



Chapter 1

Introduction

Advanced Computing applications of current interest consume computing power of the order of several GigaFLOPS. Vector Processing Systems, Massively Parallel Systems and Cluster systems are the means of achieving such orders of computing performance.

Cluster computing systems consist of relatively inexpensive workstations connected by means of a network. Cluster computing systems, apart from being economical, also offer a high degree of performance scaling.

Systems Managers of large cluster systems often find it difficult to monitor such systems as it requires the manager to initiate sessions on each such workstation in the cluster that needs to be monitored. In addition, lengthy commands with specific options need to be remembered to perform such tasks. Though several monitoring tools are available, they do not cater to the needs of cluster systems. PARMON is a tool that has been designed to provide Managers with a graphical interface and allows monitoring of the systems in the cluster from a single workstation.

1.1 About PARMON

PARMON is meant for managers/administrators of Cluster computing systems. However, it can also be utilised by system managers of single workstations.

In order to achieve cross-platform compatibility, the tool has been divided into two parts - the server program and the client program. The server program runs on the nodes in the cluster that need to be



monitored. The client program is based on the java and thus ensures compatibility across hardware and software platforms.

The tool allows a manager to monitor the nodes in a cluster from a single client doing away with the requirement of opening sessions on each of the nodes.

1.2 Features of the Tool

- Cross platform compatible client
- User-friendly GUI
- Multi-user support
- Comprehensive monitoring support

1.3 Manual Organization

The manual has been divided into several chapters:

- Introduction
 - This chapter briefly introduces the cluster computing model and the need for cluster monitoring systems.
- Resource & Setup Requirements
 - This chapter details the software and hardware requirements and the procedure to setup the application.
- Working with PARMON
 - This chapter details the components of the user interface along with essential details about configuration and use of PARMON.



Chapter 2

Resource & Setup Requirements

The PARMON tool is based on the Client-Server model in keeping with the concept of the Open architecture. The server program is an ELF 32-bit MSB SPARC executable binary and the client program is a java-based client.

2.1 Resource Requirements

2.1.1 Server Program

The Server Program requires the following minimum requirements:

- A SuperSPARC / UltraSPARC workstation
- About 250 KB of free disk space for the server program
- At least 32 MB of system RAM (64 MB recommended)

2.1.2 Client Program

A system can function as a client if it fulfils the following minimum requirements:

- Capable of running Java run-time ver 1.1.4 or higher
- The client (system) is connected to the network and TCP/IP stack supporting sockets is configured on it
- About 500 KB of free disk space for the client program



2.2 Setting up the tool

2.2.1 Server Program

The Server Program can be setup by copying the binary named 'parmond' available in the distribution. This program can be invoked so as to bind to a specified port.

The syntax for executing the server is

```
parmond [ port-id ] [ & ]
```

Where

`port-id` - is an integer number that specifies a valid unused port. In case the port number is not specified, port 4000 is assumed by default.

The optional '&' character forces the server to be run in the background.

Eg: `parmond 45`

In case the port is in use by another application, the server program exits immediately. In such a case, the server program has to be rerun so as to bind to an unused port.

It is not necessary to run several instances of the server program on a particular workstation in a cluster. Multiple clients can access a running server program. However, it is necessary to bind the server to the same port on all the nodes in a cluster if they have to be monitored from a single instance of the client program.

If the server has to be bound to a privileged port (0-1023), then the server program has to be invoked by a privileged user. It is recommended a privileged user invoke the server, thereby granting the server access to kernel tables. The absence of such privileges would render many of the features of the tool unusable.



2.2.2 Setting up the client

Setting up the client consists of the following steps:

1. The client distribution consists of a java archive named 'parmon.jar' and several GIF images. To setup the client, copy the java archive to a directory of choice. The GIF images should then be copied to a directory named 'images' (case sensitive) under the directory where the java archive is stored. For example, if the java archive is stored in a directory named 'parmon', then 'images' should be a sub-directory under 'parmon'. Also create a directory named 'setup_files' under the directory where the java archive is stored. The setup_files directory is used to hold configuration information critical for the functioning of PARMON.
2. Set/update the following environment variables:
 - a. CLASSPATH

Set the value of this variable to point to the directory of the default java class library (classes.lib) and the java archive 'parmon.jar'. The format of definition of the value is as follows:

```
<pathname>[<separator> <pathname>]
```

where

pathname - is a path specification to the directory that java searches to locate the default libraries, user-compiled classes and archives.

Separator - is a character that separates multiple path specifications. Under the UNIX environment, the separator is a colon(:) and under the DOS/Windows95 environment, the separator is a semi-colon(;).



Eg:

Since Java is dependent on the native operating system, please refer to your Operating System documentation for details regarding the setting up of environment variables. However, the following examples are shown to give an idea.

Under the C shell (csh) in UNIX

```
setenv CLASSPATH $CLASSPATH\ : \  
$HOME/parmon/parmon.jar
```

Under the Windows95/DOS Environment

```
SET  
CLASSPATH=C:\JAVA\LIB;C:\PARMON\PARM  
ON.JAR
```

b. PATH

Alter this variable to include the path of the java runtime. The procedure for defining/updating this variable also is Operating System dependent.

Under the C shell (csh) in UNIX

```
setenv PATH $PATH\:/opt/java/bin or  
set path=( $path /opt/java/bin)
```

Under the DOS/Windows95 environment:

```
SET PATH=%PATH%;C:\JAVA\BIN
```

The 'setup_files' directory contains certain configuration (host & group) information. It is recommended that these files be not edited as it may lead to improper functioning of the client. In case the configuration information is corrupt, the files in this directory could be deleted and the tool could easily be configured again.

The client program can then be invoked by executing the following command:

```
java parmon [port-id]
```



where

`port-id` - is the port at which the client program talks with the server program. Note that this port must be the same as that specified for the server.

Tip:

The environment variables can be initialized in the default boot-up scripts or in the login scripts.

For more details regarding the environment variables, please refer to your Operating System's documentation.

Chapter 3

Working with PARMON

The PARMON client program brings up a GUI when invoked. The GUI comprises of menus and icons that may be selected to configure and monitor the cluster environment.

3.1 Components of the GUI

The main screen presented by the tool is similar to that depicted in the figure below (Figure 1)

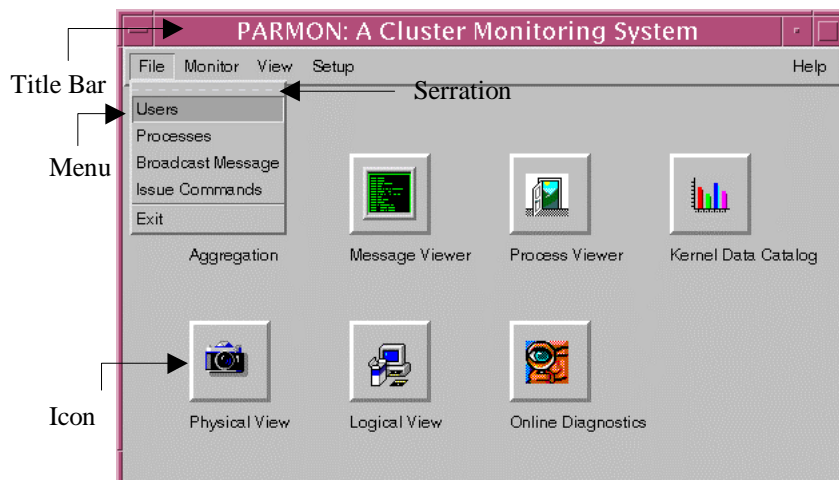


Figure 1 :The Main Screen



3.1.1 The Title Bar

The contents of the title bar indicate the monitoring feature that is currently being accessed.

3.1.2 The Icons

The monitoring functions of the tool have been divided into the following categories:

- Aggregation
- Message Viewer
- Process Viewer
- Kernel Data Catalog
- Logical View

Selecting the corresponding icon on the main screen can access each of these categories of functions.

Apart from the monitoring functions, the tool provides for integration with SunVTS diagnostics tool from Sun Microsystems*.

To invoke the SunVTS tool, simply click on the icon named On-line diagnostics.

3.1.3 Menus

All the functions accessible from icons are also accessible from the menus on the main screen. Note however, that setting up of the tool can be accomplished only from the Setup menu.

An interesting feature of the menu is that it can be detached (torn off) into a separate window. A menu can be torn off by dragging (pull with the left mouse button down) it at the serration. This feature allows for easier and faster access to menus that are frequently selected. Remember that clicking on the tear-off menu option is the same as clicking the respective menu from the main window and then selection the corresponding option. Figure 2 depicts the appearance of the file tear-off menu.



Figure 2 : File Tear-off menu

Each of the other menus too can be torn-off in a similar manner. Figure 3 depicts the appearance of the setup, view and monitor tear-off menus.



Figure 3 : Main Screen with Tear-off Menus

3.2 Setting up PARMON

The GUI has been designed to aid comprehensive monitoring of the nodes of a cluster. However, Prior to employing the tool for monitoring purposes, the client program must be configured to work in the target cluster environment. Setting up PARMON involves defining the nodes/groups of systems available in the cluster. Follow the instructions detailed below to configure the client program.



3.3 Configuring the PARMON client

3.3.1 Node

A cluster is defined as a group of independent workstations connected by means of a high-speed network. Each independent workstation is called a node. The client program requires certain information about the nodes to be monitored like its name to ip-address (Internet Protocol address) map, user disk area etc. Note that without this basic information, the tool cannot be used for monitoring.

3.3.1.1 Node Addition

Node information can be defined by the following the steps detailed below:

On the main screen select (click the left mouse button) the Setup menu. From the sub-menu (Figure 4) that is displayed, select Node.

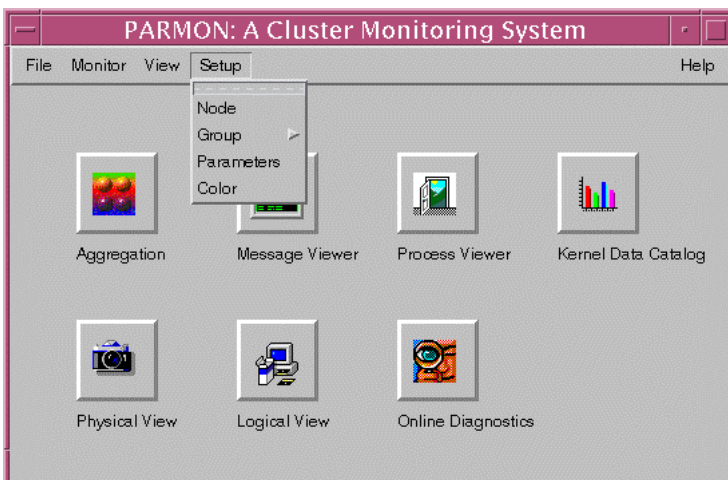


Figure 4 : Setup sub-menu

On selecting Node, the following window (Figure 5) is displayed where information relating to the particular node can be provided. The information required are :

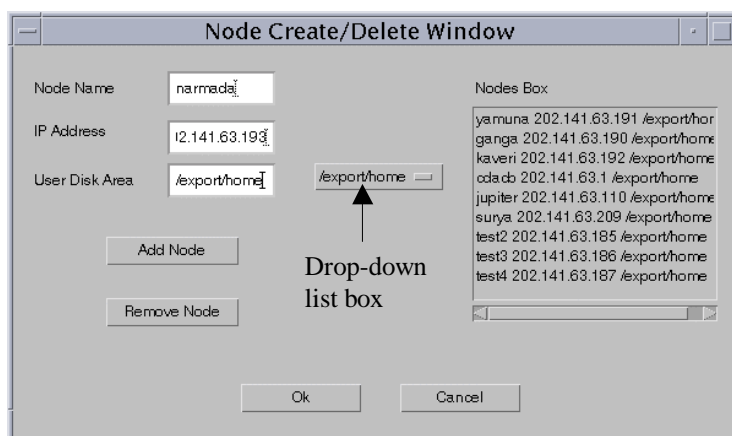


Figure 5 : Node Create/Delete window

Node name – the symbolic name by which a node is referred to. Use of the actual hostname assigned to the node is strongly recommended though not compulsory. Node Name cannot be an empty string (null). Also, node names cannot have embedded spaces.

IP address – the IP (internet protocol) address of the node.

User Disk Area – partition of the disk that needs to be continuously monitored. Refer to the section on aggregation (Refer 3.4.1) for more information. The user disk area can either be keyed in or selected from the dropdown listbox provided.

Select Add Node to add the definition of the node.

The window on the right (Nodes Box) shows the details of the nodes that have already been defined.

3.3.1.2 Node Deletion

It is also possible to delete nodes that have already been defined. To delete a node(s), simply select the node(s) to be deleted in the Nodes Box and select Remove Node. Whenever a node is selected, that node is shaded (blue) in the Nodes Box. Clicking on a selected node deselects it. Note that it is possible to select multiple nodes to delete. The following figure (Figure 6) depicts several selected nodes:

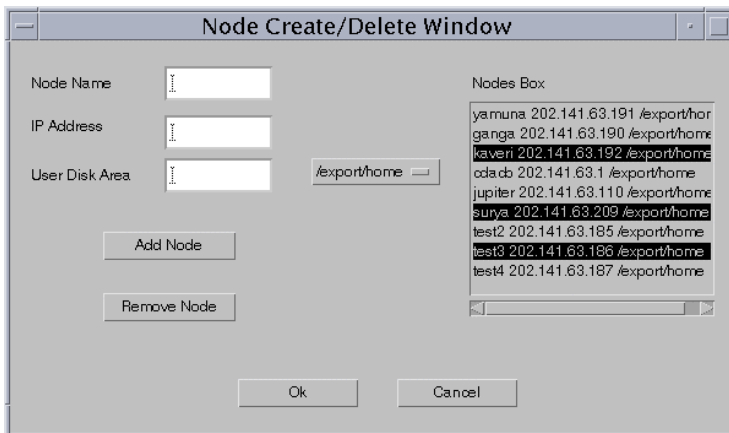


Figure 6 : The Node Create/Delete window (Selected Nodes)

Any number of such nodes can be defined. Note however that each of the nodes should have a unique name and ip-address.

3.3.1.3 Node modification

To change the definition of a node, first delete the node and then create a new one by the same name.

3.3.2 Group

In a cluster environment, there is a frequent need to select certain combinations of nodes for comparing statistics. In order to simplify

the selection of such combinations, the group definition feature has been provided. Using this feature, it is possible to define symbolic group names to combinations of nodes. Consequently, selecting a group will enable the selection of all its constituent nodes of that group. A single node may belong to several/all groups.

3.3.2.1 Group Creation

Selecting the menu sequence Setup → Group → Create (Figure 1) pops up the window depicted below (Figure 7) where groups may be created. Creating groups involves providing a symbolic name to a group and selecting the nodes that ought to be referred by that name. In the screen depicted below simply provide a symbolic name and select the nodes that belong to that group and select the OK button.

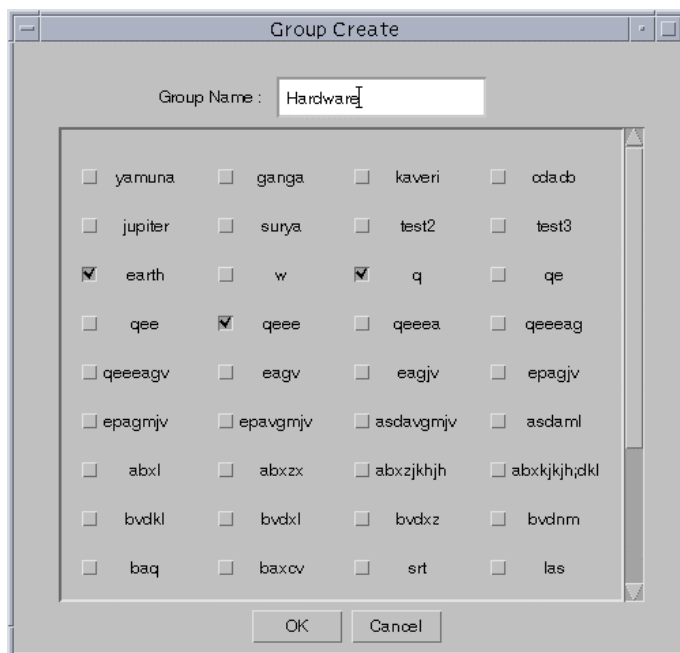


Figure 7 : The Group Create window



Group Name – is the symbolic name assigned to a group of nodes. Similar to node names, group names too cannot be empty strings. Also, group names cannot have embedded spaces and, two group cannot have identical names. If a group name is null or has embedded spaces, an error message is displayed. The error message is similar to that depicted in the following picture.

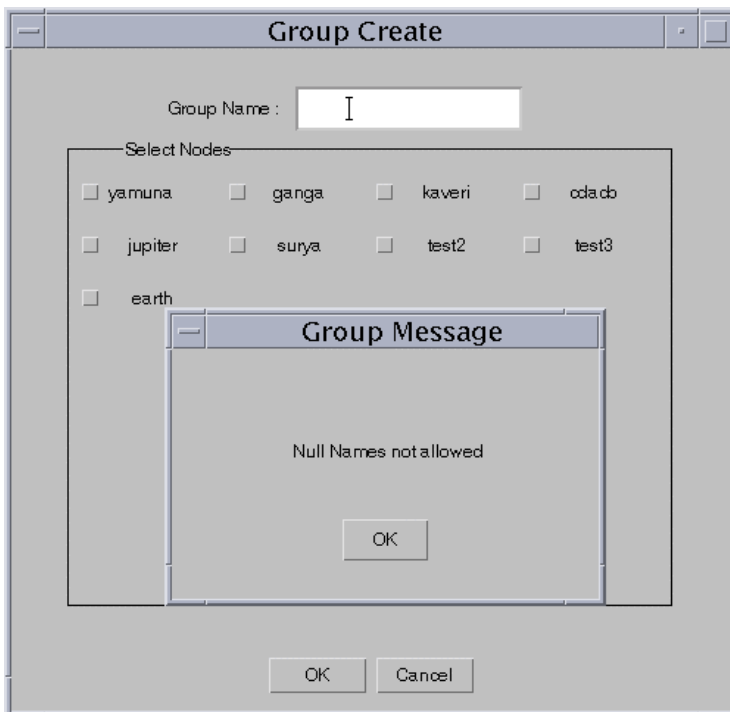


Figure 8 : Group Name error message (Null Names)

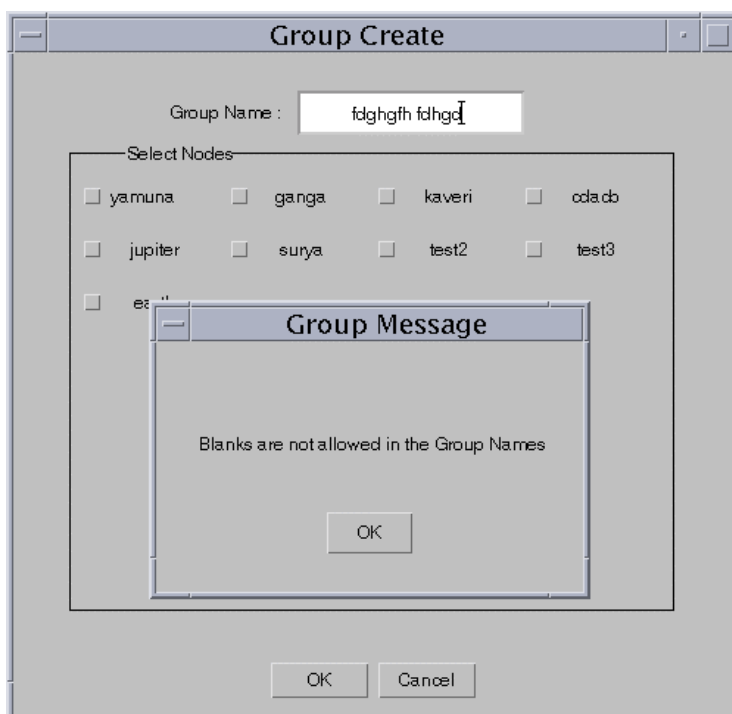


Figure 9 : Group Name error message (embedded blanks)

To select the nodes, select the checkbox (toggle) towards the left of the node. A selected node is indicated by a tick mark in the checkbox. To deselect a node, simply click on a selected one.

Note:

There are two predefined groups namely 'Entire System' and 'Individual Nodes'. These groups may be used to select all the nodes that have been defined or just a single node respectively.

3.3.2.2 Group Modification

To view and/or modify the definition of a group, select the menu sequence Setup → Group → View/Modify. The figure (Figure 10)



that follows depicts the appearance of the Group View/Modify window.

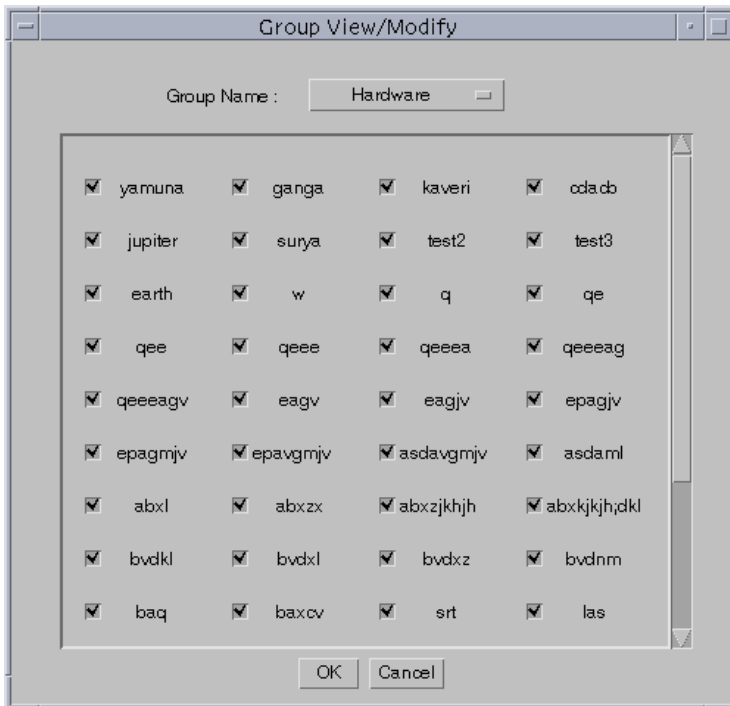


Figure 10 : Group View/Modify window

3.3.2.3 Group Deletion

Selecting the menu sequence Setup → Group → Remove pops up the Group Remove window (Figure 11). The Groups window on the left displays the list of defined groups. Select the group definition(s) that needs to be deleted/removed from the list. The selected group definitions are shaded(blue). Selecting the Remove button will delete the definition. Clicking on a selected group deselects it. Selecting the Remove button deletes the selected groups. Note that multiple groups can be selected and deleted at once.

To view the definition of a group, select the required group in the Group Name list box. The Selected Nodes window updates with the nodes belonging to that particular group.

To modify the definition of a group, select the required group and select/deselect the nodes in the Selected Nodes window.

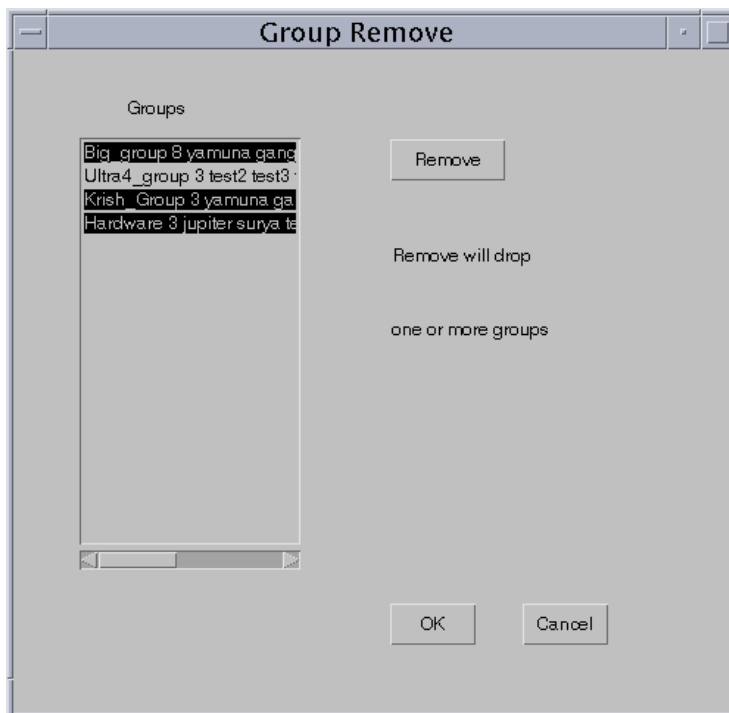


Figure 11 : The Group Remove window

3.3.3 Specifying Parameters

At present the tool supports only a single parameter, i.e., the Chart Type parameter. Selecting the menu sequence Setup → Parameters, pops up the Chart type window (Figure 12) where the preferred char type may be selected. The preferred type



may be either Pie Chart or Bar Chart may be selected. To select a type, click on the radio button beside the preferred chart type. The preferred chart type is indicated by a filled (black) radio button.

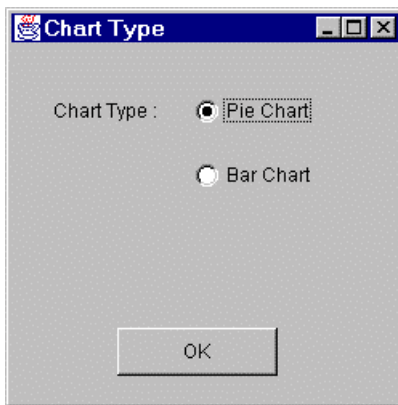


Figure 12 : The Chart Type window

3.3.4 Changing color and other preferences.

##

3.4 Monitoring the cluster

Several important statistics of the nodes in a cluster could be monitored using the tool. The statistics may be monitored under the following categories:

- Aggregation
- Messages
- Processes
- Logical View
- Kernel Data Catalog

3.4.1 Aggregation

This set of monitoring functions can be used to continuously monitor the usage of resources available on the node namely the CPU, Memory, Disk and Network. This statistic can be viewed either in the form of a PIE chart or a BAR chart. During a monitoring session it is also possible to change the type of chart that is projected. To opt for Aggregation statistics, select the icon named Aggregation in the main screen. A pop up menu similar to the one depicted in the figure (Figure 13) below is displayed.

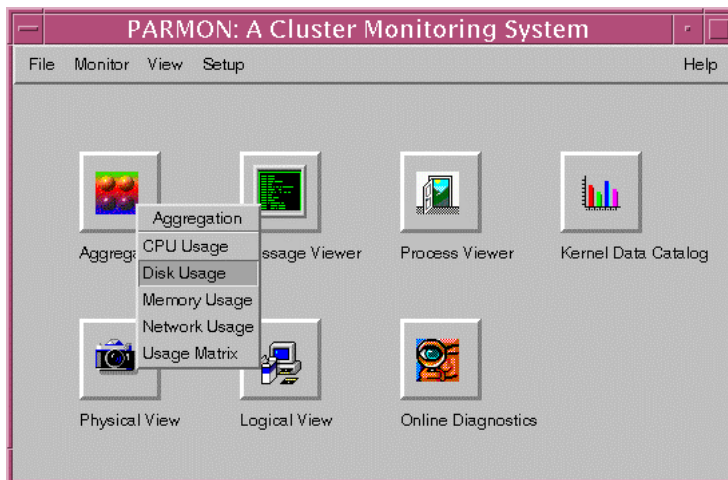


Figure 13 : Popup menu under Aggregation

| To view | Select on popup menu |
|--------------------|----------------------|
| CPU Statistics | CPU Usage |
| Disk Statistics | Disk Usage |
| Memory Statistics | Memory Usage |
| Network Statistics | Network Usage |
| Usage Matrix | Usage Matrix |

Selecting one of the pop-up menu options displays a node selection window similar to that shown in Figure 14.



Node Selection Window

The Node Selection Window allows the selection of one or more nodes/groups from those already defined. A group can be selected from the drop down list box. Selecting a group selects all the nodes that belong to that group. Further, it is possible to deselect nodes. Follow the procedure defined in Creating Groups (refer 3.3.2.1) to select or deselect nodes. Selecting the Individual Nodes group lists all the defined nodes, of which at most one node could be selected. Deselection is effective only so long as another group is not selected. Also, the group/node(s) on which resource usage needs to be monitored should have the server running and bound to a valid port.

Select the group that needs to be monitored. In case a group or multiple nodes are selected, the statistics of all the selected nodes are projected within a single window.

Selecting the OK button displays the appropriate aggregation (CPU, Disk, Network or Memory usage) window.

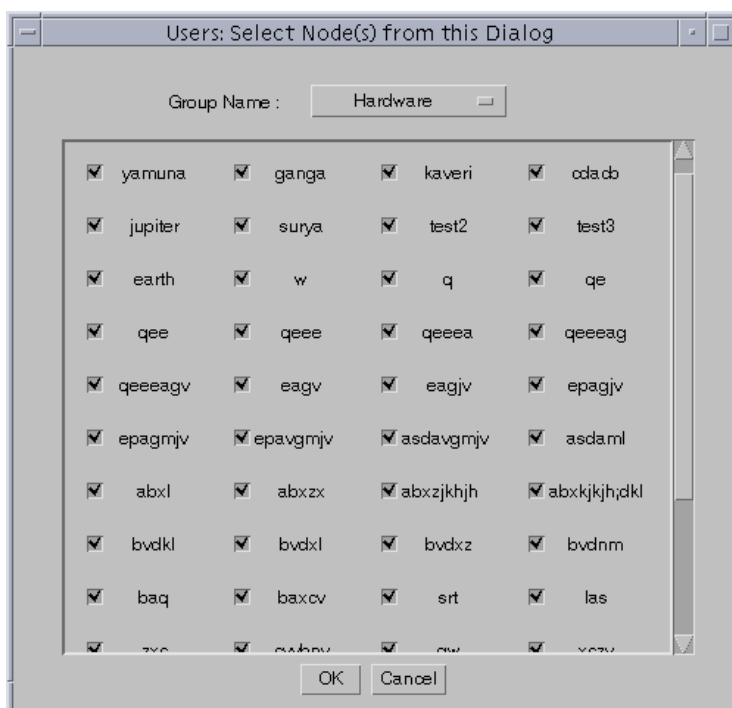


Figure 14 : The Node/Group Selection Window

Note:

In case the client program fails to communicate with PARMON server on the specified port on a given node (workstation), the client program displays the following message window.



Figure 15 : Server Error Message window

3.4.1.1 CPU Usage

The CPU Usage Statistic presents the usage in percentage of the CPU(s) on the specific node/group. This statistic is presented in the form of a chart (either PIE or BAR) and is refreshed periodically. The refresh period could be one of standard 5, 10, 15 and 20 seconds or user-defined. The CPU Usage chart is similar to the ones (Figure 16, Figure 17) shown below.

The Bar (Figure 16) and Pie (Figure 17) charts depict the usage graphically with the lower and higher extremes being 0 and 100% respectively. The view can be switched from Pie to Bar and back dynamically. To do so, select the menu option Chart and opt for the desired chart type. Refer also section 3.3.3 for defining the default chart type.

The refresh time interval defines the mean time between successive updates - lower the refresh time, faster the update. The refresh time specified is in seconds. The refresh interval can be redefined by selecting the TimeSet menu option and choosing the required refresh interval.

If desired, the refresh cycle can be suspended and resumed. The Stop button can be selected if one wishes to suspend the refresh. When the refresh is suspended, the Stop button is replaced by the Start button.

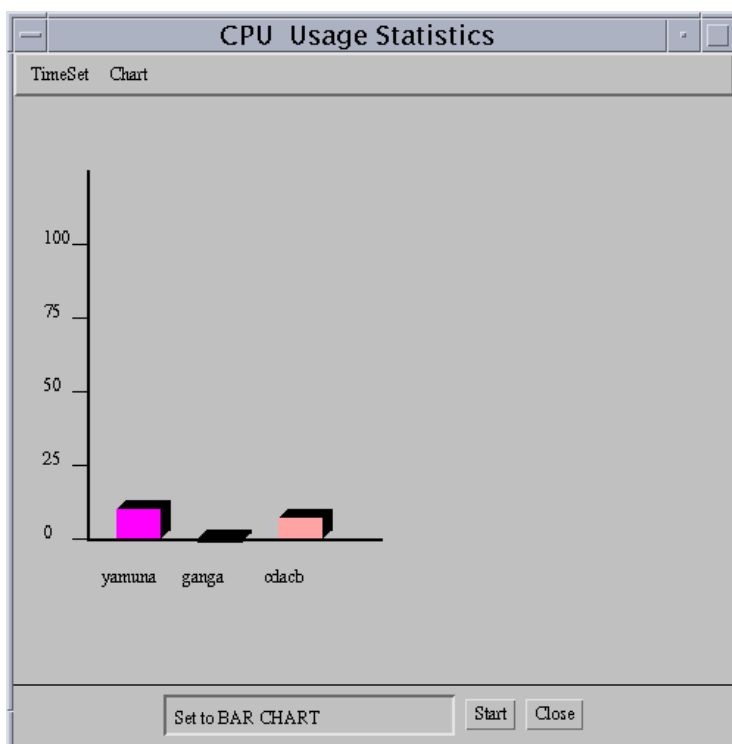


Figure 16 : CPU Usage statistics as a Bar Chart

The dialog window at the bottom of the window displays the status of the latest change (refresh interval, stop, start etc) affected on the window.




Figure 17 : CPU Usage as a Pie Chart

Figure 18 shows the Timeset sub-menu and Figure 21 depicts the Chart sub-menu. From the sub-menu it is possible to opt for a standard refresh time period or specify a user-defined time interval.




Figure 18 : CPU usage chart depicting Timeset sub-menu

Selecting the key-in sub-menu pops up a dialog window where the required refresh period can be specified. The dialog window is similar to that depicted below (Figure 19).

The time input should be an integer value. In case there is an error in input, the message window similar to that depicted in Figure 20 is displayed.

Note:

Changes to refresh time take effect only after the following refresh,i.e, the very next refresh after the time is redefined occurs at the period active currently.

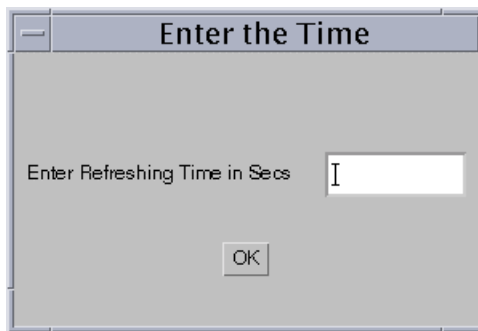


Figure 19 : Key-in dialog window



Figure 20 : Invalid input error message window

The chart sub-menu allows the selection of the preferred type of chart. Figure 21 depicts the appearance of the Chart sub-menu.




Figure 21 : CPU usage chart depicting Chart sub-menu

Selecting the Close button closes the window.

Note:

In case a node has multiple CPUs, the average usage of all the CPUs will be computed and displayed.

3.4.1.2 Disk Usage

The Disk Usage Statistic presents the usage in percentage of the User Disk Area (refer section 3.3.1.1) on the specific node/group. This statistic is presented in the form of a chart (either PIE or BAR) and is refreshed periodically. The refresh period could be one of

standard 5, 10, 15 and 20 seconds or user-defined. The Disk Usage chart will be similar to the ones (Figure 22, Figure 23) shown below.

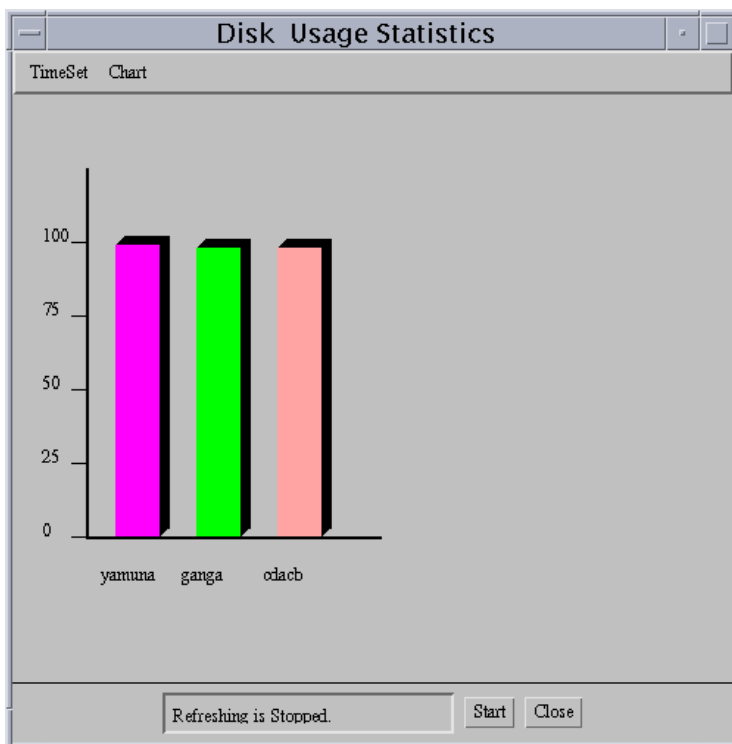


Figure 22 : Disk Usage Statistics as a Bar Chart

The meaning, behaviour and functionality of the components of the window are similar to those described in the section on CPU Usage (refer section 3.4.1.1).




Figure 23 : Disk Usage Statistics as a Pie Chart

3.4.1.3 Memory Usage

The Memory Usage Statistic presents the usage in percentage of the system memory (RAM) on the specific node/group. This statistic is presented in the form of a chart (either PIE or BAR) and is refreshed periodically. The refresh period could be one of 5, 10, 15 and 20 seconds. The Memory Usage chart will be similar to the ones (Figure 24, Figure 25) shown below.




Figure 24 : Memory Usage Statistics as a Bar Chart

The meaning, behaviour and functionality of the components of the window are similar to those described in the section on CPU Usage (refer section 3.4.1.1).




Figure 25 : Memory Usage Statistics as a Pie Chart

3.4.1.4 Network Usage

The Network Usage Statistic presents the ratio between the total number of packets transmitted successfully and the number of packets in error for the default network interface on the specified node/group. This statistic is presented in the form of a chart (either PIE or BAR) and is refreshed periodically. The refresh period could be one of 5, 10, 15 and 20 seconds. The Network Usage chart will be similar to the ones (Figure 26, Figure 27) shown below.

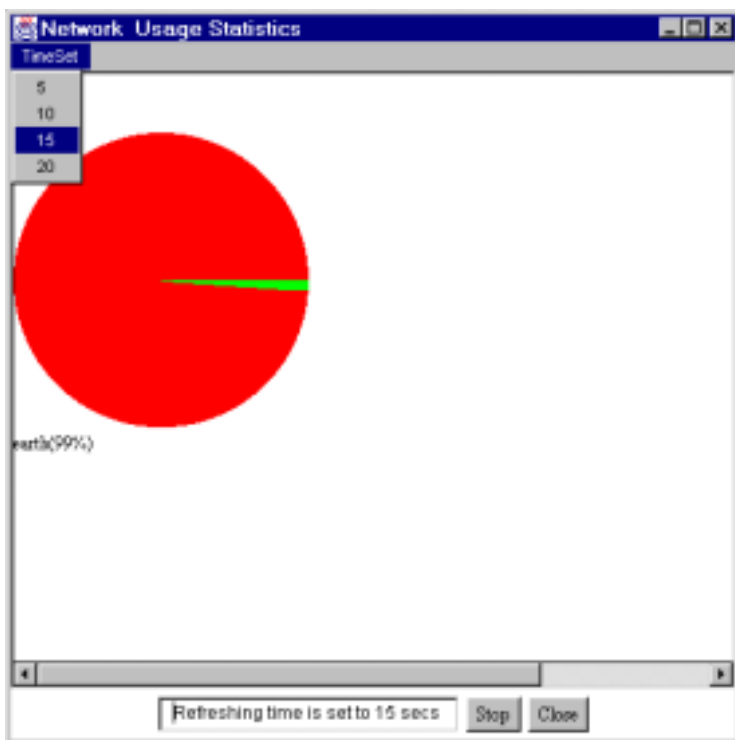


Figure 26 : Network Usage Statistics as a Pie Chart

The meaning and behaviour of the components of the window is similar to those described in the section on CPU Usage (refer section 3.4.1.1).

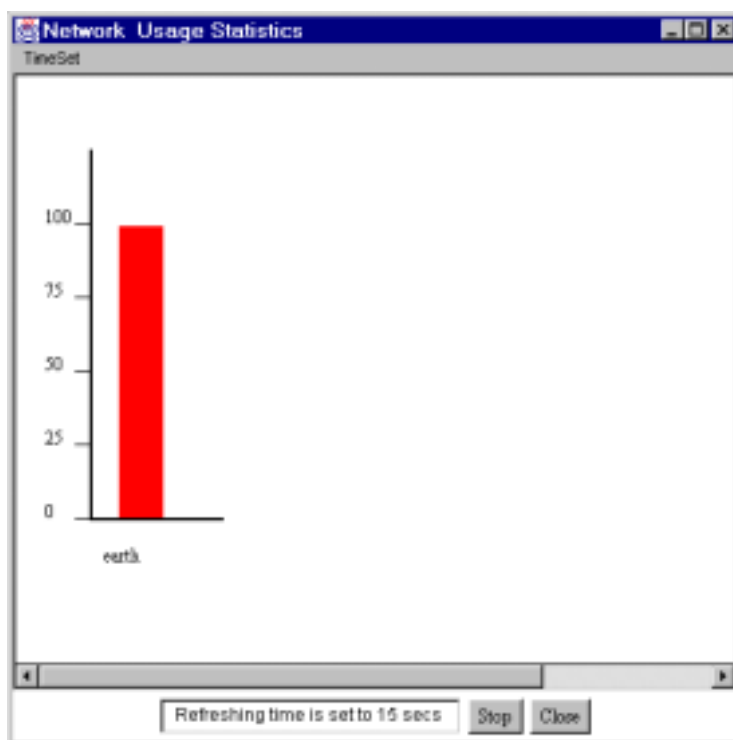


Figure 27 : Network Usage Statistics as a Bar Chart

3.4.1.5 Usage Matrix

System managers of large cluster environments prefer to observe statistics of the entire system at a glance and then, proceed to enquire about the statistics of groups, nodes etc. Usage matrix provides exactly this functionality.

PARMON is capable of depicting the usage of a cluster system in three ways:

- a. Group-wise,
- b. Node-wise within a group and,



c. Processor-wise within a node.

The group-wise usage statistics window is similar to that depicted in Figure 28.

| Group Name | Num Of Nodes | CPU Usage | Disk Usage | Mem Usage |
|------------------|--------------|-----------|------------|-----------|
| Krish_Group | 4 | 14% | 98% | 84% |
| Hardware | 3 | 12% | 98% | 84% |
| HPCC_Group | 4 | 12% | 98% | 84% |
| SSDG_Group | 5 | 11% | 98% | 84% |
| Networking_Group | 3 | 12% | 98% | 84% |
| RealTimeSystem | 4 | 12% | 98% | 84% |

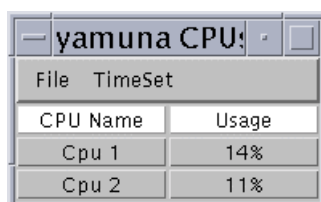
Figure 28 : Group-wise Usage Matrix

In the figure depicted above, selecting a group (shaded rose) displays another matrix where the usages of the individual nodes in the selected group are displayed (Figure 29).

| Node Name | CPU Usage | Disk Usage | Mem Usage |
|-----------|---------------|---------------|---------------|
| yamuna | 20% | 99% | 92% |
| ganga | 0% | 98% | 65% |
| cdacb | 8% | 98% | 95% |
| surya | Server Not Up | Server Not Up | Server Not Up |

Figure 29 : Node-wise usage statistics of a group

To view the usage statistics of the processors within a node, select the CPU usage statistic (shaded rose) displayed adjacent to the required node. The CPU usage statistics of the processors in a node appear similar to that depicted in Figure 30.



| yamuna CPU: | |
|-------------|---------|
| File | TimeSet |
| CPU Name | Usage |
| Cpu 1 | 14% |
| Cpu 2 | 11% |

Figure 30 : Processor-wise usage statistics of a node

3.4.2 Messages

Messages critical to a system are usually displayed on the console. These messages indicate the health of a system. The messages displayed on the console are logged by the operating system to aid in diagnosing problems with the system. System Managers frequently refer to these logs to diagnose the overall health of a system and take corrective action, whenever necessary.

In large cluster systems it is an inhuman task to constantly keep a tab on the messages generated by the system. The option to scan through the logs of an individual node(s) in a cluster system can go a long way in easing the managerial task of such managers.

Selecting the icon Message Viewer on the main screen pops up the Node selection window.

To view the messages from a node or group, select the appropriate group or nodes. Note that even after selecting a group, it is possible to select or deselect members of that group. Follow the procedure defined in Creating Groups (refer section 3.3.2.1) to select or deselect nodes. Selecting the Individual Nodes group lists all the defined nodes, of which at most one node could be selected. Selecting the OK button displays a window similar to that depicted below:

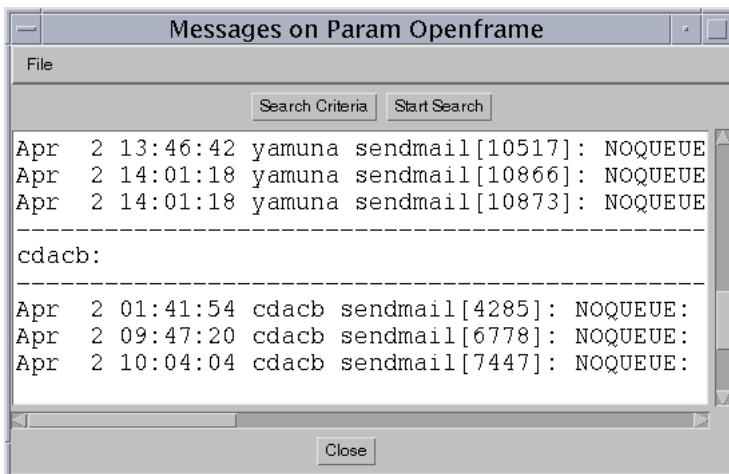


Figure 31 : The Message window

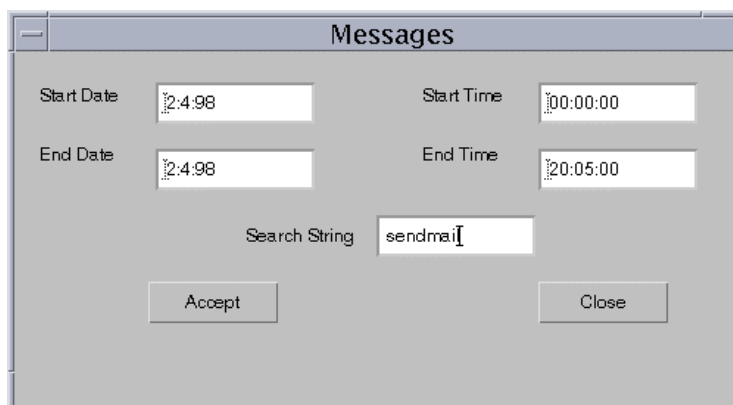
The Messages window displays the messages logged in the default messages file (/var/adm/messages).

By default, only messages logged since 0:00 hours on the current day are displayed. In case other messages are required, a search criterion has to be defined. To define search criterion follow the procedure detailed below:

1. Select the button Search Criteria. A window similar to that depicted in Figure 32 is shown.
2. Enter the start date, end date, start time and the end time. If required, a search string can be specified.
3. To accept the search criterion select the Accept button. Selecting the Close button causes the search criterion to revert to default.
4. Select the Start Search button.

Messages logged since the start time on the start date till the stop time on the stop date are displayed. The search can be stopped in the middle by selecting the Stop Search button. In case a search string is specified, only those messages that have been logged

during the specified period (date and time) and having the specified search string are displayed.



The image shows a dialog box titled "Messages". It contains four input fields for date and time: "Start Date" (2:4:98), "Start Time" (00:00:00), "End Date" (2:4:98), and "End Time" (20:05:00). Below these is a "Search String" field containing "sendmail". At the bottom are "Accept" and "Close" buttons.

Figure 32 : Criteria Dialog box

The messages shown on screen may be saved on to a file on the local disk. Follow the procedure listed below to save the messages to a local file.

1. Select the menu sequence File → Save As. A window similar to that in Figure 33* is displayed.




Figure 33 : The SaveAs Window

2. Specify the required filename and select the Save* button.
- * Since java uses the native windowing system, the appearance/functionality of the window may be entirely different on different operating systems.

3.4.3 Processes

The processes running on a specified group/node(s) can be viewed using the Process Viewer. To view the processes running on a specific group/node(s), follow the procedure listed below:

1. Select the icon Message Viewer. The Node Selection window similar to that depicted in Figure 14 is displayed.



2. Select the required group/node(s) and select the OK button. A window similar to that shown in Figure 34 is displayed.

The processes are by default ordered on the process names. It is also possible to order the processes on process number, owner's username, owner's user ID or size, by selecting the Sort menu option and one of By Number, By Username, By UID or By Size respectively from the pop-up menu.

It is possible to specify a refresh interval for updating the processes running on the specified group/node(s). Specifying the refresh interval is similar to that defined in the section on CPU Usage. It is also possible to suspend and restart the refresh cycle.

The list of processes displayed on the window can be saved on to a file on the local disk by selecting the menu sequence File → Save As.

Apart from the Save As and the Timeset menus whose behaviour is the same as that explained in the previous section (refer section 3.4.1.1), the processes window provides for an additional Sort menu. The Sort menu allows for sorting processes based on it names, ID, owner name, owner ID and size (memory image size).

Note:

It is recommended that the refresh be stopped before ordering the processes because refresh will cause the ordering to be reset to default.

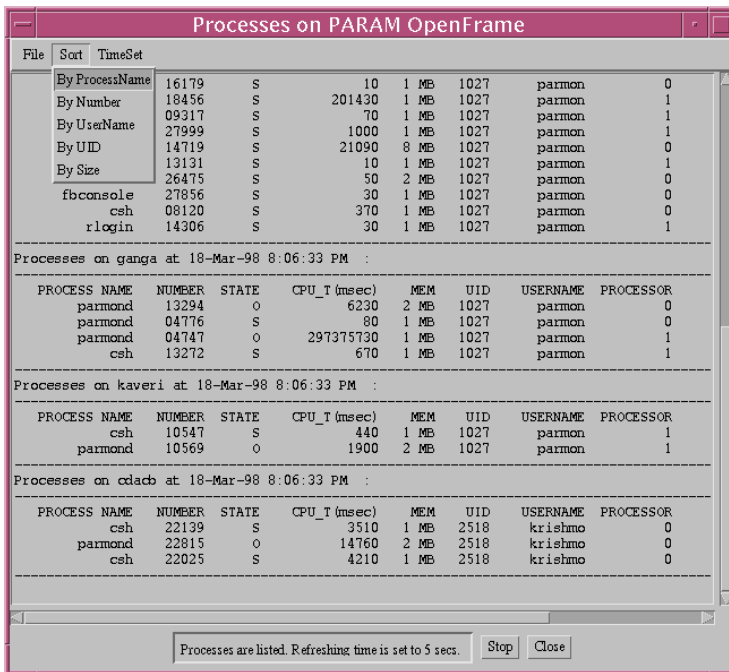


Figure 34 : Process Window depicting the Sort sub-menu

3.4.4 Logical View

The logical view of a system consists of a list of resources available on the system. The resources have been classified under the following four heads:

- CPU
- Disk
- Memory
- Network

To obtain the logical view of a specific group/node(s), select the icon named Logical View on the main screen. From the Node Selection window select a group or nodes of which the logical

view is required and select the OK button. A window similar to that shown in Figure 36 is displayed.

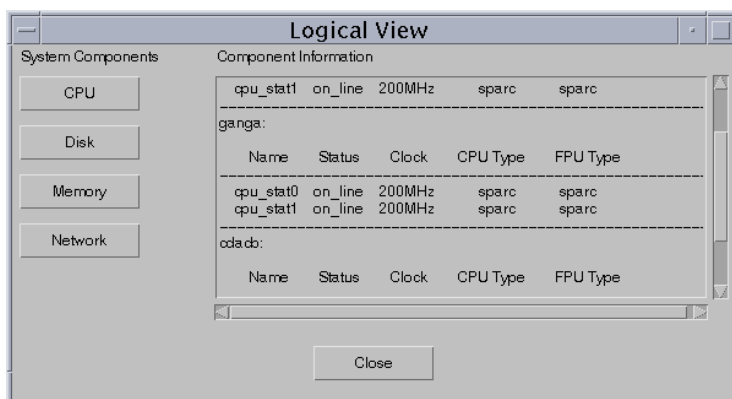


Figure 35 : The Logical View window

The following table lists the actions and the corresponding output.

| Selecting... | Output |
|--------------|---|
| CPU button | Number of processors Name assigned to each processor Clock frequency Type of CPU for each processor Type of FPU of each processor |
| Disk button | File system information like <ul style="list-style-type: none"> • Mount point • Preferred Block size • Fundamental Block Size • Total number of blocks in the partition • Total number of free blocks • Free blocks available to an user • Free blocks available to a Superuser • Maximum length of a file name • Maximum number of file nodes • Type of file system • The File system ID assigned to that file system |



| Selecting... | Output |
|----------------|---|
| Memory button | Maximum available memory Available Free memory Reserved swap size Free swap size Swap size available for swapping purposes Available resident memory Physical memory size Physical page size |
| Network button | Number of network interfaces and their names |

Figure 37 to Figure 39 list the sample outputs of the logical view of a node.

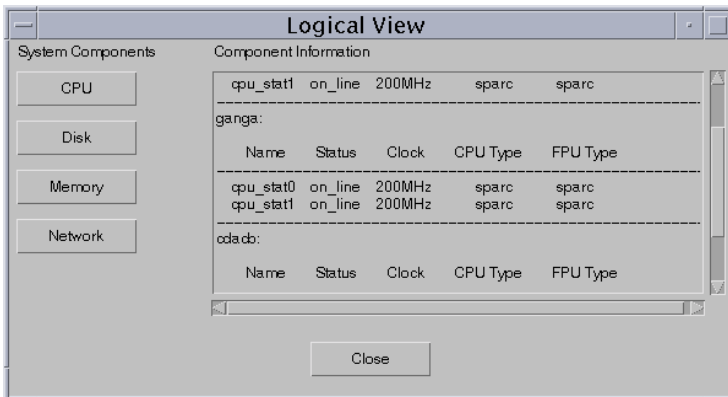
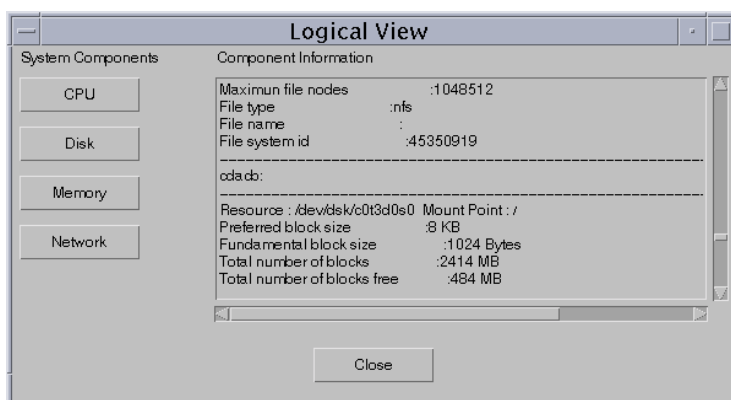
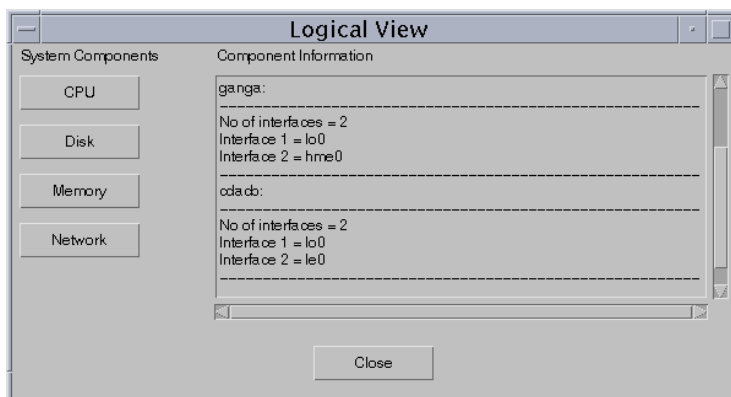


Figure 36 : Logical View - CPU

**Figure 37 : Logical View - Disk****Figure 38 : Logical View - Network**

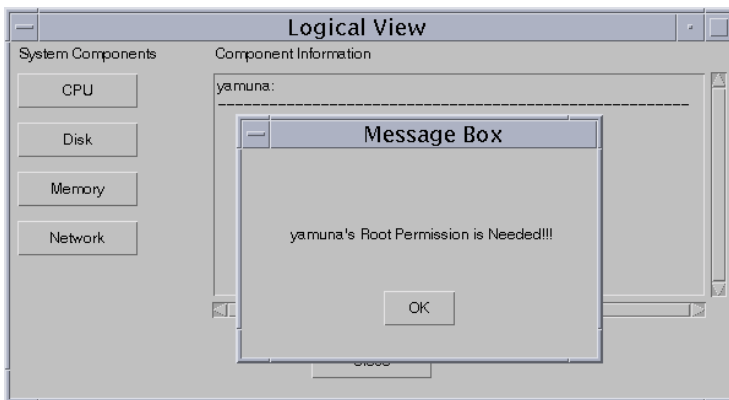


Figure 39 : Logical View – Memory (Error Message)

Note:

The PARMON server must possess superuser rights to display the Logical View of memory. The absence of superuser rights causes an error message window similar to that depicted in figure ## to be popped up.

3.4.5 Kernel Data Catalog

Kernel Data catalog statistics present managers with a catalog of the resources (CPU, Disk, Memory and Network) available on the nodes being monitored. To access these statistics select the icon named Kernel Data Catalog from the main screen. Figure 40 to Figure 44 depict the appearance of the windows for monitoring the resource utilization of different resources namely the CPU, Disk, Network and Memory.

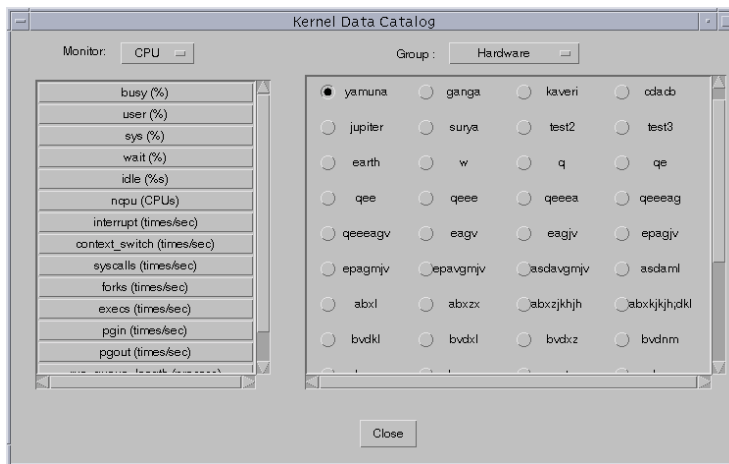


Figure 40 : kernel Data Catalog – CPU

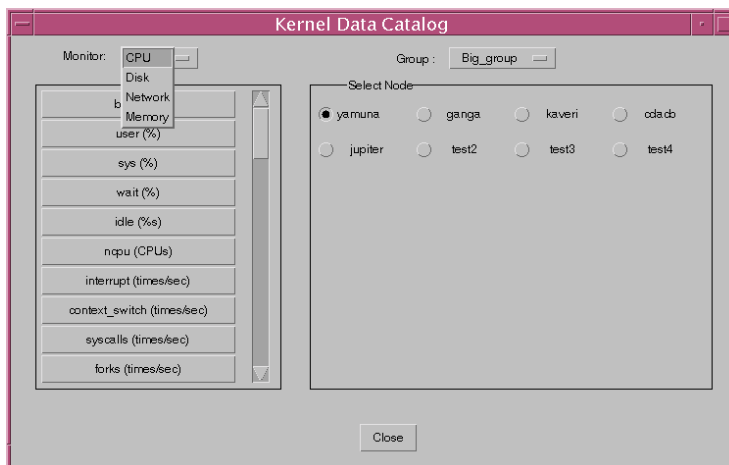


Figure 41 : Kernel Data Catalog with drop down list

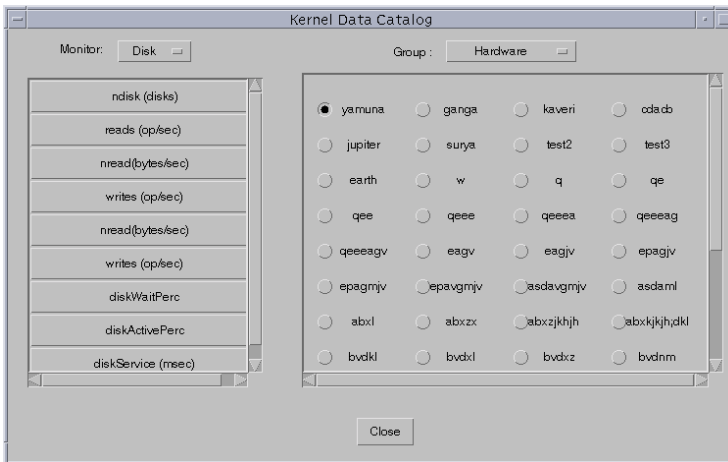


Figure 42 : Kernel Data Catalog - Disk

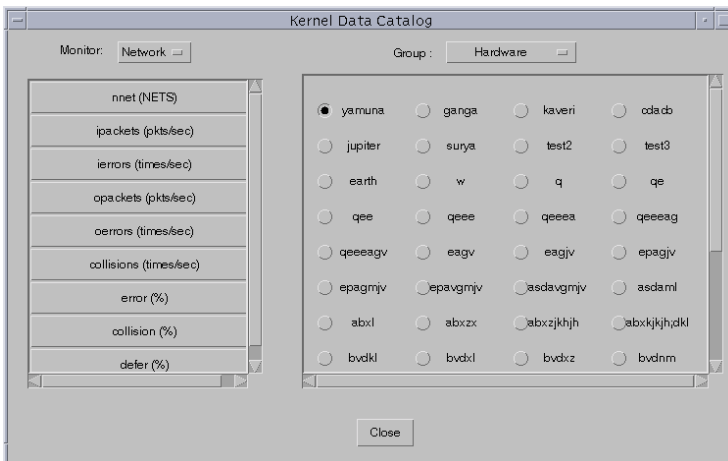


Figure 43 : Kernel Data Catalog - Network

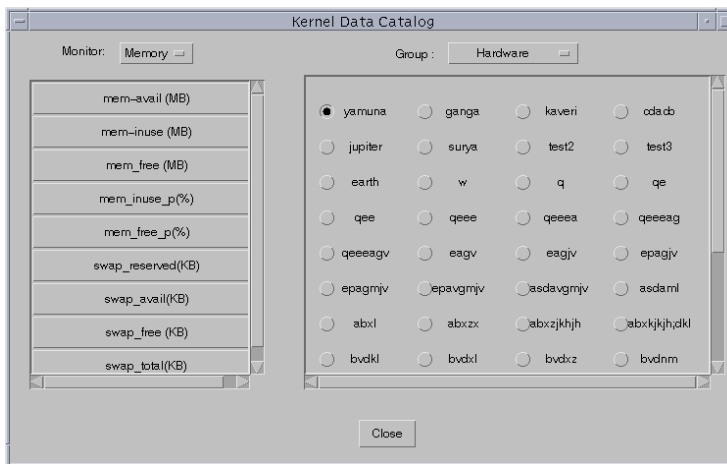


Figure 44 : Kernel Data Catalog - Memory

The right half of the window lists the group and the nodes within that group. Unlike other statistics, cataloging does not support multiple nodes in a single window. In case multiple nodes have to be cataloged, each node should be separately selected and the desired parameter chosen. In such cases, the statistics too are projected on separate windows. The top left corner of the

3.4.5.1 CPU Catalog

The CPU catalog option catalogs the utilization statistics of the CPU over a period of 5 minutes with refresh rates of 2, 3, 4 or 5 seconds. Figure 40 depicts a sample CPU catalog screen.

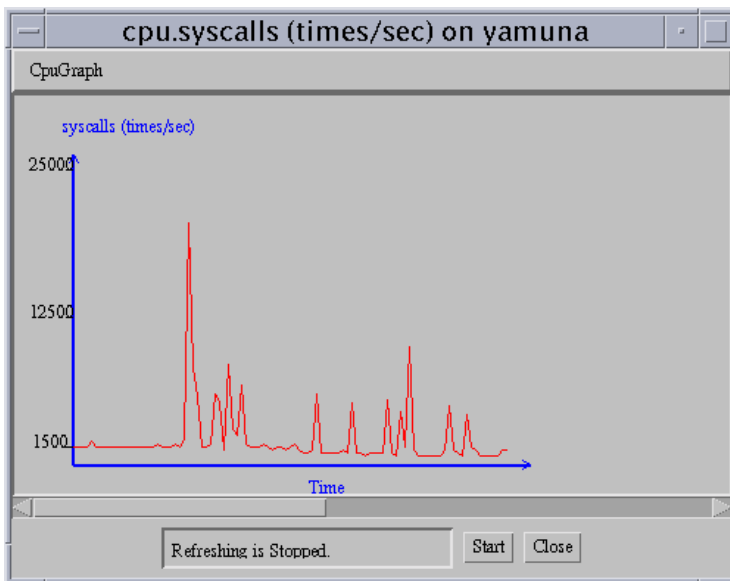


Figure 45 : CPU Utilization Curve - Syscalls

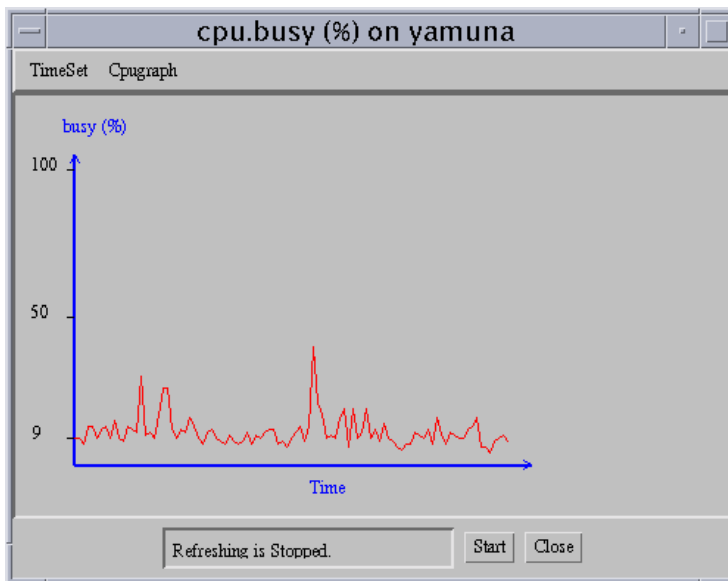


Figure 46 : CPU utilization curve

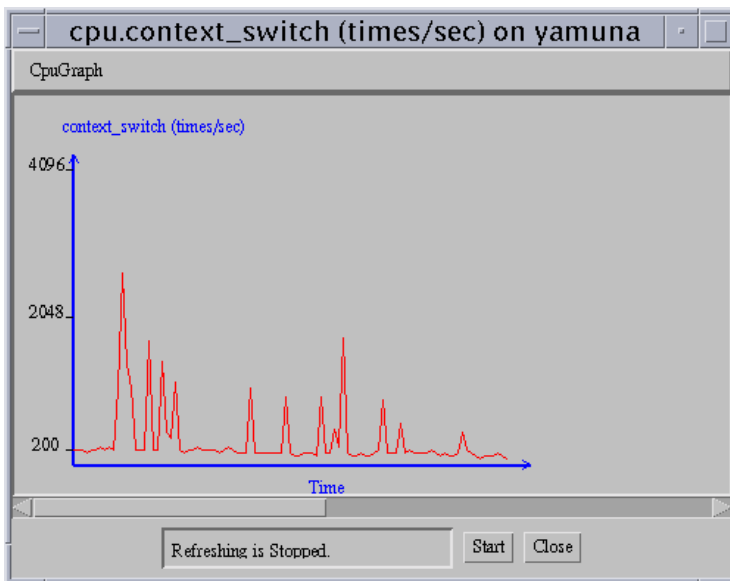


Figure 47 : CPU Utilization – Context Switch

The functioning of the Stop button is similar to that described in the section on CPU Usage. The window can be closed by selecting the Close button.

3.5 Miscellaneous Features

Apart from the monitoring functions, the PARMON tool can provide certain information like a node's hardware and software configuration. The tool also allows managers/users to execute general UNIX commands on the nodes.

3.5.1.1 The View Menu

The View menu has options that enable a user/manager to view the System Configuration and System Information. The view menu has several sub-menu options as depicted in Figure 48.



Selecting the menu sequence View → SysConf and subsequently selecting a group/node(s) displays the system configuration of the selected nodes. Figure 49 depicts the sample output of the SysConf option.

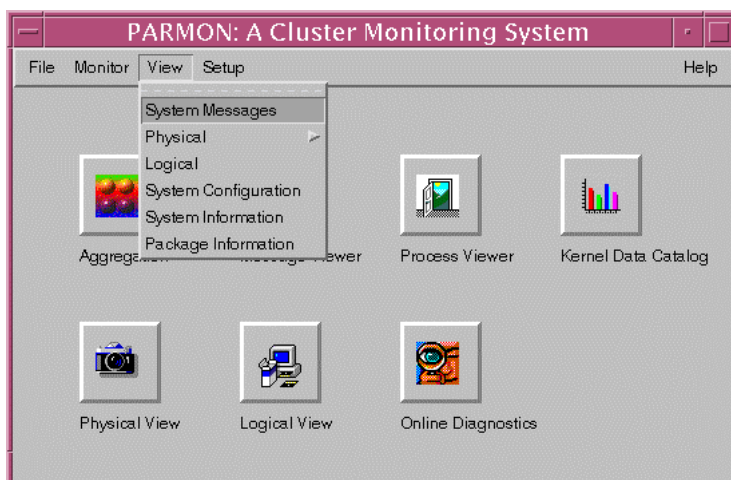


Figure 48 : View sub-menu

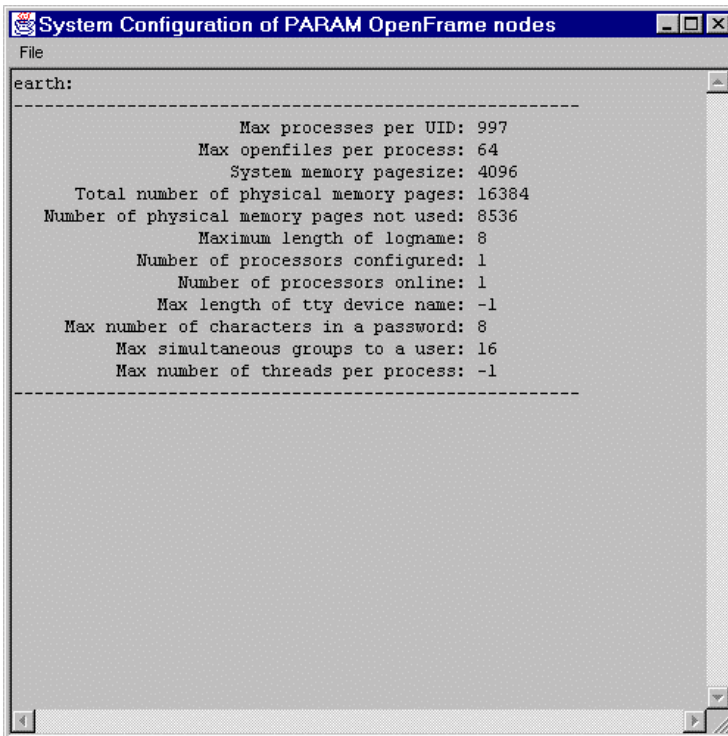


Figure 49 : System Configuration window

To view the System information, select the menu sequence View → SysInfo and subsequently choosing a group/node(s) displays the System Information. Figure 50 depicts a sample System Information window.

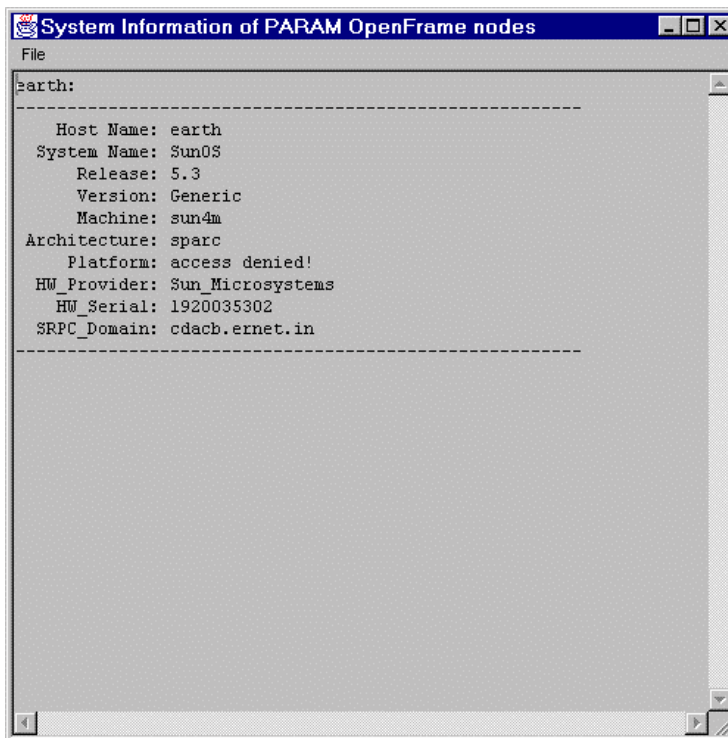


Figure 50 : System Information window

3.5.2 Executing UNIX commands

To execute standard UNIX commands on a group/node(s), follow the procedure described below:

1. Select the menu sequence File → Issue Commands. A node selection window pops up.
2. Select the group/node(s) on which the command is to be executed. Select the OK button. A window similar to that depicted in Figure 51 is displayed.



3. Key in the command to be executed and press the OK button. To aid administrators/managers, several-privileged UNIX commands are listed in the listbox, one of which can be selected.




Figure 51 : Issue Commands window

In case a privileged command is issued, a window similar to that depicted in Figure 52 is displayed. To successfully execute such commands, enter the root password and select the OK button.



Figure 52 : Password window

If the root password entered is not valid, an error message is displayed within the same window and re-entry of password is allowed. The error message displayed is similar to that shown in Figure 53.



Figure 53 : Password Error message box

Note:

In case a privileged command needs to be executed on multiple nodes, the tool prompts for the root passwords of each of the nodes separately.

3.5.3 Broadcasting Messages

To broadcast messages, select the procedure described below:

1. Select the menu sequence File → Broadcast Message.
2. From the Node selection window select the group/node(s) where the message should be broadcast. A window similar to that depicted in Figure 54 is displayed.
3. Enter the message and select the Send button.



Figure 54 : Broadcast message window

Select the Cancel button to close the window.

3.5.4 Online Help

Online help gives a detailed description about using the tool. Online help can be access under the Help menu on the main screen. Figure 55and Figure 56 depict sample Help windows.

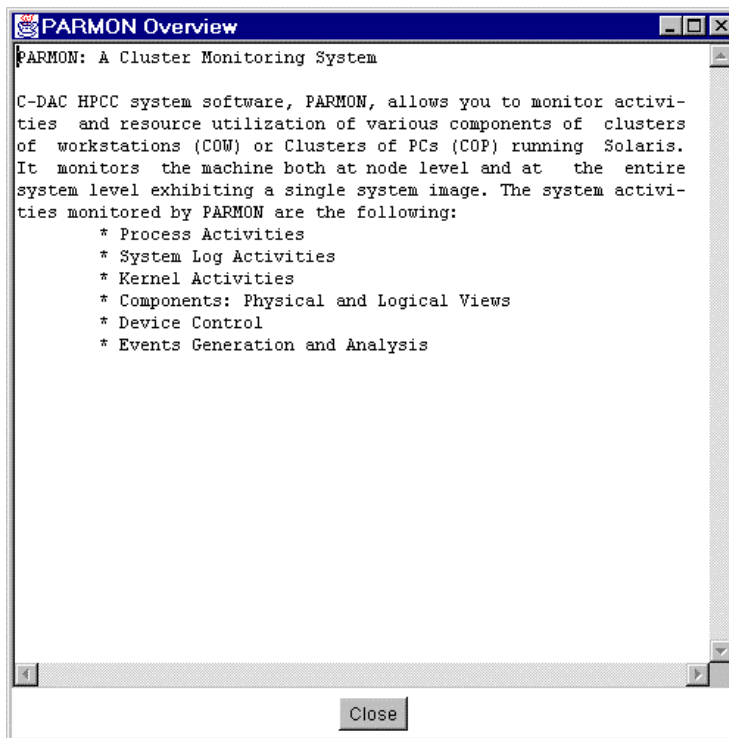


Figure 55 : Help - Overview

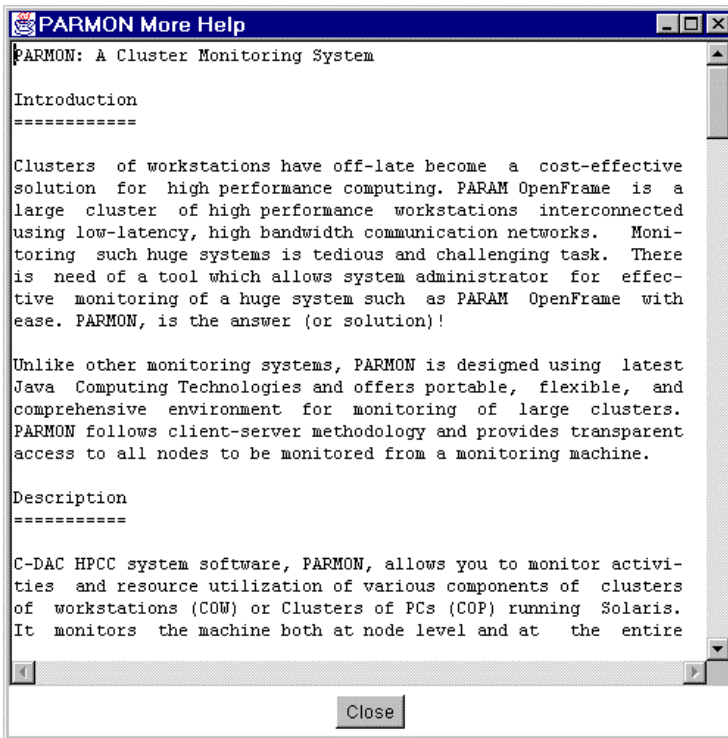


Figure 56 : Help - More Help

Figure 57 gives copyright information.

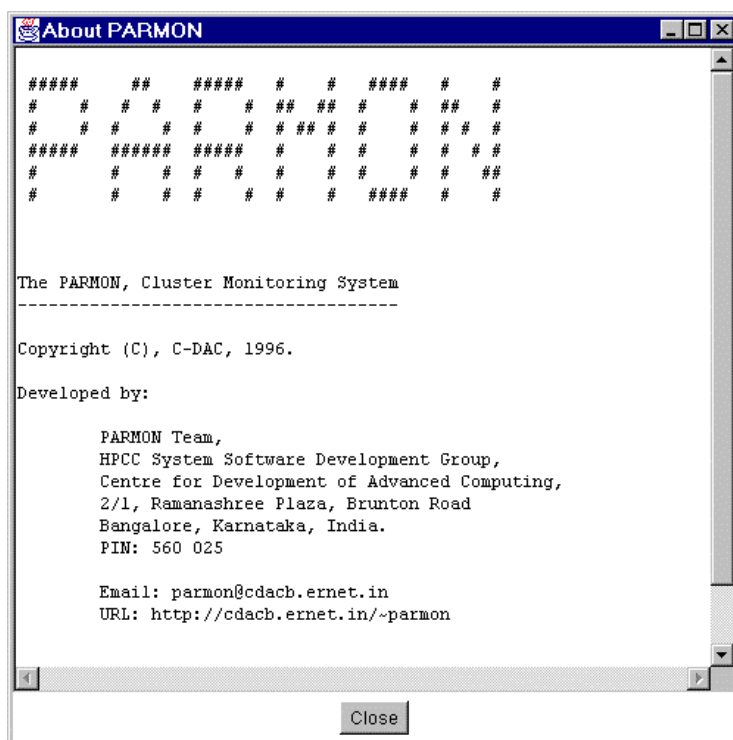


Figure 57 : Help - About