

Packet Voice

Because We Need to Talk
by Linda Kennard



Despite the widespread use of data networks, the ability to transmit written words will never eliminate the need to hear and express a voice. Sometimes you just need to talk. Frequently, talking requires using the telephone, and almost as frequently, using the telephone incurs long-distance charges. Of course, telecommunications companies (commonly called *carriers*) usually offer companies special rates, charging well under 10 cents per minute for domestic long-distance calls and as little as 50 cents per minute for international calls. But when a company has hundreds of employees placing 10- to 15-minute calls each day, 5 cents here and 50 cents there adds up quickly. Consequently, companies are looking for ways to minimize long-distance costs.

One way to minimize long-distance costs seems to be packet voice. *Packet voice*, also called *voice over packet*, refers to the practice of routing telephone calls and faxes over an existing data network rather than over the traditional Public Switched Telephone Network (PSTN). (See Figure 1 on p. 10.) As you know, a data network transports messages from point A to point B in units of information called *packets*, which typically contain data portions to which headers are appended as the packets traverse the network. Replace the data portion of the packets with digitized and usually compressed voice, and you have *packet voice*. (For more information about how packet voice works, see "Where the Lines Cross" on p. 8.)

The promises associated with packet voice sound appealing. Marketers and press members promise that whether your voice is smooth, gruff, high, or squeaky when it is delivered over a circuit-switched network (such as the PSTN), your voice will sound the same when delivered over a packet-switched network. Assuming you have spare bandwidth on

your company's WAN, marketers and press members would have you believe that you can piggyback voice over data networks for free.

Furthermore, some marketers and press members like to imply (and sometimes state outright) that packet voice actually reduces management complexity. After all, with packet voice, companies have only one network to manage. Today companies have at least two networks to manage—one for voice and one for data.

Marketers and press members make these promises regardless of the specific packet voice option they are discussing: voice over Internet Protocol (VoIP), voice over frame relay (VoFR), or voice over Asynchronous Transfer Mode (ATM), which is called *voice and telephony over ATM (VToA)*.

These promises are all very alluring—but are they legitimate? *NetWare Connection* researched packet voice technologies and here's what we concluded: Regardless of which packet voice technology you choose, you can expect good quality and long-term savings. However, nothing is failproof or free, and packet voice is no exception.

SUPERIORITY COMPLEXES

Of the three packet voice options, VoIP certainly attracts the most attention but not the most customers. In fact, Sara Hofstetter, director of public relations for International Discount Telecommunications Inc., says that "only two or three Fortune 1000 companies have implemented VoIP on a corporate level."

According to Tom Jenkins, TeleChoice consultant, companies are extremely interested in VoIP, but they are

Inside This Article

The Packet Voice Market p. 7

Where the Lines Cross (How Voice Packet Technologies Work) p. 8

Quality—How Good Is "Good"? p. 10

Cost—How Free Is "Free"? p. 12

Complexity—How Easy Is "Easy"? p. 14

Voice Over Frame Relay: Let's Talk Options p. 14

Voice and Telephony Over ATM: Let's Talk Options p. 16

What's MOS Got to Do With It? (Rating Voice Quality) . . p. 18

approaching this relatively new technology with extreme caution. (TeleChoice is a market strategy firm that consults service providers and equipment vendors in the telecommunications industry.) According to Jenkins, companies have several reasons for being cautious: For example, companies are concerned about voice quality over public IP networks. Companies are also concerned that deploying a VoIP solution may increase management costs, which, in turn, may negate the potential savings associated with packet voice.

Furthermore, VoIP-related technologies that may improve voice quality and ensure call features (such as voice mail and call forwarding) are changing rapidly. Consequently, companies are concerned that if they "purchase something today, it will be outdated tomorrow," Jenkins says. The result of these concerns, concludes Jenkins, is "decision paralysis."

Despite these concerns, the ubiquity of IP networks makes VoIP an option worth considering and an option *NetWare Connection* discussed at length in the February 1999 issue. (See "Sound Solutions: Defensible VoIP Options for Companies," pp. 6-18. You can download this article from <http://www.nwconnection.com/past>.)

THE PACKET VOICE MARKET

The attraction to VoIP stems in part from speculations about its future. Both providers of packet voice solutions and telecommunications consultants agree that VoIP may ultimately win the largest share of the corporate market for packet voice.

Despite this prediction, VoIP isn't the standard yet. Today, VoFR and VToA have a larger market share than VoIP, although the percentage of companies using packet voice at all is still quite small. Bill Callahan, services vice president for networked computing at AT&T Solutions, suggests that only about 10 percent of U.S. companies are using VoFR or VToA today. Callahan suggests that this estimate is high, and perhaps it is. However, 10 percent is substantially higher than the estimated 1 percent using VoIP.

Between VoFR and VToA, VoFR has a larger share of the market simply because more companies have frame relay networks than ATM networks. (ATM has the fortunate reputation of being an

excellent network transport service but the unfortunate reputation of being complex and expensive.) Although there are more frame relay networks carrying voice if you count sheer numbers, companies using ATM networks are more likely to implement a voice packet solution than companies using frame relay networks. Jenkins speculates that roughly 20 percent of frame relay

networks currently carry voice packets. In contrast, Jenkins conjectures that as many as four out of five of ATM networks, or 80 percent of ATM networks, currently carry voice packets.

MAKE THE CONNECTION

In addition to having a general unease about VoIP, more companies opt for VoFR or VToA because frame relay

Kyocera 1/2 Page Island AD

4 7/8" x 7 3/8"
(4.875" x 7.375")

Where the Lines Cross

All packet voice technologies follow the same basic model. (See Figure 1 on p. 10.) Packet voice equipment sits on the edge of a transport network, which might be frame relay, Asynchronous Transfer Mode (ATM), or IP. Packet voice equipment might be a voice-enabled frame relay access device (VFRAD), a router or switch, a concentrator, or even a Private Branch Exchange (PBX). (This PBX would have to be equipped with the software and interface board necessary to direct calls over the data network to which it is attached.)

When you dial a number, the call is routed through the packet voice equipment servicing the PBX to which your telephone (or fax machine) is attached. This packet voice equipment converts signals received from your telephone into a packet form suitable for transport over the network. The packet voice equipment also recognizes traditional voice signals such as the number dialed, which the equipment would then convert to an IP, frame relay, or ATM address.

The voice packets then make their way across the packet network to the packet voice equipment serving the telephone or fax machine at the destination you originally dialed. This packet voice equipment recognizes the traditional signal embedded in the packets and directs the telephone call accordingly. The packet voice equipment also converts the packets back into a form that the telephone or fax machine you dialed can accept.

VOICE CODING

All of the packet voice equipment on a packet network needs to understand and to use the same voice coding/decoding (codec) method(s). The International Telecommunication Union (ITU) has standardized several codec methods, including G.711, G.726, G.728, and G.729. (The ITU is an organization within the United Nations Educational, Scientific, and Cultural Organization, more commonly known as UNESCO. For more information about ITU, visit <http://www.itu.int>.)

All of these standards describe different methods for converting voice into digital and sometimes compressed signals. These different approaches are based on the following:

- Pulse Code Modulation (PCM)
- Adaptive Differential PCM (ADPCM)
- Linear Predictive Coding (LPC)

PCM

PCM converts voice into digital form by sampling a voice wave signal 8,000 times per second, waiting only 125 microseconds between each sample. Each sample is then converted into digital code that represents the amplitude of the sample wave signal. Over the Public Switched Telephone Network (PSTN), this code is 8 bits, so PCM voice consumes a total of 64 kbit/s:

$$8 \text{ bits} \times 8,000 \text{ samples per second} = 64,000 \text{ bit/s or } 64 \text{ kbit/s}$$

The G.711 standard describes PCM, which does not affect voice quality in any way. PCM-encoded voice can be interchanged between packet voice, PSTN, and PBX networks.

ADPCM

ADPCM also converts voice by sampling the voice wave signal 8,000 times per second and converting that signal into digital code that represents the amplitude of the sample. However, ADPCM uses only 4 bits for the digital code. As a result, ADPCM voice consumes a total of 32 kbit/s:

$$4 \text{ bits} \times 8,000 \text{ samples per second} = 32,000 \text{ bit/s or } 32 \text{ kbit/s}$$

ADPCM can be reduced to as little as 16 kbit/s by using a 2-bit digital code. As the bit size of the digital code becomes smaller, however, this code is less like the original, and voice quality can suffer.

The G.726 standard describes ADPCM coding at 40, 32, 23, and 16 kbit/s. ADPCM-encoded voice can be interchanged between packet voice, PSTN, and PBX networks if the PBX networks are configured to support ADPCM.

LPC

LPC mathematically analyzes a voice wave and then transmits a small set of parameters that describe the nature of the wave. A receiving LPC device reproduces the voice wave based on the parameters the device receives. This set of parameters can be quite small, as low as only 2,400 bit/s or 2.4 kbit/s.

The G.728 and G.729 standards describe versions of what is called codebook excited linear predictive (CELP) coding and require only 16 kbit/s and 8 kbit/s of bandwidth respectively. CELP coding cannot be interchanged with packet voice, PSTN, and PBX networks. Before going over the PSTN or a PBX network, CELP coding must first be transcoded into a format suitable for delivery across these networks. Transcoding can negatively affect voice quality.

COMPRESSING VOICE AND SUPPRESSING SILENCE

In addition to compressing voice, some packet voice equipment can compress silence and, thus, minimize the amount of bandwidth that voice traffic consumes. In a white paper titled "A Discussion of Voice Over Frame Relay," the Frame Relay Forum (FRF) states that only 22 percent of a "typical dialog consists of essential speech components that need to be transmitted for complete voice clarity." (See Figure 2 on p. 10.) (You can download this white paper from <http://www.frforum.com/4000/4017.html>.)

Because codecs do not need to packetize silences, codecs can reduce the amount of bandwidth that voice consumes by as much as 56 percent based on FRF estimates. (For more information about silence suppression, see "Sound Solutions: Defensible VoIP Options for Companies," NetWare Connection, Feb. 1999. You can download this article from <http://www.nwconnection.com/past>.)

For example, if you use G.729 to compress voice to 8 kbit/s and you also use silence compression, you can reduce the amount of bandwidth voice traffic consumes to only 4 kbit/s—cutting in half the bandwidth your voice traffic consumes. The network can then support twice the number of calls that G.729 alone would enable. And the resulting voice quality? "A lot better than that cellular telephone you're using," says Don Choi, working group chair for the voice and telephony over ATM (VToA) subcommittee in the ATM Forum. ●

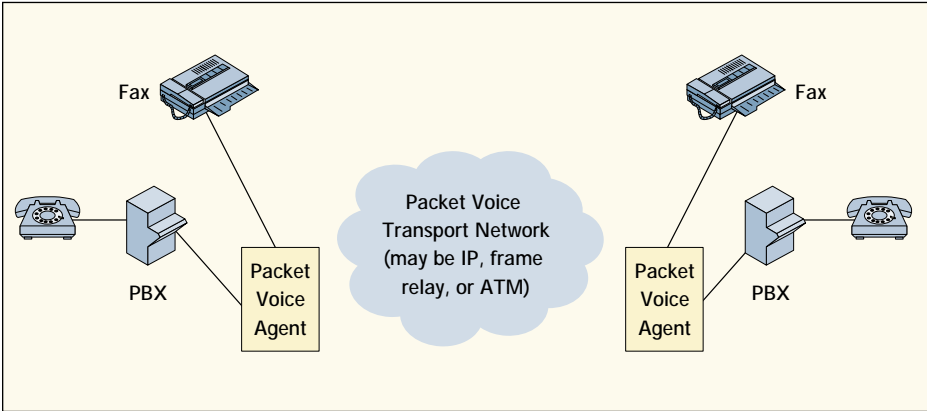


Figure 1. All packet voice technologies follow the same basic model.

and ATM are actually better suited than IP for carrying voice. IP is a connectionless network service as opposed to frame relay and ATM, which are connection-oriented network services.

As a connectionless network service, IP transmits data without first establishing a path between the source and destination stations. Consequently, each data packet is independent and must include source and destination addresses, which must be processed by the routers and other network nodes the packet encounters. Furthermore, packets headed for the same destination can travel different routes, and these packets can arrive at different times and in a different order.

In contrast, frame relay and ATM are connection-oriented network services. Connection-oriented network services transmit data only after establishing a path through the network between the source and destination stations. Packets headed for the same destination travel the same route and are delivered in the order in which they were sent.

As connection-oriented network services, frame relay and ATM are similar to the circuit-switched networks that typically carry telephone calls. Like circuit-switched networks, frame relay and ATM networks establish a connection between your telephone and the telephone you're calling. Your voice and the voice of the person you're calling then go back and forth across that connection—whether established by frame relay, ATM, or PSTN—until you hang up. Because connection-oriented networks are so similar to traditional telephone networks, they are "inherently superior to connectionless networks" at carrying voice, says Don Choi, working group chair for the VToA subcommittee in the ATM Forum. (The ATM Forum is an international nonprofit organization

committed to promoting the use of ATM products and services.)

QUALITY—HOW GOOD IS "GOOD"?

Regardless of which packet voice technology is "inherently superior," each of the technologies have the potential to deliver high-quality voice. Of course, the quality of packet voice depends on several factors. For example, if the network you are using doesn't have enough bandwidth and, consequently, becomes overly congested during peak times, or if the network doesn't have quality of service (QoS) mechanisms that ensure minimal delay and delay variation (called *jitter*), then the "quality can just stink," says Jenkins.

Melanie Hanssen, vice president of marketing for the Frame Relay Forum (FRF), agrees. (The FRF is an association of vendors, carriers,



users, and consultants committed to promoting frame relay.) Voice traffic is extremely sensitive to delay, Hanssen explains. "So it's a good idea to make sure your service provider has guarantees on maximum delay in the network." Most major providers do offer guarantees on maximum delay, but on networks without such guarantees, voice quality could suffer, says Hanssen.

Delay, Delay—Is There an Echo in Here?

The maximum round-trip, end-to-end delay you can afford for voice transmissions is no more than 250 ms, according to Daniel Minoli, who has researched packet voice for more than 20 years. (See Daniel Minoli and Emma Minoli, *Delivering Voice Over Frame Relay and ATM*, New York: John Wiley

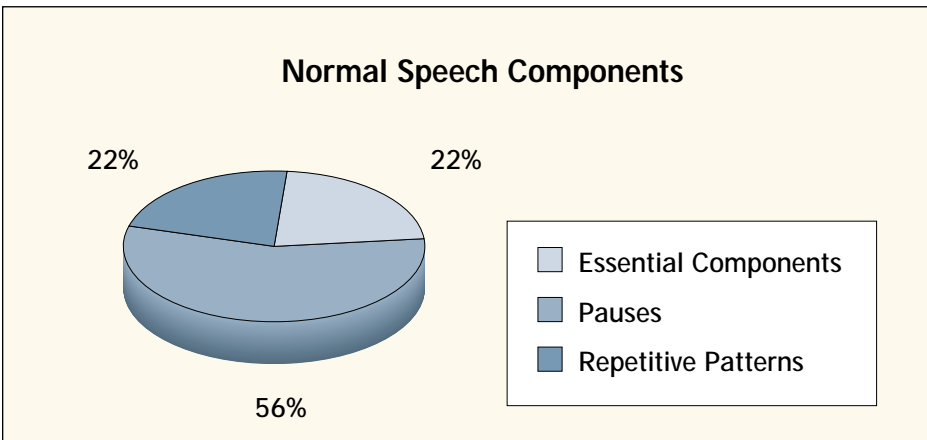


Figure 2. More than 50 percent of normal speech is silence (figure courtesy of the Frame Relay Forum [FRF]).



& Sons Inc., 1998, p. 206.) Network delays that extend to more than 250 ms are at best annoying and at worst can cause problems that affect voice quality. For example, network delay can cause an echo. When an echo occurs, you can hear your own voice, usually at about the same time the person with whom you are speaking begins to respond.

An echo becomes perceptible when delays exceed 20 ms. To ensure good voice quality without an echo, the network or the equipment you use on your sites must support the ability to erase or cancel that echo.

E-Quality

Of course, if you are considering running voice over ATM, you have less to worry about in terms of delay than if you are using frame relay. As you may know, ATM was designed to support all types of traffic—including voice—and as a result has intrinsic QoS guarantees that define various levels of service.

For example, ATM defines a constant bit rate (CBR) service that specifies a peak cell rate, a cell loss ratio, and a cell delay variation. (A *cell* is the unit of information, or packet, that ATM networks carry.) The ATM CBR service delivers traffic with very low delay and low jitter. Because specifications such as CBR are built into ATM, explains Jim Daugherty, general manager in data services marketing at AT&T Solutions, “you can select QoS that provides the delay and jitter you need so you can know exactly what performance you’ll get.”

In contrast, frame relay does not have intrinsic QoS guarantees, so the equipment you use on your network and the equipment on your carrier’s network

must support traffic prioritization and QoS techniques that can ensure low delay and jitter for voice traffic. (For more information about traffic prioritization, see “Traffic Problems? Making Way for Important Network Packets,” *NetWare Connection*, July 1999, pp. 6–21. You can download this article from <http://www.nwconnection.com/past>.)

Assuming you are delivering voice over a well-designed ATM or frame relay network, you won’t be able to tell the difference between voice carried over that network and the voice you hear when you call your mom from home.

In addition, you should look for equipment that supports the FRF standards FRF.11 and FRF.12. (For more information about FRF.11 and FRF.12, see “Voice Over Frame Relay: Let’s Talk Options” on p. 14.) FRF.11 and FRF.12 “go a long way toward catching up” with ATM in terms of defining ways to deliver voice with good quality over frame relay, says Byron Henderson, director of marketing in Cisco System Inc.’s enterprise line of business.

In fact, the quality of voice over frame relay is considered comparable to the voice quality over the public telephone network, says Callahan. As you would expect, the quality of voice over ATM “is excellent,” says Choi.

Jenkins agrees that you can expect very good quality voice with both VoFR and VToA. In fact, Jenkins says that the quality will reach or even extend above “a 4.0 on the mean opinion scale” (MOS). (For more information about MOS, see “What’s MOS Got to Do With It?” on p. 18.) In other words, assuming you are delivering voice over a well-designed ATM or frame relay network, you won’t be able to tell the difference between voice carried over that network and the voice you hear when you call your mom from home.

COST—HOW FREE IS “FREE”?

Voice quality is not usually the first concern that companies interested in packet voice discuss with potential service or equipment providers. Daugherty says that when AT&T customers want to talk about packet voice, the discussion “typically starts with ‘I want to save money by putting voice over packet.’ ”

In fact, companies can save money with packet voice. When asked whether delivering voice over frame relay or ATM is really less expensive than carriers’ business rates, Jenkins’ reply fairly represents how most providers and consultants in this corner of the networking industry respond: “The answer is actually, ‘Generally yes.’ ”

The real question is how much less expensive? According to Jenkins, carrying voice over frame relay or ATM is anywhere from 20 percent to 60 percent less expensive than carrying voice over the PSTN at regular business rates. The exact percentage varies between companies.

The percentage point at which moving voice traffic to the data network becomes cost effective also varies between companies. For many companies, 20 percent savings is not enough to make converging their voice and data networks worth the worry. The point at which many companies believe convergence is worth the worry is generally 30 percent, says Jenkins.

Of course, before your company can begin to save money, your company first has to pay for the VoFR or VToA equipment that you’ll need to install on all

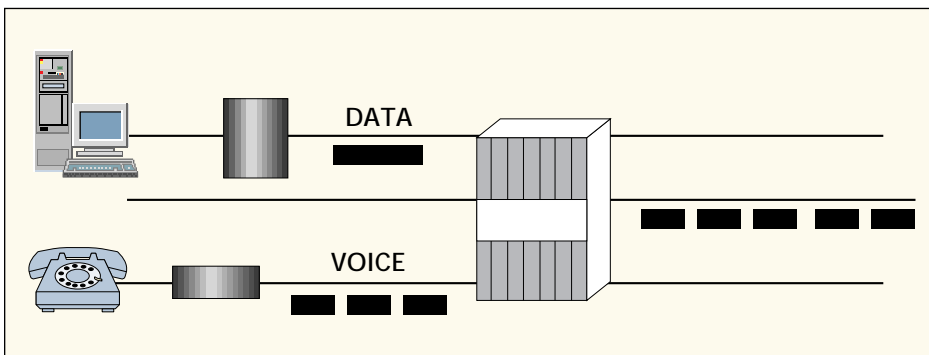


Figure 3. Equipment that supports FRF.12 fragmentation can break up large data packets to ensure that voice traffic is not subject to delay (figure courtesy of the Frame Relay Forum [FRF]).

of the sites between which packet voice will be exchanged. To deliver voice over frame relay, you can use a voice-enabled frame relay access device (VFRAD) or a voice-enabled router with a frame relay interface.

Prices for VFRADs range between U.S. \$1,000 and U.S. \$10,000, Hanssen estimates. However, as Henderson points out, some companies feel that it isn't cost effective to use both a VFRAD and a router. Consequently, many companies opt for a voice-enabled router, which could cost anywhere from U.S. \$2,500 to U.S. \$25,000, Henderson estimates.

To deliver voice over an ATM network, you can generally use a concentrator or a voice-enabled router with an ATM interface. The cost for ATM equipment is only "slightly more" than the cost of frame relay equipment, says Henderson.

To measure the relative cost (including equipment) and savings of packet voice, many companies measure the time it takes to recover the initial cost of deploying a packet voice solution. In other words, these companies measure

the return on investment (ROI). Henderson says that the ROI on packet voice solutions for sites within the United States is typically between 18

Although "there are clear benefits of carrying voice over ATM," says Choi, "the disadvantage is the cost of introducing complexity."

and 36 months. For transoceanic links, ROI can be as few as seven months.

COMPLEXITY—HOW EASY IS "EASY"?

From what you've read so far, packet voice probably seems like an attractive

alternative to the PSTN or to Private Branch Exchange (PBX) networks—maybe even an obvious alternative. After all, packet voice solutions can deliver toll-quality voice, cut long-distance bills by as much as 60 percent, and return their initial cost within as few as seven months.

If packet voice is so great, however, why doesn't every network administrator with bandwidth to spare on a WAN throw voice on it? As Callahan explains, "voice over packet is not the panacea you might think." Callahan says that you can't expect to magically reap huge rewards by simply "throwing" voice on your company's data network. That's not to say the rewards aren't possible—only that reaping the rewards is not that simple.

As you may expect, the relative complexity of deploying a packet voice solution is a matter of opinion. According to Hanssen, VoFR is "very simple" to install and to configure. Gary Cordeiro, marketing manager for Lucent's DEFINITY ATM Networking solution,

Voice Over Frame Relay: Let's Talk Options

The Frame Relay Forum (FRF) has developed two standards that are particularly relevant to delivering voice over frame relay (VoFR): FRF.11 and FRF.12. These standards are described below.

FRF.11

FRF.11 defines the frame format and procedures required to deliver voice over frame relay permanent virtual circuits (PVCs) to connect corporate Private Branch Exchanges (PBXs). With equipment that supports FRF.11 installed at multiple sites, you can concurrently transmit voice and fax traffic with your company's data traffic over a single access link.

FRF.11 notes that voice should be compressed when delivered over frame relay PVCs and, to this end, supports a variety of voice compression standards. (For more information about voice compression standards, see "Where the Lines Cross" on p. 8.) FRF.11 defines two classes of voice compression:

- Class 1 voice is typically transmitted at 32 kbit/s, using a 2-to-1 compression ratio.
- Class 2 voice is transmitted at 8 kbit/s, using an 8-to-1 compression ratio.

In addition, FRF.11 supports silence suppression and a method for carrying multiple voice samples in a single frame. FRF.11 also supports multiplexing, enabling a single frame relay circuit to support as many as 255 calls simultaneously.

FRF.12

FRF.12 defines fragmentation, which is a procedure that VoFR equipment can use to break up large data frames into smaller frames. In this way, VoFR equipment can ensure an even flow of voice frames into the network, thereby minimizing delay. (See Figure 3 on p. 12.) Fragmentation also reduces jitter because voice packets are sent and received more regularly.

EQUIPPING YOUR COMPANY'S NETWORK FOR VOFR

You can deliver voice over any carrier's frame relay services, assuming that you have ensured that the carrier's frame relay network can guarantee a delay and jitter that adequately meets your company's needs. Most of the major carriers offer frame relay services, including MCI WorldCom and AT&T Solutions. Few carriers offer managed voice over frame relay services. However, in June, AT&T Solutions announced a managed voice over frame relay service called Managed Multiservice Networking (MMN).

Unless you use a managed voice over frame relay service, such as that offered by AT&T Solutions, you will have to purchase your own VoFR equipment. Most of the major router vendors, including 3Com Corp., Cisco Systems Inc., Lucent Technologies, and Nortel Networks Corp., offer routers, switches, or voice-enabled frame relay access devices (VFRADs).

For example, Cisco offers the MC 3810, the IGX 8400, and the 2600, 3600, and 7200 router platforms that support FRF.11 trunking and FRF.12 fragmentation. You may be interested to learn that the Cisco 2600 and 7200 router platforms will soon integrate with Novell Directory Services (NDS) by way of Cisco's Complete QoS Policy Manager, which is software you use to manage many of Cisco's routers. ●

Voice and Telephony Over ATM: Let's Talk Options

The Asynchronous Transfer Mode (ATM) Forum has developed several specifications that describe various options related to voice and telephony over ATM (VToA). Two of these options should particularly interest you—the first option because it is the most commonly used and the second option because it is the newest:

- You can deliver voice over ATM Adaptation Layer 1 (AAL 1) through circuit emulation services (CES).
- You can deliver voice over AAL 2.

TECHNICALLY SPEAKING, AAL 1 ISN'T A PACKET

Voice over AAL 1 is the most common approach to VToA in part because this option has been available longer than other voice over ATM options and in part because this option is arguably the easiest to deploy. However, technically speaking, voice over AAL 1 is not a packet voice solution. CES over AAL 1 "emulates the same characteristics of a circuit-oriented connection," explains Don Choi, working group chair for the VToA subcommittee within the ATM Forum. "AAL 1 gives you a circuit connection. It's not a packet connection," Choi adds.

AAL 1 delivers voice using the ATM constant bit rate (CBR) service, which is a reserved bandwidth service that generates a permanently allocated bit stream. Because AAL 1 uses CBR to transmit voice, AAL 1 uses bandwidth inefficiently—using bandwidth even when no voice traffic is being carried. Furthermore, AAL 1 defines no standard mechanism for compressing voice or suppressing silence.

In contrast, AAL 2 is a packet voice solution and makes much better use of bandwidth by using the ATM variable bit rate (VBR) service for delivering voice. Like CBR, VBR is a reserved bandwidth service, but rather than generating a constant bit stream, VBR establishes a peak rate, a sustainable rate, and a maximum burst size for the traffic it carries. The end result is a more efficient use of bandwidth.

In addition, AAL 2 uses bandwidth more efficiently than AAL 1 because AAL 2 supports multiplexing. (Multiplexing is the ability to take several signals and combine them into one transmission.) As a result, one ATM virtual circuit connection (VCC) over AAL 2 can support several calls, as many as 248 in fact. In addition, AAL 2 supports voice compression standards and silence suppression methods.

EQUIPPING YOUR COMPANY'S NETWORK FOR VTOA

You can deliver voice over any carrier's ATM service, such as that offered by Sprint, AT&T, and Equant. Unless you find man-

aged services specifically for carrying voice over ATM, however, you must purchase your own equipment to meet that purpose. Products that support voice over AAL 2 are beginning to emerge, although at this point, few such products are targeted toward the corporate market. However, most of the major router vendors, including 3Com Corp., Cabletron Systems Inc., Cisco Systems Inc., Lucent Technologies, and Nortel Networks Corp. offer products that support voice over AAL 1.

For example, Cisco's MC 3810 concentrator delivers voice over ATM as do the Cisco 2600, 3600, and 7200 router platforms. In addition, Cisco is integrating its 2600 and 7200 router platforms with Novell Directory Services (NDS). (These routers are integrated by way of Cisco's Complete QoS Policy Manager, which is software you use to manage many of Cisco's routers and will be available soon.)

NOW THAT'S A SWITCH

Lucent offers VToA concentrators and switches and a unique approach to voice over AAL 1 called DEFINITY ATM Networking. DEFINITY ATM Networking is unique because the voice-into-ATM-cell conversion occurs on a Private Branch Exchange (PBX).

DEFINITY ATM Networking is made possible through ATM software and an ATM interface board that you install on a DEFINITY ECS G3, which is the largest of Lucent's DEFINITY PBX product line. With this VToA-enabled PBX at your company's headquarters, you can integrate voice traffic with the data traffic you route between sites on an ATM-based campus or WAN. Of course, these sites must also have a DEFINITY PBX (any model) equipped with an ATM interface board.

This fall Lucent will begin to beta test new directory-enabled tools that will make it easier to manage its DEFINITY ECS G3. These tools use Lucent's real-time Directory Synchronization Technology (DST) to synchronize the master directory that DEFINITY uses with NDS and other directory and application databases. The result is that when the database for one application or directory is modified, fields in the databases for the other applications and directories are also modified. For example, when you create a new user in NDS, this action could trigger the automatic creation of a voice extension and a voice mailbox.

Novell, which is promoting directory-enabled networking standards with Lucent, will test Lucent's directory-enabled solutions with NDS at its corporate headquarters this fall. Novell will conduct the trial with the DEFINITY ECS and Lucent's INTUITY AUDIX Multimedia Messaging System. Commenting on the planned trial, Michael Simpson, Novell director of marketing, says that Novell already uses these Lucent products and needs to integrate them with its NDS tree. "We like to use all of our technology internally," Simpson adds. ●

also implies that installing and configuring a VToA solution need not be difficult. (For more information about DEFINITY ATM Networking, see "Voice and Telephony Over ATM: Let's Talk Options.")

However, Callahan expresses a different opinion when he says that setting up VoFR "takes a lot of work."

"Believe me," Callahan adds, deploying a VoFR solution "is not a trivial task." Choi expresses a similar sentiment in response to a question about the potential disadvantages of VToA. Although "there are clear benefits of carrying voice over ATM," says Choi, "the disadvantage is the cost of introducing complexity."

Dial Plans and Hop-Off Strategies

In fact, packet voice introduces several complexities. For example, when using the standard telephone service, your company receives a monthly bill from its carrier. These bills feature itemized lists that make it relatively easy for your company to track its expenditures and to bill separate departments for



their portion of the long-distance charges. If your company uses packet voice services, billing becomes your company's problem.

As a network administrator, packet voice solutions may also create problems for you. For example, you must create a dial plan, or you must revisit the existing plan, working with the department responsible for managing the PBX network. This dial plan ensures that all of the telephone calls that you want to go over your company's data network are routed through the appropriate devices.

You and others must also develop hop-off strategies. A *hop-off strategy* is a plan that dictates when calls should hop off the data network and on to the PBX or public telephone network. A hop-off strategy is necessary if you want to reduce long-distance charges for calls to people who are not on your company's network or to branch offices that are on the PBX network but will not exchange packet voice. Carrying calls for as long as possible on the data network and having them hop off at some point onto the public telephone or PBX network is a practice called *off-net calling*.

Who's Going to Do The Dirty Work?

Creating hop-off strategies for every possible dialing situation can be a complex, time-consuming task. In fact, according to Callahan, the task sounds so unpleasant that it prompts some companies to just say "No" to packet voice. When potential packet voice customers hear that they can make international

calls for cents per minute, the idea is very attractive, Callahan says. But "when they realize how much work is required to get that savings, [packet voice] becomes much less attractive," explains Callahan.

Of course, you don't necessarily have to deal with this complexity yourself. Someone can do it for you—for a price. If you purchase DEFINITY ATM Net-

If the threat of complexity hasn't squelched your initial interest in packet voice, one final question remains: How do you decide which packet voice technology to use?

working, for example, Lucent supplies a project team for that product at each location. This project team will help resolve your dial plan, hop-off strategies, and signaling issues.

Similarly, AT&T Solutions is just beginning to offer a managed multimedia network service called *Managed Multi-service Networking (MMN)*, which enables you to deliver voice over frame relay. In the initial release of MMN, you will have to deal with the dial plan and

hop-off strategy and any other complex configurations yourself. But eventually, AT&T Solutions plans to handle those complexities for you.

DECISIONS AND SWEET SPOTS

If the threat of complexity hasn't squelched your initial interest in packet voice, one final question remains: How do you decide which packet voice technology to use? Most likely, you'll base that decision on the type of data network that your company already has in place. Even if your company doesn't already have a WAN, you will select WAN technologies based on which technology handles your company's data—not voice—the most efficiently and reliably for the best price. You will select a packet voice solution based on whatever data network you choose.

If your company's network is entirely ATM, the decision is an easy one—likewise if your company's network is primarily IP or frame relay. But most networks are heterogeneous. For example, your company's network may consist of IP over frame relay with an ATM backbone. In this case, you'll have to look at the predominant transport technology your company is using between domestic sites that exchange large volumes of calls and the plan your company has in place for international offices. Those are the "sweet spots," as Daugherty calls them, where packet voice is most likely to pay off.

Linda Kennard works for Niche Associates, which is located in Sandy, Utah. ●

What's MOS Got to Do With It?

Packet voice service providers and equipment vendors claim that frame relay and ATM have the potential to deliver toll-quality voice. But how do these service providers and equipment vendors know the voice delivered is toll quality? Isn't voice quality subjective? In fact, no one empirically knows what voice quality you are hearing because, yes, voice quality is subjective. That is why voice quality is typically measured by taking an average of subjective ratings from a large number of listeners scoring a wide variety of speakers and utterances.

These ratings, called Mean Opinion Scale (MOS) ratings, are based on a 5.0 scale, with 5.0 representing the highest possible quality. Voice quality with a 5.0 MOS rating matches the quality of voice you hear when you are in the same room with a person. The values of MOS ratings are as follows:

Value	Rating
4.0 to 5.0	Toll quality
3.0 to 4.0	Communication quality
Less than 3.0	Synthetic quality

MOS ratings are typically shown in connection to the voice compression standard being used. (For more information about voice compression standards, see "Where the Lines Cross" on p. 8.) Both VoFR and VToA support the ITU G-series voice compression standards with MOS ratings between 4.2 and 4.4.

If you are using a packet voice solution that supports one of the ITU G-series standards on a well-designed network, you won't know whether the voice you hear during a telephone call is going over the data network or over the public telephone network. You will only know that the voice sounds good—just like you expect it to sound. ●