Architecture for Voice, Video and Integrated Data

Overview

This paper discusses Cisco AVVID (Architecture for Voice, Video and Integrated Data). Cisco AVVID brings to multiservice networking a standards-based, open-systems architecture for converged networking. Cisco AVVID is the continuing evolution of the five-phase plan depicted in Figure 1 for enterprise multiservice networking that has successfully delivered the framework for an open multiservice architecture.

Figure 1 Cisco Systems Five-Phase Multiservice Strategy

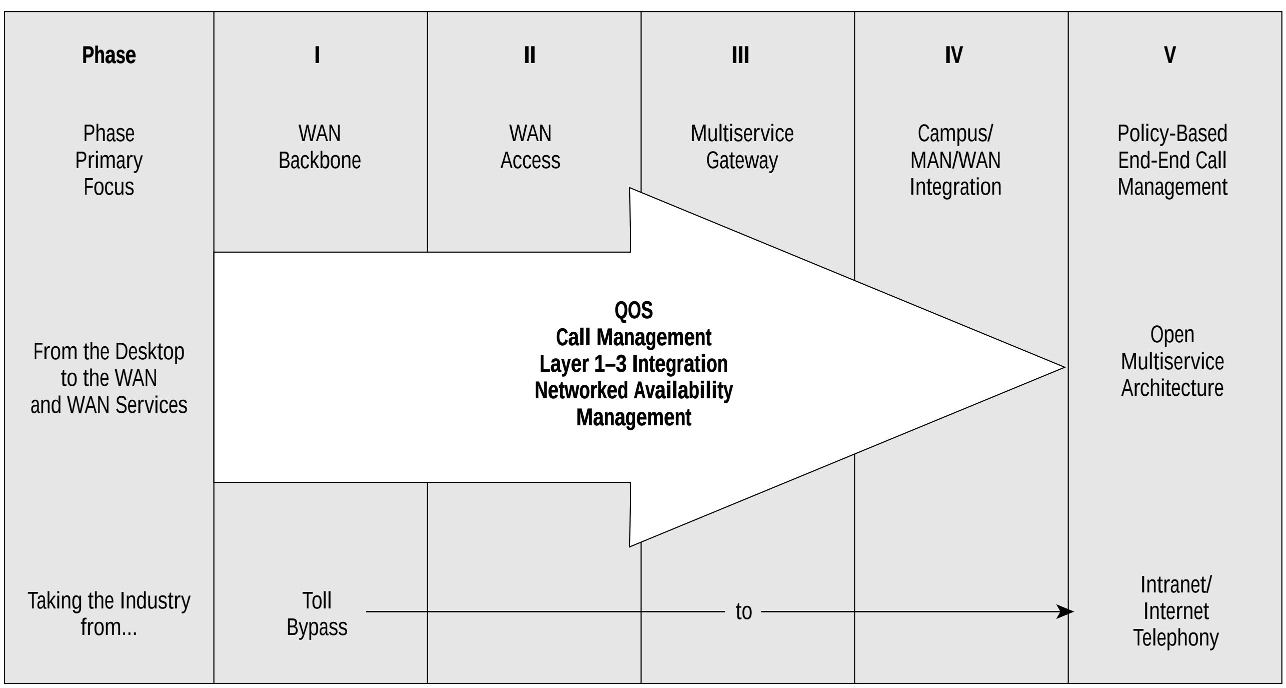

Now that the five-phase strategy has successfully delivered an open multiservice architecture, Cisco AVVID springboards enterprise multiservice networking into the next millennium. Not content with simply overlaying existing legacy communications systems on a common IP infrastructure, Cisco AVVID eliminates in many cases the need for legacy systems entirely. In addition, the rapid introduction and adoption of new and innovative applications such as unified messaging, Cisco IP Contact Centers, end-to-end IP telephony and video are now possible over a quality-of-service (QoS) enabled IP infrastructure.

The architecture comprises three distinct building blocks: infrastructure such as switches and routers, applications such as call control, and clients such as fixed and wireless IP telephones, H.323 videoconferencing equipment, and PCs. Each of these building blocks will be discussed in more detail in a subsequent section. The end result of such an architectural model is a multiservice ecosystem that is scalable, highly available and resilient, open and adaptable.

Cisco AVVID is complementary and synergistic with the Open Packet Telephony (OPT) initiative. Where Cisco OPT focuses on service providers and the benefits of converged data and voice over a common packet transport, Cisco AVVID is an enterprise initiative for integrated data, voice, and video over a common IP transport.

As new applications and technologies are brought to market, additional benefits will emerge and those networks best positioned to leverage and adopt these technologies rapidly will have a distinct competitive advantage. This paper will focus on the following areas.

- Converged networking trends
- Introducing Cisco AVVID
- Enterprise network architecture
- Enterprise network migration strategies

Converged Networking

Empirically separate networks have been provisioned within the enterprise for data, voice, and video applications. These have been deployed autonomously and operated in isolation, often implemented and managed by separate teams.

These separate networks encompass the enterprise local and wide-area networks, and have been built to interconnect private branch exchange (PBX) equipment, H.320 videoconferencing equipment and routers. The networks have been provisioned over dedicated leased lines for PBX and H.320 video, with a combination of leased lines, Frame-Relay, and ATM for data. Figure 2 below depicts a typical deployment of these disparate networks.

Figure 2 Separate Data, Voice, and Video Networking

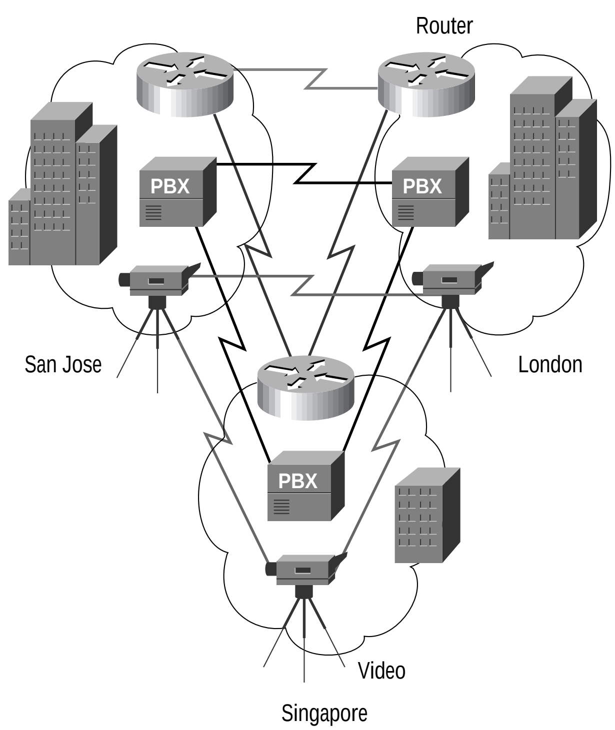

This use of disparate facilities for each application transport is extremely inefficient. The volume of data traffic is growing faster than that of voice, driven by emerging and evolving technological innovations such as the World Wide Web (WWW), e-commerce, and applications such as videoconferencing or video streaming utilizing IP multicast. While growth rates vary by country and carrier, it is certain that data transport will dominate telephony networks. In the United States, data traffic will surpass voice traffic by the year 2000 (DataQuest, October, 1998). Data has already surpassed voice on some U.S. service provider networks. It is the driving force behind global network growth. The challenge for the enterprise is to optimize networking to carry data, voice, and video traffic.

It is widely accepted and acknowledged by the communications industry and industry analysts as a whole, that the Internet Protocol (IP) will become the universal transport of the future. The rapid adoption and migration of vendors to the utilization of IP as a transport for data, voice, and video applications further endorses this transition to a converged networking paradigm. This includes those vendors who have historically used time-division multiplexing (TDM) infrastructures and relied upon “old world” practices. The message is clear: move towards IP or risk being left behind.

The converged network is shown in Figure 3. Utilizing IP as the ubiquitous transport offers the enterprise significant statistical gains in bandwidth efficiency, lower overall bandwidth requirements, ease of management, and the ability to deploy new applications rapidly. On the LAN, data, voice, and, video share a common infrastructure.

Figure 3 The Converged Enterprise Network for Voice, Video and Integrated Data\xfa˙

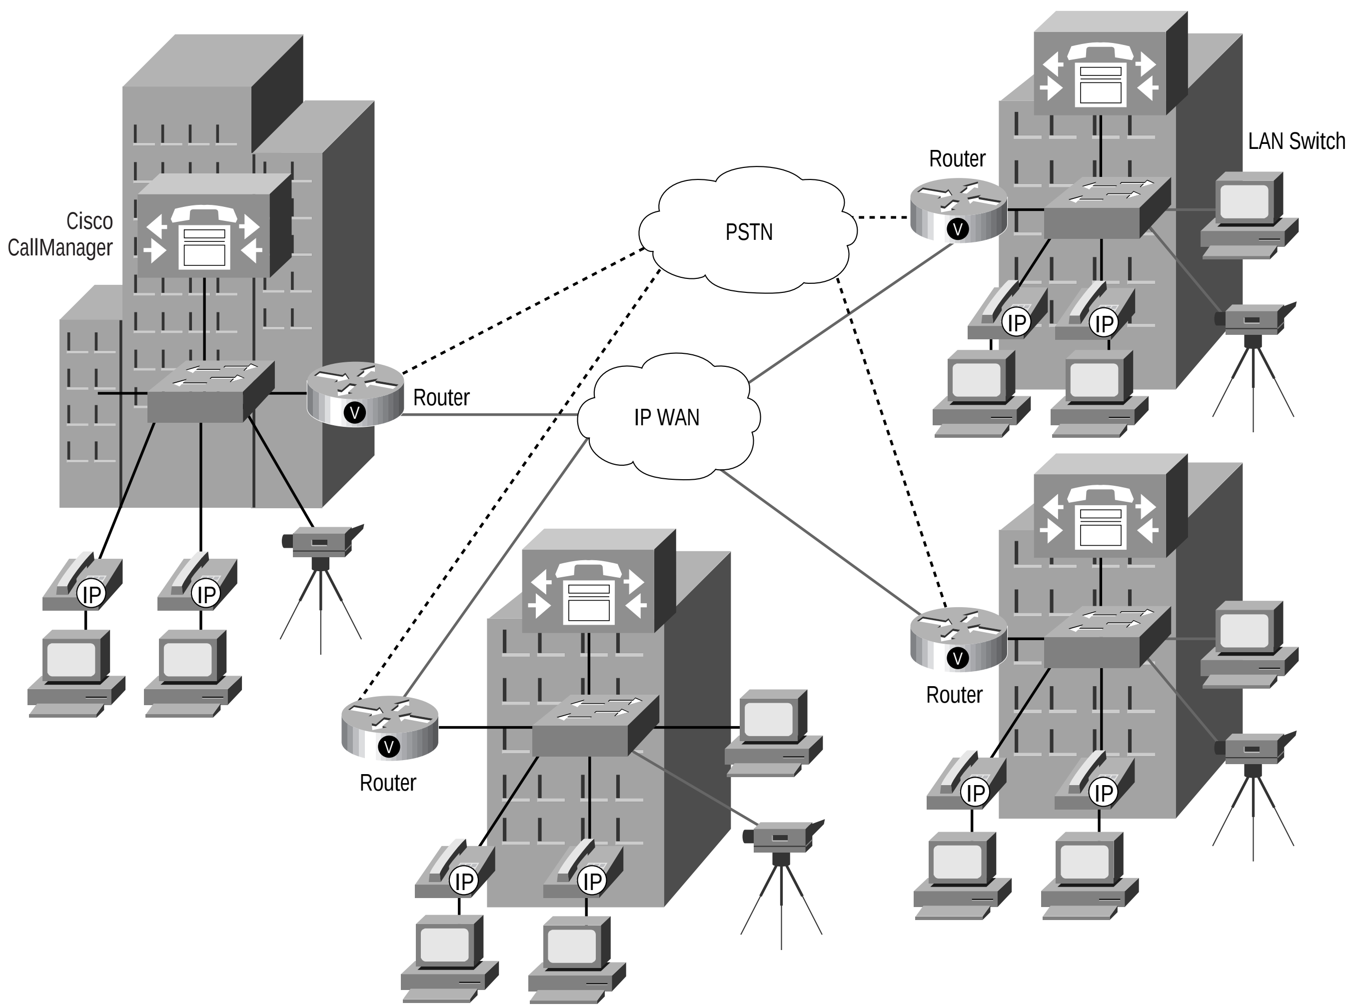

As can be seen above, Cisco AVVID allows the enterprise network to converge over a common IP transport. The number of WAN facilities is reduced, as is the number of devices required to terminate those facilities. Bandwidth can be added incrementally and shared statistically between applications adding efficiency and reducing complexity. This does not present the use-it or lose-it model typical of disparate networks. When voice is quiescent, data can utilize the available bandwidth; when voice or video applications are active, they can be guaranteed the bandwidth required.

Converged networks are a continuing trend and this consolidation of data, voice, and video is the natural evolution for multiservice networking. We have seen similar evolutions before. One of the most poignant being the migration to IP as the transport of choice for Systems Network Architecture (SNA) traffic, a proprietary networking architecture introduced by IBM in 1974. Where once a separate network infrastructure was required, SNA is typically transported over a common IP network.

Introducing Cisco Architecture for Voice, Video and Integrated Data

The following sections will discuss the features and benefits of Cisco AVVID (Architecture for Voice, Video and Integrated Data). The architecture delivers an Internet ecosystem, which thrives on open standards, encouraging the development and interoperability of multivendor, multiproduct solutions. Contrast this to the proprietary and monolithic solutions of the past; the benefits will transcend simple economics, creating innovative new business solutions and rapid deployment of these solutions.

An End-to-End Architecture

Cisco AVVID is an end-to-end architecture that includes three distinct components. These are infrastructure, applications, and clients. Figure 4 depicts the components of the architecture.

Figure 4 Cisco AVVID—An End-to-End Architectural Model

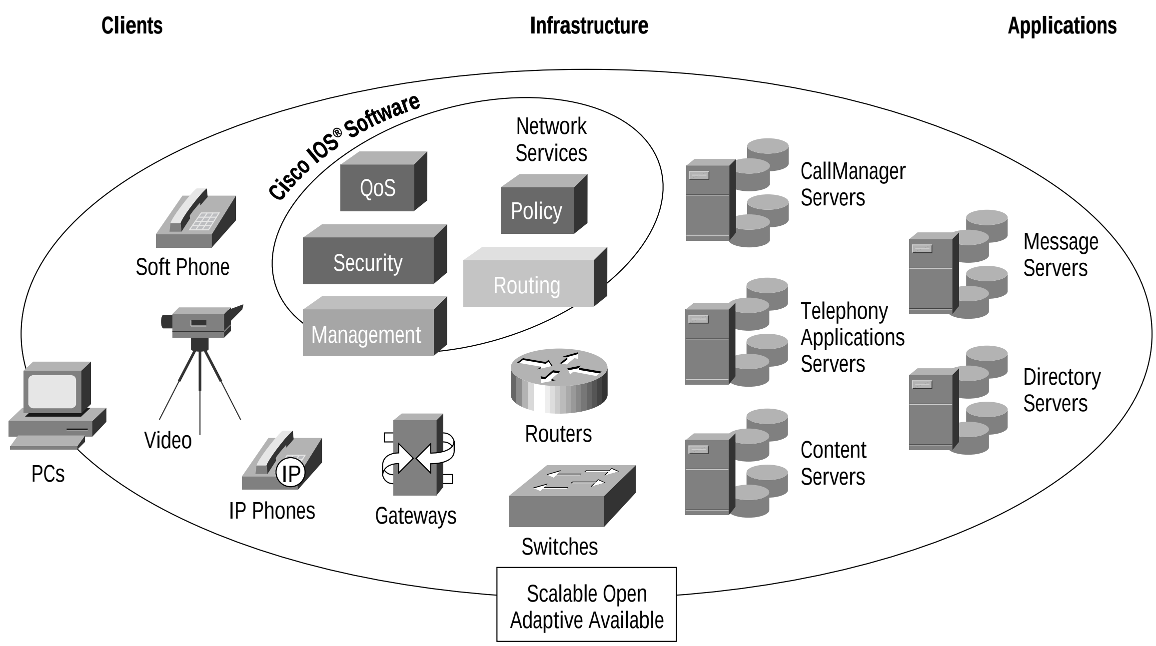

Infrastructure

As with any architecture, Cisco AVVID relies upon a strong and stable foundation. This foundation is built upon the multiprotocol routers and multilayer LAN switches that are used as building blocks for enterprise networks.

As the industry leader in multiprotocol routing, Cisco Systems has an unparalleled array of voice-capable platforms at its disposal. These include a range of products that have the ability to terminate both analog and digital voice interfaces for integration with a legacy PBX or connection to the Public Switched Telephone Network (PSTN). Routers and LAN switches that can terminate voice traffic today include the Cisco 700, 800, 1750, 2600, 3600, MC3810, 5300, 5800, 7100, 7200, and 7500 series products, and the Catalyst 4000, 6000 products.

With IP telephony expanding beyond simple toll bypass applications to the desktop, local-area networks (LANs) need to provide the prerequisite quality of service (QoS) and bandwidth required to support converged network applications such as voice and video. The Catalyst® family of multilayer LAN switches offers the required features and functionality to achieve this goal.

Switched versus shared media is a must for IP telephony and the price per port of switched media is now comparable to that of legacy shared devices. Advanced classification queuing, multicast, and buffering techniques are also required to ensure voice and videos are effectively transported. Once again Catalyst switches provide the required functionality and will include the ability to support line power for the next generation of IP telephone.

Applications

Perhaps the most exciting facet of converged networking is the enabling of new applications. Such emerging applications include desktop IP telephony, unified messaging, and the Cisco IP Contact Centers, all of which will be discussed further. The applications for this technology however have no limits; new and innovative applications will continue to emerge. A converged network will offer the framework that permits rapid deployment of these new technologies.

IP Telephony to the Desktop

By using the Cisco CallManager, a PBX can be eliminated and replaced with IP telephony over a converged network. As shown in Figure 5, the Cisco CallManager provides call-control functionality and, when used in conjunction with the IP telephone sets or a soft telephone application, can provide the PBX functionality in a distributed and scalable fashion. Cisco CallManagers can be networked via IP and provide fall back to the PSTN if required.

Figure 5 IP Telephony to the Desktop

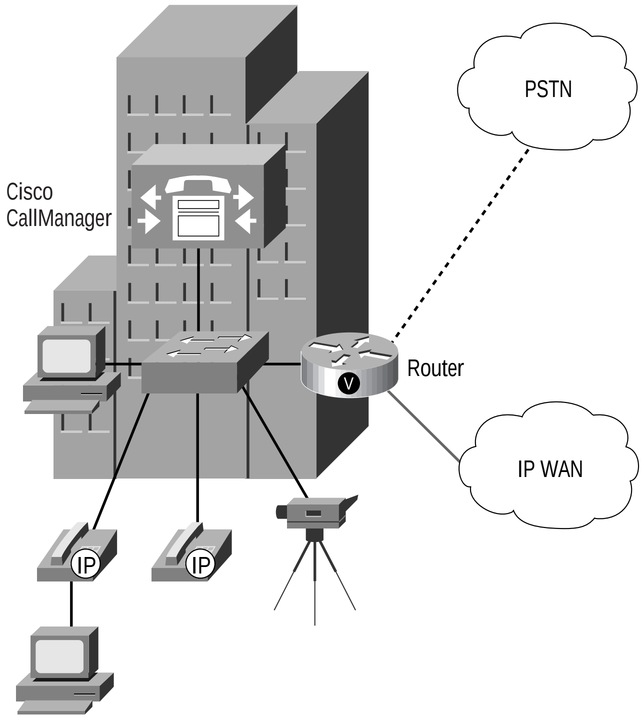

Unified Messaging

Today users have a wide range of communication and messaging mediums available to them: telephones, cell phones, pagers, fax, voice mail, and e-mail. Each of these requires distinct hardware and software components to function. Unified messaging combines voice mail, e-mail, and fax into a single application suite.

With unified messaging a single application can be used to store and retrieve an entire suite of message types. Voice-mail messages stored as WAV files can be downloaded as e-mail attachments while traveling, a response recorded and returned to the sender, all recipients, or an expanded list. E-mail can be retrieved via a telephony user interface (TUI), converted from text to speech, and reviewed from an airport lobby phone or cell phone. Infrastructure is decreased as now a single application can provide voice, e-mail, and fax. Productivity is increased because what were once disparate message types can be retrieved via the most convenient, or the user's preferred, interface.

Cisco Systems through its acquisition and integration of the Amteva innovations is able to offer unified messaging via its GateServer series of products. These provide scalable solutions for service providers and the enterprise via open standards-based interfaces. Figure 6 depicts this unified messaging model.

Figure 6 Unified Messaging

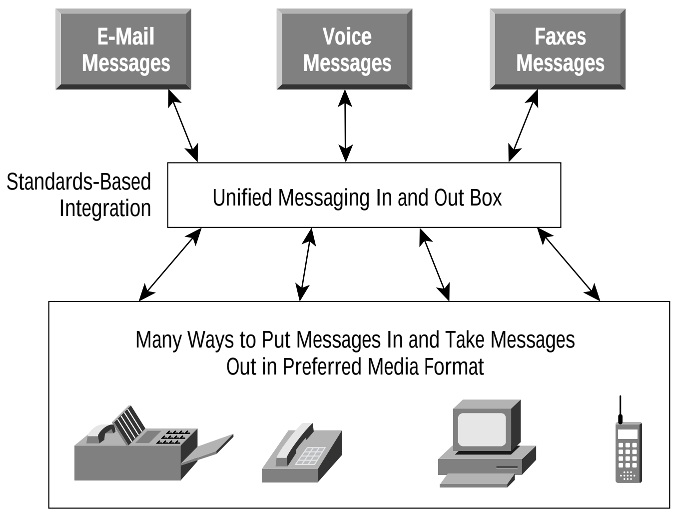

Cisco IP Contact Center

The Cisco IP Contact Center (IPCC) solution combines data and voice technologies to facilitate geographic independent multimedia customer interaction. This includes customer interactions originating from multiple diverse contact channels including IP voice, TDM voice, Web, e-mail, and fax. Regardless of transport, whether the Internet or the traditional PSTN, the IPCC fully integrated contact center architecture depicted in Figure 7 below services all media types. The Cisco IPCC architecture also provides a seamless migration path from the legacy call-center infrastructure to the IP-empowered, multimedia contact center.

Figure 7   Converged Architecture Supporting Data, Voice, and Video

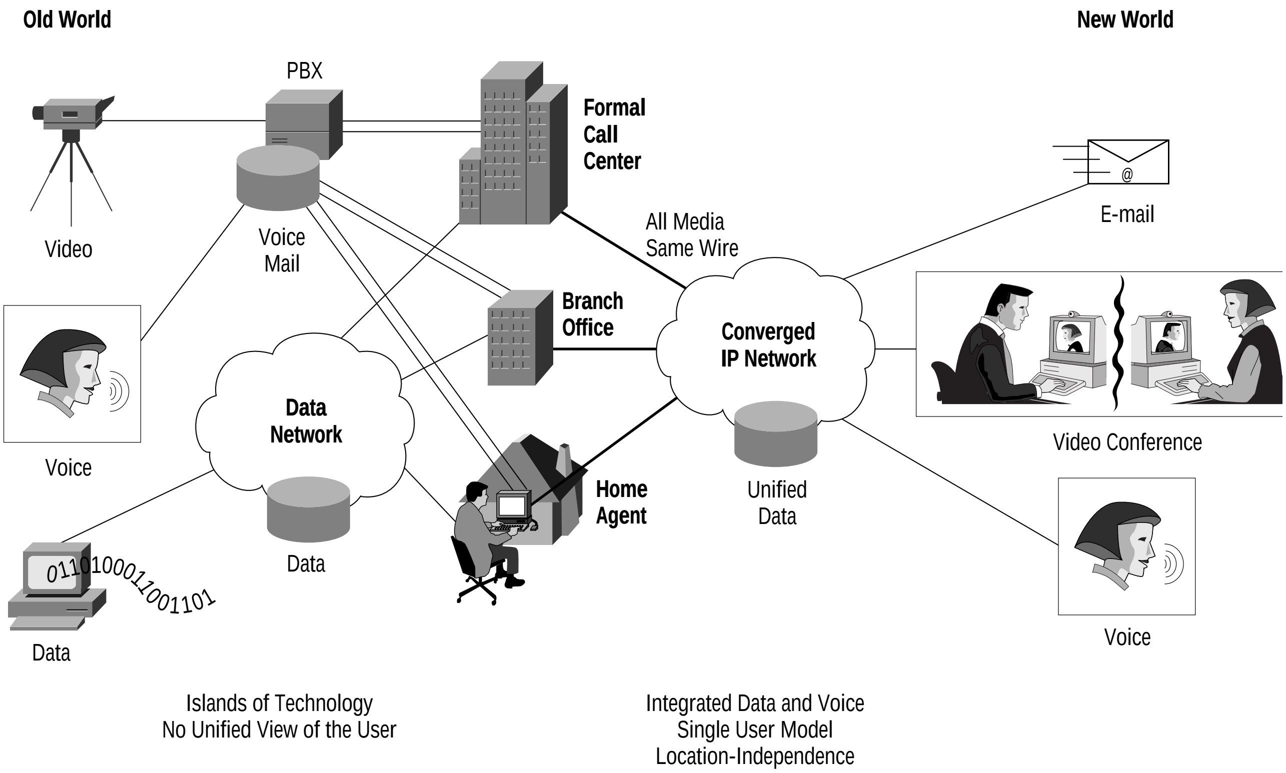

The New World IPCC solution can de deployed for new, “green-field,” IP-based contact centers. More importantly, the IPCC solution also enables server and agent-level IP telephony to coexist with traditional time-division multiplexing (TDM) based networks, existing automated call distribution (ACD)/private branch exchanges (PBX) and installed desktop systems. The Cisco IPCC solution enables an organization to take advantage of new IP-based applications at its own pace while preserving heterogeneous legacy investments and leveraging existing IP data infrastructure. Thus IPCC deployment can be incremental, adding IP telephony, new media channels, and new IP-based services at a rate that meets business demands and budget constraints. As existing solutions mature and new customer-interaction requirements emerge, such as Web collaboration, Internet voice, and e-commerce, the IPCC architecture will provide a seamless and flexible path for migration.

Specific immediate IPCC business benefits beyond heterogeneous support include:

- Enhanced productivity and profitability through new IP-based applications such as integrated multimedia queuing
- Enterprise-wide contact management based on a single set of business rules and supported by normalized consolidated reporting
- Increased customer satisfaction through personalized customer interaction
- Geographic independence of both agent resources and IP-based application servers through the ubiquity of IP transport
- Carrier-quality fault tolerance and system reliability
- Near infinite solution scalability from single-site to multisite to network service provider services
- Significantly enhanced competitive advantage through rapid solution deployment, many times faster than traditional TDM solutions
- Lower total cost of ownership, lower capital-equipment investment, single network, and single support staff eliminating the overhead of multiple diverse data, voice, and video networks

The physical topology of the IPCC as shown in Figure 8 is a solution fabric of interwoven servers, applications all interconnected via IP. Physical proximity of agents and application servers is no longer required. Via a single IP thread, the Cisco IP Contact Center can now carry high-fidelity voice to agents throughout the enterprise network, and, on that same connection, provide standard computer telephony integration (CTI) applications as well as features such as Web collaboration, chat, and unified messaging.

Figure 8 New World Contact Center Topology

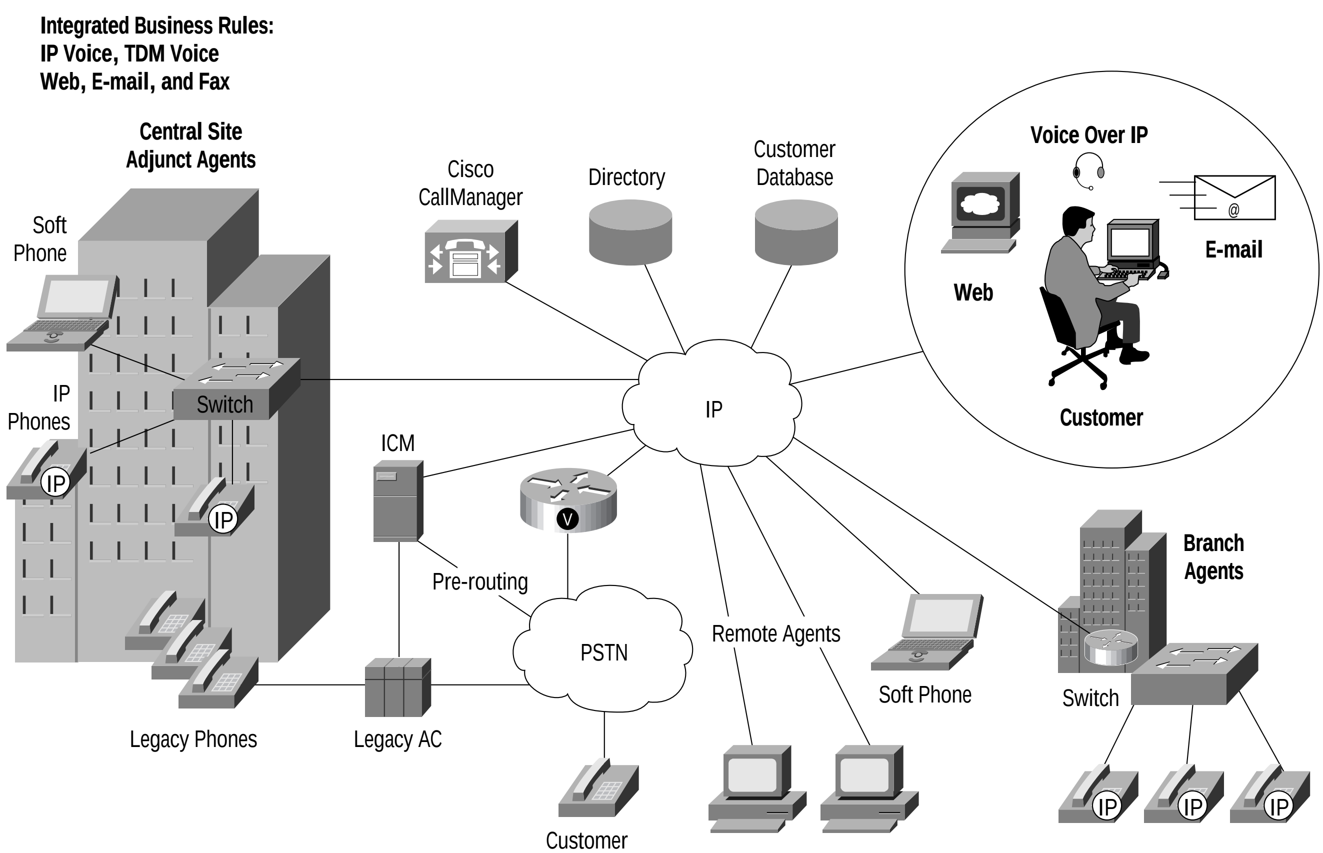

Technical advantages to the IPCC topology include:

- Intelligent contact management for personalized service and customer loyalty

- Enterprise-wide command and control
- Network-level customer queuing, customer segmentation, and contact distribution
- Consistent service standards across diverse media channels
- Proactive technical support with remote system monitoring
- Product, not promise—installs in 60 to 90 days; return on investment (ROI) in 6 to 12 months
- Scalable applications—augment services by adding servers anywhere in the network
- Leverages current, heterogeneous technology investments (no forklift upgrade)
- Seamless migration path to IP-based voice applications
- Easy and rapid deployment of remote agents
- Carrier-class, distributed fault tolerance

In addition to the improvement in development and deployment of new services and applications through a converged IP infrastructure, a further benefit of this technology is to offer users a simpler, integrated interface that presents information consistently across multiple media channels. This will result in a superior customer experience while simultaneously increasing agent productivity.

Users of this technology can, for example, utilize a Web interface to research a product. When questions arise, they can “click to talk” and be connected to a knowledgeable and highly skilled agent who is familiar with that product and has the same screen display, avoiding a multistep and multicall process. Similar schemes can be used for banking over the Web, avoiding the need to enter the account number multiple times and connecting to an agent only when a customer desires or an automated response is insufficient.

The Cisco IP Contact Center represents the next generation of customer care melding IP telephony technology, intelligent contact management technology, as well as legacy call-center applications and hardware into a unified platform to implement an organization’s business rules and objectives to deliver personalized, superior customer service for every customer contact.

Clients

The third component of the architecture is client devices. Cisco Systems is shipping IP telephones and a software only application that allows a fully functional telephone instrument to be emulated on a personal computer (PC) or laptop. These will evolve to include IP telephones that allow messages to be sorted and reviewed based upon user preference rather than the serial fashion currently available with legacy systems. Client applications such as directory lookup and click to dial are only the beginning of innovations in this area.

Open Standards-Based Architecture

Cisco Systems is promoting the use and adoption of open standards and is participating actively in the definition and approvals process for a number of standards and open protocols in this arena. Cisco Systems adopts these standards as they emerge and mature and also offers prestandard implementations to the market where no defined standard exists. Every effort is made to ensure that the gateways, applications, and clients produced integrate and operate seamlessly with third party products.

Examples of these protocols include the existing and emerging standards-based protocols for call control: H.323, the Simple Gateway Control Protocol (SGCP), the Media Gateway Control Protocol (MGCP), and the Session Initiation Protocol (SIP). The above are already standards or are in the process of being ratified as standards at the time of this document, and are all designed to provide call control for media gateways such as routers and switches. Cisco Systems plans to implement all of the above standards and provide best-in-class product for all applications that use them.

Further examples of open standards currently being adopted by the telecommunications industry are the Telephony Application Programmable Interface (TAPI) and the Java Telephony Application Programmable Interface (JTAPI). These protocols are used to communicate between applications such as the Cisco CallManager providing IP PBX functionality and unified messaging products such as the GateServer products acquired through the acquisition of Amteva. This open and standards-based interface model, depicted in Figure 9, is in direct contrast to the proprietary interfaces of legacy PBX equipment. Already third party innovators such as Active Voice and Telekol have products that can interoperate with the Cisco CallManager.

Figure 9 Open Systems Architecture

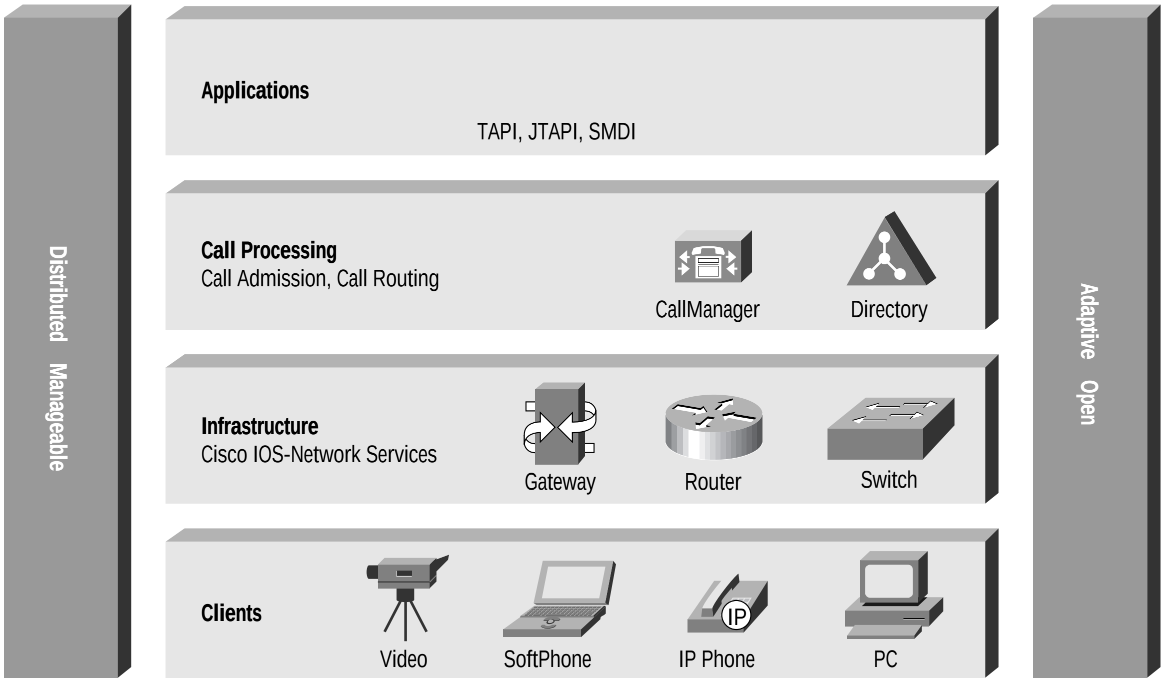

The use of open standards and the promotion of multivendor collaboration and interoperability are a key benefit of the Cisco AVVID architecture. The architecture creates an environment that fosters competition; this in turn lowers prices for the consumer. It also allows the integration of products from multiple vendors to create a customized solution. No single vendor can provide a solution that fits all requirements for data, voice, and video. Often specialized applications are designed and implemented only by a single company and need to be integrated with the overall solution. The adoption of open standards creates an ecosystem that actively promotes a model of integration.

Enabling Rapid Application Deployment

Another key benefit of Cisco AVVID is its ability to enable innovative applications to be developed and deployed more rapidly than their old world counterparts. A recent and irrefutable example of this is the advances in electronic commerce.

Billions of dollars of business is now conducted over the World Wide Web, creating a new industry. Cisco Systems alone processes almost $1 billion dollars of orders online per month. Contrast this to a manual scheme where orders are faxed to a sales representative, entered locally, then reentered at headquarters. Errors in such a process are the rule rather than the exception. The benefits of the new technology are many. Now customers can enter the configuration required on line, verify via automatic tools that the configuration is error free, and order the product. This saves time and money, and errors are drastically reduced.

The rapid innovation was possible only because the application runs over a converged network based on IP. Applications can be written independent of operating system; connectivity and compatibility with any other IP-based applications is assured. Contrast this to an application-requiring integration with a legacy PBX where the architecture is closed and proprietary, stifling innovation and increasing costs.

Highly Available and Scalable

A prerequisite for voice networking is high availability, and voice networks are often quoted as being more critical and available than their data-centric counterparts. This situation is rapidly changing; e-mail surpassed voice in 1998 as the primary business communications media (Frost & Sullivan). The concept of data tone is emerging, where the availability and reliability of the data network is every bit as important as the voice network. Cisco AVVID is a distributed architecture that is inherently available and scalable. The ability to seamlessly provision additional capacity for both infrastructure, services, and applications is a unique benefit of the architecture. Cisco CallManager clusters can now scale to 10,000 users. We depend on this technology to run our own business.

Within Cisco Systems, today 80 percent of orders and more than 82 percent of customer inquiries are transacted over the Web, making Cisco a leading example of how to use Internet applications for competitive advantage. As mentioned earlier, Cisco Systems processes almost $1 billion a month in online orders, making Cisco Connection Online (CCO) one of the largest e-commerce sites in the world. For the third consecutive year, Cisco was awarded the Web Business 50/50 Award from CIO Magazine for both its internal and external Web sites. None of this would be possible if the reliability and availability of the Cisco data network had not been assured.

Manageability

Cisco AVVID is complementary to the CiscoAssure Policy Networking initiative. CiscoAssure Policy Networking enables business users and applications to use the intelligence that is embedded in a network. Simply put, CiscoAssure Policy Networking makes it easier for a network manager to take advantage of distributed network intelligence features.

The CiscoAssure Policy Networking architecture is based upon four building blocks: Intelligent Network, Policy Services, Registration and Directory Services, and Policy Administration. Products currently available in this management suite are CiscoWorks 2000, Cisco Network Registrar™, and Cisco QoS Policy Manager.

The Old World—Box Reliability

The PBX has evolved over many years to a device perceived as highly reliable, performing basically one function and one function only, the switching of voice calls with some added services such as transfer and conference. Each vendor maintains a proprietary architecture to ensure that once a customer is using a particular brand of switch, the customer needs to continue using that same brand of device to maintain feature parity.

The system is deemed available by the user if dial tone is present when the handset is lifted. This is only a subjective measurement of the availability of dial tone. If a busy signal is received when a long distance number is dialed, there is no way to know whether all circuits are busy, the person called is busy, or the link to the PSTN is down. All these are situations that affect availability of the system to the user. In the data arena, they would be factored into system, as opposed to box, availability.

High availability in these cases is obtained at huge cost. The PBX will include as options redundant (standby) components for every major function such as power, CPU, and memory. The result is a cost-prohibitive system that has limited scalability and flexibility, often requiring a fork lift upgrade to expand capacity.

The New World—Network Availability

In contrast to the above model, the world of data networking presents a picture where availability is designed into a distributed system rather than a box. Redundancy is available in the individual hardware components for services such as power and supervisor modules. Network redundancy is, however, achieved with a combination of hardware, software, and intelligent network design practices. Figure 10 is typical of an enterprise network topology.

Figure 10   Network Availability

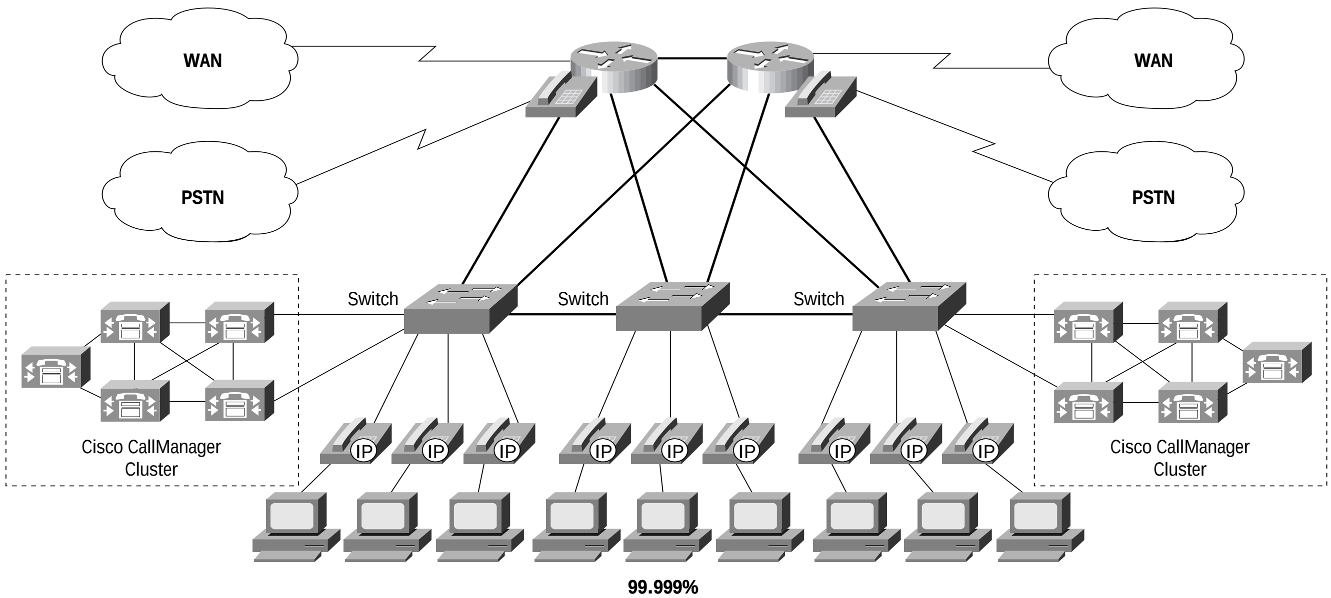

In the above diagram, network redundancy is achieved at many levels. Physical connections exist from the edge devices where IP telephones and PCs are attached to two spatially diverse aggregation devices. Should an aggregation device fail, or connectivity be lost for any reason (such as fiber cut or power outage), cut over of traffic to the other device is possible essentially without loss. This is true also of the wide-area network and PSTN connections. Clusters of Cisco CallManagers can be provisioned to provide resilient call control; should any device within the cluster fail, the other servers pick up the load. These designs can provide 99.999% reliability.

Advanced Layer 3 protocols such as Hot Standby Routing Protocol (HSRP) or, fast converging routing protocols such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP) can be used to provide optimum network layer convergence around failures. The above protocols can be configured to provide extremely fast convergence and also scale to many thousands of routes easily due to their distributed form.

Multicast routing can be used to optimize traffic that is required at many locations. Examples of this would be a video broadcast or an IP telephony conference call. The availability of advance multicast routing protocols such as Protocol Independent Multicast (PIM) as well as legacy Multicast protocols such as Distance Vector Multicast Routing Protocol (DVMRP) ensure these services can be deployed efficiently.

Moving down the protocol stack, advanced tools are also available for the MAC layer, Layer 2. Tunable spanning-tree parameters and the ability to supply a spanning tree per virtual LAN (VLAN) allow fast convergence. Value-added features such as uplink-fast and backbone-fast allow intelligently designed networks to further optimize network convergence.

Enhanced Voice Quality Options

For decades now voice networking has revolved around the Digital Signal Level 0 (DS0). This is a 64 kbps signal that limits the effective telephony audio range to 4 KHz. While efficiency schemes that lower the required bandwidth have been available for some time and are useful for WAN applications, little has been done to expand the audio range and enhance quality.

IP telephony and, in particular, future Cisco products that complement the architecture will include not only bandwidth-efficient CODECs such as G.729 and G.723.1, but also high-fidelity CODECs such as G.722 that deliver high fidelity audio to the handset. This is an innovation simply not possible with today’s TDM-based platforms.

Lower Total Cost of Ownership

While data networking has evolved to open, distributed, standards-based systems, the telephony infrastructure has changed little in the past 20 years, and those economies and efficiencies associated with open standards and competition have been impossible to attain. Once a PBX vendor has been selected, and the product implemented, the proprietary and closed architecture of that PBX effectively prevents multivendor interoperability at anything other than basic levels. This has kept the price per port of PBX systems relatively flat for recent years, and also shackled customers to the PBX vendor.

Contrast the above to data networking and the picture is very different. Moore's Law has demonstrated that the price/ performance of semiconductors doubles every 16 months. These savings have caused the cost of data networking equipment to fall rapidly over time, while performance has increased exponentially. These benefits translating into reduce prices for customers.

Consider for example the shared 10 M Ethernet connection or 4 M Token Ring to the desktop that was the norm until recently. Now a 100 M connection to the desktop is typical and the price is less than that paid for shared 10 M Ethernet only a few years ago. This represents a 20-plus-fold increase in available bandwidth, for a fraction of the cost.

Other factors that lower costs include the reduction of wide-area facility requirements, fewer devices to manage and maintain, and simpler moves, additions, and changes. This results in a lower training and staffing cost associated with a simplified and converged infrastructure. Figure 11 depicts this duplicity in economy between old and New World architectures.

Figure 11 Price/Performance Comparison between Ethernet Switching and PBX Ports

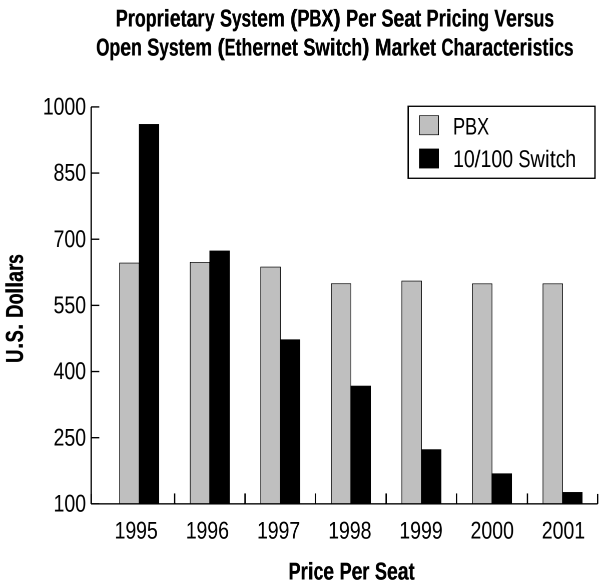

Enterprise Network Architecture

This section will detail the architectural model of implementation for Cisco AVVID. As with all initiatives, this is to be an evolution rather than a revolution and will require a phased approach and implementation. Three enterprise building blocks will be discussed: the branch office, the campus/MAN, and the wide-area network (WAN).

Branch Office Network

Cisco AVVID will first be implemented in the enterprise branch office locations. Figure 12 below depicts a typical branch office of 100 users or less.

Figure 12 Enterprise Branch Office Network with Cisco AVVID

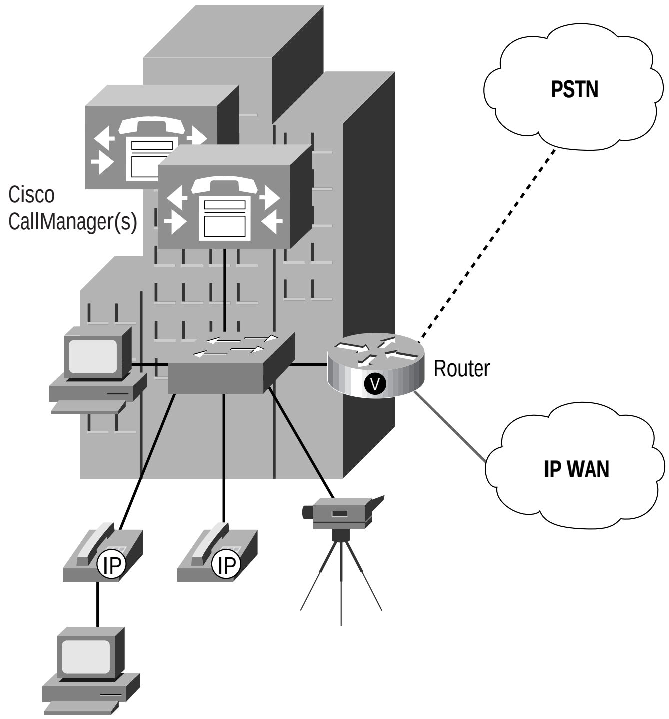

In the branch location, the infrastructure comprises a router with voice capabilities to interface with the PSTN for of-net and overflow voice calls. The primary transport for intersite on-net voice calls is the IP WAN. The router provides advanced QoS capabilities to ensure high voice quality as well as native multicast capabilities for video applications. Because IP is independent of the WAN media, leased lines, Frame Relay, Asynchronous Transfer Mode (ATM), or emerging last-mile technologies such as cable and digital subscriber line (DSL) could be used.

LAN infrastructure can be provisioned using a Catalyst multilayer switch. Catalyst line cards provide inline power to next-generation IP telephones and customers can choose to connect both IP telephone and PC to a switched port (via an integrated 10/100 switch in the telephone) or use separate ports for PC and telephone. Again the prerequisite classification, queuing, and buffer management features are available with the Catalyst switches.

Call control and integrated voice mail can be provided for the branch office in packaged solution. This can be optionally augmented with a second platform for additional redundancy. Client applications can be provided by IP telephones or softphones installed on PCs.

Campus Network

The next logical step is to migrate to a converged network in the regional office, campus, and MAN networks. The requirements for QoS and reliability do not change, only the scale of the solution. Figure 13 depicts a typical enterprise campus network design and will be used to discuss the necessary building blocks for a Cisco AVVID deployment. While the diagram depicts a large campus, the modular design allows design philosophy to be scaled from 100s to 10,000s of stations with no loss of performance or resiliency.

Figure 13   Enterprise Campus Network Design Model

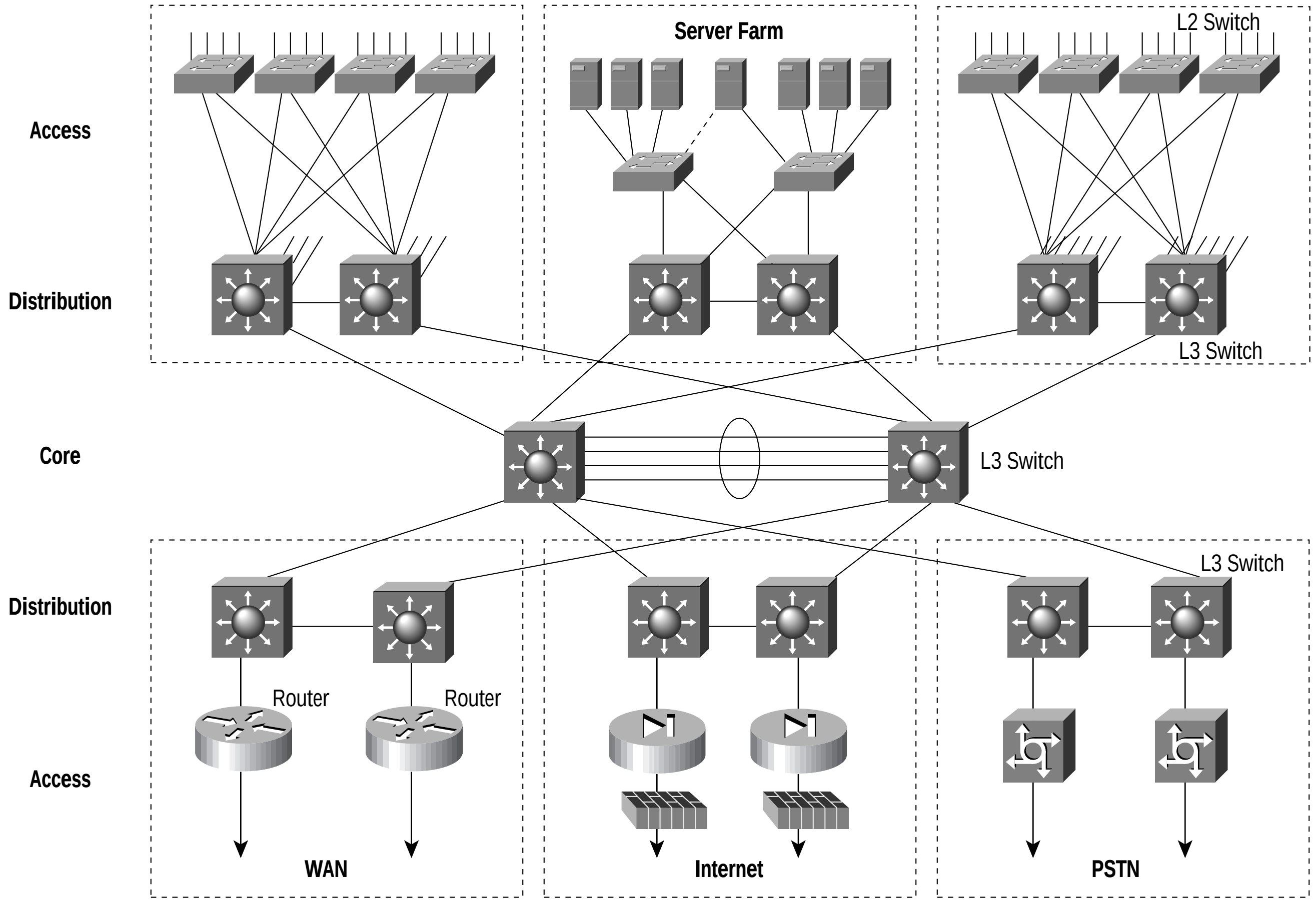

A detailed discussion of the above design is outside the scope of this paper, so the discussion here will be limited to the functionality of the components. As in the branch example, Cisco routers provide the required WAN and PSTN access. Here these routers are required to scale to support the incoming and tandem traffic from many smaller branch locations. Typical platforms used for this function would include the Cisco 7200 and 7500 families of multiprotocol routers.

LAN infrastructure is provisioned using the Catalyst family of switches. Here again the required classification, queuing, and buffering schemes are available. The design model shown is hierarchical in nature. This gives predictable and scalable performance, and also ensures fast convergence in the event of a component failure. Telephones can be connected to this switched infrastructure at the access layer either in series with a PC or via dedicated switched ports.

The Cisco CallManager and unified messaging applications in this instance can be located on separate, dedicated servers. This is depicted as the server farm building block. The Cisco CallManager will support a clustering scheme that provides a distributed scalable and highly available model. As additional users are brought on line, simply adding a new server to the cluster adds capacity to the system. In a similar fashion, voice messaging or unified messaging can be provided via the Amteva products installed on dedicated servers. As increased capacity is added, more servers are added to the system. Messages would be stored on an industry standard message store.

Wide-Area Network

For multiservice traffic to traverse a converged wide-area network (WAN), the network must support and supply the prerequisite QoS features. These are discussed in more detail in a subsequent section. In addition, the design and dimensioning of the WAN must be synergistic with the traffic profiles, business requirements, and circuit tarrifs.

Cisco Systems recommends that WANs are built using a hierarchical model to allow the most cost-effective platforms to be provisioned at the edge. At regional and headquarter locations, higher performance platforms can be deployed to allow the scaling of throughput and Layer 3 services. Figure 14 below represents the WAN model.

Figure 14 Enterprise Hierarchical Wide-Area Design

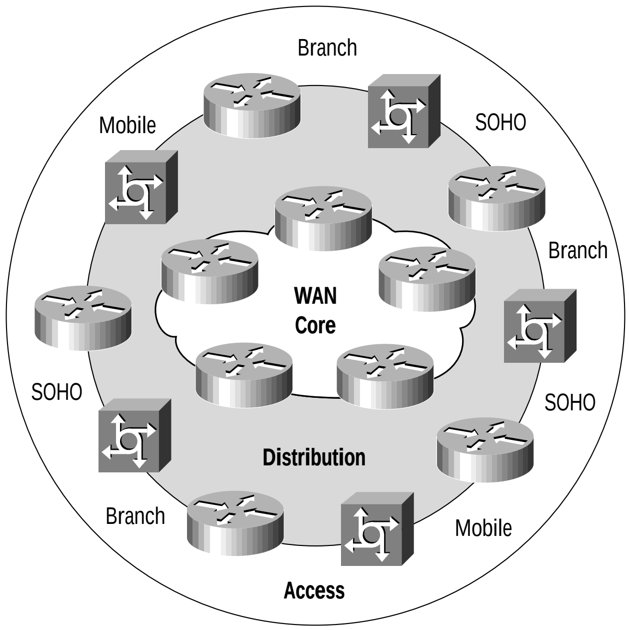

In addition to the above design philosophy, the WAN bandwidth requirements need to be adequately provisioned. As the requirements for data traffic outstrip those of voice, the percentage of the wide-area bandwidth required for voice decreases, lowering costs. It is imperative that the WAN links are provisioned to support the minimum requirements for data plus the bandwidth required for voice and video traffic. When other applications are quiescent, the bandwidth is available for data.

Enterprise Network Migration Requirements and Strategies

The following sections discuss the steps required to migrate to a converged multiservice network and reap the benefits of the Cisco AVVID architecture.

WAN QoS-Enabling Technologies

The following sections address these requirements in more detail. They will reference the network below when considering QoS for WAN applications. If prerequisite tools are applied, voice and video can with excellent quality traverse an IP-based WAN irrespective of media even at very low data rates. Figure 15 shows a typical enterprise wide-area network.

Figure 15   Typical Enterprise Wide-Area Network

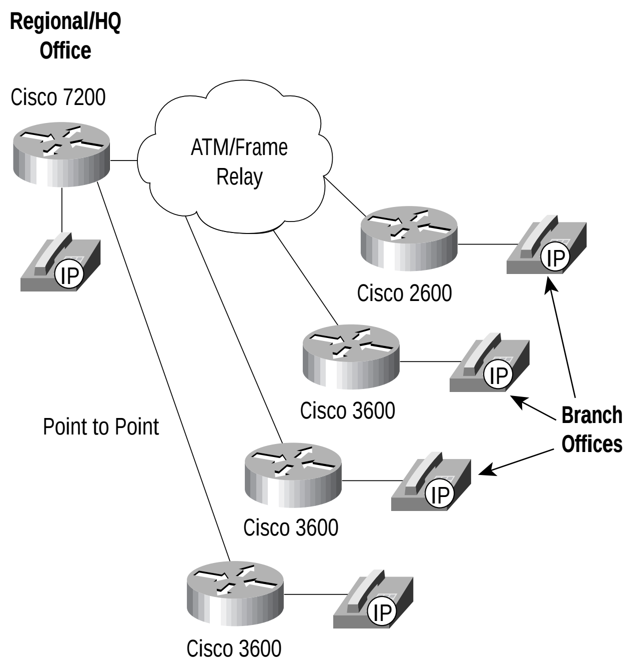

Classification

Before traffic can be handled according to its unique requirements, it must be identified or labeled. Numerous classification techniques exist which include Layer 3 (IP) schemes such as IP precedence or the use of the Differentiated Services Code Point (DSCP). Layer 2 (MAC) schemes such as 802.1P and also the use of an implicit characteristic of the data itself such as the traffic type using the Real-Time Protocol and a defined port range.

Cisco Systems uses all of the above, with the Layer 3 schemes providing the unique ability to classify from end to end. By default, the IP telephones use IP precedence level of 5 to provide critical priority to voice traffic. This will evolve to expedited forwarding when the emerging DSCP standard is implemented.

Prioritization

Prioritization schemes are also numerous, and the one which is chosen will vary dependent upon the traffic offered to the network and the wide-area media to be traversed. For multiservice traffic provisioned over an IP WAN, the recommendation for low-speed links is to use a combination of priority queuing for voice, and weighted fair queuing (WFQ) or class-based weighted fair queuing (CBWFQ) for the data traffic. Figure 16 below depicts this in operation where voice as a predictable and well-behaved traffic type is placed in a priority queue to be transmitted will low delay and delay variation (jitter). Voice is identified as RTP traffic within the standard audio range of User Datagram Protocol (UDP) ports 16383 through 32767. Data can be further subdivided using CBWFQ to offer differing levels of service to, for example, SNA traffic. Video traffic can be easily added into this scheme by inserting additional queues and adding LAN bandwidth to accommodate the higher traffic levels. The same traffic identification and classification methods can be applied, just as they are for voice.

Figure 16 Optimized Queuing for Voice over IP (VoIP) over the WAN

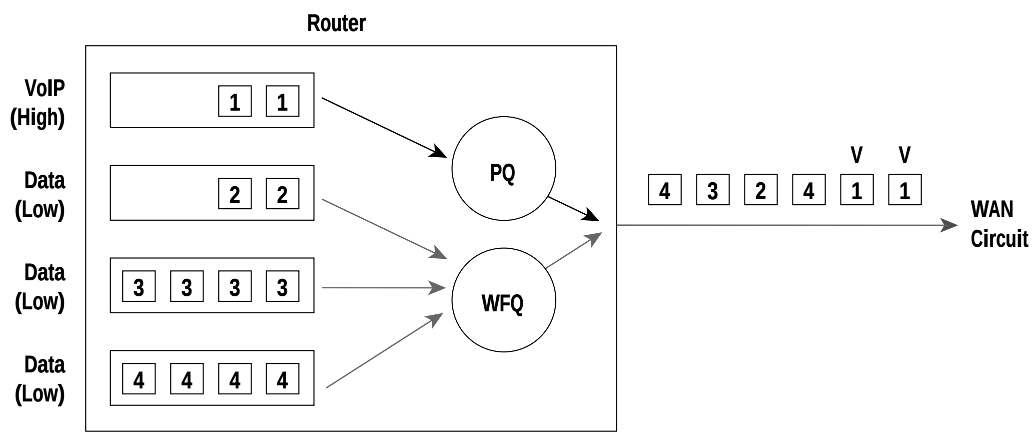

Link-Efficiency Techniques

Often wide-area bandwidth is still prohibitively expensive and only low-speed circuits are available or cost-effective when linking remote sites. In these cases, it is important to achieve maximum savings by transmitting as many voice calls as possible over the low-speed link. Many compression schemes such as G.729 allow a 64 kbps call to be compressed to an 8 kbps payload. Cisco gateways and IP telephones support a range or CODECs that allow efficiency to be enhanced on these low-speed links.

Link efficiency can be further increased by using compressed RTP which compresses a 40 byte IP + UDP + RTP header to approximately 2 to 4 bytes. In addition the use of voice activity detection (VAD) takes advantage of the fact that in most conversations only a single party is talking at a time. VAD recovers this “empty” time and allows data to use the bandwidth.

For very low-speed links (those with a link speed of less than 768 K), it is necessary to use techniques that provide link fragmentation and interleaving of packets. This prevents voice traffic from being delayed behind large data frames and hence bounds jitter. Two techniques exist for this, FRF.12 for Frame Relay and Multilink PPP (MLP) for Serial links.

Traffic Shaping

Traffic shaping is required with multiple access nonbroadcast media such as ATM and Frame Relay where the physical access speed varies between two end points. This technology enables the mismatch access speeds to be accommodated, and in the case of Frame-Relay, when combined with FRF.12, also allows delay variation or jitter to be bounded appropriately. For ATM, data rates are such that fragmentation is typically not required. Figure 17 below demonstrates this technique.

Figure 17 Traffic Shaping for Use with Frame-Relay and ATM

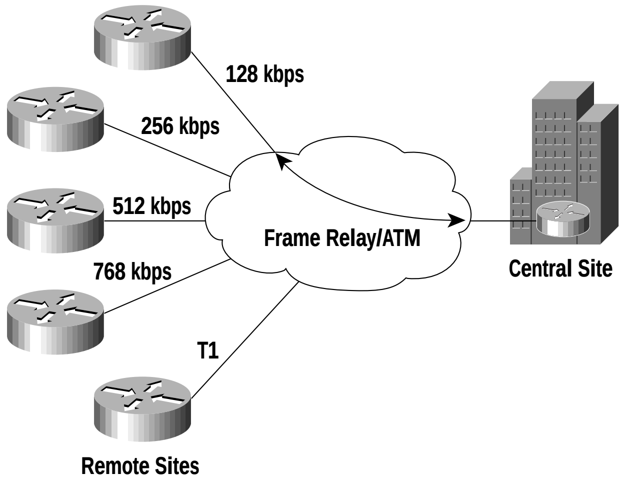

Quality of Service for the Campus

Quality of service (QoS) is an end-to-end requirement and must include all infrastructure components including the multilayer campus switches, routers, and edge devices. The area where resources are most constrained, however, is where the majority of the focus is required, hence the above discussions focused primarily upon the WAN.

Within a campus network two areas of concern exist. These are aggregation points where many links aggregate into an uplink, or where speed mismatches occur, such as from 10 M to 100 M, or 100 M to 1000 M connections. Figure 18 demonstrates this, where a number of edge switches are connected into an aggregation switch at 1000 M, with a 100 M uplink to a router.

Figure 18 Campus QoS Considerations

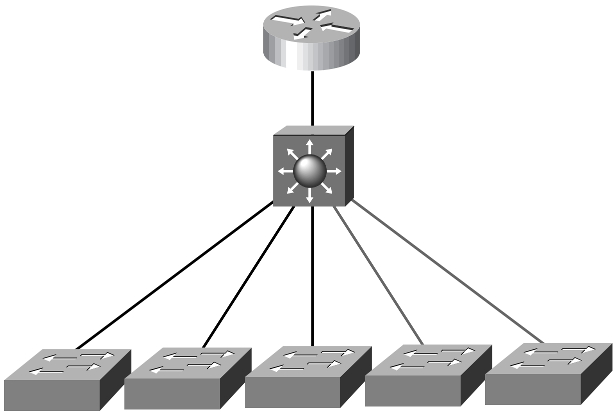

Within the campus, therefore, the primary concern is buffer management and the ability to provide the required queuing schemes to ensure voice gets the appropriate bounds for loss, delay, and jitter. A number of techniques exist that will be discussed below.

Of note, however, is the fact that WAN networks are often designed to run at or near maximum link capacity and hence congestion is the norm, due to the expense of the bandwidth. LAN networks in contrast are overprovisioned and undersubscribed with respect to bandwidth provisioning. Consequently congestion is the exception rather than the rule in the LAN.

Classification here is a key attribute. While the IP telephones manufactured by Cisco Systems classify voice traffic, those from other vendors may not. The multilayer LAN switches can classify or reclassify traffic accordingly. This feature is a must for an edge device and applies equally well to the attached PCs or servers or H.323 videoconferencing equipment that may require classification to a defined policy.

Queuing techniques are then available within the switches to ensure that data, voice, and video get the appropriate levels of service. These act upon the above classification and use schemes such as weighted round robin (WRR) to ensure that resources are allocated to achieve the desired results.

Buffer management is also used to ensure that excess jitter is not added for voice traffic. While large buffers can be beneficial to control loss in data streams, a potential side effect is added delay or jitter under congested operating conditions. To combat this, the multilayer LAN switches provide either two queues or a single queue with definable drop thresholds to bound delay.

Initial Migration Steps—Parallel Networks

For new installations and small branch office locations, a solution based upon Cisco AVVID can be implemented immediately or flash cut as a single event. For larger enterprise customers and those with a large installed base of legacy telephony, the process will require an evolutionary approach where the New World integrates with, expands, and encompasses the old, slowly removing the obsolete equipment.

The first instances of this will be technology trials where customers evaluate the equipment in a lab or on the desktops of a limited number of users. Once the technology has been evaluated and the operation and benefits verified, a parallel deployment will occur where users are migrated slowly to the new system. This will be done typically on number ranges to reduce the translations between the two systems. Figure 19 demonstrates this migration.

Figure 19 Migration Steps—Coexistence

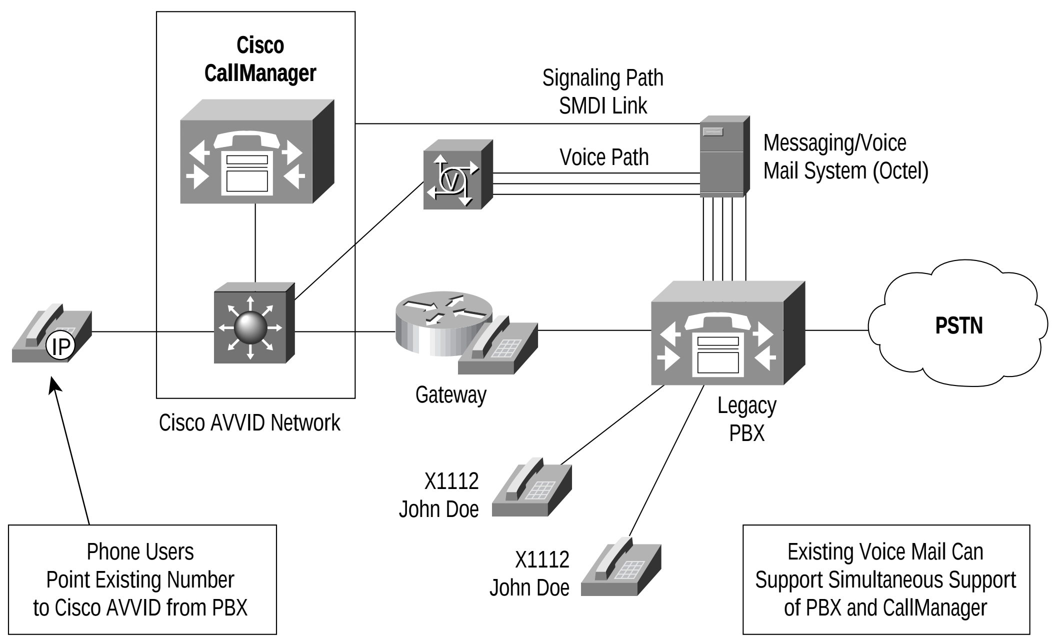

Once the process of migrating users to the Cisco AVVID architectural model has begun, it will also be necessary to integrate with the existing voice mail system pending a migration to a standards-based unified messaging solution.

Legacy Voice Mail System Integration

The vast majority of installed voice mail systems are proprietary implementations that prohibit the integration of systems from multiple vendors. The Cisco CallManager is able to interoperate with any voice mail system that supports the Simple Message Desk Interface (SMDI). The integrated Amteva products also have this functionality. In addition, the integration with proprietary implementations is possible with products that bridge the gap from the old to the New World.

The Converged Enterprise Data, Voice, and Video Network

The converged enterprise network for data, voice, and video will require the appropriate infrastructure and design as described above. Figure 20 depicts a converged network where all data, voice, and video utilize IP as the transport; between sites the IP WAN is the primary interconnect, with the PSTN being used as a secondary connectivity method.

Figure 20 The Converged Enterprise Network

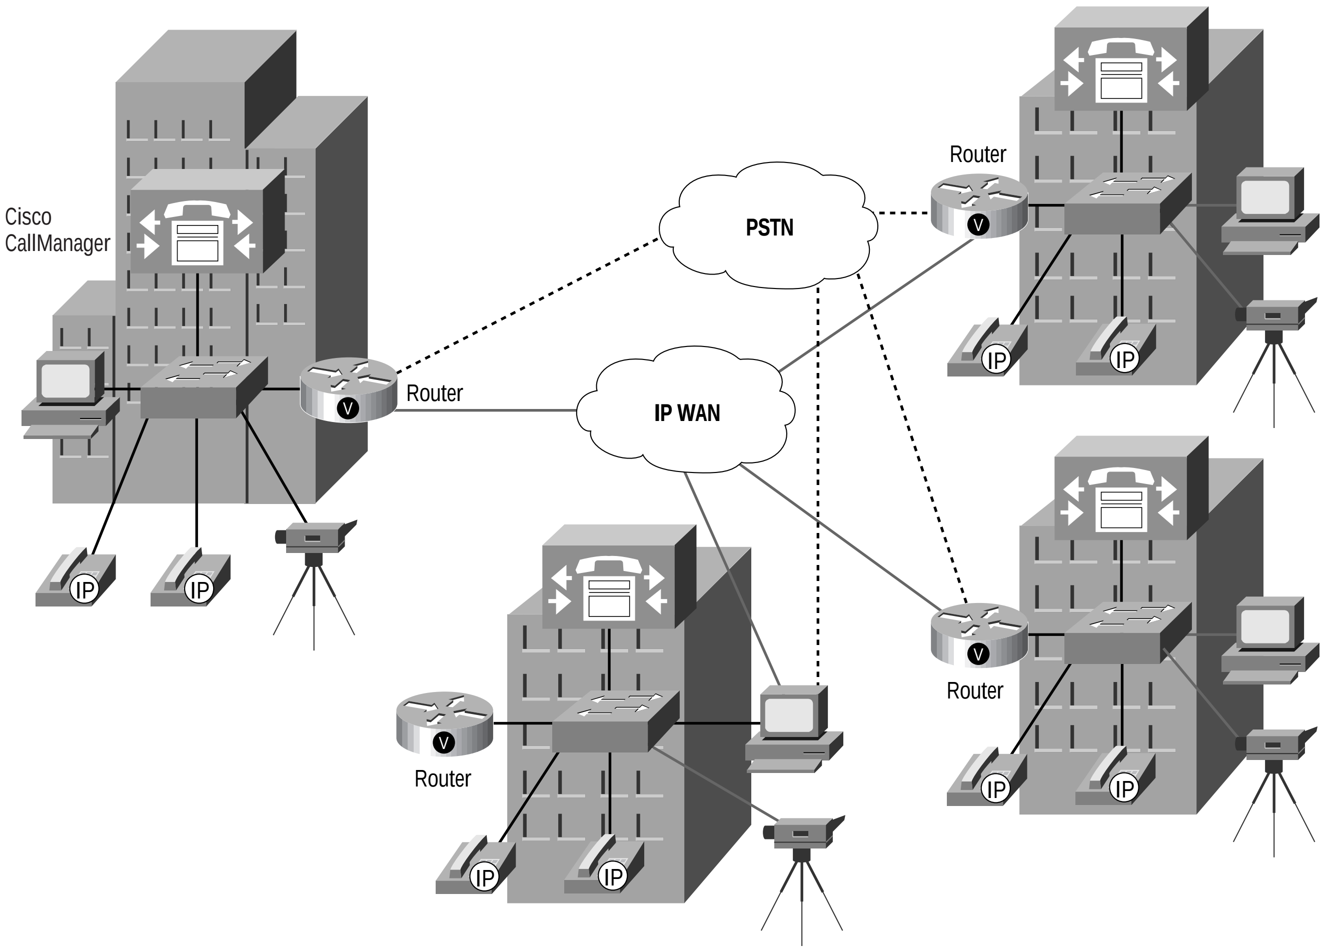

Such a converged network will lower costs and provide enhanced quality options for voice networking. It provides a highly scalable, reliable, and available network that is adaptable and permits the rapid deployment of new and innovative applications. Because the above network is based upon standards and open competition, interoperability with other New World applications is assured.

Service and Support

Cisco support solutions are designed for one purpose—to ensure customer success by delivering a suite of proactive support solutions. Cisco support solutions offer presales network audit planning, design consulting, network implementation, operational support, and network optimization. Customer success is enhanced through interactive knowledge transfer solutions that leverage
Cisco expertise and experience. By including support with Cisco equipment purchases, customers immediately gain access to a wealth of support and resources.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

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

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

3/00 LW