Get Your Data Network Ready for Voice

by Jay R. Brandstadter
Delphi, Inc.

Despite the economic blues and distress in the telecom industry, interest in IP–based enterprise phone systems (IP-PBXs) is expanding rapidly. Although there is reason to question whether the complex array of voice over IP (VoIP) technologies and products is fully ready and whether these technologies and products offer a compelling value proposition for the user community, all PBX vendors are now delivering VoIP or IP telephony systems. Many users in higher education are doing trials, network analysis, or operational deployments—and many more are trying to decide whether and how VoIP can benefit their institution. *(See sidebar on page 38 for one such example.)*

However, the fact that IP networks can carry voice doesn't mean that all IP networks will do it with acceptable reliability and quality. Many VoIP installations fail—or run up large unbudgeted expenses—because the LAN, WAN, or both are unable to handle voice successfully due to needed infrastructure and other changes. Customers often overlook, and vendors often downplay, the need to ensure that the data network is up to the challenge of voice.

If your institution is considering VoIP, you must know whether your data networks are voice-ready. If the proposed installation will compromise call quality, it's essential to know that in advance—so you can decide either that reduced quality is acceptable for your application or that you need and can afford an upgrade.

Start with Inventory and a Plan

Suppose your institution has approved some sort of initial step toward convergence of the voice and data networks. You're probably going to start migrating to VoIP through a pilot project in a portion or subnetwork of your data network infrastructure and then expand the rollout over time, based on results, budget, and other factors. Interoperability between your IP-based system and legacy voice networks could be a big problem, especially if you are mixing vendors.

Selecting the initial data network for migration to VoIP is not a trivial matter. If you're fortunate to have latitude in that decision, you should opt for a controlled or bounded situation, preferably a new or *greenfield* site or infrastructure. If you're not that lucky, then you'll have to upgrade an in-place data network for voice.

You will need to assemble traffic data for the voice calls running over this network, decide on a method of digitizing and compressing the voice signals, and convert this information into estimated bandwidth requirements.

Taking Stock

Start by taking stock of your data network assets and develop a road map or plan of where you want to go. You will have to learn everything you can about the network: configuration, equipment, data traffic, applications, users, interfaces, bandwidth/capacity, performance characteristics, error rates, and so on.

You will need to assess growth for the network—not only through the addition of voice but also for video, fax, Web interaction, and other media and applications. LANs and WANs pose different challenges for VoIP and must be analyzed separately.

As part of this initial assessment, the relationship between potential voice applications and the data network needs to be considered. (*Voice* should be defined to include related applications, such as voice mail/ messaging and voice-activated processes.) This relationship defines some boundary conditions on numerous aspects of a converged solution.

For example, how do the server requirements of the voice portion affect the IP network's configuration and bandwidth? What cooperative, interoperable functions and mechanisms are feasible in a multivendor architecture, where, for example, PBXs or voice servers are from different vendors than routers and data switches? How can the data network be strengthened to provide the comfort and robustness of familiar legacy voice and voice-related capabilities?

The Perils of Packets

Packetizing voice for transmission and handling by IP networks makes significant demands on the network. Packetizing generally begins with digitization and compression of the voice signal by a CODEC using a standard like G.711 or G.729.

Users are accustomed to phone conversations with *toll-quality* sound and performance. Getting that quality over a data network is far from guaranteed. In a traditional voice network, every call is assured a fixed amount of bandwidth; in an IP network, voice packets must contend for network resources with other traffic using the system—and the amount of contention varies millisecond by millisecond.

As a result, packets can be delayed (latency), the amount of packet delay may vary (jitter), packets are lost or dropped, and some arrive out of sequence. Beyond a certain level, all these problems degrade voice quality and can even cause dropped calls. These problems can essentially be eliminated, but ultimately it's a matter of aural perception.

Formula for Trouble

These problems don't just affect human-to-human communication. Other calls that can use VoIP connections—most notably dial-up modem and fax calls—are particularly sensitive to dropped packets. In such cases, losing a packet or two will generally drop the call.

Even with Gigabit Ethernet LAN backbones and switched 100 Mbps to the desktop, you cannot ignore the potential for trouble. Consider an e-mail with a long attachment (e.g., a large PowerPoint file) that is broadcast to most of the users on a corporate LAN. Now try to get even a normal volume of LAN telephony calls through at that moment. Any questions?

The amount of damage caused by these packet perils is, of course, closely tied to the available bandwidth and the nature and volume of the traffic handled by the IP network. As traffic increases, packets may enter the network faster than they can be forwarded, causing congestion. Data and video such as streaming IP video in distance-learning applications will, of course, add to the traffic load. When the network or network segment approaches a congested state, packet delay increases, packets may be dropped, and the network appears to slow down.

There are numerous tradeoffs between key measures that affect voice quality (delay, jitter, and packet loss), bandwidth, network architecture, and network policies. These are complex relationships marked by a delicate balance among many factors and parameters.

Navigating these waters to provide voice on a data network is no simple matter, but you can begin by examining how to tune the data network and whether to employ quality of service (QoS) policies.

Tuning

IP-PBX providers and integrators will assess your data network's voice readiness—often for a price—or you can do it on your own using network assessment tools of providers such as NetIQ. Once you determine that some degree of upgrade is required, you can take a number of measures to improve call quality.

You can add more bandwidth, upgrade or replace existing network equipment, improve your network architecture, or reconfigure the network. Bandwidth can be made available by acquiring higher-speed links—more expensive in the WAN than the LAN—and by more efficient use of the bandwidth available. Options in the latter category include increasing voice compression, varying packet size and framing, using silence suppression and voice-activity detection, and compressing the RTP headers of packets. These methods try to minimize packet overhead and maximize the efficiency of packet payload.

There are lots of trade-offs here. For example, using lower data-rate CODECs for voice compression will save bandwidth but may reduce voice quality below the level required for your application. An alternative approach may be to use call-admissions control software that limits the number of concurrent VoIP conversations to an acceptable predefined number, overflowing additional calls to a non-IP network.

It is also possible to improve performance and capacity without buying more or bigger pipes by upgrading or replacing network equipment. Among the available techniques are the following:

- Assure that you have up-to-date high-speed switches and no hubs in LANs.
- Replace software-based processes or functions (e.g., firewalls) with hardware-based versions to reduce delay times and increase capacities.
- Increase memory (RAM) for router queues and other traffic-sensitive points in the system.
- Assess whether your network backbone switches/routers are in need of modernization or upgrade to meet the new voice requirements.

Also, a review of the network layout and architecture can be helpful. Can shorter, more direct routes be established for VoIP calls, thus reducing propagation and transport delays? Can the number of router hops be reduced? Will a better understanding of traffic patterns or "communities of interest" lead to improved network performance?

Quality of Service

Another way to use the data network more efficiently is to create different categories or classes of traffic and give priority to some categories over others. This approach, generally called quality of service (QoS), does not eliminate congestion, but it allows traffic to be prioritized so that the traffic most sensitive to delay goes to the front of the line.

There are differing opinions on whether a QoS mechanism is essential to a successful VoIP installation. It's often argued (especially by data-centric vendors) that assigning priority to every voice packet is the only way to be sure that voice will get through.

Others (including some traditional PBX makers) point out that when the data network is less than cutting-edge and operates in mixed-vendor environments, QoS may be prohibitively expensive. Here, incremental improvements to bandwidth and equipment, particularly in the LAN, may provide an effective and less costly solution.

There are certainly examples of VoIP installations that have achieved acceptable call quality without QoS, simply by throwing more bandwidth at the congestion problem. Increasingly, vendors are including QoS technology in their LAN switches at no extra cost, giving users the option to turn it on. But upgrading to such equipment isn't free—nor is the management software needed to control and monitor it.

Most seriously, in QoS, vendors generally *talk the talk* of open standards but *walk the walk* of proprietary technology— which means that QoS just doesn't work very well in multi-vendor environments. If all your networking gear carries one vendor's logo, QoS can be an elegant, though still complex and time-consuming, solution. Otherwise, it's yet another immature technology associated with VoIP—implying frequent changes, and possibly inadequate knowledge, experience, and tools.

Keeping Track of Costs

If the new IP-based phone systems are to be a viable alternative to circuit-switched voice, they must reproduce the comfort, robustness, and overall familiarity of legacy voice systems and the public switched network. The system must offer the network owner and network engineer reliability, security, manageability, and availability of qualified support people and tools. Right now these items are primarily works-in-progress.

Tuning and QoS mechanisms will help prepare data infrastructures to support voice, but they aren't free. The box on this page presents a checklist of some major components in readying a data network for voice. This cost summary profile is not intended to deter you from considering VoIP for your organization. It does highlight, however, that there are many potential data-network costs that need to be factored into your VoIP value proposition.

Training, education, and conferences on convergence and VoIP are available from numerous sources, and the topic is beginning to appear in industry publications. PBX makers, independent vendors, and telecom consultants can all offer help, but at this point everyone is still in learning mode.

Conclusion

Is it still too early to enter the shark-infested waters of VoIP migration? Not necessarily. But be patient and do your homework as best you can. The readiness of data networks to support voice is a moot point if your organization is not prepared and the payback is not clear. Think of the complexities and the machinations needed for the data infrastructure to handle voice, which the circuit-switched world has been doing quite well for decades, and calculate carefully whether beginning this transition now is beneficial.

Jay R. Brandstadter is a partner of Delphi, Inc., a Washington, D.C.-based consulting firm. He can be contacted at 301/871-1021 or *jbrand@rcn.com*.

Reprinted with permission from The ACUTA Journal of Communications Technology in Higher Education, Volume 7, No. 1, Spring, 2003. To learn more about ACUTA, visit us at www.acuta.org.