

Peregrine

Network Discovery Reference Manual

Version 5.2

Copyright © 2004 Peregrine Systems, Inc. or its subsidiaries. All rights reserved.

Information contained in this document is proprietary to Peregrine Systems, Incorporated, and may be used or disclosed only with written permission from Peregrine Systems, Inc. This book, or any part thereof, may not be reproduced without the prior written permission of Peregrine Systems, Inc. This document refers to numerous products by their trade names. In most, if not all, cases these designations are claimed as Trademarks or Registered Trademarks by their respective companies.

Peregrine Systems® and Network Discovery® are registered trademarks of Peregrine Systems, Inc. or its subsidiaries. Microsoft, Windows, Windows NT, Windows 2000, and other names of Microsoft products referenced herein are trademarks or registered trademarks of Microsoft Corporation. DB2 is a registered trademark of International Business Machines Corp.

This document and the related software described in this manual are supplied under license or nondisclosure agreement and may be used or copied only in accordance with the terms of the agreement. The information in this document is subject to change without notice and does not represent a commitment on the part of Peregrine Systems, Inc. Contact Peregrine Systems, Inc., Customer Support to verify the date of the latest version of this document.

The names of companies and individuals used in the sample database and in examples in the manuals are fictitious and are intended to illustrate the use of the software. Any resemblance to actual companies or individuals, whether past or present, is purely coincidental.

If you need technical support for this product, or would like to request documentation for a product for which you are licensed, contact Peregrine Systems, Inc. Customer Support by email at support@peregrine.com.

If you have comments or suggestions about this documentation, contact Peregrine Systems, Inc. Technical Publications by email at doc_comments@peregrine.com.

This edition of the document applies to version 5.2 of the licensed program.

Peregrine Systems, Inc.
3611 Valley Centre Drive San Diego, CA 92130
Tel 800.638.5231 or 858.481.5000
Fax 858.481.1751
www.peregrine.com



Contents

Chapter 1	How Network Discovery Works	11
	Exploration and Discovery.	11
	Network Explorer	13
	Update Network Model.	13
	XML Enricher	14
	Modeling	14
	Filtering Out of Database	14
	Rulebase	15
	Checkpoint	19
	Poll Device	19
	Schedule Scan	19
	Table Reader	20
	Mapper	20
	How long does all this take?	20
	Network Discovery is always discovering	21
	Detecting broken devices	22
	All Devices	24
	High Priority Devices.	24
	Alarms and Events	25
	Aggregate Health and Events.	26
	How much network bandwidth does the appliance need?.	27
	Discovery Ping Rate	28
	Table Reading and Polling.	29
	WAN	30

	RFCs supported by Network Discovery	32
	Data from Peregrine Desktop Inventory	33
	Communication models.	33
	Frame relay	33
	FDDI	36
	HSRP	39
	Presenting Information	40
	Network Map	40
	Scheduled Events.	45
Chapter 2	Terms and Concepts	47
	Network Terms and Concepts	48
	Domain names	48
	Address types	48
	Netmask notation	50
	Community strings.	51
	Bridge aging.	51
	OSI model layers.	52
	Management workstation	53
	Network Discovery Terms and Concepts.	53
	Objects	53
	Devices.	55
	Real devices	55
	Virtual devices.	58
	Packages	61
	Priority.	63
	Special input syntax	64
	Events and Alarms	65
	Event Thread	66
	Panel Elements	67
Chapter 3	Recorded Events	71
	Line Breaks	71
	Utilization	71
	Delay	74
	Jitter	74

Collisions	74
Broadcasts	75
Errors	75
Frame Relay	75
Device Breaks	75
Packet Loss	76
Disk Utilization	76
CPU Utilization	76
Load Average	76
Memory Utilization	77
Backplane Utilization	77
Printer	77
UPS	77
Port MTTR	77
Port MTBF	77
Port Add/Deletes	77
Port Moves	78
Port Changes	78
Device MTTR	78
Device MTBF	78
Device Adds/Deletes	78
Device Moves	79
Device Changes	79
Exceptions	79
Not Recently Seen	80
Open Tickets	80
Chapter 4 Device Manager	81
Introduction	82
Toolbar	83
Configuration	85
State	92
Reports	93
Diagnosis	95
Diagnostic Information	95
IP Ping	99

	Traceroute	99
	SNMP Ping	100
	Listener Ping	101
	DNS Query	101
	Statistics	102
	Ports	105
	Events	105
	Locate	106
	Service Analyzer	106
	Manage	106
	Browse MIB	107
	View Scan Data	107
	Web	107
	Telnet	108
	Update Model <i>[Administrator or IT Manager]</i>	108
	Device Visibility <i>[Administrator or IT Manager]</i>	110
	Properties	110
	Refresh	111
	Print	111
	Text	112
	Close	112
Chapter 5	Port Manager	113
	Introduction.	114
	Toolbar	115
	Configuration	116
	State	118
	Reports	119
	Diagnosis	121
	Statistics	123
	Events	124
	Locate	124
	Purge Port <i>[Administrator or IT Manager]</i>	124
	Create Connection <i>[Administrator or IT Manager]</i>	125
	Break Connection <i>[Administrator or IT Manager]</i>	126

	Port Properties <i>[Administrator or IT Manager]</i>	126
	Refresh	128
	Print	128
	Text	128
	Close.	128
	Port number	129
Chapter 6	Line Manager	131
	Single Line Manager	131
	Toolbar	132
	About	132
	Break Connection <i>[Administrator or IT Manager]</i>	134
	Refresh.	134
	Print.	134
	Text	134
	Close.	135
	Multiple Line Manager	135
Chapter 7	Attribute Manager	137
	Introduction	138
	Toolbar.	139
	Configuration	140
	Statistics	142
	Locate	142
	Manage <i>[Administrator or IT Manager]</i>	142
	Purge Attribute <i>[Administrator or IT Manager]</i>	143
	Refresh	143
	Print	143
	Text	143
	Close.	144
Chapter 8	MIB Browser	145
	Opening the MIB Browser.	146
	Parts of the MIB Browser	147
	Tree View.	148
	Pull-down list of Devices	149

	Find Function	149
	Locate on Map	149
	Get Next	149
	Refresh	149
	Folder Tab	149
	Variable Tab	151
	MIB Description	152
	Read and Write Community Strings	152
	Walking the MIB	154
	Using Multiple MIB Browser Sessions	154
	Watching an OID with MIB Radar	155
	Saving MIB Data as a Text file	155
	Save Table Data	155
	MIB Walk	156
Chapter 9	Reports	157
	Report periods	157
	Report statistics	159
	Executive/Summary Network Reports	160
	Scanned Machine Reports	163
	WAN Reports	164
	LAN Reports	166
	Device Reports	167
	Microsoft Word documents	168
	You can use Network Discovery data with other applications	169
Chapter 10	Desktop Inventory Reports	171
	The Reports database	172
	Open Database Access	172
	Reports available on the appliance for scanned machines	173
	Accessing the Scanned Machine Reports on the appliance	173
	Scanned Machine Summaries	174
	Application Reports	176
	Scan File Status Reports	179
	Producing your own reports	179

Chapter 11	Need more help?	181
	Peregrine’s CenterPoint Web Site.	181
Chapter 12	Copyright	183
Chapter 13	Glossary of Abbreviations	205
	Index.	215

1 How Network Discovery Works

CHAPTER

You can use Network Discovery without ever having to read or refer to this section of the manual. However, experienced Network Discovery Administrators may find it easier to understand certain aspects of the behavior of Network Discovery after reading this section.

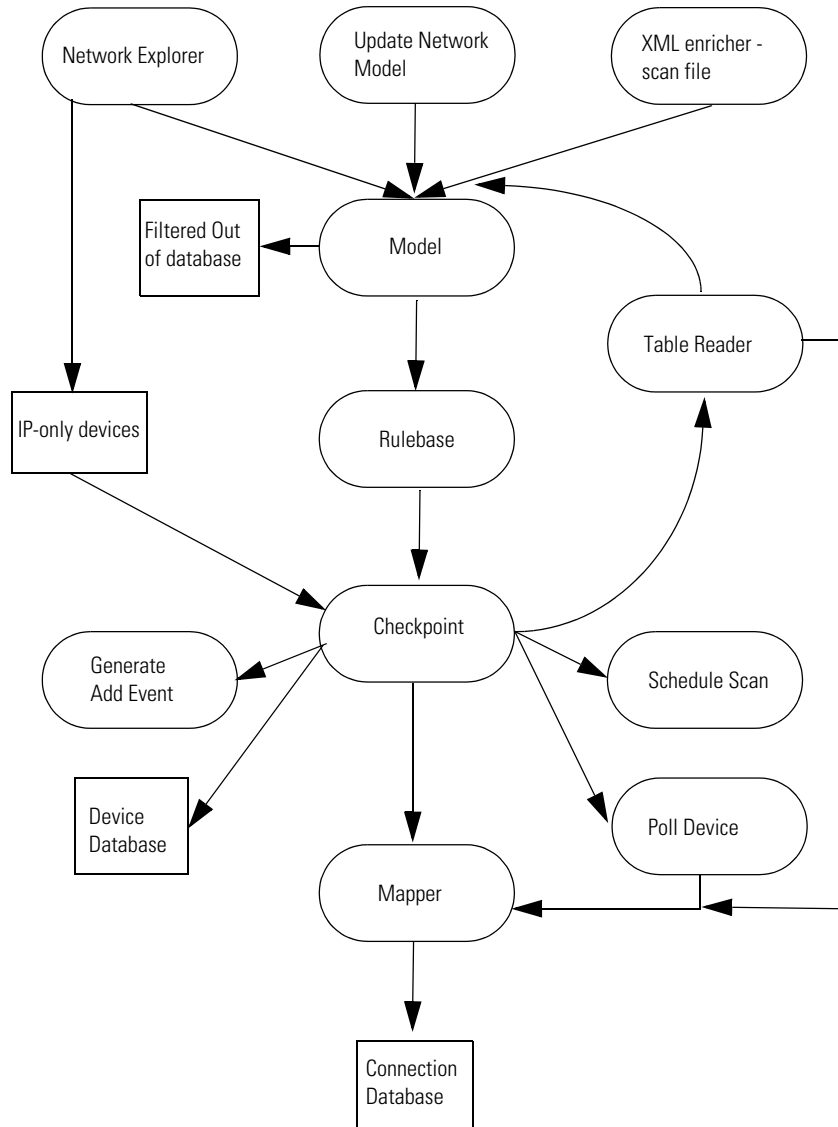
Exploration and Discovery

There are a few basic ways Network Discovery will discover new devices:

- from the Network Explorer
- from a scan file (courtesy of the XML enricher)
- through an update network model function

The flow chart on page 12 shows a basic representation of how devices are discovered by Network Discovery. Each section of that chart will be described below.

Figure 1-1: Discovery flow chart



When you prepare Network Discovery for exploration, then set it going, the Network Explorer begins by exploring the IPv4 ranges you have set up for “Active discovery” in **Administration > Network Configuration**.

Network Explorer

Discovery is strictly a yes-or-no proposition. The Network Explorer goes through each IP address in random order and pings it once to see whether or not there is a response. Is there a device at this address or not? When there is a positive response, the IP address is added to the database (if it is not already in the database) and is sent to the modeler.

Note: By default, the Network Explorer will ping each address once. You can change the number of pings at **Administration > System preferences > Network Devices > Number of Explorer pings**.

This means that the average time to discover a device can be calculated as the number of IP addresses in the IPv4 range divided by ping rate.

The Explorer works continuously to find any new devices in the configured IP ranges. For faster discovery, the Explorer also tracks devices that have responded positively and omits them the next time.

Update Network Model

The Update Network Model feature forces discovery of a device that Network Discovery has not found on its own.

This command also requests that the device have its network model updated immediately. This is helpful if you have made any physical changes to a device.

From the Device Manager, click the **Update Model** button and select **Query Network**. You can also select this command from the Network Map **Object** menu. You can also update a model in **Administration > Data management > Update network model**.

XML Enricher

The XML Enricher works if you are also using Peregrine Desktop Inventory (PDI). You can add scan files to a shared directory on the Peregrine appliance, or you can use PDI with Network Discovery to regularly scan devices in your network. For more information on these topics, see *Using Network Discovery with Desktop Inventory and Desktop Administration*.

Modeling

Every device has a “device model” in the Network Discovery database. When Network Discovery is finding information about a device, it creates a model based on the device MIB (if possible).

Network Discovery attempts to find out the device’s community strings, its domain name, NetBIOS name, how many ports each device has, and what type of device it is—whether it supports bridge tables, arp tables, Cisco CDP, source address capture, and so on.

Network Discovery will supplement that information with data from ARP caches from other devices (routers) in the network for unmanaged devices or if you have enabled the “accumulate IP addresses” in **Administration > Network Configuration > Network Property Groups**.

Filtering Out of Database

Network Discovery can filter out certain unmanaged devices that you do not want to see on your map or in your Network Discovery database.

You can change these filter settings in **Administration > System Preferences > Input filters > Devices not to add to the database**:

- Unmanaged devices which are MAC plus IP
 - Unmanaged devices which are MAC plus IP and not pingable
- Unmanaged devices which are MAC-only
 - Unmanaged devices which are MAC-only with unknown OUIs
- Unmanaged devices which are IP-only
- Scanned-only devices

All of these filter settings change which devices can be monitored by Network Discovery. If a device is filtered so it does not appear on the Network Map, the device data will be found at **Status > Filtered Devices**.

Rulebase

Data Going Into the Rulebase

The following fields are passed to the Rulebase so it can make a determination about the device:

- system.sysDescription
- system.sysObjectID
- DNS name
- MAC address
- NetBIOS name
- forwarding information (whether or not the device is routing IPv4 or IPv6)
- read community string
- write community string
- number of ports

The following fields are available for scanned devices:

- | | |
|--------------------------------|----------------------------|
| ■ Operating System | hwHostOS |
| ■ BiosProductName | hwBiosMachineModel |
| ■ BiosProductManufacturer | hwsmbiosSystemManufacturer |
| ■ BiosChassis | hwsmbiosChassisType |
| ■ BiosDescription | hwBiosMachineDescription |
| ■ Operating System ServicePack | hwOSServiceLevel |

Creating a Device Type

Network Discovery assigns a device type to each device based on the Network Discovery Rulebase. The device type can include the following characteristics:

- Model
- Model URL
- Model Manufacturer

- Family
- Family URL
- Family Manufacturer
- OS
- OS URL
- OS Manufacturer
- Network Function
- Network Function URL
- Network Function Manufacturer
- Software Historical Manufacturer
- Software Historical Manufacturer URL
- Software Present Manufacturer
- Software Present Manufacturer URL
- Historical Manufacturer
- Historical ManufacturerURL
- Present Manufacturer
- Present Manufacturer URL
- Icon
- Title
- Tag
- Priority

Adding Rules to the Rulebase

If a specific rule is not in the Rulebase, Peregrine Systems will add a rule for you, provided that:

- your Peregrine appliance is under warranty
- you can identify the devices by model or family (Peregrine would appreciate the URL for the manufacturer's web site whenever possible)
- you provide Peregrine with a CSV copy of your inventory containing the device.

Some devices and device families may not match a specific identification rule. Such rules are likely to make good assignments for companies with small product lines and less accurate assignments for companies with large product lines. This is because these rules are based on:

- advance classification of product lines; that is, some devices belonging to certain product lines can be identified by the beginning of the OUI
- pre-identification of specific devices; that is, some devices can be identified because the manufacturers make only switches

Note: These rules may make incorrect assignments. You should contact Peregrine to request additional specific rules for devices if this happens.

For devices with no SNMP management, the Rulebase can apply rules based only on information about the MAC address and OUI of the device. For each MAC address, the Rulebase identifies the most probable device class, based mostly on the OUI. (Very occasionally, manufacturers assign blocks of MAC addresses to specific products, which allows the Rulebase to make more specific identifications.)

The Rulebase also identifies the probability that each non-SNMP device may actually be an SNMP managed device providing network connectivity (such as a gateway, router, concentrator, or switch). SNMP managed devices can appear not to be managed when the device's IP address has been included in the list of Property Groups (see **Administration > Network configuration > Network Property Groups**) or when the community string for the device has not been included in the Network Discovery list of community strings (see (see **Administration > Network configuration > Community Property Groups**)). In such a case, you should install or enable the SNMP agent for that device. (You may also need to modify the address scope or community strings.)

As with devices with SNMP management, class assignments for devices with no management work well for companies with small product lines and poorly for companies with large, varied products lines. Larger companies sometimes employ the same OUI for different products, but also use different OUIs for one product.

There is also a capacity to assign icons to unmanaged devices based on information contained in the NetBIOS and domain names has been added. For example:

- a device named “PRINTER3RDFLOOR” or “PRT3RDFLOOR” could be assigned a printer icon
- a device named “marysworkstation” could be assigned a workstation icon
- a device named “webserver.example.com” could be assigned a web server icon

Only the initial segment of the domain name is considered.

Domain and NetBIOS name interpretation is low priority. It never takes precedence in a situation where more accurate information is available. The rules used are not case sensitive.

Printers

Finally, Network Discovery can identify printers attached to printer servers. Many printer servers (both internal and external) do not provide enough information in their System Description to allow for accurate identification of the specific model of printer attached.

The Network Discovery Rulebase uses information that may be found elsewhere in the device MIB. For example, the System Description of this Hewlett-Packard printer server contains the following:

- HP ETHERNET MULTI-ENVIRONMENT,ROM H.08.01,JETDIRECT EX,JD34,EEPROM H.08.05

Note that this does not provide any information about the printer. The Enterprise MIB contains additional information that allows the Rulebase to identify the printer server as J3263A and the printer model as a HP LaserJet 5.

Checkpoint

Network Discovery runs database checkpoints at least every hour. Checkpoints are run more often when there a lot of new devices being discovered.

This checkpoint is the process of committing the network information to the database, and also taking a “snapshot” of the network. This snapshot can be used as a starting point in the event of the Peregrine appliance failing.

Note: A newly discovered device is not visible until it has been included in a checkpoint.

When the checkpoint is complete, the device model has been created, and the data has been saved into the Network Discovery database.

Poll Device

Note: Scanned-only devices are not polled.

Once there are device models in the database, the pollers begin collecting data from the devices. There are three pollers: the Realtime Poller, the Resource Poller, and the Environmental Poller.

The Realtime Poller also reads the device’s MIB to get information about the device’s traffic and connectivity on all of its ports. The resulting data is passed to the Mapper to determine connectivity.

Note: The poller will start polling a device within 10 sampling periods after a checkpoint.

Schedule Scan

For more information on creating scan schedules, see *Using Network Discovery with Desktop Inventory and Desktop Administration*. You must have both Network Discovery and Desktop Inventory in order to use this feature.

Table Reader

Note: The Table Reader applies only to SNMP-managed devices.

Unmanaged MAC-only devices (and MAC-only devices with unknown OUIs) will be discovered after Network Discovery completes modeling devices with bridge table reader.

Connectivity information comes from the Table Reader. If Network Discovery has identified the device as a bridge, the Table Reader reads its bridge tables. If the device has been identified as a router, or if “Force ARP Table Read” has been enabled in **Administration > Network Configuration > Network Property Groups**, the Table Reader reads its ARP table.

Mapper

The device will appear on the map after the checkpoint and up to two sampling periods have passed.

The Network Mapper takes the list of devices and ports from the Device Poller and the information from the Table Reader. Using this data, the Mapper deduces and how the devices should be connected.

How long does all this take?

The time it takes for a device to appear on the map depends on many circumstances.

When you first install your Peregrine appliance, and set up Network Discovery, you will start to see devices appear on the map within 10 minutes. The first device to appear will normally be the Peregrine appliance itself, and it will first be linked to the Logical View virtual device, which in turn will be connected to an Unmapped IP virtual device (representing the subnet where the Peregrine appliance resides).

For all other devices, the time to appear on the map varies. Discovery is a continual process, so keep in mind that when a new device appears on the Network Map, there may be additional data coming that will change its appearance on the map, or its location and connectivity.

Generally, it can work like this:

- A device will be discovered based on the discovery ping rate and the IPv4 ranges you have set.
- Creating a device model will take anywhere from 30 seconds to 2 hours. The time will vary depending on the type of device, and the number and order of the community strings.
- A device from the database checkpoint appears on the Network Map - up to two sampling periods, each sampling period is less than one hour.

The time may be increased for some devices if Network Discovery needs to try several community strings to access the device MIB.

Network Discovery is always discovering

This process runs the entire time Network Discovery is in operation.

Also, every device is re-modeled at the device remodeling interval specified in **Administration > Network configuration > Network Property Groups**.

This way, Network Discovery constantly strives to present you with an updated view of your network, and constantly strives to improve the accuracy and depth of that view.

Detecting broken devices

Break fault analysis is an important part of network monitoring. Network Discovery starts break fault analysis during the first sampling period after a device has been polled. If the device is not responding, Network Discovery will consider it to be broken:

- For SNMP-managed devices, or unmanaged devices that are connected to an SNMP-managed device, the device will be classified as broken if it has not responded to a ping or poll for two poll cycles.
- For unmanaged devices, Network Discovery relies on data from the tables of other SNMP-managed devices. If the table information in the other devices indicates that an unmanaged device has not responded for 48 hours, the device will be classified as broken.

Network Discovery cannot diagnose breaks on the following types of devices:

- virtual devices
- scanner-only devices
- devices with no IP or MAC address, such as printers connected to print servers

The representation of a break depends on the line alarm type. The following examples show a connection between two routers, and will explain how different line alarm types can influence the type of break alarm.

In Figure 1-2, there is an example of a port being disconnected. If the connection is an Ethernet line, Network Discovery will report a break on Router B. If the connection is an ATM line, Network Discovery will report a line break, which originates at Port 1.

Figure 1-2: Example: Unplug port 1



In Figure 1-3, the same port has been disconnected, but there is a redundant connection between the routers. Regardless of the line type, Network Discovery will report a line break, which originates at Port 1. If both Port 1 and Port 2 are down, there will be a device break reported on Router B, because it is no longer reachable.

Note: If there is a break on the port that controls SNMP management on Router B, it will have a break alarm.

Figure 1-3: Example: 2 connections to the device, unplug port 1


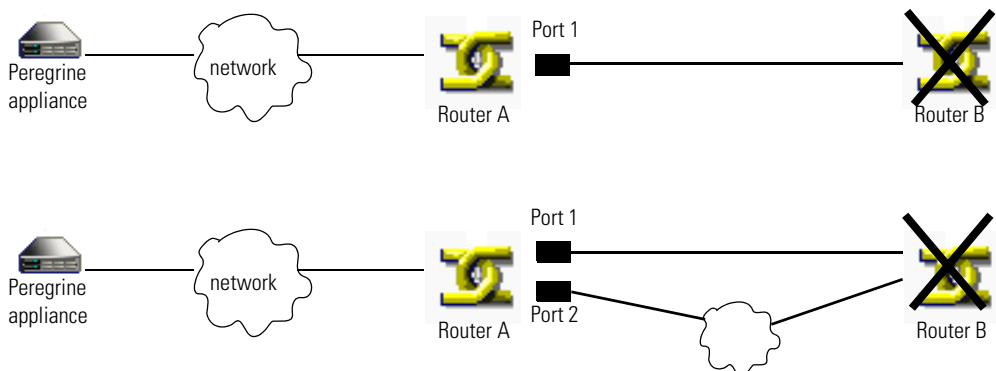


Figure 1-4 shows scenarios similar to the previous two examples. The difference here is that Router B has been shut down, while the ports on Router A are functioning normally. In this scenario, Network Discovery will report a device break for the Router B, regardless of the line alarm type.

Figure 1-4: Example: Broken Router



All Devices

Major break alarms indicate a 99.5+% certainty that a device is broken, and take some time to appear. Minor break alarms give more rapid alerts of possible breaks.

There are many reasons for Network Discovery to consider a device to be broken:

Reason	Break Alarm Type
The device is physically broken.	Major ^a
The device is powered off.	Major ^a
The line is broken (a line is no longer capable of carrying traffic).	Major ^a
No longer responding to SNMP: <ul style="list-style-type: none"> ■ management failure ■ community string has been changed ■ the Peregrine appliance is no longer on the device's access list 	Minor
The device is not responding to ping.	Minor
The device cannot be seen by Network Discovery because it is being blocked by another device that has a break alarm (i.e., the device is in a “fault shadow” and may not be broken).	Minor

^a Major alarms may first be Minor alarms until Network Discovery confirms the break.

High Priority Devices

It is important to know when high priority devices are failing. For example, if a router fails, more users will notice the problem than if it was a broken workstation or printer.

Network Discovery uses “fast break detection” on high priority devices to quickly detect when a device is broken. For Network Discovery to use fast break detection on a device, the device must:

- be priority 3 or higher

- have a known IP address
- have been responding to pings in the last 12 hours - 4 days (depending on the length of the sampling period)

If a high priority device has not responded to Network Discovery in the last 20 seconds, Network Discovery will send up to 3 extra fast break pings to see if that device is still functional. If the device does not respond to any of the pings, Network Discovery determines that the device is broken, and raises an alarm.

Note: Typically, most networks have a relatively small number of high priority devices that would use this feature. If you have a large number of high priority devices in your network, it may take a few minutes for Network Discovery to detect a failure with fast break detection.

Alarms and Events

Note: When reading this explanation, assume that for each device there are: $(1 \text{ device} + 6 \text{ ports}) * 20 \text{ attributes} = 140 \text{ attributes/device}$. This means an appliance with 15,000 devices has 2,100,000 attributes.

Network Discovery detects a problem, and sends the data to the “internal thresholding logic.”

Each attribute recorded by Network Discovery has thresholds associated with it. The thresholds relate to alarm types: OK, Info, Minor, Major, Critical. An event is triggered each time an attribute changes state (for example, from OK to Minor, or Minor to OK). An alarm is created if the attribute state value is within an alarm threshold (Info, Minor, Major, Critical).

For example, you could set your Disk Space thresholds to the following settings:

- OK - 0-10%
- Info - 10-20%
- Minor - 20-40%
- Major - 40-75%
- Critical - 75-100%

If the value of the event falls within an alarm range (for example, Disk space is 90% full), an alarm will be created (in this example, a Critical alarm).

Once Network Discovery creates an alarm, this triggers a message to the internal “event notification system.” If the user has created an Event Filter (**Administration > Event Filter Configuration**) for this change of alarm state, the user will be notified immediately through an e-mail, pager message, SNMP trap, or with a new ticket in ServiceCenter.

Note: When you are creating Event Filters, you can input a notification delay. This means that Network Discovery will not notify you of the alarm until after the delay period, in case the problem is quickly corrected.

All new values are cached to show the current state of the device. Once a minute (see Note below), that state is updated in the Network Discovery internal database. Every 30 seconds, the Health Panel checks the database for updates and refreshes the user interface. The user can press the F5 key to request an immediate database query. The refresh time should be less than 1 sec.

Note: How often the states are updated in the database depends on how many attributes you are monitoring. If you have <500,000, the database is updated every minute. If you have 1,000,000 attributes, the database is updated every two minutes.

The Health Panel also shows reporting data:

- Add/Delete/Move/Change/Exceptions are updated in the database every 15 minutes.
- MTTR/MTBF are updated once per day

Aggregate Health and Events

The Aggregate Health Panel works a bit differently. There are a few more steps involved.

First, it relies on the above being completed on each appliance. The Health Panel data is then extracted from the realtime database and made available to the aggregator. Like the regular Health Panel, this also has a time delay depending on the number of attributes in the database. If you have <500,000 attributes, the database is updated every 5 minutes. For each additional 100,000 attributes, the database takes an additional minute to update (for example, if you have 700,000 attributes, the database will update every 7 minutes).

Once the data is imported into the Reports database, the data is then available for aggregation (the Reports database is visible via ODBC).

Every 5 minutes, the aggregator appliance will check if there is any data available on each of the configured remote appliances. If so it will transfer the data and put it in an incoming queue.

Also every 5 minutes the aggregator checks the incoming queue for each appliance configured to see if there are any files waiting. If so they are processed. If there a lot of files to process from many appliances, the time waiting in the queue will increase.

In an situation where the aggregator is not overloaded, and there network is not slow in transferring the file from the remote appliance to the aggregator appliance, you should see your data within 15 minutes.

How much network bandwidth does the appliance need?

There are many factors that contribute to the network traffic caused by Network Discovery. The best practice is to connect the appliance to a major backbone switch. It is estimated that the traffic would total 3-4% on the 10MB dedicated link between the appliance and the switch. On a 100MB or 1GB link, the impact is proportionally smaller.

The traffic initiated from the appliance is heaviest on that link to the switch. From the switch, the traffic going to the network is dispersed. It is impossible to say exactly how much bandwidth will be taken by Network Discovery, but in this section, you can read about some of the influences you may want to consider.

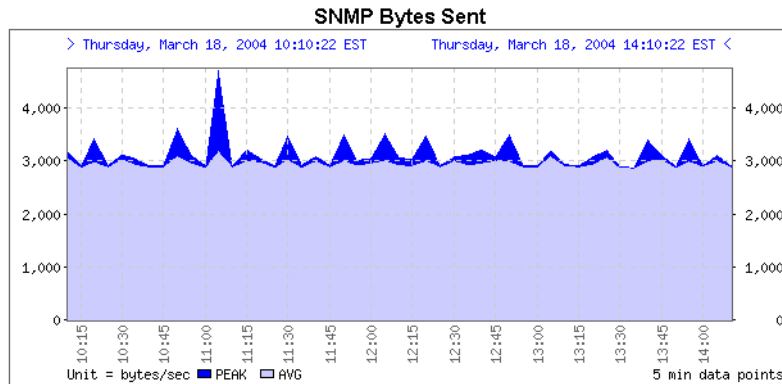
Network Discovery must contend with many different types of devices, each of which will contain varying amounts of data to be collected. Also, the many settings in Network Discovery can change how often the data is collected. This makes it difficult to offer an idea of how much bandwidth will be needed in any particular situation.

To check the traffic on your dedicated link, open a Port Manager for the Peregrine appliance's port eth0, and check Frames Out in the Statistics panel.

Also, on the Device Manager, check some of the Statistics graphs such as:

- SNMP Bytes
- SNMP Frames
- ICMP Frames

Figure 1-5: Example of SNMP Bytes graph from the Device Manager



Discovery Ping Rate

The Discovery Ping Rate (**Administration > System preferences > Network devices**) is one source of traffic. The ping sweep occurs in the background to look for new devices that have not been found previously. If you turn the ping rate down, it will take longer to discover new devices in your network.

If you turn Network Discovery's Ping feature off (**Administration > System preferences > Appliance services**), new devices will not be found through this method, and discovery will not generate any traffic on your network.

If you have configured Network Discovery to ping a large IPv4 range containing very few devices, there may be some network impact as pinging non-existent IP addresses will cause ARP broadcast requests.

Network Discovery has been configured to limit the ARP broadcasts it generates. However, Network Discovery may ping devices on the far side of a router. You should check your router configuration if broadcast levels become unacceptable for your network. On your router, consider the following:

- increase ARP cache size
- increase ARP aging time
- reduce ARP retry rate

Table Reading and Polling

Table reading and polling produce the majority of network traffic from the Peregrine appliance. These functions provide:

- connectivity information
- discovery of devices (for example, MAC-only devices)
- collection of statistics
- break fault analysis

A poll is really one frame out and one frame back in most cases. The number of polls for a device will depend on the number of ports in the device, and the number of attributes collected for each port (for example, collisions, broadcasts, etc.). The device itself is also pinged in each poll cycle.

Consider the fact that collecting statistics on a router with 200 ports requires a lot more effort than collecting statistics from a workstation with one port.

WAN

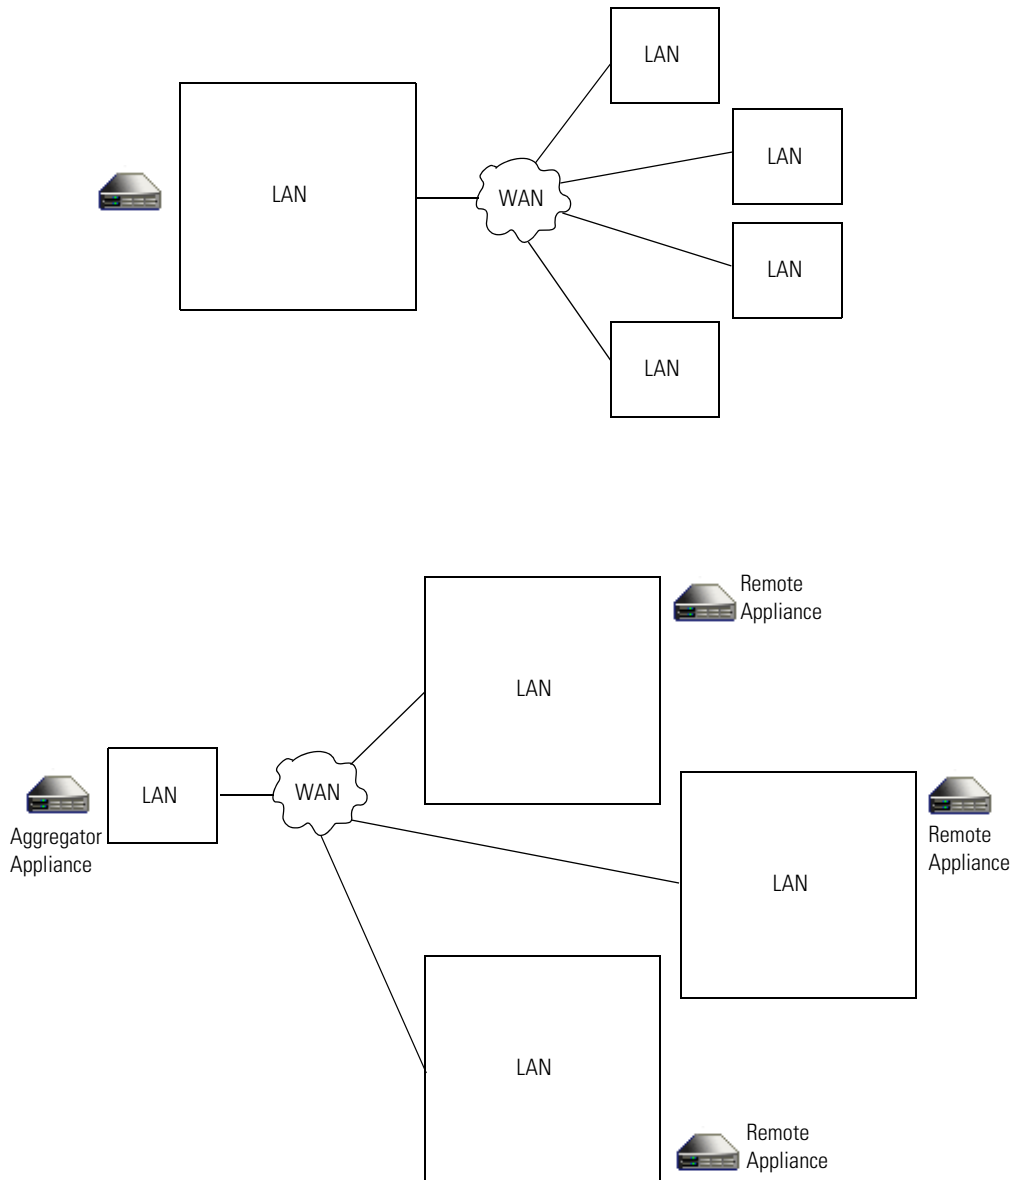
The bandwidth on a WAN link depends on the number of devices in your IPv4 ranges, the types of devices, and the many settings available in Network Discovery.

However, the real impact depends on the amount of bandwidth available, how much is used by normal network traffic, and what levels of extra traffic you are willing to accept.

Over a slow link, it would not be practical to fully manage a large network. Depending on the various parameters, you may only want to monitor the core devices. If you need to manage your entire network, it may be wise to get a second Peregrine appliance on the other side of your WAN link. The two appliances could be aggregated and share data.

There are many other considerations in a WAN scenario. In a large network (for example, 10,000 devices) with high rates, assume that the Peregrine appliance will use 5% of a 10MB link for network management.

Figure 1-6: Using Network Discovery in a WAN



RFCs supported by Network Discovery

Table 1-1: RFCs and specifications supported by Network Discovery

RFC number	Name
RFC 1155	Structure and Identification of Management Information for TCP/IP-based Internets
RFC 1157	A Simple Network Management Protocol (SNMP)
RFC 1213	<i>see</i> RFC 2011, RFC 2012, RFC 2013
RFC 1285	FDDI MIB (SMT 6.2); <i>see also</i> RFC 1512
RFC 1315	<i>see</i> RFC 2115
RFC 1354	<i>see</i> RFC 2096
RFC 1398	<i>see</i> RFC 1643
RFC 1406	Definitions of Managed Objects for the DS1 and E1 Interface Types
RFC 1407	Definitions of Managed Objects for the DS3/E3 Interface Type
RFC 1493	Definitions of Managed Objects for Bridges (Bridge MIB)
RFC 1512	FDDI MIB (SMT 7.3)
RFC 1513	Token Ring Extensions to the Remote Network Monitoring MIB
RFC 1514	Host Resources MIB
RFC 1516	Definitions of Managed Objects for IEEE 802.3 Repeater Devices
RFC 1643	Definitions of Managed Objects for the Ethernet-Like Interface Types (Ethernet Interface MIB)
RFC 1695	Definitions of Managed Objects for ATM Management Version 8.0 using SMIPv2 (ATM MIB)
—	ATM Forum 3.1 UNI specification
RFC 1748	IEEE 802.5 MIB using SMIPv2
RFC 1759	Printer MIB
RFC 2011	SNMPv2 Management Information Base for the Internet Protocol using SMIPv2
RFC 2012	SNMPv2 Management Information Base for the Transmission Control Protocol using SMIPv2
RFC 2013	SNMPv2 Management Information Base for the User Datagram Protocol using SMIPv2
RFC 2020	Definitions of Managed Objects for IEEE 802.12 Interfaces (100VG AnyLAN MIB)

Table 1-1: RFCs and specifications supported by Network Discovery (Continued)

RFC number	Name
RFC 2096	IP Forwarding Table MIB (Router MIB)
RFC 2115	Management Information Base for Frame Relay DTEs Using SMIv2 (Frame Relay MIB)
RFC 2233	Interfaces Group MIB using SMIv2
RFC 2668	Definitions of Managed Objects for IEEE 802.3 Medium Attachment Units (MAUs)
RFC 2674	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
RFC 2737	Entity MIB Version 2

Data from Peregrine Desktop Inventory

You can drop FSF or xml.gz files from PDI into the shared directory on the Peregrine appliance for processing.

Network Discovery accesses the files and presents the data:

- on the Network Map
- in Reports

The XML Enricher processes the XML scan files and passes information to the modeler. The remainder of the process continues as above except that scan-only devices are never polled.

Also, Network Discovery can schedule PDI scans with the help of Desktop Administration Listeners. For more information, see *Using Network Discovery with Desktop Inventory and Desktop Administration*.

Communication models

Frame relay

Network Discovery supports frame relay devices that conform to:

- RFC 2115, which supersedes RFC 1315

Each physical frame relay port may have one or more circuits associated with it. For some devices, Network Discovery is able to identify the circuits related to each physical port and gather traffic statistics both for the physical port and for each circuit. Network Discovery can also make connections between devices connected by these frame relay circuits.

The Device Manager Ports panel presents the ports so as to make apparent the association between a physical port and its circuits. For devices on which Network Discovery is able to do a physical port mapping, each port is displayed in the form *x.y.z*, where *x* represents the slot or card number on which the port *y* is located, and *z* represents the frame relay circuit.

Using a Cisco 7200 router as an example, here's how Network Discovery arranges the port structure:

```
...
1.5          —
1.6          —
1.7          frame relay physical port
1.7.27      frame relay circuit
1.7.32      frame relay circuit
2.1         —
2.2         —
...
```

If a device supports frame relay but Network Discovery is not able to map the exact physical ports, each port is displayed in form $x.y$, where x represents the MIB-II object `ifIndex` and y represents to frame relay circuit. Using a Cisco 2500 router as an example:

1	—
2	—
3	—
4	frame relay physical port
4.75	frame relay circuit
4.76	frame relay circuit
4.78	frame relay circuit
5	—
6	frame real physical port
6.21	frame relay circuit
6.27	frame relay circuit

The line speed is set for each frame relay circuit. Each circuit should report a Committed Information Rate (CIR).

The CIR has meaning only for frame relay lines. It is used in service-level agreements and contracts for supply of communications bandwidth over frame relay lines. CIR has no functional impact on the performance of frame relay devices. For Network Discovery to read the CIR from the device, it must have been entered into the device's MIB. If Network Discovery cannot find the CIR in the MIB, it sets the frame relay circuit CIR to the line speed for that frame relay physical port.

If Network Discovery has determined the CIR incorrectly, you can use the Port Manager's Port Properties button to redefine it (see [Properties](#) on page 110). You may change the interface rate at either end or at both ends.

The following examples and rules describe the effect of setting the interface rate to set the CIR.

Suppose a frame relay line connects device A port 1 and device B port 2. The CIR (A1-B2) is defined from A1 to B2. The CIR (B2-A1) is defined from B2 to A1 and can have a different value from the CIR (A1-B2).

Table 1-2: Effects of setting CIR (example)

A1		B2		CIR A1 to B2	CIR B2 to A1
line speed (kb/sec.)	set by user	line speed (kb/sec.)	set by user	line speed (kb/sec.)	line speed (kb/sec.)
100	no	200	no	100	100
100	no	50	no	50	50
100	no	100	no	100	100
100	yes	50	no	100	50
100	yes	200	no	100	100
100	yes	50	yes	100	50
100	yes	200	yes	100	200

The rules that constructed this table are:

- The line speed is read from the device's MIB unless overridden by the user setting it.
- If the line speed is set by the user at one end, the CIR from this end is defined as that line speed.
- If the line speed is not set by the user at an end, the lower speed at either end defines the CIR for an end.

FDDI

Network Discovery has limited support for FDDI:

- support for the SMT v6.2 MIB (specified by RFC 1285)
- support for the SMT v7.3 MIB (specified by RFC 1512)

Network Discovery makes FDDI connections based on the MAC address and MIB variables for each device, not based on the FDDI port.

SMT (Station Management) is an integral part of any FDDI implementation. SMT v6.2 can determine the upstream neighbor for an object. SMT v7.3 can determine both the upstream and downstream neighbors for an object.

Note: If you have a device that supports only SMT v6.2, check with the vendor or manufacturer for SMT v7.3 support. This will improve your FDDI connectivity.


Network Discovery uses the SMT instance—not the FDDI ports—when mapping FDDI objects. For example, if you have an FDDI concentrator with 8 ports, there is a single SMT instance, so Network Discovery shows only one uplink port and one downlink port for that concentrator.

If Network Discovery cannot always close the logical ring for your network, it is because:

- all your FDDI objects have no SNMP management
- all your FDDI objects have SNMP management but support SMT implementations other than v7.3 or v6.2
- at any point in the ring, you have an FDDI object with no SNMP management immediately downstream of an object that supports only SMT v6.2.

To understand this last case, you must realize that the object with no SNMP management (X) is providing no “ring information” about itself to the FDDI ring.

Figure 1-7: FDDI ring that cannot be closed



The only way for the ring to remain unbroken is for the next object upstream (Z) to be able to look back downstream and ask object X about itself. If Z supports only SMT v6.2, then Z cannot see downstream, and therefore the ring cannot be closed.

If Network Discovery is ever unable to close the ring, check for objects with no SNMP management followed by an object with support for SMT v6.2 only. This is the only likely cause of a broken ring that will not be immediately obvious.

HSRP

Hot Standby Router Protocol (HSRP) is specifically for Cisco routers. There are other, similar protocols, like the Virtual Router Redundancy Protocol (VRRP) that work with other products. This section is dedicated to HSRP, but Network Discovery works similarly with VRRP. For more information on how Network Discovery handles these protocols, contact Customer Support.

HSRP is a routing protocol that allows multiple routers to act as a single “virtual router.” If the main router fails, there are other routers available in “hot standby” mode that will immediately take over the traffic, ensuring constant network connectivity.

Note: For more detailed information on these routing protocols, contact the product manufacturer.

The basic HSRP configuration would be to have a single virtual IP (a.b.c.1) and a set of routers that will respond to this virtual IP (a.b.c.2, a.b.c.3, a.b.c.4, etc.). Only one active router will respond to the virtual IP at any given time.

There are special virtual MAC addresses that are reserved for use with HSRP, in the form of 0000C07ACxx. The same virtual HSRP MAC address can appear in any of the routers (main or standby) that are responding to the virtual IP address.

Note: The routers will all have their own individual MAC addresses as well (appearing in the device MIB).

Network Discovery may see any combination of IP and MAC addresses for the HSRP routers. It could see the real and virtual MAC for each router, depending on how the routers have been configured.

Since all the routers should respond to Network Discovery pings and SNMP queries, each physical router (active and standby) should appear on the Network Map. The device model for each router should include its real MAC address.

The virtual IP address may appear on the Network Map if the virtual IP has been seen in an ARP cache and if the routers are SNMP-managed. If the virtual IP does not appear on the Network Map, there will be an unmanaged IP+MAC device, likely connected to the active physical router by a diamond-shaped IP virtual device icon.

If the active router is SNMP-managed, the virtual IP will not appear on the Network Map.

Note: If Network Discovery finds the virtual IP in an ARP cache before it finds the active router, the virtual IP will appear on the Network Map until it is merged with the active router. Once the active router appears on the Network Map, the virtual IP will no longer appear on the Network Map.

Presenting Information

As a user, you may find it easier to remember what conventions Network Discovery uses when displaying data if you understand a little about how Network Discovery operates.

Network Map

The Network Mapper has very little to do with showing you the Network Map. The Network Mapper merely calculates the Network Map. The task of displaying the map is divided between two parts of a single process: map servers and the map client.

The map client—that is, the Network Map and other map windows—is the only part of the Network Discovery map system that you ever see.

When you click the **Network Map** button on the main Toolbar, Network Discovery performs four consecutive actions:

- begins a Health Panel session (even if you do not have the Health Panel open)
- begins a map session
- opens a map configuration
- opens a map window

These actions and concepts are separate, even though the actions are linked the first time they are performed. You will sometimes want to perform each action separately, so it helps to realize that each concept is distinct.

Map Session

The following session-based commands appear in the File menu:

- Disconnect
- Reconnect
- Close Map
- Connection Info

When you begin a map session, you start receiving data from a Network Discovery map server. You continue to receive data from a Network Discovery map server until you exit the map session, or until you disconnect from the map session.

Each map session places demands on the resources of the Peregrine appliance. For this reason, the total number of map sessions per appliance is limited.

Each account is limited to a single map session per appliance. There are frequently more accounts than there are map sessions. You may be asked to leave your map session by another account who needs a map session.

Administrator accounts can also disconnect an account from a map session.

Map Configuration

The following configuration-based commands appear in the File menu:

- New
- Open
- Open Copy of Prime
- Save
- Save As
- Save As Prime
- Revert

Any account can open or save a map configuration at any time during a map session. A map configuration file contains your settings for:

- layout, including the top object for each window
- packaging, including package icons

You can use map configuration files to view the map from different perspectives. For example, one view might show the network by geography, while another might show the network logically, by subnet.

Network Discovery automatically opens a map configuration file at the start of each map session. The first time a new account starts a map session, the session always opens with a copy of the Prime configuration. All other times, the map configuration file that Network Discovery opens depends the type of account you are using.

Table 1-3: Default configuration files and accounts

Account type	Subsequent default file
Demo	Copy of Prime
IT Employee	last opened or designated
IT Manager	last opened or designated
Administrator	last opened or designated

When you end a map session, Network Discovery takes note of what map configuration file is in use. The next time you start a map session, Network Discovery opens that file. There are two exceptions:

- You can designate a different configuration file to be opened next time using the Administration menu. See the *User Guide*.
- Demo accounts always start a map session with a configuration called “Copy of Prime”. This is so that each user of a Demo account can start fresh, unaffected by other accounts.

Demo accounts can open a saved configuration if they want to pick up where they left off.

If you end your session with “Copy of Prime”, you will get a fresh copy of Prime the next time you start a map session.

If you forget to save your map configuration before you end a map session, Network Discovery reminds you that your configuration has not been saved and offers you the chance to save it.

Each account has its own space for configuration files. You cannot overwrite or delete configurations belonging to others. For instance, you can have a configuration file named “test” and so can every other account—the files will not overwrite one another.

Administrator and IT Manager: The Prime configuration is a special default configuration customized for use in your system. This configuration is customized and maintained by Administrator and IT Manager-level accounts.

Autosave

Configuration files are saved automatically approximately every 2 minutes. This makes it possible for you to recover your configuration in the event of an abnormal occurrence, such as a power outage, or a disconnection from the map session or from the Peregrine appliance.

If a session ends abnormally, the recovery file will be opened the next time you start a map session, and you will be notified of the recovery with a dialog box: “Restored configuration from autosave”.

Even when Network Discovery loads the recovery file, you can still discard the recovery. Just re-open the configuration file that you last saved.

Note: Autosave never overwrites any configuration file that you have created. The autosave file is deleted any time you answer “No” to the question “Do you want to save the changes?”. The autosave file is also deleted every time you save a configuration.

Important: Always Save your map configuration before you “Close Map”. Do not rely on Network Discovery being able to recover the autosave file.

Prime configuration

The Prime configuration is a special configuration not associated with a particular account. Any Administrator or IT Manager account can overwrite the Prime configuration. To do so, click **File > Save As Prime**.

The Prime configuration includes:

- layout, including the top object for each window
- packaging, including package icons and titles

The default Prime configuration has end node packaging—all core devices are in the Network Map window. Layout, device priorities, and titles are all set to the default.

If you end your session with “Copy of Prime”, you will get a fresh copy of Prime the next time you start a map session.

The Prime configuration is automatically updated every night just before Reports are generated. This ensures that the package names that appear in Reports match the Network Map.

Scheduled Events

The majority of data that Network Discovery uses is constantly being collected. However, some information is collected at a set time every day, while other information is summarized once a day.

This is a list of major events, not a complete list.

Table 1-4: Major events in the 24-hour timetable

Time	System event
0005–1900 ^a	<ul style="list-style-type: none"> ■ summarize statistics for each attribute ■ perform internal backup to the Peregrine appliance's internal hard disk drive^b ■ perform external backup (if configured) to an FTP server and/or a tape drive^b ■ update Prime configuration ■ summarize events for reports ■ compile and calculate reports
0005–23:30 ^c	<ul style="list-style-type: none"> ■ check devices for deactivating and purging ■ update Health Panel reports (Exceptions, Last Seen, Open Tickets, Adds, Deletes, Moves, Changes)
0010	check evaluation license for expiry
0015	backup Prime configuration
0059	age out bridge tables of Plaintree WaveSwitch devices ^d

a If this series of events is not successfully completed, it will restart in 30 minutes and attempt to complete only the unsuccessful events from the series.

b Backups are performed only when the Peregrine appliance has been in operation for 2 hours.

c This series only begins once the previous series has finished.

d Is only done if you have provided a valid write community string for each device.

2 Terms and Concepts

CHAPTER

- *Network Terms and Concepts* on page 48 (to review terms and concepts common to network management)
- *Network Discovery Terms and Concepts* on page 53 (to learn terms and concepts unique to this product)
 - *Objects* on page 53
 - *Devices* on page 55
 - *Real devices* on page 55
 - *Virtual devices* on page 58
 - *Packages* on page 61
 - *Priority* on page 63
 - *Special input syntax* on page 64
 - *Events and Alarms* on page 65
 - *Event Thread* on page 66
 - *Panel Elements* on page 67

Network Terms and Concepts

These terms and concepts are common to networks and network management. They are not unique to Network Discovery.

Domain names

Example: website.example.com

A domain name such as “website.example.com” is easier to remember than an IP address such as “192.168.96.1”. This ease of remembering is the chief reason for the existence of domain names.

The term “domain name” and “host name” are sometimes used interchangeably. A domain name is a name in the Domain Name System (DNS) format as registered with a DNS server. A host name is purely an internal name, used by a device to refer to itself.

Address types

The two main types of numeric address are the IP address and the MAC address.

IP address

An IP address was intended to be a unique number identifying a unique device or port of a device.

When you see the term “IP address” with no qualifiers in Network Discovery, it means that either an IPv4 address or an IPv6 address is acceptable. The 32-bit address space of IPv4 addresses puts severe limits on the number of unique addresses available, and the supply is fast running out. The IPv6 128-bit address space was created to address this problem.

IPv4 address

An IPv4 address contains four sections separated by periods (or “dots”). Each section, called an octet, contains 8 bits expressed in decimal (0–255).

Example: 192.168.96.1

IPv6 address

An IPv6 address contains eight sections separated by colons. Each section contains 16 bits expressed in hexadecimal (0000–FFFF).

Example: 1234:5678:9ABC:DEF0:1234:5678:9ABC:DEF0

To make it easier to remember and type an IPv6 address, you can use a double colon (::) to indicate multiple contiguous sections of zeros. You can also omit leading zeroes. For example, you can simplify address 0123:0000:0000:0000:0004:0056:789A:BCDE to 123::4:56:789A:BCDE.

MAC address

A MAC (Media Access Control) address is a unique number identifying a unique device or port of a device.

When you see the term “MAC address”, it means a numeric MAC address.

Numeric MAC address

A MAC address contains six sections. Each section contains 8 bits expressed as a hexadecimal number (00–FF).

Sometimes the first three sections and last three sections are separated by one space; sometimes all sections are presented as one, without spaces; sometimes each section is separated by a colon or a space.

Examples: 010203 FDFEFF, 010203FDFEFF, 01:02:03:FD:FE:FF

MAC address including OUI

This type of MAC address is sometimes (inaccurately) referred to simply as an OUI. In fact, the Organization Unique Identifier (OUI) comprises the first three sections of a MAC address. If Network Discovery recognizes the numeric form of the OUI, it replaces the numbers with a short form of the organization name. This makes it easier to identify a device. If Network Discovery uses an alphabetic short form for a device’s OUI, the device is said to have a recognized OUI. Having a recognized OUI is sometimes abbreviated to “having” an OUI.

Example: DELL 59FC91

Netmask notation

Network masks, often referred to as netmasks, can usually be expressed in two formats in IPv4—either the familiar octet notation (also called dotted decimal notation) or CIDR notation.

Example of octet notation: 255.255.255.248

Example of CIDR notation: 29

The shorter CIDR notation is based on the binary equivalent of the octet notation, and refers to the numbers of contiguous 1's.

Table 2-1: Example netmask notation—octet and CIDR

255.255.255.255	11111111.11111111.11111111.11111111	32 1's
255.255.255.248	11111111.11111111.11111111.11110000	29 1's
255.255.0.0	11111111.11111111.00000000.00000000	16 1's

In IPv6, netmasks can only be written in CIDR notation.

Community strings

A community string is a kind of device-based password that controls access to the SNMP MIB of a device. A device controls its own community strings, but you must tell Network Discovery about them.

If Network Discovery is not given the correct community strings and access to devices on your network, Network Discovery will be unable to read device MIBs. Network Discovery will then assume that each device it cannot read has no SNMP management available.

With directed community strings, the device not only stores a “password”, but a list of trusted devices. If the Peregrine appliance is not on the list of trusted devices, the device will not recognize Network Discovery and Network Discovery will fail to read the device’s MIB—even though Network Discovery knows a valid string. Therefore, it is not enough to configure Network Discovery to know about your network devices. You may also need to configure your network devices to know about the Peregrine appliance if you are using access lists or directed community strings on your devices.

Bridge aging

To obtain the best results with Network Discovery, turn bridge aging on. Also, set the aging interval for 2–6 hours, although some circumstances may call for an aging interval as long as 12 or even 24 hours. (Longer aging intervals are not always possible. A common maximum aging interval is 32767 seconds, or just over 9 hours.)

Bridges, routers, and switches generally have tables in which they store the addresses of devices on the network. The tables are periodically purged and relearned in order to keep the list of devices current. The aging interval defines the frequency with which tables are purged and relearned.

When there is no table entry for the address of an incoming packet, the bridge, router, or switch must learn the location of the address. To learn the location, the device sends the incoming packet to all its own ports. (This is often referred to as “flooding” or “leakage”.) When the destination device with the corresponding address responds, the bridge, router, or switch learns the location and makes an entry in the address table.

If the table is full and a new entry must be made, the “oldest” entry is usually replaced by the new entry. Device manufacturers commonly strive to include a table large enough to hold the addresses of all active sessions, but space in a table is always finite.

Network Discovery reads the tables of bridges, routers, and switches to learn the addresses of all the connected devices. Many bridge, router, and switch vendors use a standard aging interval of 300 seconds (5 minutes), which is too short.

If the bridge aging interval is too short:

- Network Discovery may never discover devices that are connected to the network for short periods—for example, laptops.
- Network Discovery may take longer to determine connections between devices that it has discovered.
- Tables will be purged so frequently that flooding will occur regularly, using bandwidth unnecessarily.

If bridge aging is not turned on for a device, or if the bridge aging interval is too long:

- Tables will contain old addresses of devices that may have been removed from the network or devices that are broken. As a result, Network Discovery will work from an outdated and possibly confused representation of what is in your network and how it is connected.

OSI model layers

The Open Systems Interconnection (OSI) model has seven layers. Layers 2 and 3 are the most important to Network Discovery:

- Layer 2 is the Data Link layer, at which level MAC addresses are used. Bridges and some switches are layer 2 devices.
- Layer 3 is the Network layer, at which level IP addresses are used. Routers are layer 3 devices.

Some switches are both layer 2 and layer 3.

The seven layers are:

Layer number	Layer
7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

Management workstation

Any workstation or personal computer capable of running a supported web browser. There is more detail on requirements for a management workstation in the *Setup Guide*.

Network Discovery Terms and Concepts

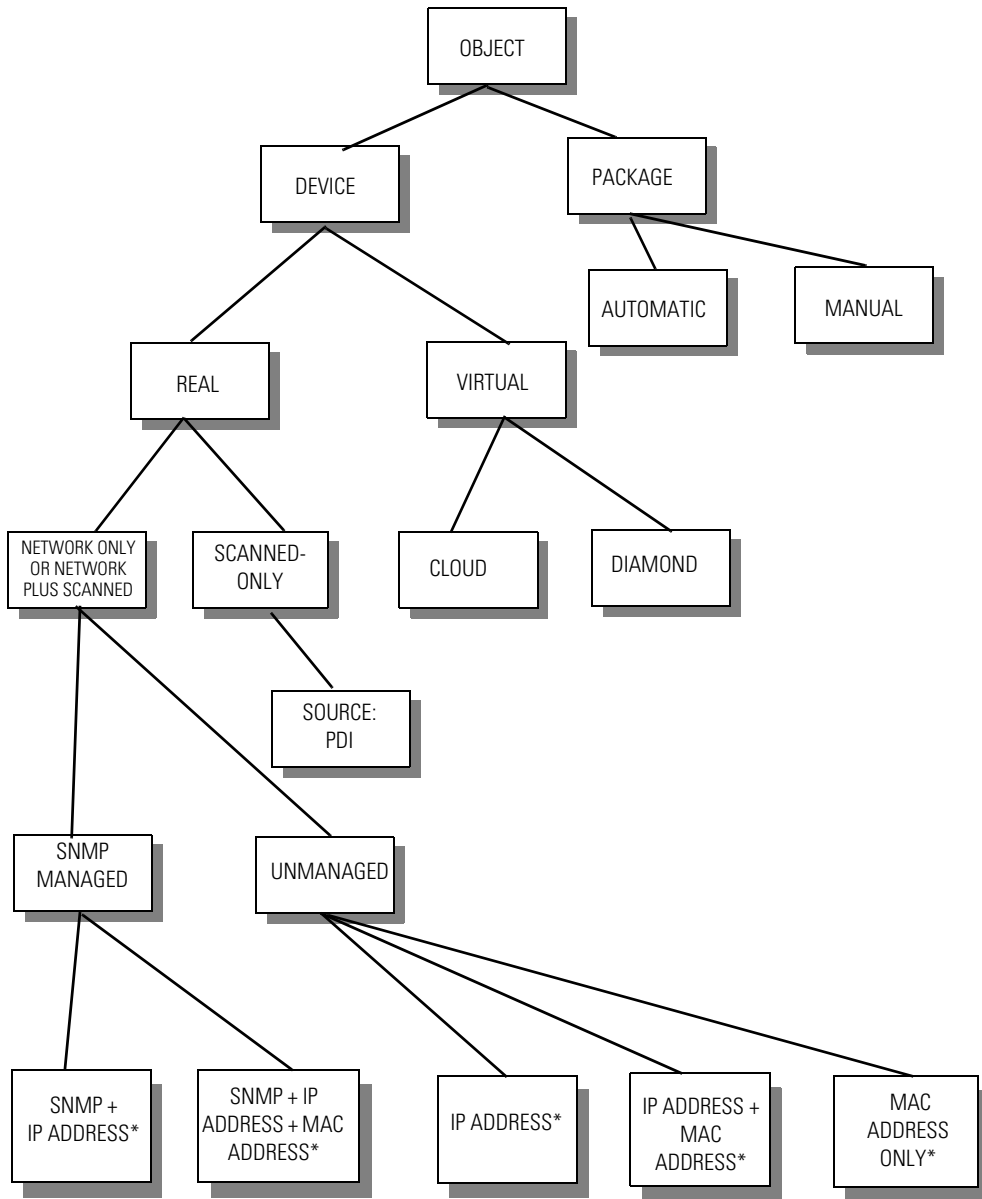
These terms and concepts are either unique to Network Discovery, or have a special meaning in this context.

Objects

A map displays icons and lines. Icons represent objects. Objects comprise devices and packages.

(Two sorts of objects are not displayed on a map—ports and attributes.)

Figure 2-1: Object type hierarchy



* May also have PDI scan file data.

Devices

Devices come in two classes, real and virtual.

Devices also come in two connectivity classes, network connectivity devices (NCDs) and end nodes. Connectivity class is discussed under *Packages* on page 61.

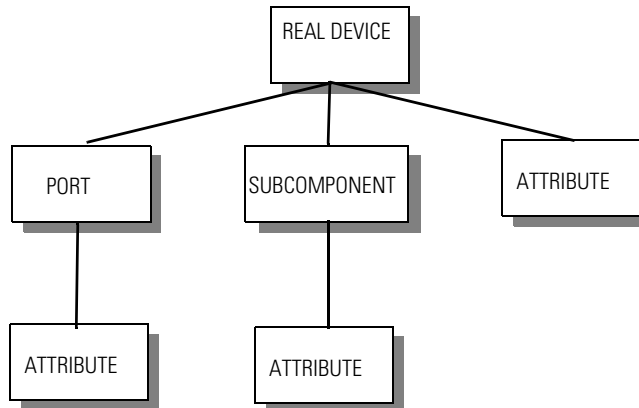
Real devices

Real devices can be found in two ways. Network devices are found automatically, in one of three ways—they respond when Network Discovery pings its IP address, their IP address appears in a bridge table, or they poll another device. Scanner-only devices are found by having their scan files sent to the Peregrine appliance.

Network devices can be SNMP-managed or unmanaged:

- SNMP-managed device: a device with SNMP information from the network (may also contain scan data)
- Unmanaged device: a device with no SNMP information from the network (may also contain scan data)
 - MAC+IP device: an unmanaged device with a MAC address and an IP address
 - IP-only device: an unmanaged device with no MAC address and no scan data
 - MAC-only device: an unmanaged device with no IP address and no scan data

Figure 2-2: Real Devices



SubComponents

SubComponents are items that a device contains such as CPU, Disk, Memory, Toner and Paper Trays. There can be many sub-components per device. Sub-Components are collected if the IP address of the device is in the Resource/Environment network configuration range.

The information about the SubComponents is merged into the Device Manager and Attribute Manager. You will only find SubComponents in the Database Schema.

Ports

A port is an interface on a device that can be controlled. On most devices this is an ethernet, token ring, FDDI, ATM, PPP or serial port. Devices such as SNMP managed power bars have each power outlet represented as a port, in addition to the network interface.

Network Discovery will try to find all the ports on the device and add them to the device model, with some exceptions:

- loopback ports are not shown
- dynamically created PPP ports are limited to 10

Attributes

An attribute is an element that can be polled on a device for data. Attributes can be broken down into the following types:

- polled by Device Poller
- polled by Resource Poller
- polled by Environment Poller
- polled on demand

The values associated with Attributes that are On Demand are never stored. When Device Manager, Port Manager, or Attribute Manager is opened, the device is polled for its value at that time. No graphs can be created on these attributes. If the device is unreachable within 3 seconds, no value will be shown. The Service Analyzer does not show these attributes.

The remaining attributes can have state, generate events, and have historical graphs with predictions.

For more details on the capabilities of each attribute, see **Help > Classifications > Supported Device/Port Attributes**. Devices that do not have a checkmark in either “Resource Poller”, “Environmental Poller”, or “Retrieved in Realtime” columns are polled or computed by the Device Poller.

Device Poller

To ensure that you have up-to-date information about your network, Network Discovery continually polls every device in your network, one by one. This is called a poll cycle. The time that it takes to complete one poll cycle is called the sampling period.

You can also have your Peregrine Systems Customer Support representative calculate the poll cycle for you.

Question Mark Icons

Question mark icons are a special case in Network Discovery. There are two types of question mark icons. The yellow question mark icon represents a device that Network Discovery can partially identify and believes to be a network connectivity device (such as a router or switch). A gray question mark icon represents an Unknown Network Connection Device (Unknown NCD).

The gray question mark icon (Unknown) represents a device for which Network Discovery has only an IPv4 address.

Question mark icons indicate a lack of information. If question mark icons remain for a few hours, a Network Discovery Administrator should check the Exceptions report to see what problems can be solved to make the network and Network Discovery work better.

The Network Discovery Administrator may decide to change the icons for Unknown devices or Unknown Network Connection Devices. If the Network Discovery Administrator changes the icon (system property), all views of the map will be updated.

Figure 2-3: Question mark icons



Virtual devices



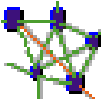

Virtual devices are principally map connectivity tools. When Network Discovery has determined that two devices are somehow connected but cannot be certain of the exact path of the connection, it inserts a virtual device between the two network devices, as a sort of placeholder. The placeholder represents either an unidentified device, or an unidentified connection between devices.

Virtual devices come in two types: the cloud and the diamond. The cloud is used to represent a real device or group of real devices that Network Discovery cannot yet identify. The diamond is used to represent connectivity that Network Discovery has not yet determined. Since a diamond is an object that does not usually represent an object, the diamond is more theoretical than a cloud.

Clouds

Clouds represent one or more devices or MAC systems that provide connectivity in the network.



There are 4 cloud icons.





Icon	Description
	The “Cloud” icon represents one or more unmanaged devices that somehow connects two or more devices. Network Discovery has determined that two or more devices are indirectly connected to each other, but cannot get any information on the device or devices that implements the connection.
	The “Radio Cloud” icon represents a wireless network connection.
	The “Carrier Network” icon represents a carrier service provider, such as a third party network that is composed entirely of unmanaged devices.
	The “Unmanaged Hub” icon represents the single hub device that connects two or more devices, but Network Discovery cannot get any information from the MIB of that hub. The hub may not have SNMP management enabled, or perhaps Network Discovery has not been configured with its community string.

Diamonds

Diamonds do not represent actual network devices; they indicate connectivity. Sometimes, Network Discovery knows that there is connectivity without being able to specify the devices.

There are 6 diamond icons.

Icon	Description
	The yellow “Unmapped IP” icon collects devices that belong on the same logical subnet, connected to a router, but for which Network Discovery has not yet determined the physical connection. The “Unmapped IP” icon has as its title the IP address of the subnet.
	The yellow “Approximate” icon is used to connect devices that Network Discovery believes to be connected to a switch or hub although Network Discovery does not know the ports to which the connections should be made. Network Discovery connects the hub or switch to an “Approximate” connection diamond and connects the diamond to the devices.

Icon	Description
	<p>The yellow “Shared Port” icon collects together the multiple devices that Network Discovery has seen attached to a single port. For example, you could see a yellow “Shared Port” icon when one port is used to test the operation of many different workstations, linked in sequence—perhaps in a test lab. Another example could be many people, each with a different laptop, sharing a common office.</p>
	<p>The orange “Logical View” icon collects all devices that are connected to other orange diamond-shaped icons, and connects them to the Peregrine appliance icon on your map. Device icons are not normally connected directly to the “Logical View” icon, rather they are connected to one of the two other orange icons described below.</p> <p>Note: The Logical View object always has at least two icons attached: one LV Unmapped IP icon and the Peregrine appliance icon. The Logical View icon will always be on the map, it cannot be removed.</p>
	<p>The orange “LV Unmapped IP” icon collects devices that belong on the same logical subnet, but for which Network Discovery has not yet determined the physical connection, nor can it determine the appropriate router for that subnet. What is the difference between yellow and orange “Unmapped IP” icons? Yellow ones are attached to the appropriate router and orange ones are attached to the “Logical View” icon.</p>
	<p>The orange “LV Unmapped” icon collects together devices which Network Discovery has no idea how to connect; not even the subnet of the devices can be determined.</p>

Packages

Packages are groups of objects. You use packages to simplify the viewing and visualization of a map. Clicking a package icon reveals the contents of the package in a separate map window.

Network Discovery automatically creates some packages, depending on the settings in **Administration > System preferences > Automatic packaging**. Users can create packages of both connectivity classes (that is, not just end nodes), or can request that Network Discovery create these types of packages.

All objects have a packaging status. They are either locked or unlocked. By default, objects are unlocked. You lock an object to prevent automatic packaging.

Effects of manual packaging

Once an object is locked, Network Discovery no longer automatically does any automatic packaging on it. Specifically, Network Discovery no longer:

- packages the object
- unpackages the object
- destroys the package containing the object

Locked objects may still be packaged manually, by the user. Locked objects may also be unlocked manually.

Sometimes manually packaging an object has unexpected effects. For example, a single device may remain unpackaged when all others around it are packaged. When packaging status is made visible, the reason for the lone unpackaged device often becomes obvious.

How an object can become locked

You can lock an object with the **Lock** command. You also lock an object as a side effect of doing any manual packaging. That is, using the commands **Package**, **Unpackage**, or **Promote** can cause objects to become locked. You cause an object to be locked, most commonly, when you move an object into or out of an end node package.

The rationale is this: It would be very confusing and frequently counterproductive if objects that you deliberately packaged were then automatically unpackaged by Network Discovery.

How to unlock an object

You unlock an object by using the **Unlock** command. You unlock all objects by using the **Unpack All** command.

In some circumstances, promoting objects—particularly into the Main Map window—will also unlock the objects.

Visibility of packaging status

Packaging status refers to whether an object is locked or unlocked as far as packaging operations are concerned. The packaging status of objects can be made visible or invisible by using the **Underline Locked Objects** command.

A blue line underneath an icon indicates that an object is locked. By default, this blue line is not shown.

Priority

In Network Discovery, each device and line on the Network Map has a priority.

Note: Packages do not have priority.

Devices and lines can have priorities 1–6. Devices and lines with priority 1 are the least important. The higher the number, the higher the priority and greater the importance.

The highest priority that Network Discovery assigns is 4. The highest priority that the user can assign is 6.

Devices

When Network Discovery identifies a device, it assigns a priority to the device. All users can change the priority of a device.

Table 2-2: Priority of objects

Default	Example	Notes
<i>more important</i>		
6 —	user-assigned only	default e-mail notification
5 —	user-assigned only	—
4 Network-connectivity objects	switches, hubs, routers, gateways, clouds	—
3 Servers	—	—
2 Common-use and auxiliary devices	printers, analyzers, UPSs	—
1 Workstations	—	—
<i>less important</i>		

By default, only devices and lines with priority 6 generate an e-mail notification of breaks.

Lines

A line inherits its priority from the devices it connects. The device with the lower priority determines the priority for the line. You can only change the priority of a line by changing the priority of the devices at its endpoints.

Special input syntax

You can create SNMP system variables for system name, system location, and system contact that will appear as hyperlinked within the Device Manager.

Network Discovery gives you a shorthand for entering URLs: “<URL: >”. You must include an appropriate prefix with the URL—such as “http://” or “mailto:”—or the link will not work.

Limits

Acceptable prefixes: mailto: | news: | http:// | https:// | telnet:// | ftp:// | gopher://

Table 2-3: URL syntax in system variables

	What you type	Results in Device Manager
CORRECT	sysadmin@example.com	sysadmin@example.com
CORRECT	<URL:mailto:sysadmin@example.com>	sysadmin@example.com
INCORRECT	<URL:sysadmin@example.com>	sysadmin@example.com ^a
INCORRECT	sysadmin@example.com (System Admin)	sysadmin@example.com (System Admin) ^a
CORRECT	<URL:mailto:sysadmin@example.com> (System Admin)	sysadmin@example.com (System Admin)
CORRECT	<URL:http://www.example.com>	http://www.example.com
CORRECT	http://www.example.com ^b	http://www.example.com
INCORRECT	<URL:www.example.com>	www.example.com ^a

^a No text will be hyperlinked.

^b Restricted to “http://” only. Not available to other prefixes.


Events and Alarms

To see what alarms are raised by report data or attribute data, see [Help > Classifications > Alarms](#).

There are 5 kinds of event types:

Event Type	Alarms Generated	Example
Thresholdable Attribute	critical, major, minor, info	Utilization
Add	info	new device added to network
Delete	info	the device is hidden or has been deactivated
Move	info	connectivity change, physically moving a device
Property Change	info	changing a device icon, priority

Figure 2-4: Alarm hierarchy







Thresholdable Attributes can cause alarms (critical, major, minor, info). These are thresholds that the user can change. These events are reflected in the Health Panel the next time a value is collected for that attribute (the database is refreshed once a minute). To see what attributes can be “alarmed” see the “Can be Thresholded” column in the table at [Help > Classifications > Supported Device/Port Attributes](#). To change the thresholds, from the Network Map or Health Panel, click [Edit > Alarm Thresholds](#).

Reports (ex, MTTR, MTBF) are accumulated data, and summarize data for the past 24 hours. You can change the default “time period” for these alarms in [Administration > System preferences > MTTR and MTBF](#).

Summaries (ex, Adds, Deletes) are summaries of events. You can change the default “time period” for these alarms in **Administration > System Preferences > Adds/Deletes/Changes/Moves**.

Here is a list of the alarm indicators visible in the Health Panel and elsewhere in the user interface:

Alarm Type	Indicator
n/a (not an alarm state, this indicates that the attribute is not being monitored)	blank
OK	—
Info	 green asterisk
Minor Alarm	 gold triangle
Major Alarm	 orange diamond
Critical Alarm	 red square

Event Thread

The event thread is a set of related events for a particular attribute on a device during a specific time period.

When opening a ticket in ServiceCenter, Network Discovery maintains an event thread to keep all the events for that attribute together in one ticket. Once you set the timeout (default is 2 hours), Network Discovery will wait that length of time before opening a new ticket for that device attribute.

For example, port 1 on a router may be experiencing intermittent packet loss alarms. If you have an event filter set up to open a ticket in ServiceCenter for that event, all the packet loss alarm state changes (ok > minor, minor > major, major > ok, ok > major, etc.) will be listed in the same ServiceCenter ticket for easy tracking.

However, once the alarm changes to an OK state, Network Discovery will wait a specified time to see if more events occur on this attribute before closing the event thread. You can change the timeout period at **Administration > System preferences > Network devices > Event thread timeout**.

After the Event Thread Timeout has expired, if there are more events on that port, a new event thread will be created, and a new ticket will open in ServiceCenter.

Panel Elements

Certain elements are common to all Device Manager, Port Manager, Line Manager, or Attribute Manager panels:

- When data in a table has a gray background, the data shown is considered stale, because it was obtained before the beginning of your selected time period. (In some cases, data may be shown in parentheses rather than with a gray background.) To change the time before data is considered stale, see the section on Account Properties in the *User Guide*.
- A blank space indicates that data is not available for a device or port.
- The final line on each panel is the date and time that the panel was refreshed. (This refers to rendering the panel itself, not when the data shown in the panel was last read.) This date can be useful for when you print a panel. To change the format of this date, see the section on Account Properties in the *User Guide*.

Banner

The banner that appears at the top of all Device Manager, Port Manager, Line Manager, or Attribute Manager panels consists of four elements.

Table 2-4: Device Manager banner

Element	Example	Notes
Attribute Name	Total Breaks	This only appears in the Attribute Manager
Device title and IP address	website.example.com / 192.168.96.1	<ul style="list-style-type: none"> ■ this does not appear in Line Manager banners ■ see <i>Device Title</i> on page 68 ■ if the device title is the IP address, the IP address is shown once ■ if there is no IP address, only the device title is shown
Manager name	Device Manager	—
System name of Peregrine appliance	ExampleCorp	see the <i>Setup Guide</i>
Web browser name	Netscape Internet Explorer	—

Device Title

The title displayed in the banner of Device Manager, Port Manager, Line Manager, or Attribute Manager (and in some panels of those managers) will be the first available of:

- system-level assigned name
- *virtual devices only*: Network Discovery generated name
- a device title chosen by the Network Discovery Administrator in **Administration > System preferences > Display preferences**. The Network Discovery Administrator can choose one or several of the following and choose their order too:
 - Asset Tag
 - BIOS Asset Tag
 - NetBIOS Name (scan)
 - Last Name
 - First name
 - Device-specific title
 - Domain name

- NetBIOS name (network)
- Operating system
- Family
- Model
- Network function
- System description
- System name
- System location
- System contact
- IPv6 address
- IPv4 address
- MAC address including OUI
- MAC address (all-numeric)

Note: Only Administrator or IT Manager accounts can change device titles (**Object > Device Properties**). The device titles are global. To determine the default title for a device, see the Diagnosis panel on the Device Manager.

3 Recorded Events

CHAPTER

The Health Panel summarizes connectivity changes to the network. Each device should contribute only a single alarm. (If there is more than one alarm per device or per port, they will be displayed in the Device Manager or Port Manager.)

To see what alarms are raised by report data or attribute data, see [Help > Classifications > Alarms](#).

Line Breaks

Identifies lines that are broken. A line is broken when its status is down and the line break is not due to devices at either end being broken.

Utilization

Identifies lines that are used heavily—that is, that have a lot of traffic. Describes the amount of traffic on the line as a percentage of capacity.

Available only to devices with byte counters or frame counters. For media with variable length packets, utilization is calculated by directly reading bytes counts from every interface. For media with fixed length packets (for example, ATM cells), utilization is derived from frame counts.

On the Port Manager Statistics panel, you will see three Utilization graphs:

- Utilization In

- Utilization Out
- Line Utilization

Utilization In and Out

In and Out utilization for Half Duplex and Full Duplex ports are calculated using a common algorithm.

In both cases, Network Discovery reads the byte/frame count and port speed for each interface from the device MIB. Network Discovery calculates the bit count per second by computing the difference between two successive byte/frame counts and dividing this by the poll period. The port utilization is calculated by dividing the bit count per second by the port speed.

For variable: In Utilization is calculated by taking the InByte count per port, converting this value to the Bit per second value, dividing by the Port speed, and multiplying by 100 to get a percentage. i.e. $[(\text{InByte} \times 8\text{Bits}) / \text{Port Speed}] \times 100$.

For fixed length: In Utilization is calculated by taking the InFrame count per port, converting this value to the Bit per second value, multiplied by the number of bytes per frame, dividing by the Port speed, and multiplying by 100 to get a percentage. i.e. $[(\text{InFrame} \times 8\text{Bits} \times \text{BytesperFrame}) / \text{Port Speed}] \times 100$.

Out Utilization is similarly calculated by taking the OutByte (or OutFrame) count per port.

Line Utilization

Line utilization is determined using different methods, depending on whether the ports are Half Duplex or Full Duplex.

Half Duplex

Line Utilization is equal to the total of InUtilization and OutUtilization.

Unlike Full Duplex, the IN and OUT on a Half Duplex line occur on the same channel, and the value for one limits the value of the other. Given a 100 Mbps Half Duplex line, if the IN traffic is 90 Mbps, the OUT traffic can be at most 10 Mbps. In all cases, the sum of both values cannot exceed 100 Mbps. As a result, it makes sense to use the sum of the IN and OUT to calculate the line utilization. If this total reaches a threshold, an alarm is triggered.

Full Duplex

Line Utilization for a full duplex port is the maximum value of either the In or Out Utilization.

Network Discovery calculates line utilization this way because the primary objective is to inform the user whenever the line utilization becomes a concern. It becomes a concern whenever it approaches the maximum utilization (or reaches a threshold and triggers an alarm). For a Full Duplex line, the In and Out occur on two separate channels, and the value for one does not limit the value of the other. The values are independent of each other, and their values are limited only by the line speed. Given a 100 Mbps Full Duplex line, for example, the In traffic may be 90 Mbps while the Out traffic may be 30 Mbps, but each cannot exceed 100 Mbps.

Delay

Delay is calculated for lines, however the delay values are associated with the ports at both ends of the line. The delay values for the two ports will be the same.

Note: It is possible to alarm only one end of the line if the interface types (and alarm thresholds) are different.

Delay alarms identify lines with long queuing delays. The response time, measured in milliseconds, is a portion of the time taken by a device to respond to a ping. This includes the following:

- the time a packet waits in the router queue before being transmitted
- the time to process the packet at the other end after being received

Note: The calculation does not include the delay across the link.

The values are calculated based on round trip delays (ping responses) from the devices at both ends of the line. The round trip delay for a device is the mean of the ping responses received during a poll cycle.

The delay value is half the difference between the mean round trip delays of the two devices minus the lowest value recorded for this line.

Note: Network Discovery receives one ping per poll cycle for the scheduled request and will receive some additional pings for fast break requests.

Jitter

Jitter is calculated as the difference between the delay value for the current poll cycle and the delay value for the previous poll cycle. It is measured in milliseconds, and is bi-directional.

Note: Jitter does not appear in the Health Panel, but it does appear in the Service Analyzer “summary graphs” panel, and in the Device Manager State panel.

Collisions

Identifies the number of collisions per second detected on every line in the network with values above the thresholds.

Broadcasts

Identifies the number of broadcasts per second detected on every line in the network with values above the thresholds. Broadcasts are part of normal network operation, but large numbers of broadcasts must be investigated and the cause rectified.

Network Discovery creates broadcast alarms only for devices that are broadcast sources such as servers, routers, and workstations. Other devices such as switches and FDDI devices are not seen as a source of broadcast but as devices that forward the broadcast inputted to them. For these devices, Network Discovery creates minor broadcast alarms. The minor broadcast alarms for these devices will not be seen in the event log due to the potentially large volume of events.

Note: When you are looking for the source of a broadcast, set the priority on your map window to 1. This will help you determine which workstation might be causing the broadcasts.

Errors

Identifies the number of errors per second detected on every line in the network with values above the thresholds.

Exactly what errors are reported depends on the MIBs of the devices at either end of the line. Not all devices detect all errors.

Frame Relay

Attributes relating to your frame relay connections: Data Delivery Ratio, Frame Delivery Ratio, Discard Eligibility In, Discard Eligibility Out, BECN, FECN.

Device Breaks

Device Breaks are discussed in *Detecting broken devices* on page 22.

Packet Loss

Identifies managed core network devices (for example, routers and switches) that are dropping frames. Describes the percentage of frames that are dropped by each managed device. Calculated on unicast data, inbound and outbound, for all ports of the device. Percentage is calculated over the past 5 sampling periods.

The criteria for Network Discovery to calculate “Packet Loss” is as follows:

- The device must be router or a switch.
- All ports of the device must have unicast counter attributes.
- All ports must have valid port statistics (i.e. if a port does not have a valid counter, such as if the statistics filtered out by the mapping module for any reason, the “Packet Loss” attribute will not be calculated).
- If a port has status bit up but has frames and unicast counters “0”, the “Packet Loss” attribute will not be calculated. Even if it has out statistics different than “0”.

Note: If a port has in or out frames 0 but in or out unicast = 0, the “Packet Loss” attribute will not be calculated.

Packet loss is calculated on a per device basis, more specifically if there are backplane connections to your devices the packet loss statistics might be slightly “off”.

Disk Utilization

The percentage used on each disk partition.

CPU Utilization

The percentage of the cycles used by each processor.

Load Average

A calculation of how much of the CPU is being used.

Example: If the CPU is all being used, it's maxed out for scheduled processes, then the load average will be 1. If there are twice as many processes scheduled than the CPU can handle, the load average will be 2.

Memory Utilization

The memory used by all running processes.

Backplane Utilization

For some Cisco devices, taken from a MIB variable.

Printer

Paper count.

UPS

UPS battery capacity and UPS battery time remaining.

Port MTTR

Mean time to repair (MTTR) identifies ports that take a long time to repair. A running average of the number of hours broken against how many times it was broken.

Example: A port has failed twice. The first time, it was broken for 4 hours. The second time, it was broken for 8 hours. The MTTR for this device is $(4 + 8) / 2 = 6$ hours.

Port MTBF

Mean time between failures identifies ports that fail frequently. A running average of the number of days in operation measured against the number of times a ports has failed.

Example: A port has been in operation for 100 days and Network Discovery has seen it fail twice. The MTBF for this device is 50 days.

Port Add/Deletes

Identifies ports recently added to or deleted from a device on the map. (An added port may or may not be recently discovered.)

Limits Does not include ports on virtual devices

Port Moves

Identifies if the connection to a device has changed from one port to another.

Example: A workstation is attached to a switch at port 2. You detach the workstation from port 2, and reattach it to port 8. That would create a Port Move event. The change is recorded on the workstation, not the switch. The workstation has not necessarily changed location, only the switch port to which it connects.

Port Changes

Line Alarm type, interface rate, duplex.

Device MTTR

Mean time to repair (MTTR) identifies devices that take a long time to repair. A running average of the number of hours broken against how many times it was broken.

Example: A device has failed twice. The first time, it was broken for 4 hours. The second time, it was broken for 8 hours. The MTTR for this device is $(4 + 8) / 2 = 6$ hours.

Device MTBF

Mean time between failures identifies devices that fail frequently. A running average of the number of days in operation measured against the number of times a device has failed.

Example: A device has been in operation for 100 days and Network Discovery has seen it fail twice. The MTBF for this device is 50 days.

Device Adds/Deletes

Identifies devices recently added to or deleted from the map. (An added device may or may not be recently discovered.)

Device Moves

Identifies devices recently moved, or had a change in connection to the network.

Moves are not reported for devices that have been added recently. If a device appears in Adds, it will not appear in Moves.

The time to detect changes in connectivity depends on the sampling period for the network.

Example: A workstation is attached to a switch. You detach the workstation, and reattach it to another switch. That would create a Device Move event, because Network Discovery thinks the workstation has changed location. The change is recorded on the workstation, not the switch. The workstation has not necessarily changed location, only the switch to which it connects.

Limits

- Does not include virtual devices as anchor devices—connections to a virtual device are not considered relevant
- Does not include cases where the current and previous connections are to virtual devices
- Does not include cases where the current and previous connections are the same, or are probably the same (as in the cases where the only one connection is known).

Device Changes

Changing icon, priority, title, or tag of a device.

Exceptions

Devices with exceptions. See [Help > Classifications > Exceptions](#).

Not Recently Seen

There are two types of “not recently seen” events:

- Network Not Recently Seen
- Scan Not Recently Seen

“Network Not Recently Seen” devices are those with which Network Discovery has lost contact and which may soon disappear from the Network Map.

Note: Once Network Discovery has not had contact with a device for a period greater than the threshold (by default, 6 hours), it will be displayed with a green ring. Once the “not seen” period has exceeded 24 hours, the device will also be appear faded.

“Scan Not Recently Seen” devices are devices for which Network Discovery has not received an updated scan file (by default, 4 weeks and 2 days).

Note: You can change these defaults at **Administration > System preferences > Report time periods**.

- Limits**
- Does not include virtual devices

Open Tickets

Tickets opened in ServiceCenter.

4 Device Manager

CHAPTER

- To explore icon buttons in the toolbar menus:
 - *Toolbar* on page 83
 - *Statistics* on page 102
 - *Events* on page 105
- To interpret data in the Device Manager window, see *Panel Elements* on page 67.

Important: Many of the panels in the Device Manager feature data in table form. Not all tables will look the same for all devices, because the tables will only show data that is available for that device.

Introduction

The Device Manager provides you with detailed information about a device, in several panels.

Ways of opening

- From a map window, double-click a device icon.
- From a map window, right-click the device icon and click **Open Device**.
- From a map window, click a device icon. From the **Object** menu, click **Open Device**.
- From a map window, the Health Panel, Alarms Viewer, or Events Browser: From the **Tools** menu, click Find (or Ctrl-F). Enter a device address or title, then click **Find**.
- From the Alarms Viewer or Events Browser, double-click on an event, or right-click on the device icon, title, or IP address.
- From a Service Analyzer path diagram, double-click a device icon.
- From the Toolbar, click the **Find** button. Enter a device address or title, then click **Find**.
- Click a hyperlinked device title. (Hyperlinked devices appear in Manager panels, in reports, and in **Network Map Sessions** status.)

Default panel

- *initial*: State
- *subsequent*: from Account Properties (see the *User Guide*.)

Toolbar

Availability of buttons in the Device Manager toolbar.

Table 4-1: Available toolbar buttons

Icon	Button name	No IP address	Not in database	Not on Network Map or Unknown	Virtual device	Demo or IT Employee user
	Configuration	✓	✓	✓	✓	✓
	State	✓	—	✓	—	✓
	Reports	✓	—	✓	—	✓
	Diagnosis	✓	✓	✓	✓	✓
Buttons on the Diagnosis Panel						
	Diagnostic Information	✓	—	✓	✓	✓
	IP Ping	—	✓	✓	—	✓
	Traceroute	—	✓	✓	—	✓
	SNMP Ping	—	✓	✓	—	✓
	Listener Ping	—	—	✓	—	✓
	DNS Query	—	✓	✓	—	✓
	Statistics	✓	—	✓	—	✓
	Ports	✓	—	✓	✓	✓
	Events	✓	—	✓	—	✓
	Locate	✓	—	—	✓	✓

Table 4-1: Available toolbar buttons (Continued)

Icon	Button name	No IP address	Not in database	Not on Network Map or Unknown	Virtual device	Demo or IT Employee user
	Service Analyzer	✓	—	—	✓	✓
	Manage	—	✓	✓	—	✓
	Browse MIB	—	✓	✓	—	✓
	View Scan Data	✓	—	—	—	✓
	Web	—	✓	✓	—	✓
	Telnet	—	✓	✓	—	✓
	Update Model [Administrator or IT Manager]	—	✓	✓	—	—
	Device Visibility [Administrator or IT Manager]	✓	—	—	—	—
	Properties	✓	✓	—	✓	✓
	Refresh	✓	✓	✓	✓	✓
	Print	✓	✓	✓	✓	✓
	Text	✓	✓	✓	✓	✓
	Close	✓	✓	✓	✓	✓

Configuration

Identifies a device and presents an overview of the device's identity, position, and status.

Limits This panel is blank if the device is not in the Network Discovery database.

Details This panel is divided into the following principal sections:

- Heading
- Identity table (real devices only)
- Device structure (Serial number and description)
- Address table (real devices only)

Heading

The heading also appears in the State, Reports, and Diagnosis panels (when available).

Table 4-2: Heading

Element	Notes	Type
Icon	for a complete list, see Help > Classifications > Device Types/Package Types	all
Descriptive prefix	for example, "SNMP-managed device"	
Device type	for a complete list, see Help > Device Types	all
Device tag	see the <i>User Guide</i> .	real
Virtual device map title	see <i>Virtual devices</i> on page 58	virtual
No. of ports	the number Network Discovery uses for the port may not match the physical port	all
No. of connections	includes user-assigned connections	all
Object title	first title available; see <i>Device Title</i> on page 68	all
Address	IP address; does not appear if identical to object title	real
Priority	see <i>Priority</i> on page 63	all
Virtual device number	created by Network Discovery or by an Administrator or IT Manager account	virtual

Identity

The information in this table can come from three sources: the Network Discovery Rulebase and the SNMP MIB of the object and, if you are using them, Peregrine Desktop Inventory (PDI). For more information on using Network Discovery with PDI, see *Using Network Discovery with Desktop Inventory and Desktop Administration*.

The Rulebase determines the device's operating system, application, device family, and model. It determines as many of these as are available.

Some of the information collected from the SNMP MIB has been set by the device manufacturer; other information can be customized. For information on how to enter MIB data so that Network Discovery interprets it as a link, see *Special input syntax* on page 64.

More elements of identity appear for the Peregrine Appliance than for any other device.

Note: All these elements are optional.

Table 4-3: Elements of Identity table in Configuration panel

Data	Example	Creator	Administrator or IT Manager
Package [*]	Main Map	Network Discovery/ account	—
Family	Cisco 2600 Series Modular Access Routers	Network Discovery Rulebase	—
Family current manufacturer	Cisco Systems Inc	Network Discovery Rulebase	—
Model	Cisco 2621XM Modular Access router	Network Discovery Rulebase	—
Model current manufacturer	Cisco Systems Inc	Network Discovery Rulebase	—
Model historical manufacturer [†]	Cisco Systems Inc	Network Discovery Rulebase	—
Operating system	Cisco IOS Version 12.2 (8) T5	Network Discovery Rulebase	—
Operating system current manufacturer	Cisco Systems Inc	Network Discovery Rulebase	—

Table 4-3: Elements of Identity table in Configuration panel (Continued)

Data	Example	Creator	Administrator or IT Manager
Operating system historical manufacturer	Cisco Systems Inc	Network Discovery Rulebase	—
Network Function	—	Network Discovery Rulebase	—
Network Function current manufacturer	—	Network Discovery Rulebase	—
Network Function historical manufacturer	—	Network Discovery Rulebase	—
Operating system	Linux	Desktop Inventory	—
Service pack	—	Desktop Inventory	—
NetBIOS name (network)	—	device owner	—
NetBIOS workgroup	MARKETING	device owner	—
rulebase extra info	—	Network Discovery Rulebase	—
Device-specific title	—	scripts	—
System OID	.1.3.6.1.4.1.295.5.1.1.2	manufacturer	—
System OID manufacturer	PlainTree Systems Inc	Network Discovery Rulebase	—
System description	Ethernet Switch	manufacturer	—
System contact	test@example.com	device owner	set [‡] link
System name	ws1216-2	device owner	set [‡] link
System location	Server Room	device owner	set [‡] link
Read community string	public	device owner	view
Write community string	n/a	device owner	view
Asset tag	78LL996	Desktop Inventory	—
BIOS asset tag	—	Desktop Inventory	—
BIOS product name	eserver xSeries 330 -[867441X]-	Desktop Inventory	—
BIOS product manufacturer	IBM	Desktop Inventory	—

Table 4-3: Elements of Identity table in Configuration panel (Continued)

Data	Example	Creator	Administrator or IT Manager
BIOS serial number	78LL996	Desktop Inventory	—
BIOS chassis	—	Desktop Inventory	—
CPU	Pentium III 1133 MHz (Genuine Intel)	Desktop Inventory	—
NetBIOS name (scan)**††	DUPONT	device owner	—
Last name	DUPONT	Desktop Inventory	—
first name	MARIE	Desktop Inventory	—
Memory (MB)	1024	Desktop Inventory	—
Windows/NIS domain	—	Desktop Inventory	—

*This is optional if you have not opened a map configuration since this object was discovered.

†Appears only when different from the current manufacturer.

‡A shortcut to the MIB Browser.

**On Windows workstations, frequently the same as the system name.

††NetBIOS data is blank unless the device has an IP address.

Package:

- Displays the position of a device within the packaging of the Network Map. Click on a hyperlink to open a corresponding map window.
- If you have a map open, this row reflects the packaging of your current configuration. If you open the Device Manager and then make packaging changes that affect the device, click the **Refresh** button to have this row updated.
- If you do not have a map open, this row reflects the packaging of the configuration you were using in your previous map session.
- If you have never had a map open, this row does not appear.
- If the device has been added to the network since the last time you saved your configuration, this row does not appear.

An admin or IT Manager user will also see a read and a write community string for a device. These values are taken from the list of community strings; however:

- strings from the list appear here only if they are valid.

- only a single valid string appears here even if the list has multiple valid strings for this device.
- the read string that appears here is the string that Network Discovery is currently using to poll the device.

Device Structure

Provides information on the serial number of the chassis and modules in a device. You will see the following information about each module:

- Type:
 - backplane
 - container
 - misc
 - other
 - powerSupply
 - stack
 - chassis
 - fan
 - module
 - port
 - sensor
 - unknown
- Name
- Hardware
- Firmware
- Software

- SerialNumber
- Description

Figure 4-1: Device Structure example

Type	Name	Hardware	Firmware	Software	SerialNumber	Description
[-]	Switch System (Cisco Systems WS-C6509 9 slot switch)	3.0		7.6(3)	TSC07190058	WS-C6509
[-]	slot 1 (WS-C6509 9 slot switch chassis slot)					
[-]	1 (1000BaseX Supervisor 2 port WS-X6K-S2U-MSFC2 Rev. 4.2)	4.2	7.1(1)	7.6(3)	SAL0717CDW8	WS-X6K-S2U-MSFC2
	env temp (Module Intake Temp Sensor)					
	env temp (Module Exhaust Temp Sensor)					
	env temp (Module Device 1 Temp Sensor)					
	env temp (Module Device 2 Temp Sensor)					
[-]	L3 Switching Engine II Container					
[-]	L3 sub-module (L3 Switching Engine II)	3.3			SAL0718CERH	WS-F6K-PFC2
	env temp (L3 SE Intake Temp Sensor)					
	env temp (L3 SE Exhaust Temp Sensor)					
	env temp (L3 SE device1 Temp Sensor)					
	env temp (L3 SE device2 Temp Sensor)					
?	CPU of supervisor					
[-]	Container of Router Switch Feature Card					
[-]	15 (Router Switch feature Card)	2.5	12.1(19) E1	12.1(19) E1	SAL0718CE1U	WS-F6K-MSFC2
	env temp (RSFC Intake Temp Sensor)					
	env temp (RSFC Temp Sensor)					
	env temp (RSFC device1 Temp Sensor)					
	env temp (RSFC device2 Temp Sensor)					
[-]	slot 2 (WS-C6509 9 slot switch chassis slot)					
[-]	Container of Power Supply Group					
[-]	Container of Power Supply					
	2500 watt AC supply	1.0			ART070800DC	WS-CAC-2500W
	Fan Sensor (Power Supply Fan Sensor)					
[-]	Container of Power Supply					

Address

Provides information about the IP addresses and/or MAC addresses of a device's ports. The information comes from the Network Explorer.

This table has hyperlinks for all the ports with addresses. If a port does not have an address, it does not appear. To open a Port Manager, click a port hyperlink. Each table row contains either:

- a MAC address, an OUI abbreviation (if known), and a manufacturer (if known)
- an IP address, a netmask (if known), and a domain name (if known)

A special port of “Device” is used:

- for the IP or MAC address that Network Discovery identifies as the primary IP or MAC address for the device
- when Network Discovery does not know which port an IP or MAC address is associated with

Table 4-4: Ports table

Data	Notes
Port index	port number and description
MAC/IP address	—
OUI/Netmask	netmask in octet notation
Manufacturer/Domain name	usually hyperlinked to an external web site

The address table is particularly useful:

- When the device is
 - a router
 - a device with multiple IP addresses and domain name aliases (such as a web server)
- When you want to know a device’s domain name (and domain name is not included in the list of **Device Title Preferences**)

VLANs

VLANs are software-defined broadcast domains, created by your System Administrator at the switch. If your device has any VLANs configured, you will see them in the Device Manager.

Figure 4-2: VLAN example

Virtual LANs:

VLAN ID	Description
1	default
100	ComputerRoom
200	VLAN0200
1002	fdci-default
1003	token-ring-default
1004	fdinet-default
1005	trnet-default

State

Displays current values for attributes. Displays alarms even when the device priority is less than the minimum priority for a configuration (unlike in map windows).

Limits This panel is not available if the object is not in the Network Discovery database.

Details **Heading**

The heading also appears in the Configuration, Reports, and Diagnosis panels (when available). See Table 4-2 on page 85.

Attributes

See [Help > Classifications > Supported device/port attributes](#).

The displayed attributes will differ depending on whether or not the device is managed, the type of device, and if resource management is configured.

Note: The Peregrine appliance itself will have the most attributes because you will see attributes for the appliance, and for the network as a whole.

Note: The Peregrine appliance itself will not have any Breaks or Total Breaks data.

Information on attributes is collected from the network regularly (during each poll cycle). The information is the latest available, and so may be different each time you view it. Network Discovery only shows you attributes that are relevant.

When data in a table has a gray background, the data shown is considered stale, because it was obtained before the beginning of your selected time period. To change the time before data is considered stale, see the *User Guide*.

A blank space indicates that data is not available for a device or port.

Note: If a device had a partitioned disk, each partition will appear as a separate “Disk” attribute. You can open an Attribute Manager for each partition. Each partition may have a different disk serial number (assigned by the device OS).

Reports

Displays current values for report (MTTR, MTBF) and summary historical data (utilization, availability, and so on). Displays alarm signals even when the device priority is less than the minimum priority for a configuration (unlike in map windows).

This 'report' data is historical information. You can use the data on this panel in conjunction with the State panel to look for problem trends in your device. For example, you can see an alarm in the State panel for the device CPU, you can check the 'reports panel' to see if there was a problem yesterday, or over the past week or month.

Limits This panel is not available if the object is not in the Network Discovery database.

Heading

The heading also appears in the Configuration, State, and Diagnosis panels (when available). See Table 4-2 on page 85.

State

Lists 'report' data like MTTR, MTBF, adds, deletes, and changes, lets the user know if any of these are in an alarm state (info, minor, major, critical).

If there are any exceptions for the device, they are noted in this table. For a complete list of exceptions in your network, see the Health Panel and Alarms Viewer.

Table 4-5: State

Data	Notes
Report name	Exceptions, MTTR, MTBF, Device Adds, Device Deletes, Device Moves, Device Changes, Not Recently Seen
State	OK, Info, Minor, Major, Critical
Value	For exceptions: <ul style="list-style-type: none"> ■ description ■ effect ■ action

Note: Exceptions cannot always be reported for a device.

For a complete listing of Network Discovery exceptions, see [Help > Classifications > Exceptions](#).

Metrics

Table 4-6: Metrics

Data	Notes
Name of metric	Device attributes with historical data, like memory utilization, CPU utilization, availability, and so on.
Time Period	When these events were recorded.
Mean Value	The average of the average data points for this time period.
Mean Peak Value	The average of the peak data points for this time period.
Peak Value	The highest data point in this time period.
Peak Time	What time the attribute hit its peak.
Mean Min Value	The average of the minimum data points in this time period.
Min Value	The lowest data point in this time period.
Min Time	What time the attribute hit its lowest point.

Events

Table 4-7: Events

Data	Notes
Name	Events on this device.
Time Period	When these events were recorded.
State	OK, Info, Minor, Major, Critical
Count	How many events in this category.
Duration	How long the event was active.

Diagnosis

Displays information about the current state of the device that can be helpful in diagnosing problems. Has buttons that give you access to diagnostic tools. Opens with a configuration panel.

Diagnostic Information

Heading

The heading also appears in the Configuration, State, and Reports panels (when available). See Table 4-2 on page 85.

Beneath the heading, this panel is divided into four main sections:

- Main Diagnosis
- Network Configuration
- Property Assignment

Main Diagnosis

The main table indicates the data flow for this device—when the device was first and most recently seen by various parts of Network Discovery—plus the current values for several parameters.

Table 4-8: Main Diagnosis table

Data	Output	Notes
First discovered	elapsed time* / absolute date & time	Reset if database is cleared.
Added to map	elapsed time / absolute date & time	Resets if the device is deactivated, but then returns to the map.
Last seen	elapsed time / absolute date & time	in ping or poll by Network Discovery
Last changed	elapsed time / absolute date & time	the last time a connection to this device changed
Network model last updated	elapsed time / absolute date & time	the last time the model changed; determines whether or not the model has been for this device

Table 4-8: Main Diagnosis table (Continued)

Data	Output	Notes
Device checked for existence	elapsed time / absolute date & time	the last time a device was pinged for discovery; should be “n/a” or a time before “Model last updated”
Last deactivated or hidden	elapsed time / absolute date & time	the last time a device was deactivated
Mean break diagnosis time	minutes for alarms	Mean break diagnosis time is approximate. Diagnosing a break may take longer, if communication with the device is unreliable.
ARP tables seen	elapsed time / absolute date & time	the last time the ARP tables were seen, to the nearest 30 minutes.
Port ARP tables seen	elapsed time / absolute date & time	the last time the ARP tables were seen by this port, to the nearest 30 minutes.
Port Bridge tables seen	elapsed time / absolute date & time	the last time the Bridge tables were seen by this port, to the nearest 30 minutes.
Port Source address capture seen	elapsed time / absolute date & time	the last time the Source address capture was seen by this port, to the nearest 30 minutes.
Port Radio link seen	elapsed time / absolute date & time	the last time a Radio link was seen by this port, to the nearest 30 minutes.
Port Bus link seen	elapsed time / absolute date & time	the last time a Bus link was seen by this port, to the nearest 30 minutes.
Port carrier link seen	elapsed time / absolute date & time	the last time a carrier link was seen by this port, to the nearest 30 minutes.
Port link training seen	elapsed time / absolute date & time	the last time link training was seen by this port, to the nearest 30 minutes.
Port detailed link training seen	elapsed time / absolute date & time	the last time detailed link training was seen by this port, to the nearest 30 minutes.
Device modeler interval	either “Default as set in Network Configuration.” or time (in days, hours, minutes, seconds)	If custom, is shown here.

Table 4-8: Main Diagnosis table (Continued)

Data	Output	Notes
Mean device modeler update run time	elapsed time	the mean length of time it takes to update the model for this device the previous 4 times
Recent device modeler update run times	elapsed time	the length of time it took to update the model for this device the previous 4 times
Rulebase ID	—	an internal number

*Elapsed time is reported in at least two of the following units: weeks, days, hours, minutes, and seconds. As elapsed time increases, the finer units of measure are not reported.

Network Configuration

The Network Configuration table shows what parameters have been set up for the range in which the device resides and what their values are. It shows the Network Properties (**Administration > Network Configuration**).

Options Network Properties include:

- Allow devices
- Actively ping
- Net BIOS query
- Resource manage
- Force ARP table read
- Accumulate IP addresses
- Allow IP addresses
- Allow ICMP and SNMP
- Service manage
- Device Modeler interval

Network Properties may be set to On, Off or Inherit

Community read strings and write strings include names of Community Property Groups or Sets that have been applied to the range in which the device occurs.

- Bandwidth
- Frequency

- Scanner run schedule
- Scanner upgrade schedule
- Scan file download schedule
- Listener communication ports

Property Assignment

The property assignment table helps you to determine the rules Network Discovery has used to assign the title, icon, and priority to the device.

Certain object properties, such as device titles, cascade from system preferences as set by an Administrator or IT Manager account, provided the current map configuration has not had a property assigned by the user. User-assigned properties always take precedence, even over the cascade. One property that does not cascade under any circumstances is priority.

Table 4-9: Property Assignment table

Parameter	Notes
Default title	—
Custom title	takes precedence over default
Actual title	the title as it appears for your account
Default icon	—
Custom icon	the icon assigned by an Administrator or IT Manager account in Object > Device Properties .
Actual icon	the icon as it appears for your account
Default priority	—
Custom priority	for information only; never affects active configuration; useful if you receive e-mail or a page from Network Discovery
Actual priority	the priority as it appears for your account
Default tag	—
Custom tag	the tag assigned by an Administrator or IT Manager account in Object > Device Properties
Actual tag	the tag as it appears for your account

If no value has been assigned, an asterisk (*) appears in this table, indicating that the value for the property comes from the previous row of the table.

IP Ping

Pings the device to see if it responds, and how quickly. The IP address pinged is the address identified by Network Discovery as the primary IP—see *State* on page 92.

- Limits**
- 1–20 pings
 - The device must have an IP address. If not, this button is dimmed.

Default 5 pings

Traceroute

Displays the path that data takes to get from the Peregrine appliance to the selected device by listing the gateway devices associated with each hop of the journey. The device identifier is often the host name, where available, but can also be the IP address. Each device title is hyperlinked to a Device Manager.

Traceroute also displays the amount of time each hop took. This time is the round trip in milliseconds. Traceroute includes two retry hops for each try, so the times for all three hops are shown.

Traceroute helps you to understand where on the network problems are occurring. It is often used after *IP Ping* on page 99 has been used to confirm the existence of a device.

Note: The path displayed by traceroute is at OSI layer 3 and may not match the connectivity on the Network Map or in the *Service Analyzer* on page 106, which map at layer 2.

- When to use it**
- If you suspect that you are losing packets due to a large hop count. In a TCP/IP network, where data are transmitted in packets, the header for a packet tracks the hop count. If the hop count grows too large, the packet is discarded.
 - If you are trying to determine the point along the path where traffic is slowing down or getting lost altogether.
 - If you are trying to determine the precise path taken—not so much to solve a problem as for general information.

Limits The device must have an IP address. If not, this button is dimmed.

Output Results of an asterisk for the device and for all three times (i.e. the result * * *) indicates that data is not available for that hop of the journey, and usually indicates a trouble spot along the path. The following table explains codes you may see when you attempt a Traceroute.

Table 4-10: Traceroute special results

Chars.	Meaning
*	no response within a 3-second timeout interval
!	ttl <= 1 *
!H	host is unreachable
!N	network is unreachable
!P	protocol is unreachable
!S	source route failed
!F	fragmentation needed
!X	communication is prohibited administratively
!V	a host precedence violation has occurred
!C	precedence cutoff is in effect

*The ttl value is supposed to start at 1 and increase by 1 until the host is reached.

Related To see the OSI layer 2 path between any two devices, see also *Service Analyzer* on page 106.

SNMP Ping

Queries the device for basic SNMP information and displays this information. The IP address pinged is the address identified by Network Discovery as the primary IP—see *State* on page 92.

Limits The device must have an IP address. If not, this button is dimmed.

Default

- Demo, IT Employee, IT Manager: “public”
- Administrator: the read community string for the device as defined in **Administration > Network configuration > Community Property Groups**.

Listener Ping

Makes a connection to the listener agent running on the device to see if:

- the port number you have is correct (you can set this in **Administration > System preferences > Listener communication**)
- the security keys are correct

Limits The device must be in the Network Discovery database.

DNS Query

Sends a host query to the domain name server and displays a table that highlights configuration errors. A highlighted line indicates that the next line in the progression is missing.

The highlighted configuration errors are:

- a pointer (PTR) without an IP address (A or AAAA)
- duplicate pointer (PTR) records for the same IP address (A or AAAA)
- a mail exchanger (MX) directed to a canonical name (CNAME)
- a canonical name (CNAME) directed to anything that doesn't exist

Highlighted information also includes an explanation in the “Exceptions” column. You will see one of the following explanations:

- Duplicate
- Target does not exist
- n/a

If no information in the table is highlighted, Network Discovery did not detect any problems with the DNS configuration of the device.

Limits If the device does not have an IP address, the button is dimmed.

Procedural alerts If Network Discovery displays the message “Non-existent domain”, it means that the device has not been assigned a domain name.

Statistics

Provides a second toolbar with which to view or export detailed historical statistics of the device. The statistics may be viewed in graph or table form, and may be exported in Comma Separated Value (CSV) form.

Not all statistics are available for all devices. Only available statistics appear in the list box. Statistics are a subset of Attributes (*Attributes* on page 92).

Statistics for the past two to three days are averaged every five minutes, statistics for the past 33 days are averaged every hour, and statistics for the past 365 days are averaged every day.

Note: The y-axis maximum pull-down list only applies when graphing data. It allows you to change the topmost data point on the y-axis. Some of the options may have no effect on the display depending on the actual data. The highest data point is always shown, regardless of your selection.

Options **Graph**

Whenever a graph contains multiple averages, the data is adjusted to the lowest common denominator and the data points used are indicated on the graph. For example, a graph of the past seven days contains only one-hour data points.

At the beginning of data collection, Network Discovery shows whatever data it has at five-minute intervals—even if you want to see statistics for long periods of time. This can be changed if you select a different option from the granularity pull-down list.

Gray portions of the graph indicate that data was not available for a period. Darker gray is used for unavailable data plotted in dark blue, lighter gray for unavailable data plotted in light blue. Also shown on the graph are horizontal lines representing alarm thresholds (depending on the option you have selected in the pull-down list).

You can change the graph by changing the selection in any of the pull-down lists. You can change the statistic, the interval, the maximum levels, and the granularity of data displayed.

Note: Every account can have its own default settings for the statistic, interval, maximum levels, and granularity. See **Administration > Account administration > Account properties**.

Table

The table shows a tabular view of the statistics.

Export

Creates a Comma Separated Value (CSV) file of the data. Popular spreadsheets such as Microsoft Excel can import CSV files if you want to sort or graph the statistics in a way that is beyond the capabilities of Network Discovery.

Statistics

Available statistics depend on the device model.

Notes on some statistics:

- *Total Breaks:* This statistic reports cumulative values.
- *Downtime:* This statistic reports cumulative values.
- *Total In Bytes:* Some devices do not report traffic in bytes, so this menu item may not appear. For such devices, try Total In Frames.
- *Total Errors:* Includes only errors in that the device stores in its MIB. Network Discovery does not control which errors are stored, and cannot report errors that the device does not chose to store.
- *Total Collisions:* Only available for Ethernet half duplex. Also restricted to devices that report collisions in the dot3StatsEntry object of their MIB.

Interval

Past 2 hours | Past 4 hours | Past 6 hours | Past 12 hours | Past 24 hours | Past 48 hours | Past 7 days | Past 30 days | Past+Next 30 days | Past 90 days | Past+Next 90 days | Past 180 days | Past+Next 180 days | Past 365 days | Past+Next 365 days | Today | This week | This month | This quarter | This half | This year | Last week | Last month | Last quarter | Last half | Last year

Note: Some of these intervals include future predictions (for example, Past+Next 30 days). You can use this feature to see how Network Discovery predicts your statistics to change. For more information on this and the Forecast feature, see the *User Guide*.

Maximum

These attributes show the Max value of the vertical axis. If you have set alarm thresholds (Utilization, Delay, etc.), the Threshold Max is shown with a red line.

Selection	Description
Threshold Max	The vertical axis will show the maximum threshold value, if the thresholds have been configured.
Data Max	The vertical axis will show the maximum value of the data gathered.
Data Max with Thresholds	Same as “Data Max.” However, if thresholds have been configured, they will be displayed as horizontal lines.
Attribute Max	Used for graphs such as Utilization, Availability, or Disk Space so that the Vertical axis is adjusted according to the Max value of these Attributes. For example, the maximum level for Utilization is 100%.
Attribute Max with Thresholds	Same as “Attribute Max.” However, if thresholds have been configured, they will be displayed as horizontal lines.
Forecast Max	For attributes that can have a forecast calculated, the vertical axis will display the maximum value of the data gathered or the maximum forecasted value, whichever is greater.
Forecast Max with Thresholds	Same as “Forecast Max.” However, if thresholds have been configured, they will be displayed as horizontal lines.

Note: The y-axis maximum drop down list only applies when graphing data. It allows you to change the topmost data point on the y-axis. Some of the options may have no effect on the display depending on the actual data. The highest data point is always shown, regardless of your selection.

Granularity

Default granularity | 5-minute granularity | Hourly granularity | Daily granularity

Ports

Lists ports for this device and summarizes the information available for them. Displays 24 ports at a time (by default, you can change this in **Administration > Account administration > Account properties**). There are also Previous and Next buttons and an All button that shows all ports in a single panel.

Note: Ports do not always support all the attributes listed on the Device Manager Ports panel. If an attribute is not supported, the table column will be blank.

You can create different views for this panel, so you can concentrate on the data most important to you. See **Administration > System Preferences > Device Manager ports display preferences**. Read the inline help for definitions of all the preference properties.

The Configuration panel and Ports panel are the most commonly used ways of starting the Port Manager.

Events

Opens the Events Browser with this device in context.

For detailed information, see the *User Guide*.

Locate

Highlights in a map window the location of the device.

If a map session was not open already, one will open now. Within the map window, the device you are locating has a yellow circle around it.

If the window containing the device was already open, that window becomes the front-most window, and the window scrolls so that the highlighted icon can be seen.

Service Analyzer

Opens the Service Analyzer query window with the current device already selected as Device 1, to allow the user to view the state of the path between this device and any other device on the Network Map. For more information, see the *User Guide*.

Manage

Launches an element manager of your choice.

Limits

The URL or application must be defined in **Administration > System preferences > Element management**. If not, this button is dimmed.

Note: *for Aggregator*—This button is never dimmed when you are viewing a remote appliance from the Aggregator appliance.

Procedural alerts

Note: *for Aggregator*—Definitions for **Element management** are supplied from the Aggregator appliance, not the remote appliance.

Browse MIB

Opens the MIB Browser to allow the user to view the device's SNMP MIB.

The MIB Browser also allows an expert user with an Administrator or IT Manager account to manipulate the device on a more detailed level.

- Limits**
- The device must have an IP address. If not, this button is dimmed.
 - The device must support basic SNMP functionality.

View Scan Data

Opens an Asset Viewer window to show information about the device collected by Peregrine Desktop Inventory (PDI).

Also opens an Asset Viewer window to show information about the Peregrine appliance collected by Peregrine Desktop Inventory scanner. For more information, see your Desktop Inventory documentation.

- Limits** If there is no scan data, the View Scan data button is dimmed.

Note: The Peregrine appliance always has scan data.

Web

Attempts to open a web browser window for the device.

- When to use it** If the device supports web-based management or other web services.

- Limits**
- The device must have an IP address. If not, this button is dimmed.
 - The device must support HTTP sessions. (Network Discovery does not check before attempting a connection.)

- Related** To control how HTTP connections are made, see the section on Proxy Services in the *User Guide*.

Note: for Aggregator—See also **Administration > Remote appliance administration > Remote appliance properties**.

Telnet

Attempts to open a Telnet session. Many network devices provide Telnet as a means to set up and configure the device.

- Limits**
- The device must have an IP address. If not, this button is dimmed.
 - The device must support Telnet sessions. (Network Discovery does not check before attempting a connection.)

Related To control how Telnet connections are made, see the section on Proxy Services in the *User Guide*.

Note: for Aggregator—See also [Administration > Remote appliance administration > Remote appliance properties](#).

Update Model [*Administrator or IT Manager*]

For Network Discovery users, there are two options available in this panel: Network and Rulebase. If you are also using Peregrine Desktop Inventory (PDI), there are additional options available.

At the top of the panel, there is a pull-down list so you can select the command you want to perform:

Command	Explanation
Query Network	Puts the device at the top of the device modeler's queue, and automatically runs a Rulebase update as well
Upgrade Scanner	Only available with Peregrine Desktop Inventory. Transfers the relevant scanner executable and configuration files to the device, then execute the scanner and finally transfer the resulting scan file to the appliance.
Configure Scanner	Only available with Peregrine Desktop Inventory. Executes the scanner and then transfers the resulting scan file to the appliance.
Run Scanner	Only available with Peregrine Desktop Inventory. Requests the immediate Desktop Inventory scan of this device.
Retrieve Scan File	Only available with Peregrine Desktop Inventory. Transfers the result of the latest scan from the device to the appliance.
Enrich XML	Only available with Peregrine Desktop Inventory. Requests immediate enrichment of the scan file associated with this device. For more information on the XML Enricher, see <i>Using Network Discovery with Desktop Inventory and Desktop Administration</i> .
Run Rulebase	Allows you to only re-check the Network Discovery rulebase for this device.

The Network command tries all valid community strings for this device, in the order specified in **Administration > Network Configuration > Community Property Groups**. The command does not begin with the currently active community string, it begins with the first string in the list of community strings.

Note: If doing an Update Model on a new device, there may be a delay of as much as 1–2 hours before the device appears on the Network Map.

Network Discovery checks several conditions before updating a device model.

Table 4-11: Conditions for updating device model

State	Message
major alarm	IP address is not in scope
major alarm	no read community strings have been specified
minor alarm	no write community strings have been specified
minor alarm	IP address is not in scope for resource management
info	current discovery process
info	list of read community strings to be tried
info	list of write community strings to be tried
info	update interval
info	mean time to update model

- When to use it**
- When you've made changes to a device that affect connectivity—for example, when you've changed cards in a router.
 - When you've made changes to a device's community strings.

Limits The device must have an IP address. If not, this button is dimmed.

Related To determine when a model was last updated, see *Diagnosis* on page 95, under “Network model last updated”.

Device Visibility [Administrator or IT Manager]

You can activate, deactivate, hide, or purge devices on this panel.

For information on how to activate, deactivate, hide, or purge devices, see the *User Guide*.

Properties

For information, requirements, and procedures on how to use this feature, see the *User Guide*.

Refresh

Refreshes the contents of the panel.

When used with IP Ping and SNMP Ping panels, uses the last entered value instead of prompting you for a value.

- Limits** Does not re-read the data in the panel from the network. Re-reads the data only from the Network Discovery database.
- Does not affect Properties or Locate panels, or any of the interactive session windows (Browse MIB, Web, Telnet).

Print

Sends the contents of the panel to a printer attached to the management workstation.

Text

Displays the contents of the Device Manager as text that can be copied and pasted.

Note: If the Statistics Graph panel is displayed, the Text button displays a text version of the Table, since there can be no text version of a graph.

Note: May cause the panel to be refreshed with new data.

Procedural alert

To return to non-text mode, click the currently depressed button again.

Close

Closes the window and exits the Device Manager.

5 Port Manager

CHAPTER

- To explore icon buttons in the toolbar menus, see:
 - *Toolbar* on page 115
 - *Statistics* on page 123
 - *Events* on page 124
- To interpret data in the Port Manager window, see *Panel Elements* on page 67.
- To select a different port for the same device, use the port list box—see *Port number* on page 129.

Important: Many of the panels in the Port Manager feature data in table form. Not all tables will look the same for all ports, because the tables will only show data that is available for that port.

Introduction

Provides you with detailed information about a device's ports, in one of several panels.

Administrator or IT Manager: Also enables you to change the way Network Discovery perceives a connection.

Note: The Port Manager enables you to change only Network Discovery's perception of a connection. The Port Manager does not change the physical connection.

Important: The Port Manager options that are only for Administrator or IT Manager accounts require you to make changes to all accounts and all map configurations.

Ways of opening

- Click a port hyperlink from:
 - the Device Manager State or Ports panel
 - the Line Manager
 - a report
- Right-click on a line on the Network Map
- Right-click a port number on the Events Browser or Alarms Viewer

Default panel

- *initial:* State
- *subsequent:* from **Administration > Account administration > Account properties**

Toolbar

Availability of buttons in the Port Manager toolbar.

Table 5-1: Available toolbar buttons


















Icon	Button name	IT Employee or Demo user
	Configuration	✓
	State	✓
	Reports	✓
	Diagnosis	✓
	Statistics	✓
	Events	✓
	Locate	✓
	Purge Port [Administrator or IT Manager]	—
	Create Connection [Administrator or IT Manager]	—
	Break Connection [Administrator or IT Manager]	—
	Port Properties [Administrator or IT Manager] <ul style="list-style-type: none"> ■ Interface Rate ■ Interface Type ■ Line Alarm Type ■ Duplex Mode 	—
	Refresh	✓
	Print	✓
	Text	✓

Table 5-1: Available toolbar buttons

Icon	Button name	IT Employee or Demo user
	Close	
	Port number	

There are also buttons for Refresh, Print, Text, and Close. These are the same in all manager windows.

At the right side of the Port Manager toolbar is a pull-down list of ports. Use this list to toggle between different ports on a device.

Configuration

Identifies a port and presents an overview of the port's identity and connections.

Details This panel is divided into three main sections:

- Heading
- Connectivity table
- Identity table
- VLAN table

Heading

The heading also appears in the *State* and *Diagnosis* panels (when available).

Table 5-2: Heading elements

Element	Notes	Type
Device Icon	for a complete list, see Help > Classifications > Device Types/Package Types	all
Device type	for a complete list, see Help > Classifications > Device Types	all
Device tag	see the <i>User Guide</i>	real
No. of ports	the number Network Discovery uses for the port may not match the physical port	all

Table 5-2: Heading elements (Continued)

Element	Notes	Type
No. of connections	includes user-assigned connections	all
Object title	first title available; see <i>Device Title</i> on page 68	all
Port no./ description	number of port / description of port	all
Device priority	see <i>Priority</i> on page 63	all

Connectivity

Most information in this table comes from the Network Discovery Rulebase.

Table 5-3: Elements of Identity table in Configuration panel

Data	Example	Notes
Connected to	the selected port is connected to another device on this port	hyperlinked to Device Manager, Port Manager, and Line Manager
Description	100Base-TX Port	from device manufacturer
Interface type	Ethernet CSMA/CD	from device MIB/Network Discovery Rulebase
Alarm type	Ethernet 100 HD	from device MIB/Network Discovery Rulebase
Interface rate	100 Mbits/sec.	from device MIB/Network Discovery Rulebase
Duplex	Half	Half Full

Identity

This table identifies the port and the manufacturer of the device:

- MAC address of the port
- OUI of the device (alphabetic abbreviation of the device manufacturer)
- Manufacturer of the device, hyperlinked to manufacturer's web site
- IP address of the port
- Netmask of the port
- Domain Name of the port

VLAN data

VLANs are software-defined broadcast domains, created by your System Administrator at the switch. By showing VLAN information in Network Discovery, the Administrator can see how the devices in that virtual domain are functioning.

State

Heading

The heading also appears in the *State* and *Diagnosis* panels (when available). See Table 5-2 on page 116.

Table

There is a list of supported device and port attributes in **Help > Classifications > Supported Device/Port Attributes**.

These values are collected from the network regularly and may change each time they are viewed. The values shown are the latest information available.

When data in a table has a gray background, the data shown is considered stale, because it was obtained before the beginning of your selected time period. To change the time before data is considered stale, see the section on Account Properties in the *User Guide*.

The left-most column is for attributes that are associated with the alarm categories on the Health Panel. The alarm icon in this column tells you at a glance if the port is experiencing problems. The column also includes Operational Status.

A blank space indicates that data is not available for a device or port.

Unlike in map windows, the State panel displays alarm signals even when the priority for the port's device is less than the minimum priority for a configuration.

Reports

Displays current values for report (MTTR, MTBF) and summary historical data (utilization, availability, and so on). Displays alarm signals even when the device priority is less than the minimum priority for a configuration (unlike in map windows).

This 'report' data is historical information. You can use the data on this panel in conjunction with the State panel to look for problem trends in your device. For example, you can see an alarm in the State panel for the device CPU, you can check the 'reports panel' to see if there was a problem yesterday, or over the past week or month.

Limits This panel is not available if the object is not in the Network Discovery database.

Heading

The heading also appears in the Configuration, and State panels (when available). See Table 5-2 on page 116.

State

Lists 'report' data like MTTR, MTBF, adds, deletes, and changes, lets the user know if any of these are in an alarm state (info, minor, major, critical).

Table 5-4: State

Data	Notes
Report name	Exceptions [*] , MTTR, MTBF, Device Adds, Device Deletes, Device Moves, Device Changes, Not Recently Seen
State	OK, Info, Minor, Major, Critical
Value	For exceptions: <ul style="list-style-type: none"> ■ description ■ effect ■ action

^{*}For a full list of possible exceptions, see **Help > Classifications > Exceptions**.

Metrics

Table 5-5: Metrics

Data	Notes
Name of attribute	Device attributes with historical data, like memory utilization, CPU utilization, availability, and so on.
Time Period	When these events were recorded.
Mean Value	The average of the average data points for this time period.
Mean Peak Value	The average of the peak data points for this time period.
Peak Value	The highest data point in this time period.
Peak Time	What time the attribute hit its peak.
Mean Min Value	The average of the minimum data points in this time period.
Min Value	The lowest data point in this time period.
Min Time	What time the attribute hit its lowest point.

Events

Table 5-6: Events

Data	Notes
Name	Events on this device.
Time Period	When these events were recorded.
State	OK, Info, Minor, Major, Critical
Count	How many events in this category.
Duration	How long the event has been active.

Diagnosis

Displays information about the current state of the port that can be helpful in diagnosing problems with Network Discovery.

This panel is divided into three main sections:

- Heading
- Main table
- Property Assignment

Heading

The heading also appears in the *State* and *Diagnosis* panels (when available). See Table 5-2 on page 116.

Main table

The main table indicates the data flow for this port—when the device was first and most recently seen by various parts of Network Discovery—plus the current values for several parameters.

Table 5-7: Main Diagnosis table

Data	Output	Notes
First discovered	elapsed time* / absolute date & time	resets if database is cleared
Added to map	elapsed time / absolute date & time	resets if the device is deactivated/hidden, but returns to the map
Last moved	elapsed time / absolute date & time	the last time a connection to this device changed
Network model last updated	elapsed time / absolute date & time	the last time the model changed; determines whether or not the model has been for this device
Scanner model last updated	elapsed time / absolute date & time	—
Last deactivated or hidden	elapsed time / absolute date & time	the last time a device was deactivated or hidden
Mean break diagnosis time	time for alarms	—

Table 5-7: Main Diagnosis table (Continued)

Data	Output	Notes
ARP tables seen	elapsed time / absolute date & time	—
Bridge tables seen	elapsed time / absolute date & time	—
Source address capture seen	elapsed time / absolute date & time	—
Radio link seen	elapsed time / absolute date & time	—
Bus link seen	elapsed time / absolute date & time	—
Carrier link seen	elapsed time / absolute date & time	—
Link training seen	elapsed time / absolute date & time	—
Detailed link training seen	elapsed time / absolute date & time	—
Connection method	<ul style="list-style-type: none"> ■ bridge tables ■ source address capture ■ traffic ■ link training ■ logical subnet ■ approximate; see <i>Virtual devices</i> on page 58 ■ user-defined; see <i>Create Connection [Administrator or IT Manager]</i> ■ unknown 	—
Previously connected to	<ul style="list-style-type: none"> ■ none ■ device (real or virtual), hyperlinked to Device Manager ■ device and port, hyperlinked to the Device Manager and Port Manager 	if blank, the device is no longer in the database, or the connection has never changed

* As elapsed time increases, the finer units of measure are not reported.

Property Assignment

The property assignment table helps you to determine how Network Discovery sees the port.

Table 5-8: Property Assignment table

Parameter	Notes
Default Interface Rate	as generated automatically by Network Discovery
Custom Interface Rate	as set by the Administrator or IT Manager account
Actual Interface Rate	the interface rate as it appears for your account
Default Interface Type	as generated automatically by Network Discovery
Custom Interface Type	as set by the Administrator or IT Manager account
Actual Interface Type	the interface type as it appears for your account
Default Line Alarm Type	as generated automatically by Network Discovery
Custom Line Alarm Type	as set by the Administrator or IT Manager account
Actual Line Alarm Type	the line alarm type as it appears for your account
Default Duplex Mode	as generated automatically by Network Discovery
Custom Duplex Mode	as set by the Administrator or IT Manager account
Actual Duplex Mode	the duplex mode as it appears for your account

Statistics

Provides a second toolbar with which to view or export detailed historical statistics for the port. The statistics may be viewed in graph or table form, and may be exported in Comma Separated Value (CSV) form.

Inbound and outbound data is displayed for several statistics. Average values and peak values are available for several statistics.

Not all statistics are available for all ports. Only available statistics appear in the list box. Statistics are a subset of Attributes (see **Help > Classifications > Supported Device/Port Attributes**).

The Statistics panel in the Port Manager works the same as in the Device Manager. See the description in *Statistics* on page 102 in the *Device Manager* chapter.

Events

Opens the Events Browser with this device and port in context.

For detailed information, see the *User Guide*.

Locate

Highlights in a map window the location of the device to which this port is attached.

If a map session was not open already, one will open now. Within the map window, the device you are locating has a yellow circle around it.

If the window containing the device was already open, that window becomes the front-most window, and the window scrolls so that the highlighted icon can be seen.

Purge Port *[Administrator or IT Manager]*

Removes the port from the device's model as created by Network Discovery.

Warning: This action cannot be undone.

Important: You are *not* making a physical change to the port. If you purge a port but the port is still operational, the port will be rediscovered and will reappear.

When to use it When a port has been removed from the network and you wish to update Network Discovery's representation of the device.

- Effects**
- Deletes the statistical history associated with the port. This in turn affects the graphs and reports for this port.
 - Deletes the events associated with the port from the event log.
 - breaks the connection on the port

- Related**
- To break a connection between ports, see *Break Connection [Administrator or IT Manager]* on page 126.
 - To purge an attribute, see *Purge Attribute [Administrator or IT Manager]* on page 143.
 - To purge a device, see *Device Visibility [Administrator or IT Manager]* on page 110 and the *User Guide*.

Create Connection *[Administrator or IT Manager]*

Forces a new connection. You can create a connection to a real device or to a virtual device.

Create connection does not change the physical connection on the device; they change only how Network Discovery represents the connection on the Network Map.

Tip: You can create a virtual device by creating a connection to a nonexistent virtual device.

Connections changes take effect at the end of the current sampling period.

Effects

Important: Do not create a connection to another real device except as a last resort. If you force a connection prematurely, you could slow Network Discovery down or even make it impossible for Network Discovery to correctly connect to your network. Never use forcing a connection as a quick fix.

Warning: Do not create a connection without consulting your Network Discovery Customer Support representative. If you force a connection, Network Discovery may not be able to correctly connect your network devices.

Note: An exception: you may create connections to ports external to your network (for example, to your ISP) to ensure that the line break is reported.

Forcing a new connection first breaks any existing connection.

When to use it When Network Discovery has made incorrect assumptions about connectivity.

Break Connection *[Administrator or IT Manager]*

Breaks an existing connection.

Break connection does not change the physical connection on the device; they change only how Network Discovery represents the connection on the Network Map.

When to use it When Network Discovery has made incorrect assumptions about connectivity.

Related See also *Chapter 6, Line Manager*.

Port Properties *[Administrator or IT Manager]*

Interface Rate

Sets rate for a line interface.

When to use it

- When you want to set a custom line speed
- When Network Discovery has set the wrong line speed.

Limits 0 bit/sec.–1 Tbit/sec.

Effects Interface rate affects utilization statistics.

Interface Type

Sets the media type used for the line.

When to use it

- When Network Discovery does not recognize the type of interface for the line.
- When Network Discovery has set the wrong interface type for the line.

Limits Network Discovery assigns a default duplex to each interface type. Full duplex and half duplex are listed separately.

Related To change the duplex, see *Duplex Mode* on page 127.

Line Alarm Type

Sets the line alarm type for the connection. The line alarm type is normally associated with the interface type, but may be changed independently.

Table 5-9: Abbreviations used in alarm types

Abbreviation	Expanded form
ATM	asynchronous transfer mode
DSL	digital subscriber line
FD	full duplex
FDDI	fiber distributed data interface
HD	half duplex
LAN	local area network
SPN	switched packet network

When to use it When the default line alarm type associated with the interface is inappropriate.

Effects ■ Collision thresholds are valid only when the alarm type is for an Ethernet half duplex connection (alarm types 3, 5, and 7)

Duplex Mode

Sets the duplex to full or half. Full duplex allows for two-way communication over a line; half duplex permits only one-way communication.

When to use it When Network Discovery has set the wrong duplex.

Limits Full | Half

Effects Duplex affects utilization statistics.

Refresh

Refreshes the contents of the panel.

Limits Does not re-read the data in the panel from the network. Re-reads the data only from the Network Discovery database.

Print

Sends the contents of the panel to a printer attached to the management workstation.

Text

Displays the contents of the Port Manager as text that can be copied and pasted.

Note: If the Statistics **Graph** panel is displayed, the Text button displays a text version of the Table, since there can be no text version of a graph.

Note: Not available to **Interface Rate**, **Interface Type**, **Alarm Type**, **Duplex Mode**, or **Connection**.

Note: May cause the panel to be refreshed with new data.

Procedural alert To return to non-text mode, click the currently depressed button again.

Close

Closes the window and exits the Port Manager.

Port number

Allows you to select from the valid port numbers for this device.

Note: The number Network Discovery uses for the port may not match the physical port.

On your Cisco devices, the Cisco naming convention is displayed (for example, “Tu1” for Tunnel 1, or “Fa0/1” for Fast Ethernet 1).

6 Line Manager

CHAPTER

Network Discovery has two different Line Managers:

- *Single Line Manager* on page 131
- *Multiple Line Manager* on page 135

Single Line Manager

The single Line Manager provides you with detailed information about the two devices on either side of a connection.

The line can be between:

- the ports on two known devices
- a port on a known device and an unknown port on a device
- unknown ports on two devices

Ways of opening

- From a map window, double-click a line.
- From a map window, point cursor at a line, and right-click.
- From a Service Analyzer path diagram, click on a line.
- Click a [line] hyperlink. Line hyperlinks appear in Manager panels.
- a report





Effects

Selecting a line on the map opens either a single line window or a multiple line window.

Toolbar

The Line Manager has two panels. Its principal panel is About.

Table 6-1: Available toolbar buttons

Icon	Button name	Virtual device
	About	
	Break Connection [Administrator or IT Manager]	

About

The About panel shows two columns. In each column are a device and the relevant port for that device. If the Line Manager was opened by the Device Manager or Port Manager, the left column contains the device that was in context for the other Manager.

This panel is divided into two main sections:

- Heading
- State and Attribute Name (with Unit and Value)

Heading

Displays enough data to allow you to identify any device with which the port is associated.

Table 6-2: Heading elements

Element	Notes	Type
Icon	for a complete list, see Help > Classifications > Device Types/Package Types	all
Object type	for a complete list, see Help > Classifications > Device Types	all
Device tag	see the <i>User Guide</i>	real
Port (optional)	the number that Network Discovery uses for the port may not match the physical port	real
Connections	includes user-assigned connections	all

Table 6-2: Heading elements (Continued)

Element	Notes	Type
Object title	first title available	all
Port title	port index and port description; hyperlink to Port Manager	all
Priority	see <i>Priority</i> on page 63	all
[locate] hyperlink	hyperlink to map window	all
Cloud number	created by Network Discovery or by an Administrator or IT Manager account user	virtual
Port properties (not labelled)	from the MIB, includes: <ul style="list-style-type: none"> ■ interface type ■ interface speed ■ duplex ■ alarm type (from the Network Discovery Rulebase) 	real

Underneath the heading is a single line that explains how the connection was made. This is identical to the “Connection method” row in the Port Manager panel for *Diagnosis* on page 121.

State

The left-most column for each device tells you at a glance if either device or the port of either device is experiencing any problems for any Attribute.

Unlike in map windows, the Line Manager displays alarm signals even when the priority for the device (and its ports) is less than the minimum priority for a configuration.

A blank space indicates that data is not available for a device or port.

Attribute name

Displays the current statistics for any attribute available.

These values are refreshed at the end of each poll cycle and may change each time they are viewed.

The metrics tables presented here is similar to the ones that would appear in the Device Manager and Port Manager’s State panel (see page 92 and page 118) for each device port. The only difference here is the absence of the “update time column.”

Note: It is important to understand that metrics for the two device ports will probably not match exactly. This is because the statistics for each device are not collected at the same time. Although there is rarely an exact match, the two sets of statistics should however be approximately equal, with in/out values reversed.

Break Connection *[Administrator or IT Manager]*

Breaks an existing connection.

When to use it When Network Discovery has made incorrect assumptions about connectivity.

Related See also the Port Manager *Break Connection [Administrator or IT Manager]* on page 126.

Refresh

Refreshes the contents of the panel.

Limits Does not re-read the data in the panel from the network. Re-reads the data only from the Network Discovery database.

Print

Sends the contents of the panel to a printer attached to the management workstation.

Text

Displays the contents of the Line Manager as text that can be copied and pasted.

Note: May cause the panel to be refreshed with new data.

Procedural alert ■ To return to non-text mode, click **About** again.

Close

Closes the window and exits the Line Manager.

Multiple Line Manager

The multiline window opens when a line represents multiple connections between:

- two devices
- a device and a package

At the top of the multiline window is a graphic (looking like the Network Map), showing you the connected devices. Below the graphic is a table detailing all the connections.

The screenshot illustrates the Multiple Line Manager interface. At the top, a network map shows two nodes: 'SSII Hub 500 hub 500' (with a '100' icon) and '8 devices for hub 500'. A context menu is open over the '8 devices for hub 500' node, listing options: 'Open Port' (Ctrl+Shift-D), 'Open Line', 'Open Multiline', 'Browse Port Events', and 'Port Properties'. Below the network map is a table with the following data:

Device	Port	Device	Port
hub 500	1.24	172.23.0.17	3
hub 500	1.4	NT4_ENT1	
hub 500	1.3	NT4_ENT	
hub 500		hub 13	
hub 500		B	0.0

A context menu is also open over the '1.3' port in the table, listing options: 'Open Port' (Ctrl+Shift-D), 'Open Line', 'Browse Port Events', and 'Port Properties'. Labels with arrows point to the 'Connections' header and the context menu in the table.

By right-clicking on the port number, you can open a Port Manager or Single Line Manager. By right-clicking on a device name, you can open the Device Manager.

If any of the lines shown have alarms, you will see the alarm icons in the middle column.

7 Attribute Manager

CHAPTER

- To explore icon buttons in the toolbar menus, see:
 - *Toolbar* on page 139
 - *Statistics* on page 142
- To interpret data in the Attribute Manager window, see *Panel Elements* on page 67.

Introduction

Provides you with detailed history of an attribute associated with a device or a port.

Note: Virtual devices cannot have attributes.

Administrator or IT Manager: Also enables you to change the state of an attribute, and to change the way Network Discovery perceives an attribute.

Ways of opening

- Click a port hyperlink from:
 - the Device Manager State panel
 - the Port Manager State panel
 - the Line Manager About panel
- Right-click an attribute on the Events Browser or Alarms Viewer









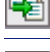
Default panel

- *initial:* Configuration
- *subsequent:* from **Administration > Account Administration > Account properties.**

Toolbar

Availability of buttons in the Attribute Manager toolbar.

Table 7-1: Available toolbar buttons

Icon	Button name	IT Employee or Demo user
	Configuration	✓
	Statistics	✓
	Locate	✓
	Manage [Administrator or IT Manager]	—
	Purge Attribute [Administrator or IT Manager]	—
	Refresh	✓
	Print	✓
	Text	✓
	Close	✓

Configuration

Identifies an attribute and presents details of its most recently observed state.

Heading

Table 7-2: Heading

Element	Notes
Icon	for a complete list, see Help > Classifications > Device Types
Descriptive prefix	for example, “SNMP-managed device”
Device type	for a complete list, see Help > Classifications > Device Types
Device tag	see the <i>User Guide</i>
No. of ports	the number Network Discovery uses for the port may not match the physical port
No. of connections	includes user-assigned connections
Object title	first title available; see Device Title on page 68
Address	IP address; does not appear if identical to object title
Port no./ description	number of port / description of port
Priority	see Priority on page 63

Identity

Table 7-3: Identity

Element	Notes	Optional
Name	for a complete list, see Help > Supported Device/Port Attributes	—
Description	there can be multiples of an attribute (for example, disk, CPU, memory, toner)	✓
Volume label	—	✓

Table 7-3: Identity (Continued)

Element	Notes	Optional
Serial number	—	✓
Units	varies according to the attribute, for example, time, percent, bytes/sec., frames/sec., milliseconds, days and hours, gigabytes. Not applicable for Breaks	✓
Minimum value	—	✓
Maximum value	—	✓
System threshold	available only for those attributes tracked on the Health Panel	✓
Default threshold	available only for those attributes tracked on the Health Panel	✓
State	available only for those attributes tracked on the Health Panel	✓
State Time	available only for the Break attribute	✓
Value	—	—
Ticket number	—	✓
Update time	—	—
Forecast sample count	available only when using the Forecast feature	✓
Forecast first sample	available only when using the Forecast feature	✓
Forecast last sample	available only when using the Forecast feature	✓

For the Break attribute, there are two different times listed:

- The State Time represents the time when the attribute changed state (when the break occurred).
- The Update Time represents the most recent time Network Discovery has seen the problem (i.e. the time of the last poll cycle where the problem was still present).

Note: If a device had a partitioned disk, each partition will appear as a separate “Disk” attribute. You can open an Attribute Manager for each partition. Each partition will have a different disk serial number (assigned by the device OS).

Statistics

Provides a second toolbar with which to view or export detailed historical statistics for the attribute. The statistics may be viewed in graph or table form, and may be exported in Comma Separated Value (CSV) form.

Note: “No data available” means that no data has yet been collected for the attribute. This is normal if the device or port was discovered less than 48 hours before.

The Statistics panel in the Attribute Manager works the same as in the Device Manager. See the description in *Statistics* on page 102 in the *Device Manager* chapter.

Locate

Highlights in a map window the location of the device to which this attribute refers.

If a map session was not open already, one will open now. Within the map window, the device you are locating has a yellow circle around it.

If the window containing the device was already open, that window becomes the front-most window, and the window scrolls so that the highlighted icon can be seen.

Manage [*Administrator or IT Manager*]

Manages the attribute.

Examples: In the case of ports, Administrative Status can be turned on or off. In the case of the Bridge Aging Interval, the length of the interval can be changed.

- Limits**
- Available only when Network Discovery has a write community string for the attribute.
 - Not all attributes can be managed.

Purge Attribute *[Administrator or IT Manager]*

Removes an attribute and its historical statistics from the Network Discovery database.

Warning: This action cannot be undone.

Important: You are *not* making a physical change. If you purge an attribute but the attribute is still present—that is, still associated with a device or port that is still present in your network—Network Discovery will discover the attribute and the attribute will reappear.

- When to use it**
- When an attribute is no longer associated with a device or port.
 - When you no longer wish to retain or examine the history of an attribute.

Refresh

Refreshes the contents of the panel.

- Limits**
- Does not re-read the data in the panel from the network. Re-reads the data only from the Network Discovery database.

Print

Sends the contents of the panel to a printer attached to the management workstation.

Text

Displays the contents of the Attribute Manager panel as text that can be copied and pasted.

Note: If the Statistics **Graph** panel is displayed, the Text button displays a text version of the Table, since there can be no text version of a graph.

Note: May cause the panel to be refreshed with new data.

**Procedural
alert**

To return to non-text mode, click the currently depressed button again.

Close

Closes the window and exits the Attribute Manager.

8 MIB Browser

CHAPTER

The MIB Browser is a tool for the SNMP expert who knows what details to look for and how to look for them.

The MIB (Management Information Base) is a set of data that can be managed with SNMP. If you have the proper community strings for a device, you can use the Network Discovery MIB Browser to read or write data to the device MIB.

Network Discovery has a database of MIB definitions that the MIB Browser uses. The MIB Browser's private enterprises sub-tree contains placeholders for many of the vendors of network equipment who have non-standard or proprietary MIBs.

Note: In order to work, the device must have an IP address, and it must support basic SNMP functionality.

Opening the MIB Browser

You can open a MIB Browser with or without a device in context. In other words, you can open a MIB Browser for a specific device, or you can open the MIB Browser and use its **Find** function to locate the device you want to see.

When you open a MIB Browser with a device in context, you will see the device icon, label and IP address in the right panel. It also shows the value of the “sysName” object from the MIB of that device.

There are three ways to open a MIB Browser with a device in context:

- From the Device Manager, click the **Browse MIB** button.
- Admin: From the Device Manager’s Configuration panel, click a [set] hyperlink.
- From a map window, click **Object > Browse MIB**.

There are three ways to open a MIB Browser without a device in context:

- From the main Toolbar, click the **MIB Browser** button.
- From the MIB Browser, click **File > New MIB Browser**.
- From any applet window (Network Map, Health Panel, and so on), click **Tools > MIB Browser**.

Parts of the MIB Browser

Figure 8-1: MIB Browser example - Folder Panel

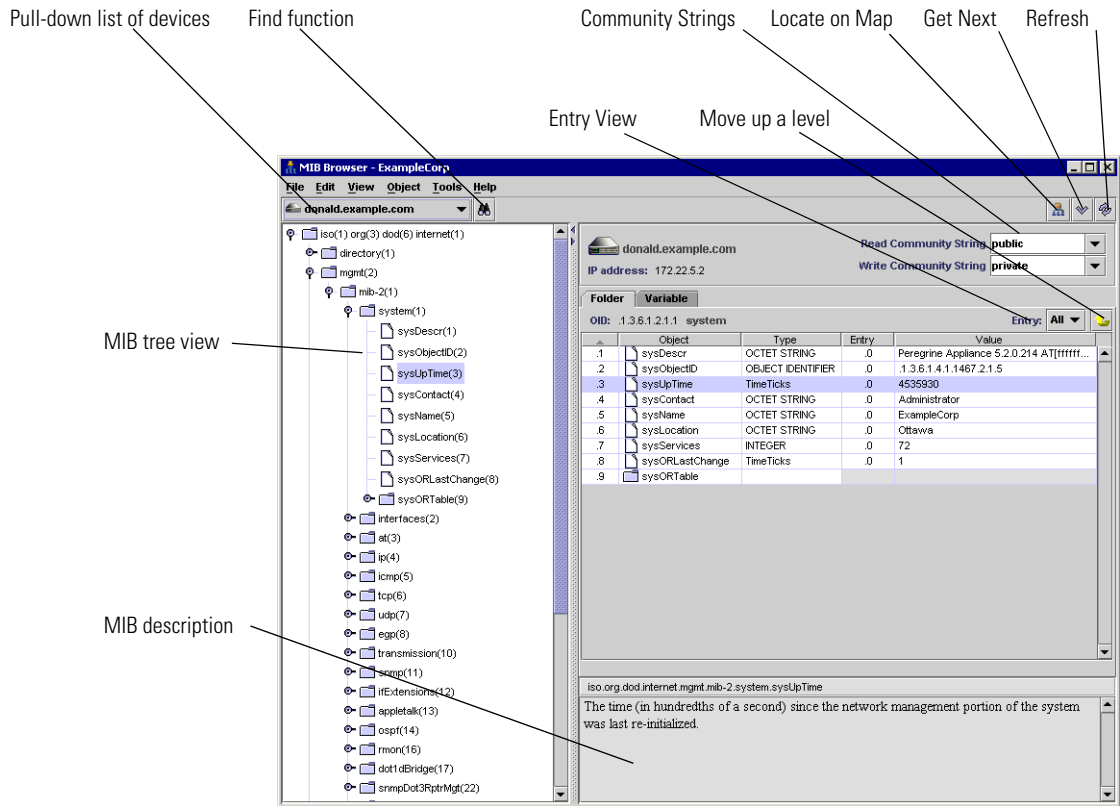
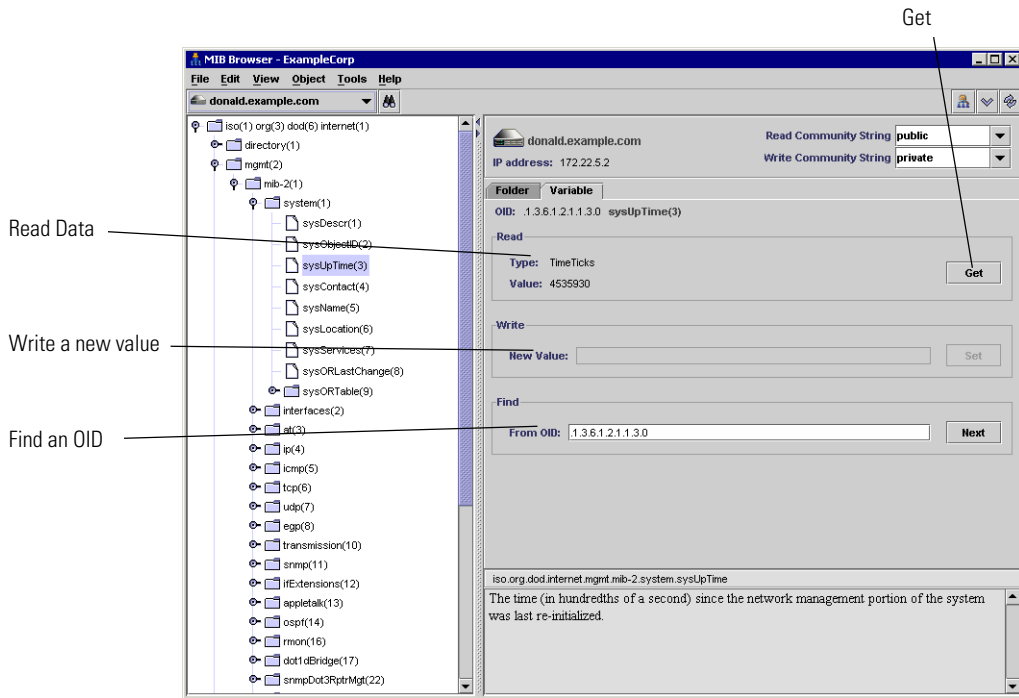


Figure 8-2: MIB Browser example - Variable Panel



Tree View

The left hand side of a MIB Browser shows a tree view of all the MIBs for which Network Discovery has definitions. These definitions are stored within the Peregrine appliance, independent of any one device in your network. You can browse through the definition tree even without an SNMP device in context, by clicking on those tree nodes. Each node represents one SNMP object, and the hierarchy of nodes reflects the SNMP object hierarchy. The name, object ID, type and description of each SNMP object is displayed on the right hand side.

Note: When there is an SNMP device in context (i.e. the device icon and IP address appear), clicking on a tree node will “Get” the value for that object from the device, if that object is supported. No one device supports all the objects in the MIB definition tree.

Pull-down list of Devices

You can toggle between devices in the MIB Browser with this pull-down list. All the devices you have seen in this MIB Browser session appear in this list.

Find Function

If you want to find a particular device to check its MIB, you can use the MIB Browser **Find** button. It works like the Find in the Network Map and other Network Discovery features. Click the button, and a dialog appears. Enter the device name in the dialog and press **Enter**.

Locate on Map

The **Locate** button works like the Locate button in other Network Discovery features. Click this button and you will see where this device is located on the Network Map.

Get Next

The **Get Next** button will assist you in moving through the list of MIB objects.

Note: For a given SNMP device, it is not possible to efficiently determine which MIB definitions it supports. It is only by “MIB walking” (using the **Get Next** button) a device that its supported objects can be determined. Thus the MIB Browser displays the same tree of MIB definitions for all devices, but not all of it is valid for any one device.

Refresh

The **Refresh** button should be used after you change a community string, or after you change the device in context by using the device pull-down list.

Folder Tab

Entry

Within an SNMP device, each supported object can have one or more entries, each of which has a value. For example, an object may define a column of a table, but each row of the column is a different value.

Each entry has an ID too, which defines how to index that entry within the object. The entry ID has the same syntax as an object ID. When a value is displayed, the “OID” field is actually the object ID followed by the entry ID. The “Name” field shows the name of the object followed by the entry ID.

Try thinking of it this way. The tree view on the left hand side shows objects and their hierarchy. The right hand side shows the values of instances of objects. As you press **Get Next**, you may see several successive instances of the same object.

The pull-down list will let you choose how the entries are displayed in the Folder Tab. By choosing **All**, you will see a list of all the entries, in numerical order. Also, you can select an entry number (for example, “.2”) and see only the object values for .2 entries.

Figure 8-3: Entry examples

Entry setting to “All”

Entry setting to “.2”

OID	Object	Type	Entry	Value
.1	ifIndex	INTEGER	.1	1
.1	ifIndex	INTEGER	.2	2
.1	ifIndex	INTEGER	.3	3
.1	ifIndex	INTEGER	.4	4
.2	ifDescr	OCTET STRING	.1	ni0 Hewlett Packard Network Interfac...
.2	ifDescr	OCTET STRING	.2	ni1 Hewlett Packard Network Interfac...
.2	ifDescr	OCTET STRING	.3	lo0 Hewlett-Packard Software Loopb...
.2	ifDescr	OCTET STRING	.4	lan0 Hewlett-Packard 10/100 TV M...
.3	ifType	INTEGER	.1	other(1)
.3	ifType	INTEGER	.2	other(1)
.3	ifType	INTEGER	.3	softwareLoopback(24)
.3	ifType	INTEGER	.4	ethernetCsmacd(6)
.4	ifMtu	INTEGER	.1	0
.4	ifMtu	INTEGER	.2	0
.4	ifMtu	INTEGER	.3	4608
.4	ifMtu	INTEGER	.4	1500
.5	ifSpeed	Guage	.1	1000000
.5	ifSpeed	Guage	.2	1000000
.5	ifSpeed	Guage	.3	10000000

OID	Object	Type	Entry	Value
.1	ifIndex	INTEGER	.2	2
.2	ifDescr	OCTET STRING	.2	ni1 Hewlett Packard Network Interfac...
.3	ifType	INTEGER	.2	other(1)
.4	ifMtu	INTEGER	.2	0
.5	ifSpeed	Guage	.2	1000000
.6	ifPhysAddress	OCTET STRING	.2	
.7	ifAdminStatus	INTEGER	.2	down(2)
.8	ifOperStatus	INTEGER	.2	down(2)
.9	ifLastChange	TimeTicks	.2	0
.10	ifInOctets	Counter32	.2	0
.11	ifInUcastPkts	Counter32	.2	0
.12	ifInNUcastPkts	Counter32	.2	0
.13	ifInDiscards	Counter32	.2	0
.14	ifInErrors	Counter32	.2	0
.15	ifInUnknownProtos	Counter32	.2	0
.16	ifOutOctets	Counter32	.2	0
.17	ifOutUcastPkts	Counter32	.2	0
.18	ifOutNUcastPkts	Counter32	.2	0
.19	ifOutDiscards	Counter32	.2	0

Move Up a Level

This button will move you up one level in the MIB tree view.

Variable Tab

Read/Get

The Read area of the Variable tab displays the currently selected OID. If you want to update the view of that OID, click the **Get** button.

Write

IT Manager and Administrator accounts can write to a device MIB.

Some devices may have a directed community string, which means they accept SNMP operations from only certain devices. The network administrator may have set directed community strings allowing only the Peregrine appliance access to the devices on your network.

Warning: Remember that although the MIB Browser GUI runs on a user's workstation, it is actually the Peregrine appliance that performs the SNMP **Set** and **Get**. A malicious user with the MIB Browser could leverage the Peregrine appliance to effectively bypass the protection of a directed community string. Thus it is a potential security breach to allow a user other than Administrator or IT Manager to do a **Set** with the MIB Browser.

To change a MIB entry

- 1 In the MIB Browser **Folder** panel, select an OID.
- 2 Click the **Variable** panel.
The **Variable** panel will show you the current definition for the OID.
- 3 Select the correct Write community string.
- 4 In the **Write** section, enter a new definition for that OID.
- 5 Click **Set**.

Find OID

When you select an OID in the tree view, the OID will also appear in the “Find OID” text box. Click the **Next** button to move through the MIB, like you would with the **Get Next** button at the top of the panel.

To go to a specific OID, you can change the OID in the “Find OID” text box, and click **Next**. The MIB Browser will go directly to that object entry.

MIB Description

This area provides the standard description for each MIB object.

Sometimes, especially when first learning about MIBs, it is educational to view just the description of an object. Open a MIB Browser without a device in context, and you can see all the MIB descriptions available.

Read and Write Community Strings

Your ability to read or write MIB data depends on your account type.

Demo and IT Employee accounts can only read with the “public” read community string on devices.

IT Manager and Administrator accounts have full read/write access, as long as you have the correct community strings.

The MIB Browser pull-down lists of community strings are populated by the network community strings you created in **Administration > Network configuration > Community Property Groups**. This procedure is described in the *Setup Guide*.

The MIB Browser requests the complete list of community strings, as supplied on the **Community Property Groups** page, and takes from that list all strings that apply to the device in context. For example, if the **Community Property Groups** page lists the following strings for the network being explored:

- public (r), for 0.0.0.0-255.255.255.255
- private (w), for 192.168.0.0-192.168.9.255
- su_only (r/w), for 192.168.0.0-192.168.0.255
- OnTheHalves (r/w), for 192.168.1.0-192.168.1.255

then the device 192.168.1.32 will show the following strings:

- public (r)

- private (w)
- OnTheHalves (r/w)

The string “su_only (r/w)” will not appear in this window since the device's IP address (192.168.1.32) is outside the range of the string (192.168.0.0-192.168.0.255).

Note: Community strings are case-sensitive. “Public” and “public” are two different strings.

If necessary, you can enter a new community string into the text box in the MIB Browser.

As Network Discovery discovers the managed devices in your network, it uses the read community strings that you have configured to read the MIBs of those devices. The MIB Browser automatically uses the read community string that Network Discovery has determined is valid for that device.

However, if that device is not yet known to Network Discovery, then Network Discovery does not know the valid read community string for that device, and you need to enter a community string in the text box.

The situation is a bit different for write community strings. Network Discovery must have a valid read community string to discover a managed device, but a valid write community string is optional. Network Discovery tries to determine a valid write community string for each managed device from the list of strings in the **Network Configuration** menu. If it finds one, the MIB Browser uses it, but otherwise the MIB Browser has no current write string.

Note: If at any time you change the community string that you are using to view the MIB, click the **Refresh** button.

Walking the MIB

The **Get Next** button requests the value of the next SNMP object instance supported by this device. You may have noticed that as you push **Get Next**, the tree is expanded as necessary to keep the current object highlighted. Sometimes the next object a device supports is not the next item in the MIB tree because no one device supports all the objects in the MIB definition tree; some parts of the MIB tree are not relevant for any given device. These irrelevant sections of the MIB tree get skipped.

If there is no data available for a MIB object, the Value column will appear gray. If the community string does not allow you access to the MIB, you will see “no response.”

If there is a “no response” message, the SNMP device did not respond to a **Get**, **Get Next**, or **Set** request. There are a few reasons for this error:

- The device does not have an SNMP agent; in other words the device is not managed.
- The device has an SNMP agent, but it did not respond to the request within a certain amount of time. Some devices drop management requests when they are too busy handling their network traffic.
- The community string being used by the MIB Browser is incorrect. Perhaps someone recently changed the device's community string and Network Discovery has not yet, or cannot, determine a valid one.

Unfortunately, the SNMP protocol does not distinguish amongst these conditions.

Using Multiple MIB Browser Sessions

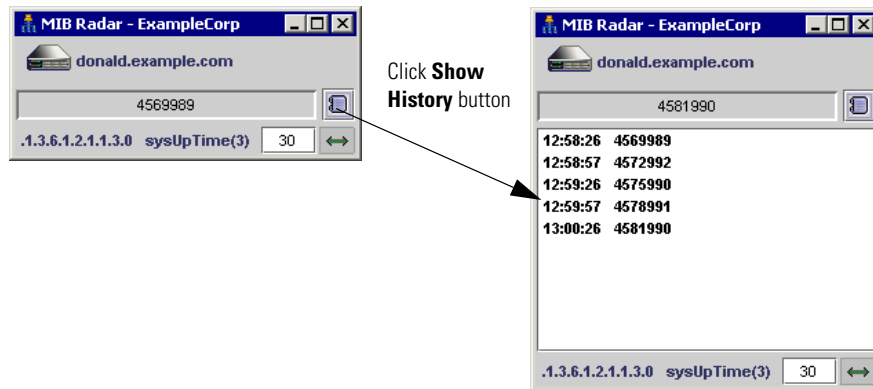
You can have more than one MIB Browser window open at any time. Also, you can toggle between several devices in the “found device” pull-down list at the top-left of the MIB Browser window.

To open a new MIB Browser session from your current MIB Browser window, click **File > New MIB Browser**.

Watching an OID with MIB Radar

You can use the MIB Radar feature to watch a particular MIB object in a separate small window on your screen. If you want to monitor one counter in the MIB (for example, sysUpTime), you can select that OID and then click **File > Open Radar**. The radar window will appear, which looks like this:

Figure 8-4: MIB Radar example



Note: By default, the data refreshes every 30 seconds.

You can change the refresh rate by entering a number higher than 5 in the text box. Also, you can view the history of the OID by clicking on the **Show History** button.

Saving MIB Data as a Text file

There are two ways you can save MIB Browser data to a text file:

- Save Table Data
- MIB Walk

Save Table Data

You can use the **Save Table Data** feature to save selected info into a tab-separated-value (.tsv) file. This feature can also be found in the Health Panel, MIB Browser, Alarms Viewer, Service Analyzer, and Events Browser.

You can save the entire contents of a MIB Browser table, or you can Ctrl-click to select the OIDs you want to save.

Saving data to a text file

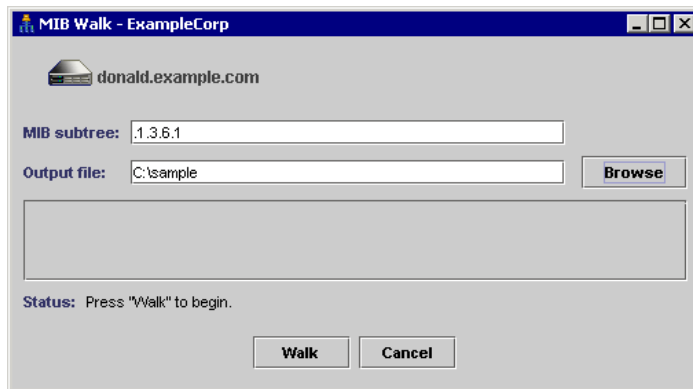
- 1 Select the MIB OIDs you want to save.
- 2 Click **File > Save Table Data**.
A Save Table Data dialog appears.
- 3 Select a file name and location for the text files.
- 4 Click **Save**.

MIB Walk

You can use the MIB Walk feature to save a sub-tree from the MIB Browser as a text (.txt) file.

Saving a MIB Walk

- 1 Click **File > Open Walk**.
The MIB Walk dialog appears.



- 2 Enter an OID that represents the start of the sub-tree you want to save.
- 3 Click **Browse** to select a name and location for your output file.
- 4 Click **Walk**.

9 Reports

CHAPTER

Network Discovery reports comprise the following groups:

- *Executive/Summary Network Reports* on page 160
- *Scanned Machine Reports* on page 163
- *WAN Reports* on page 164
- *LAN Reports* on page 166
- *Device Reports* on page 167

Note: All reports will reflect the Prime map configuration and its packaging.

Report periods

There are two types of report, summary and detail. Both report types have a different group of reporting periods.

Table 9-1: Reporting period groups for summary and detail

Summary	Detail
Today	All Periods*
Last 7 Days	Yesterday
Last Week	Last 7 Days
This Month	Last Week
Last Month	Last Month

*All periods in this group.

Table 9-2: Reporting periods

Period	Contents	Generated	Summary	Detail
Today	data for today and yesterday	each hour*	✓	—
Yesterday	data for the previous 24 hours	each day after midnight	—	✓
Last 7 Days	data for the previous 7 days, starting yesterday (not including today)	each day after midnight	✓	✓
Last Week	data for the previous week (weeks begin each Monday)	each Monday	✓	✓
This Month	data for the days in the current month, starting yesterday (not including today)	each day after midnight	✓	—
Last Month	data for the previous calendar month	on the first day of each month	✓	✓

*For a restricted period: 0600–2000 (6 AM–8 PM).

Report statistics

Many reports feature bar graphs and values for three statistics: the peak, the mean peak, and the mean.

A mean value and a peak value are collected for every sample. At the end of the report period, all peak values are used to calculate the mean peak.

Imagine that we record a mean value and a peak value three times a day:

Table 9-3: Examples for mean and peak values

Value	first	second	third
Mean	2.0	2.0	3.0
Peak	6.0	7.0	6.0

To obtain the mean peak for the day, we take the peak values of 6.0, 7.0, and 6.0, and find the mean of those three values, which is 6.3.

Different report periods have different statistical sampling periods. For example, a report with the period “Yesterday” takes samples every five minutes. A report for “Last 7 Days” takes samples every hour.

Executive/Summary Network Reports

Executive Summary reports are about the network as a whole. Here is one example of how you might use them—just to see what’s in your network.

To view the Executive Summary Network Inventory Reports

- ▶ Reports > Network Documentation > Device Inventory Summary
- ▶ Or Reports > Network Documentation > Device Inventory

You may have very little idea of what is actually in your network beyond the core network devices:

- There may be several people responsible for the network.
- Someone or several people may be adding equipment without informing you.
- You may be new to the job and the last person didn’t keep complete records or records you can understand.
- Some or all of the network management may have been delegated to someone outside your organization.
- You may be outside the organization whose network you must manage.

You may not even know what you don’t know! The Device Inventory Summary report tells you and the Device Inventory report tells you in more detail.

Table 9-4: Executive/Summary Reports reports

Folder	Report	Type
Network Documentation	Network Classification	pie graph, table
	Network Devices by Function	pie graph, table
	End Nodes by Function	pie graph, table
	Device Inventory Summary	table
	Device Inventory by Category	list
	Resource Managed Devices Inventory	list
	Resource Inventory and Usage	table
	Frame Relay PVC Inventory	table
	Possible Modems Report	list
	Underutilized Equipment	table

Table 9-4: Executive/Summary Reports reports (Continued)

Folder	Report	Type
Aggregate Inventory*	Network Classification	pie graph, table
	Network Devices by Function	pie graph, table
	End Nodes by Function	pie graph, table
	Device Inventory Summary	table
	Device Inventory by Category	list
	Device Inventory	list
	Resource Managed Devices Inventory	list
	Resource Managed Devices Attributes	table
	Frame Relay PVC Inventory	table
	Possible Modems Report	list
	Underutilized Equipment	table
Performance Summaries	Network Summary Reports	line/bar graphs
	WAN Summary Reports	line/bar graphs
	Frame Relay Summary Reports	line/bar graphs
	DSL Summary Reports	line/bar graphs
	Point to Point Summary Reports	line/bar graphs
	Serial to SPN Summary Reports	line/bar graphs
	ATM Summary Reports	line/bar graphs
	LAN Backbone Summary Reports	line/bar graphs
	FDDI Summary Reports	line/bar graphs
Token Ring Summary Reports	line/bar graphs	

Table 9-4: Executive/Summary Reports reports (Continued)

Folder	Report	Type
Alarm Summaries [†]	All Alarms	table
	Line Breaks	table
	Line Utilization [†]	table
	Delay	table
	Collisions	table
	Broadcasts	table
	Errors	table
	Device Breaks	table
	Packet Loss	table
Network Wide	Network Availability	line/bar graphs
	Mean Network Utilization [‡]	line/bar graphs
	Peak Network Utilization ^{**}	line/bar graphs
	Mean Network Throughput [‡]	line/bar graphs
	Peak Network Throughput ^{**}	line/bar graphs
	Inventory	list
	Availability Details	table
	Utilization Details	table

* These reports are only available to Aggregator appliances.

† Line alarm reports include connected lines only.

‡ For the Mean graphs, only mean values are used, so the lines represent the “mean of the means.”

** For the peak graphs, only peak values are used. This means that the bars on the graphs represent the highest peak utilization value of all the peak values from the network devices.

Performance Summaries

These folders contain the following periods:

- Today
- Last 7 Days
- Last Week

- This Month
- Last Month

Alarm Summaries

These folders contain the following periods:

- All Periods
- Yesterday
- Last 7 Days
- Last Week
- Last Month

Network Wide

The folders for Availability Details and Utilization Details contain the following periods:

- All Periods
- Yesterday
- Last 7 Days
- Last Week
- Last Month

Scanned Machine Reports

These reports display summary counts of the scanned machines grouped by different machine properties. For example, machines at the top level may be grouped by their company division, in turn by their office location within that division, and finally by the department to which they belong.

The summary reports provide drill-down to details of those machines which belong to the summary group clicked on.

For more information, see *Desktop Inventory Reports* on page 171.

WAN Reports

Frame Relay reports, as an example, can tell you if you are getting the service you are paying for. Note, for instance, the Data Delivery Ratio Report, one of the detailed reports. The Data Delivery Ratio Report tells you which Permanent Virtual Circuits (PVCs) are dropping data and is a good guide to whether or not you are getting the Frame Relay service you are paying for and whether you could do with less.

To view a Data Delivery Ratio Report

- ▶ Reports > WAN Reports > Frame Relay Service > Data Delivery Ratio

There are two report structures for WAN Reports:

- Frame Relay folder
- all other folders

Table 9-5: Frame Relay reports

Frame Relay Summary Reports

Frame Relay Availability

Frame Relay Mean Utilization

Frame Relay Peak Utilization

Frame Relay Mean Throughput

Frame Relay Peak Throughput

Frame Relay Detail Reports

Inventory

Connected DLCI Inventory

Availability Details

Mean Time Between Service Outage (MTBSO)

Mean Time To Service Repair (MTTSR)

PVC Utilization

Over Utilized PVCs

Under Utilized PVCs

Interface/DLCI Utilization

Table 9-5: Frame Relay reports (Continued)

Congested PVCs

Data Delivery Ratio (DDR)

Frame Delivery Ratio (FDR)

Unconnected Frame Relay Ports

Delay Details

In Table 9-6, *<WAN_type>* stands for one of the following:

- Point to Point (Serial)
- Serial to SPN (Service Provider Network)
- DSL (Digital Subscriber Line)
- ATM (Asynchronous Transfer Mode)
- WAN (WAN Wide)

Table 9-6: All other WAN Reports

Report/Folder	Type
Summary Reports	
<i><WAN_type></i> Availability	line/bar graphs
Mean <i><WAN_type></i> Utilization	line/bar graphs
Peak <i><WAN_type></i> Utilization	line/bar graphs
Mean <i><WAN_type></i> Throughput	line/bar graphs
Peak <i><WAN_type></i> Throughput	line/bar graphs
Detail Reports	
Inventory	list
Availability Details	table
Utilization Details	table

LAN Reports

LAN reports give you inventory information and information about the availability, throughput and utilization of your Local Area Network whether you have a LAN backbone, FDDI, or Token Ring.

In Table 9-7, *<LAN_type>* stands for one of the following:

- LAN Backbone
- FDDI
- Token Ring

Table 9-7: LAN Reports

Report/Folder	Type
Summary Reports	
<i><LAN_type></i> Availability	line/bar graphs
Mean <i><LAN_type></i> Utilization	line/bar graphs
Peak <i><LAN_type></i> Utilization	line/bar graphs
Mean <i><LAN_type></i> Throughput	line/bar graphs
Peak <i><LAN_type></i> Throughput	line/bar graphs
Detail Reports	
Inventory	table
Availability Details	table
Utilization Details	table

The folders for Availability Details and Utilization Details contain the following periods:

- All Periods
- Yesterday
- Last 7 Days
- Last Week
- Last Month

Device Reports

Device reports give you inventory information and information about availability, throughput and utilization, broken down by category of device. They can also give you such information as what servers are using the most memory for a given time.

Note: The Inventory report exported to a CSV file reflects the default map configuration for the current account.

All other reports reflect the Prime map configuration and its packaging.

Device reports are available for the following groupings of devices:

- Servers
- Routers
- Input and Output Devices
- Resource Managed Workstations
- Web Servers

Table 9-8: Device Reports

Report/Folder	Type
All Devices	
Inventory	list
Availability Details	table
Utilization Details	table
Resource Managed	
Top CPU Utilization	table
Top Memory Utilization	table
Top Load Average	table
Top Disk Utilization	table
Top Virtual Memory Utilization	table

The Resource Managed folders contain the following:

- Inventory
- All Periods

- Yesterday
- Last 7 Days
- Last Week
- Last Month

Note: The Web Servers reports will reflect the default map configuration for the current account.

Microsoft Word documents

Network Discovery comes with two documents for Microsoft Word that allow you to print reports with graphs of your network. The first document is an example report framework for the intermediate Microsoft Word user. To use it, you use cut and paste to rearrange the built-in graphs. The second document is a report template for the advanced Microsoft Word user. To use the second document, you should be comfortable with Word field codes and macro substitution.

Each document contains links to an Network Discovery graph on your Peregrine appliance. Once you customize the report with the name of your Peregrine appliance, you can easily update the graphs—to present at weekly meetings, for example.

The exact steps for setting up Microsoft Word to use the Network Discovery templates are described in the *Data Export Guide*.

Compatibility

Compatible with:

- Microsoft Word 97
- Microsoft Word 2000
- Microsoft Word 2003

You can use Network Discovery data with other applications

You can use your own data access application to customize the presentation of Network Discovery data, if your data access application operates on the Open Database Connectivity (ODBC) standard. ODBC applications into which you can export Network Discovery data comprise (but are not limited to):

- Crystal
- Cognos Impromptu
- PowerPlay OLAP tool

For more information on how to use Network Discovery data with data access applications, see the *Data Export Guide*.

10 Desktop Inventory Reports

CHAPTER

In this chapter you will find information on the following topics:

- *The Reports database* on page 172
- *Reports available on the appliance for scanned machines* on page 173
- *Producing your own reports* on page 179

The Reports database

The Peregrine appliance contains an inventory database which is accessible via ODBC.

This database is used for reporting and export purposes.

The data includes:

- Inventory of devices
- Ports
- Connections
- Desktop hardware and installed software applications

The schema of this database is optimized for reporting performance and human understandability rather than real-time access and statistic storage. In particular, the tables and columns are clearly named and the schema is somewhat normalized.

Further information

For further information about the database schema please refer to the documents entitled:

- *Network Discovery Data Export Guide* available in pdf format.
- *Desktop Inventory Data collected by the Scanners* available in HTML format from the Desktop Inventory Start menu.

Open Database Access

The reports database is open. This means the schema is published using standard description languages and the data can be accessed and exported in a number of ways using standard protocols.

- The schema of the reports database is described using English, SQL and XML DTDs. All descriptions are available from the Peregrine appliance.
- The data can be accessed directly on the appliance (queries only, no updates) using the Open Database Connectivity (ODBC) protocol. MySQL provides an ODBC driver for Windows. This driver is available on the appliance:
 - From **Home**, click **Download**. The file is called **MyODBC-Windows-3.51.06.exe**.

Using this driver, Windows-based applications can remotely access the reports database.

- Third-party reporting tools, for example Crystal Reports, Cognos Impromptu and Cognos Powerplay, can generate custom reports from the reports database via ODBC. See the *Producing your own reports* on page 179 for further information on how to do this.
- Several Connect-It scenarios have been provided to make use of this database. You can find further information about the Connect-It scenarios in the Connect-It documentation.

Reports available on the appliance for scanned machines

Accessing the Scanned Machine Reports on the appliance

Several pre-defined reports have been provided on the appliance.

To access the Scanned Machine Reports on the appliance:

- 1 Login to the appliance.
- 2 From **Home** click the **Reports** link. The **Reports** page is displayed.



- 3 There are three types of Scanned Machine Reports. Click on the link for the type of report you require.
 - a Scanned Machine Summaries
 - b Applications
 - c Scan File Status

Scanned Machine Summaries

These reports display summary counts of the scanned machines grouped by different machine properties. For example, machines at the top level may be grouped by their company division, in turn by their office location within that division, and finally by the department to which they belong.

The summary reports provide drill-down to details of those machines which belong to the summary group clicked on.

The following Scanned Machine Summaries based on Asset Data are available:

Note: If collection of the relevant Asset Data fields is not enabled, the data will be categorized as N/A, making the reports less useful).

Summary report by Division, Location, Department

This report lists summary counts for all scanned machines by Division, Location, and Department.

Clicking on a summary count for a Division, Location, or Department will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a location count will display all machines at that location and division sorted by department.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by division, location, and department.

Summary Report By Location, Division, Department

This report lists summary counts for all scanned machines by Location, Division, and Department.

Clicking on a summary count for a Location, Division, or Department will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a division count will display all machines at that division and location sorted by department.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by location, division, and department.

Summary Report By Department, Location

This report lists summary counts for all scanned machines by Department and Location.

Clicking on a summary count for a Department or Location will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a department count will display all machines at that department sorted by location.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by department and location.

Summary Report By Location, Building, Floor

This report lists summary counts for all scanned machines by Location, Building, and Floor.

Clicking on a summary count for a Location, Building, or Floor will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a building count will display all machines at that building and location sorted by floor.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by location, building, and floor.

Summary Report By Location, Cost Center

This report lists summary counts for all scanned machines by Location and Cost Center.

Clicking on a summary count for a Location or Cost Center will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a location count will display all machines for that location sorted by cost center.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by location and cost center.

Summary Report By Cost Center, Location

This report lists summary counts for all scanned machines by Cost Center and Location.

Clicking on a summary count for a Cost Center or Location will display a report of machines belonging to those categories summarized and sorted by any lower level categories. For example, clicking on a cost center count will display all machines for that cost center sorted by location.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by cost center and location.

Summary Report By Operating System Category

This report lists summary counts for all scanned machines by Operating System Category.

Clicking on a summary count for an Operating System category will display a detailed report of machines belonging to that Operating System category.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by Operating System category.

Summary Report By Hardware Chassis Type

This report lists summary counts for all scanned machines by hardware chassis type.

Clicking on a summary count for a chassis type will display a detailed report of machines belonging to that chassis type.

Clicking on the Total Scanned Machines count at the bottom of the report will display a report of all scanned machines sorted by hardware chassis type.

Application Reports

These reports display software application licence and installation counts for all installed applications grouped by application, application version, and application publisher.

The reports also provide links to detailed reports of those scanned machines where the individual applications are installed.

These reports are based on Application Recognition performed by the XML Enricher.

To ensure that the data presented is sufficiently accurate, make sure that

- Application Recognition is enabled in the XML Enricher. See the ‘Configuring the XML Enricher on the appliance’ section in the document entitled *Using Network Discovery with Desktop Inventory and Desktop Administration*.
- The Application Library used is up to date. See the ‘Updating the application library used by the Enricher’ section in the document entitled *Using Network Discovery with Desktop Inventory and Desktop Administration*.

Application Licence Reports

Licence Counts by Application

This summary report displays all applications by publisher and application with counts of licences required and installations for each application.

Clicking in the Publisher column will display detailed version information about all of the publisher's applications. Clicking in the Application Name column will display the details of the different versions of that application.

Top 20 Applications

This report lists those applications requiring the largest number of licences, sorted by number of licences.

Clicking in the Publisher column will display detailed version information about all of the publisher's applications. Clicking in the Application Name column will display the details of the different versions of that application.

Top 50 Applications

This report lists those applications installed on the greatest number of scanned machines which require a licence.

Clicking in the Publisher column will display detailed version information about all of the publisher's applications. Clicking in the Application Name column will display the details of the different versions of that application.

Top 100 Applications

This report lists those applications installed on the greatest number of scanned machines which require a licence.

Clicking in the Publisher column will display detailed version information about all of the publisher's applications. Clicking in the Application Name column will display the details of the different versions of that application.

Publishers Licence Report

Publishers Summary

This summary report lists all publishers sorted by name together with their application licences required and installed application counts.

Clicking on the publisher's name will display detail information for all of the publisher's applications.

Top 20 Publishers

This report lists up to 20 publishers with the greatest number of installed applications that require licences.

Clicking on the publisher's name will display detail information for all of the publisher's applications.

Top 50 Publishers

This report lists up to 50 publishers with the greatest number of installed applications that require licences.

Clicking on the publisher's name will display detail information for all of the publisher's applications.

Top 100 Publishers

This report lists up to 100 publishers with the greatest number of installed applications that require licences.

Clicking on the publisher's name will display detail information for all of the publisher's applications.

Scan File Status Reports

These reports display the following details for the Scanner and Listener, grouped by Scanner Status and Scan File Status.

- Machine
- Listener Version
- Listener Operating System
- Listener Port Number
- Scanner Software Status
- Scan File Date
- Scan File Status

The reports provide drill-down to details of those machines which belong to the summary group clicked on.

If collection of the relevant fields is not enabled, the data will be categorized as N/A, making the reports less useful.

Producing your own reports

You can use your own data access application to customize the presentation of the data, if your application operates on the Open Database Connectivity (ODBC) standard. ODBC applications into which you can export data comprise (but are not limited to):

- Microsoft Access
- Crystal
- Cognos Impromptu
- PowerPlay OLAP tool

Note: The data imported into the data access application is read-only. You cannot manipulate it or your network from within the application.

Further information

Refer to the *Data Export Guide* for more information on how to produce your own reports.

11

Need more help?

CHAPTER

Peregrine is committed to ensuring your success with our products. We offer a number of ways for you to provide product feedback, suggest enhancements, and receive technical assistance with any issues you encounter.

For further information and assistance contact Peregrine's CenterPoint Web Site.

Peregrine's CenterPoint Web Site

Current details of local support offices are available through Peregrine's CenterPoint Web site at <http://support.peregrine.com>.

To find Peregrine worldwide contact information:

- 1 Log on with your login user name and password.
- 2 Click **Go** for **CenterPoint**.
- 3 Select **Whom Do I Call?** in the navigation bar on the left side of the page. Peregrine worldwide information is displayed for all products.

12 Copyright

CHAPTER

Peregrine Systems acknowledges the copyrights belonging to the following third parties. (This page constitutes a continuation of the copyright page.)

agetty

Copyright Wietse Venema 1987, 1988, 1989, 1990, 1991, 1992, 1993, 1994, 1995. Written by Wietse Venema. Portions are Copyright University of Toronto, 1994, 1995. Written by John DiMarco, Matthew Lee. Portions are Copyright (C) 1992-98 Digi International. All rights reserved.

Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it freely, subject to the following restrictions: (1) The origin of this software must not be misrepresented, either by explicit claim or by omission; (2) Credits must appear in the documentation; (3) Altered versions must be plainly marked as such, and must not be misrepresented as being the original software; (4) This notice may not be removed or altered.

am-utils

Copyright (c) 1997-2003 Erez Zadok

Copyright (c) 1989 Jan-Simon Pendry

Copyright (c) 1989 Imperial College of Science, Technology & Medicine

Copyright (c) 1989 The Regents of the University of California.

All rights reserved.

This code is derived from software contributed to Berkeley by Jan-Simon Pendry at Imperial College, London.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

This product includes software developed by the University of California, Berkeley and its contributors, as well as the Trustees of Columbia University.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

apache httpd

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>).

cmu-snmp

Copyright 1998 by Carnegie Mellon University All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

CMU DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CMU BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

file

Copyright © Ian F. Darwin 1986, 1987, 1989, 1990, 1991, 1992, 1994, 1995.

Software written by Ian F. Darwin and others; maintained by Christos Zoulas.

gd

Portions copyright 1994, 1995, 1996, 1997, 1998, 1999, by Cold Spring Harbor Laboratory. Funded under Grant P41-RR02188 by the National Institutes of Health.

Portions copyright 1996, 1997, 1998, 1999, by Boutell.Com, Inc.

Portions relating to GD2 format copyright 1999 Philip Warner.

Portions relating to PNG copyright 1999, Greg Roelofs.

Portions relating to libtiff copyright 1999, John Ellson (ellson@lucent.com).

Permission has been granted to copy and distribute gd in any context without fee, including a commercial application, provided that this notice is present in user-accessible supporting documentation.

This does not affect your ownership of the derived work itself, and the intent is to assure proper credit for the authors of `gd`, not to interfere with your productive use of `gd`. If you have questions, ask. “Derived works” includes all programs that utilize the library. Credit must be given in user-accessible documentation.

This software is provided “AS IS.” The copyright holders disclaim all warranties, either express or implied, including but not limited to implied warranties of merchantability and fitness for a particular purpose, with respect to this code and accompanying documentation.

Although their code does not appear in `gd` 1.6.3, the authors wish to thank David Koblas, David Rowley, and Hutchison Avenue Software Corporation for their prior contributions.

gnuplot

Copyright 1986–1993, 1998 Thomas Williams, Colin Kelley

Permission to use, copy, and distribute this software and its documentation for any purpose with or without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation.

Permission to modify the software is granted, but not the right to distribute the complete modified source code. Modifications are to be distributed as patches to the released version. Permission to distribute binaries produced by compiling modified sources is granted, provided you

1. distribute the corresponding source modifications from the released version in the form of a patch file along with the binaries,
2. add special version identification to distinguish your version in addition to the base release version number,
3. provide your name and address as the primary contact for the support of your modified version, and
4. retain our contact information in regard to use of the base software.

Permission to distribute the released version of the source code along with corresponding source modifications in the form of a patch file is granted with same provisions 2 through 4 for binary distributions.

This software is provided “as is” without express or implied warranty to the extent permitted by applicable law.

iputils/ping

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the following copyright notice (Copyright (c) 1989 The Regents of the University of California. All rights reserved), this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution (THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) All advertising materials mentioning features or use of this software must display the following acknowledgement (This product includes software developed by the University of California, Berkeley and its contributors); and (3) Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.netkit-telnet.

jpeg

This software is copyright (C) 1991-1998, Thomas G. Lane. All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

(1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.

- (2) If only executable code is distributed, then the accompanying documentation must state that “this software is based in part on the work of the Independent JPEG Group”.
- (3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

libicu

Copyright © 1995–2001 International Business Machines Corporation and others. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

libogg

Copyright (c) 2002, Xiph.Org Foundation

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Xiph.Org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

libvorbis

Copyright (c) 2002, Xiph.org Foundation

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- Neither the name of the Xiph.org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

lilo

Authored by Werner Almesberger (wa@almesberger.net) and John Coffman (johninsd@san.rr.com).

BSD license required.

lsof

Copyright 1994, 1998, 2000 Purdue Research Foundation, West Lafayette, Indiana 47907. All rights reserved.

MySQL

Network Discovery includes software whose copyright is owned by MySQL, A.B.

ncftp

Used under the following terms:

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain, or those made Freely Available, or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
 - a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major network archive site allowing unrestricted access to them, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 - b) use the modified Package only within your corporation or organization.
 - c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 - d) make other distribution arrangements with the Copyright Holder.
 - e) permit and encourage anyone who receives a copy of the modified Package permission to make your modifications Freely Available in some specific way.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 - a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 - b) accompany the distribution with the machine-readable source of the Package with your modifications.
 - c) give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
 - d) make other distribution arrangements with the Copyright Holder.
 - e) offer the machine-readable source of the Package, with your modifications, by mail order.

5. You may charge a distribution fee for any distribution of this Package. If you offer support for this Package, you may charge any fee you choose for that support. You may not charge a license fee for the right to use this Package itself. You may distribute this Package in aggregate with other (possibly commercial and possibly nonfree) programs as part of a larger (possibly commercial and possibly nonfree) software distribution, and charge license fees for other parts of that software distribution, provided that you do not advertise this Package as a product of your own. If the Package includes an interpreter, You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called “undump” or “unexec” methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of the Standard Version of the Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

netaddr-ip

This software is (c) Luis E. Muñoz. It can be used under the terms of the perl artistic license provided that proper credit for the work of the author is preserved in the form of this copyright notice and license for this module.

net-snmp

Various copyrights apply to this package, listed in separate parts below.

Part 1: CMU/UCD copyright notice: (BSD like)

Copyright 1989, 1991, 1992 by Carnegie Mellon University; Derivative Work - 1996, 1998-2000 Copyright 1996, 1998-2000 The Regents of the University of California. All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

Part 2: Networks Associates Technology, Inc copyright notice (BSD)

Copyright (c) 2001-2003, Networks Associates Technology, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED

TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Part 3: Cambridge Broadband Ltd. copyright notice (BSD)

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

Part 4: Sun Microsystems, Inc. copyright notice (BSD)

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

Part 5: Sparta, Inc copyright notice (BSD)

Copyright (c) 2003, Sparta, Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED

TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

OpenSSH

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OpenSSL

Copyright (c) 1998-2000 The OpenSSL Project. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.); (2) All advertising materials mentioning features or use of this software must display the following acknowledgment: This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org>); (3) The names “OpenSSL Toolkit” and “OpenSSL Project” must not be used to endorse or promote products derived from this software without prior written permission; (4) Products derived from this software may not be called “OpenSSL” nor may “OpenSSL” appear in their names without prior written permission of the OpenSSL Project; (5) Redistributions of any form whatsoever must retain the following acknowledgment: “This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org>)”.

PPP

Copyright © 1989 Carnegie Mellon University.

Copyright © 1991 Gregory M. Christy.

Copyright © 1993 The Australian National University.

Copyright © 1994 Philippe-Andre Prindeville.

Copyright © 1995 Eric Rosenquist, Strata Software Limited.

Copyright © 1995 Pedro Roque Marques

All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by the Authors listed above. The names of the Authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

psmisc

Copyright 1993-1998 Werner Almesberger. All rights reserved.

Redistribution and use in source and binary forms of parts of or the whole original or derived work are permitted provided that the original work is properly attributed to the author. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission. This work is provided “as is” and without any express or implied warranties.

python-optick

Copyright (c) 2001-2002 Gregory P. Ward. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

strace

Copyright © 1991, 1992 Paul Kranenburg <pk@cs.few.eur.nl>

Copyright © 1993 Branko Lankester <branko@hacktic.nl>

Copyright © 1993 Ulrich Pegelow <pegelow@moorea.uni-muenster.de>

Copyright © 1995, 1996 Michael Elizabeth Chastain
<mec@duracef.shout.net>

Copyright © 1993, 1994, 1995, 1996 Rick Sladkey <jrs@world.std.com>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

tcpdump and libpcap

Redistribution and use in source and binary forms, with or without modification, are permitted provided that: (1) source code distributions retain the following copyright notice (Copyright (c) 1988, 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997 The Regents of the University of California. All rights reserved) and this paragraph in its entirety, (2) distributions including binary code include the above copyright notice and this paragraph in its entirety in the documentation or other materials provided with the distribution, and (3) all advertising materials mentioning features or use of this software display the following acknowledgement: "This product includes software developed by the University of California, Lawrence Berkeley Laboratory and its contributors."

telnet

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the following copyright notice (Copyright (c) 1983, 1995 Eric P. Allman

Copyright (c) 1995, 2000 David A. Holland. Copyright (c) 1988, 1989, 1990, 1991, 1993 The Regents of the University of California. All rights reserved), this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution (THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) All advertising materials mentioning features or use of this software must display the following acknowledgement (This product includes software developed by the University of California, Berkeley and its contributors); and (3) Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

traceroute

Redistribution and use in source and binary forms, with or without modification, are permitted provided that: (1) source code distributions retain the following copyright notice (Copyright (c) 1988, 1989, 1991, 1994, 1995, 1996, 1997 The Regents of the University of California. All rights reserved.) and this paragraph in its entirety, (2) distributions including binary code include the above copyright notice and this paragraph in its entirety in the documentation or other materials provided with the distribution, and (3) all advertising materials mentioning features or use of this software display the following acknowledgement: "This product includes software developed by the University of California, Lawrence Berkeley Laboratory and its contributors."

unzip

This is version 2002-Feb-16 of the Info-ZIP copyright and license. The definitive version of this document should be available at <ftp://ftp.info-zip.org/pub/infozip/license.html> indefinitely.

Copyright (c) 1990-2002 Info-ZIP. All rights reserved.

For the purposes of this copyright and license, "Info-ZIP" is defined as the following set of individuals:

Mark Adler, John Bush, Karl Davis, Harald Denker, Jean-Michel Dubois, Jean-loup Gailly, Hunter Goatley, Ian Gorman, Chris Herborth, Dirk Haase, Greg Hartwig, Robert Heath, Jonathan Hudson, Paul Kienitz, David Kirschbaum, Johnny Lee, Onno van der Linden, Igor Mandrichenko, Steve P. Miller, Sergio Monesi, Keith Owens, George Petrov, Greg Roelofs, Kai Uwe Rommel, Steve Salisbury, Dave Smith, Christian Spieler, Antoine Verheijen, Paul von Behren, Rich Wales, Mike White

This software is provided "as is," without warranty of any kind, express or implied. In no event shall Info-ZIP or its contributors be held liable for any direct, indirect, incidental, special or consequential damages arising out of the use of or inability to use this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. Redistributions of source code must retain the above copyright notice, definition, disclaimer, and this list of conditions.
2. Redistributions in binary form (compiled executables) must reproduce the above copyright notice, definition, disclaimer, and this list of conditions in documentation and/or other materials provided with the distribution. The sole exception to this condition is redistribution of a standard UnZipSFX binary as part of a self-extracting archive; that is permitted without inclusion of this license, as long as the normal UnZipSFX banner has not been removed from the binary or disabled.

3. Altered versions--including, but not limited to, ports to new operating systems, existing ports with new graphical interfaces, and dynamic, shared, or static library versions--must be plainly marked as such and must not be misrepresented as being the original source. Such altered versions also must not be misrepresented as being Info-ZIP releases--including, but not limited to, labeling of the altered versions with the names "Info-ZIP" (or any variation thereof, including, but not limited to, different capitalizations), "Pocket UnZip," "WiZ" or "MacZip" without the explicit permission of Info-ZIP. Such altered versions are further prohibited from misrepresentative use of the Zip-Bugs or Info-ZIP e-mail addresses or of the Info-ZIP URL(s).
4. Info-ZIP retains the right to use the names "Info-ZIP," "Zip," "UnZip," "UnZipSFX," "WiZ," "Pocket UnZip," "Pocket Zip," and "MacZip" for its own source and binary releases.

util-linux

Copyright (c) David L. Mills 1992, 1993, 1994, 1995, 1996. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

Copyright (c) 1980, 1987, 1989, 1993, 1994 The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: (1) Redistributions of source code or in binary form must retain the above copyright notice, this list of conditions and the following disclaimer (THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)

HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE); (2) All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the University of California, Berkeley and its contributors; and (3) Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

xntp

Copyright © David L. Mills 1992, 1993, 1994, 1995, 1996

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. The University of Delaware makes no representations about the suitability this software for any purpose. It is provided “as is” without express or implied warranty.

13 Glossary of Abbreviations

CHAPTER

A—Ampere. Unit of electric current.

AC—Alternating Current. Electric current that reverses direction, as opposed to direct current, which always flows in the same direction. In most countries, the household electric current is AC.

ARP (pronounced “arp”)—Address Resolution Protocol. ARP allows a device to find the physical address of a target device on the same physical network, given the IP address of the target device. An ARP request is broadcast to all devices on the same physical network, but only the target device replies with its IP address. Each device that uses ARP has an ARP cache of recently acquired IP-to-physical address bindings or pairings.

ATM—Asynchronous Transfer Mode. A networking technology with the capacity to transmit voice and video in real time as well as data, including frame relay traffic.

Note: ATM also stands for Automated Teller Machine. Network Discovery has a device icon dedicated to point-of-sale and automated tellers machines, POS/ATM.

CD—Compact Disc. A metal disc with a plastic coating. The disc is small enough to be held in the hand. A compact disc is used for storing digital data. Network Discovery software and licenses are supplied on compact discs.

- CIDR**—Classless Inter-Domain Routing. A format that allows you to abbreviate network mask (for example, 16) instead of typing it out (for example, 255.255.0.0).
- CIR**—Committed Information Rate. In a frame relay network, the bandwidth associated with a virtual circuit. The higher the CIR, the more priority given to the traffic for that circuit.
- CPU**—Central Processing Unit. A part of a computer that interprets and carries out instructions, usually a microprocessor chip.
- CRC**—Cyclic Redundancy Check, or Cyclic Redundancy Code. Cyclic redundancy checking is a method of examining data for errors by performing a computation on the data both before and after it is sent, and verifying that the computation yielded the same result each time.
- CSV**—Comma Separated Value. CSV files contain values from a spreadsheet, table, or database with each value separated by a comma. CSV files are extremely transportable—that is, they can easily be used by many kinds of software on many kinds of computers.
- DC**—Direct Current. Electric current that always flows in the same direction, as opposed to alternating current.
- DHCP**—Dynamic Host Configuration Protocol. DHCP assigns IP addresses to devices automatically, and can reassign them dynamically when there are more devices than there are IP addresses available. It also helps to manage IP addresses by having one location from which they are tracked and assigned.
- DIN**—Deutsche Industrie Norm. DIN is a German standards organization. A cable that has a DIN connector (for example, DIN-6) conforms to these standards.
- DLCI**—Data Link Connection Identifier. Part of the header in frame relay, the DLCI is used to route the frame.

DNS—Domain Name System. DNS is the system whereby domain names—such as “starter.example.com”—are first located (usually on a DNS server) and then translated into the less ambiguous but more difficult to remember IP addresses—such as “192.168.2.129”. A DNS server is a computer that exists primarily to maintain a listing of which domain names correspond to which IP addresses. What this means to you is that you are allowed to work with and remember names, which are more likely to be meaningful than collections of numbers.

DTE—Data Terminal Equipment. Any device that can transmit digital information over a cable—for example, a microcomputer workstation. One of two types of computer hardware connected by an RS-232-C connection. The other type is DCE, or Data Communications Equipment—for example, a switch or modem.

ESD—ElectroStatic Discharge. The release of static electricity. Static electricity can damage or even destroy electronic equipment.

FAQ—Frequently Asked Questions—*see Knowledge Base*.

FCS—Frame Check Sequence. A method of checking the integrity of a frame. A FCS error indicates that the frame has somehow become corrupted, since the frame failed its CRC.

FDDI (sometimes pronounced “*fid-dee*”)—Fiber Distributed Data Interface. FDDI is a set of rules for sending and receiving data on fiber optic lines to a local area network (LAN). FDDI is based on the token ring protocol, but uses two tokens instead of one. FDDI networks have a range of 124 miles / 200 km.

FS—File System. A file system is concerned with the naming and storing/retrieving of files, and comprises files (collections of data) and directories (collections of files).

FTP—File Transfer Protocol. FTP is a method for sending and receiving files from one place in a network to another.

HTML—HyperText Mark-up Language. HTML is a documentation standard intended to enhance the display of a document in a World Wide Web browser such as Netscape or Internet Explorer. For example, HTML codes for displaying subscript text (as in H₂O) look like this: “H₂>O”

HTTP—HyperText Transfer Protocol. HTTP is a method for exchanging files on the World Wide Web.

HSRP—Hot Standby Routing Protocol. A routing protocol that allows more than one router to act as a single virtual router. If one router fails, the next router assumes its identity immediately. As a result, as far as the rest of the network is concerned, the virtual router is still working.

Hz—Hertz. A unit of frequency of one cycle per second. This unit of measure is named for German physicist Heinrich Hertz.

ICMP—Internet Control Message Protocol. ICMP provides communication between the Internet Protocol (IP) software on one machine with the IP software on another. It is a simple protocol (or “set of rules and standards”) that every IP-based device must support. ICMP is used to communicate control, information, and error messages among IP devices. Probably the best known ICMP messages are the echo request and echo reply messages of a ping.

IDE—Integrated Device Electronics. An interface for disk drives. Early Peregrine appliances used IDE drives. See also SCSI.

IE—Internet Explorer. Microsoft’s Web browser software. Network Discovery is compatible with Internet Explorer and Netscape.

IP—Internet Protocol. The Internet Protocol (IP) handles the address part of each data packet that is transmitted from one computer to another on the Internet.

When you see the term “IP address” with no qualifiers in Network Discovery, it means that either a version 4 IP address (IPv4) or a version 6 IP address (IPv6) is acceptable.

- IPv4 address

An IPv4 address contains four sections separated by periods (or “dots”). Each section, called an octet, contains 8 bits expressed in decimal (0–255).

Example: 192.168.2.129

- IPv6 address

An IPv6 address contains eight sections separated by colons. Each section contains 16 bits expressed in hexadecimal (0000–FFFF).

Example: 1234:5678:9ABC:DEF0:1234:5678:9ABC:DEF0

Knowledge Base—The Knowledge Base at Peregrine Systems Customer Support has answers to questions customers have asked. It is available from the Help menu. The Knowledge Base covers topics best addressed in question-and-answer format rather than through conventional documentation, and acts as a catch-all place to describe quirks and common misunderstandings.

LAN (pronounced “lan”)—Local Area Network. A LAN is a network of workstations sharing resources, usually within a restricted geographic area, such as an office building. LANs typically serve tens or hundreds of people rather than thousands. If your network has a single central site, you probably have a LAN. The main LAN technologies are Ethernet, token ring, and FDDI.

LED—Light Emitting Diode. A small light, sometimes round in shape. This term is also used for part of the Network Discovery user interface: a colored circle used to indicate alarm state—visible, for example, on the Health Panel.

LSN—Logical SubNet. A subnet (short for “sub-network”) is a segment of a network. A logical subnet is a segment organized not by geography (where a device physically resides) but by netmask (short for “network mask”).

MAC (pronounced “mac”)—Media Access Control. A MAC address is a computer’s unique hardware number. A MAC address looks like this—0040E5010025 —or like this—00 40 E5 01 00 25 (spaces added for readability). The six numbers are hexadecimal (base 16) values.

MB—MegaByte. A measurement of capacity, applied to such computer components as memory and disk storage.

MIB (pronounced “mib”)—Management Information Base. A MIB is a collection of data that can be read and written using a network management protocol such as SNMP. The MIB is structured as a hierarchy of “objects”. There are both standard MIBs (supported by many vendors) and proprietary MIBs (vendor-specific).

MTBF—Mean Time Between Failures. The average time that a device is operational. In Network Discovery, MTBF is measured in days. Devices that break frequently can cause you aggravation. In this context, the term refers to network devices such as switches and workstations, but both MTBF and MTTR could equally well refer to an automobile or telephone.

MTTR—Mean Time To Repair. The average time it takes to repair a device. In Network Discovery, MTTR is measured in hours. Devices that take a long time to repair can cause you considerable concern, particularly if they are important to the operation of your network.

NAT—Network Address Translation. NAT is an Internet standard that enables a local-area network to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. NAT servers have two main purposes: hiding private IP addresses from external addresses, and enabling a private network to use more private IP addresses.

NTP—Network Time Protocol. An Internet standard used to synchronize the clock of a computing device to a time server with a degree of accuracy of milliseconds. The Peregrine appliance can take advantage of NTP to set its internal clock accurately.

ODBC—Open Data Base Connectivity. An open standard for accessing a database.

OID—Object IDentifier. The number which identifies an object in a MIB. MIBs are made up of objects. Each object in a MIB has unique object ID, which is a series of numbers separated by dots. For example, the OID of system.sysName is “.1.3.6.1.2.1.1.5”. This ID defines the location of that object within the MIB tree hierarchy.

OS—Operating System. A core computer program that manages all application programs. Often, OS is used as synonymous with DOS, or Disk Operating System.

OSI—Open Systems Interconnection. Usually in reference to the OSI network model. The OSI model has seven layers. Layers 2 and 3 are the most important to Network Discovery.

- OUI** (sometimes pronounced “ow-ee”)—Organization Unique Identifier. The OUI is the first three octets of a MAC address, and it identifies the organization that manufactures the device associated with the MAC address. For example, in the MAC address “00 40 E5 01 00 25”, the actual OUI is “00 40 E5”. If the OUI is one that Network Discovery recognizes, then the first three numeric octets are replaced by either an abbreviation of the organization’s name or the full organization name, depending on the context. For example, the Device Manager represents the MAC address “00 40 E5 01 00 25” as “PRGRIN 010025”. Although Network Discovery has an extensive database of OUIs, there may be some it doesn’t recognize, in which case all octets of the OUI are displayed numerically.
- POS**—Point of Sale. The point of sale is the place in a store where purchases are made. Point of sale devices include electronic cash registers and debit card readers.
- PVC**—Permanent Virtual Circuit. A logical (rather than physical) connection in a network, particularly in a frame relay network. With a PVC, you define connection points and let someone else worry about how the data physically moves from one point to the other.
- RAM**—Random Access Memory. Memory that can be used to store data. Different from ROM, read-only memory.
- RH**—Relative Humidity. The amount of water vapor actually in the air divided by the maximum amount of water vapor the air can hold at its current temperature.
- RJ**—Registered Jack. Modular wiring (receptacles and plugs) used to connect equipment over telephone lines. Examples are RJ-11 and RJ-45.
- RS**—Recommended Standard. The Electronic Industries Association has adopted such standards as RS-232 for serial communications.
- SCSI**—Small Computer System Interface. An interface for disk drives and other peripheral devices. Peregrine appliance use SCSI hard disk drives.
- SMT**—Station Management. This FDDI module operates at the data link and physical layers to monitor and manage both the FDDI ring and the devices on it.

SMTP—Simple Mail Transfer Protocol. SMTP is a set of rules concerning the sending and receiving of electronic mail (e-mail). An SMTP server is a device that exists primarily to perform the service of directing e-mail.

SNMP—Simple Network Management Protocol. SNMP is a set of rules that allow networks to be managed and devices on those network to be examined. This set of rules is a broadly accepted and implemented open standard. A device on which SNMP can perform actions (such as reading and writing the device's MIB) is said to be “managed”. The chief benefit of SNMP is that it allows network managers to administer networks and devices remotely; that is, without having to physically locate and adjust each device.

Network Discovery uses SNMP to obtain information about your network and the devices within it. The SNMP standard allows Network Discovery to obtain this information in a manner that is independent of a specific network device implementation or its vendor.

TCP—Transmission Control Protocol. TCP is a communications protocol that sends data between network devices in the form of message units, or packets. TCP divides the data into packets at one end, and reassembles the packets once they arrive.

TCP/IP—Transmission Control Protocol/Internet Protocol. TCP/IP is the pairing of two protocols, TCP and IP. TCP/IP forms the basic communication language of the Internet. TCP/IP is not a program that you use; it is a pair of protocols required by programs that you use. For example, FTP, HTTP, SMTP, and Telnet use TCP/IP to make their connections.

U—Unit. Unit of measurement for rackmount equipment (U is 1.75in or 4.44cm)

UDP—User Datagram Protocol. An alternative communications protocol to TCP. Useful for network applications that have small data units to exchange.

UPS—Uninterruptible Power Supply. A piece of equipment that connects a device to a power source so that the principal source provides power when all is well, and the UPS's secondary source—a battery—provides power when the primary power source is not available, such as a blackout. Peregrine Systems strongly recommends the use of an uninterruptible power supply with the Peregrine appliance.

URL (sometimes pronounced “erl”)—Uniform Resource Locator or Universal Resource Locator. The address of a file accessible through the World Wide Web. A URL looks like this: “http://www.example.com”. Sometimes the “http://” is left off, and the URL is given simply in the form “www.example.com”.

VA—Volt-Ampere. A measurement of power in an AC circuit.

VAC—Volts Alternating Current. Measurement of voltage swing.

VM—Virtual Machine. Any software that imitates the performance of a hardware device, such as a CPU.

WAN (pronounced “wan”)—Wide Area Network. A WAN is essentially a network like a LAN, except with a broader structure and larger geographic area. If your network has a single central site, but also one or more remote sites (such as sales offices in other parts of the country), you probably have a WAN.

WWW—World Wide Web. If you create a file that can be transmitted using HyperText Transport Protocol (HTTP), and put it where others can view it (by having their web browser temporarily transfer and display a copy of the file), then you are part of the World Wide Web.

Index

Numerics

- 24 hours
 - bridge aging interval 51
 - not seen 80

A

- About
 - Line Manager panel 132
- address types (definitions) 48
- aging interval 51
- Alarm Type panel (Port Manager) 127
- Approximate 59
- approximate connection 122
- ATM cells 71
- Attribute Manager 137–144
 - Configuration 140
 - Locate button 142
 - Manage 142
 - panel elements 67
 - Statistics 142
- autosave 43

B

- Backplane Utilization 77
- banner
 - Device Manager 68
- brackets around numerals 67
- Break Connection panel
 - Line Manager 134
 - Port Manager 126
- break fault 63

- breaking a connection 125
- Breaks
 - Device *see* Device Breaks
 - Line *see* Line Breaks
- bridge aging 51
- Broadcasts 75, 162
- broken line *see* Line Breaks
- Browse MIB
 - Device Manager button 107

C

- capacity, line 71
- Carrier Network 59
- CenterPoint Web site 181
- CIDR notation 50
- CIR, redefining 35
- clearing the database 95, 121
- Close
 - button
 - Attribute Manager 144
 - Device Manager 112
 - Line Manager 135
 - Port Manager 128
- cloud icon 59
- clouds 58
 - carrier network icon 59
 - cloud icon 59
 - radio cloud icon 59
 - unmanaged hub icon 59
- clouds *see also* virtual device
- Collisions 74, 103, 162

- comma separated value. See CSV
- communication models
 - FDDI 36–38
 - frame relay 33–36
- community string
 - definition 51
- community strings
 - in MIB Browser 152
- Configuration
 - Attribute Manager panel 140
 - Device Manager panel 85
 - Port Manager panel 116
- configuration (map)
 - autosave 43
 - Prime 43
- connection
 - breaking single (conceptual) 126, 134
 - forcing single new (conceptual) 125
 - types of 122
- contact, system 64, 87
- CPU Utilization 76
- Create Connection panel (Port Manager) 125
- creating a connection 125
- CSV 16, 102, 123, 142, 167
- customer support, how to contact 181

D

- data
 - clearing 95, 121
- default map configuration 42, 167, 168
- definitions
 - Network Discovery terms and concepts 53–64
 - network terms and concepts 48–53
- Delay 74, 162
- device
 - definition 55
 - icon 15, 53, 82, 116
 - model
 - see also Update Model
 - virtual 58
 - virtual devices 58
- Device Adds/Deletes 78
- Device Breaks 22, 162
- Device Changes 79

- Device Manager
 - banner 68
 - Browse MIB 107
 - Configuration 85
 - Diagnosis 95
 - DNS Query 101
 - Events 105
 - Export (Statistics) 103
 - Graph (Statistics) 102
 - IP Ping 99
 - Locate 106
 - Manage 106
 - MIB Browser 107
 - panel elements 67
 - Ports 105
 - Properties 110
 - Scan Data 107
 - Service Analyzer 106
 - SNMP Ping 100
 - State 92
 - Statistics 102–104
 - Table (Statistics) 103
 - Telnet 108
 - title 68
 - toolbar 83
 - Traceroute 99
 - Update Model 108
 - Web 107
- Device Moves 79
- Device MTBF 78
- Device MTTR 78
- Device Reports 167
- Diagnosis
 - Device Manager panel 95
 - Port Manager panel 121
- diamonds 59
 - approximate icon 59
 - logical view 60
 - LV unmapped icon 60
 - LV unmapped IP icon 60
 - shared port icon 60
 - unmapped IP icon 59
- directed community string 51
- disconnecting
 - Peregrine appliance 41

- disconnecting a port or line 126, 134
- Disk Utilization 76
- DNS Query button (Device Manager) 101
- domain name
 - definition 48
- Duplex Mode panel (Port Manager) 127

E

- Entry 149
- Errors 75, 162
- event thread 66
- Events
 - Device Manager panel 105
 - Port Manager panel 124
- Exceptions 79
- Executive/Summary Reports 160
- Export
 - Device Manager Statistics button 103

F

- FDDI 36–38
- Find
 - MIB Browser 149
- Find OID 151
- Folder tab 149
- forcing a connection 125
- Frame Relay 75
- frame relay 33–36

G

- Get 151
- Get Next 149, 154
- Graph
 - Device Manager button 102
- gray background
 - Manager data 67, 92, 118

H

- Health Panel
 - Backplane Utilization 77
 - Broadcasts 75
 - Collisions 74
 - CPU Utilization 76
 - Delay 74
 - Device Adds/Deletes 78

- Device Breaks 22
- Device Changes 79
- Device Moves 79
- Device MTBF 78
- Device MTTR 78
- Disk Utilization 76
- Errors 75
- Exceptions 79
- Frame Relay 75
- Line Breaks 71
- Load Average 76
- Memory Utilization 77
- Not Recently Seen 80
- Open Tickets 80
- Packet Loss 76
- Port Add/Deletes 77
- Port Changes 78
- Port Moves 78
- Port MTBF 77
- Port MTTR 77
- Printer 77
- UPS 77
 - Utilization 71
- HSRP (Hot Standby Router Protocol) 39
- HTML in SNMP MIB 64
- HTTP session 107

I

- icons
 - question mark 57
 - virtual devices 58
- Interface Rate panel (Port Manager) 126
- Interface Type panel (Port Manager) 126
- IP address
 - definition 48
 - multiple 91
- IP Ping
 - Device Manager button 99
- IPv4 address (definition) 48, 208
- IPv6 address (definition) 49, 208

L

- LAN Reports 166
- layers 2 and 3 (OSI) 52
- line (definition) 64

- Line Breaks 71, 162
- line capacity 71
- Line Manager 131
 - About 132
 - Break Connection 134
 - panel elements 67
 - toolbar 132
- line priority 64
- line, multiple 135
- link training 122
- Load Average 76
- Locate
 - MIB Browser button 149
- Locate button
 - Attribute Manager 142
 - Device Manager 106
 - Port Manager 124
- location, system 87
- logical subnet 122
- Logical View
 - icon 60
- LV Unmapped 60
- LV Unmapped IP 60

M

- MAC address (definition) 49
 - numeric 49
 - with OUI 49
- Manage
 - Attribute Manager button 142
 - Device Manager button 106
- management workstation requirements 53
- map configuration 41
 - autosave 43
 - default 42
- map objects (definition) 53
- map session 41
- mask, network *see* netmask
- mean peak statistics 159
- mean statistic 159
- Memory Utilization 77
- MIB Browser 145
 - via Device Manager 107
 - Entry 149
 - find a device 149

- Find OID 151
- Folder tab 149
- Get 151
- Get Next 149, 154
- locate 149
- MIB description 152
- MIB radar 155
- opening 146
- pull-down list of devices 149
- refresh 149
- set 151
- tree view 148
- Variable tab 151
- write 151
- Microsoft Word documents 168
- multiple IP addresses 91
- multiple lines 135

N

- name, system 87
- negative statistics 67, 92, 118
- netmask notation (definition) 50
- Network Map
 - autosave 43
 - configurations *see* map configuration
 - Prime configuration 43
 - question mark 57
 - sessions *see* map session
- Network Mapper 20, 40
- New MIB Browser
 - MIB Browser
 - New MIB Browser 154
- Not Recently Seen 80

O

- object properties 110
- objects 53–61
- Open Tickets 80
- Organization Unique Identifier (OUI) 49
- OSI model 52
- OUI 49

P

- package
 - definition 61

- Packet Loss 76, 162
- parentheses around numerals 67
- PC requirements 53
- peak statistics 159
- period, sampling 76, 79
- ping button (Device Manager) 99
- Port Add/Deletes 77
- Port Changes 78
- port index list box (Port Manager) 129
- Port Manager
 - Alarm Type 127
 - Break Connection 126
 - Configuration 116
 - Create Connection 125
 - Diagnosis 121
 - Duplex Mode 127
 - Events 124
 - Interface Rate 126
 - Interface Type 126
 - Locate button 124
 - panel elements 67
 - port index 129
 - Purge Port 124
 - State 118
 - toolbar 115
- Port Moves 78
- Port MTBF 77
- Port MTTR 77
- Ports panel (Device Manager) 105
- Prime configuration 43
- Print button
 - Attribute Manager 143
 - Device Manager 111
 - Line Manager 134
 - Port Manager 128
- Printer 77
- priority
 - definition 63
 - device
 - definition 63
 - display 98
 - Prime 98
 - range 63
 - line
 - definition 64

- Properties
 - Device Manager button 110
- properties
 - object 110
- Purge Port (Port Manager) 124

Q

- question mark 57

R

- Radar 155
- Radio Cloud 59
- real cloud *see* virtual device
- recorded events
 - Backplane Utilization 77
 - Broadcasts 75
 - Collisions 74
 - CPU Utilization 76
 - Delay 74
 - Device Adds/Deletes 78
 - Device Breaks 22
 - Device Changes 79
 - Device Moves 79
 - Device MTBF 78
 - Device MTTR 78
 - Disk Utilization 76
 - Errors 75
 - Exceptions 79
 - Frame Relay 75
 - Line Breaks 71
 - Load Average 76
 - Memory Utilization 77
 - Not Recently Seen 80
 - Open Tickets 80
 - Packet Loss 76
 - Port Add/Deletes 77
 - Port Changes 78
 - Port Moves 78
 - Port MTBF 77
 - Port MTTR 77
 - Printer 77
 - UPS 77
 - Utilization 71
- Refresh button
 - Attribute Manager 143

- Device Manager 111
- Line Manager 134
- Port Manager 128
- Reports
 - business
 - device 167
 - executive/summary 160
 - LAN 166
 - WAN 164
 - periods 157
 - statistics 159
- RFCs 32
- S**
- sampling period 76, 79, 125
- Scan Data
 - Device Manager button 107
- schedule, Network Discovery 45
- Service Analyzer
 - button (Device Manager) 106
- Set 151
- Shared Port 60
- SMT 37
- SNMP Ping
 - Device Manager button 100
- source address capture 122
- special input syntax 64
- speed, line
 - see* Interface Rate
- spreadsheets, exporting to. *See* CSV
- stale data 67, 92, 118
- State panel
 - Device Manager 92
 - Port Manager 118
- Statistics
 - Attribute Manager panel 142
 - Device Manager panel 102–104
- support, how to contact 181
- system contact 64, 87
- system location 87
- system name 87

T

- Table
 - Device Manager button 103

- table-based connection 122
- Telnet
 - Device Manager button 108
- terminology definitions
 - address types 48
 - community string 51
 - devices (map) 55
 - domain name 48
 - IP address 48
 - IPv4 address 48, 208
 - IPv6 address 49, 208
 - layers 2 and 3 (OSI) 52
 - MAC address (numeric) 49
 - MAC address (with OUI) 49
 - map objects 53
 - netmask notation 50
 - network 48–53
 - Network Discovery 53–64
 - OSI 52
 - package 61
 - priority 63, 64
 - special input syntax 64
- Text button
 - Attribute Manager 143
 - Device Manager 112
 - Line Manager 134
 - Port Manager 128
- title
 - icon 68
 - object 68
- title bar *see* banner
- toolbar
 - Device Manager 83
 - Line Manager 132
 - Port Manager 115
- Traceroute
 - Device Manager button 99
- traffic-based connection 122
- U**
- Unmanaged Hub 59
- Unmapped IP 59
- Update Model button (Device Manager) 108
- UPS 77
- URL in SNMP MIB 64

Utilization 71

V

Variable Tab 151

virtual cloud *see* virtual device

virtual device 58

 clouds 58

 creating 125

 diamonds 59

W

WAN Reports 164

Web

 Device Manager button 107

Word documents 168

Write 151



July 23, 2004