VoIP and IP Telephony

Abstract

Today, corporations and service providers are creating single networks that serve multiple purposes and save critical dollars. Currently there are two views of a converged network. The first is the creation of ISDN and Switched 56. This view is one of "Build a voice network and data will ride for free." "Legacy" data rarely required transmission rates greater than 56 Kbps, it was easy to consider transporting it along with the already establish voice network. The other was created when IP traffic accounted for a greater volume than voice on carrier networks. This view was, and still is, "Build a data network, and the voice will ride for free." This is the view of convergence that is the basis of voice over IP (VoIP).

Hill Associates offers a course that expands on this paper. For more information visit our website at www.hill.com.

Introduction

Imagine you are at home, connected to the Internet using your high-speed, always-on connection. You are listening to CNN or your local news on your computer; or you are listening to your favorite selection of music via MusicMatch. (If you have not done this, get online and go to http://www.cnn.com or http://www.npr.org. Click on one of the choices for live or taped broadcasts.) This is voice over IP (VoIP).

Now, imagine you call a customer service center to place an order, and the representative is located in a call center in Bombay, India. How does voice traffic make it cost-effectively from the United States to Bombay? It was probably an IP call. This is IP telephony.

VoIP generically refers to any technology that carries digitized human voice in IP packets. This includes trunking services as well as streaming audio and video applications. IP telephony adds to that basic capability all of the capabilities necessary to have a true phone system, including signaling and interworking with the public switched telephone network (PSTN).

According to the latest quarterly update to TeleGeography's U.S. VoIP report, VoIP subscribers grew 248 percent from 1.3 million in the fourth quarter of 2004 to 4.5 million subscribers at the end of 2005. And, "even more impressive was operator revenue growth of over 300 percent, from less than $200 million in 2004 to over $1 billion in 2005," This growth has occurred despite the fact that popular voice-enabled programs such as Skype and Yahoo IM are free. TeleGeography predicts that "by 2010, almost 19 million voice over broadband lines will be in service. This growth will largely come at the expense of incumbent telephone companies."

Vonage and Time Warner are currently the VoIP market leaders. Cablevision and Comcast are currently in the next two positions. Cable companies in general are the big winners, accounting for 52 percent of the market. Enhanced service providers, such as Vonage, account for 37 percent of subscribers, and telecommunications carriers, such as AT&T and Verizon account for the rest.

The bottom line is that voice over IP, and now IPT, is here to stay.

Why Combine Voice and Data?

Today, we can digitize voice and move it over a packet network. We have demonstrated this capability over frame relay, ATM, and now IP networks. We can move it quickly and reliably across wide distances, and users do not even know we are doing it.

Convergence underlies this movement. Convergence is the integration of voice telephony and data communications over a single network infrastructure. Convergence must consider all types of traffic—video, voice, image, and data (ViVID). As a result, there are considerations for business, personnel, technology, and service offerings, as well as implications for access and transport alternatives. Convergence must also consider the merger of wireless and wireline technologies.

ViVID transmission is moving to an IP infrastructure because businesses are under tremendous pressure to improve bottom line performance. To do so, corporations can do two things: improve productivity and reduce operating expenses. Since it is less expensive to support one network than it is to support multiple separate networks, carriers are replacing many overlapping networks with one infrastructure that can provide services seamlessly to end-users. Such a multiservice infrastructure enables a telecommunications company to diversify its product offerings while reducing its operating costs.

When building such a network, a carrier must decide on the fundamental principle of the base network. The choice is to use circuit or packet technology. Although a circuit network can be used to build a multiservice infrastructure, packet networks are much better suited to play this role.

Moving into a converged network infrastructure has its challenges, though. To prepare for them, a company must ask itself the questions below once it decides to move to a converged network.

- Why are we doing this? What is the goal?
- What will the voice audio quality be like?
- What is the infrastructure quality?
- What are the traffic and usage patterns on the existing voice and data networks?
- Will we implement just voice, or perhaps some features of telephony too?
- What are the security concerns and challenges to consider?
- What is the expected return on investment (ROI)? Have we considered the total cost of ownership (TCO), which includes soft cost savings from things like moves, adds, and changes?

What Is VoIP?

VoIP technologies use an IP-based packet-switched network to carry voice traffic. VoIP users can place and receive "telephone" calls via the Internet or via a private IP network. A private IP network can be completely on premises, or can use frame relay, ATM, or even MPLS-based virtual private networks to connect corporate locations. A packet-switched network is less expensive than the global PSTN because it does not dedicate resources (i.e., fixed path and bandwidth) to a specific user for the duration of a given call. As with data traffic, users get the resources they need when they need it, and all resources are flexibly shared.

The Basics of Packet Networks

Circuit networks establish a fixed path and fixed amount of bandwidth between two points for the duration of a communication. This has been the predominant model for telephony since its inception. Such networks tend to be all or nothing networks: the network accepts the call if it has the resources, and rejects the call if it does not. Although this makes for high quality communication, it is a waste when such a network is used to transmit data packets.

Computer communication is bursty; thus when circuit technology connects two computers, the circuit is idle for a significant amount of time, wasting network resources. Packet networks were developed to address this problem. In a packet network, multiple systems share the transmission path on a first come, first served basis. Each packet is labeled to indicate its source and its destination.

Interactive voice communication has long been associated with circuit networks because it has been viewed as a communication type that creates a fixed amount of traffic with tight timing and low delay requirements, all of which match the circuit model. But we have learned that many of those characteristics are linked to how human voice was digitized. In North America, pulse amplitude modulation (PAM)/pulse code modulation (PCM) encoding produces an 8-bit sample every 8000th of a second for a constant bit rate of 64 kbps. However, human voice is actually highly bursty. There are large pauses when we speak, which do not need to be transmitted. We also don't need to encode voice in such a static manner. We can compress the information into much smaller bit streams.

When we compress the voice signal and eliminate the pauses we end up with a very bursty bit stream that requires significantly less than 64 kbps of bandwidth. Timing must still be provided, but buffering the transmission at either end deals with most timing issues. The one characteristic that remains critical is delay. Interactive voice requires a network with very small end-to-end delays and relatively little variation in that delay.

In a packet network, delay can be highly variable, mostly due to the resource contention dilemma. One user's traffic must wait until another's has been transmitted on the shared link (queuing delay). Propagation delay (the amount of time it takes bits to travel on a wire) still exists, but it can be variable if different packets travel different routes between the endpoints of a call. These problems are magnified when voice traffic is intermingled with data traffic. Data traffic can produce large packets that can "clog" a switching node. The total delay, then, is unpredictable and variable based on the packet's destination and the amount of network traffic.

Packet networks are excellent for data traffic, since data is bursty and delay-insensitive, but they have generally been considered unacceptable for voice traffic because of this queuing delay problem. However, modern packet switches can handle 100 million packets per second at "wireline" speeds, which reduces queuing delay significantly. Delay is still possible and variable, but total end-to-end delay is often within acceptable bounds.

VoIP and QoS

Modern switches are fast, and creating networks with enormous capacity will help minimize delay. However, none of these things provides guarantees of end-to-end delay. The packet network remains a risky place to carry voice traffic. To make voice over IP (VoIP) acceptable to users in a more stable fashion, delay, jitter, echo, and noise must be explicitly controlled. Jitter is essentially variation in delay. Echo is the phenomenon we occasionally experience on our cell phones, when we hear our own voice reflected back to us. Noise is any random change made to the signal between transmission and reception. In a packet network, the latter translates to packet loss or bit errors.

These controls typically fall into the realm of quality of service (QoS). QoS refers to the requirement that a network provide a level of performance commensurate with the needs of the applications it supports. Each application expects some degree of consistency and predictability from the network.

QoS techniques must consider that different applications expect different handling from the network. Toll quality voice is a narrowband application, historically requiring either 3 kHz in an analog environment or 64 kbps in a digital environment. We no longer demand that 64 kbps of bandwidth be used to transmit voice traffic. Nonetheless, a network that carries voice must guarantee some level of minimum bandwidth, maximum level of end-to-end delay, near-constant delay variability (i.e., minimal jitter), and minimal packet loss or bit error.

Specifically, a packet network supporting an interactive voice application should meet the criteria below.

- The delay from sender to receiver must be small—less than 200 milliseconds (msec).
- The predictability of the delay (jitter) from sender to receiver must be minimal (10 ms).
- Senders should not hear their own voices echoed (near end crosstalk of -40 dB).
- The signal-to-noise ratio on the line must be high (S/N 1000).

Something to Ponder

Try a small experiment. Find a colleague, friend, or family member who has a mobile phone and use your mobile phone to call theirs. When you have a connection, speak into one phone and listen to the other. (This might be more effective if one end is on speakerphone.) You will notice a very small delay between the time you speak and the time your friend hears it. Why? The signal leaves one hand set, goes out to the cell site, through the mobile switching center (MSC), is routed around while it figures out where to go, out perhaps a different MSC and a different cell tower to your friend's phone.

Mobile phones have conditioned us to accept some amount of delay and noise, and most mobile phone users have experienced areas with poor coverage. But there are limits to our tolerance. Studies have shown that the average person will tolerate delays up to about 250 milliseconds. After that, conversation becomes stilted and people start talking over one another. If delay exceeds 500 msec (as it can in some satellite communications) we have to adopt a new voice protocol (e.g., by saying "over" at the end of a transmission) to compensate for the total delay. Echo cancellation might fail on our connection, so we would consistently hear our own voice reflected back. Such a service would never be acceptable. Most of us would hang up and try again.

QoS Alternatives for Voice over IP Implementations

Many options are available to achieve QoS in an IP network. The originators of IP did not add QoS mechanisms to the original protocol, so QoS must be "added" to IP networks. The easiest way to achieve the required QoS is to provide sufficient bandwidth for all voice traffic. "Sufficient," however, typically means significantly over-engineering the network to minimize the probability of queuing delay—not a cost-effective strategy. Other mechanisms are required. Generally, QoS mechanisms for use in an IP environment fall into one of three categories: One approach operates end-to-end; one operates hop-by-hop; and the last approach bypasses IP altogether and supports QoS at Layer "2.5."

End-to-end QoS approaches employ a call setup procedure to establish a path for transporting traffic. The IETF Intserv (Integrated Services) is an example of the end-to-end approach. It includes protocols such as Resource Reservation Protocol (RSVP) and Real-Time Transport Protocol (RTP). One concern with this approach is scalability in large networks.

Hop-by-hop IP QoS approaches do not require calls to be set up before transmitting session data. These approaches use predefined rules that result in packets following one another through network paths. The packets carry an identifier that associates them with one flow or another. The IETF Diffserv (Differentiated Services) is an example of a hop-by-hop QoS approach.

Multiprotocol Label Switching (MPLS) is an example of the Layer 2.5 approach. This network-based technology provides packet guarantees and low latency in a modified IP network. Different network paths can handle traffic with different QoS requirements. In many cases today, carriers provide QoS in MPLS-based networks using CoS markings translated from the Diffserv code points placed on IP packets at the egress router connecting to the MPLS network.

Functional Components of a VoIP/IPT Network

The components of an IP voice network are about as varied as the number of VoIP implementations. The list below covers the more generic items in most implementations.

- *Routed backbone:* The central component of a VoIP network, or any IP network, is a routed backbone. Routers can support lower layer switching, special routing and/or reservation protocols, the ability to communicate, and quality of service (QoS).

- *VoIP gateways:* These devices sit at the edges of the routed cloud and connect the IP network to the PSTN. The gateway's purpose is to packetize the voice signal and send it to the desired destination IP address, and vise versa. Gateways convert the call's signaling information and voice component.

- *Telephones:* The user devices—the phones and PBXs—comprise the customer premises equipment (CPE). These devices might be PSTN devices connected by gateways or IP-capable devices (e.g., sophisticated VoIP business telephone or a PC running VoIP software).

- *Call servers:* This optional component of the VoIP environment consolidates features and capabilities. The servers offer location, registration, and negotiation services.

- *Protocols:* VoIP implementations typically use two different types of protocols. One sets up and terminates calls, and the other transports the voice.

Voice over IP Protocols

Traditional voice is primarily considered a Layer 1, or Physical Layer, service. Consider the B-channel of an ISDN PRI. In contrast, VoIP, like any Application Layer service of the TCP/IP protocol stack, is a combination of protocols. IP and possibly MPLS focus on getting the individual packets across the network. These packets could carry either voice or the information that sets up and terminates calls, so IP is involved in carrying both the encoded voice and the call control information.

At the Transport Layer, the User Datagram Protocol (UDP) identifies the applications that source and sink the packets and is primarily used to carry voice samples. UDP, however, is not a complete enough protocol to support voice communication. The current standard protocol for transporting voice traffic (digitized voice samples) is RTP. RTP is an enhancement to UDP since it actually rides within UDP. The Transmission Control Protocol (TCP), if used, provides flow control and recovery, so it ensures that call setup, control, and termination information (i.e., signaling) is successfully transported end-to-end.

The telephony protocols that set up and terminate a voice call on an IP network—H.323, Session Initiation Protocol (SIP), and Megaco—originate from multimedia conferencing/networking applications. They have been refined and specialized for voice communications, but they still retain the capabilities for a true multimedia event.

Early implementations of IP voice were largely proprietary. Different vendors' equipment could not interoperate, meaning that all VoIP hardware and software had to be purchased from a single manufacturer. Today, companies are forming alliances to ensure that their equipment is interoperable, and the result is the emergence of standards. As IP telephony evolves, so will the standards, and increasing adherence to standards.

All parts of the IP telephony infrastructure need standards. Standards are required inside the IP network, at the interface to customer premises equipment, and at the interface to the PSTN; gateways must interoperate with Signaling System 7 (SS7), IP phones must support standard compression techniques, and QoS implementations, as they emerge, must be based on standards.

H.323

Chief among the major standards that relate to IP telephony is the ITU-T H.323 specification entitled Packet-Based Multimedia Communications Systems. H.323 is an umbrella standard for IP voice and videoconferencing, and many other standards. The view here is from the telephony, or voice, perspective. Anyone familiar with videoconferencing or audio conferencing has likely used H.323 call signaling to establish the connections.

H.323 relies on the H.225 protocol for call signaling that is similar to that found in an Integrated Services Digital Network (ISDN) service, while the H.245 protocol controls the actual media streams flowing between the terminals. H.323 also calls upon the T.120 data conferencing protocols to support applications such as shared white boards. H.323 provides the interface between IP voice and video devices and an IP or internetwork packet exchange (IPX) network. The actual voice/video devices and the IP/IPX networks are outside the scope of H.323. Gateway services allow H.323 to interface with either POTS or ISDN. H.323 signaling messages can define the use of a variety of voice encoding standards, including G.711 (pulse code modulation, at 64 kbps), G.726 (adaptive differential pulse code modulation, at 32 kbps), and G.723.1 (algebraic code-excited linear prediction, at 5.3 kbps). Video encoding standards include H.261 and H.263, both of which operate around 28.8 kbps.

Both H.225 (call signaling) and H.245 (transport control) use TCP as their Transport Layer protocol to ensure reliable delivery of messages. Most H.323 implementations carry data/voice/video packets with RTP. This protocol is mapped to UDP to ensure efficient delivery of voice samples.

Session Initiation Protocol

Defined in RFC 3261, the Session Initiation Protocol (SIP) is a multimedia terminal signaling protocol. Some view it as a competitor to the more complex H.323 protocols. Since the IETF is responsible for most things IP, the view here is from the data, or packet switching, perspective. The Domain Name System (DNS) allows computers to find other hosts based on names and recursive look-ups in databases, which also have IP addresses. Similarly, SIP, like H.323, defines the mechanism by which an IP telephone call is established and terminated. SIP is not used to carry the voice samples.

SIP operates in a client/server fashion similar to other Internet protocols such as the Simple Mail Transfer Protocol (SMTP) for email or the Hypertext Transfer Protocol (HTTP) for Web surfing. A SIP architecture consists of two parts: the user agent, which supports multimedia sessions, and the network server part, which provides address resolution and user location functions. Figure 1 provides a high-level overview of SIP operation between a pair of user agents (PCs) and a server. We could easily replace the PCs with an IP phone. The principles are the same.

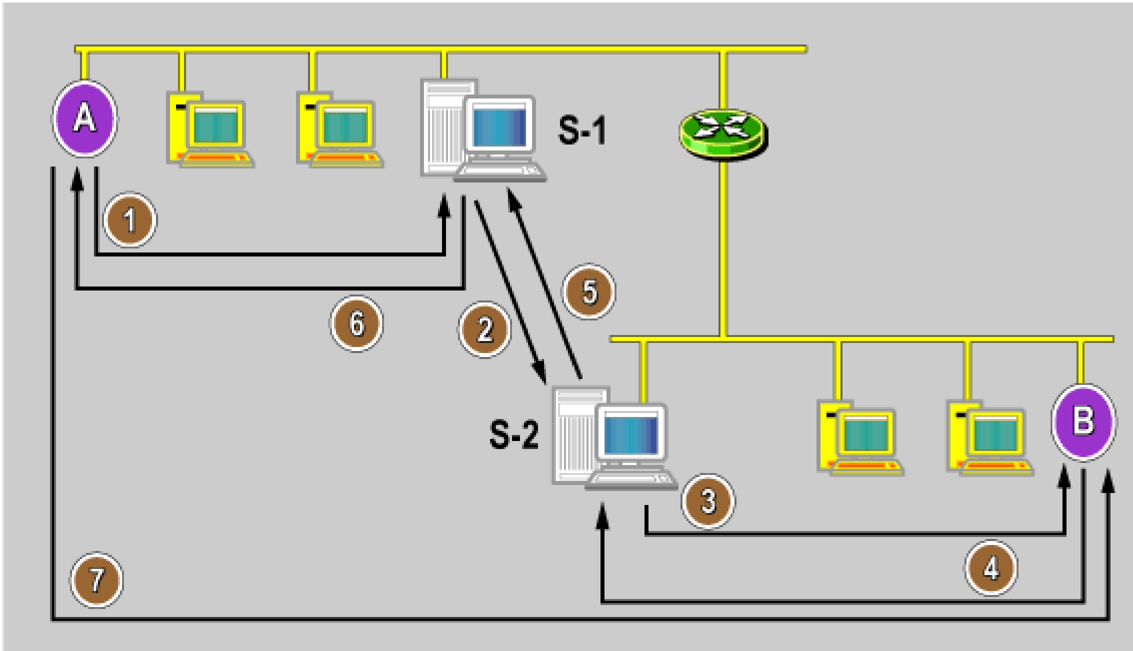

Figure 1. SIP Signaling Diagram

We begin at step 1 where user A calls user B. User A's phone sends an Invite message to its SIP server (S-1). The SIP Invite message contains the call description (e.g., codecs and protocols). The server (S-1) does not know user B's location. As a result, it forwards the invitation to all other SIP servers it knows, S-2 (step 2), where user B is registered.

S-2 forwards the invitation request to user B's phone (step 3). User B's phone checks the call description and accepts the call. When user B answers, the phone sends back an OK message, which follows the same path as user A's invitation (steps 4, 5, and 6). User B's call description is carried in the OK message.

Since both ends accept the other's call descriptions, the two phones can now begin the voice session (step 7). Note that all signaling messages flow through the servers while the call traffic flows directly between the terminals.

There is some debate as to the merits of using SIP as opposed to H.323 for multimedia terminal signaling. On one hand, H.323 is an international standard, is more mature, and is somewhat transport agnostic (e.g., IPX is mentioned in the H.323 series of specifications). On the other hand, SIP is simpler, IP centric, and enjoys far better support for mobile users due to its client registration and use of SIP servers. Both H.323 and SIP are widely implemented today, and almost any VoIP or IPT solution will likely support both. However, many reports suggest that SIP will be the ultimate long-term winner in the signaling game since it is better aligned with IP.

Megaco

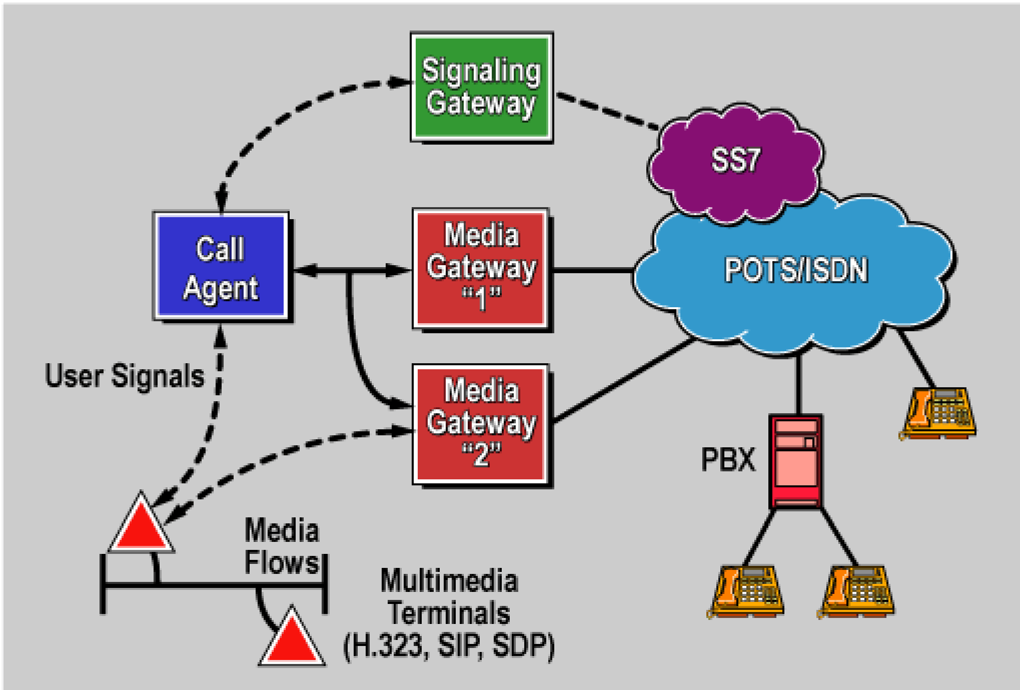

Figure 2. Megaco Signaling Diagram

This visual provides an overview of the Megaco/H.248 environment. Megaco also comes from the IETF, and is specified in RFC 2805. Megaco describes a distributed set of gateways controlled by one or more call agents. The call agent's role is to provide call control for packet-to-packet or packet-to-circuit media connections. The signaling gateway's function is to convey the signaling received from the PSTN to and from the call agent; the signaling gateway and the call agent are often collocated. The media gateway's role is to convert audio formats and transmission techniques between the PSTN and a packet-based VoIP network.

As shown in Figure 2, the call agent communicates with the gateways using Megaco, an asymmetric protocol in that the call agent always issues commands while the gateways always issue responses. The call agent can interface with the end-user via several different VoIP signaling protocols, but H.323 and SIP are the primary ones.

SIP and Megaco complement each other and interwork well due to the Session Description Protocol (SDP). In effect, the call agent simply converts incoming connection requests (either SS7/ISUP or H.323) into the appropriate SIP messages (Invite and OK).

VoIP Implementations

Coexistence

As networks everywhere embrace IP, it makes sense to IP-enable traditional voice equipment. Consider voice mail, which lives on a card, that has been added to the PBX or Centrex. It's essentially a fancy digital answering machine. If I can add cards to a PBX to enable new capabilities such as voice mail, why not add another card that enables the PBX to connect to an Ethernet network, and then place on that card call management functions tied to IP network understanding?

Such capabilities allow companies to converge their voice and data traffic on the wide area and keep existing equipment. Creating an IP-enabled solution does not have to be a forklift upgrade; such an implementation is often considered VoX (voice over some protocol). The difference between this and a full IP telephony system is that the call control and features are still controlled by the circuit-switched devices. Only the voice is passed over the packet network.

An IP-Centric World

In an IP-centric converged network, all devices are IP-capable. The PSTN becomes the backup to the IP network. Calls are routed across the IP network based on decisions made by a server—the brains for the new IP voice switch. Databases handle the logistics of what ties to where.

This converged network also allows companies to leverage common wiring infrastructure to save costs, bring new intelligence to the desktop devices, use unified messaging, and save on toll calls for internal corporate dialing.

Other VoIP Applications

VoIP by itself is not what it is all about. It's about all the new functionality that VoIP brings with it. VoIP applications are potentially as numerous as those for traditional telephony, but they come with a caveat. While roughly one billion people worldwide have easy access to a telephone (and therefore the worldwide telephone network), the number with access to a PC with an Internet connection is far smaller—by some estimates, only 150 million, many of which do not have daily access. If we can identify what those people want to do with their newfound access to VoIP, we will have discovered the applications that will drive the beginnings of the IP telephony world.

This is a daunting task to be sure, but there are a few assumptions we can make. While VoIP offers many possibilities for interesting value-added services, for most early adopters the primary driver for the technology will be one of simple economics; customers want to save money on telephone calls by replacing some of the calls with VoIP services, or lower their costs associated with moves/adds/changes of their existing phone infrastructure. Over time customers might also see a lower total cost of ownership (TCO) through reduction of capital equipment, consolidation of staff, and improved productivity.

For carriers and service providers, the money in VoIP applications is in the advanced services, not just the calls. As users incorporate IPT solutions, the feature set will expand to include capabilities brought to a handset not thought of before—integration with Microsoft Outlook and "click to call" capability, unified messaging, and the ability to get weather or stock information. Perhaps users will be able to reserve conference rooms without checking with an administrator, or a law office will be able to automatically track call accounting functions and billing. Corporations that manually track employee time sheets can now use the phone to allow employees to logon and check in and out. Employee bulletins are no longer posted as paper on a cork board, but pushed to an IP phone as messages. Managers can verify that their staff members have read the message as well. Colleges are using IPT and IP-enabled phones to advertise pizza parlors. If a student orders a pizza from a particular parlor, a percentage of the sale is returned to the college. How many PBXs do you know that actually generate revenue? This is merely a partial list of examples of what is possible. With the ability to create applications using XML (HTML for universal and smaller devices), almost anything is possible. The following section discusses some of the more traditional uses for VoIP and IPT solutions.

Voice-Enabled Websites

The World Wide Web and its protocols—Hypertext Markup Language (HTML) and Hypertext Transfer Protocol (HTTP)—have been the major agents behind the evolution of the Internet into the global marketplace it is becoming. The ability to display catalog information, offer goods and services, take orders, and provide customer support over a data network has changed radically with the advent of electronic commerce.

Some customers, however, will always be uncomfortable doing business through a website, and some aspects of commerce will not work well through a browser interface. For example, when the first retail websites appeared, a major obstacle to Internet commerce was the issue of payment; people were uncomfortable providing credit card information over the Web. So, many early businesses asked shoppers to place their orders via the Web, and then phone in their credit card information. Similarly, while a well designed website can provide an immense amount of information about a supplier's product or service, sometimes the sale can only be made by the intercession of a salesperson or customer service agent who can answer questions and provide personal advice—something very difficult to do via a browser.

So, the next generation of websites are those that integrate VoIP technology to allow the shopper to go beyond the graphics and text of HTML and click on a "talk to someone" button at the bottom of the page. Integrating voice services into the stream of IP data with the PC as communications device (rather than making users call the company on the conventional telephone) means that the sales transaction can now move along as a multimedia experience; the salesperson leads customers through the site to areas of interest, makes suggestions, and helps them complete their order. Marketing professionals can barely contain their excitement at the thought of such a sales channel, as it combines the ability to provide large amounts of information to the customer with a personal touch.

Enhanced ACD Functionality

Customer service centers have used automatic call distribution (ACD) functionality for years. Customer service representatives use computers connected to virtual private networks to access corporate databases about the clients that call. In addition, when clients call a customer service agent, the call could be routed from the PSTN to a call manager in an IP PBX at headquarters, then routed across the corporate network to a call center, or possibly even through the Internet and a VPN connection to an agent in another city, or country. Such an arrangement would make it easier for a company to treat its ACD resources as a single integrated pool, thereby enhancing its bottom line and improving productivity.

Much of the recent growth in commercial IP telephony deployments has been in call centers, and the combination of applications like customer relationship management software and incoming caller ID to create specific screen-pops. With VoIP, the phone set, PC, and various applications can be integrated, delivering better features and functionality to call center staff. Call centers can be virtualized, often locating personnel in different locations (including their homes), while maintaining or enhancing functionality. Intra-call center transfers can be eliminated, often justifying the investment in organizations with multiple call center locations and high intra-center transfer traffic. All of these result in improved customer care.

This is the ultimate in convergence technology. The marriage of the PSTN (even if it were a mobile call) with an IP network using VPN technology occurred seamlessly.

Voice Mail over IP Leads to Unified Messaging

From the network's point of view, digitized voice stored on a disk is just another data file to be transferred from one place to another. That is, once stored, its sensitivity to network delay almost completely disappears. When you receive your voice mail do you mind if it takes an extra two seconds? Yet voice mail retrieved from afar over traditional telephony networks is treated by those networks with all the reverence and concern for delay that an interactive voice conversation is, and it incurs the same costs. It would make much more sense to treat stored voice—and any voice data that has no real-time urgency, such as broadcast audio—like the error-sensitive, delay-insensitive data that it is.

RFC 1911, entitled Voice Profile for Internet Mail, sets forth a standards-based approach to IP-based voice mail. The next logical step in this migration strategy would be to have your email read to you through the phone after having been stored there, or accessed by the call manager. Integrating messaging services—voice, email, bulletins, pages—would help eliminate the need for multiple storage facilities required to reach someone. Imagine that regardless of the means used, the "message" is received on the phone, or forwarded to an email account.

VoIP Conclusions

Who is best positioned to make money from implementing a VoIP network? National and international ISPs and telecommunications service providers (the RBOCs and IXCs) with an existing IP network would benefit immediately.

However, the service providers must innovate to take advantage of this new way to move voice traffic. Managed services offer great potential for carriers. Today we see several managed VPN, storage, and security offerings, in addition to the traditional managed network services. As noted, no one manufacturer of voice technology is building the next-generation time division multiplexing PBX. Customers and carriers want improvements in productivity, efficiency, and cost. Moving to a single network infrastructure that can carry all types of traffic in IP packets with QoS makes sense.

It is not about the technology. Technology is merely a tool that enables services and work to occur. IP telephony is a tool that provides a path for equipment vendors, service providers, and businesses to innovate and develop new, integrated offerings that improve efficiencies, profitability, and perhaps even the quality of life for many who use it. Unfortunately, IPT is still considered a disruptive technology to those who believe only the PSTN can carry voice. However, IPT is on the verge of reshaping the entire telecommunications industry. It will not be a revolution but an evolutionary path that will take time to mature.

References

TeleGeography. (2006). U.S. VoIP Revenue & Subscribers Post Triple Digit Gain in 2005. Retrieved May 11, 2006, from *http://www.telegeography.com/press/releases/2006-03-07.php*.

About the Authors

Mark Steinberg

Mark Steinberg has spent more than 20 years in the business of technology. The telecommunications industry has been experiencing dramatic changes over the past three decades. Mark often provides consulting services regarding the strategic implications of technological change.

In addition to teaching the Hill Associates core curricula, he facilitates interactive discussions with a variety of audience types. These have included middle and senior management of Hill Associates' client companies and training partners. Over the past several years, he has used his expertise in communications and business strategy to create and deliver courses that focus on the issues facing companies involved in the communications marketplace. Business leaders rely on Mark's insights as they determine their business and technology strategies. His knowledge of convergence, both voice and data and wireline and wireless, has allowed him to create a variety of programs that address these business strategies. His engagements have been with companies such as CANTV, Telstra, AT&T, Qwest, Global Crossing, BellSouth, Verizon, Cingular, Ernst & Young, and Sprint—in the U.S., South America, Europe, Asia, and Australia.

A dynamic and energetic presenter with a broad perspective on technology and business, Mark is capable of dealing with technology and learning challenges in the classroom. Mark often combines real world analogies and humor to help explain the technical details. He has developed educational programs for technical and non-technical audiences that range from five days to over sixty days. A Senior Member of the Technical Staff, Mark has been with Hill Associates since 1994.

Mark is a Certified Information Systems Security Professional (CISSP #49990). He holds an MBA in International Finance from the University of Santa Clara, CA, a BA in Pre-med/Biology from Hartwick College in Oneonta, NY, and has done graduate work in Astronomical Physics at the University of Colorado, Colorado Springs, CO. He was previously a consultant to TeleCommunications, Inc. (TCI, now Comcast) and a Product Manager at both McDATA Corporation (now part of EMC) and Honeywell's Solid State Electronics Division.

Mark's hobbies include travel, photography, sailing, skiing, and theater.

Peter Southwick

For over 25 years, Peter has been in telecommunications - designing, implementing, and training on voice and data systems. He specializes in the hands-on training of wide area technologies, network design, telephony, ISDN, SS7, LANs, and SNA. As a Senior Member of Technical Staff at Hill Associates and a key member of the workshop group, he has been instrumental in creating courseware and exercises that allow participants not only to see and hear about the technology, but set up and test real-life customer scenarios.

Before joining Hill Associates in 1990, he was a Senior Systems Engineer at GTE Government Systems Corporation. Peter is a contributing author to Telecommunications: A Beginner's Guide (McGraw-Hill/Osborne) and The Handbook of Local Area Networks (CRC Press) and coauthor of ISDN: Concepts, Facilities, and Services (McGraw-Hill). His speaking engagements have included ComNet and ICA/SUPERCOMM.

Peter holds a B.S. in Electrical Engineering from Clarkson University. He is a member of IEEE and holds the Cisco Certified Network Associate and Cisco Certified Design Associate certifications.

His hobbies include sailing, skiing, mountain biking, and scuba diving.

About Hill Associates, Inc.

Hill Associates is a premier training and marketing services firm specializing in the field of telecommunications. For over 20 years, we have demystified the complex world of data networking, network security, and telephony. Our clients leverage our expertise, content, and development capabilities to create customized solutions that bring benefit and deliver results. Today, in addition to knowledge-based technology courses, we offer solution selling, a variety of e-learning options, and hands-on workshops. Our Marketing Services bring industry and subject matter experts to your sales events, customer meetings, Webinars, and white papers.

At Hill Associates, our consultative approach addresses each client's unique needs and challenges. Whether it's executive consulting, training for field teams, network engineers or corporate IT, or marketing programs designed to bring knowledge to your prospects, Hill Associates delivers your solution. Come visit us at www.hill.com and see what we can do for you.

.

Hill Associates, Inc.
Attn: Copyright Coordinator
106 Highpoint Center
Colchester, VT 05446
USA

Phone	802-655-0940
Fax	802-655-7974
E-mail	info@hill.com
Web Site	www.hill.com