Unified Communications

POCKET GUIDE

www.shoretel.com

Unified Communications

POCKET GUIDE

Author: Barry Castle

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Internet protocol telephony (IPT) provides the foundation for what, without question, will become a major driver for enterprise productivity improvement – Unified Communications (UC). The basic premise of UC is that if a means of communication is available to two or more parties, then they should be able use it intuitively.

This list illustrates the concept:

1. Video: Children use video intuitively to communicate with each other. Yet expensive, enterprise-class phone systems either ignore this technology or make it so cumbersome to use that nobody bothers. With UC, users can switch to video at the click of a button.
2. Presence: Who is available right now to help me get the answer my customer needs?
3. Instant Messaging (IM): I rang John, but he's in a meeting. Is there any way to get real-time answers from him? By switching to IM, I receive the answers without interrupting his meeting.
4. Video Again: We need to quickly evaluate the extent of damage caused by a recent storm. Our agent switches the call to video, so we can all see what she sees, in real time.
5. Productivity: I want to know when members of the key account sales team come out of their meeting, so I can debrief them and send the customer information she has requested, without waiting until everyone returns to their offices.

And that's just a fraction of what unified communication is expected to deliver. Every organization will find new ways to leverage a single, coherent communication system for their own needs.

However, there is a catch: It is true that some of the systems on the market today were designed from the ground up to deliver the productivity gains we expect, which makes things considerably easier to put in place. However, other approaches have appeared through vendor acquisition. These require a lot more work to deliver a meaningful benefit to end users.

This means that evaluating the different architectures on offer, making a selection and then successfully delivering on the promise of UC requires a basic understanding of at least five technologies:

1. Voice/telephony systems and services
2. Data communication networks
3. IT systems
4. Mobile telephony
5. Video conferencing technology

ShoreTel recognizes the challenges of defining a communication strategy that spans multiple domains and this guide is intended to help you with this task. Following the introduction is a review of the typical components and features of an enterprise telephony system. Then, an introduction to voice over IP (VoIP), the underlying network infrastructure, and various issues take must be taken into account when designing the network to carry voice. Next is a description of advanced applications such as collaboration, presence, customer relationship management, unified messaging, Fixed Mobile Convergence and Video. The conclusion compares these VoIP solutions to traditional centralized private branch exchange (PBX), and provides examples of what you can expect from a well-designed UC implementation.

1. INTRODUCTION

This guide is intended equally for technical IT staff, voice system managers and CIOs. It describes the issues decision makers need to understand as they set out to build a winning IP business communication strategy. Key discussion areas include:

- Breakdown of the critical components of an enterprise voice system
- VoIP technologies and standards
- Underlying network elements that can affect the voice system
- Business applications such as unified messaging, converged conferencing, call centers and customer resource management (CRM)
- Emerging technologies, including fixed mobile convergence (FMC), presence, IM, video – the components of UC
- Advantages and disadvantages of various architectures
- Review of total cost of ownership (TCO)
- Glossary of terms

IP data communications is already the global standard, and the transition to a pure IP environment has important implications for IT organizations. There are many reasons to implement an IP-based voice communication system: reduced long-distance telephony charges; lower capital costs; decreased management and administrative costs; reduced complexity; improved integration of distributed business entities; and a greater ease with which voice applications may be combined with other business systems.

But for many decision makers, a key driver is the opportunity to gain competitive advantage by deploying these applications. For example, by improving the quality and value of integrated voice and data communication, the promise is that businesses can more effectively leverage internal business processes, leading to more effectively managed external customer relationships. To define a coherent strategy, business decision makers, IT managers, and communication professionals need a firm grasp of voice and data communication technologies. They need to understand how such technologies and standards support emerging applications to deliver a converged enterprise communication platform.

1.1 EXPECTATIONS AND DESIRED OUTCOME

Many companies have successfully made the jump from legacy systems to IPT – so, what do these companies like most about their IPT implementations? This section summarizes many of IPT's advantages and can serve as a checklist for evaluating vendor offerings.

Distributed Intelligence – By distributing call processing intelligence (the ability to set up and manage calls) across the network, the voice system eliminates single points of failure, including a failure of the IP wide area network (WAN) itself. This is critical to delivering reliable voice calls.

Single Management Interface – The ability to incorporate every element of a multi-site voice system (media gateways, gateway controllers, telephones, productivity applications) into a single homogeneous management system dramatically reduces administrative costs.

Application Rich – A system that delivers a range of customer interaction solutions that can be activated at the click of a button and enables powerful multi-site collaboration creates a better customer experience. Such a system allows your organization to appear more coordinated and more professional, because calls and conferences are seamlessly transferred and shared between team members, sites and mobile employees.

Ease of Use – Today's systems deliver more features while eliminating the guesswork about how to use the phone system. Your employees should have access to the full range of advanced telephony features and internal/ external phone directories, without having to become phone experts. These productivity features should be 100-percent transparent across your enterprise network. When it is time to evaluate different vendor solutions, we recommend testing the applications available on the desktop interface, to ensure they are intuitive and consistent across the full range of analog and IP phones. The more easily your staff can use the phone system, the more productive they'll be.

Outstanding Clarity – Digital phone systems were introduced in the late 1970's and technology has evolved considerably since that time. Rather than copying the technologies of yesterday, today's best systems leverage additional network capacity and offer improved design ergonomics that provide improved sound quality. Your voice is more easily heard because the system delivers the full range of audible tones to the human ear. Because calls sound better, less time is spent trying to communicate and conversations are more productive.

Simple Expansion Capability – Legacy PBX systems can be expensive and complex to grow. Some IP-PBX vendors make matters worse by deploying multiple management interfaces for related data-networking components. That is why it is important to carefully consider how many steps are required to expand a system. Does the vendor solution require a long lead time and expensive, highly trained personnel, or is it so easy to upgrade that an unskilled staff member in your headquarters office can physically connect a voice switch at a remote site and bring it online in a matter of minutes? Speed and flexibility are critical in today's business environment; if your voice system can be adapted to your business imperatives, your new team can get up and running faster. So your company is more competitive.

Smooth Migration Path – The new system should be able to co-exist with legacy systems and applications, as well as provide backward compatibility with legacy trunks, extensions and voicemail. Examine whether the system has a set of interfaces to enable a stepwise migration from your legacy PBX systems to IP voice. A smooth migration path allows you to go live with new locations and teams at will. You drive the project, rather than allowing the technology to drive you.

The transition to IP voice technology offers a rare opportunity to improve business functions within your company. By including the criteria outlined above in your evaluation process, you are more likely to ensure that this opportunity delivers a better voice system that benefits your employees. It is critical that the key telephony stakeholders be involved in the decision making process early on:

- IT and communication team
- CEO/director's office
- Marketing and sales force
- Key administrative personnel
- Customer service

By including each group in the process of selecting a vendor, you can ensure that their requirements are taken into account and the project's goals are tightly aligned with the company's business objectives. The next section provides an introduction to voice telephony. *(Note: This section is intended primarily for readers with limited experience in voice communication. Experts in this area can move on to the third section.)*

2. ENTERPRISE VOICE SYSTEMS

Let's start with an overview of a typical voice system. What are its components and what functions should you expect from such a system?

PBX systems – based on time-division multiplexing (TDM) – were traditionally deployed by large enterprises, until the arrival of professional- grade IP telephony systems in the late 1990s.

Before we drill down into the details, let's review the PBX in general terms. These systems typically include:

Telephone handsets
Cables connecting the telephones
Line interfaces to the phone cables
Switching and call processing to make calls
Trunk interfaces to communicate with the outside world
Management console and ability to track and account for calls
Applications and enhanced services

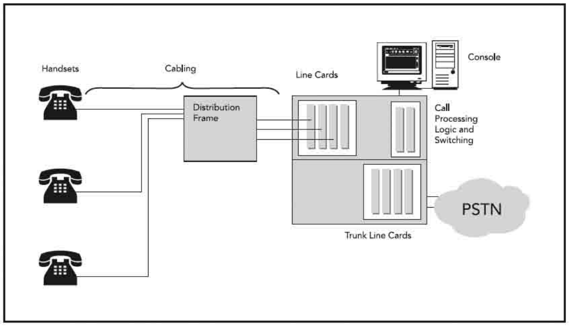

Figure 1: Components of the Legacy PBX

For these components to fulfill their tasks, software and signaling capabilities are also required. The next section explores the functionality of each element and discusses its contribution to the overall solution.

2.1 CALL SWITCHING, PROCESSING AND SIGNALING

To understand what PBX switches do, it helps to travel back in time, to before switches existed. At that time, a switchboard operator was required to set up a call between two phones. A call could take place only after a continuous connection of wire had been established from the calling party to the called party, to form a circuit. The switchboard consisted of a wooden panel with cables and jacks, and an operator connected a cable to the plug of each party in order to set up a call. Things could get fairly involved when setting up a long distance or international call. Operators had to talk to each other as they established a continuous circuit across many of these switchboards.

Just like the switchboard operators of the past, today's PBX switches must remember what everyone is doing at each moment in time, and connect telephone calls between the appropriate places. The switch effectively establishes a circuit between the called parties, and the act of establishing this circuit (i.e., setting up and terminating calls) is referred to as call processing.

Call processing is accomplished using specific signaling protocols between the PBX and attached handsets, adjacent PBXs, and Public Switched Telephone Network (PSTN). In some cases, these protocols tend to be vendor-specific and proprietary, while in other cases, the protocols are based on national or international standards. The list shows the protocols used to communicate between various devices:

- PBX and Analog Handsets – Standard signaling protocols
- PBX and Digital Handsets – Proprietary, vendor-specific signaling protocols
- PBX and Central Office (CO) Exchange – Standard signaling protocols
- PBX to PBX – Both proprietary and standard signaling (feature loss with standard signaling)

In general, customers moved to non-standard signaling to take advantage of enhanced (though often unused) functionality. This strategy worked well enough when a customer was using products from a single vendor. However, the downside was that the customer was locked into a permanent relationship with that vendor, losing interoperability with products that relied on existing industry standards.

2.2 LINE INTERFACES

As mentioned earlier, there are two types of line interfaces for legacy PBX systems. These are trunk-line interfaces that connect the PBX to the CO exchange, and terminal-line interfaces that connect the PBX to telephone handsets.

2.2.1 TERMINAL LINE CARDS

Every telephone handset connects directly to at least one corresponding port on a line card, although multi-line handsets and attendant consoles (Direct Station Select/Busy Lamp Field or DSS/BLF) may use up additional line card ports. Terminal line cards fall into two categories – analog and digital – and each supports only the corresponding analog or proprietary digital handsets.

The type of telephone handsets provided typically depends on the user's role and status within the organization. A manager might expect a full-featured phone. Department secretaries or administrative assistants often require specialized multi-line sets and a broader set of telephony features. In such cases, the telephony team is faced with increased cost and administrative issues.

2.2.2 TRUNK LINE INTERFACES

Trunk interfaces connect the PBX to the PSTN, enabling communication to the outside world. Trunks were the first part of the telephone network to adopt digital technology –deployments began in the early 1960s. Prior to this, telephone connections were entirely analog. Many of us use analog telephones at home or even at work, yet the phone systems to which we connect are almost exclusively digital, so it makes sense at this point to explain why we shifted to digital.

If you throw a pebble into a calm lake, it generates waves. These waves emanate out from the place where the pebble hit the water. As they travel further, the waves begin to flatten and attenuate. In a similar action, when we talk into a telephone handset, the microphone converts the sound waves generated by our vocal chords into electrical waves, which are transmitted down the line. Electrical waves behave like waves on the lake – as they move further away from the source, they flatten and attenuate, eventually becoming impossible to decipher. Early phone systems boosted (amplified) the signal, but this caused minor pops and crackles due to interference from other power sources, which were also amplified. In the days of analog telephony, long distance calls suffered from these hissing, crackling and popping sounds, which often made call quality very poor.

Enter digital telephony. When the electrical signal reaches the telephony exchange, it is sampled very quickly. Each sample is converted to a numerical value representing the frequency of the sound at the moment the sample was taken. This number is sent as a pattern of ones and zeros all the way through the network. If the signal becomes weak, then the ones and zeros simply regenerate along the way – without the hissing and crackling, of course.

Traditional digital trunks are often sold in terms of multiple channels, each with a capacity of 64 kilobits per second (kbps). These channels form the basis of the global telephony network – so where did the number 64 come from? The answer relates to the way we convert analog sound waves into the ones and zeros carried over the digital network. The frequency range for the human voice has a size of 4000 hertz (Hz). To render this into numerical values that can be converted back into something representative of the original sound waves, we need to sample the wave at twice the highest frequency value, i.e. 8,000 samples per second. Eight bits represents the frequency values numerically, so each time the sound wave is sampled, we use eight bits. To sample 8,000 times per second times eight bits per sample (8 x 8000), equals 64,000 bits per second: 64 kbps. Today's codecs can improve on these numbers, but 64 kbps is still found throughout the telephony network.

It is important to understand the needs of your organization and associated costs when selecting and ordering a trunk connection from your local telecommunication provider. Technology in this area moves very rapidly, so care should be taken not to sign up for long-term contracts that may lock you into outdated technology. Finally, IP trunks are now widely available commercially. They compete with traditional digital or analog connections, keeping the call on-net to the carrier, which then uses its own gateways to break out to the PSTN.

Depending on your business requirements, it may make sense to establish service level agreements as part of your service provider contract. These can include: time to respond, time to fix and latency over a data or VoIP link. Third-party applications and appliances can be used to independently gather statistics concerning availability and service quality.

Another factor is that signaling can be in-band (robbed bit) or out of band with the use of a separate, dedicated channel. In Table 1, the letter D stands for a dedicated signaling channel. The channels used to carry voice calls are known as bearer or B channels. So the formula 2B+D describes an integrated services digital network (ISDN) Basic Rate Interface (BRI) providing 2 x 64kbps channels.

BRI interfaces are still widely used outside North America, although DSL is increasingly being used to carry on-net VoIP using open protocols like the session initiation protocol (SIP). One thing to keep in mind about BRI is that for historical reasons, two interfaces are available. North America utilizes the U interface, which connects directly to the local exchange. In Europe, the S/T interface connects the ISDN device behind a small network terminal owned that is operated by the service provider.

Trunk Type	Channels
Analog (FXO)	1
ISDN BRI	2+D
T1	24
T1-PRI	23+D
E1	32
E1 PRI	30+D
DS3	672
SIP trunking	Bandwidth dependent

Table 1: Trunk Options

Trunks can be analog (like the foreign exchange office or FXO) or digital: T-1 with ISDN PRI; or E-1, which is used in Europe. The various options for each of these trunk types offer tradeoffs in terms of cost, capacity and features. Most vendors support the full range of trunk options available; so tradeoffs are based on cost and which features are required by the customer. These are discussed in the next section. The technical details of how PBX systems and central office exchanges initiate calls and present audio streams over trunks are beyond the scope of this guide.

2.2.3 TRUNK FEATURES

A common feature deployed by nearly all businesses is Caller Identification (Caller ID). This allows the called party to see the calling party's name and telephone number before picking up the phone (unless the calling party has specifically blocked this feature). There are two Caller ID formats for delivering this information—Single Data Message Format (SDMF) and Multiple Data Message Format (MDMF). SDMF provides the calling number, while MDMF provides any combination of calling name and number. Note: If you are leveraging a complex call center application, be sure to work closely with your vendor to determine which other trunk features may be necessary.

Two additional mechanisms deliver caller ID:

1. Automatic Number Identification (ANI), similar to Caller Line Identification (CLI)
2. Dialed Number Identification Service (DNIS), an enhancement of 800-number services that enable the use of CLI intelligence for sophisticated routing of calls into the organization.

Another feature delivered by your telecommunication provider (telco) is used for inbound call routing. In North America, it is called direct inward dial (DID); in the U.K., it is DDI. This feature enables external callers to contact a user directly at his or her unique phone number, without intervention by an automated attendant or operator.

DID trunks are ordered in blocks consisting of 20 or more 10-digit telephone numbers. These numbers are assigned by the telco to each customer, and are routed to DID trunks connected to the PBX. When a call is made to a DID number, the telephone company sends the last three or four digits of the 10-digit number via the DID trunks at call set-up time. The PBX monitors for the digits and routes the calling party to the called party's extension. "Wink start" is a mechanism for initiating an inbound call and passing the extension number to the PBX using a specific signal. Analog DID trunks are inbound only and cannot be configured as two-way trunks. Connecting PBX systems across the WAN or within the same office location can be accomplished using either T- 1 or analog interfaces. These interfaces were designed to interact with the telco's CO switches; therefore, one of the PBX systems must simulate CO signaling to enable the two PBXs to communicate effectively. Similar schemes are often used when configuring a gateway or IP telephony system to connect to a legacy PBX.

2.2.4 TRAFFIC CALCULATIONS

To decide the exact number of telephone lines and trunks your company requires, first determine the number of telephone users, calling traffic and acceptable percentage of call blocking (failure of calls completed due to an insufficient number of available trunks). A sample traffic calculator for determining the number of telephone lines and trunks can be found at www.erlang.com. If no data is available for determining your telephone line and trunk requirements, you can follow the recommendations given in Table 2.

In general, smaller installations require more trunks per telephones (typical configuration), whereas larger installations do not need as many trunks per telephones (light configuration).

Telephone Traffic	Trunks per Telephones	Trunk Factor
Heavy	3 trunks per 6 telephone users	3/6
Typical	2 trunks per 6 telephone users	2/6
Light	1 trunk per 6 telephone users	1/6

Table 2: Trunk Ratios

Note: These numbers are not applicable to call center implementations, which are much more intensive users of trunk capacity. In call centers, calls are often held in queue prior to passing them through to agents. Please consult your vendor for suggested ratios.

2.3 CABLING

The cables pulled between telephone devices represent a significant portion of the investment in the phone system. It is important to ensure that the cabling is appropriate for that location and is installed correctly. Today, category (CAT) 5e twisted pair cable is the most popular cabling system. It carries both voice and data traffic at gigabit-per-second speeds.

The jack linking a cable to the desktop varies, depending on whether a telephone or a network device (such as a PC Network Interface Card (NIC) is connected. The Ethernet NIC uses an RJ-45 plug, but a standard analog telephone utilizes an RJ11 plug. When the cabling system is installed, the vendor tests each line for integrity. It is important to ensure that this testing is performed and that test reports are provided on each line.

The other end of the cable terminates close to the PBX, normally at a distribution frame or punch down block. The distribution frame is a rack-like structure where cables are threaded from an entry point to the appropriate exit point. The telephone engineer establishes the connection using a special-purpose tool that pushes the copper wire into a receiving contact. A dedicated corporate telephone network—where phones are connected directly to the PBX through a structured cabling system—increases reliability, but decreases flexibility. Moves, adds and changes (MACs) in the legacy PBX environment often require reconfiguring the wiring infrastructure. According to many enterprise telecom managers, a typical mid-size enterprise experiences MACs that involve approximately 12 percent of its users every year, with an average cost of $150 per user. Therefore, MACs in a traditional PBX environment are a significant, yet hidden, cost of ownership.

In contrast, data network cabling terminates desktop wires on a patch panel, so that an Ethernet drop cable can link the desktop device to its corresponding Ethernet port. This same scheme is increasingly being used for voice cabling, because it significantly reduces the costs of handling MACs.

2.4 BASIC FEATURES AND FUNCTIONS

A telephony system is expected to deliver basic features and functions, and we expect these features to behave in a predictable and familiar manner. Following is a list of the typical features available to users and administrators:

- Speaker button
- Mute call button
- Call forward
- Call transfer
- Blind transfer
- Call park
- Conference
- Hunt groups
- LCD displaying calling information
- Support for DTMF codes
- Programmable keys
- Redials
- Music on hold
- Last number redial
- Call pickup
- Shared line ringing
- Line hold (Hold)
- Speed dial

Value-added features are often embedded in telephone handsets to encourage customers to upgrade in order to gain access to these functions (handsets represent a large portion of the overall cost of owning a PBX). The feature lists associated with these handsets are fairly similar from one vendor to the next. Unfortunately, adding features in new handsets requires significant engineering in the central PBX every time a new feature is added. And even more problematic, the list of required features is exploding. The good news is that as the market continues to move to a all-pervasive IP environment, adding new features is similar to loading a new plug-in for a Web browser. This ability to increase the functionality of voice communications is a critical driver for the adoption of next-generation telephony.

2.5 ENHANCED FEATURES AND APPLICATIONS

Beyond the basic feature list, PBX vendors are scrambling to develop additional application components that can be added to the system, in order to significantly increase the types of services provided by the phone system. The final section of this guide provides an in-depth view of some of the more strategic next-generation applications, such as unified messaging, voice recognition and CRM. These adjunct systems are frequently listed by PBX vendors in their solution offerings:

- Voice mail
- Automated attendant
- CTI connectivity
- Conference bridge

Often, these systems are not fully integrated within the PBX itself, but are part of an increasing number of system adjuncts that reside outside the chassis and are linked via several line interfaces. The cost of such applications is beyond the basic PBX purchase and significantly increases the price of the overall system.

The model of adding value to the system using third-party devices is made easier when the voice system and applications are designed from the ground up to share a common IP infrastructure. The advantages of different architectures are covered in Section 3.

2.6 CALL FLOWS AND DIAL PLANS

When installing a voice communication system, one of the most important decisions that must be made is how calls are routed, even when the person is not available to take the call. Will calls be transferred to the auto-attendant, operator, assistant, off-site number, pager or cellular phone?

In evaluating how to determine call routing policies, it is imperative to seek input from system users, particularly high-volume users and groups. For service centers and customer reps, “hunt groups” and workgroups often must be defined. The term “hunt group” describes the way a call might be handled by the phone system. For example, if a call is not answered by a customer agent after a few rings, it is forwarded to the next available phone in the agent group until it is picked up. If the call reaches the end of the available extensions without being picked up, it may be passed on to the group's voicemail. Understanding and configuring such functionality is critical to building a successful system.

The call handling process also must be carefully planned for outbound calls in such a way that, for any number dialed, a corresponding route is available for it. For very large multi-site systems with local hop off, the dial plan information can become quite complex. Here are a few examples of call handling policies:

- Your New York office is linked to the Dallas office. You would like to save money by routing long-distance calls over your company network. For example, if someone places a call from the Dallas office to an external number in New York, the call transits between internal PBX systems in those two offices. The external call only has to make a local hop to the destination number, saving long distance charges.

- You make a deal with a long distance carrier for calls made to London, so that all international calls to the country code 44 are prefixed with the alternate carrier's prefix number.

- When a staff member calls an internal extension using the full external number prefixed by 9, the system automatically strips off all but the extension number and routes the call internally.

Whether an external service organization is involved or not, it's clear that defining dial plans requires careful analysis and thought.

2.7 AUTOMATED ATTENDANT

The auto-attendant provides a customizable way for incoming calls to be quickly routed to their destinations. This application uses in-band signaling called Dual Tone Multi-Frequency (DTMF) codes. DTMF assigns a certain sound frequency to each telephone key, so when the dial pad keys are pressed, the auto-attendant "hears" these frequencies, interprets the information contained in these frequencies and acts on the information.

For small businesses, the immediate advantage of an auto-attendant is the cost savings of not hiring an operator. However, keep in mind that this feature can frustrate potential customers if the menu levels get too deep. When comparing features of an auto-attendant application, consider these questions:

1. How many menu levels can the system provide?
2. Can I design different menus depending on time of day and year? How many?
3. Can these menus be programmed to automatically update themselves on a particular time/date?
4. How does the system handle incorrect user input?
5. Can the seasoned user go straight to a destination?
6. Can the user bypass a long prompt or are they forced to wait?
7. Does the system provide directory search with name lookup?
8. Can the system forward calls to workgroups or call center agents?

2.8 CALL DETAIL RECORDS AND BILLING

It is important to manage the overall cost of running the telephone system. PBX systems typically generate detailed logs of calls on the system. These call logs can be outputted from the management console and saved to file for processing and analysis.

Because PBXs are isolated from the IT infrastructure, the generated call detail information is fed into a report engine that produces more structured reports by department, group or usage cost. This information can be used to answer questions like:

- What calls are being made outside office hours and where are calls being placed?
- Which extensions are costing the organization the most money?
- What are the phone usage costs by department?

Caller Line Identification (CLI) can be used to determine the duration of calls from specific customers. This information can be useful for basic customer billing or service level review. More detailed statistics require a call center-type system. The process of generating such reports can be outsourced to third parties that take the basic PBX data and convert it into useful reports. Service organizations like legal, advertising, etc. that bill by the hour use such call detail records as input into customer billing systems.

2.9 NEXT STEPS

The information presented to this point provides a basic understanding of business telephony – at least, the way it used to be. But the reality is that the world of voice communication is changing, and as a result, next-generation IP technologies are replacing outdated TDM technologies in the enterprise. This section includes an introduction to data communication technologies and explains how IP voice communication is delivered on top of this infrastructure. Data networking professionals may want to skip ahead to Section 4.

3. DATA NETWORKING

The fundamental building blocks of a typical enterprise data network are Ethernet, switching, IP and routing.

3.1 LAN INFRASTRUCTURE

Today, a majority of LANs are based on Ethernet technology. Increasingly, IP runs over Ethernet, replacing protocols like Systems Network Architecture (SNA), IPX and AppleTalk. Ethernet has moved from shared bus to shared hub, and today, switched Ethernet dominates the enterprise. Speeds have also improved from 10 Mbps to 10,000 Mbps (10 Gbps IEEE 802.3ae). One of the earliest evolutions was the shift from shared coaxial cable to twisted pair cable. Ethernet is standardized as IEEE 802.3. Each device on the network has a unique six-byte media access control (MAC) address. Three bytes identify the vendor and a different three bytes identify the specific device. Information is sent between network devices using a predefined format known as a frame (see Figure 2). Frame formats continue to evolve but are backward compatible.

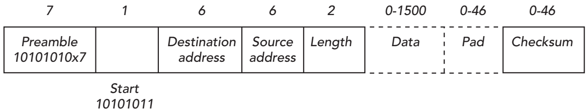

Figure 2: Ethernet Frame Format

The NIC device resides on all networked devices in one form or another. It sends and receives all the signals to and from the device, and is responsible for packaging raw information produced by network devices into frames, before the data is sent to the cable that connects the device to the network. If the target address of a specific communication does not match that of the NIC, then it simply ignores the frame.

To communicate with another device, the sending device first listens for a quiet period, then begins transmitting. It listens to make sure that its transmission has been correctly sent, i.e.: the checksum matches the transmitted data. To prevent two devices from communicating at precisely the same time, Ethernet employs a scheme known as Carrier Sense, Multiple Access/Collision Detection (CSMA/CD).

Here's how CSMA/CD works. Imagine two computers that hear silence on the media and determine that it is safe to transmit. They both transmit and listen at the same time, so if another device heard silence and started transmitting at exactly the same time, they would immediately recognize it, because the information they detect coming back on the network would not match what they sent. They have detected a collision.

As a result, each computer then backs off – quickly flooding bits onto the cable and ceasing transmission for a random amount of time. They resume listening for a quiet slot and the cycle begins again – though this time, hopefully without a collision.

This simple mechanism has been found to be fairly scalable and works well on LANs with a small number of users. Over the years, Ethernet's CSMA/CD has outperformed competing approaches, such as IBM's Token Ring.

As you design a network infrastructure, keep in mind that Ethernet has some important constraints in terms of cable lengths:

- The maximum length of twisted pair cables that connect Ethernet switches to devices (other switches, computers, IP phones, and so on) is 100 meters.
- The maximum fiber cable length is 420 meters.

Today's cabling is typically UTP type 5 or 5e. A test certificate should be obtained from the cable contractor to ensure the cable conforms to Ethernet requirements. (RJ-45 pin layouts are defined in TIA 568B.) When you connect an Ethernet switch or hub to another switch, you need a crossover cable (also defined in TIA 568B), unless the uplink port of the switch is used.

Although the hub and spoke topology (one of a number of different topologies that can be used with Ethernet) created by a LAN switch (the hub) and several NICs (the spokes) is superficially similar to the PBX and twisted pair cable that connects telephone handsets, there are fundamental differences between the two systems. These differences include:

1. Unlike PBX systems, LAN devices can be easily interconnected and daisy-chained to extend the network's capacity.
2. Unlike PBX systems, Ethernet devices are backward compatible, so older NICs continue to work with newer switch ports.
3. Unlike PBX systems, Ethernet is an open standard (IEEE 802.3), and any compliant device can be added to the network, irrespective of vendor.
4. Unlike PBX systems, addressing schemes with Ethernet are relatively easy to implement, because Ethernet-compliant devices have unique MAC addresses built into the hardware, enabling network managers to deploy Ethernet without having to manage the addressing scheme. Devices simply "declare themselves" on the network. IP addresses above this layer must still be managed, but even these can be allocated automatically using a scheme like Dynamic Host Configuration Protocol (DHCP).

3.1.1 ETHERNET SWITCHING

Over the last 25 years, Ethernet has evolved. Today's Ethernet networks are built from both chassis-based and stackable switches, rather than shared media hubs.

A switch provides each device connected to one of its ports with a dedicated bi-directional (full duplex) connection. This means that a device connected to a switch port communicates at the maximum speed supported by that device. This differs from shared Ethernet topologies (such as hubs or, more often today, wireless), where the bandwidth is shared.

To achieve this improvement, Ethernet switches must know the addresses of as many of the devices connected to them as possible, and identify the ports used to reach these addresses. Switches automatically do this using a protocol defined in IEEE 802.1, transparent bridging. The switch stores the source address and switch port of every frame it receives in a table, and finds the destination devices by flooding its other ports with a request for the destination device. When the destination device responds, its address and port number are added to the table. After source/destination addresses are known, the switch uses that information to begin forwarding frames. Shared hubs forward frames to every device connected to them, reducing overall throughput for every device.

In some circumstances, it makes sense to segment traffic either by department or by application using a virtual local area network (VLAN) in order to enhance security or optimize bandwidth. Given that voice and data are sharing the same switch infrastructure, it may make sense to segment the LAN into smaller groups of users to protect real-time voice traffic from unpredictable data traffic (which can create spikes of high-volume traffic over brief time periods). One method of ensuring optimum voice quality is to run voice traffic on a separate VLAN. This virtual segmentation allows voice traffic to share the same physical infrastructure as bursty data traffic, but voice traffic is protected at a logical level from interacting with data traffic.

Ethernet switch architecture can also be designed to eliminate points of failure – uplinks, specific switch ports – that could impact everyone in a department or office floor. Redundant links can be built between switches, but this introduces the problem of a logical loop, where switches keep claiming they are responsible for devices that are, in fact, connected to some other part of the network. Or worse, the redundancy could lead to broadcast storms, where switches continue forwarding broadcasts and network devices respond to those broadcasts, until the responses feed back on themselves causing a network meltdown.

The spanning tree algorithm IEEE 802.1d provides a way to benefit from the redundancy, while avoiding the problems described above. Each link is weighted. For any path, a switch uses only the lowest path for a link, ignoring the others. It should be noted that spanning tree information takes time to update in the event of catastrophic failure, and this process of updating spanning tree information can impact a call in progress. Switch manufacturers have developed proprietary solutions for providing rapid spanning tree updates, and these solutions largely address the challenge of maintaining call quality when the network is experiencing technical issues.

When designing enterprise networks, it is important to recognize that in spite of efforts to segment traffic, much of the traffic still transits certain links, resulting in bottlenecks. The IEEE 802.3ad standard addresses this bottleneck by providing a standard mechanism for aggregating multiple links between switches.

In concluding this section on LAN infrastructure, we would like to point out that while Ethernet's plug-and-play design makes it easy to implement, as the leader of your organization's migration to a fully converged voice/data network, you are seeking to implement advanced networking capabilities, such as redundancy, link aggregation and quality of service (QoS), which require careful planning and fine tuning.

3.1.2 POWER OVER ETHERNET

One of the advantages of an IP-based PBX system is that it enables the use of a converged network (as opposed to maintaining two separate networks for data and voice). IP telephones plug directly into the Ethernet network, and interact with a media gateway controller (MGCP) or a gatekeeper (H.323) for call control over the LAN.

However, the challenge of a single, converged network is that the phone (which is seen as a lifeline to emergency services) may not be available during a power outage. With users expecting dial tone no matter what else is going on around them, this can create problems.

Even though in many cases, digital sets were not line powered, the legacy PBX vendors—seeing an opportunity to hold back the inevitable— accused the IP-PBX community of cutting corners on fundamentals. This led to the myth that data networks need to be upgraded so that low voltage devices like wireless hubs (see next section) and IP telephones could be powered through the LAN.

The IEEE 802.3af standard defines two ways to provide power to IP phones:

1. End Span – Replace Ethernet switches with new devices that utilize DC current over the pairs used for data 1/2 and 3/6 (on the RJ-45 jacks). This approach is most appropriate for a new building or as part of a major network upgrade, because it requires new Ethernet switches.

2. Mid Span – This device inserts power onto the unused 4/5 and 7/8 pairs on the RJ-45 jacks. The device has two ports and sits between the Ethernet switch and the device it is powering. It is a less expensive option than upgrading the Ethernet infrastructure.

Note: In both cases, the system is non-destructive, because a non-compliant device can be plugged into the powered line without damaging the device.

3.1.3 WIRELESS LANS

In the USA, wireless premises for voice has remained vendor-specific, in contrast with Europe, where regulated spectrum was established and a standard approach has been widely adopted – the European Telecommunications Standards Institute (ETSI) Digital Enhanced Cordless Telecommunications (DECT) technology. In 2005, the Federal Communications Commission (FCC) opened up a spectrum in the 1900 MHz band, which effectively means DECT can now be employed in the U.S.

Whatever air interface is selected, the challenge with wireless has always been money. The cost of providing full building coverage, campus-wide roaming and establishing the necessary access points or base stations is often too high. It is recommended that a careful financial analysis accompany any large-scale wireless project, because technology decisions can lead to vastly different price points.

WiFi or 802.11 technologies are the current leaders in mindshare for wireless premises in the U.S. market. Broad market adoption has helped drive access point prices down, and many enterprises are intrigued by the potential to provide employees with a single mobile device for applications, as well as voice.

The 802.11 standards continue to evolve, and the 802.11n standard is expected to be ratified in late 2008. It improves on previous 802.11a and g standards (54Mbit/s) to deliver a maximum data rate of 248 Mbit/s.

Key factors to keep in mind when planning a wireless implementation:

1. QoS – Delays for enterprise voice should not exceed 150ms. Given that *WiFi* is a contention protocol (like the original shared Ethernet), when a particular access point is heavily used, voice quality suffers. Ratified in 2005, IEEE 802.11e defines traffic classes, assigning time-sensitive voice traffic to a higher class relative to other traffic types. Not all wire less solutions implement 802.11e. We recommend favoring solutions that include this standard when designing voice over WiFi solutions.
2. Reliability – Except for 802.11a, which specifies the 5GHz band, 802.11 standards use the 2.4 GHz band and suffer from interference with other wireless devices.
3. Security – A voice-specific encryption standard, IEEE 802.11i, was developed. However, data encryption can slow down voice delivery.
4. Standards – Multiple signaling standards (802.11a, 11b, 11g and the emerging 11n), should be compared based on distance, capacity and frequency interference. For international markets, keep in mind that not all standards or in some countries, channel slots within the standards, are available.
5. Handsets – Running voice through personal digital assistant (PDA) devices may be desirable, but there are implications for battery life, overheating and emissions that should be carefully evaluated prior to an organizational adoption. Traditional wireless premises phones are more successful.
6. Coverage – At the high frequencies used by wireless local area networks (WiFi), establishing coverage throughout a building is not a trivial task, and should be carried out by experienced wireless network designers.

3.2 IP

As mentioned earlier, Ethernet is the dominant LAN technology. It scales well and works over twisted pair, fiber optic cable and even wireless. Why do we need anything else? Let's consider some of the challenges of running Ethernet:

- How do we tell applications running on computers the best way to find the addresses of other computers with which they need to communicate?
- How do we communicate over very long distances? For example, WAN connections spanning the globe.
- How do we deal with circumstances where economic constraints force us to use some other transmission technology? ("I could sell you Ethernet between those two sites, but frame relay is half the price.")
- How do we communicate with other organizations that do not necessarily want us to know the specific addresses of machines on their internal networks?
- How should we package information, like e-mails or graphic-rich documents, so that other machines can download and display the information?

The answer to all these questions is that a protocol was needed that could provide an open interface between applications and the various physical networks underneath (e.g. Ethernet on the LAN, frame relay, digital subscriber line (DSL) and ISDN in the WAN, and so on.)

Various protocols have been proposed, including:

1. IP Internetworking Protocol – IETF standard
2. Open Systems Interconnect (OSI) – ITU standard
3. Systems Network Architecture (SNA) – IBM
4. DECnet – Digital Equipment Corporation (DEC)
5. Internet Packet eXchange (IPX) - Novell

But interestingly, one of the oldest protocols finally won the race. IP was originally developed by an agency of the United States Department of Defense, the Advanced Research Projects Agency (ARPA), as a peer-to-peer communication protocol that would While this level of reliability is extremely valuable, the true drivers for IP's success were:

1. IP protocols are simple to understand and implement
2. IP protocol specifications are freely available
3. New applications like the Web encouraged broad market adoption

The IP community had an enthusiastic, open and sharing philosophy that fundamentally drove adoption through universities and beyond. Today, IP is employed across the globe as the basis for the Internet and most enterprise networks.

3.2.1 IP ADDRESSES AND NAMES

IP addresses identify devices connected to the network. The 32-bit addresses are written as four decimal values (with each group capable of representing less than 256 values) separated by a dot. The first groups of numbers identify the network where the machine is located, and the other groups of numbers identify a specific machine within that sub network.

Here is an example of an IP address:

169.254.70.213

When installing a computer, the network manager defines the machine's address and the address of the default gateway – a machine on the local network that provides connectivity to other networks. In the early days of the Internet, this machine was a computer with two network interfaces, and it performed the task of forwarding packets between the two interfaces.

Today that machine is more likely to be a router. For machines that are connected directly to the Internet, IP addresses are carefully allocated in blocks and managed by the owner of the addresses, to ensure that no two machines have the same address. However, within an enterprise, the IETF allocates certain address ranges for internal machines only. As we discuss in the upcoming Security section, addresses of internal machines are protected by a firewall that provides network address translation (NAT). This not only protects the identity of internal machines, but also ensures that addresses belonging to one enterprise do not conflict with similar addresses used by another enterprise.

Like telephone numbers, the numbers to the right of the address identify a specific machine or device. Legacy voice managers should learn the details of the addressing scheme used within the organization before connecting new IP devices to the network.

To make it easier for people to find services on their network and the Internet, a name-to-number translation service was introduced, the Domain Name System (DNS). This global system allows the use of a name like http://www.louvre.fr/, rather than its IP address, http://160.92.103.98/. The DNS service simplifies things by allowing the end user to memorize a Web site's name, instead of trying to remember its IP address, which could change if the Louvre changed to a different service provider.

When you consider how phone numbers are entered in your cell phone, it's clear that here too, people's names are used, rather than their numbers. This facility is being introduced by the current generation of IP voice systems.

3.2.2 ALLOCATING AND MANAGING ADDRESSES

Given that each device has its own IP address, does this mean that you must manually assign an IP address to every phone installed on your network? Fortunately, the answer to this question is "no." An automatic address allocation system, Dynamic Host Configuration Protocol (DHCP), eases the administrative overhead associated with IP address allocation:

The DHCP protocol resides on a server and manages a pool of IP addresses. DHCP keeps track of which addresses are currently in use and which are available for allocation. For devices that conform to certain criteria, IP addresses can be leased from an address range. For example, an Ethernet MAC address range and leasing can be time-bound; to handle applications such as shared desks for mobile employees.

3.2.3 WAN INFRASTRUCTURE

Because IP is designed to function independently over lower topology layers (i.e., closer to the physical layer in the OSI model), the options for interconnecting IP devices over the WAN are almost limitless.

In today's market, options include:
1. T-1 and fractional T-1
2. ISDN
3. Frame relay
4. Asynchronous transfer mode (ATM)
5. DSL
6. Cable
7. Wireless local loop technology, such as Wimax/802.16, 802.11 and Local Multipoint Distribution Service (LMDS)
8. Synchronous Optical Network (SONET)
9. Wide area Ethernet (delivered by metropolitan area networks)

Previously, you may have managed an access router with a rack of modems, but this service has been replaced by the combination of service providers and VPN technology. Enterprises connect their LANs to the WAN using a router. The router combines IP routing intelligence with knowledge of the appropriate lower-level network signaling used on the WAN. Such devices range from low-end access devices for connecting home users to large enterprise routers with redundant links for failover and load balancing.

3.2.4 ROUTERS AND ROUTING

In the telephony world, conversations are carried out over circuits, which are end-to-end connections between the caller and the person being called. All the work of determining how to route the call is done during call setup. Once the circuit is in place, no further route decisions are required (assuming there are no catastrophic problems on the network).

Data networks, however, do not work like this. Packets are forwarded and forgotten, so that each intermediary router between the source and the destination performs these steps:

1. Reads the packet's destination address
2. Checks which route to use
3. Forwards the packet to the interface associated with that route
4. Forgets it

You can think of routers like a group of children standing in a circle, throwing hot potatoes to each other. Each child catches the potato, throws it before his fingers burn and forgets it. There is no notion of a circuit, although as discussed in Section 5.2, which looks at Multi-protocol Label Switching (MPLS) and quality, voice is placing new requirements on router protocols.

The primary function of a router is to direct packets to the best path across a network. Each router maintains a route table, which is a network roadmap that is kept up to date by exchanging information with other routers about the status of each link and the network. When an incoming packet is received, the router identifies the destination address, checks a route table to determine the best route, and then forwards the packet to the next router along the path. A router should be connected to multiple forwarding paths so that if one path fails, packets are rerouted around the failed connection.

Due to their strategic location in the network (at the LAN/WAN boundary), routers are frequently used to prioritize and filter traffic. WAN links are shared resources, and therefore suffer from similar contention problems inherent with the early Ethernet. This can cause serious problems for real-time voice communication.

Routers play an important role in guaranteeing QoS for voice quality on the network. QoS is discussed in section five. VRouters also act as firewalls, filtering packets to protect the network from unwanted attempts to gain access. Firewalls rely on techniques similar to traffic prioritization; that is, they identify and filter traffic based on source or destination address, protocol type or IP port numbers. Port and socket numbers, in particular, may indicate application functions such as telnet or file transfer protocol (FTP). Because these applications can be used to break into a corporate network, identifying types of traffic before it enters the network can offer valuable protection. There are literally hundreds of well-known techniques for breaking into a network, such as IP spoofing, Denial of Service (DoS) attacks, or synchronize (SYN) message floods. In all cases, the router or firewall must be capable of identifying and filtering these types of traffic.

Use of traffic prioritization or firewall technology can become an issue when transporting voice over the network, due to the additional processing required for these functions. However, the newest generation of routers and stand-alone firewall devices are much more powerful, making use of custom, application-specific integrated circuits (ASICs) to simultaneously classify, queue, filter and forward packets with minimal latency.

3.2.5 TCP/UDP THE TRANSPORT LAYER

At the network layer where IP resides, issues such as addressing and routing of packets arise. Above the network layer, there is a choice of two protocols:

TCP: Transmission Control Protocol (includes error checking and correction)
UDP: User Datagram Protocol (no error checking)

While it may be counter-intuitive, voice data actually performs better without error correction. The reason for this is that notifying the sender that a packet has failed to arrive takes time. So, it actually works better to discard the offending packet and play an approximation of the missing data. Usually, the previous packet is played again (although other, more sophisticated mechanisms are available).

Recently, another protocol was proposed at this layer, the Stream Control Transmission Protocol (SCTP). It was designed specifically to reliably carry the time-sensitive signaling used for carrier networks. It is unlikely this protocol will be used in enterprise applications unless your organization provides sophisticated call center applications that must interact directly (at the Signaling System 7 or SS7 level) with public networks.

3.2.6 THE REAL TIME PROTOCOL (RTP)

During a conversation, people can make sense of imperfect sound in which small (barely perceptible) bits of audio are missing. Audio CDs and the pointillism painting technique are two examples of how the mind is capable of putting together a complete picture from fragments. But if a communication system must stop processing to ask for a packet to be retransmitted, the delay can cause a serious degradation in the conversation's perceived quality. In fact, for real-time communication, it actually makes more sense to discard the missing packet rather than ask for it again.

The RTP protocol is designed to handle the needs of real-time communication. Among the fields defined for an RTP message format are:

- Sequence Number: Incremented for each RTP packet
- Time Stamp: Records the sample rate and therefore, playback rate

A variety of different approaches (in a VoIP implementation) can be taken when handling a lost packet during a conversation. For example, we can simply replay the sounds from the previous packet, discard the missing packet, then play the next one when it arrives.

Another aspect of real-time communication is that in order to achieve higher quality sound, the sampling rate must be increased, which leads to smaller packet sizes. These avoid the increased processing latency caused by large packet sizes, but small packets cause a new inefficiency—namely, packet headers take up as much bandwidth as the payload (i.e., the actual sound being transmitted). RTP solves this by using header compression. RTP does not handle voice signaling or define the format for transporting voice packets, but it provides an important solution to the challenges of real-time voice communication. RTP provides the necessary infrastructure and protocol foundation upon which to build an IP voice system.

Problems of latency (perceptible delays between talker and listener) and jitter (a gradual loss of synchronization between the two endpoints) can still appear. Given the speed and reliability of most enterprise LANs, such problems are rarely seen for conversations at a single site. Instead, they tend to crop up in WAN communication. Careful planning should be taken to ensure the WAN is correctly sized for the organization's needs, and appropriate jitter buffers must be set and verified. This is discussed in more detail in Section 4.

4. VOIP – TECHNOLOGIES AND STANDARDS

4.1 HOW IP VOICE ACTUALLY WORKS

To understand the components of a typical VoIP system, let's begin by looking at what goes on while two parties talk to each other. Imagine you are holding a traditional analog phone. When you speak, your voice excites the airwaves, literally creating a wave pattern in the air – an analog signal. This sound wave is picked up by the microphone in the phone's handset, which replicates the analog signal as electrical impulses. The signal is sampled and converted to a numerical representation of the wave pattern – ending up as a series of ones and zeros. The resulting digital signal can then be compressed, if desired.

Thus far, the process described is exactly the same for a digital telephone connected via a line interface on a legacy PBX. But that's where the similarity ends. The next step consists of preparing the signal for transmission over the network infrastructure. For VoIP, this means dividing the stream into IP packets (see Figure 3). At this point, a key difference between legacy telephony and IP telephony emerges. Legacy PBX systems use proprietary, non-interoperable signaling, while IP telephony leverages standard signaling protocols.

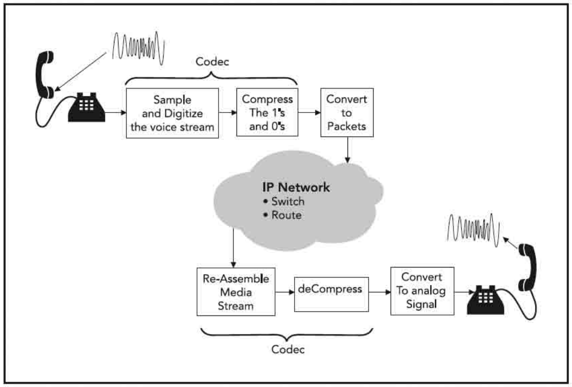

Figure 3: VoIP Media Stream

4.2 VOIP COMPONENTS

Having established the process for voice communication across any data network, the next challenge is to provide sophisticated call control (call set up, tear down, etc.) capabilities.

The act of setting up, tearing down and routing calls requires intelligence that goes beyond the simple transmission issues we have covered so far. Much effort has been put into standardizing these rules, so that IP voice systems can accomplish call processing in the same predictable manner as Web servers deliver information to browsers anywhere in the world. This section focuses on the functionality required.

Softswitch: To set up a call, the system must act on signals from the calling phone. One way to accomplish this is with specialized call processing software that tracks and manages call progress. It also handles conversion between the addressing schemes used on a data network (IP addresses) and telephone numbers (defined in ITU E.164). There are different names for this function: call server, call processor, gatekeeper, media gateway controller or softswitch (see Figure 4).

This device is an automated operator, handling all the tasks the switchboard operator used to handle. Figure 4 shows an imaginary dialog between the call control software located in the softswitch or IPBX, and the telephone.

Telephone	Call Server/SoftSwitch
I've gone off hook	OK, here's your dialtone
Here's the phone number I dialed	OK, I'm routing this; here's some progress info
I've just put this caller on hold	OK, I'll remember that call for you, here's another dial tone
I've replaced the receiver	OK, I'll store the call detail record and terminate the call

Figure 4: Dialog between telephone and softswitch

These examples – which are typical of signals that might be sent between a telephone and a call server – show how the call server performs the same functions as a PBX. So if the call server or softswitch can manage call set up, call routing and call tear down, does that mean this is a fully functional IP-based alternative to the PBX? Not quite. We are still missing an important interface with the legacy PBX. Specifically, we need a gateway between the IP world and the legacy, circuit-switched world.

Gateways: The gateway accomplishes this with three components:

1. Trunk or line interface on one side
2. VoIP transmission capability on the other side
3. In between, logic is necessary to convert between the two media formats and ask the call server for help in setting up the call

In practice, it often makes sense to combine the functionalities of call processing and gateway into single network elements. But for the purposes of this discussion, they are treated as separate components.

IP Phones: While a complete system with a softswitch and specialized media gateways can potentially support existing analog handsets, in practice, most implementations only support IP phones. These phones can either be hardware devices that plug into the Ethernet network (and look just like a normal legacy phone), or softphones that run on user PCs.

IP phones provide the functionality of a single user gateway, converting the analog speech pattern into digitized voice packets which are sent over the IP network. These are some of the characteristics you should consider when selecting IP phones:

1. Which signaling standard is used?
2. Does the phone provide a second Ethernet port, so a PC can use the same uplink as the phone (offering savings on cabling cost)?
3. Does the phone support Power over Ethernet (PoE), so it works without interruption during a power outage?
4. Does the phone provide a mechanism to classify traffic, so voice can be prioritized through the network?
5. Does the phone provide easy access to advanced features through an intuitive interface?
6. Is the phone easy to install and configure?
7. Does it deliver good sound quality?

As a mature market, differentiation was critical for legacy PBX vendors. Over time, phones and handsets became the focal points for vendor competition. Specialist phones were developed, including:

- Operator consoles
- Administrative assistants
- Key systems
- Conference phones
- Additional phones for various hierarchical levels within organizations

Today's IP phones offer software hooks for customization and are far more flexible than their legacy counterparts. Even the simplest phone design can be extended with applications that reside on user PCs. IP phones provide intuitive interfaces with access to application-rich features and are capable of leveraging recent improvements in sound quality to provide a better experience for users and the people with whom they communicate.

A complete phone system includes:
1. Softswitch or media gateway controller
2. Gateways
3. IP phones

Let's review how these components work during a call:

- A phone transmits state changes (off hook, on hook, etc.) to the call server or softswitch.
- The softswitch sets up calls, finds routes and keeps track of everyone's state.
- The softswitch automatically converts between telephone numbers and IP addresses.
- After a call route is established, the softswitch gets out of the way so the path for the voice stream is independent of the softswitch. (This is important because it prevents delay from being introduced into the conversation.)
- If the call leaves the IP network and is routed to the PSTN or a legacy PBX, a gateway converts the IP packets back into the appropriate media stream for the trunk.
- If the call is sent to another IP device, the call may be managed by multiple softswitches. But eventually, the VoIP packets reach the called party's phone and are converted back into voice.

Figure 5 shows how all of these elements work together in a complete business telephony system.

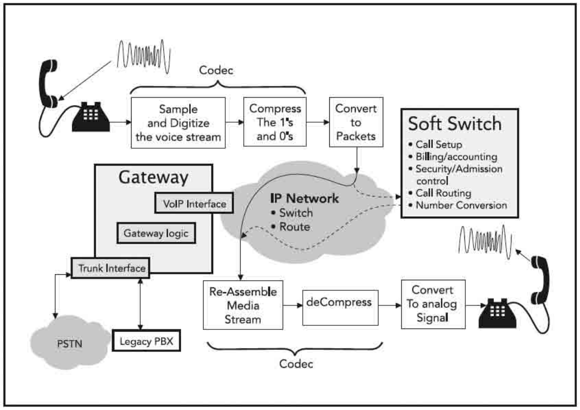

Figure 5: VoIP Architecture

Moving beyond the active components shown in Figure 5, a key point is the total flexibility of VoIP architecture delivered by the IP cloud. This IP network can span both LAN and WAN, independent of both location and service provider. Devices can simply be plugged into the IP cloud to become visible to the entire enterprise.

Management may be conducted from any point on the IP network and encompass every element of the business telephony system. This is a fundamental difference between legacy PBX and IP voice communication. Businesses can benefit from the inherent scalability of IP infrastructure, as well as interoperability between voice and computer components.

Having reviewed the functional elements that must be provided by an enterprise-class IP voice communication system, let's look at some of the available frameworks and standards.

4.3 VOIP STANDARDS

Before reviewing the specifics of VoIP standards, it is worth spending a little time discussing why standards are so important to the communication industry.

Fundamentally, standards provide the basis for communicating between vendor systems. Standards solve two requirements in a non-standard communication environment: vendor interoperability and service provider interoperability.

In Section 2.1, the discussion of signaling in the legacy PBX environment mentioned several areas where proprietary protocols lock customers into vendor-specific solutions that ultimately lead to increased cost of ownership (e.g. proprietary handsets, costly application integration and proprietary signaling between PBX systems).

Standards offer relief from vendor lock in through the use of IP, which forms the foundation of any standards-based implementation. However, this is only the beginning of the process to open up the business telephony world. The market is rapidly moving to an open systems architecture where standards-based phones, call servers, gateways and application servers will interoperate. This is the real promise and power of IP voice communication.

In the communication world, standards discussions often become political. Competition is based on different views of how problems can best be solved. In addition, standards specifications are complex and require a solid technical background to understand. With this in mind, Section 4.3.1 provides a brief overview of three key standards upon which IP telephony is based. (Refer to the *Glossary* for definitions of commonly used acronyms.)

4.3.1 SESSION INITIATION PROTOCOL (SIP)

SIP has emerged as a lightweight, extensible alternative to H.323. SIP defines standard objects/components and a limited message hierarchy for communicating between these elements.

SIP COMPONENTS
1. SIP User Agent Clients (UACs)
2. SIP Registrar Server – tracks which IP addresses clients are currently using
3. SIP Proxy Servers – forwards requests to other servers on behalf of SIP clients and provides them with target addresses, but retains knowledge of calls in progress
4. Redirect Servers – communicates target addresses of called parties to calling parties, then backs off
5. The SIP protocol defines a set of basic messages to signal events: *SIP Methods*
   i. Invite – to join a session/call
   ii. Ack – to accept this invitation
   iii. Options – determine the capabilities of a server
   iv. Register – register with a server
   v. Cancel – cancels a previously issued request
   vi. Bye – to end a call

Developed by the Internet Engineering Task Force (IETF), SIP focuses on session initiation, modification and termination – leaving the session and connection details to be negotiated by the end systems. SIP uses a simple, text-based command structure, with HTTP syntax and URL addressing. Thus, it is well suited for Internet and Web-based applications where, for example, phone calls and Web pages work together in customer call center environments.

Where terminals are concerned, SIP's emphasis – like H.323 – is still on end-point intelligence, and this has some significant implications for handset costs. However, SIP's key advantage is that it offers a well-defined mechanism for device-to-device signaling beyond the handset itself. Specifically, SIP is well-architected for communication between multiple proxy and location servers. This part of the SIP specification makes it very scalable and manageable.

SIP uses the Session Description Protocol (SDP) to negotiate capabilities for a range of features like audio and video codecs, encryption schemes, ports and addresses used between call participants. SDP content is included in the INVITE part of the caller's request.

SIP is also the basis for the Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE), one of two standards currently being proposed for instant messaging and presence. A voice system based on SIP is better able to integrate presence and telephony and delivers a richer suite of applications. For these reasons, we expect SIP signaling to a requirement of next-generation IP voice communication systems.

4.3.2 MGCP/MEGACO/H.248

Two more standards are worth mentioning: the Media Gateway Control Protocol (MCGP) and the closely related MEGACO protocol. Unlike SIP and H.323, MGCP assumes that edge systems are unintelligent gateways. So the gateway controller handles all aspects that go beyond media conversion. Central management of less-intelligent gateway devices is a reality in some business telephony implementations today. Reduced-cost IP phones can act as simple gateway devices (analog-to-IP converters) and the intelligence of call control can be handled by the gateway controller (see Figure 6).

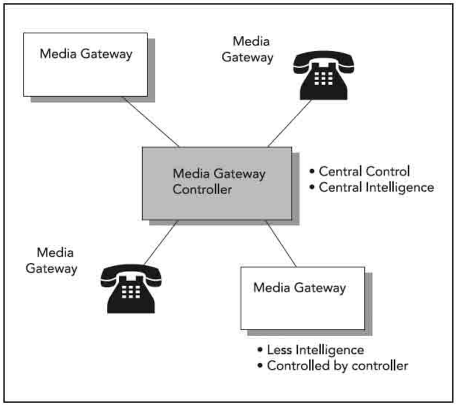

Figure 6: MGCP Components

When MGCP was initially introduced to the IETF for standardization, the name was changed to MEGACO and an agreement was reached with the International Telecommunications Union (ITU) to work on a parallel standards activity: H.248. The key difference between MEGACO and H.248 is that H.248 mandates the support of H.323 (see Section 4.3.3). To simplify things and reduce costs, vendors are implementing systems that use the original MGCP signaling proposal.

4.3.3 ITU H.323

The earliest VoIP standard was the ITU H.323 standard, which evolved from H.320—the standard for video conferencing. H.323 facilitates multimedia conferencing over packet-based networks. As such, it offers a complete suite of protocols for audio, video and data conferencing. The H.323 standard contains the following modules:

Terminals: Telephones (software or hardware)
Gateway: Translates between packet and telephony media streams
Gatekeeper: Performs address translation, admission control and band width management
MCU: Multi-point conferencing unit, which supports multi-party conferences for voice and video

In addition to IP voice communication, H.323 supports collaborative applications, such as white-boarding and video conferencing. This has important implications for IT organizations:

1. Protocol stacks are large and therefore, expensive to develop
2. End points must accommodate more than just basic telephony signaling to be compliant

Since its original introduction in the mid 1990s, H.323 has gone through six versions and provided an important foundation for the VoIP community, but has lost ground to the newer SIP and MGCP protocols. An umbrella standard, H.323 groups multiple sub-standards together into a single specification. Because the actual standards documents are cross-referenced with each other, they can be quite challenging to new readers. This list clarifies roles and responsibilities of the main H.323 components:

G.711 Codec: Pulse code modulation (PCM) of voice frequencies
G.723.1 Codec: Dual-rate speech coder for multimedia communication transmitting at 5.3 and 6.3 kbit/s
G.729a Codec: Coding of speech at 8 kbit/s using Conjugate-Structure Algebraic-Code-Excited Linear-Prediction (CS-ACELP)
H.225.0: Call signaling protocols and media stream packetization for packet-based multimedia communication systems
H.245: Control protocol for multimedia communication
H.323: Packet-based multimedia communication systems
H.248: ITU equivalent of IETF MEGACO
H.450: Generic functional protocol for the support of supplementary services in H.323

Consider the many technologies at work in an H.323 call. First, when a person picks up the phone and dials, an H.323-enabled phone uses H.245 to negotiate a channel and exchange capabilities. H.225.0 handles call signaling and call set-up, and finally, the registration/admission/status (RAS) component channel signals the gatekeeper that coordinates calls within the zone. In version 2 of H.323, a mechanism was established for setting up a call more rapidly, because version 1 took too many steps (particularly from the perspective of carriers and service providers). The new approach is Fast Connect (Fast Start).

If the destination is on or over the PSTN, a gateway must be used to translate H.323 packets to circuit-switched telephony. Although technically transmitting voice over packets has been feasible for decades, H.323 has served an important role in establishing an early framework for how this might be achieved.

Since its first release, H.323 has been constantly enhanced. This has led to a new challenge with H.323 – as the protocol has evolved, not all vendors have been able to keep pace. At this time, H.323 version 6 is approved, yet vendors are still supporting versions 2, 3 and 4, which are not necessarily backward-compatible.

4.3.4 Codecs

A codec (an acronym for coder/decoder) device converts sound from an analog format to digital/numerical representation in the form of ones and zeros. Codecs may also handle compression and decompression. Many codec options are available, and the choice often involves a tradeoff between voice quality, processing speed and data size. Figure 7 includes some examples of codecs. Each is listed with the associated data stream produced, compression delay and quality of the voice stream as measured by mean opinion score (MOS). This is an open test, where a variety of listeners judge the quality of voice samples on a scale of 1 (low) to 5 (high).

Compression Method	Bit Rate (kbps)	Delay	MOS
Wideband	128	1	4.8
Linear (no compression)			4.5
G.711	64	0.75	4.1
G.723.1	5.3 or 6.3	30	3.65
G.726	32	1	3.85
G.728	16	3 to 5	3.61
G.729a	8	10	3.27
G.729a (annexe b)	*Includes comfort noise for silence suppression*		
G.729b	*Includes silence suppression*		

Figure 7: Codec comparison

Determining which codec option is best for your company depends on your requirements. For example, G.711 provides excellent voice quality, useful for call center or sales environments. In contrast, G.729a might be preferable for WAN voicemail, where bit rate is more important than voice quality. The codec war has more or less died down, with vendors selecting and licensing the codec that best serves their purposes (cost and technical). Part of the reason is that network capacity grew to such an extent that network performance is no longer a bottleneck (particularly on LANs). Vendors took advantage of this situation by introducing systems that utilize wideband codecs, which significantly improve voice quality without affecting either throughput or cost. Systems that support wideband codecs actually improve voice quality beyond that achievable by the digital PBX.

Another codec issue is transcoding. Let's imagine two parties talking – one on a global system for mobile communication (GSM) or other cell phone, and the other using an IP phone system, so the call must transit a gateway. If the IP telephony party is using a typical VoIP codec, let's say G.729a, then successfully communicating with the other party requires decoding the G.729a stream and transmitting to the PSTN, which recodes it with a GSM full-rate or half-rate code. Each time the call is coded or decoded, at least 12 milliseconds of delay is introduced (often as much as 50-100 ms). Clearly, this kind of transcoding should be avoided, because it impacts overall, end-to-end delay, as described in earlier sections of this guide. Systems from a single vendor are designed to make the correct codec decision at call set-up, but network designers must be diligent. As the telco/service provider market evolves, the enterprise system will likely be required to interact with service provider equipment.

4.3.5 WHICH STANDARD?

When selecting an IP voice system, choose a vendor that has committed to certain standards. The following standards have experienced a significant level of adoption in today's market:

1. H.323, with its roots in ISDN-based video conferencing, has served its purpose of helping to transition the industry to IP telephony. One of the biggest challenges of H.323 is the number of versions of the standard: six. Some of these versions must be supported in parallel on the same equipment – increasing the cost and complexity of this standard.
2. SIP is ideal for IP voice and will play an important role for next-generation service providers and distributed enterprise architectures. However, SIP suffers from some of the limitations of H.323 in that it has become a collection of IETF specifications, some of which are still under definition. Another similarity with H.323 is that SIP defines intelligent end points. Vendors have found this approach to be more costly and less reliable. However, the biggest concern with SIP today is its poor security record (see Section 5.4.1 and 5.4.3 on VoIP security).
3. In contrast to SIP, the MGCP/MEGACO standards centralize the control of simple phones. This is popular in environments where both cost and control are important issues, which is certainly the case in enterprise environments, where PCs can be used to augment features and functionality. Moving forward, the market will likely support multiple standards for IP voice communication, with certain standards optimized for specific areas – such as carrier markets or communication with end-point devices. But as previously mentioned, the current trend toward delivering presence along with instant messaging using SIP as the transport make it a strong contender for delivering application-rich voice systems.

5. DEPLOYMENT ISSUES

Having described the fundamentals of VoIP and underlying IP infrastructure, let's address deployment issues.

5.1 LEGACY INTEGRATION

In most cases, enterprises own legacy PBX systems, so the migration to IP voice communication will be carried out in steps. This may be for financial reasons, because PBX assets depreciate over seven- or even ten-year periods, and a large number of legacy PBXs were sold in the lead up to the Year 2000 (Y2K). Equally important, the logistics involved in deploying simultaneous, multi-site cutovers can be daunting. Although it may be only a temporary requirement, the new IP-PBX must interact with various types of legacy systems.

5.1.1 BASIC CONNECTIVITY

The easiest way to connect two PBX systems is with a digital trunk, like T-1 or E-1. Some legacy systems expect a trunk to act as the local telephony exchange. To handle the broadest range of connectivity scenarios, the IP PBX should be able to emulate the network side of a trunk connection (network primary rate interface or PRI).

After linking the IP-PBX to the legacy system, begin the process of configuring the legacy system to deliver extension-to-extension dialing between the two systems. If the IP-PBX has been designed with ease of deployment in mind, then setting up the extension-to-extension mappings will take a matter of a few minutes on the IP-PBX side.

After verifying that it's possible to dial extensions between each system, verify that other features are implemented correctly. Caller ID, for example, can be forwarded from one voice system to another, so the caller's identity can be displayed on the handset. Over the years, legacy vendors have pushed proprietary solutions for delivering these features, but standards like ISDN-PRI and Q signaling (QSIG) provide the necessary signaling and are implemented in most PBX and IP-PBX systems.

The final decision concerning connectivity is which system owns the trunk connection to the PSTN. There are three ways to construct the trunk connection:

Legacy PBX owns the PSTN trunk
IP-PBX owns the PSTN trunk
Both systems have a trunk connection

The approach selected often depends largely on budget. If the line card used for the PSTN connection links the legacy system to the IP-PBX, then an additional line card must be purchased for the legacy system. However, most organizations prefer to purchase equipment that works even after the legacy switch is decommissioned, and this implies that additional line cards should be purchased for the IP-PBX, which can also act as a trunk gateway.

5.1.2 VOICEMAIL INTEGRATION

You must decide how the IP-PBX will interact with the legacy voice-mail system. There are two approaches to address this requirement:

1. Keep them separate: The legacy PBX and the IP-PBX systems have their own separate voicemail systems. Voicemails may be passed between the two systems using Audio Messaging Interchange Specification (AMIS) or Voice Profile for Internet Mail (VPIM) protocols.
2. Host voicemail on a single system: Either the legacy voicemail system is preserved for all users, or they all move to the IP-PBX voicemail system. The Simplified Message Desk Interface (SMDI) is the appropriate protocol for this scenario.

First, let's look at the two protocols available for interconnecting separate voicemail systems:

1. AMIS mimics the handset's DTMF signals to forward voice mails between two AMIS-compliant systems over trunks. Of the two approaches, AMIS is simpler to implement on the legacy system and requires significantly lower investment.
2. VPIM provides a standard Multipurpose Internet Mail Extensions (MIME) encoding, so voicemails can be sent as multimedia emails over LAN/WAN data connections.

While VPIM is the more advanced standard, it is often sold as an add-on feature and therefore requires considerable investment to implement. When migrating to IP voice, you may need to make decisions regarding which system houses the company-wide voicemail storage. No matter which system is finally selected, message-waiting lamps on the alternate system (IP-PBX or legacy PBX) must be activated and deactivated. In other words, the IP-PBX handsets should light up when voicemail arrives, even though that voicemail is stored on the legacy voicemail system. Similarly, legacy phones should indicate the availability of voicemail on telephone handsets connected to the PBX, even though the actual voicemail resides on the IP-PBX system. The SMDI protocol was developed for serial connections (TIA-232 cables) between voicemail and PBX systems. The protocol signals the availability of voicemail for a specific extension to the legacy PBX system's management. The legacy PBX uses this information to signal the handset that it should light a message-waiting lamp. After the message has been played, the lamp is switched off using the same mechanism. For some PBX systems, an additional external interface may be required.

5.2 SUPPORTING VOICE QUALITY (QOS) IN THE NETWORK

Voice quality is based on user experience and expectation. In general, today's business telephony networks have virtually no noise, echo, or delay. In IP voice communication networks, noise and echo problems can be easily addressed, but delay may still be a problem.

The ITU specifies an acceptable delay of not more than 300 ms round trip, or 150 ms one way. Delays of more than one-fourth of a second (250 ms) are noticeable and unacceptable. In the U.S., acceptable delays in the PSTN are below 100 ms. Cellular phones cannot deliver this level of delay quality, but users accept the delay in exchange for improved mobility. There are several sources of latency in transmitting voice over IP networks. The first comes from encoding analog voice into a digital data stream. Remember that each codec has an associated data stream, processing delay and voice quality rating. After the voice stream is digitized, it must be packetized and transmitted onto the network.

The second place where latency can be added to the overall delay is when the voice is sampled. This is usually in 20 ms intervals, but that is not a given, and transmission speed depends on both line speed and the packet size that results from the sample rate. For example, transmitting a 64-byte packet onto 56-Kbps line takes eight microseconds (ms). The formula: 64 bytes x 8 bits/byte = 512 bits/56,000bps = .0008 seconds). However, transmitting a 64-byte packet onto a 10 Mbps LAN takes only 51 ms. In today's world of wire-speed, switched-LAN infrastructures, transmission delays introduced by switches and routers are negligible. Delays are more likely to occur at the WAN access point, where line speeds are low and the potential for congestion is high. To ensure voice quality, determine that there is sufficient bandwidth. As packet sizes get smaller and compression techniques improve, the packet header portion takes up the majority of what is transmitted. Using G.729a compression, a single voice conversation is likely to require a minimum of 26 kbps, including the IP packet overhead.

Remember that sitting on top of IP are two transport layer protocols: TCP and UDP. Note that where voice is concerned, it is wasteful to pause in the middle of reconstituting the media stream to re-request a packet that failed to arrive – better to discard it and move on. In practical terms, this means that UDP rather than TCP (which has error correction) is the optimum transport protocol for voice. Statistically, there is a high likelihood that the next packet will be similar to the previous one, so simply repeating the packet resolves most situations.

Packets can also be prioritized, so that voice traffic is transmitted ahead of other types of less time-bound information. This can be accomplished in several ways.

One method is to prioritize IP packets based on source or destination address, such as the IP call server's or softswitch's address, rather than the addresses of every end system. An easier method is to prioritize packets based on predefined UDP port numbers. After the packets are identified, the WAN access router can prioritize accordingly. Unfortunately, H.323, MGCP and SIP dynamically select destination port numbers from a range, which makes routing more difficult, because all of them have to be tracked. Further, it raises concerns that the security firewall is inadequate. Some of the better VoIP implementations send all time-sensitive voice packets to a single destination port, thereby solving this problem. As you consider how to prioritize voice traffic over your network, give some thought not only to which approach is more sophisticated, but also how easily it is to manage. In general, a simpler approach is the least costly.

This list includes technologies that can have a positive effect on real-time traffic placed on the network:

1. Classify the traffic: IEEE 802.1p and Differentiated Service Code Points (DSCP)
2. Segment the local network using VLANs: IEEE 802.1q
3. Establish policies for which traffic should be handled first
4. Queue traffic and prioritize in the access router: random early detection (RED), Weighted Random Early Detection (WRED) and weighted fair queuing (WFQ)
5. Traffic Shaping: delay the way traffic is loaded onto the WAN connection, in order to ensure timely transmission for priority packets (avoid bottle necks in the network) and reserve bandwidth for real-time traffic
6. Caching: Store Web and application content – even the applications themselves – locally, to avoid repeated transmission of the same material over the WAN
7. Map LAN traffic classification to the scheme used by the service provider: DSCP or MPLS

The means to classify traffic has been available in the IP header from inception, but was rarely used because real-time traffic is not usually transmitted over IP. In fact, the IP header contains a Type of Service (ToS) field that indicates priority or handling characteristics (e.g., high priority, low loss, low delay).

In private WANs with dedicated leased lines, prioritizing voice over data at the access point is all that is required. In public or shared WANs such as VPNs and the Internet, voice traffic must be prioritized throughout the shared network space. Rather than using routers to classify and prioritize packets at every hop, the IETF is developing several standards to identify and prioritize different types of traffic.

To request resources along the communication path, two approaches are available. The resource reservation protocol (RSVP) employs each intermediate router along a communication path to allocate resources based on current usage and keep track of which resources have already been allocated. This essentially converts the router into a stateful device, i.e., the router is tracking virtual circuits for packets. The challenge of RSVP is that routers were not designed to for this function. After unsuccessful attempts to make RSVP scale for large deployments, an alternative was defined, DiffServ (DSCP). In this scheme, different packets are "labeled" based on their desired behavior and are managed accordingly by each router along the path. This type of hop-by-hop decision making is much closer to the way routers were designed to work. However, because of the lack of quality guarantee, this approach was considered a Class of Service (CoS) rather than a QoS solution.

While it is true that RSVP encountered scalability problems, another scheme called MPLS successfully uses RSVP Traffic Engineering (RSVP-TE) to reserve bandwidth for whole trunk groups, and has been widely deployed by carriers and service providers.

MPLS introduces an additional mechanism to improve quality. (The protocol gets its name, multi-protocol label switching, from the fact that it works with a number of protocols: IP, ATM and frame relay.) Without MPLS, routers must perform a lookup on the header of each packet that enters the router. With MPLS, a label is applied at the edge router to each packet in a flow, and subsequent routers along the path simply forward the packets along a predetermined path, without wasting time examining the full packet header. This process is carried out at a lower or less sophisticated level of the router's execution so it uses up fewer processing cycles – in short, it is faster. The labels are similar to those proposed in DiffServ. MPLS can be combined with another device, the packet shaper, which has knowledge of specific applications (SAP, voice, video, HTML, e-mail, etc.), as well as individuals and organizations. By defining packet shaping policies that link to MPLS labels in the routers, voice traffic can be assigned capacity on route paths, so the voice quality is maintained no matter what other types of traffic are being generated on the network.

In summary, QoS must be evaluated based on user perceptions of call quality. That quality is influenced by codec selection, echo cancellation and silence suppression. In addition, latency above 100ms sounds uncomfortable for both parties. Network delay on the WAN introduces additional delay, so simply throwing more bandwidth at the problem will not necessarily improve the quality. For this reason, technologies such as WAN optimization and protocols like MPLS, as well as packet shaping solutions, should be evaluated for high traffic links.

5.3 RELIABILITY

For both a legacy PBX and a next-generation, IP voice communication system, the issue of reliability is dependent on the system's ability to ensure access to dial tone, voicemail, administrative functions and value-added applications. The most important considerations include:

- Where is dial tone and call processing accomplished?
- What operating system is used by this device (Windows or embedded real-time)?
- What does it cost to protect this device from failure?
- If the WAN goes down, can I still make phone calls? Will the system failover automatically to the PSTN?
- If the device providing call control to the IP phones fails, is the overall system intelligent enough to cause the phones to failover to another call control device elsewhere on the network?

(Note: Embedded operating systems for real-time devices are employed in mission-critical applications, such as vehicle breaking systems, control systems for airplanes and pace-makers. They are designed and tested to run without interruption for years. Be aware of which operating systems are used to deliver dial tone to these devices.)

The answers to these questions will indicate the architectural robustness of each system being reviewed. For some vendors, reliability comes at additional cost, which will be incurred as soon as senior managers find they are unable to make phone calls because of an OS virus, for example. It is worth getting this right from the start, so spend time looking carefully at the architecture of each system and base your decision on which one offers cost-effective, reliable call control.

Moving beyond issues of fundamental architecture and operating system selection, the best way to ensure reliability of IP voice communication is through redundancy. This can be done with a server that has redundant processors, hard-drives, power supplies and fans. It can also be achieved externally, through redundant servers and sub-second failover between systems. Just as important, the IP voice system's architecture should allow you to distribute voice communication functions throughout the network, so a user on the network can place a call without needing help from a central voice call manager.

It is also important to design the IP network backbone for resiliency, in case one of the networking devices fails. While Layer 3 Ethernet switches and routers are capable of re-routing around failures, a delay of even a few seconds is long enough to terminate most phone calls. An alternative technique is to use trunking or link aggregation, which delivers sub-second failover between physically redundant links. Several networking vendors extend this technique, connecting redundant links to parallel systems.

	Typical	Advanced
Project	*Place a value/cost on downtime. Identify and eliminate single points of failure. Document actions/ procedures in case of emergency.*	*Determine whether this should be part of a broader disaster recovery program. Hire a third-party to audit reliability measures.*
Power	*Provide a UPS for server, PBX and network components.*	*Install a backup power generator. Separate power cables should enter the data center from different sides of the building.*
Servers	*Eliminate desktop PCs that are being used as servers. Utilize server mirroring and load balancing. Install multiple server NICs with load balancing and failover.*	*Store archives in a secure site outside the building where the computers reside.*
Data backup	*Establish a periodic (daily) backup schedule. Carry out full backups prior to any system changes. Archive backups (weekly).*	
Network	*Select network switches designed for reliability:* • *Redundant backplanes* • *Redundant power supplies and fans* • *Support for hot-swappable components* • *Redundant links between network devices that support rapid (sub-second) failover* • *Management software with remote notification of alerts*	
Personnel	*Train staff, including first aid and fire supervisors, on emergency response. Place a phone (attached to PSTN-powered line) in the computer room.*	*Require at least two people to be in computer room after work hours.*
Access	*Establish access control for the computer room.*	*Build protective grids around racks. Install video monitoring in computer room.*
Environmental	*Review computer room air cooling system to verify it can handle the heat generated by machinery.*	
WAN and Telco	*Establish multiple links to PSTN and critical WAN links, preferably with multiple carriers. Configure dial plan and router software to support cutover to redundant links. Install secure firewall.*	
Voice System/PBX	*Link "stay alive" voice switch ports to emergency phones distributed around the premises. Use a distributed voice architecture that limits failure points. Comply with E.911 requirements to provide location information for emergency services. Deploy phones that take power from the voice switch.*	
Service/ Support	*Subscribe to service contracts that provide the appropriate level of critical support.*	*Establish service level agreements (SLAs) with clearly defined penalties that are links to lost revenues (insurance contract).*

Figure 8 is a non-exhaustive list of things to consider when designing a reliable communication network.

5.4 SECURITY

Many of the reliability issues described in Figure 8 also apply to the issue of security. This section takes a quick look at network attacks and telephony fraud to make you aware of the issues, rather than prescribe solutions to specifics. Many companies focus on securing themselves from outside abuse, but may overlook breeches of security that occur from inside the organization.

5.4.1 TELEPHONY SYSTEM SECURITY ISSUES

With today's security spotlight firmly pointed at IT systems, it may surprise you to learn that phone hackers have been around for decades. A good example is the amazing story of Joe Engressia. In 1957, he discovered that when he whistled a certain note (the fourth E above middle C) corresponding to 2600 Hz, it caused a telephone switch to reset its trunk, effectively enabling him to use the switch as he wanted.

Enterprise voice systems have also suffered from misuse. However, a good accounting package should be able to pinpoint the worst cases of phone service misuse. Often, such problems can be resolved by a discussion with the manager. It is important to demonstrate to employees that the system is monitored and abuse is acted on. Some examples of things to watch out for:

- Automatically forwarding an extension to an external number in such a way that the extension can be used from outside the organization for free long distance phone calls.
- Party line or kiosk services (for example, sex chat lines) can also be abused by employees. The use of these high-cost services can lead to out- of-control (and unbudgeted) phone expenses. The solution is simple – ban those numbers and monitor the employee. The challenge is to discover these abuses in the first place.
- Intrusive and time-wasting telemarketing calls to employees (fax or voice). The solution is for employees to request removal from the caller's list.
- Features that facilitate day-to-day use of the voice system can be misused by the intended users or auxiliary staff. Inappropriate calls can be inhibited automatically on a per-user or time-of-day basis. Some organizations resolve problems by issuing personal ID numbers that must be used for personal calls. However, this can irritate employees, so you must determine the appropriate balance between protecting the organization versus trusting the employee.

5.4.2 NETWORK AND COMPUTER SECURITY

Unauthorized entry into corporate networks and computer systems can result in downtime and lost information. As a result, many organizations have evaluated the cost of such attacks and established protection against them. Total security is probably an unattainable goal, so try to strike a balance between securing the network and what it costs to do so.

There are many sources of information on network security and many tools to help you identify known vulnerabilities. A useful Web reference is maintained by the Computer Emergency Response Team (CERT) at www.cert.org. The CERT Coordination Center (CC) is a major reporting center for Internet security problems. Staff members provide technical assistance and coordinate responses to security compromises, identify trends in intruder activity, work with other security experts to identify solutions to security problems and disseminate information to the broad community. The CERT/CC also analyzes product vulnerabilities, publishes technical documents and presents training courses. The Internet was designed as a peer-to-peer network, where any connected device could see and directly address any other device. Over the years, hackers have turned this design to their advantage. But the problem is not unique to IP. When a computer system is connected to a network, it becomes vulnerable to attack – no matter what protocol is used.

Where IP is concerned, several techniques are used by hackers to gain access to your network. A variety of tools (some of which are not covered in this guide) are available to help address security challenges. For now, this list provides a summary of some common techniques:

- IP Spoofing – One of the earliest techniques, where packets appear to be sent by a trusted source address or port number.
- Denial of Service (DoS) Attack – Several techniques that can be used to bombard the network with enough packets to disable network elements or gain access to the network.
- Synchronize Message Flood (SYN Flood)– A type of DoS attack where a target machine, often a router, is flooded with TCP connection requests, resulting in either a slow down or complete crash of the target machine. Once the IP addresses are known, a hacker can then attempt to penetrate these computers using known security holes, default pass words, etc. Phone systems that rely on Windows for call management are vulnerable to Windows security threats and thus require a higher level of vigilance.

The first level of defense is a firewall (see Figure 9), preferably with integrated Network Address Translation (NAT) to ensure internal IP addresses are not visible to the outside world. NAT acts as a source address proxy for internal machines. RFC 1918 establishes that the address ranges 192.168.*.*, 172.16.*.* through 172.31.*.*, and 10.*.*.* are not routable over the Internet. Internal machines with these addresses must use an address translation mechanism when they want to visit a Web site and their request transits through the NAT machine, which keeps track of their internal addresses but substitutes its own addresses. The firewall also integrates a filter, so that the only routes in and out of the internal network are via the NAT machine. Just outside the firewall, an area is often set up called a de-militarized zone (DMZ) where systems are located that interact directly with the Internet (e.g., public Web server, e-mail forwarder). A gateway firewall protects internal systems from any flow of information that is not between predefined systems and in predefined directions. Such rules might be:

1. Internal systems can get to Web sites, but only if they transit via the Web proxy, so that source addresses are hidden
2. E-mail that is being sent must transit via the email proxy, which also disguises the source address information
3. E-mail that is being received must first go to the proxy, then takes a predefined path between the proxy and internal e-mail server
4. No e-mails trying to bypass this route can get through the gateway
5. Additional functionality can be added in the e-mail proxy, to open and verify incoming e-mails for known viruses

With such rules in place, the task of monitoring security can be focused on specific machines. The security of the proxy and gateway devices can be locked down, so that they do not accept network logins or other means of penetration. Any attempts to log in or attach to these systems should generate alerts or even send misleading information to the potential hacker. An obvious way to accomplish this is to disguise the OS and vendor names with misleading names and banners.

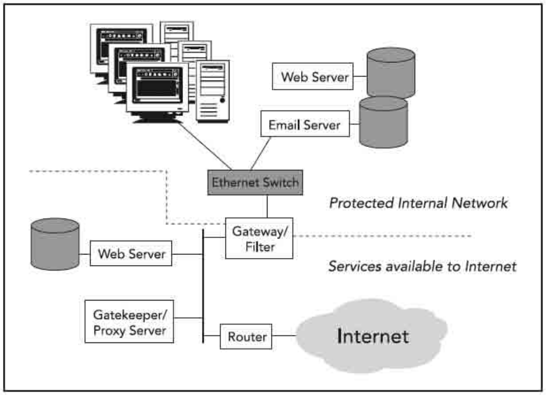

Figure 9: Network Security Firewall

Personal and small business firewalls are now available that are ideally suited to the teleworker and branch office environment, where IP voice travels over a public network connection. For public networks or VPNs, the same techniques used to secure data (IPSec, data encryption standard or DES, and triple DES encryption) may be used with voice. The only issue that needs to be considered is the potential delay associated with encryption techniques. However, many vendors provide solutions capable of delivering wire-speed encryption filtering and forwarding.

Intrusion Detection System (IDS): In spite of our best efforts, it is still likely that enterprise network security will be compromised at some time. Most frequently, this occurs because a user somewhere inside the corporate network has managed to infect her PC – this may have occurred because she used a portable PC on a public network, and did not have time to do a virus scan before connecting at the office. In any case, the fact of the matter is that there is now an intruder inside the network – and within seconds, he may have gained knowledge of internal systems and communicated this information to a cracker located outside.

Our next line of defense is the IDS. This monitors the traffic inside the firewall, to determine whether security has been compromised. There are essentially two types of IDS:

- Statistical: Establishes a profile of typical (uncompromised) network traffic, and uses statistical analysis to determine whether traffic patterns imply something unexpected is happening on the network.
- Signature-based: Employs a set of rules describing packet contents – or packet signatures – which identify known attack sequences. An open source IDS engine is available at www.snort.org, and Snort rules are maintained at www.bleedingthreats.net.

The IDS can be configured to act as an Intrusion Prevention/Protection Systems (IPS). In this capacity, it blacklists the source of offending packets – preventing them from communicating with the network for a pre-defined amount of time. This can have an adverse affect on user satisfaction, because intrusion prevention is more of an art than a science and can create false positives – benign packets that are blocked because they have a profile or signature that matches up with a rule.

Network Access Control (NAC)
Network Access Protection (NAP)
Since 2006, there has been a shift in the way networks are designed. They have become more intelligent at the edge – where your PC is connected. This might be a wireless access

point or an Ethernet cable connected to a switch, but the current thinking is that more analysis of the current state of the user's PC should take place, before the network allows that PC to actually connect to corporate resources. In the example earlier in this section, if virus scans need to be carried out, then this task (as well as a virus update) are imposed on users before they can go any further. Two communities are currently driving such architectures, with broad support from vendors for one or the other model. It is clear that the network infrastructure is taking on more ownership for general, application-level security, and this trend will continue.

It is worth mentioning that many security issues are general, and not specific to IP telephony. However, voice communications exposes any fundamental weaknesses in the design of the data infrastructure. Ask your vendor for guidelines and instructions to tune the IP infrastructure in preparation for the VoIP system, and carry out regular security audits after the system goes live.

5.4.3 VOIP SECURITY/NAT TRAVERSAL

Having reviewed the separate issues of telephony and network security, the next section discusses additional security issues specifically relating to VoIP systems. Failing to address these issues exposes your business to security vulnerabilities that can cause increased cost, loss of confidentiality and even system failure. As we will see, help is at hand both from more secure protocols and also from the careful selection of VoIP system architectures that are designed to avoid single points of failure.

VoIP Security issues are grouped into the following topics:

1. Signaling security: Capturing and manipulating the signaling information of a call, in order to learn about or simulate the user's identity. To address this issue we can:
   a. Encrypt the signaling information using Transport Layer Security (TLS)
   b. Authenticate endpoints, so an attacker cannot simulate a valid user
2. Media stream security: Eavesdropping by listening to private conversations. Media is encrypted at the:
   a. IP Layer: VPNs with IPSEC
   b. RTP Layer: SRTP
3. Denial of Service: Flooding specific components of the voice system until they cannot perform their tasks. A distributed, rather than a centralized, architecture considerably reduces the effect of this type of attack. Also deploy IDS.

As we have seen in the general security section, most enterprises use NAT to hide their internal topologies from the general public. Internal addresses using NAT cannot, by definition, be routed over the Internet. This presents a serious challenge for organizations seeking to deploy SIP as an internal signaling protocol. When a caller sets up a SIP call using the INVITE method, it fills out informational fields using the Session Description Protocol (SDP). One such field – Connection Data – provides the address of the device participating in the call. This is an example of its use:

C=IN IP4 192.168.10.10

IN means it is an Internet call, and IP4 means it is using IP V4, rather than IP V6. Finally, the IP address (four numbers separated by dots) is the phone's address.

Here's the problem – when this hits, the NAT simply changes the actual source IP address to its own address, creating a miss-match. There are various ways to get around this problem. First, these three protocols can be implemented on SIP clients:
1. Simple Traversal of UDP through NATs (STUN) – provides a proxy that delivers address information prior to the call being set up. This requires that STUN be supported in the SIP clients. In certain cases, particularly for large enterprises that use bi-directional NATs, STUN is not enough
2. Traversal Using Relay Nat (TURN) – This solution is very resource- intensive and should be avoided, if possible
3. Interactive Connectivity Establishment (ICE) – Provides a way of selecting whether to use STUN or TURN for NAT traversal

Another approach is the use of SIP Application Layer Gateway (ALG) on a firewall. This solution processes SIP packets as they transit through the firewall – converting addresses as needed. Unfortunately, SIP ALGs tend to be very limited in the features they support, and require enterprises to replace and reconfigure firewalls. Many organizations would find such a requirement too excessive.

Session Border Controllers (SBCs) are deployed by operators to resolve a range of real-time issues related to VoIP sessions (QoS, transcoding, proxy routing, protocol conversion, topology hiding). Smaller-footprint enterprise SBCs that incorporate STUN, TURN and ICE can offer similar real-time processing for NAT traversal.

Because of its ease of use, SIP has very rapidly exposed VOIP security issues. Luckily, new solutions are emerging. But while support for SIP continues to grow on the Internet, within the enterprise it should be deployed as part of an overarching security strategy – and with full understanding of its tradeoffs.

6. TELEPHONY APPLICATIONS

This section moves beyond the basic features and functions of making phone calls to explore features, applications and solutions designed to enhance your business's productivity. The first section provides an introduction to standards and technologies used for integrating applications with the phone system. Next comes an exploration of the applications themselves.

6.1 CONVERGENCE: COMPUTER TELEPHONY INTEGRATION (CTI)

One advantage of IP-based voice communication systems is the ease with which voice can be enhanced with additional information about the parties involved in a call. IPT also facilitates the linking of voice to different enterprise applications, such as e-mail, fax and CRM. In the PBX model, voice cards were installed in dedicated application servers to bridge the two worlds. Alternatively, CT extensions were purchased for every user telephone. Accomplishing this was price-prohibitive and very complex, so generalized CTI never took off in the legacy PBX world. You may encounter standards from the Enterprise Computer Telephony Forum (ECTF), such as S.100, S.200, S300 and H.110, which define standard APIs for developing CT applications. (Note: Most PBX vendors have proprietary APIs.) Some vendors offer an extended CTI model to develop their own CT-rich alternatives to PBXs.

In the IP world, the link between the telephone and computer is already available over the IP network. Open standards, such as the Telephony Application Programming Interface (TAPI), enable programmers to interact directly with the phone system, from within the PC applications they develop.

6.1.1 TAPI

TAPI, Microsoft's Telephony API, allows third-party call control applications to handle telephony functions from the client application. Outlook, for example, is a TAPI-compliant application that enables users to dial the phone numbers of people directly from the Outlook contact database by clicking a button. This means that when using a contact database such as Microsoft Outlook or Symantec ACT, you no longer have to look up the contact information and manually dial from a phone. The TAPI interface lets you select a contact manager and initiate a phone dial directly from your list of contacts. This simplifies dialing and eliminates dialing mistakes.

Unlike the previous releases, Microsoft TAPI, version 2.1 fully addressed client-server needs for call monitoring and control. Its 32-bit architecture supports existing 16-bit TAPI applications. The most recent version from Microsoft, TAPI 3.0, came with Windows 2000. General-purpose applications such as Outlook are not telephony-specific, and most users prefer specialist clients that deliver important productivity gains.

6.2 PERSONAL PRODUCTIVITY

Most VoIP users find that the productivity applications really drive the benefits of a new system. When employees can focus on their jobs, rather than on searching for (and dialing) phone numbers, they are less stressed and more relaxed with customers. Contrast this with the traditional approach. In the PBX world, when joining a new company, employees received:

1. Phone
2. Guide for using the phone's features
3. Company phone directory

In this model, the following factors would impact your productivity with the phone:

1. Your willingness to read the guide and remember how to use the phone's features
2. Your ability to touch-type on a phone's number pad
3. How many phone numbers you could remember in your head

To speed things up, you might ask someone nearby, “Hey, how do you set up a three-way conference again?” and then state to the person on the other end of the line: “Listen if I lose you, could you call me back?” In one breath, two employees are involved, while simultaneously establishing with a customer that we don't value their time enough to learn how to set up a conference call. Enterprise PBXs were designed half a century ago, and things have changed from a technology standpoint, but how do we leverage today's technologies to improve productivity? Can we make things easier when delivering enhanced applications? The PC is a resource that can be leveraged to improve the voice system usability and increase productivity. These are some of the features to look for in an IP-PBX:

Accessing Features with Personal Call Control: In an IP-PBX environment, the user's PC can be logically linked to an analog, digital or IP phone, enabling the PC to interact with the phone. The application controls the phone, and the phone can pass information to the PC application. The IP network provides the “glue” between these two devices.

The model can be extended so that all features are readily available through buttons and menus, eliminating the need for the legacy model's handset guide. The full range of features is now easy to use and accessible.

Call History: The system can also track the user's call history – both incoming and outgoing calls. In this way, the phone system can be used much like an automatic log or notepad. Using the call history, there's no need to search for a number again.

Simply scan through the history and click dial. The number can even be added to our contact list (as explained in the next paragraph).

System-wide Directory Services: Production of the organization's phone book could be an expensive activity, so in the past it was done in full only about once a year. Employees really need a dynamic, always-current contact list that incorporates everyone they interact with – not just other employees, but also suppliers and customers. This information should also include cell phone numbers.

A well-designed IP-PBX provides a system-wide directory that combines personal and corporate contacts in a single database, which can be easily located by typing in fragments of people's names or numbers. Users can type in first names, last names, initials, names that sound similar, departments or any other criteria that makes sense. Database searches like this make telephone usage much easier, because they are dynamic and always up to date. Users are more productive as they search for colleagues and even external contacts.

Application Integration: Personal calendars integrated with the voice system could be employed so that call handling modes are set based on knowledge of the user's meeting schedule. For example, an employee is in a meeting, so her calls are automatically sent directly to voicemail. CRM client applications can be enhanced with call buttons that allow agents to dial directly from the customer record, thereby improving agent productivity.

Presence: The concept that if users have recently typed on the keyboard or moved the mouse, then they are located close to that computer is a model that is gaining support in a wide variety of applications: messaging, cell phone usage and voice systems.

The real-time knowledge of people's current status provided by a combined voice/data system can be used to determine how a call should be routed – before the call is made. For example: the individual appears to be at her desk, but is on a conference call –forward the call somewhere else.

With real-time knowledge of the phone's state, calls can be routed to a person who is at their desk and can take the call immediately, rather than forwarding valuable customers to an absent person's phone (where it ends up in voicemail). The net result is a more responsive and productive organization. There are two presence standards: SIMPLE, which is based on the SIP protocol, and the XML-based XMPP. SIP was originally designed as a VoIP standard, so most implementations of presence for voice systems use SIMPLE.

Web Conferencing: After a call is in place, we have immediate knowledge of the PCs logically linked to the participants of the call. The system tracks the state of the phones and associated computers. Because of this, if participants in a call or a conference want to review a document they are working on, a user simply clicks on a button and the other participants immediately see the document. No additional set up is required, because the overall system has knowledge of what each user is doing and who is attending the conference. When people do not have to figure out how to accomplish tasks with technology, productivity is increased.

In fact, the standard TAPI programmer's interface can be used to customize and enhance the system, so there is no limit to the time-saving applications that can be delivered. In fact, to determine how TAPI-compliant your vendor is, ask whether they use TAPI to deliver their own applications. The answer to this question says a lot about how committed they are to your business productivity.

6.3 COLLABORATION

In the old days, the corporation was a brick and mortar building where employees went to work. In fact, it was for this application that the original PBX was designed. However, things have changed dramatically:

1. Enterprises nearly always have more than one site
2. Sales staff are expected, even encouraged, to spend most of their time with customers and often don't have an office with a desk
3. Many key contributors to an enterprise actually work for supplier organizations
4. Customer service requirements have increased so dramatically that the successful enterprise must seamlessly collaborate with their customers. From specific phone calls to formal meetings, the voice system is a critical foundation for addressing these new requirements:
   - Establish computerized phone directories that include customer and supplier contacts
   - Enable conference calls (on the fly or scheduled) where off-site employees, customers and suppliers can easily participate, sharing documents and other information
   - Provide recordings of audio conferences so that absent team members (employees, suppliers and customers) can review actions and decisions.
   - Deliver distributed applications so customer interaction can be spread across multiple sites and can dynamically include any employee at the click of a button
   - Offer presence management integrated with company- wide scheduling applications, so you know whether an employee is online and available, or busy with a customer on the phone – before you try the number

- Make people easier to reach with follow-me find-me rules, so their customers can find the right representative by typing a single extension number
- Enable sales staff located temporarily at their home offices or customer premises to remotely leverage the full range of features. For example, with softphones on their portable computers

Of course, all of these features should be easy to use, so employees can spend their time doing their jobs, rather than figuring out how to schedule a conference call. With renewed focus on collaboration between employees, customers and suppliers, the IPBX moves beyond the monolithic single-site systems of yesterday to address the needs of the distributed organizations that our businesses have become.

6.4 VOICEMAIL AND UNIFIED MESSAGING

A voicemail system is used for store-and-retrieve voice communication. Initially designed to replace proliferating answering machines, the voicemail system has taken on a broader role of bringing the advantages of ubiquitous communication tools like e-mail to the telephony world.

The basic process of using a voicemail system is as follows:

An unanswered incoming call is redirected to the voicemail account of the called party, where the called party's prompt is played to the caller. After a beep is played, the voicemail system begins recording. The caller leaves a message, then presses a predefined key for options or simply hangs up. The called party receives an immediate notification on a predefined device (pager, message light on phone, e-mail). The called party dials into a central voicemail number or simply presses a voicemail key on the desktop phone to access the system. The system prompts for number and ID number, and then provides menu options. The called party listens to messages, then deletes, saves, forwards or replies to them.

Beyond the basics, certain enhancements are available. For example, in a multi-site environment, this functionality can be extended to include the ability to forward voicemails between sites using low-cost communications links. Also, unified messaging can provide a single in-box for all message types (voicemail, e-mail, fax).

6.4.1 UNIFIED MESSAGING

With a typical voicemail system, the challenge is that it takes time for users to learn the multi-layer menu system and where the shortcuts are located. Users often must keep a visual representation of the menu layout next to their phones.

By providing a unified view of all media in the e-mail system, users can leverage a more intuitive interface and enjoy the freedom to handle and store messages in the same way as they do e-mail. For example, users can view and listen to voicemail messages in any order they choose, instead of tabbing through a series of messages to get to important message number 11. The key to unified messaging is establishing a very tight relationship between PC and voice system. In the past, this kind of integration could be achieved only through the use of separate, loosely linked systems – a PC full of voice processing cards that effectively provided a bridge between the worlds of voicemail and e-mail.

With IP telephony, this problem is solved. Because it is integrated by design, the IP network effectively provides a common link between any network component: e-mail server, PC, voicemail and phone.

6.5 SUPPORTING TELEWORKERS AND ROAD WARRIORS

During the past few years, teleworking has emerged as an increasingly important component of the distributed enterprise workforce. This trend was driven largely by the need to retain skilled employees by allowing them to work non-standard schedules at locations outside the office.

Communication and collaboration play an essential role in employee productivity, so the ability to support teleworkers with both voice and data network integration is becoming a critical requirement for IT organizations. While most teleworkers have some level of data network access and integration from their home offices, they have historically lacked voice network integration. Softphones can deliver the full range of telephony features and enable employees to feel like active members of their organizations, even if they are geographically far away. Such phones run on users' computers, which might connect through home offices, customer sites or coffee shop WiFi hotspots.

This type of flexibility was never available with legacy PBXs, and demonstrates the advantage of a distributed approach to voice communication.

6.6 MULTI-SITE CONNECTIVITY

One of the main advantages of IP-based voice is that the communication system goes where the IP network goes. It's that simple. This allows companies to extend their voice networks from LAN to WAN, central site to remote office, or even to home offices. The system is totally location-independent, yet it functions as a single, cost-effective network.

These presence and messaging capabilities add considerably to employee productivity.

When a system is designed as a multi-site application, rather than as a set of independent PBX systems, the user experience is dramatically improved. A single dial plan covers the entire organization irrespective of site/location, and if users travel from one site to another, they simply log into the system, which forwards calls to their new location without the user having to make any configuration changes.

The benefits are even more dramatic when users are located at multiple sites. For example, the ability to know whether a colleague is available before making a call maximizes efficiency and reduces frustration. In a geographically dispersed team, this type of presence information helps decrease the distance gap between employees in different time zones, facilitating better global employee communication. Even when a voice system is distributed across a LAN, metropolitan area network (MAN) or WAN, all of its elements and advanced voice services are easily managed as a seamless, unified whole. IP telephony systems with Web-based management provide a comprehensive, single-system view of all users, sites, equipment, features and services, enabled by a single-system database. This unified approach is in stark contrast to the legacy method, where each site and various voice services (such as voicemail, auto attendant or workgroup automatic call distributor or ACD) are managed as separate entities, with separate databases and interfaces. This archaic approach makes legacy systems extremely complex to manage – not to mention expensive (due to the impact and cost of specialized training and staffing).

6.7 CALL CENTERS AND CUSTOMER RELATIONSHIP MANAGEMENT

The advantages of IP-based voice communication systems over traditional PBX systems for call center or CRM applications include greater flexibility in distributing calls and easier integration of voice and data. As described above, in an IP-based system, voice goes where the IP network goes. In call center environments, calls may be distributed to users anywhere on the IP network. While distributing calls between local and remote offices is possible with traditional PBX systems, the ability to easily move users from one location to another is a strength of IP-based voice systems. This ability allows, for example, a customer support expert to move from the central headquarters to a home office while remaining connected to the call distribution group. This flexibility is becoming extremely important as businesses define their CRM strategies.

The Internet is making this requirement even more urgent. With less business being done in person and more attention being paid to customer satisfaction, it is critical that your voice communication system work with key CRM applications to provide the highest possible level of customer service, which is key to building customer loyalty. Solutions for managing customer contact and ongoing relationships range from informal contact center capabilities to a highly sophisticated call center solution built around an ACD module. The ACD enables queuing and managing of call distribution to the appropriate agents. Calls entering the contact center are initially handled by the interactive voice response (IVR) module, which helps to determine how to best service the customer. Often, the customer is handed off to an available agent with the appropriate information. At that point, a screen pop-up window delivers customer database information to the agent's PC.

Today, the call center is transitioning and taking on a much broader role in the enterprise. CRM enhances call center technology with applications tailored to specific business functions, such as sales force automation and customer support services. The expanded role of CRM is driving improvements in the way call centers (or, more appropriately, call "non-centers") need to work:

CRM must expand to include distributed employees: With the growing number of remote workers, mobile workers and teleworkers, it is important that the CRM system be able to leverage agent skills, regardless of where they sit in the organization. Distributed, location-independent CRM delivered wherever IP access exists enables more employees to take responsibility for managing customer relationships.

Systems must be able to handle any media type: With the rise of the Internet, it is no longer possible to assume that voice (telephony) is the only way of communicating with businesses. E-mail, Web forms, instant messaging and chat are now widely used. All these tools provide opportunities to exceed customer expectations and build loyalty.

Queues must link to agent skills: Customers with a particular interest or requirement do not want to be routed from one agent to another; they want to deal with one person who can provide the assistance they need. Similarly, businesses must use the most appropriate person for each call, based on business objectives (cost, workload and results). Finding the right person to handle the call ultimately costs less and enhances the customer experience.

6.8 IVR AND VOICEXML

Providing telephony-based access to information is now a critical requirement for many enterprises. One of the challenges is to ensure commonality across the various interfaces into back-end information. Specifically, if a customer accesses an IVR system to find out where their package is currently located, then when they sign in on the Web they should get the same information, presented in a format that is as similar as possible.

Standardized in 2004, the VoiceXML language defines voice menus for IVRs. The language has much in common with the eXtended Markup Language (XML), which is used for developing Web pages and is supported by various application generators to make it relatively straightforward to create the interface. This interface to corporate databases is exactly the same for Web applications, enabling the entire system to be managed as a single entity.

Work has recently been completed to define compatible protocols for speech recognition, including the Speech Synthesis Markup Language (SSML) and Speech Recognition Grammar Specification (SRGS).

This means that whether you are pressing the Update button on a Web page, pressing the number three to update your phone, or speaking the word "update" directly into your handset – the request is handled in exactly the same way on the back end. And of course, it is then presented to you in whichever way makes the most sense: screen update on the Web page or audio message for two-telephony interfaces.

6.9 FIXED MOBILE CONVERGENCE

Within the enterprise, fixed mobile convergence (FMC) allows employees to use a single handset for fixed and mobile communication without losing access to the enterprise telephony applications on their desktop phones.

Also, FMC can potentially provide cost savings by eliminating expensive roaming charges for making calls from a cell phone to other mobile networks. These expenses are particularly burdensome outside North America.

FMC can be delivered in the following ways:

1. Find Me Follow Me: Personalizes the call handling feature of enterprise phone systems so it always forwards calls to your cell phone. This approach provides the advantage of a single number which, while useful, does not deliver additional applications or cost savings.
2. Off System Extension with Intelligent Call Handling: An addition to the previous feature. An intelligent call manager client runs on the PC, or is actually on the phone, to control in-call features. This can be achieved only if the voice system has the ability to manage off-system extensions and reach them with outbound calls, as if they were direct extensions. Today, we have the applications (conferencing, directory dialing, park etc.) and have effectively converted cell phones into enterprise-feature phones. Of course, the issue of cost savings still needs to be addressed.
3. Mobile Extension Gateways (MOBEX GW): Cell phone charges vary greatly from one region to another. For example, in Europe, cell phone users are required to pay expensive roaming charges when they travel to another region within the same country and also when they travel to another country. Regular travelers generally try to find packages that offer better rates for roaming. This can help reduce, but does not eliminate, steep charges that are applied for both inbound and out bound calls. Enter the mobile extension gateway. This appliance takes advantage of special rates or free calls between cell phones that are subscribed to the same mobile network. The network manager basically inserts SIM cards for into the mobile extension gateway, attaches the gateway to the enterprise network and routes calls to the appliance using an open standard like SIP. When purchasing a mobile extension gate way, it is important to understand how in-call signaling is handled and passed through to the IPBX (a flash key, for example), because some current-generation MOBEX gate ways now offer their own intelligent clients. This means they may not provide direct access to your usual voice system features, but only a subset. These gateways and intelligent forwarding go a long way toward eliminating the unnecessary separation between mobile networks and enterprise telephone systems.
4. FMC with Bluetooth: Some network operators that do not own mobile networks have begun offering FMC solutions within the DSL access device installed at the customer's premises. The device uses Bluetooth as an air interface for signaling and media streaming. The advantage to the service provider is that revenue is generated for cell phone calls without having to purchase a license or build a cell phone network. The advantage to the subscriber is that calls are much cheaper (possibly charged at a flat rate) when made in the proximity of the access device. The solution requires very little configuration or change of behavior by the cell phone user. However, the system is limited to a small number of users.
5. WiFi and GSM or CDMA with backend handover: Many business users now have mobile phones that incorporate both WiFi and third-generation (3G) signaling. Such phones run the IP and SIP protocols, so they can effectively act as SIP user agents, as well as standard cell phones. This moves the FMC discussion to the next level – it is more or less a given that once registered as a SIP endpoint on the enterprise voice system, a user can select whichever protocol makes the most sense in terms of cost, reliability and quality of service. But the more exciting scenario is the ability to move seamlessly between SIP/WiFi and 3G, even during the call. Imagine you are on a call covered by your enterprise WiFi network. The call is free, but you need to leave the building and head out to a customer meeting. No problem. With FMC, the phone,

network and enterprise phone system share information about signal strength. The call is forwarded through a trunk and merged back with the other party using 3G. Note that there are issues relating to advice of charge, so the calling parting has the option to step back into better coverage or accept the new path for this call, which is then charged at the standard rate.
6. WiMAX IEEE 802.16m – The Worldwide Interoperability for Microwave Access (WiMAX) interface is receiving strong industry support. In particular, the IEEE 802.16m standard is proposed as a viable alternative to 3G. WiMAX uses a similar air interface to WiFi at a different frequency spectrum and with SIP signaling. For FMC, the advantage of this approach is that a single protocol is used for both signaling requirements, premises-based and mobile. It effectively provides a single, end-to-end IP transport that can be used by enterprises for all of their communication needs. This would provide renewed impetus for a range of new voice and video-enabled enterprise applications.

6.10 VIDEO TELEPHONY

In the next ten years, video telephony will play an important role in enterprise communication. It will impact our business travel patterns, the way we build call centers, and already is having a significant affect on how we meet and establish relationships with each other.

This section assumes that you have read the previous sections on audio coding and call setup and focuses on the challenges of adding video to telephony. First, let's review the challenges of video communication as a basis for understanding what the various standards and technologies are designed to do.

How We Interpret Moving Images

Earlier in this guide, the ability of the human mind to construct a complete media flow made up of discrete parts was discussed. It is the same for video. If our eyes receive static images (like photos) very rapidly, at least 10 images per second, then our brains combine the separate images into a single moving picture – a video.

The discrete images in video are called frames. To illustrate the point just made, look at current standards for broadcast transmission: the National Television Standards Committee (NTSC) standard used in North America refreshes the TV image at 30 frames per second (fps), and European Phase Alternating Line (PAL)/ Sequential Color Memory (SECAM) standard refresh rate is 25 fps.

To move from moving images to the discrete pictures that make up moving images, we need to a way of representing colors so that each image appears real. The human eye reacts when stimulated with light waves. The retina is comprised of more than 100 million rods (which react to brightness) and cones (which react to three colors: red, green and blue or RGB, as well as saturation and brightness).

The human eye is capable of processing light with wavelengths between 400 to 700 nanometers. The NTSC standard defines wavelengths of the primary colors in nanometers (nm) as follows:

- Red: 610 nm
- Green: 534 nm
- Blue: 472 nm

The eye produces images that are comprised of more than 100 million small dots, which are combined in the brain with color (hue), saturation and brightness information provided by the cones to construct a complete (and apparently solid) image.

When the image is transmitted from eye to brain, it is done in parallel – meaning the brain receives data from each rod or cone simultaneously. If networks were built this way, the cost would be prohibitive. So they must combine hue (RGB information), saturation (how much of this color is present) and brightness, using enough pixels (or dots) so that receivers know what they're seeing. QCIF uses 176 _ 144 pixels over a network at 25 fps. To accomplish this for uncompressed video requires about 25 mbps. A technology is needed to compress all this information, so it works with today's networks.

6.10.1 VIDEO COMPRESSION

In order to resolve the huge bandwidth problems described above, the following techniques can be adopted:

- Transmit only the parts of an image that have changed since the last frame – differential pulse code modulation (DPCM)
- When pixels (dots) are moving, transmit their direction and speed – Motion compensation
- Compress the image's binary representation by eliminating repeated digit patterns and sending a code and quantity (i.e. 50 pixels of yellowish orange).
- Compress the encoding by using shorter codes for the most frequently used binary patterns

There is an enormous number of video codecs that are proprietary, license-based or free. The following list includes some of these codecs:

- H.261 (1990): Used in early ISDN-based, H.320 video conferencing equipment
- MPEG-1 part 2: Widely supported by PCs everywhere
- MPEG-2 part 2 (H.262): Used on DVDs
- H.263 (March 1996): Significant improvement on H.261, now used for H.323 (has been employed by popular sites such as YouTube and MySpace)
- MPEG-4 part 10 or H.264 (May 2003): The current state of the art, used in various game consoles and the Blu-ray Disc

Display Size

Previously, we mentioned that the amount of data to transmit is a factor of the number of pixels supported by the displaying application on the receiver's device. When setting up a call, it is important to select a display resolution that works over the network we're using, one that will display on the device in use by the called party. In fact, there must be a negotiation of display resolution at call setup time.

SQCIF	Sub Quarter CIF	128 x 96
QCIF	Quarter CIF	176 x 144
CIF	Common Intermediate Format	352 x 288
4CIF	4 x CIF	704 x 576
16 CIF	16 x CIF	1408 x 1152

Figure 10 lists the standard display resolutions currently in use.

Call Setup

Using an intelligent enterprise call manager package, it is relatively easy to set up a video call between two users within the same organization. From within the application, a video call is requested and the appropriate video client is launched on the caller's machine. The called party answers and the video call is in session. To enhance ease of use, the call manager displays which users and conference rooms are video-equipped.

Let's have a look at how video clients might set up this call using the SIP protocol. When the INVITE method is sent from the calling party to the called party, a number of capabilities are described using the Session Description Protocol (SDP). This includes which audio codec is used for the call. In the case of a combined voice and video call, we the audio and video codecs that will be used for the call must be added within the SDP section.

What happens next is that the acknowledgment code (ACK) is received and the video call commences, or a non ACK is received, at which point the caller can continue with just voice, if desired. If the two parties start the call using voice and decide to add video later, they can simply issue a new invite message while maintaining the existing audio call. It's that simple. This ease of signaling is one reason why SIP is so compelling.

H.324m

While it is certain that the flexibility of SIP makes it a contender for more generalized signaling in the future, carriers are now deploying the H.324m standard for mobile-based video telephony. So let's take a look at how this works, because through the efforts of the 3G community, it will play an important role for some years.

Like H.323, H.324 is an evolution of the original H.320 video conferencing standard that has been customized to address low-bandwidth connections. The H.324m standard has been specifically designed for 3G mobile networks. It leverages the existing H.323 protocols for call setup and teardown, H.245 and H.225, and contains mechanisms for converting audio calls into video calls in mid-call.

H.263 (see codecs above) is mandated for H.324, and H.261 is supported optionally. In terms of display size, the QCIF is mandated, with optional support for other CIF-based display sizes.

The Media Stream

After the call has been set up and information is being run through the codec, the rest of the work is similar to what is needed to accomplish voice communication. While the codec is now a video codec, like H.263, the media is still wrapped in RTP. There are a number of IETF specifications describing RTP payload formats, for example, the RFC 2429 RTP payload format for H.263.

Video Applications

From Conferencing to Tele-Presence: Video conference rooms used to be niche applications for large corporations. They were used when high levels of collaboration were required between geographically-dispersed groups. With the Internet, things have changed – personal video conferencing is widely available and freeware applications deliver video conferencing outside corporate control. But a challenge still exists. There is increasing pressure to address and reduce the costs of traveling to meet people, and the technology is already available to address this challenge.

Tele-presence provides sound and video quality that is "as good as being there." With high-definition video, you can literally see the minute changes in facial expression that make up such a large portion of our non-verbal communication. While telepresence still requires dedicated equipment and services, the cost compare favorably to the cost of airplane flights, hotels and ground transportation, making the technology an attractive evolution for the communication industry.

Changes in Expectation: As has been seen in the past, new technologies are often adopted by younger generations. The real challenge for businesses and service providers is to anticipate the adoption cycle at the right time. Video telephony is one such technology. Using the latest codecs as described above, it is quite possible to carry on a video call within the confines of a 64 kbps channel. Costs are still too high for broad adoption, but the use of 3G handsets shows that the installed base is quite ready for reasonably-priced video communication. When it comes to applications, this implies increasing demand for:

1. Video messaging – Expect users to want to leave more than just voice messages. "Hey, look at what I'm looking at right now."
2. Video ring-back tone – Ring-back tone is widely used throughout Asia and is a rapidly growing market. It allows callers to play a tune or message, instead of the standard ring-back tone. The marketing potential of this feature makes it fairly inevitable that businesses will want to feed corporate video to callers.
3. Video of the caller: Call centers will be equipped with video capability, creating a number of implications on future call center designs.

6.10.2 UNIFIED COMMUNICATIONS

By fully integrating enterprise voice communication with the other tools available today – databases, search engines, instant messaging, company knowledge bases – enterprises can make better use of the spoken word as a resource. But early positioning of the benefits of unified communication is firmly focused on reachability and real-time productivity.

Microsoft recently entered the unified communication market, delivering:

- NAT traversal for IM, video and voice – Discussed this at length in the Security section. Microsoft successfully addressed this challenge through the use of TCP, rather than UDP, for transport. While this raises compatibility issues with existing gateways, it is clear that Microsoft is receiving general industry support for its approach
- Web-based conferences (voice, video and applications)
- Client application enhancements to support real-time voice communication. For example, auto-attendant and unified messaging defined within a Microsoft Exchange client
- Enterprise-class instant messaging (secure, auditable communications) that can interact with the broad installed base of IM communities: MSN, Yahoo, AOL and so on
- Enhanced video for telepresence
- IP PBX integration
- Presence management
- Unified reporting of IM, voice and conferencing

Many of these features are already available in today's generation of enterprise IP PBX systems, but enterprises find the notion of an implied single-vendor solution (even through partnership) compelling. While improved integration between the IT and voice world is good news for enterprises, there are still a number of considerations to be weighed while defining your strategy:

1. Budget implications – costs of servers/clients/headsets/gateways
2. Telephony mainstays – powered phones, support for emergency services
3. Reliability – how to evaluate mean time between failure (MTBF) stats for the overall solution and whether the cost of downtime is affordable
4. Cost – in the past, mainstream vendors designed systems that had huge impacts on IT budgets. Architecture matters, so make sure the one you select does not create unplanned additional expense

7. OPTIONS FOR ENTERPRISE VOICE COMMUNICATIONS

7.1 KEY SYSTEMS

Key systems are positioned as the first step for a small business that has outgrown an installation of separate disconnected phone lines from the telco (e.g., Centrex service). The key system achieves this by linking every phone in the organization to every other phone through fairly complex cabling. Key system phones are able to display the state of any phone on the small network, so calls can be forwarded from one person to another.

This provides small businesses with an affordable, entry-level telecommunication system. The downside is that key systems become obsolete when the business grows beyond 100 employees. In addition, they are extremely limited in terms of adding applications or managing system changes. To overcome the perception that key systems are limited in functionality, some vendors describe their key systems as hybrids, positioning them as "key systems with the smarts of a PBX." In practice, however, key systems do not typically support business growth, interact with computer telephony integration (CTI) applications, incorporate sophisticated call routing capabilities or link to other voice systems.

7.2 PBX

Because of the legacy PBX system's centralized architecture, certain considerations about the ongoing management of the system should be kept in mind when comparing the TCO of various approaches to enterprise telephony. For example, legacy PBX architecture is based on centralized system intelligence. Management and control are carried out from a single console, much like the mainframes that once dominated the IT world. The console is often connected directly to the PBX. Therefore, management of the system requires an engineer's physical presence. Although this may work in a small, single-site configuration, this centralized model becomes unmanageable and expensive for multi-site systems.

In the legacy PBX world, MACs require specialized personnel. The task of managing them is often outsourced to service organizations, which increases cost and reduces flexibility. Digital phones are proprietary; meaning any phone not manufactured by the PBX vendor will not function with that system. It also means companies pay a premium for these devices while sacrificing freedom of choice. Because management consoles use proprietary interfaces that are not integrated across applications, changing a user's location may involve use of three or more separate management consoles with three different interfaces.

Other issues are also driving the market away from legacy PBX systems. Perhaps the biggest problem is scalability, a constant source of frustration for growing businesses. These systems are designed for a specific-size customer. If a customer expands beyond its vendor-defined capacity, it must move to a higher-end PBX. One area where this problem becomes apparent is with smaller companies, which require a robust telecom system to be able to compete with the level of customer service provided by large companies. That requirement often leads to the purchase of a high-end PBX system – perhaps more than the smaller business wants or can afford – at a premium price. A better option would be to start off with a scalable system that can continue to grow in capacity and features along with the needs of the business.

In addition, PBX connectivity across multiple sites is a chronic problem for many growing customers. Although it would simplify the issue, vendors have been reluctant to use standard signaling protocols to link sites, because that would enable customers to switch to another vendor more easily. Another problem caused by PBX systems is that the handsets for low-end systems are often not supported when the company migrates to a higher-end system. Consequently, customers are forced to purchase new handsets when they upgrade.

Application integration is not easy with a legacy PBX, because of the overall design of the system. To integrate PBX voice systems with enterprise IT applications, the PBX must be supported by a separate, dedicated server loaded with voice processing cards. In an IP world, where voice and data already share the same transport network, this kind of awkward CTI equipment is unnecessary.

7.3 IP-PBX

IP telephony is clearly the future of enterprise voice communication. The most obvious benefits of an IP-based voice system include lower cost, more flexibility and improved usability and manageability. A significant cost saving comes from being able to use the existing data infrastructure, rather than a separate dedicated network. In addition, as was the case in the transition from mainframe-based computing to standards-based open systems, IP-based voice communication equipment represents a shift to commodity computing platforms available at significantly lower cost.

As we have seen, MACs are a significant expense in administering the voice system. In an IP-based system, the voice function is logically separated from the underlying network. This means that moving a user from one location to another does not require reconfiguration of the physical infrastructure. As a result, MACs in an IP-based voice system are typically one-third the cost of MACs in a legacy PBX system.

There are other benefits, as well. In the PSTN and PBX systems, users are intimately tied to their telephone numbers or extensions. In IP-based networks, the association between users and their IP addresses is through a DNS. In addition, IP addresses are usually assigned when a user logs into the network through a Dynamic Host Configuration Protocol (DHCP) server. Extending this even further, some innovative voice systems allow users to log onto any physical phone on the system, eliminating the expense related to MACs.

As mentioned earlier, legacy PBX systems tended to focus on telephone handsets, because they could generate significant revenue by locking the customer into high-margin, proprietary phones. With IP voice communication, the handset is no longer the focus. Instead, the desktop PC, with its intuitive interface, has taken center stage by simplifying voice communication and adding massive scalability to related converged applications.

Application development is done using standard interfaces, such as Microsoft's TAPI. This makes applications truly portable and interoperable with other standards-based systems. In addition, IP-based systems are not hindered by legacy evolution, so enhanced services (e.g., voicemail and automated call distribution) can be integrated into a single interface. Moving a user's location requires only a change in the associated physical port (the user's voicemail and automated call distribution profiles need not change), with the end result of lower administration cost.

Functionally, the IP-based voice system is similar to a traditional telephony system. It provides basic call management and enhanced services. A call server, such as a softswitch, provides internal and external call management, as well as translation between telephone numbers and IP addresses. Standard analog phones, IP phones or PC-based phones connect users to the network. The obvious difference with IP-based systems is that the voice server and phones are connected to the IP network, rather than a separate dedicated network. The transition is analogous to the migration from centralized mainframe networks to distributed, standards-based IP data networks. This evolution ultimately fueled tremendous growth of new applications and services, and resulted in the lower-cost computing platforms available today.

IP has become the strategic communication protocol for business – even the legacy vendors admit this. Therefore, customers who are encouraged to continue investing in proprietary, old-world PBX systems should be extremely wary. With the rapid adoption of converged IP voice and data infrastructures, a new PBX purchase is viewed as an unnecessary expense, rather than a strategic investment.

7.4 IP CENTREX

When this guide was initially written, Centrex was considered by the industry to be a failed experiment, but things have changed considerably. First, let's establish what Centrex offers, so we can discuss whether the approach makes sense.

Centrex moves call control off site. The enterprise still installs desktop telephones, but a high-capacity connection provides a link from phones to the actual switch and trunk interfaces, which are located at the operator's premises. The customer pays a rental fee, rather than making a capital investment.

The first time around, Centrex did not fair well because, as discussed earlier, moves, adds and changes (MACs) are far-from-insignificant considerations. With Centrex, the customer had to book MACs well in advance or put up with frustrating delays, contending with all the other Centrex customers for attention.

So how have things changed?

1. MACs are much simpler to use with Web-based self-provisioning
2. The cost of WAN connectivity has fallen dramatically

Ultimately, the lines between a hosted solution and a premises-based solution are almost entirely blurred:

- Money: Lease or buy?
- Skills: Hire or contract?
- Location: On or off site?

If you are seriously considering the IP Centrex model, or even moving parts of your voice solution off site for economic, security or reliability reasons, remember that as with any IT service, you must make sure you receive guaranteed service levels.

7.5 TOTAL COST OF OWNERSHIP (TCO)

Much has already been said about the expenses associated with PBX systems. Traditionally, these legacy vendors have avoided the issue of TCO, dodging the issue by offering substantial discounts on the initial purchase to create the illusion of affordability. However, they quickly make up for lost revenues by selling proprietary enhancements, applications, and services, further locking customers into long-term investments. The advent of IP telephony provides the opportunity for a new generation of vendors to challenge not just the technology, but also the overall value proposition for customers who purchase and use that technology. As you meet with vendors to discuss deployment of an IP voice communication system in your company, make sure that they provide clear, straightforward answers to the questions below. Doing so will save you time, money and a lot of frustration.

INSTALLATION
What is the charge for installing the system?
Is it simple enough to do it myself? How long does it take?

MANAGEMENT
Do I need skilled personnel to manage the system full time?
What kind of training is required to manage the system?
Do I need someone to manage each office location?

TELEPHONES
What kind of phones can I use?
How much do they cost?
Expansion: How much does it cost to expand the system?
How many users will it support?
How many sites?
What if I outgrow it?

MACS
What are the costs of handling moves, adds, and changes anywhere in my company?

Multiple Sites
How do I interconnect multiple sites?
How much does it cost?
What is the management impact?
What is the service and support impact?

Having clear answers to these questions will help ensure that the hidden costs are identified and understood prior to purchase.

Although cost is important, we believe that the factors driving the shift to VoIP are based on strategic business changes, described in Section 8.

8. CONCLUSION

VoIP systems today cannot only match the features of legacy PBX systems, but they have been built with today's communication environment in mind. When most legacy PBX architectures were launched, the Internet was irrelevant to mainstream business activity. Today, of course, the Internet is a crucial tool in facilitating business, and IP forms the foundation for many of the applications and systems that continue to drive productivity to new levels.

IP telephony is inherently designed to leverage the Internet phenomenon, providing a distributed communications infrastructure that businesses will use to both scale and simplify their activities simultaneously. The legacy vendors have clearly stated that IP telephony is the future, but they lack the focus of the pure IP players. Those organizations that embrace this technology will succeed while their competitors continue to watch and wait.

If you require additional copies of this guide, please contact ShoreTel at (408) 331-3300, or e-mail your contact details to info@shoretel.com.

9. TERMS AND ABBREVIATIONS

10Base-T: 10 Mbps Ethernet over twisted pair copper cable
100Base-T: 100 Mbps Ethernet over twisted pair copper cable
1000Base-T: 1000 Mbps Ethernet over twisted pair copper cable
10GBase-T: Proposed 10Gbps over copper cable
4e: Class 4 switch from Alcatel/Lucent
5e: Class 5 switch Alcatel/Lucent
ACD: Automatic Call Distribution
ACK: Acknowledgement Code
ADPCM: Adaptive Differential Pulse Code Modulation
ALG: Application Layer Gateway
AMIS: Audio Messaging Interchange Specification
ANI: Automatic Number Identification
API: Application Programming Interface
ARPA: Advanced Research Projects Agency
ASIC: Application Specific Integrated Circuit
ATM: Asynchronous Transfer Mode
BRI: Basic Rate Interface (2B +D)
CAT 5e: Category 5e twisted pair cable
CDMA: Code Division Multiple Access
CDR: Call Detail Record
CERT: Computer Emergency Readiness Team
CIF: Common Intermediate Format
CLI: Calling Line ID
CO: Central Office of a telecommunications operator
Codec: Coder/Decoder
CoS: Cost of Service
CRM: Customer Relationship Management
CSMA/CD: Carrier Sense Multiple Access/Collision Detection
CSU: Channel Service Unit
CS-ACELP: Conjugate Structure Algebraic Code Excited Linear Prediction
CTI: Computer Telephony Integration
D4: A T-1 framing scheme
DEC: Digital Equipment Corporation
DECT: Digital Enhanced Cordless Telecommunications
DES: Data Encryption Standard
DHCP: Dynamic Host Configuration Protocol
DID (DDI in the U.K.): Direct Inward Dial
DiffServ: Differentiated Services
DMS CO: Switches from Nortel Networks
DMZ: De-militarized Zone
DNS: Domain Naming System
DNIS: Dialed Number Identification Service
DoS: Denial of Service
DPCM: Differential Pulse Code Modulation
DSCP: Differentiated Services Code Point
DSL: Digital Subscriber Line
DSP: Digital Signal Processor
DS-0: 64 Kbps channel
DS-1: 1.544 Mbps = T-1 = 24 x 64 Kbps channels
DS-3: 44.736 Mbps = T-3 = 28 x T-1s = 672 x 64 Kbps channels
DSS/BLF: Direct Station Select Busy Lamp Field (Attendant Console)
DTMF: Dual Tone Multi Frequency
E-1 = 32 x 64kbps
E.164: International public telecommunication numbering plan
ECTF: Enterprise Computer Telephony Forum
ESF: Extended Super Frame – a T-1 framing scheme
FCC: Federal Communications Commission
FMC: Fixed Mobile Convergence
FPS: Frames Per Second
FXO: Foreign Exchange Office
FXS: Foreign Exchange Station
G.711: Pulse code modulation (PCM) of voice frequencies
G.723.1: Dual rate speech coder for multimedia communications transmitting at 5.3 and 6.3 kbit/s
G.729a: Coding of speech at 8 Kbps using CS-ACELP
G.729a (annexe B): G.729a with silence suppression
Ground Start: A way of signaling initiation of a call from a PBX to the CO by briefly grounding one side of a line
GSM: Group System Mobile
H.100: ECTF-standard CT bus implementation with PCI
H.110: ECTF-standard CT bus implementation with Compact PCI
H.225.0: Call signaling protocols and media stream packetization for packet-based multimedia communication systems

H.245: Control protocol for multimedia communication
H.248: ITU equivalent of IETF MEGACO
H.263: Video Coding Standard
H.264: Advanced Video Coding Standard
H.320: Narrow-band visual telephone systems and terminal equipment
H.323: Packet-based multimedia communications systems
H.324m: Video conferencing standard for low bit-rate (mobile) connections
H.450: Generic functional protocol for the support of supplementary services
HTTP: HyperText Transfer Protocol
ICE: Interactive Connectivity Establishment (Assists with NAT traversal)
IDS: Intrusion Detection System
IEEE 802.1: Transparent Bridging
IEEE 802.1d: Spanning Tree Algorithm
IEEE 802.3: CSMA/CD (Ethernet)
IEEE 802.3ad: Link Aggregation
IEEE 802.3ae: 10-Gbps Ethernet
IEEE 802.3af: Power over Ethernet
IEEE 802.11 a,b,g,n: Wireless LAN standards
IEEE 802.11e: QoS for WiFi
IEEE 802.16e: Mobile Wimax standard
IETF: Internet Engineering Task Force
IM: Instant Messaging
IP: Internet Protocol
IPBX: IP PBX
IPS: Intrusion Protection System
IPSEC: IP Security
IPT: IP Telephony
ISDN: Integrated Services Digital Network
ITU: International Telecommunications Union
IVR: Interactive Voice Recognition
LCD: Liquid Crystal Display
LDAP: Lightweight Directory Access Protocol
LED: Light Emitting Diode
Loop Start: A way of signaling call initiation by creating a loop across the two wires of a telephone pair
MAC: Media Access Control (Ethernet)
MAC: Moves, Adds and Changes (Telephony)
MAN: Metropolitan Area Network
MDMF: Multiple Data Message Format
MGCP: Media Gateway Control Protocol
MEGACO: Media Gateway Control, also known as H.248 signaling protocol
MIME: Multipurpose Internet Mail Extensions
MGCP: Media Gateway Controller
MOS: Mean Opinion Score
MPEG: Motion Picture Experts Group
MPEG-4 Advanced Video Coding Standard
MPLS Multi-Protocol Label Switching
MTBF: Mean Time Between Failures
NAC: Network Access Control
NAP: Network Access Protection
NAT: Network Address Translation
NIC: Network Interface Card
NI2: Standard ISDN Signaling scheme
NM: Nanometer
NTSC: National Television Standards Committee
OC-1: 51.840 Mbps
OC-3: 155 Mbps
OC-12: 622 Mbps
OC-48: 2.4 Gbps
OUI: Organizationally Unique Identifier
P.831: Subjective performance evaluation of network echo cancellers
PABX: Private Automatic Branch Exchange
PAL: Phase Alternating Line
PBX: Private Branch Exchange
PCM: Pulse Code Modulation
PoE: Power over Ethernet
PRI: Primary Rate Interface
Q.931: ISDN user-network interface layer 3 specification for basic call control
QoS: Quality of Service
QSIG: Q reference point Signaling
RAS: Registration Admission Status
RED: Random Early Detection
RGB: Red, green, blue
RSVP-TE: Resource Reservation Protocol – Traffic Engineering
RTP: Real-Time Transport Protocol
SBC: Session Border Controller
SCTP: Stream Control Transmission Protocol
SDH: Synchronous Digital Hierarchy (European Equivalent of SONET)
SDMF: Single Data Message Format
SDP: Session Description Protocol
SECAM: Sequential Color Memory
SIMPLE: SIP for Instant Messaging and Presence
SIP: Session Initiation Protocol (RFC 2543)
SLA: Service Level Agreement
SMDI: Simplified Message Desk Interface
SNA: Systems Network Architecture
SONET: Synchronous Optical Network
SRGS: Speech Recognition Grammar Specification
SRTP: Secure RTP
SSML: Speech Synthesis Markup Language
SS7: Signaling System number 7
STS: Synchronous Transport Signal
STS-1/STM-1: 51.840 Mbps

STS-3/STM-3: 155.52 Mbps
STS-12/STM-12: 622.08 Mbps
STS-48/STM-48: 2.488 Gbps
STUN: Simple Traversal of UDP through NATs
SYN: Synchronize Message Flood
T-1 1.544 Mbps = DS-1 = 24 x 64 Kbps channels
T-3 44.736 Mbps = DS-3 = 28 x T-1s =672 x 64 Kbps channels
TAPI: Telephony Application Programming Interface,
TCO: Total Cost of Ownership
TCP/IP: Transmission Control Protocol/Internetworking Protocol
TDM: Time Division Multiplexing
TELCO: Telecommunication service provider
TIA 568B: Pin layouts for RJ45 plugs
TLS: Transport Layer Security
TURN: Traversal Using Relay NAT
UAC: User Agent Client
UC: Unified Communications
UDP: User Datagram Protocol
UTP: Unshielded Twisted Pair
URL: Universal Resource Locator
VLAN: Virtual LAN
VoIP: Voice over IP
VoiceXML: Voice Extended Markup Language
VPIM: Voice Profile for Internet Messaging
VPN: Virtual Private Network
WAN: Wide Area Network
WEP: Wired Equivalent Privacy
WFQ: Weighted Fair Queuing
WiFi: Wireless Frequency
WIMAX: Worldwide Interoperability for Microwave Access
WPA: WiFi Protected Access
WRED: Weighted Random Early Detection
XML: eXtensible Markup Language
XMPP: Extensible Messaging and Presence Protocol
Y2K: Year 2000

REQUEST INFORMATION FROM SHORETEL

Your Name:______________________________

Your Company:____________________________

Title:__________________________________

Address:________________________________

Email Address:___________________________

Telephone Number:_______________________

Number of Employees:_____________________

Number of Sites:__________________________

Current Voice System(s):___________________

Please describe or check your specific areas of interest:
❍ Reliability through distributed call control
❍ Single system management
❍ Organizational productivity applications
❍ Collaboration and presence
❍ Flexible expansion capability
❍ Legacy migration
❍ Improved sound quality
❍ Cost reduction

Mail To:
ShoreTel
960 Stewart Drive
Sunnyvale, CA 94085
USA

World Headquarters:
960 Stewart Drive
Sunnyvale, CA
94085 USA

+1 (800) 425-9385 Toll Free
+1 (408) 331-3300 Tel
+1 (408) 331-3333 Fax

info@shoretel.com

www.shoretel.com

EMEA:
+44 800 652 8645 Tel

Asia Pacific:
+61 (0)2 9959 8000 Tel