



Enlighten**DSM**/Workgroup

Enlighten**DSM**/Advanced

Reference Manual

Release 3B

Enlighten Software Solutions, Inc.

Corporate Headquarters

999 Baker Way, Fifth Floor
San Mateo, CA 94404 U.S.A.

Telephone: (650) 578-0700

Facsimile: (650) 578-0118

E-mail: support@sftw.com
<http://www.EnlightenDSM.com>

Copyright

© 1994, 1995, 1996, 1997, 1998, 1999 by Enlighten Software Solutions, Inc. All rights reserved.

This manual is designed as a guide to using various features and functions of the Enlighten**DSM** software package to monitor and administrate UNIX and Windows-based computer systems.

All information contained in this manual is proprietary to Enlighten Software Solutions.

Enlighten Software Solutions makes no representation or warranties of any kind whatsoever with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Enlighten Software Solutions shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this publication or its contents.

Trademarks or Service Marks

Enlighten**DSM** is a registered trademark of Enlighten Software Solutions, Inc. Enlighten**DSM** for Unix is a trademark of Enlighten Software Solutions, Inc.

All other product names mentioned herein are the trademarks or registered trademarks of their respective owners.

Contents

CHAPTER 1: Introduction

License Advisory Window	1-3
-------------------------------	-----

CHAPTER 2: Configure

Session Preferences	2-2
Fields	2-2
Buttons	2-7
Alter Menu	2-8
Pool Configuration	2-10
Add	2-11
Modify	2-12
Delete	2-13
View	2-13
Select	2-13
User Authorization	2-14
Add	2-15
Modify	2-17
Delete	2-17
Copy	2-17
EMD Data Expiration	2-19
Fields	2-19
Buttons	2-20
About Enlighten Window	2-21

CHAPTER 3: User

Configure	3-2
Add	3-3
Multi-User Add	3-9
LDAP	3-14
Modify	3-17
Copy	3-17
Delete	3-18
Password	3-18
Lock	3-19
Unlock	3-19
Expire	3-20
View	3-20
Processes	3-22
NIS	3-22
Query	3-23
Fields	3-24
Buttons	3-26
Groups	3-27
Add	3-28
Modify	3-29
Delete	3-29
Assign Users	3-30
Activity Monitor	3-32
Who is Logged In	3-32
Process Status	3-34
CPU Summary	3-35
New User Templates	3-37
Add	3-38
Modify	3-43
Delete	3-44
Copy	3-44
User Management Scripts	3-45
Adding a User	3-45
Modifying a User	3-46
Deleting a User	3-46

CHAPTER 4: Printer

Configure (all) / Configure (local).....	4-2
Add	4-2
Delete	4-7
Enable	4-8

Disable	4-8
Jobs	4-8
Queue.....	4-9
Cancel	4-9
Enable	4-9
Disable	4-9
Move	4-10
Custom Driver	4-11
TTY Configure	4-16
Add	4-16
Modify	4-19
Enable	4-19
Disable	4-20

CHAPTER 5: Network

Host Entries	5-2
Add	5-3
Modify	5-4
Delete	5-5
Hosts File Import	5-5
NIS	5-6
Host Configuration	5-7
Host Naming Services	5-8
Modify Fields	5-9
Modify Buttons	5-11
Naming Service Overrides.....	5-12
Add	5-13
Modify	5-17
Delete	5-17
DNS	5-18
Server Configuration	5-18
Mail Exchangers	5-24
NIS and NIS+	5-28
Add NIS Server	5-29
Add NIS+ Server	5-31
Delete	5-34
NFS.....	5-37
Mounted Directories	5-37
Exported Directories	5-44

CHAPTER 6: Security

Vital Files	6-3
Filesystem Devices	6-4
Boot and Shutdown Script Checks	6-5
Crontab Contents	6-6
Password File Integrity	6-7
Group File Integrity	6-8
User Home Directories	6-9
Attempted Break-ins	6-10
Obvious Passwords	6-11
Levels of Checking	6-11
Who is Checked	6-12
What is Checked	6-13
Audit Log	6-14
Enabling and Disabling the Audit Log	6-14
Data Storage and Display	6-15

CHAPTER 7: Disk

Building Snapshots	7-2
Usage by Filesystem	7-3
User Sum.	7-3
Group Sum.	7-5
File Search	7-6
Processes	7-6
Graph	7-7
Usage by Username	7-8
View Files	7-9
Compare Snapshots	7-14
Save Snapshots	7-14
Graph	7-14
Usage by Groupname	7-15
View Files	7-16
Compare Snapshots	7-16
Save Snapshots	7-16
Graph	7-17
File Search	7-18
Fields	7-19
Buttons	7-22

Processes by Filesystem	7-23
Terminate	7-24
Hangup	7-24
Suspend	7-25
Continue	7-25
Priority	7-25
Save Current Snapshots	7-26
Compare Snapshots	7-29
Summary	7-31
File List	7-32
Search	7-34

CHAPTER 8: Archive

Configure Devices	8-2
Add	8-3
Modify	8-6
Delete	8-6
Tape Operations.....	8-7
Retention	8-7
Erase	8-7
Backup Partitions.....	8-8
Full Backup	8-9
Incremental Backup	8-9
Schedule Backup	8-10
Restore Files	8-11
Viewing an Archive	8-11
Restoring an Archive	8-12
Backup Catalog.....	8-14
Fields	8-15
Buttons	8-17
Scheduled Backups.....	8-18
Add	8-19
Modify	8-22
Delete	8-23
Copy	8-23
Execute Now	8-23
Status Logs	8-24

CHAPTER 9: System

CPU Summary by TTY	9-2
Processes	9-2
Graph	9-3
Swap Space	9-4
Processes	9-4
Graph	9-4
Shutdown/Reboot	9-5
Fields	9-5
Buttons	9-7
Cron Management	9-9
Configure	9-9
Query	9-13
Cron Users	9-15
Clock Management	9-18
Clock Synch	9-18
Clock Set	9-21
Mail Aliases	9-23
Configure	9-23
Query	9-29
Remote File Distribution	9-32
Add	9-33
Modify	9-37
Delete	9-38
Copy	9-38
Execute Now	9-38
Status Logs	9-38
Login Status	9-40

CHAPTER 10: Events

Configure	10-2
How EVENTS Works	10-3
Components	10-6
Configuration	10-8
Configuring Tests	10-8
The testtab.<hostname> File	10-11
Group Tests	10-14
Cache Tests	10-27
General Tests	10-31
Item Tests	10-32

API Tests	10-37
Generating Reports	10-39
Status Map	10-40
Map Navigation	10-41
Map Editing	10-41
Window Buttons	10-42
Information Areas	10-44
Host Overview	10-45
Combo Areas	10-45
Buttons	10-47
Events List	10-48
Query Events	10-49
Fields	10-49
Buttons	10-51

APPENDIX A: Wizard

Disk Wizard	A-4
Full Disk Partitions	A-4
Set UID/GID Problems	A-4
Devices Not in /dev	A-5
User Wizard	A-6
Duplicate UserID's	A-6
Duplicate Names	A-6
Vulnerable Directories	A-6
Obvious Passwords	A-6
Printer Wizard.....	A-7
Downed Printers	A-7
Long Print Queues	A-7
Large Print Jobs	A-7
Host Wizard.....	A-8
Host Address Conflicts	A-8
Host Entry Name Conflicts	A-8
Host Aliasing Conflicts	A-8

APPENDIX B: Altering the Menu

Overriding the Default Configuration File	B-3
Where to Edit	B-4
Defining Menus	B-5
Menu Attributes	B-5
Sample UserMenu1 Menu	B-8

Defining Menu Items	B-9
Menu Item Attributes	B-9
A Utilities Submenu	B-12
The Main Menu Window	B-14

APPENDIX C: Time Formats

Examples:	C-2
Absolute Date/Time BNF	C-3
Relative Date/Time BNF	C-4

APPENDIX D: Password Aging

Berkeley UNIX and Pre-system V.4	D-2
System V.4.....	D-3

APPENDIX E: SNMP Managed Groups

SNMP Groups.....	E-2
MIB II Groups	E-5
Enterprise-Specific Traps	E-6

APPENDIX F: Events Commands

AgentENL (8N)	F-2
AgentMon (8N)	F-3
EventsCli (8N)	F-5
FieldCk	F-8
testtab (1)	F-9

APPENDIX G: Sample Events Files

Sample testtab Configuration File	G-2
Sample EventsCli Use	G-9

APPENDIX H: O/S Compatibility

APPENDIX I: SYSADMIN Commands

netutar	I-2
p_nsnaps	I-4
p_snapdiff	I-5
Filesystem Comparisons.....	I-7

APPENDIX J: EMD

Basic Design	J-2
Relational Database	J-3
Directory Structure	J-4
Database Access	J-4
Manually Killing/Invoking EMD	J-6
Terminating EMD	J-6
Starting EMD	J-6
Audit Log Data	J-7
Audit Log Reports	J-7
Purging Audit Log Data	J-10

APPENDIX K: Macros

User Account Macros	K-2
Format	K-2
Examples	K-4
Form Letter Macros	K-6

APPENDIX L: Windows Support

Global Differences.....	L-3
Unix and Windows File Pathnames	L-3
File Ownership	L-4
Directory and File Access Times	L-4
File Permissions	L-5
Filesystem Limitations	L-5
Configure Menu.....	L-6
User Menu	L-7
User and Group ID's	L-7
Configure Options	L-7
Query Options	L-7
Groups Options	L-8
Login Status Command	L-8
Activity Monitor Commands	L-8
Printer Menu	L-10
Network Menu	L-11
Local vs. Domain Accounts	L-11
DNS Files	L-11
Host Entries	L-11
Naming Services Overrides Options	L-11

NIS and NIS+ Options	L-11
NFS Options	L-11
Security Menu.....	L-12
Disk Menu	L-13
Usage by Filesystem Options	L-13
File Search Options	L-13
Processes by Filesystem Options	L-14
Compare Snapshots	L-14
Partition Differences Query	L-14
Archive Menu	L-15
Configure Devices Options	L-15
System Menu	L-16
TTY Configure Options	L-16
Shutdown/Reboot Options	L-16
Cron Management Options	L-16
Mail Aliases	L-16
Remote File Distribution	L-16
Login Status	L-16
Events Menu	L-17
Events Monitoring	L-17
Enlighten DSM Events	L-17
Process Tests	L-17
Events E-mail Notification	L-18
MIBII Tests	L-18
Status Map	L-18



Introduction

This manual contains detailed information on Enlighten**DSM**'s features for system administration. This includes instructions for creating and modifying Events tests, and using the Status Map as a visual aid to track and fix problems.

These features are accessed using Enlighten**DSM**'s graphical user interface (GUI). Each main menu item in this GUI tool bar is described in a separate chapter:

- [Chapter 2, “Configure”](#) describes how to set up Enlighten**DSM** program preferences, create network pools, and configure user authorizations.
- [Chapter 3, “User”](#) describes how to configure user and group accounts, manage processes, check user security, and create new user templates.
- [Chapter 4, “Printer”](#) describes how to configure local and network printers, enable and disable printers, manage print queues, and manage TTY ports.
- [Chapter 5, “Network”](#) describes how to configure and manage networks, host entries, NFS file systems, and DNS, NIS, and NIS+ servers.

- [Chapter 6, “Security”](#) describes how to perform security checks and check for obvious passwords.
- [Chapter 7, “Disk”](#) describes how to search for files and observe disk utilization by groups, users, and file systems.
- [Chapter 8, “Archive”](#) describes how to configure archive devices, schedule backups, restore files and perform immediate backups.
- [Chapter 9, “System”](#) describes how to use EnlightenDSM to shut down and reboot systems, view summaries of process activity by tty, manage cron jobs, set system clocks, manage mail aliases, observe swap space usage, control remote file distribution and control logins.
- [Chapter 10, “Events”](#) describes how to configure Events parameters and how to use the Status Map GUI.

For further instructions on using Events to monitor networked systems, refer to Chapter 5 of the *EnlightenDSM User Guide*.

For general details on installing EnlightenDSM, navigating the GUI, or using the PEP processor, see the *EnlightenDSM User Guide*.

License Advisory Window

There are two versions of EnlightenDSM—Advanced and Workgroup. A License Advisory window similar to [Figure 1-1](#) will appear if you attempt to access an Advanced feature with the Workgroup version.

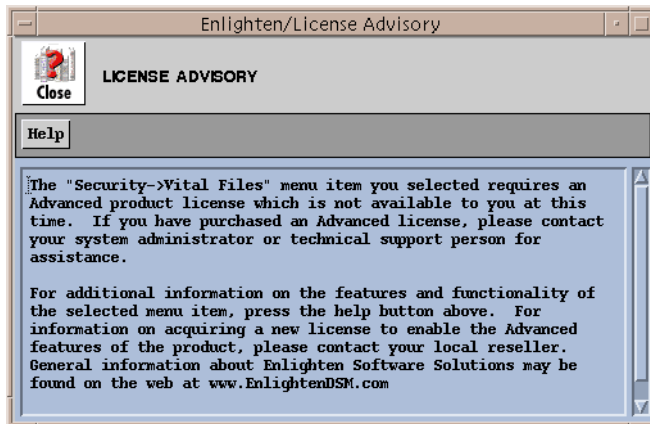


Figure 1-1 License Advisory window

For information on purchasing EnlightenDSM Advanced Version, please visit Enlighten Software's web site at <http://www.EnlightenDSM.com>.

Help Button

Click the Help button for information about the feature that caused the License Advisory window to appear ([Figure 1-2](#)).

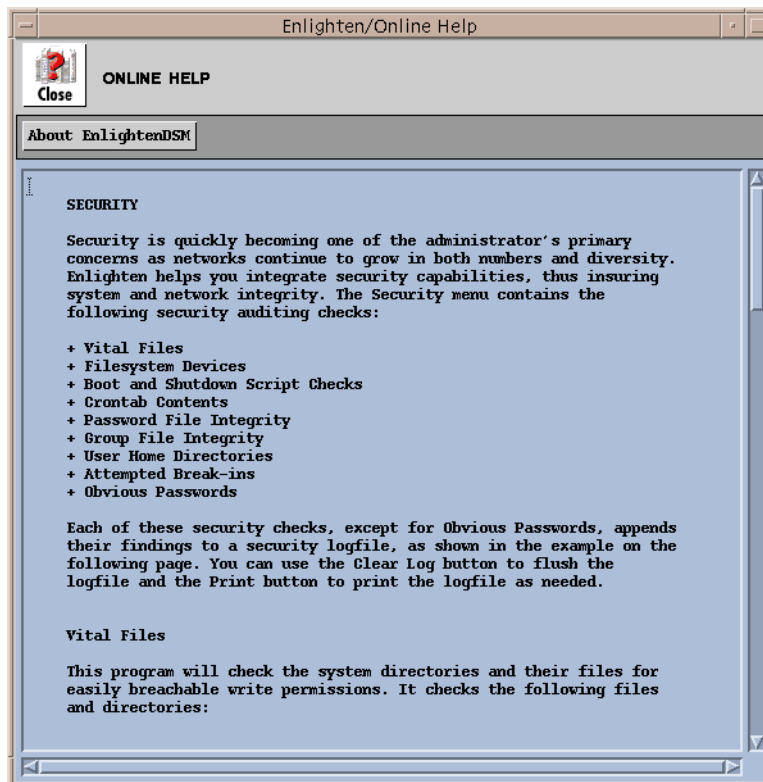


Figure 1-2 Online Help window



Configure

Enlighten**DSM** can be easily customized to meet individual or company needs. Use the Configure menu to set basic program preferences, alter or extend Enlighten**DSM** menu items, set EMD data expiration parameters, create network pools, and set user authorizations.

The Configure menu options are:

- Session Preferences
- Alter Menu
- Pool Configuration
- User Authorization
- EMD Data Expiration
- About Enlighten**DSM**

Session Preferences



You can use this module to set a number of parameters that affect how EnlightenDSM works during your session. Choose Session Preferences from the Configure menu. The Session Preferences window appears (Figure 2-1).

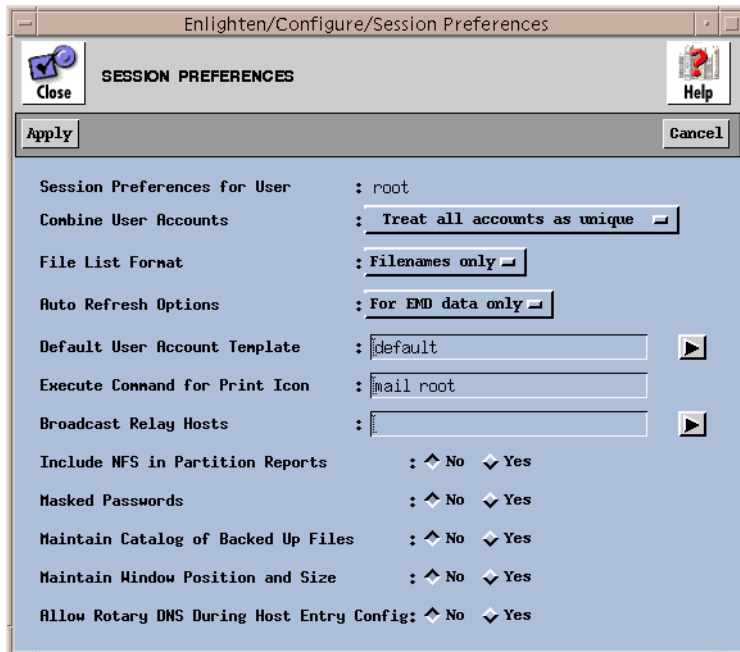


Figure 2-1 Session Preferences window

The rest of this section details how to use this window's functionality.

Fields

The Session Preferences window contains the following fields:

Session preferences for user

This view-only field shows the user for which these session preferences are being defined.

Combine User Accounts

Normally, Enlighten**DSM** combines accounts with identical parameters across the current system pool, and reports the account as a single entry with an incident count. For example:

```
caesar (3) Encrypted 101 20 /home/caesar
```

means user caesar has accounts on three of the systems in the current system pool, all of which have the same Username, Userid, primary Groupid, Login Shell, Home Directory, and so on. For many administration tasks, this form of reporting is easier to scan since it eliminates redundant entries.

In some instances, you may need to see a more detailed report of user configurations. You can change the results of Enlighten**DSM** reports by selecting that all accounts be treated as separate entries or that each account's uniqueness is determined by three parameters: Username, Userid, and primary Groupid.



Be careful using the latter form; the reported accounts may not be complete duplicates, so making modifications to the configuration might remove important differences.

Use this three-way toggle to decide how Enlighten**DSM** reports account configurations. The options are:

- Treat all accounts as unique (the default)
- With identical UserID & GroupID
- With all parameters identical (the first case shown above)

File list format

Use this toggle to determine how much information is shown when file names are displayed in a report list box. The Filenames Only setting (default) displays the absolute name of the file. Choose the Long Listing format to display additional information, such as the protection mode, file size, and date of last modification.

Auto refresh options

In a multi-user/tasking system, files are created and deleted, print requests are made, and processes start and terminate, and so system information displayed may no longer reflect the current state of the system. You can update a screen with current information by clicking the Refresh button. EnlightenDSM also automatically refreshes a screen after any action is taken on a displayed item in that window.

Use this three-way toggle to choose which auto refresh option you want to use during your session:

- Turned off
- For EMD data only
- Turned on

Turned off

Choose this option to prevent the current screen from being updated automatically. You will have to manually click the Refresh button to update the current list in *any* window.

For EMD data only

When this option is selected, EnlightenDSM will automatically update screens only where the data required to build the report is local EnlightenDSM data (that is, all EMD data).

Turned on

When this option is selected, EnlightenDSM will automatically refresh the current screen after taking an action on items in the display list. The screen refreshes when all actions have been completed in that window. This option is particularly useful when killing a process. Once EnlightenDSM has sent the desired signal to the process, the report is regenerated and you can quickly determine if the process still exists.



Some reports are generated based on the selections made from a parent window. For example, a list of processes may have been generated based on the selected users from the Process Summary report. The parent window is first regenerated and the same items are reselected. The criterion for reselection is solely based on position. So, if the first and third entries are selected in the original report, they will also be selected in the newly generated report. After the parent window has been regenerated, then the current window can be regenerated.

Disk-related windows do not have any refresh capabilities since they are based on previously generated disk snapshots.

Default User Account Template

User account templates simplify the creation of user accounts. The User account template allows you to set as the default a specific user template. Click the arrow button to the right of the field to display a list of previously defined templates from which to choose. For more information on defining these templates, see [“New User Templates” on page 3-37](#).

Execute command for Print Icon

The upper right corner of most list windows contains a set of buttons with icons representing global actions. When you click the Print icon, the full list is piped to the command entered in this field.

The default is to email the report list to the user `root`. You may use this field to create any executable UNIX command list here. For example, to print to the printer `text`, you could specify the following command:

```
lp -d text
```

You can use a maximum of 100 characters to specify your custom printing command.

Broadcast Relay Hosts

You can use this field to specify any hosts that will be used as broadcast relays for remote file distribution jobs. If you are using multiple entries, leave a blank between each entry.

Include NFS in Partition Reports

Use this toggle to select whether EnlightenDSM should include NFS-mounted partitions when it displays a list of disk partitions (you can use for further manipulation). The default setting is No.

Although including these partitions gives you a more complete set of choices, the system information available for these partitions is not as detailed as it is for local partitions. Also, building and searching NFS partitions are also relatively slow processes. Finally, within an EnlightenDSM system pool, the NFS partitions on different systems may be redundant mounts of a system within the Pool that already has the same partitions mounted locally. In this case, the local daemon will operate more efficiently on that partition.

Masked Passwords

Use this toggle to select whether EnlightenDSM echoes asterisks in the Password fields when you are assigning passwords. If you select Yes, the password will be masked when you enter it. If you select No (the default), the newly assigned password will be displayed as you enter it.



No confirmation is required when you assign a password. Be careful when assigning masked passwords to ensure what you type will match what the system will expect.

Maintain catalog of backed up files

Use this toggle to select whether EnlightenDSM will keep a catalog of all files it backs up. The default is No.

If you choose Yes as the option, EnlightenDSM creates a catalog of all files it has backed up every time a partition backup is made. It stores the date and time the backup was made, the name of the file, the file size, the file owners, and the date of last modification. The default name of the catalog is NONAME.

You can use the Archive menu to search through the catalogs for instances of backed up files and/or restore selected files. See [“Backup Catalog” on page 8-14](#) for more details on how to use these options.

Maintain window position and size

EnlightenDSM windows will pop up in a default location with a default size. If you select the Yes setting, any changes to window size and location are saved. If you choose the No setting, the window will revert to the default size and location settings the next time you open the window.

Allow Rotary DNS During Host Entry Config

By default, EnlightenDSM will check for duplicate host IP addresses. Checking NO in this field will allow duplicate IP addresses to be entered into the DNS table.

Buttons

The Session Preferences window contains the following buttons:

Apply

Click this button to save the session preferences defined by the rest of this window’s fields.

Cancel

Click this button to close the window without making any changes.

Alter Menu



You can use this module to incorporate your own scripts or restrict accessibility to specific EnlightenmentDSM menu items by providing a user-alterable interface. Choose Alter Menu from the Configure menu. The Alter Menu window appears (Figure 2-2).

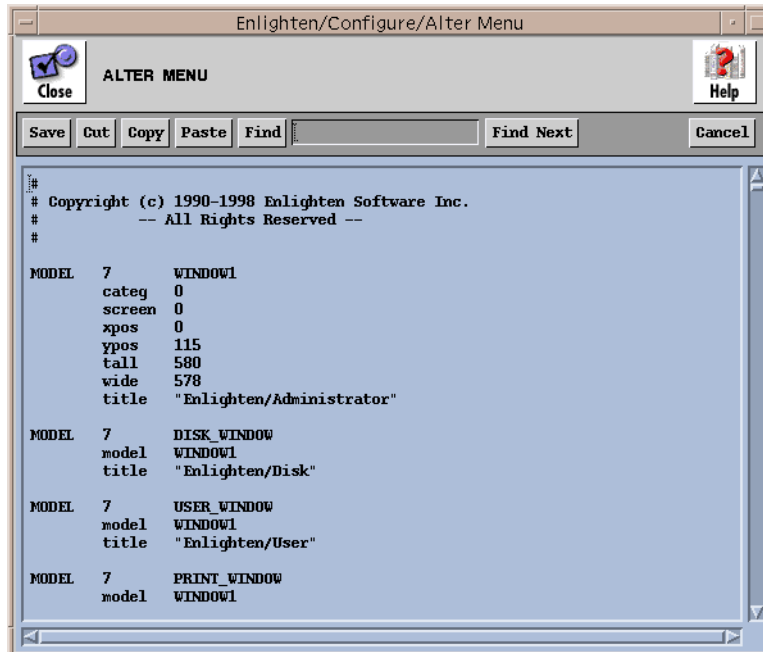


Figure 2-2 Alter Menu window

The following items may be altered:

- The order of menu icons
- The order of menu items
- The actions of menu items
- The number of icons and menu items

The EnlightenmentDSM graphical user interface is built from a small text database (file) containing objects and attribute lists. This file contains information describing the layout of the windows. By editing this file, you can create alternate interfaces. For example, you may wish to add a menu item with your own script as the action item.

When you select the Alter Menu menu item, a window appears showing the text database. You can edit the object and attribute lists within this as you need. See [Appendix B, “Altering the Menu,”](#) for more details on using this option.

Pool Configuration



You can use this module to display a list of all currently configured network pools, as shown in the Network Pool Configuration window (Figure 2-3). The rest of this section describes how to use this window's buttons to perform a variety of pool-related actions.

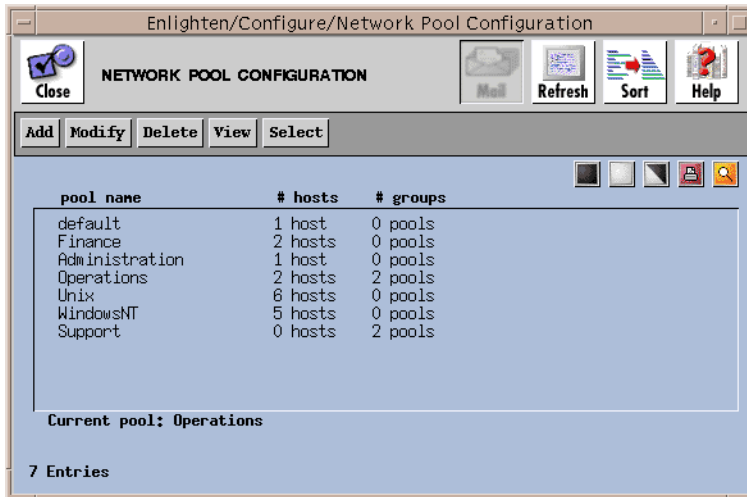


Figure 2-3 Network Pool Configuration window

From the Network Pool Configuration window you can:

- **Add** Create a new network pool
- **Modify** Modify the highlighted pool
- **Delete** Delete the highlighted pool
- **View** See a list of all (effective) hosts in the highlighted pool
- **Select** Make the first highlighted pool the default (or focus) pool

Add

Click the Add button to create a new network pool. The Network Pool Add window will appear (Figure 2-4).

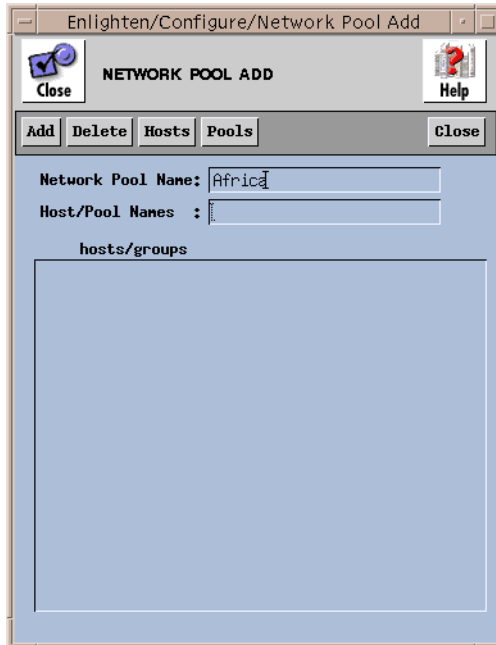


Figure 2-4 Network Pool Add window

Buttons

This window contains the following buttons:

Add

When you configure the pool with the hosts and sub-pools you want, click the Add button to save it.

Delete

To delete selected host or pool entries from the current pool, click the Delete button. EnlightenDSM will prompt you to confirm your action.

Hosts

Click the Hosts button to bring up a window showing all the currently configured Hosts. EnlightenDSM generates the Hosts database from the `/etc/hosts` file the first time EnlightenDSM is executed. Highlight the selections you want and then click the Apply button to add the hosts to the new pool.

Pools

Click the Pools button to bring up a window showing all the currently configured Pools. Highlight the desired selections and then click the Apply button to add the pools to the new pool.

Cancel

Click this button to close the window without making any changes.

Fields

This window contains the following fields:

Network Pool Name

Use this field to specify the name of the new pool.

Host/Pool Names

Use this field to add hosts or pools or any combination thereof to the newly created pool. You can enter these names or use the Hosts and/or Pools button(s). If you are using multiple entries, leave a blank between each entry.

Modify

Click this button to modify a pool's configuration. A pop-up window similar to the Network Pool Add window will appear, except the Network Pool Name field is view-only.

You can click the Modify button (rather than the Add button) after you've made all your changes.

For a description of the rest of the buttons and fields in this window, see ["Add" on page 2-15](#).

Delete

Select one or more pools from the list box and click the Delete button to remove it or them from the recognized set of configured pools. Enlighten**DSM** will prompt you to confirm your action.

View

Select one or more pools from the list and click the View button to see a list of all (effective) hosts in the named network pool(s). If you have highlighted more than one selection, click the Next button to view the hosts in the next selected pool.

Select

Click the Select button to make the first highlighted pool the default (or focus) pool. Once a pool has been selected as a focus, all further Enlighten**DSM**-related action and inquiries will be performed on all hosts in that pool.

User Authorization



You can use this module to create sub-administrators who can easily manage their specific environment. Choose User Authorization from the Configure menu. The User Authorization Configuration window will appear (Figure 2-5).

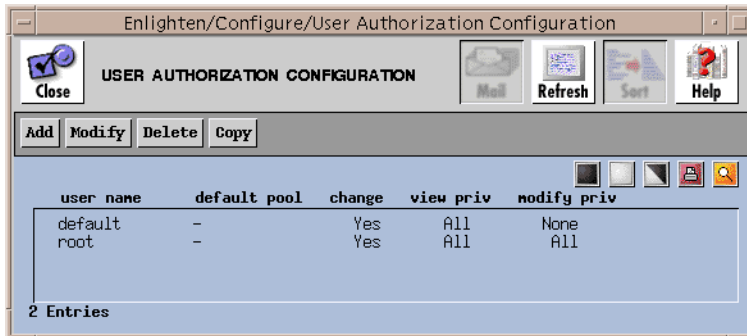


Figure 2-5 User Authorization Configuration window

From here, you have the option to:

- Add Add a user to the user authorization list
- Modify Modify the user privileges
- Delete Delete the user from the user authorization list
- Copy Create a second set of user privileges using the selected user's privileges as a starting point



For security reasons, users authorized to access EnlightenDSM must also be added to the enldsm group. The root user or any user previously authorized to Add/Modify User Account information can add the user to the enldsm group through the User Modify functionality of EnlightenDSM. See [“Modify” on page 3-17](#).

Add

Click the Add button to add a user to the user authorization list. The Add User Authorization window will appear (Figure 2-6). Use this window to select which manageable objects a user can access and what type of access the user will have.



Figure 2-6 Add User Authorization window

Fields

The Add User Authorization window has the following fields:

User Name

Use this field to specify which user will receive these authorization privileges. You can also click the arrow button on the right to select from a list of known users.

Default Pool

Use this field to specify which pool of hosts the user should belong to when he or she starts EnlightenDSM. You can also click the arrow button on the right to display a list of previously defined pools and select one.

Change Pool

Use this toggle button to specify if the user may change his or her pool. The default setting is Yes.

View Privileges

The Yes setting (default) allows the user to view the network task listed. To revoke the viewing privilege for any task listed, select the No setting.

Modify Privileges

The Yes setting (default) allows the user to modify the network task listed. To revoke the modifying privilege for any task listed, select the No setting.



If the user's View Privileges for any object are disabled (No), the user's Modify Privileges for that object will also automatically be disabled.

Buttons

The Add User Authorization window has the following buttons:

Add

After you've made all the selections for the new user authorization, click this button to save the current user configuration. The new authorized user will then appear in the list box within the User Authorization Configuration window.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Modify

Click this button to modify a user's privileges. A pop-up window similar to the Add User Authorization window will appear, except the User Name field is view-only.

There are also two button differences in the Modify window:

- Click the Modify button (rather than the Add button) after you've made all your changes, or
- Click the Next button to modify additional user privileges if you've selected more than one user to modify from the User Authorization list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 2-15](#).

Delete

Click this button to delete the selected user(s) from the user authorization list. EnlightenDSM will prompt you to confirm your action.

Copy

Click this button to copy a set of user privileges to a second user. The Add User Authorization window will appear showing the highlighted user's privilege settings in each of the fields. You can edit this window as needed and then click the Add button to complete the copy.

See [“Add” on page 2-15](#) for a description of how to use this window's fields and buttons.



Enlighten**DSM** is shipped with a `default` user configuration for all users. `root` always retains the ability to modify the privileges on all other users.

EMD Data Expiration



EnlightenDSM stores the information it may need to access later in a common repository called Enterprise Management Database (EMD). Some of this information, especially the log files, can grow quite large over time.

You can use this module to set default parameters for how long you want EnlightenDSM to keep some of this log information around. Choose EMD Data Expiration from the Configure menu. The EMD Data Expiration Parameters window will appear (Figure 2-7).

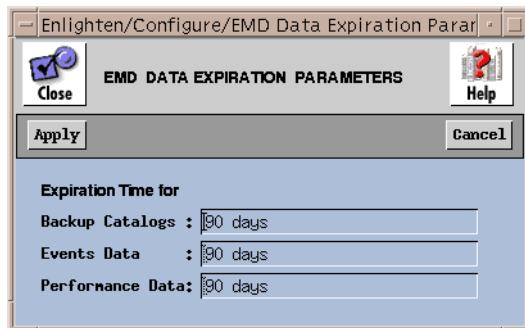


Figure 2-7 EMD Data Expiration Parameters window

Fields

The default setting for all fields is 90 days. The minimum time increment you can specify for any of these fields is one day. See [Appendix C, “Time Formats,”](#) for more information about the time formats you can use in these fields.

The window has the following fields:

Backup Catalogs

Use this field to specify a data expiration time for all data associated with maintaining the Backup Catalogs. This includes the data stored in the `backup_tbl`, the `cat_index_tbl`, and the `file_index_tbl`.

Events Data

Use this field to specify a data expiration time for all data associated with the event messages that PEP sends to the EMD. This includes the data stored in the `events_tbl`.

Performance Data

Use this field to specify a data expiration time for all log messages AgentMon sends to the EMD. This includes unused Hardware and Software index data. Data will also be expired from `event_log_tbl`, `hrdw_index_tbl`, and `sftw_index_tbl`.

Audit Log Data

Use this field to specify a data expiration time for audit log entries sent to the EMD from the distributed agents and the GUI.



Every Events test can potentially be logged to the EMD.

Buttons

Click the Apply button once you've chosen the times you want for these categories. EnlightenDSM will then remove any of this type of data as it becomes older than the times you've specified (expires). Click the Cancel button to ignore any changes and close this window.

About Enlighten Window



The About Enlighten window (Figure 2-8) contains information for contacting Enlighten Software Solutions in both the United States and Europe. After you read the information in the About EnlightenDSM window, click Close to dismiss it.



Figure 2-8 About Enlighten window



User

With Enlighten**DSM**'s User menu, administrators can easily manage user and group accounts on either a single file server or across a network of heterogeneous systems.

This chapter describes how to configure users and groups, query for user information, monitor user activity, and create new user templates.

The User menu options are:

- Configure
- Query
- Groups
- Activity Monitor
- New User Templates

Note: Enlighten**DSM** supports `/etc/passwd` file, NIS, and NIS+ naming services.

Configure



You can use this module to manage user access and monitor user processes. When you activate this program, a window will appear displaying a list of all current user accounts (Figure 3-1).

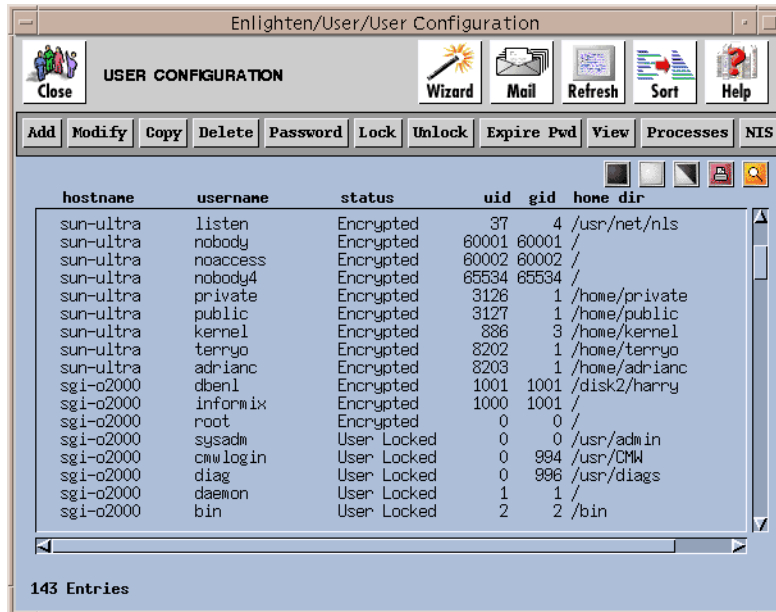


Figure 3-1 Current User Accounts List

You can use this window to:

- Add a user
- Modify a user
- Copy a User
- Delete a user
- Set or modify a password for a user
- Lock a user account
- Unlock a user account
- Expire a user's password
- View a user's information
- Monitor user processes
- Push NIS maps on any active NIS servers

The rest of this section details each of these options.

Add

Click the Add button to create user accounts. You can do this individually or en masse. The Create New User Account window (Figure 3-2) also gives you the option to use default parameters, as defined in the Configure —> New User Templates window. These defaults ensure all new user accounts in a particular type are created with consistent information (see “New User Templates” on page 3-37 for more details).

Figure 3-2 Create New User Account window

If you leave one of the fields blank when adding a user account, EnlightenDSM will fill in the field with the default value (where possible) or prompt you for any additional information. After you’ve filled in all the fields properly and clicked on the Add button, EnlightenDSM will create the user account and home directory, update the appropriate naming service, and create the home directory with the required shell scripts.

The rest of this subsection describes the functionality of this window's fields and buttons.

Fields

You can fill in all of these fields manually or have EnlightenDSM set defaults in most of the fields for you. Either way, you need to specify the following five fields before trying to add the user account.

Hostnames

If you want to limit the creation of this account to specific hostnames within a pool, enter those hostnames in this field. If you are using multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select the available hosts from the current pool.

Login name

The Login name is a unique name that identifies a user account. A user needs a login name, and a password to gain access to the system. The login name must be unique and the first character should be set in lower case (e.g., cHE034). You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Real name

This is the user's actual name. This field can identify a person, company, or organization. You can use this field to associate the sometimes cryptic Login name with the actual user's name. On larger systems, this name matching is invaluable.

Office

This field should contain the location where the user will normally access the system. The field can be an office number, an assigned location code, or a department name. This field does not allow punctuation characters.

Telephone

This field can contain the telephone number, telephone extension, or some other information to help locate and contact the user. The format of the field is user definable.

At this point you can have Enlighten**DSM** automatically fill in the remainder of the form by clicking on the Add button. The default values entered depend on the values set in the Configure → New User Templates window.

Or, you can fill in the rest of the fields manually. The fields and their defaults are:

User ID

This field contains the UID number. The UNIX operating system uses this number to identify each user. This number does not have to be unique, but should be for easier system administration. Enter the UserID number to be assigned to this user account.

If the New User Template you're using is set to automatically generate a unique UID, then this field will be initialized with that number. See [“New User Templates” on page 3-37](#) for more details.



If this field is left blank, Enlighten**DSM** will prompt you for an entry in this field when it tries to create the user.

Home dir.

The Home dir. field specifies where the user's start-up directory is attached. Most systems will use `/usr` or `/home` as the default parent directory. Each user will have their own directory created under the default home directory, for example, `/usr/charlie`. Enlighten**DSM** stores `home/%u` as the default HOME directory unless you change this field.

Enlighten**DSM** initializes the HOME directory by executing the *HOME directory initialization* script `$ENLIGHTEN/scripts/adduser.rc`. This copies the initialization files (e.g., `.profile` and `.cshrc`) into the HOME directory. You can modify this script to perform other initialization procedures. See User Management Scripts at the end of this chapter for instructions on doing so.

You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Shell

The Shell is a program that acts as a translator between the user and the UNIX operating system. There may be several standard shells on a system, including the C Shell, the Bourne Shell, the Korn Shell, and the restricted Bourne, C, and Korn Shells. This, however, does not stop experienced users from creating custom shells. Each shell has its own start-up executable program:

Bourne Shell	/bin/sh
Restricted Bourne Shell	/bin/rsh
C Shell	/bin/csh
Korn Shell	/bin/ksh
Restricted Korn Shell	/bin/rksh
Other	user defined

EnlightenDSM uses the Bourne shell as the default shell in this field (/bin/sh). All users on the system do not have to use the same shell; this is a matter of user preference, as each shell offers different functionality. The number of standard shells available for this field depends on your system. Click the arrow button to display a pick list of available shells. This list of available shells comes from the GUI host, *not* the host to which the user will be added.

Group name

Every user account must be assigned to at least one User Group. The User Group is part of the UNIX security system; each user group has specific read, write, and execute privileges associated with every file created on the system.

Enter the primary Group name for this user or leave this field blank to have the default assigned. Click the arrow button for a pick list of defined User Groups. If you enter a Group Name that does not exist, EnlightenDSM will prompt you for a correct Group Name when it tries to create the user.

Add Groups

You can use this field to put the user in additional user groups. This will effectively give the user more group privileges beyond those set in the primary user group. Enter the names of the user groups to which the user should also belong. Click the arrow button for a pick list of defined User Groups. Select the User Groups and click the Apply button to import them.

Password

The password is part of the UNIX security system. Each user must have a password. You can make the password the same as the user login, and the user can then change the password. Don't use passwords such as spouse names, pet names, or addresses, they are easy for another user to guess. You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.



If the entry in the password field is not long enough, an error message will be displayed when you try to add this user account.

Expires

For systems that support password aging, this optional field contains the time when the password will expire. Once a password has expired, the system will force the user to change their password the next time they log in to the system.

The password expiration time may be rounded up depending on the minimum password expiration interval for the OS.

If this field is left blank, the appropriate default parameter will be used. If no default is set in the User Add window (when you're adding a user), password aging will be turned off for that user. For more details, refer to [Appendix D, “Password Aging.”](#)

To specify a time format, see [Appendix C, “Time Formats.”](#)

Span

In conjunction with the Expires field, this field defines the minimum period of time that needs to elapse before the user can change his or her password. You

can, however, change the user's password at any time by using the Password button in the User Configuration window. See [“Password” on page 3-18](#) for more details. To specify a time format, see [Appendix C, “Time Formats.”](#)

Mail Alias (Create New User Account Window only)

This field specifies the default mail alias the account will use. An alias designates a short name as the substitute for the full pathname for a particular user or group. You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Mail Lists

This field specifies which mailing lists will be updated to include this user. A mail list defines a group of people under one alias, so sending mail to that alias will distribute a copy to everyone on that list. Click the arrow button for a pick list of defined Mail Lists.

Buttons

This window contains the following buttons:

Add

Once you've entered the relevant information, click the Add button to create the user. If the information is incorrect, use the mouse to navigate through the fields and make the necessary changes. When complete, click the Add button again.

Clear Fields

Click the Clear Fields button to clear the existing choices in all fields.

Multi-User Add

Refer to [“Multi-User Add” on page 3-9](#) to use this button and the window it generates.

Template

Click this button for a pick list of defined New User Templates. You can use one of these templates to populate default values in the new user's account, make any necessary changes, and then Add the user. Only those fields that are currently blank (in the Create New User Account window) will be initialized from your selected template.

Close

Click this button to discard any changes and close the window.

Multi-User Add

You can also use the Create New User Account window to add multiple users in a single operation. Click the Multi-User Add button. The pop-up window shown in [Figure 3-3](#) is then activated for creating new users en masse.

The screenshot shows a window titled "Enlighten/User" with the subtitle "ADD MULTIPLE USERS FROM TEXT FILE". The window has a "Close" button on the top left and a "Help" button on the top right. Below the title bar are three buttons: "Apply", "Clear Fields", and "Template", followed by another "Close" button on the far right. The main area contains the following fields and options:

- Filename of Accounts DB: [text input]
- Hostnames: [text input] ▶
- Starting User ID Number: [text input with value "100"] ▶
- Default User Group: [text input] ▶
- Default Home Directory: [/usr] ▶
- Default Startup Shell: [/bin/sh] ▶
- Mail Alias: [text input]
- Mail Lists: [text input] ▶
- Default Password Configuration:
 - Add all users with same initial password
 - Generate random password for each user
 - Password same as user name
- Initial Password: [text input]
- Password to Expire: [text input]
- Span Between Changes: [text input]

Figure 3-3 Creating Multiple New Users

You must first create a text file before you can use this program properly. Each line is the record for one user account. See [“Record Format” on page 3-13](#) for details on creating this file.

The rest of this subsection describes the functionality of this window’s fields and buttons.

Fields

This window contains the following fields:

Filename of Accounts DB

Use this field to specify which database file contains the user account information.

Hostnames

If you want to limit this user's access to specific hostnames within a pool, enter those hostnames in this field. If you are entering multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select the available hosts from the current pool.

Starting User ID Number

If the Userid field (field #5) is left blank in any of the records, EnlightenDSM will assign the first available (unique) Userid found (higher than the UID entered in this field). Click the arrow button on the right to view the already assigned UserIDs. A pop-up menu will appear with a current list of UserIDs. Highlight the desired ID number and then click the Apply button. EnlightenDSM will automatically insert the starting Userid into the appropriate record.

Default User Group

If the GroupID Number field (field #6) is left blank in any of the records, EnlightenDSM will use this field to assign those users a default Group Name. Click the arrow button on the right to view a current list of User Groups. Highlight the desired entry in the pop-up menu and then click the Apply button to make the selection. EnlightenDSM will automatically insert the Group Name into the appropriate record.

Default Home directory

If the Home Directory field (field #7) is left blank in any of the records, EnlightenDSM will use this field to assign those users a default Home directory. Use the full pathname of the parent directory where all user's HOME directories will be created.

If you specify a HOME directory that does not exist, EnlightenDSM creates and initializes it. EnlightenDSM initializes the HOME directory by executing the *HOME directory initialization* script

`$ENLIGHTEN/scripts/adduser.rc`. This copies the initialization files (e.g., `.profile` and `.cshrc`) into the `HOME` directory. You can modify this script to perform other initialization procedures. See [User Management Scripts](#) at the end of this chapter for instructions on doing so.

You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Default Startup Shell

If the Shell field (field #8) is left blank in any of the records, EnlightenDSM will use this field to assign those users a default shell. Click the arrow button on the right to view a current list of shells. Highlight the desired entry in the pop-up menu and then click the Apply button to make the selection. EnlightenDSM will automatically insert the shell name into the appropriate record.

Mail Alias

This field specifies the default mail alias the account will use. An alias designates a short name as the substitute for the full pathname for a particular user or group. You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Mail Lists

This field specifies which mailing lists will be updated to include this user. A mail list defines a group of people under one alias, so sending mail to that alias will distribute a copy to everyone on that list. Click the arrow button for a pick list of defined Mail Lists.

Default Password Configuration

When a user is created in UNIX, he or she must be assigned a password. If the password field (field #9) is left blank in the database record, the system must be told which method to use to create a password. Choose one of the following options as the default:

- Add all users with the same initial password (the default)
- Generate random password for each user
- Password same as user name

Initial Password

Use this field if you choose to add all users with the same initial password. Enter the password to be assigned to all the user accounts created from this process. Remember to check later and make sure each new user has changed his or her initial password. You may also use macros in this field; see [“User Account Macros” on page K-2](#) for more details.

Password to expire

For systems that support password aging, this optional field contains the time when the password will expire. Once a password has expired, the system will force the user to change their password the next time they log in to the system.

If this field is left blank, the appropriate default parameter will be used. If no default is set in the User Add window (when you’re adding a user), password aging will be turned off for that user. For more details, refer to [Appendix D, “Password Aging.”](#)

To specify a time format; see [Appendix C, “Time Formats.”](#)

Span between changes

In conjunction with the Expires field, this field defines the minimum period of time that needs to elapse before the user can change his or her password. You can, however, change the user’s password at any time via EnlightenDSM.

To specify a time format, see [Appendix C, “Time Formats.”](#)

Buttons

This window contains the following buttons:

Add

Once you’ve entered the relevant information, click the Add button to start the user account creation process. EnlightenDSM first checks that all field entries are valid and then starts creating the new user accounts.

EnlightenDSM will then output a list of generated passwords after all the accounts have been created. Click the Print button if you want to print the list.

Clear Fields

Click this button to clear the existing choices in all fields.

Template

Click this button for a pick list of defined New User Templates. You can use one of these templates to populate default values in the new user's account, make any necessary changes, and then Add the user. Only those fields that are currently blank (in the Create New User Account window) will be initialized from your selected template.

LDAP

The LDAP icon will query the LDAP server(s) specified within the params configuration file every time it is depressed and load the configured input fields with addition user information found on the LDAP server.

See [“LDAP” on page 3-14](#) for further information on using this feature.

Close

Click this button to discard any changes and close the window.

Record Format

The format for each record in the text file is:

Field #	Description	Comment
1	Login Name	Mandatory
2	User Real Name	Mandatory
3	Office	Optional
4	Telephone	Optional
5	Userid Number	Optional
6	Groupid Number	Optional
7	Home Directory	Optional
8	Shell	Optional
9	Password	Optional

Each record (line) in the file represents one user. Each user can have up to nine fields in the record. Use a comma (,) to separate each field specified. Each record is terminated with a carriage return (<CR>). There is no limit to the amount of records you can specify in the file.

The following example sets the account details for five people.

```
fred,The Fredster
barney,The Barnarama,Head office
wilma,Wilma Fraglerock,,233-7625
betty,Betty Boomrock,,,401,other
bambam,Boom,,,403,other,/usr/bambam,/bin/sh,bedrock
```

LDAP

The LDAP icon will query the LDAP server(s) specified within the params configuration file every time it is depressed and load the configured input fields with addition user information found on the LDAP server.

LDAP - Lightweight Directory Access Protocol

The interface provided allows for information about a user to be retrieved from a LDAP server to aid in the creation of the user's UNIX account. Thru the use of keyword parameters specified in the PARAMS configuration file, located in the /config subdirectory under the \$ENLIGHTEN directory tree, the ADD USER window is able to load field values about a user when creating an account. The following keywords are used in the PARAMS file to configure this functionality:

LDAPSERVER: ON | OFF

This entry enables or disables the LDAP functionality within the ADD USER window.

LDAPSERVERURL: <a URL string>

A LDAP URL string has a format of
LDAP://<host>[:<port>]/<dn>

Host may be a single entry or a space separated list of entries being either hostnames or IP addresses. :<Port> is an optional alternate port specifier which can be used if the LDAP server uses a

non-standard port number. <dn> is a LDAP Distinguished Name string which specifies the top of the search tree for locating users.

LDAPREALNAME: <attribute>[<sep><attribute>...]

LDAPOFFICE: <attribute>[<sep><attribute>...]

LDAPTELEPHONE: <attribute>[<sep><attribute>...]

LDAPLOGINNAME: <attribute>[<sep><attribute>...]

LDAPUSERID: <attribute>[<sep><attribute>...]

These keywords are optional and when present specify a string to be placed in their corresponding field in the ADD USER window. Attribute is a LDAP Attribute name which is replaced by its value for the user located. Sep is any non-alpha character. This syntax allows for a list of attributes using any non-alpha character, including a space, for separating the attributes.

Within the ADD USER window the contents of the Real Name field is used for locating a user via the LDAP server(s). The value may be a user's full name separated with spaces or a partial name using strings that begin the user's first and/or last names. The LDAP icon in the ADD USER window will use this value and query the LDAP server(s) every time it is depressed. If an exact match can not be found a dialog box will be displayed asking if another search should be performed using the value as a partial specification of the user's name. If more than one user is located or if the user located was not an exact match then a list box will be displayed containing the entries returned by the LDAP server(s). When an exact match is found or an entry is selected from the list box the keywords in the params file that represent input fields which have also not had any value previously entered, will have the evaluated attribute list specified loaded for the LDAP user located.

CONFIGURATION

To use the LDAP facility you will need to obtain the name or IP address of your LDAP server, the Distinguished Name of your organization within the LDAP server, and the LDAP Attribute names that you wish to use as field values for the ADD USER window. You may also need a port number if your organization has located their LDAP server on a non-standard port. Your System Administrator should be able to supply you with all of the above mentioned information.

First you need to modify the params configuration file for EnlightenDSM. This file is located in the /config subdirectory under the \$ENLIGHTEN directory tree and can be edited with any standard text editor like vi or emacs using the file name \$ENLIGHTEN/config/params. From with the text editor search for the string LDAP to find the section containing the LDAP keywords. The first keyword should be LDAPSERVER and will need its value of OFF changed to ON. Next you will need to add your organizations LDAP server name and Distinguished Name to the LDAPSERVERURL keyword. For example if your LDAP server was located on a machine named boris and had a Distinguished Name of o=Enlighten, c=US then you would append the string:

```
LDAP://boris/o=Enlighten,c=US
```

to the keyword LDAPSERVERURL. And if there was a non-standard port number being used like 89:

```
LDAP://boris:89/o=Enlighten,c=US
```

Last, you will need to decide which input fields from the list of allowable fields (REALNAME,OFFICE,TELEPHONE,LOGINNAME,USERID) that you would like to have filled in with attribute values from the LDAP server. For example, if you want to have the input field OFFICE to be filled in with the value from the LDAP attributes building and roomnumber you would append the string:

```
building,roomnumber
```

to the keyword LDAPOFFICE.

Using the above examples your params configuration file should like as follows:

```
LDAPSERVER          ON
LDAPSERVERURL       LDAP://boris/o=Enlighten,c=US
LDAPOFFICE          building,roomnumber
```

You are configured to use the LDAP facility from within the ADD USER window. To test your configuration settings run the EnlightenDSM product and navigate to the ADD USER window from the Configure Users menu entry under the USERS icon. The correct window will have a LDAP icon

second from the end in the middle box of the window. Type a single letter into the Real Name input field within this window and depress the LDAP icon. You should receive a dialog box asking to repeat the search as a exact match was not found. Answer yes by clicking the Yes icon within the dialog box and you should see a list box displayed with the first 20 user's whose names start with the letter type. Select one of the users. The building number and room number should now be displayed in the office input field. You may now either cancel the window or continue to add the user account to the system.

Modify

Click this button to modify user accounts. A pop-up window similar to the Create New User Account window will appear, except you cannot modify the Login name or Password fields.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional user accounts if you've selected more than one account to modify from the User Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 3-3](#).

Copy

Click this button to copy user accounts from one host to another host. A pop-up window similar to the Create New User Account window will appear. There are three differences between these windows:

- The Login Name and Password fields are read-only in the Copy window.
- Use the Copy button to save changes (rather than the Add button).
- Use the Next button to modify additional accounts if more than one account was chosen for copying.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 3-3](#).

Delete

Click this button to easily delete a user from the system. This program will delete the user account and also (optionally) remove the HOME directory and all of its contents (files and subdirectories) associated with the user. Highlight the username(s) you wish to delete from the system and then click the Delete button. EnlightenDSM will prompt you to confirm your action.



When deleting local user accounts across multiple hosts, the user's home directory will not be deleted if the home directory is the same on all hosts. EnlightenDSM assumes they are shared.

EnlightenDSM will NOT delete the users HOME directory if:

- It is considered to be a required directory. The directories are:

```
/usr/etc/bin  
/lib/home/dev/sbin  
/usr/bin/usr/adm/usr/lib/usr/sbin  
/usr/man/usr/lib/uucp
```
- It is shared by more than one user. In this case, the HOME directory is not deleted since this would also delete the HOME directory of any sharing co-users.
- It is not owned by the user whose HOME directory it is supposed to be. In this case, EnlightenDSM will then ask for your confirmation before deleting it.

Furthermore, EnlightenDSM will not delete the root user account; this is considered to be too dangerous to allow.

Password

Click this button to change a currently defined User Account password. You don't need to know the password. A pop-up window will ask you to set the new password for the highlighted user ([Figure 3-4](#)).

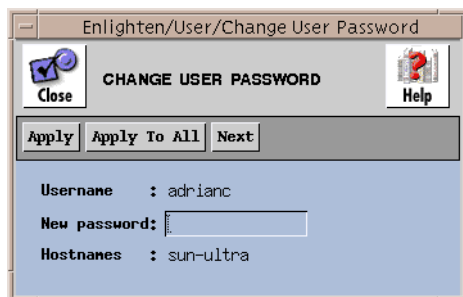


Figure 3-4 Changing User Account Password

If you have selected multiple usernames, clicking the Apply button will set the current password and bring up the next user. Clicking the Next button causes EnlightenDSM to skip to the next selected user without changing the user's password.

Click the Close button to discontinue changing passwords. Any previous changes made by using the Apply button remain; they are not undone by clicking this button.

Lock

Click this button to deny a specific user access to the system without deleting the user's account. This procedure is part of the EnlightenDSM security module; be cautious about who can access it. A pop-up window will prompt you to confirm this action.



You must assign a new password to unlock a user—the existing password will be lost! See the previous section to assign a password.

Unlock

Click this button to remove a lock from a user account. Because locking and unlocking user accounts is considered part of system security, EnlightenDSM will ask for a new password when it unlocks a User Account. See [“Password” on page 3-18](#) to create the new password.

Expire

Click this button to expire the password for the selected user account(s) immediately. This option will only work on systems that support password aging. The next time the user attempts to log in to the system, he or she will be forced to change his/her password. You must confirm this option before it will be executed.

For more details on password aging, see [Appendix D, “Password Aging.”](#)

View

Click this button to view existing User Account information. A pop-up window will appear ([Figure 3-5](#)).

The View User window is similar to the Create New User Account window. It has two additional fields that contain:

- The date and time the user last logged in to the system (if available), and
- The date and possibly the time the user last changed their password (if available).



Figure 3-5 User Account Information

This is a read-only window; no modifications may be made. If you have highlighted multiple usernames, click the Next button to see the next highlighted username.

See [“Add” on page 3-3](#) for a description of the display fields in this window. Refer to [“Modify” on page 3-17](#) to modify this information.

Processes

This command displays a window of all processes currently running that belong to the highlighted users. To view the processes, highlight the users you wish to view and then click the Processes button. A window will appear displaying all processes for the highlighted users. To further manipulate this information, see [“Process Status” on page 3-34](#).

NIS

Click this button to push/remake the NIS maps on any active NIS servers. This updates the NIS servers and hosts on your network with any Host Entry changes you've made to the Configuration list. If there are no NIS servers to update for your choice(s), EnlightenDSM will display a dialog box telling you so.

Query



You can use this module to query desired User account information on one or more users depending on your search criteria. You can search for active or inactive accounts, accounts using a specific shell, accounts assigned to specific office locations, and so on. Once your query is successful, you can then modify the accounts, monitor the accounts, print the account configurations, or send mail to the marked users.

When you activate this program, the following window will appear (Figure 3-6).

Figure 3-6 Query User Account Information window

This window gives you the ability to search for information about the user that otherwise might be tedious and time-consuming to gather. The rest of this section describes the functionality of this window's fields and buttons. For additional descriptions of the field names, see [“Add” on page 3-3](#).

Fields

This window contains the following fields:

User names

You can use this field to specify that one or more user names will be used in the search criteria. Multiple user names may be entered, but must be separated by spaces.

Group names

You can use this field to specify that one or more group names will be used in the search criteria. Multiple group names may be entered, but must be separated by spaces.

UID between ... and

You can use these fields to limit the range of the User IDs in the search criteria. This data must be numeric. When EnlightenDSM performs the search, user accounts whose UserID is less than or greater than the respective values are ignored.

GID between ... and

You can use these fields to limit the range of the Group IDs in the search criteria. This data must be numeric. When EnlightenDSM performs the search, user accounts whose GroupID is less than or greater than the respective values are ignored.

Home dirs.

You can use this field to find all users whose home directory matches this search criteria. Multiple directory names may be entered, but must be separated by spaces.

The standard UNIX wild cards '*', '[]', and '?', along with the negation operator '!', can be used in this field (e.g., /home/*).

Shell(s)

You can use this field to find all shell programs that match this search criteria. Multiple shell names may be entered, but must be separated by spaces.

The standard UNIX wild cards '*', '[]', and '?', along with the negation operator '!', can be used as in this field (e.g., /bin/[ck]sh).

Real name

You can use this field to find all users whose Real Name matches this search criterion. Multiple names may be entered, but must be separated by spaces.

Office

You can use this field to find all users whose Office description matches this search criterion. Multiple locations may be entered, but must be separated by spaces.

Telephone

You can use this field to find all users whose telephone number matches this search criterion. Multiple telephone numbers may be entered, but must be separated by spaces.

Logged in since

Every time a user gains access to the system, the system date is recorded. You can use this field to search for users based on the time of their last system login on the local system. If the user's date of last login is earlier than the date in this field, the search will skip the user.

This field can be used in two ways. The first is as a static date, where the date entered is the date used in the search. The second is a relative date, against the system clock. To specify a time format, see [Appendix C, "Time Formats."](#)

Not logged in since

Every time a user gains access to the system, the system date is recorded. You can use this field to specify a cutoff date the search program will use to match users. If the date of last login of a user is later than the date entered in this field, the search will skip the user.

To specify a time format, see [Appendix C, "Time Formats."](#)

Password

You must indicate which Password type is used in the search criteria: Encrypted, Locked, or None. Any combination of one or more of these three attributes may be selected (the default is all three).

Buttons

This window contains the following buttons:

Execute Search

Once you've selected your search criteria, click the Execute Search button. A User Query Results window appears with the results listed. If you click this button without filling in any of the fields, all current user's information will be displayed.

Then you can use the following options to act further on the results:

- Modify
- Delete
- Password
- Lock
- Unlock
- Expire
- View
- Processes

See [“Configure” on page 3-2](#) for more details.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Groups



You can use this module to add, modify, delete, or assign users to a group(s). EnlightenDSM will display the list of currently configured user groups, as shown in [Figure 3-7](#).

The screenshot shows a window titled 'Enlighten/User/Group Configuration'. It has a toolbar with icons for 'Close', 'Mail', 'Refresh', 'Sort', and 'Help'. Below the toolbar are buttons for 'Add', 'Modify', 'Delete', 'Assign Users', and 'NIS'. The main area contains a table with the following columns: 'hostname', 'group name', 'id', and 'members'. The table lists 20 entries for the 'sun-ultra' host. At the bottom left of the table area, it says '94 Entries'.

hostname	group name	id	members
sun-ultra	users	20	0
sun-ultra	root	0	1
sun-ultra	other	1	2
sun-ultra	bin	2	3
sun-ultra	sys	3	5
sun-ultra	adm	4	3
sun-ultra	uucp	5	2
sun-ultra	mail	6	1
sun-ultra	tty	7	3
sun-ultra	lp	8	3
sun-ultra	nuucp	9	2
sun-ultra	staff	10	0
sun-ultra	daemon	12	2
sun-ultra	sysadmin	14	0
sun-ultra	nobody	60001	0
sun-ultra	noaccess	60002	0
sun-ultra	nogroup	65534	0
sun-ultra	esadm	60007	5
sun-ultra	esdomadm	60008	5
sun-ultra	esops	60009	0

Figure 3-7 Currently Configured User Groups

The rest of this section details how to use the functionality of this module.

Add

Click the Add button to create a new user group. A pop-up window will appear prompting you for the Hostnames, Group Name, and Group ID (Figure 3-8).

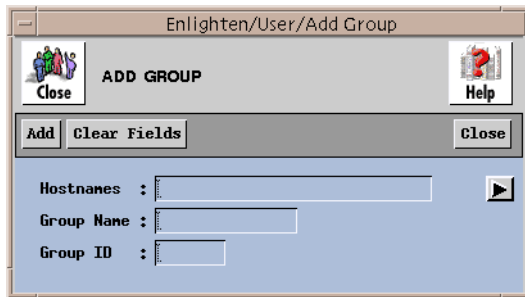


Figure 3-8 New User Group window

Fields

This window contains the following fields:

Hostnames

If you want to limit the creation of this group to specific hostnames within a pool, enter those hostnames in this field. If you are entering multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select the available hosts from the current pool.

Group Name

The Group Name represents how this group will be identified. In most cases the name is related to the activities of the group. For example, the general accounting department group might be known as `acctgen` while the accounting supervisors group may be called `acctsup`. Enter the new User Group Name.

Group ID

The Group ID (or `GID`) is how the operating system refers to the User Group. The Group ID is a numeric way of referring to a User Group, while the Group Name is the more user friendly way of referring to the same Group. Enter the `GID` number. Decimals are not allowed in this field.

Buttons

This window contains the following buttons:

Add

Click the Add button to add the new user group.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Modify

Click this button to modify a group's parameters. A pop-up window similar to the Add Group window will appear, except you cannot modify the Hostnames field. To modify more than one group at a time, you must highlight each group individually before proceeding.

There are two additional buttons in the Modify window:

- Click the Modify button (rather than the Add button) after you've made all your changes, and
- Click the Next button to modify additional user accounts if you've selected more than one account to modify from the User Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 3-28](#).

Delete

Click this button to delete a group definition. Enlighten**DSM** will prompt you to confirm your action. This command will remove the currently highlighted User Group from the system and also delete any links between the User Group and members of the group.

Assign Users

Click this button to add users, delete users, or move users to another group. A window will appear (Figure 3-9).

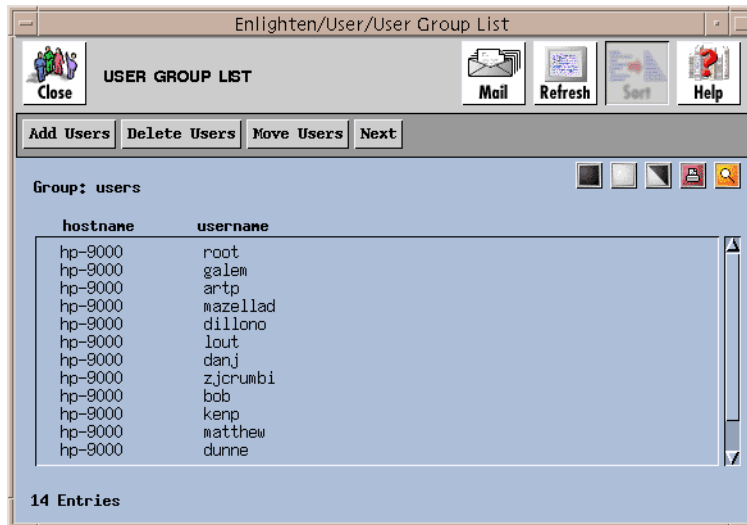


Figure 3-9 User Group List window

The rest of this subsection details how to use these window buttons.

Add Users

Click this button to add users to a group. After you do so, a pop-up menu will be brought up showing all current users on the system. Highlight the users to add to this group and click the Apply button.

Delete Users

Click this button to delete one, some, or all of the users from the currently selected User Group. A pop-up window will appear asking you to confirm this action.

Move Users

Click this button to move one or more User Accounts from the currently selected User Group to another defined User Group. A pop-up window will appear requesting the name of the new user group. You can also click the arrow button on the right to display a pick list of available groups. Highlight the group you wish to move the users to and then click Apply.

When you've selected a new user group, click the Apply button in the New User Group window. Enlighten**DSM** will then prompt you to confirm that each marked user should be moved to the new user group.

Activity Monitor



You can use this module to monitor login activity, process statuses, and CPU usage by user. The actual options are:

- Who is Logged In
- Process Status
- CPU Summary

The rest of this section details how to use each of these options.

Who is Logged In

You can use this program to see at a glance which users are currently logged in (accessing) the system (Figure 3-10). The program displays the following information: Hostname, Username, TTY, Login Time, Idle Time, Process ID, and the location of the TTY (if available).

hostname	username	tty	login time	idle	pid	location
sgi-o2000	adrianc	tttyq0	Dec 15 12:08	.	10185	
sun-ultra	root	dtremote	Dec 14 13:02	3:16	1616	(horton:0)
sun-ultra	root	pts/10	Dec 15 08:41	2:41	18387	
sun-ultra	root	pts/11	Dec 14 14:43	20:37	21351	(:0,0)
sun-ultra	root	pts/12	Dec 14 16:54	18:48	9731	(:0,0)
sun-ultra	root	pts/13	Dec 15 08:51	2:41	19654	
sun-ultra	root	pts/3	Dec 14 14:19	21:47	16131	(horton:0,0)
sun-ultra	root	pts/6	Dec 14 13:12	21:52	3270	(horton:0,0)
sun-ultra	root	pts/7	Dec 14 14:51	19:47	22341	(:0,0)
sun-ultra	root	pts/8	Dec 14 13:38	0:47	7789	(horton:0,0)
sun-ultra	root	console	Dec 14 13:07	.	2627	(:0)
sun-ultra	root	pts/9	Dec 14 13:38	18:44	7791	(horton:0,0)
hp-server	root	console	Dec 15 10:57	1:10	2720	(:0)
hp-server	root	pts/3	Dec 15 10:58	.	2815	(hp-server:0,0)
hp-server	root	pts/4	Dec 15 10:58	1:08	2827	(hp-server:0,0)
hp-9000	root	console	Dec 9 11:04	old	984	
hp-9000	root	ttyp2	Dec 9 11:04	old	1024	
hp-9000	root	ttyp3	Dec 9 11:04	old	1031	hp-9000:0,0
hp-9000	root	ttyp4	Dec 9 11:04	old	1032	hp-9000:0,0
sun-ultra	terryo	pts/14	Dec 15 12:07	.	16677	(hp-server)
hp-9000	terryo	ttyp5	Dec 15 12:23	.	20813	sgi-indy

Total: 21 sessions logged into 4 hosts (3 users)
21 Entries

Figure 3-10 Logged in users

You may select one or more user accounts for further processing by highlighting the desired users. Then select one of the following menu options:

Write

Use the Write command to write a message directly to the highlighted users. When you click the Write button, a window will appear. You can now write a message of any length to each of the highlighted users (one at a time). When you have completed your message, press the <return> key and the message will be sent. To close down the window, press <Control>-C or the interrupt key. The recipient can respond to this message.

Message

The Message command is similar to the Mail command (see Appendix A, “Enlighten**DSM** Basics,” in the *Enlighten**DSM** User Guide*), except a predefined or custom form letter is sent directly to the user’s screen instead of the user’s mailbox. Messages can only be sent to users logged in through the console. The recipient cannot reply to this message.

Logout

This command will terminate all highlighted work sessions by killing the initial Shell process belonging to the marked users.



This command can be dangerous as it may also cause related user processes to be killed. Be careful!

Processes

This command displays a window of all processes currently running that belong to the highlighted users. To view the processes, highlight the users you wish to view and then click the Processes button. A window will appear displaying all processes for the highlighted users. To further manipulate this information, see the next section, “[Process Status.](#)”

Process Status

You can use this menu item to display a list of all active processes (Figure 3-11).

hostname	username	pid	ppid	c	stime	tty	time	command
sun-ultra	root	0	0	0	13:00:23	?	0:00	sched
sun-ultra	root	1	0	0	13:00:26	?	1:53	/etc/init -
sun-ultra	root	2	0	0	13:00:28	?	0:00	pageout
sun-ultra	root	3	0	0	13:00:26	?	0:22	fsflush
sun-ultra	root	341	1	0	13:01:18	?	0:00	/usr/lib/dm1/dm1spd
sun-ultra	root	12061	1	0	17:11:32	?	1:56	/usr/sbin/nscc
sun-ultra	root	523	1	0	13:01:24	console	0:00	/usr/lib/saf/ttymon -g -h -p
sun-ultra	root	104	1	0	13:00:45	?	1:28	/usr/sbin/rpctind
sun-ultra	root	106	1	0	13:00:45	?	0:00	/usr/sbin/keyserv
sun-ultra	root	154	1	0	13:00:51	?	0:00	/usr/sbin/in_named
sun-ultra	root	172	1	0	13:00:58	?	0:00	/usr/lib/nfs/lockd
sun-ultra	root	165	1	0	13:00:58	?	0:00	/usr/sbin/inetd -s
sun-ultra	root	170	1	0	13:00:58	?	0:00	/usr/lib/nfs/statd
sun-ultra	root	139	1	0	13:01:03	?	0:02	/usr/lib/autofs/automountd
sun-ultra	root	193	1	0	13:01:03	?	0:02	/usr/sbin/syslogd
sun-ultra	root	244	1	0	13:01:09	?	0:00	/usr/lib/sendmail -bd -q1h
sun-ultra	root	207	1	0	13:01:07	?	0:00	/usr/sbin/cron
sun-ultra	root	223	1	0	13:01:07	?	0:00	/usr/lib/psched
sun-ultra	root	265	1	0	13:01:10	?	0:00	/usr/sbin/vold

Total: 103 processes (1 user on 1 host)
103 Entries

Figure 3-11 Active processes list

Then you can select a process and use the menu buttons to impact it, as described in the following subsections.

Terminate

This command is very powerful and can be extremely dangerous. It will immediately kill the highlighted process. This command will not kill related processes, so if there are child processes running they will become orphans and will have to be terminated separately. A pop-up window will prompt you for verification to terminate the process.

Hangup

This command is similar to the Terminate command, except it provides enough time for the process to shut down properly. This means the process can close any files and terminate any child processes. A pop-up window will prompt you for verification to hang up the process.

Suspend

This command stops a process from working, but it does not terminate the process. Essentially, this command puts a process on hold; it can be activated again at a later time. You must use the Continue command to re-activate a suspended process. A pop-up window will prompt you for verification to suspend the process.

Continue

This command allows you to re-activate a process that was previously put on hold by a Suspend command. A pop-up window will prompt you for verification to resume the process.

Priority

This command allows you to change the priority of a process. This priority determines when the CPU acts on a process. It may have a value from -19 to $+19$; the smaller the number, the higher the priority. Clicking the Priority button activates a Process Priority window. You can enter the desired priority or use the arrow buttons to make your selection.

CPU Summary

This program allows you to view a breakdown of CPU usage by user. The Summary of Process window will show all currently logged in users, the current number of processes, and the total cumulative CPU usage for each active user ([Figure 3-12](#)).

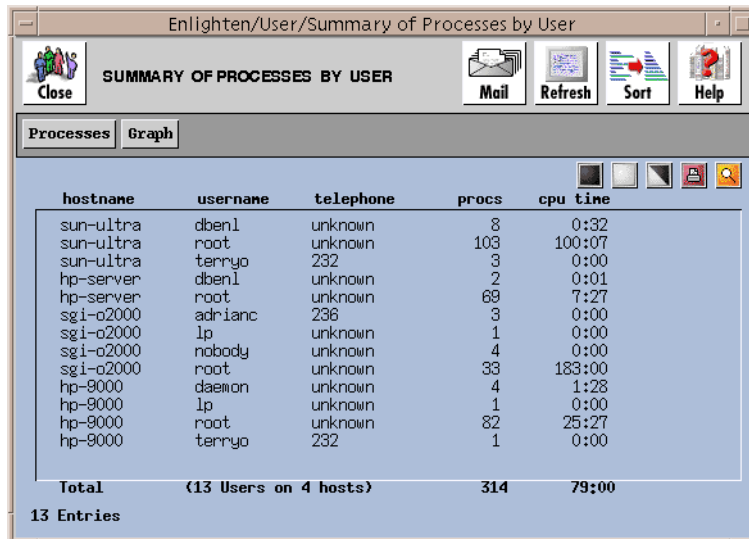


Figure 3-12 CPU Summary window

You now have the option to Graph all or selected processes, or to view individual processes.

Graph

To graph the processes, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

Processes

To view the processes, highlight the users you wish to view and then click the Processes button. A window will appear displaying all processes for the highlighted users. To further manipulate this information, see [“Process Status” on page 3-34](#).

New User Templates



You can use User account templates to facilitate the creation of new user accounts.

Each template specifies the defaults EnlightenDSM will use when it adds a new user to the system. If you leave a field blank when you're adding a new user account, EnlightenDSM will substitute the default parameters set up for that field.

Choose New User Templates from the User menu to display the New User Template Configuration window (Figure 3-13).

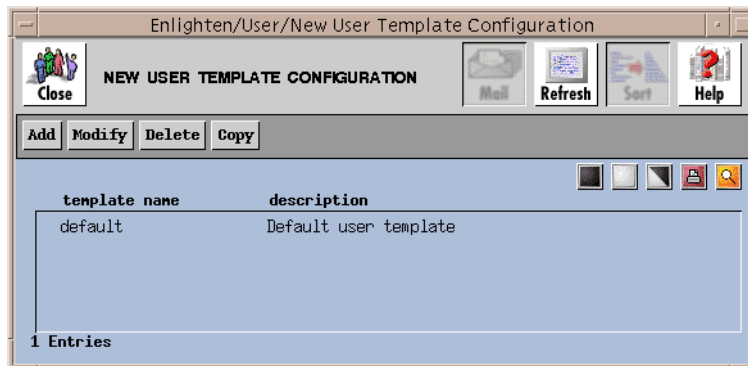


Figure 3-13 New User Template Configuration window

From here, you have the option to:

- **Add** Create a new template. No fields are initialized
- **Modify** Modify the highlighted template
- **Delete** Delete the highlighted template
- **Copy** Create a second template using the values in the selected template as a starting point

Add

Click this button to create a new template. The New User Template Add window will appear (Figure 3-14).

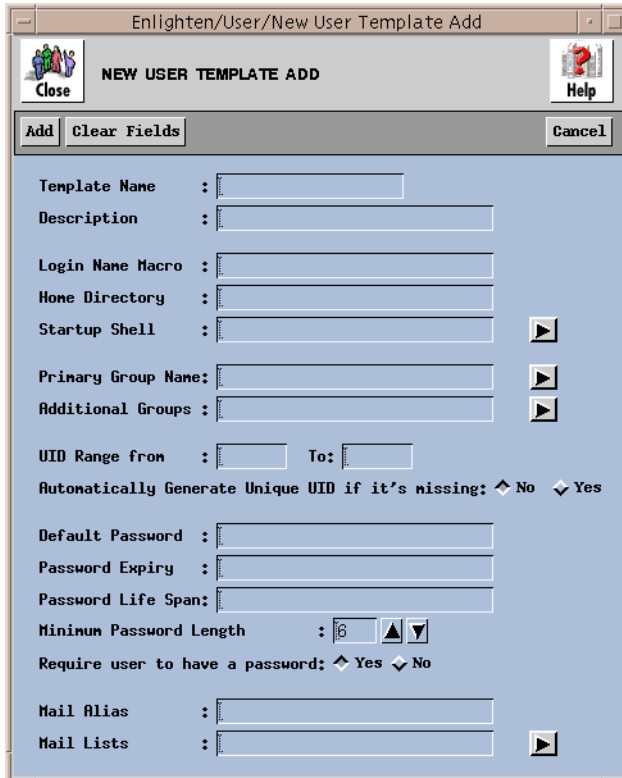


Figure 3-14 New User Template Add window

Fields

The window has the following fields:

Template Name

Use this field to specify the template's name.

Description

Use this field to briefly describe this template's purpose. This is used by other commands, such as Session Preferences and Adding Users, when you want a pop-up list of available templates displayed.

Login Name Macro

Use this field to specify a macro for the user name. See [“User Account Macros” on page K-2](#) for more information.

Home directory

The Home directory field specifies where the user’s start-up directory is attached. Most systems will use `/usr` or `/home` as the default parent directory. Each user will have their own directory created under the default home directory; for example, `/usr/charlie`. EnlightenDSM stores `/usr` as the default HOME directory unless you change this field. If the default Home Directory is `/eng`, and the login name of the new user account is `fred`, the name of the HOME directory EnlightenDSM generates would be:

```
/eng/fred
```

You can also use macros in this field. See [“User Account Macros” on page K-2](#) for more information.

Startup Shell

The Shell is a program that acts as a translator between the user and the UNIX operating system. There may be several standard shells on a system, including the C Shell, the Bourne Shell, the Korn Shell and the restricted Bourne, C, and Korn Shells. Each shell has its own start-up executable program as shown below.

Bourne Shell	<code>/bin/sh</code>
Bourne Shell	<code>/bin/sh</code>
Restricted Bourne Shell	<code>/bin/rsh</code>
C Shell	<code>/bin/csh</code>
Korn Shell	<code>/bin/ksh</code>
Restricted Korn Shell	<code>/bin/rksh</code>
Other	user defined

EnlightenDSM uses the Bourne shell as the default shell in this field (`/bin/sh`). The number of standard shells available for this field depends on your system. Click the arrow button to display a pick list of available shells.

Primary Group Name

Every user account must be assigned to at least one User Group. The User Group is part of the UNIX security system; each user group has specific read, write, and execute privileges associated with every file created on the system.

Enter the Primary Group Name for this template. You can also click the arrow button for a pick list of defined User Groups and make your selection from there. If you enter a Primary Group Name that does not exist, EnlightenDSM will prompt for a correct Group Name when it tries to create the user.

Additional Groups

You can use this field to put the user in additional user groups. This will effectively give the user more group privileges beyond those set in the primary user group. Enter the names of the user groups to which the user should also belong. If you are using multiple entries, leave a blank between each entry.

You can also click the arrow button for a pick list of defined User Groups. Select the User Groups and click the Apply button to import them.

UID Range from... To

You can use these fields to define a UID range for the template. This range is used if you set the next field to have a UID automatically generated for the user.

Automatically Generate Unique UID if it's missing

Each user account created must have a UserID number. Use this toggle to choose whether a unique UID is generated for a new user if you don't specify one during creation of the new user account. The default is not to (No).

Default Password

The password is part of the UNIX security system. Each user must have a password. If you make the password the same as the user login, the user can then reset the password to one of his or her own choice. Don't use passwords, such as spouse names, pet names, or addresses, that are easy for another user to guess.

You can also use macros in this field. See [“User Account Macros” on page K-2](#) for more information.



If the entry in the password field is not long enough, an error message will be displayed when you try to add this user account.

Password Expiry

For systems that support password aging, this optional field allows you to set the time when the password will expire. Once a password has expired, the system will force the user to change their password the next time they log in to the system.

If this field is left blank, the appropriate default parameter will be used. If no default is set in the User Add window (when you're adding a user), password aging will be turned off for that user. For more details, refer to [Appendix D, “Password Aging.”](#)

To specify a time format, see [Appendix C, “Time Formats.”](#)

Password Life Span

In conjunction with the previous field, this field defines the minimum period of time that needs to elapse before the user can change his or her password. To specify a time format, see [Appendix C, “Time Formats.”](#)



You can change the user's password at any time by using the Password button in the User Configuration window. See [“Password” on page 3-18](#) for more details.

Minimum Password Length

Each user account has a password associated with it. This password is defined when the user account is created and may be changed any time thereafter. This field accepts a numeric entry defining the minimum length that the (new) password must be whenever the password is created or modified. You can also use the counter buttons on the right to increment or decrement the number shown.



If you set this value to zero, the user does not need to use a password when logging in to the system; however, the user will be asked by the system to enter a password. The user can then log in by hitting the <return> key.

Require user to have a password

It may be necessary for a user to have no password. By default, EnlightenDSM forces each user to have a password. To configure user accounts with no password, set this toggle to No.

Mail Alias

This field specifies the default mail alias the account will use. An alias designates a short name as the substitute for the full pathname for a particular user or group, for example, `laura` for `laura_lombardo@desk.com`.

You can also use macros in this field. See [“User Account Macros” on page K-2](#) for more information.

Mail Lists

This field specifies which mailing lists will be updated to include this user account. A mail list defines a group of people under one alias, so sending mail to that alias will distribute a copy to everyone on that list. If you are using multiple entries, leave a blank between each entry. You can also click the arrow button for a pick list of defined Mail Lists and make your selection(s) from there.

Buttons

The New User Template Add window has the following buttons:

Add

Once you've specified all the values for the new user template, click this button to save it. Then you can access this template when you are setting your session preferences (see [“Session Preferences” on page 2-2](#)) or adding users (see [“Add” on page 3-3](#)).

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click this button to modify a template's values. A pop-up window similar to the New User Template Add window will appear, except the Template Name field is view-only.

There are also two button differences in the Modify window:

- Click the Modify button (rather than the Add button) after you've made all your changes, or
- Click the Next button to modify additional templates if you've selected more than one to modify from the template list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 3-38](#).

Delete

Click this button to delete a template from the template list. EnlightenDSM will prompt you to confirm your action.

Copy

Click this button to copy the set of values in the selected template to a second template. The User New Template Add window will appear showing the highlighted template's values in each of the fields. You can edit this window as needed and then click the Add button to complete the copy.

See [“Add” on page 3-38](#) for a description of how to use this window's fields and buttons.

User Management Scripts

These scripts allow you to customize the actions taken when creating, modifying or deleting user accounts.

Adding a User

`adduser.rc`. XenIn calls this script when it creates a user, after the password entry is added. Because this script uses environment variables (instead of command line options like `user.rc` used to), the information does not have to be renumbered when one is added or deleted.

Supported variables:

`pw_createhomedir` copies the `CREATEHOMEDIR` params file entry. If the params file entry is missing, then `pw_creathomedir` is set to "1".

`pw_realname`, `pw_telephone`, `pw_shell`, `pw_uid`, `pw_gid` and `pw_lname` correspond to Real name, Telephone, Shell, User ID, Group name, and Login name, respectively.

`pw_home` corresponds to Home directory. This field will have the `$HOMEDIR` macro expansion done with the `HOMEDIRLOCAL` params file entry (that is, `$HOMEDIR` in the GUI field will be replaced by `HOMEDIRLOCAL` in this parameter, but with `HOMEDIRPASSWD` when the entry is being written to the password file).

if the home directory field of the GUI contains "`$HOMEDIR`", then that string will be replaced by the params file entry `HOMEDIRPASSWD` when the password file entry is written. "`$HOMEDIR`" will be replaced by the params file entry `HOMEDIRLOCAL`, and the result will be passed to `adduser.rc` as the `pw_home` environment variable.

These scripts can be used to set up automounter maps automatically for new users by modifying the commented examples in the scripts. The example assumes that the `$HOMEDIR` substitution sets `/home/username` in the password map and sets `$pw_home` to `/host/realhomedir`. The `sed` commands strip out the host and the real home directory into separate variables, then create the automount mapping. Finally, the example creates a home directory as `user.rc` used to do.

Modifying a User

`moduser.rc` is similar to `adduser.rc`, except that it also passes in the following variables:

```
pw_old_uid
pw_old_gid
pw_old_shell
pw_old_home
```

These variables contain the old UID, GID, shell and home directory.

It is left to the user to check these variables against their non-old (that is, check whether `pw_gid` is the same as `pw_old_gid`) counterparts to detect changes. For example, a user may do a recursive `chown/chgrp` if the UID or GID is all that changed, or perhaps a recursive copy if the home directory changed.

No indication of the old contents of the GECOS (full name, office, telephone) field is given. If none of the above 4 variables indicates any change, then this script does nothing.

Deleting a User

There are two scripts to handle deleting a user. After the `passwd` file entry is removed, `deluser.rc` runs. Then a dialog box asks if the home directory should be removed. Deleting the home directory is done through `delhome.rc`.

`deluser.rc` deletes a user. It takes one parameter, `pw_lname`, the username to delete. This script is called AFTER the `passwd` file entry has been removed, so the login name is likely the only information that is available.

`delhome.rc` deletes the home directory. It takes 3 parameters: `pw_home`, which is the old `passwd` file home directory entry, `pw_homedirlocal`, which is the related params file entry, and `pw_homedirpasswd`, which is also the related params file entry.



Printer

The Printer menu is used to set up local and network printers, manage print jobs, define custom printer drivers, and configure TTY ports.

The Printer menu options are:

- Configure (all)
- Configure (local)
- Queue
- Custom Driver
- TTY Configure

Configure (all) / Configure (local)



These two programs pop up a window displaying all the currently configured printers (Figure 4-1).

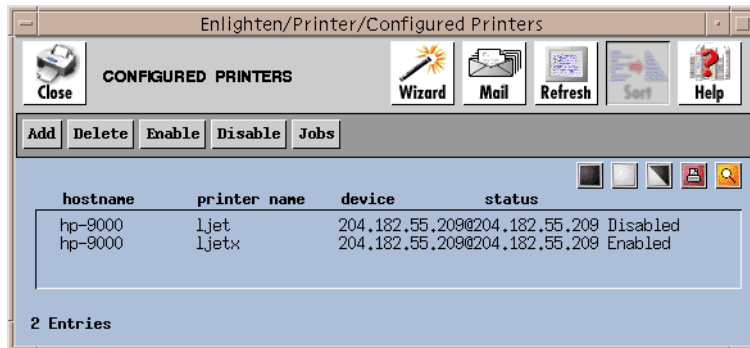


Figure 4-1 Configured Printers window

The “Configure (local)” option is similar to the “Configure (all)” option, except it reports only local printer destinations. The Configure (local) option does not display network printers. This can be useful to gain a better overview of printer resources.

From the Configured Printers window, you have the option to:

- Add Add a printer
- Delete Delete a printer
- Enable Enable a printer
- Disable Disable a printer
- Jobs View and modify print jobs

Add

You can define a printer as a print server or a print client. Altogether, there are three options for configuring a printer when you click the Add button: as a Print Client, a Local Print Server, or a Network Print Server.

Each of these windows contains buttons and fields that are common to both print client and print server configurations.

Common Buttons

Apply

Click this button to save the configuration specified.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Common Field

Print Service Type

This set of toggle buttons specifies whether to configure a printer client or printer server. The setting you choose will determine the look of the rest of the window.

The following subsections detail the window and field formats for each option.

Print Client

This section describes configuring printers as print clients. To do so, select the Client option from the Add a New Print Service window (Figure 4-2).

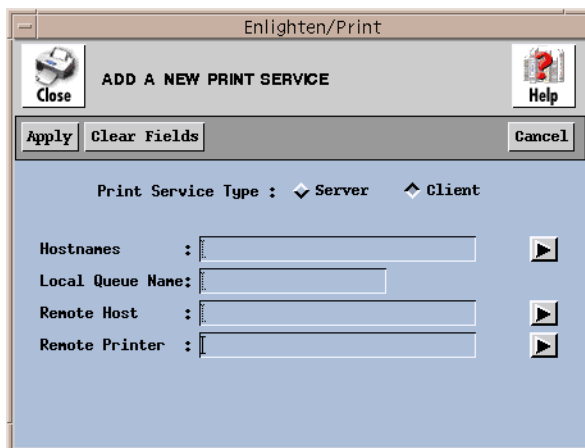


Figure 4-2 Add a New Print Service window

The following subsections describe the format for the rest of the fields in the Add a New Print Service window (Add Print Client option).

Hostname

Type the hostname to configure the new printer on a subset of the currently managed hosts. Leave a blank space between each hostname if you are specifying multiple entries. You can also click the arrow button to the right of the field to bring up a pick list of all currently managed hosts and make your selections from there.

Local Queue Name

Type the full pathname for the local printer queue.

Remote Host

Type the hostname you want to use as the print server for this particular client. You can also click the arrow button to the right of the field to display a list of all hosts in the current pool.

Remote Printer

Type the remote printer you want to use to service the print jobs. Click the arrow button to the right of the field to display a list of all known printers defined on hosts in the current pool.

Local Print Server

This section describes configuring printers as local print servers. To do so, you must select the Server and Local options from the Add a New Print Service window ([Figure 4-3](#)).

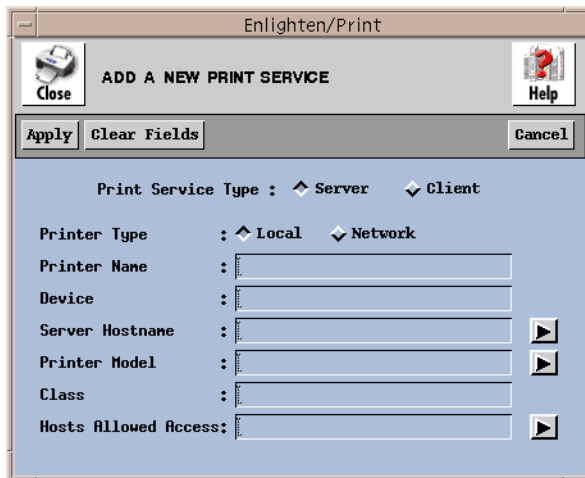


Figure 4-3 Add a New Print Service window

The following subsections describe the formats for the rest of the fields in this Print Service window (Add Local Print Server option).

Printer Name

Type the printer name users will use to refer to the printer. Use a meaningful name, such as the printer's location, type, or purpose. This name is also referred to as a printer *queue* or *destination*. These expressions can be used interchangeably.

This field is mandatory; enter the name of the new printer.

Device

Type the device, such as `/dev/lp` (parallel) or `/dev/ttyh7` (serial), where the printer will be attached. This device must already exist.

Server Hostname

Type the hostname from the current pool that you want to use as the print server for this printer. The arrow button to the right of the field displays a list of all hosts in the current pool.

Printer Model

Specify the type of printer driver you want to use. You must enter a hostname in the Server Hostname field before selecting a driver. The arrow button to the right of the field displays a listing of all the available printer models.

This field is mandatory. See [“Custom Driver” on page 4-11](#) for more information on creating your own printer drivers.



When you add a printer to the system, it is automatically enabled.

Class

For systems using the System V type spooling system, you can better distribute print jobs on multiple printers of the same type. When you enter a printer in the Class field, the print job is sent to the first available printer in the printer class.

This field is an optional field. If you don't enter a printer class, then the printer (queue) will not be put into a printer class. If you enter a printer class that does not exist, the class will automatically be created.

This field is not used by systems using the Berkeley style printer spooler (lpd) or by IBM AIX systems.

Hosts Allowed Access

This field lets you specify which hosts are allowed access to the print server. The arrow button to the right of the field displays a list of all hosts in the current pool.

Network Print Server

To configure printers as network print servers, select the Server and Network options from the Add a New Print Service window ([Figure 4-4](#)).

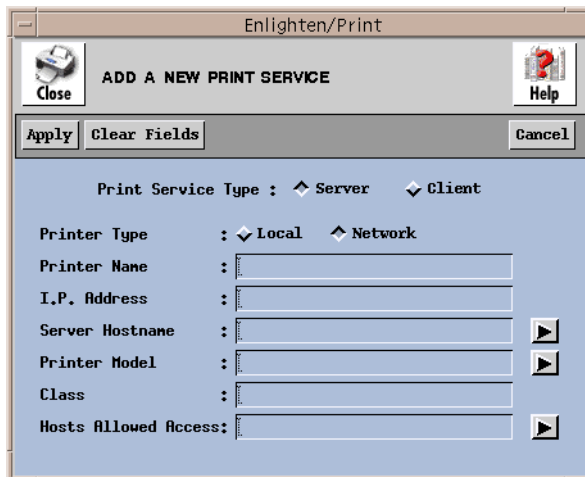


Figure 4-4 Add a New Print Service window

These fields are the same as the ones for the Local Print Server, except you need to specify an I.P. Address for the network printer rather than a Device name for a local printer. See [“Local Print Server” on page 4-4](#) for a description of the rest of the fields in this Print Service window.

Delete

To delete a printer from the system, click the Delete button. From the Configure Printers window, highlight the printers you want to delete from the system. EnlightenDSM will prompt you to confirm your choice in order to ensure the selection was correct.



EnlightenDSM cannot delete a printer configuration if there are pending print requests. Use the ‘Printer → Queue’ menu item to cancel all pending print requests on the printer you want to delete.

Enable

Click the Enable button to enable a printer that has been disabled. The Enable command is similar to physically turning the printer on; it tells the print queue administrator daemon that the printer is now available for printing. Any print jobs that might have been waiting in that print queue will now begin to print.

Disable

Click the Disable button to stop the highlighted printers from receiving print requests from the spooler. Any job that is currently printing will be added back to the print queue to be reprinted once the printer is enabled. This command does not stop users from sending print jobs to the disabled printer; those print requests will be added to the print queue.

Jobs

Click the Jobs button to display all print jobs currently printing or waiting to be printed for the highlighted printer destinations. A pop-up window will notify you if there are no jobs. If there are print jobs, the Printer Queue window will appear. This window and its features are described in [“Queue” on page 4-9](#).

Queue



This program displays a list of all outstanding print requests (jobs). EnlightenDSM allows you to issue specific commands, such as canceling print requests and enabling or disabling a printer. The Printer Queue window (Figure 4-5) appears only if there are pending print jobs. If not, a message window appears stating that there are no pending jobs.

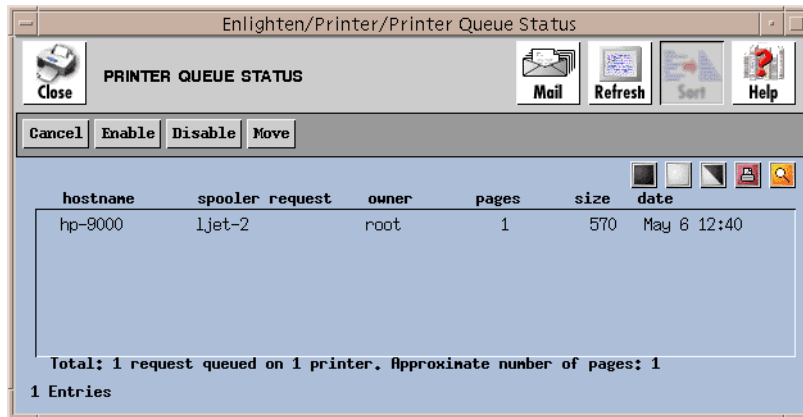


Figure 4-5 Printer Queue Status window

The following sections detail the operations of the buttons within the Queue Status window.

Cancel

Click the Cancel button to delete the selected print job from the system. If the job is currently printing, the Cancel command will remove the print request from the system; however, the printer may continue printing until its buffer is empty.

Enable

The Enable button is similar to turning the printer on. It tells the print queue administrator that the printer is now available for printing. Any print jobs that might have been waiting in that print queue will now begin to print.

Disable

Click the Disable button to temporarily stop printing. The printer queue will still accept print jobs. Typically, you use this command when the printer is out

of paper or will only be out of service for a short while. When a printer is disabled, its print queue is kept intact.

Disabling stops the marked printers from receiving print requests from the spooler. Any job currently printing will be added back to the print queue to be reprinted when you re-enable the printer. This command does not stop users from sending print requests to this printer; those print requests will be added to the print queue.



On Sun O.S. 4.1.x, or other systems on which the Berkeley style printing system is running, the disabling of a printer will cause it to decline print requests until it is re-enabled. Also, all printers on the same device are enabled/disabled even if only one is highlighted.

Move

Click the Move button to move print jobs from one printer queue to another if, for example, a printer breaks down or has many jobs queued. EnlightenDSM displays a window showing a list of all available printer destinations. Select the new printer destination (queue).

EnlightenDSM will then ask for confirmation to move each of the highlighted print jobs to the other printer. If you click the Yes button, the print job will be moved to the selected print queue.

You can move print jobs to print queues on other hosts even if the new destination host has a different spooling system running.



When moving a print job, you should select a new printer destination compatible with the old printer destination. Failing to do so may produce unexpected results.

Custom Driver



You can use this program to define custom printer models or drivers. Each type of printer has its own characteristics; for example, a laser printer needs to map the print commands differently than a dot matrix style of printer. Use the Custom Driver window to create different types of printer drivers.

The following sections detail the fields and their formats for the options shown in the Create Custom Printer Model window (Figure 4-6).

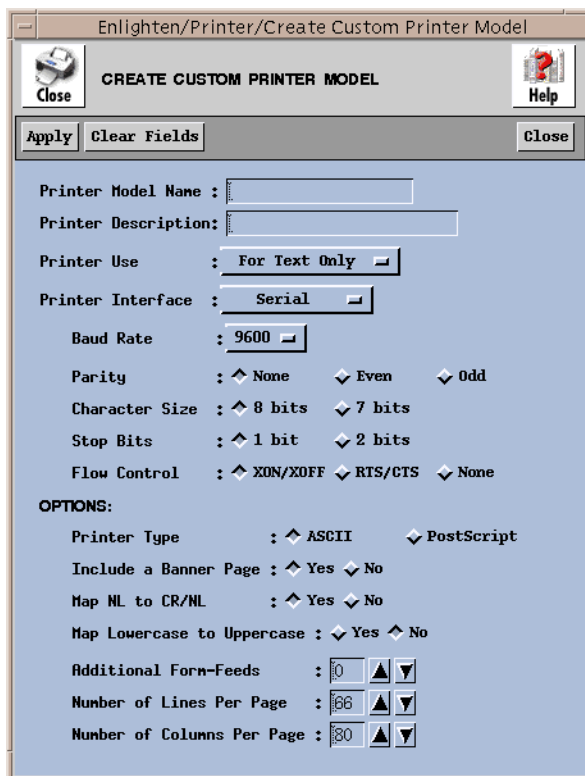


Figure 4-6 Create Custom Printer Model window

Fields

This window contains the following fields:

Printer Model Name

Type a unique name that refers to this particular printer driver. This name is usually related to the model of printer, for example a “Raven PR 9101” might be called `rav9101`.

Printer Description

You can optionally enter a brief description of the driver in this field.

Printer Use

This field allows you to specify how the printer will be used. If you want to use the printer for printing text fields and graphic designs or charts, the interface will have to include a graphic module. If you want to use the printer only for text, then only the text module is required.

Using the left mouse button, click the option button to display the following list of options:

- For Text Only (the default)
- Text or Graphics
- Other

If data sent to the printer is already formatted for that printer, choose the Text or Graphics option. This is typical of applications such as word processors or drawing programs where the printer type can be selected at print time. The application prepares its output for the chosen printer.

If, however, you want to print only plain text, then choose the For Text Only option. If you have to manipulate the text before it is sent to the printer, you can use the `OPTIONS` section to specify pre-printing manipulations. See the following subsections for more details.

Printer Interface

The Printer Interface option allows you to choose the type of hardware interface that is needed for the printer to work properly. Different types of printers require different types of interfaces. The interface may be:

- Serial (the default)
- Parallel
- Data Products
- Other

The most common of these is Serial or Parallel. Using the left mouse button, click the appropriate printer interface. Refer to your printer manual for more information about what interface is required.

If the printer interface selected is `Serial`, then additional information will be required to properly set up the serial interface. The information required is:

- Baud rate—The speed at which the printer will send and receive data. Click the appropriate baud rate.
- Parity—Select the parity the printer will be using. The options are Odd, Even, or None.
- Character Size—The size of the byte the printer uses. The two options are 7 bits or 8 bits.
- Stop Bits—The stop bits setting for the printer. The two options are 1 bit or 2 bits.
- Flow Control—This is how the printer recognizes signals from the operating system. The flow control tells the system if the printer is ready to receive information or if it is busy printing. This field is hardware dependent on the serial boards installed in the machine. Most serial printers use Xon/Xoff flow control.



All of the flow control options are dependent on the serial board's capabilities. Not all flow control options are supported on all serial (TTY) boards.

Please check your printer manual for complete information regarding the serial interface settings.

Options

You can use the options area for any pre-printing manipulation this custom printer driver needs.

- **Printer type**

Click the PostScript button if the printer is running PostScript emulation. Otherwise, click the ASCII button.
- **Include a banner page**

To identify individual printer jobs, you can add an optional banner page by clicking the Yes button.
- **Map NL to CR/NL**

For some printer emulations, such as HPCL, the carriage return (CR) character must be mapped to CR/NL (New line) for proper output of lines of text. For these emulators, click the Yes button.
- **Map lowercase to uppercase**

If the printer only prints in uppercase mode (typically band printers) or if you want all uppercase output, click the Yes button.
- **Additional form-feeds**

When the print job is finished, the printer driver will output up to five additional form feeds to advance the paper. The default value is zero. Enter or select a new value as desired.
- **Number of lines per page**

Enter the number of lines of text for a single page. This information is used in determining when a page ends. The default value is 66. Enter or select a new value as needed.
- **Number of columns per page**

Enter the number of characters to fit onto a single line of text. This information is used in determining where to wrap text to a new line. The default value is 80. Enter or select a new value as required.

Buttons

This window contains the following buttons:

Apply

When you have completed filling in this window, click the Apply button. A window appears confirming the creation of the new printer model. Then you can assign it to a device in the Printer Configuration window.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Pop down the window without performing any other action.

TTY Configure



You can use this program to display a list of all configured ports on the local UI machine ([Figure 4-7](#)). You can add new ports and modify, enable, or disable existing ports. You can also access port information by port number or TTY name.

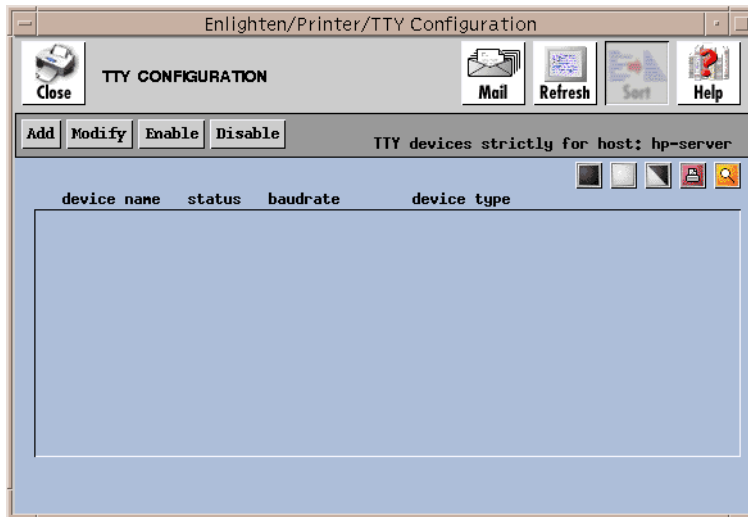


Figure 4-7 TTY Configuration window

From here, you have the option to:

- Add Add a new device to the system
- Modify Modify an existing TTY port
- Enable Enable an existing TTY port
- Disable Disable an existing TTY port

Add

Click the Add button to quickly and easily add a TTY port. See [Figure 4-8](#).

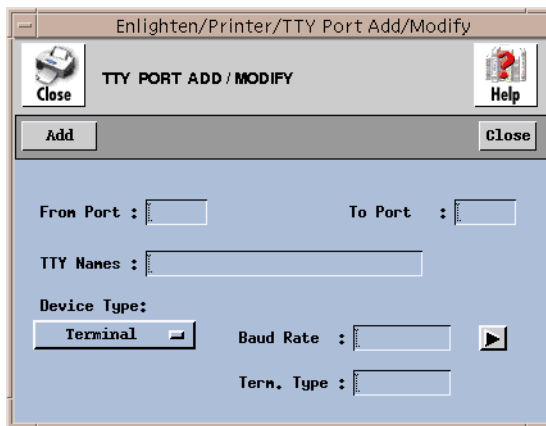


Figure 4-8 Adding a TTY port

Fields

This window contains the following fields:

From Port

Use this field to specify the first port number to be configured when you're defining new TTYs. If the TTY is already defined, you can leave this field blank and enter the TTY name field instead. The sequence for port numbering always starts at 0 (zero), not 1 (one). For example, if there are 16 ports on the machine, the port numbers will be 0 to 15.

Check with your hardware vendor for more information on the number of ports and their setup.

To Port

Use this field to specify the last port number to be configured when you're defining new TTYs. If you are configuring only one TTY, you can leave this field blank.

TTY Names

You can use this field to specify the exact TTY name of a port. Every port on a system has a port number and a corresponding TTY name. If you're configuring new TTYs, leave this field blank and the system will assign the TTY name to each port as it is defined.

If the previous two fields (From Port and To Port) are blank, you must specify the TTY name(s) in this field. Leave a blank space between each TTY name for multiple entries.

Device Type

Use this field to choose the type of device that will be added. Click the left mouse button and a list of available device types will appear. You can choose one of the four standard device types:

- Terminal (the default)
- Serial Printer
- Dialin Modem
- Dialout Modem



Because each device type has unique characteristics, it is important that you select the correct device type. If the wrong device type is chosen, the port may not react as expected.

Baud Rate

The baud rate is the speed at which the port will send and receive signals. Click the arrow button on the right to display a pick list of the baud rates and select one of them. EnlightenDSM will insert your selection into this field.

Term. Type

When a user accesses the system from this port, the system will set the terminal type to whatever you specify here. Each terminal has its own set of characteristics, for example, an IBM 3151 terminal emulation is not the same as a VT100 emulation. If the terminal type is unknown, specify the vt100 type.



It is important that you select the correct terminal type. If the wrong emulation is chosen, the screen display will not work properly.

Buttons

This window contains the following buttons:

Add

Click the Add button to add your TTY selection(s).

Close

Click this button to discard any changes and close the window.

Modify

Click this button to modify a TTY port. The TTY Port Add/Modify window appears, as shown in [Figure 4-8](#). The first three fields in this window are not available, though; the only fields available for modification are the Device Type, Baud Rate, and Term. Type fields. Once you've made all your changes, click the Modify button to change the TTY port settings.

Enable

When you enable a port, you are telling the operating system this port will be used to log in to the system. Since this will allow users to log in the system via the named port, make sure you do not enable a printer or dialout port.

Highlight the desired device names you wish to enable and click the Enable button. A pop-up window will appear asking for confirmation. Click the appropriate button.



Do not enable a port that has been designated as a serial printer port. The system will try to send the login sequence to the printer and all sorts of strange things will happen.

Disable

If a port is disabled, the operating system will not offer users the ability to log in via that port. Printer and dialout ports are already disabled when configured. You can selectively disable other ports to manage which ports provide access to the system.

Highlight the desired device names you wish to disable and click the Disable button. A pop-up window will appear asking for confirmation. Click the appropriate button.



Disabling a port where a user is currently logged in will *not* log out the user. The user will be able to continue working until he or she does log out. After all users are logged out of the port, it is then disabled from further logins.



Network

The Network menu is used to establish and update DNS and NIS servers. It is also used to configure and manage NFS servers.

The Network menu options are:

- Host Entries
- Host Configuration
- Host Naming Services
- Naming Service Overrides
- DNS
- NIS and NIS+
- NFS

Host Entries



Each system has a host database containing a list of known hosts and their TCP/IP address on the network. You can use this module to concurrently manage the host database of multiple machines. When you choose this command, the Host Entry Configuration window will appear (Figure 5-1).

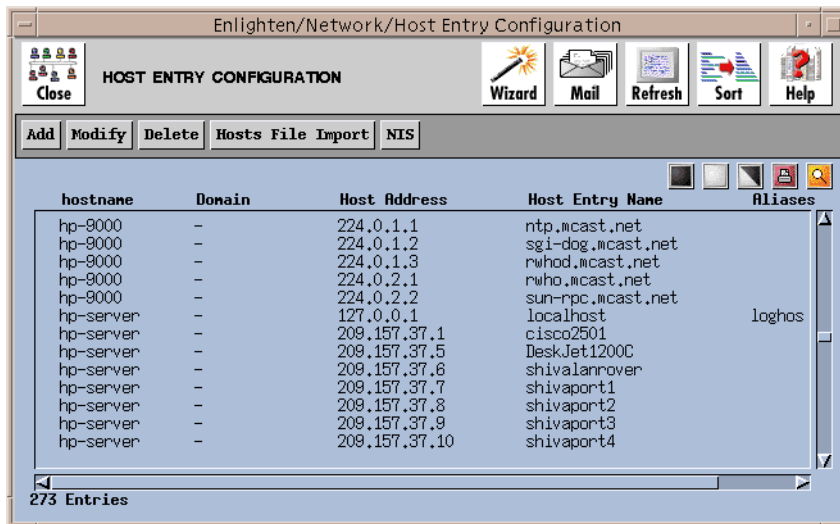


Figure 5-1 Host Entry Configuration window

The list box will show all known host entries in the current pool, their domains, TCP/IP addresses, host entry names, and any aliases. From here, you can:

- Add Add new host entries
- Modify Modify existing host entries
- Delete Delete existing host entries
- Hosts File Import Import a file defining the host entries for any active hostname servers
- NIS Push the NIS maps on active NIS servers

The rest of this section details how to use this window's buttons.

Add

Click the Add button to bring up the Add Host Entry window (Figure 5-2).




Figure 5-2 Add Host Entry window

Then you can create new host entries by entering the name of the new host, its TCP/IP address, and any host alias names it has.

Fields

This window contains the following fields:

On hosts

You can use this field to limit which hostname server(s) will be configured. Leave a blank space between hostnames for multiple entries. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool and make your selection(s) from there.

Hostname

The hostname should be unique for the server where you're adding this host. Use this field to specify the host's name.

Address

This field is initialized with the base address of the current system. You can select a different network from the one shown or simply add the last part of the address.

Aliases

You can use this field if you want to specify any alias(es) this host uses. Leave a blank space between alias names for multiple entries.

Buttons

This window contains the following buttons:

Add

After you've entered all the information, click the Add button. EnlightenDSM will create the new host entries on the hostname servers for all hosts in the current pool. If you specified an incorrect address, or if the address is already being used, EnlightenDSM will display an appropriate error message telling you so.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

If you wish to change the address of a host or modify the alias list of a host, highlight the desired entries and click the Modify button. A window similar to the Add Host Entry window will appear, except you cannot modify the Hostname field.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional host entry configurations if you've selected more than one to modify from the Host Entry Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-3](#).

Delete

Click the Delete button to delete the highlighted host entries. EnlightenDSM will prompt you to confirm your action.

Hosts File Import

If you want to load the host entries from an existing file into the current host naming services server(s) (or this host database from a local file you've already created), click this button to import that file. The Import Host Entries File window appears (Figure 5-3).




Figure 5-3 Import Host Entries File window

Fields

This window contains the following fields:

Hostname

You can use this field to limit which server(s) should be configured when importing this file. Leave a blank space between hostnames for multiple entries. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool and make your selection(s).



The Host File Import feature interprets a fully qualified host name as a separate entry from the non-qualified equivalent. Duplicate entries may result.

Local File

Use this field to specify the full pathname for the local file containing the host entry information.

Buttons

This window contains the following buttons:

Import

After specifying your choices, click this button to import your local host entries file to the selected host's database(s).

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

NIS

Click this button to push/remake the NIS maps on any active NIS servers. This updates the NIS servers and hosts on your network with any Host Entry changes you've made to the Configuration list. If there are no NIS servers to update for your choice(s), EnlightenDSM will display a dialog box telling you so.

Host Configuration



You can use this module to view network parameters such as the I.P. address, netmask, and broadcast address for a particular host. When you choose this command, the Network Host Configuration window appears (Figure 5-4).

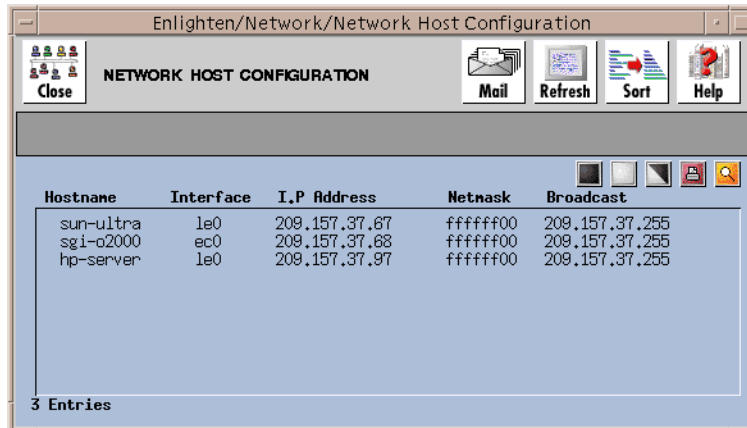


Figure 5-4 Network Host Configuration window

Host Naming Services



You can use this module to configure which NIS, NIS+, and/or DNS servers a client host should use. When you use this module, the Host Naming Services Configuration window will appear (Figure 5-5).

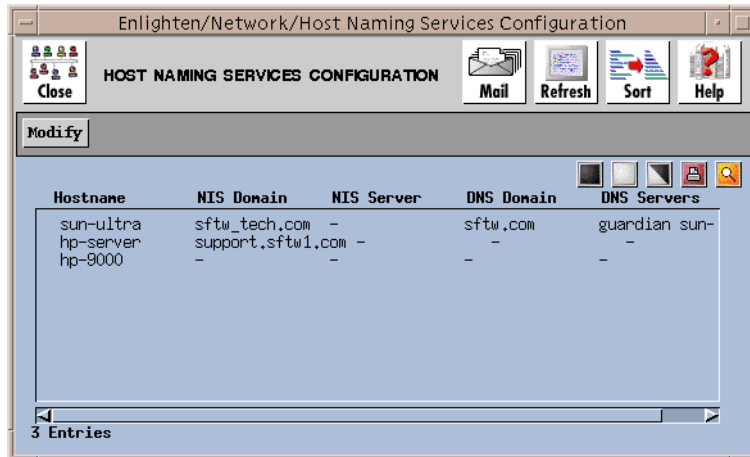


Figure 5-5 Host Naming Services Configuration window

Highlight which host-server configuration(s) you want to modify and click the Modify button to bring up the Modify Host Naming Services window (Figure 5-6).



DNS clients display Host Entry information from the DNS database files, regardless of the name service lookup order specified in a host's `/etc/resolv.conf` file.

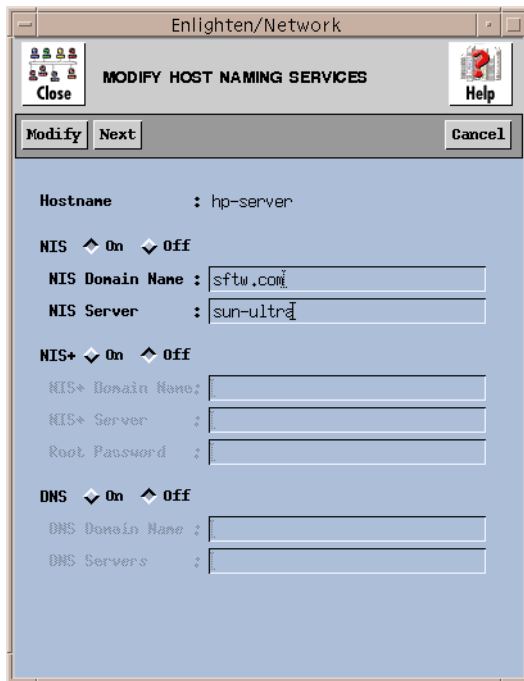


Figure 5-6 Modify Host Naming Services window

The rest of this section describes how to modify the combination of your host-server configuration(s) parameters.

Modify Fields

This window contains the following fields:

Hostname

This view-only field shows which host's naming service configuration is currently being modified.

NIS

You can use these toggles to specify if you're using the NIS naming service for this configuration. If you turn this field on, the NIS+ fields will be automatically turned off; these two naming services are mutually exclusive for any particular host.

These fields are only available if the NIS toggle is set to on.

NIS Domain Name

Use this field to specify the domain name for your NIS server.

NIS Server

Use this field to specify which NIS Server this host will use.

NIS+

You can use these toggles to specify if you're using the NIS+ naming service for this configuration. If you turn this field on, the NIS fields will be automatically turned off; these two naming services are mutually exclusive for any particular host.

These fields are only available if the NIS+ toggle is set to on.

NIS+ Domain Name

Use this field to specify the domain name for your NIS+ server.

NIS+ Server

Use this field to specify which NIS Server this host will use.

Root Password

Use this field to specify the root password for the server shown in the NIS+ Server field.

DNS

You can use these toggles to specify if you're using the DNS naming service for this configuration. These fields are only available if the DNS toggle is set to on.

DNS Domain Name

Use this field to specify the domain name for your DNS server.

DNS Servers

Use this field to specify which DNS Servers this host will use.

Modify Buttons

This window contains the following buttons:

Modify

Click this button to add the specified combination of NIS and/or DNS host-server connections. Enlighten**DSM** will prompt you to confirm your action, especially if you're disabling NIS, NIS+, and/or DNS.

Next

Click the Next button to view and/or modify additional host-server configurations if you've selected more than one to modify. No changes will be made to the configuration you're currently viewing.

Cancel

Click this button to close the window without making any changes.

Naming Service Overrides



EnlightenDSM auto-detects the naming services any particular host uses. You can use this module to create and update any naming service overrides an individual host or pool should use. You can establish overrides for password and group files, host entry files, and mail alias files within this option.

When you choose this command, the Naming Services Overrides window will appear ([Figure 5-7](#)).

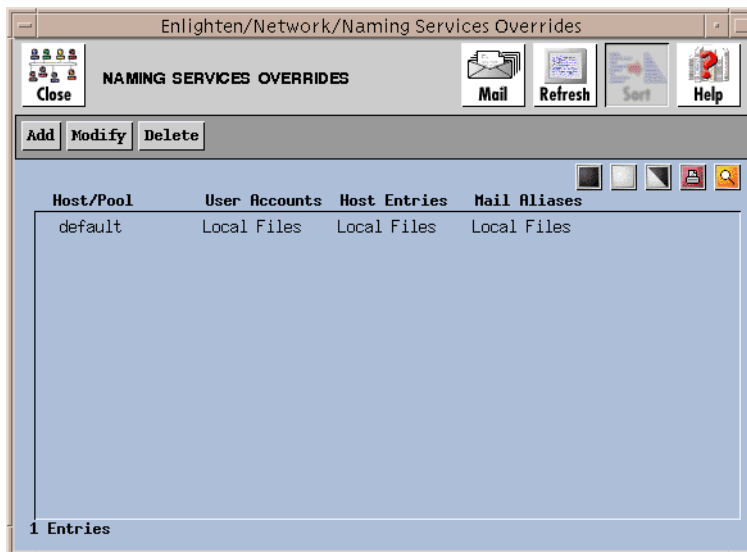


Figure 5-7 Naming Services Overrides window

The list box will show all host/pools that already have overrides assigned and the type of user accounts, host entry, and mail aliases overrides. From here, you can:

- **Add** Create a new set of overrides for a host/pool
- **Modify** Modify an existing set of overrides for a host/pool
- **Delete** Remove the set of overrides for a host/pool

The rest of this section details how to use this window's buttons.

Add

Click the Add button to bring up the Create Naming Service Override window (Figure 5-8).

The screenshot shows a window titled "Enlighten/Network/Create Naming Service Override". The window has a toolbar with "Close", "Help", "Add", "Clear Fields", and "Cancel" buttons. The main content area is divided into three sections:

- Host/Pool Name:** A single text input field.
- Password/Group Overrides:** A "None" dropdown menu, followed by fields for "Domain Name", "Master Host", "Password File", and "Group File".
- Host Entry Overrides:** A "None" dropdown menu, followed by fields for "Domain Name", "Master Host", "Hosts File", and "Named Boot".
- Mail Alias Overrides:** A "None" dropdown menu, followed by fields for "Domain Name", "Master Host", and "Aliases File".

Figure 5-8 Create Naming Service Override window

Then you can create a new set of overrides for a host/pool by filling in the form as needed. The rest of this section describes this window's fields and buttons.

Fields

This window contains the following fields:

Host/Pool Name

Use this field to specify which host or pool will use these overrides.

Password/Group Overrides

You can use this toggle to set up the overrides for user accounts. Use this if you want to use an alternate account naming service rather than those determined by EnlightenDSM for your host/pool. You can use overrides from the following types of naming services:

- None (the default; no overrides are used in this area)
- local files (/etc/... files)
- NIS
- NIS+

Domain Name

If you've selected the NIS or NIS+ option in the Password/Group Overrides field, use this field to specify the Masterhost domain name.

Master Host

If you've selected the NIS or NIS+ option in the Password/Group Overrides field, use this field to specify the NIS server providing the user account and user group services.

Password File

If you've selected the NIS option in the Password/Group Overrides field and you want to use a different password file (besides `/etc/passwd`), use this field to specify that password file pathname.

Group File

If you've selected the NIS option in the Password/Group Overrides field and you want to use a different group file (besides `/etc/group`), use this field to specify that group file pathname.

Host Entries Overrides

You can use this toggle to set up the overrides for host entries. Use this if you want to use an alternate host naming service rather than those determined by EnlightenDSM for your host/pool. You can use overrides from the following types of naming services:

- None (the default; no overrides are used in this area)
- local files (`/etc/hosts` files)
- NIS
- NIS+
- DNS



Overriding a NIS master to a second NIS master will result in the local `/etc` files of the second master being used for password, group, alias and host information.

Domain Name

If you've selected the NIS+ or DNS option in the Host Entry Overrides field, use this field to specify the Masterhost domain name.

Masterhost

If you've selected the NIS, NIS+, or DNS option in the Host Entry Overrides field, use this field to specify the DNS server providing the host entry services.

Hosts file

If you've selected the NIS option in the Host Entry Overrides field, use this field to specify an alternate to the `/etc/hosts` file.

Named boot

If you've selected the DNS option in the Host Entry Overrides field, use this field to specify an alternate to the `/etc/named.boot` file. This file is used by the BIND daemon to locate the DNS configuration files at system start-up.

Mail Alias Overrides

You can use this toggle to set up the overrides for mail aliases. Use this if you want to use an alternate aliases service rather than those determined by EnlightenDSM for your host/pool. You can use overrides from the following types of naming services:

- None (the default; no overrides are used in this area)
- local files (`/etc/aliases` files)
- NIS
- NIS+

Domain Name

If you've selected the NIS+ option in the Mail Aliases Overrides field, use this field to specify the Masterhost's domain name.

Masterhost

If you've selected the NIS or NIS+ option in the Mail Aliases Overrides field, use this field to specify the NIS server name providing the aliases services.

Aliases file

If you've selected the NIS option in the Mail Aliases Overrides field and you want to use a different set of aliases, use this field to specify that alias file pathname.

Buttons

This window contains the following buttons:

Add

After you've entered all the relevant information, click the Add button. EnlightenDSM will create your specified overrides for the designated host or pool.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click the Modify button to modify an existing set of overrides for a host/pool. A window similar to the Create Naming Services Override window will appear, except you cannot modify the Host/pool name field.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional sets of naming services overrides if you've selected more than one to modify from the Naming Services Overrides list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-13](#).

Delete

Click this button to remove the selected set of overrides for a host/pool. Enlighten**DSM** will prompt you to confirm your action.

DNS



You can use this module to configure and manage Domain Name Servers (DNS) on a network. You can configure Primary Domain Name Servers, Secondary Domain Name Servers, Mail Exchangers, and Hosts in a Domain. The options are:

- Server Configuration
- Mail Exchangers

The rest of this section details how to use these options.

Server Configuration

You can use this program to configure primary and secondary domain name servers. When you activate this program, the window shown in [Figure 5-9](#) appears.

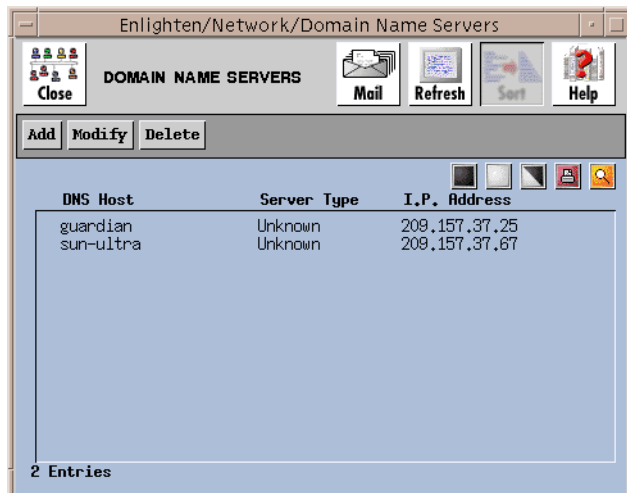


Figure 5-9 Domain Name Servers window

The list box will display all known Domain Name Server hostnames, the server type (either primary or secondary), and the host's I.P. Address. A primary DNS server is displayed in the Domain Name Servers window only if one or more hosts in the current pool are clients of it.

From here, you have the option to:

- Add Add a Domain Name Server to your network
- Modify Modify an existing DNS configuration
- Delete Delete all selected DNS configurations

The rest of this section details how to use this window's buttons.

Add

Click the Add button to add a Domain Name Server to your network. The Add Domain Name Server window will appear ([Figure 5-10](#)).

Figure 5-10 Add Domain Name Server window

You can configure both primary and secondary Domain Name Servers. If you configure both and the primary server is busy or down, EnlightenmentDSM will utilize the secondary server. Use the Server Type field to select which type

you want to configure. [Figure 5-10](#) shows the fields for configuring a primary server (the default setting).

The rest of this subsection describes the fields for adding a primary domain name server. See “[Secondary Server Add Fields](#)” on page 5-21 for a description of the fields used when adding a secondary server.

Primary Server Add Fields

This window contains the following fields:

Server Type

Use this toggle button to specify whether a primary or secondary DNS server is being configured. The remaining fields in this window assume you selected the Primary option.

Domain Name

Use this field to specify the domain name of the new DNS server.

Primary Server Hostname

Use this field to enter which host will be the primary domain name server. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool and select one.

Network I.P. Numbers

You can use this field to enter the network I.P. Address of the subnets the DNS server will manage within the domain. For example, if you’re managing hosts with the address range 192.57.36.01 through 192.57.36.254, then you should enter the value `192.57.36` in this field.

Dir. for DNS Database Files

Use this field to specify the full directory pathname for the DNS database files. The default is `/etc/named`.

Pathname to DNS Config. Files

Use this field to enter the full pathname for the `named.boot` file. This file is used by the BIND daemon to locate the DNS configuration files at system start-up. The default is `/etc/named.boot`.

Secondary Name Servers

You can use this field to specify the hostnames of any secondary name servers. Leave a blank space between server names for multiple entries.

Administrator e-mail Address

Use this field to specify the DNS administrator's email address.

Database Serial Number

Use this field to enter the serial number associated with the DNS database. The default serial number is 1 (one). This serial number is incremented every time the primary database is altered.

DNS Database Refresh Frequency

You can use this field to choose how often the secondary DNS database serial numbers are checked against the primary DNS database numbers. If the serial numbers are different, the secondary DNS updates itself. The default value is every three hours.

DNS Query Retry Frequency

You can use this field to choose how frequently the secondary DNS server will attempt to reconnect with the primary DNS server if the initial contact with the primary server failed. The default value is every hour.

DNS Data Expiry

You can use this field to choose the maximum time the hostname and IP address information will live in a secondary server when the Primary server is unreachable. The default value is one week.

Minimum TTL

You can use this field to choose the maximum time a cached hostname and IP address will live in a DNS server before the secondary server contacts the primary server to resynchronize the data. The default value is one day.

Secondary Server Add Fields

This subsection describes the fields for adding a secondary domain name server. See [“Primary Server Add Fields” on page 5-20](#) for a description of the fields used when adding a primary server.

The default Add Domain Name Server window is shown in [Figure 5-11](#).

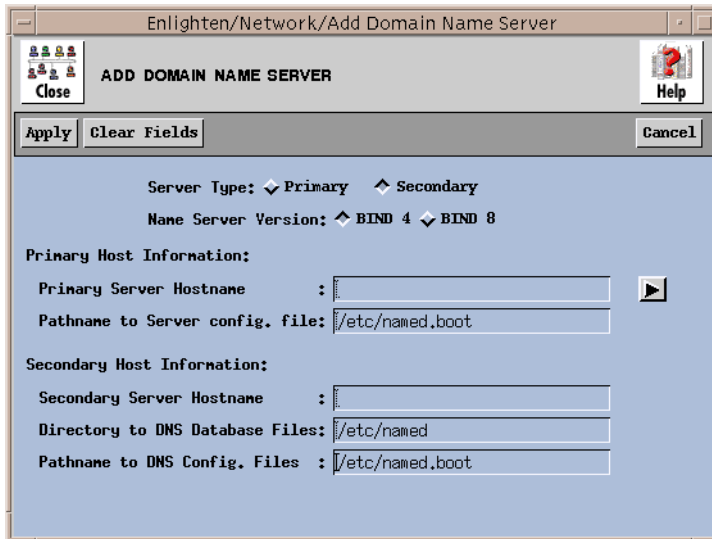


Figure 5-11 Add Domain Name Server window

This window contains the following fields:

Server Type

Use this toggle button to specify whether a primary or secondary DNS server is being configured. The remaining fields in this window assume you selected the Secondary option.

Primary Server Hostname

Use this field to enter the hostname of the primary domain name server for which the secondary server is being configured. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool and select one.

Pathname to Server config. file

Use this field to enter the full pathname for the `named.boot` file. This file is used by the BIND daemon to locate the server configuration files at system start-up. The default is `/etc/named.boot`.

Secondary Server Hostname

You can use this field to specify the name of one or more hosts on which to configure secondary servers. Leave a blank space between server names for multiple entries.

Directory to DNS Database Files

Use this field to specify the full directory pathname for the DNS database files. The default is `/etc/named`.

Pathname to DNS Config. Files

Use this field to enter the full pathname for the `named.boot` file. This file is used by the BIND daemon to locate the DNS configuration files at system start-up. The default is `/etc/named.boot`.

Add Buttons (both types of servers)

This window contains the following buttons:

Apply

Click this button to begin the configuration of the new DNS server.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click the Modify button to modify an existing DNS configuration. The Modify Domain Name Server window is similar to the Add Domain Name Server window, except you cannot modify the Domain Name, Primary Server Hostname, Network IP Numbers, Dir. for DNS database files, and Pathname to DNS config. files fields.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and

- You can use the Next button to modify additional DNS configurations if you've selected more than one to modify from the Domain Name Servers list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-19](#).

Delete

Click the Delete button to disable the DNS server on the selected host(s). EnlightenDSM will prompt you to confirm your action.

Mail Exchangers

You can use this program to configure mail exchange information. When you activate this program, the window in [Figure 5-12](#) will appear.

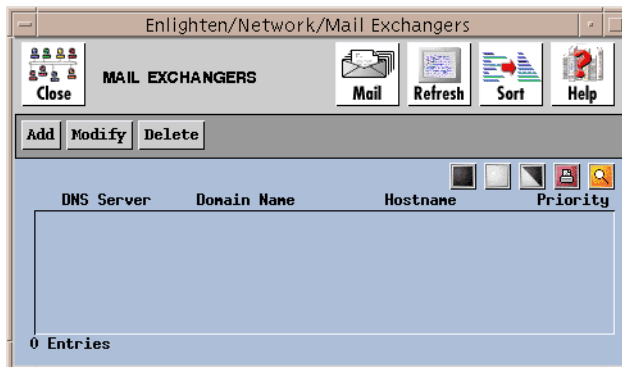


Figure 5-12 Configuring mail exchange information

The list box will display the DNS Server name, Domain name, Hostname, and priority for all configured routes.

From here, you have the option to:

- Add Add a Mail Exchanger host to your network
- Modify Modify an existing Mail Exchanger configuration
- Delete Delete all selected Mail Exchanger configurations

The rest of this section details how to use this window's buttons.

Add

Click the Add button to add a Mail Exchanger host to your network. The Add Mail Exchanger Host window will appear ([Figure 5-13](#)).

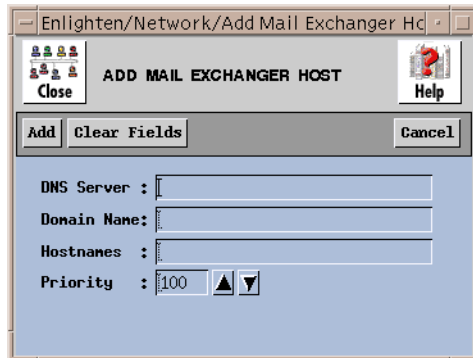


Figure 5-13 Add Mail Exchanger Host window

The rest of this section details how to use this window's functionality.

Add Fields

This window contains the following fields:

DNS Server

You can use this field to specify the DNS Server for this mail exchanger. If you leave this field blank, it defaults to the default DNS Servers value set in Network —> Host Naming Services. Refer to [“Host Naming Services” on page 5-8](#) to change this value.

Domain Name

You can use this field to specify the Internet domain name for this mail exchanger. If you leave this field blank, it defaults to the default domain name value set in Network —> Host Naming Services. See [“Host Naming Services” on page 5-8](#) for more information on changing this value.

Hostnames

Use this field to specify which host machine(s) will act as mail host, or mail dispatcher, for the machines in the Domain Name field. Leave a blank space between hostnames for multiple entries.

Priority

You can use this field to associate a priority, or preference, to using this hostname as a mail host. The number can range from 0 to 65536 inclusive. The lower the number, the higher the priority. You can also use the counter buttons to increment or decrement this number.

Add Buttons

This window contains the following buttons:

Add

Click this button to save the current information as a mail exchanger for the Internet domain set in the Domain Name field.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click the Modify button to modify an existing mail exchanger configuration. The Modify Mail Exchanger Host window is similar to the Add Mail Exchanger Host window, except you cannot modify the DNS Server and Domain Name fields.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional mail exchanger configurations if you've selected more than one to modify from the Mail Exchangers list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-25](#).

Delete

Click the Delete button to remove the selected mail exchanger configurations from the Mail Exchangers list. Enlighten**DSM** will prompt you to confirm your action.

NIS and NIS+



You can use this module to configure Network Information Servers NIS and NIS+ servers. You can configure master servers (NIS and NIS+), slave servers (NIS), and replica servers (NIS+). When you choose this command, the Network Information Servers window will appear ([Figure 5-14](#)).

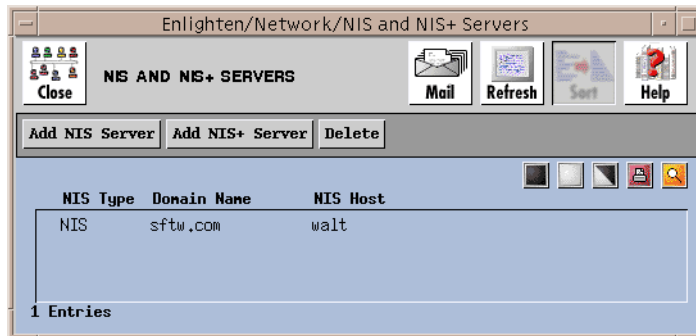


Figure 5-14 Configuring servers

The list box will display all configured NIS and NIS+ servers. Each line in the list will show the type, domain, and hostname for each server. An NIS server is displayed in the NIS and NIS+ Servers window only if one or more hosts in the current pool are clients of it.

From here, you have the option to:

- Add NIS Server Add an NIS server or slave to your network
- Add NIS+ Server Add an NIS+ server or replica server to your network
- Delete Delete all selected NIS configurations



You must first disable a host as an NIS server or client before reconfiguring that host as an NIS+ server or client.

The rest of this section details how to use this window's buttons.

Add NIS Server

Click this button to configure NIS master servers or NIS slave servers. The Add NIS Server window will appear ([Figure 5-15](#)).

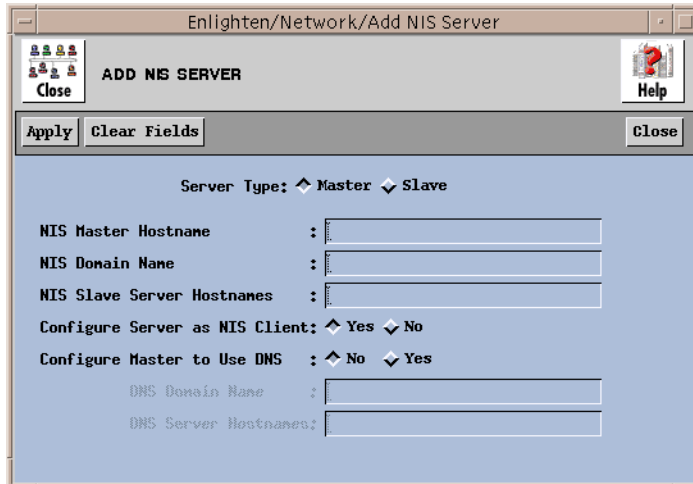


Figure 5-15 Add NIS Server window

The rest of this section details how to use this window’s functionality.

Master Fields

This window has the following fields:

Server Type

Use this toggle to specify what type of NIS server to configure, a Master (the default) or Slave server. If you set this to be a Slave server, the window will be redrawn to show the fields listed in [“Slave Fields” on page 5-30](#).

NIS Master Hostname

Use this field to specify the hostname of the NIS master server.

NIS Domain Name

Use this field to specify the domain name for this server.

NIS Slave Server Hostnames

You can use this field to specify any NIS slave servers for this domain.

Configure Server as NIS Client

Use this toggle to specify if this server should also be configured as an NIS client. The default is Yes.

Configure Master to Use DNS

Use this toggle to specify if the NIS master server should use DNS as a fallback. The default is No.

DNS Domain Name

You can only use this field if you set the Configure Master to use DNS field to Yes. If so, you can use this field to specify the domain name of the DNS servers the NIS master should reference.

DNS Server Hostnames

You can only use this field if you set the Configure Master to use DNS field to Yes. If so, you can use this field to specify the hostnames of the DNS servers the NIS master should reference.

Slave Fields

This window has the following fields:

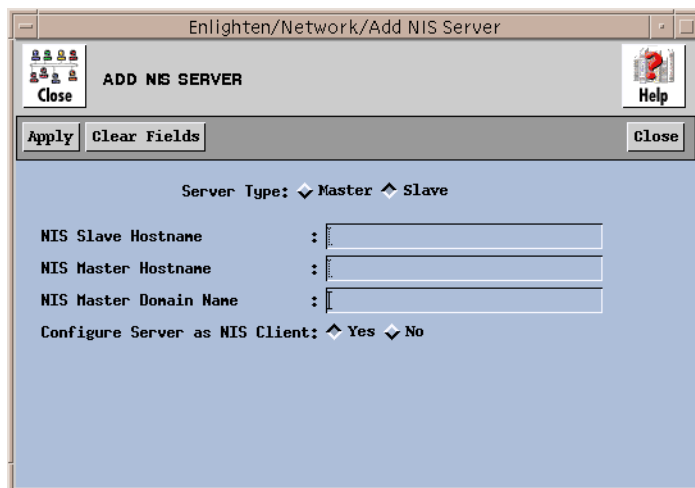


Figure 5-16 Add NIS server

Server Type

Use this toggle to specify what type of NIS server to configure, a Master (the default) or Slave server. If you set this to be a Slave server, the window will be redrawn to show the fields listed in [“Master Fields” on page 5-29](#).

NIS Slave Hostname

Use this field to specify the hostname of the NIS slave server.

NIS Master Hostname

Use this field to specify the hostname of the NIS master server.

NIS Master Domain Name

Use this field to specify the domain name for the NIS master server.

Configure Server as NIS Client

Use this toggle to specify if this server should also be configured as an NIS client. The default is Yes.

Buttons (both types of NIS servers)

This window contains the following buttons:

Apply

Click this button to add the new NIS server shown in the window.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Add NIS+ Server

Click this button to configure NIS+ master servers or NIS+ replica servers. The Add NIS+ Server window will appear ([Figure 5-17](#)).

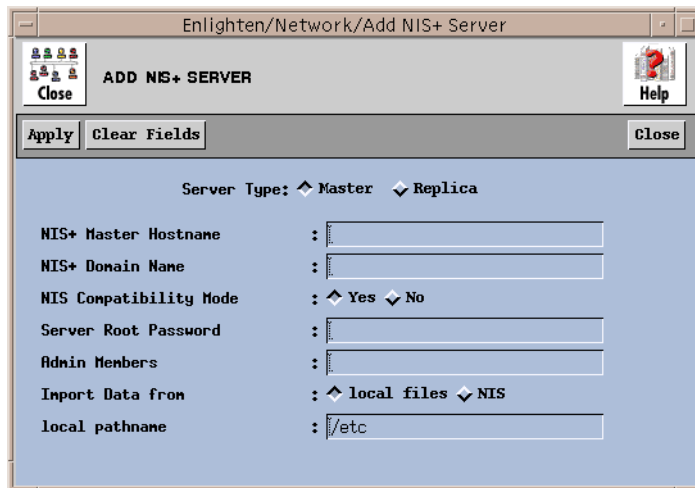


Figure 5-17 Add NIS+ server

The rest of this section details how to use this window's functionality.

Master Fields

This window has the following fields:

Server Type

Use this toggle to specify what type of NIS+ server to configure, a Master (the default) or Replica server. If you set this to be a Replica server, the window will be redrawn to show the fields listed in [“Replica Fields” on page 5-33](#).

NIS+ Master Hostname

Use this field to specify the hostname of the NIS+ master server.

NIS+ Domain Name

Use this field to specify the domain name for the NIS+ master server.

NIS Compatibility Mode

Use this toggle to specify if this NIS+ should be made backward compatible with NIS. The default is Yes.

Server Root Password

Use this field to specify the root password for the NIS+ master server.



NIS+ server setup will fail if the server root password contains spaces.

Admin Members

Use this field to specify which users are allowed to administer this NIS+ host. Leave a blank space between names for multiple entries.

Import Data from

Use this toggle to specify if any existing imported data should come from the servers' `/etc` files or from an NIS server. The default is `/etc`.

local pathname

You can only use this field if you set the Import Data from field to `/etc files`. If so, you can use this field to specify these files' pathname.

NIS Domain

You can only use this field if you set the Import Data from field to NIS. If so, you can use this field to specify the domain name of the NIS server from which to import the data.

NIS Server

You can only use this field if you set the Import Data from field to NIS. If so, you can use this field to specify the hostname of the NIS server.

Replica Fields

This window is similar to the Add NIS+ Server window for a Master server shown on [page 5-32](#). It contains the following fields:

Server Type

Use this toggle to specify what type of NIS+ server to configure, a Master (the default) or Replica server. If you set this to be a Master server, the window will be redrawn to show the fields listed in [“Master Fields” on page 5-32](#).

NIS+ Master Hostname

Use this field to specify the hostname of the NIS+ master server.

NIS+ Domain Name

Use this field to specify the domain name for the NIS+ master server.

Replica Hostname

Use this field to specify the hostname of the NIS+ replica server.

Replica ServerRoot Password

Use this field to specify the root password for the NIS+ replica server.

NIS Compatibility Mode

Use this toggle to specify if this NIS+ server can also interact with NIS servers. The default is Yes.

Buttons (both types of NIS+ servers)

This window contains the following buttons:

Apply

Click this button to add the new NIS server shown in the window.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Delete

Click this button to delete a highlighted server configuration from the NIS and NIS+ Servers list. The Delete NIS/NIS+ Server window will appear (Figure 5-18).



Enlighten**DSM** does NOT automatically detect if a server is a master or slave/replica server.




Figure 5-18 Delete NIS and NIS+ Servers

The rest of this section details how to use this window's functionality.

Fields

This window has the following fields:

NIS Type

Use this toggle to specify what type of server to delete, an NIS (the default) or NIS+ server.

Server Type

Use this toggle to specify what type of server to delete, a Master (the default) or Slave/Replica server.

Domain Name

Use this field to specify the domain name for the server.

Server

Use this field to specify the hostname of the server.

Master Host

You can only use this field if you set the Server Type field to Slave/ Replica. If so, you can use this field to specify the hostname of the NIS/ NIS+ master server.

Remain as NIS Client

You can only use this field if you set the Server Type field to Slave/ Replica. If so, you can use this toggle to specify if this server should remain as an NIS client after deletion. The default is Yes.

Buttons

This window contains the following buttons:

Delete

Click this button to remove the specified server configuration from the NIS and NIS+ Servers list. EnlightenDSM will prompt you to confirm your action.

Close

Click this button to discard any changes and close the window.

- **Modify** Modify the mount parameters for an existing NFS configuration
- **Delete** Remove the marked networked directories from the Network Filesystem Configuration list
- **Copy** Copy the selected mount parameters and create a new partition to be mounted
- **Mount** Mount an NFS partition immediately
- **Unmount** Unmount the marked NFS partition(s)
- **Processes** View the processes for the marked filesystem(s)

Add

Use this command to add an NFS partition on one or more hosts. You can mount this partition immediately, at boot time, or both. Click the Add button to bring up the Add an NFS Partition window (Figure 5-20).

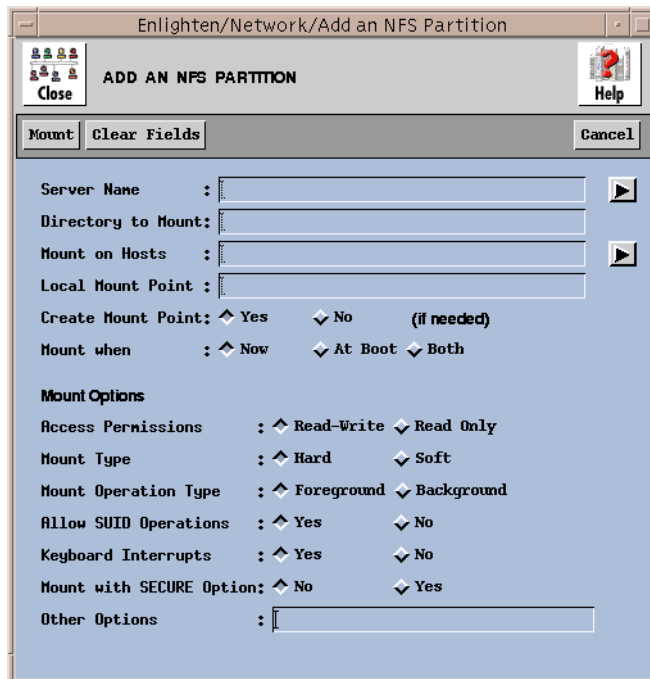


Figure 5-20 The NFS Partition window

Add Fields

This window has the following fields:

Server Name

Use this field to specify the NFS server from which the partition will be mounted. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool that may have partitions that you can mount.

Directory to Mount

Use this field to specify the remote directory to be mounted. Leave a blank space between directory names for multiple entries.



EnlightenDSM does not support the double slash syntax (e.g. //filesystem).

Mount on Hosts

Use this field to specify the hosts where the directory will be mounted. Leave a blank space between host names for multiple entries. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool that may mount the partition specified above and make your selection(s) from there.

Local Mount Point

Use this field to specify the directory name where the exported partition will be mounted.

Create Mount Point

Use this toggle to specify if the local mount point directory should be created during a mount if that directory does not already exist. The default is Yes.

Mount when

Use these toggles to specify when the mount should take place. The options are Now (the default), On Reboot, or Both.

Access Permissions

Use this toggle to specify how the remote partition should be mounted: Read-Write (the default) or Read Only.

Mount Type

Use this toggle to specify if the mount should be a Hard (the default) or Soft mount. Filesystems that are mounted read-write or that contain executable files should be mounted with the hard setting. Applications using soft-mounted filesystems may produce unexpected I/O errors. Soft-mounted filesystems return errors on request; hard-mounted filesystems display a warning message and continue to try the request.

Mount Operation Type

Use this toggle to specify whether the mount should be performed in the Foreground (the default) or Background.

Allow SUID Operations

Use this toggle to specify if setuid execution is allowed for this mount, but silently ignores requests to set the SUID bit.

Keyboard Interrupts

Use this toggle to specify if keyboard interrupts are allowed to kill a process that is hung up while waiting for a response on a hard-mounted filesystem. The default is Yes.

Mount with SECURE Option

Use this toggle to specify whether to mount with the authentication of RPC option. The default is No.

Other Options

You can use this field to specify other NFS options. This field is not prechecked for syntax errors.

Add Buttons

This window has the following buttons:

Mount

Click this button to initiate the mount defined in the window.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click the Modify button to modify the mount parameters for an existing NFS configuration. You can do so only if the partition is mounted automatically at boot time (see the list box entry).

The Modify An NFS Partition window is similar to the Add An NFS Partition window, except you cannot modify the Server Name, Directory to mount, or Mounted on Host fields.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional NFS partitions if you've selected more than one to modify from the Network Filesystem Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-39](#).

Delete

Click the Delete button to remove directories from the Network Filesystem Configuration list. EnlightenDSM will prompt you to confirm your action.

Copy

Use this command to copy the selected mount parameters and create a new partition to be mounted. Select the desired partition and click the Copy button. The Add An NFS Partition window will appear showing the highlighted partition's settings in each of the fields.

You can edit this window as needed and then click the Mount button to initiate the new mount defined in the window.

See [“Add” on page 5-39](#) for a description of how to use this window's fields and buttons.

Mount

Click this button to mount an NFS partition immediately. The Mount An NFS Partition window will appear. This window is similar to the Add An NFS Partition window, except you do not have a Mount when field here.

For a description of the buttons and fields in this window, see [“Add” on page 5-39](#).

Unmount

Click this button to unmount the highlighted NFS partition(s). A pop-up window will appear asking you to confirm this action.

Processes

Click this button to bring up the Processes by Filesystem window. Only processes being run on the selected mounted partitions will be displayed. You can use this option to view any processes before forcefully unmounting a filesystem.

See [“Processes by Filesystem” on page 7-23](#) for details on how to use the Processes by Filesystem window and its buttons.



The Processes By Filesystem functionality does not work correctly in version 4.1x of the Sun operating system. The resulting report may be incomplete and/or display processes that are NOT associated with the selected partitions.

Exported Directories

Use this menu item to view a list of all directories available for export and/or export one or more of them. The window shown in [Figure 5-21](#) will appear, displaying the exportable directories on all hosts in the current pool.

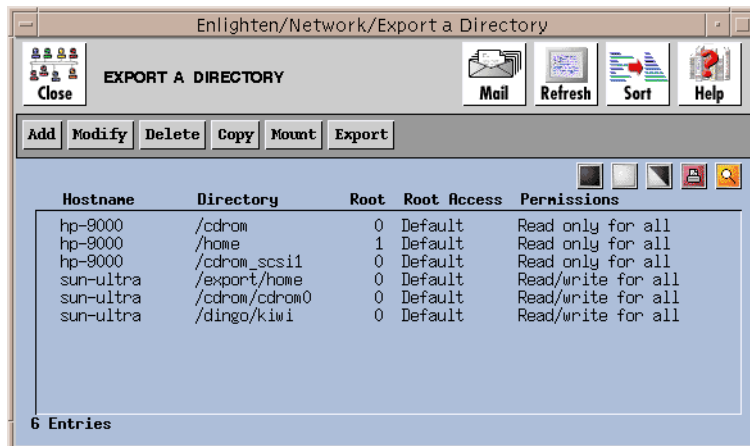


Figure 5-21 Exportable directories

Each line in the list will display the hostname, exported partition, root access, and permissions. You may select one or more partitions for further processing by highlighting it and using the following operations to further manipulate your choice(s):

- **Add** Make a directory exportable and export it
- **Modify** Modify the export parameters for an existing NFS partition
- **Delete** Remove the marked directory from the Export a Directory list
- **Copy** Copy the selected export parameters and create a new exportable directory
- **Mount** Mount an NFS partition immediately
- **Export** Re-export the selected directory

Add

Use this command to define a new export criterion for a directory. Click the Add button to bring up the Export a Filesystem window (Figure 5-22).

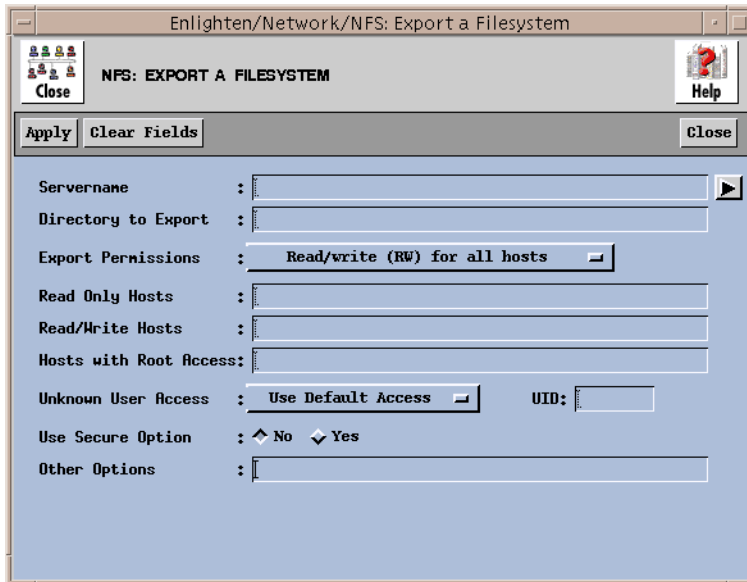


Figure 5-22 Export a Filesystem window

Add Fields

This window has the following fields:

Servername

Use this field to specify the NFS server from which the partition will be mounted. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool.

Directory to Export

Use this field to specify the full pathname for the directory to export. Leave a blank space between directory names for multiple entries.

Export Permissions

Use these toggles to specify what access permissions the client machines will have to the directory. The options are:

- Read/write (RW) for all hosts (the default)
- Read only (RO) for all hosts
- RW for specific hosts and RO for rest
- Read/write for specific hosts
- Read only for specific hosts

Read Only Hosts

This field is active only if you selected the Read only for specific hosts option in the Export Permissions field. Use this field to specify the read only hosts. Leave a blank space between hostnames for multiple entries.

Read/Write Hosts

This field is active only if you selected RW for specific hosts and RO for rest or Read/write for specific hosts options in the Export Permissions field. Use this field to specify the read/write hosts. Leave a blank space between hostnames for multiple entries.

Hosts with Root Access

Use this field to specify which hosts have root equivalency permissions on this partition. Leave a blank space between hostnames for multiple entries.

Unknown User Access

Use these toggles to determine how unknown user (UID) access should be handled. This will also determine how any root access (other than those hosts specified in the previous Hosts with Root Access field) should be handled. The options are to:

- Use Default Access (the default). UID_NOBODY if it is unknown.
- Disable Unknown Access. -1 access is denied.
- Select own UID. Only specified UID has access.

UID

This field is active only if you selected the Select own UID option in the Unknown User Access field. Use this field to enter a UID to which unknown users will be mapped.

Use Secure Option

Use this toggle to specify if the filesystem should be exported with the Secure option. The default is No.



This does not work on all systems (e.g., SCO and HP). Both the server and the client must be configured to be “secure” for this to work properly.

Other Options

You can use this field to specify other NFS options. This field is not prechecked for syntax errors.

Add Buttons

This window has the following buttons:

Apply

Click this button to export the directory.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Modify

Click the Modify button to modify the export parameters for an existing NFS partition. The Modify An Exported Filesystem window is similar to the Export a Filesystem window, except you can not modify the Servername or Directory to Export fields.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional NFS partitions if you've selected more than one to modify from the Network Filesystem Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 5-45](#).

Delete

Click the Delete button to remove directories from the Export a Directory list. EnlightenDSM will prompt you to confirm your action.

Copy

Use this command to copy the selected export parameters and create a new exportable directory. Select the desired directory and click the Copy button. The Export a Filesystem window will appear showing the highlighted export definition.

You can edit this window as needed and then click the Apply button to create the new export definition.

See [“Add” on page 5-45](#) for a description of how to use this window's fields and buttons.

Mount

Click the Mount button to mount an NFS partition immediately. The Mount An NFS Partition window will appear, as detailed in the section “[Mount](#)” on [page 5-43](#). The window’s fields will already be initialized based on the exported directory’s permissions.

For a description of the buttons and fields in this window, see “[Add](#)” on [page 5-39](#).

Export

Click the Export button to immediately re-export the highlighted directories. A pop-up window will appear asking you to confirm this action.



Security

As networks grow in size and diversity, security becomes one of the administrator's primary concerns. The features in EnlightenDSM's Security menu help ensure system and network integrity by auditing the following:

- Vital Files
- Filesystem Devices
- Boot and Shutdown Script Checks
- Crontab Contents
- Password File Integrity
- Group File Integrity
- User Home Directories
- Attempted Break-ins
- Obvious Passwords

The features described in this chapter are part of the EnlightenDSM/Advanced product. A License Advisory window similar to the one shown on [page 1-3](#) will appear if you attempt to access this menu with the Workgroup version.

Each of these security checks, except for Obvious Passwords, appends its findings to a security logfile, shown in [Figure 6-1](#). You can use the Clear Log button to flush the logfile and the Print button to print the logfile as needed.

Vital Files



This program will check the system directories and their files for easily breachable write permissions. It checks the following files and directories:

- /etc
- /usr/adm
- /usr/bin
- /usr/etc
- /usr/lib
- /usr/local
- /usr/spool

and creates a report listing any potential breaches it found.

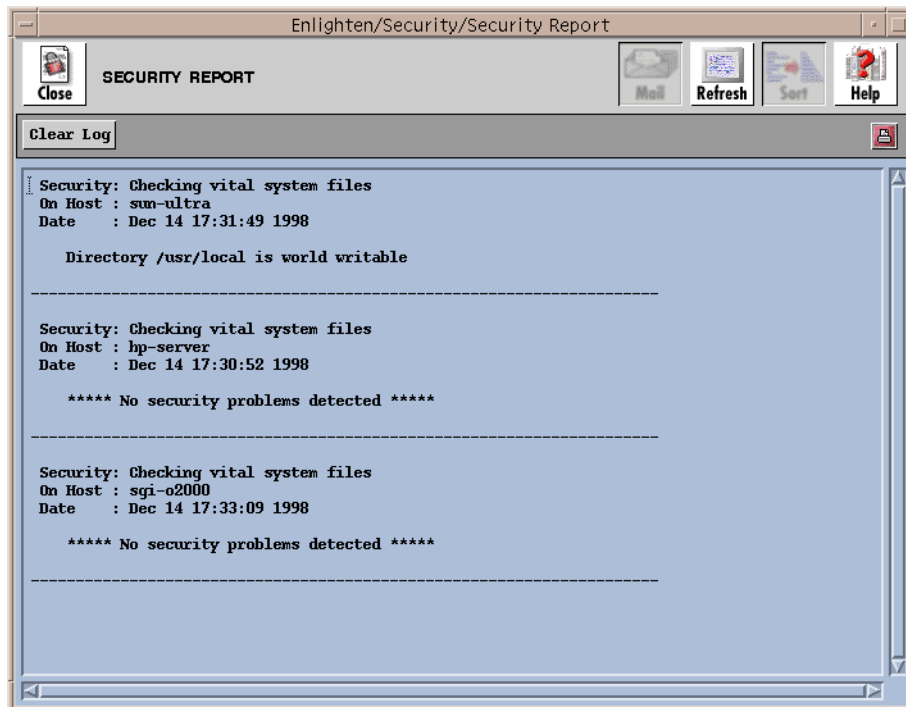


Figure 6-1 Security Report—Vital files

Filesystem Devices



This program will check that the raw device name of each mounted filesystem is in order. It looks at:

- The ownership of the device name (filesystem)
- The user group ownership of the device name (filesystem)
- Read and write permissions on the device name (filesystem)

and creates a report listing any potential problems it found.

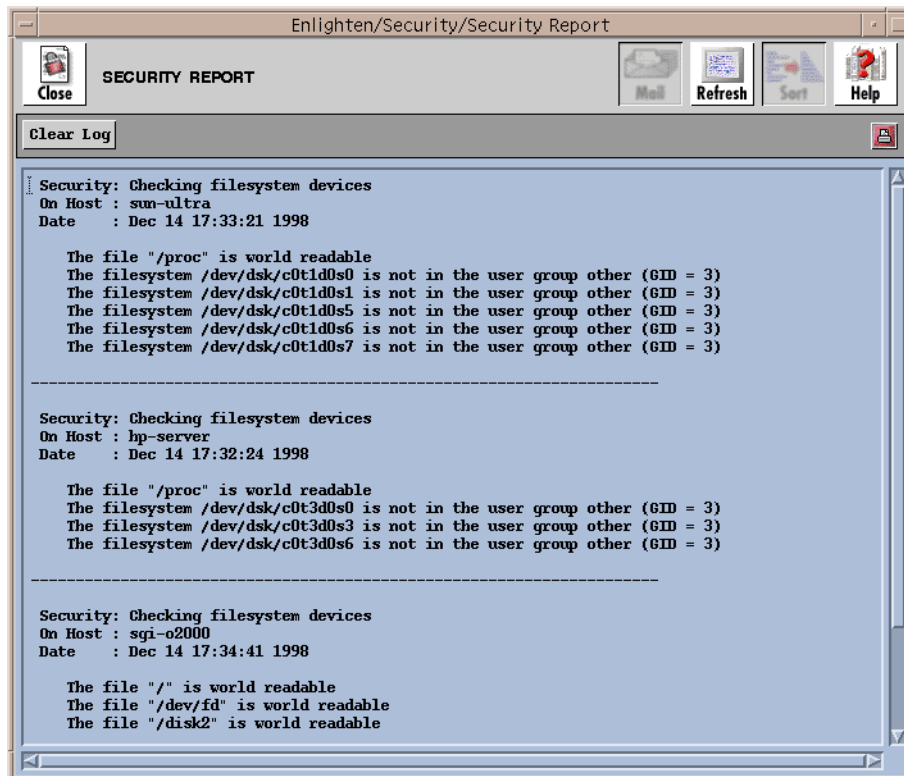


Figure 6-2 Security Report—Filesystem devices

Boot and Shutdown Script Checks



This program will check the contents of all files in the `/etc/rc*.d` directories and the `/etc/rc` files for potential back doors that could be activated at boot or shutdown time. These files are executed at boot time to initialize/shut down the system.

All the start-up and shutdown files referenced within the boot and shutdown scripts are checked for:

- The existence of the file
- Write permissions on the file
- Write permissions on the directory containing the file

and a report listing any problems found is then created. This process may be time-consuming when a large number of files is being checked.

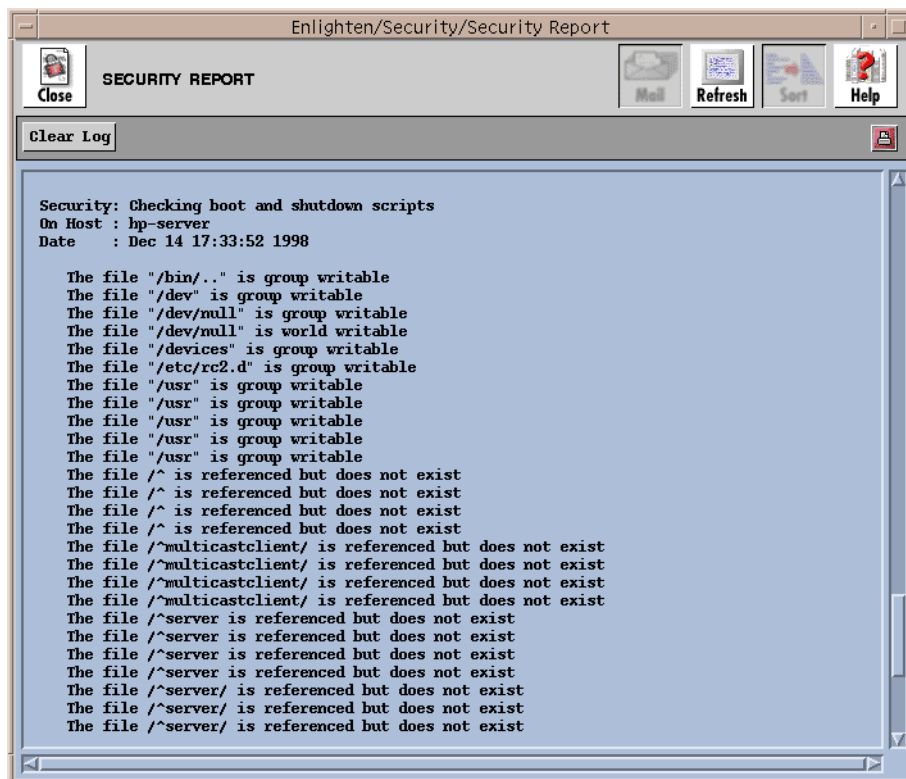


Figure 6-3 Security Report—Boot and shutdown script checks

Crontab Contents



This program will check the contents of the crontab files for potential back doors via the `cron` utility. All executable programs referenced in the cron tables are checked for:

- The existence of the file
- Write permission on the file
- Write permission on the directory containing the file
- Crontab permissions on the file

and a report listing any problems found is then created.

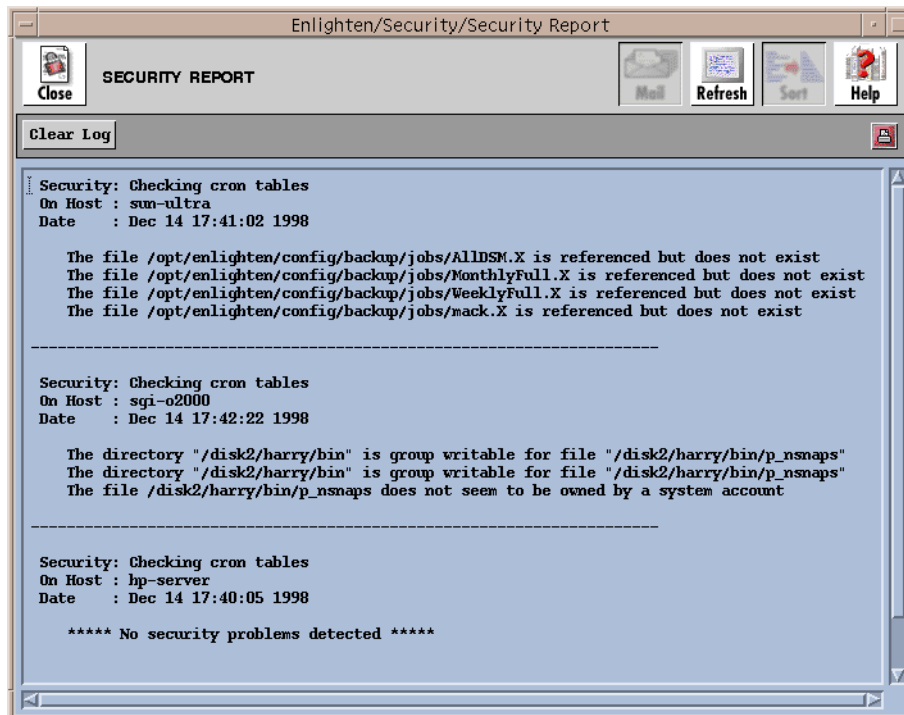


Figure 6-4 Security Report—Crontab contents

Password File Integrity



This program will check the password file `/etc/passwd` for the following security breaches:

- No blank lines in the file
- Each entry has seven fields separated by colons (':')
- The username is alphanumeric
- The user has a password
- The Userid is numeric
- If the Userid is 0, the login name is `root`
- The Groupid is numeric and exists in the group file
- The `HOME` directory of each entry exists
- The starting `SHELL` of each entry exists

and creates a report listing any problems it found.

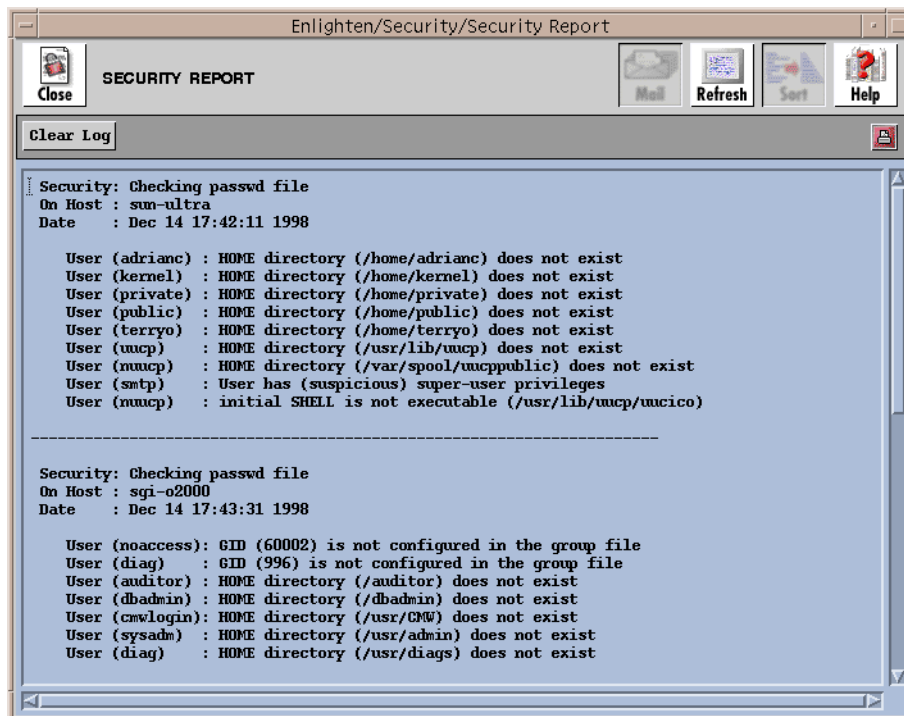


Figure 6-5 Security Report—Password file integrity

Group File Integrity



This program checks the group file `/etc/group` for the following security breaches:

- No blank lines are in the file.
- Each entry has four fields separated by colons (':').
- The groupname is alphanumeric.
- The GID is numeric.
- The user group does not have a password.
- Check each group member to see if: the username is alphanumeric, they do not exist in duplicate within the group, and they are valid users in the `/etc/passwd` file.

A list report of any problems found is created.

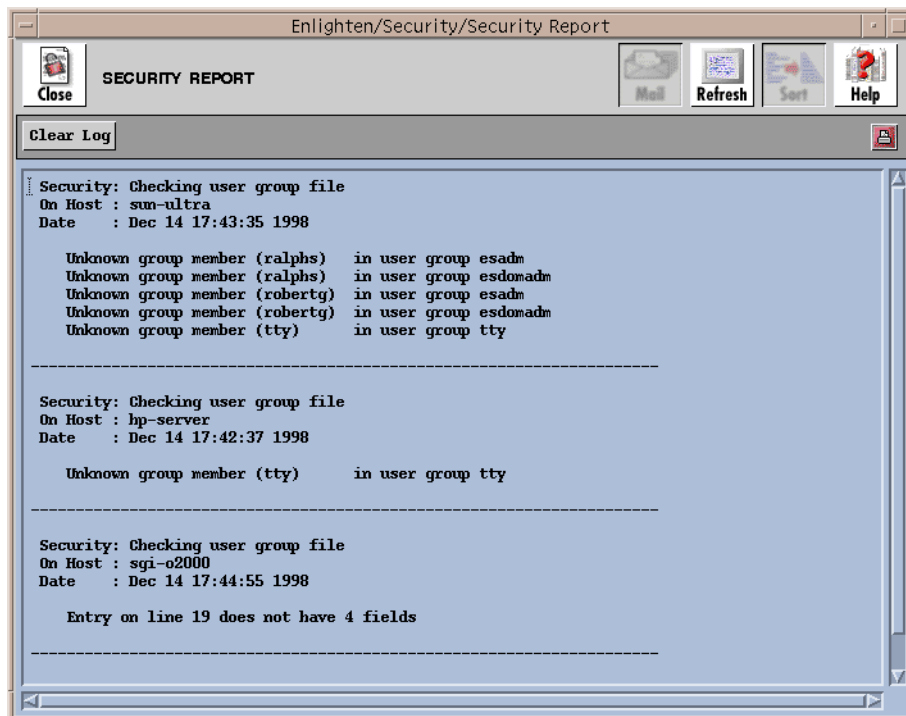


Figure 6-6 Security Report—Group file integrity

User Home Directories



This program checks each user's HOME directory for the following potential problems:

- Make sure each user has a home directory
- Make sure the directory is *not* world writable
- Checks start-up scripts for world write permission
- Checks start-up scripts for group write permission

and creates a report listing any problems it found.

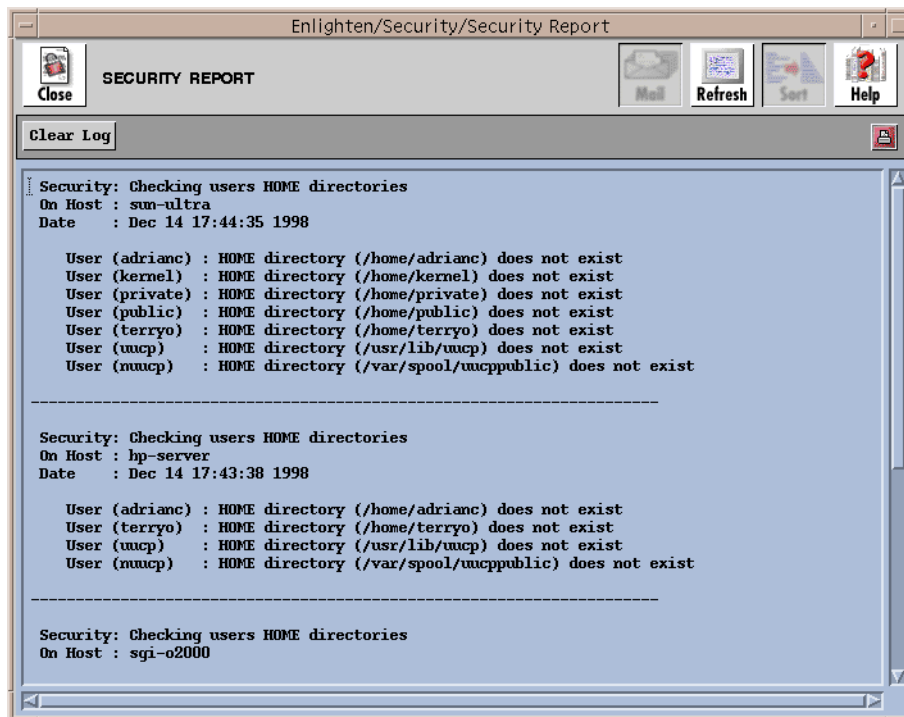


Figure 6-7 Security Report—User home directories

Attempted Break-ins



This program finds all the users who failed a second attempt to become superuser or some other user through the `su` command. It then creates a report listing any problems it found.

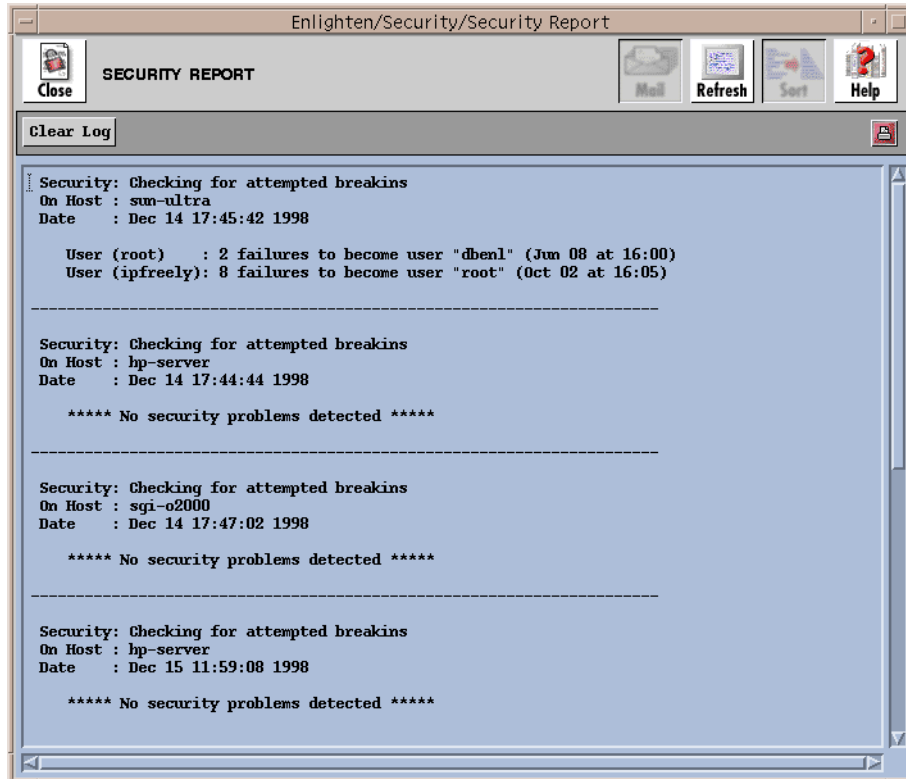


Figure 6-8 Security Report—Attempted break-ins

Obvious Passwords



One of the most frequent invasions of computer systems is not through a complex or back door trick, but through the front door via the login program. So, simple account passwords can be a high security risk.

Users frequently choose very simplistic passwords. While some systems have password programs to force some variation of simple passwords, the protection is usually minimal. You can use the Obvious Password security check to get a list of users with easily detectable passwords.

This program runs obvious password checks on:

- The user's login name
- The `passwd` file on the machine (hostname)
- All words in the user's Realname, Phone number, and Office user account fields of the `passwd` file

after you fill in the relevant data fields in the Password Check Parameters window (Figure 6-9).

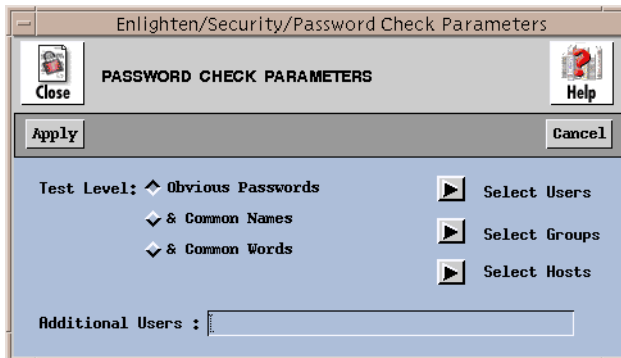


Figure 6-9 Password Check Parameters window

Since users have the ability to change their passwords at any time, you should run this security check frequently. Keeping a historical reference will also provide some insight into your own users' password habits.

Levels of Checking

There are three levels of Obvious Password Checking you can select, with level one being the lowest and each subsequent level encompassing the previous levels:

1) Obvious Passwords

Passwords that closely match the account name. Each user takes approximately 5 seconds to check.

2) Common Names

Passwords matching proper names, a frequent choice of novice users. Each user takes several minutes to check.

3) Common Words

Passwords found in a common-use dictionary. Each user takes several minutes to check.

Who is Checked

Click the right arrows in the Password Check Parameters window for lists of the users, groups, and hosts on the system. You can use these to target specific users and/or groups for the checking process.

For the Users and Groups windows, you will be presented with the users and groups in the current system pool. Choosing a user or group will cause all occurrences of that user and/or group to be checked across the hosts in that system pool. For the Hosts window, selecting a host will check all users on that host.

Instead of using the above process, you may also choose to examine only specific targets by using the Additional Users field. Use the format *hostname:username* to specify these entries. The following example tests user *johnf* on host *rome* and all users on host *paris*.

```
Additional Users:           rome:johnf paris:
```

Whichever method you use, when you are finished, click the Apply button. The time required to execute the testing process depends on the number of users to check and the level of checking. When this check is done, a list of users with obvious passwords appears in the form of the Configure Users window.

At this point, you could select a subset of the users and lock them out, obtain a list of their current processes, change their passwords, e-mail an “Obvious Password” letter to them, or even delete them.

What is Checked

All Options

EnlightenDSM will check each word in the generated list of potentially obvious passwords as follows:

- As entered
- In all lowercase letters
- In all uppercase letters
- Capitalized first letter, the rest in lowercase

Each of these checks will be performed on the actual word as well as its reverse. For example, 'Mirror' will be checked as Mirror, rorriM, mirror, rorrim, MIRROR, RORRIM.

This program will also check for accounts with no password, empty passwords and other obvious passwords such as "secret", "computer", "hello", etc.

The Common Names Option

The second level of password checking checks a list of 4000 names for possible matches. This list of names is found in the file `$ENLIGHTEN/config/names`. You may modify this file to fit your working environment.

Entries in the file are checked as is, from first to last, and any deviations such as capitalization or backwards spelling are not made. If you want these deviations to be checked, you need to add them sequentially as entries in the file.

Audit Log

The audit log utility is implemented in order to facilitate the monitoring of transactions performed with the use of our product. The log keeps track of the transaction performed, the time, the user, and other imperative information.

Here are two examples of why it is necessary:

Example 1

EnlightenDSM is deployed in a Help Desk environment where the Help Desk operators have access to user account information. A corporate computer user calls the Help Desk supervisor to complain that he has called into the Desk 3 times to have his password reset, but the password still is not reset. The Help Desk supervisor review the audit log and determines that Help Desk Operator #4 has attempted to reset the password 3 times. Furthermore, the audit log revealed that the Operator was resetting the password on the wrong machine. With this useful information, the Help Desk supervisor re-trains Operator #4 on the proper procedure of resetting a user password. Operator #4 then proceeds to successfully reset the password for the Help Desk caller.

Example 2

An NFS partition that contains the Engineering source code has been wiped cleaned. The only way this would be possible is via root access to the machine hosting the NFS partition. The company keeps the root password a closely guarded secret, so the only way the NFS partition could have been erase would be via EnlightenDSM. The System Administrator examines the audit log and discovers that a cronjob was created by disgruntled employee 5 days ago to perform the evil deed. The Administrator alerts the employees supervisor, and the supervisor takes the appropriate action.

Enabling and Disabling the Audit Log

By default, the Audit Log is disabled when EnlightenDSM is installed. The Audit Log is controlled by the `AUDIT_LOG` parameter in the `params` file on each Agent and GUI host.

The `params` file is located in the `config` sub-directory of the EnlightenDSM install directory. To enable Auditing for operations on a particular host, edit the `AUDIT_LOG` entry in the `params` file on that host to read:

```
AUDIT_LOG 1
```


You will have to restart the Enlighten**DSM** daemons and the GUI before the change will take effect. You can restart the daemons by executing

```
# <install-dir>/bin/start_enl_daemons -restart
```

To disable the Audit Log, simply change the `AUDIT_LOG` entry to read:

```
AUDIT_LOG 0
```

0 and 1 are the only valid values for `AUDIT_LOG`. Any other value will disable the `AUDIT_LOG`.

Enlighten**DSM** can be used to enable the Audit Log on all the managed hosts via the GUI in a few mouse clicks. This is accomplished using the Custom command button. For an example on the use of this button, please see [“View Files” on page 7-9](#).

Data Storage and Display

Audit information is logged to a centralized location, which is the Enterprise Management Database (EMD). Examples of the types of operations recorded in the audit log are: when a new pool is created, when a user’s account is changed, and when a print job is deleted from the printer Queue.

The audit information may be accessed through a supplied command line utility. For information on the utility, see [Appendix J, “EMD.”](#)



Disk

The Disk menu is used to view summaries of disk space usage by user, group or filesystem. These reports help network administrators manage the space on network partitions.

To decrease the time needed to generate these summaries, Enlighten**DSM** creates a database, or *snapshot*, of the files in a filesystem. That *snapshot* is used to perform searches or generate reports.

The Disk menu can also be used to search for specific files.

The Disk menu options are:

- Usage by Filesystem
- Usage by Username
- Usage by Groupname
- File Search
- Processes by Filesystem
- Save Current Snapshots
- Compare Snapshots

The features described in this chapter are part of the Enlighten**DSM**/Advanced product. A License Advisory window similar to the one shown on [page 1-3](#) will appear if you attempt to access this menu with the Workgroup version.

Building Snapshots

To increase the speed of these programs and cut down on disk access time, EnlightenDSM creates a database (disk snapshot) of all files in the requested filesystem once and then uses this database to perform the required tasks. The amount of time needed to create this database is dependent on the size of the filesystem you wish to monitor.

To update the snapshot of a filesystem, highlight it in the list of filesystems and click the Rebuild button located on the right side of the menu bar. Then click the button for the process you want to update: User Sum., Group Sum., or File Search.

EnlightenDSM will display a window indicating when the last snapshot of your request was taken. If you want to refresh this disk snapshot, click the appropriate button. To save the new snapshot, see [“Save Current Snapshots” on page 7-26](#).



Disk summaries and searches are performed on this snapshot. Consequently, you should rebuild your snapshots periodically to reflect a more current state of the disk.

Usage by Filesystem



You can use this command to view Disk Usage by Filesystem. A window showing all mounted file systems as well as their current usage of disk space and inodes will appear (Figure 7-1).

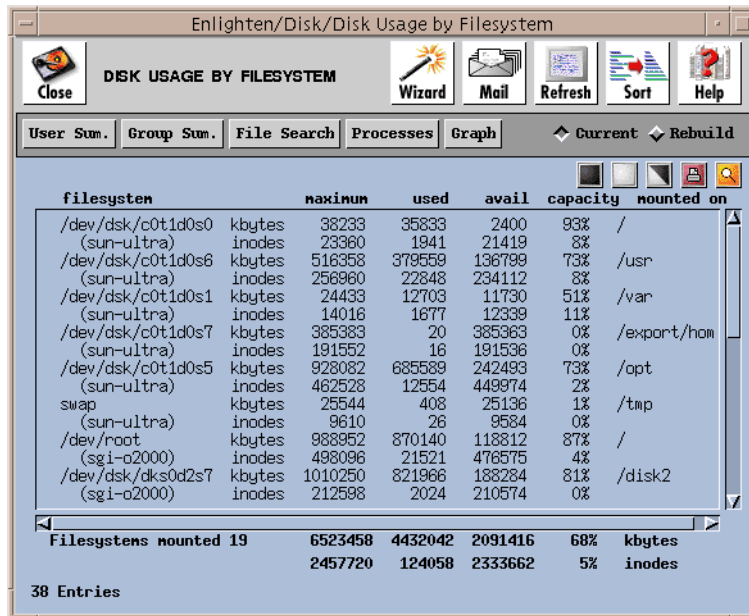


Figure 7-1 Disk Usage By Filesystem window

You now have the option to look at:

- User Sum. Summarize disk usage by username
- Group Sum. Summarize disk usage by user group
- File Search Search for files based on search criteria
- Processes View the processes using a filesystem
- Graph Graph the disk usage

User Sum.

Highlight the filesystems you wish to summarize, then click the User Sum. button. The Disk Usage Information By Users window will appear summarizing the disk usage by user for the highlighted filesystems. This

window shows you the user name, real name, telephone number, number of files owned by the user, and total disk usage (in KB).

The disk usage summaries are based on previously taken disk snapshots. Consequently, the displayed data may not be current. To rebuild the disk snapshots before the summary is made, click the Rebuild button in the previous window. See [“Building Snapshots” on page 7-2](#) for more details.

Within the Disk Usage Information By Users window, you now have the following options:

- View Files View a list of files owned by the user(s) and perform further operations on them
- Compare Snapshots Compare the current snapshot to a previously saved one
- Save Snapshots Save the current snapshot
- Graph Graph the disk usage of the user(s)

View Files

By highlighting specific Username entries and clicking on the View Files button, you can view a list of files owned by the highlighted users and then manipulate the file characteristics. For more details, refer to [“View Files” on page 7-9](#).

Compare Snapshots

You can compare the state of the current partition (snapshot) to that of one from a previous moment in time. You are able to view summaries of changes, see detailed lists of changes, and look for changes based on search criteria. To compare an existing snapshot to a previously saved one, highlight the disk partitions (filesystems) you want to compare and click the Compare Snapshot button to run the comparison. For more details, see [“Compare Snapshots” on page 7-14](#).

Save Snapshots

To save existing snapshots, highlight the disk partitions you want to save and click the Save Snapshot button to create a dynamic new snapshot. For more details, see [“Save Current Snapshots” on page 7-26](#).

Graph

To graph the disk usage summary by user, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

Group Sum.

Highlight the filesystems you wish to summarize, then click the Group Sum. button. The Disk Usage Information By Group window will appear summarizing the disk usage by user for the highlighted filesystems. This window shows you the group name, number of files owned by the group, and total disk usage (in KB).

The disk usage summaries are based on previously taken disk snapshots. Consequently, the displayed data may not be current. To rebuild the disk snapshots before the summary is made, click the Rebuild button in the previous window. See [“Building Snapshots” on page 7-2](#) for more details.

Within the Disk Usage Information By Group window, you now have the following options:

- View Files View a list of files owned by the group(s) and perform further operations on them
- Compare Snapshots Compare the current snapshot to a previously saved one
- Save Snapshots Save the current snapshot
- Graph Graph the disk usage of the group(s)

View Files

By highlighting specific Groupname entries and clicking on the View Files button, you can view a list of files owned by the highlighted users and then manipulate the file characteristics. For more details, refer to [“View Files” on page 7-9](#).

Compare Snapshots

You can compare the state of the current partition (snapshot) to that of one from a previous moment in time. You are able to view summaries of changes, see detailed lists of changes, and look for changes based on a search criterion. To compare an existing snapshot to a previously saved one, highlight the disk partitions (filesystems) you want to compare and click the Compare Snapshot button to run the comparison. For more details, see [“Compare Snapshots” on page 7-29](#).

Save Snapshots

To save existing snapshots, highlight the disk partitions you want to save and click the Save Snapshot button to create a dynamic new snapshot. For more details, see [“Save Current Snapshots” on page 7-26](#).

Graph

To graph the disk usage summary by group, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

File Search

There will be times when you need to perform disk file searches for information not readily apparent, such as file size, file types, creation date of a file, and who owns a particular file. Highlight the disk partitions (filesystems) you wish to query and click the File Search button to find, and then view, these files through a query process.

The File Search Query window will appear where you can enter your search criteria. If you do not fill in a field, it is not used as part of the search pattern.

Processes

You can use this command to display all currently running processes that are accessing the highlighted filesystems. To view the processes, highlight the filesystems you wish to view and then click the Processes button. A window will appear displaying all processes accessing the highlighted filesystems.



The Processes By Filesystem functionality does not work correctly in version 4.1x of the Sun operating system. The resulting report may be incomplete and/or display processes that are NOT associated with the selected partitions.

Graph

To graph the disk usages, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

Usage by Username



You can use this command to see, at a glance, a summary of disk usage by user for currently mounted file systems. The Select Disk Partitions window will appear displaying all currently mounted file systems (Figure 7-2).

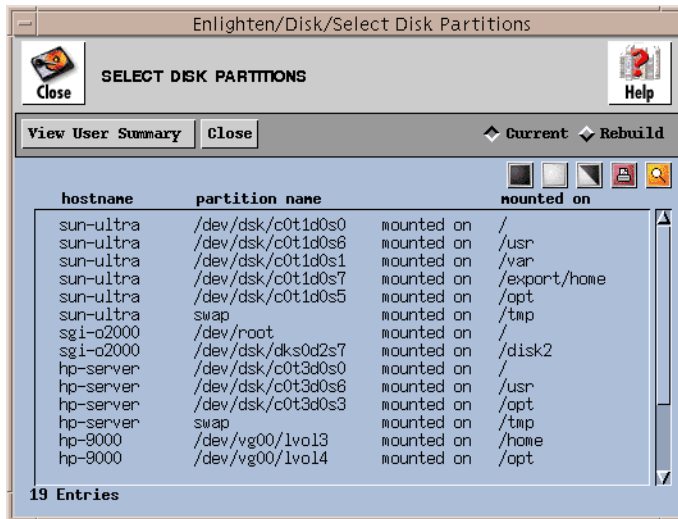


Figure 7-2 Select Disk Partitions window

Highlight the appropriate disk partitions for summarizing from the list box and click the View User Summary button. The Disk Usage Information by Users window will appear (Figure 7-3). You can sort this information, mail it to selected users, or print it. Plus, you now have the following options:

- View Files View a list of files owned by the user(s) and perform further operations on them
- Compare Snapshots Compare the current snapshot to a previously saved one
- Save Snapshots Save the current snapshot
- Graph Graph the disk usage summary of the user(s)

username	realname	telephone	files	total (KB)
root	Super-User	unknown	1630	10226
adm	Admin	unknown	4	1379
lp	Line Printer Admin	unknown	28	11
bin	unknown	unknown	11	5
nobody	Nobody	unknown	3	2
dben1	Informix user	unknown	3	1
uucp	uucp Admin	unknown	2	1
sys	unknown	unknown	1	1
Total Users		8 (0 unknown)	1680 files	11624 KB of data
Unknown users have		0 files	0 KB of data	

8 Entries

Figure 7-3 Disk Usage Information by Users window

View Files

By highlighting specific Username entries and clicking the View Files button, you can further view a list of files owned by the highlighted users, as shown in the File Systems Detail window (Figure 7-4).

```

filename
sun-ultra:/var/sadm
sun-ultra:/var/sadm/install/admin/default
sun-ultra:/var/sadm/install/contents
sun-ultra:/var/sadm/install/.lockfile
sun-ultra:/var/sadm/install/.pkg.lock
sun-ultra:/var/sadm/pkg
sun-ultra:/var/sadm/pkg/SUNWkwrt1
sun-ultra:/var/sadm/pkg/SUNWkwrt1/install
sun-ultra:/var/sadm/pkg/SUNWkwrt1/install/copyright
sun-ultra:/var/sadm/pkg/SUNWkwrt1/save
sun-ultra:/var/sadm/pkg/SUNWkwrt1/pkginfo
1630 files using 10226 KB
1630 Entries

```

Figure 7-4 File Systems Detail window

Now you can perform the following operations to further manipulate the files in the list box:

- Edit Change the contents of the marked files
- View View the contents of the marked files
- Stats Display statistics about the marked files
- Delete Remove the marked files from the system
- Backup Back up the marked files
- chmod Change the read, write, and/or execute permissions of the marked files
- chgrp Change the group(s) to which the marked files belong
- chown Change the owner of the marked files
- Custom Define and execute a UNIX command on the marked files

Edit

Use this command to edit files ('vi'). Highlight the files you wish to edit and click the Edit button. A pop-up confirmation window will appear. For each file, select Yes to edit the file, No to skip that file and go to the next one, All to edit all files, or Cancel to abort the command. You MUST utilize standard 'vi' commands to edit any files (or a hex editor if this is a binary file).

When you quit 'vi' with the :q command, the following message will appear: "Hit the Return key to continue". This action will close the window and allow you to continue using other EnlightenDSM options.

View

Use this command to view the contents of files. Highlight the files you wish to view and then click the View button. A pop-up confirmation window will appear. For each file, select Yes to view the file, No to skip that file, All to view all files, or Cancel to abort the command. If you select Yes or All, EnlightenDSM will activate a window showing the contents of the file.

Once the (last) file is displayed, the following message will appear: “Hit the Return key to continue”. This action will close the window and allow you to continue using other Enlighten**DSM** options. You cannot use the View command to modify files.

Stats

Use this command to display a statistical summary for each file (Figure 7-5). Highlight the files you want and click the Stats button. A window will appear with the statistics for the highlighted file. If you selected more than one file, click the Next button to view statistics on the next file you selected.

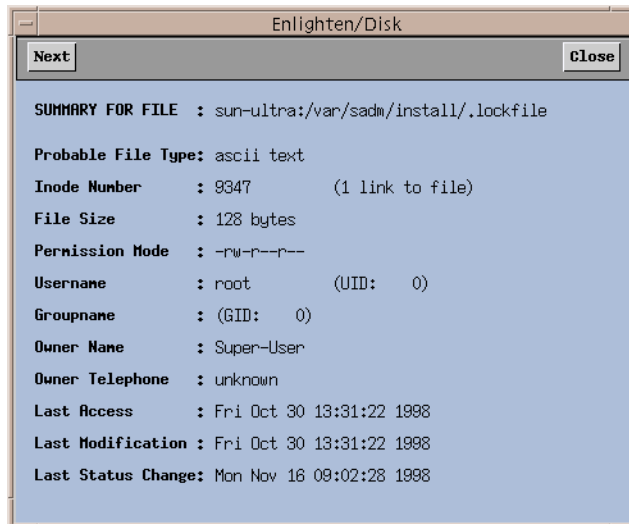


Figure 7-5 Statistical File Summary

Delete

Use this command to remove files from the system. Highlight the files you wish to delete and click the Delete button. Enlighten**DSM** will prompt you to confirm your action. Once a file has been deleted, it cannot be recovered.



If a selected file is actually a directory, Enlighten**DSM** will not remove it unless it is empty. If the directory is not empty, a status message will appear stating this.

Backup

Use this command to make a back up copy of each highlighted file. Highlight the files you wish to back up and click the Backup button. A pop-up window will appear showing the different backup devices available. Select one of them. EnlightenDSM will then ask for confirmation before performing the backup.

During the backup, a window will display the progress of the backup. When the backup is complete, a summary log will be displayed in the same window. Press the <return> key to continue using other EnlightenDSM options.



If a selected file is actually a directory, only the directory itself will be backed up. Subfiles and/or subdirectories of the highlighted directory will *NOT* be backed up.

Chmod

Use the chmod command to manipulate the read, write, and/or execute permissions of a file. Highlight the files you wish to change and click the Chmod button. A pop-up window will appear showing the current set of permissions for that file (Figure 7-6).

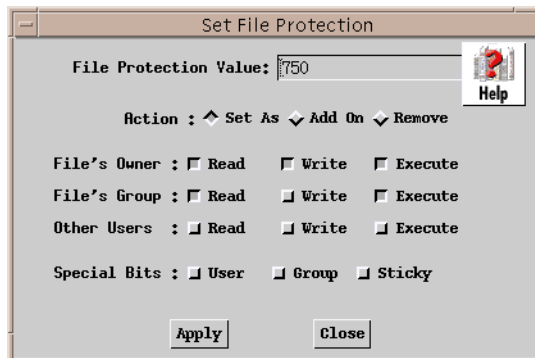


Figure 7-6 File Permissions

You can enter the new permissions numerically in the first field or select the desired permissions from the available check boxes. Furthermore, you can

specify the permissions be set as is, added on to, or subtracted from the current mode.

Once you've made your choices, click the Apply button. Another pop-up window will ask you to confirm your changes.

Chgrp

Use the `chgrp` command to change the group(s) to which a file or files belong. Highlight the files you wish to change and click the Chgrp button. A pop-up window will appear. Use this to specify the new group(s).

You can click the arrow button to access a list of valid User Groups. A window will appear showing all Groups. Highlight the desired Group(s) and then click the Apply button. You must specify a valid User Group name. Another pop-up window will ask you to confirm your changes.

Chown

Use the `chown` command to change the ownership of a file. Highlight the files you wish to change and click the Chown button. A pop-up window will appear. Use this to specify the new username.

You can click the arrow button to access a list of valid usernames. A window will appear showing all Users. Highlight a single user and click the Apply button. You must specify a valid username. Another pop-up window will ask you to confirm your changes.

Custom

Use this command to directly apply UNIX commands to files without having to exit Enlighten**DSM**. The commands entered can be any valid UNIX commands. Highlight the files you need to use for your custom command and click the Custom button.

A pop-up window will appear. Enter the appropriate UNIX command and click the Apply button. Another pop-up window will ask you to confirm your action. Once the command has executed successfully, the following message will appear: "Hit the Return key to continue". This action will close the window and allow you to continue using other Enlighten**DSM** options.



EnlightenDSM supports using the character macro % as part of the command. If you use this macro, EnlightenDSM will replace it with the name of the current file. If not, EnlightenDSM will automatically append the name of the current file to the command before executing it.

Compare Snapshots

You can compare the state of the current partition (snapshot) to that of one from a previous moment in time. You are able to view summaries of changes, see detailed lists of changes, and look for changes based on a search criterion. To compare an existing snapshot to a previously saved one, highlight the disk partitions (filesystems) you want to compare and click the Compare Snapshot button to run the comparison.

Save Snapshots

To save existing snapshots, highlight the disk partitions you want to save and click the Save Snapshot button to create a dynamic new snapshot. For more details, see [“Save Current Snapshots” on page 7-26](#).

Graph

To graph the disk usages, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

Usage by Groupname



You can use this command to see, at a glance, a summary of disk usage by user group for currently mounted file systems. The Select Disk Partitions window will appear displaying all currently mounted file systems (Figure 7-7).

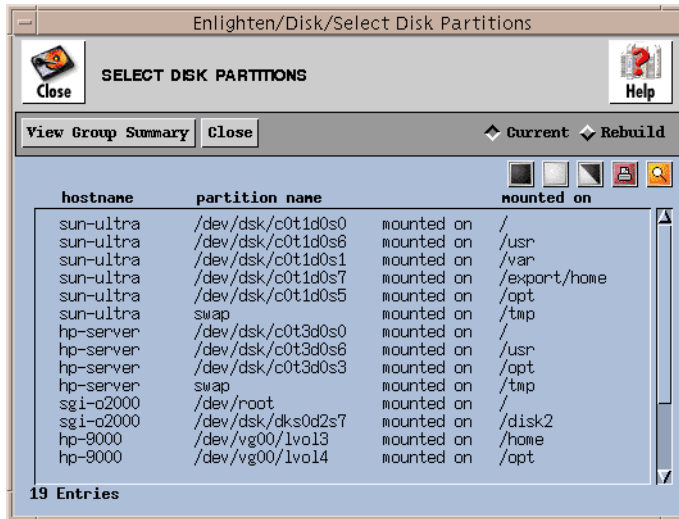


Figure 7-7 Select Disk Partitions window

Highlight the disk partitions you want to monitor from the list box and click the View Group Summary button. The Disk Usage Information By Group window will appear (Figure 7-8). Within this window, you now have the following options:

- View Files View a list of files owned by the group(s) and perform further operations on them
- Compare Snapshots Compare the current snapshot to a previously saved one
- Save Snapshots Save the current snapshot
- Graph Graph the disk usage of the group(s)

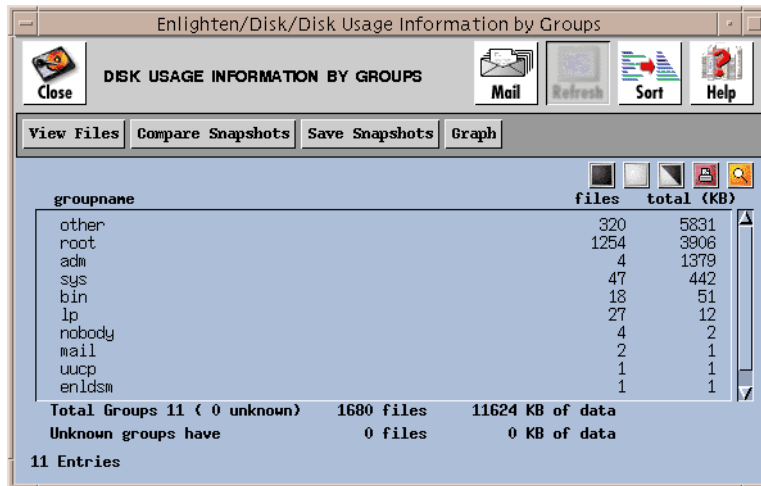


Figure 7-8 Disk Usage Information By Group window

View Files

By highlighting specific Groupname entries and clicking the View Files button, you can further view a list of files owned by the highlighted users and then manipulate the file characteristics. For more details, refer to [“View Files” on page 7-9](#).

Compare Snapshots

You can compare the state of the current partition (snapshot) to that of one from a previous moment in time. You are able to view summaries of changes, see detailed lists of changes, and look for changes based on a search criterion. To compare an existing snapshot to a previously saved one, highlight the disk partitions (filesystems) you want to compare and click the Compare Snapshot button to run the comparison. For more details, see [“Compare Snapshots” on page 7-29](#).

Save Snapshots

To save existing snapshots, highlight the disk partitions you want to save and click the Save Snapshot button to create a dynamic new snapshot. For more details, see [“Save Current Snapshots” on page 7-26](#).

Graph

To graph the disk usages, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

File Search



There will be times when you need to perform disk file searches for information not readily apparent, such as file size, file types, creation date of a file, and who owns a particular file. You can use this command to find these files.

A window will appear listing all the currently mounted filesystems (Figure 7-9). Highlight the disk partitions (filesystems) you wish to query and then click the Search Parameters button.

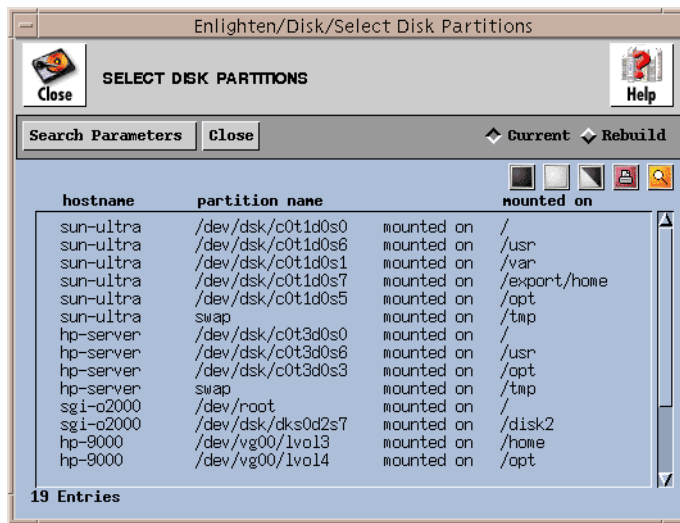


Figure 7-9 Currently mounted filesystems

The File Search Query window will then appear (Figure 7-10). You can now fill in the fields to use in your search pattern. If you do not fill in a field, it is not used as part of the search.

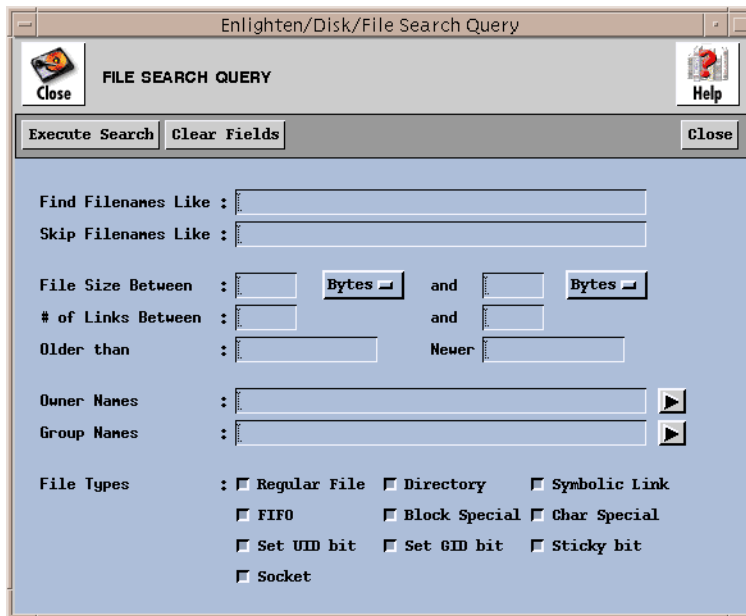


Figure 7-10 File Search Query window

The rest of this section describes the functionality of this window's fields and buttons.

Fields

This window contains the following fields:

Find Filenames Like

Use this field to specify filenames to use in the search pattern. The filenames can contain the standard UNIX wild card matching characters '*', '?', and '['. This field will accept more than one entry. Leave a blank space between each unique filename for multiple entries.

Skip Filenames Like

This is the opposite of the Find Filenames like field. The filenames entered here create an exclusion list. Any filename entered here will NOT be matched by the search. This field uses the same format as the previous field.

File Size Between ... and

Every file residing on the system requires space. This space is referred to as the file size. Any input for these fields specifies the minimum and/or maximum file size a file can have as part of the search routine. If a file is smaller than the amount entered in the first field, the search routine will exclude it. Conversely, if the file is larger than the amount entered in the second field, the search routine will also exclude it. You can also use these two fields together to specify a range of file sizes in your search.

To change the default size, enter the data in the format `N unit`, where `N` is the numeric size to search for and `unit` is one of the following units of measure:

- KB = Kilobytes
- MB = Megabytes
- GB = Gigabytes

You can either add the unit directly in the field or just add the numeric size and click the Bytes button to select the unit. The file size must be specified as a whole number.

of Links Between ... and

Files on a UNIX system can have more than one name (every file is considered to have one link to its original name). If they do have more than one name they are referred to as linked files. In many cases a file may have more than one link.

Any input for these fields specifies the minimum and/or maximum links a file can have as part of the search routine. If the number of links for a file is smaller than the amount entered in the first field, the search routine will exclude it. Conversely, if the number of links is greater than the amount entered in the second field, the search routine will also exclude it. You can also use these two fields together to narrow the search to a specific link count range.

Older than

When a file is created, its creation date is stored by the filesystem in what is called an inode. Every time the file is modified the date of the modification is updated in the inode. Specifying an entry in this field will limit the search to any files created *after* the date entered. You can use one of two formats for the date. See [Appendix C, "Time Formats,"](#) for more details.

Newer

Specifying a date in this field (which has the same format as the previous field) will limit the search to any files created before the date entered. You can also use the Older than and Newer fields together to specify a range of creation times in your search.

Owner Names

Every file on the system has an owner. In most cases the owner is the user who created the file. Use this field to find only those files whose owners match the names of the users entered. Click the arrow button and a pop-up menu will appear listing all Owner names. Highlight the Owner names you wish to search on, then click the Apply button.

This field will accept more than one entry. Leave a blank space between each unique owner for multiple entries. To exclude any files owned by certain users, place a '!' before those usernames.

Group Names

Every file on the system belongs to a user group. Use this field to find only those files whose group ownership matches the user group names entered in this field. Click the arrow button and a pop-up menu will appear listing all Group Names. Highlight the Group Names you wish to search for, then click the Apply button.

This field will accept more than one entry. Leave a blank space between groups for multiple entries. To exclude any files owned by certain groups, place a '!' before those group names.

File Types

Use this field to select specific file types for the search. To do so, click in the box preceding the file description. You may select more than one file type. The default is all file types are used:

Regular file	Directory	Symbolic Link
FIFO (pipe)	Block Special	Char Special
Set Userid Bit	Set Groupid Bit	Sticky bit
Socket		

Buttons

This window contains the following buttons:

Execute Search

The search criteria may use one or any combination of the preceding fields. Once you've entered your search criteria, click the Execute Search button. The File Systems Detail window will appear with a list of the files that met the criteria. You may then execute more commands such as edit, delete, and chown to manipulate the displayed files' characteristics. See [“View Files” on page 7-9](#) for more details on how to use the File Systems Detail window.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Processes by Filesystem



You can use this command to display a report on all currently running processes associated with the highlighted Disk Partitions. The Select Disk Partitions window will appear displaying all currently mounted filesystems, along with where they are mounted (Figure 7-11).

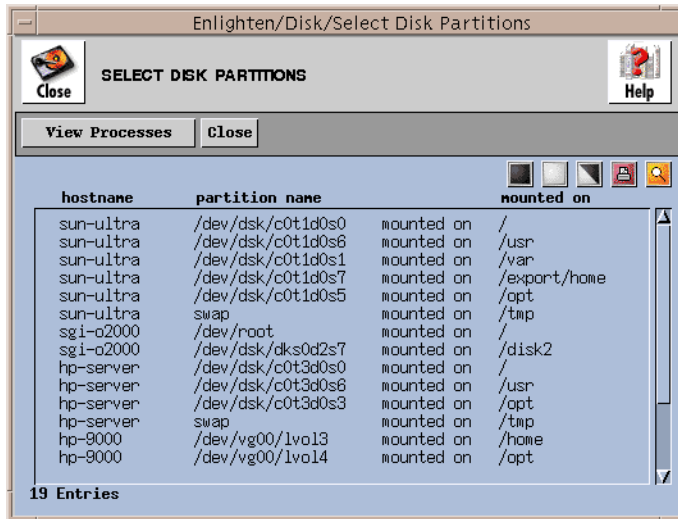


Figure 7-11 Select Disk Partitions window

Highlight the desired disk partitions you want to use to further examine their associated processes and click the View Processes button. The Processes By Filesystem window will appear (Figure 7-12).



The Processes By Filesystem functionality does not work correctly in version 4.1x of the Sun operating system. The resulting report may be incomplete and/or display processes that are NOT associated with the selected partitions.

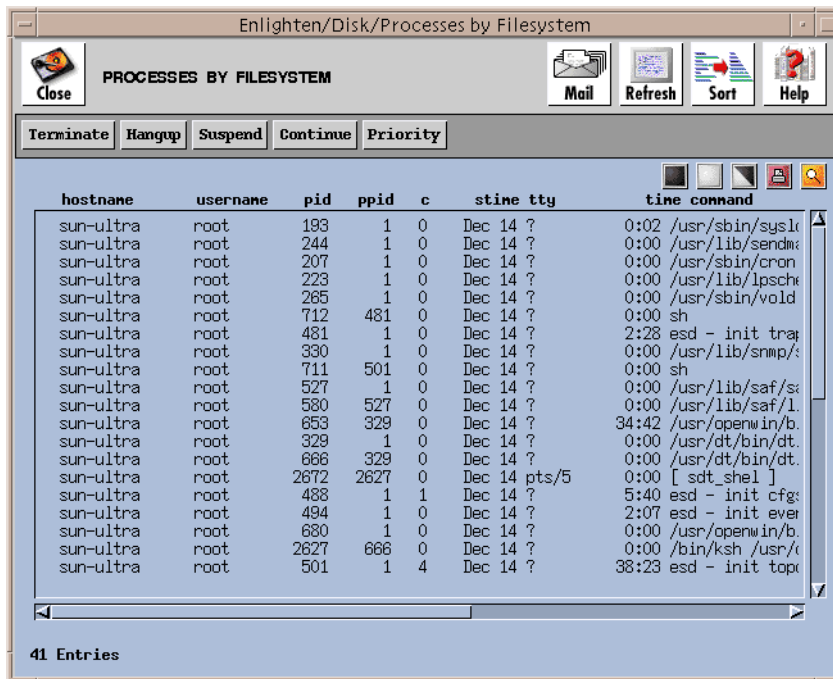


Figure 7-12 Processes by Filesystem window

Then you can select a process and use the menu buttons to impact it, as described in the following subsections.

Terminate

This command is very powerful and will immediately kill the highlighted process. This command will not kill related processes, so if there are child processes running they will become orphans and will have to be terminated separately. A pop-up window will prompt you for verification to terminate the process.

Hangup

This command is similar to the Terminate command, except it provides enough time for the process to shut down properly. This means the process can close any files and terminate any child processes. A pop-up window will prompt you for verification to hang up the process.

Suspend

This command stops a process from working, but it does not terminate the process. Essentially, this command puts a process on hold; it can be activated again at a later time. You must use the Continue command to re-activate a suspended process. A pop-up window will prompt you for verification to suspend the process.

Continue

This command allows you to re-activate a process that was previously put on hold by a Suspend command. A pop-up window will prompt you for verification to resume the process.

Priority

This command allows you to change the priority of a process. This priority determines when the CPU acts on a process. It may have a value from -20 to +20; the smaller the number, the higher the priority. Clicking the Priority button activates a Process Priority window. You can enter the desired priority or use the arrow buttons to make your selection.

Save Current Snapshots

EnlightenDSM creates disk snapshots (databases) for use in its monitoring of disk usage. These snapshots can be saved for historical purposes and then used to determine if changes have been made to a partition. They can also be used for security-related issues (see Appendix A, “EnlightenDSM Basics,” in the *EnlightenDSM User Guide*).

There are two ways to store a snapshot. You can store it as the Master Snapshot. This is a snapshot of what you believe is a secure system state. This is usually made before users are given access to the partition or after a system cleanup is done. Then you can use this Master state to gauge if changes are made over time to that partition.

The second way is similar to the first, except you can assign a name to the saved snapshot. These snapshots can then be used to represent intermediate states.

Either way, once you select this menu item, a window will appear listing all currently mounted filesystems ([Figure 7-13](#)). Highlight the disk partitions for which you want to make a current snapshot and click the Save Snapshot button.

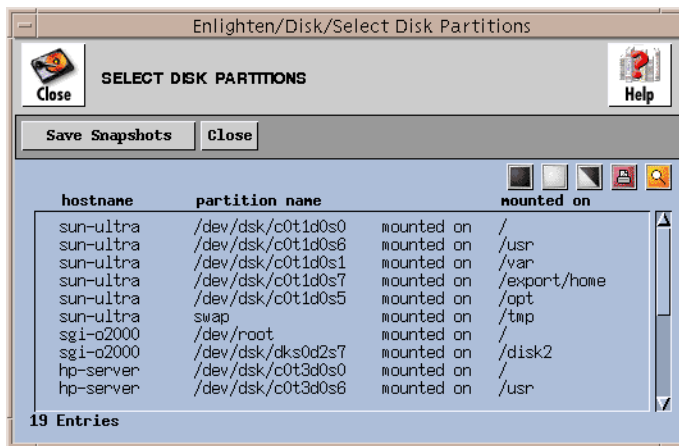


Figure 7-13 Currently mounted filesystems

The Save Selected Disk Snapshots window will appear (Figure 7-14). The names of the selected partitions are automatically imported into the Snapshot from the partition column in the window. You may also add other partition names or modify the existing ones.



Figure 7-14 Save Selected Disk Snapshots window

Then you must choose for each partition whether to save the snapshot as the master snapshot (default) or not. If you do, click the Yes button in the Save as Master column. If not, click the No button and enter a unique name for this snapshot in the appropriate Save snapshot as field.

Once you've made all your selections, click the Save Snapshot button. Enlighten**DSM** will then notify you as each snapshot is saved. If you've requested an existing snapshot be overwritten, you will be prompted for confirmation before doing so.

Compare Snapshots



Disk snapshots provide a detailed description of a disk partition at a given moment in time. You can use this menu item to compare the state of the current partition (snapshot) to that of one from a previous moment in time. A window will appear listing all currently mounted filesystems ([Figure 7-15](#)).

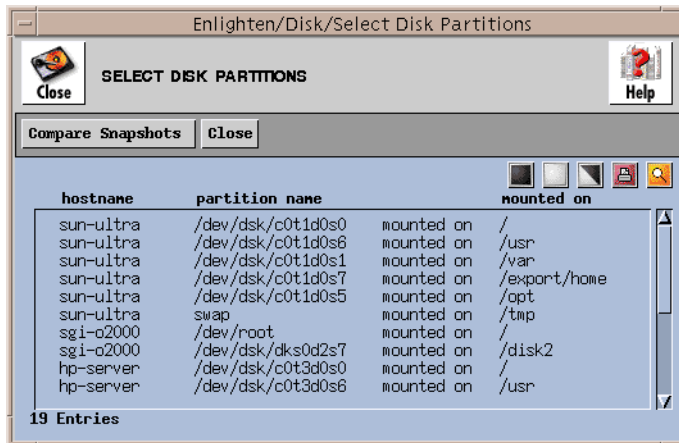


Figure 7-15 Currently mounted filesystems

Highlight the disk partitions (filesystems) for which you would like to perform a comparison and click the Compare Snapshot button. The Select Saved Snapshots window will appear ([Figure 7-16](#)). You can use this window to look at a summary of the changes, view a detailed list of the changes, or search for changes based on a query process.

Within this window, you must decide, for each partition, if the comparison should be made against the partition's master snapshot or against another named snapshot for this partition.

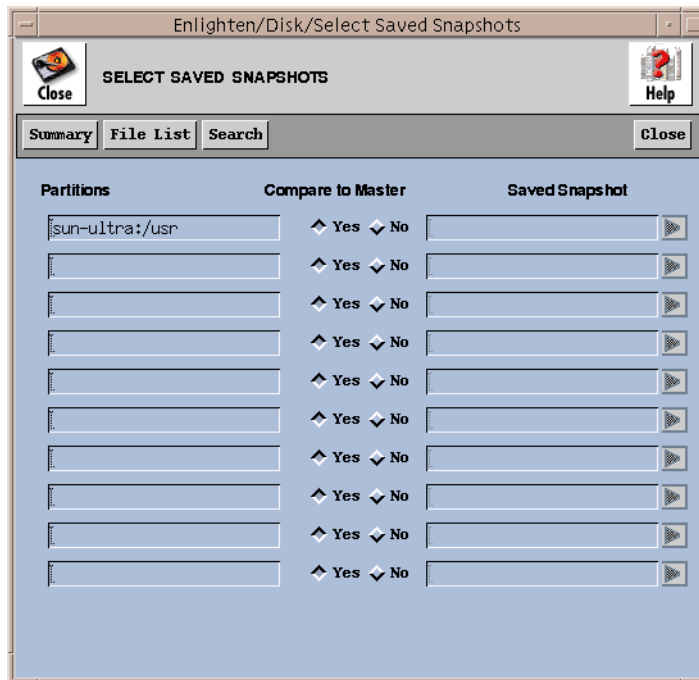


Figure 7-16 Select Saved Snapshots window

To compare against a named snapshot, click the No button under the Compare to Master column. Then, enter the name of the previously saved snapshot in the appropriate Saved Snapshot field. You can also click the arrow button on the right to bring up a window containing a list of all saved snapshots for that partition and select one of them.



Partitions of the same name, but from different systems, can be compared. For example, the / partition from the host *athens* can be compared to the / partition from the host *paris*. You can use this to help maintain consistency of static partitions (partitions that should not change much) across hosts.

Finally, you also need to choose how you want to compare the snapshots by selecting one of this window's buttons:

- Summary Get a summary of the changes
- File List Get a list of all changes
- Search Search further for certain types of changes

Summary

Click the Summary button to view a summary of the changes. Each user who had a file changed is listed in this window ([Figure 7-17](#)).

USERNAME	NEW	LOST	REN	+SIZE	-SIZE	MODE	READ	WRITE	+USER	-USER	GROUP
root	8743	638	0	0	1	1	0	19	0	0	20
bin	80	0	0	0	8	0	0	15	0	0	23
dben1	0	197	0	0	0	0	0	0	0	0	0
informix	0	86	0	0	0	0	0	0	0	0	0

4 Entries

Figure 7-17 User Files Changed

Enlighten**DSM** monitors file that are/have

- New files
- Lost files
- Renamed files
- Increased in size
- Decreased in size
- Changed modes
- Been read
- Been updated

- New owners
- New group owners

and shows in the list box how many files have changed in each category for every user displayed there.



The summary of snapshot comparisons is a time-consuming action, as are all disk comparison functions. It may take several minutes to gather all the information for large or multiple partitions.

The rest of this subsection details how to use the Summary of Changes for Disk Partitions window buttons.

View Files

By highlighting specific Username entries and clicking the View Files button, you can view a more detailed list of file changes and then manipulate the file characteristics. To use the resulting List of Changed Files window, refer to [“Type of Changes” on page 7-33](#).

Graph

To graph the disk changes, highlight the information you wish to view and then click the Graph button. A window will appear displaying the highlighted items in a graphical format.

File List

Click the File List button to see a detailed list of all files that have been changed in some way. Every file that has changed (as listed in [“Summary” on page 7-31](#)) is detailed in this window ([Figure 7-18](#)).

Often, this report consists mostly of files that have been read. Consequently, the report can be very large. You may want to first run a Summary (see [“Summary” on page 7-31](#)) or Search for changes (see [“Search” on page 7-34](#)) before using the File List option.

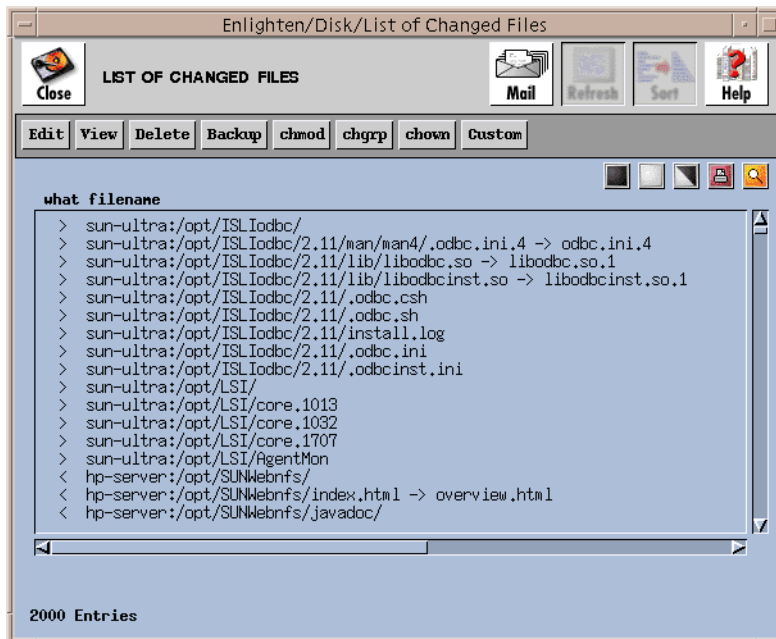


Figure 7-18 Changed Files

Type of Changes

This report is similar to that of previously described file lists, except an extra field is now included in the first column to describe the type of change that has occurred to the file.



If you've set the File list format field in your Session Preferences window to the Long Listing option, the filename in this listing is annotated to describe the change itself. For details on how to set your session preferences, refer to [“Session Preferences” on page 2-2](#).

The possible change types, their meanings, and the optional annotations (if the Long Listing option is set) are:

Code	Meaning	Annotation
>	New file	None
<	Lost file	None
>>	File increased in size	<i>(old size -> new size)</i>
<<	File decreased in size	<i>(old size -> new size)</i>
>=	File has been updated	<i>(Updated: date)</i>
==	File has been read	<i>(Last read: date)</i>
U	New user ownership	<i>(old owner-> new owner)</i>
G	New group ownership	<i>(old owner-> new owner)</i>
M	New mode (permissions)	<i>(old mode-> new mode)</i>
R	File has been renamed	<i>(old filename is now new filename)</i>

Mode changes (M) include changes to file permissions, setuid, setgid, and programs with the sticky bit turned on.

Buttons

Refer to [“View Files” on page 7-9](#) for a detailed description of how to use the List of File Changes window buttons.

Search

Since many changes can occur to a partition over time, you may only want to look for specific types of changes. Click the Search button to define search criteria for the comparison you want to make. A window will appear providing a comprehensive set of search fields ([Figure 7-19](#)).



Figure 7-19 Search fields

The rest of this subsection describes how to use this window's fields and buttons.

General Fields

Find Filenames Like

Use this field to specify filenames to use in the search pattern. The filenames can contain the standard UNIX wild card matching characters '*', '?', and '[']. This field will accept more than one entry. Leave a blank space between each unique filename for multiple entries. This is a good mechanism to use if you only want to check for changes in a certain area of a partition.

Skip Filenames Like

This is the opposite of Find Filenames Like field. The filenames entered here create an exclusion list. Any filename entered here will NOT be matched by the search. This field uses the same format as the previous field. This is a good way to ignore files that change on a regular basis and omit unnecessary entries in the report.

Owner Names

Use this field to find only those files whose owners match the names of the users entered. This field will accept more than one entry. Leave a blank space between each unique owner for multiple entries. To exclude any files owned by certain users, place a ! before those usernames. For example:

<code>root bin uucp</code>	Only files belonging to users <code>root</code> , <code>bin</code> , or <code>uucp</code>
<code>!root !bin</code>	All files except those belonging to users <code>root</code> and <code>bin</code>

Group Names

This field is similar to the Owner Names field, except you can specify user group names to use (or exclude) instead of user names.

Min. File Size

Use this field to limit the search to files larger than a certain size. The default unit is bytes, but by appending one of the letters `k`, `m`, or `g`, the corresponding unit size of kilobytes (KB), megabytes (MB), or gigabytes (GB) will be used. This can be useful when you are looking for logfiles growing out of control.

Max. File Size

This field is similar to the Min. File Size field, except you can specify the maximum file size allowed in the search.

File Type Fields

New Files

Click the Yes button (the default) to search for new files that did not exist in the previously saved snapshot.

Include contents of new/lost dirs

When EnlightenDSM searches for new or lost files and it comes across a new/lost directory, this selection determines if the contents of the new/lost directory should be included in the search. Click the Yes button (the default) if you want EnlightenDSM to look at the complete contents of the new/lost directory in this case. Otherwise, click the No button to just put the name of the new/lost directory into any report results.

Lost Files

Click the Yes button (the default) to search for files that did exist in the previously saved snapshot and now no longer exist.

Attempt to determine renamed files

Sometimes, files are renamed. Normally this is considered to be the creation of a new file and the loss of the old one. Click the Yes button (the default) if you want EnlightenDSM to attempt to determine if the file has been renamed or not.

This check is not foolproof; it will only find a file being renamed (moved) within the same directory in which it was originally found.

Altered Files

Click the Yes button (the default) to search for files that have changed since the previously saved snapshot. There are many different ways a file can change. These can be selected and/or limited by using the following fields. The default value for all of these fields is also set to Yes.

Type of Change Fields

Size

Use this file to search for files that have either increased or decreased in size and/or to search for changes within certain ranges. In the second case, you may enter the range of size changes that should be searched for in the corresponding Ranges field.

Each size range should be separated by a blank space. The range can be prefixed by a minus (-) sign to denote a decrease in the size or a plus (+) sign to denote an increase in size. If no sign is given, then an increase in size is assumed. A range has the following format:

[*+/-*]*<minsize>*[*<unit>*]-[*<maxsize>*[*<unit>*]]

The square brackets ([]) denote that the field is optional. The fields are:

+	Search for files that have increased in size
-	Search for files that have decreased in size
<i><minsize></i>	The minimum size that file has changed in size
<i><maxsize></i>	The maximum size that file has changed in size
<i><unit></i>	The unit of size, with the possible units being:
B	Bytes
K	Kilo bytes
M	Mega bytes
G	Giga bytes
%	Relative increase in size in terms of percent

The following are all valid range examples:

1000	File that increased in size by 1000 bytes
-2K	File that decreased in size by 2 Kb (2048 bytes)
10%	File that increased in size by 10 percent
100-10K	File that increased by 100 bytes but not more than 10 Kb
-10K-2M	File that decreased in size by 10Kb but not more than 2Mb

User Ownership

Click the Yes button to search for files that have changed user ownership. Sometimes you may want to focus on changes in ownership for just a few users or a specific group of users (e.g., chown to `root`). In this case, you can specify those entries in the corresponding Ranges field. Leave a blank space between each unique username or user ID (UID) for multiple entries.

Group Ownership

This field is similar to the User Ownership field, except you can specify user group names or group IDs (GIDs) to use instead of user names or user IDs (UIDs).

Mode

Click the Yes button to search for file permission changes. If you only want to look for specific mode changes (e.g., new setuid programs), you can specify those entries in the corresponding Ranges field. Leave a blank space between each set of mode changes for multiple entries.

You can use two formats to search for specific mode changes:

```
[+/-]<numeric mode>  
[+/-]<who><modes>
```

Use the first format to search for added or lost permissions followed by UNIX style octal *<numeric mode>*. Enlighten**DSM** will search for files that have had one of the specified bits changed in the file permissions, for example, use `+007` to look for any gaining of permissions for the *other* user type. Only experienced UNIX administrators should use this format.

Use the second format to search for where a type of user (*<who>*) and/or a type of mode (*<modes>*) permissions were added or lost.

The *<who>* permission values are:

- | | |
|---|---------------------------------------------------------------------------------------------------|
| u | Look for changes to the user's permissions |
| g | Look for changes to the group's permissions |
| o | Look for changes to the others' permissions |
| a | Look for changes to any of these permissions (user, group, and other). This is the default value. |

The *<modes>* permission values are:

r	Look for changes to the read permissions
w	Look for changes to the write permissions
x	Look for changes to the execute permissions
s	Look for changes to the setuid permissions
t	Look for changes to the sticky bit permissions

The following are examples of the second format:

+orwx	Look for any gaining of permissions for <i>other</i> user types
+s	Look for any files gaining setuid/setgid privileges
-ax	Look for any loss of executable permissions for all user types

Updated

Updated files are files that have been modified without having the size of the file increase or decrease. This often happens to database files where data in the middle of the file is replaced with new data. Click the Yes button if updated files should be part of the search.

Accessed (read)

Every time a file/program is read or executed, the access date is updated. Since many files are typically read or executed between snapshots, the list of read files may become quite long. Click the No button if read files should not be included in the search.

Buttons

Execute Search

The search criteria may use one or any combination of the fields described in the previous section. Once you've entered your search criteria, click the Execute Search button. The List of Changed Files window will appear with a list of the files that met the criteria ([Figure 7-20](#)).

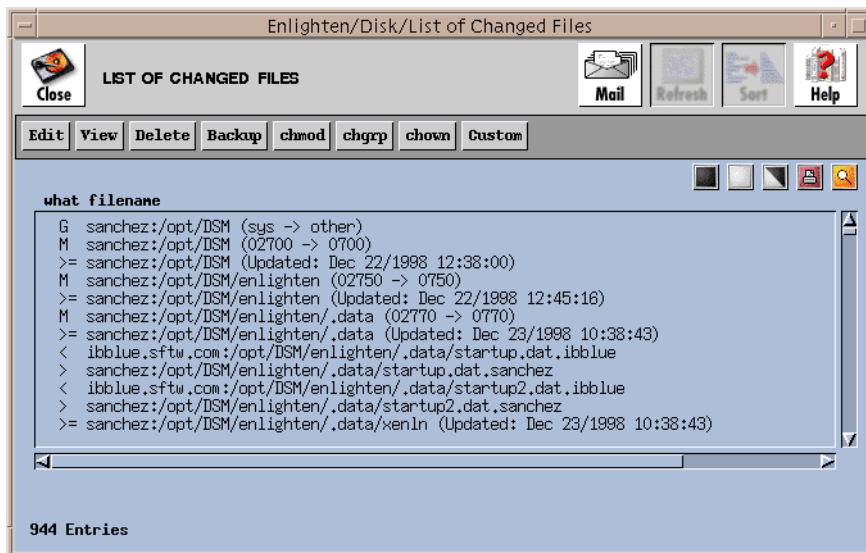


Figure 7-20 List of Changed Files window

As mentioned previously, disk snapshot comparisons are time-intensive operations that may take time to complete.


From here, you can execute more commands such as `edit`, `backup`, and `chown` to manipulate the displayed files' characteristics. See [“File List” on page 7-32](#) for more details on using this window and [“View Files” on page 7-9](#) to use its buttons.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.



Archive

The Archive menu is used to configure backup devices, erase and prepare tapes, restore files, and perform full and incremental backups.

Catalogs of backups are used by Enlighten**DSM** to quickly locate and retrieve backed-up files. This chapter also describes how to create and search those catalogs.

The Archive menu options are:

- Configure Devices
- Tape Operations
- Backup Partitions
- Restore Files
- Backup Catalogs
- Scheduled Backups

The features described in this chapter are part of the Enlighten**DSM**/Advanced product. A License Advisory window similar to the one shown on [page 1-3](#) will appear if you attempt to access this menu with the Workgroup version.

Configure Devices



You can use this module to add backup devices to the system. The device can be a custom device, a reel-to-reel tape, a 1/4-inch cartridge tape, or a raw device type. The Archive Device Configuration window (Figure 8-1) lists which devices are available for archiving data.

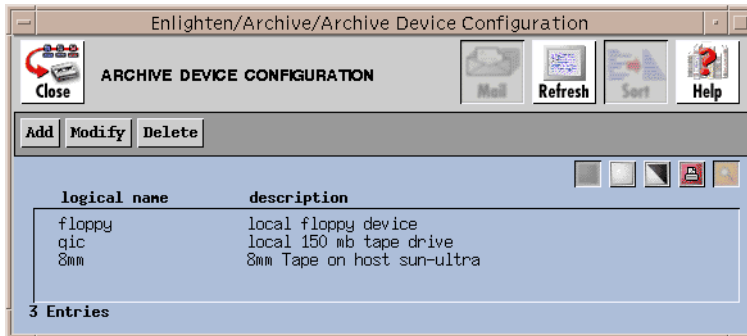


Figure 8-1 Archive Device Configuration window

From this window, you have the following options:

- Add a device
- Modify a device
- Delete a device

Add

Click the Add button to add a backup device. Enter the appropriate information in the Archive Device Add window (Figure 8-2).

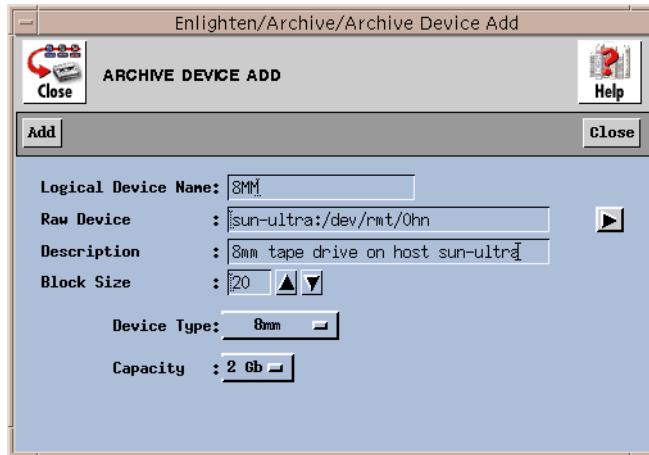


Figure 8-2 Archive Device Add window

The rest of this section describes the fields and their parameters within this window.

Logical Device Name

This field represents the common name a user will specify when accessing this backup device. This field is required.

Raw Device

In the Raw Device field, enter a UNIX device name using the form: `hostname:/dev/rmt0`. Click the right arrow button to choose from a list of possible tape device names. The device name must be preceded with its hostname, followed by a colon (:).

For example, if you have a device `/dev/rmt/0h` on host `athens`, the name of the device would be:

```
athens:/dev/rmt/0h
```

Note that this host must have an active administration daemon for a successful backup.

This field is required.

Description

This can be a friendly name for the device or a description of the device. This field is required.

Block Size

Each type of device has a blocking factor associated with it. The blocking factor is the size of each data block written to the device. The most common blocking factor is 20 blocks; however, you can back up data using a blocking factor of up to 2000 blocks. This field is required.



Block size is O/S dependent. On Solaris, the block size is 1024 bytes (1K).

Device Type

The Device Type specifies the device's capacity. This field is required. The following options are available:

- Cartridge Tape
- Reel
- Custom
- Raw
- 8 MM

Cartridge

When you click this button, tape length and density descriptions are displayed. For length and density, click the left mouse button and then drag the cursor to the appropriate value.

Reel

When you click this button, tape length and density line descriptions are displayed. For length, click the left mouse button and then drag the cursor to the appropriate value: 600, 1200, or 2400 feet. For density, click 1600 (low) or 6250 (high).

Custom

You can click the Custom button to define your own device. You will be prompted to enter the capacity for the new device in KB.

Raw

Click the Raw button to add an entry when you don't have all the information available. Eventually, you should redefine this storage device.

8 MM

Click the 8 MM button to define an 8mm cartridge tape device. You will be prompted to select the capacity of the tape. The options are 2 GB (default) and 5 GB. If neither of these options applies, create a Custom backup device instead, and manually set the device's capacity. Or you can create a Raw device and define the device capacity later.

Modify

You can use this to view and/or modify the backup devices configured on your system. Highlight any logical device name you wish to modify, then click the Modify button. A window will appear similar to the Archive Device Add window, except the Logical Device Name field is view-only.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional backup schedules if you've selected more than one backup to modify from the Add Device Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 8-3](#).

Delete

Use the delete command to remove a backup device from EnlightenDSM. Once an entry has been removed, the only way to re-add it is to recreate the entry, so use care with this command. To ensure the correct backup device is being deleted, EnlightenDSM will prompt you for confirmation before deleting it.

Tape Operations



You can use this module to check retention of a logical device or erase all data on the backup media.

This program works only if the backup media is a tape device. The Tape Operations window (Figure 8-3) displays all configured backup devices and their descriptions.

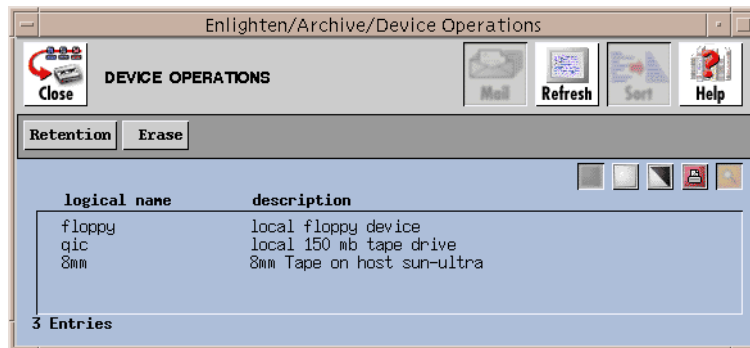


Figure 8-3 Tape Operations window

The rest of this section describes how to use the buttons in this window.

Retention

You should use this option prior to using a tape for the first time or if it has been quite a while since the tape has been used. Highlight the logical device and click the Retention button. A window will appear confirming the action.



Tape retention is only supported on a 1/4 inch tape device

Erase

You can use this program to delete all data on the backup media, for example, to erase or clean tapes. Highlight the logical device you wish to erase, then click the Erase button. A window will appear asking for confirmation of the action.

Backup Partitions



You can use this module to perform incremental or full backups, or to schedule backups on filesystems (disk partitions). You do not need to know specific UNIX backup commands; EnlightenDSM will prompt you for all the information needed to create the backup. The Backup Partitions window is shown in [Figure 8-4](#).

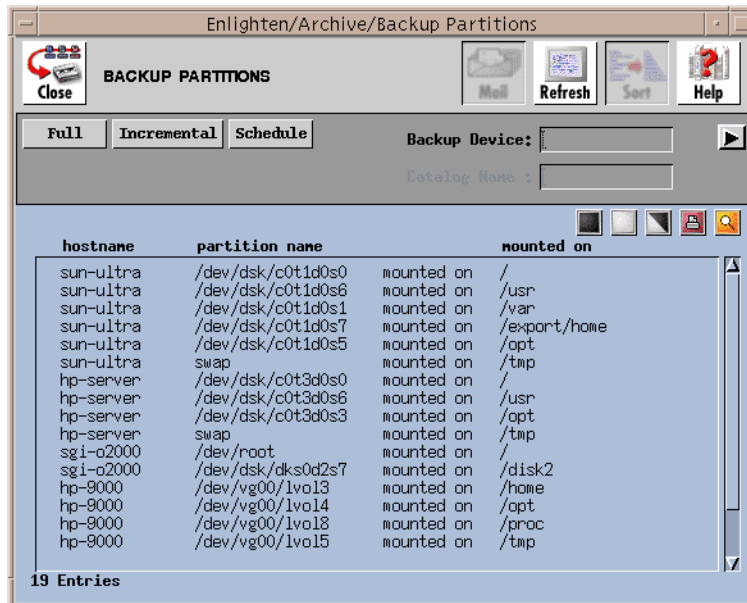


Figure 8-4 Backup Partitions window

You can have EnlightenDSM catalog the backed-up files. To do so, turn on the Catalog function in the Session Preferences menu item (see [“Session Preferences” on page 2-2](#)). These catalogs can be queried at a later time to help find instances of backed-up files. The default is NOT to keep a catalog of backed up files.

If the catalog option is enabled, you should enter a catalog name in the Catalog Name field. This can be a filename or just an identifying tag. If no name is entered, the default name of NONAME is used.

The rest of this section details the other backup options.

Full Backup


To perform a full backup:

- 1) Choose Backup Partitions from the Archive menu. A list of current local mounted partitions is displayed. If Session Preferences is set to Include NSF Parts in Partition Reports, then NFS mounted parts are also displayed.
- 2) Highlight the partitions to be backed up.
- 3) Enter the logical device name or click the right arrow button to choose from a pick list of defined backup devices.
- 4) Click the Full button. As EnlightenDSM performs the backup, a listing of each file will be displayed in the window.

Incremental Backup

To perform an incremental backup:

- 1) Choose Backup Partitions from the Archive menu. A list of local mounted partitions is displayed. If the Include NSF Parts in Partition Reports option is turned on in the Session Preferences window (Configure menu), then NFS mounted parts are displayed also.
- 2) Highlight the partitions to be backed up and enter the desired backup device. You can also click the right arrow button to choose from a pick list of defined backup devices.
- 3) Click the Incremental button. A window will prompt you for the type of incremental backup to run.



4) Choose one of the following incremental backup options:

- Since the last full backup (the default setting).
- Since the last incremental backup.
- Within a specific time period (the last n days).

Click the Apply button. As EnlightenmentDSM performs the backup, a listing of each file will be displayed in the window.

Schedule Backup

This program helps you automatically schedule a full or incremental backup. For further details on how this process works, see [“Scheduled Backups” on page 8-18.](#)

Restore Files



You can use the Restore Files window (Figure 8-5) to view the table of contents and/or extract files from an archive of a backup device. If you do not know where a specific file is located or its backup device name, you can use a search pattern here to extract the file. Furthermore, you may view or extract any tar format archive, even those not created through EnlightenDSM.

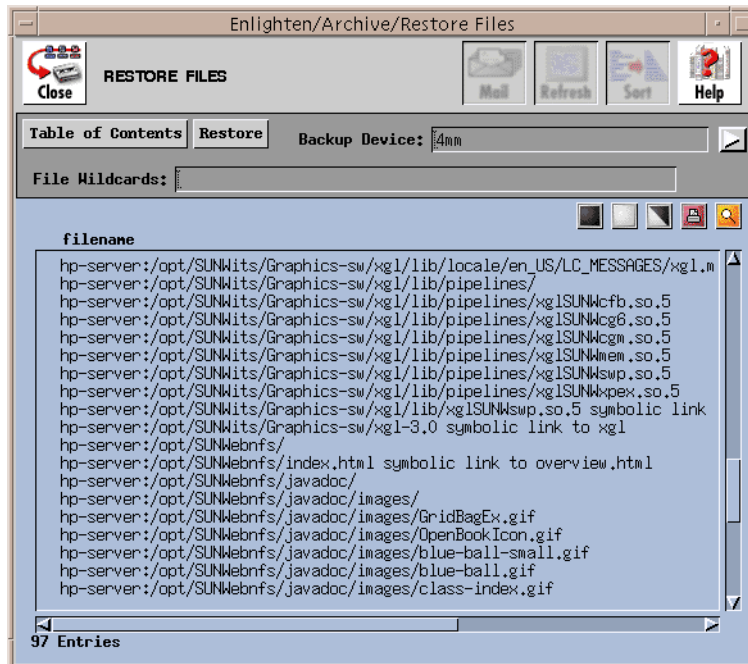


Figure 8-5 Restore Files window

The rest of this section defines the buttons and fields for this window.

Viewing an Archive

First, select the archival device's logical name as set up in the Configure Archive Devices window (see “Configure Devices” on page 8-2) by typing the name in (it must precisely match an existing name) or by selecting it from the list of available devices.

Next, if you place a file-matching pattern in the File Wildcards text field, this pattern acts as a filter to viewing an archive. For example, the pattern `*pa*` will find any file with the letters `pa` embedded in the name, such as

`/etc/passwd` and `partition.save`. If you leave this field blank, all the files in the archive are listed in the results.

Finally, click the Table of Contents button. The full archive is scanned. If EnlightenDSM encounters a multi-volume archive generated through the Backup facility, it prompts you to exchange the media. To interrupt the search, select Cancel. The results of the search will be displayed in the list box.

The list box is similar to that of the File Detail window; details about the file (its size, owner, etc.) appear at the top of the list box when you click on an entry. In this manner, you can preview which files to restore should you decide to extract files from the archive.



Large archives may take several minutes to scan.

Restoring an Archive

First, select the archival device's logical name as set up in the Configure Archive Devices window (see [“Configure Devices” on page 8-2](#)) by typing the name in (it must precisely match an existing name) or by selecting it from the list of available devices.

Then, click the Restore button to extract the contents of an archive. The extraction is based on the following conditions, tested in the order shown.

- If any entries are highlighted in the list box, only those files are retrieved.
- Otherwise, if a pattern is entered in the File Wildcards field, only those files matching that pattern are retrieved. If this pattern matches a directory, then the full directory—with all its subdirectories—is retrieved.
- Otherwise, if none of the above conditions is true, the full archive is restored.

The status of the extract procedure is displayed in a text window. You may stop the progress by pressing “Ctrl-C” (or your appropriate interrupt key) within the window. If multi-volume archives are encountered, Enlighten**DSM** will prompt you to exchange the media. At this point, you may Cancel the process, which will NOT impact those files already fully restored during this process.



Any tar files that were created without absolute pathnames are restored in the directory where Enlighten**DSM** resides.

Backup Catalog



You can have EnlightenDSM catalog the backed-up files. To do so, turn on the Catalog function in the Session Preferences menu item (see “[Session Preferences](#)” on page 2-2). These catalogs can be queried at a later time to help find instances of backed-up files. The default is NOT to keep a catalog of backed-up files.

You can use the Backup Catalog Query window (Figure 8-6) to search through the catalog of backed-up files. There are many search criteria available for finding files, including modification date and time, name, owner, and size.

Filenames longer than 120 characters will be truncated when writing to a catalog (filenames include the full path). This does not affect the ability to backup and restore files. The Restore Files feature should be used to search for backed up files.

Figure 8-6 Backup Catalog Query window

The rest of this section defines the fields and parameters for this function. If a field is left blank, all records for that criteria are searched.

Fields

This window contains the following fields:

Catalog Name

Use this field to limit the search to the specified catalog names. Make sure you enter the correct catalog names and place a space between each catalog name if you are specifying multiple catalogs.

Backup Made Since

Each catalog has a backup date and time associated with it. Specifying an entry in this field will limit the search to any backups made after the date entered. You can use one of two formats for the date. See [Appendix C, “Time Formats,”](#) for more details.

Backup Made Before

Specifying a date in this field (which uses the same format as the previous field) will limit the search to any files backed up *before* the date entered. You can also use the two Backup fields (since and before) together to specify a range of backup times in your search.

Find Filenames Like

Use this field to specify filenames to use in the search pattern. The filenames can contain the standard UNIX wild card matching characters ‘*’, ‘?’, and ‘[]’. This field will accept more than one entry. Leave a blank space between each unique filename for multiple entries.

Skip Filenames Like

This is the opposite of Find Filenames like. The filenames entered here create an exclusion list. Any filename entered here will NOT be matched by the search. This field uses the same format as the previous field.

User Owners

Every file on the system has an owner. In most cases the owner is the user who created the file. Use this field to find only those files whose owners match the names of the users entered. To exclude any files owned by a user, place a ‘!’ before those usernames. This field will accept more than one entry. Leave a blank space between each unique owner for multiple entries.

Group Owners

Every file on the system belongs to a user group. Use this field to find only those files whose group ownership matches the user group names entered in this field. To exclude any files owned by certain user groups, place a ‘!’ before those group names. This field will accept more than one entry. Leave a blank space between each unique group name for multiple entries.

File Modified Since File Modified Before

When a file is created, the filesystem stores its creation date in what is called an inode. This inode is updated every time the file is modified. You can use these fields to look for files modified before and/or since certain inode dates.

These fields work similarly to the Backup catalog fields, except these fields define a search pattern for file inode dates rather than catalog modification dates. The input format for these fields is the same as the Backup formats described earlier. See [Appendix C, “Time Formats,”](#) for more details.

Minimum File Size Maximum File Size

Every file residing on the system requires space. This space is referred to as the file size. Any input for these fields specifies the minimum and/or maximum file size a file can have as part of the search routine. If a file is smaller than the amount entered in the first field, the search routine will exclude it. Conversely, if the file is larger than the amount entered in the second field, the search routine will also exclude it.

To change the default size, enter the data in the format *N unit*, where *N* is the numeric size to search for and *unit* is one of the following units of measure:

- KB = Kilobytes
- MB = Megabytes
- GB = Gigabytes

Files to Find

This field shows two values. The first (and default) option is to find only the most recently backed-up version of the requested files. The second option is to find all instances of the requested files. Use this second option with care since it can easily generate very long lists.

Buttons

Click the Clear Fields button if you want to clear all the fields before defining and executing another search pattern. After you've finalized your search criteria, click the Execute Search button to begin the search. A list of all matching files is displayed in a new window ([Figure 8-7](#)).

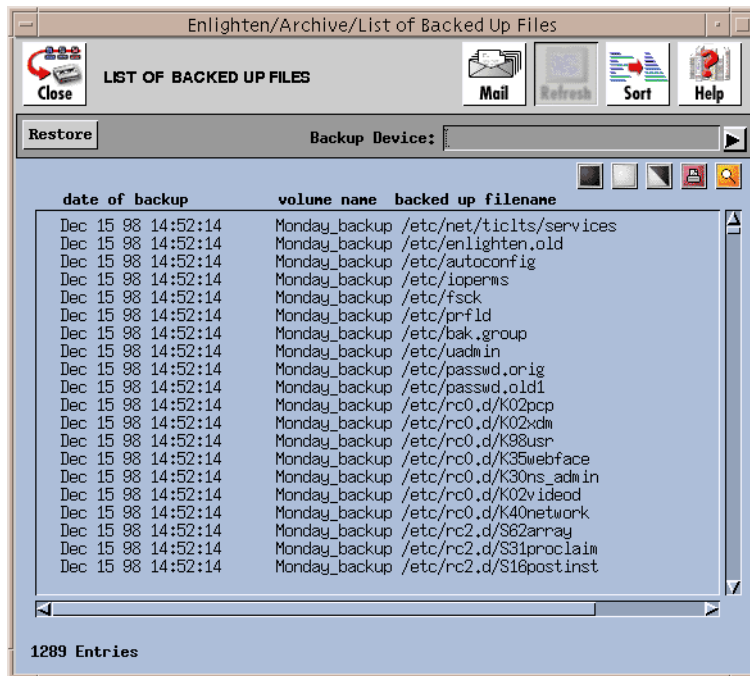


Figure 8-7 List of matching files

To restore any of the files in this list, enter the name of the backup device, highlight the desired files, and click the Restore button. If the selected files were backed up on more than one catalog, you will be prompted to load each backup catalog as required.

Scheduled Backups



You can use the Scheduled Backup Configuration window (Figure 8-8) to have the system automatically back up selected files and partitions.

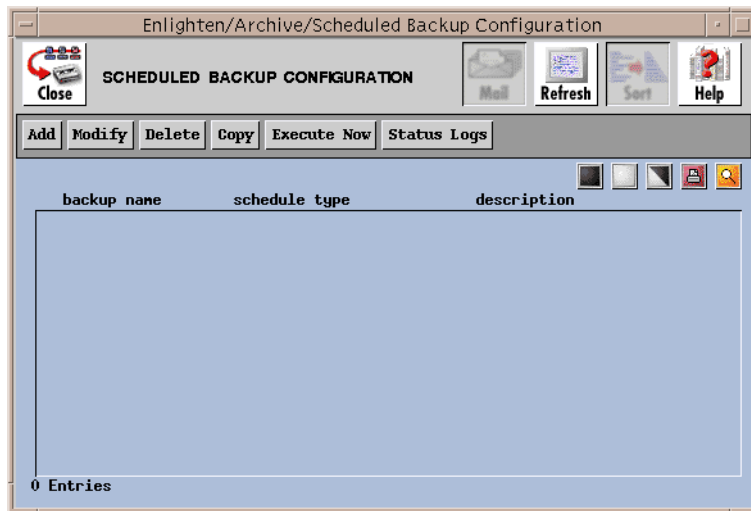


Figure 8-8 Scheduled Backup Configuration window

The Scheduled Backup Configuration window displays a list of currently defined scheduled backups, types, and descriptions. The rest of this section describes the buttons and fields in this window and other backup windows.

Add

Click this button to use the Add Scheduled Backup window (Figure 8-9).

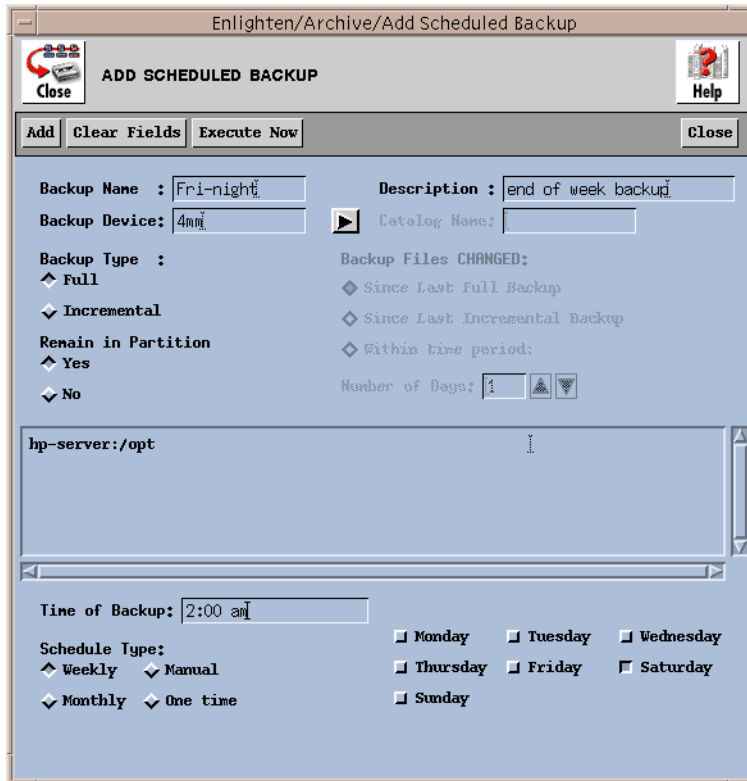


Figure 8-9 Add Scheduled Backup window

The rest of this subsection details this window's functionality.

Fields

The Add Scheduled Backup window contains the following fields:

Backup Name

Use this field to specify the backup's name.

Description

Use this field to provide a brief description of the backup.

Backup Device

Use this field to specify the logical device name for the backup. Logical device names are defined via the Configure Device window (see [“Configure Devices” on page 8-2](#)). You can also click the arrow button to the right to bring up a list of all defined device names and select one of them.

Catalog Name

If you have enabled backup catalogs, specify in this field a catalog name that will be used to keep a record of this backup. See the Catalog function in the Session Preferences menu item (see [“Session Preferences” on page 2-2](#) for more details).

Backup Type

These toggle buttons specify whether to run an Incremental or Full backup. The Incremental backup choice will also activate the Backup Files CHANGED toggle buttons.

Backup Files CHANGED

This set of toggle buttons specifies which files should be backed up during an Incremental backup. The options are files that have changed:

- Since Last Full Backup (the default)
- Since Last Incremental Backup
- Within time period

The Within time period choice will also activate the Number of Days field so you can specify a time period.

Number of Days

You can use this field to specify how many days back EnlightenDSM will check to see if a file has changed and therefore should be backed up. The arrow buttons to the right will automatically increment or decrement the number displayed in the text field.

Remain in Partition

Use this option to select whether you want the backup to cover only the immediately defined partition for a particular directory (Yes) or all partitions listed under that directory (No). The default is Yes.

File List

Use this list box to enter which files will be checked for changes. Those that have changed will then be backed up as specified in the preceding fields.

Time of Backup

You can use this field with the Schedule Type toggle buttons to determine when a backup should be executed. See the following Schedule Type field description for more details on how this field is used.

To specify a time/date value, see [Appendix C, “Time Formats.”](#)

Schedule Type

This set of toggle buttons determines how and when the backup will occur. The choices are:

- **Weekly** The days of the week will be displayed. Select which days the backups should occur. Enter when the backup will occur in the Time of Backup field.
- **Monthly** The days of the month will be displayed. Select which days the backups should occur. Enter when the backup will occur in the Time of Backup field.
- **Manual** This option does not schedule a time to perform the backup; you run it by selecting the Execute Now button in the Scheduled Backup Configuration parent window.
- **One time** Specify a (one time) date and time to run the backup in the Time of Backup field.

To specify a time/date value, see [Appendix C, “Time Formats.”](#)



Backups are scheduled using the local GUI's host Cron capabilities. The Cron job will be owned by the user specified by the Local Cron User parameter in the EnlightenDSM Params file located on the local GUI host.

Buttons

The Add Scheduled Backup window contains the following buttons:

Add

Click this button to create the scheduled backup defined by the rest of this window's fields.

Clear Fields

Click this button to clear the existing choices in all fields.

Execute Now

Click this button to execute the highlighted backup immediately. This occurs regardless of the scheduled backup time.

Close

Click this button to close the window without saving any scheduled backup information.

Modify

Click this button to modify the schedule of a backup configuration. A window similar to the Add Scheduled Backup window will appear, except the Backup Name field is now a display-only field.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and

- You can use the Next button to modify additional backup schedules if you've selected more than one backup to modify from the Scheduled Backup Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 8-19](#).

Delete

Click this button to delete the highlighted backup from the Scheduled Backup Configuration list. To ensure the correct backup schedule is being deleted, Enlighten**DSM** will prompt for confirmation before deleting it.

Copy

Click this button to copy the highlighted backup schedule. This action will pop up the Add Scheduled Backup window and initialize it with the values from your highlighted selection. Change the field values as needed and click the Add button to create the new copy backup schedule.

See [“Add” on page 8-19](#) for details on the buttons and fields in the Add Scheduled Backup window.

Execute Now

Click this button to execute the highlighted backup immediately. This occurs regardless of the backup's defined schedule time.

Status Logs

When backups are run (in the background) at their scheduled times, they create a backup log. Click the Status Logs button to access these status logs in the Backup Logs Configuration window ([Figure 8-10](#)).

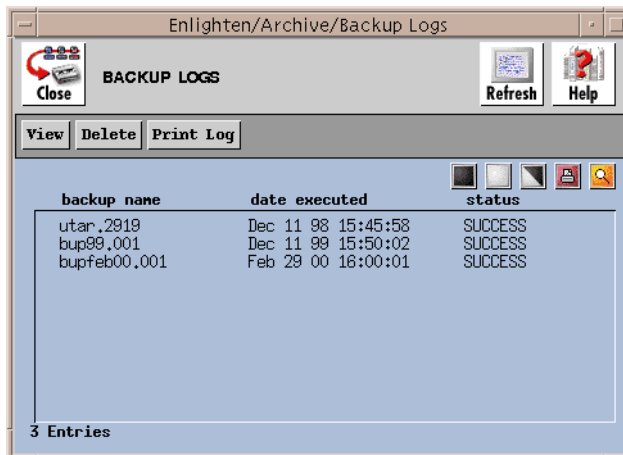


Figure 8-10 Backup Logs Configuration window

The window contains the following buttons:

View

Click this button to view the selected backup logs.

Delete

Click this button to delete the selected backup logs. EnlightenDSM will prompt you to confirm your action.

Print Log

Click this button to print a copy of the selected backup logs. The output is sent to the printer defined in the Print field in the Configure —> Session Preferences window. See "[Session Preferences](#)" on [page 2-2](#) for further information.



System

Enlighten**DSM**'s System menu contains numerous tools for basic network management. With the System menu, administrators can view CPU summaries for any given TTY port, monitor and manipulate swap space processes, shut down or reboot the system, manage crontab entries, change a host's clock settings, maintain and update mail aliases, manage files and directories, and disallow or allow logins.

The System menu options are:

- CPU Summary by TTY
- Swap Space
- Shutdown/Reboot
- Remote File Distribution
- Cron Management
- Clock Management
- Mail Aliases
- Login Status

The features described in this chapter are part of the Enlighten**DSM**/Advanced product. A License Advisory window similar to the one shown on [page 1-3](#) will appear if you attempt to access this menu with the Workgroup version.

CPU Summary by TTY



You can use this program to view a summary of all process activity by TTY (Figure 9-1).

hostname	tty	username	procs	cpu time
sun-ultra	??	(TTY)	3	0:04
sun-ultra	console	root	1	0:00
sun-ultra	pts/10	root	1	0:00
sun-ultra	pts/11	root	2	0:31
sun-ultra	pts/12	root	1	0:00
sun-ultra	pts/13	root	2	0:04
sun-ultra	pts/14	terryo	1	0:00
sun-ultra	pts/5	(TTY)	5	1:08
sun-ultra	pts/7	root	1	0:00
hp-server	??	(TTY)	2	0:03
hp-server	console	root	1	0:00
hp-server	pts/2	(TTY)	4	0:04
hp-server	pts/3	root	4	0:47
hp-server	pts/4	root	2	0:55
hp-server	ttyp0	(TTY)	3	0:00
hp-server	ttyp1	(TTY)	3	0:00
hp-9000	console	root	1	0:00
hp-9000	ttyp2	root	1	0:00
hp-9000	ttyp3	root	2	0:02
Total	(20 Users on 3 hosts)		41	1:12

20 Entries

Figure 9-1 Summary of Processes by TTY window

The information shown is the TTY number, the login user, the number of processes running on the TTY, and the CPU time used. You can further investigate one or more of the TTYs listed by using the Processes button. The rest of this section describes how to use the buttons in the window.

Processes

To view the processes, highlight the TTYs you wish to view and click the Processes button. A window will appear displaying all processes for the highlighted TTYs. To further manipulate this information, see [“Process Status” on page 3-34](#).

Graph

To graph the processes, highlight the information you wish to view and then click the Graph button. A window will appear displaying the amount of CPU utilization on each highlighted TTY in a stacked bar graph format. The left axis shows the amount of processes.

Swap Space



You can use this program to check the swap space usage ([Figure 9-2](#)).

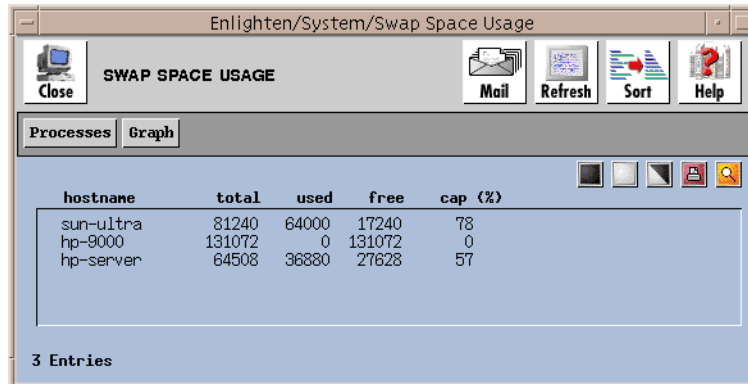


Figure 9-2 Swap Space Usage window

The information shown is the total amount of swap space in KB, the amount used, the amount available, and the amount used as a percentage. The rest of this section describes how to use the buttons in this window.

Processes

To view the processes, highlight the hosts you wish to view and click the Processes button. A window will appear displaying all processes for the highlighted hosts. To further manipulate this information, see [“Process Status” on page 3-34](#).

Graph

To graph the swap space usage, highlight the information you wish to view and then click the Graph button. A window will appear displaying the amount of free and used swap space for each highlighted host in a stacked bar graph format. The left axis shows the number of KB of swap space.

Shutdown/Reboot



You can use this program to either shut down or reboot all hosts in the current pool (Figure 9-3).

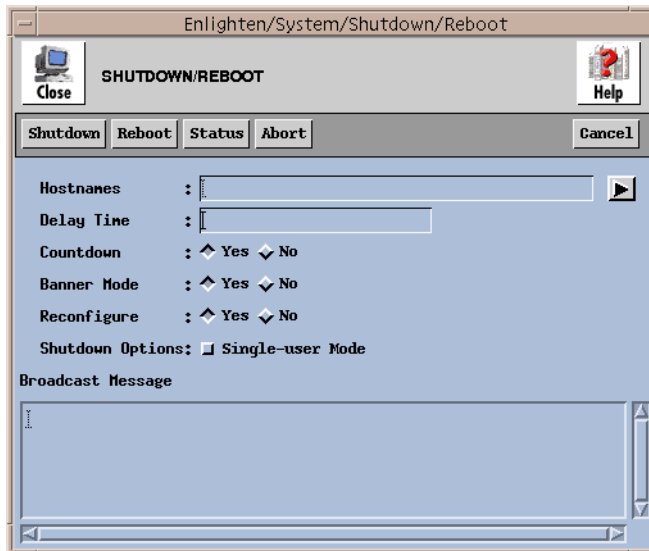


Figure 9-3 Shutting down and rebooting hosts

The rest of this section details how to use this window's functionality.

Fields

This window has the following fields:

Hostnames

Use this field if you want to perform the shutdown/reboot operation only on a subset of the currently managed hosts. You can also click the arrow button to the right to bring up a pick list of all currently managed hosts and make your selections. Leave a blank space between each hostname if you are specifying multiple entries.

If you leave this field blank, all hosts in the current pool will be shutdown/reboot.

Delay Time

You can use this field to specify a delay before the shutdown/reboot occurs. The delay time should give users enough time to prepare for the system shutdown/reboot. The default is 60 seconds.

This field will accept both relative and absolute time formats. See [Appendix C, “Time Formats,”](#) for more details.

Countdown

Use this toggle to choose whether a countdown should occur prior to starting the shutdown or reboot. A countdown will display warning messages at regular intervals to the terminals of all hosts being taken down or rebooted. The default is Yes.

Banner Mode

Use this toggle to choose whether the (user-definable) broadcast message should be displayed in extremely large text. This message should be relatively brief. The default is Yes.

Reconfigure

Use this toggle to choose whether the hosts should be reconfigured when they are rebooted. The default is Yes.

Shutdown Options

Use this toggle to choose whether the hosts should be put in an initialized state of a single-user mode when they are shut down. The default is Off. This choice only matters if you then use the Shutdown button.



Some systems, like Solaris 2.x, will cut power to the CPU during a shutdown.

Broadcast Message

You can use this text field to type in the message to be broadcast to all hosts prior to shutdown. Otherwise, the default message “System Going Down in *n* Minutes” will be used.

Buttons

This window has the following buttons:

Shutdown

Click this button to selectively begin to shut down the host specified in the Hostnames field. You will be prompted for confirmation of this action. If you leave the Hostnames field blank, all hosts in the current pool will be shut down.

Reboot

Click this button to selectively begin to reboot the host specified in the Hostnames field. You will be prompted for confirmation of this action. If you leave the Hostnames field blank, all hosts in the current pool will be shut down.

Status

Click this button to display the status of all pending shutdowns and/or reboots. The following window will appear (Figure 9-4).

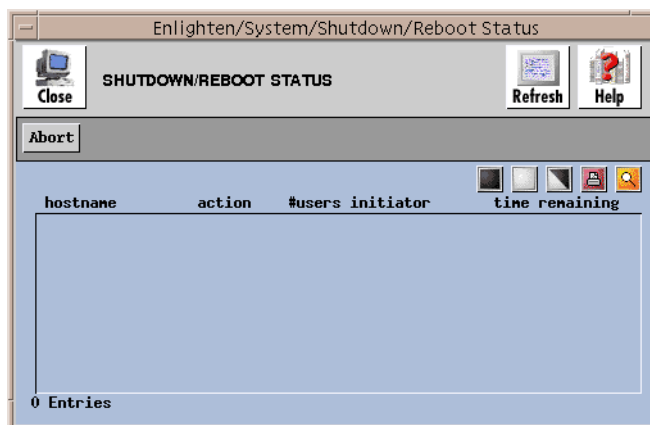


Figure 9-4 Status of pending shutdowns or reboots

You can then highlight any pending shutdowns and/or reboots and click the Abort button. Enlighten**DSM** will abort the pending shutdowns and/or reboots you selected.

Abort

Click this button to abort all pending shutdowns and/or reboots immediately. You will be prompted for confirmation of this action. Enlighten**DSM** will then abort all pending shutdowns and/or reboots on hosts in the current pool.

Cancel

Click this button to close the window without making any changes.

Cron Management



You can use this module to manage crontab entries on multiple heterogeneous hosts. You can add, change, delete, or copy cron jobs; perform a query on all cron jobs; or manage user access to crontab. The actual options are:

- Configure
- Query
- Cron Users

The rest of this section details how to use each of these options.

Configure

You can use this program to view all crontab jobs for all hosts in the current pool. The Crontab Configuration window will appear (Figure 9-5). The list box will show the current cron jobs for all users on all hosts in the current pool.

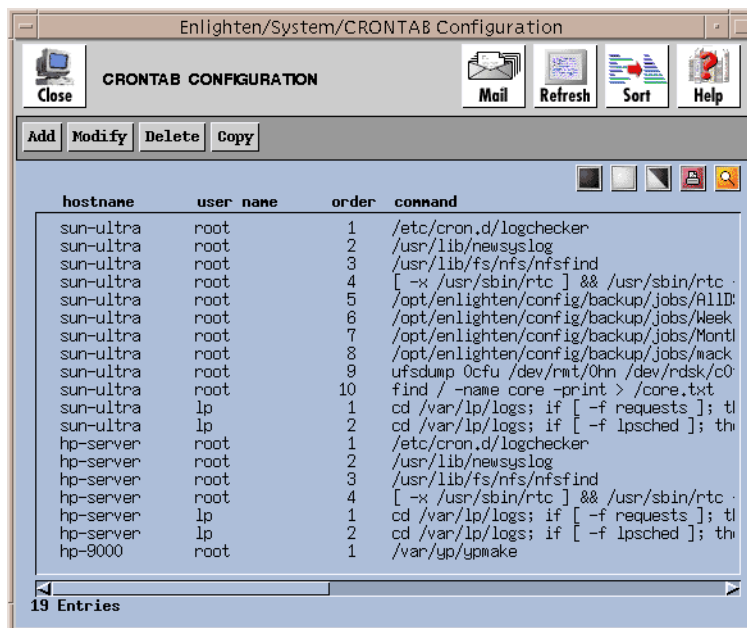


Figure 9-5 Crontab Configuration window

From here, you have the option to:

- Add Create a new crontab job
- Modify Modify an existing crontab entry
- Delete Selectively delete all marked crontab jobs
- Copy Create another cron job using the settings of the first selected cron job

Add

Click this button to create a new crontab job. The Add Crontab Entry window will appear (Figure 9-6).

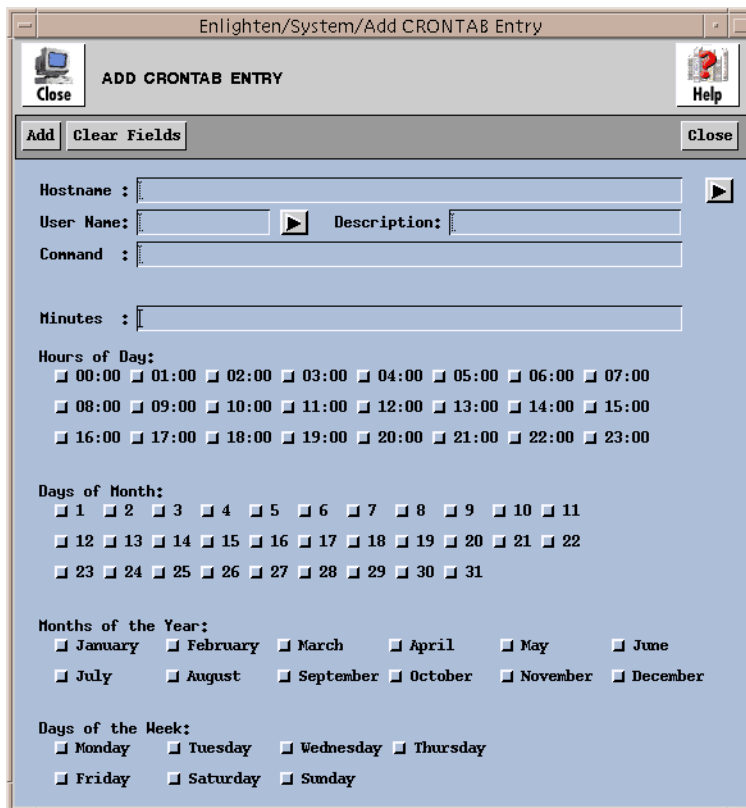


Figure 9-6 Add Crontab Entry window

A crontab entry consists of six fields. Any combination of the five time fields (minutes, hours, days, months, and/or days) defines when the command (sixth field) will execute. By default, the job will run every hour if all the time fields are left blank.



This feature is system dependent; you may not be able to use all combinations of time settings on your system. Check your local cron for its precise workings.

Add Fields

The Add Crontab Entry window has the following fields:

Hostname

If you want to limit this job to specific hostnames within a pool, enter those hostnames in this field. If you are using multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select the available hosts from the current pool.

User Name

You can use this field to specify which user will run this job. You can also click the arrow button on the right to display a list of users and select one.

Description

You can use this field to briefly describe this job's purpose.

Command

Use this field to specify the cron command (and arguments) to be run at the time established by the remaining five fields. You can use a maximum of 256 characters to specify this command.

Minutes

Use this field to specify what point in the specified hour(s) the job will run. Enter a whole number between 0 and 59. Or you can specify a range by using a dash (-), for example, 0-20. If you are using multiple entries, leave a blank between each entry.

Hours of Day
Days of Month
Months of the Year
Days of the Week

Click on the appropriate boxes in these fields to further define when this job will run. You may select multiple values per field.

Add Buttons

This window has the following buttons:

Add

Once you've made all the selections for the new job, click this button to selectively create the crontab job.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Modify

Click this button to modify a selected crontab entry. A pop-up window similar to the Add Crontab Entry window will appear, except the User Name field is view-only.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional jobs if you've selected more than one job to modify from the Crontab Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 9-10](#).

Delete

Click this button to delete a job from the Crontab Configuration list. EnlightenDSM will prompt you to confirm your action. EnlightenDSM will not delete a Cron job for a non-existent user.

Copy

Click this button to copy the settings in a cron job. The Add Crontab Entry window will appear showing the highlighted job's settings in each field. You can edit this window as needed and then click the Add button to complete the copy.

See [“Add” on page 9-10](#) for a description of how to use this window's fields and buttons.

Query

You can use this program to find specific cron jobs. You can search for hostnames, user names, the command itself, the time settings, and so on. Once your query is successful, you can then modify, delete, or copy the cron job.

When you activate this program, the window in [Figure 9-7](#) will appear.

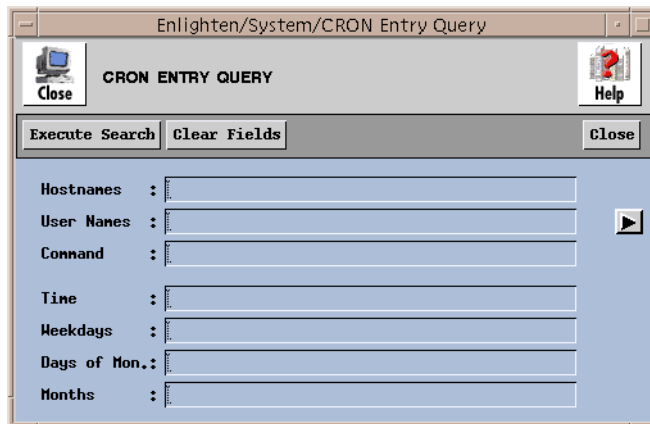


Figure 9-7 Finding cron jobs

The cron job must match all fields to get a successful match. The fields in this window have a logical AND relationship. Blank fields always match and you can also negate fields with the ! character (e.g., !Monday).

Fields

This window contains the following fields:

Hostnames

If you want to limit the search to specific hostnames within a pool, enter those hostnames in this field. If you are using multiple entries, leave a blank between each entry.

User Names

If you want to limit the search to specific users, specify which one(s) in this field. If you are using multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select from a pick list of the users.

Command

Use this field to limit the search for a specific cron command (or any of its options). You can use a maximum of 256 characters in this field, including the standard UNIX wildcards.

Time

Weekdays

Days of Mon.

Months

You can use these fields to specify times or time ranges. This limits the search to find jobs executing at/between the times entered in these field. See [Appendix C, "Time Formats,"](#) for more details on using these fields.

Buttons

This window contains the following buttons:

Execute Search

Once you've selected your search criteria, click the Execute Search button. If you click this button without filling in any of the fields, all current cron jobs will be displayed. A window similar to the Crontab Configuration window appears with the results listed.

Then you can use the following options to act further on the results:

- Modify
- Delete
- Copy

See [“Configure” on page 9-9](#) for more details.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Cron Users

cron has built-in security features allowing it to control which users are allowed to execute cron jobs. There are five possible user access states:

- Only root is allowed to execute cron jobs
- Only selected users are allowed to execute cron jobs
- All users are allowed to execute cron jobs
- All users except specified users are allowed to execute cron jobs
- No users are allowed to execute cron jobs

You can use this program to display the cron security status of all hosts, as shown in [Figure 9-8](#).

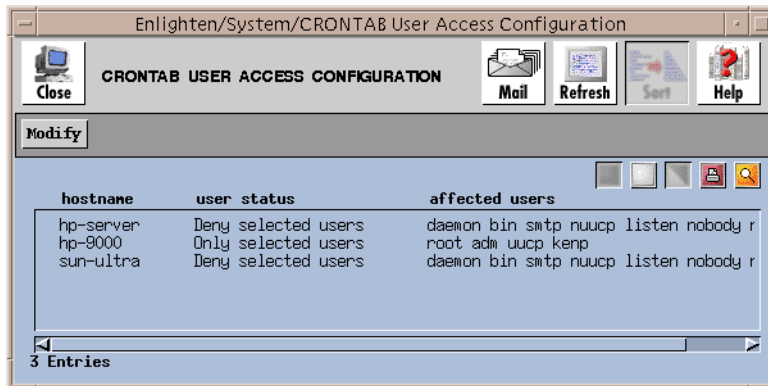


Figure 9-8 Cron security status of hosts

From here, you can highlight which users' accessibility you want to change and click the Modify button. This action will bring up the Modify Crontab User Access window (Figure 9-9). You can then use this window to modify the cron user access for multiple hosts. The rest of this subsection details how to use the Modify Crontab User Access window.

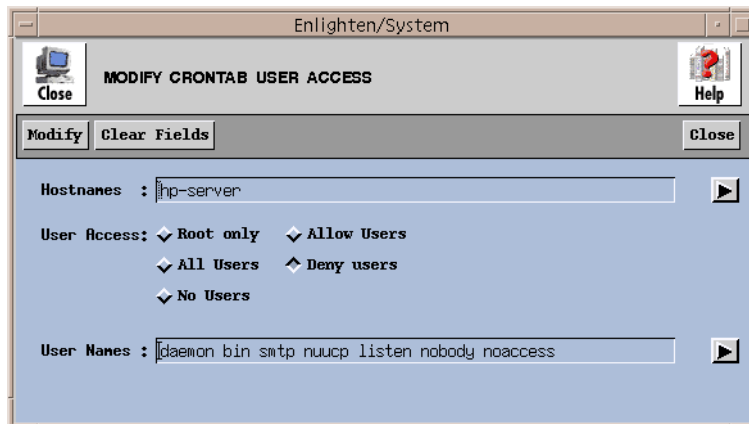


Figure 9-9 Modify Crontab User Access window

Fields

This window contains the following fields:

Hostnames

If you want to limit this modification to specific hostnames within a pool, enter those hostnames in this field. If you are using multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select the available hosts from the current pool.

User Access

Use this set of toggles to specify what type of user cron access you want to allow.

User Names

If you want this modification to allow or exclude (deny) specific users, enter the user name(s) in this field. If you are using multiple entries, leave a blank between each entry. You can also use the arrow button to the right to select from a pick list of the users.

Buttons

This window contains the following buttons:

Modify

After making your selections, click the Modify button to make the crontab user access modifications.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Clock Management



You can use this module to synchronize or set a host's clock. The options are:

- Clock Synch
- Clock Set

The rest of this section describes how to use each of these options.

Clock Synch

You can use this program to choose when to synchronize a host's clock on the time kept by a master clock. When you activate this program, the window in [Figure 9-10](#) will appear.

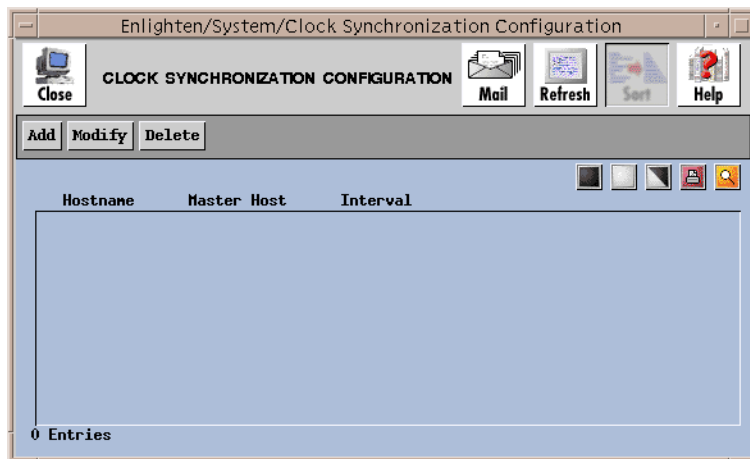


Figure 9-10 Synchronizing a host clock

From here, you have the option to:

- Add Create a clock synchronization process
- Modify Modify an existing clock synchronization
- Delete Selectively delete all marked clock synchronization processes

The rest of this subsection details how to use this window's features.

Add

Click the Add button to set up a synchronization process. The window shown in [Figure 9-11](#) appears.

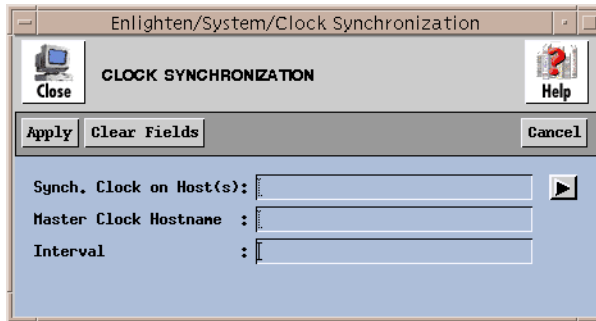


Figure 9-11 Setting up a synchronization process

Add Fields

This window contains the following fields:

Synch. Clock on Host(s)

Use this field if you want to limit the clock synchronization to a subset of the currently managed hosts. If you are using multiple entries, leave a blank between each entry. You can also click the arrow button to the right to bring up a pick list of all currently managed hosts and make your selections from there.

Master Clock Hostname

Use this field to specify which host will act as the master timekeeper.

Interval

You can use this field to specify when the clock synchronizations should occur. This field will accept both relative and absolute time formats. See [Appendix C, “Time Formats,”](#) for more details.

Add Buttons

This window contains the following buttons:

Apply

Once you've made your selections, click this button to establish the synchronization for all selected hosts' clocks in the current pool. You will be prompted to confirm your action.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Modify

Click this button to modify an existing clock synchronization process. A pop-up window similar to the Clock Synchronization window will appear, except:

- You also have access to a pick list option in the Master Clock Hostname field, and
- You can use the Next button to modify additional jobs if you've selected more than one process to modify from the Clock Synchronization Configuration list.

For a description of the rest of the buttons and fields in this window, see ["Add" on page 9-19](#).

Delete

Click this button to delete a process from the Clock Synchronization Configuration list. EnlightenDSM will prompt you to confirm your action.

Clock Set

You can use this program to reset a host's clock. When you activate this program, the window shown in [Figure 9-12](#) will appear. The rest of this subsection details how to use this window's functionality.

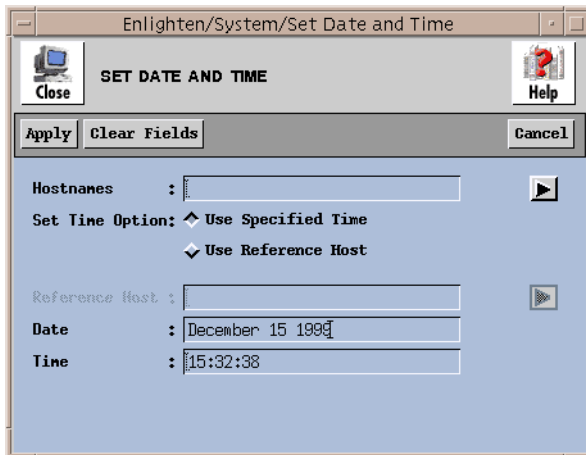


Figure 9-12 Resetting a host clock

Fields

This window contains the following fields:

Hostnames

Use this field to specify which hosts in the current pool will have their clocks reset. If you are using multiple entries, leave a blank between each entry. You can also click the arrow button to the right to bring up a pick list of all currently managed hosts and make your selections from there.

Set Time Option

Use these toggles to choose whether the new time setting should be based on a specific time (default) or a reference host's clock setting.

Reference Host

If you set the Set Time Option toggle to Use Reference Host, use this field to specify which host's clock setting should be referenced for the correct time. You can also click the arrow button to the right to bring up a pick list of all currently managed hosts and select one of them.

This can be *any* host, but it cannot be in the current pool.

Date Time

If you set the Set Time Option toggle to Use Specified Time, you can use these fields to specify the date and time. The default values are the current date and time. See [Appendix C, "Time Formats,"](#) for more details on using these fields.

Buttons

This window contains the following buttons:

Apply

After making your selections, click this button to reset the clock for all selected hosts' clocks in the current pool. You will be prompted to confirm your action.

Clear Fields

Click this button to clear the existing choices in all fields.

Cancel

Click this button to close the window without making any changes.

Mail Aliases



You can use this module to manage mail aliases. You can add, change, delete, view, copy, or perform a query on all mail aliases. The actual options are:

- Configure
- Query

The rest of this section details how to use each of these options.

Configure

When you create user accounts, you can have a mail alias created for those accounts and/or have them included in a mail list. Use this program to set up the mail lists and mail aliases.

You can also create these when you create user accounts. See [“Add” on page 3-38](#) for more details.

When you activate this program, the window shown in [Figure 9-13](#) will appear.

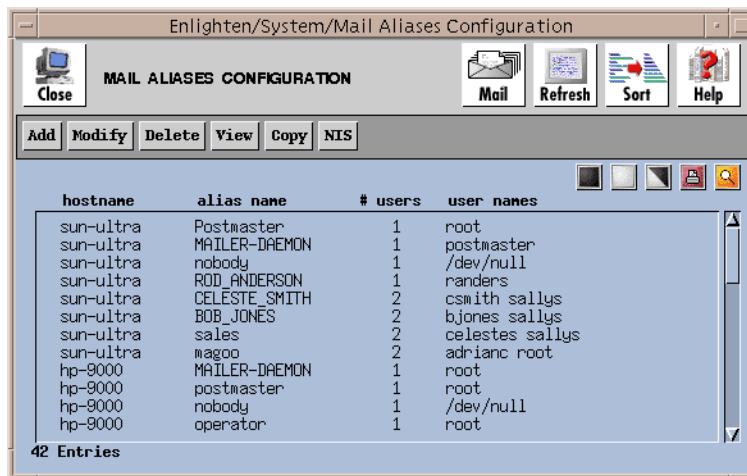


Figure 9-13 Configuring mail aliases

This window displays a list of all mail aliases for the hosts in the current pool. If the alias is a mail list (has more than one entry), the number of users in the list is displayed; otherwise, the user name and the respective Realname are displayed.

From here, you have the option to:

- Add Create a new mail alias
- Modify Modify an existing mail alias (on multiple hosts)
- Delete Delete an existing mail alias from a host
- View View a summary of the mail alias
- Copy Copy a selected mail alias to one or more hosts
- NIS Push the NIS maps on any active NIS servers

The rest of this subsection details how to use the Mail Aliases Configuration window's features.

Add

Click the Add button to create a new mail alias. The window in [Figure 9-14](#) appears.

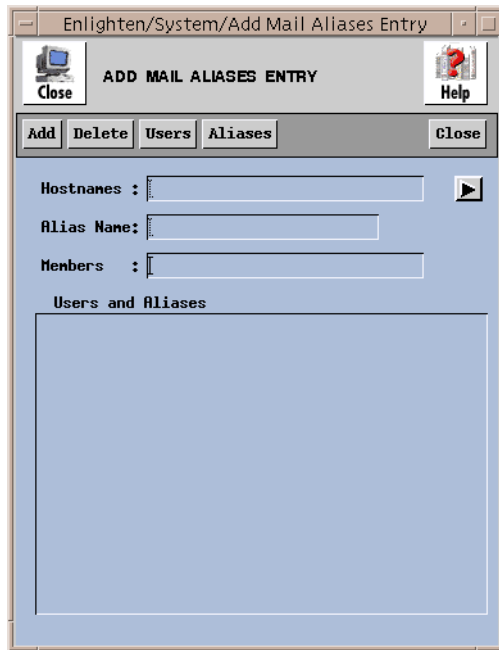


Figure 9-14 Creating a new mail alias

Add Fields

This window has the following fields:

Hostnames

You can use this field to specify which hosts will have their alias list updated. If you are using multiple entries, leave a blank between each entry. You can also click the arrow button to the right to bring up a pick list of all the available hosts and make your selections from there.

Alias Name

Use this field to specify the name of the mail list alias.

Members

Use this field to specify which users and/or other aliases are members of *this* mail list. If you are using multiple entries, leave a blank between each entry. You can also use the Users and/or Aliases buttons to make your selections.

Add Buttons

This window has the following buttons:

Add

Once you've made all the selections for the new mail alias, click this button to create it.

Delete

Click this button to delete a mail alias from the Mail Aliases Configuration list. EnlightenDSM will prompt you to confirm your action.

Users

Click this button to pop up a pick list of all recognized users and select the one(s) you want to be members of this mail alias.

Aliases

Click this button to pop up a pick list of all currently known mail alias lists and select the one(s) you want to be included in this mail alias list.

Close

Click this button to discard any changes and close the window.

Modify

Click this button to modify an existing mail alias entry. A pop-up window similar to the Add Alias Entry window will appear, except the Alias Name field is view-only.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional jobs if you've selected more than one job to modify from the Crontab Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 9-25](#).

Delete

Click this button to delete a highlighted mail alias from the Mail Aliases Configuration list. Enlighten**DSM** will prompt you to confirm your action.

View

Click this button to view a summary of the mail alias list. The Mail Alias View window will appear ([Figure 9-15](#)). The list box will display a list of each user and/or alias in this mail alias list.

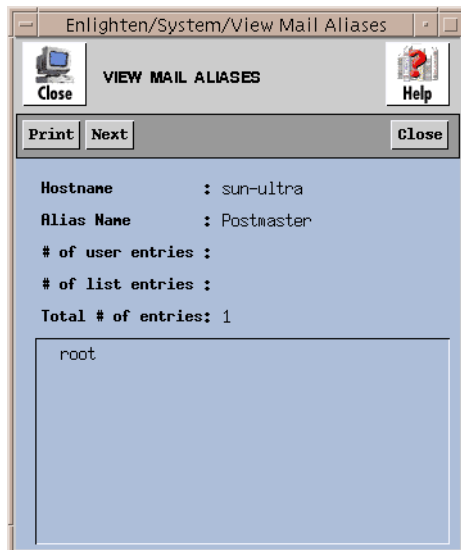


Figure 9-15 Mail alias summary

View Fields

This window contains the following (display-only) fields:

Hostname

The name of the host(s) that use this alias.

Alias Name

The name of the mail alias.

of user entries

The number of direct users in the alias.

of list entries

The number of direct mail lists in the alias.

Total # of entries

The total number of effective users in the alias.

User names

This list box shows which users are a part of this alias.

View Buttons

This window contains the following buttons:

Print

Click this button to print the mail alias information.

Next

If you selected multiple aliases in the Mail Aliases Configuration window, click this button to view the next highlighted alias.

Close

Click this button to finish viewing the mail aliases and close the View window.

Copy

Click this button to copy the settings in one mail alias list to a second list. The Mail Aliases Configuration window will appear showing the highlighted mail aliases settings in each field except the Alias Name field. You can edit this window as needed and then click the Add button to complete the copy.

See [“Add” on page 9-25](#) for a description of how to use this window’s fields and buttons.

NIS

Click this button to push/remake the NIS maps on any active NIS servers. This updates the NIS servers and hosts on your network with any Host Entry changes you’ve made to the Configuration list. If there are no NIS servers to update for your choice(s), Enlighten**DSM** will display a dialog box telling you so.

Query

You can use this program to find specific mail aliases. You can search for hostnames, alias names, user names, and so on. Once your query is successful, you can then modify, delete, view, or copy the mail alias.

When you activate this program, the window in [Figure 9-16](#) will appear.

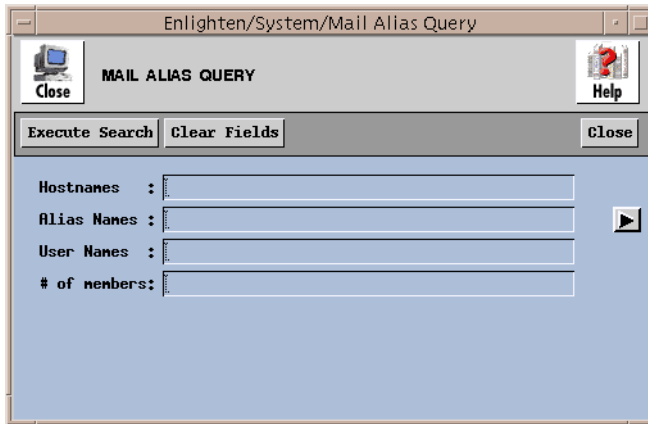


Figure 9-16 Finding mail aliases

This search must match all fields to get a successful match. The fields in this window have a logical AND relationship. Blank fields always match. You can also negate fields with the ! character (e.g., !Joe).

Fields

This window contains the following fields:

Hostnames

To limit the search to specific hostnames within a pool, enter those hostnames in this field. If you are using multiple entries, leave a blank between each entry.

Alias Names

To limit the search to specific mail aliases, enter those names in this field. If you are using multiple entries, leave a blank between each entry. You can also click the arrow button to the right to bring up a pick list of all the recognized mail aliases and make your selections from there.

User Names

To limit the search to specific users, enter those user names in this field. If you are using multiple entries, leave a blank between each entry.

of members

Use this field to specify how many members are in the alias. Enter a whole number between 1 and 100. Or you can specify a range by using a dash(-), for example, 10-20.

Buttons

This window contains the following buttons:

Execute Search

After selecting your search criteria, click the Execute Search button. If you click this button without filling in any of the fields, all mail aliases will be displayed. A window similar to the Mail Aliases Configuration window appears with the results listed.

Then you can use the following options to act further on the results:

- Modify
- Delete
- View
- Copy
- NIS

See [“Configure” on page 9-23](#) for more details.

Clear Fields

Click this button to clear the existing choices in all fields.

Close

Click this button to discard any changes and close the window.

Remote File Distribution

EnlightenDSM also supports Remote File Distribution. You can create files and directories; set file and directory ownerships and permissions; and run downloaded scripts that return TTY output and program exit codes. When you choose this command, the Remote File Distribution Configuration window appears (Figure 9-17).

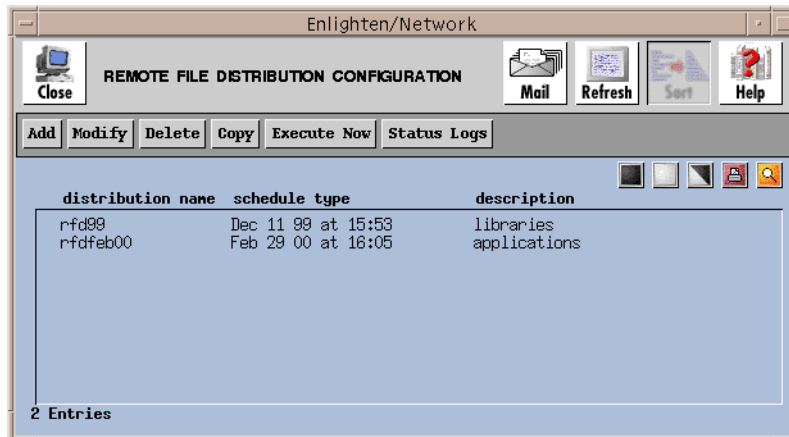


Figure 9-17 Remote File Distribution Configuration window

From here you can:

- **Add** Create a file distribution job
- **Modify** Modify an existing file distribution job
- **Delete** Delete selected file distribution jobs
- **Copy** Copy an existing file distribution job
- **Execute Now** Execute (run) the first selected file distribution job now regardless of the actual scheduling status
- **Status Logs** View the Remote File Distribution Logs

The rest of this section details how to use this window's functions.

Add

Click the Add button to define a new file distribution job. The Add Remote File Distribution Job window will appear (Figure 9-18).

Figure 9-18 Add Remote File Distribution Job window

The rest of this section details how to use this window's features.

Fields

This window contains the following fields:

Job Name

Use this field to assign a name to the remote file distribution job.

Description

Use this field to give a brief description of the job. This is used along with the name to help identify jobs in the status logs.

Install to Hosts

You can use this field to limit which hosts will be accessed during the job. Leave a blank space between hostnames for multiple entries. You can also click the arrow button to the right to pop up a pick list of all hosts in the current pool with the selected O/S type(s) and make your selection(s) from there.

Target Directory

Use this field to specify the name of the directory where the files should be copied. Do not use hard or soft links in this field.

Min. Disk

You can use this field to define the minimum amount of free disk space the partition containing the target directory must have before this remote file distribution job will execute. You can specify the units in bytes or Kbytes.

Pre-Install Script

You can use this field to specify any scripts or programs that should be executed on the destination hosts before transferring the files. Enter the full pathname to the script and, optionally, an input data filename. Multiple scripts should be separated by semicolons.

Source Directory

You can use this field to specify the directory where the Pre-Install Script will be put and then run. Do not use hard or soft links in this field.



Remote File Distribution fails if the source directory is read-only.

Post-Install Script

You can use this field to specify any scripts or programs that should be executed on the destination hosts after the files are transferred. Enter the full pathname to the script and, optionally, an input data file name. Multiple scripts should be separated by semicolons.

Directory

You can use this field to specify the directory where the Post-Install Script will be put and then run.

Source Directory

You can use this field to specify the directory where this job will be started on the local host.

Number of Retries

Use this field to choose how many times a remote copy to a given host should be attempted before the copy operation fails. This can be 0 (the default) through 5 times.

Rollback on

You can use this toggle to select if the copied files should be removed (or not) from the receiving hosts when an error occurs during the remote copy process. The choices are to remove the copied files from:

- Ignore Errors (the default; no copied files are removed)
- Failed Hosts Only (remove copies from failed hosts only)
- All Hosts (if one host fails, remove copies from all hosts)

Overwrite Files

Use this toggle to specify if existing files in the destination directory should be overwritten by the new files being copied AND if new directories should be created as needed during the distribution job. The default is Yes.

Use Relay Host

Use this toggle to specify if this distribution of files should be broadcast to any Relay Hosts for distribution to hosts on other subnets. The default is No.

If you set this to Yes, you should also specify the Broadcast Relay Hosts in the Session Preferences window. See [Chapter 2, “Configure.”](#) for more details.

Distribution Files

This field is an editable scrolled list. You can click in this list once to activate it and then type in the full pathname of the files you want to be copied to other hosts.

Time of RFD Job

You can use this field, along with the Schedule Type field, to schedule when a file distribution should be executed. See [Appendix C, “Time Formats,”](#) for details on specifying this value.

Schedule Type

Use this toggle to specify how and when the file distribution will occur. The choices are:

- **Weekly** When you click this option, the days of the week will be displayed. Select the days of the week the file distribution should occur. The Time of RFD job field should contain the time of day the file distribution will occur.
- **Manual** This is the default setting. You can run this type of job by clicking the Execute Now button.
- **Monthly** When you click this option, the days of the month will be displayed. Select the days of the month the file distribution should occur. The Time of RFD job field should contain the time of day the file distribution will occur.
- **One time** Use this option to specify in the Time of RFD job field a one-time date and time to run the file distribution. See [Appendix C, “Time Formats,”](#) for details on specifying this value.

Buttons

This window contains the following buttons:

Apply

Click this button to create the file distribution process.

If you didn't enter the Target Directory and Distribution Files information, EnlightenDSM will prompt you to supply the missing information. If any of the Distribution Files do not exist, EnlightenDSM will pop up an error box informing you which files could not be found and the remote distribution will stop. Otherwise, a status window will pop up to show the progress of the remote file distribution.

Clear Fields

Click this button to clear the existing choices in all fields.

Execute Now

Click this button to execute the file distribution job immediately without having to save it or schedule it. See [“Execute Now” on page 9-38](#) for more details.

Close

Click this button to discard any changes and close the window.

Modify

Click the Modify button to modify an existing file distribution job. A window similar to the Add Remote File Distribution Job window will appear, except you cannot modify the Job name field.

There are also two button differences in the Modify window:

- You can use the Modify button (rather than the Add button) after you've made all your changes, and
- You can use the Next button to modify additional remote distribution jobs if you've selected more than one to modify from the Remote File Distribution Configuration list.

For a description of the rest of the buttons and fields in this window, see [“Add” on page 9-33](#).

Delete

Click this button to delete the selected file distribution jobs. EnlightenDSM will prompt you to confirm your action.

Copy

Click this button to copy an existing file distribution job and create a second job. The Add Remote File Distribution Job window will appear showing the highlighted job's settings in each field except the Job name field. You can edit this window as needed and then click the Apply button to complete the copy.

See [“Add” on page 9-33](#) for a description of how to use this window's fields and buttons.

Execute Now

Click this button to execute (run) the first file distribution job you've selected now, regardless of its actual scheduling status. The progress of the job distribution will be displayed in an xterm.

Status Logs

When file distribution jobs are run (in background) at their scheduled times, a log is created showing their execution time and completion status. Click the Status Logs button to view this log. The Remote File Distribution Logs window will appear ([Figure 9-19](#)).

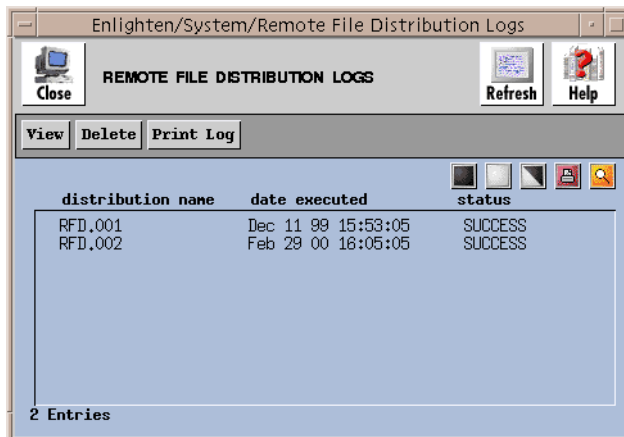


Figure 9-19 Remote File Distribution Logs window

This window will display a list of available file distribution logs. From here, you can use the window's buttons to perform the following operations:

- **View** View the backup log in a window
- **Delete** Delete the selected backup logs
- **Print Log** Print (make a hardcopy) a backup log

The rest of this subsection details how to use each of this buttons.

View

Click this button to view the contents of the selected status log in a separate window. Click the Enlighten**DSM** icon when you are done viewing the log.

Delete

Click this button to delete the selected distribution job status logs. Enlighten**DSM** will prompt you to confirm your action.

Print Log

Click this button to print the selected distribution job status logs to your default printer. Enlighten**DSM** will prompt you to confirm your action.

Login Status



You can use this menu choice to allow or disallow logins. The second option stops any users from gaining access to the system. This only works on systems that recognize the `/etc/nologin` mechanism. Typically you use the second option during system maintenance or when a large job is running that requires most of the CPU's resources.

Highlight the desired action and EnlightenDSM will prompt for confirmation, if needed. For example, if you highlighted Allow Logins and no confirmation appeared, you are already allowing logins. Otherwise a pop-up window will appear and you can click OK to execute the desired action.



The only way to remove a temporary lockout from the system is to go back to the System menu and choose Allow Logins. Or you can reboot the system.

10

Events

The Events menu is used to create and monitor event tests. The Events feature can be configured to collect and save status, configuration, performance, and capacity information. Warning messages can then be automatically sent to system or network managers when an alarm condition occurs. Events can also take action under specific conditions using a process you specify.

The information collected by Events is also used by the Status Map GUI to graphically display the status of all systems managed on your network. The map identifies which hosts and/or pools are having problems, allowing network managers to anticipate errors and take quick action.

The Events menu options are:

- Configure
- Status Map

See also Chapter 5, “Monitoring Network Systems With Events” of the *EnlightenDSM User Guide*.

The features described in this chapter are part of the EnlightenDSM/Advanced product. A License Advisory window similar to the one shown on [page 1-3](#) will appear if you attempt to access this menu with the Workgroup version.

Configure



The Events Configuration features allow you to set up tests that collect information about your systems that you can use to assist in tasks such as system tuning, load balancing, resource planning, and upgrade analysis. This chapter describes the basics of building testtab files.

For an overview of events capabilities and features and information about using the Events Configuration window (Figure 10-1) to add, modify, and delete tests using the Events graphical user interface, refer to Chapter 5, “Monitoring the Network,” in the *EnlightenDSM User Guide*.

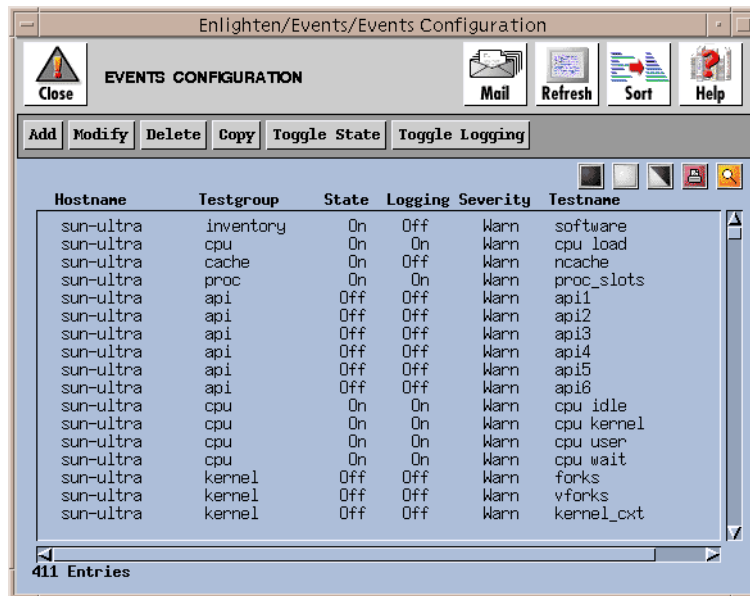


Figure 10-1 Events Configuration window

Events is a distributed systems management feature that provides for the unattended monitoring of your systems. It provides extensive automated data collection for use by both local System Administrators and Network Managers. Events can help you predict when a problem is about to occur, where it will occur, and report the event while taking user-definable corrective action.

How EVENTS Works

You can use the data collected by Events to assist in tasks, such as:

- System Tuning
- Load Balancing
- Resource Planning/Justification
- Upgrade Requirement Analysis

Events collects this data by monitoring the following:

- Memory Subsystems
- Individual Files
- Directory Queues
- File Systems
- Printer Queues
- Critical Processes
- Network Statistics
- Hardware Inventory
- Software Inventory
- User Provided Data

An appropriate message can be sent to Network Managers, System Managers, or both when an alarm condition occurs. Events can send alarms using one or more of the following methods:

- SNMP (Simple Network Management Protocol) trap messages
- Email
- PEP (Programmable Events Processor) messages

Events can also pass the alarm to a process you've defined for possible corrective action. You can specify the same process for all tests or a separate process for each test.

Inventory Tracking

One of the unique features of Events is its hardware inventory tracking mechanism. At system start-up, Events assembles an inventory record, including hard disks, tape drives, RAM, network interfaces, and software where applicable. If a list exists from a previous start-up, then the two lists are compared. Additions and deletions are reported via email and to the EMD.

On most systems, Events also includes a software inventory. The software inventory process is done similarly to the hardware inventory tracking mechanism. You can set how frequently the software inventory is generated in the `testtab` file. See [“The testtab.<hostname> File” on page 10-11](#) for more details.

Communications

Events communicates to SNMP management systems via SNMP. As an SNMP agent, Events initiates error messages and alerts, and provides information to the SNMP management system. By responding to inquiries from the SNMP management system, Events makes workstation monitoring an interactive process.

Practical Use

For the System Manager, Events is a process that runs in the background and looks after the system. The System Manager can specify whether or not measured values are stored to a database, who should be notified if a test fails, and how to notify someone when a test fails for each test that is performed. Also, other tests can easily be added to the built-in tests suite.

For the Network Manager, Events is an SNMP agent that emits enterprise-specific traps to notify the appropriate Network Manager of a failed test condition. All tests are also manageable via SNMP. MIB II (Management Interface Base) is also supported.

Alarm Thresholds

Each Events test measures some numeric value. It is easier to figure out which alarm threshold is the most appropriate if you know what the value represents. The following examples use tests that are the most often misunderstood.

The paradigm is:

- 1) Select a test.
- 2) If there is an alarm threshold specified or if logging is enabled:
 - Execute the test (go and count something).
 - Log the count if it changed significantly from the last logged value.
 - Compare the count with each alarm threshold.
 - Send alarms if any of the thresholds were exceeded.
- 3) Schedule the time of the next test.

File Clamping

This test looks in ASCII log files for the recent addition of message types you have defined. Though an alarm message may contain text from the monitored file, the threshold itself is not the message text; it's the number of matched string entries found.

Since the test value (count) is the number of matches found (regular expression matches), you must set one of the alarm thresholds if you desire an alarm. By setting the High Level limit equal to one, an alarm occurs every time an offending message is found. Only the first instance of any given threshold is reported.

Let's say, for example, that you want to look for the presence of a hardware failure by "scanning" the `/var/adm/syslog`. The value matches should be set to 1. An alarm will occur as soon as Events detects the presence of the desired message string.

In another instance you may not be concerned about a particular message per say, but rather the frequency of a particular message, such as the number of hardware error messages in `/var/adm/syslog` within a given period of time. Set the Threshold value to the number of messages to look for and the test frequency to the amount of time in which that number of messages constitutes an alarm condition.

File Accessed

This test checks the time stamp associated with the file. The time stamp is actually a number. If the number increases, the file has been accessed. The number should never decrease. A decrease in value would be suspicious and may indicate a security breach.

A value of 1 for the Pos Jump Threshold is sufficient to constitute a change in the timestamp being checked.

Process Instances

This test counts the number of processes containing the same name that are running.

Events are not sent for the first and last instance of a process. A low threshold value of zero tests for the absence of a process. Setting the low threshold value higher than the high threshold will disable the low threshold testing.

This test counts the number of processes running as specified by the name value used in **testname** and arguments field. Minimally, the testname must be the filename portion of the binary used as the process to test. Use the **arguments** field to further refine process matching.

Process Size

This test looks at the size of processes you've specified. Some platforms provide identification of the process in long name (fully qualified pathname) as well as short name (process name) formats as specified in the **Testname** field of the Add or Modify Events Test window. Use the **arguments** field to further refine process matching.

Components

Events consists of the following parts:

- **AgentENL**—An SNMP Agent for sites not currently running a multi-MIB SMUX (SNMP Multiplexer) compliant agent.
- **AgentMon**—A subagent required for each workstation you want to monitor.

- **EventsCli**—A command line interface syntax to the Events management framework. You can use this to generate events from your third-party applications.

Component Relationship

AgentMon performs all tests. When an alarm condition occurs, AgentMon may notify someone via traditional methods, such as email specified by the local user, and/or it may also notify network management via an enterprise-specific SNMP trap.

AgentENL is an interface between the AgentMon SubAgent and the network management application, such as SunNet Manager or HP OpenView. AgentENL must be started before AgentMon starts, if you intend to use the SNMP interface.

You can also use EventsCli with any current monitoring scripts and programs to handle any special needs that AgentMon can't cover. By calling EventsCli from the alarm notification section of your script, you gain much greater flexibility and control over how the alarm is treated. EventsCli will send an SNMP trap, it will log the alarm to the EMD, and it will notify PEP of its activity. Refer to [Appendix F, "Events Commands,"](#) for more information.

Standards Compliance

Events complies with the following Internet standards:

RFC 1155	Structure and identification of management information of TCP/IP-based internets: MIB-II
RFC 1157	Simple Network Management Protocol (SNMP)
RFC 1212	Concise MIB definitions
RFC 1213	Management Information Base for network management of TCP/IP-based internets: MIB-II
RFC 1215	Convention for defining traps for use with the SNMP
RFC 1227	SNMP MUX (SMUX) protocol and MIB

Configuration

Events has three different types of stored tests you can configure:

- **Group Tests**

Each of these tests is predefined with a purpose and a name, such as `cpu load`, `cache_usr`, or `kernel_traps`. See [“Group Tests” on page 10-14](#) for more information.

- **Item Tests**

These tests are `File`, `Directory`, and `Processes`. Each of these three tests types is predefined with a purpose and may have further subcategories of tests. Each test name will be the name of the particular item being tested.

You can have as many of each of the three types of tests as you want. For example, you could have 18 `File` tests, 24 `Directory` tests, and 26 `Processes` tests. See [“General Tests” on page 10-31](#) for more information.

- **API Tests**

There are six of these tests; each is predefined with a name and nothing else. You can use these test names, `api1` through `api6`, to create your own tests with shell scripts, SQLs, or even compiled programs, and incorporate them into Events. See [“API Tests” on page 10-37](#) for more information.

Configuring Tests

All the information that comprises an Events test is described in (a file called) the `testtab.<hostname>` file as an entry. Each entry must begin with a test name, and the subsequent lines must contain the test’s parameters. Tests are not required to have any associated alarm thresholds; you may, for instance, want to enable logging.

A test is ignored if:

- There are no alarm thresholds and/or
- Logging is not enabled.

You can disable a test by using the `off` parameter. AgentMon modifies the `testtab.<hostname>` file whenever the configuration is changed via SNMP or the Events GUI.

If the default `testtab` values are acceptable, the simplest test entry need only contain the following:

- A valid test name AND
- An alarm threshold or the logging is enabled.

Tests, or entries, can be added, modified, or deleted using any editor or the Events GUI. Try to use the Events GUI where possible. See the section “Monitoring Systems with Events” in Chapter 5, “Monitoring Network Systems with Events,” of the *EnlightenDSM User Guide*.



Deleting the entry for a built-in test will not turn the test off, but merely makes that test run with the default values. To turn a test off, reconfigure the entry by changing the `on` capability to `off`.

In reality, the `testtab.<hostname>` file need only contain changes from the default settings.

Data Logging

Logging to the EMD (Enterprise Management Database) is enabled by specifying the parameter `log` for each test from which you wish to collect data. Also, you can specify a value for `delta` for each test you want to log. Logging will then occur if the previously logged value varies by more than `delta` units from the current value. Setting `delta=0` results in the measured value always being logged.

Delta

You can use `delta` to set your tests so they monitor more frequently, log less often, and still not lose any logging information. This is most useful in one of two situations:

- Monitoring an object whose value seldom changes
- Monitoring an object whose value has more precision than a log needs

An example of the first case: Monitoring a file that seldom changes size

With traditional logging (log every measurement), if you monitor this file once per minute and its size changes only once each hour, the other 59 out of 60 loggings are identical. With `delta` logging, you can specify that logging should only occur if the value changes significantly.

In this case, you could specify you want logging to occur only if the file size changes by one or more bytes (set `delta=1`). Now the log only contains the *same* two critical points (and not the other 58 instances of the repeated measurement).

An example of the second case: Monitoring an object whose value has a greater precision than is required

Suppose you are monitoring a filesystem's free space and logging the data so you could later use the data to help predict when to buy additional storage. If you set `delta=1000`, then logging will only occur when the available disk space changes by `>1000 units` from the previously logged value. In effect, you are specifying the resolution of your data.

Alarm Messages

If the test generates an alarm message, logging will always occur (even if the test has logging disabled). AgentMon also logs itself each time it is started, restarted (warm start), or normally terminated.

Automated Corrective Action

You can use a command parameter, which is a pathname, in the `testtab` entry for a test to perform some automated corrective activity. The pathname is assumed to be the name of a user-provided program. When an alarm condition occurs, the named process is started with the following arguments (listed in order):

- 1) Test name
- 2) Value
- 3) Unit of measure
- 4) Alarm type

- 5) Time of measurement
- 6) Total Measured Value

Examples of corrective action scripts can be found in the subdirectory `$ENLIGHTEN/contrib`.

You can also include the `pep` capability in the `testtab` entry to notify PEP of the alarm.

The `testtab.<hostname>` File

The file `$ENLIGHTEN/config/testtab.hostname` contains the entries defining AgentMon's current configuration. If this file does not exist, it will be created upon start-up. If AgentMon's configuration is changed via SNMP or the Events GUI, this file will be rewritten to reflect those changes. If the file is manually edited, AgentMon will do a "warm" restart and configure itself in accordance with the new contents.

The rest of this chapter details the basics of building a `testtab` file. You can also use the Events GUI to easily build and alter the configuration of any of your tests. See the section "Monitoring UNIX Systems with Events" in Chapter 5, "Monitoring Network Systems With Events," of the *EnlightenDSM User Guide*.

Configuration File Parameters

Each test, or `testtab` entry, is composed of a test name and optional parameters that serve as "keywords" to define the scope and behavior of the test. The parameters are:

Name	Type	Default	Comment
testfreq	int	5	Test interval, in minutes.
alarmfreq	int	60	Minimum time in minutes between alarms. For process monitoring, the default is zero.
command	str	nonw	Pathname of any processes to start when an alarm occurs.

Name	Type	Default	Comment
mailer	str	/bin/mail	Program that will deliver the alarm if a username was given for 'notify'.
notify	str	root	Where alarms are sent. May be set to 'nobody' for no notification, except on Solaris systems.
log	boolean	false	Indicates logging is enabled. Alarms are always logged.
!log	boolean	true	Turns off logging.
delta	int/float	0	If logging is enabled, the most recent value measured will be recorded if it differs by at least this amount from the previous value.
pep	boolean	varies	Notify PEP when an alarm occurs.
!pep	boolean	varies	Do not notify PEP when an alarm occurs.

The following is an example test entry using some of these parameters.

```
cpu load |:\
:on:testfreq=1:alarmfreq=60:mailer=/bin/mail:notify=root:log:\
:high=5.0:units=units:delta=50:pep:command=/policy/myproc:\
:/* end cpu load */:
```

Refer to [Appendix G, “Sample Events Files,”](#) for a sample of a complete testtab file.

Alarm Thresholds

Threshold types used in the testtab.<hostname> file can be specified as absolute values, percentage change, or incremental change based on the unit of measure for the particular test. Any combination of threshold types can be set. AgentMon checks the current values for each test and compares the measured value against these thresholds and sends alarms based on the following order of precedence:

Name	Type	Default	Comment
high	int/float	0	Absolute value—high-level alarm threshold.
low	int/float	0	Absolute value—low-level alarm threshold.
+rate	float	0.0	Percent change value—positive rate change threshold.
-rate	float	0.0	Percent change value—negative rate change threshold.
+jump	int/float	0	Incremental value—positive shift threshold.
-jump	int/float	0	Incremental value—negative shift threshold.
age	int	0	For monitoring directory queues only, in minutes.

Most tests measure integer values; others measure floating point values. Ensure that the values supplied for threshold types are consistent with the **Unit of Measure** for the particular test being configured. Failure to do so, however, will only result in rounding errors.

The jump and rate thresholds compare the current test value with the last measured value (change over time). For process monitoring, the process built-in first/last alarm is checked first, then the alarm thresholds are checked in the order listed in the preceding table (see [“Processes Tests” on page 10-36](#) for more details).

Group Tests

This section describes all the modifiable tests that can appear in the `testtab.<hostname>` file. Each test name is listed with the corresponding Events MIB Group name and a brief description.

Refer to [Appendix H, “O/S Compatibility,”](#) to determine if any particular test is supported on your operating system.

O/S Tests

The following table shows the CPU tests.

Test Name	Test Group	Description
cpu load	CPU	<p>CPU load average, one minute average. This shows the average number of jobs in a run queue. Default Alarm: high=5.0</p> <p>This is a relative indication of how busy the system is. Some slowness in system response may be noticed when the load exceeds a value of approximately 5.0. Interactive use, like editing, can become aggravating when loads are heavy. A low value of cpu load factor is preferred. No value is too low. A high value indicates that the system is being overworked. The most common causes are:</p> <ul style="list-style-type: none"> Too many people logged on One or more CPU intensive programs running System needs more RAM Something causing an excessive number of interrupts <p>When any cpu load alarm threshold has been breached, a TRAP PDU will be sent to the Network Management System (NMS) if the NMS has SET this trap to ENABLE (see the TrapManage Group). The TRAP message will include the current value of cpu load factor.</p>
cpu user	CPU	Percentage of time spent handling user processes.
cpu idle	CPU	Percentage of time the CPU is idle.
cpu kernel	CPU	Percentage of time spent handling system processes.
cpu wait	CPU	Percentage of time spent waiting for I/O procedures to complete.

The following Kernel Group tests may be of occasional usefulness to your local UNIX performance expert and can also be an excellent troubleshooting aid for certain kinds of problems. These tests do not normally need to be active.

Test Name	Test Group	Description
kernel_cxt	Kernel	Number of kernel context switches since the last reboot.
kernel_traps	Kernel	Number of kernel traps since the last reboot.
kernel_syscalls	Kernel	Number of kernel mode system calls made since the last reboot.
kernel_devints	Kernel	Number of device interrupts since the last reboot.
forks	Kernel	Number of Forks since the last reboot.
vforks	Kernel	Number of VForks since the last reboot.
fork_pages	Kernel	Number of Forked Pages since the last reboot.
vfork_pages	Kernel	Number of VForked pages since the last reboot.

File System Tests

The following table shows the File System tests.

Test Name	Test Group	Description
/<fs> blocks free	File System	<p>Amount of free space on each logical (partition) disk drive. For example: /usr blocks free</p> <p>For each file system, the default low threshold limit for blocks free value is set to 10% of that file system's size (in 512-byte blocks). Blocks free refers to the number of blocks available on the disk. There will be as many tests as there are filesystems. These tests are created dynamically.</p> <p>Each filesystem is automatically discovered at program start-up time. Alarm thresholds are also automatically computed. Since near-full disks are typical, the default low-level limit for this test will be adjusted if the filesystem is already in an alarm condition at installation. When the amount of free space has breached an alarm threshold, a TRAP PDU will be sent to the NMS if the TRAP for that particular filesystem has been ENABLED (see the TrapManage Group).</p>
where <fs> is a File System name		

Test Name	Test Group	Description
<p>/<fs> inodes free</p> <p>where <fs> is a File System name</p>	<p>File System</p>	<p>Maximum number of new files that can be added to the disk. For example: /usr inodes free</p> <p>For each filesystem, the default low threshold limit for inodes free is set equal to 10% of the total allocated for that system. Since near-full disks are typical, the default low level-limit for this test will be adjusted if the filesystem is already in an alarm condition at installation.</p> <p>When the number of free inodes has breached an alarm threshold, a TRAP PDU will be sent to the NMS if the TRAP for that particular filesystem has been ENABLED (see the TrapManage Group).</p>

Printer Tests

The following table shows the Printer test.

Test Name	Test Group	Description
printers	Printer	<p>Reports changes in printer status for each monitored printer. By default, only local printers are monitored.</p> <p>An alarm will be sent each time the monitored printer changes state. Some printers are smarter than others, so an alarm could consist of anything from “not printing” to “out of toner.” At start-up time for AgentMon, printers are automatically discovered and the tests are automatically configured. When the status of a printer changes state, a TRAP will be sent if the NMS has ENABLED the trap for that particular printer (see the TrapManage Group). The TRAP message will include the printer’s name and a description of the new status.</p> <p>For SunOS only: To monitor a remote printer, the printer must be defined in <code>/etc/printcap</code> and contain the boolean capability: “enlightened”.</p>

Process Tests

The following table shows the Process test.

Test Name	Test Group	Description
proc_slots	Process	<p>Number of additional processes that may be started. The default low-level alarm threshold for this is set to a value equivalent to 20% of the total number of process slots for which your kernel was configured.</p> <p>This test refers to the maximum number of new processes, applications, and programs that can be started (assuming other adequate resources exist). At program start time, AgentMon computes an alarm threshold based on your system's resources. A large number indicates you have relatively few programs running and can have many more started. A small number, especially one that became small quickly, could indicate a problem is developing.</p> <p>When this alarm occurs, you must act quickly to find the cause before the number of available slots reaches zero. When the number of available slots reaches zero, there is nothing to do but reboot. Even the simplest commands will fail to load.</p> <p>When the number of available slots breaches the alarm threshold, a TRAP PDU will be sent if the NMS has ENABLED this trap. The TRAP message includes the current number of available slots.</p>

Inventory Tests

The following table shows the Inventory tests.

Test Name	Test Group	Description
hardware	Inventory	<p>At start-up, an inventory list of the host machine is made. This is a list of display strings listing the hardware items found in the kernel at boot time. Each listed item is a device name followed by a description. The list of hardware is stored in the text file <code>\$/ENLIGHTEN/data/hardware.hostname/inventory</code>.</p> <p>For each start-up, a new list is made and compared to the previous one, if it exists. Whenever the current inventory list differs from the previous list, an alarm message is issued indicating the detected hardware addition(s) and/or subtraction(s). There are no alarm thresholds or TRAPS associated with this group.</p> <p><i>This test cannot be turned off.</i></p> <p>The following are example <code>inventory</code> files.</p> <p>For SunOS:</p> <pre> mach:Sun 4/40 ID# 289102334 sdo:sd0: Hard Disk, 3662 RPM, Intrlv 1:1 450Mbytes zso:zs0: Serial com chip (Zilog 8530) RAM:RAM: 8335360 bytes OS:sunOS 4.1.3 </pre>

Test Name	Test Group	Description
hardware (cont'd)	Inventory	For HP/UX: CPU FPU CORE-GRAPHICS-L on /dev/diag/crt100 CORE-SCSI SEAGATETEST11200N on /dev/diag/dsk/c201d6 CORE-LAN on /dev/diag/lan202 CORE-RS232-1 CORE-CENT CD-NB-AUDIO PC-FLOPPY-INTERFACE FD235HG on /dev/diag/pcflpyc20ad1 CORE-PS2-1 CORE-PS2-2 RAM: 33554432 bytes OS: HP-UX A. 09.05 IP: 129.1.2.130
software	Inventory	Similar to hardware, but only detects software installed by 'custom' and/or 'pkgadd'. This test can be run periodically. It can also be turned off.

RPC Tests

The following table shows the Remote Procedure Call (RPC) tests. These tests report various types of errors that can occur with the RPC protocol. This information is not normally required, but can be very useful when tracing network problems related to RPCs. Alarm thresholds may be set for each test, but there are no TRAPs associated with them.

Your local network specialist and O/S provider can provide more specific information about these tests. Specifics vary from O/S to O/S and most O/Ss do not support everything on this list.



The RPC statistics are *not* available on HP/UX.

Test Name	Test Group	Description
rpcc_calls	RPC	Number of client RPC calls since the last reboot.
rpcc_badcalls	RPC	Number of bad client RPC calls since the last reboot.
rpcc_retrans	RPC	Number of client RPC retransmissions since the last reboot.
rpcc_badxid	RPC	Number of unexpected packets received (client).
rpcc_timeout	RPC	Number of timeouts (client) since the last reboot.
rpcc_wait	RPC	Number of client waits since the last reboot.
rpcc_newcred	RPC	Number of times client authentication refreshed since the last reboot.
rpcc_timers	RPC	Number of client timers.
rpcc_calls	RPC	Number of server calls received since the last reboot.

Test Name	Test Group	Description
rpcs_badcalls	RPC	Number of server calls rejected since the last reboot.
rpcs_nullrecv	RPC	Number of server calls not available, though received.
rpcs_badlen	RPC	Number of server truncated packets received since the last reboot.
rpcs_xdrCALL	RPC	Number of server undecodable headers since the last reboot.

VM Tests

The following table shows the Virtual Memory (VM) tests. The virtual memory system uses a portion of disk, called SWAP, as though it were RAM memory. Virtual memory is organized into “pages,” typically 4096 bytes per page. This size varies greatly from system to system.

A machine that frequently runs out of virtual memory often requires more RAM. Another solution may be to off-load some of its work to other, less burdened systems. When the number of available virtual memory pages breaches an alarm threshold, a TRAP will be sent to the NMS if the NMS has ENABLED this trap (see the TrapManage Group). The TRAP message includes the current number of pages free.



The VM group is *not* available on HP/UX systems.

Test Name	Test Group	Description
vm_locked	VM	<p>Number of virtual memory pages currently locked. The default alarm is:</p> <p>The high limit is set to an integer value equal to 80% of the total number of pages on your system.</p> <p>Special note for SCO systems:</p> <p>This is initially set to ((total memory+total swap) -vmClaimed). This indicates the amount of free pages for user page storage.</p> <p>As this value nears zero, processes begin to fail as the malloc() and calloc() system calls refuse new memory allocation requests.</p>
vm_claimed	VM	<p>Number of virtual memory pages claimed.</p> <p>Special note for SCO systems:</p> <p>This tracks the number of memory pages not currently “locked down.” This represents all of memory minus whatever the kernel is using.</p> <p>Typically, this value starts at some value and then decreases slightly for awhile. If it continues to decrease, or decreases by a large increment, then there is probably a memory leak in the kernel or in a device driver.</p>

Test Name	Test Group	Description
vm_free	VM	<p>Number of free vm blocks (30-second moving average).</p> <p>Special note for SCO systems:</p> <p>This represents the number of memory pages not currently in use. If it is large, then little RAM is being used. If it is near zero, then the system is having to use swap space.</p>
cache_ctx	VM	Number of Ctx Cache flushes since the last reboot.
cache_seg	VM	Number of Segment Cache flushes since the last reboot.
cache_pag	VM	Number of Page Cache flushes since the last reboot.
cache_par	VM	Number of Partial Page Cache flushes since the last reboot.
cache_usr	VM	Number of User Cache flushes since the last reboot.
cache_reg	VM	Number of Region Cache flushes since the last reboot.

MBUF Tests

The following table shows the MBUF tests. These tests refer to the BSD UNIX buffer monitoring group. This group is not supported on System V based Operating Systems.

Test Name	Test Group	Description
mbufs	MBUF	Current number of mbufs obtained from the page pool.
mbuf_clusters	MBUF	Number of mbuf clusters obtained from the page pool.
mbuf_clfree	MBUF	Number of free clusters.
mbuf_drops	MBUF	Number of times failed to find space.
mbuf_space	MBUF	Number of interface pages obtained from the page pool.
mbuf_wait	MBUF	Number of times waited for space.
mbuf_drain	MBUF	Number of times drained protocols for space.
mbufs	MBUF	Current number of mbufs obtained from the page pool.

Cache Tests

The following table shows the Ncache test. This test refers to the Name Cache, which is another memory subsystem. Its size is tunable on some versions of UNIX.

The size of this cache, like any cache, will affect its hit/miss ratio and, to a lesser degree, its purge frequency. Adjusting the size of your name cache may require recompiling the kernel and is not recommended. You can also use alarm thresholds and data logging to verify that any changes in size had the desired result.

Test Name	Test Group	Description
ncache	cache	Percent of name cache misses. Specifically: misses/(hits+misses)

MIB II Tests

This section details the objects AgentMon can manage (which conform to the MIBII standard).

System Group

This group is used to store basic information about the workstation, who should be contacted, the system's location, and other administrative details.



For Solaris only the systems group is implemented. The network statistics normally available on other O/Ss are not available on Solaris.

Interfaces Group

These groups list the various network interfaces and information relevant to their current state. They contain various statistics for the different networking protocols that are in use on the server. The information in these groups can help pinpoint host-based network problems, aid in bandwidth utilization, and assist in resource planning. The groups include:

- **IP**—the Internet Protocol
- **ICMP**—the Internet Control Message Protocol
- **TCP**—the Transmission Control Protocol
- **UDP**—the Unreliable Datagram Protocol

Typically, these groups contain:

- The number of packets (and/or bytes) sent and received

- The number of packets that were bad for various reasons
- The number of protocol errors

The following table shows which MIB II tests AgentMon can manage.

Test Name	Test Group	Description
ip_total	IP	Total IP packets received.
ip_badsum	IP	Total IP packets having the wrong checksum.
ip_tooshort	IP	Number of IP packets that were “too short.”
ip_toosmall	IP	Number of IP packets that were “too small.”
ip_badhlen	IP	Number of IP headers having a bad length.
ip_badlen	IP	Number of IP packets of wrong length.
ip_fragments	IP	Number of fragmented IP packets.
ip_fragdropped	IP	Number of IP fragments discarded.
ip_fragtimeout	IP	Number of timeouts waiting for an IP fragment.
ip_forward	IP	Number of IP packets forwarded.
ip_cantforward	IP	Number of IP packets that could not be forwarded.
ip_redirectsend	IP	Number of redirected IP packets.
icmp_error	ICMP	Number of ICMP errors since the last reboot.
icmp_badcode	ICMP	Number of ICMP packets having an unknown icmp code.
icmp_tooshort	ICMP	Number of short ICMP packets received.
icmp_checksum	ICMP	Number of ICMP packets having a wrong checksum.

Test Name	Test Group	Description
icmp_badlen	ICMP	Number of ICMP packets having a bad length.
icmp_reflect	ICMP	Number of ICMP packets received.
tcp_psent	TCP	Number of TCP packets sent since the last reboot.
tcp_bsent	TCP	Number of TCP bytes sent since the last reboot.
tcp_pgot	TCP	Number of TCP packets received since the last reboot.
tcp_bgot	TCP	Number of TCP bytes received since the last reboot.
tcp_dropped	TCP	Number of TCP connections dropped since the last reboot.
udp_badhead	UDP	Number of udp packets that arrived with bad headers.
udp_badsum	UDP	Number of udp packets that arrived with a wrong checksum.
udp_badlen	UDP	Number of udp packets that arrived with a bad length.
udp_overflow	UDP	Number of udp socket overflows that have occurred.

General Tests

The following table shows the general Events tests. These tests are only available to SNMP-based network management software.

Test Name	MIB Group	Description
N/A	Limits	<p>This group is accessed via your NMS (Network Management Software) and allows the network manager to turn tests on or off, adjust alarm thresholds, and control data logging. Changes made in this group take effect immediately and become part of AgentMon's new start-up configuration. If an item does not appear in the limits group, then the network manager cannot edit it via the NMS.</p> <p>The path for perspective is:</p> <pre>testtab file > limits group > NMS GUI > Network Manager</pre>
N/A	TrapManage	<p>The SNMP protocol used by AgentMon and many NMS applications defines an alarm messaging facility called TRAPs. AgentMon has several traps and allows the network manager to enable and disable them.</p> <p>Changes made in this group take effect immediately and become part of AgentMon's new start-up configuration.</p>

Item Tests

The Group and API built-in tests have static names and functionality. EnlightenDSM also provides support for certain types of additional tests you can define. These tests are File, Directory, and Processes tests.

Each of these three tests types is predefined with a purpose and may have further subcategories of tests. You can have as many of each of the three types of tests as you want; each test name will be the name of the particular item being tested.

The following list shows the types of these “item” tests and their associated naming conventions:

- **File size**

The test will monitor the size of the specified file, for example, `/myfile size`. Use this to monitor the many files that are allowed to grow without bound. See the file `example1.sh` and the `command` capability for one possible way of automating corrective action for files that get too big.

- **File accessed**

Use this test to monitor any files that have been read.

- **File modified**

Use this test to monitor any files that have been modified.

- **File clamped**

Use this test to monitor logfiles for specific message patterns that you define.

- **Directories**

If a test name refers to a directory, the test will monitor the number of files in a directory, or queue. This can be used with an alarm set point to provide notification of a queue that is filling. Another possibility would be to use the data to show how queues fill and empty throughout the day.

If the `age =` parameter is specified, then only files more than `age` minutes old will be counted.

- **Processes instances**

A test name preceded by an exclamation point (“!”) is assumed to refer to a process. Alarm thresholds can be used to issue notification when the process is started, stopped, or if the number of instances of the process changes.

- **Processes size**

This test monitors the size of the named process in pages of the swappable process’s image in main memory.

- **Processes time**

This test monitors the total amount of CPU time used by the process.

File Tests

This section details the File subcategory tests. When the file size has breached an alarm threshold, a TRAP PDU will be sent to the NMS if the NMS has ENABLED the TRAP associated with that file. To create a test entry in the `testtab` file, enter the full pathname to the particular file you want AgentMon to monitor.

File Size

You can use the files group to monitor files. Since many UNIX files can grow without bound, AgentMon provides a method of automatically archiving, truncating, or otherwise averting a file size problem. For example, monitoring the size of `sudo` and other system files could aid in preserving the security of your system.

For example:

```
/var/adm/syslog/syslog.log
: on: testfreq=1:alarmfreq=60:\
: command=/opt/ENLIGHTEN/example.sh:\
: !log:high=1000000:units=bytes:\
: /*end /var/adm/syslog/syslog.log */:
```

If a test name is a full pathname and it does not refer to a directory, the test is assumed to refer to a file of the same name. The size of the file is then monitored and compared with alarm thresholds. If the named file does not exist, the test is ignored (until the file does exist).

File Accessed

Tests the “last-accessed” time of a file. AgentMon checks the last-accessed time and compares it with the last test’s last-accessed timestamp. If there is any change in the last access timestamp by any value in any of the threshold types, AgentMon will issue an alarm.

Unit of measure for test: **time**

For example:

```
/etc/shadow accessed | :\  
:on:testfreq=1:alarmfreq=1:log:delta=1:\br/>:+jump=1:\br/>:/*alarms if anyone reads the file */:
```

File Modified

Tests the “last-modified” time of a file. AgentMon checks the last-modified time and compares it with the last test’s last-modify timestamp. If there is any change in the timestamp as compared with any of the tests threshold types, AgentMon will issue an alarm.

Unit of measure for test: **time**

Since timestamps are in UDT format, any change in the timestamp means that the file has been modified or accessed. Therefore, the most effective way to test if a file has been accessed or modified is to set the **+jump threshold** type to a `value=1`.

For example:

```
!/usr/adm/sulog modified | :\  
:on:testfreq=1:+jump=1:\br/>:/*end /usr/adm/sulog modified */:
```

File Clamped

This test evaluates recent additions in ASCII files to string pattern regular expressions. An alarm is generated if one or more of the regular expressions match one or more of the “new” log entries. A “new” entry is any entry added to the log since this test was last run. Up to 32 regular expressions (string matches) can be configured for a single file-clamp test.

The **value** parameter passed by AgentMon to the process or script in **command** will be the number of matches for the test. The **Total Value** parameter will be set to the extended string that matched the regular expression. The actual string returned in **Total Value** may vary by platform type and OS version. Refer to the `regex(3)` man page for the platform on which you wish to configure the file clamp test.

Unit of measure for test: **matches**

For example:

```
/usr/adm/sulog clamped| :\
: on:testfreq=5:alarmfreq=5:\
: regx1 = (root)*(fail)*:high=1: \
: /*end /usr/adm/sulog clamped */:
```

Directories Tests

The `directories` group is used to monitor the number of files in a directory. The directory may be a print queue, an email queue, or any other queue where files in transition are temporarily stored. An individual queue can be monitored by watching the number of files and/or by watching the number of “old” files.

Since queues are bottlenecks, it can become a problem if the number of files becomes large because this could then also lead to a shortage of disk space. An email queue that has many old files could mean that a remote site is no longer reachable.

When the number of files in a queue breaches an alarm threshold, a TRAP PDU will be sent to the NMS if the NMS has `ENABLED` the trap associated with that queue. The TRAP message will contain the name of the queue and the number of (old) files in it.

To create a test entry in the `testtab` file, use the full pathname of a directory as the test name. Then the number of files in that directory is tracked. By setting a high alarm threshold, you can be notified when queues are getting too big (backlogged).

For example:

```
/usr/spool/ps| :\
: on:testfreq=5:alarmfreq=60:mailer=/bin/mail:\
: age=5:notify=root:!log:high=20:units=old_files:\
: /*end /usr/spool/ps */:
```

Since some queues have accounting files (and other non-queued files) in the queue directory, be sure to consider their number when setting alarm thresholds. If the named directory does not exist, the test is ignored. If the test description for this test contains the `age` parameter, then only files more than `age` minutes old will be counted.

Processes Tests

The processes group measures the number of currently running processes having the specified name. Given a list of processes and alarm thresholds, AgentMon can tell you when your daemons die, your programs terminate, when you don't have enough instances of a program running, how long a process has been executing, and the amount of memory a process has consumed.

In addition to the usual six types of alarms, two others are available: first instance started and last instance terminated. This means if you are monitoring a process called ABC, an alarm will occur when the first instance of ABC starts up, and when the last one terminates.

For each named process, a TRAP PDU will be sent if an alarm threshold for that process is breached and if the NMS has ENABLED the corresponding trap. The TRAP message will include the name of the process that went into the alarm condition and the number of processes with the same name that are currently running.



On BSD flavors of UNIX, you must specify the real process name, that is, the name returned by the `ps -c` command.

Processes Instances

To create a test entry in the `testtab` file, enter the process name (preceded by an exclamation point (!)) you want AgentMon to monitor.

For example:

```
!syslogd instances|\
: on: testfreq=5: mailer=/bin/mail: notify=root: !log:\
: units=process(es): +jump=1: -jump=1:\
: /*end!syslogd*/:
```


An alarm will be issued if any of the following are true:

- The named process starts and no other processes by that name were previously running (first instance start-up). This is enabled five minutes after AgentMon starts up to prohibit alarm at reboot time.
- The named process terminates and there are no other currently running processes having the same name.
- There is more than one instance of the named process running, and the number of instances changes by more than some user-specified alarm threshold (low limit, high limit, etc.).

Processes Size

This test monitors the size of the named process.

Processes Time

This test monitors the total amount of CPU time used.

API Tests

In addition to the many built-in tests, AgentMon also has six general-purpose test names you can use for tests you create. You can write the tests using any shell script or SQL, or you can use a compiled program. You can use any method you want.

All you have to do is have the test write results to an ordinary ASCII file. The file may contain many columns of data and many data records. AgentMon will watch the data, compare it to the alarm thresholds you have set, and notify you of any faults.

Designing an API Test

To use the API tests, follow these steps:

- 1) Write a script or program that will collect the data, for example:

```
ls -l /usr/spool/mail > mydata
```

- 2) Use the Cron management feature of EnlightenDSM to make your script run periodically. For information on EnlightenDSM's Cron Management feature, refer to [Chapter 9, "System."](#)

3) Use the Events menu in the GUI or edit the `testtab` file directly and define a NEW test. You can chose from the following names:

- `api1`
- `api2`
- `api3`
- `api4`
- `api5`
- `api6`

4) The “api” tests are set up just like any other test, except they have three additional fields. These fields are:

- *filename*

Specifies the full pathname of the file where your test writes data.

- *data*

Specifies which field or column the data is in. The value assigned is a digit prefaced by either an ‘f’ for field number or a ‘c’ for column number. In the absence of a qualifier, the default is ‘f’ for field.

The field/column delineator is any blank space. Each character in a row is considered a column.

- *label*

Specifies the field or column containing a descriptive word or label.

At each `testfreq` interval, AgentMon will check the file pointed to by `filename` to see if its modification time has changed. If the file has changed, AgentMon will reread it. For each line it reads, AgentMon will find the value stored in field (column) `data` and compare it to any alarm thresholds you have defined for this test (`api1`, `api2`, ..., `api6`).

An Example API Test

The following is an example for creating an API Group test.

Scenario: You have several databases and want to know when one or more begin to run out of space. Have cron execute a script that will record the free space and database name to a file called `dbsizes`. Suppose the file looks like:

```
1289 Kbytes parts.db
9023 Kbytes customer.db
389 Kbytes phones.db
```

You could create the new test with the following additional parameters:

```
Test name: apil(pick any unused API test name)
:file=/dbsizes:(full pathname to file holding the data)
:data=f1:(monitor the data in field one (column one if c1))
:label=f3:(the label in field #3 is the database name)
:low=400:(set a low-level alarm at 400 Kbytes)
```

as shown in the following example:

```
apil |:\
:on:file=/dbsizes:data=f1:label=f3:low=400:\
:testfreq=1:alarmfreq=60:\
:mailer=/bin/mail:notify=root:!log:\
:/*end apil*/:
```

Generating Reports

You can use the Status Map to view all logs and events and the Query Events function to search for relevant event logs. See the next section for information about using both of these features.

You can also use SQL to generate Events reports from the EMD. Refer to your *SQL User's Manual* for more details.

Status Map



The Status Map helps you navigate among individual managed hosts and/or pools on the network and it shows the status of each host and/or pool. The latter is done through changes in the color and color state (blinking or solid) of the icon for each host and/or pool.

To display the Status Map GUI, choose the Status Map option in the Events menu. If you don't want the map displayed, choose the Close Map option. The rest of this chapter details how to use the Status Map GUI and its supporting windows.

An example of the Status Map window is shown in [Figure 10-2](#).

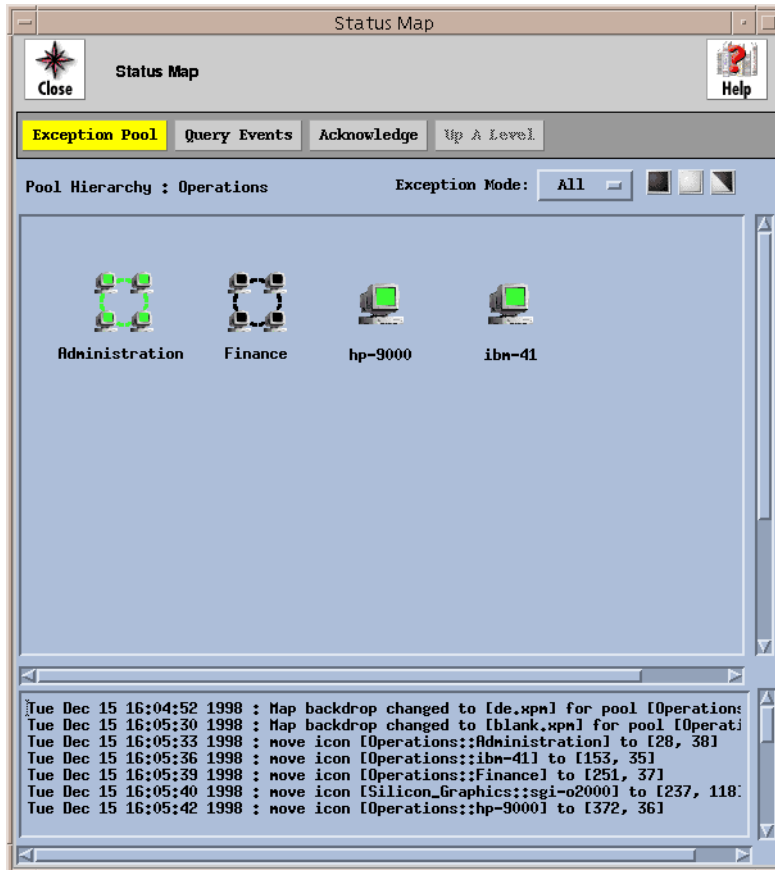


Figure 10-2 Status Map window

The Status Map uses the following color scheme to show the current highest priority event for a host/pool. This list is shown from lowest to highest priority:

- Green—OK
- White—Information
- Yellow—Warning
- Blue—Error
- Red—Severe

See Chapter 5, “Monitoring Network Systems With Events,” of the *EnlightenDSM User Guide* for a description of the types of events that occur (acknowledged, unacknowledged, cleared, and uncleared) and how the Status Map will reflect these changes.

The rest of this section describes the basic functionality of the main Status Map window. See the Chapter 5, “Monitoring Network Systems With Events,” of the *EnlightenDSM User Guide*.

Map Navigation

Select an host or pool icon by clicking on it. All *EnlightenDSM Sys Admin* functions act upon any hosts selected in the Status Map.

Map Editing

You can perform various operations on an icon by positioning the cursor over the desired icon and then using the right and/or left mouse buttons.

Right Mouse Button

Use the right mouse button while the cursor is over a host or pool icon to pop up a menu to further manipulate that icon.

The host pop-up menu offers the following choices:

- Host Overview
- Remove Host from Current Pool

Select the `Host Overview` option to bring up the Host Overview window. See “[Host Overview](#)” on [page 10-45](#) for more details on how to use this window. Select the `Remove Host from Current Pool` option to delete the host

from the current pool and thus remove the host icon from the map (if you have the correct permissions to do so).

The pool pop-up menu offers the following choices:

- Zoom into Current Pool
- Delete Pool

Select the Zoom into Current Pool option to view the contents of the selected network pool.

Select the Delete Pool option to delete the pool definition from your list of configured pools. EnlightenDSM will prompt you to confirm this action.

Left Mouse Button

Select a host or pool icon by left-clicking on it. You can also move icons by clicking and dragging (if you have Modify permissions).

You can view the contents of a pool by double-clicking on the pool icon. The Status Map will be redrawn to represent the configuration of the selected pool.

You can also bring up the Host Overview window for a host by double-clicking on its icon. See [“Host Overview” on page 10-45](#) for more details on how to use this window.

Window Buttons

This window contains the following buttons:

Exception Pool

Regular Pool

You can use this button to switch the display from the Exception Pool status icons to the Regular Pool status icons or vice versa.

An Exception Pool is a dynamic pool containing all the hosts that generated an alarm through Events or Sys Admin. The Exception Pool only shows unacknowledged events for those hosts.

The Exception Pool toggle button always shows the most severe exception (color) in its pool. You can click the Exception Pool option to display all the hosts in the Exception Pool on the Status Map. The Exception Pool then becomes your current active pool.

Query Events

You can click this button to bring up the Query Event Messages window and search for Events messages of a specific type. See [“Query Events” on page 10-49](#) for more details on how to use this window.

Acknowledge

You can click this button to acknowledge the unacknowledged events for the selected host or for all hosts in the selected pool. The flashing color of the icon will become a solid color, indicating there are no longer any unacknowledged events associated with the host(s). You can execute this option only if you have Modify permissions.



If you are in the Exception Pool and acknowledge events for a host, the icon for that host will be removed from the Status Map display.

Up A Level

You can click this button to traverse out of a sub-pool. This action will reset the network map, placing you outside of the sub-pool you are currently viewing. This is useful since EnlightenDSM has the ability to place pools inside of pools, creating many layers of sub-pools.

Exception Mode

You can use this toggle to switch the display mode for the Exception Pool. The options are to show All exceptions in the Exception Pool (the default) or to only show the ones Hidden from your current point of view (those exceptions found in the pools in the level above, below, or alongside the pool you're currently viewing).

Information Areas

This window contains the following information areas: Pool Hierarchy, Status Map, and Event Notification. This window contains the following information areas:

Pool Hierarchy

This line shows the complete pool or sub-pool pathname for the icons being displayed in the Status Map. This indicates what depth level the pool or sub-pool is within its hierarchy.

Status Map

The Status Map is the large scrollable region below the action buttons. This map is used by both Sys Admin and Events functions. For Sys Admin, you can use the map to easily view which hosts are in what pools and then target some administration functions to only a subset of hosts in a pool. For Events, you can use the map to see when an event has been triggered on a host.

The Status Map consists of a background drawing (which can be blank) and (any) icons placed on the drawing. These icons represent individual hosts and/or sub-pools in the selected pool. You can modify the background drawing by placing the mouse in the Status Map area, clicking the right-mouse button, and selecting the Change Background Map option. Then choose the new background to Apply. You can only execute this option if you have Modify permissions.

Event Notification

The list box at the bottom of the window will display event messages as they come in. Any Status Map activity, such as when you change what the top-level pool view is, will also be reflected in this log. Events messages from all hosts are logged here, regardless of what pool you're currently viewing.

You can clean out this log by placing the mouse in the Event Notification area, clicking the right-mouse button, and selecting the Clear Messages option.

Host Overview

You can use this function to get an overview of the activity for your selected host. Refer to [“Map Editing” on page 10-41](#) to select a host for this function. The Host Overview window ([Figure 10-3](#)) will appear.

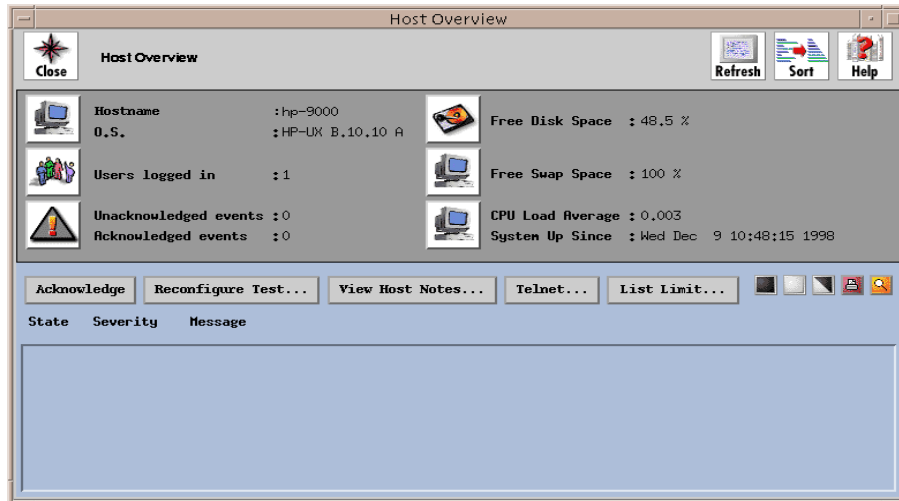


Figure 10-3 Host Overview window

Combo Areas

This window will display the following information and icons:

Hostname

O.S.

This area displays the selected hostname and its O/S type. You may click on the System icon next to this block of information to bring up the Processes window for the current host. Refer to [“Process Status” on page 3-34](#) to modify the host configuration.

Free Disk Space

This area displays the amount of free disk space currently available on the current host. You may click on the Disk icon next to this area to bring up the Disk Usage By Filesystem window for the current host. See [“Usage by Filesystem” on page 7-3](#) for details on how to use this window.

Users logged in

This area displays the number of users currently logged on to the current host. You may click on the User icon next to this area to bring up the Who Is Logged In? window for the current host. See [“Who is Logged In” on page 3-32](#) for details on how to use this window.

Free Swap Space

This area displays the amount of free swap space currently available on the current host. You may click on the System icon next to this area to bring up the Swap Space Usage window for the current host. See [“Swap Space” on page 9-4](#) for details on how to use this window.

Unacknowledged events

Acknowledged events

This area displays the number of unacknowledged and acknowledged events currently associated with the selected host. You may click on the Events icon next to this block of information to bring up the Query Event Messages window for the current host. Refer to [“Query Events” on page 10-49](#) for details on how to use this window.

Cpu Load Average

System Up Since

This area displays the current CPU load average and the most recent time the system was brought up for the selected host. You may click on the User icon next to this block of information to bring up the Summary of Processes by User window for the current host. See [“CPU Summary” on page 3-35](#) for details on how to use this window.

Buttons

This window contains the following buttons:

Acknowledge

Select one or more events from the Events List box and click this button to acknowledge it/them. This action will change the state of the selected event(s) in the Events List. You can execute this option only if you have Modify permissions.

This action could also update the host(s) icon status state in the Status Map, depending on the status for the rest of the events associated by that host.

Reconfigure Test ...

Select one or more unacknowledged events from the Events List box and click this button to bring up the Modify Events Test window for the selected event(s).

View Host Notes ...

Click this button to bring up the Host Overview - Host Notes window for the current host. This window will contain any notes associated with the current host, as shown in [Figure 10-4](#).

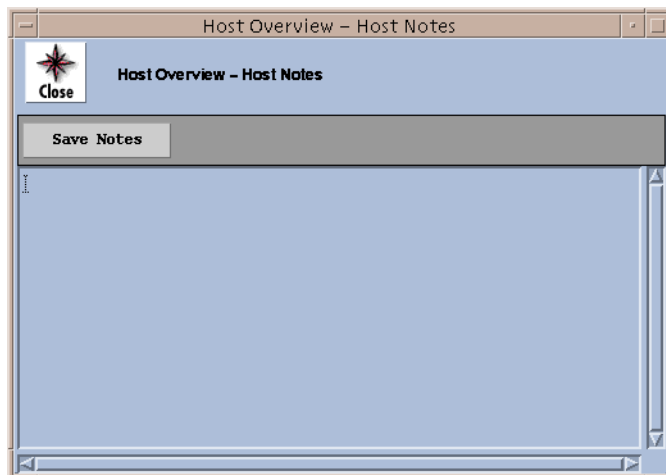


Figure 10-4 Host Overview-Host Notes window

You can also add more information to this window. To do so, move the cursor to the relevant point, type in your notes, and click the Save Notes button to update the current Host Notes.

Telnet ...

Click this button to pop up an xterm window running a remote session on the current host. From here you can log in to the system and then use any standard commands.

Events List

This list box shows all the current events, acknowledged and unacknowledged, received from the current host.

This list is sorted alphabetically by test name and then chronologically by event time. The most recent event will be inserted at the top of the sub-list for the relevant test as the event is received.

Each line in the list will display the event status (Ack = acknowledged, Unack = unacknowledged), severity, and message for that test/event.

You can change how many messages show in the list box by placing the mouse in the list, clicking the right-mouse button, and selecting from the following options:

- View All Events (the default)
- View Acknowledged Events (shows only acknowledged events, if any)
- View Unacknowledged Events (shows only unacknowledged events, if any)
- View Unclear Events (shows only unclear events, if any)

Refer to Chapter 5, “Monitoring Network Systems With Events,” in the *EnlightenDSM User Guide* for a definition of these “event types.”

Query Events

You can use this function to search for and view both alarm and normal logged messages from Events. The Query Events Messages window will appear (Figure 10-5).

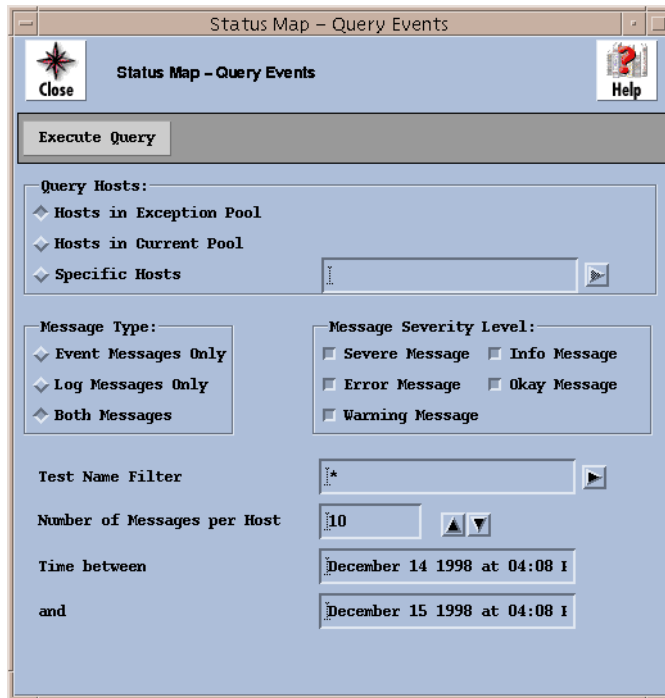


Figure 10-5 Query Events Messages window

Fields

This window contains the following fields:

Query Hosts

Use these toggles to specify which hosts to query for reports. The options are:

- Hosts in Exception Pool (the default)
- Hosts in Current Pool
- Specific Hosts

If you choose the Specific Hosts option, use the text field to the right of that option to specify which host(s) to check for messages. Leave a blank space

between host names for multiple entries. You can also click the arrow button to the right to pop up a pick list of all hosts within the current pool and make your selection(s) from there.

Message Type

Use these toggles to specify which type of messages to search for in the reports. The options are:

- Event Messages Only
- Log Messages Only
- Both Messages (the default)

Event Messages Only are Events alarm messages generated by a test crossing a user-defined threshold. Log Messages Only are those Events informational messages generated when a test runs without generating an event and that (successful) result is logged.

Message Severity Level

You can use these toggles to choose what severity level of the messages to use in the search. The options are:

- Severe Message
- Info Message
- Error Message
- Okay Message
- Warning Message

You can select more than one severity level to use in the search; you must select at least one severity level. The default is all levels of message severity will be queried during the search.

Test Name Filter

You can use this field to limit the search to specific test names. You may type in the entire test name or just the first few letters of the test name. All tests whose name contains part or all of the specified string will be queried. Leave a blank space between test names for multiple entries. You can also click the arrow button to the right to pop up a pick list of all pre-defined standard Events tests and make your selection(s) from there.

You can also use the standard wild cards '*', '[]', and '?' in this field (e.g., /home/*).

Number of Messages per Host

Events logs can become large very quickly and therefore take more time to search. You can use this field to help speed up the search time and view only the most recent X messages per host. You can also use the counter buttons to the right to increment or decrement the number displayed. The most recent message is always displayed first.

Timestamp between ... and

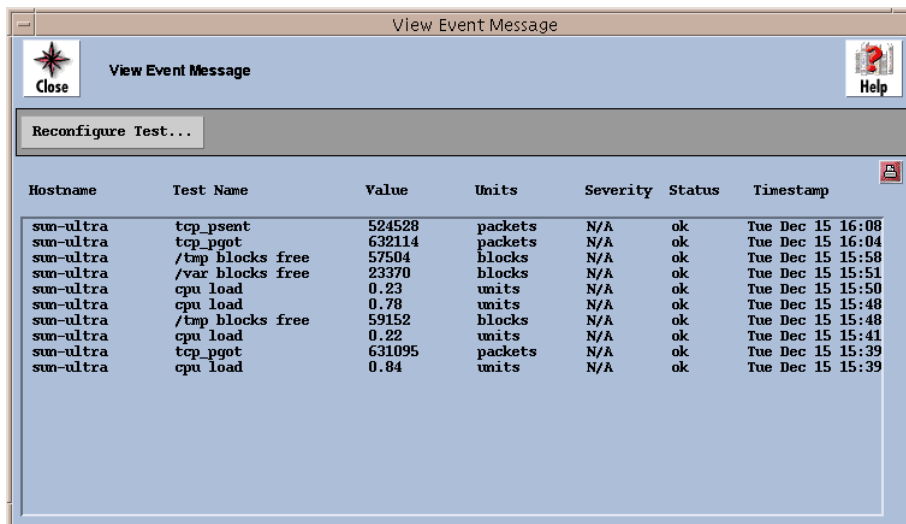
You can use these two fields to limit the search to messages logged between the specified times. See [Appendix C, “Time Formats,”](#) for details on specifying this time-date value.

Buttons

This window contains the following button:

Execute Query

After making your selections, click this button to begin the search process. When the query is completed, the results will be displayed ([Figure 10-6](#)).



The screenshot shows a window titled "View Event Message" with a "Close" button and a "Help" icon. Below the title bar is a "Reconfigure Test..." button. The main area contains a table with the following data:

Hostname	Test Name	Value	Units	Severity	Status	Timestamp
sun-ultra	tcp_psent	524528	packets	N/A	ok	Tue Dec 15 16:08
sun-ultra	tcp_pgot	632114	packets	N/A	ok	Tue Dec 15 16:04
sun-ultra	/tmp blocks free	57504	blocks	N/A	ok	Tue Dec 15 15:58
sun-ultra	/var blocks free	23370	blocks	N/A	ok	Tue Dec 15 15:51
sun-ultra	cpu load	0.23	units	N/A	ok	Tue Dec 15 15:50
sun-ultra	cpu load	0.78	units	N/A	ok	Tue Dec 15 15:48
sun-ultra	/tmp blocks free	59152	blocks	N/A	ok	Tue Dec 15 15:48
sun-ultra	cpu load	0.22	units	N/A	ok	Tue Dec 15 15:41
sun-ultra	tcp_pgot	631095	packets	N/A	ok	Tue Dec 15 15:39
sun-ultra	cpu load	0.84	units	N/A	ok	Tue Dec 15 15:39

Figure 10-6 New Event Messages

The list box shows all the messages matching your search criteria. Each line in the list will display the hostname, test name, logged value, units, severity, status, and time stamp.

From here you can select one of the tests and click the Reconfigure Test button to bring up the Modify Event Test window for that test.



Wizard

The Wizard function exists in the Disk Usage Information by Filesystem, the User Configuration, the Configured Printers, and the Host Configuration windows. These Wizards help administrators maintain system security and usability by checking the status of various functions and settings.

Click the Wizard button to view the appropriate analysis window. Select the type of security evaluations you want, then click the Apply button.

A summary of findings will be displayed in the top right portion of the window (as shown in [Figure A-1](#)). Use the window buttons for further analysis.



The evaluations marked by a star (*) are relatively time-consuming.

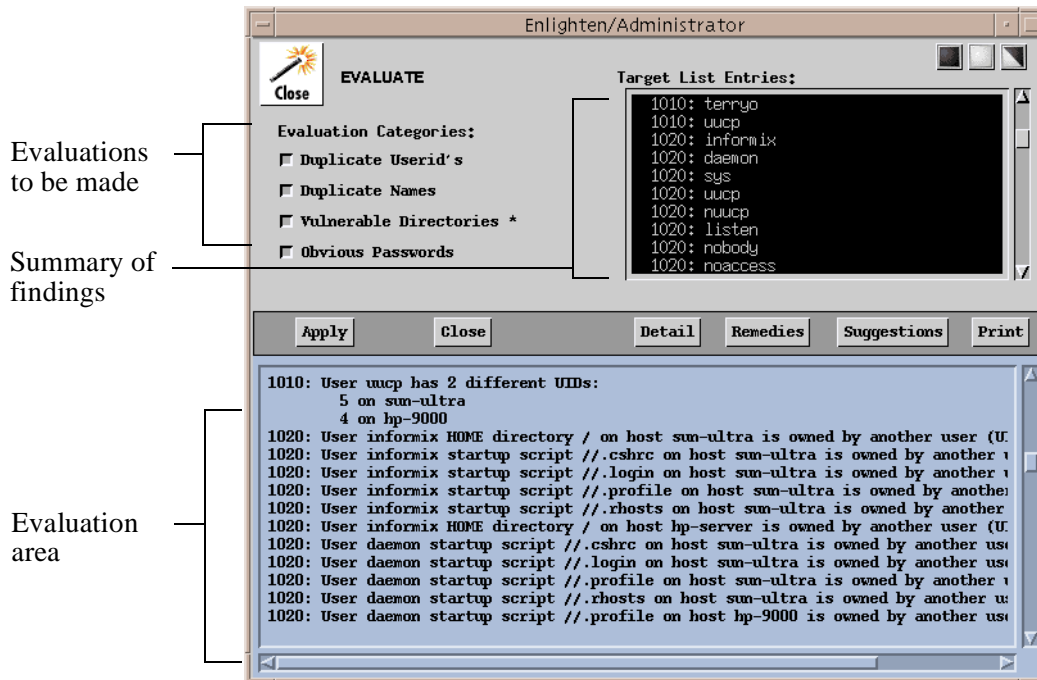


Figure A-1 Evaluate window

Detail

Use this button to display (in the evaluation area) a detailed list of findings for those entries highlighted in the summary window.

Remedies

Just knowing that there is a potential problem does not necessarily solve it. You can use this button to get a brief description of the problems encountered for the entries highlighted in the summary window. Furthermore, the Wizard will also make suggestions on how to remedy the problems using EnlightenDSM.

Suggestions

Use this button to see suggestions on how to avoid the problem in the future for any entries highlighted in the summary window.

Print

You can use this button to print the contents of the explanation area, as defined in the default print command field in the Session Preferences window. See [“Session Preferences” on page 2-2](#) for more details on setting up this command.

The remaining sections of this chapter detail the security check options within each of the four Wizards.

Disk Wizard



The Disk Wizard performs a series of disk checks with an emphasis on security needs.

Full Disk Partitions

This is a simple test to check if disk usage on any of the partitions has surpassed 85% capacity.

Set UID/GID Problems

When a set UID or GID (setuid) program is executed, it gains the permissions of the owner/group of the program. Consequently, a setuid program owned by root will give the user super user privileges when executed. This concept is fine as long as these programs have the proper built-in safety mechanisms. Without the proper safety mechanisms, these programs can create a major security risk.

Wizard runs several checks on setuid programs when this option is selected. It performs the checks against the current disk snapshots of all partitions. Furthermore, some of the checks involve the use of the associated Master snapshot. Make sure the disk snapshots for all partitions are relatively new. See [“Save Current Snapshots” on page 7-26](#) for more information on re-creating snapshots and creating Master Snapshots.

First, the Wizard checks if each setuid program found is in what it considers to be a system directory. If not, this will be reported. The current system directories are:

```
/bin/sbin/etc  
/usr/bin/usr/sbin/usr/lib  
/usr/ucb/usr/kvm/etc
```

Next, the Wizard checks the permissions of the program. If the program is world-writable, that is, may be replaceable, this will be reported.

Finally, if a master snapshot for the partition being checked exists, the Wizard checks if the setuid program existed at the time the master snapshot was created. If the setuid program has been newly created since the master snapshot was taken, this will be reported. If the program did previously exist,

Wizard checks if it is still the same file by looking for a difference in the file's size, permissions, or date of last modification.

Devices Not in /dev

Device files are used in UNIX as an interface to the hardware and kernel memory. These device files should only be found in the directory `/dev` (and on some systems `/devices`). If a device file is found outside of these directories, it should be considered a serious threat.

Wizard scans the existing disk snapshots looking for devices not in the `/dev` or `/devices` (where appropriate) directories when this option is selected. If any files are found, they are reported.

User Wizard



The User Wizard helps with user-related problems.

Duplicate UserID's

This operation searches for users with different names sharing the same UID. There may be times when this usage is necessary.

Duplicate Names

This operation searches for users with the same name using different UIDs on different hosts. Typically a user should have a consistent UID on all hosts.

Vulnerable Directories

This operation is similar to the User Home Directory security check in the security menu. See [“User Home Directories” on page 6-9](#) for more information on this check.

Obvious Passwords

This operation is a subset of the obvious password checking in the security menu. It checks if the user has a password, if the password is the same as the username, or if the password is zero length (just press Enter to get in). For more detailed password checking, see [“Obvious Passwords” on page 6-11](#).

Printer Wizard



The Printer Wizard checks for printer-related problems.

Downed Printers

This checks if any of the printers are disabled. If any are, they are reported.

Long Print Queues

This checks if a print queue for any of the printers has more than 15 print jobs queued. If so, this is reported. If a print queue gets too long you might want to move some of the jobs to other compatible printers. See [“Move” on page 4-10](#) for more details.

Large Print Jobs

This checks if any queued print job is larger than 400 KB. If so, this is reported.

Host Wizard



The Host Wizard checks for potential host conflicts.

Host Address Conflicts


This check looks at all hostnames and their associated TCP/IP addresses and verifies all hosts of a given name have the same network address. If a hostname is found on one host with a different address than on another host, this is reported.

Host Entry Name Conflicts

This check looks at all hostnames and their associated TCP/IP addresses and verifies a given address only has one hostname associated with it. If an address is found to have more than one hostname associated with it on any host, this is reported.

Host Aliasing Conflicts

This check looks at all host alias names and verifies each alias name is used uniquely for a single hostname. If an alias is found to apply to more than one hostname, this is reported.



Altering the Menu

Enlighten**DSM** menus and icons can be altered to suit your needs. The options are:

- The order of menu icons
- The order of menu items
- The actions of menu items
- The number icons and menu items



You may *not* remove the Configure menu from the main menu toolbar. Doing so would cause improper placement of the remaining menu icons in the main menu.

The Enlighten**DSM** GUI is built through a small text database file containing objects and attribute lists. Alternative interfaces can be created by editing this file. A menu item could be added, for example, that would run a customized script. Another possibility would be to replace the Archive configuration menu item with your own backup tool.



During the Enlighten**DSM** start-up, the interface file will be scanned for syntax and for attribute matching. If an error is detected, Enlighten**DSM** terminates the start-up process and reports the line number and the reason for rejection on `stderr`.

Overriding the Default Configuration File

The Enlighten**DSM** GUI is defined in a configuration file. This file contains information describing the layout of the windows. You can edit the configuration file to make desired additions or modifications to the GUI. Then you can override the default configuration file

`$ENLIGHTEN/config/winn.out` by setting the shell variable `WINOUT_NET` to the name of **YOUR** configuration file.

If Enlighten**DSM** finds the shell variable `WINOUT_NET` set and it contains the name of a readable file, this file will be used instead of the default configuration file. If the shell variable is set, but points to a file that does not exist or is unreadable, a warning will be issued. Enlighten**DSM** will then use the default configuration file.

Enlighten**DSM** is shipped with a backup copy of the configuration file called `$ENLIGHTEN/config/winn.org`. You can copy this over its corresponding `.out` file if you need to return the interface to its original state.

Warning!

The format of the interface file is crucial to the successful start-up of Enlighten**DSM** and the consistency of the product. *Be careful when making any modifications to this file.* You should also test your modified interface before replacing the current version.

Where to Edit

While there is no order of declaration of menus within the interface file, the current group of Menu definitions occurs after the set of MODEL initializations. If you wish to alter the existing menu definitions, or add your own, you should do so at this location (between the MODELS and the Window definitions). You can quickly locate the first menu definition by searching for the keyword Define.

The only required order in the file is that menus must be defined prior to any reference to their use. Thus, a Menu must be defined before it is attached to an icon in a window, and a submenu must be defined before it is attached to a menu item.



A submenu is not a unique object; it is just the term for a Menu that's attached to a menu item as a pull-right style menu.

Defining Menus

The interface file uses a simple, yet concrete, syntax. Any line with a # as its first character is a comment. Any line with no characters (zero-length) is a break line. Otherwise, every line has three fields separated by tabs.

To create a menu, use the syntax:

```
DEFINE<tab>12<tab><menu-name>
```

where *<menu-name>* is a single-word name that is unique among the defined menus and used internally for reference purposes only. For example, to add a “UserMenu1” menu, add the definition:

```
DEFINE      12      USERMENU1
```

Subsequent lines following the DEFINE statement set the menu attributes and menu items (and their attributes) for the current DEFINE menu. This occurs until the next DEFINE statement is reached.

Menu Attributes

The syntax for an attribute is:

```
<tab><attribute-type><tab><value>
```

The valid attributes for menus are (order is not important):

- **categ** *<number>* General EnlightenDSM sphere of monitoring
- **context** *<string>* Mini-help text that appears when the mouse moves over the menu icon
- **pinned** *<value>* Whether window is pinnable
- **label** *<string>* Menu title appearing at the top of the menu window

categ <number>

Each category is in charge of windows and actions associated with it. The default interface shows the full range of windows accessible from a given category; it is not possible to change a window's accessibility from one menu category to another. Thus, it is not possible to call the User Query window from the Printer category. This restriction occurs since the window generation is determined by the Menu Category and the Return Code (as generated by menu items).

The valid (reserved) menu categories for Enlighten**DSM** are:

- 2 = Configuration
- 3 = System
- 4 = Disk
- 5 = User
- 7 = Printer
- 8 = Archive
- 9 = Events
- 10 = Security
- 11 = User-Defined Section 1
- 12 = User-Defined Section 2
- 15 = Network

The valid (reserved) submenu categories for Enlighten**DSM** are:

51 = Spool

53 = Activities

55 = Statistics

57 = Logins

59 = cron

60 = NFS

61 = DNS

62 = Aliases

63 = Clocks

context *<string>*

You may give a string, enclosed in quotation marks, that Enlighten**DSM** uses to display at the bottom of the Main Menu Window as the mouse is dragged into the menu's associated icon. This is a form of context-sensitive help. If the text is longer than the window, it is "lopped off" at the window edge.

pinned *<value>*

This attribute determines whether a menu is pinnable; it will remain visible by moving the mouse over the window push-pin, so its menu items are always accessible without going back to the menu icon. The value is either 0 or 1 (False or True). By default, a menu is not pinned, so you only need to specify this attribute if the menu will be pinnable.

label *<string>*

Whenever a menu is displayed, the title of the menu appears at the top. A title is required for all menus and it must have quotation marks around it.

Sample UserMenu1 Menu

For our UserMenu1 menu example, we might set the following attributes:

```
DEFINE      12          USERMENU1
            categ      11
            pinned     0
            label      "User Scripts 1"
            context    "\"Help subsystem"
```

This example defines a Menu we can refer to later in the interface file as `USERMENU1`, within the User-Defined section, not pinned (so it will disappear after a selection is made), with the title “User Scripts 1”, and some help text showing that this menu comprises menu items that call up internally created executables.

Defining Menu Items

Menus also contain Menu Items, which comprise a list of entries from which the user may select. A menu item may be a pull-right item, meaning that a small arrow appears to the right of the item label. If the user “follows” the arrow, a submenu appears.

A menu item is simply an attribute assigned to a menu. It is created using the attribute type ‘add’. Thereafter, any further attributes listed apply not to the menu itself, but to the menu item most recently added. The syntax for adding a menu item is:

```
<tab>add<tab><entry text>
```

where *<entry text>* is the actual text (enclosed in quotation marks) the user will see in the list of menu items. Items will appear in the menu in the order specified within the file, with the first item appearing at the top of the menu.

Menu Item Attributes

Attributes for menu items use the same syntax as those for Menus (see “[Menu Attributes](#)” on page B-5). The valid attributes for a menu item are (order is not important):

- **exec** *<string>* Executes the given program/script
- **state** *<value>* Shows the display style of the menu item
- **submenu** *<value>* Attaches a pull-right menu to the current item
- **return** *<value>* Returns the “key value” associated with the menu entry

exec *<string>*

A “system” call will be made for *<string>*, which is enclosed in quotation marks. Any program that could be called from a command-line interface (complete with arguments) is legal here. The program name must be ≤ 220 characters. If you do not wish to have the program block further execution of EnlightenDSM, run the program in the background by using the ampersand (&), just as you would from the command line.

state *<value>*

The default state of an added menu item is active; the menu item is displayed clearly and is selectable by the user. Or you can have the menu item displayed but grayed-out; it is viewable but not selectable by the user. To make a menu item inactive, use the value eight (8). This option also will cancel the display of any subsequent submenus for this item.

submenu *<name>*

You attach a previously defined submenu to the current menu item by giving its internal name (from the DEFINE statement) embedded in quotation marks. This attribute overrides any exec associated with the menu item; it can only display a pull-right menu.

return *<value>*

The Return value links a menu item to a particular set of EnlightenDSM windows. The menu Category, along with the Return value, determines which window is then initiated. In a sense, these values are “reserved” codes. The default interface file has all the combinations of the Category values and Return codes. For example:

#2 - Configuration Section

1 = Session Preferences

2 = Alter Menu

3 = Pool Configuration

4 = User Authorization

5 = New User Templates

6 = EMD Data Expiration

In addition to the reserved Return codes for Menus, there are also reserved Return codes for submenu categories. For example:

#53 - Activity Submenu

1 = Who is Logged In

2 = Process Status

3 = CPU Summary

Although a Submenu, by default, belongs in an associated Menu category, since it carries its own Category and Return codes, it can be placed anywhere. For example, you could place both the User Security and System Integrity submenus on their own User-Defined menu and call that menu “System Security.”

To change the names of any existing menu items, change either the ‘label’ or ‘add’ attribute arguments. The order of the menu items is not fixed, but be sure to keep all of a menu item’s attributes with its corresponding add attribute.

A Utilities Submenu

To expand on the previous example of a user-defined UserMenu1 menu, you could add the following items:

```
#
# The following user utility menu is a predefined example.
# It can be used to add you own scripts into the ENL main menu.
#
# The definition exists here, but is currently deactivated in the
# DEFINE 7 MAINMENU definition further down. To activate it, the
# comment markers (#) must be removed from the following lines
#
# #      add MENUBUTTON1
# #      name      "USERSCRIPT"
# #      icon  12
# #      menu     USERMENU2
#
DEFINE 12  USERMENU2
    categ  12
    context "This is a mini-help line as the mouse moves over the
icon"
    label   "User Scripts 2"
    pinned  1
    add "Script 1"
    return  1
    exec   "xterm -e $ENLIGHTEN/bin/script1"
    add "Script 2"
    return  2
    exec   "xterm -e $ENLIGHTEN/bin/script2"
#
# The following user utility menu is a predefined example.
# It can be used to add you own scripts into the ENL main menu.
#
# The definition exists here, but is currently deactivated in the
# DEFINE 7 MAINMENU definition further down. To activate it, the
# comment markers (#) must be removed from the following lines
#
```

```
DEFINE 12 USERMENU1
  categ 11
  context "Help subsystem, by category and action"
  label "User Scripts 1"
  pinned 1
  add "Script 1"
  return 1
  exec "xterm -e $ENLIGHTEN/bin/script1"
  add "Script 2"
  return 2
  exec "xterm -e $ENLIGHTEN/bin/script2"
  .....
  .....
  .....
  ..
  .

# add MENUBUTTON1
# name "USER MENU 1"
# icon 11
# menu USERMENU1
# add MENUBUTTON1
# name "USER MENU 2"
# icon 12
# menu USERMENU2
```

The Main Menu Window

All menus are attached to icons in the Main Window. This definition is shown in the interface file as the first **DEFINE 7** statement. Use the attribute `add MENUBUTTON1` to attach a menu to its window (left to right, in order of declaration). An internal `name` attribute is not necessary, but is usually included.

Use the menu `menu-name` attribute to attach the previously defined menu to the menu button (the `menu-name` is shown in the line `DEFINE 11 name`). Use the icon `icon-type` attribute to assign the icon to be displayed as the menu button. The icons are pre-defined:

2 = The configure icon

3 = A picture of a terminal

4 = A picture of a disk

5 = A picture of users

6 = A picture of a terminal

7 = A picture of a laser printer

8 = A picture of a 1/4-inch tape


9 = The Events icon

10 = A picture of a lock

11 = A picture of tools (User-Defined Menu)

12 = A picture of tools (User-Defined Menu)

You can also have multiple User-Defined menu buttons, as long as the return codes for the menu items are unique across all of them. For example, you can define two utility menus, “UTILITY1” and “UTILITY2,” with three menu items each, with UTILITY1 having menu items with return codes 1–3 and UTILITY2 having menu items with return codes 4–6.



Time Formats

Several fields throughout Enlighten**DSM** require the entry of a date and time. This date and time can be specified in two ways: *absolutely* or *relatively*.

The *absolute* method is used to enter a specific period of time. To search for a file created since May 10, 1996, at 12 noon, you could use the entry: `MAY 10 96 at 12:00`. To search for a file created between two specific times and dates, you would enter both those times and dates in the appropriate fields. The *absolute* method is the only way to search for a combined date/time format.

The *relative* method specifies a span of time, such as days or months, to be used in a search against the current system clock. To find users who have been logged in during the last three days, you could use the entry: `3 days`. The *relative* method is strictly a time OR date format: for example, `2 seconds OR 1 month`.

With either method, the format is case-insensitive. For example, you could use `JAN`, `Jan`, or `jan` as the abbreviation for the month of January.

This appendix shows all the BNF time measurements in initial caps (e.g., `Jan` or `January`, `2 Sec` or `2 Seconds`, etc.).

Examples:

Absolute:

15 1996 (defaults to the current month)
Jan 1 97
July 4 (defaults to the current year)
June 15 at 3 PM (defaults to the current year)
June 16 at 2 AM (defaults to the current year)
June 17 96 at 15:31
18 at 9:31:14 AM (defaults to current month AND year)

Relative:

1 second
5 Min
1 day
2 weeks
1 Year

The rest of this appendix defines the time formats used to fill in a date and/or time field.

Absolute Date/Time BNF

<Absolute Date> = (<Month> <Day Of Month> (<Year>)
("at" <Time>)

<Month> = <Long Month> | <Abr Month>

<Long Month> = "January" | "February" | "March" | "April" | "May" |
"June" | "July" | "August" | "September" | "October" |
"November" |
"December"

<Abr Month> = "Jan" | "Feb" | "Mar" | "Apr" | "May" | "Jun" | "Jul"
| "Aug" | "Sep" | "Oct" | "Nov" | "Dec"

<Day Of Month> = 1 ... 31

<Year> = 19xx | xx

<Time> = <Hour> (":"< Minutes>) (":"< Seconds>)

<Hour> = <12 Hour> | <24 Hour>

<12 Hour> = <12 Hours> "AM" | "PM"

<24 Hour> = <24 Hours>

<12 Hours> = 00 ... 00

<24 Hours> = 00 ... 23

<Minutes> = 00 ... 59

<Seconds> = 00 ... 59

Relative Date/Time BNF

<Relative Date> = <N> <Units>

<N> = 1 ... *infinity*

<Units> = <Seconds> | <Minutes> | <Hours> | <Days> | <Months> |
<Years>

<Seconds> = "Seconds" | "Sec"

<Minutes> = "Minutes" | "Min"


<Hours> = "Hours" | "Hour"

<Days> = "Days" | "Day"

<Weeks> = "Weeks" | "Week"

<Months> = "Months" | "Month"

<Years> = "Years" | "Year"



Password Aging

With password aging, a user's password is valid only for a specific period of time. When the user attempts to log in after that period of time, they are forced to change their password. Enlighten**DSM** supports password aging for user accounts.

Two time periods are generally associated with password aging:

- Length of time a password is valid
- Minimum period of time during which the user is not allowed to change their password

Both of these times are relative to the last time the password was changed. Consequently, every time the user changes their password, the expiration timer is reset.

Different operating systems may use different formats to implement password aging. The methods of storing information may differ along with the granularity of keeping time. This appendix describes two common types of password aging implementation.

Berkeley UNIX and Pre-system V.4

This password implementation uses n weeks as the unit of time. If a user changes his or her password on two consecutive days, unless these days happened to fall on two different “weeks,” the date of the last password change would be the same as the first password change. A password can be valid for a maximum of 64 weeks.

If your system supports this type of password aging, you cannot expire a user’s password without first turning password aging on for that user.

System V.4

With the advent of UNIX Release 5.4, a secondary method of storing passwords was created — the shadow password file. This file also includes fields for password aging. The granularity of time for a shadow password is one day. This file is updated as a function of EnlightenDSM's user account management.



SNMP Managed Groups

The following is a list of Events management objects (SNMP Groups) that can be monitored or altered. This group of objects is defined in the AgentMon MIB.

SNMP Groups

This section lists the SNMP Groups and their test results

Group	Action
limits	OBJECT IDENTIFIER::={ Events20 1 } This SNMP group is used to manage Events from a network management station. It allows the Network Manager to reconfigure Events in many of the same ways that the System Administrator can with the Configure Events GUI.
trapmanage	OBJECT IDENTIFIER::={ Events20 2 } This group is also used by the Network Manager. This group is used for enabling and disabling TRAPs and for changing the severity level associated with each TRAP.
os	OBJECT IDENTIFIER::={ Events20 3 } Displays read-only values for the operating system name, local time, and time of last reboot.
kernel	OBJECT IDENTIFIER::={ Events20 4 } Displays values for various kernel-related statistics. These values are useful for an OS kernel guru who likes to tweak the O/S for performance.
mbuf	OBJECT IDENTIFIER::={ Events20 5 } Displays statistics for various network-related message buffers.
ncache	OBJECT IDENTIFIER::={ Events20 6 } Displays statistics on the Name Cache (if present).
vm	OBJECT IDENTIFIER::={ Events20 7 } Displays statistics about the virtual memory system.
cpu	OBJECT IDENTIFIER::={ Events20 8 } Displays the CPU load (average growth of a run queue) and CPU time for the following states: kernel, user, system, and idle (as a percentage; total % of four states = 100%).

Group	Action
proc	OBJECT IDENTIFIER::={ Events20 9 } Displays the number loaded, CPU time used, and the process size for a user-selected process. The process status is also available on some operating systems.
fs	OBJECT IDENTIFIER::={ Events20 10 } Displays a list of local filesystems, their size, and the amount of free space available. This may also include swap space.
files	OBJECT IDENTIFIER::={ Events20 11 } For each user-specified file, this reports one or more of the user-specified items: filename, size, last accessed time, last time file was modified, or last string matched if 'file clamping' is selected and the user has specified one or more regular expressions.
qu	OBJECT IDENTIFIER::={ Events20 12 } For each user-specified directory, this reports the total number of files and number of old files in the directory. This also has a variable field (for each directory) so the user can specify what 'old' is.
hardware	OBJECT IDENTIFIER::={ Events20 13 } Displays a list of all hardware that was discovered on the system when Events was last started.
software	OBJECT IDENTIFIER::={ Events20 14 } Displays a list of all software currently installed on the system, if such software was (de)installed with pkgadd. On SCO UNIX, the list also includes all software installed with custom.
printer	OBJECT IDENTIFIER::={ Events20 15 } Displays the printer daemon's view of the current status for each printer being monitored.
disk	OBJECT IDENTIFIER::={ Events20 16 } This group is reserved for future use.
rpc	OBJECT IDENTIFIER::={ Events20 17 } Displays RPC client and server statistics.

Group	Action
api1, ..., api6	<p data-bbox="500 323 1176 524">OBJECT IDENTIFIER ::= { Events18-23 } This group allows Events to monitor user-provided data. If you have test scripts that already collect data, you could modify the scripts to write any data to a file. Then you could use one or more of the API groups to monitor the data in the file and check for High and Low thresholds.</p> <p data-bbox="500 566 1176 693">Like all other tests, a threshold violation could then send email or a TRAP, notify PEP, <i>and</i> be logged to the EMD. Logging of non-alarm data is also provided, just as it is with other tests.</p> <p data-bbox="500 735 1176 833">See Chapter 5, “Monitoring Network Systems With Events,” in the <i>EnlightenDSM User Guide</i> for more details on using API tests.</p>
usertraps	<p data-bbox="500 862 1176 1135">OBJECT IDENTIFIER ::= { Events20 24 } This group shows the last TRAP PDU that was initiated through the EventsCli command line interface. The EventsCli is an “open” interface between your existing shell scripts (and other monitoring programs) and the Events alarm notification methods (PEP, TRAPs, EMD). This is best suited to users who are already monitoring site-specific data and have their own threshold logic.</p> <p data-bbox="500 1177 1176 1268">When your tests detect an alarm condition, they could call the EventsCli. This then gives you easy access to our event management framework.</p> <p data-bbox="500 1310 1176 1343">See “EventsCli (8N)” on page F-5 for more details.</p>

MIB II Groups

This section lists the MIB II Groups and their object identifiers.

Group	Action
mib-2	OBJECT IDENTIFIER ::= { mgmt 1 } Events20 also provides MIB-2 functionality in addition to UNIX MIB. MIB-2 primarily relates to statistics concerned with various network protocols.

Enterprise-Specific Traps

AgentMon can issue Enterprise-specific traps. Each trap is manageable via SNMP. A trap is issued whenever the corresponding measured value traverses an alarm threshold (if one is specified) and the particular trap is enabled. Our Internet Enterprise number is:

1.3.6.1.4.1.548.1.2.2

The first part “1.3.6.1.4.1.548” identifies “our” SNMP domain and the remaining numbers refer to a specific MIB of ours, in this case, our UNIX MIB: Events20.

To enable and disable traps, use the trapmanage group from your Network Management Application. Changes take effect immediately and become the new start-up values should Events be restarted later. Only Network Administrator can enable and disable Traps, via SNMP. If a new test is added to Events, and that test has an associated trap, the trap will be disabled.

For instance, the blocksFree trap refers to the amount of free space on a filesystem. There is a single trap type (trap #12), but you can enable and disable each “instance type” (i.e., enable/disable per filesystem). The Enterprise-specific traps and the reasons they are sent (alarm conditions) are defined in the following table.

Trap Name	Alarm Condition
1. mbufclfree	Number of message buffer clusters
2. vmfree	Amount of virtual memory remaining
3. swapfree	Amount of free swap space
4. cpuload	Average length of the CPU run queue
5. cpuuser	Percentage of time CPU has spent in user mode
6. cpukernel	Percentage of time CPU has spent in kernel mode
7. cpuwait	Percentage of time CPU has spent in wait mode
8. procinstance	Number of instances of a process has changed
9. procsz	Size of a process has changed
10. proctime	CPU usage (time) of a process
11. procfree	Number of free process slots in kernel
12. fsbfree	Filesystem space (number of blocks free)
13. fsifree	Filesystem inodes (number of inodes free)
14. filesize	Size of a file
15. filemtime	Modification time of a file
16. fileatime	Time of last access of a file
17. queuedOldFiles	Number of old files in a directory
18. queuedFiles	Total number of files in a directory
19. hwchange	Change in the hardware inventory
20. swchange	Change in the software inventory
21. printerState	Change in the status of a printer
22. api1trap	User-provided data
23. api2trap	User-provided data
24. api3trap	User-provided data

Trap Name	Alarm Condition
25. api4trap	User-provided data
26. api5trap	User-provided data
27. api6trap	User-provided data
28. fileclamp	Regular expression match occurred in a file
500. utrap500	TRAP originating from an instance of the EventsCli
501. utrap501	TRAP originating from an instance of the EventsCli
502. utrap502	TRAP originating from an instance of the EventsCli
503. utrap503	TRAP originating from an instance of the EventsCli
504. utrap504	TRAP originating from an instance of the EventsCli
505. utrap505	TRAP originating from an instance of the EventsCli
506. utrap506	TRAP originating from an instance of the EventsCli
507. utrap507	TRAP originating from an instance of the EventsCli
508. utrap508	TRAP originating from an instance of the EventsCli
509. utrap509	TRAP originating from an instance of the EventsCli
510. utrap510	TRAP originating from an instance of the EventsCli
511. utrap511	TRAP originating from an instance of the EventsCli



Events Commands

The following pages are the UNIX-equivalent manpages for the various Events commands. This section describes the command-line options for each command and what that command will do. These descriptions are in the following order:

- AgentENL (8N)
- AgentMon (8N)
- EventsCli (8N)
- FieldCk
- testtab (1)

AgentENL (8N)

AgentENL - dynamically supports multiple SMUX subagents

SYNOPSIS

AgentENL [config_file]

DESCRIPTION

AgentENL is an SNMP/SMUX-compliant agent daemon. Subagents may be dynamically registered and de-registered without having to bring the AgentENL down. The UDP and IP port numbers defined within the relevant standards will be used unless the file /etc/services contains entries specifying alternate ports.

COMMUNITY entries determine which community strings are valid. If no COMMUNITY entries are present, ANY community will be accepted by the agent. For each community, the set of supported operations is specified by the ALLOW clause.

The MEMBERS clause in a community entry restricts the set of managers allowed to use that community. If the MEMBERS clause is absent from a COMMUNITY entry, any manager may use that community.

MANAGER entries determine where traps should be sent. If the PORT clause is absent, then TRAPS are sent to port 162, as specified in RFC 1157. If the WITH COMMUNITY clause is absent, the community string "public" is used.

OPTIONS

config_file

The filename (or pathname) of the configuration file. The configuration file should be edited to specify the SNMP manager hostnames, Community names, and privileges. The default is to use the configuration file \$ENLIGHTEN/config/AgentENL.config, where \$ENLIGHTEN represents the top-level product installation directory.

DIAGNOSTICS

Complains and exits if no configuration file is found or if the socket/port is already in use.

AgentMon (8N)

AgentMon - workstation monitoring program

SYNOPSIS

```
AgentMon [ -t testname ]
```

```
AgentMon -x
```

DESCRIPTION

AgentMon monitors tests as defined in the text file `testtab.{HOSTNAME}`. This `testtab` file is expected to be in `$/ENLIGHTEN/config`, where `$/ENLIGHTEN` represents the top-level product installation directory. Each test may also be remotely monitored and managed via SNMP. When a test value exceeds one or more of the user-defined thresholds, AgentMon may emit an Enterprise-Specific Trap, send an alarm via email, start a user-specified program that could possibly take corrective action, and/or notify PEP.

The local user may define and alter tests by editing `testtab.{HOSTNAME}`. The network manager may manage tests by using the SNMP SET commands from almost any SNMP-compliant network management application. All enterprise-specific traps may be individually enabled and disabled via SNMP. AgentMon does not rely on SNMP and will run without hindrance in non-SNMP environments.

OPTIONS

`-t testname`

Disregard the normal testing interval for the named test and run the test every 15 seconds. (See **RESTRICTIONS** later in this section.)

`-x`

Causes AgentMon to print (to `stdout`) its impression of the current test's configuration. The output is identical in form to the `testtab` file and represents the logical sum of the test defaults and the current contents of the `testtab` file. This is a good way of verifying AgentMon understands your `testtab` edits.

NOTES

The `testtab` file does not exist prior to the first invocation of AgentMon. It will be created immediately after start-up, however.

DIAGNOSTICS

Complains and exits if a connection with the AgentENL (or other agent) cannot be made. AgentMon will exit gracefully if communication with the agent goes away.

RESTRICTIONS

Only one test may be specified with the -t option. This is intended to be primarily a systems troubleshooting aid. The -t option has no effect on process monitoring; processes are always checked every 30 seconds.

SEE ALSO

AgentENL.8N testtab.5

EventsCli (8N)

EventsCli - Command Line Interface for sending alarms to PEP, SNMP, and so on.

SYNOPSIS

```
EventsCli -u units -v measured_value -n name -s severity -t trap_num -T  
threshold
```

DESCRIPTION

EventsCli is a user interface to the Events alarm dispatching facilities. It can be used to send SNMP traps, notify PEP, and log events. These events can be routed from external processes that call EventsCli each time a condition warrants reporting.

EventsCli is intended to be an interface between your existing system monitoring programs and scripts. These programs and scripts can call EventsCli and then rely on the SNMP traps, Events data logging, and PEP to take the corrective action you desire.

Any command line option that contains spaces should be quoted.

OPTIONS

- h hostname
The name of the event for which the event applies. The default is the local host name.
- u units
The units value specifies the unit of measure for the value measured.
- v measured_value
The value being reported. This can be any data type including a short string.
- n name
A brief name, which must be quoted if it contains white space. This could be a test name or other descriptive text.
- S setting state
Setting state describes the severity values; for example, “high”, “low”, or “OK”.
- s severity
Severity is a number between a low of 1 and a high of 5.

-t trap_num

Events has 12 general-purpose TRAPs for EventsCli. These are numbered from **500 to 511**. The trap number you choose can be arbitrary, but must be within the range indicated.

-T threshold

The name of the threshold exceeded. This is a very short description of what happened; that is, "low alarm threshold." In real life, any brief text is acceptable.

-V Total Value string

-q quiet mode-verbose off

When specified, EventsCli will act silently

-V total_value_string

A string or number representing the total value for the test. The default is -1.

-c application_class

The application class associated with the event. Default is 0.

The following application classes are used by AgentMon:

- 0 General product class
- 1 Kernal tests
- 2 MBUF tests
- 3 Cache tests
- 4 Virtual memory tests
- 5 Cpu tests
- 6 Process tests
- 7 Filesystem tests
- 8 File tests
- 9 Directory queue tests
- 10 Printer tests
- 11 Disk tests
- 12 Rpc tests
- 13 Api tests

- 14 Tcp tests
- 15 IP tests
- 16 UDP tests
- 17 ICMP tests
- 18 Inventory tests

-x Explanation flag.

When specified, EventsCli will output a usage statement and exit.

EXAMPLE

EventsCli -n DatabaseA -T "high tps" -u "tps" -v 12345 -s 5 -t 501

would produce the following comment and report:

DatabaseA is doing 12345 tps and has exceeded the high tps threshold.
The event will be reported with the highest severity level (5) and will
be sent using trap number 501.

RESTRICTIONS

EventsCli must be invoked on a machine that has AgentMon running.

FieldCk

FieldCk - a simple utility for verifying an API test.

SYNOPSIS

FieldCk

DESCRIPTION

When configuring an API test, you must specify a filename, which field (or column) contains the data, and which field (or column) contains the optional label. To check your data entry, you may use FieldCk to read the file and display the data and label fields.

FieldCk is started without options and will prompt you for all needed information. The file will be parsed with the fields you designated and the check-results will be sent to stdout.

If FieldCk parses the file correctly, you may be assured AgentMon will also do so.

testtab (1)

testtab - AgentMon capability database

SYNOPSIS

\$ENLIGHTEN/config/testtab.{HOSTNAME}

DESCRIPTION

testtab is the capability database for the AgentMon subagent. The format is similar to other UNIX system administrative files, but is simpler to use. AgentMon accesses this database when it starts up and frequently checks it for any changes.

At time of installation, this file may not yet exist. When no testtab file exists, all default tests and all test capability defaults are assumed and the testtab file is automatically created.

Each entry in this file represents a test description. A test description consists of a line containing the test name followed by additional lines containing a list of keywords and, possibly, values. There are three types of keywords (capabilities):

Boolean

Capabilities that have only two states (on/off, enable/disable, etc.) are specified by simply using that capability name to activate it or by prepending the name with an exclamation point ("!") to deactivate it.

Numerics and Strings

A capability that requires a numeric (or alphabetic) value is specified by separating the capability name and value with an equal sign.

All capabilities are delimited by colons (":"). Values that must contain a colon must be escaped with a backslash. Any test description line not terminated with a backslash ("\") is assumed to be the last line of that test description.

CAPABILITIES

Name	Type	Default	Comment
testfreq	int	5	Minimum testing interval (in minutes).
alarmfreq	int	60	Minimum time between alarms (in minutes).
notify	str	root	Specifies who receives any alarms.
mailer	str	/bin/mail	Specifies which program will deliver the alarm if a username was given in the notify field.
log	boolean	false	Indicates whether logging is enabled or not.
!log	boolean	true	Indicates logging is disabled.
delta	int/float	0	If logging is enabled, the most recent value measured will be logged if it differs by at least this amount from the previously logged value.
pep	boolean	false	If true, alarms are also sent to PEP.
age	int	0	For monitoring queues only. If not zero, only files more than "age" minutes old are counted as "old."
high	int/float	0	High-level alarm set point.
low	int/float	0	Low-level alarm set point.
+rate	float	0.0	Positive rate change set point.
- rate	float	0.0	Negative rate change set point.
+jump	int/float	0	Positive jump set point.
jump	int/float	0	Negative jump set point.

Name	Type	Default	Comment
command	str	null	This is a pathname to a script/executable.
regx_	R/E	null	regx1-regx31 are for File Clamping tests. They hold lists of regular expressions used in the test.
units	str	“units”	These values are preset and “changes” are silently ignored.

NOTES

Most tests measure integer values. Some measure floating-point values. Set point parameters should be given by using the same variable type as the test result. However, failure to comply will only result in rounding errors.

The jump and rate thresholds compare the current test value with the last measured value. Therefore, they are comparing change over time (as defined by testfreq).

The values for the unit’s capability are built in and cannot be changed. It is shown here for completeness, as AgentMon includes this field whenever it re-writes the testtab file and whenever it is started with the -x command line option.

If the notify field is set to “nobody”, that is, :notify=nobody:, then no email will be sent.



Sample Events Files

This appendix contains a sample Events testtab file and some examples of using the Events Command Line Interface (EventsCli) to send alarms to PEP, SNMP, and so on.

Refer to [Appendix F, “Events Commands,”](#) for more details on using the syntax for the testtab and EventsCli Events commands.

Sample testtab Configuration File

The following is an example testtab file with many of the optional test entries.

```
#
# This is an example testtab file. Not all test may be on all
# systems. Normally this file is automatically created. Changes
# to this file should be made using the GUI or via SNMP (with your
# network management app). However you may also edit this file.
# When editing this file, you may add tests anywhere, there is
# no specific order required. When Events updates this file, it
# WILL reorder the tests and any comments you added will be lost.
#

*****
  cache
*****
ncache | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:\
  :!pep:severe=3:!log:high=50.00:units=percent:delta=5.00:\
  :+rate=10.00:\
  :/* end ncache */:
*****
  cpu
*****
cpu load | :\
  :on:testfreq=1:alarmfreq=1:mailer=/bin/mail:notify=jerry:pep:\
  :severe=2:!log:high=5.00:units=units:delta=0.00:\
  :/* end cpu load */:
cpu idle | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=jerry:\
  :!pep:severe=3:log:units=percent:delta=1:\
  :/* end cpu idle */:
cpu kernel | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=jerry:pep:\
  :severe=3:log:high=70:units=percent:delta=1:\
  :/* end cpu kernel */:
cpu user | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=jerry:pep:\
  :severe=3:log:high=70:units=percent:delta=1:\
  :/* end cpu user */:
```

```

cpu wait | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=jerry:pep:\
  :severe=3:log:high=70:units=percent:delta=1:\
  /* end cpu wait */:
*****
  proc
*****
proc_slots | :\
  :on:testfreq=1:alarmfreq=60:mailer=/bin/mail:notify=root:pep:\
  :severe=3:!log:low=98:units=slots:delta=15:\
  /* end proc_slots */:
# restart a process if it dies
!nfsd instances | :\
  :on:mailer=/bin/mail:notify=root:pep:\
  :severe=3:log:low=3:high=5:command=/scripts/nfsd_restart:
  /* end !nfsd */:
# alarm if a process starts taking up too much memory,
# maybe it has a memory leak.
!myprog size |:\
  :on:high=500:+jump=25:notify=alice:
  /* end !myprog size */:
# alarm if a process becomes a cpu hog
!his_prog time |:\
  :on:high=200:log:delta=10:\
  /* end !his_prog time */:
*****
  fs
*****
/ blocks free | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\
  !pep:severe=2:log:low=48079:units=blocks:delta=100:\
  /* end / blocks free */:
/ inodes free | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\
  !pep:severe=2:log:low=12083:units=inodes:delta=0:\
  /* end / inodes free */:
/usr blocks free | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\
  !pep:severe=2:log:low=48079:units=blocks:delta=100:\
  /* end /usr blocks free */:
/usr inodes free | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\
  !pep:severe=2:log:low=12083:units=inodes:delta=0:\
  /* end /usr inodes free */:

```

```
/var blocks free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=9595:units=blocks:delta=0:\  
/* end /var blocks free */:  
/var inodes free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=2643:units=inodes:delta=0:\  
/* end /var inodes free */:  
/tmp blocks free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=4085:units=blocks:delta=0:\  
/* end /tmp blocks free */:  
/tmp inodes free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=612:units=inodes:delta=0:\  
/* end /tmp inodes free */:  
/opt blocks free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=76969:units=blocks:delta=0:\  
/* end /opt blocks free */:  
/opt inodes free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:\  
:!pep:severe=2:log:low=19257:units=inodes:delta=0:\  
/* end /opt inodes free */:  
# The /SWAP is a psuedo name and refers to total system  
# swap space regardless of the actual number of swap devices.  
/SWAP blocks free | :\  
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=tet:pep:\  
:severe=5:log:low=2:units=blocks free:delta=2000:\  
/* end /SWAP blocks free */:  
*****  
files  
*****  
# The size test could be used to automatically archive the many  
# system and other log files that can grow without bound.  
/var/adm/myapp.log size | :\  
:on:testfreq=60:alarmfreq=60:mailer=/bin/mail:notify=root:pep:\  
:severe=2:log:units=size:delta=0:high=200000:\  
:command=~/.policy/archive.sh:  
/* end /var/adm/sulog size */:  
# Alarm each time a file is accessed.  
/secure_data accessed | :\  
:on:testfreq=1:alarmfreq=1:\  
:-jump=1:+jump=1:pep:delta=1:log:  
/* end /secure_data accessed */
```

```

# Alarm each time a file is modified
/mystuff/setuid_script modified |:\
  :on:testfreq=1:alarmfreq=1:\
  :-jump=1:+jump=1:
  /* end /mystuff/setuid_script modified */
*****
  QUEUES
*****
# This group is good for detecting when an email, printer, or
# other directory based queue is getting backed up.

#This test alarms when there are two many files in the directory.
# It might be used to alert you when it's time to start shifting
# print jobs to another printer.
/usr/spool/lp/laser24 |:\
  :on:testfreq=5:alarmfreq=30:\
  :log:delta=10:\
  :high=100:\
  /* end /usr/spool/lp/laser24 */:
#This test is usefull when care about how many 'old' files are
# stuck in the queue. It is like the above test, but contains a
# value for 'age' (minutes).
# Specifically, this test will alarm when the mail queue has more
# than 20 files that are more than 30 minutes old.
/usr/spool/mqueue |:\
  :on:testfreq=5:alarmfreq=30:command=/usr/gordon/vickers:\
  :age=30:high=20:\
  /* end /usr/spool/mqueue */:
*****
  printers
*****
printers | :\
  :on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:\
  :!pep:severe=3:!log:units=...:delta=0:\
  /* end printers */:
*****
  rpc
*****
rpcc_calls | :\
  :off:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:\
  :!pep:severe=3:!log:units=calls:delta=100:\
  /* end rpcc_calls */:

```

```
rpcc_badcalls | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: !log: units=bad calls: delta=30: +jump=10:\
: /* end rpcc_badcalls */:
rpcc_retrans | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=retransmissions: delta=30: +jump=10:\
: /* end rpcc_retrans */:
rpcc_badxid | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=units: delta=0: +jump=10:\
: /* end rpcc_badxid */:
rpcc_timeout | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root: pep:\
: severe=3: log: units=units: delta=50: +jump=10:\
: /* end rpcc_timeout */:
rpcc_wait | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=units: delta=0: +jump=10:\
: /* end rpcc_wait */:
rpcc_newcred | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=units: delta=0: +jump=20:\
: /* end rpcc_newcred */:
rpcc_timers | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=units: delta=0: +jump=20:\
: /* end rpcc_timers */:
rpcc_calls | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=calls: delta=100:\
: /* end rpcc_calls */:
rpcc_badcalls | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=bad calls: delta=20: +jump=10:\
: /* end rpcc_badcalls */:
rpcc_nullrecv | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=units: delta=0: +jump=10:\
: /* end rpcc_nullrecv */:
rpcc_badlen | :\
: on: testfreq=5: alarmfreq=60: mailer=/bin/mail: notify=root:\
: !pep: severe=3: log: units=bad_length: delta=30: +jump=10:\
: /* end rpcc_badlen */:
```



```

rpcs_xdrccall | :\
:on:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:\
:!pep:severe=3:log:units=units:delta=100:\
/* end rpcs_xdrccall */:
*****
  api
*****
# to use the API's, collect data yourself and write it to a file.
# Then edit one of the entries below, setting "file=" to the
# pathname of your file, "data" to the field number your data is
# in, and "label" to the field number an optional label is in. You
# can use one file for one or more api tests or use separate files
# for each one.
api1 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api1 */:
api2 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api2 */:
api3 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api3 */:
api4 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api4 */:
api5 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api5 */:
api6 | :\
:off:file=:data=:label=:\
:testfreq=5:alarmfreq=60:mailer=/bin/mail:notify=root:!pep:\
:severe=3:log:units=units:delta=0:\
/* end api6 */:

```

```
*****
inventory
*****
software | :\
: on: testfreq=60: alarmfreq=1400: mailer=/bin/mail: notify=root:\
: !pep: severe=3: !log: units=units: delta=0:\
: /* end software */:
# The hardware test only gets done when Events is started.
hardware | :\
: on: mailer=/bin/mail: notify=root: !pep: severe=2: log:\
: units=units: delta=0:\
: /* end hardware */:
```

Sample EventsCli Use

This section shows examples of using the Events Command Line Interface (EventsCli) to send alarms to PEP, SNMP, and so on.

The first example:

```
EventsCli -c DatabaseA -T "high tps" -u "tps" -v 12345 -s 5 -t 501
```

means the DatabaseA is doing 12345 tps and has exceeded the high setting for the tps threshold. This event will be reported with the highest severity level (5). It is automatically reported to PEP and EMD, and an SNMP trap is sent using trap number 501.

The following example is a more general use of how to run EventsCli in a script.

```
#!/bin/sh
# This is one way the EventsCli might be used. This example program
# sends an alarm when it finds users that have been idle for more
# than one day. This script is OS specific, so you might have to
# make subtle changes for it to work on your system.
# This script could be called via crontab at what ever interval you
# deem appropriate, perhaps daily.
# This script is only an example, we don't claim it has any real
# usefulness nor is it warrented in any way.This is AS IS software.
#
TESTNAME="idle_user"
TESTVALUE=
UNITS=days
USERNAME=
SEVERITY=2 # we've arbitrarily chose this value.
TRAP=501 # any valid number will work, perhaps YOUR policy says
#         use 501 for this type of test.
export

# Begin Subroutines
# Alarm is called recursively
Alarm()
{
    case $1 in
        "" ) exit ;; # when there are no more arguments to process
        * ) USERNAME=$1 ;
            TESTVALUE=$2 ;
            echo EventsCli -n $TESTNAME -T $USERNAME -v $TESTVALUE
                -u $UNITS -s $SEVERITY -t $TRAP ;
            shift ;
            shift ;
            Alarm $* ;;
    esac
}

# end of subroutines
#
# MAIN starts here
#get a list of users and how long they've been logged in
#Pass the list to the Alarm() subroutine
# Syntex is Alarm username idletime username idletime ...
Alarm `w | awk '{ printf "%s %s\n", $1, $4 }' | grep $UNITS `
```



O/S Compatibility

This appendix contains a functional compatibility matrix illustrating the availability of AgentMon tests and functions across various operating system platforms.

Y means that the test is supported on that platform; **N/A** means that it is not supported.


Test	Group	Sun OS	Solaris	HP/UX	SCO	AIX	SGI	DEC
time of last reboot	OS	Y	Y	Y	Y	Y	Y	Y
local time	OS	Y	Y	Y	Y	Y	Y	Y
cpu load	OS	Y	Y	Y	Y	Y	Y	Y
kernel_cxt	OS	Y	Y	Y	Y	Y	Y	Y
kernel_traps	OS	Y	Y	Y	N/A	Y	N/A	N/A
kernel_syscalls	OS	Y	Y	Y	Y	Y	Y	Y
kernel_devinits	OS	Y	Y	Y	N/A	Y	Y	Y
forks	OS	Y	Y	Y	Y	Y	Y	Y
vforks	OS	Y	Y	Y	Y	Y	Y	Y
fork_pages	OS	Y	N/A	Y	N/A	N/A	N/A	N/A
vfork_pages	OS	Y	N/A	Y	N/A	N/A	N/A	N/A
blocks free	File System	Y	Y	Y	Y	Y	Y	Y
inodes free	File System	Y	Y	Y	Y	Y	Y	Y
printer state change	Printer	Y	Y	Y	Y	Y	Y	Y
proc_slots	Process	Y	Y	Y	Y	Y	Y	Y
hardware info	Inventory	Y	Y	Y	Y	Y	Y	Y
software info	Inventory	N/A	Y	Y	Y	Y	Y	Y
rpcc_calls	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_badcalls	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_retrans	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_badxid	RPC	Y	Y	N/A	Y	Y	Y	Y

Test	Group	Sun OS	Solaris	HP/UX	SCO	AIX	SGI	DEC
rpcc_timeout	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_wait	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_newcred	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcc_timers	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcs_calls	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcs_badcalls	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcs_nullrecv	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcs_badlen	RPC	Y	Y	N/A	Y	Y	Y	Y
rpcs_xdr call	RPC	Y	Y	N/A	Y	Y	Y	Y
vm_locked	VM	Y	Y	N/A	Y ^a	Y	N/A	Y
vm_claimed	VM	Y	Y	N/A	Y ^a	Y	N/A	Y
vm_free	VM	Y	Y	N/A	Y ^a	Y	Y	Y
cache_ctx	VM	Y	Y	N/A	N/A	Y	N/A	N/A
cache_seg	VM	Y	Y	N/A	N/A	Y	N/A	N/A
cache_pag	VM	Y	Y	N/A	N/A	Y	N/A	N/A
cache_par	VM	Y	Y	N/A	N/A	Y	N/A	N/A
cache_usr	VM	Y	Y	N/A	N/A	Y	N/A	N/A
cache_reg	VM	Y	Y	N/A	N/A	Y	N/A	N/A
mbufs	MBUF	Y	N/A	Y	N/A	Y	Y	Y
mbuf_clusters	MBUF	Y	N/A	Y	N/A	Y	Y	Y
mbuf_clfree	MBUF	Y	N/A	Y	N/A	Y	Y	N/A
mbuf_drops	MBUF	Y	N/A	Y	N/A	Y	Y	Y

Test	Group	Sun OS	Solaris	HP/UX	SCO	AIX	SGI	DEC
mbuf_space	MBUF	Y	N/A	Y	N/A	Y	Y	N/A
mbuf_wait	MBUF	Y	N/A	Y	N/A	Y	Y	N/A
mbuf_drain	MBUF	Y	N/A	Y	N/A	Y	Y	Y
name cache stats	Ncache	Y	Y	Y	Y	N/A	Y	Y
API Groups	API	Y	Y	Y	Y	Y	Y	Y
File Size	Files	Y	Y	Y	Y	Y	Y	Y
File Mod time	Files	Y	Y	Y	Y	Y	Y	Y
Directory Queue	Queue	Y	Y	Y	Y	Y	Y	Y
processes	Process	Y	Y	Y	Y	Y	Y	Y
ip_total	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_badsum	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_tooshort	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_toosmall	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_badhlen	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_badlen	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_fragments	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_fragdropped	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_fragtYmeout	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_forward	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_cantforward	MIB-II	Y	Y	Y	Y	Y	Y	Y
ip_redYrectsent	MIB-II	Y	Y	Y	Y	Y	Y	Y
icmp_error	MIB-II	Y	Y	Y	Y	Y	Y	Y
icmp_badcode	MIB-II	Y	Y	Y	Y	Y	Y	Y

Test	Group	Sun OS	Solaris	HP/UX	SCO	AIX	SGI	DEC
icmp_tooshort	MIB-II	Y	Y	Y	Y	Y	Y	Y
icmp_checksum	MIB-II	Y	Y	Y	Y	Y	Y	Y
icmp_badlen	MIB-II	Y	Y	Y	Y	Y	Y	Y
icmp_reflect	MIB-II	Y	Y	Y	Y	Y	Y	Y
tcp_psent	MIB-II	Y	Y	Y	Y	Y	Y	Y
tcp_bsent	MIB-II	Y	Y	Y	Y	Y	Y	Y
tcp_pgot	MIB-II	Y	Y	Y	Y	Y	Y	Y
tcp_bgot	MIB-II	Y	Y	Y	Y	Y	Y	Y
tcp_dropped	MIB-II	Y	Y	Y	Y	Y	Y	Y
udp_badhead	MIB-II	Y	Y	Y	Y	Y	Y	Y
udp_badsum	MIB-II	Y	Y	Y	Y	Y	Y	Y
udp_badlen	MIB-II	Y	Y	Y	Y	Y	Y	Y
udp_overflow	MIB-II	Y	Y	Y	Y	Y	Y	Y
file clamping	Files	Y	Y	Y	Y	Y	Y	Y
file accessed	Files	Y	Y	Y	Y	Y	Y	Y
cpu user	OS	Y	Y	Y	Y	Y	Y	Y
cpu kernel	OS	Y	Y	Y	Y	Y	Y	Y
cpu idle	OS	Y	Y	Y	Y	Y	Y	Y
cpu wait	OS	Y	Y	Y	Y	Y	Y	Y
swap blocks free	File System	Y	Y	Y	Y	Y	Y	Y
processes cpu time	OS	Y	Y	Y	Y	Y	Y	Y
processes mem size	OS	Y	Y	Y	Y	Y	Y	Y

a. Monitors similar resource. Exact resource not present.



SYSADMIN Commands

This appendix describes the various Sys Admin commands you can use via the Command Line Interface (CLI). These descriptions include an explanation of what that command will do and a summary of the command-line options for each command. The descriptions cover the following:

- netutar
- p_nsnaps
- p_snapdiff
- Filesystem Comparisons

netutar

You can also initiate backups and restores via the Command Line Interface (CLI). The program to invoke a backup or restore is called `netutar` and is located in the `$ENLIGHTEN/bin` directory. You can use the CLI version of Archive to run unattended nightly backups of all the systems when the system/network load is at the lowest use.

`netutar` is fully compatible with `tar`, but `netutar` has the added feature of verification. By specifying the `-v` option, `netutar` will compare the size of the image on tape with that of the original on disk. If the size matches, `netutar` will compare the backup with the original image byte for byte.

The actual invocation for `netutar` is

```
# ./netutar x
```

where `x` can be any combination of the following options:

`c` - Create tape archive

`f` - Input/output device name

File. This causes `netutar` to use the device argument as the name of the tarfile. If the name of the tarfile is '-', `netutar` writes to the standard output or reads from the standard input, whichever is appropriate. Thus, `netutar` can be used as the head or tail of a pipeline.

`t` - Get table of contents

Table of Contents. The names of the specified files are listed each time they occur on the tarfile. If no filename arguments are given, all the names on the tarfile are listed. With the `-v` function modifier, additional information for the specified files is displayed. The listing is similar to the format produced by the `ls -l` command.

`v` - Verbose output contents

Verbose. Normally, `netutar` does its work silently. This option causes `netutar` to output the name of each file it treats, preceded by the function letter.

`x` - Extract files from archive

Extract or restore. The named filenames are extracted from the tarfile and written to the current directory. If a named file matches a directory whose contents had been written onto the tarfile, this directory is

(recursively) extracted. Use the file or directory's relative path when appropriate or `netutar` will not find a match.

The owner, modification time, and mode are restored (if possible); otherwise, to restore an owner, you must be the superuser. If no filename argument is given, the entire content of the tarfile is extracted.



If several files with the same name are on the tarfile, the last one overwrites all earlier ones.

b - Device block size

Blocking Factor. This causes `tar` to use the block argument as the blocking factor for tape records. The default size is 20, the maximum is 2000.

The maximum blocking factor is determined only by the amount of memory available to `tar` when it is run. Larger blocking factors result in better throughput, longer blocks on nine-track tapes, and better media utilization.

k - Device size in Kb

K - Like 'k' but files are never split across volumes

F - Filename containing list of input files. ('-' stdin)

I - Incremental backup based on date of last modification

L - Name of logfile to be created

P - Only files from named partitions are to be archived

S - Skip files named `"*.o"`, `"a.out"`, and `"core"`

SS - Skip all the 'S' type files and skip any SCCS directories

T - Name of table of contents file

X - Output this explanation

p_nsnaps

This creates new disk snapshots from remote hosts. This command is found in:

```
$ENLIGHTEN/bin/p_nsnaps
```

The usage is

```
./p_nsnaps [-i] [-n] [-v] [-x] [-p partitions] [host | pool ...]
```

where the options mean:

-i	Include NFS partitions
-n	Output list of partitions to be captured only
-p parts	List of partitions to builds snapshots for
-v	Output list of partitions to be captured
-x	Output the usage of this command
hosts	List of hosts or network pool name
pool	Your network pool name

These snapshots are then stored in:

```
$ENLIGHTEN/work/<username>/__<hostname>__/<partition_name>
```

p_snapdiff

This shows the difference between two snapshots. This command is found in:

```
$ENLIGHTEN/bin/p_snapdiff
```

The usage is

```
./p_snapdiff [-F filename] [-d] [-e filename] [-i filename]
[-n|-l|-m] [-r] [-s val1-val2] [-u user] [-x] [-L] [-y]
[-G [+|-]group] [-M [+|-]mode] [-P] [-R] [-g group]
[-S [+|-]size] [-U [+|-]user] [-p file]
snap1 snap2 ... ..
```

where the options mean:

-F filename	Read arguments from filename
-d	Show content of new/lost directories
-e filename	Exclude files of this names
-g group	File must belong to user group
-i filename	Only include files of this name
-l	Find lost (removed) files
-m	Find modified files
-n	Find new files
-p file	Use file as alternate password file
-r	Look for (possibly) renamed files
-s val1-val2	Only look for file in this size range
-u user	File must belong to user
-x	Output this explanation
-y	Output summary of changes only
-L	Output filename in long format
-G +/-group	Group ownership changed to/from group
-M +/-mode	Permission modes changed
-P	File has been updated but has retained same size
-R	File has been read

-S +/-size	Size has changed by size bytes
-U +/-user	User ownership changed to/from user
filename	Name of file containing list of snapshot names
snap1	Name of first (newer) disk snapshot file
snap2	Name of second (older) disk snapshot file

Filesystem Comparisons

You can also compare file information from various heterogeneous hosts via the Command Line Interface (CLI). You can see the contents of lost or new directories; view modified, new, and removed files; or even read files.

You can use Sys Admin to check for user and group modified files or any file that has changed permissions. It can compare the current partition snapshot to a previously saved snapshot, or it can compare filesystems to other hosts' partitions running the `renld` Sys Admin agent daemon.

To take a snapshot, use the `p_nsnaps` command. To do a comparison of filesystems on different hosts, use the `p_snapdiff` command.

The following is an example of using the CLI to make snapshots of the `/usr` filesystem on three systems and then using the CLI to compare the new snapshots against a single `/usr` partition located on the system `apple`. This simple shell script can be started by cron at night and the output mailed to the specified user.


```
#!/bin/sh
ENLIGHTEN=`egrep "^enlighten=" /etc/enlighten | cut "-d=" -f2-`
SNAP=$ENLIGHTEN/bin
WORK=$ENLIGHTEN/work
LOG=$WORK/log
$SNAP/p_nsnaps -p /usr apple taz grumpy
echo `Summary of differences in /usr between apple and taz` > $LOG
$SNAP/p_snapdiff -y $WORK/__apple__/usr $WORK/__taz__/usr >> $LOG
echo `Summary of differences in /usr between apple and grumpy` >>
$LOG
$SNAP/p_snapdiff -y $WORK/__apple__/usr $WORK/__grumpy__/usr >>
$LOG
/bin/mail root < $LOG
```

The previous script mailed a copy of the summary output only (-y flag) to admin. The following mail message is the output of this script:

```
From root Tue Jun 20 18:37 PDT 1995
Date: Tue, 20 Jun 1995 18:37:18 -0700
From: root (0000-Admin(0000))
Apparently-To: admin
Summary of differences between taz and the master
USERNAMENEWLOSTREN+SIZE-SIZEMODEREADWRITE+USER-USERGRP
root 15107311341276102
bin60023501541327010
lp020610226000
uucp00000051000
adm00000001000
(UID=117)10000000000
sys00000001000
```

```
Summary of differences between apple and the master
USERNAMENEWLOSTREN+SIZE-SIZEMODEREADWRITE+USER-USERGRP
root11912408471091203367225
bin43712402090760406013221
(UID=101)01000000000
lp060261300184000
uucp000172008000
(UID=222)03000000000
adm00000001000
(UID=117)00020002000
sys00000001000
(UID=645)01000000000
```

The snapshots can also be compared to a single Master by specifying the directory of the Master snapshot instead of the current snapshot. To do so, change the pathname of the snapshot to the master directory:
`$ENLIGHTEN/work/master/hosts/taz/usr.`



EMD

The Enterprise Management Database (EMD) application is in charge of all communication between the database and other Enlighten**DSM** components. EMD consists of three components:

- Informix Database Engine with an ANSI-compliant database
- ODBC drivers (currently supporting only Informix)
- A daemon emdd that handles all requests from other Enlighten**DSM** components

An RPC-based server, emdd communicates with the following Enlighten**DSM** components: Events agent, Sys Admin DSA, Status Map, PEP, and the user interface.

The communication layer between emdd and the database is ODBC compliant. Our ODBC drivers are supplied by Visigenic Software Inc. The drivers are Core and Level 1 API-conformant and Minimum and Core SQL-conformant.

Basic Design

EMD registers all its procedures when it comes up and then waits for incoming requests. The emdd daemon also checks daily to see if any data in the database needs to be expired.

When servicing incoming requests, EMD (the emdd daemon) forks itself so it can handle more than one request at a time.

Relational Database

EnlightenDSM includes a relational database as part of EMD. The following data will be stored in the database:

- Pool configurations
- Session preferences
- Host overrides
- User authorizations
- Add user templates
- Archive device configurations
- Events log data
- Events alarm data
- Software and hardware inventory lists generated by events
- Acknowledgment of events from the status map
- Backup catalogs
- The scheduling of backups
- Host notes

The following data will not be stored in a database:

- Events host-specific `testtab` files
- Events `AgentENL.config` file data
- Snapshots of disks/filesystems

Directory Structure

The EMD directory structure is:

- *install-path/bin*
- *install-path/dbtables*
- *install-path/odbc*
- *install-path/informix*
- *install-path/msg*
- *install-path/log*

where:

<code>bin</code>	Contains the emdd daemon and its logfile <code>emdd.log</code>
<code>dbtables</code>	Contains the database tables, the database's transaction log, and a script to create the database tables
<code>odbc</code>	Contains the ODBC drivers and other related components
<code>informix</code>	Contains the Informix SE products
<code>msg</code>	Contains any message catalog files
<code>log</code>	Contains the emdd'd logfile <code>emdd.log</code>

Database Access

The database and its tables are owned by the user `dbenl`. This is a new user that is created at installation time. User `dbenl`'s home directory will contain an odbc initialization file `.odbc.ini`, and its `.cshrc` or `.profile` file will contain the environment variables necessary to run Informix and the ODBC driver.

By default, only user `dbenl` has access to the database. `dbenl` must start the daemon `emdd` or the daemon will not be able to access the database. The database administrator at your installation site may change the database privileges if he/she so chooses.

Warning!

Do not modify or rename the `.odbc.ini` initialization file. Doing so will compromise all EMD usage and connections fail.

Manually Killing/Invoking EMD

Sometimes you may want to manually kill or invoke EMD, for example, when you're bringing down the system for maintenance or upgrading to a new release of EnlightenDSM.



You *must* terminate EMD before backing up any database tables and then manually re-invoke them afterwards.

The rest of this section details how to do this for EMD.

Terminating EMD

You can bring down the `emdd` daemon with the following CLI command:

```
stop_enl_daemons
```

This sends an interrupt to the daemon and allows it to shut down gracefully. The daemon waits for any child processes to exit before shutting down. This may take up to 10 seconds.

Starting EMD

You can invoke the `emdd` daemon with the following CLI command:

```
start_enl_daemons
```

This starts the daemon in the background and sets the log level to 0. The logfile will generated, but only error messages will be logged (if any occur).

Audit Log Data

Audit Log Reports

e_audit_report

SYNOPSIS

\$EMD/bin/e_audit_report [-from *time-string* | **beginning**] [-to *time-string* | **now**] [-guihost *hostname*] [-managedhost *hostname*] [-dsmuser *username*]

DESCRIPTION

e_audit_report generates a report of dsm user audit information stored in the GMS. There are several arguments to this command that enable a user to query the database for specific information based on the time a action was inserted into the audit log, the host the audited operation was execute from, the host that the audited operation was execute on. and the dsm user name of the operator responsible for the audited operation. Any arguments that are specified are combined with logical *ands* to form a single search expression.

Executing **e_audit_report** without arguments prints all the audit information stored in the GMS sorted in chronological order.

Executing **e_audit_report** with a **-from** and/or a **-to** argument selects and prints the audit log entries that match audited operations executed during the times specified by a *time-string* in Enlighten Time Format. The *time-string* can represent an absolute time, such as Apr 15 1998 at 11:59pm, or it can represent a relative time in the past, such as '1 day' or '2 weeks' ago. Note that the command operates in the time zone set for the system.

Executing **e_audit_report** with a **-guihost** argument selects and prints the audit log entries that match audited operations executed from the gui host specified by *hostname*. A gui host is a machine which runs the EnlightenDSM gui, xenln.

Executing **e_audit_report** with a **-managedhost** argument selects and prints the audit log entries that match audited operations executed on the host specified by *hostname*. A managed host is a machine which runs the EnlightenDSM agents.

Executing **e_audit_report** with a **-dsmuser** argument selects and prints the audit log entries that match audited operations executed by a EnlightenDSM user specified by *username*.

e_audit_report prints the report to standard out and has the following format:

Time	User@Host	Target Host	Event
Nov 15 1998 at 8:03:15 am	johnd@chico	kansas	User account 'janed' successfully created.
Nov 15 1998 at 8:03:37am	bobbys@nypd	lapd	Pool 'cops' successfully modified.
Nov 15 1998 at 1:35:00 pm	johnd@chico	kansas	Printer 'anvil' successfully enabled.
Nov 15 1998 at 1:35:01 pm	johnd@chico	kentucky	Printer 'anvil' successfully enabled.

USAGE

The following options are supported:

-from *time-string*

-from beginning

Select audit log entries inserted since the time specified in the time-string. The time is specified in Enlighten Time Format. The keyword **beginning** may be used to specify the first chronological entry in the audit log. If this argument is used with the **-to** argument then the audit log entries occurring between the two times will be selected.

-to *time-string*

-to now

Select audit log entries inserted before the time specified in the time-string. The time is specified in Enlighten Time Format. The keyword **now** may be used to specify the last chronological entry in the audit log. If this argument is used with the **-from** argument then the audit log entries occurring between the two times will be selected.

-guihost *hostname*

Select audit log entries of audited operations executed from the gui host specified by *hostname*.

-managedhost *hostname*

Select audit log entries of audited operations executed on the managed host specified by *hostname*.

-dsmuser *username*

Select audit log entries of audited operations executed by the EnlightenDSM user specified by *username*.

EXAMPLES

example% e_audit_report

example% e_audit_report -from beginning

example% e_audit_report -from beginning -to now

All three of these commands are equivalent. They print all the connects stored in the Audit Log Table in the GMS.

example% e_audit_report -from Nov 30 1998 at 8:00 am

This command prints all the contents stored in the Audit Log Table in the GMS inserted since Nov 30 1998 at 8:00 am.

example% e_audit_report -from 1 week

This command prints all the contents stored in the Audit Log Table in the GMS inserted since 1 week ago.

example% e_audit_report -from 10 days -to 5 days

This command prints all the contents stored in the Audit Log Table in the GMS inserted between 10 and 5 days ago.

example% e_audit_report -from 1 hour

This command prints all the contents stored in the Audit Log Table in the GMS inserted within the last hour.

example% e_audit_report -dsmuser johnd

This command prints all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action taken by dsmuser johnd.

example% e_audit_report -guihost chico

This command prints all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action executed from gui host chico.

example% e_audit_report -managedhost kansas

This command prints all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action executed on managed host kansas.

example% e_audit_report -from Oct 1 1998 -to 3 days -dsmuser johnd -guihost chico

This command prints all the entries stored in the Audit Log Table in the GMS that were inserted between midnight Oct 1 1998 exactly 3 days ago as a result of an action taken by dsmuser johnd from gui host chico.

EXIT STATUS

- 0** All information was printed successfully.
- >0** Invalid arguments. Command failed.

SEE ALSO

Appendix C, Time Formats in the EnlightenDSM Reference Manual.

Purging Audit Log Data

e_rm_audit_events

SYNOPSIS

\$EMD/bin/e_rm_audit_events [-from *time-string* | **beginning] [-to *time-string* | **now**] [-guihost *hostname*] [-managedhost *hostname*] [-dsmuser *username*]**

DESCRIPTION

e_rm_audit_events selectively removes audit entries from the audit log store in the GMS. There are several arguments to this command that enable a user to query the database for specific information based on the time a action was inserted into the audit log, the host the audited operation was execute from, the host that the audited operation was execute on. and the dsm user name of the operator responsible for the audited operation. Any arguments that are specified are combined with logical *ands* to form a single search expression.

Executing **e_rm_audit_events** without arguments deletes all the audit information stored in the GMS.

Executing **e_rm_audit_events** with a **-from** and/or a **-to** argument selects and deletes the audit log entries that match audited operations executed during the times specified by a *time-string* in Enlighten Time Format. The *time-string* can represent an absolute time, such as Apr 15 1998 at 11:59pm, or it can represent a relative time in the past, such as '1 day' or '2 weeks' ago. Note that the command operates in the time zone set for the system.

Executing **e_rm_audit_events** with a **-guihost** argument selects and deletes the audit log entries that match audited operations executed from the gui host specified by *hostname*. A gui host is a machine which runs the EnlightenDSM gui, xenln.

Executing **e_rm_audit_events** with a **-managedhost** argument selects and deletes the audit log entries that match audited operations executed on the host specified by *hostname*. A managed host is a machine which runs the EnlightenDSM agents.

Executing **e_rm_audit_events** with a **-dsmuser** argument selects and deletes the audit log entries that match audited operations executed by an EnlightenDSM user specified by *username*.

USAGE

The following options are supported:

-from *time-string*

-from beginning

Select audit log entries inserted since the time specified in the *time-string*. The time is specified in Enlighten Time Format. The keyword **beginning** may be used to specify the first chronological entry in the audit log. If this argument is used with the **-to** argument then the audit log entries occurring between the two times will be selected.

-to *time-string*

-to now

Select audit log entries inserted before the time specified in the *time-string*. The time is specified in Enlighten Time Format. The keyword **now** may be used to specify the last chronological entry in the audit log. If this argument is used with the **-from** argument then the audit log entries occurring between the two times will be selected.

-guihost *hostname*

Select audit log entries of audited operations executed from the gui host specified by *hostname*.

-managedhost *hostname*

Select audit log entries of audited operations executed on the managed host specified by *hostname*.

-dsmuser *username*

Select audit log entries of audited operations executed by the EnlightenDSM user specified by *username*.

EXAMPLES

example% e_rm_audit_events

example% e_rm_audit_events -from beginning

example% e_rm_audit_events -from beginning -to now

All three of these commands are equivalent. They delete all the entries stored in the Audit Log Table in the GMS.

example% e_rm_audit_events -from Nov 30 1998 at 8:00 am

This command deletes all the entries stored in the Audit Log Table in the GMS inserted since Nov 30 1998 at 8:00 am.

example% e_rm_audit_events -from 1 week

This command deletes all the entries stored in the Audit Log Table in the GMS inserted since 1 week ago.

example% e_rm_audit_events -from 10 days -to 5 days

This command deletes all the entries stored in the Audit Log Table in the GMS inserted between 10 and 5 days ago.

example% e_rm_audit_events -from 1 hour

This command deletes all the entries stored in the Audit Log Table in the GMS inserted within the last hour.

example% e_rm_audit_events -dsmuser johnd

This command deletes all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action taken by dsmuser johnd.

example% e_rm_audit_events -guihost chico

This command deletes all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action executed from gui host chico.

example% e_rm_audit_events -managedhost kansas

This command deletes all the entries stored in the Audit Log Table in the GMS that were inserted as a result of an action executed on managed host kansas.

example% e_rm_audit_events -from Oct 1 1998 -to 3 days -dsmuser johnd -guihost chico


This command deletes all the entries stored in the Audit Log Table in the GMS that were inserted between midnight Oct 1 1998 exactly 3 days ago as a result of an action taken by dsmuser johnd from gui host chico.

EXIT STATUS

- 0** All information was printed successfully.
- >0** Invalid arguments. Command failed.

SEE ALSO

[Appendix C, “Time Formats”](#)



Macros

This appendix describes the use of macros for building user templates and accounts. A description of macros for mailing from form letters is also included.

User Account Macros

You can use macros to help initialize the form fields when you create a User Account. You can enter these macros manually. Or you can have them fill (some of) the form fields automatically by first defining them in a New User Template and using that template when you create the user account.

The following types of fields accept macros in these two windows:

- Login Name
- Home dir.
- Password
- Mail Alias

For working examples, see [“Macro Name” on page K-4](#).

Format

A macro has the following format:

`%<modifiers><macro name>`

Modifiers

<modifiers> are optional (you can define them multiple times). You can use the *<modifiers>* to specify only specific parts of the field, such as the first letter of the field, or to manipulate and convert the given input.

A modifier has the following format:

`[<converter><range>]`

where *<converter>* is one or more of the following characters:

- l - Convert the resulting string to lowercase
- L - Same as 'l'
- u - Convert the resulting string to uppercase
- U - Same as 'u'
- s - Do not include SPACES
- S - Same as 's'
- _ - Replace SPACES with underscores (_)

and the *<range>* applies to the complete string in the field (by default). Or you can apply a range to a specific word in the input string.

Consequently, a *<range>* has the following format:

$\$$ *<word number><character range>*

where *<word number>* shows the word's sequence in the string (this is optional). The first word would be addressed as \$1, the second word as \$2, and so on. You can also use \$0 to address the first word.

<character range> is a comma-separated list of character ranges in the form. If you specify a *<word number>*, you must first use a comma (',') in the range list.

The character range itself has the following format:

<start char>-<end char>

where *<start char>* and *<end char>* are character numbers. You can omit one end of the range to signify the "start of string" and "end of string", respectively. Ranges that go beyond the end of the string are truncated to match the string.

Macro Name

<macro name> is one of the following characters:

- g - replace with the users' GID
- u - replace with the user's UID
- G - replace with the users' group name
- U - replace with the user's login name (username)
- R - replace with data from the user's Real Name field
- O - replace with data from the user's Office field
- T - replace with data from the user's Telephone field

Examples

These definitions make more sense when they are viewed as examples. Assume for the following example, the following data is true:

```
Real Name:      Thomas Kraus
Login Name:    tmk
UID:          302
```

You could then define the following macros and see the following substitutions.

Field Name	Macro	Value
Home Dir	/home/%U	/home/tmk
Password	%U-%u	tmk-302
Mail Alias	%[u\$1]R_%[u\$2]R	THOMAS_KRAUS
Login Name	%[l\$1,1][l\$2,-6]R	tkraus


Some explanations:

<code>/home/%U</code>	Create the directory <code>/home/<login name></code>
<code>%U-%u</code>	Get the User name and add the UID (separated by a dash)
<code>%[u\$1]R_%[u\$2]R</code>	Get the first word of the user's real name and add the second word of the user's real name (separated by an underscore). Convert both strings to uppercase.
<code>%[l\$1,1][l\$2,-6]R</code>	Take the first letter of the user's first name and the first six letters of the user's last name. Convert both strings to lowercase.

Form Letter Macros

You can also use macros to define various text strings in a form letter. You need to enter these macros manually in the appropriate form letter. The following table shows the name and the effective substitution for each of these macros.

Macro	Function	Effect
%D	System Date	The system date in the format of Sun Mar16 10:32:01 1996 will be substituted for the macro.
%G	Group Name	The User's Group Name will be substituted for the macro.
%g	Groupid	The User's GID will be substituted for the macro.
%H	Home Dir	The User's HOME directory will be substituted for the macro.
%O	Office	The User's Office location will be substituted for the macro.
%R	Real Name	The User's Real Name will be substituted for the macro.
%T	Telephone	The User's Telephone number will be substituted for the macro.
%U	User Name	The User's Login Name will be substituted for the macro.
%u	Userid	The User's UID will be substituted for the macro.
%X	Report	The report displayed in the window will be included in the letter.
%x	Sub Report	Only the selected lines of the report displayed in the window will be included in the letter.



Windows Support

This appendix lists the features and options that Enlighten**DSM** either does not support on the Windows NT/95/98 platform, or that are implemented differently from the Unix platform.

The first section in this appendix describes global differences between systems running Enlighten**DSM** on Windows vs. those running Enlighten**DSM** on Unix. These differences affect the entire product, not just a single feature or module.

The rest of this appendix describes Windows differences in specific menus. This information is in the same order as both the Enlighten**DSM** GUI and the chapters in this manual.

Features not listed in this appendix work as documented by the rest of this manual. Refer to the appropriate chapter for more information about that feature.

The sections are:

- Global Differences
- Configure Menu
- User Menu
- Printer Menu

- Network Menu
- Security Menu
- Disk Menu
- Archive Menu
- System Menu
- Events Menu



Because changes have been made to Windows agents, it is necessary to use the latest release of Enlighten**DSM** on Unix workstations to administer Windows hosts. Some operations may appear to be functional from older software, but are, in fact, not supported.

Global Differences

Some global differences between Windows and Unix systems affect the entire product, not just a single feature or module. These global differences are described in the following section. The topics are:

- Unix and Windows File Pathnames
- File Ownership
- Directory/ File Access Times
- File Permissions
- File System Limitations

Unix and Windows File Pathnames

Unix and Windows operating systems have different and incompatible methods of specifying pathnames. EnlightenDSM supports only Unix pathnames. You can specify Windows pathnames in a Unix style by typing two forward slashes (//) before the drive designations and typing forward slashes throughout the rest of the pathname instead of back slashes. For example,

Windows and UNIX pathnames

Windows Pathname	Unix-compatible Pathname
C:\phonenet\spool	//c/phonenet/spool
\phonenet\spool	/phonenet/spool
d:\	//d/

Use this Unix-compatible pathname anywhere in the GUI that requires a pathname. The pathname is automatically converted as needed by Windows agents. Pathnames from agents (for example, in disk snapshots) are displayed in Unix format.

The only exception to entering Unix-compatible pathnames occurs on Windows NT in the user administration Home Directory text field. In this field, enter a Unix directory and a Windows NT directory name, separated by a vertical bar (|). You can enter Unix or Windows pathnames. The target machine will use the directory filename formatted in the same style as its operating system.

Windows filenames, user names, and group names may contain spaces; however, you have to add double quotes around the name (for example, “//c/my files”). The program will automatically display filenames, user names, or group names that contain spaces with double quotation marks.

Windows pathnames are case-insensitive, which can lead to unexpected results in some operations, like compare disk snapshots.

File Ownership

Unix and Windows NT operating systems use different and incompatible means of managing file ownership. NT systems uses Access Control Lists (ACL) that determine what an individual user can do with particular files. Two users with complete ACL privileges for a particular file each “own” the file, since there is no information attached to the file giving one individual ultimate ownership. Unix systems mandate that every file has a single owner, though others may have access privileges.

Some of the file scanning operations, like Snapshot, work differently on NT systems. Since there is no single owner, each file will appear to be owned by User 0 and Group 0.

Also, you cannot determine if an Windows file has been renamed because they do not have inode numbers which provide information about the file (except the name) and “handles” to the Unix filesystem that disk snapshots store and use.

Directory and File Access Times

Directory access times are not available in Windows disk snapshots because directory access times are updated when a snapshot is traversed on Windows. Disk snapshots on Windows were changed to store the last modification time instead of the last access time, which means that you cannot determine if a disk has been scanned, but just modified since the last snapshot. The last access times of Windows files work the same as Unix files, as long as you are using NTFS partition. The FAT file system does not track access times, but only modification times for both files and directories.

File Permissions

Unix files have three levels of permissions that can be set independently: read, write, and execute. As mentioned previously in the section, “File Ownership,” Windows uses the ACL method and does not modify the ACL. On Windows, EnlightenDSM will only modify the owner write permission, and not change the execute or read setting

Filesystem Limitations

Windows supports several filesystem types, including NTFS which was developed specifically for the NT operating system.

Most Microsoft filesystems that run on Windows are case-insensitive. See the beginning of this appendix for information about how to convert Windows filenames to Unix file names.

Configure Menu

The Session Preferences features are not supported. All other Configure menu features are supported.

User Menu

None of the User functions are supported on Windows 95/98 because that version of Windows does not support Users. EnlightenDSM for Windows supports all of the User menu features on Windows NT with the following exceptions:

User and Group ID's

NT does not use numeric IDs to uniquely identify users and groups. This may cause unexpected results for some operations because EnlightenDSM maps between names and numbers internally. To alleviate this problem, the program stores the provided UID number along with other information when creating an NT account. The “fake” ID is reported when the account is queried. For NT accounts not created with EnlightenDSM, a UID is created on the fly, creating a large number. This number serves as a place holder for missing information.

A similar process is used for Group ID's, but these GID numbers are not usually visible.

Configure Options

- The Span, Mail Alias, Mail Lists fields are not supported.
- The User ID and Expires options work in the Create New User Account menu.
- The Multi-user Add button is not functional because of Group Id issues.
- The Push NIS Maps option is not supported.

Query Options

- The GID Between, Shell, Password None fields are not supported.
- The UID Between field is functional for accounts created by the GUI, and for UIDs 0 and 1 (that is, administrator and guest IDs).

Groups Options

- The Group ID field is not supported.
- The NIS button is not supported.

Login Status Command

Login Status is not supported

Activity Monitor Commands

The Activity Monitor activates the following:

- Who Is Logged In
- Process Status
- CPU Summary

Who is Logged In

Logged in users are defined differently by Windows than Unix systems. NT reports remote mounts of a local filesystem as a user session and differentiates between local interactive users and remote users. EnlightenDSM makes a determination of when to report remote users, or not. The Status Map/Host Overview counts remote users, but the logout user screen only reports interactive users. If EnlightenDSM reports users on an NT machine, but they appear to be idle, then it may be providing file services to remote users.

- The Write button is not supported.
- The Message button is not supported.
- Domain user accounts will not be displayed in this window. Only locally defined accounts are shown.

Logging Out Users

EnlightenDSM can logout interactive users from an NT machine, but it does not provide any warning. Use this option carefully as it may cause lost data. You may want to reboot the machine instead as it provides a countdown warning.

Process Status

- The Hangup, Suspend, and Continue buttons are not supported.
- The Priority button is supported. Windows process priorities are displayed using the words High, Medium and Low instead of the priority numbers shown on UNIX. When changing the priority of a Windows process, use the following values to relate numeric values to High/Medium and Low priorities.
 - High priority is -19 to -10
 - Normal priority is -9 to 0
 - Idle priority is 1 to 19

Printer Menu

You can manage print queues (kill jobs, turn the printer on and off, etc.); however, the Move command and Custom Driver options are not supported. Creating a Windows printer client service and network printer server are not supported.

Network Menu

Local vs. Domain Accounts

EnlightenDSM does not recognize domains on Windows. Windows domains are not the same as Domain Name Service (DNS) domains. All user account administration by EnlightenDSM is performed on local accounts. To make changes to Windows domain accounts, you have to change the account domain controller. The account list of a Windows domain is the account list of the domain controller. Because the program does not recognize Windows domains, it is not possible to query a host to determine if it is a domain controller, or to find out the controller of a given Windows machine.

DNS Files

The Domain Name Servers (DNS) configuration screen provides a default pathname of `/etc/named` as a location for DNS files. This is correct for Unix machines, but when you configure an Windows NT machine as a DNS server, the correct path is `//c/winnt/system32/dns`. Because the EnlightenDSM GUI is unaware of the type of remote machine it will be talking to, it cannot provide a different default path for Windows NT targets.

Windows 95/98 machines cannot be configured as DNS servers.

Host Entries

Host Entries is fully supported only for DNS.

Naming Services Overrides Options

None of the Naming Services Overrides options are supported.

NIS and NIS+ Options

NIS and NIS+ is not a concept supported by Windows.

NFS Options

The NFS command features are not supported. NFS is not provided for Windows by Microsoft, so there is no standardized interface for administering it. Mounting and Exporting directories is not supported.

Security Menu

EnlightenDSM for Windows does not support any Security menu features.

Disk Menu

Usage by Filesystem Options

- The Processes button is not supported.
- Inode Number is not supported.
- Group Summary is not supported.
- chmod command:
 - File's owner execute command is not supported.
 - File's Group is not supported.
 - Other Users is not supported.
 - Special bits is not supported.
- chgrp, chown command not supported by the Windows operating system.
- Edit and View commands are not supported by Windows.
- The Processes command is not supported.

File Search Options

File Search Query

- # of Links Between field is not supported.
- Groupnames are not supported.
- File types options are not supported:
 - FIFO
 - Set UID bit
 - Socket
 - Block Special
 - Set GID bit
 - Symbolic link
 - Char special
 - Sticky bit

Processes by Filesystem Options

This feature is not supported.

Compare Snapshots

- Summary of renamed files not supported.
- Summary of new owners and new group owners is not supported.

Partition Differences Query

- Owner and Group names is not supported.
- User ownership and group ownership is not supported.

Archive Menu

Configure Devices Options

The tape operation of retention and erase are not supported.

Raw device should be set to `//a/filename` when configuring a floppy device for backup.

System Menu

TTY Configure Options

None of the TTY Configuration options are supported, as Windows does not support controlling TTY's. All Windows processes will report the TTY as N/A. Also note that TTY Configure is only available on the host running the GUI.

Shutdown/Reboot Options

Banner mode, Reconfigure, and Single-user mode shutdown options are not supported on Windows. Countdown is not supported on Windows NT.

Cron Management Options

Crontab management is not supported. NT provides a less sophisticated scheduled execution facility that is compatible with the EnlightenDSM interface. Future releases may provide an NT version of cron.

Mail Aliases

Mail aliases features are not supported. Windows does not use the Internet/sendmail model of e-mail distribution, rather it uses the Mail Application Programming Interface (MAPI) which uses the Address Books model.

Remote File Distribution

Remote File Distribution status logs may be incomplete for failed jobs.

The use of pre and post install scripts on Windows is not supported.

Login Status

Allowing and disallowing logins is not supported on Windows.

Events Menu

Events Monitoring

Events requires that the Microsoft SNMP package (included as a standard part of the Windows Workstation and Server distribution) to be installed. If this package is not loaded, all network related events will report as zero. The SNMP service daemon which is provided in the Microsoft SNMP package should **not** be enabled. Enlighten**DSM** provides its own SNMP agents, AgentENL and Agent Mon, which provide the actual SNMP service. Since the Enlighten**DSM** and Microsoft SNMP agents use the same SNMP network port, only one can run at a time. Enlighten**DSM** provides the same basic SNMP functionality as the Microsoft service, as well as enhanced DSM SNMP services.

Enlighten**DSM** Events

The list of built-in events is quite different for Windows and Unix. A subset of the very large performance statistics on the Windows has been selected. The majority of these events are TCP/IP related; the other are mostly CPU and memory related. The `cpu_busy` event approximately corresponds to the `cpu_load` event on Unix, but is reported as a percentage of non-idle time rather than the an average process run queue length, as on Unix.

There are also events for the user, system and interrupt time, reported as percentages. Note that these CPU values are samples, not slow moving averages. They tend to fluctuate more quickly than the Unix load average.

Process Tests

On UNIX, process test names may contain arguments to check for similar processes that are started with different arguments. This information is not available on Windows, so this type of test is not supported.

Process test filenames are case-insensitive on Windows.

Events E-mail Notification

Although the program does not support MAPI, Events provides a limited mailer that can send e-mail from an Windows machine to Internet addresses (e.g., user @host). This utility is called `smtpmail` and is called by AgentMon. It only supports internet style addresses, and cannot send to other types of addresses. This mailer does not queue and so if the target host is unreachable or not listening for mail, the message is lost.

MIBII Tests

These tests are not supported.

Status Map

The Host Overview Telnet feature is not supported on Windows.

Index

A

- About Enlighten window 2-21
- Absolute Date/Time BNF C-3
- Absolute Set Points
 - high 10-13
 - low 10-13
- Activity Monitor 3-32–3-36
- Add
 - a Group 3-7, 3-28, 3-40
 - Logical Device Name 8-3
 - Pools 2-11
 - Scheduled backup 8-19
- Add Domain Name Server 5-19
- Add Remote File Distribution Window 9-33
- Add User 3-3
 - Add Groups 3-7, 3-40
 - Default Start-up Shell 3-6
 - Expires 3-7
 - Office 3-4
 - Password 3-7, 3-41
 - Real name 3-4
 - Span 3-7
 - Telephone 3-4
 - User ID 3-5
- AgentENL 10-6
- AgentENL Command F-2
- AgentMon 10-6, 10-7
 - tests and functions by OS H-1
- AgentMON (8N) Command F-3
- alarm 10-37

- Alarm Messages 10-10
- Alarm Set Points 10-13
- Allow/Disallow Logins 9-40
- Alter Menu 2-8
- API Tests 10-8, 10-37
- Apply button 2-12, A-1
- Archive
 - Backup catalogs 8-14–8-17
 - Backup partitions 8-8–8-10
 - Configure Devices 8-2–8-6
 - Modify device configuration 8-6
 - Restore files 8-11–8-13
 - Restoring 8-12
 - Scheduling backups 8-18–8-24
 - Tape Operations 8-7–??
 - Viewing 8-11
- Archive Device Configuration window 8-2
- Assign Users
 - Add 3-30
- Auto refresh
 - options 2-4
- Automated Corrective Action 10-10

B

- Backing Up 8-18–8-24
- Backup 2-7
 - catalog name 8-14
 - Changed files 8-20
 - Command 7-12
 - Executing 8-23

- Full 8-8
- Incremental 8-8
- List of files 8-17
 - made before 8-14
 - made since 8-14
- Modifying schedule 8-22
- Print log 8-24
- Scheduling 8-21
- Status logs 8-24

Backups

- Full 8-9
- Incremental 8-9

Boot scripts

- Checks 6-5

C

- Chgrp Command 7-13
- Chmod Command 7-12
- Chown Command 7-13
- Clock Management 9-18–9-22
 - Setting 9-21
 - Synchronizing 9-18
- Communicating with SNMP systems 10-4
- Compare Snapshots 7-4, 7-14, 7-16, 7-29–7-41
- Components of Events 10-6–10-7
- Configuration 10-8
- Configuration File B-3
- Configuration File Parameters 10-11
- Configuration, Overriding Default B-3
- Configure 10-2–10-39
 - Add Devices 8-2
- Configuring
 - About Enlighten window 2-21
 - Add a Printer 4-2
 - Alter menu 2-8
 - Crontab jobs 9-9
 - EMD Data Expiration 2-19–2-20
 - Pools 2-10–2-13
 - Session Preferences 2-2–2-7
 - User Authorization 2-14–2-17
- Configuring Tests 10-8
- Continue 3-22, 3-33, 3-35, 9-2, 9-4
- Continue Command 3-35
- Copy
 - Print Destinations 4-2

- CPU Summary 9-2
 - Terminate 3-35
- Cron Management 9-9–9-17
 - Adding a job 9-10
 - Querying jobs 9-13
 - Users 9-15
- Crontab Contents 6-6
- Custom
 - Driver 4-11
- Custom Command 7-13

D

- Data Logging 10-9
- Default Configuration File B-3
- Delete
 - a device 8-2
 - a Group 3-29
 - a Printer 4-2
 - Directories from export list 5-48
 - Users from a group 3-30
- Delete button
 - Pool 2-13
- Devices
 - Delete 8-2
 - Modify 8-2
 - Type 8-3
- Directories Tests 10-35
- Directory Queues 10-37
- Disable
 - Printer 4-2
 - TTY Configuration 4-16
- Disk
 - Compare Snapshots 7-29–7-41
 - Exporting directories 5-44–5-49
 - File search 7-6, 7-18–7-22
 - Graph summary 7-6
 - Processes by Filesystem 7-23–7-25
 - Save Snapshots 7-26–7-28
 - Set UID/GID problems A-4
 - Usage by filesystem 7-3–7-7
 - Usage by groupname 7-15–7-17
 - Usage by username 7-8–7-14
 - Wizard A-4
- DMS
 - Adding mail exchangers 5-25
 - Deleting mail exchangers 5-27

- Modifying mail exchangers 5-26
- DNS 5-18–5-27
 - Adding 5-19–5-23
 - Adding a secondary server 5-21
 - Adding primary server 5-20
 - Deleting existing configuration 5-24
 - Mail exchangers 5-24
 - Modifying existing configuration 5-23
 - Server configuration 5-18
- Domain Name Servers. *See* DNS

E

- EMD J-1–J-6
 - Basic design J-2
 - Invoking and killing J-6
 - Relational database J-3
- EMD Data Expiration 2-19–2-20
- Enable
 - Printer 4-2
- Enterprise Management Database. *See* EMD
- ENTERPRISE Specific Traps E-6
- Erase 8-7
- Events
 - Alarm messages 10-10
 - Communications 10-4
 - Components 10-6
 - Configuration 10-8
 - Configuring Tests 10-8
 - Generating Reports 10-39
 - Host Overview 10-45–10-48
 - Inventory Tracking 10-4
 - Notification with Status Map 10-44
 - Overview 10-3
 - Practical use 10-4
 - Query Events 10-49–10-50
 - Querying 10-49–10-50
 - Status Map 10-40–10-44
 - testtab file 10-11
- Events Commands F-1–F-11
 - AgentENL F-2
 - AgentMON F-3
 - EventsCli F-5
 - FieldCk F-8
 - testtab F-9
- Events List 10-48
- EventsCli Command F-5

- Exporting Directories
 - Adding a directory 5-45
 - Copying parameters 5-48
 - Deleting 5-48
 - Export button 5-49
 - Modifying parameters 5-48
 - Mounting 5-49
- Exporting directories 5-44–5-49

F

- FieldCk Command F-8
- File
 - Decreased in Size 7-32
 - Increased in Size 7-32
 - Lost 7-32
 - Minimum/maximum size 8-14
 - Modified since/before 8-16
 - Read 7-32
 - Renamed 7-32
 - Search 7-3, 7-6, 7-18–7-22
 - to find 8-14
 - Types 7-22
 - Updated 7-32
 - Viewing 7-9
- File Lists
 - New File 7-32
- File Size 10-33
 - TRAP PDU 10-33
- File System 10-18
- File System Devices 6-4
- File System Tests 10-17
- File Tests 10-33
- Filenames
 - Find Like 8-14
 - Skip Like 8-14
- Files
 - View 7-4, 7-5, 7-16, 7-32
- files group 10-33

G

- Graph 7-3, 7-5, 9-4
- Group 3-28
 - Assign Users 3-30
 - Summary 7-3, 7-5
- Group File Integrity 6-8

Group Tests 10-8, 10-14

- / blocks free 10-17
- / inodes free 10-18
- cache_ctx 10-26
- cache_pag 10-26
- cache_par 10-26
- cache_reg 10-26
- cache_seg 10-26
- cache_usr 10-26
- cpu_load 10-15
- fork_pages 10-16
- forks 10-16
- icmp_badcode 10-29
- icmp_badlen 10-30
- icmp_checksum 10-29
- icmp_error 10-29
- icmp_reflect 10-30
- icmp_tooshort 10-29
- ip_badhlen 10-29
- ip_badlen 10-29
- ip_badsum 10-29
- ip_cantforward 10-29
- ip_forward 10-29
- ip_fragdropped 10-29
- ip_fragments 10-29
- ip_fragtimeout 10-29
- ip_redirectsend 10-29
- ip_tooshort 10-29
- ip_toosmall 10-29
- ip_total 10-29
- kernel_cxt 10-16
- kernel_devints 10-16
- kernel_syscalls 10-16
- kernel_traps 10-16
- mbuf_clfree 10-27
- mbuf_clusters 10-27
- mbuf_drain 10-27
- mbuf_drops 10-27
- mbuf_space 10-27
- mbuf_wait 10-27
- mbufs 10-27
- ncache 10-28
- printers 10-19
- proc_slots 10-20
- rpcc_badcalls 10-23
- rpcc_badxid 10-23
- rpcc_calls 10-23

- rpcc_newcred 10-23
- rpcc_retrans 10-23
- rpcc_timeout 10-23
- rpcc_timers 10-23
- rpcc_wait 10-23
- rpcs_badcalls 10-24
- rpcs_badlen 10-24
- rpcs_calls 10-23
- rpcs_nullrecv 10-24
- rpcs_xdr call 10-24
- tcp_bgot 10-30
- tcp_bsent 10-30
- tcp_dropped 10-30
- tcp_pgot 10-30
- tcp_psent 10-30
- udp_badhead 10-30
- udp_badlen 10-30
- udp_badsum 10-30
- udp_overflow 10-30
- vfork_pages 10-16
- vforks 10-16
- vm_claimed 10-25
- vm_free 10-26
- vm_locked 10-25

Groups 3-27–3-31

- Adding 3-28
- Adding users 3-30
- Deleting 3-29
- Deleting users 3-30
- IDs 3-28
- MIB II E-5
- Modifying 3-29
- Moving Users 3-31
- Names of 3-28
- SNMP E-2

GUI

- modifying B-3

H

- Hangup 3-22, 3-33, 3-35, 9-2, 9-4
- Hangup Command 3-34, 7-24
- Host Address Conflicts A-8
- Host Aliasing Conflicts A-8
- Host Configuration 5-7–??
- Host Entries 3-22, 5-2–5-6, 9-29
 - Adding 5-3–5-4

- Deleting 5-5
- Modifying 5-4
- NIS maps 5-6
- Overrides 5-15
- Host Entry Configuration Window 5-2
- Host Entry Name Conflicts A-8
- Host File Import 5-5
- Host Naming Services 5-8–5-11
 - DNS 5-10
 - Modifying 5-9
 - NIS naming 5-9
- Host Naming Services Configuration Window 5-8
- Host Overview 10-45–10-48
- Host Wizard A-8
- Hosts
 - Add 2-12
 - Address conflicts A-8
 - Aliasing conflicts A-8
 - Button 2-12
 - Entry name conflicts A-8

I

- ICMP 10-28
- Incremental Set Points
 - jump 10-13
- Interfaces group 10-28
- Inventory Tests 10-21
- Inventory Tracking 10-4
- IP 10-28
- Item tests 10-8, 10-32
 - Directory Queues 10-32
 - File Size 10-32
 - Processes 10-33

L

- License Advisory window 1-3–1-4
- limits E-2
- Lock a User 3-19
- Logical Device Name 8-3
- Login
 - allowing 9-40
 - disallowing 9-40
 - Who is logged in 3-32
- Logout 3-32
 - Command 3-33

M

- Mail Aliases 9-23–9-26
 - Adding 9-25
 - Modifying 9-27
 - Overrides 5-16
 - Querying 9-29
 - Viewing 9-29
- Mail Exchangers 5-24
 - Adding 5-25
 - Deleting 5-27
 - Modifying 5-26
- Main Menu Window B-14
- Masked Passwords 2-6
- Master Snapshots A-4
- MBUF Tests 10-27
- Menu Items B-9
 - Altering B-1
 - Creating B-5
- Menus
 - Attributes B-5
 - Defining B-5
 - Defining items B-9
 - Item attributes B-9
 - Main B-14
 - Sample Utility B-8
 - Utilities submenu B-12
- Message 3-32
 - Command 3-33
- MIB Group 10-15, 10-16, 10-17, 10-19,
10-20, 10-21, 10-23, 10-25, 10-27,
10-28, 10-29, 10-31
 - File System 10-17
 - Inventory 10-21, 10-22
 - Limits 10-31
 - MBUF 10-27
 - Ncache 10-28
 - OS 10-15
 - Printer 10-19
 - Process 10-20
 - RPC 10-23
 - TrapManage 10-31
 - VM 10-25
- MIB II Groups E-5
- MIB II Tests 10-28
- Modify 4-16
- Modify Host Naming Services Window 5-8

Modifying Tests 10-11

Monitor User Processes

 Terminate 3-22

monitoring 10-11

Move 3-30

Multi-User Add 3-9–3-14

 Default HOME directory 3-10

 Default Password Configuration 3-11

 Default Start-up Shell 3-11

 Default User Group 3-10

 Initial Password 3-12

 Password to expire 3-12

 Span between changes 3-12

 Starting Userid 3-10

N

Name Cache, 10-27

Names

 Duplicate A-6

Naming Service Configuration 5-12–5-17

 Adding 5-13–5-16

 Alias file 5-16

 Deleting 5-17

 Group overrides 5-14

 Host overrides 5-15

 Mail alias overrides 5-16

 Modifying 5-17

 Password 5-14

Naming Service Overrides Window 5-12

Ncache Tests 10-27

Network

 DNS 5-18–5-27

 Host Configuration 5-7–??

 Host entries 5-2–5-6

 Host file import 5-5

 Host Naming Services 5-8–5-11

 Naming Service Configuration 5-12–5-17

 NIS 5-28–5-36

 Remote File Distribution 9-32–9-39

Network File System. *See* NFS

Network Host Configuration Window 5-7

Network Information Service. *See* NIS

network manager 10-31

NFS 2-6, 5-37–5-49

 Adding a partition 5-39

 Create a new partition 5-42

 Modifying mount parameters 5-42

 Mounted directories 5-37

 removing directories 5-42

NIS 5-6, 5-28–5-36

 Adding a server 5-29–5-31

 Deleting server configuration 5-34

 Master server 5-29

 Slave server 5-30

NIS Maps 5-6

NIS+

 Adding a server 5-31

 Deleting server configuration 5-34

 Master server 5-32

 Replica server 5-33

number of packets 10-28

 bad 10-29

 sent and received 10-28

number of protocol errors 10-29

O

OBJECT IDENTIFIER E-2

Obvious Passwords 6-11

P

Password

 Masked 2-6

 Obvious A-6

Passwords

 Add user 3-7

 Aging D-1

 Berkeley UNIX and pre-System V.4 D-2

 Common names 6-12

 Common words 6-12

 Default 3-41

 expire user 3-20

 Expiry 3-41

 File Integrity 6-7

 Life Span 3-41

 Masked 2-6

 Minimum length 3-42

 Modify 3-18

 Naming service configuration 5-14

 Obvious 6-12

 Security checks 6-13

 System V.4 D-3

PDU 10-35

Percent Change Set Points
 rate 10-13

Pool
 Adding 2-11
 Button 2-12
 Hierarchy 10-44
 Modifying 2-12
 Selecting 2-13
 Viewing 2-13

Pool Configuration 2-10–2-13

Port Configuration 4-17

Preferences 2-2–2-7

Printer
 Adding Local 4-2
 Canceling jobs 4-9
 Class 4-2
 Cross-Platform 4-6
 Delete 4-7
 Description 4-11, 4-12
 Devices 4-2
 Disable 4-8, 4-9
 Disabled A-7
 Driver 4-11
 Enable 4-8, 4-9
 Interface 4-11, 4-13
 Jobs 4-8
 Large jobs A-7
 Local Queue Name 4-6
 Long queues A-7
 Model 4-2
 Model Name 4-12
 Move (Print Job) 4-10
 Name 4-2
 Options 4-14
 Queue 4-9–4-10, A-7
 Remote Host 4-6
 Use 4-11, 4-12
 Wizard A-7

Printer Configuration 4-2

Printer Tests 10-19

Priority 3-22, 3-33, 3-35, 9-2, 9-4

Priority Command 3-35, 7-25

Process Status 3-34–3-35

Process Tests 10-20

Processes 3-32, 7-3, 10-36
 CPU 3-36
 Status 3-34

 Terminate 3-33, 9-2, 9-4

Processes by Filesystem 7-23

Processes Command 7-6

processes group 10-36

Processes Tests 10-36

Q

Query 3-23–3-26
 GID between ... and 3-24
 Home dirs 3-24
 Hosts 10-49
 Logged in 3-25
 Messages 10-50
 Not Logged in 3-25
 Office 3-25
 Password 3-25
 Real Name 3-25
 Shell(s) 3-24
 Telephone 3-25
 Test name filter 10-50
 UID 3-24
 User Group 3-24
 User Names 3-24

Query Events 10-49–10-50

Query Events Messages Window 10-43

queue group 10-35

Queue, printer 4-9

R

Raw Device Name 8-3

Rebooting the System 9-5–9-8

Rebuild 7-3

Refresh
 Button 2-4
 Options 2-4

Relational Database
 Access J-4
 Data types stored J-3
 Directory structure J-4

Relative Date/Time BNF C-4

Remote File Distribution 9-32–9-39
 Adding 9-33–9-37
 Copying 9-38
 Deleting 9-38
 Executing 9-38
 Install scripts 9-34

- Install to host 9-34
 - Installing to host 9-33
 - Modifying 9-37
 - Print log 9-39
 - Relay hosts 9-35
 - Scheduling 9-36
 - Status logs 9-38
 - Reports
 - Events 10-39
 - Restore
 - Archives 8-11, 8-12
 - Viewing 8-11
 - RFC 1155 10-7
 - RFC 1157 10-7
 - RFC 1212 10-7
 - RFC 1213 10-7
 - RFC 1215 10-7
 - RFC 1227 10-7
 - RPC Tests 10-23
- S**
- Save Current Snapshots 7-26
 - Scheduled Backup Configuration Window 8-18
 - Search 7-34
 - Altered files 7-37
 - Attempt to determine renamed files 7-37
 - Button 7-34
 - Files 7-18–7-22
 - Group names 7-36
 - Group Ownership 7-39
 - Maximum file size 7-36
 - Minimum file size 7-36
 - Mode 7-39
 - New files 7-37
 - New/lost dirs 7-37
 - Owner names 7-36
 - Ownership 7-39
 - Size 7-38
 - Security
 - Attempted break-ins 6-10
 - Boot/shutdown scripts 6-5
 - Cron users 9-15
 - Crontab contents 6-6
 - Disk Wizard A-4–A-5
 - File system devices 6-4
 - Group file integrity 6-8
 - Home directories 6-9
 - Logfile 6-2
 - Obvious passwords 6-12
 - Password file integrity 6-7
 - Vital files 6-3
 - Select button
 - Pool 2-13
 - Server Configuration 5-18
 - Session Preferences 2-2–2-7
 - Set or modify a password 3-18
 - Shell
 - start-up 3-39
 - Shutdown scripts
 - Checks 6-5
 - Snapshots
 - Building 7-2
 - Compare 7-3, 7-4, 7-5, 7-6, 7-14, 7-16, 7-29–7-41
 - Graph button 7-5
 - Save 7-3, 7-4, 7-5, 7-6, 7-14, 7-16, 7-26
 - Search 7-34–7-41
 - Summary 7-31
 - SNMP 10-4
 - SNMP Groups E-2
 - api1,...,api6 E-4
 - cpu E-2
 - files E-3
 - fs E-3
 - hardware E-3
 - limits E-2
 - mbuf E-2
 - ncache E-2
 - os E-2
 - printer E-3
 - proc E-3
 - queue E-3
 - rcp E-3
 - software E-3
 - trapmanage E-2
 - Traps E-6
 - vm E-2
 - SNMP Managed Groups E-1
 - SNMP trapmanage group E-6
 - Starting EMD J-6
 - Status Logs 8-24, 9-38

- Status Map 10-40–10-44
 - Editing 10-41
 - Event notification 10-44
 - Information areas 10-44
 - Navigating 10-41
- Suspend 3-22, 3-33, 3-35, 9-2, 9-4
- Suspend Command 3-35, 7-25
- Swap Space 10-46
- Swap Space Usage 9-4
- System
 - Clock Management 9-18–9-22
 - CPU Summary 9-2
 - Cron Management 9-9–9-17
 - Mail aliases 9-23–9-31
 - Reboot 9-5–9-8
 - Swap space 9-4
 - TTY Configuration 4-16–4-20
- System group 10-28

T

- Tape Operations
 - Retention 8-7
- TCP 10-28
- Terminate 3-34
- Terminate Command 3-34, 7-24
- test off, 10-9
- tests 10-8
- testtab Command F-9
- testtab file 10-11
- testtab file configuration parameters
 - !log 10-12
 - age 10-13
 - alarmfreq 10-11
 - log 10-12
 - mailer 10-12
 - notify 10-12
 - testfreq 10-11
- Time Formats C-1–C-4
- Trap types
 - cpuLoad E-7
 - mbufClfree E-7
 - quOldFiles E-7
 - quQuedFiles E-7
 - vmFree E-7
- Traps E-6

- TTY Configuration 4-16
 - Add 4-16
 - Baud rate 4-18
 - Device Type 4-18
 - Disabling the port 4-20
 - Enabling the port 4-19
 - From Port 4-17
 - Modifying the port 4-19
 - Terminal Type 4-18
 - To Port 4-17
 - TTY Name 4-18

U

- UDP 10-28
- Usage By Filesystem
 - User summary 7-3
- Usage by Groupname 7-15–7-17
- Usage by Username 7-8–7-14
- User
 - Activity Monitor 3-32–3-36
 - Add 3-3
 - Adding a user 3-3–3-9
 - Adding multiple 3-9
 - Delete 3-18
 - Home Directories 6-9
 - Lock 3-19
 - Login status 9-40
 - Modify 3-17
 - Query 3-23–3-26
 - Summary 7-3
 - Unlock 3-19
 - User Templates 3-37–3-44
 - View Information 3-20
- User Authorization 2-14–2-18
- User Templates 3-37–3-44
- User Wizard A-6
- Userids
 - Duplicate A-6

V

- View
 - Files 7-4, 7-5, 7-9, 7-16, 7-32
 - Print Jobs 4-2
 - User Information 3-20
- Vital Files 6-3
- VM Tests 10-24

W

Who is Logged In

Write command 3-32

WINN.ORG B-3

WINN.OUT B-3

WINOUT_NET B-3

Wizard

button A-1

Disk A-4–A-5

Host A-8

Printer A-7

User A-6

Write command 3-33