

Performance Manager for Tru64 UNIX
Version 5.0A

User's Guide

© 1999 Compaq Computer Corporation

COMPAQ, the Compaq logo, and the Digital logo are Registered in the U.S. Patent and Trademark Office.

Microsoft and Windows NT are registered trademarks of Microsoft Corporation. Intel, Pentium, and Intel Inside are registered trademarks of Intel Corporation. UNIX is a registered trademark and The Open Group is a trademark of The Open Group in the US and other countries. Other product names mentioned herein may be the trademarks of their respective companies.

Possession, use, or copying of the software described in this publication is authorized only pursuant to a valid written license from Compaq Computer Corporation or an authorized sublicensor.

Compaq Computer Corporation shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change without notice.

Revision/Update Information: This is an update to the Version 4.0x manuals.

Operating System and Version: Tru64 UNIX Version 5.0A. Future maintenance releases may require higher versions.

Software Version: Performance Manager Version 5.0A

Date: March 2000

Contents

Preface	vii
Structure of This Document	vii
Related Information	viii
Related Manuals	viii
Conventions	viii
1 Overview	1
Management Station	1
Managed Node	2
Metrics Server Information	2
2 Getting Started	3
Starting Performance Manager	3
Exiting Performance Manager	3
Displaying the Performance Manager GUI	4
Setting the DISPLAY Environmental Variable	4
Displaying Performance Manager on a PC	4
Main Window Overview	4
The Performance Manager Main Window	5
Work Area	5
Icons	6
Main Window Toolbar and Menu Bar	8
Modifying the Main Window	13
Performance Manager Popup Menu	13
3 Managing Nodes	15
Creating Groups	15
Deleting Groups	16
Adding Nodes	16
Deleting Nodes	16
Moving Nodes	17
Adding Clusters	17
Deleting Clusters	17
Moving Clusters	18

4	Displaying Clusters	19
	Auto-Discovery for Clusters	19
	Display Representation of Clusters	19
	Possible Anomalies for TruCluster Production Server	20
5	Monitoring	21
	Sessions	21
	Creating a Session	22
	Managing Sessions	23
	Displays	24
	Floating Displays	26
	Consolidating Displays	26
	Manipulating Displays	27
	Setting Display Styles	27
	Other Monitoring Methods	29
	Monitoring from the Command Line	29
	Monitoring with SNMP Network Management Systems	30
6	Metrics	33
	Displaying Metrics	33
	Showing Hidden Metric Categories	33
	Hiding Metric Categories	34
7	Thresholds	35
	Threshold Notifications	36
	Setting Thresholds	36
	CPU Thresholds	37
	System Thresholds	37
	Processes Thresholds	37
	Buffer Cache Thresholds	37
	Network Thresholds	37
	File System Thresholds	37
	Memory Thresholds	38
	AdvFS Thresholds	38
	TruCluster Thresholds	38
	Environmental Thresholds	38
	Advanced Threshold (more...) Dialog Box	38
	Threshold Environment Variables	39
8	Commands	41
	Performance Analysis Commands	41
	CPU Commands	41
	Memory Commands	41
	Network Commands	42
	Disk I/O Commands	42
	System Management Commands	43
	Cluster Performance Analysis Commands	46
	Threshold Management Commands	47
	AdvFS Performance Analysis Commands	47

Command Operations	47
Executing Commands	48
Adding Commands to the Execute Menu	49
Deleting Commands from the Execute Menu	50
Modifying Commands	50
Adding Command Categories	50
Deleting Command Categories	50
Moving Commands Between Categories	51

9 Archiving 53

 Archive Recording 53

 Archive Playing 54

10 Troubleshooting 55

 Log Files 55

 Example Log File Entry 56

 Nodes Not Responding 56

 Metrics Servers or GUI Will Not Start 57

 No Log Files 58

 No Startup Files 58

 Commands Not Running 58

 Disks Not Visible to Performance Manager 58

 Reporting Bugs 58

 Software Performance Reports 59

Glossary 61

Index 63

Preface

Compaq® Performance Manager for Tru64™ UNIX® is an SNMP-based, user-extensible, real-time performance monitoring and management tool that allows you to detect and correct performance problems from a central location. Performance Manager has a graphical user interface (GUI) called `pmgr` that runs locally and can display data from the managed nodes in your Compaq Tru64 UNIX network. Performance Manager operates through interaction between nodes assigned as management stations and managed nodes.

Note It is possible for a managed node to also be the management station. For more information on management stations and managed nodes, read the Overview chapter.

Performance Manager is an optional subset of Tru64 UNIX (formerly known as DIGITAL™ UNIX).

Performance Manager for Tru64 UNIX comprises two primary components: Performance Manager GUI (`pmgr`) and Performance Manager metrics server (`pmgrd`). Additional metrics servers are used in monitoring Compaq TruCluster™ systems (`clu_mib`) and Advanced File System (`advfsd`), supplied in the AdvFS Utilities subset.

Structure of This Document

This manual includes the following chapters, followed by a glossary and an index:

- Chapter 1, Overview, provides a general description of Performance Manager's purpose and capabilities.
- Chapter 2, Getting Started, describes setting up the environment, learning the terminology, and using the interface.
- Chapter 3, Managing Nodes, describes using Performance Manager to manage and monitor the nodes in your network.
- Chapter 4, Displaying Clusters, describes how Performance Manager displays clusters using auto-discovery.
- Chapter 5, Monitoring, describes creating, saving, and recalling sessions for monitoring data in real time, and customizing displays.
- Chapter 6, Metrics, describes arranging your metrics in categories, and choosing which metrics to display or hide.
- Chapter 7, Thresholds, describes limits you can set on metrics. Crossing these thresholds triggers an alert, notifying you of computer or network problems.
- Chapter 8, Commands, describes running commands with Performance Manager (its own or yours) on remote nodes and displaying the results.

- Chapter 9, Archives, describes Performance Manager scripts that enable storing files of performance data.
- Chapter 10, Troubleshooting, describes creating log files, restarting daemons, solving problems, and reporting problems.
- Glossary describes terms specific to Performance Manager.
- Index

Related Information

In addition to this guide, you can use the following manuals and documents to learn more about Performance Manager:

- *Performance Manager Installation Guide*
- *Performance Manager Release Notes*
- *Performance Manager Web Site*

For updates and the latest information about Performance Manager, see the PM web site at this URL:
<http://www.unix.digital.com/performance-manager/>

Related Manuals

The following manuals are part of the base operating system documentation set and may help you with your use of Performance Manager:

- *Tru64 UNIX Installation Guide*
- *Tru64 UNIX Software License Management*

Conventions

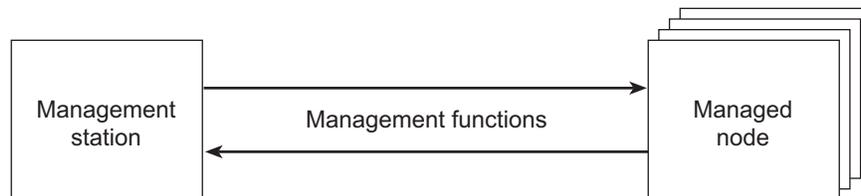
The following conventions are used in this guide:

Convention	Meaning
UPPERCASE and lowercase	The Tru64 UNIX system differentiates between lowercase and uppercase characters. Literal strings that appear in text, examples, syntax descriptions, and function descriptions must be entered exactly as shown.
<i>variable</i>	This italic typeface indicates system variables.
user input	This bold typeface is used in interactive examples to indicate input entered by the user.
system output	This typeface is used in code examples and other screen displays. In text, this typeface indicates the exact name of a command, option, partition, path name, directory, or file.
%	The percent sign is the default user prompt.
#	A number sign is the default root user prompt.
Ctrl/X	In procedures, a sequence such as Ctrl/X indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.

Chapter 1 Overview

Performance Manager interacts between nodes assigned as *management stations* and *managed nodes*. Their features are described in the following sections.

Figure 1 PM Overview



Management Station

Management stations are the operating centers for managing and monitoring the nodes in the system. With Performance Manager, you can monitor the state of one or more managed nodes in real-time. Tables and graphs, such as plot, area, bar, stack bar, and pie charts, show you hundreds of different system values, including:

- CPU performance
- Memory usage
- Disk transfers
- File-system capacity
- Network efficiency
- AdvFS-specific metrics
- Cluster-specific metrics

In addition to monitoring, Performance Manager provides these features for actively managing your network:

- **Thresholding:** Thresholds can be set to alert you when a potential problem occurs by triggering a response when a threshold is crossed. This response can be notification through a GUI window, an email, pager, or FAX message, or the response can be an actual command execution for system management or archiving.
- **Archiving:** Metrics can be archived to a file and then played back, showing resource usage trends and historical analysis. Performance Manager includes these archiving scripts: `pm_archiver`, `pm_delta_archiver`, and `rc_archiver`.

- **Commands:** Performance analysis, system management, and/or cluster analysis and AdvFS commands (yours and those supplied with Performance Manager) can be run simultaneously on multiple nodes using the GUI.
- **For analysis:** You can run commands that analyze the state of managed nodes. Commands can be run on the management station or on the managed nodes.
- **To take actions:** You can run commands that take actions on managed nodes from the management station.
- You can add your own administration tasks to the extensible GUI.

Managed Node

Managed nodes are those that run one or more *metrics servers* recognized by Performance Manager. Cluster nodes are recognized and displayed as such. A metrics server is a daemon process that implements *management information base (MIB)* variables that the Performance Manager GUI knows about.

A metrics server listens for and services requests for operating system metric information. These requests are issued by management applications such as the Performance Manager GUI. Upon receipt of such a request, a metrics server queries the operating system and returns the appropriate value(s) to the requester. The following are examples of metrics servers supported by Performance Manager:

- `pmgrd` — Provides general Tru64 UNIX metrics
- `clu_mib` — Provides TruCluster-related metrics
- `os_mibs` — Provides MIB-II metrics
- `svrclu_mib` — provides common cluster metrics
- `advfsd` — Provides AdvFS-related metrics

The `pmgrd` metrics server is provided with Performance Manager. Other metrics servers either come with the operating system (such as `os_mibs`) or are provided by other products (such as `advfsd`).

PM-provided metrics servers are subagents of the Tru64 UNIX extensible SNMP agent (`snmpd`). In addition, they support extensions for bulk data transfer of metric data. Because metrics servers support SNMP, you can use other SNMP applications to access their data. In addition, a set of UNIX commands for command-line metrics server access is provided.

The nodes and metrics you choose to monitor can be saved as a session, then played back or modified later.

Metrics Server Information

Chapter 10, Troubleshooting, contains information on server startup, possible problems, and references to more detailed information.

Chapter 2

Getting Started

This chapter tells how to start and exit Performance Manager, and explains the GUI's main window.

Starting Performance Manager

Log in to a node where Performance Manager has been installed. If the `rehash` command has not been issued since Performance Manager was installed, type this command to recreate the internal command tables used by the shell:

```
# rehash
```

Before starting Performance Manager, be sure the `DISPLAY` environment variable on the starting system is set for the display you wish to use.

There are additional considerations if you wish to display Performance Manager on a PC. To start Performance Manager, issue the `/usr/bin/x11/pmgr` command at a root prompt (see the `pmgr(8)` reference page for details):

```
# /usr/bin/x11/pmgr
```

Performance Manager can be started from a non-root account, but the log file (`/var/opt/pm/1/pmgr_gui.log`) must first have its permissions changed to allow non-root users to write to it; for example, issue the following command as `root` to make the log file writable by everyone:

```
# chmod 666 /var/opt/pm/1/pmgr_gui.log
```

When Performance Manager starts, it opens its main window on the workstation defined by the `DISPLAY` environment variable.

Exiting Performance Manager

To exit Performance Manager, from the File menu, choose Exit. Your current session will not be saved when exiting. To save a session, choose Save Session or Save Session As from the main window's File menu. Save Session As opens a file selection dialog box.

Displaying the Performance Manager GUI

These topics explain how to display the Performance Manager GUI.

Setting the DISPLAY Environmental Variable

To set the DISPLAY environment variable in a C shell (`csh`), issue the following command, where *workstation* is the node name of your workstation:

```
setenv DISPLAY workstation:0.0
```

To set the DISPLAY environment variable in a Bourne shell (`sh`), issue the following commands, where *workstation* is the node name of your workstation:

```
DISPLAY=workstation:0.0
```

The system output will be as follows:

```
export DISPLAY
```

Note Your workstation should be a Tru64 UNIX node running the Common Desktop Environment (CDE). Nodes running other operating systems and other window managers might work, but only Tru64 UNIX and CDE have had full quality assurance testing for Performance Manager.

If you are running Performance Manager remotely, be sure your workstation supports the GUI display.

Displaying Performance Manager on a PC

Performance Manager can be displayed on most PCs. Either start Performance Manager through a PC X server program (such as Compaq eXcursion™), or start Performance Manager on a server node whose DISPLAY environment variable (in either the C shell or Bourne shell) is set to the PC. Either TCP/IP or DECnet™ will work, but consider the following when displaying Performance Manager on a PC:

- 1 The PC and the Tru64 UNIX server node must know about each other. The PC's network name and address must be in the server node's `/etc/hosts` or DUS database file (TCP/IP), or NCP/NCL database (DECnet). The server node's network name must be in the PC's TCP/IP file or NCP/NCL database (DECnet).
- 2 When starting Performance Manager on a PC using an X server program (such as eXcursion), there can be error messages that the X server program cannot report, such as your user name not being authorized to run Performance Manager, LMF license check failure, and so forth. To check for such errors, start Performance Manager on the server node after setting DISPLAY to the PC.
- 3 Depending on how your PC's resources are configured, it is possible to overload eXcursion by displaying too many applications, especially large ones such as Performance Manager (as compared to small ones such as `dxclock`, `dxterm`, and `dxcalendar`). Overloading an X server program can cause odd, nonintuitive errors. If you see such errors, try closing a few applications and restarting Performance Manager.

Main Window Overview

The main window is the first window you see when starting Performance Manager. This window consists of the menu bar, toolbar, nodes area, work area, message area, and Start Session and Stop Session buttons.

The nodes area, on the left side of the main window, displays icons for the nodes you can monitor. By default, the local node is displayed and belongs to the group World.

Clicking on a node, cluster, or group in Performance Manager's initial main window causes the work area to appear. The work area contains selection buttons for tasks and categories, and a scroll window for metric selection.

The message area displays status, warning, and error messages.

The Performance Manager Main Window

This is the opening window, and is the starting place for all your tasks.

Figure 2 Main Window



Work Area

Use the work area, on the right side of the main window, to configure displays and thresholds for nodes or clusters you have selected in the nodes area. Your view of the work area depends on whether you have selected the Display or Threshold buttons; each has a specific work area, showing related categories, metrics, and options.

Figure 3 Display Work Area

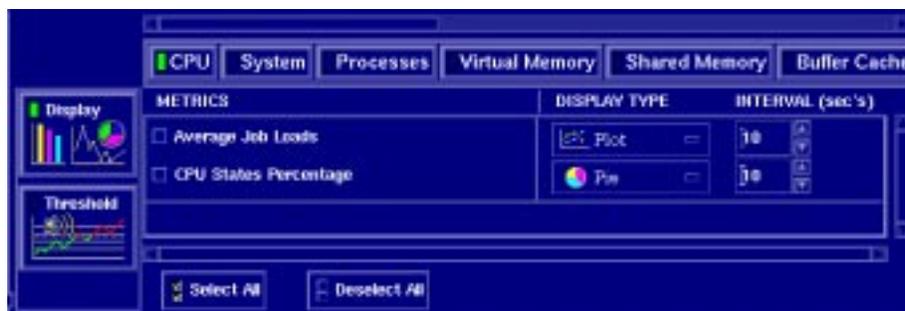
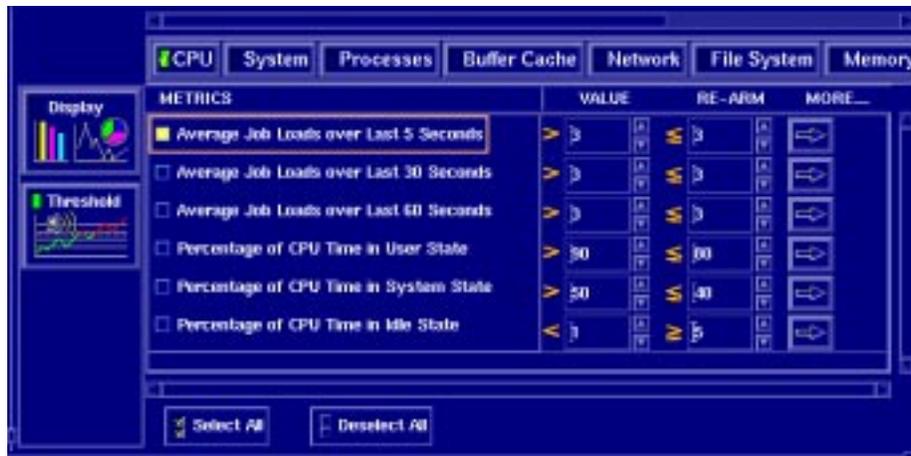


Figure 4 Threshold Work Area



Icons

The icons are sensitive. Click them to perform the operations in this section.

Main Window Icons

The nodes area, on the left side of the main window, displays icons for nodes you can monitor. By default, the local node is displayed and belongs to the group World.

To manage the nodes, clusters, and groups appearing in the nodes area, use the toolbar or go to the main window's Tasks menu and choose Node Management.

Nodes

A node is a computer system that is uniquely addressable on a network. A node can have more than one CPU. Single globes represent individual nodes in various states. Note that a node icon may take a few moments to reflect the state of the node after the node is newly added or comes up. A node icon changes to reflect one of the following three node states:



Hand is holding world down: Node is down or invalid.



Hand is holding world up: Node is up.



Hand is holding world up, with check mark: Node is up, metrics have been selected for monitoring.

A check mark indicates that metrics have been selected for monitoring. In addition, when a node is selected, the background color of the node icon will change.

Clusters



A cluster is a collection of nodes that appear as a single-server system. Clusters offer application availability and scalability greater than is possible with a single system.

A check mark indicates that metrics have been selected for monitoring. When a cluster is selected the background color of the cluster icon changes.

Groups



A group is a collection of nodes and/or clusters that are frequently managed together. Globes in a container represent these collections.

If the group icon shows a check mark, metrics have been selected for monitoring for every cluster and node in the group. When a group is selected the background color of the group icon changes.

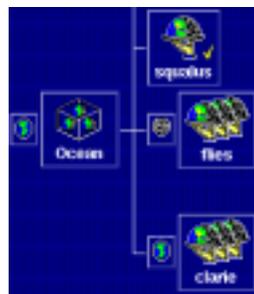
Globes



A globe appears next to each container (group) and set of three globes (cluster). A globe displaying the continent side shows that all nodes in the group or cluster are exposed. A globe showing the darker, latitude and longitude grid side shows that all nodes are hidden. Clicking on this icon exposes or hides all the nodes and clusters inside.



Figure 5 Nodes Display



Main Window Buttons

Buttons are sensitive. Click them to perform the operations in this section.



Each category of metrics has its own button. This is the button for the CPU metric category. Click on it to display the CPU metrics available for threshold monitoring. Each metric category presents its choices in a similar manner.



A metric category button looks like this when it is selected. The LED on the button shows bright green.

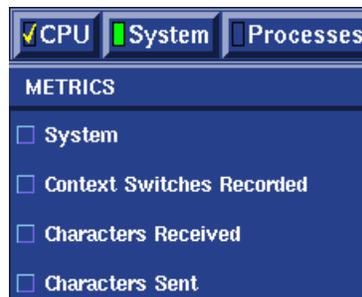


A metric category button looks like this when it is no longer selected, but metrics within that category are selected.



A metric category button looks like this when both the category and the metrics within that category are selected.

Figure 6 Metrics Selection



When this button is on, the display work area is shown. The display view of the work area provides controls for selecting metric categories, individual metrics, display types, and sampling intervals. The type of display used depends on the display type chosen from the option menu to the right of each metric.



When this button is on, the threshold work area is shown. The threshold view of the work area provides controls for selecting threshold categories, setting individual thresholds, and choosing notification methods.



This button (more...Advanced) is active only when the threshold work area is shown.



Click on this button to start the session currently specified. The displays and thresholds you have selected become active as soon as you click on this button. This button is active only when no session is running.



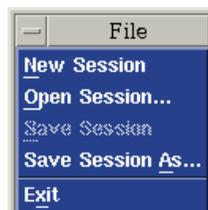
Click on this button to stop the current session. All metric displays close. This button is active only when a session is running.

Main Window Toolbar and Menu Bar

The toolbar and menu bar provide quick access to functions.

The main window has both a menu bar and a toolbar. Together they provide quick access to the functions of Performance Manager. The menu bar contains the following items, which are tear-off menus. If you click the underscored letter in each item, that menu will “tear off” and display separately.

Menus and Menu Commands



File

Use the commands on the File menu to start a new session, open a previously saved session, save as another name, or quit the session and exit Performance Manager.

- New Session
 - Opens a new session.
- Open Session

Displays the Open Session dialog box, providing a choice of existing session files.

- Save Session

Saves an open session.

- Save Session As

Displays the Save Session As dialog box, providing a means to preserve the existing session file and begin a new session file with the same characteristics.

- Exit

Quits the session and exits Performance Manager.



View

Use the commands on the View menu to choose the area of the main window displayed.

- Toolbar

Selects the toolbar for display.

- Nodes

Selects the node area for display.

- Work Area

Selects the work area for display.

- Messages

Selects the message area for display.



Options

Use the commands on the Options menu when you want to customize the interface.

- Enable Tool Bar Label

When turned on, displays a label as the cursor passes over each toolbar icon.

- Show Domain Names in Nodes Area

When turned on, displays the fully qualified domain names for each node, instead of the simple name. This is an example:

Simple: starfish

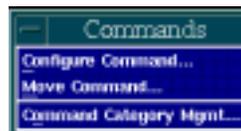
Fully qualified: starfish.bottom.PugetSound.com



Tasks

Use the commands on the Tasks menu when you want to manage nodes, metric categories, or thresholds.

- **Node Management**
Provides access to the controls for adding, deleting, and moving nodes and clusters.
- **Category Management**
Metric categories can be made visible or hidden. Visible categories are selectable for viewing.
- **Threshold Notifications**
Presents a list of activity with a reporting window.



Commands

Use the commands on the Commands menu when you want to configure commands, move commands, or manage command categories.

- **Configure**
Displays the Configure dialog box, which you can use to integrate your commands with Performance Manager.
- **Move**
Displays the Move dialog box, which enables you to regroup commands in different categories.
- **Command Category Mgmt**
Displays the Command Category Mgmt dialog box, which enables you to add or delete command categories.



Execute

The Execute menu lists categories of commands, with related submenus, showing commands that can be run on selected nodes. When you choose a command from a submenu, the Execute dialog box opens. You can also change the categories listed, move commands between categories, modify the commands, add new commands, and delete commands. The following categories are listed by default:

- **Performance Analysis**

These commands detect performance problems and offer corrective advice in four areas: CPU, memory, network, and disk I/O.

- System Management

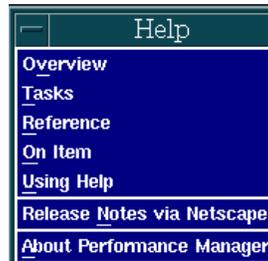
These commands perform tasks on the node they are executing on.

- AdvFS Performance Analysis

These commands analyze file system performance.

- Cluster Performance Analysis

These commands analyze cluster performance.



Help

Use the commands on the Help menu to view online help about Performance Manager, start Netscape Navigator®, and see topics about how to use CDE Help.

- Overview

Opens the first window of the help volume. From this scroll box you can navigate to any topic.

- Tasks

Opens the Using Performance Manager section of the help volume. From this scroll box you can navigate to any topic.

- Reference

Opens a section of the help volume with more information about the functions of Performance Manager than is available from On Item.

- On Item

Changes the cursor to a question mark. Placing the question mark on an area of the GUI and clicking opens a help window with specific information. This is a quick way to read the description of a metric listed in the work area.

- Using Help

Opens the CDE help volume, which explains how the help system works.

- Release Notes via Netscape

Opens the *Performance Manager Release Notes* in Netscape, the browser that ships with Tru64 UNIX.

- About Performance Manager

Opens the help window containing information about this software version, copyrights, and trademarks.

Toolbar Icons

Use the icons on the toolbar for quick access to the functions of Performance Manager. The toolbar icons are arranged by groups and represent the actions described in this section.

File Group

Use these icons to create a new session, open a saved session, or save a session.



New Session



Open Session



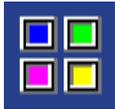
Save Session As

Task Group

The Node Management icon provides access to the controls for adding, deleting, and moving nodes and clusters. Use the Category Management icon to open a dialog box for making metric categories visible or hidden. Visible categories are selectable for monitoring. Use the Threshold Notification icon to display a list of activity with a reporting window.



Node Management



Category Management



Threshold Notification

Command Group

Use the Configure Command icon to open the Configure dialog box, which allows you to integrate your commands with Performance Manager. The Move Command enables you to regroup commands in different categories. Command Category Management enables you to add or delete command categories.



Configure Command



Move Command



Command Category Management

Help

Clicking On-Item Help changes the cursor to a question mark. Place the question mark on an area of the GUI and click to open a help window with specific information about an item. Clicking Overview Help opens the first window of the help volume. From this scroll box you can navigate to any topic.



On-Item Help



Overview Help

Modifying the Main Window

You can change the appearance of the main window. The background color can be changed by starting Performance Manager with a different background color; for example:

```
# pmgr -fg black -bg salmon
```

You might want to do this to provide greater viewing contrast, but be careful not to choose a color that will obscure text, such as a black foreground that hides black text.

You can also modify the font and the foreground and background colors used in the interface by editing the X resource file `/usr/lib/X11/app-defaults/PM`.

Performance Manager Popup Menu

Click the third (right) mouse button anywhere in the GUI to open the Performance Manager popup menu. This menu provides quick access to tasks for those who are familiar with Performance Manager. The popup menu mirrors the tasks in the toolbar, grouping them in the following sequence:

- Sessions
 - New Session
 - Open Session
 - Save Session As
- Tasks
 - Node Management
 - Category Management
 - Threshold Notifications
- Commands
 - Configure Commands
 - Move Commands
 - Command Category Management
- Options
 - Enable Tool Bar Label
 - Show Domain Names in Node Area

- GUI Session Controls
 - Start Session
 - Stop Session

Chapter 3

Managing Nodes

Manage nodes by adding nodes or clusters to and deleting nodes or clusters from the main window's nodes area, moving nodes or clusters among groups, and creating and deleting groups. From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.

See the individual task descriptions for specific procedure steps. All tasks begin from the Node Management dialog box.

Figure 7 Node Management Dialog Box



The Apply button applies any changes you made.



The OK button applies any changes you made and closes the dialog box.



The Cancel button dismisses the window without applying any changes.

Creating Groups

Create groups to organize your nodes in the main window's nodes area. Follow these steps to create a group:

- 1 From the main window's Tasks menu, choose Node Management, which opens the Node Management dialog box.
- 2 Select Create from the option menu.

- 3 Click in the Group field and type the name of the group to be added, or choose the group from the drop-down list.
- 4 Click on Apply or OK.

Deleting Groups

Deleting a group removes it from the main window's nodes area, and all nodes and clusters in that group will also be removed. Follow these steps to delete a group:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Delete from the option menu.
- 3 Click in the Group field and type the name of the group to be deleted, or choose the group from the drop-down list.
- 4 Click on Apply or OK.

Adding Nodes

Adding a node makes an icon for it appear in the main window's nodes area, which allows you to display the node's metrics and run scripts on it. Follow these steps to add a node:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Create from the option menu.
- 3 Click in the Group field and type the name of the group (new or existing) the node is to be added to, or choose the group from the drop-down list.
- 4 Click in the Node or Cluster Alias field and type the name of the node to be added.
- 5 Click on Apply or OK.

Deleting Nodes

Deleting a node removes it from the main window's nodes area. Once it is deleted, you will no longer be able to display the node metrics or run scripts on the node. Follow these steps to delete a node:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Delete from the option menu.
- 3 Click in the Group field and type the name of the group the node is to be deleted from, or choose the group from the drop-down list. If you choose a group that does not contain the node, the node is not deleted.
- 4 Click in the Node or Cluster Alias field and type the name of the node to be deleted, or choose the node from the drop-down list.
- 5 Click on Apply or OK.

Moving Nodes

You can move a node from one group to another in the main window's nodes area. Follow these steps to move a node:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Move Node from the option menu.
- 3 Click in the Group field and type the name of the group the node is to be moved from, or choose the group from the drop-down list. If you choose a group that does not contain the node, the node is not moved.
- 4 Click in the Node or Cluster Alias field and type the name of the node to be moved, or choose the node from the drop-down list.
- 5 Click in the Move to Group field and type the name of the group the node is to be moved to, or choose the group from the drop-down list.
- 6 Click on Apply or OK.

Adding Clusters

Add clusters so you can monitor their nodes in the main window's nodes area. Follow these steps to add a cluster:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Create from the option menu.
- 3 Click in the Group field and type the name of the group (new or existing) the cluster is to be added to, or choose the group from the drop-down list.
- 4 Click in the Node or Cluster Alias field and type the name of the cluster to be added; the other cluster nodes will automatically be added to the cluster.
- 5 Click on Apply or OK.

Deleting Clusters

Deleting a cluster removes it from the nodes area. Once it is deleted, you will no longer be able to display metrics or run scripts on any node in the cluster. Follow these steps to delete a cluster:

- 1 From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2 Select Delete from the option menu.
- 3 Click in the Group field and type the name of the group the node is to be deleted from, or choose the group from the drop-down list. If you choose a group that does not contain the cluster, the cluster is not deleted.
- 4 Click in the Node or Cluster Alias field and type the name of the cluster to be deleted, or choose the cluster from the drop-down list.
- 5 Click on Apply or OK.

Moving Clusters

You can move a cluster from one group to another in the main window's nodes area. Follow these steps to move a cluster:

- 1** From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management dialog box.
- 2** Select Move Node from the option menu.
- 3** Click in the Group field and type the name of the group the cluster is to be moved from, or choose the group from the drop-down list. If you choose a group that does not contain the cluster, the cluster is not moved.
- 4** Click in the Node or Cluster Alias field and type the name of the cluster to be moved, or choose the cluster from the drop-down list.
- 5** Click in the Move to Group field and type the name of the group the cluster is to be moved to, or choose the group from the drop-down list.
- 6** Click on Apply or OK.

Chapter 4

Displaying Clusters

Performance Manager displays clusters using the auto-discovery feature. There are some differences in PM's operation for TruCluster Production Server (TruCluster Version 3.5) and TruCluster Server (TruCluster Version 5.0). With TruCluster Server, PM recognizes cluster aliases and does not use director names.

Auto-Discovery for Clusters

When you add a node, Performance Manager checks to see if the node belongs to a cluster or is a cluster alias. PM does this by querying the node for a cluster name or *director name*. If a value for either cluster name or director name is returned, the cluster populates the GUI with its members. If the returned value is for cluster name, PM recognizes the cluster as a TruCluster Server cluster and populates the GUI using the cluster name (the default cluster alias) and displays all of its members. If the returned value is for director name, PM recognizes the cluster as a TruCluster Production Server cluster and creates a cluster entity using the director name for the cluster. The cluster entity queries the node's membership table and populates the GUI with the members.

PM watches the membership table and updates the GUI to reflect changes.

Note For TruCluster Production Server, if the director name changes, the cluster node changes its name to match the new director name. This changes all uses of the old name to the new name in displays and thresholds. Note that this means cluster nodes defined in old sessions will have their names changed to match the director name.

Display Representation of Clusters

When monitoring a cluster, Performance Manager discovers all the members of the cluster. When the membership changes, Performance Manager adjusts its representation of the cluster as follows:

- If a node was added to the cluster, a new icon for that node is added. If the cluster has any active displays, the display adjusts to include the new node.
- If a node was removed from the cluster, Performance Manager deletes the icon for that node from its view of the cluster. Any active displays for the cluster adjust to remove the deleted node.
- If the deleted node has any displays defined explicitly for that node, they are deleted from the session. If the deleted node subsequently returns to the session, Performance Manager adds it to the cluster view. However, node-specific displays will not be recreated. Currently, the only way to regain these node-specific displays is manually redefining them or reloading them from a saved session.

Possible Anomalies for TruCluster Production Server

Director name changes may result in two cluster nodes for the same cluster appearing in Performance Manager. This may happen if attempts to get cluster information from a node occur during the change and a node is removed from the cluster as described above. If the nodes(s) removed from the cluster notices the new director names before the cluster node notices it, the removed node will create a new cluster node with the new name.

Usually the pre-existing cluster node notices the director name change, and also notices there is already a cluster node with the same name. In that case it does the following:

- Moves its displays and thresholds to the new node.
- Removes its children, allowing the new cluster to acquire them.
- Deletes itself from the session.

If the pre-existing node removes all of its children because it could not get information from them, it will continue asking for information from the last node that it polled. If this node never responds, this cluster node will continue to exist without children even if a new cluster node has been created based on information from the other nodes.

Note To avoid conflicts between group names and cluster node (director) names, do not give group nodes the same names as cluster director names. This interferes with cluster auto-discovery.

For example, if you give the same name as a cluster director when a corresponding cluster node does not exist in the session, and then add nodes from that cluster to the session, the cluster nodes will not be created.

Chapter 5

Monitoring

Monitoring nodes means looking at performance data in real time. This chapter explains sessions and the types of displays you can choose, and includes information on additional monitoring methods. When monitoring, you are watching metrics and thresholds, as defined below:

- Metrics

Performance Manager can gather data on several hundred metrics. Performance Manager metrics servers listen for and service requests for operating system information. For a description of a particular metric, use context-sensitive help. Metrics are covered in more detail in Chapter 6.

- Thresholds

A threshold is a limit (high or low) placed on a specific monitored metric. When a limit is exceeded for more than a specified number of sampling intervals (its tolerance), that threshold is crossed. With its thresholding capability, Performance Manager can set these limits, notify you, and run commands to act on the situation. Thresholds are covered in more detail in Chapter 7.

Sessions

Everything you do in Performance Manager occurs within a *session*. A session is to Performance Manager as a file is to an editor. You can change sessions, save sessions, and recall previous sessions.

When creating a session, you can use the default session settings or select which nodes to monitor and which metrics to watch, and set up any thresholds or archives. One session window can contain both display and threshold metrics, and is identified by file name. The following image of the main window calls out the controls you use in setting up a session.

Figure 8 Creating a Session



Creating a Session

To create a session, follow these steps:

- 1 From the main window's File menu or toolbar, choose New Session.
- 2 Select a node, cluster, or group in the main window's nodes area. The work area will appear to the right.
- 3 Click on the Display or Threshold button, if not already selected.
- 4 Select a metric category from the horizontally scrolling list at the top of the work area.
- 5 Under Metrics, set a metric check box.
- 6 If you are working in the Display work area, use the metric's related option menu to choose:
 - Display type
 - Sampling interval

If you are working in the Threshold work area, use the metric's related option menu to choose:

- Value
- Re-arm point
- Notification methods

- Tolerance
 - Interval
- 7** Repeat the steps (except step 1) for every node, cluster, or group you want to monitor.
 - 8** To start the session you have just created, click on the Start Session button. Starting the session puts everything in motion: the displays you specified will open and the thresholds you specified will be set.
 - 9** After the session window opens, choose actions from the buttons on the session window toolbar:
 - Expand

Click this button to display a selected title. Display metrics are expanded by default.
 - Collapse

Click this button to close the display, showing only the title. Threshold metrics are collapsed by default. However, a visual alert icon next to the threshold title displays the state of the threshold (crossed or not crossed, waiting for data, data request timed out).
 - Float

Click this button to detach (float) this window.

Managing Sessions

Sessions can be saved and recalled later, which eliminates the need to respecify your choices, but you can change anything about a session.

After creating a new session or opening a previously saved session, you need to start it in order to open the session window and monitor data.

To start a session:

- Click on the main window's Start Session button.

To save a session:

- 1** From the File menu, choose Save Session or Save Session As.
- 2** From the main window's File menu, choose Save Session. The File Selection dialog box opens.
- 3** Provide a name for the session; the default extension is `.spm`.

To recall a previous session:

- 1** From the File menu, choose Open Session.
- 2** From the main window's File menu, choose Open Session. The File Selection dialog box opens.
- 3** Choose a session from the dialog box.

To stop a session:

- In the main window, click on the Stop Session button. You can also stop a session by choosing Stop Session from the session window's File menu.

Displays

Each performance metric can be displayed in several display types. Display types are chosen from the option menus to the right of each metric in the main window. Each display includes a charting key designating colors used for each metric. The following images are examples of each display type:

Figure 9 Chart Key



The default background color is black, and the default charting colors used in these examples are blue for 5-second intervals, yellow for 30-second intervals, and magenta for 60-second intervals.

Figure 10 Area Display

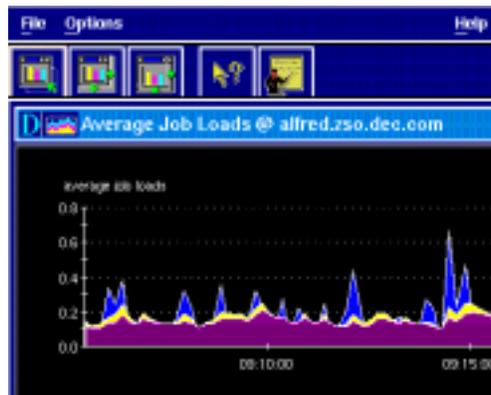


Figure 11 Bar Display



Figure 12 Pie Display



Figure 13 Plot Display

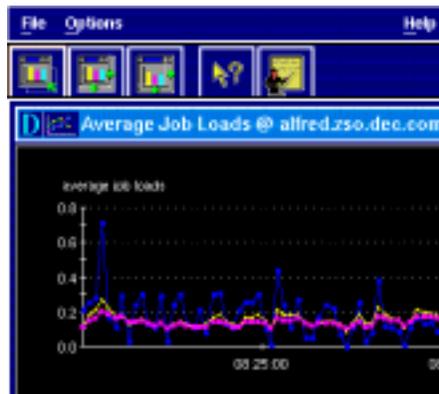


Figure 14 Stack Bar Display

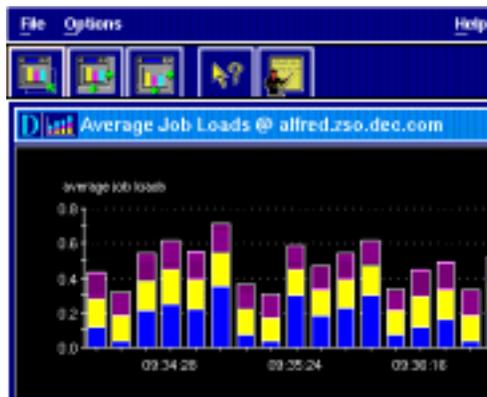


Figure 15 Table Display

Average Job Loads @ alfred.zso.dec.com			
	5 secs	30 secs	60 secs
alfred.zso.dec.com	0.08	0.17	0.16

Floating Displays

When a new session is opened, all displays are shown in the session window; however, individual displays can be expanded, collapsed, or *floated out* in their own separate windows.

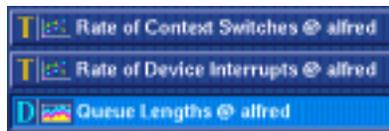
To expand or collapse a display:

- Expand: Click the expand button to display a selected title. Display metrics are expanded by default.
- Collapse: Click the collapse button to close the display, showing only the title. Threshold metrics are collapsed by default.

To float a display:

- 1 Select the metric title, which changes color to show it is selected, as shown in the figure below:

Figure 16 Metric Display Selection



- 2 From the toolbar, choose the first flag icon, Float Selected Display, or from the session window's File menu, choose Current Display, then choose Float.

The display now appears in its own window.

You must save a session after floating displays if you want the displays to appear in their own windows when the session is reopened.

Consolidating Displays

Floating displays can be closed so that they reappear in the session window.

To consolidate a floating display into the session window:

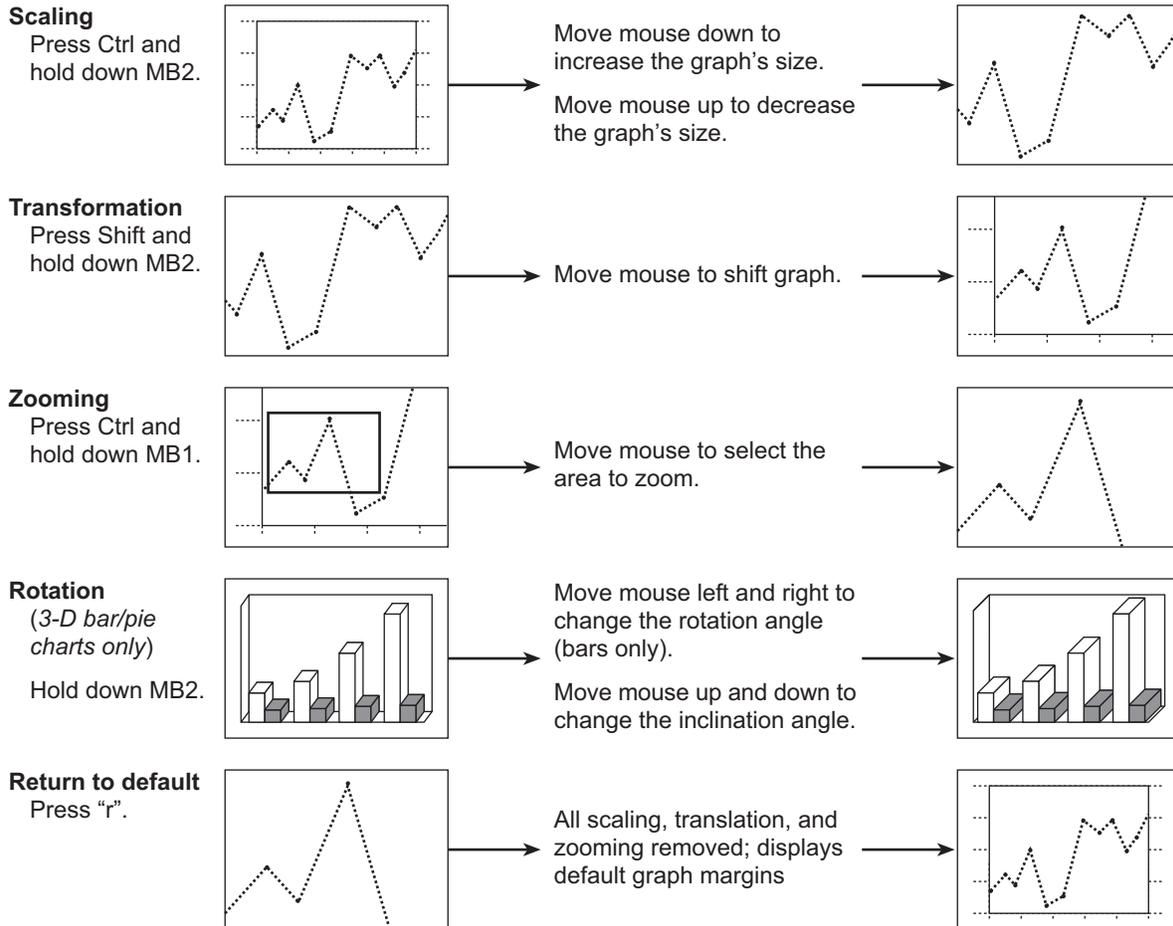
- From the display window's File menu, choose No Float.

The display now appears in the session window.

For thresholds, a visual alert icon by the title displays the state of the threshold (crossed or not crossed, waiting for data, data request timed out).

Manipulating Displays

You can interact with the graph displays in Performance Manager in the following ways:



Setting Display Styles

You can change the data styles chosen for the Performance Manager displays by modifying the PM resource file. The resource file is in this location:

```
/usr/lib/X11/app-defaults/PM
```

A copy of the resource file is included in the reference section of the Performance Manager Help Volume.

The following information may help you work with the resource file:

Default Data Styles

The XrtDataStyle data structure contains all the information about how a set of data will be represented graphically. The fields are broken down as follows:

- lpat — The line pattern used for plots.
- fpat — The fill pattern used in area graphs and bar and pie charts.
- color — The color used when drawing lines in plots and for fills in area graphs and bar charts. It is either a named color or a # character followed by two hexadecimal characters for each of the Red, Green, and Blue components.
- width — The line width used for plots. Must be greater than or equal to one.

- point — The point style used for plots.
- pcolor — The point color used for points in plots. It is either a named color or a # character followed by two hexadecimal characters for each of the Red, Green, and Blue components.
- psize — The size of points that appear in plots. Must be equal to or greater than 0. A size of 0 will result in no point being drawn. A point size is a relative measure. It should not be assumed that a point size of 12 means that the point's glyph will be exactly 12 pixels from top to bottom.

For further information, please see your Xt Intrinsic documentation.

Figure 17 Plot Line Patterns

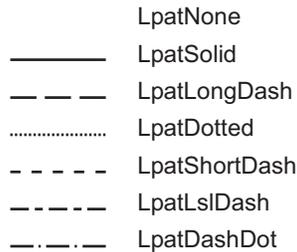


Figure 18 Fill Patterns

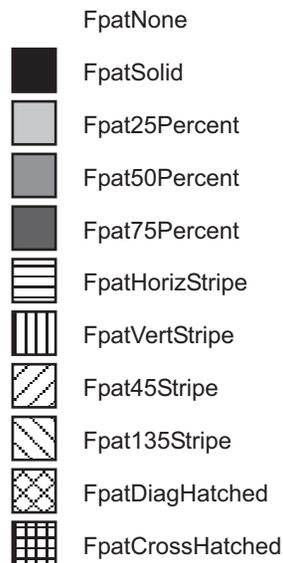
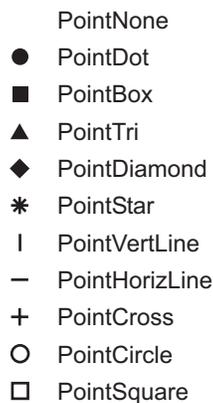


Figure 19 Point Styles



List of Data Styles

Resources of type (XtRXrtDataStyles) specified as a parenthesized list, with each member specifying a complete data style (XtRXrtDataStyle). For example:

```
! change the graph data styles
pmgr*xrtDataStyles: (( LpatSolid FpatSolid "blue" 1 PointDot "blue" 4 ) \
( LpatSolid FpatSolid "yellow" 1 PointTri "yellow" 4 ) \
( LpatSolid FpatHorizStripe "magenta" 1 PointBox "magenta" 4 ) \
( LpatSolid Fpat25Percent "cyan" 1 PointDiamond "cyan" 4 ) \
( LpatSolid FpatVertStripe "#6699ff" 1 PointStar "#6699ff" 4 ) \
( LpatSolid FpatDiagHatched "#ff9900" 1 PointCircle "#ff9900" 4 ) \
( LpatSolid Fpat45Stripe "#33cc99" 1 PointSquare "#33cc99" 4 ) \
( LpatSolid FpatCrossHatched "#cc3333" 1 PointCross "#cc3333" 4 ))
```

For further information on resource files and their usage, please see your Xt Intrinsic documentation.

Other Monitoring Methods

Performance Manager supports two additional monitoring methods:

- From the command line using UNIX commands supplied by Performance Manager
- Using third-party SNMP applications

Monitoring from the Command Line

The following UNIX commands are provided for command-line access to the metrics servers:

- `getone`
- `getnext`
- `getmany`
- `getbulk`
- `gettab`

Note The `getbulk` command uses the SNMPv1 extensions and requires that you access the metrics servers via their private SNMP request ports rather than the well-known SNMP request port. The port to be used is specified by the environment variable `PMGR_SNMP_PORT`. The appropriate port numbers should be listed in the `/etc/services` file on the management station.

The following example shows how to query `pmgrd` using the `getmany` command:

```
% getmany alfred public pm
pmCmSysProcessorType.0 = alpha(2)
pmCmSysOperatingSystem.0 = digital-unix(2)
pmCmSysOSMajorVersion.0 = 3
pmCmSysOSMinorVersion.0 = 2
pmCmSysPageSize.0 = 8192
pmCmSysNumCpusOnline.0 = 2
```

```

pmCmSysPhysMem.0 = 262136
pmCmSysPhysMemUsed.0 = 56328
pmCmSysUpTime.0 = 88677120
pmCmSysDate.0 = 7.204.1.17.17.58.57.0.--.8.0
pmCmSysNumUsers.0 = 14
pmCmSysProcesses.0 = 81
.
.
.
pmAoVmSwapInUse.0 = 57160
pmAoVmSwapDefault.0 = /dev/re3c
pmAoVmSiIndex.1 = 1
pmAoVmSiPartition.1 = /dev/re3c
pmAoVmSiPagesAllocated.1 = 256896
pmAoVmSiPagesInUse.1 = 7145
pmAoVmSiPagesFree.1 = 249751
pmAoBcReadHits.0 = 21761200
pmAoBcReadMisses.0 = 78356
pmAoIfEthIndex.1 = 1
pmAoIfEthName.1 = tu0
pmAoIfEthCollisions.1 = 13064347
End of MIB.

```

Monitoring with SNMP Network Management Systems

You can also use SNMP Network Management Systems (NMS) to access Performance Manager's metrics servers. Examples of available systems include:

Commercially Available	Freely Available
DIGITAL NetView®	scotty/tkined
IBM® NetView/6000	
HP® OpenView™	
SunNet Manager	

Note The following information is taken from the file `/usr/opt/pm/nms/README.nms`.

Using NetView

Use the following procedure to install and use NetView:

To install and uninstall NetView support:

- To use PM's NetView support, you should first install NetView and Performance Manager on your management node. Then, as superuser, use the following command:

```
# /usr/opt/pm/nms/PMGR_Netview_Setup INSTALL
```

- To uninstall NetView support, use the following command as superuser:

```
# /usr/opt/pm/nms/PMGR_Netview_Setup DELETE
```

Loading PM MIBs

To make NetView aware of the MIB variables provided by PM's metrics servers, it is necessary to load their associated MIB files into NetView. This is done using the Options Load/Unload MIBs: SNMP... menu item. The MIB files for PM's metrics servers are listed below, with the metrics server name followed by the NetView-loadable MIB file:

- pmgrd
/usr/OV/bin/snmp_mibs/pm-mib.pnv
- clu_mib
/usr/OV/bin/snmp_mibs/cluster-mib.pnv

Using the NetView MIB Browser Application

Once you have loaded Performance Manager's MIB files you should be able to browse them using the NetView MIB browser. Note that MIB browsers that were opened prior to loading a new MIB will not reflect the additional MIB information, so you will have to open new ones to get the changes.

Performance Manager's MIB files are found under `.iso.org.dec`.

Note The string `dec` appears in at least two places in the OSI naming tree (`iso.org.dod.internet.private.enterprises.dec` is another well-known place). In the NetView browser, click on Up Tree until you reach `org` and then go down `dec` to find the PM MIB variables.

Sending SNMP Traps Using `trapsend`

The script `trapsend-example` found in this directory is an example of a script that periodically monitors the value of a variable against a threshold value. Upon crossing the threshold value, it sends a trap to NetView. As described in the KNOWN BUGS section of `trapsend(1)`, the script takes care of temporarily setting and then unsetting `SR_MGR_CONF_DIR`. The Performance Manager kit installation sets up `mgr.cnf` and `snmpinfo.dat` in the `/etc/srconf/agt` directory.

The script assumes that you are running the extensible SNMP agent (`snmpd`) that ships with Tru64 UNIX version 4.0F (and later versions).

Sample MIB Applications

The following sample PNV applications are shipped with this kit. They are installed by `PMGR_NetView_Setup` and can be accessed from the Monitor-Performance Manager NetView menu.

File Name	Files Installed As
<code>ovmib.pmgr_RunQueue</code>	<code>/usr/OV/registration/C/ovmib/PMGR_RunQueue</code>
<code>ovmib.pmgr_RunQueue.help</code>	<code>/usr/OV/help/ovmib/OVW/Functions/PMGR_RunQueue</code>
<code>ovmib.pmgr_SysInfo</code>	<code>/usr/OV/registration/C/ovmib/PMGR_SysInfo</code>
<code>ovmib.pmgr_SysInfo.help</code>	<code>/usr/OV/help/ovmib/OVW/Functions/PMGR_SysInfo</code>
<code>ovmib.pmgr_SwapConfig</code>	<code>/usr/OV/registration/C/ovmib/PMGR_SwapConfig</code>
<code>ovmib.pmgr_SwapConfig.help</code>	<code>/usr/OV/help/ovmib/OVW/Functions/PMGR_SwapConfig</code>

Chapter 6 Metrics

Performance Manager can gather data on several hundred metrics. For a description of a particular metric, use context-sensitive help.

Note Context-sensitive help for metrics is only available in the work area, not the session window or displays.

From the main window's Help menu, choose On Item, then click on a metric. A Help box will appear.

Displaying Metrics

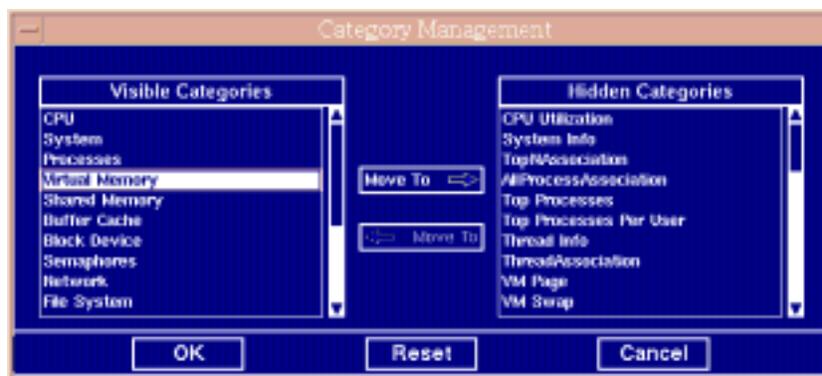
Select one of the metric categories at the top of the work area to display metrics that you can select for monitoring.



Showing Hidden Metric Categories

To display additional metric categories in the list:

- 1 From the main window's toolbar or Tasks menu, choose Category Management, which opens the Category Management dialog box.
- 2 Select a category or multiple categories in the Hidden Categories list box.



- 3 Click on the lower Move To button. The selected category now appears in the Visible Categories list box.

- 4 Click on OK.

Hiding Metric Categories

If the list of metric categories shows categories that you are not using, you can choose to temporarily remove categories from the list. To remove categories from the list:

- 1 From the main window's toolbar or Tasks menu, choose Category Management, which opens the Category Management dialog box.
- 2 Select a category or multiple categories in the Visible Categories list box.



- 3 Click on the upper Move To button. The selected category now appears in the Hidden Categories list box.
- 4 Click on OK.

Chapter 7

Thresholds

A threshold is a limit (high or low) placed on a specific monitored metric. When a limit is exceeded for more than a specified number of sampling intervals (its *tolerance*), that threshold is crossed.

For example, you could set a threshold of 5% maximum CPU time on system processes on all nodes, and give the threshold a tolerance of three. Then, if a node had more than 5% of its CPU time used for system processes for more than 3 consecutive sampling intervals, that threshold would be crossed.

You can set thresholds to notify you when they are crossed. The Threshold Notifications dialog box is the default method of notification and provides you with detailed information.

Caution Executing resource-intensive commands when a threshold is crossed causes the system load to increase. The increased load can cause more frequent threshold crossings, and in some cases, the threshold crossings are due solely to command execution. This can result in an excessive and continually growing system load.

To avoid this situation, increase the tolerance for the expression being monitored. The command will not execute until the threshold is crossed the number of times specified by the tolerance level.

Some other examples of thresholds:

- A node's I/O Queue exceeds a dozen processes for more than 10 consecutive sampling intervals.
- A node's Disk Transfers exceed 25/second for more than 5 consecutive sampling intervals.
- A node's Total Bad IP Packets exceed zero in any sampling interval.

When a threshold is crossed, the following occurs:

- 1 The event is logged (written in the Performance Manager log file: `/var/opt/pm/log/pmgr_gui.log`).
- 2 A command (if specified) is run. Performance Manager has a number of commands built in, but it is also extensible. You or your system administrator can create your own commands. This command can do anything from sending you mail about the problem, to taking steps to fix the problem.

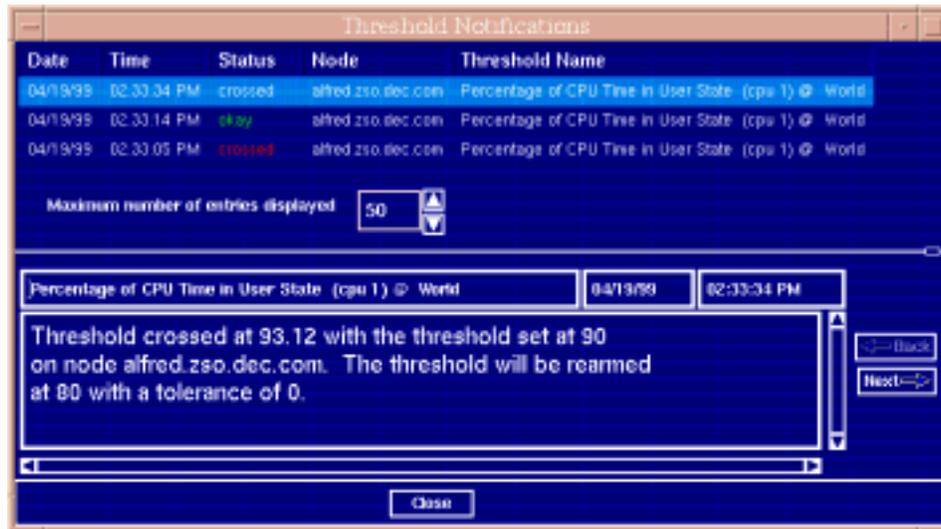
The session window displays threshold data along with monitoring data. The displays are managed in the same way, and the type is designated at the beginning of the title bar with a *D* for displays and a *T* for thresholds.

Threshold Notifications

The Threshold Notifications dialog box has a list view of threshold activity and a reporting window for information on selected thresholds. There are three action buttons:

- Back — Returns you to the previous threshold.
- Next — Moves to the next threshold.
- Display — Switches to the display mode.

Figure 20 Threshold Notifications Dialog Box



Setting Thresholds

Follow this procedure to set a threshold:

- 1 Select a node, cluster, or group in the main window's node area.
- 2 Click on the Threshold button in the work area.
- 3 Select a metric category.
- 4 Select the specific metrics for monitoring from the list.
- 5 Set the value of the threshold.
- 6 Set the rearm point. The *rearm point* occurs when the metric drops a specified amount below the threshold. If it recrosses the threshold after rearming, another alert will be sent.

These are the metric categories displayed by default in the threshold work area:

Figure 21 Default Threshold Metric Categories



Selecting the More button for a specific metric opens another dialog box for advanced settings (notification methods and additional information).

Figure 22 More... Button



CPU Thresholds

You can set thresholds on the following CPU metrics:

- Average Job Loads over Last 5 Seconds
- Average Job Loads over Last 30 Seconds
- Average Job Loads over Last 60 Seconds
- Percentage of CPU Time in User State
- Percentage of CPU Time in System State
- Percentage of CPU Time in Idle State

System Thresholds

You can set thresholds for the following system metrics:

- Rate of Context Switches
- Rate of Device Interrupts

Processes Thresholds

You can set thresholds for the following processes metrics:

- Percentage of CPU Use by Top Processes
- Percentage of CPU Use by Top Users

Buffer Cache Thresholds

You can set thresholds for the following buffer cache metric:

- Percentage of Read Misses

Network Thresholds

You can set thresholds for the following network metrics:

- Percentage of Timeouts for Calls
- Rate of Ethernet Collisions
- Percentage of Erroneous Outbound Packets
- Percentage of Erroneous Inbound Packets
- Rate of IP Datagrams Discarded
- Rate of ICMP Errors
- Rate of TCP Errors
- Rate of UDP Errors

File System Thresholds

You can set thresholds for the following file system metrics:

- Percentage of Available File Space
- Percentage of Free Inodes

Memory Thresholds

You can set thresholds for the following memory metrics:

- Percentage of Free Paging Memory
- Rate of Page Faults
- Rate of Pages Paged Out
- Number of Free Pages
- Rate of Processes Swapped Out
- Percentage of Free Swap Space

AdvFS Thresholds

You can set thresholds for the following AdvFS metrics:

- AdvFS Agent is Down
- Percentage of Free Space in AdvFS Domains
- Percentage of Free Space in Domain
- Percentage of Free Space in Fileset
- Percentage of Free Space in Domain Volume

TruCluster Thresholds

You can set thresholds for the following TruCluster metrics:

- TCR Agent is Down
- Deadlock Queue

Environmental Thresholds

You can set thresholds for the following environmental metrics:

- High Temperature Reading
- Status of Thermal Sensor
- Status of Fans
- Status of Power Supplies

Advanced Threshold (more...) Dialog Box

The advanced threshold (more...) dialog box has two sections. Use them for these tasks:

Threshold Notification Methods

- Choose one or more notification methods by clicking the check box *on*.
 - Threshold Notifications Dialog Box (default selection). This displays a dialog box on your screen when a threshold is crossed.
 - Send Email to: Type an address in this field.
 - Execute: Command - Set the Execute toggle. Choose Command to open a pull-down list of command categories, then choose a command from the submenu to open a command execution dialog box.

- Use the Notification Message text entry field to create your own notification message.

Additional Threshold Information

- Set the tolerance for this threshold. This is the number of consecutive threshold crossings permitted before a violation is reported.
- Set the interval for this threshold. This is the *sampling rate*, or time specified between samples.

Click on OK to save the settings and return to the main window, click on Reset to return the settings to their defaults, and click on Cancel close the dialog box without saving the settings.

Threshold Environment Variables

These environment variables are set up internally to retrieve threshold information from commands that you create. For example, the `./var/opt/pm/Smscripts/pm_mailer` script sends detailed mail about the crossed threshold that uses this information. You can create your own shell script that accesses these values using the dollar sign (\$) symbol in front of the variable; for example, `$PMTHRESH DESCRIPTION`. These variables are helpful in creating your own logging script that tracks thresholds and rearms of Performance Manager’s metrics.

Environment Variable	Description
PMTHRESH_DESCRIPTION	Description of the expression in the database.
PMTHRESH_CURRENT_VALUE	Value that has triggered threshold.
PMTHRESH_THRESHOLD_VALUE	Value that had to be passed to trigger threshold.
PMTHRESH_NODE	Node on which triggered threshold was detected.
PMTHRESH_USER_MESSAGE	User message from advanced threshold dialog box.
PMTHRESH_UPDATE_TIME	The update time value from the triggered expression.
PMTHRESH_REARM_VALUE	The value at which the threshold will be rearmed.
PMTHRESH_TOLERANCE_VALUE	The tolerance of the triggers.
PMTHRESH_STATE	Value is a string being either crossed or rearmed corresponding to the triggered event.
PMTHRESH_INSTANCE	Additional information about the triggered threshold, such as which file system or CPU crossed.
PMTHRESH_OPERATOR	Greater than or less than the threshold value.

Chapter 8

Commands

A command is any executable program, such as a shell script or binary file. Performance Manager can execute commands on remote nodes or the local GUI node, and display the output back to the local GUI node.

Performance Manager comes with several performance analysis, AdvFS analysis, cluster analysis and system management commands. You can execute these as they are or modify them to suit your needs. Performance Manager commands can be found below the `/var/opt/pm` directory.

You can also execute your own commands from Performance Manager by adding commands to the Execute menu, and you can organize your commands in categories. Use the Configure dialog box to integrate your commands with Performance Manager.

Performance Analysis Commands

Performance analysis commands can execute on one node, but analyze data collected from other nodes. Performance Manager's performance analysis commands are scripts that detect performance problems and offer corrective advice in four areas: CPU, memory, network, and disk I/O. To execute a performance analysis command, from the main window's Execute menu, choose Performance Analysis, then one of the following commands.

CPU Commands

These commands analyze CPU performance.

CPU Analysis

This script determines how efficiently a computer's CPU is being used. High idle time during a heavy load indicates an I/O bottleneck. High system time under a heavy load indicates excessive overhead. If inefficiency is discovered, other scripts can reveal the cause; try the Virtual Memory, Swapping, and Device I/O scripts.

Load Average

This script determines a computer's load average for the last minute, last 5 minutes, and last 15 minutes. The load average is the number of jobs in the run queue. An acceptable load average is 3 to 7 jobs for a large system, 1 to 2 jobs for a workstation. This script also reports if a computer is consumed by a small number of user processes, and lists the top CPU-using processes.

Memory Commands

These commands analyze memory performance.

Buffer Cache

This script determines if a computer's buffer cache is too large or too small. A too-small cache causes excessive I/O. A too-large cache causes excessive paging and swapping.

Excessive Paging

This script determines if there is excessive paging on a computer by checking the number of free pages, paged out pages, and page faults. Excessive paging can be caused by a new process trying to allocate pages, or by active virtual memory being too large relative to active real memory.

Excessive Swapping

This script displays virtual memory and swap space usage and detects excessive usage.

Memory Shortage

This script determines if a computer has a memory shortage. If there is much swapping during paging, and runnable processes are swapped out while the free list increases, lack of memory could cause desperation swapping (also called *thrashing*) to occur.

Virtual Memory

This script determines if a computer has virtual memory problems. This script displays swap configurations and the number of free pages, and compares the amounts of physical and virtual memory.

Network Commands

These commands analyze network performance.

Gateway Errors

This script determines if a computer has excessive gateway errors by looking at the number of bad checksum fields for IP, ICMP, TCP and UDP. Gateway errors should be less than one hundredth of a percent of the total number of packets received.

Network Errors

This script determines if a network node (a computer in a network) has exceeded the acceptable number of network output errors and collisions. This script examines the length of the send queue for all connections, and displays the number of output errors, input errors, and collisions, as well as the number of in and out packets.

Packet Retransmissions

This script determines if a node has excessive network packet retransmissions by looking at the number of retransmissions and bad xids. (Bad xids are packets that return an xid different from the one sent.) Packet retransmissions should be less than 1% of the total number of client NFS calls. Retransmissions increase when you are working with network hardware or all your computers boot at the same time.

Disk I/O Commands

These commands analyze disk I/O performance.

Excessive Transactions

This script displays the transactions per second (tps) and total transactions on each device and reports excessive activity.

File System Analysis

This script determines if there are sufficient inode and file table entries to support the number of system processes. If inode and open file usage are more than 80%, increase the system parameter to make the usage less than 80%.

System Management Commands

System management commands perform tasks on the node they are executing on. Performance Manager provides the following system management scripts. To execute one, from the main window's Execute menu choose System Management, then one of the following scripts:

CleanFilesystems

This script cleans full file systems of core files and other user-specified unneeded files.

FileModification

This script determines if files have been modified or accessed.

GrowthOfFiles

This script determines if files are growing faster than a certain rate.

MaintainFiles

This script allows you to perform the following file management tasks:

- Move files to new file systems
- Copy files to new file systems or tapes
- Make symbolic links
- Delete files
- Change file permissions
- Change user and group ownership for files
- Undelete AdvFS files

PMArchiver

This script allows you to capture all metric data on one or more nodes without having to monitor the nodes. The archived data can be replayed using Microsoft® Excel or any other graphing tool you create an interface for. PMArchiver also provides you with running averages. You can choose the sample interval for measurement granularity, the number of intervals to average over, and total sample time. The lower limit of the interval (-i) is bound by the time it takes to query the metrics.

- This script can be used for multiple CPUs, using the metrics for idle time, nice time, system time, and user time to produce average time.
- This script allows you to choose the metrics for archiving. You construct a file containing the metrics you want to average and determine whether you want the output file named by metric or machine.

Performance Manager will wait while this script runs, only closing after it has reached completion. If you set a duration longer than the time you want to run the PM GUI, you can run the script outside PM, from a command line.

PMDeltaArchiver

This script is similar to PMArchiver, but it tracks the delta of COUNTER type metrics, rather than the raw values of GAUGE type metrics.

RCArchiver

The `rc_archiver` will archive metrics from the `snmpd`, `pmgrd`, `advfsd`, and `clu_mib` daemons. It assumes the ports for the daemons are 161, 1167, 1163, and 1165 respectively. You will need to modify the script if your daemons run on different ports.

This demonstration script archives the rate in seconds or count per sample of data for a tabular metric that you specify on the command line. You can choose the sample interval, sample duration, archive field delimiter character, the port number of the daemon from which the metrics will be retrieved, and the directory where the archive files will be written.

PingNode

This script pings a node at intervals you set. When the round trip ping time between the initiating node and the node specified on the command line exceeds the set threshold, you are notified.

impact_diskmon and impact_procmon

These scripts monitor disks and processes, sending traps when a capacity threshold is crossed or a process has failed. If they are run from the PM GUI, they will close upon completion. If you wish to monitor over a period of time, run them from a command line.

- `impact_diskmon` monitors disk partitions for fill percent thresholds.
- `impact_procmon` monitors process names that should exist on `node_list`.

SignalProcess

This script sends the user-specified SIGNAL, in alphabetic or numeric form, to one or more processes. This script allows you to set the following flags:

- Signal a process directly by entering a process ID.
- Display all processes for a user and choose which to signal.
- Display all processes containing a given string and choose which to signal.

If only one process matches your entry when using the `grep` or `user` flag, it will be signaled directly.

DiskUsage

This script creates a report displaying the disk usage of each user on the file system specified. By default the display will be written to standard out. This script allows you to set the following optional flags:

- Mail the usage report to a user.
- Write the report to a file.

AddSwapFile

This script allows you to add a UFS partition as additional swap space. The script prompts you for a block special device (such as `rz4c` on a 4.0x system or `dsk1a` on a 5.0 system), creates an additional swap entry in `/etc/fstab`, and starts swapping to the newly created swap file. You will be asked to confirm items that alter your current system configuration. The script assumes that the disk is configured into the kernel, has a device special file, and that the in-memory disk label can be read.

Renice

This script alters the scheduling priority of one or more running processes. It allows you to do the following:

- Set the scheduling priority.
- Alter the priority of a process ID.
- Alter the priority of all processes for a given user
- Alter the priority of all processes for a given process group ID.

ProcessTree

This script parses the output of the UNIX `ps` command to give a tree of all processes with child processes tab indented underneath their parents.

filesize_thresh

This script makes an entry in `cron` to periodically check if a given file or directory has exceeded the specified threshold size. When a threshold is exceeded, mail will be sent to the address given with the `-m` flag and the `cron` entry will be removed automatically. The interval is limited to: 1, 5, 10, 15, 20, 30, 60 or `time_of_day` (hh:mm) in 24 hour format due to `cron` entry restrictions.

pm_fax

This script faxes a message created from the threshold environmental variables to the specified phone number. This script relies on a properly configured and functioning version of HylaFAX (see <http://www.vix.com/hylafax/> for source distribution and build information). The script was tested with `hylafax-v3.0pl1`. This script relies on the `hylafax` environmental variables being set.

pm_mail

This script will mail a threshold message read from the threshold environmental variables to the user specified on the command line. If no user is specified the message will be mail to `root`.

pm_pager

This script will send a message based on the threshold environmental variables to the specified pager phone number. This script assumes that you have a properly configured and functioning version of HylaFAX™ (see <http://www.vix.com/hylafax/> for source distribution and build information). The script was tested with `hylafax-v3.0pl1`. This script relies on the `hylafax` environmental variables being set. The pager of HylaFAX does not appear to work with the SkyTel® SkyPager® service.

pm_shutdown

This script is a wrapper for the UNIX `shutdown` command that takes a list of machines that will be shut down simultaneously. If a message is not given, a default one will be included in the shutdown invocation.

pm_broadcast

This script is a wrapper for the UNIX `rwall` command. It writes a message to all users logged on the node(s) specified in the space-separated node list.

Cluster Performance Analysis Commands

Performance Manager provides the following Cluster Performance Analysis commands. To execute one, from the main window's Execute menu choose Cluster Performance Analysis, then one of the following commands:

ClusterLoadAverage

This script determines if a cluster is working under an extreme load (3 jobs in the run queue by default) using metrics retrieved from `pmgrd` for the last 5 seconds, last 30 seconds, and the last 60 seconds. It also reports if the cluster is consumed by a small number of user processes and lists the top process.

ClusterNodeStatus

This script lists the node members of a cluster maintained by the Connection Manager. When the `-s` switch is specified, it will list the state of each node in the cluster and notify the user when a node is down or not working properly.

DLMdeadlocks

This script checks to see if the Distributed Lock Manager (DLM) locks and deadlocks exceed thresholds acceptable for a cluster system. It also compares the number of locks received with the number of locks sent to see if they are within a specified percentage of each other.

DLMlocks

This script checks to see if the Distributed Lock Manager (DLM) lock requests and messages are within a certain specified percentage of each other. The lock metrics received are compared to the number of lock metrics sent to see if the result exceeds a specified percentage.

DLMresources

This script checks to see if the Distributed Lock Manager (DLM) resources and locks exceed thresholds acceptable for a cluster system. Threshold checks made include: too many processes currently attached to the DLM, too many locks currently allocated, and too many resources currently allocated.

DRDblockingServerClient

This script checks to see if the Distributed Raw Disk (DRD) block shipping server and client operations exceed thresholds acceptable for a cluster system. These operations include number of `opens`, `closes`, `reads`, `writes`, and `ioctl`s.

DRDmemoryChannel

This script checks to see if the following Distributed Raw Disk (DRD) block shipping client memory channel operations exceed thresholds acceptable for a cluster system. These operations include number of `reads`, `writes`, and `waits` over the MC as well as number of unaligned `reads` and `writes`.

cmon

Wrapper for executing the TruCluster Version 1.0 `cmon` utility.

asemgr

Wrapper for executing the TruCluster Version 1.0 `asemgr` utility.

Threshold Management Commands

Threshold management commands can be executed when a threshold is crossed. Performance Manager provides the following threshold management commands. To execute one, from the main window's Execute menu choose Threshold Management, then one of the following commands:

SendFax

This script faxes a message created from the threshold environment variables to the specified phone number. This script relies on a properly configured and functioning version of HylaFAX (see <http://www.vix.com/hylafax/> for source distribution and build information). The script was tested with hylafax-v3.0pl1. This script relies on the hylafax environment variables being set.

SendPage

This script will send a message based on the threshold environment variables to the specified pager phone number. This script assumes that you have a properly configured and functioning version of HylaFAX (see <http://www.vix.com/hylafax/> for source distribution and build information). The script was tested with hylafax-v3.0pl1. This script relies on the hylafax environment variables being set. The pager of HylaFAX does not appear to work with the SkyTel SkyPager service.

SendMail

This script will mail a threshold message read from the threshold environmental variables to the user specified on the command line. If no user is specified the message will be mailed to `root`.

AdvFS Performance Analysis Commands

Performance Manager provides the following AdvFS Performance Analysis scripts. To execute one, from the main window's Execute menu choose AdvFS Performance Analysis, then one of the following scripts:

AdvFSDomain

This script determines if AdvFS performance can be improved by tuning some parameters. It looks at the percentage of volumes used and checks if there is any uneven usage. The balance command should be used to do any necessary balancing. The AdvFSDomain script can limit the number of volumes if necessary.

AdvFSIO

This script determines if the node has excessive AdvFS I/O problems. It looks at the number of maximum read/write blocks and the I/O write flush threshold value and checks if any of these parameters need tuning.

AdvFSTuner

This script determines if AdvFS performance can be improved by tuning some parameters. It looks at the percentage of volumes used and the buffer cache hit ratio. It checks whether the log needs to be moved to a less used volume and whether the cache needs any tuning.

Command Operations

You can execute, configure, move, add, and delete commands from the Performance Manager GUI. The example (on the following page) of an execute dialog box for CPUAnalysis shows the extent of controls you can set for command execution.

Executing Commands

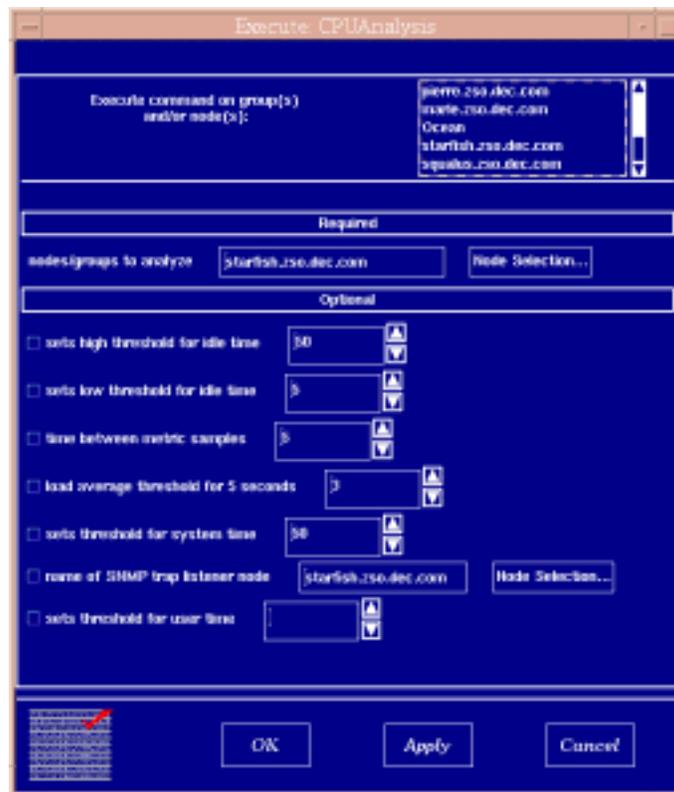
To run a command on one or more nodes, follow these steps:

- 1 Before running scripts on remote nodes, you must have a login ID and the `.rhosts` file on each remote node must give root access to the node running the Performance Manager GUI. Specify both a node alias and a fully qualified domain name. For example:

```
gui_node root
```

```
gui_node.usc.edu.com root
```

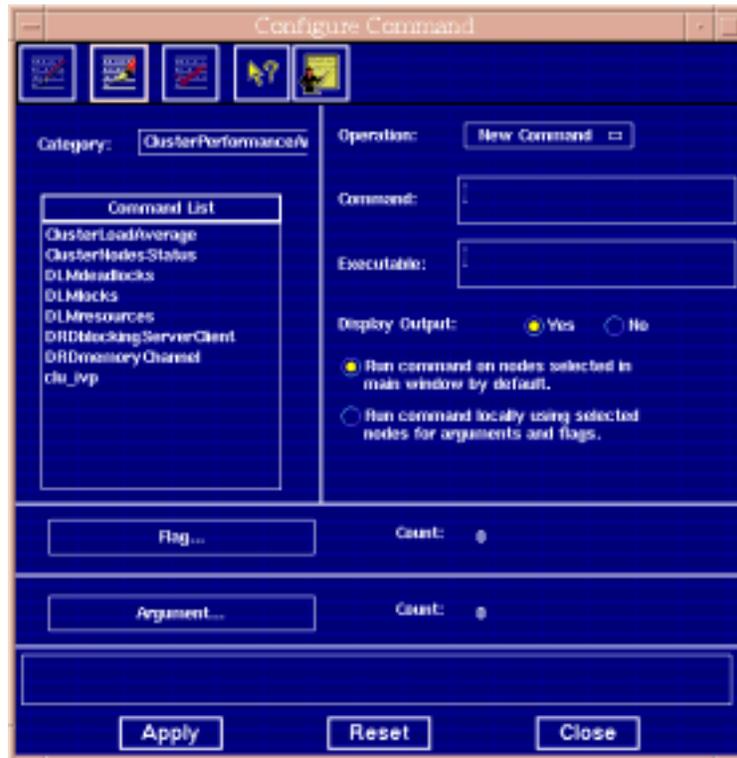
- 2 If the command does not exist on a remote node:
 - a. When the command is executed, Performance Manager copies the command from the node running the GUI to the remote node.
 - b. Executes the command.
 - c. Deletes the command on the remote node.
 - d. Any output is sent back to the node running the GUI for display in an output window.
- 3 In the main window's nodes area, select the nodes you want to run a command on. (If no nodes are selected, the command runs on the node on which the GUI is running.)
- 4 From the main window's Execute menu, choose a command to run. (You can modify these commands and add your own; from the main window's Commands menu, choose Configure.)
- 5 If the command takes any flags or arguments, an Execute window opens. Specify the flags and arguments you want, then click on the OK or Apply button to run the command.



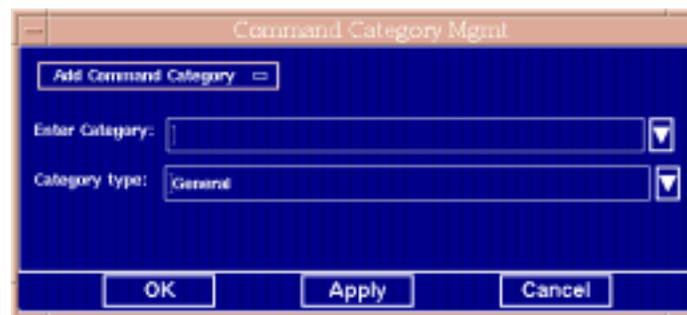
Adding Commands to the Execute Menu

To add your own commands to the Execute menu:

- 1 From the main window's Commands menu, choose Configure, which opens the Configure dialog box:



- 2 From the Category option menu, choose a command category, or choose New to create a new one. Choosing New (even if it is already visible, you must click on the word New) opens the Command Category Mgmt dialog box. Choose Add Category from the option menu, type a new category in that dialog box, and click on OK. The category you choose is the category the new command will belong to.



- 3 From the Operation option menu, choose New Command.
- 4 Click in the Command field and type a command name. Use no more than 50 characters consisting of letters, numbers, spaces, commas, underscores (_), and percent signs (%).
- 5 Click in the Executable field and type the full path of the command's executable file; for example /staff3/bin/print_page. Use no more than 50 characters consisting of letters, numbers, commas, periods, slashes (/), underscores (_), and percent signs (%).

- 6 If you choose Yes, when the command is run, a window opens containing the command's output.
- 7 Click on your choice and the radio button will change to another color.
- 8 If the command takes flags, click on the Flag button to open the Flag dialog box.
- 9 If the command takes arguments, click on the Argument button to open the Argument dialog box.

The Apply button applies any changes you made. The Reset button clears all the fields in the Configure window. The Close button closes the dialog box without applying any changes.

Deleting Commands from the Execute Menu

Follow this procedure to delete commands:

- 1 From the main window's Commands menu, choose Configure, which opens the Configure dialog box.
- 2 From the Category option menu, choose the command category containing the command to be deleted.
- 3 From the Command List, select the command to be deleted.
- 4 From the Operation option menu, choose Delete Command.
- 5 Click on the Apply button to delete the command.

Modifying Commands

Follow this procedure to modify a command:

- 1 From the main window's Commands menu, choose Configure, which opens the Configure dialog box.
- 2 From the Category option menu, choose the command category containing the command to be modified.
- 3 From the Command List, select the command to be modified.
- 4 From the Operation option menu, choose Modify Command. Make the changes to modify the command.
- 5 Click on the Apply button to modify the command.

Adding Command Categories

Follow this procedure to add a command category:

- 1 From the main window's Commands menu, choose Script Category Mgmt, which opens the Script Category Mgmt dialog box.
- 2 From the option menu, choose Add Category.
- 3 Click in the Enter Category field and type the name of the new category.
- 4 Click on the OK button.

Deleting Command Categories

Follow this procedure to delete a category:

- 1 From the main window's Commands menu, choose Script Category Mgmt, which opens the Script Category Mgmt dialog box.
- 2 From the option menu, choose Delete Category.
- 3 Click in the Enter Category field and type the name of the category to be deleted.
- 4 Click on the OK button.

Moving Commands Between Categories

Follow this procedure to move commands:

- 1 From the main window's Commands menu, choose Move, which opens the Move Command dialog box.



- 2 Choose a category from the From menu. The commands in this category will appear in the Command List.
- 3 In the Command List, select a command to be moved.
- 4 Choose a category from the To menu. This is the category the selected command will be moved into.
- 5 Click on the OK or Apply button.

Chapter 9 Archiving

Archives are files of data stored for later use. The type of data Performance Manager monitors can be saved in an archive file, then later graphed. Thus, archives allow you to capture all data on one or more nodes without having to monitor them. Should performance problems develop later, you can retrieve the archive and examine the data to see when the problem began.

Performance Manager includes scripts that store the metric data you choose in an archive file. These scripts allow you to capture all metric data on one or more nodes without having to monitor the nodes. The archived data can be replayed using Microsoft Excel or any other graphing tool you create an interface for. The information needed to archive metrics includes:

Archive duration (in minutes)

- Sample interval (in minutes)
- Type of metrics for archiving (`pmgrd`, `snmpd`, `advfsd`, `clu_mib`)
- Storage file name (the file that will contain the archived metrics)
- Storage directory (location for the `archived_host.out` archive file)
- Field delimiter used in the archive file

Later, you can graph an archive file to look at the metric data recorded.

Archive Recording

When you record an archive, Performance Manager collects all data from one or more of the nodes selected in the session and writes it to one or more files.

Archive files can become quite large. Each sample for a single-CPU, single-disk node requires 2.2 kilobytes. The total size of the file depends on the sampling interval, the number of nodes monitored, and the number of disks and CPUs on each node.

This version of Performance Manager includes sample archiving scripts for recording the metrics that Performance Manager monitors: `pm_archiver`, `pm_delta_archiver`, and `rc_archiver`. These archiver scripts are located in the `/var/opt/pm/SMScripts` directory, along with Readme files explaining their functionality.

These scripts can be executed from the command line. The `pm_archiver` script can also be executed from the Performance Manager GUI by selecting SystemManagement from the main window's Execute menu, then selecting the PMArchiver item.

Both archiver scripts archive metrics from the `snmpd`, `pmgrd`, `advfsd`, and `clu_mib` metrics servers. The archiver assumes the ports for the metrics servers are 161, 1161, 1163, and 1165, respectively. If your metrics servers run on different ports, modify the scripts accordingly.

Archive Playing

Playing an archive is like watching a recorded television show since you can skip the parts you are not interested in.

The data gathered from the `archiving` scripts can be opened directly in Microsoft Excel.

Excel will chart the data from any of the archiver scripts. When given an output file, it will allow you to choose the object that you want to plot and chart the data for all nodes. It can also plot all instances of a chosen object against time.

Chapter 10

Troubleshooting

This chapter contains information that will help you keep Performance Manager running properly.

Log Files

The Performance Manager GUI writes messages to a log file, `/var/opt/pm/log/pmgr_gui.log`. The Performance Manager metrics server (`pmgrd`) also writes messages to a log file, `/var/opt/pm/log/pmgrd.log`. These log files provide a history that is useful for troubleshooting and debugging.

The installation procedure creates initial copies of the log files with appropriate protections. For security reasons, the log directory (`/var/opt/pm/log`) is protected so that no new files can be created in it. If a log file is deleted, an appropriately protected empty file must be left in its place; otherwise, no new process (that writes to that particular log file) can be started.

To view just the last 50 lines of a log file (the GUI log file, in this example), enter the following command:

```
% tail -50 /var/opt/pm/log/pmgr_gui.log | more
```

Here is the entry format used in all log files. Each entry has three lines, the second and third lines being indented. Vertical bars separate each field in a line:

```
date_time | local_host | remote_host | user
severity | error_code | module | line_number
error_text
```

The following table describes each field in a log file entry.

Table 1 Log File Field Description

Log File Field	Description
<i>date_time</i>	The date and time the entry was written.
<i>local_host</i>	The node running the process that generated the entry.
<i>remote_host</i>	The node that originated the request. For user-interface log files, <i>remote_host</i> is always blank because there is no remote node. For metrics server log files, <i>remote_host</i> is blank only if a local event caused the entry.
<i>user</i>	The user running the application. For user-interface log files, this is the login name. For metrics server log files, this is the login name of the user on the remote node, if it is available. The field is blank if the metrics server is unable to determine the name of the application user. For metrics server messages that are not caused by a remote request, the user field is <code>Daemon</code> .
<i>severity</i>	Possible values are <code>Info</code> , <code>Warn</code> , <code>Fatal</code> , and <code>Debug</code> .
<i>error_code</i>	A string that identifies an error.

Table 1 Log File Field Description (cont.)

Log File Field	Description
<i>module</i>	The program module that generated the entry.
<i>line_number</i>	The line number in the program module where the entry originated.
<i>error_text</i>	A description of the message.

Example Log File Entry

```
October 24 11:47:03 1999|oscar.zso.dec.com||root (smith)
```

```
error|PMD_NOSUCHINST|pmdci_manager.c|line 2158
```

```
The specified instance does not exist
```

Nodes Not Responding



If a node is not responding to the Performance Manager GUI, its icon shows a hand holding the world down, as shown here.

Either the network link to that node is broken, the node has crashed, or the node doesn't exist in the network.

The installation script starts all Performance Manager metrics servers automatically after a successful installation and configuration, and these servers are started automatically at boot time. Use the startup information about these servers only if you need to restart a Performance Manager server.

Performance Manager Tru64 UNIX Metrics Server (pmgrd)

This server must run on each node managed by Performance Manager. Without `pmgrd`, the Performance Manager GUI cannot gather its data from that node.

To see if Performance Manager's Tru64 UNIX metrics server is running, issue the following command:

```
# ps awx | grep pmgrd
```

If the server is running, you should see output similar to the following:

```
329 ??S <0:16.02 bin/pmgrd
```

```
292 tty1S +0:00.03 grep pmgrd
```

If `pmgrd` is not running, it failed to start or has crashed, see the `pmgrd` log file, `/var/opt/pm/log/pmgrd.log`, for the cause. To start `pmgrd` from the Performance Manager GUI, follow these steps:

- 1 From the main window's Execute menu, choose System Management Command Category.
- 2 Choose the Start Stop Pmgrd command from this submenu.
- 3 Choose the node on which to start `pmgrd`.
- 4 Press OK or Apply to start `pmgrd` on the selected node.

To start `pmgrd` from a root account, issue the `pmgrd` command with the `start` argument:

```
# /usr/opt/pm/scripts/pmgrd start
```

If `pmgrd` is not starting at boot time, ensure that these boot-time startup files exist:

```
/sbin/rc2.d/K47pmgrd
```

```
/sbin/rc3.d/S47pmgrd
```

If they are missing, re-install the Performance Manager Daemons & Base subset (See the *Performance Manager Installation Guide*).

For more information, see the `pmgrd(8)` reference page.

Performance Manager TruCluster Metrics Server (`clu_mib`)

The TruCluster metrics server must run on each cluster where Performance Manager runs commands. Without `clu_mib`, a command cannot run on a cluster, and it cannot display its output to the Performance Manager GUI.

Beginning with Tru64 UNIX Version 5, this server ships with the operating system. In earlier releases the server shipped with the Performance Manager product. To successfully use a Version 5 system to monitor Tru64 UNIX Version 4.x systems, you must install the `clu_mib` metrics server on the monitored systems. You can ensure this configuration by installing the appropriate PM Version 4.0x on these systems.

To see if Performance Manager's TruCluster metrics server is running, issue the following command:

```
# ps awx | grep clu_mib
```

If the server is running, you should see output similar to the following:

```
329 ??S <0:16.02 bin/clu_mib
292 ttyp1S +0:00.03 grep clu_mib
```

If `clu_mib` is not running, it failed to start or has crashed, see the `clu_mib` log file, `/var/opt/pm/log/clu_mib.log`, for the cause. To start `clu_mib` from the Performance Manager GUI, follow these steps:

- 1 From the main window's Execute menu, choose System Management Command Category.
- 2 Choose the Start Stop `clu_mib` command from this submenu.
- 3 Choose the node on which to start `clu_mib`.
- 4 Press OK or Apply to start `clu_mib` on the selected node.

To start `clu_mib` from a root account, issue the `clu_mib` command with the `start` argument:

```
# /usr/opt/pm/scripts/clu_mib start
```

If `clsrtrmond` is not starting at boot time, ensure that these boot-time startup files exist:

```
/sbin/rc2.d/K47clu_mib
/sbin/rc3.d/S47clu_mib
```

If they are missing, reinstall the Performance Manager Daemons & Base subset (see the *Performance Manager Installation Guide*). The MIB file describing the metrics provided by the TruCluster metrics server is provided in this location:

```
/usr/opt/pm/data/cluster_mib
```

For more information, see the `clu_mib(8)` reference page.

Metrics Servers or GUI Will Not Start

If the GUI or metrics servers fail to start, it could be because their log files are missing. If the GUI fails to appear and there is no error message, check the `DISPLAY` environment variable and confirm that an `xhost` session is authorized.

If `pmgrd` fails to start automatically when a node is rebooted, but can be started manually, its startup files might be missing.

No Log Files

The installation procedure creates initial copies of the log files with appropriate protections. For security reasons, the log directory (`/var/opt/pm/log`) is protected so that no new files can be created in it. If a log file is deleted, an appropriately protected empty file must be left in its place; otherwise, no new process (that writes to that particular log file) can be started.

- The GUI log file is `/var/opt/pm/log/pmgr_gui.log`.
- The `pmgrd` log file is `/var/opt/pm/log/pmgrd.log`.
- The `clu_mib` log file is `/var/opt/pm/log/clu_mib.log`.

No Startup Files

The installation script writes entries in system startup files that start `pmgrd` automatically each time a node is rebooted. If `pmgrd` is not starting on a node after it is booted, check the following files and be sure they have the correct entries:

```
/sbin/rc2.d/K47pmgrd
```

```
/sbin/rc3.d/S47pmgrd
```

If they are missing, reinstall the Performance Manager Daemons & Base subset (see the *Performance Manager Installation Guide*).

Commands Not Running

If commands fail to run on certain nodes:

- 1 Make sure the nodes are up.
- 2 Before running commands on remote nodes, you must have a login ID, and the `.rhosts` file on each remote node must give `root` access to the node running the Performance Manager GUI. Specify both a node alias and a fully qualified domain name. For example:

```
gui_node root
```

```
gui_node.usc.edu.com root
```

Disks Not Visible to Performance Manager

If your kernel configuration does not match your disk configuration, Performance Manager may not recognize the disks that are not configured in the kernel. When you add disks to your system configuration, check that your kernel is configured for the new device. If needed, run the `doconfig` command to update your kernel. See the `doconfig(8)` reference page for more information.

Reporting Bugs

If an error occurs while installing or using Performance Manager, and you believe the error is caused by a problem with the product, take one of the following actions:

- If you have a basic or DECsupport™ Software Agreement, call your Customer Support Center. The Customer Support Center provides high-level advisory and remedial assistance.
- If you have a Self-Maintenance Software Agreement or you purchased Performance Manager within the past 90 days, you can submit a Software Performance Report.
- For documentation problems, casual questions, or suggestions, use the response form, or email us at `pm_feedback.zso.dec.com`.

Software Performance Reports

When you submit a Software Performance Report, please take the following steps:

- Reduce the problem to as small a size as possible.
- Describe as accurately as possible the circumstances and state of the node when the problem occurred. Include the description and version number of Performance Manager being used. Demonstrate the problem with specific examples.
- Report only one problem per Software Performance Report; this ensures a faster response.
- Mail the Software Performance Report package to Compaq.
- Many Software Performance Reports do not contain enough information to duplicate or identify the problem. Concise, complete information helps Compaq give accurate and timely service to software problems.

Glossary

archive file

A file containing data gathered by Performance Manager. Instead of watching data displayed in real time, you can capture data in an archive and graph the data later.

cron

A UNIX daemon that executes commands at a specified time. The daemon reads these commands from the `crontab` file.

cluster

A collection of nodes that appears to be a single-server system, allowing for greater application availability and scalability than would be possible with a single system.

director name

The name of one designated member of a TruCluster Production Server cluster. Performance Manager uses this value to recognize the cluster and populate the GUI with the members.

group

A collection of nodes and/or clusters that are frequently managed together.

managed node

Nodes that run one or more metrics servers recognized by Performance Manager.

management station

Nodes that are the operating centers for managing and monitoring the nodes in the network.

metric

A particular item of information about a node. For example, the average run queue length over the past 5 seconds, the number of bytes transferred to or from a disk, or the number of characters sent to a terminal. Performance Manager has several hundred metrics, divided among several categories (CPU, Disk, Network, and so on).

metrics server

A UNIX daemon process that services requests for system information. Performance Manager includes support for several metrics servers.

MIB

Management information base.

node

A computer system that is uniquely addressable on a network. A node can have more than one CPU.

rearm point

In thresholding, a specified point below the threshold. If a metric drops to this point and then recrosses the threshold, another alert will be sent.

sampling rate

In thresholding, the interval at which metric samples are taken. The interval is specified in seconds.

session

A set of choices you make using Performance Manager. A session comprises selected nodes, metrics, display types, intervals, and threshold settings. You can save as many sessions as you want, but you can only run one session at a time.

tear-off menu

A tear-off menu has an underscored key letter. If you click that letter, the menu will tear off, or float, in a separate display.

thrashing

Intensive disk activity that occurs with excessive swapping, usually indicating a memory shortage.

threshold

A limit you can set on a metric. If that limit is crossed, an action you previously specified is taken. For example, you could set a threshold of 90% capacity on some or all of your disks, with the action being to run a command that moves some files off that disk.

tolerance

A specified number of sampling intervals for which a metric must exceed its limit before a threshold is considered crossed.

Index

A

- add swap file, 44
- adding
 - clusters, 17
- adding categories, 50
- adding command categories, 50
- adding commands, 49
- Adding nodes, 16
- AddSwapFile, 44
- Advanced File System monitoring, vii
- AdvFS IO, 47
- AdvFS monitoring, vii
- AdvFS performance analysis commands, 47
 - AdvFS domain, 47
- AdvFS thresholds, 38
- AdvFS tuner, 47
- AdvFSDomain, 47
- AdvFSIO, 47
- AdvFSTuner, 47
- archive file
 - defined, 61
- archive playing, 54
- archive recording, 53
- Archiving, 53
- archiving, 1
- asemgr, 46
- auto-discovery for clusters, 19
- average archiver, 43

B

- Buffer Cache, 42
- buffer cache thresholds, 37
- buttons, 7

C

- CleanFilesystems, 43
- clu_mib
 - See* TruCluster metrics server (clu_mib)
- cluster auto-discovery, 19

- cluster name, 19
- cluster node status, 46
- Cluster performance analysis commands, 46
 - DLM resources, 46
- ClusterLoadAverage, 46
- ClusterNodeStatus, 46
- Clusters, 6
- clusters
 - defined, 61
 - discovering, 19
- cmon, 46
- Command Category Mgmt dialog box, 49
- command execution as notification method, 38
- command line monitoring, 29
- Command operations, 47
 - adding commands, 49
 - deleting, 50
 - executing commands, 47
 - moving commands, 51
- Commands, 41
 - AdvFS performance analysis
 - AdvFS domain, 47
 - cluster performance analysis
 - cluster load average, 46
 - DLM resources, 46
 - disk I/O commands
 - excessive transactions, 42
 - memory commands
 - buffer cache, 41
 - network commands
 - gateway errors, 42
 - performance analysis commands
 - CPU
 - analysis, 41
 - system management
 - archiving
 - tabular archiver, 43
 - clean file systems, 43
 - disk usage, 44
 - pm pager, 45
 - threshold management

- send fax, 47
 - send mail, 47
- commands
 - clu_mib, 57
 - pmgrd, 56
 - system management
 - PMDeltaArchiver, 44
- commands not running, 57
- comon, 46
- Configure dialog box, 49
- consolidating, 26
- CPU Analysis, 41
- CPU average archiver, 43
- CPU commands, 41
- CPU thresholds, 37
- CPU_average_archiver, 44
- creating groups, 15
- cron daemon, 45, 61

D

- deleting categories, 50
- Deleting clusters, 17
- deleting command categories, 50
- Deleting commands, 50
- Deleting groups, 16
- Deleting nodes, 16
- director name, 19, 61
- Disk I/O commands
 - excessive transactions, 42
- DiskUsage, 44
- DISPLAY environment variable, 3, 4
 - setting, 4
- displaying metrics, 33
- Displaying Performance Manager on a PC, 4
- Displays, 24
 - floating, 26
 - setting styles, 27
- DLM locks, 46
- DLMdeadlocks, 46
- DLMlocks, 46
- DLMresources, 46
- DRD blocking server client, 46
- DRD memory channel, 46
- DRDblockingServerClient, 46
- DRDmemoryChannel, 46

E

- email as notification method, 38
- environment variables
 - DISPLAY, 3, 4
 - hylafax, 47
 - PMGR_SNMP_PORT, 29
 - threshold, 47
- environmental thresholds, 38

- Excessive Paging, 42
- excessive paging, 41
- Excessive Swapping, 42
- excessive swapping, 41
- Excessive Transactions, 42
- eXcursion, 4
- Executing commands, 48
- executing commands, 2
- extensible GUI, 2

F

- file modification, 43
- file size threshold, 44
- File System Analysis, 43
- file system analysis, 42
- file system thresholds, 37
- FileModification, 43
- filesize_thresh, 45

G

- Gateway Errors, 42
- groups, 7
 - defined, 61
- growth of files, 43
- GrowthOfFiles, 43

H

- hidden metric categories, 33
- hylafax environment variables, 47

I

- Icons, 6
- impact diskmon, 43
- impact procmon, 43
- impact_diskmon, 44
- impact_procmon, 44

L

- Load Average, 41
- loadaverage, 41
- log directory, 55, 58
- log files, 55, 56
 - PM GUI (pmgr_gui.log), 3, 58
 - PM metrics server (pmgrd.log), 55, 56, 58
 - troubleshooting, 55
 - TruCluster metrics server (clu_mib.log), 57, 58

M

- maintain files, 43
- main window, 4, 6
- maintain files, 43
- MaintainFiles, 43
- managed node, 2, 61

- management information base
 - See* MIB
- management station, 1, 61
- Managing nodes
 - deleting nodes, 16
- managing nodes
 - adding clusters, 17
 - creating groups, 15
 - moving clusters, 18
- manipulating, 26
- memory commands
 - buffer cache, 41
- Memory Shortage, 42
- memory shortage, 41
- memory thresholds, 38
- menu bar, 8
- metric categories, 33
- Metrics
 - hiding categories, 34
- metrics, 21
 - defined, 61
- metrics server, 2
 - defined, 61
- MIB
 - defined, 61
 - files, 31, 57
 - variables, 2, 31
- Microsoft Excel, 43
- modifying, 50
- modifying commands, 50
- monitoring methods
 - command line, 29
 - SNMP systems, 30
- moving clusters, 18
- moving commands, 51
- Moving nodes, 17
- moving nodes, 16

N

- NetView, 30
- Network commands
 - gateway errors, 42
- Network Errors, 42
- network errors, 42
- network thresholds, 37
- Node Management dialog box, 15
- Nodes, 6
- nodes, 6
 - defined, 61
- nodes area, 6, 15
- notification methods, 38

P

- Packet Retransmissions, 42
- packet retransmissions, 42
- Performance Manager, vii
- Performance Manager daemon
 - See* Performance Manager metrics server (pmgrd)
- Performance Manager metrics server (pmgrd), 2, 31, 54, 55, 57, 58
- ping node, 43
- PingNode, 44
- PM
 - See* Performance Manager
- pm broadcast, 45
- PM Delta Archiver command, 44
- pm fax, 44
- pm mail, 44
- pm shutdown, 45
- pm_broadcast, 45
- pm_fax, 45
- pm_mail, 45
- pm_pager, 45
- pm_shutdown, 45
- PMDeltaArchiver command, 44
- PMGR_SNMP_PORT environment variable, 29
- pmgrd
 - See* Performance Manager metrics server (pmgrd)
- process tree, 44
- processes thresholds, 37
- ProcessTree, 45

R

- rearm point, 36
 - defined, 62
- Renice, 45
- renice, 44

S

- sample archiving scripts, 53
- sampling rate, 39, 62
- saving sessions, 23
- Script Category Mgmt dialog box, 50
- send page, 47
- SendFax, 47
- SendMail, 47
- SendPage, 47
- Sessions
 - creating, 22
- sessions, 21
 - defined, 62
 - managing, 23
 - recalling, 23

- saving, 23
- starting, 23
- stopping, 23
- setting, 36
- signal process, 43
- SignalProcess, 44
- SNMP Network Management Systems (NMS), 30
- System management commands
 - clean file systems, 43
 - tabular archiver, 43
- system thresholds, 37

T

- tear-off menu, 62
- thrashing, 42
 - defined, 62
- Threshold environment variables, 39
- threshold environment variables, 47
- Threshold management commands, 47
 - send fax, 47
 - send mail, 47
- Threshold Notifications dialog box, 38
- threshold work area, 36
- Thresholds
 - environmental variables, 39
 - notification, 36
 - notification methods, 38
- thresholds, 1, 21
 - AdvFS metrics, 38
 - buffer cache metrics, 37
 - CPU metrics, 37
 - defined, 62
 - environmental metrics, 38
 - file system metrics, 37
 - memory metrics, 38
 - network metrics, 37
 - processes metrics, 37
 - system metrics, 37
 - TruCluster metrics, 38
- tolerance, 35, 39
 - defined, 62
- toolbar, 8
- Troubleshooting, 55
 - disks not visible, 58
 - metric servers, 57
 - nodes not responding, 56
- troubleshooting
 - log files, 55
- TruCluster, 19
 - monitoring, vii
 - thresholds, 38
- TruCluster daemon
 - See* TruCluster metrics server (clu_mib)

- TruCluster metrics server (clu_mib), vii, 57
- TruCluster Production Server, 19
- TruCluster Server, 19

V

- Virtual Memory, 42
- virtual memory, 41

W

- Work area, 5

How Did We Do?

Please take a moment to let us know whether the online and paper documentation that we provided with this product is useful to you. We are particularly interested in comments about clarity, organization, figures, examples (or the lack thereof), the index, page layout, and, of course, accuracy.

When you send a comment, be sure to provide ample background information to help us locate the right source files. Include the product name and version, the document name, and the page number.

Internet address: doc_comments@zso.dec.com

Fax numbers: 800-DEC-FAXX (800-332-3299)
425-865-8890

Mailing address: Publications Manager
Compaq Computer Corporation
14475 NE 24th Street
Bellevue, WA 98007

