



Avaya IP Voice Quality Network Requirements

Issue 3.0

April 2004



Table of Contents

Section 1:	Executive Summary	1
Section 2:	Introduction	3
Section 3:	Defining Voice Quality	4
Section 4:	Prioritizing Voice Traffic	4
	4.1 Understanding CoS versus QoS.....	4
	4.2 Using Ports.....	5
	4.3 Using DSCP (or TOS).....	5
	4.4 Using IEEE 802.1 p/Q.....	5
	4.5 Using VLANs	6
Section 5:	Network Parameters	6
	5.1 Network Packet Delay.....	6
	5.2 Network Jitter	7
	5.3 Packet Loss	8
	5.4 Network Packet Mis-Order.....	9
	5.5 Transcoding	9
	5.6 Echo	9
	5.7 Silence Suppression and Voice Activity Detection.....	9
	5.8 Network Duplex.....	10
	5.9 Codec Selection	10
Section 6:	Bearer Bandwidth	11
	6.1 WAN Bandwidth Comparison	11
	6.2 LAN Bandwidth Comparison.....	12
Section 7:	Network Assessment	13
Section 8:	PC Considerations using Avaya IP Softphone	14
Section 9:	Bandwidth Requirements	15
	9.1 Bandwidth Requirements using IP SoftPhone (or IP Agent).....	15
Section 10:	Other Elements that Affect VoIP	16
	10.1 WAN Considerations	16
	10.2 Virtual Private Network	16
	10.3 Frame Relay	16
	10.4 Network Address Translation (NAT)	17
Section 11:	Avaya VoIP Products	17
	11.1 Avaya Communication Manager	17
	11.2 Avaya Communication Manager on an Avaya S8700 Media Server and an Avaya G600 Media Gateway (S8700/G600)	17
	11.3 Avaya Communication Manager on an Avaya S8700 Media Server and an Avaya MCC/SCC Media Gateway (S8700/MCC)	17
	11.4 Avaya S8300 Media Server/Avaya G700 Media Gateway	18
	11.5 Avaya S8300 Media Server/Avaya G350 Media Gateway	18
	11.6 Avaya IP600 Internet Protocol Communications Server.....	18
	11.7 Media Processor Circuit Pack	18
	11.8 Control LAN Circuit Pack	19
	11.9 IP SoftPhone	19
	11.10 IP Telephone.....	19
	11.11 IP Office Offerings	19
	11.12 Cajun Switches	20
	11.13 VPNet	20
Section 12:	VoIP Tools	20
	12.1 Network Tools	20
Section 13:	Resources	20
Appendix A:	21
Appendix B:	24
Appendix C:	26
Appendix D:	27



Section 1: Executive Summary

This document contains basic network requirements that are foundational for good voice quality when using Avaya IP products and solutions over a data network. No document can satisfy the detailed needs of every network, and therefore, this paper serves only as a starting point. The executive summary provides a short list of networking requirements, allowances and recommendations. Use this page as a checklist to determine if the network meets the minimum requirements for implementing Voice over Internet Protocol (VoIP) with acceptable quality. The rest of the document contains basic networking and telephony concepts for those who haven't been exposed to a converged implementation. It also explains why VoIP applications can yield poor results when data traffic on the same network doesn't seem to have problems.

Voice quality is always a subjective topic. Defining "good" voice quality varies with business needs, cultural differences, user expectations and hardware/software. The requirements below are based on the ITU-T, EIA/TIA guidelines and extensive testing at Avaya Labs. Note that while Avaya's requirements will meet or exceed most quality expectations, the final determination of acceptable voice quality lies with the user's definition of quality and the design, implementation and monitoring of the end to end data network.

Quality is not measured by one discrete value where a number like 8 is good and 9 is bad. There is a tradeoff between real-world limits and acceptable voice quality. Lower delay, jitter and packet loss values can produce the best voice quality, but may also come with a cost to upgrade the network infrastructure to get to the lower network values. Another real-world limit is the inherent WAN delay over a trunk linking the U.S. West coast to India. This link could easily add a fixed 150ms delay into the overall delay budget and is beyond the control of a business. Perfectly acceptable voice quality is attainable but will not be "toll" quality. Therefore, Avaya presents a tiered choice of values that make up the requirements.

Voice quality is made up of both objective and subjective factors. The objective factors in assessing VoIP quality are delay, jitter and packet loss. To maintain good and consistent levels of voice quality, Avaya suggests the following network requirements. Note that these suggestions hold true for both LAN only and LAN/WAN connectivity. All requirements values listed are measured between endpoints because this document assumes that IP telephony has not yet been implemented. All values therefore reflect the network's performance without endpoint consideration. This is why there is seemingly a discrepancy between the well-known ITU-T value for one-way delay and the values listed. The ITU-T values are end-to-end values; from the mouth of the transmitter to the ear of the receiver. The network requirements listed are meant for the network only—between endpoints—so that a business data system can be assessed and modified if need be for successful deployment of real time applications like voice. The requirement values are also useful for ongoing network monitoring by IT staff. Upward trends in delay, jitter or packet loss are a warning of potential voice quality problems.

Also, please note that "Business Communication Quality" is defined as slightly less than toll but far better than cell-phone quality. This is the tier where most businesses experience the best trade-off between voice quality and network infrastructure costs.



Requirements for objective factors:

- **Network delay:** One-way between endpoints should be:
 - 80ms (milliseconds) delay or less, which can—but may not—yield toll quality.
 - 80ms to 180ms delay for business communications quality. This is far better than cell phone quality and in fact is very well suited for the majority of businesses.
 - 180ms or higher, which may still be quite acceptable depending on user expectations, analog trunks used, codec type, etc.
- **Network jitter:** The measure of the variability of delay between endpoints should be:
 - An average of less than 20ms, or less than $\frac{1}{2}$ the packet payload, for toll quality. This value has some latitude depending on the type of service the jitter buffer has in relationship to other router buffers, packet size used, etc.
- **Network packet loss:** The maximum loss of packets (or frames) between endpoints should be:
 - 1% or less, which can yield toll quality depending on many factors.
 - 3% or less, which should give business communications quality. This quality is much better than cell phone quality.
 - More than 3%, which may be acceptable for voice but may interfere with signaling.

Avaya highly recommends consideration of the following list of Best Practices 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 are recommended to use IEEE 802.1p/Q. Routed networks should use DSCP (DiffServ Code Points). Mixed networks should use both as a best practice. Port priority can also be used to enhance DiffServ and IEEE 802.1p/Q. Even networks with plentiful bandwidth should implement CoS/QoS to protect voice communications from periods of unusual congestion such as from a computer virus.
- **Switched Network:** A fully switched LAN network is a network that allows full duplex and full endpoint bandwidth for every endpoint that exists on that 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. For a review and certification, contact an Avaya representative or authorized BusinessPartner.
- **VLANs:** Placing voice packets on a separate VLAN (subnet) from data packets is a generally accepted practice to reduce both broadcast and data traffic from contending for the same bandwidth as voice. Other benefits become available when using VLANs, but there may be a substantial cost for initial administration and maintenance.



Avaya recommends caution when using the following:

- **Network Address Translation (NAT):** Be very careful when using NAT. 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. It is unlikely that NAT can find and translate all of them. Avaya Communication Manager will work seamlessly with any static NAT application, even if that NAT is not H.323 aware.
- **Analog Dial-Up:** Be careful when using analog dial-up (bandwidth $\leq 56K$) 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.
- **Virtual Private Network (VPN):** Use 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.

Section 2: Introduction

Voice over Internet Protocol (VoIP) is the convergence of traditional voice onto an IP data network. VoIP provides better application integration by using a common protocol and can help lower costs by melding separate support staffs. 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 voice worlds. Historically, these worlds have used different networks, support organizations and philosophies. The protocols used have always separated the voice network from the data network. In addition, the characteristics of voice applications are very different from those of data applications.

Traditionally, voice calls have had their own dedicated bandwidth throughout the circuit switched network. This provided an environment where “five nines” of reliability became the standard. Interactive voice traffic is sensitive to delay and jitter but can tolerate some packet loss, problems that were rarely an issue with circuit switching.

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 centered on providing reliable data transmission over unreliable media, almost regardless of delay. Bandwidth in the data world is largely shared, so congestion and delay are often present and can cause problems for multimedia applications such as voice.

The factors that affect the quality of data transmission are different from those affecting the quality of voice transmission. For example, data is generally not affected by delay. Voice transmissions, on the other hand, are degraded by relatively small amounts of 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 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 configuration
- Scalability
- WAN protocols
- Encryption/Decryption

This document provides basic network guidelines that should be addressed to maintain 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.

Section 3: Defining Voice Quality

Because quality can mean different things to different organizations, defining voice quality is difficult. The values and network parameters of a small call center can be greatly different than those of a business in another culture or country. For this reason, Avaya presents tiered network requirements, because one number for delay, jitter or packet loss cannot satisfy every business in every culture. Ultimately, each business must decide if quality voice using VoIP requires the first tier of values or one of the other values.

Section 4: 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 or sufficient bandwidth must always be available. Avaya products for VoIP—Communication Manager and Avaya Cajun® 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. Avaya products are designed to work with most other popular switches and routers through open standards to provide end-to-end voice prioritization.

4.1 Understanding CoS versus QoS

Class of Service is simply a classification method. CoS does not determine a level of Quality of Service (QoS), but is the method used by 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 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 CBWFQ (Class Based Weighted Fair Queuing), RSVP (RESERVATION Protocol - RFC 2205), MPLS, (Multi Protocol Label Switching - RFC 1117 and others).

CoS, or tagging, is totally ineffective in the absence of QoS because it can only mark data. QoS relies on those tags or filters to give priority to data streams.

4.2 Using Ports

One prioritization scheme assigns priority based on the User Datagram Protocol (UDP) port numbers used by the voice packets. This scheme allows the use of network equipment to 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. By using Communication Manager, users can define a port range for voice priority. Routers and layer 3 data switches can then use these ports to distinguish priority traffic. This priority traffic can be voice packets (UDP), signaling packets (TCP) or both. This is an OSI model layer 4 solution and works on data coming to and from the specified ports or port range.

4.3 Using DSCP (or TOS)

The DSCP prioritization scheme redefines the existing 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 Communication Manager, IP Telephones, and other network elements such as routers and switches in the LAN and WAN. A DSCP of 46 (101110) is suggested for the expedited forwarding of voice and signaling packets. However, with Communication Manager, one can set any DSCP value as desired to work with a company's QoS scheme.

Please note that older routers may require a DSCP setting of 40 (101000), which is backward compatible to the original TOS byte definition of critical. But again, Avaya products and software allow users to set any of the 64 possible DSCP values to work with a voice quality policy. The TOS byte is an OSI model layer 3 solution and works on IP packets on the LAN and possibly the WAN depending on the service provider.

4.4 Using IEEE 802.1 p/Q

Yet another prioritization scheme is the IEEE 802.1Q standard, which uses four bytes to augment the layer-2 header. IEEE 802.1Q 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. With Communication Manager, users can add the 802.1Q bytes and set the priority bits as desired. Avaya suggests using a priority of six for both voice and signaling. The Avaya Cajun line of data switches can switch frames with or without these VLAN headers with no configuration time spent. IEEE 802.1p and IEEE 802.1Q are OSI layer-2 solutions and work on frames.



4.5 Using VLANs

VLANs provide limited security and create smaller broadcast domains through software by creating virtually separated subnets. Broadcasts are a natural occurrence in data networks from protocols used by PCs, servers, switches, routers, NOS, etc. Creating a separate VLAN for voice reduces the amount of broadcast traffic (and unicast traffic on a shared LAN), the telephone will receive. Separate VLANs result in more effective bandwidth utilization and reduces the processor burden on the IP telephones and PCs by freeing them from having to analyze irrelevant broadcast packets. VLANs, a layer-2 feature, are created in data switches. A voice VLAN can be manually applied to an IP telephone or given by a DHCP server. CoS tagging and QoS policies can be applied at OSI layer 2 by using VLANs. Separate voice and data VLANs are an option that makes sense for most enterprises and is highly recommended by Avaya. Note however that VLAN implementation and maintenance can be substantial, and again, is an option even as a best practice. Proper VLAN implementation is not trivial and Avaya can help with planning and implementation through its Converged Services Group.

Section 5: Network Parameters

There are a number of network parameters that affect voice quality. Some of the more important are described in this section. IP telephony quality can be engineered to several different levels to accommodate differing business needs. A small company may choose to implement IP telephony with very good sound instead of buying newer networking equipment to support excellent voice sound. A large call center company may want excellent voice sound as part of its corporate strategy. Avaya therefore presents options in network requirements to allow an enterprise to choose the “quality” level that best suits their specific business needs.

5.1 Network Packet Delay

Packet delay is the length of time it takes a packet to traverse the network. 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). Router delay depends not only on hardware, but also on configurations such as access lists, queuing methods, and transmission modes. Delay (latency) can have a noticeable affect but can be controlled somewhat in a private environment (LAN/WAN) because the company or enterprise manages the network infrastructure or SLA. When using the public network, there are inherent delays that cannot be controlled.

The following suggestions are guidelines for one-way network delay. Again, there is a trade-off between voice quality and the technical and monetary constraints that businesses confront daily. The Tiered measurements from the E-Model R-Values are as follows:

Best Quality is rated toll quality or above: R94 (very satisfied)

Business Communication Quality: R80 - R92 (satisfied)

Possibly acceptable: R70 - R80 (Some unsatisfied)



Network delay: LAN/WAN measurements, not including IP phones, between endpoints:

- 80ms (milliseconds) delay or less can, but may not, yield toll quality.
- 80ms to 180ms delay can give business communication quality. This is much better than cell phone quality and in fact is very well suited for the majority of businesses.
- Delays exceeding 180ms may still be quite acceptable depending on user expectations, analog trunks used, codec type, etc.

The ITU-T has recommended 150ms one-way delay (including endpoints) as the limit for “excellent” voice quality. This value is largely misinterpreted as the only range to calculate a network delay budget for IP telephones. A network delay budget of 230ms proved almost imperceptible in lab experiments at Avaya. One-way delays in excess of 250ms can cause the well-known problem of “talk-over,” when each person starts to talk because the delay prevents them from realizing that the other person has already started talking. Certainly long WAN transports must be considered as a major contributor to the network delay budget, as one major WAN service provider averaged 75ms delay from Los Angeles to New York. Los Angeles to Paris was found to be about 145ms. Some WAN service providers can lower delay in their network if it is negotiated and recorded as part of the companies Service Level Agreement (SLA). Even so, staying within 150ms (end to end) may not be possible.

Finally, end-to-end delay over 400ms can cause port network instability. A network assessment is highly recommended to measure latency (and other factors) and make recommendations to solve any latency issues before implementing a VoIP solution.

5.2 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). Jitter is thought of 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 (assuming an existing 20ms packet size). 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 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 data-switched network than on a hub-based network, there will be less jitter on the switched network.



The Avaya G600 and G700 media gateways, Avaya IP600 Internet Protocol Communications Server, Avaya IP SoftPhone software and Avaya 4600 series 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 and bring control to the network.

5.3 Packet Loss

Network packet loss occurs when packets are sent, but not received at the final destination due to some network problem. Qualifying problems caused by occasional packet loss are difficult to detect because each codec has its own packet loss concealment method. Therefore, it is possible that voice quality would be better using a compression codec (G.729A) compared to a full bandwidth G.711 codec. Several factors 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 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.
- Packet loss may be more tolerable for one codec over another.
- Even small amounts of packet loss can greatly affect a TTY (TDD) device's ability to work properly.
- Packet loss for TCP signaling traffic increases substantially when more than 3% loss is due to retransmissions.

Network packet loss: The maximum loss of packets (or frames) between endpoints should be:

- 1% or less, which can yield toll quality depending on many factors.
- 3% or less, which should give business communications quality. This quality is much better than cell phone quality.
- More than 3%, which may be acceptable for voice but may interfere with signaling.

Like delay, Avaya allows enterprises a tiered approach of packet loss to balance new network costs and limitations with business directives. Tools such as the Agilent (HP) Internet Advisor, Finisar's Surveyor Explorer, Radcom's Prism, and NAI's Sniffer measure packet loss.

Remember that too much delay or packet mis-order can cause dropped packets, and it may appear that the network is losing packets when in fact they have been discarded intentionally.



5.4 Network Packet Mis-Order

Network packet mis-order is, for VoIP, 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 and buffers are small. 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.

5.5 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 transcoding (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 Communication Manager 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. Shuffling and Hairpinning also reduce transcoding.

5.6 Echo

The two main types of echo are acoustic and impedance although the sources of echo can be many. Echo will result when a VoIP call leaves the LAN through a poorly 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 Time Division Multiplexing (TDM) 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 the 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 G600, G650, G700, Avaya IP SoftPhone software and Avaya 4600 series IP telephones all incorporate echo cancellers designed for VoIP to improve voice quality.

5.7 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 Avaya 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. The Avaya 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



comfort 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 Wide Area Network (WAN). Avaya Communication Manager, Avaya 4600 series IP Telephone and Avaya IP SoftPhone products all can employ silence suppression to preserve vital bandwidth.

5.8 Network Duplex

The ideal LAN network for transporting VoIP traffic is a network that is fully 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 and should be avoided.

Ethernet connections from Avaya hardware default to auto-negotiation for speed and duplex to work with the network endpoints right away. Avaya recommends, however, that connections become set values for static links because of known problems with the way auto-negotiation was designed. The CLAN, Media Processor, IPSI etc. connections should be set to 100Mbps and full duplex both in Communication Manager and at the Ethernet data switch to which it terminates. Note that IP Telephones use auto-negotiation, and specific older versions of circuit cards may require different speed and/or duplex settings.

5.9 Codec Selection

Depending upon the bandwidth availability and acceptable 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 comparisons 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 Mean Opinion Score (MOS) of four or above. MOS scoring is a long-standing subjective method of measuring voice quality.

Table 1 – Comparison of Speech Coding Standards¹

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

Generally, G.711 is used within LANs because bandwidth is abundant and inexpensive whereas G.729A is used across WAN links because of the bandwidth savings and good performing voice quality.

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



Section 6: Bearer Bandwidth

There are bandwidth requirements for both bearer (voice) and signaling. Table 2 shows the differences of various codecs with varying voice payload sizes for bearer traffic only.

6.1 WAN Bandwidth Comparison

Examples of Layer 2 protocols include Ethernet, Frame Relay, PPP, ATM and others. Note that eight bytes was used for the layer-2 calculation contribution because the most used data protocols, Frame Relay and PPP, can use eight bytes or less.

WAN bandwidth is expensive compared to LAN bandwidth driving the need to use compression in both voice payloads and the headers of the TCP/IP protocol stack. Table 2 shows bearer bandwidth required using OSI layers 2 through 5 and several codecs for WAN connections.

Table 2 – WAN Bandwidth Comparison

WAN Bandwidth (using Frame Relay or PPP L2 Protocol) in Kbps							
Codec Type		G.711m and G.711A	G.711m/A with cRTP	G.729 and G.729A	G.729A with cRTP	G.723 6.3kbps	G.723 5.3kbps
Voice	10 ms	102.4	~73.6	46.4	~17.6	N/A	N/A
	20 ms	83.2	~73.6	27.2	~12.8	N/A	N/A
Payload	30 ms	76.8	~67.2	20.8	~11.2	19.1	18.1
	40 ms	73.6	~66.4	17.6	~10.4	N/A	N/A
Size	50 ms	71.7	~65.9	15.7	~9.9	N/A	N/A
	60 ms	70.4	~65.6	14.4	~9.6	12.7	11.7

Observations:

- G.711 is inappropriate for most WAN connections because it uses more bandwidth than traditional TDM channels. The G.711 column values then become academic.
- G.723 gives very little bandwidth reduction compared to G.729A with cRTP but the voice quality is much less than the G.729A. Therefore the G.723 column values are rarely used.
- Most point-to-point links use G.729A or G.729A with cRTP as a good compromise between voice quality rendered and bandwidth required.
- The approximate bandwidth of the G.729 with cRTP headers varies slightly because even though headers are compressed to 12 bytes, a full header of 48 bytes is periodically sent. The period of full header transmission is usually somewhat configurable by the user.
- The most commonly used codecs for WAN implementations are G.729 and G729A, as well as G.729A with cRTP.
- Bandwidth required when using the Avaya default voice payload size is 27.2 for G.729 and G.729A, and approximately 12.8 for G.729A with cRTP. Use these values to begin your WAN bandwidth calculations.



6.2 LAN Bandwidth Comparison

Ethernet is the reigning layer-2 protocol in most LANs worldwide. Yet there are several seemingly conflicting bandwidth values stated for VoIP in the LAN. There are four different types of Ethernet frames, but let us confine this discussion to the most popular one—Ethernet Version 2 (EV2). The EV2 frame looks like this:

Destination Address (6 octets)	Source Address (6 octets)	Type (2)	Data (46 to 1500 octets)	Frame Check Sequence (4 octets)
---	--------------------------------------	---------------------	-------------------------------------	--

Many different values can be calculated using this frame alone depending on voice payload size and the inclusion of preamble or the 802.1Q shim. The next table displays LAN bandwidth using a G.711 codec. G.711 is most often used in the LAN because it gives the best voice quality and bandwidth is relatively inexpensive and plentiful compared to the WAN.

Table 3 – LAN Bandwidth Comparison

LAN Bandwidth using G.711 codec in kbps				
Ethernet Type		EV2 with trailer but no preamble	EV2 with trailer and preamble	EV2 with trailer and preamble and 802.1Q
G.711 Voice Payload Size	10 ms	110.4	116.8	120.0
	20 ms	87.2	90.4	92.0
	30 ms	79.5	81.6	82.7
	40 ms	75.6	77.2	78.0
	50 ms	73.3	74.6	75.2
	60 ms	71.7	72.8	73.3

Notes:

- The preamble consists of 8 bytes of alternating 1s and 0s with the last two bits both being 1s. This is the synchronization method used to alert all nodes on the asynchronous network or subnet that a frame is coming. Since the preamble takes time to impress bits onto the network, it should be counted in the overall bandwidth calculation.
- The trailer is 4 bytes called the Frame Check Sequence (FCS) containing the results of a cyclical redundancy check (CRC). The CRC is a mathematically derived value used to verify that the frame was received uncorrupted. Again, these four bytes take time to transmit on the network and should also be included in the bandwidth calculations.
- Switch throughput calculations often do not include the preamble or the trailer. The values listed for EV2 with trailer but not preamble don't really apply to the business world. Data switch manufacturers use these values for switch backplane speed calculations.
- The values for EV2 with trailer and preamble are useful if VLAN number are not included in the transmission. Before VoIP is implemented on a network, these values are used.



- The values for EV2 with trailer, preamble and 802.1Q, shown in light gray in table 3, are accurate when using separate voice and data VLANs as Avaya recommends. The bandwidth for the default payload size set in Communication Manager is 92.0.

Section 7: Network Assessment

Avaya Network Assessment for IP Telephony Solutions is designed to verify that an enterprise data network is capable of supporting Voice over IP applications before installation of any Avaya application. Network Assessment for IP Telephony Solutions is a flexible process, allowing an enterprise to provide the required network assessment data, or they can provide the data to Avaya through their network vendor.

Network Assessment for IP Telephony Solutions consist of two distinct phases.

- The first phase—Customer Infrastructure Readiness Survey (CIRS)—is focused on providing a high-level evaluation of an enterprise’s LAN/WAN infrastructure. It is used to determine if significant issues exist that must be dealt with prior to deploying the newly proposed IP solution.
- Network Analysis/Network Optimization (NANO) is typically the second phase in the process. The NANO takes information gathered from the CIRS, performs problem diagnosis and provides functional requirements for the network to implement an IP telephony solution. A NANO is required when the CIRS indicates that the network as it is configured will not support the proposed IP solution at the desired performance levels.

Working remotely (Phase 1 CIRS), and using a combination of interactive questionnaires and innovative software tools, Avaya Network Consulting Services network engineers will:

1. Identify the equipment in an enterprise’s network, as well as physical and network layer information, device connections, network topology and device configurations through a Site Configuration Survey and network Topology maps.
2. Test the enterprise’s current network infrastructure to discover any throughput and response time issues in multi-protocol networks using VitalAgent software.
3. Baseline existing throughput performance statistics for critical LAN and WAN circuits using CIRS Network Monitor software.
4. Determine 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 of pre-implementation network details to use for comparison with post-implementation network details.

Working with representatives of an enterprise on site (Phase 2 NANO), and using a combination of software and hardware tools, Avaya Network Consulting Services network engineers will perform:

1. Discovery of the customer’s network and document findings in a NANO Report delivered to the enterprise



2. Accurate Network Topology
3. Measurements of actual usability performance levels, throughput performance of the LAN, and server utilization
4. Results of traffic simulation on the network at projected volumes
5. Definition of problem areas, causes, and functional requirement recommendations to be implemented in the network design

Enterprises that do not avail themselves of this offer assume responsibility for all network-related problems with the IP voice installation. Also, Avaya personnel may be required to charge a higher Time and Materials (T&M) rate if assistance is needed, since troubleshooting will be more difficult without the assessment data.

Section 8: PC Considerations using Avaya IP Softphone

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

- Transducer Quality: The selection of speaker and microphone or headset has an impact on the sound reproduction.
- 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.
- 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.
- 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 play back 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.




Section 9: 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 enterprise. Enterprise-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.

9.1 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.

Figure 1



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. Presently, testing is underway to determine signaling bandwidth requirements.



Section 10: Other Elements that Affect VoIP

10.1 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 Compressed Real Time Protocol (CRTP) can and should be used between WAN links. Hardware based CRTP is effective with very minimal delays, but software CRTP can add significant delay.

10.2 Virtual Private Network

There are many definitions for Virtual Private Networks. 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 one milli-second to one 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. Users may be able to negotiate a service-level agreement with the VPN provider to guarantee an acceptable level of service. Before implementing VoIP with a VPN, users should test their VPN network to make sure it meets the requirements specified in the Executive Summary.

10.3 Frame Relay

Voice transported over frame relay can be 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, Avaya offers remedies to protect voice traffic from the susceptibilities of frame relay in Appendix B.



10.4 Network Address Translation (NAT)

VoIP does not work well with networks that use NAT 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. Most NAT implementations change only the address in the IP header resulting in a mismatch that prohibits the control of calls. Avaya suggests using a firewall to guard against intruders, but the 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. It is expected that Avaya products will soon work seamlessly with any static NAT application even if it is not H.323 aware.

Section 11: Avaya VoIP Products

11.1 Avaya Communication Manager

Avaya Communication Manager—evolved from DEFINITY® software—delivers Enterprise Class IP Solutions. Avaya Communication Manager runs on open platforms and standards-based commercial operating systems. This flexible, feature-rich software opens new opportunities to eliminate inefficiency and reliably integrate rich voice features and functions—proven in Avaya DEFINITY® Enterprise Communications Servers—into an organization’s Internet Protocol (IP) wide area and local area networks.

11.2 Avaya Communication Manager on an Avaya S8700 Media Server and an Avaya G600 Media Gateway (S8700/G600)

This configuration runs Avaya Communication Manager on the rack-mounted S8700 Media Server pair. It is configured with the rack-mounted Avaya G600 Media Gateway module yielding a pure rack-mount solution. All control signals are run over an enterprise’s data network, facilitating the distribution of media gateways. It can be configured as a pure IP PBX. It also supports the rich set of Avaya interface cards to operate with a wide variety of telephony facilities and endpoints. This offer supports 10,000 endpoints up to 30,000 busy-hour calls.

11.3 Avaya Communication Manager on an Avaya S8700 Media Server and an Avaya MCC/SCC Media Gateway (S8700/MCC)

This configuration runs Avaya Communication Manager on the rack-mounted S8700 Media Server pair. With the S8700 Media Server MCC/SCC configuration, the control and bearer networks are separate and the processor (control network) resides in an external 19” data rack; however, the bearer connectivity remains in the port networks (Media Gateways). Standard, High and Critical reliability configurations are all supported. This allows an enterprise to gently migrate an existing Avaya infrastructure to an IP environment with up to three times greater processing power than with the embedded processor model. This offer may support up to 36,000 stations, 64 port networks and 300,000 busy hour calls.



11.4 Avaya S8300 Media Server/Avaya G700 Media Gateway

This configuration is designed to meet the mission-critical communication needs of medium offices, branch offices, and global multi-site networked enterprises with 40-450 stations. With the full capabilities of Avaya Communication Manager and an internal Layer 2 Ethernet switch, it can be used as a total converged standalone solution for a medium business or as a remote edge device for a larger enterprise network. This configuration leverages the distributed architecture benefits of the media gateway concept and provides a survivable, standards-based, IP communications infrastructure without compromising applications, reliability, and multiservice networking. The solution is scalable, modular and supports stackable, hot swappable, redundant architectures.

11.5 Avaya S8300 Media Server/Avaya G350 Media Gateway

This gateway supports the mission-critical communication needs of small offices, branch offices, and global multi-site networked businesses with 8-32 stations or more. With the full capabilities of Avaya Communication Manager and an internal Layer 2 Ethernet switch, and a Layer 3 Router, it can be used as a total converged standalone solution for a small business or as a remote edge device for a larger enterprise network. The solution is cost effective by integrating all the data and voice needs into one small 3-U high data rack-mounted unit.

11.6 Avaya IP600 Internet Protocol Communications Server

This rack-mounted server is an all IP PBX for small and mid-sized businesses. It uses the same R10 software load as its larger counterpart, the DEFINITY® Enterprise Communications Server. It also works with all the other endpoints that Communication Manager supports. One IP600, as with any Communication Manager product, can support up to 256 R300 units.

11.7 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 or greater. If auto-negotiation is not used, 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 calls 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 Communication Manager 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 Communication Manager 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 Communication Manager.



11.8 Control LAN Circuit Pack

The TN799 Control LAN circuit pack (C-LAN) controls signaling and call setup. This circuit pack controls all IP call establishment/release and shuffled/hairpinned calls so that the Communication Manager 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.

11.9 IP SoftPhone

The Avaya IP SoftPhone is a client-based telephony application that provides excellent 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.

11.10 IP Telephone

The Avaya 46XX IP Telephone looks and feels just like a circuit-switched set because it supports most of the features of Avaya digital sets. This family of phones supports IEEE 802.1Q/p, DiffServ and a separate VLAN for voice traffic. It can withstand over 1,000 broadcasts per second making it an excellent phone in resisting broadcast storms. It also has a full-duplex speakerphone. It also supports traffic at 10 or 100 Mbps, supports silence suppression and can use Dynamic Host Configuration Protocol (DHCP) for easy setup. The 4630 IP telephone features a color touch screen, is web and LDAP enabled.

11.11 IP Office Offerings

Avaya INDeX (Intergrated Network Digital exchange) is well-established in EMEA and APAC markets as a modular business communications platform for medium-sized businesses with support for up to 120 VoIP DSPs. Scaling up to 1,088 ports, INDeX combines powerful traditional telephony capabilities with advanced productivity applications, including advanced contact center capabilities for up to 400 agents. Use of the INDeX IP Networking Cassette (IPNC) module enables organizations to support desk-to-desk connectivity for IP softphones and hardphones as well as VoIP trunking between sites.

Avaya IP Office is a global family of hybrid IP-PBX processors and powerful productivity applications for small and medium sized businesses with support for up to 60 VoIP DSPs. Designed for standalone applications up to 256 extensions and networked environments with up to 500 virtual PBX users, IP Office can also be used in large branch office deployments, typically up to 50 users per site with relatively low intensity networking and collaboration requirements. Further to this, Avaya expects to introduce a software option for IP Office units, that enables them to operate in slaved/sub-tended mode using Avaya Communication Manager. This will allow enterprises to transparently extend MultiVantage™ applications to smaller branch offices (2-24 extensions).

Avaya Small Office not only supports traditional extension and trunk types, it is also enabled for VoIP as standard, enabling the support of both the latest generation IP phones and VoIP trunks connecting other locations as needed. VPN and secure remote access options make this ideal for use with the latest broadband IP services delivered by DSL or cable modem. Small Office incorporates a WiFi (802.11b) access point, allowing the integration of roaming users with productivity applications such as Avaya Phone Manager Pro and wireless VoIP devices.



11.12 Cajun Switches

The Avaya 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 mapping DSCP support to IEEE 802.1 p/Q priority values.

11.13 VPNet

Avaya 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.

Section 12: VoIP Tools

12.1 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.

Section 13: Resources

- Network Requirements and Configuration Guidelines for Communication Manager on S87000/G600— Issue 1.1: http://www1.avaya.com/enterprise/resourcelibrary/applicationnotes/eclips_networking.html
- Avaya IP Telephony Implementation Guide—Communication Manager: <http://www1.avaya.com/enterprise/applicationnotes/av-iptel-imp-gd-acm-0503.pdf>
- Finisar Explorer: <http://www.finisar.com/home/>
- Empirix Hammer: <http://www.empirix.com>
- NetIQ Chariot: <http://www.NetIQ.com>



- Fluke Enterprise LANmeter: <http://www.fluke.com>
- OPNET IT DecisionGuru and Modeler: <http://www.opnet.com>
- Network Associates Sniffer Tools: <http://www.nai.com>

Appendix A: Network Design Recommendations

In the early days of Local Area 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 most users will not tolerate poor voice quality, administrators should implement a sound network before beginning VoIP pilots or deployments.

Best Practices

Industry best practices dictate that a network be designed with reliability/redundancy, scalability, manageability and bandwidth in mind. Additionally, voice mandates require delay, jitter, loss and duplex be considered when designing a network. 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 should 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 should 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. 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 should 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 committed information rate (CIR) 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 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, Visibility, 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 can 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 1 Gbps) 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 can quickly degrade to an unacceptable level.



- Multiple subnets on a VLAN: A network of this type can have issues with broadcasts, multicasts, and routing protocol updates. It should be avoided. It can 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: 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: 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.

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 and open standards-based QoS features for success.

Switches

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


- Avaya P130 family of switches (access layer: 24-48 Ports)
- Avaya P330 family of switches (access layer: 24-640 Ports, stackable)
- Avaya P580 switches (access distribution layer: up to 288 ports)
- Avaya P882 switches (distribution core layer: up to 768 ports)



Appendix B

Frame relay service is composed of three elements: the physical access circuit, the frame relay port, and the virtual circuit. The physical access circuit is typically a T1 or fractional T1 and is provided by the local exchange carrier (LEC) between the enterprise premise and the nearest central office (CO). The frame relay port is the physical access into the frame relay network—a port on the frame relay mux itself. The access circuit rate and the frame relay port rate must match. The virtual circuit is a logical connection between frame relay ports that can be provided by the LEC for intralata frame relay, or by the inter-exchange carrier (IXC) for interlata frame relay. The most common virtual circuit is a permanent virtual circuit (PVC), which is associated with a committed information rate. The PVC is identified at each end by a separate data-link connection identifier (DLCI).

Figure 2



This hypothetical implementation shows the Dallas corporate office (CO) connected to three branch offices in a common star topology (or hub and spoke). Each office connects to a LEC's CO via a fractional T1 circuit, which terminates onto a frame relay port at the CO and onto a frame relay capable router at the customer premise. The port rates and the access circuit rates match. PVCs are provisioned within the frame relay network between Dallas and each branch office. The CIR of each PVC is sized so that it is half the respective port rate, which is a common implementation. Each branch office is guaranteed its respective CIR, but it is also allowed to burst up to the port rate without any guarantees. The port rate at Dallas is not quite double the aggregate CIR, but it does not need to be because the expectation is that not all three branch offices will burst up to the maximum at the same time.

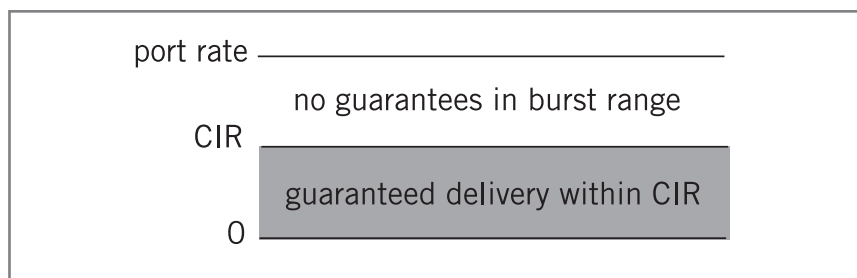
In an implementation like this the service is probably negotiated through a single vendor. But it is likely that Dallas and Houston are serviced by the same LEC and that the frame relay is intralata, even if it was negotiated through an IXC (such as AT&T or WorldCom or Sprint). The service between Dallas and the other two branch offices; however, is most likely interlata.



Issue and Alternatives

The obstacle in running VoIP over frame relay involves the treatment of traffic within the CIR and outside of CIR, commonly termed the “burst range.”

Figure 3



As Figure 3 illustrates, traffic up to the CIR is guaranteed, whereas traffic beyond the CIR typically is not. This is how frame relay is intended to work. CIR is a committed and reliable rate, whereas burst is a bonus when network conditions permit it without infringing upon any user’s CIR. For this reason, burst frames are marked Discard Eligible (DE) and are queued or discarded when network congestion exists. Although experience has shown that enterprises can achieve significant burst throughput, it is unreliable and unpredictable and not suitable for real-time applications like VoIP.

Therefore, the objective is to prevent voice traffic from entering the burst range and being marked DE. One way to accomplish this is to prohibit bursting by shaping the traffic to the CIR and setting the excess burst size (Be – determines the burst range) to zero. However, this also prevents data traffic from using the burst range as well. Another possible alternative is to size the CIR above the peak voice traffic level, and then prioritize the voice traffic so that it is always delivered first. The underlying assumption here is that the network administrator has an expectation of peak voice traffic. By sizing the CIR to meet or exceed the peak voice traffic, and then applying priority queuing on the interface so that VoIP is serviced first, we can intuitively determine that voice traffic will not enter the burst range.

The problem with the latter method, however, is that the actual queuing mechanism is not always intuitive. Even though the aggregate voice traffic throughput cannot exceed the CIR, it is possible that a voice packet could be sent in the burst range. The technical workings of this are beyond the scope of this document. But simply stated, it is possible that a voice packet would arrive right after many data packets have already been transmitted in the CIR range, such that the voice packet ends up in the burst range when the router processes it. However, the latter method is certainly worth trying.

Additional Frame Relay Information

One good piece of knowledge is that most IXCs convert the long-haul delivery of frame relay into ATM. That is, the frame relay PVC is converted to an ATM PVC at the first frame relay switch after leaving the enterprise premise. It is not converted back to frame relay until the last frame relay switch before entering the enterprise premise. This has significance because ATM has built in Class of Service (CoS). An enterprise

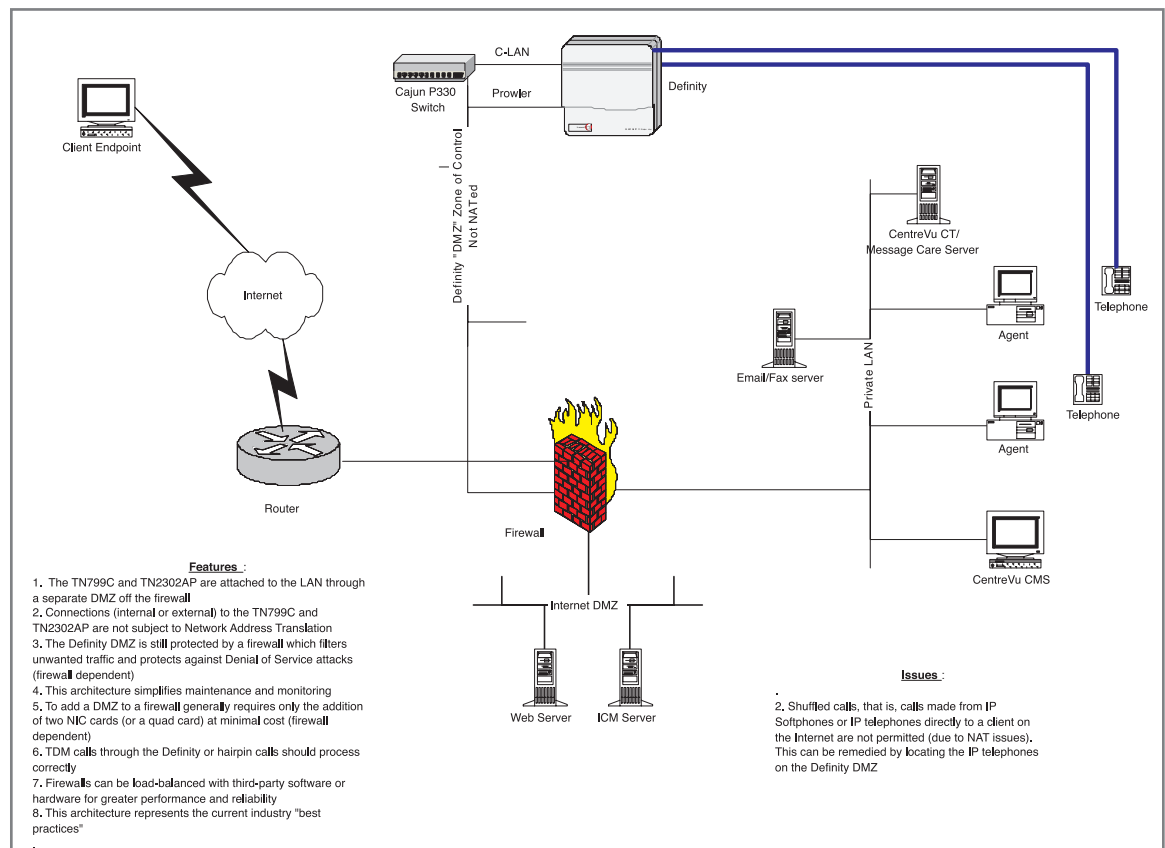


can contract with a carrier to convert the frame relay PVC into a Constant Bit Rate (CBR) ATM PVC. ATM CBR cells are delivered with lower latency and higher reliability.

As a final note, the reader should understand that under the best circumstances, frame relay is still inherently more susceptible to delay than ATM or TDM. Therefore, after applying the best possible queuing mechanism, one should still expect more delay over frame relay than would be present over ATM or TDM.

Appendix C: VoIP without using NAT

Figure 4



Note that Avaya will soon have a patented method of using VoIP with any NAT device even if the NAT device is not H.323 aware.



Appendix D: VoIP Tools

Finisar Explorer

This hardware-based tool measures delay, cell loss and jitter at wire speeds from 10 Mbps 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.

Empirix Hammer

The Hammer VQTS is a software-based solution for measuring the voice quality of next generation gateways and networks. The Hammer VQTS provides extensive Quality of Experience (QoE) metrics including: voice quality scoring algorithms (PESQ (ITU P.862), PAMS, PSQM (ITU P.861), PSQM +, and MOS correlation), Voice Activity Detection measurement (VAD), echo cancellation measurement, & speech latencies.

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.

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.

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.

These industry-standard frame-capturing tools are very handy for examining and verifying content of OSI model layers 2, 3, 4 and higher. They are portable and also analyze long-term network trends. LAN and WAN interfaces are available.

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.

About Avaya

Avaya enables businesses to achieve superior results by designing, building and managing their communications networks. Over one million businesses worldwide, including more than 90 percent of the FORTUNE 500®, rely on Avaya solutions and services to enhance value, improve productivity and gain competitive advantage.

Focused on enterprises large to small, Avaya is a world leader in secure and reliable IP telephony systems, communications software applications and full life-cycle services. Driving the convergence of voice and data communications with business applications – and distinguished by comprehensive worldwide services – Avaya helps customers leverage existing and new networks to unlock value and enhance business performance.



IP Telephony

Contact Centers

Unified Communication

Services

© 2004 Avaya Inc.

All Rights Reserved. Avaya and the Avaya Logo are trademarks of Avaya Inc. and may be registered in certain jurisdictions. All trademarks identified by the ®, SM or TM are registered trademarks, service marks or trademarks, respectively, of Avaya Inc. All other trademarks are the property of their respective owners. Printed in the U.S.A.

04/04 • EF-LB1500-02

