PHIL SHERBURNE AND
CARY FITZGERALD, CISCO

You Don't Know Jack About VoIP

The communications they are a-changin'.

Telecommunications worldwide has experienced a significant revolution over recent years. The long-held promise of network convergence is occurring at an increasing pace. This convergence of data, voice, and video using IP-based networks is delivering advanced services at lower cost across the spectrum, including residential users, business customers of varying sizes, and service providers.

One of the key technologies driving this convergence is VoIP (voice over IP), which has evolved from what many viewed as experimental to a fundamental technology on which businesses from small to Fortune 500 are running their enterprises. VoIP has moved to a level of reliability and capability such that mainstream users are adopting it at a rapidly increasing pace. For this to happen, a number of technical innovations were required to solve issues such as quality of service and reliability.

This article explores key principles and technology

innovations underlying VoIP, and describes the implications of these innovations for software developers.

FROM ANALOG TO VOIP

Telecommunications technology is entering its third wave with VoIP. It began with analog signals carried by the first telephones and evolved into digital networks decades later. Now, with the increasing sophistication of the Internet, VoIP is coming into its own.

ANALOG NETWORKS

From the invention of the telephone in 1876 to today's modern communications infrastructure, voice has been carried by analog wave signals. Human speech is an analog wave signal. In the initial telephone networks, speech was converted to electrical wave forms (microphone) and converted back to speech at the other end of the con-

You Don't Know Jack About VoIP

versation (speaker), traveling the distance between the phones as this analog wave form.

While an obvious leap forward over previous methods of communication, this early technology had severe limitations that included introduction of “noise” in the signal. This noise increases with distance traveled. Although various methods of reducing noise were developed over the years, it remained a noticeable problem (remember the amount of static on long distance calls?). Another significant problem was one of economics. As the demand for communications increased dramatically post-World War II, the need to increase the carrying capacity of a pair of copper wires was significant. This led to the development of digital transmission capabilities in the long distance network.

DIGITAL NETWORKS

The early 1950s saw the introduction of technology that converted speech into digital signals. Specifically, the invention and deployment of T1 lines allowed for transmission of voice at 1.544 megabits per second (Mbps). (This is referred to as E1 in Europe and other places outside the United States, at a rate of 2.044 Mbps). Among other benefits, T1 lines addressed the two primary problems with analog voice transmission—noise and economics. Because the digital signals contained either 0s or 1s, the digital “repeaters” that were used to regenerate the signals over distances could also re-form the signals in a near-perfect rendition of the original. Thus, the impact of distance on the quality of the speech was virtually eliminated. The issue of economics was alleviated since a T1 line carries twenty-four 64-Kbps channels (32 for E1 lines). The mechanism to place multiple calls on the T1 line is known as TDM (time division multiplexing).

T1 technology (as well as higher digital transmission rates, including fiber optics) has been deployed extensively over the past 50-plus years. With the exception of the access portion of the network (i.e., lines from homes to the phone company), virtually all voice is carried over digital lines worldwide. Again, a great step forward in communications but with its own set of limitations. Critical among these is the nature of TDM connections,

known as circuit-switched connections. Fundamentally, this means that a call from one end of the circuit-switched connection to the other always follows the same path through the network and consumes the same amount of bandwidth, whether there is useful data to be transmitted or not. For example, during a silent pause on a phone call, 64 Kbps of data are still being transmitted in each direction. From an economic viewpoint, this is obviously inefficient.

VOIP

Since the 1970s there has been an increasing use of packet networks for transmitting data. Today the most obvious use of this technology is the Internet. The nature of packet networks in general and IP (Internet protocol) in particular is that the data to be transmitted is split into small packets that include small amounts of address information added to each packet. These packets are sent out over the network—quite possibly taking different paths through the network, unlike a legacy TDM connection where data is simply data, and routing is established at call setup time. The packets are then reassembled at the destination node.

Packet-switched networks have significant advantages over circuit-switched networks. Among these is the ability of packets to take different routes through the network. In the case of network failures (transmission lines being cut, etc.), this allows the data still to reach the destination. In addition, the only bandwidth used is that required for useful data (other than a small amount of control information such as address bits).

People recognized that if a means could be found to use packet technology for the transmission of voice, then the limitations of TDM networks could be overcome. Voice packets could take different routes through the network, and only necessary bandwidth would be used rather than always transmitting even in the face of silence. Even more significantly, both data and voice could be carried on a common, packet-based network. This would simplify management by reducing the number of networks to manage, and lowering network facility and hardware costs.

By the early 1990s certain fundamental technologies were developed that allowed for initial efforts in VoIP. The rest of this article outlines those technologies and their implications for software developers.

TECHNOLOGY OVERVIEW

Any discussion of VoIP must begin with a discussion of both *bearer* and *signaling* components. Bearer refers to the

actual voice being sent over the network. Signaling refers to the information necessary for successful setup and tear-down of the call. This includes the dialed digits, off-hook and on-hook information, originating number, etc. The separation of signaling from bearer information began in the circuit-switched digital networks—for example, ISDN. The concepts behind this were leveraged for VoIP.

BEARER TRANSMISSION

One difference between data and voice transmission is the sensitivity to delay associated with transmission across the network. Data is far less sensitive to delay than voice is. Anyone who has experienced an international call over satellite will recognize this sensitivity. This is partially solved by the use of RTP (realtime protocol) for the transmission of voice.

RTP is the standard protocol designed for realtime sensitive data transmission. Because of the realtime nature of voice, all VoIP traffic is carried as RTP packets. RTP “rides” on top of the standard UDP (user datagram protocol) and provides information to the endpoints not available in UDP. Specifically, RTP provides packet sequence information so endpoints can determine arrival order and time-stamping to allow endpoints to help manage “jitter” (discussed later in this article).

Voice coding standards. A number of different voice-encoding algorithms—codecs—are used in VoIP networks. These are standardized as a set of G-series recommendations by the ITU (International Telecommunication Union). Common ones are G.711, which encodes at 64 Kbps, and G.729, which encodes at 8 Kbps. Each of the codecs has different attributes, including compression

level, quality, etc.

Considerations for different bearer traffic. Although the discussion so far has focused on voice, in reality, other types of information are transmitted over traditional voice networks. For VoIP to be practical and gain common usage, these types of traffic must also be handled effectively:

- **DTMF** (dual-tone multi-frequency). This refers to the tones generated by a common touch-tone phone. These are used for not only initiating a phone call but also communicating during a phone call—such as for voice-mail and IVR (interactive voice response) systems. When used for making a phone call, DTMF is part of the signaling information and not transmitted as part of the bearer information. When used mid-call, however, it is transmitted as part of the bearer data.
- **Fax.** The use of fax machines—although less common today than before the common use of e-mail—remains a critical form of data communication (such as in the legal profession). For broad market acceptance, VoIP networks and equipment must be able to handle traditional fax machines, given the large number deployed worldwide. The issue in handling fax on VoIP networks is that fax transmissions are much more sensitive to packet loss than voice is. Different methods (Fax Passthru and T.38 Fax Relay) have been developed to ensure successful fax transmission over VoIP.

SIGNALING METHODS

An architectural model has evolved within the VoIP industry (see figure 1). Just as with any reference model, specific products or protocols do not necessarily strictly

adhere to the model, but it has proven to be a useful framework for characterizing components and their roles. The model is a PSTN (public switched telephone network) gateway, with a set of interfaces looking into telephone networks and a set of interfaces looking into VoIP networks, but it equally applies to IP phones and other VoIP endpoints.

The heart of the system is the MGC (media gateway controller). An MGC is an “intelligent” end-

You Don't Know Jack About VoIP

point; it interacts with its peers to establish, modify, and destroy connections with its peers within a network. The manipulation of these connections results in various end-

user services: call establishment, features such as transfer and park and hold, and call forwarding. The MGC is the component that supervises calls and services from end to end. Often it is implemented as a highly reliable system component, so call-related information must be mirrored across a complex of MGCs.

The MG (media gateway) is responsible for the media interfaces to the PSTN and to the IP network. Typically, an MG is implemented with a complex of DSPs (digital signal processors) to lower system costs, but general-purpose processors are also sometimes used, depending on

Emerging Internet Voice Services: The Year Ahead

CHRISTOPHER D. LIBERTELLI,
FEDERAL COMMUNICATIONS COMMISSION

In the area of emerging Internet voice services, 2004 will be remembered as an inflection point. As more and more enterprise and mass market consumers adopt VoIP (voice over IP) services, the chart showing demand for VoIP looks like a hockey stick. But are we sitting on the cutting edge of the early-adopter blade or has consumer demand accelerated toward the hockey stick's handle? With approximately 26 million residential and small businesses using broadband Internet access, there is every reason to believe that we will face a period of accelerating demand for VoIP services in late 2004 and into 2005.

The FCC (Federal Communications Commission) will likewise need to accelerate its work to ensure that VoIP services remain free of uneconomic and innovation-stifling regulations. VoIP "un-regulation"—particularly at the state level—is critical if the nation is to reap the benefits of new products and services currently in the labs of countless entrepreneurs.

The FCC has been protective of a number of different types of VoIP services. Some providers, such as Vonage, are offering Internet voice services using the public Internet and a consumer's broadband connection to allow consumers to make calls to other broadband Internet users or to consumers connected to the PSTN (public switched telephone network). Many cable operators, on the other hand, are offering IP-based voice services using their private digital networks to interconnect with the PSTN and not using the public Internet at all to transmit voice services.

As we hockey-stick through 2004 and enter a world in which both the caller and the called party are increasingly likely to be using VoIP, cable operators and others will

handle an increasing amount of traffic "on-net" without ever handing a call off to the PSTN—thereby avoiding entirely a complex regulatory environment riddled with unsustainable subsidies. This is a good thing. The policy question presented by this migration of traffic to packet-switched networks is not whether the legacy subsidy system is sustainable—it is how much time will pass before it is completely eroded.

Off the PSTN, we are also seeing the development of peer-to-peer voice services. Pulver.com's FWD (Free World Dialup), for example, employs peer-to-peer technologies to allow those using the service to transmit calls to one another. The FCC is hard at work on this issue. In the area of peer-to-peer telephony, under Chairman Michael Powell's leadership, the FCC adopted the Pulver.com Order by a 3-2 margin. This order's primary purpose was to encourage investment and innovation in the case where the voice application rides entirely over the Internet in digital form. Specifically, the FCC looked closely at the FWD offering and concluded that it clearly fit the information service definition and could not be categorized as a telecommunications service.

Perhaps the most significant aspect of the decision was the FCC's announcement that VoIP services such as Pulver.com were subject to the exclusive jurisdiction of the federal government and that state economic regulation most likely cannot survive. With this decision, the FCC has begun laying the foundation for a "light touch" regulatory environment for Internet voice services at the federal level.

At the access layer, the FCC continues to work to bring alternative broadband Internet distribution networks to people in the United States. We are working with our colleagues elsewhere in the federal government and at the state and local levels to develop a sound policy framework in this area. The Triennial Review Remand proceeding, in which the FCC will work to craft alternative unbundling rules, will commence shortly. Working together, we hope to complete this rulemaking by year's end. Meanwhile, the FCC is actively searching for a way to make better use of license-exempt spectrum to

the application.

An MG is a simple endpoint. It does only what it is told to do. It does not understand signaling to either the PSTN or the IP network. It does not understand services or even calls. It creates, modifies, and destroys connections as instructed by an MGC. These connections can be between the PSTN and the IP network, between PSTN ports, or even between IP-based endpoints. Since an MG does not understand the end-to-end nature of a call, it needs to concern itself only with the connections it is holding up, so the system reliability requirements for an

MG can be somewhat relaxed.

MGCs and MGs interact with each other over a control plane, which can be a proprietary interface such as an internal API; standardized protocols have also been developed. Both the ITU and the IETF (Internet Engineering Task Force) saw the need for this protocol and cooperated to produce the MEGACO (H.248/Media Gateway Controller) protocol, a first-of-its-kind cooperation between the two standards bodies. It is published in the IETF as RFC3525 and in the ITU as H.248. H.248/MEGACO is in an early adoption phase of deployment. Eventually, this

provide another pathway into the home for VoIP providers.

At the same time the FCC adopted the Pulver.com Order, it released a comprehensive Notice of Proposed Rulemaking (NPRM) to examine how best to address VoIP. This IP-Enabled Services NPRM takes up a number of important policy areas. The NPRM examines some of the important definitional debates surrounding Internet voice services, with a view to existing definitions and how those definitions might apply to today's changing communications environment. The FCC is building a record to determine how social and public safety objectives can best be achieved when using IP-enabled communications.

A full and robust record will pave the way for the FCC to adopt policies that facilitate economic growth, a more secure homeland, and preserve and advance universal service and access to people with disabilities. There is near-unanimous agreement at the FCC that we must remain committed to universal service, law enforcement access, E911 (enhanced 911) capabilities, and access for people with disabilities. Against this backdrop, the FCC has hosted a series of Solutions Summits as it moves forward with the IP-enabled communications proceeding to address these important social and public safety policies.

But our work is not done. Still other petitions, including one filed by Vonage and Level 3 Communications, remain pending before the FCC. These petitions involve different varieties of IP-voice services, with different levels of digitization and interaction with the PSTN. The questions posed by these petitions highlight the telecommunications industry's continued struggle to produce a consensus plan for reform of the complex intercarrier compensation regime. In all these areas we begin with the principle of nonregulation of the Internet, because absent compelling justification for regulation, limiting government intrusions—both at the federal and state levels—maximizes the potential for innovation and increases opportunity for the nation as a whole.

VoIP will revolutionize the way consumers work and play.

The engineering practices that produced the PSTN were a proud national achievement, but the network these practices produced was not built for Internet time. It took 100 years of Ma Bell know-how to produce *69. The policy environment must recognize that VoIP services that allow users to dynamically assign new telephone numbers to nomadic workers and that translate voice-mail into e-mail, text messages into voice-mail, and spoken word into text are built upon an innovation platform that is fundamentally different than the SS7 (Signaling System 7) systems of old. In the past, it may have taken up to 18 months to introduce a new feature set into the closed architecture of a legacy SE circuit switch; but communications built on IP platforms can introduce new services as quickly as software developers can think of them. The country will be well served if we hasten the transition to "SS8."

If the FCC does its job right in the IP-enabled services proceeding and other proceedings, the legal and policy environment will foster and protect an open platform for software-based innovations in communications. The challenge for the FCC is to create the kind of environment in which these changes can flourish. Through the various petitions that are pending before the FCC, we hope to create a rational, federal policy environment for VoIP services free of legacy regulation. Doing so is the surest way to ensure that 2005 will be the year that VoIP services hockey-stick to more widespread adoption by consumers.

CHRISTOPHER LIBERTELLI is senior legal advisor to FCC Chairman Michael K. Powell. Previously he served as special counsel for competition policy in the Office of the Bureau Chief of the Wireline Competition Bureau and as an attorney-advisor in the Policy Division of the Common Carrier Bureau. Before joining the FCC in 1999, he served as an associate in the communications law firm of Dow, Lohnes, & Albertson in Washington, D.C. Libertelli received a J.D. from the Boston University School of Law, and a B.A. in international relations from Boston University.

You Don't Know Jack About VoIP

is expected to replace an earlier effort to standardize the MGC/MG control plane that has been known as MGCP (media gateway control protocol). MGCP (IETF RFC3435) has been deployed in a number of networks and has been adopted in the ITU for application in VoIP over cable.

MGCs interact with their peers using an intelligent signaling protocol. Two intelligent protocols have emerged—from the ITU, H.323, and from the IETF, SIP (session initiation protocol). They share many concepts; both suppose that the endpoints are intelligent—but they also differ in significant ways.

H.323 is derived from the PSTN protocols used to access PSTN services—Q.931. VoIP connections in H.323 follow the ISDN model: the same message sequences are used to establish and tear down calls. H.323 has been extended to support a number of services; again, these follow an established model from established TDM network architectures. For example, a number of services are described in the H.450 series; these are modeled on the corresponding services from Q.SIG.

SIP is a methods-based protocol, whose roots are in HTTP. In general, services are not explicitly exposed in the protocol; rather, the designer can use a set of well-defined methods to implement services. So, for example, SIP does not have a transfer primitive per se, but executing a set of SIP transactions will result in the user experiencing a transfer. A significant amount of work is going on in the standards communities with respect to SIP, as well as a significant increase in market adoption of SIP-based equipment. SIP-based equipment is clearly expected to achieve a significant share of installed VoIP equipment over the next few years.

A number of exciting new services and concepts are coming out of the VoIP community. We highlight just a couple as follows: the impact of IM (instant messaging) and presence on converged communications; and ENUM, a mechanism for telephone number resolution in VoIP networks.

Instant messaging and presence. A significant number of features in the telephone network are devoted to the concept of increasing the probability that a call will be completed to the right recipient at a time that is

acceptable to both the caller and the called party. IM and presence have recently emerged as important business and personal communications tools. Combining IM and presence with VoIP yields some valuable new features. Presence information can be used to determine whether offering a new call to a party is likely to be successful—there is no point in placing the call until the called party is available and willing to take the call. Instant messages can be used as part of the alerting process, which allows both called and calling parties to provide more information to each other about the nature of their communications. The two systems—VoIP and IM/presence—working in concert are more valuable than either one alone. VoIP deployments for these applications are in their very early stages.

ENUM. The best-understood and most widely deployed name resolution system today is the DNS (domain name system). In the DNS, names are written from right to left, with the most general part of the address on the right, and more specific names written to the left (e.g., www.ietf.org). In the PSTN, telephone numbers are written from left to right, with the most general part of the number written on the left and the more specific toward the right (e.g., 1.212.543.6789). ENUM calls for telephone numbers to be written DNS-style, rooted at the domain `e164.arpa`. So, 1.212.543.6789 becomes 9.8.7.6.3.4.5.2.1.2.1.e164.arpa. Interestingly, each digit is treated as a subdomain. This allows ENUM to ignore the nuances of country codes, city codes, etc. that vary broadly worldwide. When this address is queried, the DNS can return a specific IP address corresponding to the telephone number, or it can return a rule for rewriting the original number into some other form. For example, rules can be returned to rewrite 1.212.543.6789 as `sip:36789@nyc-gw.example.net`, `sip:caryfitz@service-provider.com`. ENUM offers the possibility to reuse the worldwide DNS for VoIP. ENUM is a standard set by the IETF as RFC3761.

MANAGING VOIP QUALITY OF SERVICE

Voice quality. The fundamental concern for VoIP QoS (quality of service) is voice quality. Unfortunately, objective measurements for this have been elusive. That said, the major factors that affect voice quality are delay, packet loss, and treatment at the endpoints.

Voice codecs are unevenly tolerant of packet loss, but loss above 2 to 5 percent will have a perceptible effect on quality. Loss is rarely random and is often associated with high jitter (simply defined as the variation in packet arrival times at the destination).

When one-way delay through a voice network exceeds about 150 milliseconds, natural conversational communi-

cation is strained, so most network deployments attempt to keep the delay well below that threshold. There are a number of components to delay: codecs have an intrinsic delay; it takes time to prepare and route a packet to the IP interface on the phone or gateway; various access networks have intrinsic delays; and transit networks contribute both in terms of routing delay and propagation delay.

Further, packets are generated at regular intervals, but because of the vagaries of routing across the IP network, they are delivered to the endpoint with a certain amount of jitter. Endpoints have built into the software a “jitter buffer” where packets are buffered and then played out at a constant rate. This, of course, works fine unless the amount of jitter exceeds what can be absorbed by the jitter buffer. Software-based mechanisms exist in the endpoints to automatically adjust buffer sizes, etc. as jitter increases or decreases. Jitter can be a major component of the delay budget.

QoS tools. The basic idea for controlling QoS revolves around two aspects: the first is ensuring that the network has enough capacity (bandwidth) to allow for high-quality calls; the second is establishing priorities such that the more realtime-sensitive packets are given higher priority for transit through the network.

To ensure enough capacity, there are mechanisms such as RSVP (reservation protocol, RFC 2750). This allows bandwidth to be reserved through the network. Using RSVP, the endpoints or MGCs signal through the network, reserving capacity. This is done in advance of a call being set up.

For priority management, the mechanism is different packet queueing methods within the MGs and routers. A variety of algorithms are available, with the best choice depending on the customer network and traffic types. Associated with this is the concept of TOS (type of service) bits. Within each packet there are three bits at the IP level that indicate up to eight levels of precedence. These are used to ensure that higher-priority packets make it through the network first.

SOFTWARE CONSIDERATIONS

System and software designers for VoIP equipment and networks face myriad challenges. Common concerns are QoS; security; manageability and operations, reliability, redundancy, and sparing; scalability; deployability, installability, and upgradability; serviceability, capacity management, fault detection, diagnostics, reparability, and metrics; testability and regressions; internationalization; performance and graceful degradation under fault conditions and load; extensibility; interoperability with both

IP and legacy systems; modularity; manufacturability and costs; open systems and standards compliance; ease of use for end users; consistent/normalized database use; billing and audits; and feature interactions.

RELIABILITY

Expectations for the reliability of VoIP are as high as those for the traditional voice networks. Although there are different measures of reliability (such as the oft-misused “5 9s”), for our discussion the assumption is that the VoIP system must work all the time (7 x 24 x 365). While there are occasional “maintenance windows,” the expectations are the system is always operational (think about a 911 call center in a major metropolitan area).

A number of **exciting new services and concepts** are coming out of the VoIP community.

At a high level this means the software designer must “design for failure”—that is, the designer must consider potential failures in three domains:

1. In the network, which might be caused by external events such as power failures or line cuts.
2. In the hardware, including processor, memory failures, etc.
3. In the software, which may be the result of bugs, corrupted data, etc.

Thus, the software designer must include capabilities such as:

- **No upgrade downtime.** This typically implies some form of duplicated active/standby systems with data synchronization between the active/standby systems.
- **Software audits.** There should be separate software components that audit the primary system software. This includes validating the internal data structures for accuracy as well as consistency among data structures. Corrective action may include automatic correction of the invalid data.
- **Process monitoring.** This means having system

You Don't Know Jack About VoIP

monitors that ensure the primary system software is operating correctly. This includes techniques such as watchdog timers—that is, having the primary software send a message on a regular basis to the monitoring software indicating proper functioning. Corrective action by the system monitor may range from process restart to system failure over to a standby system.

- **Automatic failover.** As a response to certain types of failures—including a full system failure—the system automatically fails to a standby system.
- **Geographic redundancy.** This is the ability to have the active and standby systems separated by hundreds of miles.

MANAGEABILITY

The need to manage their networks is critical for all customers, whether large or small. In many cases the cost of operating and managing voice systems—whether traditional TDM or VoIP—far outweighs the cost of the equipment. Therefore, the need for effective tools allowing for cost-effective management is important for successful deployments.

Manageability, as used here, covers many different areas, including accurate and flexible billing systems, error reporting and resolution, call tracing, adds/moves/changes, etc. Although VoIP does not create new concerns, manageability takes on additional roles.

Consider the need for call tracing, which typically arises when an end user complains about a dropped call, noisy lines, etc. A system administrator will then typically look at the call traces—the route the call has taken through the network—to identify the source of the trouble. As noted earlier in the context of a traditional TDM circuit-switched network, when a call is set up, the voice takes the same path through the network for the duration of the call. This makes tracing calls through the network reasonably straightforward by collating call detail records, etc.

In a VoIP network, the packets containing the voice may take very different routes through the network, which makes the issue of call tracing and diagnosing of intermittent problems much more challenging. This

requires not only good instrumentation on the MGCs, MGs, and routers in the network, but also very sophisticated management tools that provide the correlation and reporting of the information.

THE IMPACT OF VOIP

In recent years, we've seen increasing adoption of VoIP networks for customers of varying sizes on a global basis. The cost advantage resulting from convergence and the value of new applications offered by this convergence are the primary drivers of this adoption. With this comes the need for increasingly sophisticated systems and management tools to allow for the extensive adoption and deployment of VoIP.

VoIP's increasing adoption will have a significant impact on our communications and the products that provide those communications. Therefore, software developers across the industry will increasingly need to be aware of and understand the challenges that come with this latest change in the communications infrastructure. ☐

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

PHIL SHERBURNE is senior director of the Voice Technology Group, Cisco Systems. He runs the Call Control Division, and his team is responsible for the development and deployment of Call Control technologies including the Cisco Call Manager, BTS 10200, PGW 2200, and the SIP Proxy Server products. Previously, he was general manager for the Packet Telephony Call Control Business Unit responsible for the Softswitch products from Cisco. During his career with Cisco, he has been involved with a number of VoIP products and offerings.

Prior to joining Cisco in 2000, Sherburne spent more than 20 years at AT&T and Lucent Technologies Bell Laboratories, where he was involved in development of both PBX and messaging products. He has a B.Sc. in computer science from the University of Oregon and an M.Sc. in computer science from Ohio State University.

CARY FITZGERALD is senior director of the Voice Technology Group at Cisco Systems. He joined Cisco in 1996 and formed the team that built the first commercial VoIP gateway. He is a key contributor setting Cisco's VoIP architectural directions. Prior to joining Cisco, FitzGerald was a distinguished member of technical staff at AT&T Bell Laboratories, where he led architecture and design teams for voice-response and voice-mail systems. He has a B.S. in computer science from Purdue University.

© 2004 ACM 1542-7730/04/0900 \$5.00