Enlighten®
Software Solutions

EnlightenDSM**/**Workgroup

EnlightenDSM**/**Advanced

User Guide

Release 3B

Copyright

Trademarks or Service Marks

Contents

CHAPTER 6: Fine-tuning Events with PEP

APPENDIX A: EnlightenDSM Basics

APPENDIX D: De-installing Enlighten**DSM**

Index

Preface

The contents of this *User Guide*, the *Reference Manual*, and Enlighten**DSM** on-line help contain all the information you need to take full advantage of Enlighten**DSM** software. Electronic versions of both the *User Guide* and the *Reference Manual* are also located in the /doc directory of the installation CD-ROM (see the readme.txt file for information on accessing these electronic versions).

- This Enlighten**DSM** *User Guide* contains step-by-step instructions for setting up, maintaining, and monitoring systems on a network. If you are new to network administration, start at the beginning of this manual and work through the chapters sequentially. If you are an experienced network system administrator, you can use the *User Guide* as a reference tool to find specific information.

- For instructions on installing the Enlighten**DSM** software, refer to Appendix C in the Enlighten**DSM** *User Guide*. These instructions include system requirements for running the software.

- The Enlighten**DSM** *Reference Guide* provides a concise summary of all Enlighten**DSM** features.

- Enlighten**DSM**'s on-line help system provides easily accessible information on each feature. To access the on-line help system, click the Help icon that appears in the top right corner of most Enlighten**DSM** screens.

This user guide is organized into the following chapters:

Chapter 1, "Before You Begin," briefly describes the features of the Enlighten**DSM** program.

Chapter 2, "Getting Started," describes how to set up Enlighten**DSM** program preferences, create network pools, and configure user authorizations.

 Chapter 3, "Setting Up a Basic Network Configuration," describes how to add and modify user accounts, create groups, and configure network printers.

Chapter 4, "Managing the Network Systems," describes how to manage disk partitions, NFS servers, hosts, cron jobs, and archive files.

Chapter 5, "Monitoring Networked Systems with Events," describes how to use Enlighten**DSM** Events and Status Map features to gather, monitor and interpret system information. This chapter contains instructions on using these tools to view host and pool status, audit security checks, and monitor logins, processes, CPU and printer usage.

Chapter 6, "Fine-tuning Events with PEP," contains instructions for using the Programmable Events Processor (PEP) and Enterprise Management Database (EMD). The Programmable Events Processor (PEP) allows you to centrally manage Events so that it reports only on issues you predetermine. EMD is an application manager in charge of all communication between the database and other Enlighten**DSM** components

Appendix A, "EnlightenDSM Basics," describes the use of global function buttons in Enlighten**DSM** windows and how to perform actions on list items. This chapter also contains information on network pools, creating defaults and performance expectations.

Appendix B, "Troubleshooting," lists error messages, describes their significance, and suggests possible solutions. This section also describes solutions to some common problems you might encounter when using Enlighten**DSM**.

Appendix C, "Installing EnlightenDSM," contains step-by-step instructions for installing Enlighten**DSM**.

Appendix D, "De-installing EnlightenDSM," describes how to de-install Enlighten**DSM**.

The Enlighten**DSM** documentation assumes that you are familiar with basic UNIX commands and environments. For information about your workstation or operating system, refer to the manufacturer's documentation.

1

Before You Begin

Welcome to Enlighten**DSM**™, a scalable, standards-based application enabling you to easily manage complex network environments.

Enlighten**DSM**'s numerous event monitoring and system administration features are managed through the Graphical User Interface (Figure 1-1).
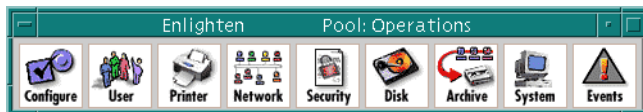
Figure 1-1   Graphical user interface

This cross-platform GUI simplifies the management of distributed, corporate intranets running various operating systems from many vendors. Complex, multi-step operations are reduced to more intuitive single functions, reducing the chance of error. You spend less time learning platform-specific protocols and more time effectively managing the network.

Using Enlighten**DSM**'s system administration features, you can manage users and disks, perform backups and retrievals, configure security features, manage processes, and control printer functions throughout your network.

Enlighten**DSM** also allows you to monitor, suspend or terminate processes on an intranet-wide basis.

Enlighten**DSM**'s Events management capabilities can automatically collect and save status, configuration, performance, and capacity information. As an SNMP agent, Enlighten**DSM** captures error messages and alerts, and provides that information to the SNMP manager (such as SunNet Manager, HP OpenView, or IBM NetView).

Enlighten**DSM**'s Events feature also enables you to monitor devices, processes, resources, and services on networked systems. This feature reduces network traffic by eliminating the need for hundreds of polling messages. Instead, you are notified via the Status Map or an SNMP-aware application only when a characteristic reaches an alert stage whose threshold you determine. Status maps can be customized with meaningful backgrounds, such as office floor plans or geographical maps.

Because Enlighten**DSM** is scalable, it can easily be configured to manage a broad spectrum of networks, from smaller ones with just a few hosts, to extensive corporate systems.

Enlighten**DSM**'s key features include:

- Enterprise Management Database (EMD)

- Centralized event configuration management

- Programmable events

- DNS, NIS, NIS+, setup and management

- NFS management

- Network printer/spooler management

- Hardware/software asset tracking

- Security auditing

- Disk/file system management

- Logical resource pooling

- User account templates

- Crontab management

- Backup scheduler

- Remote software distribution/installation

- Application logfile monitoring

- Scalable, customizable user interface

- Node health at a glance

1
Before You
Begin

2

Getting Started

Before configuring your network, you may want to customize Enlighten**DSM** by setting program preferences, creating network pools, and setting user authorizations. Follow the instructions in this chapter for directions on doing so.

If you'd rather use the program's default settings, skip to Chapter 3, "Setting Up a Basic Network Configuration." You can always change program preferences, set up network pools or modify user authorizations at a later time.

☞ Enlighten**DSM** contains certain features that are common to most of the program's windows.

Setting Session Preferences

Choose Session Preferences from the Configure menu to display the Session Preferences window (Figure 2-1).
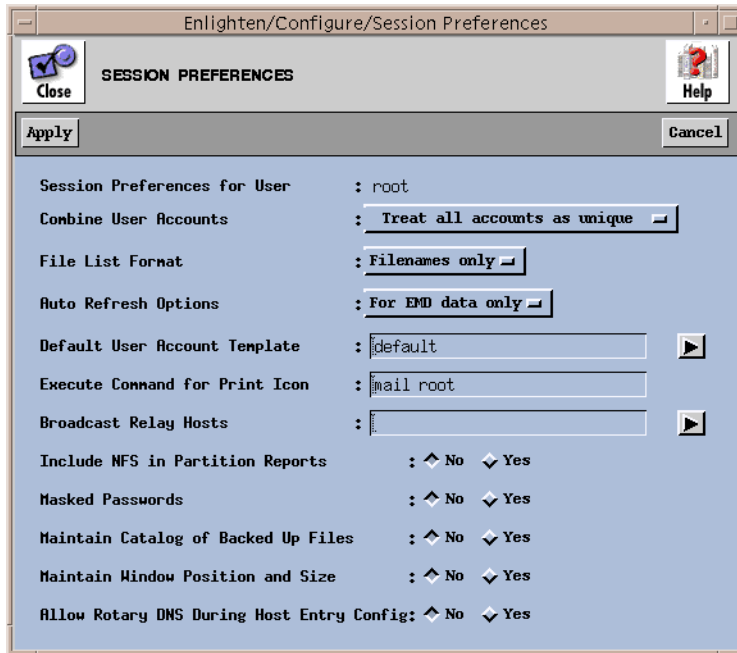
Figure 2-1 Session Preferences window

Each Session Preference option is described below. Click the Apply button to save the changes you make. Click the Cancel button to restore the default settings.

Session Preferences for User field

Displays the user for whom the Session Preferences are being defined.

Combine User Accounts option

Enlighten**DSM** can combine accounts with identical parameters across the
current System Pool and report them as a single entry.

For example, the entry

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

means that the user `caesar` has accounts on three of the systems in the
current System Pool. All three accounts have the same User name, User id,
primary Group id, 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 cases, you may want to see a more detailed report of user
configurations. To report each account separately, select the Treat All
Accounts as Unique option. The uniqueness of each account is determined by
three parameters: Username, Userid, and primary Groupid.

To determine how Enlighten**DSM** reports account configurations, choose one
of the following three options:

- Treat all accounts as unique (the default)

- With identical UserID & GroupID

- With all parameters identical

File List Format option

The File List Format option determines what information is displayed when
filenames are listed in a window:

- The Filenames Only option (default) displays only the name of the
  file.

- The Long Listing option displays additional information, such as the
  protection mode, file size, the date of last modification, and the
  absolute pathname.

2 **Getting Started**

Auto refresh options

In multi-user and multi-tasking systems, files are created and deleted, print requests are made, and processes start and terminate. Because of this, the system information displayed might not reflect the actual state of the system. Clicking the Refresh button updates the screen with current information.

Enlighten**DSM** can automatically refresh a screen after tasks are performed or changes made:

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

- Choose the For EMD Data Only option (the default setting) to automatically update screens where the data required to build the report is not local Enlighten**DSM** data (that is, all EMD data). See the section "Relational Database" on page 27 for list of data types stored in the EMD.

- Choose Turned On to automatically refresh the current screen after taking action on items in display lists. The screen refreshes when all actions have been completed in that window. This option is particularly useful when killing a process. After Enlighten**DSM** sends the kill signal to the process, the screen is automatically refreshed, allowing you to quickly determine if that 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. During Refresh, the parent window is first regenerated and the same items are reselected based on their position in the list. So, if the first and third entries are selected in the original report, they are also selected in the newly generated report. After the parent window regenerates, the current window is also refreshed.

Disk-related windows do not have refresh capabilities since they are based on previously generated disk snapshots. See "Building Snapshots" on page 8 for more details on using disk snapshots.

Default User Account Template field

User Account Templates simplify the creation of user accounts by allowing you to use a pre-defined set of parameters. Click the right arrow button to choose from a pick-list of previously defined templates.

Execute command for Print Icon field

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

The default command will email the report list to the user `root`, but you may use this field to create any executable command list. To print to the printer `text`, for example, you could specify the following command:

```
lp -d text
```

A maximum of 100 characters can be used for the custom printing command.

Broadcast Relay Hosts field

Specify any hosts to use as broadcast relays for remote file distribution jobs. For multiple entries, leave a space between each hostname.

Include NFS in Partition Reports option

Select whether Enlighten**DSM** should include NFS-mounted partitions when displaying a list of disk partitions. The default setting is No.

Although including NFS-mounted 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 snapshots for and searching on NFS partitions are relatively slow processes. Within an Enlighten**DSM** 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.

2 Getting Started

Masked Passwords option

Select whether Enlighten**DSM** displays asterisks in the Password fields when a password is entered. If you select Yes, the password will be displayed by asterisks when you enter it. If you select No (the default setting), the newly assigned password will be displayed using the characters you type.

☞ No confirmation is required when you assign a password. Be careful to avoid typing errors.

Maintain Catalog of Backed Up Files option

The Maintaining Backed Up Files option specifies whether Enlighten**DSM** will keep a catalog of all backed-up. The default setting is No.

If Yes is chosen, Enlighten**DSM** will create a catalog of all backup files for each partition. This catalog stores the backup date and time, the filenames, file sizes, the file owners, and the date of last modification. The default name of the catalog is NONAME.

The catalogs can then be quickly searched to find and restore backed up files.

Maintain window position and size

Enlighten**DSM** windows will pop-up in a default location with a default size.

- Choose the Yes setting to save any changes to window size and location.
- Choose the No setting to retain the default window position and size (the window will revert to the default size and location when the window is closed and re-opened).

Allow Rotary DNS During Host Entry Config

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

Creating Network Pools

Network pools are used to simultaneously manage groups of individual hosts. These pools can be based on functional or organizational needs without regard to network topology or the physical location of workstations. Pools can contain other pools, as well as individual hosts, to create a hierarchy of pools. Individual hosts can also belong to multiple pools.

In addition to making management and configuration of individual hosts easier, pools simplify reporting. For example, if you want to look at disk usage of hosts, you can generate a report showing disk usage for each host within the pool. From that information you can modify individual hosts, or modify all of the hosts in the pool simultaneously.

To set up a pool:

1) Choose Pool Configuration from the Configure menu. The Network Pool Configuration window appears (Figure 2-2).
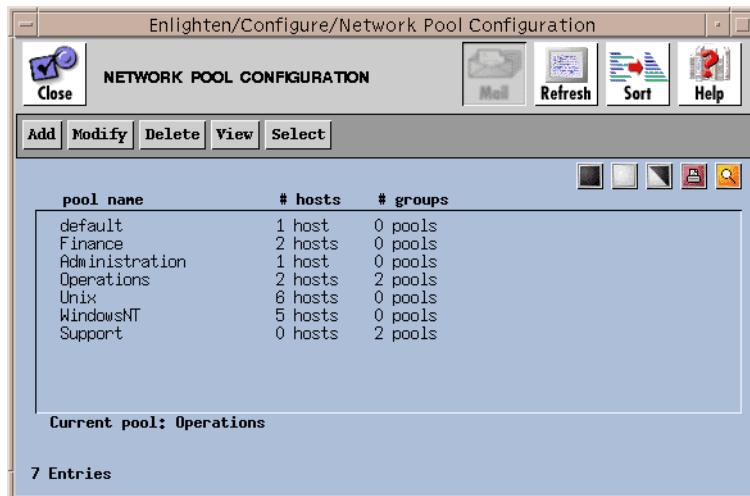
Figure 2-2 Network Pool Configuration window

2) Click the Add button. The Network Pool Add window appears
([Figure 2-3](navigation)).
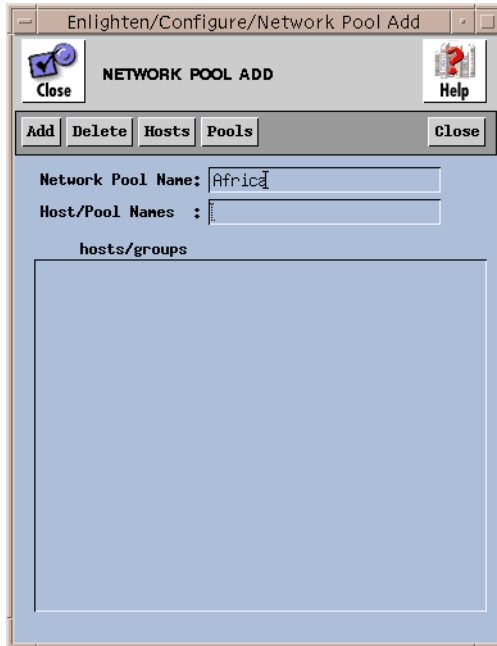
Figure 2-3 Network Pool Add window

3) Type the name of the new pool in the Network Pool Name field.
Choose a name that describes the function of the pool, such as
Marketing or Engineering.

4) Enter the hostnames to be included in the pool:

- Click the Hosts button to display a list of all currently configured
  hosts (generated from the /etc/hosts entries). Highlight the
  desired selections and click Apply to add the hosts to the new
  pool.

- You can also type the hostname(s) in the Hosts/Pool Names field.
  For multiple entries, leave a space between each hostname.

5) Click the Add button to save the new pool.

☞ To view a list of the hosts in a network pool, highlight the pool name and click the View button. If more than one pool is chosen, click the Next button to see the hosts in the next selected pool.

Changing User Authorizations

Enlighten**DSM**'s User Authorization features set and modify the privileges of others who use Enlighten**DSM**. These authorization features are used to create sub-administrators who can easily change and manage their own local environment. The root user always retains the ability to modify the privileges of all other users on the network.

☞ Enlighten**DSM**'s default user configuration allows any user to use the program. To prevent any user from using Enlighten**DSM**, delete the default User Authorization template.

2 Getting Started

Modifying a User's Authorization

To modify the default privileges for a user:

1) Choose User Authorization from the Configure menu. The User Authorization Configuration window appears with a list of all users, their viewing and modifying privileges, and their default pool.
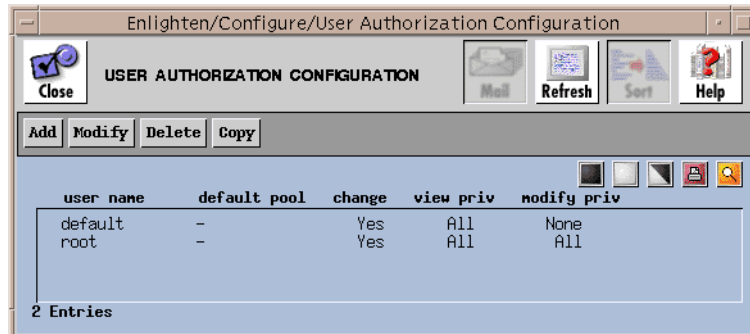
Figure 2-4 User Authorization Configuration window

2) Highlight one or more users whose privileges you want to change.

3) Click the Modify button. The Modify User Authorization window appears.

Figure 2-5 Modify User Authorization window

The user name, default pool, and a list of system tasks are displayed.

4) To change the default pool for the selected user, type the pool name in the Default Pool field (or click the right arrow button to choose from a pick-list of pools). The default pool specifies which pool of hosts the user belongs to when he or she starts up Enlighten**DSM**.

5) To allow the user to change their default pool, click the Change Pool option to Yes.

6) To allow the user to view and/or modify any of the network tasks listed, keep the default settings as Yes. To revoke the user's viewing and/or modifying privileges for any of the tasks listed, click the appropriate button and drag the setting to No.

    For a complete description of each task listed in the Modify User Authorization window, refer to Chapter 2, "Configure," in the Enlighten**DSM** *Reference Manual.*

7) After setting the options:

    •   Click the Modify button to save the changes, or

    •   Click the Clear Fields button to restore the default settings.

8) If more than one user was selected from the User Authorization window in step 2, click the Next button to modify the settings of the next selected user. Be sure to click Modify before moving on to the next user or the new settings will not be saved.

After modifying the selected user accounts, click the Close button to exit the Modify User Authorization window.

Adding and Deleting User Authorizations

Adding a new user is similar to modifying the privileges of an existing user. To add a new user:

1) Choose User Authorization from the Configure menu. The User Authorization Configuration window shown in Figure 2-4 appears.

2) Click the Add button. The Add User Authorization window appears.
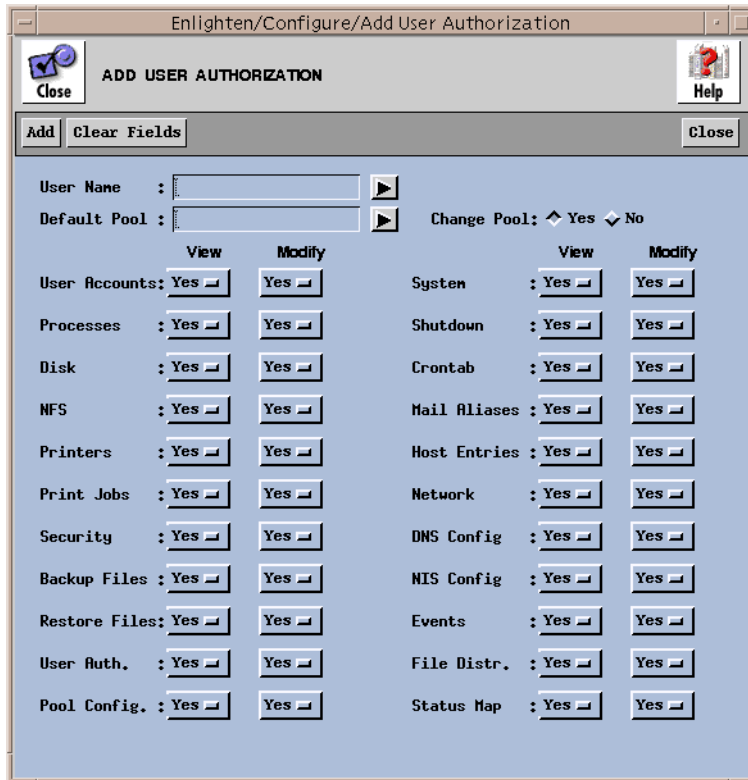
Figure 2-6 Add User Authorization window

3) Type the user name in the User Name field (or click the right arrow button to choose from a pick list of available user names).

4) Enter the Default Pool and Change Pool settings, then modify privileges for any of the tasks listed in the window. Refer to the previous section, "Modifying a User's Authorization" on page 10 and Chapter 2, "Configure," of the Enlighten**DSM** *Reference Manual* for more detailed information.

5) After completing the settings, click Add to save the changes, or click the Clear Fields button to restore the default settings.

To delete users from the user authorization list and prevent them from using Enlighten**DSM**:

1) Choose User Authorization from the Configure menu. The User Authorization Configuration window appears (see Figure 2-4).

2) Highlight the user(s) to be removed and click the Delete button. A message asks you to confirm the deletion.

3) Click Yes or All to delete the user(s).

3

Setting Up a Basic Network Configuration

Enlighten**DSM** helps you easily perform some of the more common and basic tasks of setting up and managing a network. This chapter describes how to create user accounts, set up user groups, and configure network printers.

Setting Up User Accounts

Setting up and modifying user accounts is one of the most common network management tasks. This section describes how to use the User Configuration window to add, modify, and delete user accounts and passwords. This section also describes management features for locking and unlocking user accounts, and for monitoring user information and processes.

For information about managing a user's system access, see "Managing a User's System Access," on page 4-2.

Adding a New User Account

To add a single new user account:

1) Choose Configure from the User menu. The User Configuration window appears.
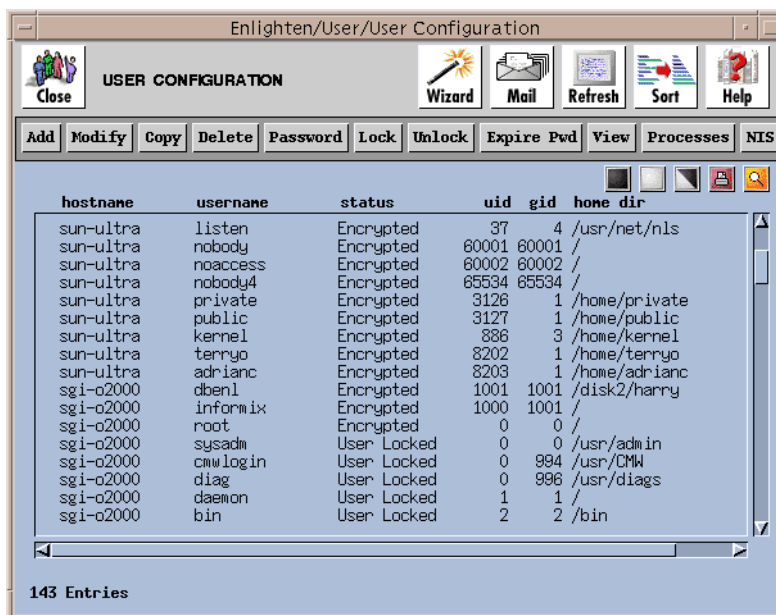
Figure 3-1 User Configuration window

The User Configuration window lists all of the current user accounts by hostname and user name. This window also lists user's password status, user id, group id, and home directory.

2) Click the Add button. The Create New User Account window appears.

Figure 3-2 Create New User Account window

The Create New User Account window allows you to add a single new user or multiple users.

3) Enter the appropriate information in each field. Refer to the section "New User Account and Modify User Fields," on page 3-6 for information about each of the fields.

4) After entering the information in the fields, click the Add button.

☞   To add multiple users in a single operation, you must first create a text file of the new users. Next, click the Multi-User Add button in the User Configuration window to bring up the Add Multiple User window. For information on using this feature, refer to Chapter 3, "User," in the Enlighten**DSM** *Reference Manual*.

Creating New User Accounts from Templates

New user accounts often contain the same configuration information. Creating a new user account from an existing template can reduce the amount of redundant typing you'll have to do.

To create a new user account from an existing record (template):

1) Click the Template button in the Create New User Account window. A list of templates appears in a pick list.

2) Select the template that best matches the parameters of the new user. The information from that template will automatically appear in the Create New User Account window.

3) Change any of the information from the template.

4) Click Add to create a new user account.

5) If running NIS or NIS+, click the NIS button to save the changes.

LADP: Lightweight Directory Access Protocol

The LADP 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 LADP server.

For more information about using the LADP feature, refer to Chapter 3, "User," in the Enlighten**DSM** *Reference Manual*.

Modifying a User Account

To modify a user account:

1) Highlight the user(s) whose account(s) to be modified.

2) Click the Modify button in the User Configuration window. A Modify User window similar to the Create New User Account window appears.

Figure 3-3 Modify User window

3) Change the information in the fields as desired. For information about each field refer to "New User Account and Modify User Fields," on page 3-6.

4) Click the Modify button to save the changes.

5) If more than one user account was selected from the User Configuration list, click the Next button to modify the next account.

Be sure to click Modify before moving on to the next user or the new settings will not be saved.

6) If running NIS or NIS+, click the NIS button to save the changes.

New User Account and Modify User Fields

The fields in the Create New User Account and Modify User Account windows are similar. The Login Name field is read-only in the Modify window. The Mail Alias and Mail Lists fields exist only in the Create window.

Hostnames field

To limit an account to specific hostnames within a pool, enter those hostnames in this field. For multiple hostnames, leave a blank between each entry. Click the right arrow button to choose from a pick list of available hosts from the current pool. If this field is left blank, the user will be added to all hosts in the pool.

Login Name field

Each user needs a unique Login Name to access the system. The first character of this Login Name should be set in lowercase (for example, cHE034). Macros can also be used in this field; see Appendix K, "Macros," in the Enlighten**DSM** *Reference Manual* for more details.

In the Modify User window, this field is read-only.

Real Name field

Enter the user's real name to help associate the sometimes cryptic Login Name with the actual user. On larger systems, this name matching is invaluable. This field can also identify a company or organization.

Office field

Enter the location where the user will normally access the system. This field can contain a department name, office number, or an assigned location code. Punctuation characters are not allowed in this field.

Telephone field

Enter the telephone number, telephone extension, or some other contact information, in this field. The text format is user definable.

Click the Add button to have Enlighten**DSM** automatically fill in the remainder of the form.

The rest of this section contains information for manually filling in these fields:

User ID field

Enter the UserID number for the user account. Although more than one user can technically have the same number, the UserID number *should* be unique for easier system administration. The number is also used by the operating system to identify each user.

If the New User Template is set to automatically generate a unique UserID, that number will be entered in this field.

☞   If this field is left blank, Enlighten**DSM** will automatically enter a unique number when the user account is saved.

Home Dir. field

The Home Dir. field specifies the parent directory in which each user's start-up directory is created. Enlighten**DSM** uses /home as the default HOME directory unless you change this field (most systems will use /usr or /home as the default parent directory). Creating a new user "Charlie", for example, will also create a user directory /home/charlie.

If you specify a HOME directory that does not exist, Enlighten**DSM** creates and initializes one. Enlighten**DSM** initializes the HOME directory by executing the HOME directory initialization script $ENLIGHTEN/scripts/user.rc. This copies the initialization files (for example; .profile and .cshrc) into the HOME directory. This script can be modified to perform other initialization procedures.

Shell field

The Shell is a program that acts as a translator between the user and the 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. You can also create a custom shell. 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

Enlighten**DSM** uses the Bourne Shell (/bin/sh) as the default shell in this field. Each shell offers different functionality and all users on the system do not have to use the same shell. The number of standard shells available for this field depends on your system. Click the right arrow button to choose from a pick list of available shells.

Group Name field

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

Enter the primary Group Name for this user (or click the right arrow button to choose from a pick list of defined User Groups). If a non-existent Group Name is entered, Enlighten**DSM** will ask for a correct Group Name when trying to create the user.

Add Groups field

To give users privileges beyond those set in the primary User Group, you can assign users to additional Groups. Enter the names of the additional Groups in the Add Groups field (or click the right arrow button to choose from a pick list of defined User Groups).

Password field

Because the password is part of system security, each user must have a password. A new password can be the same characters as the user's Login Name, allowing the user to change it later. Macros can also be used in this field to automatically generate new passwords. See Appendix K, "Macros," in the Enlighten**DSM** *Reference Manual* 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 field

For systems that support password aging, this optional field contains the time when the password will expire. When a password expires, the system forces the user to change their password at the next system login.

If the Expires field is left blank, the default parameter is used. If you don't set a default when adding a user, password aging will be turned off for that user. For more details, refer to Appendix D, "Password Aging," in the Enlighten**DSM** *Reference Manual*.

To specify a time format, see Appendix C, "Time Formats," in the Enlighten**DSM** *Reference Manual*.

Span field

The Span field defines the minimum period of time that must elapse before a user can change their password (you can, however, change the user's password at any time by using the Password button in the User Configuration window). For information on time formats, see Appendix C, "Time Formats," in the Enlighten**DSM** *Reference Manual*.

Mail Alias field (Create New User Account window only)

Enter the default mail alias for an account. This alias designates a short name as the substitute for the full user address or group of addresses. You may also use macros in this field. See Appendix K, "Macros," in the Enlighten**DSM** *Reference Manual* for more details.

3 Setting Up a Basic Network Configuration

Mail Lists field (Create New User Account window only)

Enter the mailing lists this user will belong to. A mail list allows a single message to be sent to a group of people. Click the right arrow button to choose from a pick list of defined mail lists.

Deleting a User

To delete a user from the system:

1) In the User Configuration window (Figure 3-1), highlight the username(s) to be removed from the system and click Delete.

2) Confirm the deletion of the user(s).

3) By default, EnlightenDSM will also delete the user's HOME directory (including all files and sub-directories). EnlightenDSM will prompt you to confirm the deletion of this directory. Click Yes to delete the directory. Click No to leave the directory intact while still deleting the user account.

4) If running NIS or NIS+, click the NIS button to save the changes.

EnlightenDSM will *not* delete the user's HOME directory if:

- The user's HOME directory is considered to be a system (required) directory. The system directories are:

```
/              /usr            /etc          /bin
/lib           /home           /dev          /sbin
/usr/bin       /usr/adm        /usr/lib      /usr/sbin
/usr/man       /usr/lib/uucp
```

- The user's HOME directory 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.

- The user's HOME directory is not owned by the user. In this case, you are asked to confirm before deleting it.

- It is the HOME directory of the root user account.

Changing a User Password

You do not have to know a user's password in order to change it.

To change a current user account password:

1) From the User Configuration window (Figure 3-1), highlight the user(s) whose password(s) you want to change.

2) Click the Password button. Enter the new password in the pop-up window (Figure 3-4).
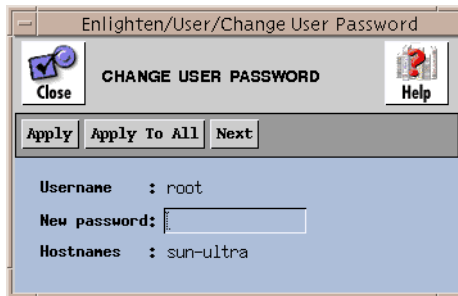
Figure 3-4 Change User Password window

3) Click the Apply button to save the password.

4) If multiple usernames were selected, click the Next button to enter the next user password. Be sure to click Apply before moving on to the next user or the password will not be saved.

5) Click Apply To All to set the same password for all the selected users.

6) Click the Close button to stop changing passwords and close the window. Any previous changes made by using the Apply button remain; they are not undone by clicking this button.

7) If running NIS or NIS+, click the NIS button to save the changes.

Expiring a Password Immediately

To immediately cancel a user's password, click the Expire Pwd button at the top of the User Configuration window. The user(s) will be forced to enter a new password at their next login.

Setting Up User Groups

User Groups allow system managers to assign access privileges to groups of individual users. Use Enlighten**DSM** to create, modify or delete User Groups.

To create a new User Group:

1) Choose Groups from the User menu. The Group Configuration window appears.
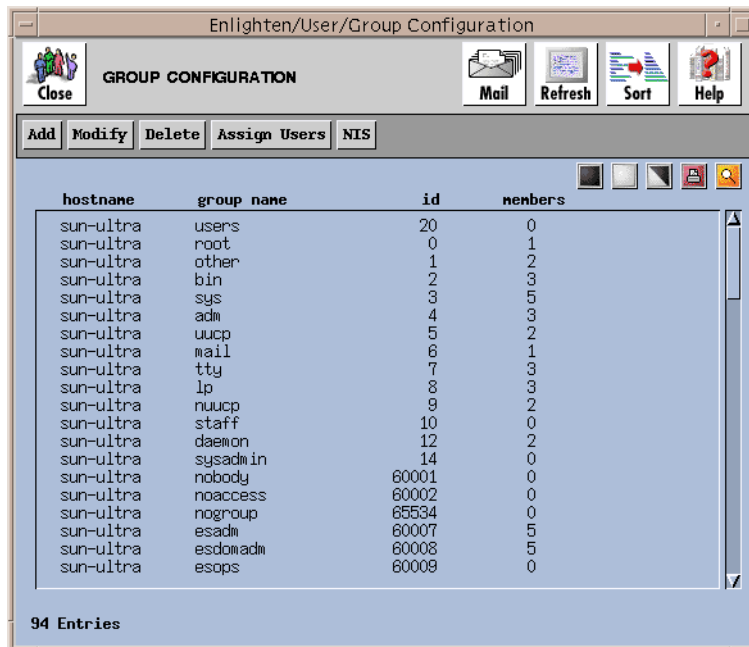
Figure 3-5 Group Configuration window

2) Click the Add button. The New User Group window appears.
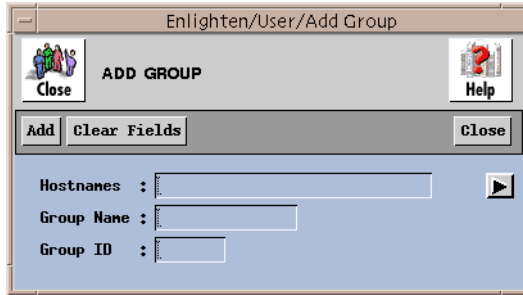
Figure 3-6 New User Group window

3) To limit the group to certain hostnames, type the hostnames in the Hostnames field. For multiple entries, leave a blank between each hostname.

4) Type the name of the group in the Group Name field. Choose a name that reflects the group's function, such as "marketing" or "bizteam".

5) Enter a unique number (no decimal points) in the Group ID field. The Group ID (GID) is used by the operating system to identify the group. SCO operating system limits the Group ID to eight digits.

6) Click Add to save the new User Group. Click Close to return to the Group Configuration window.

7) If running NIS or NIS+, click the NIS button to save the changes.

Modifying Group Parameters

To modify the parameters of a User Group:

1) Click the Modify button in the Group Configuration window (Figure 3-5). The Modify Group Configuration window appears.
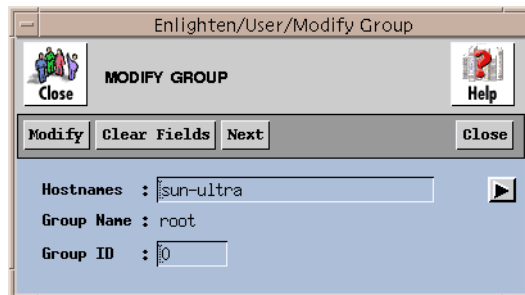
Figure 3-7 Modify Group Configuration window

2) Change the parameters as desired.

3) Click the Modify button to save the changes.

4) If multiple groups were selected, click the Next button to modify the next User Group. Be sure to click Modify before moving on to the next group or the new settings will not be saved.

5) If running NIS or NIS+, click the NIS button to save the changes.

Deleting User Groups

To delete an entire User Group:

1) Select the group(s) to be deleted from the Group Configuration window (Figure 3-5).

2) Click the Delete button. In addition to removing the highlighted User Groups, links between the group and its members are deleted.

3) If running NIS or NIS+, click the NIS button to save the changes.

Assigning Group Users

Click the Assign Users button (see Figure 3-5) to modify the users assigned to a group. The User Group List window (Figure 3-8) will appear.
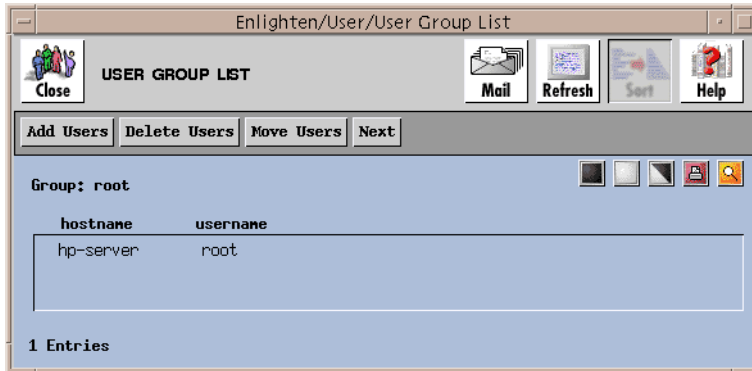
Figure 3-8 User Group List window

To add a user to a group:

1) Click the Add Users button. A menu appears listing all users on the system.

2) Highlight the users to be added to the group.

3) Click the Apply button.

To delete users from a group:

1) Highlight the users to be deleted from the group.

2) Click the Delete button.

To move a user from one group to another:

1) Highlight the users to be moved to another group. Click the Move Users button. A pop-up window appears.

2) Enter the User Group that the user is to be moved to (or click the right arrow button to choose from a pick list of available User Groups).

3) Click the Apply button to save the changes.

4) If running NIS or NIS+, click the NIS button to save the changes.

Setting Up Print Services

Use Enlighten**DSM** to define printer servers and printer clients.

Printer servers are used to efficiently manage print jobs. Print servers receive all print requests, queue the jobs, and distribute each print job to various available printer clients. The printer clients then print the jobs sent to them by the printer server.

Enlighten**DSM** has two configuration options: the Configuration Local option reports only local printer destinations. The Configuration All option displays both network and local printers. A local printer is a printer connected directly to a host. A network printer is connected to your network and so has an IP address.

This section describes how to set up network printers, enable and disable print jobs, and manage print jobs and queues. This section assumes that the printers are already physically connected to either a host or network.

Refer to Chapter 4, "Printer," in the Enlighten**DSM** *Reference Manual* for information about

- Defining custom printer models or drivers

- Displaying and managing print jobs

Setting Up a Local Print Service

To set up a local print service for a printer attached directly to a host:

1) Choose Configure (local) from the Printer menu. The Configured Printers window appears.
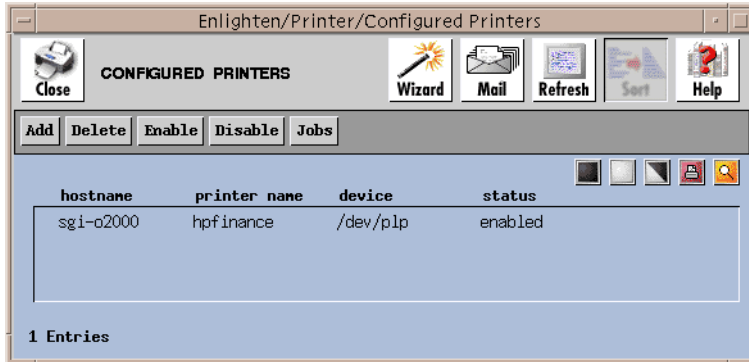
Figure 3-9 Configured Printers window

2) Click the Add button. The Add a New Print Service window appears.
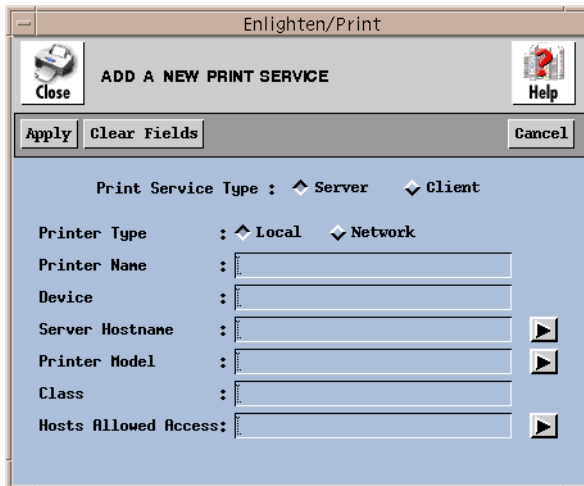
Figure 3-10 Add a New Print Service window

3) For Print Service Type, click the Server option.

4) For Printer Type, click the Local option.

5) In the Printer Name field, enter the printer name users will refer to. (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*.

6) In the Device field, enter the type of device by which the printer will be attached: for example, /dev/lp (parallel) or /dev/ttyh7 (serial). This device must already exist.

7) In the Server Hostname field, enter the hostname (from the current pool) that will act as the print server for this printer.

8) In the Printer Model field, specify the printer driver to be used. A Server Hostnames must be entered before selecting a driver.

9) The Class field is optional. By entering a printer in this field, the print job is sent to the first available printer in that printer class. For systems using the System V type spooling system, defining the class will more efficiently distribute print jobs to multiple printers of the same type.

   If a printer class in not entered, 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.

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

10) In the Hosts Allowed Access field, enter the hostnames that will have access to this printer.

11) Click the Apply button to save the settings.

Setting Up a Print Server for a Network Printer

Adding a network print server is similar to adding a local print server.

To add a network print server, follow the steps described in the previous section, "Setting Up a Local Print Service," on page 3-17. You must also:

- Click the Printer Type option to Network.

- Specify the printer's IP address in the IP Address field (this field will appear when you select the Network Printer Type option).

Setting Up a Client Print Service

To set up a client print service:

1) Choose Configure All from the Printer menu. The Configured Printers window appears.
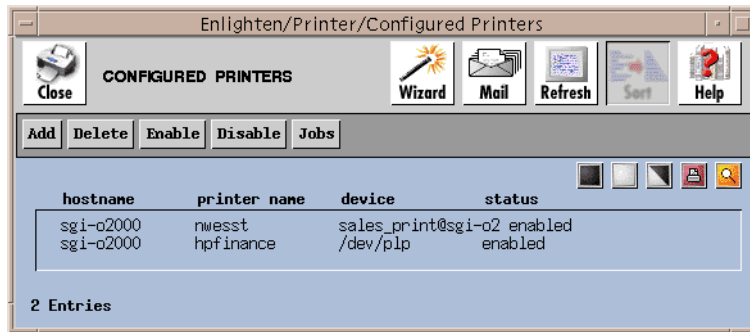
Figure 3-11 Configured Printers window

2) Click the Add button. The Add a New Print Service window appears.

3) Click the Client button in the Print Service Type option.
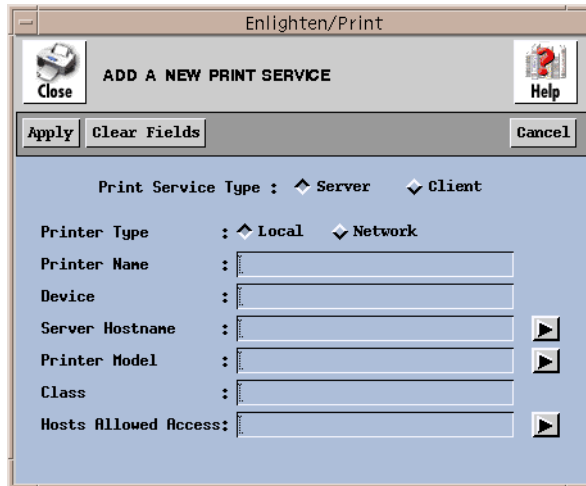
Figure 3-12 Add a New Print Service window

4) To configure the new printer on a subset of the currently managed hosts, enter the hostname(s) in the Hostnames field. For multiple entries, leave a blank space between each hostname. You can also click the right arrow button to choose from a pick list of hosts in the current pool.

5) In the Local Queue Name field, enter the full pathname for the local printer queue (for example, /var/spool/lp).

6) In the Remote Host field, enter the hostname that will act as the print server for this particular client (or click the right arrow button to choose from a pick list of hosts in the current pool).

7) In the Remote Printer field, enter the remote printer that will service the print jobs (or click the right arrow button to choose from a pick list of all printers defined on hosts in the current pool).

8) Click the Apply button to save the printer configuration.

Deleting a Printer

To delete a printer from the system:

1) Choose either Configure (all) or Configure (local) from the Printer menu. The Configured Printers window appears (Figure 3-11).

2) Highlight the printers to be deleted from the system. Click the Delete button. A Enlighten**DSM** window will ask for confirmation of the deletion.

☞   Enlighten**DSM** cannot delete a printer configuration if there are pending print requests. To move or cancel all pending print requests on the printer to be deleted, choose Queue from the Printer and click on the appropriate option.

Enabling and Disabling a Printer

The Enable command makes the printer active and available for printing. Once enabled, the printer will print any jobs waiting in the queue.

To enable a printer that has been disabled, click the Enable button in the Configured Printers window (Figure 3-11).

☞   When a printer is added to the system, it is automatically enabled.

To disable printers and stop them from receiving spooler print requests, highlight the printers and click the Disable button. Currently printing jobs will stay in the print queue and finish printing once the printer is enabled. This command does not stop users from sending print jobs to the disabled printer; those print requests are added to the queue and printed once the printer is enabled.

4

Managing the Network Systems

EnlightenDSM provides powerful tools for managing systems on your network.

This chapter contains basic information about the following EnlightenDSM network management features:

- Locking and unlocking user accounts

- Managing disk space by building disk snapshots, viewing usage, searching files, and displaying currently running processes

- Managing Network Filesystems (NFS) for clients and servers

- Managing cron jobs

- Configuring archive devices and performing immediate and scheduled backups.

• Configuring and managing Network Information Servers (NIS) on a network.

For information about other management tasks not covered in this chapter, refer to the Enlighten**DSM** *Reference Manual*.

Managing a User's System Access

This section describes how to lock and unlock a user's access to the system without deleting their account. Enlighten**DSM** also provides tools to check files, passwords, and other network functions. Refer to the Enlighten**DSM** *User Guide* and Chapter 5, "Monitoring Your Network Systems," for information about those security features.

To lock a specific user access to the system, follow the instructions below. Locking an account will deny that user access to the system without deleting their account:

1) Choose Configure from the User menu. The User Configuration window appears.
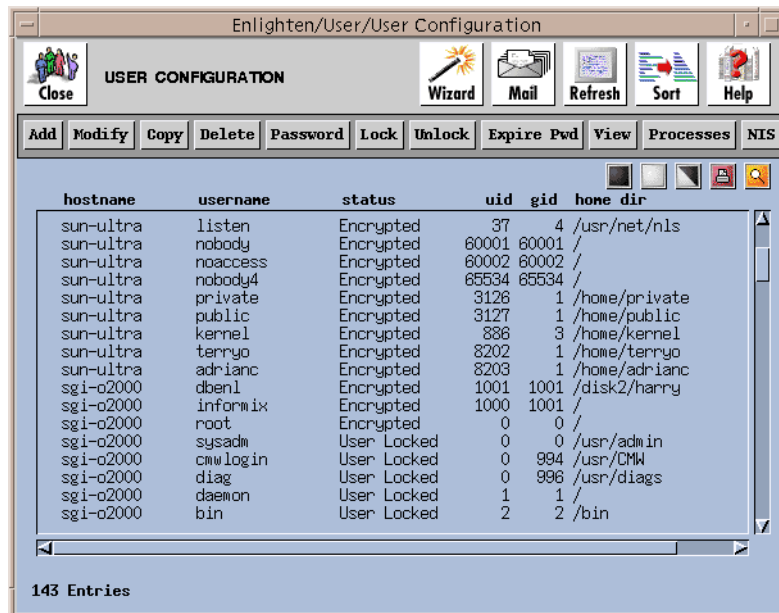
Figure 4-1 User Configuration window

2) Select the user whose access you want to deny.

3) Click the Lock button. A pop-up window will prompt you to confirm this action.

☞ Locking a user account deletes the old password. You must create a new password when unlocking the account.

To unlock a user account:

1) From the User Configuration window, highlight the user account to be unlocked.

2) Click the Unlock button. This will restore that user's access privileges.

3) Assign a new password for the user when prompted.

Because locking and unlocking user accounts is considered part of system security, a new password is required when Enlighten**DSM** unlocks a User Account

Managing Disk Partitions

With Enlighten**DSM**, network managers can easily search for files and view disk usage summaries for users, groups, and disk filesystems. To decrease the time required to generate these summaries, Enlighten**DSM** creates a database, or *snapshot*, of the files in the requested filesystem. This *snapshot* database is used to perform the required tasks.

This section briefly describes how to view usage by filesystem, username, and groupname. It also describes how to build snapshots, perform file searches, and display processes. For more detailed information about these options, refer to Chapter 7, "Disk," in the Enlighten**DSM** *Reference Manual*.

4 Managing the Network Systems

Viewing Usage by Filesystem

To view disk usage by filesystem:

1) Choose Usage by Filesystem from the Disk menu. The Usage by Filesystem window appears.

Figure 4-2 Usage by Filesystem window

2) Highlight the filesystem(s) to be summarized, then choose one of the following options:

- Click the User Sum. button for a summary of disk usage by username.

- Click the Group Sum. button for a summary of disk usage by user group.

- Click the File Search button to search files based on a search criterion.

- Click the Processes button to view the processes using a filesystem.

- Click the Graph button to graph disk usage.

Viewing Disk Usage by Username

To view a quick summary of disk usage by user for currently mounted filesystems:

1) Choose Usage by Username from the Disk menu. The Select Disk Partitions window appears.

Figure 4-3 Select Disk Partitions window

2) Highlight the disk partition to be summarized.

3) Click the View User Summary button. The Disk Usage Information by Users window appears.

**4 Managing the
Network Systems**

Figure 4-4 Disk Usage Information by Users window

4) Highlight the Usernames to be summarized, then choose one of the
    following options:

    •   Click the View Files button to view a list of files owned by the
        user(s). A File Systems Detail window will appear with
        additional options for gathering and viewing file information,
        deleting or backing up the file or changing the access privileges
        to the file.

    •   Click the Compare Snapshots button to compare the current
        snapshot to a previously saved one.

    •   Click the Save Snapshots button to save the current snapshot. See
        Building Snapshot below for more information.

    •   Click the Graph button to graph the disk usage summary of the
        user(s).

Viewing Disk Usage by Groupname

To view a quick summary of disk usage by Usergroup for currently mounted
filesystems:

1) Choose Usage by Groupname from the Disk menu. The Select Disk
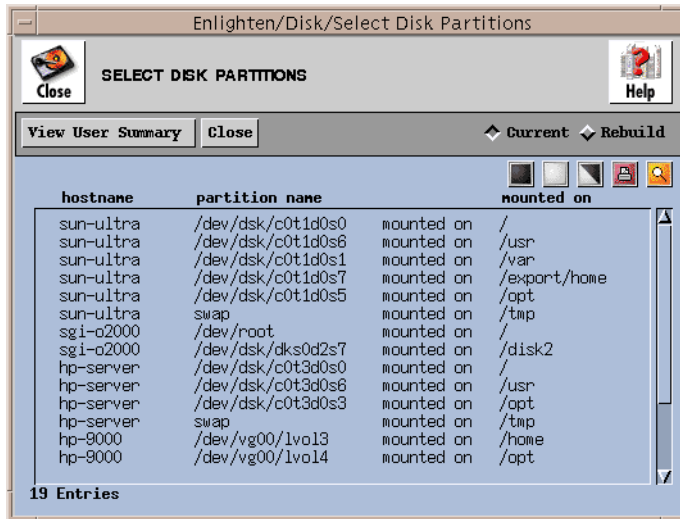   Partitions window appears.
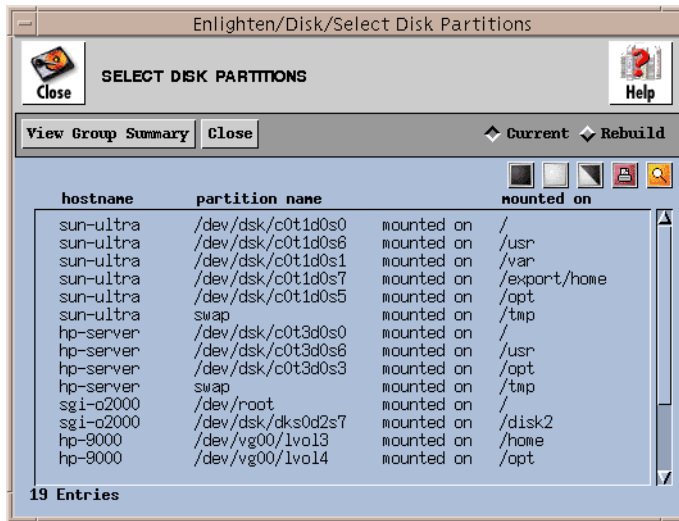
Figure 4-5 Select Disk Partitions window

2) Highlight the disk partitions you want to view, then click the View Group Summary button. The Disk Usage Information by Groups window appears.
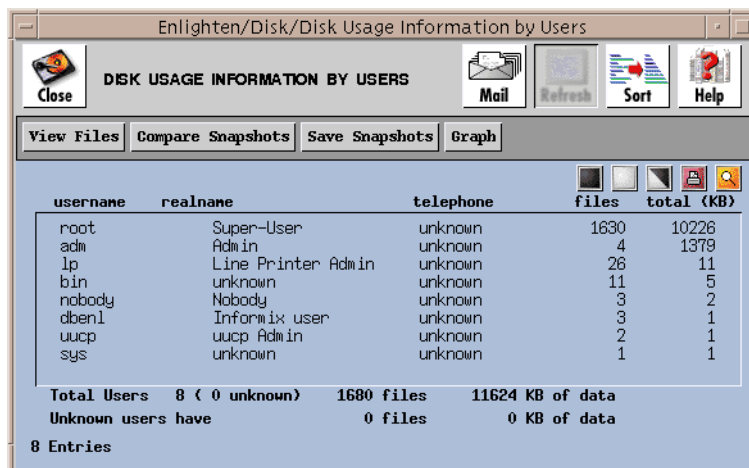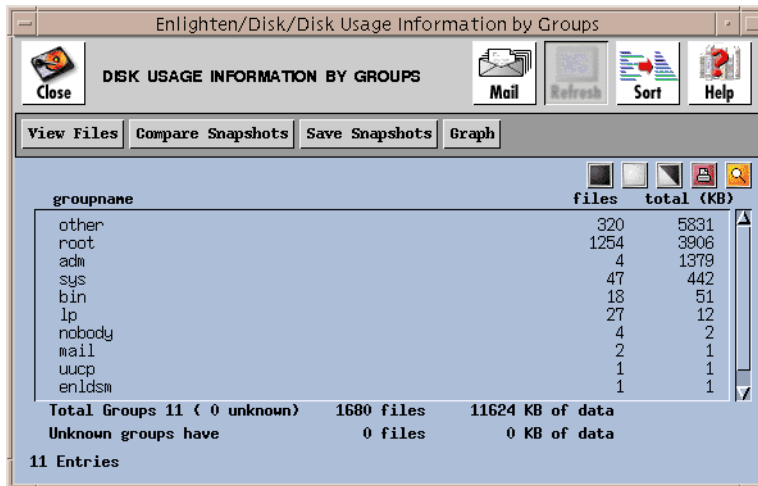

Figure 4-6 Disk Usage Information by Groups window

3) Highlight the Groupname entries to be summarized, then choose one of the following options:

4 Managing the
Network Systems

- Click the View Files button to view a list of files owned by the group(s). A File Systems Detail window will appear with additional options for gathering and viewing file information, deleting or backing up the file or changing the access privileges to the file.

- Click the Compare Snapshots button to compare the current snapshot to a previously saved one.

- Click the Save Snapshots button to save the current snapshot. See the next section, "Building Snapshots," for more information.

- Click the Graph button to graph the disk usage of the group(s).

Building Snapshots

Disk snapshots provide a detailed description of a disk partition at a given moment. These snapshot are used to monitor disk usage and determine if changes have been made to a partition over time. You can save and store these snapshots, or update them with the rebuild function.

Creating New Master and Interim Snapshots

Master snapshots are not automatically created during installation. You must first create and save a master snapshot that can be compared with later snapshots. Follow the instructions below to create a master snapshot.

A master snapshot is usually created after installing the Enlighten**DSM** software but before giving users access to partitions. This way, the master snapshot represents a secure system. A master snapshot can also be created after a system cleanup to track changes made over time.

Interim snapshots are used to compare changes to a partition over time. It's best to assign meaningful names to interim snapshots.

To create and save a new master or interim snapshot:

1) Choose Save Current Snapshots from the Disk menu. The Select Disk Partitions window appears.
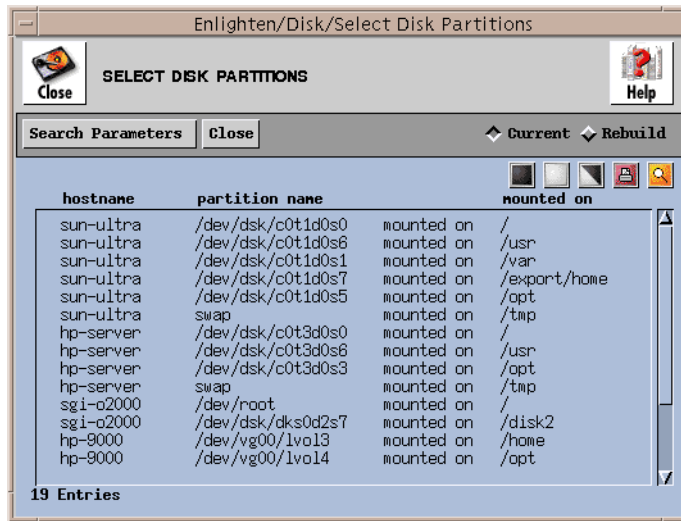
Figure 4-7 Select Disk Partitions window

2) Highlight the disk partitions for which you want to make a snapshot.

3) Click the Save Snapshot button. The Save Selected Disk Snapshots window appears. The names of the selected partitions are imported into the Snapshot from Partition column. You can add other partition names by typing them in.

Figure 4-8 Save Selected Disk Snapshots window

4) For each partition, do the following:

- To save the snapshot as the Master, leave the Yes button enabled in the Save as Master option.

- To save the snapshot as an interim snapshot, click the No button and enter a unique name in the Save Snapshot As field.

5) Click the Save Snapshot button.

Comparing Snapshots

To compare a new snapshot against a master or interim snapshot:

1) Choose one of the Disk Usage commands:

   • Disk—>Usage by Filesystem

   • Disk—>Usage by Username

   • Disk—>Usage by Groupname

2) Highlight the disk partitions (filesystems) you would like to compare to an earlier snapshot.

Figure 4-9 Currently Mounted Filesystems

3) Click the Rebuild radio button in the upper right part of the window, and then click one of the function buttons; for example, the User Sum. button.

   A new snapshot of the current filesystem is created that can be compared to the master or interim snapshots.

4) Click the Compare Snapshot button in the Disk Usage Information window. The Select Saved Snapshots window appears.
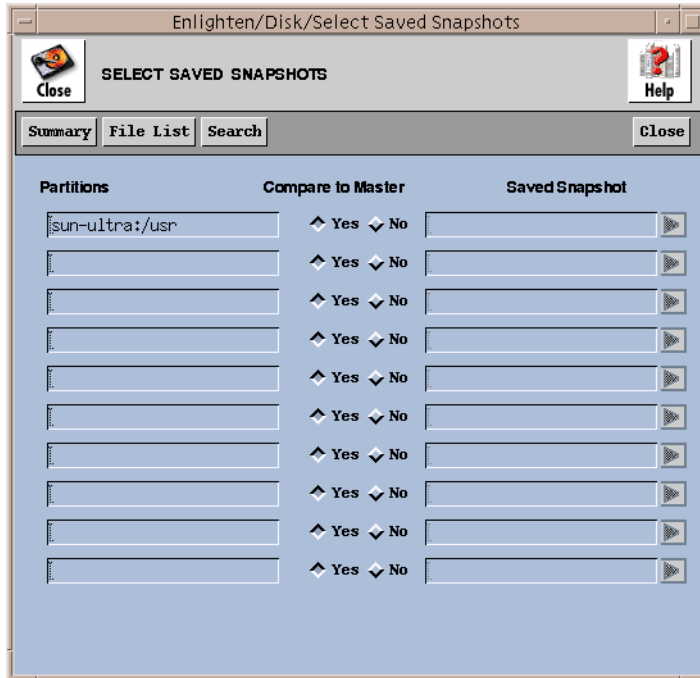
Figure 4-10 Select Saved Snapshots window

5) To compare against a master snapshot, click the Yes button under the compare to Master column (the default). To compare against an interim named snapshot click the No button and enter the name of the previously saved snapshot in the Saved Snapshot field. Click on the right arrow button to choose from a pick list of previously saved snapshots.

☞ 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 do not change) across hosts.

6) Choose one of the following snapshot comparison options:

- Click the Summary button to get a summary of the changes.

- Click the File List button to view a detailed list of the changes.

- Click the Search button to search for specific types of changes.

Updating

Snapshots of selected filesystems can be updated using the Disk Usage by Filesystem and Select Disk Partitions windows. The updated snapshots replace the old disk partition information with current information.

To update the snapshot of a filesystem:

1) Highlight the filesystem in the list and click the Rebuild button.

2) Click one of the buttons: User Sum. or Group Sum.

   Enlighten**DSM** displays a window indicating when the last request was taken.

3) Click the Save Snapshot button to save the snapshot as an updated master or interim snapshot.

☞ Because disk summaries and searches are performed on this snapshot, you should periodically rebuild snapshots to reflect a more current state of the disk. This can be done automatically with the cli and cron commands.

Managing Network File Systems (NFS)

There are two sides to Network File System (NFS) management: the client side where filesystems are mounted or unmounted on a given host, and the server side where directories are made available for export by a host. You can use Enlighten**DSM** to handle both sides. The actual options are:

- Mounted Directories
- Exported Directories

This section explains how to mount and export directories, or *partitions*. For more information about NFS options not explained here, refer to Chapter 5, "Network," in the Enlighten**DSM** *Reference Manual*.

Mounting Directories on Hosts

Partitions can be mounted on the host(s) immediately, when the hosts are booted, or both.

To add an NFS partition on one or more hosts:

1) From the Network menu, choose NFS, then Mounted Directories. The Network Filesystem Configuration window appears. Mounted filesystems on all hosts in the current pool are displayed.
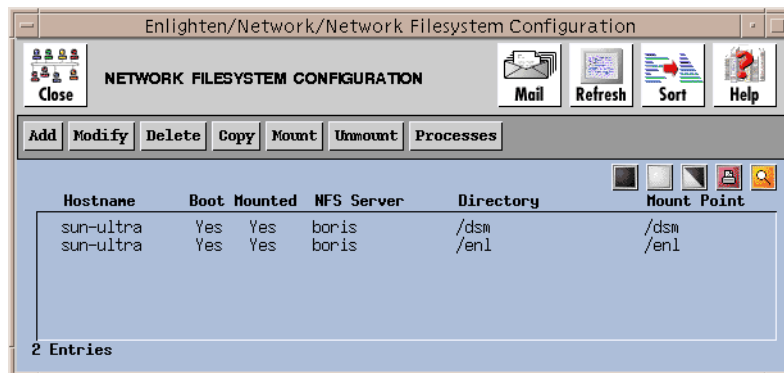
Figure 4-11 Network Filesystem Configuration window

Each line in the list box displays the hostname, whether the filesystem is currently mounted and/or automatically mounted at boot time, the NFS Server name, the remote partition (directory) name, and the local mount point.

2) Click the Add button. The Add an NFS Partition window appears.

Figure 4-12 Add an NFS Partition window

3) Enter the server and directory information in the appropriate fields and select the mount options you want (partitions can be mounted on the hosts immediately, when the hosts are booted, or both). For information about all of the options in the window, refer to the next section, "Add and Modify NFS Partition Options," on page 4-17.

4) Click the Mount button to initiate the mount you defined.

To modify the mount parameters for an existing configuration:

1) From the Network menu, choose NFS, then Mounted Directories. The Mounted Filesystems on Host window appears, displaying the mounted filesystems on all hosts in the current pool (see Figure 4-11).

2) Highlight the NFS partition(s) to be modified.

3) Click the Modify button. The Modify an NFS Partition
window appears.
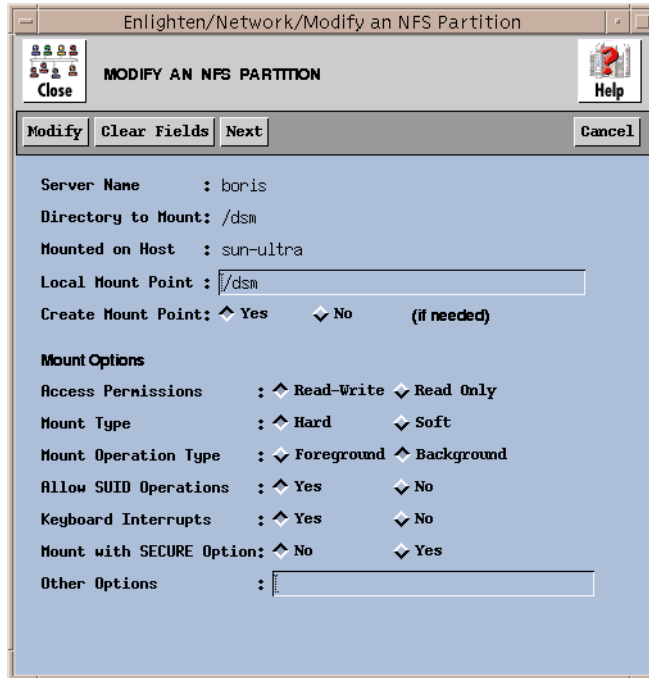
Figure 4-13 Modify an NFS Partition window

4) Change the options as needed. For information about all of the
options in the window, refer to the next section, "Add and Modify
NFS Partition Options," on page 4-17.

5) Click the Modify button to save the changes.

6) If more than one partition was selected from the list, click the Next
button to modify the next NFS partition. Be sure to click Modify
before moving on to the next partition or the new settings will not be
saved.

Add and Modify NFS Partition Options

The Add an NFS Partition and Modify an NFS Partition windows are similar. The Modify an NFS Partition window does not allow the modification of the Server Name, Directory to Mount, or Mounted on Host fields. Each field or option in these windows is described in the following sections:

Server Name field

Enter the NFS server from which the partition will be mounted (or click the right arrow button to choose from a pick list of all hosts within the current pool that may have mountable partitions).

Directory to Mount field

Enter the remote directory to mount. Leave a blank space between directory names for multiple entries.

Mount on Hosts field

Enter the hostname(s) where the directory will be mounted. Leave a blank space between hostnames for multiple entries. You can also click the right arrow button to choose from a pick list of all hosts within the current pool that may have mountable partitions.

Local Mount Point field

Enter the directory name where the partition will be mounted.

Create Mount Point option

Choose whether the local mount point directory should automatically be created during a mount if that directory does not already exist. The default setting is Yes.

4 Managing the Network Systems

Access Permissions option

Choose how the remote partition should be mounted: Read-Write (the default) or Read Only.

Mount Type option

Choose whether the mount should be a Hard mount (the default) or a 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 option

Choose whether the mount should be performed in Foreground (the default) or Background. Filesystems mounted in the Background will retry the mount in the background if the server's mount daemon does not respond.

Allow SUID operations option

Choose whether setuid execution is allowed for this mount. The default setting is Yes. The No setting causes the filesystems to be mounted but silently ignores the request to set the suid bit.

Keyboard Interrupts option

Choose whether keyboard interrupts are allowed to kill a hung process that is waiting for a response on a hard-mounted filesystem. The default setting is Yes.

Mount with Secure option

Choose whether to mount with RPC authentication. The default setting is No.

Other Options field

Use this field to specify other NFS options. This field is not pre-checked for syntax errors.

Mounting and Unmounting NFS Partitions

To manually mount or unmount NFS partitions, highlight the appropriate partition(s) in the Network Filesystem Configuration window (see Figure 4-11) and click the Mount or Unmount button.

Exporting Directories

This section describes how to export a directory, define new export criteria, and modify export parameters for an existing NFS partition.

To export a directory:

1) From the Network menu, choose NFS, and then Exported Directories. The Export A Directory window appears displaying all currently exportable directories.
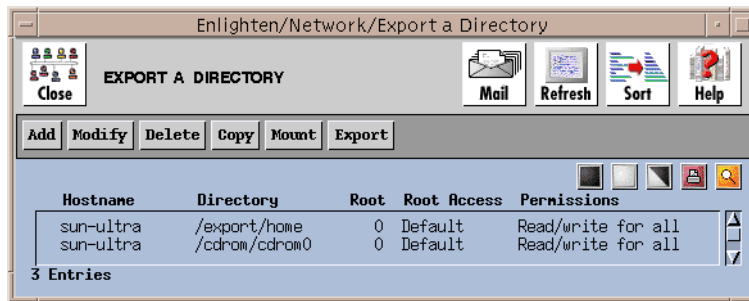
Figure 4-14 Export a Directory window

Each line in the list displays the hostname, exported partition, number of hosts with root access, root access, and permissions.

2) Select the directory to be exported.

3) Click the Export button.

To define new export criteria for a directory:

1) From the Network menu, choose NFS, and then Exported Directories. The Export a Directory window appears displaying all currently exportable directories.
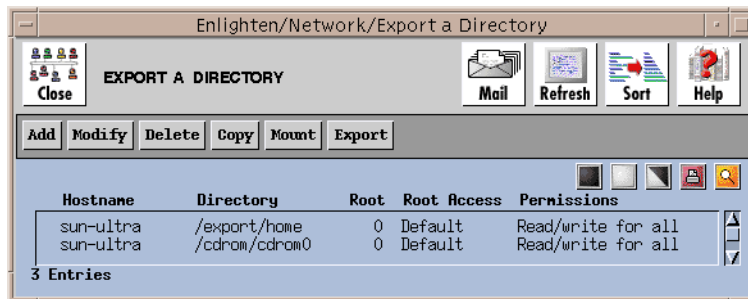
Figure 4-15 Export a Directory window

2) Click the Add button. The NFS: Export a Filesystem window appears.
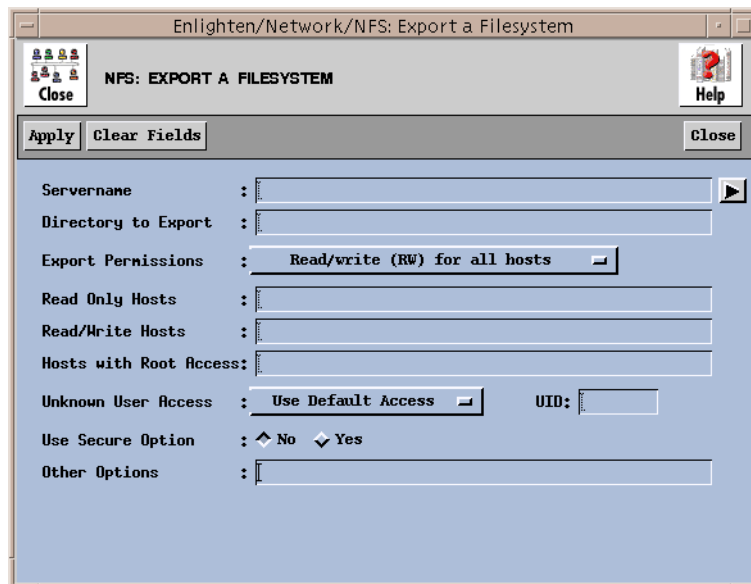
Figure 4-16 NFS: Export a Filesystem window

3) Enter the export parameters. Refer to the next section, "Add and Modify Export Fields," for information about each parameters.

4) Click Apply to export the directory.

To modify the export parameters for an existing NFS partition:

1) From the Network menu, choose NFS, and then Exported Directories. A window appears displaying the exportable directories on all hosts in the current pool.
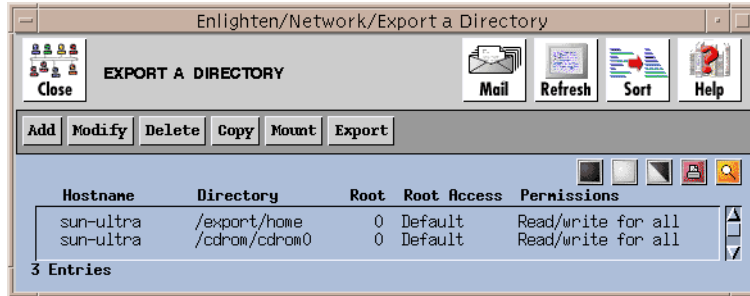
Figure 4-17 Export a Directory window

Each line in the list displays the hostname, exported partition, number of hosts exported to, root access, and permissions.

The Modify an Exported Filesystem window is similar to the Export a Filesystem window. In the Modify window, the Servername field and Directory to Export field are read-only.

4
Managing the
Network Systems

2) Highlight one or more partitions to be modified, then click the Modify button. The Modify an Exported Filesystem window appears.
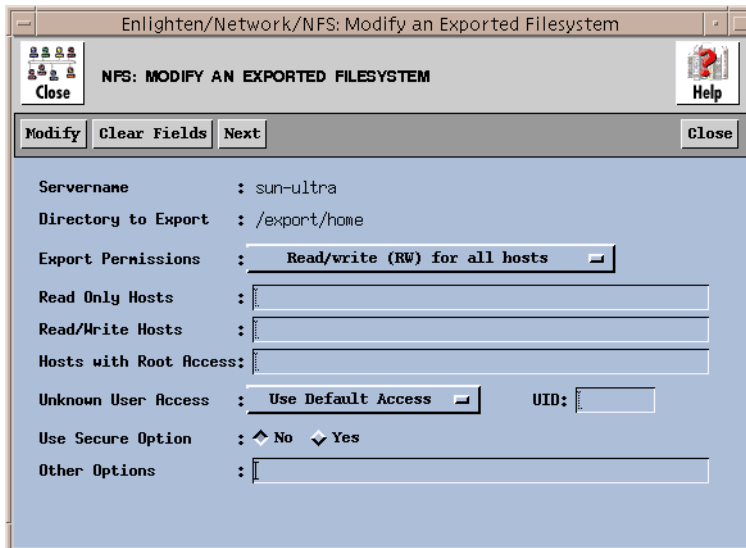
Figure 4-18 Modify an Exported Filesystem window

3) Change the parameters as desired. Refer to the next section, "Add and Modify Export Fields," for information about each parameter.

4) Click the Modify button to save any changes.

5) If more than one partition is selected in the Network File Configuration window, click the Next button to display and modify the parameters of the next selected partition. Be sure to click Modify before moving on to the next partition or the new settings will not be saved.

Add and Modify Export Fields

Server Name field

Enter the NFS server from which the partition will be exported (or click the right arrow button to choose from a pick list of all hosts within the current pool).

Directory to Export field

Enter the full pathname for the directory to export. For multiple entries, leave a blank space between each directory name.

Export Permissions option

These toggles specify what directory access permissions the client machines will have. 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 field

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 field

This field is active only if you selected the 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 field

Enter the hosts that will have root equivalency permissions on this partition. Leave a blank space between hostnames for multiple entries.

Unknown User Access option

This toggle determines how unknown user (UID) access should be handled. This setting also determines how any root access (other than those hosts specified in the previous Hosts with Root Access field) should be handled. The options are:

- Use Default Access (the default). If the default is unknown, it puts UID_NOBODY in the file.

- Disable Unknown Access. Access is denied.

- Select own UID. Only the specified UID has access.

UID field

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

Use Secure Option

Choose whether the filesystem should be exported with the Secure option. The default setting is No.

☞ This secure option does not work on all systems. Both the server and the client must be configured to be "secure" for this to work properly.

Other Options field

Use this field to specify other NFS options. This field is not pre-checked for syntax errors.

Managing Cron Jobs

This section describes how to manage crontab entries on multiple heterogeneous hosts. The Cron Management options in the System menu are used to configure cron jobs, perform queries on all cron jobs, or manage user access to crontab.

Configuring Cron Jobs

The Crontab Configuration window (Figure 4-19) displays current cron jobs for all users on all hosts in the current pool. This window also allows you to:

- Create a new crontab job.

- Modify an existing crontab entry.

- Delete selected crontab jobs.

- Create a cron job using the settings of an existing cron job (copy).
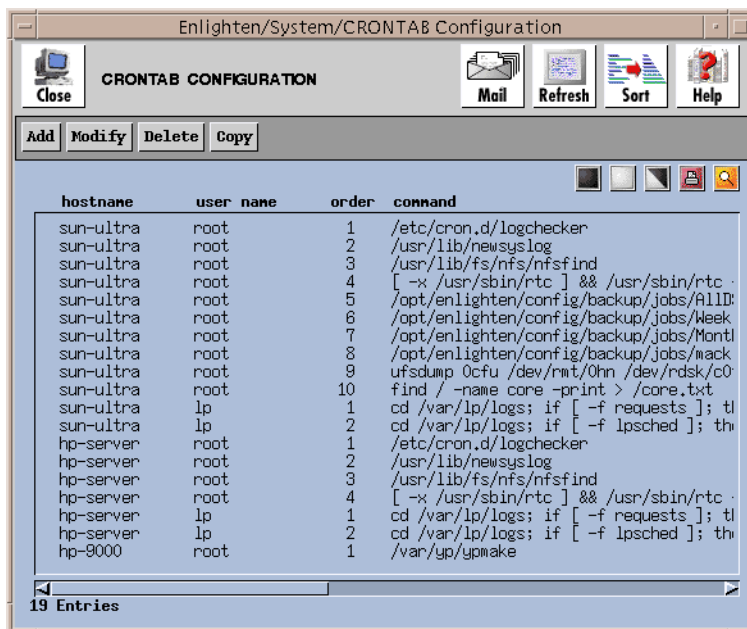
Figure 4-19 Crontab Configuration window

Adding New Crontab Jobs

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.

To create a new crontab job:

1) From the System menu, choose Cron Management and then Configure.

2) Click the Add button in the Crontab Configuration window. The Add Crontab Entry window appears (Figure 4-20).

Figure 4-20 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 of the week) define when the command (sixth field) will execute. If all the fields are left blank, the job will run by default every hour.

3) Enter the Add Crontab Entry parameters. See the section, "Add/Modify Crontab Entry Parameters," on page 4-29 for information about the fields in this window.

4) Click the Add button to create the crontab job.

Modifying a Crontab Entry

To modify a crontab entry:

1) From the System menu, choose Cron Management and then Configure.

2) Highlight the entries to be modified in the Crontab Configuration window (Figure 4-19).

3) Click the Modify button. The Modify Crontab Entry window appears.
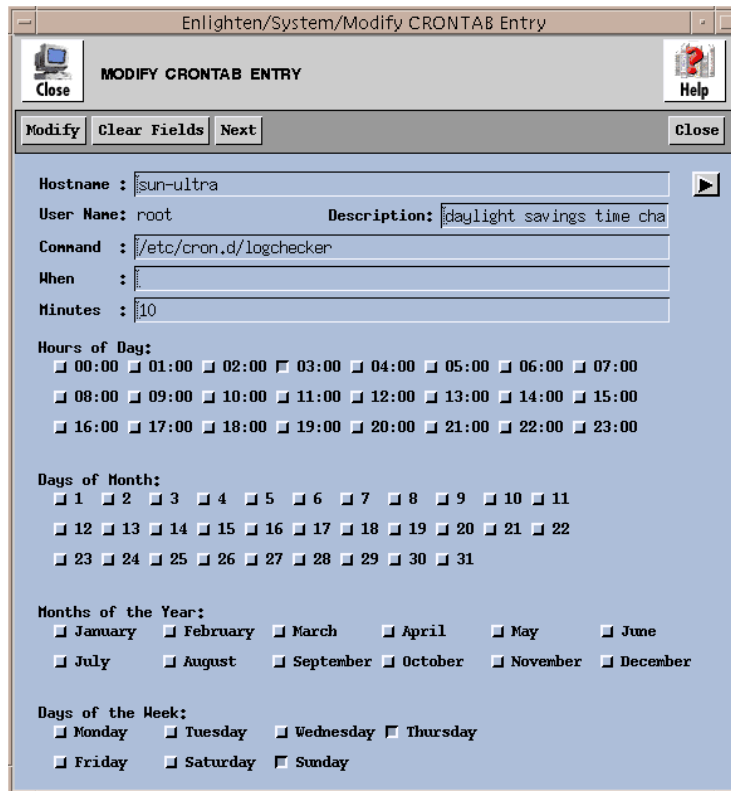
Figure 4-21 Modify Crontab Entry window

4) Change the entry parameters as desired. See the next section,
   "Add/Modify Crontab Entry Parameters," for information about each
   field.

5) Click the Modify button.

6) If more than one cron job was selected for modification, click the
   Next button to modify the next entry. Be sure to click Modify before
   moving on to the next cron job or the new settings will not be saved.

Add/Modify Crontab Entry Parameters

The Modify Crontab Entry window is similar to the Add Crontab Entry window except that the Username field is read-only.

Hostname field

To limit this job to specific hostnames within a pool, enter those hostnames in this field. For multiple entries, leave a blank between each hostname. You can also use the right arrow button to choose from a pick list of available hosts in the current pool.

User Name field

Enter the user name for the user who will run this job (or click the right arrow button to choose from a pick list of users).

Description field

Enter a brief description of the job's purpose.

Command field

Enter the cron command (and arguments) to be run at the time specified by the remaining five fields. A maximum of 256 characters are allowed for this field.

Minutes field

Enter the minute of the hour(s) the job will run. Enter a whole number between 0 and 59. You can also specify a range by using a hyphen (-); for example, 0-20. For multiple entries, leave a blank between each entry.

Hours of Day
Days of month
Months of the Year
Days of the Week options

Click the appropriate boxes to further define when this job will run. Multiple values per field are allowed.

4 Managing the Network Systems

Querying Cron Jobs

Specific cron jobs can be located by searching for hostnames, usernames, the command itself, the time settings, and so on. Once the query is successful, the cron job can be modify, deleted, or copied.

To search for a cron job:

1) From the System menu, choose Cron Management and then Query. The Cron Entry Query window appears.
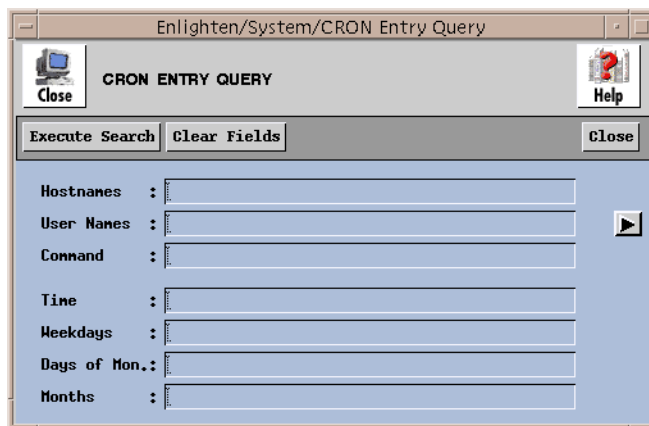
Figure 4-22 Cron Entry Query window

For a successful match, all filled-in entries in the Query fields must exactly match those in the Cron job (logical *and*). Blank fields always match. Possible matches can be negated within fields using the ! character (logical *not*): for example, `!Monday`.

2) To limit the search to specific hostnames within a pool, enter those hostnames in the Hostnames field. For multiple entries, leave a blank space between each entry.

3) To limit the search to specific users, enter those user name(s) in the User Names field. For multiple entries, leave a blank between each entry. You can also use the right arrow button to choose from a pick list of users.

4) To limit the search to a specific cron command (or any of its options), enter the command in the Command field. A maximum of 256 characters are allowed in this field, including the standard wildcards.

5) To specify times or time ranges, enter the parameters in the Time, Weekdays, Days of Mon., and Months fields. This field limits the search to jobs executing at/between the times entered.

6) After entering the search criteria, click the Execute Search button. If this button is clicked with all the fields empty, all current cron jobs will be displayed. A window similar to the Crontab Configuration window appears with the results listed.

The cron jobs listed in the Crontab Configuration list can be modified, deleted, or copied.

Managing Cron Users

Execution of cron jobs can be limited to specified users . There are five user access options:

- 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.

To control user access:

1) From the System menu, choose Cron Management and then Cron Users. The Cron User Access Configuration window appears (Figure 4-23).
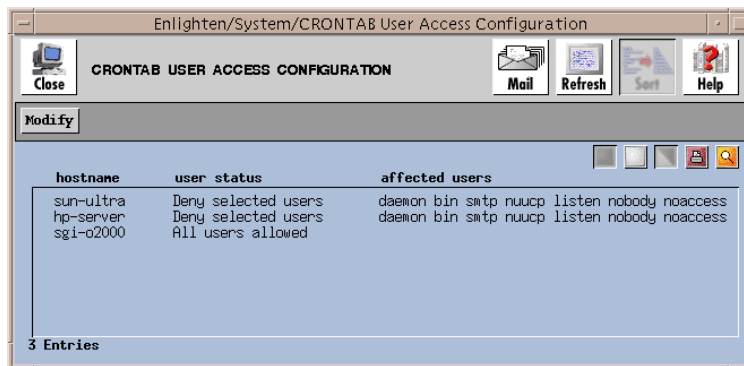
Figure 4-23 Crontab User Access Configuration window

2) Highlight the users to be modified and click the Modify button. The Modify Crontab window (Figure 4-24) appears.
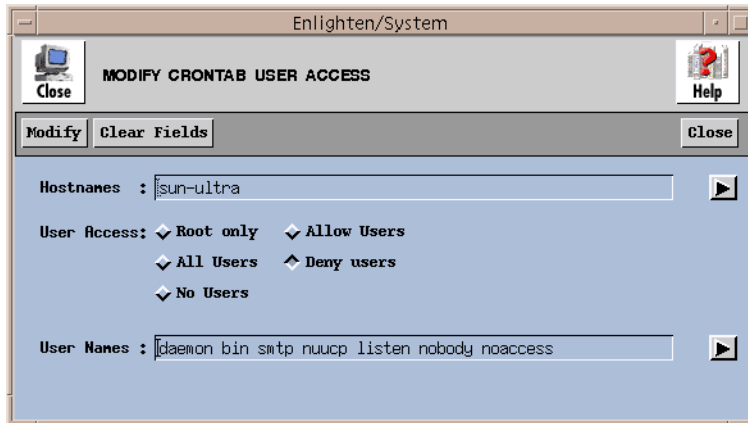
Figure 4-24 Modify Crontab window

3) Use this window to modify the cron user access for multiple hosts.

4) To limit this modification to specific hostnames within a pool, enter those hostnames in the Hostnames field. For multiple entries, leave a blank space between each hostname. You can also click the right arrow button to choose from a pick list of available hosts in the current pool.

5) The User Access option specifies what type of cron access users are allowed.

6) To allow or exclude (deny) specific users, enter the user name(s) in the User Names field (or click the right arrow button choose from a pick list of users). For multiple entries, leave a blank between each user name.

7) Click the Modify button to save the crontab user access modifications.

Archiving Files

Essential to every network manager is the ability to safeguard files through automatic backups. Enlighten**DSM** can be used to configure devices, perform tape operations, restore files, and schedule full and incremental backups.

This section provides basic information about configuring devices and performing backups. For information about backing up to tape drives, restoring files, and cataloging backup files, refer to Chapter 8, "Archive," in the Enlighten**DSM** *Reference Manual*.

Adding Backup Devices

A variety of different devices can be used to back up system files, including custom devices, reel-to-reel tapes, 1/4-inch cartridge tapes, or raw device types.

To add a new device:

1) Choose Configure Devices from the Archive menu. The Archive Device Configuration window appears, listing the available devices.
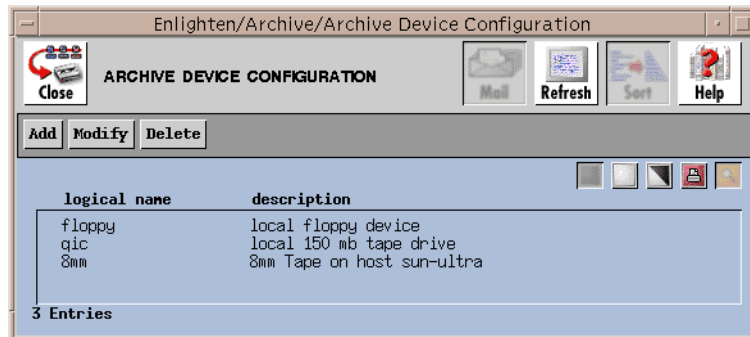
Figure 4-25 Archive Device Configuration window

2) To add a new device, click the Add button. The Archive Device Add window appears (Figure 4-26).
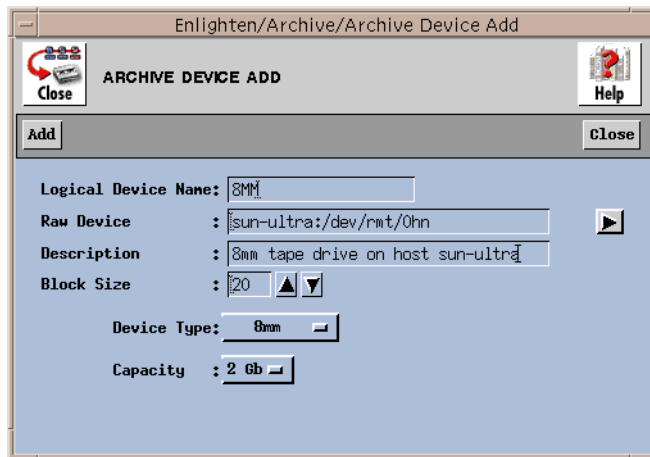
Figure 4-26 Archive Device Add Window

3) In the Logical Device Name field, enter the device name that the user specifies when accessing this backup device.

4) In the Raw Device field, enter a 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.

5) Enter a device description or name in the Description field.

6) Choose a block size from the Block Size option.

   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. Block size is O/S dependent. On Solaris 2.4, the block size is 1024 bytes (1K). On SCO 5 platform, the default block size is 1K.

7) Choose the device type from the Device Type option. The Device Type specifies the device's capacity. The options are:

- The Cartridge Tape option displays tape length and density. For length, click the left mouse button and then drag the cursor to the appropriate value: 300, 450, or 600 feet. For density, click QIC 24 (low) or QIC 120 (high).

- The Reel option displays tape length and density line. For length, click the left mouse button and drag the cursor to the appropriate value: 600, 1200, or 2400 feet. For density, click 1600 (low) or 6250 (high).

- The Custom option allows you to define your own device. You will be prompted to enter the capacity for the new device in KB.

- The Raw option allows you to quickly add an entry to the device list. Use this option to add a device with no capacity or a new piece of media. You should eventually redefine this storage device.

- Click the 8 MM button to define an 8mm cartridge tape device. You are prompted to select the capacity of the tape. The options are 2GB (default) and 5GB. If neither of these options applies, create a custom backup device and manually set the device capacity. You can also create a raw device and then define the device capacity.

8) Click the Add button to save the parameters.

4 Managing the Network Systems

Modifying Backup Devices

To view and/or modify the configured backup devices:

1) Highlight the device name to be modified from the Archive Device Configuration window (Figure 4-25).

2) Click the Modify button. A window will appear similar to the Archive Device Add window. In the Modify window, the Logical Device Name field is read-only.

3) Change the parameters as desired.

4) Click the Modify button to save the changes.

5) If multiple backup schedules were chosen, click the Next button to modify the next schedule. Be sure to click Modify before moving on to the next schedule, or the new settings will not be saved.

Deleting Backup Devices

To remove a backup device from Enlighten**DSM**, highlight the devices to be deleted from the Archive Device Configuration window (Figure 4-25), then click the Delete button. To ensure the correct backup device is being deleted, Enlighten**DSM** will prompt you for confirmation before deleting it. After an entry is removed, the only way to add it again is to recreate the entry.

Performing Backups

Either full or incremental backups can be performed on filesystems (disk partitions). Full backups include all files in a partition, while incremental backups include only files that have been modified since the last backup. Incremental backups are generally much faster and therefore more useful as a daily strategy.

The Backup Partitions window shown in Figure 4-27 is used for both full and incremental backups. Enlighten**DSM** prompts you for the information needed to create backups.
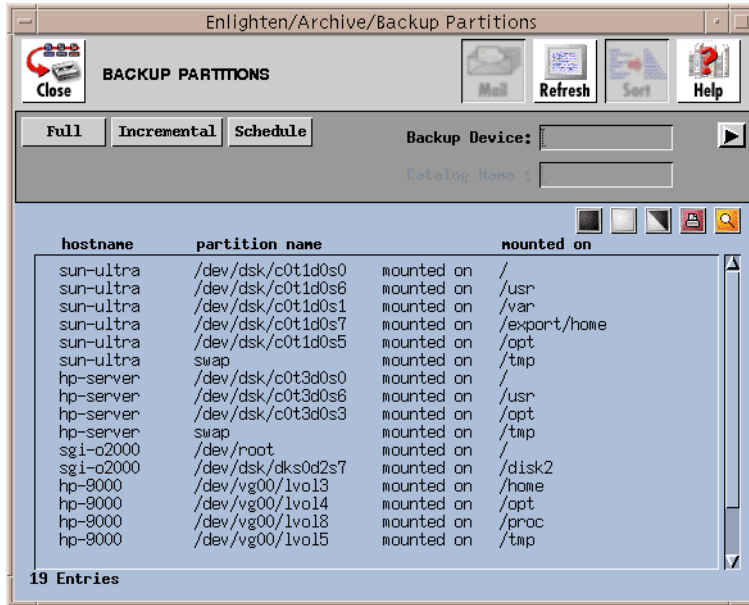
Figure 4-27 Backup Partitions window

Catalogs of backups are used to quickly locate and retrieve backed-up files. By default, Enlighten**DSM** does *not* keep a catalog of backed-up files. To turn cataloging on, go to the Configure menu, choose Session Preferences, then click Yes under the field: Maintain Catalog of Backed Up Files.

If the Catalog option is enabled, enter a catalog name in the Catalog Name field. This can be a filename or just an identifying tag. If no name is entered, Enlighten**DSM** uses the default name of NONAME.

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 Enlighten**DSM** performs the backup, a listing of each file archived 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 SessionPreferences 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 (Figure 4-28).
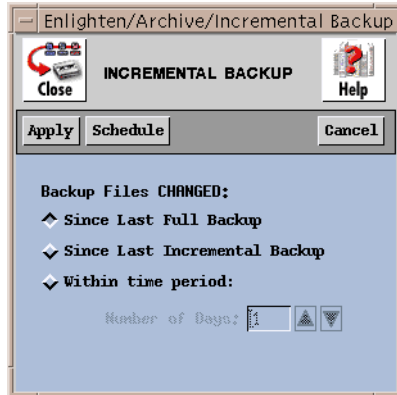
Figure 4-28 Incremental Backup window

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).

5) Click the Apply button. As Enlighten**DSM** performs the backup, a listing of each file archived will be displayed in the window.

Scheduling Backups

To automatically schedule full or incremental backups of selected files and partitions.

Choose Scheduled Backup from the Archive menu. The Scheduled Backup window (Figure 4-29) appears displaying a list of all currently defined scheduled backups, their schedule types, and descriptions.
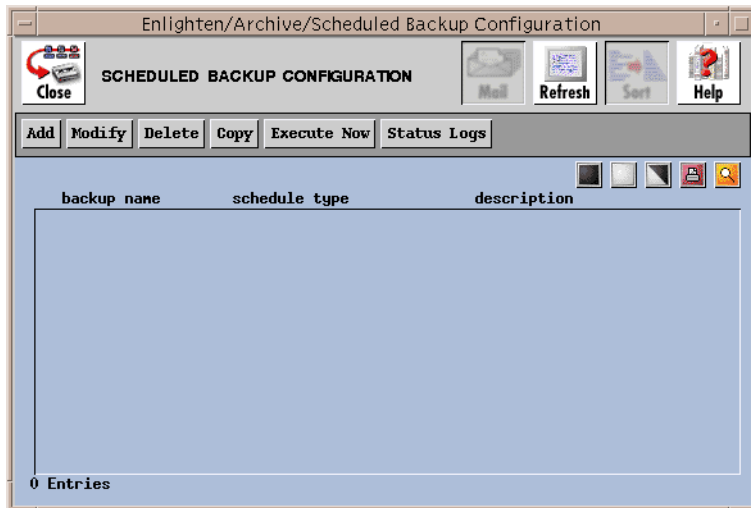
Figure 4-29 Scheduled Backup Configuration window

To add a scheduled backup:

1)   Choose Scheduled Backups from the Archive menu.

2)   Click the Add button. The Add Scheduled Backup window appears.
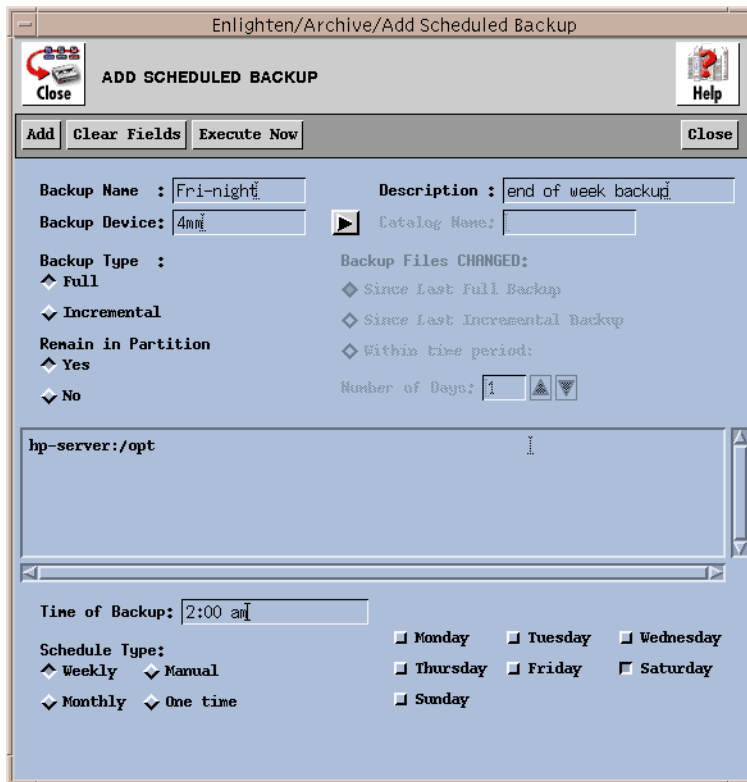
Figure 4-30 Add Scheduled Backup window

3) Enter the parameters for each field. For information about each field, see "Scheduled Backup Fields," on page 4-42.

4) Click the Add button to save the scheduled backup.

5) To immediately begin a backup, regardless of its scheduled time, click the Execute Now button.

Modifying Scheduled Backups

To modify a scheduled backup:

1) Choose Scheduled Backups from the Archive menu.

2) Highlight the backup to be changed.

3) Click the Modify button. The Modify Scheduled Backup window appears.

4) Change the parameters as desired. See the next section, "Scheduled Backup Fields," for a description of each option in the window.

5) Click the Modify button to save your changes.

6) If multiple backups were selected for modification, click the Next button. Be sure to click Modify before moving on to the next backup or the new settings will not be saved.

Scheduled Backup Fields

The Add and Modify Scheduled Backup windows contain the following fields:

Backup Name field

This field specifies the backup's name. This field is read-only in the Modify window.

Description field

Enter a brief description of the backup.

Backup Device field

Enter the logical device name for the backup (or click the right arrow button to choose from a pick list of all defined device names). Logical device names are defined via the Configure Device window.

Catalog Name field

If the Maintain Catalogs of Baked Up Files option is turned On in the Session Preferences window (Configure menu), enter the catalog name that will keep a record of this backup.

Backup Type option

Choose whether to run an Incremental or Full backup. The Incremental setting also activates the Backup Files Changed toggle buttons (see below).

Backup Files Changed options

Choose which files should be backed up during an Incremental backup. The choices are:

- Since the last full backup (the default setting).

- Since the last incremental backup.

- Within a specific time period (the last n days).

Choosing the Within Time Period setting displays the Number of Days field.

Number of Days field

Specify how many days back Enlighten**DSM** will check to see if a file has changed. Once the number of days has been exceeded, the file will be backed up. The arrow buttons to the right can be used to change the number displayed.

Remain in Partition option

Choose the Yes setting to backup only the immediately defined partition for a particular directory. Choose the No setting to backup all partitions listed under that directory. The default setting is Yes.

File List option

For incremental back-ups, this option determines the directories and files to be archived. Enter the directories and/or files to be checked for changes. Only files and directories that have changed will be backed up.

If this field is left blank, no directories or files are checked for changes and no backup is performed.

Remote files must be listed with the hostname followed by the directory name.

4 Managing the Network Systems

Time of Backup field

Use this field with the Schedule Type toggle buttons to determine when a backup will be executed. See the following description for more details on how this field is used.

To specify a time/date value, see Appendix C, "Time Formats," in the Enlighten**DSM** *Reference Manual*.

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 day(s) the backup should occur. Enter the time for the backup to occur in the Time of Backup field.

- Monthly   The days of the month will be displayed. Select which day(s) the backups will occur. Enter the time for the backup to occur in the Time of Backup field.

- Manual   The manual setting does not schedule a time to perform the backup. The backup will only run when the Execute Now button in either the current window or the parent Scheduled Backup Configuration window is clicked.

- One Time   Specify a date and time to run the backup in the Time of Backup field. This backup will only run once.

To specify a time/date value, see Appendix C, "Time Formats," in the Enlighten**DSM** *Reference Manual*.

Status Logs

When backups are run in the background at their scheduled times, a backup log is created. Click the Status Logs button in the Scheduled Backups window to access these status logs (Figure 4-31).
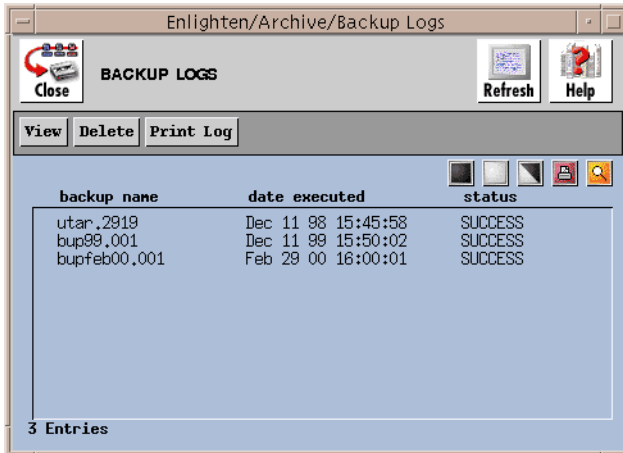
Figure 4-31 Backup Logs window

Click the Print Log button to print a copy of the selected backup logs. The output is sent to the printer defined in the Print field of the Session Preferences window (Configure menu).

Click the View button to view the selected backup logs.

Click the Delete button to remove the selected backup logs.

Configuring NIS and NIS+ Servers

This section describes how to use Enlighten**DSM** tools to configure master Network Information Servers (NIS and NIS+), slave servers (NIS), and replica servers (NIS+).

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

Adding an NIS Server

To add an NIS server or slave to your network:

1) Choose NIS and NIS+ from the Network menu. The NIS Servers window appears displaying all configured NIS and NIS+ servers. Each line in the list displays the type, domain, and hostname for each server.

Figure 4-32 NIS and NIS+ Servers window

2) Click the Add NIS Server button to configure NIS master servers or NIS slave servers. The Add NIS Server window appears.
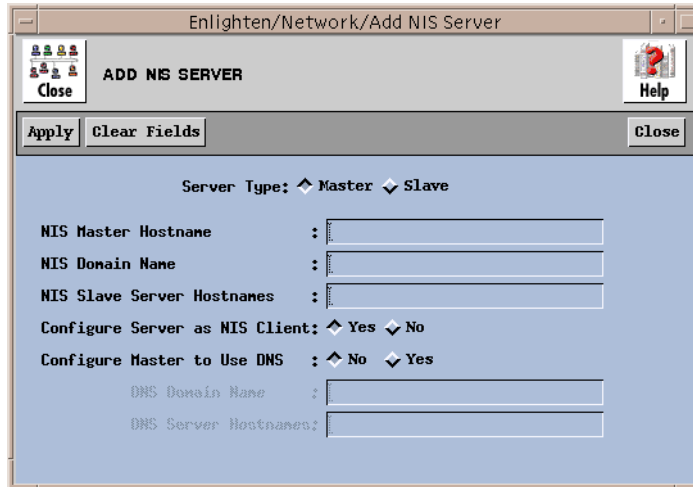
Figure 4-33 Add NIS Server window

3) Click the Master Server Type or Slave Server Type button to select the server type. The fields in the Add NIS Server window change depending on the Server Type option chosen.

4) Enter the field parameters for either the master or slave server. The master and slave fields are described in the next two sections.

5) Click the Apply button to add the new NIS server.

4 Managing the Network Systems

NIS Master Fields

When the Master Server Type option is chosen, the Add NIS Server window has the following fields (see Figure 4-33):

NIS Master Hostname field

Enter the hostname of the NIS master server.

NIS Domain Name field

Enter the domain name for this server.

NIS Slave Server Hostnames field

Enter any NIS slave servers for this domain.

Configure Server as NIS Client

Choose whether this server should also be configured as an NIS client. The default setting is Yes.

Configure Master to use DNS

Choose whether the NIS master server should use DNS as a fall back. The default setting is No.

DNS Domain Name field

This field can only be used if you set the Configure Master to use DNS field to Yes. If so, enter the domain name of the DNS servers the NIS master should reference.

DNS Server Hostnames field

This field can only be used if you set the Configure Master to use DNS field setting to Yes. If so, enter the domain name and hostnames of the DNS servers the NIS master should reference.

NIS Slave Fields

When the Slave Server Type option is chosen, the Add NIS Server window (Figure 4-34) has the following fields:
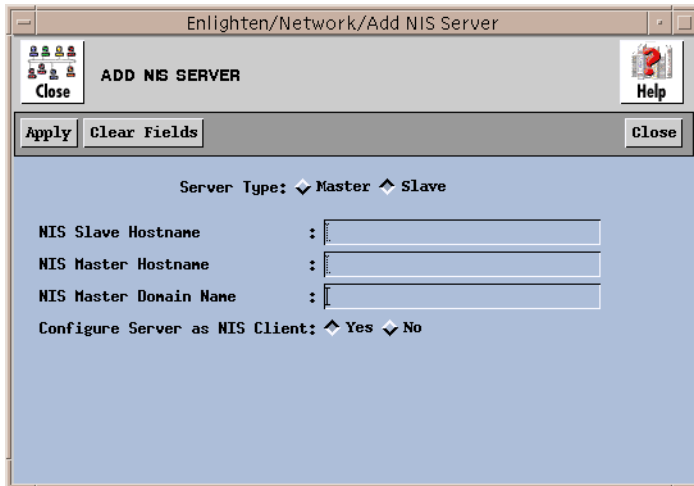
Figure 4-34 Slave Add NIS Server window

NIS Slave Hostname field

Enter the hostname of the NIS slave server.

NIS Master Hostname field

Enter the hostname of the NIS master server.

NIS Master Domain Name field

Enter the domain name for the NIS master server.

Configure Server as NIS Client option

Choose whether this server should also be configured as an NIS client. The default setting is Yes.

Adding an NIS+ Server

To add an NIS+ server,

1) Choose NIS and NIS+ from the Network menu. The NIS and NIS+ Servers window appears displaying all configured NIS and NIS+ servers. Each line in the list will show the type, domain, and hostname for each server.
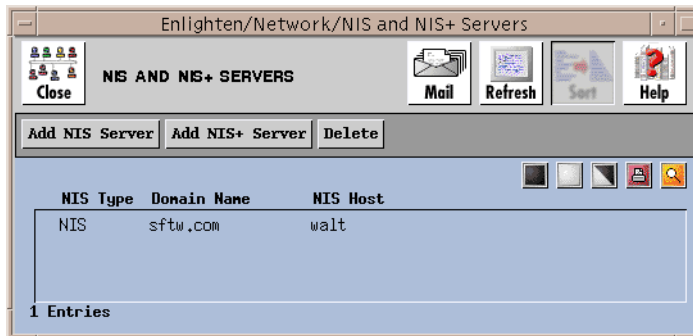
Figure 4-35 NIS and NIS+ Servers window

2) Click the Add NIS+ Server button. The Add NIS+ Server window appears.
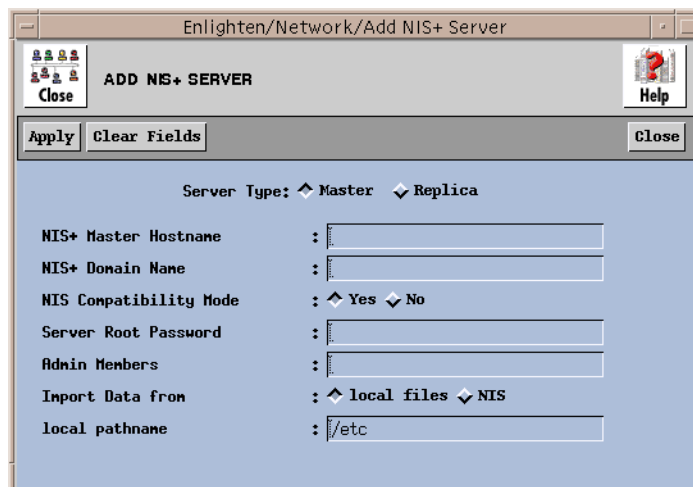
Figure 4-36 Add NIS+ Server window

3) Choose the desired NIS+ server type: a Master (the default) or Replica server. If the Replica server option is chosen, the window will be redrawn to show the fields listed in "Replica Fields," on page 4-52.

4) Enter the remaining field parameters. See the next two sections for information about NIS+ master and replica configurations.
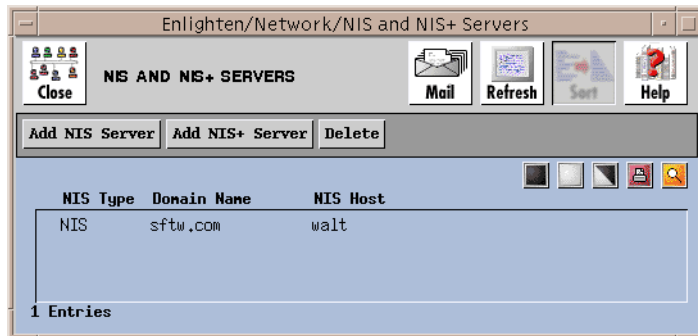
5) Click the Apply button.

NIS+ Master Fields

When the Master Server Type is chosen, the following fields are displayed on the Add NIS+ Server window:

NIS+ Master Hostname field

Enter the hostname of the NIS+ master server.

NIS+ Domain Name field

Enter the domain name for the NIS+ master server. The domain name is used to uniquely identify users and hosts that are making NIS+ requests.

NIS Compatibility mode

Choose whether this NIS+ should be made backwards compatible with NIS. The default setting is Yes.

Server Root Password field

Enter the root password for the NIS+ master server.

Admin Members field

Enter the users that are allowed to administer this NIS+ host. Leave a blank space between names for multiple entries.

Import Data From field

Choose whether any existing imported data should come from the servers' /etc files or from an NIS server. The default is /etc.

/etc pathname field

You can use this field only if you set the Import Data From field to /etc files. If so, enter the files' pathname.

NIS Domain field

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

NIS Server field

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

Replica Fields

This window is similar to the Add NIS+ Server window for a master server shown in Figure 4-36. It contains the following fields:

Server Type option

Choose what type of NIS+ server to configure: a Master (the default) or a Replica server. If you set this to be a Master server, the window will be redrawn to show the fields listed in "NIS+ Master Fields," on page 4-51.

NIS+ Master Hostname field

Enter the hostname of the NIS+ master server.

NIS+ Domain Name field

Enter the domain name for the NIS+ master server.

Replica Hostname field

Enter the hostname of the NIS+ replica server.

Replica Server Root Password field

Enter the root password for the NIS+ replica server.

NIS Compatibility mode

Choose whether this NIS+ server can also interact with NIS servers. The default is Yes.

Deleting Server Configurations

To delete an NIS or NIS+ server configuration:

1) Select the server you want to delete from the NIS and NIS+ Servers window.

2) Click the Delete button. The Delete NIS/NIS+ Server window appears.
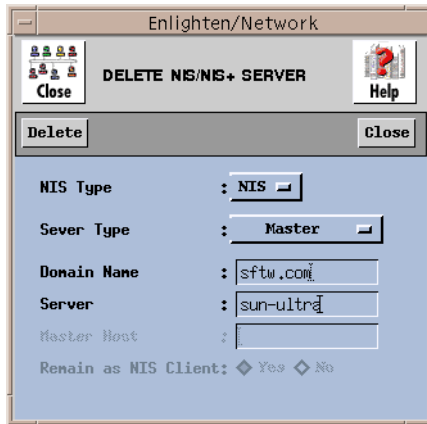
Figure 4-37 Delete NIS/NIS+ Server window

3) Click the NIS Type button to specify what type of server to delete, an NIS (the default) or NIS+ server.

4) Click the Server Type button to specify what type of server to delete, a Master (the default) or Slave/Replica server.

5) Enter the domain name for the server in the Domain Name field.

6) Enter the hostname in the Server field. If you set the Server Type field to Slave/Replica, enter the name of the NIS/NIS+ master server in the Master Host field.

7) If you set the Server Type field to Slave/Replica, select a Remain as Client option. Choose whether this server should remain as an NIS client after deletion. The default is Yes.

8) Click the Delete button.

☞ Enlighten**DSM** does not automatically detect whether a server is a master or slave/replica server.

5

Monitoring Networked Systems with Events

Enlighten**DSM** is a powerful tool for monitoring host and network problems. The Events feature collects and saves status, configuration, performance and capacity information, automatically sending warning messages 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 feature to graphically display the state of your systems. These two powerful tools identify which hosts or pools are having problems, or are about to have problems, allowing network managers to anticipate errors and take quick action when problems arise.

This chapter provides information about:

- Monitoring systems with Events

- Viewing system status with the Status Map

- Monitoring logins, processes, and CPU usage

- Auditing security checks

Using Events

This section describes how to add, modify and delete event tests using Enlighten**DSM**'s Events feature. For a detailed overview of Events capabilities and features, standards compliance, and the basics of building a testtab file, refer to Chapter 10, "Events," in the Enlighten**DSM** *Reference Manual.*

How Events Works

The Events feature helps predict problems by reporting the event and taking corrective action defined by the user.

Data collected by Events can be used to assist in tasks such as:

- system tuning
- load balancing
- resource planning and justification
- upgrade requirement analysis

Events collects this data by monitoring the following:

- memory subsystems
- individual files
- directory queues
- filesystems
- 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

- e-mail

- Programmable Events Processor (PEP) messages

Events can also pass the alarm for possible corrective action to a defined process. The same process can be assigned for all tests or a separate process can be specified for each test.

Events Test Types

There are two Events categories:

Group Tests—Also referred to as Internal tests, these tests are automatically configured with default configuration values. Refer to "Modifying and Existing Test" later in this chapter for more details. Group Tests include:

- O/S Tests
- File Systems Tests
- Printer Tests
- Process Tests
- Inventory Tests
- RPC Tests
- VM (Virtual Memory) Tests
- MBUF Tests
- Ncashe Tests
- IP Tests
- ICMP Tests
- TCP Tests
- UDP

Item Tests—Tests configured by the administrator using the Events Configuration Add feature. Refer to "Adding an Event Test" in this chapter for more details. Item tests include:

Process Tests:

- Instance of Process
- Size of Process
- Time Used by Process
- API built-in tests

File Tests:

- File Size
- File accessed
- File modified
- File clamped

Adding an Event Test

To create a new test:

1) Choose Configure from the Events menu. The Events Configuration window appears.
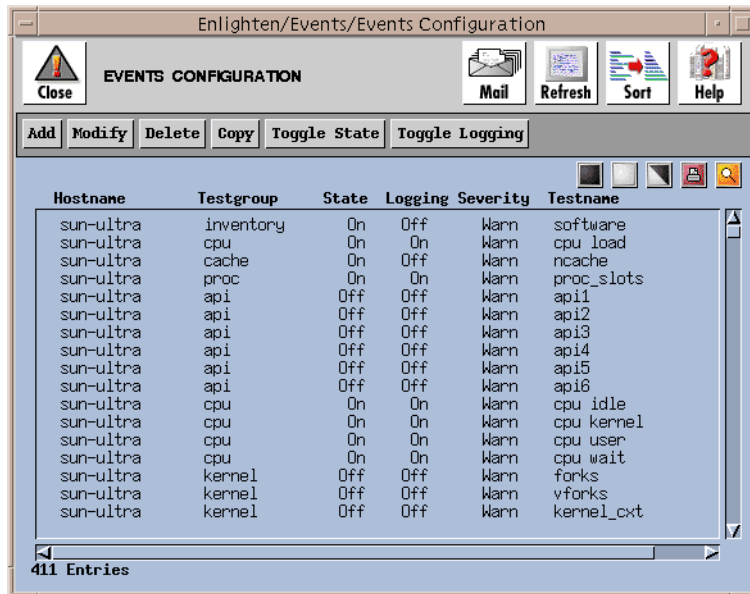
Figure 5-1 Events Configuration window

This window displays the hostname, the test group, whether the test is turned on or off, whether logging for the test is turned on or off, the severity level, and the test name for each test Enlighten**DSM** finds on your default host.

2) Click the Add button to select the test type for this new event test. A Select New Event Type window appears. Select the type of test to be created.
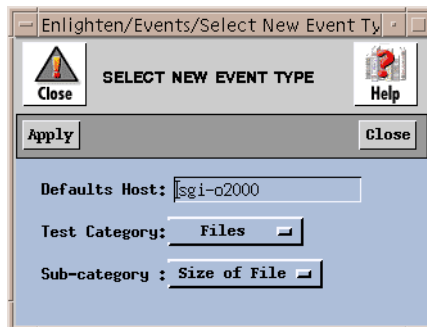
Figure 5-2 Select New Event Type window

3) Enter the hostname that will contain the default settings for the new tests.

4) From the Test Category field, choose the test type. The options are:

    •   Files (the default)

    •   Processes

    •   Directories

5) From the Sub-category field, specify a subcategory of the test types. If the Files option in the Test Category field was chosen, the type of file test can also be selected:

    •   Size of File (the default)

    •   Last Modified

    •   Last Accessed

    •   File Clamping
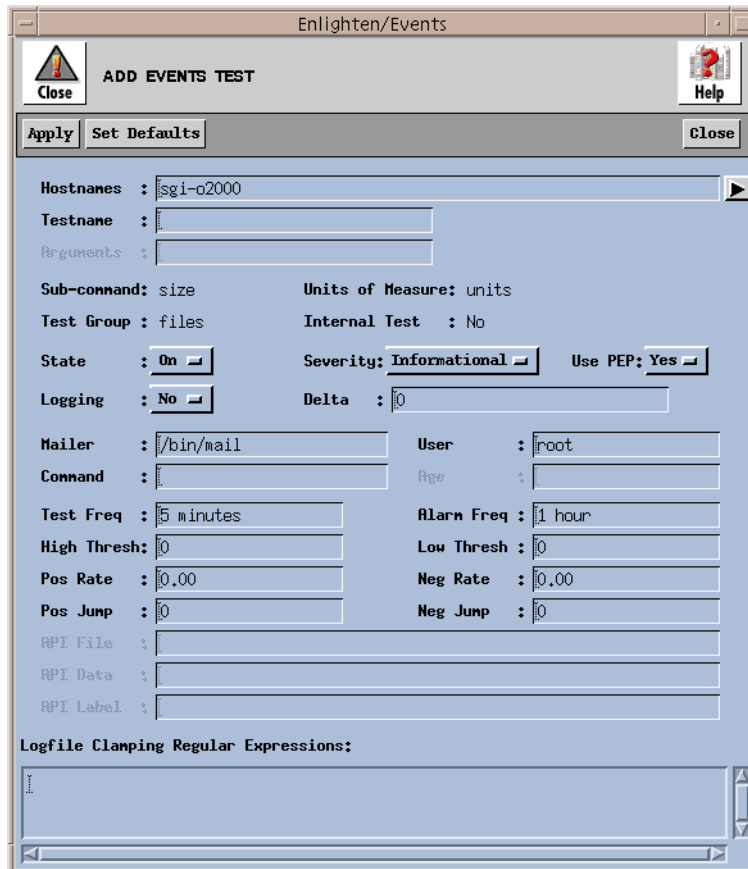
If the Processes option in the Test Category field was chosen, the type of process test can also be chosen:

- Instance of Process (the default)

- Size of Process

- Time Used by Process

If the Directories option in the Test Category field was chosen, there are no further subcategories to choose from.

6) Click the Apply button to display the Add Events Test window for the selected test type. Enter the parameters for the new event test.

Figure 5-3 Add Events Test window

Entering Test Parameters

For more information on the formats of these fields, or to run Events from a command-line mode, refer to Chapter 10, "Events," in the Enlighten**DSM** *Reference Manual*.

This section describes how to use each field and button in the Add Events Test window.

Hostnames field

Type the hostname(s) for this test (or click the right arrow button to choose from a pick list of all hosts within the current pool). Leave a blank space between hostnames for multiple entries.

Testname field

Enter the name of the test. This must be either the process name, or the full pathname of the file or directory to be monitored.

Arguments field

An optional list of command arguments is used in matching a process name. This field can be used to differentiate between two process instances by also matching the argument list used by each of the processes.

Units of Measure field

This read-only field shows what the standard units of measure are for this test.

Sub-command field

This read-only field shows the Events-defined subcommand (if any) this test will use during its execution.

Internal Test field

This read-only field determines if this test is an Events built-in test (Yes setting) or a user-defined test (No setting).

Test Group field

This read-only field shows the test group type for this test.

State field

This toggle turns the test On (the default) or Off.

Severity field

This toggle chooses the level of severity to assign this test from the following message types:

- OK

- Informational (the default)

- Warning

- Error

- Severe

Use PEP field

Choose whether this test should use PEP to report its results and/or filter any action to be taken. The default setting is Yes.

Logging field

Choose whether logging should be enabled for this test. The default is Yes.

Delta field

If Logging is enabled, enter a "changed by" (delta) value. Enlighten**DSM** will record the most recent value measured by the test if that value differs by at least this delta amount from the previously logged value.

See Chapter 10, "Events," in the Enlighten**DSM** *Reference Manual* for more details.

Mailer field

Enter the mail program that should be used to deliver alarm messages. The default is /bin/mail. If another mail program is used, it must use the same syntax as the standard mail program for the target operating system.

User field

Enter the user(s) who should receive any alarm information. The default is value is root. Leave a blank space between each user name for multiple entries. If this value is set to `nobody`, no mail will be sent.

Command field

Enter any executable this test should run when it sets an alarm. This can be a script or a compiled executable.

Age field

Enter how much time (in minutes) must elapse before a file in a directory is considered to have "aged". Only files more than "aged" minutes old are counted as "old".

This value is only available if this test will monitor a directory queue.

Test Freq field

Specify how often in minutes to run this test. The default is every five minutes.

Alarm Freq field

Specify how long to wait in minutes before sending another new alarm about this test. The default is every hour.

High Thresh field

Specify an absolute high-level alarm set point for the data you're measuring in this test. This can be an integer or floating-point value.

Low Thresh field

Specify an absolute low-level alarm set point for the data you're measuring in this test. This can be an integer or floating-point value.

Pos Rate field

Specify a positive percentage change alarm set point for the data you're measuring in this test. This threshold checks the percentage of change by comparing the current test value with the last measured value. This must be a floating-point value.

Neg Rate field

Specify a negative percentage change alarm set point for the data you're measuring in this test. This threshold checks the percentage of change by comparing the current test value with the last measured value. This must be a floating-point value.

Pos Jump field

Specify a positive incremental change alarm set point for the data you're measuring in this test. This threshold checks for the change of *n* points by comparing the current test value with the last measured value. This can be an integer or floating-point value.

Neg Jump field

Specify a negative incremental change alarm set point for the data you're measuring in this test. This threshold checks for the change of *n* points by comparing the current test value with the last measured value. This can be an integer or floating-point value.

API File field

If you're creating an API test, specify the full pathname of the file that will hold the values you are monitoring.

☞ See Chapter 10, "Events," in the Enlighten**DSM** *Reference Manual* for examples and more information on creating API tests.

API Data field

If an API test is being created, specify which field or column holds the data value to monitor. Use a digit prefaced by an 'f' for a field number or by a 'c' for a column number. The default assumes this value is a field number. If you're using a column designator, each character in any input file line/row is handled as one column.

5 Monitoring Networked Systems with Events

API Label Field

If an API test is being created, specify which field in your file contains a descriptive word or label.

Logfile Clamping Regular Expressions field

Regular expressions can be used to define "types" of messages based on pattern matching. When one or more of these message types are found in a file, an alarm is sent to the agents specified in the test. Each time this test runs, it evaluates only those files added since the last occurrence of the test.

Apply button

Click the Apply button to add the test configuration to the testtab files for all specified hosts. The Events process monitoring the data will also be updated.

☞ Enlighten**DSM** is updated immediately and the testtab file is updated two minutes later.

Set Defaults button

Click the Set Defaults button to set default values in the text fields based on the test type category you previously specified in the Select New Event Type window.

Modifying an Existing Test

Click the Modify button in the Events Configuration window (Figure 5-1) to modify an existing test configuration. A window similar to the Add Events Test window will appear.
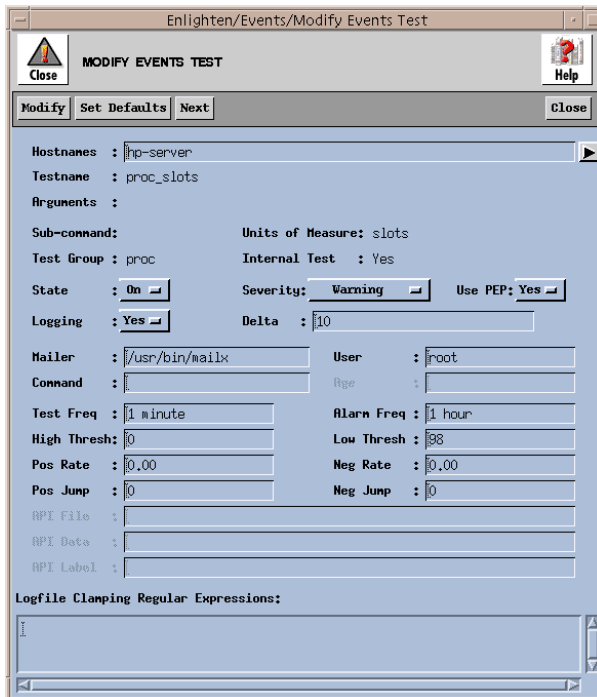
Figure 5-4 Modify Events Test window

See the previous section, "Entering Test Parameters" on page 5-8 for information about each field.

There are three differences between the Add and Modify Events Test windows:

- The Test Name field in the Modify window is read-only.

- The Modify button (rather than the Add button) is used to save changes.

- The Next button is used to modify additional test configurations if more than one test is selected for modification from the Events Configuration list. Clicking next will not save your changes so be sure to click Modify before moving on to the next test configuration.

Deleting an Event Test

Highlight the test to be deleted from the Events Configuration window. Click the Delete button. Enlighten**DSM** prompts you to confirm your action.

☞ Only tests *you have added* can be deleted. Built-in tests can be turned off, but cannot be deleted from the Events Configuration list.

Copying an Event Test

To create a new test using the parameters of an existing test, highlight the existing test and click Copy. The Add Events Test window will appear containing the original test parameters in each field. Edit this window as needed and click Apply to save the new test.

See "Entering Test Parameters" on page 5-8 for information about each field.

☞ You can only copy tests you have added.

Viewing System Status

The Status Map uses information provided by Events to graphically display the current state of hosts and pools using color-coded icons. You can ignore the status, or query events and fix the problem.

This section describes how to interpret and navigate the Status Map, query events using the Status Map window, and clear the Status Map of Events.

Interpreting the Status Map

To view the Status Map, choose Status Map from the Events menu.

A map similar to the one in Figure 5-5 appears.
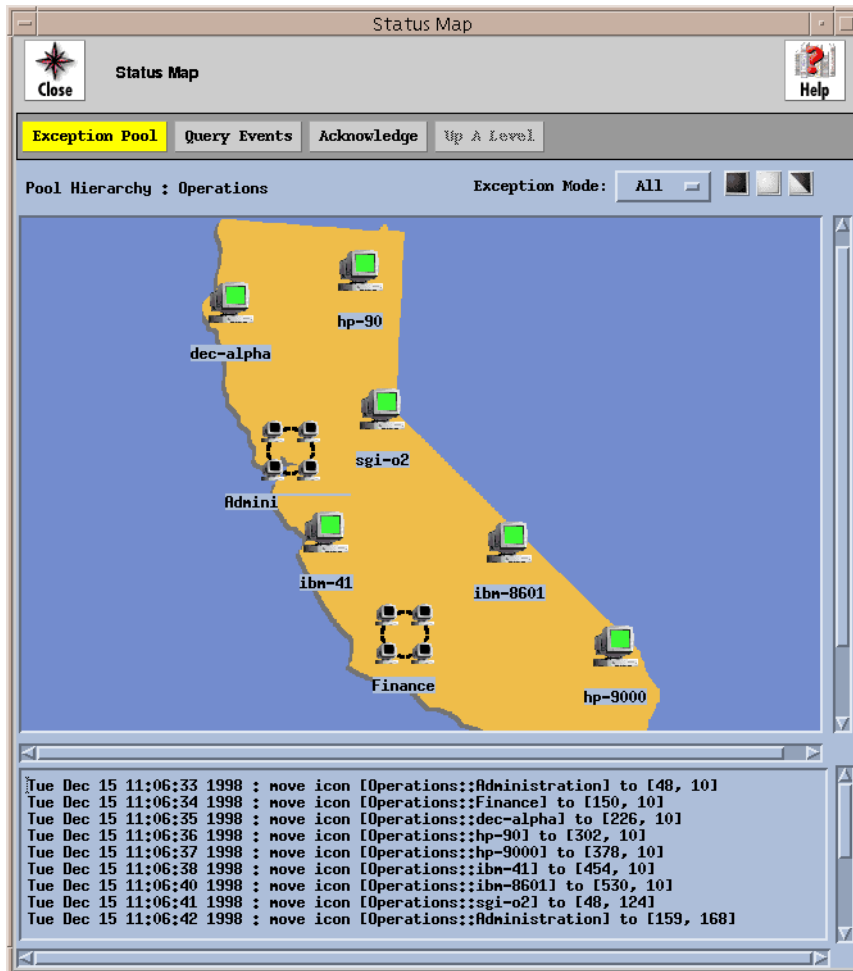
Figure 5-5 Example of a Status Map

The state of each host or pool icon is displayed according to the icon color:

- Green - OK

- White - Informational

- Yellow - Warning

- Blue - Error

- Red - Severe

The color of the host icon reflects the highest priority event that has not been cleared. The color of pool icons reflect the highest priority uncleared event for *any host* within that pool:

A blinking host or pool icon indicates the following:

- If a host icon is blinking, an unacknowledged event has occurred for that host.

- If a pool icon is blinking, an unacknowledged event has occurred for at least one host contained in the pool. An unacknowledged event is any event message that you have not yet acknowledged using the Status Map.

☞ If you haven't recently viewed the Status Map and an icon is blinking green, look at the preceding activity for that host or pool. An alarm might have occurred and cleared itself between viewings. For example, if an unauthorized user logged on to your system and then exited, and an Events test had already checked for this type of alert, a green alert would appear. All Enlighten**DSM** system administration functions act upon any hosts selected in the Status Map.

To change the Status Map background, place the cursor arrow over the map and click the right mouse button. A list of backgrounds appears from which to choose.

5 Monitoring Networked Systems with Events

Navigating the Status Map

To navigate the Status Map:

- If pool icon is selected, all hosts within that pool are selected and can be managed as a single unit.

- To perform administration functions on a subset of hosts in a pool, select those hosts by single-clicking on the host icons.

- The standard mouse "select rectangle" or "sweep" methods can also be used to define a temporary group.

Querying Events

To search and view Events alarms and messages logged:

1) Click the Query Events button in the Status Map window. The Status Map Query Events window appears.
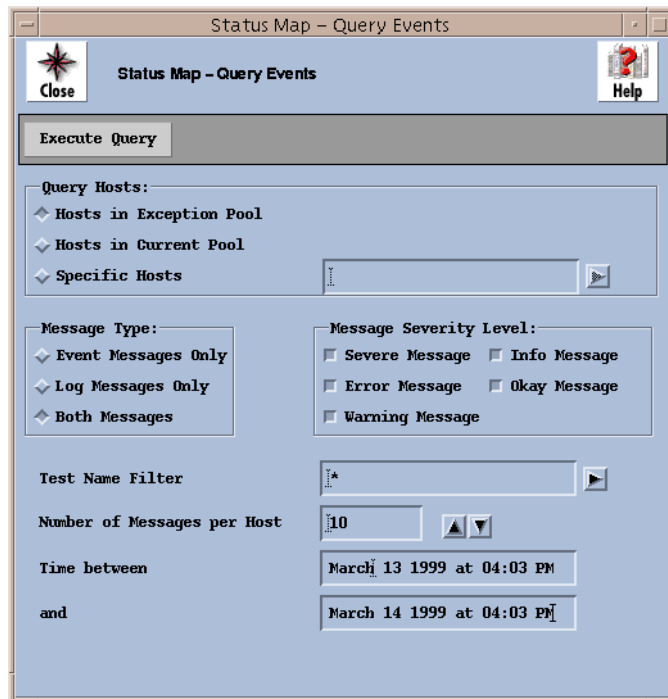
Figure 5-6 Status Map Query Events window

2) In the Query Hosts field, choose the hosts to be searched. The options are:

- Hosts In Exception Pool (the default)

- Hosts In Current Pool

- Specific Host(s). Specify which host(s) to check for messages in the text field to the right. Leave a blank space between hostnames for multiple entries.

3) Under Message Type, the options are:

- Event Messages Only: alarm messages generated when the results of a test violate a predefined threshold.

- Log Messages Only: informational messages logged when a test successfully runs without generating an alarm.

- Both Messages (the default).

4) Under Message Severity Level, the options are:

- Severe Message

- Info Message

- Error Message

- Okay Message

- Warning Message

Select one or more severity levels to use in the search. The default setting selects all levels of message severity.

5) Use the Test Name Filter field to limit the search to a few tests. Type the entire test name or just the first few letters of the test name in this field.

All tests names containing all or part of the specified string will be queried (or click the right arrow button to choose from a pick list of all pre-defined standard Events tests). Leave a blank space between test names for multiple entries. The standard regular expression wildcards '*', '[]', and '?' can also be used in this field (for example, /home/*).

---

6) In the Number of Messages per Host field, enter the number of messages per host to be searched. The most recent messages are displayed first. The counter buttons to the right also change the number displayed.

7) Use the Time Between fields to limit the search for message to those generated between a span of time. Enter the beginning and ending dates and times of messages to be searched. For a detailed description of date/time formats allowed in this field, refer to Appendix C of the Enlighten**DSM** *Reference Manual*.

8) Click the Execute Query button to begin the search process. The results are displayed in the View Event Message window (Figure 5-7).

Figure 5-7 Query results

All messages matching your search criteria are displayed in the list box. Each line includes the hostname, test name, logged value, units, severity, status, and timestamp.

To modify a test, highlight the desired test and click the Reconfigure Test button. The Modify Events Test window for that test will appear. For details on using this window to modify a test, see "Modifying an Existing Test" on page 5-13.

Clearing an Event from the Status Map

An event can be cleared from the Status Map when the condition that triggered the event no longer exists. There are two ways to clear an event:

- The event clears itself (for example, an Events CPU load test returns an OK result).

- You correct the activity that caused the original event/problem to occur (for example, by correcting whatever is overloading the CPU load).

After all events for a host or pool are cleared, the color for that icon is set to green (current status = OK).

Auditing Security Checks

Enlighten**DSM** provides a host of security features to check vital files, filesystem devices, boot and shutdown scripts, crontab contents, password integrity, group files, home directories, and break-in attempts. The findings from these security checks are placed in a logfile, allowing you to choose the functions that you want to audit.

For information on each one of the security features, refer to Chapter 6, "Security," in the Enlighten**DSM** *Reference Manual*.

Monitoring Logins, Processes, and CPUs

The Activity Monitor in the User menu is used to easily monitor login activity, process statuses, and CPU usage.

The rest of this section details how to use each Activity Monitor option.

Monitoring Logins

To quickly see which users are currently logged into the system, go to the User menu, choose Activity Monitor and then Who Is Logged In. The Who Is Logged In window appears (Figure 5-8).
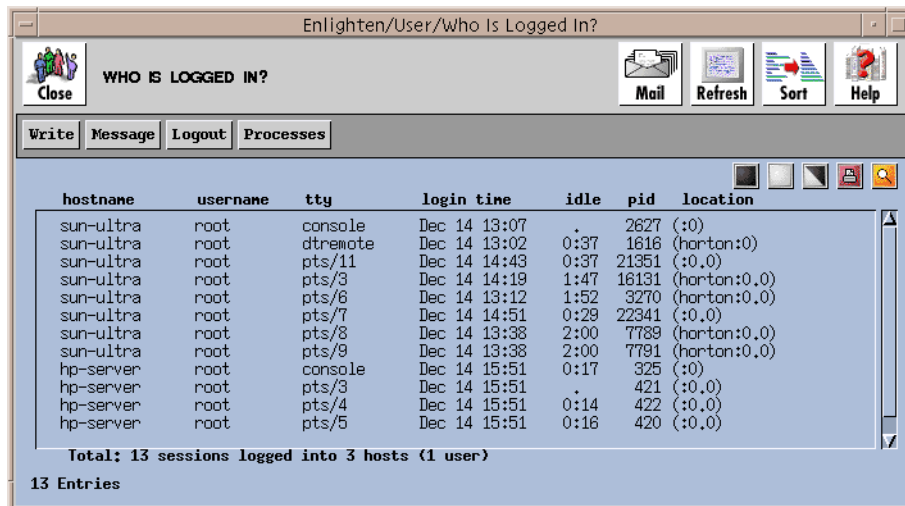
Figure 5-8 Who Is Logged In window

Each line in the window displays the hostname, user name, tty, login time, idle time, process ID, and the location of the tty (if available). To select one or more user accounts for further information, highlight the usernames and select one of the following options:

Writing Messages to Users

Click the Write button to write a new message directly to the selected user's mailbox. A window appears with a field for composing a message of any length. Press the Return key to send the message. The recipient can respond to this message, allowing for a two-way conversation.

Message

The Message command sends a predefined or custom form letter directly to the user's screen, instead of the user's mailbox. The recipient cannot reply to this message.

Logging Out

Click the Logout button to immediately terminate all highlighted work sessions. This command kills the initial Shell process belonging to the marked users.

☞ Use this command with caution as it may also cause related user processes to be killed.

Displaying Processes

Click the Processes button to display a window of all processes currently running for the highlighted users. To further manipulate this information, see the next section, "Monitoring Process Status."

5 Monitoring Networked Systems with Events

Monitoring Process Status

To view a list of all active processes, choose Activity Monitor and then Process Status from the User menu. The Processes window appears (Figure 5-9).
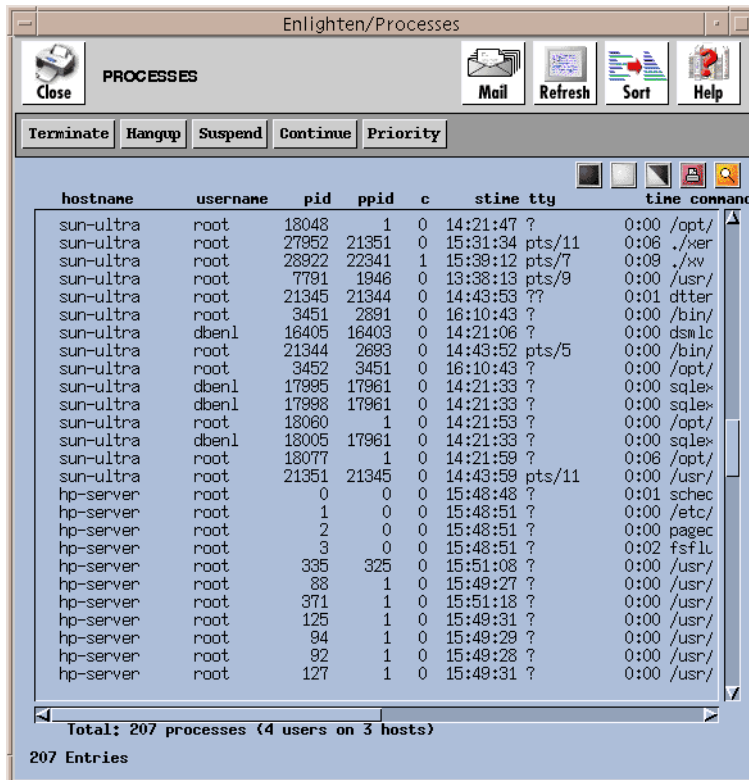
Figure 5-9 Processes window

Highlight a process and use the command buttons in the window to perform the following actions.

Killing a Process

Click the Terminate button to immediately kill the highlighted process. A pop-up window will prompt you to verify your Terminate command.

☞ This command will not kill related processes, so if there are child processes running, they will become orphans. These orphans might terminate automatically, or you may have to kill them manually.

Hanging Up a Process

The Hangup 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 to verify your Hangup command.

Suspending and Continuing a Process

Clicking the Suspend button will stop a process from working but will not terminate it. The process is put on hold and can be re-activated later. Click Continue to re-activate a suspended process.

Changing Process Priorities

Click the Priority button 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. You can enter the desired priority or use the arrow buttons to change the value.

Monitoring CPU Usage

From the User menu, choose Activity Monitor, then CPU Summary, The Summary of Process By User window appears (Figure 5-10). This window shows all currently logged-in users, the current number of processes, and the total cumulative CPU usage for each active user
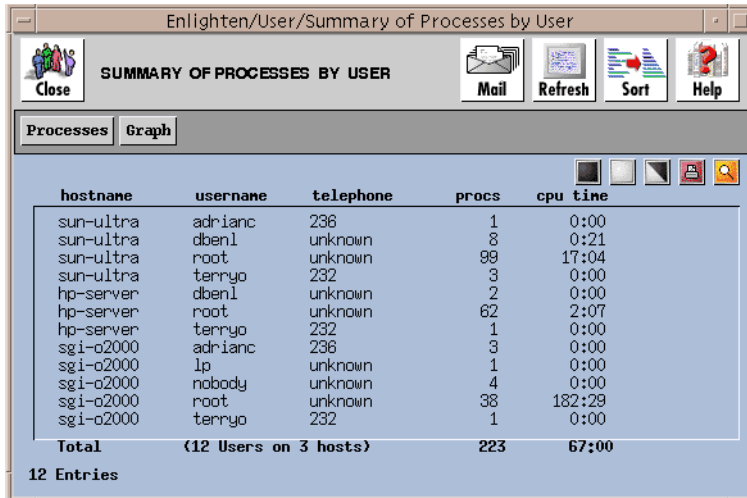
Figure 5-10 Processes window

Graphing Processes

To graph the processes, highlight the users to be graphed and click the Graph
button. A window appears displaying the highlighted items in a graphical
format. Press and hold down the middle mouse button to rotate the graph in
the direction you move the mouse.

Viewing Processes

To view the processes in detail, highlight the users to be viewed and click the
Processes button. A window appears displaying all processes for the
highlighted users. To further manipulate this information, see "Monitoring
Process Status" on page 5-24.

6

Fine-tuning Events with PEP

Enlighten**DSM** has two powerful tools for fine-tuning and managing networks. The Programmable Events Processor (PEP) is used to centrally manage Events so that it reports only on pre-determined issues.

The Enterprise Management Database (EMD) is an application manager in charge of all communication between the database and other Enlighten**DSM** components.

This chapter describes how to use these two tools.

Dispatching and Managing New Events

Enlighten**DSM** has a Programmable Events Processor (PEP) for centrally dispatching and managing new Events. Use PEP to have Events report only on issues you want to see. PEP can also communicate with external applications, such as Remedy.

Applications can inform PEP of noteworthy events by using an Remote Procedure Call (RPC)-based Application Programmers Interface (API). PEP will then perform whatever action(s) have been previously defined for that particular event. PEP can also log these events in its common logging facility.

PEP is implemented in Enlighten**DSM** as a daemon. There must be at least one instance of PEP per network, but no more than one instance per host machine. Most installations of Enlighten**DSM** will use a single PEP daemon, or a primary and secondary PEP daemon.

The diagram in Figure 6-1 shows how PEP interacts with the other main components of the Enlighten**DSM** package.
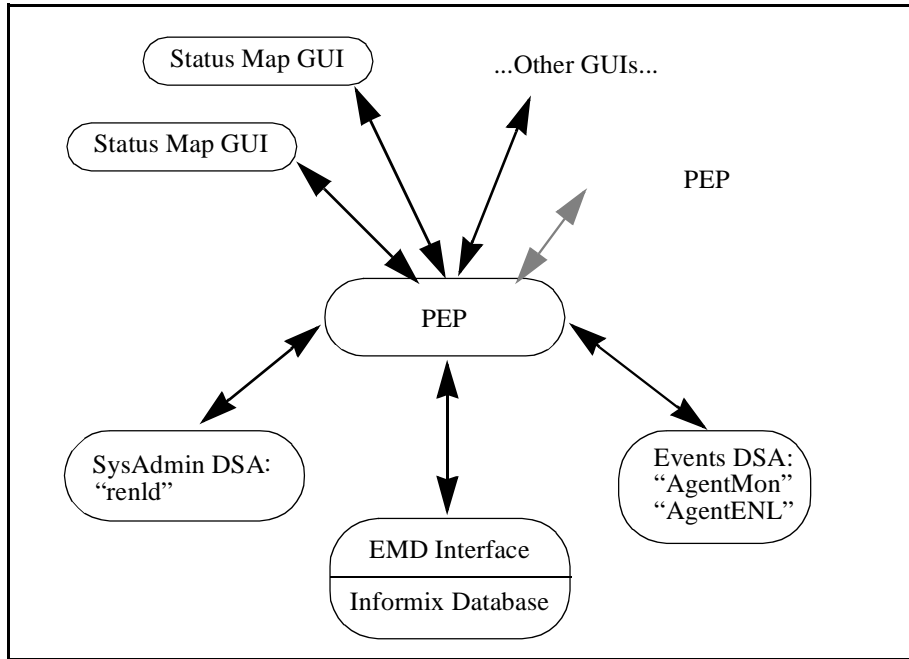
Figure 6-1 PEP interaction with Enlighten**DSM** components

Functionality

PEP plays a central role in determining how the parts of Enlighten**DSM** work. PEP can do the following:

- Accept event notifications via an RPC-based API.

- Write events to the common log via the Enterprise Management Database (EMD).

- Determine how to dispatch the event (using a rule-based policy engine).

- Dispatch events.

- Manage multiple Status Map GUIs reflecting the event activity.

- Manage the actions between GUIs, for example, acknowledging an event on one GUI will cause it to be acknowledged on all other GUIs.

- Forward events to another PEP (see Figure 6-1), allowing scalability for large networks.

- Help you query for an event's relationship to pools.

Policy Engine

PEP uses a rule-based dispatch table as its policy engine. Events are dispatched based on an ordered matching scheme. You can also add rules to define when this matching stops.

Each incoming event is matched against all defined rules (unless a stopmatch command is encountered). Events can be matched based on the following fields:

- Actual event ID (test name)
- Host where the event was generated
- Severity level of the event
- Application class of the event
- Application type
- Time the event occurred
- The description of the event

Actions

Based on the policy you define (or the defaults), PEP can take one or more of the following actions:

- Send e-mail
- Send notification to the Status Map
- Start a process you defined and run it on any managed host
- Log the event
- Forward the event to another instance of PEP (via RPC)
- Query for the relationship to a pool

Status Map GUI Dispatching

PEP maintains connections with all running GUIs to allow the inter-operability between them. This includes:

- Maintaining a list of current on-line GUIs

- Dispatching events to all GUIs

- Marshalling events' acknowledgments from the acknowledging GUI to all other connected GUIs

Policy Definition

The policy definition is a rules-based system contained in an ASCII file, so rules can be defined and changed dynamically as the system grows. Rules are scanned in the order they are defined.

Rules-Based Language

Using this rules-based method, new events can be generated and complex relationships between events can be created. The policy engine supports the following elements:

- Complex, nested if/then/else syntax

- Variables (local, global, static, and dynamic)

- Timers with one-second granularity

- Pattern matching on string fields

- Time of day, day of week, date matching

- Modifying the elements of the current event

- Querying for pool relationships

Timers can be used to set up rules based on multiple occurrences of an event within a given time period. Portions of an event can be modified, as shown in the following example (this effectively creates a new event from an existing one, which can then be further dispatched).

Common Logging

All data logs go directly into the EMD. All event logs go through PEP for additional dispatching.

Manually Killing/Invoking PEP

Occasionally, you may want to manually kill or invoke PEP. This is necessary, for example, when bringing down the system for maintenance or upgrading to a new release of Enlighten**DSM**.

The rest of this section details how to manually kill and restart PEP.

Terminating PEP

The following CLI command terminates the `pep` daemon:

```
kill -2 <pid#>
```

This command sends an interrupt to the `pep` daemon, allowing it to shut down gracefully.

Starting PEP

The following CLI command invokes the `pep` daemon:

```
./pep &
```

This command starts the daemon in the background. If you have changed the pathname for the `rules.txt` file (from `$ENLIGHTEN/contrib/rules.txt`), you need to use the following command instead:

```
./pep -f yourpathname/contrib/rules.txt &
```

where `yourpathname` shows the full pathname to the `rules.txt` file.

Program Start-up

When the program is started, a "startup" event is passed through the system. The event's name is set to the constant `EventInit`. This event should be used to explicitly initialize data such as timer states.

Program Layout

PEP programs must adhere to a certain layout. The general layout is:

```
# Comments start with "#" and
# Terminate at the end of the line
#
Global variable declarations
Timer declarations

Rule declaration # Comments can start any where
   Local Variable Declarations
   Program
;

Rule declaration
   Local Variable Declarations
   Program
;
```

More specifically:

- The program allows comments to begin anywhere on the line and extend all the way to the end of the line.
- Any global data must be declared before any rules are defined.
- Timers are implicitly global. They must come after any global variable declarations.
- Rules are terminated with a semicolon (;). You can use as many rules as necessary.

Program Execution

The engine program is executed every time PEP receives an event. External processes, such as SysAdmin agents or Events agents, can send an event using timers. PEP can also do this (see for more details).

There is a special event called a start-up event that gets generated internally at start-up. You should use this event to initialize timers.

Event Properties

Anything that is sent to PEP, most commonly by AgentMon in the form of an alarm, is considered an event.

PEP itself can generate events to itself in the form of:

- PEP Startup where the event.name = EventInit
- One for every timer expiration

Events are identified, inspected and modified by their properties. All non-timer events are prefixed by the keyword event followed by the property and are separated by the dot (.) operand. As in:

event.*property*

where *property* is one of the properties described in the table below.

Timer events can be identified by testing the event by the timer identifier. See Timer Declarations later in this chapter.

Constants

Event constants are values for some event properties. You can use constants to compare event properties and time states. Constants are strings and integers that are preset to a certain state, depending on the event. All constants begin with an uppercase letter. The table below contains constants for those event properties that use them.

Properties:

event.**name**

Data type: string

Testname that generated the event. Testname as described in Chapter 10 of the EnlightenDSM Reference Manual.

event.**time**

Data type: data/time

Time when event was generated in system local time

event.**hostname**

Data type: string

Name of the host that caused the event

---

event.**severity**

Data Type: constant

Severity level of the event. May evaluate by constant in either string or numeric representation.

Constant	Numeric Value
OK	1
Info	2
Warning	3
Error	4
Severe	5

---

event.**description**

Data type: string

Descriptive text string about the event. Consists of the following components:

"*testname measured-value threshold-type threshold-value timestamp*"

---

event.**state**

Data type: string

Threshold type such as:

- low
- high
- +jump / -jump

- +rate / -rate
- Age

For Timer events, event states have one of the following constants:

- Set
- Expired
- Cleared

---

event.**threshold**

Data type: string

The value of the threshold against which the measured value was compared
resulting in an alarm.

---

event.**measurement**

Data type: string

The value that resulted in the alarm.

---

event.**totalValStr**

Data type: string

Represents the total measured value for certain tests. The following tests
support total value.

Test	Subtest	totalValStr
File System	blocks totalValue	Max filesystem size
	inodes totalValue	Max number of inodes
File	modified totalValue	File size in bytes
	accessed totalValue	File size in bytes
Swap	totalValue	Max swap space capacity
Inventory	hardware totalValue	Total number of hardware items
	software totalValue	total number of software items

CPU	idle totalValue	total number of cpus
	user totalValue	total number of cpus
	kernel totalValue	total number of cpus
	wait totalValue	total number of cpus
CPU Load	totalValue	total number of cpus
File Clamp	totalValue	first line in the file to match the pattern searched for

---

event.**appclass**

Data type: constant

The application class which issued the event.

Constant	Source
EventClass	Events issued by AgentMon
CLIClass	Events issued via EventsCLI (Events Command-line Interface)
AdminClass	Unused

---

event.**units**

Data type: string

The unit type of the event as depicted in the Events Configure Window of the EnlightenDSM GUI.

Values such as:

- Unit
- Blocks
- Inodes
- Matches

event.**application**

Data type: constant

Application type that sent the event. Evaluate this to determine the type of operation that triggered the event

Constant	Meaning
EventAckOne	Event generated when a single event is acknowledged from the GUI.
EventAckAll	Event generated when a "Host Acknowledge" function is performed from the GUI.
EventDeletePool	Event generated when a host is deleted from a pool via the GUI.
EventChangeMap	Event generated when the background image is changed in the Status Map.
EventInit	Initialization Event upon PEP startup.
EventTimer	Event generated by PEP itself when a timer event changes state. To identify an event as a timer event, test the event.application for the value of EventTimer.

Modifying Events

Events can be modified by assigning values or constants to their event properties. A typical scenario would be to increment the severity of an event based on the measured value contained in the event. Lets say we want a CPU Load alarm, which is configured as Error severity (level 4) and a High threshold of 50, to escalate to Severe if the CPU load goes above 75. A rule to this effect would look like this:

```
rule escalateCpuLoad
    if ( event.name == "cpu load" and
event.measurement >= 75 )
```

```
        then
            event.severity = Severe
        endif
    ;
```

In the above example the event was configured to be event.severity=Error before it passed through PEP. As PEP inspects the actual measurement it changes the severity to the next severity level by using the constant Severe.

Any event property can be modified to effect a particular outcome of the event.

When evaluating and assigning values to variables and event properties using constants, <u>do not</u> enclose the constant key- word in quotes as in:

```
        If (event.severity = Error )
        then . . .

          or
        timer.state = Set
```

Data Types

The PEP language supports several different data types:

- integer numbers
- floating point numbers
- strings
- date/time
- timers
- events
- identifiers
- constants

Integer Numbers

The following are examples of valid integer declarations:

```
-1  32  42235523  09
```

The following are examples of invalid integer declarations:

```
+24  +2442533
```

Floating Point Numbers

The following are examples of valid floating point number declarations:

```
4.2    0.332    -.442    -2.44    -0.243    89.0
```

The following are examples of invalid floating point number declarations:

```
3.2e02    +5.2
```

Strings

Strings are anything in quotes ("") except the quote character itself. All string data types in PEP support regular expression matching, including wildcard matching. The following are examples of valid string declarations:

```
"Hello There"
"He said *"
"123"
"[a-z].*"
"I can punctuate...,,,"
```

The following are examples of invalid string declarations:

```
"what "?"
" No end in site
was"sup
```

Date/Time

PEP supports partial time and date declarations. You can use the keyword timeofday to read the system's local time. Add the designators PM/AM to specify absolute time designation. This is what the engine looks for to distinguish between relative and absolute time references.

Dates are declared in U.S. notations and assume the local time of the host system. The notation is as follows:

```
month day year OR
month year
```

It does not support the European or Asian date notation of:

```
day month year OR
year month day
```

A few different styles of notations and separators are supported. The following are examples of valid date/time declarations:

```
12/95   12-95   12/3/1990   12-3-90
timeofday
Dec 4, 1992
Wed 9:00 PM
3 am
12:24:31 AM
Tuesday
```

The following are examples of invalid date/time declarations:

```
12:30
1995/12/3
1:1
15/7/96
```

Timers

Timers are special data types that are an extension of the date/time data type. Timers support relative and absolute time specifications. Relative time is either *hr:min:sec* or *min:sec*. The following are examples of valid timer declarations:

```
0:30   1:00:00   1:00
```

Remember that timers support all date/time declarations as well. Timers can also be manipulated programmatically. They have two fields:

Field name	Data type	Description
state	string	Current state of the timer (Set, Cleared, Expired). See "Constants" on page 6-16.
time	timer	Absolute or relative time of timer.

Identifiers

You can use identifiers to name variables, rules, and timers. They must begin with a letter and can contain letters, numbers, or underscores ('_'). The following are examples of valid identifiers:

```
a1234  bbc  ABC  A_1234
```

The following are examples of invalid identifiers:

```
1A  A-Z  bor%d
```

As a general rule, you should use lowercase letters for identifiers, since constants start with uppercase letters. This will help avoid confusion between identifiers and constants.

Constants

You can use constants to compare event subfields and timer states to preset states. Constants are actually strings and integers that are hardwired to a certain state. All constants begin with an uppercase letter. The following is a list of constants and their appropriate uses:

Events have one of these five severity types:

```
Okay
Info
Severe
Warning
Error
```

Events have one of these seven application types:

```
EventAckOne
EventAckAll
EventDeletePool
EventMoveIcon
EventChangeMap
EventInit
EventTimer
```

Events have one of these three application classes:

```
EventClass
AdminClass
CLIClass
```

Timers have one of these three states:

```
Set
Expired
Cleared
```

Program Syntax

The PEP language consists of:

- variables
- timer declarations
- rules declarations
- expressions
- implicit conversions

Variables

Variables can be either dynamic or static. All variables must be explicitly assigned and are implicitly typed:

```
variable static x = 2
```

This variable declares a static variable x of an integer type and assigns the value of 2 to it. The following are more examples of valid variable definitions:

```
variable static y = 4.2 + 44.2
variable cc = "Hello World"
variable backup_time = 10:30 pm
variable now = timeofday
```

The variable's domain is implicit where it is declared. Global variables are always defined before any rule declarations and are accessible in any rule. Local variables are always declared after the rule declaration and only accessible to the rule that declared it. Variables that are named the same as global ones will hide the globals and continue until they go out of the local scope.

Timer Declarations

You can use timers to set the time-in and time-out of events in PEP. Time-outs generate events themselves. Timer states are implicitly set to "CLEAR" and must be initialized explicitly using the start-up event:

```
timer backup = 4:30 am
timer april_fools = Apr 1
timer one_minute = 0:60
timer one_a_day = 24:00:00
```

Rules Declarations

Rules are declared with the key word rule and followed by a name using a valid identifier. Rules must have a body:

```
rule myrule
   variable static count = 0
   if ( event.severity > 0 )
   then
       increment count
   endif
;
```

The rule myrule uses a local static variable to count the number of events that have a positive severity.

Expressions

Assignments

Assignments are performed by using the equal sign (=).

> _variable_ or _event property_ = _expression_

where

variable is one of any locally or globally declared variable. See the Variables section in this chapter.

event property is any event property described under the Event Properties section in this chapter.

<u>*expression*</u> is one of the following objects:

- Integer
- Floating point value
- String
- Variable
- Constant
- Date/Time expression
- Event property
- Complex arithmetic operation
- Implicit conversions

Integer - is a value of 2^{31} for signed integers and 2^{32} for unsigned integers.

Floating point value - PEP supports floating point values to two digits to the right of the decimal point.

String - Any alpha-numeric sequence enclosed in quotes (" ").

Variable - A valid variable type supported by PEP.

Constant - A valid constant supported by PEP. See Constants in the Event Properties section in this chapter. <u>Constants are not enclosed in quotes.</u>

Date/Time expression - Any date/time expression supported by PEP. See the Date/Time section in this chapter.

Event Property - Any event property described in the Event Properties section in this chapter.

Complex arithmetic operation - The following arithmetic expressions are supported:

> +, -, / ,*, %

Increment - Increments a counter

> `Increment mycounter` **is the same as** `mycounter = mycounter + 1`

Decrement - Decrements a counter

> `Decrement mycounter` **is the same as** `mycounter = mycounter - 1`

String Conversions

Several event properties are string objects. Even though some of these properties have a numerical nature, such as event.threshold, PEP retrieves and processes them as string values because event properties have different meanings among the different types of tests.

Therefore, it is necessary to convert these string properties to integer when performing any type of value comparison or calculation. Using PEPs left-to-right evaluation of expressions, converting a string element to an integer involves proper ordering of the elements in the expression. The first element in the expression must be of integer type:

```
variable a = 1
variable b = event.threshold

newval = a * event.threshold

              or put more simply

newval = 1 * event.threshold
```

Date/Time Comparisons

Absolute date/time comparisons compare only the common elements defined. For example, variables a, b, c, and d are set as follows:

```
variable a = 12/19/95
variable b = 2 pm

variable c = Tuesday
variable d = timeofday
```

Assume the current time is "Tuesday Dec 19, 1995 2:25pm". Comparing variable a, b, or c to variable d would evaluate to true because each comparison only compares the common date/time components. However, comparing a to b, b to c, or a to c would result in a false comparison since none share common components. If we added an additional variable e,

```
variable e = 2:25 pm
```

comparing e to b would result in a true expression since the only common components are the hour and am/pm designator.

Unsupported Conversions

The "string to time" conversion is not supported at this time.

Query Events using *in pool*

PEP supports simple queries. Use the following syntax:

```
host in pool abc_pool
```

This command searches the pool *abc_pool* recursively to see if the *host* is in that pool or its subpools.

This is useful when you set up your own pools and subpools. You can write a rule to check if an event came from a certain pool of hosts. For example, the following code fragment will check if an event came from the pool "important_pool":

```
if ( event.hostname in pool "important_pool" )
then
   event -> mail "user"
endif
```

Implicit Conversions - PEP supports manipulating different data types by using implicit data type conversion. Unlike C, which converts to the highest data type, PEP converts based on left-to-right evaluations of expressions.

Integer and floating point comparisons

```
variable  a = 2
variable  c = 4.2
variable  i = 0
```

In this first expression, $a + c$ is evaluated as an integer added to a float. Since the integer is the first expression, c is also converted to an integer. The temporary result is an integer with a value of 6, thus causing the if expression to evaluate to false.

```
if ( a + c > 6 )
then
     increment i
endif
```

In this second expression, `c + a` is evaluated as a float added to an integer. Since the float is the first expression, a is converted to a float. The temporary result is a float with a value of `6.2`. This temporary float also causes the constant 6 in the if expression to promote to a float of `6.0`, thus causing the if expression to evaluate to true.

```
if ( c + a > 6 )
then
        increment i
endif
```

Decisions

You can make decisions by using an "if/then/else/endif" expression in conjunction with logical expressions (==, !=, <, >, <=, >=, and, or). An if expression must have a then and a terminating endif statements; the else branch is optional.

```
variable count = 0
variable myval = 1

if ( event.hostname == "a*" )
then
   increment count
   if ( count >= 2 and my_val < 3 )
   then
      increment myval
   else
      decrement myval
   endif
endif
```

Use the stopmatch command to stop a rule's execution. When PEP encounters a stopmatch command, it will "break out" of the program and discontinue execution until the next event occurs.

Send Events

PEP's primary function is to direct events to an appropriate location or mechanism. PEP's logic, structure, and rule elements are designed to ultimately control, dispatch and otherwise suppress events. The control of events is referred to as Send Events.

Events can be sent to five separate targets:

- The Event Management Database
- Status Map
- Mail Recipient
- A shell command
- PEP at another host
- A SNMP destination

The syntax for sending an event is:

```
event -> destination
```

where destination is one of the following:

gui - Send the event to the Status Map portion of Enlighten**DSM**'s Graphical Interface.

log - Send the event to the Event Management Database (EMD) not store the event.

☞ It is possible to configure an event such that AgentMon will not forward alarms to PEP. In this case the event is, by default, forwarded to the EMD and the GUI by AgentMon. When PEP is employed for an event, the burden of forwarding the event to the EMD or log, as well as to the gui, is placed on PEP.

---

mail - Send the event to an e-mail recipient or recipient list where the recipient or list is enclosed in quotes (" "").

```
event -> mail "recipient [recipient]"
```

exec - Send the event to a shell command where the shell command and arguments are enclosed in quotes (" "). Any valid shell command or script can be used. Note that the shells $PATH variable must include the path of the command or script. Since the shell is launched by PEP, it will inherit the $PATH value of PEPs launching shell.

```
event -> exec Shell-command-or-script [arg1...arg2]
```

Environment Variables

When an event is sent by PEP to a destination, it will forward the event with all event properties. Particularly with exec, PEP will pass these properties to the shell in the form of environment variables. Here is a list of environment variables passed by PEP. Refer to the section Event Properties for details on these properties.

```
 Event Property     Corresponding environment variable
 event.description $ENL_DESCRIPTION
 event.time        $ENL_EVENTTIME
 event.hostname    $ENL_HOSTNAME
 event.measurement $ENL_MEASUREMENT
 event.severity    $ENL_SEVERITY
 event.state       $ENL_STATE
 event.name        $ENL_TESTNAME
 event.threshold   $ENL_THRESH
 event.units       $ENL_UNITS
 event.appclass    $ENL_APPCLASS
 event.application $ENL_APPLICATION
```

pep at *hostname* - Send the event to PEP at another host. Requires that PEP be running at the other host. All event properties will be passed along with the event so that PEPs rules at another host can be invoked and take action on the event. This is known as PEP re-direction. The host name must be enclosed in quotes (" ").

```
    event -> pep at "hostname"
```

PEP Engine Program Example

The following is an example of a PEP engine program to test the event e-mail process:

```
# Rule to log all events
#

rule log_all_events

    event -> gui
;
#
# During off-hours, page the on-call operator
#
rule page_after_hours
if ( event.severity = "Severe")
    then
        if ( timeofday <= 8:00 and timeofday >= 17:00 )
        then
         event -> exec "/opt/enlighten/scripts/PageOnCallOperator"
        endif
        event -> gui
    endif
;
#
# Uses the 'in pool' option to determine if event should go to
# all admins or just certain ones.
#
rule dispatch_to_host_owner
    if ( event.hostname in pool SENIOR_POOL )
    then
        event -> mail "SENIOR_ADMINS@company.com"
        STOPMATCH
    else
     event-> mail "SENIOR_ADMINS@company.com JUNOIR_ADMINS@company.com"
    endif
;
#
# Counts the number of icmp_error events and escalates the otherwise
# nominally sever event to Severe.
#
rule escalateICMPErrors

    variable static error_count = 0
```

```
if  ( event.hostname == "myhost" and event.name == "icmp_error" )
   then
       increment error_count
       if ( error_count >= 10 )
       then
          event.severity = Severe
          event -> gui
          event -> mail "admin@host"
          error_count = 0
          stopmatch
    endif
endif
```

Enterprise Management Database

The Enterprise Management Database (EMD) is in charge of all communication between the database and other Enlighten**DSM** components. It consists of three components:

- 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.

As an RPC-based server, the emdd communicates with the following Enlighten**DSM** components: Events agent, SysAdmin agent, Status Map, PEP, and the user interface.

The communication layer between emdd and the database is ODBC compliant. The drivers are Core and Level 1 API conformant and Minimum and Core SQL conformant.

The emdd daemon also checks daily to see if any data in the database needs to be expired. See Chapter 2, "Configure," in the Enlighten**DSM** *Reference Manual* for more details.

Relational Database

Enlighten**DSM** includes a relational database as part of EMD. The following data is 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
- audit log
- acknowledgment of events from the Status Map
- backup catalogs

- scheduling of backups
- host notes

The following data is not stored in a database:

- Events host-specific `testtab` files
- Events `AgentENL.config` file data
- snapshots of disks/file systems

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:

`bin`	contains the emdd daemon and scripts such as start_enl daemons
`dbtables`	contains the database tables, the database's transaction log, and a script to create the database tables
`odbc`	contains the ODBC drivers and other related components
`informix`	contains the database products
`msg`	contains any message catalog files
`log`	contains the emdd's logfile emdd.log

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.

Warning!   Do not modify or rename the `.odbc.ini` initialization file. Doing so will cause all subsequent EMD usage and connections to fail.

Manually Killing/Invoking EMD

Sometimes you may want to manually kill or invoke EMD; for example, when bringing down the system for maintenance or upgrading to a new release of Enlighten**DSM**.

☞   You *must* terminate EMD before backing up any database tables. You must manually re-invoke them afterward.

The rest of this section describes how to manually kill or invoke EMD.

Terminating EMD

You can bring down the `emdd` daemon with the following command:

```
stop_enl_daemons
```

This sends an interrupt to all Enlighten**DSM** daemons and allows the program 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 log file is generated, but only error messages are logged (if any occur).

A

Enlighten**DSM** **Basics**

Many of Enlighten**DSM**'s menus, buttons, icons and other tools are common throughout the program, allowing you to easily perform tasks and procedures without resorting to shell scripts. This appendix describes the following:

- Global functions

- List function buttons

- Performing actions on list items

- Creating defaults

- Network pools

- Performance expectation

- Data requests and sorts

Global Program Functions

Two sets of global function buttons are located in the top right corner of most Enlighten**DSM** windows (Figure A-1). These buttons allow you to easily search, sort or print the information listed in the window. There are also buttons for sending mail, making selections and analyzing you system. Each button is described below:
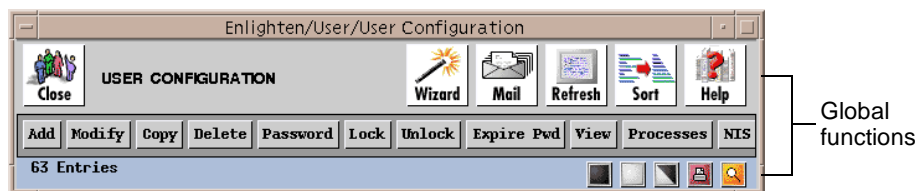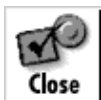

Figure A-1 Global function buttons

Close Icon

The Close icon in the upper left-hand corner closes windows without further modifying any data. Most Enlighten**DSM** windows also have a smaller Close or Cancel button for the same purpose.

Wizard Icon

The Wizard icon is available in the User Configuration, Configured Printers, Host Configuration, and Disk Usage Information By Filesystem windows. This feature helps analyze the security and consistency of your system. Refer to Appendix A, "Wizard," in the Enlighten**DSM** *Reference Manual* for information about the Wizard features.

Mail Icon

The Mail icon is used to send customized forms to highlighted user(s). Click on the Mail icon and the Mailing Instructions window appears (Figure A-2) This window includes a box for entering text.
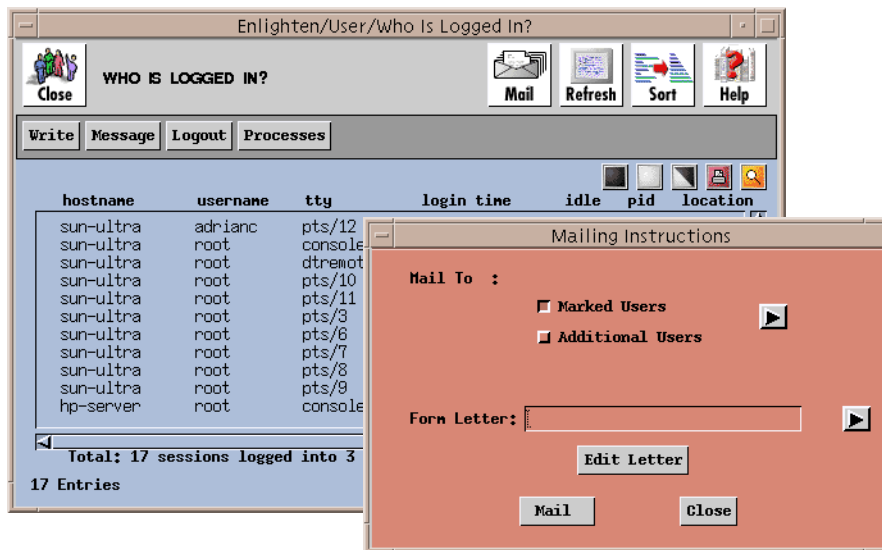
Figure A-2 Mail windows

Mail To field

Use this field to send mail to the highlighted user(s) and/or additional users who may not be on the current list. Click the right arrow button to choose from a pick list of additional users.

Form Letter field

Enter the name of the form letter to be sent or click the right arrow button to choose from a pick list of available form letters.

New form letters are created by entering a new filename in the Form Letter field and then clicking the Edit Letter button. A text editor appears. Edit the text as necessary and save the "new" letter.

Existing form letters can be modified by entering the name in the Form Letter field and clicking the Edit Letter button. The existing letter will appear in a text editor. Edit and save the letter as a different filename, then quit the original letter. Edited e-mail forms are saved in the directory where Enlighten**DSM** started.

☞ You must quit/close the form letter window before attempting to access other Enlighten**DSM** features.

Once the letter is ready, send it by clicking the Mail button. Enlighten**DSM** will prompt you to confirm your action.

Form Letters

Enlighten**DSM** is installed with three pre-defined form letters. These letter can be modified:

- Disk   Disk usage letter

- Obvious   Security letter regarding an obvious password

- System   System performance letter

These letters are stored in the file $ENLIGHTEN/config/formlet. Macros can also be used within the body of a form letter. For more details, see Appendix J, "Macros," in the Enlighten**DSM** *Reference Manual*.

Refresh Icon

When the Refresh icon is clicked, Enlighten**DSM** automatically updates the current screen as if the window was reselected. This is particularly useful when killing a process (to check if that process still exists).

Enlighten**DSM** can also be set to refresh screens automatically. See Chapter 2, "Configure," in the Enlighten**DSM** *Reference Manual* for more details.

Previous Selections

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. In such cases, the parent window is first regenerated and the same items are reselected (based on their order in the window).

The criteria for re-selection is based solely on position. If the first and third entries are highlighted, then the first and third entries will also be selected in the newly generated report. This is independent of the fact that in the first report, for example, the first entry referred to user root and the third entry to user bin.

Sort Order

When a window is refreshed using the Refresh button, the sort order previously selected for the report (via the Sort icon) will be preserved. When a report is generated from the main menu, the sort order is reset. Remember, selected items remain selected based solely on position (as described above in "Previous Selections").

☞ Except for NFS, disk-related windows that are based on previously generated snapshots do not have any refresh capabilities. Disk snapshots are used because of the heavy performance and resource load when scanning disk partitions in real time.

Sort Icon

Click the Sort icon to prioritize the list of information on the screen. Information can also be sorted in an ascending or descending order. (Figure A-3).
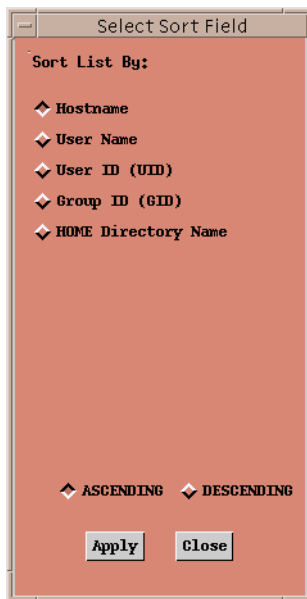
Figure A-3 Sort window

To perform a sort, select the desired criteria and click Apply. The longer the list, the longer the sort will take. Click the Close button to abort the sort. Most reports are sortable and each has its own set of sort criteria, depending on the information available.

Help Icon

The Help icon is available on all report and configuration screens. Click the Help icon for more information from an on-line Help screen. (Figure A-4).
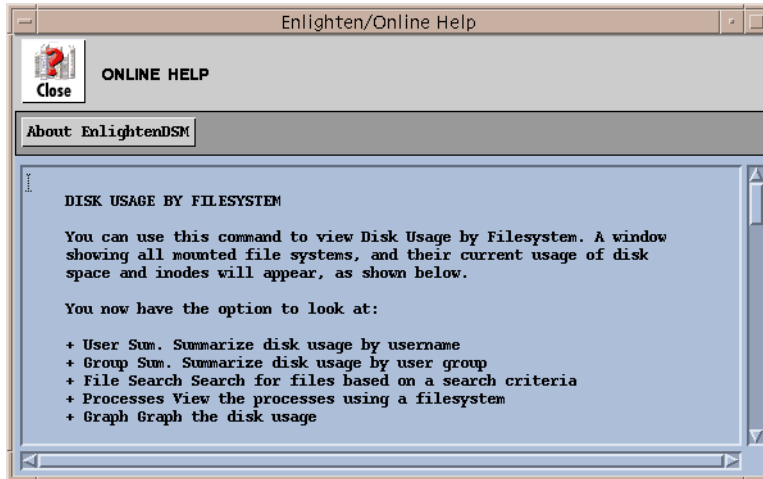

Figure A-4 Help window

List Function Buttons

The list function buttons, located on the right side of a window (Figure A-5), are used to manipulate items within a list box.
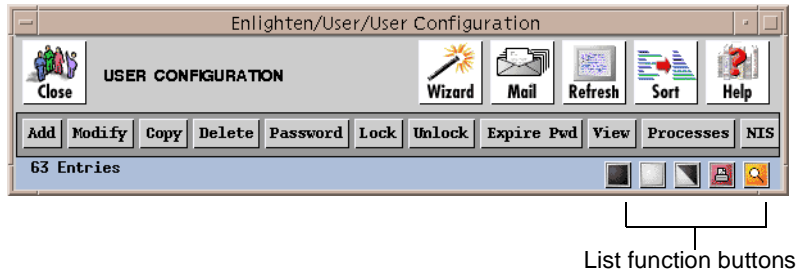
Figure A-5 List function buttons

Click the corresponding button to:

Select all items in the list

Deselect all items in the list

Invert selections

Print the report

Search the list for a given text string

Selection Options

Click the Select button to highlight the full list box. For long lists this may take some time. Click the Deselect button and all of the highlighted selections will be deselected.

Click the Invert button to highlight all list entries not currently selected, while unhighlighting those currently selected. This is especially useful when you want to select most, but not all, of a list; first select the few exceptions, then click the Invert button.

Printing

Click the Print button to print the currently displayed information. The print job will be sent to the path entered in the Session Preferences window (Configure menu). For more information on setting up the Print command, refer to Chapter 2, "Configure," in the Enlighten**DSM** *Reference Manual*.

Searching

Click the Search button to bring up the Search List window (Figure A-6). Enter your search string, set the scope of the search, and specify how the results should be presented.
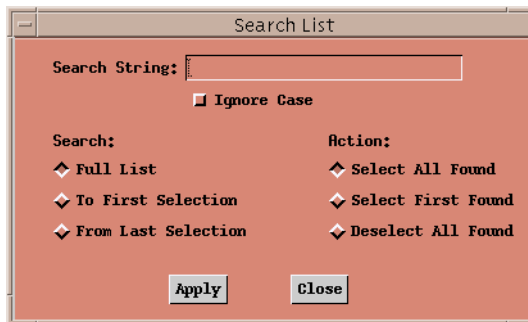
Figure A-6 Search List window

Pattern

Only simple text strings are allowed in the Search String field. Regular expressions or wildcards are not allowed. All characters are searched, including any leading or trailing spaces. Use the Ignore Case toggle perform a case-sensitive search (the default) or to ignore cases during the search. A successful match occurs whenever the pattern is found anywhere in the list entry.

Scope

A search on the Full List begins with the first item and proceeds through the list to the last item.

A search To First Selection begins with the first item and continues until a currently selected item is reached. If the first item in the list is currently selected, no search is performed.

A search From Last Selection begins with the item following the last selected item and continues to the bottom of the list. If the last item is currently selected, no search is performed.

Highlighting the Results

Choose the action to be taken after a successful match. By default, all entries successfully matched will be highlighted (Select All Found). You can also choose to have all matches unhighlighted (Deselect All Found); or, you can have the search stop after the first successful match (Select First Found), which is then highlighted if it was not already so.

Performing Actions on List Items

In Enlighten**DSM** windows, the available actions are displayed as buttons just below the global icons. Since these actions differ for each type of report, each window's option buttons are unique. See the Enlighten**DSM** *Reference Manual* for information on actions for various functions. Figure A-7 shows the action buttons for the User Configuration window.
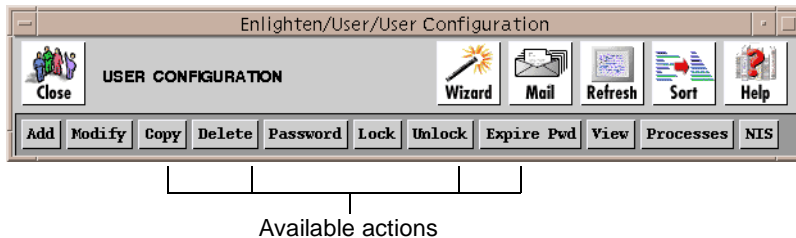

Figure A-7 Action buttons

If no items are selected when an action button is clicked, an error message similar to Figure A-8 appears, reminding you to first select a list entry.

Figure A-8 Error message

Click the OK button, select the desired items from the original list and reselect the relevant action button.

☞ Some actions, such as changing user ownership, require further
information before being completed. In such cases, a reminder
notice will appear.

Confirming Requests

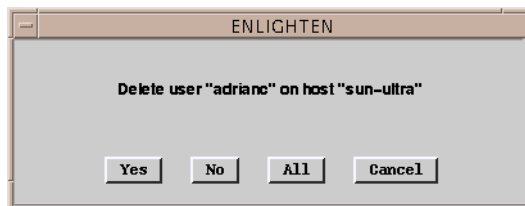Enlighten**DSM** may also prompt you for confirmation of a requested action
(Figure A-9).

Figure A-9 Confirmation request

Select from the following confirmation options:

Yes	Perform the action described in the window and proceed to the next highlighted item (if any).
No	Do not perform the action on the current item and proceed to the next highlighted item (if any).
All	Perform the action on all remaining highlighted items.
Cancel	Cancel the current command. This has the same effect as clicking the No button for all remaining highlighted items.
	Clicking the Cancel button will not reverse any actions taken against previously selected items from the current list.

Graphs

Several windows, such as Swap Space Usage, all you to create graphs of the
data (Figure A-10). Enlighten**DSM** determines the appropriate graph type for
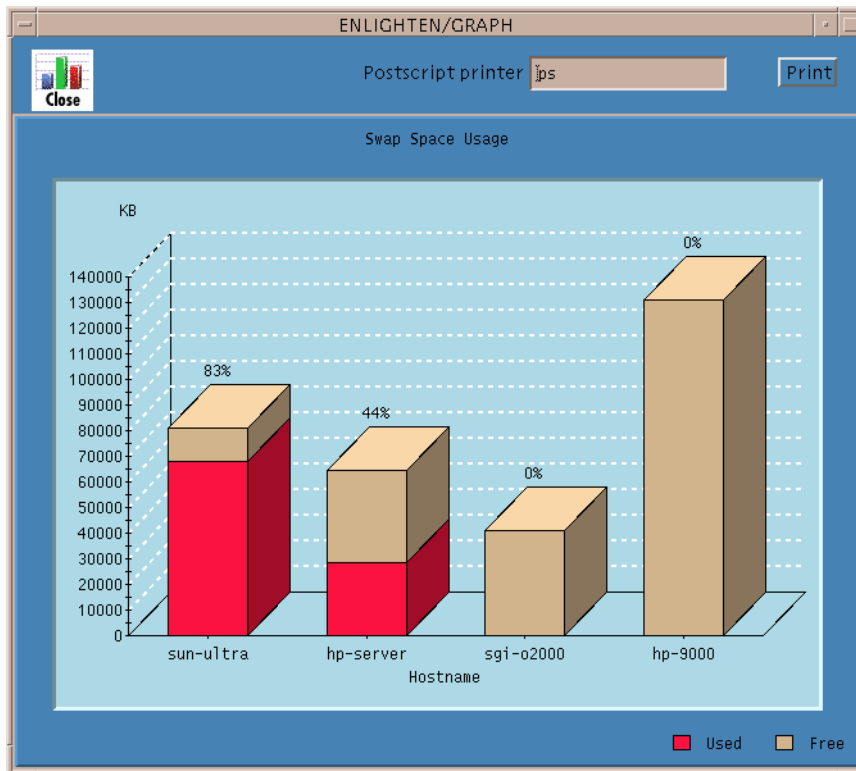each type of report.

Figure A-10 Displaying report data

Only highlighted items are graphed. If no items are selected, Enlighten**DSM** graphs all entries in the report.

Items are graphed in the same order as they appear in the report. Remember to sort the report as needed (via the Sort icon) before generating the graph.

Once the graph is displayed, use the middle mouse button to rotate the picture about its axis. Click the Print button to print the graph.

Creating Defaults

Enlighten**DSM** defaults automatically set up many of the program's Session Preferences. To revise any of these parameters, choose Session Preferences from the Configure menu. For more details, see Chapter 2, "Configure," in the Enlighten**DSM** *Reference Manual*.

Network Pools

Network Pools are used by Enlighten**DSM** to simultaneously manage multiple hosts. A "pool" treats a group of hosts as a single entity. When viewing disk usage by pool, for example, the usage of all hosts in that pool is displayed. Similarly, any monitoring request will report activities from all hosts in the pool. This allows multiple hosts to be monitored simultaneously.

☞ If the Hostnames field in a window is left empty, any actions performed from that window will affect *all* hosts in the current pool.

As each report is generated, each item in the report will also show the original hostname associated with it. Specific host entries can then be viewed and selected for performing further activities.

Performance Expectations

Enlighten**DSM** is constructed from a layered, client-server approach. Each layer is responsible for providing one phase of handling and preparing the data. Therefore, the operating system environment and the actions you request can affect the rate for obtaining results. This section highlights some performance facts to help predict the time required to perform Enlighten**DSM** actions.

Large Lists

The most time-consuming part of displaying an Enlighten**DSM** window is the building of window lists. In fact, the time required grows geometrically with the length of each entry and the number of entries in a list. On large lists (greater than 1000 entries), more than 80% of the wait time is spent building the display.

Disk Snapshots

Enlighten**DSM** maintains a database of file entry information for each disk partition examined during a Search or Summary of Usage. The first time one of these features is used to examine a given partition, or when Rebuild is used, the selected partition databases are built. This process may take up to two minutes for a partition with 10,000 files. For more details on this process, see Chapter 7, "Disk," in the Enlighten**DSM** *Reference Manual*.

Data Requests and Sorts

When you request configuration or process information (disk partitions, disk files, CPU summaries, and so on), either a text database must be searched or a series of commands performed. On a system supporting many CPU or disk requests, all processes, including Enlighten**DSM**, run slowly. Furthermore, Enlighten**DSM** combines multiple commands to present the results properly; so what you might think of as a simple query could require complex handling.

AlthoughEnlighten**DSM** has time-saving methods to speed queries, data requests are similar to database requests—the larger the database to be searched and the broader the search criteria, the longer it takes to process the query and sort the results. When information is requested across hosts, this happens more.

The first queries across a network "warm up" the connections, so these requests may take several seconds longer than usual to complete. This timing is dependent on the current usage of the network. Once connections are firmly established, however, Enlighten**DSM** processes requests almost immediately.

In general, the time it takes to fulfill a request on one machine is multiplied by the number of machines. Of course, different machines also respond at different rates based on their chip speed, disk speeds, and process load. Finally, the time for completion for a request will always vary based on the amount of result data shipped back to the user.

B
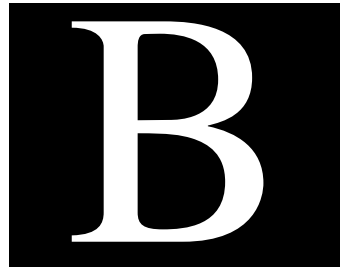
Troubleshooting

This appendix describes solutions to some common error messages and problems that you might encounter using Enlighten**DSM**.

Check the Enlighten**DSM** website at http:\\www.EnlightenDSM.com for up to date support information and answers to frequently asked questions.

Errors

Message: "Unable to access master daemon on host <hostname>."

Solution: The renld daemon is not running on the remote host. Start the daemon.

Message: "The host <hostname> has rejected your request. You do not have authorization to access the Enlighten Admin daemon on host <hostname>."

Solution: The host that you are running Enlighten**DSM** on is not in the remote host's hosts.allow file. The hosts.allow file is located in the Config subdirectory of the Enlighten**DSM** installation directory. To allow yourself access, edit this file on the remote host by placing your host in that file.

The hosts.allow file has three conditions:

- If no hostnames are listed in the host.allow file, then everyone has access to the remote daemon.

- If any hostname is listed in the hosts.allow file, then only those listed have access to the daemon.

- !hostname denies access to the specified host. (The exclamation point (!) means everything except the hostname associated with !.)

- If entries in the file contradict each other in allowing or not allowing access, access is denied. For example, if the hostnames listed in the file are !xyz and abc, efg, hij, and xyz, then !xyz takes precedence over xyz, so xyz would not have access.

Message: "Unable to access database."

Solution: Make sure that the EMD is up and running on the host it was installed on. Make sure the local /etc/enlighten file points the EMD to the correct host.

Message: "Unable to reach Events daemon on host <hostname>." (Message appears when you try to configure Events.)

Solution: Events on that particular host is not running, so start up the Events daemon. Restart all daemons with the $enlighten/bin/start_enl_daemons command.

Problems

Problem: I made a disk snapshot of a filesystem (Disk—>Save Current Snapshot) and after the files had changed, I tried to make another current snapshot of the filesystem by using Disk—>Save Current Snapshot command to compare the second snapshot with the first one. I got the message, "A master snapshot for partition <filename> already exists. Overwrite it with current snapshot?"

Solution: To create a second snapshot to compare to the first snapshot you made, you have to choose one of the Disk Usage commands from the Disk menu, select the filename, and then click the Rebuild button. Then click one of the function buttons, like the User Sum. button. This creates an updated snapshot of the filesystem that you can compare to the previous snapshots.

Problem: I changed ownership and permissions on the files, but they didn't show up.

Solution: Since the files you see listed are from the last disk snapshot, they do not reflect changes made since that last snapshot was taken. In order to view changes, rebuild the snapshot.

Problem: The icons on the Status Map sometimes fail to change back to green, even after the problem has been resolved.

Solution: The Events daemon may have been terminated before the "clear" message was sent to the host. The test that sent the trap may have also been turned off or deleted. The next time Events sends an alarm and the problem is resolved, the icon will return to green (status=OK).

Problem: When I open certain windows, like the User Configuration window, I get the message, "You do not have authorization to view this information."

Solution: You have been denied viewing privileges. Ask the administrator in charge of user authorizations to grant you viewing privileges (using the User Authorization window).

Problem: When I try to modify information, I get the message, "You do not have authorization to modify this information."

Solution: You have been denied modification privileges. Ask the administrator in charge of user authorization to grant you modification privileges (using the User Authorization window).

Problem: When I edit a host's `testtab` file, I receive an AgentMon startup events.

Solution: Editing the `testtab` file will cause AgentMon to restart so it can reread the parameters. AgentMon monitors the `testtab` file for any modifications.

Problem: Installing Enlighten**DSM** uses up too much disk space.

Solution: You can remove the unused maps from the `$ENLIGHTEN/maps` directory to free space.

Problem: Events daemons are not coming up after rebooting the system.

Solution: Make sure that there is not an SNMP daemon running. If one is running, kill it and edit the boot rc startup to remove the SNMP start entry.

Problem: Some program processes (for example, setting up DNS and NIS servers, building snapshots, or obvious password checking) seem to take a long time to complete.

Solution: It's normal for functions that process a lot of information to take a while to complete. Take a break and wait for the processes to be completed, or run these functions at night using cli. Password functions should be run at night.

Problem: When I click the Expert button to view the information about a specific problem, only the first few entries are displayed.

Solution: Fix the problems already listed in the Expert window, then click the Refresh button. Any other entries not previously displayed become visible.

C
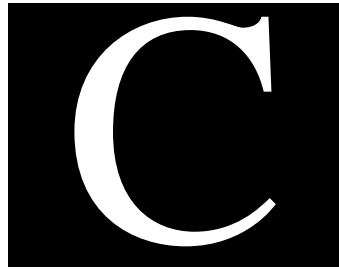
Installing Enlighten**DSM**

This chapter contains instructions for installing Enlighten**DSM**. Topics include:

- Pre-installation Considerations
- Mounting the CD-ROM
- Verifying that the CD-ROM is Mounted
- Starting the CDINSTAL Program
- Installing Enlighten**DSM**
- Starting Enlighten**DSM**

Pre-Installation Considerations

Before installing Enlighten**DSM**, gather the following information about the hosts in your intranet.

1) Choose which hosts will run the Enlighten**DSM** installer. This machine needs to have access to a CD-ROM drive and be able to connect through TCP/IP to the hosts where you want to install Enlighten**DSM**. The host will also be in charge of Enlighten**DSM** license management. Enlighten**DSM** uses flexLM for all license management. The flexLM components are installed in the `FLEXlm` sub-directory of the EMD installation directory.

2 Choose which host(s) will run the Enterprise Management Database (EMD) and Programmable Event Processor (PEP) servers. The EMD requires 20 MB of disk space for the initial installation. Depending on your particular situation, this may grow to 50 MB (average) or even 100 MB. The PEP engine requires 10 MB of free disk space. EMD is currently supported on Solaris, HP-UX, AIX, and SGI IRIX.

3 Determine the `root` password for each host that will use Enlighten**DSM**.

4 Find a suitable partition with enough free space for each host that will use Enlighten**DSM**. The default Enlighten**DSM** installation directory is `/opt/enlighten`. If there isn't enough free space on your `/opt` partition, select another partition with a minimum of 60 MB for Enlighten**DSM** agents and an additional 60-80 MB for each host running an Enlighten**DSM** GUI interface (`xenln`). Any partition will do, as long as the installation directory will only be used by Enlighten**DSM** (e.g., don't install in `/usr`, install in `/usr/enl` or `/usr/enlighten` instead).

Installing on an NFS-Imported Partition

For optimal performance, installing Enlighten**DSM** on an NFS partition is not recommended. In addition:

- EMD cannot be installed on an NFS mounted partitions.
- You must have root write permissions on the NFS partition on which you will install Enlighten**DSM**.

Enlighten**DSM** runs with root privileges and may create files within the Enlighten**DSM** home directory during its activities. If `root` write permissions are not enabled on the partition on which Enlighten**DSM** is installed, many of Enlighten**DSM**'s disk monitoring functions will not function. Be sure to install Enlighten**DSM** on a partition where `root` write permissions are enabled, such as local disk partitions.

• Writing to NFS filesystems is relatively slow. Consequently, disk snapshots may take three to five times longer to create. This should be taken into consideration even if an NFS partition has root write permissions.

SNMP Management

Consult the documentation that came with your network management application to manage Enlighten**DSM** Events Agent via SNMP. Your management application will need knowledge of our MIB. Some applications require a MIB in ASN-1 format while others require a copy of the MIB that meets the specifications of one or more RFC's. Enlighten**DSM** Events Agent provides both types of MIB formats.

☞ If you want to use the Events snmp daemon, AgentENL, delete any snmpd daemons that are running prior to installing Events on any of your hosts.

8-Bit X-Servers

Please note that the Enlighten Network Installer `enl_xinstall`, needs at least 32 free colors (viewed on an 800x700 or greater X-display) in order to start up. Some color intensive applications like Netscape will consume the entire colormap on your display, and `enl_xinstall` will not be able to come up. If you receive X-warnings regarding color, restart your x-server to clear your colormap, and try running `enl_xinstall` again.

C Installing EnlightenDSM

Mounting the CD-ROM

This section describes how to mount the CD-ROM on various systems. For all platforms, you must be the root user.

1) Type the following command:

   `su root`

2) Enter the root password when prompted. If the CD-ROM mount point, `/cdrom`, does not exist, type the command `mkdir /cdrom` to create it.

3) Type the following O/S specific commands to mount your CD-ROM (if you are unsure of your operating system, enter the command `uname -a`):

Solaris

On Solaris 2 systems, the Volume manager, vold, may automatically mount the CD-ROM. If vold does not mount the CD-ROM, type the following commands:

`/etc/mount -r -F hsfs /dev/sr0 /cdrom`

☞ You may need to use the complete device name for the cdrom drive, such as `/dev/dsk/c0t6d0s0.`

Sun OS

Type the following command:

`/etc/mount -r -t hsfs /dev/sr0 /cdrom`

HP-UX

Type the following command:

```
/etc/mount –r –u –t cdfs <device name> /cdrom
```

where `<device name>` is the device name and number associated with your local CD-ROM drive (e.g., `/dev/dsk/2s0`).

IBM AIX

Use SMIT to mount the CD-ROM:

1) Type the following commands:

   ```
   mkdir /cdrom
   ```

   ```
   smit storage
   ```

2) Select `File Systems`

3) Select `Add / Change / Show / Delete File Systems`

4) Select `CD-ROM File Systems`

5) Select `Add a CD-ROM File System`

6) Select the appropriate `DEVICE` name. If a conflict occurs, you may need to delete a previously defined CD-ROM file system.

7) Enter `/cdrom` as the `MOUNT POINT`.

8) Select `Do`.

   Alternatively, press Enter if you are using the ASCII version of SMIT.

9) Exit SMIT.

10) Type the following command:

    ```
    smit mountfs
    ```

11) Select `/dev/cd0`

    Alternatively, select a previously selected device.

12) Select `/cdrom` as the `DIRECTORY to mount`.

C Installing EnlightenDSM

13) Select cdrfs for the TYPE of file system.

14) Select yes for Mount as a READ-ONLY system.

15) Select Do.

Alternatively, press Return for the ASCII version of SMIT.

16) Exit SMIT.

SCO

Type the following command:

```
/etc/mount -r -f HS /dev/cd0 /cdrom
```

SGI IRIX

The Indigo Magic Desktop will mount the CD-ROM to /CDROM, and the CD drive icon on your desktop will show an inserted CD. If the CD drive is on another system, use the **Remote Directory** tool in your Desktop menu bar to locate the system and /CDROM directory.

Verifying that the CD-ROM is Mounted

To verify that the CD-ROM has been mounted, get a file list of the `root` directory of the CD-ROM by typing the following command:

```
ls /cdrom (or for IRIX, ls /CDROM)
```

You should see the following files in the directory:

- AIX41

- COMMON

- DOC

- HP10X

- IRIX5X6X

- SOL2X

- CDINSTAL

- DEC3240

- EXTRA

- WINTEL

- SCO5X

- SUNOS41X

If these files are not visible, then the CD-ROM was not properly mounted. Try mounting your CD-ROM again. If the CD-ROM still fails to mount, make sure that your cables are connected, that the CD-ROM has a unique SCSI ID, and that the machine was rebooted after you installed your CD-ROM drive.

C
**Installing
EnlightenDSM**

Starting the CDINSTAL Program

Once your X-Window System is running, you can run the CDINSTAL script.

1) Depending on your operating system, enter the following command:

- For Solaris using `vold`, type `/cdrom/cdrom0/cdinstal`

- For SunOS, type `/cdrom/cdinstal`

- For SGI, type `/CDROM/cdinstal`

- For all other platforms, type `/cdrom/CDINSTAL`

2) If this is your first-time installing Enlighten**DSM**, CDINSTAL will ask where you want to install Enlighten**DSM** components. The default Enlighten**DSM** installation directory is `/opt/enlighten` If there isn't enough free space on your `/opt` partition, select another partition with a minimum of 100 Mb of free space. Any partition will suffice, as long as the installation directory will only be used by Enlighten**DSM**.

☞ Don't install in `/usr`; instead, install in the `/usr/enl` or `/usr/enlighten`.

3) The CDINSTAL script will display a list of all the components being extracted into your chosen directory.

4) When the extraction is complete, the Enlighten**DSM** License Agreement dialog will appear (Figure C-1).

Carefully read and follow the instructions in the window. To accept the terms of the Licensing Agreement and continue installing Enlighten**DSM** on your system type "I Agree" followed by your name in the green text field. Then click Accept. To decline the License Agreement and quit the installation process, click the **DO NOT ACCEPT** button.

.
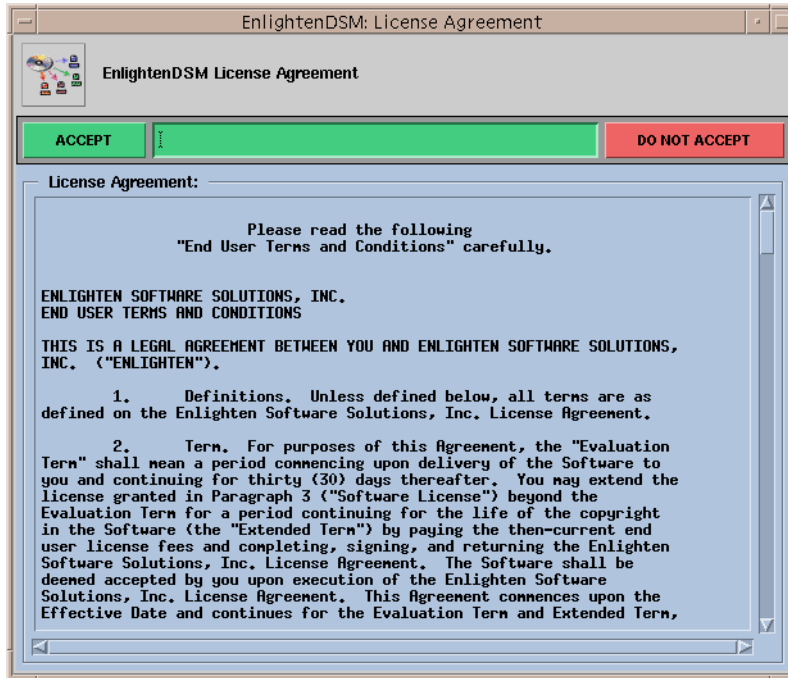
Figure C-1 Enlighten**DSM** License Agreement window

Installing Enlighten**DSM**

If you accept the License Agreement, the main Network Installer window will appear (Figure C-2).
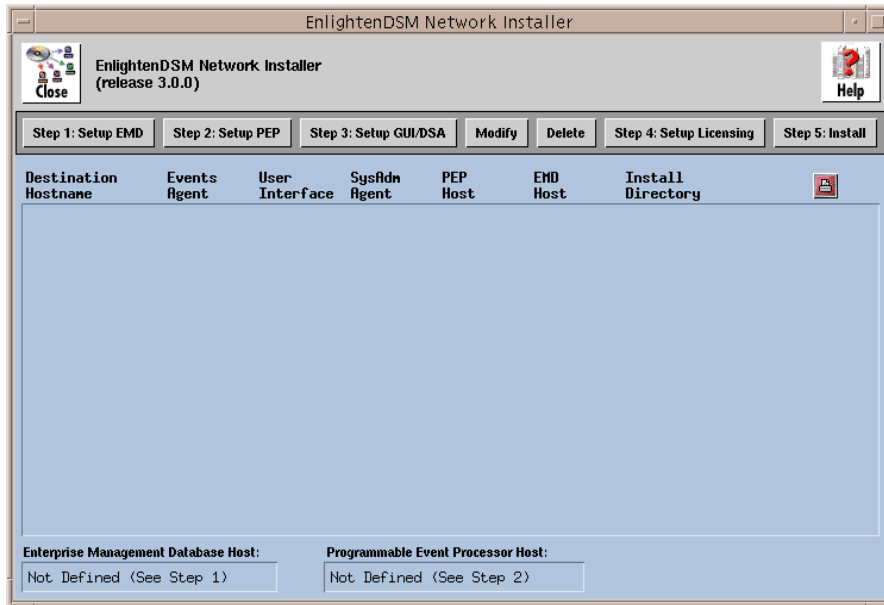
Figure C-2 Enlighten Network Installer

Basics

- At the top of the Network Installer screen are buttons for each step in the installation process. The rest of this section details how to install Enlighten**DSM** using these buttons.

- No files are transferred or modified in this installation process until you complete **"Step 5: Install"** and click the Start Installation button. If you make a mistake, just go back and type the correct information.

- Because your configuration information is verified dynamically; a slow network will cause the installation to run slowly.

☞ The Enlighten**DSM** Network Installer is SIGINT and SIGHUP aware. If you want to delete the Enlighten**DSM** Network Installer or cancel any operation, type <Ctrl>-c in the window where you started this process. Enlighten**DSM** will then prompt you to confirm your action.

Step Window Buttons

At the top of the Network Installer screen are buttons for each step in the installation process. For all the following steps, you can always click the Reset button to reset the fields to their default values. If you decide not to complete a particular step, simply click the Cancel button.

Once you have typed the necessary information for configuring a particular component, click the **OK** button to continue. If all the configuration information is acceptable, it will be added to the list in the Network Installer window.

Step 1: EMD Setup

Click the **Step 1: EMD Setup** button to configure the Enterprise Management Database (EMD).

The Setup the Enterprise Management Database window will appear, as shown in the example on page 12 (Figure C-3).

Fields and Toggles

This window contains the following fields and toggles:

EMD

Use this toggle to Install, Upgrade, or Don't Install EMD. Enlighten**DSM** components must have access to both EMD and the Programmable Events Processor (PEP) to function properly.

C — Installing EnlightenDSM

Install EMD on host

Specify where to load EMD and the database Engine. The default is your local hostname or the location of a previously installed EMD, if available. Depending on your DNS settings, you may need to specify a fully qualified host and domain name.

At most, you may have to type in the IP address of the host.

EMD Install Dir

Specify the directory that will hold EMD. This directory **must be** a local partition. The default is /opt/emd.
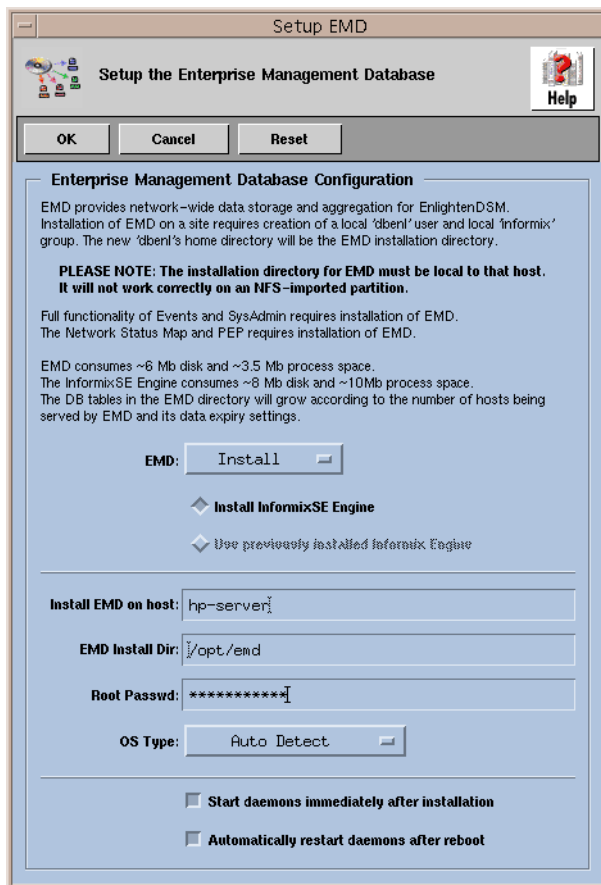
Figure C-3 Setup the Enterprise Management Database window

Root Passwd

Specify the root password for the host shown in the Install EMD on host: field.

OS Type

Specify the O/S type for the EMD host machine.

Start daemons immediately after installation

Ensure that this checkbox is selected (the default) or EMD will not be available for the other Enlighten**DSM** components.

Automatically restart daemons after reboot

Ensure that this checkbox is selected (the default) or EMD will not be available for the other Enlighten**DSM** components.

C
Installing
EnlightenDSM

Step 2: PEP Setup

Click the **Step 2: PEP Setup** button to configure the Programmable Events Processor (PEP). The Setup the Programmable Events Processor window will appear (Figure C-4).
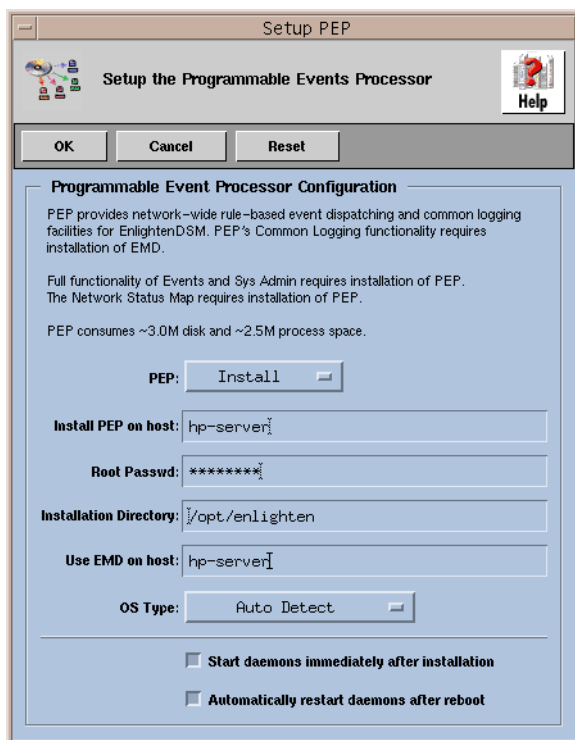
Figure C-4 Set Up the Programmable Events Processor window

This window contains the following fields and toggles:

PEP

Choose between Install, Upgrade, or Don't Install PEP. Enlighten**DSM** components must have access to both the Enterprise Management Database (EMD) and PEP to function properly.

Install PEP on host

Specify where to load PEP (hostname or IP address). The default is your local hostname or the location of a previously installed PEP, if available.

Root Passwd

Enter the root password for the host shown in the Install PEP on host: field.

Installation Directory

Specify the directory that will hold PEP.

Use EMD on host

Enter the hostname (or IP address) of the EMD host. This defaults to the value specified in "Step 1: EMD Setup"**.**

OS Type

Choose the O/S type for the PEP host machine. You may use the Auto Detect option (the default).

Start daemons immediately after installation

Ensure that this checkbox is selected (the default) or PEP will not be available for the other Enlighten**DSM** components.

Automatically restart daemons after reboot

Ensure that this checkbox is selected (the default) or PEP will not be available for the other Enlighten**DSM** components.

Step 3: Setup GUI/DSA

Click the **Step 3: Setup GUI/DSA** button to configure Enlighten**DSM** for the remainder of the hosts on your intranet. The Setup Enlighten**DSM** Destination Host(s) window will appear (Figure C-5).

Every host in your intranet must be able to resolve the hostname of the install host and the hosts running the EMD, PEP, and GUIs through DNS (including reverse lookups), NIS/NIS+, or local /etc/hosts file entries.

☞ To quickly modify the specified EMD and PEP hosts, double-click these entries in the Main Install list.
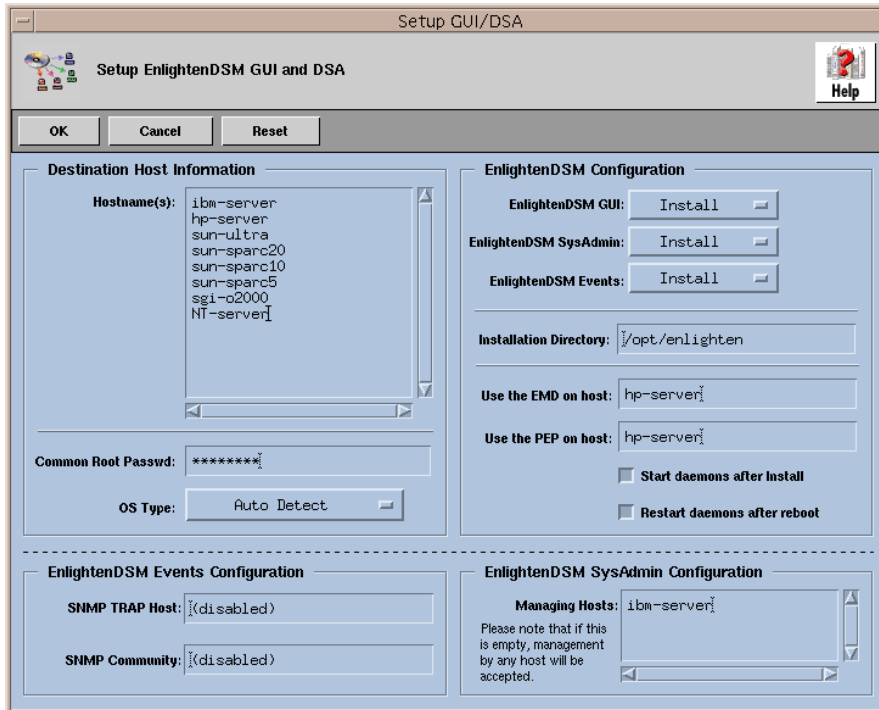
Figure C-5 Set Up Enlighten**DSM** Destination Host(s) window

Destination Host Information

This window contains the following fields and toggles:

Hostname(s)

Specify the host(s) where you will install the Enlighten**DSM** components specified in the Enlighten**DSM** Configuration area.

Common Root Passwd

Enter the common root password for the hosts shown in the Hostname(s): field. If some hosts have different root passwords, they will remain in the Hostname(s): field after you complete this step. You can then enter the root password for those hosts.

OS Type

Selects the O/S type for the host's machine. Use the Auto Detect option (the default), especially if you have a heterogeneous intranet.

Enlighten**DSM Configuration**

This area contains the following fields and toggles:

Enlighten**DSM GUI**

Select Install or Don't Install the Enlighten**DSM** GUI.

Enlighten**DSM SysAdmin**

Select Install or Don't Install the Enlighten**DSM** Sysadmin Agent.

Enlighten**DSM Events**

Select Install or Don't Install the Enlighten**DSM** Events Agent.

Installation Directory

Specify the directory that will hold the Enlighten**DSM** components. A full installation requires 100 Mb of free disk space.

Use the EMD on host

This field defaults to the hostname for your EMD. You should leave this field with its default value.

Use the PEP on host

This field defaults to the hostname for your PEP. You should leave this field with its default value.

Start daemons after install

Ensure that this checkbox is selected (the default) or the Enlighten**DSM** components will not come up automatically.

Restart daemons after reboot

Ensure that this checkbox is selected (the default) or the Enlighten**DSM** components will not come up automatically.

☞ If you want to manually configure your `rc` files, deselect the **Restart daemons after reboot**: button and add the call to `<enlighten install dir>/bin/S92startenl` in your system startup file.

Enlighten**DSM Events Configuration**

This area contains the following fields:

SNMP TRAP host

You can use this field to specify a hostname on your intranet that is running an SNMP-aware system manager. If you don't have such a manager installed, leave this set to `<disabled>` (the default). This will disable the SNMP trap generation.

SNMP Community

You can use this field to specify the COMMUNITY name of your SNMP network. The default is `<disabled>`.

Enlighten**DSM SysAdmin Agent Configuration**

This area contains the following field:

Managing Hosts

Use this list box to specify the host(s) that will be allowed to connect through the Enlighten**DSM** GUI to manage the hosts shown in the Hostname(s): list box. The default is your group management server (GMS).

Warning!	If you leave this field blank, **any** host running the Enlighten**DSM** GUI can manage the hosts shown in the Hostname(s): field. You should only use this option if you have a firewalled intranet containing "trusted" hosts.

Modify and Delete

Click the **Modify** button (in the Main window) to modify a host in the Main Install List. This will take you to the **Step 3: Setup GUI/DSA** screen.

Double-clicking a host in the Main Install List is the same as selecting that host and clicking on the **Modify** button. You can also click the **Delete** button (in the Main window) to delete the highlighted host(s) from the Main Install List.

Step 4: Setup Licensing

Once you have completed the previous three steps, click the **Step 4: Setup Licensing** button. The Enlighten**DSM** Configure License File window will appear (Figure C-6).
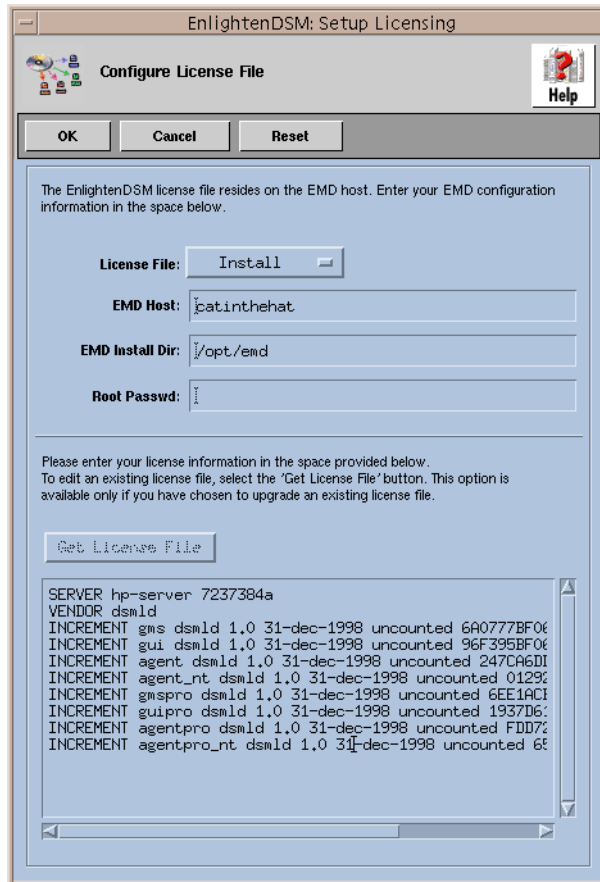
Figure C-6 Enlighten**DSM** Configure License File window

Under License File, select Install, Upgrade or Don't Install.

If this is a new installation, enter the EMD hostname and the directory where EMD should be placed (the default is /opt/emd).

The first two lines of the license file are automatically entered values based on the installation information already provided. These lines should not be changed. They specify the name and location of the license server process (which must be the same as the EMD host specified in **"Step 1: EMD Setup"**) and the name of the Enlighten**DSM** license vendor daemon that is also located on the EMD host.

The remaining license file details provided by your Enlighten**DSM** vendor must be manually entered. Depending on the license features being activated, you may have already received license keys. You may also need to send additional information to your Enlighten**DSM** vendor to receive the properly validated license keys. Please contact your Enlighten**DSM** product vendor for additional information about license keys.

To read the existing license file, click the Get License File button. The contents of the existing license will be displayed in the License Information window. If you wish to update your license file and are not upgrading or installing additional Enlighten**DSM** components, you can manually edit the `<emd_dir>/FLEXlm/license.dat` file on the EMD host.

Once the license file has been updated, click OK to continue the installation process.

Warning!	It is important that the *license file be installed correctly with valid license keys or the* Enlighten**DSM** *processes will not run.*

Step 5: Install

1) Click the **Step 5: Install** button to complete the installation process.

   The Install to Local and Remote Hosts window will appear.

2) Click the **Start Installation** button to install Enlighten**DSM** in the configuration(s) you specified in the previous four steps.

   The Installation progress information is written to you Enlighten directory as `logs/install.log` and to the Installation Log scrolled text field. Clicking on the **Clear Log** button will only clear the scrolled text field; this will not erase any entries in the `install.log` file.

3) Once the installation is complete, click **Close** to exit the Enlighten**DSM** Network Installer interface, and continue with the next section to start the Enlighten**DSM** GUI.

C Installing EnlightenDSM

Starting EnlightenDSM

1) From a shelltool, xterm, or another such window, type the following commands:

```
cd <your enlighten installation directory>
```

```
bin/xenln
```

If you have forgotten your Enlighten**DSM** installation directory, type the following command:

```
cat /etc/enlighten | grep "enlighten="
```

```
bin/xenln
```

2) Once you've found the correct directory and invoked `xenln`, the Enlighten**DSM** main menu toolbar will appear on your screen (Figure C-7).
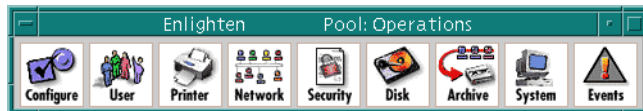
Figure C-7 Enlighten**DSM** main menu toolbar

Now you can begin using Enlighten**DSM** to help you administer your system and network needs. Refer to the Enlighten**DSM** *Reference Manual* and *User Guide* for more details on using this toolbar.

☞ An Enlighten**DSM** Sysadmin Agent daemon must be running on each host that you want to manage.

Installing EnlightenDSM for Windows

EnlightenDSM for Windows Requirements

- Windows 95, Windows, 98, Windows NT 4.0 Server or Windows NT 4.0 Workstation.

- Intel-compatible machine (486 or later) with 32 Mb RAM and 6 Mb disk space available

- Enlighten**DSM** for Windows NT must be installed on an NTFS partition.

- Please note that installations of Enlighten**DSM** must be performed by an Administrator.

Installing EnlightenDSM for Windows from CD-ROM

The following instructions will use D: to represent the letter that is assigned to your CD-ROM drive.

1) Insert the Enlighten**DSM** CD-ROM

2) From the Start menu, select Run...

3) In the Open: text field, type D:\Wintel\Disk1\Setup.exe, and click OK.

What's New with Version 3.0.0?

- Enlighten**DSM** now uses FLEXlm license management to license it's components.

- Enlighten**DSM** is now available with two different feature sets. (Workgroup and Advanced). The Workgroup set consists of the Configure, User, Printer and Network features. The Advanced set consists of the Workgroup set plus the Security, Disk, Archive, System, Events and Status Map features.

EnlightenDSM Installation Notes

- For the Enlighten**DSM** services to run properly, they must be able to communicate (through RPC) to a UNIX host running EMD and a UNIX host running PEP (which may be the same UNIX host). The configuration GUI, dsm_cfg.exe, will prompt for these hostnames.

C Installing
EnlightenDSM

- To graphically manage your enterprise, you must use the Enlighten**DSM** GUI. The GUI must run on a UNIX host and be displayed on an X/motif-compliant host. Perform the following steps to display the Enlighten**DSM** GUI on a Windows machine:

    1) Install third-party X software on the Windows machine you want to display the GUI on.

    2) On a UNIX box that has the Enlighten**DSM** GUI, `bin/xenln`, set the DISPLAY environment variable to `"<name of your Windows machine>:0.0"`

    3) Then run `bin/xenln` (from the UNIX machine).

- Enlighten**DSM** provides SNMP services via the AgentMon and AgentENL services. AgentENL primarily provides SNMP and SMUX protocol services, AgentMon performs all monitoring functions. Because many other SNMP packages are available for Windows, some of which do not support the SMUX (SNMP Multiplexor) protocol, AgentMon has been modified to operate without SNMP/SMUX if it cannot connect to AgentENL or some other SMUX-compliant service. If AgentENL interferes with other SNMP services on your system, you may now safely disable the AgentENL service and AgentMon will continue to generate EMD events, send email, etc. It will just not be reachable from SNMP applications. AgentMon will automatically re-connect to an SMUX service if one becomes available at a later time.

- Whether SNMP is in use by Enlighten**DSM** or not, you should still install the SNMP Service package. This package is needed to obtain meaningful network statistics, the Microsoft SNMP Service package contains drivers which provide network performance data. AgentMon monitors network statistics whether it will be generating SNMP traps or EMD events. After installing the SNMP package (Start->Settings->Control Panel->Network, select the Services tab, click Add..., then select "SNMP Service"). If you wish to run the SNMP service provided by Enlighten**DSM**, either run the DSM Configuration GUI, `dsm_cfg.exe`, or ensure the Microsoft SNMP service is disabled.

- The Simple TCP/IP Services package must be installed in order to use Enlighten**DSM**'s clock synch feature with a Windows machine as the time server. Install the package by opening Start->Settings->Control

Panel->Network, select the Services tab, click Add..., then select "Simple TCP/IP Services"),

- EnlightenDSM requires (and ships with) an ONCRPC-compliant portmap service. If you have installed a third-party NFS client, this service will already be registered, and the DSM Configuration GUI will not register EnlightenDSM's own ONCRPC portmapper. If you deinstall the NFS client software (and consequently remove the non-EnlightenDSM portmap service), make sure you run the DSM Configuration GUI after deinstallation - dsm_cfg.exe will then register EnlightenDSM's portmap service.

- EnlightenDSM Events' default mailer destination is root@mailhost. If your enterprise does not have a hostname alias for "mailhost," or if you want Events notification to go elsewhere, use the EnlightenDSM GUI or simply edit config\testtab.<hostname>.

- Be aware that the logout and system shutdown functionality is forceful in order to ensure the system actually shuts down at the specified time. Interactive users are warned about the impending shutdown, but cannot cancel the action.

- The child process of the EnlightenDSM SysAdmin service, renldc, runs as the System Account. It is possible that ACL permissions on a file or set of files could disallow read access to the System Account. For that reason, disk snapshot databases may include every file on disk.

- User and group administration information is always provided for the local machine - even if the interactive user is logged onto, for example, the "Workspace Domain," the EnlightenDSM GUI will list users on the local machine. To get a listing of the users on the "Workspace Domain", add the Primary Domain Controller for the "Workspace Domain" to your current Network Pool (in the EnlightenDSM GUI, click Configure->Pool Configuration, select your current pool, then click modify.)

C Installing
EnlightenDSM

Installing EnlightenDSM for Windows from a Network Drive

On the Network Drive machine:

1) Insert the Enlighten**DSM** CD-ROM

2) Copy the contents of the CD-ROM's folder into a shared filesystem (example: \\networkdrive\sharename\wintel).

On the machine you want to install Enlighten**DSM**:

1) Map the Network Drive that contains the wintel installation fileset that you just copied to a drive letter. In the above example, you could map P: to \\networkdrive\sharename by right-clicking on "My Computer," selecting Map Network Drive..., select Drive: P:, and in Shared Directories:, select \\networkdrive\sharename. Then click OK.

2) From the Start menu, select Run...

3) Click Browse..., access the mapped wintel\disk1 directory on the Network Drive, select Setup.exe, and click Open. To continue the example, you could instead type P:\Wintel\Disk1\Setup.exe in the Open: text field.

Installing EnlightenDSM for Windows from High Density (1.44Mb) Floppy Disk

The Windows installation fileset has been broken up so that it can be copied from the Enlighten**DSM** CD-ROM onto two floppy disks (to install to machines that don't have a local CD-ROM drive).

To create these installation disks, perform the following steps:

1) Copy Wintel\Disk1 to Disk 1.

2) Copy Wintel\Disk2 to Disk 2.

3) Go to the Windows machine without a CD-ROM and insert Disk 1.

4) Run a:\Setup.exe.
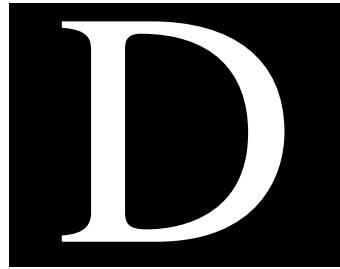
De-installing Enlighten**DSM**

You must de-install Enlighten**DSM** on a local host. There is no multiple network de-installation facility. To de-install Enlighten**DSM**,

1) Log in as a root user.

2) cd into the `enlighten` directory. If the EMD is the only component installed on the local host, cd into the `EMD` directory.

☞ If you forgot where Enlighten**DSM** is installed on the local host, type the following command:

`grep enlighten= /etc/enlighten`

Or, for the EMD directory:

`grep emd_dir= etc/enlighten`

or cat the /etc/enlighten file.

3) Type the command:

   ./.deinstall

4) Follow the instructions on screen as you are prompted.

Here's a sample de-install:

```
$ su
Password:
# grep enlighten= /etc/enlighten
enlighten=/opt/enlighten
# cd /opt/enlighten
# ./.deinstall
.deinstall: Are you sure you want to deinstall Enlighten?
[y|n|exit] y
.deinstall: Killing the following Enlighten processes:
PID TTY        TIME COMD
  2417 pts/15    0:08 xenln
stop_enl_daemons: Killing the following Enlighten
daemons:
PID TTY        TIME COMD
  2506 ?          0:04 AgentMon
  2468 pts/25   0:00 e_tokend
  2494 pts/25   0:00 AgentENL
  2407 ?          0:00 renld
stop_enl_daemons: Waiting for daemons to shut down.......
stop_enl_daemons: Finished.
stop_enl_daemons: Waiting for daemons to shut down..
stop_enl_daemons: Finished.
```

```
.deinstall: Remove Enlighten references from
/etc/rc2.d/s92startenl?

.deinstall: [y|n|exit] (default is 'y') > y

.deinstall: /etc/rc2.d/s92startenl was archived to
/tmp/s92startenl.old

.deinstall: Recursively remove the Enlighten directory
/opt/enlighten?

.deinstall: [y|n|exit] (default is 'y') > y

.deinstall: Removed /etc/enlighten.

.deinstall: Successfully finished deinstalling Enlighten.
```

Index