

IP Telephony Demystified

Ken Camp

McGraw-Hill

New York Chicago San Francisco Lisbon
London Madrid Mexico City Milan New Delhi
San Juan Seoul Singapore Sydney Toronto

Library of Congress Cataloging-in-Publication Data

Camp, Ken.

IP telephony demystified / Ken Camp.

p. cm.

ISBN 0-07-140670-0 (alk. paper)

1. Internet telephony. I. Title.

TK5105.8865 .C34 2002

621.382'12—dc21

2002032669

Copyright © 2003 by The McGraw-Hill Companies, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

1 2 3 4 5 6 7 8 9 0 DOC/DOC 0 9 8 7 6 5 4 3 2

ISBN 0-07-140670-0

The sponsoring editor for this book was Stephen S. Chapman and the production supervisor was Sherri Souffrance. It was set in New Century Schoolbook by Patricia Wallenburg.

Printed and bound by RR Donnelley.

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. For more information, please write to the Director of Special Sales, Professional Publishing, McGraw-Hill, Two Penn Plaza, New York, NY 10121-2298. Or contact your local bookstore.

Information contained in this book has been obtained by The McGraw-Hill Companies, Inc., ("McGraw-Hill") from sources believed to be reliable. However, neither McGraw-Hill nor its authors guarantee the accuracy or completeness of any information published herein, and neither McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that McGraw-Hill and its authors are supplying information, but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.



This book is printed on recycled, acid-free paper containing a minimum of 50 percent recycled, de-inked fiber.

CHAPTER **4**

**IP Telephony
Protocols**

Earlier, we reviewed the need for standards and why they exist. We've discussed several different standards organizations. Now we're going to explore some of the standards being used and developed specifically for IP telephony. To keep things in perspective, we must first set the stage.

In the previous chapter, we mentioned that standards used in the worldwide PSTN are the work of the International Telecommunications Union—Telephony sector (ITU-T), formerly known as the CCITT. This group operates under the auspices of the United Nations, and that is an important note. This is a global body, and while it has responsibility for telephony standards, it has not historically been known for speed. Many different international political agendas come into play when dealing with the United Nations, and change takes time.

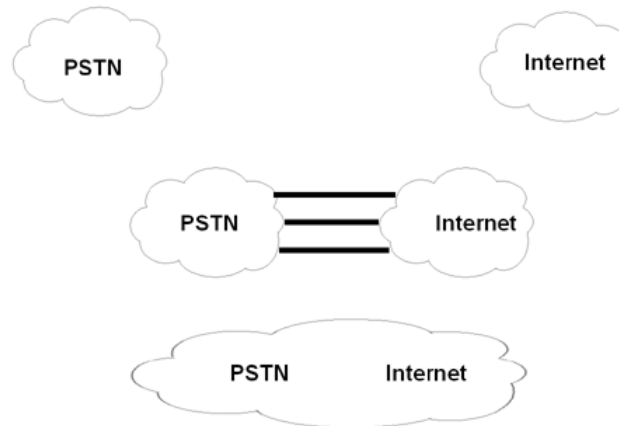
Standards in the Internet are primarily the work of the Internet Engineering Task Force (IETF). This is also a global group; however, it's primarily an organization of volunteers. Any interested party can join the IETF and participate in standards development. The IETF working groups are made up of technology specialists from colleges and universities, vendors, telecommunications providers, Internet providers, governments, and a variety of interested parties. The organization is voluntary, and the structure of developing standards is much different from that used in the PSTN.

Internet standards are developed through the Request for Comments (RFC) process, which is quite efficient. In many cases, interested parties band together to jointly present new open standards for the improvement of the network. It's clear that a proposal for a new standard jointly presented by a team of vendors and providers can carry broad support at introduction. Therefore, in the real world of standards development, change can often occur very quickly in Internet standards.

Thus, we have are two different networks, with technical standards developed and approved by two separate organizations. And while that sounds straightforward, it isn't. We often draw the network as a cloud, for the sake of simplicity, because lay people often don't need to know or understand the inner workings of the public network. Although this is true, even a concept as simple as a cloud raises complex problems where interoperability is a concern, and interoperability with the incumbent network, the PSTN, is absolutely necessary for global success and widespread deployment of IP telephony solutions.

In Figure 4.1 we see three distinct viewpoints. At the top, the PSTN and Internet are distinctly different networks, with no connection or interoperability required between the two. They perform different functions and are designed for different purposes. This model represents the

Figure 4.1
Three variations of
PSTN and Internet
networks.



relationship of these two networks in the 1970s and much of the 1980s. They operated independently of one another for the most part, and many felt that the two networks would never meet.

In the center, we see the two networks joined, or connected, not by three connections, but by thousands of links. When consumers started using modems to dial up to the Internet for access, the two networks interleaved to some degree. Later, as IP telephony began to blossom as an idea, it became clear that an Internet “phone call” would have to be compatible with the PSTN in order to provide ubiquitous service. Connectivity is not enough to allow for telephony; interoperability is a paramount issue. This conclusion presented the IETF with a mandate to work in cooperation with the ITU-T in ensuring that evolving Internet standards support and remain compatible with PSTN standards, and that the two networks function together, ideally in a seamless fashion. The phrase “transparent to the end user” has often been used to describe this interworking functionality. This perspective of cooperatively working together was the dominant viewpoint of many developers of the early IP telephony standards.

During the past two years, the word *convergence* has taken on a life all its own. At a high level, we think of convergence as the migration of services like voice, data, and video to a single consolidated network like the Internet. Some people said the Internet was growing so quickly it would absorb the PSTN, and others speculated that telephony would migrate to the Internet and that someday the PSTN would just be “turned off”—a rather bold forecast.

The real truth is that both networks are vital, growing, and critical. At one point, the common view was that the two would become linked

together at high-capacity access points, passing traffic through gateways to each other. Today, they are far more inextricably interwoven than that. The two networks aren't one today, but they are so tightly interwoven into delivering the services required by business that they are beginning to represent a single cloud. This cloud is nothing more than service capacity, with services being defined through agreements between users and providers. The network of tomorrow—perhaps today—is a cloud of capacity that provides whatever service the end user requests of it. (For a more detailed review of this concept and the evolution of the networked world see Steven Shepard's book *Telecommunications Convergence*.)

As we've already mentioned, Internet standards, as developed by the IETF, are many and varied. In this chapter, we explore a few of the most commonly used and most vital standards dealing directly with the delivery of voice or multimedia services for IP telephony.

H.323 Standards for Multimedia Over Packet Networks

In 1995, the ITU-T began work on a series of standardized signaling protocols. One outgrowth of this work was a product called the *Internet Phone* from VocalTec. Initially, this was a proprietary solution and did not focus on interworking between the PSTN and the Internet. Transmission of audio signals, and the idea of videoconferencing over the Internet, were key issues. These standards fell within the H.323 family of protocols for multimedia transmission over packet networks, and were to some extent an outgrowth and extension of H.320 ISDN videoconferencing standards.

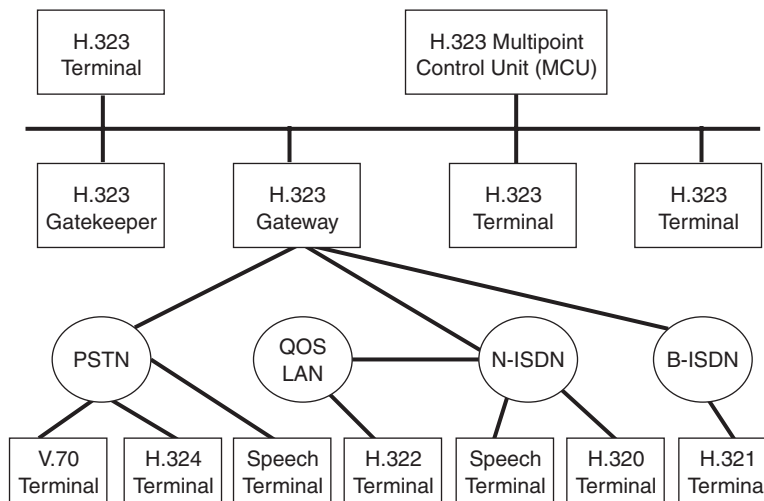
Interworking with the PSTN quickly became a major focal point of this technological development work. Designers recognized the need to incorporate some method for calls to traverse from the IP network to the PSTN. This led to efforts in gatewaying protocols and connectivity to the SS7 network, which provides extensive command and control functionality in the PSTN. *H.323* products began to appear from vendors in 1996.

H.323 embraces a set of goals that are quite simple and straightforward in principle; however, implementation has proven far more complex.

- Internetworking with the PSTN became a central theme.
- H.323 had to handle the conversion of signaling from whatever packet protocols were used to the PSTN signaling format used by SS7.
- H.323 had to have a call control mechanism for call setup and teardown.
- H.323 had to encode the media—digitize and packetize the audio voice transmission—for the IP network.

The last three functions were ideally performed in a single device referred to as a gateway, but overall H.323 encompasses a complex suite of protocols and approaches to signaling, media conversion, registration, and admission to the network as shown in Figure 4.2.

Figure 4.2
The scope of H.323.



A gateway may support one or more of the PSTN, N-ISDN, and/or B-ISDN connections.

H.323 is often referred to as an “umbrella standard” because it includes standards for a variety of purposes. Codecs are defined for the transport of both audio and video signals. Audio codecs are included for compression as low as 5.3 kbps for voice. It is important to note that “voice” can mean audio, modem data, fax messages, or touch tone signals (DTMF). Encoding schemes don’t necessarily work equally well with each of these. The *Real Time Transport Protocol* (RTP) is defined for use with both audio and video. It is used for delay-sensitive informa-

tion and includes timestamping for the sequencing and timing of packet delivery. Whenever RTP is used, *Real Time Control Protocol* (RTCP) is also used. This control protocol establishes and monitors RTP sessions.

It is important to note that RTTP uses UDP for the transmission of the audio and video packets. Because this medium is providing for delivery of real-time information, UDP is used for the quickest delivery. TCP can provide guaranteed delivery of information, but the overhead associated with TCP, coupled with the retransmission of any lost data, is too intrusive to support real-time data delivery. Therefore, a voice packet that is lost during transmission is simply lost. Testing has shown that in real-world applications, the human ear is far more tolerant of lost packets containing a fraction of a syllable than of the delays of using TCP.

The *H.225 standard* is used for *registration, admission, and status*, or RAS. An H.323-compliant terminal, upon connection with the network, registers with a *gatekeeper* in order to participate as a member of the voice network. Stations need to request network resources, make calls, and resolve the IP address of the called station. A gatekeeper often performs these functions, although the device is optional under the H.323 standards.

Q.931 signaling is used between devices for call setup and teardown. Q.931 is the identical signaling protocol used in ISDN services, complete with all the features used there. This signaling can be supported by a gateway, but can also be supported by a telco central office (CO) switch. This signaling has provided a crucial interoperability function in allowing signals to move between IP networks and the PSTN.

H.245 standards are used to provide an exchange of media capabilities between end stations. The H.323 family of protocols supports multimedia, including video. H.245 might be used when a video call is requested to negotiate for voice-only connectivity, if the called party does not have a video-capable terminal.

According to the standards, *reliable transport* is used for signaling. Thus, in an IP network, the overhead and performance of TCP are necessary. A good way to view this is to correlate this signaling to the support provided by the SS7 network in the PSTN. Without guaranteed delivery of call control and signaling message, call processing would halt and the system could not function.

The *T.120 standard* was implemented in 1996–1997 and contains protocols and services that support real-time, multipoint data communications. T.120 has often been implemented in the form of a “whiteboard” application that both users can share, but it is also used for sharing files or multiplayer gaming. One example might be two users collaborating on

a spreadsheet while talking about the changes being made. Many vendors, including Apple, AT&T, Cisco Systems, Intel, MCI, and Microsoft have implemented and support T.120-based products or services.

Figure 4.3 demonstrates the relationships between these protocols.

Figure 4.3
H.323 protocols.

Video		Audio		Control			Data
H.261 H.263 (video codec)		G.711 G.722 G.723 G.728 G.729 (audio codec)		H.225 Terminal to gatekeeper signaling	Q.931 Call signaling	H.245 Control Channel	T.120 (Data terminal sharing)
RTP	R T C P	RTP	R T C P				

Call Setup Using H.323

Many people have suggested that H.323 is cumbersome. Products supporting it have often been accused of “software bloat” because of the comprehensive functionality included and supported. Yet H.323 is widely deployed and supported today. Given that, we’ll explore what it takes to establish a telephone call between two workstations using H.323 protocols.

The scenario we review is rather simple. Two users in Figure 4.4, Bob and Alice, in the same network, on the same LAN, in the same building, wish to communicate. Bob needs to call Alice to discuss a project they are working on. We assume that there is a gatekeeper on the network for administration purposes, and step through the entire process of establishing the necessary connections and sessions to conduct a telephone call. We assume that they are both using their computers as the H.323 workstation or telephone:


1. When Bob and Alice boot up their computers in the morning, each station must send a Discovery message to locate the Gatekeeper.
2. The Gatekeeper must reply and provide its IP address

3. Each workstation must now transmit a Registration Request to the Gatekeeper and receive an Acknowledgment in return.
4. Since Bob wants to call Alice, his workstation sends a Locate Request and receives an Acknowledgment in return. This provides the IP address of Alice's workstation.
5. Bob's station now transmits an Access Request seeking resources and permission to make the call and receives an Acknowledgment in return.
6. At this point, a TCP session is established for H.225 setup.
7. Alice's workstation has received an incoming call request and now must transmit an Access Request for resource and receive Acknowledgment.
8. An H.245 Connect message is exchanged between the two workstations.
9. A second TCP session is established for the H.245 session.
10. Bob's workstation must open a Logical Channel for the media stream in the forward direction and receive an Acknowledgment.
11. An RTP media forward channel is opened, immediately followed by an RTCP control stream in the reverse direction.
12. Since the stations negotiated a full-duplex, or two-way call, at the capabilities exchange phase, Alice's workstation must open a Logical Channel for the media stream in the reverse direction and receive an Acknowledgment.
13. An RTP media reverse channel is opened, immediately followed by an RTCP control stream, in the forward direction.
14. We have established a full-duplex, two-way path for voice audio.

If this seems complex, keep in mind that H.323 supports multimedia beyond pure voice. If full-duplex video were negotiated, steps 10 through 14 would be necessary to set up those media streams. If data sharing were also required, steps 10 through 14 would have to be repeated again to establish that stream, and another TCP session would be established for T.120 data.

As you can see, establishing a call in H.323 can be overhead intensive. Although this works very quickly on a local LAN segment, where bandwidth is not a real problem, over a wide area network like the Internet, the delay in call setup can pose a serious problem. In some cases, setting up an H.323 call can take several seconds. Because TCP has timers for the retransmission of packets, the delays can be much worse if the network is suffering from congestion and packet loss.

Figure 4.4
An H.323 telephone
call, step by step.



H.323 is widely used and supported by every major equipment vendor. This support remains to maintain full compatibility with the PSTN and ITU-T standards, but designers quickly observed that the complete feature set of H.323 might not be necessary for IP telephony. And the inclusion of the complete set of protocols led to products that were large in size, and often inefficient in their coding and use of system resources.

H.323 Version 2 allows for fast connections, and supports opening media streams simultaneously, but the issue of overhead remains a concern to many designers of voice services and software.

Note that none of the discussion so far has provided a mechanism for call transfer or diverting, call hold, call parking, call pickup, call waiting, or message waiting services, all common services in business systems and expected to be present in today's working environment. Neither is there a mechanism for failover to an alternate path if a node becomes congested or does not have resources necessary to support the call.

Many designers, the author included, argue that whereas H.323 provides the necessary functionality, there are far too many scalability and performance-related issues to treat it as an acceptable solution for current IP telephony technology. It is included here because of its incumbent position in many applications and implementations, but clearly better performance is necessary for widespread success of IP telephony.

Session Initiated Protocol (SIP)

Unlike H.323, the original work on Session Initiated Protocol (SIP) was performed by the IETF as one of several different efforts. The Multiparty Multimedia Session Control (MMUSIC) working group took much of the lead in early efforts. Since 1999, the IETF-SIP working group has led this work. Their specific charter states “SIP is a text-based protocol, similar to HTTP and SMTP, for initiating interactive communication sessions between users. Such sessions include voice, video, chat, interactive games, and virtual reality. The main work of the group involves bringing SIP from proposed to draft standard, in addition to specifying and developing proposed extensions that arise out of strong requirements. The SIP working groups concentrate on the specification of SIP and its extensions, and will not explore the use of SIP for specific environments or applications.

Throughout its work, the group strives to maintain the basic model and architecture defined by SIP. In particular:

1. Services and features are provided end-to-end whenever possible.
2. Extensions and new features must be generally applicable, and not applicable only to a specific set of session types.
3. Simplicity is key.
4. Reuse of existing IP protocols and architectures, and integrating with other IP applications, is crucial.”

SIP provides protocols and mechanisms that allow both end systems and proxy servers to provide services including:

- Call forwarding under a variety of scenarios (no answer, busy, etc.)
- Calling party and called party number identification using any naming scheme
- Personal mobility allowing a single address that is location and terminal independent
- Capabilities negotiation between terminals
- Call transfer
- Instant messaging
- Event notification
- Control of networked devices

Extensions to SIP also provide for fully meshed conferences and connections to multipoint control units (MCUs).

SIP uses an addressing structure very much like email addressing. Given that a user might log on from any location and receive an IP address dynamically, there must be a way to resolve common naming conventions to the active and current IP address. Because people are familiar and comfortable with email addresses, this structure seems most appropriate and remains a popular choice.

Because SIP is a text-based protocol like HTTP or SMTP, the addresses, which are SIP uniform resource locaters (URLs), can be embedded in email messages or Web pages. Also, since this is a text protocol, the addresses are network-neutral, thus the URL might point to an email-like address, using SIP, an H.323 address, or it might point to a PSTN telephone number. The ITU-T E.164 standard defines the telephone numbering structure.

SIP provides a comprehensive set of building blocks that can be extended to allow for E911 or advanced intelligent network services.

Because it can support forking to multiple destinations, SIP can support call forwarding, automatic call distribution (ACD) techniques for call centers, and redirecting the call to multiple alternate locations.

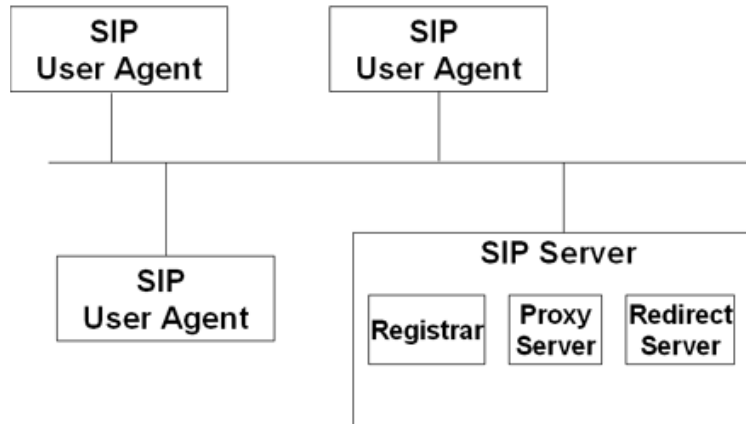
SIP operates independently of the Network Layer and requires only unreliable datagram or packet delivery. It provides its own reliability mechanism. While in the IP environment of the Internet SIP is used over UDP or TCP, it could run over IPX, Frame Relay, ATM AAL5, or X.25 without modification. Generally UDP is used to avoid the overhead associated with TCP.

The protocol model for SIP, as described in RFC 2543, is shown in Figure 4.5. It provides four different functional components:

- *User agents* (UAs) either initiate call requests or are the destination of those requests. A user agent might be IP telephony software running in a computer or an IP telephone.
- The *registrar* keeps track of users within the network or domain. User agents register with the registrar as members of the network.
- The *proxy server* is an Application Layer routing process that directs SIP requests and replies within the network.
- The *redirect server* receives requests for users (UAs) and provides the location of other SIP user agents or servers where the called party can be reached.

Within the SIP server, the registrar, proxy server, and redirect server can be implemented in the same software package.

Figure 4.5
Session initiation
protocol model.



During a SIP session, a user initiates a call, which prompts the user agent to transmit a SIP message. These messages traverse one or more SIP servers. Once the destination user agent information is obtained, actual message transfer takes place directly between the user agents. If one end of the call is located in the PSTN, a gateway between the IP-based SIP network and the PSTN is required to provide all the necessary protocol conversions between networks.

Session Description Protocol (SDP)

The MMUSIC working group of the IETF also provided RFC 2327, the Session Description Protocol (SDP). SDP is intended to describe multimedia sessions for the purposes of session announcement, session invitation, and other forms of multimedia session initiation. The Session Description Protocol is used within both SIP and Megaco implementations. SDP is not intended to support negotiation of session content or media encoding, but to act as a general-purpose tool. It also supports multicast media and can be used for broadcast environments like Internet radio or television.

SDP provides *session announcements* as the mechanism used to convey session description information between devices or nodes and proactively deliver these to users. These announcements might also be delivered via email or the Web, allowing for automatic launching of the appropriate application on the called party's workstation. SDP includes:

- The sessions name and purpose
- The time the session is active
- The type of media used in the session; this might be voice, video, data, etc.
- The format of the media (MPEG video, H.261 video, etc.)
- The transport protocol used (UDP, TCP, IP, etc.)
- Information necessary to receive the media (TCP/IP ports, addresses, and formats)

The actual syntax for the port and addressing information varies depending on the transport protocol in use. The following is an example of an SDP description:

```
v=0
o=kcamp 2890844526 2890842807 IN IP4 126.16.64.4
s=SDP Seminar
i=A Seminar on IP Telephony
u=http://www.ipadventures.com/seminar/voip.17.ps
e=ken@ipadventures.com (Ken Camp)
c=IN IP4 63.215.128.129/127
t=7944393265 8746931596
a=recvonly
m=audio 49360 RTP/AVP 0
m=video 51782 RTP/AVP 31
m=application 32416 udp wb
a=orient:portrait
```

In general terms, SDP is used to convey enough information for a user to join the session or call. This may not include encryption keys in the virtual private network (VPN) environment, which might be handled by another set of protocols like IPSec.

We won't explore the inner workings of SDP further, but it's necessary to have a high-level grasp of how the information describing multimedia sessions is transmitted between user systems.


Call Setup Using SIP

Earlier, we stepped through a telephone call between Bob and Alice. Now we'll step through that process again, using SIP as our IP telephony protocol. We'll assume each user is on a different LAN and the two networks are connected via a router. Each network has a SIP server on the local LAN, segment in this scenario.

1. At startup of the user agent software (the IP telephony package), both Bob and Alice automatically register with their local SIP server.
2. Bob initiates a telephone call. In so doing, the user agent on Bob's computer transmits an "invitation" to the SIP server on his local network. This invitation contains the session description information. In most cases, rather than using some discovery method, SIP servers are statically configured in the SIP user agent software.
3. Since Alice registered with the SIP server on another network, the SIP server on Bob's network doesn't know how to reach her. It forwards the invitation to every SIP server it knows how to reach—in this example, Alice's SIP server.
4. Since Alice is on the same LAN and registered with her SIP server, it knows how to reach her and forwards the invitation to her.
5. Since Alice also wants to talk to Bob, she answers that call, which returns and acknowledgment (ACK) back over the same path the invitation followed. Alice's session description information is included in the acknowledgment.
6. Since both ends have exchanged session description information, they have the IP address and port information to directly contact the other party on the call. They can now transmit RTP encapsulated media directly. The SIP servers do not need to participate in the call session any further.

Figure 4.6

An SIP telephone call, step by step.



- ① Registration – performed by each station
- ② Bob initiates call, sending invitation and SDP
- ③ SIP Server forwards invitation to all known SIP Servers
- ④ Alice's SIP Server delivers invitation to called party
- ⑤ Alice accepts invitation and returns SDP
- ⑥ SIP Server returns ACK to calling party
- ⑦ End-to-end telephone conversation

As you can see in Figure 4.6, the process is simpler than H.323, using fewer messages for call setup. It doesn't require TCP, which can eliminate more overhead for improved performance. Beyond that, many designers feel that the registration process with SIP servers provides better support for mobile users. And because SIP is a text-oriented protocol, a simple BYE command is used to terminate the session.

Comparison of H.323 and SIP

Table 4.1 compares some of the features and functionality of H.323 and SIP side-by-side. This is not a complete or exhaustive comparison, but gives us a quick look at some key similarities and differences between the two.

TABLE 4.1
H.323 versus SIP

	H.323	SIP
Transport Protocol(s)	Uses both TCP and UDP. Requires reliable transport.	Can use either TCP or UDP. Can also run on any unreliable packet network.
Addressing format	Allows addresses to hosts directly. Aliases resolved by a gatekeeper.	Address-neutral URL, including email, phone, H.323, and HTTP. Email-like names can be mapped to any network device.
Multicast	Supported by another H.323 set of specifications.	Caller can invite called party to join multicast sessions.
Topology	Uses gatekeeper routing and has no loop detection.	Supports fully meshed, multicast, and MCU-based conference calling, including loop detection capability.
Complexity	More complex call setup.	Simplified call setup.
Mobility support	Limited support for mobility.	Supports call redirection, call transfer, and similar telephony features.
Authentication	No user authentication is included.	User authentication can be performed via HTTP, S-HTTP, SSH, or any HTTP-like transport layer security.

continued on next page

TABLE 4.1

H.323 vs. SIP

	H.323	SIP
Protocol Encoding	H.323 uses Q.931 and the ASN.1 PER encoding.	SIP is a text-based protocol similar to HTTP.
Connection state	H.323 calls can be terminated explicitly or when the H.245 connection is torn down. Gatekeepers must monitor status for the duration of the call.	A SIP call is independent of the SIP server once established. A simple BYE command explicitly terminates a call.
Content Description	H.323 only supports H.245 to negotiate media.	SIP can use any session description format. It is not limited to only SDP.
Instant Messaging	Not supported.	Directly supported.

Both protocols reviewed so far are quite extensively supported and documented. The standards are mature enough to be widely deployed and accepted. H.323 is older and has been deployed in more mature systems. SIP, on the other hand, is very closely tied to the Internet protocols. It is fast, efficient, and, from the developer's perspective, produces tight code, which is easily readable and can produce comprehensive error codes. In the last year, SIP has overcome many hurdles and become the protocol of choice for many developers, vendors, and service providers due to its speed and efficiency.

Megaco and H.248

During the mid-1990s, Telcordia and Cisco introduced the Simple Gateway Control Protocol (SGCP). It was a single-sided model that optimized connections from traditional circuits to IP data streams. IP Device Control (IPDC) was another protocol developed by Level3, Alcatel, 3Com, Ascend, and others in a cooperative Technical Advisory Council. In the fall of 1998, features began to cross-pollinate from IPDC, and the evolution was underway.

One of the most vital points in the evolution of *H.248* or *Megaco* standards was the recognition that two distinctly different functions are performed in telephony applications. First, the call control component provides for establishing and disconnecting calls, along with cost

accounting and monitoring of the network. Second, the process of media control moves the media itself. In the PSTN, these two functions are separate, with signaling and control provided by the SS7 packet network, and call traffic handled by the trunking network between COs. With the new *Media Gateway Control Protocol*, later shortened to Megaco, the concept of multimedia calling took a different tactic, splitting functionality between a *media gateway controller* (MGC) and a *trunking media gateway* (TMG). Some engineers refer to this as a *physically decomposed multimedia gateway*. This more granular approach to functionality necessitated a protocol for communicating between these two different types of gateways.

The IETF Megaco working group worked very closely with the ITU-T SG 16 group to develop Megaco/H.248, based heavily on the roots of MGCP. The design is such that the distributed system appears as a single IP telephony gateway to the outside. Megaco consists of a *call agent* and a set of gateways, including at least one *media gateway* that performs packet-to-circuit conversion, and at least one *signaling gateway*, if connected to the SS7 network and the PSTN. Megaco can also interface with SIP and H.323 compliant gateways, as shown in Table 4.2.

TABLE 4.2Megaco
Functionality

Functional Plane	Phone switch.	Terminating entry.	H.323-conformant systems.
Signaling Plane	Signaling exchanges through SS7/ISUP.	Call agent. Internal synchronizations through Megaco and H.225/Q.931.	Signaling exchanges with the call agent through H.255/RAS Possible negotiation of logical channels and transmission parameters through H.245 with the call agent.
Bearer Data Transport Plane	Connection through high-speed trunk groups.	Telephony gateways.	Transmission of VoIP data using RTP directly between the H.323 station and the gateway.

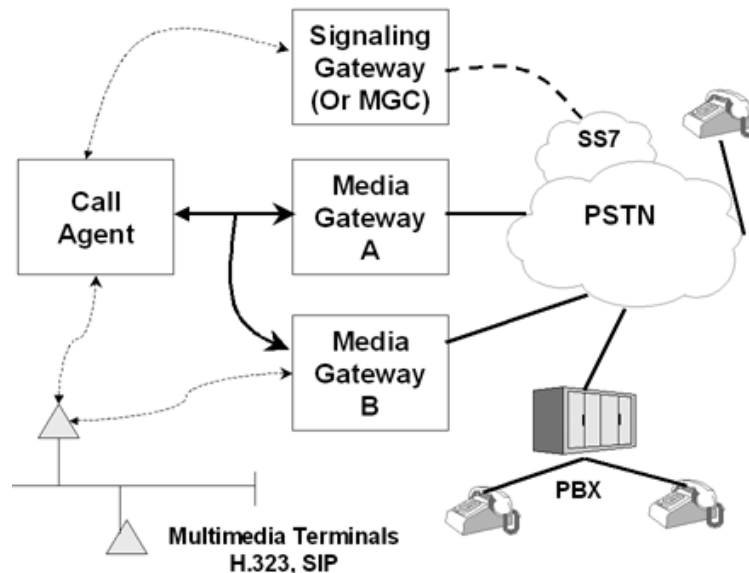
Megaco describes a distributed system consisting of one or more gateways under the control of one or more call agents (CAs). One media gateway is always necessary for transporting the information. A signaling gateway is required if the IP-based Megaco network is interconnected with the PSTN. In many cases, the call agent and signaling gateway are co-located.

The call agent plays the role of providing intelligence for call control; this is distinctly separate from the packet-to-circuit conversion provided at the gateways. The CA is responsible for signaling the gateways to either set up or disconnect a call. These signals might be sent to the IP-based packet network, the circuit-based PSTN, or both networks.

The media gateway provides the mechanisms for converting audio formats from the packet media types used in IP to those used in the circuit switched network. These conversions might include functions like converting a 64-kbps PCM-encoded conversation to a more efficient encoding scheme, such as CS-ACELP, to preserve bandwidth. This gateway might be functionally spread among several devices, or have everything contained in a single hardware device. Even if distributed among multiple devices, a single CA can control the media gateway.

As shown in Figure 4.7, the call agent always communicates with the Megaco gateways. In this communication, the SA always transmits commands, and the gateway always replies with responses. The VA can communicate with the multimedia workstation or IP telephony device using a number of protocols, including H.323 and SIP.

Figure 4.7
Megaco
communications
flow.



This arrangement eliminates any necessity for H.323 gateways. End users are free to use H.323 software applications, and the Megaco gateway is fully compatible for communication. The CA passes signaling

requirements to the signaling gateway (or MGC) and any media-specific requirements are passed to the media gateway. The MGC signaling gateway initiates any signaling necessary to interoperate with the PSTN, while the media gateway handles message coding and transmission formatting.

Megaco Terms and Definitions

Several terms and definitions must be kept in mind when working with Megaco.

- *End points (EP)* are either physical or virtual devices that function as the source of data. Physical end points require physical termination. This might be an RJ-11 jack for a telephone, an RJ-45 jack for a computer LAN adapter, an RJ-48 jack for a DS-1 interface into a gateway, or some other physical connection. Virtual end points are defined in software. In SIP, end points are identified with an email-like address.
- *Calls* are each assigned a unique ID by the call agent. Each connection associated with a call has the same ID.
- *Call agents* provide signaling intelligence.
- *Connections* are associations set up in memory between two end points. Connections might be point to point or multipoint. Connections are maintained in memory for the duration of the call. Each connection also has a unique ID.
- *Digit maps* provide a mechanism for using regular expressions to communicate so that gateways can recognize a dialing string. These prevent the timing problems associated with sending dialed information one digit at a time.
- The *Session Description Protocol (SDP)* is used in Megaco to describe sessions. It operates in the same manner as previously described.
- The *Session Initiated Protocol (SIP)* has already been discussed. The *Session Announcement Protocol (SAP)* is a multicast variation that we will not explore here.

The Gateway Commands

Megaco works because call agents issue commands and gateways respond to commands. These responses may be simple result codes, and

this might be combined with IP address and port number information for use by SDP.

The command structure used by Megaco is primarily simple and straightforward:

- **End point configuration**—Configures the end points or multiple end points within the gateway. This command is used to configure the end point parameters for requirements such as mu-law or A-law.
- **Notification request**—Sent to the gateway when the CA needs to be notified of some specific event. Examples of events might be the user going off hook at the telephone set, incoming fax or modem tones, or touch tone digit input.
- **Create connection**—Sent to the gateway to establish connections between endpoints.
- **Modify connection**—Modifies how the gateway handles an existing connection. This command can be used to alter the parameters of an existing call, to change the coding scheme for example.
- **Delete connection**—Signals the gateway to disconnect a call. If a call fails during the session, the gateway issues a response to the CA advising that the connections were deleted because of failures. The response to this command can also provide analytical information useful in monitoring network traffic load and performance, such as packets sent/received, packets lost, and average delay.

Megaco supports transport of these messages using UDP as the transport protocol. Megaco uses a combination of timers and counters to ensure information is delivered, because UDP is unreliable and does not assure delivery. By default, Megaco uses ports 2722 and 2427 to send messages. Multiple Megaco messages can be aggregated within a single UDP segment.

Because of the potentially tight integration between the IP network and the PSTN using Megaco, many common events in the PSTN may have to be monitored by the Megaco network devices. This is accommodated through the implementation of *event packages*. The event package describes what notifications the CA can request and what actions the gateway can generate when the CA issues a command. These event packages are described in Table 4.3. Detailed information is available in listed RFCs.

TABLE 4.3

Megaco Event Packages

Event Package	Name	Where Defined	Description
Generic Media Package	G	RFC 2705	The generic media package events and signals can be observed on several types of endpoints, such as trunking gateways, access gateways, or residential gateways.
DTMF Package	D	RFC 2705	Defines DTMF tones.
Line Package	L	RFC 2705	The definition of the tones is as follows: dial tone, visual message waiting indicator, alerting tone, ring splash, call waiting tone, caller ID, recorder warning tone, calling card service tone, distinctive tone pattern, report on completion, and ring back on connection.
MF Package	M	RFC 2705	The definition of the MF package events includes wink, incoming seizure, return seizure, unseize circuit, and wink off.
Trunk Package	T	RFC 2705	The trunk package signal events are continuity tone, continuity test, milliwatt tones, line test, no circuit, answer supervision, and blocking.
Handset Package	H	RFC 2705	The handset emulation package is an extension of the line package, to be used when the gateway is capable of emulating a handset.
RTP Package	R	RFC 2705	Used with RTP media streams.
Network Access Server Package	N	RFC 2705	The packet arrival event is used to notify that at least one packet has been recently sent to an Internet address that is observed by an end point.
Announcement Server Package	A	RFC 2705	The announcement action is qualified by an URL name and by a set of initial parameters.
Script Package	Script	RFC 2705	Supports scripting in Java, Perl, TCL, XML, and others.

continued on next page

TABLE 4.3
Megaco Event
Packages
(continued)

Event Package	Name	Where Defined	Description
Feature Key Package	KY	RFC 3149	The feature key package groups events and signals that are associated with the additional keys that are available on business phones.
Business Phone Package	BP	RFC 3149	The business phone package groups signals other than those related to feature keys and displays.
Display XML Package	XML	RFC 3149	The XML package contains one event/signal that is used to convey XML data to and from the phone.

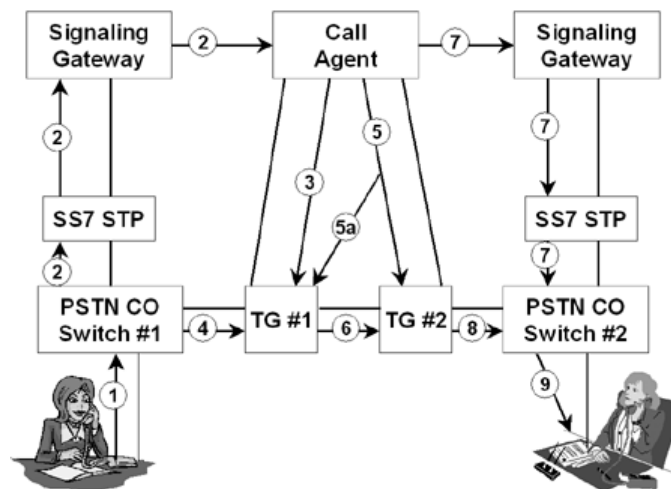
Call Setup Using Megaco/H.248

To step through a telephone call using Megaco, we'll have Alice call her sister Mary at home. In this case, we'll make the call a bit more complex, placing the users in different locations on the PSTN rather than directly on the IP network. IP will be used as a transport network between the telco COs. Figure 4.8 takes us through each step of the process.

1. Betty starts the process by going off-hook and dialing Mary's home number. PSTN CO Switch #1 makes a call routing decision and seizes a trunk leading to TG1.
2. The trunk seizure and dialed digits are signaled via a Signaling System 7 message to the SS7 signaling transfer point (STP), which passes the request on to the signaling gateway.
3. The signaling gateway passes the signaling information to the call agent (CA) using an internal protocol. In many cases, these two are co-located and this might be an internal exchange of information.
4. The CA creates the trunk connection on TG 1 by sending the Create Connection command. This results in the trunk connection between PSTN CO Switch #1 and TG 1.
5. The CA performs a call routing lookup, based on the signaling information received from the signaling gateway and sends a Create Connection command to TG 2.
 - a. Session properties of the two end points are exchanged at 5A in Figure 4.8.

6. This sets up the TG 1 to TG 2 connection and the TG 2 to PSTN CO Switch #2 trunk connection (shown at 8).
7. The CA informs the signaling gateway to SS7 STP in the PSTN about why the trunk was seized, and to send the called party's telephone number to PSTN CO Switch #2.
8. The call agent then issues a Modify Connection command to send the updated session parameters.
9. PSTN CO Switch #2 can now ring Mary's telephone, establishing end-to-end communications.

Figure 4.8
A Megaco telephone call, step by step.



This telephone call begins with Betty at a circuit-switched connection, then traverses an IP network using Megaco, and ends at Mary's telephone, which is also a circuit-switched connection on the PSTN. Their voices are encoded using standard pulse code modulation (PCM) in the PSTN. The media gateways or trunking gateways (TG 1 and 2 in Figure 4.8) convert PCM voice into IP packetized data and back.

When implementing Megaco, different media streams and types can all be handled in the same manner. The media gateway can act not only as a trunking gateway, but it might also include multimedia conferencing equipment, a voicemail system, and an interactive voice response unit (IVR). Megaco is not restricted to IP telephony either. This approach could easily incorporate a speech-to-text conversion engine for hearing-impaired subscribers.

Real-time versus Nonreal-time Traffic

We have focused primarily on voice calls and how they get made, but it is perhaps worthwhile to take a moment to consider multimedia in general. Usually, multimedia refers to a combination of audio and video, but audio and video transmissions come in different varieties.

The voice traffic we use in the telephone network is real-time interactive voice. Two people are engaged in a conversation. Issues discussed earlier, such as delay and jitter, are quality factors that must be taken into consideration when designing the network. Videoconferencing is another real-time, two-way, interactive transmission requiring special handling in delivery.

Audio and video don't have to be real time. Broadcast television, corporate announcements, and Internet radio are good examples of multimedia traffic that is nonreal-time traffic. Although the bandwidth requirements are very similar, delivery requirements are different. If a server is sending a radio transmission over the Internet, a two- or three-second delay has little or no impact on the receiver because there is no direct interaction. Delay, while present, is not a quality factor. Several other protocols are available that incorporate buffering techniques not used in IP telephony for these services. These technologies tend to deliver information from a server, or machine, to people.

Some designers tend to think of the difference in terms of one-to-one traffic versus one-to-many traffic. This is a dangerous trap to fall into because multipoint conference calling presents a variation of one-to-many, in that the "one" might change from one person or location to another during the course of the call. It is far safer to keep in mind the concept of real-time interactive communication between people when evaluating IP telephony solutions.

Which Protocol is Needed? Which is Best?

Equipment vendors are implementing H.323, SIP, and Mecago/H.248 because the standards are mature and accepted for all three solutions. In reality, all three are necessary and play a role in one place or another. Network service providers also deploy all three.

H.323 wasn't necessarily designed for the wide area network or the Internet. As a result, it can be slow, or perceived as slow. Performance problems and scalability issues follow. It performs well in the LAN environment, which means it could provide a suitable PBX replacement or campus telephone system, but a potential bottleneck exists at the connection point to the PSTN, and great care needs to be taken in implementing interoperability support.

SIP was designed specifically by Internet developers for the wide area network. It makes the assumption that there is no intelligence in the network, whereas in the PSTN, all the intelligence resides within the network. SIP is fast and efficient, but SIP alone doesn't interoperate with the PSTN. It might be the best choice for a global corporate IP network, for internal calling from site to site.

Megaco/H.248 has a limited scope, but is defined to bring the PSTN and Internet together. Because it was designed to support both H.323 and SIP terminals, it could provide the gateway and connectivity requirements for a company to integrate the two networks. The question is whether that integration is best performed at the customer premise or at the edge of the network by the provider. Both solutions have been implemented, and both work.

The choice of protocol need not be a technical decision, but rather, a business choice. Companies that are heavily dependent on the PSTN for services may be best served by a methodical approach toward integration, using H.323 where it fits. Companies that are very Internet driven, with large IP networks, may find that SIP offers a superb solution for their internal telephony requirements. Large enterprises may well negotiate with their service providers to undertake some cooperative integration efforts that provide the benefits of a single network infrastructure for the customer, but small and mid-sized companies will likely find this approach cost prohibitive.

The key to the right choice is flexibility today and extensibility for tomorrow.

CHAPTER **10**

**The Future
of IP Telephony**

IP telephony holds great promise for the future as a cost reduction measure, but most truly as a means of creating converged network infrastructures running multiple application services. We've seen a tremendous migration of the intelligence of the network in the past few years. In the mature PSTN, all the intelligence lives within the network. The end devices are telephones that have no CPU or intelligence of any kind. This environment really turned the customer into a commodity, and telephone companies focused on how many their local loops served. A direct billing connection with the customer was driven by the customer's dependence on the network. In the new Internet world, intelligence is migrating to the end stations in computers, but also to the edge of the network. Services are moving closer and closer to the user, but with dynamic IP addressing, the location of the user is irrelevant. The user doesn't have to be at a specific device or location in the new network model.

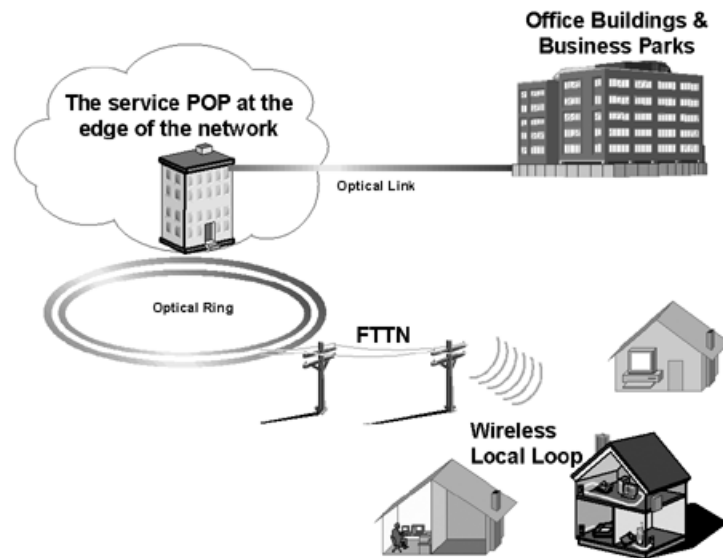
One necessity for the continued growth and success of IP telephony is a steady advance in consumer broadband solutions. The issues of the last-mile delivery of services have been the topic of raging debate for far too long.

DSL and cable modem are essentially legacy technologies, reliant on a predominantly copper local loop. Hybrid fiber-coax has aided in the advances of some cable modem solutions, but these options are both the past, not the future.

- DSL will never succeed until TA-96 regulations are modified. The ILECs have a strong incentive not to enter into widespread deployment of consumer DSL under current regulations. The Tauzin-Dingell Bill (HR1542) could impact that, but is poorly written, and unlikely to pass the Senate in its current iteration. It is not universally supported by the telecom industry.
- Relaxation of regulatory requirements alone will not assure DSL deployment. Since divestiture in 1983, the ILECs are companies with a mandate to be profitable. They have to make a legitimate business case for return on investment to deploy a service en masse. DSL often does not meet business criteria for good business. In many cases, it was deployed as a knee-jerk defense to cable company encroachment into the ILEC embedded customer base.
- DSL is too little, too late at too high a price for everyone except consumers in very densely populated areas. At best, it may have a usable life span of 5 to 7 years. If DSL ever succeeds, it will only be to extend the life of the local loop and further protect the investments made long ago by the ILECs.

Both voice and data networks are using fiber optics in the core and middle mile, with a wireless connection to the customer. This next generation of the Internet is shown in Figure 10.1. We show a “service POP” for the local provider. It’s clear that this doesn’t need to be an ILEC. In many cases, it will be an ISP, a CLEC, or some new blended provider offering voice, data, television, and any number of other possible services to subscribers. The network will use optical technology like dense wave division multiplexing (DWDM) to reach subscribers over *fiber to the neighborhood* (FTTN). The local loop will sometimes be a fiber-connected loop, particularly in the case of office buildings and business parks, where large businesses need to move massive amounts of information. For most consumers, the local loop will become a wireless connection.

Figure 10.1
The future of
networking.



802.11 wireless is extremely popular at present, yet no major players are deploying it on any great scale. There are two clear reasons for this. First, the spectrum remains unlicensed at present. That poses issues for widespread commercial viability. More important, the incumbent rarely has incentive to lead the way to a disruptive technology. Why would the ILEC move to a technology that can so clearly eliminate any revenue stream from that twisted pair local loop?

Friendly legislation and support for deployment scenarios would ease the migration to newer technologies and provide incentive for the incumbents to give serious credence to the technologies. Legislation

should foster competition, not discourage it, but good legislation must foster fair and equal competition, something we've not always seen due to political and financial influence.

Advances in optical technology and a glut of fiber have conceptually turned bandwidth into an infinite commodity, but the reality is very different. Unfortunately, the advances being realized in the core of the network have been terribly slow to pass on to consumers. DSL-like speeds at dial-up prices are the only general solution consumers will embrace.

Beefing Up IP Telephony

IP telephony has emerged from a hobby-quality niche to a competitive, viable solution that is being accepted more widely every day. The progress made over the past two years brought the technology into the "carrier class" service category, and many carriers are working at integrating IP networking into the telephony core of their networks. Areas still exist where development work is needed for the solution to grow to a healthy replacement for traditional telephony, but equipment manufacturers, standards organizations, and other players in the industry are addressing every issue identified.

Personal computer processor power still poses potential problems in real-world implementations. The typical PC sold today comes equipped with a 1.2-GHz Pentium class central processor, 128 to 256 megabytes of RAM, and a spacious hard drive. But the most common operating system in use is the Windows family of working environments, and all remain terribly processor intensive. Running H.323 has already been described as overhead intensive. SIP is far more "resource friendly," but the fact remains that using a PC to make telephone calls can still result in unacceptable performance degradation without the support of a coprocessor card. Certainly, one solution will see a more integrated telephony solution provided by Microsoft, but that isn't the best way to resolve the performance issues. A more robust approach would be to enhance PC hardware to include hardware coprocessing capabilities for encoding and decoding voice traffic. This would allow Linux and other future operating systems to include universal support for telephony. A standards-based *telephony application programmer's interface* (TAPI) can provide tools that allow tighter integration of all software. The abili-

ty to directly call a telephony software application from Microsoft Outlook, Lotus Notes, or a Palm Pilot Desktop PIM will be useful to everyone when the technology provides seamless integration.

PBX-like features have come a long way, but still have far to go. As vendors become more aggressive in the increased competition for telephony business, this area will get lots of attention. It will soon evolve into a *feature war* among salespeople out promoting their solution at your business. Evaluate features thoughtfully, and with an eye on productivity rather than “sex appeal.” Some of the hottest, sexiest features we’ve seen in software add zero value to the bottom line of a business.

After you’ve read this book, many of you will be out evaluating products and looking to implement new solutions. Before you do so, go back and review the chapter discussing performance, quality of service, and traffic engineering. Salespeople may assure you that you only need to add telephony to your existing network, no other changes are needed, and your network can support telephony with “just this one new device.” Don’t believe it. Don’t buy it. Networks are overburdened already. Many networks have more bandwidth than utilization analysis can prove a need for, yet they perform poorly. Ethernet LANs are clogged and congested due to the natural evolution of the network. Take the time to reassess the design requirements of the corporate network before adding new services. It will cost less, and your project will be far more successful.

Internet QoS must improve. Moving telephony to the Internet isn’t going to happen overnight, and it isn’t going to be a complete migration; the PSTN won’t be turned off one day when all telephone traffic has migrated. The two networks will continue to converge. The pace will increase, but they will coexist in harmony. The PSTN and Internet cannot be mutually exclusive. The differences between connection-oriented and connectionless services mandate that routers provide new services that they were not designed for. One thing that must change is that standard, out-of-the-box configurations must become more robust in support for integrated services. Packet switching has proved to be more cost effective than circuit switching, but ongoing advancements in optical networking could invalidate that reality in a very short time. Multiprotocol label switching (MPLS) and other approaches accelerate the convergence of connectionless and connection-oriented services.

Industry Changes Affecting IP Telephony

The telecommunications industry is changing forever. Several defining moments have occurred over the evolutionary lifetime of the industry. Two major occurrences in the past year or two have had a debilitating affect on the health of traditional telephony. In what is often referred to as the “Internet dot-com bubble,” we saw *optimistic euphoria*. To be charitable, that’s a really nice way of saying investors were blindly willing to fund some pretty hare-brained schemes in the hope of making a ton of money from an Internet business. Simple economics have since proved that nothing in technology can replace the value of a solid business plan.

The telecommunications industry followed in exuberant support of untold wealth in many ways. Equipment vendors overproduced and stockpiled inventory. Partnerships involving speculative future business formed. Dark fiber was leased back and forth between companies, inflating revenue projections. And because the industry is very complex, many people, even industry analysts, remained optimistic because of the successful track record of so many major players.

In 2001, the telecom industry suffered terrible losses and the wildly optimistic projections, one by one, failed to come to fruition. Thousands of workers, as many as half a million, found themselves laid off, hitting the job market with a special skill set that suddenly had far greater supply than demand. Companies’ stock values declined and at a heart-stopping pace. We saw what was perhaps, at least to that point, the worst decline any industry had ever encountered in so short a time.

Most recently, we hear stories of greed and avarice in the industry. Investigations by law enforcement and charges of corporate fraud have rocked the investment communities. Chief Executive Officers are being called to testify before the Senate and the Securities Exchange Commission (SEC), and, rather than openly discuss the situations, many have invoked their rights under the Fifth Amendment and refused to testify. Whether this is an indication of any actual wrongdoing is irrelevant, but what it does signify is an industry in the throes of meltdown. Investor confidence is at the lowest point seen in many years.

It is of interest to step back and take a look at the overall health of the industry. Major vendors have all struggled with changes in technology. Equipment manufacturers have cut back production, laid off workers, and scaled down estimates of earnings. Providers aren’t buying

products and, without exponential leaps in innovation fueled by substantive research and development work, that scenario won't change any time soon. The salvation of companies making equipment lies primarily in innovation. As costs have declined and profit margins have shrunk, innovation has been left to smaller companies playing in niche markets. These small companies are often swallowed up in corporate acquisitions as the larger vendors look for new products.

A look at the service providers and the major carriers isn't any more optimistic. WorldCom and Qwest alone represent a tremendous portion of telephony revenues, and their struggles have captured the public eye because of potential wrongdoing. Global Crossing has filed Chapter 11 bankruptcy and is trying to find a reorganization approach that works. These carriers all sold capacity to one another at a time when nobody was selling in volume to end customers, and that's where the real revenue comes from. AT&T stock recently reached the lowest point since the breakup of the Bell System in 1984. Telecommunications is simply not a healthy industry.

Interestingly enough, three notably healthy companies remain as of this writing. Bell South, SBC, and Verizon remain strong and financially viable. One former executive in the industry noted that the industry might really be a monopoly business just because of the huge amount of capital required to succeed. The local giants that remain, whether by coincidence or not, are vestiges of the former Bell System. They may not even need to buy out their failed competition, but the recession has cut their revenues, and they've been consistently losing the battle for broadband customers. The question now becomes whether or not they can innovate and bring the new solutions required to the market.

All this leads to a climate of innovation that is ripe for change. And radical change comes from without, rather than within.

Disruptive Technology Takes Center Stage

Perhaps the most defining text lately addressing disruptive technologies has been *The Innovator's Dilemma* by Clayton Christensen. It describes the differences identified between sustaining and disruptive technologies. The telecommunications industry has provided innovation over the course of its history, but for many years now, has rested firmly on the laurels of prior achievements and settled into the comfortable rut of sus-

tainability. Sustaining technologies provide incremental progress and small enhancements to existing technologies. This is precisely what the traditional telecommunications providers have delivered to the market for the past 20 years or more.

Disruptive technologies are by nature, very different. They arrive significantly lower in the market chain, perhaps even focused in other markets. They are simple and inexpensive. Disruptive technologies are often lower-performance solutions than the incumbent sustaining technologies, and they offer lower profit margins. Historically, mainstream companies tend to look down their noses at these disruptive upstarts, often ignoring them entirely.

The classic example of a disruptive technology is the 5.25-inch Winchester disk drive introduced in 1981, which dominated the PC market for several years. Prior to this, disk drives were much larger, in 8- and 14-inch varieties. Large computer companies ignored the 5.25-inch drive because it was slower, of lower quality, and had a smaller storage capacity. This was clearly not a product that could sustain the mainframe computer market. Rather than fade away, we saw these inexpensive disk drives find a fit in the desktop PC, where the lower cost was a benefit, the performance acceptable, and the capacity not an issue (at that time). In 1981, there were four leading makers of the 8-inch drive, but by 1985, three were gone from the business. Of the 14-inch drive manufacturers, not a single one survived. What changed wasn't the sustaining technology, but rather an evolution of computing in general, fueled by smaller, less-expensive components and a fundamental change in the way people used computers.

Market-leading companies experience tremendous difficulty in embracing disruptive technologies. It's the nature of their business to support existing customers and sustain an ongoing revenue stream. This requires the adoption of widely deployed standard technologies. The motivation isn't to innovate and do something completely new and untried; the motivation is not to rock the boat, but rather, to maintain status quo.

Disruptive technologies lead to new markets, which are difficult to analyze and harder still to project. The traditional old-school planning methods used in the telecommunications industry make this task nearly insurmountable because the value created by the disruptive new technology is incompatible with the business model and processes, even the corporate culture, of the incumbent providers.

The telecommunications industry didn't begin as the mature, even lethargic behemoth it is today. It began as a disruptive technology, dis-

The Future of IP Telephony

placing the telegraph system with a completely new, real-time, interactive communications tool. But no technology can remain disruptive. Once it becomes mainstream, as the PSTN did long ago, it becomes a sustaining technology. And once the sustaining technology matures, incremental advancements slow, and the door to disruption from outside opens.

Let's perform a quick side-by-side comparison of traditional telephony and IP telephony and see just how disruptive this solution might be; see Figure 10.2.

Figure 10.2
Mainstream telephony versus disruptive IP telephony.

Traditional Telephony PSTN	IP Telephony VoIP
Deeply entrenched in mainstream PSTN	Overlooked as unnecessary by many PSTN providers
Well established value model	Requires a new value model
High performance	Lower performance
Low cost due to years of experience	High cost to deploy. New systems, training, methods
Large established market	Unknown and untested market
Large companies cannot ignore the technology	Large companies can easily ignore the technology
Mature, well-developed products	New, evolving products

IP telephony pretty clearly sits in a position to seriously disrupt the telecommunications industry. Called minutes billed to IP telephony are rising at dramatic rates, with expectations of an ongoing multibillion dollar revenue stream. Over the past year or two, ongoing advancements have moved IP up the market chain, making it more and more viable and worth serious consideration.

Instant messaging (IM) applications play a role in this disruption activity as well. This tiny application has become one of the most widely used utilities on the Internet. As companies discover employees using IM in productive business ways, they will embrace it more and more. There's a fairly tight synergy between IM and telephony, with the ability to shift from short messages by simply clicking a button to speak to someone. These applications will converge to some degree over the next year or so as both continue to mature.

IP telephony services and applications are poised to move quickly from a minor, yet significant portion of the telephony market to a major and growing share. As the existing providers struggle to find a new identity in the market, and fight to regain consumer and investor confidence, the providers that only offer IP-related service have an opportunity to seize major market share. The advancements in features and technology are continuing at a rapid pace. The products are maturing, and while they haven't fully matured, an analysis of cost versus benefits will often tip the scales solidly toward IP telephony as the best solution for many businesses. That trend will continue, and IP telephony will very likely become at least equal to traditional telephony in the most technology-oriented markets (North America, Europe, and much of Asia).

This is a pretty clear wake-up call to the telecommunications industry: there's a barbarian at the gate, and it's called IP telephony. As users find new ways to implement and use the technology, this particular barbarian could overrun the city, leaving few profitable survivors in the traditional telephony sector.

Internet Call Centers

One implementation of IP telephony that will grow in the next year is the *call center*. We've seen many changes in how people work; more and more people focus on family and friends and want to change how they participate in their work environment. Internet call centers provide an approach that offers a work-at-home job, but with far more connectivity than we've previously seen.

Call centers have historically driven jobs to areas that offered a lower tax base to businesses. Many U.S. companies now use offshore call centers in other countries, where the labor rate is lower; however, the distributed call center approach alters the cost structure and provides an effective method to hire domestic staff around the country. With the growing e-commerce environment of the Web, call centers have proved to be a growing segment of the market, yet traditional telephone technologies still present a barrier.

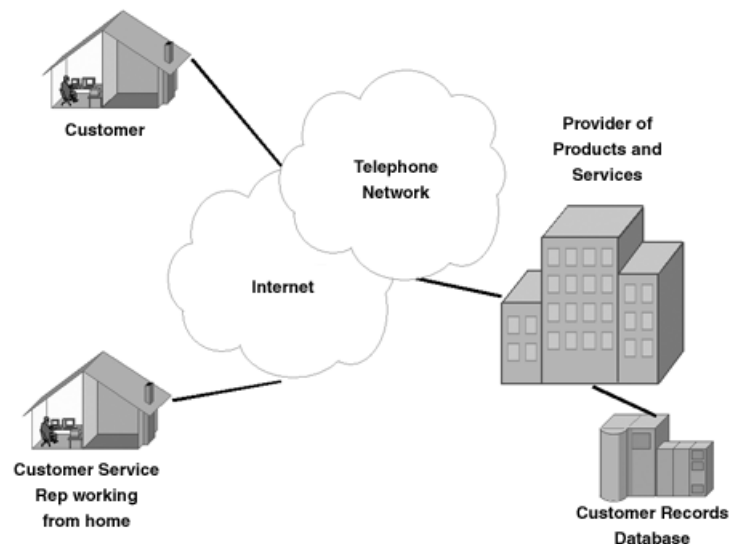
The leading approach for home agents has historically been ISDN. The integrated services digital network still doesn't provide integrated services. The bottom-line costs wind up being too high for most companies to invest in this architecture. A shift to IP telephony reduces the

cost and provides better integration of services than ISDN has ever offered. Many companies are now looking at IP telephony, coupled with DSL and cable modem solutions, to fill future call center requirements.

Perhaps the greatest vision of promise in call center technologies was driven by the burst of the dot-com bubble. A large number of companies tried to enter the Internet market and sell all manner of goods on the Internet. These businesses had no “brick and mortar,” which was one of the attractions, but many also had no history in providing customer service. Some succeeded, but even Amazon, one of the most successful examples, has yet to become a profitable company. Most attempts failed, and many failures related directly to the inability to provide proper customer service. The focus of any company doing business on the Internet has become customer service. The demand for customer service is best filled by the ability to speak with a live person. The call center, in any form, provides that ability.

Consider Figure 10.3; the telephone network and Internet are both represented, but linked together much as they are in the real world. Customers might be at home or anywhere. They might contact the provider via telephone or via a Web site. According to the Gartner Group, more than 70 percent of transactions take place over the telephone. Those Web sites that have been able to implement live voice support for potential customers report as much as 50 percent increases in sales. Online worldwide revenues from retail sales in the United States are anticipated to hit \$35.3 billion this year.

Figure 10.3
Merging the PSTN
and the Internet.



The provider of products or services receives a query for customer support and through distributed call center technology, is able to redirect that call to a customer service representative working from home.

The staffing of call centers has always been a difficult issue. The ability to hire remote staff, perhaps even part-time remote staff, allows the provider, regardless of location, to find good, qualified employees. Time zones become a non-issue. Even customer service reps with special linguistic skills become obtainable resources. The benefits to the telecommuter work force have been studied time and again in the past 10 years or more, and everything suggests continued migration to work-at-home efforts.

This solution also provides an opportunity for a pool of workforce candidates that may have been inaccessible in the traditional call center. Stay-at-home mothers, retirees, even people without transportation now become potential job candidates, participating in the workforce in ways they may not have been able to previously.

Distributed call centers do not require IP telephony, but it does provide the greatest level of integration at the lowest cost. The distributed call center has often been implemented using PBX solutions and off-premise stations or ISDN lines. ISDN can be expensive, often more so than DSL or cable modem, and it is not at all ubiquitous. The newer generation of IP and Internet technologies make the distributed call center more cost effective to implement than it has ever been in the past. The use of IP technology as a “PBX extender” creates a virtual call center environment that can physically be anywhere. Employees needn’t be passed over because of their geographic location. To customers, the company presents a single unified point of presence.

A broadband or high-speed connection capable of supporting voice and data simultaneously is optimal, with DSL the technology of choice for most companies implementing these call centers today. The “home office” or heart of the distributed call center is typically traditional call center technology but, as we’ve already seen, it is easily replaced by an IP solution.

Distributed call centers use a job performance technology that requires managers take a more “hands-off” approach to supervising workers than traditional workflow methods. Supervisors learn to rely on the systems, both telephone and computer networks, to measure and monitor productivity and worker activity. The idea of a worker in the corporate office where work can be directly observed is transformed into a measurement of productivity and results rather than activity.

For some companies, the challenge of remote teleworkers becomes one of combating isolation. Because workers don't have the social interaction with colleagues (i.e., water cooler chat), supervisors must adapt a management style that encourages interaction at all levels to minimize the chance that remote workers are left "out of the loop" and disenfranchised from the business of the company. Conference calls, computer video conferencing, and regular visits to the office or meetings with colleagues have proved effective in overcoming this issue.

The distributed call center provides a very attractive alternative for companies in large metropolitan areas that often have alternative commuter requirements because of air-quality management regulations. This solution can aid in bringing a business into compliance with these regulations.

Although there has been a trend to move call centers offshore, today many companies have become very security conscious and are more reluctant to engage in offshore arrangements. The distributed IP telephony call center allows for substantial savings above traditional costs without sending jobs outside the United States, although that option can also work using IP telephony.

The benefits of distributed call centers can be measured in a variety of ways. Some benefits cited by companies implementing this technique for managing a call center business include:

- **Reduced cost of office space**—Because teleworkers do not require a cubicle or workspace in the corporate office, that office can shrink, thus reducing real estate and associated costs. In reviewing real-life implementations, the companies interviewed estimate that building costs alone were recouped within 3 years by shifting to a distributed model.
- **Location**—From the mid-1970s through the 1980s, a migration of call centers from major metropolitan areas to more rural settings provided a readily available workforce and a reduced-tax incentive to business. In particular, Omaha, Nebraska provided incentives to companies considering relocation and successfully attracted many call center businesses to the area. Using the distributed model, corporate location is irrelevant, and so is the location of the teleworkers.
- **Tax benefits**—In major metropolitan areas particularly, mandates are in place that require businesses to encourage alternate commuting methods. Car pooling, bicycling to work, and telecommuting are all proven alternative commutes in the proper setting.

Any organization that interacts with a distributed base of customers over the telephone may have need for a call center:

- Healthcare providers, particularly large organizations and HMO/PPO groups, needing to handle patient calls
- Insurance carriers dealing with customer calls regarding policies
- Catalog or Internet-based merchants with a high volume of telephone transactions
- Airline, hotel, event registration, and ticketing agencies
- Financial investment firms, particularly those dealing in high volumes of telephone calls
- Social services organizations of many types providing call-in services for their clients

In general terms, the call center is probably not an appropriate technology for the small business. This solution is the best fit for high-volume, transaction-based services that require interaction with a customer service representative. On the other hand, a small business with an IP Centrex solution can implement call center services easily to provide new levels of customer service. In short, IP telephony brings the call center into the reach of a small company without the burden of a full-blown, expensive implementation. IP telephony not only disrupts the telecommunications industry, but also can level the playing field for smaller companies, providing tools that, in the past, have only been available to businesses with a large budget and staff.

Fax Over IP

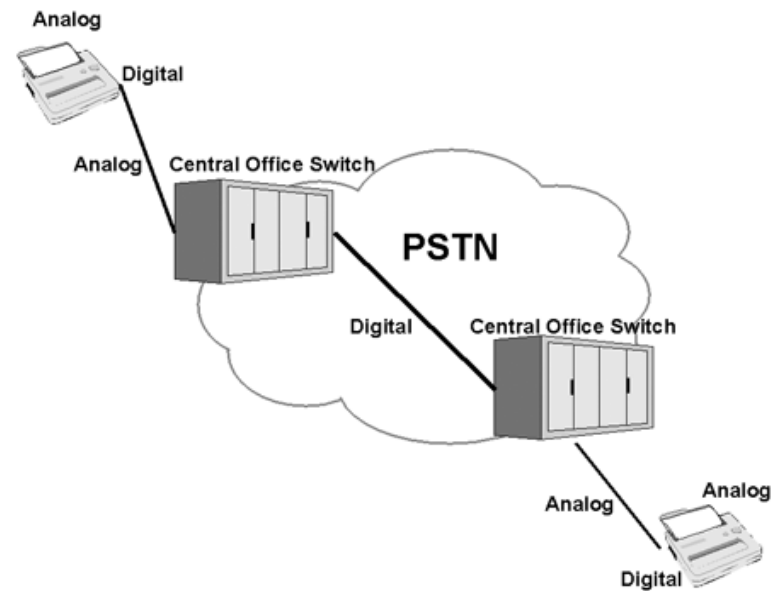
As estimated 80 million fax machines are in use in the world today. In many small businesses, the fax transmission of orders is standard routine. Yes, when we think about fax traffic, it really isn't real-time traffic requiring a circuit-switched connection. Fax traffic isn't delay sensitive, at least in the sense that a few seconds delay won't damage the integrity of the transmission. The question is whether or not it really makes sense to tie up circuit-switched connections in the PSTN to send fax documents.

Matters are complicated further by the document scanning and transmission process. At the scanner, the fax machine reads an analog document from a piece of paper and digitizes (or rasterizes) it into digital form. Now that it's been converted to a digital bit stream, it's ready to

The Future of IP Telephony

send, but the medium we use to transmit fax is a plain old analog phone line. To send digital information over the analog line, we have to modulate it with a carrier using a modem. When it hits the local CO switch, it then has to be sampled and converted for transmission over digital carrier trunk facilities through the network. At the receiving end, a modulated carrier is sent out over the analog local loop, which the modem in the fax machine must convert back into digital data. This digital data is then used to recreate a facsimile of the original by printing out on an analog piece of paper. As you can see in Figure 10.4, this entails six analog-to-digital conversions to transmit one page of information. Why not digitize the document once and complete the whole process digitally?

Figure 10.4
Analog-to-digital
conversion for fax.



In many cases, we avoid the fax machine entirely and send documents as digital attachments in email. When we do this, we eliminate the conversion process and the need for circuit-switched resources. But we don't always have an electronic version of the document we need to send. Sometimes it's a copy of a magazine article, a sales brochure, or a preprinted contract. Since we're spending several billion dollars a year sending this sort of information, there's clearly an opportunity to take advantage of the technology.

New solutions reduce this problem: fax servers and fax gateways on LANs are more and more commonplace; and fax-service Web sites are

very popular and becoming more so. (Readers are encouraged to look at www.efax.com and www.jfax.com for examples that can provide packetized fax service, which not only provides saving on paper manufacture, but provides a means to move fax off the PSTN and transmit via the Internet.)

There is still room for enhancement in IP fax technologies. This is an area where billions of dollars are spent every year, as much as 20 percent of toll and long distances revenues. The opportunity for cost savings is huge for companies that transmit or receive high volumes of faxed material.

Voice Mail over IP

Voice mail has become such a commonplace tool that many people take it for granted. The days of tape drives to store messages, even on answering machines, are long gone. Stored voice mail messages are now digitized voice stored as data files on some form of memory device. In corporate voice mail systems and telco service provider networks, messages are stored on a hard disk.

One question, as the Internet and IP telephony progress, is just how many voice mail systems do we need? How many do you have? Many people have voice mail at home and at work, but then there's the home office. And what about voice mail as a standard cellular service? The technology issue is that, given multiple voice mailboxes where this information is all stored in digital form, how many telephone numbers does the user need to call to retrieve messages? Two? Three? Four? And why? Why not just build a user-programmable interface that delivers the digitized file to a central point for retrieval? Why not let the user program in an email address and deliver messages to email as an audio attachment?

Like fax service, voice mail delivery isn't delay sensitive. The message needs to be delivered intact, but even a few seconds delay in delivery won't harm message integrity, as long as the entire message is delivered. Voice mail isn't an interactive service and need not be treated the same as real-time traffic. Like other forms of data traffic, it is sensitive to errors, but not to delay.

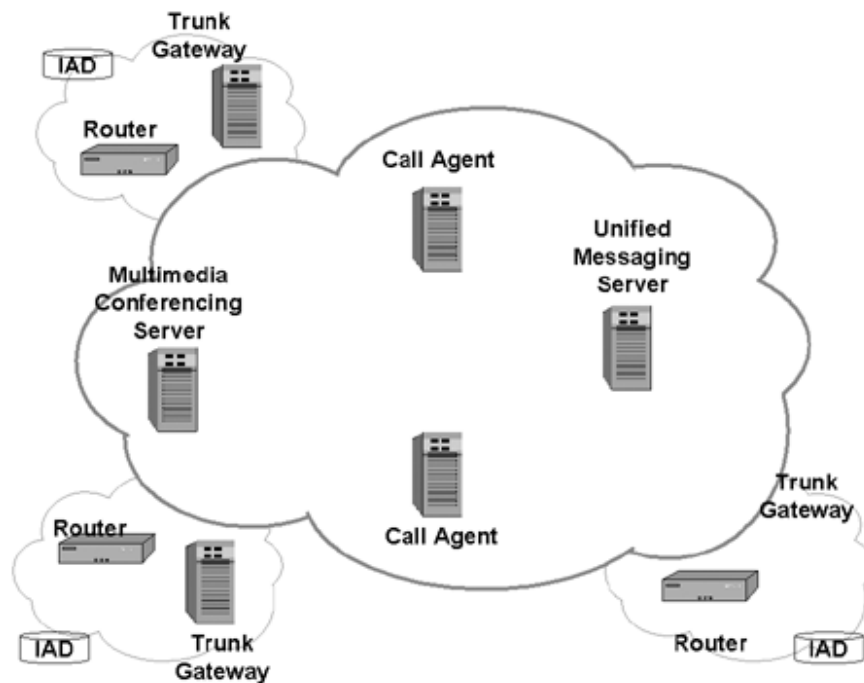
To some people, this still sounds a bit farfetched, but readers might wish to surf the Web to www.onebox.com or www.address.com to see live examples of just this type of service.

The Future of IP Telephony

Several streaming audio technologies work very nicely at delivering this sort of information, both proprietary and openly published. Anyone who has ever listened to Internet radio has heard the technology. It may not be interactive voice quality, but with some buffering, even dial-up modem connections to the Internet allow for the delivery of voice mail messages using these protocols and techniques.

Unified messaging is discussed, but not truly available...yet. It will be. This service offers promise in the next cycle of advancements to IP telephony-related technologies. This advanced next generation network of IP telephony is shown in Figure 10.5. SIP and H.248/Megaco are the control protocols within the IP network and operate between the call agents, the servers, the trunking gateways, and the integrated access devices (IADs). Media streams are handled by a combination of RTP, MPLS, and other protocols (like MPLS in conjunction with ATM) between the routers, gateways, and integrated access devices. SS7 is used between the Gateways in the IP network and the PSTN. Connecting lines have not been included for all of these relationships because the illustration would become too busy, but as you envision each, you can see that the next generation IP telephone network is a very sophisticated and complex environment.

Figure 10.5
Next generation
IP telephony.



IP Telephony and Denial of Service (DoS) Attacks

Some concern exists in the Internet community about a recent series of problems referred to as denial of service (DoS) attack. This discussion is not designed to convince anyone that DoS attacks are not a real and growing problem.

In the IP network, every device has an individual address that allows it to become a fully participating member of the network. This host address is unique for every connected device. In the Internet in particular, addresses are widely known, even advertised in some fashion. The result is that if someone with malicious intent knows the IP address of a potential victim, it is possible to flood the victim's address with more traffic than can be processed, thereby denying service to legitimate users. DoS attacks are often much more sophisticated, but for purposes of discussion, this simple description is adequate.

The danger to the Internet user is that often there is nearly no defense against denial of service attacks. If a company needs a public Web site, that site must have an address users can reach. The network has no way of identifying undesirable traffic, thus making it very difficult to eliminate threats before problems occur. It's difficult, but not impossible. Many intrusion detection systems, firewalls, and routers can detect some predefined conditions that may indicate an attack is in progress and take defensive action. Nonetheless, the bottom line is that any host connected to the Internet is a potential victim of this type of attack.

Denial of service attacks are further complicated by the attackers' tendencies to use falsified, or *spoofed*, IP addresses to hide the true source of the attack. Attackers have gone so far as to initiate *distributed denial of service* (DDoS) by loading a program onto remote computers that allow the attacker to seize control. These programs, referred to as *zombies*, can even allow an attacker to use an unwitting victim's computer to attack a third party.

Although this sounds like a potentially fatal flaw in using the IP environment or the Internet for telephony, that simply isn't the case. It has been argued that users on the PSTN can't be blocked from service in this manner, but caller ID information can be used to block the offending caller and prevent blockage of the user's telephone. And while this is true enough, it's a reasonably safe assumption that given a large enough bank of pay phones, rolls of quarters, and accomplices, denial of traditional telephony services is actually quite achievable.

Detractors speculate that this vulnerability presents an insurmountable barrier to IP telephony and makes the solution no longer useful. Not really. It presents a challenge and a responsibility to network designers and administrators. Designers are responsible for the intelligent, fault-tolerant design of networks. Networks cannot be allowed to just multiply and grow without planning, as they did in the past. Network administrators must also embrace the need for closer monitoring and auditing of network traffic. No longer can the network be turned on and ignored, with the assumption it will run error free. The network is no longer merely a simple Ethernet LAN, but a complex organism providing multiple services and requiring much of the same attention to security, traffic engineering, and monitoring as the PSTN has long required.

The Future for Equipment Vendors

There's really only one guiding principle for equipment manufacturers to focus on for the future of IP telephony: innovation. Don't duplicate; innovate. Incremental improvements are expected and necessary, but don't expect them to sustain an equipment product line for years. Issuing the next "dot release" of software that adds two new features isn't innovation; it's stagnation. Push for new ideas and encourage an entrepreneurial spirit within your company.

Even the largest, most widely accepted products are subject to disruption. Several vendors have already seen the impact of leaps in improvement by competitors. Exponential advances quickly become incremental sustaining advances in technology. Research and development work must be encouraged and cultivated, even knowing that the old funding models for such activity no longer work. Venture capital for speculative efforts won't come easily. Cultivate small pockets of innovations within your company. And don't punish mistakes. They will happen as part of the development process. To paraphrase Confucius, he who doesn't make mistakes, doesn't make anything at all.

Let's consider a simple theory in technology equipment: acquisition speeds the path to stagnation and stifles innovation. The trend of the past ten years is to let small companies innovate, then, as technologies prove themselves, to purchase the innovators. It sounds reasonable, but there's a two-fold penalty to pay for this approach. First, the new technologies come from different processes and may not integrate well into

your existing product lines. We've seen this scenario played out time and time again as a vendor's product line identity shifts, trying to accommodate new solutions that were developed with a different approach. A far more dangerous penalty comes from the message this sends to your own internal team: we don't innovate here; we look to the outside. This approach stifles the creative spirit of your staff and, in the long run, motivates creative designers and entrepreneurs to move on to other opportunities. Given the speed of technological advancements, driving innovation out is the surest way to propagate mediocrity. It might not hit today, but it's a sign that a company is falling into the sustaining technology mindset and becoming a candidate to be overrun by another disruptive technology.

Equipment manufacturers no longer make just one or two components. Now, single vendors provide the router, gateway, softswitch, integrated access device, the LAN hubs and switches, the firewall, even the telephones. Manufacturers must attend to not only the design and management of their systems, but to the simplest things. Standard configurations in this equipment should reflect a reasonable and realistic approach to design. Make the defaults more like what customers really do. Make it easier to fine tune configurations to fit the implementation. Provide simple voice services in even the low-end products, and simplify their configuration everywhere possible. Simplify and integrate more tightly, while adhering to open standards of interoperability.

Don't give your customers a goeey GUI. So many vendors think the more complex a GUI is, the more complete it looks. Some products in the market appear to have more development work applied to the user interface software than to the actual product, so here's a word of caution. Treat your GUI interface as if it were the Web. Two or three clicks should take a user to where they need to be quickly and easily. Don't make customers hunt through layers of menus and nonintuitive options trying to find out how to add a user or other simple task. Simplify!

Customers don't want to buy more of the same thing. For equipment vendors, telecommunications carriers and ISPs represent the largest customer segment, so look not to their success, but to their failure in technology. Give them solutions that do the things they can't do today. Resources galore identify where end user needs are not being met. Long-term success is not found by creating new business models, but by creating the products that serve customer needs.

The Future for Telecommunications Service Providers

Like equipment manufacturers, service providers must innovate. The telecommunications service industry has become a gargantuan, lethargic beast. It's true that we've seen large major carriers undertake widespread deployment of cellular services to provide telephone connectivity anywhere, and we would argue that cellular services are sustaining advances in the traditional telephone service model. They're an improvement, an increase in availability. They may even be something of a large step, but it's still an incremental large step. Cellular service in the United States has become a sustaining technology, and new advances are incremental chest-thumping attempts to gain market share. While integrating every square foot of the United States into a coverage area improves the service and increases the customer base, it isn't innovation.

Managed service offerings hold great potential for the future. Today we see managed VPNs, managed firewalls, and managed networks. Web-hosting services are offered, but in most cases they are still immature. Providing a building with power, fiber access, and a secure environment isn't really hosting—it's real estate management. Sell services for what they are, and explore the new requirements. Offer a managed data center integrated with remote hosted IP PBX services that scales from a company of 10 employees to a company of 100,000 employees. Manage it wisely and well. Don't overcommit. Don't sell smoke and mirrors. Build a valuable managed service by asking customers what they need, then provide it. Building solutions based on what looks profitable and manageable creates nice, clean service offerings that fail to meet customer needs.

Speed is of the essence. A recent implementation of a major provider's offering of a managed service for fewer than 30 locations took nearly one year from initial proposal to cutover. This wasn't a new offering, but one that is widely deployed. This plodding, inefficient process hurt both the provider and the customer. The provider didn't see any revenue from the new service for nearly a year, but invested tremendous resources in the interim. The customer didn't gain any advantage for the whole timeline. Speed and efficiency in providing managed services are critical to their success. Providers must demonstrate their proficiency at managed services or lose customers to competitors.

Telecommunications carriers need to look to small competitive ISPs and CLECs. See how they're succeeding and take advantage of their efforts. Take the things they do, do them better and improve upon them. Your biggest disadvantage may be your size. A corporate telecom elephant may find dancing the jitterbug very awkward. Foster a corporate culture that is lean, mean, and responsive. Your smaller competitors will always be faster and probably always more innovative. Don't rest on the laurels of your embedded customer base and think that you can catch up later. That's living in denial. Later is too late.

The incumbent telecommunications carriers in particular should recognize DSL offerings for what they truly are: a last-ditch effort to maintain the customer relationship over the twisted pair. It's time to accept that the twisted pair is a legacy of the past. Hang on to the copper, and you're tied to the past. It's time to invest research and development efforts in wireless technologies. Some wireless solutions available today have been tried and proven questionable. Look at them again and re-evaluate service deployment models. Investigate 802.11 WiFi technologies, not for carrier class deployment but for what it is—a simple LAN technology. But the groundswell of popularity is rapidly proving that it meets customer demands. Look to the future rather than the past.

ISPs and CLECs need to look to the telecom industry. Service level agreements and guarantees of delivery are areas that Internet technologies have always been weak in. The major telco carriers have an advantage in their large infrastructure, because they can often provide services that a smaller company struggles to compete against. Your strength is your size. Smaller companies tend to have an ingrained competitive and entrepreneurial spirit that encourages new ideas and new thinking. Seize the hot new markets quickly. Work closely with your customer base to give them what they need before they need it. Understand and anticipate your customer's requirements before they do, and you will win customers faster than you can imagine.

The convergence of voice and data networks is a reality. It is what customers want. Give customers what they want to succeed. It's a very simple business tactic. The next generation of wireless, integrated tightly with IP telephony, will not only disrupt the traditional telephony market, but the cellular market as well. (The wireless PDA has struggled to find solid acceptance.) When we think of the Internet, we tend to think of the World Wide Web, but there's no reason that information can't be accessed via other methods. Couple IP telephony with wireless networks. Tie in instant messaging. Integrate data query using location-

based services. Yes, it sounds a lot like 3G wireless, but consider how disruptive it is to the traditional telecommunications industry.

The Future for Business Users

This book is about a technology, but the technology is merely a tool that brings us a step closer to full convergence of voice and data networks and services. Let's consider just why this type of convergence is important to business.

The two most valuable assets a company has are its people and its intellectual capital. Intellectual capital might be housed inside network servers and data center archives, but it's often most commonly found inside the staff of employees. Data, information, and knowledge are more valuable than inventory and cash reserves. Yet even this mindset fails to account for human capital. The information model in Figure 10.6 takes this concept one step further.

The raw data is easily turned into information that tells us Ken is 6 feet tall, lives in Vermont, and has a dog named Zoë. Even automated systems can incorporate logic rules that can extrapolate knowledge from this raw data. We know Ken is easier to reach by phone than email because we have a telephone number, but no email address. Wisdom is only gained through the human factor. Only personal conversations on a human level discover that Ken takes Zoë for a walk in the woods or to play in the lake in the late afternoons. This *wisdom* cannot be extracted by machine. It requires the one computing device that can think extemporaneously, building relationships between data elements dynamically...the human brain.

Figure 10.6
How raw data becomes knowledge.

DATA	Ken, 6, Brown, Vermont, Dog/Zoe 802-555-1234
INFORMATION	Ken is 6 feet tall, has brown eyes, lives in Vermont and his a telephone number is 802-555- 1234
KNOWLEDGE	Ken is easier to reach by phone than by email
WISDOM	It is best to not contact Ken in the late afternoon
Raw data isn't much use if you don't know what to do with it, but a good system of sharing knowledge and wisdom can be the secret to better efficiency and higher profits company-wide.	

Raw data is easily collected. In the past, it often went into archival storage, buried in a tomb from which it could never be retrieved. The past few years have focused attention on *data warehousing* concepts, which rapidly evolved into *knowledge management solutions*. Wisdom remains in the human domain, outside the reach of computers.

For a business to take advantage of the wisdom in human capital among its staff, that business must give the staff tools and knowledge upon which decisions can be based. The closer at hand the tools are, the more efficient and productive a company will be. To bring the tools closer to hand, create a converged environment where access to any piece of information and any communication need is available instantly.

Business customers have a great challenge, because new technologies cannot be overlooked. They must be evaluated, and a business case analysis performed, but the analysis must be done swiftly. Implementations must be quick and efficient. Don't place unwarranted faith in an equipment vendor or provider. You are managing your business; they are managing their business. Insist on performance commitments and service level agreements. Don't assume you're getting what you're paying for. Monitor your systems and know what you're getting.

Managed services may offer the best hope for many businesses. Evaluate them closely. IP Centrex clearly needs to be called something else. Coupled with data center hosting services, it provides the truest service convergence a provider can offer. Today, the offerings are limited, but they're improving. Work closely with service providers. Above all, be vocal. When offerings fall short, make the shortcomings known. The world of telephony is far more than a dial tone now. Don't settle for a dial tone.

A converged single network creates perhaps the greatest differentiator between competitors in business. Build your business with an eye to the future. Incorporate technologies directly into your business when they fit with the core competencies. Outsource to qualified providers when appropriate, but manage the providers lest they manage you. Give employees the tools to do the best possible job in order to receive the best possible results.

The Evolution Continues

Humanity has evolved through several different eras. Nomadic tribes of hunter-gatherers followed the migration of animal herds or the availability and growing seasons of crops. Over time, this gave way to the

Agricultural Age that gave rise to cities and towns (or castles and villages), which were later connected by roads. The Industrial Age brought the automobile, tying cities and people closer together. It also focused on productivity in the manufacture of goods. The Industrial Age was an age of commoditization.

While all of these will forever remain important, the age of information is well upon us. Data, information, and knowledge comprise the intellectual capital that makes up the value of much business today. Manufacturing processes have become so efficient that goods quickly become commodities, but the services associated with those goods also bring value.

The telecommunications industry, as we've seen it grow and mature, is now gone. It will never be restored. It's perhaps important to put Darwin's theory of evolution in perspective. The theory never suggested survival of the "best." The theory of evolution really describes how the most adaptable species will survive and rule the day. To merely survive, we must be nimble and adaptable in embracing the Information Age and all its technological challenges. To thrive, we must embrace the changes and seize opportunities.

Just as Sun Tzu's *Art of War* has been used to prepare a set of competitive guidelines for businesses to follow, we can extend them to fit the evolution of IP telephony and other disruptive and emerging technologies.

Show the way by thinking out of the box and embracing new uses of technology. IP telephony provides a tool that integrates voice and data at a level never before available. It's bringing a new set of tools, so use the new tools to do new things.

Do it right by not embracing technology for the sake of technology. Buy the steak, not the sizzle. Perform business analysis with due diligence, but efficiency. Know the facts. Don't evaluate and plan forever. Know your business. Understand your needs, Evaluate solutions that take your business forward. IP networking is the clear path for business and communications.

IP telephony is a tool, nothing more. But it's a tool that provides a path for equipment vendors, telecommunications service providers, and businesses to innovate and develop integrated solutions that improve efficiency, profitability, and even quality of life for many. Now is the time to use the tools available to our best advantage.

It's critical that we not overlook the obvious. IP telephony has emerged as a solid, viable business tool. Today, it's a disruptive technology that is on the brink of reshaping the entire telecommunications industry. To overlook IP telephony would be a serious mistake, because

there will come a day when IP telephony becomes no longer the disruptive upstart, but the sustaining technology. We must continue to look to the future. The one constant trend of the future is change. Technology will continue to evolve: the key is to take advantage of IP telephony at the right time, and that time is now.