



In *CHIPS Winter 2004*, we looked at the development of the telephone, telephone systems, and the traditional circuit-switching method of providing telephone service. In this installment, we will examine digital telephony in more detail, including a look at transmitting voice over Internet Protocol (VoIP). As with most communications technologies over the last century, the telephone has become inextricably linked with digital technologies. Computers control phone switches and all the call accounting and detail reporting that runs the multi-billion dollar worldwide telephone industry. Many telephones today are small computers themselves. Many in the telecommunications industry apparently want to go even further. The hottest topic in telephony today is converging voice and data networks and migrating traditional telephone services to VoIP. The end goal is the eventual delivery of all services — voice, data, video and teleconferencing — over the same network.

Digital Telephony: The Basics

The first successful system that supported digitized voice transmission was the T-carrier system, introduced in the United States in the 1960s by Bell Telephone. The basic connection in the T-carrier system is the "T-1" line, which is still a standard today for networking. In a T-1, voice signals are sampled 8,000 times a second and each sample is digitized into an 8-bit digital "word." The samples are then combined over 24 simultaneous digital channels to make a 192-bit frame. A T-1 transmits 8,000 of these 192-bit frames a second. In addition, each frame is separated from the next by a single bit, which makes each block transmit 193-bits. To calculate the data rate for a T-1, multiply the 192 bit frame by 8,000 and add the additional 8,000 framing bits. This gives us a T-1 speed of 1,544,000 bits per second, usually described as 1.544 megabits per second (Mbps).

The T-carrier system uses pulse code modulation and time-division multiplexing. TDM allows you to combine multiple data signals into a single stream by separating the original signals into many segments, each with a very short duration. The circuit that combines signals at the transmitting end of a communications link is called a multiplexer. The multiplexer accepts input from each individual sender, breaks each signal into segments, and inserts these segments into the stream in a rotating, repeating sequence. The resulting transmission contains segmented data from multiple senders.

For example, if you have four phone calls (A through D) traveling simultaneously over one circuit, a segment from phone call A goes

into time slot 1, a segment from call B goes into time slot 2, etc., until each call is completed. If there is only one call, it uses all the time slots. At the other end of the line, the individual signals are separated out and routed to the proper receivers.

Voice Over IP: The Next Big Thing?

While a T-1 is digital, its use in voice telephony still follows the traditional circuit-switching concept of establishing a dedicated communications circuit between the participants. VoIP, which involves sending voice information in discrete digital packets, is a significantly different way of connecting calls. VoIP is based on the International Telecommunication Union H.323 standard for real-time multimedia communications. In a VoIP system, packets are not sent in a continuous stream over a single, dedicated connection. Instead, each individual packet is sent by the fastest route available at the time each packet is sent, just like data file packets. They are then reassembled back into their proper order on the receiving end. Packet-switched voice works, if you minimize latency (delay) and jitter (variation in delay that causes packets to arrive out of order). In the original frequency division multiplexing phone systems, the 4 kHz dedicated channel guaranteed this. Digital TDM systems also do a good job of delivering packets on time and preserving their order of delivery.

Packet-switching protocols, however, were designed to transmit data files and were not developed with any special delivery guarantees for individual packets. If a data packet arrives late, it is shuffled into the correct place when it arrives. If a packet does not arrive at all, it is just sent again later. This is, for example, what happens when you open a document from a Web site that loads via a Web browser plug-in — the document isn't displayed until the entire file arrives.

Voice quality using VoIP is a whole different ball game. Network delays longer than 150 milliseconds (ms) or packet loss of 10 percent or more will significantly degrade a voice transmission. If a packet is lost, there is not enough time to send it again. If a packet doesn't arrive in time for translation it might as well have been lost. In short, if you cannot guarantee rapid delivery (within 100 ms) of each packet, you probably won't be satisfied with your VoIP service. There are practical ways to achieve the required level of service. But first, let's look at an impractical one.

Zipped VoIP

In the summer 2003 issue, I described how Zippy's mega-networked cabin in the North Woods had crashed. OK, maybe crashed is too mild a term to describe the digital meltdown we experienced. However, like the Phoenix, Zippy arose from the ashes. This time, however, he did the one thing that guaranteed success: He got his wife to help rebuild the cabin's network. While Zippy has the attention span of a goldfish, his wife has the focus of an electron microscope. All the various individual management systems were replaced with a single, end-to-end environmental management application. The new system handles all the global functions (light, temperature control, power sensing, etc.) through one application that tracks and relates all the various things going on in the house. Lights and music, for example, can now follow you as you move from room to room. Overall, they went from over 100 different individual applications down to five.

One other significant change in Zippy's system was that the phone system was now VoIP. As their home network backbone is a SONET

(Synchronous Optical Network) ring running at OC-192 (Optical Carrier Level 192), they do not spend a lot of time worrying about packet loss in their phone calls. All in all, things looked pretty good until the phone started ringing. Well, perhaps "ringing" is the wrong word; "singing" might describe it better, though only very loosely.

Zippy has many idiosyncrasies, one of which is that he owns at least one of every singing fish ever produced. There isn't a room in his cabin that does not have some form of piscatorial oratorio. Zippy had wired every one of them into the phone system. When a call came in, instead of a phone ringing, the fish would start singing. It was actually pretty amusing, at least for the first two or three calls. The fish would sing until you said, "Hello." Then the lip-sync modifications would kick in and the fish would animate in time with the caller's voice. Like the lights, phone conversations would also follow you from room to room (i.e., from carp to catfish). Zippy was quite proud of his little telefishies and his toddler twins loved them, too.

As usual with Zippy's system modifications, though, things eventually went horribly wrong. Much to his children's delight, right in the middle of a call every fish in the house started "speaking" gibberish. The cacophonous babble filled the house until we found the problem — somehow the computer controlling the fish had become completely corrupted with "Teletubbies" sounds sampled from the broadband cable television feed. We suspect the children, though nothing's been proven yet.

VoIP Issues

As illustrated by Zippy's setup, the two biggest issues I see with VoIP at the moment are service and security. On the service side, VoIP really needs gigabit Ethernet running over single-mode fiber. There are network prioritization systems, in particular asynchronous transfer mode (ATM), which can give voice traffic the priority it needs as it flows over the same network as other packets. However, most networks will require significant expansion of their existing infrastructure capacity to make room for anything resembling enterprise-level VoIP on copper wiring where giving voice packets priority can really squeeze your data traffic. I prefer my VoIP over glass for two reasons:

First, I believe our future lies in optical networking. Today's advanced wide-area networks use a combination of Internet Protocol, ATM, SONET and Dense Wavelength Division Multiplexing. DWDM, in particular, has driven a tremendous improvement in long-haul bandwidth that will soon affect metropolitan and local area networking, as well. DWDM allows a single fiber cable to carry up to 80 separate channels of data using different wavelengths of light for each channel. That is a lot of room for whatever you want to transmit. As bandwidth increases to a point where there is enough space for everything, you can spend less time and effort managing traffic.

Wide-area telecommunications operations will probably continue using SONET for the time being, as it is an established technology that performs well for voice transmission. However, as IP and DWDM evolve, the ATM and SONET layers, which impose a significant amount of overhead, will eventually disappear. For example, increasing the capacity of a SONET connection requires an upgrade to every SONET device on the fiber. This is both expensive and disruptive. With DWDM, you can increase point-to-point band-

width just by lighting up another color wavelength on the existing fiber connection. I believe that just-in-time provisioning of optical circuits using DWDM over IP will soon become the gold standard of optical networking. It will take time to evolve beyond ATM and SONET, both in terms of technology and culture, but it will happen.

Second, building optical networks to support voice traffic may help us avoid slow, painful evolutions of existing legacy data network infrastructures. We will not be ready to move voice services onto our networks until our "net tone" is as reliable as our dial tone. My cunning plan is to do it the other way around. Instead of trying to add voice traffic to legacy data networks already struggling to keep up with bandwidth requirements, why not build a discrete voice network and then gradually migrate data traffic to it? You can keep your existing voice and data services intact during the transition while you custom-build the new network from scratch without any legacy artifacts from the old data network. Once the voice services are stable, you will be able to start adding in data and video. It will not be a cheap, instantaneous process, but I believe it may be more rewarding in the long run than trying to do it the other way around.

VoIP Security

The second big issue for VoIP is security. Mention putting voice services on a data network and most network security professionals will be outraged. There are a variety of security concerns with attaching a phone switch to a network. First, there is a long list of security vulnerabilities associated with phone switches, though most of them are only threats on older key-based systems, not the newer digital PBXs (private branch exchanges). A report on PBX vulnerabilities released by the National Institute of Science and Technology in 2001 outlined a variety of maintenance, tapping and feature-related vulnerabilities.

Second, most phone switches come with modems, which might as well be made of wood and shaped like a horse as far as the network security staff is concerned. PBXs have modems to enable remote maintenance, which means the computer security officer is not likely to let you attach the PBX as long as the modem is active regardless of whether or not the switch is capable of being used as an attack platform. The best way to avoid using a modem is to do all the maintenance via IP inside the firewall, which is yet another argument for building a discrete voice network.

Third, many of the most publicized VoIP successes involve using 802.11 wireless networks. Given the perceived vulnerabilities of both wireless networks and phone switches, I cannot blame the network security community for thinking that combining the two is the network equivalent of throwing gasoline on an electrical fire. Not only has the 802.11b Wired Equivalent Privacy protocol been compromised, but so have the proprietary security protocols developed by VoIP vendors. Despite this, however, there are apparently quite a few places using VoIP over wireless network, though the people I have talked to have set up discrete wireless networks just for voice traffic. Hopefully the 802.11i secure wireless standard due out this year will address most of this. Fourth, I also have concerns about putting my voice switches on the data network primarily because anyone on the network can now try to hack the phone switch. Security is a two-way street.

Finally, as with any new thing there will be few wrinkles in how people implement it. A recent report by the United Kingdom's

National Infrastructure Security Co-Ordination Center (<http://www.uniras.gov.uk/vuls/2004/006489/h323.htm>) disclosed some vulnerabilities in products that support ITU H.323. While any vulnerability is theoretically fixable, this is why I recommend waiting until Version 3 (or even 3.5) before you invest in a new technology. Let someone else pay the company to fix their products; I will buy the one that finally works. For those with a serious interest in security, I recommend the Oulu University Secure Programming Group Web site at <http://www.ee.oulu.fi/research/ouspg/>. This is the group that developed the software used in the British group's H.323 testing.

Why Ask Why?

Now that we know what VoIP is, why would we want to use it? Traditional PBXs work just fine and provide pretty much all the same features without all the aforementioned service and security issues. Aside from being a stealth method of acquiring a new DWDM optical enterprise data network, what is the payoff for switching to VoIP? A lot of VoIP advocates propose that Internet telephony avoids the toll charges generated from traditional telephone service. This is not really a good argument. Government long distance rates in 2003 through the General Services Administration Federal Technology Service's contract are less than 3 cents a minute. To use toll avoidance as the sole cost justification for a \$100 million VoIP system, 1 million people would have to spend 3,704 hours apiece calling long distance. That comes out to 463 work days per person for 1 million people. As we are not in the telemarketing business, it does not seem likely that we will recoup much from toll avoidance, if we could even begin to provide an enterprise voice network for a million people for as little as \$100 million.

A better reason might be to reduce administrative and maintenance overhead. System upgrades and reconfigurations can be expensive for traditional phone systems. Moving people from office to office or from one building to another requires rewiring and reprogramming. VoIP phones can each be identified by a unique IP address. Unplug one, move it, and plug it back into the network and it could identify itself and go right back into service with the same phone number, even if you move it from New York to California.

However, by itself this still is not enough to justify the cost of VoIP. In fact, if you already have relatively new (within six years old) digital PBXs on an installation, you will not see much, if any, benefit from switching to VoIP right now. Key-based PBXs will also hold their value for a long time. They have all the same features as VoIP systems and they are very good at the only thing they really have to be good at: They provide dial tone every time someone picks up a phone.

The people who will benefit most from VoIP are those who can make a great leap from leased services or really old, obsolete technology. If you are providing services to 5,000 people using Integrated Services Digital Network (ISDN) lines and Tone Commander telephones, you might be ready for VoIP. If your switch is so old that its label reads "Bell Telephone," you might be ready for VoIP. If your phone services are provided by leased Centrex lines and your account representative sends you a Christmas postcard from his villa in Acapulco every year, you might be ready for VoIP.

Evolution, Not Revolution

Let's say you have decided that VoIP is where you want to go, but

you are not ready to completely replace your current systems. Here is a phased approach that might work for you.

1. Connect your digital PBXs via a network. In my region, we have 132 PBXs, with 21 models from nine vendors and almost as many different versions of the software as there are switches. They range in age from a few weeks to 19 years old and have all been managed locally as stand-alone systems. While this is not uncommon, it is a configuration management nightmare. My first objective is to be able to reach out and touch all our switches remotely, if only to get an audit of the hardware, software, number of phone lines, etc., associated with each switch. In this phase, the PBX is simply another dumb device attached to the network that has no rights or privileges, and it doesn't pass any traffic other than text-based data about itself back to whatever central management software is keeping track. Replace any switches that do not allow this functionality and gradually reduce your switch inventory down to no more than three models each from two vendors. Standardize the software version on each model of switch or, if possible, for each vendor's products. This gives you a foothold on the network and a chance to prove to the security officer that PBXs can play securely.

2. Remote maintenance and administration. Once you can pull information from each switch, the next step is to tinker with them from a distance. This is where you can do minor software patches or upgrades, and maybe even the adds, moves and changes that are the largest part of the day-to-day management of any PBX. This still is not VoIP, though you are getting closer at this point.

3. Local area VoIP. Your first steps into true VoIP should be on a small, local scale. Test and implement locally using a discrete voice network. However, unlike what we first did with data networking, try to make sure every VoIP implementation adheres to the same standards, equipment and protocols. That will make it much easier to assemble all the different VoIP networks into a cohesive system if you decide to expand it to the enterprise.

4. All VoIP – all the time. Enterprise-class VoIP, spread across the continent and around the world, including ships at sea — by the time this happens, I would also expect that we will no longer consider radios and telephones different communications mediums because it will not matter where you are standing or what device you are holding — they will all interoperate. Handoffs between telephone systems and command and control networks will be seamless and secure and your telephone number will follow you wherever you go.

I believe that we should design future systems on our expectations, not on current limitations. There have been too many systems deployed in the last 10 years that were obsolete when we flipped the "ON" switch. We have traditionally treated voice, video, text and data as different communications systems. I prefer to think of them as modular components of how we communicate. VoIP will eventually be a big step in that direction. Next issue we will look at mobile telephony.

Until then, Happy Networking!

Long is a retired Air Force communications officer who has written regularly for CHIPS since 1993. He holds a Master of Science degree in Information Resource Management from the Air Force Institute of Technology. He is currently serving as a Telecommunications Manager in the U.S. Department of Homeland Security.

