Describing Network Requirements

Through the introduction of two concepts known as the Intelligent Information Network (IIN) and Service-Oriented Network Architecture (SONA), Cisco has made new recommendations for the way networks are designed and implemented based on the particular size and business need to be met.

The IIN concept provides a means of articulating the evolving role of the network in enabling all components of an Information Technology (IT) infrastructure. The network is the common denominator that brings all the pieces together. This new view of the network's role in today's business models provides a means of reclassifying that role from mere transport into a service- and application-oriented role. SONA provides an underlying foundation (or framework) encompassing all technologies, applications, and services, combining them into a single entity focused on becoming an IIN.

In support of this, Cisco has released the Cisco Enterprise Architecture (CEA). The CEA is an enterprise-wide model that allows companies to protect, optimize, and grow their infrastructures as business needs dictate. The CEA provides a comprehensive design and implementation resource for a wide range of service offerings required in the typical network infrastructure today and into the future, including Campus, Data Center, Branch, Teleworker, and WAN architectures.

Intelligent Information Network

The Intelligent Information Network (IIN) offers companies an understanding of how the role of the network is evolving to meet business needs. The IIN vision is essentially the concept of network simplification through the alignment of technology and business priorities. Beyond evolution, the role of the network is expanding as more and more services become available network offerings. Cisco has established four technological roadmaps specific to the individual business needs of its customers. Each of the four roadmaps defines the IIN vision for a particular market segment or business type. These architectures are meant to show businesses how to look forward three to five years in planning network expansion. These four technological roadmaps are as follows:

- Service-Oriented Network Architecture (SONA)
- Service Provider Architecture (IP Next-Generation-Networks or IP-NGN)
- Commercial Architecture
- Consumer Architecture

Together these comprise the foundation of the IIN. The goal of the IIN is to build intelligence across multiple protocols and infrastructure layers to allow the network to be more aware of the needs of its users and respond efficiently to those needs by allocating needed resources and/or applications regardless of the nature of the connected device. The network aligns itself with the business priorities of an organization through services, availability, adaptivity, and resilience. The Cisco vision of the IIN composition includes these features:

- **Network resource and information asset integration into the network**—Includes video, voice, and data integration into the network infrastructure
- **Cross-platform/cross-product intelligence spanning all layers of infrastructure**—Network-wide extension of that intelligence to permit end-to-end connectivity and a common user experience regardless of access device or method
- **A network that actively participates in the delivery of services and applications**—Proactive allocation of network resources as needs demand for a particular application, service, or user

IIN is beyond the traditional concept of basic network connectivity, bandwidth allocation, and access to applications. A true IIN offers end-to-end functionality that adaptively shapes the user experience on-the-fly and promotes true business transparency and agility.

The evolutionary approach of the IIN technology model consists of the following three essential phases. In each phase, the opportunity exists to further augment the applications and services available to meet the business need.

- **Integrated transport phase**—The network is a common pathway for all traffic types. Each traffic type is classified according to the identified business priorities and/or the nature and sensitivity of the traffic to latency, jitter, and other assorted network conditions. This permits the network architect to present a modular functionality that can be customized by organizations or individual departments according to their individual needs. Network convergence also lays the foundation for a new class of IP-enabled applications delivered through Cisco IP Communications solutions.
- **Integrated services phase**—With full network convergence, IT resources can be pooled and personnel can be cross-trained and utilized more efficiently. This remedies the age-old issue of having only one “go-to” person in IT. Each IT staff member becomes a “go-to” person. Diverse resources required by individual organizations and personnel can be virtualized and moved into the network so that a new degree of flexibility can become reality. This flexibility comes into reality by using the network as the platform—a single resource capable of providing common services to all applications. Rather than having hundreds or thousands of mission-specific servers, the network becomes the platform. The servers are moved into the network as virtual services, thereby providing immense savings in hardware, power

consumption, and real estate usage in the data center. Business continuity is also enhanced because shared resources across the IIN provide services in the event of a local systems failure.

- **Integrated applications phase**—The third phase of the IIN evolution is known as Application-Oriented Networking (AON). This is where the plans come to fruition. The network reaches an "application-aware" state that allows it to optimize application performance and more efficiently deliver networked applications to the end-user community. Additional capabilities, such as content caching, load balancing, and application-level security, allow the infrastructure to add intelligence through simplification of the overall network infrastructure.

Of particular interest in this book is the technical roadmap focused on enterprise networks known as SONA. SONA is the framework that provides the evolutionary path for an enterprise network to become an IIN. While the remaining three architectures are critical for their respective market segments, they are beyond the scope of this book. They are mentioned here to illustrate that concepts similar to those discussed here are laid out for service provider (SP), small/medium business (SMB), and small office/home office (SOHO) networks.

SONA

The path of evolution for business services and applications is emerging into a more efficient, flexible, and dynamic model. This is the IIN. The network is the platform. Individual resources can be allocated dynamically, as needed by resource-hungry applications or services. Resources such as CPU, memory, and storage can be added and/or removed on-the-fly and without impact on other processes. Even better, the cost of such a model is reduced through shared resource utilization. No longer are dedicated resources needed for mission-specific applications. Instead, the network maintains resource pools that provide dynamic allocation of resources on demand.

For enterprise networks, SONA provides the architectural framework necessary to build an IIN. SONA leverages the network to allow interactive services to be added to it. This provides the additional benefit of allowing loosely connected services and/or applications to communicate, yet remain independent of each other. This collaborative capability permits provisioning of a new level of service, allowing an enterprise to offer its user community the same network experience, including applications, services, and capabilities, regardless of their location or choice of network-endpoint device.

As previously mentioned, the SONA vision is built around the enterprise network. The architecture itself is further subdivided into layers so that each can be implemented properly to support the next. SONA is the architectural framework that leads enterprise network evolutionary processes, allowing a network to reach the IIN state in order to accelerate applications, business processes, and, most importantly, profitability. Figure 1-1 illustrates the breakdown of the SONA layers.

Figure 1-1 *Cisco SONA*

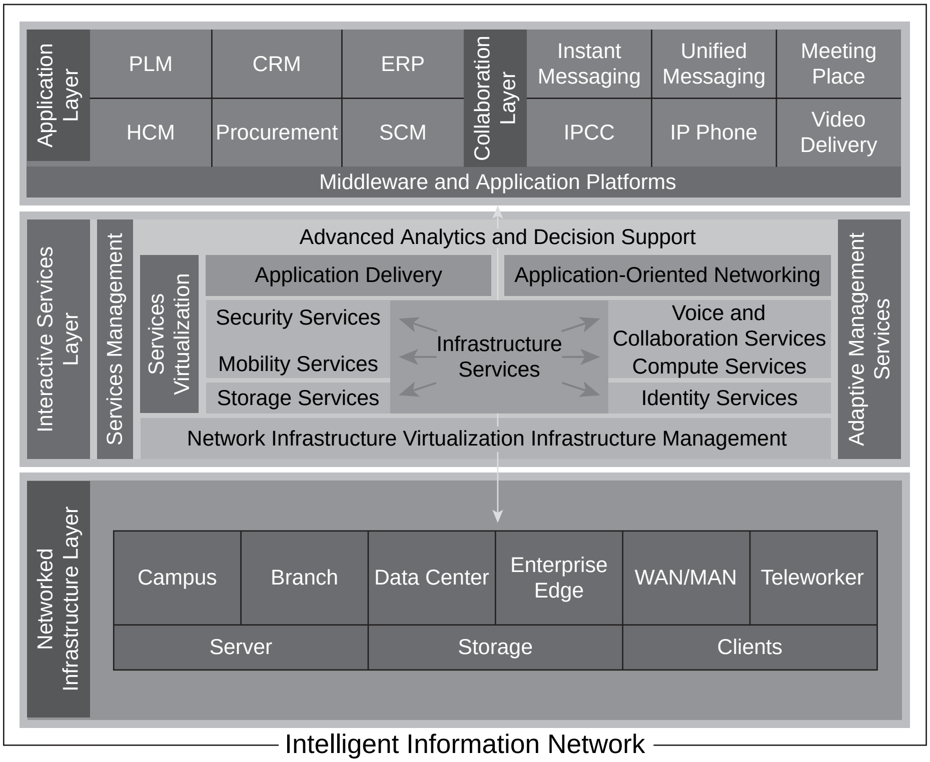

SONA makes extensive use of Cisco product lines and business partners to accomplish its goal of providing secure, flexible, adaptive, and converged network infrastructures. To aid the comprehension and to promote understanding of individual technology roles in the architecture, a layered model was created. Unlike the OSI Model, the SONA layered model consists only of three layers. As shown in Figure 1-1, these are as follows (from the bottom up):

- Networked Infrastructure Layer
- Interactive Services Layer
- Application Layer

Service integration is a key concept in the overall SONA picture. This allows common services to be provided from a single point within the infrastructure. Keeping these services in loosely coupled relationships with other services (for example, web services, XML, and so on) allows a single service or resource to be shared among multiple applications. This simplifies support, reduces maintenance costs, and potentially provides licensing savings on some applications.

Each layer has its form and function in the construction of an IIN. The sections that follow provide a brief discussion of that form and function at each layer.

Networked Infrastructure Layer

The lowest of the three SONA layers provides the point of interconnection between various IT resources. The Networked Infrastructure Layer encompasses servers, storage, and network-connected endpoints. These resources exist in various volumes and geographies throughout the network. The Networked Infrastructure Layer provides the common transport and connectivity between required services such as CPU cycles, storage, memory, and I/O. Rather than using individual, dedicated (or mission-specific) resources, SONA sees these elements simply as resource pools.

The SONA model reaches out across network geographies to pull all resources into a single, logical entity. The architecture includes specifications on the construction of all of these geographies, including the campus, branch, data center, WAN/MAN, and teleworkers. Each is addressed individually in the SONA model as each is crucial to the creation of an IIN capable of providing a common user experience anytime, anywhere and from any device.

As you might expect, TCP/IP becomes the pervasive network protocol and the network provides the shared transport for all business application traffic. This is known as *convergence*. This allows the network infrastructure to become *service ready*, allowing the offloading of application functions away from application resources through service integration.

Interactive Services Layer

A significant cause of inefficiency within an IT organization is the presence of “silos”; that is, application-specific hardware and software that cannot be reused or shared. As more and more businesses begin to rely on collaborative services, the need to more closely align IT resources and computing platforms becomes more crucial.

The Infrastructure Services Layer (ISL) pools these resources in a process known as virtualization. These resources include both the Networked Infrastructure Layer and Infrastructure Services.

The Infrastructure Services Layer sees these as resource pools as well. However, in addition, SONA sees the network infrastructure as simply one more element in a resource pool to be managed and shared.

By virtualizing these resources and defining their use through adaptive management capabilities, the business transformation becomes more dynamic and, more importantly, more simplified. By keeping these resources loosely coupled, they remain modular. That is, they can be added, removed, upgraded, and maintained individually with no impact whatsoever on other resources in the pool.

No longer are individual servers dedicated to mission-specific roles. They become part of a bigger picture and a shared resource. Flexibility is achieved when virtual resources are available on an as-needed basis over a shared infrastructure without having to make any change to the underlying network architecture. As silos are removed and hardware/software investments further leveraged as shared resources, individual components can no longer negatively impact business operations in the event of maintenance, failure, or another service-impacting event.

As resources become part of the larger shared (or virtualized) entity, the lines between the application and the network begin to blur as the network is the transport and is providing access dynamically to needed services and associated resources seamlessly.

One function of the ISL deals specifically with application networking services. Application networking refers to a set of services consisting of network-embedded technologies that improve the deployment of applications in a distributed model without impacting the responsiveness of the application and resulting user experience (as the experience will vary depending on the location of the user versus that of the resource). The goal is to remove location dependency while maintaining comparable functionality.

Breaking the location dependency is possible in the architecture through delivery of high application throughput, reduced latency, encryption, compression, and optimization of communications between client and application resources.

Examples of these resources and services include

- Voice and collaboration services
- Device mobility services
- Security and identity services
- Storage services
- Computer services
- Application networking services
- Network infrastructure virtualization

- Services management
- Adaptive management services
- Advanced analytics services
- Infrastructure management services

The list goes on, but the services identified here should provide some idea of the concept of resource virtualization.

Application Layer

The Application Layer contains the business and collaborative applications that use interactive services to function more efficiently. The interactive services allow the applications to grow dynamically, thus allowing more rapid and efficient deployment while keeping integration costs down. When a new user base, department, or branch site is added, the application can simply be allocated a larger share of the resource pools dynamically to compensate for the increased use.

The Application Layer is most concerned with two application categories:

- **Business applications**—Include those applications that are mission-specific to a business or department and are crucial to that organization's function. For example, a procurement or human resources application would be used only by the respective departmental personnel. Yet, those personnel would require use of the shared resources at all three layers.
- **Collaboration applications**—Include Instant Messaging (IM), Unified Messaging (UM), IP Contact Center (IPCC), IP Phones, and Video delivery. These are the tools that allow people to interact in the manner and time of their own choosing. The use of presence technologies allows an individual to choose the manner in which they wish to be contacted at a given time and on which device that contact should be made. The experience and functionality will be similar (for a given application type) regardless of the access device.

Cisco Network Models

Now that the basic concepts of SONA, the road to the creation of an IIN, are somewhat clearer, some discussion of network models is needed. Network models vary based on the technology being implemented; however, the goal of the models is still the same—convergence and enabling service integration.

As mentioned previously, Cisco has created a visionary architecture for its customer market segments. For the enterprise network, SONA is the architecture. At the Networked Infrastructure Layer exists a rather wide array of technologies and possibilities. These were touched upon briefly in the "Describing Network Requirements" section and are expanded upon in this section.

Typically, six distinct geographies exist in an end-to-end network architecture. These are contained within the Networked Infrastructure Layer of SONA. Refer to Figure 1-1 for an illustrated view.

- **Campus network**—Provides network access to campus-wide resources
- **Branch network**—Provides network access to remote resources
- **Data Center**—Provides access to and interconnectivity between servers and storage resources
- **Enterprise Edge**—Provides secure access to and from public and partner networks
- **WAN/MAN**—Provides connectivity between branch offices, campuses, and/or data centers
- **Teleworker**—Provides connectivity to the corporate network for home-based employees

As is readily apparent, all of these are somewhat interdependent, yet very different in terms of resource and architectural needs.

Cisco Hierarchical Network Model

Prior to any discussion of the architecture models proposed in the IIN vision, it is necessary to step back to a discussion of a somewhat older model advocated for network scalability, the Cisco Hierarchical Network Model. Figure 1-2 illustrates the model for purposes of discussion.

Figure 1-2 *Cisco Hierarchical Network Model*

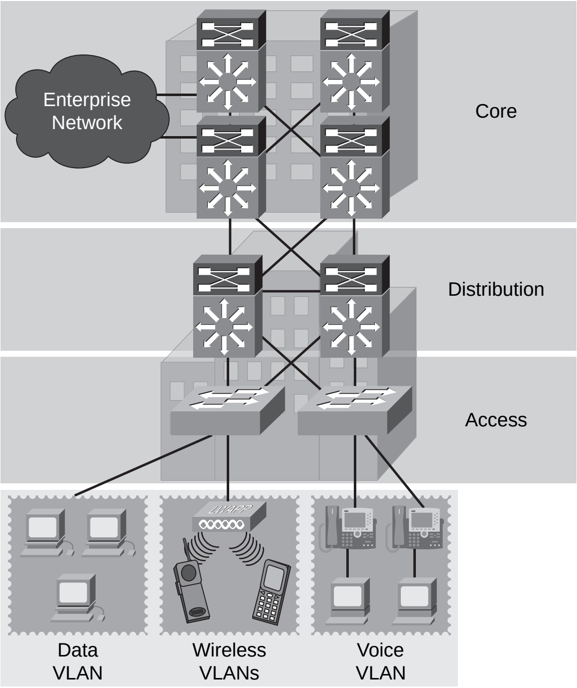

As is evident in the figure, the essential layers of the network are divided into three layers: Core, Distribution, and Access. This provides a repeatable, or “cookie-cutter,” model that is easily reproduced site to site. The model also has the benefit of being scalable from hundreds to thousands of devices in a campus network. Additionally, this model supports the integration of SONA Interactive Services Layer applications and services, facilitating an improved experience in the interaction between the clients and applications/services provided by the network.

Each layer has its prescribed function, as described here:

- **Access Layer**—Devices deployed throughout the network with the express purpose of providing user access to the network, generally through switch port access. Access layer switches are generally located near the user populous they serve.
- **Distribution Layer—**Devices deployed as aggregation points for Access layer devices. Distribution layer devices can be used to segment workgroups or departments in a campus environment. The Distribution layer devices also provide for WAN aggregation connectivity at the Campus Edge and provide policy-based connectivity.
- **Core Layer (a.k.a. Backbone Layer)**—Devices that carry the weight of the network. They are designed to switch packets as fast as possible. The Core layer must be highly available and redundant to ensure that no loss or degradation of service is experienced in the event of a network outage.

This model can be applied to any network of any size regardless of the technologies and connectivity options it presents. This includes LAN, WAN, MAN, wireless, VPN, and other networks. In smaller networks, it is feasible that one or more of these layers might be combined into a multi-functional layer. In the discussions to follow, and throughout nearly any networking technology-related book, these three layers are referenced quite frequently.

Campus Network Architecture

Campus network architecture has evolved rapidly over the last decade or more. The number of services supported in a campus environment has evolved just as quickly, if not more so. The basic infrastructure has traditionally been summed up under the Cisco Hierarchical Network Model mentioned in the previous section.

This remains the case because that model scales very well. The role has expanded somewhat on its own to include technologies such as quality of service (QoS), Multiprotocol Label Switching Virtual Private Networks (MPLS VPN), IPsec VPN, Hot Standby Router Protocol (HSRP), and more. Shifting topological ideology has seen a dramatic increase in the number of enterprise networks shifting from traditional Layer 2 switching to Layer 3 switching at the Access and Distribution layers. The campus network architecture is meant to provide enterprise corporate headquarters sites (which might mean a single building or multiple buildings in a common geography) with a means of consolidating and simplifying network support and administration while increasing service and application offerings to the user community. Figure 1-3 illustrates the campus network architecture.

Figure 1-3 *Campus Network Architecture*

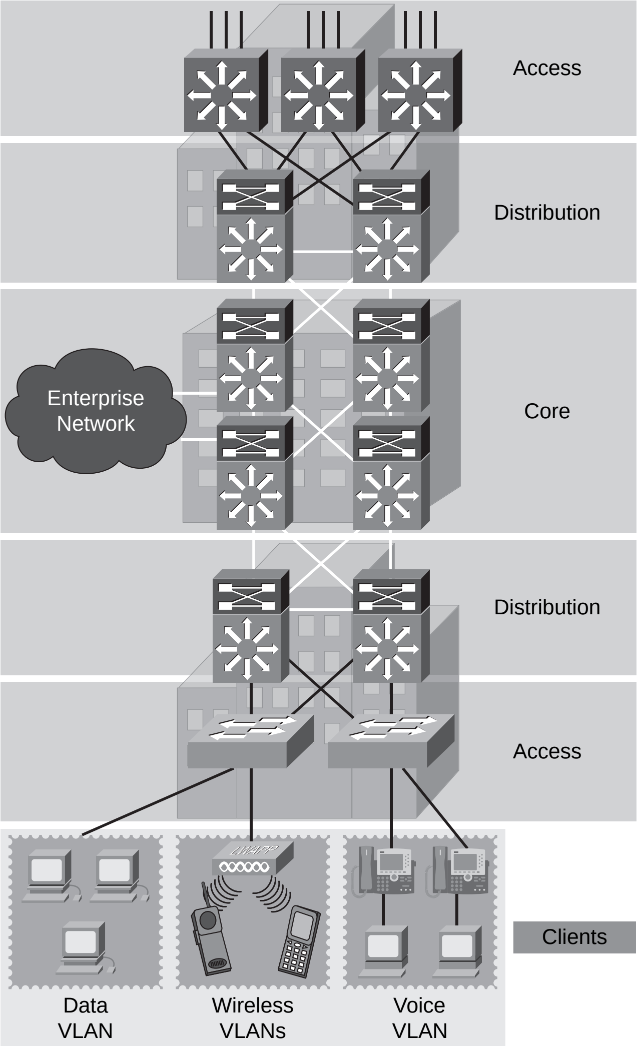

Campus services are changing in nature from traditional stateless (connection and/or session unaware and packet switching) to stateful services requiring highly available, redundant devices to track sessions and connections at all times. Meeting this need requires changes in the basic networking paradigm.

For example, a wireless client roaming throughout a wireless-enabled campus using both a laptop and a Cisco 7920 802.11b IP Phone would need to have the ability to seamlessly home from access point to access point with no interruption in service. This is especially true of voice calls in progress. The presence of voice and data together would necessitate a fully QoS-enabled network for both wired and wireless connected devices. An access point would be required to be able to exchange session state information and user/device credentials with no user interaction.

This poses only a single example of the need for campus evolution to accommodate dynamic needs of users and devices. Also, as the evolution progresses and more devices become network-dependent, the need to eliminate any/all single points of failure becomes more and more critical as a factor to success.

Branch Network Architecture

Branch network architecture provides an integrated, multiservice environment that provides connectivity to its users who are working in remote or satellite offices rather than in the primary corporate headquarters or campus. This provision requires both hardware- and software-specific integration considerations to provide the applications and services required for the users to properly perform their job functions.

The SONA branch architecture allows an enterprise to extend campus-like services and applications to the remote branch site while maintaining proper service levels and responsiveness. Advanced services such as Cisco Unified Communications, security, and more can be offered at branch sites, a traditionally unavailable option for some services due to inadequate connectivity and reachability. Figure 1-4 illustrates the branch network architecture.

Figure 1-4 *Branch Network Architecture*

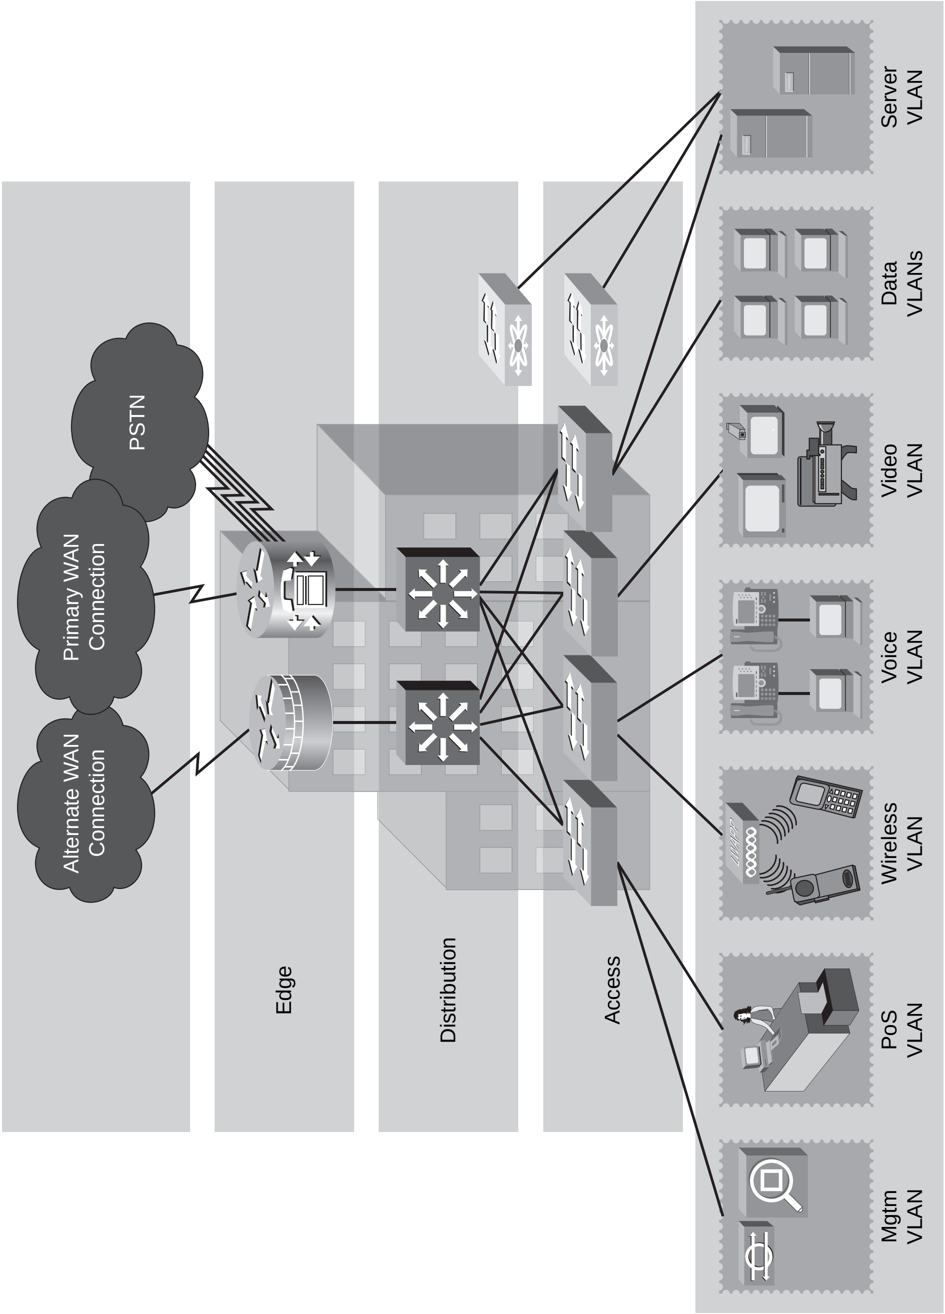

Cisco integrates security, switching, network analysis, content caching, and converged voice/video services into a series of products known as integrated services routers (ISRs). Aside from filling the role of traditional router platforms, these devices are multitalented and designed for performance. ISRs are affectionately known as "branch-in-a-box" routers. They very effectively stand up to that name. Key individual services, such as voice and security, are built into the motherboard as standard components. Additional modules provide voice messaging and content services that easily install into the same chassis.

This solution provides for local access to key services and/or applications without reliance on a WAN link or other connectivity. Call control, security, VPN connectivity, and more are simply built into the same router chassis. Configuration is also simplified through the addition of the Security Device Manager (SDM). The SDM provides an intuitive graphical, web-based interface that can be used to configure firewall services, routing, VPN, and more.

Data Center Architecture

The data center is a key point in the evolution of the network. It is rapidly evolving to take in more and more service-oriented functions. The move toward a dynamic, demand-based service offering dictates that the network be aware of server and application health at all times. This health information is then used to take appropriate action, making incremental increases in service resources that are available to a particular application or service. This can be the addition of virtual servers, application instances, or dynamic network configuration changes needed to bring newly added resources online to support increasing needs.

Resources can be provisioned for server OS needs, Layer 2 functionality (for example, switch port mode, VLAN assignment), Layer 3 functionality (for example, routing, HSRP, multicast, and so on), and services for Layers 4 through 7. Figure 1-5 illustrates the data center architecture.

Figure 1-5 *Data Center Architecture*

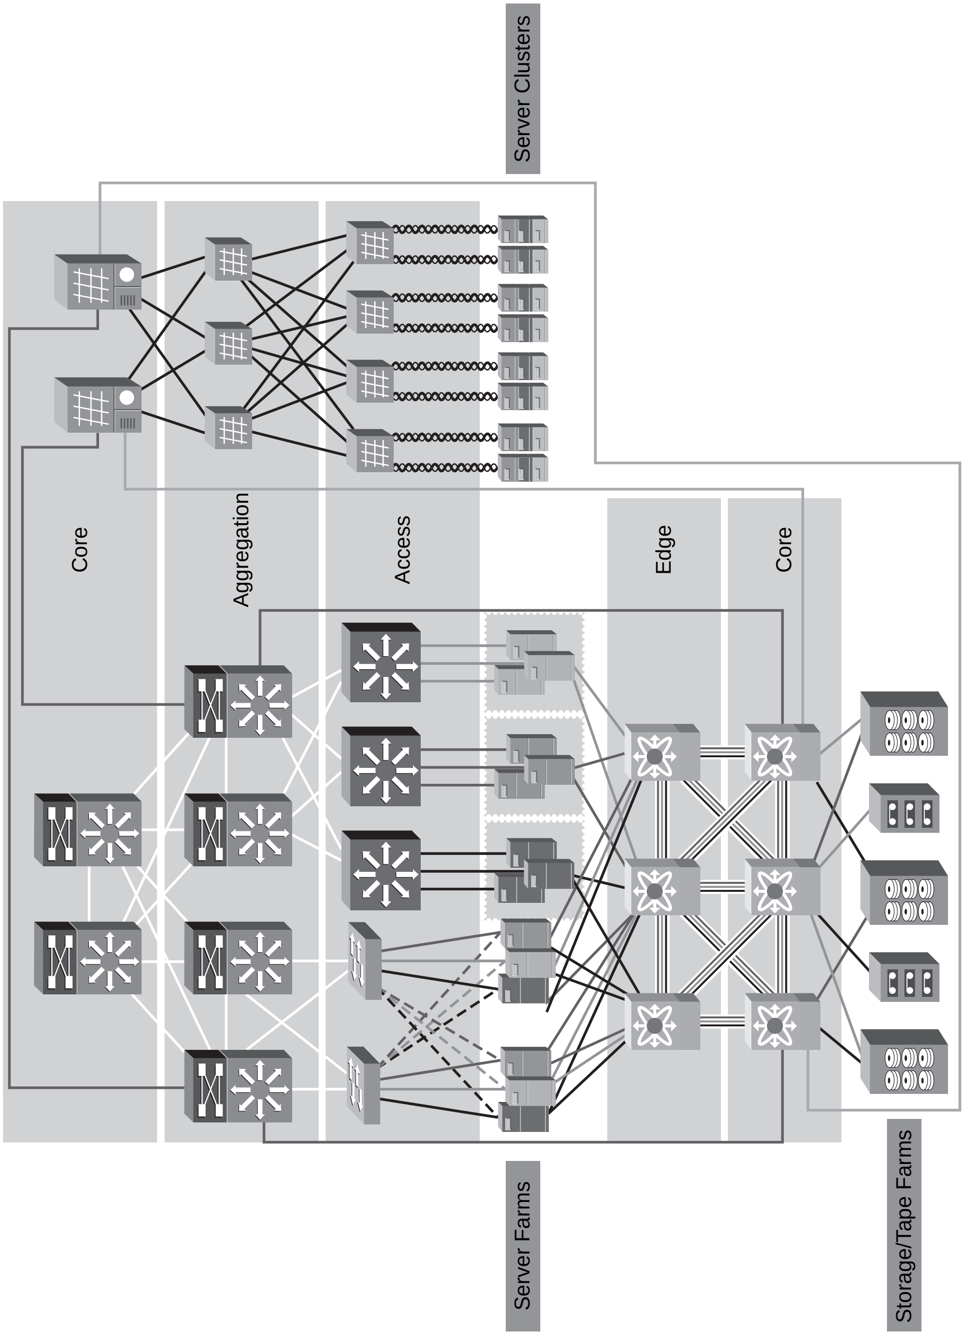

This architecture provides a cohesive, adaptive network that allows for consolidation of resources while increasing availability and business continuance. Enabling service-oriented architectures, virtualization, and on-demand services to provide a dynamic network environment for all users in all locations leads to streamlined management and reporting and more effective use of capital. This solution allows the network to scale to a significant degree without infrastructure changes that would traditionally be needed to support a diverse and varied user base.

Enterprise Edge Architecture

The enterprise edge is evolving with the need to provide more and higher-level security features as a first line of defense for the network. This is true of both internal- and external-facing server farms and services. Figure 1-6 illustrates the enterprise edge architecture.

Figure 1-6 *Enterprise Edge Architecture*

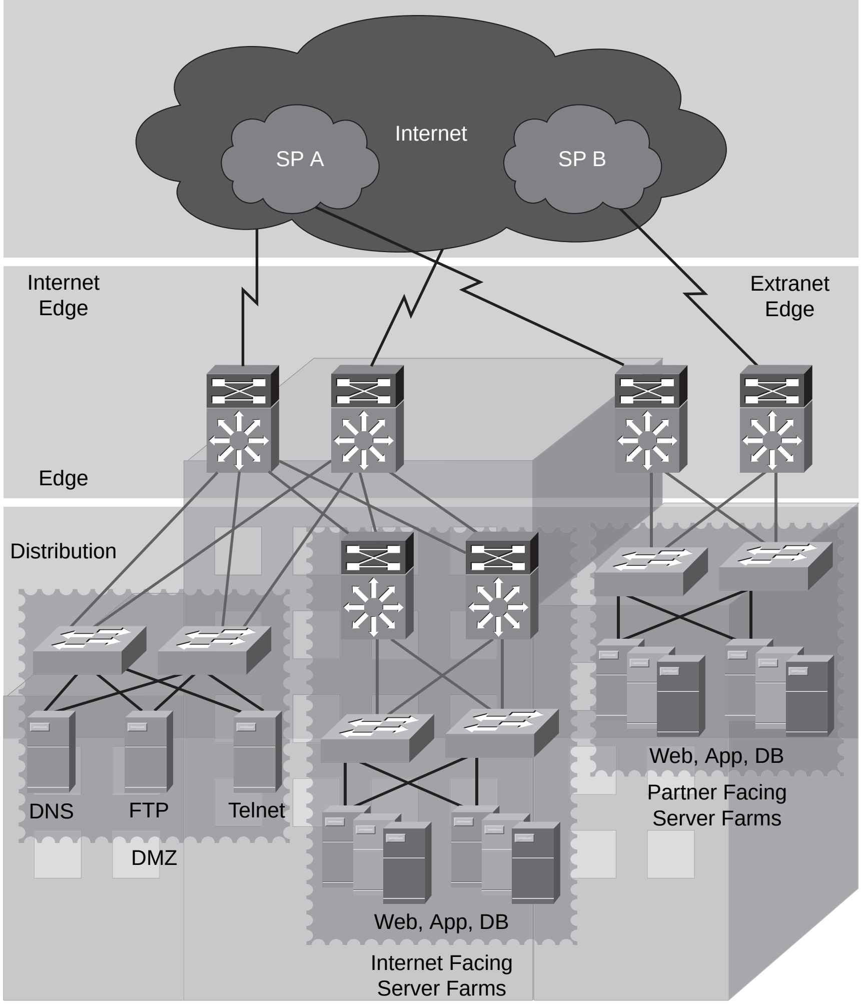

A number of server farms may be supported, each varying in function from demilitarized zone (DMZ) functions for internal or external users (DNS, FTP, web, Telnet, and so on) to Internet services or partner-access servers hosting applications shared with business partners and their employees.

Teleworker Architecture

Increasingly, due to space, real estate, employee accommodation, workforce diversification, and other factors, the population of the home-based workforce is increasing at an exceedingly high rate. Call center remote agents with access to features and functionality identical to their in-office counterparts are taking customer calls from home offices. Salespeople are making deals and booking them via VPN connections back to the corporate site. Most of these workers are using IP telephony to place their office desk phone on their home desk. Figure 1-7 illustrates the teleworker architecture.

Figure 1-7 *Teleworker Architecture*

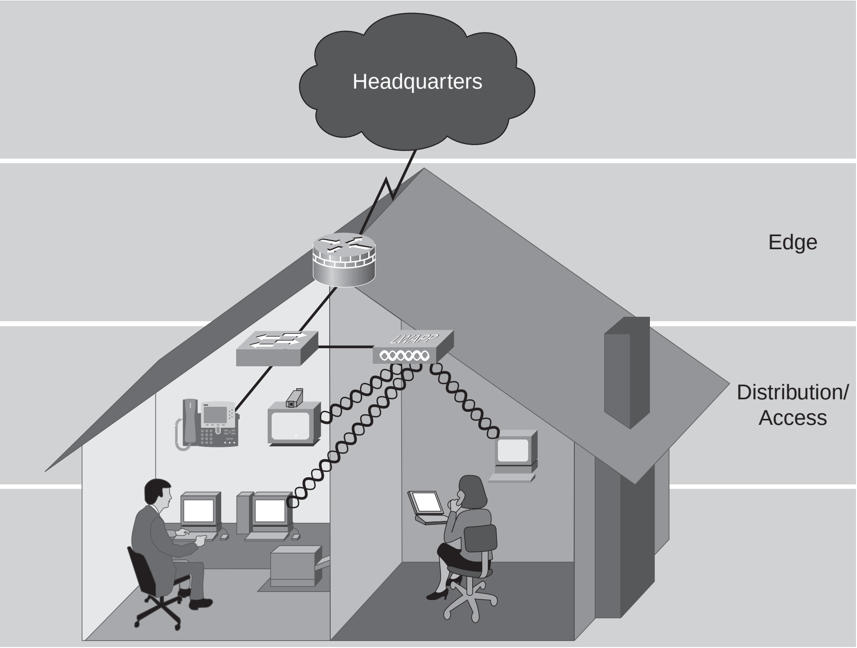

These and many other examples are out there in the world. Cisco is a very big proponent of the enterprise teleworker model and using an ISR platform to provide all the comforts, and access, of physically being in the office.

This architecture dictates the delivery of secure voice and data services to remote small or home office sites over standard, widely available broadband connections (cable, DSL, fiber optic services [FiOS], and so on). This allows for centralized management of devices and standardized application and service availability identical to that of campus-based employees. This includes "always-on" connectivity, security, and, in most cases, wireless connectivity and audio/video conferencing capabilities.

All these applications and services must be provided over "nailed-up" (always on) VPN links that are QoS enabled for the various traffic types used throughout the architecture.

WAN/MAN Architecture

With all the discussion of service-enabled networking, convergence, QoS, and more, the focus tends to be somewhat removed from an equally crucial component of the bigger picture. The design and construction of the wide-area network (WAN) and (where utilized) metropolitan-area network (MAN) can make or break the overall architectural vision.

The transport services necessary for end-to-end connectivity as viewed from the SONA perspective are somewhat different from the traditional view of "just enough bandwidth to make it function properly and no more." Equally dangerous to the vision is the outdated (and far from true) assertion that QoS can be avoided by provisioning "a big, fat pipe." Figure 1-8 illustrates the WAN/MAN architecture.

Geography and function play large roles in deciding the method and speed of connectivity between various sites. Figure 1-8 shows the various possibilities for connecting Site A to Site B.

Whether the connection is traditional Frame Relay WAN connectivity or provided via a service provider MPLS network providing full Layer 3 connectivity from end to end, the needs of the business and the costs involved have a great deal to do with connectivity selection. If sites are very close in relation to each other; for example, in adjacent or nearby buildings, a Metro Ethernet connection might be feasible. Where sites are large and business-critical, requiring high availability, a Synchronous Optical Network (SONET) ring might be the chosen connection type. Whatever the needs of the business and the users at a given site, a means of connecting those sites is available.

Figure 1-8 *WAN/MAN Architecture*

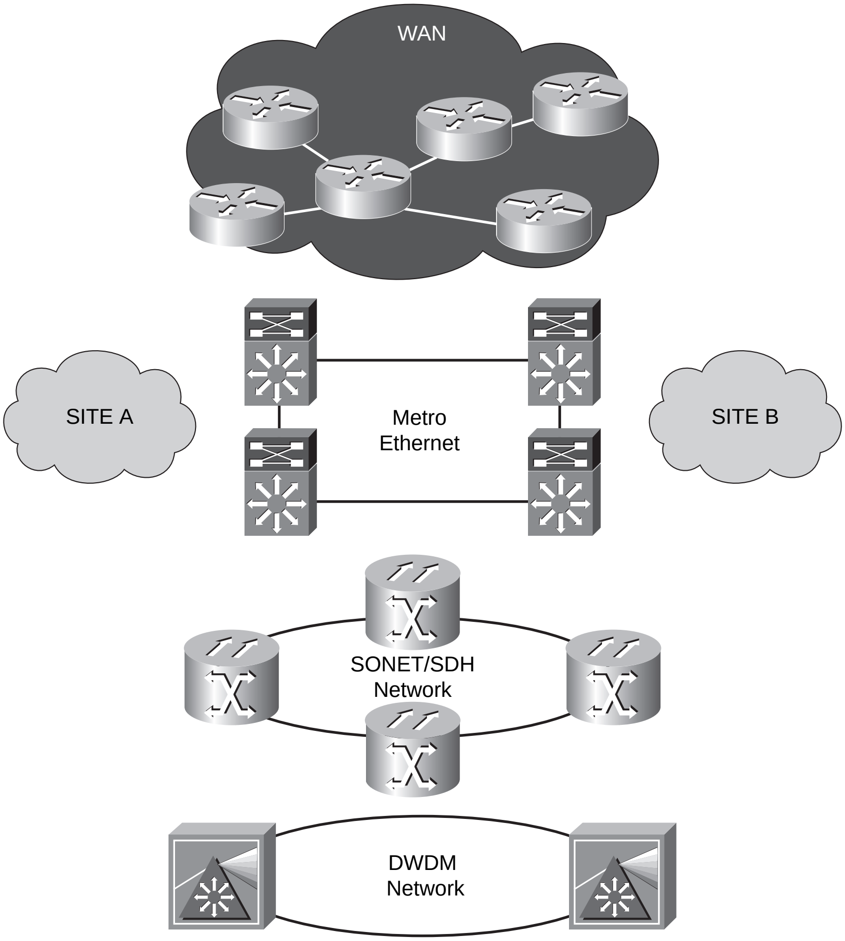

The convergence of voice, video, and data over a single IP network requires a significant degree of forethought and consideration to properly provide services over potentially large geographical areas. QoS, granular service levels, and security factor into the equation as well, to provide secure delivery of various supported traffic types. Emerging trends have seen the deployment of WAN/MAN environments to provide path isolation for traffic between clients and their destination devices, which is a requirement of traffic segmentation over shared infrastructure. Technologies supporting such deployments include MPLS, generic routing encapsulation (GRE), Virtual Routing and Forwarding (VRF), and IPsec.

Remote Connection Requirements in a Converged Network

In the process of evaluating factors and details necessary to effectively design and deploy a central site, branch office, or SOHO site, the most basic requirement is that the site must work effectively for the personnel who staff it. While that factor should be rather obvious and up-front, that is not always the case. Poor site selection can make or break a business, depending on the type of business and needs of that business.

Central Site

A central site must be capable of providing needed services and applications to its user community. Many of these services need to be scalable and flexible, as discussed in the SONA portion of this chapter. Typically, the central site is the largest site in terms of size and population. It could be a corporate headquarters site or a dedicated IT site for larger enterprise networks.

Because all users will access resources at the central site, it is crucial that proper network management practices be in place. This includes planning, design, implementation, and change control practices, to name a few. This site will also accommodate the hub of WAN connectivity, providing access to other sites, branch offices, and teleworkers. Regardless of the geographical disposition of the user population, the network should be designed to provide a consistent user experience across all sites and platforms.

Branch Office

Branch offices vary in size and purpose according to the business needs. A decision must be made about how the branch office network will be designed and what services will be provided locally versus what services will be offered via a WAN connection to the central site. Providing applications and services from the central site is typically most effective in providing a consistent experience for the users. More importantly, the central site is then well positioned to leverage a more complete business picture based on real-time information gathered from those centrally housed applications and services. The process of gathering and processing information from multiple branch office networks with locally provided applications and services can be time consuming and inefficient.

Branch offices can benefit from high-speed WAN links to the central site as well as to the Internet. When branch sites have their own locally provisioned Internet connectivity they also need locally provisioned security resources such as firewalls and content engines.

In cases where sites have only an Internet connection and no dedicated WAN connection to the central site, VPN connections can be “nailed-up” via an Internet connection to ensure a more secure connection back to the central site.

Branch offices of significant size can provide local service and applications to local teleworkers needing access to company resources from a home or satellite branch site. QoS is a concern at all points in the architecture, especially if voice and video services are being provided from the central site to remote employees.

SOHO Site

SOHO sites typically are single-user sites but may include several employees. In any event, these are the smallest sites. A smaller size does not equate to a smaller need for access to applications and services. Although providing those services from a central or branch office site to the SOHO site might be more challenging, doing so is still a crucial factor in ensuring business success.

SOHO sites will likely access resources at multiple other sites including branch offices and the central site. This presents some challenges in figuring out just how the SOHO sites will access all of these resources independently and simultaneously. Here again is the argument for centralized or virtualized applications and resources for all sites being based and hosted from the central site. The need to access resources at multiple branch offices is eliminated.

SOHO site users typically require VPN connectivity back to the central site. This access may be accomplished through a VPN client installed on a company-provided laptop or via small VPN-capable router placed at the user's home. The connectivity back to the central site will vary based on the local service provider offerings available in the user's home area. The connectivity options are relatively wide-ranging and include DSL, cable modem, satellite, and other technologies. A small router (for example, Cisco 871) will make a permanent VPN connection back to a VPN aggregator at the central site to provide access to needed services and applications. This provides the needed security as well as connectivity over which to pass voice, video, and data traffic.

Integrated Services for Secure Remote Access

The cost of providing voice and data services to all users who require them has traditionally been exceedingly high. This has made the business case for opening branch offices a rather difficult one to make. The office required a small PBX or key system to provide telephony and a router to provide data connectivity. This often required two separate departments to maintain services at a single branch office. Add to that equation the need for a third department for support and maintenance of user PCs and laptops and things could get quickly out of hand.

This is no longer the case. With SONA, the applications and services, including voice, data, and essential PC maintenance needed to support users at all sites, are built into the single platform that is the network. No longer is a PBX or key system needed at each branch site. The Cisco ISR platforms provide fallback call control when a centralized call-control model is in use. Alternately, the Cisco ISR platform can provide primary call control on a site-by-site basis. No changes are needed in hardware or software on the router to affect the change between the centralized and distributed call-control models.

Along with voice capabilities, the ISR can also provide native security functions such as VPN connectivity to the central site and firewall capabilities for the local site should they have a local Internet connection. Figure 1-4, in the “Branch Network Architecture” section, illustrates a good reference point for such a deployment. The single Cisco ISR provides a single point of administration for LAN/WAN, PSTN, call control, and security services provided to the branch. The model is most typically used to provide virtualized services at the central site with failover capabilities for each service at the branch site should the WAN connection(s) become unavailable. This provides a significant step forward over traditional telephony because there is typically no redundancy built into smaller PBXs and/or key systems.

A remote office might make use of local broadband connectivity for both Internet and VPN access back to a central site resource pool. Bandwidth, as always, is a primary consideration. With SOHO users, the residential broadband solutions include technologies such as DSL, cable modem, satellite, and fiber optic solutions such as that recently made available to residential customers by local service providers. All of these solutions are relatively affordable and easily installed.

For office sites, the solutions are not always so well laid out. Business-class DSL, traditional Frame Relay, or, in more modern terms and in line with SONA, MPLS VPN connectivity have all become viable solutions for home and office. Many cities, corporations, and even housing subdivisions have begun to offer metropolitan-area wireless connectivity to their tenants/residents. Public reaction and marketing viability will certainly dictate the course of this type of network in the next few years. MPLS will be discussed in more detail in Chapter 8, “The MPLS Conceptual Model.” For now, suffice to say that the carrier MPLS networks are Layer 3 end-to-end networks and can be QoS enabled for varied traffic types, unlike traditional WAN connectivity technologies.

Foundation Summary

The next three to five years will see a significant change in the way organizations view the network as an entity. Currently, many of the services and applications providing businesses with the means to function in their respective industries reside on dedicated hardware platforms, using dedicated resources. As these businesses grow, so does the resource demand on these dedicated platforms. Eventually, the demand outpaces the platform's ability to keep up with the needs of the business. The cycle is reset and the process repeated. This evolution to obsolescence is inefficient and needlessly costly.

SONA details various architectures common in enterprise networks, including campus, data center, enterprise edge, branch, and teleworker architectures. These individual architectures allow the IT personnel to lay out a modular path for each of the various deployments common in today's networks.

Cisco has provided a new routing platform in support of the SONA vision of integrated applications and services. The ISR line of routers is specifically positioned to provide capabilities needed in edge, branch, and teleworker architectures to match those offerings present in campus and data center architectures. The ISRs provide local call control, call-control fallback for centralized call-control models, content caching, and security and VPN capabilities, among others.

Q&A

The questions and scenarios in this book are designed to be challenging and to make sure that you know the answer. Rather than allowing you to derive the answers from clues hidden inside the questions themselves, the questions challenge your understanding and recall of the subject.

Hopefully, mastering these questions will help you limit the number of exam questions on which you narrow your choices to two options, and then guess.

You can find the answers to these questions in Appendix A. For more practice with exam-like question formats, use the exam engine on the CD-ROM.

1. In the SONA model, collaboration services are offered in which layer?
2. Corporate resources are often allocated and deployed in silo-like models. While these resources are dedicated to a department or group within the company, are there any resources they might have in common?
3. Which architecture would typically be associated with a remote user based in a residential office integration solution?
4. List the architectures addressed at the SONA networked infrastructure layer.
5. A branch site housing 50 users needs to access services and applications housed in the central site data center. Consider a solution that would allow these services and applications to be provided to duplicate the experience of central site users accessing the same resources.
6. List at least five services provided at the SONA integrated services layer.
7. Virtualization of resources for dynamic allocation provides a compelling business case in support of a SONA model. Which types of resources can be virtualized?
8. What is the difference between SONA and IIN?

Exam Topic List

This chapter covers the following topics that you need to master for the CCNP ISCW exam:

- **Facilitating Remote Connections**—Describes how to facilitate remote connections that an enterprise network has to support

- **Challenges of Connecting Teleworkers**—Describes the challenges faced in connecting teleworkers to the enterprise network, and the solutions that exist to address these challenges