Open Blueprint

Telephony Resource Manager

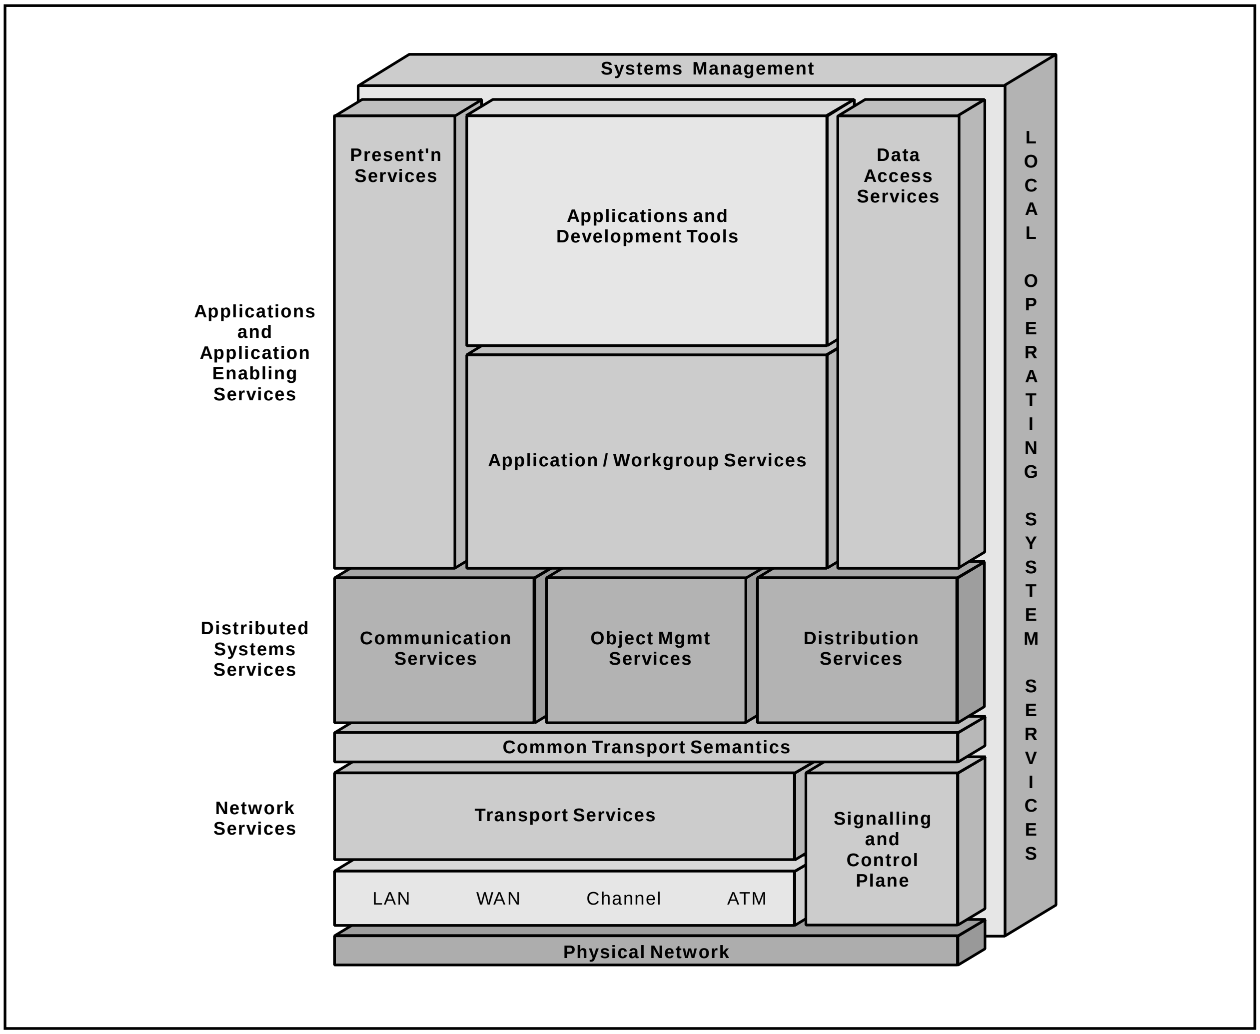

GC23-3928-01

Open Blueprint

Telephony Resource Manager

GC23-3928-01

About This Paper

Open, distributed computing of all forms, including client/server and network computing, is the model that is driving the rapid evolution of information technology today. The Open Blueprint structure is IBM's industry-leading architectural framework for distributed computing in a multivendor, heterogeneous environment. This paper describes the Telephony resource manager component of the Open Blueprint and its relationships with other Open Blueprint components.

The Open Blueprint structure continues to accommodate advances in technology and incorporate emerging standards and protocols as information technology needs and capabilities evolve. For example, the structure now incorporates digital library, object-oriented and mobile technologies, and support for internet-enabled applications. Thus, this document is a snapshot at a particular point in time. The Open Blueprint structure will continue to evolve as new technologies emerge.

This paper is one in a series of papers available in the *Open Blueprint Technical Reference Library* collection, SBOF-8702 (hardcopy) or SK2T-2478 (CD-ROM). The intent of this technical library is to provide detailed information about each Open Blueprint component. The authors of these papers are the developers and designers directly responsible for the components, so you might observe differences in style, scope, and format between this paper and others.

Readers who are less familiar with a particular component can refer to the referenced materials to gain basic background knowledge not included in the papers. For a general technical overview of the Open Blueprint, see the *Open Blueprint Technical Overview*, GC23-3808.

Who Should Read This Paper

This paper is intended for audiences requiring technical detail about the Telephony Resource Manager in the Open Blueprint. These include:

- Customers who are planning technology or architecture investments
- Software vendors who are developing products to interoperate with other products that support the Open Blueprint
- Consultants and service providers who offer integration services to customers

Contents

Figures

Summary of Changes

This revision describes:

- Support for making telephone calls over the Internet
- The relationship between the Telephony resource manager and the Collaboration resource manager

In addition, changes were made to terminology to better conform to industry usage.

Telephony Description

The Telephony resource manager provides computer telephony integration (CTI). CTI is the merger of telecommunications technology and computer technology. It provides the user with integrated voice and data services. With the rapidly changing world around us, the combination of verbal communication and computerized documentation is essential. Today, corporate call centers use computer telephony to control their telephones and to coordinate telephony functions with their data environment. On our workstations, we are continually experiencing telephony applications, such as those that allow us to click on a telephone number found in a document being displayed in a window and request that number be dialed. We continue to experience more voice processing applications, for example, to provide an audio menu for us to get information, such as an account balance, the weather, current portfolio balance, and voice-mail.

Voice is now part of the Internet. You can converse with others on the Internet using your workstation. Internet applications allow you to find and connect with any Internet user. Your voice is captured, put into packets, and sent over the Internet just as any other data item.

These are all examples of computer telephony integration.

Introduction

The three main functions of the Telephony resource manager are *call control*, *voice processing*, and *Internet voice*. The following discussions treat these functions separately. Call control and voice processing are discussed as they relate to the traditional telephone network. Then, the new dynamics of Internet voice and how all three functions provide one seamless set of integrated functions are described.

CTI Call Control

CTI call control encompasses basic telephone applications that control the traditional 12 button telephone. It also includes more sophisticated call center applications, which work closely with the telephone switch to provide multiple automated functions for the call center personnel. These functions include:

- Automatically presenting information about the caller on the screen
- Transferring or conferencing both the call and associated caller data by a simple action such as clicking on an icon or on a window field
- Routing calls automatically based on telephone-network-provided caller information matched against a computer database
- Controlling other types of calls such as fax calls and modem calls

Voice Processing

Voice processing works with the digitized voice stream. One can record it, store it, retrieve it, play it, use it for programmatic command and control, transform it into text, and provide the mechanisms for displacing voice in time such as is found in voice-mail applications. That is, an individual can leave a voice message that another individual listens to at a later time.

The primary aspects of voice processing include:

- Receiving an incoming call and dialoguing with the caller to retrieve or store information in the business's information system.

- Intelligent processing of voice data. Voice data can annotate other data or can be transformed into other types of data.
- Storing and forwarding of voice data.
- Placing outbound calls to relay information.
- Enabling voice control over a myriad of functions, such as sending alphanumeric paging messages.

Internet Voice

Internet voice allows you to speak to anybody on the Internet using the voice and sound features of your workstation. The voice is captured, compressed, and sent over the Internet. New enhancements to the Internet protocols allow the voice packets to be prioritized, expediently travel over the network, and arrive at the destination in a timely manner that preserves the fidelity of the voice. Internet voice applications provide user interfaces and access to Internet directory services to allow you to reach the party you wish to speak to.

Technical Background

Call Control

Call control is a fundamental function of the Telephony resource manager. It enables an application to influence and monitor the call setup process associated with creating a connection across a switched network. This function can be invoked to place an outgoing call, receive an incoming call, or influence a call in progress (for example, to place a call on hold).

The call control aspect of CTI is represented by two basic models. The first model consists of a view in which control is accomplished by direct attachment to the telephone system. This solution is commonly referred to as a *direct-attach* model. The call control function of the Telephony resource manager uses the Signalling and Control Plane resource manager to directly influence and monitor the setup, tear-down, and manipulation of switched service connections, when applied in the direct attach approach to call control.

Most commonly, in direct-attach telephony, the signaling to the switch and the voice itself are carried by the telephone wire that is attached to the telephone (see Figure 1 on page 5 for an illustration of the ways the computer is connected). The control signaling is provided by the telephone or a device in the workstation. The richness of the functions available depends on the functions provided by the private branch exchange (PBX) or local telephone company. The signaling is accomplished by a variety of methods, including flashhook and dual-tone multifrequency (DTMF) tones (those tones one hears when pressing the keys of a TouchTone telephone) in the simplest cases. Richer environments provide more sophisticated functions, such as mute, do not disturb, hold, conference, and transfer. Integrated services digital network (ISDN) and digital interfaces to the switch provide a wide variety of bidirectional control and signaling functions. Commands and requests are sent through the Signalling and Control Plane to the switch. The switch responds with information that indicates impending actions, the nature of incoming call media content, and network-passed information such as the caller's number, enabling a wide variety of software synergy with the telephony network.

The Signalling and Control Plane resource manager contains the various subnetwork connection protocols and specific control functions for different subnetwork-unique connection setup interfaces. For narrowband ISDN, it includes the Q.931 signaling protocol. For public switched telephone networks (PSTN), the Signalling and Control Plane includes the physical layer and the modem-related interface manipulation necessary to establish a switch connection.

The second model provides functions through a switch-attached server. Switches are either private branch exchange (PBX), which reside on the premise of the establishment, or central office (CO), which is part of the public network. This model makes use of specialized software written to allow the server to communicate with a telephone switch's proprietary CTI link. This model is referred to as a *client/server* model for controlling the telephone. In the client/server model, the telephony functions are provided to multiple clients. See Figure 1 for an illustration of this connection. The client/server model provides complete control over the features of a telephone without requiring manual intervention with the telephone set. The client/server model provides applications with information such as automatic number identification or caller ID, the state of the telephones connected to the switch, and event progress messages that indicate what actions are happening with calls (on hold, conferencing, transferring). The client/server model can enable an application to know as much about a call as the switch does, and to follow the progress of a call even after it has been transferred from a particular telephone.

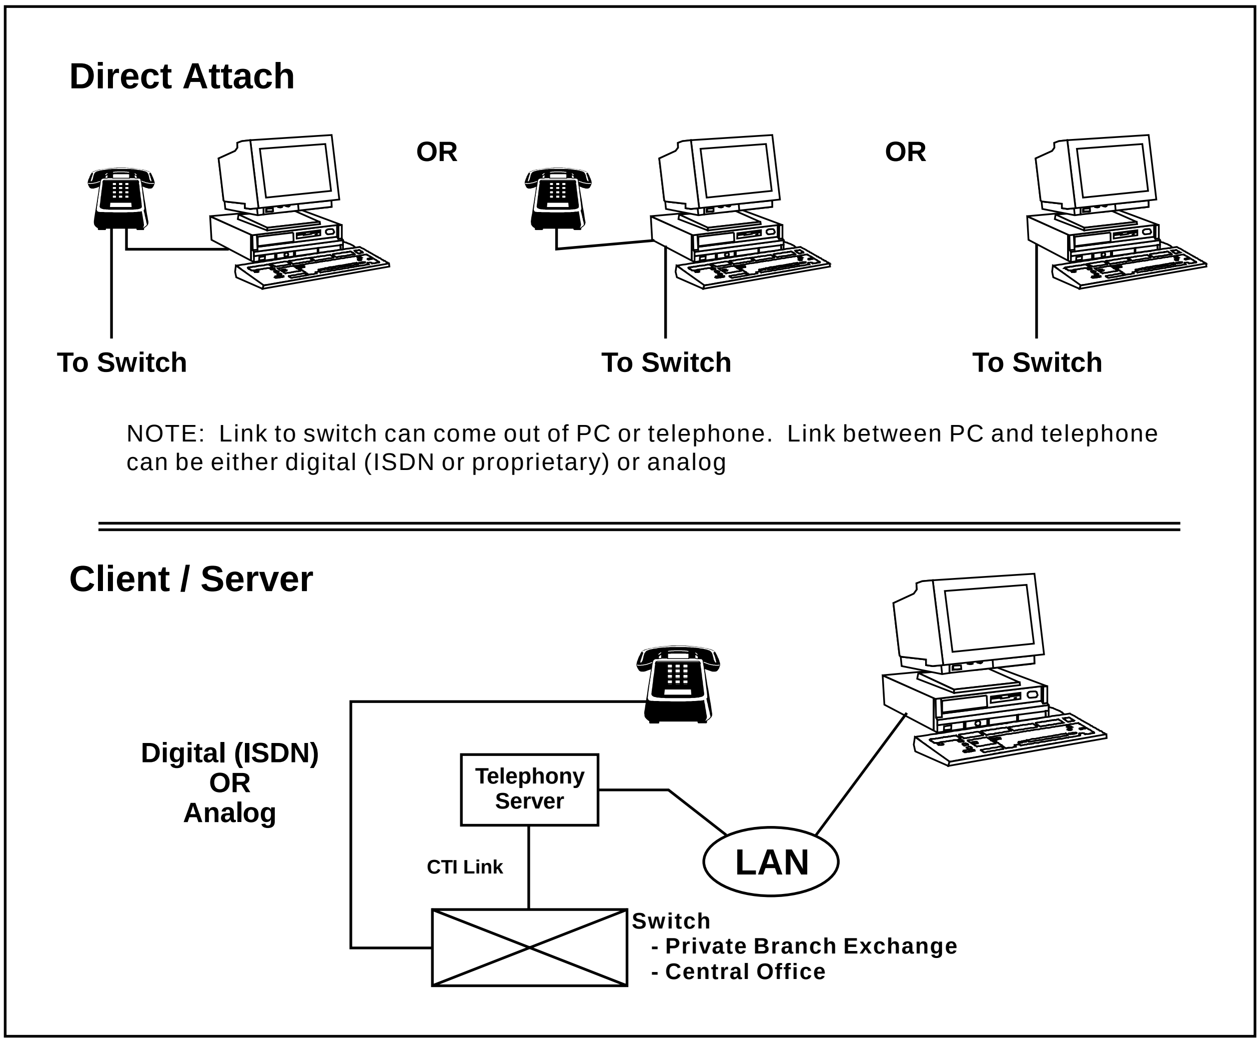

Figure 1. Direct Attach and Client/Server Telephony

Call Control Functions: Call control functions include:

- Making a call
- Answering a call (going off-hook)
- Conferencing a call
- Transferring a call

- Disconnecting or ending a call
- Placing a call on hold
- Retrieving a call from hold
- Forwarding a call
- Requesting that the server tell the switch how to route an incoming call
- Invoking telephone device features
  - Setting the do not disturb feature on or off
  - Setting the message-waiting lamp on or off
- Receiving messages from the device, switch, or Telephony resource manager.

  Messages that a program can receive include:

  - Response message to an earlier request that flowed to the switch from a program. Response messages originate in the switch.
  - Call-progress event messages that notify the program of the progress of an incoming or outgoing call. Call-progress event messages originate in the switch.
  - System message that informs the program of information that was detected by the Telephony resource manager. System messages originate in the Telephony resource manager and enable an application to perform local systems management.

  Some examples of message content are:

  - Automatic call distribution (ACD) information
  - Notification of:
    - An alerting device (ringing)
    - A conferenced call
    - A call placed on hold
    - A call torn down
  - Hardware status
  - Response to a request (positive or negative, including the reason)

Voice Processing

The voice processing component of the Telephony resource manager adds new dimensions to the use of the voice. It allows a simple telephone to be an input device to access information. It also allows voice to be treated as a data element and to associate the voice with other data in a logical grouping, such as a customer folder containing text, audio, and image data. Voice processing is based on general computing architecture rather than an integrated voice response (IVR) specific architecture. See Figure 2 for an illustration of the basic elements of voice processing.

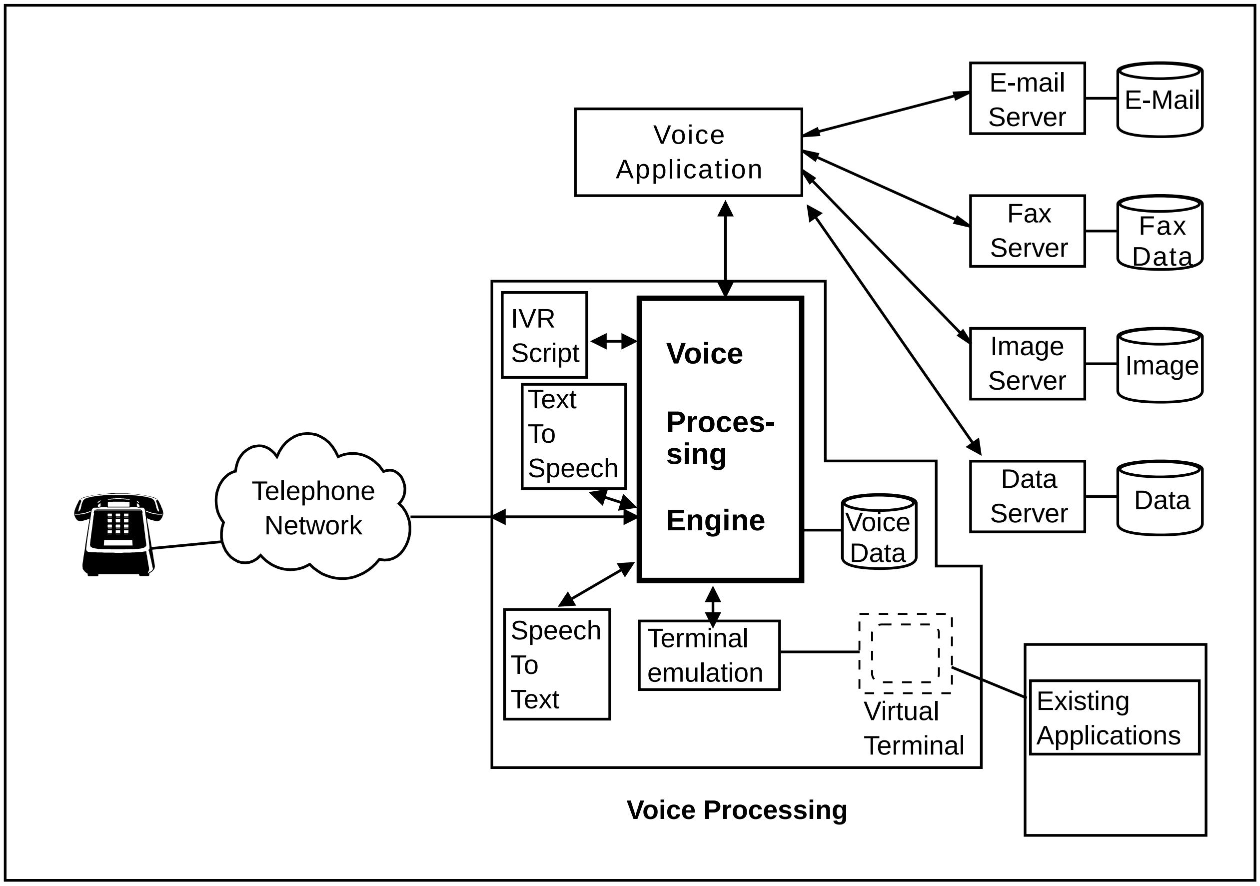

Figure 2. Basic Elements of Voice Processing

The voice processing component of the Telephony resource manager allows the deployment of voice applications that:

- Answer a call, and:
  - Play a prerecorded message
  - Begin an audio dialog with the caller based on identity or the number dialed
- Allow a user to request a service
  - Obtain information about checking balance
  - Request a fax be sent to a particular fax machine
- Provide voice-mail functions
- Place outbound calls to uniformly inform a large number of people
  - For example, informing employees on a particular work shift of a delay in reporting time
- Provide integrated mail—voice, e-mail, fax, and image using other Open Blueprint resource managers

- Maintain call statistics to handle voice processing calls more efficiently by reengineering the audio menu, forecasting future capacity needs, determining bottlenecks in processing, and determining the value of the voice processing system for the enterprise
- Enable intelligent management of voice as a data element, allowing voice annotation of text, image, and fax
- Transform speech to text and text to speech.

The technologies that underlie these capabilities are diverse and include:

- Telephony-related technologies such as analog and digital connectivity, which provide T1 and E1 digital connectivity for faster telephony throughput.
- Tone recognition, which enables simple or sophisticated applications to accept tone input from callers using push-button telephones.
- Voice-related technologies such as speech recognition and speech synthesis. Speech recognition enables all telephone users, not just those with push-button telephones that transmit tones, to use telephony functions. Speech recognition technology can handle continuous or delineated speech. It can be speaker dependent or independent. Implementations can be subject specific or general. The simplest cases of speech recognition are know as voice recognition.

  Speech synthesis transforms stored text to voice, which can be played to the caller. Thus, you can have your e-mail read to you over the telephone, for example.

The flexibility, expansion capabilities, and open design of the Telephony resource manager enable existing applications and resources to be used for simple applications and to evolve as the business environment changes.

There are *voice applications* associated with the Telephony resource manager. These applications can be customized to start a sequence of outbound calls, link to host data on any computer, add a feature such as a fax, or support signaling protocols not directly supported by the voice processor.

Internet Voice

Historically, if you wanted to speak to someone in another location, you would pick up the telephone and call them. When someone answered, you would begin speaking. Today, you can speak with someone by using your Internet connected workstation. This alternative brings many advantages, including the ability to more easily deploy integrated voice and data applications as well as a theoretical cost savings in having to deploy and maintain only one network. This emerging Internet technology shows much promise for the future. To better understand Internet voice, let us compare the traditional telephone and the telephone network with the new Internet technology.

In this simple scenario of making a call to someone, you first dial a number. This dialing information is sent to a series of switches in the telephone network. The switches set up a dedicated, direct connection between you and the person you are calling. The connection is made by "welding" all of the intervening wires between you and the person you are calling. Nobody else shares this dedicated path with you.

Figure 3 on page 9 below illustrates a traditional telephone call.

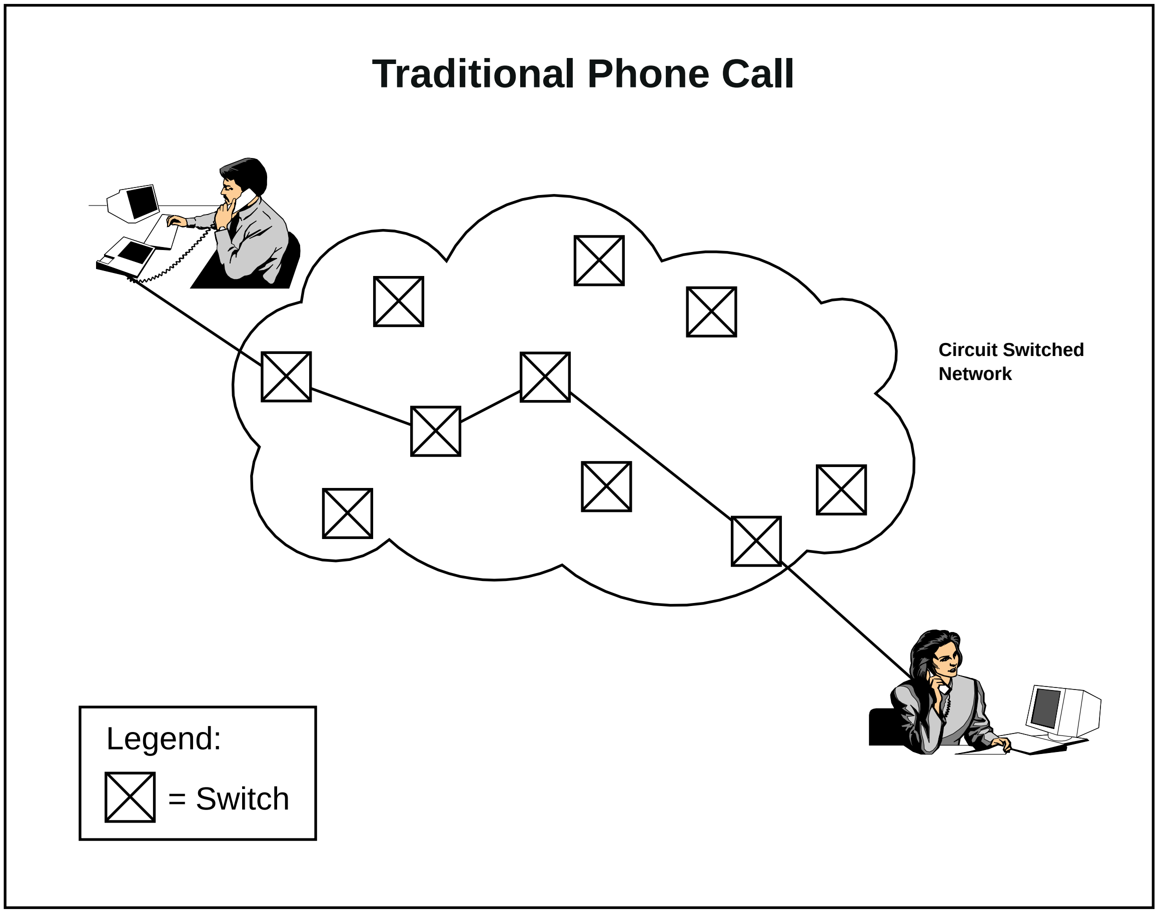

Figure 3. Traditional Telephony

Now let us look at the mechanics of an Internet voice call and relate it to a traditional telephone call. First, your workstation must contain a microphone, speaker and sound card, Internet connection software, and Internet voice software (see Figure 4 on page 10).

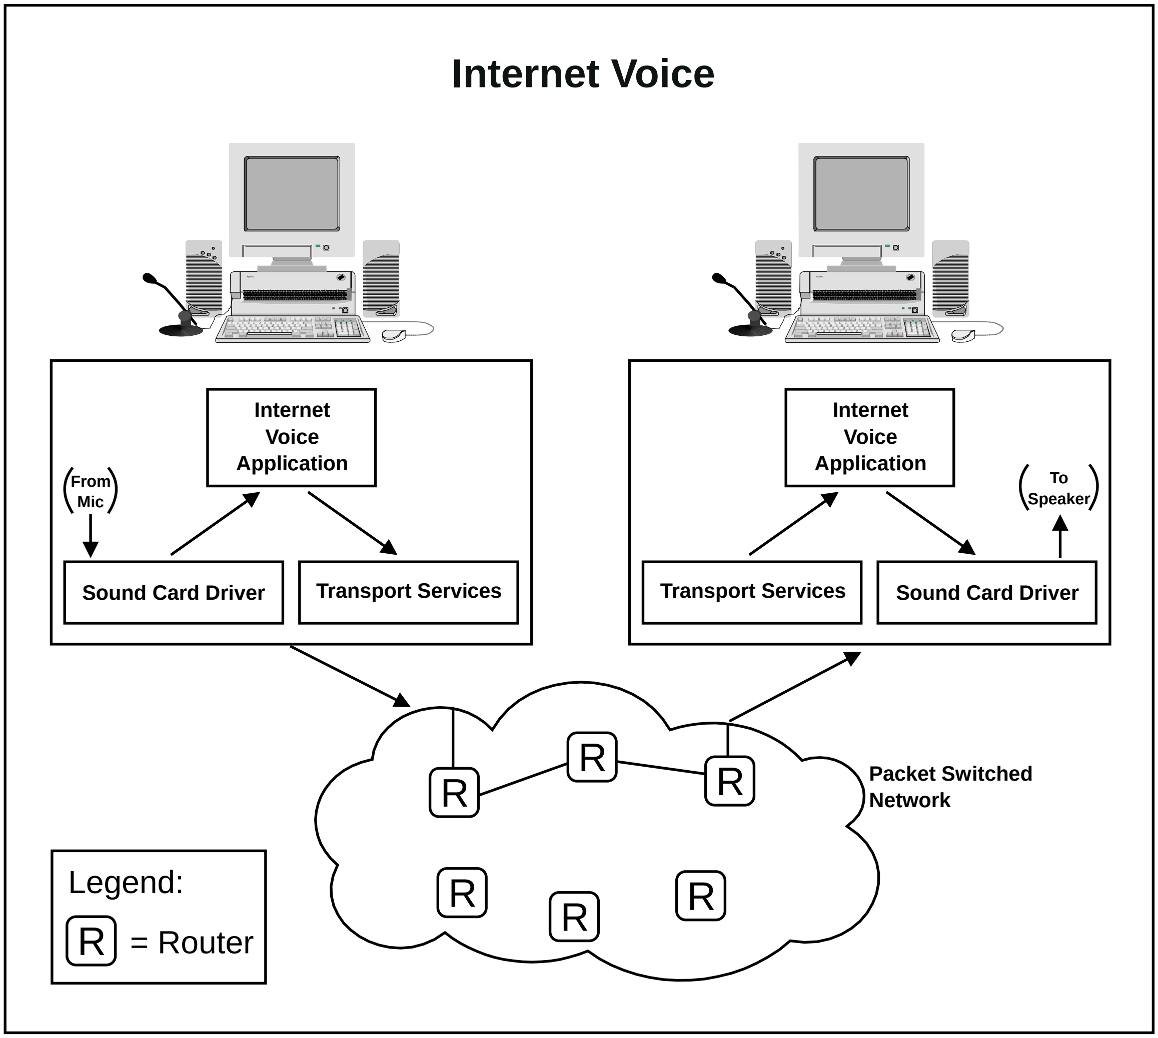

Figure 4. Internet Telephony

When you want to talk to someone on the Internet, you invoke the Internet voice software. Through the graphical user interface, you select the person you wish to speak to. The Internet voice software then sets up an Internet connection between you and the workstation of the person you are calling. This connection is a software connection. Once the connection is made, you begin speaking. The Internet voice software is given the voice packets that are collected by the microphone and digitized by the sound card. The internet voice software then compresses the voice, puts the voice data into buffers of appropriate size, provides address information, and passes the packets to the Open Blueprint Transport Services.

The Internet contains *routers* that look at the address information in the packet headers to forward the packet to the next router in the network, until it arrives at the destined workstation. Here, the voice stream is reassembled, decompressed, given to the sound card interface, and played out to the speaker. Thus, the person you are calling can hear you speak.

When that person wants to talk, the same process is initiated at that person's workstation.

Circuit-Switched Network versus Packet-Switched Network: This leads us into a discussion of different network technologies. The discussion of the traditional telephone call illustrates a circuit-switched network scenario. The session is dedicated for the duration of the call. The second discussion illustrates a packet-switched network scenario. There is no dedicated session; the session is virtual. Packets are sent to the destination by placing address information in the data packet headers. The headers are interrogated and the packets are routed in the network accordingly.

Advantages of a Packet-Switched Network: Having one network and one addressing scheme allows end-to-end connectivity for multiple data types and facilitates deploying applications that use data and voice simultaneously, as shown in Figure 5 on page 12. For example, several individuals can work concurrently on a particular document, making changes that appear simultaneously on each participant's workstation, and conversing about the concepts the document is meant to convey. This powerful workgroup scenario provides increased productivity for individuals working in separate locations and allows flexibility in work locations. For example, individuals can now effectively work from home using one communications line.

Wires are a precious and scarce commodity and therefore have high economic value. Coupled with the fact that networks are expensive to administer, there is obvious value of doing more in a single network. In a packet-switched network, wires are not dedicated to one session, but are shared among multiple sessions carrying voice, data, video, and so on between various sources and destinations. This technology holds much promise for being the dominant, cost-effective model for future networks, because it maximizes the utilization of the wires.

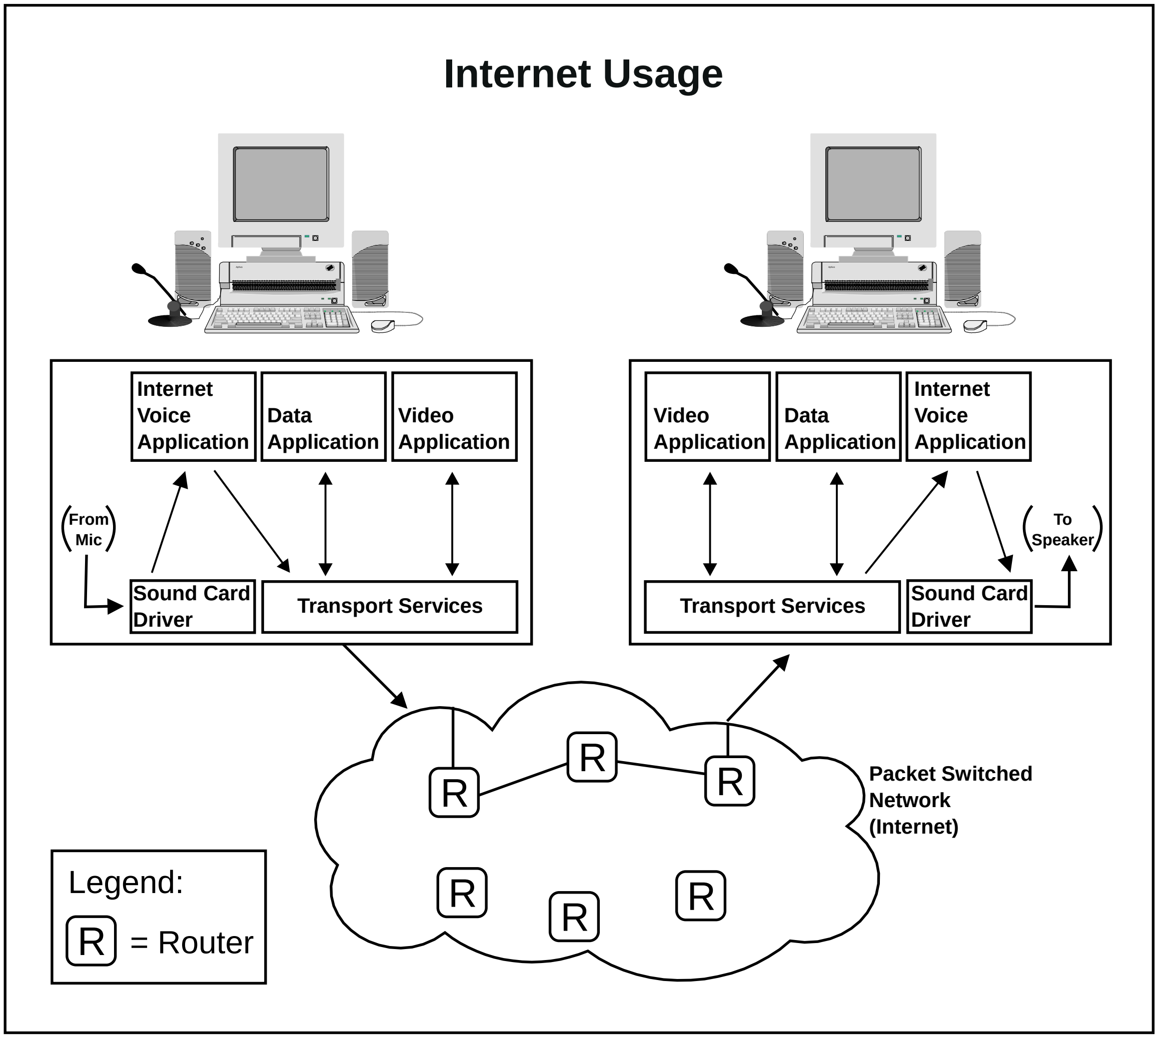

Figure 5. Internet Packet-Switched Network Usage

Benefits of Computer Telephony Integration

Everyone, from the home workstation user to the small business worker, to the call center agent, benefits from the integration of computer and telephone. Increased productivity, ease of use, and telephony functions that are integrated with other computer applications are the main benefits.

Using call control or voice processing or both, a business can change dramatically :

- Calls can be answered effectively by having the caller's information presented simultaneously with the call.
- Telephone calls can be divided evenly among groups.
- Calls can be answered around the clock, providing more business opportunities.
- Calls and data associated with the call can move from one place to another.
- Many calls can be answered simultaneously, reducing bottlenecks and delays during peak periods.
- Routine calls can be handled without human intervention.
- Callers can retrieve information they need or update their own database information anytime, anyplace, without human intervention.

The following scenarios help illustrate these benefits.

Scenario 1 - Call Control for an Office Worker

Jan is an insurance adjuster for the Duke Hazard company. Sam was involved in an automobile accident and is a policy holder with Duke Hazard. He needs to pass some information on to Jan, who has been assigned to the case. Sam calls Duke Hazard using an 800 number. Duke has deployed client/server telephony, so the PBX lets the server know that Sam is calling. The server searches a database and determines that Jan is handling the case. The server requests the switch to direct the call to Jan's telephone. The switch tells the server that Jan's phone is about to ring, and the server passes this information to Jan's desktop computer. Jan's computer does a database search and displays a summary of Sam's information to Jan. Concurrently, the telephone rings, Jan answers, verifies that Sam is calling, and begins immediately to interact with Sam.

During the conversation, Jan needs to confer with someone from the medical department. Jan does this with a click on a telephone book name and a click on the **Conference** icon.

In this scenario, Jan saved time and was able to give the customer almost instant service. Jan, Sam, and Duke Hazard benefitted.

The next day, Jan is reviewing Sam's case and needs some additional information. She highlights the number in Sam's folder displayed on the screen and clicks on the **Dial** function. Sam answers and Jan can quickly secure the needed information.

Scenario 2 - Voice Processing for a Home or Office User

Bob wants to get his current checking balance from the bank. He calls the bank's voice processing system and is lead through an audio menu to retrieve the balance. The balance is spoken to him.

Unable to reconcile his balance with the bank's, Bob requests a list of transactions between two dates be sent to his home fax machine. After reviewing it, Bob then asks for an image of one check in particular be sent to his fax machine so he can further verify this expenditure. Finally, Bob is satisfied that his own bookkeeping is at fault.

Scenario 3 - Voice Processing of Voice and Data Objects

John is an insurance claims adjuster. His system allows him to do his job in a paperless environment. All written correspondence and the images of accidents are kept on his system. In addition, voice recordings from depositions and interviews are kept in the system. An icon-based user interface allows him to go through a particular folder and work with the contents—text, image, and voice. He can record interviews over the telephone line, or he can use a tape recorder in the field and upload it to his system. In addition, he can annotate image or text with voice.

Here is an example of how he might use the system. John opens the Smith folder. Mr. Smith just had a minor car accident. In reviewing the material sent to his office, John notices a discrepancy between the image of the damaged car and the repair estimate. There appears to be only minor damage, but the repair estimate includes massive replacement of front-end parts. John records a message using a microphone attached to his multimedia workstation indicating this discrepancy, and puts a routing slip on the folder to the auto specialist.

The auto specialist opens the folder, notices the audio message icon blinking (a feature of his application), clicks on the message, and hears John's inquiry. The auto specialist investigates and finds the replacement is warranted for the particular type of automobile. The specialist then records an audio response and reroutes the folder to John.

Scenario 4 - Voice Processing and Call Control with Outbound Dialing

Jim is the manager of an assembly line. Due to a malfunction, the assembly line is down. He needs to contact his second shift workers and have them report to work an hour later than scheduled. He calls the voice processing system, records the message indicating that they should report one hour later, and requests the voice processing system contact all members on the second shift. The voice processing system begins dialing out to all members on the second shift, and uses answer detection technology to determine if a person answers the telephone (as opposed to an answering machine or no answer at all). If a person answers, it plays the manager's message and asks the employee to key in their employee number to confirm that they received the message. If the voice processing system cannot reach the employee on the first try, it periodically retries. Alternatively, it can send a message to the employee's pager. When the employee calls in, the message is played.

Scenario 5 - Coordinated Voice and Data Transfer

The Duke Hazard company handles insurance and securities. Sam calls to give Jan, who works for Duke Hazard, the details of a car accident. Sam then inquires about the current rates on CDs. Jan informs Sam that she is going to transfer him to that department. Jan clicks on the “Securities Department,” adds a note that Sam wants information on CD rates, and clicks on the **Transfer** function. The next available agent, George, is presented with Sam's profile and Jan's note. George answers the telephone and greets Sam by saying, “Hi Sam, I understand you would like the current CD rates.” George is now totally prepared to serve Sam.

Scenario 6 - An Internet Voice Scenario

Barbara wants to find out information about investments for retirement. The Duke Hazard Company has several annuity policies with various investment attributes. Barbara uses her favorite Internet browser to view the company's Web pages. At a particular point, Barbara decides that she wants to speak to an agent for additional information. The Web page allows Barbara to click on an icon for initiating an Internet call with an agent. Barbara is greeted by a message that indicates that all agents are busy and that the first available agent will handle her call. Instead of having to wait, Barbara can continue working on the Internet, going to additional Web sites.

When an agent becomes available, Barbara's personal information and the last several web pages that Barbara reviewed prior to selecting an Internet call are presented to the agent. The agent now has an idea of who is calling and what Barbara will probably inquire about. The agent can access more detailed information that is pertinent to what Barbara will probably ask about. When Barbara is finally connected to the agent, a very responsive and efficient conversation can take place. Using standard browsing techniques, the agent can control Barbara's browser to bring up the same detailed information the agent has on his screen, and the agent can talk Barbara through the information. Barbara can then print the displayed data for future reference.

Telephony Resource Manager

The Telephony resource manager provides an extensive suite of telephony functions. These functions are divided into call control, voice processing, and Internet voice for discussion purposes only. In reality, they are part of one seamless, interactive set of functions as depicted in Figure 6 on page 15. The Telephony resource manager provides end users with a full range of function anywhere—in the office with a local PBX attached, remotely connected through a wired or wireless telephone network, or over the Internet.

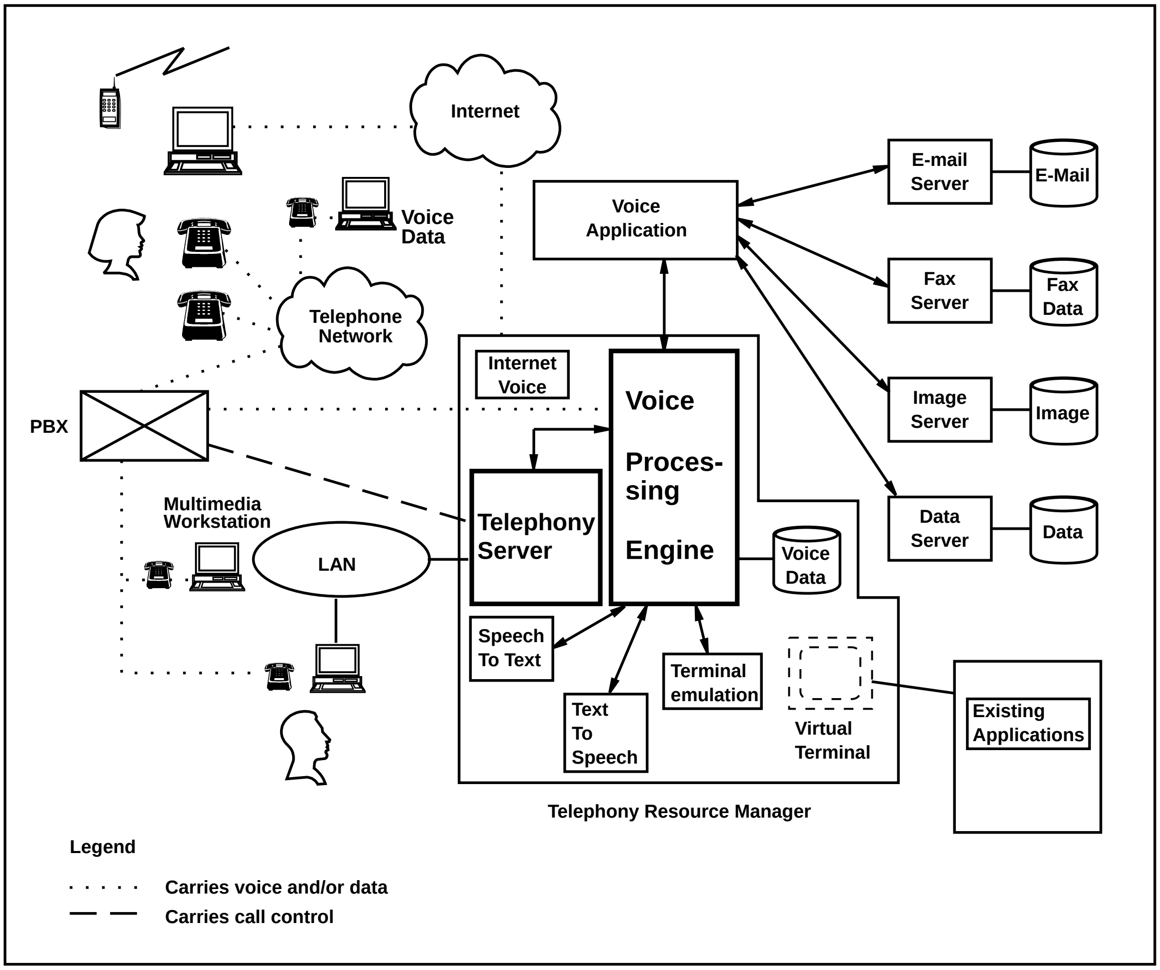

Figure 6. Call Control, Voice Processing and Internet Voice

Call Control

Again, let us look at the Direct-Attach and the client/server models.

Direct-Attach Model Support: In the direct-attach model, an adapter device in the workstation accomplishes the call control through inband signaling to the switch. Two APIs are provided in this model; one for the application programmer and one for the device driver programmer. The Telephony resource manager ties the flow of the two APIs together, as shown in Figure 7 on page 16.

The adapter device is controlled by a workstation program through the Signalling and Control Plane. This enables the workstation user to establish telephone calls with a wide variety of characteristics and quality of service. Once the call is made, the application interfaces with the Multimedia resource manager to flow the multimedia data. The data can be voice, fax, text, or an image. In short, it can be any type of data that can be transported over the specific telephone network consistent with the quality of service provided by the network.

Similarly, on inbound calls, the application is provided with information about the type of multimedia data being sent—voice, fax, text or image. (For more information about multimedia data, see the *Multimedia Resource Manager* component description paper.) The application interacts with the Telephony resource manager to establish and manipulate the inbound connection and interacts with the Multimedia resource manager to process the incoming multimedia data.

Within the Open Blueprint structure, the Telephony resource manager resides in Application/Workgroup services and the service provider code resides in the Signalling and Control Plane resource manager.

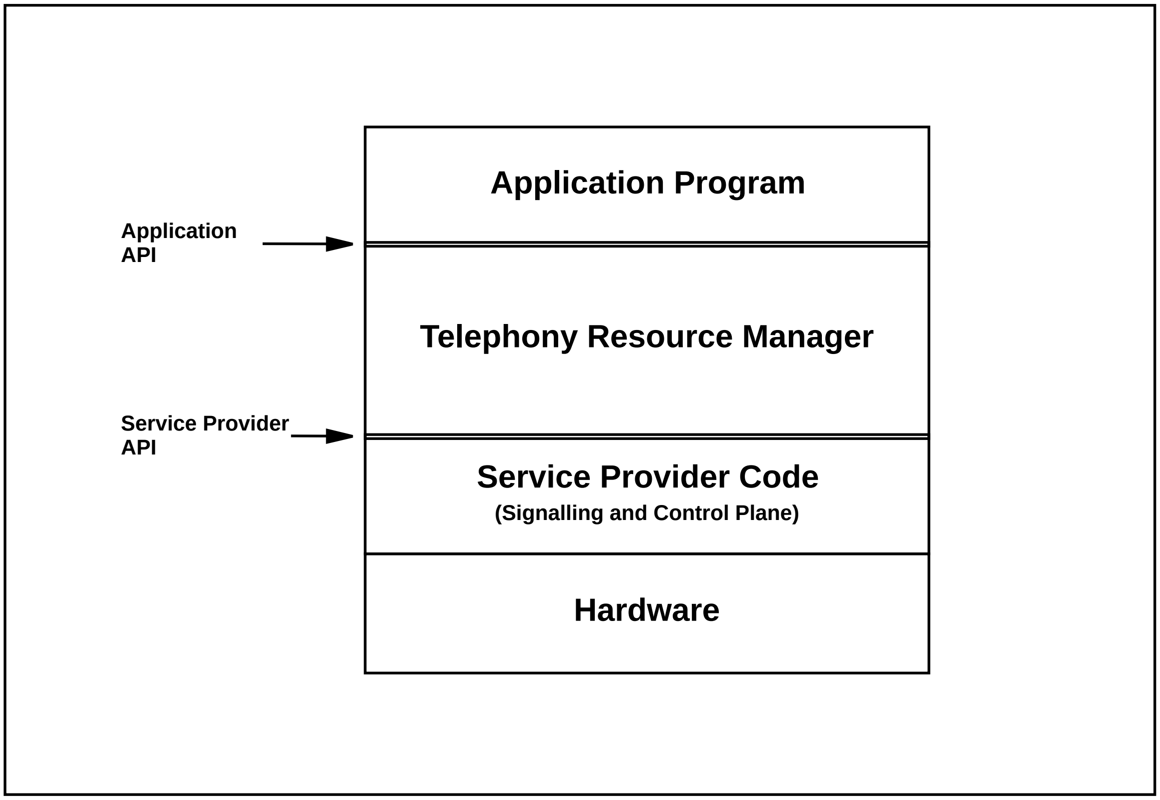

Figure 7. Direct-Attach Structure

Client/Server Model Support: The server portion of the Telephony resource manager is connected to the switch and contains most of the controlling code for passing information between an application and the switch. The server transforms requests from the client into requests that the switch understands. It takes information from the switch and transforms it into a well-defined, consistent, architected format that the application understands.

Communication between the client and the server portions of the Telephony resource manager is through remote procedure call (RPC). This mechanism allows the APIs to be deployed easily on any computing platform supporting RPC. Most of the function is performed on the server, including coordinating applications with requests and responses.

Using the client portion of the Telephony resource manager, the requesting application has access to telephony functions as well as other functions provided by Open Blueprint resource managers.

The Telephony resource manager provides consistency among switch vendors. In the client/server model, the control link to every switch— (PBX or central office) has its own method of communication. Most of these communication methods are proprietary. The server has switch-device modules for each switch. IBM has developed the concept of a virtual switch model which is a superset of all switch functions. The process of transforming individual switch flows and behaviors into the virtual switch model is called *normalization*. Applications written to this virtual switch model run unchanged and unaware of which switch the server is attached to.

These applications can be used with any switch. This protects the application development investment. Because of the rich flow available from current switches, this can equate to a major savings in time and resources and it enables telephony applications to be written and deployed more quickly, which is consistent with the Open Blueprint's goal of promoting integration of multi-vendor systems.

Voice Processing Model Support

The voice processing model takes distributed processing to a new dimension. It uses the services of other resource managers to intelligently manage and distribute audio, text, fax, image, and other data types. The voice processing system gathers various types of data from multiple resource managers, transforms that data into the appropriate format, and presents it to the user remotely attached through the telephony system. The device that interfaces with the user is the simple telephone or a telephone with a text capability, and it can be either wire-attached or wireless. More sophisticated end user devices can also be supported, such as mobile personal information managers (PIMs).

The voice processing system may also request information from another voice processing system made by another manufacturer using standard mechanisms for the interchange of audio information. A classic example is an application that consolidates the corporate voice-mail with a personal workstation voice-mail system. Or, it might combine the voice-mail of two corporations, as in the case of a contractor who receives voice-mail from both an employer and the contracting company.

The Telephony resource manager can be configured to include peer servers, such as text-to-speech and speech-to-text, and fax servers.

Associated with the resource manager are applications, such as the terminal emulation application. The design also supports user-written applications, thus opening up possibilities to access a wide variety of servers using a variety of communications protocols provided by the Communications Services resource managers. Further, fax and image modules can be incorporated to support the various industry standards in these areas. Digitized data can be stored and retrieved by anyone, anytime, anyplace, using standard interfaces.

Internet Voice

The Telephony resource manager includes the capability to transfer real-time voice over the Internet and private TCP/IP networks (intranets). Using a workstation with a sound card, speaker, and microphone, a user can make a voice call to anyone on the Internet with voice quality as good as a cellular telephone call. This enables users to deploy applications that simultaneously share data and transmit voice over a single connection using regular modems. A variety of applications now become available, such as home banking (talking to a teller while viewing your account status or check images), remote shopping (acquiring more information from an assistant while viewing a catalog), and customer help desk (a support technician can give instructions while accessing the customer's computer).

Administrative Aspects of the Telephony Resource Manager

Administrative tasks are accomplished through management applications. Administrative tasks include:

- Associating the telephone number with the client workstation to form the telephony desktop. Although this association can be established statically, it is often established dynamically by the client application when the application is started.
- Reflecting telephone group setup in the switch.
- Passing administrative messages between the client and the server and between the switch and the server. These system messages reflect the overall well-being of the system and any disruptions, their nature and severity. The applications receive these messages and take appropriate actions.

- Manipulating and configuring connections using the appropriate addressing scheme, establishing and deleting connections, activating and deactivating connections, and providing status information and trace information on these connections.
- Enabling and disabling telephone lines.
- Performing telephone line diagnostics.
- Checking the status of terminal emulation sessions.
- Enabling and disabling server traces and application program trace logging.

Alarms and alerts are sent to systems management facilities for reporting and corrective action. Support for running unattended is provided by a management application “standing in” for an operator.

Scalability

Scalability is a basic design characteristic that allows the deployment of the same solution for a wide range of processing needs. Additional hardware resources may be added to a single configuration, or processing may be distributed over multiple instances of the Telephony resource manager.

The Telephony resource manager is designed to manage a wide variety of resources, and provide several types of scalability. Not only can one specific system be scaled, but several cooperating systems, working in tandem, can be deployed, across multiple geographic locations. In addition, the Telephony resource manager is highly portable among various operating systems on various hardware.

The Telephony resource manager provides scalability on all supported computer operating systems. It uses the facilities of each of the operating systems to allow this scalability. On each platform, there is a wide range of possibilities that can accommodate just a few users, thousands of users in a centralized configuration such as a mainframe, or a distributed environment of workstations. The Telephony resource manager client and server configuration is very flexible. It requires adding the appropriate resources such as telephone lines, adapter cards, host emulation lines, or memory to fit the needs of a particular processing domain, then notifying the Telephony resource manager through systems administration facilities.

Relationship to Other Resource Managers

The Telephony resource manager is part of the Application/Workgroup Services portion of the Open Blueprint structure. The Telephony resource manager adds telephony functions to the set of functions available to applications and other resource managers. It uses the Presentation Services to provide the graphical user interface (GUI) on the desktop computer, enabling dynamic call control from screen icons or mouse clicks, and it allows the user to easily build voice processing applications using 4GL technology. Further, it enhances the Presentation Services by providing an audio interaction paradigm between applications and the caller. Prerecorded voice messages are played to the user to provide an interactive audio menu. The caller interacts with the application by use of the telephone keypad or voice responses. The Telephony resource manager then translates this input into what can be understood by the application.

The Telephony resource manager uses the Open Blueprint Communications Services to communicate with client telephony applications, server applications external to this resource manager, and the switch. It uses the Open Blueprint Network Services to communicate to telephone lines, LAN-based servers, remote servers, and remote switches. The Telephony resource manager uses the Remote Procedure Call resource manager for client/server telephony.

See “Telephone Access to Collaboration Resource Manager Data” on page 22 and “Integrated Messaging” on page 22 for a discussion of the Collaboration resource manager functions that allow integration between the Collaboration resource manager and the Telephony resource manager.

Interoperability

Interoperability means that products can work together, either because they have implemented the same set of standards, followed the new software engineering techniques and have successfully encapsulated functions through iterative design and test, or adhere to a design requiring very loose coupling.

Interoperability is provided in the following areas:

- Among different switch manufacturers
- Among a variety of network connectivity capabilities such as analog and digital lines, ISDN, and T1
- Between different data types: voice, text, fax
- Between different devices that can interoperate, such as a TouchTone telephone, a dial telephone, and a text telephone

The server portion of the resource manager provides telephony functions to a wide variety of client platforms including Apple, Hewlett-Packard, Sun, Microsoft Windows, and UNIX. All the platforms can coexist in the same server domain and interoperate to provide coordinated data and voice functions. In voice processing, digitized data can be stored and retrieved by anyone, anytime, anyplace, using standard interfaces.

Different technologies can be used to capture, store, and replay the voice data, as shown in Figure 6 on page 15. A multimedia-equipped desktop can be used to capture the voice. This digitized captured voice can then be sent to the Telephony resource manager application over the LAN for cataloging and storing the voice. Subsequently, a telephone can be used locally or remotely to retrieve that voice message using the Telephony resource manager voice-mail application, which retrieves and plays the cataloged message. Likewise, the voice can be stored by a Telephony resource manager voice-mail application. Later, a multimedia workstation application might request the digitized message from the Telephony resource manager over the LAN and play the message over the workstation speaker using functions of the Multimedia resource manager.

Performance Considerations

The CTI environment is a real-time environment. The voice must be played and captured isochronously. No data can be lost. As such, timing between the call control application and the switch must be kept under values that are real-time compliant to ensure that the CTI functions properly. Management applications do real-time statistical processing for capturing and displaying performance information. Examples of monitored items include the state of the links, message rates, number of messages sent or received on various links, mean transit time of messages transmitted, and the calling pattern of inbound versus outbound. This information can be used for timely client workload balancing and planning for future resource needs.

Other aspects require near real-time response. Response from the voice processing system to a user's request must be timely. There cannot be a long delay between the time the user requests information using the telephone, and when the information is delivered.

In essence, each segment of the entire configured system involved in this call must be highly tuned. This includes the call control component, the execution of the voice processing application, the communications to the enterprise database, the database retrieval, and the text-to-speech components.

The following scenarios illustrate the real-time nature of the CTI environment.

The Voice Processing Component: Speech recording and playback must be done at a very constant rate (isochronous operation). Failure to do so causes the voice to become unintelligible. Imagine hearing what is supposed to be "the red fox jumped over the fence" as "t-- re- -ox jum--- over --e fence." Sufficient design in the hardware and software has enabled a high quality of fidelity in the capture and playback process.

This is also the case with the Internet voice. Delays in the TCP/IP stacks on the desktop and delays in the Internet can cause the voice packets to arrive too late to preserve the voice quality. By using priority queueing mechanisms found in current Internet standards, and by using sufficient bandwidth, this problem can be eliminated.

Call Redirection: Call redirection deals with directing an incoming call to the proper telephone. The switch sends a message to the server indicating that an incoming call has arrived and passes all the information it has received from the telephone network about the incoming call. The server determines where to route the call based on some predefined information that a user or administrator, using a Telephony resource manager client, has given the server. The server has a predefined amount of time to respond to the switch before the switch takes default actions. Because the server usually has to do a database search, the amount of processing time left for non-database functions might be minimal.

Timely Screen Pop: Timely screen pop is another example where timing is critical. Here, the client code must do a database search and fill in the screen with data associated with the incoming call. It is critical that the called party can answer the telephone and begin an appropriate dialog with the customer. This ensures customer satisfaction and ensures payback for the investment in deploying the system. The more time delay in answering a call, the more time must be spent per call. This has a negative effect on productivity. The same is true for conferencing and transferring a call.

Migration and Coexistence Considerations

Migrating from one level of telephony support to another can be a nightmare if migration considerations are not included in the design. The design of the Telephony resource manager allows new levels of servers to be introduced into the establishment without disruption of existing applications, and it allows multiple levels of servers to coexist.

Another aspect of migration is in deploying a new version of a server, then deciding the establishment must revert to the previous version in the middle of prime production time. This unexpected change can be handled smoothly with the application automatically reverting to functions provided by the old level server. The primary design mechanism used in the Telephony resource manager API is versioning on both requests and events. By properly programming the application to use these versioning capabilities, migration and coexistence of various level of applications and servers is accomplished.

Associated Support

To provide a complete set of capabilities to build, deliver, and manage the complex set of function described above, other capabilities must be available. These capabilities include tools for the developer and for environment managers.

Coordinated Voice and Data Applications: Applications exist that make it very easy to integrate telephony support with an establishment's large databases. These offerings have functions that allow an application to retrieve information from inbound calls and use that information as input to transaction processing systems with auto-key input. (Auto-key involves putting information into the transaction processing system's screen programmatically to simulate the end user who is typing the information in and pressing the Enter key.) When a call comes in, the application end user sees the screen that presents information about the calling customer. The user is then able to continue inputting to the transaction system on the window where coordination offerings left off. In addition, when a call is transferred within the business establishment, these offerings ensure that the caller information goes to the new destination of the call, presenting the target of the transfer with the information about the caller. The new application end user can continue interacting with the caller without losing continuity. The caller does not have to repeat any of the information previously imparted.

Application Development: The call control environment provides a basic software development toolkit that consists of the header files and source code libraries needed to compile and link the application programs with the call control environment. Trace facilities allow real-time debugging of the environment. These traces exist on the server and client for problem identification and isolation.

The voice processing environment provides a flexible scripting capability that allows customization of the audio menu provided to the caller. It has many functions, including recognizing DTMF input from the caller's keypad, program flow macros (if, then, else; while loops), and error condition handling. In addition, it provides the capability to customize the voice prompts, and it provides a visual voice editor.

Finally, the voice processor provides custom back-end functions for use in the script builder, performs syntax and reference checks, supports multiple voice recording and playback rates, and lets the user test and debug applications interactively. Included in the built-in system functions is a terminal-emulation package to retrieve information from mainframe systems.

These functions enable the development of applications that can result in integrated computer telephony solutions.

Telephony Solutions: The Telephony resource manager provides the basis for voice solutions in the business environment. These solutions enable the user to:

- Provide the caller with better service in a business environment by providing a richer functional environment for the person answering the call
- Gain voice access to e-mail
- Use an e-mail system to view the text conversion of voice-mail
- Use the telephone as the instrument for interacting with a company's database
- Conduct business with someone who is using Internet voice

Call Control Applications: The call center segment of the marketplace is deploying increasing numbers of call control applications. These applications incorporate the telephone functions into existing business applications to increase the productivity of the call center.

In this paper, the term *agent* refers to the individual or group of individuals assigned to handle customer calls for a call center or business.

The functions provided by call-control applications include:

- Intelligent answering

  The intelligent answering functions present the agent with information about the caller and provides the first window for data entry. This allows the agent to begin working with the caller immediately, saving time for both the agent and the caller.

- Load balancing

  The load balancing function spreads the work evenly between several call centers; the customer is serviced by the next available agent in any of the call centers.

- Skills-based inbound answering and outbound dialing

  This function selects the next available agent who has the skill to talk to a customer. For example, the agent needs to be able to speak Spanish to converse with the customer.

- Estimated wait time

  This function calculates the amount of time that a new inbound caller will have to wait on a queue before someone will be available. This function can be used to send alerts to supervisors when certain thresholds are exceeded.

- Statistics collection and report generation

  These functions allow the call center to understand the effectiveness of its process and to make the proper adjustments.

Telephone Access to Collaboration Resource Manager Data: The Collaboration resource manager provides an environment that allows any TouchTone telephone to access its documents (including e-mail). This enables any Collaboration resource manager user who does not have convenient access to a PC to connect to the Collaboration resource manager database from a telephone. Information captured over the telephone can be stored in a Collaboration resource manager document, which then uses the same integral Collaboration resource manager services as any other document. That is, voice messages can be replicated across a distributed workgroup and can become essential components of workflow applications. See **1** - **5**. in Figure 8 on page 23. As a Collaboration resource manager client, the telephone can be used to retrieve, create, modify, and delete documents in any application. Documents can be read over the telephone using the text-to-speech synthesis of voice processing. The user can respond to documents or create new documents that capture the user's voice instead of capturing information from the user as text. Evolving speech recognition technologies allow the user's voice to be converted into text directly.

Integrated Messaging: Integrated messaging is a technique that allows the Collaboration resource manager and voice processing system to access each other's data. This enables a user to access voice-mail messages using the Collaboration resource manager e-mail interface and to access the Collaboration resource manager. e-mail using voice processing. See **7** - **8** and **7a** - **8a** in Figure 8 on page 23.

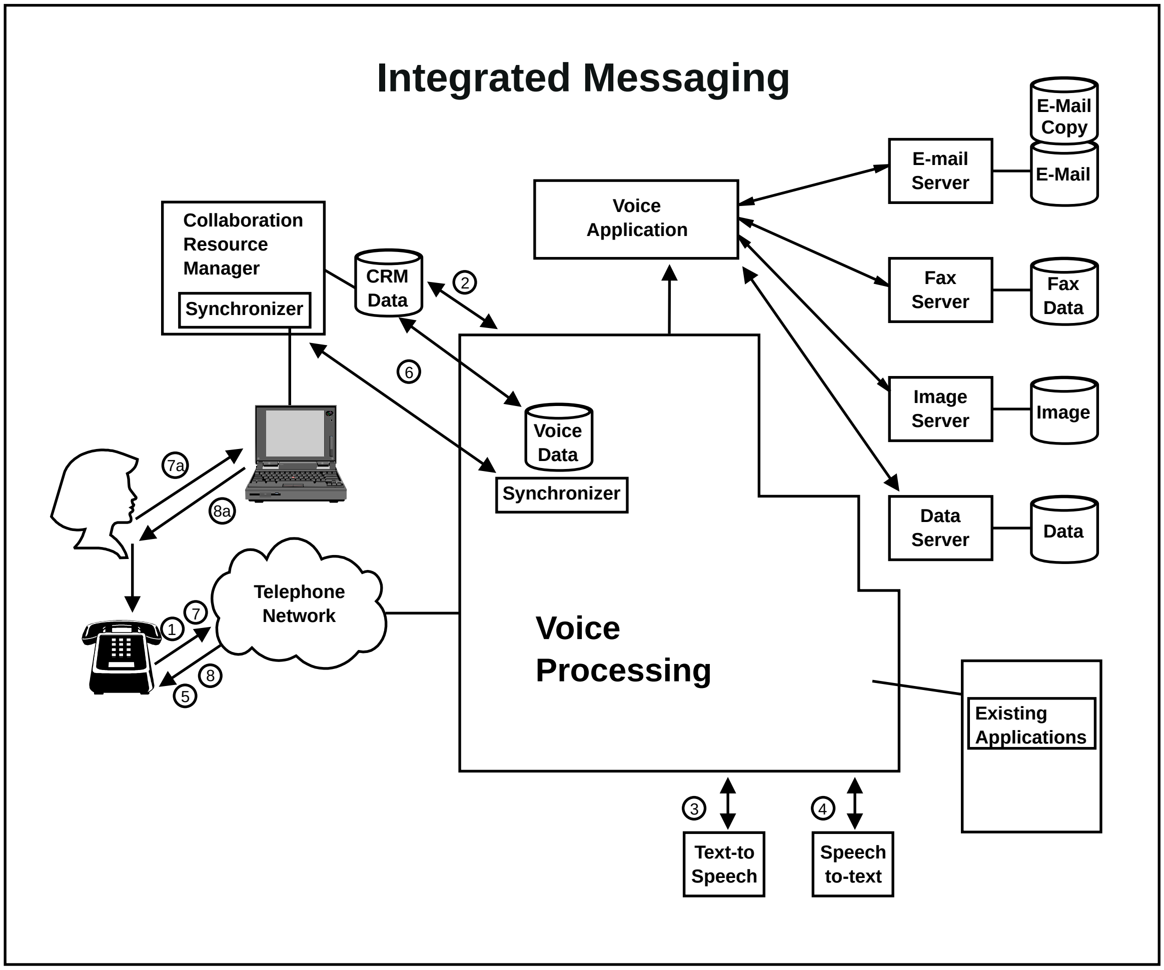

Figure 8. Telephone Access to the Collaboration Resource Manager and Integrated Messaging

Integrated messaging is targeted to people who use both electronic and voice-mail systems in a corporate environment. The user has at least one mailbox each for the Collaboration resource manager and voice processing. User access to the two types of messages is seamless. Voice messages can be browsed, categorized, and achieved in the same manner as electronic messages. The voice processing user is able to access Collaboration resource manager e-mail or voice-mail through the telephone or a computer-based voice-mail message management system. Additionally, users in either system can address and send messages to users in other e-mail and voice-mail systems.

Integrated messaging is implemented using a synchronizer that replicates voice messages in the e-mail system and e-mail messages in the voice-mail system (**6**). The synchronizer interfaces to both mail systems using the Common Messaging Call (CMC) API endorsed by the Electronic Messaging Association. In replicating messages between the two system, advanced speech processing technologies can be used to convert e-mail into voice (text-to-speech) and to convert voice into e-mail (speech-to-text). Other types of messages, such as fax, can be replicated across the two systems, allowing one to convert a fax into a voice message.

Voice Applications: Voice applications using the Telephony resource manager can answer inbound calls and interact with the callers, or make outbound calls and interact with the called party. Voice applications can also perform computations, retrieve and update information in databases, communicate information by playing it back over the telephone, let callers leave messages, transfer calls to another application, or connect callers to a human operator. Typically, voice applications are concerned with:

- Interacting with the caller or called party

- Accessing information
- Managing voice messages

The voice application, which integrates voice processing with various server functions, can provide a variety of compelling solutions as shown in Figure 9.

Examples of servers that can be integrated with voice processing are:

- Database servers
- External speech recognition servers
- Text-to-speech servers
- Fax servers
- E-mail servers
- Speaker-identification products that recognize an individual's voice

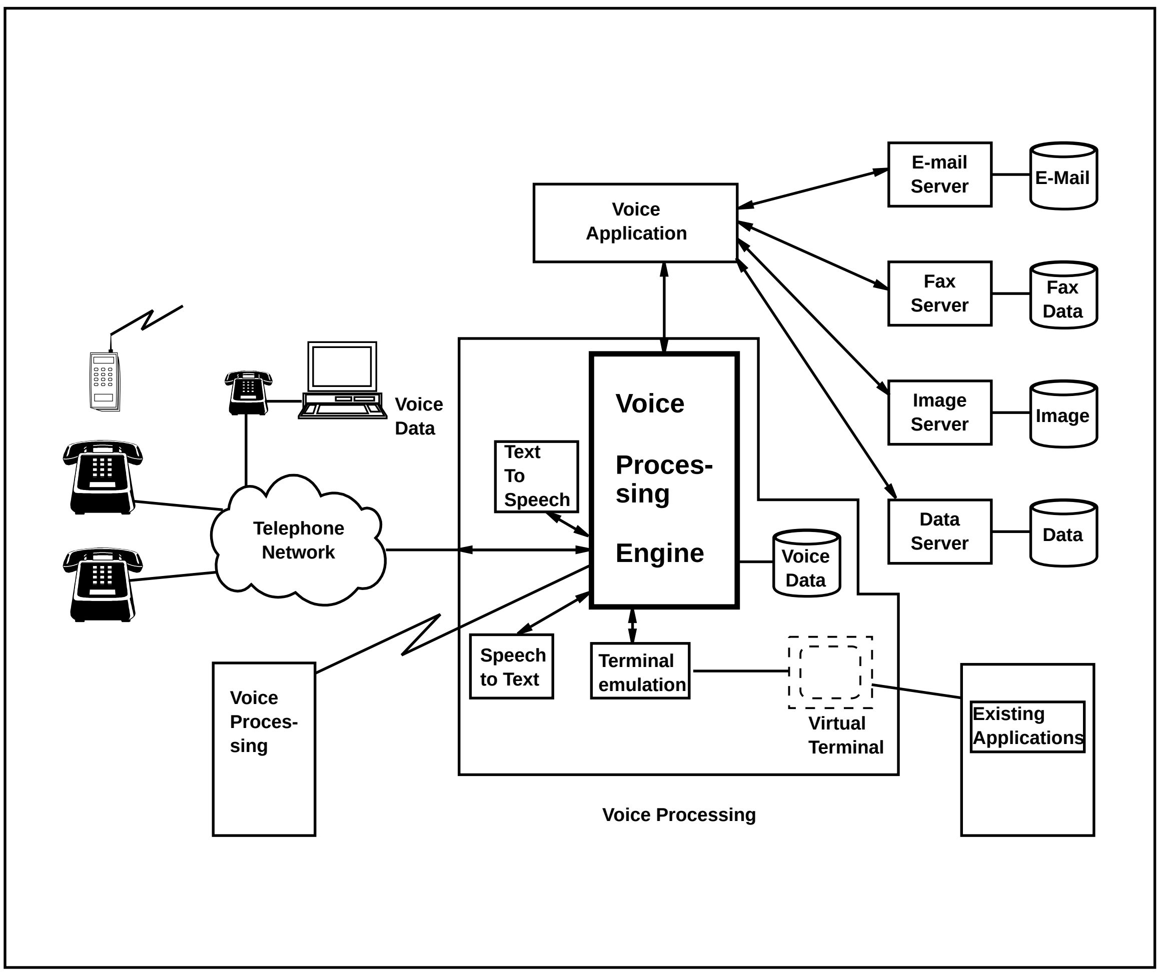

Figure 9. Voice-Application Solutions

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
USA

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries:

IBM
IBMLink
Open Blueprint

The following terms are trademarks of other companies:

Apple (name and logo)	Apple Computer, Incorporated
Hewlett-Packard	Hewlett-Packard Company
Sun	Sun Microsystems

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Communicating Your Comments to IBM

If you especially like or dislike anything about this paper, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply. Feel free to comment on specific error or omissions, accuracy, organization, subject matter, or completeness of this paper.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

- If you prefer to send comments by FAX, use this number:

  United States and Canada: 1-800-227-5088.

- If you prefer to send comments electronically, use one of these ID's:
  - Internet: **USIB2HPD@VNET.IBM.COM**
  - IBM Mail Exchange: **USIB2HPD at IBMMAIL**
  - IBMLink: **CIBMORCF at RALVM13**

Make sure to include the following in your note:

- Title of this paper
- Page number or topic to which your comment applies

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC23-3928-01