RAS 2000™
Web Interface Guide

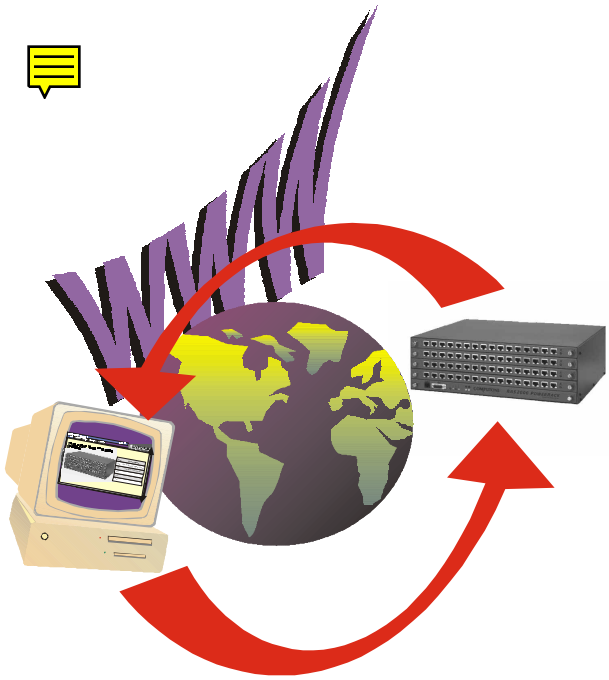

COMPUTONE
Corporation

1060 Windward Ridge Parkway, Suite 100 Alpharetta, GA, 30005-3992 (USA)
(800) 241-3946 s Outside U.S./Canada: (770) 625-0000
FAX: (770) 625-0013 email: sales@computone.com
INTERNET World Wide Web - http://www.computone.com

Computone Corporation
1060 Windward Ridge Parkway
Alpharetta, GA 30005-3992
U.S.A.

Disclaimer: Computone Corporation ("Computone") makes no representations or warranties with respect to the contents hereof, and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Computone reserves the right to revise this publication and make changes from time to time to the contents hereof, without obligation of Computone to notify any person of such revisions or changes.

FCC Statement: This equipment has been tested and found to comply with the limits of a Class A device, pursuant to Part 15 of the United States FCC regulations. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

There is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception (which can be determined by turning the offending equipment off and then on), you are encouraged to try to correct or remove the interference using one or more of the following methods: (a) reorient or relocate the receiving antenna; (b) increase the separation between the equipment and the receiver; (c) connect the equipment to an outlet on a circuit different from that of the receiver; (d) consult the dealer or an experienced radio/television technician for assistance.

Industry Canada Statement: "This Class A digital apparatus meets all requirements of the Canadian Interference Causing Equipment Regulations."

"Cet appareil numérique (de la classe A) respecte toutes les exigences du Règlement sur le materiél brouilleur du Canada.

Support Information: If you require technical support, contact your Computone dealer or Computone Technical Support. The Computone Technical Support staff can be reached by phone at the following numbers, from 8:30 a.m. to 8:00 p.m. Eastern time, Monday through Friday:

(800) 241-3946 ext. 2002
(770) 625-0000 ext. 2002
(770) 625-0013 (FAX

Technical Support can be contacted by email at the Internet address **support@computone.com**

Trademarks: Computone and IntelliServer are trademarks of Computone Corporation. All other brand names or product names are trademarks or registered trademarks of their respective corporations.

Contents

CHAPTER 1

Starting the Web Interface

The RAS 2000 provides an Internet Browser Interface to make configuration easier. To access the RAS 2000:

1. Enable your browser.
2. Enter the IP address of your RAS 2000. In this case, 160.77.25.14.

The main menu is displayed.

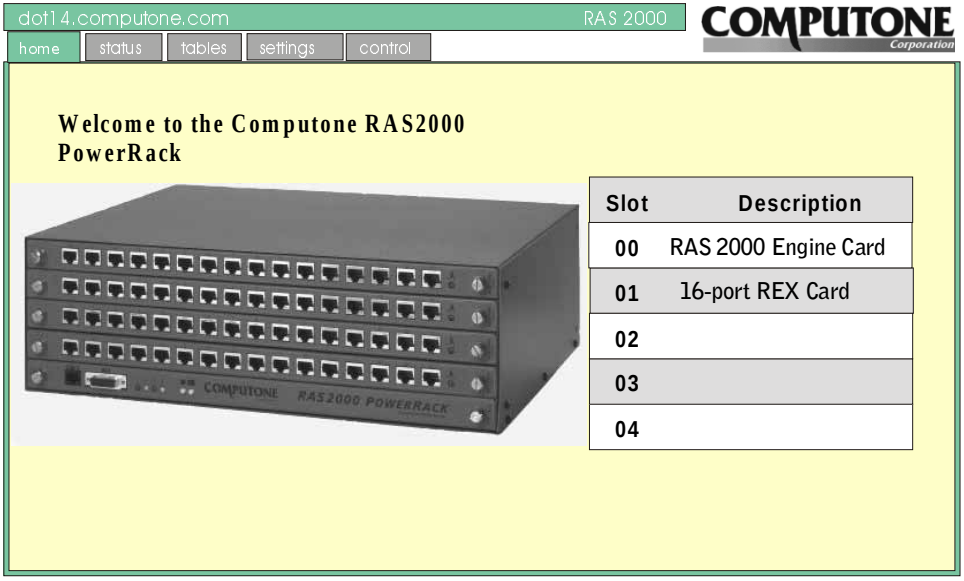

Figure 1 RAS 2000 Main Menu

The tabs across the top of the menu provides access to the following features:

Table 1 RAS 2000 Web Interface Features

Tab	Selections	Description
Home		Welcome screen and RAS 2000 hardware configuration information.
Status		
	Activity	Shows active users on the RAS 2000.
	Processes	Reports the status of all processes running on the RAS 2000. A lot of this information is meaningful only to the RAS 2000's software engineers, so details are not provided in this manual.
	ARP	Address Resolution Protocol (ARP) is a protocol for determining the correct Ethernet address of a host when its IP address is known.
	Routes	Tells the RAS 2000 where to send IP packets based on the IP address. The basics of routing are discussed in ***IP Addresses and Routing*** in the *RAS 2000 Software Configuration Guide*.
Tables		
	Name Servers	On larger networks, a single host (or small group of hosts, for redundancy) is given the responsibility of storing host names and addresses. Such a host is called a *nameserver* and its job is to listen for requests from other hosts and supply IP addresses for particular names.
	Gateways	The gateway table contains *static routes* which are automatically added when the RAS 2000 starts up and when any new SLIP or PPP links are brought up. Internet Protocol (IP) uses these routes to ensure that data reaches its proper destination. For details on routing tables, see ***IP Addresses and Routing*** in the *RAS 2000 Software Configuration Guide*.
	Hosts	The RAS 2000 uses its Host Address Table to resolve host names into IP addresses. Hosts not found in the local table are resolved through external name servers. Each host name on the left is assigned the IP address on the right. New entries to the table are available as soon as they are added.
	Services	You can display the services table. Changes to service ports take effect when the associated process starts up. For most practical purposes this means the changes don't take effect until after the changes are saved and the RAS 2000 rebooted.
	Users	Shows the RAS 2000 users that are configured and provides a means to add new users or delete existing users.
	RSP Profiles	Use to add and assign Remote Serial Port (RSP) profiles.
	RSP Configurations	Use to assign a RSP profile to each of the possible 64 ports of a RAS 2000.

Table 1 RAS 2000 Web Interface Features

Tab	Selections	Description
	IP Filters	Use to add and configure IP filters.
	Ports	Use to configure port parameters.
	IntelliSet Profiles	Use to add or configure IntelliSet Profiles.
	IntelliPrint Profiles	Use to add, delete or configure IntelliPrint profiles.
	IntelliView Profiles	Use to add, delete or configure IntelliView screen profiles.
	PPP Option Profiles	Use to add, delete or configure PPP Option profiles.
	Remote Profiles	Use to add, delete or configure remote profiles.
	Login Scripts	Use to add, configure or delete login scripts for remote devices.
	Dial Scripts	Use to add, configure or delete dial scripts for remote modems.
	Global Connections	Use to add, configure or delete a connection to a user's configuration.
Settings		
	System	Use to set or verify system parameters, such as host name, domain name, IP address, Ethernet address, IP filter, RIP type, login prompt, password prompt and user prompt.
	Applications	Use to enable or disable web server (htppd), secure shell (sshd) or insecure shell (telnetd).
	Boot	Use to select the boot type, host, source file and retry count for system booting.
	Syslog	Use to specify Syslog Host, Syslog Facility, and Syslog Priority.
	SNMP	Use to enable or disable and to configure the Simple Network Management Protocol (SNMP).
	RADIUS	Use to configure Remote Authentication Dial-In User Service (RADIUS).
	RIP	Use to enable or disable and to configure Routing Information Protocol (RIP).
	Secured Shell	Use to configure the Secure Shell (sshd).
	Web Server	Use to configure the Web Server (httpd).
Control		
	Shutdown	Use to shutdown the RAS 2000.
	Save to NVRAM	Use to save the current working configuration to NVRAM.
	Save to host: file	Use to save the current working configuration to a TFTP host.

Home Tab

The ***Home*** tab shows the hardware configuration of the RAS 2000.

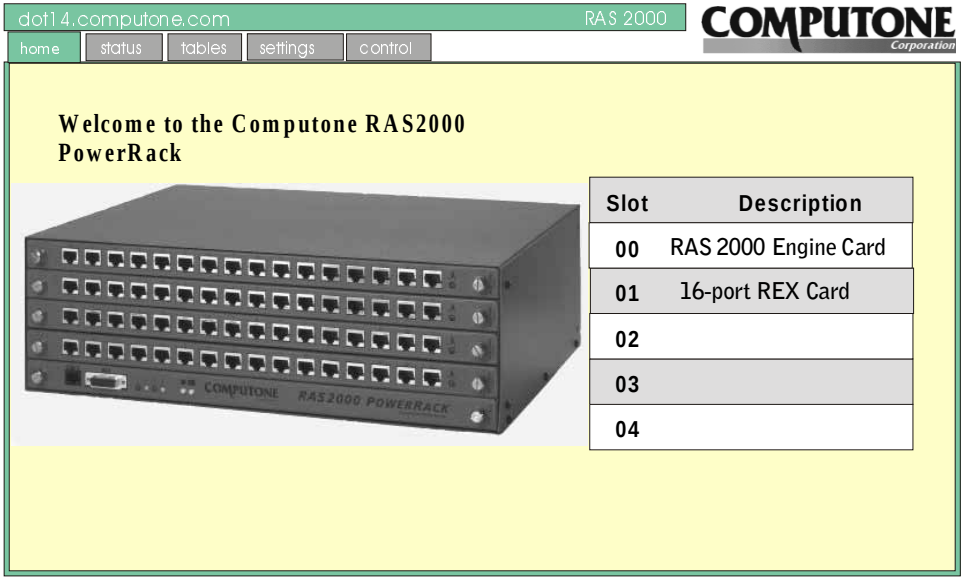

Figure 2 Home Tab Screen

The RAS 2000 comes with the Engine Card and one 16-port REX Card. The Engine Card is installed in the bottom slot (slot 0). A single REX serial interface card is also installed in slot 1. This card has the first 16 serial ports and these ports are numbered 0 - 15. Additional REX cards are sold separately and you can install up to three more, for a total of 64 serial ports.

Where to Install Additional REX Cards

The first additional REX serial interface card should be installed in slot 2, use slot 3 for the next, and slot 4 for the final card. Figure 3 shows the card slots and port numbers.

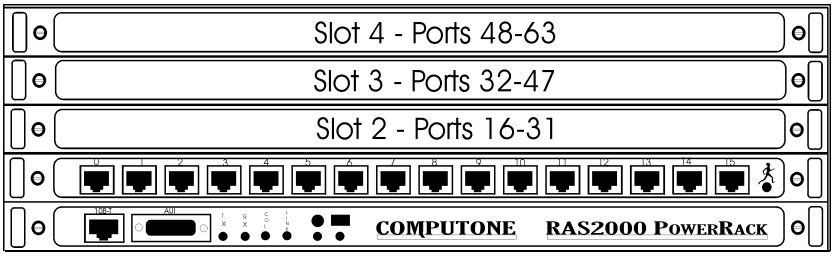

Figure 3 Card Slots and Port Numbers

It is possible to install REX cards in nonconsecutive slots. If you do this, the ports are still numbered according to the slots you have used. For example, if REX cards are in slots 1 and 3, with slot 2 empty, the RAS2000 PowerRack recognizes ports 0-15 and 32-47. Ports 16-31 are not recognized. For safety reasons, always keep unpopulated card slots covered with the card blank panels provided.

Labeling the Ports

On every REX card the ports come individually labeled 0-15, with port 0 on the left and port 15 on the right. For slot 1 these port numbers correspond to the true port numbers recognized by the RAS2000 PowerRack For the remaining slots, the marking on the port does not correspond to its real port number. Therefore, to help you remember which ports on the additional REX cards correspond to which port number, labels are provided in the accessory kit provided with the REX Card. The labels are marked “PORTS 0-15”, PORTS 16-31”, and so on. On the front panel of each REX card there is an oval outline on the left side. The label are designed to fit these outlines.

CHAPTER 2

Viewing the System Status

Sometimes it is helpful to be able to view the system status so you can understand why the system is operating in a particular fashion.

Status Tab

Selecting the ***Status*** tab displays the following screen.

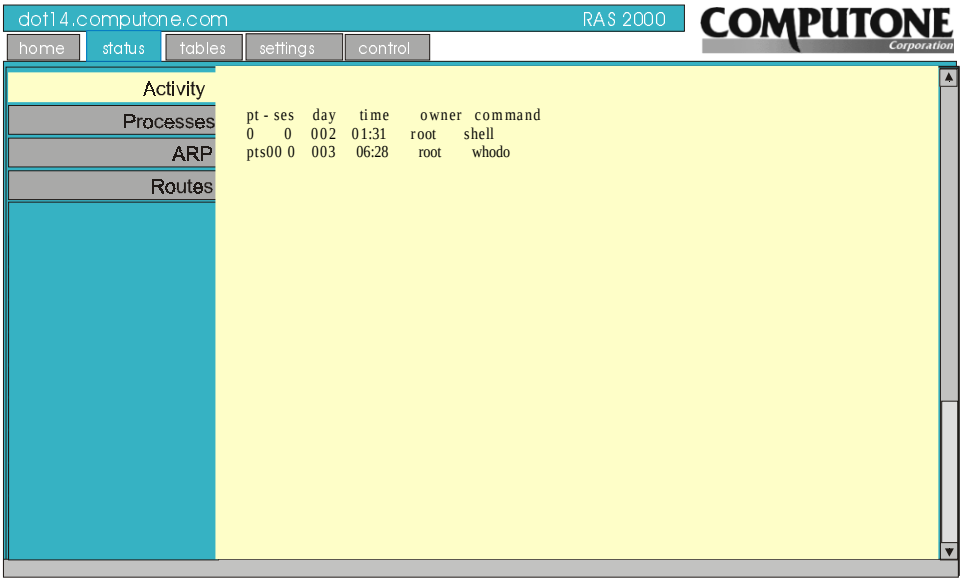

Figure 4 Status Tab Screen

The status parameters available for examination are:

Table 2 *Status* Selections

Parameter	Description
Activity	Shows active users on the RAS 2000.
Processes	Reports the status of all processes running on the RAS 2000. A lot of this information is meaningful only to the RAS 2000's software engineers, so details are not provided in this manual.
ARP	Address Resolution Protocol (ARP) is a protocol for determining the correct Ethernet address of a host when its IP address is known.
Routes	Tells the RAS 2000 where to send IP packets based on the IP address. The basics of routing are discussed in ***IP Addresses and Routing*** in the *RAS 2000 Software Configuration Guide*.

Activity

Selecting Activity displays the following screen. The information that can be obtained from this screen is what ports are active and who is the owner.

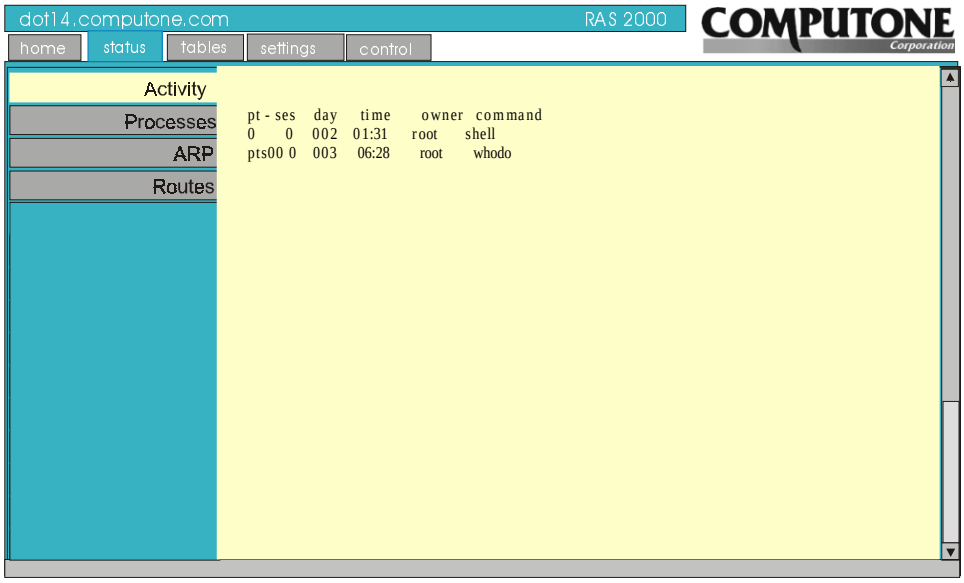

Figure 5 ***Activity*** **Screen**

Processes

Selecting ***Processes*** displays the following screen.

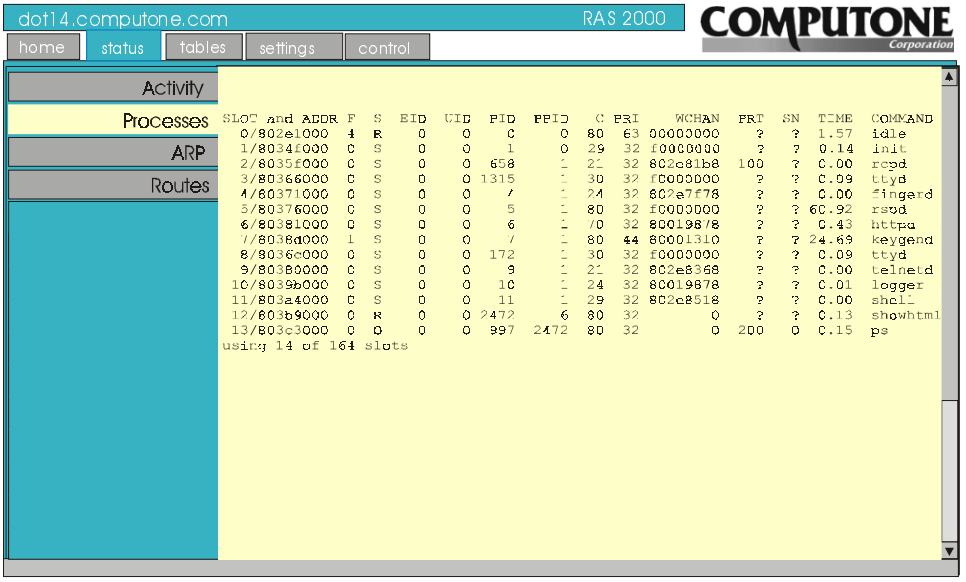

Figure 6 *Processes* Screen

A lot of this information is meaningful only to the Computone's software engineers, so details are not provided in this manual. The columns that are of interest to you are:

- **PRT** - the port number the processes is using. Port numbers 200 and 201 represent the sessions created when you telnet into the RAS 2000. A question mark under this column indicates *daemon* processes not associated with a particular port.
- **COMMAND** - the name of the process or command that is running on that port.
- **TIME** - the number of seconds of CPU time this process has used since it started.
- **PID** - the Process ID number.

ARP

Selecting ARP displays the following screen.

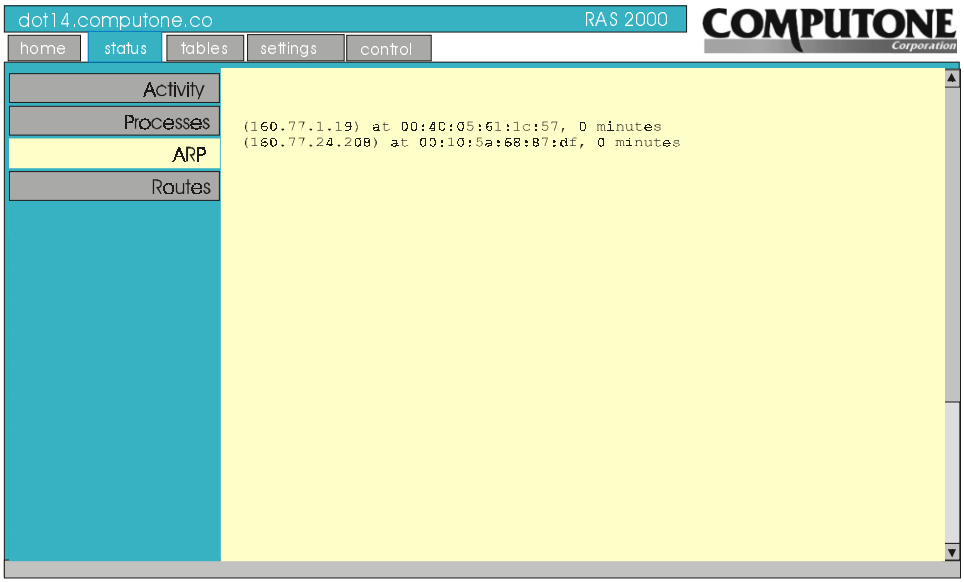

Figure 7 ***ARP*** **Screen**

The ***ARP*** screen shows the current ARP table. This includes the following data:

• Host name - If the RAS 2000 can determine it from the IP address stored in the ARP table.

• IP address.

• Ethernet address, if known. - If there was no response to the RAS 2000's AR request for this address, the address is marked **(incomplete)**.

• Number of minutes this entry has been in the table since it was last referenced.

• Option flags - **published**, **permanent**, or **trailers**. If an entry is marked *published*, the RAS 2000 will respond to any ARP requests it receives for that IP address's Ethernet Address. If marked *permanent*, the entry will not expire; other entries will be removed from the table if there is no activity for that host for a long period of time. The option flag, *trailers* is reserved for future use.

For more information, refer to ***ARPTable*** in chapter 12 of the *RAS 2000 Software Configuration Guide*.

Using ARP to Determine Ethernet Addresses

Since it would be impractical to manually maintain tables of IP addresses and corresponding Ethernet Addresses, there is a protocol, called *Address Resolution Protocol* (ARP) which does this automatically. When a host wants to send a packet to some other host on the local network but does not know its Ethernet address, it broadcasts a request to everyone on the local network, saying in essence, "Does anyone know what the Ethernet Address is for IP address 160.77.99.103?" Since the question is broadcast to all the hosts on the local LAN, it should be seen by the host you are looking for. It knows its own Ethernet address and IP address and so it sends back a reply: "160.77.99.103 can be reached via Ethernet address 80:4e:5f:ca:ff:ee".

Now that the sending host knows the Ethernet address, it can send the packet. Suppose it gets another packet for the same host. Does it start all over and send an ARP request again? That would not be wise, because each ARP request is broadcast to every host on the network. If you were going to do this for every packet, you might as well have broadcast each packet to everyone. The other hosts on your network have better things to do than read broadcast messages and throw them away. So, once your host has learned the Ethernet address for a particular IP host address, it retains the information in an *ARP Table*. It always checks its local ARP table first before sending an ARP request.

Do ARP entries stay in the ARP table forever? Generally, no. It is possible to store a permanent ARP entry in the table, but normal entries are dropped from the table if they have not been used for a long time and there is usually a way to purge entries from the table manually. This handles the case where some Ethernet card has been changed, but the IP address has stayed the same. Anyone on the network with ARP table entries made before the swapped out the card have stale information. By removing the ARP reference manually (at the command line with a `delate arp <IP address>`), you don't need to wait for the entry to expire. With the old entry gone, the host will need to perform another ARP request, and in doing so gets the new information.

Permanent ARP entries are only permanent as long as the RAS 2000 stays up. They are not stored in NVRAM the way static routes in the *Gateway Table* are, for example. This is generally not an issue, because usually the only permanent entries you deal with are the ones created automatically for proxy ARP on behalf of remote hosts.

Proxy ARP

Usually when an ARP request is sent to the network, the target of the request can answer for itself. "Yes, I have that IP address and here is the Ethernet address". But sometimes, it is necessary for a *different* host to answer on its behalf: "Yes, I know who that IP address belongs to, and here is its Ethernet address". A very useful instance of this is when a host is configured to report its *own* Ethernet address as that of the target. This is known as *Proxy ARP*.

If a host is configured for Proxy ARP and it reports its own Ethernet address a being the target host's, it had better expect to receive packets destined for that host. When it receives the packets, it will use its own routing table to send the packets to some other interface. ProxyARP, then, would not be used if the target host were actually connected to the local LAN. If it were so connected, it could answer ARP requests on its own behalf and receive its own packets. Nor can proxy ARP be used when the target's host IP address is not a member of the local LAN's network. Only host addresses that are members of the local network (as determined by the network portion of their addresses) would have been sent to the Ethernet interface in the first place.

Routes

Selecting ***Routes*** displays the following screen.

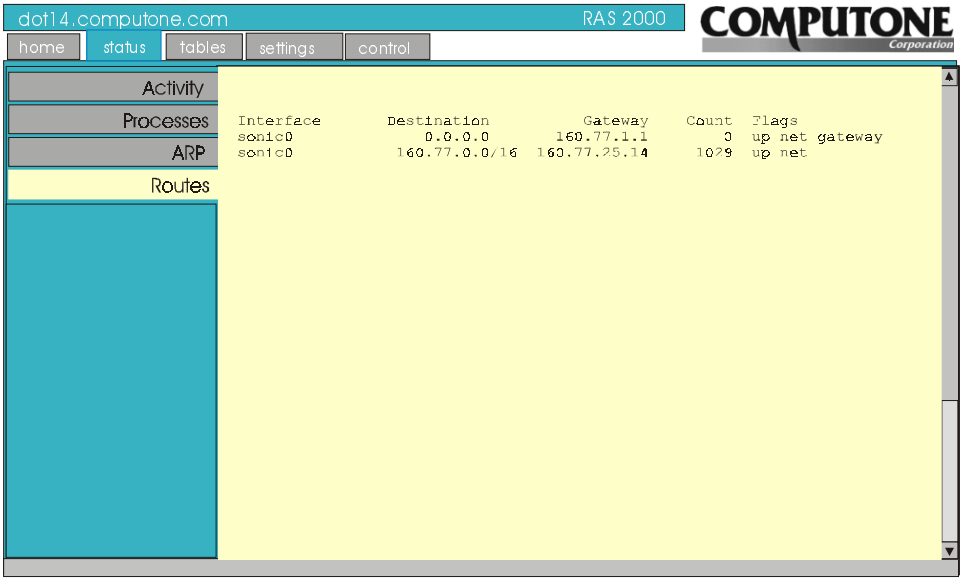

Figure 8 ***Routes*** **Screen**

Routing is the process of directing an IP packet to its proper destination. When there is only one network, routing is trivial and so it is easy to ignore the issue. When there are several networks and you need to route packets from one network to another, you can no longer ignore the issue. When a host has a packet to send (either one it has generated itself or one it received from the network), it could do one of five things with it:

1. Send the packet to an appropriate process running on this host, because the packet is addressed to the host itself.
2. Send the packet to a local network. This would include packets addressed to other hosts on the same Ethernet LAN, for example.
3. Send the packet to a host connected to a PPP or SLIP interface.

4. Send the packet to a *different* host on the local network or PPP/SLIP interface; that host being expected to forward it to the correct host.
5. Discard the packet because what to do with it is not known.

How does the host decide what to do? To determine how a packet should be disposed of, the host first considers whether the packet is for itself. This is easy because the host knows its own IP address (or IP addresses, when the host is on more than one network). Packets for *this* host are sent to the appropriate protocol or process to be dealt with locally.

For packets addressed elsewhere, the host uses a routing table, as shown in Figure 8 . Each entry (o*r**route)* in the routing table has a destination address, a gateway address, and an interface.

- If the destination address is a host address, this is a route to a specific host. A route to a specific host takes precedence over other, more general routes.
- If the destination address is a network address, this route applies to any destinations with host addresses on this network.
- If the destination address is zero, this is a default route. Packets sent to destinations not otherwise accounted for are sent via this route.

For more information on routing, see ***IP Addresses and Routing*** in the *RAS 2000 Software Configuration Guide*.

CHAPTER 3

Working with Tables

Most of the configuration that can be done to the RAS 2000 is done through the use of tables. The following table lists the selections available for configuration through the ***Tables*** tab.

Table 3 ***Tables*** **Selections**

Selection	Description
Name Servers	On larger networks, a single host (or small group of hosts, for redundancy) is given the responsibility of storing host names and addresses. Such a host is called a *nameserver* and its job is to listen for requests from other hosts and supply IP addresses for particular names.
Gateways	The gateway table contains *static routes* which are automatically added when the RAS 2000 starts up and when any new SLIP or PPP links are brought up. Internet Protocol (IP) uses these routes to ensure that data reaches its proper destination. For details on routing tables, see ***IP Addresses and Routing*** in the *RAS 2000 Software Configuration Guide*.
Hosts	The RAS 2000 uses its Host Address Table to resolve host names into IP addresses. Hosts not found in the local table are resolved through external name servers. Each host name on the left is assigned the IP address on the right. New entries to the table are available as soon as they are added.
Services	You can display the services table. Changes to service ports take effect when the associated process starts up. For most practical purposes this means the changes don't take effect until after the changes are saved and the RAS 2000 rebooted.
Users	Shows the RAS 2000 users that are configured and provides a means to add new users or delete existing users.
RSP Profiles	Use to add and assign Remote Serial Port (RSP) profiles.
RSP Configurations	Use to assign a RSP profile to each of the possible 64 ports of a RAS 2000.
IP Filters	Use to add and configure IP filters.
Ports	Use to configure port parameters.
IntelliSet Profiles	Use to add or configure IntelliSet Profiles.
IntelliPrint Profiles	Use to add, delete or configure IntelliPrint profiles.

Table 3 *Tables* Selections

IntelliView Profiles	Use to add, delete or configure IntelliView screen profiles.
PPP Option Profiles	Use to add, delete or configure PPP Option profiles.
Remote Profiles	Use to add, delete or configure remote profiles.
Login Scripts	Use to add, configure or delete login scripts for remote devices.
Dial Scripts	Use to add, configure or delete dial scripts for remote modems.
Global Connections	Use to add, configure or delete a connection to a user's configuration.

Tables Tab

Selecting the ***Tables*** tab displays the following screen and selects ***Name Servers***.

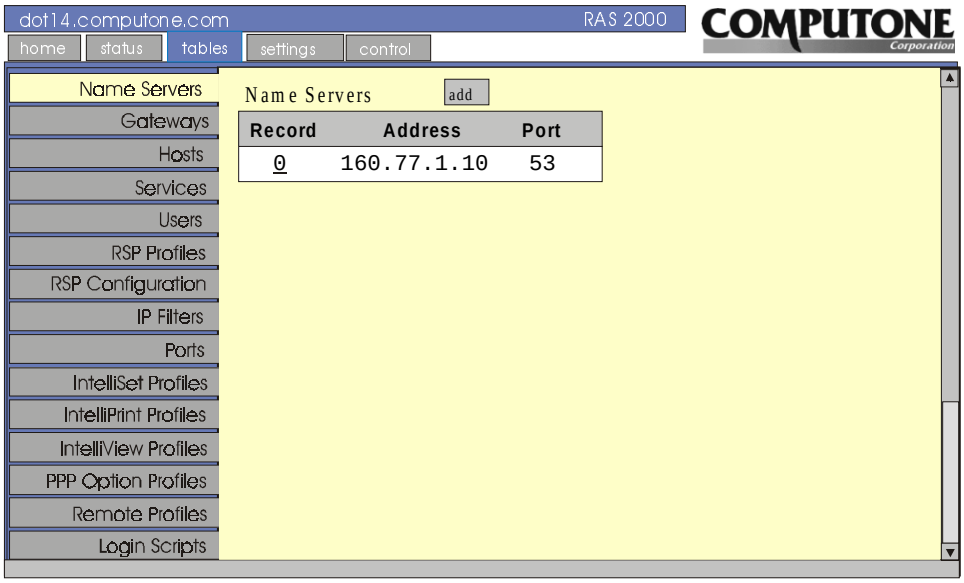

Figure 9 *Tables* and *Name Servers* Screen

On large networks, a single host (or small group of hosts, for redundancy) is given the responsibility of storing host names and addresses. Such a host is called a *name server* and its job is to listen for requests from other hosts and supply IP addresses for particular names. If the name is not found in the local table, the name server might send a request to another name server asking whether*it* might know what IP address corresponds to a particular name. The second name server might have the desired name in its table, or it might be configured to check other name servers. If an IP address is finally discovered, the name server sends it back in a reply.

The process of converting a host name to an IP address is known as *name resolution*. When names are *resolved* through an external name server, the protocol used is called *Domain Name Service*, or DNS.

If you want to add a name server, click on the ***Add*** button in this screen and the following screen is displayed. Edit the ***Address*** and ***Port*** fields, then select ***Update***.

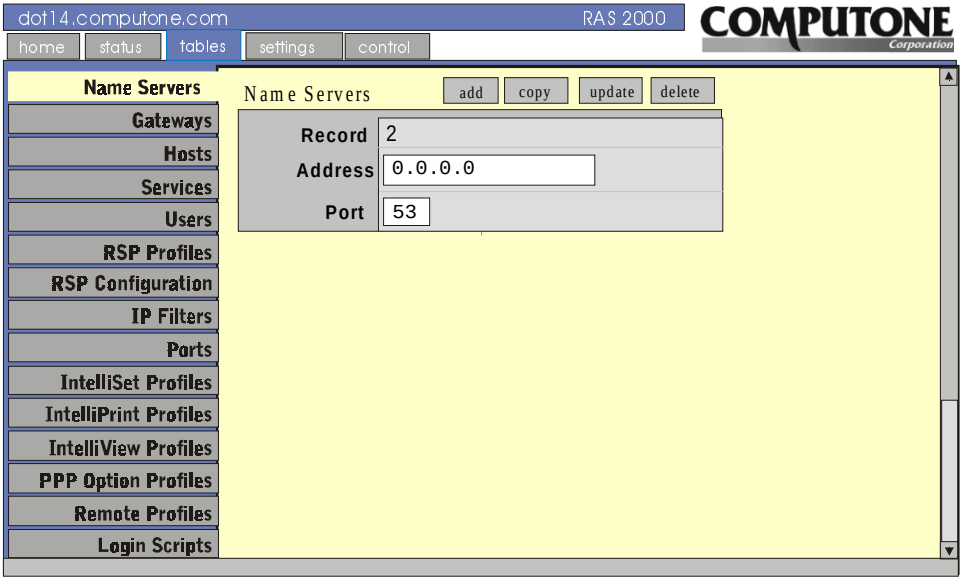

Figure 10 ***Adding a Name Server*** **Scree**

Gateways

Selecting ***Gateways*** displays the following screen.

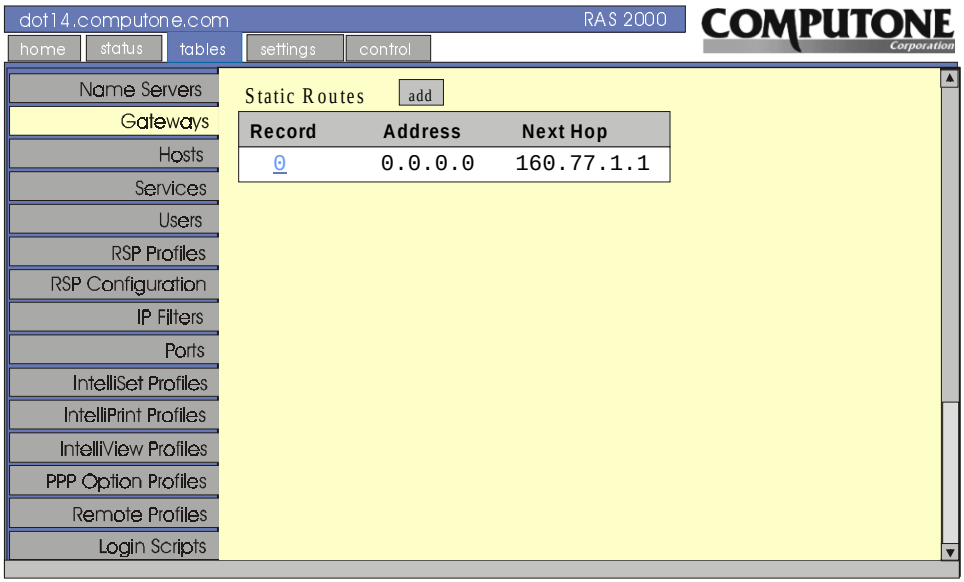

Figure 11 ***Gateways*** **Screen**

The gateway table shown in Figure11 contain***s****tatic routes* which are automatically added when the RAS 2000 starts up and when any new SLIP or PPP links are brought up. Internet Protocol (IP) uses these routes to ensure that data reaches its proper destination. For details on routing tables, see ***IP Addresses and Routing*** on page 181 of the *RAS 2000 Software Configuration Guide*.

Why is the gateway table reread when SLIP and PPP connections come up? There may be some routes in the gateway table whose destinations are unreachable when the RAS 2000 is first started up, because those destinations are reached through SLIP or PPP links that are not yet up. Such a route cannot be added at that time, but *can* be added after the required SLIP or PPP link has come up.

This screen provides the facilities to ***Add, Delete***, and ***Cop*** . After making your changes, select ***Update***.

Hosts

Selecting ***Hosts*** displays the following screen.

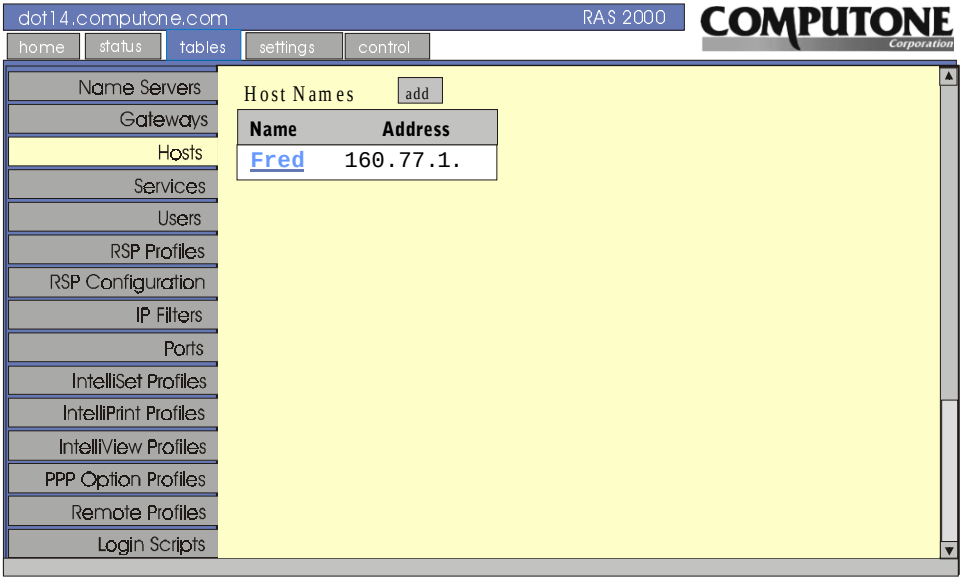

Figure 12 ***Hosts*** **Screen**

The RAS 2000 uses its Host Address Table, as shown in Figur e12 , to resolve host names into IP addresses. Hosts not found in the local table are resolved through external name servers. Each host name on the left is assigned the IP address on the right. New entries to the table are available as soon as they are added.

Services

Selecting ***Services*** displays the following screen.

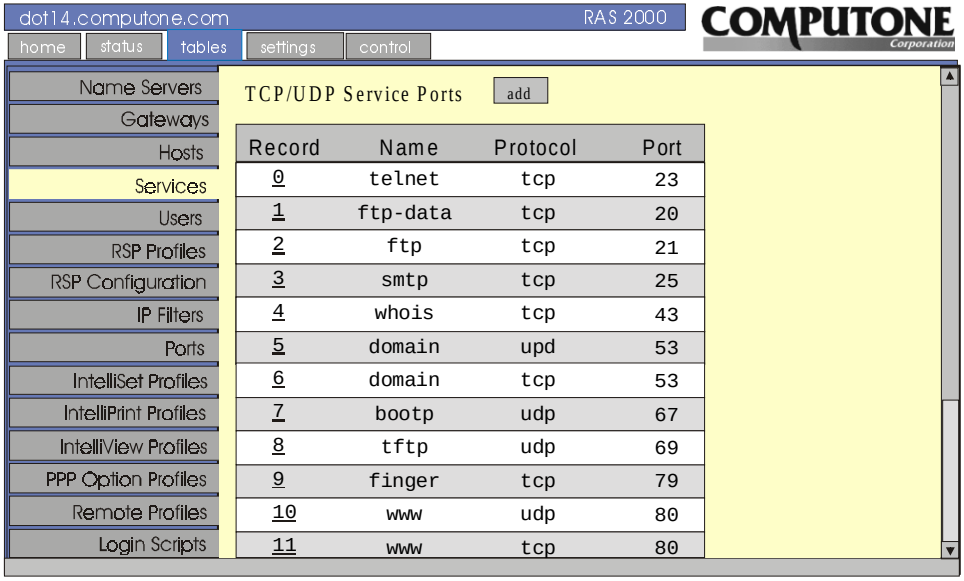

Figure 13 ***Services*** **Screen**

When client and server processes communicate with each other using Internet Protocol, IP addresses in the *IP heade* are used to ensure that the data is sent to the proper host computer. The *IP header* also contains source and destination *port numbers*, which serve to identify which *particular* client or server on a host is the source or destination of that data.

When you are talking about the RAS 2000, “ports” usually refer to *serial ports*, but *these* port numbers have nothing to do with serial ports; they are just numbers used in Internet Protocol. Processes which provide standard services listen on particular *well-known ports*. Client processes which want to get a particular type of service from a host try to make a connection to that *well-known port*. After it does, the server process can assign the client a different port number that applie to that particular session between those particular processes. Standard well-known port numbers have been assigned to standard services and are listed in the RAS 2000’s *Service Ports* table. You will probably never need to change the entries unless your network is extremely unusual, but the table is provided nonetheless.

You can display and modify the *service ports* table by clicking the ***Record*** number shown in the table in Figure 1 3 . Several protocols are listed, each with its own *well-known port*. The column marked “Protocol” shows whether TCP or UDP protocol is used for that service. Look at the first entry in the table in Figure 1 3 : “telnet, port 23, tcp”. This means that if the RAS 2000 wants to telnet into some host, it needs to contact TCP port 23 on that host. This is a multi-page table and only one page is illustrated here. *The services table may contain entries for protocols that the RAS 2000 does not support.*

Users

Selecting ***Users*** displays the following screen.

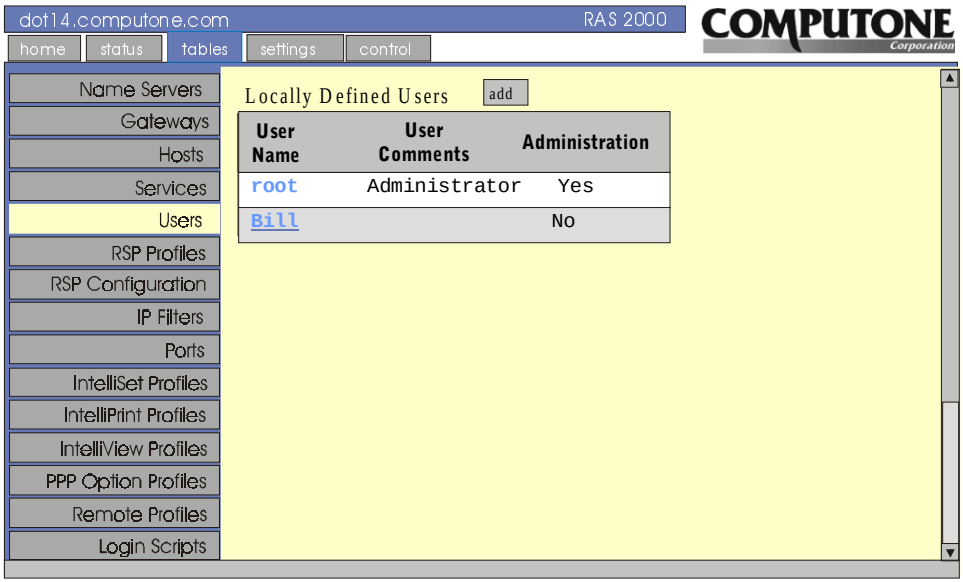

Figure 14 ***Users*** **Screen**

A user that is configured on the RAS 2000 and stored locally in its *non-volatile RAM* is called a ***NVRAM use*** . This is the easiest way to configure a small number of users on a single RAS 2000. The local NVRAM is limited to storing about a hundred users, so if you must support a larger number you must store them on another host. A user whose information is stored on another host on your network is known as a ***RADIUS user*** *(Remote Authentication Dial-In User Service)*, because the RADIUS protocol allows the user information to be sent from this host to the RAS 2000. RADIUS is discussed in more detail in **chapter 4, *RADIUS*** or in **chapter 17** of the *RAS 2000 Software Configuration Guide*.

When you click on a user name, the following screen is displayed.

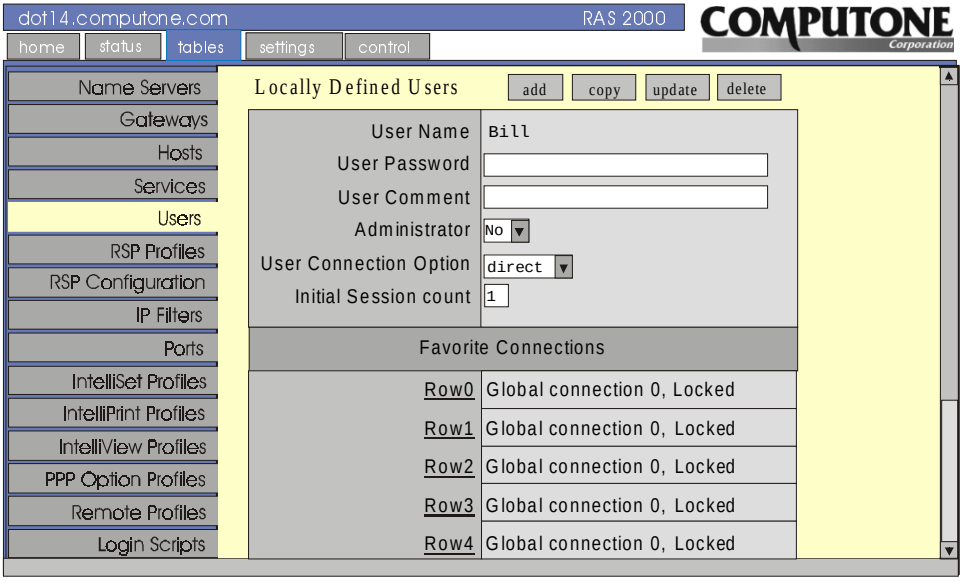

Figure 15 ***User Configuration*** **Sub-menu**

Table 4 ***Users Configuration*** **Sub-menu Entries**

Menu Entry	Description
User Name	User account you are defining.
User Password	Enter a password for this user, if one is desired.
User Comment	Enter something meaningful to identify this user.
Administrator	Select Yes or No, depending on whether you want this user to have Administrator rights.
User Connection Option	This is the master control for what happens when the user logs in. Direct Connect per Screen (***direct***), Selected Connection Menu (***select***), and Full Connection Menu (***full***) are used to support login users or to establish telnet and rlogin connections with other hosts on the network. Inbound SLIP (***SLIP***), Inbound CSLIP (***CSLIP***), Inbound PPP (***PPP***) are used to support dial-in users that want to establish PPP, SLIP, and CSLIP connections to computers or networks at their sites.

Table 4 ***Users Configuration*** **Sub-menu Entries**

Menu Entry	Description
Initial Session Count	This applies to users that have been configured as Direct Connect per Screen and limits the number of sessions that are initially started after the user logs in.
Favorite Connections	
	You can assign up to 8 global connections for each user that is assigned on the ***Users*** menu. To assign a global connection you must first find the global connection number from the global connection table for the service you want to use, and then unlock the row number that you are going to assign. Then, you enter the ***gc*** number in the **ROW** submenu.

Whenever a new selected connection is added to a user's configuration, the entry is automatically added to the *global connection table*. This is a master table that contains all the connections configured for all users. You may also add, modify, and delete entries in the global connection table directly, without working through user configuration. In most installations with lots of login users, there tends to be more users than there are places to go. If the global connection table number is known for a particular connection, then you can configure the user more quickly. More importantly, if some system-wide parameter changes, you are more likely to be able to make a single change and affect all appropriate users. For example, perhaps lots of users are configured to rlogin to a certain host in order to perform a specific function. But later, this function is moved to a different host on your network. You *could* change each user separately or do the following:

- Look at the user configuration form for one of these users. In its selected connection table is the global connection number of that connection. Remember it.
- In the global connection menu, find the entry and change it. All other users using that entry are updated as well.

This is possible because the RAS 2000 automatically forces users with identical connections to share a single global connection entry. Remember, entries must be completely identical. Even the spacing must be identical or separate entries are created. If two users were configured with identical connections and you wanted to make a change for one user only, you would have made the change using user configuration. This would automatically create a new entry in the global table for the user's new connection.

RSP Profiles

Selecting ***RSP Profiles*** displays the following screen.

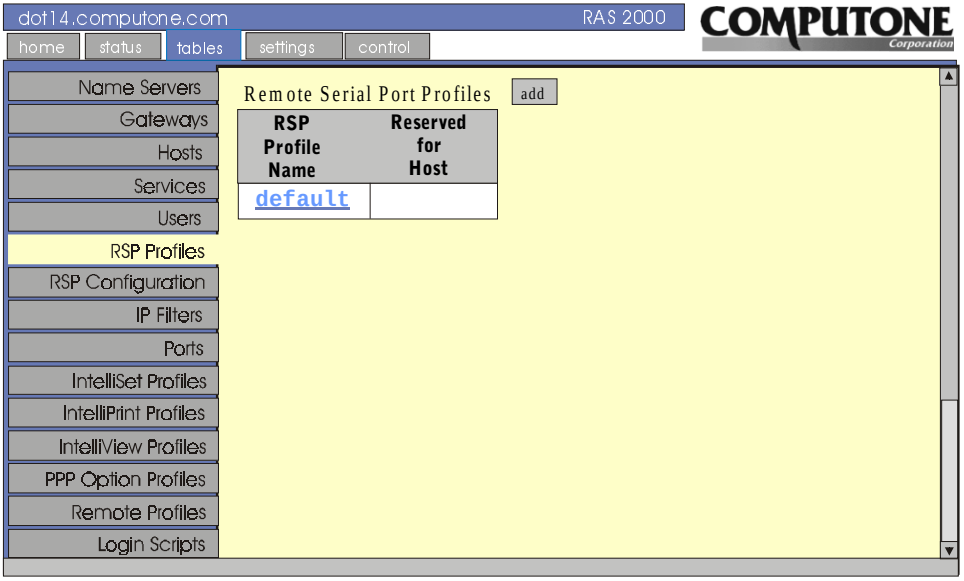

Figure 16 ***RSP Profiles*** **Screen**

Use this screen to assign profiles to devices. When a user dials into the RAS 2000 to bring up an inbound interface, there is certain information that is known right away:

- Port number used
- User name
- Protocol (SLIP, CSLIP, or PPP) needed (either from the NVRAM configuration or from the RADIUS authentication reply)
- Its IP address (for RADIUS users. RADIUS supplies an IP address using the *Framed-Address* attribute. There are two addresses which have special meaning: 255.255.255.255 means the IP address is unspecified and must be deter-mined from PPP address negotiation. 255.255.255.254 also means that the IP address is unspecified, but the RAS 2000 should assign one from a pool.)

Assignment Rules

Given the port, user, protocol, and address, the RAS 2000 must assign a Remote Profile and its interface to finish bringing up the connection. There are four rules for doing this:

1. The Remote Profile's interface cannot already be in use. One Remote Profile per interface per connection. Its interface type must be *Inbound.*
2. The Remote Profile must be compatible with the port number, user name, protocol, and IP address specified for this user.
3. More restrictive Remote Profiles (that are still compatible) are assigned in preference to less restrictive ones.
4. If other network options are specified for a RADIUS user, these option supersede the ones configured in the Remote Profile.

Rules for Compatibility

How does the RAS 2000 decide whether a Remote Profile is *compatible* with the connection's requirements?

- If the Remote Profile specifies a *Serial Port,* it must match the one receiving the connection. If set to **Any**, the serial port does not matter.
- If the Remote Profile specifies a particular *Dial-in User,* that name must match the user that logged in. If left blank, it can be used with any user name.
- If the Remote Profile specifies a particular *Protocol,* it must match the protocol this user desires. If set to **Any**, then it can be used with any protocol.
- If there is no IP address associated with the user, the Remote Profile must contain a valid *Remote Addres* . An IP address is associated with the user if an actual IP address is supplied through RADIUS (using the "Framed-Address" attribute) or if the value 255.255.255.255 is supplied. The latter option is not an address, but at least it is a promise that a real address is forthcoming during PPP address negotiation.

Assignment Priority

Considering now only the Remote Profiles that are compatible, how will the *most* suitable one be found? This is what the RAS 2000 looks for in order, from highest to lowest priority:

- A Remote Profile which specifies both the *Serial Port* and the *Dial-in User.* If this Remote weren't assigned here, it never *could* be assigned, because no one else will be using this port while this connection is using it.

- A Remote Profile which specifies the *Serial Port*, for the same reason as above.
- A Remote Profile which specifies the *Dial-in User*. This means the dedication of a Remote Profile to a particular user takes precedence over other considerations.
- A Remote Profile which species a *Remote Address*, if this address matches a specific one provided from RADIUS. This is a bit like having a matching user name.
- A Remote Profile where the *Remote Addres* is zero, and the IP address provided from RADIUS is either an actual one or 255.255.255.255 (use address negotiation); **OR**, a Remote Profile where the *Remote Address* is valid (non-zero) and the user is an NVRAM user or a RADIUS user with 255.255.255.254 specified for an address (assigned from the pool). In either of these conditions, there would be a total of one remote IP address offered from all the parties.
- A Remote Profile with a non-zero *Remote Address*, and also a valid IP address from RADIUS. The one from RADIUS supersedes the one stored in the Remote Profile.

RSP Configuration

Selecting ***RSP Configuration*** displays the following screen.

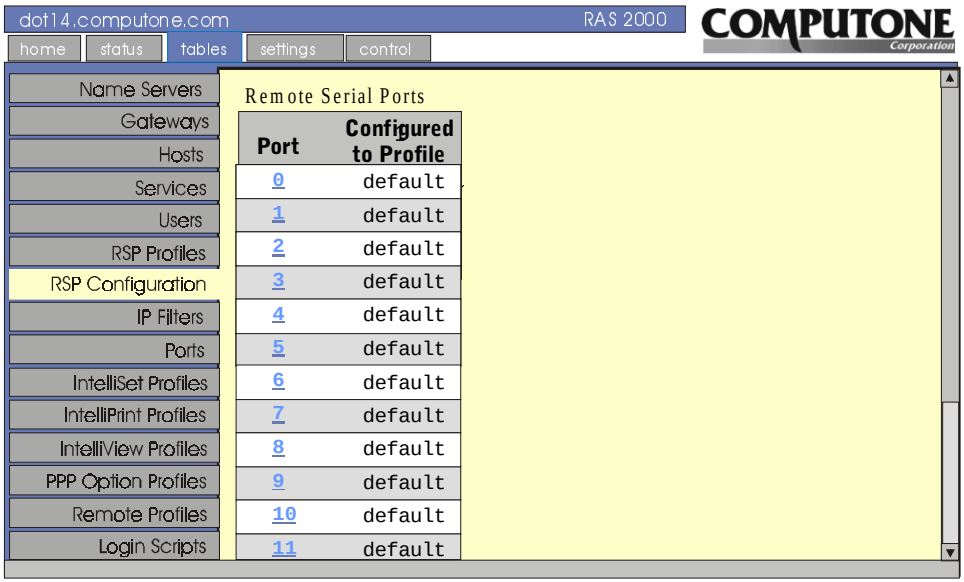

Figure 17 ***RSP Configuration*** **Screen**

This screen shows ports 0 through 63 and to what profile they are configured.

You can reassign a port's profile by clicking on the port number. The following screen is then displayed.

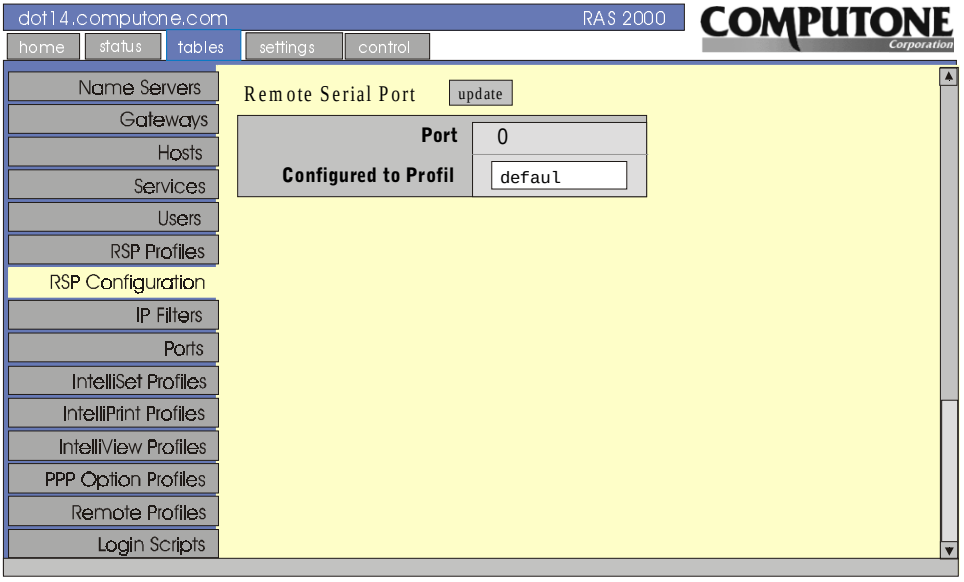

Figure 18 Changing the Port Profile

To configure this port to a new profile, type in the name of the new profile you wish to use in the ***Configured to Profile*** box.

IP Filters

Selecting ***IP Filters*** displays the following screen.

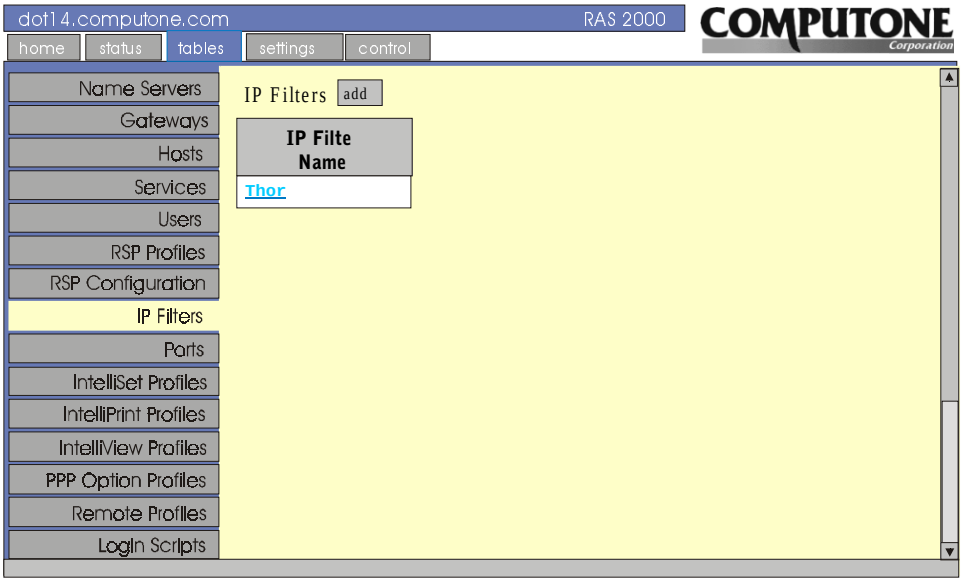

Figure 19 ***IP Filter*** **Screen**

To create a new IP filter, use the **add filter** button shown in Figure 1 9 . When this filter is first created, it has no rules associated with it. When creating a new filter, do not give it a name that is also an interface name: *ether, ppp00, ppp01…*If you do it will lead to confusion.

Once a filter has been created, you can add rules by clicking on the filter name. The following screen is displayed.

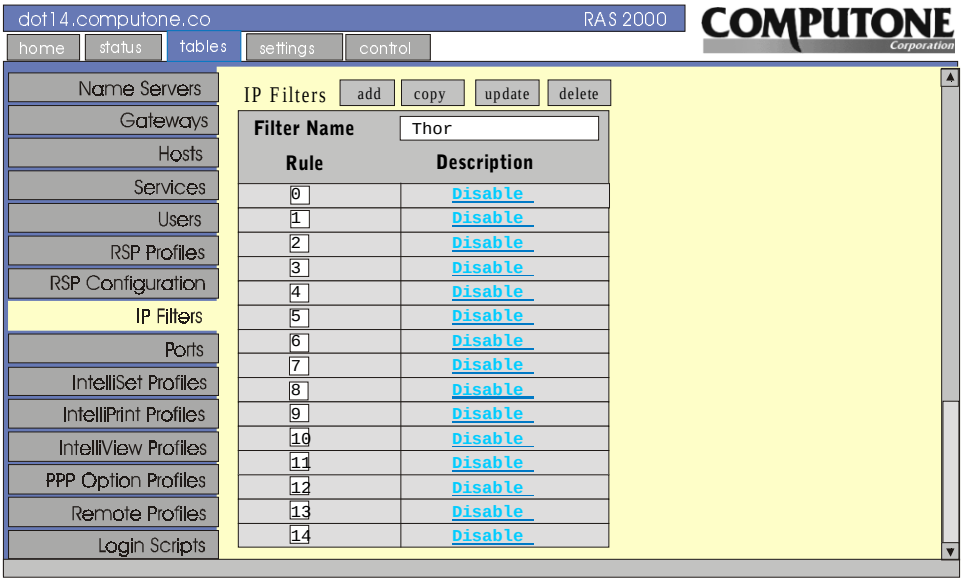

Figure 20 ***Setting the Rules*** **Screen**

From this screen, you can set the rules for the IP filter you just created. Click on the **Disabled** link beside one of the rule numbers.

The following screen is displayed.

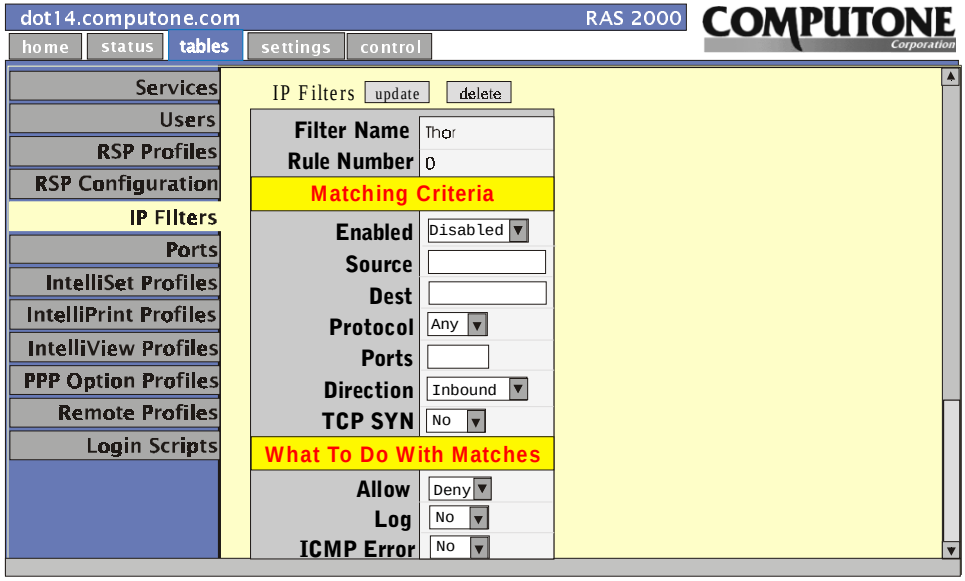

Figure 21 Setting the IP Filter Rules

A rule consists of an *action* and a *test*. As each IP packet is filtered, the rules are applied in order. If the packet matches a rule's *test* (***Matching Criteria***), the *action* (***What To Do With Matches***) associated with that rule is performed. If the action calls for the matching packet to be allowed or denied, further testing stops. For this reason the order of rules is important. More specific tests should be specified before more general ones. If one rule defines an exception to a more general rule, the exception needs to be listed first. This is why you are allowed to insert rules into specific places in a filter.

There are four possible actions you can specify. These are described first, then how you construct the tests is explained.

Table 5 IP Filtering Actions

Action	Description
allow	This packet is allowed to pass.
deny	This packet is discarded.
deny errors	This IP packet is discarded and the RAS 2000 sends an ICMP error message.
log	Do not allow or deny this packet based on the results of this test, but keep a count of how many IP packets have matched. Statistics are kept for all rules. The log action allows you to keep statistics on a condition without making an allow/deny decision based on it.

Tests are constructed of many types of building blocks. A single test may contain several conditions which must all be true for the packet to match.

Table 6 IP Filtering Tests

Keyword	Parameters	Definition	Comments
in		The test is applied to inbound packets.	You have to specify either in or out. Specify both if you want the test to apply to all packets.
out		The test is applied to outbound packets.	
src	*ip addres* *ip address/bits*	The IP packet's source address needs to match the address in this rule.	To match IP packet addresses from a particular network or subnet, specify the number of bits to be tested after the IP address from the class B network 160.77.0.0, specify it as 160.77.0.0/16. This is the same notation used for specifying subnets in routes and is explained in chapter 9, ***Network Basics*** of the *RAS 2000 Software Configuration Guide.*
dest	*ip addres* *ip address/bits*	The IP packet's destination address needs to match the address in this rule.	
tcp		This must be a TCP packet.	If a service port or range of ports is specified in this rule. UDP or TCP ports in that range are matched. If no service ports are specified, all UDP or TCP ports are included.
udp		This must be a UDP packet	
icmp		This must be an ICMP packet, for example, "ping".	

Table 6 IP Filtering Tests

Keyword	Parameters	Definition	Comments
syn		This matches any TCP packet that has the SYN flag set. This flag is always set in the first packet sent over a TCP connection, so this test could be included in a rule to prevent certain new TCP connections from being started up.	
port	*port number*	The destination port in the IP header must match this one.	Since the discussion is about network headers, port refers to the TCP or UDP service port associated with a connection; it has nothing to do with serial ports. When you use the port keyword, you must also specify either TCP or UDP, since the same port numbers could apply to either.
ports	range (e.g. 1-35)	The destination port in the IP header must match this range.	
ports reserved		The destination port in the IP header must be one of the well known reserved ports 1 - 1023.	

Figure 2 2 andFigure 2 3 are only part of the screen shown in Figure 21 and are examples of setting a rule. In the first example this rule allows all incoming packets destined for port 21 (used for FTP connections) of host 160.77.99.30. Specifically, this allows an outsider to establish an FTP connection to one particular host.

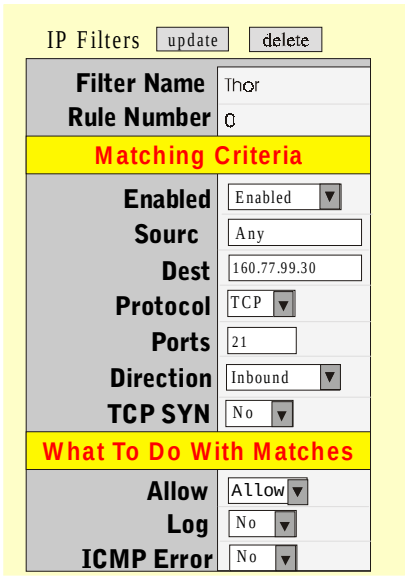

Figure 22 Example 1 of Setting a Rule

The example rule shown in Figure 23 forbids any incoming packets from host addresses in the range 160.77.128.1 - 160.77.255.254.

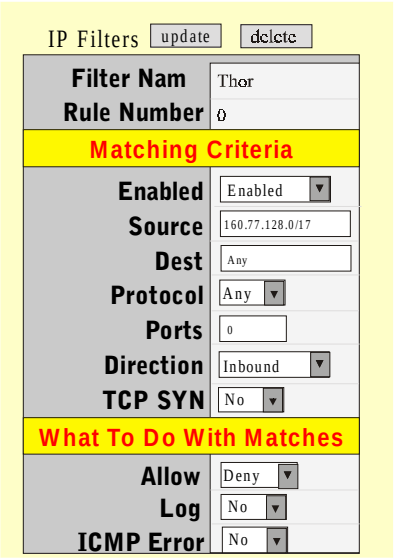

Figure 23 Example 2 of Setting a Rule

Ports

Selecting ***Ports*** displays the following screen.

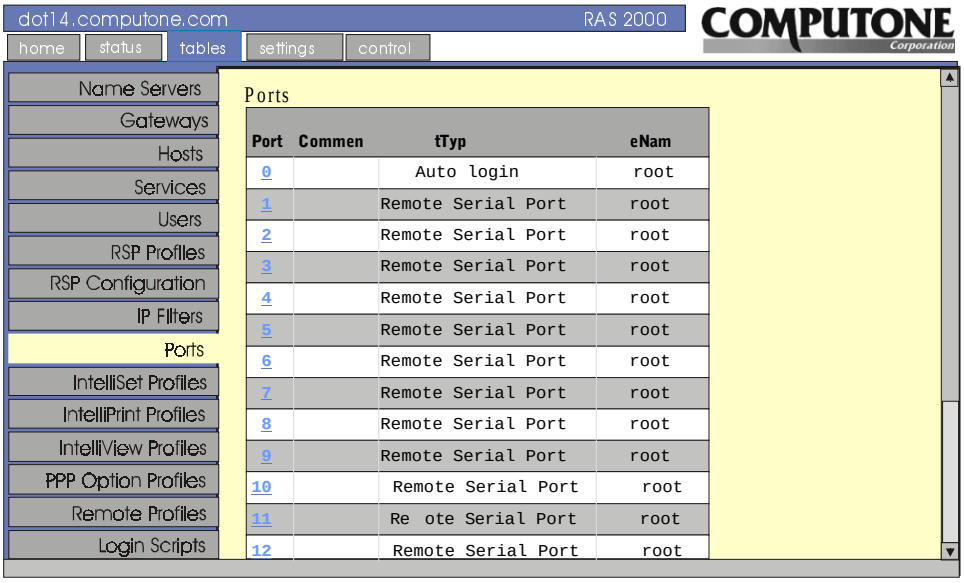

Figure 24 ***Ports*** **Screen**

You can configure serial port parameters by using the ***Ports*** configuration screen. When you change the configuration of a port, the changes take effect the *next* time the port is opened. For example, if a user is currently logged into a port and you change the line speed, the change does not take effect until that user logs off. The new login prompt is issued at the new line speed. An exception is when you change the configuration of the port you are running on. You are then given an option to allow the changes to take place immediately (but return to original values at next login), to take effect at next login, or both.

Port Type—How Will the Port be Used

Using this screen you must configure the RAS 2000 on how each port will be used. Do you want to have someone log into this port? Do you want someone on your network to be able to dial out on a modem attached to this port? Do you want to attach a printer? Do you want to start up a PPP/SLIP link?

There are four port types that support terminals and dial-in connections:

- Login by Port
- Login by Screen
- Auto-Login
- Auto-Login/Wait

For these types, the connection is started by whatever is attached to the serial port. This may be a person sitting at a local terminal, a client who dials into an attached modem, or a computer that dials in and sets up a PPP connection by running a login script. When modems are used, the ports are usually configured so they detect incoming calls by waiting for the modem to assert *carrier* (**DCD**), and when *carrier* is dropped, to recognize that the connection has been dropped. See chapter 6, ***Configuring Modems*****,** of the *RAS 2000 Software Configuraiton Guide* for more information about modems.

There are two port types which can support connections that are started by processes running elsewhere on your network, such as printing and dial-out capabilities:

- Reverse-TCP
- Printer

There is one type which supports dial-in and dial-out connections:

- Login by Port/TCP

There is one port type which supports dial-out PPP/SLIP/CSLIP connections to other networks:

- Out-bound Connection

There is one last port type which supports remote serial port access to the RAS 2000 in a Windows® NT environment:

- Remote Serial Port

Configuring a Port

To configure a port, click on the port number show in Figur e24 . The following screen is displayed.

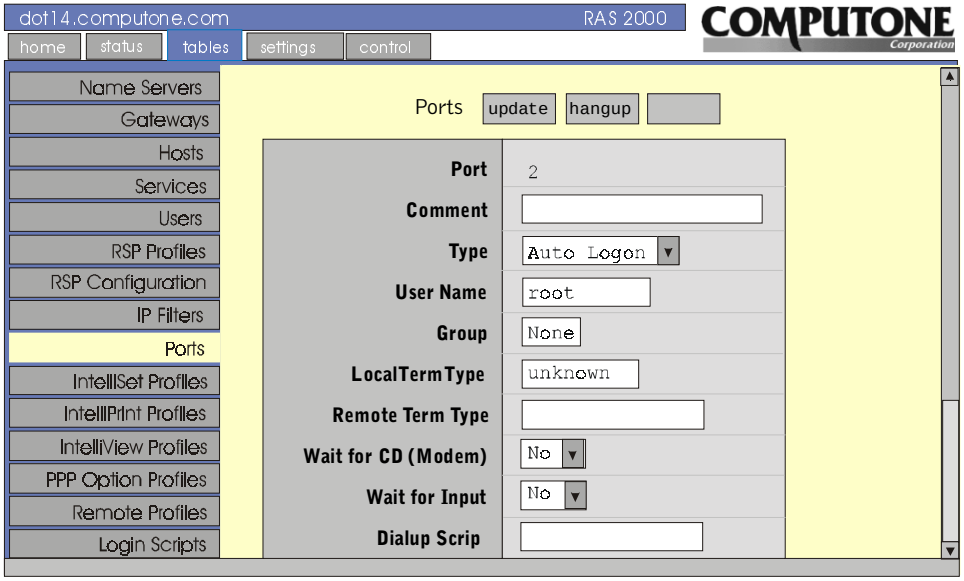

Figure 25 ***Port Configuration*** **Screen**

The following table defines the port configuration parameters

Table 7 Port Configuration Parameters

Parameters	Definition
Port	Number of port being configured
Comment	Some descriptive comment to identify the port to you.
Type	
Disabled	Nothing happens on this port, except that you can use commands such as **tip** and **output** to send data to it in order to test the port or configure a modem or other device.
Login by Port, Wait	With this selection, the port sends a login prompt to the attached terminal or modem. When the user logs in, the RAS 2000 starts up whatever connections have been configured for that user.

Table 7 Port Configuration Parameters

Parameters	Definition
Login by Virtual Screen	You generally use this setting only if you have configured the port to support multiple sessions through IntelliView. Each virtual screen is sent its own login prompt, and the user must log into each virtual screen separately. When a session is ended, a new login prompt for that virtual screen is sent, but DTR is not dropped if there is any other session on this port still active.
Auto-Login/Wait	This is almost identical to *Auto-Login*, except that instead of launching the connection immediately, the port first sends a prompt: **Press <enter> to continue**, and when the operator does this, *then* the connection is launched. This is designed to solve the quandary that occurs when a port configured as *Auto-Login* is attached to a local terminal that is always on but frequently unattended. The user logs off and walks away, and the RAS 2000 immediately launches the connection. Suppose that connection is an attempt to rlogin to some host machine. So that machine prompts for a password. Since there is no one present to enter a password, the connection soon times out and is restarted, and times out and is restarted and so on. If the port is configured as *Auto-Login, wait*, then the RAS 2000 remains at the "**Press <enter> to continue**" prompt until someone does this and you avoid the retries and time outs.
Auto-Login	Sometimes, you want to configure a port so that the RAS 2000 automatically starts a connection without prompting for a login. If the port is configured for *Auto-Login*, you must also specify a user name for this port (*myname*, in the above example). The port would behave exactly as a *Login-by-Port* but instead of sending a login prompt, he assumes that the specified user has successfully logged in, and starts up his connections accordingly. When the session is over, the RAS 2000 will (after waiting for *carrier* when appropriate) restart the sessions again.
Printer	This configuration is similar to *Reverse-TC* , except that a port configured as a *printer* can also accept connections from **rcp** and **rsh cat** clients on your network. There are other differences as well: these are discussed in chapter 18, ***Reverse TCP and Printing*** in the *RAS 2000 Software Configuration Guide*.
Reverse-TCP	When a port is configured as *Reverse-TCP*, the RAS 2000 accepts a TCP connection from some other host on the network. Data received from that host is sent out the serial port, and data received from the serial port is sent to the host. This is a common method of supporting printers and other "non-login" serial devices. In chapter 18, ***Reverse TCP and Printing***, of the *RAS 2000 Software Configuration Guide* this is discussed in much more detail, and on **page 86** (of this same manual), where a port's *group* and *tcp* parameters are discussed.
Out-bound PPP or SLIP Connection	This configuration supports outbound PPP/SLIP/CSLIP links. The RAS 2000 brings these links up automatically when it tries to route a packet to a network that it knows to be on the other side of one of these links. This process is discussed in more detail in chapter 11, ***Remote Network Configuration***, of the RAS 2000 Software Configuration Guide. Note that this port type supports only *dial-out* connections. To support clients who are dialing into the RAS 2000, you configure the port as *Login-by-Port*.

Table 7 Port Configuration Parameters

Parameters	Definition
Login by Port/TCP	This is a combination of *Login-by-Port* and *Reverse-TC* , and is designed to support bidirectional operation of a modem. You must configure the port as *modem enabled*, because the RAS 2000 uses *carrier* (DCD) to sense incoming calls and determine whether there has been a disconnection. When the port is idle and there is no incoming call, the RAS 2000 accepts TCP connections for this port from hosts on the net, just like *Reverse-TC* . If a connection is established, the client can access the modem, send dialing commands, and connect to other systems. Anyone trying to dial in gets a busy signal because the modem is off-hook. If an incoming call comes in first, the port sends out a login prompt, like *Login-by-Port*, and as long as the incoming call is connected, the RAS 2000 refuses or defers TCP connections from the network for that port. This is explained in more detail in chapter 18, ***Reverse TCP and Printing,*** of the *RAS 2000 Software Configuration Guide.*
Remote Serial Port	When a port is configured as a remote serial port, it can establish communications with another host's driver. Once a host's driver establishes communication with a RAS 2000 PowerRack, it attaches to one or more remote serial ports (RSP) and presents them as "logical" serial ports to the host operating system. In the case of NT, these appear as COM ports.
User Name	Name of port user
Group	There are 16 groups of ports, numbered 0 to 15. Any port can belong to any group, or to no group at all. When something tries to start a reverse-TCP connection to the RAS 2000, it can specify a particular port or a particular port group. When a port group is specified, the first available port in the group is used. (see chapter 18, ***Reverse TCP and Printing,*** of the *RAS 2000 Software Configuration Guide*). A port group number can also be specified in a *Remote Profile* for an outbound PPP/SLIP/CSLIP interface (see chapter 11, ***Remote Network Configuration*** of the *RAS 2000 Software Configuration Guide*). A port is configured as *Reverse-TCP* or *Login-by-Port/TCP* cannot be a member of the same group as a port configured as *Printer* or which uses *IntelliPrint*. This is because the first types suppress output processing, while the others perform it. If both types were members of the same group, the results might depend on which printer happened to be available.

Table 7 Port Configuration Parameters

Parameters	Definition
Local Term Type	This setting defines the terminal characteristics that will be used when the RAS 2000's menu interface is running on this port. This also defines the default terminal name that is sent when you telnet or rlogin from this port to a host on your network. This default value may be overridden by other settings, however. Because this information is used by the menus, the RAS 2000 needs to understand the terminal characteristics that each terminal name represents. For that reason, there are a limited number of these supported. Your choices are: unknown wyse30 xterm ansi wyse50 uterm0 uterm1 vt100 wyse60 uterm2 uterm3 The last four terminal types are user-definable. If your terminal does not emulate one of the defined terminals, a section starting on **page 90** of the *RAS 2000 Software Configuration Guide* explains how you can store your terminal's definitions under one of these four terminal types.
Remote Term Type	If you enter a name here, then by default it IS sent when you rlogin or telnet to a host, instead of using the one given for the *Local terminal name*. Since the RAS 2000 does not need to know what this name actually means, it can be any name that is understood by the login host. The telnet and rlogin commands also support command-line arguments which, if used, can override these default terminal-types. If there is no command-line argument, the *remote term type* is used, and if no *remote term type* is defined, then the *local term type* is sent. The telnet and rlogin commands are described in chapter 16, ***Connections*****,** of the *RAS 2000 Software Configuration Guide*.
Wait for CD (Modem)	The modem port waits for the modem to assert carrier (**DCD**), if the port was configured as a *modem port*. When a *modem port* is connected to a local terminal, the port' **DCD** is usually wired to the terminal' **DTR** or **RTS** (whichever is *not* being used for flow-control). In that case, the RAS 2000 would be waiting not for an incoming call, but for someone to turn the terminal on. If this is a non-modem port, it is assumed that carrier is present. The RAS 2000 waits for 1 second before continuing, after carrier (**DCD**) is detected by a *modem port*.
Wait for Input	This allows any attached modem or device to stabilize, before an attempt to send data to it. For example, there are some modems which assert carrier before coming out of command mode. Data intended for transmission to the remote modem may be interpreted as a command. After this 1 second delay, any data that might have been received so far is flushed before the preamble or login prompt is sent.

Table 7 Port Configuration Parameters

Parameters	Definition
Dialup Script	This is used by ports configured for outbound PPP/SLIP/CSLIP links (see **page 71** of the *RAS 2000 Softwae Configuration Guide*). It specifies the commands that have to be sent to the modem so it dials and establishes a connection and allows the RAS 2000 to wait for particular responses. Different modems may require different dialer scripts; that is why the dialer script is stored on a per-port basis, while the *login script* (which depends on the particular target of the call) is identified in the *remote profile*. Dialer scripts are discussed in chapter 11of the *RAS 2000 Softwae Configuration Guide*, ***Remote Network Configuration***.
Modem Init String	This setting is used by ports that are configured for terminals or dial-in connections (see **page 67** of the *RAS 2000 Software Configuration Guide*). It defines a string of commands which the RAS 2000 transmits to the modem before it waits for the next incoming call. This is not always required. Some modems can be configured ahead of time and never seem to lose their settings. In chapter 6, ***Configuring Modems***, of the *RAS 2000 Softwae Configuration Guide* examples are given of command strings you might want to send, and other ways to configure a modem. When does the string get sent: some user logs off, RAS 2000 drops **DTR** to hang up the line, waits a second, raises **DTR**, *sends the initialization string*, waits for modem to assert **DCD**, call comes in, next fellow logs in, works, logs off—and it starts all over again.
Baud Rate	This sets the line speed at which data is transmitted and received. In addition, you can define custom rates by setting up an IntelliSet profile and assigning that profile to a port. By using IntelliSet, you can also specify a *split baud-rate* where the port transmits at one speed, and receives at another. When line speeds and other parameters are defined using IntelliSet, those values override the ones chosen here. For more details, see chapter 13, ***IntelliFeatures,*** in the *RAS 2000 Software Configuration Guide.* 50 150 1200 3600 19200 64000 230.4k 75 200 1800 4800 38400 76800 307.2k 110 300 2000 7200 56000 115200 460.8k 134.5 600 2400 9600 57600 153.6k 921.6k
Data Bits	How many data bits per character. Selections are 8, 7, 6, and 5.
Parity	This controls the parity bit sent with each characters: *parity* must be one of the following: odd, even, space, mark, and none.
Stop Bits	This controls the number of stop bits that are transmitted after each character. Choices are**1**, **1.5**, or **2** bits. One stop bit is generally sufficient except when you are connecting to devices that are very old, very slow, or very unusual. This has no effect on the receiver. One stop bit is always sufficient.
Auto-sense PPP	This controls whether this port senses a PPP packet and automatically starts up a PPP link.
Input Flow Control	Since the RAS 2000 can receive data from a serial device, it must be configured to signal when its buffers start to fill up so that the serial device stops sending data for a while. The selections are none, XOFF, RTS, and XOFF & RTS.

Table 7 Port Configuration Parameters

Parameters	Definition
Input Translations	Any input processing specified here affects the port's operation when it is accepting line-based input, such as at the command prompt, or when in telnet command mode. At other times, the individual applications (telnet, rlogin, menu, etc.) force the input processing to an appropriate setting. Selections are CR to NL (carriage return to newline), none (no translations), and NL to CR (newline to carriage return).
Output Flow Control	Sometimes the serial port on the RAS 2000 is the sender, and it must avoid overrunning the terminal, modem, or printer to which it is attached. This is called *output flow control*. The *output flow control* you choose needs to match the *input flow control* of the device you are sending data to, and vice-versa. The selections are XON, None, XANY, CTS, XON & CTS, and XANY & CTS.
Output Translations	Any output processing specified here affects the operation of the port *only* when it is configured as a printer . The selections are NL to CR+NL, None, CR to NL, Strip CR, CR | NL to CR+NL, and NL to CR+NL.
Output Expand Tabs	With this setting, the port will translate ascii **tab** characters to a sequence of spaces sufficient to achieve tab stops at 8-character intervals. This tab setting corresponds to the traditional tab processing performed on UNIX systems and is useful when printing output from a UNIX system using tools that expect this processing to be performed "downstream". If this parameter is set to *No* (or *disabled* using the command), then **tab** characters are sent unchanged.
TCP Mode	A protocol which establishes a reliable connection between two processes, generally on separate computers. Higher-level protocols like telnet and rlogin rely on TCP to ensure that data is not lost in transmission and that data is not sent faster than it can be processed. The settings are Normal, CRNL ->CR, and Raw. Normally, a reverse-TCP connection uses telnet protocol. Telnet server implementations differ in their treatment of carriage-return (CR) and linefeed (or new-line, NL) characters. With some, if a CR-NL pair is received from the network, the two characters will be output. That is what the *normal* option does. With other telnet servers, if a CR-NL pair is received, the CR is sent but the NL is ignored. This ia the *CRNL->CR* option. These two options are provided for maximum compatibility. The third option, *Raw*, causes the Reverse-TCP connection on that port to not use telnet protocol at all. Instead, the data received over the TCP connection is sent to the port exactly as received, and vice-versa. This is provided for compatibility with other vendors' products, as well as providing an easy-to-use interface for special applications.
INTR Character	This defines the *interrupt ke* . Use this key to quickly terminate commands before they have finished. In this example **^c** represents *control-c*.
ERASE Character	This defines the character used to backspace a single character and erase it.
EOF Character	This defines the character used to denote the end of file character.
IntelliView Profile	This specifies the name of the IntelliView profile you want to apply to this port. It can be any profile you have already created.

Table 7 Port Configuration Parameters

Parameters	Definition
IntelliPrint Profile	This specifies the name of the IntelliPrint profile you want to apply to this portIt can be any profile you have already created.
IntelliSet Profile	This specifies the name of the IntelliSet profile you want to apply to this port.

IntelliSet Profiles

Selecting IntelliSet Profiles displays the following screen.

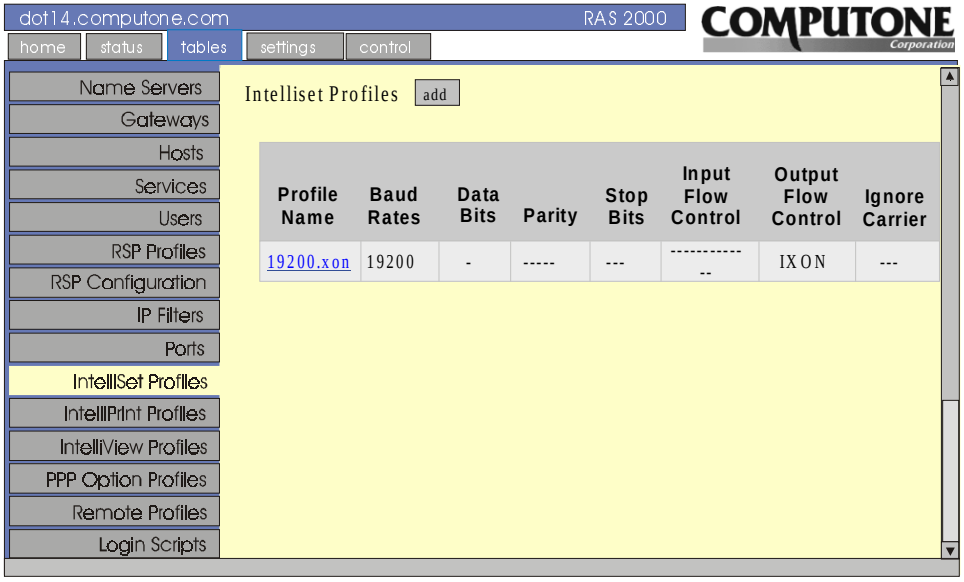

Figure 26 *IntelliSet Profiles* Screen

IntelliSet profiles include a mixed bag of specifications that can be thought of as an extension of port configuration. In fact, there are many parameters in common. The difference is that features you specify with IntelliSet *override* the settings in port configuration.

IntelliSet parameters also resist any attempts by applications to change them. For example, a telnet session normally puts a serial port into raw mode, disabling any output processing that might have been specified under port configuration. If, however, this had been specified in an IntelliSet profile, the output processing

continues. To change the IntelliSet Profile parameters, click on the profile name. The following screen is displayed.

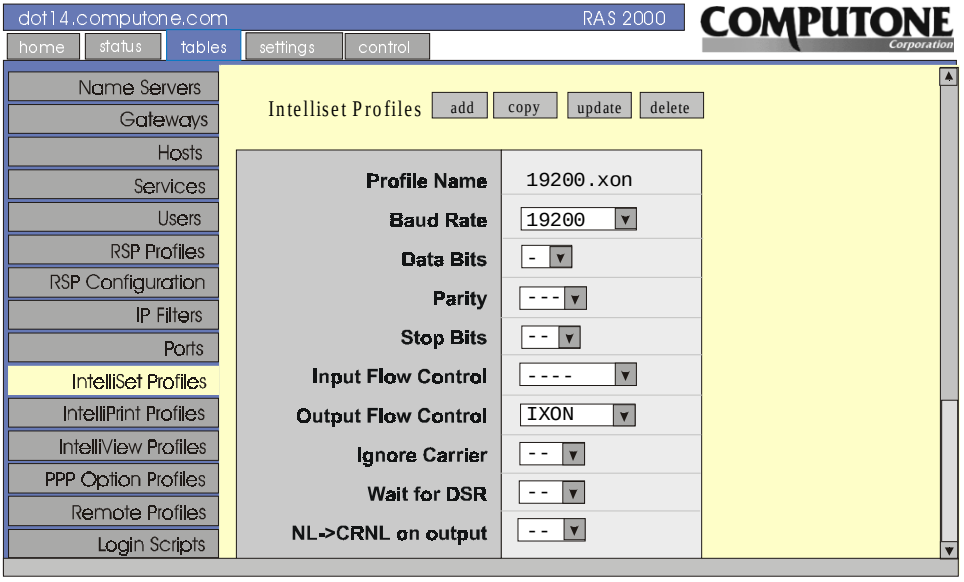

Figure 27 IntelliSet Configuration Screen

The IntelliSet Configuration Form is shown in Figure 2 7 , and illustrates profile *19200.xon* which is supplied with the factory defaults. Except for the profile name and the two custom baud rates, all the other fields are pick-lists and offer many of the same selections as the Port Configuration screen. In this example, many of the input areas contain dashes. These indicate that the corresponding parameter is not specified by this IntelliSet profile. When a value appears instead of dashes, then that parameter is specified. When a parameter is specified in an IntelliSet profile, that value overrides anything in the port configuration and it cannot be changed by whatever application might be running (telnet, rlogin, etc.).

Table 8 IntelliSet Parameter Selections

Parameter	Description
Profile Name	The name of the IntelliSet Profile under configuration.
Baud Rate	The following baud rates are available: ------ 50 150 1200 3600 19200 64000 75 200 1800 4800 38400 76800 110 300 2000 7200 56000 115200 134.5 600 2400 9600 57600
Data Bits	The following data bits selections are available: ---, 5, 6, 7,and 8.
Stop Bits	The following stop bit selections are available: --, 1, 1.5, 2.
Input Flow Control	The selections are -----, None, IXOFF, RTS, RTS+IXOFF. *IXOFF* indicates that the RAS 2000 should send an XOFF character when its receive buffers become nearly full, and send an XON character when they again have room for more data. *RTS* indicates that the RAS 2000 should drop the RTS signal when the buffers are nearly full, and raise it when they have room. RTS+IXOFF indicate that a combination of actions are taken.
Output Flow Control	The selections are -----, None, IXON, XANY, CTS, IXON+CTS, IXANY+CTS. The *Outflow* parameter allows you to specify and lock the output flow control. When the RAS 2000 is sending data to a device, and that device cannot process data quickly enough, it must signal the RAS 2000 in some way to tell it to stop transmitting. Output flow control specifies what conditions cause the RAS 2000 to stop sending data. The dash indicates that this IntelliSet profile will not affect output flow control. *CTS* indicates that the RAS 2000 will not transmit unless the CTS input is asserted. *IXON* indicates that the RAS 2000 should stop transmitting when it receives an XOFF character and resume when it receives an XON. *IXANY* is the same, except that after the XOFF character has disabled transmission, receiving *any* character (not just an XON) will re-start it.
Ignore Carrier	The selections are: -----,Yes, and No. The *Ignore Carrier* parameter specifies whether the port will be treated as a *modem* port. Modem ports are affected by the DCD (carrier detect) signal, while non-modem ports ignore it. The dash indicates that this IntelliSet profile will have no effect on whether the port is a modem or non-modem port.

Table 8 IntelliSet Parameter Selections

Parameter	Description
Wait for DSR	The selections are: ----, Yes, and No. The *data-terminal-ready* (**DTR**) signal from the RAS 2000 is asserted when the port is opened, dropped when the port is closed, and stays dropped if the port is disabled or the RAS 2000 is off. If you don't want your modems answering the phones when the RAS 2000 Communications Server is turned off, be sure **DTR** is connected and make sure that the modem has not been configured to ignore **DTR**. If you are using CTS/RTS flow control (see chapter 5, ***Output Flow Control Options,*** in the *RAS 2000 Software Configration Guide*) be sure to connect both of these signals as well. The RAS 2000 software does not require the use of the *data-set-ready* (**DSR**) and *ring-indicator* (**RI**) signals, so these do not need to be connected.
NL->CRNL on Output	The selections are ----, Yes, and No. These parameters allow you to specify and lock output processing on this port. Carriage returns can be inserted before linefeeds to prevent *barber-pole* output as show on **page 337** of the *RAS 2000 Software Configuration Guide*. The dash indicates that this IntelliSet profile should not affect output processing.
Expand Tabs	The selections are ----, Yes, and No. Tab expansion can be helpful if your output contains tab characters and your terminal doesn't understand tabs.
Use RTS in Half Duplex	The selections are ----, Yes, and No. *RTS* indicates that the RAS 2000 should drop the RTS signal when the buffers are nearly full, and raise it when they have room.

IntelliPrint Profiles

Selecting IntelliPrint Profiles displays the following screen.

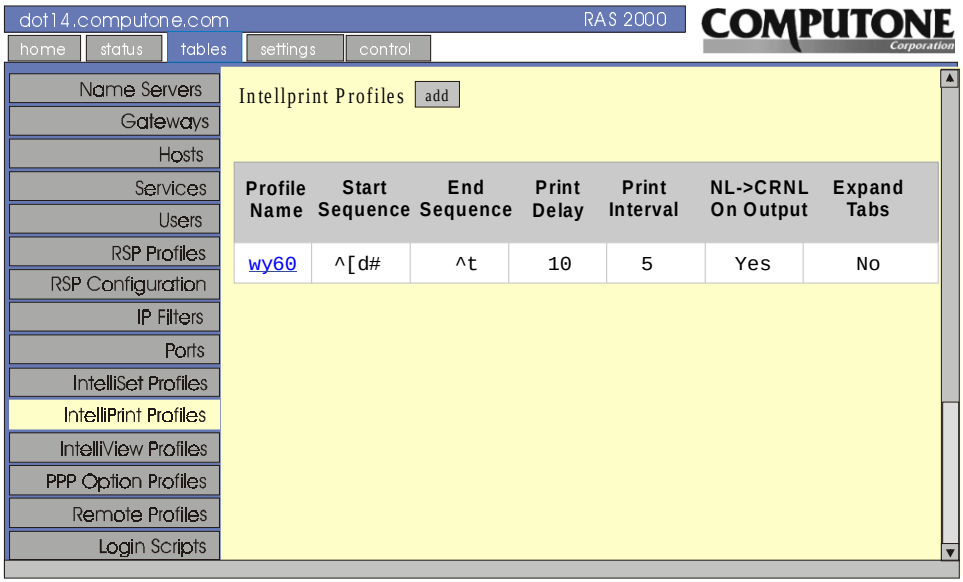

Figure 28 ***IntelliPrint Profiles*** **Screen**

To support IntelliPrint, your terminal must have an AUX port, and support commands that route subsequent data to that port instead of to the display, and also route subsequent data to the display and no longer to the Aux port. These commands are special data sequences that are sent from the RAS 2000 before and after any data directed to the printer. An IntelliPrint profile contains these sequences, as well as other information for setting the relative priorities between data for the printer and data for display

When you have associated an IntelliPrint profile with a serial port, you can configure your network hosts to send data directly to that printer, independently of what is happening on the terminal. For example, your terminal could be logged into a host and running an application. While in the application, you decide you want to print a report. The report is sent to your system's print spooler, which you have configured to send output to the printer attached to your terminal. While the output is printing, you are still able to use your terminal.

To configure the IntelliPrint profile, click on the profile name. The following screen is displayed.

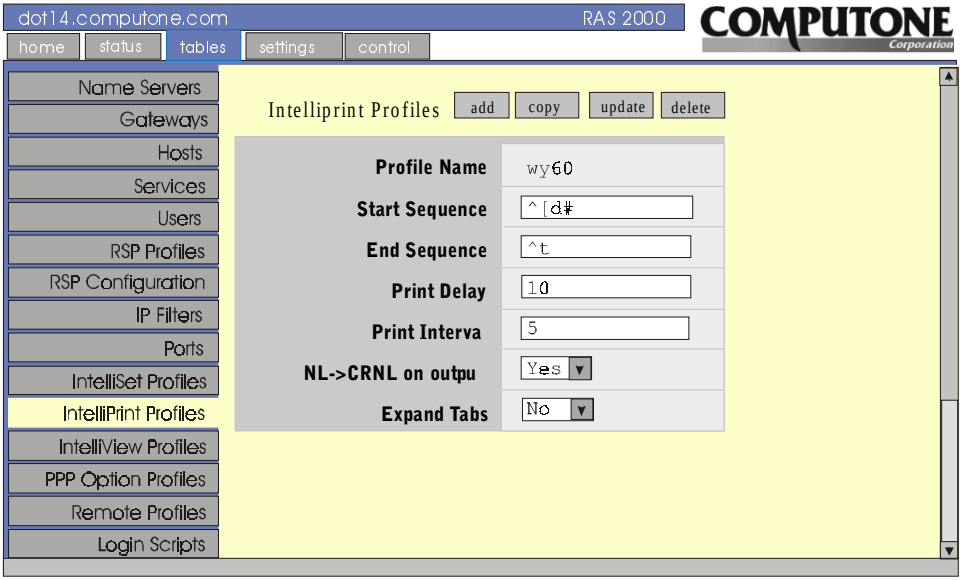

Figure 29 ***IntelliPrint Profiles*** **Configuration Screen**

The IntelliPrint Profile parameters are defined as follows:

Table 9 IntelliSet Parameter Selections

Parameter	Description
Profile Name	The name of the IntelliSet Profile under configuration.
Start Sequence	To make a terminal send data to its Auxiliary port, the RAS 2000 must first send a command saying *"send all subsequent data to the aux port, until I say otherwise"*. Not in so many words, of course, but some sort of special data sequence must be used. In this example, for the Wyse 60 the sequence consists of the three characters: **escape d #**. The escape character is represented here by the symbols **^[**. It is often necessary to represent the *escape* code and other unprintable bytes in IntelliFeatures profiles (refer to Table 5-4 on page 94 in the *RAS 2000 Software Configuration Guide* to help you remember how to do this). How do you find out what codes to use for your particular terminal? Consult your terminal's manual or programmers' guide. What is called the *start print sequence* is often called *start transparent print*. The print sequences to use are determined by which *terminal* you are using, not by your choice of printer. The printer never "sees" these sequences, they are interpreted by the terminal.

Table 9 IntelliSet Parameter Selections

Parameter	Description
End Sequence	When the RAS 2000 has been sending data to the printer and now wants to send data to be displayed again, it must send a command to the terminal. For a Wyse 60, the command to do this is the single character, **ctrl-t.** If you want to send data to a printer attached to your terminal, that data had better *not* contain the terminal's *end print* sequence. When the terminal sees such data, it will *not* send it blindly to the printer. It will rightly interpret it as a command and send the following data (that you had intended for the printer) to the display
Print Delay	The *Print Delay* tells how long the RAS 2000 must wait after any display output before it sends any data for the printer and the delay is measured in tenths of a second. Why would you want such a delay? Data for display often contains control sequences for cursor addressing, highlighting, and so on. These sequences consist of two or more bytes of data and most terminals get confused if such a sequence is interrupted by a command to start transparent printing. The delay ensures that all the bytes of any control sequence have a chance to be completely sent before a command to start printing is sent.
Print Interval	The *Print Interval* defines a delay to be inserted between successive blocks of print data and is measured in tenth-seconds. Why a delay between successive blocks of print data? To make the RAS 2000 sends its printer data more slowly. Most terminals can display faster than most printers can print.
NL->CRNL on Output	This defines whether the RAS 2000 adds Carriage-Returns before linefeeds and expand tabs in data to be sent to the printer. In most cases it is better to keep this option disabled and configure your host software to perform whatever processing is appropriate.
Expand Tabs	This defines whether the RAS 2000 adds expand tabs in data to be sent to the printer. In most cases it is better to keep this option disabled and configure your host software to perform whatever processing is appropriate.

IntelliView Profiles

Selecting IntelliView Profiles displays the following screen.

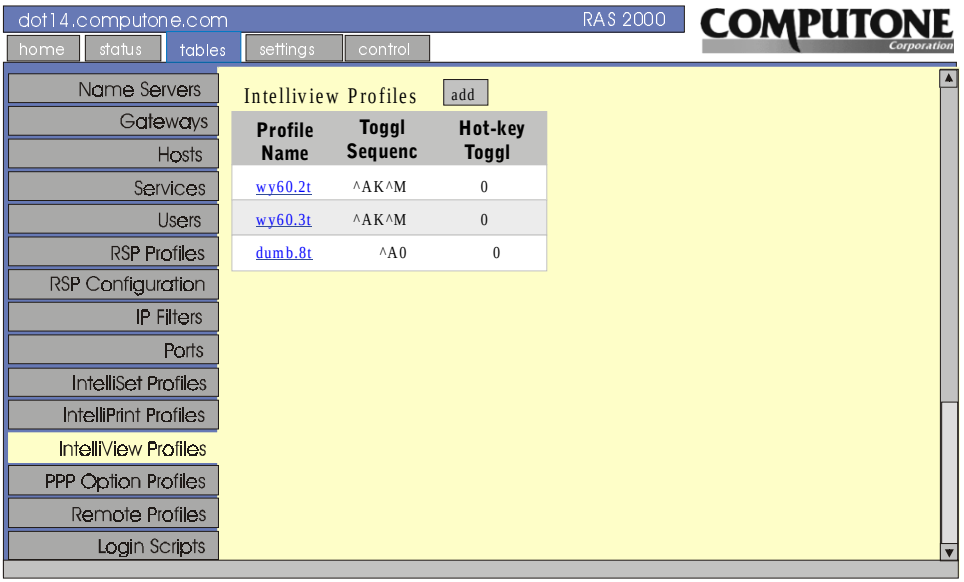

Figure 30 ***IntelliView Profiles*** **Screen**

IntelliView allows you to run a separate session on each of the terminal's display pages. On one page (or "screen") you might be logged into one host, while on another page you are logged into a different one. You would use designated function keys to flip between the two sessions. In order to use the terminal's function keys to flip between pages, the RAS 2000 needs to be told what codes these keys send. Generally this is not a problem, but there are some applications which send commands to the terminal to reprogram its function keys to specific values. If this includes a function key you want to use for IntelliView, you may find yourself unable to switch screens once you have entered that application. In that case, you need to investigate reconfiguring that application or using different function keys for IntelliView.

To support IntelliView, your terminal must be capable of maintaining two or more *pages* of display memory. Often, the number of pages supported depends on the display mode. For example, a Wyse 60 in 132x43 mode (Wy60 emulation) might support only a single *page*, where in 80x25 (Wy60 emulation) it might

support two pages, and in “Econ-80” mode it might support three or more pages. If you are not sure about your particular terminal, you should read its manual.

The WY60.2t IntelliView profile, shown in Figur e30 , supports a Wyse 60 terminal with two pages of screen memory. Three *hot keys* are defined, all of them function keys:

- Press **F12** to switch between screens. When you press it, the three-byte sequence **^A K ^M** is sent.
- Press **Ctrl-F1** to switch to screen 0 (the main screen). If your Wy60 is configured for 8-bit data, this key sends a single byte, octal value 200.
- Press **Ctrl-F2** to switch to screen 1. This key sends a single byte of octal value 201. If you are creating a new IntelliView profile, you enter the name on the first line. If you are modifying an existing one, the name is already there and you are not allowed to change it.

To edit one of the profiles, click on the profile name (wy60.2t is selected, for this example). The following screen is displayed.

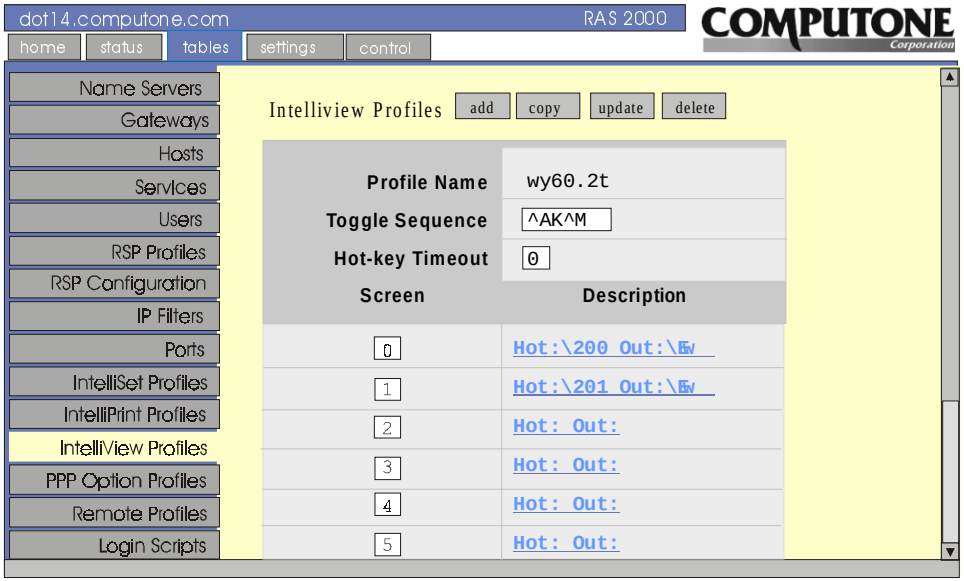

Figure 31 ***IntelliView Virtual Screens*** **Screen**

To set up 1 of the 8 screens, click on the name in the description column of the one you want to set up (for this example, the description for screen 0 is selected). The following screen is displayed.

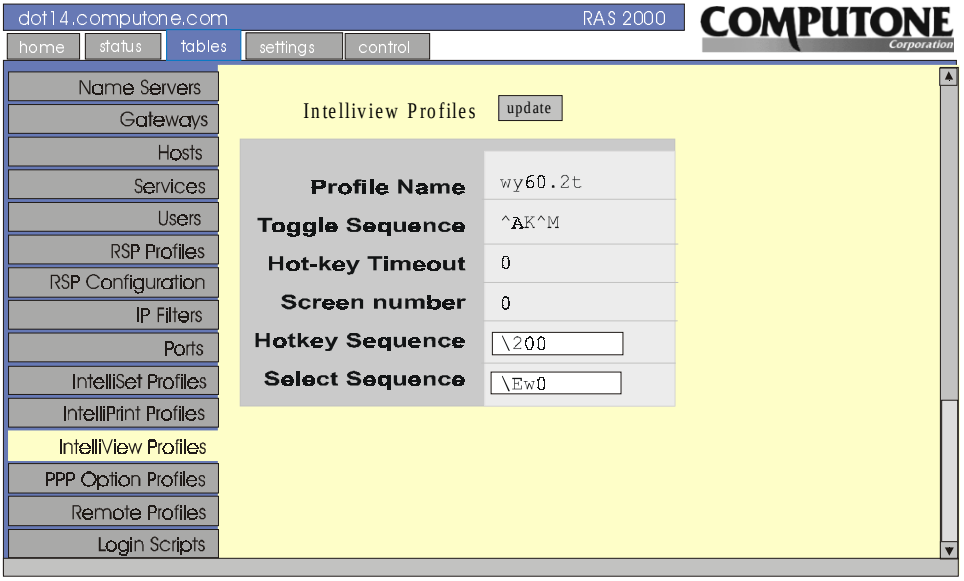

Figure 32 ***Hot Key Sequence*** **Screen**

On this screen you can enter or change the *Hotkey Sequence* and *Select Sequence* octal values.

PPP Option Profiles

Selecting PPP Option Profiles displays the following screen.

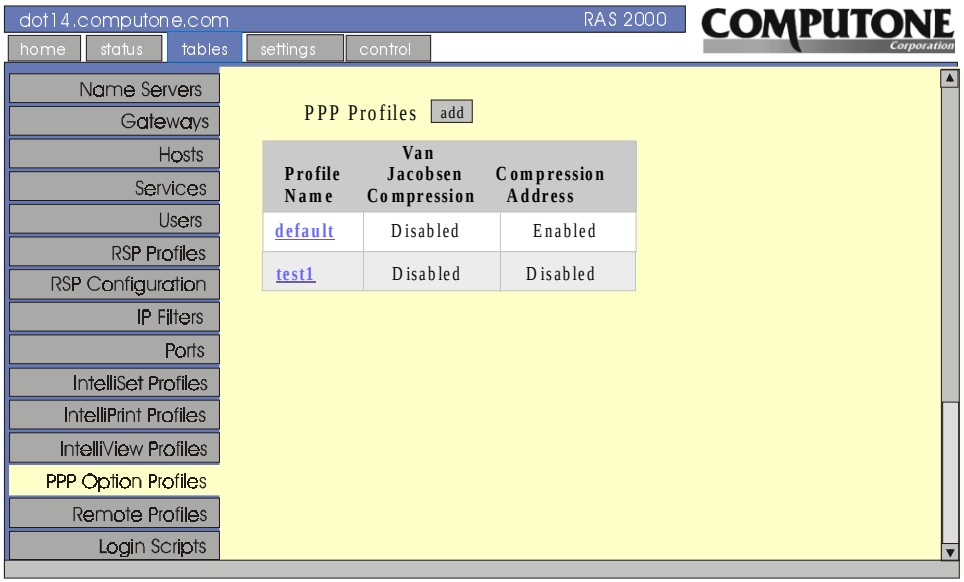

Figure 33 ***PPP Option Profiles*** **Screen**

PPP Option Profiles contain additional protocol options used in bringing up PPP and SLIP links. Options Profiles (sometimes called SLIP/PPP options) are used for storing configuration parameters that do not change very often. These parameters are also likely to be shared by a number of interfaces at a given site.

An Options Profile is created with a particular collection of settings and it i given a name. To assign these settings to a particular interface, you enter the name in that interface's Remote Profile. This reduces the number of separate parameters that an individual Remote Profile must contain.

At factory default there is a single options profile defined called **default**. To create a new profile, click on ***add.*** When the new remote profile is created (***test1***, in this example), this default options profile is assigned to it and remains until you change it.

To change PPP Option Profile ***test1***, click on it. The following screen is displayed.

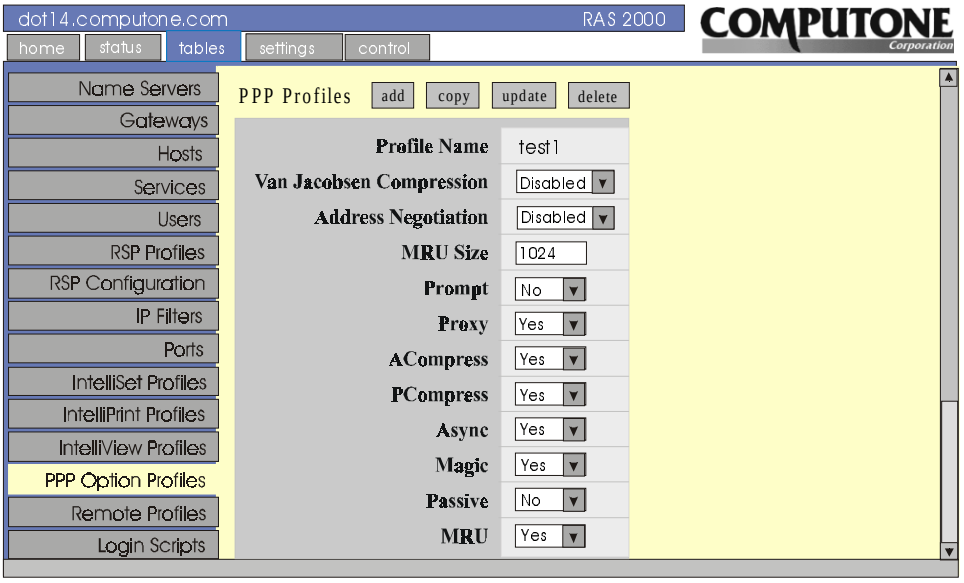

Figure 34 ***Changing a Profile*** **Screen**

Table 8 describes the parameter selections.

Table 10 PPP Option Profile Parameter Selections

Parameter	Description
Profile Name	The name of the PPP Option Profile under configuration.
Van Jacobson Compression	*Van Jacobsen (VJ) Compression* is a method of compressing TCP/IP headers in PPP or SLIP packets. With SLIP protocol, both sides must agree beforehand whether to use it. (SLIP with VJ Compression is called CSLIP). With PPP, the two sides can negotiate whether to use VJ Compression. On the RAS 2000, connections are designated PPP, SLIP, or CSLIP. Since CSLIP always uses VJ Compression, and SLIP never does, this option affects only PPP links. Set it to **Disabled** (the default) if you do not want to use VJ Compression or to **Enabled** if you do.

Table 10 PPP Option Profile Parameter Selections

Parameter	Description
Address Negotiation	Enabling *Address Negotiation* on inbound PPP connections allows the RAS 2000 to learn the caller's IP address and inform the caller of our IP address through the PPP negotiation process. While address negotiation can be enabled for outbound connections as well, the RAS 2000 needs to know the remote site's correct IP address ahead of time because it is the attempt to *access* that address which causes the RAS 2000 to bring up the interface.
MRU Size	The *MRU Size* (*Maximum Receive Unit Size*) represents the maximum number of bytes the RAS 2000 can receive in a single PPP packet. This is a partner to the MTU, or *Maximum Transmit Unit*, which is configured in the Remote Profile and defines the largest packet the RAS 2000 can *send*. Each side of the link has an MRU, usually constrained by internal buffer sizes and an MTU. The first step is making sure that one side's MTU is not greater than the other side's MRU. With PPP, this is done through *Maximum Receive Negotiation*. If *Maximum Receive Negotiation* (or *mru*) is **Yes**, (and assuming the remote side of the link is so configured) each side informs the other of its own MRU. If the recipient's MTU is larger, it temporarily reduces it accordingly. For SLIP and CSLIP connections, the effective MRU is always 1536 bytes. A large value is chosen because there is no mechanism, other than mutual agreement at configuration time, to agree on a smaller value.
Prompt	This applies to inbound SLIP and CSLIP connections. When set to **Yes**, the RAS 2000 prompts the user to enter his IP address. After the address is entered, the link is brought up using that IP address as the "Remote Address". This facilitates multiple sites being able to use a single interface at different times. This option does not apply to PPP connections which are able to use PPP address negotiation for this purpose. This option is also not required when remote dial-in users are configured on a RADIUS server because each user's IP address can be stored in the RADIUS server's user database (see chapter 17, ***User Authentication using RADIUS*** in the *RAS 2000 Software Configuration Guide*).
Proxy	If the *Proxy* option is set to **Yes** (the default), the RAS 2000 responds to ARP requests for the remote IP address on this interface, as long as the link is up. For example, suppose the RAS 2000's IP address (on the local Ethernet network) was 160.77.99.30. Suppose the remote IP address (the host at the other end of this PPP link) was 160.77.99.17. If a host on the RAS 2000's local network wanted to access this remote host it would think from the IP address that it is on the local network. So, it would perform an ARP request to learn the Ethernet Address. The RAS 2000 replies giving its *own* Ethernet address and enabling it to receive packets destined for that host. This is explained more fully under ***Proxy ARP*** on page 189 of the *RAS 2000 Software Configuration Guide*. If the option is set to **No**, Proxy ARP is not per-formed.
ACompress	The *Address Compression* controls the local compression of address and control fields in the PPP header. Specify **Yes** (the default) to compress these fields, or **No** to leave them uncompressed.
PCompress	The *Protocol (Field) Compression* controls the local compression of the protocol field in the PPP header. Specify **Yes** (the default) to compress it, or **No** to leave it uncompressed.

Table 10 PPP Option Profile Parameter Selections

Parameter	Description
Async	The *Async (Map)* is used by PPP to prevent certain control characters (such as XON and XOFF) from occurring in the data stream. The map indicates which characters are proscribed. Specify **Yes** (the default) to allow the RAS 2000 to negotiate this map with the remote system. Specify **No** to force the RAS 2000 to use the ASYNC Map specified in the Remote Profile.
Magic	The *Magic (number)* is a arbitrary 32-bit number which is randomly chosen by each side of a PPP link. During negotiation, each side sends the other its*magic number.* It would be unusual for two different hosts to randomly choose the same random number, so if the magic number received is the same as our own, it is assumed that something has gone wrong (perhaps the modem is running in loop-back mode) and the RAS 2000 must be talking to itself. Since this is a bad thing, the RAS 2000 drops the connection (hang up the modem, etc.). Choose **Yes** (the default) if you want the RAS 2000 to check the magic numbers, or**No** if you want to ignore it.
Passive	*Passive (Mode)* only affects an outbound PPP connection on the RAS 2000. Specify **No** (the default) if it should initiate the PPP negotiations, or **Yes** if it should passively wait for the other side to do so.This option has no effect on SLIP or CSLIP connections because these protocols do not involve negotiations.
MRU	Each side of the link has an MRU, usually constrained by internal buffer sizes and an MTU. The first step to harmony is making sure that one side's MTU is not greater than the other side's MRU. With PPP, this is done through *Maximum Receive Negotiation.* If *Maximum Receive Negotiation* (or *mru*) is **Yes**, (and assuming the remote side of the link is so configured) each side informs the other of its own MRU. If the recipient's MTU is larger, it temporarily reduces it accordingly
Bring Up	The *Bring Up (Slip Link Immediately)* option applies to outbound SLIP and CSLIP connections. By default this option is set to **No** and the RAS 2000 attempts to bring up the outbound link when it is first required to route network traffic to the IP address at the other end. If you choose **Yes**, then the RAS 2000 attempts to bring up the line immediately on start-up. Furthermore, if the link goes down (because of a modem disconnect, for example) the RAS 2000 attempts to bring it up immediately.

Remote Profiles

Selecting Remote Profiles displays the following screen.

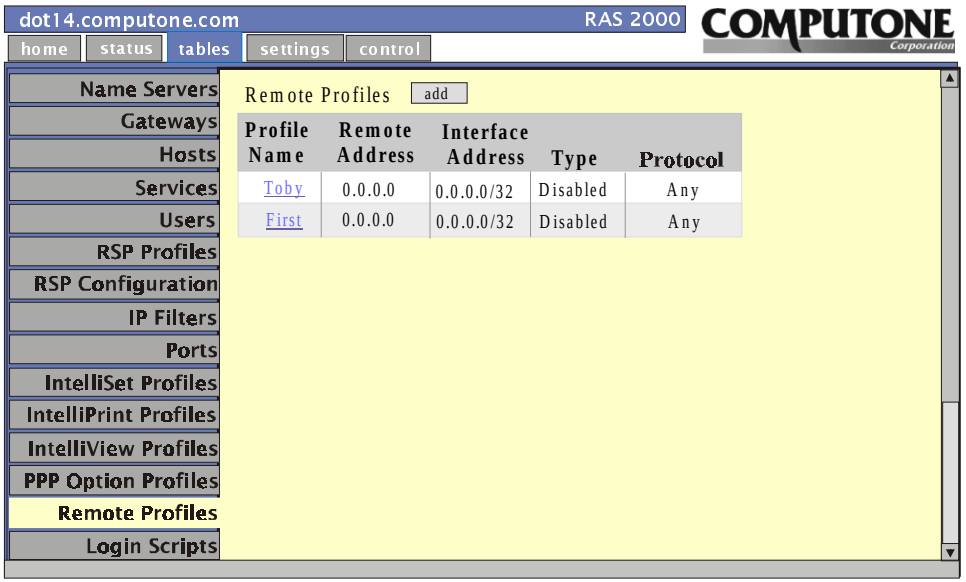

Figure 35 ***Remote Profiles*** **Screen**

Remote Profiles contain basic information about the interface: whether it is inbound (i.e., dial-up) or outbound, whether it is dedicated to a particular serial port or user, which protocols may be used, the IP address of the remote site, and so on. To add a profile, click ***add*** then enter a name for the remote profile. Once the profile is created, click on the name to change the values (in this instance, ***Toby*** is used as an example).

The following screen is displayed.

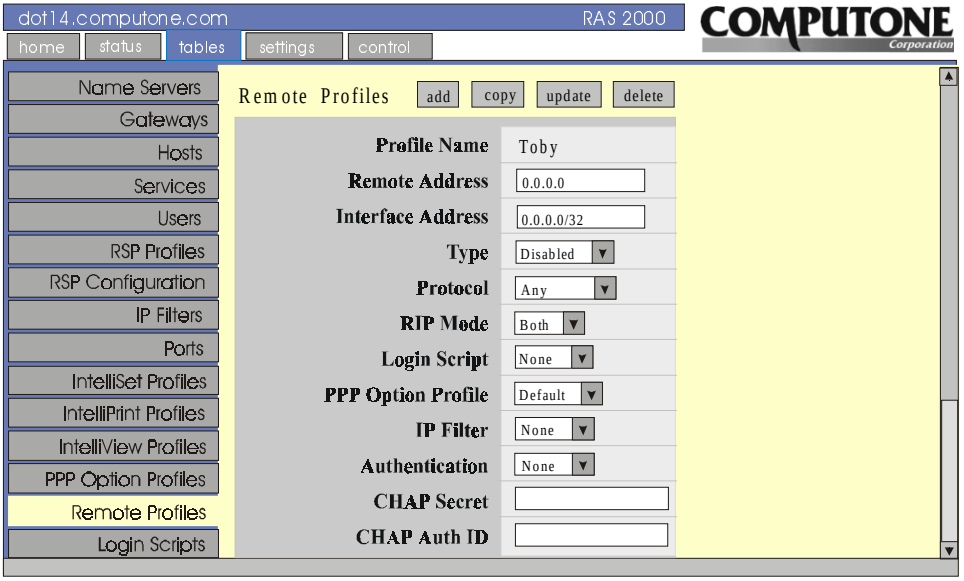

Figure 36 ***Remote Profiles Parameters*** **Screen**

Table 11 Explains parameter selections.

Table 11 Remote Profile Parameter Selections

Parameter	Description
Profile Name	The name of the PPP Option Profile under configuration.
Remote Address	The *Remote Address* is the IP address of the PPP, SLIP, or CSLIP interface at the other end of the link. When the link is brought up, this address is used unless a different one has been assigned through PPP address negotiation or information from the RADIUS user file. It is possible to leave this field set to **0.0.0.0**, in which case the correct IP address *must* be supplied by other means. Here are the rules: • For *Outbound* interfaces, the *Remote Address* must be set to the correct value, because it is the attempt to route to this address that brings up the link. This address cannot be subsequently changed by PPP address negotiation. • For *Inbound* interfaces, if the IP address of the remote interface is supplied from the RADIUS user database, or if it will be available from PPP address negotiation, the *Remote Address* in the Remote Profile can be left "open", i.e., set to **0.0.0.0**. • For *Inbound* interfaces, if the IP address of the remote interface will not be available by other means, the *Remote Address* must be set to some valid address. This technique is widely used by ISP providers to supply temporary IP addresses to dial-in users. The *Remote Addresses* you have assigned to various inbound interface comprise a pool of available IP addresses that are assigned dynamically as users dial in. See ***Assigning Remote Profiles*** on page 290 in the *RAS 2000 Software Configuration Guide* for an explanation of how this is done.
Interface Address	The *Interface Address* is the IP address of this end of the PPP, SLIP, or CSLIP link. If two RAS 2000s were connected via a PPP connection, each one's *Interface Address* would be the other's *Remote Address*. The *Interface Address* must be set to some valid address, but ***Interface Addresses* in different Remote Profiles are not required to be different.** In fact, it is common for Internet Providers to use the Ethernet's IP address (page 208 in the *RAS 2000 Software Configuration Guide*) for all interfaces. In some situations, you may need to use a different IP address. For instance, this could be an outbound interface to a site which expects you to have a particular IP address. This could happen if the remote site had another RAS 2000 and it were configured to assign a specific IP address for specific users.
Type	The *Interface Type* specifies whether this Remote Profile's interface will support inbound connections, or initiate outbound connections. The default value is **disabled**, so you must remember to set this value before the interface can be used. The selections are ***Disabled***, ***Inbound, Outbound***, or ***Both.***

Table 11 Remote Profile Parameter Selections

Parameter	Description
Protocol	The selections are Disabled, PPP, SLIP, CSLIP, and Any For an ***Outbound*** interface, you must specify either SLIP, CSLIP, or PPP, because the interface needs to know which protocol to use before it can bring up the link. For an ***Inbound*** interface, the desired protocol is learned directly from the user. If it is an NVRAM user it is configured specifically as either a SLIP, PPP, or CSLIP user. If it is a RADIUS user, similar information is stored in that database. For inbound interfaces, then, this is used to optionally restrict this Remote Profile's use, much like the *Serial Port* and *Dial-in User* are. If you want this Remote Profile to be available for any protocol, set the *Protocol* to **Any**. If you want to restrict it, specify SLIP, CSLIP, or PPP. The setting **Disabled** exists for compatibility with earlier versions of the RAS 2000. It is no longer needed because a Remote Profile's interface can be disabled by setting the *Interface Type* to **disabled** (see **page 281** in the *RAS 2000 Software Configuration Guide*.)
RIP Mode	RIP (Routing Information Protocol Mode) is used when the RAS 2000 needs to share routing information with other hosts. By *listening*, it learns routes from other hosts and by broadcasting or *sending*, it tells other hosts about the routes *it* knows.
Login Script	The *Login Script* is used only by ***Outbound*** interfaces. It allows the interface to log into the remote host. You must supply the name of one of the Login Scripts you defined earlier. See ***Login Scripts*** on page 259 of the *RAS 2000 Software Configuration Guide* for details. If this is left blank, then no login script is used. It is unusual for a dial-in connection to *not* require a login, so you will usually be specifying one.
IP Filter	This is the name of an IP filter you have defined. If this is blank, no IP filter is attached and the traffic through this interface is unrestricted. It is common for different interfaces to have different filters. Certain things might be allowed in the local Ethernet Network might not be allowed through any of the remote interfaces. Other traffic might be allowed over an outbound interface to a remote branch of the same business, that would not be allowed over the interface that goes to an Internet provider (ISP).
Authentication	The *Authentication Protocol* determines which authentication protocol can be used as part of the PPP negotiation. When set to **None** (default), no PPP authentication is performed. When set to **PAP**, PAP protocol is used for authentication, and when set to **CHAP**, CHAP protocol is used. CHAP is considered the more secure method because the authentication fields are sent encrypted.
CHAP Secret	This setting applies to ***Outbound*** as well as ***Inbound*** interfaces. For *Outbound* interfaces, the RAS 2000 supplies the information when requested by the site you are logging into. For *Inbound* interfaces, you request the information from the site that logged into yours.
CHAP Auth ID	This needs to be a matter of some pre-arrangement. If the RAS 2000 has the outbound interface and the remote site expects CHAP authentication, it had better enable it and know the correct CHAP name and secret to configure. So, this is not always a question of how you *want* to configure things.

Table 11 Remote Profile Parameter Selections

Parameter	Description
PPP User	When the user logs in, the intent may be to bring up a PPP or SLIP connection between the RAS 2000's local network and a client computer that has just dialed in. These users are called ***PPP users*** (although they may be using SLIP or CSLIP protocol instead). Sometimes these are called ***framed users*** because PPP, SLIP, and CSLIP are all protocols in which data is *framed* (i.e., separated into well-defined blocks marked by headers). When you configure a ***PPP user***, you need to provide networking information particular to this user so that routes between its network and yours will be set up correctly. A*framed user* exists only for the purpose of bringing up the PPP or SLIP link. Once the network has been extended by this connection, hosts on one side of this connection can connect to hosts on the other side. These connections may include rlogin and telnet sessions, and those users have no relationship to the *framed user* that caused the PPP/SLIP connection to be made.
Phone Number	The *Phone Number* is used only by ***Outbound*** interfaces with dial or login scripts. If the script contains the command **%p**, this phone number is inserted at that point. See ***Dial and Login Script Commands*** on page 256 of the RAS 2000 Software Configuration Guide. Being able to configure the phone number here conserves dial scripts.
ASYNC Map	The *Async Map* is used with PPP connections to prevent selected control characters from appearing in the data stream. When any proscribed character is sent, the interface substitutes a special character sequence. When any special character sequence is received, it is replaced with the corresponding original character. There are thirty-two ASCII control characters with decimal values 0 through 31. The *Async Map* is a 32-bit number written in hexadecimal notation. Each bit, starting with the **rightmost**, corresponds to different control character: A map of **00000001** would represent the character value 0 (ASCII **NULL**) and a map of **80000000** would represent the character value 31 (ASCII **US**). Remembering that each hexadecimal "digit" corresponds to four bits, consider the example **000a0000**. The first sixteen bits (from the right!) are clearly zero. Next is a hexadecimal **a**, which in binary is **1010**. The ones correspond to bits 17 and 19, and the rest of the bits are clearly zero. So, this mask traps the control characters with decimal values 17 and 19, which are the XON and XOFF characters.
MTU	The *Maximum Transmit Unit*, or *MTU*, is the maximum number of bytes that this interface sends in a single packet. For SLIP and CSLIP connections, it is important that this number not be larger than the other end's MRU, the largest packet it is prepared to *receive*. This must be done by prior configuration, because SLIP and CSLIP do not negotiate options.For PPP connections, this value can be negotiated. If Maximu Receive Negotiation has been enabled in the associated Options Profile, (see **page 268** of the *RAS 2000 Software Configuration Guide*) each side informs the other of its MRU, so that the other side can reduce *its* MTU (for this connection) accordingly.

Table 11 Remote Profile Parameter Selections

Parameter	Description
Idle Timeout	The *Idle* (*Inactivity) Timeout* is the maximum number of seconds the interface allows the connection to remain established when there is no data passing through it. If a connection remains inactive for longer than this, the RAS 2000 drops the connection — closes the serial port, drops DTR, and hangs up the modem. This timeout applies to both inbound and outbound interfaces. If the *Idle Timeout* is set to **0**, there is no timeout. A link could remain inactive forever unless shut down by other means (e.g., if the other side had an inactivity timer).
Port Group	The *Serial Port* and *Group* settings are used differently for inbound and outbound interfaces. An ***Outbound*** interface needs to know which serial port to use when it initiates the connection. Usually there will be a particular phone line and modem dedicated to this, so you provide the serial port number the modem is on. To allow more than one alternative, set the (*Serial) Port* to **Any** and enter a *group number* for the *group* instead. If a serial port is to be used for an outbound connection, it must be configured as an*Outbound Connection*, as shown on **page 71** of the *RAS 2000 Software Configuration Guide*. To configure serial ports with *group numbers*, assign this group number to all of the ports you want to be members. This is explained during serial port configuration, on **page 86** of the *RAS 2000 Software Configuration Guide*. An ***Inbound*** interface does not need to be told which serial port to use. The connection has already been initiated over one of them, the one you're using! For inbound connections, the *Serial Port* is used to restrict the use of this interface to a particular serial port. If you want this interface to be used regardless of the serial port, set its *Serial Port* to **Any** (the default). There are two cases to restrict an interface to a serial port. One is to make that port "special". Someone who dials into *that* phone number gets into *that* serial port which has its own dedicated Remote Profile with its own special settings. Ironically, the other case occurs when everything else is the same. In other words, a collection of Remote Profiles may be so configured that any connection may share them. In that case tying a particular interface to a particular port is not restrictive and it simplifies administration. If you are not covered by one of these cases, it is safer for you to keep the *Serial Port* set to **Any**. The *Group* is ignored by inbound interfaces.
Redial Delay	If an ***Outbound*** interface fails to bring up a connection, it waits this minimum amount of time before attempting again. This setting is ignored for *Inbound* connections. Enter a value in seconds.

Login Scripts

Selecting Login Scripts displays the following screen.

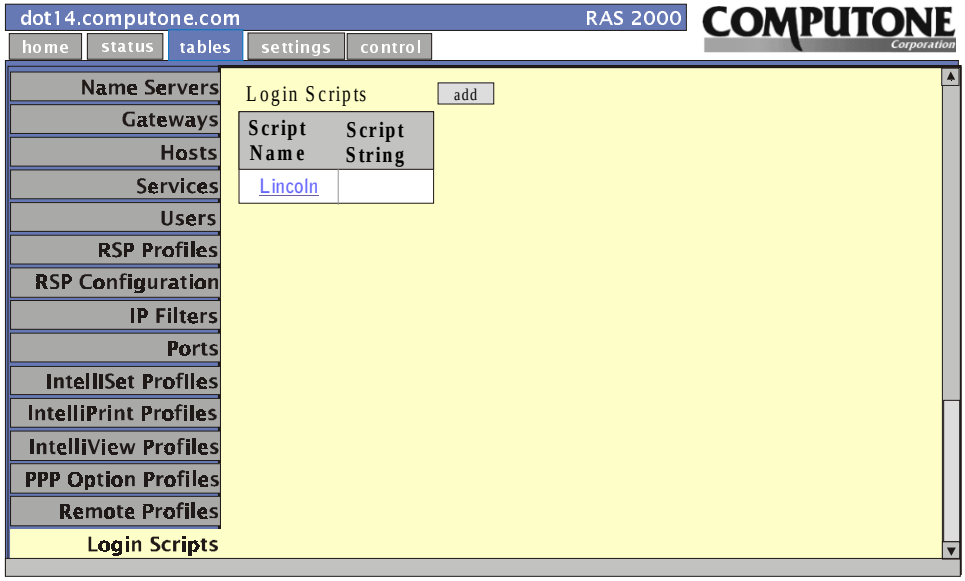

Figure 37 ***Login Scripts*** **Screen**

When the RAS 2000 starts to bring up an outbound PPP or SLIP connection, it first uses the dialer script to make the modem dial the remote site. At the remote site, a modem answers the call and the host computer may be configured to issue a login or password prompt. Then, it brings up its side of the link (or hangs up) based on what user name and password are provided.

Login scripts are run immediately after the dial scripts, allowing the RAS 2000 to provide automatically the necessary responses to a remote site's login, password, or other prompts. The nature of your login script is determined by the remote site you are contacting. Therefore, a login script is associated with a *remote profile*, not with a serial port (as are dial scripts).

Click on the script name and the following screen is displayed.

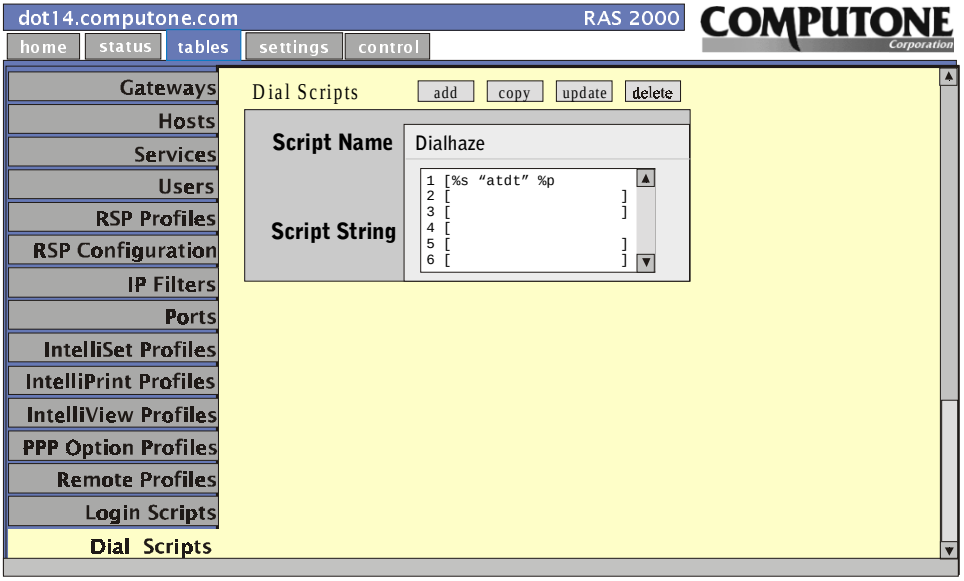

Figure 38 ***Login Script String*** **Screen**

The **Script Name** *(lincoln* in this example) is the unique name of this Login Script. During Remote Profile configuration, you can assign this to the profile's *Login Script* as described on **page 286** of the *RAS 2000 Software Configuration Guide*. The rest of the form contains the body of the script. Like Dial Scripts, it is composed of commands. The rules for forming these commands are the same a for Dial Scripts, found in **Table 11-2 on page 256** of the *RAS 2000 Software Configuration Guide.*

In the example shown in Figu re38 , the RAS 2000 waits up to thirty seconds for data matching **gin:** to come in. Presumably, this indicates the remote host has prompted us for our login name. Then, the RAS 2000 send **abraham**, our login name. Then, it waits for the password prompt, as indicated by the fragment **word:**. Finally, the RAS 2000 sends our password, **opensesme**. The login script is finished, and the RAS 2000 brings up our side of the link and so does the remote site.

Dial Scripts

Selecting Dial Scripts displays the following screen.

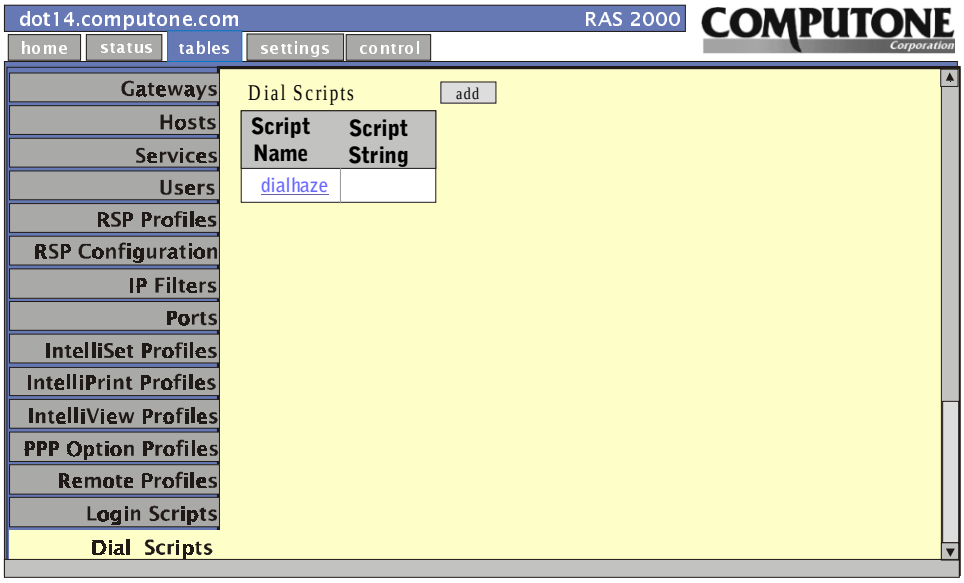

Figure 39 ***Dial Script*** **Screen**

The Dial Script screen is shown in Figure 3 9 . You can add new dial scripts or configure an existing dial script. To configure an existing dial script, click on the dial script name.

The following screen is displayed.

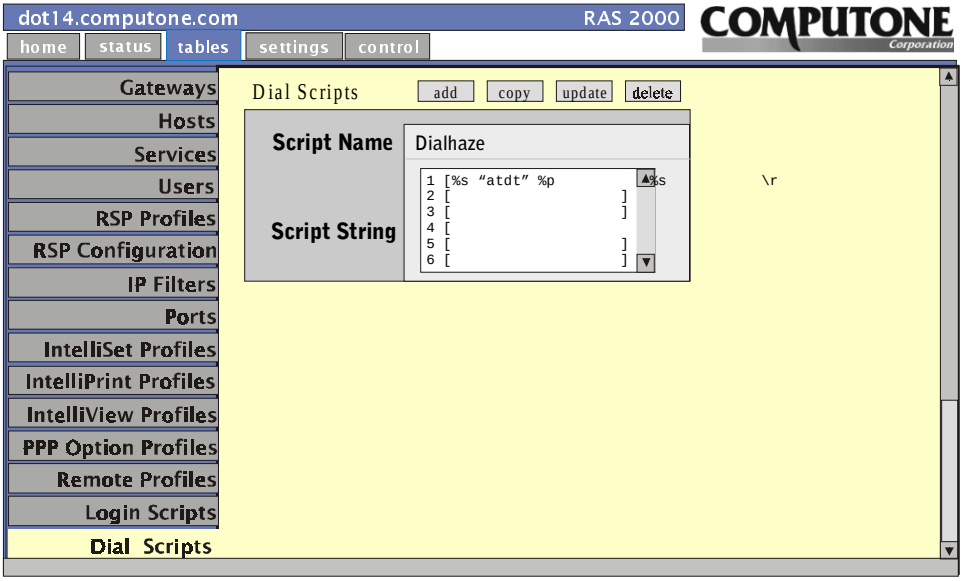

Figure 40 ***Dial Script Configuration Form*** **Screen**

The Dial Script Configuration Form is shown here in Figur e40 . Dial scripts define what needs to be sent to a modem so that it dials out and connects to another modem. The contents of your dial script depends somewhat on the modem. Since a given serial port is attached to a given modem, dial scripts are associated with serial ports. The **Script Name** *(dialhaze* in this example) is the unique name of this Dial Script. During serial port configuration, you assign thi to the port's *Dial Script.* The remainder of the form consists of six lines of forty-two columns each, but there is nothing special about the arrangement into row and columns. Before the script is run, trailing blanks are removed from each line and all the lines are run together.

Not only must the script send out strings of data, it sometimes needs to wait until certain responses are *received* before continuing. Therefore, a Dial Script is built not from simple data strings, but fro *script commands*. There may be several commands on a line, but a command must not be split across lines.

The following table shows how script commands are constructed:

Table 12 Dial and Login Script Commands

Command Definition / Examples	Description
%s `string` %s “ATDT5551212\r” %s “hello there” %s hellothere	Transmit the string to the serial port. If the string contains any spaces, enclose the entire string in quotes. Control characters can be represented using the codes in **Table 5-4 on page 94** of the *RAS 2000 Software Configuration Guide*; in this example, **\r** represents a carriage-return.
%w `time string` %w 10 connect\r %w 5 “carrier” **%w** `time` %w 10 **%w** `string` %w carrier\r %w “10” %w “1derful”	Wait until the specified string is received from the serial port, or the time (in seconds) elapse, whichever comes first. You may omit *either* the time or the string. • If the time is omitted the script will wait forever for the string. • If the string is omitted the script will wait the specified time unconditionally. • If a time and a string are *both* given, getting the string first is considered good. Timing out first is considered bad. **When** an interface is using a dial or login script to bring up a connection, if a *wait* command times out before the string is received, the connection attempt will be stopped and the line disconnected. If there is only one thing after the **%w** how does the script know whether it is supposed to be the time or the string? If it is a number, it is assumed to represent a time. Otherwise it is a string. If you want to wait forever for a certain string *and the string is a number*, then enclose it in quotes so it won’t be mistaken for a time. Control characters are represented in these strings the same as for the **%s** command.
%p	Send the phone number stored in the associated Remote Profile. This command allows the same dial script to support several outbound connections with different phone numbers. Otherwise, separate dial scripts would have been needed.

Global Connections

Selecting Global Connections displays the following screen.

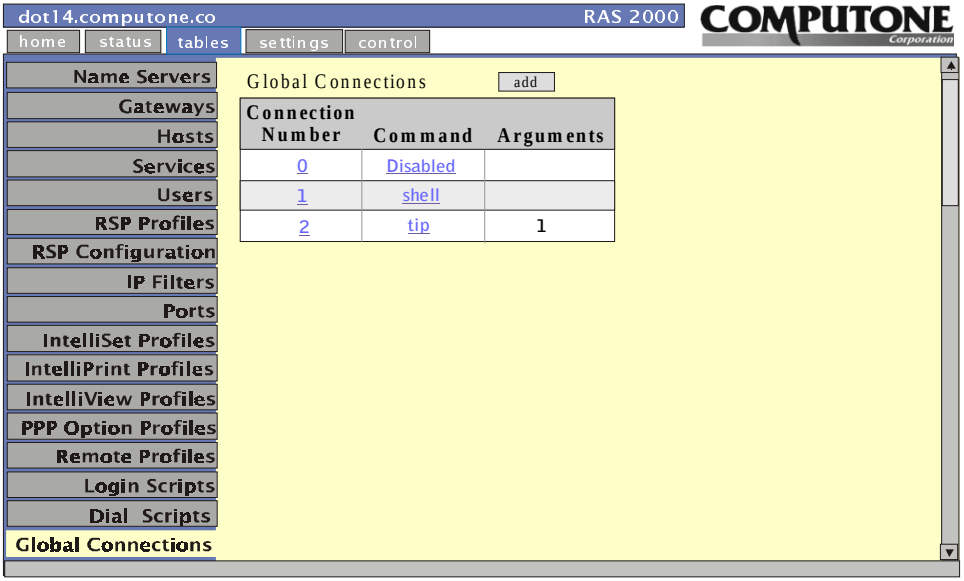

Figure 41 ***Global Connections*** **Screen**

Whenever a new selected connection is added to a user's configuration, the entry is automatically added to the *global connection table*. This is a master table that contains all the connections configured for all users. When you are using the global connection menu, you can choose to run any of the connections in this table.

You may also add, modify, and delete entries in the global connection table directly, without working through user configuration. In most installations with lots of login users, there tends to be more users than there are places to go. If the global connection table number is known for a particular connection, then you can configure the user more quickly. More importantly, if some system-wide parameter changes, you are more likely to be able to make a single change and affect all appropriate users. For example, perhaps lots of users are configured to rlogin to a certain host in order to perform a specific function. But later, this function is moved to a different host on your network.

You *could* change each user separately or do the following:

- Look at the user configuration form for one of these users. In the user's selected connection table will be the global connection number of that connection. Remember it.
- In the global connection menu, find the entry and change it. All other users using that entry will be updated as well. This is possible because the RA 2000 automatically forces users with identical connections to share a single global connection entry. Remember, entries must be completely identical. Even the spacing must be identical or separate entries are created.

If two users were configured with identical connections, and you wanted to make a change for one user only, you would have made the change using user configuration. This would automatically create a new entry in the global table for the user's new connection.

To change an entry in the global connection table, click on connection number. The following screen is displayed.

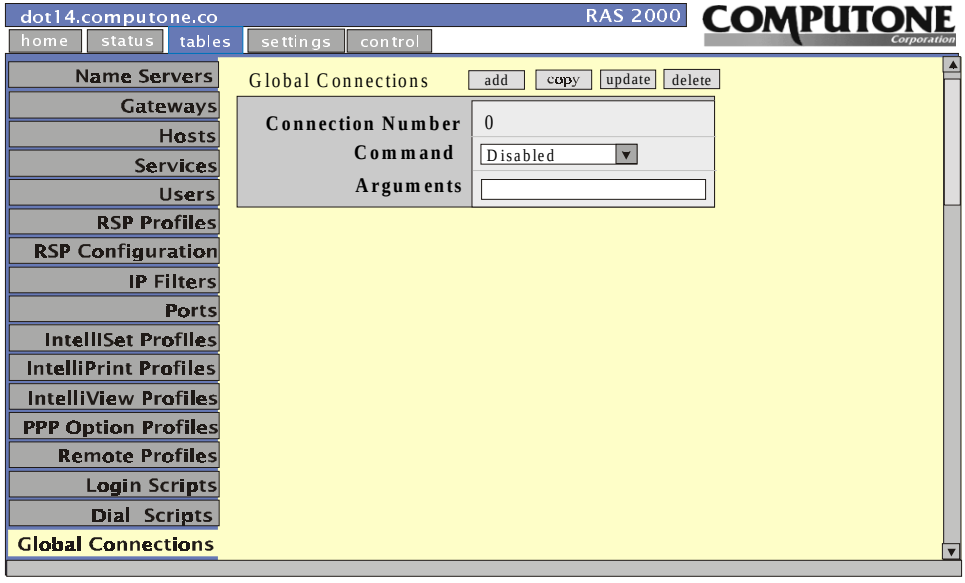

Figure 42 ***Global Connections Configuration*** **Screen**

The commands are pick-lists, you may choose between:

- *Free* - entry is available to store a new global connection.
- *Disabled* - entry is not available to store a new global connection.
- *shell* - use to administer and maintain RAS 2000 or to start a connection.
- *menu* - use to administer and maintain RAS 2000 or to start a connection.
- *rlogin* - use to start a connection to a specified host.
- *telnet* - use to start a connection to a specified host.
- *tip* - use to directly access a specified port.
- *tipmenu* - use to access all the ports.

To delete an existing global connection, you can try to set its command to *Free*. If there is a user configured with this as one of his selected connections, the RAS 2000 won't allow you to make the change. When you modify one of these entries, it affects all users whose selected connection table contains the entry.

The *Arguments* option provides a place for you to enter any arguments you feel are necessary for the command you select.

CHAPTER 4

Configuring System Settings

The following table lists the selections available for configuration through the ***Settings*** tab:

Table 13 *Settings* Selections

Selection	Description
System	Contains the host name, domain name, IP address, Ethernet address, console port, IP filter, RIP type, login prompt, password prompt, and user prompt.
Applications	Provides selection for Web Server (httpd), Secure Shell (sshd) and Insecure Shell (telnetd).
Boot	Provides selection of Boot Type, Host 1, File 1, Host 2, File 2, and Retry Count.
Syslog	Provides selection of Syslog Host, Syslog Facility, and Syslog Priority.
SNMP	Provides selection of SNMP State, Contact, Location, Trap Host 1, Trap Host 2, Get Host 1, Get Host 2, Set Host 1, Set Host 2, Access Rights 1, Access Rights 2, Community 1, Password 1, Community 2, Password 2,Community 3, Password 3, and Community 4, and Password 4.
RADIUS	Provides selection of Radius Host 1, Radius Host 2, Radius CHAP Secret, Accounting Host 1, Accounting Host 2, Accounting CHAP Secret, Retry Count, and Retry Time.
RIP	Provides selection of State, Version, Domain (RIP-II Only), Hot List Type, Host 1, Host 2, Host 3, Host 4, and Password (RIP-II Only).
Secured Shell	Provides selection of Hot Key Size, Server Key Size, Authentication Grace Time, Server Key Regen, TCP Port, Authentication Method, and Allow Root Login.
Web Server	Provides selection of TCP number.

Settings Tab

Selecting the ***Settings*** tab displays the following screen.

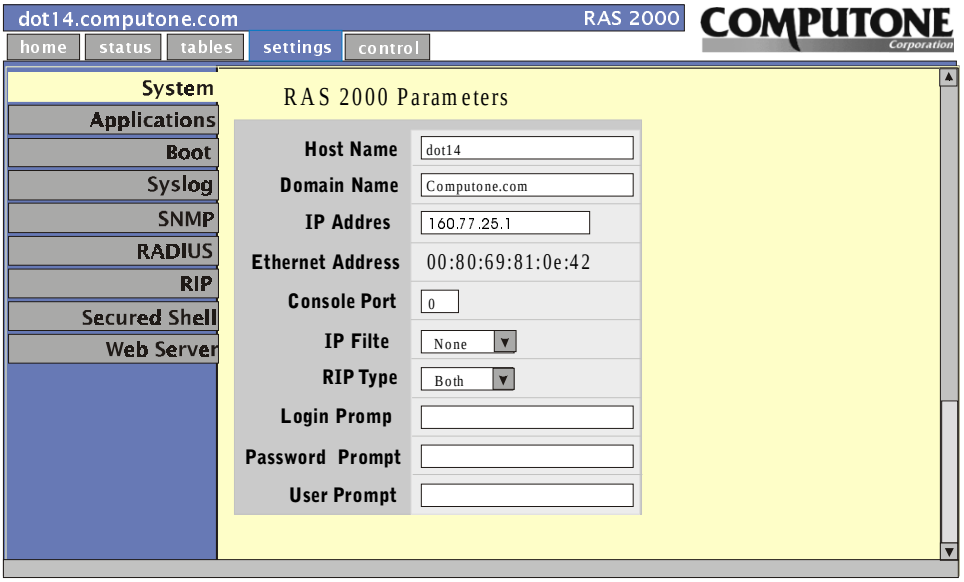

Figure 43 ***Settings*** **Screen**

The following table defines the selections.

Table 14 *System* Entries

Menu Entry	Description
Host Name	Although Internet Protocol identifies hosts and networks by their IP addresses, these addresses are not very practical for a human being to use. If you had to remember that Computone's FTP site was 160.77.1.10 and Generic General's WEB page was on 160.77.99.101, and so on, it would get difficult very quickly. That is why it is possible to identify a host by a name, rather than by its IP address. Enter your host's name here.
Domain Name	In much the same way you are given a registered Internet addresses, you can request a registered "domain name" as well. This domain name needs to be added to the end of any host name you assign. Domain names are organized in a hierarchic structure. The universe of names the "root domain") has been divided into basic groups, each with its own domain name and each potentially with its own administrator. For example, the domain**.com** assigns domains to commercial networks, **.gov** to government agencies, **.edu** to educational institutions, and **.org** to non-profit organizations. Notice that the individual elements of a domain name are separated by periods. Computone is a commercial establishment, so it has a domain name from whoever administers the **.com** domains.
IP Address	An IP address is a number of the for *nnn.nnn.nnn.nnn* where *nnn* represents some number between 0 and 255. The byte values 0 and 255 themselves have special meaning, so IP addresses for hosts use the values 1 through 254. Enter the IP address assigned by your system administrator.
Ethernet Address	This is the Ethernet address assigned to this RAS 2000 and is displayed for informational purposes. You cannot change it from the menu.
Console Port	This is the port number that is used for displaying system messages, such as the banners which are displayed when the RAS 2000 starts up. This is also the port which is used to display syslog messages when the syslog host is **console.** The default console is port 0 but you can set it to any port 0-15. To disable console messages entirely, set the console to port **255.**
IP Filter	This is the name of an IP filter you have defined (see **page 239** of the *RAS 2000 Software Configuration Guide*). The IP filter you specify here is applied only to traffic received on the RAS 2000's Ethernet interface. Separate IP filters can be applied to individual PPP and SLIP interfaces.

Table 14 *System* Entries

Menu Entry	Description
RIP Type	This specifies whether RIP (Routing Information Protocol) will be used *over the RAS 2000's Ethernet interface*. RIP can be separately enabled or disabled on each PPP or SLIP interface you configure.The selections are one of the following: • **send -** Routing information packets will be periodically broadcast over this interface and the RAS 2000 responds to specific routing information requests it receives over this interface. • **listen -** The RAS 2000 listens for any routing information packets broadcast to this interface by other hosts, updating its routing table as appropriate. • **both -** The RAS 2000 both broadcasts and listens for these information packets over this interface. • **none -** The RAS 2000 neither broadcasts nor listens for RIP packets over this interface.
Login Prompt	Allows you to change the login prompt.
Password Prompt	Allows you to change the password prompt.
User Prompt	Allows you to change the user prompt. **NOTE:** If you log in as user root, the user prompt is overwritten by a # sign.

Applications

Selecting Applications displays the following screen.

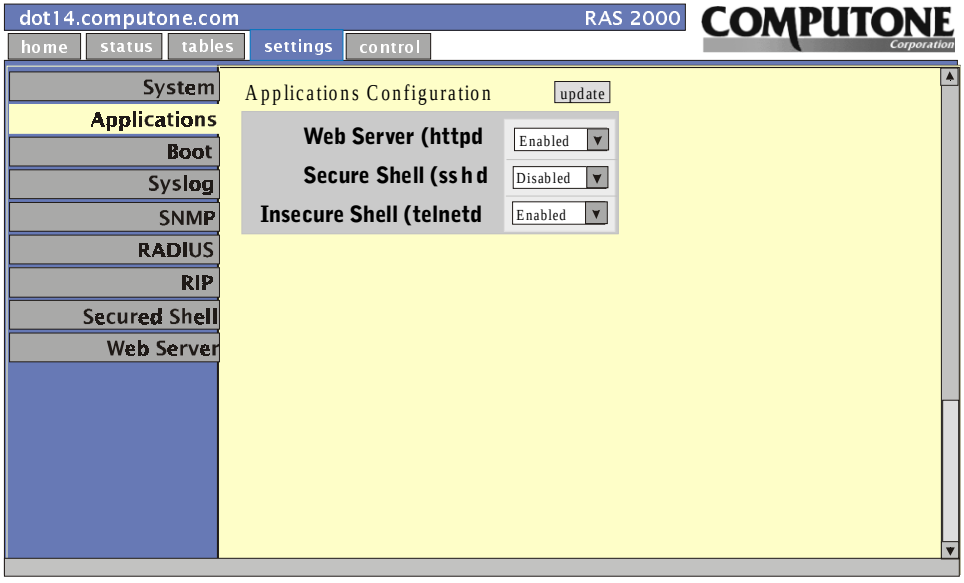

Figure 44 ***Applications*** **Screen**

This screen allows you to enable or disable the web server (httpd), secure shell (sshd), or insecure shell (telnetd).

IMPORTANT: If you disable the web server from this screen, save the configuration, and then wish to re-enable it you will have to do so from the command line by entering the following command:

```
apps set httpd enable
```

It is possible to disable the web interface, secure shell and insecure shell all at the same time. If you have previously defined a port as login-by-port, you can undo this situation by accessing this port. Otherwise, you must restore factory defaults and then reconfigure the RAS 2000.

There are three ways to restore factory defaults to NVRAM:

1. Enter *restore factory* at the shell command prompt if you can access the command line.
2. Hold the paperclip button down while powering-up the RAS 2000. Release it after 1 second.
3. Tap the ESC key a few times while powering-up the RAS 2000.

Even if you use method 2 or 3, it is recommended that you have a terminal or PC attached to port 0. The reason is that the machine has no IP address and, thus, is not accessible over the network. None of the methods above alter NVRAM so if you want the old configuration back, do a *restore*, *shutdown no* , or a power cycle. If you want to save the factory defaults, type in *save* just as you would when saving any other configuration.

Figure 4 5 shows the location of the paperclip button to restore factory defaults.

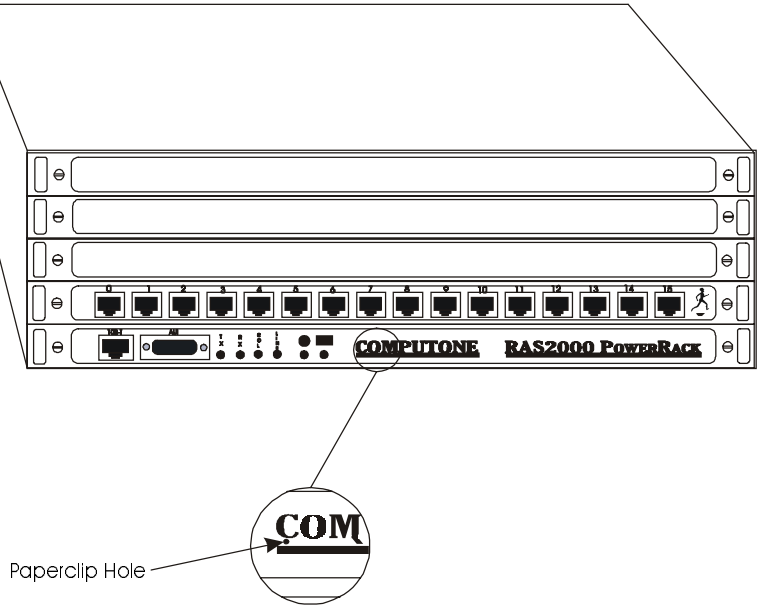

Figure 45 Restoring Factory Defaults

Boot

Selecting Boot displays the following screen.

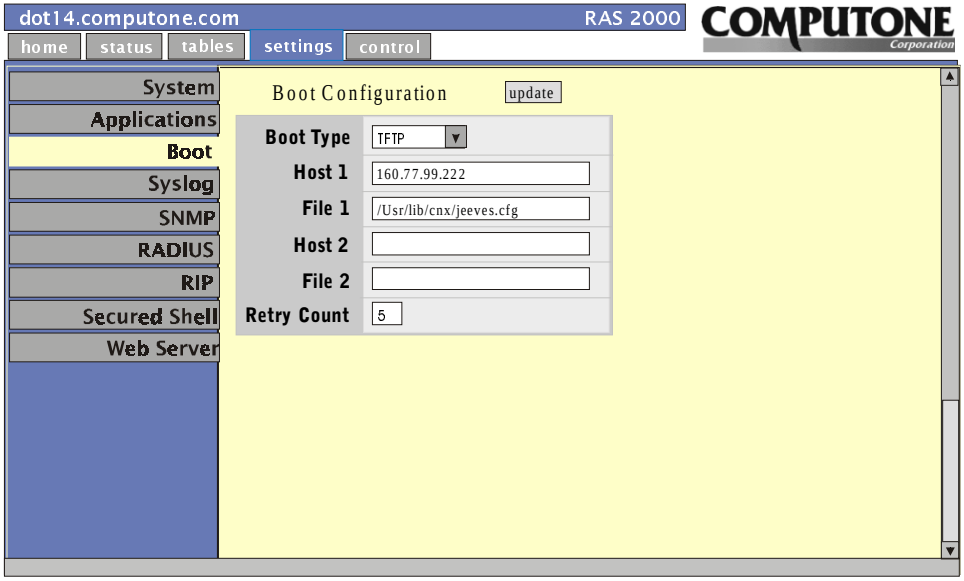

Figure 46 ***Boot*** **Screen**

The RAS 2000 has the option of running the version of software stored in its internal PROM, or of running a later version stored on one of the hosts on your local network. Booting a newer software version over the network is the most common method of upgrading when new releases of RAS 2000 software become available. The RAS 2000 also has the option of using the configurations stored in internal NVRAM, or of using configurations stored in a file on one of the network's hosts.

This option determines whether the RAS 2000 tries to get its software and configuration information from the network. Table 15 defines the boot type selections.

Table 15 *Boot* Configuration Parameters

Parameter	Description
Boot Type	
Disabled	The RAS 2000 runs using the software stored in PROM and the configuration stored in its local NVRAM. It does not attempt to get any of this information over the network.
BOOTP	With boot type **BOOT** , the RAS 2000 uses BOOTP protocol to find a host on the network which has been configured to supply this RAS 2000 with new software to run, or a configuration file to load. When BOOTP is used, you do not need to specify primary and secondary TFTP hosts, boot files, and configuration files. All configuration is done on the host that provides BOOTP services for your network. BOOTP protocol acts as a "front end" which provides configuration information to the RAS 2000 (i.e., the names of the boot file and configuration file to use). To actually download these files, the RAS 2000 uses TFTP, just as it would use with boot type **TFTP.**
TFTP	The RAS 2000 uses TFTP protocol to download a *TFTP Boot File* and *TFTP Config File* from a *TFTP Host*. You can specify *Primary* and *Secondary* files and hosts and if the primary fails, the secondary is used.
(Primary) Host 1	IP address of the host to check first for a configuration file.
(Configuration) File 1	Name of the configuration file on Host 1.
(Secondary) Host 2	IP address of the host to check second for a configuration file.
(Configuration) File 2	Name of the configuration file on Host 2.
Retry Count	This controls the number of times the RAS 2000 attempts to boot from the network before it gives up and uses the software and configuration stored locally. If the retry count is set to **0,** then it continues to retry forever.

When the RAS 2000 is configured to netboot, it first must bring up its own software in order to start up the networking code so that it can do the netbooting. If you are watching the console, you will see this older version's messages and banners. Since it knows it must netboot, the RAS 2000 configures itself to allow space in DRAM to download the new software; most serial ports are deactivated and non-essential processes are removed. After it is loaded, the new software is started and if you are watching the console you see *its* power-up messages and banners, which look almost like the first set, and then you are running.

Primary TFTP Host and Config File

These are used when the *Boot Type* is set to TFTP. Host 1 (Primary TFTP Boot Host*)* is the IP address of the first host the RAS 2000 tries to download its files from. File 1 (*Primary TFTP Config File* contains a configuration file that had earlier been saved from a RAS 2000 to this host (see chapter 14, ***Saving and Restoring Configurations***, of the RAS 2000 Software Configuration Guide).

When Net-booting Fails

If the netbooting should fail after a predetermined number of retries, it finally brings itself up using the software and configuration in PROM and local NVRAM. To do this, it cannot simply stop trying to TFTP the files. It has to actually reboot itself again, because it had previously reconfigured its DRAM for net-booting. This means temporarily deleting things which it now must reload from PROM to recover. The net result is that if you are watching the console you see a double set of banners in this case as well. A RAS 2000 which has not yet been configured with an IP address also uses BOOTP protocol in order to learn this and other information. This happens regardless of the **boot type** settings.

Syslog

Selecting Syslog displays the following screen.

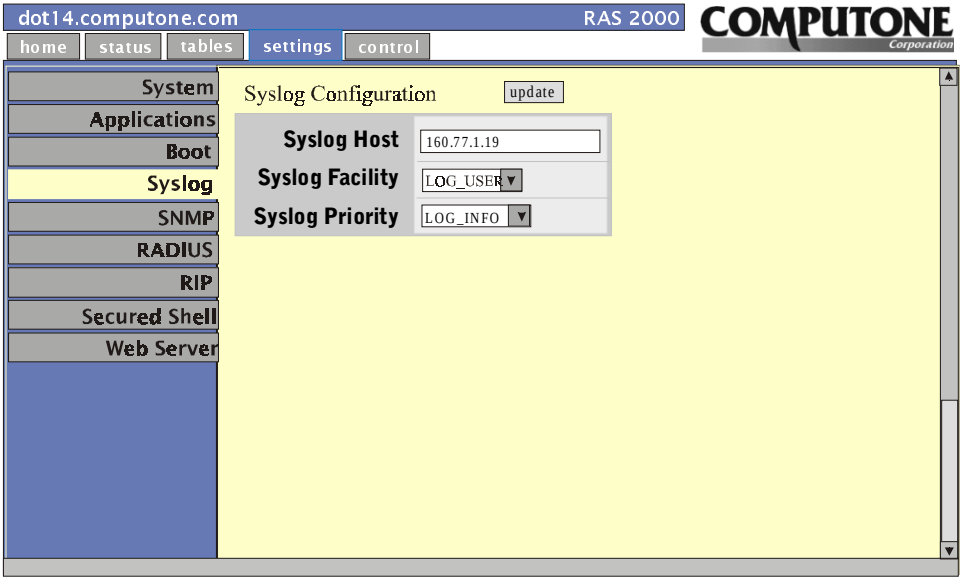

Figure 47 ***Syslog*** **Screen**

Network syslogging is a way of keeping track of what is going on within the system. There are two parts to syslogging:

1. The *syslog host* - a computer on the network running software called a syste log daemon (on UNIX hosts usually called *syslogd*).
2. One or more *syslog clients* - computers and devices configured to send syslog messages to the syslog host.

Syslog messages are sent as UDP datagrams and the syslog host does not confirm receipt by returning any acknowledgment. Therefore, syslog messages are intrinsically unreliable. While in most networks most syslog messages will be delivered most of the time, it is still possible that a message can be lost. Messages are especially apt to be lost if an extremely large number are sent to the same host very quickly, and especially if that host is otherwise busy.

Each syslog message contains three parts:

Table 16 Syslog Message Components

Syslog Message Components	Description
Message Text	The message itself. By convention this message begins with something to indicate which of the sender's processes generated the message. That is, messages generated from the *init* process might be expected to begin with *"init:"*, for example.
Priority	This identifies the urgency of the message. Syslog clients can be configured to send only messages of a certain priority or higher. Syslog hosts can be configured to store messages based on priority: messages of a certain urgency and higher being sent to a certain file, messages of a certain priority or lower are discarded, and still others are recorded elsewhere.
Facility	The facility serves to classify the source of the message. In that way, messages from user processes on the computer can be distinguished from system messages and other types. The syslog daemon can be configured to record messages from different facilities in different files. the separation that results for messages with different priorities. When messages of a certain priority are sent to a file, any messages with greater priority would also be sent to that file.

There are two more pieces of information that the syslog host will usually record in the log files:

1. The source (IP address or host name) of the syslog message. The IP address i obtained from the message's packet header.
2. The time the message was received. This is obtained from the syslog host's resident clock.

Syslogging provides an important debugging tool in many situations, so you should be comfortable with it. On the RAS 2000 it is especially useful when there is trouble bringing up a PPP connection, and when there is a problem involving Reverse-TCP ports.

Table 17 explains the Syslog screen selections.

Table 17 Syslog Screen Selections

Parameter	Description
Syslog Host	This is the IP address of some host on your network which is configured to receive *syslog* messages.
Syslog Facility	When the RAS 2000 sends syslog messages to a syslog host, it identifies them by *facilit* . Ultimately, this controls how the syslog host will log these messages: the syslog host will have been configured to record messages from some facilities into one file, and from other facilities into another.
LOG_USER	The facility is **LOG_USER** by default, but you can set it to any of the following: LOG_LOCAL1, LOG_LOCAL2, LOG_LOCAL3, LOG_LOCAL4, LOG_LOCAL5, LOG_LOCAL6, LOG_LOCAL7. There will probably be processes running on your syslog host which also syslog at the **LOG_USER** facility, so leave the RAS 2000's facility at the default if you want all the messages to be logged together. If your RAS 2000 has been configured to generate lots of syslog messages, then you may want them to appear in a file all their own. In this case, you would set it to one of the other facilities that no one else is using and see that your syslog host is configured to put those messages in a separate file.
Syslog Priority	The *priority* indicates which conditions on the RAS 2000 will generate messages, and is also used by the syslog host to determine in which file to log the messages. Below are shown the eight possible priority settings: • Higher in this table - only the most urgent conditions will generate syslog messages. • Lower in this table - less urgent errors and warnings will be sent (as well as all the conditions above). • Still lower in this table - more mundane information is added until the last entry turns on every last bit of syslogging information that can be had (far too much to leave on all the time).
LOG_ALERT LOG_EMERG	Nothing at all is presently sent from the RAS 2000 at this high of a priority: choosing either of these essentially turns off syslogging.
LOG_CRIT	"Impossible" conditions suggesting software bugs or a hard ware fault. Resource problems or other conditions likely to affect multiple ports until the problem is resolved.
LOG_ERR	All the above, plus: Normal errors, such as could be caused by configuration mistakes or other user errors.
LOG_WARNING	All the above, plus: Unusual conditions which do not necessarily mean there is a problem, but which may be informative if there *is* a problem.

Table 17 Syslog Screen Selections

Parameter	Description
LOG_NOTICE	All the above, plus: • Users logins, logouts, and disconnects. • PPP/SLIP links coming up and dropping. • Starting and stopping connections to Reverse-TCP ports.
LOG_INFO	All the above, plus: • PPP negotiation and additional information related to bringing up PPP and SLIP connections. • Additional RADIUS and Accounting information.
LOG_VERBOSE	All the above, plus: • **EXTREME** debugging output including dumps of all data sent to and from Reverse-TCP ports. This level corresponds to what is probably called **LOG_DEBUG** on your syslog host.

RAS 2000 Syslog Tips

The following information about syslog may be helpful.

Table 18 Syslog Tips

Topic	Description
File Storage	By default, your syslog host will often store syslog messages from the RAS 2000 in the same files as messages from other sources. Sometimes this is good, but at other times you want to separate the messages from a particular RAS 2000. To do this, change the syslog facility for that server, save and reboot. Configure your syslog host to send messages from that facility to a separate file.
Bringing up PPP Links	If you are trying to get a PPP link to come up, set your RAS 2000's syslog priority to LOG_INFO. This level generates syslog messages to track the PPP negotiation process.

Table 18 Syslog Tips

Topic	Description
Debugging Reverse-TCP Ports	To debug Reverse-TCP ports on the RAS 2000, configure the RAS 2000 for a priority of LOG_VERBOSE (probably called "LOG_DEBUG" on most syslog hosts). This will generate syslog messages that record all the data sent to and received by any Reverse-TCP host.
No Syslog Host	If you do not have a syslog host, the RAS 2000 can be configured to send the messages to its console instead. If you need to store the messages, attaching a plain terminal to the console would not be a good idea, because the information would just scroll off. If you are debugging and are needing to look at the syslog information, a better choice would be a computer running a terminal emulation package, one that will store whatever is displayed. In some cases even a printer has been attached.

SNMP

Selecting SNMP displays the following screen.

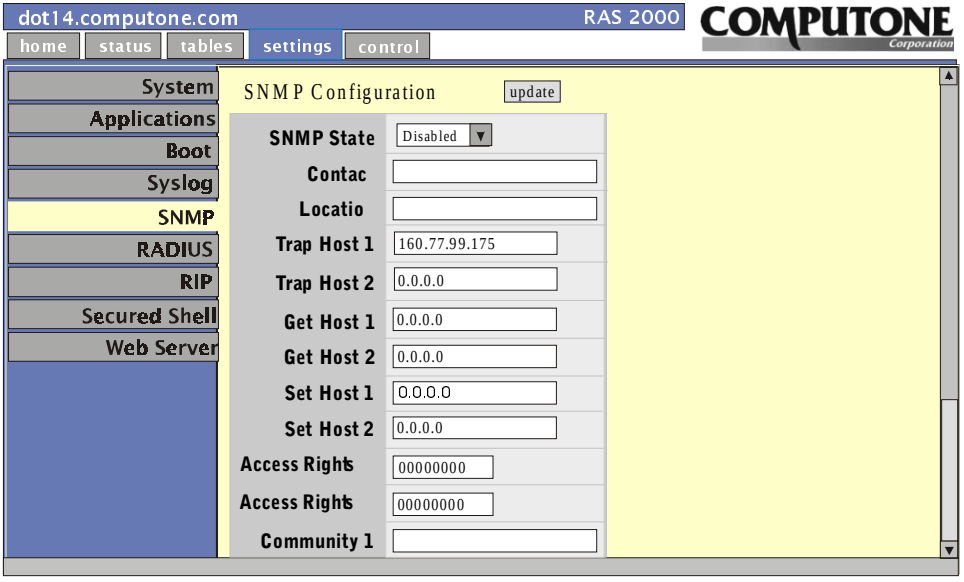

Figure 48 *SNMP* Screen

Overview

SNMP, or *Simple Network Management Protocol*, requires the following:

- One or more SNMP *managers*. A manager is a network computer that is running one or more SNMP management applications.
- One or more SNMP *agent* . Agents are network computers and devices (such as the RAS 2000) that can respond to queries from SNMP managers. The SNMP managers use UDP datagrams to send commands and queries to the agents and the agents send back responses (also using UDP). Agents also can send unsolicited messages, called *traps*, to report important conditions such as shutdown and start-up. The hosts that receive these traps are called *trap host* .

Table 19 SNMP Screen Selections

Selection	Description
SNMP State	Enables or disables SNMP. When SNMP is enabled, the RAS 2000 responds to queries from SNMP manager and sends trap messages to any trap hosts that may be configured.
Contact	Name of person who supports this function.
Location	Physical location of this machine.
Trap Host 1	IP address of Trap Host 1. A *Trap Host* is a host that sends unsolicited messages, called *traps*, to report important conditions such as shutdown and start-up.
Trap Host 2	IP address of Trap Host 2.
Get Host 1	IP address of Get Host 1. A *Get Host* is a host that is allowed to get information from the SNMP client.
Get Host 2	IP address of Get Host 2.
Set Host 1	IP address of Set Host 1. A *Set Host* is a host that is allowed to put information on the SNMP client.
Set Host 2	IP address of Set Host 2.
Access Rights 1	Access policy used for granting privileges to requesting parties for Host 1.
Access Rights 2	Access policy used for granting privileges to requesting parties for Host 2.
Community 1 - 4	Community 1 is a pairing of an SNMP agent with some arbitrary set of SNMP application entities. Each community is named by a string of octets.

Table 19 SNMP Screen Selections

Selection	Description
Password 1	Password used to authenticate requests from Community 1 Service.
Community 2	Community 2 is a pairing of an SNMP agent with some arbitrary set of SNMP application entities. Each community is named by a string of octets.
Password 2	Password used to authenticate requests from Community 2 Service.
Community 3	Community 3 is a pairing of an SNMP agent with some arbitrary set of SNMP application entities. Each community is named by a string of octets.
Password 3	Password used to authenticate requests from Community 3 Service.
Community 4	Community 4 is a pairing of an SNMP agent with some arbitrary set of SNMP application entities. Each community is named by a string of octets.
Password 4	Password used to authenticate requests from Community 4 Service.

Trap Hosts

The RAS 2000 can send *trap* messages to as many as two *trap hosts*, which you can configure using the configuration screen. When using the configuration screen, you enter IP addresses in either or both of the spaces provided. If there i only a single trap host, SNMP Trap Host2 is set to **0.0.0.0.**

When using the commands, you either **add** a new trap host or **delete** an existing one. You cannot add a new trap host if there are already two configured. You need to delete one first because this changes the IP address for that host to 0.0.0.0. To change an existing entry, you delete it first and then add a new one.

Enabling & Disabling

When SNMP is enabled, the RAS 2000 responds to queries from SNMP managers and sends trap messages to any trap hosts that may be configured. Note that the RAS 2000 responds to queries from SNMP managers even when no trap hosts are configured.

When SNMP is disabled, it does not listen for queries from SNMP managers and it does not send trap messages, even if trap hosts are configured.

If SNMP is disabled, the RAS 2000 re-allocates memory that would have been needed for SNMP support to make the memory available for other processes. This means that when you first enable SNMP, the change cannot take effect immediately because there is no way for SNMP to reclaim the resources it had sacrificed. **When you enable SNMP, the change does not take effect until** after you save the configuration and reboot.

RADIUS

Selecting RADIUS displays the following screen.

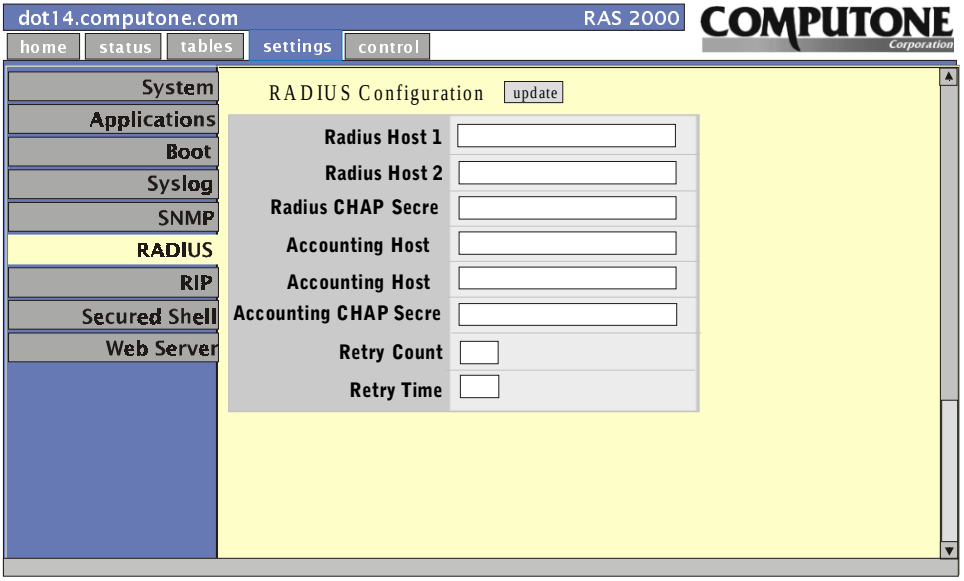

Figure 49 ***RADIUS*** **Screen**

The following table describes the RADIUS Configuration screen entries.

Table 20 RADIUS Screen Entrie

Selection	Description
Radius Host 1	Host name of main RADIUS authorization server or its IP address.
Radius Host 2	Host name of alternate RADIUS authorization server or its IP address.
Radius CHAP Secret	Authenticates authorization requests to RADIUS Server.
Accounting Host 1	Host name of main RADIUS accounting server or its IP address.
Accounting Host 2	Host name of alternate RADIUS accounting server or its IP address.
Accounting CHAP Secret	Authenticates replies from RADIUS Accounting Server.

Table 20 RADIUS Screen Entrie

Selection	Description
Retry Count	The number of times the RAS 2000 sends an authentication request to a host and waits for a reply. If there is no reply this request is re-sent a few times, and if a secondary RADIUS host is defined, it is tried as well. If there is still no reply, or if one of the replies is an *access rejection*, the connection is dropped.
Retry Time	RADIUS retry time in seconds.

The following table defines different elements of a RADIUS system.

Table 21 RADIUS Elements

Elements	Description
RADIUS Server	On some host, somewhere on your network, you have installed a software package known as a "**RADIUS server**". This package includes configuration files that have a list of users and their associated configuration, and a means for you to create and maintain this list. There will be a 'daemon" program which runs in the background and listens on the network for authentication requests from "**RADIUS clients**" (including RAS 2000s). There will also be configuration files that control which clients this RADIUS server is authorized to respond to, and additional security keys to ensure that the requests are actually coming from the authorized source.
Three RADIUS Elements	Regardless of the software implementation, there will always be: 1. The RADIUS Server software (including the RADIUS daemon). 2. A user authentication file and some means to maintain it. 3.A list of authorized clients, with associated security keys.
Logging into a Client	When a user tries to log in to a client, it sends authentication requests to the RADIUS server. When the RADIUS server gets the request, it looks up the user's information, and sends a reply back to the client. What information does the client need for this: 1. The IP address or host name of the RADIUS server. 2. The security key to be used.

Table 21 RADIUS Elements

Elements	Description
Two Parts of RADIUS	The two parts of RADIUS are: 1. RADIUS Authentication 2. RADIUS Accounting RADIUS Authentication occurs when a user tries to log into the RADIUS *client*. After prompting the user for login name and password, the client sends this information in an *authentication request* to the RADIUS server. The RADIUS server checks the validity of the request, then checks its database of user names and passwords. If they are bad it sends a *rejection* back to the client, which in turn rejects the login. If the login name and password are good, the RADIUS server sends back a packet containing information about this user and the client (i.e., the RAS 2000) uses this information to decide what type of service to supply for the user. RADIUS Accounting occurs when a user logs into or out of a RADIUS *client* after approving the login (either through an internal database or through RADIUS authentication). The client sends notification to the accounting server that this particular user has logged in. When the user logs off or is disconnected, the client also sends notification including the number of seconds the user was connected. When the RADIUS Accounting server receives these notices, it stores the information and then sends an acknowledgment back to the *client*. If the client does not receive an acknowledgment for its notices, it assumes they were lost and sends out duplicates. **You can do RADIUS authentication without doing accounting, or accounting without authentication.** If you are doing both, the accounting server can be the same host or a different one from the authentication server. Secondary authentication and accounting hosts can also be defined which the RAS 2000 uses when there is no reply from the primary servers.

Table 21 RADIUS Elements

Elements	Description
Things to Configure	Regardless of the software implementation, you must configure the following: 1. A list of authorized clients and their shared secrets. The RADIUS server needs to know the IP addresses of all the authorized RADIUS clients. Along with each client's address is a *secret*. You can pick whatever you like, but this same secret has to be configured into the *client* (RAS 2000) as well (see **page 143** of the *RAS 2000 Software Configuration Guide*). The RADIUS client and server use the secret to encrypt parts of the packets they send each other, and to guarantee that the messages and replies are authentic. Your RADIUS server might store this list in a text file and in Merit's implementation this is a file called *clients*. 2. A list of authorized users and their configuration information. The RADIUS server needs to know which users have what passwords and what these users are authorized to do after they log in. In Merit's implementation, this is a text file called *users*. Each user is listed along with password (or an indication that the UNIX password file should be consulted), and any restrictions as to which RAS 2000 or serial ports the user may be allowed to log in from. Information about the user is stored as a list of RADIUS protocol *attributes* and their associated *values*. These translate directly into the authentication reply the server sends back to the client.

For more information about RADIUS, refer to chapter 17 of the RAS 2000 Software Configuration Guide, ***User Authentication Using RADIUS.***

RIP

Selecting RIP displays the following screen.

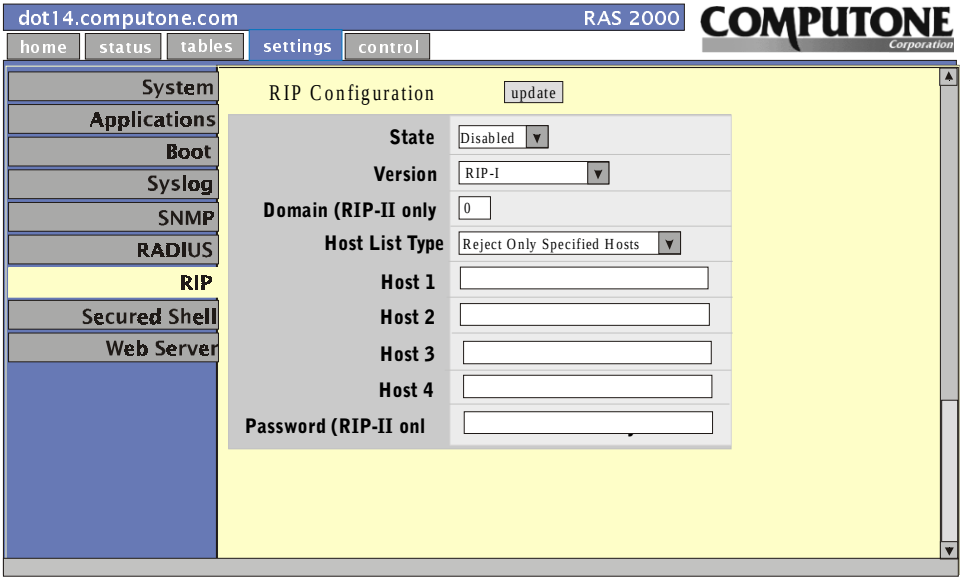

Figure 50 *RIP* Screen

RIP (Routing Information Protocol) is used when the RAS 2000 needs to share routing information with other hosts. By *listening*, it learns routes from other hosts and by broadcasting or *sending*, it tells other hosts about the route *it* knows

The following table defines the entries for the RIP screen.

Table 22 RIP Screen Entries

Selection	Description
State	Enables or disables RIP.
Version	Allows you to select which version of RIP you want. The selections are: RIP - Original version. RIP-II and broadcast- Version II of RIP which sends packets to all hosts on the link. RIP II and multicast -Version II of RIP which sends packets to hosts which have certain IP addresses.
Domain (RIP-II only)	The domain which this server is a member.
Host List Type	Sets the selection process for RIP. The selections are: Reject Only Specified Hosts or Accept Only Specified Hosts.
Host 1	Host 1 to be acted upon, depending on the Host List Type specification. Requires IP address or host name.
Host 2	Host 2 to be acted upon, depending on the Host List Type specification. Requires IP address or host name.
Host 3	Host 3 to be acted upon, depending on the Host List Type specification. Requires IP address or host name.
Host 4	Host 4 to be acted upon, depending on the Host List Type specification. Requires IP address or host name.
Password (RIP-II only)	Password for authenticating RIP II packets.

Secured Shell

Selecting Secure Shell displays the following screen.

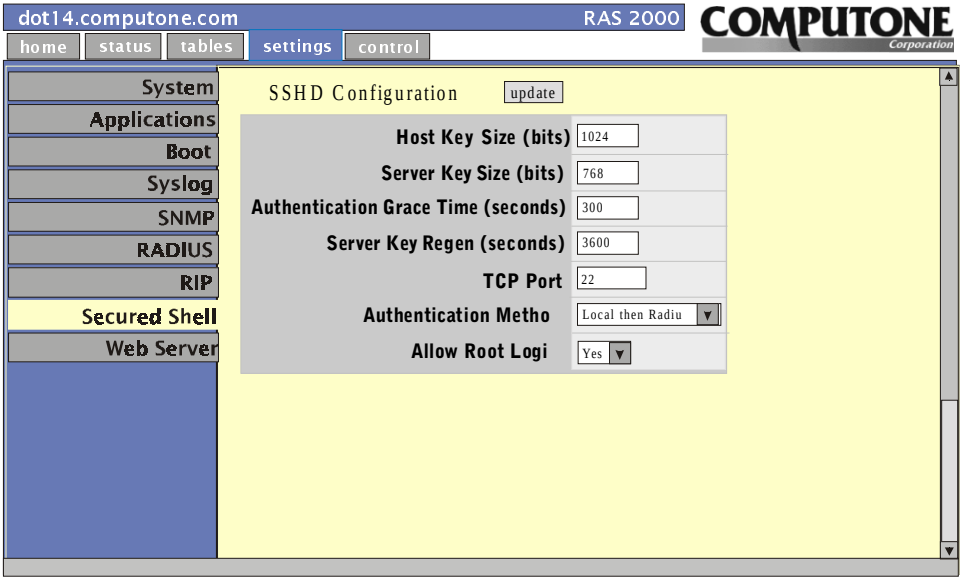

Figure 51 ***Secured Shell*** **Screen**

The RAS 2000 is shipped configured to have the web server running, as well as telnetd and sshd. However, sshd refuses to run until a host key has been generated, a process that must be started manually. Before generating a host key, how-ever, you must configure the secure shell parameters.

The Secured Shell screen entries are defined in the following table:

Table 23 Secured Shell Entries

Entry	Description
Host Key Size (bits)	The host key is generated one time and is good for the life of the machine. As such, it gets stored in an area of flash that can neither be saved nor restored.
Server Key Size (bits)	The server key is generated at run time, is stored in memory, and is never used for more than one hour (configurable).
Authentication Grace Time (seconds)	The number of seconds that a user has to successfully authenticate before being cut off.
Server Key Regen (seconds)	Defines the usable life span of a server key in seconds.
TCP Port	Defines the TCP port on which sshd listens for connection requests.
Authentication Method	Defines the method by which users are authenticated. Possible values are **Local**, **Radius**, or **Local then RADIUS**.
Allow Root Login	Either allow or disallow a user to login as root.

Key Size and Security

The host key size and the server key size are paramount in determining how secure the keys are. The following is important to remember.

Table 24 Key Size and Security

Key Type	Strong Security	Weak Security	Notes
Symmetric	128 bit	40 bits	
Asymmetric (Host & Server)	1024 bits	512 bit	It is recommended configuring the host key to be 1024 bits and the server key to be 768 bits. Since the session key is generated by the client, it is suggested that you configure your client (if possible) to generate a 128-bit key. The RAS 2000 supports up to 2048-bit host and server keys. The practicality of keys larger than that decreased substantially while CPU requirements increased exponentially. Additionally, if you are going to secure a RAS 2000, be sure to disable the web configuration (httpd) and the insecure shell (telnetd).

The following table defines symmetric and asymmetric key types.

Table 25 Encryption Key Types

Key Type	Definition	Advantage	Disadvantage
Symmetric	When the key used to decrypt the message is the same as the key used to encrypt the message, the key is said to be "symmetric".	They are fast. Encryption and decryption can be accomplished very quickly when compared to asymmetric keys.	Both the sender and the recipient must agree ahead of time on a specific key.
Asymmetric	Some algorithms derive separate encryption and decryption keys. These are said to be "asymmetric" and are also know as public/private key pairs or RSA keys.	The advantage of asymmetric keys is that you can publish one of them as your "public key" and keep the other one private. Thus someone wanting to send you a message encrypts the message with your public key, knowing that only the person with the private key namely you, can decrypt it. The reverse can also be used to send a message. If a recipient can decrypt a message using your public key, then the message must have originated and been encrypted by someone with the private key.	The disadvantage of asymmetric keys is that they are extremely math intensive and thus require lots of CPU time.

Configuring Secure Shell Parameters

To begin the secure shell parameter configuration process:

1. Do one of the following. Either telnet into the RAS 2000, or connect a serial terminal (or terminal emulator) to the console port (usually port 0).
2. From a shell prompt, run the `apps` command to display a list of applications and their current enabled or disabled status.

NOTE: If you wish to secure your RAS 2000, it is suggested that you disable httpd and telnetd and enable sshd. If you wish to leave your RAS 2000 unsecured, it is suggested that you disable sshd. For the time being, leave httpd and telnetd enabled so that if something goes wrong during the configuration process you are able to telnet in to fix it.

3. Enable the sshd applications with the following command:

```
apps set sshd enable
```

4. Run `ssh` and the ssh parameters are displayed:

Table 26 ssh Parameters

Parameter	Definition
hostkey	Number of bits that are used to generate the host key.
serverkey	Number of bits that are used to generate a server key.
authgrace	Number of seconds that a user has to successfully authenticate before being cut off.
regen	Defines the usable life span of a server key in seconds.
port	Defines the TCP port on which sshd listens for connection requests.
authmethod	Defines the method by which users are authenticated. Possible values are “local”, “radius”, or “both”. In the latter case, the local user file is checked first, then RADIUS.
allowroot	Either allow or disallow a user to login as root.

5. Set the parameters as desired using one or more commands of the form:

```
ssh set <paramName> <value>
```

For more help, type `help ssh`.

6. Save these configuration parameters to flash so that they are used the next time the machine is booted. Saving to a TFTP host is fine as well, as long as the system is configured to fetch its configuration from that same TFTP host. To save, type:

```
save
```

or

```
save <hostname> <filename>
```

All the configuration parameters are now set.

■ End of Procedure

Generating a Host Key

Once you have configured the secure shell parameters it is time to generate a host key. Use the following procedure:

1. Type `sshd gen`

Typical times for a 1024-bit key range from 21 seconds to 4 minutes. The host key is saved to a section of non-volatile memory that is not considered part of the system's configuration and, thus, is not included in saves and restores. Therefore, it is safe to save or restore your configuration via TFTP.

CAUTION: It is strongly recommended erasing the host key prior to selling the machine, shipping it back for RMA, or otherwise releasing the machine from your possession. An additional argument to `sshd` has been added to help you keep your system secure in the event you should ship it somewhere. This command is:

sshd erase

This erases the host key from non-volatile memory. During the RMA process, Computone can not guarantee the security of your machine, nor the security of your host key. Furthermore, Computone can not guarantee that you'll get the same machine you sent us; it could have a different engine card.

2. Once the host key generation is complete, reboot the machine with the following command:

```
shutdown now
```

This assures that the new configuration takes effect. Depending on the size of your server key, it may take a few minutes after a reboot for sshd to become ready. Recall that it requires both a host key and a server key, and that the server key is generated at runtime.

3. Start your ssh client and connect to the RAS 2000. Depending on the client it may or may not ask for your user name, though it should always ask for your password. After logging in you should see a shell prompt.

4. To complete securing your RAS 2000, it is recommended that httpd and telnetd be disabled. This can be accomplished from a shell prompt with the following command:

```
apps set httpd disable
apps set telnetd disable
```

5. Save this configuration to flash and/or a TFTP server:

```
save
```

6. Reboot the RAS 2000 using the following command to ensure these two services are NOT running:

```
shutdown now
```

■ End of Procedure

Web Server

Selecting Web Server displays the following screen.

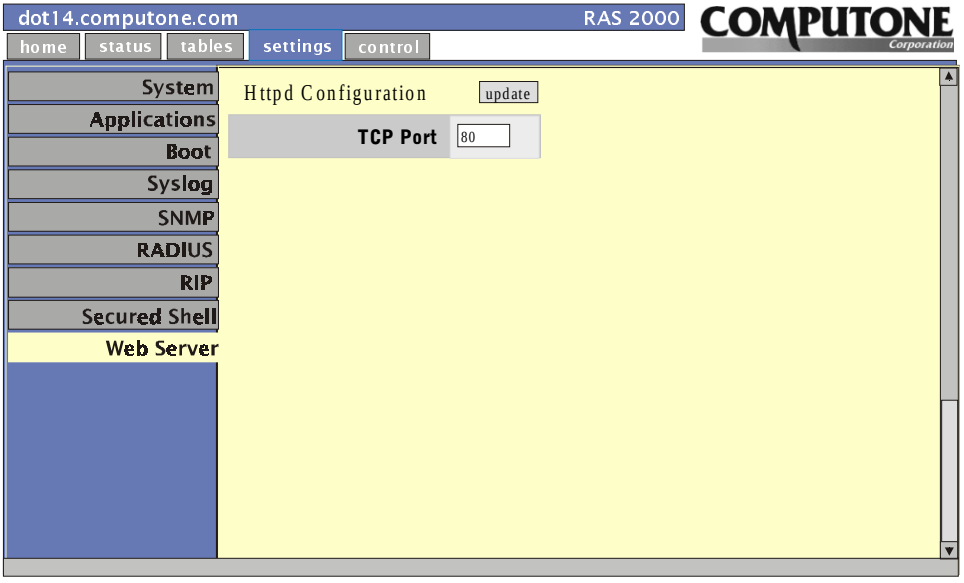

Figure 52 ***Web Server*** **Screen**

This screen allows you to set the TCP port that you want httpd to use.

CHAPTER 5

Using System Controls

The following table lists the selections available for configuration through the ***Control*** tab:

Table 27 *Control* Selections

Selection	Description
Shutdown	Allows you to reboot the system in the specified number of minutes.
Save to NVRAM	Allows you to save the current working configuration to non-volatile memory.
Save to Host: File	Allows you to save the current working configurationto a TFTP site.

Shutdown

Selecting Shutdown displays the following screen.

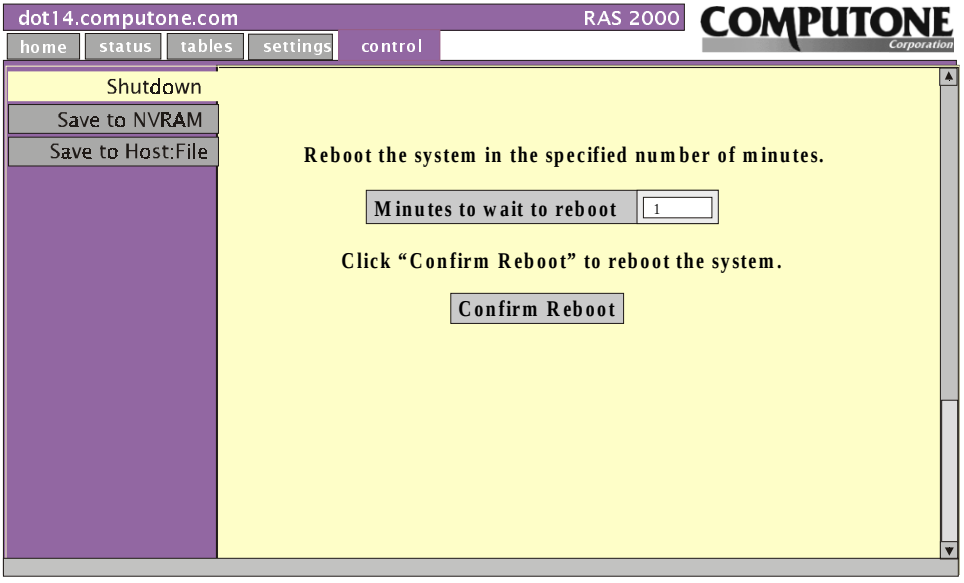

Figure 53 ***Shutdow*** **Screen**

This screen allows you to reboot the system in the number of minutes you specify.

Save to NVRAM

Selecting Save to NVRAM displays the following screen.

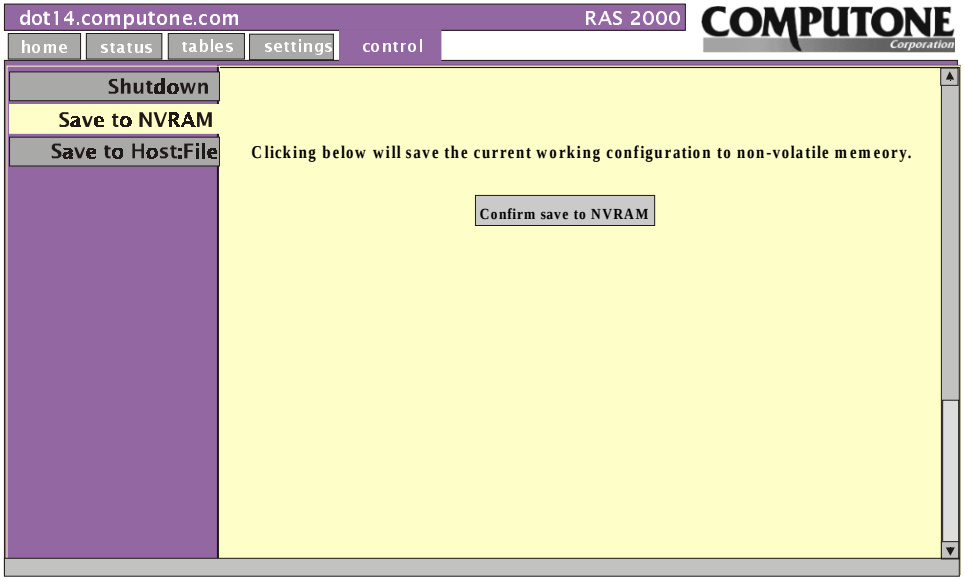

Figure 54 ***Save to NVRAM*** **Screen**

This screen allows you to save the current working configuration to non-volatile memory.

Save to Host: File

Selecting Save to Host: File displays the following screen.

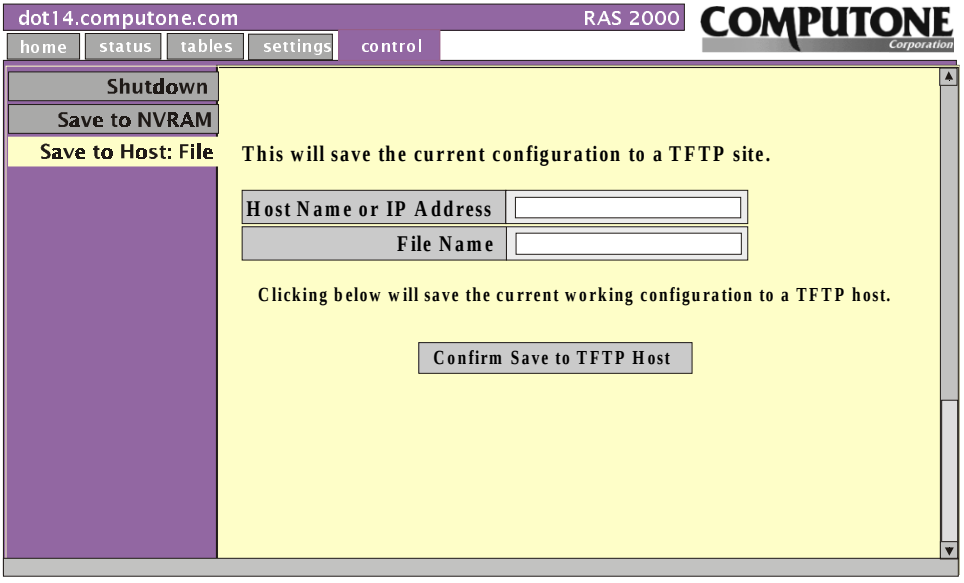

Figure 55 ***Save to Host: File*** **Screen**

This screen allows you to save the current working configurationto a TFTP site. You must enter the host name or IP address where you want the file to be stored, and you must also enter the name you want for the file.

INDEX

R

S

T

U

W