Volume
3

IMPLEMENTING THE VOICE OVER IP NETWORK

Mark A. Miller, P.E.
President
DigiNet Corporation

A technical briefing from:

Sniffer
TECHNOLOGIES

April 2002

Table of Contents

Executive Summary

This is the third of six technical briefing papers that examine the concepts, operation and analysis of *converged networks*. A converged network — one that combines voice, data, and other signal transmissions into a single, higher-speed network interface — is more complex than a single-purpose voice or data network. Because of this challenge, the design, implementation and testing of this infrastructure requires rigorous attention to details in order to present the end users with an acceptable quality of service (QoS).

The design and implementation phase can be divided into ten steps, beginning with a clear definition of the objectives for the converged network and defining the applications that this network must support. A systems requirement document provides guidance for the procurement process. Installation and testing plans include plans for testing between multivendor components. The installation phase should begin with a small pilot network in a lab or on one network segment to thoroughly test for reliable transport of all media — voice, video, fax, and data — over the converged infrastructure. Once the operation of the pilot program has been verified, the larger system can then be tested, deployed and documented.

Most end users have had experiences with voice calls over the PSTN that have been quite positive, with fast and reliable call setups and clear connections. These positive experiences have raised the bar of end user expectations, which means that the converged network must equal — if not exceed — these experiences in order to be considered successful. Many factors can influence QoS, including bandwidth of the communication channels, network loading factors, sources of echo within the network, and latency or delay. Several methods of measuring QoS, taken from both subjective and objective analysis approaches have been defined and are discussed in this paper.

Finally, a few guidelines are presented, contributed by network managers who have already installed converged networks, to make your journey down the converged network road a more positive experience.

1. Challenges, Choices and Concerns for Voice over IP Network Implementations

In the first two installments of this series of technical briefing papers, we considered the concepts and challenges of a *converged network*, and the protocols that are required to support such a network architecture. Recall that a converged network is one that has the capabilities to transmit voice, data, fax, video, or some combination of these different signals. Furthermore, that architecture incorporates elements of both the Public Switched Telephone Network (PSTN) and Internet Protocol (IP)-based networks such as virtual private networks (VPNs) and the Internet. For a successful implementation, the inherent nature of these dissimilar systems must be considered. Let's consider these issues from three perspectives: the challenge, the choices and the concerns.

The *challenges* are derived from the fundamental differences between voice and data networks — voice networks are connection-oriented, while data networks are connectionless; and therefore these two networks approach the flow of information quite differently. For example, voice networks were designed with an objective of 99.999% reliability, the "five nines" that is so frequently quoted. One of the fundamental reasons behind this requirement is life safety, such as contact with the police and other emergency agencies. If your house is on fire, you need to be able to reach the fire department quickly and reliably. The 99.999% reliability factor represents only two hours of downtime in 40 years of operation, which is an objective that few, if any, packet-switched data networks would be able to achieve.

In addition to differences in reliability, traditional voice and data networks were designed to provide different information transport services. Human communication is sensitive to end-to-end transport delay, plus variations in that delay, which is known as *jitter*. Within the voice network, fixed communication paths are established with bandwidth and other resources dedicated for the duration of the call, which minimizes the delay and jitter. In addition, any loss of information can be recovered by the end users by simply asking their counterpart to repeat the last phrase. In contrast, traditional packet-switched networks do not establish a fixed end-to-end path but provide routing on a per-packet basis. Delay and jitter are of less concern, since it is likely that the large messages will have been divided into multiple packets at the transmitter (known as *segmentation*) and subsequently put back together at the receiver (known as *reassembly*). These two processes are typically part of the packet-switched architecture. However, if one of the packets in this sequence is lost, then the original message cannot be reassembled and passed to the application. In other words, packet networks can reassemble the original message given some packet delay and jitter, but cannot reassemble if a packet in the original sequence is entirely lost.

Second, a *choice* must be made regarding the vendor architecture that will be deployed. There are two fundamental types, plus some hybrid variations on these themes. The first architecture is based on a voice switching system known as a private branch exchange, or PBX. This architecture derived its origins from voice architectures such as PBX networks connecting end users within a single company location, or interlocation trunks connecting branch offices back to headquarters. As the interest in integrated voice and data communication within a single enterprise grew, additional PBX interfaces, such as those supporting fax and data connections were added.

The next logical progression was to add support for interfaces to networks other than the PSTN, such as ISDN, Ethernet and IP. Finally, the internal architecture was migrated from a circuit-switched basis to a packet-switched basis, which allows other packet processing functions such as gateways and gatekeepers to coexist alongside station and trunk interfaces. And since many of these interfaces were compatible with IP, the term IPBX, short for IP-based PBX was born (Figure 1A).

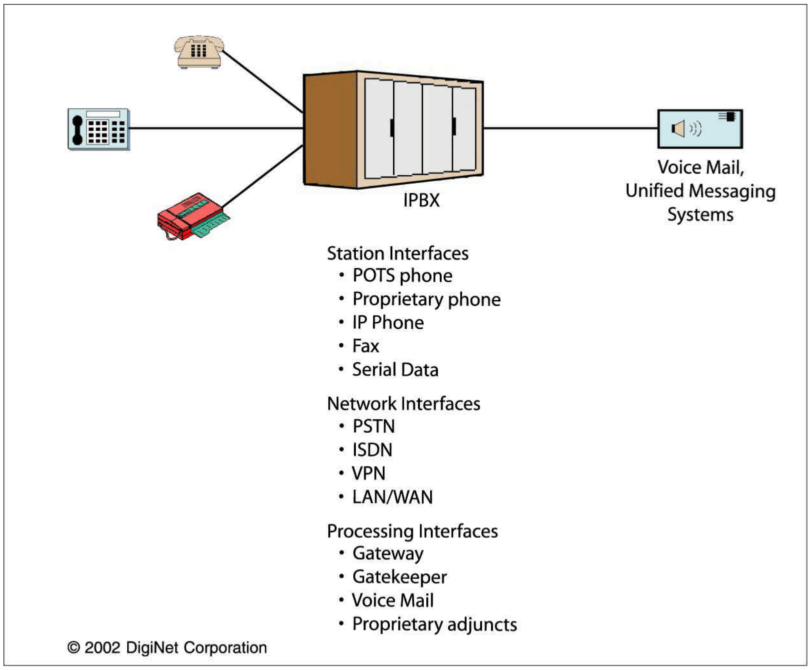

© 2002 DiGiNet Corporation

FIGURE 1A. **IPBX-based Converged Network**

Another choice is a router-based architecture that derives its origins from an IP-based data network. Routers are placed at higher density sites where traffic is aggregated for transport to other locations. Leased lines connect these core routers with each other, plus links to other networks such as the PSTN or Internet. Single or remote users are connected through dial-up lines and remote access servers into the core network. If the voice or video information is in a digital format, it can be packetized and then transported over this router-based network along with emails, file transfers, and Web pages. Multimedia processors, such as gatekeepers and gateways, can be connected into this network like any other IP-enabled device (Figure 1B).

Note that the IPBX-based architecture is more centralized, while the router-based architecture is more distributed. As would be expected, each approach has advantages and disadvantages. For the IPBX network, a reconfiguration may involve fewer subsystems and intermediate steps. For the router network, sites can be migrated individually, with less disruption likely on the network as a whole. With possibly more sites to migrate, the implementation time may be lengthened.

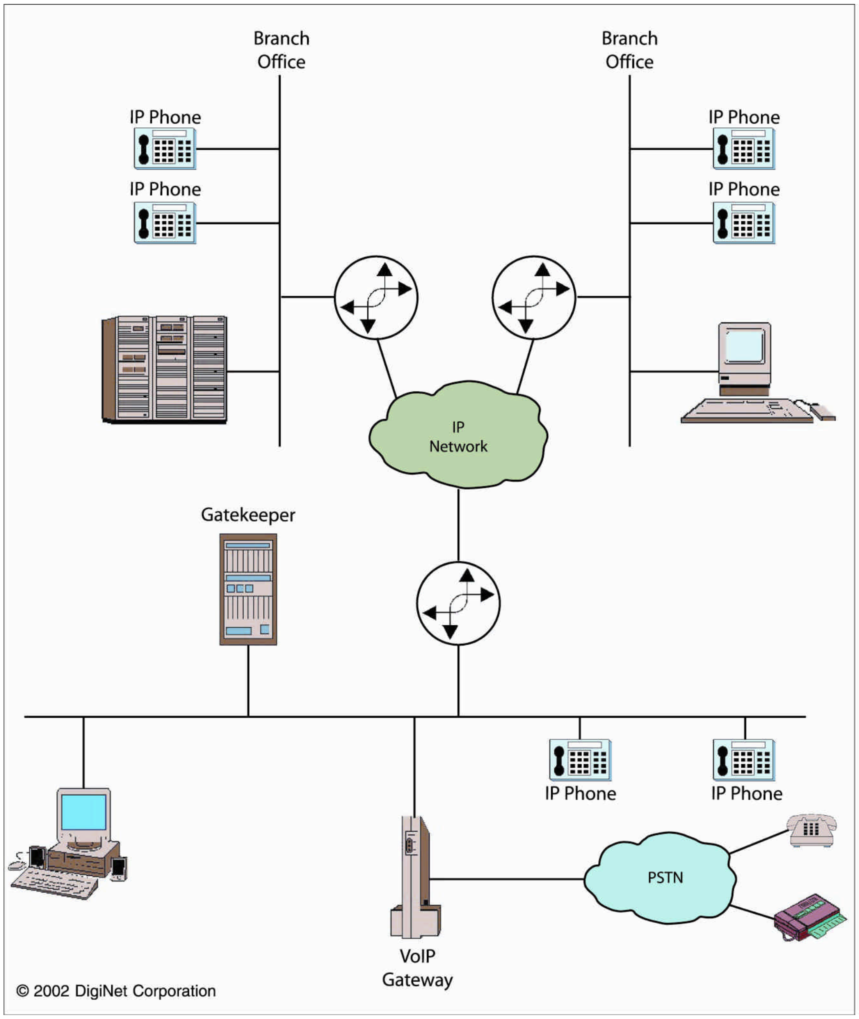

FIGURE 1B. **Router-based Converged Network**

Both of these approaches —IPBX and router-based — share a common *concern*: will the new system interoperate with the existing systems, such as a voice mail processor, automatic call distributor (ACD) or call detail recorder? We will address this concern by first considering steps necessary to design and implement a VoIP network.

2. Design and Implementation of the Voice over IP Network

The converged network is intended to support both voice and data communications, and the design process should consider both of these requirements. Each network presents unique challenges, however ten factors should be considered in all cases. Note that not all steps progress linearly, and that some feedback between steps is required (Figure 2):

1. Define the objective of the converged network, including the applications to be supported, such as video conferencing or a voice-enabled Web site. Reduce this objective to a short mission statement so that all parties involved understand the challenges at hand. In addition, document other ancillary factors that are driving this project. For example, perhaps the PBX system is scheduled for an upgrade, and now is an appropriate time to consider an overhaul of the entire networking infrastructure. Other circumstances might be support for a new enterprise application that is not possible within the confines of the existing infrastructure, or carrier contracts that are about to renew and thus afford the opportunity for voice/data network consolidation.

2. Understand the current voice and data operating environments. Include existing loading factors on LAN and WAN segments, plus the anticipated bandwidth requirements for any new data applications. Consider current call patterns and anticipated growth during your busy hour periods.

3. Prepare a System Requirements Document that delineates functions that the network must provide, and include the appropriate design objectives, such as end-to-end delay, reliability, redundancy, and so on. Determine if one architectural infrastructure, such as the ITU-T's H.323 or IETF's SIP, is preferred for this network application. Also consider any changes that will be necessary to the telephone dialing plans, WAN links, IP subnets or other issues that could impact end users.

4. Consider the adjunct systems that will support this network, such as network management, network analysis, security, and end-user help desks. Verify that any of these support systems, such the network management console or protocol analyzer, are also prepared to support the new environment. Modify the Systems Requirements Document as necessary based on the results of this study.

5. Consult with both current and prospective vendors to obtain their input on the systems requirements. Solicit input from other customers of your vendors, your colleagues, or other network managers in similar industries that may have already experienced the same challenges. Modify the Systems Requirements Document as necessary based on the results of this input.

6. Develop installation and systems acceptance plans, including interoperability testing between various components. Also verify that all signaling protocols between dissimilar networks are compatible.

7. Solicit presentations and bids from qualified vendors, evaluate these proposals, and award the appropriate contracts. In some cases, it may be necessary to hold two rounds of bids: one that is open to all interested vendors, and a second round (called the *short list*) that is limited to the most qualified. Include in the bid documents the additional support systems that have been identified, such as network management and troubleshooting tools, so that acceptance testing of the system can proceed once installation is complete.

8. Develop an Installation and Cutover schedule with the winning vendor(s), plus other organizations (such as carriers) that are also part of the project. Verify the critical paths in the schedule, such as equipment order and circuit installation, and reach agreements from all involved parties regarding the importance of this schedule.

9. Install the new system in a test lab or pilot location, and test its operation and applications before expanding to the larger network. Verify that all dialing plan and network address changes have been successfully migrated to the new environment. Modify the Installation and Cut-over Schedule as necessary based on the results of this pilot installation.

10. As the production network installation proceeds, verify operation of all network components, including interoperability between multivendor systems and adjunct systems, such as voice mail processors. Conduct end user and system manager training classes, and document procedures that these individuals will need to know in order to be effective in the new environment. Complete the as-built drawings of the network and any other appropriate documentation as required.

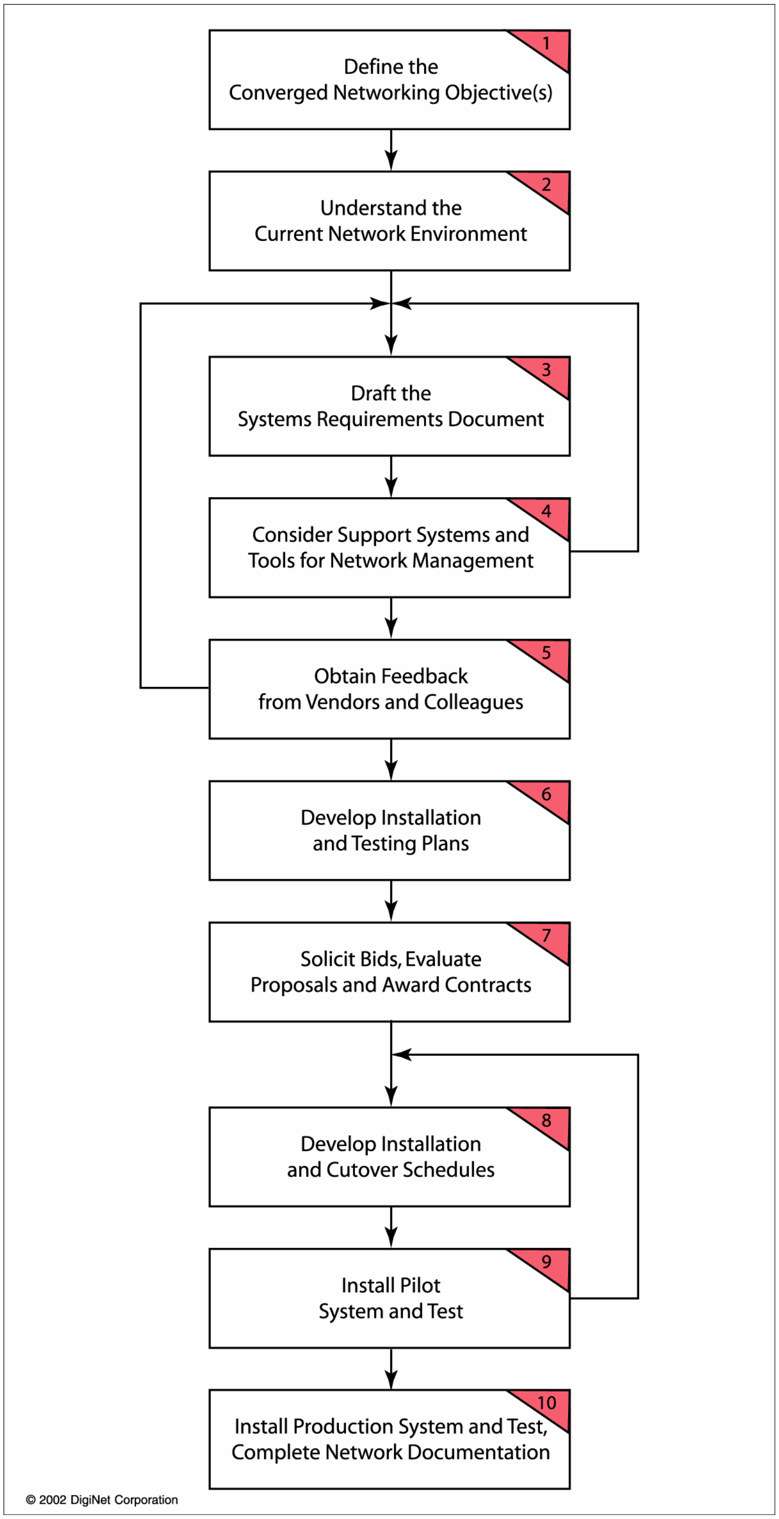
1
Define the Converged Networking Objective(s)
2
Understand the Current Network Environment
3
Draft the Systems Requirements Document
4
Consider Support Systems and Tools for Network Management
5
Obtain Feedback from Vendors and Colleagues
6
Develop Installation and Testing Plans
7
Solicit Bids, Evaluate Proposals and Award Contracts
8
Develop Installation and Cutover Schedules
9
Install Pilot System and Test
10
Install Production System and Test, Complete Network Documentation

© 2002 DigiNet Corporation

FIGURE 2. **VoIP Network Design and Implementation Plan**

3. Measuring the Quality of Service of VoIP Networks

Our previous discussions have considered the reliability differences between voice and data networks, where voice networks are designed for very high reliability — in part because of life-safety issues — and data networks operate from the perspective of "best efforts" service. But if your plans are to *upgrade* the existing voice network with a VoIP infrastructure, you cannot expect the end users to accept a *downgrade* in system performance — known as quality of service, or QoS — resulting from your efforts. Let's begin by looking at the factors that affect QoS, and then consider the industry standards that have been developed to measure those factors.

ITU-T Recommendation E.800 defines QoS as *"The collective effect of service performance, which determines the degree of satisfaction of a user of the service."* A number of network design and operations factors affect QoS. These factors include: packet loss and delays, the available bandwidth, the WAN protocols in use and their efficiencies, the presence of echo that is caused by impedance mismatches within the PSTN, the use of silence suppression that optimizes utilization of the communication facilities, and others.

However, perhaps the most important factor affecting the QoS is the amount of latency, or the delay of the end-to-end transmission. ITU-T standard G.114 recommends a maximum delay of 150 milliseconds (0.150 seconds) in one direction of transmission. If this delay is significantly exceeded, then conversation becomes more difficult. As part of the network design, a *delay budget* may be developed to identify components of this maximum value (e.g. 150 milliseconds), as shown in Figure 3:

- Signal encoding/decoding algorithm, typically 15-37.5 milliseconds at both the origin and destination.

- Protocol processing overheads in components to include RTP, UDP, and IP information plus echo cancellation, typically less than 5 milliseconds.

- Bandwidth and utilization (loading) of the LAN and WAN channels, which may introduce framing and queuing delays in the range of 5-25 milliseconds, depending on the transmission rate.

- Routing, queuing, and propagation delays across the WAN which depend on transmission media and distance, typically 10-40 milliseconds.

- Jitter, or variation in the arrival rates of the packets comprising the audio or video stream. Since the packetized voice samples may take different paths through the packet-switched network, arrival rates of those packets may vary. The receiver buffers the early-arriving packets and waits for the latecomers to catch up, which typically takes 20-40 milliseconds (or more).

Many of the above factors, such as WAN propagation delays, are fixed and cannot be improved. Others, such as the choice of signal encoder, are options that can affect network performance. Many encoding algorithms have been designed, some of which are published as ITU-T standards, and others that are proprietary to a particular vendor. These techniques differ in their underlying mathematical algorithms plus their end results. In general, algorithms that compress the voice (and thus conserve bandwidth) have higher processing delays and lower Mean Opinion Scores (MOS) than algorithms that use no compression. For example, the ITU-T G.711 standard uses a Pulse Code Modulation (PCM) technique and transmits at 64 Kbps with a negligible coding delay. The G.723.1 standard uses two techniques: Algebraic Code Excited Linear Prediction (ACELP) at 6.3 Kbps, or Multipulse Maximum Likelihood Quantization (MP-MLQ) when operating at 5.3 Kbps, with a coding delay of 37.5 milliseconds. A third standard, G.729, uses Conjugate Structure Algebraic Code Excited Linear Prediction (CS-ACELP) coding and operates at 8 Kbps with a coding delay of 15 milliseconds, providing a middle point between these two extremes.

Thus, the network manager faces a tradeoff of bandwidth for delay, and must balance the higher cost of bandwidth with the reality of lower QoS from increasing delays. This balance is sometimes achieved by using a higher bandwidth coder (such as G.711) for local communication, where bandwidth is more plentiful and less expensive. For calls that traverse the WAN, where bandwidth is scarce and more expensive, a bandwidth-conserving coder (such as G.729) is used.

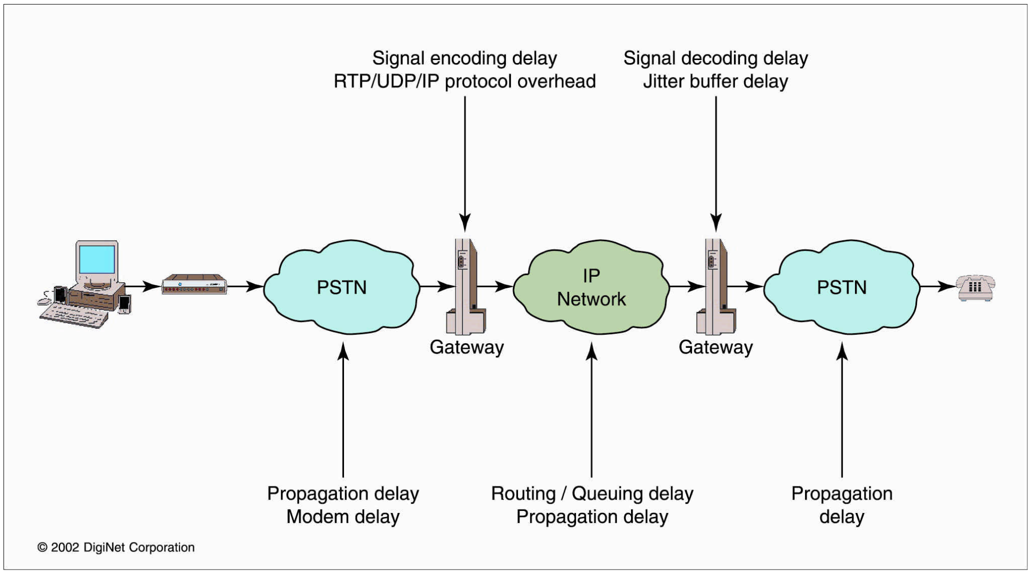

FIGURE 3. **Sources of Delays within the VoIP Network**

Fortunately, many standards are available that provide benchmark methods of evaluating voice network performance.

For many years, the telephone industry has employed a very subjective rating system known as the Mean Opinion Score, or MOS, to measure the quality of telephone connections. These measurement techniques are defined in ITU-T P.800, and are based on the opinions of many testing volunteers who listen to a sample of voice traffic and rate the quality of that transmission. In doing so, they consider a number of factors that could degrade the quality of transmission, including loss, circuit noise, sidetone, talker echo, distortion, propagation time (or delay), and other transmission problems. The most well known test, described in Annex A of P.800, is called the Conversation Opinion Test. The volunteer subjects are asked to provide their opinion of the connection they have just been using, based on a five point scale:

Quality	**Rating**
Excellent	5
Good	4
Fair	3
Poor	2
Bad	1

Since the test subjects are human, and the testing therefore subjective, some variation in the scores is expected. For that reason a large number of people are used in the test, and their individual scores are averaged (hence the term "Mean" in Mean Opinion Score). A MOS of 4 is considered "toll quality" within the telephone industry, and is generally the accepted standard that is targeted for quality VoIP implementations.

Two other ITU-T standards define telephone transmission quality tests that are more objective. The first method is called Perceptual Speech Quality Measurement (PSQM) and is defined in P.861. This measurement is used for assessing the quality of telephone-band (300–3,400 Hertz) voice encoders. The second method is called Perceptual Evaluation of Speech Quality (PESQ) and is defined in P.862. PESQ also addresses the effects of filters, variable delay and coding distortions, and is thus applicable for both speech codec evaluation and end-to-end measurements. In summary, the P.800 MOS test is a subjective evaluation, while those tests defined in P.861 and P.862 are objective measurements that are implementable in hardware for greater accuracy.

4. Interoperability Testing

Testing for interoperability between multivendor systems has been an important issue for enterprise network managers for many years. With multivendor converged networks, these test procedures take on greater complexity and importance. Fortunately, a trade organization, known as the International Multimedia Teleconferencing Association (IMTC), has developed a number of interoperability and testing profiles to make this work easier. The IMTC was formed in 1994 and is a nonprofit consortium of multimedia developers, manufacturers, and carriers. It focuses on open international multimedia standards and the need for interoperability between various systems and services. The IMTC is divided into a number of Working Groups that individually address key multimedia protocols or issues, including H.323, H.248/MEGACO, SIP, QoS, Security, and others.

One of the first efforts of the IMTC was to publish a Voice over IP Implementation Agreement (IA). This IA was developed to complement and clarify the ITU-T H.323 standard, thus making the product developer's job less ambiguous. For many of the architectural layers, several protocols are available and may be equally acceptable. For example, a number of voice encoding algorithms have been developed, including ITU-T G.711, G.723.1, G.728, G.729, and others. If one vendor chooses G.723.1, and another chooses G.728, the two devices will exchange bits, but not communicate. The IA eliminated this problem by recommending that all products support two ITU-T standardized codecs, G.711 and G.723.1, instead of leaving this selection to the discretion of the product developers. Suggestions for protocols at other architectural layers were made as well. Therefore, products that were developed using the IA guidelines are much more likely to interoperate.

Another early effort of the IMTC was to develop an interoperability profile for gateway and gatekeeper operation known as the Interoperability NOW! (iNOW!) profile. This effort was focused at gateway-to-gateway, gatekeeper-to-gatekeeper and gatekeeper-to-exchange carrier interoperability, taking into account the various vendors and carriers that might be participating in a VoIP networking infrastructure.

Recent interoperability testing initiatives from the IMTC have focused on interoperability between various H.323-based systems. The IMTC Conferencing over IP (CoIP) activity group has developed an H.323 Interoperability Test Plan that specifies various call scenarios to be tested on an end-to-end basis. For example, call scenarios are defined between stations with and without the involvement of a gatekeeper; between two network zones requiring communication between the respective gatekeepers; and from an IP network through a gateway and across a WAN. Using these testing scenarios as a starting point, the enterprise network manager can develop similar procedures to assure interoperability on that network's unique infrastructure. Further information regarding the IMTC's work can be found on their Web site at www.imtc.org.

5. Advice from the Veterans

In conclusion, consider the following suggestions from network managers who have taken the converged networking path and have implemented VoIP services:

- Make sure that both the voice and the data networking staffs are buying into the converged networking concepts. The voice folks must understand IP, and the data folks must understand telephony signaling for all systems to coexist.

- Understand your network traffic, call volumes, and calling patterns. If you do not know what the network utilization on your LAN or WAN is at the present time, what will happen when you add voice and/or video to the mix?

- The type of business that the VoIP network will support will play a factor in the overall design of the network. For example, if the network is supporting a call center, then you will have to focus on voice quality. However, if your network connects a number of remote offices or lightly used offices, then a focus on bandwidth and expense reduction may be more appropriate.

- When you are surveying the architectural choices, such as a PBX-based or router-based network, consider the experience you have had with your existing vendors. If you are uncertain which of these two options is better, look at your history of dealing with the respective vendors. Then extrapolate this experience into the more complex environment of the converged network.

- The choice of the ITU-T H.323 protocol or the IETF SIP protocol for signaling may come down to the recommendation of your preferred vendor. Take time to do your own study of this area, however, as both options have their advantages and disadvantages.

- Balance the issues of cost and QoS carefully. Many managers use cost to determine how much bandwidth to purchase, starting with the least possible amount and working up to an acceptable point of QoS. End users may not be very receptive to these experiments. Instead, start with a QoS objective on the test network, and see how much bandwidth is required to meet that objective.

- The codec selection may impact both network engineering and customer relations issues. A quality vs. bandwidth choice will be necessary, and compromises and/or tradeoffs may be necessary. For example, a low-rate codec that gives voice quality equivalent to cell phone calls may be adequate for internal calls between sites. In contrast, a higher rate and quality codec might be a better solution for calls to and from the customer response center.

- Realize that the location of application servers may impact the service provided by those servers. For example, voice mail servers sitting thousands of miles away from the end user may produce poorer voice quality and syllable breaks because of the higher round trip delays and effects of packet jitter.

- A VoIP network implementation may provide additional frustrations in dealing with the carriers involved, as the traditional carriers (both local exchange and interexchange) may view your migration to a VoIP network as a threat to their livelihood.

- Thoroughly evaluate operational parameters such as buffer sizes and silence suppression on the test network before moving to the production network.

- Schedule the implementation of the production network systems carefully, especially if a customer call center is involved. Remember that these applications are the economic fuel for your enterprise engine, and that you can't disrupt their business operations without consequences.

6. Looking Ahead

This is the third of six technical briefs on Converged Networks sponsored by Sniffer Technologies. Titles of current and future volumes in the series include:

1. Introduction to Converged Networking: A description of concepts and challenges of converged networks, including business, technical, and operational issues.

2. Protocols for the Converged Network: A look at the components of the converged network, the ITU-T and IETF multimedia protocol suites, and the protocols required by each component.

4. Managing Call Flows Using H.323: The operation of the H.323 family of multimedia protocols, illustrated with case studies and output from the *Sniffer* protocol analyzer that show converged network operation from the H.323 protocol perspective.

5. Managing Call Flows Using SIP: The operation of the Session Initiation Protocol (SIP) and the IETF multimedia protocol suite, again illustrated with case studies and output from the *Sniffer* protocol analyzer.

6. Supporting the Converged Network: This concluding paper will deal with ongoing support requirements, including: traffic prioritization, WAN bandwidth optimization, and quality of service optimization.

7. Acronyms and Abbreviations

CCITT	Consultative Committee for International Telephony and Telegraphy
CON	Connection-oriented network service
CNLS	Connectionless network service
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
ENUM	Electronic Numbers
ETSI	European Telecommunications Standards Institute
IETF	Internet Engineering Task Force
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ITSP	Internet Telephony Service Provider
ITU-T	International Telecommunication Union — Telecommunications Standards Sector
LAN	Local Area Network
MGCP	Media Gateway Control Protocol
MOS	Mean Opinion Score
PBX	Private Branch Exchange
PDA	Personal Digital Assistant
PESQ	Perceptual Evaluation of Speech Quality
PSQM	Perceptual Speech Quality Measurement
PSTN	Public Switched Telephone Network
QoS	Quality of Service
RAS	Registration, Admission, and Status
RSVP	Resource Reservation Protocol
RTCP	Real-time Control Protocol
RTP	Real-time Transport Protocol
RTSP	Real-time Streaming Protocol
SAP	Session Announcement Protocol
SCCP	Skinny Client Control Protocol
SCTP	Stream Control Transmission Protocol
SDP	Session Description Protocol
SIP	Session Initiation Protocol
TCP	Transmission Control Protocol
TIPHON	Telecommunications and Internet Protocol Harmonization Over Networks
VoIP	Voice over Internet Protocol
WAN	Wide Area Network

8. About the Author and Sponsor

Mark A. Miller, P.E., is President of DigiNet Corporation, a Denver-based consulting engineering firm providing services in internetwork design, strategic planning, network management, and new product development. Mr. Miller is the author of nineteen books on network analysis, design, and management. His latest book is titled *Voice over IP Technologies, Strategies for the Converged Enterprise*, published in 2002 by M&T Books, Inc., a division of John Wiley (Indianapolis, Indiana). He is a frequent presenter at industry events and has taught at the ComNet, CT Expo, Internet Telecom Expo, Networld+Interop, Comdex, and other conferences. He holds B.S. and M.S. degrees in electrical engineering, and is a registered professional engineer in four states. For more information, DigiNet Corporation may be reached at 303-682-5244 or on the Internet at http://www.diginet.com.

Sniffer Technologies, a division of Network Associates, is a leading provider of network and application management solutions designed to ensure e-business uptime. Supporting one of the widest ranges of network topologies in the industry, the Sniffer Total Network Visibility (TNV) suite is an integrated solution enabling enterprises and service providers to cost-effectively keep their networks and applications up and running at peak performance. As one of the most trusted solutions for monitoring, troubleshooting, reporting, and proactively managing network availability and performance, the Sniffer TNV suite meets the demanding 24x7 availability requirements of e-business Web sites, Internet applications, converged voice, video, and data networks, and high speed switched and optical networks. For more information, Sniffer Technologies can be reached on the Internet at http://www.sniffer.com.

With headquarters in Santa Clara, CA, Network Associates, Inc. is a leading supplier of network security and availability solutions. Network Associates is comprised of three product groups: McAfee, delivering world class anti-virus and security products; Sniffer, a leader in network availability and system security; and Magic, providing Web-based service desk solutions. For more information, Network Associates can be reached at 972-308-9960 or on the Internet at http://www.nai.com.

Copyright

Limit of Liability/Disclaimer of Warranty

Information contained in this work has been obtained by the author and sponsor from sources believed to be reliable. However, neither the author nor the sponsors guarantee the accuracy or completeness published herein, and neither the author nor the sponsor shall be responsible for any errors, omissions, or damages arising out of the use of this information. This work is published with the understanding that the author and sponsor are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

Trademarks

DigiNet is a registered trademark of Digital Network Corporation.

Network Associates, Sniffer, Total Network Visibility, TNV, McAfee, and Magic Solutions are registered trademarks of Network Associates, Inc. and/or its affiliates in the United States and/or other countries.

All other registered and unregistered trademarks in this document are the sole property of their respective owners.