CHAPTER 3

Structuring and Modularizing the Network

This chapter introduces a modular hierarchical approach to network design, the Cisco Enterprise Architecture. The chapter begins with a discussion of the hierarchical network structure. The next section introduces network modularization and discusses the details of the Cisco Enterprise Architecture. Following that are a detailed description of services within modular networks, and a discussion of network management protocols and features.

Network Hierarchy

This section explains the hierarchical network model, which is composed of the access, distribution, and core layers. The functions generally associated with each of these layers are discussed, as is the most common approach to designing a hierarchical network.

Historically used in the design of enterprise local-area network and wide-area network data networks, this model works equally well within the functional modules of the Cisco Enterprise Architecture. These modules are discussed later in this chapter, in the section "Using a Modular Approach to Network Design."

Hierarchical Network Model

The hierarchical network model provides a framework that network designers can use to help ensure that the network is flexible and easy to implement and troubleshoot.

Hierarchical Network Design Layers

As shown in Figure 3-1, the hierarchical network design model consists of three layers:

- The access layer provides local and remote workgroup or user access to the network.
- The distribution layer provides policy-based connectivity.
- The core (or backbone) layer provides high-speed transport to satisfy the connectivity and transport needs of the distribution layer devices.

Figure 3-1 *Hierarchical Model's Three Layers*

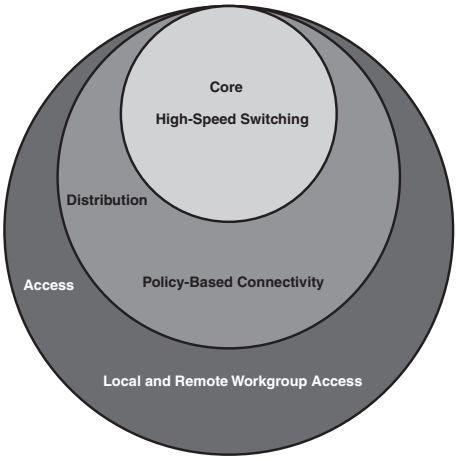

Each hierarchical layer focuses on specific functions, thereby allowing the network designer to choose the right systems and features based on their function within the model. This approach helps provide more accurate capacity planning and minimize total costs. Figure 3-2 illustrates a sample network showing the mapping to the hierarchical model's three layers.

Figure 3-2 *Sample Network Designed Using the Hierarchical Model*

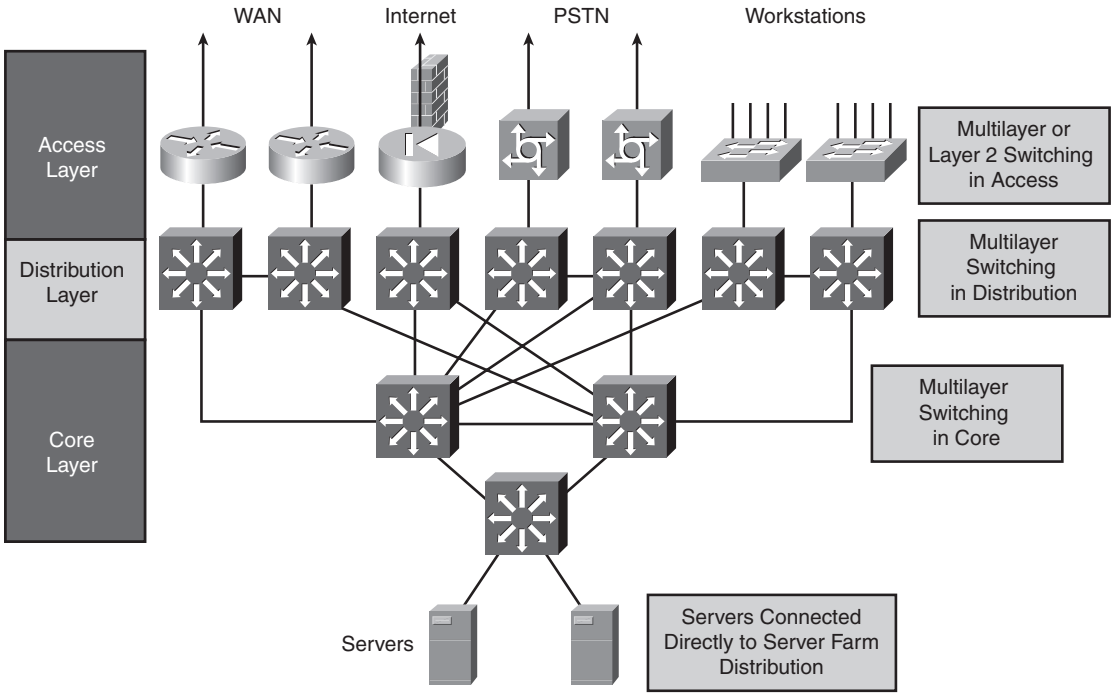

You do not have to implement the hierarchical layers as distinct physical entities; they are defined to aid successful network design and to represent functionality that must exist within a network. The actual manner in which you implement the layers depends on the needs of the network you are designing. Each layer can be implemented in routers or switches, represented by physical media, or combined in a single device. A particular layer can be omitted, but hierarchy should be maintained for optimum performance. The following sections detail the functionality of the three layers and the devices used to implement them.

Access Layer Functionality

This section describes the access layer functions and the interaction of the access layer with the distribution layer and local or remote users.

The Role of the Access Layer

The access layer is the concentration point at which clients access the network. Access layer devices control traffic by localizing service requests to the access media.

The purpose of the access layer is to grant user access to network resources. Following are the access layer's characteristics:

- In the campus environment, the access layer typically incorporates switched LAN devices with ports that provide connectivity for workstations and servers.
- In the WAN environment, the access layer for teleworkers or remote sites provides access to the corporate network across some wide-area technology, such as Frame Relay, Multiprotocol Label Switching (MPLS), Integrated Services Digital Network, leased lines, Digital Subscriber Line (DSL) over traditional telephone copper lines, or coaxial cable.
- So as not to compromise network integrity, access is granted only to authenticated users or devices (such as those with physical address or logical name authentication). For example, the devices at the access layer must detect whether a telecommuter who is dialing in is legitimate, yet they must require minimal authentication steps for the telecommuter.

Layer 2 and Multilayer Switching in the Access Layer

Access can be provided to end users as part of either a Layer 2 (L2) switching environment or a multilayer switching environment.

> **NOTE** In this book, the term *multilayer switching* denotes a switch's generic capability to use information at different protocol layers as part of the switching process; the term *Layer 3 switching* is a synonym for multilayer switching in this context.
>
> Cisco switches implement the use of protocol information from multiple layers in the switching process in two different ways. The first way is *multilayer switching (MLS)* and the second way is *Cisco Express Forwarding (CEF).* MLS and CEF are described further in Chapter 4, "Designing Basic Campus and Data Center Networks."

Using Layer 2 Switching in the Access Layer

Access to local workstations and servers can be provided using shared or switched media LANs; VLANs may be used to segment the switched LANs. Each LAN or VLAN is a single broadcast domain.

The access layer aggregates end-user switched 10/100 ports and provides Fast Ethernet, Fast EtherChannel, and Gigabit Ethernet uplinks to the distribution layer to satisfy connectivity requirements and reduce the size of the broadcast domains. You can deploy multiple VLANs, each with its own IP subnet and its own instance of Spanning Tree Protocol (STP) providing alternative paths in case of failure. In this case, Layer 2 trunking (typically using the Institute for Electrical and Electronic Engineers [IEEE] 802.1Q trunking protocol) is used between the access layer switches and the distribution layer switches, with per-VLAN STP on each uplink for load balancing and redundancy, and with a distribution layer multilayer switch providing the inter-VLAN communication for the access layer.

> **NOTE** Chapter 4 discusses STP further.

KEY POINT A recommended best practice is to implement one VLAN—thus supporting one IP subnet—per access switch and to connect the access switches to the distribution switches with Layer 3 links rather than with trunks.

> **NOTE** In small networks, the access layer is often collapsed into the distribution layer; in other words, one device might handle all functions of the access and distribution layers.

KEY POINT Using the Rapid Spanning Tree Protocol (RSTP) is a recommended best practice in the enterprise. RSTP is an evolution of the IEEE 802.1d STP standard and provides faster spanning-tree convergence after a topology change.

When RSTP cannot be implemented, Cisco IOS STP features such as UplinkFast, PortFast, and BackboneFast can be used to provide equivalent convergence improvements. These features are described as follows:

- **UplinkFast**: Enables faster failover on an access layer switch on which dual uplinks connect to the distribution layer. The failover time is reduced by unblocking the blocked uplink port on a switch immediately after root port failure, thereby transitioning it to the forwarding state immediately, without transitioning the port through the listening and learning states.
- **BackboneFast**: If a link fails on the way to the root switch but is not directly connected to the local switch, BackboneFast reduces the convergence time from 50 seconds to between 20 and 30 seconds.
- **PortFast**: Enables switch ports connected to nonswitch devices (such as workstations) to immediately enter the spanning-tree forwarding state, thereby bypassing the listening and learning states, when they come up. Ports connected only to an end-user device do not have bridging loops, so it is safe to go directly to the forwarding state, significantly reducing the time it takes before the port is usable.

NOTE Chapter 4 discusses other STP features.

Using Multilayer Switching in the Access Layer

The most common design for remote users is to use multilayer switches or routers. A multilayer switch, or router, is the boundary for broadcast domains and is necessary for communicating between broadcast domains (including VLANs). Access routers provide access to remote office environments using various wide-area technologies combined with multilayer features, such as route propagation, packet filtering, authentication, security, Quality of Service (QoS), and so on. These technologies allow the network to be optimized to satisfy a particular user's needs. In a dialup connection environment, dial-on-demand routing (DDR) and static routing can be used to control costs.

Access Layer Example

Figure 3-3 illustrates a sample network in which the campus access layer aggregates end users and provides uplinks to the distribution layer. The access layer switches are dual-attached to the distribution layer switches for high availability.

Figure 3-3 *Access Layer Connectivity in a Campus LAN*

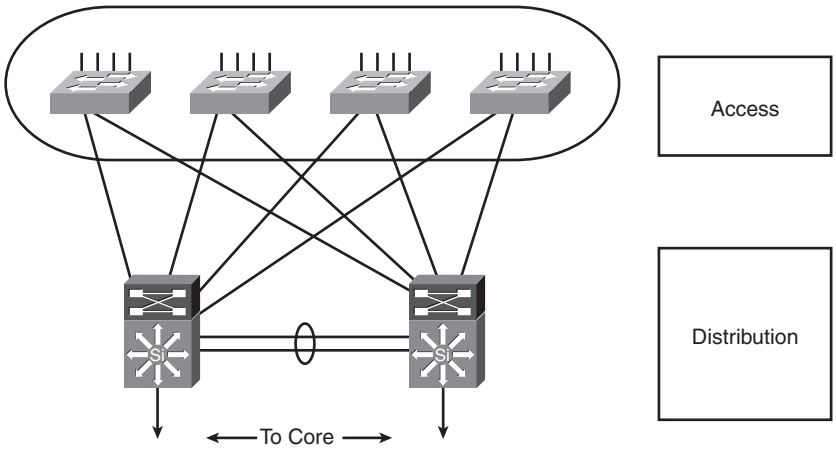

The access layer can support convergence, high availability, security, QoS, and IP multicast. Some services found at the access layer include establishing a QoS trust boundary, broadcast suppression, and Internet Group Management Protocol (IGMP) snooping.

Distribution Layer Functionality

This section describes distribution layer functions and the interaction of the distribution layer with the core and access layers.

The Role of the Distribution Layer

The *distribution layer* represents both a separation between the access and core layers and a connection point between the diverse access sites and the core layer. The distribution layer determines department or workgroup access and provides policy-based connectivity.

Following are the characteristics of the distribution layer:

- Distribution layer devices control access to resources that are available at the core layer and must therefore use bandwidth efficiently.
- In a campus environment, the distribution layer aggregates wiring closet bandwidth by concentrating multiple low-speed access links into a high-speed core link and using switches to segment workgroups and isolate network problems to prevent them from affecting the core layer.
- Similarly, in a WAN environment, the distribution layer aggregates WAN connections at the edge of the campus and provides policy-based connectivity.

- This layer provides redundant connections for access devices. Redundant connections also provide the opportunity to load-balance between devices.
- The distribution layer represents a routing boundary between the access and core layers and is where routing and packet manipulation are performed.
- The distribution layer allows the core layer to connect diverse sites while maintaining high performance. To maintain good performance in the core, the distribution layer can redistribute between bandwidth-intensive access-layer routing protocols and optimized core routing protocols. Route filtering is also implemented at the distribution layer.
- The distribution layer can summarize routes from the access layer to improve routing protocol performance. For some networks, the distribution layer offers a default route to access-layer routers and runs dynamic routing protocols only when communicating with core routers.
- The distribution layer connects network services to the access layer and implements policies for QoS, security, traffic loading, and routing. For example, the distribution layer addresses different protocols' QoS needs by implementing policy-based traffic control to isolate backbone and local environments. Policy-based traffic control prioritizes traffic to ensure the best performance for the most time-critical and time-dependent applications.
- The distribution layer is often the layer that terminates access layer VLANs (broadcast domains); however, this can also be done at the access layer.
- This layer provides any media transitions (for example, between Ethernet and ATM) that must occur.

Policy-Based Connectivity

Policy-based connectivity means implementing the policies of the organization (as described in Chapter 2, "Applying a Methodology to Network Design"). Methods for implementing policies include the following:

- Filtering by source or destination address
- Filtering based on input or output ports
- Hiding internal network numbers by route filtering
- Providing specific static routes rather than using routes from a dynamic routing protocol
- Security (for example, certain packets might not be allowed into a specific part of the network)
- QoS mechanisms (for example, the precedence and type of service [ToS] values in IP packet headers can be set in routers to leverage queuing mechanisms to prioritize traffic)

Distribution Layer Example

Figure 3-4 shows a sample network with various features of the distribution layer highlighted.

Figure 3-4 *Example of Distribution Layer Features*

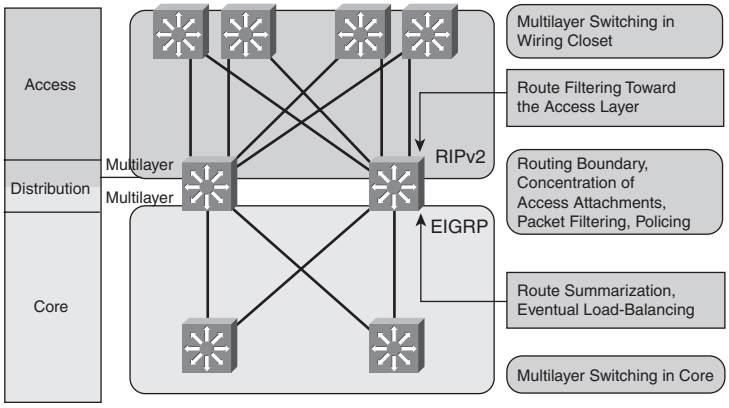

Following are the characteristics of the distribution layer in the routed campus network shown in Figure 3-4:

- Multilayer switching is used toward the access layer (and, in this case, within the access layer).
- Multilayer switching is performed in the distribution layer and extended toward the core layer.
- The distribution layer performs two-way route redistribution to exchange the routes between the Routing Information Protocol version 2 (RIPv2) and Enhanced Interior Gateway Routing Protocol (EIGRP) routing processes.
- Route filtering is configured on the interfaces toward the access layer.
- Route summarization is configured on the interfaces toward the core layer.
- The distribution layer contains highly redundant connectivity, both toward the access layer and toward the core layer.

Core Layer Functionality

This section describes core layer functions and the interaction of the core layer with the distribution layer.

The Role of the Core Layer

The function of the core layer is to provide fast and efficient data transport. Characteristics of the core layer include the following:

- The core layer is a high-speed backbone that should be designed to switch packets as quickly as possible to optimize communication transport within the network.
- Because the core is critical for connectivity, core layer devices are expected to provide a high level of availability and reliability. A fault-tolerant network design ensures that failures do not have a major impact on network connectivity. The core must be able to accommodate failures by rerouting traffic and responding quickly to changes in network topology. The core must provide a high level of redundancy. A full mesh is strongly suggested, and at least a well-connected partial mesh with multiple paths from each device is required.
- The core layer should not perform any packet manipulation, such as checking access lists or filtering, which would slow down the switching of packets.
- The core layer must be manageable.
- The core devices must be able to implement scalable protocols and technologies, and provide alternative paths and load balancing.

Switching in the Core Layer

Layer 2 switching or multilayer switching (routing) can be used in the core layer. Because core devices are responsible for accommodating failures by rerouting traffic and responding quickly to network topology changes, and because performance for routing in the core with a multilayer switch incurs no cost, most implementations have multilayer switching in the core layer. The core layer can then more readily implement scalable protocols and technologies, and provide alternate paths and load balancing.

Figure 3-5 shows an example of Layer 2 switching in the campus core.

Figure 3-5 *Layer 2 Switching in the Campus Core*

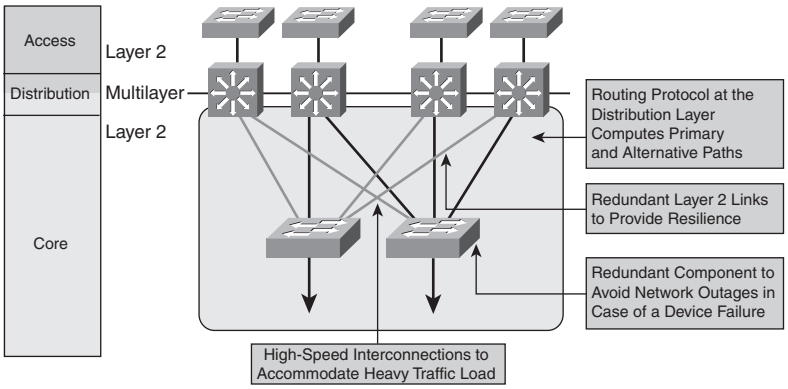

In Figure 3-5, a typical packet between access sites follows these steps:

Step 1 The packet is Layer 2–switched toward a distribution switch.

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

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

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

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

Figure 3-6 shows an example of multilayer switching in the campus core.

Figure 3-6 *Multilayer Switching in the Campus Core*

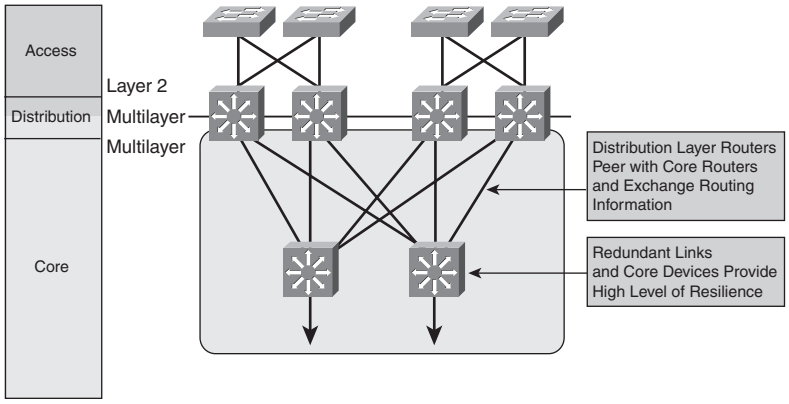

In Figure 3-6, a typical packet between access sites follows these steps:

Step 1 The packet is Layer 2–switched toward a distribution switch.

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

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

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

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

Hierarchical Routing in the WAN

Figure 3-7 shows an example of hierarchical routing in the WAN portion of a network.

Figure 3-7 *Hierarchical Routing in the WAN*

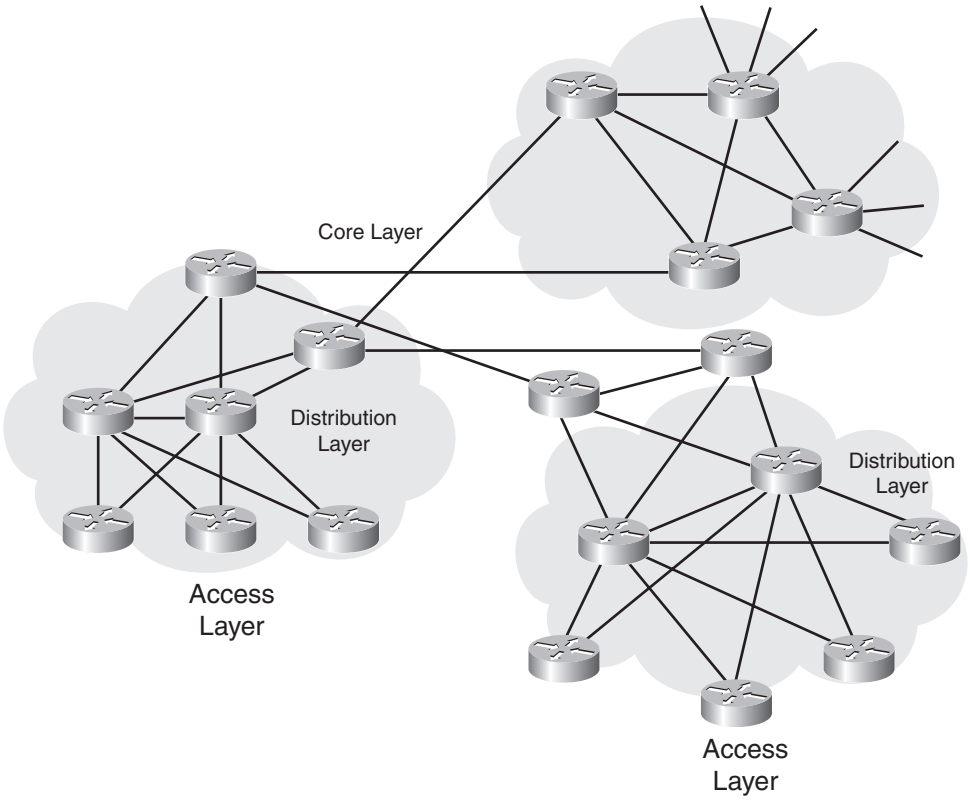

In Figure 3-7, a typical packet between access sites follows these steps:

Step 1 The packet is Layer 3–forwarded toward the distribution router.

Step 2 The distribution router forwards the packet toward a core interface.

Step 3 The packet is forwarded across the WAN core.

Step 4 The receiving distribution router forwards the packet toward the appropriate access layer router.

Step 5 The packet is Layer 3–forwarded to the destination host's access layer LAN.

Using a Modular Approach to Network Design

This section expands on the Cisco Service-Oriented Network Architecture (SONA) framework described in Chapter 2 and explores the six modules of the Cisco Enterprise Architecture, with an emphasis on the network infrastructure design considerations.

NOTE The access, distribution, and core layers can appear within each module of the Cisco Enterprise Architecture.

The modularity built into the architecture allows flexibility in network design and facilitates implementation and troubleshooting. Before the details of the architecture itself are introduced, an overview of the evolution of enterprise networks is provided.

Evolution of Enterprise Networks

You do not have to go far back in history to find a time when networks were primarily used for file and print services. These networks were isolated LANs that were built throughout the enterprise organization. As organizations interconnected, these isolated LANs and their functions grew from file and print services to include critical applications; the critical nature and complexity of the enterprise networks also grew.

As discussed in the previous section, Cisco introduced the hierarchical model to divide the enterprise network design (separately for both campus and WAN networks) into the access, distribution, and core layers. This solution has several weaknesses, especially for large networks, which are difficult to implement, manage, and, particularly, troubleshoot. Networks became complex, and it was difficult to evaluate a network solution end-to-end through the network. The hierarchical model does not scale well to these large networks.

An efficient method of solving and scaling a complex task is to break it into smaller, more specialized tasks. Networks can easily be broken down smaller because they have natural physical, logical, and functional boundaries. If they are sufficiently large to require additional design or operational separation, these specialized functional modules can then be designed hierarchically with the access, distribution, and core layers.

The Cisco Enterprise Architecture does just that: It reduces the enterprise network into further physical, logical, and functional boundaries, to scale the hierarchical model. Now, rather than designing networks using only the hierarchical model, networks can be designed using this Cisco Enterprise Architecture, with hierarchy (access, distribution, and core) included in the various modules, as required.

Designing with this Cisco Enterprise Architecture is not much different from what is already used in practice; it formalizes current practice. There have always been separate hierarchies for the

campus (with access, distribution, and core) and for the WAN (the remote office was the access layer, the regional office provided the distribution layer, and the headquarters was the core). The hierarchies tied together at the campus backbone. The Cisco Enterprise Architecture extends the concept of hierarchy from the original two modules: Campus and WAN.

Cisco SONA Framework

As illustrated in Figure 3-8, the Cisco SONA provides an enterprise-wide framework that integrates the entire network—campus, data center, enterprise edge, WAN, branches, and teleworkers—offering staff secure access to the tools, processes, and services they require.

Figure 3-8 *Cisco SONA Framework*

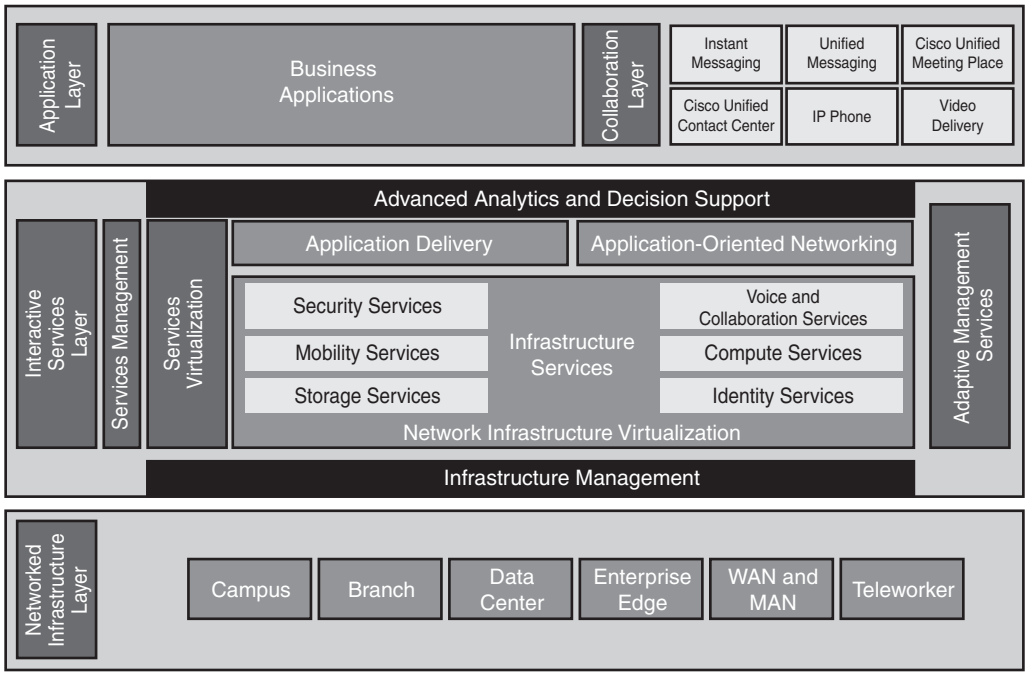

The modules of the Cisco Enterprise Architecture represent focused views of each of the places in the network described in the SONA framework. Each module has a distinct network infrastructure and distinct services; network applications extend between the modules.

Functional Areas of the Cisco Enterprise Architecture

At the first layer of modularity in the Cisco Enterprise Architecture, the entire network is divided into functional *components*—functional areas that contain network modules—while still

maintaining the hierarchical concept of the core-distribution-access layers within the network modules as needed.

NOTE The access, distribution, and core layers can appear in any functional area or module of the Cisco Enterprise Architecture.

The Cisco Enterprise Architecture comprises the following six major functional areas (also called *modules*):

- Enterprise Campus
- Enterprise Edge
- Service Provider
- Enterprise Branch
- Enterprise Data Center
- Enterprise Teleworker

KEY POINT An enterprise does not implement the modules in the Service Provider functional area; they are necessary for enabling communication with other networks.

NOTE The Cisco SONA Enterprise Edge and the WAN and metropolitan-area network (MAN) modules are represented as one functional area in the Cisco Enterprise Architecture, the Enterprise Edge.

Figure 3-9 illustrates the modules within the Cisco Enterprise Architecture.

Figure 3-9 *Cisco Enterprise Architecture*

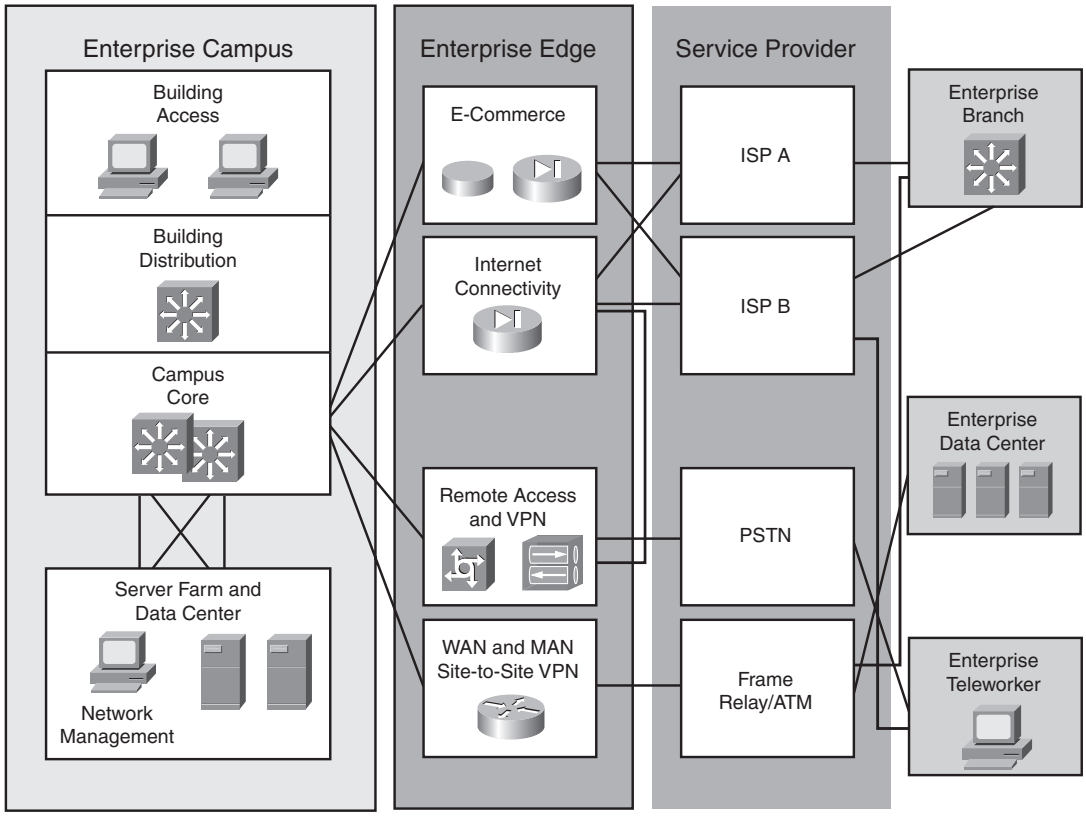

NOTE Figure 3-9 is reproduced on the inside back cover of this book for your reference.

The Cisco Enterprise Campus Architecture combines a core infrastructure of intelligent switching and routing with tightly integrated productivity-enhancing technologies, including Cisco Unified Communications, mobility, and advanced security. The architecture provides the enterprise with high availability through a resilient multilayer design, redundant hardware and software features, and automatic procedures for reconfiguring network paths when failures occur. IP multicast capabilities provide optimized bandwidth consumption, and QoS features ensure that real-time traffic (such as voice, video, or critical data) is not dropped or delayed. Integrated security protects against and mitigates the impact of worms, viruses, and other attacks on the network, including at the switch port level. For example, the Cisco enterprise-wide architecture extends support for security standards, such as the IEEE 802.1X port-based network access control standard and the Extensible Authentication Protocol. It also provides the flexibility to add Internet Protocol Security (IPsec) and MPLS virtual private networks (VPN), identity and access management, and

VLANs to compartmentalize access. These features help improve performance and security while decreasing costs.

The Cisco Enterprise Edge Architecture offers connectivity to voice, video, and data services outside the enterprise. This module enables the enterprise to use Internet and partner resources, and provide resources for its customers. QoS, service levels, and security are the main issues in the Enterprise Edge.

The Cisco Enterprise WAN and MAN and Site-to-Site VPN module is part of the Enterprise Edge. It offers the convergence of voice, video, and data services over a single Cisco Unified Communications network, which enables the enterprise to span large geographic areas in a cost-effective manner. QoS, granular service levels, and comprehensive encryption options help ensure the secure delivery of high-quality corporate voice, video, and data resources to all corporate sites, enabling staff to work productively and efficiently wherever they are located. Security is provided with multiservice VPNs (both IPsec and MPLS) over Layer 2 or Layer 3 WANs, hub-and-spoke, or full-mesh topologies.

The Cisco Enterprise Data Center Architecture is a cohesive, adaptive network architecture that supports requirements for consolidation, business continuance, and security while enabling emerging service-oriented architectures, virtualization, and on-demand computing. Staff, suppliers, and customers can be provided with secure access to applications and resources, simplifying and streamlining management and significantly reducing overhead. Redundant data centers provide backup using synchronous and asynchronous data and application replication. The network and devices offer server and application load balancing to maximize performance. This architecture allows the enterprise to scale without major changes to the infrastructure. This module can be located either at the campus as a server farm or at a remote facility.

The Cisco Enterprise Branch Architecture allows enterprises to extend head-office applications and services (such as security, Cisco Unified Communications, and advanced application performance) to thousands of remote locations and users or to a small group of branches. Cisco integrates security, switching, network analysis, caching, and converged voice and video services into a series of integrated services routers (ISR) in the branch so that the enterprises can deploy new services without buying new routers. This architecture provides secure access to voice, mission-critical data, and video applications—anywhere, anytime. Advanced routing, VPNs, redundant WAN links, application content caching, and local IP telephony call processing features are available with high levels of resilience for all the branch offices. An optimized network leverages the WAN and LAN to reduce traffic and save bandwidth and operational expenses. The enterprise can easily support branch offices with the capability to centrally configure, monitor, and manage devices located at remote sites, including tools, such as Cisco AutoQoS and the Cisco Router and Security Device Manager graphical user interface QoS wizard, which proactively resolve congestion and bandwidth issues before they affect network performance.

The Cisco Enterprise Teleworker Architecture allows enterprises to securely deliver voice and data services to remote small or home offices (known as small office, home office [SOHO]) over a standard broadband access service, providing a business-resiliency solution for the enterprise and a flexible work environment for employees. Centralized management minimizes the IT support costs, and robust integrated security mitigates the unique security challenges of this environment. Integrated security and identity-based networking services enable the enterprise to extend campus security policies to the teleworker. Staff can securely log in to the network over an always-on VPN and gain access to authorized applications and services from a single cost-effective platform. Productivity can be further enhanced by adding an IP phone, thereby providing cost-effective access to a centralized IP communications system with voice and unified messaging services.

NOTE Each of these modules has specific requirements and performs specific roles in the network; note that their sizes in Figure 3-9 are not meant to reflect their scale in a real network.

This architecture allows network designers to focus on only a selected module and its functions. Designers can describe each network application and service on a per-module basis and validate each as part of the complete enterprise network design. Modules can be added to achieve scalability if necessary; for example, an organization can add more Enterprise Campus modules if it has more than one campus.

Guidelines for Creating an Enterprise Network

When creating an Enterprise network, divide the network into appropriate areas, where the Enterprise Campus includes all devices and connections within the main Campus location; the Enterprise Edge covers all communications with remote locations and the Internet from the perspective of the Enterprise Campus; and the remote modules include the remote branches, teleworkers, and the remote data center. Define clear boundaries between each of the areas.

NOTE Depending on the network, an enterprise can have multiple campus locations. A location that might be a remote branch from the perspective of a central campus location might locally use the Cisco Enterprise Campus Architecture.

Figure 3-10 shows an example of dividing a network into an Enterprise Campus area, an Enterprise Edge area, and some remote areas.

Figure 3-10 *Sample Network Divided into Functional Areas*

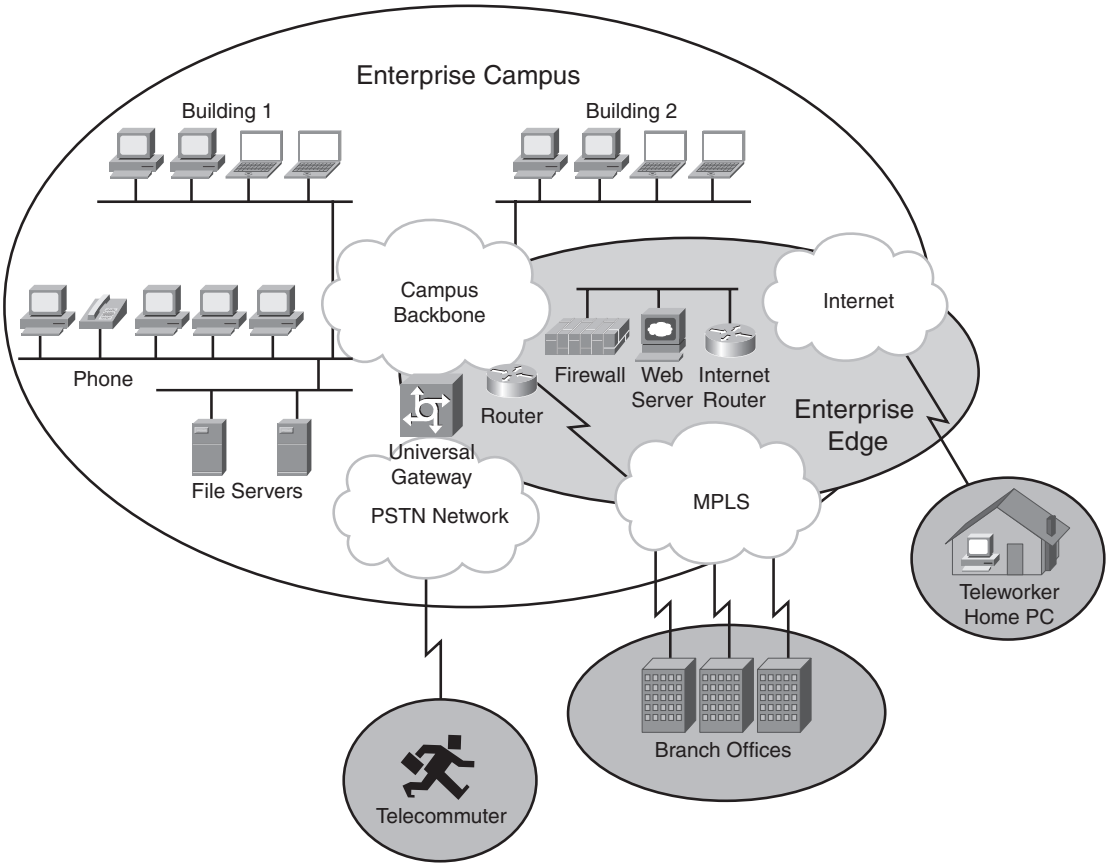

The following sections provide additional details about each of the functional areas and their modules.

Enterprise Campus Modules

This section introduces the Enterprise Campus functional area and describes the purpose of each module therein. It also discusses connections with other modules.

An *enterprise campus* site is a large site that is often the corporate headquarters or a major office. Regional offices, SOHOs, and mobile workers might have to connect to the central campus for data and information. As illustrated in Figure 3-11, the Enterprise Campus functional area includes the Campus Infrastructure module and, typically, a Server Farm module.

Figure 3-11 *Enterprise Campus Functional Area*

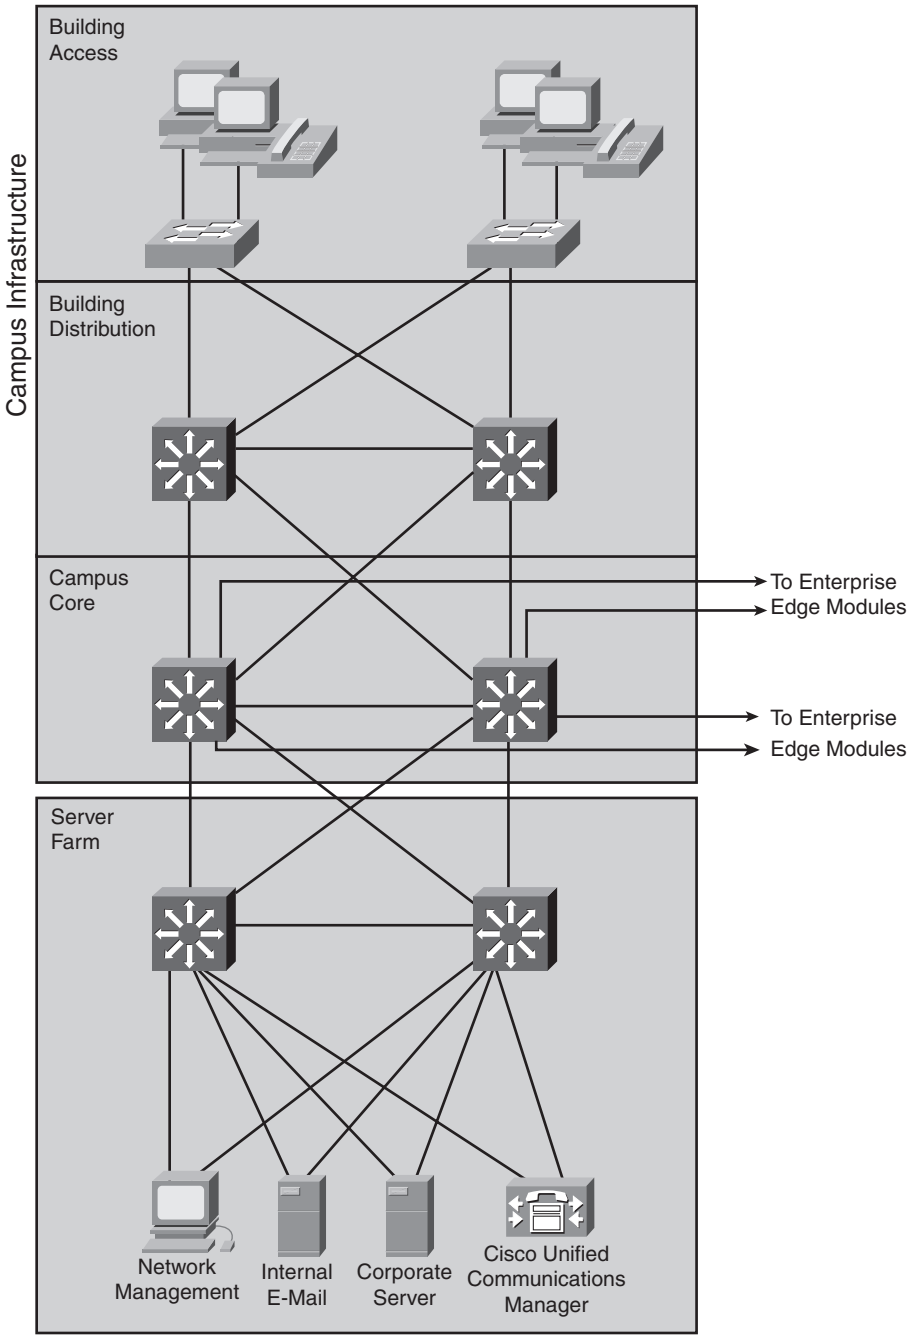

Campus Infrastructure Module

The Campus Infrastructure design consists of several buildings connected across a Campus Core. The Campus Infrastructure module connects devices within a campus to the Server Farm and Enterprise Edge modules. A single building in a Campus Infrastructure design contains a Building Access layer and a Building Distribution layer. When more buildings are added to the Campus Infrastructure, a backbone or Campus Core layer is added between buildings. The Campus Infrastructure module includes three layers:

- The Building Access layer
- The Building Distribution layer
- The Campus Core layer

NOTE In the most general model, the Building Access layer uses Layer 2 switching, and the Building Distribution layer uses multilayer switching.

Building Access Layer

The Building Access layer, located within a campus building, aggregates end users from different workgroups and provides uplinks to the Building Distribution layer. It contains end-user devices such as workstations, Cisco IP phones, and networked printers, connected to Layer 2 access switches; VLANs and STP might also be supported. The Building Access layer provides important services, such as broadcast suppression, protocol filtering, network access, IP multicast, and QoS. For high availability, the access switches are dual-attached to the distribution layer switches. The Building Access layer might also provide Power over Ethernet (PoE) and auxiliary VLANs to support voice services.

Building Distribution Layer

The Building Distribution layer aggregates the wiring closets within a building and provides connectivity to the Campus Core layer. It provides aggregation of the access layer networks using multilayer switching. The Building Distribution layer performs routing, QoS, and access control. Requests for data flow into the multilayer switches and onward into the Campus Core layer; responses follow the reverse path. Redundancy and load balancing with the Building Access and Campus Core layer are recommended. For example, in Figure 3-11, the Building Distribution layer has two equal-cost paths into the Campus Core layer, providing fast failure recovery because each distribution switch maintains two equal-cost paths in its routing table to every destination network. If one connection to the Campus Core layer fails, all routes immediately switch over to the remaining path.

Campus Core Layer

The Campus Core layer is the core layer of the Campus Infrastructure module. Within the Enterprise Campus functional area, this high-performance, switched backbone connects the buildings and various parts of the campus. Specifically, this layer interconnects the Building Distribution layer with the Server Farm and the Enterprise Edge modules.

The Campus Core layer of the Campus Infrastructure module provides redundant and fast-converging connectivity between buildings and with the Server Farm and Enterprise Edge modules. It routes and switches traffic as quickly as possible from one module to another. This module usually uses multilayer switches for high-throughput functions with added routing, QoS, and security features.

Server Farm Module

A high-capacity, centralized server farm module provides users with internal server resources. In addition, it typically supports network management services for the enterprise, including monitoring, logging, and troubleshooting, and other common management features from end to end.

The Server Farm module typically contains internal e-mail and other corporate servers that provide internal users with application, file, print, e-mail, and Domain Name System (DNS) services. As shown in Figure 3-11, because access to these servers is vital, as a best practice, they are typically connected to two different switches to enable full redundancy or load sharing. Moreover, the Server Farm module switches are cross-connected with the Campus Core layer switches, thereby enabling high reliability and availability of all servers in the Server Farm module.

The network management system performs system logging, network monitoring, and general configuration management functions. For management purposes, an out-of-band network connection (a network on which no production traffic travels) to all network components is recommended. For locations where an out-of-band network is impossible (because of geographic or system-related issues), the network management system uses the production network.

Network management can provide configuration management for nearly all devices in the network, using a combination of the following two technologies:

- Cisco IOS routers can act as terminal servers to provide a dedicated management network segment to the console ports on the Cisco devices throughout the enterprise by using a reverse-Telnet function.

- More extensive management features (software changes, content updates, log and alarm aggregation, and Simple Network Management Protocol [SNMP] management) can be provided through the dedicated out-of-band management network segment.

NOTE These Server Farm attributes also apply to a remote Data Center module.

Enterprise Campus Guidelines

Follow these guidelines for creating the modules within an Enterprise Campus functional area:

Step 1 Select modules within the campus that act as buildings with access and distribution layers.

Step 2 Determine the locations and the number of access switches and their uplinks to distribution layer switches.

Step 3 Select the appropriate distribution layer switches, taking into account the number of access layer switches and end users. Use at least two distribution layer switches for redundancy.

Step 4 Consider two uplink connections from each access layer switch to the two distribution layer switches.

Step 5 Determine where servers are or will be located, and design the Server Farm module with at least two distribution layer switches that connect all servers for full redundancy. Include out-of-band network management connections to all critical devices in the campus network.

Step 6 Design the Campus Infrastructure module's Campus Core layer using at least two switches and provide for the expected traffic volume between modules.

Step 7 Interconnect all modules of the Enterprise Campus with the Campus Infrastructure module's Campus Core layer in a redundant manner.

Enterprise Edge Modules

This section describes the components of the Enterprise Edge and explains the importance of each module. The Enterprise Edge infrastructure modules aggregate the connectivity from the various elements outside the campus—using various services and WAN technologies as needed, typically provisioned from service providers—and route the traffic into the Campus Core layer. The Enterprise Edge modules perform security functions when enterprise resources connect across public networks and the Internet. As shown in Figure 3-12 and in the following list, the Enterprise Edge functional area is composed of four main modules:

- **E-commerce module**: The E-commerce module includes the devices and services necessary for an organization to provide e-commerce applications.

- **Internet Connectivity module**: The Internet Connectivity module provides enterprise users with Internet access.
- **Remote Access and VPN module**: This module terminates VPN traffic and dial-in connections from external users.
- **WAN and MAN and Site-to-Site VPN module**: This module provides connectivity between remote sites and the central site over various WAN technologies.

Figure 3-12 *Enterprise Edge Functional Area*

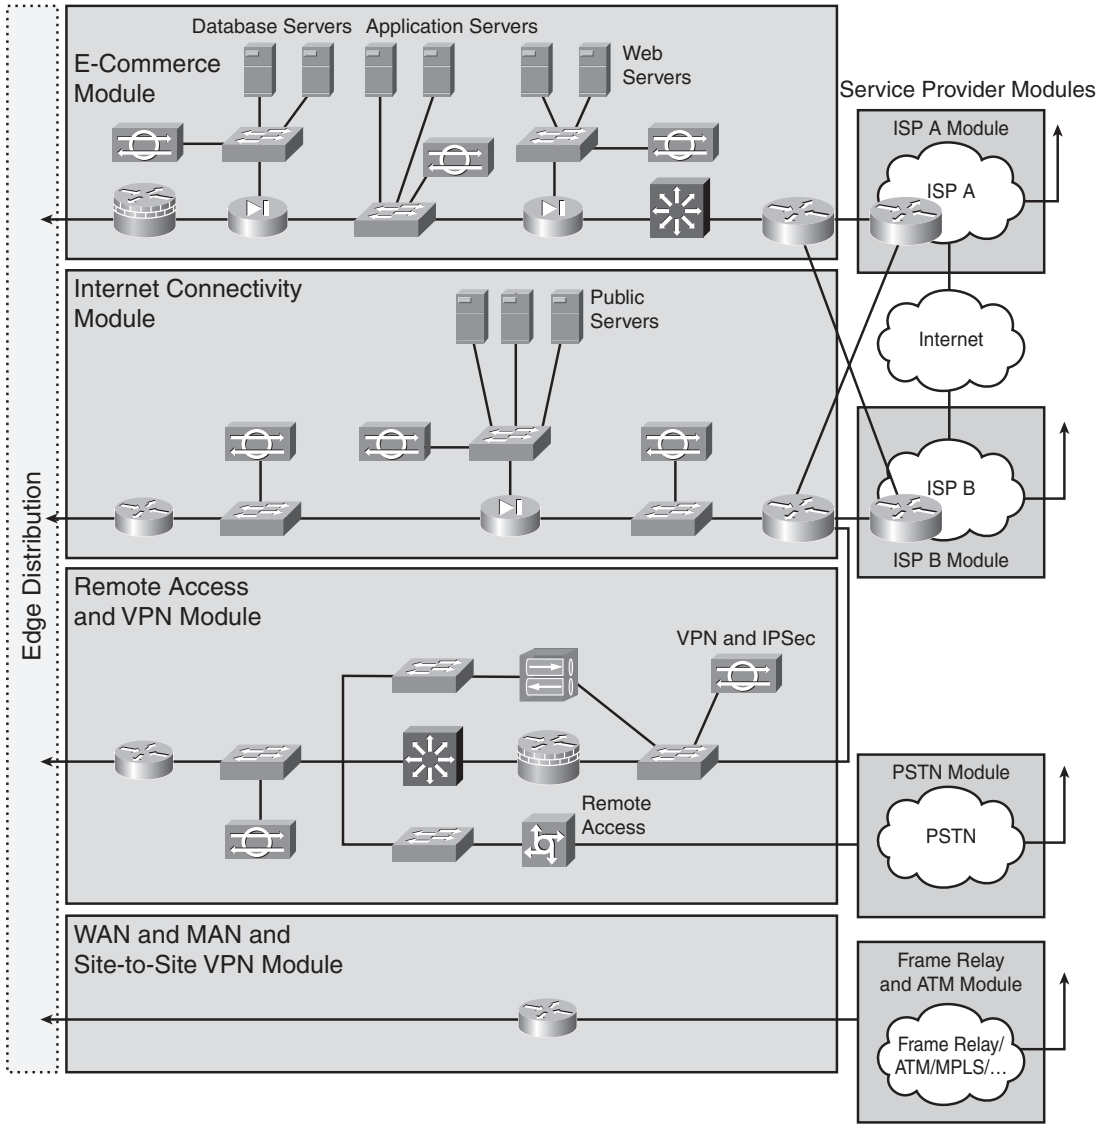

These modules connect to the Campus Core directly or through an optional Edge Distribution module. The optional Edge Distribution module aggregates the connectivity from the various elements at the enterprise edge and routes the traffic into the Campus Core layer. In addition, the Edge Distribution module acts as a boundary between the Enterprise Campus and the Enterprise Edge and is the last line of defense against external attacks; its structure is similar to that of the Building Distribution layer.

The following sections detail each of the four main Enterprise Edge modules.

E-commerce Module

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

- **Web servers**: Act as the primary user interface for e-commerce navigation
- **Application servers**: Host the various applications
- **Database servers**: Contain the application and transaction information that is the heart of the e-commerce business implementation
- **Firewalls or firewall routers**: Govern communication and provide security between the system's various users
- **Network Intrusion Detection System/Network Intrusion Protection System (NIDS/NIPS) appliances**: Monitor key network segments in the module to detect and respond to attacks against the network
- **Multilayer switch with Intrusion Detection System/Intrusion Protection System (IDS/IPS) modules**: Provide traffic transport and integrated security monitoring
- **Host-Based Intrusion Protection Systems**: Deployed on sensitive core application servers and on dedicated appliances to provide real-time reporting and prevention of attacks as an extra layer of defense

Internet Connectivity Module

The Internet Connectivity module provides internal users with connectivity to Internet services, such as HTTP, FTP, Simple Mail Transfer Protocol (SMTP), and DNS. This module also provides Internet users with access to information published on an enterprise's public servers, such as HTTP and FTP servers. Internet session initiation is typically from inside the enterprise toward

the Internet. Additionally, this module accepts VPN traffic from remote users and remote sites and forwards it to the Remote Access and VPN module, where VPN termination takes place. The Internet Connectivity module is not designed to serve e-commerce applications. Major components used in the Internet Connectivity module include the following:

- **SMTP mail servers**: Act as a relay between the Internet and the intranet mail servers.
- **DNS servers**: Serve as the authoritative external DNS server for the enterprise and relay internal DNS requests to the Internet.
- **Public servers (for example, FTP and HTTP)**: Provide public information about the organization. Each server on the public services segment contains host-based intrusion detection systems (HIDS) to monitor against any rogue activity at the operating system level and in common server applications including HTTP, FTP, and SMTP.
- **Firewalls or firewall routers**: Provide network-level protection of resources, provide stateful filtering of traffic, and forward VPN traffic from remote sites and users for termination.
- **Edge routers**: Provide basic filtering and multilayer connectivity to the Internet.

Remote Access and VPN Module

The Remote Access and VPN module terminates remote access traffic and VPN traffic that the Internet Connectivity Module forwards from remote users and remote sites. It also uses the Internet Connectivity module to initiate VPN connections to remote sites. Furthermore, the module terminates dial-in connections received through the public switched telephone network (PSTN) and, after successful authentication, grants dial-in users access to the network. Major components used in the Remote Access and VPN module include the following:

- **Dial-in access concentrators**: Terminate dial-in connections and authenticate individual users
- **Cisco Adaptive Security Appliances (ASA)**: Terminate IPsec tunnels, authenticate individual remote users, and provide firewall and intrusion prevention services
- **Firewalls**: Provide network-level protection of resources and stateful filtering of traffic, provide differentiated security for remote access users, authenticate trusted remote sites, and provide connectivity using IPsec tunnels
- **NIDS appliances**: Provide Layer 4 to Layer 7 monitoring of key network segments in the module

WAN and MAN and Site-to-Site VPN Module

The WAN and MAN and Site-to-Site VPN module uses various WAN technologies, including site-to-site VPNs, to route traffic between remote sites and the central site. In addition to traditional media (such as leased lines) and circuit-switched data-link technologies (such as Frame Relay and ATM), this module can use more recent WAN physical layer technologies, including Synchronous Optical Network/Synchronous Digital Hierarchy (SDH), cable, DSL, MPLS, Metro Ethernet, wireless, and service provider VPNs. This module incorporates all Cisco devices that support these WAN technologies, and routing, access control, and QoS mechanisms. Although security is not as critical when all links are owned by the enterprise, it should be considered in the network design.

KEY POINT The WAN and MAN and Site-to-Site VPN module does not include the WAN connections or links; it provides only the *interfaces* to the WAN.

Enterprise Edge Guidelines

Follow these guidelines for creating the modules within the Enterprise Edge functional area:

Step 1 Create the E-commerce module (for business-to-business or business-to-customer scenarios) when customers or partners require Internet access to business applications and database servers. Deploy a high-security policy that allows customers to access predefined servers and services yet restricts all other operations.

Step 2 Determine the connections from the corporate network into the Internet, and assign them to the Internet Connectivity module. This module should implement security to prevent any unauthorized access from the Internet to the internal network. Public web servers reside in this module or the E-commerce module.

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

Step 4 Determine which part of the edge is used exclusively for permanent connections to remote locations (such as branch offices), and assign it to the WAN and MAN and Site-to-Site VPN module. All WAN devices supporting Frame Relay, ATM, cable, MPLS, leased lines, SONET/SDH, and so on, are located here.

Service Provider Modules

Figure 3-13 shows the modules within the Service Provider functional area. The enterprise itself does not implement these modules; however, they are necessary to enable communication with other networks, using a variety of WAN technologies, and with Internet service providers (ISP). The modules within the Service Provider functional area are as follows:

- Internet Service Provider module
- PSTN module
- Frame Relay/ATM module

Figure 3-13 *Service Provider Functional Area*

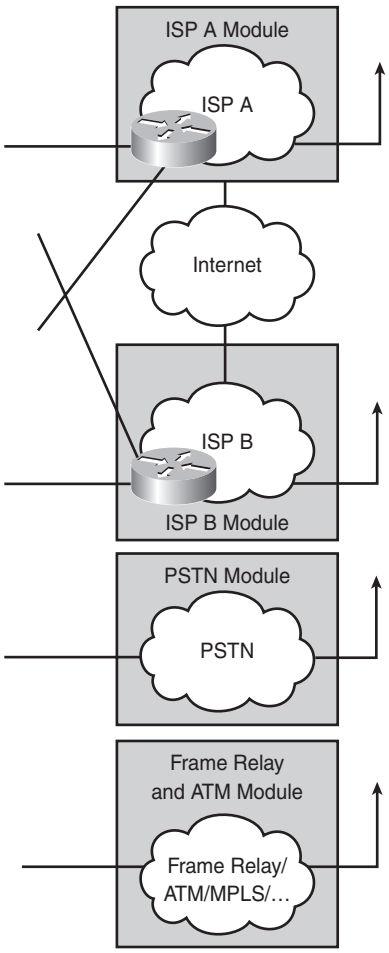

The following sections describe each of these modules.

Internet Service Provider Module

The Internet Service Provider module represents enterprise IP connectivity to an ISP network for basic access to the Internet or for enabling Enterprise Edge services, such as those in the E-commerce, Remote Access and VPN, and Internet Connectivity modules. Enterprises can connect to two or more ISPs to provide redundant connections to the Internet. The physical connection between the ISP and the enterprise can use any of the WAN technologies.

PSTN Module

KEY POINT The PSTN module represents all *nonpermanent* WAN connections.

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

Frame Relay/ATM Module

KEY POINT The Frame Relay/ATM module covers all WAN technologies for *permanent* connectivity with remote locations.

Traditional Frame Relay and ATM are still used; however, despite the module's name, it also represents many modern technologies. The technologies in this module include the following:

- Frame Relay is a connection-oriented, packet-switching technology designed to efficiently transmit data traffic at data rates of up to those used by E3 and T3 connections. Its capability to connect multiple remote sites across a single physical connection reduces the number of point-to-point physical connections required to link sites.

NOTE E3 is a European standard with a bandwidth of 34.368 megabits per second (Mbps). T3 is a North American standard with a bandwidth of 44.736 Mbps.

- ATM is a higher-speed alternative to Frame Relay. It is a high-performance, cell-oriented, switching and multiplexing technology for carrying different types of traffic.
- Leased lines provide the simplest permanent point-to-point connection between two remote locations. The carrier (service provider) reserves point-to-point links for the customer's private use. Because the connection does not carry anyone else's communications, the carrier can ensure a given level of quality. The fee for the connection is typically a fixed monthly rate.

- SONET/SDH are standards for transmission over optical networks. Europe uses SDH, whereas North America uses SONET.
- Cable technology uses existing coaxial cable TV cables. Coupled with cable modems, this technology provides much greater bandwidth than telephone lines and can be used to achieve extremely fast access to the Internet or enterprise network.
- DSL uses existing twisted-pair telephone lines to transport high-bandwidth data, such as voice, data, and video. DSL is sometimes referred to as *last-mile technology* because it is used only for connections from a telephone switching station (at a service provider) to a home or office, not between switching stations. DSL is used by telecommuters to access enterprise networks; however, more and more companies are migrating from traditional Frame Relay to DSL technology using VPNs because of its cost efficiency.
- Wireless bridging technology interconnects remote LANs using point-to-point signal transmissions that go through the air over a terrestrial radio or microwave platform, rather than through copper or fiber cables. Wireless bridging requires neither satellite feeds nor local phone service. One of the advantages of bridged wireless is its capability to connect users in remote areas without having to install new cables. However, this technology is limited to shorter distances, and weather conditions can degrade its performance.
- MPLS combines the advantages of multilayer routing with the benefits of Layer 2 switching. With MPLS, labels are assigned to each packet at the edge of the network. Rather than examining the IP packet header information, MPLS nodes use this label to determine how to process the data, resulting in a faster, more scalable, and more flexible WAN solution.

NOTE Chapter 5, "Designing Remote Connectivity," discusses WANs in more detail.

Remote Enterprise Modules

The three modules supporting remote enterprise locations are the Enterprise Branch, the Enterprise Data Center, and the Enterprise Teleworker.

Enterprise Branch Module

The Enterprise Branch module extends the enterprise by providing each location with a resilient network architecture with integrated security, Cisco Unified Communications, and wireless mobility.

A branch office generally accommodates employees who have a compelling reason to be located away from the central site, such as a regional sales office. A branch office is sometimes called a *remote site*, *remote office*, or *sales office*. Branch office users must be able to connect to the central site to access company information. Therefore, they benefit from high-speed Internet access, VPN

connectivity to corporate intranets, telecommuting capabilities for work-at-home employees, videoconferencing, and economical PSTN-quality voice and fax calls over managed IP networks. The Enterprise Branch module typically uses a simplified version of the Campus Infrastructure module design.

Enterprise Data Center Module

The Enterprise Data Center module has an architecture that is similar to the campus Server Farm module discussed earlier. The Enterprise Data Center network architecture allows the network to evolve into a platform that enhances the application, server, and storage solutions and equips organizations to manage increased security, cost, and regulatory requirements while providing the ability to respond quickly to changing business environments. The Enterprise Data Center module may include the following components:

- **At the networked infrastructure layer**: Gigabit Ethernet, 10-Gigabit Ethernet, or InfiniBand connections, with storage switching and optical transport devices

> **NOTE** InfiniBand is a high-speed switched fabric mesh technology.

- **At the interactive services layer**: Services include storage fabric services, computer services, security services, and application optimization services
- **At the management layer**: Tools include Fabric Manager (for element and network management) and Cisco VFrame (for server and service provisioning)

The remote Enterprise Data Center module also needs highly available WAN connectivity with business continuance capabilities to integrate it with the rest of the Cisco Enterprise Architecture. The Server Farm module in the campus can leverage the WAN connectivity of the campus core, but the remote Enterprise Data Center must implement its own WAN connectivity.

Enterprise Teleworker Module

The Enterprise Teleworker module provides people in geographically dispersed locations, such as home offices or hotels, with highly secure access to central-site applications and network services.

The Enterprise Teleworker module supports a small office with one to several employees or the home office of a telecommuter. Telecommuters might also be mobile users—people who need access while traveling or who do not work at a fixed company site.

Depending on the amount of use and the WAN services available, telecommuters working from home tend to use broadband or dialup services. Mobile users tend to access the company network using a broadband Internet service and the VPN client software on their laptops or via an asynchronous dialup connection through the telephone company. Telecommuters working from home might also use a VPN tunnel gateway router for encrypted data and voice traffic to and from

the company intranet. These solutions provide simple and safe access for teleworkers to the corporate network site, according to the needs of the users at the sites.

The Cisco Teleworker solution provides an easy-to-deploy, centrally managed solution that addresses both the workers' mobility needs and the enterprise's needs for lower operational costs, security, productivity, business resiliency, and business responsiveness. Small ISRs form the backbone of the Enterprise Teleworker architecture. An optional IP phone can be provided to take advantage of a centralized Cisco Unified Communications system.

Services Within Modular Networks

Businesses that operate large enterprise networks strive to create an enterprise-wide networked infrastructure and interactive services to serve as a solid foundation for business and collaborative applications. This section explores some of the interactive services with respect to the modules that form the Cisco Enterprise Architecture.

KEY POINT A network service is a supporting and necessary service, but not an ultimate solution. For example, security and QoS are not ultimate goals for a network; they are necessary to enable other services and applications and are therefore classified as network services. However, IP telephony might be an ultimate goal of a network and is therefore a network *application* (or *solution*), rather than a service.

Interactive Services

Since the inception of packet-based communications, networks have always offered a forwarding service. Forwarding is the fundamental activity within an internetwork. In IP, this forwarding service was built on the assumption that end nodes in the network were intelligent, and that the network core did not have intelligence. With advances in networking software and hardware, the network can offer an increasingly rich, intelligent set of mechanisms for forwarding information. Interactive services add intelligence to the network infrastructure, beyond simply moving a datagram between two points.

For example, through intelligent network classification, the network distinguishes and identifies traffic based on application content and context. Advanced network services use the traffic classification to regulate performance, ensure security, facilitate delivery, and improve manageability.

Network applications such as IP telephony support the entire enterprise network environment—from the teleworker to the campus to the data center. These applications are enabled by critical network services and provide a common set of capabilities to support the application's networkwide requirements, including security, high availability, reliability, flexibility, responsiveness, and compliancy.

Recall the layers of the Cisco SONA framework, illustrated in Figure 3-14. The SONA interactive services layer includes both application networking services and infrastructure services.

Figure 3-14 *Cisco SONA Framework*

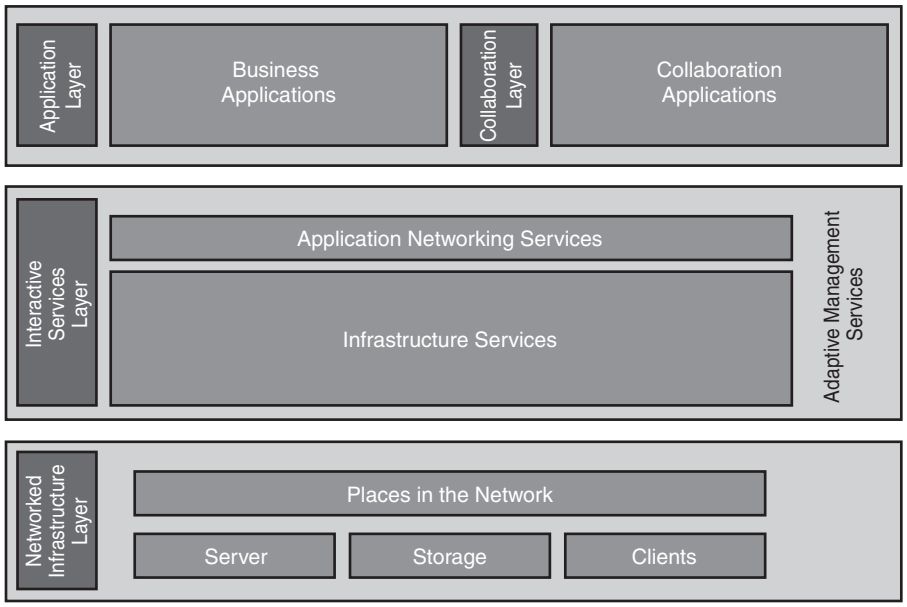

For example, the following infrastructure services (shown earlier in Figure 3-8) enhance classic network functions to support today's applications environments by mapping the application's requirements to the resources that they require from the network:

- **Security services**: Ensure that all aspects of the network are secure, from devices connecting to the network to secured transport to data theft prevention
- **Mobility services**: Allow users to access network resources regardless of their physical location
- **Storage services**: Provide distributed and virtual storage across the infrastructure
- **Voice and collaboration services**: Deliver the foundation by which voice can be carried across the network, such as security and high availability
- **Compute services**: Connect and virtualize compute resources based on the application
- **Identity services**: Map resources and policies to the user and device

Examples of network services imbedded in the infrastructure services include the following:

- **Network management**: Includes LAN management for advanced management of multilayer switches; routed WAN management for monitoring, traffic management, and access control to administer the routed infrastructure of multiservice networks; service management for managing and monitoring service level agreements (SLAs); and VPN security management for optimizing VPN performance and security administration.
- **High availability**: Ensures end-to-end availability for services, clients, and sessions. Implementation includes reliable, fault-tolerant network devices to automatically identify and overcome failures, and resilient network technologies.
- **QoS**: Manages the delay, delay variation (jitter), bandwidth availability, and packet loss parameters of a network to meet the diverse needs of voice, video, and data applications. QoS features provide value-added functionality, such as network-based application recognition for classifying traffic on an application basis, Cisco IOS IP SLAs (previously called the *service assurance agent*) for end-to-end QoS measurements, Resource Reservation Protocol signaling for admission control and reservation of resources, and a variety of configurable queue insertion and servicing functions.
- **IP multicasting**: Provides bandwidth-conserving technology that reduces network traffic by delivering a single stream of information intended for many recipients through the transport network. Multicasting enables distribution of videoconferencing, corporate communications, distance learning, software, and other applications. Multicast packets are replicated only as necessary by Cisco routers enabled with Protocol Independent Multicast and other supporting multicast protocols that result in the most efficient delivery of data to multiple receivers.

To support network applications efficiently, deploy the underlying infrastructure services in some or all modules of the enterprise network as required. These design elements can be replicated simply to other enterprise network modules as the network changes. As a result, modularization to small subsets of the overall network simplifies the network design and often reduces the network's cost and complexity.

The following sections explore some of the infrastructure services and application networking services. Network management services are described in the "Network Management Protocols and Features" section later in this chapter.

Security Services in a Modular Network Design

KEY POINT Security is an infrastructure service that increases the network's integrity by protecting network resources and users from internal and external threats.

Without a full understanding of the threats involved, network security deployments tend to be incorrectly configured, too focused on security devices, or lacking appropriate threat response options.

Security both in the Enterprise Campus (internal security) and at the Enterprise Edge (from external threats) is important. An enterprise should include several layers of protection so that a breach at one layer or in one network module does not mean that other layers or modules are also compromised; Cisco calls deploying layered security *defense-in-depth*.

Internal Security

Strongly protecting the internal Enterprise Campus by including security functions in each individual element is important for the following reasons:

- If the security established at the Enterprise Edge fails, an unprotected Enterprise Campus is vulnerable. Deploying several layers of security increases the protection of the Enterprise Campus, where the most strategic assets usually reside.
- Relying on physical security is not enough. For example, as a visitor to the organization, a potential attacker could gain physical access to devices in the Enterprise Campus.
- Often external access does not stop at the Enterprise Edge; some applications require at least indirect access to the Enterprise Campus resources. Strong security must protect access to these resources.

Figure 3-15 shows how internal security can be designed into the Cisco Enterprise Architecture.

Figure 3-15 *Designing Internal Security into the Network*

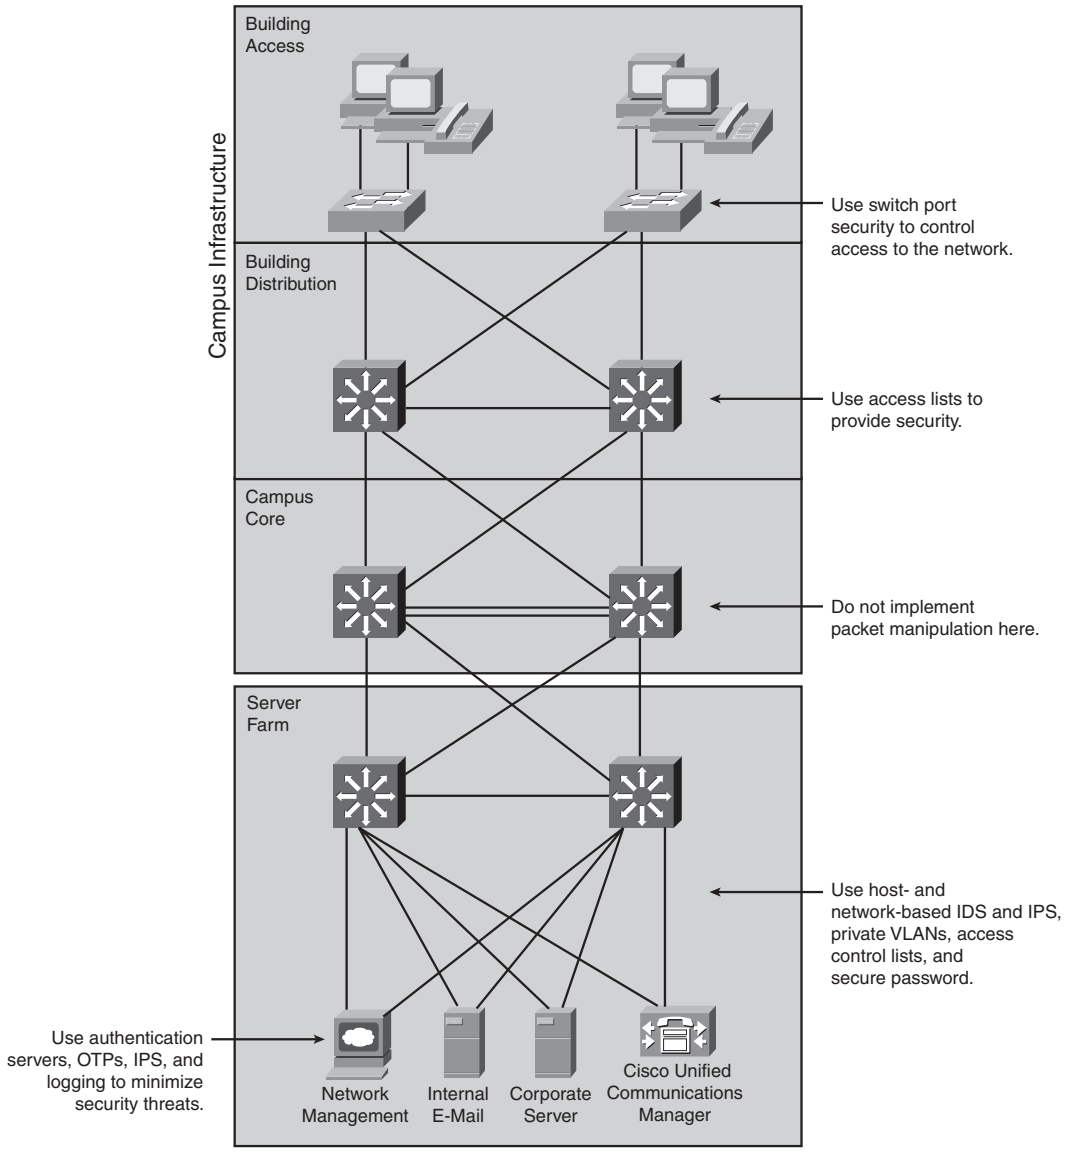

The following are some recommended security practices in each module:

- At the Building Access layer, access is controlled at the port level using the data link layer information. Some examples are filtering based on media access control addresses and IEEE 802.1X port authentication.

- The Building Distribution layer performs filtering to keep unnecessary traffic from the Campus Core. This packet filtering can be considered a security function because it does prevent some undesired access to other modules. Given that switches in the Building Distribution layer are typically multilayer switches (and are therefore Layer 3–aware), this is the first place on the data path in which filtering based on network layer information can be performed.
- The Campus Core layer is a high-speed switching backbone and should be designed to switch packets as quickly as possible; it should not perform any security functions, because doing so would slow down the switching of packets.
- The Server Farm module's primary goal is to provide application services to end users and devices. Enterprises often overlook the Server Farm module from a security perspective. Given the high degree of access that most employees have to these servers, they often become the primary goal of internally originated attacks. Simply relying on effective passwords does not provide a comprehensive attack mitigation strategy. Using host-based and network-based IPSs and IDSs, private VLANs, and access control provides a much more comprehensive attack response. For example, onboard IDS within the Server Farm's multilayer switches inspects traffic flows.

NOTE Private VLANs provide Layer 2 isolation between ports within the same broadcast domain.

- The Server Farm module typically includes network management systems to securely manage all devices and hosts within the enterprise architecture. For example, syslog provides important information on security violations and configuration changes by logging security-related events (authentication and so on). An authentication, authorization, and accounting (AAA) security server also works with a one-time password (OTP) server to provide a high level of security to all local and remote users. AAA and OTP authentication reduces the likelihood of a successful password attack.

IPS and IDS

IDSs act like an alarm system in the physical world. When an IDS detects something it considers an attack, it either takes corrective action or notifies a management system so that an administrator can take action.

HIDSs work by intercepting operating system and application calls on an individual host and can also operate via after-the-fact analysis of local log files. The former approach allows better attack prevention, and the latter approach is a more passive attack-response role.

Because of their specific role, HIDSs are often more effective at preventing specific attacks than NIDSs, which usually issue an alert only on discovering an attack. However, this specificity does not allow the perspective of the overall network; this is where NIDS excels.

Intrusion prevention solutions form a core element of a successful security solution because they detect and block attacks, including worms, network viruses, and other malware through inline intrusion prevention, innovative technology, and identification of malicious network activity.

Network-based IPS solutions protect the network by helping detect, classify, and stop threats, including worms, spyware or adware, network viruses, and application abuse. Host-based IPS solutions protect server and desktop computing systems by identifying threats and preventing malicious behavior.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks,* available at http://www.cisco.com/go/safe/, and the *Cisco Intrusion Prevention System Introduction,* available at http://www.cisco.com/en/US/products/sw/secursw/ps2113/index.html.

Authentication, Authorization, and Accounting

AAA is a crucial aspect of network security that should be considered during the network design. An AAA server handles the following:

- **Authentication—*Who?*** Authentication checks the user's identity, typically through a username and password combination.
- **Authorization—*What?*** After the user is authenticated, the AAA server dictates what activity the user is allowed to perform on the network.
- **Accounting—*When?*** The AAA server can record the length of the session, the services accessed during the session, and so forth.

The principles of strong authentication should be included in the user authentication. *Strong authentication* refers to the two-factor authentication method in which users are authenticated using two of the following factors:

- **Something you know**: Such as a password or personal identification number (PIN)
- **Something you have**: Such as an access card, bank card, or token
- **Something you are**: For example, some biometrics, such as a retina print or fingerprint
- **Something you do**: Such as your handwriting, including the style, pressure applied, and so forth

As an example, when accessing an automated teller machine, strong authentication is enforced because a bank card (something you have) and a PIN (something you know) are used.

Tokens are key-chain-sized devices that show OTPs, one at a time, in a predefined order. The OTP is displayed on the token's small LCD, typically for 1 minute, before the next password in the sequence appears. The token is synchronized with a token server, which has the same predefined list of passcodes for that one user. Therefore, at any given time, only one valid password exists between the server and a token.

This information was derived from Cisco Press's *Campus Network Design Fundamentals* by Diane Teare and Catherine Paquet, 2006.

External Threats

When designing security in an enterprise network, the Enterprise Edge is the first line of defense at which potential outside attacks can be stopped. The Enterprise Edge is like a wall with small doors and strong guards that efficiently control any access. The following four attack methods are commonly used in attempts to compromise the integrity of the enterprise network from the outside:

- **IP spoofing**: An IP spoofing attack occurs when a hacker uses a trusted computer to launch an attack from inside or outside the network. The hacker uses either an IP address that is in the range of a network's trusted IP addresses or a trusted external IP address that provides access to specified resources on the network. IP spoofing attacks often lead to other types of attacks. For example, a hacker might launch a denial of service (DoS) attack using spoofed source addresses to hide his identity.
- **Password attacks**: Using a packet sniffer to determine usernames and passwords is a simple password attack; however, the term *password attack* usually refers to repeated brute-force attempts to identify username and password information. Trojan horse programs are another method that can be used to determine this information. A hacker might also use IP spoofing as a first step in a system attack by violating a trust relationship based on source IP addresses. First, however, the system would have to be configured to bypass password authentication so that only a username is required.
- **DoS attacks**: DoS attacks focus on making a service unavailable for normal use and are typically accomplished by exhausting some resource limitation on the network or within an operating system or application.
- **Application layer attacks**: Application layer attacks typically exploit well-known weaknesses in common software programs to gain access to a computer.

DoS Attacks

DoS attacks are different from most other attacks because they are not generally targeted at gaining access to a network or its information. Rather, these attacks focus on making a service unavailable for normal use. They are 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 focus on acquiring and keeping open all the available connections supported by that server, thereby effectively locking out valid users of the server or service. DoS attacks are also implemented using common Internet protocols, such as TCP and Internet Control Message Protocol (ICMP).

Rather than exploiting a software bug or security hole, most DoS attacks exploit a weakness in the overall architecture of the system being attacked. However, some attacks compromise a network's performance 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, because it requires coordinating with the upstream network provider. If traffic meant to consume the available bandwidth is not stopped there, denying it at the point of entry into your network does little good, because the available bandwidth has already been consumed. When this type of attack is launched from many different systems at the same time, it is often referred to as a *distributed denial of service attack*.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks*, available at http://www.cisco.com/go/safe/.

Application Layer Attacks

Hackers perform application layer attacks using several different methods. One of the most common methods is exploiting well-known weaknesses in software commonly found on servers, such as SMTP, HTTP, and FTP. By exploiting these weaknesses, hackers gain access to a computer with the permissions of the account that runs the application—usually a privileged system-level account. These application layer attacks are often widely publicized in an effort to allow administrators to rectify the problem with a patch. Unfortunately, many hackers also subscribe to these same informative mailing lists and therefore learn about the attack at the same time (if they have not discovered it already).

The primary problem with application-layer attacks is that they often use ports that are allowed through a firewall. For example, a hacker who executes a known vulnerability against a web server often uses TCP port 80 in the attack. A firewall needs to allow access on that port because the web server serves pages to users using port 80. From a firewall's perspective, the attack appears as merely standard port 80 traffic.

This information was derived from the *SAFE Blueprint for Small, Midsize, and Remote-User Networks*, available at http://www.cisco.com/go/safe/.

Figure 3-16 shows these four attack methods and how they relate to the Enterprise Edge modules.

Figure 3-16 *External Threats*

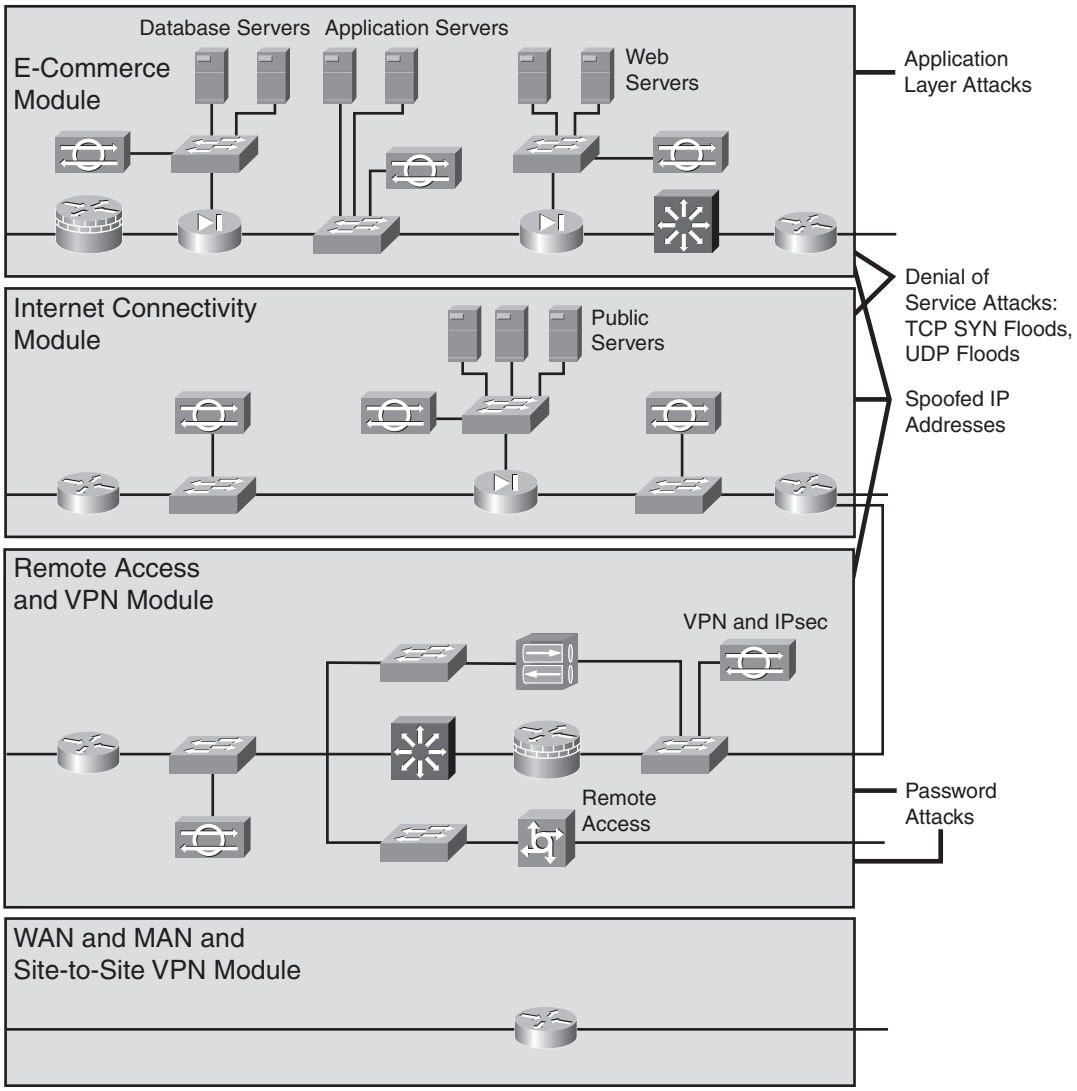

Because of the complexity of network applications, access control must be extremely granular and flexible yet still provide strong security. Tight borders between outside and inside cannot be defined, because interactions are continuously taking place between the Enterprise Edge and

Enterprise Campus. The ease of use of the network applications and resources must be balanced against the security measures imposed on the network users.

NOTE Chapter 10, "Evaluating Security Solutions for the Network," covers security in the network in more detail.

High-Availability Services in a Modular Network Design

Most enterprise networks carry mission-critical information. Organizations that run such networks are usually interested in protecting the integrity of that information. Along with security, these organizations expect the internetworking platforms to offer a sufficient level of resilience.

This section introduces another network infrastructure service: high availability. To ensure adequate connectivity for mission-critical applications, high availability is an essential component of an enterprise environment.

Designing High Availability into a Network

Redundant network designs duplicate network links and devices, eliminating single points of failure on the network. The goal is to duplicate components whose failure could disable critical applications.

Because redundancy is expensive to deploy and maintain, redundant topologies should be implemented with care. Redundancy adds complexity to the network topology and to network addressing and routing. The level of redundancy should meet the organization's availability and affordability requirements.

KEY POINT Before selecting redundant design solutions, analyze the business and technical goals and constraints to establish the required availability and affordability.

Critical applications, systems, internetworking devices, and links must be identified. Analyze the risk tolerance and the consequences of *not* implementing redundancy, and ensure that you consider the trade-offs of redundancy versus cost and simplicity versus complexity. Duplicate any component whose failure could disable critical applications.

Redundancy is not provided by simply duplicating all links. Unless all devices are completely fault-tolerant, redundant links should terminate at different devices; otherwise, devices that are not fault-tolerant become single points of failure.

KEY POINT Because many other modules access the Server Farm and Campus Core modules, they typically require higher availability than other modules.

The following types of redundancy may be used in the modules of an enterprise:

- Device redundancy, including card and port redundancy
- Redundant physical connections to critical workstations and servers
- Route redundancy
- Link redundancy
- Power redundancy, including redundant power supplies integral to the network devices and redundant power to the building's physical plant

KEY POINT The key requirement in redundancy is to provide alternative paths for mission-critical applications. Simply making the backbone fault-tolerant does not ensure high availability. For example, if communication on a local segment is disrupted for any reason, that information will not reach the backbone. End-to-end high availability is possible only when redundancy is deployed throughout the internetwork.

High Availability in the Server Farm

Improving the reliability of critical workstations and servers usually depends on the hardware and operating system software in use. Some common ways of connecting include the following:

- **Single attachment**: When a workstation or server has traffic to send to a station that is not local, it must know the address of a router on its network segment. If that router fails, the workstation or server needs a mechanism to discover an alternative router. If the workstation or server has a single attachment, it needs a Layer 3 mechanism to dynamically find an alternative router; therefore, the single-attachment method is not recommended. The available mechanisms include Address Resolution Protocol (ARP), Router Discovery Protocol (RDP), routing protocols (such as Routing Information Protocol [RIP]), Hot Standby Router Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), and Virtual Router Redundancy Protocol (VRRP). These router discovery methods are described in the "Router Discovery" sidebar on the next page.
- **Attachment through a redundant transceiver**: Physical redundancy with a redundant transceiver attachment is suitable in environments where the workstation hardware or software does not support redundant attachment options.
- **Attachment through redundant network interface cards (NIC)**: Some environments (for example, most UNIX servers) support a redundant attachment through dual NICs (primary and backup); the device driver represents this attachment as a single interface to the operating system.

- **Fast EtherChannel or Gigabit EtherChannel port bundles**: Fast EtherChannel and Gigabit EtherChannel port bundles group multiple Fast or Gigabit Ethernet ports into a single logical transmission path between a switch and a router, host, or another switch. STP treats this EtherChannel as one logical link. The switch distributes frames across the ports in an EtherChannel. This load balancing was originally done based only on MAC addresses; however, newer implementations can also load-balance based on IP addresses or Layer 4 port numbers. Source, destination, or source and destination addresses or port numbers can be used. If a port within an EtherChannel fails, traffic previously carried over the failed port reverts to the remaining ports within the EtherChannel.

Router Discovery

When a workstation has traffic to send to a station that is not local, the workstation has many possible ways of discovering the address of a router on its network segment, including the following:

- **Explicit configuration**: Most IP workstations must be configured with a default router's IP address, called the *default gateway.*

  If the workstation's default router becomes unavailable, the workstation must be reconfigured with a different router's address. Some IP stacks enable multiple default routers to be configured, but many IP stacks do not support this.

- **ARP**: Some IP workstations send an ARP frame to find a remote station. A router running proxy ARP responds with its own data link layer address; Cisco routers run proxy ARP by default.

- **RDP**: RFC 1256, *ICMP Router Discovery Messages*, specifies an extension to ICMP that allows an IP workstation and router to run RDP to allow the workstation to learn a router's address. With RDP, each router periodically multicasts a router advertisement from each of its interfaces, thereby announcing the IP address of that interface. Hosts discover the addresses of their neighboring routers simply by listening for these advertisements. When a host starts up, it multicasts a router solicitation to ask for immediate advertisements rather than waiting for the next periodic one to arrive.

NOTE RFCs are available at http://www.cis.ohio-state.edu/cs/Services/rfc/index.html.

- **Routing protocol**: An IP workstation can run RIP in passive, rather than active, mode to learn about routers. (*Active mode* means that the station sends RIP packets every 30 seconds; *passive mode* means that the station just listens for RIP packets but does not send any.) Alternatively, some workstations run the Open Shortest Path First (OSPF) protocol.

- **HSRP**: The Cisco HSRP provides a way for IP workstations to continue communicating even if their default router becomes unavailable. HSRP works by creating a virtual router that has its own IP and MAC addresses. The workstations use this virtual router as their default router.

HSRP routers on a LAN communicate among themselves to designate one router as active and one as standby. The active router sends periodic hello messages. The other HSRP routers listen for the hello messages. If the active router fails and the other HSRP routers stop receiving hello messages, the standby router takes over and becomes the active router. Because the new active router assumes the virtual router's IP and MAC addresses, end nodes do not see any change; they continue to send packets to the virtual router's MAC address, and the new active router delivers those packets. HSRP also works with proxy ARP: When an active HSRP router receives an ARP request for a node that is not on the local LAN, the router replies with the virtual router's MAC address rather than its own. If the router that originally sent the ARP reply later loses its connection, the new active router still delivers the traffic.

- **GLBP**: GLBP is similar to HSRP, but it allows packet sharing between redundant routers in a group. GLBP provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. Each host is configured with the same virtual IP address for its default gateway, and all routers in the virtual router group participate in forwarding packets.
- **VRRP**: VRRP is an election protocol that dynamically assigns responsibility for one or more virtual routers to the VRRP routers on a LAN, allowing several routers on a multiaccess link to use the same virtual IP address. A VRRP router is configured to run the VRRP protocol in conjunction with one or more other routers attached to a LAN. In a VRRP configuration, one router is elected as the virtual router master, with the other routers acting as backups in case the virtual router master fails.

Figure 3-17 shows a server-to-switch connection implemented with a redundant transceiver.

Figure 3-17 *Physical Redundancy: Redundant Transceiver*

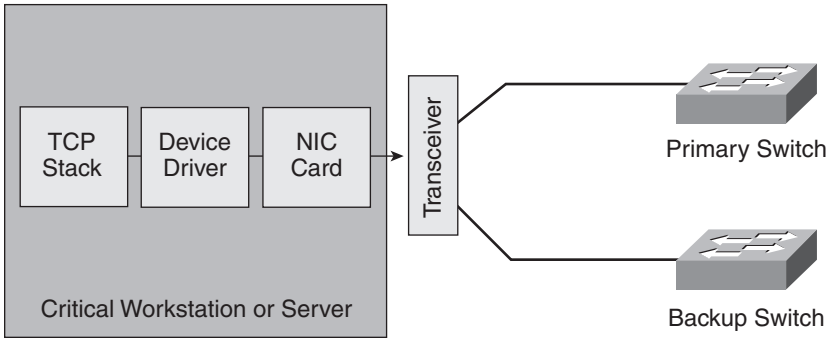

The redundant transceiver has two uplink ports that are usually connected to two access switches. The transceiver activates the backup port after it detects a link failure (carrier loss) on the primary port. The redundant transceiver can detect only physical layer failures; it cannot detect failures inside the switch or failures beyond the first switch. This type of redundancy is most often implemented on servers.

In Figure 3-18, the installation of an additional interface card in the server provides redundancy.

Figure 3-18 *Physical Redundancy: Redundant NICs*

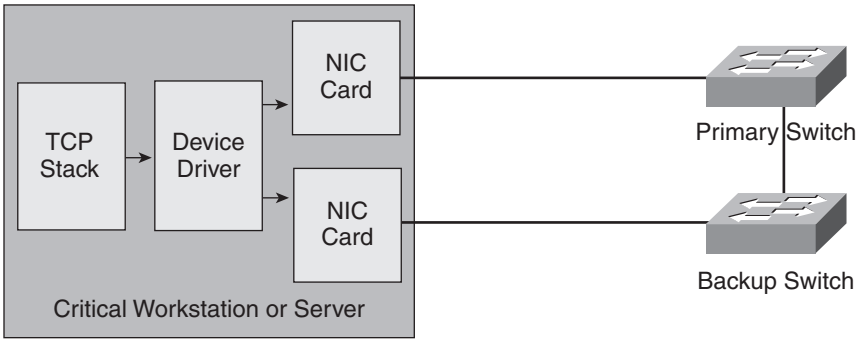

In this case, the device driver presents the configured NIC cards as a single interface (one IP address) to the operating system. If the primary link dies, the backup card activates. The two NICs might use a common MAC address, or they might use two distinct MAC addresses and send gratuitous ARP messages to provide proper IP-to-MAC address mapping on the switches when the backup interface card activates. With a redundant NIC, a VLAN shared between the two access switches is required to support the single IP address on the two server links.

NOTE The workstation sends gratuitous ARP messages to update the ARP tables and the forwarding tables on attached neighboring nodes (in this example, the Layer 2 switches).

Designing Route Redundancy

Redundant routes have two purposes:

- To minimize the effect of link failures
- To minimize the effect of an internetworking device failure

Redundant routes might also be used for load balancing when all routes are up.

Load Balancing

By default, the Cisco IOS balances between a maximum of four equal-cost paths for IP. Using the **maximum-paths** *maximum-path* router configuration command, you can request that up to 16 equally good routes be kept in the routing table (set *maximum-path* to 1 to disable load balancing).

When a packet is process-switched, load balancing over equal-cost paths occurs on a per-packet basis. When packets are fast-switched, load balancing over equal-cost paths is on a per-destination basis.

To support load balancing, keep the bandwidth consistent within a layer of the hierarchical model so that all paths have the same metric. Cisco's EIGRP includes the variance feature to load-balance traffic across multiple routes that have different metrics.

Possible ways to make the connection redundant include the following:

- Parallel physical links between switches and routers
- Backup LAN and WAN links (for example, DDR backup for a leased line)

The following are possible ways to make the network redundant:

- A full mesh to provide complete redundancy and good performance
- A partial mesh, which is less expensive and more scalable

The common approach when designing route redundancy is to implement partial redundancy by using a partial mesh instead of a full mesh and backup links to the alternative device. This protects only the most vital parts of the network, such as the links between the layers and concentration devices.

A full-mesh design forms any-to-any connectivity and is ideal for connecting a reasonably small number of devices. However, as the network topology grows, the number of links required to maintain a full mesh increases exponentially. (The number of links in a full mesh is $n(n–1)/2$, where n is the number of routers.) As the number of router peers increases, the bandwidth and CPU resources devoted to processing routing updates and service requests also increase.

A partial-mesh network is similar to the full-mesh network with some of its connections removed. A partial-mesh backbone might be appropriate for a campus network in which traffic predominantly goes into one centralized Server Farm module.

Figure 3-19 illustrates an example of route redundancy in a campus. In this example, the access layer switches are fully meshed with the distribution layer switches. If a link or distribution switch fails, an access layer switch can still communicate with the distribution layer. The multilayer switches select the primary and backup paths between the access and distribution layers based on the link's metric as computed by the routing protocol algorithm in use. The best path is placed in the forwarding table, and, in the case of equal-cost paths, load sharing takes place.

Figure 3-19 *Campus Infrastructure Redundancy Example*

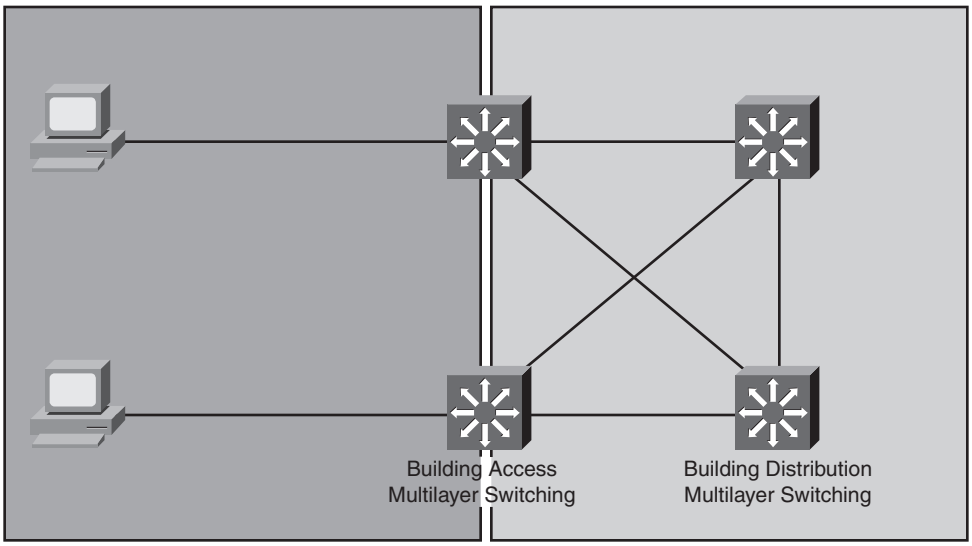

NOTE Chapter 7, "Selecting Routing Protocols for the Network," discusses routing protocols in detail.

Designing Link Redundancy

It is often necessary to provision redundant media in locations where mission-critical application traffic travels. In Layer 2–switched networks, redundant links are permitted as long as STP is running. STP guarantees one, and only one, active path within a broadcast domain, avoiding problems such as *broadcast storms* (when a broadcast continuously loops). The redundant path automatically activates when the active path goes down.

Because WAN links are often critical pieces of the internetwork, redundant media are often deployed in WAN environments. As is the case in Figure 3-20, where a Frame Relay circuit is used in parallel with a backup IPsec connection over the Internet, backup links can use different technologies. It is important that the backup provide sufficient capacity to meet the critical requirements if the primary route fails.

Figure 3-20 *Example of Enterprise Edge Link Redundancy*

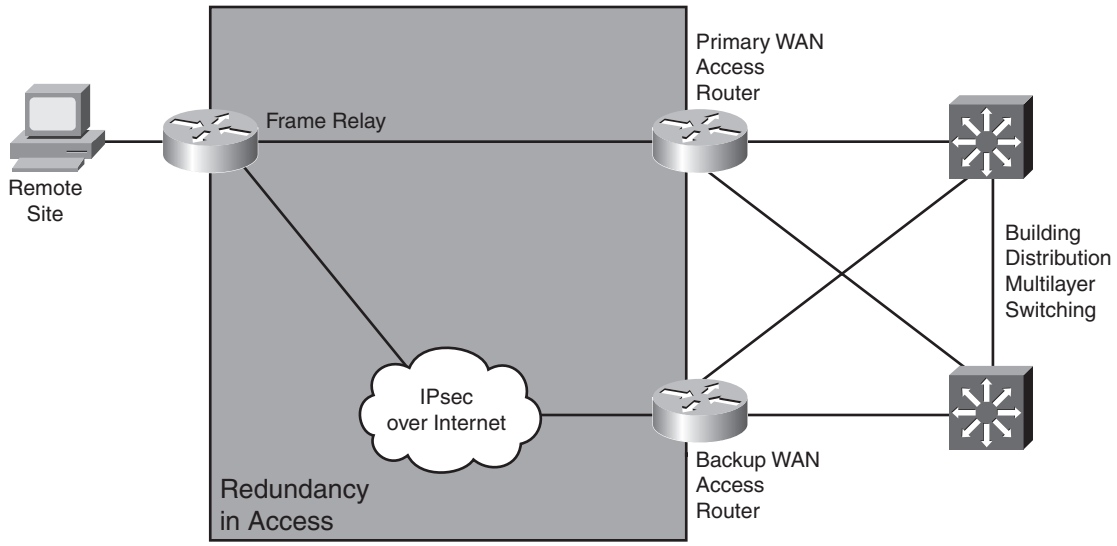

Backup links can be always-on or become active when a primary link goes down or becomes congested.

Backup Links

Backup links often use a different technology. For example, a leased line can be parallel with a backup IPsec connection over the Internet.

Using a floating static route, specify that the backup route has a higher administrative distance (used by Cisco routers to select which routing information to use) than the primary route learned from the routing protocol in use. Doing so ensures that the backup link is not used unless the primary route goes down.

When provisioning backup links, learn as much as possible about the actual physical circuit routing. Different carriers sometimes use the same facilities, meaning that your backup path is susceptible to the same failures as your primary path. Do some investigative work to ensure that the backup really is a backup.

Backup links can be used for load balancing and channel aggregation. *Channel aggregation* means that a router can bring up multiple channels (such as ISDN B channels) as bandwidth requirements increase.

Cisco supports the Multilink Point-to-Point Protocol (MLP), also referred to as *MPPP*, which is an Internet Engineering Task Force (IETF) standard for ISDN B channel (or asynchronous serial interface) aggregation. MLP does not specify how a router should accomplish the

decision-making process to bring up extra channels. Instead, it seeks to ensure that packets arrive in sequence at the receiving router. The data is encapsulated within PPP, and the datagram is given a sequence number. At the receiving router, PPP uses this sequence number to re-create the original data stream. Multiple channels appear as one logical link to upper-layer protocols.

Voice Services in a Modular Network Design

To ensure successful implementation of voice applications, network designers must consider the enterprise services and infrastructure, and its configuration. For example, to support VoIP, the underlying IP infrastructure must be functioning and robust. In other words, don't even think of adding voice to a network experiencing other problems such as congestion or network failures.

Two Voice Implementations

Voice transport is a general term that can be divided into the following two implementations:

- **VoIP**: VoIP uses voice-enabled routers to convert analog voice into IP packets or packetized digital voice channels and route those packets between corresponding locations. Users do not often notice that VoIP is implemented in the network—they use their traditional phones, which are connected to a PBX. However, the PBX is not connected to the PSTN or to another PBX, but to a voice-enabled router that is an entry point to VoIP. Voice-enabled routers can also terminate IP phones using Session Initiation Protocol for call control and signaling.
- **IP telephony**: For IP telephony, traditional phones are replaced with IP phones. A server for call control and signaling, such as a Cisco Unified Communications Manager, is also used. The IP phone itself performs voice-to-IP conversion, and no voice-enabled routers are required within the enterprise network. However, if a connection to the PSTN is required, a voice-enabled router or other gateway in the Enterprise Edge is added where calls are forwarded to the PSTN.

NOTE Earlier names for the Cisco Unified Communications Manager include Cisco CallManager and Cisco Unified CallManager.

Both implementations require properly designed networks. Using a modular approach in a voice transport design is especially important because of the voice sensitivity to delay and the complexity of troubleshooting voice networks. All Cisco Enterprise Architecture modules are involved in voice transport design.

IP Telephony Components

An IP telephony network contains four main voice-specific components:

- **IP phones**: IP phones are used to place calls in an IP telephony network. They perform voice-to-IP (and vice versa) coding and compression using special hardware. IP phones offer services such as user directory lookups and Internet access. The phones are active network devices that require power to operate; power is supplied through the LAN connection using PoE or with an external power supply.
- **Switches with inline power**: Switches with inline power (PoE) enable the modular wiring closet infrastructure to provide centralized power for Cisco IP telephony networks. These switches are similar to traditional switches, with an added option to provide power to the LAN ports where IP phones are connected. The switches also perform some basic QoS tasks, such as packet classification, which is required for prioritizing voice through the network.
- **Call-processing manager**: The call-processing manager, such as a Cisco Unified Communications Manager, provides central call control and configuration management for IP phones. It provides the core functionality to initialize IP telephony devices and to perform call setup and call routing throughout the network. Cisco Unified Communications Manager can be clustered to provide a distributed, scalable, and highly available IP telephony model. Adding more servers to a cluster of servers provides more capacity to the system.
- **Voice gateway**: Voice gateways, also called *voice-enabled routers* or *voice-enabled switches*, provide voice services such as voice-to-IP coding and compression, PSTN access, IP packet routing, backup call processing, and voice services. Backup call processing allows voice gateways to take over call processing in case the primary call-processing manager fails. Voice gateways typically support a subset of the call-processing functionality supported by the Cisco Unified Communications Manager.

Other components of an IP telephony network include a robust IP network, voice messaging and applications, and digital signal processor resources to process voice functions in hardware, which is much faster than doing it in software. These components are located throughout the enterprise network, as illustrated in Figure 3-21.

Figure 3-21 *IP Telephony Components*

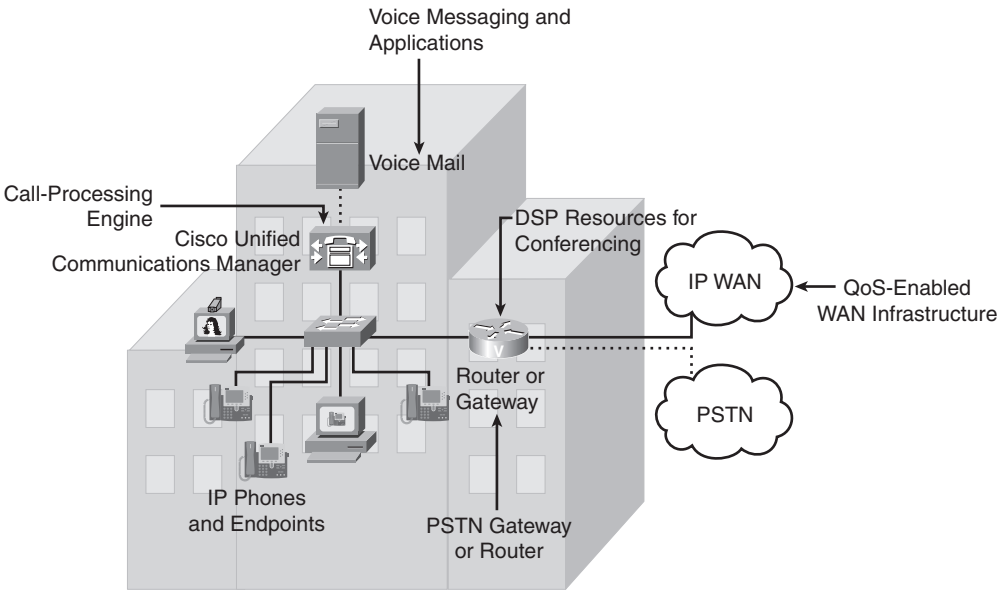

Modular Approach in Voice Network Design

Implementing voice requires deploying delay-sensitive services from end to end in all enterprise network modules. Use the modular approach to simplify design, implementation, and especially troubleshooting. Voice implementation requires some modifications to the existing enterprise network infrastructure in terms of performance, capacity, and availability because it is an end-to-end solution. For example, clients (IP phones) are located in the Building Access layer, and the call-processing manager is located in the Server Farm module; therefore, all modules in the enterprise network are involved in voice processing and must be adequately considered. Voice affects the various modules of the network as follows:

- **Building Access layer**: IP phones and end-user computers are attached to Layer 2 switches here. Switches provide power to the IP phones and provide QoS packet classification and marking, which is essential for proper voice packet manipulation through the network.
- **Building Distribution layer**: This layer performs packet reclassifications if the Building Access layer is unable to classify packets or is not within the trusted boundary. It aggregates Building Access layer switches (wiring closets) and provides redundant uplinks to the Campus Core layer.

- **Campus Core layer**: The Campus Core layer forms the network's core. All enterprise network modules are attached to it; therefore, virtually all traffic between application servers and clients traverses the Campus Core. With the advent of wire-speed multilayer gigabit switching devices, LAN backbones have migrated to switched gigabit architectures that combine all the benefits of routing with wire-speed packet forwarding.
- **Server Farm module**: This module includes multilayer switches with redundant connections to redundant Cisco Unified Communications Managers, which are essential for providing high availability and reliability.
- **Enterprise Edge**: The Enterprise Edge extends IP telephony from the Enterprise Campus to remote locations via WANs, the PSTN, and the Internet.

Figure 3-22 shows the voice network solution in the Cisco Enterprise Architecture. It illustrates how a call is initiated on an IP phone, how the call setup goes through the Cisco Unified Communications Manager, and how the end-to-end session between two IP phones is established. Note that Cisco Unified Communications Manager is involved in only the call setup.

Figure 3-22 *Voice Transport Example*

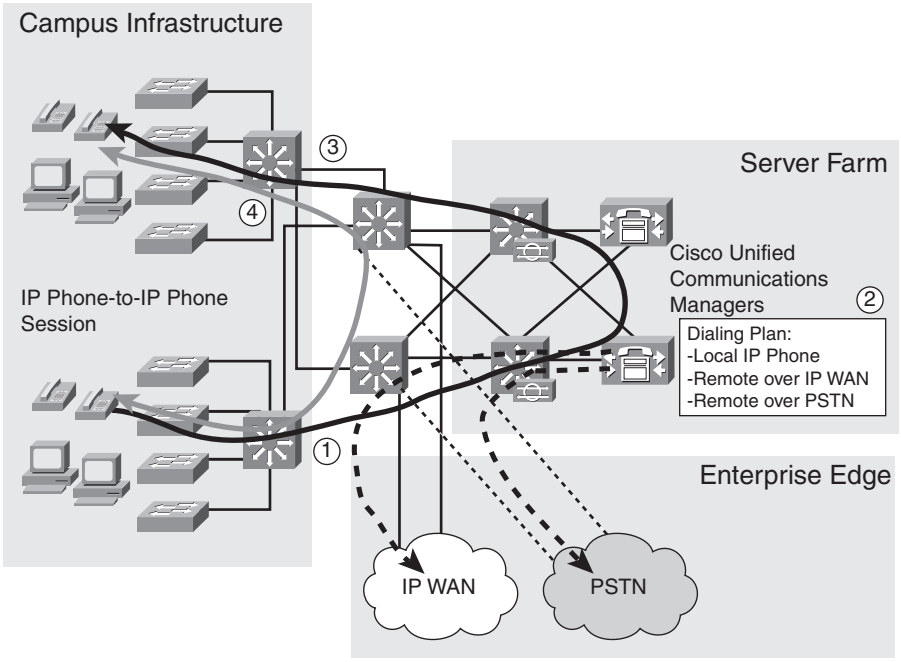

Calls destined for remote locations traverse the Enterprise Edge through the WAN and MAN and Site-to-Site VPN module or through the Remote Access and VPN module. Calls destined for

public phone numbers on the PSTN are routed over the Enterprise Edge through the Remote Access and VPN module. Calls between IP phones traverse the Building Access, Building Distribution, and Campus Core layers, and the Server Farm module. Although call setup uses all these modules, speech employs only the Building Access, Building Distribution, and, in some cases, the Campus Core layers.

Evaluating the Existing Data Infrastructure for Voice Design

When designing IP telephony, designers must document and evaluate the existing data infrastructure in each enterprise module to help determine upgrade requirements. Items to consider include the following:

- **Performance**: Enhanced infrastructure for additional bandwidth, consistent performance, or higher availability, if required, might be necessary for the converging environment. Performance evaluation includes analyzing network maps, device inventory information, and network baseline information. Links and devices such as those with high peak or busy-hour use might have to be upgraded to provide sufficient capacity for the additional voice traffic. Devices with high CPU use, high backplane use, high memory use, queuing drops, or buffer misses might have to be upgraded.
- **Availability**: Redundancy in all network modules should be reviewed to ensure that the network can meet the recommended IP telephony availability goals with the current or new network design.
- **Features**: Examine the router and switch characteristics—including the chassis, module, and software version—to determine the IP telephony feature capabilities in the existing environment.
- **Capacity**: Evaluate the overall network capacity and the impact of IP telephony on a module-by-module basis to ensure that the network meets capacity requirements and that there is no adverse impact on the existing network and application requirements.
- **Power**: Assess the power requirements of the new network infrastructure, ensuring that the additional devices will not oversubscribe existing power. Consider taking advantage of PoE capabilities in devices.

NOTE Chapter 8, "Voice Network Design Considerations," covers voice in detail.

Wireless Services in a Modular Network

A wireless LAN (WLAN) supports mobile clients connecting to the enterprise network. The mobile clients do not have a physical connection to the network because WLANs replace the Layer 1 traditional wired network (usually Category 5 cable) with radio frequency (RF)

transmissions through the air. WLANs are for local networks, either in-building, line-of-sight outdoor bridging applications, or a combination of both.

In a wireless network, many issues can arise to prevent the RF signal from reaching all parts of the facility, including multipath distortion, hidden node problems, interference from other wireless sources, and near/far issues. A site survey helps find the regions where these issues occur by defining the contours of RF coverage in a particular facility, discovering regions where multipath distortion can occur, areas where RF interference is high, and finding solutions to eliminate such issues.

Privacy and security issues must also be considered in a wireless network. Because WLANs are typically connected to the wired network, all the modules within the enterprise infrastructure must be considered to ensure the success of a wireless deployment.

Centralized WLAN Components

As illustrated in Figure 3-23, the four main components in a centralized WLAN deployment are as follows:

- **End-user devices**: A PC or other end-user device in the access layer uses a wireless NIC to connect to an access point (AP) using radio waves.
- **Wireless APs**: APs, typically in the access layer, are shared devices that function similar to a hub. Cisco APs can be either lightweight or autonomous.

  Lightweight APs are used in centralized WLAN deployments. A lightweight AP receives control and configuration from a WLAN controller (WLC) with which it is associated, providing a centralized point of management and reducing the security concern of a stolen AP. An autonomous AP has a local configuration and requires local management, which might make consistent configurations difficult and add to the cost of network management.
- **WLC**: A WLC provides management and support for wireless services such as roaming. The WLC is typically in the core layer of an enterprise network.
- **Existing switched and routed wired network**: The wireless APs connect to the wired enterprise network.

Figure 3-23 *Centralized WLAN Components*

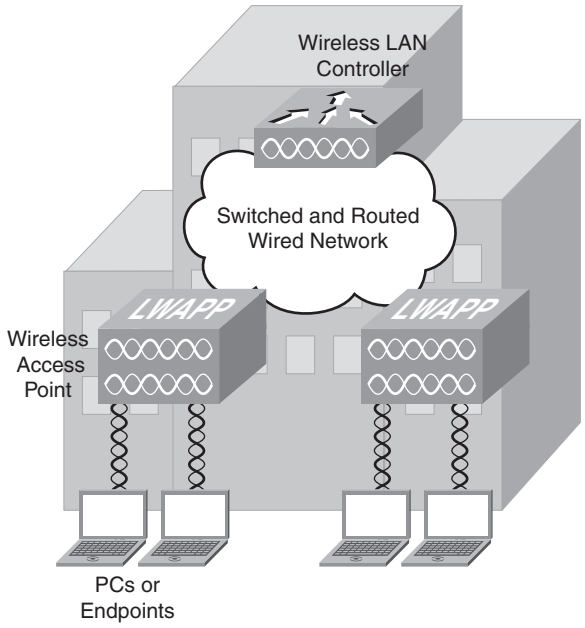

NOTE WLANs are described further in Chapter 9, "Wireless Network Design Considerations."

Application Networking Services in a Modular Network Design

Traditional networks handled static web pages, e-mail, and routine client/server traffic. Today, enterprise networks must handle more sophisticated types of network applications that include voice and video. Examples include voice transport, videoconferencing, online training, and audio and video broadcasts. Applications place increasing demands on IT infrastructures as they evolve into highly visible services that represent the face of the business to internal and external audiences.

The large amount and variety of data requires that the modern network be *application-aware*—in other words, be aware of the content carried across it to optimally handle that content. It is no longer enough simply to add more bandwidth as needs grow. Networks have had to become smarter. A new role is emerging for the network as a provider of application infrastructure services that extend the value of applications, either by improving delivery of content to users and other applications or by offloading infrastructure functions that today burden development and operations teams. Application Networking Services (ANS) provide this intelligence.

ANS Examples

Table 3-1 illustrates some sample application deployment issues that many IT managers face today and how ANS resolves these issues.

Table 3-1 *Examples of Application Deployment Issues and Solutions*

Sample Deployment Issue	Sample ANS Solution
Consolidation of data centers results in remote employees having slower access to centrally managed applications	Wide-area application services in the branch office that compress, cache, and optimize content for remote users so that they experience LAN-like responsiveness
A new web-based ordering system experiences a high proportion of abandoned orders because of poor responsiveness during the checkout process	Optimization of web streams being sent to an e-commerce portal, which reduces latency, suppresses unnecessary reloading of web objects, and offloads low-level tasks from the web server
Business partners need immediate and secure electronic access to information held in back-office applications, such as shipment information	Security and remote connectivity services that automatically validate a partner's request, route it to the appropriate back-office application, and encrypt and prioritize the response
A purchasing application needs to log and track orders over a certain value for compliance purposes	Application messaging service that intercepts purchase orders, locates the value, and logs large orders to a database according to business policy rules

ANS Components

Figure 3-24 illustrates an example of ANS deployed in offices connected over a WAN, providing LAN-like performance to users in the branch, regional, and remote offices. ANS components are deployed symmetrically in the data center and the distant offices. The ANS components in this example are as follows:

- **Cisco Wide Area Application Services (WAAS) software**: Cisco WAAS software gives remote offices LAN-like access to centrally hosted applications, servers, storage, and multimedia.
- **Cisco Wide Area Application Engine (WAE) appliance**: Cisco WAE appliances provide high-performance global LAN-like access to enterprise applications and data. WAEs use either WAAS or Application and Content Networking System [ACNS] software. WAEs help consolidate storage, servers, and so forth in the corporate data center, with only low-cost, easy-to-maintain network appliances in distant offices.

Each Cisco WAE device can be managed using the embedded command-line interface, the device Web GUI, or the Cisco WAAS Central Manager GUI. The Cisco WAAS Central Manager runs on Cisco WAE appliances and can be configured for high availability by deploying a pair of Cisco WAEs as central managers. The two central manager WAEs automatically share configuration and monitoring data.

- **Cisco 2600/3600/3700 Series Content Engine Module**: Content Engine Modules can be deployed in the data center or branch offices to optimize WAN bandwidth, accelerate deployment of mission-critical web applications, add web content security, and deliver live and on-demand business video.

Figure 3-24 *ANS Components in a WAN Environment*

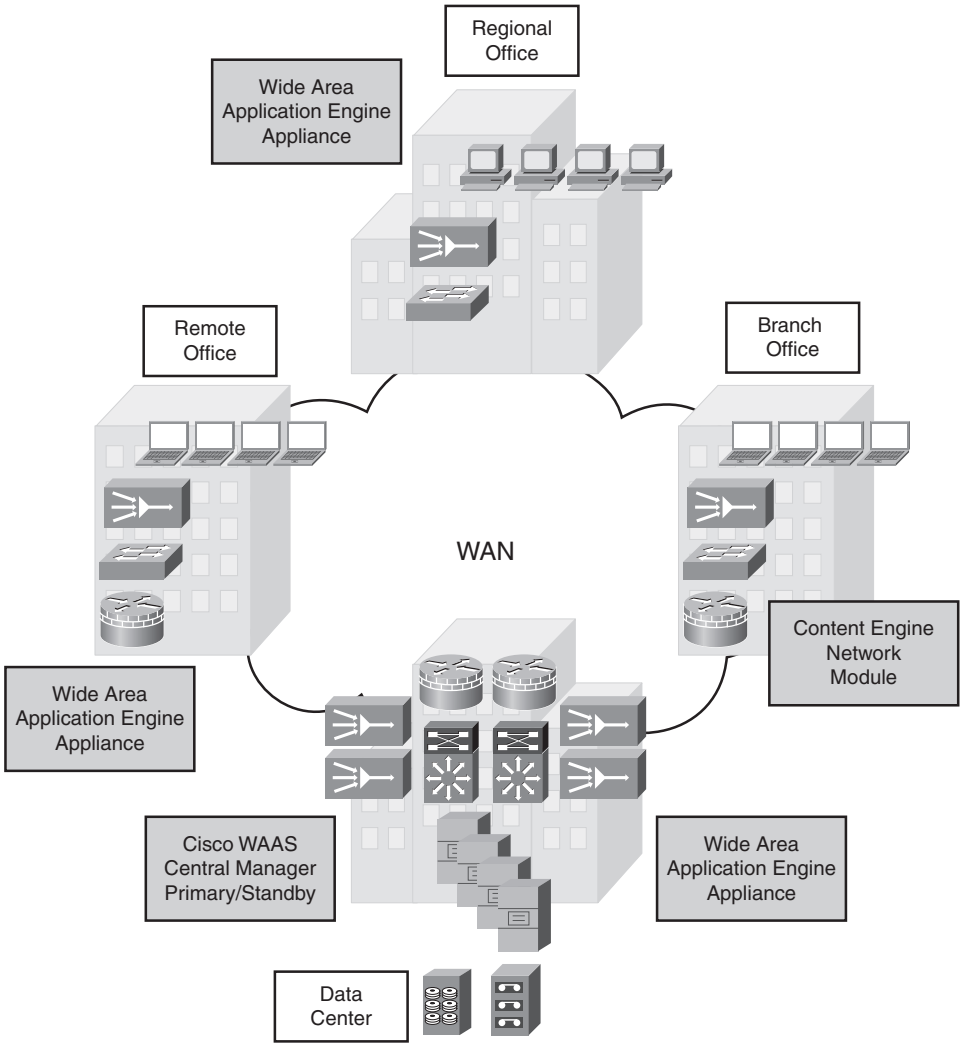

NOTE Further details on ANS are available at http://www.cisco.com/go/applicationservices/.

Network Management Protocols and Features

Proper network management is a critical component of an efficient network. Network administrators need tools to monitor the functionality of the network devices, the connections between them, and the services they provide. SNMP has become the de facto standard for use in network management solutions and is tightly connected with remote monitoring (RMON) and Management Information Bases (MIB). Each managed device in the network has several variables that quantify the state of the device. You can monitor managed devices by reading the values of these variables, and you can control managed devices by writing values into these variables.

This section introduces SNMP and describes the differences between SNMP versions 1, 2, and 3. The role of MIBs in SNMP and RMON monitoring is described, and Cisco's network discovery protocol, Cisco Discovery Protocol (CDP), is introduced. The section concludes with a description of methods for gathering network statistics.

Network Management Architecture

Figure 3-25 illustrates a generic network management architecture.

Figure 3-25 *Network Management Architecture*

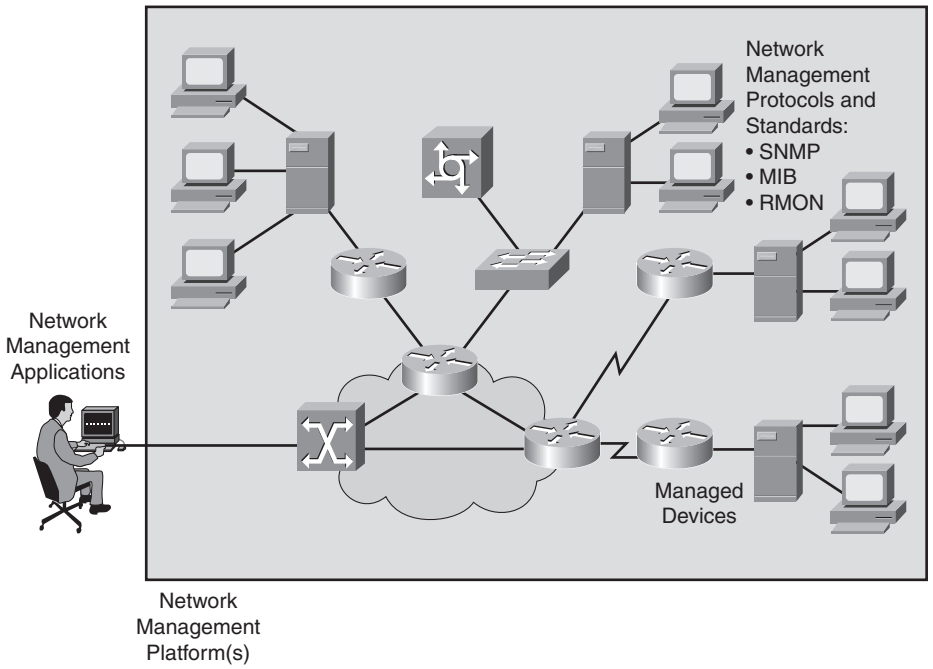

The network management architecture consists of the following:

- **Network management system (NMS)**: A system that executes applications that monitor and control managed devices. NMSs provide the bulk of the processing and memory resources that are required for network management.
- **Network management protocol**: A protocol that facilitates the exchange of management information between the NMS and managed devices, including SNMP, MIB, and RMON.
- **Managed devices**: A device (such as a router) managed by an NMS.
- **Management agents**: Software, on managed devices, that collects and stores management information, including SNMP agents and RMON agents.
- **Management information**: Data that is of interest to a device's management, usually stored in MIBs.

A variety of network management applications can be used on a network management system; the choice depends on the network platform (such as the hardware or operating system). The management information resides on network devices; management agents that reside on the device collect and store data in a standardized data definition structure known as the *MIB*.

The network management application uses SNMP or other network management protocols to retrieve the data that the management agents collect. The retrieved data is typically processed and prepared for display with a GUI, which allows the operator to use a graphical representation of the network to control managed devices and program the network management application.

Protocols and Standards

Several protocols are used within the network management architecture.

KEY POINT

SNMP is the simplest network management protocol. SNMP version 1 (SNMPv1) was extended to SNMP version 2 (SNMPv2) with its variants, which were further extended with SNMP version 3 (SNMPv3).

The MIB is a detailed definition of the information on a network device and is accessible through a network management protocol, such as SNMP.

RMON is an extension of the MIB. The MIB typically provides only static information about the managed device; the RMON agent collects specific groups of statistics for long-term trend analysis.

NOTE The ISO network management model defines the following five functional areas of network management (which are abbreviated as *FCAPS*): fault management, configuration management, accounting management, performance management, and security management.

The FCAPS model and these functional areas are rarely implemented in a single enterprise-wide network management system. A typical enterprise uses a variety of network infrastructure and service elements managed by element-specific network management systems.

NOTE Information on specific management systems for technologies such as voice, security, and wireless are provided in the relevant chapters in this book.

The following sections discuss SNMP, MIB, and RMON in detail.

SNMP

SNMP has become the de facto standard for network management. SNMP is a simple solution that requires little code to implement, which enables vendors to easily build SNMP agents for their products. In addition, SNMP is often the foundation of the network management architecture. SNMP defines how management information is exchanged between network management applications and management agents. Figure 3-26 shows the terms used in SNMP; they are described as follows:

- **Manager**: The manager, a network management application in an NMS, periodically polls the SNMP agents that reside on managed devices for the data, thereby enabling information to be displayed using a GUI on the NMS. A disadvantage of periodic SNMP polling is the possible delay between when an event occurs and when it is collected by the NMS; there is a trade-off between polling frequency and bandwidth usage.
- **Protocol**: SNMP is a protocol for message exchange. It uses the User Datagram Protocol (UDP) transport mechanism to send and retrieve management information, such as MIB variables.
- **Managed device**: A device (such as a router) managed by the manager.
- **Management agents**: SNMP management agents reside on managed devices to collect and store a range of information about the device and its operation, respond to the manager's requests, and generate traps to inform the manager about certain events. SNMP traps are sent by management agents to the NMS when certain events occur. Trap notifications could result in substantial network and agent resource savings by eliminating the need for some SNMP polling requests.

- **MIB**: The management agent collects data and stores it locally in the MIB, a database of objects about the device. Community strings, which are similar to passwords, control access to the MIB. To access or set MIB variables, the user must specify the appropriate read or write community string; otherwise, access is denied.

Figure 3-26 *NMP Is a Protocol for Management Information Exchange*

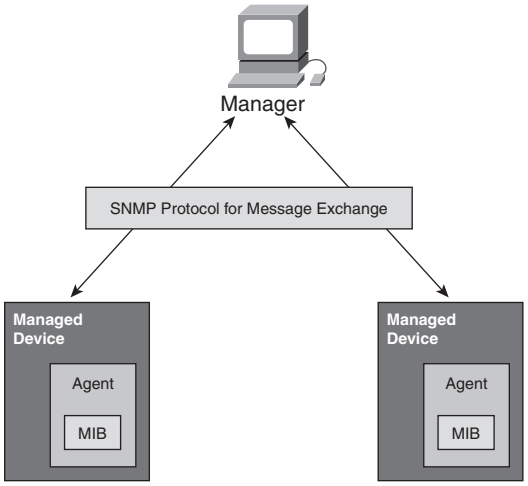

SNMPv1

The initial version of SNMP, SNMPv1 is defined in RFC 1157, *Simple Network Management Protocol (SNMP)*. The protocol's simplicity is apparent by the set of operations that are available. Figure 3-27 shows the basic SNMP messages, which the manager uses to transfer data from agents that reside on managed devices. These messages are described as follows:

- **Get Request**: Used by the manager to request a specific MIB variable from the agent.
- **Get Next Request**: Used after the initial get request to retrieve the next object instance from a table or list.
- **Set Request**: Used to set a MIB variable on an agent.
- **Get Response**: Used by an agent to respond to a manager's Get Request or Get Next Request message.
- **Trap**: Used by an agent to transmit an unsolicited alarm to the manager. A Trap message is sent when specific conditions occur, such as a change in the state of a device, a device or component failure, or an agent initialization or restart.

Figure 3-27 *SNMPv1 Message Types*

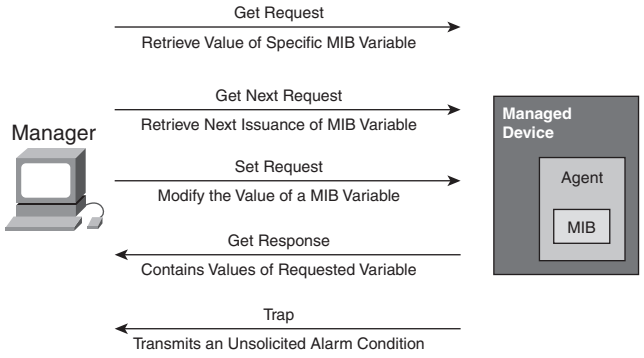

SNMPv2

SNMPv2 is a revised protocol that includes performance and manager-to-manager communication improvements to SNMP. SNMPv2 was introduced with RFC 1441, *Introduction to version 2 of the Internet-standard Network Management Framework*, but members of the IETF subcommittee could not agree on several sections of the SNMPv2 specification (primarily the protocol's security and administrative needs). Several attempts to achieve acceptance of SNMPv2 have been made by releasing experimental modified versions, commonly known as SNMPv2*, SNMPv2, SNMPv2u, SNMPv1+, and SNMPv1.5, which do not contain the disputed parts.

Community-based SNMPv2 (or SNMPv2c), which is defined in RFC 1901, *Introduction to Community-based SNMPv2*, is referred to as SNMPv2 because it is the most common implementation. The "c" stands for *community-based security* because SNMPv2c uses the same community strings as SNMPv1 for read and write access. SNMPv2 changes include the introduction of the following two new message types:

- **GetBulk message type**: Used for retrieving large amounts of data, such as tables. This message reduces repetitive requests and replies, thereby improving performance.
- **InformRequest**: Used to alert the SNMP manager of a specific condition. Unlike unacknowledged trap messages, InformRequest messages are acknowledged. A managed device sends an InformRequest to the NMS; the NMS acknowledges the receipt of the message by sending a Response message back to the managed device.

Another improvement of SNMPv2 over SNMPv1 is the addition of new data types with 64-bit counters because 32-bit counters were quickly overflowed by fast network interfaces.

On Cisco routers, Cisco IOS software release 11.3 and later versions implement SNMPv2. However, neither SNMPv1 nor SNMPv2 offers security features. Specifically, SNMPv1 and SNMPv2 can neither authenticate the source of a management message nor encrypt the message.

Because of the lack of security features, many SNMPv1 and SNMPv2 implementations are limited to a read-only capability, reducing their usefulness to that of a network monitor.

SNMPv3

SNMPv3 is the latest SNMP version to become a full standard. Its introduction has moved SNMPv1 and SNMPv2 to historic status. SNMPv3, which is described in RFCs 3410 through 3415, adds methods to ensure the secure transmission of critical data to and from managed devices. Table 3-2 lists these RFCs. Note that these RFCs make RFCs 2271 through 2275 and RFCs 2570 through 2575 obsolete.

Table 3-2 *SNMPv3 Proposed Standards Documents*

RFC Number	Title of RFC
3410	Introduction and Applicability Statements for Internet-Standard Management Framework
3411	An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
3412	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
3413	Simple Network Management Protocol (SNMP) Applications
3414	User-based Security Model (USM) for Version 3 of the Simple Network Management Protocol (SNMPv3)
3415	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)

SNMPv3 introduces the following three security levels:

- **NoAuthNoPriv**: Without authentication and without privacy (encryption).
- **AuthNoPriv**: With authentication but without privacy. Authentication is based on Hash-Based Message Authentication Code-Message Digest 5 or HMAC-Secure Hash Algorithm algorithms.
- **AuthPriv**: With authentication as described earlier and privacy using the 56-bit Cipher-Block Chaining-Data Encryption Standard encryption standard.

Security levels can be specified per user or per group of users via direct interaction with the managed device or via SNMP operations. Security levels determine which SNMP objects a user can access for reading, writing, or creating, and the list of notifications that users can receive. On Cisco routers, Cisco IOS software release 12.0 and later versions implement SNMPv3.

MIB

KEY POINT A MIB is a collection of managed objects. A MIB stores information, which is collected by the local management agent, on a managed device for later retrieval by a network management protocol.

Each object in a MIB has a unique identifier that network management applications use to identify and retrieve the value of the specific object. The MIB has a tree-like structure in which similar objects are grouped under the same branch of the MIB tree. For example, different interface counters are grouped under the MIB tree's interfaces branch.

Internet MIB Hierarchy

As shown in Figure 3-28, the MIB structure is logically represented by a tree hierarchy. The root of the tree is unnamed and splits into three main branches: Consultative Committee for International Telegraph and Telephone (CCITT), ISO, and joint ISO/CCITT.

These branches and those that fall below each category are identified with short text strings and integers. Text strings describe object names, whereas integers form object identifiers that allow software to create compact, encoded representations of the names. The *object identifier* in the Internet MIB hierarchy is the sequence of numeric labels on the nodes along a path from the root to the object. The Internet standard MIB is represented by the object identifier 1.3.6.1.2.1, which can also be expressed as iso.org.dod.internet.mgmt.mib.

Figure 3-28 *Internet MIB Hierarchy*

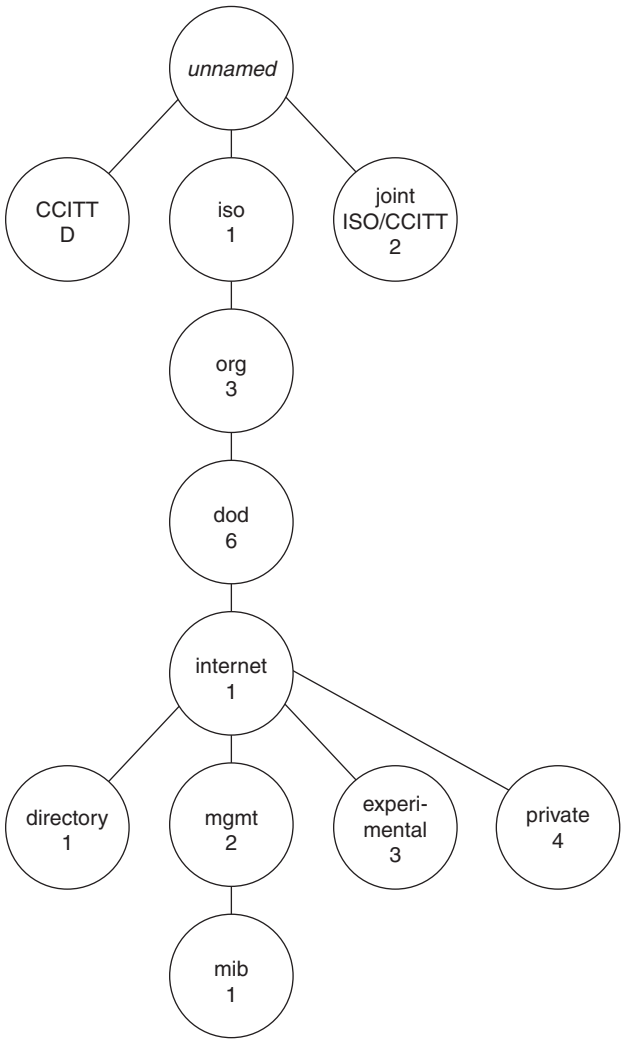

This information was adapted from the *Cisco Management Information Base (MIB) User Quick Reference*, which is available at http://www.cisco.com/univercd/cc/td/doc/product/software/ios112/mbook/index.htm.

Standard MIBs are defined in various RFCs. For example, RFC 1213, *Management Information Base for Network Management of TCP/IP-based internets: MIB-II*, defines the TCP/IP MIB.

In addition to standard MIBs, vendors can obtain their own branch of the MIB subtree and create custom managed objects under that branch. A Cisco router MIB uses both standard and private managed objects.

A Cisco router's MIB tree contains several defined standard managed objects, including from the following groups:

- Interface group (including interface description, type, physical address, counts of incoming and outgoing packets, and so forth)
- IP group (including whether the device is acting as an IP gateway, the number of input packets, the number of packets discarded because of error, and so forth)
- ICMP group (including the number of ICMP messages received, the number of messages with errors, and so forth)

The Cisco private section of the MIB tree contains private managed objects, which were introduced by Cisco, such as the following objects for routers:

- Small, medium, large, and huge buffers
- Primary and secondary memory
- Proprietary protocols

Private definitions of managed objects must be compiled into the NMS before they can be used; the result is output that is more descriptive, with variables and events that can be referred to by name.

MIB-II

MIB-II is an extension of the original MIB (which is now called *MIB-I*) and is defined by RFC 1213. MIB-II supports a number of new protocols and provides more detailed, structured information. It remains compatible with the previous version, which is why MIB-II retains the same object identifier as MIB-I (1.3.6.1.2.1).

The location of MIB-II objects is under the iso.org.dod.internet.mgmt subtree, where the top-level MIB objects are defined as follows (definitions of these objects can be found in RFC 1213):

- System (1)
- Interfaces (2)
- Address Translation (3)
- IP (4)

- ICMP (5)
- TCP (6)
- UDP (7)
- EGP (8)
- Transmission (10)
- SNMP (11)

Although the MIB-II definition is an improvement over MIB-I, the following unresolved issues exist:

- MIB-II is still a *device-centric* solution, meaning that its focus is on individual devices, not the entire network or data flows.
- MIB-II is *poll-based*, meaning that data is stored in managed devices and a management system must request (poll) it via the management protocol; the data is not sent automatically.

Cisco MIB

The Cisco private MIB definitions are under the Cisco MIB subtree (1.3.6.1.4.1.9 or iso.org.dod.internet.private.enterprise.cisco). Cisco MIB definitions supported on Cisco devices are available at http://www.cisco.com/public/mibs/.

The Cisco private MIB subtree contains three subtrees: Local (2), Temporary (3), and CiscoMgmt (9). The Local (2) subtree contains MIB objects defined before Cisco IOS software release 10.2; these MIB objects are implemented in the SNMPv1 Structure of Management Information (SMI). The SMI defines the structure of data that resides within MIB-managed objects. Beginning with Cisco IOS software release 10.2, however, Cisco MIBs are defined according to the SNMPv2 SMI and are placed in the CiscoMgmt subtree (9). The variables in the temporary subtree are subject to change for each Cisco IOS software release.

MIB Polling Guidelines

Monitoring networks using SNMP requires that the NMS poll each managed device on a periodic basis to determine its status. Frequently polling many devices or MIB variables on a device across a network to a central NMS might result in performance issues, including congestion on slower links or at the NMS connection, or overwhelming the NMS resources when processing all the collected data. The following are recommended polling guidelines:

- Restrict polling to only those MIB variables necessary for analysis.
- Analyze and use the data collected; do not collect data if it is not analyzed.

- Increase polling intervals (in other words, reduce the number of polls per period) over low-bandwidth links.
- For larger networks, consider deploying management domains, a distributed model for deploying an NMS. Management domains permit polling to be more local to the managed devices. As a result, they reduce overall management traffic across the network and the potential for one failed device or link to interrupt management visibility to the remaining network. Aggregated management data might still be centralized when management domains are used. This model is particularly appropriate for networks that already have separate administrative domains or where large campuses or portions of the network are separated by slower WAN links.
- Leverage nonpolling mechanisms such as SNMP traps, RMON, and syslog (as described in later sections of this chapter).

MIB Example

Figure 3-29 depicts SNMP MIB variable retrieval in action.

Figure 3-29 *SNMP MIB Variable Retrieval*

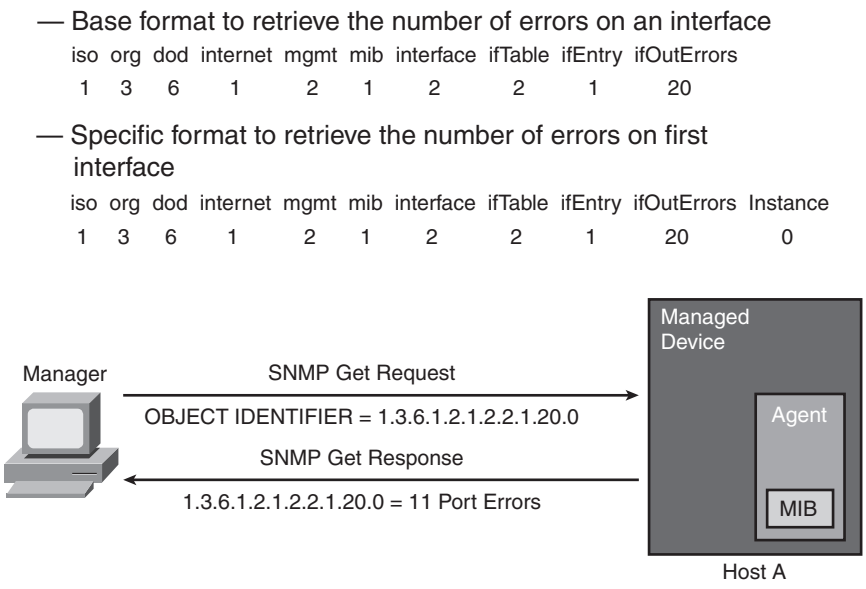

In this example, the network manager wants to retrieve the number of errors on the first interface. Starting with interface number 0, the valid range for interface numbers is *0* through *the maximum number of ports minus one*. The manager creates the SNMP Get Request message with reference to the MIB variable 1.3.6.1.2.1.2.2.1.20.0, which represents interface outgoing errors on interface

0. The agent creates the SNMP Get Response message in response to the manager's request. The response includes the value of the referenced variable. In the example, the agent returned value is 11, indicating that there were 11 outgoing errors on that interface.

RMON

KEY POINT RMON is a MIB that provides support for proactive management of LAN traffic.

The RMON standard allows packet and traffic patterns on LAN segments to be monitored. RMON tracks the following items:

- Number of packets
- Packet sizes
- Broadcasts
- Network utilization
- Errors and conditions, such as Ethernet collisions
- Statistics for hosts, including errors generated by hosts, busiest hosts, and which hosts communicate with each other

RMON features include historical views of RMON statistics based on user-defined sample intervals, alarms that are based on user-defined thresholds, and packet capture based on user-defined filters.

NOTE RMON is defined as a portion of the MIB II database. RFC 2819, *Remote Network Monitoring Management Information Base*, defines the objects for managing remote network monitoring devices. RFC 1513, *Token Ring Extensions to the Remote Network Monitoring MIB*, defines extensions to the RMON MIB for managing IEEE 802.5 Token Ring networks.

KEY POINT Without RMON, a MIB could be used to check the device's network performance. However, doing so would lead to a large amount of bandwidth required for management traffic. By using RMON, the managed device itself (via its RMON agent) collects and stores the data that would otherwise be retrieved from the MIB frequently.

RMON agents can reside in routers, switches, hubs, servers, hosts, or dedicated RMON probes. Because RMON can collect a lot of data, dedicated RMON probes are often used on routers and

switches instead of enabling RMON agents on these devices. Performance thresholds can be set and reported on if the threshold is breached; this helps reduce management traffic. RMON provides effective network fault diagnosis, performance tuning, and planning for network upgrades.

RMON1

KEY POINT RMON1 works on the data link layer (with MAC addresses) and provides aggregate LAN traffic statistics and analysis for remote LAN segments.

Because RMON agents must look at every frame on the network, they might cause performance problems on a managed device. The agent's performance can be classified based on processing power and memory.

NOTE The RMON MIB is 1.3.6.1.2.1.16 (iso.ord.dod.internet.mgmt.mib.rmon).

RMON1 Groups

RMON agents gather nine groups of statistics, ten including Token Ring, which are forwarded to a manager on request, usually via SNMP. As summarized in Figure 3-30, RMON1 agents can implement some or all of the following groups:

- **Statistics**: Contains statistics such as packets sent, bytes sent, broadcast packets, multicast packets, CRC errors, runts, giants, fragments, jabbers, collisions, and so forth, for each monitored interface on the device.
- **History**: Used to store periodic statistical samples for later retrieval.
- **Alarm**: Used to set specific thresholds for managed objects and to trigger an event on crossing the threshold (this requires an Events group).
- **Host**: Contains statistics associated with each host discovered on the network.
- **Host Top N**: Contains statistics for hosts that top a list ordered by one of their observed variables.
- **Matrix**: Contains statistics for conversations between sets of two addresses, including the number of packets or bytes exchanged between two hosts.
- **Filters**: Contains rules for data packet filters; data packets matched by these rules generate events or are stored locally in a Packet Capture group.

- **Packet Capture**: Contains data packets that match rules set in the Filters group.
- **Events**: Controls the generation and notification of events from this device.
- **TokenRing**: Contains the following Token Ring Extensions:
  - **Ring Station**—Detailed statistics on individual stations
  - **Ring Station Order**—Ordered list of stations currently on the ring
  - **Ring Station Configuration**—Configuration information and insertion/removal data on each station
  - **Source Routing**—Statistics on source routing, such as hop counts

Figure 3-30 *RMON1 Groups*

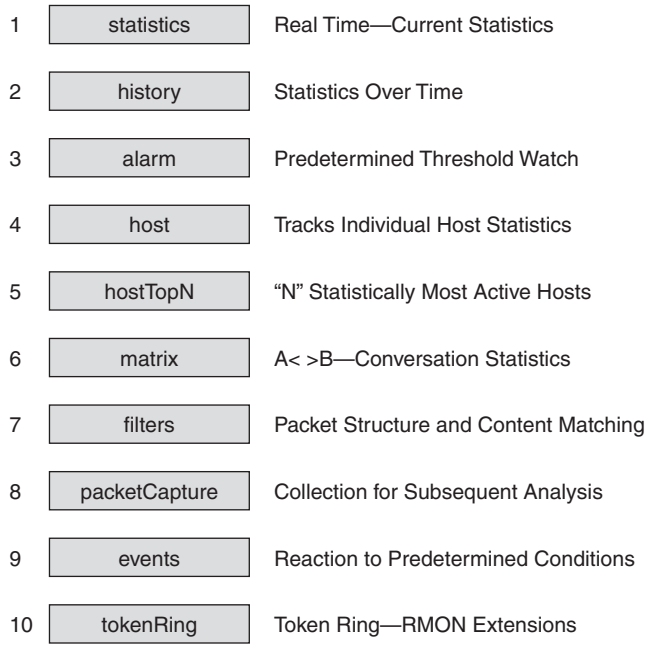

RMON1 and RMON2

RMON1 only provides visibility into the data link and the physical layers; potential problems that occur at the higher layers still require other capture and decode tools. Because of RMON1's limitations, RMON2 was developed to extend functionality to upper-layer protocols. As illustrated in Figure 3-31, RMON2 provides full network visibility from the network layer through to the application layer.

Figure 3-31 *RMON2 Is an Extension of RMON1*

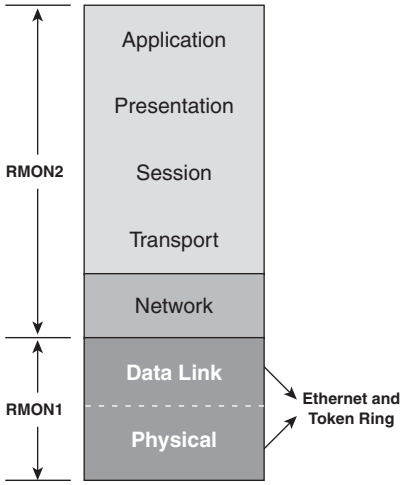

KEY POINT RMON2 is not a replacement for RMON1, but an extension of it. RMON2 extends RMON1 by adding nine more groups that provide visibility to the upper layers.

With visibility into the upper-layer protocols, the network manager can monitor any upper-layer protocol traffic for any device or subnet in addition to the MAC layer traffic.

RMON2 allows the collection of statistics beyond a specific segment's MAC layer and provides an end-to-end view of network conversations per protocol. The network manager can view conversations at the network and application layers. Therefore, traffic generated by a specific host or even a specific application (for example, a Telnet client or a web browser) on that host can be observed.

RMON2 Groups

Figure 3-32 illustrates the RMON groups that were added when RMON2 was introduced. They include the following:

- **Protocol Directory**: Provides the list of protocols that the device supports
- **Protocol Distribution**: Contains traffic statistics for each supported protocol
- **Address Mapping**: Contains network layer-to-MAC layer address mappings
- **Network Layer Host**: Contains statistics for the network layer traffic to or from each host

- **Network Layer Matrix**: Contains network layer traffic statistics for conversations between pairs of hosts
- **Application Layer Host**: Contains statistics for the application layer traffic to or from each host
- **Application Layer Matrix**: Contains application layer traffic statistics for conversations between pairs of hosts
- **User History Collection**: Contains periodic samples of user-specified variables
- **Probe Configuration**: Provides a standard way of remotely configuring probe parameters, such as trap destination and out-of-band management

Figure 3-32 *RMON2 Groups Extend RMON1 Groups*

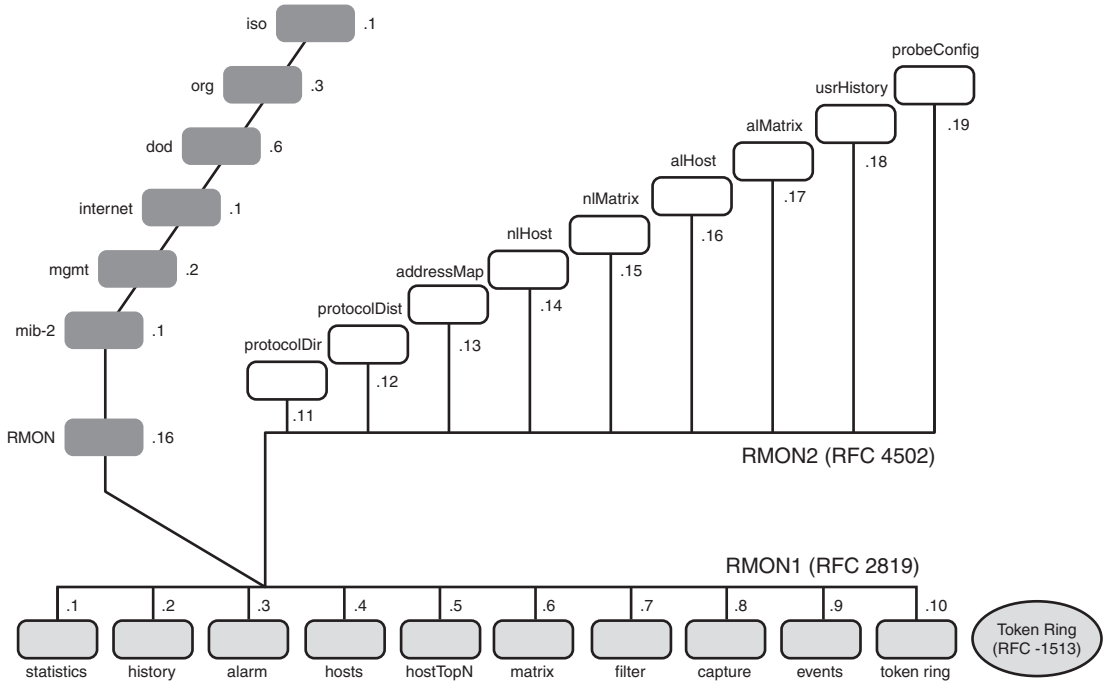

NOTE See RFC 3577, *Introduction to the Remote Monitoring (RMON) Family of MIB Modules*, for a description of RMON1, RMON2, and pointers to many of the RFCs describing extensions to RMON.

NetFlow

Cisco NetFlow is a measurement technology that measures flows that pass through Cisco devices.

NOTE NetFlow was originally implemented only on larger devices; it is now available on other devices, including ISRs.

NetFlow answers the questions of what, when, where, and how traffic is flowing in the network. NetFlow data can be exported to network management applications to further process the information, providing tables and graphs for accounting and billing or as an aid for network planning. The key components of NetFlow are the NetFlow cache or data source that stores IP flow information and the NetFlow export or transport mechanism that sends NetFlow data to a network management collector, such as the NetFlow Collection Engine.

NetFlow-collected data serves as the basis for a set of applications, including network traffic accounting, usage-based network billing, network planning, and network monitoring. NetFlow also provides the measurement base for QoS applications: It captures the traffic classification (or precedence) associated with each flow, thereby enabling differentiated charging based on QoS.

KEY POINT A *network flow* is a unidirectional sequence of packets between source and destination endpoints. Network flows are highly granular; both IP address and transport layer application port numbers identify flow endpoints. NetFlow also identifies the flows by IP protocol type, ToS, and the input interface identifier.

Non-NetFlow–enabled switching handles incoming packets independently, with separate serial tasks for switching, security services (access control lists [ACL]), and traffic measurements that are applied to each packet. Processing is applied only to a flow's first packet with NetFlow-enabled switching; information from the first packet is used to build an entry in the NetFlow cache. Subsequent packets in the flow are handled via a single, streamlined task that handles switching, security services, and data collection concurrently. Multilayer switches support multilayer NetFlow.

Therefore, NetFlow services capitalize on the network traffic's flow nature to provide detailed data collection with minimal impact on router performance and to efficiently process ACLs for packet filtering and security services. Figure 3-33 illustrates the NetFlow infrastructure.

Figure 3-33 *NetFlow Infrastructure*

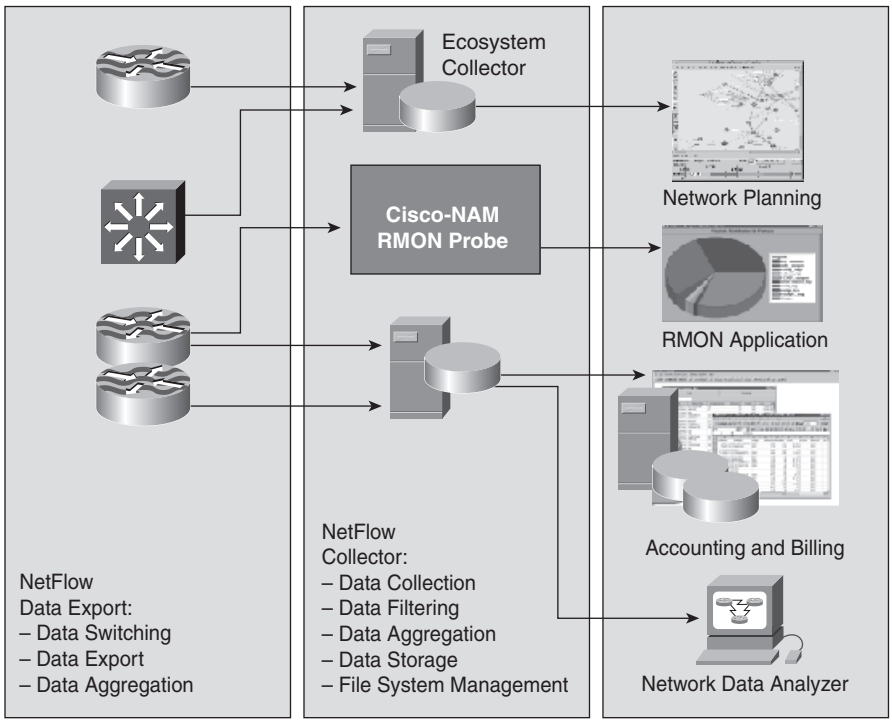

NetFlow can be configured to export data to a *flow collector*, a device that provides NetFlow export data filtering and aggregation capabilities, such as the NetFlow Collection Engine. Expired flows are grouped into NetFlow Export datagrams for export from the NetFlow-enabled device.

The focus of NetFlow used to be on IP flow information; this is changing with the Cisco implementation of a generic export transport format. NetFlow version 9 (v9) export format is a flexible and extensible export format that is now on the IETF standards track in the IP Flow Information Export (IPFIX) working group. IPFIX export is a new generic data transport capability within Cisco routers. It can be used to transport performance information from a router or switch, including Layer 2 information, security detection and identification information, IP version 6 (IPv6), multicast, MPLS, and Border Gateway Protocol (BGP) information, and so forth. NetFlow enables several key customer applications, including the following:

- **Accounting and billing**: Because flow data includes details such as IP addresses, packet and byte counts, time stamps, and application port numbers, NetFlow data provides fine-grained metering for highly flexible and detailed resource utilization accounting. For example, service providers can use this information to migrate from single-fee, flat-rate billing to more flexible

charging mechanisms based on time of day, bandwidth usage, application usage, QoS, and so forth. Enterprise customers can use the information for departmental cost recovery or cost allocation for resource utilization.

- **Network planning and analysis**: NetFlow data provides key information for sophisticated network architecture tools to optimize both strategic planning (such as whom to peer with, backbone upgrade planning, and routing policy planning) and tactical network engineering decisions (such as adding resources to routers or upgrading link capacity). This has the benefit of minimizing the total cost of network operations while maximizing network performance, capacity, and reliability.
- **Network monitoring**: NetFlow data enables extensive near-real-time network monitoring. To provide aggregate traffic- or application-based views, flow-based analysis techniques can be used to visualize the traffic patterns associated with individual routers and switches on a networkwide basis. This analysis provides network managers with proactive problem detection, efficient troubleshooting, and rapid problem resolution.
- **Application monitoring and profiling**: NetFlow data enables network managers to gain a detailed, time-based view of application usage over the network. Content and service providers can use this information to plan and allocate network and application resources (such as web server sizing and location) to meet customer demands.
- **User monitoring and profiling**: NetFlow data enables network managers to understand customer and user network utilization and resource application. This information can be used to plan efficiently; allocate access, backbone, and application resources; and detect and resolve potential security and policy violations.
- **NetFlow data warehousing and data mining**: In support of proactive marketing and customer service programs, NetFlow data or the information derived from it can be warehoused for later retrieval and analysis. For example, you can determine which applications and services are being used by internal and external users and target them for improved service. This is especially useful for service providers, because NetFlow data enables them to create a wider range of offered services. For example, a service provider can easily determine the traffic characteristics of various services and, based on this data, provide new services to the users. An example of such a service is VoIP, which requires QoS adjustment; the service provider might charge users for this service.

NetFlow Versus RMON Information Gathering

NetFlow can be configured on individual interfaces, thereby providing information on traffic that passes through those interfaces and collecting the following types of information:

- Source and destination interfaces and IP addresses

- Input and output interface numbers
- TCP/UDP source port and destination ports
- Number of bytes and packets in the flow
- Source and destination autonomous system numbers (for BGP)
- Time of day
- IP ToS

Compared to using SNMP with RMON MIB, NetFlow's information-gathering benefits include greater detail of collected data, data time-stamping, support for various data per interface, and greater scalability to a large number of interfaces (RMON is also limited by the size of its memory table). NetFlow's performance impact is much lower than RMON's, and external probes are not required.

CDP

KEY POINT

CDP is a Cisco-proprietary protocol that operates between Cisco devices at the data link layer. CDP information is sent only between directly connected Cisco devices; a Cisco device never forwards a CDP frame.

CDP enables systems that support different network layer protocols to communicate and enables other Cisco devices on the network to be discovered. CDP provides a summary of directly connected switches, routers, and other Cisco devices.

CDP is a media- and protocol-independent protocol that is enabled by default on each supported interface of Cisco devices (such as routers, access servers, and switches). The physical media must support Subnetwork Access Protocol encapsulation. Figure 3-34 illustrates the relationship between CDP and other protocols.

Figure 3-34 *CDP Runs at the Data Link Layer and Enables the Discovery of Directly Connected Cisco Devices*

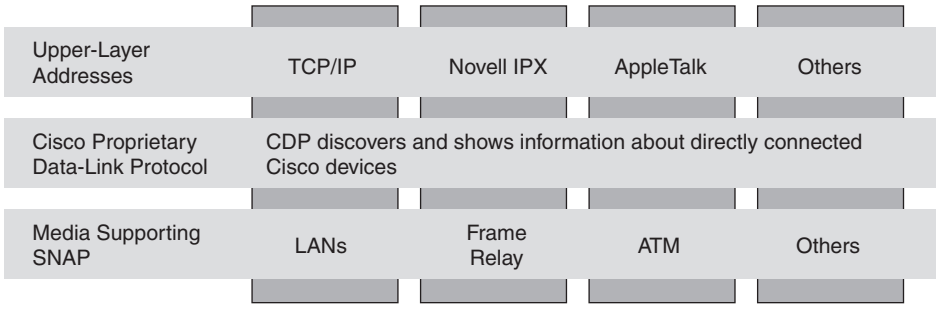

CDP Information

Information in CDP frames includes the following:

- **Device ID**: The name of the neighbor device and either the MAC address or the serial number of the device.
- **Local Interface**: The local (on this device) interface connected to the discovered neighbor.
- **Holdtime**: The remaining amount of time (in seconds) that the local device holds the CDP advertisement from a sending device before discarding it.
- **Capability List**: The type of device discovered (R—Router, T—Trans Bridge, B—Source Route Bridge, S—Switch, H—Host, I—IGMP, r—Repeater).
- **Platform**: The device's product type.
- **Port Identifier (ID)**: The port (interface) number on the discovered neighbor on which the advertisement is sent. This is the interface on the neighbor device to which the local device is connected.
- **Address List**: All network layer protocol addresses configured on the interface (or, in the case of protocols configured globally, on the device). Examples include IP, Internetwork Packet Exchange, and DECnet.

How CDP Works

As illustrated in Figure 3-35, CDP information is sent only between directly connected Cisco devices. In this figure, the person connected to Switch A can see the router and the two switches directly attached to Switch A; other devices are not visible via CDP. For example, the person would have to log in to Switch B to see Router C with CDP.

Figure 3-35 *CDP Provides Information About Neighboring Cisco Devices*

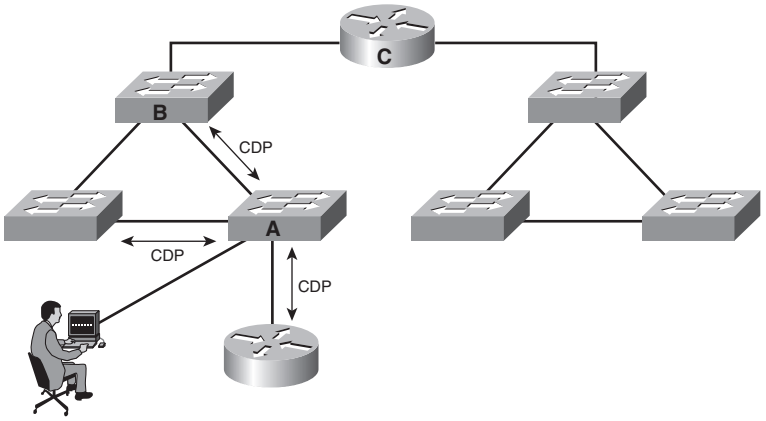

KEY POINT Cisco devices never forward a CDP frame.

CDP is a hello-based protocol, and all Cisco devices that run CDP periodically advertise their attributes to their neighbors using a multicast address. These frames advertise a time-to-live value (the holdtime, in seconds) that indicates how long the information must be retained before it can be discarded. CDP frames are sent with a time-to-live value that is nonzero after an interface is enabled. A time-to-live value of 0 is sent immediately before an interface is shut down, allowing other devices to quickly discover lost neighbors.

Cisco devices receive CDP frames and cache the received information; it is then available to be sent to the NMS via SNMP. If any information changes from the last received frame, the new information is cached and the previous information is discarded, even if its time-to-live value has not yet expired.

CDP is on by default and operates on any operational interface. However, CDP can be disabled on an interface or globally on a device. Consequently, some caveats are indicated:

- Do not run CDP on links that you do not want discovered, such as Internet connections.
- Do not run CDP on links that do not go to Cisco devices.

For security reasons, block SNMP access to CDP data (or any other data) from outside your network and from subnets other than the management station subnet.

Syslog Accounting

A system message and error reporting service is an essential component of any operating system. The syslog system message service provides a means for the system and its running processes to report system state information to a network manager.

Cisco devices produce syslog messages as a result of network events. Every syslog message contains a time stamp (if enabled), severity level, and facility.

Example 3-1 shows samples of syslog messages produced by the Cisco IOS software. The most common messages are those that a device produces upon exiting configuration mode, and the link up and down messages. If ACL logging is configured, the device generates syslog messages when

packets match the ACL condition. ACL logging can be useful to detect packets that are denied access based on the security policy that is set by an ACL.

Example 3-1 *Syslog Messages*

```
20:11:31: %SYS-5- CONFIG I: Configured from console by console


20:11:57: %LINK-5-CHANGED: Interface FastEthernet0/0, changed state to administratively
down
20:11:58: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state
to down


20:12:04: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
20:12:06: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state
to up
20:13:53: %SEC-6-IPACCESSLOGP: list internet-inbound denied udp 66.56.16.77(1029) -
> 63.78.199.4(161), 1 packet
20:14:26: %MLS-5-MLSENABLED:IP Multilayer switching is enabled
20:14:26: %MLS-5-NDEDISABLED: Netflow  Data Export disabled
20:14:26: %SYS-5-MOD_OK:Module 1 is online
20:15:47: %SYS-5-MOD_OK:Module 3 is online
20:15:42: %SYS-5-MOD_OK:Module 6 is online
20:16:27: %PAGP-5-PORTTOSTP:Port 3/1 joined bridge port 3/1
20:16:28: %PAGP-5-PORTTOSTP:Port 3/2 joined bridge port 3/2
```

Syslog messages contain up to 80 characters; a percent sign (%) follows the optional sequence number or time-stamp information if configured. Syslog messages are structured as follows:

```
seq no:timestamp: %facility-severity-MNEMONIC:description
```

The following parameters are used in the syslog messages:

- A sequence number appears on the syslog message if the **service sequence-numbers** global configuration command is configured.
- The time stamp shows the date and time of the message or event if the **service timestamps log** [**datetime** | **log**] global configuration command is configured. The time stamp can have one of three formats:
  - *mm/dd hh:mm:ss*
  - *hh:mm:ss* (for short uptimes)
  - *d h* (for long uptimes)

- **Facility**: A code consisting of two or more uppercase letters that indicate the facility to which the message refers. Syslog facilities are service identifiers used to identify and categorize system state data for error and event message reporting. A facility can be a hardware device, a protocol, or a module of the system software. The Cisco IOS software has more than 500 different facilities; the following are the most common:
  - IP
  - OSPF (OSPF protocol)
  - SYS (operating system)
  - IPsec (IP Security)
  - RSP (Route Switch Processor)
  - IF (interface)
  - LINK (data link messages)

  Other facilities include CDP, QoS, RADIUS, multicast (MCAST), MLS, TCP, VLAN trunking protocol (VTP), Telnet, and trivial file transfer protocol (TFTP).
- **Severity**: A single-digit code (from 0 to 7) that reflects the severity of the condition; the lower the number, the more serious the situation. Syslog defines the following severity levels:
  - Emergency (Level 0, which is the highest level)
  - Alert (Level 1)
  - Critical (Level 2)
  - Error (Level 3)
  - Warning (Level 4)
  - Notice (Level 5)
  - Informational (Level 6)
  - Debugging (Level 7)
- **Mnemonic**: A code that uniquely identifies the error message.
- **Description**: A text string that describes the condition. This portion of the message sometimes contains detailed information about the event, including port numbers, network addresses, or addresses that correspond to locations in the system memory address space.

NOTE For more syslog information, see http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124sup/124sms/index.htm.

Syslog Distributed Architecture

Figure 3-36 illustrates the syslog distributed architecture.

Figure 3-36 *Syslog Distributed Architecture*

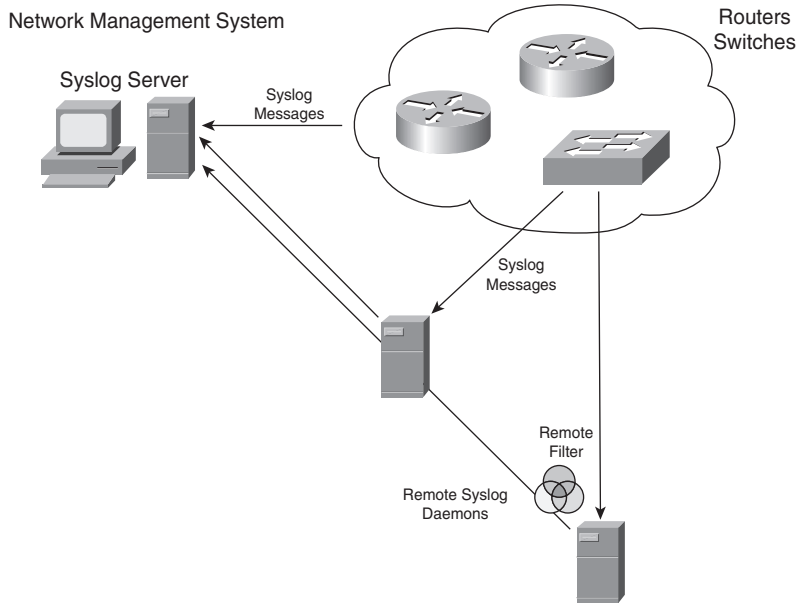

Syslog messages are sent to the console session by default. A device must be configured to send syslog messages elsewhere; the configuration includes the address of the NMS or another device. Network devices can be configured to send syslog messages directly to the NMS or to the remote network host on which a syslog analyzer is installed. A syslog analyzer conserves bandwidth on WAN links because the analyzer usually applies different filters and sends only the predefined subset of all syslog messages it receives. The analyzer filters and periodically forwards messages to the central NMS. For example, the analyzer could filter ACL logging data from other router or switch syslog entries to ensure that the ACL logging data does not overwhelm the syslog reporting tool.

The Syslog Analyzer is a CiscoWorks Resource Manager Essentials application that supports a distributed syslog server architecture for localized collection, filtering, aggregation, and forwarding of syslog data to a central syslog server for further processing and analysis. The Syslog Analyzer also supports reporting functions to automatically parse the log data into predefined or custom formats for ease of use and readability.

When it receives a syslog message, the NMS applies filters to remove unwanted messages. Filters can also be applied to perform actions based on the received syslog message, such as paging or e-mailing the network manager.

Syslog data can consume large amounts of network bandwidth and might require a very large storage capacity based on the number of devices sending syslog messages, the syslog facility and severity levels set for each, and any error conditions that may trigger excessive log messages. Therefore, it is important to enable logging only for network facilities of particular interest and to set the appropriate severity level to provide sufficient, but not excessive, detail.

KEY POINT If the collected data will not be analyzed, do not collect it.

Selectively filter and aggregate syslog data that the distributed or centralized syslog servers receive based on the requirements.

Summary

In this chapter, you learned about modularizing the network, with a focus on the following topics:

- The hierarchical network model's three layers: access, distribution, and core
- The Cisco SONA framework that integrates the enterprise-wide network
- The Cisco Enterprise Architecture functional areas:
  - Enterprise Campus: Including the Campus Infrastructure module (composed of the Campus Core layer, the Building Distribution layer, and the Building Access layer) and the Server farm module
  - Enterprise Edge: Including the E-commerce module, the Internet Connectivity module, the Remote Access and VPN module, and the WAN and MAN and Site-to-Site VPN module
  - Service Provider: Including the Internet Service Provider module, the PSTN module, and the Frame Relay/ATM module
  - Enterprise Branch
  - Enterprise Data Center
  - Enterprise Teleworker

- The infrastructure services and application networking services used within the Cisco Enterprise Architecture modules
- Security services to protect network resources and users from internal and external threats
- High-availability services to ensure adequate connectivity for mission-critical applications
- Voice services to support VoIP and IP telephony
- Wireless services to support mobile clients connecting to the enterprise network
- ANS to make the network aware of the content carried across it and to optimally handle that content
- Network management protocols and features, including SNMP, MIBs, RMON, NetFlow, CDP, and syslog

References

See the following resources for additional information:

- "Service-Oriented Network Architecture: Introduction," http://www.cisco.com/go/sona/
- *Top-Down Network Design,* Second Edition, Priscilla Oppenheimer, Cisco Press, 2004
- "Internetworking Design Basics," Cisco Internetwork Design Guide, http://www.cisco.com/univercd/cc/td/doc/cisintwk/idg4/nd2002.htm
- "SAFE Blueprint for Small, Midsize, and Remote-User Networks," http://www.cisco.com/go/safe/
- "Enterprise Architectures: Introduction," http://www.cisco.com/en/US/netsol/ns517/networking_solutions_market_segment_solutions_home.html
- *NetFlow Services Solutions Guide*, http://www.cisco.com/en/US/products/sw/netmgtsw/ps1964/products_implementation_design_guide09186a00800d6a11.html

Case Study: ACMC Hospital Modularity

This case study is a continuation of the ACMC Hospital case study introduced in Chapter 2.

Case Study General Instructions

Use the scenarios, information, and parameters provided at each task of the ongoing case study. If you encounter ambiguities, make reasonable assumptions and proceed. For all tasks, use the initial customer scenario and build on the solutions provided thus far. You can use any and all documentation, books, white papers, and so on.

In each step, you act as a network design consultant. Make creative proposals to accomplish the customer's business needs. Justify your ideas when they differ from the provided solutions. Use any design strategies you feel are appropriate. The final goal of each case study is a paper solution.

Appendix A, "Answers to Review Questions and Case Studies," provides a solution for each step based on assumptions made. There is no claim that the provided solution is the best or only solution. Your solution might be more appropriate for the assumptions you made. The provided solution helps you understand the author's reasoning and allows you to compare and contrast your solution.

In this case study, you apply the Cisco Enterprise Architecture to the ACMC Hospital network requirements and develop a high-level view of the planned network hierarchy. Complete the following steps:

Step 1 Consider each of the functional areas of the Cisco Enterprise Architecture:

- Enterprise Campus: Including the Campus Infrastructure module (composed of the Campus Core layer, the Building Distribution layer, and the Building Access layer) and the Server farm module
- Enterprise Edge: Including the E-commerce module, the Internet Connectivity module, the WAN and MAN and Site-to-Site VPN module, and the Remote Access and VPN module
- Enterprise Branch
- Enterprise Data Center
- Enterprise Teleworker

Mark up the existing network diagram, provided in Figure 3-37, indicating where each of the modules would be at a high level.

Figure 3-37 *Existing ACMC Hospital Network*

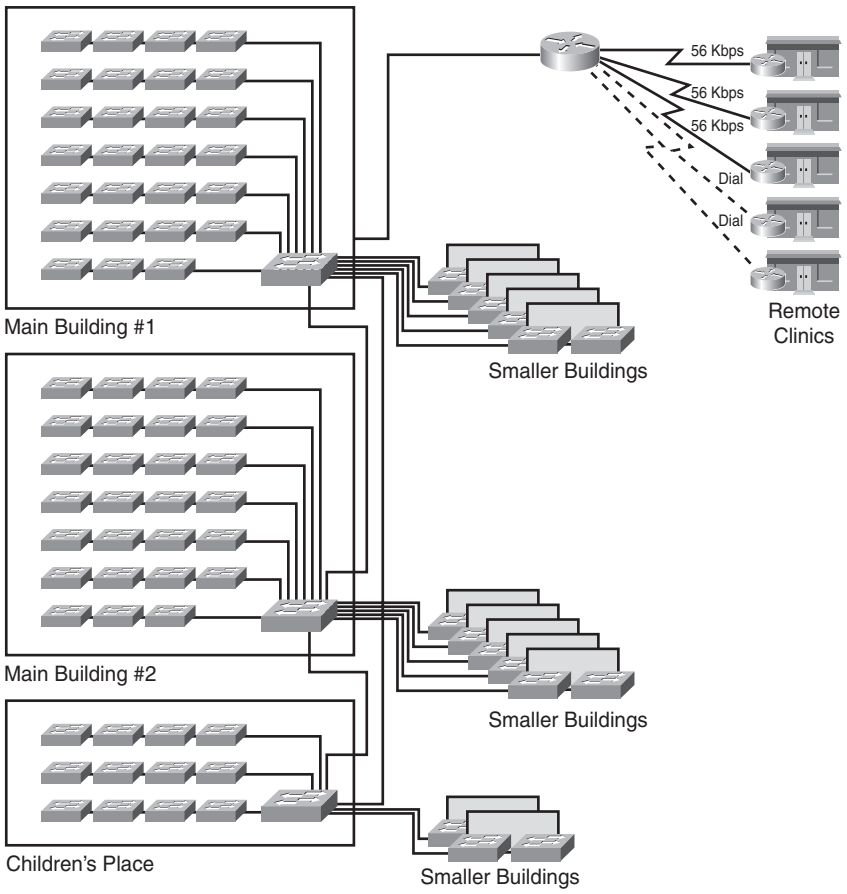

Step 2 List some key considerations or functions for each of the modules in the Cisco Enterprise Architecture. Indicate whether each module is used in the ACMC Hospital network.

Step 3 Since the time initial discussions with ACMC occurred, the following additional requirements have surfaced:

- The staff needs Internet access for purchasing supplies and reviewing research documents and new medical products.
- There has been some discussion about allowing employees to telecommute.

- ACMC has a web server for a patient communications and community relations service called "Text a Nurse." This for-fee service allows a patient to send a text message to the hospital, requesting medical advice.

How does this new information change the design? Incorporate the changes into your high-level design, and update the list of modules and considerations.

Step 4 Which of the following infrastructure or network services are immediately applicable to your design?

- Security services
- Voice services
- Wireless
- Network management
- High availability
- QoS
- Multicast

Are there specific locations or modules where some of these services are particularly relevant?

Step 5 Indicate where redundancy should be supported in the design.

Review Questions

Answer the following questions, and then refer to Appendix A for the answers.

1. Figure 3-38 presents a sample hierarchically structured network. Some of the devices are marked with letters. Map the marked devices to the access, distribution, and core layers in this figure.
2. Describe the role of each layer in the hierarchical network model.
3. True or false: Each layer in the hierarchical network model must be implemented with distinct physical devices.
4. Which two statements are true?
   a. UplinkFast immediately unblocks a blocked port after root port failure.
   b. PortFast immediately puts a port into the forwarding state.
   c. UplinkFast immediately puts a port into the forwarding state.
   d. PortFast immediately unblocks a blocked port after root port failure.

5. What features of a multilayer switch could be used in the access layer?
6. Which layer in the hierarchical model provides media translation?

Figure 3-38 *Hierarchical Network*

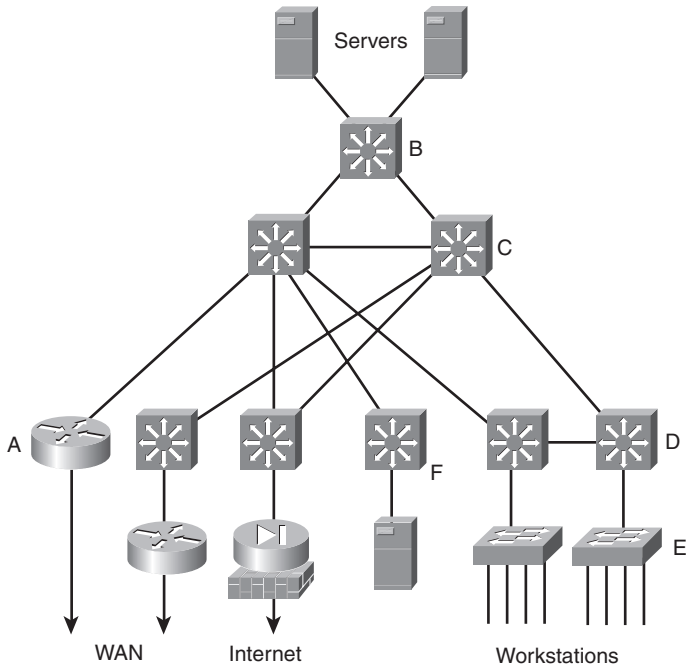

7. Why might the distribution layer need to redistribute between routing protocols?
8. What are three roles of the hierarchical model's core layer?
   a. Provide fast and efficient data transport
   b. Provide maximum availability and reliability
   c. Provide access to the corporate network via some wide-area technology
   d. Implement security policies
   e. Delineate broadcast domains
   f. Implement scalable routing protocols
9. What is a benefit of using multilayer switching in the core network layer?
10. What are the six major functional areas in the Cisco Enterprise Architecture?
11. What are the modules and layers within the Enterprise Campus functional area?
12. The Enterprise Edge functional area includes which modules?

13. The Service Provider functional area is composed of which modules?

14. Which module of the Cisco Enterprise Architecture includes wireless bridging connectivity to remote locations?

15. What is an advantage of using the Cisco Enterprise Architecture?

16. What is the Campus Core layer's role?

17. Indicate which types of devices would be found in each of these modules (note that some devices are found in more than one module).

    Modules:

    - E-commerce module
    - Internet Connectivity module
    - Remote Access and VPN module

    Devices:

    - Web servers
    - SMTP mail servers
    - Firewalls
    - Network Intrusion Detection System (NIDS) appliances
    - DNS servers
    - ASAs
    - Public FTP servers

18. What is the role of the Service Provider functional area?

19. Which other module has a design similar to that of the Enterprise Branch module?

20. Which other module has an architecture similar to that of the Enterprise Data Center module?

21. Which module of the Cisco Enterprise Architecture provides telecommuter connectivity?

22. The SONA interactive services layer includes both _____ services and _____ services.

23. How can the Server Farm module be involved in an organization's internal security?

24. High availability from end to end is possible only when _____ is deployed throughout the internetwork.

25. What is the purpose of designing route redundancy in a network?

26. A full-mesh design is ideal for connecting a _____ number of devices.
    a. small
    b. large
27. True or false: Backup links can use different technologies.
28. What components are required for IP telephony?
29. What role does the Building Access layer play in voice transportation?
30. What should you consider when evaluating an existing data infrastructure for IP telephony?
31. What are the main components of a centralized WLAN deployment?
32. What is a Cisco WAE appliance?
33. What is a network management agent?
34. How does an SNMPv1 manager request a list of data?
35. How does an SNMPv2 manager request a list of data?
36. What is the MIB structure?
37. How are private MIB definitions supported?
38. What are the RMON1 groups?
39. What groups are added to the RMON1 groups by RMON2?
40. How does RMON simplify proactive network management?
41. What is a NetFlow network flow?
42. How does NetFlow compare to RMON?
43. At which layer does CDP work?
44. Two routers are connected via Frame Relay, but ping is not working between them. How could CDP help troubleshoot this situation?
45. What are the syslog severity levels?
46. What syslog severity level is indicated by the messages in Example 3-2?

Example 3-2 *Sample Message for Question 46*

```
20:11:58: %LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/0, changed state
to down
20:12:04: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to up
```