



IP TELEPHONY POCKET GUIDE

BY BARRY CASTLE

2nd Edition September 2004

ShoreTel, Inc.
960 Stewart Drive
Sunnyvale, CA 94085

408.331.3300
1.800.425.9385

www.shoretel.com
info@shoretel.com

TABLE OF CONTENTS

	Executive Summary	4
1	Introduction	4
1.1	Expectations and Desired Outcome	6
2	Enterprise Voice Systems	8
2.1	Call Switching, Processing and Signaling	9
2.2	Line Interfaces	11
2.2.1	Terminal Line Cards	11
2.2.2	Trunk Line Interfaces	11
2.2.3	Trunk Features	12
2.2.4	Traffic Calculations	13
2.3	Cabling	14
2.4	Basic Features and Functions	15
2.5	Enhanced Features and Applications	16
2.6	Call Flows and Dial Plans	17
2.7	Automated Attendant	18
2.8	Call Detail Records and Billing	19
2.9	Next Steps	20
3	Data Networking	20
3.1	LAN Infrastructure	20
3.1.1	Ethernet Switching	23
3.1.2	Power over Ethernet	25
3.1.3	Wireless LANs	26
3.2	IP	27
3.2.1	IP Addresses and Names	29
3.2.2	Allocating and Managing Addresses	30
3.2.3	WAN Infrastructure	31
3.2.4	Routers and Routing	32
3.2.5	The Real Time Protocol (RTP)	33
4	VoIP – Technologies and Standards	35
4.1	How IP Voice Actually Works	35
4.2	VoIP Components	36
4.3	VoIP Standards	41
4.3.1	Session Initiation Protocol (SIP)	42
4.3.2	MGCP/MEGACO/H.248	43

4.3.3	ITU H.323	44
4.3.4	Codecs	46
4.3.5	Which Standard?	48
5	Deployment Issues	49
5.1	Legacy Integration	49
5.1.1	Basic Connectivity	49
5.1.2	Voicemail Integration	50
5.2	Supporting Voice Quality (QoS) in the Network	51
5.3	Reliability	55
5.4	Security	58
5.4.1	Telephony System Security Issues	59
5.4.2	Network and Computer Security	59
6	Telephony Applications	63
6.1	Convergence: Computer Telephony Integration (CTI)	63
6.1.1	TAPI	64
6.2	Personal Productivity	64
6.3	Collaboration	68
6.4	Voicemail and Unified Messaging	69
6.4.1	Unified Messaging	70
6.5	Supporting Teleworkers and Road Warriors	71
6.6	Multi-Site Connectivity	71
6.7	Call Centers and Customer Relationship Management	73
7	Options for Enterprise Voice Communications	75
7.1	Key Systems	75
7.2	PBX	75
7.3	IP-PBX	77
7.4	Total Cost of Ownership (TCO)	79
8	Conclusion	80
9	Terms and Abbreviations	82

EXECUTIVE SUMMARY

IP telephony (IPT) offers a viable alternative to the legacy voice exchange, delivering improved application integration, scalability, and multi-site management. It is precisely these features that make this rapidly-maturing technology so attractive to organizations seeking to reduce costs, increase productivity and improve customer relations. However, voice is critical to business, and a major change to emerging technologies like IPT requires an understanding of a broad range of technologies, careful planning and thoughtful implementation. This guide has been compiled to help decision makers to understand the relevant technologies as they define their strategies.

Following the introduction, we will review the typical components and features of an enterprise telephony system. We then introduce voice over IP (VoIP), the underlying network infrastructure, and then we address the various issues one must take into account when designing the network to carry voice. We conclude by comparing these VoIP solutions to the traditional centralized PBX, but before we do so, we describe advanced applications such as collaboration, presence, customer relationship management and unified messaging. ShoreTel recognizes the challenges of defining a communication strategy that spans multiple domains and this guide is intended to help you with this task.

1. INTRODUCTION

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

- A breakdown of the critical components of an enterprise voice system
- Voice over Internet Protocol (VoIP) technologies and standards

- Business applications such as unified messaging, converged conferencing, call centers and CRM
- Advantages and disadvantages of different architectures
- A review of the total cost of ownership (TCO)
- Glossary of terms

IP data communications is already the global standard. The transition to a pure IP environment will obviously have 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; lower 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, the key driver is the opportunity to gain competitive advantage by deploying these applications. For example, with the promise of improving the quality and value of integrated voice and data communications, businesses will 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 communications professionals will need a firm grasp of both voice and data communications. They will need to understand how such technologies and standards support emerging applications with the end result of delivering a converged enterprise communications platform.

This guide is brought to you by ShoreTel, the innovation leader in IP telephony for the enterprise. This is not intended to be a sales or product brochure. Its primary objective is to provide IT and voice professionals with the information required to prepare themselves and their IT infrastructures to more effectively leverage IP voice communications.

1.1 Expectations and Desired Outcome

Many companies have successfully made the jump from legacy systems to IP telephony – so what do these companies most like about their IPT implementations? The list below summarizes many of the advantages IPT offers and can serve as a checklist that can be used to evaluate the many vendor offerings:

Distributed Intelligence – By distributing call processing intelligence (the ability to set up and manage calls) across the network, the voice system eliminates any single point of failure, including the failure of the IP 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, applications) into a single view 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 will create a better customer experience. Such a system will allow your organization to appear more coordinated and more professional as calls and conferences are seamlessly transferred and shared between team members, sites, and mobile employees.

Ease of Use – Today's systems should deliver more 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. When it comes time to evaluating the different vendor solutions, we recommend checking 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 your staff gets from 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, the better systems of today have leveraged the additional capacity on the network and have improved design ergonomics to deliver improved sound quality. Your voice is more easily heard and the system is able to deliver the full range of audible tones to the human ear. In short, calls sound better so less energy is spent trying to communicate and more time can be focused on productive conversations.

Simple Expansion Capability – Legacy PBXs can be expensive and complex to grow. Some IP-PBX vendors manage to make matters worse by deploying multiple management interfaces for related data-networking components. Thus, we emphasize the importance of carefully considering how many steps it takes to expand the system. Does the vendor solution require lots of time and expensive highly trained personnel to grow the system, or is it so easy to use that you can have an unskilled staff member physically connect a voice switch at a remote site and bring it online from the headquarters in a matter of minutes? Speed and flexibility are critical in today's business environment; if you can adapt your voice system to your business imperatives you can get a new team up and running faster. You're more competitive.

Smooth Migration Path – The new system should be able to co-exist with legacy systems and applications and provide backwards compatibility with legacy trunks, extensions and voicemail. Examine whether the system has a set of interfaces to enable a stepwise migration from your legacy PBXs to IP voice. A smooth migration path will allow you to go live with new locations and teams at will. You drive the project rather than having the technology 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 you leverage this opportunity to deliver a better voice system from which your employees will directly benefit. It is critical that the key telephony stakeholders be involved in the decision making process early on:

- IT and Communications team
- CEO/Director's office
- Marketing and sales
- Key administrative personnel
- Customer service

By including each group in the process of selecting a vendor, you ensure that their requirements are taken into account and that 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 communications. Experts in this area can move on to the section that follows.)

2. ENTERPRISE VOICE SYSTEMS

A typical voice system has many components provides a number of functions.

The private branch exchange (PBX) is the telephony system most often deployed in larger companies today. The PBX is based on time-division multiplexing (TDM) technology – a technology that has been around for about 20 years ago and hasn't evolved much during this period.

A PBX system typically includes the following:

- 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
- A management console and the ability to track and account for calls
- Applications and enhanced services

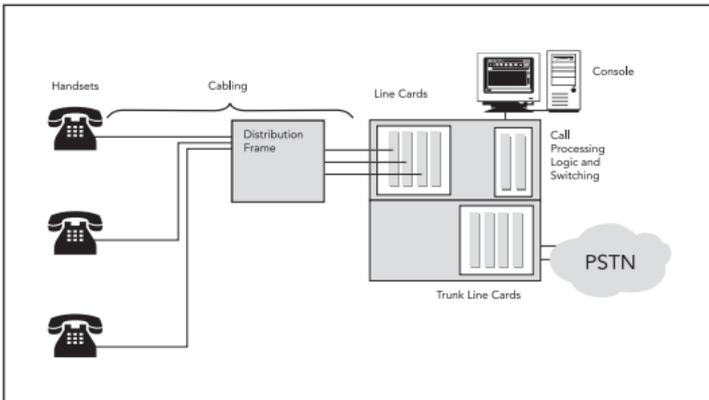


Figure 1: Components of the Legacy PBX

Note that this lists describes the physical components of a PBX system. To fully understand how these components fulfill their tasks, we also need software and signaling capabilities.

As we look at each consideration in the sections below, we will explore the functionality of the element and discuss the contribution of each element 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 only take place once 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 would connect 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 would have to talk to each other as they established a continuous circuit across many of these switchboards.

Just like the switchboard operators of the past, the PBX switches of today must remember what everyone is doing at any moment in time and must switch telephone calls between the appropriate places. The switch effectively establishes a circuit between the call 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 the attached handsets, adjacent PBXs, and the public telephone network.

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 below 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 the CO exchange – Standard signaling protocols
- PBX to PBX – Both proprietary and standard signaling (feature loss with standard signaling)

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

2.2 Line Interfaces

As mentioned earlier, there are two types of line interfaces for legacy PBXs. These are trunk-line interfaces used for connecting the PBX to the central office (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 (DSS/BLFs) 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 to users typically depends on the user's role and status within the organization. A manager might expect to use 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 are used to connect the PBX to the Public Switched Telephone Network (PSTN) and enable communications to the outside world. Note that there are many trunk options and it is important to understand the needs of your organization when selecting and ordering a trunk connection from your local telecommunications provider.

Trunks can be analog or digital T-1 with ISDN PRI or E-1 used in Europe. The various options for each of these trunk types have tradeoffs in terms of cost, capacity and features. Most vendors support the full range of trunk options available; therefore, the tradeoffs will be

based on cost and features required by the customer (covered in the next section). Note: the technical details of how PBXs and central office exchanges initiate calls and present audio streams over trunks are beyond the scope of this guide.

2.2.3 Trunk Features

Many trunk features are available. One common feature that almost all businesses will deploy is Caller Identification (Caller ID). Caller ID lets the called party see the calling party's name and telephone number before picking up the telephone handset (unless the calling party has specifically blocked this feature).

There are two different Caller ID formats that are used to deliver caller 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 exist for delivering caller ID: (1) Automatic Number Identification (ANI – similar to CLI), and (2) Dialed Number Identification Service (DNIS), which is an enhancement of 800# services that enable you to use CLI intelligence for sophisticated routing of calls into the organization.

An additional feature to order from your telecommunications provider (telco) is direct inward dial (DID) for inward call routing. DID essentially provides a way for external callers to contact a user directly at his or her unique phone number without any automated attendant or operator intervention. DID trunks are ordered in blocks consisting of 20 or more 10-digit telephone numbers. These numbers are assigned by the telco to a customer and are routed to DID trunks connected to the PBX. When a call is made to a DID number, the telephone company will send the last 3 or 4 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 refers to a mechanism for initiating an inbound call and passing the extension number to the PBX using a specific signal. Note that 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 basically designed to interact with the telco's CO switches; therefore, one of the PBX systems has to simulate CO signaling so that the two PBXs can communicate effectively. Similar schemes often are used when configuring a gateway or IP telephony system to connect to a legacy PBX.

2.2.4 Traffic Calculations

To determine the exact number of telephone lines and trunks your company requires, you must first determine the number of telephone users, calling traffic, and the 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 on the internet at <http://www.erlang.com/>. If no data is available for determining your telephone line and trunk requirements, you can follow the recommendations given below.

In general, smaller installations need 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

Note: These numbers are not applicable to a call center implementation, which typically have much different requirements. 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 has been installed correctly. Today, category five (CAT 5) twisted pair cable is the most popular cabling system being installed. It can carry both voice and data traffic.

The jack linking the cable to the desktop varies depending on whether a telephone or a network device (such as a PC Network Interface Card (NIC)) will be connected. The Ethernet NIC uses an RJ-45 plug and a standard analog telephone uses 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 cables terminate 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 one entry point to an appropriate exit point. The telephone engineer can establish the connection using a special-purpose tool used to push or punch the copper wire into a receiving contact.

A dedicated corporate telephone network (in which 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 has a move, add, or change (MAC) involving

approximately 12% of its users per year, with an average cost of \$150 per user. Therefore, MACs in a traditional PBX environment can be a significant, yet hidden, cost of ownership.

In contrast, data networking cabling tends to terminate the desktop wires on a patch panel so that an Ethernet drop cable can be used to link the desktop device to its corresponding Ethernet port. This same scheme is increasingly being used for voice cabling because it significantly reduces the cost of handling moves, adds and changes.

2.4 Basic Features and Functions

There are certain features and functions a telephony system is expected to deliver. 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's displaying calling information
- Support for DTMF codes
- Programmable keys
- Redials
- Music on Hold

- Last Number Redial
- Call Pickup
- Shared line ringing

Value-added features are most often embedded in telephone handsets to encourage customers to upgrade their handsets 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 along with a new handset requires significant engineering in the central PBX each time a new feature is added. And even more problematic, in today's business environment, the list of required features is exploding. The good news is that as the market continues to move to a pervasive IP-communications environment, adding new features will be similar to loading a new plug-in for a web browser. This ability to increase the functionality of voice communications will be 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 the phone system can provide. 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. The list below covers the adjunct systems typically listed by PBX vendors in their solution offering:

- Voicemail
- 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.

It's worth noting that the model of adding value to the system using third party devices is made easier when the voice system and the applications are designed from the ground up to share a common IP infrastructure. We will investigate the advantages of different architectures in a later section.

2.6 Call Flows and Dial Plans

When installing a voice communications 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, the operator, an assistant, an 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. Hunt groups and workgroups will often need to be defined for service centers and customer reps. The term hunt group is used to describe 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 will be 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, then it may be passed on to the group's voicemail. Understanding and configuring such functionality is critical to a successful telephony system installation.

A call handling process also needs to be carefully planned out for outbound calls in such a way that, for any number dialed, a corresponding route is available for the call. 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:

- You have an office in New York linked to the office in Dallas. You would like to save money by having long-distance calls routed over your company network. For example, if you were to place a call to an external number in New York from the Dallas office, the call would transit between your two internal PBXs from the Dallas office to the New York office and then the external call would only have to make a local hop to the destination number – this will save your organization money in long distance charges.
- You have a deal with a long distance carrier for calls made to London so that all international calls to the country code 44 will be prefixed with the alternate carrier's prefix number.
- When a staff member makes a call to an internal extension but uses the full external number prefixed by 9, you want to strip off all but the extension number and route the call internally.

Whether an external service organization is involved or not, defining dial plans requires careful analysis and thought.

2.7 Automated Attendant

The auto-attendant provides a customizable way for incoming callers to quickly route themselves to the people they need. This application uses in-band signaling called Dual Tone Multi-Frequency (DTMF) codes. DTMF assigns a certain sound frequency to each telephone key, so that when the dial pad keys are pressed, the auto-attendant “hears” these frequencies, interprets the information contained in these frequencies, and acts on that information.

For small businesses, the immediate advantage of having an auto-attendant is the cost savings of not having to pay an employee to act as an operator. However,

please 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 I have?
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 I forward 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 typically fed into a report engine that can produce more structured reports by department, by group, by 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 who take the basic PBX data and convert it into useful reports. Note that some service organizations (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 should provide you with a basic understanding of business telephony – at least, the way it used to be. But the reality is that the world of voice communications is changing, and as a result, next generation IP technologies are replacing outdated TDM technologies in the enterprise. With that in mind, we will now shift our focus to “voice over IP.” We include an introduction to data communications technologies and explain how IP voice communications are delivered on top of this infrastructure.

3. DATA NETWORKING

This section introduces the fundamental building blocks of a typical enterprise data network – Ethernet, switching, IP and routing (Data networking professionals may want to skip ahead to the next section).

3.1 LAN Infrastructure

The majority of local area networks today are based on Ethernet technology. Increasingly, IP runs over Ethernet, replacing protocols like 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). And, of

course, one of the earliest evolutions was the shift from shared coaxial cable to twisted pair cable.

Ethernet was 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 three-bytes identify the specific device. Information is sent between network devices using a predefined format known as a frame. Frame formats continue to evolve but are backward compatible.

The network interface card (NIC) is a device which 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 the raw information produced by the network device into frames before the data are sent to the cable used to connect 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, and then begins transmitting. It then 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 uses 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 then they would immediately recognize this because the information they detect coming back on the network does not match what they sent. They have detected a collision.

As a result, each computer will then back off – quickly flood bits onto the cable and then cease transmission for a random amount of time. They then 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 beaten competing approaches, such as IBM's Token Ring.

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

- Maximum length of twisted pair cables that connect Ethernet switches to devices (other switches, computers, IP phones etc ...) = 100 Meters
- Maximum fiber cable Length = 420 Meters

Today's cabling is typically UTP type 5 and a test certificate should be obtained from the cable contractor to ensure the cable conforms to the requirements of Ethernet. RJ-45 pin layouts are defined in TIA 568B. One thing to note is that when you connect an Ethernet switch or hub to another switch you will need a crossover cable (defined in 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 the twisted pair cable connecting 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 capacity of the network.
2. Unlike PBX systems, Ethernet devices are backward compatible so older NICs will 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 these 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 still need to be managed above this layer, 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 and today Ethernet networks are built from both chassis-based and stackable switches rather than shared media hubs. Switches provide each device connected to one of their ports with a dedicated bi-directional (full duplex) connection. This means that a device connected to a switch port will communicate at the maximum rated speed for that device. This differs from shared Ethernet topologies (such as hubs or, more often today, wireless) where the bandwidth must be 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 do this automatically using a protocol defined in IEEE 802.1 called **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 then this address and port number are also added to the table. Once source/destination addresses are known the switch can use that information to begin to forward frames. Shared hubs must forward frames to every device connected to them, which reduces the overall throughput for every device.

In certain circumstances it makes sense to segment traffic either by department or by application using a

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 in order to protect the 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 the voice traffic to share the same physical infrastructure as the bursty data traffic, but the voice traffic is protected at a logical level from interacting with the data traffic.

An 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 would keep claiming they were responsible for devices that were 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, and the responses start feeding back on themselves until a network meltdown occurs.

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 and for any path, a switch will only use the lowest path for a link and ignore 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 the spanning tree information can impact a call in progress. Switch manufacturers have developed proprietary solutions for providing rapid spanning tree updates and these solutions have largely addressed the challenge 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 (see

paragraph above on VLANs) much of the traffic will still transit certain links, resulting in bottlenecks. IEEE 802.3ad is a standard that 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 will be seeking to implement advanced networking capabilities, such as redundancy, link aggregation and QoS, which require careful planning and a modicum of 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, one of the challenges of a single, converged network is that the phone (which is seen as the lifeline to emergency services) may not be available during a power outage. With users expecting that they will have 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 - started accusing the IP-PBX community of cutting corners on the fundamentals. This has 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.

IEEE 802.3af is the standard that defines two ways to provide power to IP phones:

End Span – Replace Ethernet switches with new devices that use 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, as it requires an investment in new Ethernet switches.

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

Note that in both cases the system is non-destructive in that a non-compliant device can be plugged into the powered line without damage to the device.

3.1.3 Wireless LANs

In the USA, wireless premises for voice has remained vendor specific, which contrasts with Europe where regulated spectrum was established and a standard approach has been widely adopted – ETSI's DECT. The challenge with wireless has always been cost. The cost of providing full building coverage and possibly campus-wide roaming has been prohibitive; establishing the necessary access points or base-stations has been too expensive.

This situation is changing with the broad adoption of Wi-Fi networks based on IEEE 802.11, the wireless Ethernet standard. Broad market adoption has helped to drive access point prices down and many enterprises are intrigued by the potential to provide employees a single mobile device for applications as well as voice.

There are, however, some challenges that should be considered when planning a wireless implementation:

1. QoS – Delays for Enterprise voice should not exceed 150ms. Given that Wi-Fi is a contention protocol (like the original shared Ethernet), when a particular

access point is heavily used, voice quality will suffer. IEEE 802.11e is expected to be ratified in Sept 2004 and addresses wireless quality issues.

2. Reliability – Except for 802.11a, which specifies the 5GHz band, the 802.11 standards use the 2.4 GHz band and suffer from interference with other wireless devices.
3. Security – Voice can be slowed down by data encryption. A voice-specific encryption standard has been developed: IEEE 802.11i
4. Standards – There are multiple signaling standards (802.11a, 11b and 11g), which should be compared based on distance, capacity and frequency interference.
5. Handsets – Running voice through PDAs may be desirable, but ultimately, users will still need wireless phones for the foreseeable future. The number of suppliers is still relatively limited.
6. Coverage - Establishing coverage throughout a building is no 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 would we need anything more than just Ethernet? Let's consider some of the challenges of running just Ethernet:

- How do we tell applications running on computers how 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 will be 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 emails 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, DSL and ISDN in the WAN, etc..) There have been various proposed protocols including:

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

But interestingly, it is one of the oldest protocols that finally won the race. IP was originally developed by an agency of the United States Department of Defense, called Advanced Research Projects Agency (ARPA) as a peer-to-peer communications protocol to protect US military and R&D computer sites from being cut off in the event of nuclear attack. The protocol removed the hierarchical structure and there was no single point of failure that could bring down the nation’s communications. While this level of reliability was extremely valuable, the true drivers for IP’s success were:

1. IP protocols are simple to understand and implement
2. The protocol specifications are freely available

3. Applications like the World Wide 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 used across the globe as the basis for the Internet and the majority of enterprise networks.

3.2.1 IP Addresses and Names

IP addresses are used to 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) separate by a dot. The first groups of numbers identify the network upon which the machine is located, and the other groups of numbers identify the specific machine within that subnetwork. Here is an example:

169.254.70.213

When installing a computer, the network manager defines the machine's address and the address of the default gateway – the machine on the local network that provides connectivity to other networks. In the early days of the Internet, this machine would be a computer with two network interfaces and it would perform the task of forwarding packets between the two interfaces. Today however, it is more likely to be a router.

For use on machines 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 has allocated certain address ranges for internal machines only. The addresses of the internal machines are protected by a firewall that provides Network Address Translation (NAT) that 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 the specific machine or device. When planning an IP voice implementation, a legacy voice manager 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 on the Internet, a name-to-number translation service was introduced called the Domain Name System (DNS). This global system allows one to type in 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 website name instead of having to remember its IP address, which could change if the Louvre were to change to a different service provider.

When you consider how you use phone numbers in your cell phone, you will recognize that here too, you use people's names rather than their numbers. Later we will show how 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 to have its own IP address, does this mean that you must manually assign an IP address to every phone you install on your network? Fortunately, the answer to this question is “no.” An automatic address allocation system was developed to ease the administrative overhead associated with IP address allocation: Dynamic Host Configuration Protocol (DHCP).

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. IP addresses can be leased from an address range for devices that conform to certain criteria. 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 was 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 wide area network are almost limitless.

Here's what you will have to deal with in today's market:

1. T-1 and fractional T-1
2. ISDN
3. Frame Relay
4. ATM
5. xDSL
6. Cable
7. Wireless Local Loop technology such as 802.11 and LMDS
8. SONET
9. Wide Area Ethernet (delivered by metropolitan area networks)

Previously you may also have had to manage an access router with a rack of modems, but this service has been replaced by the combination of service providers and Virtual Private Network (VPN) technology.

Enterprises will typically connect their LAN 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 – end to end connections between the caller and callee. 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:

1. Reads the destination address of the packet
2. Checks which route to use
3. Forwards the packet to the interface associated with that route
4. And forgets it

So you can think of routers like a group of children standing in a circle throwing hot potatoes to each other. You catch it, throw it before your fingers burn, and forget it. There is no notion of a circuit, although as we will later see when we look at MPLS and quality, voice is placing new requirements on router protocols.

The primary function of a router is to route packets along the best path across a network. Each router maintains a route table, essentially a roadmap of the network which is kept up to date by exchanging information with other routers about the status of each link and the status of the network. When an incoming packet is received, a router identifies the destination address, checks the route tables 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 will be re-routed around the failed connection.

Due to their strategic location in the network (at the LAN/WAN boundary), routers are frequently used to pri-

ortimize and filter traffic. Note that 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 communications. QoS will be discussed in the following section.

Routers are also used to act as a firewall, filtering packets to protect the network from unwanted attempts to gain access to the network. Firewalls use 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), and because these applications can be used to break into a corporate network, identifying these types of traffic before the traffic 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 attacks, or SYN 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 could 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 has become much more powerful, making use of custom ASICs to simultaneously classify, queue, filter, and forward packets with minimal latency.

3.2.5 The Real Time Protocol (RTP)

During a conversation, the human mind can make sense of imperfect sound in which small (barely perceptible) bits of audio are missing – Audio CDs and the pointillism technique in painting show how the mind is capable of putting together a complete picture from fragments. But if a communication system has to stop processing to ask for a packet to be retransmitted, then the delay can cause a serious degradation in the per-

ceived quality of the conversation. In fact, for real-time communications it actually makes more sense to discard the missing packet rather than re-request it.

RTP is a protocol designed specifically 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: Used to record the sample rate and therefore playback rate

One can take a variety of different approaches (in a voice over IP implementation) to handling a lost packet during a conversation. For example, one approach might be to simply replay the sounds from the previous packet, discard the missing packet, and then play the next one when it arrives.

Another aspect of real-time communications is that in order to achieve higher quality sound, the sampling rate must be increased, and this leads to smaller packet sizes. The smaller packet sizes have the benefit of avoiding the increased processing latency caused by large packet sizes, but the small packets cause a new inefficiency—namely packet headers end up taking up as much bandwidth as the payload (i.e., the actual sound being transmitted). RTP provides a solution by using header compression. So RTP doesn't handle the voice signaling nor does it define the format to be used for the transporting the voice packets, but it provides important solutions to the challenges of real-time voice communications.

Problems of latency (perceptible delays between talker and listener) and jitter (a gradual loss of synchronization between the two endpoints) can still appear. However, given the speed and reliability of most enterprise LANs, such problems are rarely seen for conversations at a single site and tend to be crop up in WAN communications. Careful planning should be taken to ensure the WAN is correctly sized for the needs of the

organization, and care should be taken to verify that appropriate jitter buffers have been set. We discuss this more in the following section on voice quality.

With RTP we now have the necessary infrastructure and protocol foundation upon which to build our IP voice system.

4. VOIP – TECHNOLOGIES AND STANDARDS

4.1 How IP Voice Actually Works

To understand what makes up a typical VoIP system, let's begin by looking at what goes on while two parties are talking to each other. Imagine yourself holding a traditional analog telephone. 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 telephone handset, and the handset 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 1's and 0's. 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. (This process is shown below in Figure 2: VoIP Media Stream.)

It is at this point that one of the key differences between legacy telephony and IP telephony begins to emerge. Legacy PBX systems use proprietary, non-interoperable signaling, while IP telephony is designed to leverage standard signaling protocols.

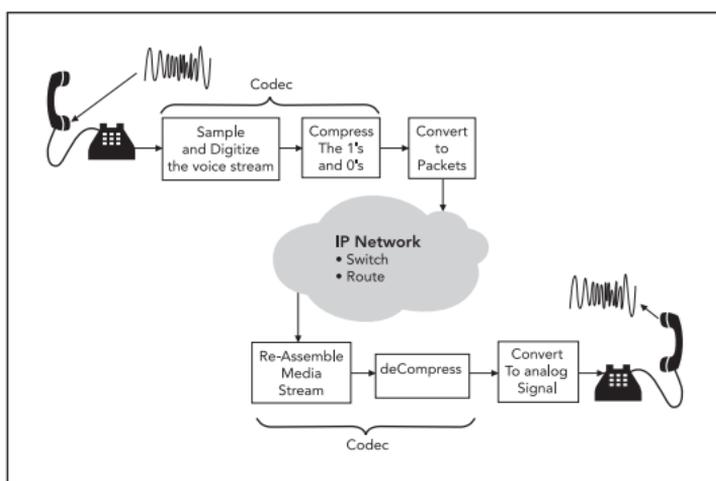


Figure 2: VoIP Media Stream

4.2 VoIP Components

Having established the process for voice communications 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. A lot of effort has been put into standardizing these rules so that IP voice systems can accomplish call processing in the same predictable manner web servers deliver information to web browsers anywhere in the world. We will hold our discussion of standards for a later section and, for now, focus on the functionality required.

Softswitch: To set up a call, the system must to act on signals from the calling phone. One way to accomplish this is with specialized call processing software that tracks and manages call progress, and 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.

Think of this device as an automated operator, handling all the tasks the switchboard operator used to handle.

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

These examples – which are typical of signals that might be sent between a telephone and a call server – show how a call server can perform the same functions as those of a PBX. So if the call server or softswitch can manage call set up, call routing and call tear down - does that mean we now have a fully functional IP-based alternative to the PBX?

Gateways: 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. 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, the gateway must have the necessary logic to convert between the two media formats and ask the call server for help setting up the call

In practice, it often makes sense to combine the functionality of call processing and gateway into single network elements – but for the purposes of this discussion, they are being 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. Phones can either be hardware devices that plug into the Ethernet network (and look just like a normal legacy telephone) or soft-

phones that run on the user's PC.

IP phones actually provide the functionality of a single user gateway, converting the analog speech pattern into digitized voice packets which are then sent over the IP network.

Here 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 that a PC can use the same uplink as the phone (offering savings on cabling cost)?
3. Does the phone support Power over Ethernet so it would work without interruption during a power outage?
4. Does the phone provide a mechanism to classify traffic so that 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 the phone deliver good sound quality?

As a mature market, differentiation was critical for the legacy PBX vendors and over time, the phone or handsets became the focal point for vendor competition. Specialist phones were developed:

- Operator consoles
- Administrative assistants
- Key systems
- Conference phones
- And a range of phones for the different hierarchical levels within an organization.

Today's IP phones have software hooks for customization and are far more flexible than their legacy counter-

parts. Even the simplest of telephone designs can be extended with applications that reside on the user's PC.

Today's IP phone provides an intuitive interface with access to application-rich features and is also capable of leveraging recent improvements in sound quality to provide a better experience for users and the people with whom they communicate.

The Complete System: Using these components:

1. Softswitch or media gateway controller
2. Gateways
3. IP Phones

We have the necessary components to build a complete system. 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, keeps track of everyone's state.
- The softswitch automatically converts between telephone numbers and IP addresses.
- Once a call route is established, the softswitch gets out of the way so that 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 is leaving the IP network and being 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 being 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 3 shows how all of these elements work together in a complete business telephony system.

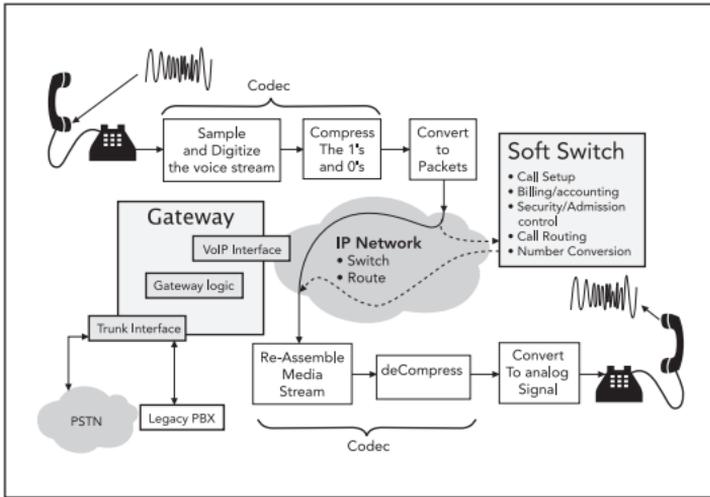


Figure 3: VoIP Architecture

Moving beyond the active components shown in this diagram, a key point to note is the total flexibility of the VoIP architecture delivered by the IP cloud. This IP network can span both LAN and WAN, and can be geography-independent as well as service provider-independent. Devices can simply be plugged into the IP cloud and be visible to the entire enterprise. Management can potentially be conducted from any point on the IP network and encompass every element of the business telephony system. We stress the importance of this fundamental difference between the legacy PBX and IP voice communications. Businesses can benefit from the inherent scalability that the IP infrastructure provides, as well as the interoperability between voice and computer components.

Having reviewed the functional elements that need to be provided by an enterprise-class IP voice communications system, we will now look at some of the frameworks and standards available.

4.3 VoIP Standards

Before we review the specifics of VoIP standards, it is worth spending a little time discussing why standards are so important to the communications 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.

We discussed signaling in the legacy PBX environment in an earlier section, and you will recall that there were several areas where proprietary protocols locked the customer into vendor-specific solutions that ultimately led to increased cost of ownership (e.g. proprietary handsets, costly application integration, and proprietary signaling between PBXs).

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 from one vendor to another. This is the real promise and power of IP voice communications.

In the communications world, standards discussions often get political and the 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, the following section provides a brief overview of three key standards upon which IP telephony is based. The glossary at the end of the guide can provide you with a reference for some of the commonly used acronyms.

4.3.1 Session Initiation Protocol (SIP)

SIP has emerged as a lightweight and 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 Location Server – track which IP address a client is currently using
3. SIP Proxy Servers – forwards requests to other servers on behalf of SIP clients

Redirect Servers – communicates the target address of the called party to the calling party.

The SIP protocol defines a set of basic messages to signal events:

SIP Messages

1. Invite – to join a session/call
2. Ack – to accept this invitation
3. Options – determine the capabilities of a server
4. Register – Register with a server
5. Cancel – cancels a previously issued request
6. 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 centers environments.

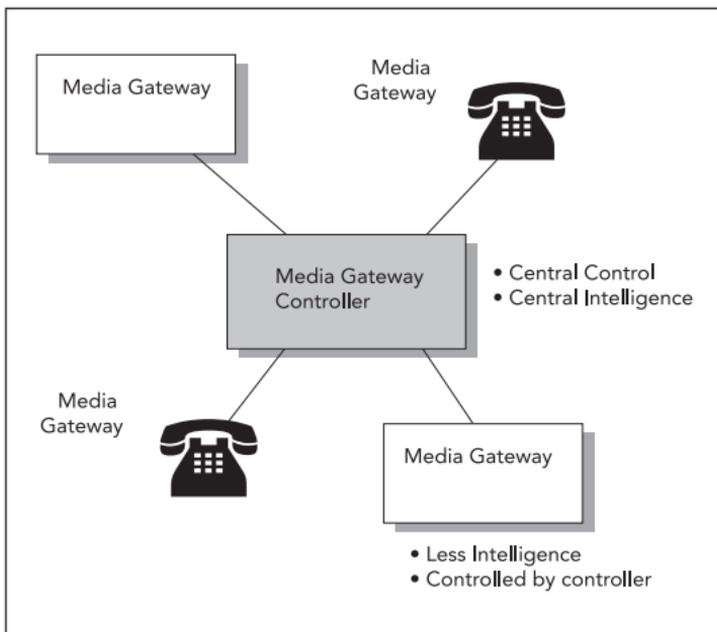
Where terminals are concerned, SIP's emphasis – like H.323 – is still on end-point intelligence, and this has

some significant implications for the cost of telephone handsets. 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 communications between multiple proxy and location servers. This part of the SIP specification makes it very scalable and manageable.

SIP is also the basis for SIMPLE, which is one of the two standards currently being proposed for instant messaging and presence. A voice system based on SIP will therefore be better able to integrate presence and telephony and will be able to deliver a richer suite of applications. For these reasons, we expect SIP signaling to be one of the requirements of next generation IP voice communications systems.

4.3.2 MGCP/MEGACO/H.248

There are two other standards worth mentioning: MGCP and closely related MEGACO. Unlike SIP and H.323, MGCP assumes that edge systems are unintelligent gateways and, therefore, 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.

When MGCP was initially introduced for standardization to the IETF, the name was changed to MEGACO and an agreement was reached with the 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 below). To simplify things and reduce costs, vendors have been implementing systems that use the original MGCP signaling proposal.

4.3.3 ITU H.323

The earliest VoIP standard was the ITU's H.323 standard, which evolved from H.320—the standard for video-conferencing. H.323 defines a set of standards to facilitate 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 bandwidth management
- MCUs: Multi-point Conferencing Units, support multi-party conferences for voice and video

In addition to IP voice communications, H.323 supports collaborative applications, such as whiteboarding 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 1990's H.323 has provided an important foundation for the VoIP community but has since lost ground to the more recent SIP and MGCP protocols.

H.323 is an umbrella standard that groups multiple sub-standards together into a single specification. Because the actual standards documents are cross-referenced with each other, it can be quite challenging to a new reader. The list below is intended to clarify roles and responsibilities of some 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 communications 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 communications systems
- H.248** ITU equivalent of IETF MEGACO
- H.450** Generic functional protocol for the support of supplementary services in H.323

Consider this example of the many technologies at work in a H.323 call. First, when a user picks up the phone and dials, an H.323-enabled phone uses H.245 to negotiate a channel and exchange capabilities. Then, H.225.0 handles call signaling and call set-up, and finally, a component called the RAS (registration/admission/status) channel signals the gatekeeper that coordinates the calls

within a zone. In addition, 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. However, it is still generally perceived to be overly complex and “too heavy” with functionality that is not required for IP voice communications.

4.3.4 Codecs

A codec is a device that converts sound from an analog format to a digital/numerical representation in the form of 1's and 0's (which is why the term itself is an acronym for coder/decoder). Codecs may also handle compression and decompression. There are many codec options available, and the choice often involves a tradeoff between voice quality, processing speed, and data size.

Below are some examples of codecs. Each codec is listed with the associated data stream produced, compression delay, and quality of the voice stream as measured by mean opinion score (MOS), an open test where a variety of listeners judge the quality of the voice sample 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

Determining which codec option is best for your company depends on your requirements. For example, G.711 provides excellent voice quality and would likely be 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.

Today the codec war has more or less died down with vendors selecting and licensing the codec they feel best serves their purposes (cost and technical). Part of the reason for this issue dying down is that network capacity has grown to such an extent that network performance is no longer the bottleneck (particularly on the LAN). Vendors have taken advantage of this situation by introducing systems that make use of wideband codecs that significantly improve voice quality without affecting either throughput or cost. Systems that support wideband codecs actually improve voice quality beyond what was achievable with the digital PBX.

Another codec issue to bear in mind is that of transcoding. Let's imagine two parties talking – one on a GSM or other cellular phone and the other using an IP telephone system. The call has to transit a gateway. If the IP telephony party is using one of the typical VoIP codecs, let's say G.729a, then to successfully communicate with the other party we need to decode the G.729a stream, transmit to the PSTN, which recodes the stream with a GSM full-rate or half-rate code. Each time we code or decode we introduce at least 12 milliseconds of delay (but often as high as 50-100 ms) – and this kind of transcoding should be avoided as it will impact overall end-to-end delay as described in earlier sections. 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 the service provider's equipment.

4.3.5 Which Standard?

When selecting an IP voice system, you should select a vendor which has clearly committed to certain standards. Each of the standards described below has 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. Today, however, its circuit switched heritage makes H.323 complex to implement, resource intensive, and difficult to scale. Vendors and service providers are now de-emphasizing H.323's role in their IP voice communications strategies.
2. SIP is ideal for IP voice and will play an important role for next generation service providers and distributed enterprise architectures. 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. The other similarity with H.323 is that SIP defines intelligent end points and vendors have found this approach to be more costly and less reliable.
3. In contrast to SIP, the MGCP/MEGACO standards both centralize the control of simple telephones. This is popular in environments where both cost and control are important issues, which is certainly the case in the enterprise environment where the PC can be used to augment features and functionality.

Moving forward, the market will likely support multiple standards for IP voice communications, with certain standards optimized for specific areas – such as carrier markets or communication with end point devices.

But as we have previously mentioned the current trend towards delivering presence along with instant messaging using SIP as the transport make it a strong contender for the delivery of application-rich voice systems over time.

5 DEPLOYMENT ISSUES

Having described the fundamentals of VoIP and the underlying IP infrastructure, we now turn to some of the deployment issues that need to be addressed.

5.1 Legacy Integration

In most cases, enterprises have legacy PBX systems and the migration to IP voice communications will be carried out in steps. This may be for financial reasons, as the PBX asset depreciates over a seven- or even ten-year period and a large number of legacy PBXs were sold in the lead up to Y2K. Equally important, the logistics involved in deploying simultaneous multi-site cutovers can be daunting. Thus, though it may only be a temporary requirement, it is important that the new IP-PBX can interact with various types of the legacy system.

5.1.1 Basic Connectivity

The easiest way to connect two PBXs is to use a digital trunk like T-1 or E-1. By linking an IP-PBX to a legacy system, we can 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 be a matter of a few minutes on the IP-PBX side.

After verifying that we can dial extensions between each system, we should verify that other features are implemented correctly. Caller ID, for example, can be forwarded from one voice system to another so that the identity of a caller can be displayed on the handset. Over the years, the legacy vendors have pushed proprietary solutions for delivering these features but, in fact, standards like ISDN-PRI and QSIG will provide the necessary signaling and are implemented in most PBXs and IP-PBXs.

The final decision concerning connectivity will be which system owns the trunk connection to the public

telephony network. Clearly there are three options for the location of the trunk connection:

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

The approach selected will depend on budget as much as anything else. If the line card used for the PSTN connection is to be employed to link the legacy system to the IP-PBX, then clearly an additional line card will have to be purchased for the legacy system. However, most organizations would prefer to purchase equipment that will work even after the legacy switch has been decommissioned, and this implies that any additional line cards should be purchased for the IP-PBX, which would act as a trunk gateway.

5.1.2 Voicemail Integration

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

Keep them Separate: the legacy PBX and the IP-PBX each have their own separate voicemail systems. Voicemails may be passed between the two systems using AMIS or VPIM

Host Voicemail on a Single System: either the legacy voicemail system is preserved for all users, or all users move to the voicemail system provided by the IP-PBX. SMDI is the appropriate protocol for this scenario.

First we will look at the two protocols available for interconnecting separate voicemail systems:

1. Audio Messaging Interchange Specification (AMIS) – Mimics the DTMF signals of the handset to forward voicemails between two AMIS-compliant systems over trunks between the two systems. Of the two approaches, AMIS is simpler to implement on the legacy system and requires significantly lower investment.

2. Voice Profile for Internet Messaging (VPIM) – provides a standard Multipurpose Internet Mail Extensions (MIME) encoding so that voicemails can be sent as multimedia emails over LAN/WAN data connections. While VPIM is the more modern of the standards, 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 company-wide voicemail storage. No matter which system is finally selected, the message-waiting lamps of the alternate system (IP-PBX or Legacy PBX) will need to be activated and deactivated. In other words, the handsets of the IP-PBX will light up when a voicemail arrives even though that voicemail gets stored on the legacy voicemail system. Similarly the legacy phones should indicate the availability of voicemail on telephone handsets connected to the PBX even though the actual voicemail resides, this time, on the IP-PBX system.

Simplified Message Desk Interface (SMDI) is a protocol developed for serial connections (TIA-232 cables) between voicemail systems and PBXs. The protocol signals the availability of voicemail for a specific extension to the management interface of the legacy PBX. The legacy PBX uses this information to signal to the handset that it should light its message-waiting lamp. When the message has been played, the lamp is switched off using the same mechanism. For some PBXs, an additional external interface may be required.

5.2 Supporting Voice Quality (QoS) in the Network

Voice quality is based on users' experiences and expectations. In general, today's business telephony networks have virtually no noise, echo, or delay. In IP voice communications networks, noise and echo problems can be easily addressed, but delay can still be a problem.

The ITU specifies an acceptable delay of not more than 300ms round trip or 150ms one way. Delays over 1/4 of a second (250ms) are noticeable and unacceptable. In the US, the acceptable delays in the PSTN are below 100ms. 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 source of delay 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.

Once the voice stream has been digitized, it must be packetized and transmitted onto the network. And this is the second place where latency can be added to the overall delay. Voice is typically sampled in 20ms intervals, but this is not a given, and transmission is dependent both on line speed and the packet size resulting from the sample rate. For example transmitting a 64 byte packet onto 56 Kbps line takes 8ms, (formula: $64 \text{ bytes} \times 8 \text{ bits/byte} = 512 \text{ bits} / 56,000\text{bps} = .0008 \text{ seconds}$). Transmitting a 64 byte packet onto a 10 mbps LAN takes only 51 microseconds.

In today's world of wire-speed switched LAN infrastructures, transmission delays introduced by switches and routers are negligible. Delays are likely to occur at the WAN access point where line speeds are low and the potential for congestion is high. To ensure voice quality, you first need to 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 about 26 kbps including the IP packet overhead.

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 address of the IP call server or softswitch, rather than of every end system.

An easier method implemented by some vendors is to prioritize packets based on predefined UDP port numbers. Once the packets have been 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 you have to track a range of port numbers. 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 you will prioritize voice traffic over your network we recommend you give some thought not only to which of the approaches is the more sophisticated but also how easily will you be able to manage the scheme. In general, the simpler approach will be the least costly.

Let's continue looking at some of the other ways vendors prioritize voice packets.

Looking deeper into the IP header, routers may also prioritize packets based on protocol type, such as time-sensitive protocols like SNA or SDLC. In addition, 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 where there is a dedicated leased line, prioritizing voice over data at the access point is all that is required. In public or shared wide area networks (Internet, VPNs), voice traffic must be prioritized throughout the shared network space. Rather than have routers classify and prioritize packets at every hop, the IETF is developing several standards to identify and prioritize different types of traffic.

Two approaches are available to request resources along

the communications path. One is the RSVP protocol, where each intermediate router along a communications path allocates resources based on current usage and keeps track of which resources have already been allocated. This essentially converts a 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 work that way. So, following unsuccessful attempts to get RSVP to scale for large deployments, an alternative was defined called DiffServ. In this scheme, different packets are “labeled” based on their desired behavior and are managed accordingly by each router along the path. This kind 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 QoS solution.

Multi-Protocol Layer Switching (MPLS) introduces an additional mechanism to improve quality. (The protocol gets its name 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. However, with MPLS, a label is applied at the edge router to each packet in a flow and any subsequent routers along the path simply forward the packets along a predetermined path without having to waste time examining the full packet header. This process is carried out at a lower or less sophisticated level of the router’s execution and uses up fewer processing cycles – In short, it is faster. The labels, of course, are similar to those proposed in DiffServ.

MPLS can be combined with another device known as a packet shaper, which has knowledge of specific applications (SAP, Voice, Video, HTML, email, etc.) and of individuals and organizations. By defining packet shaping policies that link to the MPLS labels in the routers, voice traffic can be assigned capacity on route paths such that 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 a user's perception of the quality of the call. The quality is influenced by codec selection, echo cancellation and silence suppression. In addition, latency above 100ms will sound uncomfortable for both parties. Network delay on the WAN introduces additional delay. Simply throwing lots of bandwidth at the problem will not necessarily eliminate quality problems. For this reason 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 communications 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 will be:

- Where is dial tone and call processing done?
- What operating system does this device use (Windows NT or embedded real-time)?
- What is the cost of protecting this device from failure?
- If I lose the WAN can I still make phone calls and 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 for the IP phones to fail over to another call control device somewhere on the network?

Note: Embedded operating systems for real-time devices are used in mission-critical applications, like vehicle braking 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.

Clearly the answers to these questions will indicate the

architectural robustness of each system being reviewed. Beware that for some vendors reliability comes at additional cost—a cost that will be incurred as soon as senior management find they are unable to make phone calls because of an operating system virus for example. It is clearly worth getting it right from the start, so spend time looking carefully at each architecture and base your decision on which architecture offers cost-effective and reliable call control.

Moving beyond issues of fundamental architecture and operating system selection, the best way to ensure reliability of IP voice communications is through redundancy. This can be done within a server with redundant processors, hard-drives, power supplies, and fans. And it can be done externally through redundant servers and sub-second failover between systems.

Just as important, the architecture of the IP voice system should allow you to distribute the voice communication functions throughout the network so that a user on the network can place a call without the need for a central voice call manager.

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

Project

Typical - Place a value/cost on downtime. Identify and eliminate any single point of failure. Document actions/procedures in case of emergency.

Advanced - Determine whether this should be part of a broader disaster recovery program. Have a third party audit the measures established for reliability.

Power

Typical - Provide a UPS for server, PBX and network components.

Advanced - Install a backup power generator. Have separate power cables enter the data center from different sides of the building.

Servers

Typical - Eliminate any desktop PCs being used as servers. Use server mirroring and load balancing. Install multiple server network interface cards with load balancing and failover.

Advanced - If the server is used for delivering call control or routing then you should implement clustering and find some way to isolate these systems from potential hacker attack.

Data backup

Typical - Establish a periodic (daily) backup schedule. Carry out fully backups prior to any system changes. Archive backups (weekly).

Advanced - Store archives in a secure site outside the building in which the computers reside

Network

Typical - 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
- Use management software with remote notification of alerts between network devices that support rapid (sub-second) failover
- Use management software with remote notification of alerts.

Personnel

Typical - Train staff on actions in case of emergency. Train first aid and fire supervisors. Have a phone (attached to PSTN powered line) available in computer room.

Advanced - Require at least two people to be in computer room after work hours.

Access

Typical - Establish access control for the computer room.

Advanced - Place protective gridding around racks. Place video monitoring in computer room.

Environmental

Typical - Review computer room air cooling system to verify it can handle the heat generated by machinery.

WAN and Telco

Typical - 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

Typical - Link “stay alive” ports of the voice switch to emergency phones distributed around the premises. Use a distributed voice architecture that limits failure points. Comply with E.911 requirements providing location information to emergency services. Deploy phones that take power from the voice switch.

Service/Support

Typical - Subscribe to service contracts that provide the appropriate level of critical support.

Advanced - Establish SLAs with clearly defined penalties that are links to lost revenues (insurance contract).

5.4 Security

Note that many of the reliability issues described in the previous table also apply to the issue of security. In this section we will quickly look at network attacks and telephony fraud. The goal is to make you aware of the issues rather than prescribe solutions to specifics. Many companies focus on securing themselves from outside abuse, but they may overlook breaches of security that occur from inside the organization.

5.4.1 Telephony System Security Issues

A good accounting package should be able to pinpoint misuse of phone services. Often, such problems can be resolved by a simple discussion with a 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 to request to be removed from the caller's list.
- The features that facilitate our day to day use of our 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 the problem with 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 down time 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 you need to consider tradeoffs between securing the network and the cost of doing so.

There are many sources of information on network security and there are many tools to help you identify known vulnerabilities. One useful reference on the Internet is <http://www.cert.org/>.

The CERT/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 used this feature of the network to their advantage. But this problem is not unique to IP. When a computer system is connected to a network, it becomes vulnerable to attack – whether IP or any other 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, here is a summary of a few 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 Attack – there are several DoS attack techniques that can be used to bombard the network with enough packets to disable network elements or gain access to the network.

Syn Floods – an example of a DoS attack where a target machine, often a router, is flooded with TCP connection requests, typically resulting in either a

slowdown 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 passwords, 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 4) to ensure internal IP addresses are not visible to the outside world. Just outside the firewall, an area is often set up where systems are located that interact directly with the Internet (e.g., public web server, email forwarder). A gateway firewall protects the 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 websites but only if they transit via the web proxy so that source addresses are hidden
 2. Email being sent, transits via the email proxy, which also disguises the source address information
 3. Email being received first goes to the proxy, then has a predefined path between the proxy and the internal email server
- No emails trying to bypass this route can get through the gateway
 - Additional functionality can be added in the email proxy to open and verify incoming emails 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 don't accept network logins and other means of penetration. Any attempts to login or attach to these systems should have the affect of generating alerts or even generate misleading information to the potential hacker. One obvious way to accomplish this is to disguise the OS and vendor names with misleading names and banners.

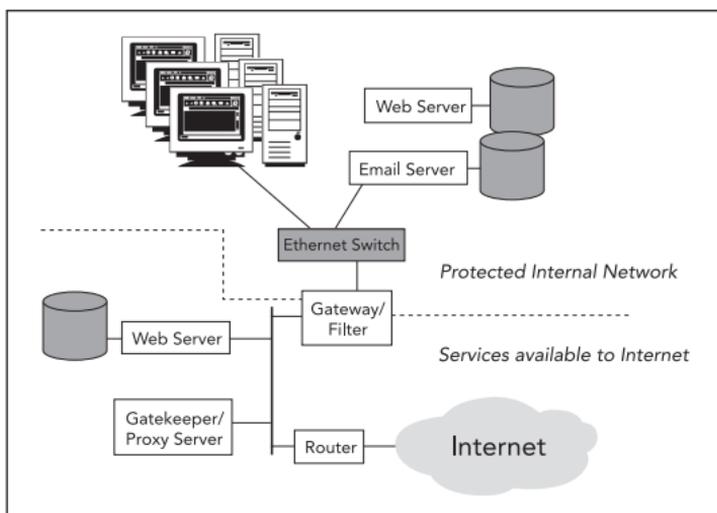


Figure 4: Network Security Firewall

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

As we conclude this section it is worth mentioning that many of these issues are actually already addressed by the existing data communications infrastructure. In other words, these issues are not specific to voice. However, you should be aware that voice communications will expose any fundamental weaknesses in the design of the data infrastructure.

Ask your vendor for guidelines and instructions for tuning the IP infrastructure in preparation for the VoIP system.

6 TELEPHONY APPLICATIONS

In this section we move beyond the basic features and functions of making phone calls to explore enhanced 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, and then we explore 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 email, fax and CRM.

In the PBX model, voice cards had to be installed in dedicated application servers to bridge the two worlds. Alternatively CT extensions had to be purchased for each user telephone to be bridged. 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 that define standard APIs for developing CT applications. (Note: Most PBX vendors have proprietary APIs.) Some vendors have an extended CTI model to develop their own CT-rich alternatives to PBXs.

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

6.1.1 TAPI

TAPI is Microsoft's Telephony API that allows third-party call control applications to control telephony functions from the client application. Outlook, for example, is a TAPI-compliant application that allows users to dial the phone numbers of people directly from the Outlook contacts database by clicking a button. This means that whether you are a user of a contact database such as Microsoft Outlook or Symantec ACT, you no longer have to look up your contact information and manually dial from the phone. The TAPI interface lets you select your contact manager and initiate a telephone 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 is TAPI 3.0 released with Windows 2000.

General purpose applications such as Outlook are not telephony-specific, and most users prefer the specialist clients that deliver important productivity gains.

6.2 Personal Productivity

Most VoIP users find it is the productivity applications that really drive the benefits of the new system. When employees can focus on their jobs rather than on searching for (and dialing) a phone number, they are less stressed in their work and more relaxed with their customers. Contrast this with the traditional approach.

In the PBX world, when you joined a new company you expected to receive:

1. A phone
2. A guide for using the phone's features
3. The 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 request to the person on the other end of the line: "Listen if I lose you, could you call me back?" In one breath, we involve two employees and simultaneously establish 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? How can we make things easier and how do we deliver enhanced applications? Clearly, the PC is a resource we can leverage to improve the usability of the voice system and increase productivity. Here are some of the features you should 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 or digital phone in such a way that the PC can interact with the phone. The application can control the phone and the phone can pass information to the PC application. The IP network provides the "glue" between the two devices.

We can now extend the model so that all features are readily available through buttons and menus, thereby eliminating the need for the telephone handset guide of the legacy model. 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

we can use the phone system much like an automatic log or notepad. Using the call history we never need to search for a number again, we simply scan through the history and click dial – we can even add the number to our contacts list (next paragraph).

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

A well-designed IP-PBX will provide a system-wide directory that combines personal contacts with corporate contacts in a single database that can be easily located as you type in fragments of people's names or numbers. You might type in first names, last names, peoples' initials, names that sound similar, departments or any other criteria that makes sense. Database searches like this make telephone usage so much easier by being 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 used so that call handling mode can be set based on knowledge of the user's meeting schedule. For example, she's in a meeting, so automatically send her calls 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 a user has recently typed at the keyboard or moved the mouse on the computer 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 a combined voice/data system

provides of people's current status 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 – we should forward the call somewhere else.

With real-time knowledge of what state the phone is in, calls can now be routed to a person who is at their desk and can take the call immediately rather than forwarding a valuable customer to someone's phone number (where it ends up in voicemail). The net result is a more responsive, more productive organization.

There are two presence standards you should be aware of: SIMPLE based on the SIP protocol and the XML-based XMPP –SIP was originally designed as a VoIP standard so it is more than likely that implementations of presence for voice systems will use SIMPLE.

Web Conferencing: Once 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 the associated computers. Because of this, if callers 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. Whenever we free people from the time they spend working out how to accomplish tasks with their technology we increase productivity.

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

So we are now more productive but can we work as a team?

6.3 Collaboration

In the old days, the concept of a corporation was of a bricks and mortar building where the employees went to work. In fact it was for this application that the original PBXs were designed. Today however, things have dramatically changed:

1. Enterprises nearly always have more than one site
2. Sales staff are expected, even encouraged, to spend the majority of their time with customers and often don't have an office with a desk
3. Many key contributors for an enterprise actually work for supplier organizations
4. Customer service requirements have increased so dramatically that the successful enterprise must integrate in seamless collaboration with their customers

Whether it's for specific phone calls or for formal meetings, the voice system is the 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.
- Provide **distributed applications** so that customer interaction can be spread across multiple sites and can dynamically include any employee at the click of a button
- Provide **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 his number

- Make people easier to reach with **follow-me find-me** rules so that their customers can find their representative just by typing a single extension number
- Enable sales staff located temporarily at their home office or at a customer's premises to leverage the full range of features remotely, for example with a **soft-phone** on their portable computer
- And all of these features should be **easy-to-use** in such a way that employees can spend their time on their jobs rather than on working out how to schedule a conference call

With renewed focus on **collaboration** between employees, customers and suppliers – the IPBX moves us beyond the monolithic single-site systems of yesterday, to address the needs of the distributed organizations our businesses have become.

6.4 Voicemail and Unified Messaging

A voicemail system is used for store-and-retrieve voice communications. 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 email 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 gets an immediate notification on predefined device (pager, message light on phone, email). The called party dials into a central voicemail number or simply presses voicemail key on desktop phone to access the system. The system prompts for number and ID number, and then provides

menu options. The called party listens to message, then deletes, saves, forwards or replies to message.

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, email, fax).

6.4.1 Unified Messaging

The challenge with a typical voicemail system is that it takes time for the user to learn the multi-layer menu system and know where the shortcuts are. Users often have to keep a visual representation of the menu layout next to the phone.

By providing a unified view of all media in the email system, users can leverage a more intuitive interface and have the freedom to handle and store messages in the same way as they do with email. For example, users can view and listen to voicemail in any order they choose, instead of tabbing through a series of messages on the phone to get to important message number 11.

The key to unified messaging is to establish a very tight relationship between PC and voice system. In the past this kind of integration could only be achieved through the use of separate, loosely linked systems – a PC full of voice processing cards that effectively provides a bridge between the worlds of voicemail and email.

With IP telephony this problem goes away. Because it is integrated by design, the IP network effectively provides a common link between any network component: email server, PC, voicemail, and telephone.

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 has been driven 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 office, they have historically lacked any level of voice network integration.

Softphones can deliver the full range of telephony features and enable the individual to feel like an active member of their organization even if they are geographically far away. Such phones can run on the user's computer which might be located in:

1. The home office
2. A customer site
3. A WiFi hotspot located in a coffee shop

This kind of flexibility was never available with legacy PBXs and demonstrates the advantage of a distributed approach to voice communications.

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 network from LAN to WAN, central site to remote office, or even to home offices. It's totally location-independent, yet it functions as a single, cost-effective network.

When a system has been designed as a multi-site application rather than as a set of independent PBXs, the user experience is dramatically improved. There is a single dial plan for the whole organization irrespective of site/location, and if users travel from one site to another, they simply log into the system and their calls are forwarded to their new location without their having to make any configuration changes.

As we have already mentioned, presence and messaging capabilities add considerably to employee productivity. This benefit is even more dramatic when users are located at multiple sites. For example, the ability to know the availability of a colleague 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 communications.

Even when a voice system is distributed across the LAN, MAN, or WAN, all of its elements and all of its 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 ACD) are managed as separate entities, with separate databases, and separate interfaces. This archaic approach makes it 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 an IP-based voice communication system over a traditional PBX system 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. Therefore, in call center environments calls may be distributed to users anywhere on the IP network. While distributing calls between local and remote offices is also possible with traditional PBX systems, the ability to easily move users from one location to another is a strong advantage of IP-based voice systems. This allows, for example, a customer support expert to move from a central site to a home office, and easily remain 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 communications system work with key CRM applications to provide the highest level of customer service possible. This 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 Automatic Call Distribution (ACD) module. The ACD provides the function of queuing and managing the distribution of calls to the appropriate agents. Calls entering the contact center can initially be 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. This new expanded role is being described within the umbrella concept of CRM, which 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 there is IP access 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) will be the only way of communicating with businesses. Email, web forms, instant messaging and chat are all now widely used and all these tools provide an opportunity 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 a call, based on their business objectives (cost, workload and results). Getting the right person to handle the call ultimately costs less and enhances the customer experience.

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 (cabling is fairly complex). Key system phones are able to display the state of any phone on this small network so that calls can be forwarded from one person to another. These systems provide small businesses with an affordable, entry-level telecommunication system.

The downside is that key systems are quickly rendered obsolete by any business growing beyond 100 employees. In addition, they tend to be extremely limited in terms of adding applications or managing system changes. To try 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 the PBX.”

In practice, however, key systems will not typically support business growth, will not interact with CTI applications, will not incorporate sophisticated call routing capabilities, and will not link to other voice systems.

7.2 PBX

Because of the centralized architecture of the legacy PBX, there are certain considerations concerning the ongoing management of the system to keep in mind when comparing the total cost of ownership (TCO) of different approaches to enterprise telephony.

For example, the legacy PBX architecture is based on centralized system intelligence, so that management and control is 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 the physical presence of the engineer. 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 and the task of managing them is often outsourced to service organizations, which increases cost and decreases flexibility. Digital phones are proprietary, meaning any phone not manufactured by the PBX vendor will not function with that system. It also means you end up paying a premium and sacrificing freedom of choice. And management consoles use proprietary interfaces that are not integrated across applications. As a result, changing a user's location may involve use of three or more separate management consoles with three different interfaces.

There are other issues driving the market away from legacy PBX systems. Perhaps the biggest problem is one of PBX scalability, which is a constant source of frustration for growing businesses. These systems are typically designed for a specific size customer, and if that customer expands beyond the vendor-defined capacity, the customer must move to a higher end PBX.

One area where this problem is apparent is with smaller companies that require a robust telecom system to be able to provide the same level of customer service as 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 provide customers with a scalable system that would continue to grow in capacity and features with the needs of the business.

In addition, PBX connectivity across multiple sites is a chronic problem for many growing customers. Vendors have been reluctant to use standard signaling protocols to link the sites and simplify the problem

because such protocols would enable the customer to move more easily from one vendor to another.

Another problem with PBX systems is that the handsets customers purchase for their low-end systems are often not supported as they migrate to a higher end system. Consequently, customers are forced to purchase new handsets when they upgrade.

Finally, 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 typically be supported with 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 communications. 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 pointed out above, MACs are a significant expense in administering the voice system. In an IP-based system, because the voice function is logically separated from the underlying network, 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 with a legacy PBX system.

There are other benefits as well. In the PSTN and in

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 domain name server (DNS). In addition, IP addresses are usually assigned when a user logs into the network through a DHCP (Dynamic Host Configuration Protocol) server. Extending this even further, some innovative voice systems allow users to log onto any physical phone on the system thereby eliminating any expense related to MACs.

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

Applications 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), and the end result is lower cost of administration.

Functionally, the IP-based voice system is similar to that of a traditional telephony system. It provides basic call management and enhanced services. Typically, a call server such as a softswitch provides internal and external call management and 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 connect-

ed to the IP network rather than a separate dedicated network. The transition is, again, 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 communications protocol for business – even the legacy vendors will admit to that. Therefore, customers being encouraged by those vendors 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 would be viewed as an unnecessary expense rather than a strategic investment.

7.4 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 the customer into a long term investment.

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 the customers who purchase and use that technology.

As you meet with vendors to discuss deployment of an IP voice communications system in your company, make sure that they provide you with 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 well identified and understood prior to purchase.

Although cost is important, we believe that the factors driving the shift to VoIP are based on the more strategic business changes around us. We describe these changes in our final words below.

8 CONCLUSION

VoIP systems today cannot only match the features of legacy PBXs, but they have been built with today's communications environment in mind. When most of the legacy PBX architectures were launched, the Internet was irrelevant to mainstream business activity. Today, however, 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 our 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.

We hope you have found this guide useful. If you require additional copies please contact ShoreTel at (408) 331-3300, or email 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 Lucent Technologies
5e	Class 5 switch from Lucent Technologies
ACD	Automatic Call Distribution
ADPCM	Adaptive Differential Pulse Code Modulation
AMIS	Audio Messaging Interchange Specification
ANI	Automatic Number Identification
ASIC	Application Specific Integrated Circuit
DiffServ	Differentiated Services
DNS	Domain Naming System
DSP	Digital Signal Processor
CLI	Calling Line ID
CO	Central Office of a telecommunications operator
Codec	Coder/DeCoder
CRM	Customer Relationship Management
CSMA/CD	Carrier Sense Multiple Access/Collision Detection

CSU	Channel Service Unit
D4	A T-1 framing scheme
DMS	CO Switches from Nortel Networks
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
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.164	The international public telecommunications numbering plan
ECTF	Enterprise Computer Telephony Forum
ESF	Extended Super Frame – a T-1 framing scheme
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 Conjugate-Structure Algebraic-Code-Excited Linear-Prediction (CS-ACELP)

Ground Start	A way of signaling initiation of a call from a PBX to the CO by briefly grounding one side of a line
H.100	ECTFs standard CT bus implementation with PCI
H.110	ECTFs 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.320	Narrow-band visual telephone systems and terminal equipment
H.323	Packet-based multimedia communications systems
H.450	Generic functional protocol for the support of supplementary services in H.323
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
IP	Internet Protocol
LCD	Liquid Crystal Display
LED	Light Emitting Diode

Loop Start	A way of signaling call initiation by creating a loop across the two wires of a telephone pair.
MGCP	Media Gateway Control Protocol
MEGACO	Media Gateway Control, also know as H.248 signaling protocol
MOS	Mean Opinion Score
MPLS	Multi-Protocol Label Switching
NAT	Network Address Translation
NI2	Standard ISDN Signaling scheme
OC-1	51.840 Mbps
OC-3	155 Mbps
OC-12	622 Mbps
OC-48	2.4 Gbps
OUI	Organizationally Unique Identifier
PABX	Private Automatic Branch Exchange
P.831	Subjective performance evaluation of network echo cancellers
PBX	Private Branch Exchange
PCM	Pulse Code Modulation
Q.931	ISDN user-network interface layer 3 specification for basic call control
QSig	Q reference point Signaling
RTP	Real-Time Transport Protocol
SDH	Synchronous Digital Hierarchy (European Equivalent of SONET)
SIMPLE	SIP for Instant Messaging and Presence Leveraging Extensions
SIP	Session Initiation Protocol (RFC

	2543)
SLA	Service Level Agreement
SMDI	Simplified Message Desk Interface
SONET	Synchronous Optical Network
STS	Synchronous Transport Signal (SONET)
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
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
TIA 568B	Pin layouts for RJ45 plugs
WAN	Wide Area Network
UTP	Unshielded Twisted Pair
VoIP	Voice over IP
VPIM	Voice Profile for Internet Messaging
VPN	Virtual Private Network
XMPP	Extensible Messaging and Presence Protocol

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

© 2004 ShoreTel, Inc. All rights reserved.

ShoreTel and the ShoreTel logo are trademarks of ShoreTel, Inc. in the United States and/or other countries. Other product and brand names may be trademarks or registered trademarks of their respective owners. All specifications are subject to change without notice.

ShoreTel, Inc.
960 Stewart Drive
Sunnyvale, CA 94085

408.331.3300
1.800.425.9385

www.shoretel.com
info@shoretel.com