Introduction

The Cisco Designing for Cisco Internetwork Solutions (DESGN) exam is the required exam for the Cisco Certified Design Associate (CCDA) certification. Objectives for the DESGN exam include the following:

- Describe a systematic and modular approach to design.
- Design enterprise campus, enterprise data center, enterprise edge, and remote modules.
- Assign an appropriate IP addressing scheme.
- Select an appropriate routing protocol.
- Specify security solutions.
- Provide support for voice traffic.
- Offer a solution for basic wireless connectivity.

These Quick Reference Sheets summarize the main topics presented on the DESGN exam. The information presented represents the content covered on exam number 640-863.

Strategic Network Design

This section introduces you to the Cisco Service-Oriented Network Architecture (SONA) framework for network design. In addition, you learn how to examine characteristics of an existing network, while determining design requirements. Finally, this section discusses Cisco's top-down approach to network design.

Cisco Service-Oriented Network Architecture

Cisco recently updated its Architecture for Voice Video and Integrated Data (AVVID) design approach to the Intelligent Information Network (IIN). IIN is a complete architecture that is more all encompassing than AVVID.

The three phases of constructing an IIN are as follows:

- **Integrated transport**—Voice, data, and video are all converged onto a single transport.
- **Integrated services**—Services, such as VoIP or storage networking, rely on the underlying network transport mechanisms.
- **Integrated applications**—Applications (for example, Cisco IP Communicator) leverage services (for example, VoIP), which rely on the network transport.

The Cisco architectural approach to designing an IIN is their SONA framework. Figure 1-1 shows individual IIN components and how those components are categorized by SONA's three layers: networked infrastructure layer, infrastructure services layer, and application layer.

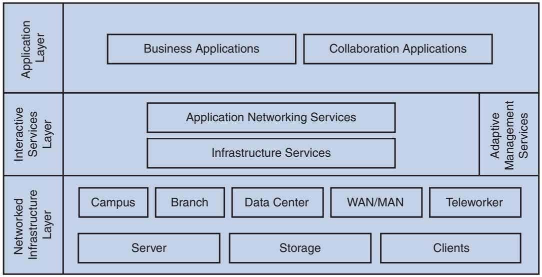
Application Layer
Business Applications
Collaboration Applications
Interactive Services Layer
Application Networking Services
Infrastructure Services
Adaptive Management Services
Networked Infrastructure Layer
Campus
Branch
Data Center
WAN/MAN
Teleworker
Server
Storage
Clients

FIGURE 1-1 SONA layers.

SONA offers the following benefits to a network design:

- Functionality
- Scalability
- Availability
- Performance
- Manageability
- Efficiency

Identifying Design Requirements

Cisco categorizes a network's life cycle into six phases identified with the acronym PPDIOO. The components of PPDIOO are as follows:

1. **Prepare**—This phase involves determining the network's requirements, formulating a network strategy, and suggesting a conceptual architecture of the network.
2. **Plan**—This phase compares the existing network with the proposed network to help identify tasks, responsibilities, milestones, and resources required to implement the design.
3. **Design**—This phase clearly articulates the detailed design requirements.
4. **Implement**—This phase integrates equipment into the existing network (without disrupting the existing network) to meet design requirements.
5. **Operate**—This phase entails the day-to-day network operation, while responding to any issues that arise.
6. **Optimize**—This phase gathers feedback from the Operate phase to potentially make adjustments in the existing network. Changes might be implemented to address ongoing network support issues.

PPDIOO's life-cycle approach offers the following benefits:

- PPDIOO reduces total cost of ownership (TCO).
- PPDIOO improves network availability.
- PPDIOO allows business networks to quickly respond to changing needs.
- PPDIOO accelerates access to network applications and services.

Designing a network in conjunction with the PPDIOO approach involves three steps:

1. Identify customer requirements.

   To identify customer requirements, obtain the following pieces of information:

   - Network applications
   - Network services
   - Business goals
   - Constraints imposed by the customer
   - Technical goals
   - Constraints imposed by technical limitations

2. Identify characteristics of the current network.

   To identify characteristics of the current network, perform the following tasks:

   - Collect existing network documentation (with the understanding that the documentation might be somewhat dated and unreliable), and interview organizational representatives to uncover information not available in the documentation.

- Conduct a network audit to identify information such as network traffic types, congestion points, and suboptimal routes.
- Supplement the information collected in the two previous tasks by performing a network traffic analysis with tools such as Cisco Discovery Protocol (CDP), Network Based Application Recognition (NBAR), NetFlow, Cisco CNS NetFlow Collection Engine, Open Source Cacti, Network General Sniffer, WildPackets EtherPeek and AiroPeek, SolarWinds Orion, Wireshark, and Remote Monitoring (RMON) probes.

3. Design the network topology.

Using information collected in Steps 1 and 2, you are ready to begin your network design. Although designing a network can be a daunting task, Cisco's recommended top-down design approach assists the designer by breaking the design process into smaller and more manageable steps. The term ***top-down*** refers to beginning at the top of the OSI reference model (that is, the application layer) and working your way down through the underlying layers, as shown in Figure 1-2.

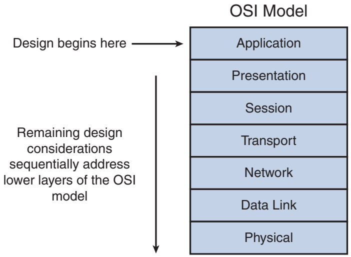

FIGURE 1-2 Top-down design strategy.

Using a top-down design strategy as opposed to a bottom-up design strategy (that is, where the design begins at the physical layer of the OSI model and works its way up) provides the following benefits:

- Does a better job of including specific customer requirements
- Offers a more clearly articulated "big picture" of the desired network for both the customer and the designer
- Lays the foundation for a network that not only meets existing design requirements but provides for scalability to meet future network enhancements

When using the OSI reference model in the top-down design approach, the designer should determine what design decisions, if any, are required for each of the seven layers. For example, when considering the application layer, the designer might determine that voice applications such as the Cisco IP Contact Center and the Cisco Unity converged messaging system are applications needed for the design.

Network layer design decisions might include the selection of a routing protocol (for example, Enhanced Interior Gateway Routing Protocol [EIGRP] or Open Shortest Path First Protocol [OSPF]). Also, when analyzing the network layer, the designer might need to determine an appropriate IP addressing scheme for the network (for example, the use of private versus public IP addresses and subnet masks to be used) to provide for future network scalability.

Physical layer and data link layer design decisions might involve the selection of LAN/WAN technologies (for example, Gigabit Ethernet, Fast Ethernet, Frame Relay, ATM, or PPP) to provide for media transport.

With the multitude of design decisions required in larger networks, network designers often benefit from network design tools such as the following:

- **Network modeling tools**—Generate suggested configurations based on input information, which can then be further customized (for example, adding redundancy or support for additional sites)
- **Strategic analysis tools**—Enable a network designer to experiment with various "what-if" scenarios and observe resulting network effects
- **Decision tables**—Record design decisions based on network requirements
- **Simulation and verification tools/services**—Verify design decisions in a simulated environment to reduce the need to implement a pilot network

Even with the availability of simulation tools, some network designs still benefit from building a small prototype network to serve as a proof of concept. Such prototype networks are commonly known as ***pilot networks***.

Modular Network Design

For many years, Cisco recommended a three-layer network design model: access layer, distribution layer, and core layer. However, to provide for enhanced scalability and flexibility, Cisco later introduced the ***Cisco Enterprise Architecture***, which categorizes enterprise networks into six modules. The three layers of the Cisco Service-Oriented Network Architecture (SONA) can be found in each of these six modules. Specifically, each module can contain its own network infrastructure, services, and applications. This section explores the design considerations surrounding the modules that comprise the Cisco Enterprise Architecture.

Designing the Network Hierarchy

Traditionally, Cisco prescribed a three-layer model for network designers. Those three layers, as shown in Figure 2-1, are as follows:

- **Access layer**—Typically, wiring closet switches connecting to end-user stations
- **Distribution layer**—An aggregation point for wiring closet switches, where routing and packet manipulation occur
- **Core layer**—The network backbone where high-speed traffic transport is the main priority

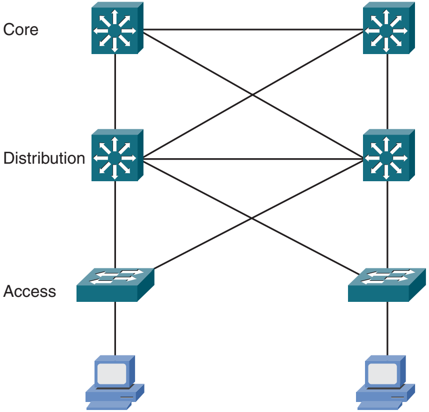

FIGURE 2-1 Three-layer hierarchical model.

Modularizing Network Design

The three-layer hierarchical approach suffers from scalability limitations. For today's enterprise networks, Cisco developed the Cisco Enterprise Architecture. The functional areas that comprise the Enterprise Architecture, as illustrated in Figure 2-2, include the following:

- **Enterprise campus**—The portion of the network design providing performance, scalability, and availability that defines operation within the main campus
- **Enterprise edge**—An aggregation point for components at the edge of the network (for example, Internet and MAN/WAN connectivity) that routes traffic to and from the Enterprise Campus functional area
- **WAN and Internet**—The portion of the network made available by a service provider (for example, Frame Relay or ATM)
- **Enterprise branch**—Remote network locations that benefit from extended network services, such as security
- **Enterprise data center**—A consolidation of applications, servers, and storage solutions (similar to a campus data center)
- **Enterprise teleworker**—A collection of small office/home office (SOHO) locations securely connected to the enterprise edge via an Internet service provider (ISP) or public switched telephone network (PSTN)

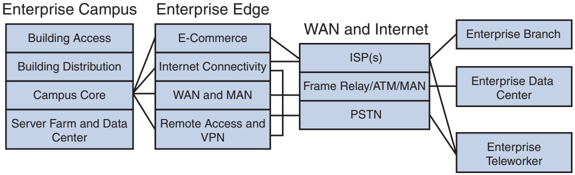

FIGURE 2-2 Cisco Enterprise Architecture.

When designing the enterprise campus functional area, as diagramed in Figure 2-3, in the enterprise architecture, four primary areas need to be addressed:

- **Building access**—Connects end-user devices to the network
- **Building distribution**—Aggregates building access switches and performs Layer 3 switching (that is, routing) functions
- **Campus core**—Provides high-speed, redundant connectivity between buildings
- **Server farm and data center**—Consolidates application servers, e-mail servers, domain name servers, file servers, and network management applications

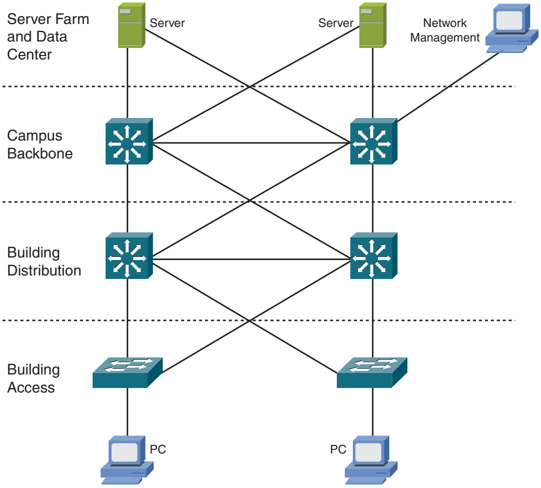

FIGURE 2-3 Enterprise campus.

The enterprise edge connects the enterprise campus with the WAN and Internet functional area. The four modules comprising the enterprise edge are as follows:

- **E-commerce**—Contains the servers used to provide an e-commerce presence for an organization, including the following:

  Web servers

  Application servers

  Database servers

  Security servers

- **Internet connectivity**—Provides Internet-related services, including the following:

  E-mail servers

  Domain Name System (DNS) servers

  Public web servers

  Security servers

  Edge routers

- **WAN and MAN site-to-site VPN (virtual private network)**—Interconnects a main office with remote offices over various transport technologies, such as the following:

  Frame Relay

  ATM

  PPP

  SONET

- **Remote access and VPN**—Provides secure access for remote workers (for example, telecommuters) or remote offices and includes components such as the following:

  Dial-in access concentrators

  VPN concentrators

  Cisco Adaptive Security Appliances (ASA)

  Firewalls

  Intrusion detection system (IDS) appliances

The WAN and Internet modules are sometimes referred to as *service provider modules*. These modules are the areas of the Enterprise Composite Network module not explicitly designed because the service provider modules are designed, owned, and operated by a service provider. However, the enterprise network designer can specify the type of connection to use in connecting to the service provider(s). Specifically, the service provider modules include the following types of connectivity:

- Frame Relay
- ATM
- Point-to-point leased line
- SONET and Synchronous Digital Hierarchy (SDH)
- Cable modem
- Digital subscriber line (DSL)
- Wireless bridging

Enterprise locations are supported via the following previously described modules:

- Enterprise branch
- Enterprise data center
- Enterprise teleworker

Identifying Infrastructure Services

Layered on top of an enterprise's network infrastructure are infrastructure services, which enable business applications. Examples of these infrastructure services include the following.

Security

The security service helps protect a network from both internal and external attacks. These threats might vary depending on the attack target (for example, the campus core or the e-commerce module). Therefore, security threats should be evaluated on a module-by-module basis.

Security services in enterprise edge can mitigate many attacks originating outside the enterprise network. However, some attacks might get through, and some attacks might originate internally. Therefore, critical devices in the enterprise campus need to be independently protected.

Examples of attacks that originate from outside the network include the following:

- IP spoofing
- Password attacks
- Denial-of-service (DoS) attacks
- Application layer attacks
- High-availability attacks

Today's enterprise networks often carry mission-critical traffic. Therefore, one of your design goals should be to include a degree of redundancy in a design, such that traffic can continue to flow through the enterprise network even if there is a link or component failure. However, adding redundancy (for example, redundant WAN links) not only adds to the complexity of the network, but it can also dramatically increase the cost to implement the design. With these factors in mind, consider which specific areas of the network would benefit most from a redundant design.

Approaches to providing redundancy include the following:

- **Adding redundant devices**—You could add redundant switches/routers to your design, as demonstrated in Figure 2-4, so that traffic continues to flow even if a router or switch fails.

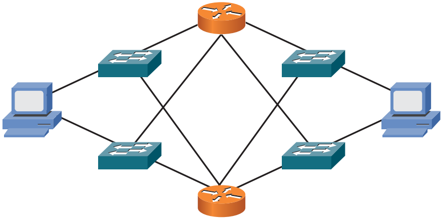

FIGURE 2-4 Redundant devices.

- **Adding redundant physical connections to end stations**—In a server farm, for example, you could have more than one network interface card (NIC) for each server. Each NIC could be connected to a different switch. Therefore, the server maintains network connectivity in the event of a single switch failure.
- **Advertising multiple routes to reach a destination network**—When you include physical redundant paths in your design, those routes should be advertised by a routing protocol with fast convergence (for example, Open Shortest Path First Protocol [OSPF] or Enhanced Interior Gateway Routing Protocol [EIGRP]).
- **Adding redundant links for load balancing and to accommodate for a link failure**—You can add more than one link between switches/routers, as depicted in Figure 2-5. These redundant links can not only improve network availability, but also provide load balancing for increased throughput.

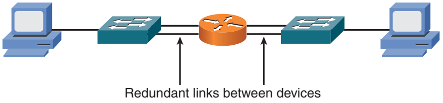

FIGURE 2-5 Redundant links.

Voice

Modern enterprise network designs need to support the transmission of voice traffic. This voice traffic can come from both analog phones (much like the phones typically found in homes) and IP phones, which are Ethernet devices that transmit voice IP packets. Because the analog phones cannot generate IP packets, they connect to analog gateways (such as Cisco routers), which convert the analog waveforms into IP packets.

The term ***Voice over IP***, or VoIP, is used to describe the transmission of voice over a network using voice-enabled routers. However, the term ***IP telephony*** refers to the use of IP phones and a call-processing server (for example, Cisco Unified CallManager).

Figure 2-6 shows the basic components of an IP telephony network.

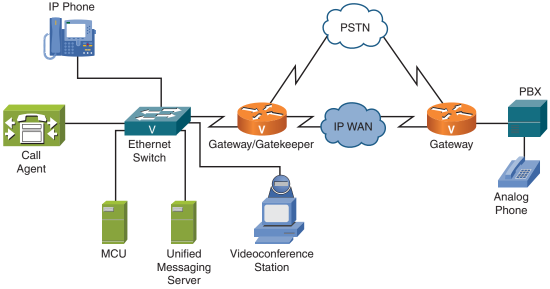

FIGURE 2-6 IP telephony network.

- **IP phone**—Provides IP voice to the desktop.
- **Gatekeeper**—Provides call admission control (CAC), bandwidth control and management, and address translation.
- **Gateway**—Provides translation between VoIP and non-VoIP networks, such as the PSTN. A gateway also provides physical access for local analog and digital voice devices, such as telephones, fax machines, key sets, and PBXs.
- **Multipoint control unit (MCU)**—Mixes audio/video streams, thus allowing participants in multiple locations to attend the same conference.
- **Call agent**—Provides call control for IP phones, CAC, bandwidth control and management, and address translation.

- **Application server**—Provides services such as voice mail, unified messaging, and Cisco CallManager Attendant Console.
- **Videoconference station**—Provides access for end-user participation in videoconferencing. The videoconference station contains a video capture device for video input and a microphone for audio input. The user can view video streams and hear the audio that originates at a remote user station. Cisco targets its VT Advantage product at desktop videoconferencing applications.

Other components, such as software voice applications, interactive voice response (IVR) systems, and softphones, provide additional services to meet the needs of enterprise sites.

Wireless

Not all devices in an enterprise network are necessarily wired into the network. Today, wireless connectivity is growing in popularity, allowing users to roam throughout the enterprise with their wireless device, such as a laptop.

However, because wireless networks send data through radio waves, as opposed to using physical cabling, security becomes a concern. Improper wireless designs might have the radio waves extended out of the building, into neighboring buildings or a parking lot. This type of radio frequency coverage provides an opportunity for attackers to infiltrate the enterprise network.

These Quick Reference Sheets address wireless design considerations in much more detail in a different section. However, for now, understand that wireless LANs are made up of four primary components:

- **End devices**—For example, laptops and PCs that have a wireless network adapter
- **Wireless access points**—Devices that act much like a shared hub for wireless clients and serve as an interconnection between the wireless and wired networks
- **Existing routed and switched wired network**—The enterprise network to which wireless access points connect
- **Wireless LAN controller**—A device that adds management and support capabilities to a wireless LAN, in addition to services (for example, roaming)

Application Networking

Application Networking Services (ANS) can use caching and compression technologies to make LAN-like responsiveness available to application users at remote offices. For example, when a web page is downloaded to a remote office, the images that make up the web page can be locally cached. Then, if a subsequent request is made for that web page, the initially downloaded graphics can be retrieved from the local cache, providing better response time and less demand on the WAN bandwidth. Also, security services validate application requests and provide confidentiality through encryption.

Primary components of a Cisco ANS network include the following:

- **Cisco Wide Area Application Engine (WAE)**—An appliance that provides LAN-like responsiveness to enterprise applications and data
- **Cisco Wide Area Application Services (WAAS)**—Software that provides high-performance access to centralized applications, servers, and storage resources
- **Cisco 2600/3600/3700 Series Content Engine Module**—A module installed in certain Cisco router platforms that contributes to WAN bandwidth optimization

Specifying Network Management Protocols and Features

When designing a network, remember to include network management protocols and features to allow network administrators to monitor their network devices, network connections, and network services. A network management solution can contain the following elements:

- **Network Management System (NMS)**—An NMS is a server that runs some sort of network management software, such as CiscoWorks.
- **Network Management Protocols**—Commonly used protocols that support network management functionality include the following:

  Simple Network Management Protocol (SNMP)—SNMP acts as the protocol used to transfer network management information between a managed device and a network management server. SNMP uses an SNMP agent that stores statistical information about a managed device inside of a Management Information Base (MIB). The three most popular implementations of SNMP are SNMPv1, SNMPv2c, and SNMPv3. The latest incarnation of SNMP (that is, SNMPv3) adds additional security levels.

  Management Information Base (MIB)—A MIB defines specific types of information about a device that an SNMP server can retrieve using a network management protocol, such as SNMP.

  Remote Monitoring (RMON)—RMON extends the information available in a MIB. Specifically, RMON collects and stores information locally on a device, and this information can be retrieved by an NMS to, for example, provide trend analysis.

  Many network devices support two levels of RMON, named RMON1 and RMON2. RMON1 only provides information about the physical and data link layers, whereas RMON2 can collect upper-layer information, as shown in Figure 2-7.

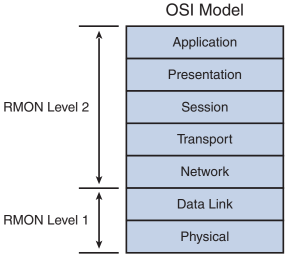

FIGURE 2-7 RMON levels.

Managed network elements include the following:

- **RMON**—RMON extends the information available in a MIB. Specifically, RMON collects and stores information locally on a device, and this information can be retrieved by an NMS.
- **Managed device**—A managed device is an endpoint (such as a server) that can be monitored, and perhaps controlled, by an NMS.
- **Management agent**—A management agent is a piece of software that runs on a managed device. Management agents include both SNMP agents and RMON agents.
- **Management information**—Data stored in MIBs are commonly referred to as management information.

Other applications that can assist in network management include the following:

- **NetFlow**—The Cisco NetFlow technology offers another approach to monitoring network statistics. NetFlow can store information about network flows, which are unidirectional communications paths between two devices. This stored information can then be exported to a network management collector, such as a NetFlow Collection Engine. Because of the way NetFlow analyzes specific flows, its information gathering places minimal overhead on a router's processor. Also, the data collected by NetFlow provides more detailed information than the data collected by RMON.

  NetFlow data can be used by various applications, such as

  - Billing applications based on network usage
  - Applications used for network planning
  - Security monitoring applications
  - Applications that need to know the network's quality of service (for example, amount of delay and percentage of dropped packets)
- **Cisco Discovery Protocol (CDP)**—Another protocol that can provide visibility into a network's topology is CDP. CDP functions at Layer 2 of the OSI model and can dynamically discover adjacent Cisco devices. For example, a Cisco router could discover information about Cisco Catalyst switches connected to that router. Because CDP is a Layer 2 technology, adjacent devices do not need to have a Layer 3 IP address to be discovered.

- **Syslog**—Network managers can also benefit from the System Message and Error Reporting Service, commonly known as syslog. Cisco's network devices can generate syslog messages to log various events to a syslog server. Each of these syslog messages contains a severity level and a facility.

  The severity level provides a measure of how serious an event is considered to be. For example, the debugging severity level (that is, Level 7) causes syslog messages to be sent for all routine operations, which can generate a large amount of output. However, a severity level of emergency (that is, Level 0) only generates a syslog message for the most serious events.

  A syslog facility identifies the service associated with the event. Examples of syslog facilities include IP, OSPF, and IPsec.

Exploring Basic Campus and Data Center Network Design

The multilayer design strategy uses a modular approach, which adds scalability to a design. This section examines how the multilayer design approach can be applied to both the enterprise campus and the enterprise data center.

Understanding Campus Design Considerations

As illustrated in Figure 3-1, an enterprise campus might be composed of multiple buildings that share centrally located campus resources.

Enterprise campus design considerations fall under three categories:

- **Network application considerations**—A network's applications might include the following:
  - **Peer-to-peer applications** (for example, file sharing, instant messaging, IP telephony, videoconferencing)
  - **Client/local server applications** (for example, applications on servers located close to clients or servers on the same LAN)

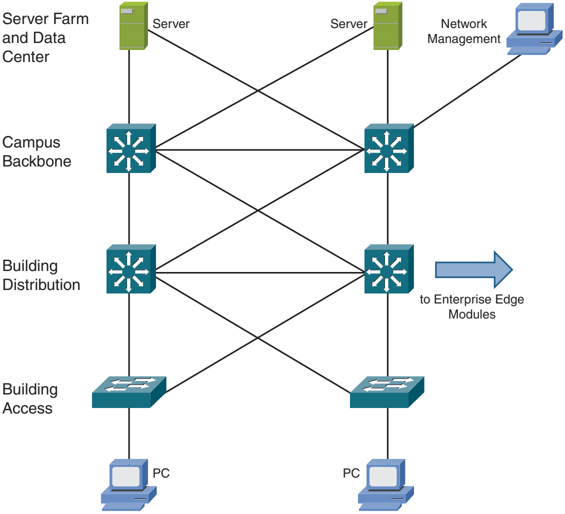

FIGURE 3-1 Enterprise campus.

  - **Client/server farm applications** (for example, e-mail, file sharing, and database applications)
  - **Client/enterprise edge server applications** (for example, Internet accessible web and e-commerce applications)

- **Environmental considerations**—Network environmental considerations vary with the scope of the network. Three scopes are as follows:
  - **Intrabuilding**—An intrabuilding network provides connectivity within a building. The network contains both building access and building distribution layers. Typical transmission media includes twisted pair, fiber optics, and wireless technology.
  - **Interbuilding**—An interbuilding network provides connectivity between buildings that are within two kilometers of each other. Interbuilding networks contain the building distribution and campus core layers. Fiber optic cabling is typically used as the transmission media.
  - **Remote Buildings**—Buildings separated by more than two kilometers might be interconnected by company-owned fiber, a company-owned WAN, or by service provider offerings (for example, metropolitan-area network [MAN] offerings).

Common transmission media choices include the following:

- **Twisted pair**

  1000-m distance limit

  10-Gbps speed limit

  Low cost

- **Multimode fiber** (as illustrated in Figure 3-2)
  - 2-km distance limit (Fast Ethernet) or 550-m distance limit (Gigabit Ethernet)
  - 10-Gbps speed limit
  - Moderate cost

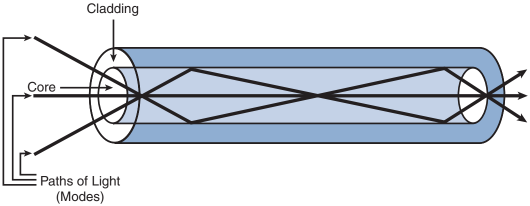

FIGURE 3-2 Multimode fiber.

NOTE

The core diameter in a multimode fiber is large enough to permit multiple paths (that is, modes) for light to travel. This might cause different photons (that is, light particles) to take different amounts of time to travel through the fiber. As distance increases, this leads to multimode delay distortion. Therefore, multimode fiber has a distance limitation of approximately 2 km.

- **Single-mode fiber** (as illustrated in Figure 3-3)
  - 80-km distance limit (Fast Ethernet or 10 Gigabit Ethernet)
  - Speed limit of 10-Gbps or greater
  - High cost

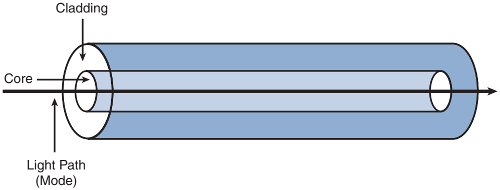

FIGURE 3-3 Single-mode fiber.

NOTE

The core diameter in a single-mode fiber is only large enough to permit one path for light to travel. This approach eliminates multimode delay distortion, thus increasing the maximum distance supported.

- **Wireless**

  500-m distance limit (at a rate of 1 Mbps)

  Speed limit of 54 Mbps

  Moderate cost

Infrastructure device considerations include the following:

- When selecting infrastructure devices, Layer 2 switches are commonly used for access layer devices, whereas multilayer switches are typically found in the distribution and core layers.
- Selection criteria for switches include the need for QoS, the number of network segments to be supported, required network convergence times, and the cost of the switch.

Understanding the Campus Infrastructure Module

When designing the enterprise campus, different areas of the campus (that is, building access, building distribution, campus core, and server farm) require different device characteristics (that is, Layer 2 versus multilayer technology, scalability, availability, performance, and per-port cost).

- **Building access best practices**

  Limit the scope of most VLANs to a wiring closet. A VLAN is a single broadcast domain.

  If you use the Spanning Tree Protocol (STP), select Rapid Per VLAN Spanning Tree Plus (RPVST+) for improved convergence.

  When using trunks to support the transmission of traffic from multiple VLANs across a single physical link, set both ends of the trunk to *desirable*, which causes the switches at each end of the link to send Dynamic Trunk Protocol (DTP) frames in an attempt to negotiate a trunk. Also, set the DTP mode to *negotiate*, to support DTP protocol negotiation.

Remove (that is, “prune”) unneeded VLANs from trunks.

Set the VLAN Trunking Protocol (VTP) mode to *transparent* because a hierarchical design has little need for a VLAN to span multiple switches.

When using an EtherChannel, set the Port Aggregation Protocol (PAgP) mode to *desirable* to cause both sides of the connection to send PAgP frames, in an attempt to create an EtherChannel.

Consider the potential benefits of implementing routing at the access layer to achieve, for example, faster convergence times.

- **Building-distribution considerations**

Switches selected for the building distribution layer require wire-speed performance on all their ports. The need for such high performance stems from the roles of a building distribution layer switch: acting as an aggregation point for access layer switches and supporting high-speed connectivity to campus core layer switches.

The key roles of a building distribution layer switch demand redundant connections to the campus core layer. You should design redundancy such that a distribution layer switch could perform equal-cost load balancing to the campus core layer. However, if a link were to fail, the remaining link(s) should have enough capacity to carry the increased traffic load. Redundancy technologies such as Stateful Switchover (SSO) and Nonstop Forwarding (NSF) offer failover times in the range of one to three seconds. Also, some platforms support the In Service Software Upgrade (ISSU) feature, which allows you to upgrade a switch’s Cisco IOS image without taking the switch out of service.

Building distribution layer switches should support network services such as high availability, quality of service (QoS), and policy enforcement.

- **Campus core considerations**

Evaluate whether a campus core layer is needed. Campus core layer switches interconnect building distribution layer switches, and Cisco recommends that you deploy a campus core layer when interconnecting three or more buildings or when interconnecting four or more pairs of building distribution layer switches.

Determine the number of high-speed ports required to aggregate the building distribution layer.

For high-availability purposes, the campus core should always include at least two switches, each of which can provide redundancy to the other.

Decide how the campus core layer connects to the enterprise edge and how WAN connectivity is provided. Some designs use edge distribution switches in the core to provide enterprise edge and WAN connectivity. For larger networks that include a data center, enterprise edge and WAN connectivity might be provided through the data center module.

- **Server farm considerations**—Determine server placement in the network. For networks with moderate server requirements, common types of servers can be grouped together in a separate server farm module connected to the campus core using multilayer switches. Access control lists (ACL) in these multilayer switches offer limited access to these servers.

  All server-to-server traffic should be kept within the server farm module and not be propagated to the campus core.

  For large network designs, consider placing the servers in a separate data center. This data center could potentially reside in a remote location.

  Consider using network interface cards (NIC) in servers that provide at least two ports. One NIC port could be active, with the other port in standby mode. Alternatively, some NICs support EtherChannel, which could increase the effective throughput between a server and the switch to which it connects.

  For security, place servers with similar access policies in the same VLANs, and then limit interconnections between servers in different policy domains using ACLs on the server farm's multilayer switches.

  Understand the traffic patterns and bandwidth demands of applications deployed on the servers. Some applications (for example, backup applications or real-time interactive applications) place a high bandwidth demand on the network. By understanding such application characteristics, you can better size the server farm uplinks to prevent oversubscription.

Understanding Enterprise Data Center Considerations

An enterprise data center's architecture uses a hierarchical design, much like the campus infrastructure. However, there are subtle differences in these models. Large networks that contain many servers traditionally consolidated server resources in a data center. However, data center resources tended not to be effectively used because the supported applications required a variety of operating systems, platforms, and storage solutions. These diverse needs resulted in multiple ***application silos***, which can be thought of as separate application "islands."

Today, the former server-centric data center model is migrating to a service-centric model. The main steps in this migration are as follows:

1. Use virtual machine software, such as VMware, to remove the requirement that applications running on different operating systems must be located on different servers.
2. Remove network storage from the individual servers, and consolidate the storage in shared storage pools.
3. Consolidate I/O resources, such that servers have on-demand access to I/O resources, to reach other resources (for example, other servers or storage resources).

The Cisco enterprise data center architecture consists of two layers:

- **Networked Infrastructure Layer**—The Networked Infrastructure Layer contains computing and storage resources, which are connected in such a way to meet bandwidth, latency, and protocol requirements for user-to-server, server-to-server, and server-to-storage connectivity design requirements.
- **Interactive Services Layer**—The Interactive Services Layer supports such services as Application Networking Services (ANS) (for example, application acceleration) and infrastructure enhancing services (for example, intrusion prevention).

Data centers can leverage the Cisco enterprise data center architecture to host a wide range of legacy and emerging technologies, including *N*-tier applications, web applications, blade servers, clustering, service-oriented architecture (SOA), and mainframe computing.

An enterprise data center infrastructure design requires sufficient port density and L2/L3 connectivity at the access layer. The design must also support security services (for example, ACLs, firewalls, and intrusion detection systems [IDS]) and server farm services (for example, content switching and caching). Consider the following design best practices for an enterprise data center's access, aggregation, and core layers:

- **Data center access layer design best practices**

  Provide for both Layer 2 and Layer 3 connectivity.

  Ensure sufficient port density to meet server farm requirements.

  Support both single-attached and dual-attached servers.

  Use RPVST+ as the STP approach for loop-free Layer 2 topologies.

  Offer compatibility with a variety of uplink options.
- **Data center aggregation layer design best practices**

  Use the data center aggregation layer to aggregate traffic from the data center access layer.

  Provide for advanced application and security options.

  Maintain state information for connections, so that hardware failover can occur more rapidly.

  Offer Layer 4 through 7 services, such as firewalling, server load balancing, Secure Sockets Layer (SSL) offloading, and IDS.

  Provision processor resources to accommodate a large STP processing load.
- **Data center core layer design best practices**

  Evaluate the need for a data center core layer by determining whether the campus core switches have sufficient 10-Gigabit Ethernet ports to support both the campus distribution and data center aggregation modules.

  If you decide to use a data center core, use the separate cores (that is, the campus core and the data center core) to create separate administrative domains and policies (for example, QoS policies and ACLs).

If you decide that a data center core is not currently necessary, anticipate how future growth might necessitate the addition of a data center core. Determine whether it would be worthwhile to initially install a data center core, instead of adding one in the future.

Designers commonly use modular chassis (for example, Cisco Catalyst 6500 or 4500 series switches) in an enterprise access layer. Although this design approach does offer high performance and scalability, challenges can emerge in a data center environment. Server density has increased thanks to 1RU (one rack unit) and blade servers, resulting in the following issues:

- **Cabling**—Each server typically contains three to four connections, making cable management between high-density servers and modular switch more difficult.
- **Power**—Increased server and switch port density requires additional power to feed a cabinet of equipment.
- **Heat**—Additional cabling under a raised floor and within a cabinet can restrict the airflow required to cool equipment located in cabinets. Also, due to higher-density components, additional cooling is required to dissipate the heat generated by switches and servers.

One approach to address these concerns is just to not deploy high-density designs. Another approach is to use rack-based switching, with 1RU top-of-rack switches, which allows the cables between the servers and switches to be confined within a cabinet. If you prefer to use modular switches, an option is to locate modular switches (for example, Cisco Catalyst 6500 series switches) much like “bookends” on each end of a row of cabinets. This approach reduces administration overhead because you have fewer switches to manage compared to using multiple 1RU switches.

Remote Connectivity Design

Remote office locations, such as branch offices or the homes of teleworkers, connect to the enterprise campus via the enterprise edge and enterprise WAN. When selecting an appropriate WAN technology to extend to these remote locations, design considerations include ownership (that is, private, leased, or shared ownership) of the link, reliability of the link, and a backup link if the primary link were to fail. This section explores various WAN technologies and provides guidance for designing the enterprise WAN and the enterprise branch.

Considering WAN Technology Options

In the Cisco Enterprise Architecture, the enterprise edge allows the enterprise campus to connect to remote offices using a variety of WAN, Internet access, and remote-access technologies (for example, secure virtual private network [VPN] access). A WAN spans a relatively broad geographical area and a wide variety of connectivity options exist. Therefore, designing a WAN can be a complex task. To begin a WAN design, first understand the following network characteristics:

- **Service level agreement (SLA)**—This document is an agreement between a customer and service provider that specifies acceptable levels of bandwidth, latency, and packet loss across a WAN.
- **Cost and usage**—Understanding how the WAN will be used can help determine a cost-effective technology to meet the design requirements.

The primary goals of WAN design include the following:

- The WAN must achieve the goals, meet the characteristics, and support the policies of the customer.
- The WAN must use a technology to meet present requirements, in addition to requirements for the near future.
- The expense of the WAN (one-time and recurring expenses) should not exceed customer-specified budgetary constraints.

Today's WAN designer can select from a plethora of technologies. Consider the characteristics of the following modern WAN technologies:

- **Time-division multiplexing (TDM)**—A TDM circuit is a dedicated point-to-point connection that is constantly connected. T1 and E1 circuits are examples of TDM circuits.
- **Integrated Services Digital Network (ISDN)**—ISDN uses digital phone connections to support the simultaneous transmission of voice, video, and data. ISDN is considered to be a circuit-switched technology because an ISDN call is set up much the same way a telephone call is set up.
- **Frame Relay**—Frame Relay is considered to be a packet-switched technology, which uses the concept of permanent virtual circuits (PVC) and switched virtual circuits (SVC) to potentially create multiple logical connections using a single physical connection.

- **Multiprotocol Label Switching (MPLS)**—MPLS is considered to be a label-switching technology, where packets are forwarded based on a 32-bit label, as opposed to an IP address. Service providers often use MPLS to engineer traffic through the network based on an initial route lookup, quality of service (QoS) classification, and application bandwidth requirements.
- **Metro Ethernet**—Metro Ethernet uses Ethernet technology to provide high-speed, yet cost-effective, links for some metropolitan-area networks (MAN) and WANs.
- **Digital subscriber line (DSL)**—DSL provides high-bandwidth links over existing phone lines. A variety of DSL implementations exist. The most popular type of DSL found in homes is asynchronous DSL (ADSL), which allows home users to simultaneously use their phone line for both high-speed data connectivity and traditional analog telephone access.
- **Cable**—Cable technology leverages existing coaxial cable, used for delivery of television signals, to simultaneously deliver high-speed data access to the WAN, and optionally to the public switched telephone network (PSTN), as illustrated in Figure 4-1.
- **Wireless**—Wireless technologies use radio waves to connect devices, such as cell phones and computers. As an example of a wireless application, wireless bridges can connect two buildings that are less than 1 mile apart and have a line-of-site path between them, as shown in Figure 4-2.

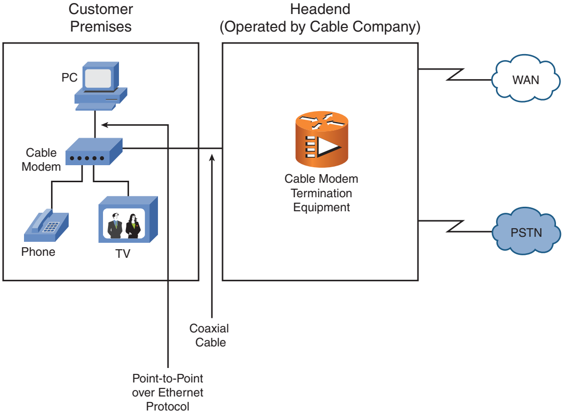
Customer Premises
Headend (Operated by Cable Company)
PC
Cable Modem
Phone
TV
Cable Modem Termination Equipment
WAN
PSTN
Coaxial Cable
Point-to-Point over Ethernet Protocol

FIGURE 4-1 Data and voice over cable.

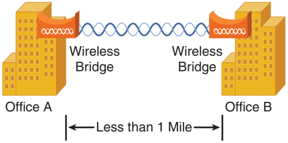

FIGURE 4-2 Wireless bridges.

- **Synchronous Optical Networking (SONET) and Synchronous Digital Hierarchy (SDH)**—SONET and SDH both use TDM technology to provide services over an optical network, as demonstrated in Figure 4-3. Thanks to the optical transport used by these technologies, relatively high-bandwidth solutions are available. Some of the popular SONET/SDH access speeds include 155 Mbps and 622 Mbps, with a maximum bit rate of 10 Gbps.

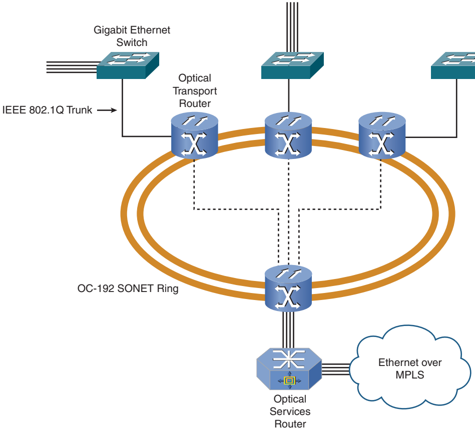

FIGURE 4-3 SONET network example.

- **Dense wavelength division multiplexing (DWDM)**—DWDM increases the bandwidth capacity of an optical cable by sending multiple traffic flows over the same fiber, with each flow using a different wavelength.

When selecting a WAN technology, be aware that provisioning a circuit can require 60 days or more. Therefore, sufficient lead time must be built in to the schedule. Also, Metro Ethernet coverage is limited compared to other technologies. Be sure to negotiate an SLA that meets your design requirements, and be conscious of the contract period. Typically, WAN contract periods are in the range of one to five years.

Enterprise edge design uses the PPDIOO approach discussed earlier. Specifically, you should do the following:

- **Determine network requirements**—Network requirements are influenced by the volume and patterns of traffic generated by networked applications.
- **Evaluate existing network technology**—When documenting current network technology, include not only the types of equipment connected to the network (for example, hosts and servers), but also the location of the equipment.
- **Design the network topology**—The network topology design should preserve the customer's existing investment by leveraging existing technology, with the understanding that upgrades might be required. Also, the proposed topology should accommodate not only existing traffic patters, but projected traffic patterns.

When you are designing networks to traverse the WAN, a primary design consideration is making the most efficient use of the relatively limited WAN bandwidth. Fortunately, Cisco provides a variety of QoS mechanisms that can help:

- **Compression**—By compressing the header/payload of a packet, that packet requires less bandwidth for transmission across a WAN. Therefore, compressing traffic is much like adding WAN bandwidth. However, there is a drawback. Compression requires processing resources from the router. Therefore, although more information can be sent across the same link speed, the router's processor bears an additional burden.
- **Link aggregation**—Cisco routers support the bonding together of physical links into a virtual link. For example, if you have two serial interfaces, each running at a speed of 256 kbps, you can use a technology such as Multilink PPP (MLP) to create a virtual multilink interface running at a speed of 512 kbps.
- **Window size**—TCP traffic uses the concept of a "sliding window." A ***window*** is the number of segments that a TCP sender can transmit before receiving an acknowledgment from the receiver. Network delay can be reduced by increasing the window size (that is, sending more TCP segments before expecting an acknowledgment). However, on unreliable links that suffer from high error rates, the number of retransmissions could increase dramatically.
- **Queuing**—When a router is receiving traffic (for example, from a LAN interface) faster than it can transmit that traffic (for example, out of a WAN interface), the router delays the excess traffic in a buffer called a *queue*. To prevent bandwidth-intense applications from consuming too much of the limited WAN bandwidth, various queuing technologies can place different types of traffic into different queues, based on the traffic priority. Then, different amounts of bandwidth can be given to the different queues, allowing more important applications to receive the bandwidth they need, as illustrated in Figure 4-4.

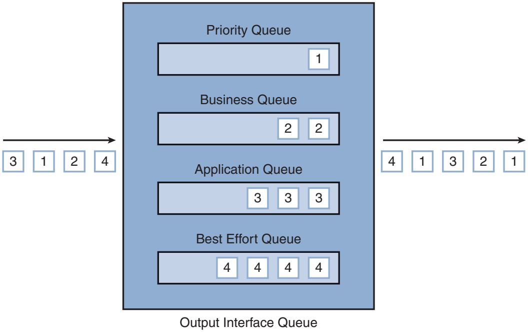

FIGURE 4-4 Queuing.

- **Traffic conditioning**—To prevent some types of traffic (for example, music downloads from the Internet) from consuming too much WAN bandwidth, a traffic conditioner called *policing* can be used to set a "speed limit" on those specific traffic types, and drop any traffic exceeding that limit. Similarly, to prevent a WAN link from becoming oversubscribed (for example, oversubscribing a remote office's 128 kbps link when receiving traffic from the headquarters that is transmitting at a speed of 768 kbps), another traffic conditioner, called *shaping*, can be used to prevent traffic from exceeding a specified bandwidth. With shaping, compared to policing, excessive traffic is delayed and transmitted when bandwidth becomes available, instead of being dropped. Unlike shaping, policing mechanisms can also re-mark traffic, giving lower-priority QoS markings to traffic exceeding a bandwidth limit. Policing mechanisms include Committed Access Rate (CAR) and class-based policing; examples of shaping mechanisms include Frame Relay Traffic Shaping (FRTS) and class-based shaping.

Performing the Enterprise WAN Design

When considering design elements for the enterprise WAN, be aware of possible WAN design choices. Consider the following WAN design categories:

- **Traditional WAN design**—Most traditional WAN designs could be categorized under one of three options:
  - **Leased Lines**—A leased line is a point-to-point connection that provides a reserved amount of bandwidth for a customer. An example of a leased line WAN is a T1 link between two sites using PPP.
  - **Circuit switched**—A circuit-switched design uses circuits that are brought up on an as-needed basis and then torn down. ISDN falls under the category of a circuit-switched network.
  - **Packet/cell switched**—A packet-switched (for example, Frame Relay) or cell-switched (for example, ATM) network can use permanent virtual circuits (PVC) and switched virtual circuits (SVC) to connect multiple sites. These networks can leverage a variety of topologies, such as full mesh or hub and spoke.
- **Remote-access network design**—Remote-access networks allow remote employees (for example, telecommuters or traveling salespeople) to access the corporate network. Besides data, a remote-access network might also need to support voice calls. Typical technologies offering remote access include dial-up (using a traditional modem or an ISDN connection), DSL, cable, and wireless.
- **Virtual private network (VPN) design**—A VPN can provide security to a remote connection by creating a virtual tunnel through which all traffic is sent, even though the connection might be traversing an untrusted network. One type of VPN is a site-to-site VPN, which might connect a remote office with the headquarters office over the publicly accessible Internet. In such a design, each site typically has hardware to terminate each end of the VPN tunnel. Another option is to have VPN client software on a user's

PC, allowing them to connect to the headquarters' VPN equipment and set up a secure VPN connection, by providing credentials, such as a username and password. Figure 4-5 shows sample topologies of these VPN types.

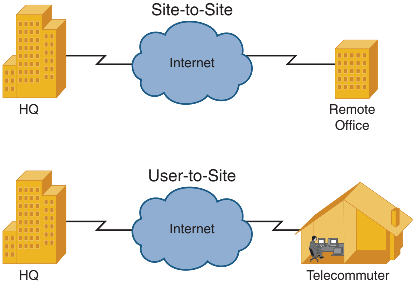

FIGURE 4-5 VPN types.

- **WAN backup design**—WAN links tend to be less reliable than LAN connections. Therefore, a good WAN design provides for fault tolerance in the form of a WAN backup. Consider the following options:
  - **Dial Backup Routing**—Dial backup routing uses dial-up technologies, such as modem and ISDN technologies, to bring up a backup link if the primary link fails.
  - **Redundant WAN link**—Instead of having a backup link that comes up only when needed, a secondary WAN link can be a permanent link. One option for using this permanent secondary link is to use a floating static route, or a routing protocol, to send traffic over that secondary link only when the primary link is unavailable.

    Another option is to leverage the extra bandwidth provided by the secondary link and perform load balancing across both links, when both links are available. Then, if one link goes down, the other link can carry all the traffic.
  - **Shadow PVC**—A shadow PVC is made available by your service provider, typically at an extra charge. This shadow PVC becomes active only if your primary PVC becomes unavailable.
  - **IPsec tunnel**—Because most networks already have Internet access, in addition to WAN links that connect office locations, the Internet can act as a backup WAN link. However, because the Internet is a public network, security becomes a concern. IPsec tunneling can alleviate that concern by protecting sensitive corporate traffic inside a secure VPN tunnel.

At this point, you understand remote connectivity requirements, and you have been exposed to various WAN architectures. You are now ready to select an appropriate WAN architecture for your design. Following are design considerations for the enterprise WAN architecture:

- **Network growth**—Your design should not only accommodate existing bandwidth requirements but should also allow the customer to grow their network along with their business.

- **Availability**—A common availability design goal is for the network to be up 99.999 percent of the time. This metric is commonly referred to as “the five nines of availability.” The five nines of availability translates into only five minutes of downtime per year. A key design factor that influences availability is redundancy. Redundancy should be built in to the design, such that no major component (for example, a router or a WAN link) represents a single point of failure. In addition to equipment and link backups, also consider a power backup. Do you have sufficient UPS (uninterruptible power supply) and generator equipment in your design to sustain key network components if an extended power outage occurs?
- **Recurring expenses**—Companies pay regular subscription fees to their service provider for their WAN service. This type of recurring expense (in addition to equipment leases) can influence your decision in selecting a WAN technology. For example, Frame Relay and ATM WANs usually cost more than using an IPsec VPN over the public Internet. However, performance trade-offs might come with cost savings. For example, if you select an IPsec VPN over the Internet, as opposed to a Frame Relay network, your network might suffer from QoS issues.
- **Network complexity**—Your customer might have their own IT staff for maintaining their WAN connection. Therefore, you need to understand the skill set of the IT staff and their ability to work with complex network designs, because different WAN technologies require differing levels of technical expertise.
- **Multimedia support**—Determine whether the customer is going to use the WAN link to transmit voice/video. If these types of multimedia applications are going to be transmitted over the network, your design must include QoS mechanisms to ensure appropriate treatment for these latency-sensitive traffic types.
- **Migration expense**—Migrating one MAN/WAN technology to another MAN/WAN technology often necessitates a significant initial investment (for example, to cover the expenses of the new equipment, installation labor, and employee training). However, this initial investment might very well be recovered from future cost savings. Therefore, your design should include a return on investment (ROI) calculation for your proposed expenditures.
- **Network segmentation**—Instead of having multiple autonomous networks, having a single network that is logically segmented can reduce the expenses (for example, equipment and maintenance expenses) of supporting multiple physical networks. The single physical network can be logically segmented into multiple network segments, thus providing security between the different segments.

After identifying the remote connectivity requirements and architecture for a design, the next step is to select the specific WAN components to be used in the design. This step involves the selection of hardware and software components:

- **Hardware selection**—When selecting hardware for your design, examine the product documentation looking for such product specifics as port density, throughput, enhanced capabilities, and redundancy.
- **Software selection**—Cisco IOS Software supports a wide variety of features, services, and platforms. For example, consider the following "trains" of IOS Software:
  - **T Train**—T train Cisco IOS Software supports IP services such as IP communications, security, and mobility. Such services are well suited for the enterprise core and service provider edge.
  - **S Train**—S train Cisco IOS Software is appropriate for high-end enterprise core networks. The S train offers various IP services and infrastructure features such as MPLS, video, and multicast.
  - **XR Train**—XR train Cisco IOS Software is appropriate for large-scale networks. The XR train offers high availability features such as in-service software upgrades.

When selecting an appropriate Cisco IOS version, you might need to select from various IOS ***feature sets***. As a reference, Table 4-1 provides a sampling of features included in various feature sets.

TABLE 4-1 IOS Feature Sets

IOS Feature Sets	Data Connectivity	VoIP and VoFR	ATM, VoATM, and MPLS	AppleTalk IPX, IBM Protocols	Firewall, IDS, and VPN
IP Base	X				
IP Voice	X	X			
Advanced Security	X				X
Enterprise Base	X			X	
SP Services	X	X	X		
Advanced IP Services	X	X	X		X
Enterprise Services	X	X	X	X	
Advanced Enterprise Services	X	X	X	X	X

Performing the Enterprise Branch Design

The Cisco enterprise branch architecture seeks to extend enterprise services (for example, voice, video, and security services) to smaller branch locations. An employee's residence can also serve as a branch office.

Following are devices commonly found in enterprise branch architectures:

- WAN routers
- LAN switches
- Security appliances
- Wireless access points
- Call-processing servers for voice/video calls (for example, Cisco Unified CallManager)
- Endpoints (for example, IP phones and computers)

When designing the enterprise branch, consider the following issues:

- Total number of branch locations
- Total number of connected devices
- Anticipated growth
- Level of required security

Server farm requirements

- Location of network management system
- Impact of wireless networking (if used)
- Available budget

While a branch office is considered to be a "smaller" remote office, different degrees of smallness exist. Specifically, branch offices can be categorized as one of the following:

- **Small branch office**—A branch office is considered small if it has fewer than 50 users. The network supporting a small branch office is typically a single-tier design, as opposed to a hierarchical design. Therefore, Spanning Tree Protocol (STP) design is not an issue, although STP should be enabled to prevent the accidental creation of a Layer 2 switching loop. Design recommendations might include the integration of switch ports into an Integrated Services Router (ISR) or a multiservice router, using a Cisco EtherSwitch module.
- **Medium branch office**—A branch office is considered medium sized if it supports 50 to 100 users. This type of network can benefit from a two-tier design. Therefore, STP becomes a design issue. Because of the increased number of devices to be supported on the network, instead of integrating switch ports into a router, external stackable switches might be used.

- **Large branch office**—A branch office is considered large if it supports at least 100 users, but no more than 200 users. With this number of users, the network design can start to benefit from a three-layer hierarchical design. Redundant components (for example, redundant distribution layer switches and redundant WAN routers running Hot Standby Router Protocol [HSRP]) can improve the network's availability. Access layer switches tend to be higher-density stackable switches, whereas distribution layer switches might run enhanced Cisco IOS images to support, for example, multiple routing protocols and policy-based routing.

Other than the small, medium, and large sized branch offices, some networks support teleworkers, which are sometimes considered to be a "branch of one." Enterprise teleworkers, however, can be distinguished from typical telecommuters in that enterprise teleworkers enjoy access to networking services typically available to clients of a corporate network (for example, VoIP, videoconferencing, and real-time collaboration applications). These services are usually available to teleworkers over a secure VPN connection because the link between a teleworker's home and the corporate office is via the public Internet. Access to the Internet leverages widely available broadband services, such as DSL and cable. If the broadband link becomes unavailable, a traditional dial-up modem can be used as a backup link.

IP Addressing and Routing Protocols

Efficiently assigning IP addresses to your network is a critical design decision, impacting the scalability of the network and the routing protocol that can be used. This section reviews IP Version 4 addressing, introduces IP Version 6 addressing, and analyzes characteristics of various routing protocols.

IP Addressing

Before discussing design decisions surrounding IP addressing, first review the following characteristics of Internet Protocol Version 4 (IPv4) addressing:

- IPv4 addresses are 32 bits in length.
- IPv4 addresses are divided into various classes (for example, Class A networks accommodate more than 16 million unique IP addresses, Class B networks support more than 65 thousand IP addresses, and Class C networks permit 254 usable IP addresses). Originally, organizations applied for an entire network in one of these classes. Today, however, subnetting allows a service provider to give a customer just a portion of a network address space, in an attempt to conserve the depleting pool of IP addresses. Conversely, service providers can use supernetting (also known as classless interdomain routing [CIDR]) to aggregate the multiple network address spaces that they have. Aggregating multiple network address spaces into one reduces the amount of route entries a router must maintain.
- Devices, such as PCs, can be assigned a static IP address, by hard-coding the IP address in the device's configuration. Alternatively, devices can dynamically obtain an address from, for example, a DHCP server.
- Because names are easier to remember than IP addresses, most publicly accessible web resources are reachable by their name. However, routers must determine the IP address with which the name is associated to route traffic to that destination. Therefore, a Domain Name System (DNS) server can perform the translation between domain names and their corresponding IP addresses.
- Some IP addresses are routable through the public Internet, whereas other IP addresses are considered private and are intended for use within an organization. Because these private IP addresses might need to communicate outside the local network, Network Address Translation (NAT) can translate a private IP address into a public IP address. In fact, multiple private IP addresses can be represented with a single public IP address using NAT. This type of NAT is called Port Address Translation (PAT) because the various communication flows are identified by the port numbers they use to communicate with outside resources.

When beginning to design the IP addressing for a network, determine the following:

- The number of network locations that need IP addressing
- The number of devices requiring an IP address at each location
- Customer-specific IP addressing requirements (for example, static IP addressing versus dynamic IP addressing)
- The number of IP addresses that need to be contained in each subnet (for example, a 48 port switch in a wiring closet might belong to a subnet that supports 64 IP addresses)

Proper address planning can minimize the number of entries in a routing table through the use of aggregation. For example, suppose that Building 1 has a network address space of 10.1.1.0/24 (that is, 10.1.1.0 with a 24-bit subnet mask) and Building 2 has a network address space of 10.1.2.0/24. Instead of advertising both of those networks separately to the core layer, a distribution layer switch or router could aggregate those two addresses into a single route advertisement of 10.1.0.0/16 (that is, 10.1.0.0 with a 16-bit subnet mask). This approach to aggregating routes is called ***route summarization***.

Figure 5-1 illustrates how subnets within individual buildings can be summarized by distribution layer switches before the routes are advertised to a core switch. In the figure, even though there are a total of four building subnets, the core switch maintains only two entries in its routing table for those four networks.

NOTE

In the preceding example, the summarized route of 10.1.0.0/16 encompassed more networks than the two being discussed. Therefore, a more appropriate subnet mask might have been chosen for a real-world design. However, the actual calculation of variable-length subnet masks (VLSM) is beyond the scope of the DESGN course, and as a result, only classful subnet masks (that is, 8-bit, 16-bit, or 24-bit subnet masks) are used for the examples in these Quick Reference Sheets.

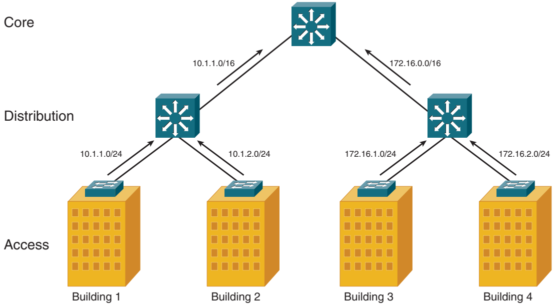

FIGURE 5-1 Route summarization.

A major challenge with IPv4 is the limited number of available addresses. A newer version of IP, specifically IPv6, fixes this concern. An IPv6 address is 128 bits long, compared to the 32-bit length of an IPv4 address.

To make such a large address more readable, an IPv6 address uses hexadecimal numbers, and the 128-bit address is divided into eight fields. Each field is separated by a colon, as opposed to the four fields in an IPv4 address, which are separated by a period.

As an example, consider the following IPv6 address:

4071:0000:130F:0000:0000:09C0:D76A:9801

Notice the use of hexadecimal numbers and the eight colon-separated fields.

To further reduce the complexity of the IPv6 address, leading 0s in a field are optional, and if one or more consecutive fields contain all 0s, those fields can be represented by a double colon (that is, ::). A double colon can be used only once in an address; otherwise, it would not be possible to know how many 0s are present between each pair of colons.

To illustrate these techniques, consider the IPv6 address presented in the previous example. There are three fields consisting of all 0s:

4071:**0000**:130F:**0000:0000**:09C0:D76A:9801

Because a double colon can be used only one time, you want to replace the two consecutive all 0s fields with the double colon:

4071:0000:130F**::**09C0:D76A:9801

Next, the remaining field that contains all 0s can be represented with a single 0, because leading 0s are optional:

4071:**0**:130F**::**09C0:D76A:9801

By the same reasoning, the leading 0 in the 09C0 field can be removed, leaving a resulting IPv6 address of

4071:**0**:130F**::**9C0:D76A:9801

Consider some of the benefits offered by IPv6:

- IPv6 dramatically increases the number of available addresses (that is, approximately 3.4 * 10^{38} addresses).
- Hosts can have multiple IPv6 addresses, allowing those hosts to multihome to multiple Internet service providers.
- Other benefits include enhancements relating to quality of service (QoS), security, mobility, and multicast technologies.

Unlike IPv4, IPv6 does not use broadcasts. Instead, IPv6 uses the following methods of sending traffic from a source to one or more destinations:

- **Unicast (one-to-one)**—Unicast support in IPv6 allows a single source to send traffic to a single destination, just as unicast functions in IPv4.
- **Anycast (one-to-nearest)**—A group of interfaces belonging to nodes with similar characteristics (for example, interfaces in replicated FTP servers) can be assigned an anycast address. When a

host wants to reach one of those nodes, the host can send traffic to the anycast address, and the node belonging to the anycast group that is closest to the sender will respond. For example, imagine a company has replicated FTP servers in countries throughout the world. A host in the United States can send a packet out to the anycast address (which all the FTP servers are associated with), and an FTP server in the United States will respond, rather than an FTP server in Japan, for example, because the United States FTP server is the closest server.

- **Multicast (one-to-many)**—Like IPv4, IPv6 supports multicast addressing, where multiple nodes can join a multicast group. The sender sends traffic to the multicast IP address, and all members of the multicast group receive the traffic.

Migrating an IPv4 network to an IPv6 network can take years because of the expenditures of upgrading equipment. Therefore, during the transition, IPv4-speaking devices and IPv6-speaking devices need to peacefully coexist on the same network. Consider three popular solutions for maintaining both IPv4 and IPv6 devices in the network:

- **Dual stack**—Some systems (including Cisco routers) can simultaneously run both IPv4 and IPv6, allowing communication to both IPv4 and IPv6 devices.
- **Tunneling**—To send an IPv6 packet across a network that only uses IPv4, the IPv6 packet can be encapsulated and tunneled through the IPv4 network.
- **Translation**—A device, such as a Cisco router, could sit between an IPv4 network and an IPv6 network and translate between the two addressing formats.

Enterprise Routing Protocols

Routing protocols fall under one of two major categories:

- **Distance vector**—Distance vector routing protocols, such as Routing Information Protocol (RIP), RIPv2, and Interior Gateway Routing Protocol (IGRP), make routing decisions based on information learned from neighbors. Therefore, distance vector routing protocols are said to use “routing by rumor.” Most distance vector routing protocols advertise their entire routing table to their neighbors on a periodic basis (with the exception of RIPv2 which uses triggered updates). Slow convergence is another common characteristic of these protocols. Therefore, distance vector routing protocols are not appropriate for large enterprise networks.
- **Link state**—Link-state routing protocols cause a router to flood information about itself (that is, the state of its links) to all the other routers in a network, or routers in part of a network (for example, an area). Based on the information received, each router can independently calculate what it believes to be the shortest path to a given destination network. Examples of link state routing protocols include Open Shortest Path First Protocol (OSPF) and Integrated Intermediate System-to-Intermediate System Protocol (IS-IS).

A network under a single administrative control is said to be an ***autonomous system***. Routing protocols running within an autonomous system are called ***interior gateway protocols*** (IGP). However, routing protocols are also needed to connect autonomous systems. For example, you might use OSPF as your IGP within an enterprise network, but you might need a separate routing protocol to connect your enterprise network to your service providers. This type of routing protocol that connects different autonomous systems is called an ***exterior gateway protocol*** (EGP). The only EGP in widespread use today is the Border Gateway Protocol (BGP).

The most popular routing protocols found in today's enterprise networks are as follows:

- **Enhanced IGRP (EIGRP)**—EIGRP is a Cisco-developed routing protocol that is considered to be an advanced distance vector protocol, because it is based on IGRP but also has link-state characteristics. Unlike some distance vector routing protocols, EIGRP uses triggered updates (as opposed to periodic updates). EIGRP uses a topology table to keep track of all the routes received from its neighbors. VLSM is supported, in addition to multiple network layer protocols, including IPv4, IPv6, AppleTalk, and IPX. EIGRP also offers fast convergence times if a router or link fails.
- **OSPF**—Like EIGRP, OSPF is well suited for enterprise networks due to its fast convergence and VLSM support. OSPF also uses the concept of ***areas*** to limit the number of route advertisements sent through the network. Specifically, OSPF has a backbone area (that is, Area 0), and all other areas must connect to Area 0. If you allocate IP addresses appropriately, the routers sitting at the borders between the Area 0 and the nonbackbone areas can summarize the routes within their area and send summary route information into Area 0. Figure 5-2 shows an example of an OSPF network. Notice that an Autonomous System Boundary Router (ASBR) connects the OSPF network with an external autonomous system.

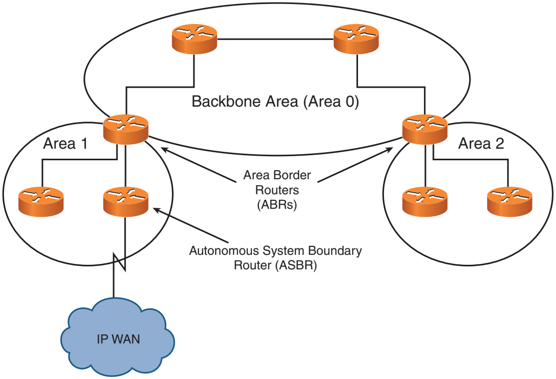
Backbone Area (Area 0)
Area 1
Area 2
Area Border Routers (ABRs)
Autonomous System Boundary Router (ASBR)
IP WAN

FIGURE 5-2 OSPF network example.

- **IS-IS**—Similar to OSPF, the IS-IS routing protocol is a link-state routing protocol that uses the concept of network areas. With IS-IS, the backbone area is called a Level 2 area and a nonbackbone area is called a Level 1 area. Routers that sit at the border between the backbone and a nonbackbone area are called Level 1/Level 2 (L1/L2) routers. IS-IS offers support for VLSM. However, IS-IS usually is deployed in service provider networks rather than enterprise networks.

- **BGP**—BGP is the routing protocol used on the Internet to connect different autonomous systems (for example, connecting an enterprise network's autonomous system to a service provider's autonomous system). However, some large enterprises use BGP to connect their network locations.

  BGP is highly tunable, allowing network administrators to influence BGP's path selection. For example, if your enterprise network connects to two service providers, each at different speeds, BGP could be manipulated to prefer the higher speed route. Figure 5-3 shows an example of a BGP network.

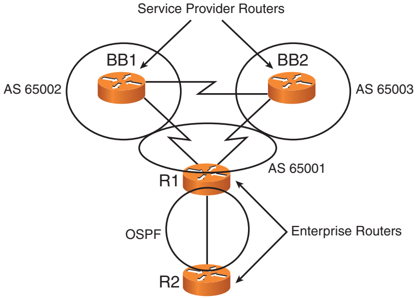

FIGURE 5-3 BGP network example.

Notice that the enterprise network contains routers R1 and R2. Within the enterprise network, OSPF is used as the IGP. The enterprise network connects with two service providers (that is, to routers BB1 and BB2) via BGP. The enterprise has an autonomous system number of 65001. The autonomous system numbers of the service providers are 65002 and 65003.

Routing Protocol Deployment

Enterprise network design requires quick convergence. Therefore, network designers often choose either OSPF or EIGRP, as previously described, for their IGP. Before selecting one routing protocol over the other, consider the following limitations:

- OSPF requires a hierarchical design, with all areas connecting to the backbone area. OSPF areas should map to a hierarchical addressing scheme. These requirements might not be practical or possible in all circumstances.
- EIGRP is a Cisco proprietary protocol. Therefore, EIGRP might not be appropriate in a mixed-vendor environment.

In addition to routing protocol selection, network designers should evaluate the following route manipulation techniques for their networks:

- **Route redistribution**—Route redistribution allows one routing protocol (for example, OSPF) to communicate its route information to another routing protocol (for example, EIGRP). As an example, this approach could support a mixed-vendor environment with Cisco routers using EIGRP and third-party routers using OSPF.
- **Route filtering**—Cisco routers support the filtering of selected routes in their routing updates. In some circumstances, route filtering can prevent routing loops and help provide optimal routing. In addition, a design might require that specific routes not enter a certain area of the network.
- **Route summarization**—The more routes a router must maintain in its routing table, the more router resources are consumed. Fortunately, route summarization can combine (that is, aggregate) multiple network addresses into a single network advertisement. For example, instead of advertising the individual networks 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24, these routes could be summarized as a single advertisement for network 10.1.0.0/16, which encompasses all the individually listed networks.

Network Security

An enterprise network design must include security measures to mitigate network attacks. Fortunately, with the modularity of the Cisco Enterprise Architecture, you can address security concerns on a module-by-module basis. This section introduces the concept of a security policy, reviews various types of network attacks, discusses the elements of the Cisco Self-Defending Network, and helps you select appropriate security design components for the various locations in an enterprise network.

Network Security Concepts

Organizational requirements and potential threats drive the scope of a security design. At its essence, network security measures should not only defend against attacks and guard against unauthorized access, these measures should also prevent data theft and comply with security legislation, industry standards, and company policy.

Consider the following threats and risks facing today's enterprise networks:

- **Threats:**
  - **Reconnaissance**—A reconnaissance attack gathers information about the target of an attack (for example, the customer's network). For example, a reconnaissance attack might use a port-scanning utility to determine what ports (for example, Telnet or FTP ports) are open on various network hosts.
  - **Gaining system access**—After attackers gather information about their target, they often attempt to gain access to the system. One approach is to use ***social engineering***, where they convince a legitimate user of the system to provide their login credentials. Other approaches for gaining access include exploiting known system vulnerabilities or physically accessing the system.
  - **Denial of service (DoS)**—A DoS attack can flood a system with traffic, thereby consuming the system's processor and bandwidth. Even though the attacker does not gain system access with a DoS attack, the system becomes unusable for legitimate users.
- **Risks:**
  - **Data confidentiality**—Companies should ensure that sensitive data on their systems is protected against theft. Without such protection, the company might be subject to legal liabilities and damage to the organization.
  - **Data integrity**—Besides stealing data, attackers could also modify sensitive data. Therefore, security measures should only allow authorized users to alter data.
  - **Data availability**—As previously mentioned, a DoS attack could make a system (and therefore the system's data) inaccessible by legitimate users. Therefore, security measures should be used to maintain system and data availability.

When designing a network security solution, realize that although hosts are the primary targets of an attack, other potential network targets also need protection. Other potential attack targets include routers, switches, DHCP/DNS (Dynamic Host Configuration Protocol/Domain Name System) servers, user PCs, IP phones, and IDS/IPS (intrusion detection system/intrusion prevention system) devices, in addition to the bandwidth available in the network infrastructure.

To guide security design decisions and provide a guideline to future security enforcement, organizations need to formulate a ***security policy***. A security policy is a documented set of rules that specify how people are allowed, or not allowed, to access an organization's technology and data.

Other considerations in a security design include the following:

- **Business needs**—Determine what the organization wants to accomplish with their network.
- **Risk analysis**—Determine the risk/cost ratio for the design.
- **Industry best practices**—Evaluate commonly accepted industry best practices for securing a network.
- **Security operations**—Define the process for monitoring security, performing security audits, and responding to security incidents.

In addition to a security policy, organizations might need to prepare the following documents to address specific risk categories:

- **Network access control policy**—This document defines levels of data security (for example, confidential or top secret) in the network and outlines procedures for gaining access to different security levels.
- **Acceptable-use policy**—This document should be distributed to all end users and be clear for what purposes a user is allowed to use the system and what types of data can be retrieved by the user.
- **Security management policy**—This document describes how an organization manages its network security.
- **Incident-handling policy**—For when security incidents occur, this document describes an orderly set of procedures for responding to the incident or an emergency situation.

The previously described security policy is a continually evolving document that changes in response to technology and organizational requirements. Like the continually evolving security policy, the process of securing the network is also continuous. Specifically, designers use the following four steps to continually secure the network, as illustrated in Figure 6-1:

- **Secure**—Securing the network involves such measures as authorizing and authenticating users, filtering unwanted traffic, encrypting data, and providing secure remote access using virtual private networks (VPN).
- **Monitor**—Monitoring the network involves the use of detection mechanisms (for example, IDSs) to send notifications if a security incident occurs.

- **Test**—Testing the network involves proactive verification of the network's security capabilities. For example, administrators might periodically perform vulnerability scanning on the network.
- **Improve**—Based on newly emerging security risks and analysis of the network's current ability to mitigate attacks, improved security measures are instated.

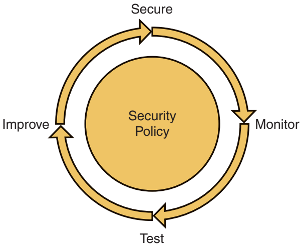

FIGURE 6-1 Network security process.

Cisco Self-Defending Network

Security needs to be fully integrated into a network to combat data theft. Fortunately, Cisco has defined the concept of the ***Self-Defending Network*** to leverage the security abilities of network components to protect the network from both internal and external threats. Network security integration consists of three components:

- **Trust and identity management**—Access is limited based on a user's access level. The three components of trust and identity management are as follows:
  - **Trust**—Defines how two or more network entities are allowed to communicate.
  - **Identity**—Validates the user accessing network resources. Identity can be proven by means such as passwords, tokens, or certificates.
  - **Access control**—Limits access to specific resources by specific users. The main concepts of access control are ***authentication*** (which determines the identity of the user) and ***authorization*** (which defines what a user is allowed to do on a network).
- **Threat defense**—Security breaches are minimized and mitigated through three primary approaches:
  - **Physical security**—Limits physical access to network resources.
  - **Infrastructure protection**—Takes measures to ensure network devices are not accessed or altered by an attacker.
  - **Threat detection and mitigation**—Threat detection and mitigation use technologies that provide proactive notification of suspicious network traffic patterns.

- **Secure connectivity**—Cryptography features provide the following protections for data flowing across a network:
  - **Privacy**—Privacy provides confidential communication through the network. The cryptographic service that offers confidentiality is encryption. Encryption scrambles data such that if an attacker were to intercept the data, the data would not be readable. However, the legitimate recipient of the data can decrypt the data into a readable form.
  - **Data integrity**—Cryptography mechanisms such as hashing algorithms and digital signatures can verify data was not manipulated in transit.

The Cisco Self-Defending Network is based on an underlying secure network platform (for example, Cisco routers, Cisco Catalyst switches, and Cisco Adaptive Security Appliances [ASA]). Layered on top of the network platform are advanced security technology and services. The use of these technologies is then governed by security policies and security management applications. These security management applications are used by network administrators to monitor and control the network.

If you properly plan security measures to protect your network architecture, the primary security risk is an error in security policies. Network managers and administrators must be intimately familiar with security policies and predefined procedures to respond to a security breach. A thorough understanding of these policies can help provide efficient incident response.

Cisco offers a suite of security management solutions, including the following:

- **Cisco Router and Security Device Manager (SDM)**—SDM offers a graphic user interface (GUI) to Cisco router configuration for features such as VPNs, quality of service (QoS), IPS, and Cisco IOS Firewall.
- **Cisco Adaptive Security Device Manager (ASDM)**—ASDM offers security management and monitoring features for devices such as the Cisco ASA 5500 series, Cisco PIX 500 series security appliances, and the Cisco Catalyst 6500 series Firewall Services Module (FWSM).
- **Cisco Intrusion Prevention System Device Manager (IDM)**—IDM is a Java application that supports the configuration and management of intrusion prevention sensors (IPS) through a web-based interface.
- **Management Center for Cisco Security Agents**—The Cisco Security Agent (CSA) is a Host Intrusion Prevention System (HIPS) that runs on hosts' machines, such as servers and personal workstations. The Management Center for Cisco Security Agents allows hosts to be classified into different groups and have different policies applied to the different groups.
- **Cisco Secure Access Control Server (ACS)**—Cisco Secure ACS is an application that supports identity-based services for a wide range of Cisco devices (for example, routers, switches, and firewalls). For example, instead of creating a username entry in every

router in the network for a newly hired administrator, the administrator could simply have an account added in an ACS server, which could be referenced by all routers in an organization.

- **Cisco Security Manager**—The Cisco Security Manager is a GUI-based application that aids in the configuration of firewalls, VPNs, and IPS policies on a variety of Cisco devices (for example, routers, switches, and firewalls).
- **Cisco Security Monitoring, Analysis, and Response System (Cisco Security MARS)**—Cisco Security MARS is a network appliance that allows network administrators to monitor, identify, contain, and combat network attacks.

The Cisco Self-Defending Network consists of three layers:

- **Integrated security**—Security technology is built in to network components such as routers, switches, and wireless devices.
- **Collaborative security systems**—Network security elements work in a collaborative fashion to enable the network as a whole to meet the goals of an organization's security policy.
- **Adaptive threat defense**—Behavior-recognition tools defend against emerging security threats and dynamic network conditions. These tools can defend against threats such as worms, viruses, spyware, and distributed DoS (DDoS) attacks.

Figure 6-2 shows an example of a network containing many of the elements of a Cisco Self-Defending Network.

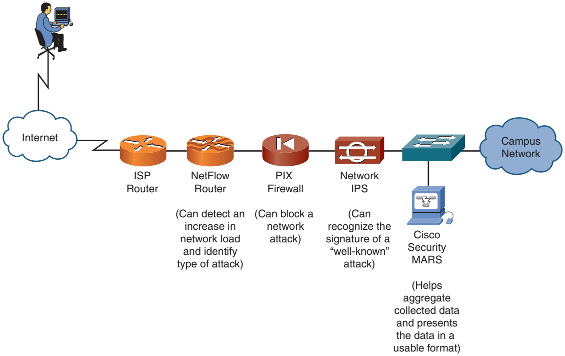
Internet
ISP Router
NetFlow Router
(Can detect an increase in network load and identify type of attack)
PIX Firewall
(Can block a network attack)
Network IPS
(Can recognize the signature of a "well-known" attack)
Cisco Security MARS
(Helps aggregate collected data and presents the data in a usable format)
Campus Network

FIGURE 6-2 Cisco Self-Defending Network example.

Network Security Solutions

To secure a network, integrate security solutions into all parts of the network. Consider how the following network elements integrate security solutions:

- **Cisco IOS router**—Depending on the feature set, a Cisco IOS router can act as a firewall/IPS. Also, a router can be used to set up an IPsec tunnel. Trust and identity solutions include authentication, authorization, and accounting (AAA), public key infrastructure (PKI), Secure Shell Protocol (SSH), and Secure Sockets Layer (SSL).

- **VPN concentrator**—Cisco VPN 3000 series concentrators are appliances that can be used for remote-access VPNs. For example, remote offices can use VPN concentrators to provide secure connectivity between a remote office and an organization's headquarters, as depicted in Figure 6-3.

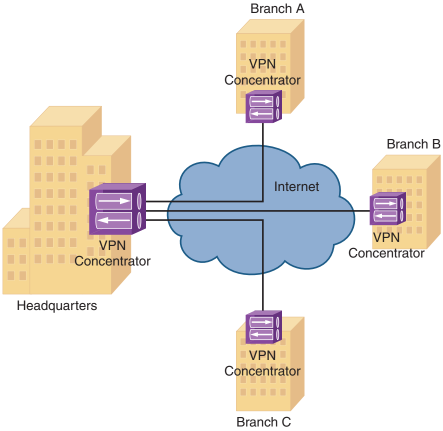

FIGURE 6-3 Site-to-site VPNs using VPN concentrators.

- **PIX Security Appliance**—A Cisco PIX Security Appliance is at its essence a firewall, which can provide application and protocol inspection to traffic flowing through the device.
- **Adaptive Security Appliance (ASA)**—Cisco ASA provides PIX-like firewall features, in addition to application security and support for advanced integration modules (for example, Cisco WebVPN Services Modules and Advanced Integration Modules [AIM]).
- **Intrusion Prevention System (IPS)**—Cisco offers a series of sensor appliances (for example, IPS 4215, 4240, 4255, and 4260 sensors) to provide IPS or IDS services. IPS works inline with data, and can stop suspicious traffic before the traffic reaches its destination. However, IDS is passive and receives a copy of network traffic, which it can analyze.
- **Cisco Catalyst Service Modules**—The Cisco Catalyst 6500 series switch is a modular switch offering support for a wide variety of service modules that can help enhance network security. Examples of these modules include Cisco Firewall Services Module (FWSM), Cisco Intrusion Detection System Services Module (IDSM-2), and the Cisco SSL Services Module.
- **Cisco Security Agent (CSA)**—CSA is software that can be installed on a host to defend against known and unknown (that is, "day zero") attacks. For example, CSA can protect a host from spyware and adware, in addition to protecting the integrity of the host's underlying operating system.

The Cisco Self-Defending Network allows network designers to specify security features throughout the network. Consider typical security solutions for the following enterprise network modules:

- **Enterprise campus**—The enterprise campus can benefit from the following security measures:
  - **Identity and access control**—802.1x, Network Access Control (NAC), access control lists (ACLs), and firewalls
  - **Threat detection and mitigation**—NetFlow, syslog, Simple Network Management Protocol (SNMP), Cisco Security MARS, Network IPS (NIPS), and Host IPS (HIPS)
  - **Infrastructure protection**—AAA, SSH, SNMPv3, Interior/Exterior Gateway Protocol (IGP/EGP) message digest 5 algorithm (MD5), and Layer 2 security features
  - **Security management**—Cisco Security Manager and MARS
- **Enterprise data center**—Similarly, the enterprise data center can leverage these security technologies:
  - **Identity and access control**—802.1x, ACLs, and firewalls
  - **Threat detection and migration**—NetFlow, syslog, SNMP, Cisco Security MARS, NIPS, and HIPS
  - **Infrastructure protection**—AAA, SSH, SNMPv3, IGP/EGP MD5, and Layer 2 security features
  - **Security Management**—Cisco Security Manager and Cisco Security MARS
- **Enterprise edge**—The enterprise edge module can benefit from such security measures as
  - **Identity and access control**—Firewalls, IPsec, SSL, VPN, ACLs
  - **Threat detection and mitigation**—NetFlow, syslog, SNMP, Cisco Security MARS, NIPS, and HIPS
  - **Infrastructure protection**—AAA, Control Plane Policing (CoPP), SSH, RFC 2827 (an approach for defeating DoS attacks that use IP source address spoofing), SNMPv3, and IGP/EGP MD5
  - **Security management**—Cisco Security Manager and Cisco Security MARS

Identifying Voice Networking Considerations

Many of today's enterprise network designs must accommodate the transmission of voice traffic in addition to data traffic. The transmission of voice over a data network is often referred to as Voice over IP (VoIP). The inclusion of VoIP in a network design typically requires integration with existing telephony services and connectivity into the public switched telephone network (PSTN). Therefore, this section reviews existing telephony networks, discusses traffic engineering, and offers design guidance for VoIP networks.

Reviewing Traditional Voice Architectures and Features

Before recommending VoIP network design solutions, a designer should first become familiar with traditional telephony networks. A fundamental concept in traditional telephony networks is the conversion of human speech into a digital signal.

When you speak into an analog phone, your voice is converted into an analog waveform. However, telephony networks cannot maintain voice quality when sending analog waveforms over long distances. Therefore, telephony networks convert analog waveforms into digital signals, which can be transmitted over great distance.

The steps for converting an analog waveform into a digital signal include the following:

- **Filtering**—Approximately 90 percent of the frequencies required to understand human speech are in the range of 300 Hz to 3400 Hz. Therefore, to filter out extraneous noise, a coder-decoder (codec) filters out frequencies greater than 4000 Hz.
- **Sampling**—Based on the Nyquist theorem, which says an analog waveform needs to be sampled at a rate that is at least double the highest frequency being sampled, the analog waveform is sampled at a rate of 8000 samples per second (that is, twice the highest frequency of 4000 Hz), as shown in Figure 7-1.

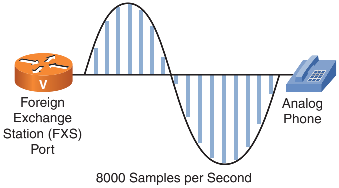

FIGURE 7-1 Sampling an analog waveform.

- **Digitizing**—When the analog waveform is sampled, the amplitude (that is, the volume) of each sample is represented as a number. This process is called ***quantization***. However, because each possible amplitude does not have an associated number, the measure of each amplitude is rounded off to the nearest number on a scale. For example, consider Figure 7-2, which shows how these amplitudes are rounded off on a linear scale. This rounding off can cause ***quantization noise***.

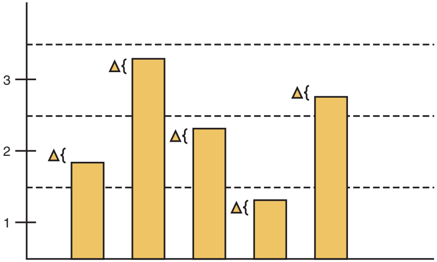

FIGURE 7-2 Linear quantization.

Instead of using a linear scale, quantization typically uses a logarithmic scale, so that more accurate measurements can be made at lower volumes. Accuracy at lower volumes is more important than accuracy at higher volumes because most samples have lower volumes, and higher volumes tend to mask the noise. The methods for constructing the logarithmic scale are called ***companding*** (that is, compressing and expanding) types. The two primary companding types in use today are a-law, which is most popular in Europe, and mu-law (sometimes written as u-law), which is most popular in North America and Japan.

Just as a router can make decisions about how packets should be forwarded through a network (for example, based on an IP address), a telephone switch makes call routing decisions (for example, based on a dialed telephone number) for forwarding a voice call through a telephony network. Although the PSTN contains a series of telephone switches (sometimes referred to as central office [or CO] switches), organizations can have their own telephone switches. An example of a privately owned switch is a Private Branch Exchange (PBX). Although a PBX does not scale to the degree a PSTN switch does, a PBX does offer enhanced telephone features to organizations (for example, call hold, conferencing, transferring, music on hold, call forwarding, call park, and voice mail). Also, PBX vendors often use their own proprietary call signaling protocols, whereas PSTN switches use standards-based signaling protocols.

Figure 7-3 demonstrates how telephone switches are connected with the following trunk types:

- **Tie trunk**—Tie trunks interconnect PBXs.
- **PBX-to-CO trunk**—PBX-to-CO trunks (sometimes just called “CO trunks”) connect an organization’s PBX to the PSTN.
- **Interoffice trunk**—Interoffice trunks connect the CO switches that make up the PSTN.
- **Local loop**—A local loop is the connection from a CO switch to a telephony device at the subscriber’s location (for example, an analog phone in a residence).

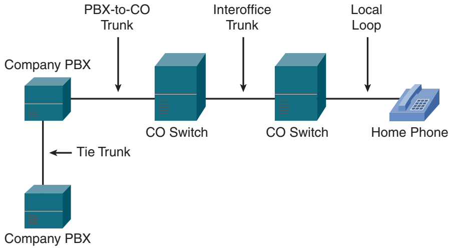

FIGURE 7-3 Trunk types.

Telephone switches use various forms of signaling to set up, maintain, monitor, and tear down calls. The three fundamental categories of signaling are as follows:

- **Supervisory signaling**—Supervisory signaling allows, for example, a telephone switch to determine whether an attached phone is in the on-hook or off-hook condition. Sending ringing voltage to a phone is another example of supervisory (also called "supervision") signaling.
- **Address signaling**—Address signaling is used to transmit dialed digits (for example, dual-tone multifrequency [DTMF] tones generated when you press keypad buttons on a phone).
- **Information signaling**—Information signaling provides feedback about the state of a call to the caller. For example, if you call someone and hear a busy signal, the busy signal is information signaling letting you know that the called party is not available.

Signaling information can be communicated over analog or digital connections. Common types of analog signaling include the following:

- **Loop start signaling**—A traditional home phone is an example of a phone that uses loop start signaling. When you pick up the handset, loop current begins to flow, telling the telephone switch that the phone is off-hook. However, loop start signaling can suffer from ***glare***, where someone is calling you and you pick up the handset to place a call before you hear the phone ring. You expect to hear dial tone, but instead you hear the calling party. Although this might occur infrequently in a home environment, because a PBX shares lines, the use of loop start signaling could lead to excessive glare in PBX environments.
- **Ground start signaling**—Ground start signaling prevents glare. Therefore, ground start signaling is preferred for PBXs, as opposed to loop start.
- **E&M (ear and mouth)**—E&M signaling (sometimes called "recEive and transMit" or "earth and magneto") is used to connect PBXs. Whereas both loop start and ground start each use two wires (that is, tip and ring) to carry both voice and signaling, E&M uses separate wires (that is, the *E & M* wires) for signaling, while still using the tip and ring wires to carry voice.

Although analog connections might be appropriate for lower port densities, if you need many connections coming into a PBX or between

PBXs, digital circuits often offer a more cost-effective alternative. The two main types of digital signaling are as follows:

- **Channel-associated signaling (CAS)**—Consider a T1 circuit. A T1 is a digital circuit with 24 64-kbps channels. With CAS, all 24 channels can be used to carry voice traffic. The signaling information is transmitted by using specific framing bits, which are not needed because most T1s send 24 T1 frames together in what is called a *superframe*. Because these unneeded framing bits are used for signaling instead of framing (that is, to indicate the beginning of a frame), CAS is sometimes called "robbed bit signaling."
- **Common channel signaling (CCS)**—With CCS, one or more channels in a digital circuit (for example, a T1) are used solely to carry signaling information. Therefore, with most T1 CCS implementations, the T1 circuit can carry 23 voice calls, with the twenty-fourth channel used to carry signaling information. ISDN is an example of CCS. ISDN sends voice, data, and video traffic in bearer channels (that is, B channels), with the signaling being carried in a D channel.

Some PBX vendors use their own proprietary signaling protocols. Therefore, connecting PBXs in a mixed-vendor environment can be a challenge. However, many PBX vendors support the ***Q Signaling*** protocol (that is, QSIG), which allows PBXs from different vendors to communicate with one another. Similarly, CO switches also have a common signaling protocol, called Signaling System 7 (SS7).

Just as data networks can benefit from hierarchical IP addressing, telephony networks often benefit from a hierarchical numbering plan. A ***numbering plan*** is a set of rules that dictate how telephone numbers are assigned and how voice calls are routed. For example, consider the North American Numbering Plan (NANP). NANP numbers use a numbering format of *NXX-NXX-XXXX*, where *N* can be any digit from 2 through 9 and *X* can be any digit from 0 through 9. Notice the first *NXX*. In North America, this digit pattern is an area code. The next *NXX* pattern represents the local office code, and the final *XXXX* pattern represents the subscriber's number. Notice that in North America, neither an area code nor an office code can begin with a 0 or a 1.

Integrating Voice Architectures

Packet telephony network designers must familiarize themselves with new terms and standards not typically encountered in data network design. Specifically, designers need understanding of integrated voice architecture concepts, standards, and design challenges.

Traditionally, organizations kept their voice, data, and video networks separate. As a result, a data burst on the data network had no adverse effect on voice traffic. However, with the advent of higher bandwidth, more reliable, quality of service (QoS)-enabled networks, network designers are beginning to see the wisdom of combining voice, data, and video on the same converged network.

The two primary approaches of sending voice over a data network are as follows:

- **VoIP**—VoIP networks allow traditional telephony devices (for example, analog phones, PBXs, key systems, and the PSTN) to attach to a voice-enabled router. The router packetizes the voice and signaling traffic from the traditional network and transports that traffic over an IP network.
- **IP telephony**—An IP telephony network, like a VoIP network, transmits voice and signaling traffic in IP packets. However, the distinction between an IP telephony network and a VoIP network is an IP telephony network includes IP-based voice devices (for example, IP phones that contain an Ethernet port and connect directly to a network).

Both VoIP and IP telephony networks require gateways to convert voice and signaling information between the traditional telephony environment (such as a PBX or the PSTN) and the IP environment. These gateways communicate using gateway control protocols (sometimes called call control protocols).

The most mature of the gateway control protocols is H.323. The H.323 standard not only defines a suite of protocols, but it also includes hardware specifications for physical components in an H.323 network.

Among the H.323 protocols used for call setup are the following:

- **H.225.0**—The H.225.0 protocol (often written as H.225) has a couple of functions. H.225.0 can use TCP to send the initial call setup message between a couple of H.323 endpoints. Also, H.225.0 can use User Datagram Protocol (UDP) for communication with an H.323 gatekeeper (which can be used to resolve phone numbers to IP addresses and grant or deny a call to be placed, based on bandwidth availability).
- **H.245**—When the H.225.0 protocol initiates the call setup process between two H.323 endpoints, the H.245 protocol negotiates the parameters of the call (for example, how the voice will be encoded and which UDP ports to use when sending voice traffic).

H.323 hardware specifications include the following:

- **Terminal**—An H.323 terminal acts as an endpoint in a call (for example, a user's PC running H.323-enabled software).
- **Gateway**—An H.323 gateway converts voice and signaling information between different environments (for example, the traditional telephony environment and the IP environment).
- **Gatekeeper**—Two of the most important jobs of an H.323 gatekeeper are the following:
  - **Number resolution**—H.323 uses IP addresses to set up calls. However, users typically dial phone numbers rather than specify IP addresses. The gatekeeper can perform phone number to IP address resolution.
  - **Admission control**—If too many calls are simultaneously placed over an IP WAN, the quality of all calls suffers.

Fortunately, an H.323 gatekeeper can be used to reject a call attempt if that call would oversubscribe the IP WAN's available bandwidth.

- **Multipoint control unit (MCU)**—H.323 networks support conference calls. However, processing power is required to mix together multiple audio streams. An H.323 MCU can perform that mixing.

An IP telephony network, such as the one pictured in Figure 7-4, has the following core components:

- **Infrastructure**—An IP telephony network runs on an underlying infrastructure composed of network layer switches and voice-enabled routers.
- **Call processing**—Cisco Unified CallManager software (available for either a Windows 2000 or Linux platform) performs PBX-like functions (for example, call routing) for an IP telephony network.
- **Applications**—Other than basic call setup, IP telephony networks can offer a wide variety of applications, such as unified messaging, interactive voice response, Cisco Unified Contact Center, and Auto Attendant.
- **Client devices**—Users interface with an IP telephony network via client devices such as Cisco IP Phones. However, a client device could be a software-based phone, such as Cisco IP Communicator.

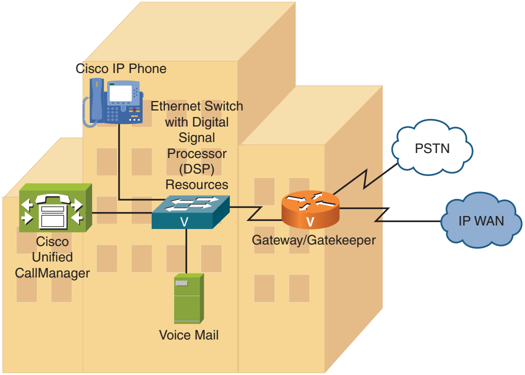
Cisco IP Phone
Ethernet Switch with Digital Signal Processor (DSP) Resources
PSTN
IP WAN
Gateway/Gatekeeper
V
V
Cisco Unified CallManager
Voice Mail

FIGURE 7-4 IP telephony network.

Because many organizations have multiple locations, their IP telephony networks might span those locations. When determining how IP telephony components should be deployed, consider the following deployment models:

- **Single-site deployment**—If an IP telephony network is contained within a single location, as illustrated in Figure 7-5, and has fewer than 30,000 phones, a single-site deployment model is often appropriate.

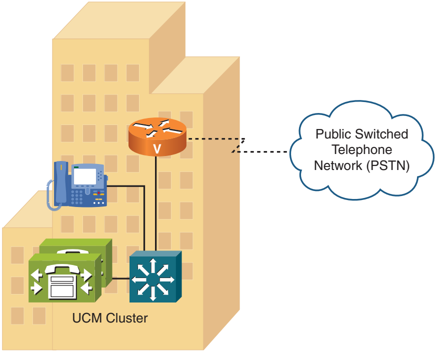

FIGURE 7-5 Single-site deployment.

- **Multisite WAN with centralized call processing deployment**—Some organizations might have smaller remote sites that do not contain enough IP phones to justify the purchase of UCM servers for those locations. In those instances, the UCM servers could be located at the headquarters, and IP phones at the remote offices could then register with the centralized UCM servers over the IP WAN. If there is an IP WAN outage, IP phones could register with the local Survivable Remote Site Telephony (SRST) routers located at each remote site, for basic call processing functionality. Figure 7-6 shows an example of this multisite WAN with centralized call processing deployment model.

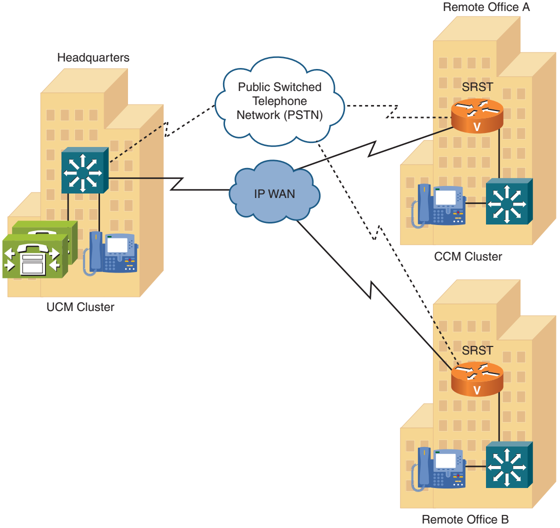

FIGURE 7-6 Multisite WAN with centralized call processing deployment.

- **Multisite WAN with distributed call processing deployment**—When designing a large IP telephony network with multiple locations, the expense of placing UCM servers at each location might be justified. As an example, Figure 7-7 provides a sample IP telephony topology using the multisite WAN with distributed call processing deployment model.

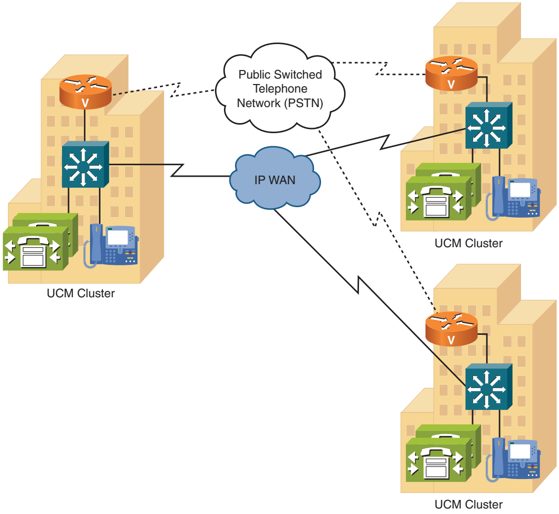

FIGURE 7-7 Multisite WAN with distributed call processing deployment.

Although H.323 is a very popular gateway control protocol for IP telephony and VoIP networks, consider some of the other protocols you might encounter in IP telephony or VoIP networks:

- **Real-time Transport Protocol (RTP)**—Voice packets are carried inside of RTP segments. RTP is a Layer 4 protocol that is encapsulated inside UDP segments.
- **Skinny Client Control Protocol (SCCP)**—By default, Cisco IP Phones use SCCP to exchange signaling messages with Cisco Unified CallManager. Unlike H.323 (which is considered a peer-to-peer protocol), SCCP is considered to be a client/server protocol.
- **Session Initiation Protocol (SIP)**—SIP is a peer-to-peer gateway control protocol that is popular in many mixed-vendor environments. When you are adding Cisco IP telephony components to an existing third-party IP telephony network, SIP might serve as an appropriate gateway control protocol.
- **Media Gateway Control Protocol (MGCP)**—MGCP is a client/server gateway control protocol. In a Cisco IP telephony environment, a Cisco Unified CallManager server acts as the "server," and a port on a router (for example, an analog Foreign Exchange Station [FXS] port) acts as the "client."

Identifying the Requirements of Voice Technologies

When designing a network to accommodate voice traffic, consider what could impact the quality of the voice and which mechanisms might be used to maintain voice quality.

When voice and data traffic are contending for limited bandwidth, the following quality issues might arise:

- **Delay**—The ITU G.114 recommendation for voice traffic specifies a maximum one-way delay of 150 ms for voice traffic. Some types of delay are considered ***fixed***, in that they do not change during a phone call. Examples of these fixed delay components include propagation delay (the time it takes a packet to traverse a network link), serialization delay (the time it takes to send a frame out of a serial link), and processing delay (the time required by the router to encode/decode, compress/decompress, and packetize voice).
- **Jitter**—Variable delay might vary during a phone call. One example of variable delay is ***jitter***. Specifically, jitter is the uneven arrival of packets at a destination router. Cisco routers use dejitter buffers to help smooth out packet playout, thus concealing the jitter experienced by those packets. Another type of variable delay, which can contribute to jitter, is ***queuing delay***. Queuing delay is the amount of time a packet must spend in a queue as it waits to be forwarded out of an interface.
- **Packet drops**—If an interface's output queue fills to capacity, newly arriving packets might be dropped. This occurrence is called ***tail drop***. Although digital signal processors can correct a maximum of approximately 30 ms of lost voice, additional voice packet drops can severely compromise voice quality.

Although not related to limited bandwidth, echo causes another serious problem for voice quality. You experience the symptom of echo when you speak and hear your own voice reflected back to you, or when you speak and the other party hears your voice twice. The issue of echo typically stems from an impedance mismatch in a two-wire to four-wire circuit, which can be found in an analog phone or in telephony switching equipment.

To combat echo, Cisco voice-enabled routers can use ***echo cancellation***, which allows a voice port to "memorize" waveforms being sent out of the interface for a period of time (typically 8–32 ms). If the voice port sees the same waveform coming back in the interface within that period of time, the voice-enabled router can cancel the echo waveform by superimposing the same waveform, which has been phase-shifted 180 degrees. Silence results from playing the same waveform twice, when those waveforms are 180 degrees out of phase.

However, because most quality issues on IP telephony and VoIP networks result from limited bandwidth, network designers use a variety of approaches to make the best use of this limited bandwidth, such as the following:

- **Codec selection**—One approach is to use a codec requiring less bandwidth per call. For example, the G.711 codec does not perform any compression, and it requires 64 kbps of bandwidth (not including overhead) for a single voice call. However, over an IP WAN, where bandwidth is at a premium, Cisco networks often use the G.729a codec, which only requires 8 kbps of bandwidth (not including overhead). Because G.729a performs compression, whereas G.711 does not, voice quality is somewhat compromised when using G.729a.
- **The mean opinion score (MOS)**—The MOS metric is used to measure voice quality, on a five-point scale, with larger numbers representing better quality. The G.711 codec has an MOS score of 4.1; G.729a's MOS score is 3.9. This slight, and barely perceptible, quality difference is often an acceptable trade-off to reduce bandwidth demand.
- **RTP header compression (cRTP)**—When using G.729a, voice packets contain 20 bytes of voice payload, while the packet contains 40 bytes of header information. However, because most information in these headers is identical (for example, the same source/destination IP address/UDP port numbers and the same RTP payload type), cRTP does not send this redundant header information in each frame. Therefore, cRTP reduces the 40-byte header down to only 2 or 4 bytes, allowing more calls to be placed over the same link speed.
- **Voice activity detection (VAD)**—Statistics show that approximately 35 percent of all voice calls are silence. Instead of consuming bandwidth to send "the sound of silence," VAD can detect the silence and suppress the transmission of silence.

Because network designers are concerned with bandwidth use, they must understand how to calculate required bandwidth. The following formula shows how to calculate a network's required voice bandwidth:

Bandwidth = ((Layer 2 header) + (IP/UDP/RTP header)) * (Codec bit rate) / (Voice payload size)

When working with this formula, make the following assumptions:

- IP/UDP/RTP header = 40 bytes
- With cRTP, the header = 2 or 4 bytes
- A WAN's Layer 2 header = 6 bytes
- An easier, and more detailed, bandwidth calculation can be performed using the Cisco Voice Codec Bandwidth Calculator, available at http://tools.cisco.com/Support/VBC/do/CodecCalc1.do.

NOTE

Your Cisco.com account must have appropriate access permissions to reach the Voice Codec Bandwidth Calculator URL.

To combat the quality issues described earlier, you can implement various QoS mechanisms available on Cisco routers and switches. For example, on wiring closet Catalyst switches, voice and data traffic can be placed in separate queues. Also, these Catalyst switches can be configured not to trust priority markings originating from a PC connected to a Cisco IP Phone.

Router QoS mechanisms include the following:

- **Classification and marking**—Classifying traffic recognizes characteristics of traffic and categorizes that traffic. As an example, access control lists (ACL) can be used to classify traffic. Once categorized, the traffic can be marked by, for example, altering bits in a packet's header to indicate the packet's relative level of priority.
- **Congestion management**—Congestion management defines the queuing algorithm used by an interface's output queue. The queuing algorithm can specify which type of traffic receives priority treatment (that is, forwarded out of the interface ahead of other traffic) and how much bandwidth is available to various traffic types during periods of network congestion. Cisco's recommended queuing mechanism for voice networks is low-latency queuing (LLQ).
- **Congestion avoidance**—To prevent an interface's output queue from filling to capacity, after which newly arriving packets are discarded, routers can use a congestion avoidance mechanism (such as weighted random early detection [WRED]) to increase the probability that lower-priority traffic will be discarded as the queue begins to fill.
- **Traffic conditioning**—Traffic-conditioning mechanisms (for example, policing and shaping) limit the amount of bandwidth that can be consumed by specific traffic types.
- **Link efficiency**—Link-efficiency mechanisms, such as link fragmentation and interleaving (which fragments larger packets and interleaves voice packets in among the fragmented data packets, thus reducing the serialization delay experienced by the voice traffic) and RTP header compression, attempt to make the most efficient use of limited WAN bandwidth.

As mentioned earlier, if too many simultaneous calls are sent across an IP WAN, and the IP WAN becomes oversubscribed, all calls experience poor voice quality. Therefore, IP telephony and VoIP networks require call admission control (CAC) tools to prevent this oversubscription. One approach to CAC is to use the previously described gatekeeper. Another approach is to use the Resource Reservation Protocol (RSVP). With RSVP, a Cisco voice-enabled router, or a Unified CallManager server (Version 5.0 or later), can reserve network bandwidth for a voice call that no other application can encroach on, thus preventing IP WAN oversubscription.

Because most QoS issues described result from insufficient bandwidth, a network designer needs to provision enough bandwidth to support projected traffic loads during a network's busiest hour of the day. The process of calculating the required amount of bandwidth is called

traffic engineering. The concept of traffic engineering dates back to PBX design, where designers needed to calculate the number of trunks between a PBX and the local CO. With IP telephony and VoIP networks, you take traffic engineering a step further by converting the calculated number of trunks into a bandwidth amount.

Although the mathematics behind traffic engineering can be quite rigorous, the following steps present a simplified approach:

1. Determine the grade of service (GoS).

   Because designing a voice network with enough trunks to prevent any incoming calls from receiving a busy signal is typically not cost effective, the designer must determine what percentage of calls can be rejected (that is, receive a busy signal) during the busiest hour of the day for an organization's telephone system. This percentage is called the grade of service, or GoS. Most designs use a GoS of 1 percent, which is written *P(.01)*.

2. Determine the busy hour traffic (BHT).

   The call volume experienced by an organization's telephone system (for example, a PBX) is measured in Erlangs, where an Erlang equals one solid hour of phone usage. Statistically, the number of Erlangs a corporate phone system experiences during the busiest hour of the day can be approximated by getting the number of hours of phone use during the previous month from your organization's telephone bill and using the following formula:

   Busy hour Erlangs = [Monthly_call_hours / 22] * .15

3. Calculate the number of required trunks.

   Usually, after you have determined the GoS and the number of Erlangs experienced during an organization's busiest hour of the day, you can use an Erlang B table to determine the number of required trunks (that is, simultaneous connections). You can refer to an Erlang B table to calculate the number of required trunks, or you can use a web-based Erlang B calculator, such as the one available at http://erlang.com/calculator/erlb.

4. Convert the number of required trunks to the amount of required bandwidth.

   Use the Cisco Voice Codec Bandwidth Calculator, as described earlier, to convert the number of required trunks into the amount of required bandwidth.

Identifying Wireless Networking Considerations

Wireless networks are experiencing widespread growth because of their availability, flexibility, and service offerings. This section introduces the Cisco unified wireless network architecture. Specifically, after an introduction of the Cisco unified wireless network, this section examines network controller technologies and presents guidelines for wireless network design in enterprise networks.

Introducing the Cisco Unified Wireless Network

Wireless local-area networks (WLAN) offer network access via radio waves. Wireless clients (such as a PC or PDA) access a wireless access point, using half-duplex communication. The wireless access point allows a wireless client to reach the rest of the network.

Traditional WLANs use an access point in ***autonomous mode***, where the access point is configured with a service set identifier (SSID), radio frequency (RF) channel, and RF power settings. However, having an autonomous access point tasked with all these responsibilities can limit scalability and can hinder the addition of advanced wireless services.

Five primary components comprise the Cisco unified wireless network architecture:

- **Clients**—A wireless client device is typically an end-user device (such as a PC) that accesses a wireless network.
- **Access point**—Wireless access points offer network access for wireless clients.
- **Network unification**—To offer wireless clients access to an organization's resources, the wireless network must be integrated (that is, unified) with the wired LAN.
- **Network management**—Just as enterprise LANs benefit from network management solutions, a wireless LAN can also use network management solutions to enhance security, reliability, and to offer assistance in WLAN deployments. An example of a wireless network management solution is the Cisco Wireless Control System (WCS).
- **Mobility**—Wireless mobility services include security threat detection, voice services, location services, and guest access.

Aside from autonomous mode, Cisco unified wireless networks can alternatively operate in ***split-MAC*** mode. With split-MAC operation, an access point is considered to be a "lightweight" access point, which cannot function without a wireless LAN controller (WLC).

Specifically, a wireless LAN client sending traffic to the wired LAN sends a packet to a lightweight access point, which encapsulates the packet using the Lightweight Access Point Protocol (LWAPP). The encapsulated traffic is sent over an LWAPP tunnel to a WLC. LWAPP sends packets in a Layer 2 frame with an Ethertype of 0xBBBB. LWAPP data traffic uses a destination port of 12222; LWAPP control traffic uses a destination port of 12223.

The lightweight access point, as shown in Figure 8-1, performs functions such as beaconing, packet transmission, and frame queuing; the WLC assumes roles such as authentication, key management, and resource reservation.

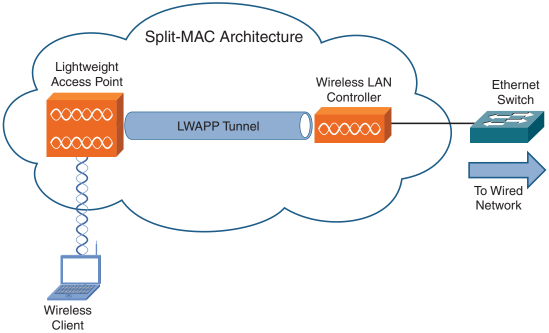

FIGURE 8-1 LWAPP architecture.

The operation of the wireless access point discussed thus far is referred to as *local mode*. However, several other access point modes exist:

- **Remote edge access point (REAP) mode**—REAP allows an access point and a WLC to be separated by a WAN, as opposed to being connected on the same LAN.
- **Rogue detector mode**—Route access points can be monitored by a wireless access point operating in rogue detector mode.
- **Monitor mode**—Wireless access points can be set to a receive-only mode, called ***monitor mode***, and act as sensors for location-based services (LBS).
- **Sniffer mode**—Wireless access points operating in sniffer mode can act as a protocol sniffer and capture packets, which are forwarded to a PC running the AiroPeek software.
- **Bridge mode**—Geographically separated wireless access points can be connected using a high-bandwidth, cost-effective wireless link, by running in bridge mode.

After a wireless client, such as a PC, associates with its access point, the access point only allows the client to communicate with the authentication server until the client successfully logs in and is authenticated, as illustrated in Figure 8-2. The WLC uses the Extensible Authentication Protocol (EAP) to communicate with the authentication server. Cisco Secure Access Control Server (ACS) could, for example, act as the authentication server.

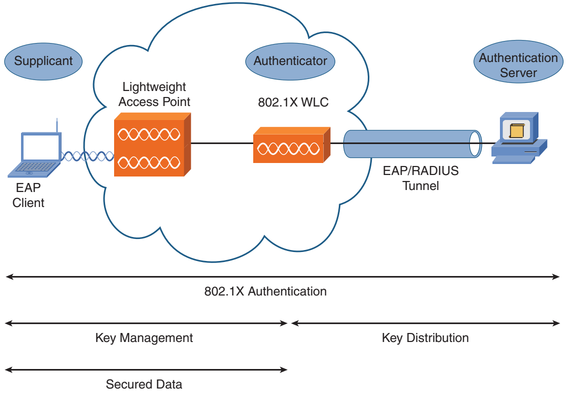

FIGURE 8-2 Wireless authentication.

Supported EAP types include the following:

- **EAP-Transport Layer Security (EAP-TLS)**—Wireless clients and authentication servers mutually authenticate using digital certificates.
- **EAP-Protected EAP (EAP-PEAP)**—The authentication server (that is, a RADIUS server) is authenticated over a Transport Layer Security (TLS) tunnel using a digital certificate; wireless clients are authenticated via EAP-GTC or EAP-MSCHAPv2.
- **EAP Tunneled Transport Layer Security (EAP-TTLS)**—The RADIUS server is authenticated over a TLS tunnel using the server's certificate, and wireless clients authenticate using username and password credentials.
- **Cisco Lightweight Extensible Authentication Protocol (LEAP)**—Cisco developed LEAP as an early and proprietary EAP method. However, LEAP's vulnerability to a dictionary attack represents a major LEAP weakness.
- **Cisco EAP-Flexible Authentication via Secure Tunneling (EAP-FAST)**—Cisco proposed EAP-FAST to address LEAP's weaknesses.

Designers should understand the following three WLAN controller components:

- **Ports**—A port on a WLAN controller physically connects the WLAN controller to the wired network (for example, to a Cisco Catalyst switch port).
- **Interfaces**—An interface of a WLAN controller logically maps to a VLAN on the wired network.
- **WLANs**—A wireless LAN can be configured with security features, quality of service (QoS) mechanisms, and other wireless network parameters. Also, a WLAN associates an SSID to a WLC's interface.

Cisco offers an array of WLCs. Different controllers support a different number of access points, as shown in Table 8-1.

TABLE 8-1 Access Point Support for WLCs

WLC Model	Number of Supported Wireless Access Points
Cisco 2000 series WLC	6
Cisco WLC module for ISRs	6
Cisco Catalyst 3750G integrated WLC	Up to 50
Cisco 4400 series WLC	Up to 100
Cisco Catalyst 6500 series wireless services modules	Up to 300

Understanding Wireless Network Controller Technologies

Lightweight access points do not require direct configuration and are therefore considered to be ***zero touch*** devices. After installing a lightweight access point, the access point goes through the following discovery process to discover a WLC:

1. The lightweight access point sends a DHCPDISCOVER request to dynamically obtain an IP address, unless it already had a statically configured IP address.
2. The lightweight access point broadcasts an LWAPP discovery message in a Layer 2 LWAPP frame, if the access point supports Layer 2 LWAPP transport mode.
3. If step 1 was unsuccessful or if the access point lacks Layer 2 LWAPP transport mode support, the access point attempts Layer 3 LWAPP WLC discovery.
4. If all steps were unsuccessful, the process begins again.

Based on the results of the discovery process, the lightweight access point selects which WLC to join. During the join process, the WLC validates the access point, and an encryption key is derived. This key is then used to encrypt and decrypt messages exchanged between the access point and the WLC.

Next, the lightweight access point and WLC perform the following steps:

1. The WLC configures the lightweight access point with an SSID, security parameters, QoS settings, and other such parameters.
2. Periodically, the WLC checks the status of the access point via query messages.
3. Every 30 seconds the access point transmits an LWAPP heartbeat, and if no acknowledgment is received after five attempts, the access point seeks a new WLC to join.

Wireless networks offer users ***mobility***, where the users can physically move throughout a campus. As the users move, their wireless clients update their access point association to the most appropriate access point, based on location.

Low-quality roaming requires that wireless clients obtain a new IP address (via DHCP), and potentially receive new security settings, as the clients move through the WLAN. This type of wireless environment can suffer from noticeable delays during the reassociation period, which makes such a solution inappropriate for voice calls.

High-quality roaming (that is, the mobility feature) does not require wireless clients to obtain a new IP address or update their security settings, thus providing seamless roaming. Mobility requires the seamless roaming experience to be maintained even if the access points, between which a client roams, are associated with different WLCs. The mobility feature also needs to support Layer 2 or Layer 3 roaming.

With Layer 2 roaming, the WLCs with which the access points associate are in the same subnet. However, with Layer 3 roaming, the access points associate with WLCs on different subnets.

When a wireless client associates with a new access point, the new access point's WLC exchanges mobility messages with the old access point's WLC. The client entry is not moved from the client database of the old WLC to the new WLC. Instead, the old WLC marks the client with an *anchor* entry, and the database entry is copied to the new WLC client database where it is marked as a *foreign* entry.

Wireless mobility groups allow WLCs in a network to form peering relationships. These peering relationships allow a mobility group to support seamless roaming between WLCs, wireless access point load balancing, and WLC redundancy. Keep the following requirements in mind when designing a mobility group:

- The management interfaces of all WLCs must be able to reach each other via IP.
- All WLCs in a mobility group must be configured with the same mobility group name, which is case sensitive.
- The same virtual IP address must be configured on all WLCs.
- The MAC addresses and IP addresses of all mobility group members must be configured on all WLCs.
- WLCs must be able to communicate with one another using UDP port 16666 for unencrypted messages or using UDP port 16667 for encrypted messages.

When designing a wireless network to support roaming, consider the following recommendations from Cisco:

- Use roaming only when necessary.
- Ensure the route-trip time between WLCs is less than or equal to 10 ms.
- When possible, use Layer 2 roaming rather than Layer 3 roaming.
- Implement Proactive Key Caching (PKC) or Cisco Centralized Key Management (CCKM) to help speed up and secure the roaming process.

Because a WLC could become a single point of failure, when designing WLANs, consider adding WLC redundancy. WLCs support either *dynamic* or *deterministic* redundancy. Specifically, an access point selects a WLC using the following sequence:

- **Deterministic**—An access point can be preconfigured with a primary, secondary, or tertiary WLC. The access point can then attempt to join those controllers in the specified order. Consider the following deterministic redundancy designs:

  ***N* + 1**—One controller backs up *N* controllers.

  N* + *N—*N* controllers back up *N* controllers.

  ***N* + *N* + 1**—*N* controllers back up *N* controllers as secondary, and one controller backs up all *N* controllers as tertiary.

- **Initializing**—Typically used only for the initial access point deployment, the WLC can attempt to join the WLC configured as a master controller.
- **Dynamic**—The access point uses a decision making algorithm to select a WLC based on the greatest availability for access point associations. Dynamic WLC redundancy uses LWAPP to perform load balancing across WLCs and to provide backup WLC information to the access points. This approach is often appropriate for a design where WLCs are clustered together at a central location.

The number of devices supported by an access point varies depending on the application being used. For example, Cisco recommends no more than seven or eight voice over WLAN (VoWLAN) devices be associated with the same access point, because of the likelihood of collisions and the issue of dropped voice packets not being retransmitted. However, as many as 20 data devices (for example, PCs) could be associated with the same access point, because most data applications can retransmit dropped packets and are more tolerant of latency, as compared to voice.

Be aware that WLAN performance depends on the structure and materials used in a building's construction, which impacts how radio waves are propagated throughout the building. These building characteristics can impact connection speeds and error rates. Fortunately, Cisco's Radio Resource Management (RRM) allows Cisco wireless devices to monitor RF conditions and dynamically make adjustments to access point power and channel configurations to help accommodate for issues such as channel interference and signal coverage.

Specifically, a designer can specify an ***RF group***, which defines a cluster of WLCs that coordinate their RRM calculations. RF groups are created via the following process:

- Access points transmit neighbor messages, which include the access points' WLC IP addresses and hashed message integrity checks (MIC).
- Access points validate each other using the MIC, and an RF group is formed when access points on different WLCs hear validated neighbor messages at a signal strength of –80 dBm or stronger.
- The RF group members or controllers then elect an RF group leader, which is responsible for maintaining a master power and channel scheme for the group.

Cisco access points also support ***self-healing***. With self-healing, a WLC uses RRM to adjust access point power levels, to accommodate for the failure of a neighboring access point.

Designing Wireless Networks with Controllers

When designing a wireless network, one of the first steps in the design process is to conduct an RF site survey. A site survey provides the designer with a better understanding of an environment's RF characteristics (for example, coverage areas and RF interference). Based on the results of the RF site survey, the designer can strategically position the wireless infrastructure devices.

Conducting an RF site survey involves these procedures:

- Determine the number of customer devices to be supported, the required service level, and peak traffic-level requirements.
- Acquire a structural building diagram, which can be used to identify potential RF obstacles.
- Perform an on-site inspection, looking for structural components (for example, metal racks or elevator shafts) that might impair the wireless signal.
- Specify preliminary locations for access points, keeping in mind that the access points need power and access to the wired network.
- Conduct the actual RF site survey, which maps out RF coverage areas. A tool such as the Cisco WCS can import a floor plan and graphically display RF coverage areas and signal strengths. This type of composite graphic is often referred to as a ***heat map***.
- Document the results of the RF site survey. The documentation should include information such as the access point models used, locations of access points, signal strength levels, and bandwidth available at the outer boundaries of the coverage areas.

Many wireless networks also need to support connectivity for guests, without permitting guests full access to network resources. One approach to guest access is to isolate guest traffic on a separate VLAN. However, in large enterprise environments, this approach might not be deemed adequately secure.

Therefore, another option is to use a Layer 2 tunnel to send all guest traffic to a controller dedicated for guest use. This controller is located in a demilitarized zone (DMZ), which uses a firewall to separate the guest network from the organization's internal network.

Wireless network design might also need to address outdoor wireless connectivity (for example, wirelessly interconnecting buildings). Traditionally, buildings were wirelessly interconnected using point-to-point bridging or point-to-multipoint bridging.

A newer approach is wireless mesh networking, as illustrated in Figure 8-3. An outdoor mesh uses multiple access points which interconnect, thus providing numerous redundant connections between nodes. These access points can dynamically discover one another and select an optimal path through the mesh.

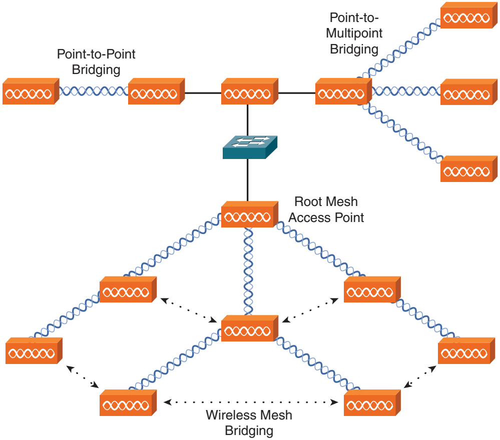

FIGURE 8-3 Outdoor wireless mesh.

The Cisco unified wireless network, which is the basis of the wireless mesh network, is composed of the following elements:

- **Cisco Wireless Control System (WCS)**—WCS provides a graphical user interface (GUI) for networkwide policy configuration.
- **Cisco Wireless LAN Controller (WLC)**—WLCs manage multiple access points, manage wireless network security, and offer Layer 3 mobility features.
- **Rooftop access point (RAP)**—Typically located on a rooftop, the RAP provides wireless connectivity into a wired network.
- **Pole-top mesh access point (MAP)**—A MAP is typically located on a pole, such as a lamp post, and serves as an access point for wireless clients.

Although the connection from a MAP a RAP can support eight hops, Cisco recommends four or fewer hops. Also, be aware that a RAP can connect up to 32 MAPs, but Cisco recommends that a RAP connect no more than 20 to 25 MAPs.

When designing a wireless network for an enterprise campus, a designer should determine the following:

- The number of required access points
- The location of the access points
- The power source for the access points
- The number of required WLCs
- The location of the WLCs

Some of these same design considerations (for example, the number of access points needed) are also relevant for branch office wireless networks. However, branch offices might not be able to justify the expense of separate lightweight access point and WLC devices. One approach for branch offices is to use local MAC, which supports full 802.11 functionality in the access point.

Another option is to point the branch access points back to a centralized controller. If a centralized controller is used, the round-trip time (RTT) between an access point and its controller should not be greater than 200 ms. Also, designs using centralized controllers should implement one of the following technologies:

- **Remote Edge Access Point (REAP)**—REAP extends LWAPP control timers, thus offering more compatibility for branch offices. Although control traffic is still encapsulated using LWAPP and sent to the centralized WLC, data is locally bridged. However, IEEE 802.1Q trunking is not supported by REAP, and REAP requires that all WLANs terminate on a single local VLAN or subnet.
- **Hybrid REAP (H-REAP)**—Unlike REAP, H-REAP allows wireless network administrators to configure and control two or three access points, located in a branch office, over the IP WAN. Also, H-REAP access points have the ability to locally switch data traffic and locally authenticate clients if connectivity to the WLC is lost.