

Performance Manager for Tru64 UNIX
Version 4.0F

User's Guide

Restricted Rights: Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

© Compaq Computer Corporation 1999
All rights reserved.

The information in this document is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN; NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THE MATERIAL.

This publication contains information protected by copyright. No part of this guide may be photocopied or reproduced in any form without prior written consent from Compaq Computer Corporation.

Compaq and the Compaq logo registered United States Patent and Trademark Office.

DEC, DIGITAL, OpenVMS, Tru64, TruCluster, and the DIGITAL logo are trademarks of Compaq Computer Corporation.

Adobe, Acrobat, and Acrobat Reader are registered trademarks of Adobe Systems Incorporated. Microsoft, MS-DOS, Windows, Windows 95, and Windows NT are registered trademarks of Microsoft Corporation. Java and JDK are trademarks and Sun is a registered trademark of Sun Microsystems, Incorporated. UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

The software described in this publication is furnished under a license agreement or nondisclosure agreement, the software may be used or copied only in accordance with the terms of the agreement.

Operating System and Version: Tru64 UNIX, Version 4.0F. Future maintenance releases may require higher versions.

Software Version: Performance Manager, Version 4.0F

Date: February 1999

Revision/Update Information: This is a new manual.

Operating System and Version: Tru64 UNIX Version 4.0F. Future maintenance releases may require higher versions.

Software Version: Performance Manager, Version 4.0F

Date: February 1999

Preface

Performance Manager 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, or GUI, called `pmgr` that runs locally and can display data from the managed nodes in your Tru64UNIX 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.

For updates and the latest information about Performance Manager, see the PM Web Site at this URL:
http://www.digital.com/unix/sysman/perf_mgr/

Performance Manager is an optional subset of Tru64 UNIX (formerly known as DIGITAL UNIX) but requires a license for use in a distributed fashion. Without the license, this product acts only on the node running the graphical user interface (GUI). With the distributed license, it can act on any number of nodes in your network.

Performance Manager for Tru64 UNIX® comprises two primary components: Performance Manager GUI (`pmgr`), and Performance Manager daemon (`pmgrd`). Additional daemons are used in monitoring TruCluster™ (`clstrmond`) and the Advanced File System (`advfsd`), supplied in the AdvFS Utilities subset

Structure of This Document

This manual is organized as follows:

- 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, Oracle database support, describes GUI support and agent configuration for Oracle7.
- Chapter 11, Troubleshooting, describes creating log files, restarting daemons, solving problems, and reporting problems to Compaq.
- Glossary describes terms specific to Performance Manager.
- Index.

Related information

In addition to this guide, the Performance Manager documentation set includes the following manuals and other documents:

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

Related manuals

The following manuals are part of the base operating system documentation set:

- *Tru64 UNIX Installation Guide*
- *Tru64 UNIX Software License Management*
- *Compact Disc User's Guide*

Conventions

The following conventions are used in this guide:

Convention	Meaning
UPPERCASE and lowercase	The DIGITAL 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.

Contents

Preface	iii
Structure of This Document	iii
Related information	iv
Related manuals	iv
Conventions	iv
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
More introductory information	3
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	7
Modifying the main window	10
3 Managing nodes	11
Creating groups	11
Deleting groups	12
Adding nodes	12
Deleting nodes	12
Moving nodes	13
Adding clusters	13
Deleting clusters	13
Moving clusters	14
4 Displaying clusters	15
Display representation of clusters	15

Auto-discovery for clusters	15
Possible anomalies	16
5 Monitoring	17
Sessions	17
Creating a session	18
Displays	19
Floating displays	20
Consolidating displays	21
Manipulating displays	21
Setting display styles	22
Other monitoring methods	23
Monitoring from the command line	23
Monitoring with SNMP network management systems	24
6 Metrics	27
Showing hidden metric categories	27
Displaying additional categories in the list	27
Hiding metric categories	28
7 Thresholds	29
Threshold notifications	29
Setting thresholds	30
CPU thresholds	30
System thresholds	30
Buffer cache thresholds	31
Network thresholds	31
File system thresholds	31
Memory thresholds	31
AdvFS thresholds	31
TruCluster thresholds	31
Oracle database thresholds	31
Environmental thresholds	31
Advanced threshold dialog (more...) box	32
Threshold environment variables	32
8 Commands	33
Performance analysis commands	33
CPU commands	33
Memory commands	33
Network commands	34
Disk I/O commands	34
System management commands	35
Cluster performance analysis commands	37
Threshold management commands	38
AdvFS performance analysis commands	39
Command operations	39
Executing commands	39
Adding commands to the Execute menu	40

Deleting commands from the Execute menu	41
Modifying commands	41
Adding command categories	42
Deleting command categories	42
Moving commands between categories	42
9 Archiving	43
Archive recording	43
Archive playing	44
10 Oracle database support	45
Oracle7 GUI support	45
Display categories	45
Threshold category:	45
Agent configuration	45
Configuration procedure	46
Oracle commands	47
11 Troubleshooting	51
Log files	51
Log file field descriptions	51
Example log file entry	52
Nodes not responding	52
Metrics servers or GUI will not start	53
Commands not running	54
Disks not visible to Performance Manager	54
Reporting bugs	54
Glossary	57
Index	59

Chapter 1 Overview

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




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
- Database performance

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. This is currently implemented with the `pm_archiver` command.
- 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.
- Database support: Oracle7 support in Performance Manager enables you to manage Oracle7 database server and its services. This support includes:
 - A set of Oracle7-related displays and thresholds in the Performance Manager GUI.
 - A recommended configuration and procedure for managing the coexistence of Peer master SNMP agent used by Oracle7 and eSNMP master agent native to Tru64 UNIX.
 - 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 metric 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
- `clstrmond` Provides TruCluster-related metrics
- `os_mibs` Provides MIB-II metrics
- `advfsd` Provides AdvFS-related metrics

The `pmgrd` and `clstrmond` metrics servers are 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 also 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 11, 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.

More introductory information

These topics will help introduce you to Performance Manager's operations.

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
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 Compaq 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

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 Performance Manager main window

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



Figure 2 Entry window

Work area

The Display and Threshold buttons each have a specific work area, showing related categories, metrics, and options.

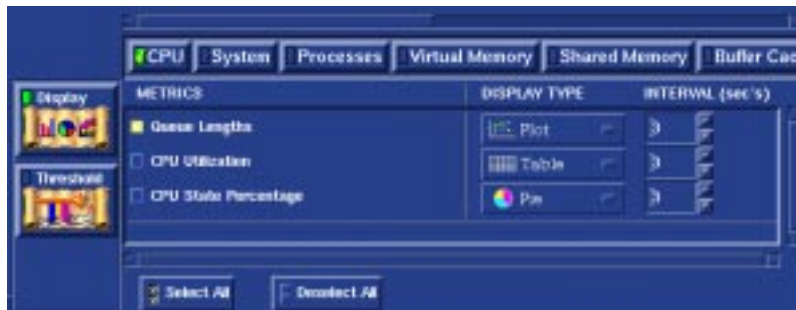


Figure 3 Display work area



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 horses 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:



No knight: Node is down or invalid.



Knight standing: Node is up.



Knight on horse, 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 cluster icon changes to reflect one of the following three states:

- No knights: All nodes in the cluster are down or invalid.
- Knights standing: At least one node in the cluster is up.
- Knights on horses: Nodes are up.

A check mark indicates that metrics have been selected for 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. Castles 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.



Shields. A shield appears next to each castle (group) and set of three horses (cluster). A shield displaying a minus sign (–) shows that all nodes in the group or cluster are exposed. A shield showing the plus sign (+) shows that all nodes are hidden. Clicking on that shield exposes 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.

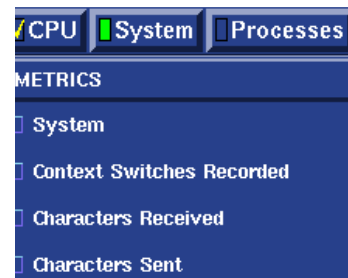


Figure 6 Metrics Selection



When this button is on, the display work area is shown.



When this button is on, the threshold work area is shown.



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

- New Session

Clicking this link displays information on sessions.

- Open Session

Clicking this link displays information on sessions.

- Save Session


Clicking this link displays information on sessions.

- Save Session As

Clicking this link displays information on sessions.

- Exit

Clicking Exit quits the session.



- View

- 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

- Enable Tool Bar Label

Displays a label as the cursor passes over each toolbar icon.



- Tasks

- 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**

- **Configure**

The Configure dialog box integrates your commands with Performance Manager.

- **Move**

This dialog enables you to regroup commands in different categories.

- **Command Category Mgmt**

This dialog enables you to add or delete command categories.



- **Execute**

- **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.

- **Cluster Performance Analysis**

These commands analysis cluster performance.



- **Help**

- **Overview**

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

- **Tasks**

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

- **Reference**

Clicking 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**


Clicking 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.

- **Using Help**

Clicking Using Help opens the CDE help volume, which explains how the help system works.

- **About Performance Manager**

Clicking About Performance Manager opens the help window containing information about this software version, copyrights, and trademarks.



Toolbar icons

The toolbar icons are arranged by groups and represent these actions:

- File Group

Clicking these links displays information on sessions.



New Session



Open Session




Save Session As

- Task Group

- Node management provides access to the controls for adding, deleting, and moving nodes and clusters.
- Use category management to make metric categories visible or hidden. Visible categories are selectable for monitoring.
- Threshold notification presents a list of activity with a reporting window.



Node Management



Category Management



Threshold Notification

- Command Group

- The Configure dialog box integrates 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 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.
- Clicking Overview 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
```

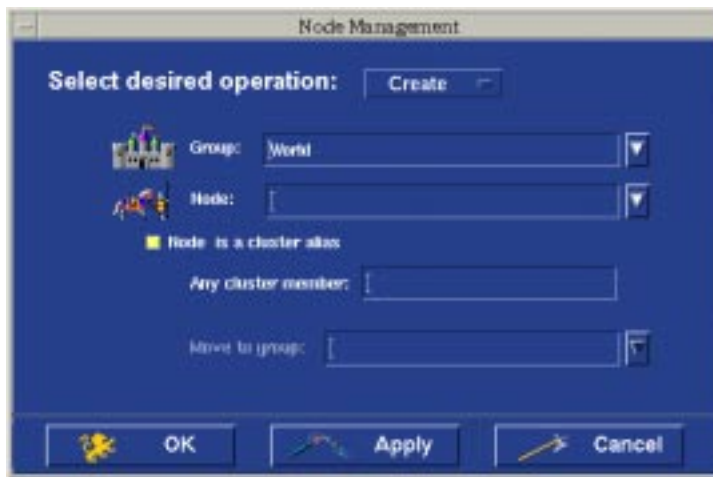
You might want to do this to provide greater viewing contrast, but be careful not to choose a color that will obscure text, such as a black foreground that hides black text.

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`.

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. Monitoring nodes means looking at performance data in real time. From the main window's Tasks menu or toolbar, choose Node Management, which opens the Node Management window.

See the individual task descriptions for specific procedure steps. All tasks begin from the node management window.



Node management window

The Apply button applies any changes you made.



Apply button

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



OK button

The Cancel button dismisses the window without applying any changes.



Cancel button

Creating groups

Create groups to organize your nodes in the main window's nodes area.

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.

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.

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 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.

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 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.

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 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:

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 field and type the name of the cluster to be added.
- 5 Set the Node Is a Cluster Alias check box to the *on* position.
- 6 Click in the Any Cluster Member field and type the name of one node in the cluster; the other cluster nodes will automatically be added to the cluster.
- 7 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.

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 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.

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 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 new auto-discovery feature.

Display representation of clusters

When monitoring a cluster, Performance Manager accesses the members of the cluster from the Connection Manager on the cluster. When the membership changes, Performance Manager adjusts its representation of the cluster as follows:

- If a node was added, 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.

Auto-discovery for clusters

When you add a node, the checks to see if it belongs to a cluster. If it does, then it creates and names a cluster node using the director name for the cluster. If the create succeeds or the cluster node already exists, the cluster node determines which nodes are in the cluster and creates them. If the node already exists it makes them children of itself.

The cluster node then requests cluster information (director name and cluster members) from one of its children. It saves the name of this node when the session is saved. If the cluster node fails to get the cluster information from this child node, it will remove it as described below, and switch to another child for information. If there are no other children it will still remove the failing child, but will continue attempting to get the information from that node, using the saved node name. Occasionally a node removed from the cluster node in this way is still actually part of the cluster. In that case the cluster node will bring it back in when it gets an updated list of members from one of the other nodes.

If membership changes, the cluster node changes its children as follows:

- New members are created or moved from their current group, becoming children of the cluster node.
- Deleted members are removed from the cluster and moved to the group that the cluster belongs to.

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.

Possible anomalies

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 node(s) removed from the cluster notices the new director name 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 section introduces metrics and thresholds, explains sessions and the types of displays you can choose, and includes information on additional monitoring methods.

- 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.

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 have in a session.

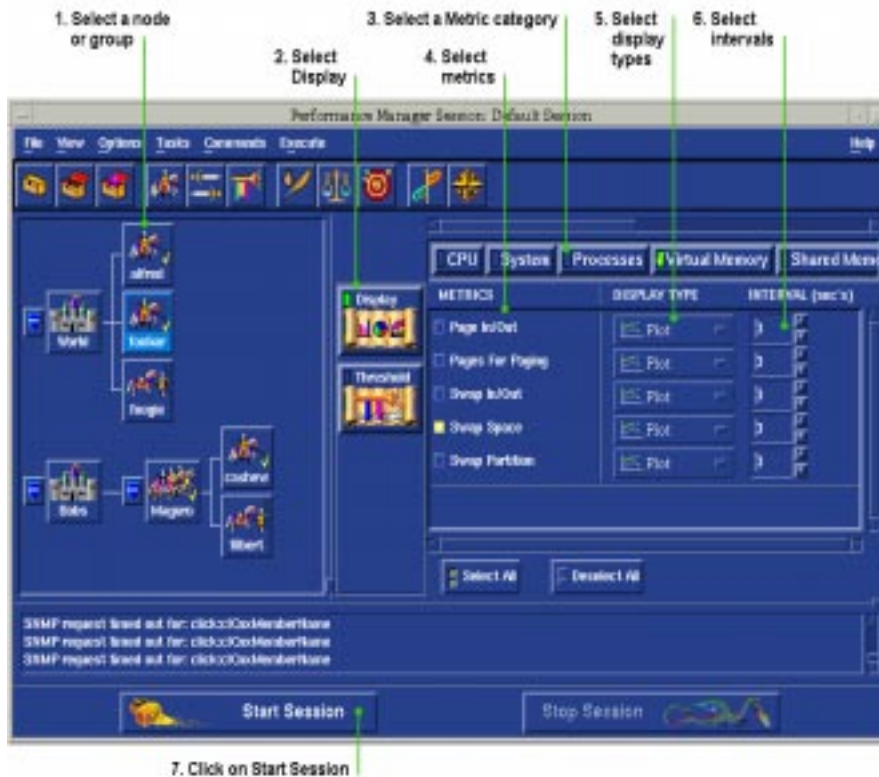


Figure 7 Session view

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.
- 7 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
- 8 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.

- Float

Click this button to detach (float) this window.

A visual alert icon by the title displays the *State* (crossed or not crossed, waiting for data, data request timed out) of the threshold.

9 Repeat the steps (except step 1) for every node, cluster, or group you want to monitor.

10 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. Starting the session puts everything in motion: the displays you specified will open and the thresholds you specified will be set.

When you are through, stop the session. 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

From the File menu, choose Save Session or Save Session As.

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

To recall a previous session

From the File menu, choose Open Session.

- From the main window's File menu, choose Open Session. The File Selection dialog box opens.
- 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. The following images are examples of each display type:




Figure 8 Area display

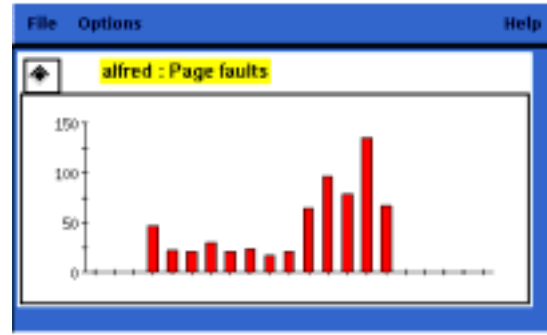


Figure 9 Bar display

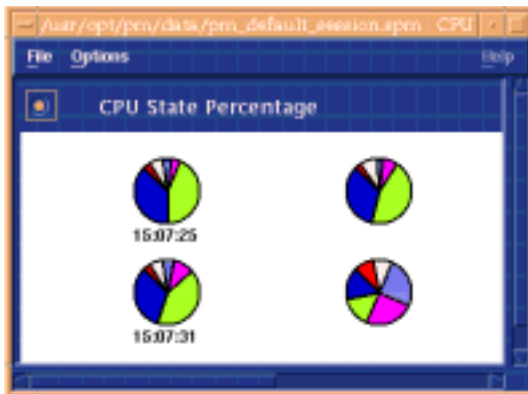


Figure 10 Pie display

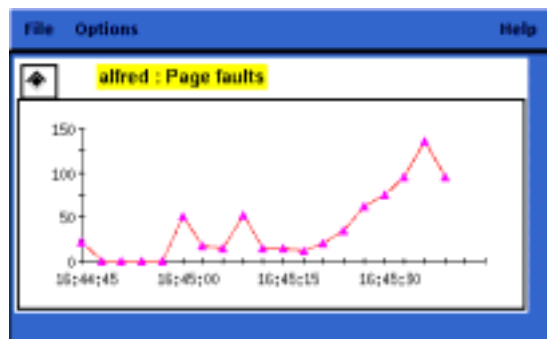


Figure 11 Plot display

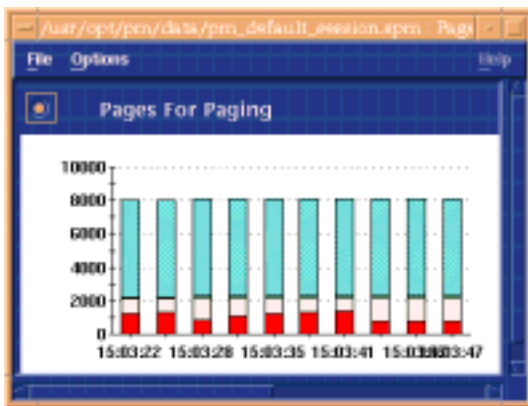


Figure 12 Stack Bar display

Name	User	Nic	System	Idle	Total
cc.com	cpu 0	0	327	2516	0
cc.com	cpu 1	0	584	2092	0

Figure 13 Table display

Floating displays

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

To expand or collapse a display

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

To float a display

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



Figure 14 Metric display selection

- 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

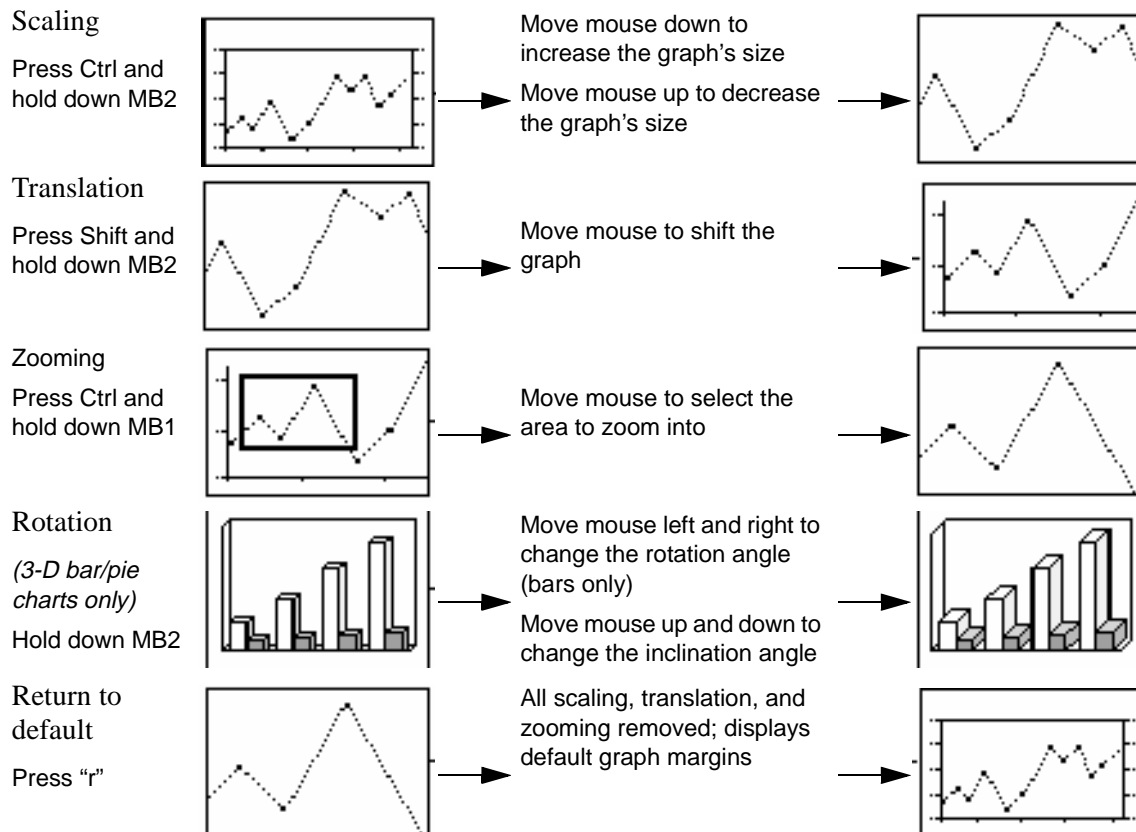
Floating displays can be closed so that they reappear in the session window.

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.
- A visual alert icon by the title displays the State (crossed or not crossed, waiting for data, data request timed out) of the threshold.

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.
- `fpattern` — 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 Intrinsics documentation.




Figure 15 Plot line pattern, fill pattern, and point style

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` ▪ `getbulk`
- `getnext` ▪ `gettab`
- `getmany` ▪

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
Compaq 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`

NetView

1 Installing and deinstalling NetView support

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

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

To uninstall NetView support, run the following as super user:

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

2 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 metric server name followed by the NetView-loadable MIB file:

```

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

```

3 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 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 are found under `.iso.org.dec`

Note *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 Up Tree until you reach *org* and then go down *dec* to find the PM MIB variables.

4 Sending SNMP Traps using `trapsend(1)`

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 the `mgr.cnf` and `snmpinfo.dat` in the `/etc/srconf/agt` directory.

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

5 Sample MIB Applications

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

Filename	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.

Showing hidden metric categories

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



Displaying additional 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:

- The event is logged (written in the Performance Manager log file: `/var/opt/pm/log/pmgr_gui.log`).
- 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.

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 16 Threshold notification

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 17 Threshold metric categories

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



Figure 18 More... button

CPU thresholds

You can set the 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 metrics:

- 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

Oracle database thresholds

You can set thresholds for the following Oracle server metrics:

- Oracle Agent is Down
- Redo Log Space Wait Ratio
- Transaction Rate
- Cache Hit Ratio
- Row Source Ratio
- User Call Rate

Environmental thresholds

You can set thresholds for the following environmental metrics:

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

Advanced threshold dialog (more...) 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 checkbox *on*.
 - Threshold Notification 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 setting 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 `$` symbol in front of the variable, for example, `$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. The Configure dialog box is used 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.

CPU commands

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

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

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

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 files' permissions
- Change files' user and group ownership
- 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 idletime, 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 wish to average and determine whether you want the output file named by metric or machine.

Performance Manager will wait while these scripts run, only closing after they have reached completion. If you set a duration longer than the time you wish to run the PM GUI, you can run the scripts 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 clstrmond 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`), 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-v3.0p11`. 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-v3.0p11`. This script relies on the `hylafax` environmental variables being set. The pager of HylaFAX does not appear to work with the SkyPage 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 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-v3.0p11`. This script relies on the `hylafax` environmental variables being set.

SendPage

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-v3.0p11`. This script relies on the `hylafax` environmental variables being set. The pager of HylaFAX does not appear to work with the SkyTel SkyPage 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
```

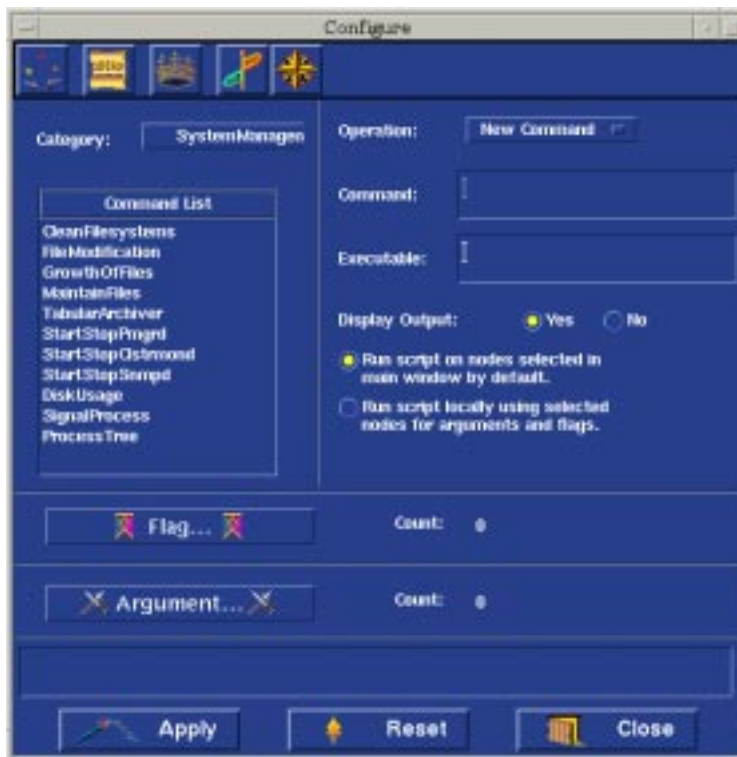
- 2 If the command does not exist on a remote node:
 - When the command is executed, Performance Manager copies the command from the node running the GUI to the remote node.
 - Executes the command.
 - Deletes the command on the remote node.
 - Any output is sent back to the node running the GUI for 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 wish, 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:



- 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 Script 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.



Category management window

- From the Operation option menu, choose New Command.
- 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 (%).
- 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 (%).
- If you choose **Yes**, when the command is run, a window opens containing the command's output.
- Click on your choice and the radio button will change to another color.
- If the command takes flags, click on the Flag button to open the Flag dialog box.
- 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:

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

Modifying commands

Follow this procedure to modify a command:

- From the main window's Commands menu, choose Configure, which opens the Configure window.
- From the Category option menu, choose the command category containing the command to be modified.
- From the Command List, select the command to be modified.
- 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

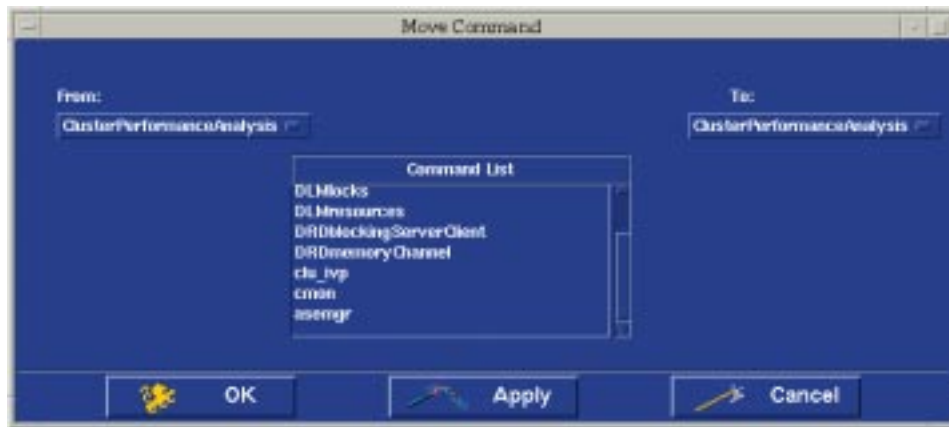
Follow this procedure to delete a category:

- 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`, `smnpsd`, `advfsd`, `clstrmond`)
- 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 `clstrmond` daemons. The archiver assumes the ports for the daemons are: 161, 1161, 1163, and 1165 respectively. If your daemons 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

Oracle database support

Performance Manager's support of Oracle7 enables you to manage the database server and its services with GUI tools and SNMP configuration. Oracle SNMP support is a new feature introduced in Oracle7 Server release 7.2 with SQLNet 2.2; Performance Manager supports release 7.2 and later.

Oracle7 GUI support

There are two display categories and one threshold category:

- Oracle Ratios (display)
- RDBMS Tables (display)
- Oracle Ratios (threshold)

Display categories

Oracle Ratios

This category displays the most useful ratios for tuning the performance of an Oracle7 database instance.

RDBMS Tables

These are easy-to-read tables that contain the most important information about an Oracle7 database.

Threshold category:

Oracle ratios

Set up thresholds over some of the most useful ratios for tuning the performance of an Oracle7 database instance.

Agent configuration

Oracle7 SNMP services are based on Peer Networks' SNMP master technology; however, in Tru64 UNIX, eSNMP master technology is native. These two technologies are not compatible in nature, but Peer offers a mechanism called Peer encapsulator that can connect any third-party SNMP agents to Peer master agent.

This method puts Tru64 UNIX's native eSNMP master agent on a private port, and then reregisters all its MIB registrations to Peer master agent through Peer encapsulator. These relationships are diagrammed in the following drawing.

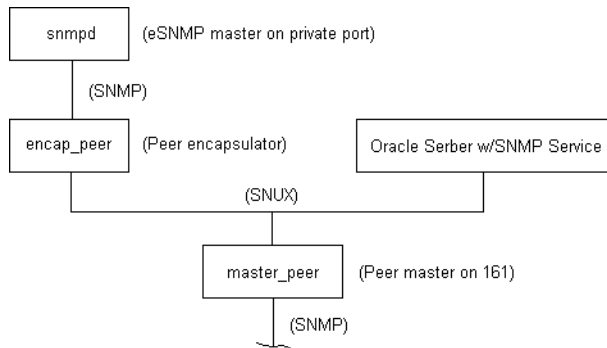


Figure 19 SNMP diagram

Configuration procedure

To configure the agent:

- 1 Log in as `root` with Oracle7 installed.
- 2 Stop eSNMP master agent if it's up running by entering the following command:
/sbin/init.d/snmp stop
- 3 Start Peer master agent by entering the following command:
start_peer -m

Note If you are not concerned about configuring eSNMP, skip to step 10.

- 4 Start eSNMP master agent on a private port by entering the following command:
 For Tru64 UNIX version 3.2 or lower:

native_snmpd.sh

See Oracle commands for the full script.

For Tru64 UNIX version 3.2C or higher:

native_esnmpd.sh

See Oracle commands for the full script.

Note You need the eSNMP patch to sit the eSNMP master agent on a private port if you are using Tru64 UNIX versions 3.2C, 3.2D, 3.2E, 3.2F or 4.0. The patch is not needed for versions 3.2G, 4.0x or later versions.

- 5 Restart all the eSNMP subagents, if necessary.
- 6 Dump out all the MIB registrations of the eSNMP master agent, which causes the eSNMP master agent (`snmpd`) to dump the contents of its registry into `/var/tmp/snmpd_dump.log`. Do this by running the following command:
/sbin/init.d/snmp dump
- 7 Reformat the dumped contents into `CONFIG.encap` file format with the same port number as the new `NEW_SMNPD_PORT` above.
- 8 Start Peer encapsulator by running the following command.:

start_peer -e

- 9 Start the desired SNMP services offered by Oracle7 Server.

Note Peer encapsulator does not always function correctly in Tru64 UNIX. If this problem occurs, you have to give up the services offered by eSNMP master agent if you choose to monitor Oracle7 Service via SNMP.

For more information, see the following:

- Oracle SNMP Support Reference Guide, release 2.2 or higher.
- Oracle7 for DIGITAL UNIX, Installation and Configuration Guide, release 7.2.3 or higher.

Oracle commands

These are the scripts you need for configuring Oracle7 support.

The native_esnmp.sh script for DIGITAL UNIX 3.2C or higher

```
#!/bin/sh
#

# variables for modifying /etc/services
#
ORG_SNMPD_PORT=161# port that snmpd listens to. Don't modify.
ORG_TRAPD_PORT=162# port that snmpd send trap to. Don't modify.
NEW_SNMPD_PORT=1611# new port that snmpd listens to. Make sure this is # the
same as the one you use in CONFIG.encap
NEW_TRAPD_PORT=1612 # peer encapsulator port that snmpd sends trap to.

ETC_SERVICES=/etc/services
SAV_SERVICES=/etc/services.org.$$
TMP_SERVICES=/tmp/services.$$
CLRCMD="/bin/rm -f $TMP_SERVICES $SAV_SERVICES"
TRPCMD="if [ -x $SAV_SERVICES ]
then
mv $SAV_SERVICES $ETC_SERVICES;
fi
$CLRCMD;
exit 1"

$CLRCMD
trap "$TRPCMD" 1 2 3 9 15

sed "s/$ORG_SNMPD_PORT\/udp/$NEW_SNMPD_PORT\/udp/g;
s/$ORG_TRAPD_PORT\/u dp/$NEW_TRAPD_PORT\/udp/g" $ETC_SERVICES
$TMP_SERVICES
```

```

echo "Modifying $ETC_SERVICES for snmpd startup..."
echo "cp $ETC_SERVICES $SAV_SERVICES"
cp $ETC_SERVICES $SAV_SERVICES
echo "cp $TMP_SERVICES $ETC_SERVICES"
cp $TMP_SERVICES $ETC_SERVICES

echo "Starting snmpd ..."
echo "/sbin/init.d/snmpd start 2&1"
/sbin/init.d/snmpd start 2&1

echo "Restoring $ETC_SERVICES ..."
echo "mv $SAV_SERVICES $ETC_SERVICES"
mv $SAV_SERVICES $ETC_SERVICES
echo "Done!"
echo

$CLRCMD

```

The native_snmpd.sh script for DIGITAL UNIX 3.2 or lower

```

#!/bin/sh
#
# variables for modifying /etc/services
#
ORG_SNMPD_PORT=161# port that snmpd listens to. Don't modify.
ORG_TRAPD_PORT=162# port that snmpd send trap to. Don't modify.
NEW_SNMPD_PORT=1611# new port that snmpd listens to. Make sure this is
# the same as the one you use in CONFIG.encap
NEW_TRAPD_PORT=1612# peer encapsulator port that snmpd sends trap to.

ETC_SERVICES=/etc/services
SAV_SERVICES=/etc/services.org.$$
TMP_SERVICES=/tmp/services.$$
CLRCMD="/bin/rm -f $TMP_SERVICES $SAV_SERVICES"
TRPCMD="if [ -x $SAV_SERVICES ]
then

```



```
mv $SAV_SERVICES $ETC_SERVICES;
fi
$CLRCMD;
exit 1"

$CLRCMD
trap "$TRPCMD" 1 2 3 9 15

sed "s/$ORG_TRAPD_PORT\/udp/$NEW_TRAPD_PORT\/udp/g" $ETC_SERVICES
$TMP_SERVICES

echo "Modifying $ETC_SERVICES for snmpd startup..."
echo "cp $ETC_SERVICES $SAV_SERVICES"
cp $ETC_SERVICES $SAV_SERVICES
echo "cp $TMP_SERVICES $ETC_SERVICES"
cp $TMP_SERVICES $ETC_SERVICES

echo "Starting snmpd ..."
echo "/usr/sbin/snmpd -p $NEW_SNMPD_PORT 2&1"
/usr/sbin/snmpd -p $NEW_SNMPD_PORT 2&1

echo "Restoring $ETC_SERVICES ..."
echo "mv $SAV_SERVICES $ETC_SERVICES"
mv $SAV_SERVICES $ETC_SERVICES
echo "Done!"
echo

$CLRCMD
```


Chapter 11 Troubleshooting

Log files

The Performance Manager GUI writes messages to a log file, `/var/opt/pm/log/pmgr_gui.log`. The Performance Manager daemon (`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), issue 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.

Log file field descriptions

Log file entries have this form:

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

Log file field	Description
<code>date_time</code>	The date and time the entry was written.
<code>local_host</code>	The node running the process that generated the entry.
<code>remote_host</code>	The node that originated the request. For user interface log files, <code>remote_host</code> is always blank because there is no remote node. For daemon log files, <code>remote_host</code> is blank only if a local event caused the entry.

Log file field	Description
user	The user running the application. For user interface log files, this is the login name. For daemon log files, this is the login name of the user on the remote node, if it is available. The field is blank if the daemon is unable to determine the name of the application user. For daemon messages that are not caused by a remote request, the user field is Daemon.
severity	Possible values are Info, Warn, Fatal, and Debug.
error_code	A string that identifies an error.
module	The program module that generated the entry.
line_number	The line number in the program module where the entry originated.
error_text	A description of the message.

Example log file entry

```
Oct 24 11:47:03 1997|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 has a riderless horse:



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 aux | grep pmgrd
```

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

```
329 ??S <0:16.02 bin/pmgrd
292 ttyplS +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 the 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 (clstrmond)

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

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

```
# ps aux | grep clstrmond
```

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

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

```
292 ttyplS +0:00.03 grep clstrmond
```

If `clstrmond` is not running, it failed to start or has crashed, see the `clstrmond` log file, `/var/opt/pm/log/clstrmond.log`, for the cause. To start `clstrmond` 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 Clstrmond command from the this submenu.
- 3 Choose the node on which to start `clstrmond`.
- 4 Press OK or Apply to start `clstrmond` on the selected node.

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

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

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

```
/sbin/rc2.d/K47clstrmond
```

```
/sbin/rc3.d/S47clstrmond
```

If they are missing, re-install 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 `clstrmond(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 `pmgrd` fails to start automatically when a node is rebooted, but can be started manually, its startup files might be missing.

No log file

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 `clstrmond` log file is `/var/opt/pm/log/clstrmond.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, re-install 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

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.

group

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

managed node

Managed nodes are those that run one or more metrics servers recognized by Performance Manager.

management station

Management Stations are the operating centers for managing and monitoring the nodes in the system.

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.

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.

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.

Index

A

- add swap file, 36
- adding categories, 41
- Adding clusters, 13
- adding command categories, 42
- adding commands, 40
- Adding nodes, 12
- adding nodes, 11
- AddSwapFile, 36
- advanced dialog, 31
- AdvFS, 31
- AdvFS IO, 39
- AdvFS performance analysis commands, 39
 - AdvFS domain, 39
- AdvFS tuner, 39
- AdvFSDomain, 39
- AdvFSIO, 39
- AdvFSTuner, 39
- Agent configuration, 45
- archive playing, 44
- archive recording, 43
- Archiving, 43
- asemgr, 38
- Auto-discovery for clusters, 15
- average archiver, 35

B

- Buffer Cache, 33
- buttons, 7

C

- CleanFilesystems, 35
- cluster node status, 37
- Cluster performance analysis commands, 37
 - DLM resources, 38
- ClusterLoadAverage, 37
- ClusterNodeStatus, 37
- Clusters, 6
- cmon, 38

- command line monitoring, 23
- Command operations, 39
 - adding commands, 40
 - deleting, 41
 - executing commands, 39
 - moving commands, 42
- Commands, 33
 - AdvFS performance analysis
 - AdvFS domain, 39
 - cluster performance analysis
 - cluster load average, 37
 - DLM resources, 38
 - disk I/O commands
 - excessive transactions, 34
 - memory commands
 - buffer cache, 33
 - network commands
 - gateway errors, 34
 - performance analysis commands
 - CPU
 - analysis, 33
 - system management
 - archiving
 - tabular archiver, 35
 - clean file systems, 35
 - disk usage, 36
 - pm pager, 37
 - threshold management
 - send fax, 38
 - send mail, 38
- commands not running, 53
- comon, 38
- configuration procedure, 46
- consolidating, 20
- CPU Analysis, 33
- CPU average archiver, 35
- CPU commands, 33
- CPU_average_archiver, 36
- Creating groups, 11

D

- deleting categories, 41
- Deleting clusters, 13
- deleting command categories, 42
- Deleting commands, 41
- Deleting groups, 12
- deleting groups, 11
- Deleting nodes, 12
- Disk I/O commands
 - excessive transactions, 34
- DiskUsage, 36
- Displaying Performance Manager on a PC, 4
- Displays, 19
 - floating, 20
 - setting styles, 22
- DLM locks, 37
- DLMdeadlocks, 37
- DLMlocks, 38
- DLMresources, 38
- DRD blocking server client, 38
- DRD memory channel, 38
- DRDblockingServerClient, 38
- DRDmemoryChannel, 38

E

- environmental, 31
- Excessive Paging, 34
- excessive paging, 33
- Excessive Swapping, 34
- excessive swapping, 33
- Excessive Transactions, 34
- Executing commands, 39

F

- file modification, 35
- file size threshold, 36
- file system, 30
- File System Analysis, 34
- file system analysis, 34
- FileModification, 35
- filesize_thresh, 37

G

- Gateway Errors, 34
- Groups, 6
- growth of files, 35
- GrowthOfFiles, 35

H

- hidden metric categories, 27

I

- Icons, 6
- impact diskmon, 35
- impact procmon, 35

- impact_diskmon, 36
- impact_procmon, 36

L

- Load Average, 33
- loadaverage, 33
- log files, 52

M

- maintain files, 35
- maintain files, 35
- MaintainFiles, 35
- Managing nodes
 - creating groups, 11
 - deleting nodes, 12
- manipulating, 20
- memory commands
 - buffer cache, 33
- Memory Shortage, 34
- memory shortage, 33
- menu bar, 7
- Metrics, 17
 - hiding categories, 28
 - modifying, 41
 - modifying commands, 41
- monitoring methods
 - command line, 23
 - SNMP systems, 24
- Moving clusters, 14
- moving commands, 42
- Moving nodes, 13
- moving nodes, 12

N

- NetView, 24
- network, 30
- Network commands
 - gateway errors, 34
- Network Errors, 34
- network errors, 34
- Nodes, 6
- notification methods, 32

O

- Oracle, 31, 45
- Oracle commands, 47
- Oracle database support, 45
 - Agent configuration, 45
 - agent configuration
 - Oracle commands, 47
 - procedure, 46
- Oracle ratios, 45
- Oracle7 GUI support
 - display categories
 - Oracle ratios, 45
 - threshold category

Oracle ratios, 45

P

Packet Retransmissions, 34
packet retransmissions, 34
ping node, 35
PingNode, 36
pm broadcast, 37
pm fax, 36
pm mail, 36
pm shutdown, 37
pm_broadcast, 37
pm_fax, 37
pm_mail, 37
pm_pager, 37
pm_shutdown, 37
process tree, 36
process.buffer cache, 30
ProcessTree, 36

R

RDBMS, 45
RDBMS tables, 45
recalling, 19
Renice, 36
renice, 36

S

saving, 19
send page, 38
SendFax, 38
SendMail, 38
SendPage, 38
Sessions, 17
 creating, 18
 starting, 19
setting, 29
Shields, 7
signal process, 35
SignalProcess, 36
stopping, 19
system, 30
System management commands
 clean file systems, 35
 tabular archiver, 35

T

Threshold environment variables, 32
Threshold management commands, 38
 send fax, 38
 send mail, 38
Thresholds, 17
 CPU, 30
 environmental variables, 32
 memory, 31

notification, 29
notification methods, 32
toolbar, 7
Troubleshooting, 51
 disks not visible, 54
 metric servers, 53
 nodes not responding, 52
TruCluster, 31

V

Virtual Memory, 34
virtual memory, 33

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
Digital Equipment Corporation
14475 NE 24th Street
Bellevue, WA 98007

