CIRCUIT CELLAR®
THE MAGAZINE FOR COMPUTER APPLICATIONS

FEATURE ARTICLE

Shawn Arnold

Designing a DSP-Based RAS Server

Part 1—RAS Server Background

Learn all you need to know about Remote Access System servers in Shawn's introduction to RAS servers. If you've ever considered implementing a DSP-based RAS server, you're going to want to start by reading this article.

In Part 1, I present introductory material on RAS servers. This material begins by taking a general look at RAS servers and introducing the most basic concepts and terminology. I include a discussion on RAS servers are used and what services and functionality they provide. I then look specifically at the Port block of a RAS Server presenting its most basic concepts and terminology. Finally, to help solidify our thoughts, I present some common examples of were RAS servers are used and in what context certain services and functionality are required and used.

The acronym "RAS" stands for Remote Access System. In the simplest terms, a RAS server is a network gateway used specifically for remote access to a network. The RAS server provides remote access to LANs, WANs, and the Internet.

Some of today's RAS servers go beyond providing just access services to remote network users. Some RAS servers also provide off-network services (i.e., they provide the local network user off-network access services).

Note that when I use the general term "network," I am specifically referring to either a LAN, WAN, or the Internet. Also, when I use the term "gateway," I'm referring to a network's access point or port.

THE LITTLE PICTURE

At the heart of a RAS server is the RAS port. The RAS port is the point in the gateway where the translation of the data between the network and Telco system clouds occurs. RAS ports can serve as one of several types of translators.

In the most common case, the RAS server provides a modem port to a remote user for access to a given network. This network could be an office LAN, an ISP WAN, a retail store or bank WAN, or even the Internet.

For example, your ISP subscription permits you access to the ISP's network. In almost all cases, the ISP's network provides access to the Internet. In order to gain access to the ISP's network, the user must call, via your PC's modem, the local ISP access number. This call actually goes to a RAS server at the ISP's local point of presence. Via the RAS server at the local access point, the ISP subscriber gains access to the ISP's network and thus all the ISP services such as e-mail, chat rooms, newsgroups, and the Internet.

As a similar example, telecommuting works in much the same way. An employee can gain access to the company network remotely from home by calling a special access number which connects them to the company's RAS server for access to the office network (see Figure 1).

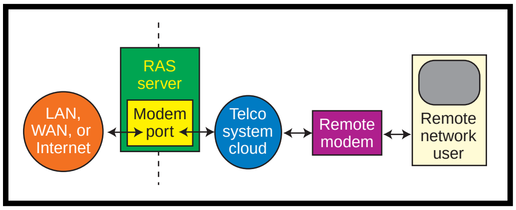

Figure 1—*In a typical telecommuting scenario, a home office PC is connected to the main office LAN via a RAS server modem port.*

Some RAS servers also provide outgoing services sucha as outgoing fax services. Not only can a RAS server provide outgoing fax services, but it can also provide incoming fax services.

An office worker might be editing a document within a word-processing application running on their PC or workstation. Their document may need to be faxed to an office in another location. Normally, they would have to print the document and manually feed it into the fax machine. The RAS server could provide a button within the word-processing application that could be clicked to fax the document directly from their application to the desired fax machine in the other office (see Figure 2).

A new and growing application of RAS servers is for voice services. Basically, the RAS server is used as a voice port to a network. This RAS port application is referred to as "voice over IP" or VoIP. In this case, the RAS port is a voice coder. It compresses/decompresses digitized voice data for passage over a network.

Let's say you call a friend who lives in another state. Instead of calling them directly and incurring long-distance Telco fees, you place this call via their ISP's local point of presence number. The ISP could route the voice call over their WAN or the Internet to the ISP's point of presence at the friend's location. From this point of presence, the call would then be routed over the local Telco system to the friend's phone—without incurring long-distance Telco charges (see Figure 3).

THE BIG PICTURE

Let's examine how the RAS server works. I'm going to break down the RAS server into its basic building blocks and examine how each block functions and how each of the blocks play together.

As I said earlier, the RAS server provides a gateway for remote access to a network. You can see in Figure 4 that the RAS server is a gateway that sits on the boundary between the Telco system and a network. In most cases, the remote access is made via the Telco system. As mentioned earlier, this could be an ISP subscriber, telecommuter, and so on.

At the heart of the RAS server is the RAS port shown in Figure 5. The port provides the transformation of different data representations in the Telco system and in the network. Because of the nature of the Telco system, information passed over this media must be transformed into a special modulated form. The port provides the transformation service from data passing over the Telco system/network boundary. The RAS port performs the majority of the data processing required in a RAS Server.

Depending on the type of connection being supported, the port will provide modem, fax, or voice services. Because the data processing is so numerically intense, DSPs become an attractive solution for the implementing of this so-called "data pump." I'll show you later that DSPs indeed meet the critical design specifications for RAS ports and become the likely choice for RAS port implementation.

The RAS server stands at the boundary of the Telco system, which provides the media for making a remote access and a network (whether it be a LAN, WAN, or the Internet). On the Telco system side of the port is the Telco I/O interface. On the network side of the port is the Network I/O interface. I'll look at both of these I/O interfaces in more detail later.

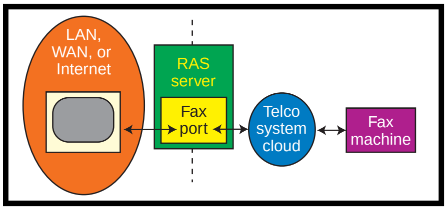

Figure 2—*For an off-network service, an application running on a company's network provides fax service via a RAS server fax port.*

RAS PORT

Let's take a closer look at what each of the services offered by the RAS port really includes. A modem port provides remote access to networks. A modem port must provide:

- data modulations standards—V.90, K56 Flex, V.34, V.32bis, V.32, V.22bis, V.22
- data protocol standards—error correction: V.42, LAPM, MNP 1-4; data compression: V.42bis, MNP 5
- some DLC processing—PPP packetization support

A fax port provides outgoing and incoming fax services. More specifically, a fax port needs to provide:

- data modulation standards—V.17, V.29, V.27ter, V.21 ch2
- fax command set—class 2 T.30
- optional fax command set extensions—ECM, in-line format conversion, BFT, extra fine mode
- image compression—T.4 1D, T.4 2D, T.6

A voice port provides access to networks for routing voice calls, so a voice port must provide voice compression/codec standards (e.g., G.711, G.722, G.728, G.723, G.723.1, G.729). In many cases, the RAS port must provide services to assist in answering or originating calls over the Telco system. The RAS port must provide DTMF and MF tone generation and detection, as well as call progress services.

An obvious question becomes, why the RAS port must support such a variety of standards and protocols? In general, the RAS port must be able to connect to any possible remote device that could call and attempt to connect to the RAS port. A caller may have a 10-year-old modem, or they may have a state-of-the-art modem that supports the newest and fastest modem standards. One important feature of

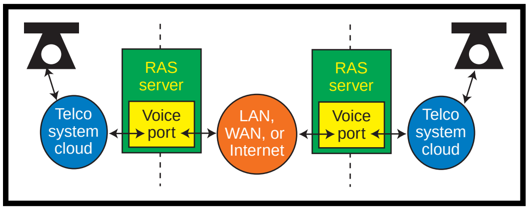

Figure 3—*With VoIP services, a RAS server voice port provides access to a network for voice routing.*

the RAS server is its compatibility with the widest variety of devices that could possibly attempt connection to the port.

NETWORK I/F BLOCK

The Network I/F refers to the RAS server block, which provides the interface to the network. The network data in and out of the RAS port device passes through this I/F. The network provides a byte/character data to the RAS server. As you can see in Figure 6, the Network I/F is made up of two main sub-block areas—the host port controller and the network bus.

The host refers to the main processor or controller that supports the operation of each port and directs the data flow between ports and the network. This host controller can usually handle and control several RAS ports. In this case, "several" means anywhere from 24 to as many as 100 port devices. The host controller processor is usually connected in parallel to the RAS port devices that it supports.

If the RAS port is the heart of the RAS server, then the host controller is the brains of the RAS server. The host is responsible for many things. Depending on the specific design, the host can be responsible for any combination of the following:

- driving the data transfers
- setting up a RAS port device; modem, fax, voice
- answering incoming or originating outgoing calls
- controlling the Telco system I/F transceiver
- dynamically downloading firmware to the RAS port device

The network bus refers to the physical connection between the RAS port devices and the host controller. The host controller supports multiple ports via a parallel connection to the RAS port devices. This parallel connection is usually to a port device parallel port that allows shared memory access or DMA to port device's memory.

Data on the network side of the RAS port is the byte/character form of the data being transferred between the network and the remote user. These bytes/characters often make up the network packets being transferred between the network and the remote application.

HOST CONTROLLER SERVICE DETAILS

The host controller must drive the network data transfers to and from the RAS port device. The RAS port provides the network data in a byte/character format that the network protocol layers understand. Embedded in the stream of byte/characters are the network IP/TCP and PPP frames, which are used by the network devices.

Depending on the chosen design, the host controller will either have to poll the RAS port devices to determine when data transfers are required or it can be interrupt driven via a hardware signal by the RAS port devices whenever host controller data transfer services are required.

Also, depending on the chosen design and modem chipset, the data transfers can consist of single word transfers to/from a RAS port device register or, more preferably, complete buffer transfers to shared memory or via DMA transfers directly to/from the RAS port device memory.

The host controller must configure the port for the correct type of call service. Assuming the RAS port is a reprogrammable device, the host ensures the appropriate setup of the required port service, depending on the type of call. As mentioned previously, the call type could require modem/data, fax, or voice configurations.

Because in most cases the call type is not known *a priori*, the host controller must dynamically configure the RAS port device in real time. This is an important consideration in the design of the RAS server.

Usually it is the host controller's responsibility to set up the physical connection between the RAS port device and the remote device (this includes both incoming and outgoing calls). In the case of an incoming call, the actual detection comes from the Telco system I/F. However, the Telco system I/F will notify the host controller of the incoming call. The host controller must then respond by allocating a RAS port to the call and setting up the port device to handle the specific call type. The host must provide firmware to the port device to configure it for the specific call type (modem, fax, or voice). Also,the host must be sure that the Telco and network data are directed to and from the allocated RAS port.

For an outgoing call, the host performs the proper tasks to place the call and ensure a connection to a remote device. Although the RAS port may provide services such as DTMF tone generation and detection, the host must orchestrate these services to ensure the call is successfully connected.

Similar to incoming calls, the host controller allocates a RAS port to the call and sets up the port device to handle the specific call type. The host does so by providing firmware to the port device to configure it for the specific call type. And again, after the connection is established, the host makes sure the Telco and network data are directed to and from the allocated RAS port.

As you'll see, the Telco system I/F includes a sophisticated transceiver that not only provides for a physical connection to the Telco system trunk

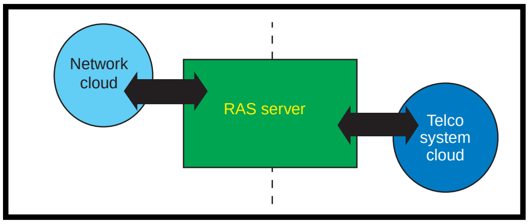

Figure 4—*The RAS server acts as a gateway for remote access to a network. On one side of the gateway lies the network cloud, and on the other, the Telco system.*

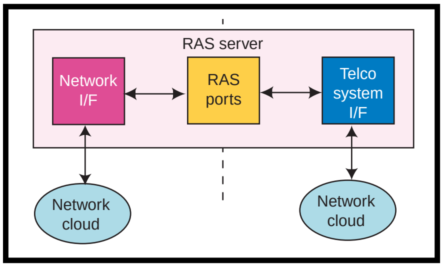

Figure 5—*A RAS server is made-up of several parts. The most important part of the RAS server is the RAS port, which lies at the center of the system.*

line, but also provides all the Telco system data in a more readily usable form. This level of sophistication comes at the cost of added control. These sophisticated Telco system I/F transceivers usually require an additional controller.

This control can be provided by a separate controller, specifically designed in to control the transceiver. To reduce the complexity of the hardware, the control functions could be performed by the RAS port host controller. In this case, the host controller would have to perform all the tasks to maintain the I/F transceiver.

Note that even if a separate controller is chosen, the host controller still has to maintain and interface to the separate controller. Although the controlling burden may be significantly reduced, this is still a required task of the host controller.

Assuming the RAS port is chosen to be a reprogrammable port device, the host must provide for the loading of firmware into the port device. This could range from a static one-time load at powerup, to dynamic loading where specific modules of firmware are dynamically requested by the port and supplied by the host controller.

There may be cases where the RAS port has a level of sophistication that enables it to access a shared EEPROM or flash-memory resource on its own. In that case, the level of host intervention and support is eliminated.

TELCO I/F BLOCK

The Telco I/F is the RAS server block that provides the interface to the Telco system. The Telco data in and out of the RAS port device passes through this I/F.

A trunk line to the RAS server is provided by the Telco system in the form of serially channelized data. The I/F trunk usually takes one of two forms—T1/E1 trunk line or ISDN trunk line.

As you can see in Figure 7, the Telco I/F is made up of three main sub-blocks: the trunk line transceiver, the transceiver controller, and the serial bus. Because the physical signals on the Telco trunk lines are robustly designed for adverse conditions of distance, interference, and minimal cost of material, the physical signals are not compatible with standard digital electronic logic circuits. For this reason, special transceiver chips are available that allow the physical connection of the trunk lines to an electronic device such as a port device.

The trunk line signals are also sophisticated. A lot of information is packed into a single 2-wire serial signal. This information includes clock timing, framing information, the actual data, and switch signaling information. Therefore, the transceiver chip not only provides the physical connection, but also breaks down the sophisticated serial signal into a signal more manageable for the delivery of the serial data to the RAS port device.

Depending on the direction of data flow, the transceiver breaks down the sophisticated trunk line signal and provides the following signals to the RAS Port device:

- serial clock
- frame sync
- PCM Tx & Rx data
- telco switch signaling information

Because the trunk line signals are so sophisticated, the analysis of these signals by the transceiver chip is also sophisticated. So much so that a controller processor is usually required to manage and set up the operation of the transceiver chip.

Note that in some cases, the controller tasks can be moved onto the network side host controller.

The Telco bus refers to the physical connection between the Telco I/F and the RAS port devices. The Telco bus supports multiply ports. This connection is a serial connection to the ports. Usually this is in the form of a connection to a serial port in the RAS port device. The serial bus signal is in TDM form. Each port is assigned a timeslot(s) on the serial bus.

The data on the trunk lines is samples of a digitized analog signal. The digitized signal is that of an analog modulated signal carrying the data being transferred between the remote site and the network over the Telco system media. The analog signal is digitized at an 8.0-kHz rate and is logarithmically compressed down to an 8-bit sample. This digitized stream of samples is referred to as the pulse-code modulated (PCM) stream.

RAS PORT DESIGN

Now that I've examined the major blocks of a RAS server from a functional perspective/standpoint, it's now time to examine some of the critical design goals and design criteria of a RAS server. Our focus is on the RAS port block of the RAS server, so I'll focus on the critical specification and important features of the device that will implement the RAS port portion of the server.

As a manufacturer of networking equipment or, more specifically, RAS servers, in order to make money, I have to offer "more" in the final product, but at the same time, I must build it with fewer resources! In the case of a RAS server the most important "more" is ports.

The critical design goal becomes

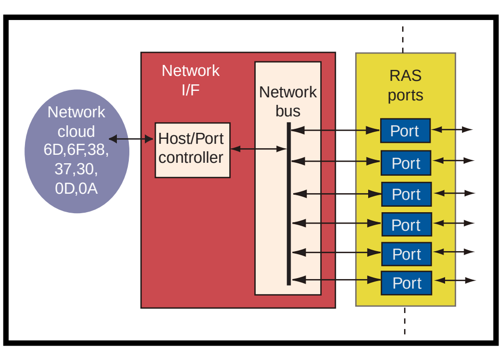

Figure 6—*On the network side of the RAS server is the Network I/F block. This block is made-up of two important pieces—the host controller and the network bus.*

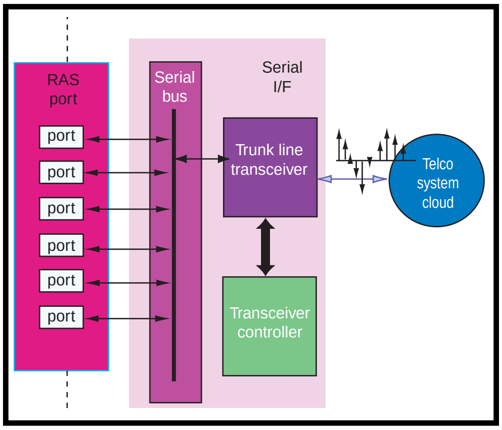

Figure 7—*On the Telco side of the RAS server, the Telco I/F block is made-up of three important pieces—the trunk line transceiver, the transceiver controller, and a serial bus.*

maximum port density (i.e., I want the RAS server to contain the greatest number of ports in the smallest amount of space). RAS server manufacturers are pushing for maximum port density, and thus, for a greater number of ports on each new generation of RAS server products.

Coupled with the desire to reduce port size, is also the desire to minimize the power consumption of each port. Even though the next generation RAS server product will have a greater number of RAS ports in the same amount of space, the manufacturer does not want an increase in the power-consumption specifications of their product. Increasing the power-consumption specifications may cause serious difficulty for their customers, who may have to redesign power systems.

FLEXIBILITY

Along with lowering the cost of each port by reducing physical size and lowering power consumption, I also have to offer more features within each port. As briefly mentioned earlier, the port must be able to handle the various type of RAS services such as data, fax, and voice.

Ultimately, I would like to have these various services supported in a dynamic sense. I don't want the restriction of having to dedicate certain ports to certain services (i.e., I do not want the service built into the port). Rather, I'd like each port to be reprogrammable or reconfigurable in real time so I can configure the port on demand to support the service required by the remote or network user. Such a real-time configurable port solution is referred to as a universal port.

COMPLETENESS

In today's highly competitive market place, the shortest time to market is paramount. This means that the design cycle must be as short as possible.

Designing a RAS server not only includes the physical construction of the hardware that implements the physical system, but it also includes designing the software that must be developed to orchestrate and breath life into the hardware. This software includes code that runs on the various controllers and data processors in the RAS server, including RAS port firmware for modem, fax, and voice services; host controller code for port and network I/F services; and trunk line transceiver code for the controller code.

Any portion of the above software that can be supplied as part of a product offering is of great value. Many RAS port devices are offered with firmware that implements all aspects of each RAS port service. This is of great value to a RAS server manufacturer because it significantly shortens their design cycle.

The implementation of modem, fax, and voice services is a complicated task. The technology required to implement and test these communication standards is difficult to acquire. Thus, it is highly desirable for a RAS server manufacturer to purchase or license this technology in a ready-made firmware form that will run on the platform chosen to implement their RAS server port. A universal port that includes this technology as part of the total solution is quite desirable.

In a smaller sense, procurement of host code that interfaces to the RAS port can also help to reduce costs and time to market. This is more difficult to find, and in many cases, the host code portion of a RAS server is where the differentiation lies between one manufacturer's RAS server and another's.

So, in this brief introduction, I've discussed the RAS server's hardware block diagram and presented several important design criteria. I covered the three main blocks of the RAS server and the fact that the RAS Port is the heart of the system.

Stay tuned for Part 2 where I'll narrow the focus of the discussion and use a DSP to implement the RAS port. You'll see why a DSP is a great choice for a Port device and I'll discuss, in detail, the hardware and software design issues.

Shawn Arnold works at Analog Devices in the VoN Group of the RAS Product Line as an embedded software development engineer. Shawn has held various engineering positions at ADI for the past 14 years, ranging from MCM and DSP Product Engineering to DSP and RAS Applications support. Shawn holds a BSEE and MSEE, both from Northeastern University in Boston, MA. When not at Analog Devices, he spends time with his 5-year-old son, Alex, and 1.5-year-old daughter, Ashley. You may reach him at shawn.arnold@analog.com.