AVAYA IP VOICE QUALITY NETWORK REQUIREMENTS

White paper

Issue 1.4

June 2001

Developed by:
Avaya, Inc.
Westminster, Colorado

TRADEMARK NOTICE

This document may contain references to the following Avaya trademarked or registered trademark products**: DEFINITY®, CajunView™,** and **CajunRules™**.

All other product names mentioned herein are the trademarks of their respective owners.

NOTICE

While reasonable efforts were made to ensure the information in this document was complete and accurate at the time of printing, Avaya can assume no responsibility for any errors. Changes and corrections to the information contained in this document may be incorporated into future releases.

Contents

Avaya IP Voice Quality Network Requirements

Document Summary

This document contains minimum basic network requirements to ensure good voice quality when using Avaya IP products and solutions. No document can satisfy the detailed needs of every network, and therefore this paper serves only as a starting point. The document summary provides a short list of networking requirements and recommendations. Use this page as a checklist to determine if your network meets the minimum requirements for implementing Voice over Internet Protocol (VoIP) with acceptable quality. The rest of the document contains networking basics for those who haven't been exposed to networking. It also explains why VoIP applications can yield poor results on data networks that typically run well.

The critical factors in assessing VoIP quality are delay, jitter and packet loss. To ensure good and consistent levels of voice quality, we suggest the following minimum network requirements:

- **Network delay:** Between endpoints, delay should be less than 50ms (milliseconds).
- **Network jitter:** Jitter is a measure of the variability of delay. Between endpoints, jitter should be less than 20ms. This value has some latitude depending on the type of service the jitter buffer has in relationship to other router buffers.
- **Network packet loss:** The maximum loss of packets (or frames) between endpoints should be 0.2% or less.

We also recommend consideration of the following list of options when implementing VoIP.

- **QoS/CoS:** Quality of Service (QoS) for voice packets is obtained only after a Class of Service (CoS) mechanism tags voice packets as having priority over data packets. Networks with periods of congestion can still provide excellent voice quality when using a QoS/CoS policy. Switched networks should use IEEE 802.1p/Q. Routed networks should use DSCP (DiffServ Code Points). Mixed networks should use both. Port priority can also be used to enhance DiffServ and IEEE 802.1p/Q. See sections 2.1 - 2.4 for more information.
- **Switched Network:** A fully switched network is a network that allows full duplex and full endpoint bandwidth for every endpoint that exists on the LAN. Although VoIP systems can work in a shared (hubs or bussed) LAN, Avaya recommends the consistently high results a switched network lends to VoIP.
- **Network Assessment:** A Basic Network Readiness Assessment Offer from Avaya is vital to a successful implementation of VoIP products and solutions. Contact your Avaya representative or your authorized dealer to review or certify your network. Section 4 "Network Assessment" explains the options available with this offer.
- **VLANs:** Although an option, placing voice packets on a separate VLAN (subnet) from data packets is a generally accepted practice to reduce broadcast and data traffic from contending for the same bandwidth as voice. Section 2.5 "Using VLANs" further explains this concept. Note that Avaya IP telephones provide industry leading broadcast storm protection.

We recommend caution when using the following:

- **NAT:** Be very careful when using NAT (Network Address Translation). Most implementations using VoIP endpoints behind NAT fail because many H.323 messages (the protocol carrying the voice information) contain multiple instances of the same IP address in a given message, but NAT is unlikely to find and translate all of them. See section 7.4, "NAT (Network Address Translation)" and Appendix C for more information on using NAT with VoIP.
- **Analog Dial-Up:** Be careful in using analog dial-up (bandwidth ≤ 56K) to connect two locations. Upstream bandwidth is limited to a maximum of 33.6K, and in most cases is less. This results in insufficient bandwidth to provide toll-quality voice. Some codecs and network parameters provide connections that are acceptable, but consider each connection individually.
- **VPN:** Use Virtual Private Network (VPN) cautiously with VoIP applications. Large delays are inherent in some VPN products due to encryption, decryption and additional encapsulation. Some hardware-based products encrypt at near wire speed and can be used. In addition, if the VPN is run over the Internet, sufficient quality for voice cannot be guaranteed unless delay, jitter and packet loss are contained within the parameters listed above. See section 7.4 "VPN (Virtual Private Network)" for more information.

This document changes frequently to keep up with advances in networking and VoIP technology. Consult your Avaya representative or authorized dealer for updates or visit us at http://support.avaya.com/comsys/definity/dolan/ta/4564.jhtml

Comments or questions may be emailed to: afunguy@Avaya.com

Avaya IP Voice Quality Network Requirements

Explanations

1 Introduction

Voice over Internet Protocol (VoIP) is the convergence of traditional voice onto the data network to lower costs and mainly, to provide many more applications to telephony through the integration of open-standards computer programming. Other real-time traffic, such as uncompressed video and streaming audio, is also converging onto data networks.

VoIP is very complex because it involves components of both the data and the voice world. Historically, these worlds have used two different networks, two different support organizations and two different philosophies. The voice network has always been separate from the data network because the characteristics of voice applications are very different from the characteristics of data applications.

The traditional voice network is circuit switched. Interactive voice traffic is sensitive to delay and jitter but can tolerate some packet loss. The voice philosophy was to ensure the "five nines" of reliability – 99.999%, because the lack of communication might be life threatening, (i.e. the inability of placing a "911" call for help). Voice calls have their own dedicated bandwidth throughout the circuit-switched network, so delay is rarely an issue.

The data network, on the other hand, is packet switched. Data is less sensitive to delay and jitter, but cannot tolerate loss. The data philosophy has been concerned with providing reliable data transmission over unreliable media, regardless of delay. Bandwidth in the data world is largely shared, so congestion and delay are often present for multimedia applications like voice.

The factors that affect the quality of data transmission are different from the factors that affect the quality of voice transmission. For example, data is generally not affected by delay. Voice transmissions are degraded by even small amounts of extra delay and cannot be retransmitted. Additionally, a tiny amount of packet (data) loss does not affect voice quality at the receiver's ear. But even a small loss of data can corrupt an entire file or computer application. So in some cases, introducing VoIP to a high performing data network can yield very poor voice quality.

Therefore, implementing VoIP requires attention to many factors, including:

- Delay
- Packet loss
- Available bandwidth
- Network design
- Duplex
- Echo
- Codec selection
- Reliability
- Manageability
- QoS/CoS policy
- Jitter
- Packet mis-order
- Packet prioritization
- Endpoint audio characteristics (sound card, microphone, earpiece, etc.)
- Transcoding
- Silence suppression
- Router and data-switch setup
- Scalability
- WAN protocols
- Encryption/Decryption

This document provides basic network guidelines that must be addressed to ensure good voice quality when implementing VoIP. This document also examines some of the more important components that affect VoIP and gives suggestions to help avoid problems during implementation.

2 Prioritizing Voice Traffic

In order for a VoIP solution to function well, the network must be able to give voice packets priority over ordinary data packets. Avaya's products for VoIP—the flagship DEFINITY® Enterprise Communication Server (ECS) R9, the new IP 600, and our Cajun® line of data switches—all include several standard strategies to prioritize voice traffic. These strategies include using class of service (CoS), prioritizing ports, prioritizing services, and using IEEE 802.1p/Q to set the priority bits. Our products also work with all the other popular switches and routers through open standards to provide end-to-end voice prioritization.

2.1 Understanding QoS versus CoS

Class of Service (CoS) is a classification method only. CoS does NOT ensure a quality of service (QoS), but is the method used for queuing mechanisms to limit delay and other factors to improve QoS. Most CoS strategies assign a priority level, usually 0–7 or 0-63, to a frame or packet respectively. Common CoS models include the IP TOS (Type Of Service) byte, Differentiated Services Code Point (DiffServ or DSCP, defined in RFC 2474 and others) and the IEEE 802.1p/Q.

Quality of Service (QoS) involves giving preferential treatment through queuing, bandwidth reservation, or other methods based on attributes of the packet, such as CoS priority. A service quality is then negotiated. Examples of QoS are ATM (Asynchronous Transfer Mode), RSVP (RESERVATION Protocol (RFC 2205)) and MPLS (Multi Protocol Label Switching (RFC 1117 and others)).

2.2 Using Ports

One prioritization scheme assigns priority based on the UDP (User Datagram Protocol) port numbers that the voice packets use. This scheme allows one to use network equipment that can prioritize all packets from a port range. UDP is used to transport voice through the LAN because, unlike TCP, it is not connection-based. Because of the human ear's sensitivity to delay, it is better to drop packets rather than retransmit voice in a real time environment. So, a connectionless protocol is preferable to a connection-based protocol. With DEFINITY ECS R9 and IP600, one can define any port range for voice priority. Routers and layer 3 data switches can use these ports to distinguish priority traffic. This priority traffic can be voice packets or signaling packets. This is an OSI model layer 4 solution and works on data coming to and from the specified port range.

2.3 Using DSCP (or TOS)

The Differentiated Services Code Point (DSCP) prioritization scheme redefines the Type of Service (TOS) byte in the IP header by combining the first six bits into 64 possible combinations. This use of the TOS byte is still evolving but can be used now by the DEFINITY ECS R9, IP600, IP Telephone, and other network elements such as routers and switches in the LAN. A DSCP of 40 (101000) is suggested for the expedited forwarding of packets. DSCP 40 is backward compatible with a TOS precedence definition of critical (5). With DEFINITY ECS R9, IP600, and IP Telephone, one can set any of these bits as desired.

The original TOS method uses the TOS byte to assign a class of service as defined in RFC 795. This IP byte can be used in the traditional way by setting the precedence bits (3) giving eight classes of service. Four other bits define delay (normal or low), throughput (high or normal), reliability (high or normal), and cost (normal or low). With DEFINITY ECS R9 and IP600, one can set any of these bits as desired. The TOS byte is an OSI model layer 3 solution and works on IP packets.

2.4 Using IEEE 802.1 p/Q

Yet another prioritization scheme is the IEEE 802.1 standard, which uses four bytes to augment the layer 2 header. The IEEE 802.1Q standard defines the open standard for VLAN tagging. Two bytes house 12 bits used to tag each frame with a VLAN identification number. The IEEE 802.1p standard uses three of the remaining bits in the 802.1Q header to assign one of eight different classes of service. Again, with DEFINITY ECS R9, the IP600, and the IP Telephone, one can add the 802.1Q bytes and set the priority bits as you wish. The Avaya Cajun line of data switches can switch frames with or without these VLAN headers. IEEE 802.1p and IEEE 802.1Q are OSI layer 2 solutions and work on frames.

2.5 Using VLANs

VLANs provide security and create smaller broadcast domains through software by creating virtually separated subnets. Broadcasts are a natural occurrence in most data networks from protocols used by PCs, servers, switches and routers. Creating a separate VLAN for voice reduces the amount of broadcast traffic the telephone will receive. Separate VLANs result in more effective bandwidth utilization and reduces the processor burden on endpoints by freeing them from having to analyze irrelevant packets. VLANs, a layer-2 feature, are defined in data switches and a VLAN for voice can also be specified from a list on the switch port from the IP telephone. Separate voice and data VLANs are an option that makes sense for most customers.

3 Network Parameters

There are a number of network parameters that affect voice quality. This section lists some of the more important ones.

3.1 Network Packet Delay

Packet delay is the length of time it takes a packet to traverse the network. Users will experience difficulties in carrying on a normal conversation when the one-way network delay exceeds 50 milliseconds (ms). Each element of the network adds to packet delay including switches, routers, distance traveled through the network, firewalls, and jitter buffers (such as those built into H.323 audio applications like the Avaya IP SoftPhone or Microsoft® NetMeeting®). Packet delay in excess of 50 ms can have a noticeable effect. However, some applications or users may elect to tolerate it, just as many people accept substandard quality when using cell phones.

Router delay depends not only on hardware, but also on configurations such as access lists, queuing methods, and transmission modes. Delay (latency) can be controlled somewhat in a private environment (LAN) because the company or enterprise controls the network infrastructure. When using the public network, there are inherent delays that one cannot control. Networks with more than 50ms of one-way delay between endpoints will most likely yield voice quality that is distracting to most users. We highly recommend a network assessment from Avaya to measure latency and make recommendations to solve any latency issues before implementing a VoIP solution.

3.2 Network Jitter

Jitter is a measure of variance in the time it takes for communications to traverse from the sender (application) to the receiver, as seen from the application layer (from RFC 2729 Taxonomy of Communication Requirements for Large-scale Multicast Applications). We tend to think of jitter as the statistical average variance in delivery time between packets or datagrams.

Jitter can create audible voice-quality problems if the variation is greater than 20ms. Symptoms of excessive jitter are very similar to symptoms of high delay, because in both cases packets are discarded if the packet delay exceeds half the jitter buffer size.

To compensate for network jitter, many vendors implement a jitter buffer in their H.323 voice applications. The purpose of the jitter buffer is to hold incoming packets for a specified period of time before forwarding them to the decompression process. A jitter buffer is designed to smooth packet flow. In doing so, it can also add significant packet delay.

Jitter buffers should be dynamic to give the best quality, or if static, should generally be sized to twice the largest statistical variance between packets. Router vendors have many queuing methods that alter the behavior of the jitter buffer. It is not enough to just select the right size of jitter buffer, one must also pair an appropriate queue-unloading algorithm type with the jitter buffer. The network topology can also affect jitter. Because there are fewer collisions on a hierarchical data-switched network than on a flat hub-based network, there will be less jitter on the switched network.

The Avaya DEFINITY ECS, IP600, IP SoftPhone software and IP telephone have all incorporated dynamic jitter buffers to minimize delay by reducing the jitter buffer size as the network allows. Note that this feature can exacerbate problems in an uncontrolled network. Many good tools are commercially available to measure jitter, delay, and packet loss to help monitor network trends and bring control to your network.

3.3 Packet Loss

Network packet loss is when packets are sent, but not received at the final destination due to some network problem. To ensure good quality voice in a VoIP network, packet loss should be less than 0.2% in the network – again between endpoints. There are several factors that make packet loss requirements somewhat variable, such as the following:

- Packet loss requirements are tighter for tones (other than DTMF) than for voice. The ear is less able to detect packet loss during speech (variable-pitch), than it is during a tone (consistent pitch).
- Packet loss requirements are tighter for short, continuous packet loss than for random packet loss over time. Losing ten contiguous packets is worse than losing ten packets evenly spaced over an hour time span.
- Packet loss may be more noticeable for larger voice payloads than for smaller ones, because more voice is lost in a larger payload.

Tools such as the Agilent (HP) Internet Advisor, Shomiti Systems™ Explorer, Radcom's Prism, and others measure packet loss. Remember that too much delay *can cause* dropped packets, and it may appear the network is losing packets when in fact they have been discarded intentionally.

3.4 Network Packet Mis-Order

Network packet mis-order is, for voice over IP, very much like packet loss. If a packet arrives out of order, it is generally discarded, as it makes no sense to play it out of order. Specifically, packets are discarded when they arrive later than the jitter buffer can hold them. Mis-order can occur when networks send individual packets over different routes. Planned events like load-balancing or unplanned events such as re-routing due to congestion, or other transient difficulties can cause packet mis-order. Packets traversing the network over different routes may arrive at their destination out of order. Network latency over multiple yet unequal routing paths can also force packet mis-order.

3.5 Transcoding

Transcoding is a voice signal converted from analog to digital or digital to analog (possibly with or without compression and decompression). If calls are routed using multiple voice coders, as in the case of call coverage on an intermediary system back to a centralized voice mail system, the calls may experience multiple transcodings (including the one in and out of the voice mailbox). Each transcoding episode results in some degradation of voice quality. These problems may be minimized by the use of the DEFINITY ECS feature called DCS with Rerouting (Path Replacement). This feature detects that the call coming through the main ECS has been routed from one tandem ECS, through the main, and back out to a third switch. In these cases, the system then re-routes the call directly, thus replacing the path through the main system with a more direct connection. Avaya products minimize transcoding while non-Avaya products may cause slight to excessive transcoding.

3.6 Echo

Sources of echo are many. The two main types of echo are acoustic and impedance. Echo will result when a VoIP call leaves the LAN through a mis-administered analog trunk into the PSTN. Another major cause is from an impedance mismatch between four-wire and two wire systems. Echo also results when an impedance mismatch exists in the conversion between the TDM (Time Division Multiplexing) bus and the LAN, or the impedance mis-match between a headset and its adapter. Impedance mis-match causes inefficient energy transfer. The energy imbalance must go somewhere and so it is reflected back in the form of an echo. Usually the speaker hears an echo but the receiver does not.

Echo cancellers, which have varying amounts of memory, compare the received voice with the current voice patterns. If the patterns match, the canceller cancels the echo. Echo cancellers aren't perfect, however. Under some circumstances, the echo gets past the canceller. The problem is exacerbated in VoIP systems. If the one-way trip delay between endpoints is larger than the echo canceller memory, the echo canceller won't ever find a pattern to cancel. Avaya's DEFINITY ECS, IP600, and IP SoftPhone software incorporate echo cancellers designed for VoIP to improve voice quality.

3.7 Silence Suppression and Voice Activity Detection

Voice Activity Detection (VAD) monitors the received signal for voice activity. When no activity is detected for the configured period of time, the software informs the Packet Voice Protocol. This prevents the encoder output from being transported across the network when there is silence, resulting in additional bandwidth savings. This software also measures the idle noise characteristics of the telephony interface. It reports this information to the Packet Voice Protocol to relay this information to the remote end for noise generation when no voice is present. Aggressive VADs cause voice clipping and can result in poor voice quality, but the use of VAD can greatly conserve bandwidth and is therefore a very important detail to consider when planning network bandwidth – especially in the WAN (Wide Area Network). Avaya's DEFINITY ECS, IP600, and IP SoftPhone products all can employ silence suppression to preserve vital bandwidth.

3.8 Network Duplex

The ideal network for transporting VoIP traffic is a network that is fully LAN switched from end-to-end because it significantly reduces or eliminates collisions. A network that has shared segments (hub-based) can result in lower voice quality due to excessive collisions.

Although there are many different brands and models of data switches available, the Avaya Cajun line of switches are specifically designed to enable and enhance VoIP quality throughout your network.

3.9 Codec Selection

Depending upon the bandwidth availability and required voice quality, it might be worthwhile to select a codec that produces compressed audio.

- A G.711 codec produces audio uncompressed to 64 kbps
- A G.729 codec produces audio compressed to 8 kbps
- A G.723 codec produces audio compressed to approximately 6 kbps

The following table provides comparison of several voice quality considerations associated with some of the codecs supported by Avaya products. It should be noted that toll-quality voice must achieve a MOS (Mean Opinion Score) of 4 or above. MOS scoring is a long-standing subjective method of measuring voice quality.

Table 1. Comparison of Speech Coding Standards [1]

Standard	Coding Type	Bit Rate (kbps)	MOS
G.711	PCM	64	4.3
G.729	CS-ACELP	8	4.0
G.723.1	ACELP MP-MLQ	6.3 5.3	3.8

4 Network Assessment

The Avaya Basic Network Readiness Assessment Offer (OA2000-064) is designed to provide assurance to Avaya customers that their data network is capable of supporting Voice over IP (VoIP) applications before installing the Avaya application. This Basic Network Readiness Assessment Offer is a flexible process, allowing the customer to choose to:

- Provide the required network assessment data themselves to Avaya or through their network vendor
- Have Avaya collect the required information using the Basic Network Assessment tools
- Commission an Avaya Professional Services Network Assessment and Optimization (NANO) Study

Working remotely, and using a combination of interactive questionnaires and innovative software tools, Avaya eCommunication Professional Services (eComProS) network engineers will:

1. Identify all equipment in the customer's network, as well as physical and network layer information, device connections, network topology and device configurations through a Site Configuration Survey.
2. Test the customer's current network infrastructure to discover any throughput and response time issues in multi-protocol networks using VitalAgent software.
3. Baseline existing application performance measures.

[1] Table 1: Rudkin, S. Grace, A., and Whybray, M. W., "Real-Time Applications on the Internet," BT Journal, Col. 15, No. 2, April 1997.

4. Ensure that voice traffic will receive proper prioritization in the network by verifying existing prioritization schemes and recommending improvements when the existing schemes are insufficient.
5. Provide a baseline to compare pre-implementation and post-implementation views of the customer's network.

Customers who do not avail themselves of this offer assume responsibility for all network-related problems. Also, Avaya personnel may be required to charge a higher T&M rate if assistance is requested later, since troubleshooting will be more difficult without the assessment data.

You can initiate a network assessment by contacting your Avaya representative or authorized dealer.

5 PC Considerations using IP Softphone

IP SoftPhone is software on a PC that simulates a telephone. The "perceived" audio/voice quality at the PC endpoint is a function of at least four factors:

1. Transducer Quality
   The selection of speaker and microphone or headset has an impact on the reproduction of the sound.

2. Sound Card Quality
   There are several parameters that affect sound card quality. The most important is whether or not the sound card supports full-duplex operation.

3. End-to-End Delay
   A PC can be a major component of delay in a conversation. PC delay consists of the jitter buffer and sound system delays, as well as the number of other processes running and the speed of the processor.

4. Speech Breakup
   Speech breakup may be the result of a number of factors:

   - Network jitter in excess of the jitter buffer size
   - Loss of packets (due to excessive delay, etc.)
   - Aggressiveness of Silence Suppression

     In an effort to reduce network load, silence suppression is used to eliminate the transmission of silence. However, some silence suppression algorithms may clip speech and have an effect on perceived audio quality.

   - Performance bottleneck in the PC

     Lower speed PCs (or PCs with slow hard drives) may have adverse interactions with sound playback and recording. This can cause breaks in received or transmitted audio. The best thing to do in this situation is to increase the processor speed, increase the amount of RAM and/or reduce the number of applications competing for the processor or hard drive resources. One notable resource consumer is the Microsoft® Find Fast program that launches from the Startup folder (and runs in the background). This application periodically re-indexes the hard drive and consumes significant PC resources in the process.

6 Bandwidth Requirements

The bandwidth available to the user is very important. Access to the network using slower connections, such as dial-up connections, will degrade voice quality. The best voice quality is achieved in both LANs and WANs when the bandwidth is “owned” by the customer.

Customer-owned bandwidth can be shaped to optimize VoIP traffic. Conversely, bandwidth that is not controlled, like the Internet, cannot give consistent sound quality because it cannot be optimized for VoIP. Because factors of delay, jitter, and packet loss are exacerbated over the Internet, we do not recommend using the Internet for voice applications at this time.

The following table summarizes some common bandwidth requirements. The first charts is for the LAN and the second and third are used for the WAN. **Note that all numbers are approximated** because there are so many variables like payload size, actual transmission rates, duplex, etc.

Table 2. LAN Bandwidth Requirements

Codec Type	Data Rate	Data Bytes per 30ms voice frame	Total L2 Frame Size in Bytes	No Silence Suppressed in Kbps	1 Side Silence Suppressed in Kbps	Both Sides Silence Suppressed in Kbps
G.711	64 Kbps	240	298	158.93	119.20	79.47
G.729	8 Kbps	30	88	46.93	35.20	23.47
G.723	5.3 Kbps	20	78	41.60	31.20	20.80

Table 3. WAN Bandwidth Requirements

Codec Type	Payload Size in Bytes	Full Rate PPP or FRF.12	PPP with CRTP	PPP with VAD	PPP, VAD & CRTP
G.711 64 Kbps	240	76 Kbps	66 Kbps	50 Kbps	43 Kbps
G.729 8 Kbps	20	26.4 Kbps	11.2 Kbps	17.2 Kbps	7.3 Kbps
G.723.1 5.3 Kbps	20	17.5 Kbps	7.4 Kbps	11.4 Kbps	4.8 Kbps

Table 4. WAN Transfer Protocol Bandwidth Requirements

All values use a 60ms voice sample.

	Codec Rate	VoIP over 802.3	Voice over AAL-2	VoIP over AAL-5	Voice over Frame	VoIP over Frame
G.711	64 Kbps	71.7 Kbps	N/A	77.7 Kbps	65.3 Kbps	70.7 Kbps
G.729	8 Kbps	15.7 Kbps	9.5 Kbps	21.2 Kbps	9.2 Kbps	14.5 Kbps
G.723	6.4 Kbps	14.1 Kbps	7.7 Kbps	14.1 Kbps	7.6 Kbps	13 Kbps

The table above suggests the differences in using WAN transport protocols like ATM (AAL-2 and AAL-5), Voice over Frame Relay (VoFR), and VoIP over Frame Relay. The chart is not exhaustive but highlights the fact that bandwidth varies by using different protocols. VAD and header compression are not factors in this chart, but can complicate the requirements significantly.

6.1 Dual Connect Bandwidth Requirements using IP SoftPhone (or IP Agent)

A dual connect system is commonly used in a Call Center for users working remotely. The PC and the telephone can transmit frames across the same telephone line or on two lines. Questions concerning the amount of bandwidth the PC uses and its effect on voice are answered here. The bandwidth used by the PC for signaling is very low. However, it is difficult to express this value in bits per second due to the variability in how quickly the buttons are pressed and how many feature buttons are used during a call. The following graph is a 50 second “average” call showing

the bandwidth needed with several buttons pushed. Remember that even with a 56K (V.90) modem the upstream bandwidth is no greater than 33.6K and the downstream is anywhere from 28.8K to 53K. The speed of each connection is determined by the PSTN line conditions at the time the call is placed.

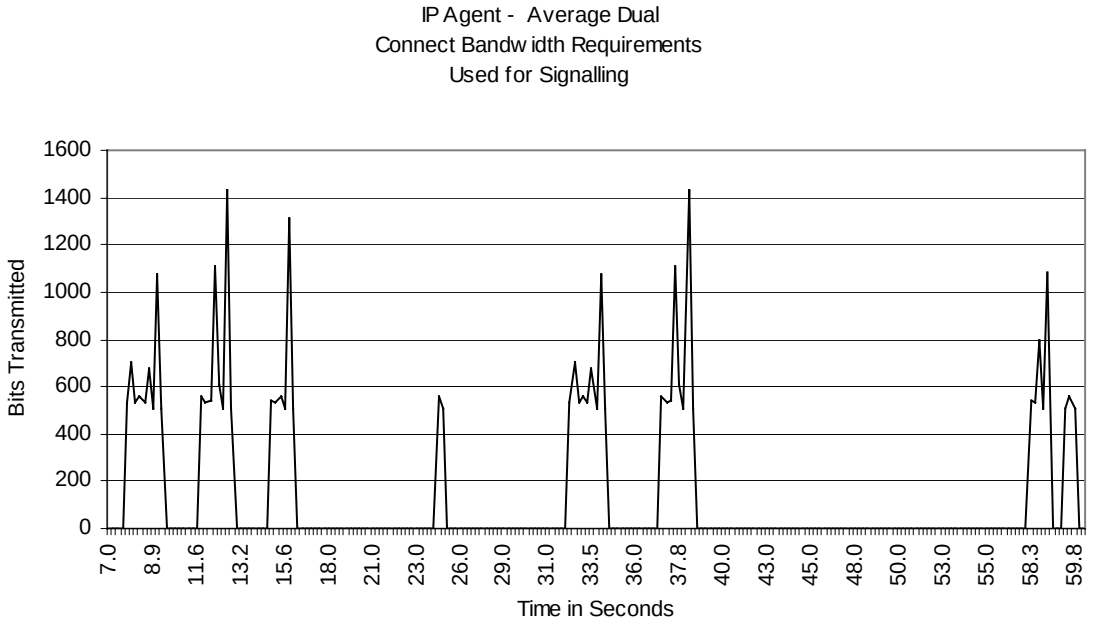

Note that during most of this call the bandwidth required is zero (X Axis). The maximum bandwidth needed is never greater than 1.450 Kilobits at any one point in time. This is small compared to even a slow 28.8 Kilobit transfer rate as it represents less than 5% of the 28.8Kbs available bandwidth at any point in time. Bandwidth required for signaling is almost moot compared to the available bandwidth for voice.

7 Other Elements that Affect VoIP

7.1 WAN Considerations

Until WAN bandwidth becomes affordable at any speed, delivering bandwidth to applications over the WAN will remain a formidable task. When voice traffic is carried on packet networks, different labeling or queuing schemes function to give voice packets priority over data packets. The presence of large data packets may result in added serialization delay for VoIP packets across WAN links. This is due to the fact that smaller VoIP packets are held in queue while larger data packets are processed onto the WAN link. To avoid excessive delay, there may be benefit to fragmenting the larger data packets and interleaving them with the smaller voice packets.

One technique is to adjust the packets by adjusting the Maximum Transmission Unit (MTU) size. Minimum MTU size should be no smaller than 300 bytes and no larger than 550 bytes. LAN based MTUs can be as large as 1500 bytes. Note: reducing the size of the MTU will add overhead and reduce the efficiency of data applications. Other techniques, such as Multilink PPP (MPP) Link Fragmenting and Interleaving (LFI), and Frame Relay Fragmentation (FRF12) allow network managers to fragment larger packets, and allow queuing mechanisms to speed the delivery of Real Time Protocol (RTP) traffic without significantly increasing protocol overhead

or reducing data efficiency. Also, header compression protocols like CRTP (Compressed Real Time Protocol) can and should be used between WAN links.

7.2 VPN (Virtual Private Network)

There are many definitions for Virtual Private Networks (VPN). In this white paper, VPNs refer to encrypted tunnels carrying packetized data between remote sites. VPNs can use private lines or use the Internet via one or more Internet Service Providers (ISP). VPNs are implemented in both dedicated hardware and software, but can also be integrated as an application to existing hardware and software packages. A common example of an integrated package is a firewall product that can provide a barrier against unauthorized intrusion, as well as perform the security features needed for a VPN session.

The encryption process can take from less than 1 milli-second to 1 second or more, at each end. Obviously, VPNs can represent a significant source of delay and, therefore, negatively affect voice performance. Also, as most VPN traffic runs over the Internet and there is little control over QoS parameters for traffic crossing the Internet, voice quality may suffer due to excessive packet loss, delay, and jitter. You may be able to negotiate a service-level agreement with your VPN provider to guarantee an acceptable level of service. Before implementing VoIP with a VPN, you should test your VPN network to make sure it meets the requirements specified in the Document Summary. For more information, see Avaya's VPN white paper, available from your Avaya representative.

Appendix C contains more information on VPN usage and impacts.

7.3 Frame Relay

Information transported over frame relay is subject to more delay and jitter when compared to ATM or point-to-point TDM circuits. This is due to many factors, which are not covered in detail here. Instead we offer remedies to protect voice traffic from the susceptibilities of frame relay in Appendix B.

7.4 NAT (Network Address Translation)

VoIP does not work well with networks that use NAT (Network Address Translation) because most NAT implementations do not support H.323 protocols. The destination IP address is encapsulated in more than one header: the Q.931, H.225, and IP headers. NAT changes only the address in the IP header resulting in a mismatch that prohibits the control of calls. We suggest that using a firewall to guard against intruders, but your firewall should not provide NAT functions for VoIP packets unless it is Q.931 friendly like the Lucent 201 Brick. Appendix C shows an approved sample implementation of a firewall using selective NAT.

8 Avaya VoIP Products

8.1 DEFINITY ECS Release 9 (G3r and G3si)

Our flagship DEFINITY Enterprise Communication Server (ECS) is IP enabled through the use of Release 9 (R9) software and two IP telephony circuit packs. The DEFINITY ECS R9 allows the use of our IP SoftPhone® software on PCs and works seamlessly with our 4600-series IP telephones, as well as all the other digital and analog endpoints you may already have. This award-winning package includes administration and troubleshooting software. R9 supports voice priority using ports, TOS/DSCP, and IEEE 802.1 p/Q.

8.2 IP600

This rack-mounted server is an all IP PBX for small and mid-sized customers. It uses the same R9 software load as its bigger brother, the DEFINITY ECS. It also works with all the other endpoints that the DEFINITY ECS supports. One IP600, as with any Definity product, can support up to 256 R300 units. For more information go to http://www1.avaya.com/enterprise/who/docs/ip600/

8.3 Media Processor Circuit Pack

The TN2302AP media processor circuit pack can be connected to either a 10BaseT or 100BaseT network. However, to fully utilize the port capacity of the TN2302AP circuit pack, it must be connected to a 100BaseT data-switched network port and use a Category 5 compliant cable. If auto-negotiation is not used, please set the switched port to 100Mbps / full duplex. This circuit pack is a media processor board that converts analog voice to data packets, supports up to 64 simultaneous ports using G.711 (32 ports using G.729/723 and somewhere between if using a mix of G.711 and G.729/723), and dynamically switches between codecs.

Perhaps the most exciting feature of this circuit pack is the ability to shuffle or hairpin calls. To shuffle means to reroute the voice channel connecting two IP endpoints so that the voice, which previously was endpoint to DEFINITY to endpoint, now goes directly from endpoint to endpoint. This also works in reverse for reasons such as conferencing, placing on hold, etc. Hairpinning reroutes the voice channel connecting two IP endpoints so that the voice goes through the TN2302AP board in IP format, without having to go through the DEFINITY TDM bus. Both endpoints must use the same codec for this feature to work. Hairpinning and shuffling mean less delay between endpoints and fewer resources used on the TN2302AP. A shuffled call that is IP-to-IP incurs no delay or transcoding whatsoever from the DEFINITY ECS or IP600.

8.4 Control LAN Circuit Pack

The TN799 Control LAN circuit pack (C-LAN) controls signaling and call setup. This board may be connected to a hub (half-duplex, 10Mbps), but works better in an all switched environment (still at 10Mbps, half-duplex). This circuit pack controls all IP call establishment/release and shuffled/hairpinned calls so that the DEFINITY ECS or IP600 can keep voice quality high even with changing network conditions. One C-LAN board can support over 400 port connections. Multiple C-LAN boards provide load balancing and will automatically cover for a failed C-LAN board in the group.

8.5 R300 Remote Office

This rack-mounted, pizza-sized box connects a remote office to the main office with IP over a WAN protocol. It supports up to 24 digital phones and 2 analog phones. It also offers the ability to connect local CO trunks. The R300 is able to route IP over Ethernet-based LANs and several WAN transport protocols such as PPP, ISDN, Frame Relay, T-1, E-1, or BRI circuits. It also features RIP and OSPF routing protocols, DHCP and DNS caching. For more information go to http://www1.avaya.com/enterprise/who/docs/r300/

8.6 IP SoftPhone

The Avaya IP SoftPhone is a client-based telephony application that provides best-in-class voice over IP quality. This LDAP client enables CTI/TAPI and supports dual or single connect for toll-class audio quality. Load balancing across multiple C-LAN cards and receiving QoS parameters is now available. For more information go to http://www1.avaya.com/enterprise/solutions/eclips/product_3e.html

8.7 IP Telephone

The Avaya 46XX IP Telephone looks and feels just like a circuit-switched set because it supports most of the features of our digital sets. This family of phones supports IEEE802.1Q/p, DiffServ and a separate VLAN for voice traffic. It can withstand over 1,000 broadcasts per second making it an industry leading phone in resisting broadcast storms. It also has a full-duplex speakerphone. It also supports traffic at 10 or 100 Mbps and can use DHCP (Dynamic Host Configuration Protocol) for easy setup. If auto-negotiation is not used, the switched port should be 10 or 100 Mbps and half-duplex. More information can be found at http://support.avaya.com/elmodocs2/avayaip/

8.8 Cajun Switches

Avaya's Cajun line of data switches won top honors in a competitive review of high-end LAN switching products. The P330 line is stackable up to ten switches that act as one virtual switch. The P580 and P882 support larger enterprises. Cajun switches have multiple priority queues, use SMON and support mapping DSCP to IEEE 802.1 p/Q priority values. Cajun Rules is a policy management system that defines QoS policies for users, groups of users and applications. Cajun View monitors QoS and other parameters of the Cajun line of switches. For more information go to http://www1.avaya.com/enterprise/who/docs/product12.html

8.9 VPNet

Avaya's VPN solutions are found in this series of products for small business, enterprise business and even carrier class ISPs. These products are hardware based to give performance that is truly wire speed using 3DES and other security measures. Software interfaces allow monitoring of any existing VPN connection and easy setup of new connections. Although these products are hardware based, a software client is available for "Road Warriors" or people that need to work securely from home. For more information go to http://www1.avaya.com/enterprise/who/docs/product12.html

9 VoIP Tools

9.1 Network Tools

Many tools are available to determine latency, jitter, and packet loss on IP networks. Tools fall into several categories: reactive and proactive, passive and active. Passive tools are reactive and "sniff" the network to display or capture existing (real) traffic. Active tools inject packets into the network to test network characteristics (proactive) or to stress test specific network elements (proactive). Modeling tools are also proactive because they model future "what-if" scenarios without inducing a load on the network. A partial list of these commercial tools is listed in Appendix D. These tools are available for purchase through their respective vendors and have been found to be very useful for diagnoses, analysis, modeling and monitoring networks and VoIP conversations. None of these tools are specifically endorsed or explicitly warranted by Avaya Inc. They merely represents a starting list of tools that fit the active, passive and modeling categories that are needed to properly assess networks and network products. Other tools exist that may be a better fit for your organization.

Appendix A

Network Design Recommendations

In the early days of networking, network designers used hubs to attach servers and workstations, and routers to segment the network into manageable pieces. Because of the high cost of router interfaces and the inherent limitations of shared-media hubs, network design was generally well done. In recent years, with the rise of switches to segment networks, designers could hide a number of faults in their networks and still get good performance. As a result, network design has suffered.

VoIP will place new demands on the network. Sub-optimal designs will not be able to cope with these demands. Even with switches installed, a company must pay attention to industry "best practices" in order to have a properly functioning voice network. Because users will not tolerate poor voice quality, administrators must implement a sound network before beginning VoIP pilots or deployments.

Best practices

Industry best practices dictate that a network be designed with the following factors in mind:

- Reliability/redundancy
- Scalability
- Manageability
- Bandwidth

Voice mandates the following additional considerations when designing a network:

- Delay
- Jitter
- Loss
- Duplex

Generally speaking, these concerns dictate a hierarchical network consisting of at most three layers: core, distribution, and access. Some smaller networks can collapse the functions of several layers into one device.

The core layer is the heart of the network. Its purpose is to forward packets as quickly as possible. It needs to be designed with high availability in mind. Generally, these high-availability features include redundant devices, redundant power supplies, redundant processors, and redundant links. In the current era, core interconnections increasingly use Gigabit Ethernet.

The distribution layer links the access layer with the core. It is here that QoS feature and access-lists are applied. Generally, Gigabit Ethernet connects to the core and either Gigabit Ethernet or 100base-TX/FX links connect the access layer. Redundancy is important at this layer, but not as important as in the core.

The access layer connects servers and workstations. Switches at this layer are smaller, usually 24-48 ports. Desktop computers and workstations are usually connected at 10 Mbps, (or 100Mbps) and servers are connected at 100 Mbps, (or 1 Gbps). Limited redundancy is used. Some QoS and security features can be implemented here.

For VoIP to work well, WAN links must be properly sized with sufficient bandwidth for voice and data traffic. Each voice call uses between 6.3 Kbps and 80 Kbps, depending on the desired codec, quality and header compression used. G.729 is one of the most promising standards today, using 24 Kbps of bandwidth uncompressed. Interoffice bandwidth demands can be sized using traditional phone metrics such as average call volume, peak volume, and average call length.

Quality of Service also becomes increasingly important with WAN circuits. In this case, Quality of Service can be taken to mean classification and prioritization of voice traffic. Voice traffic must be given absolute priority through the WAN, and if links are not properly sized or queuing strategies are not properly implemented, it will become evident both with the quality and timeliness of voice and data traffic.

There are three technologies that work well with VoIP: ATM, Frame Relay, and point-to-point (PPP) circuits. These technologies all have good throughput, low latency, and low jitter. ATM has the added benefit of enhanced QoS. Frame Relay and PPP links are more economical, but lack some of the traffic-shaping features of ATM.

Of the three technologies, Frame Relay is the most difficult WAN circuit to use with VoIP. Congestion in Frame Relay networks can cause frame loss, which can significantly degrade the quality of VoIP conversations. With Frame Relay, proper sizing of the CIR (committed information rate) is critical. In a Frame Relay network, any traffic exceeding the CIR is marked discard eligible, and will be discarded at the carrier's option if it experiences congestion in its switches. It is very important that voice packets not be dropped. Therefore, CIR should be sized to average traffic usage. Usually, 25% of peak bandwidth is sufficient. Also, Service Level Agreements (SLAs) should be established with the carrier that defines maximum levels of delay and frame loss, and remediation should the agreed-to levels not be met.

Network management is another important area to consider when implementing VoIP. Because of the stringent requirements imposed by VoIP, it is critical to have an end-to-end view of the network and ways to implement QoS policies globally. Products such as HP OpenView Network Node Manager, CajunRules™, CajunView™, Concord NetHealth, and MRTG will help administrators maintain acceptable service. Should a company not have the resources to implement and maintain network management, outsource companies are springing up to assist with this need.

Avaya offers network assessment and redesign services, should they be necessary.

Common issues

Some common "bad habits" that will severely impact network performance, especially when using VoIP include:

- Using a flat, non-hierarchical network (e.g. cascading small workgroup switches together): This technique quickly results in bottlenecks, as all traffic must flow across the uplinks (at maximum 1Gbps) versus traversing switch fabric (up to 256 Gbps). The greater the number of small switches (layers), the greater the number of uplinks, and the lower the bandwidth for an individual connection. Under a network of this type, voice performance quickly degrades to an unacceptable level.
- Multiple subnets on a VLAN: A network of this type will have issues with broadcasts, multicasts, and routing protocol updates. It should be avoided. It will greatly impact voice performance and complicate troubleshooting issues.
- Hub-based network: Hubs in a network create some interesting challenges for administrators. It is advisable not to link more than four 10baseT hubs or two 100baseT hubs together. Also, the collision domain, the number of ports connected by hubs without

a switch or router in between, should be kept as low as possible. Finally, the effective (half-duplex) bandwidth available on a shared collision domain is approximately 35% of the total bandwidth available.

- Too many access lists: Access lists slow down a router. While they are appropriate for voice networks, care must be taken not to apply them to unnecessary interfaces. Traffic should be modeled beforehand, and access lists applied only to the appropriate interface in the appropriate direction, not all interfaces in all directions.

Additional concerns when implementing VoIP include:

- Network Address Translation (NAT): Due to limitations in the H.323 VoIP standard, VoIP conversations rarely work across NAT boundaries. It is important to route voice streams around routers or firewalls running NAT or use a H.323 friendly NAT.
- Virtual Private Networks (VPN): VPNs present interesting challenges to VoIP implementations. First, the encryption used with VPNs adds significant latency to voice streams, adversely affecting the user experience. Second, VPNs generally run over the Internet. Because there is no control over QoS parameters for traffic crossing the Internet, voice quality may suffer due to excessive packet loss, delay, and jitter. For more information, please refer to Avaya's VPN white paper.

Recommended platforms

While many vendors offer VoIP products, this white paper deals only with the Avaya product line. One can substitute other vendors' products, but be sure that the products offer sufficient QoS features for success.

Switches

The following switches were designed with IP telephony in mind, and incorporate QoS features:

- Avaya Cajun P-120 family (access-layer: 24 Ports)
- Avaya Cajun P-130 family (access-layer: 24 – 48 Ports)
- Avaya Cajun P-330 family (access-layer: 24-640 Ports, stackable)
- Avaya Cajun P-550 family (access-distribution layer, up to 288 ports)
- Avaya Cajun P-880 family (distribution-core layer, up to 768 ports)

Routers

The following routers support QoS sufficient for VoIP:

- Avaya AP450 (branch office)
- Avaya AP1000 (central site)

Network Management

The following network management tools help administrators maintain a properly functioning VoIP network:

- Avaya Cajun Rules
- Avaya Cajun View

Appendix B

Frame Relay

Frame Relay remedies for known voice traffic fatalities

Ensure that the Committed Information Rate (CIR) is sufficient to support peak voice traffic or some threshold for voice traffic (as determined by the network administrator), and then prioritize the voice traffic. The underlying assumption here is that the network administrator has an expectation of peak voice traffic. You can usually determine this traffic level, because there are a limited number of IP phones that can make calls over a WAN link, or a statistically derived maximum number of calls expected to traverse a WAN link. By sizing the CIR to meet or exceed the peak voice traffic level, the network administrator can prioritize the voice traffic to ensure it is always delivered within the CIR. One method of doing this is to enable priority queuing at both ends of the frame relay link, so that voice traffic is always processed and delivered out the WAN link first.

The reason for maintaining voice traffic within the CIR is that typically, delivery of burst traffic is not guaranteed while delivery of CIR traffic is guaranteed. ILECs will all contract to meet a customer's traffic delivery requirements within the CIR. This is done through Service Level Agreements (SLA). The carriers usually will not contract to such an extent, or at all, for burst traffic. And indeed this is how frame relay is intended to work. As the name suggests, CIR is a committed and reliable rate. Burst is sort of a bonus to the customers when network conditions permit it without infringing on any user's CIR. For this reason, burst frames (frames marked Discard Eligible (DE)) are either queued until network congestion subsides, or the frames are discarded entirely. Although experience has shown that customers can achieve significant burst throughput, it is unreliable and unpredictable, especially for critical applications.

It is important to understand that ILECs convert the long-haul delivery of frame relay into ATM. That is, the frame relay PVC is converted to an ATM PVC for long-haul transport at the first frame relay switch after leaving the customer's premise. It is not converted back to frame relay until the last frame relay switch before entering the customer's premise. This has significance because ATM has built in Class of Service (CoS). A customer can contract with a carrier to convert the frame relay PVC into a Constant Bit Rate (CBR) ATM PVC. ATM CBR cells are delivered with lower latency and higher reliability.

As a final note, the network administrator should understand that under the best circumstances, frame relay is still inherently more susceptible to delay than ATM or TDM. Therefore, after following these suggestions one should still expect more delay over frame relay than would be present under ATM or TDM.

Appendix C

VoIP without using NAT

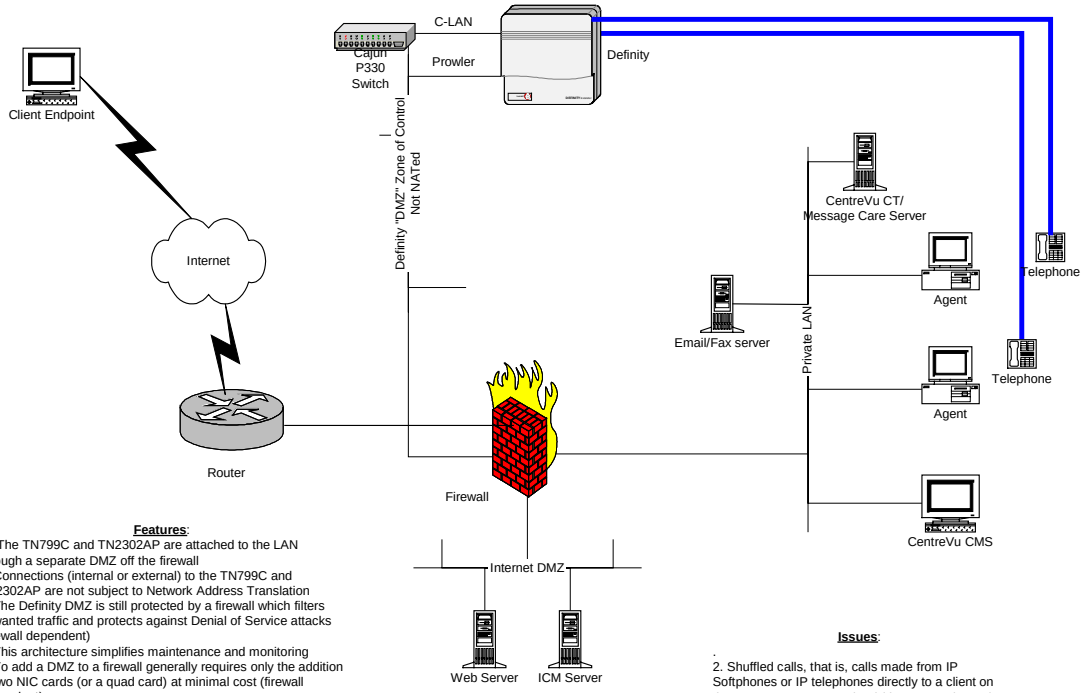

Features:

1. The TN799C and TN2302AP are attached to the LAN through a separate DMZ off the firewall
2. Connections (internal or external) to the TN799C and TN2302AP are not subject to Network Address Translation
3. The Definity DMZ is still protected by a firewall which filters unwanted traffic and protects against Denial of Service attacks (firewall dependent)
4. This architecture simplifies maintenance and monitoring
5. To add a DMZ to a firewall generally requires only the addition of two NIC cards (or a quad card) at minimal cost (firewall dependent)
6. TDM calls through the Definity or hairpin calls should process correctly
7. Firewalls can be load-balanced with third-party software or hardware for greater performance and reliability
8. This architecture represents the current industry "best practices"

.

Issues:

.

2. Shuffled calls, that is, calls made from IP Softphones or IP telephones directly to a client on the Internet are not permitted (due to NAT issues). This can be remedied by locating the IP telephones on the Definity DMZ

Appendix D

VoIP Tools

Shomiti Systems Explorer

This hardware-based tool measures delay, cell loss and jitter at wire speeds from 10Mbps to 1000 Mbps and provides a seven-layer decoding of captured frames. Because it has a dedicated processor for sensing traffic, results are more accurate than with software-based tools. It normally acts in passive mode by "sniffing" traffic. It can also be an active device by injecting packets into the network. Information is available at http://www.shomiti.com.

Ixia™ 100™ QoS Performance Tester (also the 400™ and 1600™)

This hardware-based tool is both a traffic generator and performance analyzer. It measures delay, jitter and cell loss from 10 to 1000 Mbps. It is scaleable to higher speeds as it can use OC3 through OC192 for generating traffic. The results are available for pre- and post-processing by the user and the software code is open for users to customize. Information is available at http://www.ixiacom.com.

NetIQ™ Chariot™

This software tool allows customized traffic generation controlled from a server between two PC endpoints. Traffic is created by selecting pre-made scripts or writing your own and represents data from the application level. Lower level (OSI layers 4, 3 and 2) traffic is also available to configure and send. Information is available at http://www.NetIQ.com.

Fluke® Enterprise LANmeter®

This all-purpose hardware instrument can be used as a traffic generator and diagnostic tool, or to check Category 3 and 5 cables, simulate an endpoint, etc. It will not test fiber (yet), but it is very portable and capable of troubleshooting a LAN. Results can be viewed from a web browser and an online database option is available. Information is available at http://www.fluke.com.

OPNET® IT DecisionGuru and Modeler

OPNET produces "Cadillac" software products that will discover network elements and model the behavior of a LAN. This predictive feature is a good way to test changes to the network before implementing the actual hardware. The accuracy of results to real world experience ranges from 80 to 95 percent, which is higher than most mathematical only models because each element performs like the physical unit it represents. The code is partially open and users can create new objects or modify existing ones. This is a good proactive tool for network analysis. Information is available at http://www.opnet.com.

Network Associates® Sniffer® tools

These industry-standard frame-capturing tools are very handy for examining and verifying content of OSI model layers 2, 3 and 4. They cannot measure latency or cell loss. They are portable and also analyze long-term network trends. Information is found at http://www.nai.com.