Voice/Fax Over IP:

Internet, Intranet, and Extranet

Technology Overview

Voice/Fax Over IP:

Internet, Intranet, and Extranet

MICOM Communications Corp.,
A Nortel (Northern Telecom) Company

4100 Los Angeles Ave.
Simi Valley, CA 93063-3397
(800) 642-6687
www.micom.com

Contents

INTRODUCTION

Telecommunications, historically involved only with analog voice and fax connectivity, is increasingly overlapping with data communications/ data networking, historically involved only with digital data connectivity. This merging of the two has occurred because of a special type of microprocessor, the DSP (Digital Signal Processor). DSPs have special, very powerful computer instructions that allow voice, video, fax and other analog signals to be processed into a variety of digital formats. The “bell heads” of telecommunications and “net heads” of data communications have found common ground through the DSP.

With the right software, the DSP can convert analog voice and fax into digital data for transport over data networks. DSPs have fallen in price by over a hundred-fold, allowing the development of a new type of hybrid networking — voice/data integrated networks. In most cases, voice and fax traffic, often from a PBX, is added to a data network using some type of DSP-based voice/data integrator (see Figure 1). A second voice/data integrator is placed elsewhere on the data network, and voice and fax communications can take place over the data network between the two PBXs.

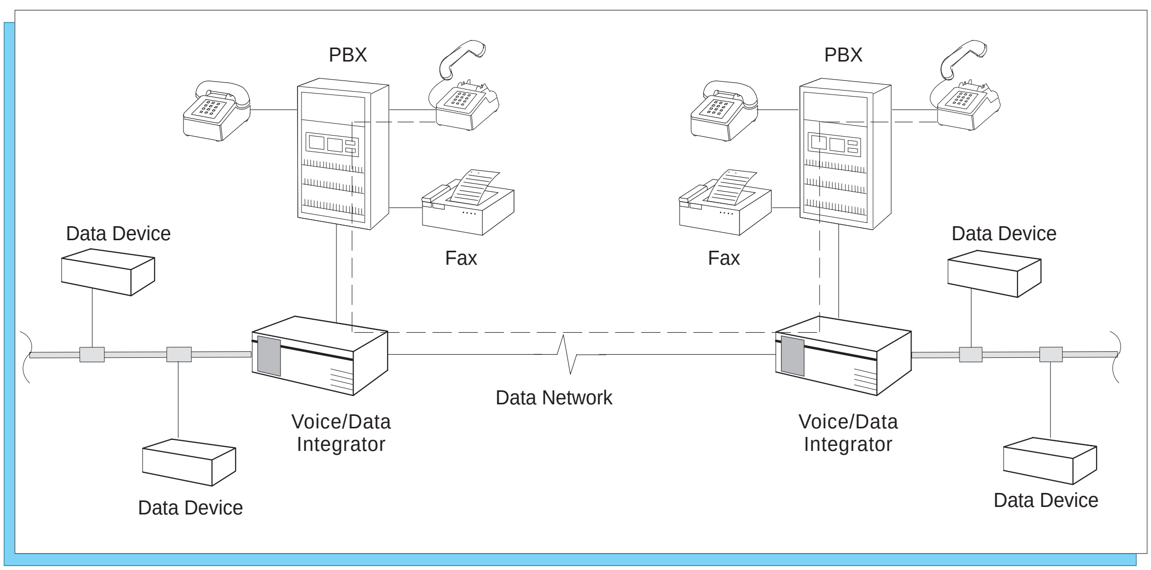

Figure 1 - Voice/Data Integration

Data networks are designed to carry data from one place to another. The data network can be local (LAN, Local Area Networking) or widely dispersed (WAN, Wide Area Networking). Voice/data integration has its strongest application over WANs, because long distance voice/fax toll rates are much higher than local toll rates.

With voice/data integration, voice and fax are converted into data and placed on a WAN for transport to a remote location. The voice/fax traffic “rides for free” over the data network, which is already paid for and usually has enough extra carrying capacity — “bandwidth” — to accommodate voice/fax traffic at no extra charge. The long distance voice/fax call is essentially toll-free.

Key Network Characteristics

Among the many different types of standardized data networking technologies — X.25, SNA, Frame Relay, ATM (Asynchronous Transfer Mode), TCP/IP, Novell IPX, and TDM (Time Division Multiplexing) — only Frame Relay, ATM, TCP/IP and TDM, along with proprietary technologies, do well carrying voice. The key characteristics a data network must have to carry voice well are low delay, predictable delivery of the voice information, a means of prioritizing the voice traffic ahead of data, and high enough efficiency to carry the extra voice traffic. (See Table 1.)

Table 1. Key Data Network Characteristics

Data Network Technology	Delay	Predictability	Priority Scheme	Efficiency
X.25	High	Poor	No	High
SNA	Moderate	Good	No	Moderate
Frame Relay	Low	Good	Coming	High
ATM	Low	Good	Yes	High
TCP/IP	Low	Good	Yes	Moderate-High
Novell/IPX	Variable	Fair	No	High
TDM	Low	Good	Yes	Low
Proprietary	Low	Good	Yes	High

The last column in the chart, Efficiency, although only indirectly related to voice transport, is so important that it is usually considered when discussing voice/fax integration. Without high efficiency, extra bandwidth must be added to carry the voice. This extra bandwidth increases costs, resulting in little or no net cost savings.

Because of its low efficiency, TDM is disqualified for voice/data integration in many scenarios, especially new applications. TDM is inefficient because it sends data at a constant rate, even when the sending device has no actual data to send. In this case, TDM sends “fill” data to keep the constant data stream full, causing significant inefficiencies in most applications.

The remaining qualifying technologies — Frame Relay, ATM, TCP/IP and Proprietary — achieve their high efficiency ratings by collecting the data to be transmitted into packets (like envelopes) and only sending them when they are reasonably full. By not sending empty or nearly envelopes, waste is minimized.

Micom, a division of Nortel (Northern Telecom), has been active in supplying voice/data integration products for over ten years. It is a leader in the Voice-over-Frame Relay (VoFR) market and the proprietary integration market with its very efficient ™*MicroBand* ATM cell relay technology. Starting in 1996, Micom entered the voice-over-TCP/IP (VoIP) market, and quickly achieved the leadership position with 56% unit share of the VoIP market.

Micom’s expertise in VoFR and *MicroBand* ATM contributed to its early success in the VoIP market. From 1991 through 1997, Micom shipped over 120,000 voice/data integration products to over 20,000 customers. These customers place over five million calls per day over their Micom-based voice/data networks, and have saved over $4 billion in voice and fax toll costs.

This paper is one of several Micom white papers on voice/data integration. Others include *Bandwidth Management Technologies, Compression and Demodulation Technologies*, *Enterprise Networking Technologies*, and *Voice over Frame Relay*. This paper focuses on voice/data integration for TCP/IP networks, that is, VoIP integrated networking.

TCP/IP NETWORKING

TCP/IP is a common networking technology made popular by the Unix computer operating system and the Internet. It is now available across virtually all operating systems and is widely used by many enterprises for a variety of applications. In fact, there are estimates that upwards of 90% of the corporate world will be using TCP/IP by 2000, with similar figures for governmental and educational sectors. Any network that connects to the Internet must run TCP/IP, and the Internet is growing rapidly.

TCP/IP is, specifically, a data communications protocol. A protocol is a set of agreed-upon conventions or standards for interactive behavior. Communications protocols define how two devices on a network are to behave when communicating with each other; when to talk (send), when to listen (receive), for how long, how to correct errors, the proper way of addressing each other, etc. The various aspects of a communications protocol are divided into layers of activity, where each higher layer builds on the foundation of the lower layers.

For example, the bottom layer designates the physical communications medium such as copper wire, fiber optic cable, microwaves, satellite, etc. The next layer specifies the electrical or optical signaling conventions for the medium. The "stack" of layers can go up to layer 7, where the behavior of the actual user application, such as e-mail or web browsing, is specified.

The term "TCP/IP" is named for a combination of layer 4 (TCP) and layer 3 (IP) of the so-called "IP" stack. "TCP" means Transport Control Protocol, and "IP" means Internet Protocol. All so-called "IP networks" use IP at layer 3; most also use TCP at layer 4. Figure 2 shows the relationship of the various IP layers in the IP stack.

5-7 Application + Layer	Telnet, FTP, SMTP, e-mail, WWW, SNMP	
4 Transport Layer	TCP or UDP	TCP is connection-oriented UDP is connectionles or datagram
3 Network Layer	IP	
2 Data Link Layer	Frame Relay, ATM, SMDS, HDLC, SDLC	
1 Physical Layer	Modem, 56K, T1, E1, BRI, PRI, RS-232, V.35, Sonet	
Medium	Fiber, Copper, Radio, Satellite, ...	

Figure 2 - IP Protocol Stack

Layer 3 is called the Network layer, and builds on layer 2, the Data Link Layer (DLL). The DLL is a set of protocols that moves data from one point in a data network to the next adjacent point. The various DLL protocols include Frame Relay, ATM, HDLC (High Data Link Control), etc. The Network layer is a set of protocols that move data from one end of a network to the other end, across one or more (usually many) successive intermediate network points.

A Cargo Transportation System

Layers 1, 2 and 3 operate in a manner similar to a cargo transportation system. Layer 1 might include highways, train tracks, and shipping lanes. Layer 2 would be the rules to go from one city to the next city, one rail yard to the next, or one port to next. And layer 3 would be the rules to go from city to rail yard to port to city to rail yard to Öetc. The packet of data mentioned earlier could be a unit of containerized cargo, for example. Taken together, they form a system to move cargo from any originating point across a broad expanse to any desired destination.

Layer 4 operates to ensure that the correct cargo arrives at the correct destination and is in good condition, somewhat like the receiving and inspection part of a cargo shipment. Layers 5, 6 and 7 serve to deliver the cargo to the end-user at the destination in a way that the end-user understands and is comfortable with.

If the cargo includes a letter, for example, these layers would be concerned about the language of the enclosed letter (English, Spanish), the type font, type size, style of stationary, and so on. Taken together, all seven layers can provide for the transportation of the letter from the sender to the receiver for reading in a precise format and style.

For IP networks (and several others), the communications device that manages the network layer is called a router. Routers are layer-3 datacomm devices responsible for routing data from one end of a network to the other; they can be likened to dispatchers and transportation managers in our cargo example. VoIP applications assume a router-based IP data network supporting both TCP and its cousin, UDP (Universal Datagram Protocol).

UDP, like TCP, functions to send data so that it arrives at the correct destination and in good condition. UDP, however, unlike TCP does not cause the originating device to retransmit the data if an error is detected. And UDP does not first “handshake” with the receiving end, allowing for faster initiation of the sending process — e-mail, for example, rides on top of UDP, unlike file transfer (FTP), which rides on top of TCP.

In a VoIP network, the IP layer can ride over a mixture of different layer 2 and 1 protocols. For example, IP might run over frame relay running at layer 2 on the Sonet (Synchronous optical network) optical light layer 1 protocol, running over fiber optic cable. Also, because the DLL can have a limited ability to do single multi-successive point-to-point networking (using the “high” part of the DLL protocols), it is quite possible to have IP running on frame relay (high part of DLL) on ATM (low part of DLL) over, say, microwave. The key here is to focus on the IP layer and let Layers 2, 1 and the physical medium take care of themselves.

Three Types of IP Networks

It is convenient to group all IP networks as one or a combination of these three basic types: the Internet; the corporate or Enterprise IP network, sometimes called the Intranet; and the IP Virtual Private Network (IP-VPN), sometimes called the Extranet. Some IP networks are comprised of all three types. The Internet and its World Wide Web are probably the most prominent, but for VoIP, the other two are equally important. (See Figure 3.)

An enterprise IP network consists of one or more LANs at each corporate location connected together by a corporate WAN. The LANs support IP usually on one or more of the various types of Ethernet, Token Ring, ATM, or FDDI technologies. The WAN supports IP typically over leased lines, public frame relay, ATM, satellite, or ISDN connections.

At each remote location, the enterprise will have routers connecting the remote site’s LAN to the IP WAN; see Figure 4. The enterprise IP network is a managed IP network, that is, data movement between sites is done expeditiously with low delay and high predictability, because the corporation’s operation is so dependent on the timely delivery of data.

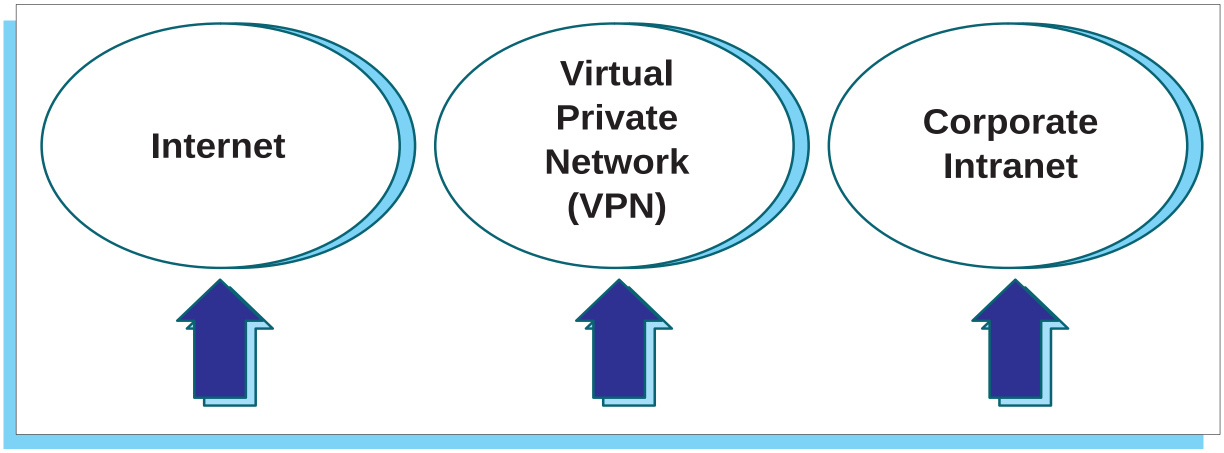

Figure 3 - Three Basic Types of IP Networks

This is in contrast to the Internet, where delays can be large and arrival time highly unpredictable.

The IP-VPN is a fairly new type of public network offering, intended to provide the managed IP network characteristics that a corporation needs. It can be thought of as an “industrial strength” Internet. Because IP-VPNs are built with plenty of bandwidth and high-speed routers, delays are low and arrival times are predictable. In fact, some IP-VPNs offer low delay guarantees, security features, and/or bandwidth guarantees in an effort to mimic enterprise IP network characteristics.

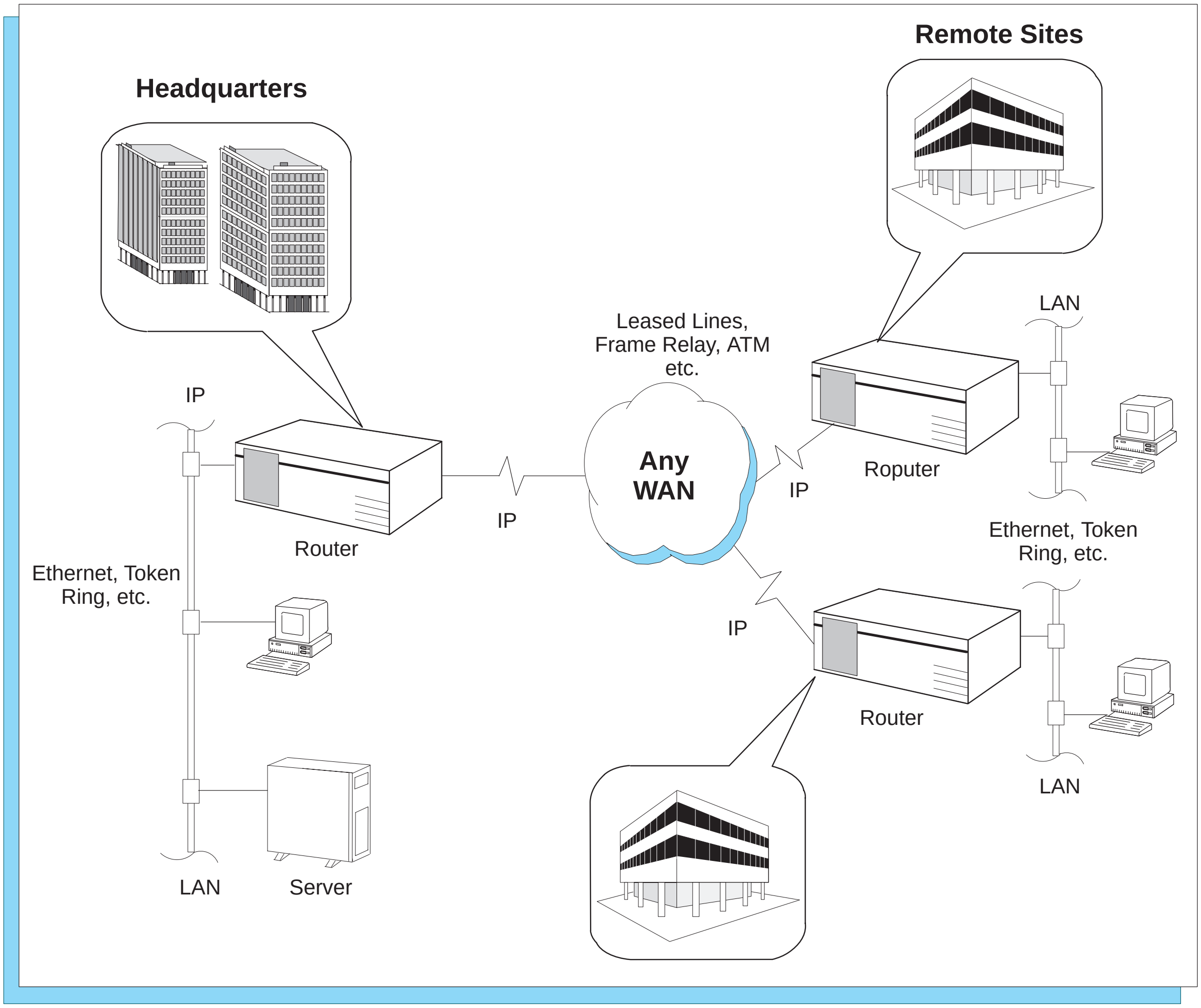

Figure 4 - Wide Area IP Connectivity

These guarantees make up a performance characteristic called Quality of Service (QoS). QoS is the key differentiation between the IP-VPNs and the Internet. The goal of IP-VPN service providers is to convince corporations to build their IP networks using the IP-VPN instead of frame relay, ATM, etc. It remains to be seen how well they compete.

SEVEN VOICE/FAX-OVER-IP MARKETS

The main goal of VoIP is to piggy-back voice and fax calls over an IP data network to save on long distance charges. A secondary goal is to incorporate IP voice and fax into certain applications for enhanced services. These two goals are the primary focus in seven main VoIP market applications.

Table 2. Seven VoIP Markets and Applications

Market	Application
1. Corporate Toll-Bypass	Toll-free intra-company voice and fax between corporate locations.
2. Fax over the Internet	Toll-free or reduced-rate fax-machine fax between any two locations.
3. PC-phone to PC-phone	Toll-free voice between any two PCs on the Internet (fax between PCs is just data; fax-machine fax is addressed in 1 and 2).
4. IP-Based Public Phone Service	New public phone services, at reduced rates (especially internationally), where voice is sent over the Internet or over new public IP networks. Voice is phone-to-phone, PC-to-phone, or phone-to-PC.
5. Call-Center IP Telephony, Agent-click	A new IP voice application that allows a PC user on the Internet to click on a phone icon in a catalog at a customer service home page and talk to an agent via the PC as a phone.
6. IP Line Doubler	A PC user at home or in a hotel, etc. with just one connection to the Internet would subscribe to a new service that allows the single phone line to carry one or more phone calls in addition to the PC data.
7. Premise IP Telephony	PCs in a building on an IP LAN would be able to make phone calls to ordinary phones in the same building or to make outside calls, using special VoIP equipment on the premise.

Markets 1, 4, 6, and 7 rely on a new type of communications equipment called the voice/fax IP gateway (or Voice-over-IP telephony gateway). This VoIP gateway converts voice and fax into IP packets and puts the packets onto a LAN, while at the same time taking packets off the LAN and converting them back to voice or fax signals. This two-way operation is called full duplex (FDX) communication. VoIP gateways can handle multiple FDX conversations, usually up to 24 or 30, with some going to 96 or higher.

Market 2 also relies on an IP gateway, but one that packetizes only fax. This gateway may also have certain added features such as fax store-and-forward, for delayed transmission and long-delay tolerance, and fax broadcast, for faxing one document to multiple recipients.

Market 3 does not use an IP gateway since the PCs perform the gateway functionality. The PC has a sound card, speakers and microphone, or a phone card and ordinary telephone. The PC does the voice packetizing. PC-fax is just data; only fax-machine fax needs a gateway. Aside from the above hardware, the telephony technology is all software.

Market 5 usually uses an IP telephony gateway, but if the call-center agent or customer service representative is using a PC as a phone, no gateway is required. IP fax is usually not involved. The call-center application is a promising VoIP application that is not driven primarily by toll-cost avoidance, rather by internet electronic commerce (E-Commerce) and on-line catalog purchases.

This white paper discusses Market 1, the Corporate Toll-Bypass Market. This VoIP market has been the fastest to mature, and led the industry in total sales in 1996. We estimate that by the end of 1997 there were over 1,000 corporate VoIP networks installed worldwide.

THE CORPORATE TOLL-BYPASS MARKET

The popularity of the IP protocol in corporate data networks has soared in the past few years. Influenced by the Internet, networking managers have increasingly adopted IP as the layer 3 protocol foundation for their networks. This has manifested itself as a huge shift away from the Novell IPX layer 3 protocol, and as a lesser but still significant shift away from the IBM SNA protocol at the networking layer.

IP is now supported over a wide variety of layer 1 and layer 2 network technologies: Ethernet (10, 100, 1000 M bps), Token Ring, FDDI, Frame Relay, ATM, Sonet, leased lines and satellites. There are virtually no routers, FRADs or network switches that do not support IP. A group of techno-wizards have even demonstrated a toaster connected to a network using IP!

Just about any type of corporate data network can be built on IP. It scales to millions of nodes and users, and does equally well at remote sites or corporate headquarters. Figure 5 illustrates the broad applicability of IP. Note that the data network can run other protocols, such as Digital

Equipment Corporation's DECNet and Novell's NetWare, as long as it also runs IP.

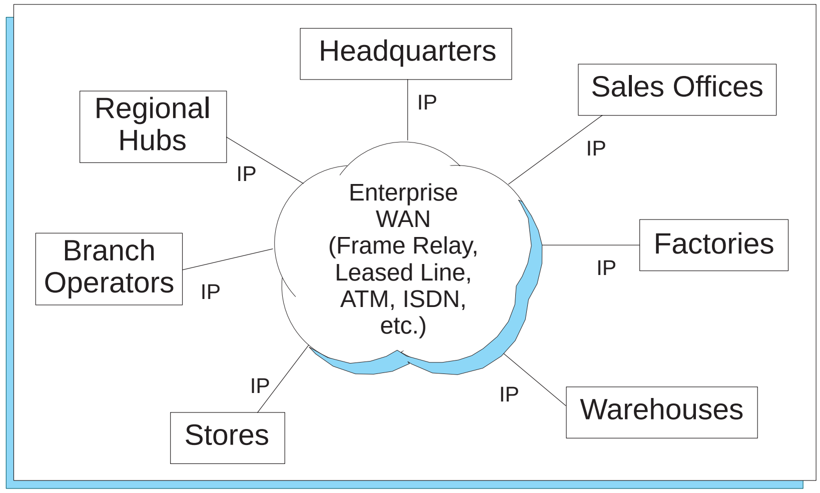

Figure 5 - Compare IP Networking

By running intracompany, inter-site voice/fax over its IP data network, a company can expect to cut intracompany telephony costs by 70%–80%. It can expect to save $10,000–$20,000 or more per remote site over a 5-year period, and it can expect to average a net toll-rate for those calls of 1 to 2 cents per minute over that same period. The installation of VoIP gateway equipment usually has a payback or break-even period of under one year in countries with low telephony costs (U.S., U.K., Chile, China and example), and is well under a year in higher cost countries (Argentina, Japan and Germany for example).

These attractive savings opportunities have made the corporate VoIP market the largest market segment of the seven VoIP markets. There are over two million installed remote routers and FRADs, with over 50,000 shipping per month. This equates to a potential $2–$3 billion VoIP gateway market growing at $50–$100 million/month. Several industry reports are projecting annual sales of VoIP products in this corporate VoIP market, in the 2000–2001 time frame, to be over $1 billion/year. As a result, there are now over a dozen vendors of corporate VoIP products.

Figures 6 and 7 illustrate how a company can deploy VoIP gateways in its network. In Figure 6, without voice/data integration, all voice and fax communications between and among the remote sites and headquarters go over the PSTN, while the data travels over LANs, router/FRADs/switches/hubs, and WANs of any makeup, as long as they support IP.

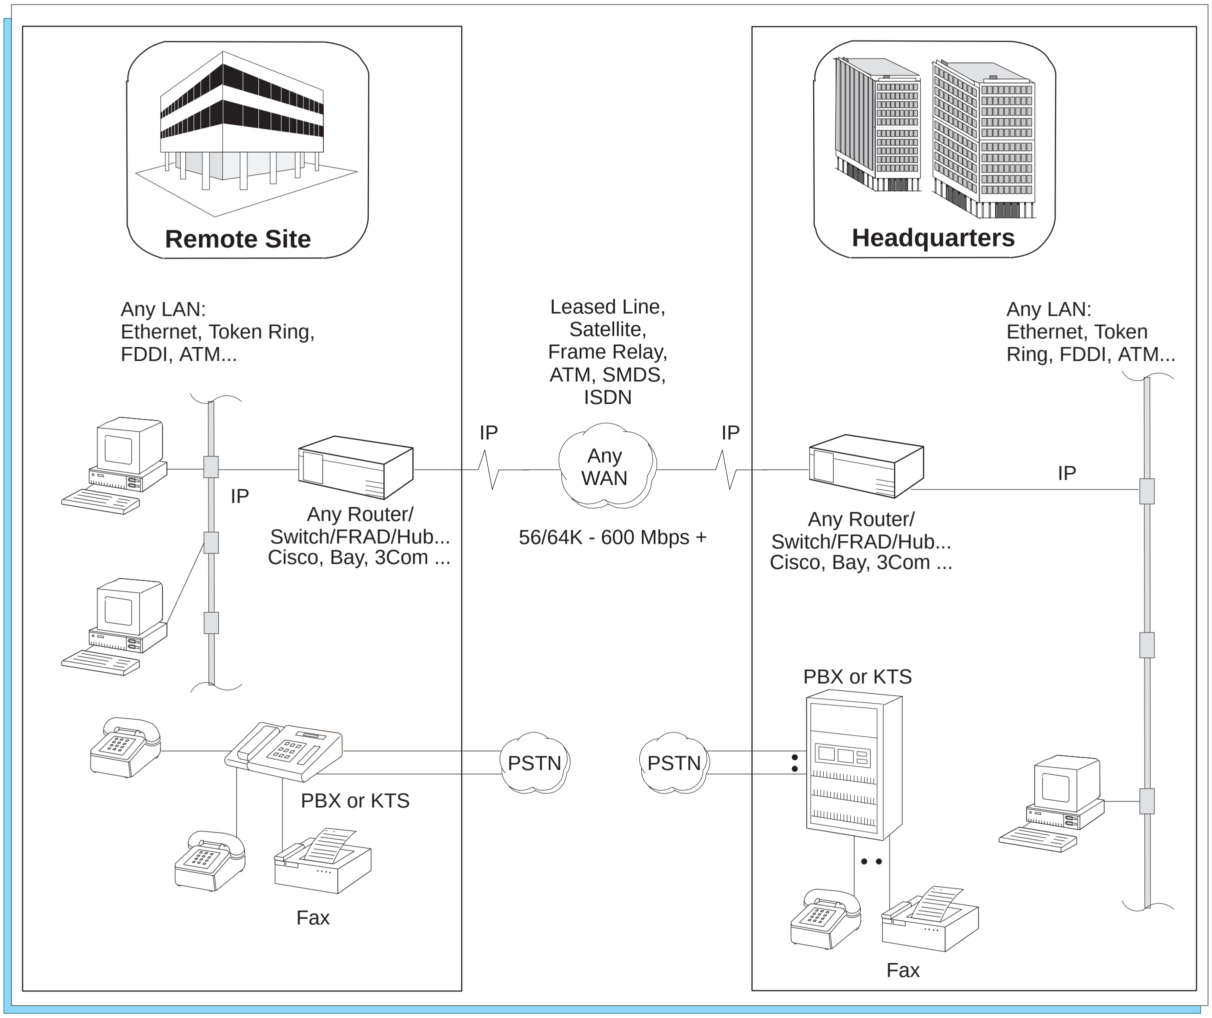

Figure 6 - Separate Data, Voice/Fax

In Figure 7, a VoIP gateway has been installed in a remote site, connected on one side to the remote site LAN (any kind) and on another side to the PBX or Key Telephone System (KTS) at the remote site. The remote site phones and fax machines connect to the PBX/KTS, and the PBX/KTS connects to the PSTN and gateway. The connections between the gateway and the PBX/KTS are usually analog in North America and ISDN BRI (Basic Rate Interface) or analog elsewhere. They can also be T1/E1 or ISDN PRI (Primary Rate Interface) connections.

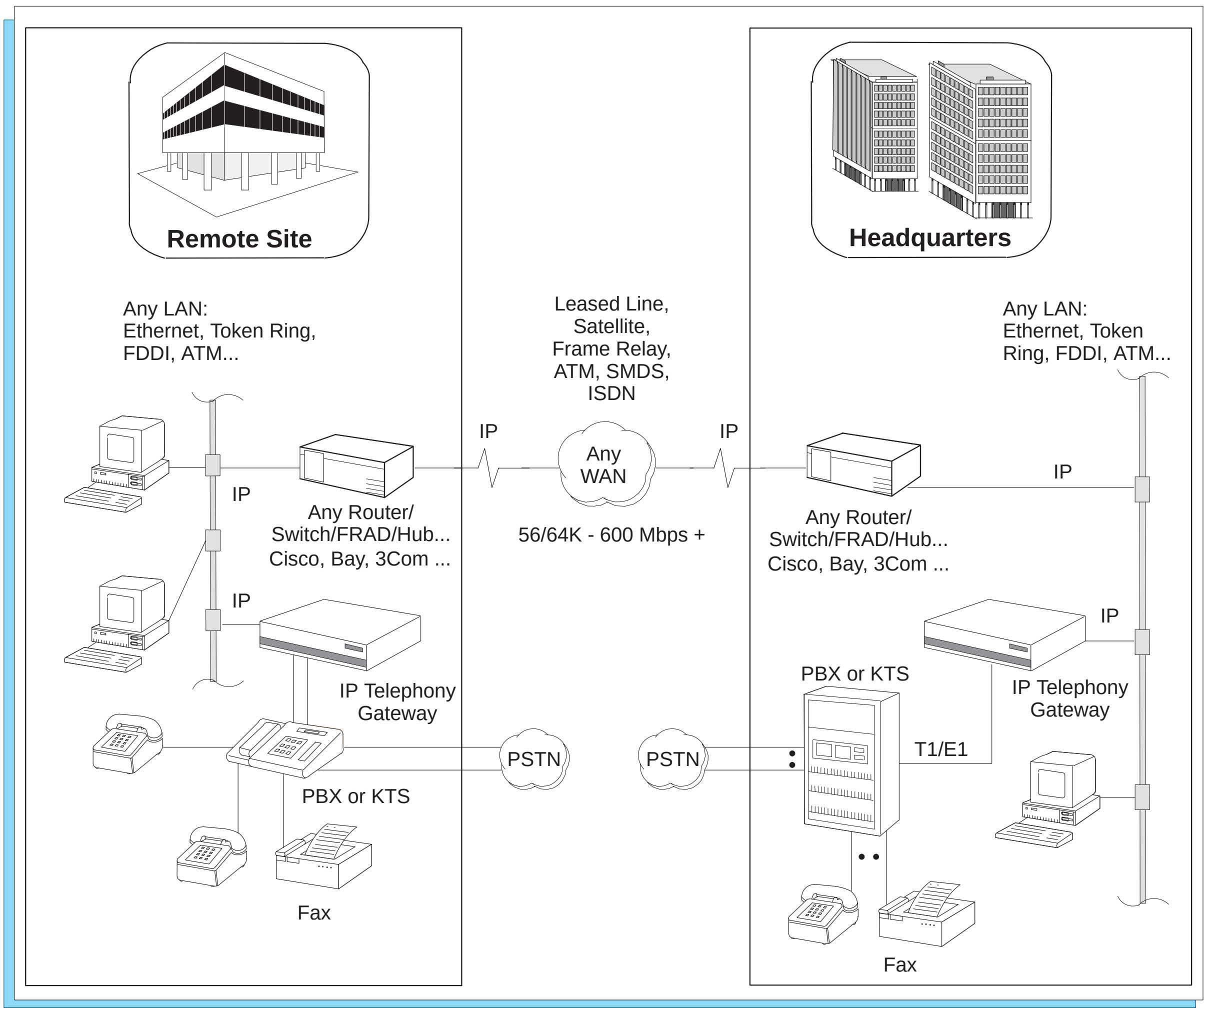

Figure 7 - Integrated Data, VoIP Voice/Fax

THE VOICE/FAX GATEWAY

In corporate VoIP applications, the VoIP gateway performs six basic functions:

1. LOOK-UP FUNCTION: When an IP gateway (originating gateway) is used to place a call across an IP network, it receives a called-party phone number from the calling phone and converts it into the IP address of the far-end or called-party's gateway (destination gateway). This can be a table lookup in the originating gateway or in a centralized directory server. Table lookups in the originating gateway are often much faster, reducing the call-connect to 1–2 seconds, compared to 4–5 seconds for a centralized server.

2. CONNECTION FUNCTION: The originating gateway establishes a connection to the destination gateway, exchanges call-setup and compatibility information, and performs any option negotiation and security handshake.

3. DIGITIZING FUNCTION: Analog telephony signals coming into a trunk on the gateway are digitized by the gateway into a format useful to the gateway, usually 64Kbps PCM (Pulse Code Modulation). This requires the gateway to interface to a variety of analog telephony signaling conventions.

   In many cases, it may also require the gateway to accommodate digital ISDN and T1/E1 Interfaces. ISDN and T1/E1 interfaces are already in the 64K PCM format, so the gateway bypasses the analog-to-PCM conversion step. ISDN BRI has one or two PCM channels, T1 has up to 24 PCM channels, and E1 has up to 30 PCM channels. ISDN PRI can have up to 24 or 30 PCM channels.

4. DEMODULATION FUNCTION: With some gateways, the gateway trunk can accept only a voice signal or a fax signal, but not both. The trunk must be dedicated to either voice or fax connectivity. In more sophisticated gateways, the gateway can handle both types, automatically determining whether the digitized signal is voice or fax, and processing the signal according to its type.

   When the signal is fax, it is demodulated by the DSP back into the original 2.4–14.4Kbps digital format that it was in before it left the fax machine (as fax leaves the fax machine, the original digital signal is converted to an analog format). This demodulated information is then put into IP packets for transmission to the destination gateway (see Figure 8). The demodulated info is remodulated back into the original analog fax signal by the remote gateway, for delivery to the remote fax machine.

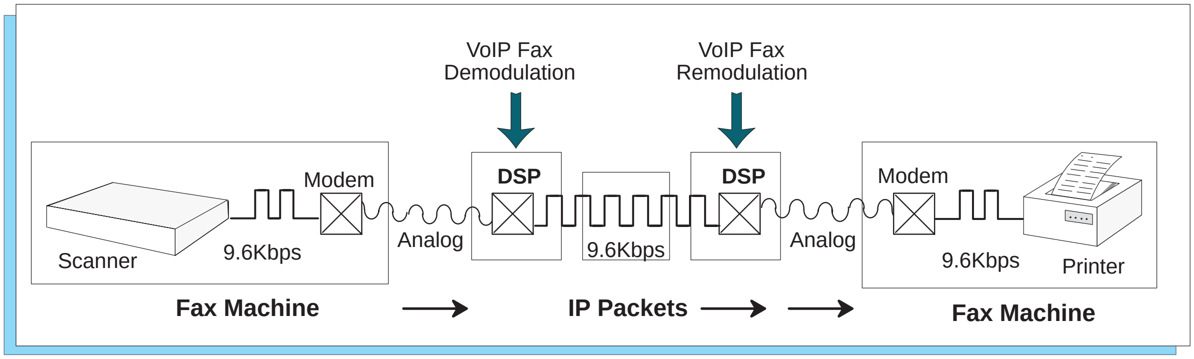

Figure 8 - VoIP Fax Domodulation/Remodulation at 9.6 Kbps

The fax transmission can be sent using either UDP/IP or TCP/IP format. UDP/IP, unlike TCP/IP, does not try to correct transmission errors by retransmitting the packet. So it would seem that UDP/IP might be preferable since a corrupted fax packet might affect only one line of the fax. However, if lost packets occur during the negotiation of a page, the fax could be terminated. When TCP/IP is used and the host software hides the TCP retransmission from the fax machine, there will be no impact.

Store-and-forward fax is sent using TCP/IP to take advantage of TCP's error correcting facility. For a more detailed discussion of fax demodulation, see also the MICOM white paper *Compression and Demodulation Technologies.*

5. COMPRESSION FUNCTION: When the signal is determined to be voice, it is usually compressed by a DSP from 64K PCM to one of several compression/decompression (CODEC) formats — see Table 3 below — and put into IP packets. PCM is shown for reference. Good voice quality and low digitizing delay are the goals.

Table 3. CODEC Characteristics

VoIP CODECs	Compressed Voice Digitizing Rate (Kbps)	Complexity	Quality	Digitizing Delay
G.711 PCM	64	N/A	Very Good	Negligible
G.726 ADPCM	40/32/24	Low (8 MIPS)	Good (40K) to Poor (16K)	Very Low
G.729 CS-ACELP	8	High (30 MIPS)	Good	Low
G.729A CA-ACELP	8	Moderate	Fair	Low
G.723 MP-MLQ	6.4/5.3	Moderate-High (20 MIPS)	Good (6.4K) Fair (5.3K)	High
G.723.1 MP-MLQ	6.4/5.3	———	———	———
G.728 LD-CELP	16	Very High (40 MIPS)	Good	Low

The voice packet is constructed as a UDP/IP packet, to avoid TCP/IP's attempt to correct a corrupted packet by retransmitting the packet. Any attempt to retransmit a voice packet would introduce too much delay and be of no value. If FEC (forward error correction) is enabled, a corrupted or missing voice packet can be recreated from the FEC data stored in the previous voice packet. If FEC is not enabled, then corrupted packet is simply discarded, and the previous good packet is reused by the destination gateway.

This packet replay is rarely noticed by the listener if kept to a low percentage of packets (<5%). If the degree of packet corruption is too great, IP telephony relies on the time-tested voice correction technique, called "What" in English, "Was" in German, etc. It usually works every time.

The CODEC's digitizing rate does not include the additional bandwidth consumed by the IP packet's addressing and control information (called "header" info) (see Figure 9). This is usually around 7Kbps, unless the IP router separately performs data header compression, which can reduce the header overhead to 2Kbps–3Kbps.

28 bytes	6 bytes	40 bytes
IP/UDP Header	Voice Header	Voice Info

Figure 9 - VoIP Packet Format

Complexity is determined by the typical amount of DSP processing power, as measured in MIPS (Millions of Instructions Per Second), required to support FDX communications. It does not include the MIPS required to perform auxiliary voice processing functions, such as echo cancellation and silence suppression. Lower complexity implies lower DSP costs. The G.7xy designators are assigned by the ITU (International Telephony Union) and refer to a standardized CODEC.

The advantages of voice compression are obvious: voice consumes less bandwidth, leaving more for data or other voice or fax communications. WAN bandwidth is expensive in proportion to the amount used and, sometimes, in proportion to its distance. If voice is left at 64K PCM, a 56K connection or even an ISDN 64K connection would not be able to carry one IP telephony conversation. For a more detailed discussion of voice compression, see Micom's white paper *Compression and Demodulation Technologies*.

6. DECOMPRESSION/REMODULATION FUNCTION: At the same time as it performs steps 1–4, the gateway is also receiving packets from other IP gateways and decompressing voice information back to the format, to be connected to the appropriate analog telephony interface, or to an ISDN or T1/E1 interface. Or, it can be remodulating a digital fax signal back to the incoming format, and then into the appropriate telephony interface.

Note that the originating telephony interface (analog, ISDN, or T1/E1) need not match the terminating telephony interface. The gateway-to-gateway connection can usually perform the necessary conversions.

Comparing VoIP Gateways with VoFR Router/FRADs

The VoIP gateway process described above is different from the voice-over-frame relay approach used by voice/fax routers and voice/fax FRADs (frame relay access devices) (see Figure 10). The VoIP gateway is fundamentally a LAN device. Although its main application is in reducing the costs of WAN phone and fax calls, the WAN connectivity is supplied by a router, not the gateway. In fact, any WAN access device that will carry IP will work. As ATM becomes more common as a WAN access technology, IP over ATM via an ATM access device (ATM-AD) will materialize, as will voice /fax over IP over ATM

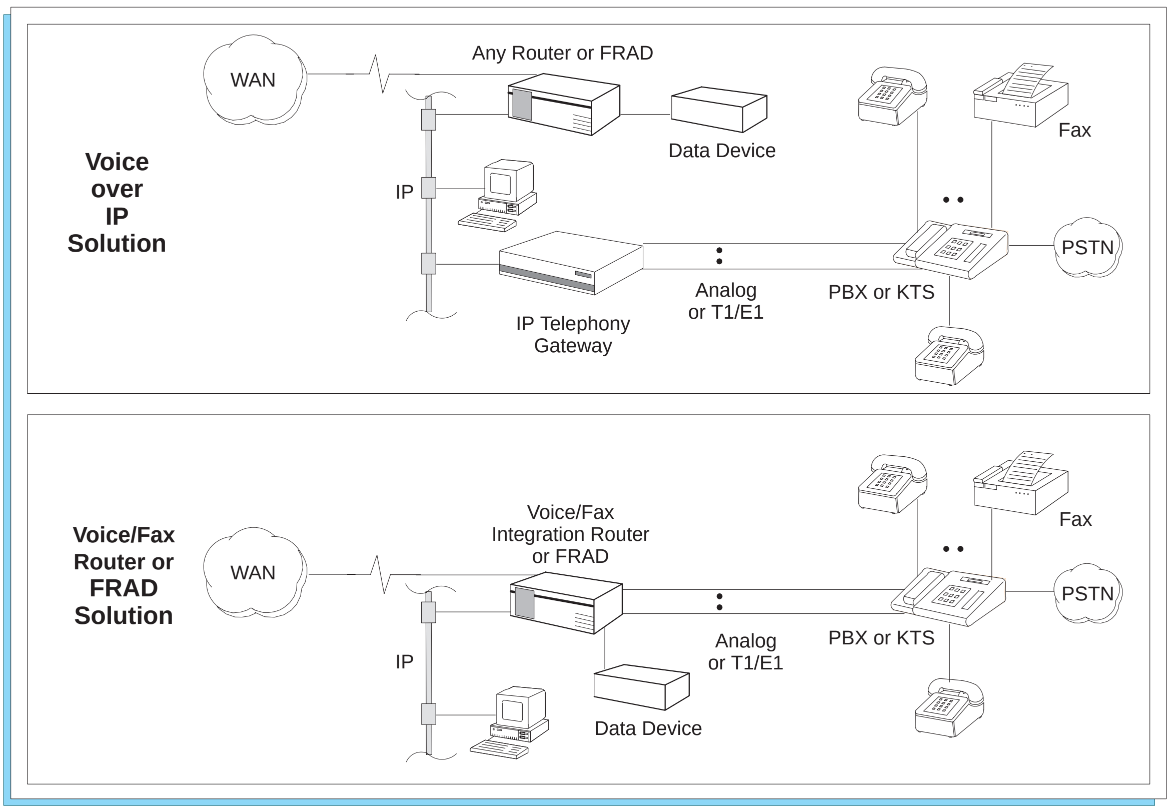

Figure 10 - Comparing VoIP Gateway with Voice/Fax Router or FRAD

The voice/fax router or FRAD, however, is fundamentally a WAN device. It connects the LAN to the WAN, and much of its technology is devoted to this function. It also connects data devices, such as SNA cluster controllers and async terminals, to the WAN. It is a more sophisticated device whose operation and position in the network require a high level of robustness, expertise and manageability. By comparison, the gateway is much less intrusive and pivotal to the operation of the network.

A failure or misbehavior in the gateway affects only the voice/fax communications, and the ordinary PSTN (Public Switched Telephone Network — sometimes called POTS, for Plain Old Telephone System) is always available as a backup. A failure or misbehavior in the router or FRAD affects the entire network. This gives the gateway an advantage over the FRAD/router in simplicity, unobtrusiveness, and non-criticality. The gateway may also be a lower cost integration solution.

The voice/fax FRAD/router, however, usually has an advantage over the gateway in efficiency, because the IP header overhead can be eliminated by the FRAD/router's sending voice and fax directly over the leased line or frame relay connection in proprietary packet format, not IP packet format. The FRAD/router can also have an advantage, depending on the capabilities of the comparable VoIP gateway, in segmenting and prioritizing the voice/fax traffic ahead of data traffic.

DEPLOYING THE VoIP GATEWAY

At a remote site, there are usually two to four VoIP connections (or trunks) from the VoIP gateway to the PBX/KTS, allowing for two to four simultaneous phone/fax connections between the remote site and other corporate locations. The actual number of trunks needed depends on the number of calls made during the day, and the total amount of time the calls consume. Most remote sites (bank branches, supermarkets, stores, and sales offices, for example) have between two and four hours total calling time, which is usually handled by two trunks. Table 4 shows the percentage of time the caller will get dial-tone when attempting to place a VoIP phone call, for gateways with two, three and four trunks.

Table 4. Off-Hooks Receiving Dial-Tone (%)

Number of Trunks	Number of Hours of Calling Per 8-hour Day						
	2	**3**	**4**	**5**	**6**	**7**	**8**
2	98	95	92	89	86	83	80
3	100	99	99	98	97	95	94
4	100	100	100	100	99	99	99

An availability of 94% means, for example, that when a remote-site user attempts to place a call over the VoIP gateway, the user will get dial-tone 94% of the time, on average. Any availability in the 90% and higher range is acceptable. Below 90%, the caller can become frustrated by too-frequent busy signals and dial directly over the PSTN, thwarting the VoIP gateway's main function, to reduce WAN telephony costs.

This table assumes an equal distribution of calls throughout the 8-hour day, with no time zone effects. If calls tend to bunch at certain times during the day, or time zone differences create a shortened overlap period, the percentages would be reduced. For example, two trunks can handle four hours per day evenly distributed, but would be insufficient if all calling happened in, say, a 6-hour period. In this case, three trunks would be a better strategy.

At the headquarters location in Figure 7, the IP gateway connects to the headquarters PBX/KTS. For most companies, this will probably be a PBX, and the connection will probably be a T1/E1 or ISDN PRI line if there is a need for more than eight VoIP trunks at headquarters; otherwise, the connection could be analog or ISDN BRI.

Trunk Contention Ratios

The number of headquarters trunks is determined by the total number of phone calls between headquarters and the remote sites and the total number of simultaneously active remote-headquarters calls. It is uncommon to have a one-to-one trunk count, that is, one trunk at headquarters for each remote trunk; this would be needed only if each remote-site trunk would be expected to be active to headquarters at the same time. Typically, headquarters will have only a fraction of the total remote trunk count, and the larger the number of remote trunks, the smaller the fraction. Table 5 shows some contention ratios. The contention ratio is the ratio of total remote-site trunks to headquarters trunks. The second column is for reference only.

Table 5. Contention Ratios of Total Remote Trunks to HQ Trunks

Total Number of Remote Trunks	Number of Remote Sites at 2 Trunks per site	Typical Number of Headquarters Trunks	Typical Contention Ratios
2	1	2	1:1
4	2	3-4	1.3:1 - 1:1
6	3	4	1.5:1
8	4	5-6	1.6:1 - 1.8:1
10	5	5-6	1.7:1 - 2:1
16	8	8-9	1.8:1 - 2:1
24	12	10-13	1.8:1 - 2.4:1
32	16	13-16	2:1 - 2.5:1
64	32	24-29	2.2:1 - 2.7:1

A necessary IP gateway function is the ability of the gateway to support rotary trunk capability. Since the number of headquarters trunks is usually less than the total number of remote-site trunks, there must be a way to share the gateway's headquarter's trunks. This is done through a rotary trunk provision. The next VoIP call from a remote site to headquarters is automatically assigned the next available headquarters gateway trunk (that is, it rotaries to the next available headquarters trunk), without the caller having to specify a trunk number.

A related important capability is cross-gateway rotary provision, whereby the next call is assigned the next available trunk from among a bank of VoIP gateways. This assures VoIP calls are placed the fastest way at the least equipment cost.

MAKING A VoIP PHONE CALL

One of the appealing aspects of IP telephony via an IP gateway over the corporate IP network is the preservation of the existing phones, fax machines, PBXs, and KTSs. Since there is no need to use a PC to make an IP phone call, there is no retraining employees on how to place calls over the gateway — a single one-page memo usually suffices. By using the already-paid-for legacy telephony equipment, the costs to implement IP telephony are minimized.

During the installation of the VoIP telephony system, an IP gateway is placed at each remote site and one or more at headquarters. The network administrator assigns a location number to each gateway, usually easily remembered. For example, in a small network, headquarters might be 10, the sales office 20, warehouse 30 and factory 40.

In a larger network, area codes can be used, provided there is only one company site per area code. In networks with more than one site in an area code, area codes plus 1, 2, or 3 digit prefixes can be used. Once the numbers are assigned, a memo is circulated to all employees identifying each location's number, with brief call placement instructions.

There are three different scenarios for placing a call over the VoIP network:

1. From a telephone attached to a PBX to a remote phone attached to a PBX, or

   to a remote phone attached to a KTS or directly attached to the remote gateway.

2. From a telephone attached to a KTS to a remote phone attached to a PBX, or

   to a remote phone attached to a KTS or directly attached to the remote gateway.

3. From a phone directly attached to local gateway to a remote phone attached to a PBX, or

   to a remote phone attached to a KTS or directly to the remote gateway.

   Figures 11, 12, and 13 illustrate these scenarios, and Table 6 summarizes the connection sequences.

Table 6. VoIP Call Connection Sequences

Local or Originating Telephony Equipment	Remote or Destination Telephony Equipment: To Phone Attached to PBX	Remote or Destination Telephony Equipment: To Phone Attached to KTS, or to Phone Directly Connected to Remote Gateway
From Phone Attached to PBX	Take phone off hook, hear PBX dial-tone, key 6, hear VoIP dial-tone, key 408234, hear ringing or busy	Take phone off-hook, hear PBX dial-tone, key 6, hear VoIP dial-tone, key 408, hear ringing or busy
From Phone Attached to KTS	Depress VoIP key, take phone off-hook, hear dial-tone, key 408234, hear ringing or busy	Depress VoIP key, take phone off-hook, hear VoIP dial-tone, key 408, hear ringing or busy
From Phone Directly Connected to Gateway	Take phone off-hook, hear VoIP dial-tone, key 408234, hear ringing or busy	Take phone off-hook, hear VoIP dial-tone, key 408, hear ringing or busy

Phone Call Examples

In Figure 11 a caller in the Los Angeles corporate office using a phone connected to a PBX goes off-hook (lifts up the receiver) and hears PBX dial tone. At this point, the caller could dial a local PBX extension (e.g., 542), a “9” to get an outside PSTN line (or some other digit, e.g., “0” in some countries), or a special digit to access the VoIP Gateway. This special digit is set up by programming the PBX when the gateway is installed. Often a “6” or “8” is used. When the digit is keyed, the PBX is instructed to rotary to the next available gateway trunk and connect the caller to the gateway.

The caller hears dial-tone from the gateway and keys in (or dials-in on a rotary phone — a good IP gateway should accept both) the called-party’s location number as defined when the network was set up. In Figure 11, this is “408” for the company’s remote location in the San Jose area, the same as San Jose’s area code, for easy recall. At the same time, the caller keys in the extension of the called party (e.g., 234) if the remote site uses a PBX, or for a KTS or directly connected telephone, waits to hear ringing or busy.

When the local gateway receives the 408, it translates the 408 (via a table) into the IP address of the remote gateway and creates a TCP/IP or UDP/IP session with the remote gateway. Optimally, the table-lookup is done in the local gateway for the fastest connection time. Some gateways do a centralized table-lookup at a distant “master” gateway, which slows down the connection time, sometimes by 3–4 seconds. The centralized lookup also creates a single point of failure — if the master gateway is unreachable or misbehaving, no connections can be made.

For gateways designed for local lookups, changes to the lookup data base are downloaded across the IP network from a central administrative gateway to each local gateway. If for some reason the local lookup fails, the local gateway updates its tables from the administrative gateway on the fly. This is done without operator or user intervention.

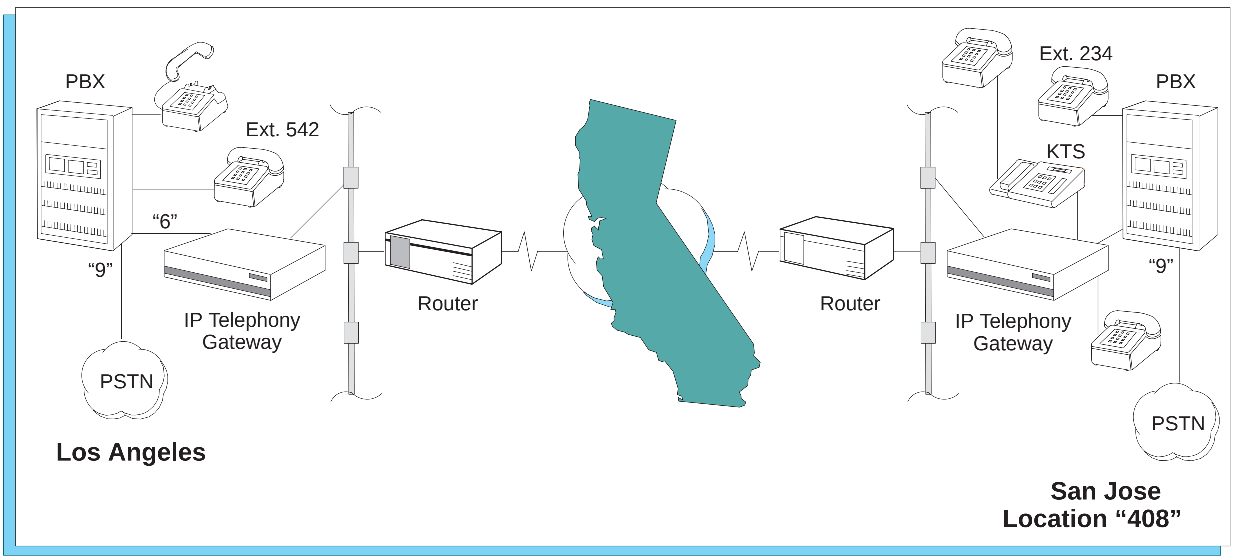

Figure 11 - Calling from a PBX using VoIP

Once the two gateways establish that a VoIP call can now be placed across the WAN (this takes less than one second in the local lookup case), the local gateway sends the extension digits 234 to the remote gateway for delivery to the PBX, and the PBX rings extension 234. For a remote KTS or directly connected telephone, the remote gateway rings the KTS or phone. In either case, the subsequent ringing or busy signal is passed back to the calling party.

The ringing or busy signal is generated by the remote PBX, or for a remote KTS or remote telephone, by the remote gateway. Once the called party answers, the gateways begin sending the voice or real-time fax transmissions using UDP/IP. The session is terminated when either party hangs up.

If the two gateways cannot link up, the local gateway returns a fast-busy signal to the caller. This can occur when a WAN link to the remote gateway is not available. A fast busy is also generated when there are no remote gateway trunks available, or when the remote Gateway does not respond.

When the originating telephony equipment is a KTS, a similar scenario is followed except the PBX dial-tone and 6-keying are replaced by depressing a key on the KTS phone that might be labeled something like "VoIP." If there are multiple gateway trunks, there will be one key for each trunk ("VoIP-1," "VoIP-2"), and the caller pushes the next unlighted key to seize a VoIP trunk. There is no trunk rotary function provided by a KTS. The caller immediately hears VoIP dial tone from the local gateway. The remaining sequence is the same as above. Figure 12 illustrates these scenarios.

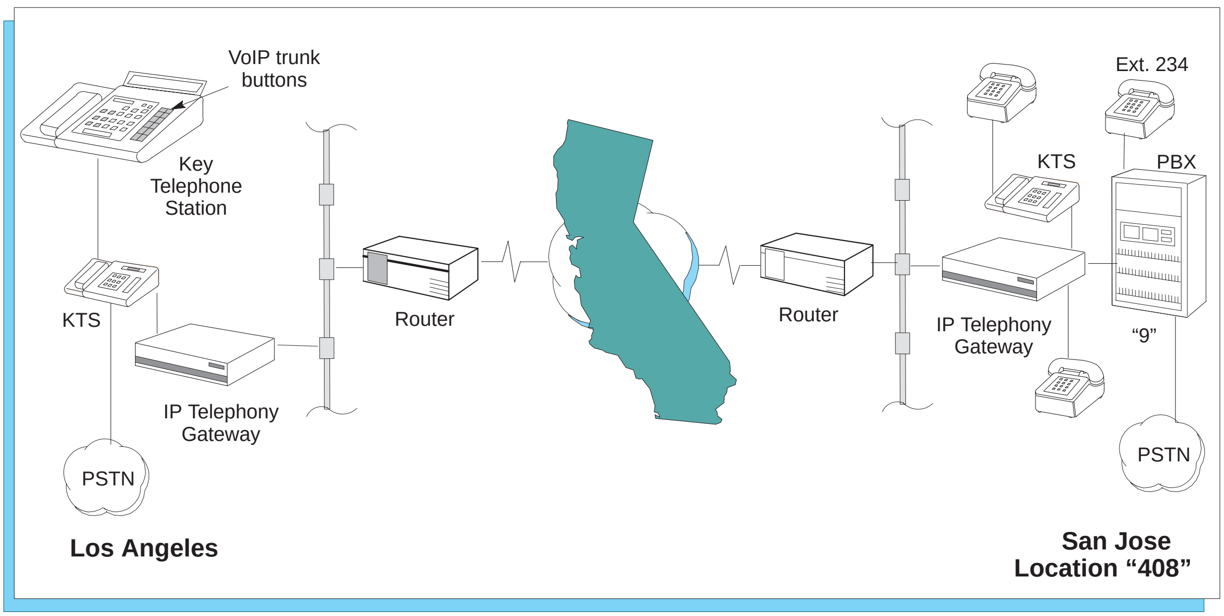

Figure 12 - Calling from a KTS using VoIP

For a call from a phone directly connected to a gateway, the caller goes off-hook and immediately hears VoIP dial-tone from the local gateway. The remaining sequence is the same as above. Figure 13 illustrates these scenarios.

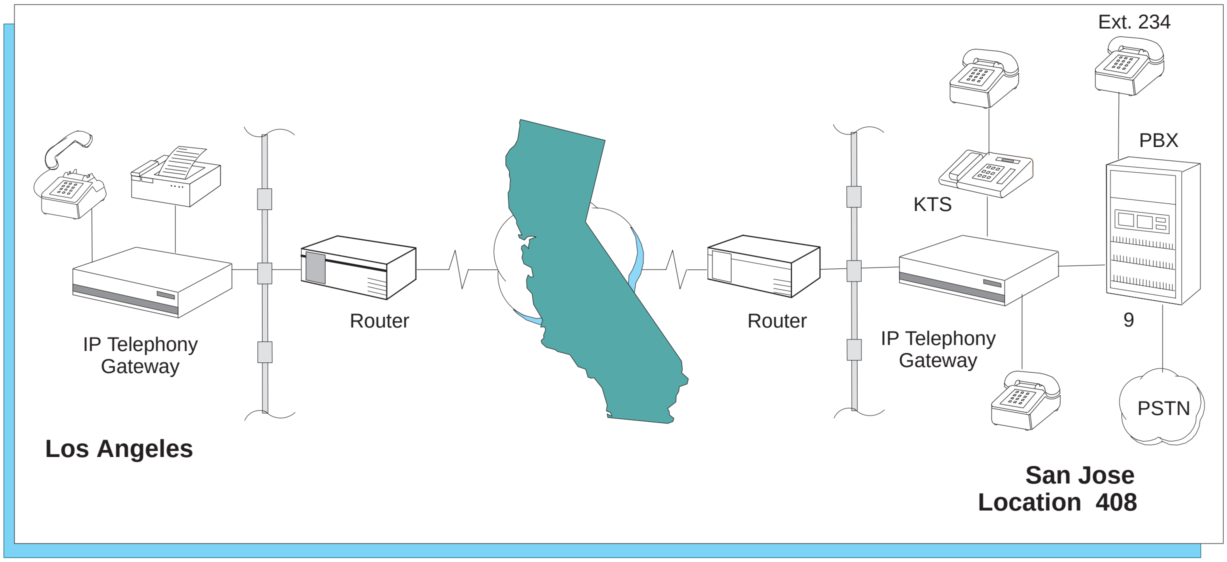

Figure 13 - Calling from a directly connected phone using VoIP

International Applications and Off-Net Calling

Because the IP gateway must be able to generate dial-tone, busy and ringing tones to the calling party, it is important for the gateway to provide a variety of different country tones to support networks in different countries. For a network in the U.K., for example, it is best for the caller to hear U.K.-style calling tones rather than, say, U.S. tones. The gateway should support the ten or so popular country-tone sets.

There are two types of VoIP calls that can be placed, on-net and off-net. On-net calls are between corporate locations, like the calls discussed above. Off-net calls are between a corporate location and a non-company site.

Off-net calling allows the caller to place a VoIP call from a local corporate location to a remote company site, and then place a second call over the remote-site's PSTN to the called party. The called party is off the VoIP network, hence it is an off-net call. Off-net calling requires a PBX at the remote company site. In Figure 14, for example, an off-net call is initiated from Toronto to Miami by dialing “63069” which instructs the Toronto gateway to connect to the Miami gateway and deliver the “9” to the PBX. This tells the Miami PBX to seize an outside line to the PSTN, so the caller hears Miami PSTN dial-tone in the Toronto corporate location.

Once the Miami PSTN dial-tone is heard, the caller dials the Miami area number, and talks to the local Miami party at local Miami metro-area toll rates. The phone company calls this a "leaky PBX" and tolerates its use; in most countries it is not illegal. Off-net cost savings are in addition to the savings from intracompany calls, and can be major in international VoIP networks.

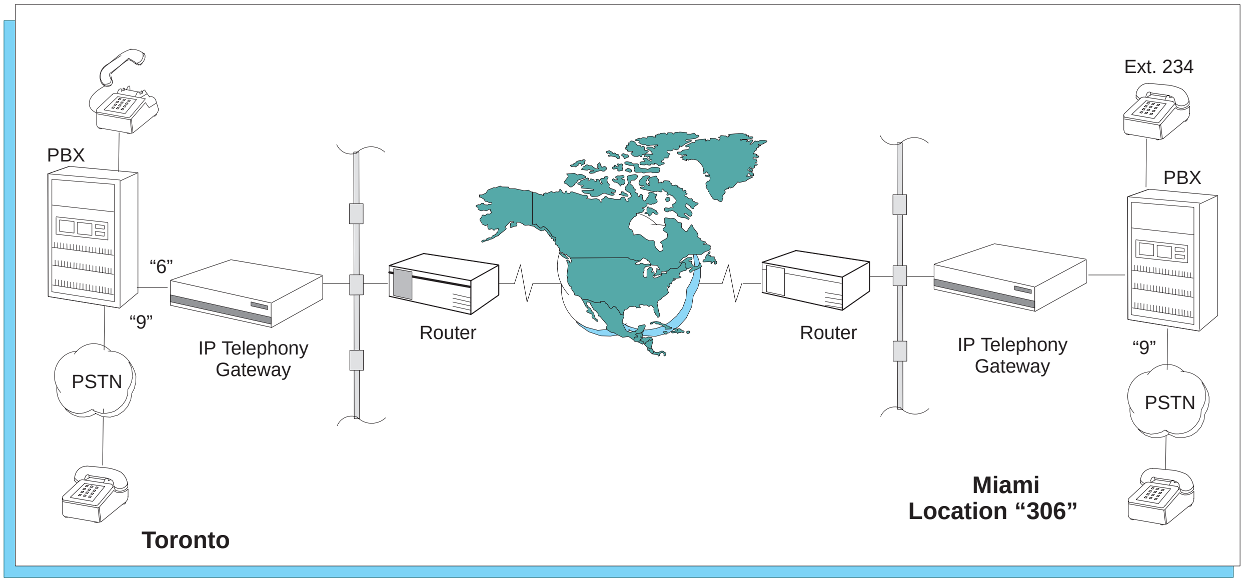

Figure 14 - Off-net calling using VoIP

If the Miami PBX allows it, a call into the Miami PBX from a Miami metro-area phone over the PSTN could be routed over the VoIP network to the Toronto gateway and connect to a phone in the Toronto corporate site. Moreover, if the PBX at the Toronto corporate site is also leaky, it is possible to make local toll-rate Miami-area to Toronto-area VoIP calls using two leaky PBXs. Double-leaky PBX setups are frowned on by the telcos, and may be illegal. But they can really reduce telco charges.

VoIP VOICE QUALITY

The keys to VoIP's success in business applications are high voice quality, low bandwidth consumption, and gateway payback periods of less than 18–24 months. Because VoIP can save a company hundreds of thousands of dollars, it is tempting to believe that the voice quality only has to be "OK", and employees will eagerly use VoIP to save the company money. In fact, if VoIP's voice quality is not toll or near-toll quality, the only persons in the company using it could be the controller and chief financial officer.

The desk worker usually has a choice when making a long distance phone call: dial “9”, for example, and use the PSTN, or dial “6” and use VoIP. The worker has a task to accomplish and business to conduct, and is not measured by how much he or she saves in telephone bills. Indeed, phone bill charges are usually buried in a general allocations budget item. The phone is a tool for the worker, and its quality is taken for granted.

To deliver the best voice quality, a VoIP gateway must use a CODEC with good voice quality and low delay. As Table 3 recapped, the characteristics of G.729 do very well, making it the choice for many gateways. It is also the ITU’s officially recommended CODEC for all wide area networking applications.

There are several other technologies necessary to ensure good voice quality: two of them are echo cancellation and packet prioritization. Echo cancellation is a function of the DSP, while packet prioritization is a function of the router and gateway.

Echo Cancellation

When a two-wire telephone cable connects to a four-wire PBX interface or a telco central office (CO) interface, a special electrical circuit called a hybrid is used to convert between two wires and four wires. Although hybrid circuits are very efficient in their conversion ability, a small percentage of telephony energy is not converted but instead is reflected back to the caller. This is called echo.

If the caller is near the PBX or CO switch, the echo comes back so quickly it cannot be discerned. However, if the delay is more than about 10 milliseconds, the caller can hear an echo. To prevent this, gateway vendors include special code in the DSPs that listens for the echo signal and subtracts it from the listener’s audio signal (see Figure 15). Echo cancellation is especially important for gateway vendors because the IP network delay can easily be 40–50 milliseconds, so the echo from the far-end hybrid would be quite pronounced at the near end. Far-end echo cancellation eliminates this.

Network Delay and Jitter

A major contributor to reduced voice quality is IP network packet delay and network jitter. Network delay describes the average length of time for a packet to traverse a network. Network jitter describes the variability in arrival time of a packet. Delay is like the average, jitter is like the standard deviation. Both are important in determining voice quality.

As the network delay (total time, including CODEC processing time) exceeds about 200 ms, the two speakers will increasingly adopt a half-duplex communications mode, where one speaks, the other listens and pauses to make sure the speaker is done. If the pauses are ill timed, they end up “stepping” on each other’s speech. This is the problem that occurs when two people converse over a satellite telephony connection. The result is a reduction in perceived voice quality.

When a VoIP voice packet is inordinately delayed and does not arrive at the far-end in time to fit into the voice stream going out of the far-end gateway, it is discarded, and the previous packet is replayed. If this happens too often, or twice in a row, the listener will perceive reduced voice quality.

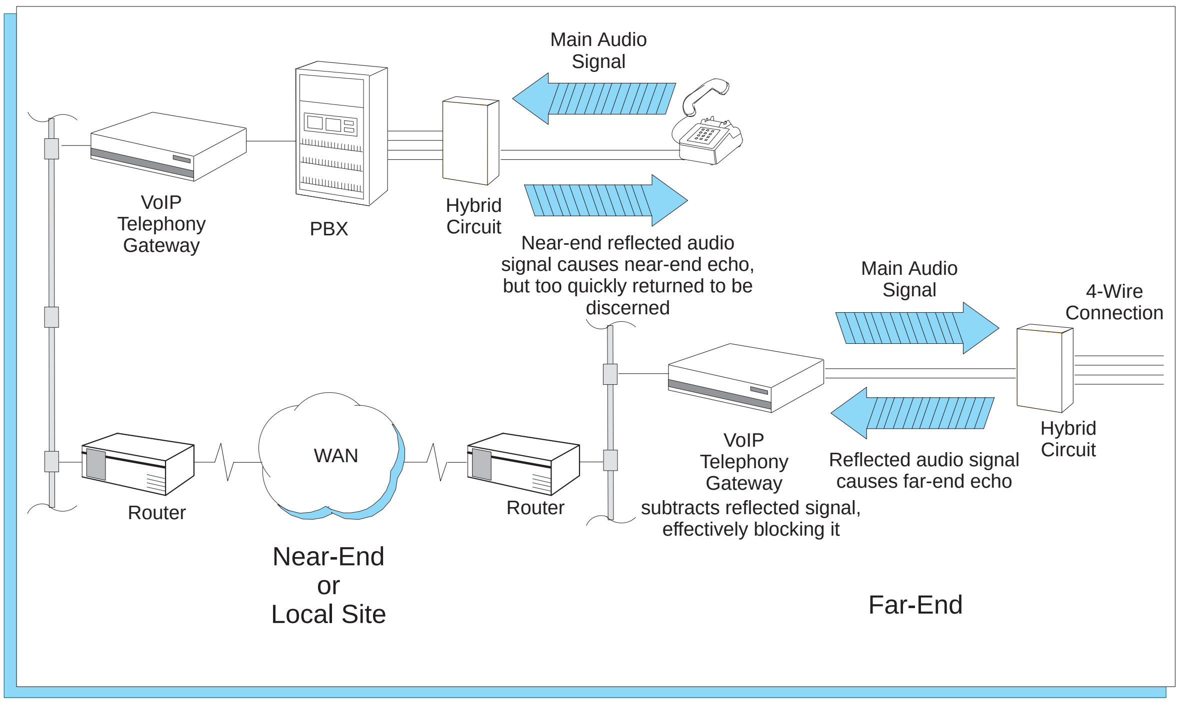

Figure 15 - VoIP Echo Cancellation

To allow for variable packet arrival time and still produce a steady outgoing stream of speech, the far-end gateway does not play out the speech as soon as the first packet arrives. Instead, it holds it for a certain time in part of its memory called the jitter buffer, and then plays it out. The amount of this hold time is the measure of the jitter buffer, e.g., a 50-ms hold time implies a 50-ms jitter buffer.

The jitter buffer hold time adds to the overall delay, so if the network has high jitter, the overall effect will be a long perceived delay in the voice stream. For example, a network might have a moderately average delay of 50 ms and a variability of 5 ms. The network is said to have 5 ms of jitter, a low figure. The jitter buffer hold time is only 5 ms, so the VoIP effective network total delay will only be 55 ms, still moderate.

On the other hand, if a network has a low average delay of 15 ms, but 10% of the time the delay goes out to a long 100 ms, while 90% of the time the delay is a brief 4 ms, the jitter buffer would have to be 100 ms and the total VoIP effective network delay would be 115 ms, a long delay. Network jitter can be more important than average delay in many VoIP applications.

VoIP Packet Prioritization

The reason VoIP works well over a corporate IP network is due more to the corporate network's low VoIP packet jitter than to its low average network delay. Corporate voice packet jitter is low because corporate routers prioritize voice IP packets. The corporate router is instructed to look for voice IP packets and put them ahead of any data packets waiting in the router's transmit queue. This way, a string of outgoing data packets will not add to the variability of the arrival time of voice packets (see Figure 16). Voice packet prioritization is especially important at WAN access speeds of 56/64Kbps—512Kbps. At T1/E1 speeds, it may not be required.

The router is instructed to prioritize voice/fax IP packets, either by the network administrator explicitly programming the router to look for the gateway's "well known UDP port number" (this is a reserved port number registered by the gateway manufacturer for its exclusive use worldwide), or by using a prioritization protocol called RSVP. RSVP is a new prioritization standard that certain router vendors are now including in the operating software for some of their routers. The gateway must also run RSVP.

When the gateway determines it needs to place or receive a voice/fax call, it establishes an RSVP session with the router, using the LAN to pass information. The gateway instructs the router to prioritize the voice/fax packets for the duration of the call. RSVP also includes provisions for constraining packet delay and guaranteeing bandwidth availability, but on a managed corporate IP network only the prioritization feature needs to be used by the gateway.

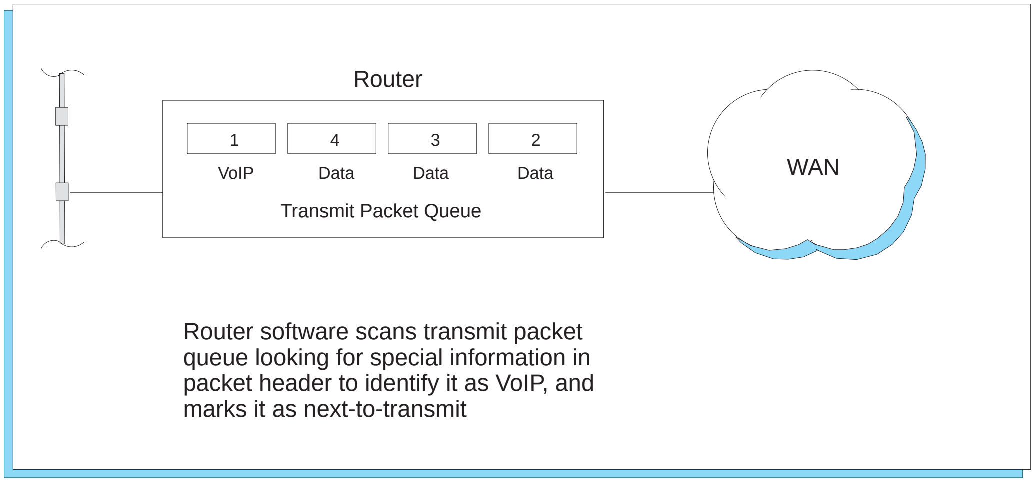

Figure 16 - VoIP Packet Prioritization in Routers

The prioritization of VoIP packet ahead of data packets will cause the data packets to be delayed in WAN transmission wherever there are active conversations on gateway trunks. The effects on LAN data packets, however, are negligible, because the bandwidth per active trunk is so low compared to LAN speeds (15Kbps vs 10–100Mbps) and because there is no VoIP packet prioritization on the LAN.

The delay effects on WAN data traffic vary according to the WAN bandwidth and the data type. For low speed WAN access links (28.8Kbps–256/384Kbps) where the data is usually nonreal-time, such as e-mail and file transfer, data delay goes unnoticed since these are background applications. Who would notice, for example, an e-mail delay of 1 second? For higher speed links (256/384Kbps–T1/E1), where the data may be real-time, such as record-level file access across the WAN, the amount of voice traffic (15K peak per trunk when active) is negligible compared to the WAN bandwidth — only a few percent per trunk at most.

In fact, the marriage of voice/fax and data over WAN links has a mutually advantageous relationship. For low-speed links, data is background and voice/fax is real-time. For high-speed links, data is a mixture of real-time and background, and voice/fax is real-time; but there is no material real-time data vs. real-time voice/fax conflict because voice/fax consumes negligible bandwidth. Prioritizing it first does not materially slow down any real-time data.

This way of combining voice/fax and data has changed considerably from the past. Before the advent of DSP and packet technologies, voice/fax typically consumed a constant 64Kbps without the possibility of silence suppression or the 25% per day duty-cycle factor. Data consumed a small amount of bandwidth and was added onto voice/fax networks. Today, voice/fax consumes 1–2Kbps, around 2% of 64K (a 50-fold improvement). And data bandwidth requirements have grown tremendously. So the old model — voice was the dog, data was the tail — has now switched. Voice is the tail and data is the dog, and voice/fax is being added onto data networks.

IP Packet Segmentation

Another important VoIP delay-management step is to ensure that a very long data packet does not delay the voice packet from exiting the router in a timely manner. This is achieved by programming the router to segment all outbound data packets according to the speed of the WAN access link. Table 7 shows the packet sizes for various WAN speeds. The combination of voice/fax prioritization and packet segmentation produces a VoIP-friendly IP network. Such corporate IP networks are called “managed” IP networks.

Table 7. Maximum IP Packet Sizes

WAN Access Speed (Kbps)	Maximum WAN IP Packet Size (bytes)
56/64	256
128	512
192	768
256	1024
384	1536
512	2048*
1544	6144*

* Ethernet packets do not exceed 1536 bytes. In an Ethernet LAN environment, packet segmentation is not needed above WAN access speeds of 256Kbps.

The router is instructed to segment voice/fax IP packets, either by the network administrator explicitly programming the router to segment all packets, both voice/fax and data, or by using RSVP. With RSVP, when the gateway determines it needs to place or receive a voice/fax call, it establishes an RSVP session with the router. The gateway instructs the router to segment both the voice/fax and data packets, but only for the duration of the call. Using RSVP is preferred because packet segmentation can add 5–10% extra overhead to the data stream, reducing bandwidth efficiency by a similar amount. RSVP improves network efficiency when there are no active calls.

VoIP Forward Error Correction

Another technology used by some gateways to ensure good voice quality is forward error correction (FEC). During IP packet transmission across an IP network, the IP packet can become corrupted or discarded. If this happens only occasionally, the IP packet replay technique discussed earlier usually suffices to maintain good voice quality. Almost all well managed corporate IP networks have very low packet corruption and loss, requiring only packet replay.

The public Internet, however, can have substantial packet corruption and loss, and packet replay is usually inadequate to maintain good voice quality (not withstanding the degradation in voice quality that results from long and unpredictable Internet delays). For voice over the Internet, FEC can compensate for the corrupted or missing packet.

FEC can operate at two levels, intra-packet and extra-packet. Intra-packet, FEC adds extra bits that allow the receiving gateway to determine which of the packet's bits were corrupted and to restore the bits to their proper setting, yielding a packet ready for playout. Extra-packet, FEC adds extra information to each packet that allows the receiving gateway to extrapolate from the previously received good packet and reconstruct the missing or severely corrupted packet.

With FEC, packet loss of 10–20% can be absorbed while still producing acceptable voice quality. However, because both techniques consume extra bandwidth, sometimes as much as an additional 30%, FEC is not usually used in corporate VoIP.

VoIP WAN BANDWIDTH COMSUMPTION

The second key to VoIP's success in business applications is minimizing VoIP WAN bandwidth consumption. Beyond speech compression and fax demodulation, the best bandwidth reducing technology is "silence suppression." Silence suppression technology recognizes the periods of silence in a conversation or fax transmission, and stops sending IP speech packets during those periods.

Telco studies show that in a typical phone conversation, only about 36–40% of a full-duplex conversation is active. When one person talks, the other listens (this is called half-duplex). And there are significant periods of silence during speaker pauses between words and phrases. By applying silence suppression, FDX bandwidth consumption is reduced by the same amount, freeing up bandwidth for other voice/fax or data communications.

Figure 17 illustrates how silence suppression allows two conversations to fit in the bandwidth otherwise used by one. This 60% bandwidth reduction develops over a 20–30 second period as the conversation switches from one direction to another. The results are significant, as Table 8 summarizes for the G.729 8Kbps case. These figures also apply to fax transmissions because fax is also half-duplex and has periods of silence during pauses as well.

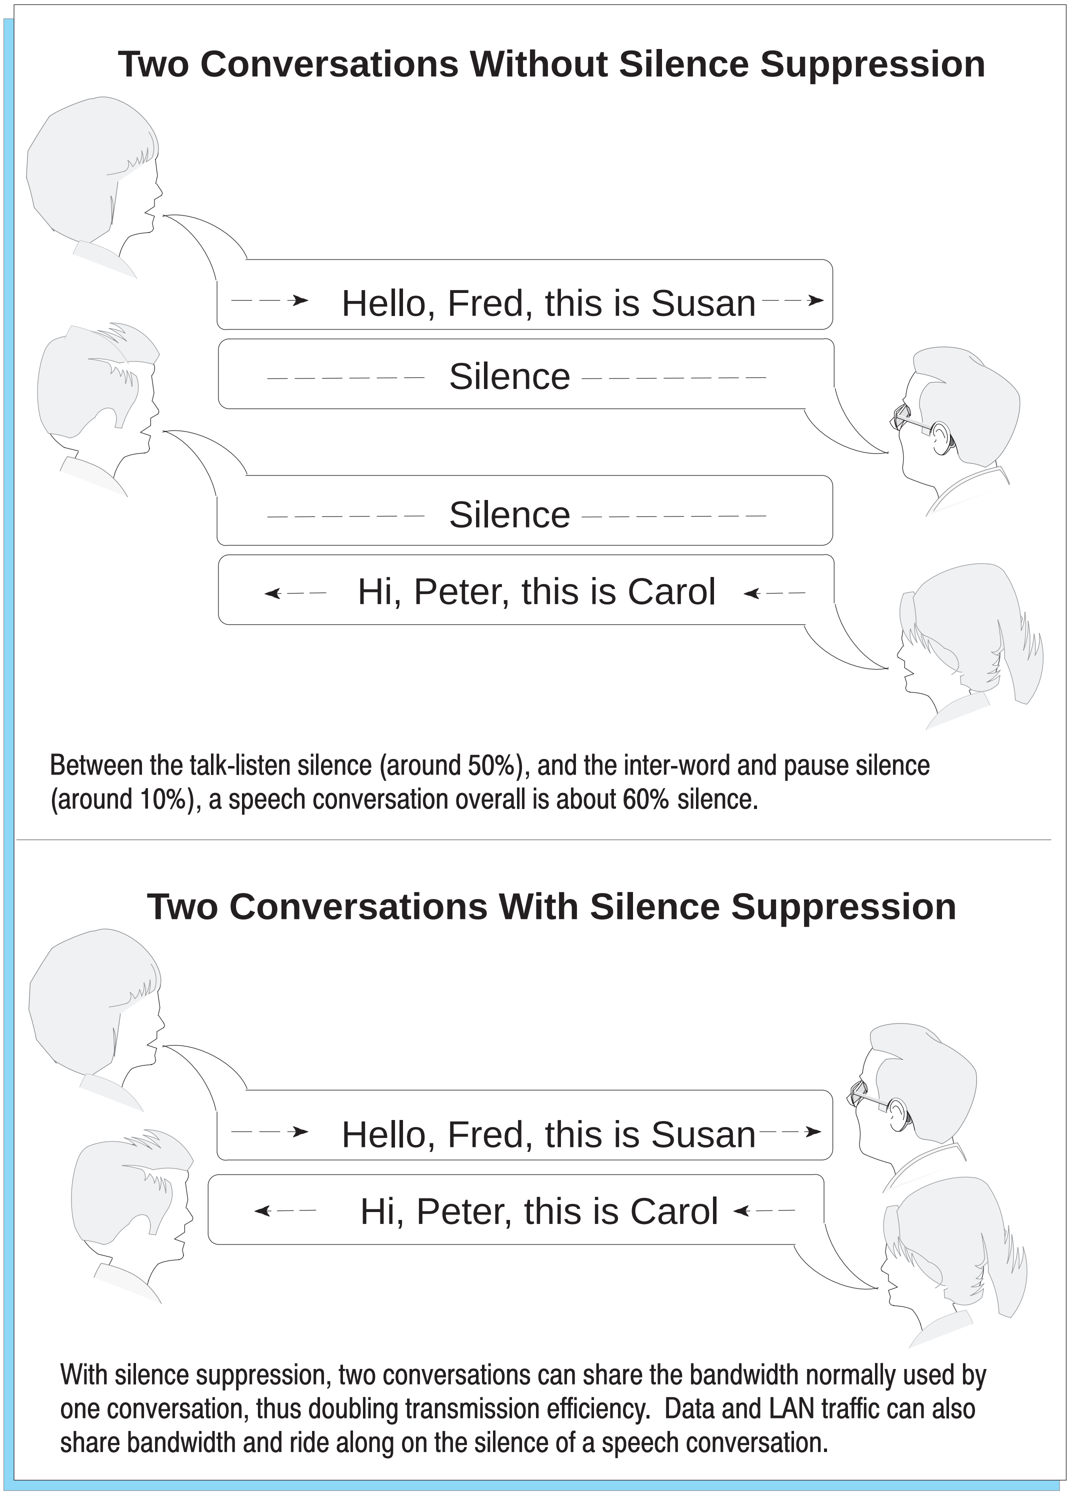

Figure 17 - Conversations with and without Silence Suppression

Table 8. Effects of Silence Suppression on Average Bandwidth Consumption

G.729 CODEC Bandwidth	8Kbps
IP Router Overhead	2–7Kbps*
Total Bandwidth	10–15Kbps
Less 60% Silence	— 6–9Kbps
Net CODEC Bandwidth consumption averaged over a 20–30 second period	4–6Kbps

* 2K with router header compression, 7K without

This means that an active VoIP Gateway trunk averages 4–6Kbps WAN bandwidth consumption. There are, however, significant periods when the gateway trunk is inactive. In most applications, a remote-site gateway trunk is active less than 25%, or inactive over 75% of the time. The periods of inactivity and activity tend to even out over a 20- to 30-minute period. So, the net average WAN bandwidth consumption per gateway trunk over a 20- to 30-minute period is about 25% x (4–6Kbps), or 1–1.5Kbps. This is about 1.7–2.7% of a 56/64K WAN access link. Figure 18 depicts this bandwidth-reduction sequence.

With this low average bandwidth consumption figure, it is tempting to think that 20, 30 even 40 voice conversations could be carried over a 56/64K link. But the fact is that only 3–4 G.729 conversations will fit on a 56/64K link. This is because all speakers may for a moment speak in the same direction at the same time on all four trunks, and this would momentarily consume 40–60K, leaving no bandwidth to accommodate a fifth, sixth, etc. speaker speaking in the same direction at that moment.

So, even though four trunks would average about 7–11Kbps, or 11–19% of a 56/64K link, the remaining 45–57Kbps of average bandwidth availability cannot be selectively used for a real-time application such as voice or fax, only for background data transmissions such as e-mail or file transfer. Table 9 shows the VoIP trunk-carrying capacity of different WAN access bandwidths.

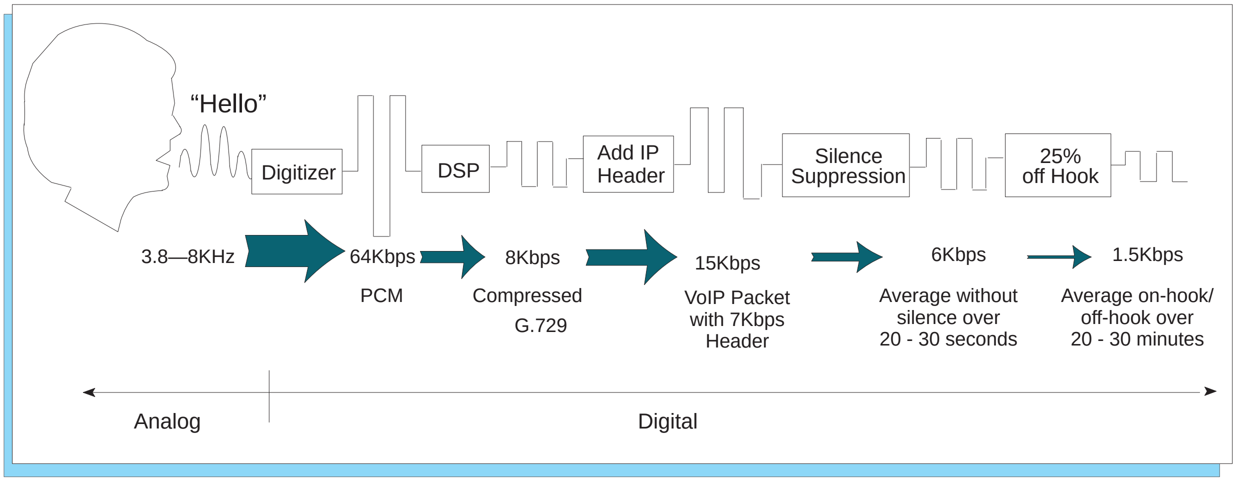

Figure 18 - VoIP Bandwidth Reduction Sequence

Table 9. WAN Access Link VoIP Trunk Carrying Capacity

With Router Header-Compression			Without Router Header-Compression	
WAN Access Bandwidth (Kbps)	**# Gateway Trunks**	**Residual Bandwidth (Kbps)***	**# Gateway Trunks**	**Residual Bandwidth (Kbps)***
28.8	2	26.8	1	27.3
33.6	2	31.6	2	30.6
56	5	51	3	51.5
64	6	58	4	58
128	12	116	8	116
256	25	231	17	231
384	38	346	25	346
512	51	461	34	461
1024	102	922	68	922
T1—1536	153	1383	102	1383
E1—1920	192	1728	128	1728

* Residual bandwidth averaged over a 20—30 minute period.

VoIP and Real-Time Data Applications

If there are real-time data applications running over a low speed WAN link, such as SNA, subtract the real-time data's bandwidth from the link's bandwidth to determine how much residual bandwidth is available for voice and fax. For example, if a 19.2Kbps SNA data stream is running over a 56Kbps link, the residual 36.8Kbps could be used for VoIP. This would allow two to three gateway trunks, depending on header compression. With two trunks peaking at 30Kbps (with no header compression) and averaging 3Kbps over a 20- to 30-minute period, an average of 33.8Kbps bandwidth would be left for nonreal-time data.

Moreover, because of gaps in the SNA data stream — usually about 50%, or 9.6Kbps in this case — another 9.6Kbps would be available for background data communications. The voice/fax would average 3Kbps, SNA would average 9.6Kbps, and a total of 43.4Kbps would be available for background data communications.

While saving substantial bandwidth, silence suppression is not without a couple of potential downsides. The first is the potential for "first-word clipping". This occurs when the speaker begins talking, and happens because the silence-suppression technology does not recognize quickly enough the suddenly increasing audio energy and misses the first part of the first word. This can be avoided by using advanced silence suppression technology.

The second problem can occur if the technology does not include a provision for background noise regeneration. Silence suppression renders the line absolutely silent to the listener, so much so that the line sounds dead. But, by inserting "comfort noise" or, even better, by periodically sampling the true far-end background noise and regenerating it for the listener, the line sounds active. The bandwidth savings of silence suppression are retained, and the listener needs are satisfied too.

VoIP COST SAVINGS

Attractive cost savings is the third key to VoIP's success in business applications. The cost savings can come from several areas:

- reduced intracompany inter-remote-site voice/fax communications costs
- reduced selected PSTN communications costs by using off-net calling
- reduced selected 800 number calling costs by making local numbers that ride over the corporate IP network available to outside callers (somewhat the reverse of off-net calling)
- reduced costs from reducing the number of telco Central Office (CO) phone lines in proportion to the offloading of calling volume to the VoIP network.

Cost savings are calculated by taking the total toll and 800 charges and CO-line costs that are expected to be avoided over a certain period and subtracting the total costs of all VoIP equipment and related costs (installation, service agreements and costs for any increased bandwidth). The period for calculating the available costs is usually five years. The VoIP solution is usually easily justified if the net savings are 80% or more of the avoided costs, and only justified with difficulty if they are less than 50–60%.

Another way to look at VoIP savings is by means of the payback period or break-even point, usually measured in months, when the avoided costs will have paid for the VoIP equipment and related costs. This calculation usually plays a role in the cost-justification process. Twelve-month or less payback periods are usually easy to justify, 24-month and longer are substantially more difficult, and 12-to-24 month periods fall in between.

A third way to view the VoIP savings is by comparing the net effective toll rate the VoIP system delivers over the life of the system, again usually five years, with the anticipated telco toll rates. The net effective toll rate is the total number of minutes of all types of cost-avoided voice/fax calls over the period, divided into the total cost over the period for all VoIP equipment and related costs.

Net-effective VoIP toll rates in the one- to three-cents per minute range are usually much better than telco rates even in highly competitive markets (e.g., U.S. and U.K.), while those in the 6- to 7-cent range are considerably less favorable. On the other hand, net-effective rates in the 10- to 12-cents/minute range compare very favorably in certain markets like Japan, Germany, Argentina, Russia and India, and in most international calling markets. And, of course, in these markets rates in the 6- to 7-cents range are much lower than telco rates.

Company Savings Example

To illustrate these calculations, consider a sample case study of the Ajax Company. North American telco conditions apply. They have an existing IP network with 10 remote sites using a mixture of leased lines and frame relay (see Figure 19). Ajax has recently estimated that they currently spend about $25,400 per year on intra-company inter-site voice and fax. This is about 2 hours per day to/from each site, at 8 cents/minute, using 22 business days/month.

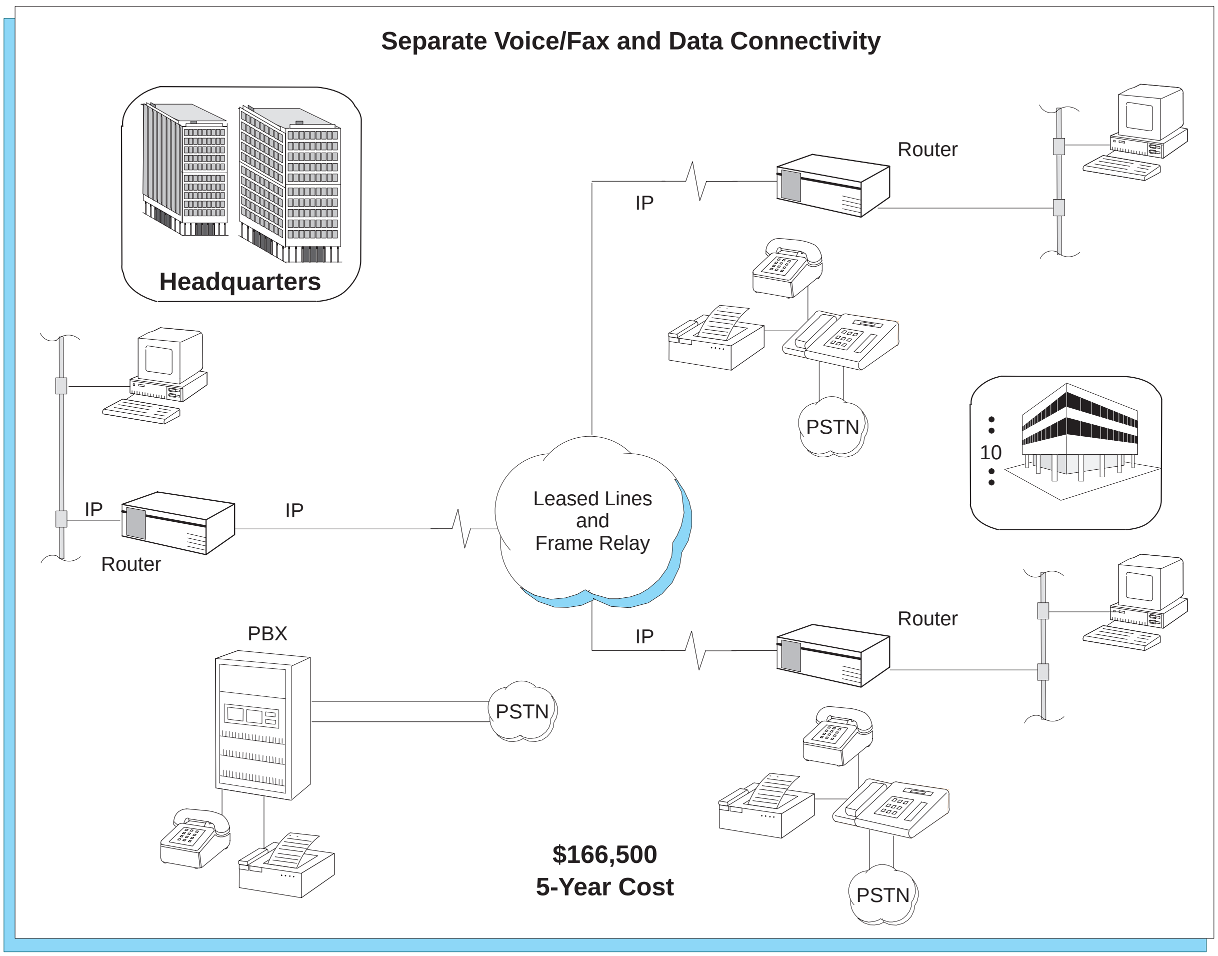

Figure 19 - VoIP Cost Savings Example

They also estimate they are spending about 20 minutes per day making calls into the general metro area of each remote site, also at 8 cents/minute. This amounts to about $4,200 per year. And they believe that by using VoIP, they can reduce the CO-line count by one line at each remote site, and by eight lines at headquarters. This amounts to about $3,000 per year for the remote sites, but only about $800 at headquarters, because they have a T1 access line to the CO.

The total estimated telco-related cost savings adds up to $33,300 per year, or $166,500 over a 5-year period. This is about $16,650 per remote site. The total number of on-net and off-net minutes is 1,848,000. Cost per minute, including CO line charges, is 9 cents.

The VoIP gateway solution calls for a 2-trunk gateway at each remote site and a 10-trunk gateway at headquarters. This provides for a 2-to-1 contention ratio. The headquarters gateway will use a T1 interface to the headquarters PBX. This allows a growth potential of 14 more T1 channels, enough for 20 more remote sites. Each remote gateway will connect to a KTS or PBX using analog interfaces. Figure 20 illustrates this VoIP solution.

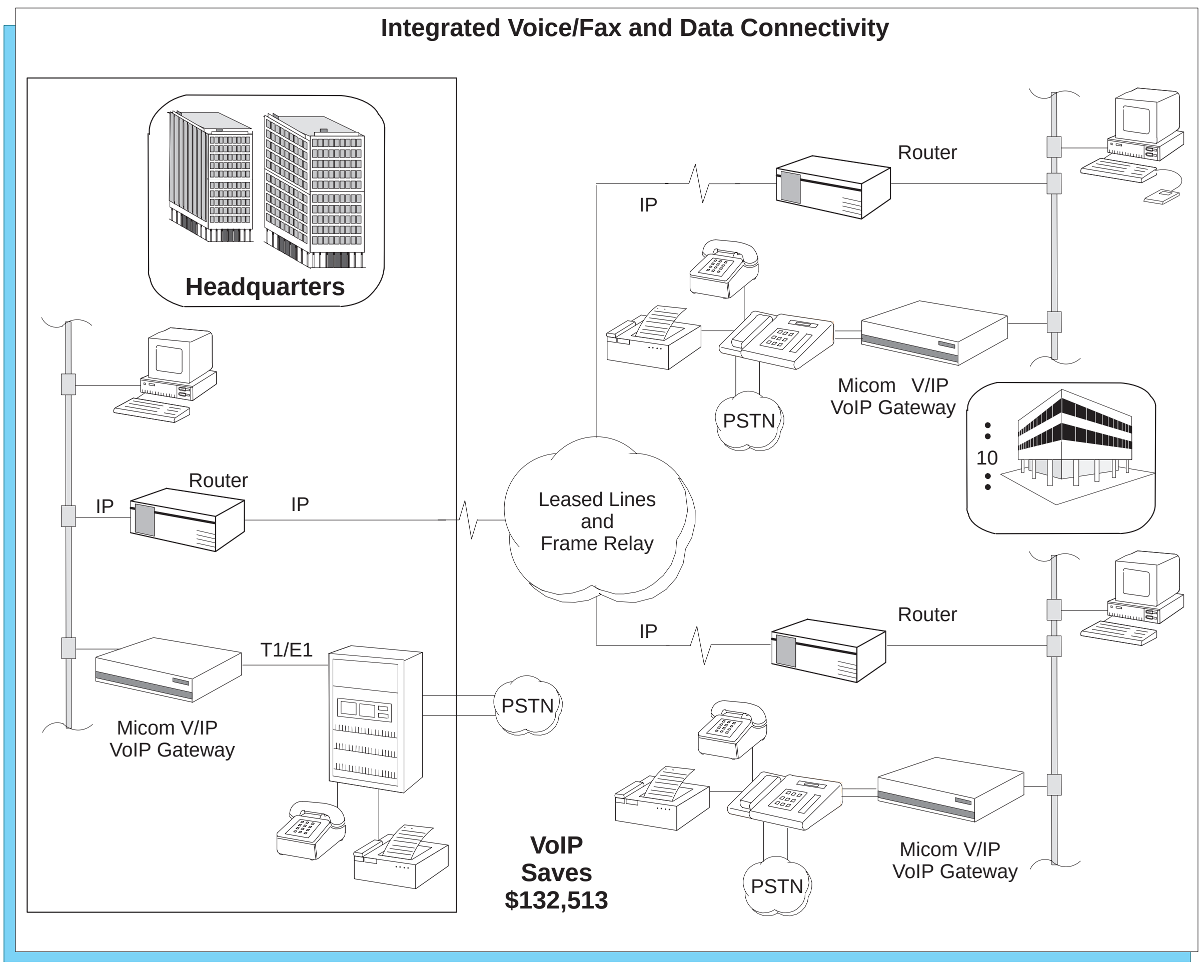

Figure 20 - VoIP Cost Savings Example

By using Micom's V/IP VoIP product family, Ajax will spend $25,985 in VoIP and related costs. Table 10 shows the cost breakdown. Because the V/IP solution uses G.729 and silence suppression, the bandwidth consumption per remote site for voice and fax will average only about 3Kbps. This is low enough to avoid having to add remote-site bandwidth.

Table 10. Micom V/IP VoIP Gateway Solution

Description	List Price	Quantity	Extension
Micom 2-Trunk Analog V/IP Module	$1,750	10	$17,500
Micom T1 Card With 4 Channels	$2,995	1	$2,995
Micom T1 Expansion Card With 4 Channels	$2,495	1	$2,495
Micom T1 Expansion Card With 2 Channels	$1,495	1	$1,495
PC for Housing T1 Card Set	$1,500	1	$1,500
		Total VoIP	$25,985

Because the Micom V/IP cards can run in a shared PC, there is no need to include the cost of a PC at each remote site. The V/IP cards' DSPs do all the processing, leaving the PC to do other work. In fact, V/IP will even run on a 486 SX20 with just 8MB of memory, by itself a zero cost chassis in many circumstances. At headquarters, to better secure the integrity of the environment for the 10 headquarters trunks, Ajax will use a dedicated PC, configured to support Windows 95 or NT.

The installation costs charged by the Micom reseller for the V/IP system are $300 for each remote site, and $1000 at headquarters, totaling $4000. The maintenance and service agreement from the reseller is $2000/year for five years, or $10,000. Total VoIP costs are $39,985. Because of the competitive bidding situation, the reseller has discounted the package 15%, and will sell the system to Ajax for $33,987, effectively doing the installation for free and throwing in the first year's service at no charge.

VoIP Payback Period

The VoIP payback period for Ajax's network is calculated to be $33,987/ $33,300 x 12 months, or about 12 months, 1 week. The five-year savings is $166,500 - $33,987 = $132,513, or $13,251 per remote site. The net effective toll rate is $33,987/1,848,000 minutes, or 1.84 cents/minute. The net reduction in telco-related charges was $132,513/$166,500 = 80%.

Figure 21 illustrates some of these calculations in a graphical format. The upward sloping line represents Ajax's cost of telephony business as usual. The horizontal line represents the cost of the Micom V/IP solution, paid for at year zero. The intersection is the break-even point, and the shaded area between the lines represents the five-year savings. The savings, of course, don't stop at five years, they continue to grow at the rate of $33,300 per year.

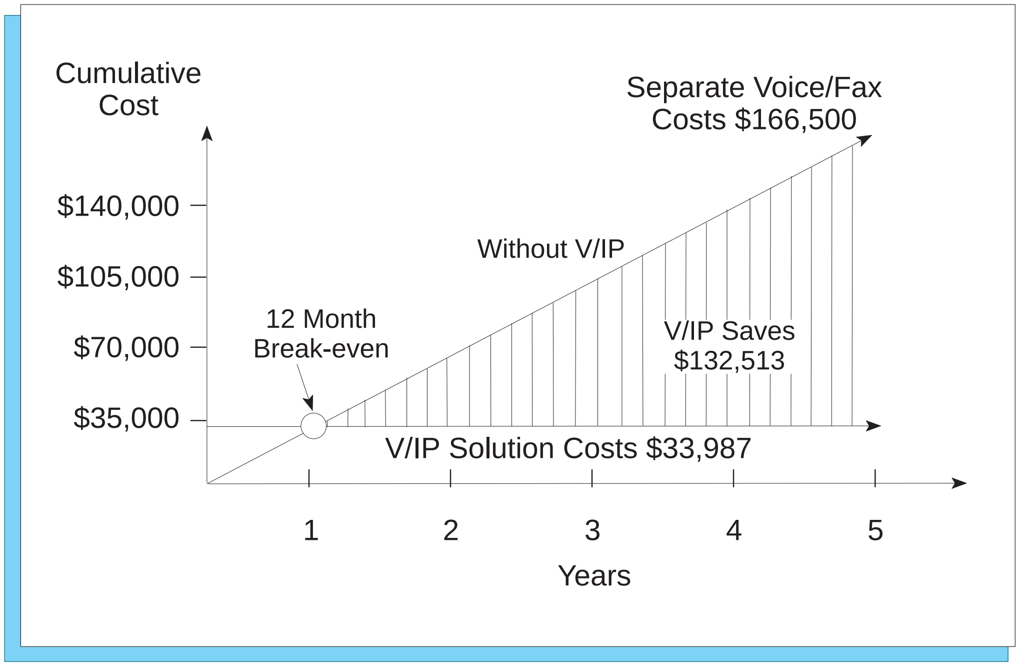

Figure 21 - Cost Savings Illustrated Graphically

Table 11 shows the five-year savings and payback periods for the Ajax VoIP solution corresponding to different pre-VoIP net-effective telephony per minute costs. The figures shown also apply for smaller and larger networks, which also have similar payback periods and proportional savings for a given per-minute cost.

Table 11. Five-Year Savings and Payback Periods

	Per-Minute Costs (U.S. cents/minute)								
	6	**9**	**12**	**15**	**20**	**25**	**30**	**50**	**80**
Savings ($U.S. 000)	77	133	188	243	336	429	522	891	1446
Payback (months)	18	12	9	7	6	4	4	2	1

VoIP OVER THE INTERNET

The internet is by far the best known IP network. Indeed, its name comes from IP, Internet Protocol. It is ubiquitous and inexpensive, so the obvious question to ask is, can a corporation use the Internet for VoIP? The answer is, generally no but sometimes yes. It depends on several factors.

As we discussed earlier, corporate VoIP voice quality must be near toll quality, simple to use and highly available, or the employee will avoid using the VoIP network and defeat any cost savings. VoIP in the consumer arena is entirely different. With personal VoIP telephony, the consumer is saving money directly into his or her pocket, and is very willing to trade off voice quality and availability for cost.

In fact, the greater the savings, the more tolerant the consumer is of reduced quality. The reduced voice quality is primarily caused by the Internet's long delay and high jitter, not so much by the actual CODEC quality. To overcome the Internet delay problems, the consumer can learn to talk "CB" style, similar to talking over a half-duplex radio.

For smaller businesses, where the individual employee may actually see some tangible benefit to reducing telephony costs, Internet VoIP may succeed. This is especially true if the company has inordinately high telephony costs, as in some international markets. And for nonreal-time voice, such as voice mail or IVR (Interactive Voice Response), Internet VoIP can work quite well in the business environment.

As we discussed earlier, the business community can opt to connect remote sites together with IP-VPNs (or "Extranets") delivering good enough QoS and low enough delay to carry VoIP in the business environment. IP-VPNs cost about the same as frame relay networking solutions, but offer the added benefit of Internet access on top of good quality corporate data connectivity. If the Internet is thought of as including IP-VPNs, then VoIP over this Internet should work well for corporate applications.

The PC as a Telephone

VoIP over the Internet is often associated with using the PC as a telephone. In the corporate setting, using the PC as a telephone is only in the experimental stages, especially when the PC replaces a desktop telephone. Going forward, however, the use of the desktop PC as a telephone will become more popular. Table 12 compares some of the characteristics of VoIP using PC phones, and VoIP using the IP gateway with legacy telephony equipment.

Table 12. Comparing PC Phones and VoIP Gateways

PC Phone	VoIP Gateway
Replaces paradigm of centrally wired phones	Retains paradigm of centrally wired phones
Requires user retraining on new phoning techniques. Introduces a new calling paradigm	Retains all legacy phone/PBX/KTS functions. Operation is transparent to user
Can offer feature-rich telephony functions, conferencing and collaborative consulting	Does not extend telephony functionality beyond existing PBX/KTS capability
Requires a VoIP Gateway to call legacy phones or make PSTN calls	Consumes LAN bandwidth only for intracompany inter-remote-site calls
Consumes LAN bandwidth for every call: in-building, PSTN, intracompany and long distance	No gateway required for ordinary telephony
Requires the installation and maintenance of hardware and software in each PC at a site	Requires the installation and maintenance of hardware and software in one PC per site
Does not accommodate fax machines	Accommodates fax machines as well as telephones

For PC phones to talk to ordinary telephones and the PSTN, the PC phone must use a VoIP gateway. Figure 22 illustrates this scenario. The PC phones can communicate with each other without a gateway, but must use a gateway to convert the IP voice packets to standard telephony signals for communicating with outside phones or making an outside call.

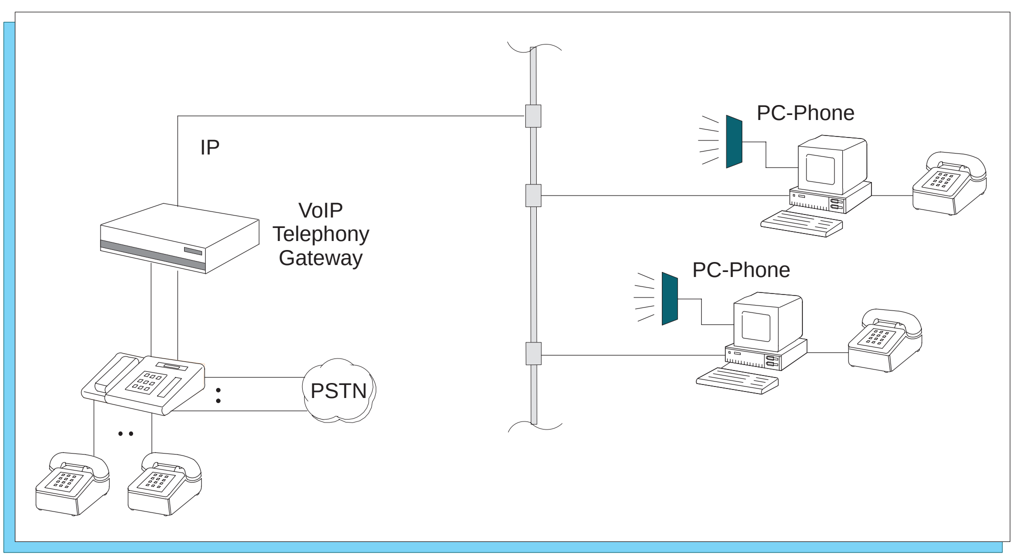

Figure 22 - PC Phone using VoIP Gateway

VoIP Standards Forum

In order for the VoIP Gateway to be able to convert the PC phone's IP voice packets into telephony signals, and vice versa for FDX communications, they must both use the same speech compression CODEC as well as, for example, the same call connection, call supervision and device addressing protocols. To achieve this level of interoperability between the various brands of PC phones and VoIP Gateways, the VoIP industry has established a standards-setting group called the "VoIP Forum," part of the IMTC (International Multimedia Teleconferencing Consortium).

Micom and Nortel both participate in the VoIP Forum, and are committed to implementing the VoIP interoperability standards as they are made available. Among the relevant standards Micom is implementing are the G.723 voice compression CODEC and the H.323 conferencing protocols. Micom's V/IP VoIP products will be enhanced to interoperate with PC phones using these standards during 1998. Starting in 1997, Micom began testing Microsoft's NetMeeting 2.0 release as a PC phone for interoperability with V/IP using G.723 and H.323, and has demonstrated NetMeeting PC phone-to-V/IP connectivity at various trade shows.

The practical application of PC phone-to-gateway interoperability in the corporate VoIP environment is the enhanced connectivity this extends to the telecommuter, road warrior (hotel traveler) and SOHO (Small-Office-Home-Office) users. Figure 23 illustrates a telecommuter PC running MS NetMeeting with an IP connection over the PSTN or the Internet back into the corporate headquarters' access hub or router. The PC user has dialed up over a single connection and is passing e-mail, files, and other data.

In order to talk to someone at an ordinary telephone in the company (at headquarters or a remote site), the user would have to make a separate phone call over the PSTN, at an additional cost. With IP telephony using H.323, however, the user can now simultaneously connect to the VoIP gateway at the desired location and call the intended legacy telephone. The call shares the same single dial-up connection with the data, both transmitting as IP packets.

The data continues passing to/from its original destination(s), not the gateway. This line-doubling advantage of VoIP reduces corporate telephony costs, and also helps the road warrior in the hotel room with only one phone line, unable to place a second call.

With the right gateway enhancement, the PC phone can also receive phone calls via the VoIP gateway, again while passing data over a single connection. The VoIP Gateway not only reduces corporate telephony costs, it also serves as a legacy telephony and PSTN gateway for PC phones, and multiplies the line carrying capacity of dial-up connections. The voice quality of the PC phone call should be near toll quality, unless the connection uses the Internet.

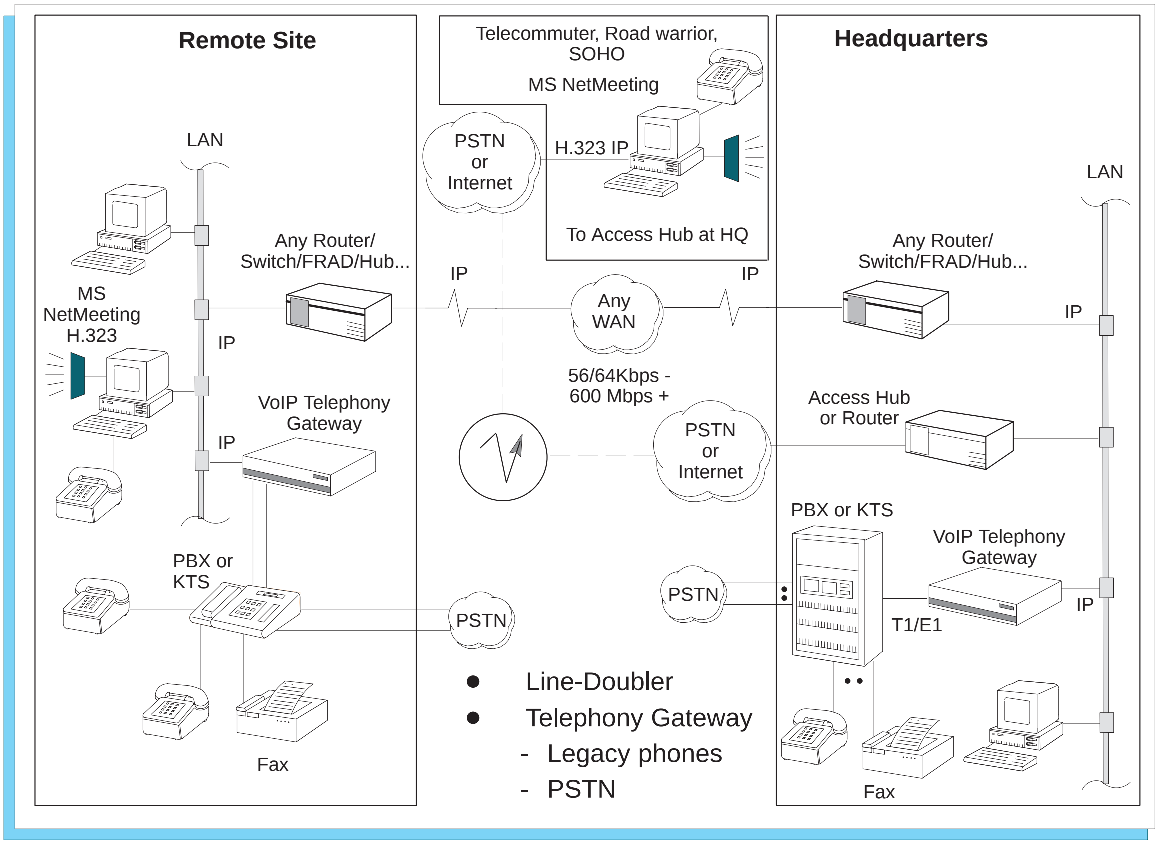

Figure 23 - PC Phone to Gateway Interoperability

SUMMARY

Voice/fax over IP has several different applications. At the present the foremost of these is sending voice and fax over corporate IP WANs, using VoIP gateways connected to LANs and legacy PBX, KTS, phone and fax machines. Voice quality is at or near toll-quality, there is little or no user retraining, and the voice traffic consumes little WAN bandwidth. Other LAN/WAN data applications are essentially unaffected by adding voice/fax to the network. Corporate IP WANs can use leased lines, public frame relay, satellite, ATM, ISDN, and the new IP-VPN services.

The savings in long distance intracompany, intersite toll calls can be large, with rapid payback periods. Savings are augmented by taking advantage of off-net calling. Net-effective toll rates are in the 1- to 2-cent/minute range. The prospect of using the Internet for corporate voice is enticing, but the required voice quality is not yet there. Internet delay and jitter are too large. Using the PC as a telephone is also interesting, and may be utilized in the corporate environment initially for the SOHO, telecommuter, and road-warrior applications. A VoIP gateway is required for the PC phone to call an ordinary telephone.

MICOM Communications Corp.
A Nortel (Northern Telecom) Company
4100 Los Angeles Avenue
Simi Valley, California 93063-3397
Tel: (805) 583-8600
Fax: (805) 583-1997

Internet Web page: http//www.micom.com
E-mail address: info@micom.com
24-Hour Fax-on-Demand System: (805) 582-7478

For more information or the location of the
Micom Certified Distributor closest to you,
call our Applications Hotline.

In the United States call:
Tel: 800-642-6687 • Fax (800) 343-0329

In Europe, Middle East and Africa call:
Tel: 44 (0) 1252 302 200 • Fax: 44 (0) 1252 781 717

In Asia-Pacific Rim call:
Tel: 65 334 6557 • Fax: 65 334 5845

In Latin America call:
Tel: (954) 858-7200 • Fax: (954) 851-8300

850-0330B