

**IBM WebSphere Voice Response**  
with DirectTalk technology

An IBM White Paper

*This white paper introduces the reader to the IBM WebSphere Voice Response offerings. It describes the features of WebSphere Voice Response for AIX and WebSphere Voice Response for Windows. It was written for anyone interested in learning about WebSphere Voice Response and how it can be used.*

## Introduction

WebSphere® Voice Response is IBM's integrated voice response system which is designed to meet the ever-expanding requirement for business information to be available over the telephone, and can support those solutions where a voice processing application is used to supplement a web application.

By bringing together existing telephone networks, business logic servers and data communications networks, voice processing provides customers with the information they need via the telephone.

## Voice Processing

Voice processing provides a caller with the ability to use a telephone connected to a voice processing system to access information. The information accessed can be contained entirely within the voice processing system or sources linked to the voice processing system, or a combination of these environments can be used. The voice processing system provides the link to bring together existing telephone network and data communications network to provide access to a range of computer based information using the telephone.

The voice processing system can be a PC or an IBM p-series server based voice processing system. Callers are connected via the normal telephone networks to the voice processing system which can be designed to offer a range of functions and features based on specific business requirements.

By using the touch tone keys on a standard telephone for communication, the telephone becomes an input and output device for the business application. The transmitted DTMF (dual tone multi frequency) tones are recognized by the voice processing system and treated as commands to perform certain functions defined within the application running on the system. Information is normally delivered to the caller using prerecorded voice. The information spoken to the caller is determined by the voice response application.

The voice processing system may be designed to accept spoken words as input. In this case, the caller speaks a response to a question from the application and the voice data is processed using the speech recognition technology to determine what the caller said. Depending on the capabilities of the technology, the input from the caller may be an isolated word, phrase or a complete sentence. Implementing a speech recognition solution requires specialized expertise, as many factors can affect recognition accuracy including the size of the required vocabulary, the speaker's accent, hesitation markers and background noise.

## **WebSphere Voice Response services**

WebSphere Voice Response can support single systems or large, networked environments supporting thousands of telephone lines, such as telecom providers or call centres. Up to 16 T1/E1 trunks (384 (T1) or 480 (E1) channels) can be supported on a single WebSphere Voice Response system (IBM AIX®) and runs on two major operating platforms: IBM AIX and Microsoft® Windows NT® and Windows® 2000.

WebSphere Voice Response offers many benefits to business, this is provided by the following services which could be developed for WebSphere Voice Response. Alternatively these could be acquired from a business partner.

### **Automated attendant**

An automated attendant can direct incoming calls to different customer service representatives or departments, or to different automated applications. Callers can make choices by using the telephone keypad or by speaking.

### **Telephone access to multiple systems and applications**

WebSphere Voice Response gives the caller access to a complete range of automated applications and information in a single call. Even if the information is stored on several different types of host computer, perhaps as a result of merging one company with another, the caller can access it with the convenience of a single call.

### **Voice response**

Voice response, the use of prerecorded speech to provide information in response to input from a telephone caller. Brief pieces of information are retrieved from a database and spoken to the caller using audio recordings of words and phrases strung together by the voice application. WebSphere Voice Response applications can also speak lengthier information, either prerecorded or stored as text and converted to speech on the fly.

### **Fax response**

For more detailed or even graphical information, WebSphere Voice Response's fax support can be used to allow callers to receive information by fax.

### **Transaction-related voice messaging**

Transaction-related voice messaging means that callers can leave voice messages related to transactions that they have conducted using WebSphere Voice Response.

### **Coordinated voice and data transfer**

The automated attendant or other voice processing application can let the caller talk directly to a service agent. While transferring the call, any data that WebSphere Voice Response has already gathered can be displayed on the agent's computer display using a computer-telephony integration (CTI) product. The agent is immediately prepared to handle the inquiry.

### **Access to paging systems**

If the person the caller wants to speak to is on the move, WebSphere Voice Response can transfer the call to a paging system.

### **Automated outbound calling**

WebSphere Voice Response doesn't just help to handle inbound calls; it can also make outbound calls. Automation of outbound application allows, with ease, the ability to improve customer service and introduce potentially revenue-generating applications. Voice applications can be designed to answer inbound calls and interact with callers or make outbound calls and interact with the called party. These voice applications can also perform computations, read and retrieve information from databases, update databases, communicate information by speaking it back over the telephone, let callers leave a message, invoke other applications, or transfer callers to a human operator.

## ***What WebSphere Voice Response applications do***

The heart of a WebSphere Voice Response application is its interaction with a calling or a called party. Typically, a voice application either handles inbound calls or makes outbound calls.

### **Inbound calls**

WebSphere Voice Response automatically answers each inbound call, and uses information about the call to pass it to the appropriate voice application. WebSphere Voice Response can run many different applications simultaneously. The application then plays a

recorded greeting to the caller, and then prompts the caller to indicate what information they want. The caller can interrupt the prompt if they already know what they want to do. The application can be designed to play greetings in different languages depending on information such as the number that the caller dialed, or the caller's own number. The application waits for the caller's response for a set period of time. The caller can either respond by speaking (if speech recognition is implemented) or by pressing keys on a DTMF phone, ranging from a single digit to multiple digits. If the response does not match the criteria that has been defined, the voice application can prompt the caller to enter the response again. The application takes whatever action is appropriate to the caller's response: for example, retrieving information from a database, host system, or file system and speaking it to the caller; updating information in a database; storing or retrieving a voice message; or making another call. After taking action, the application should prompt the caller again to tell it what to do next. The caller may indicate, in response to a prompt, that the interaction is over. The application can respond by saying good-bye, and then disconnecting, or the caller may just hang up. If the caller hangs up, the application can detect this, and automatically disconnect itself.

#### **Outbound calls**

A voice application could be triggered to make an outbound call. The application would make the telephone call, to a telephone number which is stored in a database. If WebSphere Voice Response makes the call, it can determine whether the call is answered or if the line is busy. WebSphere Voice Response can play a recorded greeting. It can then prompt the called party to tell it what to do next. After the initial greeting, an outbound call follows much the same sequence as an inbound call. The called party should have the opportunity to ask for other actions to be taken. Again, just as with an inbound call, either the called party or the WebSphere Voice Response application may end the interaction. The called party may either hang up, or tell the application not to take any further action, in which case the application must disconnect.

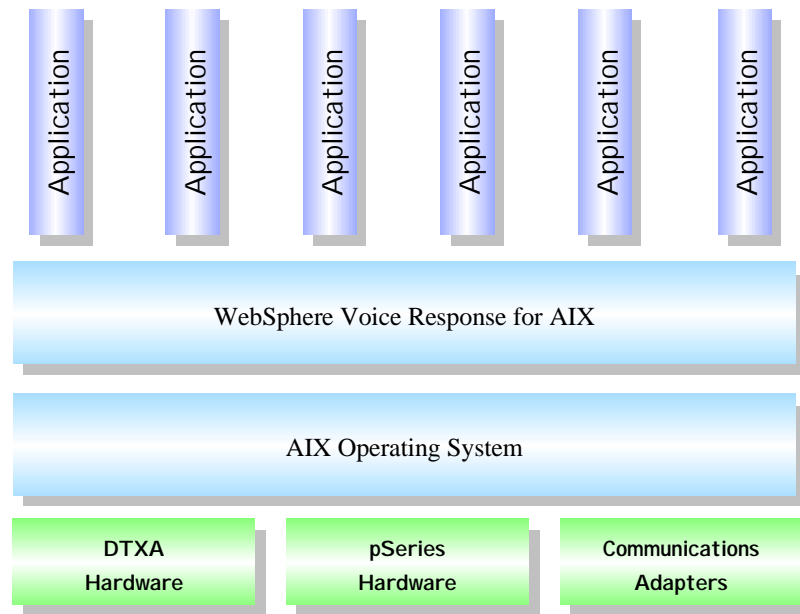
#### **Transferring calls**

An application can transfer either an inbound or outbound call to another extension.

#### **Information access**

WebSphere Voice Response can take advantage of any computer application that already retrieves, stores, or manipulates information. Because WebSphere Voice Response communicates with existing programs, information in an existing database can be accessed without writing new host programs for that purpose. The database would normally be on another host or another type of computer, although it can be on the same host as WebSphere Voice Response. If the database is accessed via a 3270 or 5250 based application or transaction server, WebSphere Voice Response can use its built-in 3270 or 5250 terminal emulation (5250 emulation available on WebSphere Voice Response for Windows only) support to exchange information. For example, a voice application can start a CICS ® transaction that accesses a DB2 ® database, without requiring alterations to the CICS transaction.

## WebSphere Voice Response for AIX



The WebSphere Voice Response for AIX voice processing system comprises, one or more pSeries processors, the AIX operating system, the IBM WebSphere Voice Response for AIX licensed program product, supporting signaling protocols to connect to a PABX or central office switch and specialized hardware for the pSeries to communicate with the telephone network(DTXA).

### IBM Hardware - DTXA

In order to connect the IBM p-series server family to the telephony network some physical hardware is required. To provide this interface IBM developed the Digital Trunk eXtended Adapter (DTXA)

Up to four DTXAs can be installed in a PCI bus IBM p-series server. By using the maximum of four DTXA cards in the system unit, up to 16 T1 or E1 trunks can be configured on a single pSeries.

WebSphere Voice Response for AIX supports H.100 connectivity on the DTXA card. This allows other PCI bus cards with an H.100 or SCBus connection to be connected to the DTXA card and therefore used within the telephony system, i.e. fax cards etc.

### Fax

WebSphere Voice Response for AIX Version 2.3 supports the Brooktrout TR114-P16S 16 port fax card, allowing up to three fax cards to be installed in selected pSeries machines. This enables voice applications written for WebSphere Voice Response for AIX to support inbound and outbound, simultaneous, two-way, channel fax communication. Up to 48 fax channels are supported when three fax cards are installed with a minimum of one DTXA card.

### Connection to the telephone network

There are various ways in which WebSphere Voice Response for AIX can be connected to the telephone network, such as through the local telephone exchange or central office (using T1 or E1 standard business lines or trunks), or by using a PABX system. WebSphere Voice Response for AIX can use either channel associated signaling (CAS) or common channel signaling (CCS) protocols.

The choice of protocol is dependent on which subset of protocols is supported by the switch or PABX and which provide the functionality required by the voice application. A single WebSphere Voice Response for AIX system supports either T1 or E1 trunks, but not both. All channels, however, do not have to share the same T1 or E1 protocol. CCS protocols can coexist with CAS protocols on the same system, some trunks can use CCS whilst others use CAS. They cannot be mixed within a single trunk.

WebSphere Voice Response for AIX supports two types of common channel signaling protocol, Integrated Services Digital Network (ISDN) and Signaling System #7 (SS7). WebSphere Voice Response for AIX also supports the development of specialized common channel signaling protocols for intelligent network (IN) applications, through the use of the signaling interface.

## **ISDN**

The optional ISDN features supports primary rate interface (PRI) vendor-specific implementations, PRI switching and signaling capabilities and calling number identification services for PRI. There is an optional feature of WebSphere Voice Response for AIX to interface to the network over an ISDN primary rate interface that conforms to each of the following standards:

- Euro-ISDN Lucent 5ESS 5E8
- Lucent 5ESS 5E9 National
- Lucent TR41449/TR41459
- Nortel DMS100 BCS34/36
- Nortel DMS100 National 2 — NA007/008
- Nortel DMS250 — IEC05
- National ISDN2/5ESS 5E12/13
- INS Net Service 1500

## **SS7**

There is a Signaling System #7 feature (SS7) for WebSphere Voice Response for AIX. SS7 allows WebSphere Voice Response for AIX to be connected directly to telephone company networks, providing tighter integration of WebSphere Voice Response for AIX applications to the network. This gives customers the benefits of common channel signaling, such as very fast call setup and tear down, reliable network signaling, and the delivery of called and calling number information. The WebSphere SS7 call control is written to the ETSI standard ISUP, most country variants are supported by the underlying SS7 stack but this needs to be verified with WebSphere Voice Response for AIX call control. Current customer investment in applications, is protected, by allowing existing applications to run unchanged over SS7. A maximum of 1440 concurrent voice channels with SS7 are supported, for E1 on multiple PCI-based systems. The coexistence of SS7 trunks, ISDN trunks, and channel associated signaling trunks are supported in the same system for ease of migration. The SS7 processes can be monitored remotely using an SNMP based network manager.

## **VoIP**

Voice Over Internet Protocol allows the sending of telephony voice over IP (Internet Protocol) data connections rather than dedicated voice networks, switching and transmission equipment. WebSphere Voice Response for AIX can be connected to VoIP networks using one or more external VoIP/PSTN Gateways. There is a Beta program under way in which WebSphere Voice Response for AIX can be directly connected to a VoIP network. For further information on using VoIP with WebSphere Voice Response for AIX, please contact an IBM representative.

### **Voice storage**

Stored voice includes voice segments used in prompts and voice messages. Voice segments can be recorded compressed or uncompressed. Uncompressed voice segments require more disk space and memory, also more I/O operations are required. Voice messages are always compressed. Voice segments may be uncompressed or compressed.

### **Custom servers**

Using a custom server, access can be provided to information that resides on any type of computer. Even if the information is held in a database on a host that uses the 3270 data stream (zSeries or iSeries) it may be more convenient to write a custom server to access it rather than to create a 3270 server. A custom server consists of a main function that calls user functions, the custom server is written as a C or C++ program. First, user functions are developed to provide the interface to information on remote systems. Then, either develop, or have WebSphere Voice Response for AIX generate, a custom server main function to pass required information between the user functions and the calling state table. Custom servers can be of two types:

- Applications that wait to be called by a state table. This type is the most common server, typically used for processing functions requested from inbound calls.
- Applications that are initiated by other means, such as another program or a timed event. This type is useful for processes like outbound calling, or invoking a state table which might then call other custom servers and state tables.

The open architecture of WebSphere Voice Response for AIX can be exploited further by using custom servers to access:

- External speech recognition servers
- Text-to-speech servers
- Fax servers
- Telecommunications devices for the deaf
- CTI Server, to provide computer-telephony integration
- Call Tromboning using the H.100 TDM (Time Division Multiplexing) bus

This openness enables the best solutions to be chosen which are suited to an organization's requirements. IBM Business Partners and other solution providers make use of the flexibility of the custom server interface to provide complete applications for speech recognition, fax integration, speech synthesis, switch integration, and internet access. In addition, some custom servers are provided with WebSphere Voice Response for AIX, or as optional features.

### **Single System Image (SSI)**

WebSphere Voice Response for AIX can be implemented as one or more standalone systems or clustered in a network architecture called Single System Image (SSI). In the SSI

implementation, a number of systems can share the application data (such as state tables and custom servers) and voice data (such as voice segments and voice messages). Each system is configured as either a client or a server.

#### **Client Nodes**

A client node interfaces with the caller, it handles the caller interactions. It runs WebSphere Voice Response for AIX (configured as a client) and must be connected to the telephone network. The client node does not contain any application data; that data is provided by the server which is connected to the Local Area Network (LAN). Only log files, statistics and configuration data are stored locally; all other WebSphere Voice Response for AIX data is stored on the server node. WebSphere Voice Response for AIX must be installed locally on each node and started.

#### **Database Server Node**

The database server node contains the application objects stored in a DB2 database. The database contains all the application profiles and mailboxes, state tables, prompts, 3270 servers and custom servers used by all the WebSphere Voice Response for AIX systems in the single system image. The database server node also holds all the voice (segments and messages) used by the client nodes, unless a voice server node is configured. WebSphere Voice Response for AIX must be installed on the server node, configured as a server and started. If the server needs to handle interactions with the caller directly, then the server needs telephony hardware installed on the system, otherwise it does not.

#### **Voice Server Node**

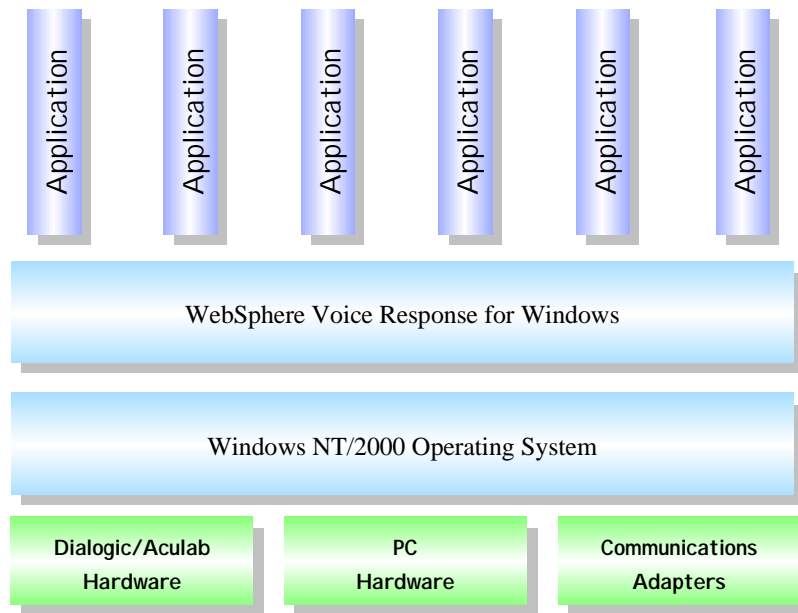
The voice server node is optional, but if configured contains the voice data for all the voice applications that run on the SSI image. The voice server node does not need to have WebSphere Voice Response for AIX installed unless there is a requirement to handle interactions with callers. If this is a requirement then the voice server node will need telephony hardware installed.

The database server and the voice server will usually be the same system and will generally only be installed on two separate systems if a large voice messaging application is being implemented thus allowing the processing load to be shared.

The nodes of a single system image must be connected together using a local area network. The type of network used will depend upon the size of the voice applications being run. A small cluster running a simple voice application may require only a token ring network. However, a larger cluster that is used to support many voice applications or a voice messaging server, may require a network that can provide a higher capacity and performance such as an Asynchronous Transfer Mode (ATM) network.



## WebSphere Voice Response for Windows



For a WebSphere Voice Response for Windows system the following components are required, one or more Intel® PCs, the Windows NT or Windows 2000 operating system, the IBM WebSphere Voice Response for Windows NT and Windows 2000 licensed program product, an Aculab network interface (NIF) card and one or more Dialogic® voice processing cards or, alternatively one or more Dialogic combined network interface and voice processing cards.

### Dialogic & Aculab hardware

The cards supported by WebSphere Voice Response for Windows are described in Chapter 6 of the General Information and Planning publication for WebSphere Voice Response for Windows NT and Windows 2000, Appendix A provides details of the cards supported in different countries and regions. Alternatively, for the latest status for country support the Aculab ([http://www.aculab.com/products\\_main/approcals\\_main.htm](http://www.aculab.com/products_main/approcals_main.htm)) or Dialogic (<http://resource.intel.com/globalapproval/globalapproval.asp>) websites can be referred to. WebSphere Voice Response supports up to 4 T1 or E1 trunks (96 (T1) or 120 (E1) channels) or 48 analog telephone lines on a single personal computer.

Additionally, WebSphere Voice Response for Windows supports a number of Dialogic VFX cards, which provide fax support. The VFX range of cards supported by WebSphere Voice Response for Windows can be used in two ways, as combined fax, analog network, and voice cards or as a fax resource shared between network and voice cards installed in the personal computer. As a shared resource, fax can be made available on digital trunks as well as analog lines.

The telephone lines, which are connected to the WebSphere Voice Response system, can come directly from the public telephone network or through a private branch exchange (PBX). If call volumes require (or grow to require) more than 4 T1 or E1 trunks (96 (T1) or 120 (E1) channels) or 48 analog telephone lines, additional WebSphere Voice Response

systems can be created and then connected together through a LAN. All systems connected together through a LAN can be managed from a single node.

### **Telephone Network connectivity**

WebSphere Voice Response for windows connections are either analog or a T1 or E1 digital trunk connection. The particular model of personal computer and the number of network adapters installed affects the number of trunks and channels which can be used. For digital connection, WebSphere Voice Response can use either channel associated signaling (CAS) or common channel signaling (CCS) protocols. Common Channel Signaling is supported in WebSphere Voice Response for Windows by ISDN and other non-ISDN protocols such as AMSTEL, DASS2, DPNSS, FETEX.

### **WebSphere Voice Response for Windows System Configurations**

WebSphere Voice Response for windows can be installed as a standalone system in a network configuration with multiple systems being connected together by a LAN. Systems connected together by a LAN in this way are known as nodes. There are three types of system, in this network topology, which have different capabilities.

#### **Runtime Node**

To provide the connection to the telephony network and the voice processing functions required by all voice applications, at least one runtime system is required. A runtime system can also be used to manage the system. In the simplest case, the entire service can be run from one runtime system. However, as the service grows, more runtime systems can be added, for extra production capacity or for testing applications before putting them into production.

#### **Non-runtime Node**

To manage the systems, a non-runtime system can be installed.

#### **Remote application Node**

To run Java or VoiceXML applications remotely, a remote application system can be installed. This can help to spread the application load over multiple systems and can also be useful for testing applications before putting them into production.

If multiple nodes are required, they must be connected using a LAN with the capacity and performance required by the applications. The LAN must have the capacity to handle the voice data transmitted over it and the voice response systems must have the capacity to handle the collection and playing of this voice data.

## Developing applications for WebSphere Voice Response

To exploit the power and functionality of WebSphere Voice Response to provide phone access to business data and logic, dialogs need to be developed that will control the interaction between the caller and the data. There are 3 application programming environments supported by WebSphere Voice Response, which can enable voice applications to access resources as well as presenting voice output.

### Java

WebSphere Voice Response support for Java voice applications is in the form of JavaBeans. These are reusable software components written in 100% pure Java, with which it is possible to quickly and easily build the interactive voice interface for a voice application. JavaBeans such as Menu, EntryField, and Form encapsulate much of the programming needed to implement typical voice applications. A visual application builder can be used, such as WebSphere Studio Application Developer that supports JavaBeans, or Java can be coded by using a text editor.

Instead of writing state tables, prompts, and 3270 servers or T-REXX using the proprietary WebSphere Voice Response programming languages, and custom servers using C or C++, complete voice applications can be created using the Java programming language.

Using the WebSphere Voice Response Java bean support for developing voice applications enables the developer to:

- Use industry standard graphical development tools such as WebSphere Studio Application Developer.
- Integrate voice applications with multi-tier business applications more easily
- Create voice applications capable of interacting with other computer and telephony systems such as:
  - (3270) IBM Host on Demand
  - JDBC database access
  - IBM Secureways (LDAP directory access)
- Develop voice applications that can run on more than one platform
- Simplify and speed up voice application development by using off-the-shelf beans
- Develop voice applications independently of a WebSphere Voice Response production system by using the Java Simulator.

An application written using the Java support requires the services of WebSphere Voice Response for AIX or WebSphere Voice Response for Windows. The Java Beans Environment provides the connection between the Java Beans and the base WebSphere Voice Response system. The base WebSphere Voice Response system can be running traditional state tables or T-REXX scripts at the same time.

A voice application written using these Java Beans can run on any operating system that runs Java, and can use any WebSphere Voice Response platform to provide telephony support: not only WebSphere Voice Response for AIX, but also WebSphere Voice Response for Windows. Using Java's Remote Method Invocation (RMI) interface, applications can run either locally, on the main WebSphere Voice Response system, or remotely, on one or more application client systems.

## **VoiceXML**

A VoiceXML browser application acts in a similar way to a web browser in that it is simply a rendering device, but uses audio rather than visual means to present and receive information from a user. A VoiceXML browser application, like a web browser, communicates with web servers using standard internet protocols and downloads pages, and the resources referenced by them (audio clips rather than images), on demand. Rather than rendering pages defined in HTML, a VoiceXML browser application deals with pages defined in VoiceXML.

The major advantages of using VoiceXML are:

- Only one application model is required for both web servers and voice services. Applications for both channels are created and deployed in the same manner on web servers.
- System management of the DirectTalk system is reduced. Using VoiceXML, there are no resources such as applications, voice segments, or grammars to manage on the WebSphere Voice Response system. All resources are centralized on the web server.

## **Native**

WebSphere Voice Response for AIX and WebSphere Voice Response for Windows both have a proprietary application development environments. The WebSphere Voice Response for AIX uses state tables and WebSphere Voice Response for Windows uses T-REXX. Each environment is separate to the other, for instance, a state table written for WebSphere Voice Response for AIX will not run on a WebSphere Voice Response for Windows system. These native environments do, however, allow access to advanced functionality such as fax, telecommunications devices for the deaf (TDD), and analog display services interface (ADSI), and system extensions coded as custom servers (on AIX) and User actions (on Windows).

### ***WebSphere Voice Response for AIX State Tables***

Essentially, a voice application on WebSphere Voice Response is a collection of state tables, prompts, voice segments, and servers, which together provide the desired function. The 'top-level' state table is the one specified in an application profile, and hence it is convenient to refer to the voice application by the name of the state table.

The state table is the program that specifies the basic logic of the interaction with callers. A state table can invoke other state tables, and uses prompts and voice segments to communicate with the caller. A voice segment is an audio recording, usually of spoken words. A prompt is a small program that specifies the logic of the voice output that callers hear, allowing the efficient combining of voice segments, rather than recording every complete utterance separately.

State tables can optionally call servers to provide other functions. A server is a program that performs as a 'bridge' between WebSphere Voice Response for AIX and other software components. WebSphere Voice Response for AIX supports two types of servers:

- A 3270 server enables information access on remote computers that use the 3270 data stream.
- A custom server is a C language or C ++ language program with which access to data or applications on any type of computer is provided, including the pSeries on which WebSphere Voice Response for AIX is running. Other functions for which custom servers

can be used include internal fax, speech recognition, or text-to-speech servers, and integration with CTI servers, support of ADSI and TDD devices.

### ***WebSphere Voice Response for Windows T-REXX applications***

Traditionally, users have written applications in Telephony REXX (T-REXX), which incorporates the REXX programming language. In particular, T-REXX applications are used to access WebSphere Voice Response options such as fax, Telecommunications Devices for the Deaf (TDD) and analog display services interface (ADSI). Voice applications consist of one or more voice programs that control the interactions between callers and the various functions of WebSphere Voice Response. Voice programs use modules that control the playing of recorded voice segments or synthesized text segments.

## **WebSphere Voice Response and Speech Technology**

WebSphere Voice Response works with speech recognition, including IBM WebSphere Voice Server. Speech recognition can replace the telephone key pad as the input device, and enables callers to access information without laboring through layers of touch-tone menus. Instead, callers can use spoken commands to retrieve the information they are seeking. Text-to-Speech can enable voice applications to present dynamic data, that is, information which is not known at the time of development of the voice application, to be spoken. Speech recognition and Text-to-Speech is supported by WebSphere Voice Server in each of the Java, VoiceXML, state table and T-REXX programming environments. The runtime components of WebSphere Voice Server work with the speech recognition engines to perform the recognition and with the text-to-speech engines to perform the text synthesis. Recognition and synthesis engines can be distributed across multiple systems so that resources can be shared, and also to provide redundancy.

## **Conclusion**

IBM WebSphere Voice Response with DirectTalk Technology is a versatile, powerful voice processing platform bringing expanded functionality to Interactive Voice Response (IVR) applications, including advanced speech recognition and VoiceXML for the Web. WebSphere Voice Response integrates information from multiple sources and can deliver direct access to services and information all day, every day. Businesses using WebSphere Voice Response can answer and screen a large number of calls simultaneously and promptly, reducing caller wait time and improving customer satisfaction.

## Notices and trademarks

This publication was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. The products offered by IBM in this publication are platforms or middleware that provides the capability for applications you may elect to add, design, or create to take advantage of the products. Notwithstanding the terms of any other agreements You have with International Business Machines Corporation, or any of its related or affiliated companies, IBM shall not be liable to You for any and all claims of patent infringement, including inducement or contributory infringement, or any claims for indemnification for such claims brought against the products or based on the combination, use or operation of the products with software or hardware. You also should be aware that it may be necessary for You to obtain a patent license from one or more third parties, including, but not limited to, Ronald A. Katz or Ronald A. Katz Technology Licensing, L.P. (commonly referred to as RAKTL), before using certain applications designed and built to run on the IBM products listed in this publication.

IBM may have patents or pending patent applications covering subject matter described 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  
North Castle Drive  
Armonk, NY 10504-1784  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law:

**INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.**

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new

editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of any of the products mentioned herein who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Department T01B  
3039 Cornwallis Road  
Research Triangle Park, NC 27709-2195  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

#### Trademarks

The following terms are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX  
CICS  
DB2  
DirectTalk  
IBM  
the IBM logo

ViaVoice  
Visual Age  
WebSphere

Dialogic and Intel are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Copyright 2002. International Business Machines Corporation. All rights reserved.