switches in the Building Distribution module	The selected switches provide cost-effective solution, supporting effective performance, scalability, and availability.

Design Question	Decision	Justification
Which Cisco products and options will be used in the Building Access submodule?	Stacks of Catalyst 3524 switches in the Building Access submodule (24-port 10/100 with integrated inline power + two-port 1000Base-X, Enterprise Edition)	The selected switches provide cost-effective solution, supporting effective performance, scalability, and availability.
What IP addressing scheme will be used? Is NAT/PAT required?	Class B addresses (RFC1918) Private Class B address with NAT to the Internet	The company wants a private address space for ease of implementation. The class B addresses allow for simple segmentation on an eight bit boundary.
Which routing protocols will be used in each area of the network?	OSPF throughout the network	Each building is a separate area and area 0 is comprised of the Campus Backbone Catalyst 3550 switches. The interface from the Campus Backbone switch into each building is the OSPF boundary. Each building is a separate OSPF area to simplify management issues.
What type of switching will be deployed at the Edge Distribution module?	Multilayer switching	Multilayer switching in the Edge Distribution module provides flexibility.

Task 3: Design the Headquarters Server Farm

Following are suggested solutions.

Step 1 Create a Server Farm network diagram for the OCSIC Bottling Company data center. This network diagram shows the physical layout and how the data center relates to the Campus Backbone module. Label each device and location.

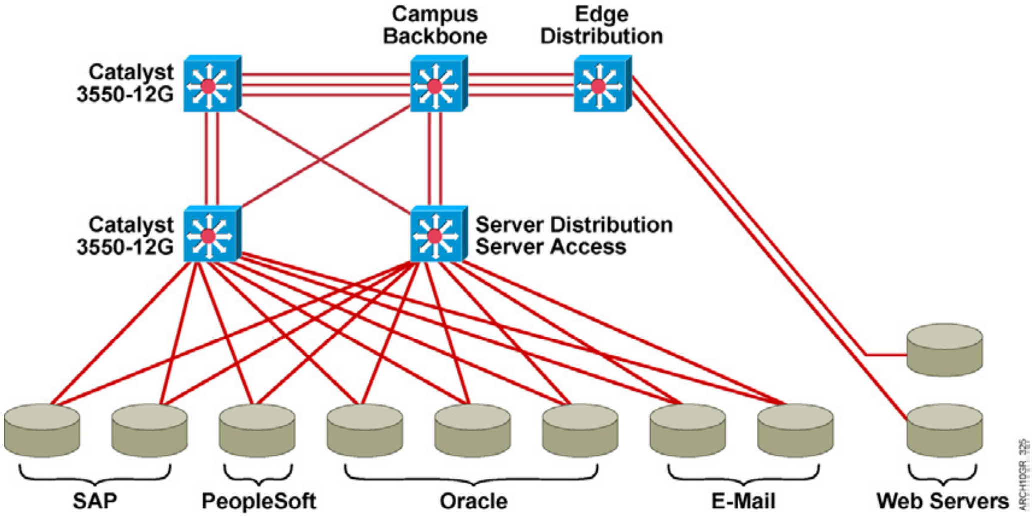

Step 2 Complete the table to design the details about your Server Farm network.

The table summarizes one possible set of design decisions that meet the OCSIC Bottling Company’s server farm requirements.

Design Question	Decision	Justification
What is the logical network design?	Use logical Layer 2 segmentation by VLAN for each type of server as follows: ■ Web server network: VLANFD31 ■ Server farm network: VLANFD32 ■ Edge distribution: VLANFD33	The network is logically segmented based on the application.
What type of VLAN trunking will be used?	802.1Q	802.1Q is an industry standard that provides interoperability with other vendors’ equipment
What physical network media will be used?	Multimode fiber for all links in the Server Farm module	The physical network media provides the bandwidth required at each location.
What data-link layer protocol will be used?	Gigabit Ethernet	Gigabit Ethernet provides the performance required.
What spanning-tree deployment will be used?	Spanning tree will be used, with a Server Distribution switch as the root	For simplicity with the rest of the network, the Catalyst 3550 is selected as the STP root.
What is the data link layer/multilayer strategy for the Server Distribution layer?	Multilayer switched in the Server Distribution layer	Multilayer switching in the Server Distribution layer provides flexibility.
What is the data link layer/multilayer strategy for the Server Access layer?	Data link layer switched in the Server Access layer with inline power	Data link layer switching provides performance and simplicity at the wiring closet.
Which Cisco products and options will be used in the Server Distribution layer?	Server Distribution and Server Access combined in one switch: Catalyst 3550-12T	The company wants to deploy a cost-effective solution that provides the optimal performance, scalability, and availability.
Which Cisco products and options will be used in the Server Access layer?	Server Distribution and Server Access combined in one switch: Catalyst 3550-12T	The company wants to deploy a cost-effective solution that provides the optimal performance, scalability, and availability.
What IP addressing scheme will be used? Is NAT/PAT required?	Class B addresses (RFC1918) Private Class B address with NAT to the Internet	The Server Farm module has its own set of subnets within the class B address range.
Which routing protocols will be used?	OSPF	The Server Farm module will be its own OSPF area.

Task 4: Design a Typical North American Plant Network (Optional)

Following are suggested solutions.

Step 1 Create a campus network diagram for a typical plant in North America. This network diagram shows the physical layout and the Campus Backbone module. Label each location or area of the building.

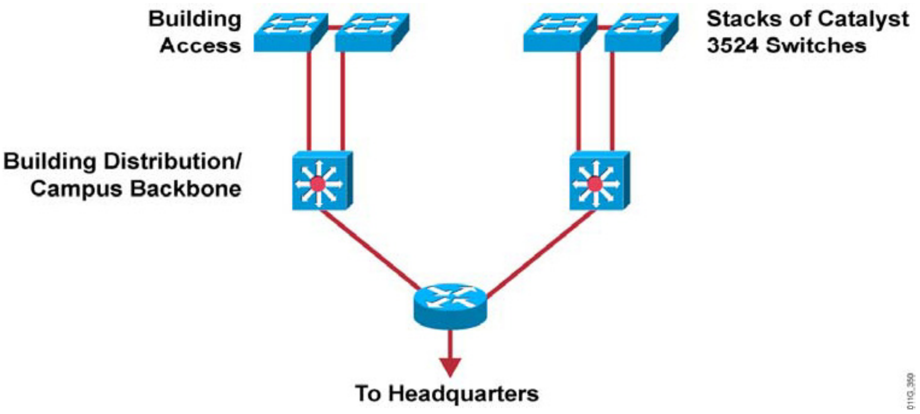
Building Access
Stacks of Catalyst 3524 Switches
Building Distribution/ Campus Backbone
To Headquarters
011G_350

Step 2 Complete the table to design the details about a typical North American plant network.

The table summarizes one possible set of design decisions that meet the OCSIC Bottling Company’s North American plant requirements.

Design Question	Decision	Justification
What is the logical network design?	Use logical Layer 2 segmentation by VLAN for each plant	The network is logically segmented based on a plant.
What type of VLAN trunking will be used?	802.1Q	802.1Q is an industry standard that provides interoperability with other vendors’ equipment
What physical network media will be used in the Campus Backbone submodule?	Given the collapsed Campus Backbone/Building Distribution submodule, multimode fiber will be used	Multimode fiber provides the bandwidth required in the Building Distribution submodule.
What physical network media will be used in the Building Distribution submodule?	Given the collapsed Campus Backbone/Building Distribution submodule, multimode fiber will be used	Multimode fiber provides the bandwidth required in the Building Distribution submodule.
What physical network media will be used in the Building Access submodule?	Category 5 cabling to the desktop	Category 5 cabling provides the bandwidth required in the Building Access submodule.
What data-link layer protocol will be used?	Fast Ethernet	Fast Ethernet provides the performance required.
What spanning-tree deployment will be used?	Spanning tree will be used The Campus Backbone/Building Distribution switches will be the root of the spanning tree	For simplicity, the Campus Backbone/Building Distribution switches are selected.

Design Question	Decision	Justification
What is the data link layer/multilayer strategy for the Campus Backbone submodule?	Multilayer switched out of each plant through a router to headquarters	Multilayer switching in the Campus Backbone submodule provides flexibility.
What is the data link layer/multilayer strategy for the Building Distribution submodule?	Included in the Campus Backbone submodule	Multilayer switching in the Building Distribution submodule provides flexibility.
What is the data link layer/multilayer strategy for the Building Access submodule?	Data link layer switched in the Building Access submodule with inline power	Data link layer switching provides performance and simplicity at the wiring closet.
Which Cisco products and options will be used in the Campus Backbone submodule?	Catalyst 3550-12G in the Campus Backbone submodule	The company deploys a cost-effective solution that provides the optimal performance, scalability, and availability.
Which Cisco products and options will be used in the Building Distribution submodule?	Included in the Campus Backbone submodule	The company deploys a cost-effective solution that provides the optimal performance, scalability, and availability.
Which Cisco products and options will be used in the Building Access submodule?	Stacks of Catalyst 3524 switches in the Building Access submodule (24-port 10/100 with integrated inline power + two-port 1000Base-X, Enterprise Edition)	The company deploys a cost-effective solution that provides the optimal performance, scalability, and availability.
What IP addressing scheme will be used? Is NAT/PAT required?	Class B addresses (RFC1918) Private Class B address with NAT to the Internet	Each regional and branch office has their own class C address from the corporate Class B addresses for ease of address manipulation.
Which routing protocols will be used?	OSPF	Each plant is a separate OSPF area to simplify management issues.

Module 3: Designing Enterprise Edge Connectivity

Case Study 3-4: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design the Wide Area Network

Following are suggested solutions.

Step 1 Refer to the global network diagram for the company that includes the headquarters location, district offices, regional offices, and international plants.

Refer to your network diagram.

Step 2 Refer to the country-level network diagram for the company that identifies the locations in North America and the WAN links between the locations. Make a copy of the diagram for this exercise.

Refer to your network diagram.

Step 3 Complete the table to design the details about the WAN. Assume that the service provider selected offers all of the popular wide-area networking services available on the market today, and their service level agreement is acceptable.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
What topology will be used for the WAN?	Should be a simple hub and spoke configuration unless planning for any redundancy	The simple hub and spoke topology works with the headquarters hub connecting to the district plants, and the district plant hubs connecting to the regional plants.
What service provider will be selected?	National carrier	A national carrier is required to provide geographical coverage.
What data-link layer protocol will be used?	Frame Relay	Frame Relay provides the bandwidth required at an effective price point.
What physical network media will be used?	Headquarters: T1 (see notes following the table) Each district plant: T1 (see notes following the table) Each regional plant: T1 (see notes following the table)	See notes following the table
What additional services would you select for each WAN link? ■ If you selected Frame Relay, choose the number of ports, committed information rate (CIR), committed burst (Bc), excess burst (Be), transmission convergence (TC), and maximum burst size. ■ If you selected ATM, choose the service class, either CBR, ABR, UBR, RT-VBR, or NRT-VBR ■ If you selected PPP, the services depend on the Layer 1 technology you selected.	The company is implementing the following features: ■ Permanent virtual circuits ■ Quality of service ■ BECN/FECN acknowledgement	The features provide the quality of service that the company requires.
Which Cisco products will be used?	Headquarters: Cisco 3640 District plants: Cisco 1760 Regional plants: Cisco 1750	Each product provides the capacity and features required.
Which routing protocols will be used?	OSPF hierarchical design	Given the number of sites and the way the design leads to an easy division of areas, OSPF is chosen.
What IP addressing scheme will be used?	Access to the Internet and NAT are required	A single Class C address provides Internet connectivity. A Class B is required used internally and NAT is used outside the corporate network.

Note **Headquarters** requires **two T1 access lines** to the Frame Relay provider. Each T1 has a two Frame Relay PVCs, each to a district office in a different time zone, at 640 Kbps CIR with a committed burst up to 768 kbps and an excess burst up to 1024 kbps.

Eastern district requires **two T1 access lines** to the Frame Relay provider. There are three PVCs provisioned: one to headquarters with a 640 kbps CIR with a committed burst up to 768 kbps and an excess burst up to 1024 kbps on a dedicated T1, and two to the associated regional offices at a CIR of 512 Kbps with a committed burst up to 640 kbps and an excess burst up to 768 kbps on the other T1.

Midwestern district requires **two T1 access lines** to the Frame Relay provider. There are three PVCs provisioned. one to headquarters with a 640 kbps CIR with a committed burst up to 768 kbps and an excess burst up to 1024 kbps on a dedicated T1, and two to the associated regional offices at a CIR of 512 kbps with a committed burst up to 640 kbps and an excess burst up to 768 kbps on the other T1.

Southern district requires **one T1 access line** to the Frame Relay provider. There are two PVC provisioned: one to headquarters with a 640 kbps CIR with a committed burst up to 768 kbps and an excess burst up to 1024 kbps, and one to the associated regional office at a CIR of 512 kbps with a committed burst up to 640 kbps and an excess burst up to 768.

Western district requires **2 T1 access lines** to the Frame Relay provider. There are four PVC provisioned: one to headquarters with a 640 Kbps CIR with a committed burst up to 768 Kbps and an excess burst up to 1024 kbps, and three to the associated regional offices at a CIR of 512 kbps with a committed burst up to 640 kbps and an excess burst up to 768. The PVC to headquarters and one west coast regional office share a T1, and the two PVC to the other regional offices share the other T1.

Each **regional office** has a single **768 kbps fractional T1 access line** to the Frame Relay provider. There is a single PVC at a CIR of 512 kbps with a committed burst up to 640 kbps and an excess burst up to 768.

Step 4 Update the WAN network diagram to indicate the products and services you selected at each location.

The figure shows an example design for the Edge Distribution module and the WAN module, with connections to the district offices.

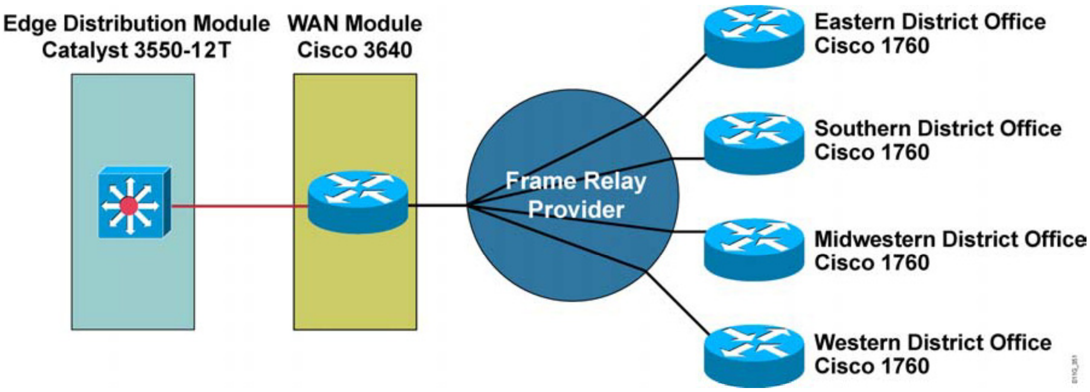

The figure shows an example design for the WAN connections to the regional offices.

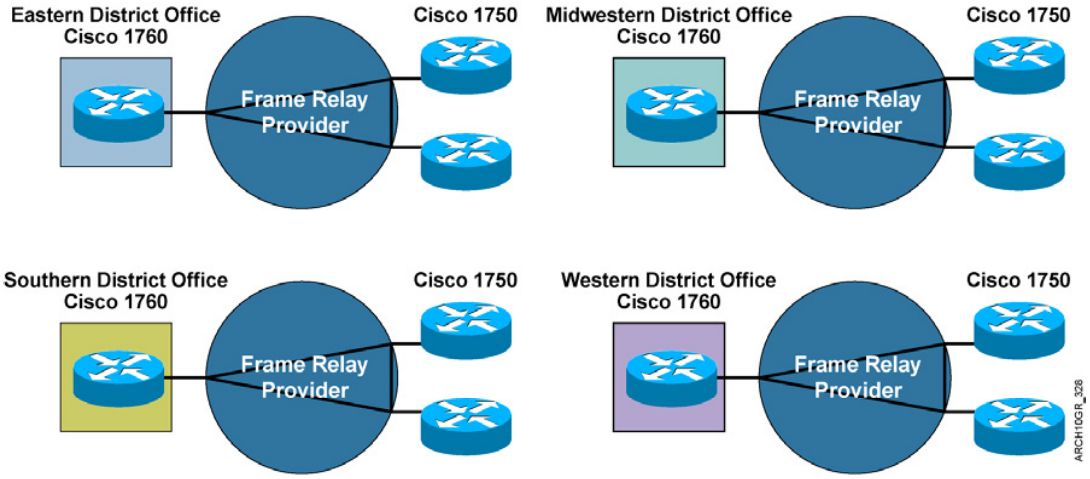

Task 2: Design the Remote-Access Network

Following are suggested solutions.

Step 1 Complete the table to design the details about the remote-access network. Assume that the service provider selected offers all of the popular remote-access services available on the market today, and their service level agreement is acceptable.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What topology will be used for the WAN?	Centralized	ISDN is deployed. Each user authenticates with the access servers at the hub to use the services in the network.
What service provider will be selected?	Local phone company	The local phone company provides the services and geographic coverage required.
What data-link layer protocol will be used?	ISDN Dial-up	ISDN provides bandwidth for users with access to ISDN. Dial-up provides access for other remote users.
What physical network media will be used? How many trunks are required?	T1 with 5 PRIs	The physical media meets the need for 120 ports to support remote users.
Which Cisco products will be used?	Cisco AS5350 with 120 ports at headquarters Client devices require appropriate client software for dial-up and applications	Each product provides the capacity and features required.
Which routing protocols will be used?	Not applicable	
What IP addressing scheme will be used?	Dynamic Host Configuration Protocol (DHCP) is used to dynamically assign addresses	Automatic IP addressing simplifies administration of the remote locations. A Class C address derived from the Class B address is dedicated to remote access.

Step 2 Update the WAN network diagram to indicate the products and services you selected at each location.

The figure shows an example design for the Edge Distribution module and the Remote Access module.

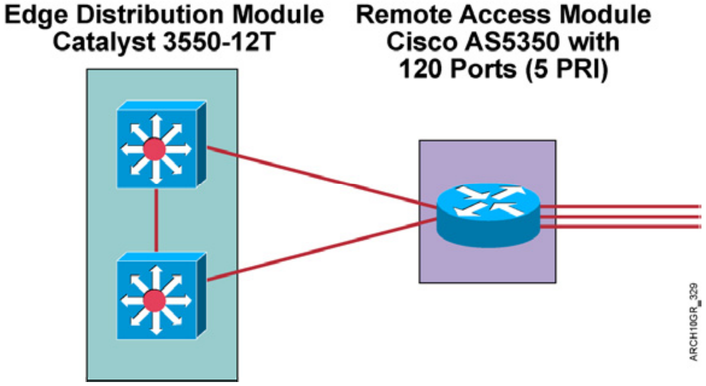

Step 3 Is authentication required for remote users? Why or why not?

Yes. Remote users should authenticate before accessing network resources.

Step 4 Are access control lists required for remote users? Why or why not?

Yes. Access control lists identify which users have access to specified network resources.

Step 5 Are firewalls or intrusion detection systems required? Why or why not?

Firewalls are required because some traffic will traverse the public network.

Task 3: Design the Internet Connectivity Module

Following are suggested solutions.

Step 1 Complete the table to design the details about the Internet Connectivity module network. Assume that the service provider selected offers all of the popular WAN and Internet services available on the market today, and their service level agreement is acceptable.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What topology will be used for the WAN?	Point-to-point	Point-to-point provides the link required.
What service provider will be selected?	Local Internet service provider	The local phone company provides the services and geographic coverage required.
What data-link layer protocol will be used?	Point-to-point protocol (PPP)	Along with the data link layer protocol, the following features are implemented: ■ Cisco IOS Firewall ■ Access lists
What physical network media will be used?	DSL	DSL provides a point-to-point link to the service provider at a cost-effective price point.
Which Cisco products will be used?	Catalyst 3350 in the Edge Distribution module Cisco 1760 to the Internet	Each product provides the capacity and features required.
Which routing protocols will be used?	Default routes throughout the OSPF network that point to the Cisco 1760 as the Internet gateway	
What IP addressing scheme will be used?	Two class C networks NAT addressing	Two Class C networks are acquired from the provider to use with NAT. NAT provides public and private addressing.

Step 2 Update the WAN network diagram to indicate the products and services you selected at each location.

The figure shows an example design for the Edge Distribution module and the Internet Connectivity module.

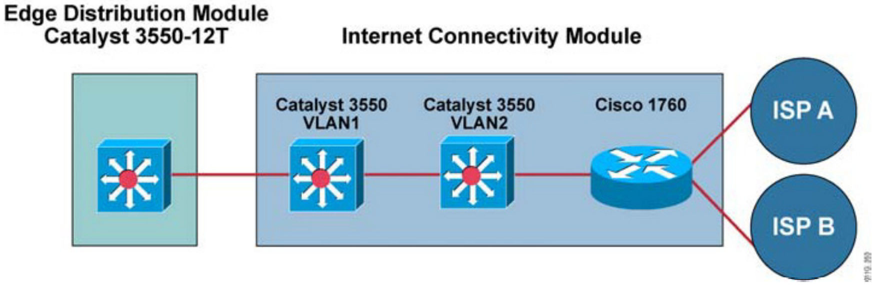

Module 4: Designing Network Management Services

Case Study 4-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Develop a Network Management Strategy for the Company

Following are suggested solutions.

Step 1 Complete the table to design the details about the Network Management solution for the OCSIC Bottling Company.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
How many management domains does the enterprise require?	1 domain for the company	The domain is managed centrally.
How many devices need to be managed?	Fewer than 500 network devices and fewer than 50 servers	
What are the key components and functions required?	LAN Management Solution Routed WAN Management Solution	
How many servers are required?	Two servers at headquarters with: ■ LAN Management Solution ■ Routed WAN Management Solution One LMS server at each district office	Two servers are required at headquarters to manage the network. The LMS servers at the district offices manage the regional offices.
What administrative grouping of network devices will work for this enterprise?	All network devices within a single administrative grouping	Given the size of the network, only one administrative grouping of network devices is required.
What administrative grouping of network management users will work for this enterprise?	All management users within a single administrative grouping	Given the size of the network, only one administrative grouping of management users is required.

Step 2 Describe policies and procedures to implement for the enterprise network management.

The company implements a proactive network management strategy that incorporates the FCAPS model for network management.

Step 3 What type of polling will you deploy on the network? How will polling affect network performance? How long would you recommend to maintain polling information?

Periodic polling of only key managed network devices is implemented in an effort to minimize the effect on the network. Polling information is stored for 7 days and then purged to minimize storage requirements.

Step 4 Update your campus network diagram to indicate the components of the Network Management module.

The figure shows an example network management design for the OCSIC Bottling Company.

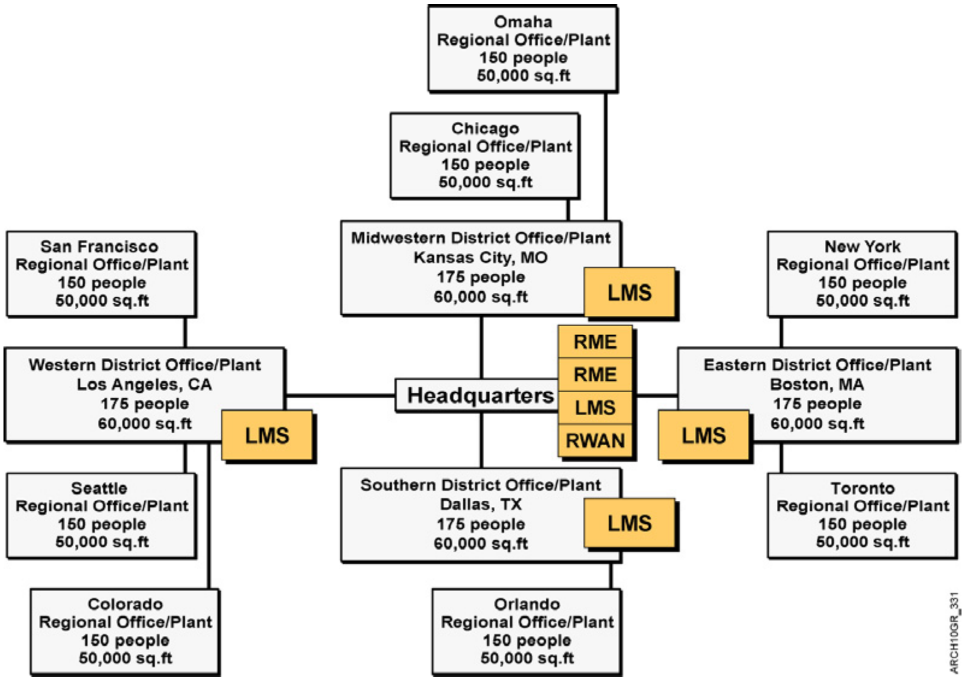

Omaha Regional Office/Plant 150 people 50,000 sq.ft
Chicago Regional Office/Plant 150 people 50,000 sq.ft
San Francisco Regional Office/Plant 150 people 50,000 sq.ft
Midwestern District Office/Plant Kansas City, MO 175 people 60,000 sq.ft
LMS
New York Regional Office/Plant 150 people 50,000 sq.ft
Western District Office/Plant Los Angeles, CA 175 people 60,000 sq.ft
LMS
Headquarters
RME
RME
LMS
RWAN
Eastern District Office/Plant Boston, MA 175 people 60,000 sq.ft
LMS
Seattle Regional Office/Plant 150 people 50,000 sq.ft
Southern District Office/Plant Dallas, TX 175 people 60,000 sq.ft
LMS
Toronto Regional Office/Plant 150 people 50,000 sq.ft
Colorado Regional Office/Plant 150 people 50,000 sq.ft
Orlando Regional Office/Plant 150 people 50,000 sq.ft
ARCH10GR_331

Module 5: Designing High-Availability Services

Case Study 5-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Develop a High-Availability Strategy for the Headquarters Campus Network

Following are suggested solutions.

Step 1 Complete the table to design a high availability solution for the headquarters campus network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
Which devices should be fault tolerant?	None	It is deemed not cost-effective to add fault tolerant devices in the campus network.
Which devices should be redundant?	For the Catalyst 3550-12G Building Distribution switches, include a Catalyst 4006 with 2 Supervisor IIIs, and 2 8-port GB Ethernet (4908G) For every Catalyst 3550-12G in the design, a second 3550-12G switch is added to provide device redundancy	Device redundancy provides high availability as needed in the campus network.
Which links should be redundant?	Catalyst 3524 stacks have redundant links to the Building Distribution switches	Redundant links provide backup in case of a link failure.
What spanning-tree implementation and root devices are required?	Spanning-tree root at the Building Distribution switches using RSTP/MST	For simplicity, the Building Distribution is used as the STP root because it provides a logical break between the data link and network layers.
What is the router availability strategy?	HSRP	HSRP implemented in the multilayer switches provides high availability.

Task 2: Develop a High-Availability Strategy for the Server Farm Module

Following are suggested solutions.

Step 1 Complete the table to design a high availability solution for the Server Farm module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
Which devices should be fault tolerant?	All devices	Fault tolerance is critical in the Server Farm.
Which devices should be redundant?	None	Fault tolerance is preferred to device redundancy in the Server Farm
Which links should be redundant?	Redundant links throughout the Server Farm module	Redundant links are required for high availability.
What spanning-tree implementation and root devices are required?	Spanning-tree root at the Server Distribution switches using RSTP/MST	For simplicity, the Server Distribution is used as the STP root because it provides a logical break between the data link and network layers.
What is the router availability strategy?	HSRP	HSRP implemented in the multilayer switches provides high availability.

Task 3: Develop a High-Availability Strategy for the WAN Module

Following are suggested solutions.

Step 1 Complete the table to design a high availability solution for the WAN module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
Which devices should be fault tolerant?	None	Fault tolerance is not cost-effective in the WAN module.
Which devices should be redundant?	Cisco 3640 WAN router includes a second device for WAN redundancy	The second Cisco 3640 WAN router provides the necessary high availability for the WAN module.
Which links should be redundant?	Redundant links to the Edge Distribution module	Redundant links provide backup in case of a link failure.
What spanning-tree implementation and root devices are required?	None	Not applicable
What is the router availability strategy?	HSRP will run on the Cisco 3640 routers in the WAN module	HSRP provides high availability.

Task 4: Develop a High-Availability Strategy for the Remote Access Module

Following are suggested solutions.

Step 1 Complete the table to design a high availability solution for the Remote Access module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
Which devices should be fault tolerant?	None	Fault tolerance is not cost-effective in the Remote Access module.
Which devices should be redundant?	None	Device redundancy is not cost-effective in the Remote Access module.
Which links should be redundant?	Redundant links to the Edge Distribution module	Redundant links provide backup in case of a link failure.
What spanning-tree implementation and root devices are required?	None	Not applicable
What is the router availability strategy?	HSRP	HSRP provides high availability.

Task 5: Develop a High-Availability Strategy for the Internet Connectivity Module

Following are suggested solutions.

Step 1 Complete the table to design a high availability solution for the Internet Connectivity module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Question	Decision	Justification
Which devices should be fault tolerant?	None	Fault tolerance is not cost-effective in the Internet Connectivity module.
Which devices should be redundant?	None	Device redundancy is not cost-effective in the Internet Connectivity module.
Which links should be redundant?	Redundant links to the Edge Distribution module	Redundant links provide backup in case of a link failure.
What spanning-tree implementation and root devices are required?	None	Not applicable
What is the router availability strategy?	HSRP	HSRP provides high availability.

Step 2 Update your network diagrams to reflect your high-availability strategy.

The figure shows a revised network diagram for the headquarters location with high availability services.

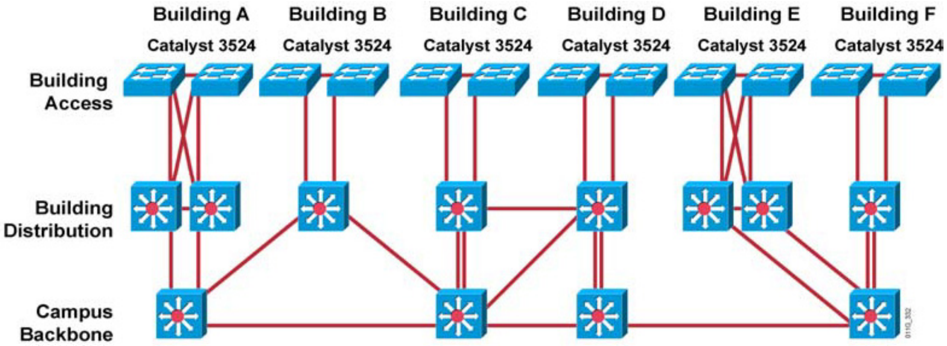

The figure shows a network diagram for the wide area network with redundant links for load sharing and high availability.

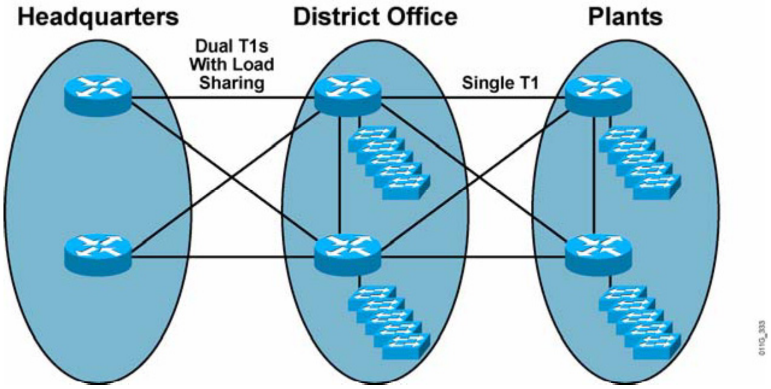

Module 6: Designing Security Services

Case Study 6-3: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Develop a Security Policy for the Network

Following are suggested solutions.

Step 1 Propose a comprehensive security policy for the company network.

Your security policy should contain these elements:

- Definition: All corporate data and devices will be covered by this policy.
- Identity: Hosts and applications must be authorized to access the network.
- Trust: A multilevel trust system, based on level within the organization, will define the conditions under which a user is trusted to perform an action.
- Enforceability: Hardware and software features will be used to enforce the security policy.
- Risk assessment: The risk assessment identifies all assets within the corporation and assigns a relative risk.
- Incident response: All incidents will be handled based on a hierarchy that defines severity of the infraction.

Task 2: Develop a Security Design for the Headquarters Campus Network

Following are suggested solutions.

Step 1 Complete the table to design your security solution for the headquarters campus network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?	None	Firewalls are not required on the campus network.
What is your intrusion detection strategy? What features would you implement?	None	Intrusion detection is not required on the campus network.
What software features would you implement?	Authentication Host-based virus scanning	Authentication is provided for network device access. Host-based virus scanning prevents most viruses and many Trojan horses.

Task 3: Develop a Security Design for the Server Farm Module

Following are suggested solutions.

Step 1 Complete the table to design your security solution for the Server Farm module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?	None	Firewalls are not necessary in the Server Farm.
What is your intrusion detection strategy? What features would you implement?	Use host intrusion protection systems (HIPS)	HIPS mitigates unauthorized access. Operating systems, devices, and applications are kept up to date with the latest security fixes and protected by HIPS. HIPS prevents port redirection agents from being installed.
What software features would you implement?	Implement AAA security with RADIUS authentication RFC 2827 filtering	RFC 2827 filtering prevents source address spoofing.

Task 4: Develop a Security Design for the WAN Module

Following are suggested solutions.

Step 1 Complete the table to design your security solution for the WAN module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?	Implement a firewall Within the firewall, implement NAT and access control features	Firewalls prevent unauthorized access. TCP setup controls at the firewall limit denial of service attacks.
What is your intrusion detection strategy? What features would you implement?	An Intrusion Detection System installed on both sides of the firewall Include IDS monitoring on the inside network and outside network Use the software application built into the firewall to parse log files	Intrusion detection mitigates application layer attacks, password attacks, attacks by packet sniffers, reconnaissance, and port redirection.
What software features would you implement?	Authentication provided for WAN access	Authentication ensures that only authorized users have access to network resources.

Task 5: Develop a Security Design for the Remote Access Module

Following are suggested solutions.

Step 1 Complete the table to design your security solution for the Remote Access module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?	Implement a firewall Within the firewall, implement NAT and access control features	Firewalls prevent unauthorized access. TCP setup controls at the firewall limit denial of service attacks.
What is your intrusion detection strategy? What features would you implement?	An Intrusion Detection System installed on both sides of the firewall Include IDS monitoring on the inside network and outside network Use the software application built into the firewall to parse log files	Intrusion detection mitigates application layer attacks, password attacks, attacks by packet sniffers, reconnaissance, and port redirection.
What software features would you implement?	Authentication provided for remote access IP spoofing security runs inside the Cisco AS5350	Authentication ensures that only authorized users have access to network resources. IP spoofing detects unwanted guests.

Task 6: Develop a Security Design for the Internet Connectivity Module

Following are suggested solutions.

Step 1 Complete the table to design your security solution for the Internet Connectivity module.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?	Implement a firewall Within the firewall, implement NAT and access control features	Firewalls prevent unauthorized access. TCP setup controls at the firewall limit denial of service attacks.
What is your intrusion detection strategy? What features would you implement?	An Intrusion Detection System installed on both sides of the firewall Include IDS monitoring on the inside network and outside network Use the software application built into the firewall to parse log files	Intrusion detection mitigates application layer attacks, password attacks, attacks by packet sniffers, reconnaissance, and port redirection.
What software features would you implement?	Authentication provided for Internet access Only HTTP access to the Internet is provided Only web-based traffic is allowed to and from the Internet	Authentication ensures that only authorized users have access to network resources.

Step 2 Update your campus network diagram to reflect your security design.

The figure shows a security design for the company network.

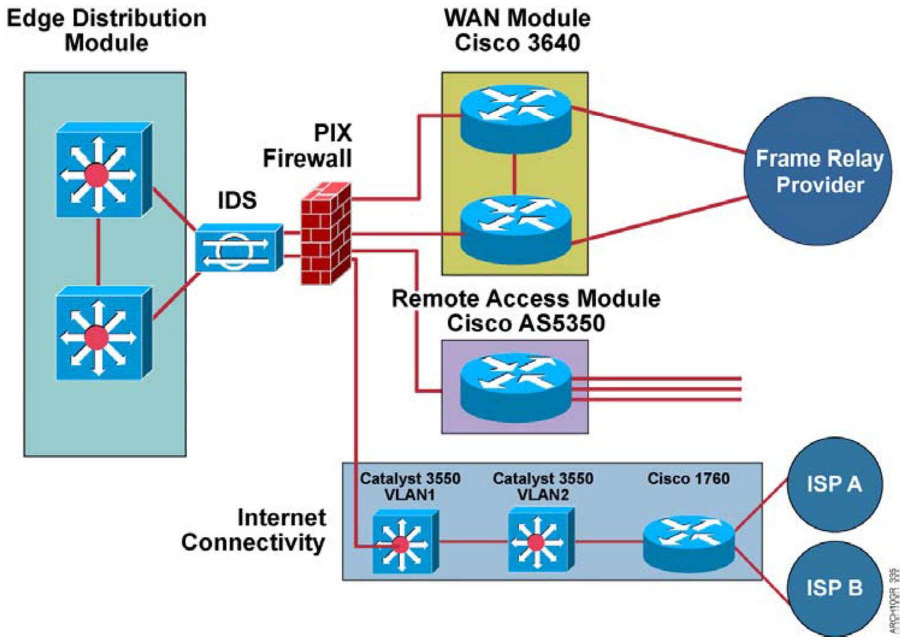

Module 7: Designing QoS

Case Study 7-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Develop a QoS Design for the Site-to-Site WAN

Following are suggested solutions.

Step 1 Complete the table to design your QoS solution.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What classification and marking tools and settings are required?	Each type of traffic is classified as close to the source as possible with marking as follows: ■ Voice packets will be marked with a precedence of 5, class of service (CoS) of 5 ■ Signaling traffic will be marked with a precedence of 4, CoS of 4 ■ Mission-critical data will be marked with a precedence of 4, CoS 4 ■ Important data will be marked with a precedence of 3, CoS 3 ■ Corporate data will be marked with a precedence 2, CoS 2 ■ All other traffic will be marked with a precedence of 0, CoS 0 ■ Classification for the data is done via CAR utilizing access lists to determine which traffic goes in which queue	Packets are classified and then scheduled on the egress from every device that forwards traffic. The WAN devices are more critical than the LAN devices because of the bandwidth restrictions.
What congestion avoidance tools and settings are required?	Use Weighted Random Early Detection (WRED)	WRED is used to drop packets from lower priority data queues.
What congestion management tools and settings are required?	Use low latency queuing (LLQ) with alternate priority is used as the scheduling mechanism Use three queues to support data traffic	SAP/Oracle goes in queue 3; PeopleSoft and e-mail goes in queue 2; and intranet traffic goes in queue 1. All other traffic goes in default queue 0). The company uses a 25, 50, and 75 packet threshold for dropping packets out of the queues with a CoS of 0, 1, and 2.

Design Questions	Decision	Justification
What traffic conditioning tools and settings are required?	Use Frame Relay traffic shaping	Frame Relay traffic shaping is used to ensure that the far ends can handle the amount of traffic sent to them from the headquarters site.
What signaling tools and settings are required?	Classify the voice signaling as a precedence 4, CoS 4	The precedence values ensure that voice receives priority on the network.
What link efficiency tools and settings are required?	No link efficiency to the district office Use Frame Relay traffic shaping from the district plants to the regional offices	Because all of the links to the district offices are T1, the solution does not need LFI going to the districts. From the district offices to the regional plants, Frame Relay traffic shaping will be used. The MTU to the remote plants will be set to 512 bytes.

Task 2: Develop a QoS Design for the Campus Network

Following are suggested solutions.

Step 1 Complete the table to design your QoS solution.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What classification and marking tools and settings are required?	Each type of traffic is classified as close to the source as possible with marking as follows: ■ Voice packets will be marked with a precedence of 5, class of service (CoS) of 5 ■ Signaling traffic will be marked with a precedence of 4, CoS of 4 ■ Mission-critical data will be marked with a precedence of 4, CoS 4 ■ Important data will be marked with a precedence of 3, CoS 3 ■ Corporate data will be marked with a precedence 2, CoS 2 ■ All other traffic will be marked with a precedence of 0, CoS 0 ■ Classification for the data is done via CAR utilizing access lists to determine which traffic goes in which queue	Packets are classified and then scheduled on the egress from every device that forwards traffic. The WAN devices are more critical than the LAN devices because of the bandwidth restrictions.
What congestion avoidance tools and settings are required?	Use Weighted Random Early Detection (WRED)	WRED is used to drop packets from lower priority data queues.

Design Questions	Decision	Justification
What congestion management tools and settings are required?	Use low latency queuing (LLQ) with alternate priority is used as the scheduling mechanism Use three queues to support data traffic	SAP/Oracle goes in queue 3; PeopleSoft and e-mail goes in queue 2; and intranet traffic goes in queue 1. All other traffic goes in default queue 0). The company uses a 25, 50, and 75 packet threshold for dropping packets out of the queues with a CoS of 0, 1, and 2.
What traffic conditioning tools and settings are required?	None	Traffic conditioning is not needed on the campus network.
What signaling tools and settings are required?	None	Signaling is not needed on the campus network.
What link efficiency tools and settings are required?	None	Link efficiency is not needed on the campus network.

Module 8: Designing IP Multicast Services

Case Study 8-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note	Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Develop an IP Multicast Design

Following are suggested solutions.

Step 1 Complete the table to design your IP multicast solution.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
Where will you implement multicast applications on the network?	Data applications are unicast applications Voice music on hold and video are multicast applications	
What IP multicast control mechanism will you use?	Cisco Group Management Protocol (CGMP) is implemented on the switches Protocol Independent Multicast-Sparse Mode (PIM-SM) and Internet Group Management Protocol (IGMP) version 3 are implemented on the routers Non RPF traffic is routed via the rules of the routing table.	
Will you use PIM-DM or PIM-SM? If you selected PIM-SM, determine the location of rendezvous points.	Use PIM-SM with the rendezvous point at the Enterprise Edge router	
What security is needed to support IP multicasting on the network?	Connections to the district and plants are a “private” network	Security depends on how much the service provider is trusted. If need be, encrypted tunnels could be created to send the traffic across the WAN.
Which network devices require upgrade (memory, IOS version, new hardware) to support the IP multicast applications?	Ensure that the proper IOS is selected to support all of the features for multicasting and the other features	Since video is being added the load will not diminish on any of the systems. It will actually increase to support the video streams.

Module 9: Designing Virtual Private Networks

Case Study 9-3: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design a Site-to-Site VPN Solution Between the Headquarters and Each International Plant

Following are suggested solutions.

Step 1 Complete the table to design your site-to-site VPN solution.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What topology will be used for the site-to-site VPN?	Hub-and-spoke topology	The corporate LAN is the hub and each individual site becomes a spoke.
What type of tunneling will be deployed?	GRE tunnels	Each International site creates a GRE tunnel to the headquarters site and then encrypts the information through IPSec tunnel mode.
What type of security will be deployed?	IPSec tunnels Authentication with TACACS+ server	The information is encrypted through the IPSec tunnel. The site-to-site tunnels are authenticated through a separate TACACS+ server.
Is NAT required?	No	All addresses are public and the remote users are not allowed access to the Internet through the corporate network.
What VPN hardware will be used?	Central site: Cisco 3030 VPN Concentrator Remote sites: Cisco 1740 or larger	The Cisco 3030 VPN Concentrator handles both the site-to-site and remote site VPN access. The remote site routers support VPN tunnel end-points.
What type of high availability will be deployed?	None, except services offered by the ISP	The only resiliency incorporated into the design is that the ISP offers the VPN service and has many numbers that users can call to access their global system. If the corporate network fails, there is no corporate access.

Step 2 Update your global network diagram to reflect your VPN strategy.

Refer to your network diagram.

Task 2: Design a Remote-Access VPN Solution for U.S.-Based Telecommuters to the Headquarters Location

Following are suggested solutions.

Step 1 Complete the table to design your remote-access VPN solution.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What topology will be used for the remote-access VPN?	Hub-and-spoke topology	The corporate LAN is the hub and each individual user becomes a spoke.
What type of tunneling will be deployed?	L2TP tunnels	Each remote user will create an L2TP tunnel to the headquarters site and then encrypt the information through IPSec tunnel mode.
What type of security will be deployed?	IPSec tunnels Authentication with TACACS+ server	The information is encrypted through the IPSec tunnel. The site-to-site tunnels are authenticated through a separate TACACS+ server.
Is NAT required?	Provided by VPN concentrator	The VPN concentrator will perform any NAT functions needed.
What VPN hardware will be used?	Central site: Cisco 3030 VPN Concentrator Cisco VPN client for remote users	The Cisco 3030 VPN Concentrator handles both the site-to-site and remote site VPN access.
What type of high availability will be deployed?	None, except services offered by the ISP	The only resiliency incorporated into the design is that the ISP offers the VPN service and has many numbers that users can call to access their global system. If the corporate network fails, there is no corporate access.

Step 2 Update your global network diagram to reflect your VPN strategy.

Refer to your network diagram.

Module 10: Designing Enterprise Wireless Networks

Case Study 10-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design a Wireless Network for a North American Plant

Following are suggested solutions.

Step 1 On an overhead transparency, create a campus network diagram indicating your wireless LAN design for one of the North American plants. Label each location.

The figure shows a wireless LAN design for the North American plants.

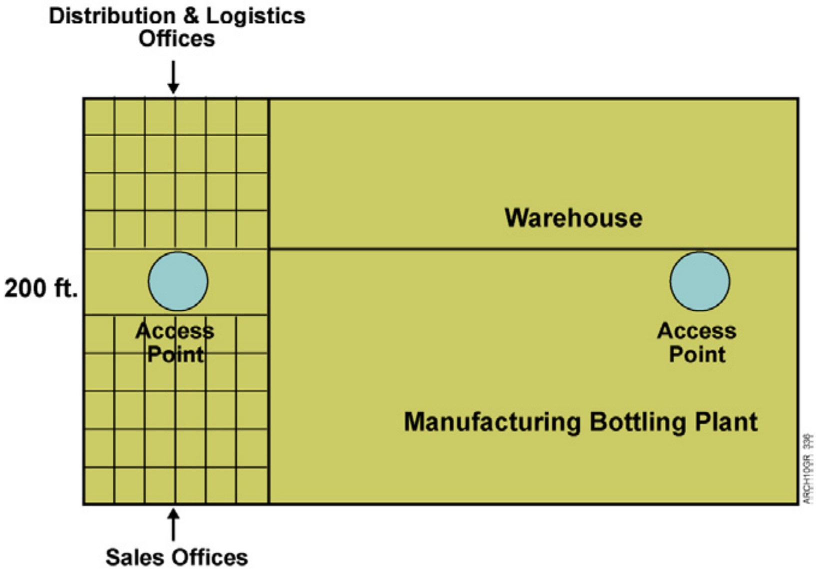

Step 2 Complete the table to design the details about the wireless network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
How many access points are required within a typical OCSIC 60,000-square-foot district office or plant? Where should the access points be placed?	Within a typical OCSIC 60,000 sq. foot district office/plant, the company will place two access points, located equidistant across the facility. Access points are placed without major obstruction between the access points and the wireless devices that use the access points. The Cisco 1200 access point is selected because it can deliver 802.11b and will support 802.11a.	Coverage of the sales offices and manufacturing/bottling plant is most important.
How many active devices can each access point support?	If there are 24 users, each user gets approximately 280 kbps worth of bandwidth. If there are more active users, each user gets less bandwidth.	The number of users varies depending on the amount of bandwidth each user needs.
How are channels identified for the design?	There are two access points in each building, each with its own channel.	Currently there are plans to put two access points in each building that will have some overlap but not total overlap. The corners of the building may receive weak signals. Each access point can broadcast about 130 feet indoors.
How will you meet the inline power requirements for the design?	The Cisco 1200 access point accepts inline power or can be powered from a power brick. Preferred power will be through the inline power.	
What is the high-availability (redundancy) strategy for the design?	There will be no redundancy for the wireless design.	The wireless network is not considered to be mission-critical. If one access point fails, users can access the second one.

Module 11: Designing IP Telephony Solutions

Case Study 11-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design an IP Telephony Network for OCSIC Bottling Company

Following are suggested solutions.

Step 1 Complete the table to design the details about the IP telephony network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
Is centralized or distributed call processing most appropriate?	Distributed call processing	Distributed call processing enables local call processing, providing the same level of features and capabilities whether the IP WAN is available or not.
How many CallManager servers will you deploy? In what locations? Which server will be the publisher? Which servers will be the subscribers? What other servers would you include in the cluster?	The company selected a multisite telephony solution for the headquarters campus location. The solution includes the following components: ■ Centralized Cisco CallManager ■ Eight Cisco CallManager servers including four subscribers, two backups, one publisher, and one TFTP server	The CallManager servers can provide the call processing capability and back-up features required for a network of this size.
What gateways will you deploy? Where will the gateways be located? What function will each gateway serve?	Each district and plant has a voice-enabled router with a two-port multiflex trunk card Headquarters uses Catalyst 6500 T1 ports for access to the PSTN	Each voice-enabled router routes traffic between the campus voice network and WAN facilities.
What QoS strategy will you deploy to support the solution?	Implement QoS at the campus and enterprise edge Voice will be prioritized over all other traffic priority 5. Voice signaling traffic will be given a priority of 4.	Classification and marking, congestion avoidance, congestion management, traffic conditioning, signaling, and link-efficiency mechanisms are all required.
What DSP resources are required for the solution?	Each voice-enabled router has the appropriate DSPs to terminate 24 voice calls Each T1 card on the headquarters Catalyst 6500 switches has DSPs	DSPs provide the appropriate resources to terminate 24 G.711 calls.

Design Questions	Decision	Justification
What transcoding resources are required for the solution?	Conferencing is done from the Catalyst 6500 located in the headquarters campus network. Transcoding resources are located in the headquarters campus network.	Transcoding supports conferencing across the headquarters campus network.
What are the network bandwidth and traffic engineering considerations?	All district sites have a T1 to the PSTN All regional sites have a T1 to the PSTN The corporate office has 14 T1s to the PSTN (288 trunks for 5700 people) Voice will use G.729 codecs in the WAN, and G.711 codecs for calls within each site All devices can use either G.711 or G.729 as requested	The traffic engineering analysis incorporates sufficient bandwidth for the anticipated call volume.

Step 2 Update your network diagrams to indicate the location of each CallManager server at headquarters and at the North American plants.

The figure shows a CallManager design for the North American plants.

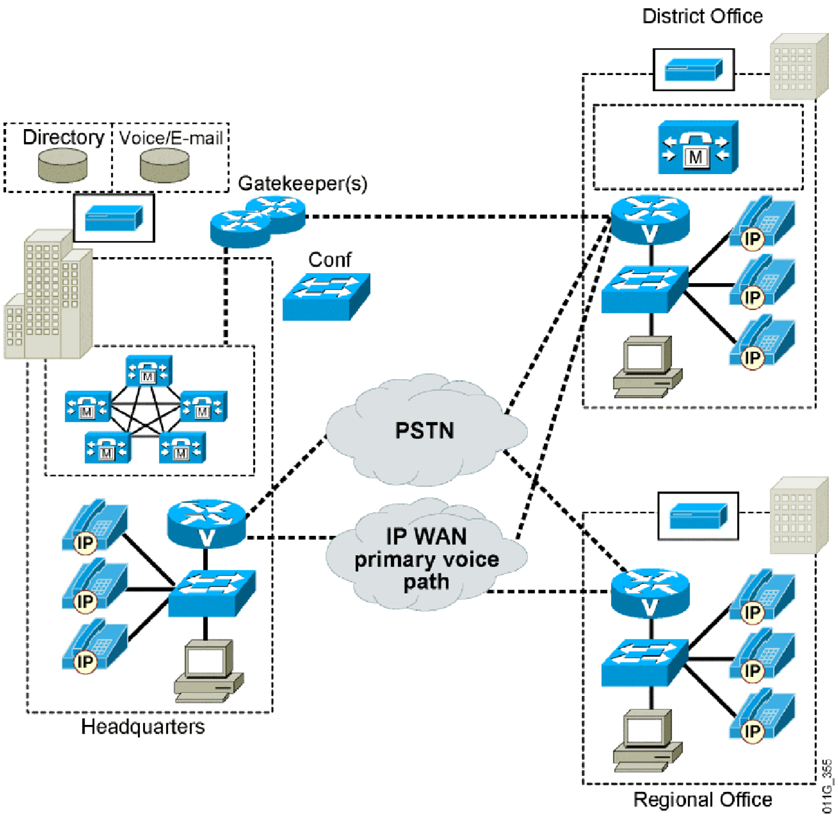

Module 12: Designing Content Networking Solutions

Case Study 12-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design a Content Networking Solution for the Company Network

Following are suggested solutions.

Step 1 Determine each location where users will need to access cached content. On an overhead transparency, create a network diagram for the headquarters and North American plants indicating the location of each content networking device including content switches, content routers, content managers, and content distribution managers.

The figure shows a content networking design for the North American plants.

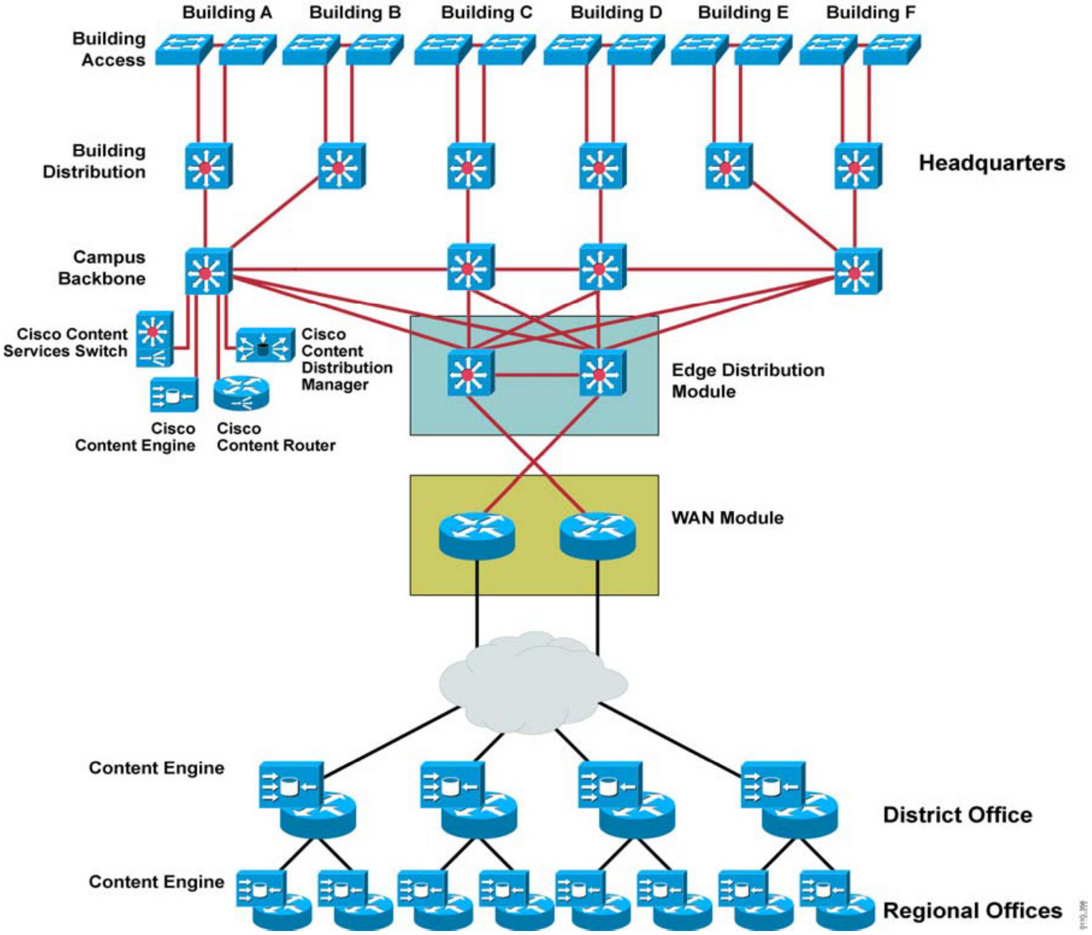

To support content networking between the headquarters and North American plant sites, the following options are included:

- Content Engines at each plant bring content closer to the users. The Content Engines support specific Java-based applications. The Content Engine at each remote site allows each plant to download the applet once for all users to access.
- Content Engines are deployed on the district and regional plant networks. The content is then delivered to the local area network, which accelerates content delivery and saves WAN bandwidth between headquarters and the district plants and between the district plants and the regional plants.
- The company has an IP-based television service for employees worldwide. Components of the content delivery network included Cisco Content Distribution Manager, Cisco Content Engine, Cisco Content Services Switch, and Cisco Content Router.

Step 3 Complete the table to design the details about the content network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your content networking solution?	Content switches and routers provide high availability for content	Redundant content engines, switches, and routers can further ensure availability for the most critical content, if desired
What security strategy will you deploy to support your content networking solution?	Applications are secured with authentication and authorization based on sensitivity of the data	The content being served can be highly sensitive so additional care is taken secure the content networking devices themselves.
What QoS strategy will you deploy to support your content networking solution?	Audio and video data are given priority over other data types	When content networking involves the delivery of delay-sensitive data such as audio or video, QoS features allow such traffic to be given priority over other data traffic.

Module 13: Designing Storage Networking Solutions

Case Study 13-2: OCSIC Bottling Company

Based on the scenario, the following tables and diagrams include the proposed solutions.

Note Your solutions may vary.

Each step that requires a solution is listed below.

Task 1: Design a Storage Networking Solution

Following are suggested solutions.

Step 1 Determine the location of each storage device on the network. On an overhead transparency, create a campus network diagram for the headquarters location and the North American plants, indicating the location of each storage networking solution.

The figure shows a storage networking design for the North American plants.

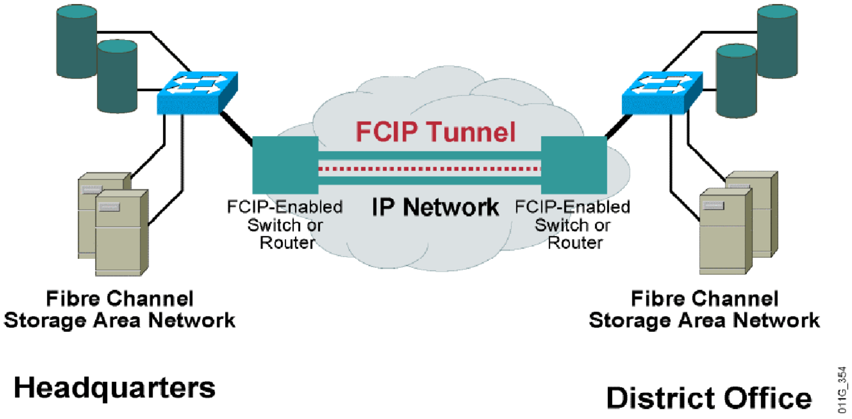
FCIP Tunnel
FCIP-Enabled Switch or Router
IP Network
FCIP-Enabled Switch or Router
Fibre Channel Storage Area Network
Fibre Channel Storage Area Network
Headquarters
District Office
01IG_354

To support IP access to storage, the company selected the following options:

- The company selected a Cisco storage router that delivers redundant iSCSI paths to a pair of Fibre Channel switches. iSCSI takes advantage of the connection-oriented TCP protocol for reliable service. Ethernet was already part of the IT network. The end result used the scalability and cost advantages of iSCSI via Gigabit Ethernet for storage networking and retained Fibre Channel for storage access.
- The applications interface with the generic SCSI layer in the Windows hosts, which see only SCSI. The storage router shielded the host from any Fibre Channel considerations.
- To support storage over the WAN, the company implemented an FCIP solution that allowed them to add storage at each plant, then use FCIP to perform asynchronous backup of data from one site to another. This allowed them to take advantage of their existing IP infrastructure and to provide additional locations for all critical data. The company put in a T3 between the regional offices and headquarters to address the bandwidth guarantees needed.

Step 2 Complete the table to design the details about the storage network.

The table summarizes the design decisions that the enterprise made to meet their requirements.

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your storage networking solution?	Use 802.1w for spanning tree Use Open Shortest Path First (OSPF) Use HSRP	802.1w provides fast spanning-tree reconvergence in Ethernet environments. OSPF provides very fast convergence following a network link or router failure. HSRP creates redundancy default gateways for iSCSI initiator hosts to ensure fast recovery if a gateway fails.
What security strategy will you deploy to support your storage networking solution?	Use IP Security (IPSec) hardware encryption Use IP and VLAN access control lists	IPSec hardware encryption encrypts FCIP tunnels across the WAN. IP and VLAN access control lists isolate storage within a LAN and storage-router-based access control lists to restrict access to storage.
What QoS strategy will you deploy to support your storage networking solution?	Prioritize iSCSI traffic Use QoS to throttle non-IP storage traffic	Prioritize iSCSI traffic within the LAN for higher-priority queuing and switching. Use QoS to throttle non-IP storage traffic to protect FCIP traffic in a WAN/MAN.

C

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.

Module 2: Designing Enterprise Campus Networks

Application Characteristics

Name of Application	Building or Location	Type of Application	Number of Users	Number of Servers	Bandwidth/ Delay Tolerance/ Loss Characteristics

Campus and Server Farm Design

Design Question	Decision	Justification
What is the logical network design?		
What physical network media will be used?		
What data link layer protocol will be used?		
What spanning-tree deployment will be used?		
What is the Layer 2/Layer 3 strategy for the network?		
Which Cisco products will be used?		
What IP addressing scheme will be used?		
Which routing protocols will be used?		

Module 3: Designing Enterprise Edge Connectivity

Application Characteristics

Name of Application	To/From Location	Type of Application	Number of Users	Number of Servers	Bandwidth/Delay/ Loss Characteristics

WAN Module Design

Design Question	Decision	Justification
What topology will be used for the WAN?		
What service provider will be selected?		
What data link layer protocol will be used		
What physical network media will be used?		
Which Cisco products will be used?		
Which routing protocols will be used?		
What IP addressing scheme will be used?		

Remote Access Design

Design Questions	Decision	Justification
What topology will be used for the WAN?		
What service provider will be selected?		
What data link layer protocol will be used		
What physical network media will be used?		
Which Cisco products will be used?		
Which routing protocols will be used?		
What IP addressing scheme will be used?		

Module 4: Designing Network Management Services

Network Management Design

Design Question	Decision	Justification
How many management domains does the enterprise require?		
How many devices need to be managed?		
What are the key components and functions required?		
How many servers are required?		
What administrative grouping of network devices will work for this enterprise?		
What administrative grouping of network management users will work for this enterprise?		

Module 5: Designing High Availability Services

High Availability Design

Design Question	Decision	Justification
Which devices should be fault tolerant?		
Which devices should be redundant?		
Which links should be redundant?		
What spanning-tree implementation and root devices are required?		
What is the router availability strategy?		

Module 6: Designing Security Services

Security Design

Design Questions	Decision	Justification
What is your firewall strategy? What features would you implement?		
What is your intrusion detection strategy? What features would you implement?		
What software features would you implement?		

Module 7: Designing QoS

QoS Design

Design Questions	Decision	Justification
What classification and marking tools and settings are required?		
What congestion avoidance tools and settings are required?		
What congestion management tools and settings are required?		
What traffic conditioning tools and settings are required?		
What signaling tools and settings are required?		
What link efficiency tools and settings are required?		

Module 8: Designing IP Multicast Services

IP Multicast Design

Design Questions	Decision	Justification
Where will you implement multicast applications on the network?		
What IP multicast control mechanism will you use?		
Will you use PIM-DM or PIM-SM? If you selected PIM-SM, determine the location of rendezvous points.		
What security is needed to support IP multicasting on the network?		
Which network devices require upgrade (memory, IOS version, new hardware) to support the IP multicast applications?		

Module 9: Designing Virtual Private Networks

Site-to-Site VPN Design

Design Questions	Decision	Justification
What topology will be used for the site-to-site VPN?		
What type of tunneling will be deployed?		
What type of security will be deployed?		
Is NAT required?		
What VPN hardware will be used?		
What type of high availability will be deployed?		

Remote-Access VPN Design

Design Questions	Decision	Justification
What topology will be used for the remote-access VPN?		
What type of tunneling will be deployed?		
What type of security will be deployed?		
Is NAT required?		
What VPN hardware will be used?		
What type of high availability will be deployed?		

Module 10: Designing Enterprise Wireless Networks

Wireless LAN Design

Design Questions	Decision	Justification
How many access points are required for the location? Where should the access points be placed?		
How many active devices can each access point support?		
How are channels identified for the design?		
How will you meet the inline power requirements for the design?		
What is the high-availability (redundancy) strategy for the design?		

Module 11: Designing IP Telephony Solutions

IP Telephony Design

Design Questions	Decision	Justification
Is centralized or distributed call processing most appropriate?		
How many CallManager servers will you deploy? In what locations? Which server will be the publisher? Which servers will be the subscribers? What other servers would you include in the cluster?		
What gateways will you deploy? Where will the gateways be located? What function will each gateway serve?		
What QoS strategy will you deploy to support the solution?		
What DSP resources are required for the solution?		
What transcoding resources are required for the solution?		
What are the network bandwidth and traffic engineering considerations?		

Module 12: Designing Content Networking Solutions

Content Networking Design

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your content networking solution?		
What security strategy will you deploy to support your content networking solution?		
What QoS strategy will you deploy to support your content networking solution?		

Module 13: Designing Storage Networking Solutions

Storage Networking Design

Design Questions	Decision	Justification
What high-availability strategy will you deploy to support your storage networking solution?		
What security strategy will you deploy to support your storage networking solution?		
What QoS strategy will you deploy to support your storage networking solution?		