PortMaster® 4

Configuration Guide

Copyright and Trademarks

Disclaimer

Contents

10. Configuring T1, E1, and ISDN PRI

13. Using Synchronous Leased Lines

A. Networking Concepts

B. TCP and UDP Ports and Services

Command Index

Subject Index

About This Guide

The *PortMaster 4 Configuration Guide* provides configuration instructions and examples and software troubleshooting instructions for the PortMaster® 4 Integrated Access Concentrator from the Remote Access Business Unit of Lucent Technologies, Inc.

This configuration guide is one of three manuals that make up the comprehensive *PortMaster 4 User Manual*:

- *PortMaster 4 Installation Guide*

- *PortMaster 4 Configuration Guide*

- *PortMaster 4 Command Line Reference*

Consult the contents and indexes in each of these three manuals for detailed lists of topics and specific page references.

See the additional manuals listed under "PortMaster Documentation" for configuration, maintenance, and troubleshooting information common to all PortMaster products.

Audience

This guide is designed to be used by qualified system administrators and network managers. Knowledge of basic networking concepts is required.

PortMaster Documentation

The following manuals are available from Lucent Remote Access. They can be ordered through your PortMaster distributor or directly from Lucent.

The manuals are also provided as PDF and PostScript files on the *PortMaster Software CD* shipped with your PortMaster.

In addition, you can download PortMaster information and documentation from **http://www.livingston.com**.

- *ChoiceNet® Administrator's Guide*

    This guide provides complete installation and configuration instructions for ChoiceNet server software.

- *PortMaster Routing Guide*

    This guide describes routing protocols supported by PortMaster products, and how to use them for a wide range of routing applications.

- *PortMaster Troubleshooting Guide*

    This guide can be used to identify and solve software and hardware problems in the PortMaster family of products.

- *RADIUS for UNIX Administrator's Guide*

  This guide provides complete installation and configuration instructions for Lucent Remote Authentication Dial-In User Service (RADIUS) software for UNIX operating systems.

- *RADIUS for Windows NT Administrator's Guide*

  This guide provides complete installation and configuration instructions for Lucent RADIUS software for Microsoft Windows NT.

Additional References

RFCs

To find a Request for Comments (RFC) online, visit the website of the Internet Engineering Task Force (IETF) at **http://www.ietf.org/**.

RFC 768, *User Datagram Protocol*
RFC 791, *Internet Protocol*
RFC 792, *Internet Control Message Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specification*
RFC 950, *Internet Standard Subnetting Procedure*
RFC 1058, *Routing Information Protocol*
RFC 1112, *Host Extensions for IP Multicasting*
RFC 1144, *Compressing TCP/IP Headers for Low-Speed Serial Links*
RFC 1157, *A Simple Network Management Protocol (SNMP)*
RFC 1166, *Internet Numbers*
RFC 1212, *Concise MIB Definitions*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1256, *ICMP Router Discovery Messages*
RFC 1321, *The MD5 Message-Digest Algorithm*
RFC 1331, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*
RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
RFC 1334, *PPP Authentication Protocols*
RFC 1349, *Type of Service in the Internet Protocol Suite*
RFC 1413, *Identification Protocol*
RFC 1490, *Multiprotocol Interconnect Over Frame Relay*
RFC 1541, *Dynamic Host Configuration Protocol*
RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*
RFC 1552, *The PPP Internet Packet Exchange Control Protocol (IPXCP)*
RFC 1587, *OSPF NSSA Options*
RFC 1597, *Address Allocations for Private Internets*
RFC 1627, *Network 10 Considered Harmful (Some Practices Shouldn't be Codified)*
RFC 1634, *Novell IPX Over Various WAN Media (IPXWAN)*
RFC 1661, *The Point-to-Point Protocol (PPP)*
RFC 1700, *Assigned Numbers*
RFC 1723, *RIP Version 2*
RFC 1771, *A Border Gateway Protocol 4 (BGP-4)*
RFC 1812, *Requirements for IP Version 4 Routers*

RFC 1814, *Unique Addresses are Good*
RFC 1818, *Best Current Practices*
RFC 1824, *Requirements for IP Version 4 Routers*
RFC 1825, *Security Architecture for the Internet Protocol*
RFC 1826, *IP Authentication Header*
RFC 1827, *IP Encapsulating Payload*
RFC 1828, *IP Authentication Using Keyed MD5*
RFC 1829, *The ESP DES-CBC Transform*
RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
RFC 1878, *Variable Length Subnet Table for IPv4*
RFC 1918, *Address Allocation for Private Internets*
RFC 1962, *The PPP Compression Control Protocol (CCP)*
RFC 1965, *Autonomous System Confederations for BGP*
RFC 1966, *BGP Route Reflection, An Alternative to Full Mesh IBGP*
RFC 1974, *PPP Stac LZS Compression Protocol*
RFC 1990, *The PPP Multilink Protocol (MP)*
RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*
RFC 1997, *BGP Communities Attribute*
RFC 2003, *IP Encapsulation within IP*
RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*
RFC 2125, *The PPP Bandwidth Allocation Protocol (BAP), The PPP Bandwidth Allocation Control Protocol (BACP)*
RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*
RFC 2139, *RADIUS Accounting*
RFC 2153, *PPP Vendor Extensions*
RFC 2328, *OSPF Version 2*
RFC 2400, *Internet Official Protocol Standards*
RFC 2453, *RIP Version 2*

Books

Building Internet Firewalls. D. Brent Chapman and Elizabeth D. Zwicky. Sebastopol, CA: O'Reilly & Associates, Inc., 1995. (ISBN 1-56592-124-0)

DNS and BIND, 2nd ed. Paul Albitz and Cricket Liu. Sebastopol, CA: O'Reilly & Associates, Inc., 1992. (ISBN 1-56592-236-0)

Firewalls and Internet Security: Repelling the Wily Hacker. William R. Cheswick and Steven M. Bellovin. Reading, MA: Addison-Wesley Publishing Company, 1994. (ISBN 0-201-63357-4) (Japanese translation: ISBN 4-89052-672-2). Errata are available at **ftp://ftp.research.att.com/dist/internet_security/firewall.book**.

Internet Routing Architectures. Bassam Halabi. San Jose, CA: Cisco Press, 1997. (ISBN 1-56205-652-2)

Internetworking with TCP/IP, Volume 1: Principles, Protocols, and Architecture. Douglas Comer. Upper Saddle River, NJ: Prentice Hall, Inc. 1995. (ISBN 0-13-216987-8 (v.1))

Routing in the Internet. Christian Huitema. Upper Saddle River, NJ: Prentice Hall PTR, 1995. (ISBN 0-13-132192-7)

TCP/IP Illustrated, Volume 1: The Protocols. W. Richard Stevens. Reading, MA: Addison-Wesley Publishing Company. 1994. (ISBN 0-201-63346-9)

TCP/IP Network Administration. Craig Hunt. Sebastopol, CA: O'Reilly & Associates, Inc. 1994. (ISBN 0-937175-82-X)

Document Conventions

The following conventions are used in this guide:

Convention	Use	Examples
Bold font	Indicates a user entry—a command, menu option, button, or key—or the name of a file, directory, or utility, except in code samples.	• Enter **version** to display the version number. • Press **Enter**. • Open the **permit_list** file.
Italic font	Identifies a command-line placeholder. Replace with a real name or value.	• **set** *Ether0* **address** *Ipaddress* • Replace *Area* with the name of the OSPF area.
Square brackets ([ ])	Enclose optional keywords and values in command syntax.	• **set nameserver** [**2**] *Ipaddress* • **set** *S0* **destination** *Ipaddress* [*Ipmask*]
Curly braces ({ })	Enclose a required choice between keywords and/or values in command syntax.	**set syslog** *Logtype* {[**disabled**] [*Facility.Priority*]}
Vertical bar (|)	Separates two or more possible options in command syntax.	• **set** *S0*|*W1* **ospf on**|**off** • **set** *S0* **host default**|**prompt**|*Ipaddress*

Document Advisories

Note – means take note. Notes contain information of importance or special interest.

Caution – means be careful. You might do something—or fail to do something—that results in equipment failure or loss of data.

Warning – means danger. You might do something—or fail to do something—that results in personal injury or equipment damage.

Contacting Lucent Remote Access Technical Support

The PortMaster comes with a 1-year hardware warranty.

For all technical support requests, record your PortMaster ComOS version number and report it to the technical support staff or your authorized sales channel partner.

New releases and upgrades of PortMaster software are available by anonymous FTP from **ftp://ftp.livingston.com/pub/le/**.

In North America you can schedule a 1-hour software installation appointment by calling the technical support telephone number listed below. Appointments must be scheduled at least one business day in advance.

For the EMEA Region

If you are an Internet service provider (ISP) or other end user in Europe, the Middle East, Africa, India, or Pakistan, contact your local Lucent Remote Access sales channel partner. For a list of authorized sales channel partners, see the World Wide Web at **http://www.livingston.com/International/EMEA/distributors.html**.

If you are an authorized Lucent Remote Access sales channel partner in this region, contact the Lucent Remote Access EMEA Support Center Monday through Friday between the hours of 8 a.m. and 8 p.m. (GMT+1), excluding French public holidays.

- By voice, dial +33-4-92-92-48-48.

- By fax, dial +33-4-92-92-48-40.

- By electronic mail (email) send mail to **emea-support@livingston.com.**

For North America, Latin America, and the Asia Pacific Region

Contact Lucent Remote Access Monday through Friday between the hours of 7 a.m. and 5 p.m. (GMT –8).

- By voice, dial 800-458-9966 within the United States (including Alaska and Hawaii), Canada, and the Caribbean, or +1-925-737-2100 from elsewhere.

- By fax, dial +1-925-737-2110.

- By email, send mail as follows:

    - From North America and Latin America to **support@livingston.com**.

    - From the Asia Pacific Region to **asia-support@livingston.com**.

- Using the World Wide Web, see **http://www.livingston.com/**.

PortMaster Training Courses

Lucent Remote Access offers hands-on, technical training courses on PortMaster products and their applications. For course information, schedules, and pricing, visit the Lucent Remote Access website at **http://www.livingston.com/tech/training/index.html**.

Subscribing to PortMaster Mailing Lists

Lucent Remote Access maintains the following Internet mailing lists for PortMaster users:

- **portmaster-users**—a discussion of general and specific PortMaster issues, including configuration and troubleshooting suggestions. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-users** in the body of the message.

  The mailing list is also available in a daily digest format. To receive the digest, send email to **majordomo@livingston.com** with **subscribe portmaster-users-digest** in the body of the message.

- **portmaster-radius**—a discussion of general and specific RADIUS issues, including configuration and troubleshooting suggestions. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-radius** in the body of the message.

  The mailing list is also available in a daily digest format. To receive the digest, send email to **majordomo@livingston.com** with **subscribe portmaster-radius-digest** in the body of the message.

- **portmaster-announce**—announcements of new PortMaster products and software releases. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-announce** in the body of the message. All announcements to this list also go to the portmaster-users list. You do not need to subscribe to both lists.

Introduction 1

This chapter discusses the following topics:

- "PortMaster Software" on page 1-1

- "Preconfiguration Planning" on page 1-2

- "Setting the View" on page 1-3

- "Configuration Tips" on page 1-3

- "Basic Configuration Steps" on page 1-4

PortMaster Software

All PortMaster products are shipped with the following software:

- ComOS®—The communication software operating system already loaded in nonvolatile RAM (also called Flash RAM) on each PortMaster. You can use the ComOS command line interface to configure your PortMaster through a console.

- PMVision™—A graphical user interface (GUI) companion to the ComOS command line interface for Microsoft Windows, UNIX, and other platforms that support the Java Virtual Machine (JVM). Because PMVision also supports command entry, you can use a combination of GUI panels and ComOS commands to configure, monitor, and debug a PortMaster. When connected to one or more PortMaster products, PMVision allows you to monitor activity and edit existing configurations. PMVision replaces the PMconsole interface to ComOS.

  This application and other Java-based configuration tools for the PortMaster are available via anonymous FTP at **ftp://ftp.livingston.com/pub/livingston/software/java/**.

- **pmd** or **in.pmd**—The optional PortMaster daemon software that can be installed on UNIX hosts to allow the host to connect to printers or modems attached to a PortMaster. The daemon also allows the PortMaster to multiplex incoming users onto the host using one TCP stream instead of multiple streams like **rlogin**. The daemon is available for SunOS, Solaris, AIX, HP-UX, and other platforms.

  For installation and configuration instructions, copy the PortMaster software to the UNIX host as described on the *PortMaster Software CD* package.

- RADIUS—The RADIUS server daemon, **radiusd**, runs on UNIX systems, providing centralized authentication for dial-in users. The **radiusd** daemon is provided to customers in binary and source form for SunOS, Solaris, Solaris/X8.6, AIX, HP-UX, IRIX, Alpha OSF/1, Linux, and BSD/OS platforms.

  For installation and configuration instructions, see the *RADIUS for Windows NT Administrator's Guide* and *RADIUS for UNIX Administrator's Guide*.

- ChoiceNet—ChoiceNet is a security technology invented by Lucent to provide a traffic filtering mechanism for networks using dial-up remote access, synchronous leased line, or Ethernet connections. When used with RADIUS, ChoiceNet provides exceptional flexibility in fine-tuning the level of access provided to users.

  For installation and configuration instructions, see the *ChoiceNet Administrator's Guide*.

Preconfiguration Planning

Before the PortMaster can be used to connect wide area networks (WANs), you must install the hardware using the instructions in the *PortMaster 4 Installation Guide*.

This configuration guide is designed to introduce the most common configuration options available for the PortMaster 4. Review this material before you configure your PortMaster and, if possible, answer the following questions:

- What general configuration do you want to implement?

- Do you want to use a synchronous connection to a high-speed line?

- Will your high-speed lines use Frame Relay, ISDN, switched 56Kbps, or PPP?

- If you want dial-on-demand routing, do you want multiline load-balancing?

- Do you want Multilink Point-to-Point Protocol (PPP) (RFC 1717)?

- Do you want packet filtering for Internet connections?

- Do you want packet filtering for connections to other offices?

- Do you want dial-in users to use Serial Line Internet Protocol (SLIP), PPP, or both?

- If you use PPP, do you want Password Authentication Protocol (PAP) or Challenge Authentication Protocol (CHAP) authentication?

- Are you using a name service—Domain Name System (DNS) or Network Information Service (NIS)?

- Have you obtained the necessary network addresses?

- Do you want to enable Simple Network Management Protocol (SNMP) for network monitoring?

- Do you want dial-in only, dial-out only, or two-way communication on each port?

- What characteristics do you want to assign to the dial-out locations?

- How do you want to configure dial-in users?

- Do you want to use RADIUS or the internal user table on the PortMaster to authenticate dial-in users?

- Do you want to use ChoiceNet to filter network traffic?

- Do you want to use the console port for administration functions, or do you want to attach an external modem to the port?

- For dial-in users, do you receive service on analog lines, ISDN Primary Rate Interface (PRI), channelized T1, or E1?

Many other decisions must be made during the configuration process. This guide discusses the various configuration options and their implications.

Setting the View

The PortMaster 4 operates via the modules and boards installed in its slots. The system manager module installed in slot 4 provides overall (global) management for the entire chassis.

To monitor and configure a particular module or board, you use the **set view** command to set the view to the slot of the installed board or module. The default view is slot 4, which is the manager view.

Because the Ethernet interfaces on a PortMaster 4 are numbered uniquely, you can configure them from any view. However, you must reboot Ether0 and reset the appropriate slot for the other Ethernet interfaces to activate configuration settings.

Except for the manager module, for which the command line prompt displays no number, the prompt indicates the view you are in. For example:

```
Command> set view 3
View changed from 4 to 3
Command 3> set view 4
View changed from 3 to 4
Command>
```

The **save all** command saves all configuration information for all boards regardless of what view is set.

Configuration Tips

PortMaster configuration can be confusing because settings can be configured for a port, a user, or a remote location. Use Table 1-1 to determine how to configure your PortMaster.

Table 1-1   PortMaster Configuration Tips

If You Are Configuring . . .	Then Configure Settings on . . .
A network hardwired port or hardwired multiline load balancing	The port
One or more ports for dial-out operation	Dial-out locations using the location table
One or more ports for dial-in operation	Dial-in users using the user table or RADIUS
A callback network user	The callback location in the location table (refer to the location name in the user table)

Basic Configuration Steps

The exact PortMaster configuration steps you follow depend upon the hardware you are installing and your network configuration. However, the following general configuration steps are the same for all PortMaster products:

1. **Install the PortMaster hardware and assign an IP address and a password as described in the** *PortMaster 4 Installation Guide***.**

2. **Boot the system and log in with the administrative password.**

   You can configure the PortMaster from a terminal attached to the console port, through an administrative Telnet session, or through a network connection.

   Note – This configuration guide assumes that you have completed Step 1 and Step 2 and does not give details on hardware installation, IP address assignment, or administrative password assignment.

3. **If you want to use PMVision software to configure your PortMaster, install it on a workstation anywhere on your network.**

   PMVision is available via anonymous FTP at **ftp://ftp.livingston.com/pub/livingston/software/java/.** See the PMVision online help for information on using PMVision.

4. **Configure the global settings.**

   PortMaster global settings are described in Chapter 2, "Configuring Global Settings."

5. **Configure the Ethernet settings, and configure the IP protocol settings for your network.**

   PortMaster Ethernet settings are described in Chapter 4, "Configuring an Ethernet Interface."

6. **Configure the synchronous ports.**

   PortMaster synchronous port settings are described in Chapter 6, "Configuring a Synchronous WAN Port."

7. **Configure T1, E1, and ISDN PRI connections.**

   ISDN PRI connection configuration is described in Chapter 10, "Configuring T1, E1, and ISDN PRI" and Chapter 11, "Configuring a T3 Mux Board."

8. **Configure dial-in users in the user table, or configure RADIUS.**

   The user table is described in Chapter 5, "Configuring Dial-In Users." If you are using RADIUS security instead of the user table, see the *RADIUS for Windows NT Administrator's Guide* or *RADIUS for UNIX Administrator's Guide*.

9. **Configure ChoiceNet, if you are using it.**

   ChoiceNet is a traffic filtering mechanism for networks using dial-up remote access, synchronous leased line, or Ethernet. Refer to the *ChoiceNet Administrator's Guide* for more information.

10. **Configure dial-out locations in the location table.**

    The location table is described in Chapter 7, "Configuring Dial-Out Connections."

11. **Configure filters in the filter table.**

    Once the filters are created, they can be assigned as input or output filters for the Ethernet interface, users, locations, or hardwired ports. Filters are described in Chapter 8, "Configuring Filters."

12. **Configure the Layer 2 Tunneling Protocol (L2TP) if you are setting up an L2TP tunnel to an L2TP-compatible router.**

    See the *PortMaster 4 Command Line Reference* for information on the L2TP commands.

13. **Configure OSPF, if you are using this protocol.**

    OSPF is described in the *PortMaster Routing Guide*.

14. **Configure BGP, if you are using this protocol.**

    BGP is described in the *PortMaster Routing Guide*.

15. **Troubleshoot your configuration, if necessary, and back it up.**

    See the troubleshooting information in this guide and the *PortMaster Troubleshooting Guide* for instructions.

Once you have correctly configured all the settings necessary for your circumstances, your PortMaster is ready to provide communication service and routing for your network.

Configuring Global Settings 2

This chapter describes how to configure settings that the PortMaster 4 uses across all its ports and interfaces.

This chapter discusses the following topics:

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Setting the View

You configure global settings from the manager view. If you are not already in slot 4 (the default), use the following command to set the view to slot 4:

```
Command 3> set view 4
View changed form 3 to 4
Command>
```

Configuring Name Resolution

You can use either a network name service or the host table on the PortMaster 4 to map hostnames to IP addresses.

Using the Host Table

Each host attached to an IP network is assigned a unique IP address. Every PortMaster supports a local host table to map hostnames to IP addresses. If your network lacks a computer that can perform hostname resolution, the PortMaster allows entries in a local host table. Hostnames are used by the PortMaster only for your convenience when using the command line interface, or if you require users to enter hostnames at the host prompt.

To avoid confusion and reduce administrative overhead, Lucent recommends using the Domain Name System (DNS) or Network Information Service (NIS) for hostname resolution rather than the local host table. The PortMaster always checks the local host table before using DNS or NIS. For information on setting the NIS or DNS name service, refer to "Setting the Name Service."

Setting the Name Service

The PortMaster can work with network name services such as the Network Information Service (NIS) or the Domain Name System (DNS). Appendix A, "Networking Concepts," describes these name services. You must explicitly identify any name service used on your network.

The PortMaster stores all information by address rather than name. As a result, configuring the name server is useful only if you are using the command line interface for administration or if you prompt a login user for a host. If you are not using either of these features, you do not need to set the name service.

To set the name service, use the following command:

```
Command> set namesvc dns|nis
```

Once the name service is set, you must set the address of your NIS or DNS name server and enter the domain name of your network. See "Setting the Name Server" for instructions.

Setting the Name Server

The PortMaster supports RFC 1877, which allows remote hosts also supporting RFC 1877 to learn a name server through PPP negotiation. You must provide the IP address of the name server if you use a name service.

You must set a name service before you set a name server. See "Setting the Name Service." If you are not using a name service, you do not need a name server.

To set the name server, use the following command:

    Command> **set nameserver** *Ipaddress*

You can set an alternate name server with the following command:

    Command> **set nameserver 2** *Ipaddress*

You must set a domain name for your network after you set a name server. See "Setting the Domain Name."

You can disable the use of a name service by setting the name server's IP address to 0.0.0.0.

Setting the Domain Name

The domain name is used for hostname resolution. If you are using DNS or NIS, you must set a domain name for your network.

To set the domain name of your network, use the following command:

    Command> **set domain** *String*

Setting the Telnet Port

The Telnet access port can be set to any number between 0 and 65535. The Telnet port enables you to access and maintain the PortMaster using a Telnet connection to this TCP port. If 0 (zero) is used, Telnet administration is disabled. The default value is 23. Ports numbered 10000 through 10100 are reserved and should not be used for this function. Up to four administrative Telnet sessions at a time can be active.

To set the Telnet access port to port number *Tport*, use the following command:

    Command> **set telnet** *Tport*

Using the Telnet Port as a Console Port

If the console port is set from a Telnet session, the current connection becomes the console. This feature is useful for administrators who log in to a port using Telnet and need to access the console for debugging purposes.

Note – Only one Telnet session can receive console messages at a time.

To set the current Telnet access port as a console port, enter the following command:

```
Command> set console
```

Setting Management Application Connections

PMVision, ChoiceNet, and the ComOS utilities **pmdial, pmcommand**, **pminstall**, **pmreadconf**, **pmreadpass**, and **pmreset** all use port 1643. For more than one of these applications to connect at the same time, you must set the maximum number of connections to two or more. The maximum is 10 connections.

If you use ChoiceNet to download filters dynamically, be sure to set the maximum number of connections to 10.

To set the maximum number of concurrent connections for management applications into the PortMaster, use the following command:

```
Command> set maximum pmconsole Number
```

Setting System Logging

PortMaster products enable you to log authentication information to a system log file for network accounting purposes.

Setting the Loghost

To set the IP address of the loghost—the host to which the PortMaster sends **syslog** messages—use the following command:

```
Command> set loghost Ipaddress
```

Note – Do not set a loghost at a location configured for on-demand connections, because doing so keeps the connection up or brings up the connection each time a **syslog** message is queued for the **syslog host**.

Setting the loghost's IP address to 0.0.0.0 disables **syslog** on the PortMaster. This change requires a reboot to become effective.

RADIUS accounting provides a more complete method for logging usage information. Refer to the *RADIUS for Windows NT Administrator's Guide* and *RADIUS for UNIX Administrator's Guide* for more information on accounting.

Disabling and Redirecting Syslog Messages

By default, the PortMaster logs five types of events at the informational (**info**) priority level using the authorization (**auth**) facility on the loghost. You can disable logging of one or more types of events and change the facility and/or priority of log messages.

To disable logging of a type of event, use the following command:

```
Command> set syslog Logtype disabled
```

Use the *Logtype* keyword described in Table 2-1 to identify the type of event you want to disable—or enable again.

Table 2-1   Logtype Keywords

Logtype Keyword	Description
admin-logins	**!root** and administrative logins.
user-logins	Nonadministrative logins; you might want to disable this logtype if you are using RADIUS accounting.
packet-filters	Packets that match rules with the **log** keyword.
commands	Every command entered at the command line interface.
termination	More detailed information on how user sessions terminate.

You can change the facility, the priority, or both, of log messages.

To change the facility or priority of log messages, use the following command. Be sure to separate the *Facility* and *Priority* keywords with a period (.).

    Command> **set syslog** *Logtype Facility.Priority*

The facility and priority can be set for each of the five types of logged events listed in Table 2-1.

Table 2-2 and Table 2-3 show the keywords used to identify facilities and priorities. Lucent recommends that you use the **auth** facility or the **local0** through **local7** facilities to receive **syslog** messages from PortMaster products, but all the facilities are provided. See your operating system documentation for information on configuring **syslog** on your host.

Table 2-2   Syslog Facility Keywords

Facility	Facility Number	Facility	Facility Number
kern	0	**cron**	15
user	1	**local0**	16
mail	2	**local1**	17
daemon	3	**local2**	18
auth	4	**local3**	19
syslog	5	**local4**	20
lpr	6	**local5**	21
news	7	**local6**	22
uucp	8	**local7**	23

Table 2-3   Syslog Priority Keywords

Priority	Number	Typically Used For
emerg	0	Messages indicating the system is unusable
alert	1	Messages announcing action that must be taken immediately
crit	2	Critical messages
err	3	Error messages
warning	4	Warning messages
notice	5	Normal but significant messages
info	6	Informational messages
debug	7	Debug-level messages

To determine current **syslog** settings, enter the following command:

    Command> **show syslog**

Setting Administrative Logins to Serial Ports

When you log in using **!root**, administrative logins to the serial ports are enabled by default. You can enable or disable administrative logins them by using the following command:

    Command> **set serial-admin on|off**

If administrative login is disabled, you can still use port C0 by setting the console (bottom) DIP switch to the left (on) position.

Setting the Chassis

When you use the PortMaster 4 as an AnyMedia™ MultiService Module (MSM), you must specify the chassis type for PMVision to be able to display it. Use the following command to set the PortMaster 4 as an MSM:

    Command> **set chassis msm-rac**

Use the **save all** command to save changes to nonvolatile RAM. The chassis is identified as a PortMaster 4 by default.

Configuring Local IP Addresses

The PortMaster 4 supports up to four internal routable IP addresses, which the PortMaster advertises as host routes through RIP-2 and the Open Shortest Path First (OSPF) routing protocol. When you configure a local IP address, it becomes the PortMaster global address for network handles such as RADIUS, the Domain Name System (DNS), SNMP, the intermachine trunk (IMT), and **bootp**. By referencing an IP address instead of an interface, you do not lose the service if the interface goes down.

With the local IP address feature, you can specify the Ethernet interface the PortMaster uses as the default service address. For example, if RADIUS and the Signaling System 7 (SS7) gateway are on a private network range attached to Ether0, you can use the Ether0 address as the first local IP address.

IPCP Negotiation

During PPP negotiations for the IP Control Protocol (IPCP), the PortMaster 4 uses the following order of precedence when choosing an IP address to identify itself:

1. The Local IP address configured in the user profile, if set

2. The global reported IP address, if set

3. The first global local IP address, if set

4. The second global local IP address, if set

5. The third global local IP address, if set

6. The fourth global local IP address, if set

7. The IP address of Ether1

8. The IP address of Ether0

Main IP Address

When the PortMaster creates an IP packet, it must identify itself by placing a source address in the IP header. To do so, the PortMaster chooses either the main IP address or the nearest IP address, depending on the service used. The main IP address is chosen in the following order, but the nearest IP address is the IP address of the interface on which the packet exits the PortMaster 4:

1. The first global local IP address, if set

2. The second global local IP address, if set

3. The third global local IP address, if set

4. The fourth global local IP address, if set

5. The IP address of Ether1

6. The IP address of Ether0

The following services use the main IP address:

- **syslog**

- **traceroute**

- **telnet**

- DNS

- RADIUS authentication and accounting

- ChoiceNet

The following services use the nearest IP address:

- **ping**

- OSPF

- RIP

- **rlogin**

The global local IP address settings can be displayed with the **show global** and **show routes** commands.

Setting the Local IP Address

To assign the PortMaster 4 IP addresses that are not limited by network interfaces, use the following command:

```
Command> set local-ip-address [1|2|3|4] Ipaddress
```

For example, to set the local IP address to 10.112.34.17, enter the following command:

```
Command> set local-ip-address 10.112.34.17
Local IP Address (1) changed from 0.0.0.0 to 10.112.34.17
```

To set 192.168.54.6 as the second local IP address on the same PortMaster, enter the following:

```
Command> set local-ip-address 2 192.168.54.6
Local IP Address (2) changed from 0.0.0.0 to 192.168.54.6
```

Use the **show global** command to view local IP addresses.

Configuring an IP Address Pool

You can dynamically assign IP addresses to PPP or SLIP dial-in users. By assigning addresses as needed from a pool, the PortMaster requires fewer addresses than if each user is assigned a specific address. When a dial-in connection is closed, the address goes back into the pool and can be reused.

When creating an address pool, you explicitly identify the first address in the sequence of addresses available for temporary assignment. The PortMaster allocates one address in the pool of addresses for each port configured for network dial-in.

To set the value of the first IP address to assign for dial-in ports, use the following command:

```
Command> set assigned_address Ipaddress
```

The default number of addresses available for the address pool is equal to the number of ports configured for network dial-in. The address pool size is determined during the boot process. You can also set the number of IP addresses assigned to the pool with the **set pool** command.

To limit the size of the IP address pool, use the following command:

> Command> **set pool** *Number*

Note – If you decrease the number of addresses in the pool, you must reboot the PortMaster for the change to take effect.

Setting the Reported IP Address

Some sites require a number of different PortMaster devices to appear as a single IP address to other networks. You can set a reported address different from the Ether0 or Ether1 address. For PPP connections, this address is reported to the outside and placed in the PPP startup message during PPP negotiation. For SLIP connections, this address is reported and placed in the SLIP startup message during SLIP startup.

To set a reported IP address, use the following command:

> Command> **set reported_ip** *Ipaddress*

Configuring Named IP Pools

With the IP pool feature, you can set up multiple dynamically assigned address pools on the PortMaster. Each IP pool contains four elements.

- **Name**—a character string that uniquely identifies an IP pool. By identifying an IP pool by name instead of by base IP address, you can use a single name for an entire network system but assign different base IP addresses for each network access server in the system.

- **IP address**—the base IP address of a pool. When dynamically assigning addresses to users, the PortMaster begins with the base address and increments up to the size of the pool.

- **Netmask**—the size of the address pool.

- **Gateway**—the IP address of the pool gateway.

Note – Configuration information for IP pool is stored in the file **/manager/ippools**. If you use the **erase** command to delete this file, you remove the entire IP pool.

The named IP pools feature introduces a new RADIUS attribute (193) that takes a string corresponding to a name in the IP pool table. You must configure a user profile for named IP pools through RADIUS. The PortMaster does not support IP pools in the local user table.

This section describes how to set up named IP pools and includes the following topics:

- "How PortMaster Address Assignment Works" on page 2-10

- "Displaying Named IP Pool Information" on page 2-10

- "Creating Named IP Pools" on page 2-10

- "Creating a Default IP Pool" on page 2-11

- "Resetting the IP Pool" on page 2-11

- "Deleting Named IP Pools" on page 2-11

- "Setting Address Ranges" on page 2-12

- "Setting a Named IP Pool Gateway" on page 2-13

- "Setting Named IP Pools in RADIUS" on page 2-13

How PortMaster Address Assignment Works

The order of priority for address assignment is as follows for a user dialing in and expecting to receive an address from an assigned pool:

1. If a named IP pool is configured in the pool table **and** the RADIUS user profile has the IP-Pool-Name attribute configured for the user, the PortMaster assigns an address from the named IP pool.

2. If the IP-Pool-Name attribute is not configured in the RADIUS user profile **and** an address range is configured for the Quad T1 or Tri E1 board that the user comes in on, the PortMaster assigns the user an address from the address range configured for the Quad T1 or Tri E1 board.

3. If the IP-Pool-Name attribute is not configured in the RADUS user profile **and** the Quad T1 or Tri E1 board's assigned range is set to 0.0.0.0, **and** a default IP pool is configured in the pool table, the PortMaster assigns the user an address from the address range specified for the default IP pool.

Displaying Named IP Pool Information

Use the **show table ippool** command to display IP pool configuration information. For example, to display the configuration for an entire IP pool and to view all entries, enter the following command:

```
Command> show table ippool
Name:  livermore                     Default Gateway: 10.23.45.56
Address/netmask        Gateway
------------------     -----------------
192.168.1.0/29          0.0.0.0
192.168.2.253/30       0.0.0.0
192.168.3.50/25        0.0.0.0
10.4.5.0/24            192.168.222.3
```

Refer to your RADIUS documentation for information about modifying a RADIUS dictionary.

Creating Named IP Pools

To add a named IP pool to the pool table, use the following command:

```
Command> add ippool Name
```

An IP pool name can contain up to 31 characters. There is no limit to the number of IP pool entries you can configure. When you add a named IP pool to the pool table on the PortMaster, you must also add the IP-Pool-Name attribute to the RADIUS user profile.

(See "Setting Named IP Pools in RADIUS" on page 2-13.) If you do not want to configure a RADIUS user profile, you can create a default IP pool. (See "Creating a Default IP Pool.")

Creating a Default IP Pool

When you configure a named IP pool, you must also add the IP-Pool-Name attribute to the RADIUS user profile. If you do not want to configure a RADIUS user profile with a named IP pool, you can create a default IP pool. When you create a default IP pool, a user dialing in receives an address from the address range specified in the default IP pool, unless you also have an IP address range configured on the Quad T1 or Tri E1 board the user comes in on.

To add a default IP pool to the pool table, enter the following command:

    Command> **add ippool default**

Resetting the IP Pool

Whenever you make changes to the IP pool table, you must reset the pool for the changes to take effect.

    Command> **reset ippool**

Resetting the IP pool causes the PortMaster to convert address ranges into summarized routes for propagation through the routing protocols.

Note – After you issue the **reset ippool** command, the routing protocols can take a short while to replace the old routes.

Deleting Named IP Pools

To remove an address range from a named IP pool, or to remove the IP pool entirely, use the following command:

    Command> **delete ippool** *Name* **address-range** *Ipaddress*|**all**

For example, to delete an IP pool named *livermore* with the address range 192.168.1.0, enter the following command:

    Command> **delete ippool livermore address-range 192.168.1.0**
    Range 192.168.1.0 in livermore successfully deleted

To remove the entire IP pool entry, for example, *livermore*, enter the following command:

    Command> **delete ippool livermore all**
    Pool livermore successfully deleted

Remember to enter the **reset ippool** command to make the changes take effect.

Setting Address Ranges

The PortMaster assigns addresses to users from address ranges that you set for named IP pools with the following command:

```
Command> set ippool Name Ipaddress/NM|Ipaddress Netmask [Gateway]
```

You can specify up to eight address ranges for each IP pool. When you specify multiple ranges, the earlier ranges are preferred over later ranges.

As the syntax of the **set ippool** command indicates, an address range must have a netmask associated with it. The address-netmask pair can be expressed as a dotted decimal base IP address followed by a mask number between 1 and 30 (for example, 192.168.1.0/24), or by the older dot-separated netmask notation (for example, 192.168.1.0 255.255.255.0). Because the first and last addresses in a range are used for the network and for broadcast and are not assigned, netmasks of /31 and /32 (255.255.255.254 and 255.255.255.255) are not valid.

For example, to assign a range of 254 address to an IP pool named *livermore*, enter the following command:

```
Command> set ippool livermore address-range 192.168.1.0/24
Range 192.168.1.0/24 256 with gateway 0.0.0.0 add to livermore
```

Although the output to this command indicates a range size of 256 address as specified by the /24 netmask, only 254 of these addresses are available to be assigned to users. The first and last addresses are not assigned. The base (second) address in the range is incremented as addresses are assigned. Remember to enter the **reset ippool** command whenever you make changes to the IP pool.

This same address range can be expressed using the dot-separated netmask notation as follows:

```
Command> set ippool livermore address-range 192.168.1.0 255.255.255.0
Range 192.168.1.0/24 256 with gateway 0.0.0.0 add to livermore
```

Setting an Address Range Gateway

As the syntax of the **set ippool** command indicates, you can optionally assign a default gateway address to an address range. For example, to set 10.34.56.78 as the default gateway for IP pool *livermore* with address range 192.168.1.0/24, enter the following command:

```
Command> set ippool livermore address-range 192.168.1.0/24 10.34.56.78
Range 192.168.1.0/24 256 with gateway 10.34.56.78 add to livermore
```

Always reset the pool when you make changes to the named IP pool.

```
Command> reset ippool
```

The default gateway functions as a crossbar IP address. See the *PortMaster 4 Command Line Reference* for details about how to configure crossbar IP address for an interface, user, or location.

When a packet comes in from a user whose address includes an assigned gateway, the PortMaster does not consult the forwarding table but forwards the packet to the gateway address. If a gateway address is not assigned to a range, the range uses the default gateway address of the IP pool. If the IP pool is not assigned a default gateway address, no crossbar IP address is used and the PortMaster consults the forwarding table.

Setting a Named IP Pool Gateway

Use the following command to set a default gateway for the entire named IP pool:

```
Command> set ippool Name default-gateway Gateway
```

Always reset the pool when you make changes to the named IP pool.

```
Command> reset ippool
```

When a packet comes in from a user whose address includes an assigned gateway, the PortMaster does not consult the forwarding table but forwards the packet to the gateway address. If a gateway address is not assigned to a range, the range uses the default gateway address of the IP pool. If the IP pool is not assigned a default gateway address, no crossbar IP address is used and the PortMaster consults the forwarding table.

The default gateway functions as a crossbar IP address. See the *PortMaster 4 Command Line Reference* for details about how to configure crossbar IP for an interface, user, or location.

Setting Named IP Pools in RADIUS

You must modify the RADIUS dictionary to enable named IP pools. You cannot configure the local user table on the PortMaster for named IP pools. To enable named IP pools, add the following line to the RADIUS dictionary:

```
ATTRIBUTE      Ip-Pool-Name      193      string
```

The following example shows a RADIUS user profile using an IP pool named *livermore*:

```
homers   Password = "kwyjibo"
            Service-Type = Framed-User,
            Framed-Protocol = PPP,
            Framed-IP-Address = 255.255.255.254,
            Framed-IP-Netmask = 255.255.255.255,
            Ip-Pool-Name = livermore
```

Setting the Dynamic Host Control Protocol (DHCP) Server

The **set dhcp server** command supports the Cable Modem Telephone Return Interface Specification (CMTRIS) developed by the Multimedia Cable Network System (MCNS) Partners Limited. The CMTRIS solves the problem of limited upstream bandwidth in a cable modem system by providing for the use of a standard telephone interface for upstream traffic. Downstream traffic travels on the coaxial cable.

The specification requires that a cable modem be able to use the telephone interface to request and receive the cable interface address and configuration information via a dynamic host control protocol (DHCP) request.

Use the following command to configure the PortMaster 4 to forward a DHCP request from a cable modem to the DHCP server:

    Command> **set dhcp server** *address*

Note – ComOS does not support DHCP requests over Ethernet—nor requests from a PortMaster 2Ei or a PortMaster ISDN Office Router (OR-U) used for dial-up.

How the Cable Modem Telephone Return System Works

After you set the IP address of the DHCP server on the PortMaster 4, the cable modem dynamically configures itself so that all subsequent data travels upstream via the telephone interface, and downstream on the coaxial cable.

Figure 2-1, using sample IP addresses, illustrates the series of events that begin upon startup and culminate in the dynamic configuration of the cable modem.

Figure 2-1   Cable Modem Telephone Return Interface Startup
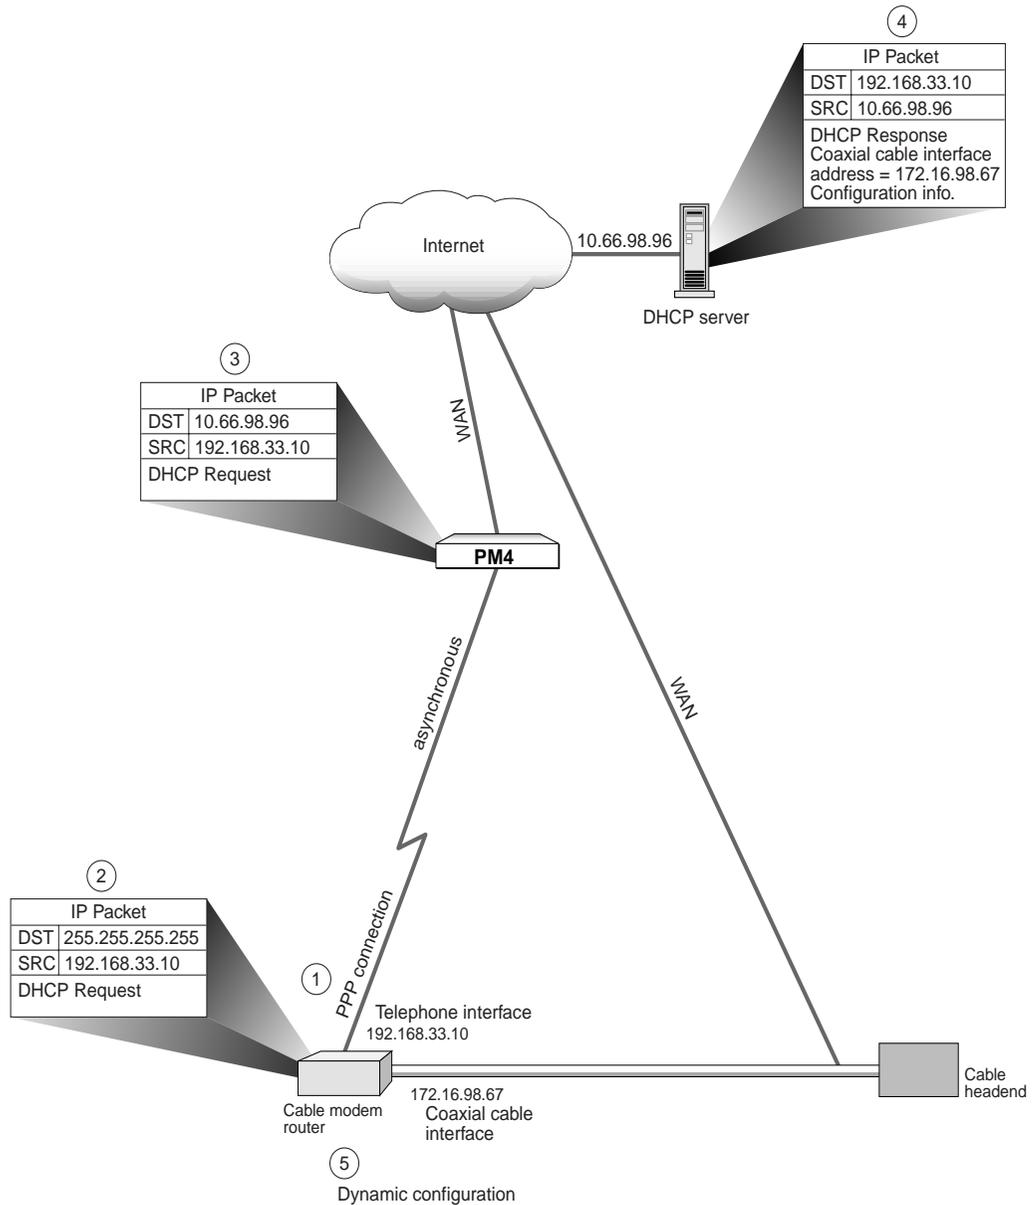
1. Using the telephone interface, the cable modem dials the PortMaster 4 and establishes a PPP connection. The PortMaster 4 assigns IP address 192.168.33.10 to the telephone interface of the cable modem.

2. Using the telephone interface, the cable modem broadcasts a DHCP request. The destination of the request is 255.255.255.255 and the source is 192.168.33.10.

3. The PortMaster 4 forwards the request to the DHCP server by substituting the IP address of the DHCP server (10.66.98.96) for the broadcast destination address.

4. The DHCP server responds with configuration information for the cable modem and an IP address (172.16.98.67) for the coaxial cable interface on the cable modem.

5. Using the configuration information received from the DHCP server, the cable modem dynamically assigns 172.16.98.67 to the cable interface, and configures the cable modem so that upstream IP packets leave the cable modem via the telephone interface with the IP address of the cable interface (172.16.98.67) as the source address. Because packets now carry the source address of the cable interface, response to these packets travels via the coaxial cable.

ComOS does not add routes to its table when forwarding or returning DHCP requests. It transparently forwards and returns DHCP requests from dial-in clients to the specified server.

To view DHCP relaying information, use the **set consol**e command followed by the **set debug 0x81** command. See the *PortMaster Troubleshooting Guide* for debugging information.

To disable DHCP reply information, enter the following command:

    Command> **set dhcp server 0.0.0.0.**

The PortMaster 4 does not forward packets to the address 255.255.255.255.

Displaying the Routing Table

Use the following command to display the IP routing table entries:

    Command> **show routes** [*String*|*Prefix*/*NM*]

You can replace *String* with **ospf** or **bgp** to display only OSPF or BGP routes. Replacing *Prefix*/*NM* with an IP address prefix and netmask displays only routes to that destination. Enter the IP address prefix in dotted decimal format and the netmask as a number from 1 to 32, preceded by a slash—for example, /24. The netmask indicates the number of high-order bits in the IP prefix.

To display the IPX routing table entries, enter the following command:

    Command> **show ipxroutes**

Note – The PortMaster 4 supports the IPX protocol when running ComOS 4.1 or later. IPX is not supported in ComOS 4.0.

The routes appear in the following order:

1. Default route

2. Host routes

3. Network routes

4. Expired routes that are no longer being advertised

Setting Static Routes

Static routes provide routing information unavailable from the Routing Information Protocol (RIP), Open Shortest Path First (OSPF) protocol, or Border Gateway Protocol (BGP). RIP, OSPF, or BGP might not be running for one of the following two reasons.

- Network administrators choose not to run RIP, OSPF, or BGP.

- Hosts connected to the PortMaster do not support RIP, OSPF, or BGP.

Separate static routes tables are maintained for IP and for IPX, which you display with the **show routes** and **show ipxroutes** commands.

You construct a static route table manually on a PortMaster by adding and deleting static routes as described in the following sections. Refer to the *PortMaster Routing Guide* for information about routing and static routes.

Adding and Deleting a Static Route for IP

A static route for IP contains the following items:

- **Destination**—The IP address prefix of the host or the number of the IPX network to which the PortMaster will be routing.

- **Netmask** —The static netmask in use at the destination. See "Modifying the Static Netmask Table" on page 2-18 for more information about netmasks.

- **Gateway**—The address of a locally attached router where packets are sent for forwarding to the destination.

- **Metric**—The number of routers (or hops) a packet must cross to reach its destination. The metric represents the cost of sending the packet through the gateway to the specified destination.

Note – Never set the gateway for the PortMaster to an address on the same PortMaster; the gateway must be on another router.

Use the following commands to add a static route for IP:

```
Command> add route Ipaddress[/NM] Ipaddress(gw) Metric
Command> save all
```

Use the following commands to delete a static route for IP:

```
Command> delete route Ipaddress[/NM] Ipaddress(gw)
Command> save all
```

You can delete only static routes.

Adding and Deleting a Static Route for IPX

A static route for IPX contains the following items:

- **Destination**—The number of the IPX network to which the PortMaster will be routing.

- **Gateway**—The address of a locally attached router where packets are sent for forwarding to the destination.

  For IPX networks, the gateway address consists of 8 hexadecimal digits for the network address, a colon (:) and the node address of the gateway router expressed as 12 hexadecimal digits—for example, 00000002:A0B1C2D3E4F5.

  The IPX node address is usually the media access control (MAC) address on a PortMaster.

- **Metric**—The number of routers (or hops) a packet must cross to reach its destination. The metric represents the cost of sending the packet through the gateway to the specified destination.

- **Ticks**—The time required to send the packet to its destination. Ticks are measured in 50ms increments.The ticks metric is used in addition to the hops metric only on IPX networks.

Note – Never set the gateway for the PortMaster to an address on the same PortMaster; the gateway must be on another router.

Use the following commands to add a static route for IPX:

```
Command> add route Ipxnetwork Ipxaddress Metric Ticks
Command> save all
```

Use the following commands to delete a static route for IPX:

```
Command> delete route Ipxnetwork Ipxaddress
Command> save all
```

Use the following command to set a static default route for all IPX packets not routed by a more specific route:

```
Command> set ipxgateway Network|Node Metric
```

Note – You can delete only static routes.

Modifying the Static Netmask Table

Note – ComOS 4.1 and later releases support both RIP-1 and RIP-2 on the PortMaster 4. Earlier releases of ComOS support only RIP-1.

The netmask table is provided to allow routes advertised by RIP-1 to remain uncollapsed on network boundaries in cases where you want to break a network into noncontiguous subnets. The PortMaster normally collapses routes on network boundaries as described in RFC 1058. However, in certain circumstances where you do not want to collapse routes, the netmask table is available.

Caution – Do not use the static netmask table unless you thoroughly understand and need its function. In most circumstances its use is *not* necessary. Very large routing updates can result from too much use of the netmask table, adversely affecting performance. In most cases it is easier to use RIP-2 or OSPF instead of using the netmask table and RIP-1. Lucent strongly recommends you use OSPF if you require noncontiguous subnets or variable-length subnet masks (VLSMs).

For example, suppose the address of Ether0 is 172.16.1.1 with a 255.255.255.0 subnet mask (a class B address subnetted on 24 bits) and the destination of PTP1 is 192.168.9.65 with a 255.255.255.240 subnet mask (a class C address subnetted on 28 bits). If routing broadcast is on, the PortMaster routing broadcast on Ether0 claims a route to the entire 192.168.9.0 network. Additionally, the broadcast on PTP1 claims a route to 172.16.0.0.

Sometimes, however, you want the PortMaster to collapse routes to some bit boundary, other than the network boundary. In this case, you can use the static netmask table. However, RIP supports only host and network routes, because it has no provision to include a netmask. Therefore, if you set a static netmask in the netmask table, the PortMaster collapses the route to that boundary instead, and broadcasts a host route with that value. Other PortMaster routers with the same static netmask table entry convert the host route back into a subnet route when they receive the RIP packet.

This approach works only if all the routers involved are PortMaster products, with the following two exceptions:

- You use a netmask table entry of 255.255.255.255. In this case, the routes broadcast as host routes really are host routes, so other vendors' routers can use them. Keep in mind that not all routers accept host routes.

- The other vendor's router can convert host routes into subnet routes through some mechanism of its own.

Uses for Static Netmasks

The most common use for the static netmask table is to split a single class C network into eight 30-host subnets for use in assigned pools. Subnetting allows each PortMaster to broadcast a route to the subnet instead of claiming a route to the entire class C network. An example of that use is provided below.

The next most common use for the static netmask table is to allow dial-in users to use specified IP addresses across multiple PortMaster products in situations where assigned IP addresses are not sufficient. This use can result in very large routing tables and is not recommended except where no other alternative is possible.

The netmask table can be accessed only through the command line interface. To add a static netmask, use the **add netmask** command. To delete a static netmask, use the **delete netmask** command. The **show table netmask** command shows both dynamic netmasks and static netmasks, marking them accordingly.

Note – Static routes use the netmask table entries that are in effect when the routes are added. If the netmask table is changed, the static route must be deleted from the route table and added again.

Example of Applying Static Netmasks

Note – Lucent recommends that you use RIP-2 or OSPF in this circumstance instead of static routes.

This static netmask example assumes the following:

- You have anywhere between 8 and 250 PortMaster routers.

- You assign all the user addresses from the dynamic address assignment pools on the PortMaster routers.

- You are using 27-bit subnets of these three class C networks: 192.168.207.0, 192.168.208.0, and 192.168.209.0.

- You are using the 192.168.206.0 network for your Ethernet.

- All PortMaster routers involved are running ComOS 3.1.2 or later.

- You do not use proxy ARP. Instead, you use your 192.168.206.0 network for the Ethernet, and divide your other networks up among the PortMaster routers.

- Each network provides 30 addresses for the assigned pool of each PortMaster.

To create the subnets defined in this example, enter the following commands on all the PortMaster routers:

```
Command> set Ether0 address 192.168.206.X (for some value of X)
Command> set gateway 192.168.206.Y (where Y points at your gateway)
Command> add netmask 192.168.207.0 255.255.255.224
Command> add netmask 192.168.208.0 255.255.255.224
Command> add netmask 192.168.209.0 255.255.255.224
Command> set Ether0 rip on
Command> save all
```

The netmask table collapses routes on the boundaries specified. As a result, if one PortMaster has an assigned pool starting at 192.168.207.33, it broadcasts a host route to 192.168.207.32 instead of broadcasting a route to the 192.168.207.0 network. The other PortMaster routers consult their own netmask tables and convert that route back into a subnet route to 192.168.207.33 through 192.168.207.32.

If your gateway on the Ethernet is not a PortMaster product, the netmask table is not supported. However, you can set a static route on the gateway for each of the three destination networks for your assigned pools (192.168.207.0, 192.168.208.0, and 192.168.209.0), pointing at one of the PortMaster routers. The identified PortMaster then forwards packets to the proper PortMaster.

If you are using an IRX running ComOS 3.2R or later as your gateway, you can configure the netmask table on the router also. This allows your PortMaster to listen to RIP messages from the other PortMaster routers and route directly to each of them.

Setting Authentication for Dial-In Users

You can configure the PortMaster for three authentication methods, PAP, CHAP, and username/password login.

By default, PAP and CHAP are set to **on**. Dial-in users are asked to authenticate with PAP when PPP is detected. If users refuse, they are asked to authenticate with CHAP.

If you set PAP to **off**, and CHAP to **on**, dial-in users are asked to authenticate with CHAP. PAP authentication is neither requested nor accepted. If you set both PAP and CHAP to **off**, dial-in users must authenticate with a username/password login.

To set PAP authentication, use the following command:

    Command> **set pap on|off**

To set CHAP authentication, use the following command:

    Command> **set chap on|off**

Setting Call-Check Authentication

You can enable services without authenticating the user at the point of entry on PortMaster products that support PRI or in-band signaling. Use the **show global** command to find out if call-check is enabled on your PortMaster.

To enable the call-check feature in ComOS, you must first configure call-check user entries on the RADIUS 2.1 server. Otherwise, the PortMaster issues a busy signal to every call. See the *RADIUS for UNIX Administrator's Guide* for more information about RADIUS.

To enable call checking on the PortMaster, use the following command:

    Command> **set call-check on|off**

Note – The call-check feature is **off** by default.

If the call-check feature is **on**, the PortMaster sends a ringing message to the switch while the service information is being looked up in RADIUS.

RADIUS does one of the following:

- Rejects the message with a busy signal

- Acknowledges the call and allows the call to be completed with no special service type determined during the call

- Allows the creation of a netdata clear channel TCP or L2TP connection to the destination specified in the RADIUS user profile

Call-check enables the PortMaster—via RADIUS—to check the telephone number of a caller before answering the call. The PortMaster can then hang up and call the user back with no charge incurred for connecting the user in the first place. Alternatively, the PortMaster can reject the call to limit the number of users who can call a given number, such as an 800 number, or to prevent certain users from calling the number.

You can also use call-check to support virtual points of presence (POPs) by redirecting a call. If a caller dials one number, the PortMaster can authenticate normally. If a caller dials a different number, the PortMaster can accept the call and forward the caller information through a netdata (TCP clear) or L2TP connection to an IP address and port of your choosing, where another process handles the user.

Additionally, you can provide guest access or establish tunnels based on dial number information services. Call checking can be done against the calling number ID (CNID) or calling line ID (CLID) or both. The RADIUS attributes are Called-Station-Id and Calling-Station-Id, respectively.

Setting the ISDN Switch

You can configure the switch provisioning for ISDN PRI connections to PortMaster ISDN ports. See Chapter 10, "Configuring T1, E1, and ISDN PRI," for details on PRI connections.

PortMaster Security Management

The PortMaster provides security through the user table, or if configured, RADIUS security. When a dial-in user attempts to authenticate at the login prompt, or via PAP or CHAP authentication, the PortMaster refers to the entry in the user table that corresponds to the user. If the password entered by the user does not match, the PortMaster denies access with an "Invalid Login" message. If no user table entry exists for the user and port security is off, the PortMaster passes the user on to the host defined for that port using the selected login service. In this situation, the specified host is expected to authenticate the user.

If port security is on and the user was not found in the user table, the PortMaster queries the RADIUS server, if one has been configured. If the username is not found in the user table, port security is on, and no RADIUS server is configured in the global configuration of the PortMaster, access is denied with an "Invalid Login" message. If the RADIUS server is queried and does not respond within 30 seconds (and neither does the alternate RADIUS server), access is denied with an "Invalid Login" message.

If security is off, any username that is not found in the user table is sent to the port's host for authentication and login. If security is on, the user table is checked first. If the username is not found and a RADIUS server is configured, RADIUS is consulted. When you are using RADIUS security, you must use the **set** *C0* **security** command to set security to **on**.

Access can also be denied if the specified login service is unavailable—for example, if the PortMaster Login Service has been selected for the user but the selected host does not have the **in.pmd** PortMaster daemon installed. Access is denied with the "Host Is Currently Unavailable" message if the host is down or otherwise not responding to the login request.

If an access filter is configured on the port and the login host for the user is not permitted by the access filter, the PortMaster refuses service with an "Access Denied" message. If the access override parameter is set on the port, the PortMaster instructs the user to authenticate himself, even though the default access filter is set to deny access.

Refer to the *RADIUS for Windows NT Administrator's Guide* and *RADIUS for UNIX Administrator's Guide* for more information about RADIUS.

Configuring SNMP 3

This chapter describes how to configure SNMP on the PortMaster 4 and includes the following topics:

- "Understanding SNMP" on page 3-1

- "Livingston Extensions" on page 3-5

- "Configuring SNMP" on page 3-12

If you want to configure SNMP and are already familiar with SNMP concepts and the Livingston extensions, go to the "Configuring SNMP" section.

Understanding SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that allows devices to communicate management information. You can configure the PortMaster to provide network and device information via SNMP to a network management system (NMS). You must have NMS software to use SNMP.

SNMP consists of the following parts:

- SNMP agent (provided in ComOS)

- SNMP manager (not provided)

- Management Information Base (MIB) (PortMaster variables provided by ComOS)

SNMP specifies the message format for exchanging information between the SNMP manager and an SNMP agent.

The SNMP agent returns values for Management Information Base (MIB) variables that can be changed or queried by the SNMP manager. The agent gathers information from the MIB, which resides on the target device. MIB information can include device parameters and network status. The agent is capable of responding to requests to get or set data from the manager.

PortMaster products support MIB II variables as specified in RFC 1213, along with a MIB specific to PortMaster products. SNMP management can be enabled for any PortMaster. Lucent Remote Access ships configuration files compatible with various network management packages along with the PMVision software.

PortMaster 4 MIB Information

The Lucent Technologies PortMaster products and PMVision support multiprotocol carrier capacity WAN access. The PortMaster 4 enables public carriers, ISPs, and major network providers to offer a variety of services such as dial-up; V.90, K56flex, or V.34 modems; ISDN, and T1, E1, or T3 leased line connections; and Frame Relay connections.

MIB Specification Overview

The PortMaster 4 MIB conforms to the first version of the Structure of Management Information (SMIv1) (RFCs 1212 and 1213). The private MIB for the product describes the data for configuration, fault, performance, security, and accounting management.

- Chassis MIB

- Configuration management (equipment, physical interfaces, and logical interfaces)

- Fault management (fault detection and fault isolation traps)

- Performance management (interfaces)

- Security management (MIB access control)

- Administration management (read-write community and trap community)

PortMaster products also support MIB II objects, and the enterprise-specific traps provide information about several alarm conditions that can be enabled or disabled. The traps are generated as SNMPv1 traps.

The PortMaster 4 SNMP agent consists of a multilingual master agent (supporting SNMPv1, v2, and v3 protocols) and subagents with AgentX interfaces with the master agent to handle different subtrees in the MIB. The master agent listens on UDP port 161 for SNMP protocol data units (PDUs). The master agent processes the PDUs and forwards them to appropriate subagents. The subgents are responsible for processing the data. The master agent and the subagents can operate asynchronously. The subagents in the SNMP software can dynamically register and unregister to provide and deny access to portions of the MIB.

The MIB module **LE41** specifies the first revision of the private MIB for the PortMaster 4. It includes only the physical equipment—chassis, T1, E1, and T3 lines; modems; serial ports; and Ethernet interfaces.

The **livingston.mib** file is in the SNMP directory of the ComOS software, and on the World Wide Web at **http://www.livingston.com/marketing/products/pmtempl.html**.

Examining the MIB Structure

The entire Management Information Base (MIB) hierarchy can be represented by a tree structure. In this representation, the unnamed "root" of the tree divides into the following main branches:

- Consultative Committee for International Telegraph and Telephone (CCITT)

- International Organization for Standardization (ISO)

- ISO/CCITT

Each branch and subbranch in the tree structure is known as an **object**, and each object is represented by an **object name** and an **object identifier** (OID). Figure 3-1 traces the "path" from the ISO branch of the MIB to the *Livingston* MIB.

OIDs provide compact representations of object names. An OID shows the position of an object in the MIB hierarchy. As shown in Figure 3-1, the OID for the Livingston MIB is 1.3.6.1.4.1.307.

Figure 3-1   Management Information Base (MIB) Hierarchy
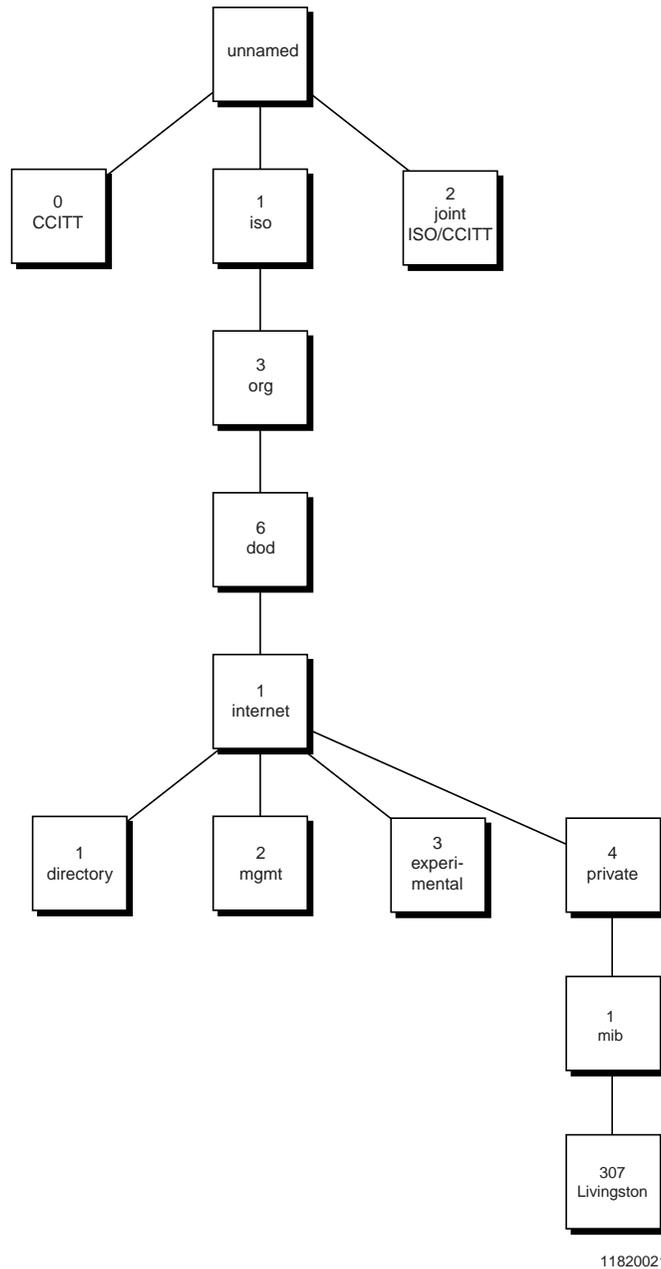
Figure 3-2 shows the tree structure of the private Livingston portion of the MIB.

Figure 3-2   Part of MIB Structure showing PortMaster Port C0


```
                        Livingston Enterprise

      305          306          307          308

 1. (not used)        2. products        3. livingstonMib

                      1. livingstonSystem    2. livingstonInterfaces

                            1. livingstonSerial        2. livingstonT1E1

                      1. livingstonSerialTable

                            1. livingstonSerialEntry

         1. Index    2. PortName    3. PhysType    4. User    ...
                                                           11820020
```
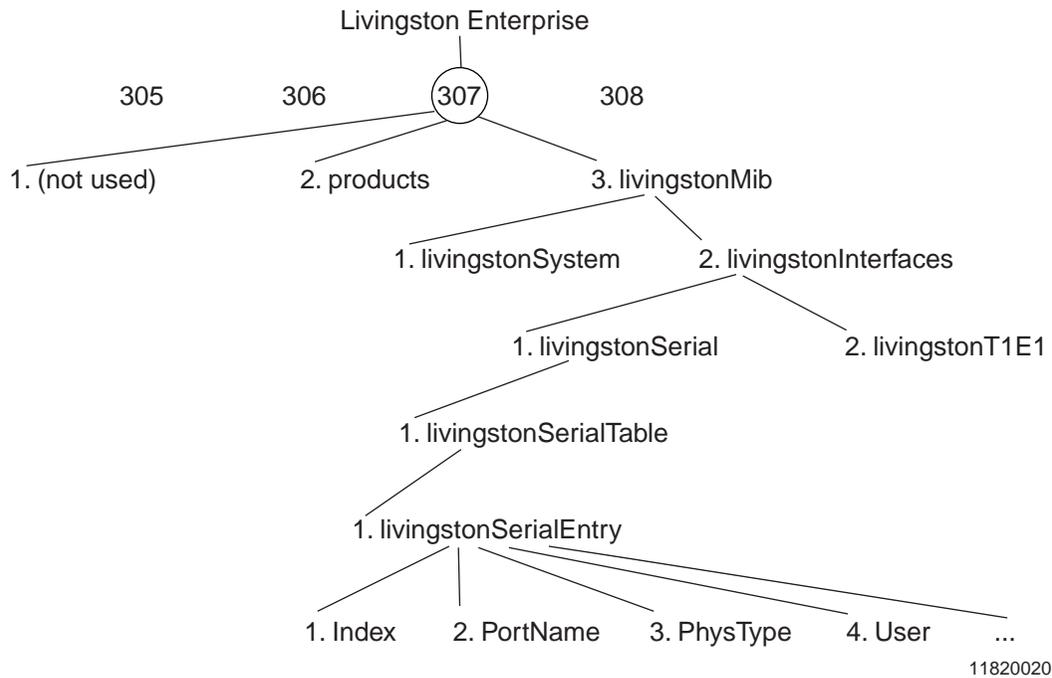
Reading from the top down, the object identifier (OID) in Figure 3-2 (307.3.2.1.1.1.2) breaks out as follows:

- 307 refers to the Livingston namespace.

- 3 refers to the MIB.

- 2 refers to interfaces.

- 1 refers to serial interfaces.

- 1 refers to the serial interfaces table.

- 1 refers to an entry in the serial interfaces table.

- 2 refers to the PortName variable.

The SNMP manager queries the agents by means of OIDs. Each OID uniquely identifies a single MIB variable. For example, the OID 307.3.2.1.1.1.2.0 returns the port name for port C0, and the OID 307.3.2.1.1.1.2.1 returns the port name for port C1 (see Table 3-1).

Table 3-1   Partial View of the Livingston Serial Interfaces Table

OID	Interface (C0 and C1)
...307.3.2.1.1.1.**1**	Index
...307.3.2.1.1.1.**2**	PortName
...307.3.2.1.1.1.**3**	PhysType
...307.3.2.1.1.1.**4**	User

Table 3-1   Partial View of the Livingston Serial Interfaces Table *(Continued)*

OID	Interface (C0 and C1)
...307.3.2.1.1.1.**5**	SessionId
...307.3.2.1.1.1.**6**	Type
...307.3.2.1.1.1.**7**	Direction

Livingston Extensions

This section lists the following tables from the Livingston Extensions section of the MIB:

- "PortMaster Serial Interfaces" on page 3-5

- "PortMaster T1 and E1 Interfaces" on page 3-7

- "PortMaster Internal Modem Table" on page 3-8

- "PortMaster Billing and Accounting Information Table" on page 3-9

- "PortMaster Call Event Status Table" on page 3-10

- "PortMaster Board Call Summary Table" on page 3-11

- "PortMaster Line Call Summary Table" on page 3-11

PortMaster Serial Interfaces

The PortMaster Serial Interfaces table (Table 3-2) in the Livingston Extensions section of the MIB lists all serial interface entries.

Table 3-2   PortMaster Serial Interfaces MIB Table

Object	Definition
livingstonSerialIndex	Unique value for each serial interface.
livingstonSerialPortName	Text string containing the name of the serial interface (for example, C0, W1, and so on).
livingstonSerialPhysType	Type of physical serial interface, distinguished according to the physical or link protocol(s) currently being used on the interface.
livingstonSerialUser	Name of the active user. Blank if not active.
livingstonSerialSessionId	Unique session identifier that matches the RADIUS session ID.
livingstonSerialType	Active type of service being provided by the serial interface.
livingstonSerialDirection	Direction in which the active session was initiated.
livingstonSerialPortStatus	Status of the serial interface.
livingstonSerialStarted	Amount of time this session has been active.

Table 3-2   PortMaster Serial Interfaces MIB Table *(Continued)*

Object	Definition
livingstonSerialIdle	Amount of time this session has been idle.
livingstonSerialInSpeed	Estimate of the current inbound bandwidth in bits per second of the serial interface.
livingstonSerialOutSpeed	Estimate of the current outbound bandwidth in bits per second of the serial interface.
livingstonSerialModemName	Text string containing the name of the digital modem in use by the serial interface.
livingstonSerialIpAddress	IP address associated with the serial interface. When characterizing a network port, this value is the IP address of the remote user. When characterizing a device or login port, this value is the IP address of the host to which the user is connected.
livingstonSerialifDescr	Text string containing information about the network interface bound to the serial interface.
livingstonSerialInOctets	Total number of octets received on the serial interface.
livingstonSerialOutOctets	Total number of octets transmitted on the serial interface.
livingstonSerialQOctets	Total number of octets queued on the serial interface.
livingstonSerialModemStatus	Status of the modem used by the serial interface.
livingstonSerialModemCompression	Compression type being used in the modem or by the serial interface.
livingstonSerialModemProtocol	Error-correcting protocol being used in the modem or by the serial interface.
livingstonSerialModemRetrains	Number of retrains attempted by the modem attached to the serial interface.
livingstonSerialModemRenegotiates	Number of renegotiations attempted by the modem attached to the serial interface.

PortMaster T1 and E1 Interfaces

The PortMaster T1 and E1 Interface table (Table 3-3) in the Livingston Extensions section of the MIB provides configuration and statistics for the T1 and E1 interfaces that connect directly to the telephone company.

Table 3-3   PortMaster T1 and E1 Interfaces MIB Table

Object	Definition
livingstonT1E1Index	Unique value for each T1E1 interface.
livingstonT1E1PhysType	Type of interface (T1 or E1).
livingstonT1E1Function	Configured function of the interface.
livingstonT1E1Status	Current operational state of the interface. Operational states include the following: • Up (1) • Down (2) • Loopback (3)
livingstonT1E1Framing	Configured line framing. Line framing types include the following: • Extended superframe (ESF) (1) • D4 (2) • Cyclic redundancy check (CRC4) (3) • Frame Alignment Signal (FAS) (4)
livingstonT1E1Encoding	Configured line signal encoding.
livingstonT1E1PCM	Configured voice modulation (pulse code modulation).
livingstonT1E1ChangeTime	Amount of time this interface has been up or down.
livingstonT1E1RecvLevel	Estimate of the current receive signal level, in decibels, of the interface.
livingstonT1E1BlueAlarms	Total number of blue alarms on the interface.
livingstonT1E1YellowAlarms	Total number of yellow alarms on the interface.
livingstonT1E1CarrierLoss	Total number of times the interface has lost the carrier signal.
livingstonT1E1SyncLoss	Total number of times the interface has lost frame synchronizations.
livingstonT1E1BipolarErrors	Total number of line code violations detected on the interface.

Table 3-3   PortMaster T1 and E1 Interfaces MIB Table *(Continued)*

Object	Definition
livingstonT1E1CRCErrors	Total number of frame-level CRC errors detected on the interface.
livingstonT1E1SyncErrors	Total number of frame synchronization errors detected on the interface.

PortMaster Internal Modem Table

The PortMaster Internal Modem table (Table 3-4) in the Livingston Extensions section of the MIB lists the objects in the internal modem table.

Table 3-4   PortMaster Internal Modem MIB Table

Object Type	Definition
livingstonModemIndex	Unique value for each modem interface.
livingstonModemPortName	Textual string containing the name of the serial interface (for example, S0, S1, and so on).
livingstonModemStatus	Current state of the modem.
livingstonModemProtocol	Error-correcting protocol being used in the modem.
livingstonModemCompression	Compression being used in the modem interface.
livingstonModemInSpeed	Estimate of the modem interface's current inbound bandwidth in bits per second.
livingstonModemOutSpeed	Estimate of the modem interface's current outbound bandwidth in bits per second.
livingstonModemInByteCount	Total number of bytes received by the modem.
livingstonModemOutByteCount	Total number of bytes transmitted by the modem.
livingstonModemRetrains	Number of retrains attempted by the modem.
livingstonModemRenegotiates	Number of renegotiations attempted by the modem.
livingstonModemCalls	Number of times a call was received by the modem.
livingstonModemDetects	Number of analog calls received by the modem.
livingstonModemConnects	Number of successful calls received by the modem.

PortMaster Billing and Accounting Information Table

The PortMaster Billing and Accounting Information table (Table 3-5) in the Livingston Extensions section of the MIB lists call events that can be used for billing.

Table 3-5   PortMaster Billing and Accounting Information MIB Table

Object	Definition
livingstonAMCEIndex	Index into the call event table. The table stores call events that can be used for billing.
livingstonAMCESessId	Session ID for the current session. This ID must be unique across all the sessions and across reboots.
livingstonAMCETimeStamp	Time stamp for this event in seconds since the last reboot.
livingstonAMCEType	Type of event associated with this entry in the call event table.
livingstonAMCESvcType	The type of service provided to the user. This field is meaningful if the event type is servicechanged(4), or namechanged(5) events. In all other cases, this object must return none(1).
livingstonAMCEUName	Username of the dial-in user. This object returns the valid username when the event type is servicechanged(4) or namechanged(5). In all other cases, it returns a NULL.
livingstonAMCEModemBoard	Board ID for the modem that handled this call. This value can be used to diagnose modem-related problems (dropping the call, retraining too frequently, and so on).
livingstonAMCEModemID	ID of the internal modem that handled this call. This object can be used to diagnose modem-related problems.
livingstonAMCEModemPort	Serial interface (S0, S1) on which the call was received.
livingstonAMCEModemName	Name of the modem interface (for example, M0...M95).
livingstonAMCEDataRate	Speed of this connection. Speed is specified as baud rate for modem calls and a receive data rate for ISDN calls. This object returns a 0 for call answered and call cleared events.
livingstonAMCECallingPartyID	Calling party ID. This object is valid only for call answered, call originated, and call cleared events. For all invalid event types, this object is set to NULL.

Table 3-5 PortMaster Billing and Accounting Information MIB Table *(Continued)*

Object	Definition
livingstonAMCEInOctets	Total octets received during this call. This object is cleared at the end of each call.
livingstonAMCEOutOctets	Total octets sent out during this call. This object is cleared at the end of each call.
livingstonAMCECallCharge	Call charge for this call. This object is valid only when the event is call cleared. For all other events this object is set to zero (0).
livingstonAMCEDisconnReason	Reason for the disconnection.

PortMaster Call Event Status Table

The PortMaster Call Event Status table (Table 3-6) in the Livingston Extensions section of the MIB lists call events that can be queried for call status on a particular modem port, and the action that can be taken to terminate the call.

Table 3-6 PortMaster Call Event Status MIB Table

Object	Definition
livingstonAMPortVTSSsnId	Session ID used by the VTS table to index and query the status of the call on a given modem port. This table can also be used to take appropriate action to terminate the session.
livingstonAMPortVTSModemBoard	Specifies the modem board number for the given session ID handling the call.
livingstonAMPortVTSModemId	Specifies the modem ID (0, 1, ...95) for the given session ID handling the call.
livingstonAMPortVTSModemName	Specifies the modem name (M0...M95) for the given session ID handling the call.
livingstonAMPortVTSSerialPort	Specifies the serial port number (S0...S95)for the given session ID handling the call.
livingstonAMPortVTSSvcType	Specifies the service type for the given session.
livingstonAMPortVTSUName	Username of the dial-in user for the given session. If the session is terminated, it returns a NULL.
livingstonAMPortVTSCallStatus	Port status. If the port is currently handling a call, it is set to active(2); if the call on this port is terminated, it is set to terminated(3). If the session ID does not match the session ID for the current call, this object is set to unknown(1).

Table 3-6   PortMaster Call Event Status MIB Table *(Continued)*

Object	Definition
livingstonAMPortVTSTerminateCall	When set to any value, this object terminates the call on the corresponding modem port.

PortMaster Board Call Summary Table

The PortMaster Board Call Summary table (Table 3-7) from the Livingston Extensions section of the MIB contains a summary of calls on a per board basis. The rows in the table correspond to the slots in the PortMaster 4, and the columns specify the type of calls as V.90, V.34, ISDN, and so on. This object is not accessible.

Table 3-7   PortMaster Board Call Summary MIB Table

Object	Definition
livingstonPMBrdCallSumBrdId	Board ID used as an index into the call summary table. The valid board IDs are the numbers of the physical slots that hold T1 or E1 boards—0 through 9 except for 4, which is reserved for the manager module.
livingstonPMBrdCallSumCapacity	Capacity of this T1 or E1 board.
livingstonPMBrdCallSumIsdnCalls	Current total of all ISDN calls handled by this T1 or E1 board.
livingstonPMBrdCallSumV90Calls	Current total of all V.90, K56flex and 56Kbps calls handled by this T or E1 board.
livingstonPMBrdCallSumV34Calls	Current total of all V.34, 33.6Kbps, and 28.8Kbps calls handled by this T1 or E1 board.
livingstonPMBrdCallSumOther	Current total of all other types of calls not handled by the other objects in this table.

PortMaster Line Call Summary Table

The PortMaster Line Call Summary table (Table 3-8) from the Livingston Extensions section of the MIB contains a summary of calls on a per line basis. The rows in the table correspond to the lines, and the columns specify the type of calls as V.90, V.34, ISDN, and so on. This object is not accessible.

Table 3-8   PortMaster Line Call Summary MIB Table

Object	Definition
livingstonPMT1E1CallSumIfId	Index into the call summary table. The valid line IDs are the T1 or E1 lines, which can range from 1 through 36 for T1 or 1 through 27 for E1.

Table 3-8   PortMaster Line Call Summary MIB Table *(Continued)*

Object	Definition
livingstonPMT1E1CallSumCapacity	Sum of all types of calls handled by this T1 or E1 line.
livingstonPMT1E1CallSumV90Calls	Sum of all V.90, K56flex and 56Kbps calls handled by this T1 or E1 line.
livingstonPMT1E1CallSumV34Calls	Sum of all V.34, 33.6Kbps, and 28.8Kbps calls handled by this T1 or E1 line.
livingstonPMT1E1CallSumOther	Sum of all other types of calls not handled by the other objects in this table.

Configuring SNMP

The rest of this chapter describes how to configure SNMP using the command line interface, and includes the following topics:

- "Setting SNMP Monitoring" on page 3-12

- "Setting SNMP Read and Write Community Strings" on page 3-12

- "Adding SNMP Read and Write Hosts" on page 3-13

- "Viewing SNMP Settings" on page 3-13

- "Monitoring SNMP Alarms" on page 3-14

Setting SNMP Monitoring

Simple Network Management protocol (SNMP) monitoring is used to set and collect information on SNMP-capable devices. This feature is most often used to monitor network statistics such as usage and error rate.

If SNMP monitoring is on, the PortMaster accepts SNMP queries. If SNMP monitoring is off, all SNMP queries are ignored.

To turn SNMP monitoring on or off, use the following commands:

```
Command> set snmp on|off
Command> save all
Command> reboot
```

Setting SNMP Read and Write Community Strings

Community strings allow you to control access to the MIB information on selected SNMP devices. The read and write community strings act like passwords to permit access to the SNMP agent's information. Every device allowed to access or read the MIB information must know the read community string. The default read community string is **public**. Before information can be set on the SNMP agent, the write community

string must be known by the device. The default write community string is **private**. Community strings must be set on SNMP agents so that configuration information is not changed by unauthorized users.

To use this feature, you must set **both** a read community string and a write community string for your network.

To set SNMP read and write community strings, use the following command:

> Command> **set snmp readcommunity|writecommunity** *String*

Note – Use of the default write community string—**private**—is strongly discouraged. Because it is the default, it is known to all users and therefore provides no security. Use another value for the write community string.

Adding SNMP Read and Write Hosts

PortMaster products allow you to control SNMP security by specifying the IP addresses of the hosts that are allowed to access SNMP information. The specification of read and write hosts allows another level of security beyond the community strings. If SNMP hosts are specified, each host attempting to access SNMP information must not only possess the correct community string, it must also be on the read or write host list. This additional level of security allows only authorized SNMP managers to access or change sensitive MIB information.

You can also specify a list of hosts allowed to read or write SNMP information. You can permit all hosts or you can deny all hosts.

Note – Permitting all hosts to read and write SNMP information can compromise security and is not recommended.

To add SNMP read and write hosts, use the following command:

> Command> **add snmphost reader|writer any|none|**_Ipaddress_

To delete read and write hosts, use the following command:

> Command> **delete snmphost reader|writer** *Ipaddress*

Viewing SNMP Settings

Settings for SNMP monitoring, read and write community strings, and read and write hosts are stored in the SNMP table.

To display the SNMP table, enter the following command:

> Command> **show table snmp**

Monitoring SNMP Alarms

When an interface or modem fails, the SNMP agent traps the error message generated by the failure and sends it to the SNMP manager.

To view the status of failed modems or interfaces from the command line interface, enter the following command:

    Command> **show alarms**

The output of this command lists alarm messages and associated alarm identification numbers. For details about a specific alarm, enter the following command:

    Command> **show alarm** [*alarm-id*]

To clear alarms from the SNMP alarm table, enter the following command:

    Command> **clear alarm** *alarm-id*|**all**

Refer to the *PortMaster 4 Command Line Reference* for more information.

Configuring an Ethernet Interface   4

This chapter describes how to configure Ethernet interfaces on the PortMaster 4 and includes the following topics:

- "Overview of PortMaster 4 Ethernet Interfaces" on page 4-1

- "Setting General Ethernet Parameters" on page 4-3

- "Setting Ethernet IP Parameters" on page 4-4

- "Setting Ethernet IPX Parameters" on page 4-5

- "Configuring Ethernet Subinterfaces" on page 4-7

- "Configuring Standalone Ethernet Boards" on page 4-8

- "Setting OSPF on an Ethernet Interface" on page 4-10

Before configuring an Ethernet interface, you must make the appropriate Ethernet connections for your needs. Refer to the *PortMaster 4 Installation Guide* for information about installing the system manager module and standalone Ethernet boards, and connecting Ethernet interfaces.

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Overview of PortMaster 4 Ethernet Interfaces

The PortMaster 4 supports two Ethernet interfaces on the manager module: **Ether0** and **Ether1**. Each interface has its own media access control (MAC) address and is fully routeable. The 10/100BaseT Ether1 interface has alternative RJ-45 and media-independent interface (MII) connections. Lucent recommends that you configure Ether1 if you configure only one Ethernet interface. If you configure both, you must connect them to separate Ethernet segments.

Understanding Ether0

Ether0 operates at 10Mbps and is physically on the manager board. Use Ether0 for netboots and SNMP. Ether0 supports subinterfaces (see "Configuring Ethernet Subinterfaces" on page 4-7).

Whenever you make changes to the Ether0 interface, you must reboot the PortMaster 4 for the changes to take effect.

Understanding Ether1

The Ethernet board (Ether1) in the manager module is accessed in logical slot 10 and gets its power directly from the manager board. Ether1 can operate at 10Mbps or 100Mbps full duplex. Ether1 is physically on the Ethernet board and communicates with the manager board over the passive ATM backplane.

Ether1 is supported by two CPUs. One CPU processes inbound data, the other processes outbound data. Ether1 does not shut down in a low power situation or due to overheating. Ether1 maintains its own forwarding table, which it learns from the manager board. You cannot configure Ethernet subinterfaces on Ether1.

Whenever you make changes to the Ether1 interface, you must reset it for the changes to take effect. Because Ether1 resides in logical slot 10, you reset the Ether1 interface with the following command:

```
Command> reset slot10
```

Resetting slot 10 reboots the Ethernet board connected to the manager board in slot 4.

During PPP negotiations for the IP Control Protocol (IPCP), the PortMaster 4 uses the following order of precedence when choosing an IP address to identify itself:

1. The Local IP address configured in the user profile, if set

2. The global reported IP address, if set

3. The first global local IP address, if set

4. The second global local IP address, if set

5. The third global local IP address, if set

6. The fourth global local IP address, if set

7. The IP address of Ether1

8. The IP address of Ether0

Note – RADIUS packets leaving the PortMaster 4 have the source IP address of Ether1, even if the packet exits through Ether0.

Understanding the Interfaces on the Standalone Ethernet Boards

The 10Mbps or 100Mbps full-duplex Ethernet interfaces on standalone Ethernet boards are identified by a numbering scheme that refers to the slot in which the board is installed. The single-interface board can be installed in any slot except slot 4. A single-interface board installed in slot 3, for example, is designated **Ether30**. If the board is installed in slot 5, it is designated **Ether50**.

The dual-interface Ethernet board can be installed in slot 3 only, and the two interfaces on the board are always **Ether30** and **Ether31**. See "Configuring Standalone Ethernet Boards" on page 4-8 for more information.

Setting General Ethernet Parameters

The commands described in this section allow you to configure an Ethernet interface. In addition to specifying the protocol type (IP, IPX, or both) and address, you must specify any routing and filtering you want on the Ethernet interface.

This section describes the general Ethernet settings that apply to your network regardless of the protocol you use.

Setting the View

Because the Ethernet interfaces on a PortMaster 4 are numbered uniquely, you can configure them from any view.

Configuring RIP Routing

As described in the *PortMaster Routing Guide,* PortMaster products automatically send and accept route information as RIP messages.

Note – ComOS 4.1 and later releases support both RIP-1 and RIP-2 on the PortMaster 4. Earlier releases of ComOS support only RIP-1.

To configure RIP routing, use the following command:

```
Command> set Ether0 rip on|off|broadcast|listen|v2
{broadcast|multicast|on|v1-compatibility}
```

Refer to the PortMaster 4 Command Line Reference for a description of the keywords in this command. Refer to the *PortMaster Routing Guide* for a discussion of routing with RIP, and for OSPF and BGP routing configuration instructions.

Applying Filters

Filters enable you to control network traffic. After you have created filters in the filter table, you can apply them to the Ethernet interface as either input or output filters. For more information about filters, see Chapter 8, "Configuring Filters."

Filters applied to the Ethernet interface take effect immediately. If you change the filter, the change will not take effect until you set the filter on the interface again or you reboot the PortMaster.

Input Filters

When an input filter is used, all traffic coming into the PortMaster on the Ethernet interface is compared to the input filter rules. Only packets permitted by the filter rules are accepted by the PortMaster.

To apply an input filter to the Ethernet interface, use the following command:

```
Command> set Ether0 ifilter Filtername
```

To remove the input filter, omit the filter name when entering the command.

Output Filters

When an output filter is used, all traffic going out of the PortMaster on the Ethernet interface is compared to the output filter rules. Only packets permitted by the filter rules are sent by the PortMaster.

Note – ICMP and UDP packets generated by the PortMaster are never blocked by the output filter.

To apply an output filter to the Ethernet interface, use the following command:

    Command> **set** *Ether0* **ofilter** *Filtername*

To remove the output filter, omit the filter name when entering the command.

Setting Ethernet IP Parameters

This section describes the IP commands, keywords, and values that must be entered for IP protocol support.

Setting the IP Address

During the PortMaster installation process, you set the IP address for the Ether0 and Ether1 interfaces. If you have one or more standalone Ethernet boards installed, you must configure an IP address and netmask and set broadcast on the Ethernet interfaces on those boards as well. See "Configuring Standalone Ethernet Boards" on page 4-8.

To set or change the IP address of an Ethernet interface, use the following command:

    Command> **set** *Ether0* **address** *Ipaddress*

Note – If you change the IP address of an Ethernet interface, you must reboot the PortMaster for the change to take effect.

Setting the Subnet Mask

The default subnet mask is 255.255.255.0. If you have divided your network into subnets, enter the subnet mask that identifies how your network addresses are divided between the network portion and the host portion.

To set the subnet mask, use the following command:

    Command> **set** *Ether0* **netmask** *Ipmask*

See Appendix A, "Networking Concepts," for more information about using subnet masks.

Setting the Broadcast Address

You can define the IP address used as the local broadcast address. The RIP routing protocol uses this address to send information to other hosts on the local Ethernet network. The actual broadcast address is constructed from the IP address of the Ethernet

interface and the netmask. The two valid values are **high**, where the host part of the address is all 1s (ones), such as 192.168.1.255, and **low**, where the host part of the address is all 0s (zeros), such as 192.168.1.0. The PortMaster default is **low**. The standard for hosts is to broadcast high, but some hosts still use the low broadcast address, including hosts running SunOS 4.*x* (Solaris 1.*x*) and earlier.

The broadcast address you set for an Ethernet interface on the PortMaster must match the broadcast address set for other hosts on your local Ethernet segment.

To set the broadcast address, use the following command:

    Command> **set** *Ether0* **broadcast high|low**

Enabling or Disabling IP Traffic

IP traffic is sent and received through a PortMaster Ethernet interface. IP is enabled by default on PortMaster Ethernet ports. If the setting has been changed, you must enable IP on the Ethernet interface of all PortMaster products attached directly to a local Ethernet. Disable IP traffic on Ethernet ports only if the PortMaster is not attached to a local Ethernet network.

To enable or disable IP traffic, use the following command:

    Command> **set ether0 ip enable|disable**

Note – This command is currently available only on the Ether0 port.

Setting Ethernet IPX Parameters

Note – The PortMaster 4 supports the IPX protocol if it is running ComOS 4.1 or later. IPX is not supported in ComOS 4.0.

You must set the following values to send IPX traffic on an Ethernet interface. IPX routing is enabled when routing is enabled.

- Network address

- Protocol

- Frame type

Setting the IPX Network Address

You must identify the IPX network of your local Ethernet segment. An IPX network address is a number entered in hexadecimal format, described in Appendix A, "Networking Concepts."

To set the IPX network address, use the following command:

    Command> **set** *Ether0* **ipxnet** *Ipxnetwork*

Note – If you change the IPX network address of an Ethernet interface, you must reboot the PortMaster for the change to take effect.

Enabling or Disabling IPX Traffic

Ethernet IPX traffic is sent and received through the PortMaster Ethernet interface. You can enable IPX on the Ethernet interface of any PortMaster products attached directly to a local Ethernet. Disable IPX traffic on Ether0 only if the PortMaster is not attached to a local Ethernet network.

To enable or disable IPX traffic, use the following command:

```
Command> set ether0 ipx enable|disable
```

Note – This command is available only on the Ether0 port.

Setting the IPX Frame Type

The IPX frame type must be identified and set to the value used on the local IPX network. The frame type identifies the encapsulation method used on your IPX ports. The IPX protocol can be implemented with one of the four commonly used IPX encapsulation and frame types shown in Table 4-1.

Table 4-1   Novell IPX Encapsulation and Frame Types

IPX Frame Type	Encapsulation
Ethernet_802.2	Consists of a standard 802.3 media access control (MAC) header followed by an 802.2 Logical Link Control (LLC) header. This is the default encapsulation used by Novell NetWare 4.0.
Ethernet_802.2_II	Not commonly used.
Ethernet_802.3	Consists of a standard 802.3 MAC header followed directly by the IPX header with a checksum of FFFF. This is the default encapsulation used by Novell NetWare 3.11.
Ethernet_II	Uses Novell's Ethernet_II and is sometimes used for networks that handle both TCP/IP and IPX traffic.

The encapsulation method and frame type were selected when your IPX network servers were installed. The IPX frame type you set on the PortMaster must match the frame type set for your network. Contact your IPX network administrator for information about the frame type used on your network.

To set the IPX frame type, use the following command—entered on one line:

```
Command> set Ether0 ipxframe
ethernet_802.2|ethernet_802.2_ii|ethernet_802.3|ethernet_ii
```

Configuring Ethernet Subinterfaces

With the subinterface feature of ComOS, you can create up to 512 subinterfaces (the total number of interfaces available on a PortMaster) on the Ether0 interface on the PortMaster 4. Because you have the bandwidth of only a single Ethernet interface, however, efficiency begins to degrade significantly when you add more than eight subinterfaces.

Note – The PortMaster 4 supports Ethernet subinterfaces only on Ether0.

Subinterfacing is essentially the segmenting of a single wire, or port, into multiple IP networks. Instead of subnetting and routing, you can create a subinterface and then set it up as you would a standard Ethernet interface. To avoid routing loops, however, you must be sure not to create two subinterfaces in the same TCP/IP network on the same port. Each Ethernet subinterface must have a unique network.

A drawback to subinterfacing is that it supports static routing only; IPX, RIP, OSPF, packet filtering, and route propagation are not supported on subinterfaces.

You must configure the primary Ethernet interface before adding subinterfaces. (See "Setting General Ethernet Parameters" on page 4-3 for details.) After you configure the primary Ethernet interface, follow this procedure to add a subinterface.

1. **Create a subinterface.**

   Command> **add subinterface** *Name*

   This command adds an entry to the subinterface table, which you can then view with the **show table subinterface** command. Remove a subinterface from the subinterface table with the **delete subinterface** command.

2. **Associate the subinterface with a physical port.**

   Command> **set subinterface** *Name* **port** *Portlabel*

3. **Assign an IP address or an IP address and netmask to the subinterface.**

   Command> **set subinterface** *Name Ipaddress* [*/NM*]|[*Ipaddress/NM*]

   You can specify the netmask in the */NM* or dotted decimal format. You can also configure the IP address and netmask separately (see the *PortMaster 4 Command Line Reference* for details*).*

4. **Set the broadcast for the interface.**

   Command> **set subinterface** *Name* **broadcast high**|**low**

5. **Save the setting to nonvolatile RAM, and reset the interface.**

   Command> **save all**
   Command> **reset slot10**

Because Ethernet subinterfaces are rebuilt every time a new subinterface is added, you can view but not modify an Ethernet subinterface using the **ifconfig** command (see the *PortMaster 4 Command Line Reference*).

Configuring Standalone Ethernet Boards

This section assumes you have installed a standalone single-interface Ethernet board or a dual-interface Ethernet module as described in the *PortMaster 4 Installation Guide*.

Interface Numbering

The 10/100BaseT interfaces on a standalone Ethenet board or module have two-digit numbers that correspond to the slot in which they are installed and the Ethernet port (Ether0 or Ether1) for that board or module.

*   On a dual-interface Ethernet module, the interfaces are always numbered **Ether30** and **Ether31** because the module must be installed in slot 3.

    Although physically installed in slot 3, the Ether31 interface is monitored and reset through virtual slot 11.

*   On a single-interface Ethernet board, the interface can have any of the following numbers because this board can be installed in any slot except slot 4: **ether00**, **Ether10**, **Ether20**, **Ether30**, **Ether50**, **Ether60**, **Ether70**, **Ether80**, or **Ether90**.

Note – The Ethernet interfaces on the manager module are always labeled Ether0 and Ether1.

Before You Begin

Before a standalone Ethernet board can function, you must configure an Ethernet interface on the manager module. Configure Ether1 (or Ether0—see "Overview of PortMaster 4 Ethernet Interfaces" on page 4-1) with an IP address and reset the slot of the Ethernet board to make configuration changes take effect. Because Ether1 is in logical slot 10, use the following command to reset the Ether1:

    Command> **reset slot10**

Setting the View

To configure a standalone Ethernet board, you must first set the view to the slot the board is installed in. If you are not sure what slot the boards resides in, use the **show boards** command to locate it and to verify that it is properly installed. The ID number (the number in the far left column) is the same as the slot number.

When you have determined the correct slot, set the view to that slot with the following command:

    Command> **set view** *Slotnumber*

The dual-interface Ethernet module is always installed in slot 3.

You can now configure the standalone Ethernet board as you would configure a regular Ethernet interface, being careful to replace *Ether0* in each command with the appropriate Ethernet interface number (see "Interface Numbering" on page 4-8). See "Setting General Ethernet Parameters" on page 4-3 for configuration guidelines.

Note – Ether0 or Ether1 must be configured for the PortMaster 4 to function normally.

IPCP Negotiation

During PPP negotiations for the IP Control Protocol (IPCP), the PortMaster 4 uses the following order of precedence when choosing an IP address to identify itself:

1. The local IP address configured in the user profile, if set

2. The global reported IP address, if set

3. The first global local IP address, if set

4. The second global local IP address, if set

5. The third global local IP address, if set

6. The fourth global local IP address, if set

7. The IP address of Ether1

8. The IP address of Ether0

Main IP Address

When the PortMaster creates an IP packet, it must identify itself by placing a source address in the IP header. To do so, the PortMaster chooses either the main IP address or the nearest IP address, depending on the service used. The main IP address is chosen in the following order, but the nearest IP address is the IP address of the interface on which the packet exits the PortMaster 4:

1. The first global local IP address, if set

2. The second global local IP address, if set

3. The third global local IP address, if set

4. The fourth global local IP address, if set

5. The IP address of Ether1

6. The IP address of Ether0

The following services use the main IP address:

- **syslog**

- **traceroute**

- **telnet**

- DNS

- RADIUS authentication and accounting

- ChoiceNet

The following services use the nearest IP address:

- **ping**

- OSPF

- RIP

- **rlogin**

The global local IP address settings can be displayed with the **show global** and **show routes** commands.

You specify the IP address that BGP uses with the **set bgp peer** command. See the *PortMaster 4 Command Line Reference* for details. The source address you set with this command is the interface address BGP uses when forming its packets.

Setting OSPF on an Ethernet Interface

You can enable or disable Open Shortest Path First (OSPF) routing protocol on an Ethernet interface.

To set OSPF on the interface, use the following command—entered all on one line:

    Command> **set** *Ether0* **ospf on|off** [**cost** *Number*] [**hello-interval** *Seconds*]
    [**dead-time** *Seconds*]

The **on** keyword enables OSPF on the specified Ethernet interface; **off** disables OSPF on that interface.

You can specify the cost of sending a packet on the interface with a link state metric by using the **cost** *Number* keyword and value. The *Number* metric is a 16-bit number between 1 and 65535; the default is 1. Refer to the *PortMaster Routing Guide* for more information about OSPF routing.

Routers in OSPF networks continually exchange hello packets with their neighbor routers. You can set the interval that elapses between the transmission of hello packets on the interface by using the **hello-interval** *Seconds* keyword and value. *Seconds* can range from 10 to 120 seconds; the default is 10 seconds.

If the PortMaster stops receiving hello packets from a neighbor, it treats that router as inactive, or down. You can specify how long the PortMaster waits for hello packets from neighbors by using the **dead-time** *Seconds* keyword and value. *Seconds* can range from 40 to 1200 seconds; the default is 40 seconds.

Note – You must set the same **cost** value, the same **hello-interval** value, and the same **dead-time** value on all routers attached to a common network.

To enable acceptance of RIP packets on the OSPF network, use the following command:

    Command> **set** *Ether0* **ospf accept-rip on|off**

See the *PortMaster Routing Guide* for more information about OSPF.

Configuring Dial-In Users   5

This chapter describes how to configure the PortMaster 4 user table to support dial-in connections. The user table settings define how each dial-in user is authenticated and how dial-in connections are made.

To configure network dial-in connections from other routers, you must define each remote router as a user on the PortMaster.

If you are using RADIUS, you must configure user attributes in individual user files in the RADIUS user database rather than in the PortMaster user table. Refer to the *RADIUS for Windows NT Administrator's Guide* and *RADIUS for UNIX Administrator's Guide* for more information.

This chapter discusses the following topics:

- "Configuring the User Table" on page 5-1

- "User Types" on page 5-2

- "Configuring Settings for Network and Login Users" on page 5-3

- "Configuring Network Users" on page 5-4

- "Configuring Login Users" on page 5-8

Note – Only 100 to 200 users can be configured in the user table and stored in the nonvolatile memory of the PortMaster. Therefore, use RADIUS for user authentication when you must configure multiple PortMaster products to handle more than a few dozen users.

See the *PortMaster 4 Command Line Reference*, the *RADIUS for Windows NT Administrator's Guide,* and *RADIUS for UNIX Administrator's Guide* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Configuring the User Table

This section describes how to display user information and how to add users to or delete them from the user table.

Displaying User Information

You can display the current users in the user table or the complete configuration information for a specified user.

To display the current users in the user table, for example, enter the following command:

```
Command> show table user
Name            Type            Address/Host    Netmask/Service RIP
-----------------------------------------------------------------------
jozef           Netuser         negotiated      0000000000
adele           Login User      default         Telnet
elena           Netuser         assigned        255.255.255.255 No
taffy           Login User      defaults        PortMaster
john            Netuser         192.168.7.8     0000000000      No
```

To display configuration information for a particular user, for example, use the following command:

```
Command> show user elena
Username:       elena           Type:           Dial-in Network User
Address:        Assigned        Netmask:        255.255.255.255
Protocol:       PPP             Options:        Quiet, compressed
MTU:            1500            Async Map:      00000000
```

Adding Users to the User Table

You must add users to the user table before configuring any settings for them. The username is a string of from 1 to 8 printable, nonspace ASCII characters. The optional user password is a string of from 0 to 16 printable ASCII characters. You cannot add users with blank usernames.

To add a login user to the user table, use the following command:

Command> **add user** *Username* [**password** *Password*]

To add a network user to the user table, use the following command:

Command> **add netuser** *Username* [**password** *Password*]

Note – To add a network user, you must use the **netuser** keyword. Thereafter, you can use either the **netuser** or the **user** keyword to configure settings for the network user. You must always use the **user** keyword when configuring login users.

Deleting Users from the User Table

To delete a user from the user table, use the following command:

Command> **delete user** *Username*

User Types

User settings define the nature and behavior of dial-in users. The user table contains entries for each defined dial-in user along with the characteristics for the user.

The user table provides login security for users to establish login sessions or network dial-in connections. If you want to allow a network dial-in connection from another router, the router must have an entry in the user table or in RADIUS.

PortMaster products allow you to configure two types of users, network users and login users.

Network Users

Network users dial in to an asynchronous serial, synchronous serial, or ISDN port on the PortMaster. A connection is established as soon as the user logs in. A PPP or SLIP (on asynchronous ports) session is started. This type of connection can be used for dial-in users or for other routers that need to access and transfer data from the network. Define this type of user when network packets must be sent through the connection.

Login Users

Login users are allowed to establish PortMaster (**in.pmd**), **rlogin**, Telnet, or **netdata** (TCP clear) connections through an asynchronous serial or ISDN port. A connection is established to the specified host as soon as the user logs in. This type of connection is useful for users who need to access an account on a host running TCP/IP.

Configuring Settings for Network and Login Users

The following settings can be configured for either network or login users.

Setting a Password

To set a password for either a login or network user, use the following command:

	Command> **set user** *Username* **password** *Password*

The password can contain between 0 and 16 printable ASCII characters.

Setting the Idle Timer

The idle timer defines the number of minutes or seconds the line can be idle—in both directions—before the PortMaster disconnects the user. You can set the idle time in seconds or minutes, with any value between 2 and 240. The default setting is 0 minutes. The idle timer is not reset by RIP, keepalive, or SAP packets.

To set the idle timer, use the following command:

	Command> **set user** *Username* **idle** *Number* [**minutes**|**seconds**]

To disable the idle timer, set the time to 0 minutes.

Setting the Session Limit

You can define the maximum length of a session permitted before the PortMaster disconnects the user. The session length can be set to between 0 and 240 minutes.

To set the session limit, use the following command:

Command> **set user** *Username* **session-limit** *Minutes*

To disable the session limit, set the time to 0.

Configuring Network Users

Network users establish PPP or SLIP connections with the network as soon as they have been authenticated.

Setting the Protocol

You can set the network protocol for the network user to PPP, SLIP, or X.75. Select a protocol that is compatible with the rest of your network configuration and the user's capabilities.

To set the network protocol for a network user, use the following command:

Command> **set user** *Username* **protocol slip|ppp|x75-sync**

If you set a nonzero IP address for the user, IP is automatically routed. If you set a nonzero IPX network number for the user, IPX is automatically routed.

Do not set a value of all 0s (zeros) or all Fs for the IPX network number.

Setting the User IP Address

You must define the IP address or hostname of the remote host or router. Table 5-1 describes three different ways that the user IP address can be determined.

Table 5-1   User IP Address Options

IP Address Type	Description
assigned	This option allows the PortMaster to assign a temporary IP address that is used for the current session only. The address used comes from a pool of addresses set up during global configuration.
	This method for assigning IP addresses to users is most commonly used when a large number of users are authorized to dial in.
negotiated	This option is used only for PPP sessions. Here, the PortMaster learns the IP address of the remote host using IPCP negotiation.

Table 5-1 User IP Address Options *(Continued)*

IP Address Type	Description
Ipaddress	This option allows you to define a specific IP address for the remote host or router. This method for assigning an IP address to a user is most commonly used for routers that establish a connection with the PortMaster.

To set the user IP address for a normal network user, use the following command:

> Command> **set user** *Username* **address**|**destination assigned**|**negotiated**|*Ipaddress*

The **address** and **destination** keywords are synonymous.

Setting the Subnet Mask

Do not set a subnet mask for a network user unless the user is routed to another network from your network. In that case, set the subnet mask to 255.255.255.255.

To set the subnet mask, use the following command:

> Command> **set user** *Username* **netmask** *Ipmask*

Setting the IPX Network Number

Note – The PortMaster 4 supports the IPX protocol if it is running ComOS 4.1 or later. IPX is not supported in ComOS 4.0.

If you are using the IPX protocol for this user, you must assign a unique IPX number to the network connection between the remote user device and the PortMaster. Each user's connection requires a different IPX network number. If you use **ffffffe** as the IPX network number, the PortMaster assigns the user an IPX network number based on an IP address from the IP address pool.

Note – Do not set a value of all 0s (zeros) or all Fs for the IPX network number.

To set the IPX network number, use the following command:

> Command> **set user** *Username* **ipxnet** *Ipxnetwork*

Configuring RIP Routing

As described in the *PortMaster Routing Guide*, PortMaster products automatically send and accept route information as RIP messages.

Note – ComOS 4.1 and later releases support both RIP-1 and RIP-2. Earlier releases of ComOS support only RIP-1.

To configure RIP routing for a network user, use the following command:

Command> **set user** *Username* **rip on|off|broadcast|listen|v2
{broadcast|on|v1-compatibility|multicast}**

Refer to the PortMaster 4 Command Line Reference for a description of the keywords in this command. Refer to the *PortMaster Routing Guide* for a discussion of routing with RIP, and for OSPF and BGP routing configuration instructions.

Setting the Asynchronous Character Map

The PPP protocol supports the replacement of nonprinting ASCII data in the PPP stream. These characters are not sent through the line, but instead are replaced by a special set of characters that the remote site interprets as the original characters. The PPP asynchronous map is a bit map of characters that must be replaced. The lowest-order bit corresponds to the first ASCII character NUL, and so on. In most environments, the asynchronous map must be set to zero to achieve the maximum data transfer rate.

To set the PPP asynchronous character map, use the following command:

Command> **set user** *Username* **map** *Hex*

Setting the MTU Size

The maximum transmission unit (MTU) defines the largest frame or packet that can be sent without fragmentation. A packet that exceeds this value is fragmented, if IP, or discarded if IPX. PPP connections can have a maximum MTU of 1520 bytes. SLIP connections can have a maximum MTU of 1006 bytes. PPP can negotiate smaller MTUs when requested by the calling party.

The MTU size is typically set to the maximum allowed for the protocol being used, either 1500 bytes (for PPP) or 1006 bytes (for SLIP). However, smaller MTU values can improve performance for interactive sessions. If you are using IPX, the MTU must be set to at least 600.

To set the MTU for a network user, use the following command:

Command> **set user** *Username* **mtu** *MTU*

Setting the Maximum Number of Dial-In Ports

You can define the number of dial-in ports that a user can use on the PortMaster for Multilink V.120, Multilink PPP (only on ISDN), and multiline load balancing.

If the maximum number of ports is unconfigured, port limits are not imposed and PortMaster multiline load balancing, Multilink V.120, and Multilink PPP sessions are allowed. You can also set the dial-in port limit using the RADIUS Port-Limit attribute.

To set the maximum number of dial-in ports, use the following command:

Command> **set user** *Username* **maxports** *Number*

The *Number* variable can be set to between 0 and the number of available ports—up to 95.

Setting Compression

Compression of TCP/IP headers can increase the performance of interactive TCP sessions over network hardwired asynchronous lines. Lucent implements Van Jacobson TCP/IP header compression and Stac LZS data compression. Compression is on by default.

Compression cannot be used with multiline load balancing, but can be used with Multilink PPP.

Compression must be enabled on both ends of the connection if you are using SLIP. With SLIP, TCP packets are not passed if only one side of the connection has compression enabled. For PPP connections, the PortMaster supports both bidirectional and unidirectional compression. Refer to RFC 1144 for more information about header compression.

The PortMaster supports Stac LZS data compression only for PPP connections with bidirectional compression. Stac LZS data compression cannot be used for SLIP connections.

To set header compression for a network user, use the following command:

> Command> **set user** *Username* **compression on|off**

Table 5-2 describes the results of using each keyword.

Table 5-2   Keywords for Configuring Compression

Keyword	Description
on	Enables compression. The PortMaster tries to negotiate both Van Jacobson and Stac LZS compression. This is the default.
off	Disables compression.

To find out what type of compression was negotiated for the user, enter the following command:

> Command> **show** *S0*

Setting Filters

Input and output packet filters can be applied to each network user. If an input filter is applied to a user, when the user dials in and establishes a connection, all packets received from the user are evaluated against the rule set for the applied filter. Only packets allowed by the filter can pass through the PortMaster. If an output filter is applied to a user, packets going to the user are evaluated against the rule set for the applied filter. Only packets allowed by the filter are sent out of the PortMaster to the user.

If either filter is changed while a user is logged on, the change does not take effect until the user disconnects and logs in again.

Note – You must define a filter in the filter table before you can apply it. For more information about filters, see Chapter 8, "Configuring Filters."

To apply an input filter for a network user, use the following command:

Command> **set user** *Username* **ifilter** [*Filtername*]

To apply an output filter for a network user, use the following command:

Command> **set user** *Username* **ofilter** [*Filtername*]

Omitting the *Filtername* removes any filter previously set for this user.

Note – Filters are applied to the user the next time the user dials in.

Specifying a Callback Location

You can configure the user for callback connections to enhance network security or to simplify telephone charges. When a network user logs in, the PortMaster disconnects the user and then calls back to the location specified for that user. The location is stored in the location table. The PortMaster always calls back using the same port on which the user called in. Network users have PPP or SLIP sessions started for them, as defined in the user table.

To specify the callback location for a network user, use the following command:

Command> **set user** *Username* **dialback** *Locname*|**none**

To disable callback connections for the user, use the **none** keyword.

Configuring Login Users

To configure a login user, you must set the login host, apply an optional access filter, set the login service type, and specify a callback telephone number.

Setting the Login Host

You must define the host to which the user is connected. The login host can be defined in one of three ways. Table 5-3 shows the login host options.

To set the login host for a login user, use the following command:

Command> **set user** *Username* **host default**|**prompt**|*Ipaddress*

Table 5-3   Login Host Options

Host Option	Description
default	This option allows the user to log in to the default or alternate host specified for this PortMaster. You can specify the default host with the **set host** command. For more information see the *PortMaster 4 Command Line Reference*.
prompt	This option allows the user to log in to a host by IP address or name at the time the login session is established.

Table 5-3   Login Host Options *(Continued)*

Host Option	Description
Ipaddress	This option allows the user to connect only to the host specifically named. A valid 39-character hostname or IP address must be entered. This configuration is used when you want to allow a user to access a specific host. For example, this configuration can be used to allow the user *carmela* to always be connected with the host *sales*.

Applying an Optional Access Filter

An access filter is an input filter that restricts hosts users can log in to. Access filters work as follows:

• The user logs in and specifies a host.

• The host address is compared against the access filter.

• If the address is permitted by the filter, the connection is established.

• If the address is not permitted, the connection is denied.

To apply an access filter to a login user, use the following command:

    Command> **set user** *Username* **ifilter** [*Filtername*]

Note – You must define a filter in the filter table before you can apply it. For more information about filters, see Chapter 8, "Configuring Filters."

Setting the Login Service Type

All login users must have an associated login service that determines the nature of their connection with the host.

The **login service** specifies how login sessions are established. Four types of login service are available as described in Table 5-4.

Table 5-4   Types of Login Service

Login Services	Function
portmaster	PortMaster is the default login service and can be used to access any host that has the PortMaster **in.pmd** daemon installed. This type of login service is preferred because it makes the PortMaster port operate like a serial port attached to the host. This service is the most cost-effective in terms of host resources.

Login Services	Function
rlogin	The remote login service **rlogin** uses the rlogin protocol to establish a login session to the specified host. Generally, **rlogin** is used on mixed UNIX networks where the PortMaster login service is impractical to use.
telnet	Telnet is supported on most TCP/IP hosts. This login service should be selected when the PortMaster and rlogin protocols are not available. The default port number is 23, but you can enter another number.
netdata	The **netdata** login service creates a virtual connection between the PortMaster port and another serial port on another PortMaster, or between the PortMaster port and a host. This login service creates a clear-channel TCP connection. To connect to another PortMaster port using **netdata**, you must configure that port as **/dev/network** with the **netdata** device service and the same TCP port number. The default **netdata** port is 6000; however, you can specify any TCP port number between 1 and 65535. This range allows TCP/IP to be used with a hardwired connection using an RS232 cable. However, some serial communications protocols, such as FAX, might have potential latency problems.

To set the login service type for a login user, use the following command:

> Command> **set user** *Username* **service portmaster|rlogin|telnet|netdata** [*Tport*]

Specifying a Callback Telephone Number

You can configure the login user for callback connections to enhance network security or to simplify telephone charges. When a user logs in, the PortMaster disconnects the user and then dials out to the telephone number specified for that user. The user is reconnected to the host specified in the user table, via the same port on which the user dialed in.

To enter the callback telephone number for a login user, use the following command:

> Command> **set user** *Username* **dialback** *String*|**none**

To disable callback connections for the user, use the **none** keyword.

Configuring a Synchronous WAN Port   6

This chapter describes the steps required to configure a PortMaster 4 synchronous wide area network (WAN) port.

This chapter discusses the following topics:

- "Synchronous Port Uses" on page 6-1

- "Configuring WAN Port Settings" on page 6-2

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Synchronous Port Uses

Synchronous WAN ports are used for high-speed dedicated connections between two remote local area networks (LANs). Once a connection is established between two remote sites, a wide area network (WAN) is created. Synchronous WAN connections can be achieved through the use of dedicated leased lines, Frame Relay connections, switched 56Kbps lines, or ISDN lines. Connection rates can range from 9600bps to 2.048Mbps (E1). The PortMaster 4 supports any of these connection types using one or more synchronous ports.

For most applications, a dedicated line connects two PortMaster routers, each located on a separate remote network

The following examples describe various uses for synchronous ports.

Routing over Leased Lines. A synchronous port can be used to connect to synchronous leased lines from 9600bps to T1 (1.544Mbps) or E1 (2.048Mbps) for continuous operation. The Quad T1 boards on a PortMaster 4 have a built-in channel service unit/digital service unit (CSU/DSU). For more information, see Chapter 13 "Using Synchronous Leased Lines."

Routing over Frame Relay. Frame Relay provides connectivity using a packet-switched network. Its two advantages over a leased line network are lower cost and the ability to have multiple permanent virtual circuits (PVCs) come into a single physical port. It is especially popular for hub-and-spoke network arrangements. For example, a dozen field offices with 56Kbps or fractional T1 Frame Relay connections can connect to a central office using a fractional T1 or T1 Frame Relay connection. The central office requires only one CSU/DSU and synchronous port on the PortMaster, instead of 12. For more information, see Chapter 12 "Using Frame Relay."

Routing over ISDN. Integrated Services Digital Network (ISDN) provides fast dial-up connectivity for applications where the expense of a dedicated Frame Relay or leased line connection is not called for by the amount and nature of the traffic. For more information, see Chapter 10 "Configuring T1, E1, and ISDN PRI."

Once you have determined the type of synchronous connection to use between your remote locations, the synchronous port on each end of the connection must be configured.

Configuring WAN Port Settings

The WAN port settings described in this section enable you to configure your synchronous port for your needs. "General Synchronous Settings" on page 6-2 includes settings that are available for all connection types. The settings in "Settings for Hardwired Connections" on page 6-5 are available only for network hardwired connections.

Setting the View

To configure a synchronous serial line as a WAN port, you must first set the view to the slot containing the board for the line that you want to configure. To set the view, enter the **show boards** command to determine the identification number of the line board you want to configure.

The board identification number is the same as the number of the slot in which the T1 or E1 line board is installed.

    Command> **show boards**

Use the following command to set the view to a slot with an installed T1 or E1 line board:

    Command> **set view** *Slotnumber*

Setting the view for a specific board gives you administrative access to that board.

General Synchronous Settings

The following settings can be used on synchronous ports configured for all connection types.

Displaying Extended Port Information

The PortMaster can display synchronous port information in brief or extended modes. The default setting is **off**.

To enable or disable extended information for a port, use the following command:

    Command> **set** *W1* **extended on|off**

Note – This command affects only the display of port information. It does not affect port behavior.

Setting the Port Type and Connection Type

Use the following command to set the port and connection type:

Command> **set** *W1* **network dialin**|**dialout**|**twoway**|**hardwired**

The port type for synchronous ports is always **network**, but you must explicitly set it. You also must specify the kind of connection to use on the synchronous port. Although you can configure a network port to allow dial-in and dial-out connections, a **network** port is typically used for a dedicated connection between two points known as **hardwired**. A hardwired connection does not use modem control.

To configure a port for a dedicated network connection, use the following command:

Command> **set** *W1* **network hardwired**

Table 6-1 describes the four connection types available on synchronous ports.

Table 6-1   Port and Network Types

Type	Description
hardwired	Allows you to establish a dedicated network connection between two sites without modem dialing or authentication. In this mode, the port immediately begins running the specified protocol. If the port is set for a hardwired connection, it cannot be used for any other purpose. A hardwired connection must be used for a leased line or Frame Relay connection.
dialin	Allows the port to accept dial-in network connections, for use with switched 56Kbps or ISDN connections. The dial-in user is required to enter a username and password before the connection is established. Authorized users are managed through the user table described in Chapter 5 "Configuring Dial-In Users," or through RADIUS.
	PPP users wishing to authenticate with PAP or CHAP can start sending PPP packets. When the packets are received, the PortMaster automatically detects PPP and requests PAP or CHAP authentication.
dialout	Allows dial-out users to establish connections with remote locations. Dial-out network destinations are managed through the location table. This network type can be used for ISDN and switched 56Kbps connections.
twoway	Allows the port to accept dial-in users and use dial-out locations. This network type can be used for ISDN and switched 56Kbps connections.

Setting the Port Speed Reference

The true port or line speed is set either by the external clock signal on the device to which the PortMaster is connected, or by the telephone company. You can record this value as a reference associated with a synchronous port, but it has no effect on PortMaster behavior.

To record the port speed, use the following command:

```
Command> set W1 speed Speed
```

You can substitute any of the following for *Speed*:

9600	**19200**	**56000**	**64000**	**115200**	**1536k**	**t1**	**e1**
14400	**38400**	**57600**	**76800**	**1344k**	**2048k**	**t1e**	

Setting Modem Control

When modem control is on, the PortMaster uses the condition of the carrier detect (DCD) signal from an attached modem to determine whether the line is in use.

Modem control is off for synchronous connections by default. With modem control set off, the PortMaster assumes the carrier detect line is always asserted. Table 6-2 describes the effects of DCD condition on port behavior.

Table 6-2   Effects of Carrier Detect Condition on Port Behavior

Connection Type	Carrier Detect Asserted	Carrier Detect De-asserted
hardwired	Port attempts to establish a network connection.	Port is unavailable.
dialin	PortMaster initiates authentication and displays a login prompt.	Port is unavailable.
dialout	No effect.	Transition from asserted to de-asserted resets the port.
twoway	Port attempts to establish a network connection.	Port is available.

Set modem control on only if you want to use the DCD signal from the attached device. In general, set modem control on for network dial-in or dial-out configurations. Modem control is usually off for leased line or Frame Relay connections, but you can use it if the channel service unit/digital service unit (CSU/DSU) is configured accordingly.

To set modem control, use the following command:

```
Command> set W1 cd on|off
```

Assigning a Port to a Dial Group

You can create modem pools for dial-out connections by associating ports and dial-out locations with dial groups. Dial groups can be used to reserve ports for dial-out to specific locations, or to differentiate among different types of modems that are compatible with the remote location. Dial groups are numbered 0 to 99. The default dial group is 0.

To assign a port to a dial group, use the following command:

    Command> **set** *W1* **group** *Group*

Setting Hangup Control

You can control whether the data terminal ready (DTR) signal on the synchronous port is dropped after a user session terminates. Hangup is set to **on** by default. In this state, DTR is dropped for 500 milliseconds, causing a hangup on the line.

To set the hangup control, use the following command:

    Command> **set** *W1* **hangup on**|**off**

The **reset** command always drops the DTR signal.

Setting the Port Idle Timer

The idle timer indicates how long the PortMaster waits after activity stops on a synchronous port before disconnecting a dial-in or dial-out connection.

You can set the idle time in seconds or minutes, to any value from 0 to 240. The default setting is 0 minutes. If the value is set to 2 seconds or a longer interval, the port is reset after having no traffic for the designated time. The idle timer is not reset by RIP, keepalive, or SAP packets. To disable the idle timer, set the value to 0.

To set the idle timer, use the following command:

    Command> **set** *W1* **idle** *Number* [**minutes**|**seconds**]

Settings for Hardwired Connections

The following settings can be used only when the synchronous port is configured for network hardwired connections.

Setting the Transport Protocol

The transport protocol for synchronous connections must be set for a network hardwired synchronous port. Choose PPP for leased line, switched 56Kbps, and ISDN connections, or Frame Relay for a Frame Relay connection. Additional Frame Relay settings must be configured for Frame Relay connections, described in Chapter 12 "Using Frame Relay."

To set the transport protocol, use the following command:

    Command> **set** *W1* **protocol slip|ppp|frame|x75-sync**

Setting the Port IP Address

You can set the local IP address of the network hardwired synchronous port to create a numbered interface.

You can use any IP address. If you set the local address of the WAN port to 0.0.0.0 for PPP, the PortMaster uses the Ether0 address for the end of the serial link. If you set the WAN port address to 0.0.0.0 for a Frame Relay connection, the port is disabled.

To set the IP address, use the following command:

    Command> **set** *W1* **address** *Ipaddress*

Setting the Destination IP Address

The destination IP address or hostname of the machine on the other end of the connection is used for leased line connections only. The destination IP address can also be set to 255.255.255.255 for PPP users. This setting allows the PortMaster to learn the IP address of the system on the other end of the connection using PPP IPCP address negotiation.

Do not set a destination IP address for Frame Relay connections. Instead, use the data link connection identifier (DLCI) list to link IP addresses to DLCIs, or use LMI or Annex-D and Inverse ARP to discover Frame Relay addresses dynamically. See Chapter 12 "Using Frame Relay," for more information.

For network dial-in or dial-out connections, do not set a destination IP address for the port. Instead, you set the destination address in the user table or RADIUS for dial-in, or in the location table for dial-out. See Chapter 5 "Configuring Dial-In Users" for more information.

To set the destination IP address for a leased-line connection only, use the following command:

    Command> **set** *W1* **destination** *Ipaddress* [*Ipmask*]

Setting the Subnet Mask

The default subnet mask is 255.255.255.0. If you have divided your network into subnets, enter the subnet mask that identifies how your network addresses are divided between the network portion and the host portion. The value of *Ipmask* is dependent upon the size of the IP subnet of which the IP address is a member. This setting is used on network hardwired ports only.

To set the subnet mask, use the following command:

    Command> **set** *W1* **netmask** *Ipmask*

See Appendix A, "Networking Concepts," for more information about using subnet masks.

Setting the IPX Network Address

Note – The PortMaster 4 supports the IPX protocol if it is running ComOS 4.1 or later. IPX is not supported in ComOS 4.0.

When using IPX, you must identify an IPX network number of the serial link that is unique from every other IPX number on the network. An IPX network address is entered in hexadecimal format, as described in Appendix A, "Networking Concepts."

Note – The serial link itself must have an IPX network number that is different from4 those at either end of the connection.

To set the IPX network address, use the following command:

    Command> **set** *W1* **ipxnet** *Ipxnetwork*

Configuring RIP Routing

As described in the *PortMaster Routing Guide,* PortMaster products automatically send and accept route information as RIP messages.

Note – ComOS 4.1 and later releases support RIP-1 and RIP-2. Earlier releases of ComOS support only RIP-1.

Turn on RIP routing for the port for network hardwired connections only, such as leased lines or Frame Relay. Routing is set in the user table for dial-in connections and in the location table for dial-out connections.

To configure RIP routing, use the following command:

    Command> **set** *W1* **rip on|off|broadcast|listen|v2**
    {**broadcast|multicast|on|v1-compatibility**}

Setting Input and Output Filters

Input and output packet filters can be attached to a synchronous port for network hardwired ports. Filters allow you to monitor and restrict network traffic. If an input filter is attached, all packets received from the interface are evaluated against the rule set for the attached filter. Only packets permitted by the filter are passed through the PortMaster. If an output filter is attached, packets going to the interface are evaluated against the rule set in the filter and only packets permitted by the filter are sent out of the interface.

Note – You must define a filter in the filter table before you can apply it. For more information about filters, see Chapter 8 "Configuring Filters."

To apply an input filter to a synchronous port, use the following command:

    Command> **set** *W1* **ifilter** [*Filtername*]

To apply an output filter to a synchronous port, use the following command:

    Command> **set** *W1* **ofilter** [*Filtername*]

You can remove filters from the port by entering the command without a filter name. If a filter is changed, you must reset the port for the change to take effect.

For example, to remove the output filter from a synchronous port, use the following commands:

```
Command> set W1 ofilter
Command> reset W1
Command> save all
```

Note – You must reset the port and re-establish the connection for the new settings to take effect.

Setting Compression

You can set Van Jacobson TCP/IP header compression and/or Stac LZS data compression on the port. To set compression, use the following command:

```
Command> set compression on|off|stac|vj
```

Van Jacobson TCP/IP header compression and Stac LZS data compression improve performance on asynchronous lines but can degrade performance on high-speed synchronous lines.

Configuring Dial-Out Connections   7

This chapter discusses how to create locations—settings for dial-out destinations—for dial-out connections.

This chapter discusses the following topics:

- "Configuring the Location Table" on page 7-1

- "Setting Multiline Load Balancing" on page 7-9

- "Setting Filters" on page 7-10

- "Testing Your Location Configuration" on page 7-11

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Configuring the Location Table

A location defines a dial-out destination and the characteristics of the dial-out connection. Locations control dial-out network connections in much the same way the user table controls dial-in network connections.

Locations are stored in the location table. All dial-out locations have the following minimum settings:

- Location name

- Name and password that the local PortMaster uses to authenticate itself to the remote host

- Telephone number of the remote host

- IP address and netmask of the remote host

- Protocol used for the connection

- Dial group that associates the location with a particular dial-out port

- Maximum number of ports

Locations can also optionally have the following settings:

- Connection type (dial-on-demand, continuous, or manual)

- Routing protocol

- IPX network number

- MTU size

- Compression

- Idle timer

- Data-over-voice for ISDN connections

- CHAP authentication

- Asynchronous character map

- Multiline load balancing

Note – The location table is not used for dialing out with the **tip** command or UNIX-to-UNIX Copy Protocol (UUCP).

To display the location table, enter the following command:

```
Command> show table location
```

A location table display looks like the following. The location table entries shown here are examples only. PortMaster products have empty location tables by default.

Location	Destination	Netmask	Group	Maxcon	Type
hq	172.16.1.1	255.255.255.0	1	4	On Demand
sf	192.168.1.21	255.255.255.0	99	1	Manual
sub1	192.168.3.1	255.255.255.0	2	0	Manual
bsp	172.16.1.21	255.255.255.0	99	1	Manual

Creating a Location

You must create a unique dial-out location for each remote host or router you want to access. Location table entries are identified by this unique location name, which can contain up to 12 characters.

To create a location, use the following command:

```
Command> add location Locname
```

Setting the Connection Type

Because the default method of initiating a connection is **manual**, you need to use the **dial** command to cause the PortMaster to manually dial out to a location. You can change the connection type as shown in Table 7-1. If you are changing an existing location's connection type, verify that the connection is not active.

Table 7-1   Dial-Out Connection Types

Connection Type	Description
on_demand	This type of connection is automatically started when packets for the remote location are queued by the PortMaster.
automatic	This type of connection is always active. If the telephone connection is dropped, the PortMaster initiates a new connection with the location after a 30-second waiting period.
manual	This type of connection is started when you request a connection. You can use this configuration to test a connection or for network callback users. This is the default.

To configure the connection type, use the following command:

    Command> s**et location** *Locname* **on_demand|automatic|manual**

On-Demand

Dial-on-demand connections to selected locations can save money because the telephone line is used only when traffic needs to be transmitted. The dial-on-demand configuration can also be used as a backup for other types of connections such as those using high-speed synchronous lines. A dial-on-demand connection usually has the idle timer set so that the connection is closed when no longer needed.

Note – When configuring a dial-on-demand location, be careful not to have the on-demand location be the route to the loghost, RADIUS server, RADIUS accounting server, or any host for a port using the PortMaster login or device service, unless you understand the effect of these services upon dial-on-demand.

If routing for a dial-on-demand location is set to **on**, **listen,** or **broadcast**, the PortMaster dials out to that location when it boots, to update routing information. The PortMaster hangs up when the idle timer expires because RIP traffic does not reset the idle timer.

To configure a location to support a dial-on-demand connection, use the following command:

    Command> **set location** *Locname* **on_demand**

Automatic

To establish an automatic dial-out connection, you must set the location type to **automatic**. In this configuration, the PortMaster dials out after it boots and establishes a network connection to the specified location. If the connection is dropped for any reason, the PortMaster dials out again and establishes the connection again after a 30-second wait.

To configure a location to support a automatic connection, use the following command:

> Command> **set location** *Locname* **automatic**

Manual Dial-Out

Use manual dial-out to test the connection or if you want the connection to be established only when you or a network callback user requests. You should test any connection before configuring it as a continuous or on-demand location.

To configure a location to support a manual connection, use the following command:

> Command> **set location** *Locname* **manual**

Note – Disconnect dial-out connections by resetting the port before switching a connection type from **manual** to **on demand**.

Setting the Telephone Number

The telephone number setting is used to dial out to the remote location.

To set the telephone number of the remote location, use the following command:

> Command> **set location** *Locname* **telephone** *String*

Setting the Username and Password

The username and password are what the PortMaster uses to authenticate itself to the remote host. Note that the username and password you enter here must also be resident on the remote host in the user table, RADIUS, or other authentication mechanism.

To set the username and password, use the following commands:

> Command> **set location** *Locname* **username** *Username*
> Command> **set location** *Locname* **password** *Password*

Setting the Protocol

The network protocol for a dial-out location is typically set for PPP packet encapsulation, SLIP encapsulation, or X.75-sync (used in Europe). PPP can be used with IP packet routing, IPX packet routing, or both. Select a protocol that is compatible with the remote location.

Note – New location table entries default to PPP.

To set the protocol for a location, use the following command:

> Command> **set location** *Locname* **protocol slip|ppp|x75-sync**

For more information about setting the location protocol to a Frame Relay subinterface, see "Frame Relay Subinterfaces" on page 12-8.

Setting the Destination IP Address

The destination IP address is the IP address expected on the system at the remote end of the dial-out connection.

For PPP connections, you can either specify an IP address or have it negotiated. If you enter 255.255.255.255 (negotiated) for the destination IP address, the PortMaster learns the IP address of the remote system during PPP IPCP negotiation.

For SLIP connections and locations set for on-demand dialing, enter the IP address or a valid hostname of up to 39 characters for the system at the remote end of the connection.

Note – Assigned addresses are not supported for dial-out locations.

To set the destination IP address for a location, use the following command:

```
Command> set location Locname destination Ipaddress
```

Setting the Destination Netmask

If the host or network on the remote end of the connection requires a netmask, you must define it in the location table.

To set the destination netmask for a location, use the following command:

```
Command> set location Locname netmask Ipmask
```

Setting the IPX Network Number

Note – The PortMaster 4 supports the IPX protocol if it is running ComOS 4.1 or later. IPX is not supported in ComOS 4.0.

If you use the IPX protocol, you must assign a unique IPX network number to the network connection between the remote host and the PortMaster. Enter the IPX network number in the hexadecimal format described in Appendix A, "Networking Concepts." The number can consist of up to eight characters. The number is used only for the serial link, and must be different from the IPX network numbers used for Ethernets at either end.

To set the IPX network number for a location, use the following command:

```
Command> set location Locname ipxnet Ipxnetwork
```

Note – Do not set a value of all 0s (zeros) or all Fs for the IPX network numbers.

Setting RIP Routing

You can associate RIP routing with locations—for example, a dial on-demand connection where the remote router is defined as a location on the local PortMaster.

As described in the *PortMaster Routing Guide*, PortMaster products automatically send and accept route information as RIP messages, unless configured otherwise.

Note – ComOS 4.1 and later releases support both RIP-1 and RIP-2 on the PortMaster 4. Earlier releases of ComOS support only RIP-1.

Refer to the *PortMaster Routing Guide* for OSPF and BGP configuration instructions.

To set RIP routing for a location, use the following command:

```
Command> set location Locname rip on|off|broadcast|listen|v2
{broadcast|multicast|on|v1-compatibility}
```

Refer to the *PortMaster 4 Command Line Reference* for a description of the keywords in this command. Refer to the *PortMaster Routing Guide* for a discussion of routing with RIP, and for OSPF and BGP routing configuration instructions.

Setting the Dial Group

Dial groups associate locations with specific dial-out ports. By default, all ports and locations belong to dial group 0 (zero). You can configure locations and ports into dial groups numbered from 0 to 99. Dial group numbers can be used to reserve ports for dial-out to specific locations, or to differentiate among different types of modems that are compatible with the remote location.

The dial group associated with a location works with the dial group specified for each port. For example, you create a dial-out location called *home* and specify that the dial group for *home* is 2. When you configure each port, you can assign the port to a dial group. Only ports assigned to group 2 are used to dial the location *home*, while other ports are not.

To associate a location with a dial group number, use the following command:

```
Command> set location Locname group Group
```

Setting the MTU Size

The maximum transmission unit (MTU) defines the largest frame or packet that can be sent through this port, without fragmentation. If an IP packet exceeds the specified MTU, it is automatically fragmented. An IPX packet that exceeds the specified MTU is automatically dropped. PPP connections can have a maximum MTU of 1500 bytes. SLIP connections can have a maximum MTU of 1006 bytes. With PPP, the PortMaster can negotiate smaller MTUs when requested during PPP negotiation.

The MTU is typically set to the maximum allowed for the protocol being used. However, smaller MTU values can improve performance for interactive sessions. During PPP negotiation, the smaller number is used. If you are using IPX, the MTU must be set to at least 600.

To set the MTU for a location, use the following command:

Command> **set location** *Locname* **mtu** *MTU*

Configuring Compression

Compression of TCP/IP headers can increase the performance of interactive TCP sessions over network hardwired asynchronous lines. Lucent implements Van Jacobson TCP/IP header compression and Stac LZS data compression. Compression is on by default.

Compression cannot be used with multiline load balancing, but can be used with Multilink PPP.

Compression must be enabled on both ends of the connection if you are using SLIP. With SLIP, TCP packets are not passed if only one side of the connection has compression enabled. For PPP connections, the PortMaster supports both bidirectional and unidirectional compression. Refer to RFC 1144 for more information about header compression.

The PortMaster supports Stac LZS data compression only for PPP connections with bidirectional compression. Stac LZS data compression cannot be used for SLIP connections.

To configure compression for a location, use the following command:

Command> **set location** *Locname* **compression on|off|stac|vj**

Table 7-2 describes the results of using each keyword.

Table 7-2   Keywords for Configuring Compression

Keyword	Description
on	Enables compression. The PortMaster tries to negotiate both Van Jacobson and Stac LZS compression. This is the default.
off	Disables compression.
stac	Enables Stac LZS data compression only.
vj	Enables Van Jacobson TCP/IP header compression only.

To display compression information about a location, enter the following command:

Command> **show** *S0*

Setting the Idle Timer

You can set the idle timer for a location with manual or on-demand connections. This timer defines the length of time the line can be idle, with no network traffic in either direction, before the PortMaster disconnects the connection. You can set the idle time in seconds or minutes, to any value from 0 to 240. The default setting is 0 minutes. If the

value is set to 2 seconds or a longer interval, the port is reset after having no traffic for the designated time. The idle timer is not reset by RIP or keepalive packets. To disable the idle timer, set the value to 0.

Note – Idle timers for dial-in connections are set on each port or for specific users. Idle timers for dial-out connections are set in the location table.

To set the idle time for a location with a manual or on-demand connection, use the following command:

    Command> **set location** *Locname* **idletime** *Number* [**minutes**|**seconds**]

Setting Data over Voice

The PortMaster supports data-over-voice for inbound and outbound ISDN connections. The PortMaster automatically accepts inbound voice calls and treats them as data calls. You can force a data-over-voice call on an outbound ISDN connection by setting the capability to **on**.

To turn on the data-over-voice capability for ISDN connections to a location, use the following command:

    Command> **set location** *Locname* **voice on**|**off**

For more information on ISDN connections, see Chapter 10, "Configuring T1, E1, and ISDN PRI."

Setting CHAP

When you enter a username and password into the location table, they are used as the system identifier and message-digest algorithm 5 (MD5) secret for CHAP authentication. You can turn on outbound CHAP authentication and eliminate the need to use the **sysname** identifier and user table configurations for CHAP, unless the device being dialed also dials in to the PortMaster. The default setting is **off**.

To set CHAP authentication for a location, use the following command:

    Command> **set location** *Locname* **chap on**|**off**

Setting the Asynchronous Character Map

The PPP protocol supports the replacement of nonprinting ASCII data in the PPP stream. These characters are not sent through the line, but instead are replaced by a special set of characters that the remote site interprets as the original characters. The PPP asynchronous map is a bit map of characters that must be replaced. The lowest-order bit corresponds to the first ASCII character NUL, and so on. Most environments must set the asynchronous map to 0 (zero) to achieve maximum throughput.

To set the PPP asynchronous map for a location, use the following command:

    Command> **set location** *Locname* **map** *Hex*

Setting Multiline Load Balancing

You can set several ports to connect to a single location to distribute heavy traffic loads. This capability is called multiline load balancing. You can define a threshold—known as a high-water mark—for a location. The high-water mark triggers the PortMaster to bring up an additional connection to the location when the amount of data specified by the high-water mark is queued. The PortMaster examines the queue several times a minute to determine if the high-water mark has been reached.

Load balancing is useful for on-demand routing because additional ports for the location are added as the load exceeds what can be handled by one port. When the ports are idle for the time specified by the **set location idletime** command (see "Setting the Idle Timer" on page 7-7), all ports used for that connection are timed out simultaneously.

Load balancing can save you money because you do not need to configure your network to handle the maximum load between locations. Periods of heavy traffic can be handled by additional ports on an as-needed basis. At other times, the additional ports can be used for other purposes.

When multiple ports are in use, each packet is queued on the port with the least amount of traffic in the queue. Ports with very different speeds must not be combined for load balancing purposes. The overall throughput for a given number of ports is approximately equal to the number of ports multiplied by the throughput of the slowest port.

The following settings are used to configure load balancing and define when additional lines to this location are dialed.

Setting the Maximum Number of Dial-Out Ports

To configure load balancing, you must define the number of dial-out ports that can be used to dial and establish a connection with this location. This setting creates a pool of ports that can be used at the same time to establish a connection with this location.

If the maximum number of ports is set to 0, no connection with this location is established. If the maximum number of ports is set to any number greater than one, the high-water mark is used to determine when additional connections are established with this location.

When more than one line is open to a given location, the PortMaster balances the load across each line. When the ports are idle for the time specified by the **set location idletime** command (see "Setting the Idle Timer" on page 7-7), all ports used for that connection are timed out simultaneously.

To set the maximum number of dial-out ports for a location, use the following command:

```
Command> set location Locname maxports Number
```

The *Number* variable is a value between 0 and 95—the total number of available ports.

Setting Bandwidth-on-Demand

The bandwidth-on-demand feature provides a way to specify a point at which the PortMaster establishes an additional line to a location. You use the **high_water** keyword to specify the number of bytes of network traffic that must be queued before the PortMaster opens an additional connection. The PortMaster examines the queue several times a minute to determine if the high-water mark has been reached.

If you set a very low threshold number, the PortMaster quickly opens the maximum number of ports you specify for this location. When selecting a threshold, bear in mind that interactive traffic from login users queues a relatively small number of bytes—only several hundred—while network users doing file transfers can queue several thousand bytes of traffic.

This high-water value is used only when the maximum number of ports is greater than one. The default high-water mark is 0 (zero).

To set the high-water mark in bytes for a location, use the following command:

> Command> **set location** *Locname* **high_water** *Number*

Setting Filters

You can attach input and output filters to each location. Filters must be defined in the filter table before they can be added to the location table. For more information about filters, see Chapter 8, "Configuring Filters." When a filter is changed, all ports in use by the location must be reset to have the changes take effect.

Note – If a matching filter name is not found in the filter table, this command is not effective and all traffic is permitted.

Input Filters

Input filters cause all packets received from the interface to be evaluated against the filter rule set. Only packets allowed by the filter are accepted.

To set an input filter for a location, use the following command:

> Command> **set location** *Locname* **ifilter** *Filtername*

Output Filters

Output filters cause all packets going out to the interface to be evaluated against the filter rule set. Only packets allowed by the filter are passed out to the interface.

To set an output filter for a location, use the following command:

> Command> **set location** *Locname* **ofilter** *Filtername*

Testing Your Location Configuration

When you are configuring a location, you can set a manual connection for the location so that you can test the configuration before resetting the connection to on-demand or automatic. To test the configuration, you must initiate a connection with the remote location by using the **dial** command from the command line.

To display the chat script (if you are using one) during dialing, use the optional **-x** keyword. You can watch the connection process to ensure that location-specific settings are configured correctly. This keyword also resets some debugging values previously set with **set debug**.

When your location is configured correctly, change the connection type from manual to automatic or on-demand.

To test your configuration, use the following command:

```
Command> dial Locname [-x]
```

Configuring Filters 8

This chapter describes how to configure input and output packet filters. IP, IPX, and Service Advertising Protocol (SAP) rules are reviewed, and filter examples are given. You can also use the ChoiceNet application to filter IP packets by lists of sites rather than by individual IP addresses. For more information on ChoiceNet, see the *ChoiceNet Administrator's Guide*.

This chapter discusses the following topics:

- "Overview of PortMaster Filtering" on page 8-1

- "Creating Filters" on page 8-4

- "Displaying Filters" on page 8-7

- "Deleting Filters" on page 8-7

- "Example Filters" on page 8-7

- "Restricting User Access" on page 8-12

Each topic in this chapter includes examples of filters used to accomplish the goal described.

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

The FilterEditor application provides a graphical interface to construct and edit filters for both PortMaster 4 Remote Access Concentrators and ChoiceNet servers.

PMVision, FilterEditor, and other Java-based configuration tools for the PortMaster are available via anonymous FTP at **ftp://ftp.livingston.com/pub/livingston/software/java/**.

Overview of PortMaster Filtering

Packet filters can increase security and decrease traffic on your network. You use filters to limit certain kinds of internetwork communications by permitting or denying the passage of packets through network interfaces. By creating appropriate filters, you can control access to specific hosts, networks, and network services.

You can enhance security on your network by limiting authorized activities to certain hosts. For example, you can restrict the DNS and SMTP interchange with the Internet to a well-secured host on your network. All Internet hosts can then access only this single server for those services. If you have several name servers or mail servers, you can use additional rules to allow access to these servers.

You use Ethernet filters to constrain the types of packets that can enter the local Ethernet port, and you can set filters on asynchronous ports configured for hardwired operation when security with another network is an issue.

The packet filtering process analyzes the header information in each packet sent or received through a network interface. The header information is evaluated against a set of rules that either allow the packet to pass through the interface or cause the packet to be discarded.

A maximum of 256 filter rules per filter is allowed for the PortMaster 4. The PortMaster generates an error message when the number of filter rules exceeds the limit.

If a packet is discarded by a filter, an appropriate "ICMP unreachable" message is returned to the source address. This message provides immediate feedback to the user attempting the unauthorized access. Packets permitted or denied can optionally be logged to a host.

Filters can also be used for packet selection—for example, you can use a packet trace filter to do troubleshooting. The packets permitted by the **ptrace** filter are displayed, while packets not permitted by the filter are not displayed. For more information about the **ptrace** facility, see the *PortMaster Troubleshooting Guide*.

Filter Options

Table 8-1 shows different filter options.

Table 8-1   Filter Options

Option	Description
Restricting packet traffic	Each user, location entry, and network hardwired port can be assigned both an input packet filter and an output packet filter. Having both input and output filters can decrease the number of rules needed and can provide better tuning of your security policy.
Restricting access based on source and destination address	You can create filters that evaluate both the source and destination addresses of a packet against a rule list. The number of significant bits used in IP address comparisons can be set, allowing filtering by host, subnet, network number, or group of hosts whose addresses are within a given bit-aligned boundary.
Restricting access to particular protocols	Packets of certain protocols can be permitted or denied by a filter, including IPX, SAP, TCP, UDP, and ICMP packets.
Restricting access to network services	You can create filters that use the source and destination port numbers to control access to certain network services. The evaluation can be based upon whether the port number is less than, equal to, or greater than a specified value.

Table 8-1   Filter Options *(Continued)*

Option	Description
Restricting access based on TCP status	You can create filters that use the status of TCP connections as part of the rule set. This feature can allow network users to open connections to external networks without allowing external users access to the local network.

Filter Organization

Filters are stored in a filter table in the PortMaster nonvolatile configuration memory. Filters can be created or modified at any time, and the changes are not applied to an active use of the filter. Filter names must be between 1 and 15 characters.

Each packet filter can contain three sets of rules: IP, IPX, and SAP. Within each set, the rules are numbered starting at one. Newly created packet filters contain zero rules, or an empty set of rules.

An empty set of rules is equivalent to the permit rule. If a filter contains one or more rules in the set, any packet not explicitly permitted by a rule is denied at the end of the rule set.

How Filters Work

IP and IPX packet filters are attached to users, locations, Ethernet interfaces, or network hardwired ports as either input or output filters. SAP filters are attached as output filters only. The Ethernet interface filter is enabled as soon as the name of the input or output filter is set.

Input and output are defined relative to the PortMaster interface. As shown in Figure 8-1, an input filter is used on packets entering the PortMaster and an output filter is used on packets exiting the PortMaster.

Figure 8-1   Input and Output Filters
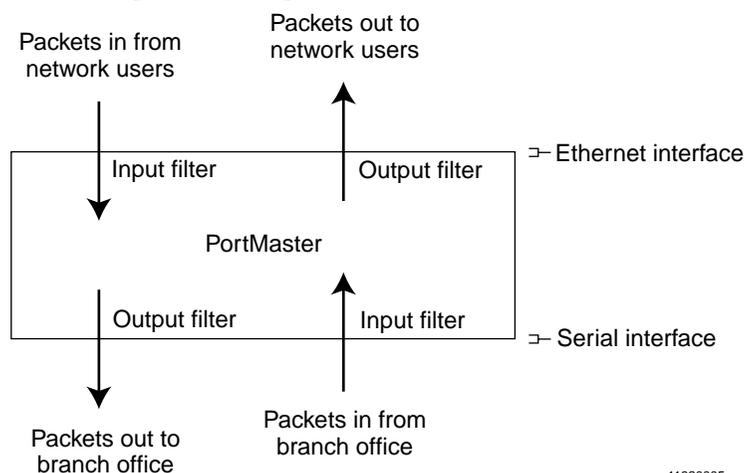
All packets entering a PortMaster through an interface with an input filter are evaluated against the rules in the filter. As soon as a packet matches a rule, the action specified by that rule is taken. If no rules match the specific packet, the packet is denied and is discarded. Whenever an IP packet is discarded, the PortMaster generates an "ICMP Host Unreachable" message back to the originator.

For interfaces with output filters attached, all packets exiting the interface are evaluated against the filter rules and only those packets permitted by the filter are allowed to exit the interface.

Creating Filters

You construct a filter by creating the filter and then adding rules that permit or deny certain types of packets. A maximum of 256 filter rules per filter is allowed for the PortMaster 4. The PortMaster generates an error message when the number of filter rules exceeds the limit.

Because the PortMaster evaluates packets in the order in which rules are listed, you can avoid bottlenecks and maximize throughput by specifying early those rules representing your highest security concerns, followed by a rule limiting the volume of traffic.

User filters are attached to users configured for dial-in SLIP or PPP access. When a user makes a PPP or SLIP connection, the designated filters are attached to the network interface created for that connection.

Location filters are attached to dial-out locations by means of SLIP or PPP connections. When the connection is established to a remote site, the designated filters are attached to the network interface used.

You can attach filters for incoming packets, or for outgoing packets or for both. It is usually more effective to filter incoming packets so that you can protect the PortMaster itself.

For more detailed instructions on using the filter commands, see the *PortMaster 4 Command Line Reference*.

To create a filter, use the following command:

```
Command> add filter Filtername
```

You must then use the appropriate **set** command to add rules that permit or deny packets. A maximum of 256 filter rules per filter is allowed. The PortMaster generates an error message when the number of filter rules exceeds the limit.

See the following sections for instructions:

- "Creating IP Filters" on page 8-4

- "Filtering TCP and UDP Packets" on page 8-5

Creating IP Filters

You can create a rule that filters IP packets according to their source and destination IP addresses. For more information on the command syntax for creating filters, see the *PortMaster 4 Command Line Reference*.

To create an IP filter rule that filters by address, use the following command—entered on one line:

```
Command> set filter Filtername RuleNumber permit|deny [Ipaddress/NM
Ipaddress(dest)/NM] [protocol Number] [log] [notify]
```

You can replace **protocol** *Number* with one of the following keywords:

- **esp**—matches packets using Encapsulation Security Payload (ESP) protocol. See RFC 1827 for more information on this protocol.

- **ah**—matches packets using Authentication Header (AH) protocol. See RFC 1826 for more information on this protocol.

- **ipip**—matches packets using the IP Encapsulation within IP (IPIP). See RFC 2003 for more information on this protocol.

If you are using ChoiceNet, you can also replace either the source or destination IP address with the value =*ListName*, which specifies a list of sites in the **/etc/choicenet/lists** directory in the ChoiceNet server. The equal sign (=) must immediately precede the value.

Filtering ICMP Packets

Internet Control Message Protocol (ICMP) packets—commonly known as ping packets—report errors and provide other information about IP packet processing. You can filter ICMP packets by source and destination IP address, or by ICMP packet type. Packet types are identified in RFC 1700.

To create an ICMP filter rule, use the following command—entered on one line:

```
Command> set filter Filtername RuleNumber permit|deny [Ipaddress/NM
Ipaddress(dest)/NM] icmp [type Itype] [log]
```

If you are using ChoiceNet, you can also replace either the source or destination IP address with the value =*ListName*, which specifies a list of sites in the **/etc/choicenet/lists** directory in the ChoiceNet server. The equal sign (=) must immediately precede the value.

Filtering TCP and UDP Packets

If you are using ChoiceNet, you can also replace either the source or destination IP address in a TCP or IDP filter with the value =*ListName*, which specifies a list of sites in the **/etc/choicenet/lists** directory in the ChoiceNet server. The equal sign (=) must immediately precede the value.

TCP Packets

You can filter TCP packets by source and destination IP address, or by TCP port number. Appendix B, "TCP and UDP Ports and Services," lists port numbers commonly used for UDP and TCP port services. For a more complete list, see RFC 1700.

To create a TCP filter rule, use the following command—entered on one line:

```
Command> set filter Filtername RuleNumber permit|deny [Ipaddress/NM
Ipaddress(dest)/NM] tcp [src eq|lt|gt Tport] [dst eq|lt|gt Tport]
[established] [log]
```

UDP Packets

You can filter UDP packets by source and destination IP address, or by UDP port number. Appendix B, "TCP and UDP Ports and Services," lists port numbers commonly used for UDP and TCP port services. For a more complete list, see RFC 1700.

To create a UDP filter rule, use the following command—entered on one line:

```
Command> set filter Filtername RuleNumber permit|deny [Ipaddress/NM
Ipaddress(dest)/NM] udp [src eq|lt|gt Uport] [dst eq|lt|gt Uport]
[established] [log]
```

Creating IPX Filters

You can filter IPX packets in the following ways:

• Source and/or destination IPX network number

• Source and/or destination IPX node address

• Source and/or destination IPX socket number

To create an IPX filter rule, use the following command—entered on one line:

```
Command> set ipxfilter Filtername RuleNumber permit|deny [srcnet Ipxnetwork]
[srchost Ipxnode] [srcsocket eq|gt|lt Ipxsock] [dstnet Ipxnetwork]
[dsthost Ipxnode] [dstsocket eq|gt|lt Ipxsock]
```

Creating SAP Filters

The Service Advertising Protocol (SAP) is an IPX protocol used over routers and servers that informs network clients of available network services and resources. SAP packets can be filtered only on output. You can filter SAP packets according to the following information about the server that is advertising the service via SAP:

• Name

• IPX network number

• IPX node address

• IPX socket number

To create a SAP filter rule, use the following command—entered on one line:

```
Command> set sapfilter Filtername RuleNumber permit|deny
[server String][network Ipxnetwork] [host Ipxnode] [socket eg|gt|lt Ipxsock]
```

Displaying Filters

To display the filter table, use the following command:

Command> **show table filter**

To display a particular filter, use the following command:

Command> **show filter** *Filtername*

Deleting Filters

To delete a filter, use the following command:

Command> **delete filter** *Filtername*

Example Filters

Because filters are very flexible, you must carefully evaluate the types of traffic that a specific filter permits or denies through an interface before attaching the filter. If possible, a filter should be tested from both sides of the filtering interface to verify that the filter is operating as you intended. Using the **log** keyword to log packets that match a rule to the loghost is useful when you are testing and refining IP filters.

Some of the following examples use the 192.168.1.0 network as the public network. Substitute the number of your network or subnetwork if you use these examples.

Note – Any packet that is not explicitly permitted by a filter is denied, except for the special case of a filter with no rules, which permits everything.

Simple Filter

A simple filter can consist of the following rules:

```
Command> set filter simple 1 permit udp dst eq 53
Command> set filter simple 2 permit tcp dst eq 25
Command> set filter simple 3 permit icmp
Command> set filter simple 4 permit 0.0.0.0/0 192.168.1.3/32 tcp dst eq 21
Command> set filter simple 5 permit tcp src eq 20 dst gt 1023
```

Table 8-2 describes, line by line, each rule in the filter.

Table 8-2   Description of Simple Filter

Rule	Description
1.	Permits Domain Name Service (DNS) UDP packets from any host to any host.
2.	Permits SMTP (mail) packets.
3.	Permits ICMP packets.
4.	Permits FTP from any host, but only to the host 192.168.1.3.

Table 8-2   Description of Simple Filter *(Continued)*

Rule	Description
5.	Permits FTP data to return to the requesting host. This rule is required to provide a reverse channel for the data portion of FTP.

Input Filter for an Internet Connection

The filter in this example is designed as an input filter for a network hardwired port that connects to the Internet. You can use this filter for a dial-on-demand connection by attaching it to the location entry.

The rules for the filter are set as follows:

```
Command> set filter internet.in 1 deny 192.168.1.0/24 0.0.0.0/0 log
Command> set filter internet.in 2 permit tcp estab
Command> set filter internet.in 3 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 25
Command> set filter internet.in 4 permit 0.0.0.0/0 172.16.0.4/32 tcp dst eq 21
Command> set filter internet.in 5 permit tcp 0.0.0.0/0 192.168.0.5/32 dst eq 80
Command> set filter internet.in 6 permit tcp src eq 20 dst gt 1023
Command> set filter internet.in 7 permit udp dst eq 53
Command> set filter internet.in 8 permit tcp dst eq 53
Command> set filter internet.in 9 permit icmp
```

Table 8-3 describes, line by line, each rule in the filter.

Table 8-3   Description of Internet Filter

Rule	Description
1.	Denies any incoming packets from the Internet claiming to be from—or **spoofing**—your own network (192.168.1.0). This rule blocks IP spoofing attacks. This rule also logs the header information in the spoofing packets to **syslog**.
2.	Permits already established TCP connections that originated from your network—packets with the ACK bit set.
3.	Permits SMTP connections to 10.0.0.3 (the mail server).
4.	Permits FTP connections to host 172.16.0.4.
5.	Permits Hypertext Transfer Protocol (HTTP) access to host 192.168.0.5.
6.	Permits an FTP data channel.
7.	Permits DNS.
8.	Permits DNS zone transfers. (You can write this rule to allow only connections to your name servers.)
9.	Permits ICMP packets.

Input and Output Filters for FTP Packets

Filters can be used to either permit or deny File Transfer Protocol (FTP) packets. You must understand how this protocol works before you develop FTP filters.

FTP uses TCP port 21 as a control channel, but it transfers data on another channel initiated by the FTP server from TCP port 20 (FTP-data). Therefore, if you want to allow your internal hosts to send out packets with FTP, you must allow external hosts to open an incoming connection from TCP port 20 to a destination port above 1023. Allowing this type of access to your network can be very risky if you are running Remote Procedure Call (RPC) or X Windows on the host from which you are transmitting FTP packets. As a result, many sites use FTP proxies or passive FTP, neither of which is discussed in this guide.

Consult *Firewalls and Internet Security: Repelling the Wily Hacker* by Cheswick and Bellovin and *Building Internet Firewalls* by Chapman and Zwicky for information on FTP proxies and passive FTP.

Likewise, if you want to allow external hosts to connect to your FTP server and transfer files, you must allow incoming connections to TCP port 21 on your FTP server and allow outgoing connections from TCP port 20 of your FTP server.

In the following examples, 172.16.0.2 is the address of your FTP server and 192.168.0.1 is the address of the host from which you allow outgoing FTP.

⚠ **Caution –** This configuration is not recommended if you run any of the following protocols on any of the hosts from which you allow FTP access: NFS, X, RPC, or any other service that listens on ports above 1023.

The rules for the input filter are as follows:

```
Command> set filter internet.in 1 permit 0.0.0.0/0 192.168.0.1/32 tcp src eq
20 dst gt 1023
Command> set filter internet.in 2 permit 0.0.0.0/0 192.168.0.1/32 tcp src eq
21 estab
Command> set filter internet.in 3 permit 0.0.0.0/0 172.16.0.2/32 tcp dst eq 21
Command> set filter internet.in 4 permit 0.0.0.0/0 172.16.0.2/32 tcp src gt
1023 dst eq 20 estab
```

The rules for the output filter are as follows:

```
Command> set filter internet.out 1 permit 192.168.0.1/32 0.0.0.0/0 tcp dst eq
21
Command> set filter internet.out 2 permit 192.168.0.1/32 0.0.0.0/0 tcp src gt
1023 dst eq 20 estab
Command> set filter internet.out 3 permit 172.16.0.2/32 0.0.0.0/0 tcp src eq
20 dst gt 1023
Command> set filter internet.out 4 permit 172.16.0.2/32 0.0.0.0/0 tcp src eq
21 dst gt 1023 estab
```

If you allow any internal host to send out packets with FTP, replace 192.168.0.1/32 with 0.0.0.0/0 or *your network_number*/24. Take appropriate precautions to reduce the risk this configuration creates.

Rule to Permit DNS into Your Local Network

If the DNS name server for your domain is outside your local network, add the following rule to your input filter:

> Command> **set filter** *Filtername RuleNumber* **permit udp src eq 53**

This rule permits DNS replies into your local network.

Rule to Listen to RIP Information

To permit incoming RIP packets, add the following rule to your input filter:

> Command> **set filter** *Filtername RuleNumber* **permit 172.16.0.0/32 192.168.0.0/32 udp dst eq 520**

In this example, 172.16.0.0/32 is the other end of the Internet connection and 192.168.0.0/32 is the local address of the connection.

Rule to Allow Authentication Queries

To allow authentication queries used by some mailers and FTP servers, add the following rule to your input filter:

> Command> **set filter** *Filtername RuleNumber* **permit tcp dst eq 113**

For more information about these types of queries, refer to RFC 1413.

Rule to Allow Networks Full Access

To allow some other network to have complete access to your network, add the following rule. In the example below, 172.16.12.0 is granted full access to 192.168.1.0/24:

> Command> **set filter** *Filtername RuleNumber* **permit 172.16.12.0/24 192.168.1.0/24**

Caution – Beware of associative trust. If you allow a network complete access to your network, you might unknowingly allow other networks complete access, as well. Any network that can access a network having complete access privileges to your network, also has access to your network. For example, if Network 1 trusts Network 2 and Network 2 trusts Network 3, then Network 1 trusts Network 3.

Restrictive Internet Filter

This example filter allows any kind of outgoing connection from the server, but blocks all incoming traffic to any host but your designated Internet server. This filter also limits incoming traffic on your Internet server to SMTP, Network News Transfer Protocol (NNTP), DNS, FTP, and ICMP services.

Note – Even if you have the latest versions of the daemons **ftpd**, **httpd**, and **sendmail** you might be vulnerable to attacks through these services. Check the latest CERT Coordination Center advisories, available on **ftp.cert.org**, for the vulnerabilities of these services.

If you use the following example, replace the name **server** with the IP address or hostname of your Internet server:

```
Command> set filter restrict.in 1 deny 192.168.1.0/24 0.0.0.0/0 log
Command> set filter restrict.in 2 permit 0.0.0.0/0 10.0.0.3/32 tcp estab
Command> set filter restrict.in 3 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 21
Command> set filter restrict.in 4 permit 0.0.0.0/0 10.0.0.3/32 tcp src eq 20
dst gt 1023
Command> set filter restrict.in 5 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 119
Command> set filter restrict.in 6 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 25
Command> set filter restrict.in 7 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 80
Command> set filter restrict.in 8 permit 0.0.0.0/0 10.0.0.3/32 udp dst eq 53
Command> set filter restrict.in 9 permit 0.0.0.0/0 10.0.0.3/32 tcp dst eq 53
Command> set filter restrict.in 10 permit 0.0.0.0/0 10.0.0.3/32 icmp
```

Table 8-4 describes, line by line, each rule in the filter.

Table 8-4   Description of Restrictive Internet Filter

Rule	Description
1.	Denies any incoming packets from your own network (192.168.1.0) and makes a log.
2.	Permits packets from any established TCP connection to 10.0.0.3 (the Internet server).
3.	Permits FTP from any IP address to 10.0.0.3 (the server).
4.	Permits the FTP data back channel.
5.	Permits incoming NNTP (news) to 10.0.0.3 (the Internet server).
6.	Permits incoming SMTP (mail) to 10.0.0.3 (the Internet server).
7.	Permits HTTP requests to 10.0.0.3 (the Internet server).
8.	Permits DNS queries to 10.0.0.3 (the Internet server).
9.	Permits DNS zone transfers from 10.0.0.3 (the Internet server).
10.	Permits ICMP to 10.0.0.3 (the Internet server). You can further limit ICMP packet types to types 0, 3, 8, and 11 using four rules instead of one.

To log all packets that are denied, add the following rule to the end of your filter:

```
Command> set filter Filtername RuleNumber deny log
```

Restricting User Access

Access filters enable you to restrict Telnet or **rlogin** connections to a specific host or network, or a list of hosts or networks. You can create an access filter that restricts user access to particular hosts.

Access filters work as follows:

1. The user specifies a host.

2. The host address is compared against the access filter.

3. If the address is permitted by the filter, the connection is established.

4. If the address is not permitted, the connection is denied unless access override is enabled.

If you want a user to be able to override a port's access filter, enable access override on that port. In this case, the process is as follows:

1. Access is denied by the access filter.

2. The user is prompted for a username and password.

3. The user is verified by the user table or RADIUS.

4. The access filter defined for this user is used to determine if the user has permission to access the specified host.

To enable users to override a port access filter with their own filter, use the following command:

```
Command> set S0 access on
```

Using External Modems 9

This chapter explains how to configure external modems to work with the PortMaster 4. The information in this chapter does not apply to the internal digital modems that come installed on Quad T1 and Tri E1 boards with modems.

Note – ComOS 4.0 and ComOS 4.1 do not support the **add modem**, **delete modem**, and **show modem** commands for external modems. If you are running ComOS 4.0 or ComOS 4.1 on the PortMaster 4, you cannot use the commands in this chapter to configure external modems on asynchronous ports C0 and C1. Instead, you can attach a previously configured modem to the C0 or C1 ports. Although you cannot display the modem's settings with the **show modem** command, it will function if properly configured.

This chapter discusses the following topics:

- "Modem Ports" on page 9-1

- "Modem Functions" on page 9-2

- "Using Automatic Modem Configuration" on page 9-2

- "Configuring Ports for Modem Use" on page 9-5

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

Because the PortMaster is a DTE device, a straight-through RS-232 cable is used to connect modems to it. Straight-through cables for modems use pins 2, 3, 4, 5, 6, 7, 8, and 20. See the *PortMaster 4 Installation Guide* for modem cable information.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMConsole interface to ComOS.

Modem Ports

The PortMaster 4 supports external modems on two asynchronous ports (C0 and C1), which are physically located on the manager module. Although you can attach a modem to port C0, it is primarily designed to be used for a console connection with a null modem cable. Port C1 is a standard asynchronous interface designed for connection to an external modem.

To connect these ports to a terminal or other DTE, use a null modem cable, typically male-to-female. Directions (input/output) are with respect to the PortMaster. The PortMaster does not use the Data Set Ready (DSR) signal.

Note – When the console port is connected to a terminal, it uses software flow control and therefore requires pins 2, 3, and 7 only.

Null modem cables can be obtained from most suppliers of computer equipment.

Dial-up modems that operate over normal telephone lines at speeds of 28,800bps or higher are now available. These modems do not operate at a guaranteed throughput, but rather at a speed dependent on the quality of the line, the effectiveness of data compression, and other variables. These modems use hardware flow control to stop the data from the host by raising and lowering the Clear to Send (CTS) signal.

PortMaster products support hardware flow control using the RTS output signal and the CTS input signal, which is also used by the normal modem handshake.

Modem Functions

Configure modems to do the following for use with the PortMaster:

- Raise DCD when a call comes in

- Reset itself when DTR is dropped

- Lock the DTE speed

- Use hardware flow control (RTS/CTS)

Using Automatic Modem Configuration

PortMaster products use a modem table to automate the external modem configuration process. The modem table is user-configurable and includes long and short modem names, preferred DTE rate, and the modem initialization string. For convenience, the table is preconfigured by Lucent for many common modems.

When you specify the name of the modem and the attached port, the PortMaster automatically configures the modem for you, provided the modem is in the factory default state when it is initialized.

After a modem type has been specified, the PortMaster automatically sets the port for hardware flow control, the correct speed, and modem control when the port is reset.

Displaying Modem Settings and Status

To display the external modems currently configured in your modem table, use the following command:

    Command> **show table modem**

A modem table display looks like the following:

```
Short Name          Long Name                         Type
--------------      -------------                     ------
cardinal            Cardinal MVP288XF                 System
mega                Massive MegaFast                  User
supra-288           Supra V.34                        System
```

The modem **type** is either system or user. *System* indicates that the configuration settings are the factory default settings. *User* indicates that the user has configured the modem table settings for that modem.

To display the settings for a particular modem, use the following command:

    Command> **show modem** *ModemName(short)*

The display for a modem looks like this:

```
   Short Name: supra-fax-288
    Long Name: SupraFax 28.8
Optimal Speed: 115200
         Type: User Defined
  Init Script: Send Command                              Wait for Reply
               ---------------------------------------   ------------------
               AT&F2&C1&D3S0=1S2=129s10=20&W             OK
```

Adding a Modem to the Modem Table

To add a modem to the modem table, use the following command:

    Command> **add modem** *ModemName(short)* **"***ModemName(long)***"** *Speed* **"***String***"**

For example, to add a Paradyne 3811+ modem to the modem table, enter:

    Command> **add modem para3811 "Paradyne 3811+" 115200 "AT&FS0=1&W\r^OK"**

Note – Use a **\r** for a carriage return, and a caret (**^**) to separate the send and expect characters in the string. In the example above, the PortMaster expects **OK**. Never use **on** or **off** for a modem short name.

Table 9-1 shows the current factory default settings for commonly used modems.

Table 9-1   Factory Default Modem Table Entries

Modem Name (Short)	Modem Name (Long)	DTE Rate	Initialization String
at&t-v32	AT&T Keep In Touch	57600	AT&F&D3&T5&R0\\D1S0=1&W^OK
cardinal	Cardinal MVP288XF	115200	AT&F1&C1&D2&K3S0=1S2=129S10=20&W0&W1
card-v34-p	Cardinal MVP288CC PCMCIA	115200	AT&F&C1&D3S0=1s2=129S10=20&W
eiger-v32-p	Eiger 14.4 PCMCIA	57600	AT&F&C1&D3S0=1S10=20&W
eiger-v34-p	Eiger 28.8 PCMCIA	115200	AT&F&C1&D3S0=1S10=20&W
gvc-14.4	GVC/Maxtech V.32	57600	AT&F&C1&D3S0=1S10=20&W0
gvc-28.8	GVC/Maxtech V.34	115200	AT&F&C1&D3S0=1S10=20&W0

Table 9-1   Factory Default Modem Table Entries *(Continued)*

Modem Name (Short)	Modem Name (Long)	DTE Rate	Initialization String
hay-cent2	Hayes Century 2 Rack V.32bis	115200	AT&F&C1&D2&K3S0=1S10=20&W0
intel-v32-p	Intel V.32bis PCMCIA	115200	AT&F&C1&D3S0=1&W&W1^\rOK
megahz-v32-p	Megahertz XJ2288 V.34bis PCMCIA	115200	AT&F&C1&D3S0=1&W
megahz-v32-p	Megahertz XJ2288 V.34bis PCMCIA	115200	AT&F&C1&D3S0=1&W
micro-desk	Microcom 28.8	115200	AT&F&C1&D2$B115200\\Q3%U1&T5S0=1S10=20*W0&Y0
mot-uds	Motorola UDS V.34	115200	AT&F&C1&D2\\Q3S0=1S10=20S80=18&W
mot-bit	Motorola Bitsurfr	115200	AT&F&C1&D2%A4=1%A2=95&m0@P2=115200@P1=a&W
mot-pwr-p	Motorola Power 14.4 PCMCIA	57600	AT&F&C1&T5&C1&D2&W
mot-life-p	Motorola Lifestyle 14.4 PCMCIA	57600	AT&FS0=1&C1&D2\\Q3&T5&W^OK
multizdx	MultiTech Z/DX fax/data v.32	115200	AT&F^ATM0&E1&C1&D3$SB115200S0=1S10=20%E0&W0
multi-v34	MultiTech MT2834 28.8k	115200	AT&F^AT&C1&D3S0=1&W0
multi-v34	MultiTech MT2834 28.8k	115200	AT&F^AT&C1&D3S0=1&W0
pp-v32	Practical Peripherals PP9600SA	57600	AT&F&C1&D3S0=1S2=129&W
pp-v34	Practical Peripherals PM288T II	115200	AT&F0M0S0=1V1&C1&D3&K3&W0&W1
para3811	Paradyne 3811+	115200	AT&FS0=1&W
ppi-v34-p	PPI ProClass V.34 PCMCIA	115200	AT&F&C1&D3&K3S0=1&W&W1
premax-v32-p	Premax V.32bis PCMCIA	115200	AT&F&C1&D3S0=1&W&W1
scout-v32-p	DSI Scout V.32bis PCMCIA	115200	AT&F&C1&D3S0=1&W
supra-288	Supra V.34	115200	AT&F2S0=1&W
supra-fax-288	SupraFax 28.8	115200	AT&F2&C1&D3S0=1S2=129s10=20&W
tdk-288-p	TDK DF2814 V.Fast PCMCIA	115200	AT&F&C1&D3S0=1&W

Table 9-1   Factory Default Modem Table Entries *(Continued)*

Modem Name (Short)	Modem Name (Long)	DTE Rate	Initialization String
usr-v32-p	USR Courier/Sportster V.32bis PCMCIA	57600	AT&F1&W
usr-v34-p	USR Courier/Sportster V.34 PCMCIA	115200	AT&F1S0=1&W
usr-v32	USR Courier/Sportster V.32bis	57600	AT&F1S0=1&W
usr-v34	USR Courier/Sportster V.34	115200	AT&F1S0=1&W
usr-spt-v32	USR Sportster V.32bis	57600	AT&F1S0=1S10=20S13.0=1&W0
usr-spt-336	USR Sportster 33.6	115200	AT&F1S0=1S10=20S13.0=1&W0
zyxel	Zyxel U1496E	57600	AT&FM0&D2S0=1S2=1

Associating a Modem with a Port

To automatically configure a modem and associate it in the modem table with the port it is attached to, use the following commands:

```
Command> set C0|all modem-type ModemName(short)
Command> reset C0|all
```

For example; to associate a U.S. Robotics V.34 modem with port C0 and configure the modem, enter the following commands:

```
Command> set C0 modem-type usr-v34
Command> reset C0
```

To configure all ports for the same modem type, use **all** instead of the port number in the previous example. After the modem is attached to the port, configure the other modem settings described in "Configuring Ports for Modem Use" on page 9-5.

To configure the modem **not** to answer when users dial in, set *C0*=**0** in the initialization string.

Configuring Ports for Modem Use

The modem settings described in this section are configured for each port and must match the configuration on the attached modem.

Setting the Port Speed

The speed of a port is defined as the DTE baud rate. The PortMaster allows you to specify three different baud rates for each port and one baud rate for host device ports. Port speeds are sequentially matched from the first baud rate through the third baud rate.

For example, when a connection with this port is established, the PortMaster uses the first baud rate value to try to synchronize the connection speed. If no synchronization is possible, the PortMaster tries to synchronize speeds using the second baud rate value. If this fails, the third baud rate value is used. Each speed can be set between 300bps and 115200bps. The default speed is 9600bps.

Modern modems and terminals must always be set to run at a fixed rate. To define a fixed rate, lock the DTE rate by setting all three speeds to the same value.

To set the port speed, use the following command—entered on one line:

    Command> **set** *C0*|**all speed [1|2|3]** *Speed*

You can substitute any of the following for *Speed*:

300	**1200**	**4800**	**19200**	**57600**	**115200**
600	**2400**	**9600**	**38400**	**76800**	

You can set the speed for all the asynchronous ports simultaneously by using the **set all speed** command.

Setting Modem Control

Set modem control on if you want to use the DCD signal for modem connections. When modem control is on, the PortMaster uses the condition of the carrier detect line to determine whether the line is in use. Modem control must be on for PortMaster outbound traffic. If modem control is off, the PortMaster assumes the carrier detect line is always asserted. As a result, the PortMaster cannot attach to the modem for outbound traffic because it regards the line as busy.

To set modem control, use the following command:

    Command> **set** *C0* **cd on|off**

Setting Parity

The parity setting must be configured to match the parity setting on the attached modem. The parity default value is **none** and must be used for ports configured for network dial-in or dial-out operation.

Table 9-2 describes the parity options.

Table 9-2   Parity Options

Option	Description
none	Assumes 8 data bits, 1 stop bit, and no parity bit. This is the default.
even	Assumes 7 data bits, 1 stop bit, and even parity.
odd	Assumes 7 data bits, 1 stop bit, and odd parity.
strip	Assumes 8 data bits and 1 stop bit. The parity bit is stripped from the data stream when it is received by the PortMaster

To set the parity for a modem and its port, use the following command:

Command> **set** *C0* **parity even|none|odd|strip**

Setting the Flow Control

The PortMaster supports both software flow control and hardware flow control. Software flow control uses the ASCII control characters DC1 and DC3 to communicate with the attached device and to start and stop the flow of data.

To set software flow control for a modem, use the following command:

Command> **set** *C0* **xon/xoff on|off**

Hardware flow control allows the PortMaster to receive data from the attached device by raising the Request to Send (RTS) signal on pin 4 of the RS-232 connector. The PortMaster sends information to the attached device only when the Clear to Send (CTS) modem line on pin 5 of the RS-232 connector is raised.

To set hardware flow control for a modem, use the following command:

Command> **set** *C0* **rts/cts on|off**

Note – Because it is more reliable, you should always use hardware flow control if it is available. Do not use both hardware and software flow control on the same port.

Hanging Up a Line

You can specify whether the DTR signal is dropped and the modem disconnected after a session is terminated. If line hangup is enabled and the session is terminated, DTR is held low, signaling the modem to disconnect. If line hangup is disabled, the DTR signal does not drop and the modem does not hang up when the user session terminates.

To set line hangup for a modem, use the following command:

Command> **set** *C0* **hangup on|off**

Note – Resetting the port administratively with the **reset** command always drops DTR.

Configuring T1, E1, and ISDN PRI 10

This chapter describes how to configure T1 or E1 lines on the PortMaster 4 for the following kinds of service:

T1 Line0 through Line3	E1 Line0 through Line2
Full T1	Full E1
Fractional T1	Fractional E1
Channelized T1	Multifrequency R2 (MFR2) signaling for channelized E1
ISDN Primary Rate Interface (PRI)	ISDN PRI
	Fractional PRI

Quad T1 and Tri E1 boards with True Digital modems have the digital modems and T1 or E1 circuits physically present on the same board.

Note – After making any configuration changes to a line (Line0 through Line3), you must use the **save all** and **reset slot** commands for the changes to take effect.

This chapter discusses the following topics:

- "PortMaster 4 Quick Setup Guide for ISDN PRI" on page 10-2

- "Configuring General Settings" on page 10-5

- "Configuring Fractional Settings" on page 10-6

- "Configuring ISDN PRI Settings" on page 10-7

- "Configuring True Digital Modems" on page 10-10

- "Configuring Channelized T1 or E1" on page 10-11

- "Using NFAS for ISDN PRI" on page 10-15

- "Configuring SS7" on page 10-20

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

PortMaster 4 Quick Setup Guide for ISDN PRI

This section provides a procedure to help you quickly get T1 or E1 lines configured for ISDN PRI service. The procedure uses an example to configure Line0 for PRI. After completing this configuration, you can configure the line for channelized T1, Frame Relay, or any supported protocol.

Substitute your IP addresses, system names, passwords, and so on for the italicized variables in the sample commands.

See "Configuring General Settings" on page 10-5 for details.

The PRI configuration involves the following three procedures:

- "Configuring the Ethernet Interface" on page 10-2

- "Configuring Global Parameters" on page 10-3

- "Configuring the Quad T1 Boards" on page 10-4

Configuring the Ethernet Interface

Use the console port to log in to the PortMaster, or log in as **!root** and press the **Return** key twice to get to the command prompt. Then follow this procedure, substituting your own information for variables.

1. **Set the system name of the PortMaster 4.**

   Command> **set sysname** *ISP-PM4-1*

2. **Set the address of Ether0 10BaseT port.**

   Command> **set ether0 address** *192.168.10.1*

3. **Set the netmask of Ether0.**

   Command> **set ether0 netmask** *255.255.255.0*

4. **Set Ether0 broadcast to high.**

   Command> **set ether0 broadcast high.**

5. **(Optional) Set Ether0 to listen for and broadcast RIP-1 packets.**

   Command> **set ether0 rip on**

   Refer to the *PortMaster 4 Command Line Reference* for additional RIP-1 options, or for configuring RIP-2.

6. **Save the settings and reboot.**

   Command> **save all**
   Command> **reboot**

7. **Set the address of the Ether1 10/100BaseT port.**

   Command> **set ether1 address** *192.168.100.1*

8. **Set the netmask of Ether1.**

   Command> **set ether1 netmask** *255.255.255.0*

9. **Set Ether1 broadcast to high.**

   Command> **set ether1 broadcast high**

10. **(Optional) Set Ether0 to listen for and broadcast RIP-1 packets.**

    Command> **set ether1 rip on**

11. **Save the settings and reset the slot.**

    Command> **save all**
    Command> **reset slot10**

12. **Set the default gateway address.**

    Command> **set gateway** *192.168.100.254*

Proceed to global configuration.

Configuring Global Parameters

You can now connect to the PortMaster 4 via Telnet or PMVision to continue the configuration process. You can perform the following command line configuration via the console or by using Telnet. Substitute your own information for variables.

1. **If you are not already in slot 4, set the view to slot 4.**

   Command> **set view 4**

2. **Set the IP address for** syslog **authentication information.**

   Command> **set loghost** *192.168.1.2*

3. **(Optional) Set a domain name.**

   Command> **set domain** *ISP.net*

4. **(Optional) Set the Domain Name System (DNS).**

   Command> **set namesvc dns**

5. **(Optional) Set the IP address of the name server.**

   Command> **set nameserver** *192.168.25.16*

6. **(Optional) Set RIP broadcast to high.**

   Command> **set ether0 broadcast high**

7. **Do not broadcast default routes.**

   Command> **set default off**

8. **Set the IP address of the RADIUS server.**

   Command> **set authentic** *192.168.120.10*

9. **Set the RADIUS password.**

   Command> **set secret** *String*

10. **Set the IP address of the RADIUS accounting server.**

   Command> **set accounting** *192.168.120.10*

11. **Use RADIUS-provided netmasks.**

   Command> **set user-netmask on**

12. **Set the ChoiceNet server address.**

   Command> **set choicenet** *192.168.120.10*

13. **Set the ChoiceNet password.**

   Command> **set choicenet-secret** *String*

14. **Save the configuration to nonvolatile memory.**

   Command> **save all**

Configuring the Quad T1 Boards

To configure a Quad T1 board, you set the view to the slot where the Quad T1 board is installed. Substitute your own information for variables.

1. **Set the slot to 0.**

   Command> **set view 0**

2. **Set the base address of an address pool to slot 0.**

   Command 0> **set assigned_address** *192.168.0.1*

3. **Set the number of IP addresses in the pool.**

   Command 0> **set pool** *92*

4. **Set the switch type for slot 0.**

   Command 0> **set isdn-switch dms-100**

5. **Set Line0 to PRI, slot 0.**

   Command 0> **set line0 isdn**

6. **Set Line1 to PRI, slot 0.**

   Command 0> **set line1 isdn**

7. **Set Line2 to PRI, slot 0.**

   Command 0> **set line2 isdn**

8. **Set Line3 to PRI, slot 0.**

   Command 0> **set line3 isdn**

9. **Save the configuration to nonvolatile RAM.**

   Command 0> **save all**

10. **Reset the slot.**

    Command 0> **reset slot0**

11. **Repeat the previous 10 steps for each installed Quad T1 board.**

    You must change the view to the appropriate slot.

Configuring General Settings

Configure the following general settings for T1, E1, or ISDN PRI lines on PortMaster 4.

Setting the View

To configure a T1, E1, or ISDN PRI line, you must first set the view. Enter the **show boards** command to determine the identification number of the line board you want to configure.

The board identification number is the same as the number of the slot in which the T1 or E1 board is installed.

   Command> **show boards**

Use the following command to set the view to a slot with an installed T1 or E1 line board:

   Command> **set view** *Slotnumber*

Setting the view for a specific board gives you administrative access to that board.

Displaying Line Status

To display the status of a E1 or T1 line, use the following command:

   Command> **show** *Line0*

Configuring Line Use

You can use a line as a single E1 or T1 line; as PRI B channels; as a fractional E1, ISDN, or T1 line divided into channel groups; or for in-band signaling for channelized T1 or E1.

Note – T1 and E1 lines require an external clock signal provided by the device to which the PortMaster is connected, or by the telephone company network.

To configure a line, use the following command. Table 10-1 explains the line use options.

> Command> **set** *Line0*|**isdn**|**t1**|**e1**|**fractional**|**isdn-fractional**|**inband**

Table 10-1   Line Use Options

Options	Descriptions
isdn	Configures the line as ISDN B channels. This is the default.
t1	Configures the entire line as a T1 line.
e1	Configures the entire line as an E1 line.
fractional	Allows a channelized T1 or E1 line to be divided into groups (see "Setting Channel Groups" on page 10-6).
isdn-fractional	Allows an ISDN PRI line to be divided into groups (see "Setting Channel Groups" on page 10-6).
inband	Sets the channelized T1 or E1 line for inband signaling.

⚠ **Caution –** If you configure a line for fractional T1 or fractional ISDN and reset the board before configuring the group and channels, you will no longer be able to see and configure the line.

You use the **fractional** keyword in this command to break up a channelized T1 line into groups. The **isdn-fractional** keyword refers to PRI only.

Configuring Fractional Settings

The PortMaster 4 supports fractional service on T1, E1, or ISDN PRI lines. To configure a line for fractional use, you must create channel groups and assign channel numbers. You must also set the channel rate for a fractional T1 or E1 line.

Setting Channel Groups

You can divide the channels of a T1, E1, or ISDN PRI line into numbered groups after the line type has been set to fractional with the **set** *Line0* **fractional** command.

To set the channel group for a T1, E1, or ISDN PRI line, use the following command. Table 10-2 explains the channel group options.

> Command> **set** *Line0* **group** *Cgroup* **channels** *Channel-list*

Table 10-2   Channel Group Options

Option	Description
Line0	Line0 through Line3 (T1), or Line0 through Line2 (E1).

Table 10-2   Channel Group Options *(Continued)*

Option	Description
Cgroup	Group number from 1 to 63 that designates a port number on each T1, E1, or ISDN PRI line, or **none** to unassign channels.
Channel-list	Space-separated list of one or more channel numbers, from 1 through 24 for T1, or 1 through 30 for E1. The channel numbers do not have to be contiguous.

Setting the Channel Rate

Table 10-3   To set the channel rate to 56Kbps or 64Kbps for a channel group, use the following command. Table 10-4 explains the channel rate options.

```
Command> set Line0 group Cgroup 56k|64k
```

Table 10-4   Channel Rate Options

Option	Description
Line0	Line0 through Line3 (T1), Line0 through Line2 (E1).
Cgroup	Defined channel group from 1 to 63.
56k	56Kbps, typically used for D4 framing.
64k	64Kbps, used for framing types other than D4. This is the default.

Configuring ISDN PRI Settings

Use the following settings to configure ISDN PRI on the PortMaster 4.

Setting the ISDN PRI Switch

The switch type information is available from your ISDN PRI service provider. To set the switch type for ISDN connections to the PortMaster ISDN PRI virtual ports, use the following command—entered on one line. Table 10-5 explains the ISDN switch options.

```
Command> set isdn-switch ni-2|dms-100|4ess|att-5ess
|net5|vn2|vn3|1tr6|ntt|kdd|ts014
```

Table 10-5   ISDN Switch Options

ISDN Switch	Description
ni-2	National ISDN-2 (NI-2) compliant. This is the default.
dms-100	Northern Telecom DMS-100 Custom.
4ess	AT&T 4ESS.
att-5ess	AT&T 5ESS.

Table 10-5   ISDN Switch Options *(Continued)*

ISDN Switch	Description
net5	European ISDN PRI standard.
vn2	France—older switch.
vn3	France—older switch.
1tr6	Germany—older switch.
ntt	Japan.
kdd	Japan.
ts014	Australia. To use this switch type, set the port type to network hardwired, set the directory number for the port appropriately, and reset the port.

Setting the Framing Format

To set the framing format used for the E1 or T1 line, use the following command. Table 10-6 explains the framing format options.

    Command> **set** *Line0* **framing esf|d4|crc4|fas**

Table 10-6   T1 and E1 Framing Format Options

Option	Description
Line0	Line0 through Line3 (T1), Line0 through Line2 (E1).
esf	Extended superframe. This is the default format for T1 lines.
d4	D4 framing, an alternative format for T1 lines.
crc4	Cyclic redundancy check 4. This is the default format for E1 lines.
fas	Frame Alignment Signal, an alternative format for E1 lines.

Setting the Encoding Method

This command sets the encoding method used with T1 and E1 lines. Table 10-7 explains the encoding method options.

    Command> **set** *Line0* **encoding b8zs|ami|hdb3**

Table 10-7   T1 and E1 Encoding Method Options

Option	Description
Line0	Line0 through Line3 (T1), Line0 through Line2 (E1).
b8zs	Bipolar 8-zero substitution. This is the default for T1 lines.
ami	Alternate mark inversion.
hdb3	High-density bipolar 3. This is the default for E1 lines.

Setting the Pulse Code Modulation

You need to set the pulse code modulation only if you are using digital modems and your PRI service provider instructs you to change the setting to something other than the default. This command sets the method for "companding"—compressing and expanding—the amplitude of analog signals.

To set the pulse code modulation, use the following command. Table 10-8 explains the pulse code modulation options.

> Command> **set** *Line0* **pcm u-law|a-law**

Table 10-8   T1 and E1 Pulse Code Modulation Options

Option	Description
Line0	Line0 through Line3 (T1), or Line0 through Line2 (E1).
u-law	Default method for T1 PRI lines.
a-law	Default method for E1 PRI lines.

Setting the Loopback

You can test the telephone line of your T1 or E1 ISDN connection by setting the local network loopback.

To set the loopback, use the following command:

> Command> **set** *Line0* **loopback on|off**

Setting the Directory Number

Normally, a T1 or E1 line has a single telephone number. However, when the line is set up as ISDN B channels, you can set a telephone number for an individual virtual port. This feature allows you to identify the circuit telephone number associated with a specific ISDN PRI port.

To set a telephone number for an individual port when the line is configured as ISDN B channels, use the following command. Table 10-9 explains the directory number options.

> Command> **set** *S0* **directory** *Number*

Table 10-9   Directory Number Options

Options	Description
S0	One of the virtual ISDN PRI ports.
Number	Access telephone number.

Configuring True Digital Modems

Each Quad T1 board with internal modems has 34 modems installed on it plus 64 modems arranged on a daughterboard that plugs directly into the Quad T1 board. There are 96 active modems per board, plus 2 modems acting as hot spares.

Similarly, the Tri E1 board with internal modems has 98 modems per board, with 90 active modems and 8 hot spares.

You can install up to nine Quad T1 or Tri E1 boards into the PortMaster 4 chassis, for a total of 882 modems—864 active modems plus 18 hot standby modems for Quad T1 boards, or 810 active modems plus 72 hot standby modems for Tri E1 boards.

All Quad T1 and Tri E1 boards can be hot swapped.

Note – Digital modems require no configuration or initialization string.

Use the **show** *M0* and **show modems** commands to display modem status.

Setting Digital Modems to Analog Service

When analog modem service is required for dial-out network connections, you can convert the analog service to digital service.

To set the digital modems to analog modem service for the specified location, use the following command. Table 10-10 explains the analog modem options.

    Command> **set location** *Locname* **analog on|off**

Table 10-10 Analog Modem Options

Option	Description
Locname	Location name that is in the location table.
on	Enables analog modem service on dial-out.
off	Disables analog modem service on dial-out, and causes the service to revert to ISDN.

Displaying Modem Status

Use the following command to display the settings for a particular modem:

    Command> **show** *M0*

You can display the status for all digital modems. Modem states are as follows:

- ACTIVE—in use

- READY—available for use

- ADMIN—busy

- TEST—under test

- DOWN—unavailable

To display the status for all modems, use the following command:

    Command> **show modems**

Troubleshooting Digital Modems

The **debug** command is useful for troubleshooting the digital modems and Multichassis PPP events on the PortMaster 4. Output is sent to the system console set by the **set console** command. After completing the debugging process, disable the **debug** commands by using the correct **set debug off** command, and reset the console with the **reset console** command. Debug information is displayed to the console.

To set debug flags used for troubleshooting, use the following command—entered on one line:

    Command> **set debug mdp-status|mdp-events on|off**

Table 10-11 explains the debug options for the PortMaster 4

Table 10-11 Debug Options for the PortMaster 4

Option	Description
mdp-status	Set **on** to display the status of the digital modems.
mdp-events	Set **on** to display the progress of the digital modems as they initialize.

Configuring Channelized T1 or E1

The PortMaster 4 supports channelized T1 service on the Quad T1 board and channelized E1 service on the Tri E1 board.

Channelized T1 Service

Channelized T1 service provides 24 channels of 56Kbps capacity each. In contrast, an ISDN PRI line provides 23 channels of 64Kbps capacity each—plus one 64Kbps signaling channel. However, channelized T1 is available in many service areas that do not yet provide ISDN PRI. In areas where PRI is available, the cost of channelized T1 can be significantly less than the cost of PRI.

Each Quad T1 board on the PortMaster 4 has an integrated channel service unit/digital service unit (CSU/DSU). However, the other end of a T1 connection might require an external clock signal provided by the telephone company, or a CSU/DSU.

How to Order DS-1 Service from the Telephone Company

The telephone company will ask you the following two questions when you order digital service level 1 (DS-1) service:

- What signaling protocol do you use?

  You can use any one of the following signaling protocols on the PortMaster 4:

  – **E & M wink start**

  – **E & M immediate start**

  – **Foreign exchange station (FXS)**

- If you use E & M wink start, how many Directory Number Identification Service (DNIS) digits do you need?

  The PortMaster 4 requires one DNIS digit.

Record the line parameters provided by the telephone company.

Setting the In-Band Signaling Protocol for T1

To set the in-band signaling protocol and the in-band call options used with channelized T1, use the following command. Table 10-12 explains the in-band signaling protocol options.

> Command> **set** *Line0* **signaling wink|immediate|fxs**

Table 10-12 T1 In-Band Signaling Protocol Options

Option	Description
Line0	Line0 through Line3.
wink	E & M wink start protocol, an option for use with T1 lines. This is the T1 default.
immediate	E & M immediate start protocol, used with T1 lines.
fxs	Foreign exchange station (FXS) loop start protocol used with T1 lines.

Configuring the PortMaster 4 for Channelized T1

Follow these steps to configure a Quad T1 board on the PortMaster 4 to use channelized T1 service:

1. **Set the view**

   Command> **set view** *Slotnumber*

2. **Set the line for in-band signaling.**

   Command *Slotnumber*> **set** *Line0* **inband**

3. **Set the signaling protocol and the line provisioning.**

    Command *Slotnumber*> **set** *Line0* **signaling wink|immediate|fxs**

4. **Set the framing format for the line.**

    Command *Slotnumber*> **set** *Line0* **framing esf|d4**

5. **Set the encoding method for the line.**

    Command *Slotnumber*> **set** *Line0* **encoding b8zs|ami**

6. **Repeat Steps 2 through 5 for the other three T1 lines on the Quad T1 board.**

7. **Save the configuration changes and reset the slot.**

    Command *Slotnumber*> **save all**
    Command *Slotnumber*> **reset slot***Slotnumber*

8. **Use the following command to display the configuration for each line:**

    Command *Slotnumber*> **show** *Line0*

Example Channelized T1 Configuration

This example configures Line1 on a Quad T1 board for channelized T1 service using
E & M wink start, extended superframe format, and bipolar 8-zero substitution.

    Command> **set view 2**
    Command 2> **set line1 inband**
    Command 2> **set line1 signaling wink**
    Command 2> **set line1 framing esf**
    Command 2> **set line1 encoding b8zs**
    Command 2> **save all**
    Command 2> **reset slot2**

The following example display shows the output from the **show line1** command for
this configuration:

    Command 2> **show line1**

    ---------------------line1 - T1 Inband DS0 ------------------
    Status: UP Framing: ESF Encoding: 8ZS PCM: u-law
    Signaling: Trunk E&M wink start Options: inbound calls only
    Receive Level: +2dB to -7.5dB
    Alarms Violations

    Blue 0   Bipolar 0
    Yellow 1   CRC Errors 0
    Receive Carrier Loss 0   Multiframe Sync 0
    Loss of Sync 0

Channelized E1 Service

Channelized E1 service is a digital standard used outside the United States. E1 technology is an improvement over T1 because it is slightly faster and contains 32 channels of 64Kbps capacity each. One channel is reserved for administrative uses, and one channel is used for signaling.

Setting the In-Band Signaling Protocol for E1

Although PortMaster products do not require dial digits (the calling number and caller ID) when establishing a connection, most telephone companies transmit this information by default. You can use the **r2generic** signaling option if you do not require dial digits, but you must first arrange for the telephone company to not transmit these signals.

The PortMaster defaults to **r2generic** when you set the line to in-band (see "Configuring Line Use" on page 10-5).

To accept caller ID and dial digit tones, use the **mrf2** option. Because some countries implement different variations of multifrequency robbed bit signaling (MFR2), you must specify a profile with the **mfr2** option.

To set the in-band signaling protocol and in-band call options for channelized E1, use the following command. Table 10-13 explains the in-band signaling protocol options and profiles.

    Command> **set** *line0* **signaling r2generic|mfr2** *Profile*

Table 10-13 E1 In-Band Signaling Protocol Options

Option	Profile	Description
Line0		Line0 through Line2.
r2generic		Generic R2, the default; no caller ID and dial digit tones are exchanged.
mfr2		Accept caller ID and dial digit tones.
	0	ITU standard; used in Argentina, Saudi Arabia, and other countries. This is the default.
	1	Mexico.
	2	Brazil and Tunisia.
	3	Venezuela.
	4	Mexico. Profile 4 is a subset of profile 1 and is used with switches that do not support caller ID. This profile can be used in Mexico whenever profile 1 is used, but the reverse is not true.

Configuring the PortMaster 4 for Channelized E1

Follow these steps to configure a Tri E1 board on the PortMaster 4 for channelized E1 service:

1. **Set the view**

   Command> **set view** *Slotnumber*

2. **Set the line for in-band signaling.**

   Command *Slotnumber*> **set** *Line0* **inband**

3. **Set the signaling protocol and the line provisioning.**

   Command *Slotnumber*> **set** *Line0* **signaling r2generic|mfr2** *Profile*

4. **Set the framing format for the line.**

   Command *Slotnumber*> **set** *Line0* **framing crc4|fas**

5. **Set the encoding method for the line.**

   Command *Slotnumber*> **set** *Line0* **encoding hdb3|ami**

6. **Repeat Steps 2 through 5 for the other two lines on the Tri E1 board.**

7. **Save the configuration changes and reset the slot.**

   Command *Slotnumber*> **save all**
   Command *Slotnumber*> **reset slot***Slotnumber*

8. **Use the following command to display the configuration for each E1 line:**

   Command *Slotnumber*> **show** *Line0*

Using NFAS for ISDN PRI

Non-facility associated signaling (NFAS) is an ISDN PRI protocol that allows you to define one or two D channels to carry signaling messages for up to 20 T1 lines, or **interfaces**. This feature relieves telephone companies and Internet service providers (ISPs) of the need to provide D channel signaling for each T1 interface, and increases bandwidth by making those D channels available to carry data.

Most telephone companies offer two varieties of NFAS:

• **Standard NFAS**—A PRI service in which one T1 interface provides D channel signaling for an NFAS group of up to 20 interfaces

• **NFAS with D channel backup (DCBU)**—A PRI service in which two T1 interfaces provide D channel signaling for an NFAS group of up to 20 T1 interfaces

  In the D channel backup system, one T1 interface is configured as the primary interface and another is configured as the secondary interface. If the primary interface fails, the secondary interface takes over signaling responsibilities for the group. When the failed primary interface returns to service, it backs up the secondary interface.

The Lucent ComOS implementation of NFAS supports both standard NFAS and NFAS with D channel backup, but recommends NFAS with backup.

Understanding Standard NFAS

If you have NFAS without backup, you gain one B channel. The drawback is that if the primary interface fails, you have no backup D channel and no calls are possible on any of the lines in the group until the primary interface reactivates.

To enable NFAS without backup, you define an NFAS group and configure one T1 line as the primary interface and the remaining T1 lines as slave interfaces. See "Standard NFAS" on page 10-20 for an example configuration.

Understanding NFAS with D Channel Backup

To enable NFAS with D channel backup, you define an NFAS group and configure one T1 line as the primary interface, another T1 line as the secondary interface, and the remaining T1 lines as slave interfaces. The primary T1 interface carries signaling messages for its own interface, the secondary interface, and all slave interfaces in the NFAS group.

When you reset the slot of a PortMaster 4 configured for NFAS with backup, the D channels on the primary and secondary interfaces initialize in "out of service" mode. The switch then puts the D channel on the primary interface in "in service" mode and the D channel on the secondary interface in "standby" mode. As call traffic commences on the T1 interfaces, the primary D channel handles signaling messages for all channels in the interface group, which typically include the primary T1 interface, the secondary T1 interface, and other T1 interfaces configured as slave interfaces.

If the primary interface fails, all calls in process are dropped on all interfaces serviced by that D channel. The D channel on the secondary interface switches to "in service" mode and begins to carry signaling messages for channels on the secondary T1 interface and all other slave interfaces previously serviced by the primary T1 interface. Call traffic does not resume until the D channel on the secondary interface switches to "in service" mode and begins carrying signaling messages.

Meanwhile, the switch attempts repeatedly to activate the primary interface. When the primary interface restarts, the D channel on that interface goes into "standby" mode and does not preempt the "in service" function from the secondary interface. Message signals for the reactivated primary T1 interface are carried by the D channel on the secondary T1 interface.

Multichassis Capacity

The ComOS implementation of NFAS is designed for use across multiple PortMaster 4 slots configured as a group on the same Ethernet. A group is an arbitrary number between 1 and 99 that you assign to an interface. You can define multiple groups of T1 interfaces on the same Ethernet segment, but each group must be supported by its own primary and secondary D channel pair.

NFAS message signaling travels over Ethernet using the User Datagram Protocol (UDP) and UDP port 1650. A reliable, proprietary protocol provides packet sequencing, acknowledgment for packets, and retransmission of lost packets.

Fault Tolerance

When you configure NFAS with D channel backup, you set one line as the primary interface, one line as the secondary interface, and the remaining lines in the group as slave interfaces. To increase the fault tolerance of the group, you can set the secondary interface on a different Quad T1 board from the primary interface.

With the primary and secondary interfaces on separate boards, all calls in the group are dropped when the primary interface board fails, but service resumes as soon as the secondary interface assumes D channel signaling responsibilities. See "Configuring NFAS with D Channel Backup" on page 10-17 for instructions on this type of fault-tolerant configuration.

NFAS Limitations

Each Quad T1 board can handle only one interface group and one type of signaling:

- The four T1 interfaces on any Quad T1 line board cannot belong to different groups.

- When you configure NFAS on one T1 interface, the other T1 interfaces cannot run in standard PRI mode.

However, you can configure more than one Quad T1 board in the same group.

Provisioning

Because NFAS requires additional control command exchanges, NFAS T1 interfaces are provisioned differently at the switch. To help you determine the kind of provisioning you require for ISDN setup, refer to the information on the Lucent website at **http://www.livingston.com**.

Configuring NFAS

To configure NFAS on a T1 line, use the following command:

> Command> **set** *Line0* **nfas primary|secondary|slave|disabled** *Identifier Group*

See the *PortMaster 4 Command Line Reference* for syntax information.

Configuring NFAS with D Channel Backup

This section describes how to configure NFAS with the D channel backup interface set, for fault tolerance, on a different Quad T1 board from the primary interface. The primary interface is set on Line0 in slot 0 for NFAS group 5. The secondary interface is set on Line0 in slot 1. All other T1 interfaces on these two boards on the PortMaster 4 are set as slave interfaces for this group.

If you configure NFAS without D channel backup, do not configure a secondary interface. The configuration is otherwise the same.

To configure NFAS with backup, follow this procedure:

1. **Set the view to the line board you want to configure.**

   ```
   Command> set view 0
   Command 0>
   ```

   The view is changed from the manager module to slot 0.

2. **Set the primary interface, the line number, and the NFAS group.**

   ```
   Command 0> set line0 nfas primary 0 5
   ```

   Line0 is set as the primary interface for NFAS group 5.

3. **Set the other interfaces as slave interfaces on the line board in slot 0.**

   ```
   Command 0> set line1 nfas slave 1 5
   Command 0> set line1 nfas slave 2 5
   Command 0> set line1 nfas slave 3 5
   ```

   Line1, Line2, and Line3 are set as slave interfaces in NFAS group 5.

4. **Save the configuration and reset the slot to make the changes take effect.**

   ```
   Command 0> save all
   Command 0> reset slot0
   ```

5. **Change the view to the slot containing the line board you want to set as the secondary interface, and set line 0 as the secondary interface.**

   ```
   Command 0> set view 1
   Command 1> set line0 nfas primary 4 5
   ```

   Line0 in slot 1 is set as the secondary interface for NFAS group 5. It is NFAS member number 4.

6. **Set the other interfaces as slave interfaces.**

   ```
   Command 1> set line1 nfas slave 5 6
   Command 1> set line2 nfas slave 6 5
   Command 1> set line3 nfas slave 7 5
   ```

7. **Save the configuration and reset the slot to make the changes take effect.**

   ```
   Command 1> save all
   Command 1> reset slot1
   ```

8. **Configure any additional interfaces for this NFAS group—up to a maximum of 20—as slave interfaces, save the configurations, and reset slots as appropriate.**

Note – When you configure a line board for NFAS, all interfaces on the board must use NFAS. You cannot configure some of the interfaces for standard PRI.

Displaying Information about NFAS Configurations

Use the **show** *Line0* command to display NFAS settings on an interface. For example, to display NFAS settings on a line board in slot 3, enter the following commands:

```
Command> set view 3
Command 3> show line3
```

Use the **show nfas** command to display a list of the members, called "neighbors," in an NFAS group. You must connect to a slot configured for NFAS to use this command. For example, to list the members of an NFAS group of which slot 2 is a member, enter the following commands:

```
Command 3> set view 2
Command 2> show nfas
```

Enter the following command to display the last 40 significant messages exchanged between a PortMaster and its neighbors:

```
Command> show nfas history
```

Troubleshooting NFAS

Use the **set debug nfas** command to diagnose problems or errors during testing. For example, to receive debug information from the line board in slot 3, enter the following commands:

```
Command 2> set console
Command 2> set view 3
Command 3> set debug nfas on
```

Refer to the *PortMaster Command Line Reference* for more information about NFAS commands.

Example NFAS Configurations

This section provides sample configurations to illustrate basic NFAS configuration for standard NFAS and NFAS without D channel backup.

NFAS with D Channel Backup

The convention, when configuring NFAS, is to set interface 0 as the primary D channel. In this example, Line0 in slot 0 is set as the primary D channel for group 5, Line1 in slot 1 is the backup D channel, and all other interfaces are set as slave interfaces.

```
Command> set view 0
Command 0> set line0 nfas primary 0 5
Command 0> set line1 nfas slave 1 5
Command 0> set line2 nfas slave 2 5
Command 0> set line3 nfas slave 3 5
Command 0> save all
Command 0> reset slot0
```

```
Command> set view 1
Command 1> set line0 nfas secondary 4 5
Command 1> set line1 nfas slave 5 5
Command 1> set line2 nfas slave 6 5
Command 1> set line3 nfas slave 7 5
Command 1> save all
Command 1> reset slot1
```

To take full advantage of NFAS, you can configure three more Quad T1 boards as slave interfaces for a maximum of 20 interfaces.

Standard NFAS

The convention, when configuring NFAS, is to set interface 0 as the primary D channel. In this example, Line0 in slot 0 is set as the primary D channel for group 5. All other interfaces are set as slave interfaces.

```
Command> set view 0
Command 0> set line0 nfas primary 0 5
Command 0> set line1 nfas slave 1 5
Command 0> set line2 nfas slave 2 5
Command 0> set line3 nfas slave 3 5
Command 0> save all
Command 0> reset slot0

Command> set view 1
Command 1> set line0 nfas slave 4 5
Command 1> set line1 nfas slave 5 5
Command 1> set line2 nfas slave 6 5
Command 1> set line3 nfas slave 7 5
Command 1> save all
Command 1> reset slot1
```

To take full advantage of NFAS, you can configure three more Quad T1 boards as slave interfaces for a maximum of 20 interfaces.

Configuring SS7

Signaling System 7 (SS7) is an out-of-band signaling system that provides fast call setup by means of high-speed connections and transactions. SS7 uses physical out-or-band signaling, which means that signaling is transmitted over a full-duplex, 64Kbps digital transmission channel on an entirely separate network from the voice and/or data information. Voice and/or data traffic flows over intermachine trunks (IMTs). SS7 makes possible features such as caller ID, call forwarding, and call waiting. When you configure SS7, you identify an SS7 gateway to the signaling network that does call management for user traffic on IMTs.

Because modem pools are managed on a slot-by-slot basis, each slot in the PortMaster 4 configured for IMTs is an SS7 client and sets up an independent session with the SS7 gateway. A PortMaster 4 supports only one SS7 gateway and can dedicate only 96 modems on its installed Quad T1 or Tri E1 boards to SS7.

You configure the PortMaster 4 for SS7 by setting the SS7 gateway address and TCP port, the TCP port on the PortMaster 4 and, occasionally, the switch type. You must configure each slot independently. The SS7 gateway address and target TCP port are provided by the SS7 gateway administrator. The TCP port of the PortMaster 4 is the base value that is used to generate the source TCP port number for the Quad T1 or Tri E1 board in each slot that has active IMTs. The actual port number is the base value plus the board's slot number.

Setting the Intermachine Trunk

To specify the SS7 gateway that does call management (signaling) for intermachine trunks (IMTs), use the following command:

Command> **set imt-parms** *Ipaddress Tport1 Tport1* [**1a**|**default**]

See the *PortMaster 4 Command Line Reference* for syntax information.

How you configure the IMT depends on whether you use the default 5ESS switch type or the optional 1A switch type.

Using a Default Switch Type

To configure the PortMaster 4 to communicate with an SS7 gateway using the default IMT switch type, you set only the IP address and port number of the gateway, and the PortMaster 4 base port number. You do not need to set the switch type—if it is not specified, the default is assumed.

For example, to set an SS7 gateway with IP address **192.168.10.10** and TCP port number **10000** to communicate with a line board in slot **0** on TCP port **7000**, enter the following commands:

```
Command> set view 0
Command 0> set imt-parms 192.168.10.10 10000 7000
Changed gateway IP address from 0.0.0.0 to 192.168.10.10
Changed gateway port from 0 to 10000
Changed local port from 7000 to 7000
```

The local port number (**7000**) is the TCP socket address the PortMaster 4 uses to communicate with the SS7 gateway. Because each Quad T1 or Tri E1 line board you configure for SS7 automatically adds its own slot number to the base port number, you use the same IMT base port number for each slot you configure. Always configure slot 0 first.

You must then save the configuration and reboot the PortMaster for the changes to take effect.

```
Command 0> save all
Command 0> reset slot0
```

Using the Optional 1A Switch Type for T1

To configure the PortMaster 4 to communicate with an SS7 gateway using the 1A IMT switch type, you must specify the switch type when you set the IP address and port number of the gateway, and the PortMaster 4 base port number. The **1A** setting allows the PortMaster 4 to interpret the loopback command from the SS7 gateway as a 1A continuity check request. Because 1A switches require continuity checks, you must group 1A IMTs together on a board-by-board basis.

When you use the 1A setting, you must also set robbed bit signaling (RBS) on the line and configure the appropriate line encoding and framing. With this setting, the PortMaster 4 treats the IMT as a robbed bit signaled line with twenty-four 56Kbps channels. It sets the bits to indicate an on-hook condition as specified by the E & M wink start protocol.

For example, to set an SS7 gateway with IP address **192.168.10.10** and TCP port number **10000** to communicate with a line board in slot **0** on TCP port **7000** using a 1A switch, enter the following commands:

```
Command > set view 0
Command 0> set imt 192.168.10.10 10000 7000 1a
Changed gateway IP address from 0.0.0.0 to 192.168.10.10
Changed gateway port from 0 to 10000
Changed local port from 7000 to 7000
Changed switch type from default to 1a
```

Save the configuration and reboot the PortMaster to make the changes take effect.

```
Command 0> save all
Command 0> reset slot0
```

Because you must group 1A IMTs together on a board-by-board basis for continuity check requests, you must then set robbed bit signaling on all active lines in the slot. For example, to configure Line0 on a Quad T1 board in slot 0 to use robbed bit signaling, enter the following commands:

```
Command 4> set view 0
Command 0> set line0 imt
Command 0> set line0 signaling rbs
```

Note – You set line signaling only for 1A type switches. When robbed bit signaling is not set, the IMT supports 24 64Kbps channels.

After you set line signaling for each line you want to use for SS7, you must reset the slot. For example, to reset slot 0, enter the following commands:

```
Command 0> set view 4
Command 4> reset slot 0
```

Viewing SS7 Configurations

To display SS7 configuration information on a slot-by-slot basis, use the following command:

```
Command 0> show imt
```

To display configuration information for a specified line, use the following command and substitute the desired line number:

```
Command 0> show Line0
```

Troubleshooting SS7 Configurations

This section provides information on checking activity between the slot on the PortMaster 4 and the SS7 gateway. To troubleshoot modems, see "Troubleshooting Digital Modems" on page 10-11.

Checking SS7 Gateway Initialization

Caution – Because debug commands can significantly slow throughput on line boards, Lucent recommends that you do not perform debugging operations on production lines.

To view SS7 debug information for a Quad T1 or Tri E1 board, set the view to the slot and enter the following commands:

```
Command 1> set console
Command 1> set debug imt on
```

For example, the following sample output shows a partial stream of debug messages that indicate slot 1 activity until the SS7 gateway is completely initialized. The output shows only "heartbeat" information when the gateway is in a ready state.

```
slot1: imt_hangup: line(1), slot(23)

slot1: imt_hangup: line(1) slot(23) flags(0) state(1)

slot1: imt_hangup: state 1, done

slot1: imt_hangup: done

slot1: imt_receive_data: got heartbeat

slot1: imt_proc: IMT_HEARTBEAT

slot1: imt_process_heartbeat: sent heartbeat

slot1: imt_send_data: sent packet

slot1:    00

slot1:    00

slot1: imt_receive_data: got heartbeat

slot1: imt_proc: IMT_HEARTBEAT

slot1: imt_process_heartbeat: sent heartbeat
```

```
slot1: imt_send_data: sent packet

slot1:   00

slot1:   00
```

The following example shows output from slot 2 when the SS7 gateway is not working. This kind of output indicates that the gateway is not "listening" or is not configured with the correct IP address and socket number.

```
slot2: imt_gw_conn_established: not!

slot2: imt_gw_conn_established: not!

slot2: imt_gw_conn_established: not!

slot2: imt_gw_conn_established: not!

slot2: imt_gw_conn_established: not!

slot2: imt_gw_conn_established: not!
```

Displaying Session Information

When the SS7 gateway is active and a session is established, use the **show netconns** command to display session information. The following output to the **show netconns** command shows that slot 1 and slot 2 (handles S2 and S1) have established sessions with the gateway:

```
    Command 0> show netconns

Hnd Recv-Q Send-Q  Local Address        Foreign Address      (state)

 21  0      0      192.168.10.10.23      0.0.0.0.0            LISTEN

 20  0      0      192.168.10.10.23      172.16.10.10.40982    ESTABLISHED

 19  0      0      192.168.10.10.67      0.0.0.0.0            UDP

 18  0      2      192.168.10.10.23      172.20.23.25.21082   ESTABLISHED

 17  0      0      192.168.10.10.10100   0.0.0.0.0            LISTEN

 16  0      0      192.168.10.10.10099   0.0.0.0.0            LISTEN

 15  0      0      192.168.10.10.1643    0.0.0.0.0            LISTEN

 14  0      0      192.168.10.10.1701    0.0.0.0.0            UDP

  5  0      0      192.168.10.10.1044    192.168.10.12.1646   UDP

  3  0      0      192.168.10.10.520     0.0.0.0.0            UDP
```

```
    S2   0       0       192.168.10.10.7002      192.168.10.12.10000     ESTABLISHED

    S1   0       0       198.36.134.10.7001      192.168.10.12.10000     ESTABLISHED
```

SS7 Configuration Examples

This section provides two sample SS7 configurations: one for the default switch type, and one using the 1A switch type. Each example shows how to configure slot 0 and slot 1. The SS7 gateway is at address **192.168.10.10** and is listening for SS7 signaling on port **10000**. Port **7000** on the PortMaster 4 is the IMT base port for both slot 0 and slot 1.

Default Configuration

Configure slot 0.

```
set view 0
set imt 192.168.10.10 10000 7000
save all
reset slot0
```

Configure slot 1.

```
set view 1
set imt 192.168.10.10 10000 7000
save all
reset slot1
```

1A Switch Type Configuration

Configure slot 0.

```
set view 0
set imt 192.168.10.10 10000 7000 1a
```

Set line signaling on all lines you want to use for SS7 in slot 0.

```
set line0 imt
set line0 signaling rbs
. . .
save all
reset slot0
```

Configure slot 1.

```
set view 1
set imt 192.168.10.10 10000 7000 1a
```

Set line signaling on all lines you want to use for SS7 in slot 1.

```
set line0 imt
set line0 signaling rbs
. . .

save all
reset slot1
```

Configuring a T3 Mux Board   11

This chapter describes how to integrate a T3 Mux board with Quad T1 boards on the PortMaster 4. It includes the following sections:

- "Overview of T3 Mux Boards" on page 11-1
- "Setting the View" on page 11-1
- "Mapping T1 Lines to T3 Channels" on page 11-1
- "Setting the Clock Source" on page 11-2
- "Performing Diagnostics" on page 11-3

Overview of T3 Mux Boards

The T3 Mux board demultiplexes DS-3 signals into 28 individual DS-1 signals (also called channels), and terminates those channels on one or more Quad T1 boards. T3 Mux board configuration is primarily a matter of mapping the individual T1 lines of the installed Quad T1 boards to the 28 DS-1 channels of the T3 line. To fully demultiplex a T3 line, you must have seven Quad T1 boards installed in the PortMaster 4.

The T3 Mux board supports the M13 framing format and converts bipolar 3-zero substitution (B3ZS) line encoding to nonreturn to zero (NRZ) DS-3 signaling. A T3 Mux board can provide internal clocking or can receive clocking externally at a rate of 44.736Mbps. The DS-1 clocking rate is 1.544Mbps.

You can install a T3 Mux board in any slot except slot 4, which is reserved for the manager module. T3 Mux boards do not become active until you configure an IP address for the PortMaster 4.

Setting the View

Before configuring a T3 Mux board and any Quad T1 board that it demultiplexes, you must set the view to the slot containing the appropriate board. Enter the **show boards** command to determine the identification number for the board. The board identification number is the same as the number of the slot in which the T1 line is installed.

    Command> **show boards**

Use the following command to set the view to a slot with an installed T1 board:

    Command> **set view** *Slotnumber*

Mapping T1 Lines to T3 Channels

You can configure the lines on a Quad T1 board to use their physical ports (Line0 through Line3), or map them to DS-1 channels on the T3 line. To map the termination of a T1 line, you must first set the view to the slot with the installed Quad T1 board.

```
Command> set view Slotnumber
```

You can now use the **set** *Line0* source command to map the T1 line source. Use the **local** keyword to terminate the line at its own RJ-45 port—this is the default. Or you can specify the slot number and channel of the T3 Mux board to multiplex the T1 line to a T3 channel between 1 and 28. For example, to map Line0 on a Quad T1 board in slot 2 to DS-1 channel 25 of the T3 line of a T3 Mux board installed in slot 8, enter the following commands:

```
Command> set view 2
Command 2> set line0 source 8:25
```

Note – Before configuring the T3 Mux board, you must activate it by assigning an IP address to Ether0 or Ether1.

Setting the Clock Source

You can set clocking on a T3 Mux board in several ways. This section discusses the options for setting the clock source.

Setting the Clock Source for Each Synchronous Serial Line

To set a clock source for the T1 lines, you must first set the view to a slot with the installed Quad T1 board.

```
Command> set view Slotnumber
```

You can now use the **set** *Line0* **clock** command to set the clock source. You use the **backplane** keyword of this command to derive clocking from a T3 Mux board. If you want the internal T1 clock to supply clocking, use the **internal** keyword. Use the **external** keyword if you need an external clock source.

For example, to specify the backplane as the clock source for Line0 on a Quad T1 board installed in slot 1, enter the following command:

```
Command 1> set line0 clock backplane
```

Warning – Do not set a T1 line to **internal** if another active line on the same Quad T1 board is set to **external** or **backplane**.

If you want to derive clocking from the backplane, you must enable the T3 Mux board to provide clocking to the backplane.

Enabling Clocking on the Backplane

When you configure a T3 Mux board to provide clocking to the PortMaster 4 backplane, other boards can derive their clocking from the backplane. To set the T3 Mux board to provide clocking to the backplane, you must first set the view to the slot with the installed T3 Mux board. For a T3 Mux board installed in slot 1, for example, you enter the following command:

```
Command> set view 1
```

You can now use the following command to set the T3 Mux board to provide clocking to the backplane:

```
Command 1> set mux backplane-clock enable|disable
```

Setting T3 Mux Clocking

You can set a T3 Mux board to either provide clocking to the T3 line or obtain clocking from the T3 line.

You use the **external** keyword to specify that the T3 Mux board obtain clocking from the T3 line. When you use the **internal** keyword, the T3 Mux board provides clocking for the T3 line. To set T3 Mux clocking, you must first set the view to the slot with the installed T3 Mux board. For a T3 Mux board installed in slot 1, for example, you enter the following command:

```
Command 1> set view 1
```

You can now use the following command to set T3 Mux clocking:

```
Command 1> set mux line-clock external|internal
```

Performing Diagnostics

This section describes how to enable a PortMaster 4 to conduct loopback tests on a single DS-1 channel and on a T3 line.

Looping an Individual DS-1 Channel

To use the following command to perform diagnostic loopback tests on a single DS-1 channel within a T3 line, you must first set the view to the slot with the installed T3 Mux board. For a T3 Mux board installed in slot 1, for example, enter the following command:

```
Command> set view 1
```

You then enter the following command to enable the PortMaster to perform loopback tests on a single DS-1 channel:

```
Command 1> set mux channel-loop Channel auto|on|off
```

Looping the T3 Line

To use the following command to diagnose a T3 line, you must first set the view to the slot with the installed T3 Mux board. For a T3 Mux board installed in slot 1, for example, enter the following command:

```
Command> set view 1
```

You can then use the following command as a diagnostic to loop the T3 line:

```
Command 1> set mux line-loop auto|on|off
```

Using Frame Relay 12

Frame Relay is a method of encapsulating network information that allows for fast delivery and high line utilization. PortMaster routers support Frame Relay over synchronous ports.

This chapter uses an example to demonstrate how to configure the PortMaster 4 to connect to a synchronous line using Frame Relay. This chapter also explains how to configure Frame Relay subinterfaces.

The following topics are discussed:

- "Overview of Frame Relay" on page 12-1
- "Frame Relay Configuration on the PortMaster" on page 12-3
- "Configuring a T1 or E1 Line for Frame Relay" on page 12-4
- "Troubleshooting a Frame Relay Configuration" on page 12-7
- "Frame Relay Subinterfaces" on page 12-8

See the *PortMaster 4 Command Line Reference* for more detailed command descriptions and instructions.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Overview of Frame Relay

Frame Relay is a switched digital service that supports multiple virtual circuits, simultaneously connected to a site by a single physical circuit. Each site requires only one physical circuit into the Frame Relay network—usually referred to as a cloud—but can have several virtual circuits to reach other sites attached to the cloud.

You configure synchronous ports on the PortMaster 4 to support Frame Relay connections. As opposed to a dedicated or leased line, a Frame Relay connection can be thought of as a virtual switch.

PVCs and DLCIs

Permanent virtual circuits (PVCs) are used to form a connection between any two devices attached to a Frame Relay cloud. Each PVC is given a unique number on each physical circuit along the path between the two devices. This unique number is called a data link connection identifier (DLCI). The DLCI is automatically changed to the PVC number of the next physical circuit as it passes through each switch along the path. A DLCI is different from a network address in that it identifies a circuit in both directions, not a particular endpoint. A frame contains only one DLCI, not a source and destination.

In general, the only DLCI numbers you see are those numbers assigned to the physical circuits on the perimeter of the Frame Relay cloud.

Line Speed

The physical circuit between point A and the network must be ordered with a certain line speed. This speed is the maximum physical bandwidth for your connection to the Frame Relay network. Expansion beyond this limit is not possible without a hardware change and a new circuit installation.

Port Speed

The connection into the telecommunications provider's Frame Relay network must be ordered at a particular port speed, which is the maximum bandwidth rate that the telecommunications provider accepts from your connection. This number must be less than or equal to the line speed. This speed is the maximum rate at which you can transmit data to any of your PVCs under any circumstances. The port speed differs from line speed only in that it can be upgraded through software without a circuit installation or hardware change.

CIR and Burst Speed

Each PVC has a property known as committed information rate (CIR), which represents the guaranteed minimum bandwidth available to the particular PVC under all conditions. In some implementations, an additional property can be assigned to a PVC, known as burst speed or maximum burst. This speed represents the highest rate at which data is allowed to flow over a given PVC, regardless of bandwidth availability.

Discarding Frames

The PortMaster 4 transmits as much data on the serial port as it can for any PVC that has traffic, regardless of CIR. The Frame Relay switch passes on as much of the data as possible to the next link. However, once a particular PVC has transmitted its CIR-worth of bits each second, the switch marks any additional frames as "discard eligible." If the switch receives more frames than it can pass along, the frames are automatically discarded in the following order:

- Frames that would be marked discard eligible even if they are forwarded

- Frames received that were marked as discard eligible

If the switch must discard other frames, the behavior is undefined. In this case, the Frame Relay network is improperly configured because the CIR total exceeds the line speed or port speed.

Ordering Frame Relay Service

In general, when ordering Frame Relay service for a private network, order large-bandwidth physical circuits (for example, T1) with a port speed appropriate to your application, and a CIR that is high enough to provide minimally acceptable

performance for your application. In most cases, ordering according to these criteria provides service that is close to your port speed. The CIR is a guaranteed minimum throughput, not a maximum limit. Port speed is the maximum limit.

LMI Types

The following Frame Relay terms relate to network management. The Frame Relay specification supports automatic network status updates, which are exchanged between adjacent devices in the Frame Relay network. These status updates are known as the Local Management Interface (LMI). Two forms of LMI are available in the PortMaster: Cisco LMI, which is commonly referred to as LMI, and ANSI T1.617 Annex D LMI, which is commonly referred to as Annex-D.

Generally, your telecommunications provider offers three LMI options for your physical circuit: LMI, Annex-D, or none. Because LMI exists only between your router and the switch to which your physical circuit connects, it does not need to match what the remote ends of your PVCs are using. However, your circuit LMI must match the configuration on your PortMaster. Generally, Annex-D is recommended because it is a more feature-rich and robust version of LMI.

Frame Relay Configuration on the PortMaster

You configure Frame Relay by selecting the Frame Relay protocol, setting the IP address of the port, and specifying the DLCIs during the synchronous port configuration.

Alternatively, the PortMaster can discover DLCIs dynamically with LMI or Annex-D and learn the IP addresses of the other routers through Inverse ARP if the other routers on your Frame Relay cloud support Inverse ARP as specified in RFC 1490. In this configuration, the PortMaster sends an LMI status request every 10 configurable seconds by default. Every sixth request is a full status request, and the others are keepalives. In this configuration, the port state is CONNECTING until it receives three replies from the switch; then the port state becomes ESTABLISHED. After six unanswered requests, the PortMaster resets the port.

Note – All synchronous ports require an external clock signal–either from the device to which the PortMaster is connected or from the telephone company–to regulate the port speed.

Enabling LMI

You can specify whether the PortMaster accepts Local Management Interface (LMI) frames from the attached Frame Relay switch. If LMI is enabled on the switch, you must enable LMI on the PortMaster. The default keepalive value is 10 seconds. However, if your telephone company chooses another keepalive value, change this value as they instruct you. Enabling LMI causes the DLCI list to be completed automatically. If the attached switch uses an interval keepalive timer different from the Frame Relay default, be sure the keepalive timer on the PortMaster matches that of the attached switch.

To enable LMI, use the following command:

```
Command> set W1 lmi Seconds
```

Note – Contact your Frame Relay carrier to determine which keepalive timer they use, LMI or Annex-D.

To enable LMI, use the following command:

    Command> **set** *W1* **lmi** *Seconds*

Enabling Annex-D

The PortMaster also accepts the Annex-D polling interval. The Annex-D default value is 10 seconds. However, if your telephone company chooses another keepalive value, change this value as they instruct you. Enabling LMI causes the DLCI list to be completed automatically. Setting the keepalive value to 0 (zero) seconds, or enabling LMI, disables Annex-D.

Note – Contact your Frame Relay carrier to determine which keepalive they are using, LMI or Annex-D.

To enable Annex-D, use the following command:

    Command> **set** *W1* **annex-d** *Seconds*

Listing DLCIs for Frame Relay Access

If LMI or Annex-D is not used, you must enter the DLCI list manually. The DLCI list is a list of DLCIs that are accessible through the Frame Relay network by this interface. The PortMaster uses Inverse ARP to learn the IP addresses of routers attached to the PVCs represented by the specified DLCIs, if those routers support Inverse ARP. Alternatively, you can specify single IP addresses or IPX network numbers by appending a colon (:) and address or number after the DLCI. See the *PortMaster 4 Command Line Reference* for more information.

The DLCI list can be provided by your Frame Relay carrier. For dynamically learned lists, 32 PVCs are allowed. Only 16 PVCs can be specified if the DLCI and IP address are entered. If you specify only DLCIs, you can list 24. When the PVC and IP address are specified, the PortMaster statically configures these entries into its ARP table.

To enter the DLCI list manually, use the following command:

    Command> **add dlci|ipxdlci** *W1 Dlci* **:**[*Ipaddress|Ipxnode*]

For information on Frame Relay subinterfaces see "Frame Relay Subinterfaces" on page 12-8.

Configuring a T1 or E1 Line for Frame Relay

You configure Frame Relay on the PortMaster 4 by configuring a T1 line on the Quad T1 board or an E1 line on the Tri E1 board.

This section describes how to configure one end of a Frame Relay connection. Because configuration for both ends of the connection is the same, you can use the procedure in this section as a guide for configuring both ends of your Frame Relay connection.

Before you configure a line for Frame Relay, you must configure global settings on the PortMaster 4, set the gateway routers, and configure the Ethernet interface. Follow the procedures in this section to configure a T1 or E1 line for Frame Relay service.

Configuring Global and Ethernet Settings

You configure global and Ethernet settings on the PortMaster 4 from the manager module (slot 4), which is the default view. See Chapter 2, "Configuring Global Settings," and Chapter 4, "Configuring an Ethernet Interface," for more information.

Note – Lucent recommends that you configure Ether1 if you configure only one Ethernet interface. If you configure both, you must connect them to separate Ethernet segments. Ether0 operates at 10Mbps and is physically on the manager board. Use Ether1 for netboots, SNMP, RADIUS, and **syslog**.

Follow this proceduG8e to configure the global and Ethernet settings on the PortMaster 4.

1. **Set the IP address of the gateway router.**

   Command> **set gateway** *Ipaddress*

2. **Save the configuration to nonvolatile RAM.**

   Command> **save all**

3. **Set the IP address of the Ethernet interface.**

   Command> **set ether1 address** *Ipaddress*

4. **Set the netmask of the Ethernet interface.**

   Command> **set ether1 netmask** *Netmask*

5. **Set broadcast to high.**

   Command> **set ether1 broadcast high**

6. **Enable RIP routing on the interface.**

   Command> **set ether1 rip on**

7. **Save the configuration to nonvolatile RAM.**

   Command> **save all**

8. **Reset the slot**

   Because the Ethernet board on the manager module is in logical slot 10, you activate the Ether1 configuration by resetting slot 10.

   Command> **reset slot10**

Configuring the Synchronous WAN Port

Follow this procedure to configure the synchronous WAN port on the PortMaster 4:

1. **Enter the** show boards **command to determine the identification number of the line board you want to configure.**

    The board identification number is the same as the number of the slot in which the T1 or E1 module is installed.

    Command> **show boards**

2. **Set the view to a slot with an installed T1 or E1 line board.**

    Setting the view for a specific board gives you administrative access to that board.

    Command> **set view** *Slotnumber*

    Command *Slotnumber*>

    The slot number appears in the prompt to identify the board you are configuring.

3. **Set the network type.**

    Command *Slotnumber*> **set** *W1* **network hardwired**

4. **Set the protocol.**

    Command *Slotnumber*> **set** *W1* **protocol frame**

5. **Set the IP address of the port.**

    Command *Slotnumber*> **set** *W1* **address** *Ipaddress*

6. **Set the netmask.**

    Command *Slotnumber*> **set** *W1* **netmask** *Netmask*

7. **Monitor the carrier detect signal.**

    Command *Slotnumber*> **set** *W1* **cd on**

8. **Enable RIP on the interface.**

    Command *Slotnumber*> **set** *W1* **rip broadcast**

9. **Set the Annex-D polling interval.**

    You can use LMI instead of Annex-D.

    Command *Slotnumber*> **set** *W1* **annex-d** *Seconds*

10. **Add a data link connection identifier (DLCI).**

    You do not need to set a DLCI list if the remote router supports Inverse ARP.

    Command *Slotnumber*> **add dlci** *W1 Dlci* **:***Ipaddress*

11. **Reset the slot.**

    Command *Slotnumber*> **reset** *W1*

12. **Save the configuration.**

    Command *Slotnumber*> **save all**

If LMI or Annex-D is set, the PortMaster receives DLCI information in the full status update messages from the Frame Relay switch. The PortMaster then attempts to discover IP addresses of other routers using Inverse ARP. You can set DLCI lists statically as well. The **show arp frm1** command lists both the static and dynamic DLCI lists for the W1 port.

If Annex-D is available from your carrier for a new connection, it is preferable to LMI.

To connect to Cisco routers using Frame Relay, the Cisco router must be set to use **encapsulation frame-relay ietf** for the serial interface; otherwise, the Cisco **frame-relay map** command for your DLCI must have the **ietf** keyword appended.

For more information about synchronous ports, refer to Chapter 6, "Configuring a Synchronous WAN Port."

Troubleshooting a Frame Relay Configuration

Most synchronous configurations establish links with very little trouble if you have configured the PortMaster using information from your carrier. If you are having problems, use the information in this section to debug your configuration.

If you are having trouble with a Frame Relay connection, do the following:

- Wait a few moments. The process of establishing a Frame Relay link, learning the DLCI list, and learning the IP address through Inverse ARP can sometimes take a few moments.

- Verify that the error counters are 0 except for abort errors. If your counters are nonzero, the problem is external to the PortMaster.

- Verify that you are using the correct cables and that they are attached securely to the correct port. Not all WAN ports are capable of the same speeds.

- Enter the following two commands to view the LMI or Annex-D keepalives:

  Command> **set console** *C0*
  Command> **set debug 0x51**

  After you verify that the proper keepalives are being received, enter the following commands to turn off the debug utility:

  Command> **set debug off**
  Command> **reset console**

- If you have a Cisco router on the other end of your connection, verify that it is set for **encapsulation frame-relay ietf** for the serial interface; otherwise, the Cisco **frame-relay map** command for your DLCI must have the **ietf** keyword appended.

Frame Relay Subinterfaces

PortMaster routers support a feature called DLCI bundling to allow the splitting of one synchronous port with multiple DLCIs into a maximum of 32 Frame Relay subinterfaces. In this configuration, the DLCIs are divided between the subinterfaces through the use of the location table and the DLCI table. Each subinterface must have its own subnet or assigned network. The PortMaster has a limit of 512 total active interfaces, which can be further limited by available memory.

The port you are configuring must be set for network hardwired use and Frame Relay, and must be in the same dial group as the location.

Configuring Subinterfaces

The following sections describe how to configure a Frame Relay subinterface.

Adding a Location

To configure a Frame Relay subinterface, you add a location for each interface, configure it with the frame protocol, and associate it with a dial group. Then associate a synchronous port with the same dial group. For example, to create a location called **sub1**, enter the following commands:

```
Command> add location sub1
Command> set location sub1 protocol frame
Command> set location sub1 group 1
Command> set W1 group 1
```

The rest of the location table entries are set as described in Chapter 7, "Configuring Dial-Out Connections," including setting an IP address, routing, and filtering for each interface.

Creating a DLCI Entry

The next step in configuring the subinterfaces is to create an entry in the DLCI table. Entries can be followed with an optional IP address or hostname. The keyword **ipxdlci** is available for IPX networks.

Note – The PortMaster 4 supports the IPX protocol if its running ComOS 4.1 or later. ComOS 4.0 does not support the IPX protocol.

To create a DLCI table entry for the subinterface **sub1**, enter the following commands:

```
Command> add ipdlci sub1 16
Command> add ipdlci sub1 19 192.168.2.19
Command> add ipdlci sub1 20 192.168.2.20
Command> add ipxdlci sub1 21 0e0a001e
```

To remove an entry, enter the following commands:

```
Command> delete dlci sub1
Command> delete ipxdlci sub1 21
```

Displaying DLCI Entries

DLCI entries that are added or deleted are linked to the location table. Use the **show location** *Locname* command to display the DLCI entries.

Troubleshooting Subinterfaces

Packets received on a subinterface can be identified as belonging to that subinterface only if the DLCI is properly entered in the DLCI table for that location. If you are having problems, do the following:

- Wait a few moments. Subinterfaces come up after the primary interface. This process can take a few moments.

- Check the list of DLCIs tied to each location using the **show location** *Locname* command.

- Verify the DLCI list on a location using the **show arp** *Interface* command, replacing *Interface* with the name of the interface. A list of interfaces can be shown with the **ifconfig** command.

- Always reset the port after changing the DLCI list.

- Verify that all DLCIs are accounted for by checking the DLCI list for your primary interface. If you enter the wrong DLCI for the subinterface, the DLCI for the subinterface is applied to the primary interface if LMI or Annex-D is in use.

- Enter the following two commands to view the LMI or Annex-D keepalives:

```
Command> set console C0
Command> set debug 0x51
```

After you verify that the proper keepalives are being received, enter the following commands to turn off the debug utility:

```
Command> set debug off
Command> reset console
```

- If you have a Cisco router on the other end of your connection, verify that it is set for **encapsulation frame-relay ietf** for the serial interface; otherwise, the Cisco **frame-relay map** command for your DLCI must have the **ietf** keyword appended.

Using Synchronous Leased Lines 13

This chapter describes how to set up a synchronous leased line between a PortMaster 4 and another PortMaster product. The chapter provides guidelines for configuring both ends of the connection and includes the following topics:

- "Overview of Leased Line Connections" on page 13-1

- "Configuring a Leased Line Connection" on page 13-2

- "Troubleshooting a Leased Line Connection" on page 13-4

See the *PortMaster 4 Command Line Reference* for more information about commands used in the chapter.

You can also configure the PortMaster 4 using the PMVision application for Microsoft Windows, UNIX, and other platforms supporting the Java Virtual Machine (JVM). PMVision replaces the PMconsole interface to ComOS.

Overview of Leased Line Connections

Leased line connections use leased or dedicated lines to establish a permanent connection between two routers. Once the connection is established, it remains available on a continuous basis whether or not network traffic exists between the two locations. Leased line connections have the required channel service unit/digital service unit (CSU/DSU) built in. The CSU/DSU takes digital data in the format used by the router and translates it into the digital format used by the leased line. Leased line connections also require a carrier that provides an external clock signal.

PortMaster routers support leased line connections using synchronous ports and the PPP protocol. In this configuration, one PortMaster is usually connected to another PortMaster or other router over a leased line where each router uses its own Ethernet address for the serial link—known as IP unnumbered—and the address of the other end is discovered dynamically. In this way, a dedicated high-speed connection is established between two routers located at separate sites.

If you are connecting two networks together for the first time, make sure the networks are not overlapping subnets. For more information on network numbers and subnetting, see Appendix A, "Networking Concepts."

In the leased line configuration described in this chapter, the Ethernet address of the PortMaster routers is used as the address for the serial link in a point-to-point unnumbered serial connection. Because the PortMaster relies on an external clock signal, you do not need to set the speed on the synchronous port. The port speed is whatever the carrier sends. If you choose to set a speed, it is used for administrative notation only and does not affect the operation of the port.

PortMaster synchronous ports support leased line connections from 9600bps to T1 (1.544Mbps) or E1 (2.048Mbps) speeds. Synchronous ports used for leased line connections are configured for PPP operation and can have input and output filters for network security.

Note – The PortMaster 4 also supports numbered IP interfaces on leased lines, but Lucent does not recommended this method because it wastes IP address space.

Configuring a Leased Line Connection

Before you configure a synchronous serial port for a network hardwired leased line, you must configure global settings on the PortMaster 4, and you must configure the Ethernet interface. Follow the procedures in this section to configure a T1 or E1 line for leased line service.

Configuring Global Settings

Follow this procedure to configure the global settings on the PortMaster 4.

1. **Set the IP address of the gateway router.**

   Command> **set gateway** *Ipaddress*

2. **Set the name of the PortMaster on the other end of the connection.**

   Command> **set sysname** [*String*]

3. **Save the configuration.**

   Command> **save all**

Configuring Ethernet Interface Settings

Note – Lucent recommends that you configure Ether1 if you configure only one Ethernet interface. If you configure both, you must connect them to separate Ethernet segments. Ether0 operates at 10Mbps and is physically on the manager board. Use Ether1 for netboots, SNMP, RADIUS, and **syslog**.

Follow this procedure to configure the Ethernet interface on the PortMaster 4:

1. **Set the IP address of the Ethernet interface.**

   Command> **set ether1 address** *Ipaddress*

2. **Set the netmask of the Ethernet interface.**

   Command> **set ether1 netmask** *Netmask*

3. **Set the broadcast address.**

   Command> **set ether1 broadcast high**

4. **Enable RIP routing on the interface.**

   Command> **set ether1 rip on**

   See the *PortMaster 4 Command Line Reference* for information about RIP-2.

5. **Save the configuration.**

   Command> **save all**

6. **Reset the slot.**

   Command> **reset slot10**

Configuring the Synchronous WAN Port

Follow this procedure to configure the synchronous WAN port on the PortMaster 4 for leased line service:

1. **Enter the** show boards **command to determine the identification number of the line board you want to configure.**

   The board identification number is the same as the number of the slot in which the T1 or E1 module is installed.

   Command> **show boards**

2. **Set the view to a slot with an installed T1 or E1 line board.**

   Setting the view for a specific board gives you administrative access to that board.

   Command> **set view** *Slotnumber*

   Command *Slotnumber*>

   The slot number appears in the prompt to identify the board you are configuring.

3. **Set a line on the line board for T1 or E1.**

   Command *Slotnumber*> **set** *Line0* **t1|e1**

4. **Save the configuration.**

   Command *Slotnumber*> **save all**

5. **Reset the slot.**

   Command *Slotnumber*> **reset** *Slotnumber*

6. **Set the interface for network hardwired.**

   Command *Slotnumber*> **set** *W1* **network hardwired**

7. **Set the protocol.**

   Command *Slotnumber*> **set** *W1* **protocol ppp**

8. **Set the destination IP address.**

   If you are not sure of the IP address on the other end of the connection, you can set the IP destination to 255.255.255.255 and the PortMaster attempts to learn the address.

   Command *Slotnumber*> **set** *W1* **destination** *Ipaddress*

9. **Set the netmask.**

   Command *Slotnumber*> **set** *W1* **netmask** *Netmask*

10. **Turn off modem control.**

    Command *Slotnumber*> **set** *W1* **cd off**

11. **Turn on RIP routing.**

    Command *Slotnumber*> **set** *W1* **rip on**

    See the *PortMaster 4 Command Line Reference* for information about RIP-2.

12. **Set the maximum transmission unit size.**

    Command *Slotnumber*> **set** *W1* **mtu** *MTU*

13. **Reset the slot and save the configuration.**

    Command *Slotnumber*> **reset** *Slotnumber*
    Command *Slotnumber*> **save all**

For more information about synchronous ports, refer to Chapter 6, "Configuring a Synchronous WAN Port."

Troubleshooting a Leased Line Connection

Use the information in this section to debug your configuration.

If you have trouble with a leased line connection, verify the following:

* Enter the following commands to view the PPP negotiation on port W1, if this is the port you are using:

      Command> **set console** *W1*
      Command> **set debug 0x51**
      Command> **reset** *W1*

  For information about the interpreting the results of the debug command, refer to the *PortMaster Troubleshooting Guide*.

  After you verify that the PPP negotiation is correct, enter the following commands to turn off the debug utility:

      Command> **set debug off**
      Command> **reset console**

* Verify that the error counters are 0 (zero) except for abort errors. If your counters are nonzero, the problem is external to the PortMaster.

Note – CRC errors will occur if the leased line cable is ever unplugged from the PortMaster.

* Verify that you are using the correct cable and that it is attached securely to the correct port. Not all WAN ports are capable of the same speeds.

- If you have a Cisco router on the other end of your connection, make sure that it is running Cisco's software release 9.14(5) or later and is using PPP encapsulation, not High-Level Data Link Control (HDLC).

- If the framing errors are greater than 0, verify that the router on the other end of the connection is running the PPP protocol.

- If you still have problems, enter the following commands:

  ```
  Command> set debug 0x51
  Command> set console C0
  ```

  Then set the line for local loopback:

  ```
  Command> set Line0 loopback on
  ```

  You should see the following message:

  ```
  LCP_APPARENT_LOOP
  ```

  For more information about the interpreting the results of the debug command, refer to the *PortMaster Troubleshooting Guide*.

- If the local loopback shows network connectivity in the local router, take the line out of loopback and set line loopback on the remote end of the connection. If the remote loopback test does not show network connectivity in the remote router, the problem is either in the configuration of one of the CSU/DSUs or in the line itself.

- When you finish, enter the following commands to turn off the debug utility:

  ```
  Command> set debug off
  Command> reset console
  ```

- Contact your carrier to review your configuration and the status of their line.

Networking Concepts A

This chapter describes general network concepts that you must understand before you configure your PortMaster.

This chapter discusses the following topics:

- "Network Addressing" on page A-1

- "Using Naming Services and the Host Table" on page A-7

- "Managing Network Security" on page A-7

See the *PortMaster Routing Guide* for information on routing and how Lucent's ComOS implements routing protocols. See the glossary for unfamiliar terms.

Network Addressing

PortMaster products support packet routing using the IP protocol. The Internet Protocol (IP) is a packet-based protocol used to exchange data over computer networks. IP provides addressing and control information that allows data packets to be routed across networks.

Novell Internetwork Packet Exchange (IPX) is another protocol used to exchange data over PC-based networks. IPX uses Novell's proprietary Service Advertising Protocol (SAP) to advertise special services such as print and file servers. The PortMaster 4 supports the IPX protocol if it is running ComOS 4.1 or later. IPX is not supported on ComOS 4.0.

IP Addressing

IP address descriptions are found in RFC 1166, *Internet Numbers.* Refer to "Additional References" in the preface for more information. The Internet Network Information Center (InterNIC) maintains and distributes the RFC documents. The InterNIC also assigns IP addresses and network numbers to Internet service providers (ISPs), who in turn provide to their customers a range of addresses appropriate to the number of host devices on their network.

The sections that follow describe the various types of IP addresses, how addresses are given, and routing issues related to IP.

IP Address Notation

IP addresses are written in dotted decimal notation consisting of four numbers separated by dots (periods). Each number, written in decimal, represents an 8-bit octet (sometimes informally referred to as a byte) giving each number a range of 0 through 255, inclusive. When strung together, the four octets form the 32-bit IP address. Table A-1 shows 32-bit values expressed as IP addresses.

Table A-1   IP Address Notation

32-Bit Value	Dotted Decimal Notation
01100100.01100100.01100100.00001010	100.100.100.10
11000011.00100000.00000100.11001000	195.32.4.200

The largest possible value of a field in dotted decimal notation is 255, which represents an octet where all the bits are 1s.

IP Address Classes

IP addresses are generally divided into different classes of addresses based on the number of hosts and subnetworks required to support the hosts. As described in RFC 1166, IP addresses are 32-bit quantities divided into five classes. Each class has a different number of bits allocated to the network and host portions of the address. For this discussion, consider a network to be a collection of computers (hosts) that have the same network field values in their IP addresses.

The concept of classes is being made obsolete by classless interdomain routing (CIDR). Instead of dividing networks by class, CIDR groups them into address ranges. A network range consists of an IP address prefix and a netmask length. The address prefix specifies the high-order bits of the IP network address. The netmask length specifies the number of high-order bits in the prefix that an IP address must match to fall within the range indicated by the prefix.

For example, 192.168.42.*x* describes a Class C network with addresses ranging from 192.168.42.0 through 192.168.42.255. CIDR uses 192.168.42.0/24 to describe the same range of addresses.

RIP-1 is an example of a protocol that uses address classes. RIP-2, OSPF, and BGP-4 are examples of protocols that do not use address classes.

Class A Addresses

The class A IP address format allocates the highest 8 bits to the network field and sets the highest-priority bit to 0 (zero). The remaining 24 bits form the host field. Only 126 class A networks can exist (0 is reserved, and 127 is used for loopback networks), but each class A network can have almost 17 million hosts. No new class A networks can be assigned at this time.

For example:

10.100.232.1

Network address — Host address
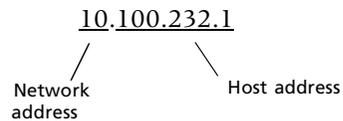
Class B Addresses

The class B IP address format allocates the highest 16 bits to the network field and sets the two highest-order bits to 1 and 0, providing a range from 128 through 191, inclusive. The remaining 16 bits form the host field. More than 16,000 class B networks can exist, and each class B network can have up to 65,534 hosts. For example:

172.16.232.121

Network address — Host address
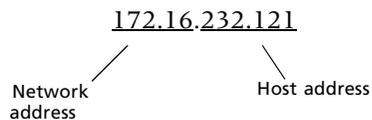
Class C Addresses

The class C IP address format allocates the highest 24 bits to the network field and sets the three highest-order bits to 1, 1, and 0, providing a range from 192 through 223, inclusive. The remaining 8 bits form the host field. More than two million class C networks can exist, and each class C network can have up to 254 hosts. For example:

192.168.20.220

Network address — Host address
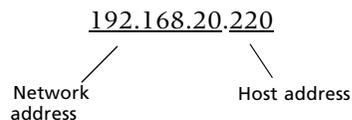
Class D Addresses

The class D IP address format was designed for multicast groups, as discussed in RFC 1112. In class D addresses, the 4 highest-order bits are set to 1, 1, 1, and 0, providing a range from 224 through 239, inclusive.

Class D addresses are currently used primarily for the multicast backbone (MBONE) of the Internet. Many routers, including those from Lucent, do not support MBONE or multicast and therefore ignore class D addresses.

Class E Addresses

The class E IP address is reserved for future use. In class E addresses, the 4 highest-order bits are set to 1, 1, 1, and 1. Routers currently ignore class E IP addresses.

Reserved IP Addresses

Some IP addresses are reserved for special uses and cannot be used for host addresses. Table A-2 lists ranges of IP addresses and shows which addresses are reserved, which are available to be assigned, and which are for broadcast.

Table A-2   Reserved and Available IP Addresses

Class	IP Address	Status
A	0.0.0.0	Reserved
	1.0.0.0 through 126.0.0.0	Available
	127.0.0.0	Loopback networks on the local host
B	128.0.0.0	Reserved
	128.1.0.0 through 191.254.255.255	Available
	191.255.0.0	Reserved
C	192.0.0.0	Reserved
	192.0.1.0 through 223.255.254.255	Available
	223.255.255.0	Reserved
D	224.0.0.0 through 239.255.255.255	Multicast group addresses
E	240.0.0.0 through 255.255.255.254	Reserved
	255.255.255.255	Broadcast

Private IP Networks

RFC 1597 reserves three IP network addresses for private networks. The addresses 10.0.0.0/8, 192.168.0.0/16, and 172.16.0.0/20 can be used by anyone for setting up their own internal IP networks.

IP Address Conventions

If the bits in the host portion of an address are all 0, that address refers to the network specified in the network portion of the address. For example, the class C address 192.31.7.0 refers to a particular network. Historically, this address was used as a broadcast.

The standard for broadcast is high, which uses all 1s in the host portion (for example, 192.168.1.255); however, many networks still use all 0s. The PortMaster can be configured either way and should be set to match the other systems on your network.

Note – Do not assign an IP address with all 0s or all 1s in the host portion of the address to a host on the network, because these are reserved as broadcast addresses.

With CIDR, networks are specified with an IP prefix and netmask length—for example, 172.16.0.0/16, 192.168.1.0/24, or 192.168.200.240/28.

IPX Addressing

An IPX address consists of 10 bytes (expressed in hexadecimal notation), which gives an IPX network host a unique identifier. IPX addresses are made up of the following two parts:

- Network segment address, expressed as 8 hexadecimal digits

  These 4 bytes (32 bits) specify on which network segment the node resides.

- Node address, expressed as dotted triplets of 4-digit hexadecimal numbers

  These 6 bytes (48 bits) provide the media access control (MAC) address of the node.

The two elements of the IPX address are separated by a colon. For example:

00000003:0001 8423 4567

Network segment address   Node address
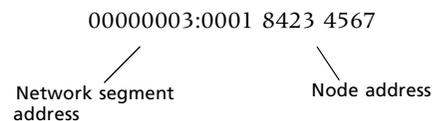
The first 8 digits represent the network segment, and the following 12 digits represent the node or MAC address of the node. All digits are expressed in hexadecimal.

Netmasks

A netmask is a four-octet number that identifies either a supernetwork (supernet) or a subnetwork (subnet). A netmask that designates a subnet is called a subnet mask.

Using Subnet Masks to Create IP Subnets

Subnet masks are used to divide networks into smaller, more manageable groups of hosts known as subnets. Subnetting is a scheme for imposing a hierarchy on hosts on a single physical network. The usual practice is to use the first few bits in the host portion of the network address for a subnet field. RFC 950, *Internet Standard Subnetting Procedure*, describes subnetting.

A subnet mask identifies the subnet field of a network address. This mask is a 32-bit number written in dotted decimal notation with all 1s (ones) in the network and subnet portions of the address, and all 0s (zeros) in the host portion. This scheme allows for the identification of the host portion of any address on the network.

Table A-3 shows the subnet masks you can use to divide a class C network into subnets.

Table A-3   Subnet Masks for a Class C Network

Length (Subnet Bits)	Number of Subnets	Number of Hosts per Subnet	Hexadecimal Subnet Mask	Dotted Decimal Subnet Mask
24	1	254	0xffffff00	255.255.255.0
25	2	126	0xffffff80	255.255.255.128
26	4	62	0xffffffc0	255.255.255.192
27	8	30	0xffffffe0	255.255.255.224
28	16	14	0xfffffff0	255.255.255.240
29	32	6	0xfffffff8	255.255.255.248
30	64	2	0xfffffffc	255.255.255.252
32	256	1	0xffffffff	255.255.255.255

Subnetting, Routing, and VLSMs

Routers and hosts can use the subnet field for routing. The rules for routing on subnets are identical to the rules for routing on networks.

Releases before ComOS 3.5. Before ComOS 3.5, correct routing required all subnets of a network to be physically contiguous. The network must be set up so that it does not require traffic between any two subnets to cross another network. Also, RFC 950 implicitly requires that all subnets of a network have the same number of bits in the subnet field. As a result, ComOS releases before ComOS 3.5 require the use of the same subnet mask for all subnets of a network. ComOS used the value of 255.255.255.255 for the user's *Framed-IP-Netmask* regardless of the value of the attribute.

ComOS 3.5 and Later Releases. ComOS 3.5 and subsequent releases support variable-length subnet masks (VLSMs); therefore, the restrictions in earlier ComOS releases no longer apply. The subnets of a network need not be physically contiguous and can have subnet masks of different lengths.

However, ComOS still ignores the *Framed-IP-Netmask* value by default. To ease the transition to use of VLSMs, ComOS sets **user-netmask** to **off** by default. This means that all netmasks specified in the user table or RADIUS are treated as if they were 255.255.255.255. To use VLSMs and have ComOS accept the value in Framed-IP-Netmask, enter the following commands:

```
Command> set user-netmask on
Command> save all
```

⚠ **Caution –** The VLSM feature affects both routing and proxy ARP on the PortMaster and must be used with caution.

Using Naming Services and the Host Table

Naming services are used to associate IP addresses with hostnames. Many networks use the Domain Name System (DNS) or the Network Information Service (NIS) for mapping hostnames to IP addresses. Both services are used to identify and locate objects and resources on the network. To use DNS or NIS, you must specify the IP address of the name server during the configuration process.

The PortMaster enables you to specify an internal host table, which can be used in addition to DNS and NIS. The host table allows each unique IP address to be aliased to a unique name. The host table is consulted when a port set for host access prompts for the name of the host. The table is used to identify the IP address of the requested host. If the user-specified hostname is not found in the host table, then NIS or DNS is consulted.

Note – The internal host table should be used only when no other host mapping facility is available. Using the host table only when necessary reduces confusion and the amount of network maintenance required.

Managing Network Security

PortMaster products allow you to maintain network security using a variety of methods. **Security** is a general term that refers to restricting access to network devices and data. To enable security features, you must identify sensitive information, find the network access points to the sensitive information, and secure and maintain the access points.

PortMaster security methods include

- Callback for remote access users

- Assignment of local passwords before connections are established

- Access control filters for host connections

- Inbound and outbound packet filtering

- IP packet filtering by protocol, source and destination address, and port

- IPX packet filtering by source and destination network, node, and socket

- SAP filtering

- PAP and CHAP authentication protocols for PPP connections

- Password security for administrative access

- Remote Authentication Dial-In User Service (RADIUS) or PortAuthority™ RADIUS support

- ChoiceNet filtering

Each of these security methods is described in more detail in this guide. All or some of these security methods can be configured as you configure the system-wide parameters and each interface. RADIUS, PortAuthority RADIUS, and ChoiceNet are described briefly in the next sections.

PortAuthority RADIUS must be purchased separately.

RADIUS

RADIUS is a nonproprietary protocol invented by Lucent and described in RFC 2138 and RFC 2139. RADIUS provides an open and scalable client/server security system for distributed network environments. The RADIUS server can be adapted to work with third-party security products. Any communications server or network hardware that supports the RADIUS protocol can communicate with a RADIUS server.

RADIUS consolidates all user authentication and network service access information on the authentication (RADIUS) server. The server can authenticate users against a UNIX password file, NIS databases, or separately maintained RADIUS database. The PortMaster acts as a RADIUS client: it sends authentication requests to the RADIUS server, and acts on responses sent back by the server. For more information about RADIUS, refer to the *RADIUS for Windows NT Administrator's Guide* and *RADIUS for UNIX Administrator's Guide*.

ChoiceNet

ChoiceNet is a client/server packet-filtering application created by Lucent. ChoiceNet provides a mechanism to filter network traffic on dial-up remote access, synchronous leased line, or asynchronous connections. Filter information is stored in a central location known as the ChoiceNet server.

ChoiceNet clients can be one or more PortMaster products. ChoiceNet clients communicate with the ChoiceNet server to determine user access.

ChoiceNet can use filter names specified by the RADIUS user record. For more information about ChoiceNet, refer to the *ChoiceNet Administrator's Guide*.

PortAuthority RADIUS

Lucent's PortAuthority RADIUS software provides enhanced RADIUS functionality and must be purchased separately.

TCP and UDP Ports and Services B

Table B-1 lists common port numbers—**well-known ports**—assigned to TCP and UDP services—**well-known services**—by the Internet Assigned Network Numbers Authority (IANA). A more complete list is available in RFC 1700, *Assigned Numbers*.

Note – If you are configuring a filter on a PortMaster from the command line interface, you must use the port number. The PortMaster does not have the **/etc/services** file and cannot use NIS to get the equivalent information.

Table B-1 TCP and UDP Port Services

Service	Port	Portocol	Description
ftp-data	20	TCP	File Transfer Protocol (FTP) (default data)
ftp	21	TCP	FTP (control)
telnet	23	TCP	Telnet
smtp	25	TCP	Simple Mail Transfer Protocol (SMTP) (email)
nicname	43	TCP	**whois** Internet directory service
nicname	43	UDP	**whois** Internet directory service
domain	53	TCP	Domain Name System (DNS)
domain	53	UDP	DNS
tftp	69	UDP	Trivial File Transfer Protocol (TFTP)
gopher	70	TCP	Gopher
gopher	70	UDP	Gopher
finger	79	TCP	Finger Protocol
finger	79	UDP	Finger Protocol
www-http	80	TCP	World Wide Web Hypertext Transfer Protocol (HTTP)
kerberos	88	TCP	Kerberos authentication
kerberos	88	UDP	Kerberos authentication
pop3	110	TCP	Post Office Protocol (POP) version 3
sunrpc	111	TCP	SUN Remote Procedure Call (RPC)
sunrpc	111	UDP	SUN RPC
auth	113	TCP	Authentication service
auth	113	UDP	Authentication service
nntp	119	TCP	Network News Transfer Protocol (NNTP)
ntp	123	TCP	Network Time Protocol (NTP)

Table B-1   TCP and UDP Port Services *(Continued)*

Service	Port	Portocol	Description
ntp	123	UDP	NTP
snmp	161	TCP	Simple Network Management Protocol (SNMP)
snmp	161	UDP	SNMP
snmptrap	162	TCP	SNMP system management messages
snmptrap	162	UDP	SNMP system management messages
imap3	220	TCP	Interactive Mail Access Protocol (IMAP) version 3
imap3	220	UDP	IMAP version 3
https	443	TCP	HTTP with Secure Sockets Layer (SSL) protocol—secure HTTP
exec	512	TCP	Remote process execution
login	513	TCP	Remote login
who	513	UDP	Remote **who** daemon (**rwhod**)
cmd	514	TCP	Remote command (**rsh**)
syslog	514	UDP	System log facility
printer	515	TCP	Line printer daemon (LPD) spooler
talk	517	TCP	Terminal-to-terminal chat
talk	517	UDP	Terminal-to-terminal chat
ntalk	518	TCP	Newer version of Terminal-to-terminal chat
router	520	UDP	Routing Information Protocol (RIP)
uucp	540	TCP	UNIX-to-UNIX Copy Protocol (UUCP)
uucp	540	UDP	UUCP
uucp-rlogin	541	TCP	Variant of UUCP/TCP
uucp-rlogin	541	UDP	Variant of UUCP/IP
klogin	543	TCP	Kerberized login
klogin	543	UDP	Kerberized login
pmd	1642	TCP	PortMaster daemon **in.pmd**
pmconsole	1643	TCP	PortMaster Console Protocol
radius	1645	UDP	Remote Authentication Dial-In User Service (RADIUS)
radacct	1646	UDP	RADIUS accounting
choicenet	1647	UDP	ChoiceNet
l2tp	1701	UDP	Layer 2 Tunneling Protocol

Command Index

A

add dlci (synchronous port) 12-4
add dlci W1 12-6
add filter 8-4
add Ippool 2-10, 2-11
add ipxdlci (synchronous port) 12-4
add location 7-2
add location sub1 12-8
add modem 9-3
add netmask 2-20
add netuser 5-2
add route Ipaddress 2-17
add route Ipxnetwork 2-18
add snmphost 3-13
add subinterface 4-7
add user 5-2

C

clear alarm 3-14

D

delete filter 8-7
delete Ippool 2-11
delete route Ipaddress 2-17
delete route Ipxnetwork 2-18
delete snmphost 3-13
delete subinterface 4-7
delete user 5-2
dial 7-11

R

reboot 10-13, 10-15
reset 6-5
reset all 9-5
reset C0 9-5
reset console 10-11, 12-7, 12-9, 13-4, 13-5
reset ippool 2-12
reset slot10 13-3
reset W1 6-8

S

save all 6-8, 10-13, 10-15
save route 2-18
set 2-2
set accounting 10-4
set all modem 9-5
set all speed 9-6
set assigned_address 2-8, 10-4
set authentic 10-3
set C0 modem 9-5
set C0 rip v2 4-3, 6-7, 7-6
set C0 rip v2 broadcast 4-3, 6-7, 7-6
set call-check 2-21
set chap 2-21
set chassis 2-6
set choicenet 10-4
set choicenet-secret 10-4
set compression 6-8
set console 2-4, 10-11, 12-7, 12-9, 13-4
set debug 12-7, 12-9, 13-4, 13-5
set debug imt 10-23
set debug mcppp 10-11
set debug mdp-events 10-11
set debug mdp-status 10-11
set debug nfas 10-19
set debug off 10-11
set dhcp server 2-14
set domain 2-3, 10-3
set Ether0 address 2-20, 4-4
set ether0 address 10-2
set Ether0 broadcast 4-5
set ether0 broadcast 10-2
set Ether0 ifilter 4-3
set Ether0 ip 4-5
set Ether0 ipx 4-6
set Ether0 ipxframe 4-6
set Ether0 ipxnet 4-5
set Ether0 netmask 4-4
set ether0 netmask 10-2
set Ether0 ofilter 4-4

Subject Index

K

keepalive timer
    Annex-D 12-4
    LMI 12-3

L

leased line connections
    Ethernet 13-2
    global settings 13-2
    overview 13-1
    synchronous WAN port 13-3
    troubleshooting 13-4
Line Call Summary table, MIB 3-11
line hangup 9-7
line speed, Frame Relay 12-2
line status 10-5
line use configuration 10-5
LMI
    enabling 12-3
    keepalives 12-3
    types 12-3
    use with DLCI 12-7
Local Management Interface. See LMI
location
    routing 7-6
locations
    defining 7-1
    dial groups, use of 6-5
    Frame Relay subinterfaces 12-8
location table
    adding a location 7-2
    CHAP 7-8
    compression 7-7
    destination IP address 7-5
    dial group 7-6
    displaying 7-2
    filters 7-10
    Frame Relay subinterfaces 12-8
    high-water mark 7-9
    idle timer 7-7
    IPX network number 7-5
    maximum dial-out ports 7-9
    MTU 7-6
    multiline load balancing 7-9
    netmask 7-5
    password 7-4
    protocol 7-4
    routing 7-6
    TCP/IP header compression 7-7
    testing 7-11
    username 7-4
loghost, setting 2-4

login host 5-8
login users
    adding to the user table 5-2
    description of 5-3
    See also user table
loopback, enabling on
    T1 or E1 lines 10-9
    T3 line 11-3

M

mailing lists, subscribing to xviii
Management Information Base. See MIB
manual connections 7-3, 7-4
maximum number of ports
    dial-in 5-6
    dial-out 7-9
messages, syslog 2-4
metrics
    hop count 2-18
    ticks 2-18
MFR2 signaling 10-14
MIB
    extensions for PortMaster 3-5
    overview 3-2
    structure of 3-2
modem ports 9-1
modems
    adding to modem table 9-3
    analog 10-10
    automatic configuration 9-2
    ComOS 4.0 limitations 9-1
    control 6-4, 9-6
    control signals 9-2
    digital to analog 10-10
    hardware flow control 9-2, 9-7
    initialization strings 9-3
    interna 3-8
    line hangup 9-7
    outbound traffic 9-6
    parity checking 9-6
    RTS/CTS 9-2, 9-7
    status 9-2, 10-10
    synchronizing speed 9-6
    table 9-3
    True Digital 10-10
Modem table
    internal, MIB 3-8
modem table
    external, adding 9-3
    external, displaying 9-2
modem table, external modems
    displaying 9-2
monitoring, SNMP 3-12

U

UDP
    for NFAS 10-17
    packets, filtering 8-6
    services and ports B-1
username, setting for dial-out 7-4
users
    deleting 5-2
    displaying configuration information 5-2
    restricting access to hosts 5-9
    session limit 5-4
user table
    access filters 5-9
    adding users 5-2
    compression 5-7
    displaying 5-1
    IP address 5-4
    IPX network number 5-5
    login host 5-8
    login service 5-9
    login users 5-2
    maximum ports 5-6
    MTU 5-6
    packet filters 5-7
    session limit 5-4
    setting the protocol 5-4
    TCP/IP header compression 5-7
    user types 5-2
utilities for concurrent connections 2-4

V

view, setting 1-3, 6-2, 10-5, 12-6, 13-3
    for global settings 2-2
virtual switch 12-1
VLSM A-6

W

WAN ports
    example configuration 12-6, 13-3
    general settings 6-2
    See also synchronous ports
warning icon xvi
well-known ports B-1
well-known services B-1
wink start 10-12

X

X.75 protocol
    dial-out locations 7-4
    for network users 5-4